11 Disaster Recovery Technology and Its Applications

www.huawei.com

HUAWEI

Foreword

- This module mainly introduces:
  - The definition of Disaster Recovery system and its importance towards Enterprise Business Continuity.
  - Common Disaster Recovery Solutions.
  - Common Disaster Recovery Technologies.
  - Success Story of Disaster Recovery Solutions.

Objectives

- Upon completion of this module, you will be able to:
  - Describe the concepts and importance of Disaster Recovery.
  - Differentiate the Advantages and Disadvantages of common Disaster Recovery solutions.
  - Understand the technical principles of Disaster Recovery solutions.
  - Learn how to deploy a Disaster Recovery solution through case study of typical application of Disaster Recovery Solution.

Contents

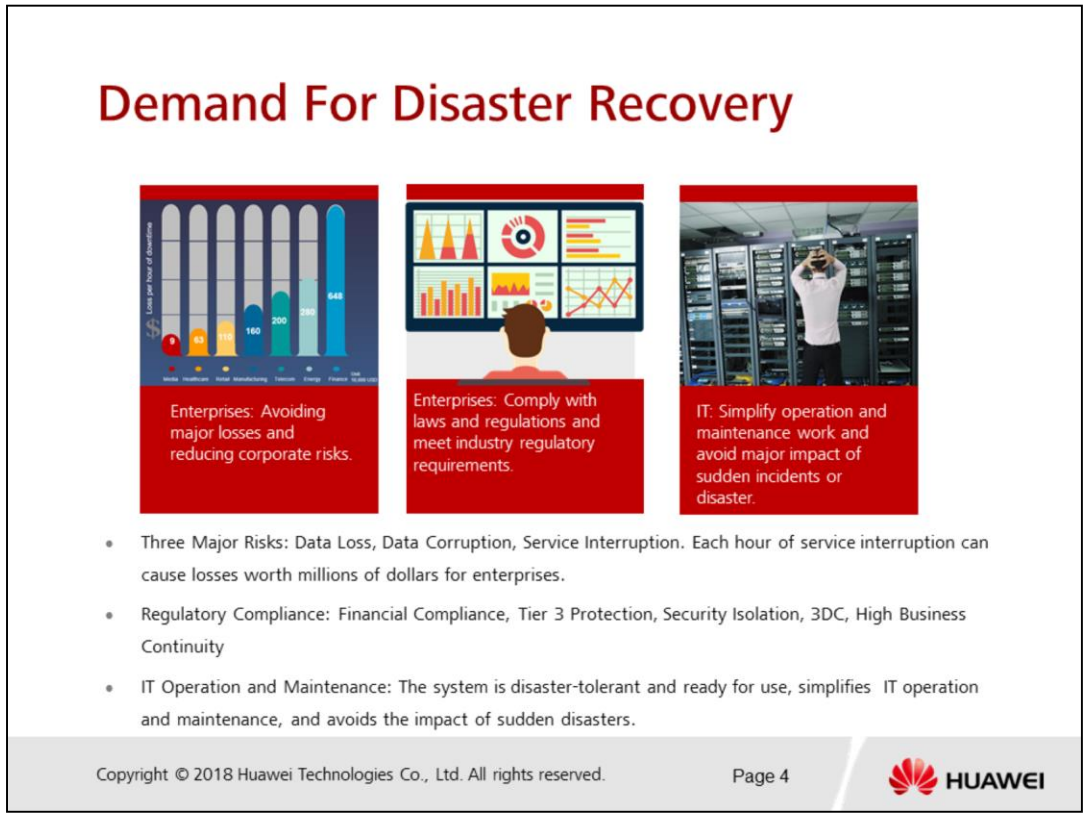

- The necessity of disaster recovery system construction: Disasters are everywhere and can happen at anytime.
- From the analysis of statistics on insurance from a respected organization in Switzerland:
  - Global direct losses due to natural disasters and human accidents reached US$123 billion in 2004/
  - About 400 catastrophes occurred in the world in 2005, with losses exceeding US$230 billion.
  - In 2006, the losses caused by natural disasters and man-made disasters were lower than the long-term trend, but still with the impact of a direct loss of 48 billion U.S. dollars.
  - Globally, compared with the 1960s, the number of catastrophic natural disasters in the world has increased threefold and the economic loss has increased by 9 times up till the 1990s.

- The impact of disasters in recent years in China
    - In the "National Eleventh Five-Year Plan for Comprehensive Disaster Reduction" issued by the General Office of the State Council in August 2007, it has been estimated that in recent years, China's comprehensive direct loss of all kinds of disasters has reached 8% of GDP (most of the developed countries only reached a few thousandths of GDP)
    - The types of natural disasters are various. Floods, typhoons, droughts, etc., bring about 200 billion Yuan in economic losses to the country each year. - National Disaster Reduction Committee
    - Affected by the Wenchuan earthquake, the net loss of the entire banking industry was about RMB 5 billion to RMB 13 billion. Urban commercial banks without disaster recovery centers suffered the greatest impact. In 2008, the bank's net profit fell by 0.7%.
- The huge losses caused by natural disasters with a low probability to occur still cannot be ignored.
    - According to IDC statistics, 55% of the companies that had experienced disasters in the United States in the decades before the year 2000 ended up with closure of business. Of the remaining 45%, due to the loss of data, 29% also closed within two years. The companies that survived are accounted for 16%.
    - Research at the University of Minnesota shows that more than 60% of companies that have disasters but do not have disaster recovery plans will exit the market after two to three years. With the increasing reliance on data processing by enterprises, this proportion or ratio has an upward trend.

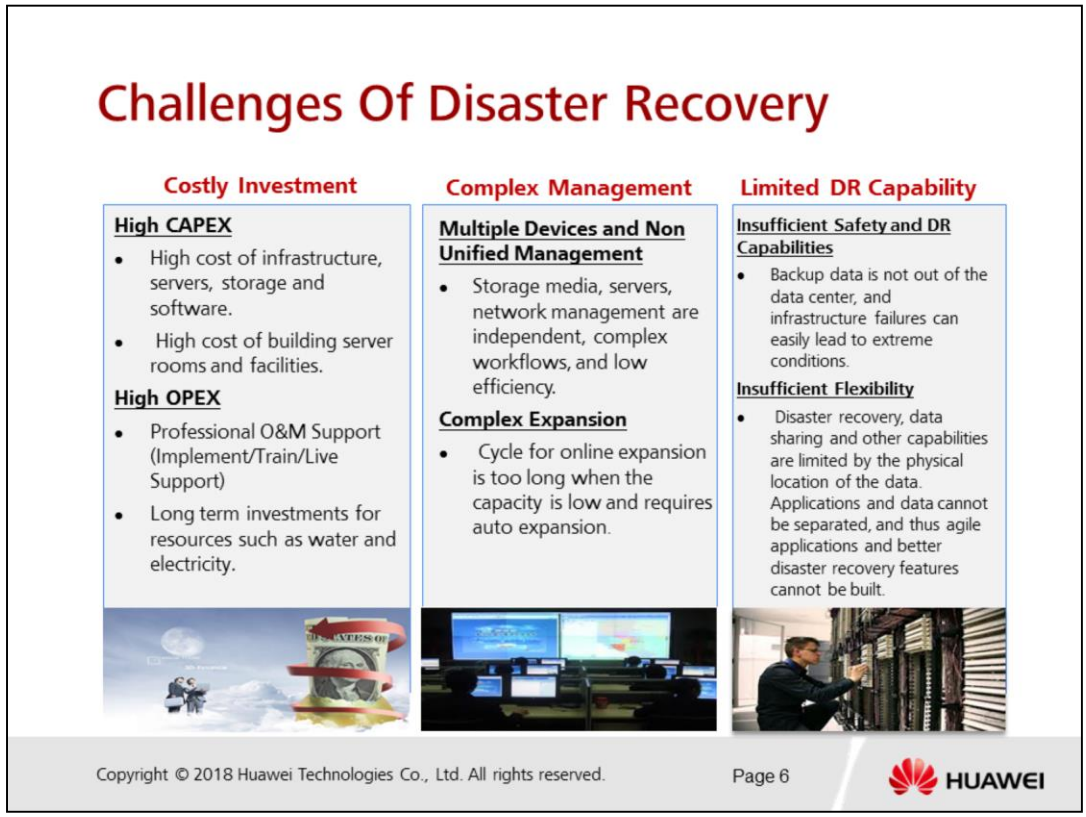

- Various number of applications and inconvenient management: More and more business systems are running in enterprise IT systems, and increasing number of applications are required for disaster recovery protection as they are part of the key business services. Common applications include Oracle, DB2, SQL Server, Exchange, etc. The cloudification trend of IT systems is becoming more obvious. Many virtual machines also need to be protected by the DR plan. Additionally, there is a lack of a unified management system.
- DR process is complicated, time-consuming and error-prone: Different application have different configurations, and their recovery processes are also different making overall DR configuration difficult. On the other hand, service failover and recovery require operations by trained professionals, which is time-consuming and error-prone. It also lacks of automated creation and deployment processes.
- Black box operation is difficult to understand: Traditional service failovers, drills and other operational processes are black box operations that occurs within the equipment, and with the lack of visual effects making it difficult to understand.

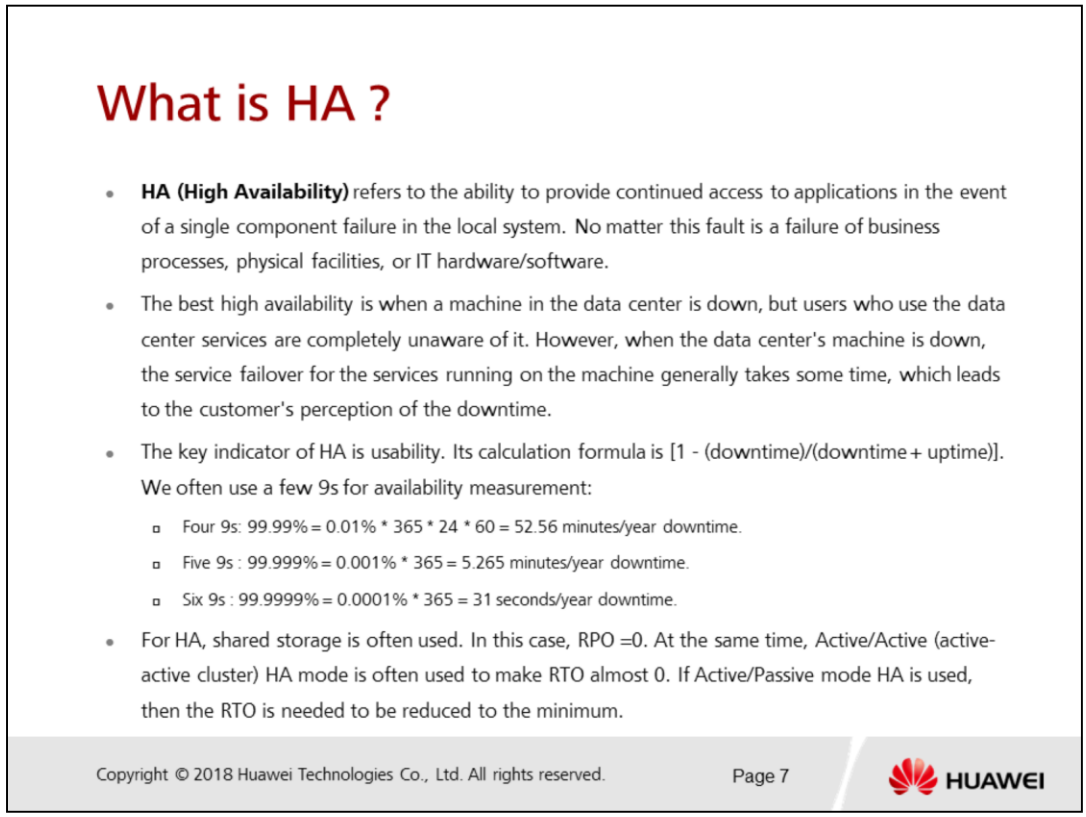

- HA needs to use redundant servers to form a cluster to run the workloads, which includes applications and services. This type of redundancy also divides HA into two categories which are:
- Active/Passive HA:
  - The cluster only includes two nodes, these nodes are referred in short as master and backup. In this configuration, the system uses primary and standby machines to provide services, and the system only provides services on the primary device.
  - When the master device fails, the service on the standby device is started to replace the service provided by the master device.
  - Typically, CRM software such as Pacemaker can be used to control the switching between active and standby devices and provides a virtual IP to provide services.

- Active/Active HA:
    - When a cluster contains only two nodes, it is referred as active-active. When it includes multiple nodes, it becomes Multi-Master.
    - In this configuration, the system runs the same load on all servers in the cluster.
    - Taking a database as an example, updates to an instance are synchronized to all instances.
    - This configuration often uses load balancing software such as HAProxy to provide virtual IP services.
- Pacemaker is a cluster manager. It utilizes the message and member capabilities provided by the preferred clustering infrastructure (OpenAIS or heartbeat). Fault detection and recovery are done by secondary nodes and systems, achieving an implementation of high availability cluster services (also known as resources). CRM: Cluster Resource Management.
- HAProxy is a free and open source software written in C language that provides high availability, load balancing, and application proxy over TCP and HTTP. HAProxy is particularly suitable for those heavily loaded web sites, which often require session maintenance or seven-tier processing.

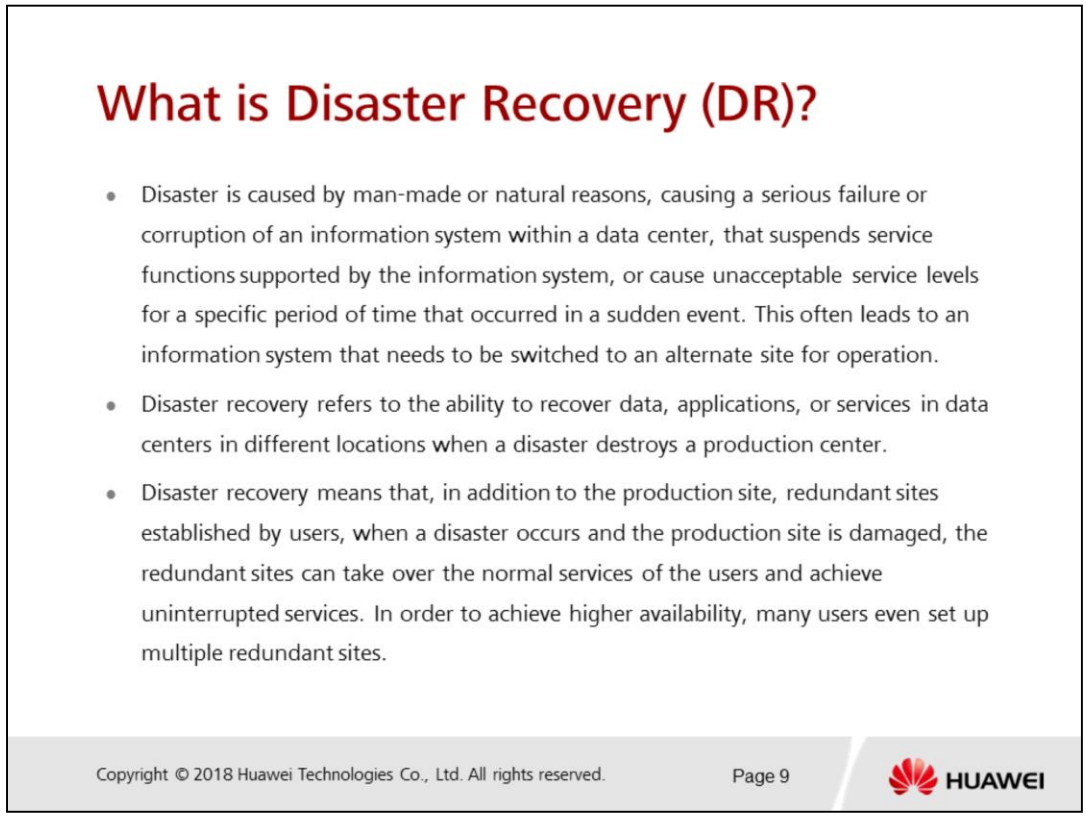

- Disaster recovery (DR) is an area of security planning that aims to protect an organization from the effects of significant negative events. DR allows an organization to maintain or quickly resume mission-critical functions following a disaster.
- A disaster can be anything that puts an organization's operations at risk, from a cyberattack to equipment failures to natural disasters. The goal with DR is for a business to continue operating as close to normal as possible. The disaster recovery process includes planning and testing, and may involve a separate physical site for restoring operations.

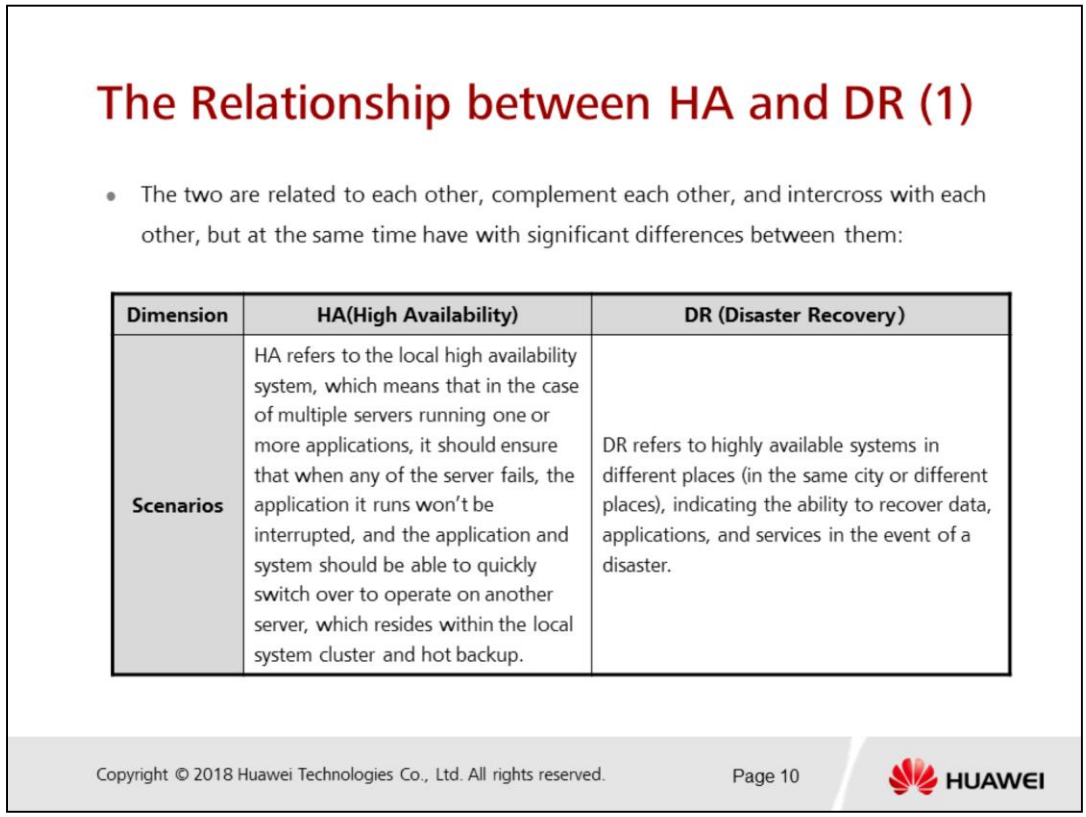

- High Availability(HA) is the process in which a system is supposed to take over when another system stops working efficiently or at all. For IT, this must occur with as minimal downtime as possible—so minimal that most users don't even know there was a problem. Data loss must also be negligible; if there is a service issue, HA takes over so that whatever employees are working on isn't lost. No system is perfect—as each has occasional hiccups, scheduled maintenance, peak usage times, upgrades, minor outages, and so on. High availability accounts for these events and guarantees they don't cut into standard operations and productivity.
- Although disaster recovery also is concerned with downtime and data security, it focuses on bigger threats, surprise crises, and potentially longer outages: beyond the normal recovery point objective (RPO) and recovery time objective (RTO) of high availability. DR deals with the unexpected—events that weren't scheduled, were apart from the norm, and often carry longer-lasting consequences. Good disaster recovery doesn't mean companies are completely caught off guard, but rather, that they are trained and prepared to respond to an event, according to set plans, in order to restore normal operations quickly, whether the systems they turn to are in house, at another facility, or in the cloud (or a combination thereof).

