509.4
Google Cloud Platform

SANS

THE MOST TRUSTED SOURCE FOR INFORMATION SECURITY TRAINING, CERTIFICATION, AND RESEARCH | sans.org

PLEASE READ THE TERMS AND CONDITIONS OF THIS COURSEWARE LICENSE AGREEMENT ("CLA") CAREFULLY BEFORE USING ANY OF THE COURSEWARE ASSOCIATED WITH THE SANS COURSE. THIS IS A LEGAL AND ENFORCEABLE CONTRACT BETWEEN YOU (THE "USER") AND SANS INSTITUTE FOR THE COURSEWARE. YOU AGREE THAT THIS AGREEMENT IS ENFORCEABLE LIKE ANY WRITTEN NEGOTIATED AGREEMENT SIGNED BY YOU.

With the CLA, SANS Institute hereby grants User a personal, non-exclusive license to use the Courseware subject to the terms of this agreement. Courseware includes all printed materials, including course books and lab workbooks, as well as any digital or other media, virtual machines, and/or data sets distributed by SANS Institute to User for use in the SANS class associated with the Courseware. User agrees that the CLA is the complete and exclusive statement of agreement between SANS Institute and you and that this CLA supersedes any oral or written proposal, agreement or other communication relating to the subject matter of this CLA.

BY ACCEPTING THIS COURSEWARE, YOU AGREE TO BE BOUND BY THE TERMS OF THIS CLA. BY ACCEPTING THIS SOFTWARE, YOU AGREE THAT ANY BREACH OF THE TERMS OF THIS CLA MAY CAUSE IRREPARABLE HARM AND SIGNIFICANT INJURY TO SANS INSTITUTE, AND THAT SANS INSTITUTE MAY ENFORCE THESE PROVISIONS BY INJUNCTION (WITHOUT THE NECESSITY OF POSTING BOND) SPECIFIC PERFORMANCE, OR OTHER EQUITABLE RELIEF.

If you do not agree, you may return the Courseware to SANS Institute for a full refund, if applicable.

User may not copy, reproduce, re-publish, distribute, display, modify or create derivative works based upon all or any portion of the Courseware, in any medium whether printed, electronic or otherwise, for any purpose, without the express prior written consent of SANS Institute. Additionally, User may not sell, rent, lease, trade, or otherwise transfer the Courseware in any way, shape, or form without the express written consent of SANS Institute.

If any provision of this CLA is declared unenforceable in any jurisdiction, then such provision shall be deemed to be severable from this CLA and shall not affect the remainder thereof. An amendment or addendum to this CLA may accompany this Courseware.

SANS acknowledges that any and all software and/or tools, graphics, images, tables, charts or graphs presented in this Courseware are the sole property of their respective trademark/registered/copyright owners, including:

AirDrop, AirPort, AirPort Time Capsule, Apple, Apple Remote Desktop, Apple TV, App Nap, Back to My Mac, Boot Camp, Cocoa, FaceTime, FileVault, Finder, FireWire, FireWire logo, iCal, iChat, iLife, iMac, iMessage, iPad, iPad Air, iPad Mini, iPhone, iPhoto, iPod, iPod classic, iPod shuffle, iPod nano, iPod touch, iTunes, iTunes logo, iWork, Keychain, Keynote, Mac, Mac Logo, MacBook, MacBook Air, MacBook Pro, Macintosh, Mac OS, Mac Pro, Numbers, OS X, Pages, Passbook, Retina, Safari, Siri, Spaces, Spotlight, There's an app for that, Time Capsule, Time Machine, Touch ID, Xcode, Xserve, App Store, and iCloud are registered trademarks of Apple Inc.

PMP and PMBOK are registered marks of PMI.

SOF-ELK® is a registered trademark of Lewes Technology Consulting, LLC. Used with permission.

SIFT® is a registered trademark of Harbingers, LLC. Used with permission.

Governing Law: This Agreement shall be governed by the laws of the State of Maryland, USA.

FOR509.4 | Enterprise Cloud Forensics and Incident Response

Google Cloud Platform

This page intentionally left blank.

This page intentionally left blank.

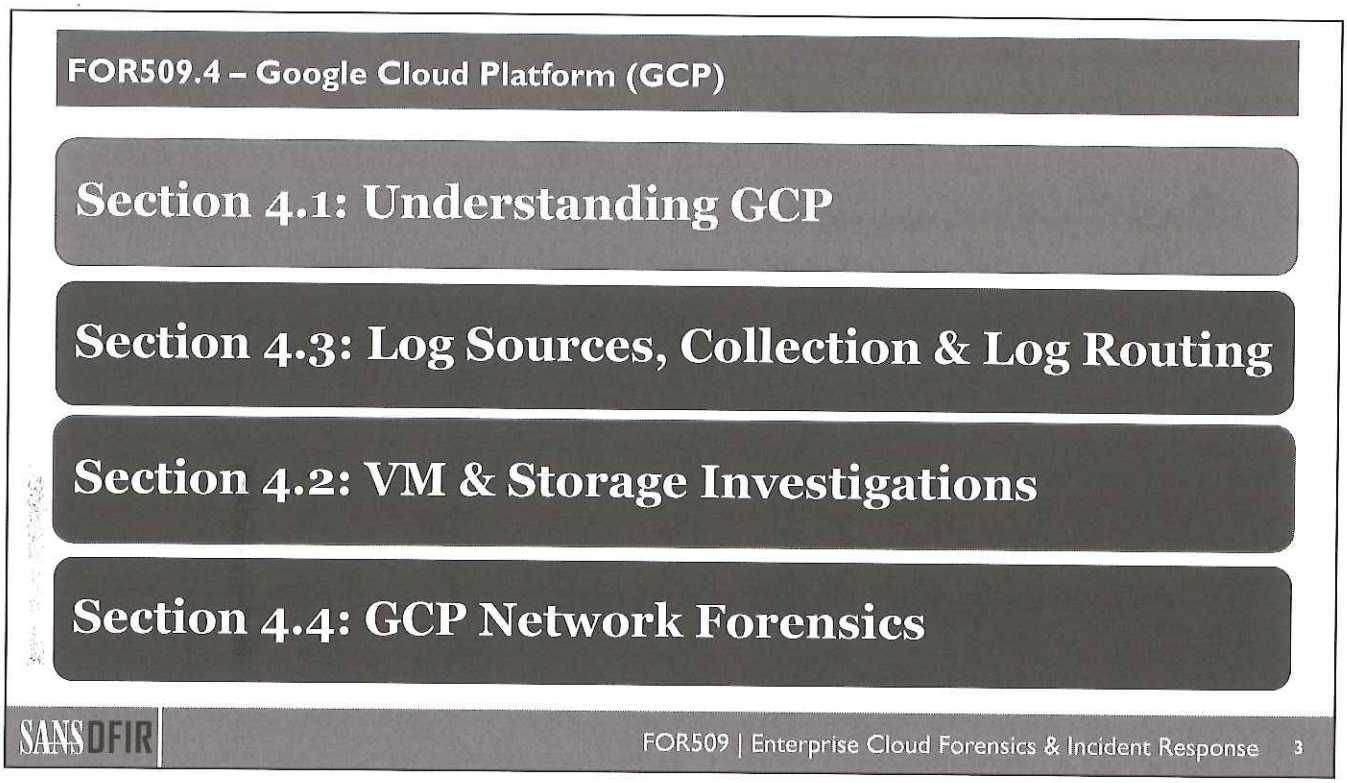

This page intentionally left blank.

Purpose of this Section

- The purpose of this section is to gain a basic understanding of key Google Cloud Platform (GCP) resources and logs to facilitate incident response and digital forensics
- Become familiar with evidence available for virtual machines, networking, and storage
- Understand how platform logs are generated, routed, and stored within GCP
- Understand Log Explorer and importing logs into SOF-ELK
- Discuss various attacks on GCP

GCP is very similar to the other three main Cloud platforms that have been shown already. However, as you have seen with the other Cloud platforms each of them have slightly different terminologies and services that are useful from an investigation perspective. In this section we will examine the logs and services available to investigators. We will also review how logs work in general within GCP and how to alter what is logged and how to forward, or ship, those logs to other locations.

This section will again focus on extracting logs out of the Google Cloud Platform and moving them into SOF-ELK for analysis. While there are many different scenarios this could work in practice, we'll continue to use SOF-ELK to show the overall principals which would allow investigators in the field to adapt these techniques as needed for their environment.

Additionally, we will also look at some common attacks that can occur in GCP. While this will not be an extensive list, it is intended to again show you techniques for leveraging GCP's infrastructure and services so you can adapt this as needed in the field.

By the end of this section, you will have a strong understanding of the logging capabilities in GCP. This will provide you a solid foundation to conduct incident response and forensics in your environment.

Why We're Covering What We're Covering

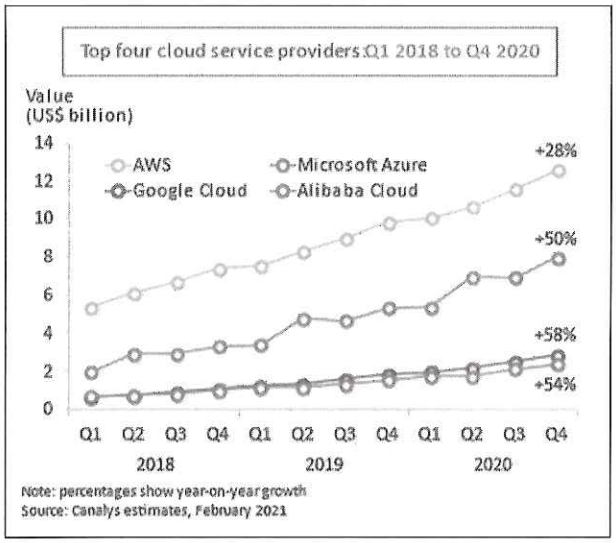

- GCP is the third largest cloud provider after AWS and Azure
- While the third largest cloud is has seen the largest growth in the past year compared to AWS or Azure.
- As with other Cloud platforms organizations are forecast to grow their cloud spend by 18.4% in 2021
- This sets up Cloud infrastructure to be in the cross-hairs of threat actors

GCP is currently seen as the third most used cloud platform as of Q4 in 2020[1], behind AWS and Azure. However, it is the cloud platform that is seeing the fastest increase in customers making the fastest growing platform as well[1].

Cloud platforms as a whole are seeing an increase growth, with end-users forecast to grow their use of the cloud by 18.4% in 2021[2]. Although SaaS (Software as a Service) is the largest growing area within the cloud, other services such as PaaS (Platform as a Service) and IaaS (Infrastructure as a Service) as followed closely behind[2]. It is due to this we try to cover a cross sections of these platforms within this section with a focus on the areas that Enterprises are most focused on from an incident response perspective.

