
AWS Well- Architected Framework

AWS Well-Architected Framework

Copyright © 2022 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Abstract and introduction	1
Introduction	1
Definitions	1
On architecture	3
General design principles	4
The pillars of the framework	5
Operational excellence	5
Design principles	5
Definition	6
Best practices	6
Resources	12
Security	12
Design principles	12
Definition	13
Best practices	13
Resources	18
Reliability	18
Design principles	19
Definition	19
Best practices	19
Resources	23
Performance efficiency	23
Design principles	24
Definition	24
Best practices	24
Resources	29
Cost optimization	29
Design principles	30
Definition	30
Best practices	31
Resources	34
Sustainability	35
Design principles	35
Definition	36
Best practices	36
The review process	41
Conclusion	43
Contributors	44
Further reading	45
Document revisions	46
Appendix: Questions and best practices	48
Operational excellence	48
Organization	48
Prepare	64
Operate	92
Evolve	110
Security	117
Security foundations	117
Identity and access management	123
Detection	137
Infrastructure protection	142
Data protection	152
Incident response	162
Reliability	169

Foundations	169
Workload architecture	184
Change management	202
Failure management	221
Performance efficiency	278
Selection	278
Review	330
Monitoring	333
Tradeoffs	340
Cost optimization	346
Practice Cloud Financial Management	346
Expenditure and usage awareness	358
Cost-effective resources	373
Manage demand and supply resources	387
Optimize over time	390
Sustainability	392
Region selection	393
User behavior patterns	393
Software and architecture patterns	398
Data patterns	402
Hardware patterns	407
Development and deployment process	410
Notices	414

AWS Well-Architected Framework

Publication date: **October 20, 2022** ([Document revisions](#) (p. 46))

The AWS Well-Architected Framework helps you understand the pros and cons of decisions you make while building systems on AWS. By using the Framework you will learn architectural best practices for designing and operating reliable, secure, efficient, cost-effective, and sustainable systems in the cloud.

Introduction

The AWS Well-Architected Framework helps you understand the pros and cons of decisions you make while building systems on AWS. Using the Framework helps you learn architectural best practices for designing and operating secure, reliable, efficient, cost-effective, and sustainable workloads in the AWS Cloud. It provides a way for you to consistently measure your architectures against best practices and identify areas for improvement. The process for reviewing an architecture is a constructive conversation about architectural decisions, and is not an audit mechanism. We believe that having well-architected systems greatly increases the likelihood of business success.

AWS Solutions Architects have years of experience architecting solutions across a wide variety of business verticals and use cases. We have helped design and review thousands of customers' architectures on AWS. From this experience, we have identified best practices and core strategies for architecting systems in the cloud.

The AWS Well-Architected Framework documents a set of foundational questions that allow you to understand if a specific architecture aligns well with cloud best practices. The framework provides a consistent approach to evaluating systems against the qualities you expect from modern cloud-based systems, and the remediation that would be required to achieve those qualities. As AWS continues to evolve, and we continue to learn more from working with our customers, we will continue to refine the definition of well-architected.

This framework is intended for those in technology roles, such as chief technology officers (CTOs), architects, developers, and operations team members. It describes AWS best practices and strategies to use when designing and operating a cloud workload, and provides links to further implementation details and architectural patterns. For more information, see the [AWS Well-Architected homepage](#).

AWS also provides a service for reviewing your workloads at no charge. The [AWS Well-Architected Tool](#) (AWS WA Tool) is a service in the cloud that provides a consistent process for you to review and measure your architecture using the AWS Well-Architected Framework. The AWS WA Tool provides recommendations for making your workloads more reliable, secure, efficient, and cost-effective.

To help you apply best practices, we have created [AWS Well-Architected Labs](#), which provides you with a repository of code and documentation to give you hands-on experience implementing best practices. We also have teamed up with select AWS Partner Network (APN) Partners, who are members of the [AWS Well-Architected Partner program](#). These AWS Partners have deep AWS knowledge, and can help you review and improve your workloads.

Definitions

Every day, experts at AWS assist customers in architecting systems to take advantage of best practices in the cloud. We work with you on making architectural trade-offs as your designs evolve. As you deploy

these systems into live environments, we learn how well these systems perform and the consequences of those trade-offs.

Based on what we have learned, we have created the AWS Well-Architected Framework, which provides a consistent set of best practices for customers and partners to evaluate architectures, and provides a set of questions you can use to evaluate how well an architecture is aligned to AWS best practices.

The AWS Well-Architected Framework is based on six pillars — operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability.

Table 1. The pillars of the AWS Well-Architected Framework

Name	Description
Operational excellence	The ability to support development and run workloads effectively, gain insight into their operations, and to continuously improve supporting processes and procedures to deliver business value.
Security	The security pillar describes how to take advantage of cloud technologies to protect data, systems, and assets in a way that can improve your security posture.
Reliability	The reliability pillar encompasses the ability of a workload to perform its intended function correctly and consistently when it's expected to. This includes the ability to operate and test the workload through its total lifecycle. This paper provides in-depth, best practice guidance for implementing reliable workloads on AWS.
Performance efficiency	The ability to use computing resources efficiently to meet system requirements, and to maintain that efficiency as demand changes and technologies evolve.
Cost optimization	The ability to run systems to deliver business value at the lowest price point.
Sustainability	The ability to continually improve sustainability impacts by reducing energy consumption and increasing efficiency across all components of a workload by maximizing the benefits from the provisioned resources and minimizing the total resources required.

In the AWS Well-Architected Framework, we use these terms:

- A **component** is the code, configuration, and AWS Resources that together deliver against a requirement. A component is often the unit of technical ownership, and is decoupled from other components.
- The term **workload** is used to identify a set of components that together deliver business value. A workload is usually the level of detail that business and technology leaders communicate about.
- We think about **architecture** as being how components work together in a workload. How components communicate and interact is often the focus of architecture diagrams.

- **Milestones** mark key changes in your architecture as it evolves throughout the product lifecycle (design, implementation, testing, go live, and in production).
- Within an organization the **technology portfolio** is the collection of workloads that are required for the business to operate.
- The **level of effort** is categorizing the amount of time, effort, and complexity a task requires for implementation. Each organization needs to consider the size and expertise of the team and the complexity of the workload for additional context to properly categorize the level of effort for the organization.
 - **High:** The work might take multiple weeks or multiple months. This could be broken out into multiple stories, releases, and tasks.
 - **Medium:** The work might take multiple days or multiple weeks. This could be broken out into multiple releases and tasks.
 - **Low:** The work might take multiple hours or multiple days. This could be broken out into multiple tasks.

When architecting workloads, you make trade-offs between pillars based on your business context. These business decisions can drive your engineering priorities. You might optimize to improve sustainability impact and reduce cost at the expense of reliability in development environments, or, for mission-critical solutions, you might optimize reliability with increased costs and sustainability impact. In ecommerce solutions, performance can affect revenue and customer propensity to buy. Security and operational excellence are generally not traded-off against the other pillars.

On architecture

In on-premises environments, customers often have a central team for technology architecture that acts as an overlay to other product or feature teams to ensure they are following best practice. Technology architecture teams typically include a set of roles such as: Technical Architect (infrastructure), Solutions Architect (software), Data Architect, Networking Architect, and Security Architect. Often these teams use [TOGAF](#) or the [Zachman Framework](#) as part of an enterprise architecture capability.

At AWS, we prefer to distribute capabilities into teams rather than having a centralized team with that capability. There are risks when you choose to distribute decision making authority, for example, ensuring that teams are meeting internal standards. We mitigate these risks in two ways. First, we have *practices* (ways of doing things, process, standards, and accepted norms) that focus on enabling each team to have that capability, and we put in place experts who ensure that teams raise the bar on the standards they need to meet. Second, we implement *mechanisms* that carry out automated checks to ensure standards are being met.

“Good intentions never work, you need good mechanisms to make anything happen” — Jeff Bezos.

This means replacing a human's best efforts with mechanisms (often automated) that check for compliance with rules or process. This distributed approach is supported by the [Amazon leadership principles](#), and establishes a culture across all roles that *works back* from the customer. Working backward is a fundamental part of our innovation process. We start with the customer and what they want, and let that define and guide our efforts. Customer-obsessed teams build products in response to a customer need.

For architecture, this means that we expect every team to have the capability to create architectures and to follow best practices. To help new teams gain these capabilities or existing teams to raise their bar, we enable access to a virtual community of principal engineers who can review their designs and help them understand what AWS best practices are. The principal engineering community works to make best practices visible and accessible. One way they do this, for example, is through lunchtime talks that focus on applying best practices to real examples. These talks are recorded and can be used as part of onboarding materials for new team members.

AWS best practices emerge from our experience running thousands of systems at internet scale. We prefer to use data to define best practice, but we also use subject matter experts, like principal engineers, to set them. As principal engineers see new best practices emerge, they work as a community to ensure that teams follow them. In time, these best practices are formalized into our internal review processes, as well as into mechanisms that enforce compliance. The Well-Architected Framework is the customer-facing implementation of our internal review process, where we have codified our principal engineering thinking across field roles, like Solutions Architecture and internal engineering teams. The Well-Architected Framework is a scalable mechanism that lets you take advantage of these learnings.

By following the approach of a principal engineering community with distributed ownership of architecture, we believe that a Well-Architected enterprise architecture can emerge that is driven by customer need. Technology leaders (such as a CTOs or development managers), carrying out Well-Architected reviews across all your workloads will allow you to better understand the risks in your technology portfolio. Using this approach, you can identify themes across teams that your organization could address by mechanisms, training, or lunchtime talks where your principal engineers can share their thinking on specific areas with multiple teams.

General design principles

The Well-Architected Framework identifies a set of general design principles to facilitate good design in the cloud:

- **Stop guessing your capacity needs:** If you make a poor capacity decision when deploying a workload, you might end up sitting on expensive idle resources or dealing with the performance implications of limited capacity. With cloud computing, these problems can go away. You can use as much or as little capacity as you need, and scale up and down automatically.
- **Test systems at production scale:** In the cloud, you can create a production-scale test environment on demand, complete your testing, and then decommission the resources. Because you only pay for the test environment when it's running, you can simulate your live environment for a fraction of the cost of testing on premises.
- **Automate to make architectural experimentation easier:** Automation allows you to create and replicate your workloads at low cost and avoid the expense of manual effort. You can track changes to your automation, audit the impact, and revert to previous parameters when necessary.
- **Allow for evolutionary architectures:** In a traditional environment, architectural decisions are often implemented as static, onetime events, with a few major versions of a system during its lifetime. As a business and its context continue to evolve, these initial decisions might hinder the system's ability to deliver changing business requirements. In the cloud, the capability to automate and test on demand lowers the risk of impact from design changes. This allows systems to evolve over time so that businesses can take advantage of innovations as a standard practice.
- **Drive architectures using data:** In the cloud, you can collect data on how your architectural choices affect the behavior of your workload. This lets you make fact-based decisions on how to improve your workload. Your cloud infrastructure is code, so you can use that data to inform your architecture choices and improvements over time.
- **Improve through game days:** Test how your architecture and processes perform by regularly scheduling game days to simulate events in production. This will help you understand where improvements can be made and can help develop organizational experience in dealing with events.

The pillars of the framework

Creating a software system is a lot like constructing a building. If the foundation is not solid, structural problems can undermine the integrity and function of the building. When architecting technology solutions, if you neglect the six pillars of operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability, it can become challenging to build a system that delivers on your expectations and requirements. Incorporating these pillars into your architecture will help you produce stable and efficient systems. This will allow you to focus on the other aspects of design, such as functional requirements.

Pillars

- [Operational excellence \(p. 5\)](#)
- [Security \(p. 12\)](#)
- [Reliability \(p. 18\)](#)
- [Performance efficiency \(p. 23\)](#)
- [Cost optimization \(p. 29\)](#)
- [Sustainability \(p. 35\)](#)

Operational excellence

The Operational Excellence pillar includes the ability to support development and run workloads effectively, gain insight into their operations, and to continuously improve supporting processes and procedures to deliver business value.

The operational excellence pillar provides an overview of design principles, best practices, and questions. You can find prescriptive guidance on implementation in the [Operational Excellence Pillar whitepaper](#).

Topics

- [Design principles \(p. 5\)](#)
- [Definition \(p. 6\)](#)
- [Best practices \(p. 6\)](#)
- [Resources \(p. 12\)](#)

Design principles

There are five design principles for operational excellence in the cloud:

- **Perform operations as code:** In the cloud, you can apply the same engineering discipline that you use for application code to your entire environment. You can define your entire workload (applications, infrastructure) as code and update it with code. You can implement your operations procedures as code and automate their execution by triggering them in response to events. By performing operations as code, you limit human error and enable consistent responses to events.
- **Make frequent, small, reversible changes:** Design workloads to allow components to be updated regularly. Make changes in small increments that can be reversed if they fail (without affecting customers when possible).
- **Refine operations procedures frequently:** As you use operations procedures, look for opportunities to improve them. As you evolve your workload, evolve your procedures appropriately. Set up regular

game days to review and validate that all procedures are effective and that teams are familiar with them.

- **Anticipate failure:** Perform “pre-mortem” exercises to identify potential sources of failure so that they can be removed or mitigated. Test your failure scenarios and validate your understanding of their impact. Test your response procedures to ensure that they are effective, and that teams are familiar with their execution. Set up regular game days to test workloads and team responses to simulated events.
- **Learn from all operational failures:** Drive improvement through lessons learned from all operational events and failures. Share what is learned across teams and through the entire organization.

Definition

There are four best practice areas for operational excellence in the cloud:

- **Organization**
- **Prepare**
- **Operate**
- **Evolve**

Your organization’s leadership defines business objectives. Your organization must understand requirements and priorities and use these to organize and conduct work to support the achievement of business outcomes. Your workload must emit the information necessary to support it. Implementing services to enable integration, deployment, and delivery of your workload will enable an increased flow of beneficial changes into production by automating repetitive processes.

There may be risks inherent in the operation of your workload. You must understand those risks and make an informed decision to enter production. Your teams must be able to support your workload. Business and operational metrics derived from desired business outcomes will enable you to understand the health of your workload, your operations activities, and respond to incidents. Your priorities will change as your business needs and business environment changes. Use these as a feedback loop to continually drive improvement for your organization and the operation of your workload.

Best practices

Topics

- [Organization \(p. 6\)](#)
- [Prepare \(p. 9\)](#)
- [Operate \(p. 10\)](#)
- [Evolve \(p. 11\)](#)

Organization

Your teams need to have a shared understanding of your entire workload, their role in it, and shared business goals to set the priorities that will enable business success. Well-defined priorities will maximize the benefits of your efforts. Evaluate internal and external customer needs involving key stakeholders, including business, development, and operations teams, to determine where to focus efforts. Evaluating customer needs will ensure that you have a thorough understanding of the support that is required to achieve business outcomes. Ensure that you are aware of guidelines or obligations defined by your organizational governance and external factors, such as regulatory compliance requirements and industry standards, that may mandate or emphasize specific focus. Validate that you have mechanisms to identify changes to internal governance and external compliance requirements. If no requirements

are identified, ensure that you have applied due diligence to this determination. Review your priorities regularly so that they can be updated as needs change.

Evaluate threats to the business (for example, business risk and liabilities, and information security threats) and maintain this information in a risk registry. Evaluate the impact of risks, and tradeoffs between competing interests or alternative approaches. For example, accelerating speed to market for new features may be emphasized over cost optimization, or you may choose a relational database for non-relational data to simplify the effort to migrate a system without refactoring. Manage benefits and risks to make informed decisions when determining where to focus efforts. Some risks or choices may be acceptable for a time, it may be possible to mitigate associated risks, or it may become unacceptable to allow a risk to remain, in which case you will take action to address the risk.

Your teams must understand their part in achieving business outcomes. Teams need to understand their roles in the success of other teams, the role of other teams in their success, and have shared goals. Understanding responsibility, ownership, how decisions are made, and who has authority to make decisions will help focus efforts and maximize the benefits from your teams. The needs of a team will be shaped by the customer they support, their organization, the makeup of the team, and the characteristics of their workload. It's unreasonable to expect a single operating model to be able to support all teams and their workloads in your organization.

Ensure that there are identified owners for each application, workload, platform, and infrastructure component, and that each process and procedure has an identified owner responsible for its definition, and owners responsible for their performance.

Having understanding of the business value of each component, process, and procedure, of why those resources are in place or activities are performed, and why that ownership exists will inform the actions of your team members. Clearly define the responsibilities of team members so that they may act appropriately and have mechanisms to identify responsibility and ownership. Have mechanisms to request additions, changes, and exceptions so that you do not constrain innovation. Define agreements between teams describing how they work together to support each other and your business outcomes.

Provide support for your team members so that they can be more effective in taking action and supporting your business outcomes. Engaged senior leadership should set expectations and measure success. Senior leadership should be the sponsor, advocate, and driver for the adoption of best practices and evolution of the organization. Empower team members to take action when outcomes are at risk to minimize impact and encourage them to escalate to decision makers and stakeholders when they believe there is a risk so that it can be addressed and incidents avoided. Provide timely, clear, and actionable communications of known risks and planned events so that team members can take timely and appropriate action.

Encourage experimentation to accelerate learning and keep team members interested and engaged. Teams must grow their skill sets to adopt new technologies, and to support changes in demand and responsibilities. Support and encourage this by providing dedicated structured time for learning. Ensure your team members have the resources, both tools and team members, to be successful and scale to support your business outcomes. Leverage cross-organizational diversity to seek multiple unique perspectives. Use this perspective to increase innovation, challenge your assumptions, and reduce the risk of confirmation bias. Grow inclusion, diversity, and accessibility within your teams to gain beneficial perspectives.

If there are external regulatory or compliance requirements that apply to your organization, you should use the resources provided by [AWS Cloud Compliance](#) to help educate your teams so that they can determine the impact on your priorities. The Well-Architected Framework emphasizes learning, measuring, and improving. It provides a consistent approach for you to evaluate architectures, and implement designs that will scale over time. AWS provides the AWS Well-Architected Tool to help you review your approach prior to development, the state of your workloads prior to production, and the state of your workloads in production. You can compare workloads to the latest AWS architectural best practices, monitor their overall status, and gain insight into potential risks. AWS Trusted Advisor is a tool that provides access to a core set of checks that recommend optimizations that may help shape your

priorities. Business and Enterprise Support customers receive access to additional checks focusing on security, reliability, performance, and cost-optimization that can further help shape their priorities.

AWS can help you educate your teams about AWS and its services to increase their understanding of how their choices can have an impact on your workload. You should use the resources provided by AWS Support (AWS Knowledge Center, AWS Discussion Forums, and AWS Support Center) and AWS Documentation to educate your teams. Reach out to AWS Support through AWS Support Center for help with your AWS questions. AWS also shares best practices and patterns that we have learned through the operation of AWS in The Amazon Builders' Library. A wide variety of other useful information is available through the AWS Blog and The Official AWS Podcast. AWS Training and Certification provides some free training through self-paced digital courses on AWS fundamentals. You can also register for instructor-led training to further support the development of your teams' AWS skills.

You should use tools or services that enable you to centrally govern your environments across accounts, such as AWS Organizations, to help manage your operating models. Services like AWS Control Tower expand this management capability by enabling you to define blueprints (supporting your operating models) for the setup of accounts, apply ongoing governance using AWS Organizations, and automate provisioning of new accounts. Managed Services providers such as AWS Managed Services, AWS Managed Services Partners, or Managed Services Providers in the AWS Partner Network, provide expertise implementing cloud environments, and support your security and compliance requirements and business goals. Adding Managed Services to your operating model can save you time and resources, and lets you keep your internal teams lean and focused on strategic outcomes that will differentiate your business, rather than developing new skills and capabilities.

The following questions focus on these considerations for operational excellence. (For a list of operational excellence questions and best practices, see the [Appendix \(p. 48\)](#).)

OPS 1: How do you determine what your priorities are?

Everyone needs to understand their part in enabling business success. Have shared goals in order to set priorities for resources. This will maximize the benefits of your efforts.

OPS 2: How do you structure your organization to support your business outcomes?

Your teams must understand their part in achieving business outcomes. Teams need to understand their roles in the success of other teams, the role of other teams in their success, and have shared goals. Understanding responsibility, ownership, how decisions are made, and who has authority to make decisions will help focus efforts and maximize the benefits from your teams.

OPS 3: How does your organizational culture support your business outcomes?

Provide support for your team members so that they can be more effective in taking action and supporting your business outcome.

You might find that you want to emphasize a small subset of your priorities at some point in time. Use a balanced approach over the long term to ensure the development of needed capabilities and management of risk. Review your priorities regularly and update them as needs change. When responsibility and ownership are undefined or unknown, you are at risk of both not performing necessary action in a timely fashion and of redundant and potentially conflicting efforts emerging to address those needs. Organizational culture has a direct impact on team member job satisfaction and retention. Enable the engagement and capabilities of your team members to enable the success of your business. Experimentation is required for innovation to happen and turn ideas into outcomes. Recognize that an undesired result is a successful experiment that has identified a path that will not lead to success.

Prepare

To prepare for operational excellence, you have to understand your workloads and their expected behaviors. You will then be able to design them to provide insight to their status and build the procedures to support them.

Design your workload so that it provides the information necessary for you to understand its internal state (for example, metrics, logs, events, and traces) across all components in support of observability and investigating issues. Iterate to develop the telemetry necessary to monitor the health of your workload, identify when outcomes are at risk, and enable effective responses. When instrumenting your workload, capture a broad set of information to enable situational awareness (for example, changes in state, user activity, privilege access, utilization counters), knowing that you can use filters to select the most useful information over time.

Adopt approaches that improve the flow of changes into production and that enable refactoring, fast feedback on quality, and bug fixing. These accelerate beneficial changes entering production, limit issues deployed, and enable rapid identification and remediation of issues introduced through deployment activities or discovered in your environments.

Adopt approaches that provide fast feedback on quality and enable rapid recovery from changes that do not have desired outcomes. Using these practices mitigates the impact of issues introduced through the deployment of changes. Plan for unsuccessful changes so that you are able to respond faster if necessary and test and validate the changes you make. Be aware of planned activities in your environments so that you can manage the risk of changes impacting planned activities. Emphasize frequent, small, reversible changes to limit the scope of change. This results in easier troubleshooting and faster remediation with the option to roll back a change. It also means you are able to get the benefit of valuable changes more frequently.

Evaluate the operational readiness of your workload, processes, procedures, and personnel to understand the operational risks related to your workload. You should use a consistent process (including manual or automated checklists) to know when you are ready to go live with your workload or a change. This will also enable you to find any areas that you need to make plans to address. Have runbooks that document your routine activities and playbooks that guide your processes for issue resolution. Understand the benefits and risks to make informed decisions to allow changes to enter production.

AWS enables you to view your entire workload (applications, infrastructure, policy, governance, and operations) as code. This means you can apply the same engineering discipline that you use for application code to every element of your stack and share these across teams or organizations to magnify the benefits of development efforts. Use operations as code in the cloud and the ability to safely experiment to develop your workload, your operations procedures, and practice failure. Using AWS CloudFormation enables you to have consistent, templated, sandbox development, test, and production environments with increasing levels of operations control.

The following questions focus on these considerations for operational excellence.

OPS 4: How do you design your workload so that you can understand its state?

Design your workload so that it provides the information necessary across all components (for example, metrics, logs, and traces) for you to understand its internal state. This enables you to provide effective responses when appropriate.

OPS 5: How do you reduce defects, ease remediation, and improve flow into production?

Adopt approaches that improve flow of changes into production, that enable refactoring, fast feedback on quality, and bug fixing. These accelerate beneficial changes entering production, limit issues

OPS 5: How do you reduce defects, ease remediation, and improve flow into production?

deployed, and enable rapid identification and remediation of issues introduced through deployment activities.

OPS 6: How do you mitigate deployment risks?

Adopt approaches that provide fast feedback on quality and enable rapid recovery from changes that do not have desired outcomes. Using these practices mitigates the impact of issues introduced through the deployment of changes.

OPS 7: How do you know that you are ready to support a workload?

Evaluate the operational readiness of your workload, processes and procedures, and personnel to understand the operational risks related to your workload.

Invest in implementing operations activities as code to maximize the productivity of operations personnel, minimize error rates, and enable automated responses. Use “pre-mortems” to anticipate failure and create procedures where appropriate. Apply metadata using Resource Tags and AWS Resource Groups following a consistent tagging strategy to enable identification of your resources. Tag your resources for organization, cost accounting, access controls, and targeting the execution of automated operations activities. Adopt deployment practices that take advantage of the elasticity of the cloud to facilitate development activities, and pre-deployment of systems for faster implementations. When you make changes to the checklists you use to evaluate your workloads, plan what you will do with live systems that no longer comply.

Operate

Successful operation of a workload is measured by the achievement of business and customer outcomes. Define expected outcomes, determine how success will be measured, and identify metrics that will be used in those calculations to determine if your workload and operations are successful. Operational health includes both the health of the workload and the health and success of the operations activities performed in support of the workload (for example, deployment and incident response). Establish metrics baselines for improvement, investigation, and intervention, collect and analyze your metrics, and then validate your understanding of operations success and how it changes over time. Use collected metrics to determine if you are satisfying customer and business needs, and identify areas for improvement.

Efficient and effective management of operational events is required to achieve operational excellence. This applies to both planned and unplanned operational events. Use established runbooks for well-understood events, and use playbooks to aid in investigation and resolution of issues. Prioritize responses to events based on their business and customer impact. Ensure that if an alert is raised in response to an event, there is an associated process to be executed, with a specifically identified owner. Define in advance the personnel required to resolve an event and include escalation triggers to engage additional personnel, as it becomes necessary, based on urgency and impact. Identify and engage individuals with the authority to make a decision on courses of action where there will be a business impact from an event response not previously addressed.

Communicate the operational status of workloads through dashboards and notifications that are tailored to the target audience (for example, customer, business, developers, operations) so that they may take appropriate action, so that their expectations are managed, and so that they are informed when normal operations resume.

In AWS, you can generate dashboard views of your metrics collected from workloads and natively from AWS. You can leverage CloudWatch or third-party applications to aggregate and present business, workload, and operations level views of operations activities. AWS provides workload insights through logging capabilities including AWS X-Ray, CloudWatch, CloudTrail, and VPC Flow Logs enabling the identification of workload issues in support of root cause analysis and remediation.

The following questions focus on these considerations for operational excellence.

OPS 8: How do you understand the health of your workload?

Define, capture, and analyze workload metrics to gain visibility to workload events so that you can take appropriate action.

OPS 9: How do you understand the health of your operations?

Define, capture, and analyze operations metrics to gain visibility to operations events so that you can take appropriate action.

OPS 10: How do you manage workload and operations events?

Prepare and validate procedures for responding to events to minimize their disruption to your workload.

All of the metrics you collect should be aligned to a business need and the outcomes they support. Develop scripted responses to well-understood events and automate their performance in response to recognizing the event.

Evolve

You must learn, share, and continuously improve to sustain operational excellence. Dedicate work cycles to making continuous incremental improvements. Perform post-incident analysis of all customer impacting events. Identify the contributing factors and preventative action to limit or prevent recurrence. Communicate contributing factors with affected communities as appropriate. Regularly evaluate and prioritize opportunities for improvement (for example, feature requests, issue remediation, and compliance requirements), including both the workload and operations procedures.

Include feedback loops within your procedures to rapidly identify areas for improvement and capture learnings from the execution of operations.

Share lessons learned across teams to share the benefits of those lessons. Analyze trends within lessons learned and perform cross-team retrospective analysis of operations metrics to identify opportunities and methods for improvement. Implement changes intended to bring about improvement and evaluate the results to determine success.

On AWS, you can export your log data to Amazon S3 or send logs directly to Amazon S3 for long-term storage. Using AWS Glue, you can discover and prepare your log data in Amazon S3 for analytics, and store associated metadata in the AWS Glue Data Catalog. Amazon Athena, through its native integration with AWS Glue, can then be used to analyze your log data, querying it using standard SQL. Using a business intelligence tool like Amazon QuickSight, you can visualize, explore, and analyze your data. Discovering trends and events of interest that may drive improvement.

The following question focuses on these considerations for operational excellence.

OPS 11: How do you evolve operations?

Dedicate time and resources for continuous incremental improvement to evolve the effectiveness and efficiency of your operations.

Successful evolution of operations is founded in: frequent small improvements; providing safe environments and time to experiment, develop, and test improvements; and environments in which learning from failures is encouraged. Operations support for sandbox, development, test, and production environments, with increasing level of operational controls, facilitates development and increases the predictability of successful results from changes deployed into production.

Resources

Refer to the following resources to learn more about our best practices for Operational Excellence.

Documentation

- [DevOps and AWS](#)

Whitepaper

- [Operational Excellence Pillar](#)

Video

- [DevOps at Amazon](#)

Security

The Security pillar encompasses the ability to protect data, systems, and assets to take advantage of cloud technologies to improve your security.

The security pillar provides an overview of design principles, best practices, and questions. You can find prescriptive guidance on implementation in the [Security Pillar whitepaper](#).

Topics

- [Design principles \(p. 12\)](#)
- [Definition \(p. 13\)](#)
- [Best practices \(p. 13\)](#)
- [Resources \(p. 18\)](#)

Design principles

There are seven design principles for security in the cloud:

- **Implement a strong identity foundation:** Implement the principle of least privilege and enforce separation of duties with appropriate authorization for each interaction with your AWS resources. Centralize identity management, and aim to eliminate reliance on long-term static credentials.

- **Enable traceability:** Monitor, alert, and audit actions and changes to your environment in real time. Integrate log and metric collection with systems to automatically investigate and take action.
- **Apply security at all layers:** Apply a defense in depth approach with multiple security controls. Apply to all layers (for example, edge of network, VPC, load balancing, every instance and compute service, operating system, application, and code).
- **Automate security best practices:** Automated software-based security mechanisms improve your ability to securely scale more rapidly and cost-effectively. Create secure architectures, including the implementation of controls that are defined and managed as code in version-controlled templates.
- **Protect data in transit and at rest:** Classify your data into sensitivity levels and use mechanisms, such as encryption, tokenization, and access control where appropriate.
- **Keep people away from data:** Use mechanisms and tools to reduce or eliminate the need for direct access or manual processing of data. This reduces the risk of mishandling or modification and human error when handling sensitive data.
- **Prepare for security events:** Prepare for an incident by having incident management and investigation policy and processes that align to your organizational requirements. Run incident response simulations and use tools with automation to increase your speed for detection, investigation, and recovery.

Definition

There are six best practice areas for security in the cloud:

- **Security**
- **Identity and Access Management**
- **Detection**
- **Infrastructure Protection**
- **Data Protection**
- **Incident Response**

Before you architect any workload, you need to put in place practices that influence security. You will want to control who can do what. In addition, you want to be able to identify security incidents, protect your systems and services, and maintain the confidentiality and integrity of data through data protection. You should have a well-defined and practiced process for responding to security incidents. These tools and techniques are important because they support objectives such as preventing financial loss or complying with regulatory obligations.

The AWS Shared Responsibility Model enables organizations that adopt the cloud to achieve their security and compliance goals. Because AWS physically secures the infrastructure that supports our cloud services, as an AWS customer you can focus on using services to accomplish your goals. The AWS Cloud also provides greater access to security data and an automated approach to responding to security events.

Best practices

Topics

- [Security \(p. 14\)](#)
- [Identity and access management \(p. 14\)](#)
- [Detection \(p. 15\)](#)
- [Infrastructure protection \(p. 16\)](#)
- [Data protection \(p. 16\)](#)

- [Incident response \(p. 17\)](#)

Security

To operate your workload securely, you must apply overarching best practices to every area of security. Take requirements and processes that you have defined in operational excellence at an organizational and workload level, and apply them to all areas.

Staying up to date with AWS and industry recommendations and threat intelligence helps you evolve your threat model and control objectives. Automating security processes, testing, and validation allow you to scale your security operations.

The following question focuses on these considerations for security. (For a list of security questions and best practices, see the [Appendix \(p. 117\)](#)).

SEC 1: How do you securely operate your workload?

To operate your workload securely, you must apply overarching best practices to every area of security. Take requirements and processes that you have defined in operational excellence at an organizational and workload level, and apply them to all areas. Staying up to date with recommendations from AWS, industry sources, and threat intelligence helps you evolve your threat model and control objectives. Automating security processes, testing, and validation allow you to scale your security operations.

In AWS, segregating different workloads by account, based on their function and compliance or data sensitivity requirements, is a recommended approach.

Identity and access management

Identity and access management are key parts of an information security program, ensuring that only authorized and authenticated users and components are able to access your resources, and only in a manner that you intend. For example, you should define principals (that is, accounts, users, roles, and services that can perform actions in your account), build out policies aligned with these principals, and implement strong credential management. These privilege-management elements form the core of authentication and authorization.

In AWS, privilege management is primarily supported by the AWS Identity and Access Management (IAM) service, which allows you to control user and programmatic access to AWS services and resources. You should apply granular policies, which assign permissions to a user, group, role, or resource. You also have the ability to require strong password practices, such as complexity level, avoiding re-use, and enforcing multi-factor authentication (MFA). You can use federation with your existing directory service. For workloads that require systems to have access to AWS, IAM enables secure access through roles, instance profiles, identity federation, and temporary credentials.

The following questions focus on these considerations for security.

SEC 2: How do you manage identities for people and machines?

There are two types of identities you need to manage when approaching operating secure AWS workloads. Understanding the type of identity you need to manage and grant access helps you ensure the right identities have access to the right resources under the right conditions.

Human Identities: Your administrators, developers, operators, and end users require an identity to access your AWS environments and applications. These are members of your organization, or external

SEC 2: How do you manage identities for people and machines?

users with whom you collaborate, and who interact with your AWS resources via a web browser, client application, or interactive command-line tools.

Machine Identities: Your service applications, operational tools, and workloads require an identity to make requests to AWS services, for example, to read data. These identities include machines running in your AWS environment such as Amazon EC2 instances or AWS Lambda functions. You may also manage machine identities for external parties who need access. Additionally, you may also have machines outside of AWS that need access to your AWS environment.

SEC 3: How do you manage permissions for people and machines?

Manage permissions to control access to people and machine identities that require access to AWS and your workload. Permissions control who can access what, and under what conditions.

Credentials must not be shared between any user or system. User access should be granted using a least-privilege approach with best practices including password requirements and MFA enforced. Programmatic access including API calls to AWS services should be performed using temporary and limited-privilege credentials such as those issued by the AWS Security Token Service.

AWS provides resources that can help you with Identity and access management. To help learn best practices, explore our hands-on labs on [managing credentials & authentication](#), [controlling human access](#), and [controlling programmatic access](#).

Detection

You can use detective controls to identify a potential security threat or incident. They are an essential part of governance frameworks and can be used to support a quality process, a legal or compliance obligation, and for threat identification and response efforts. There are different types of detective controls. For example, conducting an inventory of assets and their detailed attributes promotes more effective decision making (and lifecycle controls) to help establish operational baselines. You can also use internal auditing, an examination of controls related to information systems, to ensure that practices meet policies and requirements and that you have set the correct automated alerting notifications based on defined conditions. These controls are important reactive factors that can help your organization identify and understand the scope of anomalous activity.

In AWS, you can implement detective controls by processing logs, events, and monitoring that allows for auditing, automated analysis, and alarming. CloudTrail logs, AWS API calls, and CloudWatch provide monitoring of metrics with alarming, and AWS Config provides configuration history. Amazon GuardDuty is a managed threat detection service that continuously monitors for malicious or unauthorized behavior to help you protect your AWS accounts and workloads. Service-level logs are also available, for example, you can use Amazon Simple Storage Service (Amazon S3) to log access requests.

The following question focuses on these considerations for security.

SEC 4: How do you detect and investigate security events?

Capture and analyze events from logs and metrics to gain visibility. Take action on security events and potential threats to help secure your workload.

Log management is important to a Well-Architected workload for reasons ranging from security or forensics to regulatory or legal requirements. It is critical that you analyze logs and respond to

them so that you can identify potential security incidents. AWS provides functionality that makes log management easier to implement by giving you the ability to define a data-retention lifecycle or define where data will be preserved, archived, or eventually deleted. This makes predictable and reliable data handling simpler and more cost effective.

Infrastructure protection

Infrastructure protection encompasses control methodologies, such as defense in depth, necessary to meet best practices and organizational or regulatory obligations. Use of these methodologies is critical for successful, ongoing operations in either the cloud or on-premises.

In AWS, you can implement stateful and stateless packet inspection, either by using AWS-native technologies or by using partner products and services available through the AWS Marketplace. You should use Amazon Virtual Private Cloud (Amazon VPC) to create a private, secured, and scalable environment in which you can define your topology—including gateways, routing tables, and public and private subnets.

The following questions focus on these considerations for security.

SEC 5: How do you protect your network resources?

Any workload that has some form of network connectivity, whether it's the internet or a private network, requires multiple layers of defense to help protect from external and internal network-based threats.

SEC 6: How do you protect your compute resources?

Compute resources in your workload require multiple layers of defense to help protect from external and internal threats. Compute resources include EC2 instances, containers, AWS Lambda functions, database services, IoT devices, and more.

Multiple layers of defense are advisable in any type of environment. In the case of infrastructure protection, many of the concepts and methods are valid across cloud and on-premises models. Enforcing boundary protection, monitoring points of ingress and egress, and comprehensive logging, monitoring, and alerting are all essential to an effective information security plan.

AWS customers are able to tailor, or harden, the configuration of an Amazon Elastic Compute Cloud (Amazon EC2), Amazon Elastic Container Service (Amazon ECS) container, or AWS Elastic Beanstalk instance, and persist this configuration to an immutable Amazon Machine Image (AMI). Then, whether triggered by Auto Scaling or launched manually, all new virtual servers (instances) launched with this AMI receive the hardened configuration.

Data protection

Before architecting any system, foundational practices that influence security should be in place. For example, data classification provides a way to categorize organizational data based on levels of sensitivity, and encryption protects data by way of rendering it unintelligible to unauthorized access. These tools and techniques are important because they support objectives such as preventing financial loss or complying with regulatory obligations.

In AWS, the following practices facilitate protection of data:

- As an AWS customer you maintain full control over your data.

- AWS makes it easier for you to encrypt your data and manage keys, including regular key rotation, which can be easily automated by AWS or maintained by you.
- Detailed logging that contains important content, such as file access and changes, is available.
- AWS has designed storage systems for exceptional resiliency. For example, Amazon S3 Standard, S3 Standard-IA, S3 One Zone-IA, and Amazon Glacier are all designed to provide 99.999999999% durability of objects over a given year. This durability level corresponds to an average annual expected loss of 0.000000001% of objects.
- Versioning, which can be part of a larger data lifecycle management process, can protect against accidental overwrites, deletes, and similar harm.
- AWS never initiates the movement of data between Regions. Content placed in a Region will remain in that Region unless you explicitly enable a feature or leverage a service that provides that functionality.

The following questions focus on these considerations for security.

SEC 7: How do you classify your data?

Classification provides a way to categorize data, based on criticality and sensitivity in order to help you determine appropriate protection and retention controls.

SEC 8: How do you protect your data at rest?

Protect your data at rest by implementing multiple controls, to reduce the risk of unauthorized access or mishandling.

SEC 9: How do you protect your data in transit?

Protect your data in transit by implementing multiple controls to reduce the risk of unauthorized access or loss.

AWS provides multiple means for encrypting data at rest and in transit. We build features into our services that make it easier to encrypt your data. For example, we have implemented server-side encryption (SSE) for Amazon S3 to make it easier for you to store your data in an encrypted form. You can also arrange for the entire HTTPS encryption and decryption process (generally known as SSL termination) to be handled by Elastic Load Balancing (ELB).

Incident response

Even with extremely mature preventive and detective controls, your organization should still put processes in place to respond to and mitigate the potential impact of security incidents. The architecture of your workload strongly affects the ability of your teams to operate effectively during an incident, to isolate or contain systems, and to restore operations to a known good state. Putting in place the tools and access ahead of a security incident, then routinely practicing incident response through game days, will help you ensure that your architecture can accommodate timely investigation and recovery.

In AWS, the following practices facilitate effective incident response:

- Detailed logging is available that contains important content, such as file access and changes.
- Events can be automatically processed and trigger tools that automate responses through the use of AWS APIs.

- You can pre-provision tooling and a “clean room” using AWS CloudFormation. This allows you to carry out forensics in a safe, isolated environment.

The following question focuses on these considerations for security.

SEC 10: How do you anticipate, respond to, and recover from incidents?

Preparation is critical to timely and effective investigation, response to, and recovery from security incidents to help minimize disruption to your organization.

Ensure that you have a way to quickly grant access for your security team, and automate the isolation of instances as well as the capturing of data and state for forensics.

Resources

Refer to the following resources to learn more about our best practices for Security.

Documentation

- [AWS Cloud Security](#)
- [AWS Compliance](#)
- [AWS Security Blog](#)

Whitepaper

- [Security Pillar](#)
- [AWS Security Overview](#)
- [AWS Risk and Compliance](#)

Video

- [AWS Security State of the Union](#)
- [Shared Responsibility Overview](#)

Reliability

The Reliability pillar encompasses the ability of a workload to perform its intended function correctly and consistently when it's expected to. This includes the ability to operate and test the workload through its total lifecycle. This paper provides in-depth, best practice guidance for implementing reliable workloads on AWS.

The reliability pillar provides an overview of design principles, best practices, and questions. You can find prescriptive guidance on implementation in the [Reliability Pillar whitepaper](#).

Topics

- [Design principles \(p. 19\)](#)
- [Definition \(p. 19\)](#)

- [Best practices \(p. 19\)](#)
- [Resources \(p. 23\)](#)

Design principles

There are five design principles for reliability in the cloud:

- **Automatically recover from failure:** By monitoring a workload for key performance indicators (KPIs), you can trigger automation when a threshold is breached. These KPIs should be a measure of business value, not of the technical aspects of the operation of the service. This allows for automatic notification and tracking of failures, and for automated recovery processes that work around or repair the failure. With more sophisticated automation, it's possible to anticipate and remediate failures before they occur.
- **Test recovery procedures:** In an on-premises environment, testing is often conducted to prove that the workload works in a particular scenario. Testing is not typically used to validate recovery strategies. In the cloud, you can test how your workload fails, and you can validate your recovery procedures. You can use automation to simulate different failures or to recreate scenarios that led to failures before. This approach exposes failure pathways that you can test and fix before a real failure scenario occurs, thus reducing risk.
- **Scale horizontally to increase aggregate workload availability:** Replace one large resource with multiple small resources to reduce the impact of a single failure on the overall workload. Distribute requests across multiple, smaller resources to ensure that they don't share a common point of failure.
- **Stop guessing capacity:** A common cause of failure in on-premises workloads is resource saturation, when the demands placed on a workload exceed the capacity of that workload (this is often the objective of denial of service attacks). In the cloud, you can monitor demand and workload utilization, and automate the addition or removal of resources to maintain the optimal level to satisfy demand without over- or under-provisioning. There are still limits, but some quotas can be controlled and others can be managed (see [Manage Service Quotas and Constraints](#)).
- **Manage change in automation:** Changes to your infrastructure should be made using automation. The changes that need to be managed include changes to the automation, which then can be tracked and reviewed.

Definition

There are four best practice areas for reliability in the cloud:

- **Foundations**
- **Workload Architecture**
- **Change Management**
- **Failure Management**

To achieve reliability you must start with the foundations — an environment where service quotas and network topology accommodate the workload. The workload architecture of the distributed system must be designed to prevent and mitigate failures. The workload must handle changes in demand or requirements, and it must be designed to detect failure and automatically heal itself.

Best practices

Topics

- [Foundations \(p. 20\)](#)
- [Workload architecture \(p. 20\)](#)

- [Change management \(p. 21\)](#)
- [Failure management \(p. 22\)](#)

Foundations

Foundational requirements are those whose scope extends beyond a single workload or project. Before architecting any system, foundational requirements that influence reliability should be in place. For example, you must have sufficient network bandwidth to your data center.

With AWS, most of these foundational requirements are already incorporated or can be addressed as needed. The cloud is designed to be nearly limitless, so it's the responsibility of AWS to satisfy the requirement for sufficient networking and compute capacity, leaving you free to change resource size and allocations on demand.

The following questions focus on these considerations for reliability. (For a list of reliability questions and best practices, see the [Appendix \(p. 169\)](#).)

REL 1: How do you manage service quotas and constraints?

For cloud-based workload architectures, there are service quotas (which are also referred to as service limits). These quotas exist to prevent accidentally provisioning more resources than you need and to limit request rates on API operations so as to protect services from abuse. There are also resource constraints, for example, the rate that you can push bits down a fiber-optic cable, or the amount of storage on a physical disk.

REL 2: How do you plan your network topology?

Workloads often exist in multiple environments. These include multiple cloud environments (both publicly accessible and private) and possibly your existing data center infrastructure. Plans must include network considerations such as intra- and inter-system connectivity, public IP address management, private IP address management, and domain name resolution.

For cloud-based workload architectures, there are service quotas (which are also referred to as service limits). These quotas exist to prevent accidentally provisioning more resources than you need and to limit request rates on API operations to protect services from abuse. Workloads often exist in multiple environments. You must monitor and manage these quotas for all workload environments. These include multiple cloud environments (both publicly accessible and private) and may include your existing data center infrastructure. Plans must include network considerations, such as intrasystem and intersystem connectivity, public IP address management, private IP address management, and domain name resolution.

Workload architecture

A reliable workload starts with upfront design decisions for both software and infrastructure. Your architecture choices will impact your workload behavior across all of the Well-Architected pillars. For reliability, there are specific patterns you must follow.

With AWS, workload developers have their choice of languages and technologies to use. AWS SDKs take the complexity out of coding by providing language-specific APIs for AWS services. These SDKs, plus the choice of languages, allow developers to implement the reliability best practices listed here. Developers can also read about and learn from how Amazon builds and operates software in [The Amazon Builders' Library](#).

The following questions focus on these considerations for reliability.

REL 3: How do you design your workload service architecture?

Build highly scalable and reliable workloads using a service-oriented architecture (SOA) or a microservices architecture. Service-oriented architecture (SOA) is the practice of making software components reusable via service interfaces. Microservices architecture goes further to make components smaller and simpler.

REL 4: How do you design interactions in a distributed system to prevent failures?

Distributed systems rely on communications networks to interconnect components, such as servers or services. Your workload must operate reliably despite data loss or latency in these networks. Components of the distributed system must operate in a way that does not negatively impact other components or the workload. These best practices prevent failures and improve mean time between failures (MTBF).

REL 5: How do you design interactions in a distributed system to mitigate or withstand failures?

Distributed systems rely on communications networks to interconnect components (such as servers or services). Your workload must operate reliably despite data loss or latency over these networks. Components of the distributed system must operate in a way that does not negatively impact other components or the workload. These best practices enable workloads to withstand stresses or failures, more quickly recover from them, and mitigate the impact of such impairments. The result is improved mean time to recovery (MTTR).

Change management

Changes to your workload or its environment must be anticipated and accommodated to achieve reliable operation of the workload. Changes include those imposed on your workload, such as spikes in demand, as well as those from within, such as feature deployments and security patches.

Using AWS, you can monitor the behavior of a workload and automate the response to KPIs. For example, your workload can add additional servers as a workload gains more users. You can control who has permission to make workload changes and audit the history of these changes.

The following questions focus on these considerations for reliability.

REL 6: How do you monitor workload resources?

Logs and metrics are powerful tools to gain insight into the health of your workload. You can configure your workload to monitor logs and metrics and send notifications when thresholds are crossed or significant events occur. Monitoring enables your workload to recognize when low-performance thresholds are crossed or failures occur, so it can recover automatically in response.

REL 7: How do you design your workload to adapt to changes in demand?

A scalable workload provides elasticity to add or remove resources automatically so that they closely match the current demand at any given point in time.

REL 8: How do you implement change?

Controlled changes are necessary to deploy new functionality, and to ensure that the workloads and the operating environment are running known software and can be patched or replaced in a predictable manner. If these changes are uncontrolled, then it makes it difficult to predict the effect of these changes, or to address issues that arise because of them.

When you architect a workload to automatically add and remove resources in response to changes in demand, this not only increases reliability but also ensures that business success doesn't become a burden. With monitoring in place, your team will be automatically alerted when KPIs deviate from expected norms. Automatic logging of changes to your environment allows you to audit and quickly identify actions that might have impacted reliability. Controls on change management ensure that you can enforce the rules that deliver the reliability you need.

Failure management

In any system of reasonable complexity, it is expected that failures will occur. Reliability requires that your workload be aware of failures as they occur and take action to avoid impact on availability. Workloads must be able to both withstand failures and automatically repair issues.

With AWS, you can take advantage of automation to react to monitoring data. For example, when a particular metric crosses a threshold, you can trigger an automated action to remedy the problem. Also, rather than trying to diagnose and fix a failed resource that is part of your production environment, you can replace it with a new one and carry out the analysis on the failed resource out of band. Since the cloud enables you to stand up temporary versions of a whole system at low cost, you can use automated testing to verify full recovery processes.

The following questions focus on these considerations for reliability.

REL 9: How do you back up data?

Back up data, applications, and configuration to meet your requirements for recovery time objectives (RTO) and recovery point objectives (RPO).

REL 10: How do you use fault isolation to protect your workload?

Fault isolated boundaries limit the effect of a failure within a workload to a limited number of components. Components outside of the boundary are unaffected by the failure. Using multiple fault isolated boundaries, you can limit the impact on your workload.

REL 11: How do you design your workload to withstand component failures?

Workloads with a requirement for high availability and low mean time to recovery (MTTR) must be architected for resiliency.

REL 12: How do you test reliability?

After you have designed your workload to be resilient to the stresses of production, testing is the only way to ensure that it will operate as designed, and deliver the resiliency you expect.

REL 13: How do you plan for disaster recovery (DR)?

Having backups and redundant workload components in place is the start of your DR strategy. [RTO and RPO are your objectives](#) for restoration of your workload. Set these based on business needs. Implement a strategy to meet these objectives, considering locations and function of workload resources and data. The probability of disruption and cost of recovery are also key factors that help to inform the business value of providing disaster recovery for a workload.

Regularly back up your data and test your backup files to ensure that you can recover from both logical and physical errors. A key to managing failure is the frequent and automated testing of workloads to cause failure, and then observe how they recover. Do this on a regular schedule and ensure that such testing is also triggered after significant workload changes. Actively track KPIs, as well as the recovery time objective (RTO) and recovery point objective (RPO), to assess a workload's resiliency (especially under failure-testing scenarios). Tracking KPIs will help you identify and mitigate single points of failure. The objective is to thoroughly test your workload-recovery processes so that you are confident that you can recover all your data and continue to serve your customers, even in the face of sustained problems. Your recovery processes should be as well exercised as your normal production processes.

Resources

Refer to the following resources to learn more about our best practices for Reliability.

Documentation

- [AWS Documentation](#)
- [AWS Global Infrastructure](#)
- [AWS Auto Scaling: How Scaling Plans Work](#)
- [What Is AWS Backup?](#)

Whitepaper

- [Reliability Pillar: AWS Well-Architected](#)
- [Implementing Microservices on AWS](#)

Performance efficiency

The Performance Efficiency pillar includes the ability to use computing resources efficiently to meet system requirements, and to maintain that efficiency as demand changes and technologies evolve.

The performance efficiency pillar provides an overview of design principles, best practices, and questions. You can find prescriptive guidance on implementation in the [Performance Efficiency Pillar whitepaper](#).

Topics

- [Design principles \(p. 24\)](#)
- [Definition \(p. 24\)](#)
- [Best practices \(p. 24\)](#)
- [Resources \(p. 29\)](#)

Design principles

There are five design principles for performance efficiency in the cloud:

- **Democratize advanced technologies:** Make advanced technology implementation easier for your team by delegating complex tasks to your cloud vendor. Rather than asking your IT team to learn about hosting and running a new technology, consider consuming the technology as a service. For example, NoSQL databases, media transcoding, and machine learning are all technologies that require specialized expertise. In the cloud, these technologies become services that your team can consume, allowing your team to focus on product development rather than resource provisioning and management.
- **Go global in minutes:** Deploying your workload in multiple AWS Regions around the world allows you to provide lower latency and a better experience for your customers at minimal cost.
- **Use serverless architectures:** Serverless architectures remove the need for you to run and maintain physical servers for traditional compute activities. For example, serverless storage services can act as static websites (removing the need for web servers) and event services can host code. This removes the operational burden of managing physical servers, and can lower transactional costs because managed services operate at cloud scale.
- **Experiment more often:** With virtual and automatable resources, you can quickly carry out comparative testing using different types of instances, storage, or configurations.
- **Consider mechanical sympathy:** Understand how cloud services are consumed and always use the technology approach that aligns best with your workload goals. For example, consider data access patterns when you select database or storage approaches.

Definition

There are four best practice areas for performance efficiency in the cloud:

- **Selection**
- **Review**
- **Monitoring**
- **Tradeoffs**

Take a data-driven approach to building a high-performance architecture. Gather data on all aspects of the architecture, from the high-level design to the selection and configuration of resource types.

Reviewing your choices on a regular basis ensures that you are taking advantage of the continually evolving AWS Cloud. Monitoring ensures that you are aware of any deviance from expected performance. Make trade-offs in your architecture to improve performance, such as using compression or caching, or relaxing consistency requirements.

Best practices

Topics

- [Selection \(p. 25\)](#)
- [Review \(p. 28\)](#)
- [Monitoring \(p. 28\)](#)
- [Tradeoffs \(p. 29\)](#)

Selection

The optimal solution for a particular workload varies, and solutions often combine multiple approaches. Well-architected workloads use multiple solutions and enable different features to improve performance.

AWS resources are available in many types and configurations, which makes it easier to find an approach that closely matches your workload needs. You can also find options that are not easily achievable with on-premises infrastructure. For example, a managed service such as Amazon DynamoDB provides a fully managed NoSQL database with single-digit millisecond latency at any scale.

The following question focuses on these considerations for performance efficiency. (For a list of performance efficiency questions and best practices, see the [Appendix \(p. 278\)](#).)

PERF 1: How do you select the best performing architecture?

Often, multiple approaches are required for optimal performance across a workload. Well-architected systems use multiple solutions and features to improve performance.

Use a data-driven approach to select the patterns and implementation for your architecture and achieve a cost effective solution. AWS Solutions Architects, AWS Reference Architectures, and AWS Partner Network (APN) partners can help you select an architecture based on industry knowledge, but data obtained through benchmarking or load testing will be required to optimize your architecture.

Your architecture will likely combine a number of different architectural approaches (for example, event-driven, ETL, or pipeline). The implementation of your architecture will use the AWS services that are specific to the optimization of your architecture's performance. In the following sections we discuss the four main resource types to consider (compute, storage, database, and network).

Compute

Selecting compute resources that meet your requirements, performance needs, and provide great efficiency of cost and effort will enable you to accomplish more with the same number of resources. When evaluating compute options, be aware of your requirements for workload performance and cost requirements and use this to make informed decisions.

In AWS, compute is available in three forms: instances, containers, and functions:

- **Instances** are virtualized servers, allowing you to change their capabilities with a button or an API call. Because resource decisions in the cloud aren't fixed, you can experiment with different server types. At AWS, these virtual server instances come in different families and sizes, and they offer a wide variety of capabilities, including solid-state drives (SSDs) and graphics processing units (GPUs).
- **Containers** are a method of operating system virtualization that allow you to run an application and its dependencies in resource-isolated processes. AWS Fargate is serverless compute for containers or Amazon EC2 can be used if you need control over the installation, configuration, and management of your compute environment. You can also choose from multiple container orchestration platforms: Amazon Elastic Container Service (ECS) or Amazon Elastic Kubernetes Service (EKS).
- **Functions** abstract the execution environment from the code you want to execute. For example, AWS Lambda allows you to execute code without running an instance.

The following question focuses on these considerations for performance efficiency.

PERF 2: How do you select your compute solution?

The optimal compute solution for a workload varies based on application design, usage patterns, and configuration settings. Architectures can use different compute solutions for various components

PERF 2: How do you select your compute solution?

and enable different features to improve performance. Selecting the wrong compute solution for an architecture can lead to lower performance efficiency.

When architecting your use of compute you should take advantage of the elasticity mechanisms available to ensure you have sufficient capacity to sustain performance as demand changes.

Storage

Cloud storage is a critical component of cloud computing, holding the information used by your workload. Cloud storage is typically more reliable, scalable, and secure than traditional on-premises storage systems. Select from object, block, and file storage services as well as cloud data migration options for your workload.

In AWS, storage is available in three forms: object, block, and file:

- **Object Storage** provides a scalable, durable platform to make data accessible from any internet location for user-generated content, active archive, serverless computing, Big Data storage or backup and recovery. Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance. Amazon S3 is designed for 99.999999999% (11 9's) of durability, and stores data for millions of applications for companies all around the world.
- **Block Storage** provides highly available, consistent, low-latency block storage for each virtual host and is analogous to direct-attached storage (DAS) or a Storage Area Network (SAN). Amazon Elastic Block Store (Amazon EBS) is designed for workloads that require persistent storage accessible by EC2 instances that helps you tune applications with the right storage capacity, performance and cost.
- **File Storage** provides access to a shared file system across multiple systems. File storage solutions like Amazon Elastic File System (EFS) are ideal for use cases such as large content repositories, development environments, media stores, or user home directories. Amazon FSx makes it easy and cost effective to launch and run popular file systems so you can leverage the rich feature sets and fast performance of widely used open source and commercially-licensed file systems.

The following question focuses on these considerations for performance efficiency.

PERF 3: How do you select your storage solution?

The optimal storage solution for a system varies based on the kind of access method (block, file, or object), patterns of access (random or sequential), required throughput, frequency of access (online, offline, archival), frequency of update (WORM, dynamic), and availability and durability constraints. Well-architected systems use multiple storage solutions and enable different features to improve performance and use resources efficiently.

When you select a storage solution, ensuring that it aligns with your access patterns will be critical to achieving the performance you want.

Database

The cloud offers purpose-built database services that address different problems presented by your workload. You can choose from many purpose-built database engines including relational, key-value, document, in-memory, graph, time series, and ledger databases. By picking the best database to solve a specific problem (or a group of problems), you can break away from restrictive one-size-fits-all monolithic databases and focus on building applications to meet the performance needs of your customers.

In AWS you can choose from multiple purpose-built database engines including relational, key-value, document, in-memory, graph, time series, and ledger databases. With AWS databases, you don't need to worry about database management tasks such as server provisioning, patching, setup, configuration, backups, or recovery. AWS continuously monitors your clusters to keep your workloads up and running with self-healing storage and automated scaling, so that you can focus on higher value application development.

The following question focuses on these considerations for performance efficiency.

PERF 4: How do you select your database solution?

The optimal database solution for a system varies based on requirements for availability, consistency, partition tolerance, latency, durability, scalability, and query capability. Many systems use different database solutions for various subsystems and enable different features to improve performance. Selecting the wrong database solution and features for a system can lead to lower performance efficiency.

Your workload's database approach has a significant impact on performance efficiency. It's often an area that is chosen according to organizational defaults rather than through a data-driven approach. As with storage, it is critical to consider the access patterns of your workload, and also to consider if other non-database solutions could solve the problem more efficiently (such as using graph, time series, or in-memory storage database).

Network

Since the network is between all workload components, it can have great impacts, both positive and negative, on workload performance and behavior. There are also workloads that are heavily dependent on network performance such as High Performance Computing (HPC) where deep network understanding is important to increase cluster performance. You must determine the workload requirements for bandwidth, latency, jitter, and throughput.

On AWS, networking is virtualized and is available in a number of different types and configurations. This makes it easier to match your networking methods with your needs. AWS offers product features (for example, Enhanced Networking, Amazon EBS-optimized instances, Amazon S3 transfer acceleration, and dynamic Amazon CloudFront) to optimize network traffic. AWS also offers networking features (for example, Amazon Route 53 latency routing, Amazon VPC endpoints, AWS Direct Connect, and AWS Global Accelerator) to reduce network distance or jitter.

The following question focuses on these considerations for performance efficiency.

PERF 5: How do you configure your networking solution?

The optimal network solution for a workload varies based on latency, throughput requirements, jitter, and bandwidth. Physical constraints, such as user or on-premises resources, determine location options. These constraints can be offset with edge locations or resource placement.

You must consider location when deploying your network. You can choose to place resources close to where they will be used to reduce distance. Use networking metrics to make changes to networking configuration as the workload evolves. By taking advantage of Regions, placement groups, and edge services, you can significantly improve performance. Cloud based networks can be quickly re-built or modified, so evolving your network architecture over time is necessary to maintain performance efficiency.

Review

Cloud technologies are rapidly evolving and you must ensure that workload components are using the latest technologies and approaches to continually improve performance. You must continually evaluate and consider changes to your workload components to ensure you are meeting its performance and cost objectives. New technologies, such as machine learning and artificial intelligence (AI), can allow you to reimagine customer experiences and innovate across all of your business workloads.

Take advantage of the continual innovation at AWS driven by customer need. We release new Regions, edge locations, services, and features regularly. Any of these releases could positively improve the performance efficiency of your architecture.

The following question focuses on these considerations for performance efficiency.

PERF 6: How do you evolve your workload to take advantage of new releases?

When architecting workloads, there are finite options that you can choose from. However, over time, new technologies and approaches become available that could improve the performance of your workload.

Architectures performing poorly are usually the result of a non-existent or broken performance review process. If your architecture is performing poorly, implementing a performance review process will allow you to apply Deming's plan-do-check-act (PDCA) cycle to drive iterative improvement.

Monitoring

After you implement your workload, you must monitor its performance so that you can remediate any issues before they impact your customers. Monitoring metrics should be used to raise alarms when thresholds are breached.

Amazon CloudWatch is a monitoring and observability service that provides you with data and actionable insights to monitor your workload, respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health. CloudWatch collects monitoring and operational data in the form of logs, metrics, and events from workloads that run on AWS and on-premises servers. AWS X-Ray helps developers analyze and debug production, distributed applications. With AWS X-Ray, you can glean insights into how your application is performing and discover root causes and identify performance bottlenecks. You can use these insights to react quickly and keep your workload running smoothly.

The following question focuses on these considerations for performance efficiency.

PERF 7: How do you monitor your resources to ensure they are performing?

System performance can degrade over time. Monitor system performance to identify degradation and remediate internal or external factors, such as the operating system or application load.

Ensuring that you do not see false positives is key to an effective monitoring solution. Automated triggers avoid human error and can reduce the time it takes to fix problems. Plan for game days, where simulations are conducted in the production environment, to test your alarm solution and ensure that it correctly recognizes issues.

Tradeoffs

When you architect solutions, think about tradeoffs to ensure an optimal approach. Depending on your situation, you could trade consistency, durability, and space for time or latency, to deliver higher performance.

Using AWS, you can go global in minutes and deploy resources in multiple locations across the globe to be closer to your end users. You can also dynamically add readonly replicas to information stores (such as database systems) to reduce the load on the primary database.

The following question focuses on these considerations for performance efficiency.

PERF 8: How do you use tradeoffs to improve performance?
When architecting solutions, determining tradeoffs enables you to select an optimal approach. Often you can improve performance by trading consistency, durability, and space for time and latency.

As you make changes to the workload, collect and evaluate metrics to determine the impact of those changes. Measure the impacts to the system and to the end-user to understand how your trade-offs impact your workload. Use a systematic approach, such as load testing, to explore whether the tradeoff improves performance.

Resources

Refer to the following resources to learn more about our best practices for Performance Efficiency.

Documentation

- [Amazon S3 Performance Optimization](#)
- [Amazon EBS Volume Performance](#)

Whitepaper

- [Performance Efficiency Pillar](#)

Video

- [AWS re:Invent 2019: Amazon EC2 foundations \(CMP211-R2\)](#)
- [AWS re:Invent 2019: Leadership session: Storage state of the union \(STG201-L\)](#)
- [AWS re:Invent 2019: Leadership session: AWS purpose-built databases \(DAT209-L\)](#)
- [AWS re:Invent 2019: Connectivity to AWS and hybrid AWS network architectures \(NET317-R1\)](#)
- [AWS re:Invent 2019: Powering next-gen Amazon EC2: Deep dive into the Nitro system \(CMP303-R2\)](#)
- [AWS re:Invent 2019: Scaling up to your first 10 million users \(ARC211-R\)](#)

Cost optimization

The Cost Optimization pillar includes the ability to run systems to deliver business value at the lowest price point.

The cost optimization pillar provides an overview of design principles, best practices, and questions. You can find prescriptive guidance on implementation in the [Cost Optimization Pillar whitepaper](#).

Topics

- [Design principles \(p. 30\)](#)
- [Definition \(p. 30\)](#)
- [Best practices \(p. 31\)](#)
- [Resources \(p. 34\)](#)

Design principles

There are five design principles for cost optimization in the cloud:

- **Implement Cloud Financial Management:** To achieve financial success and accelerate business value realization in the cloud, you need to invest in Cloud Financial Management /Cost Optimization. Your organization needs to dedicate time and resources to build capability in this new domain of technology and usage management. Similar to your Security or Operational Excellence capability, you need to build capability through knowledge building, programs, resources, and processes to become a cost-efficient organization.
- **Adopt a consumption model:** Pay only for the computing resources that you require and increase or decrease usage depending on business requirements, not by using elaborate forecasting. For example, development and test environments are typically only used for eight hours a day during the work week. You can stop these resources when they are not in use for a potential cost savings of 75% (40 hours versus 168 hours).
- **Measure overall efficiency:** Measure the business output of the workload and the costs associated with delivering it. Use this measure to know the gains you make from increasing output and reducing costs.
- **Stop spending money on undifferentiated heavy lifting:** AWS does the heavy lifting of data center operations like racking, stacking, and powering servers. It also removes the operational burden of managing operating systems and applications with managed services. This allows you to focus on your customers and business projects rather than on IT infrastructure.
- **Analyze and attribute expenditure:** The cloud makes it easier to accurately identify the usage and cost of systems, which then allows transparent attribution of IT costs to individual workload owners. This helps measure return on investment (ROI) and gives workload owners an opportunity to optimize their resources and reduce costs.

Definition

There are five best practice areas for cost optimization in the cloud:

- **Practice Cloud Financial Management**
- **Expenditure and usage awareness**
- **Cost-effective resources**
- **Manage demand and supply resources**
- **Optimize over time**

As with the other pillars within the Well-Architected Framework, there are tradeoffs to consider, for example, whether to optimize for speed-to-market or for cost. In some cases, it's best to optimize for speed—going to market quickly, shipping new features, or simply meeting a deadline—rather than investing in up-front cost optimization. Design decisions are sometimes directed by haste rather than data, and the temptation always exists to overcompensate “just in case” rather than spend time benchmarking for the most cost-optimal deployment. This might lead to over-provisioned and under-optimized deployments. However, this is a reasonable choice when you need to “lift and shift” resources from your on-premises environment to the cloud and then optimize afterwards. Investing the right

amount of effort in a cost optimization strategy up front allows you to realize the economic benefits of the cloud more readily by ensuring a consistent adherence to best practices and avoiding unnecessary over provisioning. The following sections provide techniques and best practices for both the initial and ongoing implementation of Cloud Financial Management and cost optimization of your workloads.

Best practices

Topics

- [Practice Cloud Financial Management \(p. 31\)](#)
- [Expenditure and usage awareness \(p. 31\)](#)
- [Cost-effective resources \(p. 32\)](#)
- [Manage demand and supply resources \(p. 33\)](#)
- [Optimize over time \(p. 34\)](#)

Practice Cloud Financial Management

With the adoption of cloud, technology teams innovate faster due to shortened approval, procurement, and infrastructure deployment cycles. A new approach to financial management in the cloud is required to realize business value and financial success. This approach is Cloud Financial Management, and builds capability across your organization by implementing organizational wide knowledge building, programs, resources, and processes.

Many organizations are composed of many different units with different priorities. The ability to align your organization to an agreed set of financial objectives, and provide your organization the mechanisms to meet them, will create a more efficient organization. A capable organization will innovate and build faster, be more agile and adjust to any internal or external factors.

In AWS you can use Cost Explorer, and optionally Amazon Athena and Amazon QuickSight with the Cost and Usage Report (CUR), to provide cost and usage awareness throughout your organization. AWS Budgets provides proactive notifications for cost and usage. The AWS blogs provide information on new services and features to ensure you keep up to date with new service releases.

The following question focuses on these considerations for cost optimization. (For a list of cost optimization questions and best practices, see the [Appendix \(p. 346\)](#)).

COST 1: How do you implement cloud financial management?

Implementing Cloud Financial Management enables organizations to realize business value and financial success as they optimize their cost and usage and scale on AWS.

When building a cost optimization function, use members and supplement the team with experts in CFM and cost optimization. Existing team members will understand how the organization currently functions and how to rapidly implement improvements. Also consider including people with supplementary or specialist skill sets, such as analytics and project management.

When implementing cost awareness in your organization, improve or build on existing programs and processes. It is much faster to add to what exists than to build new processes and programs. This will result in achieving outcomes much faster.

Expenditure and usage awareness

The increased flexibility and agility that the cloud enables encourages innovation and fast-paced development and deployment. It eliminates the manual processes and time associated with provisioning

on-premises infrastructure, including identifying hardware specifications, negotiating price quotations, managing purchase orders, scheduling shipments, and then deploying the resources. However, the ease of use and virtually unlimited on-demand capacity requires a new way of thinking about expenditures.

Many businesses are composed of multiple systems run by various teams. The capability to attribute resource costs to the individual organization or product owners drives efficient usage behavior and helps reduce waste. Accurate cost attribution allows you to know which products are truly profitable, and allows you to make more informed decisions about where to allocate budget.

In AWS, you create an account structure with AWS Organizations or AWS Control Tower, which provides separation and assists in allocation of your costs and usage. You can also use resource tagging to apply business and organization information to your usage and cost. Use AWS Cost Explorer for visibility into your cost and usage, or create customized dashboards and analytics with Amazon Athena and Amazon QuickSight. Controlling your cost and usage is done by notifications through AWS Budgets, and controls using AWS Identity and Access Management (IAM), and Service Quotas.

The following questions focus on these considerations for cost optimization.

COST 2: How do you govern usage?

Establish policies and mechanisms to ensure that appropriate costs are incurred while objectives are achieved. By employing a checks-and-balances approach, you can innovate without overspending.

COST 3: How do you monitor usage and cost?

Establish policies and procedures to monitor and appropriately allocate your costs. This allows you to measure and improve the cost efficiency of this workload.

COST 4: How do you decommission resources?

Implement change control and resource management from project inception to end-of-life. This ensures you shut down or terminate unused resources to reduce waste.

You can use cost allocation tags to categorize and track your AWS usage and costs. When you apply tags to your AWS resources (such as EC2 instances or S3 buckets), AWS generates a cost and usage report with your usage and your tags. You can apply tags that represent organization categories (such as cost centers, workload names, or owners) to organize your costs across multiple services.

Ensure you use the right level of detail and granularity in cost and usage reporting and monitoring. For high level insights and trends, use daily granularity with AWS Cost Explorer. For deeper analysis and inspection use hourly granularity in AWS Cost Explorer, or Amazon Athena and Amazon QuickSight with the Cost and Usage Report (CUR) at an hourly granularity.

Combining tagged resources with entity lifecycle tracking (employees, projects) makes it possible to identify orphaned resources or projects that are no longer generating value to the organization and should be decommissioned. You can set up billing alerts to notify you of predicted overspending.

Cost-effective resources

Using the appropriate instances and resources for your workload is key to cost savings. For example, a reporting process might take five hours to run on a smaller server but one hour to run on a larger server that is twice as expensive. Both servers give you the same outcome, but the smaller server incurs more cost over time.

A well-architected workload uses the most cost-effective resources, which can have a significant and positive economic impact. You also have the opportunity to use managed services to reduce costs. For example, rather than maintaining servers to deliver email, you can use a service that charges on a per-message basis.

AWS offers a variety of flexible and cost-effective pricing options to acquire instances from Amazon EC2 and other services in a way that best fits your needs. *On-Demand Instances* allow you to pay for compute capacity by the hour, with no minimum commitments required. *Savings Plans and Reserved Instances* offer savings of up to 75% off On-Demand pricing. With Spot Instances, you can leverage unused Amazon EC2 capacity and offer savings of up to 90% off On-Demand pricing. *Spot Instances* are appropriate where the system can tolerate using a fleet of servers where individual servers can come and go dynamically, such as stateless web servers, batch processing, or when using HPC and big data.

Appropriate service selection can also reduce usage and costs; such as CloudFront to minimize data transfer, or completely eliminate costs, such as utilizing Amazon Aurora on RDS to remove expensive database licensing costs.

The following questions focus on these considerations for cost optimization.

COST 5: How do you evaluate cost when you select services?

Amazon EC2, Amazon EBS, and Amazon S3 are building-block AWS services. Managed services, such as Amazon RDS and Amazon DynamoDB, are higher level, or application level, AWS services. By selecting the appropriate building blocks and managed services, you can optimize this workload for cost. For example, using managed services, you can reduce or remove much of your administrative and operational overhead, freeing you to work on applications and business-related activities.

COST 6: How do you meet cost targets when you select resource type, size and number?

Ensure that you choose the appropriate resource size and number of resources for the task at hand. You minimize waste by selecting the most cost effective type, size, and number.

COST 7: How do you use pricing models to reduce cost?

Use the pricing model that is most appropriate for your resources to minimize expense.

COST 8: How do you plan for data transfer charges?

Ensure that you plan and monitor data transfer charges so that you can make architectural decisions to minimize costs. A small yet effective architectural change can drastically reduce your operational costs over time.

By factoring in cost during service selection, and using tools such as Cost Explorer and AWS Trusted Advisor to regularly review your AWS usage, you can actively monitor your utilization and adjust your deployments accordingly.

Manage demand and supply resources

When you move to the cloud, you pay only for what you need. You can supply resources to match the workload demand at the time they're needed, this eliminates the need for costly and wasteful over

provisioning. You can also modify the demand, using a throttle, buffer, or queue to smooth the demand and serve it with less resources resulting in a lower cost, or process it at a later time with a batch service.

In AWS, you can automatically provision resources to match the workload demand. Auto Scaling using demand or time-based approaches allow you to add and remove resources as needed. If you can anticipate changes in demand, you can save more money and ensure your resources match your workload needs. You can use Amazon API Gateway to implement throttling, or Amazon SQS to implementing a queue in your workload. These will both allow you to modify the demand on your workload components.

The following question focuses on these considerations for cost optimization.

COST 9: How do you manage demand, and supply resources?

For a workload that has balanced spend and performance, ensure that everything you pay for is used and avoid significantly underutilizing instances. A skewed utilization metric in either direction has an adverse impact on your organization, in either operational costs (degraded performance due to over-utilization), or wasted AWS expenditures (due to over-provisioning).

When designing to modify demand and supply resources, actively think about the patterns of usage, the time it takes to provision new resources, and the predictability of the demand pattern. When managing demand, ensure you have a correctly sized queue or buffer, and that you are responding to workload demand in the required amount of time.

Optimize over time

As AWS releases new services and features, it's a best practice to review your existing architectural decisions to ensure they continue to be the most cost effective. As your requirements change, be aggressive in decommissioning resources, entire services, and systems that you no longer require.

Implementing new features or resource types can optimize your workload incrementally, while minimizing the effort required to implement the change. This provides continual improvements in efficiency over time and ensures you remain on the most updated technology to reduce operating costs. You can also replace or add new components to the workload with new services. This can provide significant increases in efficiency, so it's essential to regularly review your workload, and implement new services and features.

The following question focuses on these considerations for cost optimization.

COST 10: How do you evaluate new services?

As AWS releases new services and features, it's a best practice to review your existing architectural decisions to ensure they continue to be the most cost effective.

When regularly reviewing your deployments, assess how newer services can help save you money. For example, Amazon Aurora on RDS can reduce costs for relational databases. Using serverless such as Lambda can remove the need to operate and manage instances to run code.

Resources

Refer to the following resources to learn more about our best practices for Cost Optimization.

Documentation

- [AWS Documentation](#)

Whitepaper

- [Cost Optimization Pillar](#)

Sustainability

The Sustainability pillar focuses on environmental impacts, especially energy consumption and efficiency, since they are important levers for architects to inform direct action to reduce resource usage. You can find prescriptive guidance on implementation in the [Sustainability Pillar whitepaper](#).

Topics

- [Design principles \(p. 35\)](#)
- [Definition \(p. 36\)](#)
- [Best practices \(p. 36\)](#)

Design principles

There are six design principles for sustainability in the cloud:

- **Understand your impact:** Measure the impact of your cloud workload and model the future impact of your workload. Include all sources of impact, including impacts resulting from customer use of your products, and impacts resulting from their eventual decommissioning and retirement. Compare the productive output with the total impact of your cloud workloads by reviewing the resources and emissions required per unit of work. Use this data to establish key performance indicators (KPIs), evaluate ways to improve productivity while reducing impact, and estimate the impact of proposed changes over time.
- **Establish sustainability goals:** For each cloud workload, establish long-term sustainability goals such as reducing the compute and storage resources required per transaction. Model the return on investment of sustainability improvements for existing workloads, and give owners the resources they need to invest in sustainability goals. Plan for growth, and architect your workloads so that growth results in reduced impact intensity measured against an appropriate unit, such as per user or per transaction. Goals help you support the wider sustainability goals of your business or organization, identify regressions, and prioritize areas of potential improvement.
- **Maximize utilization:** Right-size workloads and implement efficient design to ensure high utilization and maximize the energy efficiency of the underlying hardware. Two hosts running at 30% utilization are less efficient than one host running at 60% due to baseline power consumption per host. At the same time, eliminate or minimize idle resources, processing, and storage to reduce the total energy required to power your workload.
- **Anticipate and adopt new, more efficient hardware and software offerings:** Support the upstream improvements your partners and suppliers make to help you reduce the impact of your cloud workloads. Continually monitor and evaluate new, more efficient hardware and software offerings. Design for flexibility to allow for the rapid adoption of new efficient technologies.
- **Use managed services:** Sharing services across a broad customer base helps maximize resource utilization, which reduces the amount of infrastructure needed to support cloud workloads. For example, customers can share the impact of common data center components like power and networking by migrating workloads to the AWS Cloud and adopting managed services, such as AWS Fargate for serverless containers, where AWS operates at scale and is responsible for their efficient operation. Use managed services that can help minimize your impact, such as automatically moving infrequently accessed data to cold storage with Amazon S3 Lifecycle configurations or Amazon EC2 Auto Scaling to adjust capacity to meet demand.
- **Reduce the downstream impact of your cloud workloads:** Reduce the amount of energy or resources required to use your services. Reduce or eliminate the need for customers to upgrade their devices to

use your services. Test using device farms to understand expected impact and test with customers to understand the actual impact from using your services.

Definition

There are six best practice areas for sustainability in the cloud:

- Region selection
- User behavior patterns
- Software and architecture patterns
- Data patterns
- Hardware patterns
- Development and deployment process

Sustainability in the cloud is a continuous effort focused primarily on energy reduction and efficiency across all components of a workload by achieving the maximum benefit from the resources provisioned and minimizing the total resources required. This effort can range from the initial selection of an efficient programming language, adoption of modern algorithms, use of efficient data storage techniques, deploying to correctly sized and efficient compute infrastructure, and minimizing requirements for high-powered end-user hardware.

Best practices

Topics

- [Region selection \(p. 36\)](#)
- [User behavior patterns \(p. 36\)](#)
- [Software and architecture patterns \(p. 37\)](#)
- [Data patterns \(p. 38\)](#)
- [Hardware patterns \(p. 39\)](#)
- [Development and deployment patterns \(p. 40\)](#)
- [Resources \(p. 40\)](#)

Region selection

Choose Regions where you will implement your workloads based on both your business requirements and sustainability goals.

The following question focuses on these considerations for sustainability. (For a list of sustainability questions and best practices, see the [Appendix \(p. 392\)](#).)

SUS 1: How do you select Regions to support your sustainability goals?

Choose Regions near Amazon renewable energy projects and Regions where the grid has a published carbon intensity that is lower than other locations (or Regions).

User behavior patterns

The way users consume your workloads and other resources can help you identify improvements to meet sustainability goals. Scale infrastructure to continually match user load and ensure that only the

minimum resources required to support users are deployed. Align service levels to customer needs. Position resources to limit the network required for users to consume them. Remove existing, unused assets. Identify created assets that are unused and stop generating them. Provide your team members with devices that support their needs with minimized sustainability impact.

The following question focuses on this considerations for sustainability:

SUS 2: How do you take advantage of user behavior patterns to support your sustainability goals?

The way users consume your workloads and other resources can help you identify improvements to meet sustainability goals. Scale infrastructure to continually match user load and ensure that only the minimum resources required to support users are deployed. Align service levels to customer needs. Position resources to limit the network required for users to consume them. Remove existing, unused assets. Identify created assets that are unused and stop generating them. Provide your team members with devices that support their needs with minimized sustainability impact.

Scale infrastructure with user load: Identify periods of low or no utilization and scale resources to eliminate excess capacity and improve efficiency.

Align SLAs with sustainability goals: Define and update service level agreements (SLAs) such as availability or data retention periods to minimize the number of resources required to support your workload while continuing to meet business requirements.

Eliminate creation and maintenance of unused assets: Analyze application assets (such as pre-compiled reports, data sets, and static images) and asset access patterns to identify redundancy, underutilization, and potential decommission targets. Consolidate generated assets with redundant content (for example, monthly reports with overlapping or common data sets and outputs) to eliminate the resources consumed when duplicating outputs. Decommission unused assets (for example, images of products that are no longer sold) to free consumed resources and reduce the number of resources used to support the workload.

Optimize geographic placement of workloads for user locations: Analyze network access patterns to identify where your customers are connecting from geographically. Select Regions and services that reduce the distance that network traffic must travel to decrease the total network resources required to support your workload.

Optimize team member resources for activities performed: Optimize resources provided to team members to minimize the sustainability impact while supporting their needs. For example, perform complex operations, such as rendering and compilation, on highly utilized shared cloud desktops instead of on under-utilized high-powered single user systems.

Software and architecture patterns

Implement patterns for performing load smoothing and maintaining consistent high utilization of deployed resources to minimize the resources consumed. Components might become idle from lack of use because of changes in user behavior over time. Revise patterns and architecture to consolidate under-utilized components to increase overall utilization. Retire components that are no longer required. Understand the performance of your workload components, and optimize the components that consume the most resources. Be aware of the devices your customers use to access your services, and implement patterns to minimize the need for device upgrades.

The following questions focus on these considerations for sustainability:

SUS 3: How do you take advantage of software and architecture patterns to support your sustainability goals?

Implement patterns for performing load smoothing and maintaining consistent high utilization of deployed resources to minimize the resources consumed. Components might become idle from lack of use because of changes in user behavior over time. Revise patterns and architecture to consolidate under-utilized components to increase overall utilization. Retire components that are no longer required. Understand the performance of your workload components, and optimize the components that consume the most resources. Be aware of the devices your customers use to access your services, and implement patterns to minimize the need for device upgrades.

Optimize software and architecture for asynchronous and scheduled jobs: Use efficient software designs and architectures to minimize the average resources required per unit of work. Implement mechanisms that result in even utilization of components to reduce resources that are idle between tasks and minimize the impact of load spikes.

Remove or refactor workload components with low or no use: Monitor workload activity to identify changes in utilization of individual components over time. Remove components that are unused and no longer required, and refactor components with little utilization, to limit wasted resources.

Optimize areas of code that consume the most time or resources: Monitor workload activity to identify application components that consume the most resources. Optimize the code that runs within these components to minimize resource usage while maximizing performance.

Optimize impact on customer devices and equipment: Understand the devices and equipment your customers use to consume your services, their expected lifecycle, and the financial and sustainability impact of replacing those components. Implement software patterns and architectures to minimize the need for customers to replace devices and upgrade equipment. For example, implement new features using code that is backwards compatible with older hardware and operating system versions, or manage the size of payloads so they don't exceed the storage capacity of the target device.

Use software patterns and architectures that best support data access and storage patterns: Understand how data is used within your workload, consumed by your users, transferred, and stored. Select technologies to minimize data processing and storage requirements.

Data patterns

Implement patterns for performing load smoothing and maintaining consistent high utilization of deployed resources to minimize the resources consumed. Components might become idle from lack of use because of changes in user behavior over time. Revise patterns and architecture to consolidate under-utilized components to increase overall utilization. Retire components that are no longer required. Understand the performance of your workload components, and optimize the components that consume the most resources. Be aware of the devices your customers use to access your services, and implement patterns to minimize the need for device upgrades.

The following question focuses on these considerations for sustainability:

SUS 4: How do you take advantage of data access and usage patterns to support your sustainability goals?

Implement data management practices to reduce the provisioned storage required to support your workload, and the resources required to use it. Understand your data, and use storage technologies and configurations that best support the business value of the data and how it's used. Lifecycle data to more efficient, less performant storage when requirements decrease, and delete data that's no longer required.

Implement a data classification policy: Classify data to understand its significance to business outcomes. Use this information to determine when you can move data to more energy-efficient storage or safely delete it.

Use technologies that support data access and storage patterns: Use storage that best supports how your data is accessed and stored to minimize the resources provisioned while supporting your workload. For example, solid state devices (SSDs) are more energy intensive than magnetic drives and should be used only for active data use cases. Use energy-efficient, archival-class storage for infrequently accessed data.

Use lifecycle policies to delete unnecessary data: Manage the lifecycle of all your data and automatically enforce deletion timelines to minimize the total storage requirements of your workload.

Minimize over-provisioning in block storage: To minimize total provisioned storage, create block storage with size allocations that are appropriate for the workload. Use elastic volumes to expand storage as data grows without having to resize storage attached to compute resources. Regularly review elastic volumes and shrink over-provisioned volumes to fit the current data size.

Remove unneeded or redundant data: Duplicate data only when necessary to minimize total storage consumed. Use backup technologies that deduplicate data at the file and block level. Limit the use of Redundant Array of Independent Drives (RAID) configurations except where required to meet SLAs.

Use shared file systems or object storage to access common data: Adopt shared storage and single sources of truth to avoid data duplication and reduce the total storage requirements of your workload. Fetch data from shared storage only as needed. Detach unused volumes to free resources. Minimize data movement across networks: Use shared storage and access data from regional data stores to minimize the total networking resources required to support data movement for your workload.

Back up data only when difficult to recreate: To minimize storage consumption, only back up data that has business value or is required to satisfy compliance requirements. Examine backup policies and exclude ephemeral storage that doesn't provide value in a recovery scenario.

Hardware patterns

Look for opportunities to reduce workload sustainability impacts by making changes to your hardware management practices. Minimize the amount of hardware needed to provision and deploy, and select the most efficient hardware for your individual workload.

The following question focuses on these considerations for sustainability:

SUS 5: How do your hardware management and usage practices support your sustainability goals?

Look for opportunities to reduce workload sustainability impacts by making changes to your hardware management practices. Minimize the amount of hardware needed to provision and deploy, and select the most efficient hardware for your individual workload.

Use the minimum amount of hardware to meet your needs: Using the capabilities of the cloud, you can make frequent changes to your workload implementations. Update deployed components as your needs change.

Use instance types with the least impact: Continually monitor the release of new instance types and take advantage of energy efficiency improvements, including those instance types designed to support specific workloads such as machine learning training and inference, and video transcoding.

Use managed services: Managed services shift responsibility for maintaining high average utilization, and sustainability optimization of the deployed hardware, to AWS. Use managed services to distribute

the sustainability impact of the service across all tenants of the service, reducing your individual contribution.

Optimize your use of GPUs: Graphics processing units (GPUs) can be a source of high-power consumption, and many GPU workloads are highly variable, such as rendering, transcoding, and machine learning training and modeling. Only run GPUs instances for the time needed, and decommission them with automation when not required to minimize resources consumed.

Development and deployment patterns

Look for opportunities to reduce your sustainability impact by making changes to your development, test, and deployment practices.

The following question focuses on these considerations for sustainability:

SUS 6: How do your development and deployment processes support your sustainability goals?

Look for opportunities to reduce your sustainability impact by making changes to your development, test, and deployment practices.

Adopt methods that can rapidly introduce sustainability improvements: Test and validate potential improvements before deploying them to production. Account for the cost of testing when calculating potential future benefit of an improvement. Develop low-cost testing methods to enable delivery of small improvements.

Keep your workload up to date: Up-to-date operating systems, libraries, and applications can improve workload efficiency and enable easier adoption of more efficient technologies. Up-to-date software might also include features to measure the sustainability impact of your workload more accurately, as vendors deliver features to meet their own sustainability goals.

Increase utilization of build environments: Use automation and infrastructure as code to bring pre-production environments up when needed and take them down when not used. A common pattern is to schedule periods of availability that coincide with the working hours of your development team members. Hibernation is a useful tool to preserve state and rapidly bring instances online only when needed. Use instance types with burst capacity, spot instances, elastic database services, containers, and other technologies to align development and test capacity with use.

Use managed device farms for testing: Managed device farms spread the sustainability impact of hardware manufacturing and resource usage across multiple tenants. Managed device farms offer diverse device types so you can support older, less popular hardware, and avoid customer sustainability impact from unnecessary device upgrades.

Resources

Refer to the following resources to learn more about our best practices for sustainability.

Whitepaper

- [Sustainability Pillar](#)

Video

- [The Climate Pledge](#)

The review process

The review of architectures needs to be done in a consistent manner, with a blamefree approach that encourages diving deep. It should be a light weight process (hours not days) that is a conversation and not an audit. The purpose of reviewing an architecture is to identify any critical issues that might need addressing or areas that could be improved. The outcome of the review is a set of actions that should improve the experience of a customer using the workload.

As discussed in the “On Architecture” section, you will want each team member to take responsibility for the quality of its architecture. We recommend that the team members who build an architecture use the Well-Architected Framework to continually review their architecture, rather than holding a formal review meeting. A continuous approach allows your team members to update answers as the architecture evolves, and improve the architecture as you deliver features.

The AWS Well-Architected Framework is aligned to the way that AWS reviews systems and services internally. It is premised on a set of design principles that influences architectural approach, and questions that ensure that people don’t neglect areas that often featured in Root Cause Analysis (RCA). Whenever there is a significant issue with an internal system, AWS service, or customer, we look at the RCA to see if we could improve the review processes we use.

Reviews should be applied at key milestones in the product lifecycle, early on in the design phase to avoid *one-way doors* that are difficult to change, and then before the go-live date. (Many decisions are reversible, two-way doors. Those decisions can use a lightweight process. One-way doors are hard or impossible to reverse and require more inspection before making them.) After you go into production, your workload will continue to evolve as you add new features and change technology implementations. The architecture of a workload changes over time. You will need to follow good hygiene practices to stop its architectural characteristics from degrading as you evolve it. As you make significant architecture changes you should follow a set of hygiene processes including a Well-Architected review.

If you want to use the review as a one-time snapshot or independent measurement, you will want to ensure that you have all the right people in the conversation. Often, we find that reviews are the first time that a team truly understands what they have implemented. An approach that works well when reviewing another team’s workload is to have a series of informal conversations about their architecture where you can glean the answers to most questions. You can then follow up with one or two meetings where you can gain clarity or dive deep on areas of ambiguity or perceived risk.

Here are some suggested items to facilitate your meetings:

- A meeting room with whiteboards
- Print outs of any diagrams or design notes
- Action list of questions that require out-of-band research to answer (for example, “did we enable encryption or not?”)

After you have done a review, you should have a list of issues that you can prioritize based on your business context. You will also want to take into account the impact of those issues on the day-to-day work of your team. If you address these issues early, you could free up time to work on creating business value rather than solving recurring problems. As you address issues, you can update your review to see how the architecture is improving.

While the value of a review is clear after you have done one, you may find that a new team might be resistant at first. Here are some objections that can be handled through educating the team on the benefits of a review:

- “We are too busy!” (Often said when the team is getting ready for a big launch.)

- If you are getting ready for a big launch you will want it to go smoothly. The review will allow you to understand any problems you might have missed.
- We recommend that you carry out reviews early in the product lifecycle to uncover risks and develop a mitigation plan aligned with the feature delivery roadmap.
- “We don’t have time to do anything with the results!” (Often said when there is an immovable event, such as the Super Bowl, that they are targeting.)
- These events can’t be moved. Do you really want to go into it without knowing the risks in your architecture? Even if you don’t address all of these issues you can still have playbooks for handling them if they materialize.
- “We don’t want others to know the secrets of our solution implementation!”
- If you point the team at the questions in the Well-Architected Framework, they will see that none of the questions reveal any commercial or technical proprietary information.

As you carry out multiple reviews with teams in your organization, you might identify thematic issues. For example, you might see that a group of teams has clusters of issues in a particular pillar or topic. You will want to look at all your reviews in a holistic manner, and identify any mechanisms, training, or principal engineering talks that could help address those thematic issues.

Conclusion

The AWS Well-Architected Framework provides architectural best practices across the six pillars for designing and operating reliable, secure, efficient, cost-effective, and sustainable systems in the cloud. The Framework provides a set of questions that allows you to review an existing or proposed architecture. It also provides a set of AWS best practices for each pillar. Using the Framework in your architecture will help you produce stable and efficient systems, which allow you to focus on your functional requirements.

Contributors

The following individuals and organizations contributed to this document:

- Brian Carlson, Operations Lead Well-Architected, Amazon Web Services
- Ben Potter, Security Lead Well-Architected, Amazon Web Services
- Seth Eliot, Reliability Lead Well-Architected, Amazon Web Services
- Eric Pullen, Sr. Solutions Architect, Amazon Web Services
- Rodney Lester, Principal Solutions Architect, Amazon Web Services
- Jon Steele, Sr. Technical Account Manager, Amazon Web Services
- Max Ramsay, Principal Security Solutions Architect, Amazon Web Services
- Callum Hughes, Solutions Architect, Amazon Web Services
- Aden Leirer, Content Program Manager Well-Architected, Amazon Web Services

Further reading

[AWS Architecture Center](#)

[AWS Cloud Compliance](#)

[AWS Well-Architected Partner program](#)

[AWS Well-Architected Tool](#)

[AWS Well-Architected homepage](#)

[Operational Excellence Pillar whitepaper](#)

[Security Pillar whitepaper](#)

[Reliability Pillar whitepaper](#)

[Performance Efficiency Pillar whitepaper](#)

[Cost Optimization Pillar whitepaper](#)

[Sustainability Pillar whitepaper](#)

[The Amazon Builders' Library](#)

Document revisions

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Minor update (p. 46)	Added definition for level of effort and updated best practices in the appendix.	October 20, 2022
Whitepaper updated (p. 46)	Added Sustainability Pillar and updated links.	December 2, 2021
Major update (p. 35)	Sustainability Pillar added to the framework.	November 20, 2021
Minor update (p. 46)	Removed non-inclusive language.	April 22, 2021
Minor update (p. 46)	Fixed numerous links.	March 10, 2021
Minor update (p. 46)	Minor editorial changes throughout.	July 15, 2020
Updates for new Framework (p. 46)	Review and rewrite of most questions and answers.	July 8, 2020
Whitepaper updated (p. 46)	Addition of AWS Well-Architected Tool, links to AWS Well-Architected Labs, and AWS Well-Architected Partners, minor fixes to enable multiple language version of framework.	July 1, 2019
Whitepaper updated (p. 46)	Review and rewrite of most questions and answers, to ensure questions focus on one topic at a time. This caused some previous questions to be split into multiple questions. Added common terms to definitions (workload, component etc). Changed presentation of question in main body to include descriptive text.	November 1, 2018
Whitepaper updated (p. 46)	Updates to simplify question text, standardize answers, and improve readability.	June 1, 2018
Whitepaper updated (p. 46)	Operational Excellence moved to front of pillars and rewritten so it frames other pillars. Refreshed other pillars to reflect evolution of AWS.	November 1, 2017

Whitepaper updated (p. 46)	Updated the Framework to include operational excellence pillar, and revised and updated the other pillars to reduce duplication and incorporate learnings from carrying out reviews with thousands of customers.	November 1, 2016
Minor updates (p. 46)	Updated the Appendix with current Amazon CloudWatch Logs information.	November 1, 2015
Initial publication (p. 46)	AWS Well-Architected Framework published.	October 1, 2015

Appendix: Questions and best practices

Topics

- Operational excellence (p. 48)
- Security (p. 117)
- Reliability (p. 169)
- Performance efficiency (p. 278)
- Cost optimization (p. 346)
- Sustainability (p. 392)

Operational excellence

Topics

- Organization (p. 48)
- Prepare (p. 64)
- Operate (p. 92)
- Evolve (p. 110)

Organization

Questions

- OPS 1 How do you determine what your priorities are? (p. 48)
- OPS 2 How do you structure your organization to support your business outcomes? (p. 55)
- OPS 3 How does your organizational culture support your business outcomes? (p. 58)

OPS 1 How do you determine what your priorities are?

Everyone needs to understand their part in enabling business success. Have shared goals in order to set priorities for resources. This will maximize the benefits of your efforts.

Best practices

- OPS01-BP01 Evaluate external customer needs (p. 49)
- OPS01-BP02 Evaluate internal customer needs (p. 49)
- OPS01-BP03 Evaluate governance requirements (p. 50)
- OPS01-BP04 Evaluate compliance requirements (p. 51)
- OPS01-BP05 Evaluate threat landscape (p. 52)
- OPS01-BP06 Evaluate tradeoffs (p. 53)

- [OPS01-BP07 Manage benefits and risks \(p. 54\)](#)

OPS01-BP01 Evaluate external customer needs

Involve key stakeholders, including business, development, and operations teams, to determine where to focus efforts on external customer needs. This will ensure that you have a thorough understanding of the operations support that is required to achieve your desired business outcomes.

Common anti-patterns:

- You have decided not to have customer support outside of core business hours, but you haven't reviewed historical support request data. You do not know whether this will have an impact on your customers.
- You are developing a new feature but have not engaged your customers to find out if it is desired, if desired in what form, and without experimentation to validate the need and method of delivery.

Benefits of establishing this best practice: Customers whose needs are satisfied are much more likely to remain customers. Evaluating and understanding external customer needs will inform how you prioritize your efforts to deliver business value.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Understand business needs: Business success is enabled by shared goals and understanding across stakeholders, including business, development, and operations teams.
- Review business goals, needs, and priorities of external customers: Engage key stakeholders, including business, development, and operations teams, to discuss goals, needs, and priorities of external customers. This ensures that you have a thorough understanding of the operational support that is required to achieve business and customer outcomes.
- Establish shared understanding: Establish shared understanding of the business functions of the workload, the roles of each of the teams in operating the workload, and how these factors support your shared business goals across internal and external customers.

Resources

Related documents:

- [AWS Well-Architected Framework Concepts – Feedback loop](#)

OPS01-BP02 Evaluate internal customer needs

Involve key stakeholders, including business, development, and operations teams, when determining where to focus efforts on internal customer needs. This will ensure that you have a thorough understanding of the operations support that is required to achieve business outcomes.

Use your established priorities to focus your improvement efforts where they will have the greatest impact (for example, developing team skills, improving workload performance, reducing costs, automating runbooks, or enhancing monitoring). Update your priorities as needs change.

Common anti-patterns:

- You have decided to change IP address allocations for your product teams, without consulting them, to make managing your network easier. You do not know the impact this will have on your product teams.

- You are implementing a new development tool but have not engaged your internal customers to find out if it is needed or if it is compatible with their existing practices.
- You are implementing a new monitoring system but have not contacted your internal customers to find out if they have monitoring or reporting needs that should be considered.

Benefits of establishing this best practice: Evaluating and understanding internal customer needs will inform how you prioritize your efforts to deliver business value.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Understand business needs: Business success is enabled by shared goals and understanding across stakeholders including business, development, and operations teams.
 - Review business goals, needs, and priorities of internal customers: Engage key stakeholders, including business, development, and operations teams, to discuss goals, needs, and priorities of internal customers. This ensures that you have a thorough understanding of the operational support that is required to achieve business and customer outcomes.
 - Establish shared understanding: Establish shared understanding of the business functions of the workload, the roles of each of the teams in operating the workload, and how these factors support shared business goals across internal and external customers.

Resources

Related documents:

- [AWS Well-Architected Framework Concepts – Feedback loop](#)

OPS01-BP03 Evaluate governance requirements

Ensure that you are aware of guidelines or obligations defined by your organization that may mandate or emphasize specific focus. Evaluate internal factors, such as organization policy, standards, and requirements. Validate that you have mechanisms to identify changes to governance. If no governance requirements are identified, ensure that you have applied due diligence to this determination.

Common anti-patterns:

- You are being audited and are asked to provide proof of compliance with internal governance. You have no idea if you are compliant because you have never evaluated what your compliance requirements are.
- You have suffered a compromise resulting in financial loss. You discover that the insurance that would have covered the financial loss was contingent on your implementation of specific security controls that are not in place and required by your governance.
- Your administrative account has been compromised resulting in the defacement of your company web site and damaged to customer trust. Your internal governance requires the use of Multifactor Authentication (MFA) to secure administrative accounts. You did not secure your administrative account with MFA and subject to disciplinary action.

Benefits of establishing this best practice: Evaluating and understanding the governance requirements that your organization applies to your workload will inform how you prioritize your efforts to deliver business value.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Understand governance requirements: Evaluate internal governance factors, such as program or organizational policy, program policies, issue or system specific policies, standards, procedures, baselines, and guidelines. Validate that you have mechanisms to identify changes to governance. If no governance requirements are identified, ensure that you have applied due diligence to this determination.

Resources

Related documents:

- [AWS Cloud Compliance](#)

OPS01-BP04 Evaluate compliance requirements

Evaluate external factors, such as regulatory compliance requirements and industry standards, to ensure that you are aware of guidelines or obligations that might mandate or emphasize specific focus. If no compliance requirements are identified, ensure that you apply due diligence to this determination.

Common anti-patterns:

- You are being audited and are asked to provide proof of compliance with industry regulations. You have no idea if you are compliant because you have never evaluated what your compliance requirements are.
- Your administrative account has been compromised resulting in the download of customer data and damaged to customer trust. Your industry best practices require the use of MFA to secure administrative accounts. You did not secure your administrative account with MFA and subject to litigation by your customers.

Benefits of establishing this best practice: Evaluating and understanding the compliance requirements that apply to your workload will inform how you prioritize your efforts to deliver business value.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Understand compliance requirements: Evaluate external factors, such as regulatory compliance requirements and industry standards, to ensure that you are aware of guidelines or obligations that might mandate or emphasize specific focus. If no compliance requirements are identified, ensure that due diligence was applied to the determination.
- Understand regulatory compliance requirements: Identify regulatory compliance requirements that you are legally obligated to satisfy. Use these requirements to focus your efforts. Examples include obligations from privacy and data protection acts.
 - [AWS Compliance](#)
 - [AWS Compliance Programs](#)
 - [AWS Compliance Latest News](#)
- Understand industry standards and best practices: Identify industry standards and best practice requirements that apply to your workload, such as the Payment Card Industry Data Security Standard (PCI DSS). Use these requirements to focus your efforts.
 - [AWS Compliance Programs](#)
- Understand internal compliance requirements: Identify compliance requirements and best practices that are established by your organization. Use these requirements to focus your efforts. Examples include information security policies and data classification standards.

Resources

Related documents:

- [AWS Cloud Compliance](#)
- [AWS Compliance](#)
- [AWS Compliance Latest News](#)
- [AWS Compliance Programs](#)

OPS01-BP05 Evaluate threat landscape

Evaluate threats to the business (for example, competition, business risk and liabilities, operational risks, and information security threats) and maintain current information in a risk registry. Include the impact of risks when determining where to focus efforts.

The [Well-Architected Framework](#) emphasizes learning, measuring, and improving. It provides a consistent approach for you to evaluate architectures, and implement designs that will scale over time. AWS provides the [AWS Well-Architected Tool](#) to help you review your approach prior to development, the state of your workloads prior to production, and the state of your workloads in production. You can compare them to the latest AWS architectural best practices, monitor the overall status of your workloads, and gain insight to potential risks.

AWS customers are eligible for a guided Well-Architected Review of their mission-critical workloads to [measure their architectures](#) against AWS best practices. Enterprise Support customers are eligible for an [Operations Review](#), designed to help them to identify gaps in their approach to operating in the cloud.

The cross-team engagement of these reviews helps to establish common understanding of your workloads and how team roles contribute to success. The needs identified through the review can help shape your priorities.

[AWS Trusted Advisor](#) is a tool that provides access to a core set of checks that recommend optimizations that may help shape your priorities. [Business and Enterprise Support customers](#) receive access to additional checks focusing on security, reliability, performance, and cost-optimization that can further help shape their priorities.

Common anti-patterns:

- You are using an old version of a software library in your product. You are unaware of security updates to the library for issues that may have unintended impact on your workload.
- Your competitor just released a version of their product that addresses many of your customers' complaints about your product. You have not prioritized addressing any of these known issues.
- Regulators have been pursuing companies like yours that are not compliant with legal regulatory compliance requirements. You have not prioritized addressing any of your outstanding compliance requirements.

Benefits of establishing this best practice: Identifying and understanding the threats to your organization and workload enables your determination of which threats to address, their priority, and the resources necessary to do so.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Evaluate threat landscape: Evaluate threats to the business (for example, competition, business risk and liabilities, operational risks, and information security threats), so that you can include their impact when determining where to focus efforts.

- [AWS Latest Security Bulletins](#)
- [AWS Trusted Advisor](#)
- Maintain a threat model: Establish and maintain a threat model identifying potential threats, planned and in place mitigations, and their priority. Review the probability of threats manifesting as incidents, the cost to recover from those incidents and the expected harm caused, and the cost to prevent those incidents. Revise priorities as the contents of the threat model change.

Resources

Related documents:

- [AWS Cloud Compliance](#)
- [AWS Latest Security Bulletins](#)
- [AWS Trusted Advisor](#)

OPS01-BP06 Evaluate tradeoffs

Evaluate the impact of tradeoffs between competing interests or alternative approaches, to help make informed decisions when determining where to focus efforts or choosing a course of action. For example, accelerating speed to market for new features may be emphasized over cost optimization, or you may choose a relational database for non-relational data to simplify the effort to migrate a system, rather than migrating to a database optimized for your data type and updating your application.