The Relationship between HA and DR (2)

Dimension	HA(High Availability)	DR(Disaster Recovery)
Storage	HA often uses shared storage, which means there is no lost of data (RPO=0), more considerations are on the RTO due to the time period taken for service switching/failover.	Off-site disaster recovery data recovery uses data replication. According to the different data replication technologies used (synchronous and asynchronous), data loss often results in RPO > 0; while off-site application switching often takes longer, so RTO > 0.
Faults	Mainly handles single-component failures causing the load to switch between servers in the cluster.	Copes with large-scale failures causing the load to switch between data centers.
Network	LAN scaling tasks are within the category of HA.	WAN scaling tasks are within the category of DR.
Cloud	HA is a mechanism for ensuring business continuity within a cloud environment.	DR is a mechanism for ensuring business continuity among multiple cloud environments.
Objective	HA's main purpose is to ensure that the business/service is highly available.	DR is a business continuity solution that is based on reliable data protection.

- RPO is the maximum age of files that an organization must recover from backup storage for normal operations to resume after a disaster. The recovery point objective determines the minimum frequency of backups. For example, if an organization has an RPO of four hours, the system must back up at least every four hours.
- RTO is the maximum amount of time, following a disaster, for an organization to recover files from backup storage and resume normal operations. In other words, the recovery time objective is the maximum amount of downtime an organization can handle. If an organization has an RTO of two hours, it cannot be down for longer than that.
- The RPO and RTO helps administrators choose optimal disaster recovery strategies, technologies and procedures.

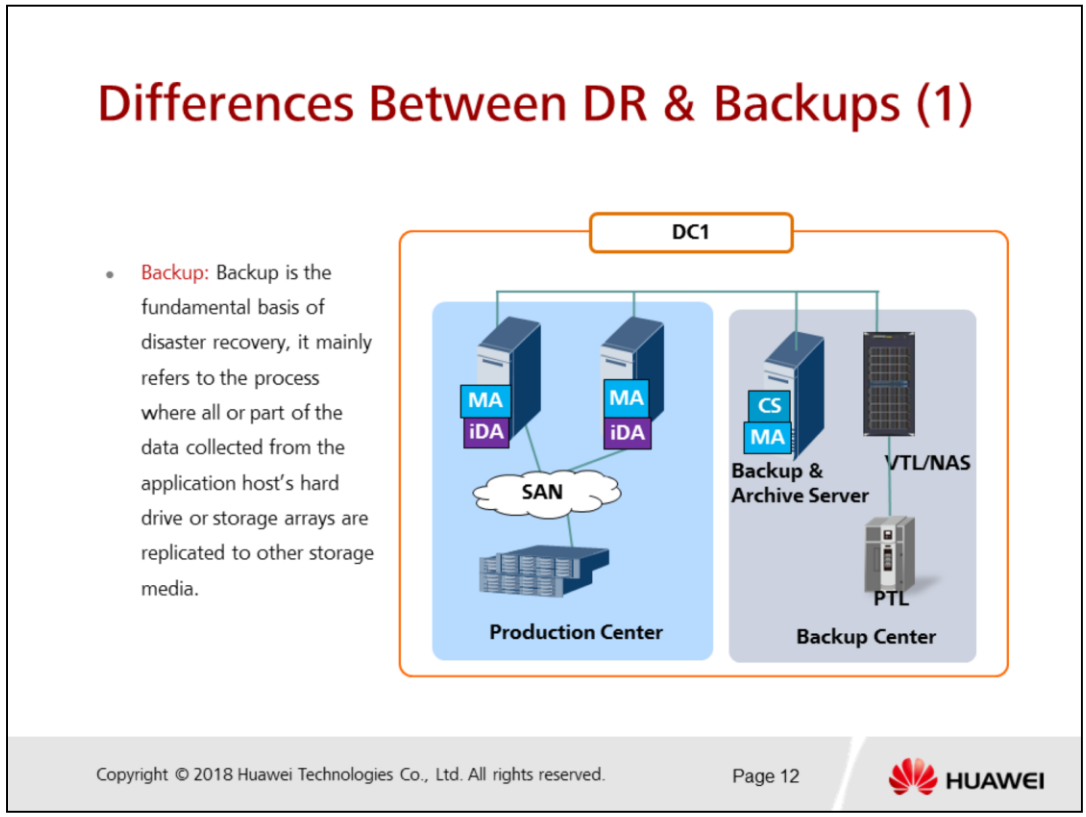

- Backup solutions are designed to provide quick data access in the event of a sudden, small-scale computing outage, such as the loss of a laptop, accidental file deletion or crashed hard drive. At their most basic form, backup could be just the process of simply copying critical files to a second storage environment. This could be a tape backup, secondary in-office workstation, or hosted cloud system. Backups excels at providing immediate access to critical files and quickly restoring data. Thus, reducing the need to rely on time-consuming OS re-installations or manually retrieving and copying files from long-term storage solutions.
- Disaster recovery, meanwhile, is designed to provide a large-scale action plan for restoring critical data and systems in the event of a serious computing outage. While the term "disaster" often calls up images of massive storms, earthquakes or tornados, a disaster covers anything that causes widespread network outages or prevents your employee from completing necessary tasks. Disaster recovery solutions typically cover the entire process necessary to restore system function — from re-establishing network connections to repopulating data to establishing recovery time objectives (RTOs).
- While backups form a critical part of your disaster recovery plan, they're not enough on their own and need to work with a proper DR plan to ensure business continuity in the event of a disaster.

Differences Between DR & Backups (2)

- **Disaster recovery:** Disaster recovery system refers to the establishment of two or more sets of IT systems with the same function in remote locations separated from each other. Health status monitoring and function switching can be performed between them. When a system stops working due to the event of disaster (such as fire, earthquake), the entire application system can be switched to another place so that the system function can continue to work normally.

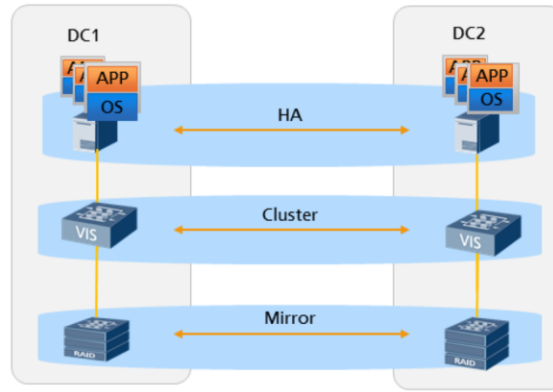

- Generally, disaster recovery refers to backup of data or application systems that are not in the same server room or location. Backup refers to local data or system backups.
- In general, DR combines disaster recovery with backup, which is, local backup combined with remote data replication to achieve perfect data protection.

HUAWEI

- Generally, backup is implemented using backup software technology, and disaster recovery is implemented through replication or mirroring software. The fundamental difference between the two is:
  - The format of the data processed by the backup software is inconsistent and must be restored before the data can be used. The data format processed by the replication or mirroring software does not change and thus can be directly mounted to the host.
  - The data protection cycle of the two is inconsistent, and the time period for copying or mirroring is shorter.
  - Generally, backup is the last line of defense for data protection and it is more biased towards archiving data.

Key Indicators For Measuring DR System Performance

- RPO (Recovery Point Objective) represents the amount of data that is allowed to be lost when a disaster occurs.
- RTO (Recovery Time Object) represents the system recovery time.
- The smaller the RPO and RTO, the higher the availability of the system and the greater the investment required by the user.

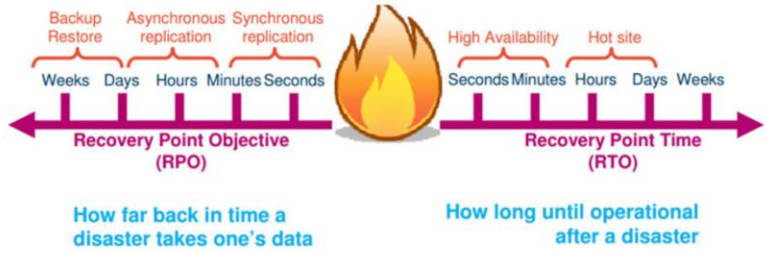

HUAWEI

- Developing a good continuity plan means deciding upon how much data your business needs to protect, and how much time it can afford to be without it.
- These distinct volume sets are referred to as your Recovery Point Objective (RPO) and Recovery Time Objective (RTO), respectively. Creating a comprehensive business continuity plan requires setting tolerances for each of them.
- Learning the difference between RPO and RTO, and setting realistic tolerances for each, is one of the most important steps you can take to ensure business continuity in the event of a data disaster.

Levels Of DR Systems (1)

Level	Definition	RTO	TCO
Data Level	Through the establishment of remote disaster recovery centers, remote backup of data will be performed to ensure that the original data will not be lost or destroyed after the disaster. In the data-level disaster recovery mode, the established disaster recovery center can be simply considered as a remote data backup center. At the data-level disaster recovery level, applications are interrupted when a disaster occurs. Data-level disaster recovery has a long recovery time, but its cost is lower than other disaster recovery levels, and its implementation is relatively simple. Data sources are the source of life for all critical business systems, so data-level disaster recovery is essential.	The longest RTO (several days) because after the disaster, the equipment needs to be redeployed to recover the business services using backup data.	Lowest

- The table above shows the detailed description of data level disaster recovery and its expected RTO and TCO.
- Different level of DR has their own method of implementation and they cover different aspects of DR which are data, application and service.
- It is important to formulate the DR plan of your business based on the DR requirement on each level of DR, investment cost and expected RTO/RPO.

Levels Of DR Systems (2)

Level	Definition	RTO	TCO
Application Level	On the basis of data-level disaster recovery, similar sets of application systems are also built at the backup site. Through synchronous or asynchronous replication technology, key applications can be guaranteed to recover within the allowable time range, minimizing the loss caused by disasters, and allows users to be unaware of the disaster. Thus, the services provided by the system are complete, reliable and safe.	Medium RTO (Few Hours)	Medium. Same or similar system or subsystems is built offsite.
Service Level	In addition to the necessary IT-related technologies, full-service disaster recovery requires full infrastructure. Most of its components are non-IT systems (such as telephones, office locations, etc.). When a catastrophe occurs, the original office space will be destroyed. In addition to the restoration of data and applications, a back-up workplace needs to be provided so business can be conducted as usual. .	Smallest RTO (Few Minutes or Seconds)	Highest

- The table above shows the detailed description of application level and service level disaster recovery and its expected RTO and TCO.
- Different DR level has different performance in terms of service failovers in the event of the disaster. Although higher performance (low RTO, high RPO) DR level ensures greatest amount of business continuity protection to the enterprise, it is important to consider the amount of investment required for implementation.
- It is best to choose the DR plan that fits the current scale of business with some rooms for growth and improve your DR plan along with the growth of your business.

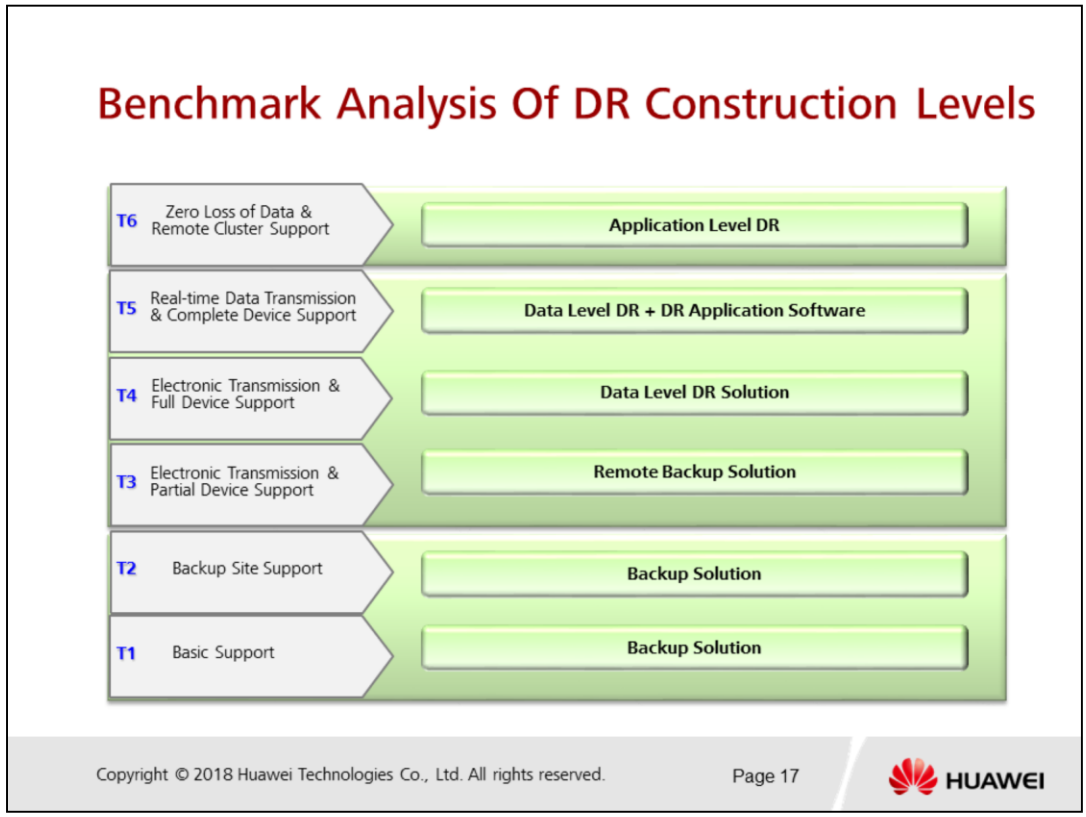

- The diagram above shows the different level for the benchmarking of the DR construction levels.

Panorama Of DR Solution

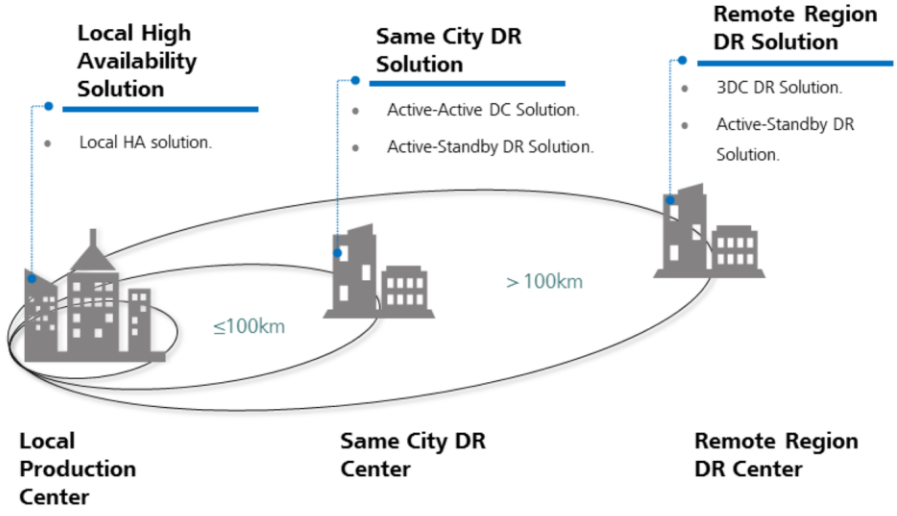

HUAWEI

Contents

HUAWEI

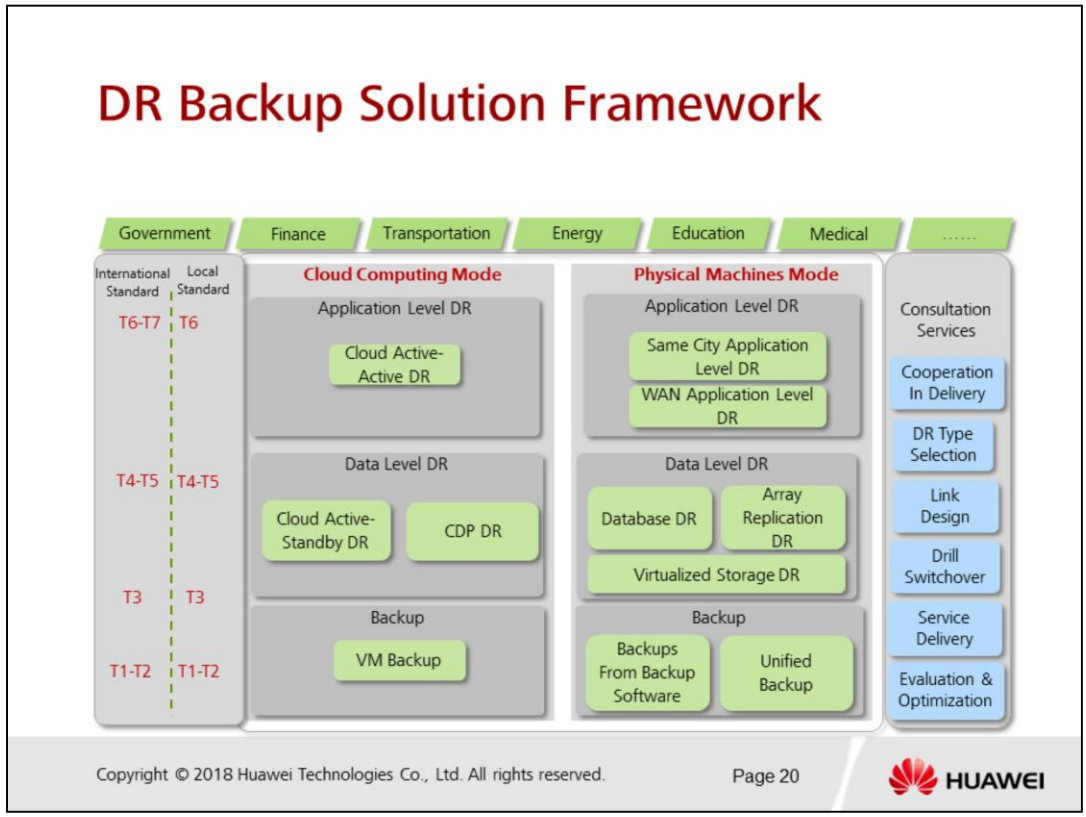

- Match customer business and development strategies, and provide professional services from strategic consulting, disaster recovery planning, and service implementation for continuous operation and management.