References:

[1] https://for509.com/xr0ai

[2] https://for509.com/yx381

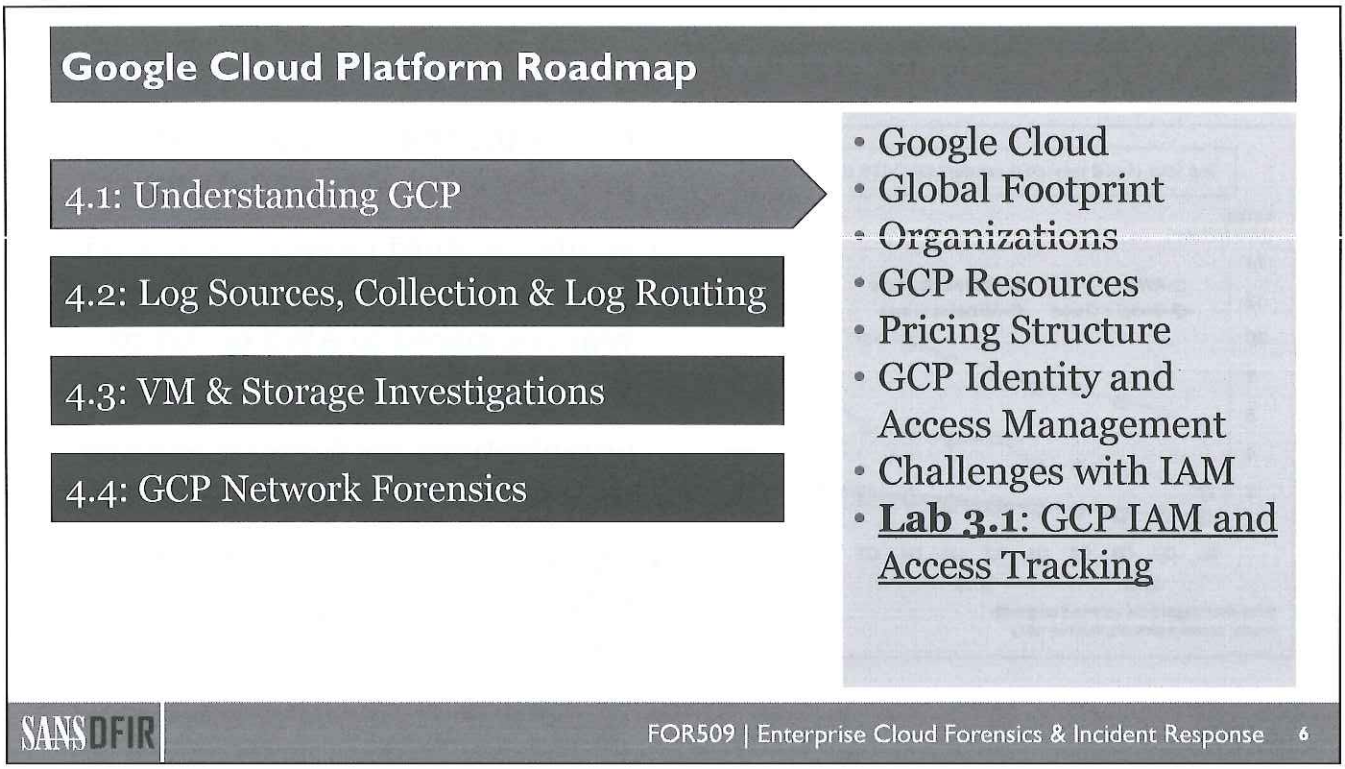

This page intentionally left blank.

Google Cloud

"Today, most of the world's enterprise computing still happens on-premise. It hasn't moved to the cloud yet, because the path forward is complex and daunting, and full of difficult decisions. How do you modernize in-place without having to jump completely to the cloud? How do you bridge incompatible architectures while you transition? And how do you maintain flexibility and avoid lock-in?"
Sundar Pichai, Google's CEO (2019)

Key services for incident response and forensics:

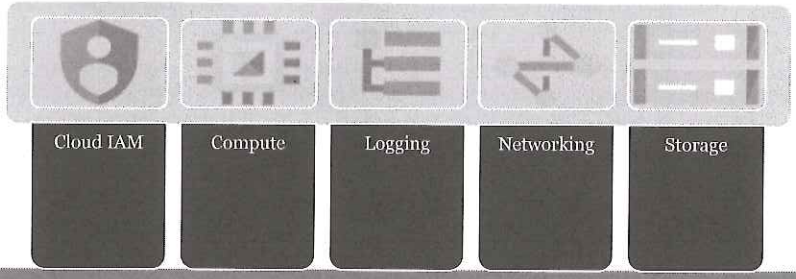

Before we can deep dive into digital forensics in GCP we need to discuss some of the fundaments to GCP to ensure everyone understands the terms and language used to describe different items and objects within GCP. Once we have some of these fundaments covered, we're going to focus on the five services that offer the most useful digital forensics value when in incident has occurred within GCP, this will include;

- GPC Identify and Access Management
- GCP Compute or also known as Virtual Machines
- GCP's Logging Platform
- GCP Networking and the Services within Networking, and
- GCP Storage Buckets

To start with, let's go back to 2019 when the current Google CEO, Sundar Pichai, addressed the Google Cloud Next 2019 Conference. During this address Sundar presented to the attendees a challenge with cloud computing as a whole – how do you produce a cloud platform without locking in customers and making it continually flexible to be used. It is this challenge that Sundar presented that you will see appear within the GCP platform to ensuing services are flexible to the end-user, this will become obvious when we start to look at host OS logging that does not even need to exist within GCP for the GCP Logging Platform to use it.

References:

[1] Sundar Pichai, Google's CEO, Google Cloud Next 2019 Conference in April 2019

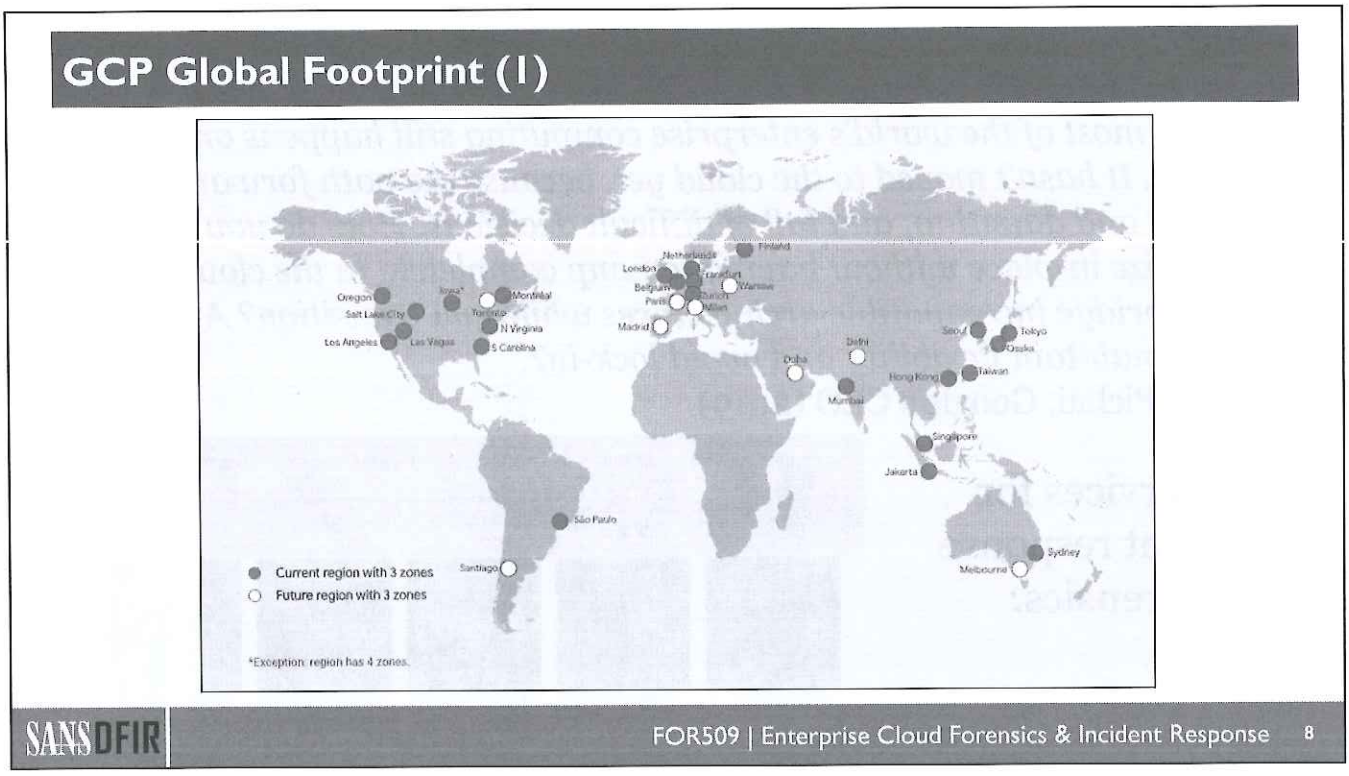

As for GCP's overall global footprint, at the time of publishing, GCP had 25 different regions which included 76 zones. Even at the time of publishing this course, GCP currently has listed additional countries and regions which will likely be in production by the time you are reading this.

While the majority of core services are in each zone and region, some specific services aren't included in every region. This can be for a verity of reasons, including licensing issues within a region, physical space to introduce additional equipment to cater for more specialized services.

While GCP's overall platform will seem relatively transparent when you're looking at the console, to be mindful of any data residency issues or regulation issues with where you are storing and processing data. This in particular is important when you're moving data around for digital forensics or evidence processing.

Reference:

[1] https://for509.com/hqvn5

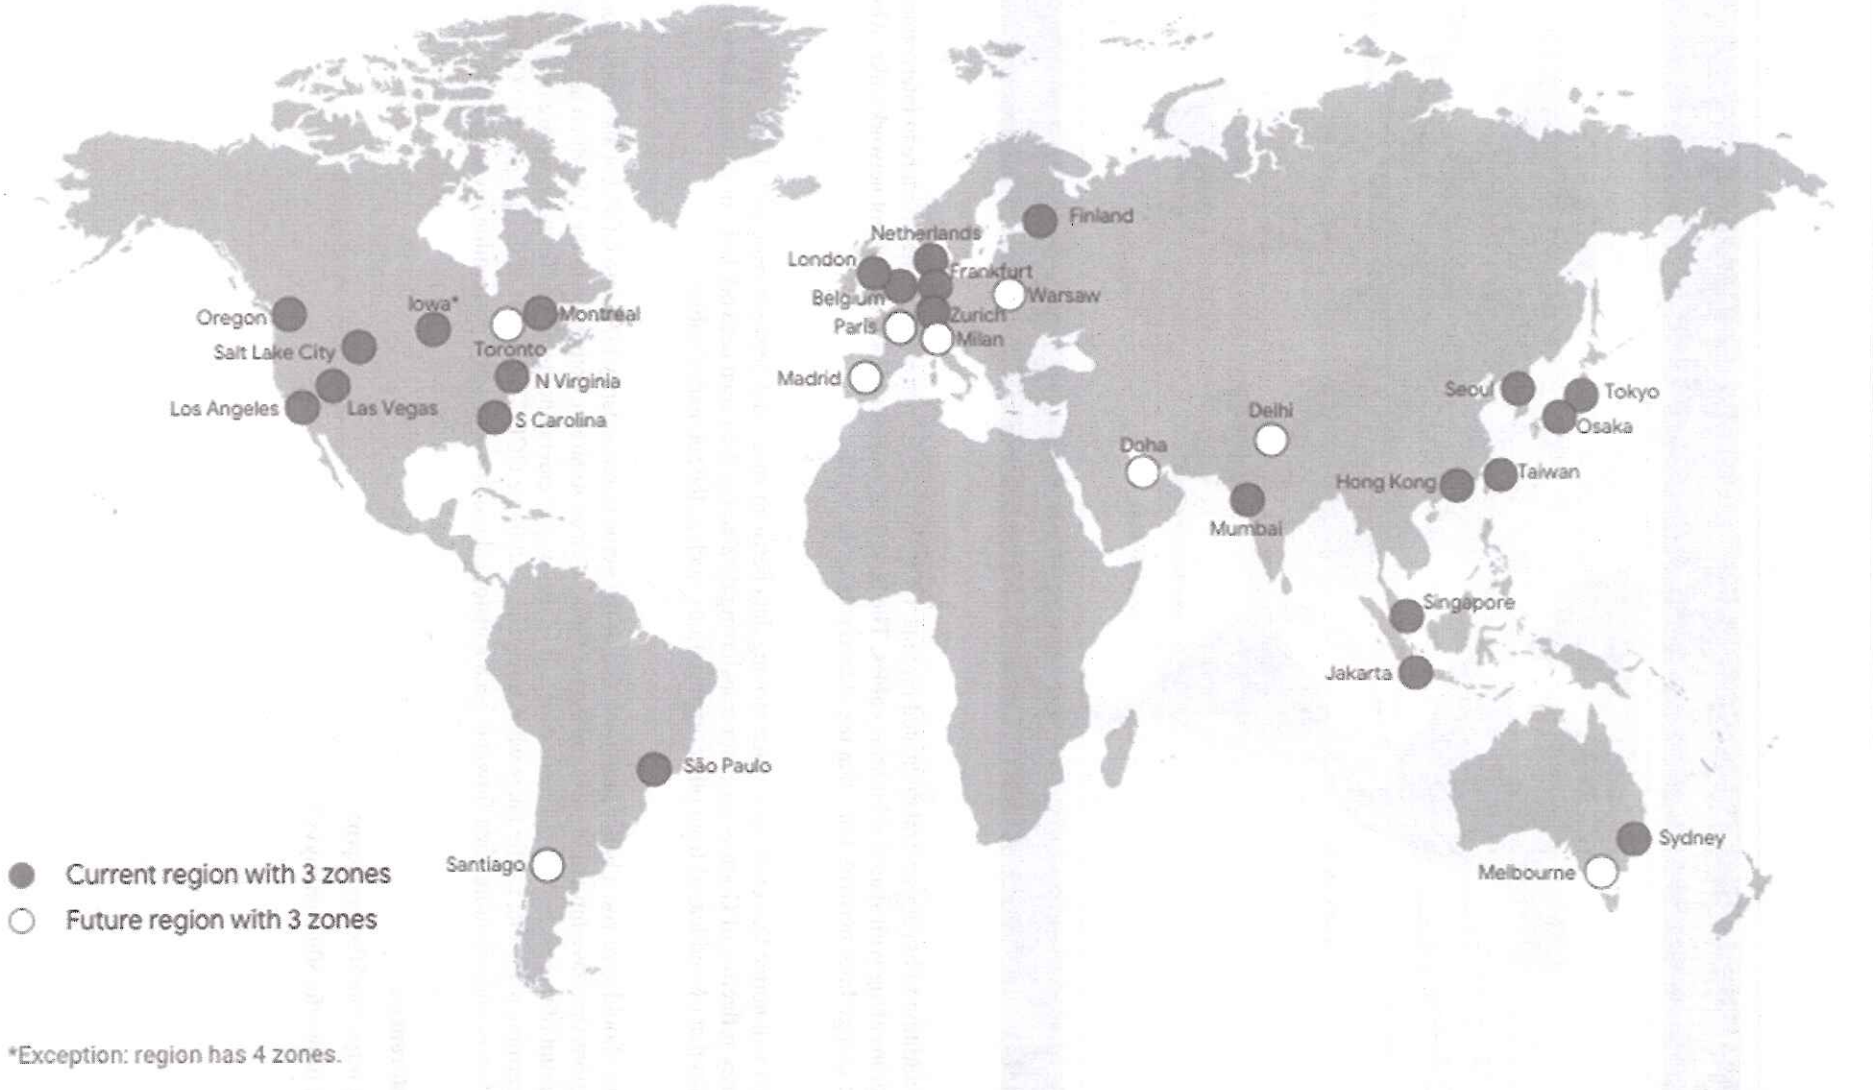

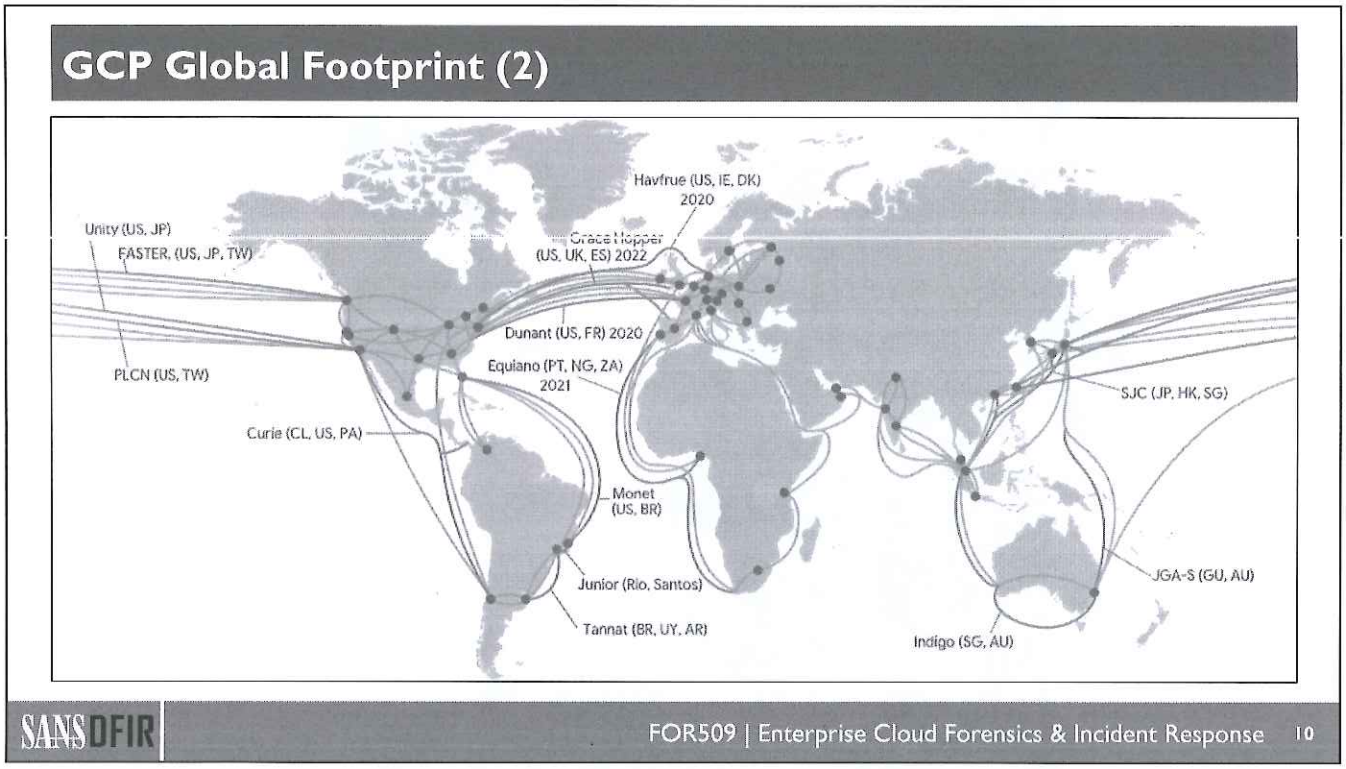

In addition to having several dedicated regions, GCP also uses its networking infrastructure to interconnect regions along with shared submarine cables. The blue lines represent Google's current network links, while the orange lines network links that use shared submarine cables.

You will notice that costs vary when moving data between zones and between regions[2]. These costs are partly reflective of Googles costs for transferring data across their own network links or utilizing network links that owned/leased from other organizations, such as the submarine cables.

This should give you a better graphical idea of the network redundancy built into GCP, but also show you regions that have limited ingress and egress links, such as Australia, compared to the US which has a large amount of ingress and egress redundancy. It is also worth understanding that as for now, we are unable to determine paths that traffic takes when traversing within the GCP. This is part of the design of a "Cloud" however understanding data flow has been a historical requirement when performing investigations.

References:

[1] https://for509.com/hqvn5

[2] https://for509.com/myvs2

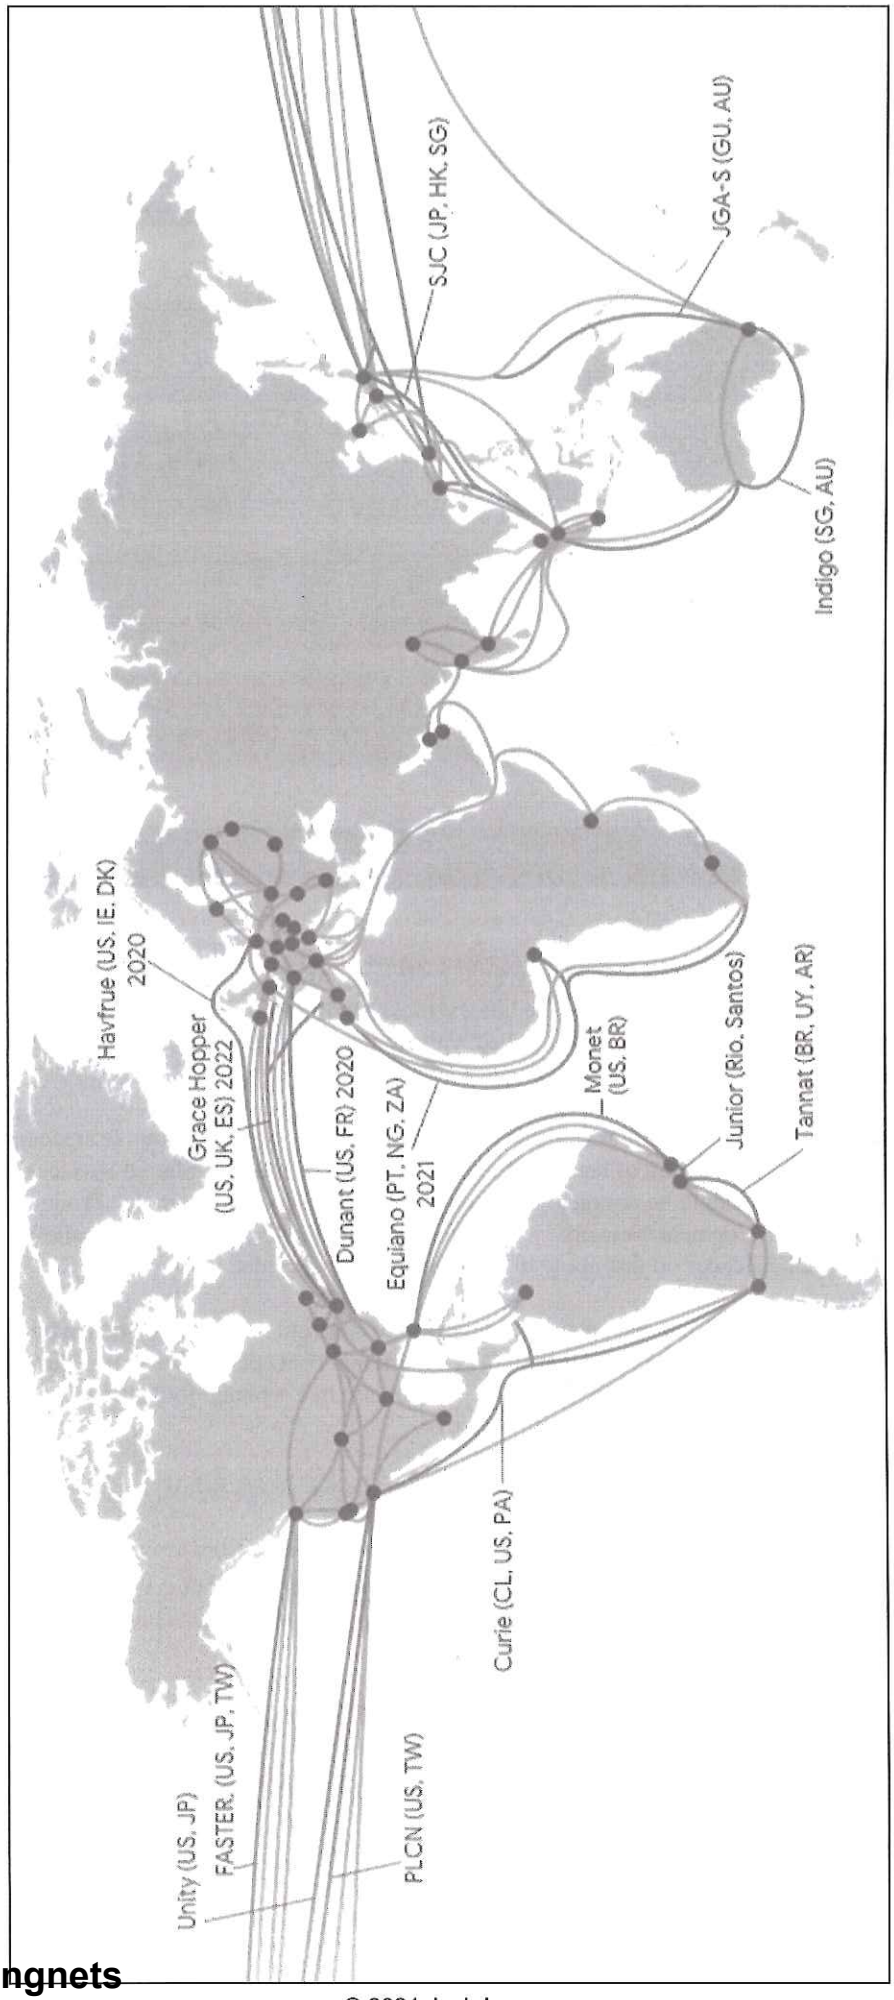

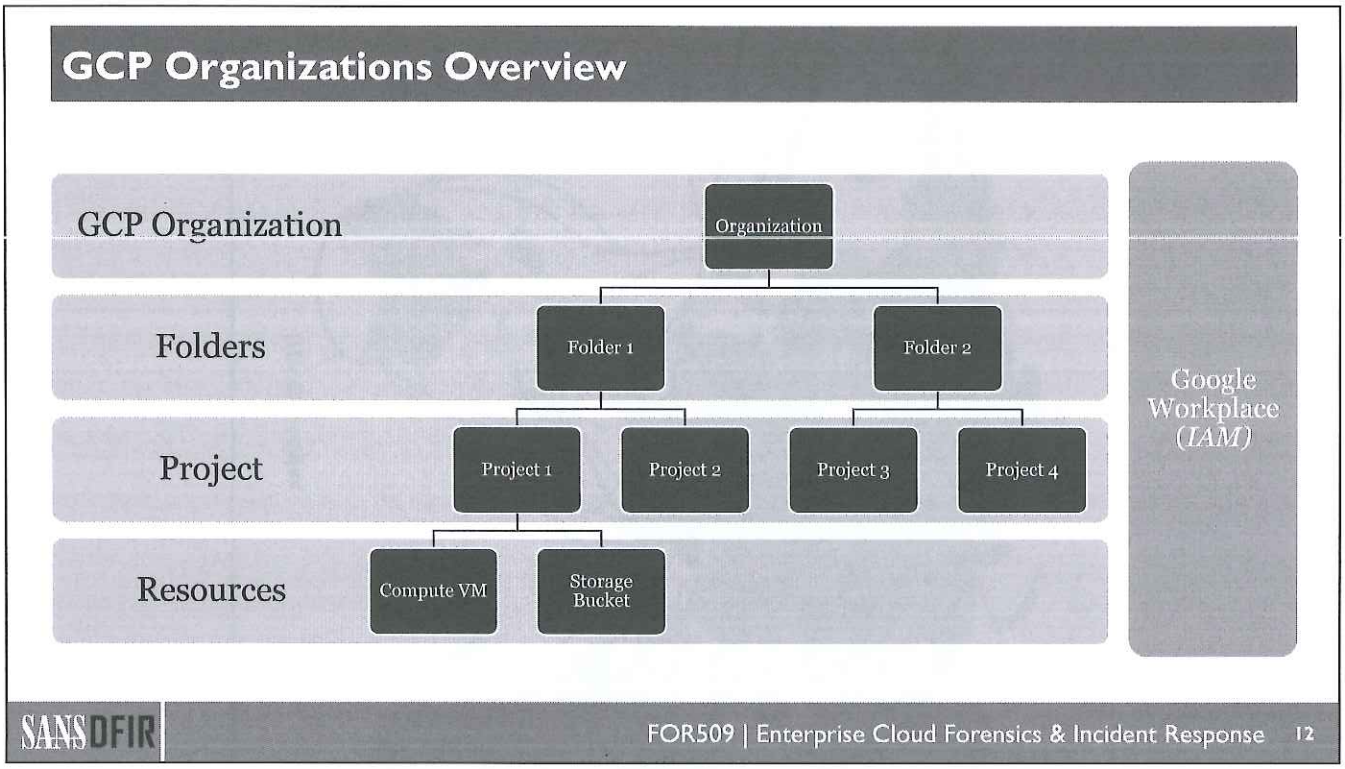

The GPC Organization is the root of how objects and services are assigned and ordered within a GCP instance. The "**Organization**" is considered the highest most parent in the overall structure. This object is usually a domain main (example.com), and if linked with Google Workspace, would also be the primary domain that is used in your Google Workspace.

Underneath an Organization are "**Folders**", these are used to logically order or group teams/departments together. These could be in the form of teams, for example an "R&D Team", "Shared Services", and "Customer Services". Or it could be arranged based on logical boundary, for example "Production Network", "Staging Network", "Development Network". The grouping of folders is significant as it allows an Organization to set policies later on that apply to everything within a folder.

For a digital forensics and incident response team, you would expect to have your evidence processing services within a folder separate form all other folders. Additional concepts on this will be discussed further into this section, however, understand that if your DFIR team are in a separate folder they can have polices separated from other teams which maybe more restrictive.

You can have multiple layers of folders to also give you sub-folders, or sub-teams, if you need to further break down the structure of teams within GCP.

"**Projects**" are logical groups of services within GCP that are run by a team that is within a Folder. Projects are the item you group your running services within GCP to. Projects are also used for billing purposes to track spend on the services consumed within GCP.

The lowest level of objects within an Organization are the GCP "**Resources**", these are the services that you use within GCP and generally have a cost associated with them. Resources inherit permissions from policies that are applied to Folders and Projects. The use of Resources are billed against a Project.

References:
[1] https://for509.com/gt8s9

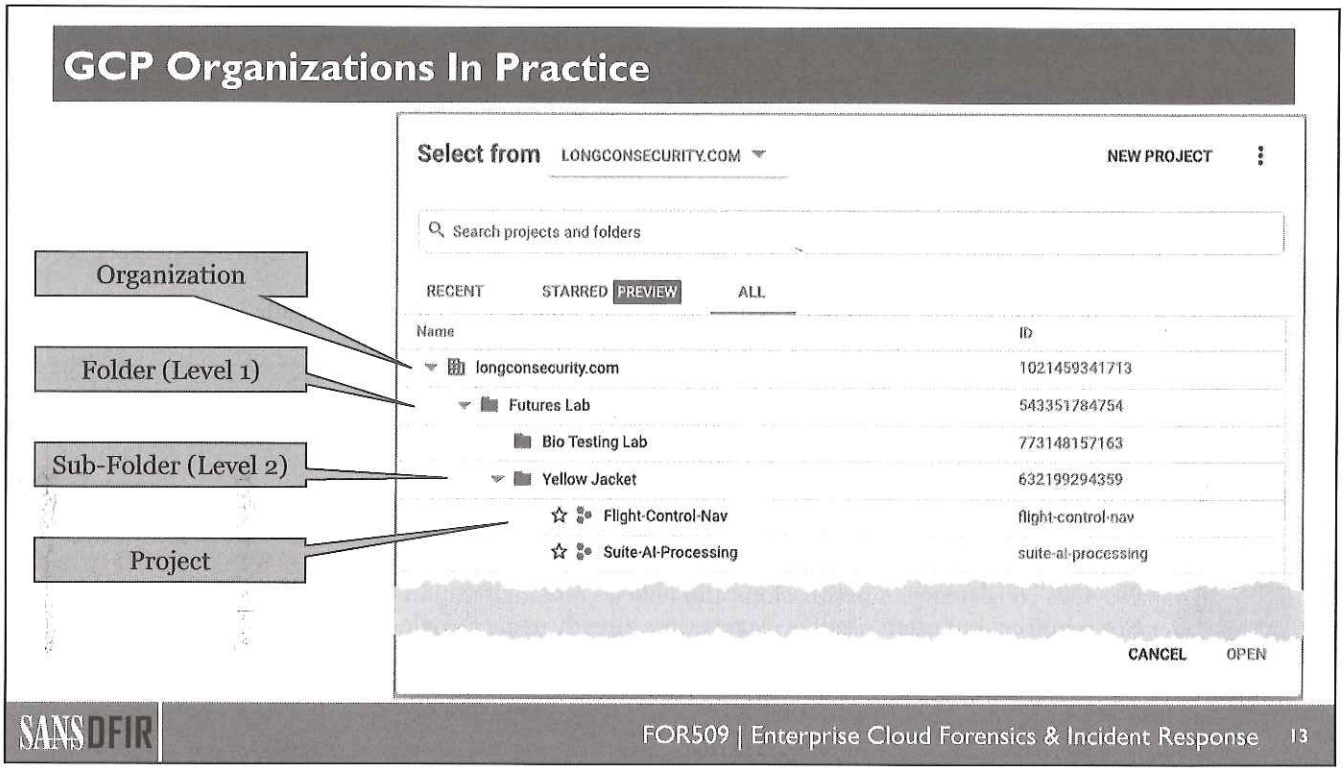

This example shows what a GCP Organization would look like within the GCP console. Conceptually the hierarchy is similar to what we have just discussed in the previous page. However, this example also shows you sub-folders as well. You can see from this example:

Organization – *longconsecurity.com*

Folder (Level 1 – *Futures Lab*

Folder (Level 2) – *Bio Testing Lab* and *Yellow Jacket*

Projects (under Yellow Jacket Folder) – *Flight-Control-Nav* and *Suite-AI-Processing*

Within the Organization view you can create new projects under folders. The Resources, discussed on the previous page, are not shown in the Organization view, however, when you select different Folders or Projects it will change the Resources visible in the GCP Console.

References:

[1] https://for509.com/gt8s9

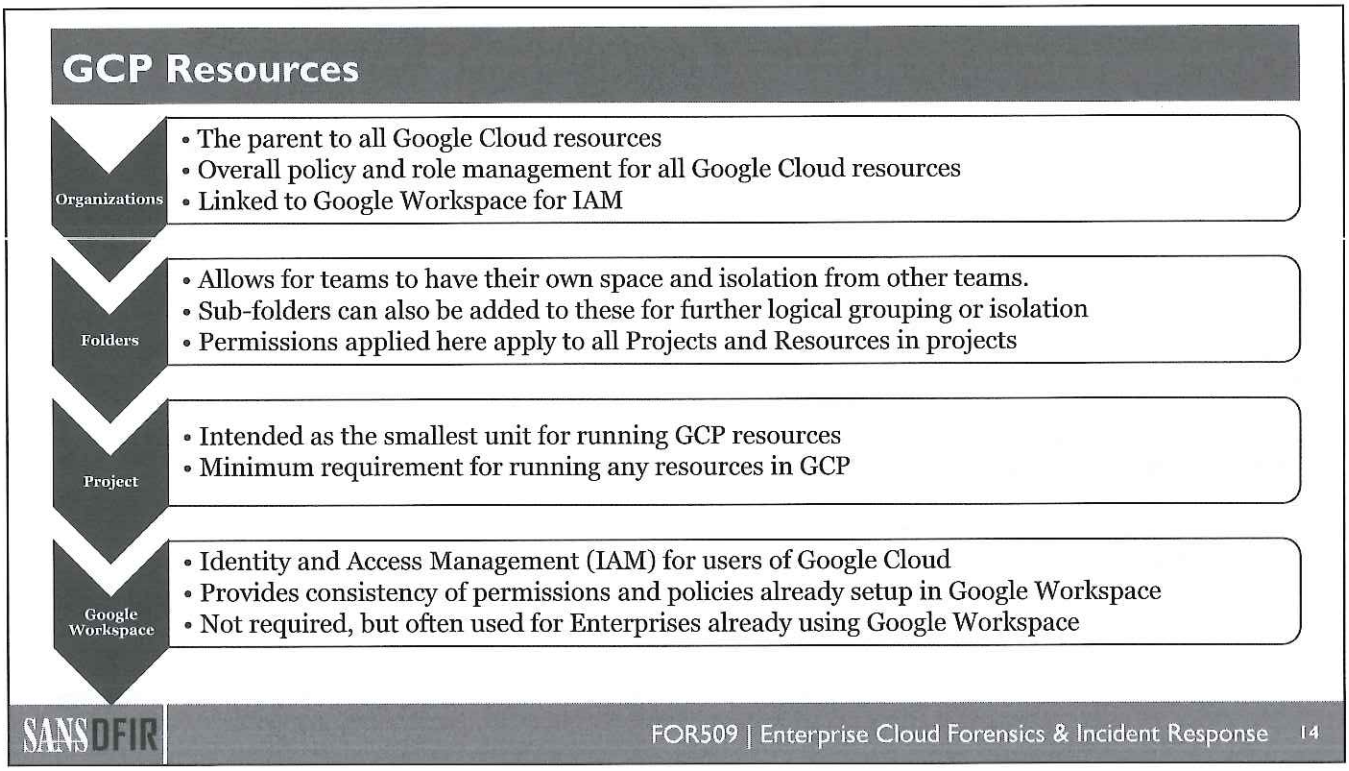

The hierarchy of the GCP resources is intended to not only make organizing your GCP resources easier, but also aid in setting polices or exclusions for different teams or resources.

Polices that are applied higher up the organization roll down to resources and objects under them. Additionally, exclusions can be applied to objects further down the organization tree to open up policy restrictions for specific groups.

The use of Folders in an Organization also allows for isolation away from other teams. While it doesn't restrict one team in a Folder sharing information with other team in a Folder, it does provide for initial isolation of Projects and Resources.