AWS can help you educate your teams about AWS and its services to increase their understanding of how their choices can have an impact on your workload. You should use the resources provided by [AWS Support](#) ([AWS Knowledge Center](#), [AWS Discussion Forums](#), and [AWS Support Center](#)) and [AWS Documentation](#) to educate your teams. Reach out to AWS Support through AWS Support Center for help with your AWS questions.

AWS also shares best practices and patterns that we have learned through the operation of AWS in [The Amazon Builders' Library](#). A wide variety of other useful information is available through the [AWS Blog](#) and [The Official AWS Podcast](#).

Common anti-patterns:

- You are using a relational database to manage time series and non-relational data. There are database options that are optimized to support the data types you are using but you are unaware of the benefits because you have not evaluated the tradeoffs between solutions.
- Your investors request that you demonstrate compliance with Payment Card Industry Data Security Standards (PCI DSS). You do not consider the tradeoffs between satisfying their request and continuing with your current development efforts. Instead you proceed with your development efforts without demonstrating compliance. Your investors stop their support of your company over concerns about the security of your platform and their investments.

Benefits of establishing this best practice: Understanding the implications and consequences of your choices enables you to prioritize your options.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Evaluate tradeoffs: Evaluate the impact of tradeoffs between competing interests, to help make informed decisions when determining where to focus efforts. For example, accelerating speed to market for new features might be emphasized over cost optimization.

- AWS can help you educate your teams about AWS and its services to increase their understanding of how their choices can have an impact on your workload. You should use the resources provided by AWS Support (AWS Knowledge Center, AWS Discussion Forums, and AWS Support Center) and AWS Documentation to educate your teams. Reach out to AWS Support through AWS Support Center for help with your AWS questions.
- AWS also shares best practices and patterns that we have learned through the operation of AWS in The Amazon Builders' Library. A wide variety of other useful information is available through the AWS Blog and The Official AWS Podcast.

Resources

Related documents:

- [AWS Blog](#)
- [AWS Cloud Compliance](#)
- [AWS Discussion Forums](#)
- [AWS Documentation](#)
- [AWS Knowledge Center](#)
- [AWS Support](#)
- [AWS Support Center](#)
- [The Amazon Builders' Library](#)
- [The Official AWS Podcast](#)

OPS01-BP07 Manage benefits and risks

Manage benefits and risks to make informed decisions when determining where to focus efforts. For example, it may be beneficial to deploy a workload with unresolved issues so that significant new features can be made available to customers. It may be possible to mitigate associated risks, or it may become unacceptable to allow a risk to remain, in which case you will take action to address the risk.

You might find that you want to emphasize a small subset of your priorities at some point in time. Use a balanced approach over the long term to ensure the development of needed capabilities and management of risk. Update your priorities as needs change

Common anti-patterns:

- You have decided to include a library that does everything you need that one of your developers found on the internet. You have not evaluated the risks of adopting this library from an unknown source and do not know if it contains vulnerabilities or malicious code.
- You have decided to develop and deploy a new feature instead of fixing an existing issue. You have not evaluated the risks of leaving the issue in place until the feature is deployed and do not know what the impact will be on your customers.
- You have decided to not deploy a feature frequently requested by customers because of unspecified concerns from your compliance team.

Benefits of establishing this best practice: Identifying the available benefits of your choices, and being aware of the risks to your organization, enables you to make informed decisions.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Manage benefits and risks: Balance the benefits of decisions against the risks involved.

- **Identify benefits:** Identify benefits based on business goals, needs, and priorities. Examples include time-to-market, security, reliability, performance, and cost.
- **Identify risks:** Identify risks based on business goals, needs, and priorities. Examples include time-to-market, security, reliability, performance, and cost.
- **Assess benefits against risks and make informed decisions:** Determine the impact of benefits and risks based on goals, needs, and priorities of your key stakeholders, including business, development, and operations. Evaluate the value of the benefit against the probability of the risk being realized and the cost of its impact. For example, emphasizing speed-to-market over reliability might provide competitive advantage. However, it may result in reduced uptime if there are reliability issues.

OPS 2 How do you structure your organization to support your business outcomes?

Your teams must understand their part in achieving business outcomes. Teams need to understand their roles in the success of other teams, the role of other teams in their success, and have shared goals. Understanding responsibility, ownership, how decisions are made, and who has authority to make decisions will help focus efforts and maximize the benefits from your teams.

Best practices

- [OPS02-BP01 Resources have identified owners \(p. 55\)](#)
- [OPS02-BP02 Processes and procedures have identified owners \(p. 56\)](#)
- [OPS02-BP03 Operations activities have identified owners responsible for their performance \(p. 56\)](#)
- [OPS02-BP04 Team members know what they are responsible for \(p. 57\)](#)
- [OPS02-BP05 Mechanisms exist to identify responsibility and ownership \(p. 57\)](#)
- [OPS02-BP06 Mechanisms exist to request additions, changes, and exceptions \(p. 57\)](#)
- [OPS02-BP07 Responsibilities between teams are predefined or negotiated \(p. 58\)](#)

OPS02-BP01 Resources have identified owners

Understand who has ownership of each application, workload, platform, and infrastructure component, what business value is provided by that component, and why that ownership exists. Understanding the business value of these individual components and how they support business outcomes informs the processes and procedures applied against them.

Benefits of establishing this best practice: Understanding ownership identifies whom can approve improvements, implement those improvements, or both.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Resources have identified owners:** Define what ownership means for the resource use cases in your environment. Specify and record owners for resources including at a minimum name, contact information, organization, and team. Store resource ownership information with resources using metadata such as tags or resource groups. Use AWS Organizations to structure accounts and implement policies to ensure ownership and contact information are captured.
- **Define forms of ownership and how they are assigned:** Ownership may have multiple definitions in your organization with different use cases. You may wish to define a workload owner as the individual who owns the risk and liability for the operation of a workload, and whom ultimately has authority to make decisions about the workload. You may wish to define ownership in terms of financial or administrative responsibility where ownership rolls up to a parent organization. A

developer may be the owner of their development environment and be responsible for incidents that its operation causes. Their product lead may own responsibility for the financial costs associated to the operation of their development environments.

- Define who owns an organization, account, collection of resources, or individual components: Define and record ownership in an appropriately accessible location organized to support discovery. Update definitions and ownership details as they change.
- Capture ownership in the metadata for the resources: Capture resource ownership using metadata such as tags or resource groups, specifying ownership and contact information. Use AWS Organizations to structure accounts and ensure ownership and contact information are captured.

OPS02-BP02 Processes and procedures have identified owners

Understand who has ownership of the definition of individual processes and procedures, why those specific process and procedures are used, and why that ownership exists. Understanding the reasons that specific processes and procedures are used enables identification of improvement opportunities.

Benefits of establishing this best practice: Understanding ownership identifies who can approve improvements, implement those improvements, or both.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Process and procedures have identified owners responsible for their definition: Capture the processes and procedures used in your environment and the individual or team responsible for their definition.
- Identify process and procedures: Identify the operations activities conducted in support of your workloads. Document these activities in a discoverable location.
- Define who owns the definition of a process or procedure: Uniquely identify the individual or team responsible for the specification of an activity. They are responsible to ensure it can be successfully performed by an adequately skilled team member with the correct permissions, access, and tools. If there are issues with performing that activity, the team members performing it are responsible to provide the detailed feedback necessary for the activity to be improved.
- Capture ownership in the metadata of the activity artifact: Procedures automated in services like AWS Systems Manager, through documents, and AWS Lambda, as functions, support capturing metadata information as tags. Capture resource ownership using tags or resource groups, specifying ownership and contact information. Use AWS Organizations to create tagging policies and ensure ownership and contact information are captured.

OPS02-BP03 Operations activities have identified owners responsible for their performance

Understand who has responsibility to perform specific activities on defined workloads and why that responsibility exists. Understanding who has responsibility to perform activities informs who will conduct the activity, validate the result, and provide feedback to the owner of the activity.

Benefits of establishing this best practice: Understanding who is responsible to perform an activity informs whom to notify when action is needed and who will perform the action, validate the result, and provide feedback to the owner of the activity.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Operations activities have identified owners responsible for their performance: Capture the responsibility for performing processes and procedures used in your environment

- Identify process and procedures: Identify the operations activities conducted in support of your workloads. Document these activities in a discoverable location.
- Define who is responsible to perform each activity: Identify the team responsible for an activity. Ensure they have the details of the activity, and the necessary skills and correct permissions, access, and tools to perform the activity. They must understand the condition under which it is to be performed (for example, on an event or schedule). Make this information discoverable so that members of your organization can identify who they need to contact, team or individual, for specific needs.

OPS02-BP04 Team members know what they are responsible for

Understanding the responsibilities of your role and how you contribute to business outcomes informs the prioritization of your tasks and why your role is important. This enables team members to recognize needs and respond appropriately.

Benefits of establishing this best practice: Understanding your responsibilities informs the decisions you make, the actions you take, and your hand off activities to their proper owners.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Ensure team members understand their roles and responsibilities: Identify team members roles and responsibilities and ensure they understand the expectations of their role. Make this information discoverable so that members of your organization can identify who they need to contact, team or individual, for specific needs.

OPS02-BP05 Mechanisms exist to identify responsibility and ownership

Where no individual or team is identified, there are defined escalation paths to someone with the authority to assign ownership or plan for that need to be addressed.

Benefits of establishing this best practice: Understanding who has responsibility or ownership allows you to reach out to the proper team or team member to make a request or transition a task. Having an identified person who has the authority to assign responsibility or ownership or plan to address needs reduces the risk of inaction and needs not being addressed.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Mechanisms exist to identify responsibility and ownership: Provide accessible mechanisms for members of your organization to discover and identify ownership and responsibility. These mechanisms will enable them to identify who to contact, team or individual, for specific needs.

OPS02-BP06 Mechanisms exist to request additions, changes, and exceptions

You are able to make requests to owners of processes, procedures, and resources. Make informed decisions to approve requests where viable and determined to be appropriate after an evaluation of benefits and risks.

Benefits of establishing this best practice: It's critical that mechanisms exist to request additions, changes, and exceptions in support of teams' activities. Without this option, current state become a constraint on innovation.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Mechanisms exist to request additions, changes, and exceptions: When standards are rigid innovation is constrained. Provide mechanisms for members of your organization to make requests to owners of processes, procedures, and resources in support of their business needs.

OPS02-BP07 Responsibilities between teams are predefined or negotiated

Have defined or negotiated agreements between teams describing how they work with and support each other (for example, response times, service level objectives, or service level agreements). Understanding the impact of the teams' work on business outcomes, and the outcomes of other teams and organizations, informs the prioritization of their tasks and enables them to respond appropriately.

When responsibility and ownership are undefined or unknown, you are at risk of both not addressing necessary activities in a timely fashion and of redundant and potentially conflicting efforts emerging to address those needs.

Benefits of establishing this best practice: Establishing the responsibilities between teams, the objectives, and the methods for communicating needs, eases the flow of requests and helps ensure the necessary information is provided. This reduces the delay introduced by transition tasks between teams and help support the achievement of business outcomes.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Responsibilities between teams are predefined or negotiated: Specifying the methods by which teams interact, and the information necessary for them to support each other, can help minimize the delay introduced as requests are iteratively reviewed and clarified. Having specific agreements that define expectations (for example, response time, or fulfillment time) enables teams to make effective plans and resource appropriately.

OPS 3 How does your organizational culture support your business outcomes?

Provide support for your team members so that they can be more effective in taking action and supporting your business outcome.

Best practices

- OPS03-BP01 Executive Sponsorship (p. 58)
- OPS03-BP02 Team members are empowered to take action when outcomes are at risk (p. 59)
- OPS03-BP03 Escalation is encouraged (p. 60)
- OPS03-BP04 Communications are timely, clear, and actionable (p. 60)
- OPS03-BP05 Experimentation is encouraged (p. 61)
- OPS03-BP06 Team members are enabled and encouraged to maintain and grow their skill sets (p. 61)
- OPS03-BP07 Resource teams appropriately (p. 62)
- OPS03-BP08 Diverse opinions are encouraged and sought within and across teams (p. 63)

OPS03-BP01 Executive Sponsorship

Senior leadership clearly sets expectations for the organization and evaluates success. Senior leadership is the sponsor, advocate, and driver for the adoption of best practices and evolution of the organization

Benefits of establishing this best practice: Engaged leadership, clearly communicated expectations, and shared goals ensures that team members know what is expected of them. Evaluating success enables identification of barriers to success so that they can be addressed through intervention by the sponsor advocate or their delegates.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Executive Sponsorship:** Senior leadership clearly sets expectations for the organization and evaluates success. Senior leadership is the sponsor, advocate, and driver for the adoption of best practices and evolution of the organization
 - **Set expectations:** Define and publish goals for your organizations including how they will be measured.
 - **Track achievement of goals:** Measure the incremental achievement of goals regularly and share the results so that appropriate action can be taken if outcomes are at risk.
 - **Provide the resources necessary to achieve your goals:** Regularly review if resources are still appropriate, or if additional resources are needed based on: new information, changes to goals, responsibilities, or your business environment.
 - **Advocate for your teams:** Remain engaged with your teams so that you understand how they are doing and if there are external factors affecting them. When your teams are impacted by external factors, reevaluate goals and adjust targets as appropriate. Identify obstacles that are impeding your teams progress. Act on behalf of your teams to help address obstacles and remove unnecessary burdens.
- **Be a driver for adoption of best practices:** Acknowledge best practices that provide quantifiable benefits and recognize the creators and adopters. Encourage further adoption to magnify the benefits achieved.
- **Be a driver for evolution of for your teams:** Create a culture of continual improvement. Encourage both personal and organizational growth and development. Provide long term targets to strive for that will require incremental achievement over time. Adjust this vision to compliment your needs, business goals, and business environment as they change.

OPS03-BP02 Team members are empowered to take action when outcomes are at risk

The workload owner has defined guidance and scope empowering team members to respond when outcomes are at risk. Escalation mechanisms are used to get direction when events are outside of the defined scope.

Benefits of establishing this best practice: By testing and validating changes early, you are able to address issues with minimized costs and limit the impact on your customers. By testing prior to deployment you minimize the introduction of errors.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Team members are empowered to take action when outcomes are at risk:** Provide your team members the permissions, tools, and opportunity to practice the skills necessary to respond effectively.
- **Give your team members opportunity to practice the skills necessary to respond:** Provide alternative safe environments where processes and procedures can be tested and trained upon safely. Perform game days to allow team members to gain experience responding to real world incidents in simulated and safe environments.
- **Define and acknowledge team members' authority to take action:** Specifically define team members authority to take action by assigning permissions and access to the workloads and components they support. Acknowledge that they are empowered to take action when outcomes are at risk.

OPS03-BP03 Escalation is encouraged

Team members have mechanisms and are encouraged to escalate concerns to decision makers and stakeholders if they believe outcomes are at risk. Escalation should be performed early and often so that risks can be identified, and prevented from causing incidents.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Encourage early and frequent escalation: Organizationally acknowledge that escalation early and often is the best practice. Organizationally acknowledge and accept that escalations may prove to be unfounded, and that it is better to have the opportunity to prevent an incident than to miss that opportunity by not escalating.
- Have a mechanism for escalation: Have documented procedures defining when and how escalation should occur. Document the series of people with increasing authority to take action or approve action and their contact information. Escalation should continue until the team member is satisfied that they have handed off the risk to a person able to address it, or they have contacted the person who owns the risk and liability for the operation of the workload. It is that person who ultimately owns all decisions with respect to their workload. Escalations should include the nature of the risk, the criticality of the workload, who is impacted, what the impact is, and the urgency, that is, when is the impact expected.
- Protect employees who escalate: Have policy that protects team members from retribution if they escalate around a non-responsive decision maker or stakeholder. Have mechanisms in place to identify if this is occurring and respond appropriately.

OPS03-BP04 Communications are timely, clear, and actionable

Mechanisms exist and are used to provide timely notice to team members of known risks and planned events. Necessary context, details, and time (when possible) are provided to support determining if action is necessary, what action is required, and to take action in a timely manner. For example, providing notice of software vulnerabilities so that patching can be expedited, or providing notice of planned sales promotions so that a change freeze can be implemented to avoid the risk of service disruption.

Planned events can be recorded in a change calendar or maintenance schedule so that team members can identify what activities are pending.

On AWS, [AWS Systems Manager Change Calendar](#) can be used to record these details. It supports programmatic checks of calendar status to determine if the calendar is open or closed to activity at a particular point of time. Operations activities can be planned around specific *approved* windows of time that are reserved for potentially disruptive activities. AWS Systems Manager Maintenance Windows allows you to schedule activities against instances and other [supported resources](#) to automate the activities and make those activities discoverable.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Communications are timely, clear, and actionable: Mechanisms are in place to provide notification of risks or planned events in a clear and actionable way with enough notice to allow appropriate responses.
- Document planned activities on a change calendar and provide notifications: Provide an accessible source of information where planned events can be discovered. Provide notifications of planned events from the same system.
- Track events and activity that may have an impact on your workload: Monitoring vulnerability notifications and patch information to understand vulnerabilities in the wild and potential risks

associated to your workload components. Provide notification to team members so that they can take action.

Resources

Related documents:

- [AWS Systems Manager Change Calendar](#)
- [AWS Systems Manager Maintenance Windows](#)

OPS03-BP05 Experimentation is encouraged

Experimentation accelerates learning and keeps team members interested and engaged. An undesired result is a successful experiment that has identified a path that will not lead to success. Team members are not punished for successful experiments with undesired results. Experimentation is required for innovation to happen and turn ideas into outcomes.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Experimentation is encouraged: Encourage experimentation to support learning and innovation.
 - Experiment with a variety of technologies: Encourage experimentation with technologies that may have applicability now or in the future to the achievement of your business outcomes. This knowledge may inform future innovation.
 - Experiment with a goal in mind: Encourage experimentation with specific goals for team members to reach for, or with technologies that may have applicability in the near future. This knowledge may inform your innovation.
 - Provide structured time to experiment: Dedicate specific times when team members can be free of their normal responsibilities, so that they can focus on their experiments.
 - Provide the resources to support experimentation: Fund the resources required to conduct experiments (for example, software, or cloud resources).
 - Acknowledge success: Recognize the value yielded by experimentation. Understand that experiments with undesired outcomes are successful and have identified a path that will not lead to success. Team members are not punished for undesired outcomes from experiments.

OPS03-BP06 Team members are enabled and encouraged to maintain and grow their skill sets

Teams must grow their skill sets to adopt new technologies, and to support changes in demand and responsibilities in support of your workloads. Growth of skills in new technologies is frequently a source of team member satisfaction and supports innovation. Support your team members' pursuit and maintenance of industry certifications that validate and acknowledge their growing skills. Cross train to promote knowledge transfer and reduce the risk of significant impact when you lose skilled and experienced team members with institutional knowledge. Provide dedicated structured time for learning.

AWS provides resources, including the [AWS Getting Started Resource Center](#), [AWS Blogs](#), [AWS Online Tech Talks](#), [AWS Events and Webinars](#), and the [AWS Well-Architected Labs](#), that provide guidance, examples, and detailed walkthroughs to educate your teams.

AWS also shares best practices and patterns that we have learned through the operation of AWS in [The Amazon Builders' Library](#) and a wide variety of other useful educational material through the [AWS Blog](#) and [The Official AWS Podcast](#).

You should take advantage of the education resources provided by AWS such as the Well-Architected labs, [AWS Support](#) ([AWS Knowledge Center](#), [AWS Discussion Forms](#), and [AWS Support Center](#)) and [AWS Documentation](#) to educate your teams. Reach out to AWS Support through AWS Support Center for help with your AWS questions.

[AWS Training and Certification](#) provides some free training through self-paced digital courses on AWS fundamentals. You can also register for instructor-led training to further support the development of your teams' AWS skills.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Team members are enabled and encouraged to maintain and grow their skill sets: To adopt new technologies, support innovation, and to support changes in demand and responsibilities in support of your workloads continuing education is necessary.
- Provide resources for education: Provided dedicated structured time, access to training materials, lab resources, and support participation in conferences and professional organizations that provide opportunities for learning from both educators and peers. Provide junior team members' access to senior team members as mentors or allow them to shadow their work and be exposed to their methods and skills. Encourage learning about content not directly related to work in order to have a broader perspective.
- Team education and cross-team engagement: Plan for the continuing education needs of your team members. Provide opportunities for team members to join other teams (temporarily or permanently) to share skills and best practices benefiting your entire organization
- Support pursuit and maintenance of industry certifications: Support your team members acquiring and maintaining industry certifications that validate what they have learned, and acknowledge their accomplishments.

Resources

Related documents:

- [AWS Getting Started Resource Center](#)
- [AWS Blogs](#)
- [AWS Cloud Compliance](#)
- [AWS Discussion Forms](#)
- [AWS Documentation](#)
- [AWS Online Tech Talks](#)
- [AWS Events and Webinars](#)
- [AWS Knowledge Center](#)
- [AWS Support](#)
- [AWS Training and Certification](#)
- [AWS Well-Architected Labs](#),
- [The Amazon Builders' Library](#)
- [The Official AWS Podcast](#).

OPS03-BP07 Resource teams appropriately

Maintain team member capacity, and provide tools and resources to support your workload needs. Overtasking team members increases the risk of incidents resulting from human error. Investments in

tools and resources (for example, providing automation for frequently performed activities) can scale the effectiveness of your team, enabling them to support additional activities.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Resource teams appropriately: Ensure you have an understanding of the success of your teams and the factors that contribute to their success or lack of success. Act to support teams with appropriate resources.
- Understand team performance: Measure the achievement of operational outcomes and the development of assets by your teams. Track changes in output and error rate over time. Engage with teams to understand the work related challenges that impact them (for example, increasing responsibilities, changes in technology, loss of personnel, or increase in customers supported).
- Understand impacts on team performance: Remain engaged with your teams so that you understand how they are doing and if there are external factors affecting them. When your teams are impacted by external factors, reevaluate goals and adjust targets as appropriate. Identify obstacles that are impeding your teams progress. Act on behalf of your teams to help address obstacles and remove unnecessary burdens.
- Provide the resources necessary for teams to be successful: Regularly review if resources are still appropriate, or if additional resources are needed, and make appropriate adjustments to support teams.

OPS03-BP08 Diverse opinions are encouraged and sought within and across teams

Leverage cross-organizational diversity to seek multiple unique perspectives. Use this perspective to increase innovation, challenge your assumptions, and reduce the risk of confirmation bias. Grow inclusion, diversity, and accessibility within your teams to gain beneficial perspectives.

Organizational culture has a direct impact on team member job satisfaction and retention. Enable the engagement and capabilities of your team members to enable the success of your business.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Seek diverse opinions and perspectives: Encourage contributions from everyone. Give voice to under-represented groups. Rotate roles and responsibilities in meetings.
- Expand roles and responsibilities: Provide opportunity for team members to take on roles that they might not otherwise. They will gain experience and perspective from the role, and from interactions with new team members with whom they might not otherwise interact. They will bring their experience and perspective to the new role and team members they interact with. As perspective increases, additional business opportunities may emerge, or new opportunities for improvement may be identified. Have members within a team take turns at common tasks that others typically perform to understand the demands and impact of performing them.
- Provide a safe and welcoming environment: Have policy and controls that protect team members' mental and physical safety within your organization. Team members should be able to interact without fear of reprisal. When team members feel safe and welcome they are more likely to be engaged and productive. The more diverse your organization the better your understanding can be of the people you support including your customers. When your team members are comfortable, feel free to speak, and are confident they will be heard, they are more likely to share valuable insights (for example, marketing opportunities, accessibility needs, unserved market segments, unacknowledged risks in your environment).
- Enable team members to participate fully: Provide the resources necessary for your employees to participate fully in all work related activities. Team members that face daily challenges have

developed skills for working around them. These uniquely developed skills can provide significant benefit to your organization. Supporting team members with necessary accommodations will increase the benefits you can receive from their contributions.

Prepare

Questions

- [OPS 4 How do you design your workload so that you can understand its state? \(p. 64\)](#)
- [OPS 5 How do you reduce defects, ease remediation, and improve flow into production? \(p. 70\)](#)
- [OPS 6 How do you mitigate deployment risks? \(p. 80\)](#)
- [OPS 7 How do you know that you are ready to support a workload? \(p. 86\)](#)

OPS 4 How do you design your workload so that you can understand its state?

Design your workload so that it provides the information necessary across all components (for example, metrics, logs, and traces) for you to understand its internal state. This enables you to provide effective responses when appropriate.

Best practices

- [OPS04-BP01 Implement application telemetry \(p. 64\)](#)
- [OPS04-BP02 Implement and configure workload telemetry \(p. 67\)](#)
- [OPS04-BP03 Implement user activity telemetry \(p. 68\)](#)
- [OPS04-BP04 Implement dependency telemetry \(p. 68\)](#)
- [OPS04-BP05 Implement transaction traceability \(p. 69\)](#)

OPS04-BP01 Implement application telemetry

Application telemetry is the foundation for observability of your workload. Your application should emit telemetry that provides insight into the state of the application and the achievement of business outcomes. From troubleshooting to measuring the impact of a new feature, application telemetry informs the way you build, operate, and evolve your workload.

Application telemetry consists of metrics and logs. Metrics are diagnostic information, such as your pulse or temperature. Metrics are used collectively to describe the state of your application. Collecting metrics over time can be used to develop baselines and detect anomalies. Logs are messages that the application sends about its internal state or events that occur. Error codes, transaction identifiers, and user actions are examples of events that are logged.

Desired Outcome:

- Your application emits metrics and logs that provide insight into its health and the achievement of business outcomes.
- Metrics and logs are stored centrally for all applications in the workload.

Common anti-patterns:

- Your application doesn't emit telemetry. You are forced to rely upon your customers to tell you when something is wrong.

- A customer has reported that your application is unresponsive. You have no telemetry and are unable to confirm that the issue exists or characterize the issue without using the application yourself to understand the current user experience.

Benefits of establishing this best practice:

- You can understand the health of your application, the user experience, and the achievement of business outcomes.
- You can react quickly to changes in your application health.
- You can develop application health trends.
- You can make informed decisions about improving your application.
- You can detect and resolve application issues faster.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Implementing application telemetry consists of three steps: identifying a location to store telemetry, identifying telemetry that describes the state of the application, and instrumenting the application to emit telemetry.

As an example, an ecommerce company has a microservices based architecture. As part of their architectural design process they identified application telemetry that would help them understand the state of each microservice. For example, the user cart service emitted telemetry about events like add to cart, abandon cart, and length of time it took to add an item to the cart. All microservices would log errors, warnings, and transaction information. Telemetry would be sent to Amazon CloudWatch for storage and analysis.

Implementation steps

The first step is to identify a central location for telemetry storage for the applications in your workload. If you don't have an existing platform [Amazon CloudWatch](#) provides telemetry collection, dashboards, analysis, and event generation capabilities.

To identify what telemetry you need, start with the following questions:

- Is my application healthy?
- Is my application achieving business outcomes?

Your application should emit logs and metrics that collectively answer these questions. If you can't answer those questions with the existing application telemetry, work with business and engineering stakeholders to create a list of telemetry that can. You can request expert technical advice from your AWS account team as you identify and develop new application telemetry.

Once the additional application telemetry has been identified, work with your engineering stakeholders to instrument your application. [The AWS Distro for Open Telemetry](#) provides APIs, libraries, and agents that collect application telemetry. [This example demonstrates how to instrument a JavaScript application with custom metrics.](#)

Customers that want to understand the observability services that AWS offers can work through the [One Observability Workshop](#) on their own or request support from their AWS account team to guide them. This workshop guides you through the observability solutions at AWS and provides hands-on examples of how they're used.

For a deeper dive into application telemetry, read the [Instrumenting distributed systems for operational visibility](#) article in the Amazon Builder's Library. It explains how Amazon instruments applications and can serve as a guide for developing your own instrumentation guidelines.

Level of effort for the implementation plan: Medium

Resources

Related best practices:

[the section called “OPS04-BP02 Implement and configure workload telemetry” \(p. 67\)](#) – Application telemetry is a component of workload telemetry. In order to understand the health of the overall workload you need to understand the health of individual applications that make up the workload.

[the section called “OPS04-BP03 Implement user activity telemetry” \(p. 68\)](#) – User activity telemetry is often a subset of application telemetry. User activity like add to cart events, click streams, or completed transactions provide insight into the user experience.

[the section called “OPS04-BP04 Implement dependency telemetry” \(p. 68\)](#) – Dependency checks are related to application telemetry and may be instrumented into your application. If your application relies on external dependencies like DNS or a database your application can emit metrics and logs on reachability, timeouts, and other events.

[the section called “OPS04-BP05 Implement transaction traceability” \(p. 69\)](#) – Tracing transactions across a workload requires each application to emit information about how they process shared events. The way individual applications handle these events is emitted through their application telemetry.

[the section called “OPS08-BP02 Define workload metrics” \(p. 93\)](#) – Workload metrics are the key health indicators for your workload. Key application metrics are a part of workload metrics.

Related documents:

- [AWS Builders Library – Instrumenting Distributed Systems for Operational Visibility](#)
- [AWS Distro for OpenTelemetry](#)
- [AWS Well-Architected Operational Excellence Whitepaper – Design Telemetry](#)
- [Creating metrics from log events using filters](#)
- [Implementing Logging and Monitoring with Amazon CloudWatch](#)
- [Monitoring application health and performance with AWS Distro for OpenTelemetry](#)
- [New – How to better monitor your custom application metrics using Amazon CloudWatch Agent](#)
- [Observability at AWS](#)
- [Scenario – Publish metrics to CloudWatch](#)
- [Start Building – How to Monitor your Applications Effectively](#)
- [Using CloudWatch with an AWS SDK](#)

Related videos:

- [AWS re:Invent 2021 - Observability the open-source way](#)
- [Collect Metrics and Logs from Amazon EC2 instances with the CloudWatch Agent](#)
- [How to Easily Setup Application Monitoring for Your AWS Workloads - AWS Online Tech Talks](#)
- [Mastering Observability of Your Serverless Applications - AWS Online Tech Talks](#)
- [Open Source Observability with AWS - AWS Virtual Workshop](#)

Related examples:

- [AWS Logging & Monitoring Example Resources](#)
- [AWS Solution: Amazon CloudWatch Monitoring Framework](#)
- [AWS Solution: Centralized Logging](#)
- [One Observability Workshop](#)

OPS04-BP02 Implement and configure workload telemetry

Design and configure your workload to emit information about its internal state and current status, for example, API call volume, HTTP status codes, and scaling events. Use this information to help determine when a response is required.

Use a service such as [Amazon CloudWatch](#) to aggregate logs and metrics from workload components (for example, API logs from [AWS CloudTrail](#), [AWS Lambda metrics](#), [Amazon VPC Flow Logs](#), and [other services](#)).

Common anti-patterns:

- Your customers are complaining about poor performance. There are no recent changes to your application and so you suspect an issue with a workload component. You have no telemetry to analyze to determine what component or components are contributing to the poor performance.
- Your application is unreachable. You lack the telemetry to determine if it's a networking issue.

Benefits of establishing this best practice: Understanding what is going on inside your workload enables you to respond if necessary.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Implement log and metric telemetry: Instrument your workload to emit information about its internal state, status, and the achievement of business outcomes. Use this information to determine when a response is required.
 - [Gaining better observability of your VMs with Amazon CloudWatch - AWS Online Tech Talks](#)
 - [How Amazon CloudWatch works](#)
 - [What is Amazon CloudWatch?](#)
 - [Using Amazon CloudWatch metrics](#)
 - [What is Amazon CloudWatch Logs?](#)
 - Implement and configure workload telemetry: Design and configure your workload to emit information about its internal state and current status (for example, API call volume, HTTP status codes, and scaling events).
 - [Amazon CloudWatch metrics and dimensions reference](#)
 - [AWS CloudTrail](#)
 - [What Is AWS CloudTrail?](#)
 - [VPC Flow Logs](#)

Resources

Related documents:

- [AWS CloudTrail](#)
- [Amazon CloudWatch Documentation](#)
- [Amazon CloudWatch metrics and dimensions reference](#)
- [How Amazon CloudWatch works](#)
- [Using Amazon CloudWatch metrics](#)
- [VPC Flow Logs](#)
- [What Is AWS CloudTrail?](#)
- [What is Amazon CloudWatch Logs?](#)
- [What is Amazon CloudWatch?](#)

Related videos:

- [Application Performance Management on AWS](#)
- [Gaining Better Observability of Your VMs with Amazon CloudWatch](#)
- [Gaining better observability of your VMs with Amazon CloudWatch - AWS Online Tech Talks](#)

OPS04-BP03 Implement user activity telemetry

Instrument your application code to emit information about user activity, for example, click streams, or started, abandoned, and completed transactions. Use this information to help understand how the application is used, patterns of usage, and to determine when a response is required.

Common anti-patterns:

- Your developers have deployed a new feature without user telemetry, and utilization has increased. You cannot determine if the increased utilization is from use of the new feature, or is an issue introduced with the new code.
- Your developers have deployed a new feature without user telemetry. You cannot tell if your customers are using it without reaching out and asking them.

Benefits of establishing this best practice: Understand how your customers use your application to identify patterns of usage, unexpected behaviors, and to enable you to respond if necessary.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Implement user activity telemetry: Design your application code to emit information about user activity (for example, click streams, or started, abandoned, and completed transactions). Use this information to help understand how the application is used, patterns of usage, and to determine when a response is required.

OPS04-BP04 Implement dependency telemetry

Design and configure your workload to emit information about the status (for example, reachability or response time) of resources it depends on. Examples of external dependencies can include, external databases, DNS, and network connectivity. Use this information to determine when a response is required.

Common anti-patterns:

- You are unable to determine if the reason your application is unreachable is a DNS issue without manually performing a check to see if your DNS provider is working.
- Your shopping cart application is unable to complete transactions. You are unable to determine if it's a problem with your credit card processing provider without contacting them to verify.

Benefits of establishing this best practice: Understanding the health of your dependencies enables you to respond if necessary.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Implement dependency telemetry: Design and configure your workload to emit information about the state and status of systems it depends on. Some examples include: external databases, DNS, network connectivity, and external credit card processing services.

- [Amazon CloudWatch Agent with AWS Systems Manager integration - unified metrics & log collection for Linux & Windows](#)
- [Collect metrics and logs from Amazon EC2 instances and on-premises servers with the CloudWatch Agent](#)

Resources

Related documents:

- [Amazon CloudWatch Agent with AWS Systems Manager integration - unified metrics & log collection for Linux & Windows](#)
- [Collect metrics and logs from Amazon EC2 instances and on-premises servers with the CloudWatch Agent](#)

Related examples:

- [Well-Architected Labs – Dependency Monitoring](#)

OPS04-BP05 Implement transaction traceability

Implement your application code and configure your workload components to emit information about the flow of transactions across the workload. Use this information to determine when a response is required and to assist you in identifying the factors contributing to an issue.

On AWS, you can use distributed tracing services, such as [AWS X-Ray](#), to collect and record traces as transactions travel through your workload, generate maps to see how transactions flow across your workload and services, gain insight to the relationships between components, and identify and analyze issues in real time.

Common anti-patterns:

- You have implemented a serverless microservices architecture spanning multiple accounts. Your customers are experiencing intermittent performance issues. You are unable to discover which function or component is responsible because you lack the traces that would allow you to pinpoint where in the application the performance issue exists and what is causing the issue.
- You are trying to determine where the performance bottlenecks are in your workload so that they can be addressed in your development efforts. You are unable to see the relationship between your application components, and the services they interact with, to determine where the bottlenecks are because you lack the traces that would allow you to drill down into the specific services and paths impacting application performance.

Benefits of establishing this best practice: Understanding the flow of transactions across your workload allows you to understand the expected behavior of your workload transactions, and variations from expected behavior across your workload, enabling you to respond if necessary.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- **Implement transaction traceability:** Design your application and workload to emit information about the flow of transactions across system components, such as transaction stage, active component, and time to complete activity. Use this information to determine what is in progress, what is complete, and what the results of completed activities are. This helps you determine when a response is required. For example, longer than expected transaction response times within a component can indicate issues with that component.
 - [AWS X-Ray](#)

- [What is AWS X-Ray?](#)

Resources

Related documents:

- [AWS X-Ray](#)
- [What is AWS X-Ray?](#)

OPS 5 How do you reduce defects, ease remediation, and improve flow into production?

Adopt approaches that improve flow of changes into production, that enable refactoring, fast feedback on quality, and bug fixing. These accelerate beneficial changes entering production, limit issues deployed, and enable rapid identification and remediation of issues introduced through deployment activities.

Best practices

- [OPS05-BP01 Use version control \(p. 70\)](#)
- [OPS05-BP02 Test and validate changes \(p. 71\)](#)
- [OPS05-BP03 Use configuration management systems \(p. 72\)](#)
- [OPS05-BP04 Use build and deployment management systems \(p. 73\)](#)
- [OPS05-BP05 Perform patch management \(p. 74\)](#)
- [OPS05-BP06 Share design standards \(p. 76\)](#)
- [OPS05-BP07 Implement practices to improve code quality \(p. 77\)](#)
- [OPS05-BP08 Use multiple environments \(p. 77\)](#)
- [OPS05-BP09 Make frequent, small, reversible changes \(p. 78\)](#)
- [OPS05-BP10 Fully automate integration and deployment \(p. 79\)](#)

OPS05-BP01 Use version control

Use version control to enable tracking of changes and releases.

Many AWS services offer version control capabilities. Use a revision or source control system such as [AWS CodeCommit](#) to manage code and other artifacts, such as version-controlled [AWS CloudFormation](#) templates of your infrastructure.

Common anti-patterns:

- You have been developing and storing your code on your workstation. You have had an unrecoverable storage failure on the workstation your code is lost.
- After overwriting the existing code with your changes, you restart your application and it is no longer operable. You are unable to revert to the change.
- You have a write lock on a report file that someone else needs to edit. They contact you asking that you stop work on it so that they can complete their tasks.
- Your research team has been working on a detailed analysis that will shape your future work. Someone has accidentally saved their shopping list over the final report. You are unable to revert the change and will have to recreate the report.

Benefits of establishing this best practice: By using version control capabilities you can easily revert to known good states, previous versions, and limit the risk of assets being lost.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use version control: Maintain assets in version controlled repositories. Doing so supports tracking changes, deploying new versions, detecting changes to existing versions, and reverting to prior versions (for example, rolling back to a known good state in the event of a failure). Integrate the version control capabilities of your configuration management systems into your procedures.
 - [Introduction to AWS CodeCommit](#)
 - [What is AWS CodeCommit?](#)

Resources

Related documents:

- [What is AWS CodeCommit?](#)

Related videos:

- [Introduction to AWS CodeCommit](#)

OPS05-BP02 Test and validate changes

Test and validate changes to help limit and detect errors. Automate testing to reduce errors caused by manual processes, and reduce the level of effort to test.

Many AWS services offer version control capabilities. Use a revision or source control system such as [AWS CodeCommit](#) to manage code and other artifacts, such as version-controlled [AWS CloudFormation](#) templates of your infrastructure.

Common anti-patterns:

- You deploy your new code to production and customers start calling because your application is no longer working.
- You apply new security groups to enhance your perimeter security. It works with unintended consequences; Your users are unable to access your applications.
- You modify a method invoked by your new function. Another function was also dependant on that method and no longer works. The issue is not detected and enters production. The other function is not invoked for some time and finally fails in production without any correlation to the cause.

Benefits of establishing this best practice: By testing and validating changes early, you are able to address issues with minimized costs and limit the impact on your customers. By testing prior to deployment you minimize the introduction of errors.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Test and validate changes: Changes should be tested and the results validated at all lifecycle stages (for example, development, test, and production). Use testing results to confirm new features and mitigate the risk and impact of failed deployments. Automate testing and validation to ensure consistency of review, to reduce errors caused by manual processes, and reduce the level of effort.
 - [What is AWS CodeBuild?](#)
 - [Local build support for AWS CodeBuild](#)

Resources

Related documents:

- [AWS Developer Tools](#)
- [Local build support for AWS CodeBuild](#)
- [What is AWS CodeBuild?](#)

OPS05-BP03 Use configuration management systems

Use configuration management systems to make and track configuration changes. These systems reduce errors caused by manual processes and reduce the level of effort to deploy changes.

Static configuration management sets values when initializing a resource that are expected to remain consistent throughout the resource's lifetime. Some examples include setting the configuration for a web or application server on an instance, or defining the configuration of an AWS service within the [AWS Management Console](#) or through the [AWS CLI](#).

Dynamic configuration management sets values at initialization that can or are expected to change during the lifetime of a resource. For example, you could set a feature toggle to enable functionality in your code via a configuration change, or change the level of log detail during an incident to capture more data and then change back following the incident eliminating the now unnecessary logs and their associated expense.

If you have dynamic configurations in your applications running on instances, containers, serverless functions, or devices, you can use [AWS AppConfig](#) to manage and deploy them across your environments.

On AWS, you can use [AWS Config](#) to continuously monitor your AWS resource configurations [across accounts and Regions](#). It enables you to track their configuration history, understand how a configuration change would affect other resources, and audit them against expected or desired configurations using [AWS Config Rules](#) and [AWS Config Conformance Packs](#).

On AWS, you can build continuous integration/continuous deployment (CI/CD) pipelines using services such as [AWS Developer Tools](#) (for example, AWS CodeCommit, [AWS CodeBuild](#), [AWS CodePipeline](#), [AWS CodeDeploy](#), and [AWS CodeStar](#)).

Have a change calendar and track when significant business or operational activities or events are planned that may be impacted by implementation of change. Adjust activities to manage risk around those plans. [AWS Systems Manager Change Calendar](#) provides a mechanism to document blocks of time as open or closed to changes and why, and [share that information](#) with other AWS accounts. AWS Systems Manager Automation scripts can be configured to adhere to the change calendar state.

[AWS Systems Manager Maintenance Windows](#) can be used to schedule the performance of AWS SSM Run Command or Automation scripts, AWS Lambda invocations, or AWS Step Functions activities at specified times. Mark these activities in your change calendar so that they can be included in your evaluation.

Common anti-patterns:

- You manually update the web server configuration across your fleet and a number of servers become unresponsive due to update errors.
- You manually update your application server fleet over the course of many hours. The inconsistency in configuration during the change causes unexpected behaviors.
- Someone has updated your security groups and your web servers are no longer accessible. Without knowledge of what was changed you spend significant time investigating the issue extending your time to recovery.

Benefits of establishing this best practice: Adopting configuration management systems reduces the level of effort to make and track changes, and the frequency of errors caused by manual procedures.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use configuration management systems: Use configuration management systems to track and implement changes, to reduce errors caused by manual processes, and reduce the level of effort.
 - [Infrastructure configuration management](#)
 - [AWS Config](#)
 - [What is AWS Config?](#)
 - [Introduction to AWS CloudFormation](#)
 - [What is AWS CloudFormation?](#)
 - [AWS OpsWorks](#)
 - [What is AWS OpsWorks?](#)
 - [Introduction to AWS Elastic Beanstalk](#)
 - [What is AWS Elastic Beanstalk?](#)

Resources

Related documents:

- [AWS AppConfig](#)
- [AWS Developer Tools](#)
- [AWS OpsWorks](#)
- [AWS Systems Manager Change Calendar](#)
- [AWS Systems Manager Maintenance Windows](#)
- [Infrastructure configuration management](#)
- [What is AWS CloudFormation?](#)
- [What is AWS Config?](#)
- [What is AWS Elastic Beanstalk?](#)
- [What is AWS OpsWorks?](#)

Related videos:

- [Introduction to AWS CloudFormation](#)
- [Introduction to AWS Elastic Beanstalk](#)

OPS05-BP04 Use build and deployment management systems

Use build and deployment management systems. These systems reduce errors caused by manual processes and reduce the level of effort to deploy changes.

In AWS, you can build continuous integration/continuous deployment (CI/CD) pipelines using services such as [AWS Developer Tools](#) (for example, [AWS CodeCommit](#), [AWS CodeBuild](#), [AWS CodePipeline](#), [AWS CodeDeploy](#), and [AWS CodeStar](#)).

Common anti-patterns:

- After compiling your code on your development system you, copy the executable onto your production systems and it fails to start. The local log files indicates that it has failed due to missing dependencies.

- You successfully build your application with new features in your development environment and provide the code to Quality Assurance (QA). It fails QA because it is missing static assets.
- On Friday, after much effort, you successfully built your application manually in your development environment including your newly coded features. On Monday, you are unable to repeat the steps that allowed you to successfully build your application.
- You perform the tests you have created for your new release. Then you spend the next week setting up a test environment and performing all the existing integration tests followed by the performance tests. The new code has an unacceptable performance impact and must be redeveloped and then retested.

Benefits of establishing this best practice: By providing mechanisms to manage build and deployment activities you reduce the level of effort to perform repetitive tasks, free your team members to focus on their high value creative tasks, and limit the introduction of error from manual procedures.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use build and deployment management systems: Use build and deployment management systems to track and implement change, to reduce errors caused by manual processes, and reduce the level of effort. Fully automate the integration and deployment pipeline from code check-in through build, testing, deployment, and validation. This reduces lead time, enables increased frequency of change, and reduces the level of effort.
 - [What is AWS CodeBuild?](#)
 - [Continuous integration best practices for software development](#)
 - [Slalom: CI/CD for serverless applications on AWS](#)
 - [Introduction to AWS CodeDeploy - automated software deployment with Amazon Web Services](#)
 - [What is AWS CodeDeploy?](#)

Resources

Related documents:

- [AWS Developer Tools](#)
- [What is AWS CodeBuild?](#)
- [What is AWS CodeDeploy?](#)

Related videos:

- [Continuous integration best practices for software development](#)
- [Introduction to AWS CodeDeploy - automated software deployment with Amazon Web Services](#)
- [Slalom: CI/CD for serverless applications on AWS](#)

OPS05-BP05 Perform patch management

Perform patch management to gain features, address issues, and remain compliant with governance. Automate patch management to reduce errors caused by manual processes, and reduce the level of effort to patch.

Patch and vulnerability management are part of your benefit and risk management activities. It is preferable to have immutable infrastructures and deploy workloads in verified known good states. Where that is not viable, patching in place is the remaining option.

Updating machine images, container images, or Lambda [custom runtimes and additional libraries](#) to remove vulnerabilities are part of patch management. You should manage updates to [Amazon Machine Images \(AMIs\)](#) for Linux or Windows Server images using [EC2 Image Builder](#). You can use [Amazon Elastic Container Registry](#) with your existing pipeline to [manage Amazon ECS images](#) and [manage Amazon EKS images](#). AWS Lambda includes [version](#) management features.

Patching should not be performed on production systems without first testing in a safe environment. Patches should only be applied if they support an operational or business outcome. On AWS, you can use [AWS Systems Manager Patch Manager](#) to automate the process of patching managed systems and schedule the activity using [AWS Systems Manager Maintenance Windows](#).

Common anti-patterns:

- You are given a mandate to apply all new security patches within two hours resulting in multiple outages due to application incompatibility with patches.
- An unpatched library results in unintended consequences as unknown parties use vulnerabilities within it to access your workload.
- You patch the developer environments automatically without notifying the developers. You receive multiple complaints from the developers that their environment cease to operate as expected.
- You have not patched the commercial off-the-shelf software on a persistent instance. When you have an issue with the software and contact the vendor, they notify you that version is not supported and you will have to patch to a specific level to receive any assistance.
- A recently released patch for the encryption software you used has significant performance improvements. Your unpatched system has performance issues that remain in place as a result of not patching.

Benefits of establishing this best practice: By establishing a patch management process, including your criteria for patching and methodology for distribution across your environments, you will be able to realize their benefits and control their impact. This will enable the adoption of desired features and capabilities, the removal of issues, and sustained compliance with governance. Implement patch management systems and automation to reduce the level of effort to deploy patches and limit errors caused by manual processes.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Patch management: Patch systems to remediate issues, to gain desired features or capabilities, and to remain compliant with governance policy and vendor support requirements. In immutable systems, deploy with the appropriate patch set to achieve the desired result. Automate the patch management mechanism to reduce the elapsed time to patch, to reduce errors caused by manual processes, and reduce the level of effort to patch.
 - [AWS Systems Manager Patch Manager](#)

Resources

Related documents:

- [AWS Developer Tools](#)
- [AWS Systems Manager Patch Manager](#)

Related videos:

- [CI/CD for Serverless Applications on AWS](#)
- [Design with Ops in Mind](#)

Related examples:

- [Well-Architected Labs – Inventory and Patch Management](#)

OPS05-BP06 Share design standards

Share best practices across teams to increase awareness and maximize the benefits of development efforts.

On AWS, application, compute, infrastructure, and operations can be defined and managed using code methodologies. This allows for easy release, sharing, and adoption.

Many AWS services and resources are designed to be shared across accounts, enabling you to share created assets and learnings across your teams. For example, you can share [CodeCommit](#) repositories, [Lambda](#) functions, [Amazon S3 buckets](#), and [AMIs](#) to specific accounts.

When you publish new resources or updates, use Amazon SNS to provide [cross account notifications](#). Subscribers can use Lambda to get new versions.

If shared standards are enforced in your organization, it's critical that mechanisms exist to request additions, changes, and exceptions to standards in support of teams' activities. Without this option, standards become a constraint on innovation.

Common anti-patterns:

- You have created your own user authentication mechanism, as have each of the other development teams in your organization. Your users have to maintain a separate set of credentials for each part of the system they want to access.
- You have created your own user authentication mechanism, as have each of the other development teams in your organization. Your organization is given a new compliance requirement that must be met. Every individual development team must now invest the resources to implement the new requirement.
- You have created your own screen layout, as have each of the other development teams in your organization. Your users are complaining about the difficulty of navigating the inconsistent interfaces.

Benefits of establishing this best practice: Use shared standards to support the adoption of best practices and to maximize the benefits of development efforts where standards satisfy requirements for multiple applications or organizations.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Share design standards: Share existing best practices, design standards, checklists, operating procedures, and guidance and governance requirements across teams to reduce complexity and maximize the benefits from development efforts. Ensure that procedures exist to request changes, additions, and exceptions to design standards to support continual improvement and innovation. Ensure that teams are aware of published content so that they can take advantage of content, and limit rework and wasted effort.
 - [Delegating access to your AWS environment](#)
 - [Share an AWS CodeCommit repository](#)
 - [Easy authorization of AWS Lambda functions](#)
 - [Sharing an AMI with specific AWS accounts](#)
 - [Speed template sharing with an AWS CloudFormation designer URL](#)
 - [Using AWS Lambda with Amazon SNS](#)

Resources

Related documents:

- [Easy authorization of AWS Lambda functions](#)
- [Share an AWS CodeCommit repository](#)
- [Sharing an AMI with specific AWS accounts](#)
- [Speed template sharing with an AWS CloudFormation designer URL](#)
- [Using AWS Lambda with Amazon SNS](#)

Related videos:

- [Delegating access to your AWS environment](#)

OPS05-BP07 Implement practices to improve code quality

Implement practices to improve code quality and minimize defects. Some examples include test-driven development, code reviews, and standards adoption.

On AWS, you can integrate services such as [Amazon CodeGuru](#) with your pipeline to automatically [identify potential code and security issues](#) using program analysis and machine learning. CodeGuru provides recommendations on how to implement the AWS best practices to address these issues.

Common anti-patterns:

- To be able to test your feature sooner, you have decided to not integrate your standard input sanitization library. After testing, you commit your code without remembering to complete incorporation of the library.
- You have minimal experience with the dataset you are processing and are unaware that there are a series of edge cases that can exist in your dataset. Those edge cases are not compatible with the code that you have implemented.

Benefits of establishing this best practice: By adopting practices to improve code quality, you can help minimize issues introduced to production.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Implement practices to improve code quality: Implement practices to improve code quality to minimize defects and the risk of their being deployed. For example, test-driven development, pair programming, code reviews, and standards adoption.
 - [Amazon CodeGuru](#)

Resources

Related documents:

- [Amazon CodeGuru](#)

OPS05-BP08 Use multiple environments

Use multiple environments to experiment, develop, and test your workload. Use increasing levels of controls as environments approach production to gain confidence your workload will operate as intended when deployed.

Common anti-patterns:

- You are performing development in a shared development environment and another developer overwrites your code changes.
- The restrictive security controls on your shared development environment are preventing you from experimenting with new services and features.
- You perform load testing on your production systems and cause an outage for your users.
- A critical error resulting in data loss has occurred in production. In your production environment, you attempt to recreate the conditions that lead to the data loss so that you can identify how it happened and prevent it from happening again. To prevent further data loss during testing, you are forced to make the application unavailable to your users.
- You are operating a multi-tenant service and are unable to support a customer request for a dedicated environment.
- You may not always test, but when you do it's in production.
- You believe that the simplicity of a single environment overrides the scope of impact of changes within the environment.

Benefits of establishing this best practice: By deploying multiple environments you can support multiple simultaneous development, testing, and production environments without creating conflicts between developers or user communities.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use multiple environments: Provide developers sandbox environments with minimized controls to enable experimentation. Provide individual development environments to enable work in parallel, increasing development agility. Implement more rigorous controls in the environments approaching production to allow developers to innovate. Use infrastructure as code and configuration management systems to deploy environments that are configured consistent with the controls present in production to ensure systems operate as expected when deployed. When environments are not in use, turn them off to avoid costs associated with idle resources (for example, development systems on evenings and weekends). Deploy production equivalent environments when load testing to enable valid results.
 - [What is AWS CloudFormation?](#)
 - [How do I stop and start Amazon EC2 instances at regular intervals using AWS Lambda?](#)

Resources

Related documents:

- [How do I stop and start Amazon EC2 instances at regular intervals using AWS Lambda?](#)
- [What is AWS CloudFormation?](#)

OPS05-BP09 Make frequent, small, reversible changes

Frequent, small, and reversible changes reduce the scope and impact of a change. This eases troubleshooting, enables faster remediation, and provides the option to roll back a change.

Common anti-patterns:

- You deploy a new version of your application quarterly.
- You frequently make changes to your database schema.
- You perform manual in-place updates, overwriting existing installations and configurations.

Benefits of establishing this best practice: You recognize benefits from development efforts faster by deploying small changes frequently. When the changes are small, it is much easier to identify if they have unintended consequences. When the changes are reversible, there is less risk to implementing the change as recovery is simplified.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Make frequent, small, reversible changes: Frequent, small, and reversible changes reduce the scope and impact of a change. This eases troubleshooting, enables faster remediation, and provides the option to roll back a change. It also increases the rate at which you can deliver value to the business.

OPS05-BP10 Fully automate integration and deployment

Automate build, deployment, and testing of the workload. This reduces errors caused by manual processes and reduces the effort to deploy changes.

Apply metadata using [Resource Tags](#) and [AWS Resource Groups](#) following a consistent [tagging strategy](#) to enable identification of your resources. Tag your resources for organization, cost accounting, access controls, and targeting the execution of automated operations activities.

Common anti-patterns:

- On Friday you, finish authoring the new code for your feature branch. On Monday, after running your code quality test scripts and each of your unit tests scripts, you will check in your code for the next scheduled release.
- You are assigned to code a fix for a critical issue impacting a large number of customers in production. After testing the fix, you commit your code and email change management to request approval to deploy it to production.

Benefits of establishing this best practice: By implementing automated build and deployment management systems, you reduce errors caused by manual processes and reduce the effort to deploy changes enabling your team members to focus on delivering business value.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Use build and deployment management systems: Use build and deployment management systems to track and implement change, to reduce errors caused by manual processes, and reduce the level of effort. Fully automate the integration and deployment pipeline from code check-in through build, testing, deployment, and validation. This reduces lead time, enables increased frequency of change, and reduces the level of effort.
 - [What is AWS CodeBuild?](#)
 - [Continuous integration best practices for software development](#)
 - [Slalom: CI/CD for serverless applications on AWS](#)
 - [Introduction to AWS CodeDeploy - automated software deployment with Amazon Web Services](#)
 - [What is AWS CodeDeploy?](#)

Resources

Related documents:

- [What is AWS CodeBuild?](#)
- [What is AWS CodeDeploy?](#)

Related videos:

- [Continuous integration best practices for software development](#)
- [Introduction to AWS CodeDeploy - automated software deployment with Amazon Web Services](#)
- [Slalom: CI/CD for serverless applications on AWS](#)

OPS 6 How do you mitigate deployment risks?

Adopt approaches that provide fast feedback on quality and enable rapid recovery from changes that do not have desired outcomes. Using these practices mitigates the impact of issues introduced through the deployment of changes.

Best practices

- [OPS06-BP01 Plan for unsuccessful changes \(p. 80\)](#)
- [OPS06-BP02 Test and validate changes \(p. 80\)](#)
- [OPS06-BP03 Use deployment management systems \(p. 81\)](#)
- [OPS06-BP04 Test using limited deployments \(p. 82\)](#)
- [OPS06-BP05 Deploy using parallel environments \(p. 83\)](#)
- [OPS06-BP06 Deploy frequent, small, reversible changes \(p. 84\)](#)
- [OPS06-BP07 Fully automate integration and deployment \(p. 84\)](#)
- [OPS06-BP08 Automate testing and rollback \(p. 85\)](#)

OPS06-BP01 Plan for unsuccessful changes

Plan to revert to a known good state, or remediate in the production environment if a change does not have the desired outcome. This preparation reduces recovery time through faster responses.

Common anti-patterns:

- You performed a deployment and your application has become unstable but there appear to be active users on the system. You have to decide whether to roll back the change and impact the active users or wait to roll back the change knowing the users may be impacted regardless.
- After making a routine change, your new environments are accessible but one of your subnets has become unreachable. You have to decide whether to roll back everything or try to fix the inaccessible subnet. While you are making that determination, the subnet remains unreachable.

Benefits of establishing this best practice: Having a plan in place reduces the mean time to recover (MTTR) from unsuccessful changes, reducing the impact to your end users.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Plan for unsuccessful changes: Plan to revert to a known good state (that is, roll back the change), or remediate in the production environment (that is, roll forward the change) if a change does not have the desired outcome. When you identify changes that you cannot roll back if unsuccessful, apply due diligence prior to committing the change.

OPS06-BP02 Test and validate changes

Test changes and validate the results at all lifecycle stages to confirm new features and minimize the risk and impact of failed deployments.

On AWS, you can create temporary parallel environments to lower the risk, effort, and cost of experimentation and testing. Automate the deployment of these environments using [AWS CloudFormation](#) to ensure consistent implementations of your temporary environments.

Common anti-patterns:

- You deploy a cool new feature to your application. It doesn't work. You don't know.
- You update your certificates. You accidentally install the certificates to the wrong components. You don't know.

Benefits of establishing this best practice: By testing and validating changes following deployment you are able to identify issues early providing an opportunity to mitigate the impact on your customers.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Test and validate changes: Test changes and validate the results at all lifecycle stages (for example, development, test, and production), to confirm new features and minimize the risk and impact of failed deployments.
 - [AWS Cloud9](#)
 - [What is AWS Cloud9?](#)
 - [How to test and debug AWS CodeDeploy locally before you ship your code](#)

Resources

Related documents:

- [AWS Cloud9](#)
- [AWS Developer Tools](#)
- [How to test and debug AWS CodeDeploy locally before you ship your code](#)
- [What is AWS Cloud9?](#)

OPS06-BP03 Use deployment management systems

Use deployment management systems to track and implement change. This reduces errors caused by manual processes and reduces the effort to deploy changes.

In AWS, you can build Continuous Integration/Continuous Deployment (CI/CD) pipelines using services such as [AWS Developer Tools](#) (for example, AWS CodeCommit, [AWS CodeBuild](#), [AWS CodePipeline](#), [AWS CodeDeploy](#), and [AWS CodeStar](#)).

Common anti-patterns:

- You manually deploy updates to the application servers across your fleet and a number of servers become unresponsive due to update errors.
- You manually deploy to your application server fleet over the course of many hours. The inconsistency in versions during the change causes unexpected behaviors.

Benefits of establishing this best practice: Adopting deployment management systems reduces the level of effort to deploy changes, and the frequency of errors caused by manual procedures.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use deployment management systems: Use deployment management systems to track and implement change. This will reduce errors caused by manual processes, and reduce the level of effort to deploy changes. Automate the integration and deployment pipeline from code check-in through testing, deployment, and validation. This reduces lead time, enables increased frequency of change, and further reduces the level of effort.
 - [Introduction to AWS CodeDeploy - automated software deployment with Amazon Web Services](#)
 - [What is AWS CodeDeploy?](#)
 - [What is AWS Elastic Beanstalk?](#)
 - [What is Amazon API Gateway?](#)

Resources

Related documents:

- [AWS CodeDeploy User Guide](#)
- [AWS Developer Tools](#)
- [Try a Sample Blue/Green Deployment in AWS CodeDeploy](#)
- [What is AWS CodeDeploy?](#)
- [What is AWS Elastic Beanstalk?](#)
- [What is Amazon API Gateway?](#)

Related videos:

- [Deep Dive on Advanced Continuous Delivery Techniques Using AWS](#)
- [Introduction to AWS CodeDeploy - automated software deployment with Amazon Web Services](#)

OPS06-BP04 Test using limited deployments

Test with limited deployments alongside existing systems to confirm desired outcomes prior to full scale deployment. For example, use deployment canary testing or one-box deployments.

Common anti-patterns:

- You deploy an unsuccessful change to all of production all at once. You don't know.

Benefits of establishing this best practice: By testing and validating changes following limited deployment you are able to identify issues early with minimal impact on your customers providing an opportunity to further mitigate the impact on your customers.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Test using limited deployments: Test with limited deployments alongside existing systems to confirm desired outcomes prior to full scale deployment. For example, use deployment canary testing or one-box deployments.
 - [AWS CodeDeploy User Guide](#)
 - [Blue/Green deployments with AWS Elastic Beanstalk](#)
 - [Set up an API Gateway canary release deployment](#)
 - [Try a Sample Blue/Green Deployment in AWS CodeDeploy](#)

- [Working with deployment configurations in AWS CodeDeploy](#)

Resources

Related documents:

- [AWS CodeDeploy User Guide](#)
- [Blue/Green deployments with AWS Elastic Beanstalk](#)
- [Set up an API Gateway canary release deployment](#)
- [Try a Sample Blue/Green Deployment in AWS CodeDeploy](#)
- [Working with deployment configurations in AWS CodeDeploy](#)

OPS06-BP05 Deploy using parallel environments

Implement changes onto parallel environments, and then transition over to the new environment. Maintain the prior environment until there is confirmation of successful deployment. Doing so minimizes recovery time by enabling rollback to the previous environment.

Common anti-patterns:

- You perform a mutable deployment by modifying your existing systems. After discovering that the change was unsuccessful, you are forced to modify the systems again to restore the old version extending your time to recovery.
- During a maintenance window, you decommission the old environment and then start building your new environment. Many hours into the procedure, you discover unrecoverable issues with the deployment. While extremely tired, you are forced to find the previous deployment procedures and start rebuilding the old environment.

Benefits of establishing this best practice: By using parallel environments, you can pre-deploy the new environment and transition over to them when desired. If the new environment is not successful, you can recover quickly by transitioning back to your original environment.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Deploy using parallel environments: Implement changes onto parallel environments, and transition or cut over to the new environment. Maintain the prior environment until there is confirmation of successful deployment. This minimizes recovery time by enabling rollback to the previous environment. For example, use immutable infrastructures with blue/green deployments.
 - [Working with deployment configurations in AWS CodeDeploy](#)
 - [Blue/Green deployments with AWS Elastic Beanstalk](#)
 - [Set up an API Gateway canary release deployment](#)
 - [Try a Sample Blue/Green Deployment in AWS CodeDeploy](#)

Resources

Related documents:

- [AWS CodeDeploy User Guide](#)
- [Blue/Green deployments with AWS Elastic Beanstalk](#)
- [Set up an API Gateway canary release deployment](#)

- [Try a Sample Blue/Green Deployment in AWS CodeDeploy](#)
- [Working with deployment configurations in AWS CodeDeploy](#)

Related videos:

- [Deep Dive on Advanced Continuous Delivery Techniques Using AWS](#)

OPS06-BP06 Deploy frequent, small, reversible changes

Use frequent, small, and reversible changes to reduce the scope of a change. This results in easier troubleshooting and faster remediation with the option to roll back a change.

Common anti-patterns:

- You deploy a new version of your application quarterly.
- You frequently make changes to your database schema.
- You perform manual in-place updates, overwriting existing installations and configurations.

Benefits of establishing this best practice: You recognize benefits from development efforts faster by deploying small changes frequently. When the changes are small it is much easier to identify if they have unintended consequences. When the changes are reversible there is less risk to implementing the change as recovery is simplified.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Deploy frequent, small, reversible changes: Use frequent, small, and reversible changes to reduce the scope of a change. This results in easier troubleshooting and faster remediation with the option to roll back a change.

OPS06-BP07 Fully automate integration and deployment

Automate build, deployment, and testing of the workload. This reduces errors caused by manual processes and reduces the effort to deploy changes.

Apply metadata using [Resource Tags](#) and [AWS Resource Groups](#) following a consistent [tagging strategy](#) to enable identification of your resources. Tag your resources for organization, cost accounting, access controls, and targeting the execution of automated operations activities.

Common anti-patterns:

- On Friday, you finish authoring the new code for your feature branch. On Monday, after running your code quality test scripts and each of your unit tests scripts, you will check in your code for the next scheduled release.
- You are assigned to code a fix for a critical issue impacting a large number of customers in production. After testing the fix, you commit your code and email change management to request approval to deploy it to production.

Benefits of establishing this best practice: By implementing automated build and deployment management systems you reduce errors caused by manual processes and reduce the effort to deploy changes enabling your team members to focus on delivering business value.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Use build and deployment management systems: Use build and deployment management systems to track and implement change, to reduce errors caused by manual processes, and reduce the level of effort. Fully automate the integration and deployment pipeline from code check-in through build, testing, deployment, and validation. This reduces lead time, enables increased frequency of change, and reduces the level of effort.
 - [What is AWS CodeBuild?](#)
 - [Continuous integration best practices for software development](#)
 - [Slalom: CI/CD for serverless applications on AWS](#)
 - [Introduction to AWS CodeDeploy - automated software deployment with Amazon Web Services](#)
 - [What is AWS CodeDeploy?](#)
 - [Deep Dive on Advanced Continuous Delivery Techniques Using AWS](#)

Resources

Related documents:

- [Try a Sample Blue/Green Deployment in AWS CodeDeploy](#)
- [What is AWS CodeBuild?](#)
- [What is AWS CodeDeploy?](#)

Related videos:

- [Continuous integration best practices for software development](#)
- [Deep Dive on Advanced Continuous Delivery Techniques Using AWS](#)
- [Introduction to AWS CodeDeploy - automated software deployment with Amazon Web Services](#)
- [Slalom: CI/CD for serverless applications on AWS](#)

OPS06-BP08 Automate testing and rollback

Automate testing of deployed environments to confirm desired outcomes. Automate rollback to a previous known good state when outcomes are not achieved to minimize recovery time and reduce errors caused by manual processes.

Common anti-patterns:

- You deploy changes to your workload. After you see that the change is complete, you start post deployment testing. After you see that they are complete, you realize that your workload is inoperable and customers are disconnected. You then begin rolling back to the previous version. After an extended time to detect the issue, the time to recover is extended by your manual redeployment.

Benefits of establishing this best practice: By testing and validating changes following deployment, you are able to identify issues immediately. By automatically rolling back to the previous version, the impact on your customers is minimized.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Automate testing and rollback: Automate testing of deployed environments to confirm desired outcomes. Automate rollback to a previous known good state when outcomes are not achieved to

minimize recovery time and reduce errors caused by manual processes. For example, perform detailed synthetic user transactions following deployment, verify the results, and roll back on failure.

- [Redeploy and roll back a deployment with AWS CodeDeploy](#)

Resources

Related documents:

- [Redeploy and roll back a deployment with AWS CodeDeploy](#)

OPS 7 How do you know that you are ready to support a workload?

Evaluate the operational readiness of your workload, processes and procedures, and personnel to understand the operational risks related to your workload.

Best practices

- [OPS07-BP01 Ensure personnel capability \(p. 86\)](#)
- [OPS07-BP02 Ensure consistent review of operational readiness \(p. 87\)](#)
- [OPS07-BP03 Use runbooks to perform procedures \(p. 89\)](#)
- [OPS07-BP04 Use playbooks to investigate issues \(p. 90\)](#)
- [OPS07-BP05 Make informed decisions to deploy systems and changes \(p. 91\)](#)

OPS07-BP01 Ensure personnel capability

Have a mechanism to validate that you have the appropriate number of trained personnel to provide support for operational needs. Train personnel and adjust personnel capacity as necessary to maintain effective support.

You will need to have enough team members to cover all activities (including on-call). Ensure that your teams have the necessary skills to be successful with training on your workload, your operations tools, and AWS.

AWS provides resources, including the [AWS Getting Started Resource Center](#), [AWS Blogs](#), [AWS Online Tech Talks](#), [AWS Events and Webinars](#), and the [AWS Well-Architected Labs](#), that provide guidance, examples, and detailed walkthroughs to educate your teams. Additionally, [AWS Training and Certification](#) provides some free training through self-paced digital courses on AWS fundamentals. You can also register for instructor-led training to further support the development of your teams' AWS skills.

Common anti-patterns:

- Deploying a workload without team members skilled to support the platform and services in use.
- Deploying a workload without team members available during intended hours of support.
- Deploying a workload without sufficient team members to support it if there are team members on leave or out sick.
- Deploying additional workloads without reviewing the additional impact on team members support it and other workloads.

Benefits of establishing this best practice: Having skilled team members enables effective support of your workload.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Personnel capability:** Validate that there are sufficient trained personnel to effectively support the workload.
 - **Team size:** Ensure that you have enough team members to cover operational activities, including on-call duties.
 - **Team skill:** Ensure that your team members have sufficient training on AWS, your workload, and your operations tools to perform their duties.
 - [AWS Events and Webinars](#)
 - [Welcome to AWS Training and Certification](#)
- **Review capabilities:** Review team size and skill as operating conditions and workloads change, to ensure there is sufficient capability to maintain operational excellence. Make adjustments to ensure that team size and skill match the operational requirements for the workloads that the team supports.

Resources

Related documents:

- [AWS Blogs](#)
- [AWS Events and Webinars](#)
- [AWS Getting Started Resource Center](#)
- [AWS Online Tech Talks](#)
- [Welcome to AWS Training and Certification](#)

Related examples:

- [Well-Architected Labs](#)

OPS07-BP02 Ensure consistent review of operational readiness

Ensure you have a consistent review of your readiness to operate a workload. Reviews must include, at a minimum, the operational readiness of the teams and the workload, and security requirements. Implement review activities in code and trigger automated review in response to events where appropriate, to ensure consistency, speed of delivery, and reduce errors caused by manual processes.

You should automate workload configuration testing by making baselines using [AWS Config](#) and checking your configurations using [AWS Config Rules](#). You can evaluate security requirements and compliance using the services and features of [AWS Security Hub](#). These services will aid in determining if your workloads are aligned with best practices and standards.

Common anti-patterns:

- Planning to deploy a workload without understanding if team members have the capability to support it.
- Planning to deploy a workload without understanding the procedures necessary to manage it.
- Planning to deploy a workload without knowing the processes to diagnose issues or respond to incidents.
- Planning to deploy a workload without understanding the security risks present in the workload.
- Planning to deploy a workload without understanding if the workload complies with your governance and standards.

Benefits of establishing this best practice: Having skilled team members enables effective support of your workload.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Ensure consistent review of operational readiness: Ensure you have a consistent review of your readiness to operate a workload. Review must include at a minimum the operational readiness of the teams and the workload, and security considerations. Review elements can be hard requirements or you can make a risk-based decision to operate a workload that does not satisfy all requirements. Review elements can be specific to a workload, architecture, or can be implementation dependent. Implement reviews as code and trigger reviews in response to events where appropriate, to ensure consistency, speed of execution, and reduce errors caused by manual processes.
 - [AWS Well-Architected Guidance: Operational Readiness Reviews \(ORR\) Whitepaper](#)
 - [AWS Systems Manager](#)
 - [AWS Config Rules dynamic compliance checking for cloud resources](#)
 - [How to audit your AWS resources for security compliance by using custom AWS Config Rules](#)
 - [How to track configuration changes to CloudFormation stacks using AWS Config](#)
 - [Amazon Inspector update assessment reporting, proxy support, and more](#)
- Create checklists: Ensure you have a consistent review of your readiness to operate a workload. Create operational readiness checklists and validate them against your business, development, operations, and governance requirements. Ensure they address: governance, best practices, configuration standards, restoration procedures, monitoring, maintenance procedures, IT operations procedures, and staffing.
- Use checklists: Make checklists accessible to developers so that they can develop to the appropriate standards. Evaluate checklists when moving between lifecycle stages and environments so that you can identify issues early, when the level of effort to remediate issues is lower. Use the results of checklists to make informed decisions about benefits and risks when considering promoting changes between environments.
- Implement checklists as code and trigger running the checklist in response to events: Implement checklists as code and automatically start running the checklist in response to events where possible, to enhance speed, ensure consistency, and reduce errors caused by manual processes. Integrate automated checklists into deployment pipelines.
 - [AWS Config](#)
 - [What is AWS Config?](#)
 - [AWS Config: evaluating resources with Rules](#)

Resources

Related documents:

- [AWS Config](#)
- [AWS Config Rules dynamic compliance checking for cloud resources](#)
- [AWS Config: evaluating resources with Rules](#)
- [AWS Systems Manager](#)
- [AWS Well-Architected Guidance: Operational Readiness Reviews \(ORR\) Whitepaper](#)
- [Amazon Inspector update assessment reporting, proxy support, and more](#)
- [How to audit your AWS resources for security compliance by using custom AWS Config Rules](#)
- [How to track configuration changes to CloudFormation stacks using AWS Config](#)
- [What is AWS Config?](#)

OPS07-BP03 Use runbooks to perform procedures

Runbooks are documented procedures to achieve specific outcomes. Enable consistent and prompt responses to well-understood events by documenting procedures in runbooks. Implement runbooks as code and trigger the runbooks in response to events where appropriate, to ensure consistency, speed responses, and reduce errors caused by manual processes.

Common anti-patterns:

- Planning to deploy a workload without knowing the procedures necessary to manage it.

Benefits of establishing this best practice: Capturing runbooks ensures that procedures can be consistently followed. Codifying your runbooks limits the introduction of errors from manual activity. Automating runbooks shortens the time to respond to an event by eliminating the requirement for team member intervention.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use runbooks to perform standard procedures: Runbooks are documented procedures to achieve specific outcomes. Enable consistent and prompt responses to well understood events by documenting procedures in runbooks. Runbooks must contain the minimum information for an adequately skilled person to achieve the desired outcome. For example, required permissions, required tools, constraints on performing the procedure (for example, specific maintenance windows), and execution steps.
- Implement runbooks as code: Perform your operations as code by implementing your runbooks as code to ensure consistency and reduce errors caused by manual processes
 - [AWS Systems Manager Run Command](#)
 - [AWS Systems Manager – Working with Runbooks](#)
 - [AWS Systems Manager Automation](#)
 - [What is AWS Lambda?](#)
- Trigger runbooks in response to events: Trigger the execution of runbook code in response to observed events when appropriate. This increases the speed of the response and reduces the level of effort to respond.
 - [What is Amazon CloudWatch Events?](#)
 - [Creating a CloudWatch Events rule that triggers on an event](#)
 - [Creating a CloudWatch Events rule that triggers on an AWS API call using AWS CloudTrail](#)
 - [CloudWatch Events event examples from supported services](#)
 - [Using Amazon CloudWatch Alarms](#)

Resources

Related documents:

- [AWS Systems Manager Automation](#)
- [AWS Systems Manager Run Command](#)
- [AWS Systems Manager – Working with Runbooks](#)
- [CloudWatch Events event examples from supported services](#)
- [Creating a CloudWatch Events rule that triggers on an AWS API call using AWS CloudTrail](#)
- [Creating a CloudWatch Events rule that triggers on an event](#)
- [Using Amazon CloudWatch Alarms](#)
- [What is AWS Lambda?](#)

- [What is Amazon CloudWatch Events?](#)

Related examples:

- [Automating operations with Playbooks and Runbooks](#)

OPS07-BP04 Use playbooks to investigate issues

Enable consistent and prompt responses to issues that are not well understood by documenting the investigation process in playbooks. Playbooks are the predefined steps performed to identify the factors contributing to a failure scenario. The results from any process step are used to determine the next steps to take until the issue is identified or escalated.

AWS allows you to treat your operations as code, scripting your runbook and playbook activities to reduce the risk of human error. You can use [Resource Tags](#) or [Resource Groups](#) with your scripts to selectively execute based on criteria you have defined (for example, environment, owner, role, or version).

You can use scripted procedures to enable automation by starting the scripts in response to events. By treating both your operations and workloads as code, you can also script and automate the evaluation of your environments.

You should script procedures on your instances using [AWS Systems Manager \(SSM\) Run Command](#), use [AWS Systems Manager Automation](#) to script actions and create workflows on instances and other resources, or use [AWS Lambda](#) serverless compute functions to script responses to events across AWS service APIs and your own custom interfaces. You can also use [AWS Step Functions](#) to coordinate multiple AWS services scripted into serverless workflows. Automate your responses by initiating these scripts using [CloudWatch Events](#) and route desired events to additional operations support systems using [Amazon EventBridge](#).

You should test your procedures, failure scenarios, and the success of your responses (for example, by holding game days and testing prior to going live) to identify areas you need to plan to address.

On AWS, you can create temporary parallel environments to lower the risk, effort, and cost of experimentation and testing. Automate the deployment of these environments using [AWS CloudFormation](#) to ensure consistent implementations of your temporary environments. Perform failure injection testing in safe environments where there will be acceptable or no customer impact, and develop or revise appropriate responses.

Common anti-patterns:

- Planning to deploy a workload without knowing the processes to diagnose issues or respond to incidents.

Benefits of establishing this best practice: Creating playbooks ensures that processes can be consistently followed. Codifying your playbooks limits the introduction of errors from manual activity. Automating playbooks shortens the time to respond to an event by eliminating the requirement for team member intervention or providing them additional information when their intervention begins.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use playbooks to identify issues: Playbooks are documented processes to investigate issues. Enable consistent and prompt responses to failure scenarios by documenting processes in playbooks. Playbooks must contain the information and guidance necessary for an adequately skilled person to gather applicable information, identify potential sources of failure, isolate faults, and determine contributing factors (that is, perform root cause analysis).

- Implement playbooks as code: Perform your operations as code by scripting your playbooks to ensure consistency and limit reduce errors caused by manual processes. Playbooks can be composed of multiple scripts representing the different steps that might be necessary to identify the contributing factors to an issue. Runbook activities can be triggered or performed as part of playbook activities, or may prompt for execution of a playbook in response to identified events.
 - [Automate your operational playbooks with AWS Systems Manager](#)
 - [AWS Systems Manager Run Command](#)
 - [AWS Systems Manager Automation](#)
 - [What is AWS Lambda?](#)
 - [What is Amazon CloudWatch Events?](#)
 - [Using Amazon CloudWatch Alarms](#)

Resources

Related documents:

- [Amazon EventBridge](#)
- [AWS Step Functions](#)
- [AWS Systems Manager Automation](#)
- [AWS Systems Manager Run Command](#)
- [Automate your operational playbooks with AWS Systems Manager](#)
- [Using Amazon CloudWatch Alarms](#)
- [What is AWS Lambda?](#)
- [What is Amazon CloudWatch Events?](#)

Related examples:

- [Automating operations with Playbooks and Runbooks](#)

OPS07-BP05 Make informed decisions to deploy systems and changes

Evaluate the capabilities of the team to support the workload and the workload's compliance with governance. Evaluate these against the benefits of deployment when determining whether to transition a system or change into production. Understand the benefits and risks to make informed decisions.

A pre-mortem is an exercise where a team simulates a failure to develop mitigation strategies. Use pre-mortems to anticipate failure and create procedures where appropriate. When you make changes to the checklists you use to evaluate your workloads, plan what you will do with live systems that no longer comply.

Common anti-patterns:

- Deciding to deploy a workload without understanding the security risks present in the workload.
- Deciding to deploy a workload without understanding if it complies with your governance and standards.
- Deciding to deploy a workload without understanding if your team can support it.
- Deciding to deploy a workload without understanding how it benefits the organization.

Benefits of establishing this best practice: Having skilled team members enables effective support of your workload.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Make informed decisions to deploy workloads and changes: Evaluate the capabilities of the team to support the workload and the workload's compliance with governance. Evaluate these against the benefits of deployment when determining whether to transition a system or change into production. Understand the benefits and risks, and make informed decisions.

Operate

Questions

- [OPS 8 How do you understand the health of your workload? \(p. 92\)](#)
- [OPS 9 How do you understand the health of your operations? \(p. 98\)](#)
- [OPS 10 How do you manage workload and operations events? \(p. 104\)](#)

OPS 8 How do you understand the health of your workload?

Define, capture, and analyze workload metrics to gain visibility to workload events so that you can take appropriate action.

Best practices

- [OPS08-BP01 Identify key performance indicators \(p. 92\)](#)
- [OPS08-BP02 Define workload metrics \(p. 93\)](#)
- [OPS08-BP03 Collect and analyze workload metrics \(p. 93\)](#)
- [OPS08-BP04 Establish workload metrics baselines \(p. 95\)](#)
- [OPS08-BP05 Learn expected patterns of activity for workload \(p. 95\)](#)
- [OPS08-BP06 Alert when workload outcomes are at risk \(p. 96\)](#)
- [OPS08-BP07 Alert when workload anomalies are detected \(p. 97\)](#)
- [OPS08-BP08 Validate the achievement of outcomes and the effectiveness of KPIs and metrics \(p. 97\)](#)

OPS08-BP01 Identify key performance indicators

Identify key performance indicators (KPIs) based on desired business outcomes (for example, order rate, customer retention rate, and profit versus operating expense) and customer outcomes (for example, customer satisfaction). Evaluate KPIs to determine workload success.

Common anti-patterns:

- You are asked by business leadership how successful a workload has been serving business needs but have no frame of reference to determine success.
- You are unable to determine if the commercial off-the-shelf application you operate for your organization is cost-effective.

Benefits of establishing this best practice: By identifying key performance indicators you enable achieving business outcomes as the test of the health and success of your workload.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- **Identify key performance indicators:** Identify key performance indicators (KPIs) based on desired business and customer outcomes. Evaluate KPIs to determine workload success.

OPS08-BP02 Define workload metrics

Define workload metrics to measure the achievement of KPIs (for example, abandoned shopping carts, orders placed, cost, price, and allocated workload expense). Define workload metrics to measure the health of the workload (for example, interface response time, error rate, requests made, requests completed, and utilization). Evaluate metrics to determine if the workload is achieving desired outcomes, and to understand the health of the workload.

You should send log data to a service such as CloudWatch Logs, and generate metrics from observations of necessary log content.

CloudWatch has specialized features such as [Amazon CloudWatch Insights for .NET and SQL Server](#) and [Container Insights](#) that can assist you by identifying and setting up key metrics, logs, and alarms across your specifically supported application resources and technology stack.

Common anti-patterns:

- You have defined standard metrics, not associated to any KPIs or tailored to any workload.
- You have errors in your metrics calculations that will yield invalid results.
- You don't have any metrics defined for your workload.
- You only measure for availability.

Benefits of establishing this best practice: By defining and evaluating workload metrics you can determine the health of your workload and measure the achievement of business outcomes.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Define workload metrics: Define workload metrics to measure the achievement of KPIs. Define workload metrics to measure the health of the workload and its individual components. Evaluate metrics to determine if the workload is achieving desired outcomes, and to understand the health of the workload.
 - [Publish custom metrics](#)
 - [Searching and filtering log data](#)
 - [Amazon CloudWatch metrics and dimensions reference](#)

Resources

Related documents:

- [Amazon CloudWatch metrics and dimensions reference](#)
- [Publish custom metrics](#)
- [Searching and filtering log data](#)

OPS08-BP03 Collect and analyze workload metrics

Perform regular proactive reviews of metrics to identify trends and determine where appropriate responses are needed.

You should aggregate log data from your application, workload components, services, and API calls to a service such as CloudWatch Logs. Generate metrics from observations of necessary log content to enable insight into the performance of operations activities.

On AWS, you can analyze workload metrics and identify operational issues using the machine learning capabilities of [Amazon DevOps Guru](#). AWS DevOps Guru provides notification of operational issues with [targeted and proactive](#) recommendations to resolve issues and maintain application health.

In the AWS Shared Responsibility Model, portions of monitoring are delivered to you through the [AWS Health Dashboard](#). This dashboard provides alerts and remediation guidance when AWS is experiencing events that might affect you. Customers with Business and Enterprise Support subscriptions also get access to the [AWS Health API](#), enabling integration to their event management systems.

On AWS, you can [export your log data to Amazon S3](#) or [send logs directly to Amazon S3](#) for long-term storage. Using [AWS Glue](#), you can discover and prepare your log data in Amazon S3 for analytics, storing associated metadata in the [AWS Glue Data Catalog](#). [Amazon Athena](#), through its native integration with AWS Glue, can then be used to analyze your log data, querying it using standard SQL. Using a business intelligence tool like [Amazon QuickSight](#) you can visualize, explore, and analyze your data.

An alternative [solution](#) would be to use the [Amazon OpenSearch Service](#) and [OpenSearch Dashboards](#) to collect, analyze, and display logs on AWS across multiple accounts and AWS Regions.

Common anti-patterns:

- You are asked by the network design team for current network bandwidth utilization rates. You provide the current metrics, network utilization is at 35%. They reduce circuit capacity as a cost savings measure causing widespread connectivity issues as your point-in-time measurement did not reflect the trend in utilization rates.
- Your router has failed. It has been logging non-critical memory errors with greater and greater frequency up until its complete failure. You did not detect this trend and as a result did not replace the faulty memory before the router caused a service interruption.

Benefits of establishing this best practice: By collecting and analyzing your workload metrics you gain understanding of the health of your workload and can gain insight to trends that may have an impact on your workload or the achievement of your business outcomes.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Collect and analyze workload metrics: Perform regular proactive reviews of metrics to identify trends and determine where appropriate responses are needed.
 - [Using Amazon CloudWatch metrics](#)
 - [Amazon CloudWatch metrics and dimensions reference](#)
 - [Collect metrics and logs from Amazon EC2 instances and on-premises servers with the CloudWatch Agent](#)

Resources

Related documents:

- [Amazon Athena](#)
- [Amazon CloudWatch metrics and dimensions reference](#)
- [Amazon DevOps Guru](#)
- [AWS Glue](#)
- [AWS Glue Data Catalog](#)
- [Amazon OpenSearch Service](#)
- [AWS Health Dashboard](#)
- [Amazon QuickSight](#)

- [Collect metrics and logs from Amazon EC2 instances and on-premises servers with the CloudWatch Agent](#)
- [Using Amazon CloudWatch metrics](#)

OPS08-BP04 Establish workload metrics baselines

Establish baselines for metrics to provide expected values as the basis for comparison and identification of under- and over-performing components. Identify thresholds for improvement, investigation, and intervention.

Common anti-patterns:

- A server is running at 95% CPU utilization you are asked if that is good or bad. CPU utilization on that server has not been baselined so you have no idea if that is good or bad.

Benefits of establishing this best practice: By defining baseline metric values you are able to evaluate current metric values, and metric trends, to determine if action is required.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Establish baselines for workload metrics: Establish baselines for workload metrics to provide expected values as the basis for comparison.
 - [Creating Amazon CloudWatch Alarms](#)

Resources

Related documents:

- [Creating Amazon CloudWatch Alarms](#)

OPS08-BP05 Learn expected patterns of activity for workload

Establish patterns of workload activity to identify anomalous behavior so that you can respond appropriately if required.

CloudWatch through the [CloudWatch Anomaly Detection](#) feature applies statistical and machine learning algorithms to generate a range of expected values that represent normal metric behavior.

[Amazon DevOps Guru](#) can be used to identify anomalous behavior through event correlation, log analysis, and applying machine learning to analyze your workload telemetry. When unexpected behaviors are detected, it provides the [related metrics and events](#) with recommendations to address the behavior.

Common anti-patterns:

- You are reviewing network utilization logs and see that network utilization increased between 11:30am and 1:30pm and then again at 4:30pm through 6:00pm. You are unaware if this should be considered normal or not.
- Your web servers reboot every night at 3:00am. You are unaware if this is an expected behavior.

Benefits of establishing this best practice: By learning patterns of behavior you can recognize unexpected behavior and take action if necessary.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Learn expected patterns of activity for workload: Establish patterns of workload activity to determine when behavior is outside of the expected values so that you can respond appropriately if required.

Resources

Related documents:

- [Amazon DevOps Guru](#)
- [CloudWatch Anomaly Detection](#)

OPS08-BP06 Alert when workload outcomes are at risk

Raise an alert when workload outcomes are at risk so that you can respond appropriately if necessary.

Ideally, you have previously identified a metric threshold that you are able to alarm upon or an event that you can use to trigger an automated response.

On AWS, you can use [Amazon CloudWatch Synthetics](#) to create canary scripts to monitor your endpoints and APIs by performing the same actions as your customers. The telemetry generated and the [insight gained](#) can enable you to identify issues before your customers are impacted.

You can also use [CloudWatch Logs Insights](#) to interactively search and analyze your log data using a purpose-built query language. CloudWatch Logs Insights automatically [discovers fields in logs](#) from AWS services, and custom log events in JSON. It scales with your log volume and query complexity and gives you answers in seconds, helping you to search for the contributing factors of an incident.

Common anti-patterns:

- You have no network connectivity. No one is aware. No one is trying to identify why or taking action to restore connectivity.
- Following a patch, your persistent instances have become unavailable, disrupting users. Your users have opened support cases. No one has been notified. No one is taking action.

Benefits of establishing this best practice: By identifying that business outcomes are at risk and alerting for action to be taken you have the opportunity to prevent or mitigate the impact of an incident.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Alert when workload outcomes are at risk: Raise an alert when workload outcomes are at risk so that you can respond appropriately if required.
 - [What is Amazon CloudWatch Events?](#)
 - [Creating Amazon CloudWatch Alarms](#)
 - [Invoking Lambda functions using Amazon SNS notifications](#)

Resources

Related documents:

- [Amazon CloudWatch Synthetics](#)
- [CloudWatch Logs Insights](#)

- [Creating Amazon CloudWatch Alarms](#)
- [Invoking Lambda functions using Amazon SNS notifications](#)
- [What is Amazon CloudWatch Events?](#)

OPS08-BP07 Alert when workload anomalies are detected

Raise an alert when workload anomalies are detected so that you can respond appropriately if necessary.

Your analysis of your workload metrics over time may establish patterns of behavior that you can quantify sufficiently to define an event or raise an alarm in response.

Once trained, the [CloudWatch Anomaly Detection](#) feature can be used to [alarm](#) on detected anomalies or can provide overlaid expected values onto a [graph](#) of metric data for ongoing comparison.

Common anti-patterns:

- Your retail website sales have increased suddenly and dramatically. No one is aware. No one is trying to identify what lead to this surge. No one is taking action to ensure quality customer experiences under the additional load.
- Following the application of a patch, your persistent servers are rebooting frequently, disrupting users. Your servers typically reboot up to three times but not more. No one is aware. No one is trying to identify why this is happening.

Benefits of establishing this best practice: By understanding patterns of workload behavior, you can identify unexpected behavior and take action if necessary.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Alert when workload anomalies are detected: Raise an alert when workload anomalies are detected so that you can respond appropriately if required.
 - [What is Amazon CloudWatch Events?](#)
 - [Creating Amazon CloudWatch Alarms](#)
 - [Invoking Lambda functions using Amazon SNS notifications](#)

Resources

Related documents:

- [Creating Amazon CloudWatch Alarms](#)
- [CloudWatch Anomaly Detection](#)
- [Invoking Lambda functions using Amazon SNS notifications](#)
- [What is Amazon CloudWatch Events?](#)

OPS08-BP08 Validate the achievement of outcomes and the effectiveness of KPIs and metrics

Create a business-level view of your workload operations to help you determine if you are satisfying needs and to identify areas that need improvement to reach business goals. Validate the effectiveness of KPIs and metrics and revise them if necessary.

AWS also has support for third-party log analysis systems and business intelligence tools through the AWS service APIs and SDKs (for example, Grafana, Kibana, and Logstash).

Common anti-patterns:

- Page response time has never been considered a contributor to customer satisfaction. You have never established a metric or threshold for page response time. Your customers are complaining about slowness.
- You have not been achieving your minimum response time goals. In an effort to improve response time, you have scaled up your application servers. You are now exceeding response time goals by a significant margin and also have significant unused capacity you are paying for.

Benefits of establishing this best practice: By reviewing and revising KPIs and metrics, you understand how your workload supports the achievement of your business outcomes and can identify where improvement is needed to reach business goals.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Validate the achievement of outcomes and the effectiveness of KPIs and metrics: Create a business level view of your workload operations to help you determine if you are satisfying needs and to identify areas that need improvement to reach business goals. Validate the effectiveness of KPIs and metrics and revise them if necessary.
 - [Using Amazon CloudWatch dashboards](#)
 - [What is log analytics?](#)

Resources

Related documents:

- [Using Amazon CloudWatch dashboards](#)
- [What is log analytics?](#)

OPS 9 How do you understand the health of your operations?

Define, capture, and analyze operations metrics to gain visibility to operations events so that you can take appropriate action.

Best practices

- [OPS09-BP01 Identify key performance indicators \(p. 98\)](#)
- [OPS09-BP02 Define operations metrics \(p. 99\)](#)
- [OPS09-BP03 Collect and analyze operations metrics \(p. 100\)](#)
- [OPS09-BP04 Establish operations metrics baselines \(p. 101\)](#)
- [OPS09-BP05 Learn the expected patterns of activity for operations \(p. 101\)](#)
- [OPS09-BP06 Alert when operations outcomes are at risk \(p. 102\)](#)
- [OPS09-BP07 Alert when operations anomalies are detected \(p. 102\)](#)
- [OPS09-BP08 Validate the achievement of outcomes and the effectiveness of KPIs and metrics \(p. 103\)](#)

OPS09-BP01 Identify key performance indicators

Identify key performance indicators (KPIs) based on desired business outcomes (for example, new features delivered) and customer outcomes (for example, customer support cases). Evaluate KPIs to determine operations success.

Common anti-patterns:

- You are asked by business leadership how successful operations is at accomplishing business goals but have no frame of reference to determine success.
- You are unable to determine if your maintenance windows have an impact on business outcomes.

Benefits of establishing this best practice: By identifying key performance indicators you enable achieving business outcomes as the test of the health and success of your operations.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Identify key performance indicators: Identify key performance indicators (KPIs) based on desired business and customer outcomes. Evaluate KPIs to determine operations success.

OPS09-BP02 Define operations metrics

Define operations metrics to measure the achievement of KPIs (for example, successful deployments, and failed deployments). Define operations metrics to measure the health of operations activities (for example, mean time to detect an incident (MTTD), and mean time to recovery (MTTR) from an incident). Evaluate metrics to determine if operations are achieving desired outcomes, and to understand the health of your operations activities.

Common anti-patterns:

- Your operations metrics are based on what the team thinks is reasonable.
- You have errors in your metrics calculations that will yield incorrect results.
- You don't have any metrics defined for your operations activities.

Benefits of establishing this best practice: By defining and evaluating operations metrics you can determine the health of your operations activities and measure the achievement of business outcomes.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Define operations metrics: Define operations metrics to measure the achievement of KPIs. Define operations metrics to measure the health of operations and its activities. Evaluate metrics to determine if operations are achieving desired outcomes, and to understand the health of the operations.
 - [Publish custom metrics](#)
 - [Searching and filtering log data](#)
 - [Amazon CloudWatch metrics and dimensions reference](#)

Resources

Related documents:

- [AWS Answers: Centralized Logging](#)
- [Amazon CloudWatch metrics and dimensions reference](#)
- [Detect and React to Changes in Pipeline State with Amazon CloudWatch Events](#)
- [Publish custom metrics](#)
- [Searching and filtering log data](#)

Related videos:

- Build a Monitoring Plan

OPS09-BP03 Collect and analyze operations metrics

Perform regular, proactive reviews of metrics to identify trends and determine where appropriate responses are needed.

You should aggregate log data from the execution of your operations activities and operations API calls, into a service such as CloudWatch Logs. Generate metrics from observations of necessary log content to gain insight into the performance of operations activities.

On AWS, you can [export your log data to Amazon S3](#) or [send logs directly to Amazon S3](#) for long-term storage. Using [AWS Glue](#), you can discover and prepare your log data in Amazon S3 for analytics, storing associated metadata in the [AWS Glue Data Catalog](#). [Amazon Athena](#), through its native integration with AWS Glue, can then be used to analyze your log data, querying it using standard SQL. Using a business intelligence tool like [Amazon QuickSight](#) you can visualize, explore, and analyze your data.

Common anti-patterns:

- Consistent delivery of new features is considered a key performance indicator. You have no method to measure how frequently deployments occur.
- You log deployments, rolled back deployments, patches, and rolled back patches to track your operations activities, but no one reviews the metrics.
- You have a recovery time objective to restore a lost database within fifteen minutes that was defined when the system was deployed and had no users. You now have ten thousand users and have been operating for two years. A recent restore took over two hours. This was not recorded and no one is aware.

Benefits of establishing this best practice: By collecting and analyzing your operations metrics, you gain understanding of the health of your operations and can gain insight to trends that have may an impact on your operations or the achievement of your business outcomes.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Collect and analyze operations metrics: Perform regular proactive reviews of metrics to identify trends and determine where appropriate responses are needed.
 - [Using Amazon CloudWatch metrics](#)
 - [Amazon CloudWatch metrics and dimensions reference](#)
 - [Collect metrics and logs from Amazon EC2 instances and on-premises servers with the CloudWatch Agent](#)

Resources

Related documents:

- [Amazon Athena](#)
- [Amazon CloudWatch metrics and dimensions reference](#)
- [Amazon QuickSight](#)
- [AWS Glue](#)
- [AWS Glue Data Catalog](#)

- [Collect metrics and logs from Amazon EC2 instances and on-premises servers with the CloudWatch Agent](#)
- [Using Amazon CloudWatch metrics](#)

OPS09-BP04 Establish operations metrics baselines

Establish baselines for metrics to provide expected values as the basis for comparison and identification of under and over performing operations activities.

Common anti-patterns:

- You have been asked what the expected time to deploy is. You have not measured how long it takes to deploy and can not determine expected times.
- You have been asked what how long it takes to recover from an issue with the application servers. You have no information about time to recovery from first customer contact. You have no information about time to recovery from first identification of an issue through monitoring.
- You have been asked how many support personnel are required over the weekend. You have no idea how many support cases are typical over a weekend and can not provide an estimate.
- You have a recovery time objective to restore lost databases within fifteen minutes that was defined when the system was deployed and had no users. You now have ten thousand users and have been operating for two years. You have no information on how the time to restore has changed for your database.

Benefits of establishing this best practice: By defining baseline metric values you are able to evaluate current metric values, and metric trends, to determine if action is required.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Learn expected patterns of activity for operations: Establish patterns of operations activity to determine when behavior is outside of the expected values so that you can respond appropriately if required.

OPS09-BP05 Learn the expected patterns of activity for operations

Establish patterns of operations activities to identify anomalous activity so that you can respond appropriately if necessary.

Common anti-patterns:

- Your deployment failure rate has increased substantially recently. You address each of the failures independently. You do not realize that the failures correspond to deployments by a new employee who is unfamiliar with the deployment management system.

Benefits of establishing this best practice: By learning patterns of behavior, you can recognize unexpected behavior and take action if necessary.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Learn expected patterns of activity for operations: Establish patterns of operations activity to determine when behavior is outside of the expected values so that you can respond appropriately if required.

OPS09-BP06 Alert when operations outcomes are at risk

Raise an alert when operations outcomes are at risk so that you can respond appropriately if necessary.

Ideally, you have previously identified a metric that you are able to alarm upon or an event that you can use to trigger an automated response.

You can also use [CloudWatch Logs Insights](#) to interactively search and analyze your log data using a purpose-built query language. CloudWatch Logs Insights automatically [discovers fields in logs](#) from AWS services, and custom log events in JSON. It scales with your log volume and query complexity and gives you answers in seconds helping you to search for the contributing factors of an incident.

Common anti-patterns:

- You are addressing an incident that has caused significant customer impact. Your customers are creating support cases that are not being responded to within the expected time frame. No one is taking action to address this challenge because no one has been alerted to the issue.
- It is 11:00 pm. You are supposed to deploy by midnight to support a customer event but one of the developers is correcting issues with their code. They estimate that they will have the fix within two hours. You are unlikely to deploy in time. You wait to hear from the developer.
- You are unable to restore the backup. You begin working on your resume.

Benefits of establishing this best practice: By identifying that business outcomes are at risk and alerting for action to be taken, you have the opportunity to avert or mitigate the impact of an incident.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Alert when operations outcomes are at risk: Raise an alert when operations outcomes are at risk so that you can respond appropriately if required.
 - [What is Amazon CloudWatch Events?](#)
 - [Creating Amazon CloudWatch alarms](#)
 - [Invoking Lambda functions using Amazon SNS notifications](#)

Resources

Related documents:

- [AWS Answers: Centralized Logging](#)
- [CloudWatch Logs Insights](#)
- [Creating Amazon CloudWatch alarms](#)
- [Detect and React to Changes in Pipeline State with Amazon CloudWatch Events](#)
- [Invoking Lambda functions using Amazon SNS notifications](#)

OPS09-BP07 Alert when operations anomalies are detected

Raise an alert when operations anomalies are detected so that you can respond appropriately if necessary.

Your analysis of your operations metrics over time may establish patterns of behavior that you can quantify sufficiently to define an event or raise an alarm in response.

Once trained, the [CloudWatch Anomaly Detection](#) feature can be used to [alarm](#) on detected anomalies or can provide overlaid expected values onto a [graph](#) of metric data for ongoing comparison.

[Amazon DevOps Guru](#) can be used to identify anomalous behavior through event correlation, log analysis, and applying machine learning to analyze your workload telemetry. The [insights](#) gained are presented with the relevant data and recommendations.

Common anti-patterns:

- You are applying a patch to your fleet of instances. You tested the patch successfully in the test environment. The patch is failing for a large percentage of instances in your fleet. You do nothing.
- You note that there are deployments starting Friday end of day. Your organization has predefined maintenance windows on Tuesdays and Thursdays. You do nothing.

Benefits of establishing this best practice: By understanding patterns of operations behavior you can identify unexpected behavior and take action if necessary.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Alert when operations anomalies are detected: Raise an alert when operations anomalies are detected so that you can respond appropriately if required.
 - [What is Amazon CloudWatch Events?](#)
 - [Creating Amazon CloudWatch alarms](#)
 - [Invoking Lambda functions using Amazon SNS notifications](#)

Resources

Related documents:

- [Amazon DevOps Guru](#)
- [CloudWatch Anomaly Detection](#)
- [Creating Amazon CloudWatch alarms](#)
- [Detect and React to Changes in Pipeline State with Amazon CloudWatch Events](#)
- [Invoking Lambda functions using Amazon SNS notifications](#)
- [What is Amazon CloudWatch Events?](#)

OPS09-BP08 Validate the achievement of outcomes and the effectiveness of KPIs and metrics

Create a business-level view of your operations activities to help you determine if you are satisfying needs and to identify areas that need improvement to reach business goals. Validate the effectiveness of KPIs and metrics and revise them if necessary.

AWS also has support for third-party log analysis systems and business intelligence tools through the AWS service APIs and SDKs (for example, Grafana, Kibana, and Logstash).

Common anti-patterns:

- The frequency of your deployments has increased with the growth in number of development teams. Your defined expected number of deployments is once per week. You have been regularly deploying daily. When there is an issue with your deployment system, and deployments are not possible, it goes undetected for days.
- When your business previously provided support only during core business hours from Monday to Friday. You established a next business day response time goal for incidents. You have recently started offering 24x7 support coverage with a two hour response time goal. Your overnight staff are

overwhelmed and customers are unhappy. There is no indication that there are issues with incident response times because you are reporting against a next business day target.

Benefits of establishing this best practice: By reviewing and revising KPIs and metrics, you understand how your workload supports the achievement of your business outcomes and can identify where improvement is needed to reach business goals.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Validate the achievement of outcomes and the effectiveness of KPIs and metrics: Create a business level view of your operations activities to help you determine if you are satisfying needs and to identify areas that need improvement to reach business goals. Validate the effectiveness of KPIs and metrics and revise them if necessary.
 - [Using Amazon CloudWatch dashboards](#)
 - [What is log analytics?](#)

Resources

Related documents:

- [Using Amazon CloudWatch dashboards](#)
- [What is log analytics?](#)

OPS 10 How do you manage workload and operations events?

Prepare and validate procedures for responding to events to minimize their disruption to your workload.

Best practices

- [OPS10-BP01 Use processes for event, incident, and problem management \(p. 104\)](#)
- [OPS10-BP02 Have a process per alert \(p. 105\)](#)
- [OPS10-BP03 Prioritize operational events based on business impact \(p. 106\)](#)
- [OPS10-BP04 Define escalation paths \(p. 106\)](#)
- [OPS10-BP05 Enable push notifications \(p. 107\)](#)
- [OPS10-BP06 Communicate status through dashboards \(p. 108\)](#)
- [OPS10-BP07 Automate responses to events \(p. 108\)](#)

OPS10-BP01 Use processes for event, incident, and problem management

Have processes to address observed events, events that require intervention (incidents), and events that require intervention and either recur or cannot currently be resolved (problems). Use these processes to mitigate the impact of these events on the business and your customers by ensuring timely and appropriate responses.

On AWS, you can use [AWS Systems Manager OpsCenter](#) as a central location to view, investigate, and resolve operational issues related to any AWS resource. It aggregates operational issues and provides contextually relevant data to assist in incident response.

Common anti-patterns:

- A customer calls you to complain that your application isn't working. After hanging up you decide to reboot the server and then go back to what you were working on.

- You receive a trouble ticket that one of the applications you support isn't working. After a conversation with your peers you decide to try and find documentation on how it works. You locate some design documentation and start trying to figure out what might be wrong. At the end of your working day, you stop working on the issue and head home for the weekend.
- You receive a call notifying you that the retail site is offline. You share this information with your colleagues and everyone stops what they are doing to work on the issue independently.