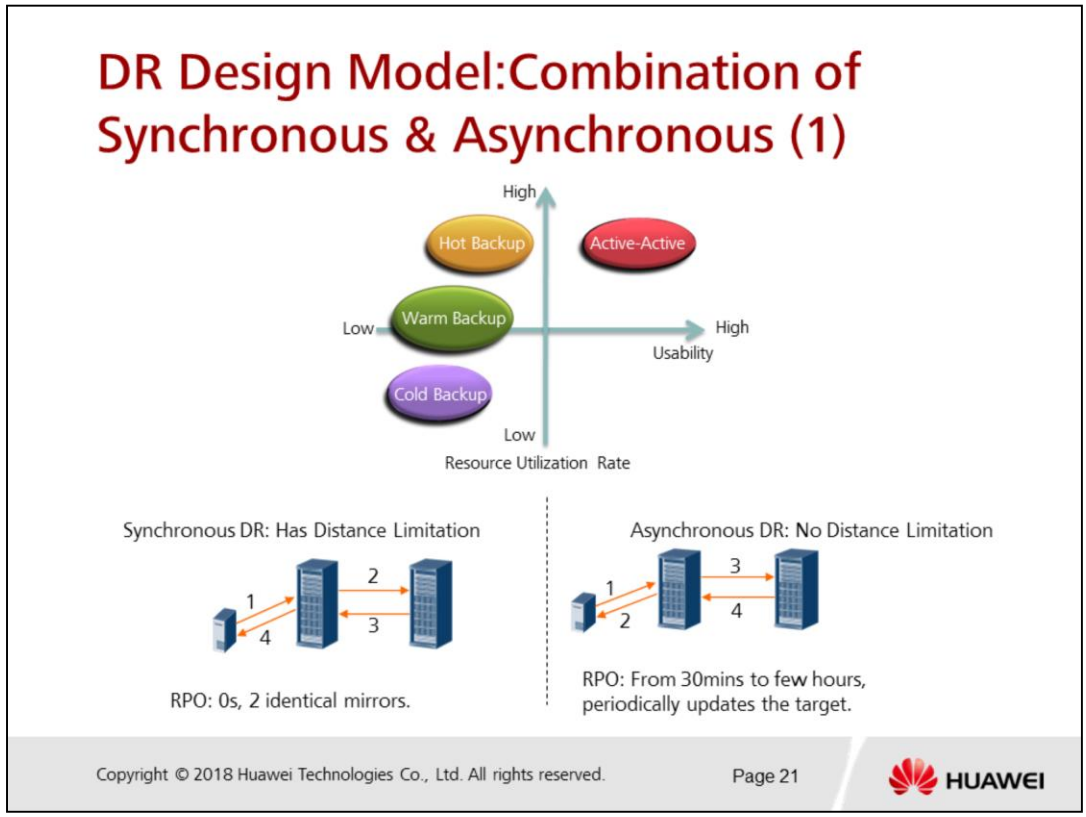

- One of the important factors to consider when constructing DR solution is to choose where and when to implement the synchronous and asynchronous DR technologies.
- Both synchronous and asynchronous DR has its challenges and limitations in implementation and requires careful planning and commissioning.
- Both these technologies can be implemented on its own for DR, but it can also be combined to provide better results and protection of the business data.

DR Design Model:Combination of Synchronous & Asynchronous (2)

DR Mode	Reliability Solution	DR Recovery	Data Backup Requirements
Active-Active	Cluster + Load Balancing	Automatic	Real-time synchronous replication (<100KM)
Hot Backup	Cluster	Automatic	Real-time synchronous replication (<100KM)
Warm Backup	Manual Intervention	Manual	Asynchronous replication (>100KM)
Cold Backup	Strong Manual Intervention	Manual	Same as above

- The table above shows the different DR mode that can be implemented with comparisons on its reliability, recovery process and data backup requirements.

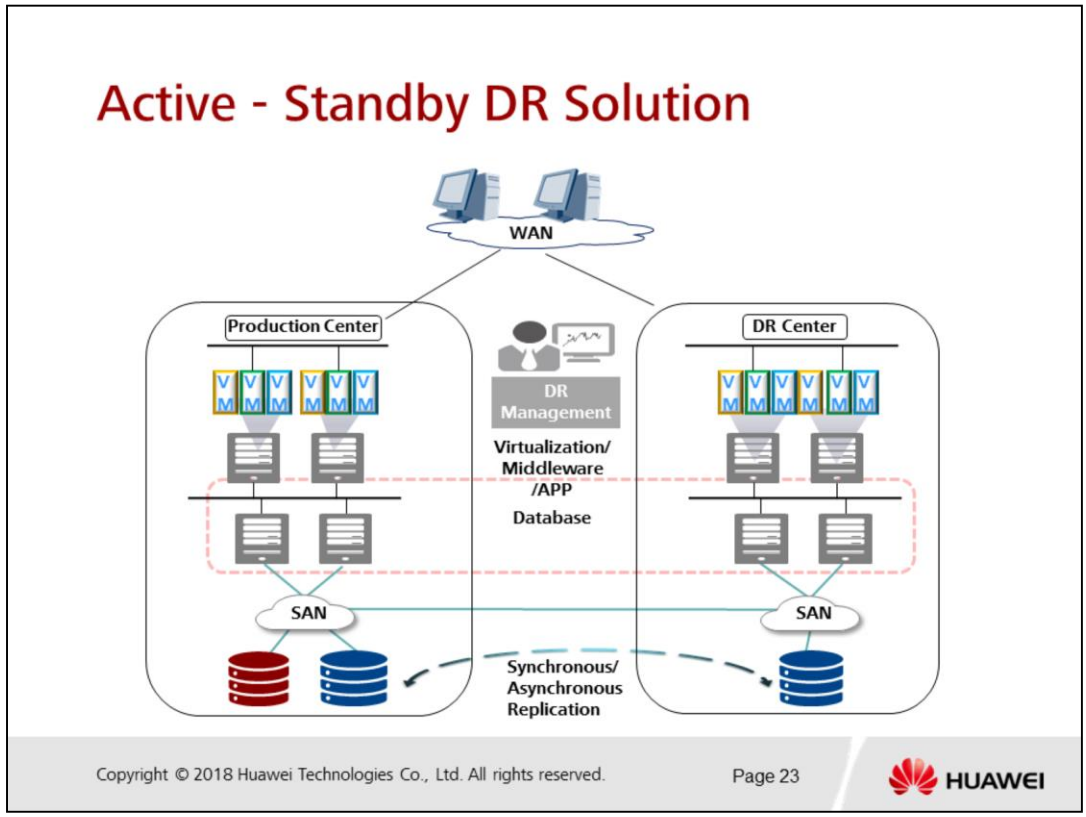

- Disaster Recovery System Management Visualization:
  - Support disaster recovery management software deployment and commissioning.
  - Supports disaster recovery drills and one-click switching, and has capabilities to assists in the completion of custom scripting tools, and assists you in recovering the standby business system in one click.
- Disaster recovery service is mature and efficient:
  - Disaster recovery system has one-stop analysis, design, delivery, and drill service.
  - Specialized disaster recovery drill process and technical plan, with 20+ disaster recovery drill project implementation experiences including finance, government, and medical.
  - Support the implementation of the disaster recovery plan for old non-Huawei equipment and reduce the cost of customers by 40%.

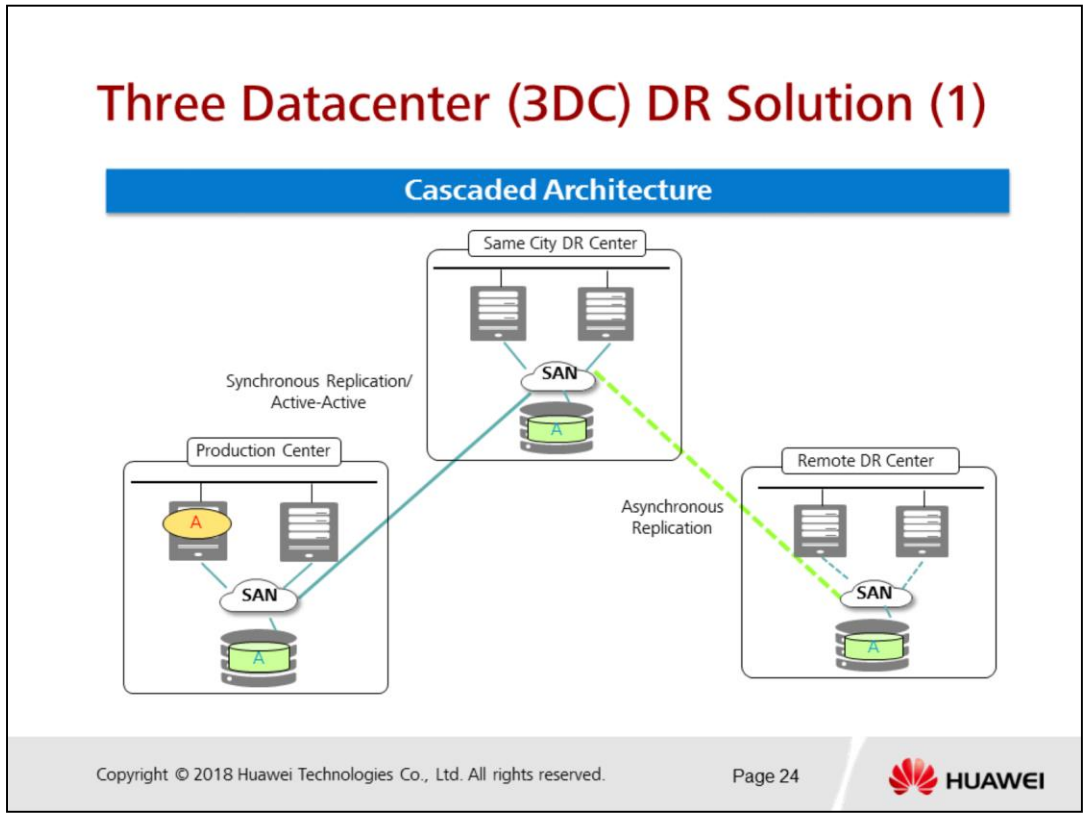

- Disaster recovery construction cycle is short and delivery risk is low:
  - Disaster recovery construction period shortened from 10 months to 7 months, which is shortened by 30%.
  - Multi-vendor cooperation with effective management, and thus shortens the project delivery cycle.
  - Effective assessment and analysis of multiple services and multiple applications to ensure rapid construction of disaster recovery systems.
  - Has Asia's largest integrated verification center to ensure effective validation of disaster recovery design and reduces the project implementation risk.

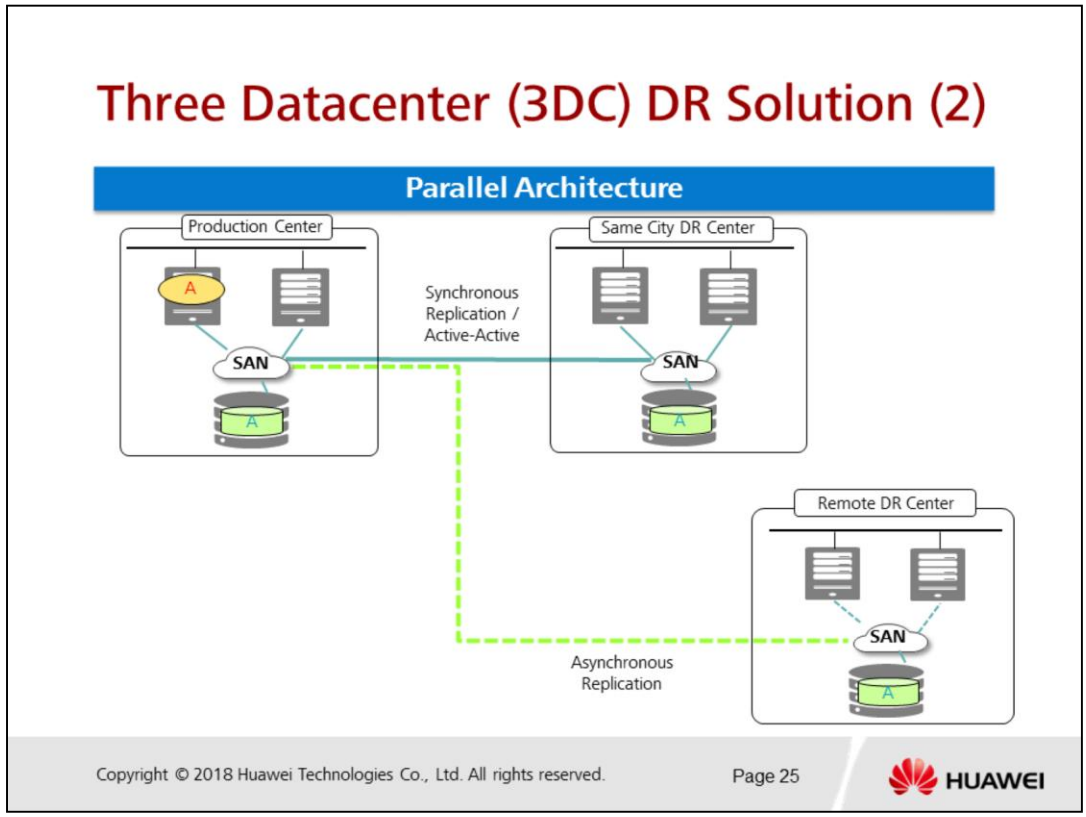

- Remote disaster recovery system management visualization:
  - Support one-click visual deployment and commissioning of disaster recovery management software.
  - Support unified management and monitoring of production centers, disaster recovery facilities in the same city, and remote disaster recovery center equipment, and thus simplifies the equipment maintenance process.
  - Support one-button disaster recovery drill and switching, and customer-customized script recovery of standby service system, and simplified management and maintenance of disaster recovery system.

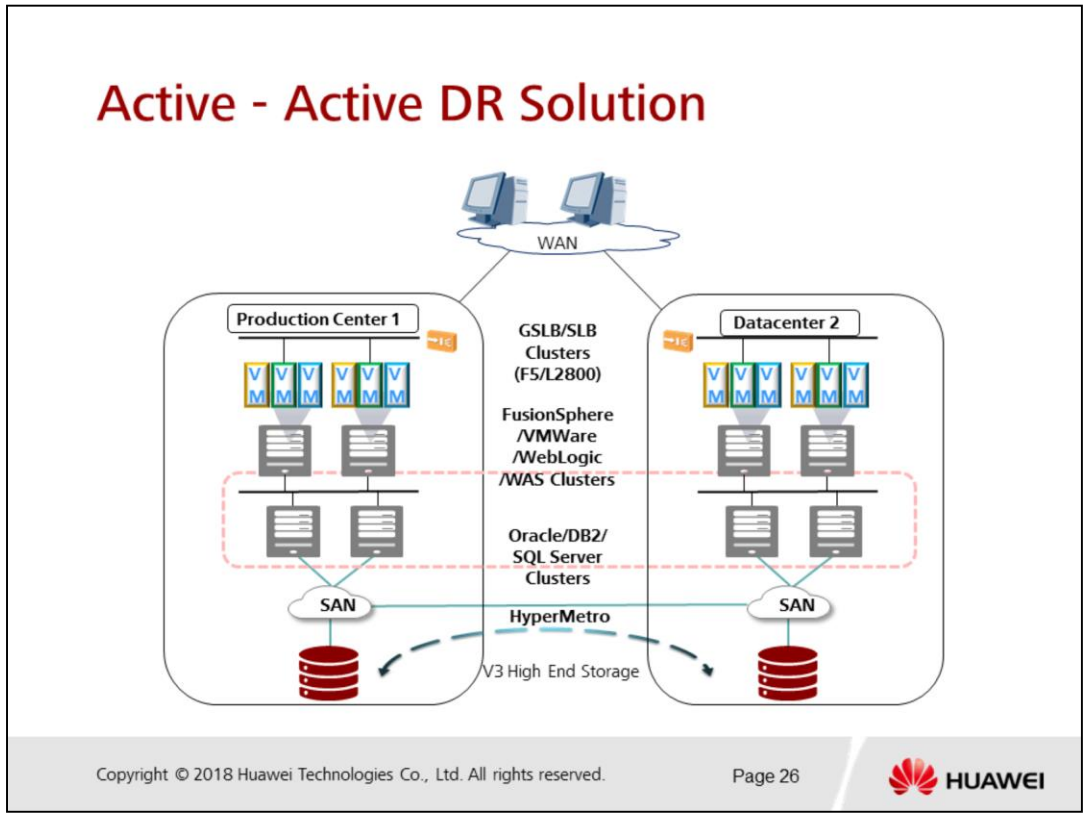

- Active-Active gateway free disaster recovery, has high efficiency in implementation:
  - Enterprise grade "Active-Active" high reliability, no service downtime even when there is data center level of fault occurs, 7*24hours online.
  - Active-Active storage layer with no virtualized gateways, reduces points of failures, and simplifies implementation and commissioning.
- Rich experiences in Active-Active DR services:
  - Rich experience in the implementation of active-active disaster recovery and implementation of more than 80+ active-active DR projects in the fields of finance, government, and medical care.
  - Completes the design and implementation of the 3DC DR plan from the existing active-active DR plan without stopping the running services.
  - With the largest integrated verification center in Asia, extensive range of IT equipment allows comprehensive test verification and remote demonstration of the disaster recovery plan to be conducted to ensure the perfect design of the plan.
  - Extensive disaster recovery implementation experience, strong professional service tools support, improve disaster recovery program implementation efficiency, and reduce TCO by up to 30%.

- Global Server Load Balancing (GSLB), refers to the technology where the Internet traffic can be distributed among different data-centers located at different locations of the planet. This technology is highly efficient in avoiding local downtimes and downtimes. There is a Master GSLB which monitors the health and responsiveness of the Slave sites, which in turn identifies the sites are available and can offer the best response. SLB refers to Server Load Balancing which is implemented locally within the network.
- WAS Cluster: WebSphere Application Cluster is a technology by IBM. It is designed to provide high availability of the services tier, by creating a cluster of application servers across two or more computers. It deploys an instance of the services tier in each application server. When a member of the cluster fails, other cluster members continue to provide services. As long as at least one member of the cluster is operational, there is no interruption in service.
- HyperMetro enables storage systems in two different data centers to process services simultaneously, by establishing a mutual backup relationship. If the storage system in one data center malfunctions, the storage system in the other data center automatically takes over services without data loss or service interruption.

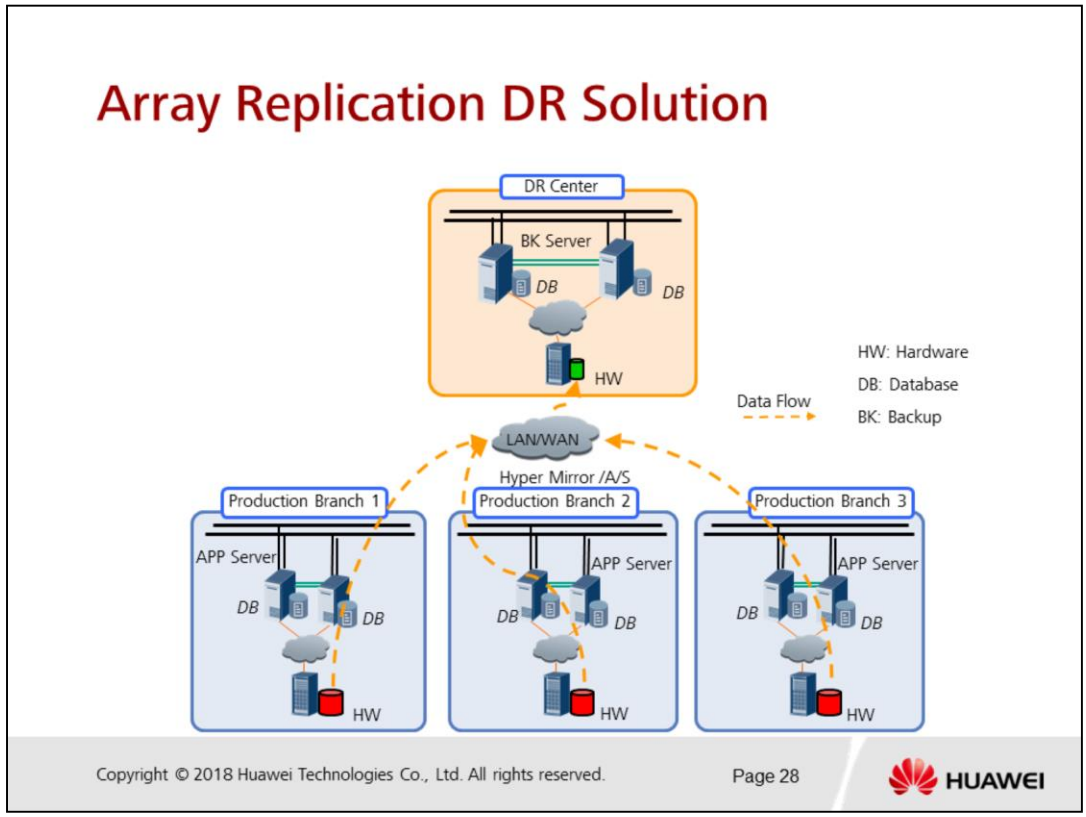

- Application scenario:
  - Support one-to-one disaster recovery or multi-to-one disaster recovery.
  - Multi-branch's data is stored independently. A unified disaster recovery system needs to be established to provide unified disaster recovery protection for each branch.
- DR Solution Features:
  - Supports up to 31 branches.
  - Support synchronous/asynchronous replication mode, flexible selection according to the characteristics of each branch.
  - Unified management of disaster recovery storage systems in each branch.
  - Different levels of arrays can replicate with each other, thus reducing TCO.
  - Supports flexible online addition/removal of branch sites.
  - With the snapshot function, continuous data protection can be achieved.
- Value to Customer:
  - Disaster recovery resource sharing, and saving investment.
  - Disaster Recovery Center maintains and manages resource sharing, and also saving maintenance and management costs.
- HyperMirror is a continuous data protection technology that creates two physical mirror copies of a LUN for redundant backup protection against host service interruption.