As we'll discuss in future pages, Identiy and Access Management (IAM) form the basis for setting up individual user permissions. At this point it's worth understanding that Google Workspace, and the groups within Google Workspace, can also provide a level of grouping users to allow managing resources and permissions. We'll cover IAM in more details, for now, understand that grouping of Resources, Users, or sub-grouping will provide stronger resource management overall.

References:

[1] https://for509.com/ahcm1

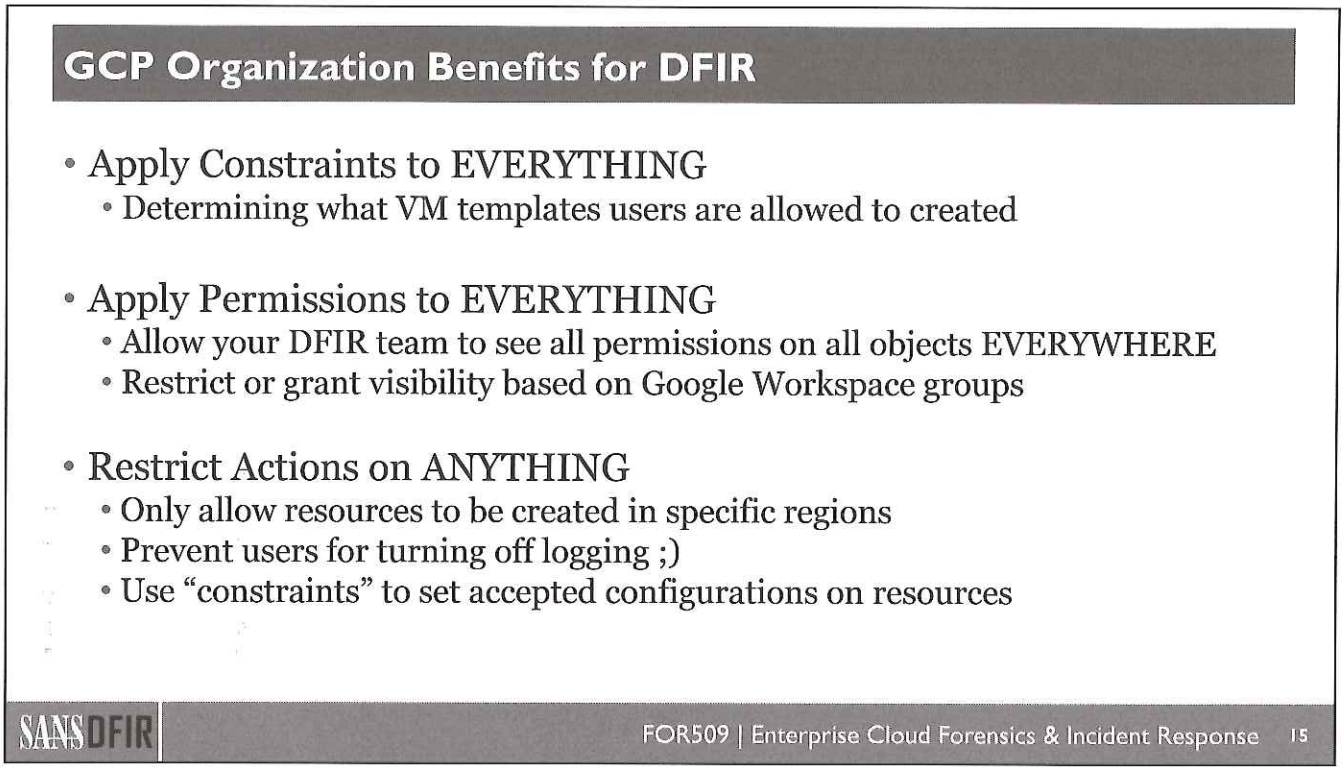

One of the strong benefits of GCP for incident response and digital forensics investigators is the ability to quickly implement policies that can take control of compromised resources or objects within an organizations GCP instance[1][3].

As mentioned in the previous pages, logically grouping resources with Folders and Projects allows you to set restrictions or open executions. However, the level of detail that can be applied vary all the way from restricting specific VM templates users can access, through to enforcing logging standards for specific resources.

When it comes to using polices for incident response, they can allow a response team to set polices for access to all resources within a folder, sub-folder, or project[3]. This can enable a response team to quickly obtain evidence or access to systems for preservation.

It is important to understand that polices allow for setting permissions on "what" an object can and cannot do, whereas IAM is used to determine "who" can and cannot access objects. All policies flow from top down within an Organisation and IAM access applies per object. If you want to add exceptions for polices, they are applied per object.

References:

[1] https://for509.com/5m3il

[2] https://for509.com/7ldt3

[3] https://cloud.google.com/resource-manager/docs/organization-policy/defining-locations

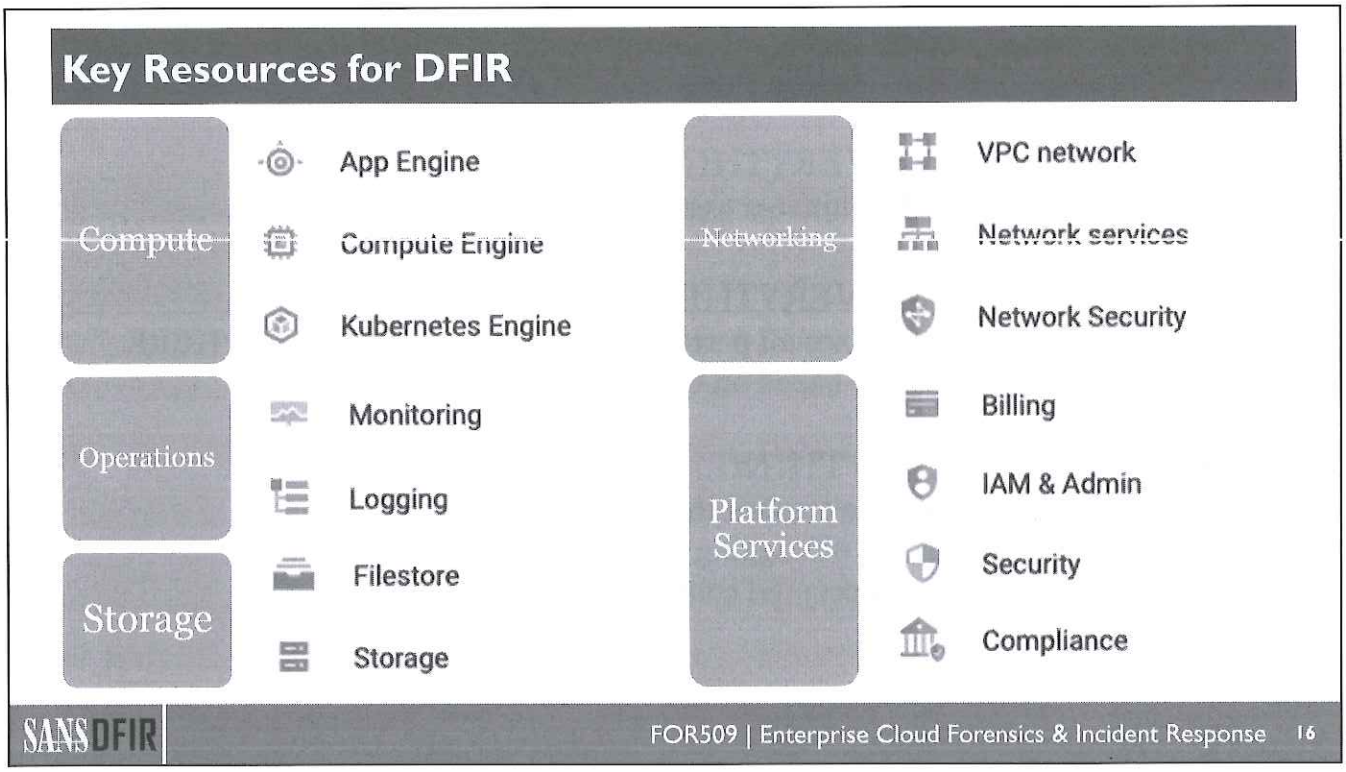

Within GCP there is an ever growing list of resources and services that can be used by Organizations. When new resources are added by Google they are often automatically pushed into a customer's GCP console ready to use. This means there is an ever growing list of resources that incident responders and digital forensics professionals need to stop on top of.

For this section of the course, we have highlighted the key resources that you should focus on first. This list is developed from the author's experience of assisting victims that have had their GCP resources compromised. While this is not an exhaustive list, it covers the large majority of resources when leveraged by threat actors.

During this section of the course, we will focus on most of these resources and how you can obtained and analyze evidence from them. The principles we will teach you for each of these resources can be transposed to similar resources. For example, the logging analysis you will learn with Storage will be applicable to Filestore.

As a starting point, it is important to understand the categories each of these resources are grouped under. These categories directly relate to the categories within GCP, although they will contain a much larger list of resources. This will help you navigate the GCP console, along with understanding conceptually what evidence maybe available to you within each category.

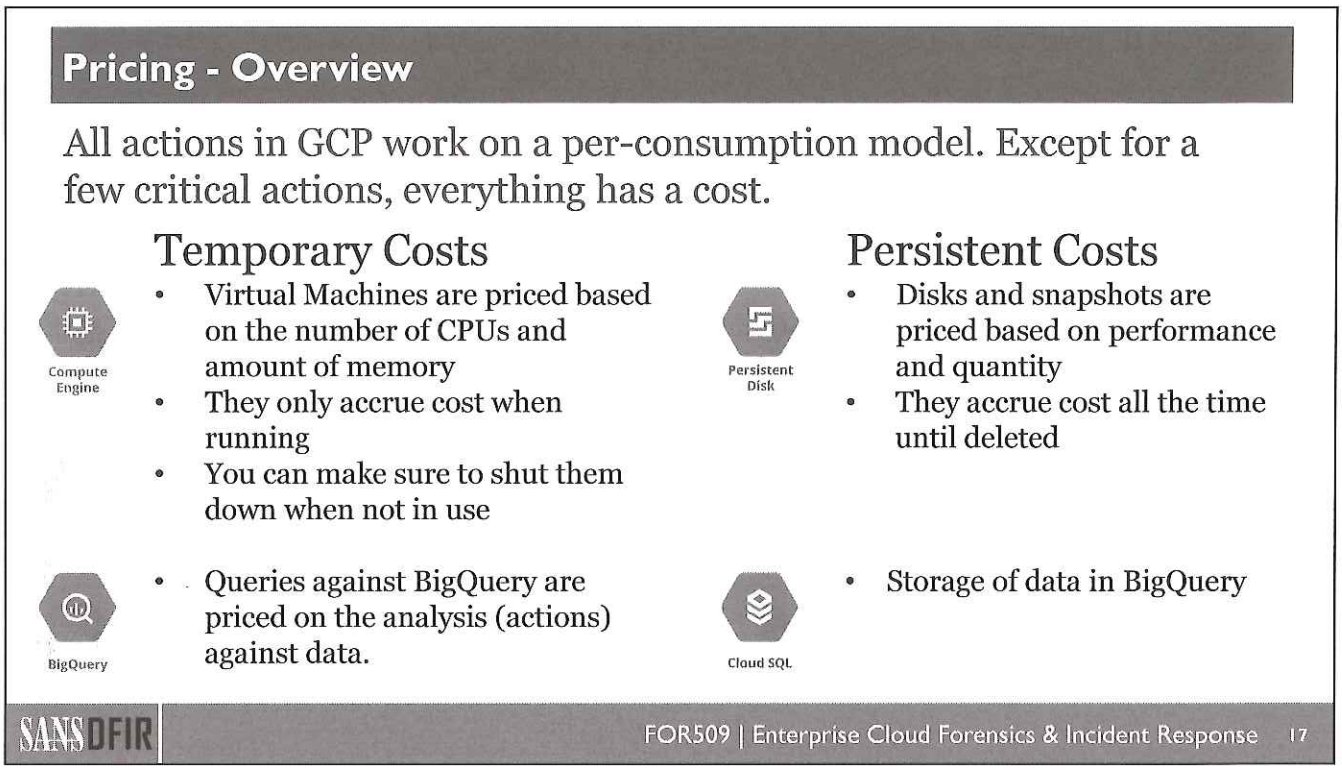

Before we can drive right in and start to enabled logs, store evidence in a Bucket, or spin up a Virtual Machine, you need to have an understand of what doing these tasks may cost you[1].