Benefits of establishing this best practice: Having predefined processes for event, incident, and problem management enables consistent and timely responses to mitigate the impact on the business and your customers.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use processes for event, incident, and problem management: Have processes to address observed events, events that require intervention (incidents), and events that require intervention and either recur or cannot currently be resolved (problems). Use these processes to mitigate the impact of these events on the business and your customers by ensuring timely and appropriate responses.

Resources

Related documents:

- [AWS Systems Manager OpsCenter](#)

OPS10-BP02 Have a process per alert

Have a well-defined response (runbook or playbook), with a specifically identified owner, for any event for which you raise an alert. This ensures effective and prompt responses to operations events and prevents actionable events from being obscured by less valuable notifications.

Common anti-patterns:

- Your monitoring system presents you a stream of approved connections along with other messages. The volume of messages is so large that you miss periodic error messages that require your intervention.
- You receive an alert that the website is down. There is no defined process for when this happens. You are forced to take an ad hoc approach to diagnose and resolve the issue. Developing this process as you go extends the time to recovery.

Benefits of establishing this best practice: By alerting only when action is required, you prevent low value alerts from concealing high value alerts. By having a process for every actionable alert, you enable a consistent and prompt response to events in your environment.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Process per alert: Any event for which you raise an alert should have a well-defined response (runbook or playbook) with a specifically identified owner (for example, individual, team, or role) accountable for successful completion. Performance of the response may be automated or conducted by another team but the owner is accountable for ensuring the process delivers the expected outcomes. By having these processes, you ensure effective and prompt responses to operations events and you can prevent actionable events from being obscured by less valuable notifications. For example, automatic scaling

might be applied to scale a web front end, but the operations team might be accountable to ensure that the automatic scaling rules and limits are appropriate for workload needs.

Resources

Related documents:

- [Amazon CloudWatch Features](#)
- [What is Amazon CloudWatch Events?](#)

Related videos:

- [Build a Monitoring Plan](#)

OPS10-BP03 Prioritize operational events based on business impact

Ensure that when multiple events require intervention, those that are most significant to the business are addressed first. Impacts can include loss of life or injury, financial loss, or damage to reputation or trust.

Common anti-patterns:

- You receive a support request to add a printer configuration for a user. While working on the issue, you receive a support request stating that your retail site is down. After completing the printer configuration for your user, you start work on the website issue.
- You get notified that both your retail website and your payroll system are down. You don't know which one should get priority.

Benefits of establishing this best practice: Prioritizing responses to the incidents with the greatest impact on the business enables your management of that impact.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Prioritize operational events based on business impact: Ensure that when multiple events require intervention, those that are most significant to the business are addressed first. Impacts can include loss of life or injury, financial loss, regulatory violations, or damage to reputation or trust.

OPS10-BP04 Define escalation paths

Define escalation paths in your runbooks and playbooks, including what triggers escalation, and procedures for escalation. Specifically identify owners for each action to ensure effective and prompt responses to operations events.

Identify when a human decision is required before an action is taken. Work with decision makers to have that decision made in advance, and the action preapproved, so that MTTR is not extended waiting for a response.

Common anti-patterns:

- Your retail site is down. You don't understand the runbook for recovering the site. You start calling colleagues hoping that someone will be able to help you.
- You receive a support case for an unreachable application. You don't have permissions to administer the system. You don't know who does. You attempt to contact the system owner that opened the case

and there is no response. You have no contacts for the system and your colleagues are not familiar with it.

Benefits of establishing this best practice: By defining escalations, triggers for escalation, and procedures for escalation you enable the systematic addition of resources to an incident at an appropriate rate for the impact.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Define escalation paths: Define escalation paths in your runbooks and playbooks, including what triggers escalation, and procedures for escalation. For example, escalation of an issue from support engineers to senior support engineers when runbooks cannot resolve the issue, or when a predefined period of time has elapsed. Another example of an appropriate escalation path is from senior support engineers to the development team for a workload when the playbooks are unable to identify a path to remediation, or when a predefined period of time has elapsed. Specifically identify owners for each action to ensure effective and prompt responses to operations events. Escalations can include third parties. For example, a network connectivity provider or a software vendor. Escalations can include identified authorized decision makers for impacted systems.

OPS10-BP05 Enable push notifications

Communicate directly with your users (for example, with email or SMS) when the services they use are impacted, and again when the services return to normal operating conditions, to enable users to take appropriate action.

Common anti-patterns:

- Your application is experiencing a distributed denial of service incident and has been unresponsive for days. There is no error message. You have not sent a notification email. You have not sent text notifications. You have not shared information on social media. You customers are frustrated and looking for other vendors who can support them.
- On Monday, your application had issues following a patch and was down for a couple of hours. On Tuesday, your application had issues following a code deployment and was unreliable for a couple of hours. On Wednesday, your application had issues following a code deployment to mitigate a security vulnerability associated to the failed patch and was unavailable for a couple of hours. On Thursday, your frustrated customers started looking for another vendor who could support them.
- Your application is going to be down for maintenance this weekend. You don't inform your customers. Some of your customers had scheduled activities involving the use of your application. They are very frustrated upon discovery that your application is not available.

Benefits of establishing this best practice: By defining notifications, triggers for notifications, and procedures for notifications you enable your customer to be informed and respond when issues with your workload impact them.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Enable push notifications: Communicate directly with your users (for example, with email or SMS) when the services they use are impacted, and when the services return to normal operating conditions, to enable users to take appropriate action.
 - [Amazon SES features](#)
 - [What is Amazon SES?](#)
 - [Set up Amazon SNS notifications](#)

Resources

Related documents:

- [Amazon SES features](#)
- [Set up Amazon SNS notifications](#)
- [What is Amazon SES?](#)

OPS10-BP06 Communicate status through dashboards

Provide dashboards tailored to their target audiences (for example, internal technical teams, leadership, and customers) to communicate the current operating status of the business and provide metrics of interest.

You can create dashboards using [Amazon CloudWatch Dashboards](#) on customizable home pages in the CloudWatch console. Using business intelligence services such as [Amazon QuickSight](#) you can create and publish interactive dashboards of your workload and operational health (for example, order rates, connected users, and transaction times). Create Dashboards that present system and business-level views of your metrics.

Common anti-patterns:

- Upon request, you run a report on the current utilization of your application for management.
- During an incident, you are contacted every twenty minutes by a concerned system owner wanting to know if it is fixed yet.

Benefits of establishing this best practice: By creating dashboards, you enable self-service access to information enabling your customers to inform themselves and determine if they need to take action.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Communicate status through dashboards: Provide dashboards tailored to their target audiences (for example, internal technical teams, leadership, and customers) to communicate the current operating status of the business and provide metrics of interest. Providing a self-service option for status information reduces the disruption of fielding requests for status by the operations team. Examples include Amazon CloudWatch dashboards, and AWS Health Dashboard.
 - [CloudWatch dashboards create and use customized metrics views](#)

Resources

Related documents:

- [Amazon QuickSight](#)
- [CloudWatch dashboards create and use customized metrics views](#)

OPS10-BP07 Automate responses to events

Automate responses to events to reduce errors caused by manual processes, and to ensure prompt and consistent responses.

There are multiple ways to automate runbook and playbook actions on AWS. To respond to an event from a state change in your AWS resources, or from your own custom events, you should create

[CloudWatch Events rules](#) to trigger responses through CloudWatch targets (for example, Lambda functions, Amazon Simple Notification Service (Amazon SNS) topics, Amazon ECS tasks, and AWS Systems Manager Automation).

To respond to a metric that crosses a threshold for a resource (for example, wait time), you should create [CloudWatch alarms](#) to perform one or more actions using Amazon EC2 actions, Auto Scaling actions, or to send a notification to an Amazon SNS topic. If you need to perform custom actions in response to an alarm, invoke Lambda through an Amazon SNS notification. Use Amazon SNS to publish event notifications and escalation messages to keep people informed.

AWS also supports third-party systems through the AWS service APIs and SDKs. There are a number of monitoring tools provided by AWS Partners and third parties that allow for monitoring, notifications, and responses. Some of these tools include New Relic, Splunk, Loggly, SumoLogic, and Datadog.

You should keep critical manual procedures available for use when automated procedures fail

Common anti-patterns:

- A developer checks in their code. This event could have been used to start a build and then perform testing but instead nothing happens.
- Your application logs a specific error before it stops working. The procedure to restart the application is well understood and could be scripted. You could use the log event to invoke a script and restart the application. Instead, when the error happens at 3am Sunday morning, you are woken up as the on-call resource responsible to fix the system.

Benefits of establishing this best practice: By using automated responses to events, you reduce the time to respond and limit the introduction of errors from manual activities.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Automate responses to events: Automate responses to events to reduce errors caused by manual processes, and to ensure prompt and consistent responses.
 - [What is Amazon CloudWatch Events?](#)
 - [Creating a CloudWatch Events rule that triggers on an event](#)
 - [Creating a CloudWatch Events rule that triggers on an AWS API call using AWS CloudTrail](#)
 - [CloudWatch Events event examples from supported services](#)

Resources

Related documents:

- [Amazon CloudWatch Features](#)
- [CloudWatch Events event examples from supported services](#)
- [Creating a CloudWatch Events rule that triggers on an AWS API call using AWS CloudTrail](#)
- [Creating a CloudWatch Events rule that triggers on an event](#)
- [What is Amazon CloudWatch Events?](#)

Related videos:

- [Build a Monitoring Plan](#)

Related examples:

Evolve

Question

- [OPS 11 How do you evolve operations? \(p. 110\)](#)

OPS 11 How do you evolve operations?

Dedicate time and resources for continuous incremental improvement to evolve the effectiveness and efficiency of your operations.

Best practices

- [OPS11-BP01 Have a process for continuous improvement \(p. 110\)](#)
- [OPS11-BP02 Perform post-incident analysis \(p. 110\)](#)
- [OPS11-BP03 Implement feedback loops \(p. 111\)](#)
- [OPS11-BP04 Perform knowledge management \(p. 112\)](#)
- [OPS11-BP05 Define drivers for improvement \(p. 112\)](#)
- [OPS11-BP06 Validate insights \(p. 114\)](#)
- [OPS11-BP07 Perform operations metrics reviews \(p. 114\)](#)
- [OPS11-BP08 Document and share lessons learned \(p. 115\)](#)
- [OPS11-BP09 Allocate time to make improvements \(p. 116\)](#)

OPS11-BP01 Have a process for continuous improvement

Regularly evaluate and prioritize opportunities for improvement to focus efforts where they can provide the greatest benefits.

Common anti-patterns:

- You have documented the procedures necessary to create a development or testing environment. You could use CloudFormation to automate the process, but instead you do it manually from the console.
- Your testing shows that the vast majority of CPU utilization inside your application is in a small set of inefficient functions. You could focus on improving them and reduce your costs but you have been tasked to create a new usability feature.

Benefits of establishing this best practice: Continual improvement provides a mechanism to regularly evaluate opportunities for improvement, prioritize opportunities, and focus efforts where they can provide the greatest benefits.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Define processes for continuous improvement: Regularly evaluate and prioritize opportunities for improvement to focus efforts where they provide the greatest benefits. Implement changes to improve and evaluate the outcomes to determine success. If the outcomes do not satisfy the goals, and the improvement is still a priority, iterate using alternative courses of action. Your operations processes should include dedicated time and resources to make continuous incremental improvements possible.

OPS11-BP02 Perform post-incident analysis

Review customer-impacting events, and identify the contributing factors and preventative actions. Use this information to develop mitigations to limit or prevent recurrence. Develop procedures for

prompt and effective responses. Communicate contributing factors and corrective actions as appropriate, tailored to target audiences.

Common anti-patterns:

- You administer an application server. Approximately every 23 hours and 55 minutes all your active sessions are terminated. You have tried to identify what is going wrong on your application server. You suspect it could instead be a network issue but are unable to get cooperation from the network team as they are too busy to support you. You lack a predefined process to follow to get support and collect the information necessary to determine what is going on.
- You have had data loss within your workload. This is the first time it has happened and the cause is not obvious. You decide it is not important because you can recreate the data. Data loss starts occurring with greater frequency impacting your customers. This also places additional operational burden on you as you restore the missing data.

Benefits of establishing this best practice: Having a predefined processes to determine the components, conditions, actions, and events that contributed to an incident enables you to identify opportunities for improvement.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use a process to determine contributing factors: Review all customer impacting incidents. Have a process to identify and document the contributing factors of an incident so that you can develop mitigations to limit or prevent recurrence and you can develop procedures for prompt and effective responses. Communicate root cause as appropriate, tailored to target audiences.

OPS11-BP03 Implement feedback loops

Include feedback loops in your procedures and workloads to help you identify issues and areas that need improvement.

Common anti-patterns:

- A single frustrated customer opens a support case for a new product feature request to address a perceived issue. It is added to the list of priority improvements.
- You do not accept customer feedback.
- There are a significant number of requests for a change to an existing feature. Its originator maintains that it works as intended and does not need to change, so no changes are made.
- When a specific team member performs a deploy it always fails. No action is taken.

Benefits of establishing this best practice: By using the data from feedback loops to identify where to target improvement, you can minimize the impact of event based motivations or emotional investment.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Feedback loops: Have procedures embedded in your operations activities to capture feedback from their execution and to identify areas for improvement.
- Immediate feedback: Immediate feedback comes from the execution of operations activities where, through review of the execution and outcomes, it is recognized that the process could be improved. Feedback can come from customers, team members, or automated output of an activity. When the improvement has a low level of effort, or significant benefit, consider implementing it immediately.

Track opportunities for improvement in your backlog or issue system as appropriate. For example, a process where data is staged on an intermediate device could be optimized by instead placing the data directly into the target environment. This would eliminate a step in the process and the requirement for the intermediate resources.

- Retrospective analysis: Perform retrospective analysis regularly to capture feedback from the review of operational outcomes and metrics over time. Use trends to identify areas that need improvement. For example, review the rate of deployment failures to identify when potential issues with development and deployment activities have emerged.
- [Serverless big data analytics - Amazon Athena and Amazon QuickSight - 2017 AWS Online Tech Talks](#)
- [View AWS CodeDeploy logs in Amazon CloudWatch console](#)
- [Analyzing VPC flow logs with Amazon Kinesis Firehose, Amazon Athena, and Amazon QuickSight](#)

Resources

Related documents:

- [Analyzing VPC flow logs with Amazon Kinesis Firehose, Amazon Athena, and Amazon QuickSight](#)
- [View AWS CodeDeploy logs in Amazon CloudWatch console](#)

Related videos:

- [Serverless big data analytics - Amazon Athena and Amazon QuickSight - 2017 AWS Online Tech Talks](#)

OPS11-BP04 Perform knowledge management

Mechanisms exist for your team members to discover the information that they are looking for in a timely manner, access it, and identify that it's current and complete. Mechanisms are present to identify needed content, content in need of refresh, and content that should be archived so that it's no longer referenced.

Common anti-patterns:

- A single frustrated customer opens a support case for a new product feature request to address a perceived issue. It is added to the list of priority improvements.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Knowledge management: Ensure mechanisms exist for your team members to discover the information that they are looking for in a timely manner, access it, and identify that it's current and complete. Maintain mechanisms to identify needed content, content in need of refresh, and content that should be archived so that it's no longer referenced.

OPS11-BP05 Define drivers for improvement

Identify drivers for improvement to help you evaluate and prioritize opportunities.

On AWS, you can aggregate the logs of all your operations activities, workloads, and infrastructure to create a detailed activity history. You can then use AWS tools to analyze your operations and workload health over time (for example, identify trends, correlate events and activities to outcomes, and compare and contrast between environments and across systems) to reveal opportunities for improvement based on your drivers.

You should use CloudTrail to track API activity (through the AWS Management Console, CLI, SDKs, and APIs) to know what is happening across your accounts. Track your AWS developer Tools deployment activities with CloudTrail and CloudWatch. This will add a detailed activity history of your deployments and their outcomes to your CloudWatch Logs log data.

[Export your log data to Amazon S3](#) for long-term storage. Using [AWS Glue](#), you discover and prepare your log data in Amazon S3 for analytics. Use [Amazon Athena](#), through its native integration with AWS Glue, to analyze your log data. Use a business intelligence tool like [Amazon QuickSight](#) to visualize, explore, and analyze your data

Common anti-patterns:

- You have a script that works but is not elegant. You invest time in rewriting it. It is now a work of art.
- Your start-up is trying to get another set of funding from a venture capitalist. They want you to demonstrate compliance with PCI DSS. You want to make them happy so you document your compliance and miss a delivery date for a customer, losing that customer. It wasn't a wrong thing to do but now you wonder if it was the right thing to do.

Benefits of establishing this best practice: By determining the criteria you want to use for improvement, you can minimize the impact of event based motivations or emotional investment.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Understand drivers for improvement: You should only make changes to a system when a desired outcome is supported.
 - Desired capabilities: Evaluate desired features and capabilities when evaluating opportunities for improvement.
 - [What's New with AWS](#)
 - Unacceptable issues: Evaluate unacceptable issues, bugs, and vulnerabilities when evaluating opportunities for improvement.
 - [AWS Latest Security Bulletins](#)
 - [AWS Trusted Advisor](#)
 - Compliance requirements: Evaluate updates and changes required to maintain compliance with regulation, policy, or to remain under support from a third party, when reviewing opportunities for improvement.
 - [AWS Compliance](#)
 - [AWS Compliance Programs](#)
 - [AWS Compliance Latest News](#)

Resources

Related documents:

- [Amazon Athena](#)
- [Amazon QuickSight](#)
- [AWS Compliance](#)
- [AWS Compliance Latest News](#)
- [AWS Compliance Programs](#)
- [AWS Glue](#)
- [AWS Latest Security Bulletins](#)
- [AWS Trusted Advisor](#)

- [Export your log data to Amazon S3](#)
- [What's New with AWS](#)

OPS11-BP06 Validate insights

Review your analysis results and responses with cross-functional teams and business owners. Use these reviews to establish common understanding, identify additional impacts, and determine courses of action. Adjust responses as appropriate.

Common anti-patterns:

- You see that CPU utilization is at 95% on a system and make it a priority to find a way to reduce load on the system. You determine the best course of action is to scale up. The system is a transcoder and the system is scaled to run at 95% CPU utilization all the time. The system owner could have explained the situation to you had you contacted them. Your time has been wasted.
- A system owner maintains that their system is mission critical. The system was not placed in a high security environment. To improve security, you implement the additional detective and preventative controls that are required for mission critical systems. You notify the system owner that the work is complete and that he will be charged for the additional resources. In the discussion following this notification, the system owner learns there is a formal definition for mission critical systems that this system does not meet.

Benefits of establishing this best practice: By validating insights with business owners and subject matter experts, you can establish common understanding and more effectively guide improvement.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Validate insights: Engage with business owners and subject matter experts to ensure there is common understanding and agreement of the meaning of the data you have collected. Identify additional concerns, potential impacts, and determine a courses of action.

OPS11-BP07 Perform operations metrics reviews

Regularly perform retrospective analysis of operations metrics with cross-team participants from different areas of the business. Use these reviews to identify opportunities for improvement, potential courses of action, and to share lessons learned.

Look for opportunities to improve in all of your environments (for example, development, test, and production).

Common anti-patterns:

- There was a significant retail promotion that was interrupted by your maintenance window. The business remains unaware that there is a standard maintenance window that could be delayed if there are other business impacting events.
- You suffered an extended outage because of your use of a buggy library commonly used in your organization. You have since migrated to a reliable library. The other teams in your organization do not know that they are at risk. If you met regularly and reviewed this incident, they would be aware of the risk.
- Performance of your transcoder has been falling off steadily and impacting the media team. It isn't terrible yet. You will not have an opportunity to find out until it is bad enough to cause an incident. Were you to review your operations metrics with the media team, there would be an opportunity for the change in metrics and their experience to be recognized and the issue addressed.

- You are not reviewing your satisfaction of customer SLAs. You are trending to not meet your customer SLAs. There are financial penalties related to not meeting your customer SLAs. If you meet regularly to review the metrics for these SLAs, you would have the opportunity to recognize and address the issue.

Benefits of establishing this best practice: By meeting regularly to review operations metrics, events, and incidents, you maintain common understanding across teams, share lessons learned, and can prioritize and target improvements.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Operations metrics reviews: Regularly perform retrospective analysis of operations metrics with cross-team participants from different areas of the business. Engage stakeholders, including the business, development, and operations teams, to validate your findings from immediate feedback and retrospective analysis, and to share lessons learned. Use their insights to identify opportunities for improvement and potential courses of action.
 - [Amazon CloudWatch](#)
 - [Using Amazon CloudWatch metrics](#)
 - [Publish custom metrics](#)
 - [Amazon CloudWatch metrics and dimensions reference](#)

Resources

Related documents:

- [Amazon CloudWatch](#)
- [Amazon CloudWatch metrics and dimensions reference](#)
- [Publish custom metrics](#)
- [Using Amazon CloudWatch metrics](#)

OPS11-BP08 Document and share lessons learned

Document and share lessons learned from the operations activities so that you can use them internally and across teams.

You should share what your teams learn to increase the benefit across your organization. You will want to share information and resources to prevent avoidable errors and ease development efforts. This will allow you to focus on delivering desired features.

Use AWS Identity and Access Management (IAM) to define permissions enabling controlled access to the resources you wish to share within and across accounts. You should then use version-controlled AWS CodeCommit repositories to share application libraries, scripted procedures, procedure documentation, and other system documentation. Share your compute standards by sharing access to your AMIs and by authorizing the use of your Lambda functions across accounts. You should also share your infrastructure standards as AWS CloudFormation templates.

Through the AWS APIs and SDKs, you can integrate external and third-party tools and repositories (for example, GitHub, BitBucket, and SourceForge). When sharing what you have learned and developed, be careful to structure permissions to ensure the integrity of shared repositories.

Common anti-patterns:

- You suffered an extended outage because of your use of a buggy library commonly used in your organization. You have since migrated to a reliable library. The other teams in your organization do not

know they are at risk. Were you to document and share your experience with this library, they would be aware of the risk.

- You have identified an edge case in an internally shared microservice that causes sessions to drop. You have updated your calls to the service to avoid this edge case. The other teams in your organization do not know that they are at risk. Were you to document and share your experience with this library, they would be aware of the risk.
- You have found a way to significantly reduce the CPU utilization requirements for one of your microservices. You do not know if any other teams could take advantage of this technique. Were you to document and share your experience with this library, they would have the opportunity to do so.

Benefits of establishing this best practice: Share lessons learned to support improvement and to maximize the benefits of experience.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Document and share lessons learned: Have procedures to document the lessons learned from the execution of operations activities and retrospective analysis so that they can be used by other teams.
- Share learnings: Have procedures to share lessons learned and associated artifacts across teams. For example, share updated procedures, guidance, governance, and best practices through an accessible wiki. Share scripts, code, and libraries through a common repository.
 - [Delegating access to your AWS environment](#)
 - [Share an AWS CodeCommit repository](#)
 - [Easy authorization of AWS Lambda functions](#)
 - [Sharing an AMI with specific AWS Accounts](#)
 - [Speed template sharing with an AWS CloudFormation designer URL](#)
 - [Using AWS Lambda with Amazon SNS](#)

Resources

Related documents:

- [Easy authorization of AWS Lambda functions](#)
- [Share an AWS CodeCommit repository](#)
- [Sharing an AMI with specific AWS Accounts](#)
- [Speed template sharing with an AWS CloudFormation designer URL](#)
- [Using AWS Lambda with Amazon SNS](#)

Related videos:

- [Delegating access to your AWS environment](#)

OPS11-BP09 Allocate time to make improvements

Dedicate time and resources within your processes to make continuous incremental improvements possible.

On AWS, you can create temporary duplicates of environments, lowering the risk, effort, and cost of experimentation and testing. These duplicated environments can be used to test the conclusions from your analysis, experiment, and develop and test planned improvements.

Common anti-patterns:

- There is a known performance issue in your application server. It is added to the backlog behind every planned feature implementation. If the rate of planned features being added remains constant, the performance issue will never be addressed.
- To support continual improvement you approve administrators and developers using all their extra time to select and implement improvements. No improvements are ever completed.

Benefits of establishing this best practice: By dedicating time and resources within your processes you make continuous incremental improvements possible.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Allocate time to make improvements: Dedicate time and resources within your processes to make continuous incremental improvements possible. Implement changes to improve and evaluate the results to determine success. If the results do not satisfy the goals, and the improvement is still a priority, pursue alternative courses of action.

Security

Topics

- [Security foundations \(p. 117\)](#)
- [Identity and access management \(p. 123\)](#)
- [Detection \(p. 137\)](#)
- [Infrastructure protection \(p. 142\)](#)
- [Data protection \(p. 152\)](#)
- [Incident response \(p. 162\)](#)

Security foundations

Question

- [SEC 1 How do you securely operate your workload? \(p. 117\)](#)

SEC 1 How do you securely operate your workload?

To operate your workload securely, you must apply overarching best practices to every area of security. Take requirements and processes that you have defined in operational excellence at an organizational and workload level, and apply them to all areas. Staying up to date with AWS and industry recommendations and threat intelligence helps you evolve your threat model and control objectives. Automating security processes, testing, and validation allow you to scale your security operations.

Best practices

- [SEC01-BP01 Separate workloads using accounts \(p. 118\)](#)
- [SEC01-BP02 Secure AWS account \(p. 118\)](#)
- [SEC01-BP03 Identify and validate control objectives \(p. 119\)](#)
- [SEC01-BP04 Keep up-to-date with security threats \(p. 120\)](#)
- [SEC01-BP05 Keep up-to-date with security recommendations \(p. 120\)](#)
- [SEC01-BP06 Automate testing and validation of security controls in pipelines \(p. 121\)](#)

- [SEC01-BP07 Identify and prioritize risks using a threat model \(p. 122\)](#)
- [SEC01-BP08 Evaluate and implement new security services and features regularly \(p. 123\)](#)

SEC01-BP01 Separate workloads using accounts

Start with security and infrastructure in mind to enable your organization to set common guardrails as your workloads grow. This approach provides boundaries and controls between workloads. Account-level separation is strongly recommended for isolating production environments from development and test environments, or providing a strong logical boundary between workloads that process data of different sensitivity levels, as defined by external compliance requirements (such as PCI-DSS or HIPAA), and workloads that don't.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use AWS Organizations: Use AWS Organizations to centrally enforce policy-based management for multiple AWS accounts.
 - [Getting started with AWS Organizations](#)
 - [How to use service control policies to set permission guardrails across accounts in your AWS Organization](#)
- Consider AWS Control Tower: AWS Control Tower provides an easy way to set up and govern a new, secure, multi-account AWS environment based on best practices.
 - [AWS Control Tower](#)

Resources

Related documents:

- [IAM Best Practices](#)
- [Security Bulletins](#)
- [AWS Security Audit Guidelines](#)

Related videos:

- [Managing Multi-Account AWS Environments Using AWS Organizations](#)
- [Security Best Practices the Well-Architected Way](#)
- [Using AWS Control Tower to Govern Multi-Account AWS Environments](#)

SEC01-BP02 Secure AWS account

There are a number of aspects to securing your AWS accounts, including the securing of, and not using the [root user](#), and keeping your contact information up-to-date. You can use [AWS Organizations](#) to centrally manage and govern your accounts as you grow and scale your workloads in AWS. AWS Organizations helps you manage accounts, set controls, and configure services across your accounts.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use AWS Organizations: Use AWS Organizations to centrally enforce policy-based management for multiple AWS accounts.
 - [Getting started with AWS Organizations](#)

- [How to use service control policies to set permission guardrails across accounts in your AWS Organization](#)
- Limit use of the AWS root user: Only use the root user to perform tasks that specifically require it.
 - [AWS Tasks That Require AWS Account Root User Credentials](#)
- Enable multi-factor-authentication (MFA) for the root user: Enable MFA on the AWS account root user, if AWS Organizations is not managing root users for you.
 - [Root user](#)
- Periodically change the root user password: Changing the root user password reduces the risk that a saved password can be used. This is especially important if you are not using AWS Organizations and anyone has physical access.
 - [Changing the AWS account root user password](#)
- Enable notification when the AWS account root user is used: Being notified automatically reduces risk.
 - [How to receive notifications when your AWS account's root access keys are used](#)
- Restrict access to newly added Regions: For new AWS Regions, IAM resources, such as users and roles, will only be propagated to the Regions that you enable.
 - [Setting permissions to enable accounts for upcoming AWS Regions](#)
- Consider AWS CloudFormation StackSets: CloudFormation StackSets can be used to deploy resources including IAM policies, roles, and groups into different AWS accounts and Regions from an approved template.
 - [Use CloudFormation StackSets](#)

Resources

Related documents:

- [AWS Control Tower](#)
- [AWS Security Audit Guidelines](#)
- [IAM Best Practices](#)
- [Security Bulletins](#)

Related videos:

- [Enable AWS adoption at scale with automation and governance](#)
- [Security Best Practices the Well-Architected Way](#)

Related examples:

- [Lab: AWS account and root user](#)

SEC01-BP03 Identify and validate control objectives

Based on your compliance requirements and risks identified from your threat model, derive and validate the control objectives and controls that you need to apply to your workload. Ongoing validation of control objectives and controls help you measure the effectiveness of risk mitigation.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Identify compliance requirements: Discover the organizational, legal, and compliance requirements that your workload must comply with.

- Identify AWS compliance resources: Identify resources that AWS has available to assist you with compliance.
 - <https://aws.amazon.com/compliance/>
 - <https://aws.amazon.com/artifact/>

Resources

Related documents:

- [AWS Security Audit Guidelines](#)
- [Security Bulletins](#)

Related videos:

- [AWS Security Hub: Manage Security Alerts and Automate Compliance](#)
- [Security Best Practices the Well-Architected Way](#)

SEC01-BP04 Keep up-to-date with security threats

To help you define and implement appropriate controls, recognize attack vectors by staying up to date with the latest security threats. Consume AWS Managed Services to make it easier to receive notification of unexpected or unusual behavior in your AWS accounts. Investigate using AWS Partner tools or third-party threat information feeds as part of your security information flow. The [Common Vulnerabilities and Exposures \(CVE\) List](#) list contains publicly disclosed cyber security vulnerabilities that you can use to stay up to date.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Subscribe to threat intelligence sources: Regularly review threat intelligence information from multiple sources that are relevant to the technologies used in your workload.
 - [Common Vulnerabilities and Exposures List](#)
- Consider [AWS Shield Advanced](#) service: It provides near real-time visibility into intelligence sources, if your workload is internet accessible.

Resources

Related documents:

- [AWS Security Audit Guidelines](#)
- [AWS Shield](#)
- [Security Bulletins](#)

Related videos:

- [Security Best Practices the Well-Architected Way](#)

SEC01-BP05 Keep up-to-date with security recommendations

Stay up-to-date with both AWS and industry security recommendations to evolve the security posture of your workload. [AWS Security Bulletins](#) contain important information about security and privacy notifications.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Follow AWS updates: Subscribe or regularly check for new recommendations, tips and tricks.
 - [AWS Well-Architected Labs](#)
 - [AWS security blog](#)
 - [AWS service documentation](#)
- Subscribe to industry news: Regularly review news feeds from multiple sources that are relevant to the technologies that are used in your workload.
 - [Example: Common Vulnerabilities and Exposures List](#)

Resources

Related documents:

- [Security Bulletins](#)

Related videos:

- [Security Best Practices the Well-Architected Way](#)

SEC01-BP06 Automate testing and validation of security controls in pipelines

Establish secure baselines and templates for security mechanisms that are tested and validated as part of your build, pipelines, and processes. Use tools and automation to test and validate all security controls continuously. For example, scan items such as machine images and infrastructure-as-code templates for security vulnerabilities, irregularities, and drift from an established baseline at each stage. AWS CloudFormation Guard can help you verify that CloudFormation templates are safe, save you time, and reduce the risk of configuration error.

Reducing the number of security misconfigurations introduced into a production environment is critical—the more quality control and reduction of defects you can perform in the build process, the better. Design continuous integration and continuous deployment (CI/CD) pipelines to test for security issues whenever possible. CI/CD pipelines offer the opportunity to enhance security at each stage of build and delivery. CI/CD security tooling must also be kept updated to mitigate evolving threats.

Track changes to your workload configuration to help with compliance auditing, change management, and investigations that may apply to you. You can use AWS Config to record and evaluate your AWS and third-party resources. It allows you to continuously audit and assess the overall compliance with rules and conformance packs, which are collections of rules with remediation actions.

Change tracking should include planned changes, which are part of your organization's change control process (sometimes referred to as MACD—Move, Add, Change, Delete), unplanned changes, and unexpected changes, such as incidents. Changes might occur on the infrastructure, but they might also be related to other categories, such as changes in code repositories, machine images and application inventory changes, process and policy changes, or documentation changes.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Automate configuration management: Enforce and validate secure configurations automatically by using a configuration management service or tool.
 - [AWS Systems Manager](#)

- [AWS CloudFormation](#)
- [Set Up a CI/CD Pipeline on AWS](#)

Resources

Related documents:

- [How to use service control policies to set permission guardrails across accounts in your AWS Organization](#)

Related videos:

- [Managing Multi-Account AWS Environments Using AWS Organizations](#)
- [Security Best Practices the Well-Architected Way](#)

SEC01-BP07 Identify and prioritize risks using a threat model

Use a threat model to identify and maintain an up-to-date register of potential threats. Prioritize your threats and adapt your security controls to prevent, detect, and respond. Revisit and maintain this in the context of the evolving security landscape.

Threat modeling provides a systematic approach to aid in finding and addressing security issues early in the design process. Earlier is better since mitigations have a lower cost compared to later in the lifecycle.

The typical core steps of the threat modeling process are:

1. Identify assets, actors, entry points, components, use cases, and trust levels, and include these in a design diagram.
2. Identify a list of threats.
3. For each threat, identify mitigations, which might include security control implementations.
4. Create and review a risk matrix to determine if the threat is adequately mitigated.

Threat modeling is most effective when done at the workload (or workload feature) level, ensuring that all context is available for assessment. Revisit and maintain this matrix as your security landscape evolves.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Create a threat model: A threat model can help you identify and address potential security threats.
 - [NIST: Guide to Data-Centric System Threat Modeling](#)

Resources

Related documents:

- [AWS Security Audit Guidelines](#)
- [Security Bulletins](#)

Related videos:

- [Security Best Practices the Well-Architected Way](#)

SEC01-BP08 Evaluate and implement new security services and features regularly

Evaluate and implement security services and features from AWS and AWS Partners that allow you to evolve the security posture of your workload. The AWS Security Blog highlights new AWS services and features, implementation guides, and general security guidance. [What's New with AWS?](#) is a great way to stay up to date with all new AWS features, services, and announcements.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Plan regular reviews: Create a calendar of review activities that includes compliance requirements, evaluation of new AWS security features and services, and staying up-to-date with industry news.
- Discover AWS services and features: Discover the security features that are available for the services that you are using, and review new features as they are released.
 - [AWS security blog](#)
 - [AWS security bulletins](#)
 - [AWS service documentation](#)
- Define AWS service on-boarding process: Define processes for onboarding of new AWS services. Include how you evaluate new AWS services for functionality, and the compliance requirements for your workload.
- Test new services and features: Test new services and features as they are released in a non-production environment that closely replicates your production one.
- Implement other defense mechanisms: Implement automated mechanisms to defend your workload, explore the options available.
 - [Remediating non-compliant AWS resources by AWS Config Rules](#)

Resources

Related videos:

- [Security Best Practices the Well-Architected Way](#)

Identity and access management

Questions

- [SEC 2 How do you manage authentication for people and machines? \(p. 123\)](#)
- [SEC 3 How do you manage permissions for people and machines? \(p. 130\)](#)

SEC 2 How do you manage authentication for people and machines?

There are two types of identities you need to manage when approaching operating secure AWS workloads. Understanding the type of identity you need to manage and grant access helps you ensure the right identities have access to the right resources under the right conditions.

Human Identities: Your administrators, developers, operators, and end users require an identity to access your AWS environments and applications. These are members of your organization, or external users with whom you collaborate, and who interact with your AWS resources via a web browser, client application, or interactive command-line tools.

Machine Identities: Your service applications, operational tools, and workloads require an identity to make requests to AWS services for example, to read data. These identities include machines running in your AWS environment such as Amazon EC2 instances or AWS Lambda functions. You may also manage machine identities for external parties who need access. Additionally, you may also have machines outside of AWS that need access to your AWS environment.

Best practices

- [SEC02-BP01 Use strong sign-in mechanisms \(p. 124\)](#)
- [SEC02-BP02 Use temporary credentials \(p. 125\)](#)
- [SEC02-BP03 Store and use secrets securely \(p. 126\)](#)
- [SEC02-BP04 Rely on a centralized identity provider \(p. 127\)](#)
- [SEC02-BP05 Audit and rotate credentials periodically \(p. 128\)](#)
- [SEC02-BP06 Leverage user groups and attributes \(p. 129\)](#)

SEC02-BP01 Use strong sign-in mechanisms

Enforce minimum password length, and educate your users to avoid common or reused passwords. Enforce multi-factor authentication (MFA) with software or hardware mechanisms to provide an additional layer of verification. For example, when using IAM Identity Center as the identity source, configure the “context-aware” or “always-on” setting for MFA, and allow users to enroll their own MFA devices to accelerate adoption. When using an external identity provider (IdP), configure your IdP for MFA.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Create an Identify and Access Management (IAM) policy to enforce MFA sign-in: Create a customer-managed IAM policy that prohibits all IAM actions except for the ones that allow a user to assume roles, change their own credentials, and manage their MFA devices on the [My Security Credentials](#) page.
- Enable MFA in your identity provider: Enable [MFA](#) in the identity provider or single sign-on service, such as [AWS IAM Identity Center \(successor to AWS Single Sign-On\)](#), that you use.
- Configure a strong password policy: Configure a strong [password policy](#) in IAM and federated identity systems to help protect against brute-force attacks.
- [Rotate credentials regularly](#): Ensure administrators of your workload change their passwords and access keys (if used) regularly.

Resources

Related documents:

- [Getting Started with AWS Secrets Manager](#)
- [IAM Best Practices](#)
- [Identity Providers and Federation](#)
- [The AWS Account Root User](#)
- [Temporary Security Credentials](#)
- [Security Partner Solutions: Access and Access Control](#)

Related videos:

- [Best Practices for Managing, Retrieving, and Rotating Secrets at Scale](#)
- [Managing user permissions at scale with IAM Identity Center](#)
- [Mastering identity at every layer of the cake](#)

SEC02-BP02 Use temporary credentials

Require identities to dynamically acquire [temporary credentials](#). For workforce identities, use AWS IAM Identity Center (successor to AWS Single Sign-On), or federation with AWS Identity and Access Management (IAM) roles to access AWS accounts. For machine identities, such as Amazon Elastic Compute Cloud(Amazon EC2) instances or AWS Lambda functions, require the use of IAM roles instead of IAM users with long-term access keys.

For human identities using the AWS Management Console, require users to acquire temporary credentials and federate into AWS. You can do this using the AWS IAM Identity Center (successor to AWS Single Sign-On) user portal. For users requiring CLI access, ensure that they use [AWS CLI v2](#), which supports direct integration with IAM Identity Center. Users can create CLI profiles that are linked to IAM Identity Center accounts and roles. The CLI automatically retrieves AWS credentials from IAM Identity Center and refreshes them on your behalf. This eliminates the need to copy and paste temporary AWS credentials from the IAM Identity Center console. For SDK, users should rely on AWS Security Token Service (AWS STS) to assume roles to receive temporary credentials. In certain cases, temporary credentials might not be practical. You should be aware of the risks of storing access keys, rotate these often, and require multi-factor authentication (MFA) as a condition when possible. Use last accessed information to determine when to rotate or remove access keys.

For cases where you need to grant consumers access to your AWS resources, use [Amazon Cognito](#) identity pools and assign them a set of temporary, limited privilege credentials to access your AWS resources. The permissions for each user are controlled through [IAM roles](#) that you create. You can define rules to choose the role for each user based on claims in the user's ID token. You can define a default role for authenticated users. You can also define a separate IAM role with limited permissions for guest users who are not authenticated.

For machine identities, you should rely on IAM roles to grant access to AWS. For Amazon Elastic Compute Cloud(Amazon EC2) instances, you can use [roles for Amazon EC2](#). You can attach an IAM role to your Amazon EC2 instance to enable your applications running on Amazon EC2 to use temporary security credentials that AWS creates, distributes, and rotates automatically through the Instance Metadata Service (IMDS). The [latest version](#) of IMDS helps protect against vulnerabilities that expose the temporary credentials and should be implemented. For accessing Amazon EC2 instances using keys or passwords, [AWS Systems Manager](#) is a more secure way to access and manage your instances using a pre-installed agent without the stored secret. Additionally, other AWS services, such as AWS Lambda, enable you to configure an IAM service role to grant the service permissions to perform AWS actions using temporary credentials. In situations where you cannot use temporary credentials, use programmatic tools, such as [AWS Secrets Manager](#), to automate credential rotation and management.

Audit and rotate credentials periodically: Periodic validation, preferably through an automated tool, is necessary to verify that the correct controls are enforced. For human identities, you should require users to change their passwords periodically and retire access keys in favor of temporary credentials. As you are moving from IAM users to centralized identities, you can [generate a credential report](#) to audit your IAM users. We also recommend that you enforce MFA settings in your identity provider. You can set up [AWS Config Rules](#) to monitor these settings. For machine identities, you should rely on temporary credentials using IAM roles. For situations where this is not possible, frequent auditing and rotating access keys is necessary.

Store and use secrets securely: For credentials that are not IAM-related and cannot take advantage of temporary credentials, such as database logins, use a service that is designed to handle management of secrets, such as [Secrets Manager](#). Secrets Manager makes it easy to manage, rotate, and securely store encrypted secrets using [supported services](#). Calls to access the secrets are logged in AWS CloudTrail for auditing purposes, and IAM permissions can grant least-privilege access to them.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Implement least privilege policies: Assign access policies with least privilege to IAM groups and roles to reflect the user's role or function that you have defined.
 - [Grant least privilege](#)
- Remove unnecessary permissions: Implement least privilege by removing permissions that are unnecessary.
 - [Reducing policy scope by viewing user activity](#)
 - [View role access](#)
- Consider permissions boundaries: A permissions boundary is an advanced feature for using a managed policy that sets the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries.
 - [Lab: IAM permissions boundaries delegating role creation](#)
- Consider resource tags for permissions: You can use tags to control access to your AWS resources that support tagging. You can also tag IAM users and roles to control what they can access.
 - [Lab: IAM tag based access control for EC2](#)
 - [Attribute-based access control \(ABAC\)](#)

Resources

Related documents:

- [Getting Started with AWS Secrets Manager](#)
- [IAM Best Practices](#)
- [Identity Providers and Federation](#)
- [Security Partner Solutions: Access and Access Control](#)
- [Temporary Security Credentials](#)
- [The AWS Account Root User](#)

Related videos:

- [Best Practices for Managing, Retrieving, and Rotating Secrets at Scale](#)
- [Managing user permissions at scale with AWS IAM Identity Center \(successor to AWS Single Sign-On\)](#)
- [Mastering identity at every layer of the cake](#)

SEC02-BP03 Store and use secrets securely

For workforce and machine identities that require secrets such as passwords to third-party applications, store them with automatic rotation using the latest industry standards in a specialized service, such as for credentials that are not IAM-related and cannot take advantage of temporary credentials, such as database logins, use a service that is designed to handle management of secrets, such as AWS Secrets Manager. Secrets Manager makes it easy to manage, rotate, and securely store encrypted secrets using supported services. Calls to access the secrets are logged in AWS CloudTrail for auditing purposes, and IAM permissions can grant least-privilege access to them.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use AWS Secrets Manager: [AWS Secrets Manager](#) is an AWS service that makes it easier for you to manage secrets. Secrets can be database credentials, passwords, third-party API keys, and even arbitrary text.

Resources

Related documents:

- [Getting Started with AWS Secrets Manager](#)
- [Identity Providers and Federation](#)

Related videos:

- [Best Practices for Managing, Retrieving, and Rotating Secrets at Scale](#)

SEC02-BP04 Rely on a centralized identity provider

For workforce identities, rely on an identity provider that enables you to manage identities in a centralized place. This makes it easier to manage access across multiple applications and services, because you are creating, managing, and revoking access from a single location. For example, if someone leaves your organization, you can revoke access for all applications and services (including AWS) from one location. This reduces the need for multiple credentials and provides an opportunity to integrate with existing human resources (HR) processes.

For federation with individual AWS accounts, you can use centralized identities for AWS with a SAML 2.0-based provider with AWS Identity and Access Management. You can use any provider—whether hosted by you in AWS, external to AWS, or supplied by the AWS Partner—that is compatible with the [SAML 2.0](#) protocol. You can use federation between your AWS account and your chosen provider to grant a user or application access to call AWS API operations by using a SAML assertion to get temporary security credentials. Web-based single sign-on is also supported, allowing users to sign in to the AWS Management Console from your sign in website.

For federation to multiple accounts in your AWS Organizations, you can configure your identity source in [AWS IAM Identity Center \(successor to AWS Single Sign-On\) \(IAM Identity Center\)](#), and specify where your users and groups are stored. Once configured, your identity provider is your source of truth, and information can be [synchronized](#) using the System for Cross-domain Identity Management (SCIM) v2.0 protocol. You can then look up users or groups and grant them IAM Identity Center access to AWS accounts, cloud applications, or both.

IAM Identity Center integrates with AWS Organizations, which enables you to configure your identity provider once and then [grant access to existing and new accounts](#) managed in your organization. IAM Identity Center provides you with a default store, which you can use to manage your users and groups. If you choose to use the IAM Identity Center store, create your users and groups and assign their level of access to your AWS accounts and applications, keeping in mind the best practice of least privilege. Alternatively, you can choose to [Connect to Your External Identity Provider](#) using SAML 2.0, or [Connect to Your Microsoft AD Directory](#) using AWS Directory Service. Once configured, you can sign into the AWS Management Console, or the AWS mobile app, by authenticating through your central identity provider.

For managing end-users or consumers of your workloads, such as a mobile app, you can use [Amazon Cognito](#). It provides authentication, authorization, and user management for your web and mobile apps. Your users can sign in directly with a user name and password, or through a third party, such as Amazon, Apple, Facebook, or Google.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Centralize administrative access: Create an Identity and Access Management (IAM) identity provider entity to establish a trusted relationship between your AWS account and your identity provider (IdP). IAM supports IdPs that are compatible with OpenID Connect (OIDC) or SAML 2.0 (Security Assertion Markup Language 2.0).
 - [Identity Providers and Federation](#)
- Centralize application access: Consider Amazon Cognito for centralizing application access. It lets you add user sign-up, sign-in, and access control to your web and mobile apps quickly and easily. [Amazon Cognito](#) scales to millions of users and supports sign-in with social identity providers, such as Facebook, Google, and Amazon, and enterprise identity providers via SAML 2.0.
- Remove old IAM users and groups: After you start using an identity provider (IdP), remove IAM users and groups that are no longer required.
 - [Finding unused credentials](#)
 - [Deleting an IAM group](#)

Resources

Related documents:

- [IAM Best Practices](#)
- [Security Partner Solutions: Access and Access Control](#)
- [Temporary Security Credentials](#)
- [The AWS Account Root User](#)

Related videos:

- [Best Practices for Managing, Retrieving, and Rotating Secrets at Scale](#)
- [Managing user permissions at scale with AWS IAM Identity Center \(successor to AWS Single Sign-On\)](#)
- [Mastering identity at every layer of the cake](#)

SEC02-BP05 Audit and rotate credentials periodically

When you cannot rely on temporary credentials and require long-term credentials, audit credentials to ensure that the defined controls for example, multi-factor authentication (MFA), are enforced, rotated regularly, and have the appropriate access level. Periodic validation, preferably through an automated tool, is necessary to verify that the correct controls are enforced. For human identities, you should require users to change their passwords periodically and retire access keys in favor of temporary credentials. As you are moving from AWS Identity and Access Management (IAM) users to centralized identities, you can [generate a credential report](#) to audit your IAM users. We also recommend that you enforce MFA settings in your identity provider. You can set up [AWS Config Rules](#) to monitor these settings. For machine identities, you should rely on temporary credentials using IAM roles. For situations where this is not possible, frequent auditing and rotating access keys is necessary.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Regularly audit credentials: Use credential reports, and Identity and Access Management (IAM) Access Analyzer to audit IAM credentials and permissions.
 - [IAM Access Analyzer](#)
 - [Getting credential report](#)

- [Lab: Automated IAM user cleanup](#)
- **Use Access Levels to Review IAM Permissions:** To improve the security of your AWS account, regularly review and monitor each of your IAM policies. Make sure that your policies grant the least privilege that is needed to perform only the necessary actions.
 - [Use access levels to review IAM permissions](#)
- **Consider automating IAM resource creation and updates:** AWS CloudFormation can be used to automate the deployment of IAM resources, including roles and policies, to reduce human error because the templates can be verified and version controlled.
 - [Lab: Automated deployment of IAM groups and roles](#)

Resources

Related documents:

- [Getting Started with AWS Secrets Manager](#)
- [IAM Best Practices](#)
- [Identity Providers and Federation](#)
- [Security Partner Solutions: Access and Access Control](#)
- [Temporary Security Credentials](#)

Related videos:

- [Best Practices for Managing, Retrieving, and Rotating Secrets at Scale](#)
- [Managing user permissions at scale with AWS IAM Identity Center \(successor to AWS Single Sign-On\)](#)
- [Mastering identity at every layer of the cake](#)

SEC02-BP06 Leverage user groups and attributes

As the number of users you manage grows, you will need to determine ways to organize them so that you can manage them at scale. Place users with common security requirements in groups defined by your identity provider, and put mechanisms in place to ensure that user attributes that may be used for access control (for example, department or location) are correct and updated. Use these groups and attributes to control access, rather than individual users. This allows you to manage access centrally by changing a user's group membership or attributes once with a [permission set](#), rather than updating many individual policies when a user's access needs change.

You can use AWS IAM Identity Center (successor to AWS Single Sign-On) (IAM Identity Center) to manage user groups and attributes. IAM Identity Center supports most commonly used attributes whether they are entered manually during user creation or automatically provisioned using a synchronization engine, such as defined in the System for Cross-Domain Identity Management (SCIM) specification.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- If you are using AWS IAM Identity Center (successor to AWS Single Sign-On) (IAM Identity Center), configure groups: IAM Identity Center provides you with the ability to configure groups of users, and assign groups the desired level of permission.
 - [AWS Single Sign-On - Manage Identities](#)
- Learn about attribute-based access control (ABAC): ABAC is an authorization strategy that defines permissions based on attributes.

- [What Is ABAC for AWS?](#)
- [Lab: IAM Tag Based Access Control for EC2](#)

Resources

Related documents:

- [Getting Started with AWS Secrets Manager](#)
- [IAM Best Practices](#)
- [Identity Providers and Federation](#)
- [The AWS Account Root User](#)

Related videos:

- [Best Practices for Managing, Retrieving, and Rotating Secrets at Scale](#)
- [Managing user permissions at scale with AWS IAM Identity Center \(successor to AWS Single Sign-On\)](#)
- [Mastering identity at every layer of the cake](#)

Related examples:

- [Lab: IAM Tag Based Access Control for EC2](#)

SEC 3 How do you manage permissions for people and machines?

Manage permissions to control access to people and machine identities that require access to AWS and your workload. Permissions control who can access what, and under what conditions.

Best practices

- [SEC03-BP01 Define access requirements \(p. 130\)](#)
- [SEC03-BP02 Grant least privilege access \(p. 131\)](#)
- [SEC03-BP03 Establish emergency access process \(p. 132\)](#)
- [SEC03-BP04 Reduce permissions continuously \(p. 133\)](#)
- [SEC03-BP05 Define permission guardrails for your organization \(p. 134\)](#)
- [SEC03-BP06 Manage access based on lifecycle \(p. 135\)](#)
- [SEC03-BP07 Analyze public and cross-account access \(p. 135\)](#)
- [SEC03-BP08 Share resources securely \(p. 136\)](#)

SEC03-BP01 Define access requirements

As you grow and manage additional workloads in AWS, you should separate these workloads using accounts and manage those accounts using AWS Organizations. We recommend that you establish common permission guardrails that restrict access to all identities in your organization. For example, you can restrict access to specific AWS Regions, or prevent your team from deleting common resources, such as an IAM role used by your central security team. You can get started by implementing [example service control policies](#), such as preventing users from disabling key services.

You can use AWS Organizations to group accounts and set common controls on each group of accounts. To set these common controls, you can use services integrated with AWS Organizations. Specifically,

you can use [service control policies \(SCPs\)](#) to [restrict access to a group of accounts](#). SCPs use the IAM policy language and enable you to establish controls that all IAM principals (users and roles) adhere to. You can restrict access to specific service actions, resources and based on specific condition to meet the access control needs of your organization. If necessary, you can define exceptions to your guardrails. For example, you can restrict service actions for all IAM entities in the account except for a specific administrator role.

Each component or resource of your workload needs to be accessed by administrators, users, or other components. Have a clear definition of who or what should have access to each component, choose the appropriate identity type and method of authentication and authorization.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Define required permissions for job function and responsibilities: Based on the user's job function, role, or responsibilities, define what resources that they need access to and the conditions that may apply. Group the users with common requirements together to make delegation of policies easier.
 - [IAM use cases](#)

Resources

Related documents:

- [Attribute-based access control \(ABAC\)](#)
- [Grant least privilege](#)
- [IAM Access Analyzer](#)
- [IAM use cases](#)
- [Remove unnecessary credentials](#)
- [Working with Policies](#)

Related videos:

- [Become an IAM Policy Master in 60 Minutes or Less](#)
- [Separation of Duties, Least Privilege, Delegation, and CI/CD](#)

SEC03-BP02 Grant least privilege access

Establishing a principle of [least privilege](#) ensures that identities are only permitted to perform the most minimal set of functions necessary to fulfill a specific task, while balancing usability and efficiency. Operating on this principle limits unintended access and helps ensure that you can audit who has access to which resources. In AWS, identities have no permissions by default with the exception of the root user, which should only be used for a few [specific tasks](#).

You use policies to explicitly grant permissions attached to IAM or resource entities, such as an IAM role used by federated identities or machines, or resources (for example, Amazon S3 buckets). When you create and attach a policy, you can specify the service actions, resources, and conditions that must be true for AWS to allow access. AWS supports a variety of conditions to help you scope down access. For example, using the PrincipalOrgID [condition key](#), the identifier of the AWS Organizations is verified so access can be granted within your AWS Organizations. You can also control requests that AWS services make on your behalf, like AWS CloudFormation creating an AWS Lambda function, by using the CalledVia condition key. This enables you to set granular permissions for your human and machine identities across AWS.

AWS also has capabilities that enable you to scale your permissions management and adhere to least privilege.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Implement least privilege policies: Assign access policies with least privilege to IAM groups and roles to reflect the user's role or function that you have defined.
 - [Grant least privilege](#)
- Remove unnecessary permissions: Implement least privilege by removing permissions that are unnecessary.
 - [Reducing policy scope by viewing user activity](#)
 - [View role access](#)
- Consider permissions boundaries: A permissions boundary is an advanced feature for using a managed policy that sets the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries.
 - [Lab: IAM permissions boundaries delegating role creation](#)
- Consider resource tags for permissions: You can use tags to control access to your AWS resources that support tagging. You can also tag IAM users and roles to control what they can access.
 - [Lab: IAM tag based access control for EC2](#)
 - [Attribute-based access control \(ABAC\)](#)

Resources

Related documents:

- [IAM Access Analyzer](#)
- [Reducing policy scope by viewing user activity](#)
- [Remove unnecessary credentials](#)
- [View role access](#)
- [Working with Policies](#)

Related videos:

- [Become an IAM Policy Master in 60 Minutes or Less](#)
- [Separation of Duties, Least Privilege, Delegation, and CI/CD](#)

Related examples:

- [Lab: IAM permissions boundaries delegating role creation](#)
- [Lab: IAM tag based access control for EC2](#)

SEC03-BP03 Establish emergency access process

A process that allows emergency access to your workload in the unlikely event of an automated process or pipeline issue. This will help you rely on least privilege access, but ensure users can obtain the right level of access when they require it. For example, establish a process for administrators to verify and approve their request, such as an emergency AWS cross-account role for access, or a specific process for administrators to follow to validate and approve an emergency request.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Pre-provision emergency access: Pre-provisioning a role for emergency access from a trusted account, for example one that is used for security team, can help you gain access quickly.
 - [Tutorial: Delegate Access Across AWS Accounts Using IAM Roles](#)

Resources

Related documents:

- [Attribute-based access control \(ABAC\)](#)
- [Grant least privilege](#)
- [IAM Access Analyzer](#)
- [Remove unnecessary credentials](#)
- [Working with Policies](#)

Related videos:

- [Become an IAM Policy Master in 60 Minutes or Less](#)
- [Separation of Duties, Least Privilege, Delegation, and CI/CD](#)

SEC03-BP04 Reduce permissions continuously

As teams and workloads determine what access they need, remove permissions they no longer use and establish review processes to achieve least privilege permissions. Continuously monitor and reduce unused identities and permissions.

Sometimes, when teams and projects are just getting started, you might choose to grant broad access (in a development or test environment) to inspire innovation and agility. We recommend that you evaluate access continuously and, especially in a production environment, restrict access to only the permissions required and achieve least privilege. AWS provides access analysis capabilities to help you identify unused access. To help you identify unused users, roles, permissions, and credentials, AWS analyzes access activity and provides access key and role last used information. You can use the [last accessed timestamp](#) to [identify unused users and roles](#), and remove them. Moreover, you can review service and action last accessed information to identify and [tighten permissions for specific users and roles](#). For example, you can use last accessed information to identify the specific Amazon Simple Storage Service (Amazon S3) actions that your application role requires and restrict access to only those. These features are available in the AWS Management Console and programmatically to enable you to incorporate them into your infrastructure workflows and automated tools.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Configure AWS Identity and Access Management (IAM) Access Analyzer: AWS IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon Simple Storage Service (Amazon S3) buckets or IAM roles, that are shared with an external entity.
 - [AWS IAM Access Analyzer](#)

Resources

Related documents:

- [Attribute-based access control \(ABAC\)](#)
- [Grant least privilege](#)
- [Remove unnecessary credentials](#)
- [Working with Policies](#)

Related videos:

- [Become an IAM Policy Master in 60 Minutes or Less](#)
- [Separation of Duties, Least Privilege, Delegation, and CI/CD](#)

SEC03-BP05 Define permission guardrails for your organization

As you grow and manage additional workloads in AWS, you should separate these workloads using accounts and manage those accounts using AWS Organizations. We recommend that you establish common permission guardrails that restrict access to all identities in your organization. For example, you can restrict access to specific AWS Regions, or prevent your team from deleting common resources, such as an IAM role used by your central security team. You can get started by implementing [example service control policies](#), such as preventing users from disabling key services.

You can use AWS Organizations to group accounts and set common controls on each group of accounts. To set these common controls, you can use services integrated with AWS Organizations. Specifically, you can use [service control policies \(SCPs\)](#) to restrict access to a group of accounts. SCPs use the AWS Identity and Access Management (IAM) policy language and enable you to establish controls that all IAM principals (users and roles) adhere to. You can restrict access to specific service actions, resources and based on specific condition to meet the access control needs of your organization. If necessary, you can define exceptions to your guardrails. For example, you can restrict service actions for all IAM entities in the account except for a specific administrator role.

Establish common controls that restrict access to all identities in your organization. For example, you can restrict access to specific AWS Regions, or prevent your operators from deleting common resources, such as an IAM role used for your central security team.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Define common restrictions that apply to all identities: Based on your organization's unique requirements, create several restrictions that you can apply using AWS Organizations, for example access only to a specific AWS Region.
 - [AWS Organizations Service Control Policies](#)
- Use AWS Control Tower to manage guardrails: Based on your organizations's unique requirements, create several restrictions that you can apply using AWS Organizations, for example access only to a specific AWS Region.
 - [AWS Control Tower Guardrails](#)

Resources

Related documents:

- [Attribute-based access control \(ABAC\)](#)
- [Grant least privilege](#)
- [IAM Access Analyzer](#)
- [Remove unnecessary credentials](#)
- [Working with Policies](#)

Related videos:

- [Become an IAM Policy Master in 60 Minutes or Less](#)
- [Separation of Duties, Least Privilege, Delegation, and CI/CD](#)

SEC03-BP06 Manage access based on lifecycle

Integrate access controls with operator and application lifecycle and your centralized federation provider. For example, remove a user's access when they leave the organization or change roles.

As you manage workloads using separate accounts, there will be cases where you need to share resources between those accounts. We recommend that you share resources using [AWS Resource Access Manager \(AWS RAM\)](#). This service enables you to easily and securely share AWS resources within your AWS Organizations and Organizational Units. Using AWS RAM, access to shared resources is automatically granted or revoked as accounts are moved in and out of the Organization or Organization Unit with which they are shared. This helps ensure that resources are only shared with the accounts that you intend.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

User access lifecycle Implement a user access lifecycle policy for new users joining, job function changes, and users leaving so that only current users have access.

Resources

Related documents:

- [Attribute-based access control \(ABAC\)](#)
- [Grant least privilege](#)
- [IAM Access Analyzer](#)
- [Remove unnecessary credentials](#)
- [Working with Policies](#)

Related videos:

- [Become an IAM Policy Master in 60 Minutes or Less](#)
- [Separation of Duties, Least Privilege, Delegation, and CI/CD](#)

SEC03-BP07 Analyze public and cross-account access

Continuously monitor findings that highlight public and cross-account access. Reduce public access and cross-account access to only resources that require this type of access.

In AWS, you can grant access to resources in another account. You grant direct cross-account access using policies attached to resources (for example, Amazon Simple Storage Service (Amazon S3) bucket policies) or by allowing an identity to assume an IAM role in another account. When using resource policies, you want to ensure you grant access to identities in your organization and are intentional about when you make a resource public. Making a resource public should be used sparingly, as this action allows anyone to access the resource. [AWS Identity and Access Management Access Analyzer \(IAM Access Analyzer\)](#) uses mathematical methods (that is, [provable security](#)) to identify all access paths to a resource from outside of its account. It reviews resource policies continuously, and reports findings of public and cross-account access to make it easy for you to analyze potentially broad access. IAM Access Analyzer also enables you to [preview IAM Access Analyzer findings](#) before deploying resource permissions. This

allows you to validate that your policy changes grant only the intended public and cross-account access to your resources.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Configure IAM Access Analyzer: IAM Access Analyzer helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity.
 - [IAM Access Analyzer](#)

Resources

Related documents:

- [Attribute-based access control \(ABAC\)](#)
- [Grant least privilege](#)
- [Remove unnecessary credentials](#)
- [Working with Policies](#)

Related videos:

- [Become an IAM Policy Master in 60 Minutes or Less](#)
- [Separation of Duties, Least Privilege, Delegation, and CI/CD](#)

SEC03-BP08 Share resources securely

Govern the consumption of shared resources across accounts or within your AWS Organizations. Monitor shared resources and review shared resource access.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Use AWS Resource Access Manager (AWS RAM): AWS RAM is a service that enables you to easily and securely share AWS resources with any AWS account or within your AWS Organizations.
 - [AWS Resource Access Manager](#)

Resources

Related documents:

- [Attribute-based access control \(ABAC\)](#)
- [Grant least privilege](#)
- [IAM Access Analyzer](#)
- [Remove unnecessary credentials](#)
- [Working with Policies](#)

Related videos:

- [Become an IAM Policy Master in 60 Minutes or Less](#)
- [Separation of Duties, Least Privilege, Delegation, and CI/CD](#)

Detection

Question

- [SEC 4 How do you detect and investigate security events? \(p. 137\)](#)

SEC 4 How do you detect and investigate security events?

Capture and analyze events from logs and metrics to gain visibility. Take action on security events and potential threats to help secure your workload.

Best practices

- [SEC04-BP01 Configure service and application logging \(p. 137\)](#)
- [SEC04-BP02 Analyze logs, findings, and metrics centrally \(p. 139\)](#)
- [SEC04-BP03 Automate response to events \(p. 140\)](#)
- [SEC04-BP04 Implement actionable security events \(p. 141\)](#)

SEC04-BP01 Configure service and application logging

Configure logging throughout the workload, including application logs, resource logs, and AWS service logs. For example, ensure that AWS CloudTrail, Amazon CloudWatch Logs, Amazon GuardDuty and AWS Security Hub are enabled for all accounts within your organization.

A foundational practice is to establish a set of detection mechanisms at the account level. This base set of mechanisms is aimed at recording and detecting a wide range of actions on all resources in your account. They allow you to build out a comprehensive detective capability with options that include automated remediation, and partner integrations to add functionality.

In AWS, services that can implement this base set include:

- [AWS CloudTrail](#) provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.
- [AWS Config](#) monitors and records your AWS resource configurations and allows you to automate the evaluation and remediation against desired configurations.
- [Amazon GuardDuty](#) is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts and workloads.
- [AWS Security Hub](#) provides a single place that aggregates, organizes, and prioritizes your security alerts, or findings, from multiple AWS services and optional third-party products to give you a comprehensive view of security alerts and compliance status.

Building on the foundation at the account level, many core AWS services, for example [Amazon Virtual Private Cloud Console \(Amazon VPC\)](#), provide service-level logging features. [Amazon VPC Flow Logs](#) enable you to capture information about the IP traffic going to and from network interfaces that can provide valuable insight into connectivity history, and trigger automated actions based on anomalous behavior.

For Amazon Elastic Compute Cloud (Amazon EC2) instances and application-based logging that doesn't originate from AWS services, logs can be stored and analyzed using [Amazon CloudWatch Logs](#). An [agent](#) collects the logs from the operating system and the applications that are running and automatically stores them. Once the logs are available in CloudWatch Logs, you can [process them in real-time](#), or dive into analysis using [CloudWatch Logs Insights](#).

Equally important to collecting and aggregating logs is the ability to extract meaningful insight from the great volumes of log and event data generated by complex architectures. See the *Monitoring* section

of the [Reliability Pillar whitepaper](#) for more detail. Logs can themselves contain data that is considered sensitive—either when application data has erroneously found its way into log files that the CloudWatch Logs agent is capturing, or when cross-region logging is configured for log aggregation and there are legislative considerations about shipping certain kinds of information across borders.

One approach is to use AWS Lambda functions, triggered on events when logs are delivered, to filter and redact log data before forwarding into a central logging location, such as an Amazon Simple Storage Service (Amazon S3) bucket. The unredacted logs can be retained in a local bucket until a reasonable time has passed (as determined by legislation and your legal team), at which point an Amazon S3 lifecycle rule can automatically delete them. Logs can further be protected in Amazon S3 by using [Amazon S3 Object Lock](#), where you can store objects using a write-once-read-many (WORM) model.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Enable logging of AWS services: Enable the logging of AWS services to meet your requirements. Logging capabilities include the following: Amazon VPC Flow Logs, Elastic Load Balancing (ELB) logs, Amazon S3 bucket logs, CloudFront access logs, Amazon Route 53 query logs, and Amazon Relational Database Service (Amazon RDS) logs.
 - [AWS Answers: native AWS security-logging capabilities](#)
- Evaluate and enable logging of operating systems and application-specific logs to detect suspicious behavior.
 - [Getting started with CloudWatch Logs](#)
 - [Developer Tools and Log Analysis](#)
- Apply appropriate controls to the logs: Logs can contain sensitive information and only authorized users should have access. Consider restricting permissions to Amazon S3 buckets and CloudWatch Logs log groups.
 - [Authentication and Access Control for Amazon CloudWatch](#)
 - [Identity and access management in Amazon S3](#)
- Configure [Amazon GuardDuty](#): GuardDuty is a threat detection service that continuously looks for malicious activity and unauthorized behavior to protect your AWS accounts and workloads. Enable GuardDuty and configure automated alerts to email using the lab.
- [Configure customized trail in CloudTrail](#): Configuring a trail enables you to store logs for longer than the default period, and analyze them later.
- Enable [AWS Config](#): AWS Config provides a detailed view of the configuration of AWS resources in your AWS account. This view includes how the resources are related to one another and how they were previously configured so that you can see how the configurations and relationships change over time.
- Enable [AWS Security Hub](#): Security Hub provides you with a comprehensive view of your security state in AWS and helps you check your compliance with the security industry standards and best practices. Security Hub collects security data from across AWS accounts, services, and supported third-party partner products and helps you analyze your security trends and identify the highest priority security issues.

Resources

Related documents:

- [Amazon CloudWatch](#)
- [Amazon EventBridge](#)
- [Getting started: Amazon CloudWatch Logs](#)
- [Security Partner Solutions: Logging and Monitoring](#)

Related videos:

- [Centrally Monitoring Resource Configuration and Compliance](#)
- [Remediating Amazon GuardDuty and AWS Security Hub Findings](#)
- [Threat management in the cloud: Amazon GuardDuty and AWS Security Hub](#)

Related examples:

- [Lab: Automated Deployment of Detective Controls](#)

SEC04-BP02 Analyze logs, findings, and metrics centrally

Security operations teams rely on the collection of logs and the use of search tools to discover potential events of interest, which might indicate unauthorized activity or unintentional change. However, simply analyzing collected data and manually processing information is insufficient to keep up with the volume of information flowing from complex architectures. Analysis and reporting alone don't facilitate the assignment of the right resources to work an event in a timely fashion.

A best practice for building a mature security operations team is to deeply integrate the flow of security events and findings into a notification and workflow system such as a ticketing system, a bug or issue system, or other security information and event management (SIEM) system. This takes the workflow out of email and static reports, and allows you to route, escalate, and manage events or findings. Many organizations are also integrating security alerts into their chat or collaboration, and developer productivity platforms. For organizations embarking on automation, an API-driven, low-latency ticketing system offers considerable flexibility when planning what to automate first.

This best practice applies not only to security events generated from log messages depicting user activity or network events, but also from changes detected in the infrastructure itself. The ability to detect change, determine whether a change was appropriate, and then route that information to the correct remediation workflow is essential in maintaining and validating a secure architecture, in the context of changes where the nature of their undesirability is sufficiently subtle that their execution cannot currently be prevented with a combination of AWS Identity and Access Management (IAM) and AWS Organizations configuration.

Amazon GuardDuty and AWS Security Hub provide aggregation, deduplication, and analysis mechanisms for log records that are also made available to you via other AWS services. GuardDuty ingests, aggregates, and analyzes information from sources such as AWS CloudTrail management and data events, VPC DNS logs, and VPC Flow Logs. Security Hub can ingest, aggregate, and analyze output from GuardDuty, AWS Config, Amazon Inspector, Amazon Macie, AWS Firewall Manager, and a significant number of third-party security products available in the AWS Marketplace, and if built accordingly, your own code. Both GuardDuty and Security Hub have an Administrator-Member model that can aggregate findings and insights across multiple accounts, and Security Hub is often used by customers who have an on-premises SIEM as an AWS-side log and alert preprocessor and aggregator from which they can then ingest Amazon EventBridge through a AWS Lambda-based processor and forwarder.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Evaluate log processing capabilities: Evaluate the options that are available for processing logs.
 - [Use Amazon OpenSearch Service to log and monitor \(almost\) everything](#)
 - [Find an AWS Partner that specializes in logging and monitoring solutions](#)
- As a start for analyzing CloudTrail logs, test Amazon Athena.
 - [Configuring Athena to analyze CloudTrail logs](#)
- Implement centralized logging in AWS: See the following AWS example solution to centralize logging from multiple sources.
 - [Centralize logging solution](#)

- Implement centralized logging with partner: APN Partners have solutions to help you analyze logs centrally.
 - [Logging and Monitoring](#)

Resources

Related documents:

- [AWS Answers: Centralized Logging](#)
- [AWS Security Hub](#)
- [Amazon CloudWatch](#)
- [Amazon EventBridge](#)
- [Getting started: Amazon CloudWatch Logs](#)
- [Security Partner Solutions: Logging and Monitoring](#)

Related videos:

- [Centrally Monitoring Resource Configuration and Compliance](#)
- [Remediating Amazon GuardDuty and AWS Security Hub Findings](#)
- [Threat management in the cloud: Amazon GuardDuty and AWS Security Hub](#)

SEC04-BP03 Automate response to events

Using automation to investigate and remediate events reduces human effort and error, and enables you to scale investigation capabilities. Regular reviews will help you tune automation tools, and continuously iterate.

In AWS, investigating events of interest and information on potentially unexpected changes into an automated workflow can be achieved using Amazon EventBridge. This service provides a scalable rules engine designed to broker both native AWS event formats (such as AWS CloudTrail events), as well as custom events you can generate from your application. Amazon GuardDuty also allows you to route events to a workflow system for those building incident response systems (AWS Step Functions), or to a central Security Account, or to a bucket for further analysis.

Detecting change and routing this information to the correct workflow can also be accomplished using AWS Config Rules and [Conformance Packs](#). AWS Config detects changes to in-scope services (though with higher latency than EventBridge) and generates events that can be parsed using AWS Config Rules for rollback, enforcement of compliance policy, and forwarding of information to systems, such as change management platforms and operational ticketing systems. As well as writing your own Lambda functions to respond to AWS Config events, you can also take advantage of the [AWS Config Rules Development Kit](#), and a [library of open source](#) AWS Config Rules. Conformance packs are a collection of AWS Config Rules and remediation actions you deploy as a single entity authored as a YAML template. A [sample conformance pack template](#) is available for the Well-Architected Security Pillar.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Implement automated alerting with GuardDuty: GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts and workloads. Enable GuardDuty and configure automated alerts.
- Automate investigation processes: Develop automated processes that investigate an event and report information to an administrator to save time.
 - [Lab: Amazon GuardDuty hands on](#)

Resources

Related documents:

- [AWS Answers: Centralized Logging](#)
- [AWS Security Hub](#)
- [Amazon CloudWatch](#)
- [Amazon EventBridge](#)
- [Getting started: Amazon CloudWatch Logs](#)
- [Security Partner Solutions: Logging and Monitoring](#)
- [Setting up Amazon GuardDuty](#)

Related videos:

- [Centrally Monitoring Resource Configuration and Compliance](#)
- [Remediating Amazon GuardDuty and AWS Security Hub Findings](#)
- [Threat management in the cloud: Amazon GuardDuty and AWS Security Hub](#)

Related examples:

- [Lab: Automated Deployment of Detective Controls](#)

SEC04-BP04 Implement actionable security events

Create alerts that are sent to and can be actioned by your team. Ensure that alerts include relevant information for the team to take action. For each detective mechanism you have, you should also have a process, in the form of a [runbook](#) or [playbook](#), to investigate. For example, when you enable [Amazon GuardDuty](#), it generates different [findings](#). You should have a runbook entry for each finding type, for example, if a [trojan](#) is discovered, your runbook has simple instructions that instruct someone to investigate and remediate.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Discover metrics available for AWS services: Discover the metrics that are available through Amazon CloudWatch for the services that you are using.
 - [AWS service documentation](#)
 - [Using Amazon CloudWatch Metrics](#)
- Configure Amazon CloudWatch alarms.
 - [Using Amazon CloudWatch Alarms](#)

Resources

Related documents:

- [Amazon CloudWatch](#)
- [Amazon EventBridge](#)
- [Security Partner Solutions: Logging and Monitoring](#)

Related videos:

- [Centrally Monitoring Resource Configuration and Compliance](#)
- [Remediating Amazon GuardDuty and AWS Security Hub Findings](#)
- [Threat management in the cloud: Amazon GuardDuty and AWS Security Hub](#)

Infrastructure protection

Questions

- [SEC 5 How do you protect your network resources? \(p. 142\)](#)
- [SEC 6 How do you protect your compute resources? \(p. 146\)](#)

SEC 5 How do you protect your network resources?

Any workload that has some form of network connectivity, whether it's the internet or a private network, requires multiple layers of defense to help protect from external and internal network-based threats.

Best practices

- [SEC05-BP01 Create network layers \(p. 142\)](#)
- [SEC05-BP02 Control traffic at all layers \(p. 143\)](#)
- [SEC05-BP03 Automate network protection \(p. 144\)](#)
- [SEC05-BP04 Implement inspection and protection \(p. 145\)](#)

SEC05-BP01 Create network layers

Group components that share reachability requirements into layers. For example, a database cluster in a virtual private cloud (VPC) with no need for internet access should be placed in subnets with no route to or from the internet. In a serverless workload operating without a VPC, similar layering and segmentation with microservices can achieve the same goal.

Components such as Amazon Elastic Compute Cloud (Amazon EC2) instances, Amazon Relational Database Service (Amazon RDS) database clusters, and AWS Lambda functions that share reachability requirements can be segmented into layers formed by subnets. For example, an Amazon RDS database cluster in a VPC with no need for internet access should be placed in subnets with no route to or from the internet. This layered approach for the controls mitigates the impact of a single layer misconfiguration, which could allow unintended access. For Lambda, you can run your functions in your VPC to take advantage of VPC-based controls.

For network connectivity that can include thousands of VPCs, AWS accounts, and on-premises networks, you should use [AWS Transit Gateway](#). It acts as a hub that controls how traffic is routed among all the connected networks, which act like spokes. Traffic between an Amazon Virtual Private Cloud and AWS Transit Gateway remains on the AWS private network, which reduces external threat vectors such as distributed denial of service (DDoS) attacks and common exploits, such as SQL injection, cross-site scripting, cross-site request forgery, or abuse of broken authentication code. AWS Transit Gateway inter-region peering also encrypts inter-region traffic with no single point of failure or bandwidth bottleneck.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Create subnets in VPC: Create subnets for each layer (in groups that include multiple Availability Zones), and associate route tables to control routing.
 - [VPCs and subnets](#)
 - [Route tables](#)

Resources

Related documents:

- [AWS Firewall Manager](#)
- [Amazon Inspector](#)
- [Amazon VPC Security](#)
- [Getting started with AWS WAF](#)

Related videos:

- [AWS Transit Gateway reference architectures for many VPCs](#)
- [Application Acceleration and Protection with Amazon CloudFront, AWS WAF, and AWS Shield](#)

Related examples:

- [Lab: Automated Deployment of VPC](#)

SEC05-BP02 Control traffic at all layers

When architecting your network topology, you should examine the connectivity requirements of each component. For example, if a component requires internet accessibility (inbound and outbound), connectivity to VPCs, edge services, and external data centers.

A VPC allows you to define your network topology that spans an AWS Region with a private IPv4 address range that you set, or an IPv6 address range AWS selects. You should apply multiple controls with a defense in depth approach for both inbound and outbound traffic, including the use of security groups (stateful inspection firewall), Network ACLs, subnets, and route tables. Within a VPC, you can create subnets in an Availability Zone. Each subnet can have an associated route table that defines routing rules for managing the paths that traffic takes within the subnet. You can define an internet routable subnet by having a route that goes to an internet or NAT gateway attached to the VPC, or through another VPC.

When an instance, Amazon Relational Database Service(Amazon RDS) database, or other service is launched within a VPC, it has its own security group per network interface. This firewall is outside the operating system layer and can be used to define rules for allowed inbound and outbound traffic. You can also define relationships between security groups. For example, instances within a database tier security group only accept traffic from instances within the application tier, by reference to the security groups applied to the instances involved. Unless you are using non-TCP protocols, it shouldn't be necessary to have an Amazon Elastic Compute Cloud(Amazon EC2) instance directly accessible by the internet (even with ports restricted by security groups) without a load balancer, or [CloudFront](#). This helps protect it from unintended access through an operating system or application issue. A subnet can also have a network ACL attached to it, which acts as a stateless firewall. You should configure the network ACL to narrow the scope of traffic allowed between layers, note that you need to define both inbound and outbound rules.

Some AWS services require components to access the internet for making API calls, where [AWS API endpoints](#) are located. Other AWS services use [VPC endpoints](#) within your Amazon VPCs. Many AWS services, including Amazon S3 and Amazon DynamoDB, support VPC endpoints, and this technology has been generalized in [AWS PrivateLink](#). We recommend you use this approach to access AWS services, third-party services, and your own services hosted in other VPCs securely. All network traffic on AWS PrivateLink stays on the global AWS backbone and never traverses the internet. Connectivity can only be initiated by the consumer of the service, and not by the provider of the service. Using AWS PrivateLink for external service access allows you to create air-gapped VPCs with no internet access and helps protect your VPCs from external threat vectors. Third-party services can use AWS PrivateLink to allow their customers to connect to the services from their VPCs over private IP addresses. For VPC assets

that need to make outbound connections to the internet, these can be made outbound only (one-way) through an AWS managed NAT gateway, outbound only internet gateway, or web proxies that you create and manage.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Control network traffic in a VPC: Implement VPC best practices to control traffic.
 - [Amazon VPC security](#)
 - [VPC endpoints](#)
 - [Amazon VPC security group](#)
 - [Network ACLs](#)
- Control traffic at the edge: Implement edge services, such as Amazon CloudFront, to provide an additional layer of protection and other features.
 - [Amazon CloudFront use cases](#)
 - [AWS Global Accelerator](#)
 - [AWS Web Application Firewall \(AWS WAF\)](#)
 - [Amazon Route 53](#)
 - [Amazon VPC Ingress Routing](#)
- Control private network traffic: Implement services that protect your private traffic for your workload.
 - [Amazon VPC Peering](#)
 - [Amazon VPC Endpoint Services \(AWS PrivateLink\)](#)
 - [Amazon VPC Transit Gateway](#)
 - [AWS Direct Connect](#)
 - [AWS Site-to-Site VPN](#)
 - [AWS Client VPN](#)
 - [Amazon S3 Access Points](#)

Resources

Related documents:

- [AWS Firewall Manager](#)
- [Amazon Inspector](#)
- [Getting started with AWS WAF](#)

Related videos:

- [AWS Transit Gateway reference architectures for many VPCs](#)
- [Application Acceleration and Protection with Amazon CloudFront, AWS WAF, and AWS Shield](#)

Related examples:

- [Lab: Automated Deployment of VPC](#)

SEC05-BP03 Automate network protection

Automate protection mechanisms to provide a self-defending network based on threat intelligence and anomaly detection. For example, intrusion detection and prevention tools that can adapt to current

threats and reduce their impact. A web application firewall is an example of where you can automate network protection, for example, by using the AWS WAF Security Automations solution (<https://github.com/aws-labs/aws-waf-security-automations>) to automatically block requests originating from IP addresses associated with known threat actors.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Automate protection for web-based traffic: AWS offers a solution that uses AWS CloudFormation to automatically deploy a set of AWS WAF rules designed to filter common web-based attacks. Users can select from preconfigured protective features that define the rules included in an AWS WAF web access control list (web ACL).
 - [AWS WAF security automations](#)
- Consider AWS Partner solutions: AWS Partners offer hundreds of industry-leading products that are equivalent, identical to, or integrate with existing controls in your on-premises environments. These products complement the existing AWS services to enable you to deploy a comprehensive security architecture and a more seamless experience across your cloud and on-premises environments.
 - [Infrastructure security](#)

Resources

Related documents:

- [AWS Firewall Manager](#)
- [Amazon Inspector](#)
- [Amazon VPC Security](#)
- [Getting started with AWS WAF](#)

Related videos:

- [AWS Transit Gateway reference architectures for many VPCs](#)
- [Application Acceleration and Protection with Amazon CloudFront, AWS WAF, and AWS Shield](#)

Related examples:

- [Lab: Automated Deployment of VPC](#)

SEC05-BP04 Implement inspection and protection

Inspect and filter your traffic at each layer. You can inspect your VPC configurations for potential unintended access using [VPC Network Access Analyzer](#). You can specify your network access requirements and identify potential network paths that do not meet them. For components transacting over HTTP-based protocols, a web application firewall can help protect from common attacks. [AWS WAF](#) is a web application firewall that lets you monitor and block HTTP(s) requests that match your configurable rules that are forwarded to an Amazon API Gateway API, Amazon CloudFront, or an Application Load Balancer. To get started with AWS WAF, you can use [AWS Managed Rules](#) in combination with your own, or use existing [partner integrations](#).

For managing AWS WAF, AWS Shield Advanced protections, and Amazon VPC security groups across AWS Organizations, you can use AWS Firewall Manager. It allows you to centrally configure and manage firewall rules across your accounts and applications, making it easier to scale enforcement of common rules. It also enables you to rapidly respond to attacks, using [AWS Shield Advanced](#), or [solutions](#) that can

automatically block unwanted requests to your web applications. Firewall Manager also works with [AWS Network Firewall](#). AWS Network Firewall is a managed service that uses a rules engine to give you fine-grained control over both stateful and stateless network traffic. It supports the [Suricata compatible](#) open source intrusion prevention system (IPS) specifications for rules to help protect your workload.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Configure Amazon GuardDuty: GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts and workloads. Enable GuardDuty and configure automated alerts.
 - [Amazon GuardDuty](#)
 - [Lab: Automated Deployment of Detective Controls](#)
- Configure virtual private cloud (VPC) Flow Logs: VPC Flow Logs is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC. Flow log data can be published to Amazon CloudWatch Logs and Amazon Simple Storage Service (Amazon S3). After you've created a flow log, you can retrieve and view its data in the chosen destination.
- Consider VPC traffic mirroring: Traffic mirroring is an Amazon VPC feature that you can use to copy network traffic from an elastic network interface of Amazon Elastic Compute Cloud (Amazon EC2) instances and then send it to out-of-band security and monitoring appliances for content inspection, threat monitoring, and troubleshooting.
 - [VPC traffic mirroring](#)

Resources

Related documents:

- [AWS Firewall Manager](#)
- [Amazon Inspector](#)
- [Amazon VPC Security](#)
- [Getting started with AWS WAF](#)

Related videos:

- [AWS Transit Gateway reference architectures for many VPCs](#)
- [Application Acceleration and Protection with Amazon CloudFront, AWS WAF, and AWS Shield](#)

Related examples:

- [Lab: Automated Deployment of VPC](#)

SEC 6 How do you protect your compute resources?

Compute resources in your workload require multiple layers of defense to help protect from external and internal threats. Compute resources include EC2 instances, containers, AWS Lambda functions, database services, IoT devices, and more.

Best practices

- [SEC06-BP01 Perform vulnerability management \(p. 147\)](#)
- [SEC06-BP02 Reduce attack surface \(p. 148\)](#)
- [SEC06-BP03 Implement managed services \(p. 149\)](#)

- [SEC06-BP04 Automate compute protection \(p. 150\)](#)
- [SEC06-BP05 Enable people to perform actions at a distance \(p. 151\)](#)
- [SEC06-BP06 Validate software integrity \(p. 152\)](#)

SEC06-BP01 Perform vulnerability management

Frequently scan and patch for vulnerabilities in your code, dependencies, and in your infrastructure to help protect against new threats.

Starting with the configuration of your compute infrastructure, you can automate creating and updating resources using AWS CloudFormation. CloudFormation allows you to create templates written in YAML or JSON, either using AWS examples or by writing your own. This allows you to create secure-by-default infrastructure templates that you can verify with [CloudFormation Guard](#), to save you time and reduce the risk of configuration error. You can build your infrastructure and deploy your applications using continuous delivery, for example with [AWS CodePipeline](#), to automate the building, testing, and release.

You are responsible for patch management for your AWS resources, including Amazon Elastic Compute Cloud(Amazon EC2) instances, Amazon Machine Images (AMIs), and many other compute resources. For Amazon EC2 instances, AWS Systems Manager Patch Manager automates the process of patching managed instances with both security related and other types of updates. You can use Patch Manager to apply patches for both operating systems and applications. (On Windows Server, application support is limited to updates for Microsoft applications.) You can use Patch Manager to install Service Packs on Windows instances and perform minor version upgrades on Linux instances. You can patch fleets of Amazon EC2 instances or your on-premises servers and virtual machines (VMs) by operating system type. This includes supported versions of Windows Server, Amazon Linux, Amazon Linux 2, CentOS, Debian Server, Oracle Linux, Red Hat Enterprise Linux (RHEL), SUSE Linux Enterprise Server (SLES), and Ubuntu Server. You can scan instances to see only a report of missing patches, or you can scan and automatically install all missing patches.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Configure Amazon Inspector: Amazon Inspector tests the network accessibility of your Amazon Elastic Compute Cloud (Amazon EC2) instances and the security state of the applications that run on those instances. Amazon Inspector assesses applications for exposure, vulnerabilities, and deviations from best practices.
 - [What is Amazon Inspector?](#)
- Scan source code: Scan libraries and dependencies for vulnerabilities.
 - [Amazon CodeGuru](#)
 - [OWASP: Source Code Analysis Tools](#)

Resources

Related documents:

- [AWS Systems Manager](#)
- [Replacing a Bastion Host with Amazon EC2 Systems Manager](#)
- [Security Overview of AWS Lambda](#)

Related videos:

- [Running high-security workloads on Amazon EKS](#)

- [Securing Serverless and Container Services](#)
- [Security best practices for the Amazon EC2 instance metadata service](#)

Related examples:

- [Lab: Automated Deployment of Web Application Firewall](#)

SEC06-BP02 Reduce attack surface

Reduce your exposure to unintended access by hardening operating systems and minimizing the components, libraries, and externally consumable services in use. Start by reducing unused components, whether they are operating system packages or applications, for Amazon Elastic Compute Cloud (Amazon EC2)-based workloads, or external software modules in your code, for all workloads. You can find many hardening and security configuration guides for common operating systems and server software. For example, you can start with the [Center for Internet Security](#) and iterate.

In Amazon EC2, you can create your own Amazon Machine Images (AMIs), which you have patched and hardened, to help you meet the specific security requirements for your organization. The patches and other security controls you apply on the AMI are effective at the point in time in which they were created—they are not dynamic unless you modify after launching, for example, with AWS Systems Manager.

You can simplify the process of building secure AMIs with EC2 Image Builder. EC2 Image Builder significantly reduces the effort required to create and maintain golden images without writing and maintaining automation. When software updates become available, Image Builder automatically produces a new image without requiring users to manually initiate image builds. EC2 Image Builder allows you to easily validate the functionality and security of your images before using them in production with AWS-provided tests and your own tests. You can also apply AWS-provided security settings to further secure your images to meet internal security criteria. For example, you can produce images that conform to the Security Technical Implementation Guide (STIG) standard using AWS-provided templates.

Using third-party static code analysis tools, you can identify common security issues such as unchecked function input bounds, as well as applicable common vulnerabilities and exposures (CVEs). You can use [Amazon CodeGuru](#) for supported languages. Dependency checking tools can also be used to determine whether libraries your code links against are the latest versions, are themselves free of CVEs, and have licensing conditions that meet your software policy requirements.

Using Amazon Inspector, you can perform configuration assessments against your instances for known CVEs, assess against security benchmarks, and automate the notification of defects. Amazon Inspector runs on production instances or in a build pipeline, and it notifies developers and engineers when findings are present. You can access findings programmatically and direct your team to backlogs and bug-tracking systems. [EC2 Image Builder](#) can be used to maintain server images (AMIs) with automated patching, AWS-provided security policy enforcement, and other customizations. When using containers implement [ECR Image Scanning](#) in your build pipeline and on a regular basis against your image repository to look for CVEs in your containers.

While Amazon Inspector and other tools are effective at identifying configurations and any CVEs that are present, other methods are required to test your workload at the application level. [Fuzzing](#) is a well-known method of finding bugs using automation to inject malformed data into input fields and other areas of your application.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Harden operating system: Configure operating systems to meet best practices.
 - [Securing Amazon Linux](#)

- [Securing Microsoft Windows Server](#)
- Harden containerized resources: Configure containerized resources to meet security best practices.
- Implement AWS Lambda best practices.
 - [AWS Lambda best practices](#)

Resources

Related documents:

- [AWS Systems Manager](#)
- [Replacing a Bastion Host with Amazon EC2 Systems Manager](#)
- [Security Overview of AWS Lambda](#)

Related videos:

- [Running high-security workloads on Amazon EKS](#)
- [Securing Serverless and Container Services](#)
- [Security best practices for the Amazon EC2 instance metadata service](#)

Related examples:

- [Lab: Automated Deployment of Web Application Firewall](#)

SEC06-BP03 Implement managed services

Implement services that manage resources, such as Amazon Relational Database Service (Amazon RDS), AWS Lambda, and Amazon Elastic Container Service (Amazon ECS), to reduce your security maintenance tasks as part of the shared responsibility model. For example, Amazon RDS helps you set up, operate, and scale a relational database, automates administration tasks such as hardware provisioning, database setup, patching, and backups. This means you have more free time to focus on securing your application in other ways described in the AWS Well-Architected Framework. Lambda lets you run code without provisioning or managing servers, so you only need to focus on the connectivity, invocation, and security at the code level—not the infrastructure or operating system.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Explore available services: Explore, test, and implement services that manage resources, such as Amazon RDS, AWS Lambda, and Amazon ECS.

Resources

Related documents:

- [AWS Website](#)
- [AWS Systems Manager](#)
- [Replacing a Bastion Host with Amazon EC2 Systems Manager](#)
- [Security Overview of AWS Lambda](#)

Related videos:

- [Running high-security workloads on Amazon EKS](#)
- [Securing Serverless and Container Services](#)
- [Security best practices for the Amazon EC2 instance metadata service](#)

Related examples:

- [Lab: AWS Certificate Manager Request Public Certificate](#)

SEC06-BP04 Automate compute protection

Automate your protective compute mechanisms including vulnerability management, reduction in attack surface, and management of resources. The automation will help you invest time in securing other aspects of your workload, and reduce the risk of human error.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Automate configuration management: Enforce and validate secure configurations automatically by using a configuration management service or tool.
 - [AWS Systems Manager](#)
 - [AWS CloudFormation](#)
 - [Lab: Automated deployment of VPC](#)
 - [Lab: Automated deployment of EC2 web application](#)
- Automate patching of Amazon Elastic Compute Cloud (Amazon EC2) instances: AWS Systems Manager Patch Manager automates the process of patching managed instances with both security-related and other types of updates. You can use Patch Manager to apply patches for both operating systems and applications.
 - [AWS Systems Manager Patch Manager](#)
 - [Centralized multi-account and multi-region patching with AWS Systems Manager Automation](#)
- Implement intrusion detection and prevention: Implement an intrusion detection and prevention tool to monitor and stop malicious activity on instances.
- Consider AWS Partner solutions: AWS Partners offer hundreds of industry-leading products that are equivalent, identical to, or integrate with existing controls in your on-premises environments. These products complement the existing AWS services to enable you to deploy a comprehensive security architecture and a more seamless experience across your cloud and on-premises environments.
 - [Infrastructure security](#)

Resources

Related documents:

- [AWS CloudFormation](#)
- [AWS Systems Manager](#)
- [AWS Systems Manager Patch Manager](#)
- [Centralized multi-account and multi-region patching with AWS Systems Manager Automation](#)
- [Infrastructure security](#)
- [Replacing a Bastion Host with Amazon EC2 Systems Manager](#)

- [Security Overview of AWS Lambda](#)

Related videos:

- [Running high-security workloads on Amazon EKS](#)
- [Securing Serverless and Container Services](#)
- [Security best practices for the Amazon EC2 instance metadata service](#)

Related examples:

- [Lab: Automated Deployment of Web Application Firewall](#)
- [Lab: Automated deployment of EC2 web application](#)

SEC06-BP05 Enable people to perform actions at a distance

Removing the ability for interactive access reduces the risk of human error, and the potential for manual configuration or management. For example, use a change management workflow to deploy Amazon Elastic Compute Cloud (Amazon EC2) instances using infrastructure-as-code, then manage Amazon EC2 instances using tools such as AWS Systems Manager instead of allowing direct access or through a bastion host. AWS Systems Manager can automate a variety of maintenance and deployment tasks, using features including [automation workflows](#), [documents](#) (playbooks), and the [run command](#). AWS CloudFormation stacks build from pipelines and can automate your infrastructure deployment and management tasks without using the AWS Management Console or APIs directly.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- **Replace console access:** Replace console access (SSH or RDP) to instances with AWS Systems Manager Run Command to automate management tasks.

- [AWS Systems Manager Run Command](#)

Resources

Related documents:

- [AWS Systems Manager](#)
- [AWS Systems Manager Run Command](#)
- [Replacing a Bastion Host with Amazon EC2 Systems Manager](#)
- [Security Overview of AWS Lambda](#)

Related videos:

- [Running high-security workloads on Amazon EKS](#)
- [Securing Serverless and Container Services](#)
- [Security best practices for the Amazon EC2 instance metadata service](#)

Related examples:

- [Lab: Automated Deployment of Web Application Firewall](#)

SEC06-BP06 Validate software integrity

Implement mechanisms (for example, code signing) to validate that the software, code and libraries used in the workload are from trusted sources and have not been tampered with. For example, you should verify the code signing certificate of binaries and scripts to confirm the author, and ensure it has not been tampered with since created by the author. [AWS Signer](#) can help ensure the trust and integrity of your code by centrally managing the code-signing lifecycle, including signing certification and public and private keys. You can learn how to use advanced patterns and best practices for code signing with [AWS Lambda](#). Additionally, a checksum of software that you download, compared to that of the checksum from the provider, can help ensure it has not been tampered with.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Investigate mechanisms: Code signing is one mechanism that can be used to validate software integrity.
 - [NIST: Security Considerations for Code Signing](#)

Resources

Related documents:

- [AWS Signer](#)
- [New – Code Signing, a Trust and Integrity Control for AWS Lambda](#)

Data protection

Questions

- [SEC 7 How do you classify your data? \(p. 152\)](#)
- [SEC 8 How do you protect your data at rest? \(p. 155\)](#)
- [SEC 9 How do you protect your data in transit? \(p. 160\)](#)

SEC 7 How do you classify your data?

Classification provides a way to categorize data, based on criticality and sensitivity in order to help you determine appropriate protection and retention controls.

Best practices

- [SEC07-BP01 Identify the data within your workload \(p. 152\)](#)
- [SEC07-BP02 Define data protection controls \(p. 153\)](#)
- [SEC07-BP03 Automate identification and classification \(p. 154\)](#)
- [SEC07-BP04 Define data lifecycle management \(p. 154\)](#)

SEC07-BP01 Identify the data within your workload

You need to understand the type and classification of data your workload is processing, the associated business processes, data owner, applicable legal and compliance requirements, where it's stored, and the resulting controls that are needed to be enforced. This may include classifications to indicate if the data is intended to be publicly available, if the data is internal use only such as customer personally identifiable information (PII), or if the data is for more restricted access such as intellectual property,

legally privileged or marked sensitive, and more. By carefully managing an appropriate data classification system, along with each workload's level of protection requirements, you can map the controls and level of access or protection appropriate for the data. For example, public content is available for anyone to access, but important content is encrypted and stored in a protected manner that requires authorized access to a key for decrypting the content.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Consider discovering data using Amazon Macie: Macie recognizes sensitive data such as personally identifiable information (PII) or intellectual property.
 - [Amazon Macie](#)

Resources

Related documents:

- [Amazon Macie](#)
- [Data Classification Whitepaper](#)
- [Getting started with Amazon Macie](#)

Related videos:

- [Introducing the New Amazon Macie](#)

SEC07-BP02 Define data protection controls

Protect data according to its classification level. For example, secure data classified as public by using relevant recommendations while protecting sensitive data with additional controls.

By using resource tags, separate AWS accounts per sensitivity (and potentially also for each caveat, enclave, or community of interest), IAM policies, AWS Organizations SCPs, AWS Key Management Service (AWS KMS), and AWS CloudHSM, you can define and implement your policies for data classification and protection with encryption. For example, if you have a project with S3 buckets that contain highly critical data or Amazon Elastic Compute Cloud (Amazon EC2) instances that process confidential data, they can be tagged with a `Project=ABC` tag. Only your immediate team knows what the project code means, and it provides a way to use attribute-based access control. You can define levels of access to the AWS KMS encryption keys through key policies and grants to ensure that only appropriate services have access to the sensitive content through a secure mechanism. If you are making authorization decisions based on tags you should make sure that the permissions on the tags are defined appropriately using tag policies in AWS Organizations.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Define your data identification and classification schema: Identification and classification of your data is performed to assess the potential impact and type of data you store, and who can access it.
 - [AWS Documentation](#)
- Discover available AWS controls: For the AWS services you are or plan to use, discover the security controls. Many services have a security section in their documentation.
 - [AWS Documentation](#)
- Identify AWS compliance resources: Identify resources that AWS has available to assist.

- <https://aws.amazon.com/compliance/>

Resources

Related documents:

- [AWS Documentation](#)
- [Data Classification whitepaper](#)
- [Getting started with Amazon Macie](#)
- [AWS Compliance](#)

Related videos:

- [Introducing the New Amazon Macie](#)

SEC07-BP03 Automate identification and classification

Automating the identification and classification of data can help you implement the correct controls. Using automation for this instead of direct access from a person reduces the risk of human error and exposure. You should evaluate using a tool, such as [Amazon Macie](#), that uses machine learning to automatically discover, classify, and protect sensitive data in AWS. Amazon Macie recognizes sensitive data, such as personally identifiable information (PII) or intellectual property, and provides you with dashboards and alerts that give visibility into how this data is being accessed or moved.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use Amazon Simple Storage Service (Amazon S3) Inventory: Amazon S3 inventory is one of the tools you can use to audit and report on the replication and encryption status of your objects.
 - [Amazon S3 Inventory](#)
- Consider Amazon Macie: Amazon Macie uses machine learning to automatically discover and classify data stored in Amazon S3.
 - [Amazon Macie](#)

Resources

Related documents:

- [Amazon Macie](#)
- [Amazon S3 Inventory](#)
- [Data Classification Whitepaper](#)
- [Getting started with Amazon Macie](#)

Related videos:

- [Introducing the New Amazon Macie](#)

SEC07-BP04 Define data lifecycle management

Your defined lifecycle strategy should be based on sensitivity level as well as legal and organization requirements. Aspects including the duration for which you retain data, data destruction processes, data

access management, data transformation, and data sharing should be considered. When choosing a data classification methodology, balance usability versus access. You should also accommodate the multiple levels of access and nuances for implementing a secure, but still usable, approach for each level. Always use a defense in depth approach and reduce human access to data and mechanisms for transforming, deleting, or copying data. For example, require users to strongly authenticate to an application, and give the application, rather than the users, the requisite access permission to perform action at a distance. In addition, ensure that users come from a trusted network path and require access to the decryption keys. Use tools, such as dashboards and automated reporting, to give users information from the data rather than giving them direct access to the data.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Identify data types: Identify the types of data that you are storing or processing in your workload. That data could be text, images, binary databases, and so forth.

Resources

Related documents:

- [Data Classification Whitepaper](#)
- [Getting started with Amazon Macie](#)

Related videos:

- [Introducing the New Amazon Macie](#)

SEC 8 How do you protect your data at rest?

Protect your data at rest by implementing multiple controls, to reduce the risk of unauthorized access or mishandling.

Best practices

- [SEC08-BP01 Implement secure key management \(p. 155\)](#)
- [SEC08-BP02 Enforce encryption at rest \(p. 156\)](#)
- [SEC08-BP03 Automate data at rest protection \(p. 157\)](#)
- [SEC08-BP04 Enforce access control \(p. 158\)](#)
- [SEC08-BP05 Use mechanisms to keep people away from data \(p. 159\)](#)

SEC08-BP01 Implement secure key management

By defining an encryption approach that includes the storage, rotation, and access control of keys, you can help provide protection for your content against unauthorized users and against unnecessary exposure to authorized users. AWS Key Management Service (AWS KMS) helps you manage encryption keys and [integrates with many AWS services](#). This service provides durable, secure, and redundant storage for your AWS KMS keys. You can define your key aliases as well as key-level policies. The policies help you define key administrators as well as key users. Additionally, AWS CloudHSM is a cloud-based hardware security module (HSM) that enables you to easily generate and use your own encryption keys in the AWS Cloud. It helps you meet corporate, contractual, and regulatory compliance requirements for data security by using FIPS 140-2 Level 3 validated HSMs.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Implement AWS KMS: AWS KMS makes it easy for you to create and manage keys and control the use of encryption across a wide range of AWS services and in your applications. AWS KMS is a secure and resilient service that uses FIPS 140-2 validated hardware security modules to protect your keys.
 - [Getting started: AWS Key Management Service \(AWS KMS\)](#)
- Consider AWS Encryption SDK: Use the AWS Encryption SDK with AWS KMS integration when your application needs to encrypt data client-side.
 - [AWS Encryption SDK](#)

Resources

Related documents:

- [AWS Key Management Service](#)
- [AWS cryptographic services and tools](#)
- [Getting started: AWS Key Management Service \(AWS KMS\)](#)
- [Protecting Amazon S3 Data Using Encryption](#)

Related videos:

- [How Encryption Works in AWS](#)
- [Securing Your Block Storage on AWS](#)

SEC08-BP02 Enforce encryption at rest

You should ensure that the only way to store data is by using encryption. AWS Key Management Service (AWS KMS) integrates seamlessly with many AWS services to make it easier for you to encrypt all your data at rest. For example, in Amazon Simple Storage Service (Amazon S3), you can set [default encryption](#) on a bucket so that all new objects are automatically encrypted. Additionally, [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) and [Amazon S3](#) support the enforcement of encryption by setting default encryption. You can use [AWS Config Rules](#) to check automatically that you are using encryption, for example, for [Amazon Elastic Block Store \(Amazon EBS\) volumes](#), [Amazon Relational Database Service \(Amazon RDS\) instances](#), and [Amazon S3 buckets](#).