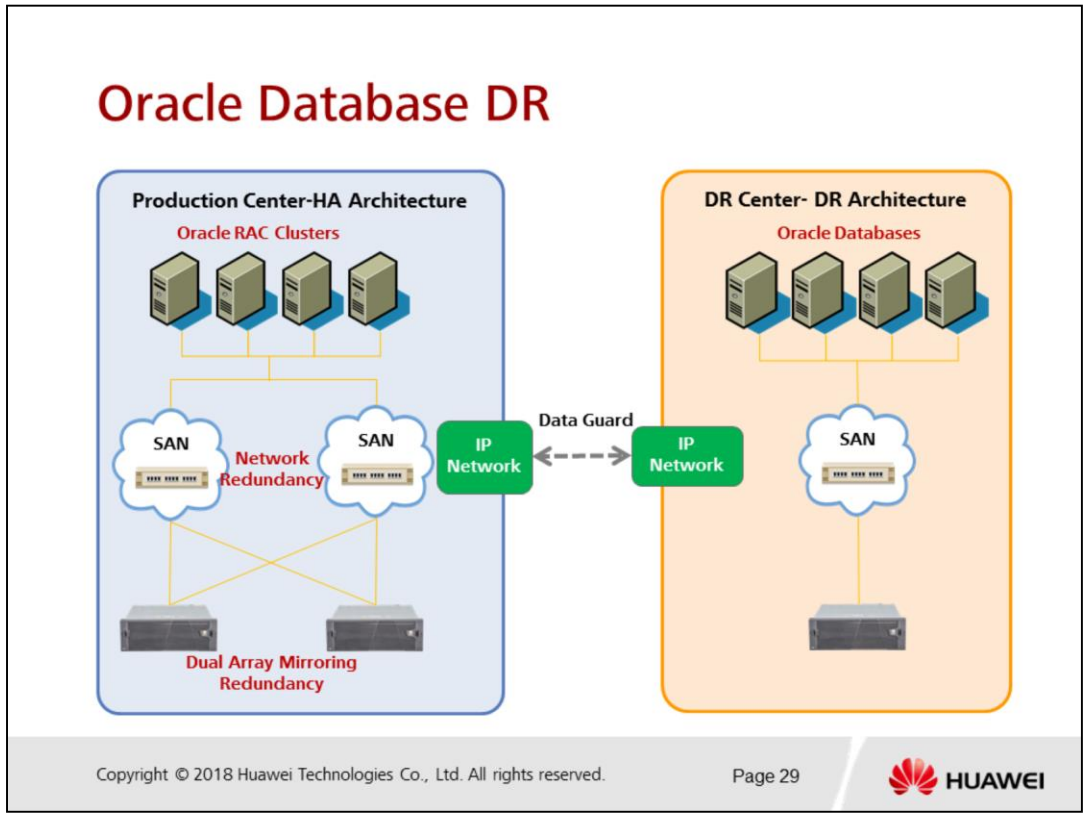

- Application scenario:
  - Adopts Oracle database.
  - Need to establish remote disaster recovery.
- Solution Features:
  - Local high availability, and remote disaster recovery.
  - Fully redundant architecture design.
  - Heterogeneous hardware compatibility.
- Value to Customer:
  - Little changes to the existing network environment.
  - Simple maintenance.
  - Redo-based replication with low bandwidth requirements.
- Oracle RAC(Real Application Cluster) is a cluster database with a shared cache architecture that overcomes the limitations of traditional shared-nothing and shared-disk approaches to provide highly scalable and available database solutions for all business applications.
- Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data. Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruptions.

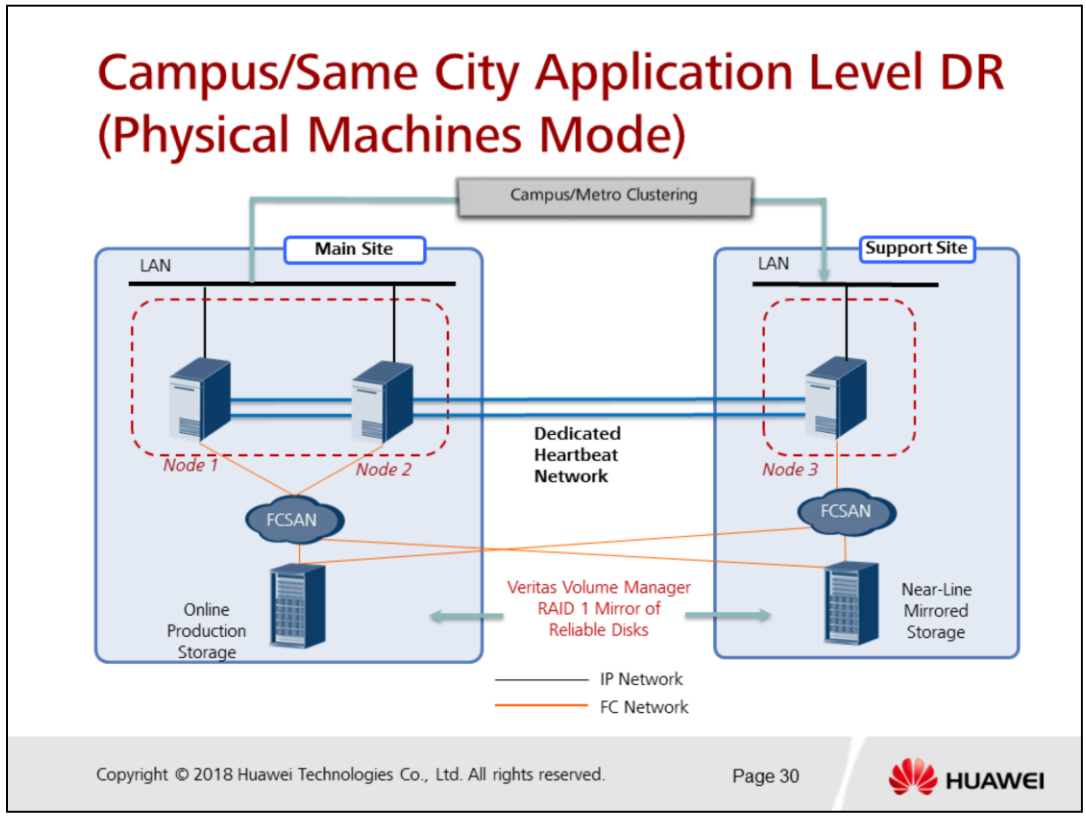

- Application scenario:
  - Physical machine mode (minicomputer, X86 server).
  - Two data center's distances must be <80KM.
- Value to Customer
  - Application automatic switching to quickly recover services.
  - Supports a wide range of application types and can be tailored to customer needs.
  - Develop a customizable switching script for each application.
- Program Features
  - Achieve local high availability and remote application level disaster recovery.
  - Equipped with openness for storages, and supports heterogeneous storage.
  - Cluster nodes can be arbitrarily combined (the total number of hosts in the cluster must be <= 64).
- Veritas Volume Manager: A Veritas software product installed on storage clients that enables management of physical disks as logical devices. Volume Manager enhances data storage management by controlling space allocation, performance, data availability, device installation, and system monitoring of private and shared systems.

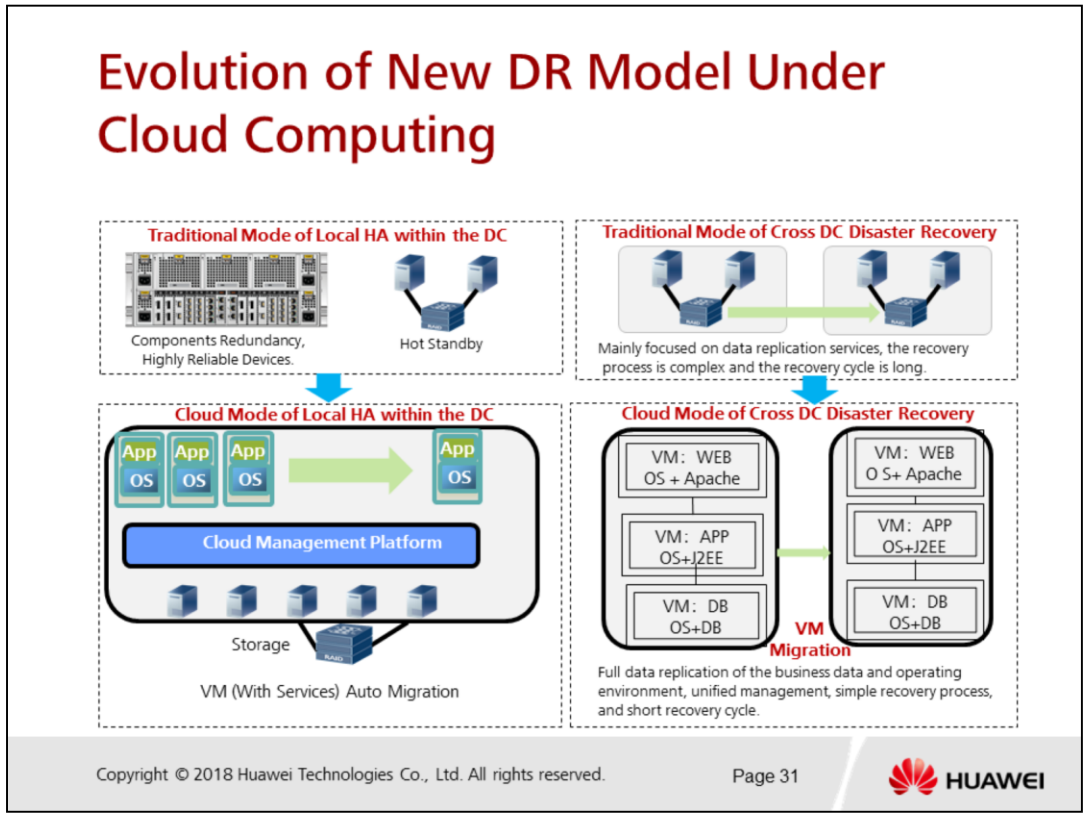

- After IT is centralized and cloudified, higher requirements are placed on business continuity, including on network requirements, data security requirements, and business reliability requirements.

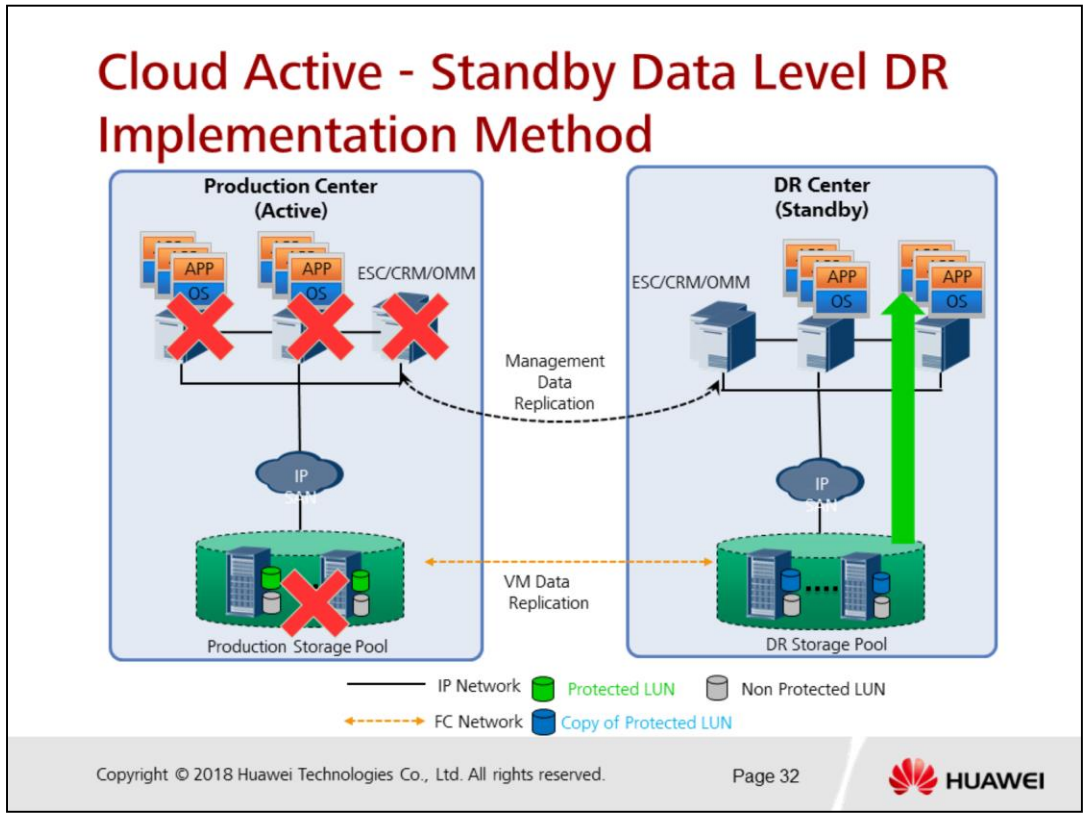

- Production Center and Disaster Recovery Center independently deploy cloud management platform (two cloud platforms deployed).
- Set up a synchronization strategy and regularly perform remote copying of the cloud management data and business data (VM) from the production center to the disaster recovery center
- In the service planning, users can divide two types of LUNs according to actual needs: protection LUNs and unprotected LUNs, create VMs that need disaster recovery on the protection LUNs, and configure array replication only for protection LUNs to save on the requirements of storage capacity within the DR center.
- When a fault occurs in the production center, the disaster recovery center uses the disaster recovery management software to restore the virtual machines in one click.

Contents

HUAWEI

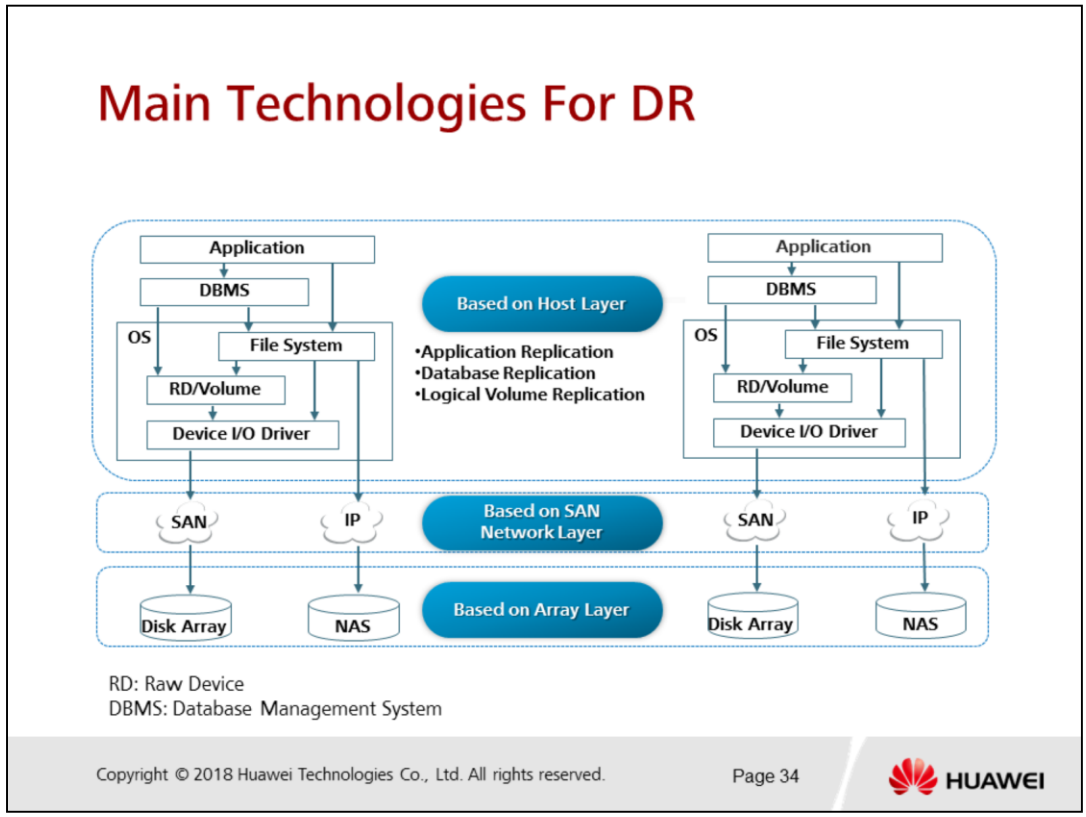

- Disaster recovery technology based on host layer :
  - Install dedicated data replication software, such as volume replication software, on servers in production centers and disaster recovery centers to enable remote replication. There must be a network connection between the two centers as data channels. The remote service switching function/service failover software can be added at the server layer to form a complete application-level disaster recovery solution.
  - This method of data replication has relatively less investment, mainly in the cost of software procurement. Its compatibility is better, and can be compatible with different brands of servers and storage devices, which is more suitable for users with complex hardware composition. However, this method requires synchronization of software operations on the server, and occupies a large amount of host resources and network resources.
- Disaster recovery technology based on network layer :
  - The SAN-based data replication technology consists of a storage area network (SAN) between the front-end application server and the back-end storage system. It joins/connected to the storage gateway, and connects to the front-end server host, and also connects to back-end storage devices.

    - The storage gateway establishes a mirror relationship between two volumes on different storage devices and writes all the data that was written to the primary volume to the backup volume.
    - When the primary storage device fails, the service is switched to the backup storage device and the backup volume is enabled to ensure that the data service is not interrupted.
- Disaster recovery technology based on array layer:
    - The storage layer disaster recovery adopts the data replication technology between the arrays. The data is copied from the local array to the disaster recovery array, and a copy of available data is generated in the disaster recovery storage array. When the primary array fails, services can be rapidly switched to the standby array, thus ensuring the greatest possible continuity of services.