Being able to use the GCP to not only extract evidence, but also analyze it, will significantly speed up your analysis time and evidence processing time. Additionally, it will enable you to keep data within specific geo-regions if needed, or even within specific network segments if you run into restrictions with moving evidence.

Within GCP there are two concepts to how resources and use of the platform are billed to a customer. These are **temporary costs** and **persistent costs**.

Temporary Costs
Resources that are used on an "as needed" basis are generally billed for as long as you required them to be available. The simplest version of this is a Virtual Machines CPU processing. While ever you need the CPU for a Virtual Machine to be processing your data you will be billed a fee for this. Once you have finished using the CPU for a Virtual Machine and you release you CPU usage, you're no longer billed for that resources.

Persistent Costs
In contrast to the above, resources that you want to hold for a long period of time or that are required for GCP to be held for an ongoing period of time, are considered persistent costs. These type of resources are billed on a quality basis, instead of a "as needed" basis. Using the example above, the storage disk of a Virtual Machine is considered a resources that will need to exist regardless if the Virtual Machine is using the CPU. In fact, the storage disk is expected to be persistent even if I reboot or temporarily suspend my Virtual Machine. The storage disk is considered a persistent cost that is billed based on the space it is consuming within GCP.

Mixing of Costs
The concept of temporary and persistent storage are intended to be simple concepts to aid in understanding how resources are bill. However, you should be aware that some services mix these concepts together when

they bill a GCP customers. An example of this would be the BigQuery resource. As a customer you are billed for the storage size you consume for holding data, you'd like to be able to query with BigQuery, this is your persistent costs. You are also billed for any queries or "actions" you make against the data you are storing; this is considered a temporary cost.

It is important to understand both these cost models as they can be useful to help decrease your GCP costs, as an analyst. They can also be cause significant billing costs if you aren't aware of the services you're enabling and the resources you leave running.

To get a complete overview on the costs associated with GCP, as they are subject to change, you should review GCP's online costing for their resources[1]. One final note on costs, there are a small number of items within GCP that have no billing cost associated with them, they include;

- Using GCP's console via a web browser – which is shown extensively throughout this section of the class.
- Cloud Shell – which is really used within the web browser while you're accessing the GCP Console.
- Logs in the "_Required" log bucket that Organizations cannot change

References:

[1] https://for509.com/rvkbh

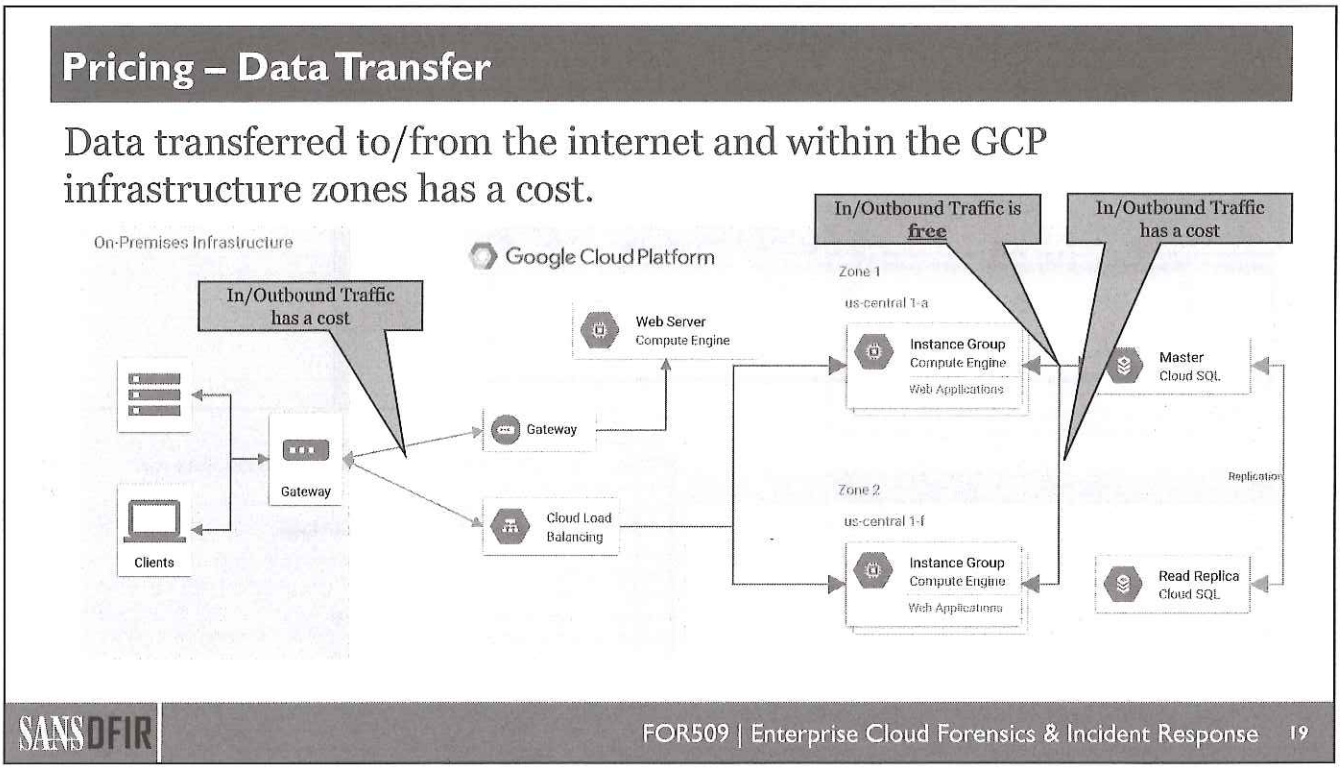

When it comes to moving data to and from your GCP instance you will also attract ingress and egress fees[1]. These fees vary depending on the region you are moving data into or out of the GCP.

The costs assigned to both ingress and egress are often based on each resource that you use. For example, if you are sending data from on VM (VM-A) to another VM (VM-B) in a different region you will incur both ingress and egress costs at both locations. This is because if "VM-A" sends traffic, you will incur an egress fee, but you will also incur an ingress fee for traffic the is going back to "VM-A" as part of any TCP communications that respond to the traffic being transferred to "VM-B".

At the time of publishing this class GCP was not chagrining for ingress traffic going into a GCP resource, however, be aware this is subject to change and should always be checked using GCP's pricing pages[1].

Traffic that moves between GCP resources within the same region and zone, while using GCP's internal assigned IP addresses, do not attract any traffic charges. Because of this you should always try and use GCP's internal IP addresses and put systems that will communicate a lot in the same zone. This applies to any of your forensic processing tools that will access other systems via GCP's network.

If you move traffic between GCP zones within a region, or between GCP regions, this will incur a fee, however, again if you use GCP's internal IP addresses you will use Google's network backbone which assists in keeping network traffic costs lower.

Traffic costs can quickly get out of hand when you're using GCP for processing evidence. Especially if you are moving large amount of data to any forensic tools within VM's. Always try to keep forensic processing VM's in the same region and zone that your evidence is originating from and try to always using GCP's internal network IP addresses. In addition to keeping costs lower, it will also keep the speed of traffic transfers over the network quick.

[1] https://for509.com/yamvd

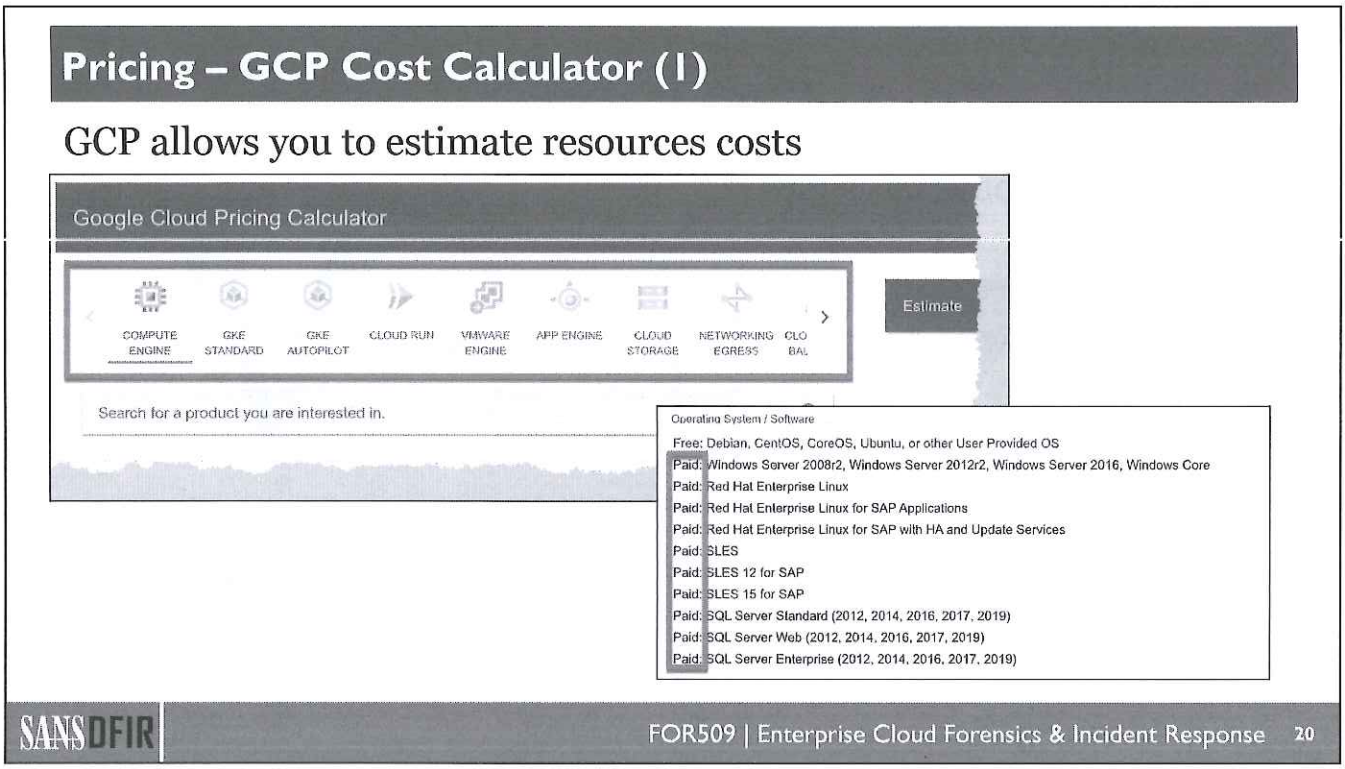

GCP provides a simple online cost calculator to estimate resources you intend to use[1]. This will allow you to select the resource, the region, the amount of data, and the length of time you intend to use GCP resources. Remember this calculator is only as accurate as the data you provide it, so ensure your estimations are realistic.

You should also be aware that some resources within GCP have additional running costs that maybe not listed on GCP's resource pricing page. An example of this could be the host operation system you are using within a VM. GCP will clearly give you details about the resource cost of using the VM, but you need to also be aware that the licensing costs for the host operating system will be charged to you.

GCP has provided useful licensing when it comes to host operating system licensing. They will provide you a license on an "on-demand" basis. So, if you create a VM running Windows Server 2016 Standard, you will be changed for this as the VM is created. Additionally, users can also use their existing Window VM licensing but adding it into the GCP console[2]. GCP's Cost Calculator will take operating system licensing into account when you are estimating you running costs.

References:

[1] https://for509.com/zqrk4

[2] https://for509.com/5u6rw

Pricing – GCP Cost Calculator (2)

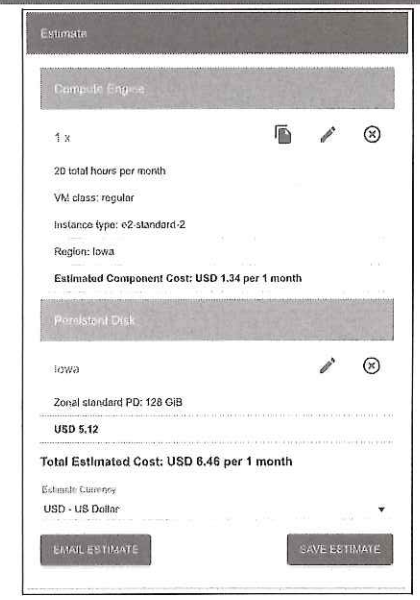

Virtual Machine Estimate:

- 2 vCPUs
- 8GB RAM
- 128GB for OS drive
- 20 hours uptime for the investigation

- Higher performance machine may be required depending on the investigation

The GCP cost calculator is perfect for estimating your running costs for a forensic analysis system within the cloud[1]. In the example above we have chosen a typical VM that would be suitable of accessing other systems in the cloud and for processing disk evidence attached to the VM. The size and requirements for an analysis system can vary significantly depending on scale and it's intended job.