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Enforce encryption at rest for Amazon Simple Storage Service (Amazon S3): Implement Amazon S3 bucket default encryption.
 - [How do I enable default encryption for an S3 bucket?](#)
- Use AWS Secrets Manager: Secrets Manager is an AWS service that makes it easy for you to manage secrets. Secrets can be database credentials, passwords, third-party API keys, and even arbitrary text.
 - [AWS Secrets Manager](#)
- Configure default encryption for new EBS volumes: Specify that you want all newly created EBS volumes to be created in encrypted form, with the option of using the default key provided by AWS, or a key that you create.
 - [Default encryption for EBS volumes](#)
- Configure encrypted Amazon Machine Images (AMIs): Copying an existing AMI with encryption enabled will automatically encrypt root volumes and snapshots.
 - [AMIs with encrypted Snapshots](#)

- Configure Amazon Relational Database Service (Amazon RDS) encryption: Configure encryption for your Amazon RDS database clusters and snapshots at rest by enabling the encryption option.
 - [Encrypting Amazon RDS resources](#)
- Configure encryption in additional AWS services: For the AWS services you use, determine the encryption capabilities.
 - [AWS Documentation](#)

Resources

Related documents:

- [AMIs with encrypted Snapshots](#)
- [AWS Crypto Tools](#)
- [AWS Documentation](#)
- [AWS Encryption SDK](#)
- [AWS KMS Cryptographic Details Whitepaper](#)
- [AWS Key Management Service](#)
- [AWS Secrets Manager](#)
- [AWS cryptographic services and tools](#)
- [Amazon EBS Encryption](#)
- [Default encryption for EBS volumes](#)
- [Encrypting Amazon RDS Resources](#)
- [How do I enable default encryption for an S3 bucket?](#)
- [Protecting Amazon S3 Data Using Encryption](#)

Related videos:

- [How Encryption Works in AWS](#)
- [Securing Your Block Storage on AWS](#)

SEC08-BP03 Automate data at rest protection

Use automated tools to validate and enforce data at rest controls continuously, for example, verify that there are only encrypted storage resources. You can [automate validation that all EBS volumes are encrypted](#) using [AWS Config Rules](#). [AWS Security Hub](#) can also verify several different controls through automated checks against security standards. Additionally, your AWS Config Rules can automatically [remediate noncompliant resources](#).

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Data at rest represents any data that you persist in non-volatile storage for any duration in your workload. This includes block storage, object storage, databases, archives, IoT devices, and any other storage medium on which data is persisted. Protecting your data at rest reduces the risk of unauthorized access, when encryption and appropriate access controls are implemented.

Enforce encryption at rest: You should ensure that the only way to store data is by using encryption. AWS KMS integrates seamlessly with many AWS services to make it easier for you to encrypt all your data at rest. For example, in Amazon Simple Storage Service (Amazon S3) you can set [default encryption](#) on a bucket so that all new objects are automatically encrypted. Additionally, [Amazon EC2](#) and [Amazon S3](#)

support the enforcement of encryption by setting default encryption. You can use [AWS Managed Config Rules](#) to check automatically that you are using encryption, for example, for [EBS volumes](#), [Amazon Relational Database Service \(Amazon RDS\) instances](#), and [Amazon S3 buckets](#).

Resources

Related documents:

- [AWS Crypto Tools](#)
- [AWS Encryption SDK](#)

Related videos:

- [How Encryption Works in AWS](#)
- [Securing Your Block Storage on AWS](#)

SEC08-BP04 Enforce access control

Enforce access control with least privileges and mechanisms, including backups, isolation, and versioning, to help protect your data at rest. Prevent operators from granting public access to your data.

Different controls including access (using least privilege), backups (see [Reliability whitepaper](#)), isolation, and versioning can all help protect your data at rest. Access to your data should be audited using detective mechanisms covered earlier in this paper including CloudTrail, and service level log, such as Amazon Simple Storage Service (Amazon S3) access logs. You should inventory what data is publicly accessible, and plan for how you can reduce the amount of data available over time. Amazon S3 Glacier Vault Lock and Amazon S3 Object Lock are capabilities providing mandatory access control—once a vault policy is locked with the compliance option, not even the root user can change it until the lock expires. The mechanism meets the Books and Records Management requirements of the SEC, CFTC, and FINRA. For more details, see [this whitepaper](#).

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Enforce access control: Enforce access control with least privileges, including access to encryption keys.
 - [Introduction to Managing Access Permissions to Your Amazon S3 Resources](#)
- Separate data based on different classification levels: Use different AWS accounts for data classification levels managed by AWS Organizations.
 - [AWS Organizations](#)
- Review AWS KMS policies: Review the level of access granted in AWS KMS policies.
 - [Overview of managing access to your AWS KMS resources](#)
- Review Amazon S3 bucket and object permissions: Regularly review the level of access granted in Amazon S3 bucket policies. Best practice is to not have publicly readable or writeable buckets. Consider using AWS Config to detect buckets that are publicly available, and Amazon CloudFront to serve content from Amazon S3.
 - [AWS Config Rules](#)
 - [Amazon S3 + Amazon CloudFront: A Match Made in the Cloud](#)
- Enable Amazon S3 versioning and object lock.
 - [Using versioning](#)
 - [Locking Objects Using Amazon S3 Object Lock](#)
- Use Amazon S3 Inventory: Amazon S3 inventory is one of the tools you can use to audit and report on the replication and encryption status of your objects.

- [Amazon S3 Inventory](#)
- Review Amazon EBS and AMI sharing permissions: Sharing permissions can allow images and volumes to be shared to AWS accounts external to your workload.
 - [Sharing an Amazon EBS Snapshot](#)
 - [Shared AMIs](#)

Resources

Related documents:

- [AWS KMS Cryptographic Details Whitepaper](#)

Related videos:

- [Securing Your Block Storage on AWS](#)

SEC08-BP05 Use mechanisms to keep people away from data

Keep all users away from directly accessing sensitive data and systems under normal operational circumstances. For example, use a change management workflow to manage Amazon Elastic Compute Cloud (Amazon EC2) instances using tools instead of allowing direct access or a bastion host. This can be achieved using [AWS Systems Manager Automation](#), which uses [automation documents](#) that contain steps you use to perform tasks. These documents can be stored in source control, be peer reviewed before running, and tested thoroughly to minimize risk compared to shell access. Business users could have a dashboard instead of direct access to a data store to run queries. Where CI/CD pipelines are not used, determine which controls and processes are required to adequately provide a normally disabled break-glass access mechanism.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Implement mechanisms to keep people away from data: Mechanisms include using dashboards, such as Amazon QuickSight, to display data to users instead of directly querying.
 - [Amazon QuickSight](#)
- Automate configuration management: Perform actions at a distance, enforce and validate secure configurations automatically by using a configuration management service or tool. Avoid use of bastion hosts or directly accessing EC2 instances.
 - [AWS Systems Manager](#)
 - [AWS CloudFormation](#)
 - [CI/CD Pipeline for AWS CloudFormation templates on AWS](#)

Resources

Related documents:

- [AWS KMS Cryptographic Details Whitepaper](#)

Related videos:

- [How Encryption Works in AWS](#)
- [Securing Your Block Storage on AWS](#)

SEC 9 How do you protect your data in transit?

Protect your data in transit by implementing multiple controls to reduce the risk of unauthorized access or loss.

Best practices

- [SEC09-BP01 Implement secure key and certificate management \(p. 160\)](#)
- [SEC09-BP02 Enforce encryption in transit \(p. 160\)](#)
- [SEC09-BP03 Automate detection of unintended data access \(p. 161\)](#)
- [SEC09-BP04 Authenticate network communications \(p. 162\)](#)

SEC09-BP01 Implement secure key and certificate management

Store encryption keys and certificates securely and rotate them at appropriate time intervals with strict access control. The best way to accomplish this is to use a managed service, such as [AWS Certificate Manager \(ACM\)](#). It lets you easily provision, manage, and deploy public and private Transport Layer Security (TLS) certificates for use with AWS services and your internal connected resources. TLS certificates are used to secure network communications and establish the identity of websites over the internet as well as resources on private networks. ACM integrates with AWS resources, such as Elastic Load Balancers (ELBs), AWS distributions, and APIs on API Gateway, also handling automatic certificate renewals. If you use ACM to deploy a private root CA, both certificates and private keys can be provided by it for use in Amazon Elastic Compute Cloud (Amazon EC2) instances, containers, and so on.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Implement secure key and certificate management: Implement your defined secure key and certificate management solution.
 - [AWS Certificate Manager](#)
 - [How to host and manage an entire private certificate infrastructure in AWS](#)
- Implement secure protocols: Use secure protocols that offer authentication and confidentiality, such as Transport Layer Security (TLS) or IPsec, to reduce the risk of data tampering or loss. Check the AWS documentation for the protocols and security relevant to the services that you are using.

Resources

Related documents:

- [AWS Documentation](#)

SEC09-BP02 Enforce encryption in transit

Enforce your defined encryption requirements based on appropriate standards and recommendations to help you meet your organizational, legal, and compliance requirements. AWS services provide HTTPS endpoints using TLS for communication, thus providing encryption in transit when communicating with the AWS APIs. Insecure protocols, such as HTTP, can be audited and blocked in a VPC through the use of security groups. HTTP requests can also be [automatically redirected to HTTPS](#) in Amazon CloudFront or on an [Application Load Balancer](#). You have full control over your computing resources to implement encryption in transit across your services. Additionally, you can use VPN connectivity into your VPC from an external network to facilitate encryption of traffic. Third-party solutions are available in the AWS Marketplace, if you have special requirements.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Enforce encryption in transit: Your defined encryption requirements should be based on the latest standards and best practices and only allow secure protocols. For example, only configure a security group to allow HTTPS protocol to an application load balancer or Amazon Elastic Compute Cloud (Amazon EC2) instance.
- Configure secure protocols in edge services: Configure HTTPS with Amazon CloudFront and required ciphers.
 - [Using HTTPS with CloudFront](#)
- Use a VPN for external connectivity: Consider using an IPsec virtual private network (VPN) for securing point-to-point or network-to-network connections to provide both data privacy and integrity.
 - [VPN connections](#)
- Configure secure protocols in load balancers: Enable HTTPS listener for securing connections to load balancers.
 - [HTTPS listeners for your application load balancer](#)
- Configure secure protocols for instances: Consider configuring HTTPS encryption on instances.
 - [Tutorial: Configure Apache web server on Amazon Linux 2 to use SSL/TLS](#)
- Configure secure protocols in Amazon Relational Database Service (Amazon RDS): Use secure socket layer (SSL) or transport layer security (TLS) to encrypt connection to database instances.
 - [Using SSL to encrypt a connection to a DB Instance](#)
- Configure secure protocols in Amazon Redshift: Configure your cluster to require an secure socket layer (SSL) or transport layer security (TLS) connection.
 - [Configure security options for connections](#)
- Configure secure protocols in additional AWS services For the AWS services you use, determine the encryption-in-transit capabilities.

Resources

Related documents:

- [AWS documentation](#)

SEC09-BP03 Automate detection of unintended data access

Use tools such as Amazon GuardDuty to automatically detect suspicious activity or attempts to move data outside of defined boundaries. For example, GuardDuty can detect Amazon Simple Storage Service (Amazon S3) read activity that is unusual with the [Exfiltration:S3/AnomalousBehavior finding](#). In addition to GuardDuty, [Amazon VPC Flow Logs](#), which capture network traffic information, can be used with Amazon EventBridge to trigger detection of abnormal connections—both successful and denied. [Amazon S3 Access Analyzer](#) can help assess what data is accessible to who in your Amazon S3 buckets.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Automate detection of unintended data access: Use a tool or detection mechanism to automatically detect attempts to move data outside of defined boundaries, for example, to detect a database system that is copying data to an unrecognized host.
 - [VPC Flow Logs](#)
- Consider Amazon Macie: Amazon Macie is a fully managed data security and data privacy service that uses machine learning and pattern matching to discover and protect your sensitive data in AWS.
 - [Amazon Macie](#)

Resources

Related documents:

- [VPC Flow Logs](#)
- [Amazon Macie](#)

SEC09-BP04 Authenticate network communications

Verify the identity of communications by using protocols that support authentication, such as Transport Layer Security (TLS) or IPsec.

Using network protocols that support authentication, allows for trust to be established between the parties. This adds to the encryption used in the protocol to reduce the risk of communications being altered or intercepted. Common protocols that implement authentication include Transport Layer Security (TLS), which is used in many AWS services, and IPsec, which is used in [AWS Virtual Private Network \(AWS VPN\)](#).

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Implement secure protocols: Use secure protocols that offer authentication and confidentiality, such as TLS or IPsec, to reduce the risk of data tampering or loss. Check the [AWS documentation](#) for the protocols and security relevant to the services you are using.

Resources

Related documents:

- [AWS Documentation](#)

Incident response

Question

- [SEC 10 How do you anticipate, respond to, and recover from incidents?](#) (p. 162)

SEC 10 How do you anticipate, respond to, and recover from incidents?

Preparation is critical to timely and effective investigation, response to, and recovery from security incidents to help minimize disruption to your organization.

Best practices

- [SEC10-BP01 Identify key personnel and external resources](#) (p. 163)
- [SEC10-BP02 Develop incident management plans](#) (p. 163)
- [SEC10-BP03 Prepare forensic capabilities](#) (p. 164)
- [SEC10-BP04 Automate containment capability](#) (p. 165)
- [SEC10-BP05 Pre-provision access](#) (p. 166)
- [SEC10-BP06 Pre-deploy tools](#) (p. 166)

- [SEC10-BP07 Run game days \(p. 167\)](#)

SEC10-BP01 Identify key personnel and external resources

Identify internal and external personnel, resources, and legal obligations that would help your organization respond to an incident.

When you define your approach to incident response in the cloud, in unison with other teams (such as your legal counsel, leadership, business stakeholders, AWS Support Services, and others), you must identify key personnel, stakeholders, and relevant contacts. To reduce dependency and decrease response time, make sure that your team, specialist security teams, and responders are educated about the services that you use and have opportunities to practice hands-on.

We encourage you to identify external AWS security partners that can provide you with outside expertise and a different perspective to augment your response capabilities. Your trusted security partners can help you identify potential risks or threats that you might not be familiar with.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Identify key personnel in your organization: Maintain a contact list of personnel within your organization that you would need to involve to respond to and recover from an incident.
- Identify external partners: Engage with external partners if necessary that can help you respond to and recover from an incident.

Resources

Related documents:

- [AWS Incident Response Guide](#)

Related videos:

- [Prepare for and respond to security incidents in your AWS environment](#)

Related examples:

SEC10-BP02 Develop incident management plans

Create plans to help you respond to, communicate during, and recover from an incident. For example, you can start at incident response plan with the most likely scenarios for your workload and organization. Include how you would communicate and escalate both internally and externally. Create incident response plans in the form of [playbooks](#), starting with the most likely scenarios for your workload and organization. These might be events that are currently generated. If you need a starting place, you should look at [AWS Trusted Advisor](#) and [Amazon GuardDuty findings](#). Use a simple format such as markdown so it's easily maintained but ensure that important commands or code snippets are included so they can be executed without having to lookup other documentation.

Start simple and iterate. Work closely with your security experts and partners to identify the tasks required to ensure that the processes are possible. Define the manual descriptions of the processes you perform. After this, test the processes and iterate on the runbook pattern to improve the core logic of your response. Determine what the exceptions are, and what the alternative resolutions are for those scenarios. For example, in a development environment, you might want to terminate a misconfigured Amazon Elastic Compute Cloud (Amazon EC2) instance. But, if the same event occurred in a production

environment, instead of terminating the instance, quarantine the instance by snapshotting volumes, and capturing memory. In addition, verify with stakeholders that critical data will not be lost, and evidence is collected before shutdown. Include how you would communicate and escalate both internally and externally. When you are comfortable with the manual response to the process, automate it to reduce the time to resolution.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Review available resources: AWS and industry resources are available for you to use.
 - [NIST: Computer Security Incident Handling Guide](#)
- Develop incident response playbooks: Easy to follow playbooks should detail steps that you would take to respond and recover from an incident.
- Develop escalation and communications plans: Escalation and communications plans should include personnel in your organization, and external parties that you must notify at each stage during an incident.
- Develop external public relations plan: Develop a plan for public relations to release information about an incident.

Resources

Related documents:

- [AWS Incident Response Guide](#)

Related videos:

- [Automating Incident Response and Forensics in AWS](#)

Related examples:

- [Lab: Incident Response Playbook with Jupyter - AWS IAM](#)

SEC10-BP03 Prepare forensic capabilities

It's important for your incident responders to understand when and how the forensic investigation fits into your response plan. Your organization should define what evidence is collected and what tools are used in the process. Identify and prepare forensic investigation capabilities that are suitable, including external specialists, tools, and automation. A key decision that you should make upfront is if you will collect data from a live system. Some data, such as the contents of volatile memory or active network connections, will be lost if the system is powered off or rebooted.

Your response team can combine tools, such as AWS Systems Manager, Amazon EventBridge, and AWS Lambda, to automatically run forensic tools within an operating system and VPC traffic mirroring to obtain a network packet capture, to gather non-persistent evidence. Conduct other activities, such as log analysis or analyzing disk images, in a dedicated security account with customized forensic workstations and tools accessible to your responders.

Routinely ship relevant logs to a data store that provides high durability and integrity. Responders should have access to those logs. AWS offers several tools that can make log investigation easier, such as Amazon Athena, Amazon OpenSearch Service (OpenSearch Service), and Amazon CloudWatch Logs Insights. Additionally, preserve evidence securely using Amazon Simple Storage Service (Amazon S3) Object Lock. This service follows the WORM (write-once- read-many) model and prevents objects from

being deleted or overwritten for a defined period. As forensic investigation techniques require specialist training, you might need to engage external specialists.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Identify forensic capabilities: Research your organization's forensic investigation capabilities, available tools, and external specialists.
- [Automating Incident Response and Forensics](#)

Resources

Related documents:

- [How to automate forensic disk collection in AWS](#)

SEC10-BP04 Automate containment capability

Automate containment and recovery of an incident to reduce response times and organizational impact.

Once you create and practice the processes and tools from your playbooks, you can deconstruct the logic into a code-based solution, which can be used as a tool by many responders to automate the response and remove variance or guess-work by your responders. This can speed up the lifecycle of a response. The next goal is to enable this code to be fully automated by being invoked by the alerts or events themselves, rather than by a human responder, to create an event-driven response. These processes should also automatically add relevant data to your security systems. For example, an incident involving traffic from an unwanted IP address can automatically populate an AWS WAF block list or Network Firewall rule group to prevent further activity.




Figure 3: AWS WAF automate blocking of known malicious IP addresses.

With an event-driven response system, a detective mechanism triggers a responsive mechanism to automatically remediate the event. You can use event-driven response capabilities to reduce the time-to-value between detective mechanisms and responsive mechanisms. To create this event-driven architecture, you can use AWS Lambda, which is a serverless compute service that runs your code in response to events and automatically manages the underlying compute resources for you. For example, assume that you have an AWS account with the AWS CloudTrail service enabled. If AWS CloudTrail is ever disabled (through the `cloudtrail:StopLogging` API call), you can use Amazon EventBridge to

monitor for the specific `cloudtrail:StopLogging` event, and invoke an AWS Lambda function to call `cloudtrail:StartLogging` to restart logging.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Automate containment capability.

Resources

Related documents:

- [AWS Incident Response Guide](#)

Related videos:

- [Prepare for and respond to security incidents in your AWS environment](#)

SEC10-BP05 Pre-provision access

Ensure that incident responders have the correct access pre-provisioned into AWS and other relevant systems to reduce the time for investigation through to recovery. Determining how to get access for the right people during an incident delays the time it takes to respond, and can introduce other security weaknesses if access is shared or not properly provisioned while under pressure. You must know what level of access your team members require (for example, what kinds of actions they are likely to take) and you must provision access in advance. Access in the form of roles or users created specifically to respond to a security incident are often privileged in order to provide sufficient access. Therefore, use of these user accounts should be restricted, they should not be used for daily activities, and usage alerted on.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Pre-provision access: Ensure that security personnel have the correct access pre-provisioned in AWS so that an appropriate response can be made to an incident.

Resources

Related documents:

- [AWS Incident Response Guide](#)

Related videos:

- [DIY guide to runbooks, incident reports, and incident response](#)

SEC10-BP06 Pre-deploy tools

Ensure that security personnel have the right tools pre-deployed into AWS to reduce the time for investigation through to recovery.

To automate security engineering and operations functions, you can use a comprehensive set of APIs and tools from AWS. You can fully automate identity management, network security, data protection, and monitoring capabilities and deliver them using popular software development methods that you already have in place. When you build security automation, your system can monitor, review, and initiate

a response, rather than having people monitor your security position and manually react to events. An effective way to automatically provide searchable and relevant log data across AWS services to your incident responders is to enable [Amazon Detective](#).

If your incident response teams continue to respond to alerts in the same way, they risk alert fatigue. Over time, the team can become desensitized to alerts and can either make mistakes handling ordinary situations or miss unusual alerts. Automation helps avoid alert fatigue by using functions that process the repetitive and ordinary alerts, leaving humans to handle the sensitive and unique incidents. Integrating anomaly detection systems, such as Amazon GuardDuty, AWS CloudTrail Insights, and Amazon CloudWatch Anomaly Detection, can reduce the burden of common threshold-based alerts.

You can improve manual processes by programmatically automating steps in the process. After you define the remediation pattern to an event, you can decompose that pattern into actionable logic, and write the code to perform that logic. Responders can then execute that code to remediate the issue. Over time, you can automate more and more steps, and ultimately automatically handle whole classes of common incidents.

For tools that execute within the operating system of your Amazon Elastic Compute Cloud (Amazon EC2) instance, you should evaluate using the AWS Systems Manager Run Command, which enables you to remotely and securely administrate instances using an agent that you install on your Amazon EC2 instance operating system. It requires the Systems Manager Agent (SSM Agent), which is installed by default on many Amazon Machine Images (AMIs). Be aware, though, that once an instance has been compromised, no responses from tools or agents running on it should be considered trustworthy.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Pre-deploy tools: Ensure that security personnel have the right tools pre-deployed in AWS so that an appropriate response can be made to an incident.
 - [Lab: Incident response with AWS Management Console and CLI](#)
 - [Incident Response Playbook with Jupyter - AWS IAM](#)
 - [AWS Security Automation](#)
- Implement resource tagging: Tag resources with information, such as a code for the resource under investigation, so that you can identify resources during an incident.
 - [AWS Tagging Strategies](#)

Resources

Related documents:

- [AWS Incident Response Guide](#)

Related videos:

- [DIY guide to runbooks, incident reports, and incident response](#)

SEC10-BP07 Run game days

Game days, also known as simulations or exercises, are internal events that provide a structured opportunity to practice your incident management plans and procedures during a realistic scenario. These events should exercise responders using the same tools and techniques that would be used in a real-world scenario - even mimicking real-world environments. Game days are fundamentally about being prepared and iteratively improving your response capabilities. Some of the reasons you might find value in performing game day activities include:

- Validating readiness
- Developing confidence – learning from simulations and training staff
- Following compliance or contractual obligations
- Generating artifacts for accreditation
- Being agile – incremental improvement
- Becoming faster and improving tools
- Refining communication and escalation
- Developing comfort with the rare and the unexpected

For these reasons, the value derived from participating in a simulation activity increases an organization's effectiveness during stressful events. Developing a simulation activity that is both realistic and beneficial can be a difficult exercise. Although testing your procedures or automation that handles well-understood events has certain advantages, it is just as valuable to participate in creative [Security Incident Response Simulations \(SIRS\)](#) activities to test yourself against the unexpected and continuously improve.

Create custom simulations tailored to your environment, team, and tools. Find an issue and design your simulation around it. This could be something like a leaked credential, a server communicating with unwanted systems, or a misconfiguration that results in unauthorized exposure. Identify engineers who are familiar with your organization to create the scenario and another group to participate. The scenario should be realistic and challenging enough to be valuable. It should include the opportunity to get hands on with logging, notifications, escalations, and executing runbooks or automation. During the simulation, your responders should exercise their technical and organizational skills, and leaders should be involved to build their incident management skills. At the end of the simulation, celebrate the efforts of the team and look for ways to iterate, repeat, and expand into further simulations.

[AWS has created Incident Response Runbook templates](#) that you can use not only to prepare your response efforts, but also as a basis for a simulation. When planning, a simulation can be broken into five phases.

Evidence gathering: In this phase, a team will get alerts through various means, such as an internal ticketing system, alerts from monitoring tooling, anonymous tips, or even public news. Teams then start to review infrastructure and application logs to determine the source of the compromise. This step should also involve internal escalations and incident leadership. Once identified, teams move on to containing the incident

Contain the incident: Teams will have determined there has been an incident and established the source of the compromise. Teams now should take action to contain it, for example, by disabling compromised credentials, isolating a compute resource, or revoking a role's permission.

Eradicate the incident: Now that they've contained the incident, teams will work towards mitigating any vulnerabilities in applications or infrastructure configurations that were susceptible to the compromise. This could include rotating all credentials used for a workload, modifying Access Control Lists (ACLs) or changing network configurations.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Run [game days](#): Run simulated [incident response events \(game days\)](#) for different threats that involve key staff and management.
- Capture lessons learned: Lessons learned from running [game days](#) should be part of a feedback loop to improve your processes.

Resources

Related documents:

- [AWS Incident Response Guide](#)
- [AWS Elastic Disaster Recovery](#)

Related videos:

- [DIY guide to runbooks, incident reports, and incident response](#)

Reliability

Topics

- [Foundations \(p. 169\)](#)
- [Workload architecture \(p. 184\)](#)
- [Change management \(p. 202\)](#)
- [Failure management \(p. 221\)](#)

Foundations

Questions

- [REL 1 How do you manage service quotas and constraints? \(p. 169\)](#)
- [REL 2 How do you plan your network topology? \(p. 175\)](#)

REL 1 How do you manage service quotas and constraints?

For cloud-based workload architectures, there are service quotas (which are also referred to as service limits). These quotas exist to prevent accidentally provisioning more resources than you need and to limit request rates on API operations so as to protect services from abuse. There are also resource constraints, for example, the rate that you can push bits down a fiber-optic cable, or the amount of storage on a physical disk.

Best practices

- [REL01-BP01 Aware of service quotas and constraints \(p. 169\)](#)
- [REL01-BP02 Manage service quotas across accounts and regions \(p. 171\)](#)
- [REL01-BP03 Accommodate fixed service quotas and constraints through architecture \(p. 171\)](#)
- [REL01-BP04 Monitor and manage quotas \(p. 172\)](#)
- [REL01-BP05 Automate quota management \(p. 173\)](#)
- [REL01-BP06 Ensure that a sufficient gap exists between the current quotas and the maximum usage to accommodate failover \(p. 174\)](#)

REL01-BP01 Aware of service quotas and constraints

You are aware of your default quotas and quota increase requests for your workload architecture. You additionally know which resource constraints, such as disk or network, are potentially impactful.

Service Quotas is an AWS service that helps you manage your quotas for over 100 AWS services from one location. Along with looking up the quota values, you can also request and track quota increases from the Service Quotas console or via the AWS SDK. AWS Trusted Advisor offers a service quotas check that displays your usage and quotas for some aspects of some services. The default service quotas per

service are also in the AWS documentation per respective service, for example, see [Amazon VPC Quotas](#). Rate limits on throttled APIs are set within the API Gateway itself by configuring a usage plan. Other limits that are set as configuration on their respective services include Provisioned IOPS, RDS storage allocated, and EBS volume allocations. Amazon Elastic Compute Cloud (Amazon EC2) has its own service limits dashboard that can help you manage your instance, Amazon Elastic Block Store (Amazon EBS), and Elastic IP address limits. If you have a use case where service quotas impact your application's performance and they are not adjustable to your needs, then contact AWS Support to see if there are mitigations.

Common anti-patterns:

- Deploying a workload with no regard of the service quotas on the AWS services used.
- Designing a workload without investigating and accommodating for AWS services' design constraints.
- Deploying a workload with significant use that replaces a known existing workload without configuring the necessary quotas or contacting AWS Support in advance.
- Planning an event to drive traffic to your workload, but not configuring the necessary quotas or contacting AWS Support in advance.

Benefits of establishing this best practice: Being aware of the service quotas, API throttling limits, and design constraints will allow you to account for these in your design, implementation, and operation of the workload.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Review AWS service quotas in the published documentation and Service Quotas
 - [AWS Service Quotas \(formerly referred to as limits\)](#)
- Determine all the services your workload requires by looking at the deployment code.
- Use AWS Config to find all AWS resources used in your AWS accounts.
 - [AWS Config Supported AWS Resource Types and Resource Relationships](#)
- You can also use your AWS CloudFormation to determine your AWS resources used. Look at the resources that were created either in the AWS Management Console or via the `list-stack-resources` CLI command. You can also see resources configured to be deployed in the template itself.
 - [Viewing AWS CloudFormation Stack Data and Resources on the AWS Management Console](#)
 - [AWS CLI for CloudFormation: list-stack-resources](#)
- Determine the service quotas that apply. Use the programmatically accessible information via Trusted Advisor and Service Quotas.

Resources

Related documents:

- [AWS Marketplace: CMDB products that help track limits](#)
- [AWS Service Quotas \(formerly referred to as service limits\)](#)
- [AWS Trusted Advisor Best Practice Checks \(see the Service Limits section\)](#)
- [AWS limit monitor on AWS answers](#)
- [Amazon EC2 Service Limits](#)
- [What is Service Quotas?](#)

Related videos:

- [AWS Live re:Inforce 2019 - Service Quotas](#)

REL01-BP02 Manage service quotas across accounts and regions

If you are using multiple AWS accounts or AWS Regions, ensure that you request the appropriate quotas in all environments in which your production workloads run.

Service quotas are tracked per account. Unless otherwise noted, each quota is AWS Region-specific. In addition to the production environments, also manage quotas in all applicable non-production environments, so that testing and development are not hindered.

Common anti-patterns:

- Allowing resource utilization in one isolation zone to grow with no mechanism to maintain capacity in the other ones.
- Manually setting all quotas independently in isolation zones.
- Not ensuring Regionally isolated deployments are sized to accommodate the increase in traffic from another Region if a deployment is lost.

Benefits of establishing this best practice: Ensuring that you can handle your current load if an isolation zone is unavailable can help reduce the number of errors that occur during failover, instead of causing a denial of service to your customers.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Select relevant accounts and Regions based on your service requirements, latency, regulatory, and disaster recovery (DR) requirements.
- Identify service quotas across all relevant accounts, Regions, and Availability Zones. The limits are scoped to account and Region.
- [What is Service Quotas?](#)

Resources

Related documents:

- [AWS Marketplace: CMDB products that help track limits](#)
- [AWS Service Quotas \(formerly referred to as service limits\)](#)
- [AWS Trusted Advisor Best Practice Checks \(see the Service Limits section\)](#)
- [AWS limit monitor on AWS answers](#)
- [Amazon EC2 Service Limits](#)
- [What is Service Quotas?](#)

Related videos:

- [AWS Live re:Inforce 2019 - Service Quotas](#)

REL01-BP03 Accommodate fixed service quotas and constraints through architecture

Be aware of unchangeable service quotas and physical resources, and architect to prevent these from impacting reliability.

Examples include network bandwidth, AWS Lambda payload size, throttle burst rate for API Gateway, and concurrent user connections to an Amazon Redshift cluster.

Common anti-patterns:

- Performing benchmarking for too short of time, utilizing the burst limit, but then expecting the service to perform at that capacity for sustained periods.
- Choosing a design that uses one resource of a service per user or customer, unaware that there are design constraints that will cause this design to fail as you scale.

Benefits of establishing this best practice: Tracking fixed quotas in AWS services and constraints in other parts of your workload, such as connectivity constraints, IP address constraints, and constraints in third-party services, allows you to detect when you are trending toward a quota and gives you the ability to address the quota before it's exceeded.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Be aware of fixed service quotas Be aware of fixed service quotas and constraints and architect around these.
 - [AWS Service Quotas](#)

Resources

Related documents:

- [AWS Marketplace: CMDB products that help track limits](#)
- [AWS Service Quotas \(formerly referred to as service limits\)](#)
- [AWS Trusted Advisor Best Practice Checks \(see the Service Limits section\)](#)
- [AWS limit monitor on AWS answers](#)
- [Amazon EC2 Service Limits](#)
- [What Is Service Quotas?](#)

Related videos:

- [AWS Live re:Inforce 2019 - Service Quotas](#)

REL01-BP04 Monitor and manage quotas

Evaluate your potential usage and increase your quotas appropriately, allowing for planned growth in usage.

For supported services, you can manage your quotas by configuring CloudWatch alarms to monitor usage and alert you to approaching quotas. These alarms can be triggered from Service Quotas or from Trusted Advisor. You can also use metric filters on CloudWatch Logs to search and extract patterns in logs to determine if usage is approaching quota thresholds.

Common anti-patterns:

- Configuring alarms for when Service Quotas are being approached, but having no process on how to respond to an alert.
- Only configuring alarms for services supported by Service Quotas and not monitoring other services.

Benefits of establishing this best practice: Automatic tracking of the AWS service quotas and monitoring your usage against those quotas will allow you to see when you are approaching a quota limit. You can also use this monitoring data to assess when you might lower quotas to save costs.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Monitor and manage your quotas Evaluate your potential usage on AWS, increase your regional service quotas appropriately, and allow planned growth in usage.
- Capture current resource consumption (for example, buckets, instances). Use service API operations, such as the Amazon EC2 DescribeInstances API, to collect current resource consumption.
- Capture your current quotas Use AWS Service Quotas, AWS Trusted Advisor, and AWS documentation.
 - Use AWS Service Quotas, an AWS service that helps you manage your quotas for over 100 AWS services from one location.
 - Use Trusted Advisor service limits to determine your current service limits.
 - Use service API operations to determine current service quotas where supported.
 - Keep a record of quota increases that have been requested, and their status After a quota increase has been approved, ensure that you update your records to reflect the change to the quota.

Resources

Related documents:

- [AWS Marketplace: CMDB products that help track limits](#)
- [AWS Service Quotas \(formerly referred to as service limits\)](#)
- [AWS Trusted Advisor Best Practice Checks for Service Limits](#)
- [AWS limit monitor on AWS answers](#)
- [Amazon EC2 Service Limits](#)
- [What Is Service Quotas?](#)
- [Monitor Service Quotas using Amazon CloudWatch alarms](#)

Related videos:

- [AWS Live re:Inforce 2019 - Service Quotas](#)

REL01-BP05 Automate quota management

Implement tools to alert you when thresholds are being approached. You can automate quota increase requests by using AWS Service Quotas APIs, you can automate quota increase requests.

If you integrate your Configuration Management Database (CMDB) or ticketing system with Service Quotas, you can automate the tracking of quota increase requests and current quotas. In addition to the AWS SDK, Service Quotas offers automation using the AWS Command Line Interface (AWS CLI).

Common anti-patterns:

- Tracking the quotas and usage in spreadsheets.
- Running reports on usage daily, weekly, or monthly, and then comparing usage to the quotas.

Benefits of establishing this best practice: Automated tracking of the AWS service quotas and monitoring of your usage against that quota allows you to see when you are approaching a quota. You

can set up automation to assist you in requesting a quota increase when needed. You might want to consider lowering some quotas when your usage trends in the opposite direction to realize the benefits of lowered risk (in case of compromised credentials) and cost savings.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Set up automated monitoring Implement tools using SDKs to alert you when thresholds are being approached.
 - Use Service Quotas and augment the service with an automated quota monitoring solution, such as AWS Limit Monitor or an offering from AWS Marketplace.
 - [What is Service Quotas?](#)
 - [Quota Monitor on AWS - AWS Solution](#)
 - Set up triggered responses based on quota thresholds, using Amazon SNS and AWS Service Quotas APIs.
 - Test automation.
 - Configure limit thresholds.
 - Integrate with change events from AWS Config, deployment pipelines, Amazon EventBridge, or third parties.
 - Artificially set low quota thresholds to test responses.
 - Set up triggers to take appropriate action on notifications and contact AWS Support when necessary.
 - Manually trigger change events.
 - Run a game day to test the quota increase change process.

Resources

Related documents:

- [APN Partner: partners that can help with configuration management](#)
- [AWS Marketplace: CMDB products that help track limits](#)
- [AWS Service Quotas \(formerly referred to as service limits\)](#)
- [AWS Trusted Advisor Best Practice Checks \(see the Service Limits section\)](#)
- [Quota Monitor on AWS - AWS Solution](#)
- [Amazon EC2 Service Limits](#)
- [What is Service Quotas?](#)

Related videos:

- [AWS Live re:Inforce 2019 - Service Quotas](#)

REL01-BP06 Ensure that a sufficient gap exists between the current quotas and the maximum usage to accommodate failover

When a resource fails, it might still be counted against quotas until it's successfully terminated. Ensure that your quotas cover the overlap of all failed resources with replacements before the failed resources are terminated. You should consider an Availability Zone failure when calculating this gap.

Common anti-patterns:

- Setting service quotas based on current needs without accounting for failover scenarios.

Benefits of establishing this best practice: When events potentially impact availability, the cloud allows you to implement strategies to mitigate or recover from these events. Such strategies often include creating additional resources to replace failed ones. Your quota strategy must accommodate these additional resources.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Ensure that there is enough gap between your service quota and your maximum usage to accommodate for a failover.
- Determine your service quotas, accounting for your deployment patterns, availability requirements, and consumption growth.
- Request quota increases if necessary. Plan for necessary time for quota increase requests to be fulfilled.
 - Determine your reliability requirements (also known as your number of 9's).
 - Establish your fault scenarios (for example, loss of a component, an Availability Zone, or a Region).
 - Establish your deployment methodology (for example, canary, blue/green, red/black, or rolling).
 - Include an appropriate buffer (for example, 15%) to the current limit.
 - Plan consumption growth (for example, monitor your trends in consumption).

Resources

Related documents:

- [AWS Marketplace: CMDB products that help track limits](#)
- [AWS Service Quotas \(formerly referred to as service limits\)](#)
- [AWS Trusted Advisor Best Practice Checks \(see the Service Limits section\)](#)
- [Amazon EC2 Service Limits](#)
- [What Is Service Quotas?](#)

Related videos:

- [AWS Live re:Inforce 2019 - Service Quotas](#)

REL 2 How do you plan your network topology?

Workloads often exist in multiple environments. These include multiple cloud environments (both publicly accessible and private) and possibly your existing data center infrastructure. Plans must include network considerations such as intra- and inter-system connectivity, public IP address management, private IP address management, and domain name resolution.

Best practices

- [REL02-BP01 Use highly available network connectivity for your workload public endpoints \(p. 176\)](#)
- [REL02-BP02 Provision redundant connectivity between private networks in the cloud and on-premises environments \(p. 177\)](#)
- [REL02-BP03 Ensure IP subnet allocation accounts for expansion and availability \(p. 180\)](#)
- [REL02-BP04 Prefer hub-and-spoke topologies over many-to-many mesh \(p. 181\)](#)
- [REL02-BP05 Enforce non-overlapping private IP address ranges in all private address spaces where they are connected \(p. 183\)](#)

REL02-BP01 Use highly available network connectivity for your workload public endpoints

These endpoints and the routing to them must be highly available. To achieve this, use highly available DNS, content delivery networks (CDNs), API Gateway, load balancing, or reverse proxies.

Amazon Route 53, AWS Global Accelerator, Amazon CloudFront, Amazon API Gateway, and Elastic Load Balancing (ELB) all provide highly available public endpoints. You might also choose to evaluate AWS Marketplace software appliances for load balancing and proxying.

Consumers of the service your workload provides, whether they are end-users or other services, make requests on these service endpoints. Several AWS resources are available to enable you to provide highly available endpoints.

Elastic Load Balancing provides load balancing across Availability Zones, performs Layer 4 (TCP) or Layer 7 (http/https) routing, integrates with AWS WAF, and integrates with AWS Auto Scaling to help create a self-healing infrastructure and absorb increases in traffic while releasing resources when traffic decreases.

Amazon Route 53 is a scalable and highly available Domain Name System (DNS) service that connects user requests to infrastructure running in AWS such as Amazon EC2 instances, Elastic Load Balancing load balancers, or Amazon S3 buckets—and can also be used to route users to infrastructure outside of AWS.

AWS Global Accelerator is a network layer service that you can use to direct traffic to optimal endpoints over the AWS global network.

Distributed Denial of Service (DDoS) attacks risk shutting out legitimate traffic and lowering availability for your users. AWS Shield provides automatic protection against these attacks at no extra cost for AWS service endpoints on your workload. You can augment these features with virtual appliances from APN Partners and the AWS Marketplace to meet your needs.

Common anti-patterns:

- Using public internet addresses on instances or containers and managing the connectivity to them via DNS.
- Using Internet Protocol addresses instead of domain names for locating services.
- Providing content (web pages, static assets, media files) to a large geographic area and not using a content delivery network.

Benefits of establishing this best practice: By implementing highly available services in your workload, you know that your workload will be available to your users.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Ensure that you have highly available connectivity for users of the workload. Amazon Route 53, AWS Global Accelerator, Amazon CloudFront, Amazon API Gateway, and Elastic Load Balancing (ELB) all provide highly available public facing endpoints. You may also choose to evaluate AWS Marketplace software appliances for load-balancing and proxying.

- Ensure that you have a highly available connection to your users.
- Ensure that you are using a highly available DNS to manage the domain names of your application endpoints.
 - If your users access your application via the internet, use service API operations to confirm the correct usage of Internet Gateways. Also confirm that the route table entries for the subnets hosting your application endpoints are correct.

- [DescribeInternetGateways](#)
- [DescribeRouteTables](#)
- Ensure that you are using a highly available reverse proxy or load balancer in front of your application.
 - If your users access your application via your on-premises environment, ensure that your connectivity between AWS and your on-premises environment is highly available.
- Use Route 53 to manage your domain names.
 - [What is Amazon Route 53?](#)
- Use a third-party DNS provider that meets your requirements.
- Use Elastic Load Balancing.
 - [What is Elastic Load Balancing?](#)
- Use an AWS Marketplace appliance that meets your requirements.

Resources

Related documents:

- [APN Partner: partners that can help plan your networking](#)
- [AWS Direct Connect Resiliency Recommendations](#)
- [AWS Marketplace for Network Infrastructure](#)
- [Amazon Virtual Private Cloud Connectivity Options Whitepaper](#)
- [Multiple data center HA network connectivity](#)
- [Using the Direct Connect Resiliency Toolkit to get started](#)
- [VPC Endpoints and VPC Endpoint Services \(AWS PrivateLink\)](#)
- [What Is AWS Global Accelerator?](#)
- [What Is Amazon VPC?](#)
- [What Is a Transit Gateway?](#)
- [What is Amazon CloudFront?](#)
- [What is Amazon Route 53?](#)
- [What is Elastic Load Balancing?](#)
- [Working with Direct Connect Gateways](#)

Related videos:

- [AWS re:Invent 2018: Advanced VPC Design and New Capabilities for Amazon VPC \(NET303\)](#)
- [AWS re:Invent 2019: AWS Transit Gateway reference architectures for many VPCs \(NET406-R1\)](#)

REL02-BP02 Provision redundant connectivity between private networks in the cloud and on-premises environments

Use multiple AWS Direct Connect connections or VPN tunnels between separately deployed private networks. Use multiple Direct Connect locations for high availability. If using multiple AWS Regions, ensure redundancy in at least two of them. You might want to evaluate AWS Marketplace appliances that terminate VPNs. If you use AWS Marketplace appliances, deploy redundant instances for high availability in different Availability Zones.

AWS Direct Connect is a cloud service that makes it easy to establish a dedicated network connection from your on-premises environment to AWS. Using Direct Connect Gateway, your on-premises data center can be connected to multiple AWS VPCs spread across multiple AWS Regions.

This redundancy addresses possible failures that impact connectivity resiliency:

- How are you going to be resilient to failures in your topology?
- What happens if you misconfigure something and remove connectivity?
- Will you be able to handle an unexpected increase in traffic or use of your services?
- Will you be able to absorb an attempted Distributed Denial of Service (DDoS) attack?

When connecting your VPC to your on-premises data center via VPN, you should consider the resiliency and bandwidth requirements that you need when you select the vendor and instance size on which you need to run the appliance. If you use a VPN appliance that is not resilient in its implementation, then you should have a redundant connection through a second appliance. For all these scenarios, you need to define an acceptable time to recovery and test to ensure that you can meet those requirements.

If you choose to connect your VPC to your data center using a Direct Connect connection and you need this connection to be highly available, have redundant Direct Connect connections from each data center. The redundant connection should use a second Direct Connect connection from different location than the first. If you have multiple data centers, ensure that the connections terminate at different locations. Use the [Direct Connect Resiliency Toolkit](#) to help you set this up.

If you choose to fail over to VPN over the internet using AWS VPN, it's important to understand that it supports up to 1.25-Gbps throughput per VPN tunnel, but does not support Equal Cost Multi Path (ECMP) for outbound traffic in the case of multiple AWS Managed VPN tunnels terminating on the same VGW. We do not recommend that you use AWS Managed VPN as a backup for Direct Connect connections unless you can tolerate speeds less than 1 Gbps during failover.

You can also use VPC endpoints to privately connect your VPC to supported AWS services and VPC endpoint services powered by AWS PrivateLink without traversing the public internet. Endpoints are virtual devices. They are horizontally scaled, redundant, and highly available VPC components. They allow communication between instances in your VPC and services without imposing availability risks or bandwidth constraints on your network traffic.

Common anti-patterns:

- Having only one connectivity provider between your on-site network and AWS.
- Consuming the connectivity capabilities of your AWS Direct Connect connection, but only having one connection.
- Having only one path for your VPN connectivity.

Benefits of establishing this best practice: By implementing redundant connectivity between your cloud environment and your corporate or on-premises environment, you can ensure that the dependent services between the two environments can communicate reliably.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Ensure that you have highly available connectivity between AWS and on-premises environment. Use multiple AWS Direct Connect connections or VPN tunnels between separately deployed private networks. Use multiple Direct Connect locations for high availability. If using multiple AWS Regions, ensure redundancy in at least two of them. You might want to evaluate AWS Marketplace appliances that terminate VPNs. If you use AWS Marketplace appliances, deploy redundant instances for high availability in different Availability Zones.
- Ensure that you have a redundant connection to your on-premises environment. You may need redundant connections to multiple AWS Regions to achieve your availability needs.
 - [AWS Direct Connect Resiliency Recommendations](#)

- [Using Redundant Site-to-Site VPN Connections to Provide Failover](#)
 - Use service API operations to identify correct use of Direct Connect circuits.
 - [DescribeConnections](#)
 - [DescribeConnectionsOnInterconnect](#)
 - [DescribeDirectConnectGatewayAssociations](#)
 - [DescribeDirectConnectGatewayAttachments](#)
 - [DescribeDirectConnectGateways](#)
 - [DescribeHostedConnections](#)
 - [DescribeInterconnects](#)
 - If only one Direct Connect connection exists or you have none, set up redundant VPN tunnels to your virtual private gateways.
 - [What is AWS Site-to-Site VPN?](#)
- Capture your current connectivity (for example, Direct Connect, virtual private gateways, AWS Marketplace appliances).
 - Use service API operations to query configuration of Direct Connect connections.
 - [DescribeConnections](#)
 - [DescribeConnectionsOnInterconnect](#)
 - [DescribeDirectConnectGatewayAssociations](#)
 - [DescribeDirectConnectGatewayAttachments](#)
 - [DescribeDirectConnectGateways](#)
 - [DescribeHostedConnections](#)
 - [DescribeInterconnects](#)
 - Use service API operations to collect virtual private gateways where route tables use them.
 - [DescribeVpnGateways](#)
 - [DescribeRouteTables](#)
 - Use service API operations to collect AWS Marketplace applications where route tables use them.
 - [DescribeRouteTables](#)

Resources

Related documents:

- [APN Partner: partners that can help plan your networking](#)
- [AWS Direct Connect Resiliency Recommendations](#)
- [AWS Marketplace for Network Infrastructure](#)
- [Amazon Virtual Private Cloud Connectivity Options Whitepaper](#)
- [Multiple data center HA network connectivity](#)
- [Using Redundant Site-to-Site VPN Connections to Provide Failover](#)
- [Using the Direct Connect Resiliency Toolkit to get started](#)
- [VPC Endpoints and VPC Endpoint Services \(AWS PrivateLink\)](#)
- [What Is Amazon VPC?](#)
- [What Is a Transit Gateway?](#)
- [What is AWS Site-to-Site VPN?](#)
- [Working with Direct Connect Gateways](#)

Related videos:

- [AWS re:Invent 2018: Advanced VPC Design and New Capabilities for Amazon VPC \(NET303\)](#)
- [AWS re:Invent 2019: AWS Transit Gateway reference architectures for many VPCs \(NET406-R1\)](#)

REL02-BP03 Ensure IP subnet allocation accounts for expansion and availability

Amazon VPC IP address ranges must be large enough to accommodate workload requirements, including factoring in future expansion and allocation of IP addresses to subnets across Availability Zones. This includes load balancers, EC2 instances, and container-based applications.

When you plan your network topology, the first step is to define the IP address space itself. Private IP address ranges (following RFC 1918 guidelines) should be allocated for each VPC. Accommodate the following requirements as part of this process:

- Allow IP address space for more than one VPC per Region.
- Within a VPC, allow space for multiple subnets that span multiple Availability Zones.
- Always leave unused CIDR block space within a VPC for future expansion.
- Ensure that there is IP address space to meet the needs of any transient fleets of EC2 instances that you might use, such as Spot Fleets for machine learning, Amazon EMR clusters, or Amazon Redshift clusters.
- Note that the first four IP addresses and the last IP address in each subnet CIDR block are reserved and not available for your use.
- You should plan on deploying large VPC CIDR blocks. Note that the initial VPC CIDR block allocated to your VPC cannot be changed or deleted, but you can add additional non-overlapping CIDR blocks to the VPC. Subnet IPv4 CIDRs cannot be changed, however IPv6 CIDRs can. Keep in mind that deploying the largest VPC possible (/16) results in over 65,000 IP addresses. In the base 10.x.x.x IP address space alone, you could provision 255 such VPCs. You should therefore err on the side of being too large rather than too small to make it easier to manage your VPCs.

Common anti-patterns:

- Creating small VPCs.
- Creating small subnets and then having to add subnets to configurations as you grow.
- Incorrectly estimating how many IP addresses a elastic load balancer can use.
- Deploying many high traffic load balancers into the same subnets.

Benefits of establishing this best practice: This ensures that you can accommodate the growth of your workloads and continue to provide availability as you scale up.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Plan your network to accommodate for growth, regulatory compliance, and integration with others. Growth can be underestimated, regulatory compliance can change, and acquisitions or private network connections can be difficult to implement without proper planning.
- Select relevant AWS accounts and Regions based on your service requirements, latency, regulatory, and disaster recovery (DR) requirements.
- Identify your needs for regional VPC deployments.
- Identify the size of the VPCs.
 - Determine if you are going to deploy multi-VPC connectivity.
 - [What Is a Transit Gateway?](#)
 - [Single Region Multi-VPC Connectivity](#)

- Determine if you need segregated networking for regulatory requirements.
- Make VPCs as large as possible. The initial VPC CIDR block allocated to your VPC cannot be changed or deleted, but you can add additional non-overlapping CIDR blocks to the VPC. This however may fragment your address ranges.

Resources

Related documents:

- [APN Partner: partners that can help plan your networking](#)
- [AWS Marketplace for Network Infrastructure](#)
- [Amazon Virtual Private Cloud Connectivity Options Whitepaper](#)
- [Multiple data center HA network connectivity](#)
- [Single Region Multi-VPC Connectivity](#)
- [What Is Amazon VPC?](#)

Related videos:

- [AWS re:Invent 2018: Advanced VPC Design and New Capabilities for Amazon VPC \(NET303\)](#)
- [AWS re:Invent 2019: AWS Transit Gateway reference architectures for many VPCs \(NET406-R1\)](#)

REL02-BP04 Prefer hub-and-spoke topologies over many-to-many mesh

If more than two network address spaces (for example, VPCs and on-premises networks) are connected via VPC peering, AWS Direct Connect, or VPN, then use a hub-and-spoke model, like that provided by AWS Transit Gateway.

If you have only two such networks, you can simply connect them to each other, but as the number of networks grows, the complexity of such meshed connections becomes untenable. AWS Transit Gateway provides an easy to maintain hub-and-spoke model, allowing the routing of traffic across your multiple networks.




Figure 1: Without AWS Transit Gateway: You need to peer each Amazon VPC to each other and to each onsite location using a VPN connection, which can become complex as it scales.




Figure 2: With AWS Transit Gateway: You simply connect each Amazon VPC or VPN to the AWS Transit Gateway and it routes traffic to and from each VPC or VPN.

Common anti-patterns:

- Using VPC peering to connect more than two VPCs.
- Establishing multiple BGP sessions for each VPC to establish connectivity that spans Virtual Private Clouds (VPCs) spread across multiple AWS Regions.

Benefits of establishing this best practice: As the number of networks grows, the complexity of such meshed connections becomes untenable. AWS Transit Gateway provides an easy to maintain hub-and-spoke model, allowing routing of traffic among your multiple networks.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Prefer hub-and-spoke topologies over many-to-many mesh. If more than two network address spaces (VPCs, on-premises networks) are connected via VPC peering, AWS Direct Connect, or VPN, then use a hub-and-spoke model like that provided by AWS Transit Gateway.
- For only two such networks, you can simply connect them to each other, but as the number of networks grows, the complexity of such meshed connections becomes untenable. AWS Transit Gateway provides an easy to maintain hub-and-spoke model, allowing routing of traffic across your multiple networks.
 - [What Is a Transit Gateway?](#)

Resources

Related documents:

- [APN Partner: partners that can help plan your networking](#)
- [AWS Marketplace for Network Infrastructure](#)
- [Multiple data center HA network connectivity](#)
- [VPC Endpoints and VPC Endpoint Services \(AWS PrivateLink\)](#)
- [What Is Amazon VPC?](#)

- [What Is a Transit Gateway?](#)

Related videos:

- [AWS re:Invent 2018: Advanced VPC Design and New Capabilities for Amazon VPC \(NET303\)](#)
- [AWS re:Invent 2019: AWS Transit Gateway reference architectures for many VPCs \(NET406-R1\)](#)

REL02-BP05 Enforce non-overlapping private IP address ranges in all private address spaces where they are connected

The IP address ranges of each of your VPCs must not overlap when peered or connected via VPN. You must similarly avoid IP address conflicts between a VPC and on-premises environments or with other cloud providers that you use. You must also have a way to allocate private IP address ranges when needed.

An IP address management (IPAM) system can help with this. Several IPAMs are available from the AWS Marketplace.

Common anti-patterns:

- Using the same IP range in your VPC as you have on premises or in your corporate network.
- Not tracking IP ranges of VPCs used to deploy your workloads.

Benefits of establishing this best practice: Active planning of your network will ensure that you do not have multiple occurrences of the same IP address in interconnected networks. This prevents routing problems from occurring in parts of the workload that are using the different applications.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Monitor and manage your CIDR use. Evaluate your potential usage on AWS, add CIDR ranges to existing VPCs, and create VPCs to allow planned growth in usage.
 - Capture current CIDR consumption (for example, VPCs, subnets)
 - Use service API operations to collect current CIDR consumption.
 - Capture your current subnet usage.
 - Use service API operations to collect subnets per VPC in each Region.
 - [DescribeSubnets](#)
 - Record the current usage.
 - Determine if you created any overlapping IP ranges.
 - Calculate the spare capacity.
 - Identify overlapping IP ranges. You can either migrate to a new range of addresses or use Network and Port Translation (NAT) appliances from AWS Marketplace if you need to connect the overlapping ranges.

Resources

Related documents:

- [APN Partner: partners that can help plan your networking](#)
- [AWS Marketplace for Network Infrastructure](#)
- [Amazon Virtual Private Cloud Connectivity Options Whitepaper](#)
- [Multiple data center HA network connectivity](#)

- [What Is Amazon VPC?](#)
- [What is IPAM?](#)

Related videos:

- [AWS re:Invent 2018: Advanced VPC Design and New Capabilities for Amazon VPC \(NET303\)](#)
- [AWS re:Invent 2019: AWS Transit Gateway reference architectures for many VPCs \(NET406-R1\)](#)

Workload architecture

Questions

- [REL 3 How do you design your workload service architecture? \(p. 184\)](#)
- [REL 4 How do you design interactions in a distributed system to prevent failures? \(p. 189\)](#)
- [REL 5 How do you design interactions in a distributed system to mitigate or withstand failures? \(p. 194\)](#)

REL 3 How do you design your workload service architecture?

Build highly scalable and reliable workloads using a service-oriented architecture (SOA) or a microservices architecture. Service-oriented architecture (SOA) is the practice of making software components reusable via service interfaces. Microservices architecture goes further to make components smaller and simpler.

Best practices

- [REL03-BP01 Choose how to segment your workload \(p. 184\)](#)
- [REL03-BP02 Build services focused on specific business domains and functionality \(p. 186\)](#)
- [REL03-BP03 Provide service contracts per API \(p. 188\)](#)

REL03-BP01 Choose how to segment your workload

Workload segmentation is important when determining the resilience requirements of your application. Monolithic architecture should be avoided whenever possible. Instead, carefully consider which application components can be broken out into microservices. Depending on your application requirements, this may end up being a combination of a service-oriented architecture (SOA) with microservices where possible. Workloads that are capable of statelessness are more capable of being deployed as microservices.

Desired outcome: Workloads should be supportable, scalable, and as loosely coupled as possible.

When making choices about how to segment your workload, balance the benefits against the complexities. What is right for a new product racing to first launch is different than what a workload built to scale from the start needs. When refactoring an existing monolith, you will need to consider how well the application will support a decomposition towards statelessness. Breaking services into smaller pieces allows small, well-defined teams to develop and manage them. However, smaller services can introduce complexities which include possible increased latency, more complex debugging, and increased operational burden.

Common anti-patterns:

- The [microservice Death Star](#) is a situation in which the atomic components become so highly interdependent that a failure of one results in a much larger failure, making the components as rigid and fragile as a monolith.

Benefits of establishing this practice:


- More specific segments lead to greater agility, organizational flexibility, and scalability.
- Reduced impact of service interruptions.
- Application components may have different availability requirements, which can be supported by a more atomic segmentation.
- Well-defined responsibilities for teams supporting the workload.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Choose your architecture type based on how you will segment your workload. Choose an SOA or microservices architecture (or in some rare cases, a monolithic architecture). Even if you choose to start with a monolith architecture, you must ensure that it's modular and can ultimately evolve to SOA or microservices as your product scales with user adoption. SOA and microservices offer respectively smaller segmentation, which is preferred as a modern scalable and reliable architecture, but there are trade-offs to consider, especially when deploying a microservice architecture.

One primary trade-off is that you now have a distributed compute architecture that can make it harder to achieve user latency requirements and there is additional complexity in the debugging and tracing of user interactions. You can use AWS X-Ray to assist you in solving this problem. Another effect to consider is increased operational complexity as you increase the number of applications that you are managing, which requires the deployment of multiple independency components.



Monolithic, service-oriented, and microservices architectures

Implementation steps

- Determine the appropriate architecture to refactor or build your application. SOA and microservices offer respectively smaller segmentation, which is preferred as a modern scalable and reliable architecture. SOA can be a good compromise for achieving smaller segmentation while avoiding some of the complexities of microservices. For more details, see [Microservice Trade-Offs](#).
- If your workload is amenable to it, and your organization can support it, you should use a microservices architecture to achieve the best agility and reliability. For more details, see [Implementing Microservices on AWS](#).
- Consider following the [Strangler Fig pattern](#) to refactor a monolith into smaller components. This involves gradually replacing specific application components with new applications and services. [AWS](#)

[Migration Hub Refactor Spaces](#) acts as the starting point for incremental refactoring. For more details, see [Seamlessly migrate on-premises legacy workloads using a strangler pattern](#).

- Implementing microservices may require a service discovery mechanism to allow these distributed services to communicate with each other. [AWS App Mesh](#) can be used with service-oriented architectures to provide reliable discovery and access of services. [AWS Cloud Map](#) can also be used for dynamic, DNS-based service discovery.
- If you're migrating from a monolith to SOA, [Amazon MQ](#) can help bridge the gap as a service bus when redesigning legacy applications in the cloud.
- For existing monoliths with a single, shared database, choose how to reorganize the data into smaller segments. This could be by business unit, access pattern, or data structure. At this point in the refactoring process, you should choose to move forward with a relational or non-relational (NoSQL) type of database. For more details, see [From SQL to NoSQL](#).

Level of effort for the implementation plan: High

Resources

Related best practices:

- [REL03-BP02 Build services focused on specific business domains and functionality \(p. 186\)](#)

Related documents:

- [Amazon API Gateway: Configuring a REST API Using OpenAPI](#)
- [What is Service-Oriented Architecture?](#)
- [Bounded Context \(a central pattern in Domain-Driven Design\)](#)
- [Implementing Microservices on AWS](#)
- [Microservice Trade-Offs](#)
- [Microservices - a definition of this new architectural term](#)
- [Microservices on AWS](#)
- [What is AWS App Mesh?](#)

Related examples:

- [Iterative App Modernization Workshop](#)

Related videos:


- [Delivering Excellence with Microservices on AWS](#)

REL03-BP02 Build services focused on specific business domains and functionality

Service-oriented architecture (SOA) builds services with well-delineated functions defined by business needs. Microservices use domain models and bounded context to limit this further so that each service does just one thing. Focusing on specific functionality enables you to differentiate the reliability requirements of different services, and target investments more specifically. A concise business problem and having a small team associated with each service also enables easier organizational scaling.

In designing a microservice architecture, it's helpful to use Domain-Driven Design (DDD) to model the business problem using entities. For example, for the Amazon.com website, entities might include

package, delivery, schedule, price, discount, and currency. Then the model is further divided into smaller models using [Bounded Context](#), where entities that share similar features and attributes are grouped together. So, using the Amazon.com example package, delivery, and schedule would be part of the shipping context, while price, discount, and currency are part of the pricing context. With the model divided into contexts, a template for how to boundary microservices emerges.



Level of risk exposed if this best practice is not established: High

Implementation guidance

- Design your workload based on your business domains and their respective functionality. Focusing on specific functionality enables you to differentiate the reliability requirements of different services, and target investments more specifically. A concise business problem and having a small team associated with each service also enables easier organizational scaling.
- Perform Domain Analysis to map out a domain-driven design (DDD) for your workload. Then you can choose an architecture type to meet your workload's needs.
 - [How to break a Monolith into Microservices](#)
 - [Getting Started with DDD when Surrounded by Legacy Systems](#)
 - [Eric Evans "Domain-Driven Design: Tackling Complexity in the Heart of Software"](#)
 - [Implementing Microservices on AWS](#)
- Decompose your services into smallest possible components. With microservices architecture you can separate your workload into components with the minimal functionality to enable organizational scaling and agility.
- Define the API for the workload and its design goals, limits, and any other considerations for use.
 - Define the API.
 - The API definition should allow for growth and additional parameters.
 - Define the designed availabilities.
 - Your API may have multiple design goals for different features.
 - Establish limits
 - Use testing to define the limits of your workload capabilities.

Resources

Related documents:

- [Amazon API Gateway: Configuring a REST API Using OpenAPI](#)
- [Bounded Context \(a central pattern in Domain-Driven Design\)](#)
- [Eric Evans "Domain-Driven Design: Tackling Complexity in the Heart of Software"](#)
- [Getting Started with DDD when Surrounded by Legacy Systems](#)
- [How to break a Monolith into Microservices](#)
- [Implementing Microservices on AWS](#)
- [Microservice Trade-Offs](#)
- [Microservices - a definition of this new architectural term](#)

- [Microservices on AWS](#)

REL03-BP03 Provide service contracts per API

Service contracts are documented agreements between teams on service integration and include a machine-readable API definition, rate limits, and performance expectations. A versioning strategy allows your clients to continue using the existing API and migrate their applications to the newer API when they are ready. Deployment can happen anytime, as long as the contract is not violated. The service provider team can use the technology stack of their choice to satisfy the API contract. Similarly, the service consumer can use their own technology.

Microservices take the concept of service-oriented architecture (SOA) to the point of creating services that have a minimal set of functionality. Each service publishes an API and design goals, limits, and other considerations for using the service. This establishes a *contract* with calling applications. This accomplishes three main benefits:

- The service has a concise business problem to be served and a small team that owns the business problem. This allows for better organizational scaling.
- The team can deploy at any time as long as they meet their API and other contract requirements.
- The team can use any technology stack they want to as long as they meet their API and other contract requirements.

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. It handles all the tasks involved in accepting and processing up to hundreds of thousands of concurrent API calls, including traffic management, authorization and access control, monitoring, and API version management. Using OpenAPI Specification (OAS), formerly known as the Swagger Specification, you can define your API contract and import it into API Gateway. With API Gateway, you can then version and deploy the APIs.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Provide service contracts per API Service contracts are documented agreements between teams on service integration and include a machine-readable API definition, rate limits, and performance expectations.
 - [Amazon API Gateway: Configuring a REST API Using OpenAPI](#)
 - A versioning strategy allows clients to continue using the existing API and migrate their applications to the newer API when they are ready.
 - Amazon API Gateway is a fully managed service that makes it easy for developers to create APIs at any scale. Using the OpenAPI Specification (OAS), formerly known as the Swagger Specification, you can define your API contract and import it into API Gateway. With API Gateway, you can then version and deploy the APIs.

Resources

Related documents:

- [Amazon API Gateway: Configuring a REST API Using OpenAPI](#)
- [Bounded Context \(a central pattern in Domain-Driven Design\)](#)
- [Implementing Microservices on AWS](#)
- [Microservice Trade-Offs](#)
- [Microservices - a definition of this new architectural term](#)
- [Microservices on AWS](#)

REL 4 How do you design interactions in a distributed system to prevent failures?

Distributed systems rely on communications networks to interconnect components, such as servers or services. Your workload must operate reliably despite data loss or latency in these networks. Components of the distributed system must operate in a way that does not negatively impact other components or the workload. These best practices prevent failures and improve mean time between failures (MTBF).

Best practices

- [REL04-BP01 Identify which kind of distributed system is required \(p. 189\)](#)
- [REL04-BP02 Implement loosely coupled dependencies \(p. 190\)](#)
- [REL04-BP03 Do constant work \(p. 192\)](#)
- [REL04-BP04 Make all responses idempotent \(p. 193\)](#)

REL04-BP01 Identify which kind of distributed system is required

Hard real-time distributed systems require responses to be given synchronously and rapidly, while soft real-time systems have a more generous time window of minutes or more for response. Offline systems handle responses through batch or asynchronous processing. Hard real-time distributed systems have the most stringent reliability requirements.

The most difficult [challenges with distributed systems](#) are for the hard real-time distributed systems, also known as request/reply services. What makes them difficult is that requests arrive unpredictably and responses must be given rapidly (for example, the customer is actively waiting for the response). Examples include front-end web servers, the order pipeline, credit card transactions, every AWS API, and telephony.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Identify which kind of distributed system is required. Challenges with distributed systems involved latency, scaling, understanding networking APIs, marshalling and unmarshalling data, and the complexity of algorithms such as Paxos. As the systems grow larger and more distributed, what had been theoretical edge cases turn into regular occurrences.
- [The Amazon Builders' Library: Challenges with distributed systems](#)
 - Hard real-time distributed systems require responses to be given synchronously and rapidly.
 - Soft real-time systems have a more generous time window of minutes or greater for response.
 - Offline systems handle responses through batch or asynchronous processing.
 - Hard real-time distributed systems have the most stringent reliability requirements.

Resources

Related documents:

- [Amazon EC2: Ensuring Idempotency](#)
- [The Amazon Builders' Library: Challenges with distributed systems](#)
- [The Amazon Builders' Library: Reliability, constant work, and a good cup of coffee](#)
- [What Is Amazon EventBridge?](#)
- [What Is Amazon Simple Queue Service?](#)

Related videos:

- [AWS New York Summit 2019: Intro to Event-driven Architectures and Amazon EventBridge \(MAD205\)](#)
- [AWS re:Invent 2018: Close Loops and Opening Minds: How to Take Control of Systems, Big and Small ARC337 \(includes loose coupling, constant work, static stability\)](#)
- [AWS re:Invent 2019: Moving to event-driven architectures \(SVS308\)](#)

REL04-BP02 Implement loosely coupled dependencies

Dependencies such as queuing systems, streaming systems, workflows, and load balancers are loosely coupled. Loose coupling helps isolate behavior of a component from other components that depend on it, increasing resiliency and agility.

If changes to one component force other components that rely on it to also change, then they are *tightly* coupled. *Loose* coupling breaks this dependency so that dependent components only need to know the versioned and published interface. Implementing loose coupling between dependencies isolates a failure in one from impacting another.

Loose coupling enables you to add additional code or features to a component while minimizing risk to components that depend on it. Also, scalability is improved as you can scale out or even change underlying implementation of the dependency.

To further improve resiliency through loose coupling, make component interactions asynchronous where possible. This model is suitable for any interaction that does not need an immediate response and where an acknowledgment that a request has been registered will suffice. It involves one component that generates events and another that consumes them. The two components do not integrate through direct point-to-point interaction but usually through an intermediate durable storage layer, such as an SQS queue or a streaming data platform such as Amazon Kinesis, or AWS Step Functions.




Figure 4: Dependencies such as queuing systems and load balancers are loosely coupled

Amazon SQS queues and Elastic Load Balancers are just two ways to add an intermediate layer for loose coupling. Event-driven architectures can also be built in the AWS Cloud using Amazon EventBridge, which can abstract clients (event producers) from the services they rely on (event consumers). Amazon Simple Notification Service (Amazon SNS) is an effective solution when you need high-throughput, push-based, many-to-many messaging. Using Amazon SNS topics, your publisher systems can fan out messages to a large number of subscriber endpoints for parallel processing.

While queues offer several advantages, in most hard real-time systems, requests older than a threshold time (often seconds) should be considered stale (the client has given up and is no longer waiting for a response), and not processed. This way more recent (and likely still valid requests) can be processed instead.

Common anti-patterns:

- Deploying a singleton as part of a workload.
- Directly invoking APIs between workload tiers with no capability of failover or asynchronous processing of the request.

Benefits of establishing this best practice: Loose coupling helps isolate behavior of a component from other components that depend on it, increasing resiliency and agility. Failure in one component is isolated from others.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Implement loosely coupled dependencies. Dependencies such as queuing systems, streaming systems, workflows, and load balancers are loosely coupled. Loose coupling helps isolate behavior of a component from other components that depend on it, increasing resiliency and agility.
 - [AWS re:Invent 2019: Moving to event-driven architectures \(SVS308\)](#)
 - [What Is Amazon EventBridge?](#)
 - [What Is Amazon Simple Queue Service?](#)
 - Amazon EventBridge allows you to build event driven architectures, which are loosely coupled and distributed.
 - [AWS New York Summit 2019: Intro to Event-driven Architectures and Amazon EventBridge \(MAD205\)](#)
 - If changes to one component force other components that rely on it to also change, then they are tightly coupled. Loose coupling breaks this dependency so that dependency components only need to know the versioned and published interface.
 - Make component interactions asynchronous where possible. This model is suitable for any interaction that does not need an immediate response and where an acknowledgement that a request has been registered will suffice.
 - [AWS re:Invent 2019: Scalable serverless event-driven applications using Amazon SQS and Lambda \(API304\)](#)

Resources

Related documents:

- [AWS re:Invent 2019: Moving to event-driven architectures \(SVS308\)](#)
- [Amazon EC2: Ensuring Idempotency](#)
- [The Amazon Builders' Library: Challenges with distributed systems](#)
- [The Amazon Builders' Library: Reliability, constant work, and a good cup of coffee](#)
- [What Is Amazon EventBridge?](#)
- [What Is Amazon Simple Queue Service?](#)

Related videos:

- [AWS New York Summit 2019: Intro to Event-driven Architectures and Amazon EventBridge \(MAD205\)](#)
- [AWS re:Invent 2018: Close Loops and Opening Minds: How to Take Control of Systems, Big and Small ARC337 \(includes loose coupling, constant work, static stability\)](#)
- [AWS re:Invent 2019: Moving to event-driven architectures \(SVS308\)](#)
- [AWS re:Invent 2019: Scalable serverless event-driven applications using Amazon SQS and Lambda \(API304\)](#)

REL04-BP03 Do constant work

Systems can fail when there are large, rapid changes in load. For example, if your workload is doing a health check that monitors the health of thousands of servers, it should send the same size payload (a full snapshot of the current state) each time. Whether no servers are failing, or all of them, the health check system is doing constant work with no large, rapid changes.

For example, if the health check system is monitoring 100,000 servers, the load on it is nominal under the normally light server failure rate. However, if a major event makes half of those servers unhealthy,

then the health check system would be overwhelmed trying to update notification systems and communicate state to its clients. So instead the health check system should send the full snapshot of the current state each time. 100,000 server health states, each represented by a bit, would only be a 12.5-KB payload. Whether no servers are failing, or all of them are, the health check system is doing constant work, and large, rapid changes are not a threat to the system stability. This is actually how Amazon Route 53 handles health checks for endpoints (such as IP addresses) to determine how end users are routed to them.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Do constant work so that systems do not fail when there are large, rapid changes in load.
- Implement loosely coupled dependencies. Dependencies such as queuing systems, streaming systems, workflows, and load balancers are loosely coupled. Loose coupling helps isolate behavior of a component from other components that depend on it, increasing resiliency and agility.
 - [The Amazon Builders' Library: Reliability, constant work, and a good cup of coffee](#)
 - [AWS re:Invent 2018: Close Loops and Opening Minds: How to Take Control of Systems, Big and Small ARC337 \(includes constant work\)](#)
 - For the example of a health check system monitoring 100,000 servers, engineer workloads so that payload sizes remain constant regardless of number of successes or failures.

Resources

Related documents:

- [Amazon EC2: Ensuring Idempotency](#)
- [The Amazon Builders' Library: Challenges with distributed systems](#)
- [The Amazon Builders' Library: Reliability, constant work, and a good cup of coffee](#)

Related videos:

- [AWS New York Summit 2019: Intro to Event-driven Architectures and Amazon EventBridge \(MAD205\)](#)
- [AWS re:Invent 2018: Close Loops and Opening Minds: How to Take Control of Systems, Big and Small ARC337 \(includes constant work\)](#)
- [AWS re:Invent 2018: Close Loops and Opening Minds: How to Take Control of Systems, Big and Small ARC337 \(includes loose coupling, constant work, static stability\)](#)
- [AWS re:Invent 2019: Moving to event-driven architectures \(SVS308\)](#)

REL04-BP04 Make all responses idempotent

An idempotent service promises that each request is completed exactly once, such that making multiple identical requests has the same effect as making a single request. An idempotent service makes it easier for a client to implement retries without fear that a request will be erroneously processed multiple times. To do this, clients can issue API requests with an idempotency token—the same token is used whenever the request is repeated. An idempotent service API uses the token to return a response identical to the response that was returned the first time that the request was completed.

In a distributed system, it's easy to perform an action at most once (client makes only one request), or at least once (keep requesting until client gets confirmation of success). But it's hard to guarantee an action is idempotent, which means it's performed *exactly* once, such that making multiple identical requests has the same effect as making a single request. Using idempotency tokens in APIs, services can receive a mutating request one or more times without creating duplicate records or side effects.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Make all responses idempotent. An idempotent service promises that each request is completed exactly once, such that making multiple identical requests has the same effect as making a single request.
- Clients can issue API requests with an idempotency token—the same token is used whenever the request is repeated. An idempotent service API uses the token to return a response identical to the response that was returned the first time that the request was completed.
 - [Amazon EC2: Ensuring Idempotency](#)

Resources

Related documents:

- [Amazon EC2: Ensuring Idempotency](#)
- [The Amazon Builders' Library: Challenges with distributed systems](#)
- [The Amazon Builders' Library: Reliability, constant work, and a good cup of coffee](#)

Related videos:

- [AWS New York Summit 2019: Intro to Event-driven Architectures and Amazon EventBridge \(MAD205\)](#)
- [AWS re:Invent 2018: Close Loops and Opening Minds: How to Take Control of Systems, Big and Small ARC337 \(includes loose coupling, constant work, static stability\)](#)
- [AWS re:Invent 2019: Moving to event-driven architectures \(SVS308\)](#)

REL 5 How do you design interactions in a distributed system to mitigate or withstand failures?

Distributed systems rely on communications networks to interconnect components (such as servers or services). Your workload must operate reliably despite data loss or latency over these networks. Components of the distributed system must operate in a way that does not negatively impact other components or the workload. These best practices enable workloads to withstand stresses or failures, more quickly recover from them, and mitigate the impact of such impairments. The result is improved mean time to recovery (MTTR).

Best practices

- [REL05-BP01 Implement graceful degradation to transform applicable hard dependencies into soft dependencies \(p. 194\)](#)
- [REL05-BP02 Throttle requests \(p. 197\)](#)
- [REL05-BP03 Control and limit retry calls \(p. 198\)](#)
- [REL05-BP04 Fail fast and limit queues \(p. 199\)](#)
- [REL05-BP05 Set client timeouts \(p. 199\)](#)
- [REL05-BP06 Make services stateless where possible \(p. 200\)](#)
- [REL05-BP07 Implement emergency levers \(p. 202\)](#)

REL05-BP01 Implement graceful degradation to transform applicable hard dependencies into soft dependencies

When a component's dependencies are unhealthy, the component itself can still function, although in a degraded manner. For example, when a dependency call fails, failover to a predetermined static response.

Consider a service B that is called by service A and in turn calls service C.



Figure 5: Service C fails when called from service B. Service B returns a degraded response to service A.

When service B calls service C, it received an error or timeout from it. Service B, lacking a response from service C (and the data it contains) instead returns what it can. This can be the last cached good value, or service B can substitute a pre-determined static response for what it would have received from service C. It can then return a degraded response to its caller, service A. Without this static response, the failure in service C would cascade through service B to service A, resulting in a loss of availability.

As per the multiplicative factor in the availability equation for hard dependencies (see [Calculating availability with hard dependencies](#)), any drop in the availability of C seriously impacts effective availability of B. By returning the static response, service B mitigates the failure in C and, although degraded, makes service C's availability look like 100% availability (assuming it reliably returns the static response under error conditions). Note that the static response is a simple alternative to returning an error, and is not an attempt to re-compute the response using different means. Such attempts at a completely different mechanism to try to achieve the same result are called fallback behavior, and are an anti-pattern to be avoided.

Another example of graceful degradation is the *circuit breaker pattern*. Retry strategies should be used when the failure is transient. When this is not the case, and the operation is likely to fail, the circuit breaker pattern prevents the client from performing a request that is likely to fail. When requests are being processed normally, the circuit breaker is closed and requests flow through. When the remote system begins returning errors or exhibits high latency, the circuit breaker opens and the dependency is ignored or results are replaced with more simply obtained but less comprehensive responses (which might simply be a response cache). Periodically, the system attempts to call the dependency to determine if it has recovered. When that occurs, the circuit breaker is closed.




Figure 6: Circuit breaker showing closed and open states.

In addition to the closed and open states shown in the diagram, after a configurable period of time in the open state, the circuit breaker can transition to half-open. In this state, it periodically attempts to call the service at a much lower rate than normal. This probe is used to check the health of the service. After a number of successes in half-open state, the circuit breaker transitions to closed, and normal requests resume.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Implement graceful degradation to transform applicable hard dependencies into soft dependencies. When a component's dependencies are unhealthy, the component itself can still function, although in a degraded manner. For example, when a dependency call fails, failover to a predetermined static response.
- By returning a static response, your workload mitigates failures that occur in its dependencies.
 - [Well-Architected lab: Level 300: Implementing Health Checks and Managing Dependencies to Improve Reliability](#)
- Detect when the retry operation is likely to fail, and prevent your client from making failed calls with the circuit breaker pattern.
 - [CircuitBreaker](#)

Resources

Related documents:

- [Amazon API Gateway: Throttle API Requests for Better Throughput](#)
- [CircuitBreaker \(summarizes Circuit Breaker from "Release It!" book\)](#)
- [Error Retries and Exponential Backoff in AWS](#)
- [Michael Nygard "Release It! Design and Deploy Production-Ready Software"](#)

- [The Amazon Builders' Library: Avoiding fallback in distributed systems](#)
- [The Amazon Builders' Library: Avoiding insurmountable queue backlogs](#)
- [The Amazon Builders' Library: Caching challenges and strategies](#)
- [The Amazon Builders' Library: Timeouts, retries, and backoff with jitter](#)

Related videos:

- [Retry, backoff, and jitter: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)

Related examples:

- [Well-Architected lab: Level 300: Implementing Health Checks and Managing Dependencies to Improve Reliability](#)

REL05-BP02 Throttle requests

Throttling requests is a mitigation pattern to respond to an unexpected increase in demand. Some requests are honored but those over a defined limit are rejected and return a message indicating they have been throttled. The expectation on clients is that they will back off and abandon the request or try again at a slower rate.

Your services should be designed to handle a known capacity of requests that each node or cell can process. This capacity can be established through load testing. You then need to track the arrival rate of requests and if the temporary arrival rate exceeds this limit, the appropriate response is to signal that the request has been throttled. This allows the user to retry, potentially to a different node or cell that might have available capacity. Amazon API Gateway provides methods for throttling requests. Amazon SQS and Amazon Kinesis can buffer requests, smooth out the request rate, and alleviate the need for throttling for requests that can be addressed asynchronously.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Throttle requests. This is a mitigation pattern to respond to an unexpected increase in demand. Some requests are honored but those over a defined limit are rejected and return a message indicating they have been throttled. The expectation on clients is that they will back off and abandon the request or try again at a slower rate.
 - Use Amazon API Gateway
 - [Throttle API Requests for Better Throughput](#)

Resources

Related documents:

- [Amazon API Gateway: Throttle API Requests for Better Throughput](#)
- [Error Retries and Exponential Backoff in AWS](#)
- [The Amazon Builders' Library: Avoiding fallback in distributed systems](#)
- [The Amazon Builders' Library: Avoiding insurmountable queue backlogs](#)
- [The Amazon Builders' Library: Timeouts, retries, and backoff with jitter](#)
- [Throttle API Requests for Better Throughput](#)

Related videos:

- [Retry, backoff, and jitter: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)

REL05-BP03 Control and limit retry calls

Use exponential backoff to retry after progressively longer intervals. Introduce jitter to randomize those retry intervals, and limit the maximum number of retries.

Typical components in a distributed software system include servers, load balancers, databases, and DNS servers. In operation, and subject to failures, any of these can start generating errors. The default technique for dealing with errors is to implement retries on the client side. This technique increases the reliability and availability of the application. However, at scale—and if clients attempt to retry the failed operation as soon as an error occurs—the network can quickly become saturated with new and retried requests, each competing for network bandwidth. This can result in a *retry storm*, which will reduce availability of the service. This pattern might continue until a full system failure occurs.

To avoid such scenarios, backoff algorithms such as the common *exponential backoff* should be used. Exponential backoff algorithms gradually decrease the rate at which retries are performed, thus avoiding network congestion.

Many SDKs and software libraries, including those from AWS, implement a version of these algorithms. However, **never assume a backoff algorithm exists—always test and verify this to be the case.**

Simple backoff alone is not enough because in distributed systems all clients may backoff simultaneously, creating clusters of retry calls. Marc Brooker in his blog post [Exponential Backoff and Jitter](#), explains how to modify the `wait()` function in the exponential backoff to prevent clusters of retry calls. The solution is to add *jitter* in the `wait()` function. To avoid retrying for too long, implementations should cap the backoff to a maximum value.

Finally, it's important to configure a *maximum number of retries* or elapsed time, after which retrying will simply fail. AWS SDKs implement this by default, and it can be configured. For services lower in the stack, a maximum retry limit of zero or one can limit risk yet still be effective as retries are delegated to services higher in the stack.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Control and limit retry calls. Use exponential backoff to retry after progressively longer intervals. Introduce jitter to randomize those retry intervals, and limit the maximum number of retries.
 - [Error Retries and Exponential Backoff in AWS](#)
 - Amazon SDKs implement retries and exponential backoff by default. Implement similar logic in your dependency layer when calling your own dependent services. Decide what the timeouts are and when to stop retrying based on your use case.

Resources

Related documents:

- [Amazon API Gateway: Throttle API Requests for Better Throughput](#)
- [Error Retries and Exponential Backoff in AWS](#)
- [The Amazon Builders' Library: Avoiding fallback in distributed systems](#)
- [The Amazon Builders' Library: Avoiding insurmountable queue backlogs](#)
- [The Amazon Builders' Library: Caching challenges and strategies](#)
- [The Amazon Builders' Library: Timeouts, retries, and backoff with jitter](#)

Related videos:

- [Retry, backoff, and jitter: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)

REL05-BP04 Fail fast and limit queues

If the workload is unable to respond successfully to a request, then fail fast. This allows the releasing of resources associated with a request, and permits the service to recover if it's running out of resources. If the workload is able to respond successfully but the rate of requests is too high, then use a queue to buffer requests instead. However, do not allow long queues that can result in serving stale requests that the client has already given up on.

This best practice applies to the server-side, or receiver, of the request.

Be aware that queues can be created at multiple levels of a system, and can seriously impede the ability to quickly recover as older, stale requests (that no longer need a response) are processed before newer requests. Be aware of places where queues exist. They often hide in workflows or in work that's recorded to a database.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Fail fast and limit queues. If the workload is unable to respond successfully to a request, then fail fast. This allows the releasing of resources associated with a request, and permits the service to recover if it's running out of resources. If the workload is able to respond successfully but the rate of requests is too high, then use a queue to buffer requests instead. However, do not allow long queues that can result in serving stale requests that the client has already given up on.
 - Implement fail fast when service is under stress.
 - [Fail Fast](#)
 - Limit queues In a queue-based system, when processing stops but messages keep arriving, the message debt can accumulate into a large backlog, driving up processing time. Work can be completed too late for the results to be useful, essentially causing the availability hit that queueing was meant to guard against.
 - [The Amazon Builders' Library: Avoiding insurmountable queue backlogs](#)

Resources

Related documents:

- [Error Retries and Exponential Backoff in AWS](#)
- [Fail Fast](#)
- [The Amazon Builders' Library: Avoiding fallback in distributed systems](#)
- [The Amazon Builders' Library: Avoiding insurmountable queue backlogs](#)
- [The Amazon Builders' Library: Caching challenges and strategies](#)
- [The Amazon Builders' Library: Timeouts, retries, and backoff with jitter](#)

Related videos:

- [Retry, backoff, and jitter: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)

REL05-BP05 Set client timeouts

Set timeouts appropriately, verify them systematically, and do not rely on default values as they are generally set too high.

This best practice applies to the client-side, or sender, of the request.

Set both a connection timeout and a request timeout on any remote call, and generally on any call across processes. Many frameworks offer built-in timeout capabilities, but be careful as many have default values that are infinite or too high. A value that is too high reduces the usefulness of the timeout because resources continue to be consumed while the client waits for the timeout to occur. A too low value can generate increased traffic on the backend and increased latency because too many requests are retried. In some cases, this can lead to complete outages because all requests are being retried.

To learn more about how Amazon use timeouts, retries, and backoff with jitter, refer to the [Builder's Library: Timeouts, retries, and backoff with jitter](#).

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Set both a connection timeout and a request timeout on any remote call, and generally on any call across processes. Many frameworks offer built-in timeout capabilities, but be careful as many have default values that are infinite or too high. A value that is too high reduces the usefulness of the timeout because resources continue to be consumed while the client waits for the timeout to occur. A too low value can generate increased traffic on the backend and increased latency because too many requests are retried. In some cases, this can lead to complete outages because all requests are being retried.
- [AWS SDK: Retries and Timeouts](#)

Resources

Related documents:

- [AWS SDK: Retries and Timeouts](#)
- [Amazon API Gateway: Throttle API Requests for Better Throughput](#)
- [Error Retries and Exponential Backoff in AWS](#)
- [The Amazon Builders' Library: Timeouts, retries, and backoff with jitter](#)

Related videos:

- [Retry, backoff, and jitter: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)

REL05-BP06 Make services stateless where possible

Services should either not require state, or should offload state such that between different client requests, there is no dependence on locally stored data on disk and in memory. This enables servers to be replaced at will without causing an availability impact. Amazon ElastiCache or Amazon DynamoDB are good destinations for offloaded state.




Figure 7: In this stateless web application, session state is offloaded to Amazon ElastiCache.

When users or services interact with an application, they often perform a series of interactions that form a session. A session is unique data for users that persists between requests while they use the application. A stateless application is an application that does not need knowledge of previous interactions and does not store session information.

Once designed to be stateless, you can then use serverless compute services, such as AWS Lambda or AWS Fargate.

In addition to server replacement, another benefit of stateless applications is that they can scale horizontally because any of the available compute resources (such as EC2 instances and AWS Lambda functions) can service any request.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Make your applications stateless. Stateless applications enable horizontal scaling and are tolerant to the failure of an individual node.
- Remove state that could actually be stored in request parameters.

- After examining whether the state is required, move any state tracking to a resilient multi-zone cache or data store like Amazon ElastiCache, Amazon RDS, Amazon DynamoDB, or a third-party distributed data solution. Store a state that could not be moved to resilient data stores.
- Some data (like cookies) can be passed in headers or query parameters.
- Refactor to remove state that can be quickly passed in requests.
- Some data may not actually be needed per request and can be retrieved on demand.
- Remove data that can be asynchronously retrieved.
- Decide on a data store that meets the requirements for a required state.
- Consider a NoSQL database for non-relational data.

Resources

Related documents:

- [The Amazon Builders' Library: Avoiding fallback in distributed systems](#)
- [The Amazon Builders' Library: Avoiding insurmountable queue backlogs](#)
- [The Amazon Builders' Library: Caching challenges and strategies](#)

REL05-BP07 Implement emergency levers

Emergency levers are rapid processes that can mitigate availability impact on your workload.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Implement emergency levers. These are rapid processes that may mitigate availability impact on your workload. They can be operated in the absence of a root cause. An ideal emergency lever reduces the cognitive burden on the resolvers to zero by providing fully deterministic activation and deactivation criteria. Levers are often manual, but they can also be automated
- Example levers include
 - Block all robot traffic
 - Serve static pages instead of dynamic ones
 - Reduce frequency of calls to a dependency
 - Throttle calls from dependencies
- Tips for implementing and using emergency levers
 - When levers are activated, do LESS, not more
 - Keep it simple, avoid bimodal behavior
 - Test your levers periodically
- These are examples of actions that are NOT emergency levers
 - Add capacity
 - Call up service owners of clients that depend on your service and ask them to reduce calls
 - Making a change to code and releasing it

Change management

Questions

- [REL 6 How do you monitor workload resources? \(p. 203\)](#)
- [REL 7 How do you design your workload to adapt to changes in demand? \(p. 210\)](#)

- [REL 8 How do you implement change? \(p. 216\)](#)

REL 6 How do you monitor workload resources?

Logs and metrics are powerful tools to gain insight into the health of your workload. You can configure your workload to monitor logs and metrics and send notifications when thresholds are crossed or significant events occur. Monitoring enables your workload to recognize when low-performance thresholds are crossed or failures occur, so it can recover automatically in response.

Best practices

- [REL06-BP01 Monitor all components for the workload \(Generation\) \(p. 203\)](#)
- [REL06-BP02 Define and calculate metrics \(Aggregation\) \(p. 205\)](#)
- [REL06-BP03 Send notifications \(Real-time processing and alarming\) \(p. 206\)](#)
- [REL06-BP04 Automate responses \(Real-time processing and alarming\) \(p. 207\)](#)
- [REL06-BP05 Analytics \(p. 208\)](#)
- [REL06-BP06 Conduct reviews regularly \(p. 209\)](#)
- [REL06-BP07 Monitor end-to-end tracing of requests through your system \(p. 210\)](#)

REL06-BP01 Monitor all components for the workload (Generation)

Monitor the components of the workload with Amazon CloudWatch or third-party tools. Monitor AWS services with AWS Health Dashboard.

All components of your workload should be monitored, including the front-end, business logic, and storage tiers. Define key metrics, describe how to extract them from logs (if necessary), and set thresholds for triggering corresponding alarm events. Ensure metrics are relevant to the key performance indicators (KPIs) of your workload, and use metrics and logs to identify early warning signs of service degradation. For example, a metric related to business outcomes such as the number of orders successfully processed per minute, can indicate workload issues faster than technical metric, such as CPU Utilization. Use AWS Health Dashboard for a personalized view into the performance and availability of the AWS services underlying your AWS resources.

Monitoring in the cloud offers new opportunities. Most cloud providers have developed customizable hooks and can deliver insights to help you monitor multiple layers of your workload. AWS services such as Amazon CloudWatch apply statistical and machine learning algorithms to continually analyze metrics of systems and applications, determine normal baselines, and surface anomalies with minimal user intervention. Anomaly detection algorithms account for the seasonality and trend changes of metrics.

AWS makes an abundance of monitoring and log information available for consumption that can be used to define workload-specific metrics, change-in-demand processes, and adopt machine learning techniques regardless of ML expertise.

In addition, monitor all of your external endpoints to ensure that they are independent of your base implementation. This active monitoring can be done with synthetic transactions (sometimes referred to as *user canaries*, but not to be confused with canary deployments) which periodically run a number of common tasks matching actions performed by clients of the workload. Keep these tasks short in duration and be sure not to overload your workload during testing. Amazon CloudWatch Synthetics enables you to [create synthetic canaries](#) to monitor your endpoints and APIs. You can also combine the synthetic canary client nodes with AWS X-Ray console to pinpoint which synthetic canaries are experiencing issues with errors, faults, or throttling rates for the selected time frame.

Desired Outcome:

Collect and use critical metrics from all components of the workload to ensure workload reliability and optimal user experience. Detecting that a workload is not achieving business outcomes allows you to quickly declare a disaster and recover from an incident.

Common anti-patterns:

- Only monitoring external interfaces to your workload.
- Not generating any workload-specific metrics and only relying on metrics provided to you by the AWS services your workload uses.
- Only using technical metrics in your workload and not monitoring any metrics related to non-technical KPIs the workload contributes to.
- Relying on production traffic and simple health checks to monitor and evaluate workload state.

Benefits of establishing this best practice: Monitoring at all tiers in your workload enables you to more rapidly anticipate and resolve problems in the components that comprise the workload.

Level of risk exposed if this best practice is not established: High

Implementation guidance

1. **Enable logging where available.** Monitoring data should be obtained from all components of the workloads. Turn on additional logging, such as S3 Access Logs, and enable your workload to log workload specific data. Collect metrics for CPU, network I/O, and disk I/O averages from services such as Amazon ECS, Amazon EKS, Amazon EC2, Elastic Load Balancing, AWS Auto Scaling, and Amazon EMR. See [AWS Services That Publish CloudWatch Metrics](#) for a list of AWS services that publish metrics to CloudWatch.
2. **Review all default metrics and explore any data collection gaps.** Every service generates default metrics. Collecting default metrics allows you to better understand the dependencies between workload components, and how component reliability and performance affect the workload. You can also create and [publish your own metrics](#) to CloudWatch using the AWS CLI or an API.
3. **Evaluate all the metrics to decide which ones to alert on for each AWS service in your workload.** You may choose to select a subset of metrics that have a major impact on workload reliability. Focusing on critical metrics and threshold allows you to refine the number of [alerts](#) and can help minimize false-positives.
4. **Define alerts and the recovery process for your workload after the alert is triggered.** Defining alerts allows you to quickly notify, escalate, and follow steps necessary to recover from an incident and meet your prescribed Recovery Time Objective (RTO). You can use [Amazon CloudWatch Alarms](#) to invoke automated workflows and initiate recovery procedures based on defined thresholds.
5. **Explore use of synthetic transactions to collect relevant data about workloads state.** Synthetic monitoring follows the same routes and perform the same actions as a customer, which makes it possible for you to continually verify your customer experience even when you don't have any customer traffic on your workloads. By using [synthetic transactions](#), you can discover issues before your customers do.

Resources

Related best practices:

- [REL11-BP03 Automate healing on all layers \(p. 244\)](#)

Related documents:

- [Getting started with your AWS Health Dashboard – Your account health](#)
- [AWS Services That Publish CloudWatch Metrics](#)
- [Access Logs for Your Network Load Balancer](#)
- [Access logs for your application load balancer](#)
- [Accessing Amazon CloudWatch Logs for AWS Lambda](#)

- [Amazon S3 Server Access Logging](#)
- [Enable Access Logs for Your Classic Load Balancer](#)
- [Exporting log data to Amazon S3](#)
- [Install the CloudWatch agent on an Amazon EC2 instance](#)
- [Publishing Custom Metrics](#)
- [Using Amazon CloudWatch Dashboards](#)
- [Using Amazon CloudWatch Metrics](#)
- [Using Canaries \(Amazon CloudWatch Synthetics\)](#)
- [What are Amazon CloudWatch Logs?](#)

User guides:

- [Creating a trail](#)
- [Monitoring memory and disk metrics for Amazon EC2 Linux instances](#)
- [Using CloudWatch Logs with container instances](#)
- [VPC Flow Logs](#)
- [What is Amazon DevOps Guru?](#)
- [What is AWS X-Ray?](#)

Related blogs:

- [Debugging with Amazon CloudWatch Synthetics and AWS X-Ray](#)

Related examples and workshops:

- [AWS Well-Architected Labs: Operational Excellence - Dependency Monitoring](#)
- [The Amazon Builders' Library: Instrumenting distributed systems for operational visibility](#)
- [Observability workshop](#)

REL06-BP02 Define and calculate metrics (Aggregation)

Store log data and apply filters where necessary to calculate metrics, such as counts of a specific log event, or latency calculated from log event timestamps.

Amazon CloudWatch and Amazon S3 serve as the primary aggregation and storage layers. For some services, such as AWS Auto Scaling and Elastic Load Balancing, default metrics are provided by default for CPU load or average request latency across a cluster or instance. For streaming services, such as VPC Flow Logs and AWS CloudTrail, event data is forwarded to CloudWatch Logs and you need to define and apply metrics filters to extract metrics from the event data. This gives you time series data, which can serve as inputs to CloudWatch alarms that you define to trigger alerts.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Define and calculate metrics (Aggregation). Store log data and apply filters where necessary to calculate metrics, such as counts of a specific log event, or latency calculated from log event timestamps
 - Metric filters define the terms and patterns to look for in log data as it is sent to CloudWatch Logs. CloudWatch Logs uses these metric filters to turn log data into numerical CloudWatch metrics that you can graph or set an alarm on.
 - [Searching and Filtering Log Data](#)

- Use a trusted third party to aggregate logs.
 - Follow the instructions of the third party. Most third-party products integrate with CloudWatch and Amazon S3.
- Some AWS services can publish logs directly to Amazon S3. If your main requirement for logs is storage in Amazon S3, you can easily have the service producing the logs send them directly to Amazon S3 without setting up additional infrastructure.
 - [Sending Logs Directly to Amazon S3](#)

Resources

Related documents:

- [Amazon CloudWatch Logs Insights Sample Queries](#)
- [Debugging with Amazon CloudWatch Synthetics and AWS X-Ray](#)
- [One Observability Workshop](#)
- [Searching and Filtering Log Data](#)
- [Sending Logs Directly to Amazon S3](#)
- [The Amazon Builders' Library: Instrumenting distributed systems for operational visibility](#)

REL06-BP03 Send notifications (Real-time processing and alarming)

Organizations that need to know, receive notifications when significant events occur.

Alerts can be sent to Amazon Simple Notification Service (Amazon SNS) topics, and then pushed to any number of subscribers. For example, Amazon SNS can forward alerts to an email alias so that technical staff can respond.

Common anti-patterns:

- Configuring alarms at too low of threshold, causing too many notifications to be sent.
- Not archiving alarms for future exploration.

Benefits of establishing this best practice: Notifications on events (even those that can be responded to and automatically resolved) allow you to have a record of events and potentially address them in a different manner in the future.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Perform real-time processing and alarming. Organizations that need to know, receive notifications when significant events occur
 - Amazon CloudWatch dashboards are customizable home pages in the CloudWatch console that you can use to monitor your resources in a single view, even those resources that are spread across different Regions.
 - [Using Amazon CloudWatch Dashboards](#)
 - Create an alarm when the metric surpasses a limit.
 - [Using Amazon CloudWatch Alarms](#)

Resources

Related documents:

- [One Observability Workshop](#)
- [The Amazon Builders' Library: Instrumenting distributed systems for operational visibility](#)
- [Using Amazon CloudWatch Alarms](#)
- [Using Amazon CloudWatch Dashboards](#)
- [Using Amazon CloudWatch Metrics](#)

REL06-BP04 Automate responses (Real-time processing and alarming)

Use automation to take action when an event is detected, for example, to replace failed components.

Alerts can trigger AWS Auto Scaling events, so that clusters react to changes in demand. Alerts can be sent to Amazon Simple Queue Service (Amazon SQS), which can serve as an integration point for third-party ticket systems. AWS Lambda can also subscribe to alerts, providing users an asynchronous serverless model that reacts to change dynamically. AWS Config continually monitors and records your AWS resource configurations, and can trigger [AWS Systems Manager Automation](#) to remediate issues.

Amazon DevOps Guru can automatically monitor application resources for anomalous behavior and deliver targeted recommendations to speed up problem identification and remediation times.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use Amazon DevOps Guru to perform automated actions. Amazon DevOps Guru can automatically monitor application resources for anomalous behavior and deliver targeted recommendations to speed up problem identification and remediation times.
 - [What is Amazon DevOps Guru?](#)
- Use AWS Systems Manager to perform automated actions. AWS Config continually monitors and records your AWS resource configurations, and can trigger AWS Systems Manager Automation to remediate issues.
 - [AWS Systems Manager Automation](#)
 - Create and use Systems Manager Automation documents. These define the actions that Systems Manager performs on your managed instances and other AWS resources when an automation process runs.
 - [Working with Automation Documents \(Playbooks\)](#)
- Amazon CloudWatch sends alarm state change events to Amazon EventBridge. Create EventBridge rules to automate responses.
 - [Creating an EventBridge Rule That Triggers on an Event from an AWS Resource](#)
- Create and execute a plan to automate responses.
 - Inventory all your alert response procedures. You must plan your alert responses before you rank the tasks.
 - Inventory all the tasks with specific actions that must be taken. Most of these actions are documented in runbooks. You must also have playbooks for alerts of unexpected events.
 - Examine the runbooks and playbooks for all automatable actions. In general, if an action can be defined, it most likely can be automated.
 - Rank the error-prone or time-consuming activities first. It is most beneficial to remove sources of errors and reduce time to resolution.
 - Establish a plan to complete automation. Maintain an active plan to automate and update the automation.
 - Examine manual requirements for opportunities for automation. Challenge your manual process for opportunities to automate.

Resources

Related documents:

- [AWS Systems Manager Automation](#)
- [Creating an EventBridge Rule That Triggers on an Event from an AWS Resource](#)
- [One Observability Workshop](#)
- [The Amazon Builders' Library: Instrumenting distributed systems for operational visibility](#)
- [What is Amazon DevOps Guru?](#)
- [Working with Automation Documents \(Playbooks\)](#)

REL06-BP05 Analytics

Collect log files and metrics histories and analyze these for broader trends and workload insights.

Amazon CloudWatch Logs Insights supports a [simple yet powerful query language](#) that you can use to analyze log data. Amazon CloudWatch Logs also supports subscriptions that allow data to flow seamlessly to Amazon S3 where you can use or Amazon Athena to query the data. It also supports queries on a large array of formats. See [Supported SerDes and Data Formats](#) in the Amazon Athena User Guide for more information. For analysis of huge log file sets, you can run an Amazon EMR cluster to run petabyte-scale analyses.

There are a number of tools provided by AWS Partners and third parties that allow for aggregation, processing, storage, and analytics. These tools include New Relic, Splunk, Loggly, Logstash, CloudHealth, and Nagios. However, outside generation of system and application logs is unique to each cloud provider, and often unique to each service.

An often-overlooked part of the monitoring process is data management. You need to determine the retention requirements for monitoring data, and then apply lifecycle policies accordingly. Amazon S3 supports lifecycle management at the S3 bucket level. This lifecycle management can be applied differently to different paths in the bucket. Toward the end of the lifecycle, you can transition data to Amazon S3 Glacier for long-term storage, and then expiration after the end of the retention period is reached. The S3 Intelligent-Tiering storage class is designed to optimize costs by automatically moving data to the most cost-effective access tier, without performance impact or operational overhead.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- CloudWatch Logs Insights enables you to interactively search and analyze your log data in Amazon CloudWatch Logs.
 - [Analyzing Log Data with CloudWatch Logs Insights](#)
 - [Amazon CloudWatch Logs Insights Sample Queries](#)
- Use Amazon CloudWatch Logs send logs to Amazon S3 where you can use or Amazon Athena to query the data.
 - [How do I analyze my Amazon S3 server access logs using Athena?](#)
 - Create an S3 lifecycle policy for your server access logs bucket. Configure the lifecycle policy to periodically remove log files. Doing so reduces the amount of data that Athena analyzes for each query.
 - [How Do I Create a Lifecycle Policy for an S3 Bucket?](#)

Resources

Related documents:

- [Amazon CloudWatch Logs Insights Sample Queries](#)
- [Analyzing Log Data with CloudWatch Logs Insights](#)
- [Debugging with Amazon CloudWatch Synthetics and AWS X-Ray](#)
- [How Do I Create a Lifecycle Policy for an S3 Bucket?](#)
- [How do I analyze my Amazon S3 server access logs using Athena?](#)
- [One Observability Workshop](#)
- [The Amazon Builders' Library: Instrumenting distributed systems for operational visibility](#)

REL06-BP06 Conduct reviews regularly

Frequently review how workload monitoring is implemented and update it based on significant events and changes.

Effective monitoring is driven by key business metrics. Ensure these metrics are accommodated in your workload as business priorities change.

Auditing your monitoring helps ensure that you know when an application is meeting its availability goals. Root cause analysis requires the ability to discover what happened when failures occur. AWS provides services that allow you to track the state of your services during an incident:

- **Amazon CloudWatch Logs:** You can store your logs in this service and inspect their contents.
- **Amazon CloudWatch Logs Insights:** Is a fully managed service that enables you to analyze massive logs in seconds. It gives you fast, interactive queries and visualizations.
- **AWS Config:** You can see what AWS infrastructure was in use at different points in time.
- **AWS CloudTrail:** You can see which AWS APIs were invoked at what time and by what principal.

At AWS, we conduct a weekly meeting to [review operational performance](#) and to share learnings between teams. Because there are so many teams in AWS, we created [The Wheel](#) to randomly pick a workload to review. Establishing a regular cadence for operational performance reviews and knowledge sharing enhances your ability to achieve higher performance from your operational teams.

Common anti-patterns:

- Collecting only default metrics.
- Setting a monitoring strategy and never reviewing it.
- Not discussing monitoring when major changes are deployed.