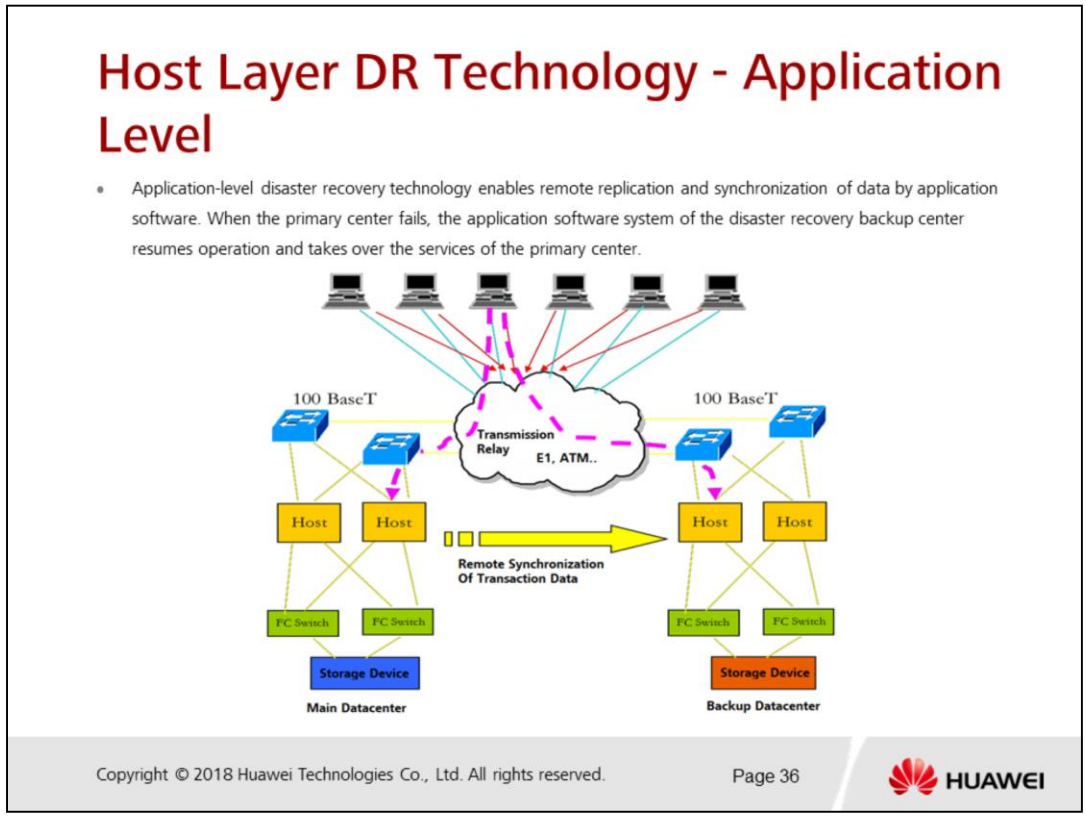

- Working principle: Two remote databases are connected through the application software. Each time the business services process data and separately stores them in the database of the main center and the backup center.
- Advantages and disadvantages:
  - Supports wide area networks, and no separate hardware or software support is required. Data logic replication prevents proliferation of human errors and these processes are transparent to disk subsystems.
  - Consistency checks need to be performed regularly. Backup data from the backup center cannot be quickly restored back to the main center and major modifications to the application are required.

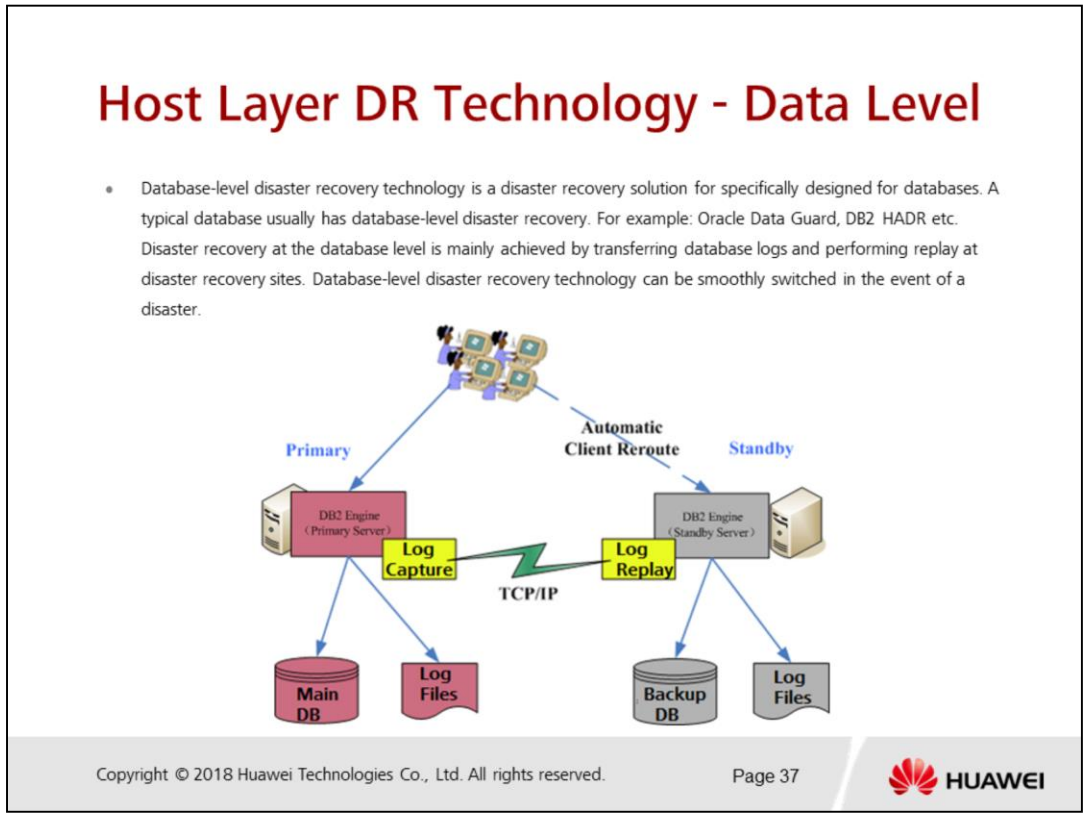

- Working principle:
  - Configure primary and standby database servers.
  - Once the primary database has transaction operations, the log files are simultaneously transferred to the standby database, and then the standby database replays the received log files to maintain consistency with the primary database.
  - When the primary database fails, the standby database server can take over the transactions of the primary database server.
- Advantages and disadvantages:
  - Support for wide area networks and no separate hardware support required. It is transparent to disk subsystems. Implementation of logical replication to reduce the risk of proliferation of human error. There is no need to modify the application. The data in the main center/disaster recovery center can be accessed at the same time.
  - The backup center's backup data cannot be quickly restored back to the main center. Non-database data can't be copied remotely. The synchronization system has a great impact on the production system. In the asynchronous mode, more data will be lost and at least one log file will be lost. Complex service switching process and production transformation is complicated.

Host Layer DR Technology - Logical Volume Level

- Logical disk volume-based remote data replication refers to the remote synchronous (or asynchronous) replication of one or more volumes based on demand. The implementation of this solution is usually done through software.

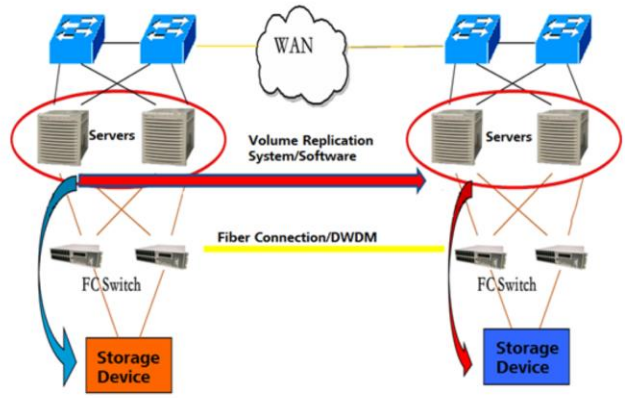

HUAWEI

- Working principle: The remote replication control management software copies the operation data of each I/O on the volume of the active node system to the corresponding volume of the remote node, in real time (or with delay) so as to achieve the data synchronization between the two remote volumes whether if it is synchronous or asynchronous.
- Advantages and disadvantages:
  - Ensures data integrity and consistency. The structure is relatively simple and transparent to the disk subsystem.
  - The write performance of the host is affected by the distance. If there is no host at the disaster recovery center, data-level disaster recovery cannot be performed and thus logical level disaster cannot be prevented.

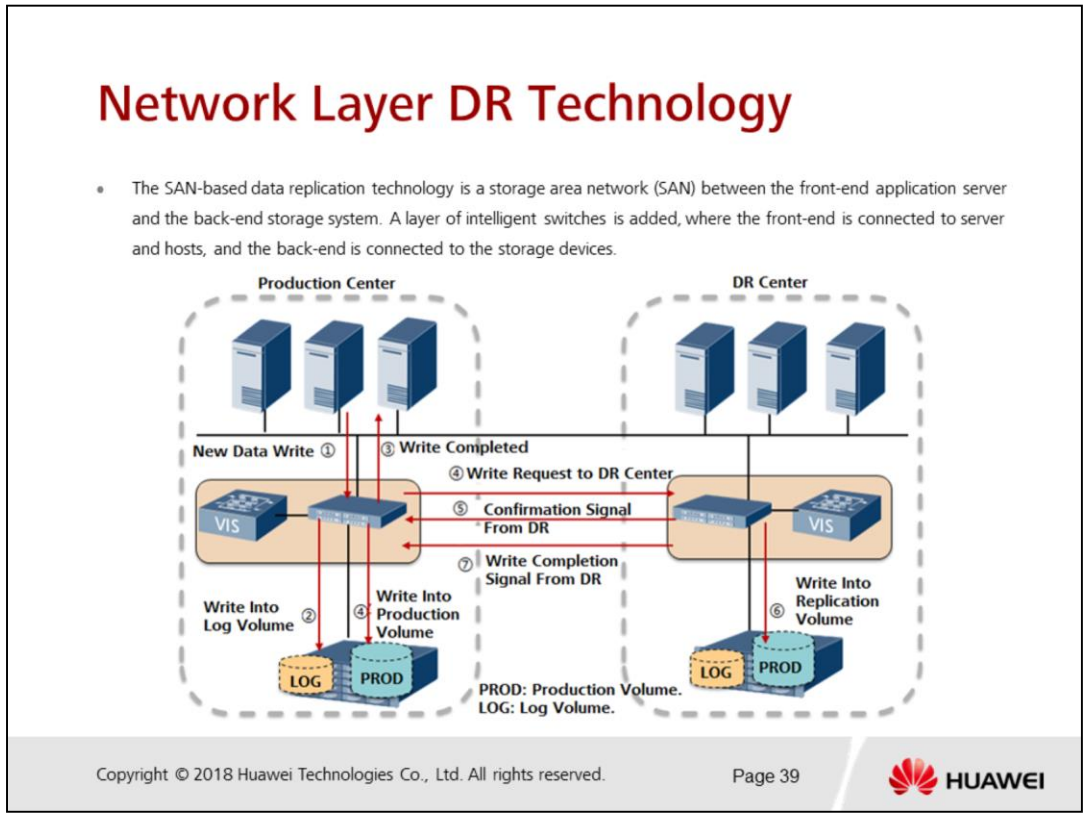

- Working principle:
  - The production center host writes data to the local virtual gateway.
  - The production-side virtual gateway writes data to the local log volume.
  - After the data is successfully written to the log volume, the virtual gateway of the production center returns the "confirmation" to the local host.
  - When the production-side virtual gateway writes data to the local production volume, it sends a data write request to the disaster recovery-side virtual gateway.
  - After the disaster recovery virtual gateway receives the write request, it returns a "confirmation" to the production virtual gateway.
  - The disaster recovery-side virtual gateway then writes data to the disaster recovery-side replication volume.
  - After the data is successfully written to the replication volume in the disaster recovery center, the virtual gateway in the disaster recovery center returns a “completed” signal to the virtual gateway in the production center.
- Advantages and disadvantages:
  - Support heterogeneous storage devices, allows virtualization integration, achieves unified management, and improve storage utilization.
  - Need to transform the SAN network.

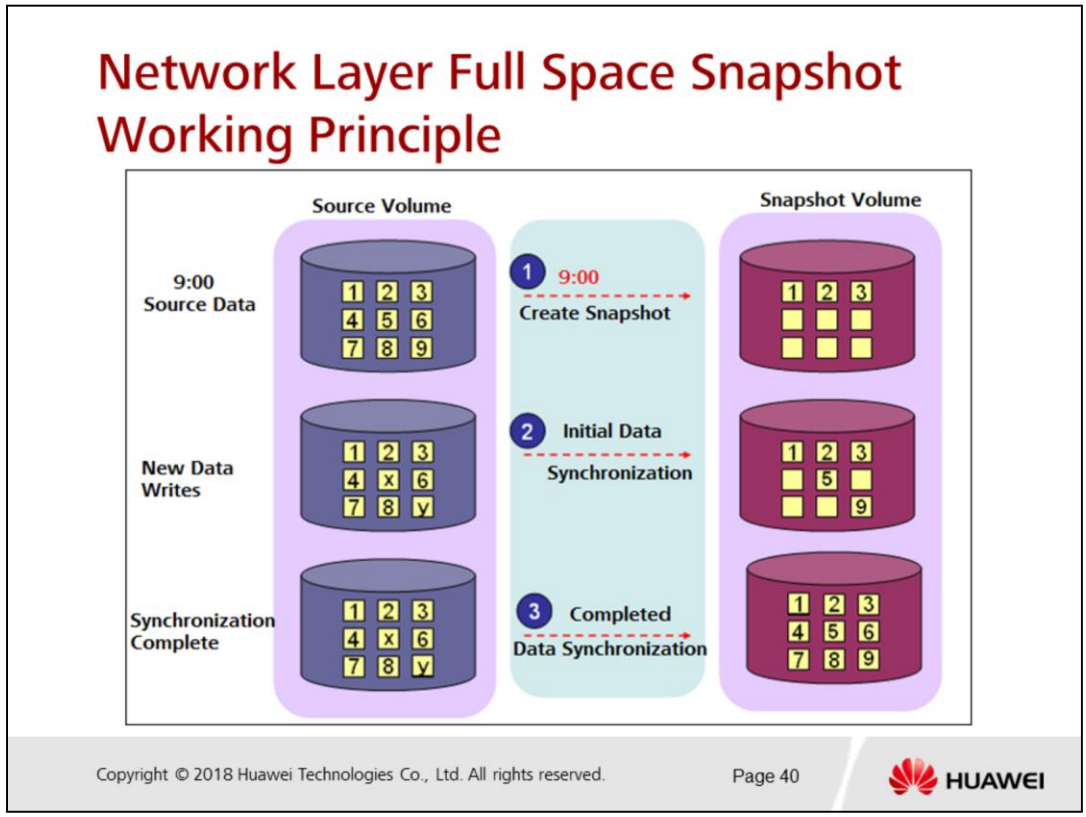

- Technical principle:
  - The working principle of the full-space snapshot technology: When the snapshot time is reached, the system allocates a physical space with the same size as the snapshot volume for the source data volume, and starts background data synchronization. After the synchronization data is completed, the snapshot is created at this point in time is considered as a success.
  - A full-sized snapshot is a point in time physical copy of the data at the source volume's snapshot.
- Steps:
  - Create a volume with the same size as the source volume as the snapshot volume and start background data synchronization.
  - If new data is written to the source volume in the data synchronization process and the data is written to the content that has not been copied to snapshot volume yet, the original data is written to the snapshot volume. Subsequently, the new data is written to the source volume to retain its most up to date status. The data is up-to-date if the location of the newly written data is part of the already synchronized copy of contents in the snapshot volume, thus only the new data is written to the source volume and the contents of the snapshot volume data are unchanged.

    - After the data is all synchronized, the snapshot volume and the 9:00 source volume data are exactly the same, and the snapshot process ends.

- Explanation:
    - The snapshot volumes in the full-space snapshot of the network layer can span heterogeneous arrays and can be placed on relatively low-end arrays that has low performance. In this way, it is possible to implement disaster recovery between arrays, while fully exploiting the old storage resources and reducing the TCO at the same time.
    - When the source volume array fails, services can be quickly pulled from the snapshot volume array.

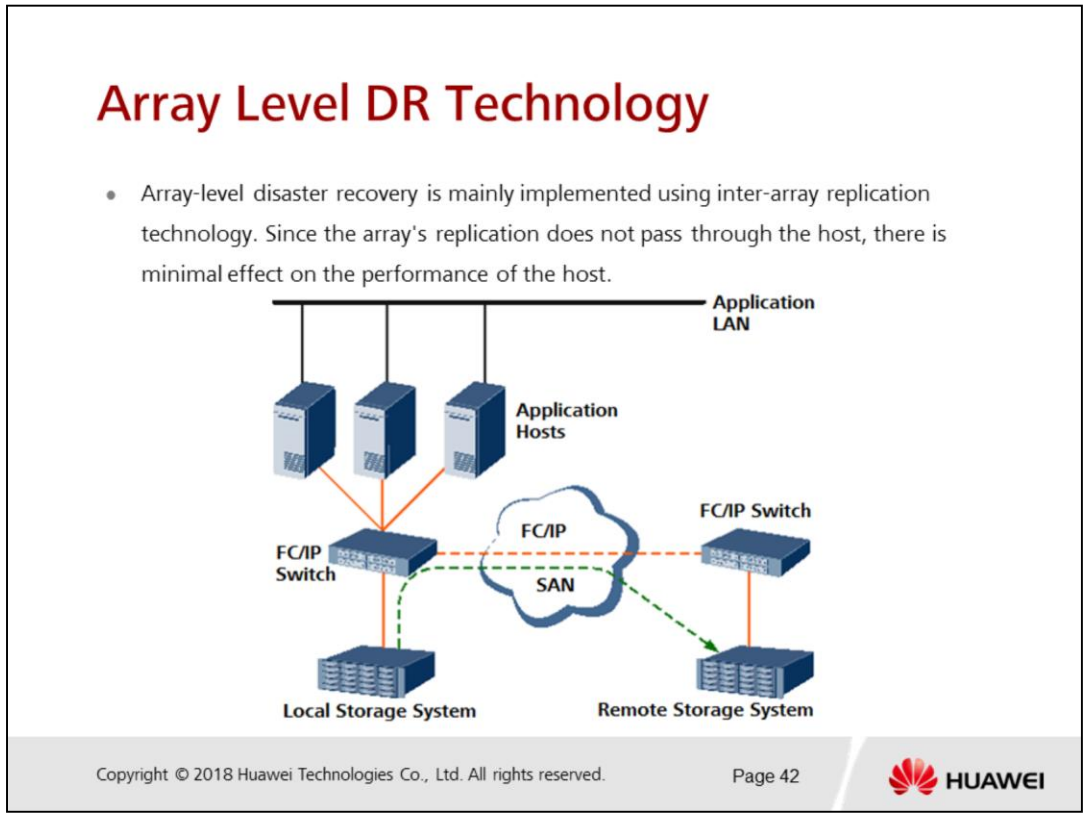

- Array Level DR technology are one of the simple DR methods used to protect data by utilizing 2 similar storage arrays to replicate data. Although it has the benefits of low impact to host performance and simple implementation, there are a few downsides.
- The Downsides to Array Based Replication
  - The first problem with array based replication, in almost all cases, is that it requires a near identical unit from the same manufacturer at the DR site. The DR site though may not need the same level of performance or redundancy as the primary, and a less expensive system may be a more appropriate fit.
  - This inability to mix storage hardware also means that adding new hardware from a new vendor must be considered carefully, since the introduction of a new vendor will mean learning another replication module and buying a second storage system for the DR site. It may also mean that a newer more innovative storage vendor may not be considered because they lack a replication capability.
  - The third problem is that array based replication software is not application aware. These modules just replicate blocks of data as it changes on the array with no understanding of the application that is changing that data or how other servers may be related to that application. Software based replication software that is application aware can ensure that clean copies of data are captured. They can also make sure that groups of servers that are dependent on each other can be kept in sync

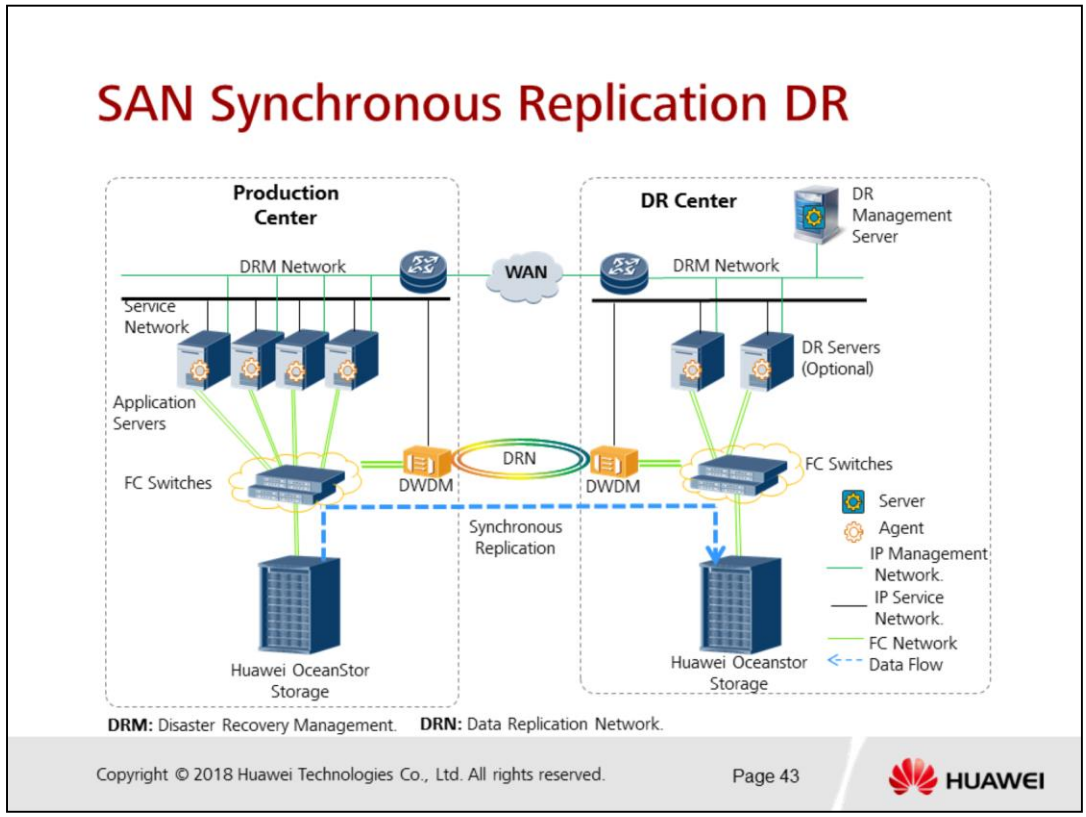

- The deployment method is shown in the figure above has the target RPO=0 and level of RTO within minutes.
- SAN-based disaster recovery replication supports synchronous replication, which is recommended to be within the distance of 100km.
- Remote Datacenter (RD) mainly provides disaster recovery management functions, including topology, disaster recovery testing, drills and disaster recovery.
- When performing application management and disaster recovery application recovery, the Agent needs to be installed on the server
- The RD management network needs to communicate with the host and storage.
- It supports FC/iSCSI links, and it is recommended to use FC links for synchronous replication.
- Dense wavelength division multiplexing (DWDM) is a technology that puts data from different sources together on an optical fiber, with each signal carried at the same time on its own separate light wavelength. Using DWDM, up to 80 (and theoretically more) separate wavelengths or channels of data can be multiplexed into a light stream transmitted on a single optical fiber.