Often when using GCP VM's for forensics analysis processing, you may setup a VM very similar to the example show, then for processing timeline data you could increase the resources for the VM so it can perform the data processing much quicker, then revert the VM back to the example shown for review once timeline processing is completed.

In the example shown we have selected "Lowa (US-Central1)", additionally the pricing was generated with a price list that was last update on 24th of March 2021.

References:

[1] https://for509.com/zqrk4

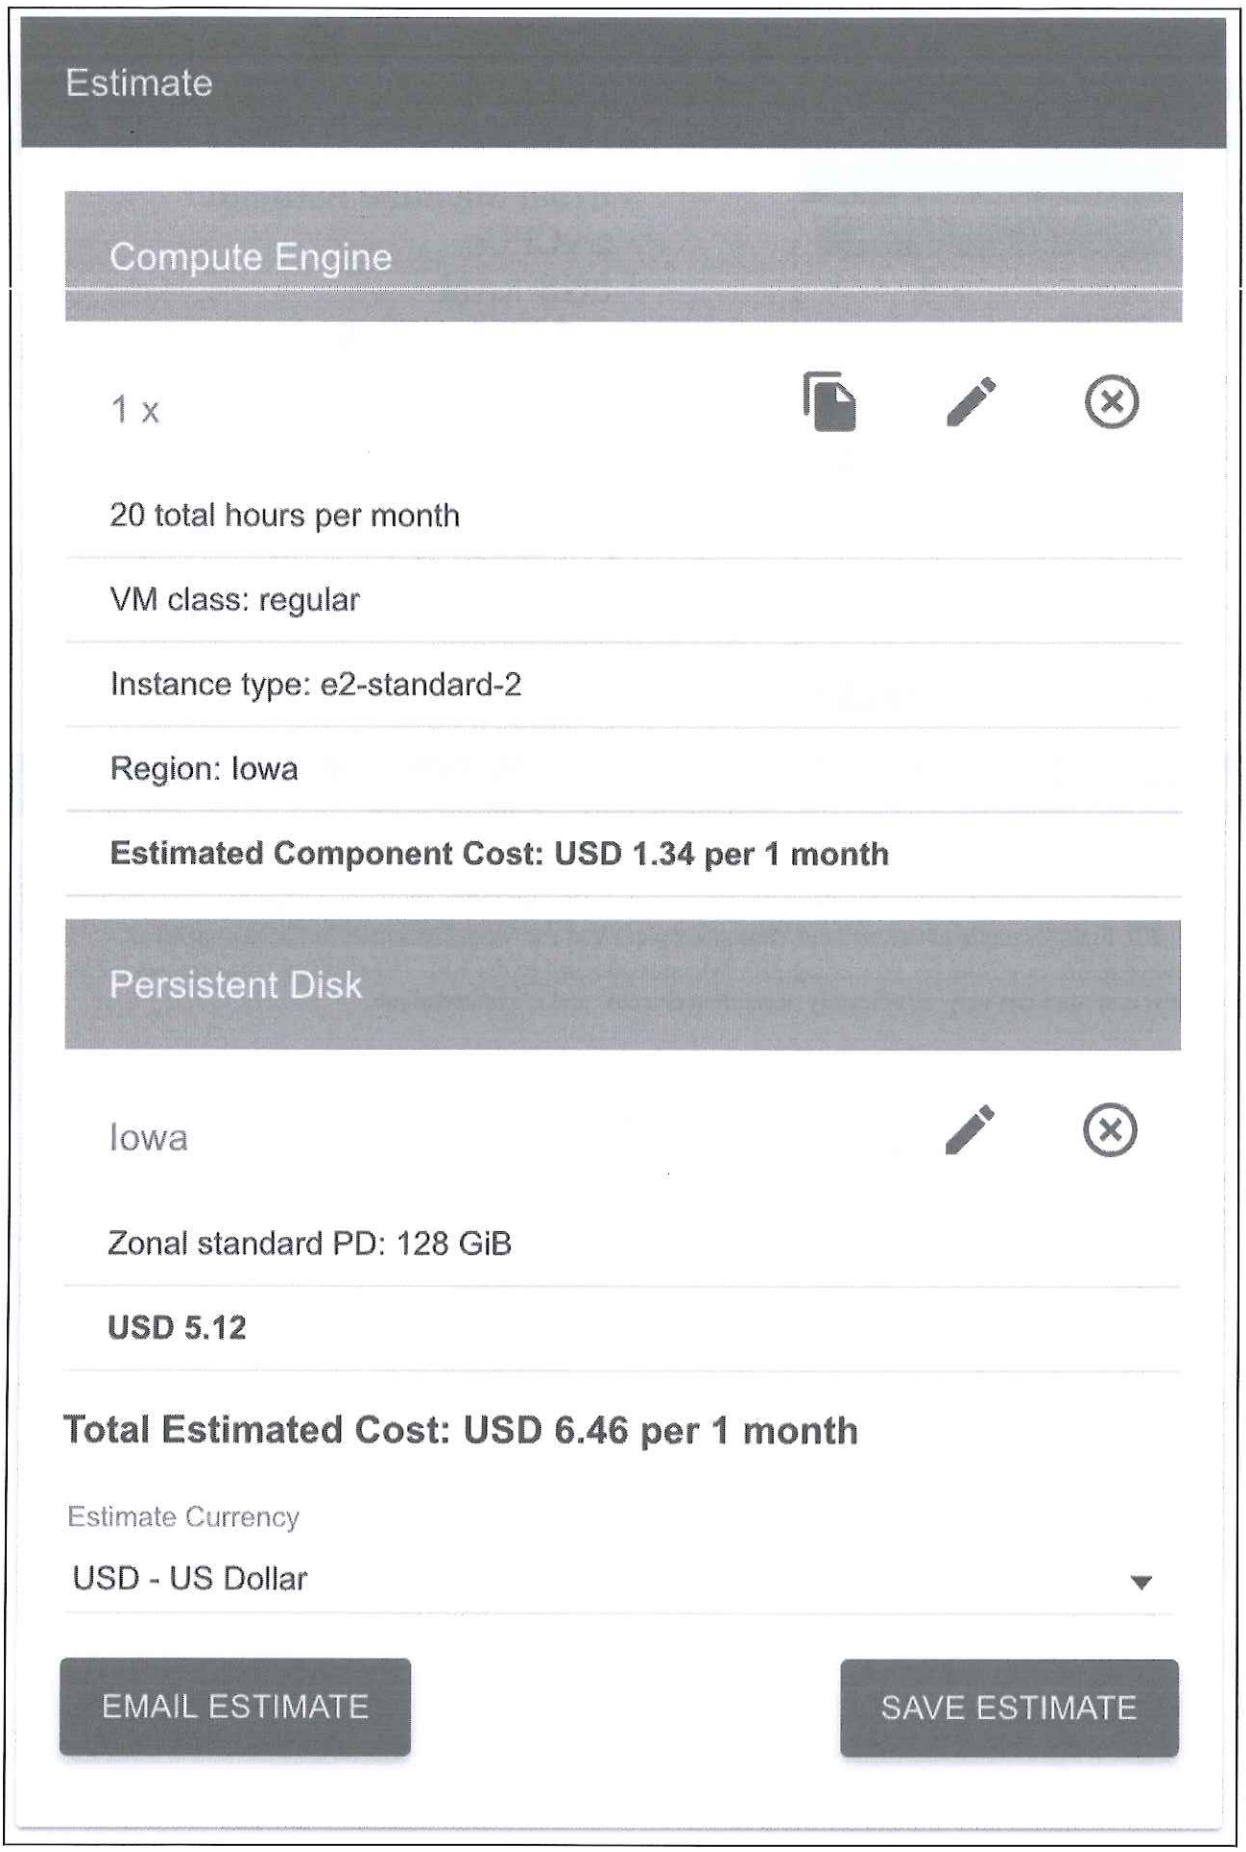

Pricing – GCP Cost Calculator (3)

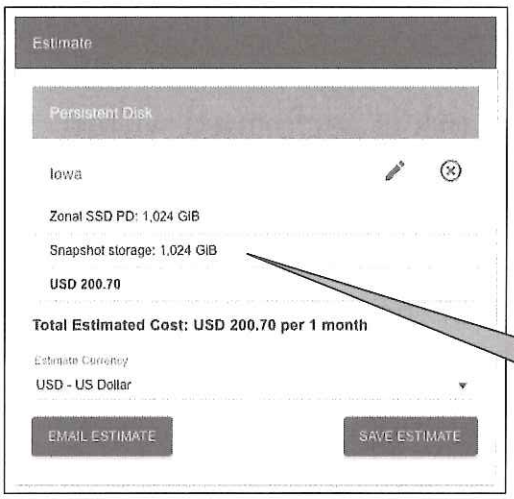

- The 1st storage cost is to apply the snapshot to and mount on the forensic VM. Performance is key, so Premium SSD is selected
- The 2nd storage cost is for holding the snapshot.

Snapshot storage is priced differently to a running disk

When performing analysis with a VM in GCP you are most often going to be access snapshots of other system when it comes to disk analysis, as mentioned earlier in this class. Because of this, you have to include the cost of hold the snapshot image as part of your disk analysis. However, you don't need to hold this snapshot on expensive storage given you will mainly be querying or extracting individual items from the image. This allows you to the snapshot on cheaper storage[1], as show in this example.

References:

[1] https://for509.com/ijn48

Pricing – GCP Cost Calculator (4)

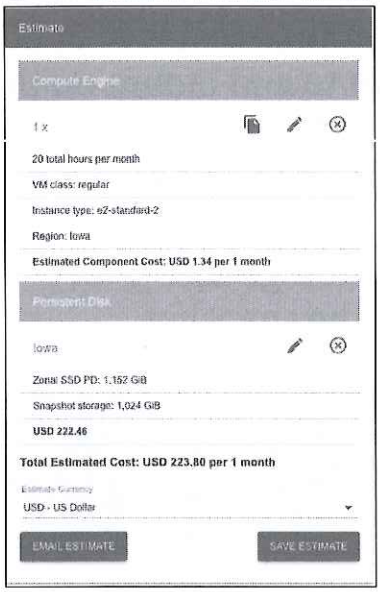

- This is a typical monthly cost for a "small" investigation where all DFIR work is performed within GCP.
- GCP's cost calculator does not break out the individual disks, it bundles them together.

Taking into consideration the size of the analysis VM we have showed so far, running that VM for approximately 20hrs to conduct analysis, along with holding 1Tb of snapshot data, it will cost you (as at the 24th of March 2021) approximately USD $222.

In terms of maintaining your own hardware, or the task of procuring hardware and setting it up, this cost is relativity inexpensive. At the time of writing this section, GCP is also the cheapest cloud, out of AWS and Azure, to run an analysis system of this size.