Benefits of establishing this best practice: Regularly reviewing your monitoring enables the anticipation of potential problems, instead of reacting to notifications when an anticipated problem actually occurs.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Create multiple dashboards for the workload. You must have a top-level dashboard that contains the key business metrics, as well as the technical metrics you have identified to be the most relevant to the projected health of the workload as usage varies. You should also have dashboards for various application tiers and dependencies that can be inspected.
 - [Using Amazon CloudWatch Dashboards](#)
- Schedule and conduct regular reviews of the workload dashboards. Conduct regular inspection of the dashboards. You may have different cadences for the depth at which you inspect.
 - Inspect for trends in the metrics. Compare the metric values to historic values to see if there are trends that may indicate that something that needs investigation. Examples of this include: increasing latency, decreasing primary business function, and increasing failure responses.

- Inspect for outliers/anomalies in your metrics. Averages or medians can mask outliers and anomalies. Look at the highest and lowest values during the time frame and investigate the causes of extreme scores. As you continue to eliminate these causes, lowering your definition of extreme allows you to continue to improve the consistency of your workload performance.
- Look for sharp changes in behavior. An immediate change in quantity or direction of a metric may indicate that there has been a change in the application, or external factors that you may need to add additional metrics to track.

Resources

Related documents:

- [Amazon CloudWatch Logs Insights Sample Queries](#)
- [Debugging with Amazon CloudWatch Synthetics and AWS X-Ray](#)
- [One Observability Workshop](#)
- [The Amazon Builders' Library: Instrumenting distributed systems for operational visibility](#)
- [Using Amazon CloudWatch Dashboards](#)

REL06-BP07 Monitor end-to-end tracing of requests through your system

Use AWS X-Ray or third-party tools so that developers can more easily analyze and debug distributed systems to understand how their applications and its underlying services are performing.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Monitor end-to-end tracing of requests through your system. AWS X-Ray is a service that collects data about requests that your application serves, and provides tools you can use to view, filter, and gain insights into that data to identify issues and opportunities for optimization. For any traced request to your application, you can see detailed information not only about the request and response, but also about calls that your application makes to downstream AWS resources, microservices, databases, and web APIs.
 - [What is AWS X-Ray?](#)
 - [Debugging with Amazon CloudWatch Synthetics and AWS X-Ray](#)

Resources

Related documents:

- [Debugging with Amazon CloudWatch Synthetics and AWS X-Ray](#)
- [One Observability Workshop](#)
- [The Amazon Builders' Library: Instrumenting distributed systems for operational visibility](#)
- [Using Canaries \(Amazon CloudWatch Synthetics\)](#)
- [What is AWS X-Ray?](#)

REL 7 How do you design your workload to adapt to changes in demand?

A scalable workload provides elasticity to add or remove resources automatically so that they closely match the current demand at any given point in time.

Best practices

- [REL07-BP01 Use automation when obtaining or scaling resources \(p. 211\)](#)
- [REL07-BP02 Obtain resources upon detection of impairment to a workload \(p. 213\)](#)
- [REL07-BP03 Obtain resources upon detection that more resources are needed for a workload \(p. 214\)](#)
- [REL07-BP04 Load test your workload \(p. 215\)](#)

REL07-BP01 Use automation when obtaining or scaling resources

When replacing impaired resources or scaling your workload, automate the process by using managed AWS services, such as Amazon S3 and AWS Auto Scaling. You can also use third-party tools and AWS SDKs to automate scaling.

Managed AWS services include Amazon S3, Amazon CloudFront, AWS Auto Scaling, AWS Lambda, Amazon DynamoDB, AWS Fargate, and Amazon Route 53.

AWS Auto Scaling lets you detect and replace impaired instances. It also lets you build scaling plans for resources including [Amazon EC2](#) instances and Spot Fleets, [Amazon ECS](#) tasks, [Amazon DynamoDB](#) tables and indexes, and [Amazon Aurora](#) Replicas.

When scaling EC2 instances, ensure that you use multiple Availability Zones (preferably at least three) and add or remove capacity to maintain balance across these Availability Zones. ECS tasks or Kubernetes pods (when using Amazon Elastic Kubernetes Service) should also be distributed across multiple Availability Zones.

When using AWS Lambda, instances scale automatically. Every time an event notification is received for your function, AWS Lambda quickly locates free capacity within its compute fleet and runs your code up to the allocated concurrency. You need to ensure that the necessary concurrency is configured on the specific Lambda, and in your Service Quotas.

Amazon S3 automatically scales to handle high request rates. For example, your application can achieve at least 3,500 PUT/COPY/POST/DELETE or 5,500 GET/HEAD requests per second per prefix in a bucket. There are no limits to the number of prefixes in a bucket. You can increase your read or write performance by parallelizing reads. For example, if you create 10 prefixes in an Amazon S3 bucket to parallelize reads, you could scale your read performance to 55,000 read requests per second.

Configure and use Amazon CloudFront or a trusted content delivery network (CDN). A CDN can provide faster end-user response times and can serve requests for content from cache, therefore reducing the need to scale your workload.

Common anti-patterns:

- Implementing Auto Scaling groups for automated healing, but not implementing elasticity.
- Using automatic scaling to respond to large increases in traffic.
- Deploying highly stateful applications, eliminating the option of elasticity.

Benefits of establishing this best practice: Automation removes the potential for manual error in deploying and decommissioning resources. Automation removes the risk of cost overruns and denial of service due to slow response on needs for deployment or decommissioning.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Configure and use AWS Auto Scaling. This monitors your applications and automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost. Using AWS Auto Scaling, you can setup application scaling for multiple resources across multiple services.

- [What is AWS Auto Scaling?](#)
 - Configure Auto Scaling on your Amazon EC2 instances and Spot Fleets, Amazon ECS tasks, Amazon DynamoDB tables and indexes, Amazon Aurora Replicas, and AWS Marketplace appliances as applicable.
 - [Managing throughput capacity automatically with DynamoDB Auto Scaling](#)
 - Use service API operations to specify the alarms, scaling policies, warm up times, and cool down times.
- Use Elastic Load Balancing. Load balancers can distribute load by path or by network connectivity.
- [What is Elastic Load Balancing?](#)
 - Application Load Balancers can distribute load by path.
 - [What is an Application Load Balancer?](#)
 - Configure an Application Load Balancer to distribute traffic to different workloads based on the path under the domain name.
 - Application Load Balancers can be used to distribute loads in a manner that integrates with AWS Auto Scaling to manage demand.
 - [Using a load balancer with an Auto Scaling group](#)
 - Network Load Balancers can distribute load by connection.
 - [What is a Network Load Balancer?](#)
 - Configure a Network Load Balancer to distribute traffic to different workloads using TCP, or to have a constant set of IP addresses for your workload.
 - Network Load Balancers can be used to distribute loads in a manner that integrates with AWS Auto Scaling to manage demand.
- Use a highly available DNS provider. DNS names allow your users to enter names instead of IP addresses to access your workloads and distributes this information to a defined scope, usually globally for users of the workload.
- Use Amazon Route 53 or a trusted DNS provider.
 - [What is Amazon Route 53?](#)
- Use Route 53 to manage your CloudFront distributions and load balancers.
 - Determine the domains and subdomains you are going to manage.
 - Create appropriate record sets using ALIAS or CNAME records.
 - [Working with records](#)
- Use the AWS global network to optimize the path from your users to your applications. AWS Global Accelerator continually monitors the health of your application endpoints and redirects traffic to healthy endpoints in less than 30 seconds.
 - AWS Global Accelerator is a service that improves the availability and performance of your applications with local or global users. It provides static IP addresses that act as a fixed entry point to your application endpoints in a single or multiple AWS Regions, such as your Application Load Balancers, Network Load Balancers or Amazon EC2 instances.
 - [What Is AWS Global Accelerator?](#)
- Configure and use Amazon CloudFront or a trusted content delivery network (CDN). A content delivery network can provide faster end-user response times and can serve requests for content that may cause unnecessary scaling of your workloads.
 - [What is Amazon CloudFront?](#)
 - Configure Amazon CloudFront distributions for your workloads, or use a third-party CDN.
 - You can limit access to your workloads so that they are only accessible from CloudFront by using the IP ranges for CloudFront in your endpoint security groups or access policies.

Resources

Related documents:

- [APN Partner: partners that can help you create automated compute solutions](#)
- [AWS Auto Scaling: How Scaling Plans Work](#)
- [AWS Marketplace: products that can be used with auto scaling](#)
- [Managing Throughput Capacity Automatically with DynamoDB Auto Scaling](#)
- [Using a load balancer with an Auto Scaling group](#)
- [What Is AWS Global Accelerator?](#)
- [What Is Amazon EC2 Auto Scaling?](#)
- [What is AWS Auto Scaling?](#)
- [What is Amazon CloudFront?](#)
- [What is Amazon Route 53?](#)
- [What is Elastic Load Balancing?](#)
- [What is a Network Load Balancer?](#)
- [What is an Application Load Balancer?](#)
- [Working with records](#)

REL07-BP02 Obtain resources upon detection of impairment to a workload

Scale resources reactively when necessary if availability is impacted, to restore workload availability.

You first must configure health checks and the criteria on these checks to indicate when availability is impacted by lack of resources. Then either notify the appropriate personnel to manually scale the resource, or trigger automation to automatically scale it.

Scale can be manually adjusted for your workload, for example, changing the number of EC2 instances in an Auto Scaling group or modifying throughput of a DynamoDB table can be done through the AWS Management Console or AWS CLI. However automation should be used whenever possible (refer to **Use automation when obtaining or scaling resources**).

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Obtain resources upon detection of impairment to a workload. Scale resources reactively when necessary if availability is impacted, to restore workload availability.
 - Use scaling plans, which are the core component of AWS Auto Scaling, to configure a set of instructions for scaling your resources. If you work with AWS CloudFormation or add tags to AWS resources, you can set up scaling plans for different sets of resources, per application. AWS Auto Scaling provides recommendations for scaling strategies customized to each resource. After you create your scaling plan, AWS Auto Scaling combines dynamic scaling and predictive scaling methods together to support your scaling strategy.
 - [AWS Auto Scaling: How Scaling Plans Work](#)
 - Amazon EC2 Auto Scaling helps you ensure that you have the correct number of Amazon EC2 instances available to handle the load for your application. You create collections of EC2 instances, called Auto Scaling groups. You can specify the minimum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes below this size. You can specify the maximum number of instances in each Auto Scaling group, and Amazon EC2 Auto Scaling ensures that your group never goes above this size.
 - [What Is Amazon EC2 Auto Scaling?](#)
 - Amazon DynamoDB auto scaling uses the AWS Application Auto Scaling service to dynamically adjust provisioned throughput capacity on your behalf, in response to actual traffic patterns. This

enables a table or a global secondary index to increase its provisioned read and write capacity to handle sudden increases in traffic, without throttling.

- [Managing Throughput Capacity Automatically with DynamoDB Auto Scaling](#)

Resources

Related documents:

- [APN Partner: partners that can help you create automated compute solutions](#)
- [AWS Auto Scaling: How Scaling Plans Work](#)
- [AWS Marketplace: products that can be used with auto scaling](#)
- [Managing Throughput Capacity Automatically with DynamoDB Auto Scaling](#)
- [What Is Amazon EC2 Auto Scaling?](#)

REL07-BP03 Obtain resources upon detection that more resources are needed for a workload

Scale resources proactively to meet demand and avoid availability impact.

Many AWS services automatically scale to meet demand. If using Amazon EC2 instances or Amazon ECS clusters, you can configure automatic scaling of these to occur based on usage metrics that correspond to demand for your workload. For Amazon EC2, average CPU utilization, load balancer request count, or network bandwidth can be used to scale out (or scale in) EC2 instances. For Amazon ECS, average CPU utilization, load balancer request count, and memory utilization can be used to scale out (or scale in) ECS tasks. Using Target Auto Scaling on AWS, the autoscaler acts like a household thermostat, adding or removing resources to maintain the target value (for example, 70% CPU utilization) that you specify.

AWS Auto Scaling can also do [Predictive Auto Scaling](#), which uses machine learning to analyze each resource's historical workload and regularly forecasts the future load for the next two days.

Little's Law helps calculate how many instances of compute (EC2 instances, concurrent Lambda functions, etc.) that you need.

$$L = \lambda W$$

L = number of instances (or mean concurrency in the system)

λ = mean rate at which requests arrive (req/sec)

W = mean time that each request spends in the system (sec)

For example, at 100 rps, if each request takes 0.5 seconds to process, you will need 50 instances to keep up with demand.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Obtain resources upon detection that more resources are needed for a workload. Scale resources proactively to meet demand and avoid availability impact.
 - Calculate how many compute resources you will need (compute concurrency) to handle a given request rate.
 - [Telling Stories About Little's Law](#)
 - When you have a historical pattern for usage, set up scheduled scaling for Amazon EC2 auto scaling.
 - [Scheduled Scaling for Amazon EC2 Auto Scaling](#)
 - Use AWS predictive scaling.

- [Predictive Scaling for EC2, Powered by Machine Learning](#)

Resources

Related documents:

- [AWS Auto Scaling: How Scaling Plans Work](#)
- [AWS Marketplace: products that can be used with auto scaling](#)
- [Managing Throughput Capacity Automatically with DynamoDB Auto Scaling](#)
- [Predictive Scaling for EC2, Powered by Machine Learning](#)
- [Scheduled Scaling for Amazon EC2 Auto Scaling](#)
- [Telling Stories About Little's Law](#)
- [What Is Amazon EC2 Auto Scaling?](#)

REL07-BP04 Load test your workload

Adopt a load testing methodology to measure if scaling activity meets workload requirements.

It's important to perform sustained load testing. Load tests should discover the breaking point and test the performance of your workload. AWS makes it easy to set up temporary testing environments that model the scale of your production workload. In the cloud, you can create a production-scale test environment on demand, complete your testing, and then decommission the resources. Because you only pay for the test environment when it's running, you can simulate your live environment for a fraction of the cost of testing on premises.

Load testing in production should also be considered as part of game days where the production system is stressed, during hours of lower customer usage, with all personnel on hand to interpret results and address any problems that arise.

Common anti-patterns:

- Performing load testing on deployments that are not the same configuration as your production.
- Performing load testing only on individual pieces of your workload, and not on the entire workload.
- Performing load testing with a subset of requests and not a representative set of actual requests.
- Performing load testing to a small safety factor above expected load.

Benefits of establishing this best practice: You know what components in your architecture fail under load and be able to identify what metrics to watch to indicate that you are approaching that load in time to address the problem, preventing the impact of that failure.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Perform load testing to identify which aspect of your workload indicates that you must add or remove capacity. Load testing should have representative traffic similar to what you receive in production. Increase the load while watching the metrics you have instrumented to determine which metric indicates when you must add or remove resources.
 - [Distributed Load Testing on AWS: simulate thousands of connected users](#)
 - Identify the mix of requests. You may have varied mixes of requests, so you should look at various time frames when identifying the mix of traffic.
 - Implement a load driver. You can use custom code, open source, or commercial software to implement a load driver.

- Load test initially using small capacity. You see some immediate effects by driving load onto a lesser capacity, possibly as small as one instance or container.
- Load test against larger capacity. The effects will be different on a distributed load, so you must test against as close to a product environment as possible.

Resources

Related documents:

- [Distributed Load Testing on AWS: simulate thousands of connected users](#)

REL 8 How do you implement change?

Controlled changes are necessary to deploy new functionality, and to ensure that the workloads and the operating environment are running known software and can be patched or replaced in a predictable manner. If these changes are uncontrolled, then it makes it difficult to predict the effect of these changes, or to address issues that arise because of them.

Best practices

- [REL08-BP01 Use runbooks for standard activities such as deployment \(p. 216\)](#)
- [REL08-BP02 Integrate functional testing as part of your deployment \(p. 217\)](#)
- [REL08-BP03 Integrate resiliency testing as part of your deployment \(p. 218\)](#)
- [REL08-BP04 Deploy using immutable infrastructure \(p. 218\)](#)
- [REL08-BP05 Deploy changes with automation \(p. 220\)](#)

REL08-BP01 Use runbooks for standard activities such as deployment

Runbooks are the predefined procedures to achieve specific outcomes. Use runbooks to perform standard activities, whether done manually or automatically. Examples include deploying a workload, patching a workload, or making DNS modifications.

For example, put processes in place to [ensure rollback safety during deployments](#). Ensuring that you can roll back a deployment without any disruption for your customers is critical in making a service reliable.

For runbook procedures, start with a valid effective manual process, implement it in code, and trigger it to automatically run where appropriate.

Even for sophisticated workloads that are highly automated, runbooks are still useful for [running game days](#) or meeting rigorous reporting and auditing requirements.

Note that playbooks are used in response to specific incidents, and runbooks are used to achieve specific outcomes. Often, runbooks are for routine activities, while playbooks are used for responding to non-routine events.

Common anti-patterns:

- Performing unplanned changes to configuration in production.
- Skipping steps in your plan to deploy faster, resulting in a failed deployment.
- Making changes without testing the reversal of the change.

Benefits of establishing this best practice: Effective change planning increases your ability to successfully execute the change because you are aware of all the systems impacted. Validating your change in test environments increases your confidence.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Enable consistent and prompt responses to well understood events by documenting procedures in runbooks.
 - [AWS Well-Architected Framework: Concepts: Runbook](#)
- Use the principle of infrastructure as code to define your infrastructure. By using AWS CloudFormation (or a trusted third party) to define your infrastructure, you can use version control software to version and track changes.
 - Use AWS CloudFormation (or a trusted third-party provider) to define your infrastructure.
 - [What is AWS CloudFormation?](#)
 - Create templates that are singular and decoupled, using good software design principles.
 - Determine the permissions, templates, and responsible parties for implementation.
 - [Controlling access with AWS Identity and Access Management](#)
 - Use source control, like AWS CodeCommit or a trusted third-party tool, for version control.
 - [What is AWS CodeCommit?](#)

Resources

Related documents:

- [APN Partner: partners that can help you create automated deployment solutions](#)
- [AWS Marketplace: products that can be used to automate your deployments](#)
- [AWS Well-Architected Framework: Concepts: Runbook](#)
- [What is AWS CloudFormation?](#)
- [What is AWS CodeCommit?](#)

Related examples:

- [Automating operations with Playbooks and Runbooks](#)

REL08-BP02 Integrate functional testing as part of your deployment

Functional tests are run as part of automated deployment. If success criteria are not met, the pipeline is halted or rolled back.

These tests are run in a pre-production environment, which is staged prior to production in the pipeline. Ideally, this is done as part of a deployment pipeline.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Integrate functional testing as part of your deployment. Functional tests are run as part of automated deployment. If success criteria are not met, the pipeline is halted or rolled back.
- Invoke AWS CodeBuild during the 'Test Action' of your software release pipelines modeled in AWS CodePipeline. This capability enables you to easily run a variety of tests against your code, such as unit tests, static code analysis, and integration tests.
 - [AWS CodePipeline Adds Support for Unit and Custom Integration Testing with AWS CodeBuild](#)
- Use AWS Marketplace solutions for executing automated tests as part of your software delivery pipeline.
 - [Software test automation](#)

Resources

Related documents:

- [AWS CodePipeline Adds Support for Unit and Custom Integration Testing with AWS CodeBuild](#)
- [Software test automation](#)
- [What Is AWS CodePipeline?](#)

REL08-BP03 Integrate resiliency testing as part of your deployment

Resiliency tests (using the [principles of chaos engineering](#)) are run as part of the automated deployment pipeline in a pre-production environment.

These tests are staged and run in the pipeline in a pre-production environment. They should also be run in production as part of [game days](#).

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Integrate resiliency testing as part of your deployment. Use Chaos Engineering, the discipline of experimenting on a workload to build confidence in the workload's capability to withstand turbulent conditions in production.
- Resiliency tests inject faults or resource degradation to assess that your workload responds with its designed resilience.
 - [Well-Architected lab: Level 300: Testing for Resiliency of EC2 RDS and S3](#)
- These tests can be run regularly in pre-production environments in automated deployment pipelines.
- They should also be run in production, as part of scheduled game days.
- Using Chaos Engineering principles, propose hypotheses about how your workload will perform under various impairments, then test your hypotheses using resiliency testing.
 - [Principles of Chaos Engineering](#)

Resources

Related documents:

- [Principles of Chaos Engineering](#)
- [What is AWS Fault Injection Simulator?](#)

Related examples:

- [Well-Architected lab: Level 300: Testing for Resiliency of EC2 RDS and S3](#)

REL08-BP04 Deploy using immutable infrastructure

Immutable infrastructure is a model that mandates that no updates, security patches, or configuration changes happen in-place on production workloads. When a change is needed, the architecture is built onto new infrastructure and deployed into production.

The most common implementation of the immutable infrastructure paradigm is the *immutable server*. This means that if a server needs an update or a fix, new servers are deployed instead of updating the

ones already in use. So, instead of logging into the server via SSH and updating the software version, every change in the application starts with a software push to the code repository, for example, git push. Since changes are not allowed in immutable infrastructure, you can be sure about the state of the deployed system. Immutable infrastructures are inherently more consistent, reliable, and predictable, and they simplify many aspects of software development and operations.

Use a canary or blue/green deployment when deploying applications in immutable infrastructures.

[Canary deployment](#) is the practice of directing a small number of your customers to the new version, usually running on a single service instance (the canary). You then deeply scrutinize any behavior changes or errors that are generated. You can remove traffic from the canary if you encounter critical problems and send the users back to the previous version. If the deployment is successful, you can continue to deploy at your desired velocity, while monitoring the changes for errors, until you are fully deployed. AWS CodeDeploy can be configured with a deployment configuration that will enable a canary deployment.

[Blue/green deployment](#) is similar to the canary deployment except that a full fleet of the application is deployed in parallel. You alternate your deployments across the two stacks (blue and green). Once again, you can send traffic to the new version, and fall back to the old version if you see problems with the deployment. Commonly all traffic is switched at once, however you can also use fractions of your traffic to dial up the adoption of the new version using the weighted DNS routing capabilities of Amazon Route 53. AWS CodeDeploy and AWS Elastic Beanstalk can be configured with a deployment configuration that will enable a blue/green deployment.




Figure 8: Blue/green deployment with AWS Elastic Beanstalk and Amazon Route 53

Benefits of immutable infrastructure:

- **Reduction in configuration drifts:** By frequently replacing servers from a base, known and version-controlled configuration, the infrastructure is **reset** to a known state, avoiding configuration drifts.
- **Simplified deployments:** Deployments are simplified because they don't need to support upgrades. Upgrades are just new deployments.
- **Reliable atomic deployments:** Deployments either complete successfully, or nothing changes. It gives more trust in the deployment process.
- **Safer deployments with fast rollback and recovery processes:** Deployments are safer because the previous working version is not changed. You can roll back to it if errors are detected.
- **Consistent testing and debugging environments:** Since all servers use the same image, there are no differences between environments. One build is deployed to multiple environments. It also prevents inconsistent environments and simplifies testing and debugging.

- **Increased scalability:** Since servers use a base image, are consistent, and repeatable, automatic scaling is trivial.
- **Simplified toolchain:** The toolchain is simplified since you can get rid of configuration management tools managing production software upgrades. No extra tools or agents are installed on servers. Changes are made to the base image, tested, and rolled-out.
- **Increased security:** By denying all changes to servers, you can disable SSH on instances and remove keys. This reduces the attack vector, improving your organization's security posture.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Deploy using immutable infrastructure. Immutable infrastructure is a model in which no updates, security patches, or configuration changes happen *in-place* on production systems. If any change is needed, a new version of the architecture is built and deployed into production.
 - [Overview of a Blue/Green Deployment](#)
 - [Deploying Serverless Applications Gradually](#)
 - [Immutable Infrastructure: Reliability, consistency and confidence through immutability](#)
 - [CanaryRelease](#)

Resources

Related documents:

- [CanaryRelease](#)
- [Deploying Serverless Applications Gradually](#)
- [Immutable Infrastructure: Reliability, consistency and confidence through immutability](#)
- [Overview of a Blue/Green Deployment](#)
- [The Amazon Builders' Library: Ensuring rollback safety during deployments](#)

REL08-BP05 Deploy changes with automation

Deployments and patching are automated to eliminate negative impact.

Making changes to production systems is one of the largest risk areas for many organizations. We consider deployments a first-class problem to be solved alongside the business problems that the software addresses. Today, this means the use of automation wherever practical in operations, including testing and deploying changes, adding or removing capacity, and migrating data. AWS CodePipeline lets you manage the steps required to release your workload. This includes a deployment state using AWS CodeDeploy to automate deployment of application code to Amazon EC2 instances, on-premises instances, serverless Lambda functions, or Amazon ECS services.

Recommendation

Although conventional wisdom suggests that you keep humans in the loop for the most difficult operational procedures, we suggest that you automate the most difficult procedures for that very reason.

Common anti-patterns:

- Manually performing changes.
- Skipping steps in your automation through emergency work flows.
- Not following your plans.

Benefits of establishing this best practice: Using automation to deploy all changes removes the potential for introduction of human error and enables the ability to test before changing production to ensure that your plans are complete.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Automate your deployment pipeline. Deployment pipelines allow you to invoke automated testing and detection of anomalies, and either halt the pipeline at a certain step before production deployment, or automatically roll back a change.
 - [The Amazon Builders' Library: Ensuring rollback safety during deployments](#)
 - [The Amazon Builders' Library: Going faster with continuous delivery](#)
 - Use AWS CodePipeline (or a trusted third-party product) to define and run your pipelines.
 - Configure the pipeline to start when a change is committed to your code repository.
 - [What is AWS CodePipeline?](#)
 - Use Amazon Simple Notification Service (Amazon SNS) and Amazon Simple Email Service (Amazon SES) to send notifications about problems in the pipeline or integrate with a team chat tool, like Amazon Chime.
 - [What is Amazon Simple Notification Service?](#)
 - [What is Amazon SES?](#)
 - [What is Amazon Chime?](#)
 - [Automate chat messages with webhooks.](#)

Resources

Related documents:

- [APN Partner: partners that can help you create automated deployment solutions](#)
- [AWS Marketplace: products that can be used to automate your deployments](#)
- [Automate chat messages with webhooks.](#)
- [The Amazon Builders' Library: Ensuring rollback safety during deployments](#)
- [The Amazon Builders' Library: Going faster with continuous delivery](#)
- [What Is AWS CodePipeline?](#)
- [What Is CodeDeploy?](#)
- [AWS Systems Manager Patch Manager](#)
- [What is Amazon SES?](#)
- [What is Amazon Simple Notification Service?](#)

Related videos:

- [AWS Summit 2019: CI/CD on AWS](#)

Failure management

Questions

- [REL 9 How do you back up data? \(p. 222\)](#)
- [REL 10 How do you use fault isolation to protect your workload? \(p. 230\)](#)
- [REL 11 How do you design your workload to withstand component failures? \(p. 240\)](#)

- [REL 12 How do you test reliability? \(p. 249\)](#)
- [REL 13 How do you plan for disaster recovery \(DR\)? \(p. 261\)](#)

REL 9 How do you back up data?

Back up data, applications, and configuration to meet your requirements for recovery time objectives (RTO) and recovery point objectives (RPO).

Best practices

- [REL09-BP01 Identify and back up all data that needs to be backed up, or reproduce the data from sources \(p. 222\)](#)
- [REL09-BP02 Secure and encrypt backups \(p. 224\)](#)
- [REL09-BP03 Perform data backup automatically \(p. 226\)](#)
- [REL09-BP04 Perform periodic recovery of the data to verify backup integrity and processes \(p. 227\)](#)

REL09-BP01 Identify and back up all data that needs to be backed up, or reproduce the data from sources

All AWS data stores offer backup capabilities. Services such as Amazon RDS and Amazon DynamoDB additionally support automated backup that enables point-in-time recovery (PITR), which allows you to restore a backup to any time up to five minutes or less before the current time. Many AWS services offer the ability to copy backups to another AWS Region. AWS Backup is a tool that gives you the ability to centralize and automate data protection across AWS services.

Amazon S3 can be used as a backup destination for self-managed and AWS-managed data sources. AWS services such as Amazon EBS, Amazon RDS, and Amazon DynamoDB have built in capabilities to create backups. Third-party backup software can also be used.

On-premises data can be backed up to the AWS Cloud using [AWS Storage Gateway](#) or [AWS DataSync](#). Amazon S3 buckets can be used to store this data on AWS. Amazon S3 offers multiple storage tiers such as [Amazon S3 Glacier](#) or [S3 Glacier Deep Archive](#) to reduce cost of data storage.

You might be able to meet data recovery needs by reproducing the data from other sources. For example, [Amazon ElastiCache replica nodes](#) or [RDS read replicas](#) could be used to reproduce data if the primary is lost. In cases where sources like this can be used to meet your [Recovery Point Objective \(RPO\)](#) and [Recovery Time Objective \(RTO\)](#), you might not require a backup. Another example, if working with Amazon EMR, it might not be necessary to backup your HDFS data store, as long as you can [reproduce the data into EMR from S3](#).

When selecting a backup strategy, consider the time it takes to recover data. The time needed to recover data depends on the type of backup (in the case of a backup strategy), or the complexity of the data reproduction mechanism. This time should fall within the RTO for the workload.

Desired Outcome:

Data sources have been identified and classified based on criticality. Then, establish a strategy for data recovery based on the RPO. This strategy involves either backing up these data sources, or having the ability to reproduce data from other sources. In the case of data loss, the strategy implemented enables recovery or reproduction of data within the defined RPO and RTO.

Cloud Maturity Phase: Foundational

Common anti-patterns:

- Not aware of all data sources for the workload and their criticality.

- Not taking backups of critical data sources.
- Taking backups of only some data sources without using criticality as a criterion.
- No defined RPO, or backup frequency cannot meet RPO.
- Not evaluating if a backup is necessary or if data can be reproduced from other sources.

Benefits of establishing this best practice: Identifying the places where backups are necessary and implementing a mechanism to create backups, or being able to reproduce the data from an external source improves the ability to restore and recover data during an outage.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Understand and use the backup capabilities of the AWS services and resources used by the workload. Most AWS services provides capabilities to back up workload data.

Implementation Steps:

1. **Identify all data sources for the workload.** Data can be stored on a number of resources such as [databases](#), [volumes](#), [filesystems](#), [logging systems](#), and [object storage](#). Refer to the **Resources** section to find **Related documents** on different AWS services where data is stored, and the backup capability these services provide.
2. **Classify data sources based on criticality.** Different data sets will have different levels of criticality for a workload, and therefore different requirements for resiliency. For example, some data might be critical and require a RPO near zero, while other data might be less critical and can tolerate a higher RPO and some data loss. Similarly, different data sets might have different RTO requirements as well.
3. **Use AWS or third-party services to create backups of the data.** [AWS Backup](#) is a managed service that enables creating backups of various data sources on AWS. Most of these services also have native capabilities to create backups. The AWS Marketplace has many solutions that provide these capabilities as well. Refer to the **Resources** listed below for information on how to create backups of data from various AWS services.
4. **For data that is not backed up, establish a data reproduction mechanism.** You might choose not to backup data that can be reproduced from other sources for various reasons. There might be a situation where it is cheaper to reproduce data from sources when needed rather than creating a backup as there may be a cost associated with storing backups. Another example is where restoring from a backup takes longer than reproducing the data from sources, resulting in a breach in RTO. In such situations, consider tradeoffs and establish a well-defined process for how data can be reproduced from these sources when data recovery is necessary. For example, if you have loaded data from Amazon S3 to a data warehouse (like Amazon Redshift), or MapReduce cluster (like Amazon EMR) to do analysis on that data, this may be an example of data that can be reproduced from other sources. As long as the results of these analyses are either stored somewhere or reproducible, you would not suffer a data loss from a failure in the data warehouse or MapReduce cluster. Other examples that can be reproduced from sources include caches (like Amazon ElastiCache) or RDS read replicas.
5. **Establish a cadence for backing up data.** Creating backups of data sources is a periodic process and the frequency should depend on the RPO.

Level of effort for the Implementation Plan: Moderate

Resources

Related Best Practices:

[REL13-BP01 Define recovery objectives for downtime and data loss \(p. 261\)](#)

[REL13-BP02 Use defined recovery strategies to meet the recovery objectives \(p. 265\)](#)

Related documents:

- [What Is AWS Backup?](#)
- [What is AWS DataSync?](#)
- [What is Volume Gateway?](#)
- [APN Partner: partners that can help with backup](#)
- [AWS Marketplace: products that can be used for backup](#)
- [Amazon EBS Snapshots](#)
- [Backing Up Amazon EFS](#)
- [Backing up Amazon FSx for Windows File Server](#)
- [Backup and Restore for ElastiCache for Redis](#)
- [Creating a DB Cluster Snapshot in Neptune](#)
- [Creating a DB Snapshot](#)
- [Creating an EventBridge Rule That Triggers on a Schedule](#)
- [Cross-Region Replication with Amazon S3](#)
- [EFS-to-EFS AWS Backup](#)
- [Exporting Log Data to Amazon S3](#)
- [Object lifecycle management](#)
- [On-Demand Backup and Restore for DynamoDB](#)
- [Point-in-time recovery for DynamoDB](#)
- [Working with Amazon OpenSearch Service Index Snapshots](#)

Related videos:

- [AWS re:Invent 2021 - Backup, disaster recovery, and ransomware protection with AWS](#)
- [AWS Backup Demo: Cross-Account and Cross-Region Backup](#)
- [AWS re:Invent 2019: Deep dive on AWS Backup, ft. Rackspace \(STG341\)](#)

Related examples:

- [Well-Architected lab: Implementing Bi-Directional Cross-Region Replication \(CRR\) for Amazon S3](#)
- [Well-Architected lab: Testing Backup and Restore of Data](#)
- [Well-Architected lab: Backup and Restore with Failback for Analytics Workload](#)
- [Well-Architected lab: Disaster Recovery - Backup and Restore](#)

REL09-BP02 Secure and encrypt backups

Control and detect access to backups using authentication and authorization, such as AWS IAM. Prevent and detect if data integrity of backups is compromised using encryption.

Amazon S3 supports several methods of encryption of your data at rest. Using server-side encryption, Amazon S3 accepts your objects as unencrypted data, and then encrypts them as they are stored. Using client-side encryption, your workload application is responsible for encrypting the data before it is sent to Amazon S3. Both methods allow you to use AWS Key Management Service (AWS KMS) to create and store the data key, or you can provide your own key, which you are then responsible for. Using AWS KMS, you can set policies using IAM on who can and cannot access your data keys and decrypted data.

For Amazon RDS, if you have chosen to encrypt your databases, then your backups are encrypted also. DynamoDB backups are always encrypted.

Common anti-patterns:

- Having the same access to the backups and restoration automation as you do to the data.
- Not encrypting your backups.

Benefits of establishing this best practice: Securing your backups prevents tampering with the data, and encryption of the data prevents access to that data if it is accidentally exposed.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use encryption on each of your data stores. If your source data is encrypted, then the backup will also be encrypted.
- Enable encryption in RDS. You can configure encryption at rest using AWS Key Management Service when you create an RDS instance.
 - [Encrypting Amazon RDS Resources](#)
- Enable encryption on EBS volumes. You can configure default encryption or specify a unique key upon volume creation.
 - [Amazon EBS Encryption](#)
- Use the required Amazon DynamoDB encryption. DynamoDB encrypts all data at rest. You can either use an AWS owned AWS KMS key or an AWS managed KMS key, specifying a key that is stored in your account.
 - [DynamoDB Encryption at Rest](#)
 - [Managing Encrypted Tables](#)
- Encrypt your data stored in Amazon EFS. Configure the encryption when you create your file system.
 - [Encrypting Data and Metadata in EFS](#)
- Configure the encryption in the source and destination Regions. You can configure encryption at rest in Amazon S3 using keys stored in KMS, but the keys are Region-specific. You can specify the destination keys when you configure the replication.
 - [CRR Additional Configuration: Replicating Objects Created with Server-Side Encryption \(SSE\) Using Encryption Keys stored in AWS KMS](#)
- Implement least privilege permissions to access your backups. Follow best practices to limit the access to the backups, snapshots, and replicas in accordance with security best practices.
 - [Security Pillar: AWS Well-Architected](#)

Resources

Related documents:

- [AWS Marketplace: products that can be used for backup](#)
- [Amazon EBS Encryption](#)
- [Amazon S3: Protecting Data Using Encryption](#)
- [CRR Additional Configuration: Replicating Objects Created with Server-Side Encryption \(SSE\) Using Encryption Keys stored in AWS KMS](#)
- [DynamoDB Encryption at Rest](#)
- [Encrypting Amazon RDS Resources](#)
- [Encrypting Data and Metadata in EFS](#)
- [Encryption for Backups in AWS](#)
- [Managing Encrypted Tables](#)

- [Security Pillar: AWS Well-Architected](#)

Related examples:

- [Well-Architected lab: Implementing Bi-Directional Cross-Region Replication \(CRR\) for Amazon S3](#)

REL09-BP03 Perform data backup automatically

Configure backups to be taken automatically based on a periodic schedule informed by the Recovery Point Objective (RPO), or by changes in the dataset. Critical datasets with low data loss requirements need to be backed up automatically on a frequent basis, whereas less critical data where some loss is acceptable can be backed up less frequently.

AWS Backup can be used to create automated data backups of various AWS data sources. Amazon RDS instances can be backed up almost continuously every five minutes and Amazon S3 objects can be backed up almost continuously every fifteen minutes, providing for point-in-time recovery (PITR) to a specific point in time within the backup history. For other AWS data sources, such as Amazon EBS volumes, Amazon DynamoDB tables, or Amazon FSx file systems, AWS Backup can run automated backup as frequently as every hour. These services also offer native backup capabilities. AWS services that offer automated backup with point-in-time recovery include [Amazon DynamoDB](#), [Amazon RDS](#), and [Amazon Keyspaces \(for Apache Cassandra\)](#) – these can be restored to a specific point in time within the backup history. Most other AWS data storage services offer the ability to schedule periodic backups, as frequently as every hour.

Amazon RDS and Amazon DynamoDB offer continuous backup with point-in-time recovery. Amazon S3 versioning, once enabled, is automatic. [Amazon Data Lifecycle Manager](#) can be used to automate the creation, copy and deletion of Amazon EBS snapshots. It can also automate the creation, copy, deprecation and deregistration of Amazon EBS-backed Amazon Machine Images (AMIs) and their underlying Amazon EBS snapshots.

For a centralized view of your backup automation and history, AWS Backup provides a fully managed, policy-based backup solution. It centralizes and automates the back up of data across multiple AWS services in the cloud as well as on premises using the AWS Storage Gateway.

In addition to versioning, Amazon S3 features replication. The entire S3 bucket can be automatically replicated to another bucket in the same, or a different AWS Region.

Desired Outcome:

An automated process that creates backups of data sources at an established cadence.

Common anti-patterns:

- Performing backups manually.
- Using resources that have backup capability, but not including the backup in your automation.

Benefits of establishing this best practice: Automating backups ensures that they are taken regularly based on your RPO, and alerts you if they are not taken.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

1. **Identify data sources** that are currently being backed up manually. Refer to [REL09-BP01 Identify and back up all data that needs to be backed up, or reproduce the data from sources \(p. 222\)](#) for guidance on this.

2. **Determine the RPO** for the workload. Refer to [REL13-BP01 Define recovery objectives for downtime and data loss \(p. 261\)](#) for guidance on this.
3. **Use an automated backup solution or managed service.** AWS Backup is a fully-managed service that makes it easy to [centralize and automate data protection across AWS services, in the cloud, and on premises](#). Backup plans are a feature of AWS Backup that enables the creation of rules which define the resources to backup, and the frequency at which these backups should be created. This frequency should be informed by the RPO established in Step 2. [This WA Lab](#) provides hands-on guidance on how to create automated backups using AWS Backup. Native backup capabilities are offered by most AWS services that store data. For example, RDS can be leveraged for automated backups with point-in-time recovery (PITR).
4. **For data sources not supported** by an automated backup solution or managed service such as on-premises data sources or message queues, consider using a trusted third-party solution to create automated backups. Alternatively, you can create automation to do this using the AWS CLI or SDKs. You can use AWS Lambda Functions or AWS Step Functions to define the logic involved in creating a data backup, and use Amazon EventBridge to execute it at a frequency based on your RPO (as established in Step 2).

Level of effort for the Implementation Plan: Low

Resources

Related documents:

- [APN Partner: partners that can help with backup](#)
- [AWS Marketplace: products that can be used for backup](#)
- [Creating an EventBridge Rule That Triggers on a Schedule](#)
- [What Is AWS Backup?](#)
- [What Is AWS Step Functions?](#)

Related videos:

- [AWS re:Invent 2019: Deep dive on AWS Backup, ft. Rackspace \(STG341\)](#)

Related examples:

- [Well-Architected lab: Testing Backup and Restore of Data](#)

REL09-BP04 Perform periodic recovery of the data to verify backup integrity and processes

Validate that your backup process implementation meets your recovery time objectives (RTO) and recovery point objectives (RPO) by performing a recovery test.

Using AWS, you can stand up a testing environment and restore your backups to assess RTO and RPO capabilities, and run tests on data content and integrity.

Additionally, Amazon RDS and Amazon DynamoDB allow point-in-time recovery (PITR). Using continuous backup, you can restore your dataset to the state it was in at a specified date and time.

Desired Outcome: Data from backups is periodically recovered using well-defined mechanisms to ensure that recovery is possible within the established recovery time objective (RTO) for the workload. Verify that restoration from a backup results in a resource that contains the original data without any of it being corrupted or inaccessible, and with data loss within the recovery point objective (RPO).

Common anti-patterns:

- Restoring a backup, but not querying or retrieving any data to ensure that the restoration is usable.
- Assuming that a backup exists.
- Assuming that the backup of a system is fully operational and that data can be recovered from it.
- Assuming that the time to restore or recover data from a backup falls within the RTO for the workload.
- Assuming that the data contained on the backup falls within the RPO for the workload
- Restoring ad hoc, without using a runbook, or outside of an established automated procedure.

Benefits of establishing this best practice: Testing the recovery of the backups ensures data can be restored when needed without having any worry that data might be missing or corrupted, that the restoration and recovery is possible within the RTO for the workload, and any data loss falls within the RPO for the workload.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Testing backup and restore capability increases confidence in the ability to perform these actions during an outage. Periodically restore backups to a new location and run tests to verify the integrity of the data. Some common tests that should be performed are checking

if all the data is available, is not corrupted, is accessible, and any data loss falls within the RPO for the workload. Such tests can also help ascertain if recovery mechanisms are fast enough to accommodate the workload's RTO.

1. **Identify data sources** that are currently being backed up and where these backups are being stored. Refer to [REL09-BP01 Identify and back up all data that needs to be backed up, or reproduce the data from sources \(p. 222\)](#) for guidance on how to implement this.
2. **Establish criteria for data validation** for each data source. Different types of data will have different properties which might require different validation mechanisms. Consider how this data might be validated before you are confident to use it in production. Some common ways to validate data are using data and backup properties such as data type, format, checksum, size, or a combination of these with custom validation logic. For example, this might be a comparison of the checksum values between the restored resource and the data source at the time the backup was created.
3. **Establish RTO and RPO** for restoring the data based on data criticality. Refer to [REL13-BP01 Define recovery objectives for downtime and data loss \(p. 261\)](#) for guidance on how to implement this.
4. **Assess your recovery capability.** Review your backup and restore strategy to understand if it can meet your RTO and RPO, and adjust the strategy as necessary. Using [AWS Resilience Hub](#), you can run an assessment of your workload. The assessment evaluates your application configuration against the resiliency policy and reports if your RTO and RPO targets can be met.
5. **Do a test restore** using currently established processes used in production for data restoration. These processes depend on how the original data source was backed up, the format and storage location of the backup itself, or if the data is reproduced from other sources. For example, if you are using a managed service such as [AWS Backup](#), [this might be as simple as restoring the backup into a new resource](#). If you used AWS Elastic Disaster Recovery you can [launch a recovery drill](#).
6. **Validate data recovery** from the restored resource (from the previous step) based on criteria you previously established for data validation in step 2. Does the restored and recovered data contain the most recent record/item at the time of backup? Does this data fall within the RPO for the workload?
7. **Measure time required** for restore and recovery and compare it to RTO established earlier in step 3. Does this process fall within the RTO for the workload? For example, compare the timestamps from when the restoration process started and when the recovery validation completed to calculate how long this process takes. All AWS API calls are timestamped and this information is available

in [AWS CloudTrail](#). While this information can provide details on when the restore process started, the end timestamp for when the validation was completed should be recorded by your validation logic. If using an automated process, then services like [Amazon DynamoDB](#) can be used to store this information. Additionally, many AWS services provide an event history which provides timestamped information when certain actions occurred. Within AWS Backup, backup and restore actions are referred to as *Jobs*, and these Jobs contain timestamp information as part of its metadata which can be used to measure time required for restoration and recovery.

8. **Notify stakeholders** if data validation fails, or if the time required for restoration and recovery exceeds the established RTO for the workload. When implementing automation to do this, [such as in this lab](#), services like Amazon Simple Notification Service (Amazon SNS) can be used to send push notifications such as email or SMS to stakeholders. [These messages can also be published to messaging applications such as Amazon Chime, Slack, or Microsoft Teams](#) or used to [create tasks as OpsItems using AWS Systems Manager OpsCenter](#).
9. **Automate this process to run periodically.** For example, services like AWS Lambda or a State Machine in AWS Step Functions can be used to automate the restore and recovery processes, and Amazon EventBridge can be used to trigger this automation workflow periodically as shown in the architecture diagram below. Learn how to [Automate data recovery validation with AWS Backup](#). Additionally, [this Well-Architected lab](#) provides a hands-on experience on one way to do automation for several of the steps here.




Figure 9. An automated backup and restore process

Level of effort for the Implementation Plan: Moderate to high depending on the complexity of the validation criteria.

Resources

Related documents:

- [Automate data recovery validation with AWS Backup](#)
- [APN Partner: partners that can help with backup](#)
- [AWS Marketplace: products that can be used for backup](#)
- [Creating an EventBridge Rule That Triggers on a Schedule](#)
- [On-demand backup and restore for DynamoDB](#)

- [What Is AWS Backup?](#)
- [What Is AWS Step Functions?](#)
- [What is AWS Elastic Disaster Recovery](#)
- [AWS Elastic Disaster Recovery](#)

Related examples:

- [Well-Architected lab: Testing Backup and Restore of Data](#)

REL 10 How do you use fault isolation to protect your workload?

Fault isolated boundaries limit the effect of a failure within a workload to a limited number of components. Components outside of the boundary are unaffected by the failure. Using multiple fault isolated boundaries, you can limit the impact on your workload.

Best practices

- [REL10-BP01 Deploy the workload to multiple locations \(p. 230\)](#)
- [REL10-BP02 Select the appropriate locations for your multi-location deployment \(p. 234\)](#)
- [REL10-BP03 Automate recovery for components constrained to a single location \(p. 237\)](#)
- [REL10-BP04 Use bulkhead architectures to limit scope of impact \(p. 238\)](#)

REL10-BP01 Deploy the workload to multiple locations

Distribute workload data and resources across multiple Availability Zones or, where necessary, across AWS Regions. These locations can be as diverse as required.

One of the bedrock principles for service design in AWS is the avoidance of single points of failure in underlying physical infrastructure. This motivates us to build software and systems that use multiple Availability Zones and are resilient to failure of a single zone. Similarly, systems are built to be resilient to failure of a single compute node, single storage volume, or single instance of a database. When building a system that relies on redundant components, it's important to ensure that the components operate independently, and in the case of AWS Regions, autonomously. The benefits achieved from theoretical availability calculations with redundant components are only valid if this holds true.

Availability Zones (AZs)

AWS Regions are composed of multiple Availability Zones that are designed to be independent of each other. Each Availability Zone is separated by a meaningful physical distance from other zones to avoid correlated failure scenarios due to environmental hazards like fires, floods, and tornadoes. Each Availability Zone also has independent physical infrastructure: dedicated connections to utility power, standalone backup power sources, independent mechanical services, and independent network connectivity within and beyond the Availability Zone. This design limits faults in any of these systems to just the one affected AZ. Despite being geographically separated, Availability Zones are located in the same regional area which enables high-throughput, low-latency networking. The entire AWS Region (across all Availability Zones, consisting of multiple physically independent data centers) can be treated as a single logical deployment target for your workload, including the ability to synchronously replicate data (for example, between databases). This allows you to use Availability Zones in an active/active or active/standby configuration.

Availability Zones are independent, and therefore workload availability is increased when the workload is architected to use multiple zones. Some AWS services (including the Amazon EC2 instance data plane) are deployed as strictly zonal services where they have shared fate with the Availability Zone they are in.

Amazon EC2 instances in the other AZs will however be unaffected and continue to function. Similarly, if a failure in an Availability Zone causes an Amazon Aurora database to fail, a read-replica Aurora instance in an unaffected AZ can be automatically promoted to primary. Regional AWS services, such as Amazon DynamoDB on the other hand internally use multiple Availability Zones in an active/active configuration to achieve the availability design goals for that service, without you needing to configure AZ placement.




Figure 9: Multi-tier architecture deployed across three Availability Zones. Note that Amazon S3 and Amazon DynamoDB are always Multi-AZ automatically. The ELB also is deployed to all three zones.

While AWS control planes typically provide the ability to manage resources within the entire Region (multiple Availability Zones), certain control planes (including Amazon EC2 and Amazon EBS) have the ability to filter results to a single Availability Zone. When this is done, the request is processed only in the specified Availability Zone, reducing exposure to disruption in other Availability Zones. This AWS CLI example illustrates getting Amazon EC2 instance information from only the us-east-2c Availability Zone:

```
AWS ec2 describe-instances --filters Name=availability-zone,Values=us-east-2c
```

AWS Local Zones

AWS Local Zones act similarly to Availability Zones within their respective AWS Region in that they can be selected as a placement location for zonal AWS resources such as subnets and EC2 instances. What makes them special is that they are located not in the associated AWS Region, but near large population, industry, and IT centers where no AWS Region exists today. Yet they still retain high-bandwidth, secure connection between local workloads in the local zone and those running in the AWS Region. You should use AWS Local Zones to deploy workloads closer to your users for low-latency requirements.

Amazon Global Edge Network

Amazon Global Edge Network consists of edge locations in cities around the world. Amazon CloudFront uses this network to deliver content to end users with lower latency. AWS Global Accelerator enables you to create your workload endpoints in these edge locations to provide onboarding to the AWS global network close to your users. Amazon API Gateway enables edge-optimized API endpoints using a CloudFront distribution to facilitate client access through the closest edge location.

AWS Regions

AWS Regions are designed to be autonomous, therefore, to use a multi-Region approach you would deploy dedicated copies of services to each Region.

A multi-Region approach is common for *disaster recovery* strategies to meet recovery objectives when one-off large-scale events occur. See [Plan for Disaster Recovery \(DR\)](#) for more information on these strategies. Here however, we focus instead on *availability*, which seeks to deliver a mean uptime objective over time. For high-availability objectives, a multi-region architecture will generally be designed to be active/active, where each service copy (in their respective regions) is active (serving requests).

Recommendation

Availability goals for most workloads can be satisfied using a Multi-AZ strategy within a single AWS Region. Consider multi-Region architectures only when workloads have extreme availability requirements, or other business goals, that require a multi-Region architecture.

AWS provides you with the capabilities to operate services cross-region. For example, AWS provides continuous, asynchronous data replication of data using Amazon Simple Storage Service (Amazon S3) Replication, Amazon RDS Read Replicas (including Aurora Read Replicas), and Amazon DynamoDB Global Tables. With continuous replication, versions of your data are available for near immediate use in each of your active Regions.

Using AWS CloudFormation, you can define your infrastructure and deploy it consistently across AWS accounts and across AWS Regions. And AWS CloudFormation StackSets extends this functionality by enabling you to create, update, or delete AWS CloudFormation stacks across multiple accounts and regions with a single operation. For Amazon EC2 instance deployments, an AMI (Amazon Machine Image) is used to supply information such as hardware configuration and installed software. You can implement an Amazon EC2 Image Builder pipeline that creates the AMIs you need and copy these to your active regions. This ensures that these *Golden AMIs* have everything you need to deploy and scale-out your workload in each new region.

To route traffic, both Amazon Route 53 and AWS Global Accelerator enable the definition of policies that determine which users go to which active regional endpoint. With Global Accelerator you set a traffic dial to control the percentage of traffic that is directed to each application endpoint. Route 53 supports this percentage approach, and also multiple other available policies including geoproximity and latency based ones. Global Accelerator automatically leverages the extensive network of AWS edge servers, to onboard traffic to the AWS network backbone as soon as possible, resulting in lower request latencies.

All of these capabilities operate so as to preserve each Region's autonomy. There are very few exceptions to this approach, including our services that provide global edge delivery (such as Amazon CloudFront and Amazon Route 53), along with the control plane for the AWS Identity and Access Management (IAM) service. Most services operate entirely within a single Region.

On-premises data center

For workloads that run in an on-premises data center, architect a hybrid experience when possible. AWS Direct Connect provides a dedicated network connection from your premises to AWS enabling you to run in both.

Another option is to run AWS infrastructure and services on premises using AWS Outposts. AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to your data center. The same hardware infrastructure used in the AWS Cloud is installed in your data center. AWS Outposts are then connected to the nearest AWS Region. You can then use AWS Outposts to support your workloads that have low latency or local data processing requirements.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use multiple Availability Zones and AWS Regions. Distribute workload data and resources across multiple Availability Zones or, where necessary, across AWS Regions. These locations can be as diverse as required.
- Regional services are inherently deployed across Availability Zones.
 - This includes Amazon S3, Amazon DynamoDB, and AWS Lambda (when not connected to a VPC)

- Deploy your container, instance, and function-based workloads into multiple Availability Zones. Use multi-zone datastores, including caches. Use the features of EC2 Auto Scaling, ECS task placement, AWS Lambda function configuration when running in your VPC, and ElastiCache clusters.
 - Use subnets that are in separate Availability Zones when you deploy Auto Scaling groups.
 - [Example: Distributing instances across Availability Zones](#)
 - [Amazon ECS task placement strategies](#)
 - [Configuring an AWS Lambda function to access resources in an Amazon VPC](#)
 - [Choosing Regions and Availability Zones](#)
 - Use subnets in separate Availability Zones when you deploy Auto Scaling groups.
 - [Example: Distributing instances across Availability Zones](#)
 - Use ECS task placement parameters, specifying DB subnet groups.
 - [Amazon ECS task placement strategies](#)
 - Use subnets in multiple Availability Zones when you configure a function to run in your VPC.
 - [Configuring an AWS Lambda function to access resources in an Amazon VPC](#)
 - Use multiple Availability Zones with ElastiCache clusters.
 - [Choosing Regions and Availability Zones](#)
- If your workload must be deployed to multiple Regions, choose a multi-Region strategy. Most reliability needs can be met within a single AWS Region using a multi-Availability Zone strategy. Use a multi-Region strategy when necessary to meet your business needs.
 - [AWS re:Invent 2018: Architecture Patterns for Multi-Region Active-Active Applications \(ARC209-R2\)](#)
 - Backup to another AWS Region can add another layer of assurance that data will be available when needed.
 - Some workloads have regulatory requirements that require use of a multi-Region strategy.
- Evaluate AWS Outposts for your workload. If your workload requires low latency to your on-premises data center or has local data processing requirements. Then run AWS infrastructure and services on premises using AWS Outposts
 - [What is AWS Outposts?](#)
- Determine if AWS Local Zones helps you provide service to your users. If you have low-latency requirements, see if AWS Local Zones is located near your users. If yes, then use it to deploy workloads closer to those users.
 - [AWS Local Zones FAQ](#)

Resources

Related documents:

- [AWS Global Infrastructure](#)
- [AWS Local Zones FAQ](#)
- [Amazon ECS task placement strategies](#)
- [Choosing Regions and Availability Zones](#)
- [Example: Distributing instances across Availability Zones](#)
- [Global Tables: Multi-Region Replication with DynamoDB](#)
- [Using Amazon Aurora global databases](#)
- [Creating a Multi-Region Application with AWS Services blog series](#)
- [What is AWS Outposts?](#)

Related videos:

- [AWS re:Invent 2018: Architecture Patterns for Multi-Region Active-Active Applications \(ARC209-R2\)](#)

- [AWS re:Invent 2019: Innovation and operation of the AWS global network infrastructure \(NET339\)](#)

REL10-BP02 Select the appropriate locations for your multi-location deployment

Desired Outcome

For high availability, always (when possible) deploy your workload components to multiple Availability Zones (AZs), as shown in Figure 10. For workloads with extreme resilience requirements, carefully evaluate the options for a multi-Region architecture.




Figure 10: A resilient multi-AZ database deployment with backup to another AWS Region

Common anti-patterns

- Choosing to design a multi-Region architecture when a multi-AZ architecture would satisfy requirements.
- Not accounting for dependencies between application components if resilience and multi-location requirements differ between those components.

Benefits of establishing this best practice

For resilience, you should use an approach that builds layers of defense. One layer protects against smaller, more common, disruptions by building a highly available architecture using multiple AZs. Another layer of defense is meant to protect against rare events like widespread natural disasters and Region-level disruptions. This second layer involves architecting your application to span multiple AWS Regions.

- The difference between a 99.5% availability and 99.99% availability is over 3.5 hours per month. The expected availability of a workload can only reach “four nines” if it is in multiple AZs.
- By running your workload in multiple AZs, you can isolate faults in power, cooling, and networking, and most natural disasters like fire and flood.
- Implementing a multi-Region strategy for your workload helps protect it against widespread natural disasters that affect a large geographic region of a country, or technical failures of Region-wide scope. Be aware that implementing a multi-Region architecture can be significantly complex, and is usually not required for most workloads.

Level of risk exposed if this best practice is not established: High

Implementation guidance

For a disaster event based on disruption or partial loss of one Availability Zone, implementing a highly available workload in multiple Availability Zones within a single AWS Region helps mitigate against

natural and technical disasters. Each AWS Region is comprised of multiple Availability Zones, each isolated from faults in the other zones and separated by a meaningful distance. However, for a disaster event that includes the risk of losing multiple Availability Zone components, which are a significant distance away from each other, you should implement disaster recovery options to mitigate against failures of a Region-wide scope. For workloads that require extreme resilience (critical infrastructure, health-related applications, financial system infrastructure, etc.), a multi-Region strategy may be required.

Implementation Steps

1. Evaluate your workload and determine whether the resilience needs can be met by a multi-AZ approach (single AWS Region), or if they require a multi-Region approach. Implementing a multi-Region architecture to satisfy these requirements will introduce additional complexity, therefore carefully consider your use case and its requirements. Resilience requirements can almost always be met using a single AWS Region. Consider the following possible requirements when determining whether you need to use multiple Regions:
 - a. **Disaster recovery (DR):** For a disaster event based on disruption or partial loss of one Availability Zone, implementing a highly available workload in multiple Availability Zones within a single AWS Region helps mitigate against natural and technical disasters. For a disaster event that includes the risk of losing multiple Availability Zone components, which are a significant distance away from each other, you should implement disaster recovery across multiple Regions to mitigate against natural disasters or technical failures of a Region-wide scope.
 - b. **High availability (HA):** A multi-Region architecture (using multiple AZs in each Region) can be used to achieve greater than four 9's (> 99.99%) availability.
 - c. **Stack localization:** When deploying a workload to a global audience, you can deploy localized stacks in different AWS Regions to serve audiences in those Regions. Localization can include language, currency, and types of data stored.
 - d. **Proximity to users:** When deploying a workload to a global audience, you can reduce latency by deploying stacks in AWS Regions close to where the end users are.
 - e. **Data residency:** Some workloads are subject to data residency requirements, where data from certain users must remain within a specific country's borders. Based on the regulation in question, you can choose to deploy an entire stack, or just the data, to the AWS Region within those borders.
2. Here are some examples of multi-AZ functionality provided by AWS services:
 - a. To protect workloads using EC2 or ECS, deploy an Elastic Load Balancer in front of the compute resources. Elastic Load Balancing then provides the solution to detect instances in unhealthy zones and route traffic to the healthy ones.
 - i. [Getting started with Application Load Balancers](#)
 - ii. [Getting started with Network Load Balancers](#)
 - b. In the case of EC2 instances running commercial off-the-shelf software that do not support load balancing, you can achieve a form of fault tolerance by implementing a multi-AZ disaster recovery methodology.
 - i. [the section called "REL13-BP02 Use defined recovery strategies to meet the recovery objectives" \(p. 265\)](#)
 - c. For Amazon ECS tasks, deploy your service evenly across three AZs to achieve a balance of availability and cost.
 - i. [Amazon ECS availability best practices | Containers](#)
 - d. For non-Aurora Amazon RDS, you can choose Multi-AZ as a configuration option. Upon failure of the primary database instance, Amazon RDS automatically promotes a standby database to receive traffic in another availability zone. Multi-Region read-replicas can also be created to improve resilience.
 - i. [Amazon RDS Multi AZ Deployments](#)
 - ii. [Creating a read replica in a different AWS Region](#)
3. Here are some examples of multi-Region functionality provided by AWS services:

- a. For Amazon S3 workloads, where multi-AZ availability is provided automatically by the service, consider Multi-Region Access Points if a multi-Region deployment is needed.
 - i. [Multi-Region Access Points in Amazon S3](#)
- b. For DynamoDB tables, where multi-AZ availability is provided automatically by the service, you can easily convert existing tables to global tables to take advantage of multiple regions.
 - i. [Convert Your Single-Region Amazon DynamoDB Tables to Global Tables](#)
- c. If your workload is fronted by Application Load Balancers or Network Load Balancers, use AWS Global Accelerator to improve the availability of your application by directing traffic to multiple regions that contain healthy endpoints.
 - i. [Endpoints for standard accelerators in AWS Global Accelerator - AWS Global Accelerator \(amazon.com\)](#)
- d. For applications that leverage AWS EventBridge, consider cross-Region buses to forward events to other Regions you select.
 - i. [Sending and receiving Amazon EventBridge events between AWS Regions](#)
- e. For Amazon Aurora databases, consider Aurora global databases, which span multiple AWS regions. Existing clusters can be modified to add new Regions as well.
 - i. [Getting started with Amazon Aurora global databases](#)
- f. If your workload includes AWS Key Management Service (AWS KMS) encryption keys, consider whether multi-Region keys are appropriate for your application.
 - i. [Multi-Region keys in AWS KMS](#)
- g. For other AWS service features, see this blog series on [Creating a Multi-Region Application with AWS Services series](#)

Level of effort for the Implementation Plan: Moderate to High

Resources

Related Documents:

- [Creating a Multi-Region Application with AWS Services series](#)
- [Disaster Recovery \(DR\) Architecture on AWS, Part IV: Multi-site Active/Active](#)
- [AWS Global Infrastructure](#)
- [AWS Local Zones FAQ](#)
- [Disaster Recovery \(DR\) Architecture on AWS, Part I: Strategies for Recovery in the Cloud](#)
- [Disaster recovery is different in the cloud](#)
- [Global Tables: Multi-Region Replication with DynamoDB](#)

Related Videos:

- [AWS re:Invent 2018: Architecture Patterns for Multi-Region Active-Active Applications \(ARC209-R2\)](#)
- [Auth0: Multi-Region High-Availability Architecture that Scales to 1.5B+ Logins a Month with automated failover](#)

Related Examples:

- [Disaster Recovery \(DR\) Architecture on AWS, Part I: Strategies for Recovery in the Cloud](#)
- [DTCC achieves resilience well beyond what they can do on premises](#)
- [Expedia Group uses a multi-Region, multi-Availability Zone architecture with a proprietary DNS service to add resilience to the applications](#)

- [Uber: Disaster Recovery for Multi-Region Kafka](#)
- [Netflix: Active-Active for Multi-Regional Resilience](#)
- [How we build Data Residency for Atlassian Cloud](#)
- [Intuit TurboTax runs across two Regions](#)

REL10-BP03 Automate recovery for components constrained to a single location

If components of the workload can only run in a single Availability Zone or in an on-premises data center, you must implement the capability to do a complete rebuild of the workload within your defined recovery objectives.

If the best practice to deploy the workload to multiple locations is not possible due to technological constraints, you must implement an alternate path to resiliency. You must automate the ability to recreate necessary infrastructure, redeploy applications, and recreate necessary data for these cases.

For example, Amazon EMR launches all nodes for a given cluster in the same Availability Zone because running a cluster in the same zone improves performance of the jobs flows as it provides a higher data access rate. If this component is required for workload resilience, then you must have a way to redeploy the cluster and its data. Also for Amazon EMR, you should provision redundancy in ways other than using Multi-AZ. You can provision [multiple nodes](#). Using [EMR File System \(EMRFS\)](#), data in EMR can be stored in Amazon S3, which in turn can be replicated across multiple Availability Zones or AWS Regions.

Similarly, for Amazon Redshift, by default it provisions your cluster in a randomly selected Availability Zone within the AWS Region that you select. All the cluster nodes are provisioned in the same zone.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Implement self-healing. Deploy your instances or containers using automatic scaling when possible. If you cannot use automatic scaling, use automatic recovery for EC2 instances or implement self-healing automation based on Amazon EC2 or ECS container lifecycle events.
- Use Auto Scaling groups for instances and container workloads that have no requirements for a single instance IP address, private IP address, Elastic IP address, and instance metadata.
 - [What Is EC2 Auto Scaling?](#)
 - [Service automatic scaling](#)
 - The launch template user data can be used to implement automation that can self-heal most workloads.
- Use automatic recovery of EC2 instances for workloads that require a single instance ID address, private IP address, Elastic IP address, and instance metadata.
 - [Recover your instance](#).
 - Automatic Recovery will send recovery status alerts to a SNS topic as the instance failure is detected.
- Use EC2 instance lifecycle events or ECS events to automate self-healing where automatic scaling or EC2 recovery cannot be used.
 - [EC2 Auto Scaling lifecycle hooks](#)
 - [Amazon ECS events](#)
 - Use the events to invoke automation that will heal your component according to the process logic you require.

Resources

Related documents:

- [Amazon ECS events](#)
- [EC2 Auto Scaling lifecycle hooks](#)
- [Recover your instance.](#)
- [Service automatic scaling](#)
- [What Is EC2 Auto Scaling?](#)

REL10-BP04 Use bulkhead architectures to limit scope of impact

Like the bulkheads on a ship, this pattern ensures that a failure is contained to a small subset of requests or clients so that the number of impaired requests is limited, and most can continue without error. Bulkheads for data are often called partitions, while bulkheads for services are known as cells.

In a *cell-based architecture*, each cell is a complete, independent instance of the service and has a fixed maximum size. As load increases, workloads grow by adding more cells. A partition key is used on incoming traffic to determine which cell will process the request. Any failure is contained to the single cell it occurs in, so that the number of impaired requests is limited as other cells continue without error. It is important to identify the proper partition key to minimize cross-cell interactions and avoid the need to involve complex mapping services in each request. Services that require complex mapping end up merely shifting the problem to the mapping services, while services that require cross-cell interactions create dependencies between cells (and thus reduce the assumed availability improvements of doing so).




Figure 11: Cell-based architecture

In his AWS blog post, Colm MacCarthaigh explains how Amazon Route 53 uses the concept of [shuffle sharding](#) to isolate customer requests into shards. A shard in this case consists of two or more cells. Based on partition key, traffic from a customer (or resources, or whatever you want to isolate) is routed to its assigned shard. In the case of eight cells with two cells per shard, and customers divided among the four shards, 25% of customers would experience impact in the event of a problem.




Figure 12: Service divided into four traditional shards of two cells each

With shuffle sharding, you create virtual shards of two cells each, and assign your customers to one of those virtual shards. When a problem happens, you can still lose a quarter of the whole service, but the way that customers or resources are assigned means that the scope of impact with shuffle sharding is considerably smaller than 25%. With eight cells, there are 28 unique combinations of two cells, which means that there are 28 possible shuffle shards (virtual shards). If you have hundreds or thousands of customers, and assign each customer to a shuffle shard, then the scope of impact due to a problem is just 1/28th. That's seven times better than regular sharding.




Figure 13: Service divided into 28 shuffle shards (virtual shards) of two cells each (only two shuffle shards out of the 28 possible are shown)

A shard can be used for servers, queues, or other resources in addition to cells.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use bulkhead architectures. Like the bulkheads on a ship, this pattern ensures that a failure is contained to a small subset of requests or users so that the number of impaired requests is limited, and most can continue without error. Bulkheads for data are often called partitions, while bulkheads for services are known as cells.
 - [Well-Architected lab: Fault isolation with shuffle sharding](#)
 - [Shuffle-sharding: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)
 - [AWS re:Invent 2018: How AWS Minimizes the Blast Radius of Failures \(ARC338\)](#)

- Evaluate cell-based architecture for your workload. In a cell-based architecture, each cell is a complete, independent instance of the service and has a fixed maximum size. As load increases, workloads grow by adding more cells. A partition key is used on incoming traffic to determine which cell will process the request. Any failure is contained to the single cell it occurs in, so that the number of impaired requests is limited as other cells continue without error. It is important to identify the proper partition key to minimize cross-cell interactions and avoid the need to involve complex mapping services in each request. Services that require complex mapping end up merely shifting the problem to the mapping services, while services that require cross-cell interactions reduce the autonomy of cells (and thus the assumed availability improvements of doing so).
- In his AWS blog post, Colm MacCarthaigh explains how Amazon Route 53 uses the concept of shuffle sharding to isolate customer requests into shards
 - [Shuffle Sharding: Massive and Magical Fault Isolation](#)

Resources

Related documents:

- [Shuffle Sharding: Massive and Magical Fault Isolation](#)
- [The Amazon Builders' Library: Workload isolation using shuffle-sharding](#)

Related videos:

- [AWS re:Invent 2018: How AWS Minimizes the Blast Radius of Failures \(ARC338\)](#)
- [Shuffle-sharding: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)

Related examples:

- [Well-Architected lab: Fault isolation with shuffle sharding](#)

REL 11 How do you design your workload to withstand component failures?

Workloads with a requirement for high availability and low mean time to recovery (MTTR) must be architected for resiliency.

Best practices

- [REL11-BP01 Monitor all components of the workload to detect failures \(p. 240\)](#)
- [REL11-BP02 Fail over to healthy resources \(p. 242\)](#)
- [REL11-BP03 Automate healing on all layers \(p. 244\)](#)
- [REL11-BP04 Rely on the data plane and not the control plane during recovery \(p. 246\)](#)
- [REL11-BP05 Use static stability to prevent bimodal behavior \(p. 247\)](#)
- [REL11-BP06 Send notifications when events impact availability \(p. 248\)](#)

REL11-BP01 Monitor all components of the workload to detect failures

Continuously monitor the health of your workload so that you and your automated systems are aware of degradation or failure as soon as they occur. Monitor for key performance indicators (KPIs) based on business value.

All recovery and healing mechanisms must start with the ability to detect problems quickly. Technical failures should be detected first so that they can be resolved. However, availability is based on the ability

of your workload to deliver business value, so key performance indicators (KPIs) that measure this need to be a part of your detection and remediation strategy.

Common anti-patterns:

- No alarms have been configured, so outages occur without notification.
- Alarms exist, but at thresholds that don't provide adequate time to react.
- Metrics are not collected often enough to meet the recovery time objective (RTO).
- Only the customer facing tier of the workload is actively monitored.
- Only collecting technical metrics, no business function metrics.
- No metrics measuring the user experience of the workload.

Benefits of establishing this best practice: Having appropriate monitoring at all layers enables you to reduce recovery time by reducing time to detection.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Determine the collection interval for your components based on your recovery goals.
 - Your monitoring interval is dependent on how quickly you must recover. Your recovery time is driven by the time it takes to recover, so you must determine the frequency of collection by accounting for this time and your recovery time objective (RTO).
- Configure detailed monitoring for components.
 - Determine if detailed monitoring for EC2 instances and Auto Scaling is necessary. Detailed monitoring provides 1-min interval metrics, and default monitoring provides 5-minute interval metrics.
 - [Enable or Disable Detailed Monitoring for Your Instance](#)
 - [Monitoring Your Auto Scaling Groups and Instances Using Amazon CloudWatch](#)
 - Determine if enhanced monitoring for RDS is necessary. Enhanced monitoring uses an agent on the RDS instances to get useful information about different process or threads on an RDS instance.
 - [Enhanced Monitoring](#)
- Create custom metrics to measure business key performance indicators (KPIs). Workloads implement key business functions. These functions should be used as KPIs that help identify when an indirect problem happens.
 - [Publishing Custom Metrics](#)
- Monitor the user experience for failures using user canaries. Synthetic transaction testing (also known as canary testing, but not to be confused with canary deployments) that can run and simulate customer behavior is among the most important testing processes. Run these tests constantly against your workload endpoints from diverse remote locations.
 - [Amazon CloudWatch Synthetics enables you to create user canaries](#)
- Create custom metrics that track the user's experience. If you can instrument the experience of the customer, you can determine when the consumer experience degrades.
 - [Publishing Custom Metrics](#)
- Set alarms to detect when any part of your workload is not working properly, and to indicate when to Auto Scale resources. Alarms can be visually displayed on dashboards, send alerts via Amazon SNS or email, and work with Auto Scaling to scale up or down the resources for a workload.
 - [Using Amazon CloudWatch Alarms](#)
- Create dashboards to visualize your metrics. Dashboards can be used to visually see trends, outliers, and other indicators of potential problems, or to provide an indication of problems you may want to investigate.

- [Using CloudWatch Dashboards](#)

Resources

Related documents:

- [Amazon CloudWatch Synthetics enables you to create user canaries](#)
- [Enable or Disable Detailed Monitoring for Your Instance](#)
- [Enhanced Monitoring](#)
- [Monitoring Your Auto Scaling Groups and Instances Using Amazon CloudWatch](#)
- [Publishing Custom Metrics](#)
- [Using Amazon CloudWatch Alarms](#)
- [Using CloudWatch Dashboards](#)

Related examples:

- [Well-Architected lab: Level 300: Implementing Health Checks and Managing Dependencies to Improve Reliability](#)

REL11-BP02 Fail over to healthy resources

Ensure that if a resource failure occurs, that healthy resources can continue to serve requests. For location failures (such as Availability Zone or AWS Region) ensure that you have systems in place to fail over to healthy resources in unimpaired locations.

AWS services, such as Elastic Load Balancing and Amazon EC2 Auto Scaling, help distribute load across resources and Availability Zones. Therefore, failure of an individual resource (such as an EC2 instance) or impairment of an Availability Zone can be mitigated by shifting traffic to remaining healthy resources. For multi-region workloads, this is more complicated. For example, cross-region read replicas enable you to deploy your data to multiple AWS Regions, but you still must promote the read replica to primary and point your traffic at it in the event of a failover. Amazon Route 53 and AWS Global Accelerator can help route traffic across AWS Regions.

If your workload is using AWS services, such as Amazon S3 or Amazon DynamoDB, then they are automatically deployed to multiple Availability Zones. In case of failure, the AWS control plane automatically routes traffic to healthy locations for you. Data is redundantly stored in multiple Availability Zones, and remains available. For Amazon RDS, you must choose Multi-AZ as a configuration option, and then on failure AWS automatically directs traffic to the healthy instance. For Amazon EC2 instances, Amazon ECS tasks, or Amazon EKS pods, you choose which Availability Zones to deploy to. Elastic Load Balancing then provides the solution to detect instances in unhealthy zones and route traffic to the healthy ones. Elastic Load Balancing can even route traffic to components in your on-premises data center.

For Multi-Region approaches (which might also include on-premises data centers), Amazon Route 53 provides a way to define internet domains, and assign routing policies that can include health checks to ensure that traffic is routed to healthy regions. Alternately, AWS Global Accelerator provides static IP addresses that act as a fixed entry point to your application, then routes to endpoints in AWS Regions of your choosing, using the AWS global network instead of the internet for better performance and reliability.

AWS approaches the design of our services with fault recovery in mind. We design services to minimize the time to recover from failures and impact on data. Our services primarily use data stores that acknowledge requests only after they are durably stored across multiple replicas within a Region. These services and resources include Amazon Aurora, Amazon Relational Database Service (Amazon RDS) Multi-

AZ DB instances, Amazon S3, Amazon DynamoDB, Amazon Simple Queue Service (Amazon SQS), and Amazon Elastic File System (Amazon EFS). They are constructed to use cell-based isolation and use the fault isolation provided by Availability Zones. We use automation extensively in our operational procedures. We also optimize our replace-and-restart functionality to recover quickly from interruptions.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Fail over to healthy resources. Ensure that if a resource failure occurs, that healthy resources can continue to serve requests. For location failures (such as Availability Zone or AWS Region) ensure you have systems in place to fail over to healthy resources in unimpaired locations.
- If your workload is using AWS services, such as Amazon S3 or Amazon DynamoDB, then they are automatically deployed to multiple Availability Zones. In case of failure, the AWS control plane automatically routes traffic to healthy locations for you.
- For Amazon RDS you must choose Multi-AZ as a configuration option, and then on failure AWS automatically directs traffic to the healthy instance.
 - [High Availability \(Multi-AZ\) for Amazon RDS](#)
- For Amazon EC2 instances or Amazon ECS tasks, you choose which Availability Zones to deploy to. Elastic Load Balancing then provides the solution to detect instances in unhealthy zones and route traffic to the healthy ones. Elastic Load Balancing can even route traffic to components in your on-premises data center.
- For multi-region approaches (which might also include on-premises data centers), ensure that data and resources from healthy locations can continue to serve requests
 - For example, cross-region read replicas enable you to deploy your data to multiple AWS Regions, but you still must promote the read replica to master and point your traffic at it in the event of a primary location failure.
 - [Overview of Amazon RDS Read Replicas](#)
 - Amazon Route 53 provides a way to define internet domains, and assign routing policies, which might include health checks, to ensure that traffic is routed to healthy Regions. Alternately, AWS Global Accelerator provides static IP addresses that act as a fixed entry point to your application, then routes to endpoints in AWS Regions of your choosing, using the AWS global network instead of the public internet for better performance and reliability.
 - [Amazon Route 53: Choosing a Routing Policy](#)
 - [What Is AWS Global Accelerator?](#)

Resources

Related documents:

- [APN Partner: partners that can help with automation of your fault tolerance](#)
- [AWS Marketplace: products that can be used for fault tolerance](#)
- [AWS OpsWorks: Using Auto Healing to Replace Failed Instances](#)
- [Amazon Route 53: Choosing a Routing Policy](#)
- [High Availability \(Multi-AZ\) for Amazon RDS](#)
- [Overview of Amazon RDS Read Replicas](#)
- [Amazon ECS task placement strategies](#)
- [Creating Kubernetes Auto Scaling Groups for Multiple Availability Zones](#)
- [What is AWS Global Accelerator?](#)

Related examples:

- [Well-Architected lab: Level 300: Implementing Health Checks and Managing Dependencies to Improve Reliability](#)

REL11-BP03 Automate healing on all layers

Upon detection of a failure, use automated capabilities to perform actions to remediate.

Ability to restart is an important tool to remediate failures. As discussed previously for distributed systems, a best practice is to make services stateless where possible. This prevents loss of data or availability on restart. In the cloud, you can (and generally should) replace the entire resource (for example, EC2 instance, or Lambda function) as part of the restart. The restart itself is a simple and reliable way to recover from failure. Many different types of failures occur in workloads. Failures can occur in hardware, software, communications, and operations. Rather than constructing novel mechanisms to trap, identify, and correct each of the different types of failures, map many different categories of failures to the same recovery strategy. An instance might fail due to hardware failure, an operating system bug, memory leak, or other causes. Rather than building custom remediation for each situation, treat any of them as an instance failure. Terminate the instance, and allow AWS Auto Scaling to replace it. Later, carry out the analysis on the failed resource out of band.

Another example is the ability to restart a network request. Apply the same recovery approach to both a network timeout and a dependency failure where the dependency returns an error. Both events have a similar effect on the system, so rather than attempting to make either event a “special case”, apply a similar strategy of limited retry with exponential backoff and jitter.

Ability to restart is a recovery mechanism featured in Recovery Oriented Computing and high availability cluster architectures.

Amazon EventBridge can be used to monitor and filter for events such as CloudWatch Alarms or changes in state in other AWS services. Based on event information, it can then trigger AWS Lambda, AWS Systems Manager Automation, or other targets to execute custom remediation logic on your workload.

Amazon EC2 Auto Scaling can be configured to check for EC2 instance health. If the instance is in any state other than running, or if the system status is impaired, Amazon EC2 Auto Scaling considers the instance to be unhealthy and launches a replacement instance. If using AWS OpsWorks, you can configure Auto Healing of EC2 instances at the OpsWorks layer level.

For large-scale replacements (such as the loss of an entire Availability Zone), static stability is preferred for high availability instead of trying to obtain multiple new resources at once.

Common anti-patterns:

- Deploying applications in instances or containers individually.
- Deploying applications that cannot be deployed into multiple locations without using automatic recovery.
- Manually healing applications that automatic scaling and automatic recovery fail to heal.

Benefits of establishing this best practice: Automated healing, even if the workload can only be deployed into one location at a time will reduce your mean time to recovery, and ensure availability of the workload.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use Auto Scaling groups to deploy tiers in an workload. Auto scaling can perform self-healing on stateless applications, and add and remove capacity.
 - [How AWS Auto Scaling Works](#)

- Implement automatic recovery on EC2 instances that have applications deployed that cannot be deployed in multiple locations, and can tolerate rebooting upon failures. Automatic recovery can be used to replace failed hardware and restart the instance when the application is not capable of being deployed in multiple locations. The instance metadata and associated IP addresses are kept, as well as the Amazon EBS volumes and mount points to Elastic File Systems or File Systems for Lustre and Windows.
 - [Amazon EC2 Automatic Recovery](#)
 - [Amazon Elastic Block Store \(Amazon EBS\)](#)
 - [Amazon Elastic File System \(Amazon EFS\)](#)
 - [What is Amazon FSx for Lustre?](#)
 - [What is Amazon FSx for Windows File Server?](#)
 - Using AWS OpsWorks, you can configure Auto Healing of EC2 instances at the layer level
 - [AWS OpsWorks: Using Auto Healing to Replace Failed Instances](#)
- Implement automated recovery using AWS Step Functions and AWS Lambda when you cannot use automatic scaling or automatic recovery, or when automatic recovery fails. When you cannot use automatic scaling, and either cannot use automatic recovery or automatic recovery fails, you can automate the healing using AWS Step Functions and AWS Lambda.
 - [What is AWS Step Functions?](#)
 - [What is AWS Lambda?](#)
 - Amazon EventBridge can be used to monitor and filter for events such as CloudWatch Alarms or changes in state in other AWS services. Based on event information, it can then trigger AWS Lambda (or other targets) to run custom remediation logic on your workload.
 - [What Is Amazon EventBridge?](#)
 - [Using Amazon CloudWatch Alarms](#)

Resources

Related documents:

- [APN Partner: partners that can help with automation of your fault tolerance](#)
- [AWS Marketplace: products that can be used for fault tolerance](#)
- [AWS OpsWorks: Using Auto Healing to Replace Failed Instances](#)
- [Amazon EC2 Automatic Recovery](#)
- [Amazon Elastic Block Store \(Amazon EBS\)](#)
- [Amazon Elastic File System \(Amazon EFS\)](#)
- [How AWS Auto Scaling Works](#)
- [Using Amazon CloudWatch Alarms](#)
- [What Is Amazon EventBridge?](#)
- [What is AWS Lambda?](#)
- [AWS Systems Manager Automation](#)
- [What is AWS Step Functions?](#)
- [What is Amazon FSx for Lustre?](#)
- [What is Amazon FSx for Windows File Server?](#)

Related videos:

- [Static stability in AWS: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)

Related examples:

- [Well-Architected lab: Level 300: Implementing Health Checks and Managing Dependencies to Improve Reliability](#)

REL11-BP04 Rely on the data plane and not the control plane during recovery

The control plane is used to configure resources, and the data plane delivers services. Data planes typically have higher availability design goals than control planes and are usually less complex. When implementing recovery or mitigation responses to potentially resiliency-impacting events, using control plane operations can lower the overall resiliency of your architecture. For example, you can rely on the Amazon Route 53 data plane to reliably route DNS queries based on health checks, but updating Route 53 routing policies uses the control plane, so do not rely on it for recovery.

The Route 53 data planes answer DNS queries, and perform and evaluate health checks. They are globally distributed and designed for a [100% availability service level agreement \(SLA\)](#). The Route 53 management APIs and consoles where you create, update, and delete Route 53 resources run on control planes that are designed to prioritize the strong consistency and durability that you need when managing DNS. To achieve this, the control planes are located in a single Region, US East (N. Virginia). While both systems are built to be very reliable, the control planes are not included in the SLA. There could be rare events in which the data plane's resilient design allows it to maintain availability while the control planes do not. For disaster recovery and failover mechanisms, use data plane functions to provide the best possible reliability.

For more information about data planes, control planes, and how AWS builds services to meet high availability targets, see the [Static stability using Availability Zones](#) paper and the [Amazon Builders' Library](#).

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Rely on the data plane and not the control plane when using Amazon Route 53 for disaster recovery. Route 53 Application Recovery Controller helps you manage and coordinate failover using readiness checks and routing controls. These features continually monitor your application's ability to recover from failures, and enables you to control your application recovery across multiple AWS Regions, Availability Zones, and on premises.
 - [What is Route 53 Application Recovery Controller](#)
 - [Creating Disaster Recovery Mechanisms Using Amazon Route 53](#)
 - [Building highly resilient applications using Amazon Route 53 Application Recovery Controller, Part 1: Single-Region stack](#)
 - [Building highly resilient applications using Amazon Route 53 Application Recovery Controller, Part 2: Multi-Region stack](#)
- Understand which operations are on the data plane and which are on the control plane.
 - [Amazon Builders' Library: Avoiding overload in distributed systems by putting the smaller service in control](#)
 - [Amazon DynamoDB API \(control plane and data plane\)](#)
 - [AWS Lambda Executions](#) (split into the control plane and the data plane)
 - [AWS Lambda Executions](#) (split into the control plane and the data plane)

Resources

Related documents:

- [APN Partner: partners that can help with automation of your fault tolerance](#)
- [AWS Marketplace: products that can be used for fault tolerance](#)

- [Amazon Builders' Library: Avoiding overload in distributed systems by putting the smaller service in control](#)
- [Amazon DynamoDB API \(control plane and data plane\)](#)
- [AWS Lambda Executions \(split into the control plane and the data plane\)](#)
- [AWS Elemental MediaStore Data Plane](#)
- [Building highly resilient applications using Amazon Route 53 Application Recovery Controller, Part 1: Single-Region stack](#)
- [Building highly resilient applications using Amazon Route 53 Application Recovery Controller, Part 2: Multi-Region stack](#)
- [Creating Disaster Recovery Mechanisms Using Amazon Route 53](#)
- [What is Route 53 Application Recovery Controller](#)

Related examples:

- [Introducing Amazon Route 53 Application Recovery Controller](#)

REL11-BP05 Use static stability to prevent bimodal behavior

Bimodal behavior is when your workload exhibits different behavior under normal and failure modes, for example, relying on launching new instances if an Availability Zone fails. You should instead build workloads that are statically stable and operate in only one mode. In this case, provision enough instances in each Availability Zone to handle the workload load if one AZ were removed and then use Elastic Load Balancing or Amazon Route 53 health checks to shift load away from the impaired instances.

Static stability for compute deployment (such as EC2 instances or containers) will result in the highest reliability. This must be weighed against cost concerns. It's less expensive to provision less compute capacity and rely on launching new instances in the case of a failure. But for large-scale failures (such as an Availability Zone failure) this approach is less effective because it relies on reacting to impairments as they happen, rather than being prepared for those impairments before they happen. Your solution should weigh reliability versus the cost needs for your workload. By using more Availability Zones, the amount of additional compute you need for static stability decreases.



Figure 14: Static stability of EC2 instances across Availability Zones

After traffic has shifted, use AWS Auto Scaling to asynchronously replace instances from the failed zone and launch them in the healthy zones.

Another example of bimodal behavior would be a network timeout that could cause a system to attempt to refresh the configuration state of the entire system. This would add unexpected load to another component, and might cause it to fail, triggering other unexpected consequences. This negative feedback loop impacts availability of your workload. Instead, you should build systems that are statically stable and operate in only one mode. A statically stable design would be to do constant work, and always refresh the configuration state on a fixed cadence. When a call fails, the workload uses the previously cached value, and triggers an alarm.

Another example of bimodal behavior is allowing clients to bypass your workload cache when failures occur. This might seem to be a solution that accommodates client needs, but should not be allowed because it significantly changes the demands on your workload and is likely to result in failures.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Use static stability to prevent bimodal behavior. Bimodal behavior is when your workload exhibits different behavior under normal and failure modes, for example, relying on launching new instances if an Availability Zone fails.
 - [Minimizing Dependencies in a Disaster Recovery Plan](#)
 - [The Amazon Builders' Library: Static stability using Availability Zones](#)
 - [Static stability in AWS: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)
 - You should instead build systems that are statically stable and operate in only one mode. In this case, provision enough instances in each zone to handle workload load if one AZ were removed and then use Elastic Load Balancing or Amazon Route 53 health checks to shift load away from the impaired instances.
 - Another example of bimodal behavior is allowing clients to bypass your workload cache when failures occur. This might seem to be a solution to accommodate client needs, but should not be allowed since it significantly changes demands on your workload and is likely to result in failures.

Resources

Related documents:

- [Minimizing Dependencies in a Disaster Recovery Plan](#)
- [The Amazon Builders' Library: Static stability using Availability Zones](#)

Related videos:

- [Static stability in AWS: AWS re:Invent 2019: Introducing The Amazon Builders' Library \(DOP328\)](#)

REL11-BP06 Send notifications when events impact availability

Notifications are sent upon the detection of significant events, even if the issue caused by the event was automatically resolved.

Automated healing enables your workload to be reliable. However, it can also obscure underlying problems that need to be addressed. Implement appropriate monitoring and events so that you can detect patterns of problems, including those addressed by auto healing, so that you can resolve root cause issues. Amazon CloudWatch Alarms can be triggered based on failures that occur. They can also trigger based on automated healing actions executed. CloudWatch Alarms can be configured to send emails, or to log incidents in third-party incident tracking systems using Amazon SNS integration.

Common anti-patterns:

- Sending alarms that no one acts upon.
- Performing auto healing automation, but not notifying that healing was needed.

Benefits of establishing this best practice: Notifications of recovery events will ensure that you don't ignore problems that occur infrequently.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Alarms on business Key Performance Indicators when they exceed a low threshold Having a low threshold alarm on your business KPIs help you know when your workload is unavailable or non-functional.
 - [Creating a CloudWatch Alarm Based on a Static Threshold](#)
- Alarm on events that invoke healing automation You can directly invoke an SNS API to send notifications with any automation that you create.
 - [What is Amazon Simple Notification Service?](#)

Resources

Related documents:

- [Creating a CloudWatch Alarm Based on a Static Threshold](#)
- [What Is Amazon EventBridge?](#)
- [What is Amazon Simple Notification Service?](#)

REL 12 How do you test reliability?

After you have designed your workload to be resilient to the stresses of production, testing is the only way to ensure that it will operate as designed, and deliver the resiliency you expect.

Best practices

- [REL12-BP01 Use playbooks to investigate failures \(p. 249\)](#)
- [REL12-BP02 Perform post-incident analysis \(p. 251\)](#)
- [REL12-BP03 Test functional requirements \(p. 251\)](#)
- [REL12-BP04 Test scaling and performance requirements \(p. 252\)](#)
- [REL12-BP05 Test resiliency using chaos engineering \(p. 253\)](#)
- [REL12-BP06 Conduct game days regularly \(p. 260\)](#)

REL12-BP01 Use playbooks to investigate failures

Enable consistent and prompt responses to failure scenarios that are not well understood, by documenting the investigation process in playbooks. Playbooks are the predefined steps performed to identify the factors contributing to a failure scenario. The results from any process step are used to determine the next steps to take until the issue is identified or escalated.

The playbook is proactive planning that you must do, to be able to take reactive actions effectively. When failure scenarios not covered by the playbook are encountered in production, first address the issue (put out the fire). Then go back and look at the steps you took to address the issue and use these to add a new entry in the playbook.

Note that playbooks are used in response to specific incidents, while runbooks are used to achieve specific outcomes. Often, runbooks are used for routine activities and playbooks are used to respond to non-routine events.

Common anti-patterns:

- Planning to deploy a workload without knowing the processes to diagnose issues or respond to incidents.
- Unplanned decisions about which systems to gather logs and metrics from when investigating an event.
- Not retaining metrics and events long enough to be able to retrieve the data.

Benefits of establishing this best practice: Capturing playbooks ensures that processes can be consistently followed. Codifying your playbooks limits the introduction of errors from manual activity. Automating playbooks shortens the time to respond to an event by eliminating the requirement for team member intervention or providing them additional information when their intervention begins.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Use playbooks to identify issues. Playbooks are documented processes to investigate issues. Enable consistent and prompt responses to failure scenarios by documenting processes in playbooks. Playbooks must contain the information and guidance necessary for an adequately skilled person to gather applicable information, identify potential sources of failure, isolate faults, and determine contributing factors (perform post-incident analysis).
- Implement playbooks as code. Perform your operations as code by scripting your playbooks to ensure consistency and limit reduce errors caused by manual processes. Playbooks can be composed of multiple scripts representing the different steps that might be necessary to identify the contributing factors to an issue. Runbook activities can be triggered or performed as part of playbook activities, or may prompt for execution of a playbook in response to identified events.
 - [Automate your operational playbooks with AWS Systems Manager](#)
 - [AWS Systems Manager Run Command](#)
 - [AWS Systems Manager Automation](#)
 - [What is AWS Lambda?](#)
 - [What Is Amazon EventBridge?](#)
 - [Using Amazon CloudWatch Alarms](#)

Resources

Related documents:

- [AWS Systems Manager Automation](#)
- [AWS Systems Manager Run Command](#)
- [Automate your operational playbooks with AWS Systems Manager](#)
- [Using Amazon CloudWatch Alarms](#)
- [Using Canaries \(Amazon CloudWatch Synthetics\)](#)
- [What Is Amazon EventBridge?](#)
- [What is AWS Lambda?](#)

Related examples:

- [Automating operations with Playbooks and Runbooks](#)

REL12-BP02 Perform post-incident analysis

Review customer-impacting events, and identify the contributing factors and preventative action items. Use this information to develop mitigations to limit or prevent recurrence. Develop procedures for prompt and effective responses. Communicate contributing factors and corrective actions as appropriate, tailored to target audiences. Have a method to communicate these causes to others as needed.

Assess why existing testing did not find the issue. Add tests for this case if tests do not already exist.

Common anti-patterns:

- Finding contributing factors, but not continuing to look deeper for other potential problems and approaches to mitigate.
- Only identifying human error causes, and not providing any training or automation that could prevent human errors.

Benefits of establishing this best practice: Conducting post-incident analysis and sharing the results enables other workloads to mitigate the risk if they have implemented the same contributing factors, and enables them to implement the mitigation or automated recovery before an incident occurs.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Establish a standard for your post-incident analysis. Good post-incident analysis provides opportunities to propose common solutions for problems with architecture patterns that are used in other places in your systems.
 - Ensure that the contributing factors are honest and blame free.
 - If you do not document your problems, you cannot correct them.
 - Ensure post-incident analysis is blame free so you can be dispassionate about the proposed corrective actions and promote honest self-assessment and collaboration on your application teams.
- Use a process to determine contributing factors. Have a process to identify and document the contributing factors of an event so that you can develop mitigations to limit or prevent recurrence and you can develop procedures for prompt and effective responses. Communicate contributing factors as appropriate, tailored to target audiences.
 - [What is log analytics?](#)

Resources

Related documents:

- [What is log analytics?](#)
- [Why you should develop a correction of error \(COE\)](#)

REL12-BP03 Test functional requirements

Use techniques such as unit tests and integration tests that validate required functionality.

You achieve the best outcomes when these tests are run automatically as part of build and deployment actions. For instance, using AWS CodePipeline, developers commit changes to a source repository where CodePipeline automatically detects the changes. Those changes are built, and tests are run. After the

tests are complete, the built code is deployed to staging servers for testing. From the staging server, CodePipeline runs more tests, such as integration or load tests. Upon the successful completion of those tests, CodePipeline deploys the tested and approved code to production instances.

Additionally, experience shows that synthetic transaction testing (also known as *canary testing*, but not to be confused with canary deployments) that can run and simulate customer behavior is among the most important testing processes. Run these tests constantly against your workload endpoints from diverse remote locations. Amazon CloudWatch Synthetics enables you to [create canaries](#) to monitor your endpoints and APIs.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Test functional requirements. These include unit tests and integration tests that validate required functionality.
 - [Use CodePipeline with AWS CodeBuild to test code and run builds](#)
 - [AWS CodePipeline Adds Support for Unit and Custom Integration Testing with AWS CodeBuild](#)
 - [Continuous Delivery and Continuous Integration](#)
 - [Using Canaries \(Amazon CloudWatch Synthetics\)](#)
 - [Software test automation](#)

Resources

Related documents:

- [APN Partner: partners that can help with implementation of a continuous integration pipeline](#)
- [AWS CodePipeline Adds Support for Unit and Custom Integration Testing with AWS CodeBuild](#)
- [AWS Marketplace: products that can be used for continuous integration](#)
- [Continuous Delivery and Continuous Integration](#)
- [Software test automation](#)
- [Use CodePipeline with AWS CodeBuild to test code and run builds](#)
- [Using Canaries \(Amazon CloudWatch Synthetics\)](#)

REL12-BP04 Test scaling and performance requirements

Use techniques such as load testing to validate that the workload meets scaling and performance requirements.

In the cloud, you can create a production-scale test environment on demand for your workload. If you run these tests on scaled down infrastructure, you must scale your observed results to what you think will happen in production. Load and performance testing can also be done in production if you are careful not to impact actual users, and tag your test data so it does not comingle with real user data and corrupt usage statistics or production reports.

With testing, ensure that your base resources, scaling settings, service quotas, and resiliency design operate as expected under load.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Test scaling and performance requirements. Perform load testing to validate that the workload meets scaling and performance requirements.
 - [Distributed Load Testing on AWS: simulate thousands of connected users](#)

- [Apache JMeter](#)
- Deploy your application in an environment identical to your production environment and execute a load test.
- Use infrastructure as code concepts to create an environment as similar to your production environment as possible.

Resources

Related documents:

- [Distributed Load Testing on AWS: simulate thousands of connected users](#)
- [Apache JMeter](#)

REL12-BP05 Test resiliency using chaos engineering

Run chaos experiments regularly in environments that are in or as close to production as possible to understand how your system responds to adverse conditions.

Desired outcome:

The resilience of the workload is regularly verified by applying chaos engineering in the form of fault injection experiments or injection of unexpected load, in addition to resilience testing that validates known expected behavior of your workload during an event. Combine both chaos engineering and resilience testing to gain confidence that your workload can survive component failure and can recover from unexpected disruptions with minimal to no impact.

Common anti-patterns:

- Designing for resiliency, but not verifying how the workload functions as a whole when faults occur.
- Never experimenting under real-world conditions and expected load.
- Not treating your experiments as code or maintaining them through the development cycle.
- Not running chaos experiments both as part of your CI/CD pipeline, as well as outside of deployments.
- Neglecting to use past post-incident analyses when determining which faults to experiment with.

Benefits of establishing this best practice: Injecting faults to verify the resilience of your workload allows you to gain confidence that the recovery procedures of your resilient design will work in the case of a real fault.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Chaos engineering provides your teams with capabilities to continually inject real world disruptions (simulations) in a controlled way at the service provider, infrastructure, workload, and component level, with minimal to no impact to your customers. It allows your teams to learn from faults and observe, measure, and improve the resilience of your workloads, as well as validate that alerts fire and teams get notified in the case of an event.

When performed continually, chaos engineering can highlight deficiencies in your workloads that, if left unaddressed, could negatively affect availability and operation.

Note

Chaos engineering is the discipline of experimenting on a system in order to build confidence in the system's capability to withstand turbulent conditions in production. – [Principles of Chaos Engineering](#)

If a system is able to withstand these disruptions, the chaos experiment should be maintained as an automated regression test. In this way, chaos experiments should be performed as part of your systems development lifecycle (SDLC) and as part of your CI/CD pipeline.

To ensure that your workload can survive component failure, inject real world events as part of your experiments. For example, experiment with the loss of Amazon EC2 instances or failover of the primary Amazon RDS database instance, and verify that your workload is not impacted (or only minimally impacted). Use a combination of component faults to simulate events that may be caused by a disruption in an Availability Zone.


For application-level faults (such as crashes), you can start with stressors such as memory and CPU exhaustion.

To validate [fallback or failover mechanisms](#) for external dependencies due to intermittent network disruptions, your components should simulate such an event by blocking access to the third-party providers for a specified duration that can last from seconds to hours.

Other modes of degradation might cause reduced functionality and slow responses, often resulting in a disruption of your services. Common sources of this degradation are increased latency on critical services and unreliable network communication (dropped packets). Experiments with these faults, including networking effects such as latency, dropped messages, and DNS failures, could include the inability to resolve a name, reach the DNS service, or establish connections to dependent services.

Chaos engineering tools:

AWS Fault Injection Simulator (AWS FIS) is a fully managed service for running fault injection experiments that can be used as part of your CD pipeline, or outside of the pipeline. AWS FIS is a good choice to use during chaos engineering game days. It supports simultaneously introducing faults across different types of resources including Amazon EC2, Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS), and Amazon RDS. These faults include termination of resources, forcing failovers, stressing CPU or memory, throttling, latency, and packet loss. Since it is integrated with Amazon CloudWatch Alarms, you can set up stop conditions as guardrails to rollback an experiment if it causes unexpected impact.



AWS Fault Injection Simulator integrates with AWS resources to enable you to run fault injection experiments for your workloads.

There are also several third-party options for fault injection experiments. These include open-source tools such as [Chaos Toolkit](#), [Chaos Mesh](#), and [Litmus Chaos](#), as well as commercial options like Gremlin.

To expand the scope of faults that can be injected on AWS, AWS FIS [integrates with Chaos Mesh and Litmus Chaos](#), enabling you to coordinate fault injection workflows among multiple tools. For example, you can run a stress test on a pod's CPU using Chaos Mesh or Litmus faults while terminating a randomly selected percentage of cluster nodes using AWS FIS fault actions.

Implementation steps

- Determine which faults to use for experiments.

Assess the design of your workload for resiliency. Such designs (created using the best practices of the [Well-Architected Framework](#)) account for risks based on critical dependencies, past events, known issues, and compliance requirements. List each element of the design intended to maintain resilience and the faults it is designed to mitigate. For more information about creating such lists, see the [Operational Readiness Review whitepaper](#) which guides you on how to create a process to prevent reoccurrence of previous incidents. The Failure Modes and Effects Analysis (FMEA) process provides you with a framework for performing a component-level analysis of failures and how they impact your workload. FMEA is outlined in more detail by Adrian Cockcroft in [Failure Modes and Continuous Resilience](#).

- Assign a priority to each fault.

Start with a coarse categorization such as high, medium, or low. To assess priority, consider frequency of the fault and impact of failure to the overall workload.


When considering frequency of a given fault, analyze past data for this workload when available. If not available, use data from other workloads running in a similar environment.

When considering impact of a given fault, the larger the scope of the fault, generally the larger the impact. Also consider the workload design and purpose. For example, the ability to access the source data stores is critical for a workload doing data transformation and analysis. In this case, you would prioritize experiments for access faults, as well as throttled access and latency insertion.

Post-incident analyses are a good source of data to understand both frequency and impact of failure modes.

Use the assigned priority to determine which faults to experiment with first and the order with which to develop new fault injection experiments.

- For each experiment that you perform, follow the chaos engineering and continuous resilience flywheel.



Chaos engineering and continuous resilience flywheel, using the scientific method by Adrian Hornsby.

- Define steady state as some measurable output of a workload that indicates normal behavior.

Your workload exhibits steady state if it is operating reliably and as expected. Therefore, validate that your workload is healthy before defining steady state. Steady state does not necessarily mean no impact to the workload when a fault occurs, as a certain percentage in faults could be within acceptable limits. The steady state is your baseline that you will observe during the experiment, which will highlight anomalies if your hypothesis defined in the next step does not turn out as expected.

For example, a steady state of a payments system can be defined as the processing of 300 TPS with a success rate of 99% and round-trip time of 500 ms.

- Form a hypothesis about how the workload will react to the fault.

A good hypothesis is based on how the workload is expected to mitigate the fault to maintain the steady state. The hypothesis states that given the fault of a specific type, the system or workload will continue steady state, because the workload was designed with specific mitigations. The specific type of fault and mitigations should be specified in the hypothesis.

The following template can be used for the hypothesis (but other wording is also acceptable):

Note

If *specific fault* occurs, the *workload name* workload will *describe mitigating controls* to maintain *business or technical metric impact*.

For example:

- If 20% of the nodes in the Amazon EKS node-group are taken down, the Transaction Create API continues to serve the 99th percentile of requests in under 100 ms (steady state). The Amazon EKS nodes will recover within five minutes, and pods will get scheduled and process traffic within eight minutes after the initiation of the experiment. Alerts will fire within three minutes.
- If a single Amazon EC2 instance failure occurs, the order system's Elastic Load Balancing health check will cause the Elastic Load Balancing to only send requests to the remaining healthy instances while the Amazon EC2 Auto Scaling replaces the failed instance, maintaining a less than 0.01% increase in server-side (5xx) errors (steady state).
- If the primary Amazon RDS database instance fails, the Supply Chain data collection workload will failover and connect to the standby Amazon RDS database instance to maintain less than 1 minute of database read or write errors (steady state).
- Run the experiment by injecting the fault.

An experiment should by default be fail-safe and tolerated by the workload. If you know that the workload will fail, do not run the experiment. Chaos engineering should be used to find known-unknowns or unknown-unknowns. *Known-unknowns* are things you are aware of but don't fully understand, and *unknown-unknowns* are things you are neither aware of nor fully understand. Experimenting against a workload that you know is broken won't provide you with new insights. Your experiment should be carefully planned, have a clear scope of impact, and provide a rollback mechanism that can be applied in case of unexpected turbulence. If your due-diligence shows that your workload should survive the experiment, move forward with the experiment. There are several options for injecting the faults. For workloads on AWS, [AWS FIS](#) provides many predefined fault simulations called [actions](#). You can also define custom actions that run in AWS FIS using [AWS Systems Manager documents](#).