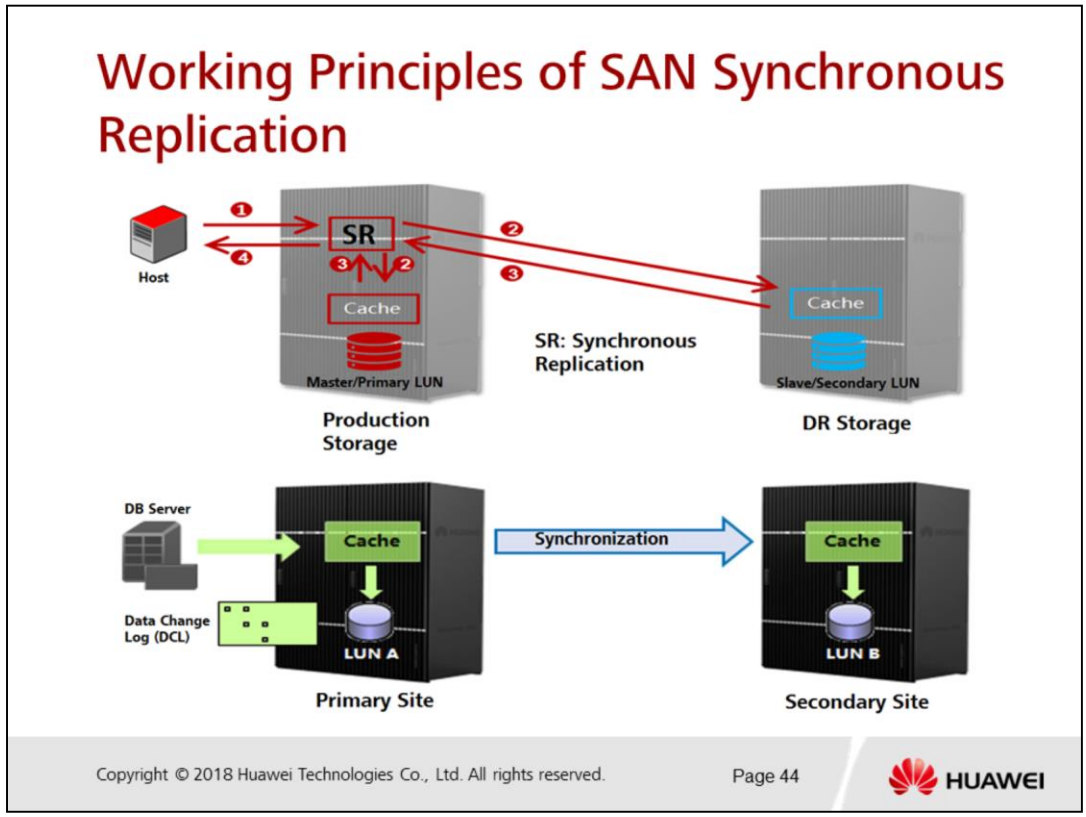

- Synchronization steps:
  - The production storage system receives a write request from host. HyperReplication logs the request. Only the address information is recorded in the log, and the data content is not recorded.
  - It write this request to the master and slave LUNs. Normally, LUNs are in write-back mode and the data is written to Cache.
  - HyperReplication returns both write processing results for the master and slave LUNs. If all the writes are successful, the log is cleared. Otherwise, the log is retained and it enters an abnormal disconnection state. When the synchronization starts back, the data block corresponding to the log address is copied again.
  - It returns the processing result of the host write request, which is based on the processing result of the data write to the master LUN.
- Split:
  - In the split state, the write request of the production host is only written to the master LUN, and the difference log records the difference between the master and slave LUN data. When you want to re-maintain data consistency between the primary and secondary LUNs, you can manually start the synchronization operation. The synchronization process is the process of incrementally copying the data block marked as "differenced" in the differential log from the primary LUN to the secondary LUN. The principle of I/O processing here is similar to the principle of initial synchronization mentioned above.

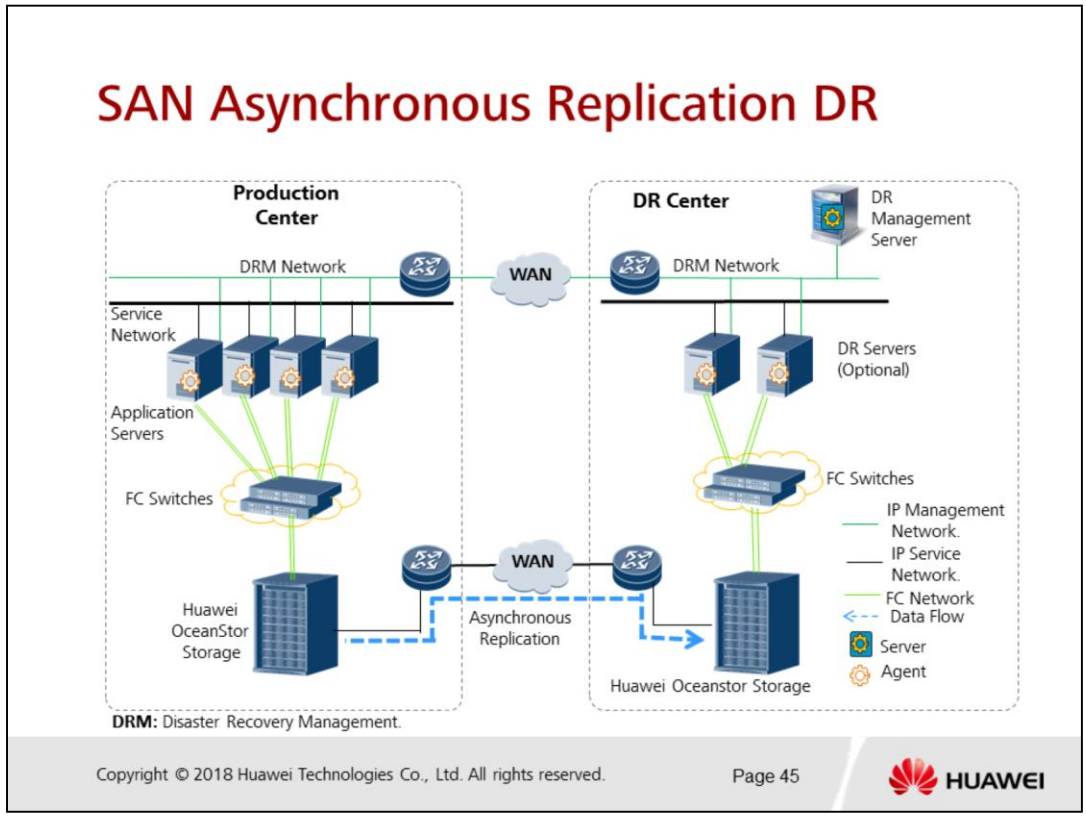

- The deployment method is shown in the figure above. The target RPO is >3s and RTO level is in minutes.
- Different from synchronous replication, asynchronous replication has time-interval replication strategies, and theoretically no limitations in distance.
- RD mainly provides disaster recovery management functions, including replication strategy, topology, disaster recovery testing, drills, and disaster recovery.
- When performing application management and disaster recovery application recovery, the Agent needs to be installed on the server
- The RD management network needs to communicate with the host and storage.
- Seconds-level replication is triggered on production storage, meanwhile replications that can tolerate for more than 15 minutes of delay can be triggered on RD.

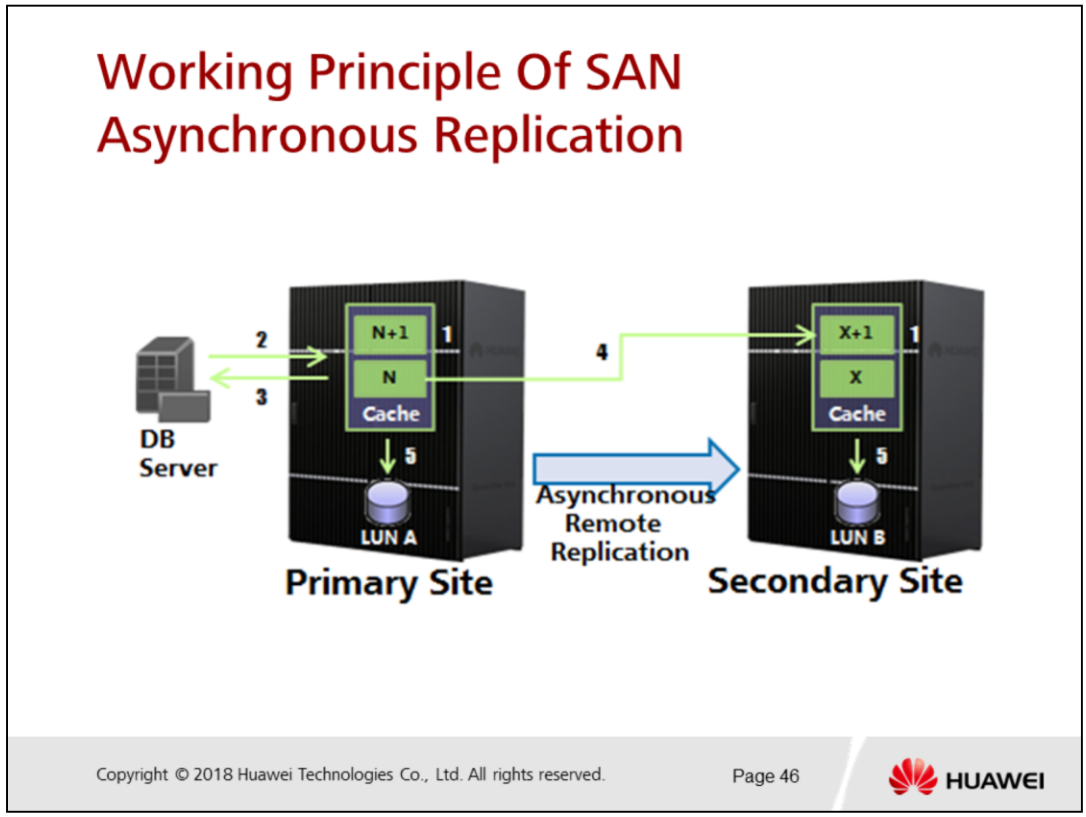

- Time slice: Logical space for writing data in the cache for a period of time (data size is not limited).
- In a low-RPO scenario, the asynchronous remote replication cycle is short. OceanStor storage system caches all data in multiple time slices. If the host service bandwidth or disaster recovery bandwidth is abnormal or fails, the replication cycle becomes longer or interrupted, at this time the data in the Cache will be automatically flushed to provide consistent protection according to the disk flushing policies. It will then read back the flushed data from the disk when the replication resumes.
- Whenever a synchronization period is set (Its set by the user, and the range is from 3s to 1440min), the system will automatically initiate a synchronization process to synchronize the master site data incrementally to the secondary site (if the synchronization type is manual, the user is required to trigger synchronization.). New time slices (TPN+1 and TPX+1) are generated in the cache of the primary LUN (LUN A) and the secondary LUN (LUN B) at the start of each replication cycle.
- The main site receives the production host's write request.

- The main site writes the requested data into the cache time slice TPN+1 and immediately responds to the host on the status of write completion.
- When data is synchronized, the data of the primary LUN (LUN A) cache time slice TPN is read, and transferred to the secondary site, and then written to the secondary LUN (LUN B) cache time slice TPX+1. If the primary site write cache reaches a high level, the data is automatically written from the cache to the hard disk. At this time, the data of the time slice TPN is generated on the disk. During synchronization, the data written to the hard disk is read from the snapshot and copied to the secondary LUN (LUN B).
- After the data synchronization is completed, the data of the time slices TPN and TPX+1 in the primary LUN (LUN A) and the secondary LUN (LUN B) Cache are written to disks (the snapshots are automatically deleted) in accordance with the disk flushing strategy, and waits for the arrival of the next synchronization cycle.

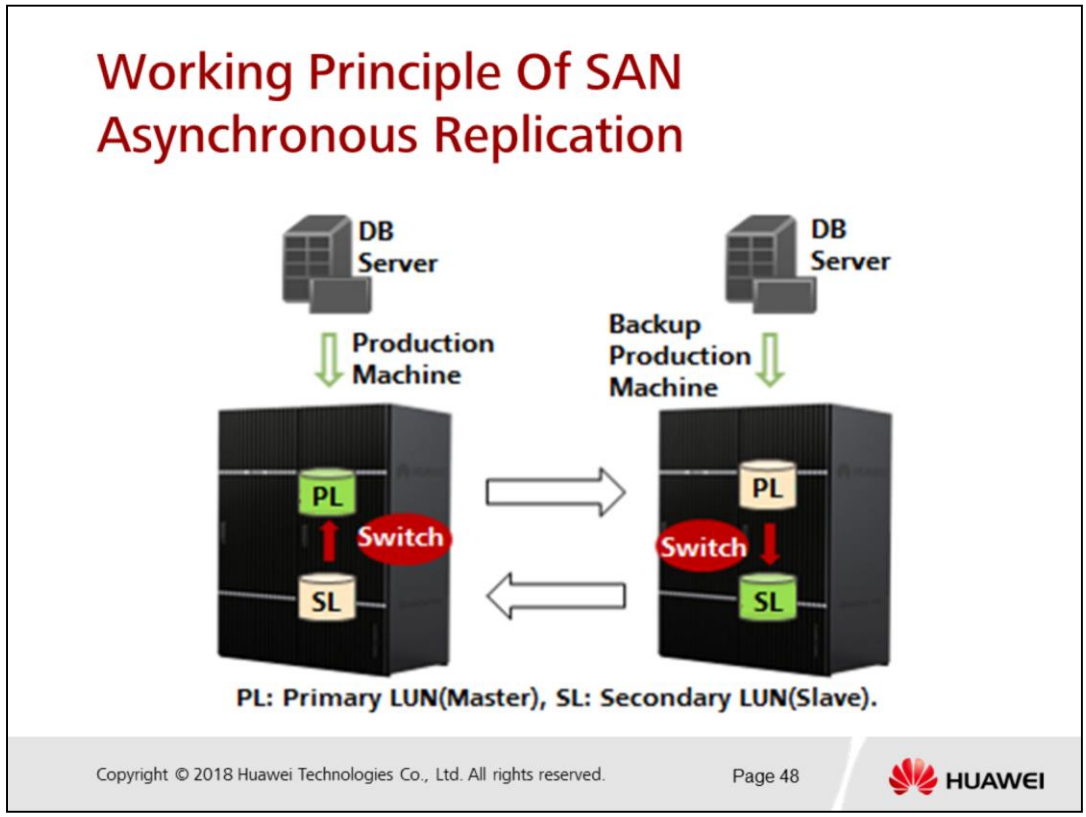

- Switching/Service Failover:
  - Synchronous remote replication can perform master-slave switching under normal conditions.
  - In the split state, you need to set the slave LUN for master-slave switchover.
  - Asynchronous remote replication is in a split state.
  - In the split state, you need to set the slave LUN to be writable.

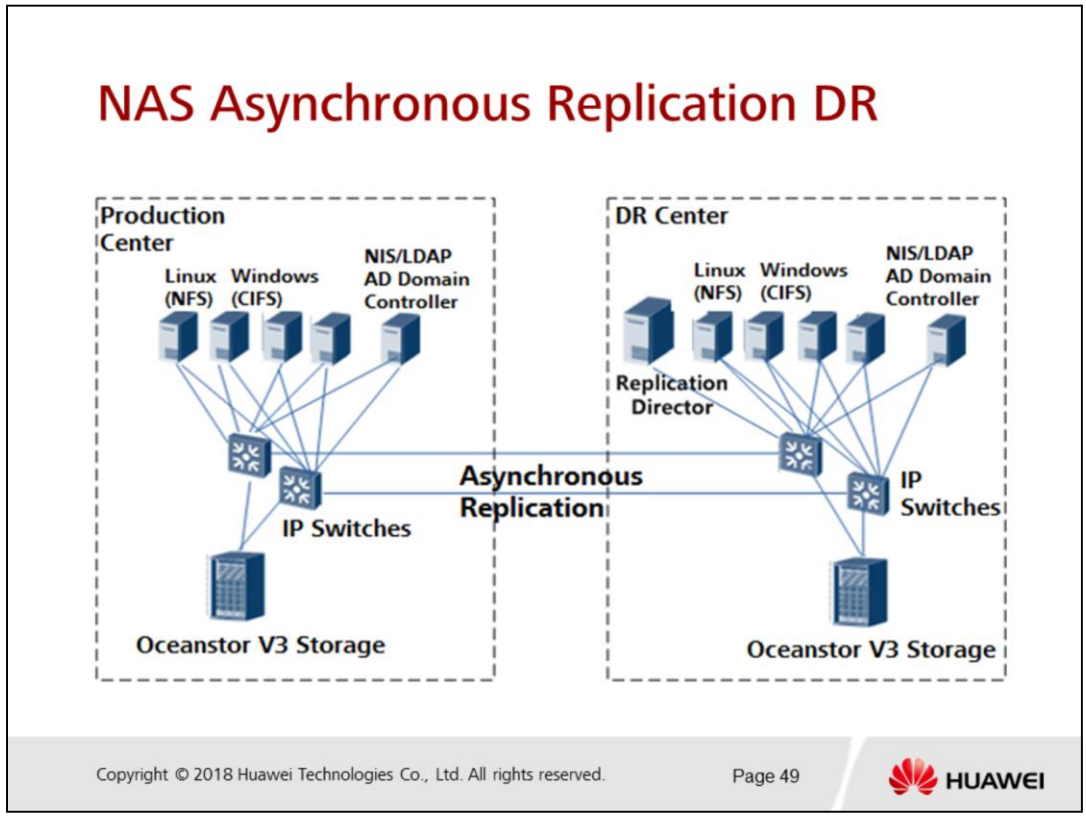

- NAS file system replication is currently available only in the V3R2C10 version and is implemented using ROW.
- The NAS's NAS disaster recovery management does not deploy the Agent on Linux or Windows. It only manages the replication policies and disaster recovery of V3 storage.
- The file system currently mainly supports NFS/CIFS. Disaster recovery management currently manages only the FS replication section, file system, and permission control section. When the system is created, it needs to be configured.
- File system replication is similar to SAN and supports FC/iSCSI links.