We discourage the use of custom scripts for chaos experiments, unless the scripts have the capabilities to understand the current state of the workload, are able to emit logs, and provide mechanisms for rollbacks and stop conditions where possible.

An effective framework or toolset which supports chaos engineering should track the current state of an experiment, emit logs, and provide rollback mechanisms to support the controlled execution of an experiment. Start with an established service like AWS FIS that allows you to perform experiments with a clearly defined scope and safety mechanisms that rollback the experiment if the experiment introduces unexpected turbulence. To learn about a wider variety of experiments using AWS FIS, also see the [Resilient and Well-Architected Apps with Chaos Engineering lab](#). Also, [AWS Resilience Hub](#) will analyze your workload and create experiments that you can choose to implement and run in AWS FIS.

Note

For every experiment, clearly understand the scope and its impact. We recommend that faults should be simulated first on a non-production environment before being run in production.

Experiments should run in production under real-world load using [canary deployments](#) that spin up both a control and experimental system deployment, where feasible. Running experiments during off-peak times is a good practice to mitigate potential impact when first experimenting in production. Also, if using actual customer traffic poses too much risk, you can run experiments using synthetic traffic on production infrastructure against the control and experimental deployments. When using production is not possible, run experiments in pre-production environments that are as close to production as possible.

You must establish and monitor guardrails to ensure the experiment does not impact production traffic or other systems beyond acceptable limits. Establish stop conditions to stop an experiment if it reaches a threshold on a guardrail metric that you define. This should include the metrics for steady state for the workload, as well as the metric against the components into which you're injecting the fault. A [synthetic monitor](#) (also known as a user canary) is one metric you should usually include as a user proxy. [Stop conditions for AWS FIS](#) are supported as part of the experiment template, enabling up to five stop-conditions per template.

One of the principles of chaos is minimize the scope of the experiment and its impact:

While there must be an allowance for some short-term negative impact, it is the responsibility and obligation of the Chaos Engineer to ensure the fallout from experiments are minimized and contained.

A method to verify the scope and potential impact is to perform the experiment in a non-production environment first, verifying that thresholds for stop conditions activate as expected during an experiment and observability is in place to catch an exception, instead of directly experimenting in production.

When running fault injection experiments, verify that all responsible parties are well-informed. Communicate with appropriate teams such as the operations teams, service reliability teams, and customer support to let them know when experiments will be run and what to expect. Give these teams communication tools to inform those running the experiment if they see any adverse effects.

You must restore the workload and its underlying systems back to the original known-good state. Often, the resilient design of the workload will self-heal. But some fault designs or failed experiments can leave your workload in an unexpected failed state. By the end of the experiment, you must be aware of this and restore the workload and systems. With AWS FIS you can set a rollback configuration (also called a post action) within the action parameters. A post action returns the target to the state that it was in before the action was run. Whether automated (such as using AWS FIS) or manual, these post actions should be part of a playbook that describes how to detect and handle failures.

- Verify the hypothesis.

[Principles of Chaos Engineering](#) gives this guidance on how to verify steady state of your workload:

Focus on the measurable output of a system, rather than internal attributes of the system. Measurements of that output over a short period of time constitute a proxy for the system's steady state. The overall system's throughput, error rates, and latency percentiles could all be metrics of interest representing steady state behavior. By focusing on systemic behavior patterns during experiments, chaos engineering verifies that the system does work, rather than trying to validate how it works.

In our two previous examples, we include the steady state metrics of less than 0.01% increase in server-side (5xx) errors and less than one minute of database read and write errors.

The 5xx errors are a good metric because they are a consequence of the failure mode that a client of the workload will experience directly. The database errors measurement is good as a direct consequence of the fault, but should also be supplemented with a client impact measurement such as failed customer requests or errors surfaced to the client. Additionally, include a synthetic monitor (also known as a user canary) on any APIs or URIs directly accessed by the client of your workload.

- Improve the workload design for resilience.

If steady state was not maintained, then investigate how the workload design can be improved to mitigate the fault, applying the best practices of the [AWS Well-Architected Reliability pillar](#). Additional guidance and resources can be found in the [AWS Builder's Library](#), which hosts articles

about how to [improve your health checks](#) or [employ retries with backoff in your application code](#), among others.

After these changes have been implemented, run the experiment again (shown by the dotted line in the chaos engineering flywheel) to determine their effectiveness. If the verify step indicates the hypothesis holds true, then the workload will be in steady state, and the cycle continues.

- Run experiments regularly.

A chaos experiment is a cycle, and experiments should be run regularly as part of chaos engineering. After a workload meets the experiment's hypothesis, the experiment should be automated to run continually as a regression part of your CI/CD pipeline. To learn how to do this, see this blog on [how to run AWS FIS experiments using AWS CodePipeline](#). This lab on recurrent [AWS FIS experiments in a CI/CD pipeline](#) enables you to work hands-on.

Fault injection experiments are also a part of game days (see [REL12-BP06 Conduct game days regularly \(p. 260\)](#)). Game days simulate a failure or event to verify systems, processes, and team responses. The purpose is to actually perform the actions the team would perform as if an exceptional event happened.

- Capture and store experiment results.

Results for fault injection experiments must be captured and persisted. Include all necessary data (such as time, workload, and conditions) to be able to later analyze experiment results and trends. Examples of results might include screenshots of dashboards, CSV dumps from your metric's database, or a hand-typed record of events and observations from the experiment. [Experiment logging with AWS FIS](#) can be part of this data capture.

Resources

Related best practices:

- [REL08-BP03 Integrate resiliency testing as part of your deployment \(p. 218\)](#)
- [REL13-BP03 Test disaster recovery implementation to validate the implementation \(p. 275\)](#)

Related documents:

- [What is AWS Fault Injection Simulator?](#)
- [What is AWS Resilience Hub?](#)
- [Principles of Chaos Engineering](#)
- [Chaos Engineering: Planning your first experiment](#)
- [Resilience Engineering: Learning to Embrace Failure](#)
- [Chaos Engineering stories](#)
- [Avoiding fallback in distributed systems](#)
- [Canary Deployment for Chaos Experiments](#)

Related videos:

- [AWS re:Invent 2020: Testing resiliency using chaos engineering \(ARC316\)](#)
- [AWS re:Invent 2019: Improving resiliency with chaos engineering \(DOP309-R1\)](#)
- [AWS re:Invent 2019: Performing chaos engineering in a serverless world \(CMY301\)](#)

Related examples:

- [Well-Architected lab: Level 300: Testing for Resiliency of Amazon EC2, Amazon RDS, and Amazon S3](#)

- [Chaos Engineering on AWS lab](#)
- [Resilient and Well-Architected Apps with Chaos Engineering lab](#)
- [Serverless Chaos lab](#)
- [Measure and Improve Your Application Resilience with AWS Resilience Hub lab](#)

Related tools:

- [AWS Fault Injection Simulator](#)
- AWS Marketplace: [Gremlin Chaos Engineering Platform](#)
- [Chaos Toolkit](#)
- [Chaos Mesh](#)
- [Litmus](#)

REL12-BP06 Conduct game days regularly

Use game days to regularly exercise your procedures for responding to events and failures as close to production as possible (including in production environments) with the people who will be involved in actual failure scenarios. Game days enforce measures to ensure that production events do not impact users.

Game days simulate a failure or event to test systems, processes, and team responses. The purpose is to actually perform the actions the team would perform as if an exceptional event happened. This will help you understand where improvements can be made and can help develop organizational experience in dealing with events. These should be conducted regularly so that your team builds *muscle memory* on how to respond.

After your design for resiliency is in place and has been tested in non-production environments, a game day is the way to ensure that everything works as planned in production. A game day, especially the first one, is an “all hands on deck” activity where engineers and operations are all informed when it will happen, and what will occur. Runbooks are in place. Simulated events are executed, including possible failure events, in the production systems in the prescribed manner, and impact is assessed. If all systems operate as designed, detection and self-healing will occur with little to no impact. However, if negative impact is observed, the test is rolled back and the workload issues are remedied, manually if necessary (using the runbook). Since game days often take place in production, all precautions should be taken to ensure that there is no impact on availability to your customers.

Common anti-patterns:

- Documenting your procedures, but never exercising them.
- Not including business decision makers in the test exercises.

Benefits of establishing this best practice: Conducting game days regularly ensures that all staff follows the policies and procedures when an actual incident occurs, and validates that those policies and procedures are appropriate.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Schedule game days to regularly exercise your runbooks and playbooks. Game days should involve everyone who would be involved in a production event: business owner, development staff, operational staff, and incident response teams.
 - Run your load or performance tests and then run your failure injection.
 - Look for anomalies in your runbooks and opportunities to exercise your playbooks.

- If you deviate from your runbooks, refine the runbook or correct the behavior. If you exercise your playbook, identify the runbook that should have been used, or create a new one.

Resources

Related documents:

- [What is AWS GameDay?](#)

Related videos:

- [AWS re:Invent 2019: Improving resiliency with chaos engineering \(DOP309-R1\)](#)

Related examples:

- [AWS Well-Architected Labs - Testing Resiliency](#)

REL 13 How do you plan for disaster recovery (DR)?

Having backups and redundant workload components in place is the start of your DR strategy. [RTO](#) and [RPO](#) are your objectives for restoration of your workload. Set these based on business needs. Implement a strategy to meet these objectives, considering locations and function of workload resources and data. The probability of disruption and cost of recovery are also key factors that help to inform the business value of providing disaster recovery for a workload.

Best practices

- [REL13-BP01 Define recovery objectives for downtime and data loss \(p. 261\)](#)
- [REL13-BP02 Use defined recovery strategies to meet the recovery objectives \(p. 265\)](#)
- [REL13-BP03 Test disaster recovery implementation to validate the implementation \(p. 275\)](#)
- [REL13-BP04 Manage configuration drift at the DR site or Region \(p. 276\)](#)
- [REL13-BP05 Automate recovery \(p. 277\)](#)

REL13-BP01 Define recovery objectives for downtime and data loss

The workload has a recovery time objective (RTO) and recovery point objective (RPO).

Recovery Time Objective (RTO) is the maximum acceptable delay between the interruption of service and restoration of service. This determines what is considered an acceptable time window when service is unavailable.

Recovery Point Objective (RPO) is the maximum acceptable amount of time since the last data recovery point. This determines what is considered an acceptable loss of data between the last recovery point and the interruption of service.

RTO and RPO values are important considerations when selecting an appropriate Disaster Recovery (DR) strategy for your workload. These objectives are determined by the business, and then used by technical teams to select and implement a DR strategy.

Desired Outcome:

Every workload has an assigned RTO and RPO, defined based on business impact. The workload is assigned to a predefined tier, defining service availability and acceptable loss of data, with an associated RTO and RPO. If such tiering is not possible then this can be assigned bespoke per workload, with the

intent to create tiers later. RTO and RPO are used as one of the primary considerations for selection of a disaster recovery strategy implementation for the workload. Additional considerations in picking a DR strategy are cost constraints, workload dependencies, and operational requirements.

For RTO, understand impact based on duration of an outage. Is it linear, or are there nonlinear implications? (for example, after four hours, you shut down a manufacturing line until the start of the next shift).

A disaster recovery matrix, like the following, can help you understand how workload criticality relates to recovery objectives. (Note that the actual values for the X and Y axes should be customized to your organization needs).

Disaster Recovery Matrix						
		Recovery Point Objective				
		< 1 Minute	< 1 Hour	< 6 Hours	< 1 Day	+ 1 Day
Recovery Time Objective	< 10 Minutes	Critical	Critical	High	Medium	Medium
	< 2 Hours	Critical	High	Medium	Medium	Low
	< 8 Hours	High	Medium	Medium	Low	Low
	< 24 Hours	Medium	Medium	Low	Low	Low
	24 + Hours	Medium	Low	Low	Low	Low

Figure 16: Disaster recovery matrix

Common anti-patterns:

- No defined recovery objectives.
- Selecting arbitrary recovery objectives.
- Selecting recovery objectives that are too lenient and do not meet business objectives.
- Not understanding of the impact of downtime and data loss.
- Selecting unrealistic recovery objectives, such as zero time to recover and zero data loss, which may not be achievable for your workload configuration.
- Selecting recovery objectives more stringent than actual business objectives. This forces DR implementations that are costlier and more complicated than what the workload needs.
- Selecting recovery objectives incompatible with those of a dependent workload.
- Your recovery objectives do not consider regulatory compliance requirements.
- RTO and RPO defined for a workload, but never tested.

Benefits of establishing this best practice: Your recovery objectives for time and data loss are necessary to guide your DR implementation.

Level of risk exposed if this best practice is not established: High

Implementation guidance

For the given workload, you must understand the impact of downtime and lost data on your business. The impact generally grows larger with greater downtime or data loss, but the shape of this growth can differ based on the workload type. For example, you may be able to tolerate downtime for up to an

hour with little impact, but after that impact quickly rises. Impact to business manifests in many forms including monetary cost (such as lost revenue), customer trust (and impact to reputation), operational issues (such as missing payroll or decreased productivity), and regulatory risk. Use the following steps to understand these impacts, and set RTO and RPO for your workload.

Implementation Steps

1. Determine your business stakeholders for this workload, and engage with them to implement these steps. Recovery objectives for a workload are a business decision. Technical teams then work with business stakeholders to use these objectives to select a DR strategy.

Note

For steps 2 and 3, you can use the [the section called "Implementation worksheet" \(p. 264\)](#).

2. Gather the necessary information to make a decision by answering the questions below.
3. Do you have categories or tiers of criticality for workload impact in your organization?
 - a. If yes, assign this workload to a category
 - b. If no, then establish these categories. Create five or fewer categories and refine the range of your recovery time objective for each one. Example categories include: critical, high, medium, low. To understand how workloads map to categories, consider whether the workload is mission critical, business important, or non-business driving.
 - c. Set workload RTO and RPO based on category. Always choose a category more strict (lower RTO and RPO) than the raw values calculated entering this step. If this results in an unsuitably large change in value, then consider creating a new category.
4. Based on these answers, assign RTO and RPO values to the workload. This can be done directly, or by assigning the workload to a predefined tier of service.
5. Document the disaster recovery plan (DRP) for this workload, which is a part of your organization's [business continuity plan \(BCP\)](#), in a location accessible to the workload team and stakeholders
 - a. Record the RTO and RPO, and the information used to determine these values. Include the strategy used for evaluating workload impact to the business
 - b. Record other metrics besides RTO and RPO are you tracking or plan to track for disaster recovery objectives
 - c. You will add details of your DR strategy and runbook to this plan when you create these.
6. By looking up the workload criticality in a matrix such as that in Figure 15, you can begin to establish predefined tiers of service defined for your organization.
7. After you have implemented a DR strategy (or a proof of concept for a DR strategy) as per [the section called "REL13-BP02 Use defined recovery strategies to meet the recovery objectives" \(p. 265\)](#), test this strategy to determine workload actual RTC (Recovery Time Capability) and RPC (Recovery Point Capability). If these do not meet the target recovery objectives, then either work with your business stakeholders to adjust those objectives, or make changes to the DR strategy is possible to meet target objectives.

Primary questions

1. What is the maximum time the workload can be down before severe impact to the business is incurred
 - a. Determine the monetary cost (direct financial impact) to the business per minute if workload is disrupted.
 - b. Consider that impact is not always linear. Impact can be limited at first, and then increase rapidly past a critical point in time.
2. What is the maximum amount of data that can be lost before severe impact to the business is incurred
 - a. Consider this value for your most critical data store. Identify the respective criticality for other data stores.
 - b. Can workload data be recreated if lost? If this is operationally easier than backup and restore, then choose RPO based on the criticality of the source data used to recreate the workload data.

3. What are the recovery objectives and availability expectations of workloads that this one depends on (downstream), or workloads that depend on this one (upstream)?
 - a. Choose recovery objectives that enable this workload to meet the requirements of upstream dependencies
 - b. Choose recovery objectives that are achievable given the recovery capabilities of downstream dependencies. Non-critical downstream dependencies (ones you can “work around”) can be excluded. Or, work with critical downstream dependencies to improve their recovery capabilities where necessary.

Additional questions

Consider these questions, and how they may apply to this workload:

4. Do you have different RTO and RPO depending on the type of outage (Region vs. AZ, etc.)?
5. Is there a specific time (seasonality, sales events, product launches) when your RTO/RPO may change? If so, what is the different measurement and time boundary?
6. How many customers will be impacted if workload is disrupted?
7. What is the impact to reputation if workload is disrupted?
8. What other operational impacts may occur if workload is disrupted? For example, impact to employee productivity if email systems are unavailable, or if Payroll systems are unable to submit transactions.
9. How does workload RTO and RPO align with Line of Business and Organizational DR Strategy?
10. Are there internal contractual obligations for providing a service? Are there penalties for not meeting them?
11. What are the regulatory or compliance constraints with the data?

Implementation worksheet

You can use this worksheet for implementation steps 2 and 3. You may adjust this worksheet to suit your specific needs, such as adding additional questions.

Step 2: Primary questions	Applies to workload?	workload RTO	workload RPO	RTO adjust.	RPO adjust.	Instructions
[1] maximum time the workload can be down						measured in time from start of outage to recovery
[2] maximum amount of data that can be lost						measured in time since last known good restorable dataset
[3a] upstream dependencies						enter the most strict upstream recovery objectives
[3b] downstream dependencies						enter the least strict downstream recovery objectives
[3a] reconciled upstream dependencies						If upstream value is less than current values and downstream value greater,
[3b] reconciled downstream dependencies						then work with dependencies to reconcile and enter reconciled values here
[3] dependencies						lower values to meet upstream dependencies or raise them based on downstream dependency capabilities
Step 2: Additional questions						Indicate if question applies. If it does not apply then skip it
Base RTO/RPO						Carry RTO and RPO values from above down to here
[4] type of outage	[] Y / [] N					Enter recovery objectives for event type with strictest requirements
[5] specific time-based objectives	[] Y / [] N					Enter recovery objectives for times with the strictest requirements
[6] customers disrupted	[] Y / [] N					Graph customers impacted as a function of time down or data lost. Use that to enter the maximum RTO and RPO permissible based on customer impact
[7] reputation impact	[] Y / [] N					Work with the business to determine maximum RTO and RPO based on impact to reputation
[8] operational impact	[] Y / [] N					Enter maximum RTO and RPO based on operational impact
[9] organizational alignment	[] Y / [] N					Enter maximum RTO and RPO for workloads of this type as per LOB and organizational requirements
[10] contractual obligations	[] Y / [] N					Enter maximum RTO and RPO based on contractual obligations
[11] regulatory compliance	[] Y / [] N					Enter maximum RTO and RPO based on applicable regulatory compliance
target based on additional questions						Take the minimum value (stricter value) from Q's 4-11 and enter it here
adjusted target						If the objectives on the above line cannot be accommodated, work with stakeholders to loosen constraints, and enter new minimum here
Adjusted RTO/RPO						Enter base RPO/RTO values, or adjusted target, whichever is lower
Step 3						
Map to predefined category or tier						Adjust both values to downward (more strict) to align to nearest defined tier

Worksheet

Level of effort for the Implementation Plan: Low

Resources

Related Best Practices:

- the section called “REL09-BP04 Perform periodic recovery of the data to verify backup integrity and processes” (p. 227)
- the section called “REL13-BP02 Use defined recovery strategies to meet the recovery objectives” (p. 265)
- the section called “REL13-BP03 Test disaster recovery implementation to validate the implementation” (p. 275)

Related documents:

- [AWS Architecture Blog: Disaster Recovery Series](#)
- [Disaster Recovery of Workloads on AWS: Recovery in the Cloud \(AWS Whitepaper\)](#)
- [Managing resiliency policies with AWS Resilience Hub](#)
- [APN Partner: partners that can help with disaster recovery](#)
- [AWS Marketplace: products that can be used for disaster recovery](#)

Related videos:

- [AWS re:Invent 2018: Architecture Patterns for Multi-Region Active-Active Applications \(ARC209-R2\)](#)
- [Disaster Recovery of Workloads on AWS](#)

REL13-BP02 Use defined recovery strategies to meet the recovery objectives

Define a disaster recovery (DR) strategy that meets your workload's recovery objectives. Choose a strategy such as: backup and restore; standby (active/passive); or active/active.

A DR strategy relies on the ability to stand up your workload in a recovery site if your primary location becomes unable to run the workload. The most common recovery objectives are RTO and RPO, as discussed in [REL13-BP01 Define recovery objectives for downtime and data loss \(p. 261\)](#).

A DR strategy across multiple Availability Zones (AZs) within a single AWS Region, can provide mitigation against disaster events like fires, floods, and major power outages. If it is a requirement to implement protection against an unlikely event that prevents your workload from being able to run in a given AWS Region, you can use a DR strategy that uses multiple Regions.

When architecting a DR strategy across multiple Regions, you should choose one of the following strategies. They are listed in increasing order of cost and complexity, and decreasing order of RTO and RPO. *Recovery Region* refers to an AWS Region other than the primary one used for your workload.




Figure 17: Disaster recovery (DR) strategies

- **Backup and restore** (RPO in hours, RTO in 24 hours or less): Back up your data and applications into the recovery Region. Using automated or continuous backups will enable point in time recovery, which can lower RPO to as low as 5 minutes in some cases. In the event of a disaster, you will deploy your infrastructure (using infrastructure as code to reduce RTO), deploy your code, and restore the backed-up data to recover from a disaster in the recovery Region.
- **Pilot light** (RPO in minutes, RTO in tens of minutes): Provision a copy of your core workload infrastructure in the recovery Region. Replicate your data into the recovery Region and create backups of it there. Resources required to support data replication and backup, such as databases and object storage, are always on. Other elements such as application servers or serverless compute are not deployed, but can be created when needed with the necessary configuration and application code.
- **Warm standby** (RPO in seconds, RTO in minutes): Maintain a scaled-down but fully functional version of your workload always running in the recovery Region. Business-critical systems are fully duplicated and are always on, but with a scaled down fleet. Data is replicated and live in the recovery Region. When the time comes for recovery, the system is scaled up quickly to handle the production load. The more scaled-up the Warm Standby is, the lower RTO and control plane reliance will be. When fully scales this is known as **Hot Standby**.
- **Multi-Region (multi-site) active-active** (RPO near zero, RTO potentially zero): Your workload is deployed to, and actively serving traffic from, multiple AWS Regions. This strategy requires you to synchronize data across Regions. Possible conflicts caused by writes to the same record in two different regional replicas must be avoided or handled, which can be complex. Data replication is useful for data synchronization and will protect you against some types of disaster, but it will not protect you against data corruption or destruction unless your solution also includes options for point-in-time recovery.

Note

The difference between pilot light and warm standby can sometimes be difficult to understand. Both include an environment in your recovery Region with copies of your primary region assets. The distinction is that Pilot Light cannot process requests without additional action taken first, while Warm Standby can handle traffic (at reduced capacity levels) immediately. Pilot Light will require you to turn on servers, possibly deploy additional (non-core) infrastructure, and scale up, while Warm Standby only requires you to scale up (everything is already deployed and running). Choose between these based on your RTO and RPO needs.

Desired outcome:

For each workload, there is a defined and implemented DR strategy that enables that workload to achieve DR objectives. DR strategies between workloads make use of reusable patterns (such as the strategies previously described),

Common anti-patterns:

- Implementing inconsistent recovery procedures for workloads with similar DR objectives.
- Leaving the DR strategy to be implemented ad-hoc when a disaster occurs.
- Having no plan for DR.
- Dependency on control plane operations during recovery.

Benefits of establishing this best practice:

- Using defined recovery strategies allows you to use common tooling and test procedures.
- Using defined recovery strategies enables more efficient sharing of knowledge between teams and easier implementation of DR on the workloads they own.

Level of risk exposed if this best practice is not established: High

- Without a planned, implemented, and tested DR strategy, you are unlikely to achieve recovery objectives in the event of a disaster.

Implementation guidance

For each of these steps, see the details below.

1. Determine a DR strategy that will satisfy recovery requirements for this workload.
2. Review the patterns for how the selected DR strategy can be implemented.
3. Assess the resources of your workload, and what their configuration will be in the recovery Region prior to failover (during normal operation).
4. Determine and implement how you will make your recovery Region ready for failover when needed (during a disaster event).
5. Determine and implement how you will reroute traffic to failover when needed (during a disaster event).
6. Design a plan for how your workload will fail back.

Implementation Steps

1. **Determine a DR strategy that will satisfy recovery requirements for this workload.**

Choosing a DR strategy is a trade-off between reducing downtime and data loss (RTO and RPO) versus cost and complexity of implementing the strategy. You should avoid implementing a strategy that is more stringent than it needs to be, as this incurs unnecessary costs.

For example, in the following diagram, the business has determined their maximum permissible RTO as well as the limit of what they can spend on their service restoration strategy. Given the business' objectives, the DR strategies Pilot Light or Warm Standby will satisfy both the RTO and the cost criteria.




Figure 18: Choosing a DR strategy based on RTO and cost

To learn more see [Business Continuity Plan \(BCP\)](#).

2. Review the patterns for how the selected DR strategy can be implemented.

This step is to understand how you will implement the selected strategy. The strategies are explained using AWS Regions as the primary and recovery sites. However, you can also choose to use Availability Zones within a single Region as your DR strategy, which makes use of elements of multiple of these strategies.

In the subsequent steps after this one, you will apply the strategy to your specific workload.

Backup and restore

Backup and restore is the least complex strategy to implement, but will require more time and effort to restore the workload, leading to higher RTO and RPO. It is a good practice to always make backups of your data, and copy these to another site (such as another AWS Region).




Figure 19: Backup and restore architecture

For more details on this strategy see [Disaster Recovery \(DR\) Architecture on AWS, Part II: Backup and Restore with Rapid Recovery](#).

Pilot light

With the *pilot light* approach, you replicate your data from your primary Region to your recovery Region. Core resources used for the workload infrastructure are deployed in the recovery Region, however additional resources and any dependencies are still needed to make this a functional stack. For example, in Figure 20, no compute instances are deployed.




Figure 20: Pilot light architecture

For more details on this strategy see [Disaster Recovery \(DR\) Architecture on AWS, Part III: Pilot Light and Warm Standby](#).

Warm standby

The *warm standby* approach involves ensuring that there is a scaled down, but fully functional, copy of your production environment in another Region. This approach extends the pilot light concept and decreases the time to recovery because your workload is always-on in another Region. If the recovery Region is deployed at full capacity, then this is known as *hot standby*.




Figure 21: Warm standby architecture

Using warm standby or pilot light requires scaling up resources in the recovery Region. To ensure capacity is available when needed, consider the use for [capacity reservations](#) for EC2 instances. If using AWS Lambda, then [provisioned concurrency](#) can ensure execution environments so that they are prepared to respond immediately to your function's invocations.

For more details on this strategy, see [Disaster Recovery \(DR\) Architecture on AWS, Part III: Pilot Light and Warm Standby](#).

Multi-site active/active

You can run your workload simultaneously in multiple Regions as part of a *multi-site active/active* strategy. Multi-site active/active serves traffic from all regions to which it is deployed. Customers may select this strategy for reasons other than DR. It can be used to increase availability, or when deploying a workload to a global audience (to put the endpoint closer to users and/or to deploy stacks localized to the audience in that region). As a DR strategy, if the workload cannot be supported in one of the AWS Regions to which it is deployed, then that Region is evacuated, and the remaining Region(s) are used to maintain availability. Multi-site active/active is the most operationally complex of the DR strategies, and should only be selected when business requirements necessitate it.




Figure 22: Multi-site active/active architecture

For more details on this strategy see [Disaster Recovery \(DR\) Architecture on AWS, Part IV: Multi-site Active/Active](#).

Additional practices for protecting data

With all strategies, you must also mitigate against a data disaster. Continuous data replication protects you against some types of disaster, but it may not protect you against data corruption or destruction unless your strategy also includes versioning of stored data or options for point-in-time recovery. You must also back up the replicated data in the recovery site to create point-in-time backups in addition to the replicas.

Using multiple Availability Zones (AZs) within a single AWS Region

When using multiple AZs within a single Region, your DR implementation uses multiple elements of the above strategies. First you must create a high-availability (HA) architecture, using multiple AZs as shown in Figure 23. This architecture makes use of a multi-site active/active approach, as the [Amazon EC2 instances](#) and the [Elastic Load Balancer](#) have resources deployed in multiple AZs, actively handling requests. The architecture also demonstrates hot standby, where if the primary [Amazon RDS](#) instance fails (or the AZ itself fails), then the standby instance is promoted to primary.




Figure 23: Multi-AZ architecture

In addition to this HA architecture, you need to add backups of all data required to run your workload. This is especially important for data that is constrained to a single zone such as [Amazon EBS volumes](#) or [Amazon Redshift clusters](#). If an AZ fails, you will need to restore this data to another AZ. Where possible, you should also copy data backups to another AWS Region as an additional layer of protection.

An less common alternative approach to single Region, multi-AZ DR is illustrated in the blog post, [Building highly resilient applications using Amazon Route 53 Application Recovery Controller, Part 1: Single-Region stack](#). Here, the strategy is to maintain as much isolation between the AZs as possible, like how Regions operate. Using this alternative strategy, you can choose an active/active or active/passive approach.

Note

Some workloads have regulatory data residency requirements. If this applies to your workload in a locality that currently has only one AWS Region, then multi-Region will not suit your business needs. Multi-AZ strategies provide good protection against most disasters.

3. Assess the resources of your workload, and what their configuration will be in the recovery Region prior to failover (during normal operation).

For infrastructure and AWS resources us infrastructure as code such as [AWS CloudFormation](#) or third-party tools like Hashicorp Terraform. To deploy across multiple accounts and Regions with a single operation you can use [AWS CloudFormation StackSets](#). For Multi-site active/active and Hot Standby strategies, the deployed infrastructure in your recovery Region has the same resources as your primary Region. For Pilot Light and Warm Standby strategies, the deployed infrastructure will require additional actions to become production ready. Using CloudFormation [parameters](#) and [conditional logic](#), you can control whether a deployed stack is active or standby with a single template. An example of such a CloudFormation template is included in [this blog post](#).

All DR strategies require that data sources are backed up within the AWS Region, and then those backups are copied to the recovery Region. [AWS Backup](#) provides a centralized view where you can configure, schedule, and monitor backups for these resources. For Pilot Light, Warm Standby, and Multi-site active/active, you should also replicate data from the primary Region to data resources in the recovery Region, such as [Amazon Relational Database Service \(Amazon RDS\)](#) DB instances or [Amazon DynamoDB](#) tables. These data resources are therefore live and ready to serve requests in the recovery Region.

To learn more about how AWS services operate across Regions, see this blog series on [Creating a Multi-Region Application with AWS Services](#).

4. Determine and implement how you will make your recovery Region ready for failover when needed (during a disaster event).

For Multi-site active/active, failover means evacuating a Region, and relying on the remaining active Regions. In general, those Regions are ready to accept traffic. For Pilot Light and Warm Standby strategies, your recovery actions will need to deploy the missing resources, such as the EC2 instances in Figure 20, plus any other missing resources.

For all of the above strategies you may need to promote read-only instances of databases to become the primary read/write instance.

For backup and restore, restoring data from backup creates resources for that data such as EBS volumes, RDS DB instances, and DynamoDB tables. You also need to restore the infrastructure and deploy code. You can use AWS Backup to restore data in the recovery Region. See [REL09-BP01 Identify and back up all data that needs to be backed up, or reproduce the data from sources \(p. 222\)](#) for more details. Rebuilding the infrastructure includes creating resources like EC2 instances in addition to the [Amazon Virtual Private Cloud \(Amazon VPC\)](#), subnets, and security groups needed. You can automate much of the restoration process. To learn how, see [this blog post](#).

5. Determine and implement how you will reroute traffic to failover when needed (during a disaster event).

This failover operation can be initiated either automatically or manually. Automatically initiated failover based on health checks or alarms should be used with caution since an unnecessary failover (false alarm) incurs costs such as non-availability and data loss. Manually initiated failover is therefore often used. In this case, you should still automate the steps for failover, so that the manual initiation is like the push of a button.

There are several traffic management options to consider when using AWS services. One option is to use [Amazon Route 53](#). Using Amazon Route 53, you can associate multiple IP endpoints in one or more AWS Regions with a Route 53 domain name. To implement manually initiated failover you can use [Amazon Route 53 Application Recovery Controller](#), which provides a highly available data plane API to reroute traffic to the recovery Region. When implementing failover, use data plane operations and avoid control plane ones as described in [REL11-BP04 Rely on the data plane and not the control plane during recovery \(p. 246\)](#).

To learn more about this and other options see [this section of the Disaster Recovery Whitepaper](#).

6. Design a plan for how your workload will fail back.

Failback is when you return workload operation to the primary Region, after a disaster event has abated. Provisioning infrastructure and code to the primary Region generally follows the same steps as were initially used, relying on infrastructure as code and code deployment pipelines. The challenge with failback is restoring data stores, and ensuring their consistency with the recovery Region in operation.

In the failed over state, the databases in the recovery Region are live and have the up-to-date data. The goal then is to re-synchronize from the recovery Region to the primary Region, ensuring it is up-to-date.

Some AWS services will do this automatically. If using [Amazon DynamoDB global tables](#), even if the table in the primary Region had become not available, when it comes back online, DynamoDB resumes propagating any pending writes. If using [Amazon Aurora Global Database](#) and using [managed planned failover](#), then Aurora global database's existing replication topology is maintained. Therefore, the former read/write instance in the primary Region will become a replica and receive updates from the recovery Region.

In cases where this is not automatic, you will need to re-establish the database in the primary Region as a replica of the database in the recovery Region. In many cases this will involve deleting the old primary database, and creating new replicas. For example, for instructions on how to do this with Amazon Aurora Global Database assuming an *unplanned* failover see this lab: [Fail Back a Global Database](#).

After a failover, if you can continue running in your recovery Region, consider making this the new primary Region. You would still do all the above steps to make the former primary Region into a recovery Region. Some organizations do a scheduled rotation, swapping their primary and recovery Regions periodically (for example every three months).

All of the steps required to fail over and fail back should be maintained in a playbook that is available to all members of the team, and is periodically reviewed.

Level of effort for the Implementation Plan: High

Resources

Related Best Practices:

- [the section called "REL09-BP01 Identify and back up all data that needs to be backed up, or reproduce the data from sources" \(p. 222\)](#)
- [the section called "REL11-BP04 Rely on the data plane and not the control plane during recovery" \(p. 246\)](#)
- [the section called "REL13-BP01 Define recovery objectives for downtime and data loss" \(p. 261\)](#)

Related documents:

- [AWS Architecture Blog: Disaster Recovery Series](#)
- [Disaster Recovery of Workloads on AWS: Recovery in the Cloud \(AWS Whitepaper\)](#)
- [Disaster recovery options in the cloud](#)
- [Build a serverless multi-region, active-active backend solution in an hour](#)
- [Multi-region serverless backend — reloaded](#)
- [RDS: Replicating a Read Replica Across Regions](#)
- [Route 53: Configuring DNS Failover](#)
- [S3: Cross-Region Replication](#)
- [What Is AWS Backup?](#)
- [What is Route 53 Application Recovery Controller?](#)
- [AWS Elastic Disaster Recovery](#)
- [HashiCorp Terraform: Get Started - AWS](#)
- [APN Partner: partners that can help with disaster recovery](#)
- [AWS Marketplace: products that can be used for disaster recovery](#)

Related videos:

- [Disaster Recovery of Workloads on AWS](#)
- [AWS re:Invent 2018: Architecture Patterns for Multi-Region Active-Active Applications \(ARC209-R2\)](#)
- [Get Started with AWS Elastic Disaster Recovery | Amazon Web Services](#)

Related examples:

- [AWS Well-Architected Labs - Disaster Recovery](#) - Series of workshops illustrating the DR strategies

REL13-BP03 Test disaster recovery implementation to validate the implementation

Regularly test failover to your recovery site to ensure proper operation, and that RTO and RPO are met.

A pattern to avoid is developing recovery paths that are rarely exercised. For example, you might have a secondary data store that is used for read-only queries. When you write to a data store and the primary fails, you might want to fail over to the secondary data store. If you don't frequently test this failover, you might find that your assumptions about the capabilities of the secondary data store are incorrect. The capacity of the secondary, which might have been sufficient when you last tested, might be no longer be able to tolerate the load under this scenario. Our experience has shown that the only error recovery that works is the path you test frequently. This is why having a small number of recovery paths is best. You can establish recovery patterns and regularly test them. If you have a complex or critical recovery path, you still need to regularly exercise that failure in production to convince yourself that the recovery path works. In the example we just discussed, you should fail over to the standby regularly, regardless of need.

Common anti-patterns:

- Never exercise failovers in production.

Benefits of establishing this best practice: Regularly testing your disaster recovery plan ensures that it will work when it needs to, and that your team knows how to execute the strategy.

Level of risk exposed if this best practice is not established: High

Implementation guidance

- Engineer your workloads for recovery. Regularly test your recovery paths Recovery Oriented Computing identifies the characteristics in systems that enhance recovery. These characteristics are: isolation and redundancy, system-wide ability to roll back changes, ability to monitor and determine health, ability to provide diagnostics, automated recovery, modular design, and ability to restart. Exercise the recovery path to ensure that you can accomplish the recovery in the specified time to the specified state. Use your runbooks during this recovery to document problems and find solutions for them before the next test.
 - [The Berkeley/Stanford recovery-oriented computing project](#)
- Use AWS Elastic Disaster Recovery to implement and launch drill instances for your DR strategy.
 - [AWS Elastic Disaster Recovery Preparing for Failover](#)
 - [What is Elastic Disaster Recovery?](#)
 - [AWS Elastic Disaster Recovery](#)

Resources

Related documents:

- [APN Partner: partners that can help with disaster recovery](#)
- [AWS Architecture Blog: Disaster Recovery Series](#)
- [AWS Marketplace: products that can be used for disaster recovery](#)
- [AWS Elastic Disaster Recovery](#)
- [Disaster Recovery of Workloads on AWS: Recovery in the Cloud \(AWS Whitepaper\)](#)
- [AWS Elastic Disaster Recovery Preparing for Failover](#)
- [The Berkeley/Stanford recovery-oriented computing project](#)
- [What is AWS Fault Injection Simulator?](#)

Related videos:

- [AWS re:Invent 2018: Architecture Patterns for Multi-Region Active-Active Applications \(ARC209-R2\)](#)
- [AWS re:Invent 2019: Backup-and-restore and disaster-recovery solutions with AWS \(STG208\)](#)

Related examples:

- [AWS Well-Architected Labs - Testing for Resiliency](#)

REL13-BP04 Manage configuration drift at the DR site or Region

Ensure that the infrastructure, data, and configuration are as needed at the DR site or Region. For example, check that AMIs and service quotas are up to date.

AWS Config continuously monitors and records your AWS resource configurations. It can detect drift and trigger [AWS Systems Manager Automation](#) to fix it and raise alarms. AWS CloudFormation can additionally detect drift in stacks you have deployed.

Common anti-patterns:

- Failing to make updates in your recovery locations, when you make configuration or infrastructure changes in your primary locations.
- Not considering potential limitations (like service differences) in your primary and recovery locations.

Benefits of establishing this best practice: Ensuring that your DR environment is consistent with your existing environment guarantees complete recovery.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Ensure that your delivery pipelines deliver to both your primary and backup sites. Delivery pipelines for deploying applications into production must distribute to all the specified disaster recovery strategy locations, including dev and test environments.
- Enable AWS Config to track potential drift locations. Use AWS Config rules to create systems that enforce your disaster recovery strategies and generate alerts when they detect drift.
 - [Remediating Noncompliant AWS Resources by AWS Config Rules](#)
 - [AWS Systems Manager Automation](#)
- Use AWS CloudFormation to deploy your infrastructure. AWS CloudFormation can detect drift between what your CloudFormation templates specify and what is actually deployed.
 - [AWS CloudFormation: Detect Drift on an Entire CloudFormation Stack](#)

Resources

Related documents:

- [APN Partner: partners that can help with disaster recovery](#)
- [AWS Architecture Blog: Disaster Recovery Series](#)
- [AWS CloudFormation: Detect Drift on an Entire CloudFormation Stack](#)
- [AWS Marketplace: products that can be used for disaster recovery](#)
- [AWS Systems Manager Automation](#)
- [Disaster Recovery of Workloads on AWS: Recovery in the Cloud \(AWS Whitepaper\)](#)

- [How do I implement an Infrastructure Configuration Management solution on AWS?](#)
- [Remediating Noncompliant AWS Resources by AWS Config Rules](#)

Related videos:

- [AWS re:Invent 2018: Architecture Patterns for Multi-Region Active-Active Applications \(ARC209-R2\)](#)

REL13-BP05 Automate recovery

Use AWS or third-party tools to automate system recovery and route traffic to the DR site or Region.

Based on configured health checks, AWS services, such as Elastic Load Balancing and AWS Auto Scaling, can distribute load to healthy Availability Zones while services, such as Amazon Route 53 and AWS Global Accelerator, can route load to healthy AWS Regions. Amazon Route 53 Application Recovery Controller helps you manage and coordinate failover using readiness check and routing control features. These features continually monitor your application's ability to recover from failures, so you can control application recovery across multiple AWS Regions, Availability Zones, and on premises.

For workloads on existing physical or virtual data centers or private clouds, [AWS Elastic Disaster Recovery](#) allows organizations to set up an automated disaster recovery strategy in AWS. Elastic Disaster Recovery also supports cross-Region and cross-Availability Zone disaster recovery in AWS.

Common anti-patterns:

- Implementing identical automated failover and failback can cause flapping when a failure occurs.

Benefits of establishing this best practice: Automated recovery reduces your recovery time by eliminating the opportunity for manual errors.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Automate recovery paths. For short recovery times, follow your [disaster recovery plan](#) to get your IT systems back online quickly in the case of a disruption.
- Use Elastic Disaster Recovery for automated Failover and Failback. Elastic Disaster Recovery continuously replicates your machines (including operating system, system state configuration, databases, applications, and files) into a low-cost staging area in your target AWS account and preferred Region. In the case of a disaster, after choosing to recover using Elastic Disaster Recovery, Elastic Disaster Recovery automates the conversion of your replicated servers into fully provisioned workloads in your recovery Region on AWS.
 - [Using Elastic Disaster Recovery for Failover and Failback](#)
 - [AWS Elastic Disaster Recovery resources](#)

Resources

Related documents:

- [APN Partner: partners that can help with disaster recovery](#)
- [AWS Architecture Blog: Disaster Recovery Series](#)
- [AWS Marketplace: products that can be used for disaster recovery](#)
- [AWS Systems Manager Automation](#)
- [AWS Elastic Disaster Recovery](#)
- [Disaster Recovery of Workloads on AWS: Recovery in the Cloud \(AWS Whitepaper\)](#)

Related videos:

- [AWS re:Invent 2018: Architecture Patterns for Multi-Region Active-Active Applications \(ARC209-R2\)](#)

Performance efficiency

Topics

- [Selection \(p. 278\)](#)
- [Review \(p. 330\)](#)
- [Monitoring \(p. 333\)](#)
- [Tradeoffs \(p. 340\)](#)

Selection

Questions

- [PERF 1 How do you select the best performing architecture? \(p. 278\)](#)
- [PERF 2 How do you select your compute solution? \(p. 285\)](#)
- [PERF 3 How do you select your storage solution? \(p. 296\)](#)
- [PERF 4 How do you select your database solution? \(p. 303\)](#)
- [PERF 5 How do you configure your networking solution? \(p. 318\)](#)

PERF 1 How do you select the best performing architecture?

Often, multiple approaches are required for optimal performance across a workload. Well-architected systems use multiple solutions and features to improve performance.

Best practices

- [PERF01-BP01 Understand the available services and resources \(p. 278\)](#)
- [PERF01-BP02 Define a process for architectural choices \(p. 279\)](#)
- [PERF01-BP03 Factor cost requirements into decisions \(p. 280\)](#)
- [PERF01-BP04 Use policies or reference architectures \(p. 281\)](#)
- [PERF01-BP05 Use guidance from your cloud provider or an appropriate partner \(p. 282\)](#)
- [PERF01-BP06 Benchmark existing workloads \(p. 283\)](#)
- [PERF01-BP07 Load test your workload \(p. 284\)](#)

PERF01-BP01 Understand the available services and resources

Learn about and understand the wide range of services and resources available in the cloud. Identify the relevant services and configuration options for your workload, and understand how to achieve optimal performance.

If you are evaluating an existing workload, you must generate an inventory of the various services resources it consumes. Your inventory helps you evaluate which components can be replaced with managed services and newer technologies.

Common anti-patterns:

- You use the cloud as a collocated data center.
- You use shared storage for all things that need persistent storage.

- You do not use automatic scaling.
- You use instance types that are closest matched, but larger where needed, to your current standards.
- You deploy and manage technologies that are available as managed services.

Benefits of establishing this best practice: By considering services you may be unfamiliar with, you may be able to greatly reduce the cost of infrastructure and the effort required to maintain your services. You may be able to accelerate your time to market by deploying new services and features.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Inventory your workload software and architecture for related services: Gather an inventory of your workload and decide which category of products to learn more about. Identify workload components that can be replaced with managed services to increase performance and reduce operational complexity.

Resources

Related documents:

- [AWS Architecture Center](#)
- [AWS Partner Network](#)
- [AWS Solutions Library](#)
- [AWS Knowledge Center](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [This is my Architecture](#)

Related examples:

- [AWS Samples](#)
- [AWS SDK Examples](#)

PERF01-BP02 Define a process for architectural choices

Use internal experience and knowledge of the cloud, or external resources such as published use cases, relevant documentation, or whitepapers, to define a process to choose resources and services. You should define a process that encourages experimentation and benchmarking with the services that could be used in your workload.

When you write critical user stories for your architecture, you should include performance requirements, such as specifying how quickly each critical story should run. For these critical stories, you should implement additional scripted user journeys to ensure that you have visibility into how these stories perform against your requirements.

Common anti-patterns:

- You assume your current architecture will become static and not be updated over time.
- You introduce architecture changes over time without justification.

Benefits of establishing this best practice: By having a defined process for making architectural changes, you enable using the gathered data to influence your workload design over time.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Select an architectural approach: Identify the kind of architecture that meets your performance requirements. Identify constraints, such as the media for delivery (desktop, web, mobile, IoT), legacy requirements, and integrations. Identify opportunities for reuse, including refactoring. Consult other teams, architecture diagrams, and resources such as AWS Solution Architects, AWS Reference Architectures, and AWS Partners to help you choose an architecture.

Define performance requirements: Use the customer experience to identify the most important metrics. For each metric, identify the target, measurement approach, and priority. Define the customer experience. Document the performance experience required by customers, including how customers will judge the performance of the workload. Prioritize experience concerns for critical user stories. Include performance requirements and implement scripted user journeys to ensure that you know how the stories perform against your requirements.

Resources

Related documents:

- [AWS Architecture Center](#)
- [AWS Partner Network](#)
- [AWS Solutions Library](#)
- [AWS Knowledge Center](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [This is my Architecture](#)

Related examples:

- [AWS Samples](#)
- [AWS SDK Examples](#)

PERF01-BP03 Factor cost requirements into decisions

Workloads often have cost requirements for operation. Use internal cost controls to select resource types and sizes based on predicted resource need.

Determine which workload components could be replaced with fully managed services, such as managed databases, in-memory caches, and ETL services. Reducing your operational workload allows you to focus resources on business outcomes.

For cost requirement best practices, refer to the *Cost-Effective Resources* section of the [Cost Optimization Pillar whitepaper](#).

Common anti-patterns:

- You only use one family of instances.
- You do not evaluate licensed solutions versus open-source solutions
- You only use block storage.
- You deploy common software on EC2 instances and Amazon EBS or ephemeral volumes that are available as a managed service.

Benefits of establishing this best practice: Considering cost when making your selections will allow you to enable other investments.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Optimize workload components to reduce cost: Right size workload components and enable elasticity to reduce cost and maximize component efficiency. Determine which workload components can be replaced with managed services when appropriate, such as managed databases, in-memory caches, and reverse proxies.

Resources

Related documents:

- [AWS Architecture Center](#)
- [AWS Partner Network](#)
- [AWS Solutions Library](#)
- [AWS Knowledge Center](#)
- [AWS Compute Optimizer](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [This is my Architecture](#)
- [Optimize performance and cost for your AWS compute \(CMP323-R1\)](#)

Related examples:

- [AWS Samples](#)
- [AWS SDK Examples](#)
- [Rightsizing with Compute Optimizer and Memory utilization enabled](#)
- [AWS Compute Optimizer Demo code](#)

PERF01-BP04 Use policies or reference architectures

Maximize performance and efficiency by evaluating internal policies and existing reference architectures and using your analysis to select services and configurations for your workload.

Common anti-patterns:

- You allow wide use of technology selection that may impact the management overhead of your company.

Benefits of establishing this best practice: Establishing a policy for architecture, technology, and vendor choices will allow decisions to be made quickly.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Deploy your workload using existing policies or reference architectures: Integrate the services into your cloud deployment, then use your performance tests to ensure that you can continue to meet your performance requirements.

Resources

Related documents:

- [AWS Architecture Center](#)
- [AWS Partner Network](#)
- [AWS Solutions Library](#)
- [AWS Knowledge Center](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [This is my Architecture](#)

Related examples:

- [AWS Samples](#)
- [AWS SDK Examples](#)

PERF01-BP05 Use guidance from your cloud provider or an appropriate partner

Use cloud company resources, such as solutions architects, professional services, or an appropriate partner to guide your decisions. These resources can help review and improve your architecture for optimal performance.

Reach out to AWS for assistance when you need additional guidance or product information. AWS Solutions Architects and [AWS Professional Services](#) provide guidance for solution implementation. [AWS Partners](#) provide AWS expertise to help you unlock agility and innovation for your business.

Common anti-patterns:

- You use AWS as a common data center provider.
- You use AWS services in a manner that they were not designed for.

Benefits of establishing this best practice: Consulting with your provider or a partner will give you confidence in your decisions.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Reach out to AWS resources for assistance: AWS Solutions Architects and Professional Services provide guidance for solution implementation. APN Partners provide AWS expertise to help you unlock agility and innovation for your business.

Resources

Related documents:

- [AWS Architecture Center](#)
- [AWS Partner Network](#)
- [AWS Solutions Library](#)
- [AWS Knowledge Center](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [This is my Architecture](#)

Related examples:

- [AWS Samples](#)
- [AWS SDK Examples](#)

PERF01-BP06 Benchmark existing workloads

Benchmark the performance of an existing workload to understand how it performs on the cloud. Use the data collected from benchmarks to drive architectural decisions.

Use benchmarking with synthetic tests and real-user monitoring to generate data about how your workload's components perform. Benchmarking is generally quicker to set up than load testing and is used to evaluate the technology for a particular component. Benchmarking is often used at the start of a new project, when you lack a full solution to load test.

You can either build your own custom benchmark tests, or you can use an industry standard test, such as [TPC-DS](#) to benchmark your data warehousing workloads. Industry benchmarks are helpful when comparing environments. Custom benchmarks are useful for targeting specific types of operations that you expect to make in your architecture.

When benchmarking, it is important to pre-warm your test environment to ensure valid results. Run the same benchmark multiple times to ensure that you've captured any variance over time.

Because benchmarks are generally faster to run than load tests, they can be used earlier in the deployment pipeline and provide faster feedback on performance deviations. When you evaluate a significant change in a component or service, a benchmark can be a quick way to see if you can justify the effort to make the change. Using benchmarking in conjunction with load testing is important because load testing informs you about how your workload will perform in production.

Common anti-patterns:

- You rely on common benchmarks that are not indicative of your workload characteristics.
- You rely on customer feedback and perceptions as your only benchmark.

Benefits of establishing this best practice: Benchmarking your current implementation allows you to measure the improvement in performance.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Monitor performance during development: Implement processes that provide visibility into performance as your workload evolves.

Integrate into your delivery pipeline: Automatically run load tests in your delivery pipeline. Compare the test results against pre-defined key performance indicators (KPIs) and thresholds to ensure that you continue to meet performance requirements.

Test user journeys: Use synthetic or sanitized versions of production data (remove sensitive or identifying information) for load testing. Exercise your entire architecture by using replayed or pre-programmed user journeys through your application at scale.

Real-user monitoring: Use CloudWatch RUM to help you collect and view client-side data about your application performance. Use this data to help establish your real-user performance benchmarks.

Resources

Related documents:

- [AWS Architecture Center](#)
- [AWS Partner Network](#)
- [AWS Solutions Library](#)
- [AWS Knowledge Center](#)
- [Amazon CloudWatch RUM](#)
- [Amazon CloudWatch Synthetics](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [This is my Architecture](#)
- [Optimize applications through Amazon CloudWatch RUM](#)
- [Demo of Amazon CloudWatch Synthetics](#)

Related examples:

- [AWS Samples](#)
- [AWS SDK Examples](#)
- [Distributed Load Tests](#)
- [Measure page load time with Amazon CloudWatch Synthetics](#)
- [Amazon CloudWatch RUM Web Client](#)

PERF01-BP07 Load test your workload

Deploy your latest workload architecture on the cloud using different resource types and sizes. Monitor the deployment to capture performance metrics that identify bottlenecks or excess capacity. Use this performance information to design or improve your architecture and resource selection.

Load testing uses your *actual* workload so that you can see how your solution performs in a production environment. Load tests must be run using synthetic or sanitized versions of production data (remove sensitive or identifying information). Use replayed or pre-programmed user journeys through your workload at scale that exercise your entire architecture. Automatically carry out load tests as part of your delivery pipeline, and compare the results against pre-defined KPIs and thresholds. This ensures that you continue to achieve required performance.

Common anti-patterns:

- You load test individual parts of your workload but not your entire workload.
- You load test on infrastructure that is not the same as your production environment.
- You only conduct load testing to your expected load and not beyond, to help foresee where you may have future problems.
- Performing load testing without informing AWS Support, and having your test defeated as it looks like a denial of service event.

Benefits of establishing this best practice: Measuring your performance under a load test will show you where you will be impacted as load increases. This can provide you with the capability of anticipating needed changes before they impact your workload.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Validate your approach with load testing: Load test a proof-of-concept to find out if you meet your performance requirements. You can use AWS services to run production-scale environments to test your architecture. Because you only pay for the test environment when it is needed, you can carry out full-scale testing at a fraction of the cost of using an on-premises environment.

Monitor metrics: Amazon CloudWatch can collect metrics across the resources in your architecture. You can also collect and publish custom metrics to surface business or derived metrics. Use CloudWatch or third-party solutions to set alarms that indicate when thresholds are breached.

Test at scale: Load testing uses your actual workload so you can see how your solution performs in a production environment. You can use AWS services to run production-scale environments to test your architecture. Because you only pay for the test environment when it is needed, you can run full-scale testing at a lower cost than using an on-premises environment. Take advantage of the AWS Cloud to test your workload to discover where it fails to scale, or if it scales in a non-linear way. For example, use Spot Instances to generate loads at low cost and discover bottlenecks before they are experienced in production.

Resources

Related documents:

- [AWS CloudFormation](#)
- [Building AWS CloudFormation Templates using CloudFormer](#)
- [Amazon CloudWatch RUM](#)
- [Amazon CloudWatch Synthetics](#)
- [Distributed Load Testing on AWS](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [Optimize applications through Amazon CloudWatch RUM](#)
- [Demo of Amazon CloudWatch Synthetics](#)

Related examples:

- [Distributed Load Testing on AWS](#)

PERF 2 How do you select your compute solution?

The optimal compute solution for a workload varies based on application design, usage patterns, and configuration settings. Architectures can use different compute solutions for various components and enable different features to improve performance. Selecting the wrong compute solution for an architecture can lead to lower performance efficiency.

Best practices

- [PERF02-BP01 Evaluate the available compute options \(p. 286\)](#)
- [PERF02-BP02 Understand the available compute configuration options \(p. 288\)](#)
- [PERF02-BP03 Collect compute-related metrics \(p. 291\)](#)
- [PERF02-BP04 Determine the required configuration by right-sizing \(p. 293\)](#)

- [PERF02-BP05 Use the available elasticity of resources \(p. 294\)](#)
- [PERF02-BP06 Re-evaluate compute needs based on metrics \(p. 295\)](#)

PERF02-BP01 Evaluate the available compute options

Understand how your workload can benefit from the use of different compute options, such as instances, containers and functions.

Desired outcome: By understanding all of the compute options available, you will be aware of the opportunities to increase performance, reduce unnecessary infrastructure costs, and lower the operational effort required to maintain your workload. You can also accelerate your time to market when you deploy new services and features.

Common anti-patterns:

- In a post-migration workload, using the same compute solution that was being used on premises.
- Lacking awareness of the cloud compute solutions and how those solutions might improve your compute performance.
- Oversizing an existing compute solution to meet scaling or performance requirements, when an alternative compute solution would align to your workload characteristics more precisely.

Benefits of establishing this best practice: By identifying the compute requirements and evaluating the available compute solutions, business stakeholders and engineering teams will understand the benefits and limitations of using the selected compute solution. The selected compute solution should fit the workload performance criteria. Key criteria include processing needs, traffic patterns, data access patterns, scaling needs, and latency requirements.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Understand the virtualization, containerization, and management solutions that can benefit your workload and meet your performance requirements. A workload can contain multiple types of compute solutions. Each compute solution has differing characteristics. Based on your workload scale and compute requirements, a compute solution can be selected and configured to meet your needs. The cloud architect should learn the advantages and disadvantages of instances, containers, and functions. The following steps will help you through how to select your compute solution to match your workload characteristics and performance requirements.

Type	Server	Containers	Function
AWS service	Amazon Elastic Compute Cloud (Amazon EC2)	Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS)	AWS Lambda
Key Characteristics	Has dedicated option for hardware license requirements, Placement Options, and a large selection of different instance families based on compute metrics	Easy deployment, consistent environments, runs on top of EC2 instances, Scalable	Short runtime (15 minutes or less), maximum memory and CPU are not as high as other services, Managed hardware layer, Scales to millions of concurrent requests

Type	Server	Containers	Function
Common use-cases	Lift and shift migrations, monolithic application, hybrid environments, enterprise applications	Microservices, hybrid environments,	Microservices, event-driven applications

Implementation steps:

1. Select the location of where the compute solution must reside by evaluating [the section called "PERF05-BP06 Choose your workload's location based on network requirements" \(p. 327\)](#). This location will limit the types of compute solution available to you.
2. Identify the type of compute solution that works with the location requirement and application requirements
 - a. [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) virtual server instances come in a wide variety of different families and sizes. They offer a wide variety of capabilities, including solid state drives (SSDs) and graphics processing units (GPUs). EC2 instances offer the greatest flexibility on instance choice. When you launch an EC2 instance, the instance type that you specify determines the hardware of your instance. Each instance type offers different compute, memory, and storage capabilities. Instance types are grouped in instance families based on these capabilities. Typical use cases include: running enterprise applications, high performance computing (HPC), training and deploying machine learning applications and running cloud native applications.
 - b. [Amazon Elastic Container Service \(Amazon ECS\)](#) is a fully managed container orchestration service that allows you to automatically run and manage containers on a cluster of EC2 instances or serverless instances using AWS Fargate. You can use Amazon ECS with other services such as Amazon Route 53, Secrets Manager, AWS Identity and Access Management (IAM), and Amazon CloudWatch. Amazon ECS is recommended if your application is containerized and your engineering team prefers Docker containers.
 - c. [Amazon Elastic Kubernetes Service \(Amazon EKS\)](#) is a fully managed Kubernetes service. You can choose to run your EKS clusters using AWS Fargate, removing the need to provision and manage servers. Managing Amazon EKS is simplified due to integrations with AWS Services such as Amazon CloudWatch, Auto Scaling Groups, AWS Identity and Access Management (IAM), and Amazon Virtual Private Cloud (VPC). When using containers, you must use compute metrics to select the optimal type for your workload, similar to how you use compute metrics to select your EC2 or AWS Fargate instance types. Amazon EKS is recommended if your application is containerized and your engineering team prefers Kubernetes over Docker containers.
 - d. You can use [AWS Lambda](#) to run code that supports the allowed runtime, memory, and CPU options. Simply upload your code, and AWS Lambda will manage everything required to run and scale that code. You can set up your code to automatically trigger from other AWS services or call it directly. Lambda is recommended for short running, microservice architectures developed for the cloud.
3. After you have experimented with your new compute solution, plan your migration and validate your performance metrics. This is a continual process, see [the section called "PERF02-BP04 Determine the required configuration by right-sizing" \(p. 293\)](#).

Level of effort for the implementation plan: If a workload is moving from one compute solution to another, there could be a *moderate* level of effort involved in refactoring the application.

Resources

Related documents:

- [Cloud Compute with AWS](#)

- [EC2 Instance Types](#)
- [Processor State Control for Your EC2 Instance](#)
- [EKS Containers: EKS Worker Nodes](#)
- [Amazon ECS Containers: Amazon ECS Container Instances](#)
- [Functions: Lambda Function Configuration](#)
- [Prescriptive Guidance for Containers](#)
- [Prescriptive Guidance for Serverless](#)

Related videos:

- [How to choose compute option for startups](#)
- [Optimize performance and cost for your AWS compute \(CMP323-R1\)](#)
- [Amazon EC2 foundations \(CMP211-R2\)](#)
- [Powering next-gen Amazon EC2: Deep dive into the Nitro system](#)
- [Deliver high-performance ML inference with AWS Inferentia \(CMP324-R1\)](#)
- [Better, faster, cheaper compute: Cost-optimizing Amazon EC2 \(CMP202-R1\)](#)

Related examples:

- [Migrating the web application to containers](#)
- [Run a Serverless Hello World](#)

PERF02-BP02 Understand the available compute configuration options

Each compute solution has options and configurations available to you to support your workload characteristics. Learn how various options complement your workload, and which configuration options are best for your application. Examples of these options include instance family, sizes, features (GPU, I/O), bursting, time-outs, function sizes, container instances, and concurrency.

Desired outcome: The workload characteristics including CPU, memory, network throughput, GPU, IOPS, traffic patterns, and data access patterns are documented and used to configure the compute solution to match the workload characteristics. Each of these metrics plus custom metrics specific to your workload are recorded, monitored, and then used to optimize the compute configuration to best meet the requirements.

Common anti-patterns:

- Using the same compute solution that was being used on premises.
- Not reviewing the compute options or instance family to match workload characteristics.
- Oversizing the compute to ensure bursting capability.
- You use multiple compute management platforms for the same workload.

Benefits of establishing this best practice: Be familiar with the AWS compute offerings so that you can determine the correct solution for each of your workloads. After you have selected the compute offerings for your workload, you can quickly experiment with those compute offerings to determine how well they meet your workload needs. A compute solution that is optimized to meet your workload characteristics will increase your performance, lower your cost and increase your reliability.

Level of risk exposed if this best practice is not established: High

Implementation guidance

If your workload has been using the same compute option for more than four weeks and you anticipate that the characteristics will remain the same in the future, you can use [AWS Compute Optimizer](#) to provide a recommendation to you based on your compute characteristics. If AWS Compute Optimizer is not an option due to lack of metrics, [a non-supported instance type](#) or a foreseeable change in your characteristics then you must predict your metrics based on load testing and experimentation.

Implementation steps:

1. Are you running on EC2 instances or containers with the EC2 Launch Type?
 - a. Can your workload use GPUs to increase performance?
 - i. [Accelerated Computing](#) instances are GPU-based instances that provide the highest performance for machine learning training, inference and high performance computing.
 - b. Does your workload run machine learning inference applications?
 - i. [AWS Inferentia \(Inf1\)](#) — Inf1 instances are built to support machine learning inference applications. Using Inf1 instances, customers can run large-scale machine learning inference applications, such as image recognition, speech recognition, natural language processing, personalization, and fraud detection. You can build a model in one of the popular machine learning frameworks, such as TensorFlow, PyTorch, or MXNet and use GPU instances, to train your model. After your machine learning model is trained to meet your requirements, you can deploy your model on Inf1 instances by using [AWS Neuron](#), a specialized software development kit (SDK) consisting of a compiler, runtime, and profiling tools that optimize the machine learning inference performance of Inferentia chips.
 - c. Does your workload integrate with the low-level hardware to improve performance?
 - i. [Field Programmable Gate Arrays \(FPGA\)](#) — Using FPGAs, you can optimize your workloads by having custom hardware-accelerated execution for your most demanding workloads. You can define your algorithms by leveraging supported general programming languages such as C or Go, or hardware-oriented languages such as Verilog or VHDL.
 - d. Do you have at least four weeks of metrics and can predict that your traffic pattern and metrics will remain about the same in the future?
 - i. Use [Compute Optimizer](#) to get a machine learning recommendation on which compute configuration best matches your compute characteristics.
 - e. Is your workload performance constrained by the CPU metrics?
 - i. [Compute-optimized](#) instances are ideal for the workloads that require high performing processors.
 - f. Is your workload performance constrained by the memory metrics?
 - i. [Memory-optimized](#) instances deliver large amounts of memory to support memory intensive workloads.
 - g. Is your workload performance constrained by IOPS?
 - i. [Storage-optimized](#) instances are designed for workloads that require high, sequential read and write access (IOPS) to local storage.
 - h. Do your workload characteristics represent a balanced need across all metrics?
 - i. Does your workload CPU need to burst to handle spikes in traffic?
 - A. [Burstable Performance](#) instances are similar to Compute Optimized instances except they offer the ability to burst past the fixed CPU baseline identified in a compute-optimized instance.
 - ii. [General Purpose](#) instances provide a balance of all characteristics to support a variety of workloads.
 - i. Is your compute instance running on Linux and constrained by network throughput on the network interface card?
 - i. Review [Performance Question 5, Best Practice 2: Evaluate available networking features](#) to find the right instance type and family to meet your performance needs.

- j. Does your workload need consistent and predictable instances in a specific Availability Zone that you can commit to for a year?
 - i. [Reserved Instances](#) confirms capacity reservations in a specific Availability Zone. Reserved Instances are ideal for required compute power in a specific Availability Zone.
- k. Does your workload have licenses that require dedicated hardware?
 - i. [Dedicated Hosts](#) support existing software licenses and help you meet compliance requirements.
- l. Does your compute solution burst and require synchronous processing?
 - i. [On-Demand Instances](#) let you use the compute capacity by the hour or second with no long-term commitment. These instances are good for bursting above performance baseline needs.
- m. Is your compute solution stateless, fault-tolerant, and asynchronous?
 - i. [Spot Instances](#) let you take advantage of unused instance capacity for your stateless, fault-tolerant workloads.
- 2. Are you running containers on [Fargate](#)?
 - a. Is your task performance constrained by the memory or CPU?
 - i. Use the [Task Size](#) to adjust your memory or CPU.
 - b. Is your performance being affected by your traffic pattern bursts?
 - i. Use the [Auto Scaling](#) configuration to match your traffic patterns.
- 3. Is your compute solution on [Lambda](#)?
 - a. Do you have at least four weeks of metrics and can predict that your traffic pattern and metrics will remain about the same in the future?
 - i. Use [Compute Optimizer](#) to get a machine learning recommendation on which compute configuration best matches your compute characteristics.
 - b. Do you not have enough metrics to use AWS Compute Optimizer?
 - i. If you do not have metrics available to use Compute Optimizer, use [AWS Lambda Power Tuning](#) to help select the best configuration.
 - c. Is your function performance constrained by the memory or CPU?
 - i. Configure your [Lambda memory](#) to meet your performance needs metrics.
 - d. Is your function timing out on execution?
 - i. Change the [timeout settings](#)
 - e. Is your function performance constrained by bursts of activity and concurrency?
 - i. Configure the [concurrency settings](#) to meet your performance requirements.
 - f. Does your function execute asynchronously and is failing on retries?
 - i. Configure the maximum age of the event and the maximum retry limit in the [asynchronous configuration](#) settings.

Level of effort for the implementation plan:

To establish this best practice, you must be aware of your current compute characteristics and metrics. Gathering those metrics, establishing a baseline and then using those metrics to identify the ideal compute option is a *low* to *moderate* level of effort. This is best validated by load tests and experimentation.

Resources

Related documents:

- [Cloud Compute with AWS](#)
- [AWS Compute Optimizer](#)
- [EC2 Instance Types](#)
- [Processor State Control for Your EC2 Instance](#)

- [EKS Containers: EKS Worker Nodes](#)
- [Amazon ECS Containers: Amazon ECS Container Instances](#)
- [Functions: Lambda Function Configuration](#)

Related videos:

- [Amazon EC2 foundations \(CMP211-R2\)](#)
- [Powering next-gen Amazon EC2: Deep dive into the Nitro system](#)
- [Optimize performance and cost for your AWS compute \(CMP323-R1\)](#)

Related examples:

- [Rightsizing with Compute Optimizer and Memory utilization enabled](#)
- [AWS Compute Optimizer Demo code](#)

PERF02-BP03 Collect compute-related metrics

To understand how your compute resources are performing, you must record and track the utilization of various systems. This data can be used to make more accurate determinations about resource requirements.

Workloads can generate large volumes of data such as metrics, logs, and events. Determine if your existing storage, monitoring, and observability service can manage the data generated. Identify which metrics reflect resource utilization and can be collected, aggregated, and correlated on a single platform across. Those metrics should represent all your workload resources, applications, and services, so you can easily gain system-wide visibility and quickly identify performance improvement opportunities and issues.

Desired outcome: All metrics related to the compute-related resources are identified, collected, aggregated, and correlated on a single platform with retention implemented to support cost and operational goals.

Common anti-patterns:

- You only use manual log file searching for metrics.
- You only publish metrics to internal tools.
- You only use the default metrics recorded by your selected monitoring software.
- You only review metrics when there is an issue.

Benefits of establishing this best practice: To monitor the performance of your workloads, you must record multiple performance metrics over a period of time. These metrics allow you to detect anomalies in performance. They will also help gauge performance against business metrics to ensure that you are meeting your workload needs.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Identify, collect, aggregate, and correlate compute-related metrics. Using a service such as Amazon CloudWatch, can make the implementation quicker and easier to maintain. In addition to the default metrics recorded, identify and track additional system-level metrics within your workload. Record data such as CPU utilization, memory, disk I/O, and network inbound and outbound metrics to gain insight

into utilization levels or bottlenecks. This data is crucial to understand how the workload is performing and how the compute solution is utilized. Use these metrics as part of a data-driven approach to actively tune and optimize your workload's resources.

Implementation steps:

1. Which compute solution metrics are important to track?
 - a. [EC2 default metrics](#)
 - b. [Amazon ECS default metrics](#)
 - c. [EKS default metrics](#)
 - d. [Lambda default metrics](#)
 - e. [EC2 memory and disk metrics](#)
2. Do I currently have an approved logging and monitoring solution?
 - a. [Amazon CloudWatch](#)
 - b. [AWS Distro for OpenTelemetry](#)
 - c. [Amazon Managed Service for Prometheus](#)
3. Have I identified and configured my data retention policies to match my security and operational goals?
 - a. [Default data retention for CloudWatch metrics](#)
 - b. [Default data retention for CloudWatch Logs](#)
4. How do you deploy your metric and log aggregation agents?
 - a. [AWS Systems Manager automation](#)
 - b. [OpenTelemetry Collector](#)

Level of effort for the Implementation Plan: There is a *medium* level of effort to identify, track, collect, aggregate, and correlate metrics from all compute resources.

Resources

Related documents:

- [Amazon CloudWatch documentation](#)
- [Collect metrics and logs from Amazon EC2 instances and on-premises servers with the CloudWatch Agent](#)
- [Accessing Amazon CloudWatch Logs for AWS Lambda](#)
- [Using CloudWatch Logs with container instances](#)
- [Publish custom metrics](#)
- [AWS Answers: Centralized Logging](#)
- [AWS Services That Publish CloudWatch Metrics](#)
- [Monitoring Amazon EKS on AWS Fargate](#)

Related videos:

- [Application Performance Management on AWS](#)
- [Build a Monitoring Plan](#)

Related examples:

- [Level 100: Monitoring with CloudWatch Dashboards](#)
- [Level 100: Monitoring Windows EC2 instance with CloudWatch Dashboards](#)
- [Level 100: Monitoring an Amazon Linux EC2 instance with CloudWatch Dashboards](#)

PERF02-BP04 Determine the required configuration by right-sizing

Analyze the various performance characteristics of your workload and how these characteristics relate to memory, network, and CPU usage. Use this data to choose resources that best match your workload's profile. For example, a memory-intensive workload, such as a database, could be served best by the r-family of instances. However, a bursting workload can benefit more from an elastic container system.

Common anti-patterns:

- You choose the largest instance available for all workloads.
- You standardize all instances types to one type for ease of management.

Benefits of establishing this best practice: Being familiar with the AWS compute offerings allows you to determine the correct solution for your various workloads. After you have selected the various compute offerings for your workload, you have the agility to quickly experiment with those compute offerings to determine which ones meet the needs of your workload.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Modify your workload configuration by right sizing: To optimize both performance and overall efficiency, determine which resources your workload needs. Choose memory-optimized instances for systems that require more memory than CPU, or compute-optimized instances for components that do data processing that is not memory-intensive. Right sizing enables your workload to perform as well as possible while only using the required resources

Resources

Related documents:

- [AWS Compute Optimizer](#)
- [Cloud Compute with AWS](#)
- [EC2 Instance Types](#)
- [ECS Containers: Amazon ECS Container Instances](#)
- [EKS Containers: EKS Worker Nodes](#)
- [Functions: Lambda Function Configuration](#)
- [Processor State Control for Your EC2 Instance](#)

Related videos:

- [Amazon EC2 foundations \(CMP211-R2\)](#)
- [Better, faster, cheaper compute: Cost-optimizing Amazon EC2 \(CMP202-R1\)](#)
- [Deliver high performance ML inference with AWS Inferentia \(CMP324-R1\)](#)
- [Optimize performance and cost for your AWS compute \(CMP323-R1\)](#)
- [Powering next-gen Amazon EC2: Deep dive into the Nitro system](#)
- [How to choose compute option for startups](#)

- [Optimize performance and cost for your AWS compute \(CMP323-R1\)](#)

Related examples:

- [Rightsizing with Compute Optimizer and Memory utilization enabled](#)
- [AWS Compute Optimizer Demo code](#)

PERF02-BP05 Use the available elasticity of resources

The cloud provides the flexibility to expand or reduce your resources dynamically through a variety of mechanisms to meet changes in demand. Combined with compute-related metrics, a workload can automatically respond to changes and use the optimal set of resources to achieve its goal.

Optimally matching supply to demand delivers the lowest cost for a workload, but you also must plan for sufficient supply to allow for provisioning time and individual resource failures. Demand can be fixed or variable, requiring metrics and automation to ensure that management does not become a burdensome and disproportionately large cost.

With AWS, you can use a number of different approaches to match supply with demand. The Cost Optimization Pillar whitepaper describes how to use the following approaches to cost:

- Demand-based approach
- Buffer-based approach
- Time-based approach

You must ensure that workload deployments can handle both scale-up and scale-down events. Create test scenarios for scale-down events to ensure that the workload behaves as expected.

Common anti-patterns:

- You react to alarms by manually increasing capacity.
- You leave increased capacity after a scaling event instead of scaling back down.

Benefits of establishing this best practice: Configuring and testing workload elasticity will help save money, maintain performance benchmarks, and improves reliability as traffic changes. Most non-production instances should be stopped when they are not being used. Although it's possible to manually shut down unused instances, this is impractical at larger scales. You can also take advantage of volume-based elasticity, which allows you to optimize performance and cost by automatically increasing the number of compute instances during demand spikes and decreasing capacity when demand decreases.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Take advantage of elasticity: Elasticity matches the supply of resources you have against the demand for those resources. Instances, containers, and functions provide mechanisms for elasticity either in combination with automatic scaling or as a feature of the service. Use elasticity in your architecture to ensure that you have sufficient capacity to meet performance requirements at all scales of use. Ensure that the metrics for scaling up or down elastic resources are validated against the type of workload being deployed. If you are deploying a video transcoding application, 100% CPU utilization is expected and should not be your primary metric. Alternatively, you can measure against the queue depth of transcoding jobs waiting to scale your instance types. Ensure that workload deployments can handle both scale up and scale down events. Scaling down workload components safely is as critical as scaling up resources when demand dictates. Create test scenarios for scale-down events to ensure that the workload behaves as expected.

Resources

Related documents:

- [Cloud Compute with AWS](#)
- [EC2 Instance Types](#)
- [ECS Containers: Amazon ECS Container Instances](#)
- [EKS Containers: EKS Worker Nodes](#)
- [Functions: Lambda Function Configuration](#)
- [Processor State Control for Your EC2 Instance](#)

Related videos:

- [Amazon EC2 foundations \(CMP211-R2\)](#)
- [Better, faster, cheaper compute: Cost-optimizing Amazon EC2 \(CMP202-R1\)](#)
- [Deliver high performance ML inference with AWS Inferentia \(CMP324-R1\)](#)
- [Optimize performance and cost for your AWS compute \(CMP323-R1\)](#)
- [Powering next-gen Amazon EC2: Deep dive into the Nitro system](#)

Related examples:

- [Amazon EC2 Auto Scaling Group Examples](#)
- [Amazon EFS Tutorials](#)

PERF02-BP06 Re-evaluate compute needs based on metrics

Use system-level metrics to identify the behavior and requirements of your workload over time. Evaluate your workload's needs by comparing the available resources with these requirements and make changes to your compute environment to best match your workload's profile. For example, over time a system might be observed to be more memory-intensive than initially thought, so moving to a different instance family or size could improve both performance and efficiency.

Common anti-patterns:

- You only monitor system-level metrics to gain insight into your workload.
- You architect your compute needs for peak workload requirements.
- You oversize the compute solution to meet scaling or performance requirements when moving to a new compute solution would match your workload characteristics

Benefits of establishing this best practice: To optimize performance and resource utilization, you need a unified operational view, real-time granular data, and a historical reference. You can create automatic dashboards to visualize this data and perform metric math to derive operational and utilization insights.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Use a data-driven approach to optimize resources: To achieve maximum performance and efficiency, use the data gathered over time from your workload to tune and optimize your resources. Look at the trends in your workload's usage of current resources and determine where you can make changes to better match your workload's needs. When resources are over-committed, system performance degrades, whereas underutilization results in a less efficient use of resources and higher cost.

Resources

Related documents:

- [Cloud Compute with AWS](#)
- [AWS Compute Optimizer](#)
- [Cloud Compute with AWS](#)
- [EC2 Instance Types](#)
- [ECS Containers: Amazon ECS Container Instances](#)
- [EKS Containers: EKS Worker Nodes](#)
- [Functions: Lambda Function Configuration](#)
- [Processor State Control for Your EC2 Instance](#)

Related videos:

- [Amazon EC2 foundations \(CMP211-R2\)](#)
- [Better, faster, cheaper compute: Cost-optimizing Amazon EC2 \(CMP202-R1\)](#)
- [Deliver high performance ML inference with AWS Inferentia \(CMP324-R1\)](#)
- [Optimize performance and cost for your AWS compute \(CMP323-R1\)](#)
- [Powering next-gen Amazon EC2: Deep dive into the Nitro system](#)

Related examples:

- [Rightsizing with Compute Optimizer and Memory utilization enabled](#)
- [AWS Compute Optimizer Demo code](#)

PERF 3 How do you select your storage solution?

The optimal storage solution for a system varies based on the kind of access method (block, file, or object), patterns of access (random or sequential), required throughput, frequency of access (online, offline, archival), frequency of update (WORM, dynamic), and availability and durability constraints. Well-architected systems use multiple storage solutions and enable different features to improve performance and use resources efficiently.

Best practices

- [PERF03-BP01 Understand storage characteristics and requirements \(p. 296\)](#)
- [PERF03-BP02 Evaluate available configuration options \(p. 300\)](#)
- [PERF03-BP03 Make decisions based on access patterns and metrics \(p. 301\)](#)

PERF03-BP01 Understand storage characteristics and requirements

Identify and document the workload storage needs and define the storage characteristics of each location. Examples of storage characteristics include: shareable access, file size, growth rate, throughput, IOPS, latency, access patterns, and persistence of data. Use these characteristics to evaluate if block, file, object, or instance storage services are the most efficient solution for your storage needs.

Desired outcome: Identify and document the storage requirements per storage requirement and evaluate the available storage solutions. Based on the key storage characteristics, your team will understand how the selected storage services will benefit your workload performance. Key criteria include data access patterns, growth rate, scaling needs, and latency requirements.

Common anti-patterns:

- You only use one storage type, such as Amazon Elastic Block Store (Amazon EBS), for all workloads.
- You assume that all workloads have similar storage access performance requirements.

Benefits of establishing this best practice: Selecting the storage solution based on the identified and required characteristics will help improve your workloads performance, decrease costs and lower your operational efforts in maintaining your workload. Your workload performance will benefit from the solution, configuration, and location of the storage service.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Identify your workload's most important storage performance metrics and implement improvements as part of a data-driven approach, using benchmarking or load testing. Use this data to identify where your storage solution is constrained, and examine configuration options to improve the solution. Determine the expected growth rate for your workload and choose a storage solution that will meet those rates. Research the AWS storage offerings to determine the correct storage solution for your various workload needs. Provisioning storage solutions in AWS increases the opportunity for you to test storage offerings and determine if they are appropriate for your workload needs.

AWS service	Key characteristics	Common use cases
Amazon S3	99.999999999% durability, unlimited growth, accessible from anywhere, several cost models based on access and resiliency	Cloud-native application data, data archiving, and backups, analytics, data lakes, static website hosting, IoT data
Amazon S3 Glacier	Seconds to hours latency, unlimited growth, lowest cost, long-term storage	Data archiving, media archives, long-term backup retention.
Amazon EBS	Storage size requires management and monitoring, low latency, persistent storage, 99.8% to 99.9% durability, most volume types are accessible only from one EC2 instance.	COTS applications, I/O intensive applications, relational and NoSQL databases, backup and recovery
EC2 Instance Store	Pre-determined storage size, lowest latency, not persisted, accessible only from one EC2 instance	COTS applications, I/O intensive applications, in-memory data store
Amazon EFS	99.999999999% durability, unlimited growth, accessible by multiple compute services	Modernized applications sharing files across multiple compute services, file storage for scaling content management systems
Amazon FSx	Supports four file systems (NetApp, OpenZFS, Windows File Server, and Amazon FSx for Lustre), storage available different per file system, accessible by multiple compute services	Cloud native workloads, private cloud bursting, migrated workloads that require a specific file system, VMC, ERP systems, on-premises file storage and backups

AWS service	Key characteristics	Common use cases
Snow family	Portable devices, 256-bit encryption, NFS endpoint, on-board computing, TBs of storage	Migrating data to the cloud, storage, and computing in extreme on-premises conditions, disaster recovery, remote data collection
AWS Storage Gateway	Provides low-latency on-premises access to cloud-backed storage, fully managed on-premises cache	On-premises data to cloud migrations, populate cloud data lakes from on-premises sources, modernized file sharing.

Implementation steps:

1. Use benchmarking or load tests to collect the key characteristics of your storage needs. Key characteristics include:
 - a. Shareable (what components access this storage)
 - b. Growth rate
 - c. Throughput
 - d. Latency
 - e. I/O size
 - f. Durability
 - g. Access patterns (reads vs writes, frequency, spikey, or consistent)
2. Identify the type of storage solution that supports your storage characteristics.
 - a. [Amazon S3](#) is an object storage service with unlimited scalability, high availability, and multiple options for accessibility. Transferring and accessing objects in and out of Amazon S3 can use a service, such as [Transfer Acceleration](#) or [Access Points](#) to support your location, security needs, and access patterns. Use the [Amazon S3 performance guidelines](#) to help you optimize your Amazon S3 configuration to meet your workload performance needs.
 - b. [Amazon S3 Glacier](#) is a storage class of Amazon S3 built for data archiving. You can choose from three archiving solutions ranging from millisecond access to 5-12 hour access with different cost and security options. Amazon S3 Glacier can help you meet performance requirements by implementing a data lifecycle that supports your business requirements and data characteristics.
 - c. [Amazon Elastic Block Store \(Amazon EBS\)](#) is a high-performance block storage service designed for Amazon Elastic Compute Cloud (Amazon EC2). You can choose from [SSD- or HDD-based](#) solutions with different characteristics that prioritize [IOPS](#) or [throughput](#). EBS volumes are well suited for high-performance workloads, primary storage for file systems, databases, or applications that can only access attached stage systems.
 - d. [Amazon EC2 Instance Store](#) is similar to Amazon EBS as it attaches to an Amazon EC2 instance however, the Instance Store is only temporary storage that should ideally be used as a buffer, cache, or other temporary content. You cannot detach an Instance Store and all data is lost if the instance shuts down. Instance Stores can be used for high I/O performance and low latency use cases where data doesn't need to persist.
 - e. [Amazon Elastic File System \(Amazon EFS\)](#) is a mountable file system that can be accessed by multiple types of compute solutions. Amazon EFS automatically grows and shrinks storage and is performance-optimized to deliver consistent low latencies. EFS has [two performance configuration modes](#): General Purpose and Max I/O. General Purpose has a sub-millisecond read latency and a single-digit millisecond write latency. The Max I/O feature can support thousands of compute instance requiring a shared file system. Amazon EFS supports [two throughput modes](#): Bursting and Provisioned. A workload that experiences a spikey access pattern will benefit from the bursting throughput mode while a workload that is consistently high would be performant with a provisioned throughput mode.

- f. [Amazon FSx](#) is built on the latest AWS compute solutions to support four commonly used file systems: NetApp ONTAP, OpenZFS, Windows File Server, and Lustre. Amazon FSx [latency, throughput, and IOPS](#) vary per file system and should be considered when selecting the right file system for your workload needs.
 - g. [AWS Snow Family](#) are storage and compute devices that support online and offline data migration to the cloud and data storage and computing on premises. AWS Snow devices support collecting large amounts of on-premises data, processing of that data and moving that data to the cloud. There are several [documented performance best practices](#) when it comes to the number of files, file sizes, and compression.
 - h. [AWS Storage Gateway](#) provides on-premises applications access to cloud-based storage. AWS Storage Gateway supports multiple cloud storage services including Amazon S3, Amazon S3 Glacier, Amazon FSx, and Amazon EBS. It supports a number of protocols such as iSCSI, SMB, and NFS. It provides low-latency performance by caching frequently accessed data on premises and only sends changed data and compressed data to AWS.
3. After you have experimented with your new storage solution and identified the optimal configuration, plan your migration and validate your performance metrics. This is a continual process, and should be reevaluated when key characteristics change or available services or options change.

Level of effort for the implementation plan: If a workload is moving from one storage solution to another, there could be a *moderate* level of effort involved in refactoring the application.

Resources

Related documents:

- [Amazon EBS Volume Types](#)
- [Amazon EC2 Storage](#)
- [Amazon EFS: Amazon EFS Performance](#)
- [Amazon FSx for Lustre Performance](#)
- [Amazon FSx for Windows File Server Performance](#)
- [Amazon FSx for NetApp ONTAP performance](#)
- [Amazon FSx for OpenZFS performance](#)
- [Amazon S3 Glacier: Amazon S3 Glacier Documentation](#)
- [Amazon S3: Request Rate and Performance Considerations](#)
- [Cloud Storage with AWS](#)
- [AWS Snow Family](#)
- [EBS I/O Characteristics](#)

Related videos:

- [Deep dive on Amazon EBS \(STG303-R1\)](#)
- [Optimize your storage performance with Amazon S3 \(STG343\)](#)

Related examples:

- [Amazon EFS CSI Driver](#)
- [Amazon EBS CSI Driver](#)
- [Amazon EFS Utilities](#)
- [Amazon EBS Autoscale](#)
- [Amazon S3 Examples](#)
- [Amazon FSx for Lustre Container Storage Interface \(CSI\) Driver](#)

PERF03-BP02 Evaluate available configuration options

Evaluate the various characteristics and configuration options and how they relate to storage. Understand where and how to use provisioned IOPS, SSDs, magnetic storage, object storage, archival storage, or ephemeral storage to optimize storage space and performance for your workload.

[Amazon EBS](#) provides a range of options that allow you to optimize storage performance and cost for your workload. These options are divided into two major categories: SSD-backed storage for transactional workloads, such as databases and boot volumes (performance depends primarily on IOPS), and HDD-backed storage for throughput-intensive workloads, such as MapReduce and log processing (performance depends primarily on MB/s).

SSD-backed volumes include the highest performance provisioned IOPS SSD for latency-sensitive transactional workloads and general-purpose SSD that balance price and performance for a wide variety of transactional data.

[Amazon S3 transfer acceleration](#) enables fast transfer of files over long distances between your client and your S3 bucket. Transfer acceleration leverages Amazon CloudFront globally distributed edge locations to route data over an optimized network path. For a workload in an S3 bucket that has intensive GET requests, use Amazon S3 with CloudFront. When uploading large files, use multi-part uploads with multiple parts uploading at the same time to help maximize network throughput.

[Amazon Elastic File System \(Amazon EFS\)](#) provides a simple, scalable, fully managed elastic NFS file system for use with AWS Cloud services and on-premises resources. To support a wide variety of cloud storage workloads, Amazon EFS offers two performance modes: general purpose performance mode, and max I/O performance mode. There are also two throughput modes to choose from for your file system: Bursting Throughput, and Provisioned Throughput. To determine which settings to use for your workload, see the [Amazon EFS User Guide](#).

[Amazon FSx](#) provides four file systems to choose from: [Amazon FSx for Windows File Server](#) for enterprise workloads, [Amazon FSx for Lustre](#) for high-performance workloads, [Amazon FSx for NetApp ONTAP](#) for NetApps popular ONTAP file system, and [Amazon FSx for OpenZFS](#) for Linux-based file servers. FSx is SSD-backed and is designed to deliver fast, predictable, scalable, and consistent performance. Amazon FSx file systems deliver sustained high read and write speeds and consistent low latency data access. You can choose the throughput level you need to match your workload's needs.

Common anti-patterns:

- You only use one storage type, such as Amazon EBS, for all workloads.
- You use Provisioned IOPS for all workloads without real-world testing against all storage tiers.
- You assume that all workloads have similar storage access performance requirements.

Benefits of establishing this best practice: Evaluating all storage service options can reduce the cost of infrastructure and the effort required to maintain your workloads. It can potentially accelerate your time to market for deploying new services and features.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Determine storage characteristics: When you evaluate a storage solution, determine which storage characteristics you require, such as ability to share, file size, cache size, latency, throughput, and persistence of data. Then match your requirements to the AWS service that best fits your needs.

Resources

Related documents:

- [Cloud Storage with AWS](#)
- [Amazon EBS Volume Types](#)
- [Amazon EC2 Storage](#)
- [Amazon EFS: Amazon EFS Performance](#)
- [Amazon FSx for Lustre Performance](#)
- [Amazon FSx for Windows File Server Performance](#)
- [Amazon Glacier: Amazon Glacier Documentation](#)
- [Amazon S3: Request Rate and Performance Considerations](#)
- [Cloud Storage with AWS](#)
- [Cloud Storage with AWS](#)
- [EBS I/O Characteristics](#)

Related videos:

- [Deep dive on Amazon EBS \(STG303-R1\)](#)
- [Optimize your storage performance with Amazon S3 \(STG343\)](#)

Related examples:

- [Amazon EFS CSI Driver](#)
- [Amazon EBS CSI Driver](#)
- [Amazon EFS Utilities](#)
- [Amazon EBS Autoscale](#)
- [Amazon S3 Examples](#)

PERF03-BP03 Make decisions based on access patterns and metrics

Choose storage systems based on your workload's access patterns and configure them by determining how the workload accesses data. Increase storage efficiency by choosing object storage over block storage. Configure the storage options you choose to match your data access patterns.

How you access data impacts how the storage solution performs. Select the storage solution that aligns best to your access patterns, or consider changing your access patterns to align with the storage solution to maximize performance.

Creating a RAID 0 array allows you to achieve a higher level of performance for a file system than what you can provision on a single volume. Consider using RAID 0 when I/O performance is more important than fault tolerance. For example, you could use it with a heavily used database where data replication is already set up separately.

Select appropriate storage metrics for your workload across all of the storage options consumed for the workload. When using filesystems that use burst credits, create alarms to let you know when you are approaching those credit limits. You must create storage dashboards to show the overall workload storage health.

For storage systems that are a fixed size, such as Amazon EBS or Amazon FSx, ensure that you are monitoring the amount of storage used versus the overall storage size and create automation if possible to increase the storage size when reaching a threshold

Common anti-patterns:

- You assume that storage performance is adequate if customers are not complaining.

- You only use one tier of storage, assuming all workloads fit within that tier.

Benefits of establishing this best practice: You need a unified operational view, real-time granular data, and historical reference to optimize performance and resource utilization. You can create automatic dashboards and data with one-second granularity to perform metric math on your data and derive operational and utilization insights for your storage needs.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Optimize your storage usage and access patterns: Choose storage systems based on your workload's access patterns and the characteristics of the available storage options. Determine the best place to store data that will enable you to meet your requirements while reducing overhead. Use performance optimizations and access patterns when configuring and interacting with data based on the characteristics of your storage (for example, striping volumes or partitioning data).

Select appropriate metrics for storage options: Ensure that you select the appropriate storage metrics for the workload. Each storage option offers various metrics to track how your workload performs over time. Ensure that you are measuring against any storage burst metrics (for example, monitoring burst credits for Amazon EFS). For storage systems that are fixed sized, such as Amazon Elastic Block Store or Amazon FSx, ensure that you are monitoring the amount of storage used versus the overall storage size. Create automation when possible to increase the storage size when reaching a threshold.

Monitor metrics: Amazon CloudWatch can collect metrics across the resources in your architecture. You can also collect and publish custom metrics to surface business or derived metrics. Use CloudWatch or third-party solutions to set alarms that indicate when thresholds are breached.

Resources

Related documents:

- [Amazon EBS Volume Types](#)
- [Amazon EC2 Storage](#)
- [Amazon EFS: Amazon EFS Performance](#)
- [Amazon FSx for Lustre Performance](#)
- [Amazon FSx for Windows File Server Performance](#)
- [Amazon Glacier: Amazon Glacier Documentation](#)
- [Amazon S3: Request Rate and Performance Considerations](#)
- [Cloud Storage with AWS](#)
- [EBS I/O Characteristics](#)
- [Monitoring and understanding Amazon EBS performance using Amazon CloudWatch](#)

Related videos:

- [Deep dive on Amazon EBS \(STG303-R1\)](#)
- [Optimize your storage performance with Amazon S3 \(STG343\)](#)

Related examples:

- [Amazon EFS CSI Driver](#)
- [Amazon EBS CSI Driver](#)
- [Amazon EFS Utilities](#)

- [Amazon EBS Autoscale](#)
- [Amazon S3 Examples](#)

PERF 4 How do you select your database solution?

The optimal database solution for a system varies based on requirements for availability, consistency, partition tolerance, latency, durability, scalability, and query capability. Many systems use different database solutions for various subsystems and enable different features to improve performance. Selecting the wrong database solution and features for a system can lead to lower performance efficiency.

Best practices

- [PERF04-BP01 Understand data characteristics \(p. 303\)](#)
- [PERF04-BP02 Evaluate the available options \(p. 307\)](#)
- [PERF04-BP03 Collect and record database performance metrics \(p. 312\)](#)
- [PERF04-BP04 Choose data storage based on access patterns \(p. 314\)](#)
- [PERF04-BP05 Optimize data storage based on access patterns and metrics \(p. 317\)](#)

PERF04-BP01 Understand data characteristics

Choose your data management solutions to optimally match the characteristics, access patterns, and requirements of your workload datasets. When selecting and implementing a data management solution, you must ensure that the querying, scaling, and storage characteristics support the workload data requirements. Learn how various database options match your data models, and which configuration options are best for your use-case.

AWS provides numerous database engines including relational, key-value, document, in-memory, graph, time series, and ledger databases. Each data management solution has options and configurations available to you to support your use-cases and data models. Your workload might be able to use several different database solutions, based on the data characteristics. By selecting the best database solutions to a specific problem, you can break away from monolithic databases, with the one-size-fits-all approach that is restrictive and focus on managing data to meet your customer's need.

Desired outcome: The workload data characteristics are documented with enough detail to facilitate selection and configuration of supporting database solutions, and provide insight into potential alternatives.

Common anti-patterns:

- Not considering ways to segment large datasets into smaller collections of data that have similar characteristics, resulting in missing opportunities to use more purpose-built databases that better match data and growth characteristics.
- Not identifying the data access patterns up front, which leads to costly and complex rework later.
- Limiting growth by using data storage strategies that don't scale as quickly as is needed
- Choosing one database type and vendor for all workloads.
- Sticking to one database solution because there is internal experience and knowledge of one particular type of database solution.
- Keeping a database solution because it worked well in an on-premises environment.

Benefits of establishing this best practice: Be familiar with all of the AWS database solutions so that you can determine the correct database solution for your various workloads. After you select the

appropriate database solution for your workload, you can quickly experiment on each of those database offerings to determine if they continue to meet your workload needs.

Level of risk exposed if this best practice is not established: High

- Potential cost savings may not be identified.
- Data may not be secured to the level required.
- Data access and storage performance may not be optimal.

Implementation guidance

Define the data characteristics and access patterns of your workload. Review all available database solutions to identify which solution supports your data requirements. Within a given workload, multiple databases may be selected. Evaluate each service or group of services and assess them individually. If potential alternative data management solutions are identified for part or all of the data, experiment with alternative implementations that might unlock cost, security, performance, and reliability benefits. Update existing documentation, should a new data management approach be adopted.

Type	AWS Services	Key Characteristics	Common use-cases
Relational	Amazon RDS, Amazon Aurora	Referential integrity, ACID transactions, schema on write	ERP, CRM, Commercial off-the-shelf software
Key Value	Amazon DynamoDB	High throughput, low latency, near-infinite scalability	Shopping carts (ecommerce), product catalogs, chat applications
Document	Amazon DocumentDB	Store JSON documents and query on any attribute	Content Management (CMS), customer profiles, mobile applications
In Memory	Amazon ElastiCache, Amazon MemoryDB	Microsecond latency	Caching, game leaderboards
Graph	Amazon Neptune	Highly relational data where the relationships between data have meaning	Social networks, personalization engines, fraud detection
Time Series	Amazon Timestream	Data where the primary dimension is time	DevOps, IoT, Monitoring
Wide column	Amazon Keyspaces	Cassandra workloads.	Industrial equipment maintenance, route optimization
Ledger	Amazon QLDB	Immutable and cryptographically verifiable ledger of changes	Systems of record, healthcare, supply chains, financial institutions

Implementation steps

1. How is the data structured? (for example, unstructured, key-value, semi-structured, relational)

- a. If the data is unstructured, consider an object-store such as [Amazon S3](#) or a NoSQL database such as [Amazon DocumentDB](#).
- b. For key-value data, consider [DynamoDB](#), [ElastiCache for Redis](#) or [MemoryDB](#).
- c. If the data has a relational structure, what level of referential integrity is required?
 - i. For foreign key constraints, relational databases such as [Amazon RDS](#) and [Aurora](#) can provide this level of integrity.
 - ii. Typically, within a NoSQL data-model, you would de-normalize your data into a single document or collection of documents to be retrieved in a single request rather than joining across documents or tables.
2. Is ACID (atomicity, consistency, isolation, durability) compliance required?
 - a. If the ACID properties associated with relational databases are required, consider a relational database such as [Amazon RDS](#) and [Aurora](#).
3. What consistency model is required?
 - a. If your application can tolerate eventual consistency, consider a NoSQL implementation. Review the other characteristics to help choose which [NoSQL database](#) is most appropriate.
 - b. If strong consistency is required, you can use strongly consistent reads with [DynamoDB](#) or a relational database such as [Amazon RDS](#).
4. What query and result formats must be supported? (for example, SQL, CSV, Parquet, Avro, JSON, etc.)
5. What data types, field sizes and overall quantities are present? (for example, text, numeric, spatial, time-series calculated, binary or blob, document)
6. How will the storage requirements change over time? How does this impact scalability?
 - a. Serverless databases such as [DynamoDB](#) and [Amazon Quantum Ledger Database](#) will scale dynamically up to near-unlimited storage.
 - b. Relational databases have upper bounds on provisioned storage, and often must be horizontally partitioned via mechanisms such as sharding once they reach these limits.
7. What is the proportion of read queries in relation to write queries? Would caching be likely to improve performance?
 - a. Read-heavy workloads can benefit from a caching layer, this could be [ElastiCache](#) or [DAX](#) if the database is [DynamoDB](#).
 - b. Reads can also be offloaded to read replicas with relational databases such as [Amazon RDS](#).
8. Does storage and modification (OLTP - Online Transaction Processing) or retrieval and reporting (OLAP - Online Analytical Processing) have a higher priority?
 - a. For high-throughput transactional processing, consider a NoSQL database such as [DynamoDB](#) or [Amazon DocumentDB](#).
 - b. For analytical queries, consider a columnar database such as [Amazon Redshift](#) or exporting the data to [Amazon S3](#) and performing analytics using [Athena](#) or [QuickSight](#).
9. How sensitive is this data and what level of protection and encryption does it require?
 - a. All [Amazon RDS](#) and [Aurora](#) engines support data encryption at rest using [AWS KMS](#). [Microsoft SQL Server](#) and [Oracle](#) also support native Transparent Data Encryption (TDE) when using [Amazon RDS](#).
 - b. For [DynamoDB](#), you can use fine-grained access control with [IAM](#) to control who has access to what data at the key level.
10. What level of durability does the data require?
 - a. [Aurora](#) automatically replicates your data across three Availability Zones within a Region, meaning your data is highly durable with less chance of data loss.
 - b. [DynamoDB](#) is automatically replicated across multiple Availability Zones, providing high availability and data durability.
 - c. [Amazon S3](#) provides 11 9s of durability. Many database services such as [Amazon RDS](#) and [DynamoDB](#) support exporting data to [Amazon S3](#) for long-term retention and archival.
11. Do [Recovery Time Objective \(RTO\)](#) or [Recovery Point Objectives \(RPO\)](#) requirements influence the solution?

- a. Amazon RDS, Aurora, DynamoDB, Amazon DocumentDB, and Neptune all support point in time recovery and on-demand backup and restore.
 - b. For high availability requirements, DynamoDB tables can be replicated globally using the [Global Tables](#) feature and Aurora clusters can be replicated across multiple Regions using the Global database feature. Additionally, S3 buckets can be replicated across AWS Regions using cross-region replication.
- 12Is there a desire to move away from commercial database engines / licensing costs?
- a. Consider open-source engines such as PostgreSQL and MySQL on Amazon RDS or Aurora
 - b. Leverage [AWS DMS](#) and [AWS SCT](#) to perform migrations from commercial database engines to open-source
- 13What is the operational expectation for the database? Is moving to managed services a primary concern?
- a. Leveraging Amazon RDS instead of Amazon EC2, and DynamoDB or Amazon DocumentDB instead of self-hosting a NoSQL database can reduce operational overhead.
- 14How is the database currently accessed? Is it only application access, or are there Business Intelligence (BI) users and other connected off-the-shelf applications?
- a. If you have dependencies on external tooling then you may have to maintain compatibility with the databases they support. Amazon RDS is fully compatible with the difference engine versions that it supports including Microsoft SQL Server, Oracle, MySQL, and PostgreSQL.
- 15The following is a list of potential data management services, and where these can best be used:
- a. Relational databases store data with predefined schemas and relationships between them. These databases are designed to support ACID (atomicity, consistency, isolation, durability) transactions, and maintain referential integrity and strong data consistency. Many traditional applications, enterprise resource planning (ERP), customer relationship management (CRM), and ecommerce use relational databases to store their data. You can run many of these database engines on Amazon EC2, or choose from one of the AWS-managed [database services](#): [Amazon Aurora](#), [Amazon RDS](#), and [Amazon Redshift](#).
 - b. Key-value databases are optimized for common access patterns, typically to store and retrieve large volumes of data. These databases deliver quick response times, even in extreme volumes of concurrent requests. High-traffic web apps, ecommerce systems, and gaming applications are typical use-cases for key-value databases. In AWS, you can utilize [Amazon DynamoDB](#), a fully managed, multi-Region, multi-master, durable database with built-in security, backup and restore, and in-memory caching for internet-scale applications.
 - c. In-memory databases are used for applications that require real-time access to data, lowest latency and highest throughput. By storing data directly in memory, these databases deliver microsecond latency to applications where millisecond latency is not enough. You may use in-memory databases for application caching, session management, gaming leaderboards, and geospatial applications. [Amazon ElastiCache](#) is a fully managed in-memory data store, compatible with [Redis](#) or [Memcached](#). In case the applications also higher durability requirements, [Amazon MemoryDB for Redis](#) offers this in combination being a durable, in-memory database service for ultra-fast performance.
 - d. A document database is designed to store semistructured data as JSON-like documents. These databases help developers build and update applications such as content management, catalogs, and user profiles quickly. [Amazon DocumentDB](#) is a fast, scalable, highly available, and fully managed document database service that supports MongoDB workloads.
 - e. A wide column store is a type of NoSQL database. It uses tables, rows, and columns, but unlike a relational database, the names and format of the columns can vary from row to row in the same table. You typically see a wide column store in high scale industrial apps for equipment maintenance, fleet management, and route optimization. [Amazon Keyspaces \(for Apache Cassandra\)](#) is a wide column scalable, highly available, and managed Apache Cassandra-compatible database service.
 - f. Graph databases are for applications that must navigate and query millions of relationships between highly connected graph datasets with millisecond latency at large scale. Many companies

use graph databases for fraud detection, social networking, and recommendation engines. [Amazon Neptune](#) is a fast, reliable, fully managed graph database service that makes it easy to build and run applications that work with highly connected datasets.