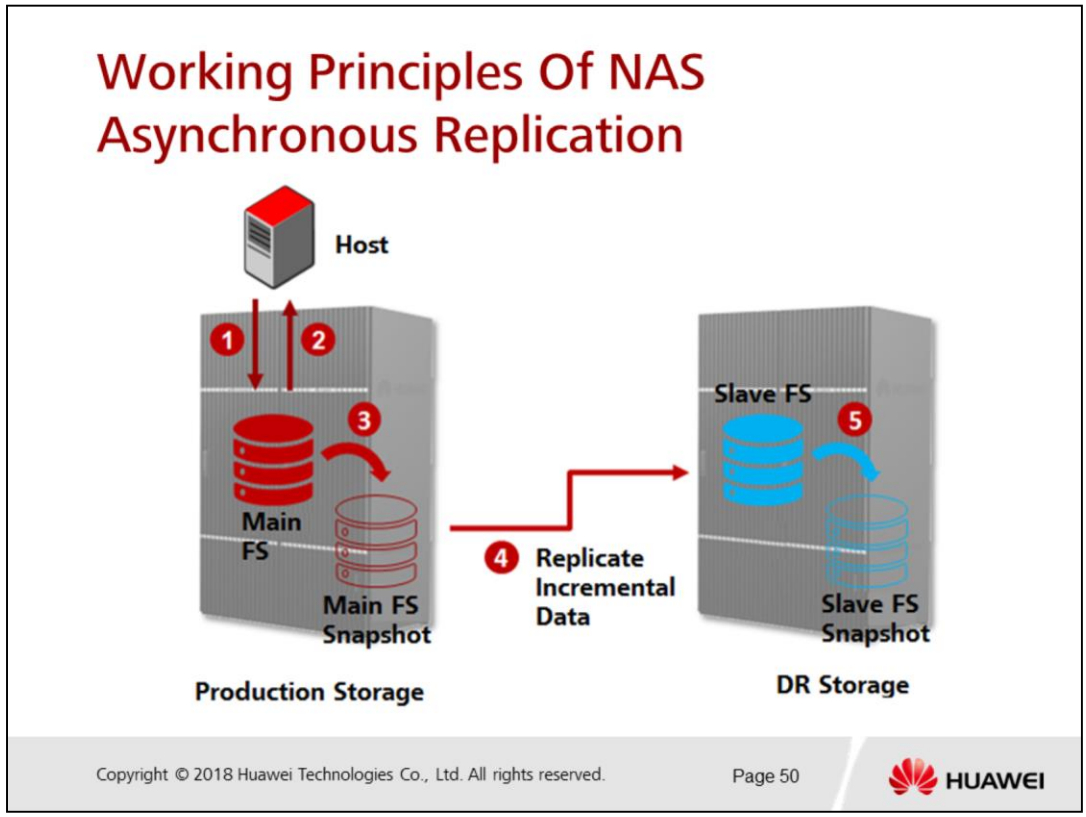

- At the beginning of each cycle, the file system asynchronous remote replication creates a snapshot of the primary FS (primary file system). Based on the delta information during the period from the completion of the previous cycle to the start of the current cycle, the snapshot data is copied to the secondary FS. After the incremental copy is completed, the content of the FS is the same as the content of the snapshot of the primary FS, and the data consistency point is formed on the FS.
- It supports long distance remote replication from one file system to another, but it does not support directory-to-directory or file-to-file remote replication.
- A single file system can only be included in one replication task, but a replication task can include many different file systems.
- File system only supports one-to-one replication, and the same file system cannot be set as the source and destination of the remote replication process. It does not support cascaded replication and also doesn't support 3DC.
- The smallest unit of incremental replication is the file system block size of (4K-64K). The shortest asynchronous replication synchronization cycle period is 5 minutes.
- Supports replication resuming if the connection was interrupted or broken during the replication process.

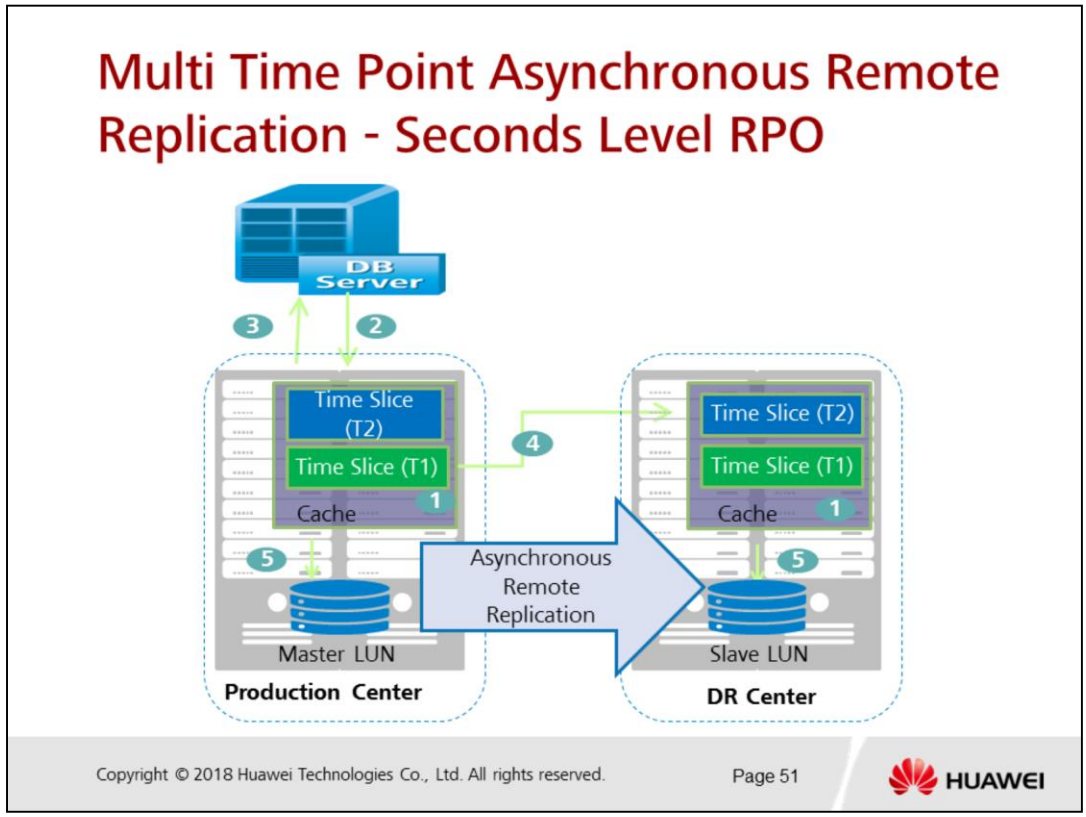

- A minimum of 3 seconds for a consistency point:
  - A new time slot*(T2, P2) is generated in the cache of the master LUN and the slave LUN at the start of each replication cycle.
  - The newly written data of the host is cached in the time slot T2 of the main LUN cache.
  - Response is sent to host on write completion status.
  - The data of the time slice T1 is directly copied to the slave LUN and written to the time slot P2 of the slave LUN.
  - The data received by the primary and secondary LUNs will be written to disks.
    - The copy reads data directly from the cache with a small delay.
    - Snapshots do not require real-time COW(Copy On Write) updates. Synchronization has a small impact on performance and the replication cycles can be reduced to a 3-second replication cycle. This means that the smallest amount of time to wait before the next cycle starts can be set to a minimum 3 seconds.

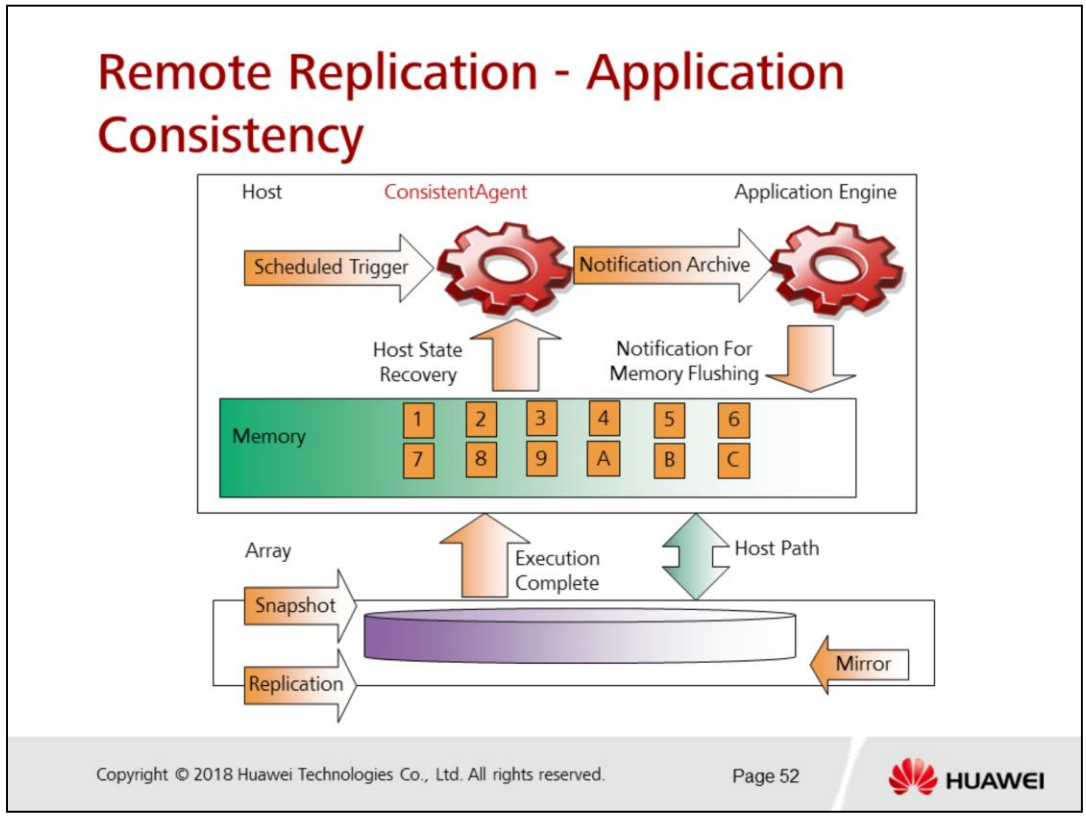

- Application consistency:
  - The consistency agent is installed on the host to implement the linking between the snapshot of the array and the database.
- When the snapshot task executes:
  - First, it puts the database in backup mode, perform checkpoints, and write all the dirty data in memory to the storage system.
  - Then, it informs the array to take a snapshot.
  - Finally, the database is taken out of backup mode.
- Advantages:
  - The disaster recovery end can pull up the data and uses it directly without the need of Roll Forward and Rollback process.

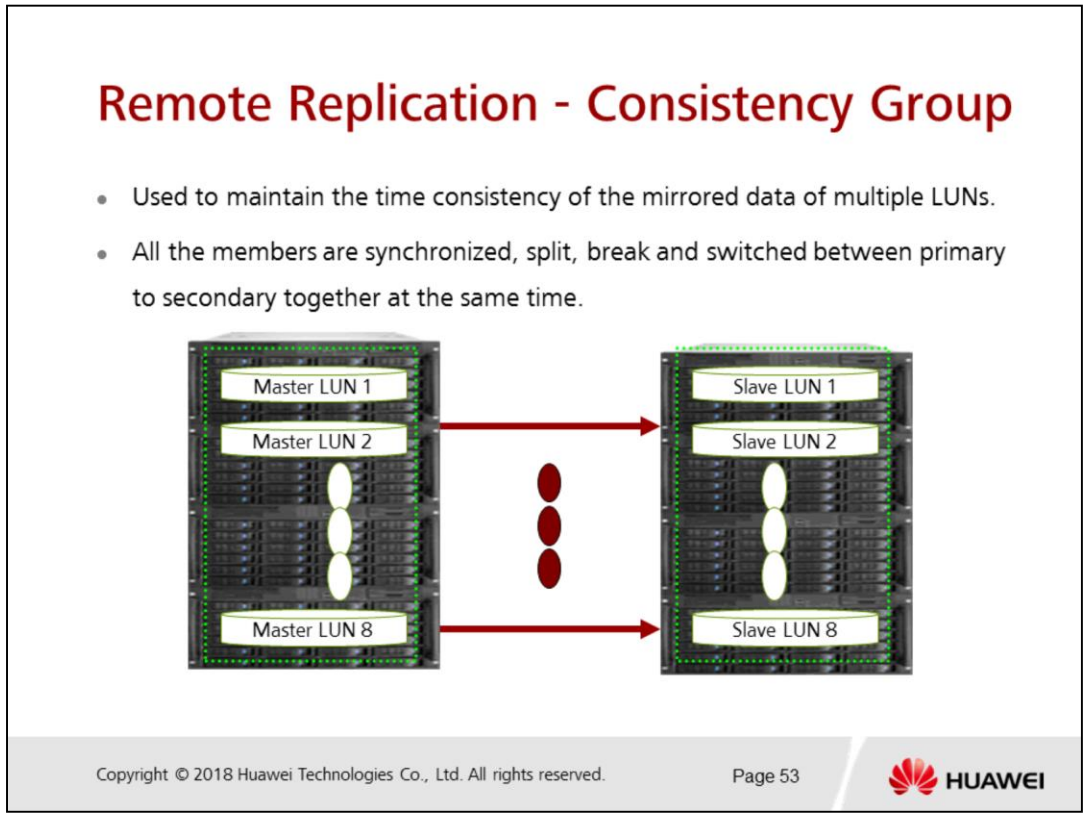

- In large and medium-sized database applications, data, logs, and modification information are stored in different LUNs of the array. Generally, the associated LUN is a non-stand-alone LUN. If the data of one LUN is missing, the data in other LUNs will be invalid. Thus these LUNs must remain consistent in time for the data to be valid and useful in the event of a disaster.
- We hope to simultaneously perform data synchronization or split operations on these LUNs to ensure that the data association between multiple secondary LUNs remains unchanged, thus ensuring the integrity and availability of disaster recovery backup data. The technology introduced to solve this issue is remote replication consistency group technology.
- The maximum number of remote replication pairs in the remote replication consistency group of the Huawei array is 8, and cross-array consistency groups are not supported.
- Note: Remote replication of associated LUNs should be placed in a consistency group. LUNs without an association should not be placed in a consistency group. In addition, synchronous remote replication and asynchronous remote replication cannot be placed in the same consistency group. All slave LUNs for remote replication must be on the same remote storage system.

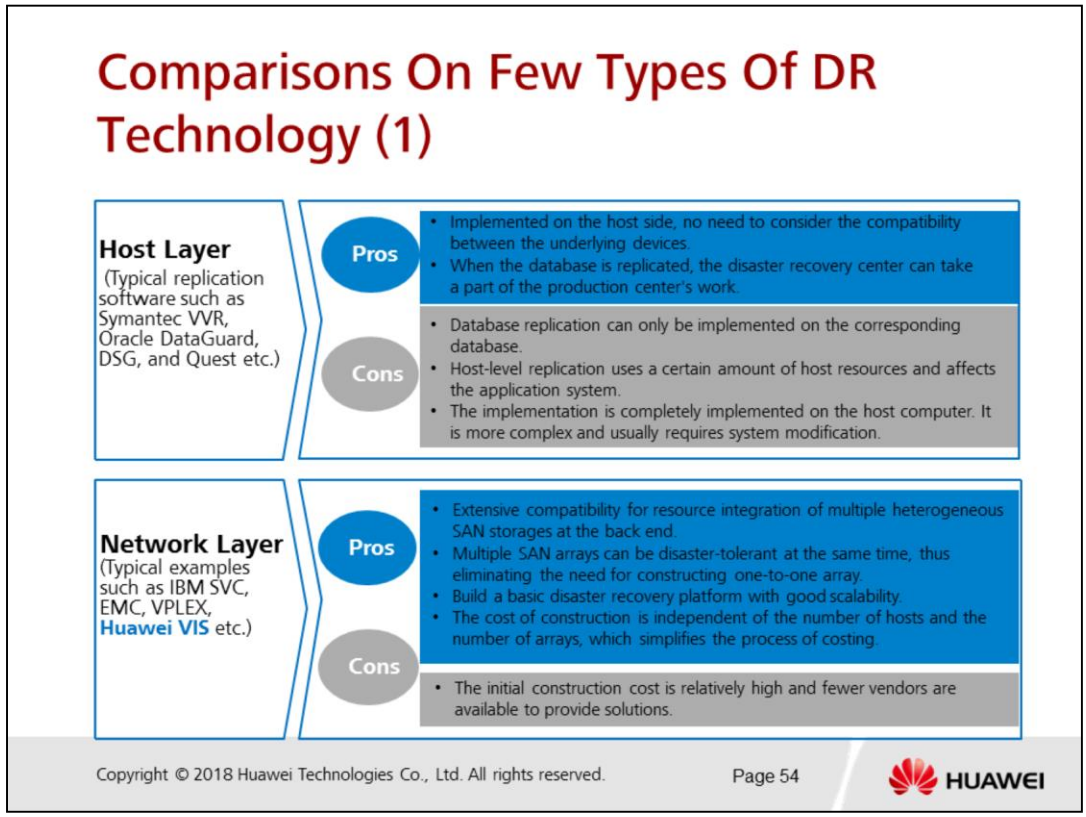

- The table above shows the comparison of DR technologies based on the host and network layer implementation.