- g. Time-series databases efficiently collect, synthesize, and derive insights from data that changes over time. IoT applications, DevOps, and industrial telemetry can utilize time-series databases. [Amazon Timestream](#) is a fast, scalable, fully managed time series database service for IoT and operational applications that makes it easy to store and analyze trillions of events per day.
- h. Ledger databases provide a centralized and trusted authority to maintain a scalable, immutable, and cryptographically verifiable record of transactions for every application. We see ledger databases used for systems of record, supply chain, registrations, and even banking transactions. [Amazon Quantum Ledger Database \(Amazon QLDB\)](#) is a fully managed ledger database that provides a transparent, immutable, and cryptographically verifiable transaction log owned by a central trusted authority. Amazon QLDB tracks every application data change and maintains a complete and verifiable history of changes over time.

Level of effort for the implementation plan: If a workload is moving from one database solution to another, there could be a *high* level of effort involved in refactoring the data and application.

Resources

Related documents:

- [Cloud Databases with AWS](#)
- [AWS Database Caching](#)
- [Amazon DynamoDB Accelerator](#)
- [Amazon Aurora best practices](#)
- [Amazon Redshift performance](#)
- [Amazon Athena top 10 performance tips](#)
- [Amazon Redshift Spectrum best practices](#)
- [Amazon DynamoDB best practices](#)
- [Choose between EC2 and Amazon RDS](#)
- [Best Practices for Implementing Amazon ElastiCache](#)

Related videos:

- [AWS purpose-built databases \(DAT209-L\)](#)
- [Amazon Aurora storage demystified: How it all works \(DAT309-R\)](#)
- [Amazon DynamoDB deep dive: Advanced design patterns \(DAT403-R1\)](#)

Related examples:

- [Optimize Data Pattern using Amazon Redshift Data Sharing](#)
- [Database Migrations](#)
- [MS SQL Server - AWS Database Migration Service \(DMS\) Replication Demo](#)
- [Database Modernization Hands On Workshop](#)
- [Amazon Neptune Samples](#)

PERF04-BP02 Evaluate the available options

Understand the available database options and how it can optimize your performance before you select your data management solution. Use load testing to identify database metrics that matter for your

workload. While you explore the database options, take into consideration various aspects such as the parameter groups, storage options, memory, compute, read replica, eventual consistency, connection pooling, and caching options. Experiment with these various configuration options to improve the metrics.

Desired outcome: A workload could have one or more database solutions used based on data types. The database functionality and benefits optimally match the data characteristics, access patterns, and workload requirements. To optimize your database performance and cost, you must evaluate the data access patterns to determine the appropriate database options. Evaluate the acceptable query times to ensure that the selected database options can meet the requirements.

Common anti-patterns:

- Not identifying the data access patterns.
- Not being aware of the configuration options of your chosen data management solution.
- Relying solely on increasing the instance size without looking at other available configuration options.
- Not testing the scaling characteristics of the chosen solution.

Benefits of establishing this best practice: By exploring and experimenting with the database options you may be able to reduce the cost of infrastructure, improve performance and scalability and lower the effort required to maintain your workloads.

Level of risk exposed if this best practice is not established: High

- Having to optimize for a *one size fits all* database means making unnecessary compromises.
- Higher costs as a result of not configuring the database solution to match the traffic patterns.
- Operational issues may emerge from scaling issues.
- Data may not be secured to the level required.

Implementation guidance

Understand your workload data characteristics so that you can configure your database options. Run load tests to identify your key performance metrics and bottlenecks. Use these characteristics and metrics to evaluate database options and experiment with different configurations.

AWS Services	Amazon RDS, Amazon Aurora	Amazon DynamoDB	Amazon DocumentDB	Amazon ElastiCache	Amazon Neptune	Amazon Timestream	Amazon Keyspaces	Amazon QLDB
Scaling Compute	Increase instance size, Aurora Serverless instances autoscale in response to changes in load	Automatic read/write scaling with on-demand capacity mode or automatic scaling of provisioned read/write capacity	Increase instance size	Increase instance size, add nodes to cluster	Increase instance size	Automatically scales to adjust capacity	Automatic read/write scaling with on-demand capacity mode or automatic scaling of provisioned read/write capacity	Automatically scales to adjust capacity

AWS Services	Amazon RDS, Amazon Aurora	Amazon DynamoDB	Amazon DocumentDB	Amazon ElastiCache	Amazon Neptune	Amazon Timestream	Amazon Keyspaces	Amazon QLDB
		in provisioned capacity mode					in provisioned capacity mode	
Scaling-out reads	All engines support read replicas. Aurora supports automatic scaling of read replica instances	Increase provisioned read capacity units	Read replicas	Read replicas	Read replicas. Supports automatic scaling of read replica instances	Automatically scales	Automatically increase provisioned read capacity units	Automatically scales up to documented concurrency limits
Scaling-out writes	Increasing instance size, batching writes in the application or adding a queue in front of the database. Horizontal scaling via application-level sharding across multiple instances	Increase provisioned write capacity units. Ensuring optimal partition key to prevent partition level write throttling	Increasing primary instance size	Using Redis in cluster mode to distribute writes across shards	Increasing instance size	Write requests may be throttled while scaling. If you encounter throttling exceptions, continue to send data at the same (or higher) throughput to automatically scale. Batch writes to reduce concurrent write requests	Increase provisioned write capacity units. Ensuring optimal partition key to prevent partition level write throttling	Automatically scales up to documented concurrency limits
Engine configuration	Parameter groups	Not applicable	Parameter groups	Parameter groups	Parameter groups	Not applicable	Not applicable	Not applicable

AWS Services	Amazon RDS, Amazon Aurora	Amazon DynamoDB	Amazon DocumentDB	Amazon ElastiCache	Amazon Neptune	Amazon Timestream	Amazon Keyspaces	Amazon QLDB
Caching	In-memory caching, configurable via parameter groups. Pair with a dedicated cache such as ElastiCache for Redis to offload requests for commonly accessed items	DAX (DAX) fully managed cache available	In-memory caching. Optionally, pair with a dedicated cache such as ElastiCache for Redis to offload requests for commonly accessed items	Primary function is caching	Use the query results cache to cache the result of a read-only query	Timestream has two storage tiers; one of these is a high-performance in-memory tier	Deploy a separate dedicated cache such as ElastiCache or Redis to offload requests for commonly accessed items	Not applicable
High availability disaster recovery	Recommended configuration for production workloads is to run a standby instance in a second Availability Zone to provide resiliency within a Region. For resiliency across Regions, Aurora Global Database can be used	Highly available within a Region. Tables can be replicated across Regions using DynamoDB global tables	Create multiple instances across Availability Zones for availability. Snapshots can be shared across Regions and clusters can be replicated using DMS to provide Cross-Region Replication / disaster recovery	Recommended configuration for production clusters is to create at least one node in a secondary Availability Zone. ElastiCache Global Datastore can be used to replicate clusters across Regions.	Read replicas in other Availability Zones serve as failover targets. Snapshots can be shared across Region and clusters can be replicated using Neptune streams to replicate data between two clusters in two different Regions.	Highly available within a Region. cross-Region replication requires custom application development using the Timestream SDK	Highly available within a Region. Cross-Region Replication requires custom application logic or third-party tools	Highly available within a Region. To replicate across Regions, export the contents of the Amazon QLDB journal to a S3 bucket and configure the bucket for Cross-Region Replication.

Implementation steps

1. What configuration options are available for the selected databases?
 - a. Parameter Groups for Amazon RDS and Aurora allow you to adjust common database engine level settings such as the memory allocated for the cache or adjusting the time zone of the database
 - b. For provisioned database services such as Amazon RDS, Aurora, Neptune, Amazon DocumentDB and those deployed on Amazon EC2 you can change the instance type, provisioned storage and add read replicas.
 - c. DynamoDB allows you to specify two capacity modes: on-demand and provisioned. To account for differing workloads, you can change between these modes and increase the allocated capacity in provisioned mode at any time.
2. Is the workload read or write heavy?
 - a. What solutions are available for offloading reads (read replicas, caching, etc.)?
 - i. For DynamoDB tables, you can offload reads using DAX for caching.
 - ii. For relational databases, you can create an ElastiCache for Redis cluster and configure your application to read from the cache first, falling back to the database if the requested item is not present.
 - iii. Relational databases such as Amazon RDS and Aurora, and provisioned NoSQL databases such as Neptune and Amazon DocumentDB all support adding read replicas to offload the read portions of the workload.
 - iv. Serverless databases such as DynamoDB will scale automatically. Ensure that you have enough read capacity units (RCU) provisioned to handle the workload.
 - b. What solutions are available for scaling writes (partition key sharding, introducing a queue, etc.)?
 - i. For relational databases, you can increase the size of the instance to accommodate an increased workload or increase the provisioned IOPs to allow for an increased throughput to the underlying storage.
 - You can also introduce a queue in front of your database rather than writing directly to the database. This pattern allows you to decouple the ingestion from the database and control the flow-rate so the database does not get overwhelmed.
 - Batching your write requests rather than creating many short-lived transactions can help improve throughput in high-write volume relational databases.
 - ii. Serverless databases like DynamoDB can scale the write throughput automatically or by adjusting the provisioned write capacity units (WCU) depending on the capacity mode.
 - You can still run into issues with *hot* partitions though, when you reach the throughput limits for a given partition key. This can be mitigated by choosing a more evenly distributed partition key or by write-sharding the partition key.
3. What are the current or expected peak transactions per second (TPS)? Test using this volume of traffic and this volume +X% to understand the scaling characteristics.
 - a. Native tools such as pg_bench for PostgreSQL can be used to stress-test the database and understand the bottlenecks and scaling characteristics.
 - b. Production-like traffic should be captured so that it can be replayed to simulate real-world conditions in addition to synthetic workloads.
4. If using serverless or elastically scalable compute, test the impact of scaling this on the database. If appropriate, introduce connection management or pooling to lower impact on the database.
 - a. RDS Proxy can be used with Amazon RDS and Aurora to manage connections to the database.
 - b. Serverless databases such as DynamoDB do not have connections associated with them, but consider the provisioned capacity and automatic scaling policies to deal with spikes in load.
5. Is the load predictable, are there spikes in load and periods of inactivity?
 - a. If there are periods of inactivity consider scaling down the provisioned capacity or instance size during these times. Aurora Serverless V2 will automatically scale up and down based on load.

- b. For non-production instances, consider pausing or stopping these during non-work hours.
- 6. Do you need to segment and break apart your data models based on access patterns and data characteristics?
 - a. Consider using AWS DMS or AWS SCT to move your data to other services.

Level of effort for the implementation plan:

To establish this best practice, you must be aware of your current data characteristics and metrics. Gathering those metrics, establishing a baseline and then using those metrics to identify the ideal database configuration options is a *low* to *moderate* level of effort. This is best validated by load tests and experimentation.

Resources

Related documents:

- [Cloud Databases with AWS](#)
- [AWS Database Caching](#)
- [Amazon DynamoDB Accelerator](#)
- [Amazon Aurora best practices](#)
- [Amazon Redshift performance](#)
- [Amazon Athena top 10 performance tips](#)
- [Amazon Redshift Spectrum best practices](#)
- [Amazon DynamoDB best practices](#)

Related videos:

- [AWS purpose-built databases \(DAT209-L\)](#)
- [Amazon Aurora storage demystified: How it all works \(DAT309-R\)](#)
- [Amazon DynamoDB deep dive: Advanced design patterns \(DAT403-R1\)](#)

Related examples:

- [Amazon DynamoDB Examples](#)
- [AWS Database migration samples](#)
- [Database Modernization Workshop](#)
- [Working with parameters on your Amazon RDS for PostgreSQL](#)

PERF04-BP03 Collect and record database performance metrics

To understand how your data management systems are performing, it is important to track relevant metrics. These metrics will help you to optimize your data management resources, to ensure that your workload requirements are met, and that you have a clear overview on how the workload performs. Use tools, libraries, and systems that record performance measurements related to database performance.

There are metrics that are related to the system on which the database is being hosted (for example, CPU, storage, memory, IOPS), and there are metrics for accessing the data itself (for example, transactions per second, queries rates, response times, errors). These metrics should be readily accessible

for any support or operational staff, and have sufficient historical record to be able to identify trends, anomalies, and bottlenecks.

Desired outcome: To monitor the performance of your database workloads, you must record multiple performance metrics over a period of time. This allows you to detect anomalies as well as measure performance against business metrics to ensure you are meeting your workload needs.

Common anti-patterns:

- You only use manual log file searching for metrics.
- You only publish metrics to internal tools used by your team and don't have a comprehensive picture of your workload.
- You only use the default metrics recorded by your selected monitoring software.
- You only review metrics when there is an issue.
- You only monitor system level metrics, not capturing data access or usage metrics.

Benefits of establishing this best practice: Establishing a performance baseline helps in understanding normal behavior and requirements of workloads. Abnormal patterns can be identified and debugged faster improving performance and reliability of the database. Database capacity can be configured to ensure optimal cost without compromising performance.

Level of risk exposed if this best practice is not established: High

- Inability to differentiate out of normal vs. normal performance level will create difficulties in issue identification, and decision making.
- Potential cost savings may not be identified.
- Growth patterns will not be identified which might result in reliability or performance degradation.

Implementation guidance

Identify, collect, aggregate, and correlate database-related metrics. Metrics should include both the underlying system that is supporting the database and the database metrics. The underlying system metrics might include CPU utilization, memory, available disk storage, disk I/O, and network inbound and outbound metrics while the database metrics might include transactions per second, top queries, average queries rates, response times, index usage, table locks, query timeouts, and number of connections open. This data is crucial to understand how the workload is performing and how the database solution is used. Use these metrics as part of a data-driven approach to tune and optimize your workload's resources.

Implementation steps:

1. Which database metrics are important to track?
 - a. [Monitoring metrics for Amazon RDS](#)
 - b. [Monitoring with Performance Insights](#)
 - c. [Enhanced monitoring](#)
 - d. [DynamoDB metrics](#)
 - e. [Monitoring DynamoDB DAX](#)
 - f. [Monitoring MemoryDB](#)
 - g. [Monitoring Amazon Redshift](#)
 - h. [Timeseries metrics and dimensions](#)
 - i. [Cluster level metrics for Aurora](#)
 - j. [Monitoring Amazon Keyspaces](#)

- k. [Monitoring Amazon Neptune](#)
- 2. Would the database monitoring benefit from a machine learning solution that detects operational anomalies performance issues?
 - a. [Amazon DevOps Guru for Amazon RDS](#) provides visibility into performance issues and makes recommendations for corrective actions.
- 3. Do you need application level details about SQL usage?
 - a. [AWS X-Ray](#) can be instrumented into the application to gain insights and encapsulate all the data points for single query.
- 4. Do you currently have an approved logging and monitoring solution?
 - a. [Amazon CloudWatch](#) can collect metrics across the resources in your architecture. You can also collect and publish custom metrics to surface business or derived metrics. Use CloudWatch or third-party solutions to set alarms that indicate when thresholds are breached.
- 5. You identified and configured your data retention policies to match my security and operational goals?
 - a. [Default data retention for CloudWatch metrics](#)
 - b. [Default data retention for CloudWatch Logs](#)

Level of effort for the implementation plan: There is a *medium* level of effort to identify, track, collect, aggregate, and correlate metrics from all database resources.

Resources

Related documents:

- [AWS Database Caching](#)
- [Amazon Athena top 10 performance tips](#)
- [Amazon Aurora best practices](#)
- [Amazon DynamoDB Accelerator](#)
- [Amazon DynamoDB best practices](#)
- [Amazon Redshift Spectrum best practices](#)
- [Amazon Redshift performance](#)
- [Cloud Databases with AWS](#)
- [Amazon RDS Performance Insights](#)

Related videos:

- [AWS purpose-built databases \(DAT209-L\)](#)
- [Amazon Aurora storage demystified: How it all works \(DAT309-R\)](#)
- [Amazon DynamoDB deep dive: Advanced design patterns \(DAT403-R1\)](#)

Related examples:

- [Level 100: Monitoring with CloudWatch Dashboards](#)
- [AWS Dataset Ingestion Metrics Collection Framework](#)
- [Amazon RDS Monitoring Workshop](#)

PERF04-BP04 Choose data storage based on access patterns

Use the access patterns of the workload to decide which services and technologies to use. In addition to non-functional requirements such as performance and scale, access patterns heavily influence the choice

of the database and storage solutions. The first dimension is the need for transactions, ACID compliance, and consistent reads. Not every database supports these and most of the NoSQL databases provide an eventual consistency model. The second important dimension would be the distribution of write and reads over time and space. Globally distributed applications need to consider the traffic patterns, latency and access requirements in order to identify the optimal storage solution. The third crucial aspect to choose is the query pattern flexibility, random access patterns, and one-time queries. Considerations around highly specialized query functionality for text and natural language processing, time series, and graphs must also be taken into account.

Desired outcome: The data storage has been selected based on identified and documented data access patterns. This might include the most common read, write and delete queries, the need for ad-hoc calculations and aggregations, complexity of the data, the data interdependency, and the required consistency needs.

Common anti-patterns:

- You only select one database vendor to simplify operations management.
- You assume that data access patterns will stay consistent over time.
- You implement complex transactions, rollback, and consistency logic in the application.
- The database is configured to support a potential high traffic burst, which results in the database resources remaining idle most of the time.
- Using a shared database for transactional and analytical uses.

Benefits of establishing this best practice: Selecting and optimizing your data storage based on access patterns will help decrease development complexity and optimize your performance opportunities. Understanding when to use read replicas, global tables, data partitioning, and caching will help you decrease operational overhead and scale based on your workload needs.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Identify and evaluate your data access pattern to select the correct storage configuration. Each database solution has options to configure and optimize your storage solution. Use the collected metrics and logs and experiment with options to find the optimal configuration. Use the following table to review storage options per database service.

AWS Services	Amazon RDS, Amazon Aurora	Amazon DynamoDB	Amazon DocumentDB	Amazon ElastiCache	Amazon Neptune	Amazon Timestream	Amazon Keyspaces	Amazon QLDB
Scaling Storage	Storage automatic scaling option available to automatically scale provisioned storage IOPS can also be scaled independently	Automatically scales. Tables are unconstrained in terms of size.	Storage automatic scaling option available to scale provisioned storage	Storage is in-memory, tied to instance type or count	Storage automatic scaling option available to automatically scale provisioned storage	Configure retention period for in-memory and magnetic tiers in provisioned days	Scales table storage up and down automatically	Automatically scales. Tables are unconstrained in terms of size.

AWS Services	Amazon RDS, Amazon Aurora	Amazon DynamoDB	Amazon DocumentDB	Amazon ElastiCache	Amazon Neptune	Amazon Timestream	Amazon Keyspaces	Amazon QLDB
	of provisioned storage when leveraging provisioned IOPS storage types							

Implementation steps:

1. Identify and document the anticipated growth of the data and traffic.
 - a. Amazon RDS and Aurora support storage automatic scaling up to documented limits. Beyond this, consider transitioning older data to Amazon S3 for archival, aggregating historical data for analytics or scaling horizontally via sharding.
 - b. DynamoDB and Amazon S3 will scale to near limitless storage volume automatically.
 - c. Amazon RDS instances and databases running on EC2 can be manually resized and EC2 instances can have new EBS volumes added at a later date for additional storage.
 - d. Instance types can be changed based on changes in activity. For example, you can start with a smaller instance while you are testing, then scale the instance as you begin to receive production traffic to the service. Aurora Serverless V2 automatically scales in response to changes in load.
1. Document requirements around normal and peak performance (transactions per second TPS and queries per second QPS) and consistency (ACID and eventual consistency).
2. Document solution deployment aspects and the database access requirements (global, Multi-AZ, read replication, multiple write nodes)

Level of effort for the implementation plan: If you do not have logs or metrics for your data management solution, you will need to complete that before identifying and documenting your data access patterns. Once your data access pattern is understood, selecting, and configuring your data storage is a *low* level of effort.

Resources

Related documents:

- [AWS Database Caching](#)
- [Amazon Athena top 10 performance tips](#)
- [Amazon Aurora best practices](#)
- [Amazon DynamoDB Accelerator](#)
- [Amazon DynamoDB best practices](#)
- [Amazon Redshift Spectrum best practices](#)
- [Amazon Redshift performance](#)
- [Cloud Databases with AWS](#)
- [Amazon RDS Storage Types](#)

Related videos:

- [AWS purpose-built databases \(DAT209-L\)](#)
- [Amazon Aurora storage demystified: How it all works \(DAT309-R\)](#)
- [Amazon DynamoDB deep dive: Advanced design patterns \(DAT403-R1\)](#)

Related examples:

- [Experiment and test with Distributed Load Testing on AWS](#)

PERF04-BP05 Optimize data storage based on access patterns and metrics

Use performance characteristics and access patterns that optimize how data is stored or queried to achieve the best possible performance. Measure how optimizations such as indexing, key distribution, data warehouse design, or caching strategies impact system performance or overall efficiency.

Common anti-patterns:

- You only use manual log file searching for metrics.
- You only publish metrics to internal tools.

Benefits of establishing this best practice: In order to ensure you are meeting the metrics required for the workload, you must monitor database performance metrics related to both reads and writes. You can use this data to add new optimizations for both reads and writes to the data storage layer.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Optimize data storage based on metrics and patterns: Use reported metrics to identify any underperforming areas in your workload and optimize your database components. Each database system has different performance related characteristics to evaluate, such as how data is indexed, cached, or distributed among multiple systems. Measure the impact of your optimizations.

Resources

Related documents:

- [AWS Database Caching](#)
- [Amazon Athena top 10 performance tips](#)
- [Amazon Aurora best practices](#)
- [Amazon DynamoDB Accelerator](#)
- [Amazon DynamoDB best practices](#)
- [Amazon Redshift Spectrum best practices](#)
- [Amazon Redshift performance](#)
- [Cloud Databases with AWS](#)
- [Analyzing performance anomalies with DevOps Guru for RDS](#)
- [Read/Write Capacity Mode for DynamoDB](#)

Related videos:

- [AWS purpose-built databases \(DAT209-L\)](#)

- [Amazon Aurora storage demystified: How it all works \(DAT309-R\)](#)
- [Amazon DynamoDB deep dive: Advanced design patterns \(DAT403-R1\)](#)

Related examples:

- [Hands-on Labs for Amazon DynamoDB](#)

PERF 5 How do you configure your networking solution?

The optimal network solution for a workload varies based on latency, throughput requirements, jitter, and bandwidth. Physical constraints, such as user or on-premises resources, determine location options. These constraints can be offset with edge locations or resource placement.

Best practices

- [PERF05-BP01 Understand how networking impacts performance \(p. 318\)](#)
- [PERF05-BP02 Evaluate available networking features \(p. 320\)](#)
- [PERF05-BP03 Choose appropriately sized dedicated connectivity or VPN for hybrid workloads \(p. 323\)](#)
- [PERF05-BP04 Leverage load-balancing and encryption offloading \(p. 325\)](#)
- [PERF05-BP05 Choose network protocols to improve performance \(p. 326\)](#)
- [PERF05-BP06 Choose your workload's location based on network requirements \(p. 327\)](#)
- [PERF05-BP07 Optimize network configuration based on metrics \(p. 329\)](#)

PERF05-BP01 Understand how networking impacts performance

Analyze and understand how network-related decisions impact workload performance. The network is responsible for the connectivity between application components, cloud services, edge networks and on-premises data and therefore it can highly impact workload performance. In addition to workload performance, user experience is also impacted by network latency, bandwidth, protocols, location, network congestion, jitter, throughput, and routing rules.

Desired outcome: Have a documented list of networking requirements from the workload including latency, packet size, routing rules, protocols, and supporting traffic patterns. Review the available networking solutions and identify which service meets your workload networking characteristics. Cloud-based networks can be quickly rebuilt, so evolving your network architecture over time is necessary to improve performance efficiency.

Common anti-patterns:

- All traffic flows through your existing data centers.
- You overbuild Direct Connect sessions without understanding the actual usage requirements.
- You don't consider workload characteristics and encryption overhead when defining your networking solutions.
- You use on-premises concepts and strategies for networking solutions in the cloud.

Benefits of establishing this best practice: Understanding how networking impacts workload performance will help you identify potential bottlenecks, improve user experience, increase reliability, and lower operational maintenance as the workload changes.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Identify important network performance metrics of your workload and capture its networking characteristics. Define and document requirements as part of a data-driven approach, using benchmarking or load testing. Use this data to identify where your network solution is constrained, and examine configuration options that could improve the workload. Understand the cloud-native networking features and options available and how they can impact your workload performance based on the requirements. Each networking feature has advantages and disadvantages and can be configured to meet your workload characteristics and scale based on your needs.

Implementation steps:

1. Define and document networking performance requirements:
 - a. Include metrics such as network latency, bandwidth, protocols, locations, traffic patterns (spikes and frequency), throughput, encryption, inspection, and routing rules
2. Capture your foundational networking characteristics:
 - a. [VPC Flow Logs](#)
 - b. [AWS Transit Gateway metrics](#)
 - c. [AWS PrivateLink metrics](#)
3. Capture your application networking characteristics:
 - a. [Elastic Network Adaptor](#)
 - b. [AWS App Mesh metrics](#)
 - c. [Amazon API Gateway metrics](#)
4. Capture your edge networking characteristics:
 - a. [Amazon CloudFront metrics](#)
 - b. [Amazon Route 53 metrics](#)
 - c. [AWS Global Accelerator metrics](#)
5. Capture your hybrid networking characteristics:
 - a. [Direct Connect metrics](#)
 - b. [AWS Site-to-Site VPN metrics](#)
 - c. [AWS Client VPN metrics](#)
 - d. [AWS Cloud WAN metrics](#)
6. Capture your security networking characteristics:
 - a. [AWS Shield, WAF, and Network Firewall metrics](#)
7. Capture end-to-end performance metrics with tracing tools:
 - a. [AWS X-Ray](#)
 - b. [Amazon CloudWatch RUM](#)
8. Benchmark and test network performance:
 - a. [Benchmark](#) network throughput: Some factors that can affect EC2 network performance when the instances are in the same VPC. Measure the network bandwidth between EC2 Linux instances in the same VPC.
 - b. Perform [load tests](#) to experiment with networking solutions and options

Level of effort for the implementation plan: There is a *medium* level of effort to document workload networking requirements, options, and available solutions.

Resources

Related documents:

- [Application Load Balancer](#)

- [EC2 Enhanced Networking on Linux](#)
- [EC2 Enhanced Networking on Windows](#)
- [EC2 Placement Groups](#)
- [Enabling Enhanced Networking with the Elastic Network Adapter \(ENA\) on Linux Instances](#)
- [Network Load Balancer](#)
- [Networking Products with AWS](#)
- [Transit Gateway](#)
- [Transitioning to latency-based routing in Amazon Route 53](#)
- [VPC Endpoints](#)
- [VPC Flow Logs](#)

Related videos:

- [Connectivity to AWS and hybrid AWS network architectures \(NET317-R1\)](#)
- [Optimizing Network Performance for Amazon EC2 Instances \(CMP308-R1\)](#)
- [Improve Global Network Performance for Applications](#)
- [EC2 Instances and Performance Optimization Best Practices](#)
- [Optimizing Network Performance for Amazon EC2 Instances](#)
- [Networking best practices and tips with the Well-Architected Framework](#)
- [AWS networking best practices in large-scale migrations](#)

Related examples:

- [AWS Transit Gateway and Scalable Security Solutions](#)
- [AWS Networking Workshops](#)

PERF05-BP02 Evaluate available networking features

Evaluate networking features in the cloud that may increase performance. Measure the impact of these features through testing, metrics, and analysis. For example, take advantage of network-level features that are available to reduce latency, packet loss, or jitter.

Many services are created to improve performance and others commonly offer features to optimize network performance. Services such as AWS Global Accelerator and Amazon CloudFront exist to improve performance while most other services have product features to optimize network traffic. Review service features, such as EC2 instance network capability, enhanced networking instance types, Amazon EBS-optimized instances, Amazon S3 transfer acceleration, and CloudFront, to improve your workload performance.

Desired outcome: You have documented the inventory of components within your workload and have identified which networking configurations per component will help you meet your performance requirements. After evaluating the networking features, you have experimented and measured the performance metrics to identify how to use the features available to you.

Common anti-patterns:

- You put all your workloads into an AWS Region closest to your headquarters instead of an AWS Region close to your end users.
- Failing to benchmark your workload performance and continually evaluating your workload performance against that benchmark.

- You do not review service configurations for performance improving options.

Benefits of establishing this best practice: Evaluating all service features and options can increase your workload performance, reduce the cost of infrastructure, decrease the effort required to maintain your workload, and increase your overall security posture. You can use the global AWS backbone to ensure that you provide the optimal networking experience for your customers.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Review which network-related configuration options are available to you, and how they could impact your workload. Understanding how these options interact with your architecture and the impact that they will have on both measured performance and the performance perceived by users is critical for performance optimization.

Implementation steps:

1. Create a list of workload components.
 - a. Build, manage and monitor your organizations network using [AWS Cloud WAN](#).
 - b. Get visibility into your network using [Network Manager](#). Use an existing configuration management database (CMDB) tool or a tool such as [AWS Config](#) to create an inventory of your workload and how it's configured.
2. If this is an existing workload, identify and document the benchmark for your performance metrics, focusing on the bottlenecks and areas to improve. Performance-related networking metrics will differ per workload based on business requirements and workload characteristics. As a start, these metrics might be important to review for your workload: bandwidth, latency, packet loss, jitter, and retransmits.
3. If this is a new workload, perform [load tests](#) to identify performance bottlenecks.
4. For the performance bottlenecks you identify, review the configuration options for your solutions to identify performance improvement opportunities.
5. If you don't know your network path or routes, use [Network Access Analyzer](#) to identify them.
6. Have you chosen the optimal network protocols and is there more that you can do to reduce your latency?
 - a. [PERF05-BP05 Choose network protocols to improve performance \(p. 326\)](#)
7. Are you using a Site-to-Site VPN across multiple locations to connect to an AWS Region?
 - a. Review [accelerated Site-to-Site VPN connections](#) for opportunities to improve networking performance.
8. Is your workload traffic spread across multiple accounts?
 - a. Evaluate your operational and performance tradeoffs between [VPC Peering](#) and [AWS Transit Gateway](#) when connecting multiple accounts. AWS Transit Gateway supports an AWS Site-to-Site VPN throughput to scale beyond a single [IPsec maximum limit](#) by using multi-path. Traffic between an Amazon VPC and AWS Transit Gateway remains on the private AWS network and is not exposed to the internet. AWS Transit Gateway simplifies how you interconnect all of your VPCs, which can span across thousands of AWS accounts and into on-premises networks. Share your AWS Transit Gateway between multiple accounts using [Resource Access Manager](#). To get visibility into your global network traffic, use [Network Manager](#) to get a central view of your network metrics.
9. Is your workload supporting a global user base?
 - a. [AWS Global Accelerator](#) is a networking service that improves the performance of your users' traffic by up to 60% using the Amazon Web Services global network infrastructure. When the internet is congested, AWS Global Accelerator optimizes the path to your application to keep packet loss, jitter, and latency consistently low. It also provides static IP addresses that simplify moving endpoints between Availability Zones or AWS Regions without needing to update your DNS configuration or change client-facing applications.

- b. [Amazon CloudFront](#) can improve the performance of your workload content delivery and latency globally. CloudFront has over 410 globally dispersed points of presence that can cache your content and lower the latency to the end user.
 - c. Amazon Route 53 offers [latency-based routing](#), [geolocation routing](#), [geoproximity routing](#), and [IP-based routing](#) options to help you improve your workload's performance for a global audience. Identify which routing option would optimize your workload performance by reviewing your workload traffic and user location.
- 10 Does your workload use Amazon S3 for content storage?
- a. [Amazon S3 Transfer acceleration](#) is a feature that lets external users benefit from the networking optimizations of CloudFront to upload data to Amazon S3. This improves the ability to transfer large amounts of data from remote locations that don't have dedicated connectivity to the AWS Cloud.
 - b. [Amazon S3 Multi-Region Access Points](#) replicates content to multiple Regions and simplifies the workload by providing one access point. When a Multi-Region Access Point is used, you can request or write data to Amazon S3 with the service identifying the lowest latency bucket.
- 11 Is your network bandwidth being throttled at the compute resource?
- a. Elastic Network Interfaces (ENA) used by EC2 instances, containers, and Lambda functions are limited on a per-flow basis. Review your placement groups to optimize your [EC2 networking throughput](#). To avoid the bottleneck at the per flow-basis, design your application to use multiple flows. To monitor and get visibility into your compute related networking metrics, use [CloudWatch Metrics](#) and [ethtool](#). ethtool is included in the ENA driver and exposes additional network-related metrics that can be published as a [custom metric](#) to CloudWatch.
 - b. Newer EC2 instances can leverage enhanced networking. [N-series EC2 instances](#), such as M5n and M5dn, take advantage of the fourth generation of custom Nitro cards to deliver up to 100 Gbps of network throughput to a single instance. These instances offer four times the network bandwidth and packet process compared to the base M5 instances and are ideal for network intensive applications.
 - c. [Amazon Elastic Network Adapters](#) (ENA) provide further optimization by delivering better throughput for your instances within a [cluster placement group](#).
 - d. [Elastic Fabric Adapter](#) (EFA) is a network interface for Amazon EC2 instances that enables you to run workloads requiring high levels of internode communications at scale on AWS. With EFA, High Performance Computing (HPC) applications using the Message Passing Interface (MPI) and Machine Learning (ML) applications using NVIDIA Collective Communications Library (NCCL) can scale to thousands of CPUs or GPUs.
 - e. [Amazon EBS-optimized](#) instances use an optimized configuration stack and provide additional, dedicated capacity to increase the Amazon EBS I/O. This optimization provides the best performance for your EBS volumes by minimizing contention between Amazon EBS I/O and other traffic from your instance.

Level of effort for the implementation plan:

To establish this best practice, you must be aware of your current workload component options that impact network performance. Gathering the components, evaluating network improvement options, experimenting, implementing, and documenting those improvements is a *low* to *moderate* level of effort.

Resources

Related documents:

- [Amazon EBS - Optimized Instances](#)
- [Application Load Balancer](#)
- [Amazon EC2 instance network bandwidth](#)
- [EC2 Enhanced Networking on Linux](#)

- [EC2 Enhanced Networking on Windows](#)
- [EC2 Placement Groups](#)
- [Enabling Enhanced Networking with the Elastic Network Adapter \(ENA\) on Linux Instances](#)
- [Network Load Balancer](#)
- [Networking Products with AWS](#)
- [AWS Transit Gateway](#)
- [Transitioning to Latency-Based Routing in Amazon Route 53](#)
- [VPC Endpoints](#)
- [VPC Flow Logs](#)
- [Building a cloud CMDB](#)
- [Scaling VPN throughput using AWS Transit Gateway](#)

Related videos:

- [Connectivity to AWS and hybrid AWS network architectures \(NET317-R1\)](#)
- [Optimizing Network Performance for Amazon EC2 Instances \(CMP308-R1\)](#)
- [AWS Global Accelerator](#)

Related examples:

- [AWS Transit Gateway and Scalable Security Solutions](#)
- [AWS Networking Workshops](#)

PERF05-BP03 Choose appropriately sized dedicated connectivity or VPN for hybrid workloads

When a common network is required to connect on-premises and cloud resources in AWS, ensure that you have adequate bandwidth to meet your performance requirements. Estimate the bandwidth and latency requirements for your hybrid workload. These numbers will drive the sizing requirements for AWS Direct Connect or your VPN endpoints.

Desired outcome: When deploying a workload that will need hybrid network connectivity, you have multiple configuration options for connectivity, such as managed and non-managed VPNs or Direct Connect. Select the appropriate connection type for each workload while ensuring you have adequate bandwidth and encryption requirements between your location and the cloud.

Common anti-patterns:

- You only evaluate VPN solutions for your network encryption requirements.
- You don't evaluate backup or parallel connectivity options.
- You use default configurations for routers, tunnels, and BGP sessions.
- You fail to understand or identify all workload requirements (encryption, protocol, bandwidth and traffic needs).

Benefits of establishing this best practice: Selecting and configuring appropriately sized hybrid network solutions will increase the reliability of your workload and maximize performance opportunities. By identifying workload requirements, planning ahead, and evaluating hybrid solutions you will minimize expensive physical network changes and operational overhead while increasing your time to market.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Develop a hybrid networking architecture based on your bandwidth requirements: Estimate the bandwidth and latency requirements of your hybrid applications. Based on your bandwidth requirements, a single VPN or Direct Connect connection might not be enough, and you must architect a hybrid setup to enable traffic load balancing across multiple connections. Direct connect may be required which offers more predictable and consistent performance due to its private network connectivity. It is great for production workloads that require consistent latency and almost zero jitter.

AWS Direct Connect provides dedicated connectivity to the AWS environment, from 50 Mbps up to 10 Gbps. This gives you managed and controlled latency and provisioned bandwidth so your workload can connect easily and in a performant way to other environments. Using one of the AWS Direct Connect partners, you can have end-to-end connectivity from multiple environments, thus providing an extended network with consistent performance.

The AWS Site-to-Site VPN is a managed VPN service for VPCs. When a VPN connection is created, AWS provides tunnels to two different VPN endpoints. With AWS Transit Gateway, you can simplify the connectivity between multiple VPCs and also connect to any VPC attached to AWS Transit Gateway with a single VPN connection. AWS Transit Gateway also enables you to scale beyond the 1.25Gbps IPsec VPN throughput limit by enabling equal cost multi-path (ECMP) routing support over multiple VPN tunnels.

Level of effort for the implementation plan: There is a *high* level of effort to evaluate workload needs for hybrid networks and to implement hybrid networking solutions.

Resources

Related documents:

- [Network Load Balancer](#)
- [Networking Products with AWS](#)
- [Transit Gateway](#)
- [Transitioning to latency-based Routing in Amazon Route 53](#)
- [VPC Endpoints](#)
- [VPC Flow Logs](#)
- [Site-to-Site VPN](#)
- [Building a Scalable and Secure Multi-VPC AWS Network Infrastructure](#)
- [Direct Connect](#)
- [Client VPN](#)

Related videos:

- [Connectivity to AWS and hybrid AWS network architectures \(NET317-R1\)](#)
- [Optimizing Network Performance for Amazon EC2 Instances \(CMP308-R1\)](#)
- [AWS Global Accelerator](#)
- [Direct Connect](#)
- [Transit Gateway Connect](#)
- [VPN Solutions](#)
- [Security with VPN Solutions](#)

Related examples:

- [AWS Transit Gateway and Scalable Security Solutions](#)

- [AWS Networking Workshops](#)

PERF05-BP04 Leverage load-balancing and encryption offloading

Distribute traffic across multiple resources or services to allow your workload to take advantage of the elasticity that the cloud provides. You can also use load balancing for offloading encryption termination to improve performance and to manage and route traffic effectively.

When implementing a scale-out architecture where you want to use multiple instances for service content, you can use load balancers inside your Amazon VPC. AWS provides multiple models for your applications in the ELB service. Application Load Balancer is best suited for load balancing of HTTP and HTTPS traffic and provides advanced request routing targeted at the delivery of modern application architectures, including microservices and containers.

Network Load Balancer is best suited for load balancing of TCP traffic where extreme performance is required. It is capable of handling millions of requests per second while maintaining ultra-low latencies, and it is optimized to handle sudden and volatile traffic patterns.

[Elastic Load Balancing](#) provides integrated certificate management and SSL/TLS decryption, allowing you the flexibility to centrally manage the SSL settings of the load balancer and offload CPU intensive work from your workload.

Common anti-patterns:

- You route all internet traffic through existing load balancers.
- You use generic TCP load balancing and making each compute node handle SSL encryption.

Benefits of establishing this best practice: A load balancer handles the varying load of your application traffic in a single Availability Zone, or across multiple Availability Zones. Load balancers feature the high availability, automatic scaling, and robust security necessary to make your applications fault tolerant.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Use the appropriate load balancer for your workload: Select the appropriate load balancer for your workload. If you must load balance HTTP requests, we recommend Application Load Balancer. For network and transport protocols (layer 4 – TCP, UDP) load balancing, and for extreme performance and low latency applications, we recommend Network Load Balancer. Application Load Balancers support HTTPS and Network Load Balancers support TLS encryption offloading.

Enable offload of HTTPS or TLS encryption: Elastic Load Balancing includes integrated certificate management, user-authentication, and SSL/TLS decryption. It provides the flexibility to centrally manage TLS settings and offload CPU intensive workloads from your applications. Encrypt all HTTPS traffic as part of your load balancer deployment.

Resources

Related documents:

- [Amazon EBS - Optimized Instances](#)
- [Application Load Balancer](#)
- [EC2 Enhanced Networking on Linux](#)
- [EC2 Enhanced Networking on Windows](#)
- [EC2 Placement Groups](#)
- [Enabling Enhanced Networking with the Elastic Network Adapter \(ENA\) on Linux Instances](#)

- [Network Load Balancer](#)
- [Networking Products with AWS](#)
- [Transit Gateway](#)
- [Transitioning to Latency-Based Routing in Amazon Route 53](#)
- [VPC Endpoints](#)
- [VPC Flow Logs](#)

Related videos:

- [Connectivity to AWS and hybrid AWS network architectures \(NET317-R1\)](#)
- [Optimizing Network Performance for Amazon EC2 Instances \(CMP308-R1\)](#)

Related examples:

- [AWS Transit Gateway and Scalable Security Solutions](#)
- [AWS Networking Workshops](#)

PERF05-BP05 Choose network protocols to improve performance

Make decisions about protocols for communication between systems and networks based on the impact to the workload's performance.

There is a relationship between latency and bandwidth to achieve throughput. If your file transfer is using TCP, higher latencies will reduce overall throughput. There are approaches to fix this with TCP tuning and optimized transfer protocols, some approaches use UDP.

Common anti-patterns:

- You use TCP for all workloads regardless of performance requirements.

Benefits of establishing this best practice: Selecting the proper protocol for communication between workload components ensures that you are getting the best performance for that workload. Connection-less UDP allows for high speed, but it doesn't offer retransmission or high reliability. TCP is a full featured protocol, but it requires greater overhead for processing the packets.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Optimize network traffic: Select the appropriate protocol to optimize the performance of your workload. There is a relationship between latency and bandwidth to achieve throughput. If your file transfer is using TCP, higher latencies reduce overall throughput. There are approaches to fix latency with TCP tuning and optimized transfer protocols, some which use UDP.

Resources

Related documents:

- [Amazon EBS - Optimized Instances](#)
- [Application Load Balancer](#)
- [EC2 Enhanced Networking on Linux](#)
- [EC2 Enhanced Networking on Windows](#)
- [EC2 Placement Groups](#)

- [Enabling Enhanced Networking with the Elastic Network Adapter \(ENA\) on Linux Instances](#)
- [Network Load Balancer](#)
- [Networking Products with AWS](#)
- [Transit Gateway](#)
- [Transitioning to Latency-Based Routing in Amazon Route 53](#)
- [VPC Endpoints](#)
- [VPC Flow Logs](#)

Related videos:

- [Connectivity to AWS and hybrid AWS network architectures \(NET317-R1\)](#)
- [Optimizing Network Performance for Amazon EC2 Instances \(CMP308-R1\)](#)

Related examples:

- [AWS Transit Gateway and Scalable Security Solutions](#)
- [AWS Networking Workshops](#)

PERF05-BP06 Choose your workload's location based on network requirements

Use the cloud location options available to reduce network latency or improve throughput. Use AWS Regions, Availability Zones, placement groups, and edge locations such as AWS Outposts, AWS Local Zones, and AWS Wavelength, to reduce network latency or improve throughput.

The AWS Cloud infrastructure is built around Regions and Availability Zones. A Region is a physical location in the world having multiple Availability Zones.

Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity, housed in separate facilities. These Availability Zones offer you the ability to operate production applications and databases that are more highly available, fault tolerant, and scalable than would be possible from a single data center

Choose the appropriate Region or Regions for your deployment based on the following key elements:

- **Where your users are located:** Choosing a Region close to your workload's users ensures lower latency when they use the workload.
- **Where your data is located:** For data-heavy applications, the major bottleneck in latency is data transfer. Application code should execute as close to the data as possible.
- **Other constraints:** Consider constraints such as security and compliance.

Amazon EC2 provides placement groups for networking. A placement group is a logical grouping of instances to decrease latency or increase reliability. Using placement groups with supported instance types and an Elastic Network Adapter (ENA) enables workloads to participate in a low-latency, 25 Gbps network. Placement groups are recommended for workloads that benefit from low network latency, high network throughput, or both. Using placement groups has the benefit of lowering jitter in network communications.

Latency-sensitive services are delivered at the edge using a global network of edge locations. These edge locations commonly provide services such as content delivery network (CDN) and domain name system (DNS). By having these services at the edge, workloads can respond with low latency to requests for content or DNS resolution. These services also provide geographic services such as geo targeting of content (providing different content based on the end users' location), or latency-based routing to direct end users to the nearest Region (minimum latency).

[Amazon CloudFront](#) is a global CDN that can be used to accelerate both static content such as images, scripts, and videos, as well as dynamic content such as APIs or web applications. It relies on a global network of edge locations that will cache the content and provide high-performance network connectivity to your users. CloudFront also accelerates many other features such as content uploading and dynamic applications, making it a performance addition to all applications serving traffic over the internet. [Lambda@Edge](#) is a feature of Amazon CloudFront that will let you run code closer to users of your workload, which improves performance and reduces latency.

Amazon Route 53 is a highly available and scalable cloud DNS web service. It's designed to give developers and businesses an extremely reliable and cost-effective way to route end users to internet applications by translating names, like `www.example.com`, into numeric IP addresses, like `192.168.2.1`, that computers use to connect to each other. Route 53 is fully compliant with IPv6.

[AWS Outposts](#) is designed for workloads that need to remain on-premises due to latency requirements, where you want that workload to run seamlessly with the rest of your other workloads in AWS. AWS Outposts are fully managed and configurable compute and storage racks built with AWS-designed hardware that allow you to run compute and storage on-premises, while seamlessly connecting to the broad array of AWS services in the cloud.

[AWS Local Zones](#) is designed to run workloads that require single-digit millisecond latency, like video rendering and graphics intensive, virtual desktop applications. Local Zones allow you to gain all the benefits of having compute and storage resources closer to end-users.

[AWS Wavelength](#) is designed to deliver ultra-low latency applications to 5G devices by extending AWS infrastructure, services, APIs, and tools to 5G networks. Wavelength embeds storage and compute inside telco providers 5G networks to help your 5G workload if it requires single-digit millisecond latency, such as IoT devices, game streaming, autonomous vehicles, and live media production.

Use edge services to reduce latency and to enable content caching. Ensure that you have configured cache control correctly for both DNS and HTTP/HTTPS to gain the most benefit from these approaches.

Common anti-patterns:

- You consolidate all workload resources into one geographic location.
- You chose the closest region to your location but not to the workload end user.

Benefits of establishing this best practice: You must ensure that your network is available wherever you want to reach customers. Using the AWS private global network ensures that your customers get the lowest latency experience by deploying workloads into the locations nearest them.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Reduce latency by selecting the correct locations: Identify where your users and data are located. Take advantage of AWS Regions, Availability Zones, placement groups, and edge locations to reduce latency.

Resources

Related documents:

- [Amazon EBS - Optimized Instances](#)
- [Application Load Balancer](#)
- [EC2 Enhanced Networking on Linux](#)
- [EC2 Enhanced Networking on Windows](#)
- [EC2 Placement Groups](#)

- [Enabling Enhanced Networking with the Elastic Network Adapter \(ENA\) on Linux Instances](#)
- [Network Load Balancer](#)
- [Networking Products with AWS](#)
- [Transit Gateway](#)
- [Transitioning to Latency-Based Routing in Amazon Route 53](#)
- [VPC Endpoints](#)
- [VPC Flow Logs](#)

Related videos:

- [Connectivity to AWS and hybrid AWS network architectures \(NET317-R1\)](#)
- [Optimizing Network Performance for Amazon EC2 Instances \(CMP308-R1\)](#)

Related examples:

- [AWS Transit Gateway and Scalable Security Solutions](#)
- [AWS Networking Workshops](#)

PERF05-BP07 Optimize network configuration based on metrics

Use collected and analyzed data to make informed decisions about optimizing your network configuration. Measure the impact of those changes and use the impact measurements to make future decisions.

Enable VPC Flow Logs for all VPC networks that are used by your workload. VPC Flow Logs are a feature that allows you to capture information about the IP traffic going to and from network interfaces in your VPC. VPC Flow Logs help you with a number of tasks, such as troubleshooting why specific traffic is not reaching an instance, which in turn helps you diagnose overly restrictive security group rules. You can use flow logs as a security tool to monitor the traffic that is reaching your instance, to profile your network traffic, and to look for abnormal traffic behaviors.

Use networking metrics to make changes to networking configuration as the workload evolves. Cloud based networks can be quickly rebuilt, so evolving your network architecture over time is necessary to maintain performance efficiency.

Common anti-patterns:

- You assume that all performance-related issues are application-related.
- You only test your network performance from a location close to where you have deployed the workload.

Benefits of establishing this best practice: To ensure that you are meeting the metrics required for the workload, you must monitor network performance metrics. You can capture information about the IP traffic going to and from network interfaces in your VPC and use this data to add new optimizations or deploy your workload to new geographic Regions.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Enable VPC Flow Logs: VPC Flow Logs enable you to capture information about the IP traffic going to and from network interfaces in your VPC. VPC Flow Logs help you with a number of tasks, such as troubleshooting why specific traffic is not reaching an instance, which can help you diagnose overly

restrictive security group rules. You can use flow logs as a security tool to monitor the traffic that is reaching your instance, to profile your network traffic, and to look for abnormal traffic behaviors.

Enable appropriate metrics for network options: Ensure that you select the appropriate network metrics for your workload. You can enable metrics for VPC NAT gateway, transit gateways, and VPN tunnels.

Resources

Related documents:

- [Amazon EBS - Optimized Instances](#)
- [Application Load Balancer](#)
- [EC2 Enhanced Networking on Linux](#)
- [EC2 Enhanced Networking on Windows](#)
- [EC2 Placement Groups](#)
- [Enabling Enhanced Networking with the Elastic Network Adapter \(ENA\) on Linux Instances](#)
- [Network Load Balancer](#)
- [Networking Products with AWS](#)
- [Transit Gateway](#)
- [Transitioning to Latency-Based Routing in Amazon Route 53](#)
- [VPC Endpoints](#)
- [VPC Flow Logs](#)
- [Monitoring your global and core networks with Amazon Cloudwatch metrics](#)
- [Continuously monitor network traffic and resources](#)

Related videos:

- [Connectivity to AWS and hybrid AWS network architectures \(NET317-R1\)](#)
- [Optimizing Network Performance for Amazon EC2 Instances \(CMP308-R1\)](#)
- [Monitoring and troubleshooting network traffic](#)
- [Simplify Traffic Monitoring and Visibility with Amazon VPC Traffic Mirroring](#)

Related examples:

- [AWS Transit Gateway and Scalable Security Solutions](#)
- [AWS Networking Workshops](#)
- [AWS Network Monitoring](#)

Review

Question

- [PERF 6 How do you evolve your workload to take advantage of new releases? \(p. 330\)](#)

PERF 6 How do you evolve your workload to take advantage of new releases?

When architecting workloads, there are finite options that you can choose from. However, over time, new technologies and approaches become available that could improve the performance of your workload.

Best practices

- [PERF06-BP01 Stay up-to-date on new resources and services \(p. 331\)](#)
- [PERF06-BP02 Define a process to improve workload performance \(p. 332\)](#)
- [PERF06-BP03 Evolve workload performance over time \(p. 332\)](#)

PERF06-BP01 Stay up-to-date on new resources and services

Evaluate ways to improve performance as new services, design patterns, and product offerings become available. Determine which of these could improve performance or increase the efficiency of the workload through ad-hoc evaluation, internal discussion, or external analysis.

Define a process to evaluate updates, new features, and services from AWS. For example, building proof-of-concepts that use new technologies or consulting with an internal group. When trying new ideas or services, run performance tests to measure the impact that they have on the efficiency or performance of the workload. Take advantage of the flexibility that you have in AWS to test new ideas or technologies frequently with minimal cost or risk.

Common anti-patterns:

- You only review new options and services when your workload is not meeting performance requirements.
- You assume all new product offerings will not be useful to your workload.

Benefits of establishing this best practice: By considering new services or product offerings, you may be able to greatly reduce the cost of infrastructure and the effort required to maintain your services as well as accelerate your time to market on deploying new services and features.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Define a process to evaluate updates, new features, and services from AWS. For example, building proof-of-concepts that use new technologies. When trying new ideas or services, run performance tests to measure the impact on the efficiency or performance of the workload. Take advantage of the flexibility that you have in AWS to test new ideas or technologies frequently with minimal cost or risk.

Resources

Related documents:

- [AWS Blog](#)
- [What's New with AWS](#)

Related videos:

- [AWS Events YouTube Channel](#)
- [AWS Online Tech Talks YouTube Channel](#)
- [Amazon Web Services YouTube Channel](#)

Related examples:

- [AWS Github](#)
- [AWS Skill Builder](#)

PERF06-BP02 Define a process to improve workload performance

Define a process to evaluate new services, design patterns, resource types, and configurations as they become available. For example, run existing performance tests on new instance offerings to determine their potential to improve your workload.

Your workload's performance has a few key constraints. Document these so that you know what kinds of innovation might improve the performance of your workload. Use this information when learning about new services or technology as it becomes available to identify ways to alleviate constraints or bottlenecks.

Common anti-patterns:

- You assume your current architecture will become static and never update over time.
- You introduce architecture changes over time with no metric justification.

Benefits of establishing this best practice: By defining your process for making architectural changes, you enable gathered data to influence your workload design over time.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Identify the key performance constraints for your workload: Document your workload's performance constraints so that you know what kinds of innovation might improve the performance of your workload.

Resources

Related documents:

- [AWS Blog](#)
- [What's New with AWS](#)

Related videos:

- [AWS Events YouTube Channel](#)
- [AWS Online Tech Talks YouTube Channel](#)
- [Amazon Web Services YouTube Channel](#)

Related examples:

- [AWS Github](#)
- [AWS Skill Builder](#)

PERF06-BP03 Evolve workload performance over time

As an organization, use the information gathered through the evaluation process to actively drive adoption of new services or resources when they become available.

Use the information you gather when evaluating new services or technologies to drive change. As your business or workload changes, performance needs also change. Use data gathered from your workload metrics to evaluate areas where you can get the biggest gains in efficiency or performance, and proactively adopt new services and technologies to keep up with demand.

Common anti-patterns:

- You assume that your current architecture will become static and never update over time.
- You introduce architecture changes over time with no metric justification.
- You change architecture just because everyone else in the industry is using it.

Benefits of establishing this best practice: To optimize your workload performance and cost, you must evaluate all software and services available to determine the appropriate ones for your workload.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Evolve your workload over time: Use the information you gather when evaluating new services or technologies to drive change. As your business or workload changes, performance needs also change. Use data gathered from your workload metrics to evaluate areas where you can achieve the biggest gains in efficiency or performance, and proactively adopt new services and technologies to keep up with demand.

Resources

Related documents:

- [AWS Blog](#)
- [What's New with AWS](#)

Related videos:

- [AWS Events YouTube Channel](#)
- [AWS Online Tech Talks YouTube Channel](#)
- [Amazon Web Services YouTube Channel](#)

Related examples:

- [AWS Github](#)
- [AWS Skill Builder](#)

Monitoring

Question

- [PERF 7 How do you monitor your resources to ensure they are performing? \(p. 333\)](#)

PERF 7 How do you monitor your resources to ensure they are performing?

System performance can degrade over time. Monitor system performance to identify degradation and remediate internal or external factors, such as the operating system or application load.

Best practices

- [PERF07-BP01 Record performance-related metrics \(p. 334\)](#)
- [PERF07-BP02 Analyze metrics when events or incidents occur \(p. 335\)](#)
- [PERF07-BP03 Establish key performance indicators \(KPIs\) to measure workload performance \(p. 336\)](#)

- [PERF07-BP04 Use monitoring to generate alarm-based notifications \(p. 337\)](#)
- [PERF07-BP05 Review metrics at regular intervals \(p. 338\)](#)
- [PERF07-BP06 Monitor and alarm proactively \(p. 339\)](#)

PERF07-BP01 Record performance-related metrics

Use a monitoring and observability service to record performance-related metrics. Examples of metrics include record database transactions, slow queries, I/O latency, HTTP request throughput, service latency, or other key data.

Identify the performance metrics that matter for your workload and record them. This data is an important part of being able to identify which components are impacting overall performance or efficiency of the workload.

Working back from the customer experience, identify metrics that matter. For each metric, identify the target, measurement approach, and priority. Use these to build alarms and notifications to proactively address performance-related issues.

Common anti-patterns:

- You only monitor operating system level metrics to gain insight into your workload.
- You architect your compute needs for peak workload requirements.

Benefits of establishing this best practice: To optimize performance and resource utilization, you need a unified operational view of your key performance indicators. You can create dashboards and perform metric math on your data to derive operational and utilization insights.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Identify the relevant performance metrics for your workload and record them. This data helps identify which components are impacting overall performance or efficiency of your workload.

Identify performance metrics: Use the customer experience to identify the most important metrics. For each metric, identify the target, measurement approach, and priority. Use these data points to build alarms and notifications to proactively address performance-related issues.

Resources

Related documents:

- [CloudWatch Documentation](#)
- [Collect metrics and logs from Amazon EC2 Instances and on-premises servers with the CloudWatch Agent](#)
- [Publish custom metrics](#)
- [Monitoring, Logging, and Performance APN Partners](#)
- [X-Ray Documentation](#)
- [Amazon CloudWatch RUM](#)

Related videos:

- [Cut through the chaos: Gain operational visibility and insight \(MGT301-R1\)](#)
- [Application Performance Management on AWS](#)

- [Build a Monitoring Plan](#)

Related examples:

- [Level 100: Monitoring with CloudWatch Dashboards](#)
- [Level 100: Monitoring Windows EC2 instance with CloudWatch Dashboards](#)
- [Level 100: Monitoring an Amazon Linux EC2 instance with CloudWatch Dashboards](#)

PERF07-BP02 Analyze metrics when events or incidents occur

In response to (or during) an event or incident, use monitoring dashboards or reports to understand and diagnose the impact. These views provide insight into which portions of the workload are not performing as expected.

When you write critical user stories for your architecture, include performance requirements, such as specifying how quickly each critical story should execute. For these critical stories, implement additional scripted user journeys to ensure that you know how these stories perform against your requirement.

Common anti-patterns:

- You assume that performance events are one-time issues and only related to anomalies.
- You only evaluate existing performance metrics when responding to performance events.

Benefits of establishing this best practice: In determine whether your workload is operating at expected levels, you must respond to performance events by gathering additional metric data for analysis. This data is used to understand the impact of the performance event and suggest changes to improve workload performance.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Prioritize experience concerns for critical user stories: When you write critical user stories for your architecture, include performance requirements, such as specifying how quickly each critical story should run. For these critical stories, implement additional scripted user journeys to ensure that you know how the user stories perform against your requirements.

Resources

Related documents:

- [CloudWatch Documentation](#)
- [Amazon CloudWatch Synthetics](#)
- [Monitoring, Logging, and Performance APN Partners](#)
- [X-Ray Documentation](#)

Related videos:

- [Cut through the chaos: Gain operational visibility and insight \(MGT301-R1\)](#)
- [Optimize applications through Amazon CloudWatch RUM](#)
- [Demo of Amazon CloudWatch Synthetics](#)

Related examples:

- [Measure page load time with Amazon CloudWatch Synthetics](#)
- [Amazon CloudWatch RUM Web Client](#)

PERF07-BP03 Establish key performance indicators (KPIs) to measure workload performance

Identify the KPIs that quantitatively and qualitatively measures workload performance. KPIs help to measure the health of a workload as it relates to a business goal. KPIs allow business and engineering teams to align on the measurement of goals and strategies and how this combines to produce business outcomes. KPIs should be revisited when business goals, strategies, or end-user requirements change.

For example, a website workload might use the page load time as an indication of overall performance. This metric would be one of the multiple data points which measure an end user experience. In addition to identifying the page load time thresholds, you should document the expected outcome or business risk if the performance is not met. A long page load time would affect your end users directly, decrease their user experience rating and might lead to a loss of customers. When you define your KPI thresholds, combine both industry benchmarks and your end user expectations. For example, if the current industry benchmark is a webpage loading within a two second time period, but your end users expect a webpage to load within a one second time period, then you should take both of these data points into consideration when establishing the KPI. Another example of a KPI might focus on meeting internal performance needs. A KPI threshold might be established on generating sales reports within one business day after production data has been generated. These reports might directly affect daily decisions and business outcomes.

Desired outcome: Establishing KPIs involve different departments and stakeholders. Your team must evaluate your workload KPIs using real-time granular data and historical data for reference and create dashboards that perform metric math on your KPI data to derive operational and utilization insights. KPIs should be documented which explains the agreed upon KPIs and thresholds that support business goals and strategies as well as mapped to metrics being monitored. The KPIs are identifying performance requirements, reviewed intentionally and are frequently shared and understood with all teams. Risks and tradeoffs are clearly identified and understood how business is impact within KPI thresholds are not met.

Common anti-patterns:

- You only monitor system level metrics to gain insight into your workload and don't understand business impacts to those metrics.
- You assume that your KPIs are already being published and shared as standard metric data.
- Defining KPIs but not sharing them with all the teams.
- Not defining a quantitative, measurable KPI.
- Not aligning KPIs with business goals or strategies.

Benefits of establishing this best practice: Identifying specific metrics which represent workload health help to align teams on their priorities and defining successful business outcomes. Sharing those metrics with all departments provides visibility and alignment on thresholds, expectations, and business impact.

Level of risk exposed if this best practice is not established: High

Implementation guidance

All departments and business teams impacted by the health of the workload should contribute to defining KPIs. A single person should drive the collaboration, timelines, documentation, and information related to an organization's KPIs. This single threaded owner will often share the business goals and strategies and assign business stakeholders tasks to create KPIs in their respective departments. Once

KPIs are defined, the operations team will often help define the metrics that will support and inform the success of the different KPIs. KPIs are only effective if all team members supporting a workload are aware of the KPIs.

Implementation steps

1. Identify and document business stakeholders.
2. Identify company goals and strategies.
3. Review common industry KPIs that align with your company goals and strategies.
4. Review end user expectations of your workload.
5. Define and document KPIs that support company goals and strategies.
6. Identify and document approved tradeoff strategies to meet the KPIs.
7. Identify and document metrics that will inform the KPIs.
8. Identify and document KPI thresholds for severity or alarm level.
9. Identify and document the risk and impact if the KPI is not met.
10. Identify the frequency of review per KPI.
11. Communicate KPI documentation with all teams supporting the workload.

Level of effort for the implementation guidance: Defining and communicating the KPIs is a *low* amount of work. This can typically be done over a few weeks meeting with business stakeholders, reviewing goals, strategies, and workload metrics.

Resources

Related documents:

- [CloudWatch documentation](#)
- [Monitoring, Logging, and Performance APN Partners](#)
- [X-Ray Documentation](#)
- [Using Amazon CloudWatch dashboards](#)
- [Amazon QuickSight KPIs](#)

Related videos:

- [AWS re:Invent 2019: Scaling up to your first 10 million users \(ARC211-R\)](#)
- [Cut through the chaos: Gain operational visibility and insight \(MGT301-R1\)](#)
- [Build a Monitoring Plan](#)

Related examples:

- [Creating a dashboard with Amazon QuickSight](#)

PERF07-BP04 Use monitoring to generate alarm-based notifications

Using the performance-related key performance indicators (KPIs) that you defined, use a monitoring system that generates alarms automatically when these measurements are outside expected boundaries.

Amazon CloudWatch can collect metrics across the resources in your architecture. You can also collect and publish custom metrics to surface business or derived metrics. Use CloudWatch or a third-party

monitoring service to set alarms that indicate when thresholds are breached — alarms signal that a metric is outside of the expected boundaries.

Common anti-patterns:

- You rely on staff to watch metrics and react when they see an issue.
- You rely solely on operational runbooks, when serverless workflows could be triggered to accomplish the same task.

Benefits of establishing this best practice: You can set alarms and automate actions based on either predefined thresholds, or on machine learning algorithms that identify anomalous behavior in your metrics. These same alarms can also trigger serverless workflows, which can modify performance characteristics of your workload (for example, increasing compute capacity, altering database configuration).

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Monitor metrics: Amazon CloudWatch can collect metrics across the resources in your architecture. You can collect and publish custom metrics to surface business or derived metrics. Use CloudWatch or a third-party monitoring service to set alarms that indicate when thresholds are exceeded.

Resources

Related documents:

- [CloudWatch Documentation](#)
- [Monitoring, Logging, and Performance APN Partners](#)
- [X-Ray Documentation](#)
- [Using Alarms and Alarm Actions in CloudWatch](#)

Related videos:

- [AWS re:Invent 2019: Scaling up to your first 10 million users \(ARC211-R\)](#)
- [Cut through the chaos: Gain operational visibility and insight \(MGT301-R1\)](#)
- [Build a Monitoring Plan](#)
- [Using AWS Lambda with Amazon CloudWatch Events](#)

Related examples:

- [Cloudwatch Logs Customize Alarms](#)

PERF07-BP05 Review metrics at regular intervals

As routine maintenance, or in response to events or incidents, review which metrics are collected. Use these reviews to identify which metrics were essential in addressing issues and which additional metrics, if they were being tracked, would help to identify, address, or prevent issues.

As part of responding to incidents or events, evaluate which metrics were helpful in addressing the issue and which metrics could have helped that are not currently being tracked. Use this to improve the quality of metrics you collect so that you can prevent or more quickly resolve future incidents.

Common anti-patterns:

- You allow metrics to stay in an alarm state for an extended period of time.
- You create alarms that are not actionable by an automation system.

Benefits of establishing this best practice: Continually review metrics that are being collected to ensure that they properly identify, address, or prevent issues. Metrics can also become stale if you let them stay in an alarm state for an extended period of time.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Constantly improve metric collection and monitoring: As part of responding to incidents or events, evaluate which metrics were helpful in addressing the issue and which metrics could have helped that are not currently being tracked. Use this method to improve the quality of metrics you collect so that you can prevent or more quickly resolve future incidents.

Resources

Related documents:

- [CloudWatch Documentation](#)
- [Collect metrics and logs from Amazon EC2 Instances and on-premises servers with the CloudWatch Agent](#)
- [Monitoring, Logging, and Performance APN Partners](#)
- [X-Ray Documentation](#)

Related videos:

- [Cut through the chaos: Gain operational visibility and insight \(MGT301-R1\)](#)
- [Application Performance Management on AWS](#)
- [Build a Monitoring Plan](#)

Related examples:

- [Creating a dashboard with Amazon QuickSight](#)
- [Level 100: Monitoring with CloudWatch Dashboards](#)

PERF07-BP06 Monitor and alarm proactively

Use key performance indicators (KPIs), combined with monitoring and alerting systems, to proactively address performance-related issues. Use alarms to trigger automated actions to remediate issues where possible. Escalate the alarm to those able to respond if automated response is not possible. For example, you may have a system that can predict expected key performance indicators (KPI) values and alarm when they breach certain thresholds, or a tool that can automatically halt or roll back deployments if KPIs are outside of expected values.

Implement processes that provide visibility into performance as your workload is running. Build monitoring dashboards and establish baseline norms for performance expectations to determine if the workload is performing optimally.

Common anti-patterns:

- You only allow operations staff the ability to make operational changes to the workload.

- You let all alarms filter to the operations team with no proactive remediation.

Benefits of establishing this best practice: Proactive remediation of alarm actions allows support staff to concentrate on those items that are not automatically actionable. This ensures that operations staff are not overwhelmed by all alarms and instead focus only on critical alarms.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Monitor performance during operations: Implement processes that provide visibility into performance as your workload is running. Build monitoring dashboards and establish a baseline for performance expectations.

Resources

Related documents:

- [CloudWatch Documentation](#)
- [Monitoring, Logging, and Performance APN Partners](#)
- [X-Ray Documentation](#)
- [Using Alarms and Alarm Actions in CloudWatch](#)

Related videos:

- [Cut through the chaos: Gain operational visibility and insight \(MGT301-R1\)](#)
- [Application Performance Management on AWS](#)
- [Build a Monitoring Plan](#)
- [Using AWS Lambda with Amazon CloudWatch Events](#)

Related examples:

- [Cloudwatch Logs Customize Alarms](#)

Tradeoffs

Question

- [PERF 8 How do you use tradeoffs to improve performance? \(p. 340\)](#)

PERF 8 How do you use tradeoffs to improve performance?

When architecting solutions, determining tradeoffs enables you to select an optimal approach. Often you can improve performance by trading consistency, durability, and space for time and latency.

Best practices

- [PERF08-BP01 Understand the areas where performance is most critical \(p. 341\)](#)
- [PERF08-BP02 Learn about design patterns and services \(p. 342\)](#)
- [PERF08-BP03 Identify how tradeoffs impact customers and efficiency \(p. 343\)](#)
- [PERF08-BP04 Measure the impact of performance improvements \(p. 344\)](#)

- [PERF08-BP05 Use various performance-related strategies \(p. 345\)](#)

PERF08-BP01 Understand the areas where performance is most critical

Understand and identify areas where increasing the performance of your workload will have a positive impact on efficiency or customer experience. For example, a website that has a large amount of customer interaction can benefit from using edge services to move content delivery closer to customers.

Common anti-patterns:

- You assume that standard compute metrics such as CPU utilization or memory pressure are enough to catch performance issues.
- You only use the default metrics recorded by your selected monitoring software.
- You only review metrics when there is an issue.

Benefits of establishing this best practice: Understand the various areas of your workload that could benefit from technical changes to increase performance efficiency. When you identify those areas, look at which cloud service you could deploy to remove those performance concerns.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Identify constrained areas of the workload: Use load testing or monitoring to identify constrained areas (memory, CPU, custom metrics and other key performance indicators).

Resources

Related documents:

- [Amazon Builders' Library](#)
- [X-Ray Documentation](#)
- [Amazon CloudWatch RUM](#)
- [Amazon DevOps Guru](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [Demo of Amazon CloudWatch Synthetics](#)

Related examples:

- [Measure page load time with Amazon CloudWatch Synthetics](#)
- [Amazon CloudWatch RUM Web Client](#)
- [X-Ray SDK for Node.js](#)
- [X-Ray SDK for Python](#)
- [X-Ray SDK for Java](#)
- [X-Ray SDK for .Net](#)
- [X-Ray SDK for Ruby](#)
- [X-Ray Daemon](#)

- [Distributed Load Testing on AWS](#)

PERF08-BP02 Learn about design patterns and services

Research and understand the various design patterns and services that help improve workload performance. As part of the analysis, identify what you could trade to achieve higher performance. For example, using a cache service can help to reduce the load placed on database systems. However, caching can introduce eventual consistency and requires engineering effort to implement within business requirements and customer expectations.

Desired outcome: Researching design patterns will lead you to choosing an architecture design that will support the best performing system. Learn which performance configuration options are available to you and how they could impact the workload. Optimizing the performance of your workload depends on understanding how these options interact with your architecture and the impact they will have on both measured performance and the performance perceived by end users.

Common anti-patterns:

- You assume that all traditional IT workload performance strategies are best suited for cloud workloads.
- You build and manage caching solutions instead of using managed services.
- You use the same design pattern for all your workloads without evaluating which pattern would improve the workload performance.

Benefits of establishing this best practice: By selecting the right design pattern and services for your workload you will be optimizing your performance, improving operational excellence and increasing reliability. The right design pattern will meet your current workload characteristics and help you scale for future growth or changes.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Learn which performance configuration options are available and how they could impact the workload. Optimizing the performance of your workload depends on understanding how these options interact with your architecture, and the impact they have on measured performance and user-perceived performance.

Implementation steps:

1. Evaluate and review design patterns that would improve your workload performance.
 - a. The [Amazon Builders' Library](#) provides you with a detailed description of how Amazon builds and operates technology. These articles are written by senior engineers at Amazon and cover topics across architecture, software delivery, and operations.
 - b. [AWS Solutions Library](#) is a collection of ready-to-deploy solutions that assemble services, code, and configurations. These solutions have been created by AWS and AWS Partners based on common use cases and design patterns grouped by industry or workload type. For example, you can set up a [distributed load testing solution](#) for your workload.
 - c. [AWS Architecture Center](#) provides reference architecture diagrams grouped by design pattern, content type, and technology.
 - d. [AWS samples](#) is a GitHub repository full of hands-on examples to help you explore common architecture patterns, solutions, and services. It is updated frequently with the newest services and examples.
2. Improve your workload to model the selected design patterns and use services and the service configuration options to improve your workload performance.

- a. Train your internal team with resources available at [AWS Skills Guild](#).
- b. Use the [AWS Partner Network](#) to provide expertise quickly and to scale your ability to make improvements.

Level of effort for the implementation plan: To establish this best practice, you must be aware of the design patterns and services that could help improve your workload performance. After evaluating the design patterns, implementing the design patterns is a *high* level of effort.

Resources

Related documents:

- [AWS Architecture Center](#)
- [AWS Partner Network](#)
- [AWS Solutions Library](#)
- [AWS Knowledge Center](#)
- [Amazon Builders' Library](#)
- [Using load shedding to avoid overload](#)
- [Caching challenges and strategies](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [This is My Architecture](#)

Related examples:

- [AWS Samples](#)
- [AWS SDK Examples](#)

PERF08-BP03 Identify how tradeoffs impact customers and efficiency

When evaluating performance-related improvements, determine which choices will impact your customers and workload efficiency. For example, if using a key-value data store increases system performance, it is important to evaluate how the eventually consistent nature of it will impact customers.

Identify areas of poor performance in your system through metrics and monitoring. Determine how you can make improvements, what trade-offs those improvements bring, and how they impact the system and the user experience. For example, implementing caching data can help dramatically improve performance but requires a clear strategy for how and when to update or invalidate cached data to prevent incorrect system behavior.

Common anti-patterns:

- You assume that all performance gains should be implemented, even if there are tradeoffs for implementation such as eventual consistency.
- You only evaluate changes to workloads when a performance issue has reached a critical point.

Benefits of establishing this best practice: When you are evaluating potential performance-related improvements, you must decide if the tradeoffs for the changes are consistent with the workload requirements. In some cases, you may have to implement additional controls to compensate for the tradeoffs.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Identify tradeoffs: Use metrics and monitoring to identify areas of poor performance in your system. Determine how to make improvements, and how tradeoffs will impact the system and the user experience. For example, implementing caching data can help dramatically improve performance, but it requires a clear strategy for how and when to update or invalidate cached data to prevent incorrect system behavior.

Resources

Related documents:

- [Amazon Builders' Library](#)
- [Amazon QuickSight KPIs](#)
- [Amazon CloudWatch RUM](#)
- [X-Ray Documentation](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [Build a Monitoring Plan](#)
- [Optimize applications through Amazon CloudWatch RUM](#)
- [Demo of Amazon CloudWatch Synthetics](#)

Related examples:

- [Measure page load time with Amazon CloudWatch Synthetics](#)
- [Amazon CloudWatch RUM Web Client](#)

PERF08-BP04 Measure the impact of performance improvements

As changes are made to improve performance, evaluate the collected metrics and data. Use this information to determine impact that the performance improvement had on the workload, the workload's components, and your customers. This measurement helps you understand the improvements that result from the tradeoff, and helps you determine if any negative side-effects were introduced.

A well-architected system uses a combination of performance related strategies. Determine which strategy will have the largest positive impact on a given hotspot or bottleneck. For example, sharding data across multiple relational database systems could improve overall throughput while retaining support for transactions and, within each shard, caching can help to reduce the load.

Common anti-patterns:

- You deploy and manage technologies manually that are available as managed services.
- You focus on just one component, such as networking, when multiple components could be used to increase performance of the workload.
- You rely on customer feedback and perceptions as your only benchmark.

Benefits of establishing this best practice: For implementing performance strategies, you must select multiple services and features that, taken together, will allow you to meet your workload requirements for performance.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

A well-architected system uses a combination of performance-related strategies. Determine which strategy will have the largest positive impact on a given hotspot or bottleneck. For example, sharding data across multiple relational database systems could improve overall throughput while retaining support for transactions and, within each shard, caching can help to reduce the load.

Resources

Related documents:

- [Amazon Builders' Library](#)
- [Amazon CloudWatch RUM](#)
- [Amazon CloudWatch Synthetics](#)
- [Distributed Load Testing on AWS](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [Optimize applications through Amazon CloudWatch RUM](#)
- [Demo of Amazon CloudWatch Synthetics](#)

Related examples:

- [Measure page load time with Amazon CloudWatch Synthetics](#)
- [Amazon CloudWatch RUM Web Client](#)
- [Distributed Load Testing on AWS](#)

PERF08-BP05 Use various performance-related strategies

Where applicable, use multiple strategies to improve performance. For example, using strategies like caching data to prevent excessive network or database calls, using read-replicas for database engines to improve read rates, sharding or compressing data where possible to reduce data volumes, and buffering and streaming of results as they are available to avoid blocking.

As you make changes to the workload, collect and evaluate metrics to determine the impact of those changes. Measure the impacts to the system and to the end-user to understand how your trade-offs impact your workload. Use a systematic approach, such as load testing, to explore whether the tradeoff improves performance.

Common anti-patterns:

- You assume that workload performance is adequate if customers are not complaining.
- You only collect data on performance after you have made performance-related changes.

Benefits of establishing this best practice: To optimize performance and resource utilization, you need a unified operational view, real-time granular data, and historical reference. You can create dashboards and perform metric math on your data to derive operational and utilization insights for your workloads as they change over time.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Use a data-driven approach to evolve your architecture: As you make changes to the workload, collect and evaluate metrics to determine the impact of those changes. Measure the impacts to the system and to the end-user to understand how your tradeoffs impact your workload. Use a systematic approach, such as load testing, to explore whether the tradeoff improves performance.

Resources

Related documents:

- [Amazon Builders' Library](#)
- [Best Practices for Implementing Amazon ElastiCache](#)
- [AWS Database Caching](#)
- [Amazon CloudWatch RUM](#)
- [Distributed Load Testing on AWS](#)

Related videos:

- [Introducing The Amazon Builders' Library \(DOP328\)](#)
- [AWS purpose-built databases \(DAT209-L\)](#)
- [Optimize applications through Amazon CloudWatch RUM](#)

Related examples:

- [Measure page load time with Amazon CloudWatch Synthetics](#)
- [Amazon CloudWatch RUM Web Client](#)
- [Distributed Load Testing on AWS](#)

Cost optimization

Topics

- [Practice Cloud Financial Management \(p. 346\)](#)
- [Expenditure and usage awareness \(p. 358\)](#)
- [Cost-effective resources \(p. 373\)](#)
- [Manage demand and supply resources \(p. 387\)](#)
- [Optimize over time \(p. 390\)](#)

Practice Cloud Financial Management

Question

- [COST 1 How do you implement cloud financial management? \(p. 346\)](#)

COST 1 How do you implement cloud financial management?

Implementing Cloud Financial Management enables organizations to realize business value and financial success as they optimize their cost and usage and scale on AWS.

Best practices

- [COST01-BP01 Establish a cost optimization function \(p. 347\)](#)
- [COST01-BP02 Establish a partnership between finance and technology \(p. 348\)](#)
- [COST01-BP03 Establish cloud budgets and forecasts \(p. 352\)](#)
- [COST01-BP04 Implement cost awareness in your organizational processes \(p. 353\)](#)
- [COST01-BP05 Report and notify on cost optimization \(p. 354\)](#)
- [COST01-BP06 Monitor cost proactively \(p. 356\)](#)
- [COST01-BP07 Keep up-to-date with new service releases \(p. 357\)](#)

COST01-BP01 Establish a cost optimization function

Create a team (Cloud Business Office or Cloud Center of Excellence) that is responsible for establishing and maintaining cost awareness across your organization. The team requires people from finance, technology, and business roles across the organization.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Establish a Cloud Business Office (CBO) or Cloud Center of Excellence (CCOE) team that is responsible for establishing and maintaining a culture of cost awareness in cloud computing. It can be an existing individual, a team within your organization, or a new team of key finance, technology and organization stakeholders from across the organization.

The function (individual or team) prioritizes and spends the required percentage of their time on cost management and cost optimization activities. For a small organization, the function might spend a smaller percentage of time compared to a full-time function for a larger enterprise.

The function requires a multi-disciplined approach, with capabilities in project management, data science, financial analysis, and software or infrastructure development. The function can improve efficiencies of workloads by executing cost optimizations within three different ownerships:

- **Centralized:** Through designated teams such as finance operations, cost optimization, CBO, or CCOE, customers can design and implement governance mechanisms and drive best practices company-wide.
- **Decentralized:** Influencing technology teams to execute optimizations.
- **Hybrid:** A combination of both centralized and decentralized teams can work together to execute cost optimizations.

The function may be measured against their ability to execute and deliver against cost optimization goals (for example, workload efficiency metrics).

You must secure executive sponsorship for this function to make changes, which is a key success factor. The sponsor is regarded as champion for cost efficient cloud consumption, and provides escalation support for the function to ensure that cost optimization activities are treated with the level of priority defined by the organization. Otherwise, guidance will be ignored and cost-saving opportunities will not be prioritized. Together, the sponsor and function ensure that your organization consumes the cloud efficiently and continues to deliver business value.

If you have a Business, Enterprise-On-Ramp, or Enterprise Support plan, and need help to build this team or function, reach out to Cloud Finance Management (CFM) experts through your Account team.

Implementation steps

- **Define key members:** You need to ensure that all relevant parts of your organization contribute and have a stake in cost management. Common teams within organizations typically include: finance, application or product owners, management, and technical teams (DevOps). Some are engaged full time (finance, technical), others periodically as required. Individuals or teams performing CFM generally need the following set of skills:

- Software development skills - in the case where scripts and automation are being built out.
- Infrastructure engineering skills - to deploy scripts or automation, and understand how services or resources are provisioned.
- Operations acumen - CFM is about operating on the cloud efficiently by measuring, monitoring, modifying, planning and scaling efficient use of the cloud.
- **Define goals and metrics:** The function needs to deliver value to the organization in different ways. These goals are defined and continually evolve as the organization evolves. Common activities include: creating and executing education programs on cost optimization across the organization, developing organization-wide standards, such as monitoring and reporting for cost optimization, and setting workload goals on optimization. This function also needs to regularly report to the organization on the organization's cost optimization capability.

You can define value-based key performance indicators (KPIs). KPIs can be cost-based or value-based. When you define the KPIs, you can calculate expected cost in terms of efficiency and expected business outcome. Value-based KPIs tie cost and usage metrics to business value drivers and help us rationalize changes in our AWS spend. The first step to deriving value-based KPIs is working together, cross-organizationally, to select and agree upon a standard set of KPIs.

- **Establish regular cadence:** The group (finance, technology, and business teams) should come together regularly to review their goals and metrics. A typical cadence involves reviewing the state of the organization, reviewing any programs currently running, and reviewing overall financial and optimization metrics. Then key workloads are reported on in greater detail.

During these regular meetings, you can review workload efficiency (cost) and business outcome. For example, a 20% cost increase for a workload may align with increased customer usage. In this case, this 20% cost increase can be interpreted as an investment. These regular cadence calls can help teams to identify value-based KPIs that provide meaning to the entire organization.

Resources

Related documents:

- [AWS CCOE Blog](#)
- [Creating Cloud Business Office](#)
- [CCOE - Cloud Center of Excellence](#)

Related videos:

- [Vanguard CCOE Success Story](#)

Related examples:

- [Using a Cloud Center of Excellence \(CCOE\) to Transform the Entire Enterprise](#)
- [Building a CCOE to transform the entire enterprise](#)
- [7 Pitfalls to Avoid When Building CCOE](#)

COST01-BP02 Establish a partnership between finance and technology

Involve finance and technology teams in cost and usage discussions at all stages of your cloud journey. Teams regularly meet and discuss topics such as organizational goals and targets, current state of cost and usage, and financial and accounting practices.


Level of risk exposed if this best practice is not established: High

Implementation guidance

Technology teams innovate faster in the cloud due to shortened approval, procurement, and infrastructure deployment cycles. This can be an adjustment for finance organizations previously used to executing time-consuming and resource-intensive processes for procuring and deploying capital in data center and on-premises environments, and cost allocation only at project approval.

From a finance and procurement organization perspective, the process for capital budgeting, capital requests, approvals, procurement, and installing physical infrastructure is one that has been learned and standardized over decades:


- Engineering or IT teams are typically the requesters
- Various finance teams act as approvers and procurers
- Operations teams rack, stack, and hand off ready-to-use infrastructure



With the adoption of cloud, infrastructure procurement and consumption are no longer beholden to a chain of dependencies. In the cloud model, technology and product teams are no longer just builders, but operators and owners of their products, responsible for most of the activities historically associated with finance and operations teams, including procurement and deployment.

All it really takes to provision cloud resources is a user account, and the right set of permissions. This is also what reduces IT and finance risk; which means teams are always a just few clicks or API calls away

from terminating idle or unnecessary cloud resources. This is also what allows technology teams to innovate faster – the agility and ability to spin up and then tear down experiments. While the variable nature of cloud consumption may impact predictability from a capital budgeting and forecasting perspective, cloud provides organizations with the ability to reduce the cost of over-provisioning, as well as reduce the opportunity cost associated with conservative under-provisioning.



Establish a partnership between key finance and technology stakeholders to create a shared understanding of organizational goals and develop mechanisms to succeed financially in the variable spend model of cloud computing. Relevant teams within your organization must be involved in cost and usage discussions at all stages of your cloud journey, including:

- **Financial leads:** CFOs, financial controllers, financial planners, business analysts, procurement, sourcing, and accounts payable must understand the cloud model of consumption, purchasing options, and the monthly invoicing process. Finance needs to partner with technology teams to create and socialize an IT value story, helping business teams understand how technology spend is linked to business outcomes. This way, technology expenditures are viewed not as costs, but rather as investments. Due to the fundamental differences between the cloud (such as the rate of change in usage, pay as you go pricing, tiered pricing, pricing models, and detailed billing and usage information) compared to on-premises operation, it is essential that the finance organization understands how cloud usage can impact business aspects including procurement processes, incentive tracking, cost allocation and financial statements.
- **Technology leads:** Technology leads (including product and application owners) must be aware of the financial requirements (for example, budget constraints) as well as business requirements (for example, service level agreements). This allows the workload to be implemented to achieve the desired goals of the organization.

The partnership of finance and technology provides the following benefits:

- Finance and technology teams have near real-time visibility into cost and usage.
- Finance and technology teams establish a standard operating procedure to handle cloud spend variance.
- Finance stakeholders act as strategic advisors with respect to how capital is used to purchase commitment discounts (for example, Reserved Instances or AWS Savings Plans), and how the cloud is used to grow the organization.
- Existing accounts payable and procurement processes are used with the cloud.
- Finance and technology teams collaborate on forecasting future AWS cost and usage to align and build organizational budgets.
- Better cross-organizational communication through a shared language, and common understanding of financial concepts.

Additional stakeholders within your organization that should be involved in cost and usage discussions include:

- **Business unit owners:** Business unit owners must understand the cloud business model so that they can provide direction to both the business units and the entire company. This cloud knowledge is critical when there is a need to forecast growth and workload usage, and when assessing longer-term purchasing options, such as Reserved Instances or Savings Plans.
- **Engineering team:** Establishing a partnership between finance and technology teams is essential for building a cost-aware culture that encourages engineers to take action on Cloud Financial Management (CFM). One of the common problems of CFM or finance operations practitioners and finance teams is getting engineers to understand the whole business on cloud, follow best practices, and take recommended actions.
- **Third parties:** If your organization uses third parties (for example, consultants or tools), ensure that they are aligned to your financial goals and can demonstrate both alignment through their engagement models and a return on investment (ROI). Typically, third parties will contribute to reporting and analysis of any workloads that they manage, and they will provide cost analysis of any workloads that they design.

Implementing CFM and achieving success requires collaboration across finance, technology, and business teams, and a shift in how cloud spend is communicated and evaluated across the organization. Include engineering teams so that they can be part of these cost and usage discussions at all stages, and encourage them to follow best practices and take agreed-upon actions accordingly.

Implementation steps

- **Define key members:** Verify that all relevant members of your finance and technology teams participate in the partnership. Relevant finance members will be those having interaction with the cloud bill. This will typically be CFOs, financial controllers, financial planners, business analysts, procurement, and sourcing. Technology members will typically be product and application owners, technical managers and representatives from all teams that build on the cloud. Other members may include business unit owners, such as marketing, that will influence usage of products, and third parties such as consultants, to achieve alignment to your goals and mechanisms, and to assist with reporting.
- **Define topics for discussion:** Define the topics that are common across the teams, or will need a shared understanding. Follow cost from that time it is created, until the bill is paid. Note any members involved, and organizational processes that are required to be applied. Understand each step or process it goes through and the associated information, such as pricing models available, tiered pricing, discount models, budgeting, and financial requirements.
- **Establish regular cadence:** To create a finance and technology partnership, establish a regular communication cadence to create and maintain alignment. The group needs to come together regularly against their goals and metrics. A typical cadence involves reviewing the state of the

organization, reviewing any programs currently running, and reviewing overall financial and optimization metrics. Then key workloads are reported on in greater detail.

Resources

Related documents:

- [AWS News Blog](#)

COST01-BP03 Establish cloud budgets and forecasts

Adjust existing organizational budgeting and forecasting processes to be compatible with the highly variable nature of cloud costs and usage. Processes must be dynamic using trend-based or business driver-based algorithms, or a combination of both.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Customers use the cloud for efficiency, speed and agility, which creates a highly variable amount of cost and usage. Costs can decrease with increases in workload efficiency, or as new workloads and features are deployed. It is possible to see the cost increase when the workload efficiency increases, or as new workloads and features are deployed. Or, workloads will scale to serve more of your customers, which increases cloud usage and costs. Resources are now more readily accessible than ever before. With the elasticity of the cloud also brings an elasticity of costs and forecasts. Existing organizational budgeting processes must be modified to incorporate this variability.

Adjust existing budgeting and forecasting processes to become more dynamic using either a trend-based algorithm (using historical costs as inputs), or using business-driver-based algorithms (for example, new product launches or regional expansion), or a combination of both trend and business drivers.

Use [AWS Budgets](#) to set custom budgets at a granular level by specifying the time period, recurrence, or amount (fixed or variable), and adding filters such as service, AWS Region, and tags. To stay informed on the performance of your existing budgets you can create and schedule [AWS Budgets Reports](#) to be emailed to you and your stakeholders on a regular cadence. You can also create [AWS Budgets Alerts](#) based on actual costs, which is reactive in nature, or on forecasted costs, which provides time to implement mitigations against potential cost overruns. You will be alerted when your cost or usage exceeds, or if they are forecasted to exceed, your budgeted amount.

AWS gives you the flexibility to build dynamic forecasting and budgeting processes so you can stay informed on whether costs adhere to, or exceed, budgetary limits.

Use [AWS Cost Explorer](#) to forecast costs in a defined future time range based on your past spend. AWS Cost Explorer's forecasting engine segments your historical data based on charge types (for example, Reserved Instances) and uses a combination of machine learning and rule-based models to predict spend across all charge types individually. Use [AWS Cost Explorer](#) to forecast daily (up to three months) or monthly (up to 12 months) cloud costs based on machine learning algorithms applied to your historical costs (trend-based).

Once you've determined your trend-based forecast using Cost Explorer, use the [AWS Pricing Calculator](#) to estimate your AWS use case and future costs based on the expected usage (traffic, requests-per-second, required Amazon Elastic Compute Cloud (Amazon EC2) instance, and so forth). You can also use it to help you plan how you spend, find cost saving opportunities, and make informed decisions when using AWS.

Use [AWS Cost Anomaly Detection](#) to prevent or reduce cost surprises and enhance control without slowing innovation. AWS Cost Anomaly Detection leverages advanced machine learning technologies to identify anomalous spend and root causes, so you can quickly take action. [With three simple steps](#), you can create your own contextualized monitor and receive alerts when any anomalous spend is detected.

Let builders build, and let AWS Cost Anomaly Detection monitor your spend and reduce the risk of billing surprises.

As mentioned in the [Well-Architected Cost Optimization Pillar's Finance and Technology Partnership](#) section, it is important to have partnership and cadences between IT, Finance and other stakeholders to ensure that they are all using the same tooling or processes for consistency. In cases where budgets may need to change, increasing cadence touch points can help react to those changes more quickly.

Implementation steps

- **Update existing budget and forecasting processes:** Implement trend-based, business driver-based, or a combination of both in your budgeting and forecasting processes.
- **Configure alerts and notifications:** Use AWS Budgets Alerts and Cost Anomaly Detection.
- **Perform regular reviews with key stakeholders:** For example, stakeholders in IT, Finance, Platform, and other areas of the business, to align with changes in business direction and usage.

Resources

Related documents:

- [AWS News Blog](#)
- [AWS Cost Explorer](#)
- [AWS Budgets](#)
- [AWS Pricing Calculator](#)
- [AWS Cost Anomaly Detection](#)
- [AWS License Manager](#)

COST01-BP04 Implement cost awareness in your organizational processes

Implement cost awareness, create transparency, and accountability of costs into new or existing processes that impact usage, and leverage existing processes for cost awareness. Implement cost awareness into employee training.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Cost awareness must be implemented in new and existing organizational processes. It is one of the foundational, prerequisite capabilities for other best practices. It is recommended to reuse and modify existing processes where possible — this minimizes the impact to agility and velocity. Report cloud costs to the technology teams and the decision makers in the business and finance teams to raise cost awareness, and establish efficiency key performance indicators (KPIs) for finance and business stakeholders. The following recommendations will help implement cost awareness in your workload:

- Verify that change management includes a cost measurement to quantify the financial impact of your changes. This helps proactively address cost-related concerns and highlight cost savings.
- Verify that cost optimization is a core component of your operating capabilities. For example, you can leverage existing incident management processes to investigate and identify root causes for cost and usage anomalies or cost overruns.
- Accelerate cost savings and business value realization through automation or tooling. When thinking about the cost of implementing, frame the conversation to include a return on investment (ROI) component to justify the investment of time or money.
- Allocate cloud costs by implementing showbacks or chargebacks for cloud spend, including spend on commitment-based purchase options, shared services and marketplace purchases to drive most cost-aware cloud consumption.

- Extend existing training and development programs to include cost-awareness training throughout your organization. It is recommended that this includes continuous training and certification. This will build an organization that is capable of self-managing cost and usage.
- Take advantage of free AWS native tools such as [AWS Cost Anomaly Detection](#), [AWS Budgets](#), and [AWS Budgets Reports](#).

When organizations consistently adopt [Cloud Financial Management](#) (CFM) practices, those behaviours become ingrained in the way of working and decision-making. The result is a culture that is more cost-aware, from developers architecting a new born-in-the-cloud application, to finance managers analyzing the ROI on these new cloud investments.

Implementation steps

- **Identify relevant organizational processes:** Each organizational unit reviews their processes and identifies processes that impact cost and usage. Any processes that result in the creation or termination of a resource need to be included for review. Look for processes that can support cost awareness in your business, such as incident management and training.
- **Establish self-sustaining cost-aware culture:** Make sure all the relevant stakeholders align with cause-of-change and impact as a cost so that they understand cloud cost. This will allow your organization to establish a self-sustaining cost-aware culture of innovation.
- **Update processes with cost awareness:** Each process is modified to be made cost aware. The process may require additional pre-checks, such as assessing the impact of cost, or post-checks validating that the expected changes in cost and usage occurred. Supporting processes such as training and incident management can be extended to include items for cost and usage.

To get help, reach out to CFM experts through your Account team, or explore the resources and related documents below.

Resources

Related documents:

- [AWS Cloud Financial Management](#)
- [AWS News Blog](#)

COST01-BP05 Report and notify on cost optimization

Configure AWS Budgets and AWS Cost Anomaly Detection to provide notifications on cost and usage against targets. Have regular meetings to analyze your workload's cost efficiency and to promote cost-aware culture.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

You must regularly report on cost and usage optimization within your organization. You can implement dedicated sessions to cost optimization, or include cost optimization in your regular operational reporting cycles for your workloads. Use services and tools to identify and implement cost savings opportunities. [AWS Cost Explorer](#) provides dashboards and reports. You can track your progress of cost and usage against configured budgets with [AWS Budgets Reports](#).

Use [AWS Budgets](#) to set custom budgets to track your costs and usage, and respond quickly to alerts received from email or Amazon Simple Notification Service (Amazon SNS) notifications if you exceed your threshold. [Set your preferred budget](#) period to daily, monthly, quarterly, or annually, and create specific budget limits to stay informed on how actual or forecasted costs and usage progress toward your

budget threshold. You can also set up [alerts](#) and [actions](#) against those alerts to run automatically, or through an approval process when a budget target is exceeded.

Implement notifications on cost and usage to ensure that changes in cost and usage can be acted upon quickly if they are unexpected. [AWS Cost Anomaly Detection](#) allows you to reduce cost surprises and enhance control without slowing innovation. AWS Cost Anomaly Detection identifies anomalous spend and root causes, which helps to reduce the risk of billing surprises. With three simple steps, you can create your own contextualized monitor and receive alerts when any anomalous spend is detected.

You can also use [Amazon QuickSight](#) with AWS Cost and Usage Report (CUR) data, to provide highly customized reporting with more granular data. Amazon QuickSight allows you to schedule reports and receive periodic Cost Report emails for historical cost and usage, or cost-saving opportunities.

Use [AWS Trusted Advisor](#), which provides guidance to verify whether provisioned resources are aligned with AWS best practices for cost optimization.

Periodically create reports containing a highlight of Savings Plans, Reserved Instances and Amazon Elastic Compute Cloud (Amazon EC2) rightsizing recommendations from AWS Cost Explorer to start reducing the cost associated with steady-state workloads, idle, and underutilized resources. Identify and recoup spend associated with cloud waste for resources that are deployed. Cloud waste occurs when incorrectly-sized resources are created, or different usage patterns are observed instead what is expected. Follow AWS best practices to reduce your waste and [optimize and save](#) your cloud costs.

Generate reports regularly for better purchasing options for your resources to drive down unit costs for your workloads. Purchasing options such as Savings Plans, Reserved Instances, or Amazon EC2 Spot Instances offer the deepest cost savings for fault-tolerant workloads and allow stakeholders (business owners, finance and tech teams) to be part of these commitment discussions.

Share the reports that contain opportunities or new release announcements that may help you to reduce total cost of ownership (TCO) of the cloud. Adopt new services, Regions, features, solutions, or new ways to achieve further cost reductions.

Implementation steps

- **Configure AWS Budgets:** Configure AWS Budgets on all accounts for your workload. Set a budget for the overall account spend, and a budget for the workload by using tags.
 - [Well-Architected Labs: Cost and Governance Usage](#)
- **Report on cost optimization:** Set up a regular cycle to discuss and analyze the efficiency of the workload. Using the metrics established, report on the metrics achieved and the cost of achieving them. Identify and fix any negative trends, and identify positive trends that you can promote across your organization. Reporting should involve representatives from the application teams and owners, finance, and management.
 - [Well-Architected Labs: Visualization](#)

Resources

Related documents:

- [AWS Cost Explorer](#)
- [AWS Trusted Advisor](#)
- [AWS Budgets](#)
- [AWS Budgets Best Practices](#)
- [Amazon CloudWatch](#)
- [AWS CloudTrail](#)
- [Amazon S3 Analytics](#)

- [AWS Cost and Usage Report](#)

Related examples:

- [Well-Architected Labs: Cost and Governance Usage](#)
- [Well-Architected Labs: Visualization](#)
- [Key ways to start optimizing your AWS cloud costs](#)

COST01-BP06 Monitor cost proactively

Implement tooling and dashboards to monitor cost proactively for the workload. Regularly review the costs with configured tools or out of the box tools, do not just look at costs and categories when you receive notifications. Monitoring and analyzing costs proactively helps to identify positive trends and allows you to promote them throughout your organization.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

It is recommended to monitor cost and usage proactively within your organization, not just when there are exceptions or anomalies. Highly visible dashboards throughout your office or work environment ensure that key people have access to the information they need, and indicate the organization's focus on cost optimization. Visible dashboards allow you to actively promote successful outcomes and implement them throughout your organization.

Create a daily or frequent routine to use [AWS Cost Explorer](#) or any other dashboard such as [Amazon QuickSight](#) to see the costs and analyze proactively. Analyze AWS service usage and costs at the AWS account-level, workload-level, or specific AWS service-level with grouping and filtering, and validate whether they are expected or not. Use the hourly- and resource-level granularity and tags to filter and identify incurring costs for the top resources. You can also build your own reports with the [Cost Intelligence Dashboard](#), an [Amazon QuickSight](#) solution built by AWS Solutions Architects, and compare your budgets with the actual cost and usage.

Implementation steps

- **Report on cost optimization:** Set up a regular cycle to discuss and analyze the efficiency of the workload. Using the metrics established, report on the metrics achieved and the cost of achieving them. Identify and fix any negative trends, and identify positive trends to promote across your organization. Reporting should involve representatives from the application teams and owners, finance, and management.
- **Create and enable daily granularity [AWS Budgets](#) for the cost and usage to take timely actions to prevent any potential cost overruns:** AWS Budgets allow you to configure alert notifications, so you stay informed if any of your budget types fall out of your pre-configured thresholds. The best way to leverage AWS Budgets is to set your expected cost and usage as your limits, so that anything above your budgets can be considered overspend.
- **Create [AWS Cost Anomaly Detection](#) for cost monitor:** [AWS Cost Anomaly Detection](#) uses advanced Machine Learning technology to identify anomalous spend and root causes, so you can quickly take action. It allows you to configure cost monitors that define spend segments you want to evaluate (for example, individual AWS services, member accounts, cost allocation tags, and cost categories), and lets you set when, where, and how you receive your alert notifications. For each monitor, attach multiple alert subscriptions for business owners and technology teams, including a name, a cost impact threshold, and alerting frequency (individual alerts, daily summary, weekly summary) for each subscription.
- **Use [AWS Cost Explorer](#) or integrate your [AWS Cost and Usage Report \(CUR\)](#) data with [Amazon QuickSight](#) dashboards to visualize your organization's costs:** AWS Cost Explorer has an easy-to-use

interface that lets you visualize, understand, and manage your AWS costs and usage over time. The [Cost Intelligence Dashboard](#) is a customizable and accessible dashboard to help create the foundation of your own cost management and optimization tool.

Resources

Related documents:

- [AWS Budgets](#)
- [AWS Cost Explorer](#)
- [Daily Cost and Usage Budgets](#)
- [AWS Cost Anomaly Detection](#)

Related examples:

- [Well-Architected Labs: Visualization](#)
- [Well-Architected Labs: Advanced Visualization](#)
- [Well-Architected Labs: Cloud Intelligence Dashboards](#)
- [Well-Architected Labs: Cost Visualization](#)
- [AWS Cost Anomaly Detection Alert with Slack](#)

COST01-BP07 Keep up-to-date with new service releases

Consult regularly with experts or AWS Partners to consider which services and features provide lower cost. Review AWS blogs and other information sources.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

AWS is constantly adding new capabilities so you can leverage the latest technologies to experiment and innovate more quickly. You may be able to implement new AWS services and features to increase cost efficiency in your workload. Regularly review [AWS Cost Management](#), the [AWS News Blog](#), the [AWS Cost Management blog](#), and [What's New with AWS](#) for information on new service and feature releases. What's New posts provide a brief overview of all AWS service, feature, and Region expansion announcements as they are released.

Implementation steps

- **Subscribe to blogs:** Go to the AWS blogs pages and subscribe to the What's New Blog and other relevant blogs. You can sign up on the [communication preference](#) page with your email address.
- **Subscribe to AWS News:** Regularly review the [AWS News Blog](#) and [What's New with AWS](#) for information on new service and feature releases. Subscribe to the RSS feed, or with your email to follow announcements and releases.
- **Follow AWS Price Reductions:** Regular price cuts on all our services has been a standard way for AWS to pass on the economic efficiencies to our customers gained from our scale. As of April 2022, AWS has reduced prices 115 times since it was launched in 2006. If you have any pending business decisions due to price concerns, you can review them again after price reductions and new service integrations. You can learn about the previous price reductions efforts, including Amazon Elastic Compute Cloud (Amazon EC2) instances, in the [price-reduction category of the AWS News Blog](#).
- **AWS events and meetups:** Attend your local AWS summit, and any local meetups with other organizations from your local area. If you cannot attend in person, try to attend virtual events to hear more from AWS experts and other customers' business cases.

- **Meet with your account team:** Schedule a regular cadence with your account team, meet with them and discuss industry trends and AWS services. Speak with your account manager, Solutions Architect, and support team.

Resources

Related documents:

- [Amazon EC2 – 15 Years of Optimizing and Saving Your IT Costs](#)
- [AWS Cost Management](#)
- [What's New with AWS](#)
- [AWS News Blog](#)
- [AWS Price Reduction](#)

Expenditure and usage awareness

Questions

- [COST 2 How do you govern usage? \(p. 358\)](#)
- [COST 3 How do you monitor usage and cost? \(p. 365\)](#)
- [COST 4 How do you decommission resources? \(p. 370\)](#)

COST 2 How do you govern usage?

Establish policies and mechanisms to ensure that appropriate costs are incurred while objectives are achieved. By employing a checks-and-balances approach, you can innovate without overspending.

Best practices

- [COST02-BP01 Develop policies based on your organization requirements \(p. 358\)](#)
- [COST02-BP02 Implement goals and targets \(p. 360\)](#)
- [COST02-BP03 Implement an account structure \(p. 360\)](#)
- [COST02-BP04 Implement groups and roles \(p. 362\)](#)
- [COST02-BP05 Implement cost controls \(p. 363\)](#)
- [COST02-BP06 Track project lifecycle \(p. 364\)](#)

COST02-BP01 Develop policies based on your organization requirements

Develop policies that define how resources are managed by your organization. Policies should cover cost aspects of resources and workloads, including creation, modification and decommission over the resource lifetime.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Understanding your organization's costs and drivers is critical for managing your cost and usage effectively, and identifying cost-reduction opportunities. Organizations typically operate multiple workloads run by multiple teams. These teams can be in different organization units, each with its own revenue stream. The capability to attribute resource costs to the workloads, individual organization, or product owners drives efficient usage behaviour and helps reduce waste. Accurate cost and usage monitoring allows you to understand how profitable organization units and products are, and allows you

to make more informed decisions about where to allocate resources within your organization. Awareness of usage at all levels in the organization is key to driving change, as change in usage drives changes in cost. Consider taking a multi-faceted approach to becoming aware of your usage and expenditures.

The first step in performing governance is to use your organization's requirements to develop policies for your cloud usage. These policies define how your organization uses the cloud and how resources are managed. Policies should cover all aspects of resources and workloads that relate to cost or usage, including creation, modification, and decommission over the resource's lifetime.

Policies should be simple so that they are easily understood and can be implemented effectively throughout the organization. Start with broad, high-level policies, such as which geographic Region usage is allowed in, or times of the day that resources should be running. Gradually refine the policies for the various organizational units and workloads. Common policies include which services and features can be used (for example, lower performance storage in test or development environments), and which types of resources can be used by different groups (for example, the largest size of resource in a development account is medium).

Implementation steps

- **Meet with team members:** To develop policies, get all team members from your organization to specify their requirements and document them accordingly. Take an iterative approach by starting broadly and continually refine down to the smallest units at each step. Team members include those with direct interest in the workload, such as organization units or application owners, as well as supporting groups, such as security and finance teams.
- **Define locations for your workload:** Define where your workload operates, including the country and the area within the country. This information is used for mapping to AWS Regions and Availability Zones.
- **Define and group services and resources:** Define the services that the workloads require. For each service, specify the types, the size, and the number of resources required. Define groups for the resources by function, such as application servers or database storage. Resources can belong to multiple groups.
- **Define and group the users by function:** Define the users that interact with the workload, focusing on what they do and how they use the workload, not on who they are or their position in the organization. Group similar users or functions together. You can use the AWS managed policies as a guide.
- **Define the actions:** Using the locations, resources, and users identified previously, define the actions that are required by each to achieve the workload outcomes over its life time (development, operation, and decommission). Identify the actions based on the groups, not the individual elements in the groups, in each location. Start broadly with read or write, then refine down to specific actions to each service.
- **Define the review period:** Workloads and organizational requirements can change over time. Define the workload review schedule to ensure it remains aligned with organizational priorities.
- **Document the policies:** Ensure the policies that have been defined are accessible as required by your organization. These policies are used to implement, maintain, and audit access of your environments.

Resources

Related documents:

- [AWS Managed Policies for Job Functions](#)
- [AWS multiple account billing strategy](#)
- [Actions, Resources, and Condition Keys for AWS Services](#)
- [Cloud Products](#)
- [Control access to AWS Regions using IAM policies](#)
- [Global Infrastructures Regions and AZs](#)

COST02-BP02 Implement goals and targets

Implement both cost and usage goals for your workload. Goals provide direction to your organization on cost and usage, and targets provide measurable outcomes for your workloads.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Develop cost and usage goals and targets for your organization. Goals provide guidance and direction to your organization on expected outcomes. Targets provide specific measurable outcomes to be achieved. An example of a goal is: platform usage should increase significantly, with only a minor (non-linear) increase in cost. An example target is: a 20% increase in platform usage, with less than a 5% increase in costs. Another common goal is that workloads need to be more efficient every 6 months. The accompanying target would be that the cost per output of the workload needs to decrease by 5% every 6 months.

A common goal for cloud workloads is to increase workload efficiency, which is to decrease the cost per business outcome of the workload over time. It is recommended to implement this goal for all workloads, and also set a target such as a 5% increase in efficiency every 6- 12 months. This can be achieved in the cloud through building capability in cost optimization, and through the release of new services and service features.

Implementation steps

- **Define expected usage levels:** Focus on usage levels to begin with. Engage with the application owners, marketing, and greater business teams to understand what the expected usage levels will be for the workload. How will customer demand change over time, and will there be any changes due to seasonal increases or marketing campaigns.
- **Define workload resourcing and costs:** With the usage levels defined, quantify the changes in workload resources required to meet these usage levels. You may need to increase the size or number of resources for a workload component, increase data transfer, or change workload components to a different service at a specific level. Specify what the costs will be at each of these major points, and what the changes in cost will be when there are changes in usage.
- **Define business goals:** Taking the output from the expected changes in usage and cost, combine this with expected changes in technology, or any programs that you are running, and develop goals for the workload. Goals must address usage, cost and the relation between the two. Verify that there are organizational programs, for example capability building like training and education, if there are expected changes in cost without changes in usage.
- **Define targets:** For each of the defined goals specify a measurable target. If a goal is to increase efficiency in the workload, the target will quantify the amount of improvement, typical in business outputs for each dollar spent, and when it will be delivered.

Resources

Related documents:

- [AWS managed policies for job functions](#)
- [AWS multiple account billing strategy](#)
- [Control access to AWS Regions using IAM policies](#)

COST02-BP03 Implement an account structure

Implement a structure of accounts that maps to your organization. This assists in allocating and managing costs throughout your organization.

Level of risk exposed if this best practice is not established: High

Implementation guidance

AWS has a one-parent-to-many-children account structure that is commonly known as a management account (the parent, formerly payer) account-member (the child, formerly linked) account. A best practice is to always have at least one management account with one member account, regardless of your organization size or usage. All workload resources should reside only within member accounts.

There is no one-size-fits-all answer for how many AWS accounts you should have. Assess your current and future operational and cost models to ensure that the structure of your AWS accounts reflects your organization's goals. Some companies create multiple AWS accounts for business reasons, for example:

- Administrative and/or fiscal and billing isolation is required between organization units, cost centers, or specific workloads.
- AWS service limits are set to be specific to particular workloads.
- There is a requirement for isolation and separation between workloads and resources.

Within [AWS Organizations](#), [consolidated billing](#) creates the construct between one or more member accounts and the management account. Member accounts allow you to isolate and distinguish your cost and usage by groups. A common practice is to have separate member accounts for each organization unit (such as finance, marketing, and sales), or for each environment lifecycle (such as development, testing and production), or for each workload (workload a, b, and c), and then aggregate these linked accounts using consolidated billing.

Consolidated billing allows you to consolidate payment for multiple member AWS accounts under a single management account, while still providing visibility for each linked account's activity. As costs and usage are aggregated in the management account, this allows you to maximize your service volume discounts, and maximize the use of your commitment discounts (Savings Plans and Reserved Instances) to achieve the highest discounts.

[AWS Control Tower](#) can quickly set up and configure multiple AWS accounts, ensuring that governance is aligned with your organization's requirements.

Implementation steps

- **Define separation requirements:** Requirements for separation are a combination of multiple factors, including security, reliability, and financial constructs. Work through each factor in order and specify whether the workload or workload environment should be separate from other workloads. Security ensures that access and data requirements are adhered to. Reliability ensures that limits are managed so that environments and workloads do not impact others. Financial constructs ensure that there is strict financial separation and accountability. Common examples of separation are production and test workloads being run in separate accounts, or using a separate account so that the invoice and billing data can be provided to a third-party organization.
- **Define grouping requirements:** Requirements for grouping do not override the separation requirements, but are used to assist management. Group together similar environments or workloads that do not require separation. An example of this is grouping multiple test or development environments from one or more workloads together.
- **Define account structure:** Using these separations and groupings, specify an account for each group and ensure that separation requirements are maintained. These accounts are your member or linked accounts. By grouping these member accounts under a single management or payer account, you combine usage, which allows for greater volume discounts across all accounts, and provides a single bill for all accounts. It's possible to separate billing data and provide each member account with an individual view of their billing data. If a member account must not have its usage or billing data visible to any other account, or if a separate bill from AWS is required, define multiple management or payer accounts. In this case, each member account has its own management or payer account. Resources

should always be placed in member or linked accounts. The management or payer accounts should only be used for management.

Resources

Related documents:

- [AWS managed policies for job functions](#)
- [AWS multiple account billing strategy](#)
- [Control access to AWS Regions using IAM policies](#)
- [AWS Control Tower](#)
- [AWS Organizations](#)
- [Consolidated billing](#)

Related examples:

- [Splitting the CUR and Sharing Access](#)

COST02-BP04 Implement groups and roles

Implement groups and roles that align to your policies and control who can create, modify, or decommission instances and resources in each group. For example, implement development, test, and production groups. This applies to AWS services and third-party solutions.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

After you develop policies, you can create logical groups and roles of users within your organization. This allows you to assign permissions and control usage. Begin with high-level groupings of people. Typically this aligns with organizational units and job roles (for example, systems administrator in the IT Department, or financial controller). The groups join people that do similar tasks and need similar access. Roles define what a group must do. For example, a systems administrator in IT requires access to create all resources, but an analytics team member only needs to create analytics resources.

Implementation steps

- **Implement groups:** Using the groups of users defined in your organizational policies, implement the corresponding groups, if necessary. Refer to the security pillar for best practices on users, groups, and authentication.
- **Implement roles and policies:** Using the actions defined in your organizational policies, create the required roles and access policies. Refer to the security pillar for best practices on roles and policies.

Resources

Related documents:

- [AWS managed policies for job functions](#)
- [AWS multiple account billing strategy](#)
- [Control access to AWS Regions using IAM policies](#)
- [Well-Architected Security Pillar](#)

Related examples:

- [Well-Architected Lab Basic Identity and Access](#)

COST02-BP05 Implement cost controls

Implement controls based on organization policies and defined groups and roles. These certify that costs are only incurred as defined by organization requirements: for example, control access to regions or resource types with AWS Identity and Access Management (IAM) policies.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

A common first step in implementing cost controls is to set up notifications when cost or usage events occur outside of the organization policies. This enables you to act quickly and verify if corrective action is required, without restricting or negatively impacting workloads or new activity. After you know the workload and environment limits, you can enforce governance. In AWS, notifications are conducted with AWS Budgets, which allows you to define a monthly budget for your AWS costs, usage, and commitment discounts (Savings Plans and Reserved Instances). You can create budgets at an aggregate cost level (for example, all costs), or at a more granular level where you include only specific dimensions such as linked accounts, services, tags, or Availability Zones.

As a second step, you can enforce governance policies in AWS through [AWS Identity and Access Management](#) (IAM), and [AWS Organizations Service Control Policies](#) (SCP). IAM allows you to securely manage access to AWS services and resources. Using IAM, you can control who can create and manage AWS resources, the type of resources that can be created, and where they can be created. This minimizes the creation of resources that are not required. Use the roles and groups created previously, and assign [IAM policies](#) to enforce the correct usage. SCP offers central control over the maximum available permissions for all accounts in your organization, ensuring that your accounts stay within your access control guidelines. SCPs are available only in an organization that has all features enabled, and you can configure the SCPs to either deny or allow actions for member accounts by default. Refer to the [Well-Architected Security Pillar whitepaper](#) for more details on implementing access management.

Governance can also be implemented through management of Service Quotas. By ensuring Service Quotas are set with minimum overhead and accurately maintained, you can minimize resource creation outside of your organization's requirements. To achieve this, you must understand how quickly your requirements can change, understand projects in progress (both creation and decommission of resources), and factor in how fast quota changes can be implemented. [Service Quotas](#) can be used to increase your quotas when required.

Implementation steps

- **Implement notifications on spend:** Using your defined organization policies, create AWS budgets to provide notifications when spending is outside of your policies. Configure multiple cost budgets, one for each account, which notifies you about overall account spending. Then configure additional cost budgets within each account for smaller units within the account. These units vary depending on your account structure. Some common examples are AWS Regions, workloads (using tags), or AWS services. Ensure that you configure an email distribution list as the recipient for notifications, and not an individual's email account. You can configure an actual budget for when an amount is exceeded, or use a forecasted budget for notifying on forecasted usage.
- **Implement controls on usage:** Using your defined organization policies, implement IAM policies and roles to specify which actions users can perform and which actions they cannot perform. Multiple organizational policies may be included in an AWS policy. In the same way that you defined policies, start broadly and then apply more granular controls at each step. Service limits are also an effective control on usage. Implement the correct service limits on all your accounts.

Resources

Related documents:

- [AWS managed policies for job functions](#)
- [AWS multiple account billing strategy](#)
- [Control access to AWS Regions using IAM policies](#)

Related examples:

- [Well-Architected Labs: Cost and Usage Governance](#)
- [Well-Architected Labs: Cost and Usage Governance](#)

COST02-BP06 Track project lifecycle

Track, measure, and audit the lifecycle of projects, teams, and environments to avoid using and paying for unnecessary resources.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Ensure that you track the entire lifecycle of the workload. This ensures that when workloads or workload components are no longer required, they can be decommissioned or modified. This is especially useful when you release new services or features. The existing workloads and components may appear to be in use, but should be decommissioned to redirect customers to the new service. Notice previous stages of workloads — after a workload is in production, previous environments can be decommissioned or greatly reduced in capacity until they are required again.

AWS provides a number of management and governance services you can use for entity lifecycle tracking. You can use [AWS Config](#) or [AWS Systems Manager](#) to provide a detailed inventory of your AWS resources and configuration. It is recommended that you integrate with your existing project or asset management systems to keep track of active projects and products within your organization. Combining your current system with the rich set of events and metrics provided by AWS allows you to build a view of significant lifecycle events and proactively manage resources to reduce unnecessary costs.

Refer to the [Well-Architected Operational Excellence Pillar whitepaper](#) for more details on implementing entity lifecycle tracking.

Implementation steps

- **Perform workload reviews:** As defined by your organizational policies, audit your existing projects. The amount of effort spent in the audit should be proportional to the approximate risk, value, or cost to the organization. Key areas to include in the audit would be risk to the organization of an incident or outage, value, or contribution to the organization (measured in revenue or brand reputation), cost of the workload (measured as total cost of resources and operational costs), and usage of the workload (measured in number of organization outcomes per unit of time). If these areas change over the lifecycle, adjustments to the workload are required, such as full or partial decommissioning.

Resources

Related documents:

- [AWS Config](#)
- [AWS Systems Manager](#)
- [AWS managed policies for job functions](#)
- [AWS multiple account billing strategy](#)
- [Control access to AWS Regions using IAM policies](#)

COST 3 How do you monitor usage and cost?

Establish policies and procedures to monitor and appropriately allocate your costs. This allows you to measure and improve the cost efficiency of this workload.

Best practices

- [COST03-BP01 Configure detailed information sources \(p. 365\)](#)
- [COST03-BP02 Identify cost attribution categories \(p. 366\)](#)
- [COST03-BP03 Establish organization metrics \(p. 367\)](#)
- [COST03-BP04 Configure billing and cost management tools \(p. 367\)](#)
- [COST03-BP05 Add organization information to cost and usage \(p. 368\)](#)
- [COST03-BP06 Allocate costs based on workload metrics \(p. 370\)](#)

COST03-BP01 Configure detailed information sources

Configure the AWS Cost and Usage Report, and Cost Explorer hourly granularity, to provide detailed cost and usage information. Configure your workload to have log entries for every delivered business outcome.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Enable hourly granularity in AWS Cost Explorer and create a [AWS Cost and Usage Report \(CUR\)](#). These data sources provide the most accurate view of cost and usage across your entire organization. The CUR provides daily or hourly usage granularity, rates, costs, and usage attributes for all chargeable AWS services. All possible dimensions are in the CUR including: tagging, location, resource attributes, and account IDs.

Configure your CUR with the following customizations:

- Include resource IDs
- Automatically refresh the CUR
- Hourly granularity
- **Versioning:** Overwrite existing report
- **Data integration:** Amazon Athena (Parquet format and compression)

Use [AWS Glue](#) to prepare the data for analysis, and use [Amazon Athena](#) to perform data analysis, using SQL to query the data. You can also use [Amazon QuickSight](#) to build custom and complex visualizations and distribute them throughout your organization.

Implementation steps

- **Configure the cost and usage report:** Using the billing console, configure at least one cost and usage report. Configure a report with hourly granularity that includes all identifiers and resource IDs. You can also create other reports with different granularities to provide higher-level summary information.
- **Configure hourly granularity in Cost Explorer:** Using the billing console, enable Hourly and Resource Level Data.

Note

There will be associated costs with enabling this feature. For details, refer to the pricing.

- **Configure application logging:** Verify that your application logs each business outcome that it delivers so it can be tracked and measured. Ensure that the granularity of this data is at least hourly so

it matches with the cost and usage data. Refer to the [Well-Architected Operational Excellence Pillar](#) for more detail on logging and monitoring.

Resources

Related documents:

- [AWS Account Setup](#)
- [AWS Cost and Usage Report \(CUR\)](#)
- [AWS Glue](#)
- [Amazon QuickSight](#)
- [AWS Cost Management Pricing](#)
- [Tagging AWS resources](#)
- [Analyzing your costs with AWS Budgets](#)
- [Analyzing your costs with Cost Explorer](#)
- [Managing AWS Cost and Usage Reports](#)
- [Well-Architected Operational Excellence Pillar](#)

Related examples:

- [AWS Account Setup](#)

COST03-BP02 Identify cost attribution categories

Identify organization categories that could be used to allocate cost within your organization.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Work with your finance team and other relevant stakeholders to understand the requirements of how costs must be allocated within your organization. Workload costs must be allocated throughout the entire lifecycle, including development, testing, production, and decommissioning. Understand how the costs incurred for learning, staff development, and idea creation are attributed in the organization. This can be helpful to correctly allocate accounts used for this purpose to training and development budgets, instead of generic IT cost budgets.

Implementation steps

- **Define your organization categories:** Meet with stakeholders to define categories that reflect your organization's structure and requirements. These will directly map to the structure of existing financial categories, such as business unit, budget, cost center, or department. Look at the outcomes the cloud delivers for your business, such as training or education, as these are also organization categories. Multiple categories can be assigned to a resource, and a resource can be in multiple different categories, so define as many categories as needed.
- **Define your functional categories:** Meet with stakeholders to define categories that reflect the functions that you have within your business. This may be the workload or application names, and the type of environment, such as production, testing, or development. Multiple categories can be assigned to a resource, and a resource can be in multiple different categories, so define as many categories as needed.

Resources

Related documents:

- [Tagging AWS resources](#)
- [Analyzing your costs with AWS Budgets](#)
- [Analyzing your costs with Cost Explorer](#)
- [Managing AWS Cost and Usage Reports](#)

COST03-BP03 Establish organization metrics

Establish the organization metrics that are required for this workload. Example metrics of a workload are customer reports produced, or web pages served to customers.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Understand how your workload's output is measured against business success. Each workload typically has a small set of major outputs that indicate performance. If you have a complex workload with many components, then you can prioritize the list, or define and track metrics for each component. Work with your teams to understand which metrics to use. This unit will be used to understand the efficiency of the workload, or the cost for each business output.

Implementation steps

- **Define workload outcomes:** Meet with the stakeholders in the business and define the outcomes for the workload. These are a primary measure of customer usage and must be business metrics and not technical metrics. There should be a small number of high-level metrics (less than five) per workload. If the workload produces multiple outcomes for different use cases, then group them into a single metric.
- **Define workload component outcomes:** Optionally, if you have a large and complex workload, or can easily break your workload into components (such as microservices) with well-defined inputs and outputs, define metrics for each component. The effort should reflect the value and cost of the component. Start with the largest components and work towards the smaller components.

Resources

Related documents:

- [Tagging AWS resources](#)
- [Analyzing your costs with AWS Budgets](#)
- [Analyzing your costs with Cost Explorer](#)
- [Managing AWS Cost and Usage Reports](#)

COST03-BP04 Configure billing and cost management tools

Configure AWS Cost Explorer and AWS Budgets inline with your organization policies.

Level of risk exposed if this best practice is not established: High

Implementation guidance

To modify usage and adjust costs, each person in your organization must have access to their cost and usage information. It is recommended that all workloads and teams have the following tooling configured when they use the cloud:

- **Reports:** Summarize of all cost and usage information

- **Notifications:** Provide notifications when cost or usage is outside of defined limits.
- **Current State:** Configure a dashboard showing current levels of cost and usage. The dashboard should be available in a highly visible place within the work environment (similar to an operations dashboard).
- **Trending:** Provide the capability to show the variability in cost and usage over the required period of time, with the required granularity.
- **Forecasts:** Provide the capability to show estimated future costs.
- **Tracking:** Show the current cost and usage against configured goals or targets.
- **Analysis:** Provide the capability for team members to perform custom and deep analysis down to the hourly granularity, with all possible dimensions.

You can use AWS native tooling, such as [AWS Cost Explorer](#), [AWS Budgets](#), and [Amazon Athena](#) with [Amazon QuickSight](#) to provide this capability. You can also use third-party tooling — however, you must ensure that the costs of this tooling provide value to your organization.

Implementation steps

- **Create a Cost Optimization group:** Configure your account and create a group that has access to the required Cost and Usage reports. This group must include representatives from all teams that own or manage an application. This certifies that every team has access to their cost and usage information.
- **Configure AWS Budgets:** Configure AWS Budgets on all accounts for your workload. Set a budget for the overall account spend, and a budget for the workload by using tags.
- **Configure AWS Cost Explorer:** Configure AWS Cost Explorer for your workload and accounts. Create a dashboard for the workload that tracks overall spend, and key usage metrics for the workload.
- **Configure advanced tooling:** Optionally, you can create custom tooling for your organization that provides additional detail and granularity. You can implement advanced analysis capability using [Amazon Athena](#), and dashboards using [Amazon QuickSight](#).

Resources

Related documents:

- [Tagging AWS resources](#)
- [Analyzing your costs with AWS Budgets](#)
- [Analyzing your costs with Cost Explorer](#)
- [Managing AWS Cost and Usage Reports](#)

Related examples:

- [Well-Architected Labs - AWS Account Setup](#)
- [Well-Architected Labs: Billing Visualization](#)
- [Well-Architected Labs: Cost and Governance Usage](#)
- [Well-Architected Labs: Cost and Usage Analysis](#)
- [Well-Architected Labs: Cost and Usage Visualization](#)

COST03-BP05 Add organization information to cost and usage

Define a tagging schema based on organization, and workload attributes, and cost allocation categories. Implement tagging across all resources. Use Cost Categories to group costs and usage according to organization attributes.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Implement [tagging in AWS](#) to add organization information to your resources, which will then be added to your cost and usage information. A tag is a key-value pair— the key is defined and must be unique across your organization, and the value is unique to a group of resources. An example of a key-value pair is the key is Environment, with a value of Production. All resources in the production environment will have this key-value pair. Tagging allows you categorize and track your costs with meaningful, relevant organization information. You can apply tags that represent organization categories (such as cost centers, application names, projects, or owners), and identify workloads and characteristics of workloads (such as test or production) to attribute your costs and usage throughout your organization.

When you apply tags to your AWS resources (such as Amazon Elastic Compute Cloud instances or Amazon Simple Storage Service buckets) and activate the tags, AWS adds this information to your Cost and Usage Reports. You can run reports and perform analysis, on tagged and untagged resources to allow greater compliance with internal cost management policies, and ensure accurate attribution.

Creating and implementing an AWS tagging standard across your organization's accounts enables you to manage and govern your AWS environments in a consistent and uniform manner. Use [Tag Policies](#) in AWS Organizations to define rules for how tags can be used on AWS resources in your accounts in AWS Organizations. Tag Policies allow you to easily adopt a standardized approach for tagging AWS resources

[AWS Tag Editor](#) allows you to add, delete, and manage tags of multiple resources.

[AWS Cost Categories](#) allows you to assign organization meaning to your costs, without requiring tags on resources. You can map your cost and usage information to unique internal organization structures. You define category rules to map and categorize costs using billing dimensions, such as accounts and tags. This provides another level of management capability in addition to tagging. You can also map specific accounts and tags to multiple projects.

Implementation steps

- **Define a tagging schema:** Gather all stakeholders from across your business to define a schema. This typically includes people in technical, financial, and management roles. Define a list of tags that all resources must have, as well as a list of tags that resources should have. Verify that the tag names and values are consistent across your organization.
- **Tag resources:** Using your defined cost attribution categories, place tags on all resources in your workloads according to the categories. Use tools such as the CLI, Tag Editor, or Systems Manager, to increase efficiency.
- **Implement Cost Categories:** You can create Cost Categories without implementing tagging. Cost Categories use the existing cost and usage dimensions. Create category rules from your schema and implement it into Cost Categories.
- **Automate tagging:** To verify that you maintain high levels of tagging across all resources, automate tagging so that resources are automatically tagged when they are created. Use the features within the service, or services such as AWS CloudFormation, to ensure that resources are tagged when created. You can also create a custom microservice that scans the workload periodically and removes any resources that are not tagged, which is ideal for test and development environments.
- **Monitor and report on tagging:** To verify that you maintain high levels of tagging across your organization, report and monitor the tags across your workloads. You can use AWS Cost Explorer to view the cost of tagged and untagged resources, or use services such as Tag Editor. Regularly review the number of untagged resources and take action to add tags until you reach the desired level of tagging.

Resources

Related documents:

- [AWS CloudFormation Resource Tag](#)

- [AWS Cost Categories](#)
- [Tagging AWS resources](#)
- [Amazon EC2 and Amazon EBS add support for tagging resources upon creation](#)
- [Analyzing your costs with AWS Budgets](#)
- [Analyzing your costs with Cost Explorer](#)
- [Managing AWS Cost and Usage Reports](#)

COST03-BP06 Allocate costs based on workload metrics

Allocate the workload's costs by metrics or business outcomes to measure workload cost efficiency. Implement a process to analyze the AWS Cost and Usage Report with [Amazon Athena](#), which can provide insight and charge back capability.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Cost Optimization is delivering business outcomes at the lowest price point, which can only be achieved by allocating workload costs by workload metrics (measured by workload efficiency). Monitor the defined workload metrics through log files or other application monitoring. Combine this data with the workload costs, which can be obtained by looking at costs with a specific tag value or account ID. It is recommended to perform this analysis at the hourly level. Your efficiency will typically change if you have some static cost components (for example, a backend database running 24/7) with a varying request rate (for example, usage peaks at 9am – 5pm, with few requests at night). Understanding the relationship between the static and variable costs will help you to focus your optimization activities.

Implementation Steps

- **Allocate costs to workload metrics:** Using the defined metrics and tagging configured, create a metric that combines the workload output and workload cost. Use the analytics services such as Amazon Athena and Amazon QuickSight to create an efficiency dashboard for the overall workload, and any components.

Resources

Related documents:

- [Tagging AWS resources](#)
- [Analyzing your costs with AWS Budgets](#)
- [Analyzing your costs with Cost Explorer](#)
- [Managing AWS Cost and Usage Reports](#)

COST 4 How do you decommission resources?

Implement change control and resource management from project inception to end-of-life. This ensures you shut down or terminate unused resources to reduce waste.

Best practices

- [COST04-BP01 Track resources over their lifetime \(p. 371\)](#)
- [COST04-BP02 Implement a decommissioning process \(p. 371\)](#)
- [COST04-BP03 Decommission resources \(p. 372\)](#)
- [COST04-BP04 Decommission resources automatically \(p. 372\)](#)

COST04-BP01 Track resources over their lifetime

Define and implement a method to track resources and their associations with systems over their lifetime. You can use tagging to identify the workload or function of the resource.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Decommission workload resources that are no longer required. A common example is resources used for testing, after testing has been completed, the resources can be removed. Tracking resources with tags (and running reports on those tags) will help you identify assets for decommission. Using tags is an effective way to track resources, by labeling the resource with its function, or a known date when it can be decommissioned. Reporting can then be run on these tags. Example values for feature tagging are `feature-X` testing to identify the purpose of the resource in terms of the workload lifecycle.

Implementation steps

- **Implement a tagging scheme:** Implement a tagging scheme that identifies the workload the resource belongs to, verifying that all resources within the workload are tagged accordingly.
- **Implement workload throughput or output monitoring:** Implement workload throughput monitoring or alarming, triggering on either input requests or output completions. Configure it to provide notifications when workload requests or outputs drop to zero, indicating the workload resources are no longer used. Incorporate a time factor if the workload periodically drops to zero under normal conditions.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [AWS Trusted Advisor](#)
- [Tagging AWS resources](#)
- [Publishing Custom Metrics](#)

COST04-BP02 Implement a decommissioning process

Implement a process to identify and decommission orphaned resources.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Implement a standardized process across your organization to identify and remove unused resources. The process should define the frequency searches are performed, and the processes to remove the resource to ensure that all organization requirements are met.

Implementation steps

- **Create and implement a decommissioning process:** Working with the workload developers and owners, build a decommissioning process for the workload and its resources. The process should cover the method to verify if the workload is in use, and also if each of the workload resources are in use. The process should also cover the steps necessary to decommission the resource, removing them from service while ensuring compliance with any regulatory requirements. Any associated resources are also covered, such as licenses or attached storage. The process should provide notification to the workload owners that the decommissioning process has been executed.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [AWS Trusted Advisor](#)

COST04-BP03 Decommission resources

Decommission resources triggered by events such as periodic audits, or changes in usage. Decommissioning is typically performed periodically, and is manual or automated.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

The frequency and effort to search for unused resources should reflect the potential savings, so an account with a small cost should be analyzed less frequently than an account with larger costs. Searches and decommission events can be triggered by state changes in the workload, such as a product going end of life or being replaced. Searches and decommission events may also be triggered by external events, such as changes in market conditions or product termination.

Implementation steps

- **Decommission resources:** Using the decommissioning process, decommission each of the resources that have been identified as orphaned.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [AWS Trusted Advisor](#)

COST04-BP04 Decommission resources automatically

Design your workload to gracefully handle resource termination as you identify and decommission non-critical resources, resources that are not required, or resources with low utilization.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Use automation to reduce or remove the associated costs of the decommissioning process. Designing your workload to perform automated decommissioning will reduce the overall workload costs during its lifetime. You can use [AWS Auto Scaling](#) to perform the decommissioning process. You can also implement custom code using the [API or SDK](#) to decommission workload resources automatically.

Implementation steps

- **Implement AWS Auto Scaling:** For resources that are supported, configure them with AWS Auto Scaling.
- **Configure CloudWatch to terminate instances:** Instances can be configured to terminate using CloudWatch alarms. Using the metrics from the decommissioning process, implement an alarm with an Amazon Elastic Compute Cloud (Amazon EC2) action. Verify the operation in a non-production environment before rolling out.

- **Implement code within the workload:** You can use the AWS SDK or AWS CLI to decommission workload resources. Implement code within the application that integrates with AWS and terminates or removes resources that are no longer used.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [AWS Trusted Advisor](#)
- [Create Alarms to Stop, Terminate, Reboot, or Recover an Instance](#)
- [Getting Started with Amazon EC2 Auto Scaling](#)

Cost-effective resources

Questions

- [COST 5 How do you evaluate cost when you select services? \(p. 373\)](#)
- [COST 6 How do you meet cost targets when you select resource type, size and number? \(p. 378\)](#)
- [COST 7 How do you use pricing models to reduce cost? \(p. 380\)](#)
- [COST 8 How do you plan for data transfer charges? \(p. 384\)](#)

COST 5 How do you evaluate cost when you select services?

Amazon EC2, Amazon EBS, and Amazon S3 are building-block AWS services. Managed services, such as Amazon RDS and Amazon DynamoDB, are higher level, or application level, AWS services. By selecting the appropriate building blocks and managed services, you can optimize this workload for cost. For example, using managed services, you can reduce or remove much of your administrative and operational overhead, freeing you to work on applications and business-related activities.

Best practices

- [COST05-BP01 Identify organization requirements for cost \(p. 373\)](#)
- [COST05-BP02 Analyze all components of the workload \(p. 374\)](#)
- [COST05-BP03 Perform a thorough analysis of each component \(p. 375\)](#)
- [COST05-BP04 Select software with cost-effective licensing \(p. 376\)](#)
- [COST05-BP05 Select components of this workload to optimize cost in line with organization priorities \(p. 376\)](#)
- [COST05-BP06 Perform cost analysis for different usage over time \(p. 377\)](#)

COST05-BP01 Identify organization requirements for cost

Work with team members to define the balance between cost optimization and other pillars, such as performance and reliability, for this workload.

Level of risk exposed if this best practice is not established: High

Implementation guidance

When selecting services for your workload, it is key that you understand your organization priorities. Ensure that you have a balance between cost and other Well-Architected pillars, such as performance

and reliability. A fully cost-optimized workload is the solution that is most aligned to your organization's requirements, not necessarily the lowest cost. Meet with all teams within your organization to collect information, such as product, business, technical, and finance.

Implementation steps

- **Identify organization requirements for cost:** Meet with team members from your organization, including those in product management, application owners, development and operational teams, management, and financial roles. Prioritize the Well-Architected pillars for this workload and its components, the output is a list of the pillars in order. You can also add a weighting to each, which can indicate how much additional focus a pillar has, or how similar the focus is between two pillars.

Resources

Related documents:

- [AWS Total Cost of Ownership \(TCO\) Calculator](#)
- [Amazon S3 storage classes](#)
- [Cloud products](#)

COST05-BP02 Analyze all components of the workload

Verify every workload component is analyzed, regardless of current size or current costs. The review effort should reflect the potential benefit, such as current and projected costs.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Perform a thorough analysis on all components in your workload. Ensure that balance between the cost of analysis and the potential savings in the workload over its lifecycle. You must find the current impact, and potential future impact, of the component. For example, if the cost of the proposed resource is \$10 a month, and under forecasted loads would not exceed \$15 a month, spending a day of effort to reduce costs by 50% (\$5 a month) could exceed the potential benefit over the life of the system. Using a faster and more efficient data-based estimation will create the best overall outcome for this component.

Workloads can change over time, and the right set of services may not be optimal if the workload architecture or usage changes. Analysis for selection of services must incorporate current and future workload states and usage levels. Implementing a service for future workload state or usage may reduce overall costs by reducing or removing the effort required to make future changes.

[AWS Cost Explorer](#) and the [AWS Cost and Usage Report \(CUR\)](#) can analyze the cost of a Proof of Concept (PoC) or running environment. You can also use [AWS Pricing Calculator](#) to estimate workload costs.

Implementation steps

- **List the workload components:** Build the list of all the workload components. This is used as verification to check that each component was analyzed. The effort spent should reflect the criticality to the workload as defined by your organization's priorities. Grouping together resources functionally improves efficiency, for example production database storage, if there are multiple databases.
- **Prioritize component list:** Take the component list and prioritize it in order of effort. This is typically in order of the cost of the component from most expensive to least expensive, or the criticality as defined by your organization's priorities.
- **Perform the analysis:** For each component on the list, review the options and services available and chose the option that aligns best with your organizational priorities.

Resources

Related documents:

- [AWS Pricing Calculator](#)
- [AWS Cost Explorer](#)
- [Amazon S3 storage classes](#)
- [Cloud products](#)

COST05-BP03 Perform a thorough analysis of each component

Look at overall cost to the organization of each component. Look at total cost of ownership by factoring in cost of operations and management, especially when using managed services. The review effort should reflect potential benefit, for example, time spent analyzing is proportional to component cost.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Consider the time savings that will allow your team to focus on retiring technical debt, innovation, and value-adding features. For example, you might need to lift and shift your on-premises environment to the cloud as rapidly as possible and optimize later. It is worth exploring the savings you could realize by using managed services that remove or reduce license costs. Managed services remove the operational and administrative burden of maintaining a service, which allows you to focus on innovation. Additionally, because managed services operate at cloud scale, they can offer a lower cost per transaction or service.

Usually, managed services have attributes that you can set to ensure sufficient capacity. You must set and monitor these attributes so that your excess capacity is kept to a minimum and performance is maximized. You can modify the attributes of AWS Managed Services using the AWS Management Console or AWS APIs and SDKs to align resource needs with changing demand. For example, you can increase or decrease the number of nodes on an Amazon EMR cluster (or an Amazon Redshift cluster) to scale out or in.

You can also pack multiple instances on an AWS resource to enable higher density usage. For example, you can provision multiple small databases on a single Amazon Relational Database Service (Amazon RDS) database instance. As usage grows, you can migrate one of the databases to a dedicated Amazon RDS database instance using a snapshot and restore process.

When provisioning workloads on managed services, you must understand the requirements of adjusting the service capacity. These requirements are typically time, effort, and any impact to normal workload operation. The provisioned resource must allow time for any changes to occur, provision the required overhead to allow this. The ongoing effort required to modify services can be reduced to virtually zero by using APIs and SDKs that are integrated with system and monitoring tools, such as Amazon CloudWatch.

[Amazon RDS](#), [Amazon Redshift](#), and [Amazon ElastiCache](#) provide a managed database service. [Amazon Athena](#), [Amazon EMR](#), and [Amazon OpenSearch Service](#) provide a managed analytics service.

[AMS](#) is a service that operates AWS infrastructure on behalf of enterprise customers and partners. It provides a secure and compliant environment that you can deploy your workloads onto. AMS uses enterprise cloud operating models with automation to allow you to meet your organization requirements, move into the cloud faster, and reduce your on-going management costs.

Implementation steps

- **Perform a thorough analysis:** Using the component list, work through each component from the highest priority to the lowest priority. For the higher priority and more costly components, perform additional analysis and assess all available options and their long term impact. For lower priority

components, assess if changes in usage would change the priority of the component, and then perform an analysis of appropriate effort.

Resources

Related documents:

- [AWS Total Cost of Ownership \(TCO\) Calculator](#)
- [Amazon S3 storage classes](#)
- [Cloud products](#)

COST05-BP04 Select software with cost-effective licensing

Open-source software eliminates software licensing costs, which can contribute significant costs to workloads. Where licensed software is required, avoid licenses bound to arbitrary attributes such as CPUs, look for licenses that are bound to output or outcomes. The cost of these licenses scales more closely to the benefit they provide.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

The cost of software licenses can be eliminated through the use of open-source software. This can have significant impact on workload costs as the size of the workload scales. Measure the benefits of licensed software against the total cost to ensure that you have the most optimized workload. Model any changes in licensing and how they would impact your workload costs. If a vendor changes the cost of your database license, investigate how that impacts the overall efficiency of your workload. Consider historical pricing announcements from your vendors for trends of licensing changes across their products. Licensing costs may also scale independently of throughput or usage, such as licenses that scale by hardware (CPU-bound licenses). These licenses should be avoided because costs can rapidly increase without corresponding outcomes.

Implementation steps

- **Analyze license options:** Review the licensing terms of available software. Look for open-source versions that have the required functionality, and whether the benefits of licensed software outweigh the cost. Favorable terms will align the cost of the software to the benefit it provides.
- **Analyze the software provider:** Review any historical pricing or licensing changes from the vendor. Look for any changes that do not align to outcomes, such as punitive terms for running on specific vendors hardware or platforms. Additionally look for how they execute audits, and penalties that could be imposed.

Resources

Related documents:

- [AWS Total Cost of Ownership \(TCO\) Calculator](#)
- [Amazon S3 storage classes](#)
- [Cloud products](#)

COST05-BP05 Select components of this workload to optimize cost in line with organization priorities

Factor in cost when selecting all components. This includes using application level and managed services, such as Amazon Relational Database Service ([Amazon RDS](#)), [Amazon DynamoDB](#), Amazon Simple

Notification Service ([Amazon SNS](#)), and Amazon Simple Email Service ([Amazon SES](#)) to reduce overall organization cost. Use serverless and containers for compute, such as AWS Lambda, Amazon Simple Storage Service ([Amazon S3](#)) for static websites, and Amazon Elastic Container Service ([Amazon ECS](#)). Minimize license costs by using open source software, or software that does not have license fees: for example, Amazon Linux for compute workloads or migrate databases to [Amazon Aurora](#).

Level of risk exposed if this best practice is not established: Low

Implementation guidance

You can use serverless or application-level services such as [AWS Lambda](#), [Amazon Simple Queue Service \(Amazon SQS\)](#), [Amazon SNS](#), and [Amazon SES](#). These services remove the need for you to manage a resource, and provide the function of code execution, queuing services, and message delivery. The other benefit is that they scale in performance and cost in line with usage, allowing efficient cost allocation and attribution.

For more information on Serverless, refer to the [Well-Architected Serverless Application Lens](#) [whitepaper](#).

Implementation steps

- **Select each service to optimize cost:** Using your prioritized list and analysis, select each option that provides the best match with your organizational priorities.

Resources

Related documents:

- [AWS Total Cost of Ownership \(TCO\) Calculator](#)
- [Amazon S3 storage classes](#)
- [Cloud products](#)

COST05-BP06 Perform cost analysis for different usage over time

Workloads can change over time. Some services or features are more cost effective at different usage levels. By performing the analysis on each component over time and at projected usage, the workload remains cost-effective over its lifetime.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

As AWS releases new services and features, the optimal services for your workload may change. Effort required should reflect potential benefits. Workload review frequency depends on your organization requirements. If it is a workload of significant cost, implementing new services sooner will maximize cost savings, so more frequent review can be advantageous. Another trigger for review is change in usage patterns. Significant changes in usage can indicate that alternate services would be more optimal. For example, for higher data transfer rates a direct connect service may be cheaper than a VPN, and provide the required connectivity. Predict the potential impact of service changes, so you can monitor for these usage level triggers and implement the most cost-effective services sooner.

Implementation steps

- **Define predicted usage patterns:** Working with your organization, such as marketing and product owners, document what the expected and predicted usage patterns will be for the workload.
- **Perform cost analysis at predicted usage:** Using the usage patterns defined, perform the analysis at each of these points. The analysis effort should reflect the potential outcome. For example, if the change in usage is large, a thorough analysis should be performed to verify any costs and changes.

Resources

Related documents:

- [AWS Total Cost of Ownership \(TCO\) Calculator](#)
- [Amazon S3 storage classes](#)
- [Cloud products](#)

COST 6 How do you meet cost targets when you select resource type, size and number?

Ensure that you choose the appropriate resource size and number of resources for the task at hand. You minimize waste by selecting the most cost effective type, size, and number.

Best practices

- [COST06-BP01 Perform cost modeling \(p. 378\)](#)
- [COST06-BP02 Select resource type, size, and number based on data \(p. 379\)](#)
- [COST06-BP03 Select resource type, size, and number automatically based on metrics \(p. 379\)](#)

COST06-BP01 Perform cost modeling

Identify organization requirements and perform cost modeling of the workload and each of its components. Perform benchmark activities for the workload under different predicted loads and compare the costs. The modeling effort should reflect the potential benefit. For example, time spent is proportional to component cost.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Perform cost modeling for your workload and each of its components to understand the balance between resources, and find the correct size for each resource in the workload, given a specific level of performance. Perform benchmark activities for the workload under different predicted loads and compare the costs. The modelling effort should reflect potential benefit; for example, time spent is proportional to component cost or predicted saving. For best practices, refer to the [Review](#) section of the [Performance Efficiency Pillar](#) [whitepaper](#).

[AWS Compute Optimizer](#) can assist with cost modelling for running workloads. It provides right-sizing recommendations for compute resources based on historical usage. This is the ideal data source for compute resources because it is a free service, and it utilizes machine learning to make multiple recommendations depending on levels of risk. You can also use [Amazon CloudWatch](#) and [Amazon CloudWatch Logs](#) with custom logs as data sources for right sizing operations for other services and workload components.

The following are recommendations for cost modelling data and metrics:

- The monitoring must accurately reflect the end-user experience. Select the correct granularity for the time period and thoughtfully choose the maximum or 99th percentile instead of the average.
- Select the correct granularity for the time period of analysis that is required to cover any workload cycles. For example, if a two-week analysis is performed, you might be overlooking a monthly cycle of high utilization, which could lead to under-provisioning.

Implementation steps

- **Perform cost modeling:** Deploy the workload or a proof-of-concept, into a separate account with the specific resource types and sizes to test. Run the workload with the test data and record the output results, along with the cost data for the time the test was run. Then redeploy the workload or change the resource types and sizes and re-run the test.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [Amazon CloudWatch features](#)
- [Cost Optimization: Amazon EC2 Right Sizing](#)
- [AWS Compute Optimizer](#)

COST06-BP02 Select resource type, size, and number based on data

Select resource size or type based on data about the workload and resource characteristics. For example, compute, memory, throughput, or write intensive. This selection is typically made using a previous (on-premises) version of the workload, using documentation, or using other sources of information about the workload.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Select resource size or type based on workload and resource characteristics, for example, compute, memory, throughput, or write intensive. This selection is typically made using cost modelling, a previous version of the workload (such as an on-premises version), using documentation, or using other sources of information about the workload (whitepapers, published solutions).

Implementation steps

- **Select resources based on data:** Using your cost modeling data, select the expected workload usage level, then select the specified resource type and size.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [Amazon CloudWatch features](#)
- [Cost Optimization: EC2 Right Sizing](#)

COST06-BP03 Select resource type, size, and number automatically based on metrics

Use metrics from the currently running workload to select the right size and type to optimize for cost. Appropriately provision throughput, sizing, and storage for services such as Amazon Elastic Compute Cloud (Amazon EC2), Amazon DynamoDB, Amazon Elastic Block Store (Amazon EBS) (PIOPS), Amazon Relational Database Service (Amazon RDS), Amazon EMR, and networking. This can be done with a feedback loop such as automatic scaling or by custom code in the workload.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Create a feedback loop within the workload that uses active metrics from the running workload to make changes to that workload. You can use a managed service, such as [AWS Auto Scaling](#), which you configure to perform the right sizing operations for you. AWS also provides [APIs](#), [SDKs](#), and features that allow resources to be modified with minimal effort. You can program a workload to stop-and-start an Amazon Elastic Compute Cloud(Amazon EC2) instance to allow a change of instance size or instance type. This provides the benefits of right-sizing while removing almost all the operational cost required to make the change.

Some AWS services have built in automatic type or size selection, such as [Amazon Simple Storage Service\(Amazon S3\) Intelligent-Tiering](#). Amazon S3 Intelligent-Tiering automatically moves your data between two access tiers: frequent access and infrequent access, based on your usage patterns.

Implementation steps

- **Configure workload metrics:** Ensure you capture the key metrics for the workload. These metrics provide an indication of the customer experience, such as the workload output, and align to the differences between resource types and sizes, such as CPU and memory usage.
- **View rightsizing recommendations:** Use the rightsizing recommendations in AWS Compute Optimizer to make adjustments to your workload.
- **Select resource type and size automatically based on metrics:** Using the workload metrics, manually or automatically select your workload resources. Configuring AWS Auto Scaling or implementing code within your application can reduce the effort required if frequent changes are needed, and it can potentially implement changes sooner than a manual process.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [AWS Compute Optimizer](#)
- [Amazon CloudWatch features](#)
- [CloudWatch Getting Set Up](#)
- [CloudWatch Publishing Custom Metrics](#)
- [Cost Optimization: Amazon EC2 Right Sizing](#)
- [Getting Started with Amazon EC2 Auto Scaling](#)
- [Amazon S3 Intelligent-Tiering](#)
- [Launch an EC2 Instance Using the SDK](#)

COST 7 How do you use pricing models to reduce cost?

Use the pricing model that is most appropriate for your resources to minimize expense.

Best practices

- [COST07-BP01 Perform pricing model analysis \(p. 381\)](#)
- [COST07-BP02 Implement Regions based on cost \(p. 381\)](#)
- [COST07-BP03 Select third-party agreements with cost-efficient terms \(p. 382\)](#)
- [COST07-BP04 Implement pricing models for all components of this workload \(p. 383\)](#)
- [COST07-BP05 Perform pricing model analysis at the master account level \(p. 384\)](#)

COST07-BP01 Perform pricing model analysis

Analyze each component of the workload. Determine if the component and resources will be running for extended periods (for commitment discounts), or dynamic and short-running (for Spot or On-Demand Instances). Perform an analysis on the workload using the Recommendations feature in AWS Cost Explorer.

Level of risk exposed if this best practice is not established: High

Implementation guidance

AWS has multiple [pricing models](#) that allow you to pay for your resources in the most cost-effective way that suits your organization's needs.

Implementation steps

- **Perform a commitment discount analysis:** Using Cost Explorer in your account, review the Savings Plans and Reserved Instance recommendations. To verify that you implement the correct recommendations with the required discounts and risk, follow the [Well-Architected labs](#).
- **Analyze workload elasticity:** Using the hourly granularity in Cost Explorer, or a custom dashboard. Analyze the workload elasticity. Look for regular changes in the number of instances that are running. Short duration instances are candidates for Spot Instances or Spot Fleet.
 - [Well-Architected Lab: Cost Explorer](#)
 - [Well-Architected Lab: Cost Visualization](#)

Resources

Related documents:

- [Accessing Reserved Instance recommendations](#)
- [Instance purchasing options](#)

Related videos:

- [Save up to 90% and run production workloads on Spot](#)

Related examples:

- [Well-Architected Lab: Cost Explorer](#)
- [Well-Architected Lab: Cost Visualization](#)
- [Well-Architected Lab: Pricing Models](#)

COST07-BP02 Implement Regions based on cost

Resource pricing can be different in each Region. Factoring in Region cost helps ensure that you pay the lowest overall price for this workload.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

When you architect your solutions, a best practice is to seek to place computing resources closer to users to provide lower latency and strong data sovereignty. For global audiences, you should use multiple locations to meet these needs. You should select the geographic location that minimizes your costs.

The AWS Cloud infrastructure is built around [Regions and Availability Zones](#). A Region is a physical location in the world where we have multiple Availability Zones. Availability Zones consist of one or more discrete data centers, each with redundant power, networking, and connectivity, housed in separate facilities.

Each AWS Region operates within local market conditions, and resource pricing is different in each Region. Choose a specific Region to operate a component of or your entire solution so that you can run at the lowest possible price globally. You can use the [AWS Pricing Calculator](#) to estimate the costs of your workload in various Regions.

Implementation steps

- **Review Region pricing:** Analyze the workload costs in the current Region. Starting with the highest costs by service and usage type, calculate the costs in other Regions that are available. If the forecasted saving outweighs the cost of moving the component or workload, migrate to the new Region.

Resources

Related documents:

- [Accessing Reserved Instance recommendations](#)
- [Amazon EC2 pricing](#)
- [Instance purchasing options](#)
- [Region Table](#)

Related videos:

- [Save up to 90% and run production workloads on Spot](#)

COST07-BP03 Select third-party agreements with cost-efficient terms

Cost efficient agreements and terms ensure the cost of these services scales with the benefits they provide. Select agreements and pricing that scale when they provide additional benefits to your organization.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

When you utilize third-party solutions or services in the cloud, it is important that the pricing structures are aligned to Cost Optimization outcomes. Pricing should scale with the outcomes and value it provides. An example of this is software that takes a percentage of savings it provides, the more you save (outcome) the more it charges. Agreements that scale with your bill are typically not aligned to Cost Optimization, unless they provide outcomes for every part of your specific bill. For example, a solution that provides recommendations for Amazon Elastic Compute Cloud(Amazon EC2) and charges a percentage of your entire bill will increase if you use other services for which it provides no benefit. Another example is a managed service that is charged at a percentage of the cost of resources that are managed. A larger instance size may not necessarily require more management effort, but will be charged more. Ensure that these service pricing arrangements include a cost optimization program or features in their service to drive efficiency.

Implementation steps

- **Analyze third-party agreements and terms:** Review the pricing in third party agreements. Perform modeling for different levels of your usage, and factor in new costs such as new service usage, or

increases in current services due to workload growth. Decide if the additional costs provide the required benefits to your business.

Resources

Related documents:

- [Accessing Reserved Instance recommendations](#)
- [Instance purchasing options](#)

Related videos:

- [Save up to 90% and run production workloads on Spot](#)

COST07-BP04 Implement pricing models for all components of this workload

Permanently running resources should utilize reserved capacity such as Savings Plans or Reserved Instances. Short-term capacity is configured to use Spot Instances, or Spot Fleet. On-Demand Instances are only used for short-term workloads that cannot be interrupted and do not run long enough for reserved capacity, between 25% to 75% of the period, depending on the resource type.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Consider the requirements of the workload components and understand the potential pricing models. Define the availability requirement of the component. Determine if there are multiple independent resources that perform the function in the workload, and what the workload requirements are over time. Compare the cost of the resources using the default On-Demand pricing model and other applicable models. Factor in any potential changes in resources or workload components.

Implementation steps

- **Implement pricing models:** Using your analysis results, purchase Savings Plans (SPs), Reserved Instances (RIs) or implement Spot Instances. If it is your first RI purchase then choose the top 5 or 10 recommendations in the list, then monitor and analyze the results over the next month or two. Purchase small numbers of commitment discounts regular cycles, for example every two weeks or monthly. Implement Spot Instances for workloads that can be interrupted or are stateless.
- **Workload review cycle:** Implement a review cycle for the workload that specifically analyzes pricing model coverage. Once the workload has the required coverage, purchase additional commitment discounts every two to four weeks, or as your organization usage changes.

Resources

Related documents:

- [Accessing Reserved Instance recommendations](#)
- [EC2 Fleet](#)
- [How to Purchase Reserved Instances](#)
- [Instance purchasing options](#)
- [Spot Instances](#)

Related videos:

- [Save up to 90% and run production workloads on Spot](#)

COST07-BP05 Perform pricing model analysis at the master account level

Use Cost Explorer Savings Plans and Reserved Instance recommendations to perform regular analysis at the management account level for commitment discounts.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Performing regular cost modeling ensures that opportunities to optimize across multiple workloads can be implemented. For example, if multiple workloads use On-Demand Instances, at an aggregate level, the risk of change is lower, and implementing a commitment-based discount will achieve a lower overall cost. It is recommended to perform analysis in regular cycles of two weeks to one month. This allows you to make small adjustment purchases, so the coverage of your pricing models continues to evolve with your changing workloads and their components.

Use the [AWS Cost Explorer](#) recommendations tool to find opportunities for commitment discounts.

To find opportunities for Spot workloads, use an hourly view of your overall usage, and look for regular periods of changing usage or elasticity.

Implementation steps

- **Perform a commitment discount analysis:** Using Cost Explorer in your account review the Savings Plans and Reserved Instance recommendations. To verify you implement the correct recommendations with the required discounts and risk, follow the Well-Architected labs.

Resources

Related documents:

- [Accessing Reserved Instance recommendations](#)
- [Instance purchasing options](#)

Related videos:

- [Save up to 90% and run production workloads on Spot](#)

Related examples:

- [Well-Architected Lab: Pricing Models](#)

COST 8 How do you plan for data transfer charges?

Ensure that you plan and monitor data transfer charges so that you can make architectural decisions to minimize costs. A small yet effective architectural change can drastically reduce your operational costs over time.

Best practices

- [COST08-BP01 Perform data transfer modeling \(p. 385\)](#)
- [COST08-BP02 Select components to optimize data transfer cost \(p. 385\)](#)
- [COST08-BP03 Implement services to reduce data transfer costs \(p. 386\)](#)

COST08-BP01 Perform data transfer modeling

Gather organization requirements and perform data transfer modeling of the workload and each of its components. This identifies the lowest cost point for its current data transfer requirements.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Understand where the data transfer occurs in your workload, the cost of the transfer, and its associated benefit. This allows you to make an informed decision to modify or accept the architectural decision. For example, you may have a Multi-Availability Zone configuration where you replicate data between the Availability Zones. You model the cost of structure and decide that this is an acceptable cost (similar to paying for compute and storage in both Availability Zone) to achieve the required reliability and resilience.

Model the costs over different usage levels. Workload usage can change over time, and different services may be more cost effective at different levels.

Use [AWS Cost Explorer](#) or the [AWS Cost and Usage Report](#) (CUR) to understand and model your data transfer costs. Configure a proof of concept (PoC) or test your workload, and run a test with a realistic simulated load. You can model your costs at different workload demands.

Implementation steps

- **Calculate data transfer costs:** Use the [AWS pricing pages](#) and calculate the data transfer costs for the workload. Calculate the data transfer costs at different usage levels, for both increases and reductions in workload usage. Where there are multiple options for the workload architecture, calculate the cost for each option for comparison.
- **Link costs to outcomes:** For each data transfer cost incurred, specify the outcome that it achieves for the workload. If it is transfer between components, it may be for decoupling, if it is between Availability Zones it may be for redundancy.

Resources

Related documents:

- [AWS caching solutions](#)
- [AWS Pricing](#)
- [Amazon EC2 Pricing](#)
- [Amazon VPC pricing](#)
- [Deliver content faster with Amazon CloudFront](#)

COST08-BP02 Select components to optimize data transfer cost

All components are selected, and architecture is designed to reduce data transfer costs. This includes using components such as wide-area-network (WAN) optimization and Multi-Availability Zone (AZ) configurations

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Architecting for data transfer ensures that you minimize data transfer costs. This may involve using content delivery networks to locate data closer to users, or using dedicated network links from your premises to AWS. You can also use WAN optimization and application optimization to reduce the amount of data that is transferred between components.

Implementation steps

- **Select components for data transfer:** Using the data transfer modeling, focus on where the largest data transfer costs are or where they would be if the workload usage changes. Look for alternative architectures, or additional components that remove or reduce the need for data transfer, or lower its cost.

Resources

Related documents:

- [AWS caching solutions](#)
- [Deliver content faster with Amazon CloudFront](#)

COST08-BP03 Implement services to reduce data transfer costs

Implement services to reduce data transfer. For example, using a content delivery network (CDN) such as Amazon CloudFront to deliver content to end users, caching layers using Amazon ElastiCache, or using AWS Direct Connect instead of VPN for connectivity to AWS.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

[Amazon CloudFront](#) is a global content delivery network that delivers data with low latency and high transfer speeds. It caches data at edge locations across the world, which reduces the load on your resources. By using CloudFront, you can reduce the administrative effort in delivering content to large numbers of users globally, with minimum latency.

[AWS Direct Connect](#) allows you to establish a dedicated network connection to AWS. This can reduce network costs, increase bandwidth, and provide a more consistent network experience than internet-based connections.

[AWS VPN](#) allows you to establish a secure and private connection between your private network and the AWS global network. It is ideal for small offices or business partners because it provides quick and easy connectivity, and it is a fully managed and elastic service.

[VPC Endpoints](#) allow connectivity between AWS services over private networking and can be used to reduce public data transfer and [NAT gateways](#) costs. [Gateway VPC endpoints](#) have no hourly charges, and support Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB. [Interface VPC endpoints](#) are provided by [AWS PrivateLink](#) and have an hourly fee and per GB usage cost.

Implementation steps

- **Implement services:** Using the data transfer modeling, look at where the largest costs and highest volume flows are. Review the AWS services and assess whether there is a service that reduces or removes the transfer, specifically networking and content delivery. Also look for caching services where there is repeated access to data, or large amounts of data.

Resources

Related documents:

- [AWS Direct Connect](#)
- [AWS Explore Our Products](#)
- [AWS caching solutions](#)
- [Amazon CloudFront](#)

- [Deliver content faster with Amazon CloudFront](#)

Manage demand and supply resources

Question

- [COST 9 How do you manage demand, and supply resources? \(p. 387\)](#)

COST 9 How do you manage demand, and supply resources?

For a workload that has balanced spend and performance, ensure that everything you pay for is used and avoid significantly underutilizing instances. A skewed utilization metric in either direction has an adverse impact on your organization, in either operational costs (degraded performance due to over-utilization), or wasted AWS expenditures (due to over-provisioning).

Best practices

- [COST09-BP01 Perform an analysis on the workload demand \(p. 387\)](#)
- [COST09-BP02 Implement a buffer or throttle to manage demand \(p. 388\)](#)
- [COST09-BP03 Supply resources dynamically \(p. 389\)](#)

COST09-BP01 Perform an analysis on the workload demand

Analyze the demand of the workload over time. Verify that the analysis covers seasonal trends and accurately represents operating conditions over the full workload lifetime. Analysis effort should reflect the potential benefit, for example, time spent is proportional to the workload cost.

Level of risk exposed if this best practice is not established: High

Implementation guidance

Know the requirements of the workload. The organization requirements should indicate the workload response times for requests. The response time can be used to determine if the demand is managed, or if the supply of resources will change to meet the demand.

The analysis should include the predictability and repeatability of the demand, the rate of change in demand, and the amount of change in demand. Ensure that the analysis is performed over a long enough period to incorporate any seasonal variance, such as end-of-month processing or holiday peaks.

Ensure that the analysis effort reflects the potential benefits of implementing scaling. Look at the expected total cost of the component, and any increases or decreases in usage and cost over the workload lifetime.

You can use [AWS Cost Explorer](#) or [Amazon QuickSight](#) with the AWS Cost and Usage Report (CUR) or your application logs to perform a visual analysis of workload demand.

Implementation steps

- **Analyze existing workload data:** Analyze data from the existing workload, previous versions of the workload, or predicted usage patterns. Use log files and monitoring data to gain insight on how customers use the workload. Typical metrics are the actual demand in requests per second, the times when the rate of demand changes or when it is at different levels, and the rate of change of demand. Ensure you analyze a full cycle of the workload, ensuring you collect data for any seasonal changes such as end of month or end of year events. The effort reflected in the analysis should reflect the workload characteristics. The largest effort should be placed on high-value workloads that have the largest changes in demand. The least effort should be placed on low-value workloads that

have minimal changes in demand. Common metrics for value are risk, brand awareness, revenue or workload cost.

- **Forecast outside influence:** Meet with team members from across the organization that can influence or change the demand in the workload. Common teams would be sales, marketing, or business development. Work with them to know the cycles they operate within, and if there are any events that would change the demand of the workload. Forecast the workload demand with this data.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [AWS Instance Scheduler](#)
- [Getting started with Amazon SQS](#)
- [AWS Cost Explorer](#)
- [Amazon QuickSight](#)

COST09-BP02 Implement a buffer or throttle to manage demand

Buffering and throttling modify the demand on your workload, smoothing out any peaks. Implement throttling when your clients perform retries. Implement buffering to store the request and defer processing until a later time. Verify that your throttles and buffers are designed so clients receive a response in the required time.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Throttling: If the source of the demand has retry capability, then you can implement throttling. Throttling tells the source that if it cannot service the request at the current time it should try again later. The source will wait for a period of time and then re-try the request. Implementing throttling has the advantage of limiting the maximum amount of resources and costs of the workload. In AWS, you can use [Amazon API Gateway](#) to implement throttling. Refer to the [Well-Architected Reliability pillar whitepaper](#) for more details on implementing throttling.

Buffer based: Similar to throttling, a buffer defers request processing, allowing applications that run at different rates to communicate effectively. A buffer-based approach uses a queue to accept messages (units of work) from producers. Messages are read by consumers and processed, allowing the messages to run at the rate that meets the consumers' business requirements. You don't have to worry about producers having to deal with throttling issues, such as data durability and backpressure (where producers slow down because their consumer is running slowly).

In AWS, you can choose from multiple services to implement a buffering approach. [Amazon Simple Queue Service \(Amazon SQS\)](#) is a managed service that provides queues that allow a single consumer to read individual messages. [Amazon Kinesis](#) provides a stream that allows many consumers to read the same messages.

When architecting with a buffer-based approach, ensure that you architect your workload to service the request in the required time, and that you are able to handle duplicate requests for work.

Implementation steps

- **Analyze the client requirements:** Analyze the client requests to determine if they are capable of performing retries. For clients that cannot perform retries, buffers will need to be implemented. Analyze the overall demand, rate of change, and required response time to determine the size of throttle or buffer required.

- **Implement a buffer or throttle:** Implement a buffer or throttle in the workload. A queue such as Amazon Simple Queue Service (Amazon SQS) can provide a buffer to your workload components. Amazon API Gateway can provide throttling for your workload components.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [AWS Instance Scheduler](#)
- [Amazon API Gateway](#)
- [Amazon Simple Queue Service](#)
- [Getting started with Amazon SQS](#)
- [Amazon Kinesis](#)

COST09-BP03 Supply resources dynamically

Resources are provisioned in a planned manner. This can be demand-based, such as through automatic scaling, or time-based, where demand is predictable and resources are provided based on time. These methods result in the least amount of over or under-provisioning.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

You can use [AWS Auto Scaling](#), or incorporate scaling in your code with the [AWS API or SDKs](#). This reduces your overall workload costs by removing the operational cost from manually making changes to your environment, and can be performed much faster. This will ensure that the workload resourcing best matches the demand at any time.

Demand-based supply: Leverage the elasticity of the cloud to supply resources to meet changing demand. Take advantage of APIs or service features to programmatically vary the amount of cloud resources in your architecture dynamically. This allows you to scale components in your architecture, and automatically increase the number of resources during demand spikes to maintain performance, and decrease capacity when demand subsides to reduce costs.

[AWS Auto Scaling](#) helps you adjust your capacity to maintain steady, predictable performance at the lowest possible cost. It is a fully managed and free service that integrates with Amazon Elastic Compute Cloud (Amazon EC2) instances and Spot Fleets, Amazon Elastic Container Service (Amazon ECS), Amazon DynamoDB, and Amazon Aurora.

Auto Scaling provides automatic resource discovery to help find resources in your workload that can be configured, it has built-in scaling strategies to optimize performance, costs or a balance between the two, and provides predictive scaling to assist with regularly occurring spikes.

Auto Scaling can implement manual, scheduled or demand-based scaling. You can also use metrics and alarms from [Amazon CloudWatch](#) to trigger scaling events for your workload. Typical metrics can be standard Amazon EC2 metrics, such as CPU utilization, network throughput, and [Elastic Load Balancing\(ELB\)](#) observed request or response latency. When possible, you should use a metric that is indicative of customer experience, which is typically a custom metric that might originate from application code within your workload.

When architecting with a demand-based approach keep in mind two key considerations. First, understand how quickly you must provision new resources. Second, understand that the size of margin between supply and demand will shift. You must be ready to cope with the rate of change in demand and also be ready for resource failures.

[ELB](#) helps you to scale by distributing demand across multiple resources. As you implement more resources, you add them to the load balancer to take on the demand. Elastic Load Balancing has support for Amazon EC2 Instances, containers, IP addresses, and AWS Lambda functions.

Time-based supply: A time-based approach aligns resource capacity to demand that is predictable or well-defined by time. This approach is typically not dependent upon utilization levels of the resources. A time-based approach ensures that resources are available at the specific time they are required, and can be provided without any delays due to start-up procedures and system or consistency checks. Using a time-based approach, you can provide additional resources or increase capacity during busy periods.

You can use scheduled Auto Scaling to implement a time-based approach. Workloads can be scheduled to scale out or in at defined times (for example, the start of business hours) thus ensuring that resources are available when users or demand arrives.

You can also leverage the [AWS APIs and SDKs](#) and [AWS CloudFormation](#) to automatically provision and decommission entire environments as you need them. This approach is well suited for development or test environments that run only in defined business hours or periods of time.

You can use APIs to scale the size of resources within an environment (vertical scaling). For example, you could scale up a production workload by changing the instance size or class. This can be achieved by stopping and starting the instance and selecting the different instance size or class. This technique can also be applied to other resources, such as Amazon Elastic Block Store (Amazon EBS) Elastic Volumes, which can be modified to increase size, adjust performance (IOPS) or change the volume type while in use.

When architecting with a time-based approach keep in mind two key considerations. First, how consistent is the usage pattern? Second, what is the impact if the pattern changes? You can increase the accuracy of predictions by monitoring your workloads and by using business intelligence. If you see significant changes in the usage pattern, you can adjust the times to ensure that coverage is provided.

Implementation steps

- **Configure time-based scheduling:** For predictable changes in demand, time-based scaling can provide the correct number of resources in a timely manner. It is also useful if resource creation and configuration is not fast enough to respond to changes on demand. Using the workload analysis configure scheduled scaling using AWS Auto Scaling.
- **Configure Auto Scaling:** To configure scaling based on active workload metrics, use Amazon Auto Scaling. Use the analysis and configure auto scaling to trigger on the correct resource levels, and ensure that the workload scales in the required time.

Resources

Related documents:

- [AWS Auto Scaling](#)
- [AWS Instance Scheduler](#)
- [Getting Started with Amazon EC2 Auto Scaling](#)
- [Getting started with Amazon SQS](#)
- [Scheduled Scaling for Amazon EC2 Auto Scaling](#)

Optimize over time

Question

- [COST 10 How do you evaluate new services? \(p. 391\)](#)

COST 10 How do you evaluate new services?

As AWS releases new services and features, it's a best practice to review your existing architectural decisions to ensure they continue to be the most cost effective.

Best practices

- [COST10-BP01 Develop a workload review process \(p. 391\)](#)
- [COST10-BP02 Review and analyze this workload regularly \(p. 392\)](#)

COST10-BP01 Develop a workload review process

Develop a process that defines the criteria and process for workload review. The review effort should reflect potential benefit. For example, core workloads or workloads with a value of over 10% of the bill are reviewed quarterly, while workloads below 10% are reviewed annually.

Level of risk exposed if this best practice is not established: High

Implementation guidance

To ensure that you always have the most cost-efficient workload, you must regularly review the workload to know if there are opportunities to implement new services, features, and components. To ensure that you achieve overall lower costs the process must be proportional to the potential amount of savings. For example, workloads that are 50% of your overall spend should be reviewed more regularly, and more thoroughly, than workloads that are 5% of your overall spend. Factor in any external factors or volatility. If the workload services a specific geography or market segment, and change in that area is predicted, more frequent reviews could lead to cost savings. Another factor in review is the effort to implement changes. If there are significant costs in testing and validating changes, reviews should be less frequent.

Factor in the long-term cost of maintaining outdated and legacy, components and resources, and the inability to implement new features into them. The current cost of testing and validation may exceed the proposed benefit. However, over time, the cost of making the change may significantly increase as the gap between the workload and the current technologies increases, resulting in even larger costs. For example, the cost of moving to a new programming language may not currently be cost effective. However, in five years time, the cost of people skilled in that language may increase, and due to workload growth, you would be moving an even larger system to the new language, requiring even more effort than previously.

Break down your workload into components, assign the cost of the component (an estimate is sufficient), and then list the factors (for example, effort and external markets) next to each component. Use these indicators to determine a review frequency for each workload. For example, you may have web servers as a high cost, low change effort, and high external factors, resulting in high frequency of review. A central database may be medium cost, high change effort, and low external factors, resulting in a medium frequency of review.

Implementation steps

- **Define review frequency:** Define how frequently the workload and its components should be reviewed. This is a combination of factors and may differ from workload to workload within your organization, it may also differ between components in the workload. Common factors include the importance to the organization measured in terms of revenue or brand, the total cost of running the workload (including operation and resource costs), the complexity of the workload, how easy is it to implement a change, any software licensing agreements, and if a change would incur significant increases in licensing costs due to punitive licensing. Components can be defined functionally or technically, such as web servers and databases, or compute and storage resources. Balance the factors accordingly and develop a period for the workload and its components. You may decide to review the full workload every 18 months, review the web servers every 6 months, the database every 12 months, compute and short-term storage every 6 months, and long-term storage every 12 months.

- **Define review thoroughness:** Define how much effort is spent on the review of the workload or workload components. Similar to the review frequency, this is a balance of multiple factors. You may decide to spend one week of analysis on the database component, and four hours for storage reviews.

Resources

Related documents:

- [AWS News Blog](#)
- [Types of Cloud Computing](#)
- [What's New with AWS](#)

COST10-BP02 Review and analyze this workload regularly

Existing workloads are regularly reviewed based on for each defined processes.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

To realize the benefits of new AWS services and features, you must execute the review process on your workloads and implement new services and features as required. For example, you might review your workloads and replace the messaging component with Amazon Simple Email Service (Amazon SES). This removes the cost of operating and maintaining a fleet of instances, while providing all the functionality at a reduced cost.

Implementation steps

- **Regularly review the workload:** Using your defined process, perform reviews with the frequency specified. Verify that you spend the correct amount of effort on each component. This process would be similar to the initial design process where you selected services for cost optimization. Analyze the services and the benefits they would bring, this time factor in the cost of making the change, not just the long-term benefits.
- **Implement new services:** If the outcome of the analysis is to implement changes, first perform a baseline of the workload to know the current cost for each output. Implement the changes, then perform an analysis to confirm the new cost for each output.

Resources

Related documents:

- [AWS News Blog](#)
- [Types of Cloud Computing](#)
- [What's New with AWS](#)

Sustainability

Topics

- [Region selection \(p. 393\)](#)
- [User behavior patterns \(p. 393\)](#)
- [Software and architecture patterns \(p. 398\)](#)
- [Data patterns \(p. 402\)](#)
- [Hardware patterns \(p. 407\)](#)

- [Development and deployment process \(p. 410\)](#)

Region selection

Question

- [SUS 1 How do you select Regions to support your sustainability goals? \(p. 393\)](#)

SUS 1 How do you select Regions to support your sustainability goals?

Choose Regions where you will implement your workloads based on both your business requirements and sustainability goals.

Best Practice:

SUS01-BP01 Choose Regions near Amazon renewable energy projects and Regions where the grid has a published carbon intensity that is lower than other locations (or Regions)

Choose Regions near Amazon renewable energy projects and Regions where the grid has a published carbon intensity that is lower than other locations (or Regions).

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

Choose Regions near Amazon renewable energy projects and Regions where the grid has a published carbon intensity that is lower than other locations (or Regions).

Resources

Related documents:

- [Amazon Around the Globe](#)
- [Renewable Energy Methodology](#)
- [What to Consider when Selecting a Region for your Workloads](#)

User behavior patterns

Question

- [SUS 2 How do you take advantage of user behavior patterns to support your sustainability goals? \(p. 393\)](#)

SUS 2 How do you take advantage of user behavior patterns to support your sustainability goals?

The way users consume your workloads and other resources can help you identify improvements to meet sustainability goals. Scale infrastructure to continually match user load and ensure that only the minimum resources required to support users are deployed. Align service levels to customer needs. Position resources to limit the network required for users to consume them. Remove existing, unused assets. Identify created assets that are unused and stop generating them. Provide your team members with devices that support their needs with minimized sustainability impact.

Best practices:

SUS02-BP01 Scale infrastructure with user load

Identify periods of low or no utilization and scale down resources to eliminate excess capacity and improve efficiency.

Common anti-patterns:

- You do not scale your infrastructure with user load.
- You manually scale your infrastructure all the time.
- You leave increased capacity after a scaling event instead of scaling back down.

Benefits of establishing this best practice: Configuring and testing workload elasticity will help reduce workload environmental impact, save money, and maintain performance benchmarks. You can take advantage of elasticity in the cloud to automatically scale capacity during and after user load spikes to make sure you are only using the exact number of resources needed to meet the needs of your customers.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Elasticity matches the supply of resources you have against the demand for those resources. Instances, containers, and functions provide mechanisms for elasticity, either in combination with automatic scaling or as a feature of the service. Use elasticity in your architecture to ensure that workload can scale down quickly and easily during the period of low user load:

Auto-scaling mechanism	Where to use
Amazon EC2 Auto Scaling	Use to verify you have the correct number of Amazon EC2 instances available to handle the user load for your application.
Application Auto Scaling	Use to automatically scale the resources for individual AWS services beyond Amazon EC2, such as Lambda functions or Amazon Elastic Container Service (Amazon ECS) services.
Kubernetes Cluster Autoscaler	Use to automatically scale Kubernetes clusters on AWS.

- Verify that the metrics for scaling up or down are validated against the type of workload being deployed. If you are deploying a video transcoding application, 100% CPU utilization is expected and should not be your primary metric. You can use a [customized metric](#) (such as memory utilization) for your scaling policy if required. To choose the right metrics, consider the following guidance for Amazon EC2:
 - The metric should be a valid utilization metric and describe how busy an instance is.
 - The metric value must increase or decrease proportionally to the number of instances in the Auto Scaling group.
- Use [dynamic scaling](#) instead of [manual scaling](#) for your Auto Scaling group. We also recommend that you use [target tracking scaling policies](#) in your dynamic scaling.
- Verify that workload deployments can handle both scale-up and scale-down events. Create test scenarios for scale-down events to ensure that the workload behaves as expected. You can use [Activity history](#) to test and verify a scaling activity for an Auto Scaling group.

- Evaluate your workload for predictable patterns and proactively scale as you anticipate predicted and planned changes in demand. Use [Predictive Scaling with Amazon EC2 Auto Scaling](#) to eliminate the need to overprovision capacity.

Resources

Related documents:

- [Getting Started with Amazon EC2 Auto Scaling](#)
- [Predictive Scaling for EC2, Powered by Machine Learning](#)
- [Analyze user behavior using Amazon OpenSearch Service, Amazon Kinesis Data Firehose and Kibana](#)
- [What is Amazon CloudWatch?](#)
- [What is AWS X-Ray?](#)
- [VPC Flow Logs](#)
- [Monitoring DB load with Performance Insights on Amazon RDS](#)
- [Introducing Native Support for Predictive Scaling with Amazon EC2 Auto Scaling](#)
- [How to create an Amazon EC2 Auto Scaling policy based on a memory utilization metric \(Linux\)](#)
- [Introducing Karpenter - An Open-Source, High-Performance Kubernetes Cluster Autoscaler](#)

Related videos:

- [Better, faster, cheaper compute: Cost-optimizing Amazon EC2 \(CMP202-R1\)](#)

Related examples:

- [Lab: Amazon EC2 Auto Scaling Group Examples](#)
- [Lab: Implement Autoscaling with Karpenter](#)

SUS02-BP02 Align SLAs with sustainability goals

Define and update Service Level Agreements (SLAs) such as availability or data retention periods to minimize the number of resources required to support your workload while continuing to meet business requirements.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Define SLAs that support your sustainability goals while meeting your business requirements.
- Redefine SLAs to meet business requirements, not exceed them.
- Make trade-offs that significantly reduce sustainability impacts in exchange for acceptable decreases in service levels.
- Use design patterns that prioritize business-critical functions, and allow lower service levels (such as response time or recovery time objectives) for non-critical functions.

Resources

Related documents:

- [AWS Service Level Agreements \(SLAs\)](#)
- [Importance of Service Level Agreement for SaaS Providers](#)

Related videos:

- [Building Sustainably on AWS](#)

SUS02-BP03 Stop the creation and maintenance of unused assets

Analyze application assets (such as pre-compiled reports, datasets, and static images) and asset access patterns to identify redundancy, underutilization, and potential decommission targets. Consolidate generated assets with redundant content (for example, monthly reports with overlapping or common datasets and outputs) to remove the resources consumed when duplicating outputs. Decommission unused assets (for example, images of products that are no longer sold) to free consumed resources and reduce the number of resources used to support the workload.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Manage static assets and remove assets that are no longer required.
- Manage generated assets and stop generating and remove assets that are no longer required.
- Consolidate overlapping generated assets to remove redundant processing.
- Instruct third parties to stop producing and storing assets managed on your behalf that are no longer required.
- Instruct third parties to consolidate redundant assets produced on your behalf.

Resources

Related documents:

- [Optimizing your AWS Infrastructure for Sustainability, Part II: Storage](#)

Related videos:

- [Building Sustainably on AWS](#)

SUS02-BP04 Optimize geographic placement of workloads for user locations

Analyze network access patterns to identify where your customers are connecting from geographically. Select Regions and services that reduce the distance network traffic must travel to decrease the total network resources required to support your workload.

Common anti-patterns:

- You select the workload's Region based on your own location.

Benefits of establishing this best practice: Placing a workload close to its customers provides the lowest latency while decreasing data movement across the network and lowering environmental impact.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Select the Regions for your workload deployment based on the following key elements:
 - **Your Sustainability goal:** as explained in [Region selection](#).

- **Where your data is located:** For data-heavy applications (such as big data and machine learning), application code should execute as close to the data as possible.
- **Where your users are located:** For user-facing applications, choose a Region close to your workload's customer base.
- **Other constraints:** Consider constraints such as security and compliance as explained in [What to Consider when Selecting a Region for your Workloads](#).
- Use [AWS Local Zones](#) to run workloads like video rendering and graphics-intensive virtual desktop applications. Local Zones allow you to benefit from having compute and storage resources closer to end users.
- Use local caching or [AWS Caching Solutions](#) for frequently used resources to improve performance, reduce data movement, and lower environmental impact.

Service	When to use
Amazon CloudFront	Use to cache static content such as images, scripts, and videos, as well as dynamic content such as API responses or web applications.
Amazon ElastiCache	Use to cache content for web applications.
DynamoDB Accelerator	Use to add in-memory acceleration to your DynamoDB tables.

- Use services that can help you run code closer to users of your workload:

Service	When to use
Lambda@Edge	Use for compute-heavy operations that are executed when objects are not in the cache.
Amazon CloudFront Functions	Use for simple use cases like HTTP(s) request or response manipulations that can be executed by short-lived functions.
AWS IoT Greengrass	Use to run local compute, messaging, and data caching for connected devices.

- Use connection pooling to enable connection reuse, and reduce required resources.
- Use distributed data stores that don't rely on persistent connections and synchronous updates for consistency to serve regional populations.
- Replace pre-provisioned static network capacity with shared dynamic capacity, and share the sustainability impact of network capacity with other subscribers.

Resources

Related documents:

- [Optimizing your AWS Infrastructure for Sustainability, Part III: Networking](#)
- [Amazon ElastiCache Documentation](#)
- [What is Amazon CloudFront?](#)
- [Amazon CloudFront Key Features](#)
- [Lambda@Edge](#)
- [CloudFront Functions](#)
- [AWS IoT Greengrass](#)

Related videos:

- [Building Sustainably on AWS](#)

Related examples:

- [AWS Networking Workshops](#)

SUS02-BP05 Optimize team member resources for activities performed

Optimize resources provided to team members to minimize the sustainability impact while supporting their needs. For example, perform complex operations, such as rendering and compilation, on highly utilized shared cloud desktops instead of on underutilized high-powered single-user systems.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Provision workstations and other devices to align with how they're used.
- Use virtual desktops and application streaming to limit upgrade and device requirements.
- Move processor or memory-intensive tasks to the cloud.
- Evaluate the impact of processes and systems on your device lifecycle, and select solutions that minimize the requirement for device replacement while satisfying business requirements.
- Implement remote management for devices to reduce required business travel.

Resources

Related documents:

- [What is Amazon WorkSpaces?](#)
- [Amazon AppStream 2.0 Documentation](#)
- [NICE DCV](#)
- [AWS Systems Manager Fleet Manager](#)

Related videos:

- [Building Sustainably on AWS](#)

Software and architecture patterns

Question

- [SUS 3 How do you take advantage of software and architecture patterns to support your sustainability goals? \(p. 398\)](#)

SUS 3 How do you take advantage of software and architecture patterns to support your sustainability goals?

Implement patterns for performing load smoothing and maintaining consistent high utilization of deployed resources to minimize the resources consumed. Components might become idle from lack of use because of changes in user behavior over time. Revise patterns and architecture to consolidate

under-utilized components to increase overall utilization. Retire components that are no longer required. Understand the performance of your workload components, and optimize the components that consume the most resources. Be aware of the devices your customers use to access your services, and implement patterns to minimize the need for device upgrades.

Best practices:

SUS03-BP01 Optimize software and architecture for asynchronous and scheduled jobs

Use efficient software designs and architectures to minimize the average resources required per unit of work. Implement mechanisms that result in even utilization of components to reduce resources that are idle between tasks and minimize the impact of load spikes.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Queue requests that don't require immediate processing.
- Increase serialization to flatten utilization across your pipeline.
- Modify the capacity of individual components to prevent idling resources waiting for input.
- Create buffers and establish rate limiting to smooth the consumption of external services.
- Use the most efficient available hardware for your software optimizations.
- Use queue-driven architectures, pipeline management, and On-Demand Instance workers to maximize utilization for batch processing.
- Schedule tasks to avoid load spikes and resource contention from simultaneous execution.
- Schedule jobs during times of day where carbon intensity for power is lowest.

Resources

Related documents:

- [What is Amazon Simple Queue Service?](#)
- [What is Amazon MQ?](#)
- [Scaling based on Amazon SQS](#)
- [What is AWS Step Functions?](#)
- [What is AWS Lambda?](#)
- [Using AWS Lambda with Amazon SQS](#)
- [What is Amazon EventBridge?](#)

Related videos:

- [Building Sustainably on AWS](#)
- [Moving to event-driven architectures](#)

SUS03-BP02 Remove or refactor workload components with low or no use

Monitor workload activity to identify changes in utilization of individual components over time. Remove components that are unused and no longer required, and refactor components with little utilization to limit wasted resources.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Analyze load (using indicators such as transaction flow and API calls) on functional components to identify unused and underutilized components.
- Retire components that are no longer needed.
- Refactor underutilized components.
- Consolidate underutilized components with other resources to improve utilization efficiency.

Resources

Related documents:

- [What is AWS X-Ray?](#)
- [What is Amazon CloudWatch?](#)
- [Using ServiceLens to monitor the health of your applications](#)
- [Automated Cleanup of Unused Images in Amazon ECR](#)

Related videos:

- [Building Sustainably on AWS](#)

SUS03-BP03 Optimize areas of code that consume the most time or resources

Monitor workload activity to identify application components that consume the most resources. Optimize the code that runs within these components to minimize resource usage while maximizing performance.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Monitor performance as a function of resource usage to identify components with high resource requirements per unit of work as targets for optimization.
- Use a code profiler to identify the areas of code that use the most time or resources as targets for optimization.
- Replace algorithms with more efficient versions that produce the same result.
- Use hardware acceleration to improve the efficiency of blocks of code with long execution times.
- Use the most efficient operating system and programming language for the workload.
- Remove unnecessary sorting and formatting.
- Use data transfer patterns that minimize the resources used based on how frequently the data changes and how it is consumed. For example, push state change information to a client instead of having it consume resources to poll and receive valueless 'no change' messages.

Resources

Related documents:

- [What is Amazon CloudWatch?](#)
- [What is Amazon CodeGuru Profiler?](#)

- [FPGA instances](#)
- [The AWS SDKs on Tools to Build on AWS](#)

Related videos:

- [Building Sustainably on AWS](#)

SUS03-BP04 Optimize impact on customer devices and equipment

Understand the devices and equipment your customers use to consume your services, their expected lifecycle, and the financial and sustainability impact of replacing those components. Implement software patterns and architectures to minimize the need for customers to replace devices and upgrade equipment. For example, implement new features using code that is backward compatible with older hardware and operating system versions, or manage the size of payloads so they don't exceed the storage capacity of the target device.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Inventory the devices your customers use.
- Test using managed device farms with representative sets of hardware to understand the impact of your changes, and iterate development to maximize the devices supported.
- Account for network bandwidth and latency when building payloads, and implement capabilities that help your applications work well on low-bandwidth, high-latency links.
- Pre-process data payloads to reduce local processing requirements and limit data transfer requirements.
- Perform computationally intense activities server-side (such as image rendering), or use application streaming to improve the user experience on older devices.
- Segment and paginate output, especially for interactive sessions, to manage payloads and limit local storage requirements.

Resources

Related documents:

- [What is AWS Device Farm?](#)
- [Amazon AppStream 2.0 Documentation](#)
- [NICE DCV](#)
- [Amazon Elastic Transcoder Documentation](#)

Related videos:

- [Building Sustainably on AWS](#)

SUS03-BP05 Use software patterns and architectures that best support data access and storage patterns

Understand how data is used within your workload, consumed by your users, transferred, and stored. Select technologies to minimize data processing and storage requirements.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Analyze your data access and storage patterns.
- Store data files in efficient file formats such as Parquet to prevent unnecessary processing (for example, when running analytics) and to reduce the total storage provisioned.
- Use technologies that work natively with compressed data.
- Use the database engine that best supports your dominant query pattern.
- Manage your database indexes to ensure index designs support efficient query execution.
- Select network protocols that reduce the amount of network capacity consumed.

Resources

Related documents:

- [Athena Compression Support file formats](#)
- [COPY from columnar data formats with Amazon Redshift](#)
- [Converting Your Input Record Format in Kinesis Data Firehose](#)
- [Format Options for ETL Inputs and Outputs in AWS Glue](#)
- [Improve query performance on Amazon Athena by Converting to Columnar Formats](#)
- [Loading compressed data files from Amazon S3 with Amazon Redshift](#)
- [Monitoring DB load with Performance Insights on Amazon Aurora](#)
- [Monitoring DB load with Performance Insights on Amazon RDS](#)
- [AWS IoT FleetWise](#)

Related videos:

- [Building Sustainably on AWS](#)

Data patterns

Question

- SUS 4 How do you take advantage of data access and usage patterns to support your sustainability goals? (p. 402)

SUS 4 How do you take advantage of data access and usage patterns to support your sustainability goals?

Implement data management practices to reduce the provisioned storage required to support your workload, and the resources required to use it. Understand your data, and use storage technologies and configurations that best support the business value of the data and how it's used. Lifecycle data to more efficient, less performant storage when requirements decrease, and delete data that's no longer required.

Best practices:

SUS04-BP01 Implement a data classification policy

Classify data to understand its significance to business outcomes. Use this information to determine when you can move data to more energy-efficient storage or safely delete it.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Determine requirements for the distribution, retention, and deletion of your data.
- Use tagging on volumes and objects to record the metadata that's used to determine how it's managed, including data classification.
- Periodically audit your environment for untagged and unclassified data, and classify and tag the data appropriately.

Resources

Related documents:

- [Data Classification Process](#)
- [Leveraging AWS Cloud to Support Data Classification](#)
- [Tag policies from AWS Organizations](#)

SUS04-BP02 Use technologies that support data access and storage patterns

Use storage that best supports how your data is accessed and stored to minimize the resources provisioned while supporting your workload. For example, Solid State Devices (SSDs) are more energy intensive than magnetic drives and should be used only for active data use cases. Use energy-efficient, archival-class storage for infrequently accessed data.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Monitor your data access patterns.
- Migrate data to the appropriate technology based on access pattern.
- Migrate archival data to storage designed for that purpose.

Resources

Related documents:

- [Amazon EBS volume types](#)
- [Amazon EC2 instance store](#)
- [Amazon S3 Intelligent-Tiering](#)
- [Using Amazon S3 storage classes](#)
- [What is Amazon CloudWatch?](#)
- [What is Amazon S3 Glacier?](#)

Related videos:

- [Architectural Patterns for Data Lakes on AWS](#)

SUS04-BP03 Use lifecycle policies to delete unnecessary data

Manage the lifecycle of all your data and automatically enforce deletion timelines to minimize the total storage requirements of your workload.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Define lifecycle policies for all your data classification types.
- Set automated lifecycle policies to enforce lifecycle rules.
- Delete unused volumes and snapshots.
- Aggregate data where applicable based on lifecycle rules.

Resources

Related documents:

- [Amazon ECR Lifecycle policies](#)
- [Amazon EFS lifecycle management](#)
- [Amazon S3 Intelligent-Tiering](#)
- [Evaluating Resources with AWS Config Rules](#)
- [Managing your storage lifecycle on Amazon S3](#)
- [Object lifecycle policies in AWS Elemental MediaStore](#)

Related videos:

- [Amazon S3 Lifecycle](#)

SUS04-BP04 Minimize over-provisioning in block storage

To minimize total provisioned storage, create block storage with size allocations that are appropriate for the workload. Use elastic volumes to expand storage as data grows without having to resize storage attached to compute resources. Regularly review elastic volumes and shrink over-provisioned volumes to fit the current data size.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Monitor the utilization of your data volumes.
- Use elastic volumes and managed block data services to automate allocation of additional storage as your persistent data grows.
- Set target levels of utilization for your data volumes, and resize volumes outside of expected ranges.
- Size read-only volumes to fit the data.
- Migrate data to object stores to avoid provisioning the excess capacity from fixed volume sizes on block storage.

Resources

Related documents:

- [Amazon EBS Elastic Volumes](#)
- [Amazon FSx Documentation](#)
- [What is Amazon CloudWatch?](#)

- [What is Amazon Elastic File System?](#)

SUS04-BP05 Remove unneeded or redundant data

Duplicate data only when necessary to minimize total storage consumed. Use backup technologies that deduplicate data at the file and block level. Limit the use of Redundant Array of Independent Drives (RAID) configurations except where required to meet Service Level Agreements (SLAs).

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Use mechanisms that can deduplicate data at the block and object level.
- Use backup technology that can make incremental backups and deduplicate data at the block, file, and object level.
- Use RAID only when required to meet your SLAs.
- Centralize log and trace data, deduplicate identical log entries, and establish mechanisms to tune verbosity when needed.
- Pre-populate caches only where justified.
- Establish cache monitoring and automation to resize cache accordingly.
- Remove out-of-date deployments and assets from object stores and edge caches when pushing new versions of your workload.

Resources

Related documents:

- [Amazon EBS snapshots](#)
- [Change log data retention in CloudWatch Logs](#)
- [Data deduplication on Amazon FSx for Windows File Server](#)
- [Features of Amazon FSx for ONTAP including data deduplication](#)
- [Invalidating Files on Amazon CloudFront](#)
- [Using AWS Backup to back up and restore Amazon EFS file systems](#)
- [What is Amazon CloudWatch Logs?](#)
- [Working with backups on Amazon RDS](#)

Related examples:

- [Lab: Optimize Data Pattern Using Amazon Redshift Data Sharing](#)

SUS04-BP06 Use shared file systems or object storage to access common data

Adopt shared storage and single sources of truth to avoid data duplication and reduce the total storage requirements of your workload. Fetch data from shared storage only as needed. Detach unused volumes to make more resources available.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Migrate data to shared storage when the data has multiple consumers.
- Fetch data from shared storage only as needed.

- Delete data as appropriate for your usage patterns, and implement time-to-live (TTL) functionality to manage cached data.
- Detach volumes from clients that are not actively using them.

Resources

Related documents:

- [Amazon FSx](#)
- [Caching strategies](#)
- [What is Amazon Elastic File System?](#)
- [What is Amazon S3?](#)

SUS04-BP07 Minimize data movement across networks

Use shared storage and access data from regional data stores to minimize the total networking resources required to support data movement for your workload.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Store data as close to the consumer as possible.
- Partition regionally consumed services so that their Region-specific data is stored within the Region where it is consumed.
- Use block-level duplication instead of file or object-level duplication when copying changes across the network.
- Compress data before moving it over the network.

Resources

Related documents:

- [Optimizing your AWS Infrastructure for Sustainability, Part III: Networking](#)
- [AWS Global Infrastructure](#)
- [Amazon CloudFront Key Features including the CloudFront Global Edge Network](#)
- [Compressing HTTP requests in Amazon OpenSearch Service](#)
- [Intermediate data compression with Amazon EMR](#)
- [Loading compressed data files from Amazon S3 into Amazon Redshift](#)
- [Serving compressed files with Amazon CloudFront](#)

SUS04-BP08 Back up data only when difficult to recreate

To minimize storage consumption, only back up data that has business value or is needed to satisfy compliance requirements. Examine backup policies and exclude ephemeral storage that doesn't provide value in a recovery scenario.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Use your data classification to establish what data needs to be backed up.

- Exclude data that you can easily recreate.
- Exclude ephemeral data from your backups.
- Exclude local copies of data, unless the time required to restore that data from a common location exceeds your service level agreements (SLAs).

Resources

Related documents:

- [Using AWS Backup to back up and restore Amazon EFS file systems](#)
- [Amazon EBS snapshots](#)
- [Working with backups on Amazon Relational Database Service](#)

Hardware patterns

Question

- [SUS 5 How do your hardware management and usage practices support your sustainability goals? \(p. 407\)](#)

SUS 5 How do your hardware management and usage practices support your sustainability goals?

Look for opportunities to reduce workload sustainability impacts by making changes to your hardware management practices. Minimize the amount of hardware needed to provision and deploy, and select the most efficient hardware for your individual workload.

Best practices:

SUS05-BP01 Use the minimum amount of hardware to meet your needs

Using the capabilities of the cloud, you can make frequent changes to your workload implementations. Update deployed components as your needs change.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Enable horizontal scaling, and use automation to scale out as loads increase and to scale in as loads decrease.
- Scale using small increments for variable workloads.
- Align scaling with cyclical utilization patterns (for example, a payroll system with intense bi-weekly processing activities) as load varies over days, weeks, months, or years.
- Negotiate service level Agreements (SLAs) that allow for a temporary reduction in capacity while automation deploys replacement resources.

Resources

Related documents:

- [AWS Compute Optimizer Documentation](#)
- [Operating Lambda: Performance optimization](#)

- [Auto Scaling Documentation](#)

SUS05-BP02 Use instance types with the least impact

Continually monitor the release of new instance types and take advantage of energy efficiency improvements, including those instance types designed to support specific workloads such as machine learning training, inference, and video transcoding.

Common anti-patterns:

- You are only using one family of instances.
- You are only using x86 instances.
- You specify one instance type in your Amazon EC2 Auto Scaling configuration.
- You use AWS instances in a manner that they were not designed for (for example, you use compute-optimized instances for a memory-intensive workload).
- You do not evaluate new instance types regularly.
- You do not check recommendations from AWS rightsizing tools such as [AWS Compute Optimizer](#).

Benefits of establishing this best practice: By using energy-efficient and right-sized instances, you are able to greatly reduce the environmental impact and cost of your workload.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Learn and explore instance types which can lower your workload environmental impact.
 - Subscribe to [What's New with AWS](#) to be up-to-date with the latest AWS technologies and instances.
 - Learn about different AWS instance types.
 - Learn about AWS Graviton-based instances which offer the best performance per watt of energy use in Amazon EC2 by watching [re:Invent 2020 - Deep dive on AWS Graviton2 processor-powered Amazon EC2 instances](#) and [Deep dive into AWS Graviton3 and Amazon EC2 C7g instances](#).
- Plan and transition your workload to instance types with the least impact.
 - Define a process to evaluate new features or instances for your workload. Take advantage of agility in the cloud to quickly test how new instance types can improve your workload environmental sustainability. Use proxy metrics to measure how many resources it takes you to complete a unit of work.
 - If possible, modify your workload to work with different numbers of vCPUs and different amounts of memory to maximize your choice of instance type.
 - Consider transitioning your workload to Graviton-based instances to improve the performance efficiency of your workload (see [AWS Graviton Fast Start](#) and [AWS Graviton2 for ISVs](#)). Keep in mind the [considerations when transitioning workloads to AWS Graviton-based Amazon Elastic Compute Cloud instances](#).
 - Consider selecting the AWS Graviton option in your usage of [AWS managed services](#).
 - Migrate your workload to Regions that offer instances with the least sustainability impact and still meet your business requirements.
 - For machine learning workloads, use Amazon EC2 instances which are based on custom Amazon Machine Learning chips such as [AWS Trainium](#), [AWS Inferentia](#), and [Amazon EC2 DL1](#).
 - Use [Amazon SageMaker Inference Recommender](#) to right size ML inference endpoint.
 - For workloads with real time video transcoding, use [Amazon EC2 VT1 Instances](#).
 - For spikey workloads (workloads with infrequent requirements for additional capacity), use [burstable performance instances](#).

- For stateless and fault-tolerant workloads, use [Amazon EC2 Spot Instances](#) to increase overall utilization of the cloud, and reduce the sustainability impact of unused resources.
- Operate and optimize your workload instance.
 - For ephemeral workloads, evaluate [instance Amazon CloudWatch metrics](#) such as `CPUUtilization` to identify if the instance is idle or under-utilized.
 - For stable workloads, check AWS rightsizing tools such as [AWS Compute Optimizer](#) at regular intervals to identify opportunities to optimize and right-size the instances.

Resources

Related documents:

- [Optimizing your AWS Infrastructure for Sustainability, Part I: Compute](#)
- [AWS Graviton Processor](#)
- [AWS Inferentia](#)
- [AWS Trainium](#)
- [Amazon EC2 DL1](#)
- [Amazon EC2 Burstable performance instances](#)
- [Amazon EC2 Capacity Reservation Fleets](#)
- [Amazon EC2 Spot Fleet](#)
- [Amazon EC2 Spot Instances](#)
- [Amazon EC2 VT1 Instances](#)
- [Amazon EC2 instance types](#)
- [AWS Compute Optimizer](#)
- [Functions: Lambda Function Configuration](#)

Related videos:

- [Deep dive on AWS Graviton2 processor-powered Amazon EC2 instances](#)
- [Deep dive into AWS Graviton3 and Amazon EC2 C7g instances](#)

Related examples:

- [Lab: Rightsizing Recommendations](#)
- [Lab: Rightsizing with Compute Optimizer](#)
- [Lab: Optimize Hardware Patterns and Observe Sustainability KPIs](#)

SUS05-BP03 Use managed services

Managed services shift responsibility for maintaining high-average utilization, and sustainability optimization of the deployed hardware to AWS. Use managed services to distribute the sustainability impact of the service across all tenants of the service, reducing your individual contribution.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Migrate from self-hosted services to managed services. For example, use managed [Amazon Relational Database Service \(Amazon RDS\)](#) instances instead of maintaining your own Amazon RDS instances on [Amazon Elastic Compute Cloud \(Amazon EC2\)](#), or use managed container services, such as [AWS Fargate](#), instead of implementing your own container infrastructure.

Resources

Related documents:

- [AWS Fargate](#)
- [Amazon DocumentDB](#)
- [Amazon Elastic Kubernetes Service \(EKS\)](#)
- [Amazon Managed Streaming for Apache Kafka \(Amazon MSK\)](#)
- [Amazon Redshift](#)
- [Amazon Relational Database Service \(RDS\)](#)

SUS05-BP04 Optimize your use of GPUs

Graphics Processing Units (GPUs) can be a source of high-power consumption, and many GPU workloads are highly variable, such as rendering, transcoding, and machine learning training and modeling. Only run GPU instances for the time needed, and decommission them with automation when not required to minimize resources consumed.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Use GPUs only for tasks where they're more efficient than CPU-based alternatives.
- Use automation to release GPU instances when not in use.
- Use flexible graphics acceleration rather than dedicated GPU instances.
- Take advantage of custom-purpose hardware that is specific to your workload.

Resources

Related documents:

- [Accelerated Computing](#)
- [AWS Inferentia](#)
- [AWS Trainium](#)
- [Accelerated Computing for EC2 Instances](#)
- [Amazon EC2 VT1 Instances](#)
- [Amazon Elastic Graphics](#)

Development and deployment process

Question

- [SUS 6 How do your development and deployment processes support your sustainability goals? \(p. 410\)](#)

SUS 6 How do your development and deployment processes support your sustainability goals?

Look for opportunities to reduce your sustainability impact by making changes to your development, test, and deployment practices.

Best practices:

SUS06-BP01 Adopt methods that can rapidly introduce sustainability improvements

Test and validate potential improvements before deploying them to production. Account for the cost of testing when calculating potential future benefit of an improvement. Develop low-cost testing methods to enable delivery of small improvements.

Level of risk exposed if this best practice is not established: Medium

Implementation guidance

- Add requirements for sustainability to your development process.
- Allow resources to work in parallel to develop, test, and deploy sustainability improvements.
- Test and validate potential sustainability impact improvements before deploying into production.
- Test potential improvements using the minimum viable representative components.
- Deploy tested sustainability improvements to production as they become available.

Resources

Related documents:

- [AWS enables sustainability solutions](#)

Related examples:

- [Lab: Turning](#) cost & usage reports into efficiency reports

SUS06-BP02 Keep your workload up-to-date

Up-to-date operating systems, libraries, and applications can improve workload efficiency and enable easier adoption of more efficient technologies. Up-to-date software might also include features to measure the sustainability impact of your workload more accurately, as vendors deliver features to meet their own sustainability goals.

Common anti-patterns:

- You assume your current architecture will become static with no updates over time.
- You do not have any systems or a regular cadence to evaluate if updated software and packages are compatible with your workload.
- You introduce architecture changes over time without justification.

Benefits of establishing this best practice: By establishing a process to keep your workload up to date, you will be able to adopt new features and capabilities, resolve issues, and improve workload efficiency.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Define a process and a schedule to evaluate new features or instances for your workload. Take advantage of agility in the cloud to quickly test how new features can improve your workload to:
 - Reduce sustainability impacts.

- Gain performance efficiencies.
- Remove barriers for a planned improvement.
- Improve your ability to measure and manage sustainability impacts.
- Inventory your workload software and architecture and identify components that need to be updated. You can use [AWS Systems Manager Inventory](#) to collect operating system (OS), application, and instance metadata from your Amazon EC2 instances and quickly understand which instances are running the software and configurations required by your software policy and which instances need to be updated.
- Understand how to update the components of your workload.

Workload component	How to update
Machine images	Manage updates to Amazon Machine Images (AMI) for Linux or Windows server images using EC2 Image Builder .
Container images	You should use Amazon Elastic Container Registry (Amazon ECR) with your existing pipeline to manage Amazon Elastic Container Service (Amazon ECS) images and manage Amazon Elastic Kubernetes Service images .
AWS Lambda	AWS Lambda includes version management features .

- Use automation for the update process to reduce the level of effort to deploy new features and limit errors caused by manual processes. Use tools such as [AWS Systems Manager Patch Manager](#) to automate the process of system updates, and schedule the activity using [AWS Systems Manager Maintenance Windows](#).

Resources

Related documents:

- [AWS Architecture Center](#)
- [What's New with AWS](#)
- [AWS Developer Tools](#)
- [AWS Systems Manager Patch Manager](#)

Related examples:

- [Well-Architected Labs: Inventory and Patch Management](#)
- [Lab: AWS Systems Manager](#)

SUS06-BP03 Increase utilization of build environments

Use automation and infrastructure-as-code to bring pre-production environments up when needed and take them down when not used. A common pattern is to schedule periods of availability that coincide with the working hours of your development team members. Hibernation is a useful tool to preserve the state and rapidly bring instances online only when needed. Use instance types with burst capacity, Spot Instances, elastic database services, containers, and other technologies to align development and test capacity with use.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

- Use automation to maximize utilization of your development and test environments.
- Use automation to manage the lifecycle of your development and test environments.
- Use minimum viable representative environments to develop and test potential improvements.
- Use On-Demand Instances to supplement your developer devices.
- Use automation to maximize the efficiency of your build resources.
- Use instance types with burst capacity, Spot Instances, and other technologies to align build capacity with use.
- Adopt native cloud services for secure instance shell access rather than deploying fleets of bastion hosts.

Resources

Related documents:

- [AWS Systems Manager Session Manager](#)
- [Amazon EC2 Burstable performance instances](#)
- [What is AWS CloudFormation?](#)

SUS06-BP04 Use managed device farms for testing

Managed device farms spread the sustainability impact of hardware manufacturing and resource usage across multiple tenants. Managed device farms offer diverse device types so you can support older, less popular hardware, and avoid customer sustainability impact from unnecessary device upgrades.

Level of risk exposed if this best practice is not established: Low

Implementation guidance

Test using managed device farms with representative sets of hardware to understand the impact of your changes, and iterate development to maximize the devices supported.

Resources

Related documents:

- [What is AWS Device Farm?](#)

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Copyright © 2021 Amazon Web Services, Inc. or its affiliates.