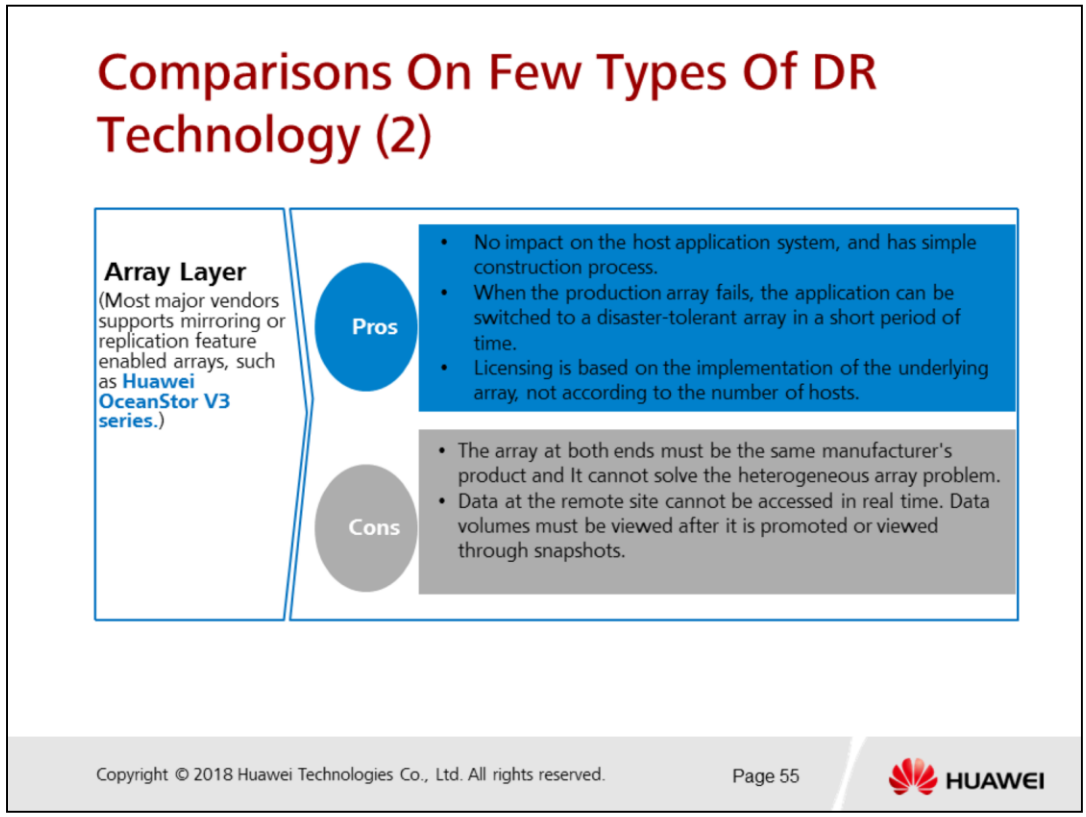

- The diagram above shows the pros and cons of array layer DR technology compared to the earlier host layer and network layer DR technologies.

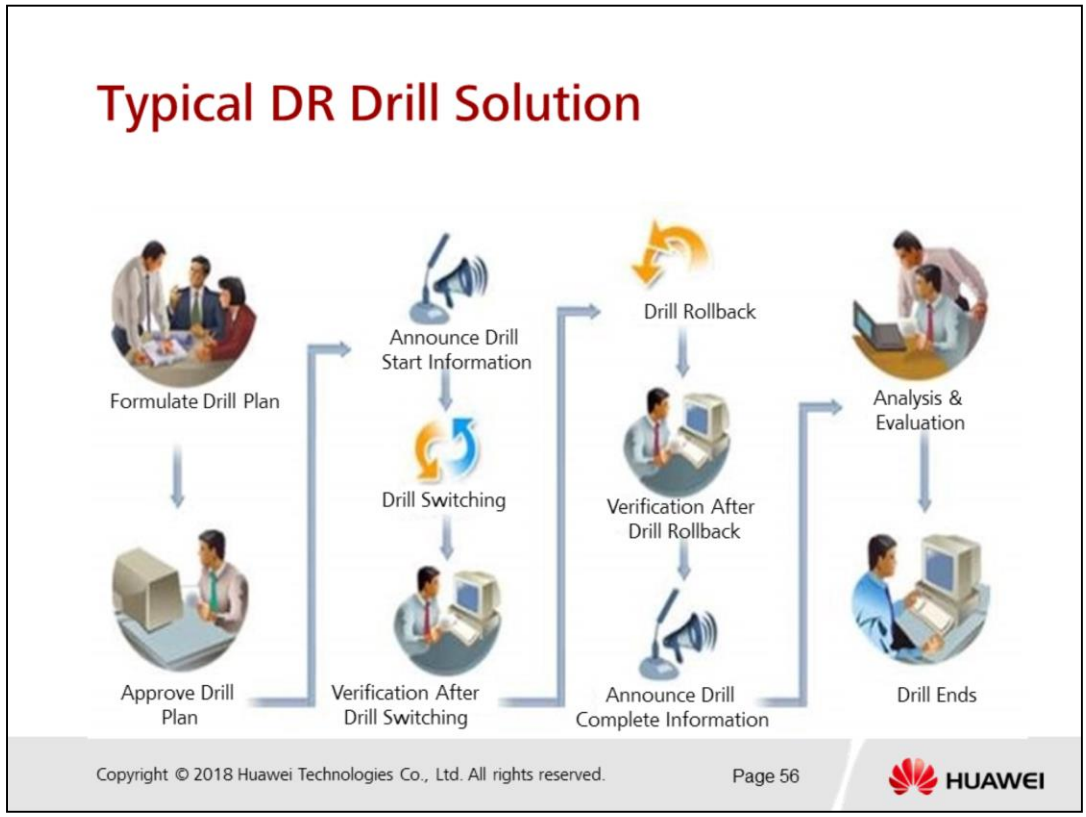

- Features and Benefits of Disaster Recovery Drill Exercise Process:
  - Real Disaster Recovery steps drill exercise.
  - Through Disaster Recovery Drill Exercise, familiarize with the Disaster Recovery Process and improve the Service Recovery Ability in the event of disaster.
  - Through disaster recovery drill exercises, check the business service and data integrity to properly prepare for handling the actual event of disasters.

Contents

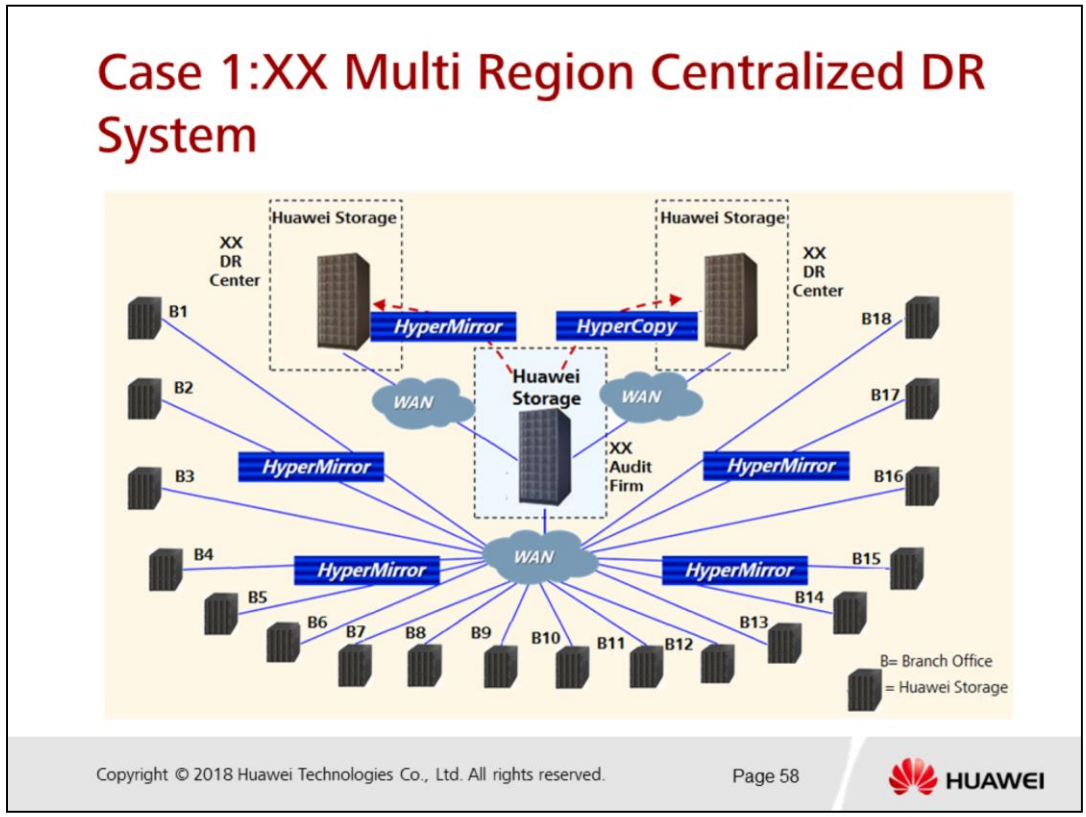

- The challenges faced:
  - The 18 special offices of the XX Audit Firm are distributed in various provinces and have a wide geographical distribution. The investment in disaster recovery and construction is extremely demanding and the implementation of the later projects is difficult.
  - The technical expertise of the various agencies within the firm is weak, making it difficult to maintain and perform DR Drill exercises of the disaster recovery system.
  - The firm has large amount of data, and high demand for disaster recovery network.
- Huawei Solution:
  - All the 18 branches were deployed with a storage system, and the disaster recovery is centralized to the disaster recovery center in Changsha. The whole DR system deployed 21 sets of medium-to-high-end arrays with a total deployment capacity of more than 1PB.
  - Utilized Huawei's 32:1 remote replication technology to achieve centralized disaster recovery on the existing network.
  - The DR Drill software is customized at the disaster recovery center in Chang XX City to meet the requirements of users for the maintenance and drill of the disaster recovery system.

- Value to Customer:
  - The branches and offices has dual benefits of achieving centralized storage, and achieving remote disaster recovery at the same time, which effectively safeguards the data security.
  - The original plan required more than 40 sets of storage system to complete the disaster recovery system, but with Huawei solution it only used 21 sets to complete, saving more than 40% of the investment.
  - With a customized disaster recovery management system, the entire network disaster recovery system is centrally managed and regularly drilled by the Audit Firm’s Information Center, eliminating the worry of management.
  - Utilizing Huawei's global service organizations ensured the progress of the entire system and the business services went live on time.

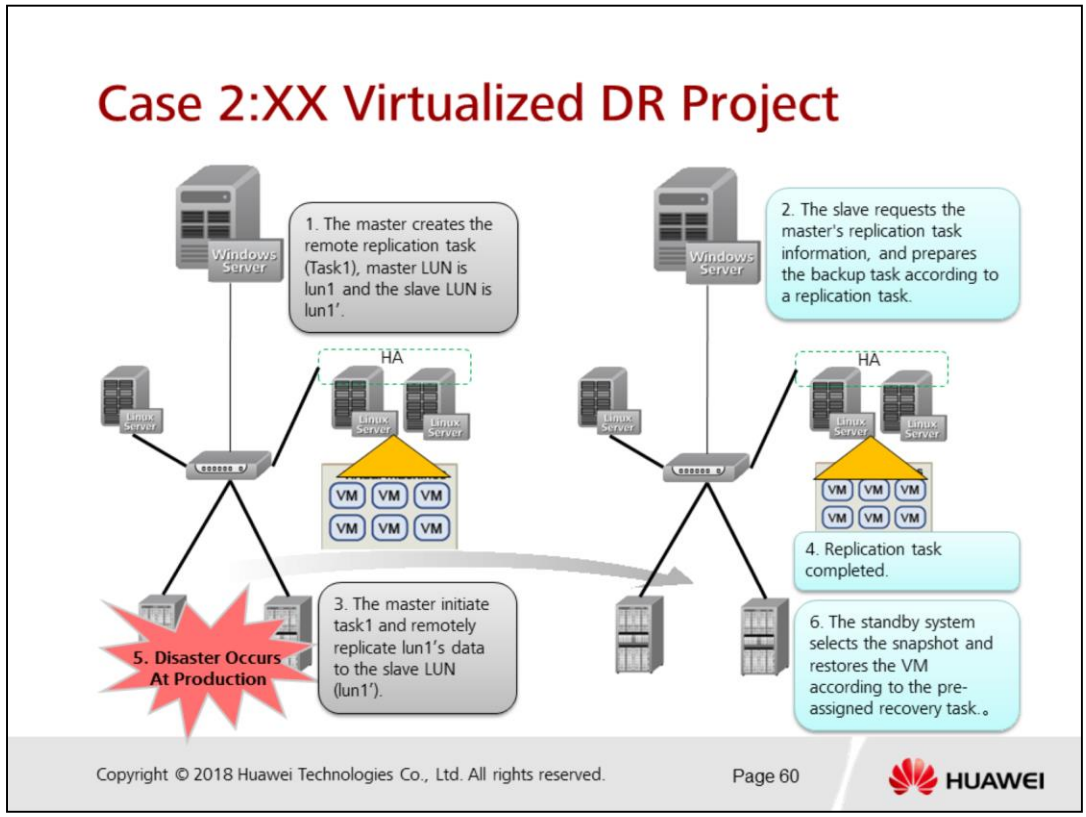

- The challenges faced:
  - The customer already has a vSphere virtual data center and needs to build a new data center for disaster recovery.
  - Customer requires Low TCO, and high return on investment.
- Huawei Solution:
  - Deploy a set of IT systems in disaster recovery centers, including storage, services, networking, and virtualization platforms.
  - Huawei UltraVR disaster recovery components is installed in production and disaster recovery centers.
  - The host of the virtual machine, which is installed with the consistency agent called "ConsistentAgent", that implements application-level protection of the virtual machines.
- Value to Customer:
  - There is no need to change the existing network architecture.
  - Flexible setting of disaster recovery strategy and policies, with one-click recovery of the systems.
  - Support disaster recovery drill exercise and disaster recovery switching/failovers.

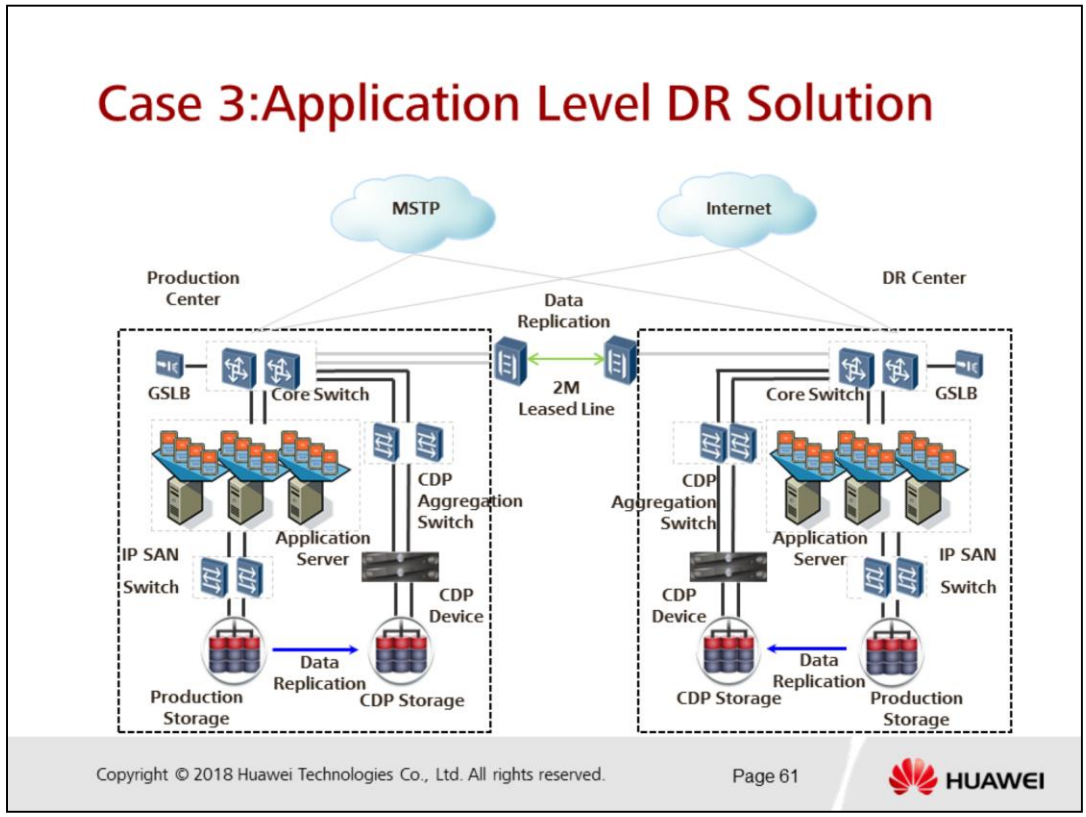

- The challenges faced:
  - The current state of IT is difficult to support business development needs, and it is difficult to ensure the continuous operation of online business.
  - IT operation and maintenance complexity, high energy consumption, and low resource utilization.
- Huawei Solution:
  - Move the business systems to Huawei Cloud Platform.
  - Deploy CDP(Continuous Data Protection) storage and CDP software in two data centers and use CDP technology to implement application-level disaster recovery for both data centers in the same city.
- Value to Customer:
  - Resilient resources, reuse of resources, improve resource utilization, and reduce operation and maintenance costs.
  - The key business RTO and RPO are 0. When the production center fails, services and data are automatically switched to the disaster recovery center.
  - Guaranteed business continuity.

- Continuous data protection (CDP), also called continuous backup or real-time backup, refers to backup of computer data by automatically saving a copy of every change made to that data, essentially capturing every version of the data that the user saves. It allows the user or administrator to restore data to any point in time.
- MSTP (Multiple Spanning Tree Protocol), refers to the network protocol that allows multiple instances of the Spanning Tree Protocol (STP) and are mainly used to create a loop-free topology in networks with multiple spanning-tree regions.

Summary

- This module mainly let us learn about:
  - The concepts of DR, and the importance of constructing the DR systems.
  - The common DR solutions and the advantage and disadvantages of these solutions.
  - The different technologies used to implement these DR solutions and the typical DR application cases.

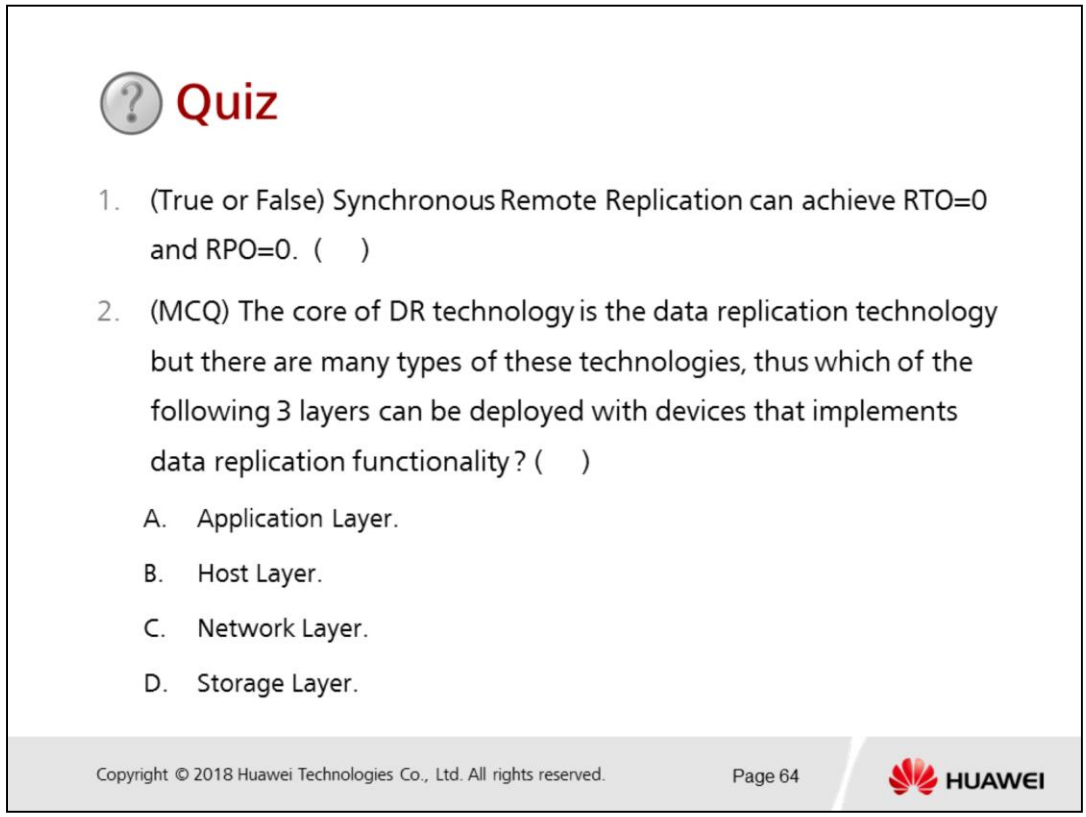

- Answers:
  - False.
  - BCD.

Thank You

www.huawei.com