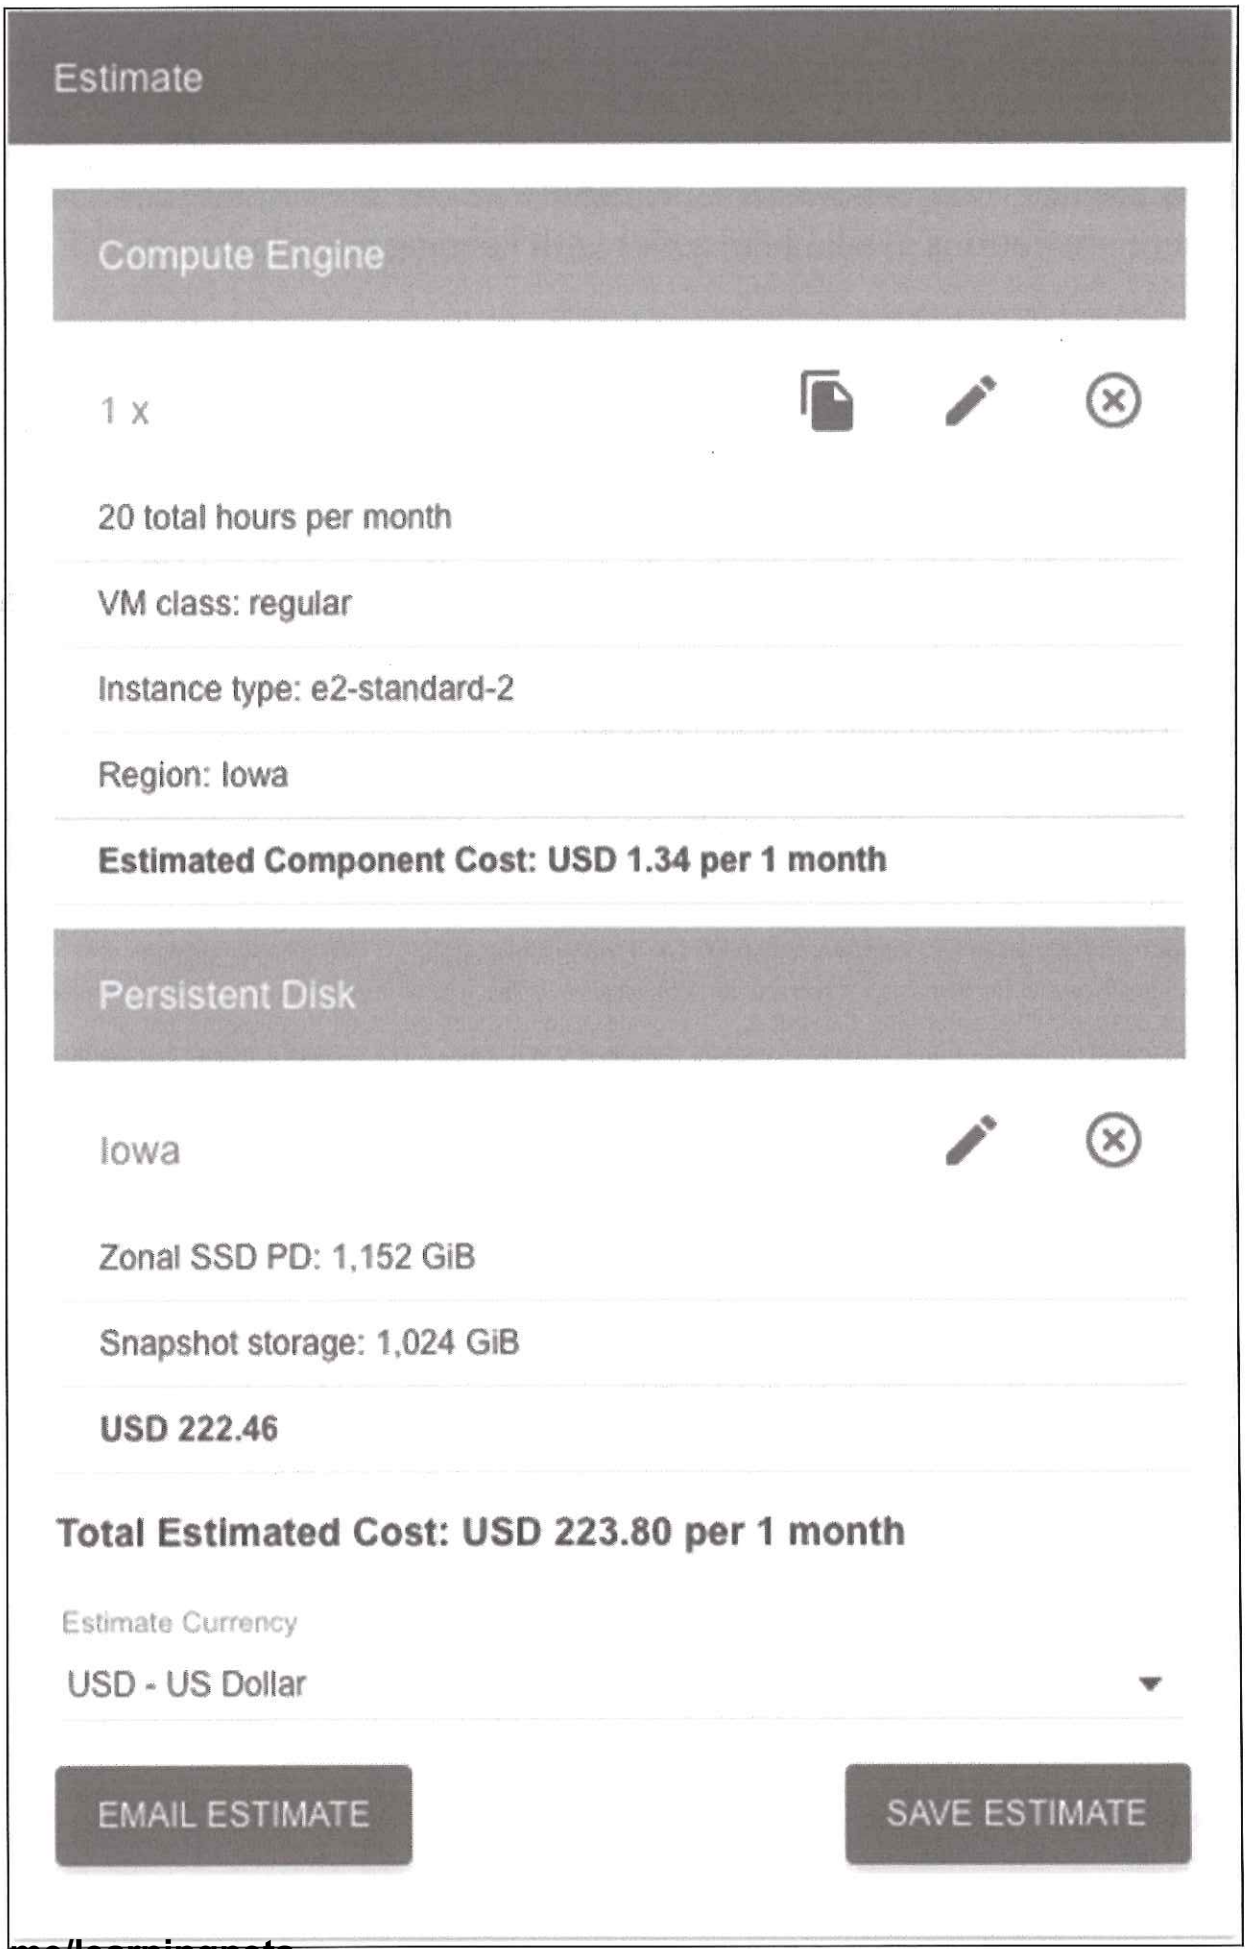

Pricing – Cost Reduction Options

There are options available to assist with keeping costs lower in GCP.

- **Preemptible VM Instance**
  - CPU is variable and changes pending demand within GCP
  - Ideal when you have multiple VMs and you don't require dedicated CPU
  - Not ideal for single system evidence processing (i.e., Plaso)
- **Data Movement within GCP**
  - Moving data within a single zone doesn't cost money – if you use an internal IP address
  - Put your DFIR Analysis VM in the same zone as your evidence source
- **Queries Against Data**
  - Queries are charged based on bytes processed
  - Holding data in smaller data sets can save on the cost of each query

There are some useful cost reduction options that the GCP provides customers. One of these, that's extremally useful for VM's, is a "Preemptible" VM instance[1]. This type of instance won't provide customers with dedicated CPU processing, but instead will provide customers with excess CPU processing that isn't being used by other customers that have paid for dedicated VM instances. The concept is that a "Preemptible" VM will flux within its CPU utilization based on CPU resources that aren't being used within the overall GCP. This is a particularly good way to reduce VM costs, however, it does mean you don't get dedicated CPU resources for processing. Where these instance maybe useful to customers and investigators, is when you have multiple VM's that are load balancing requests or computation, so they have redundancy built in. These type of VM's would not be suitable if you have an evidence process that has to reside within a single VM, for example Plaso.

When you conduct investigations within GCP ensure you get your evidence processing VM's in the same region and zone as the majority of your target systems/evidence. As discussed previously, moving data within the same region and zone has no ingress or egress costs provided you use the GCP internal IP address to communicate between systems.

Lastly, be careful of how you use data analysis tools within GCP. Most common for investigators is the use of resources like BigQuery. As mentioned previously BigQuery is billed based on holding data and also querying of data. While there is not a lot that can be done to reduce the cost of holding the data, you can reduce the cost of querying data. BigQuery is billed based on compute time used to perform a query[2]. If you can reduce the number of queries you need to make it will aid in keeping costs lower, along with ensuring you document the output of a query so you, or other investigators, don't need to run it a second time. If you have team members that have strong SQL skills, use them to help craft queries that perform sub-queries or cross-correlate data in a single query.

References:

[1] https://for509.com/5xbh3

[2] https://for509.com/ei73s

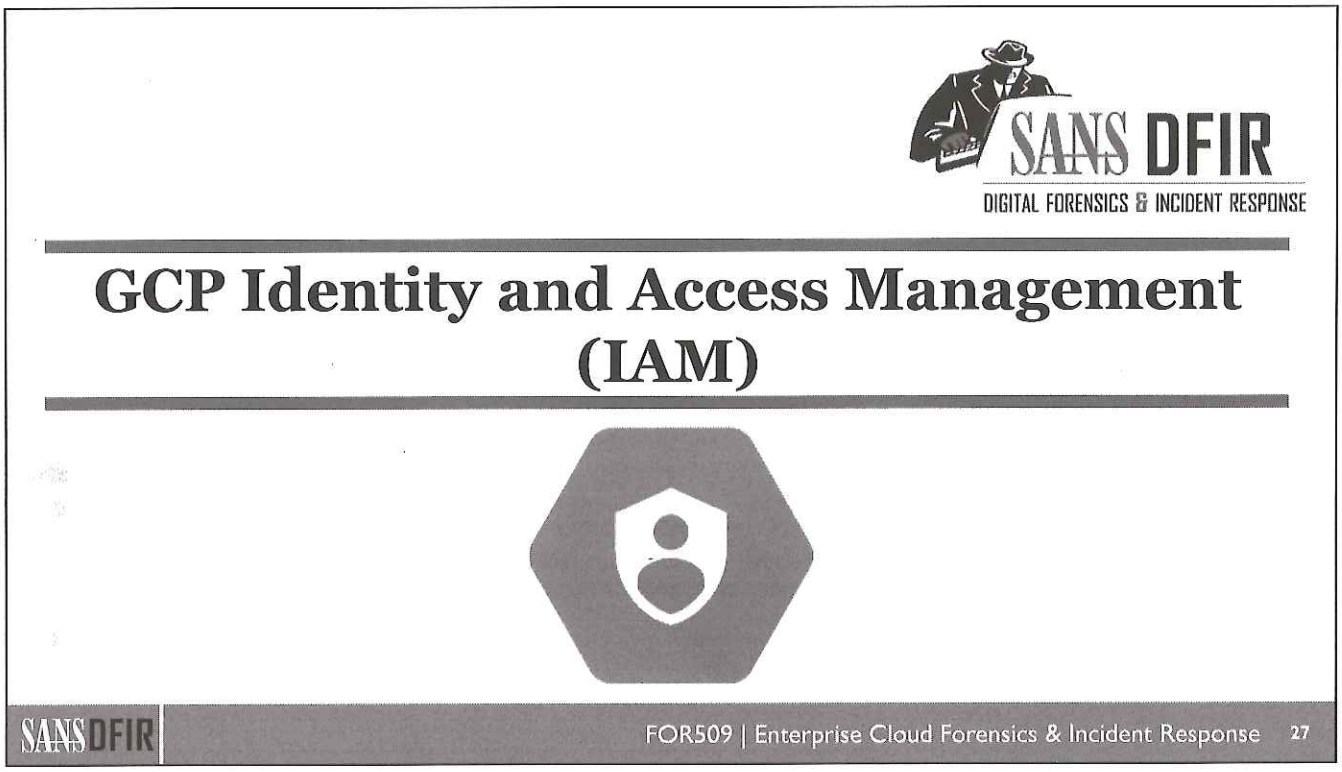

This page intentionally left blank.

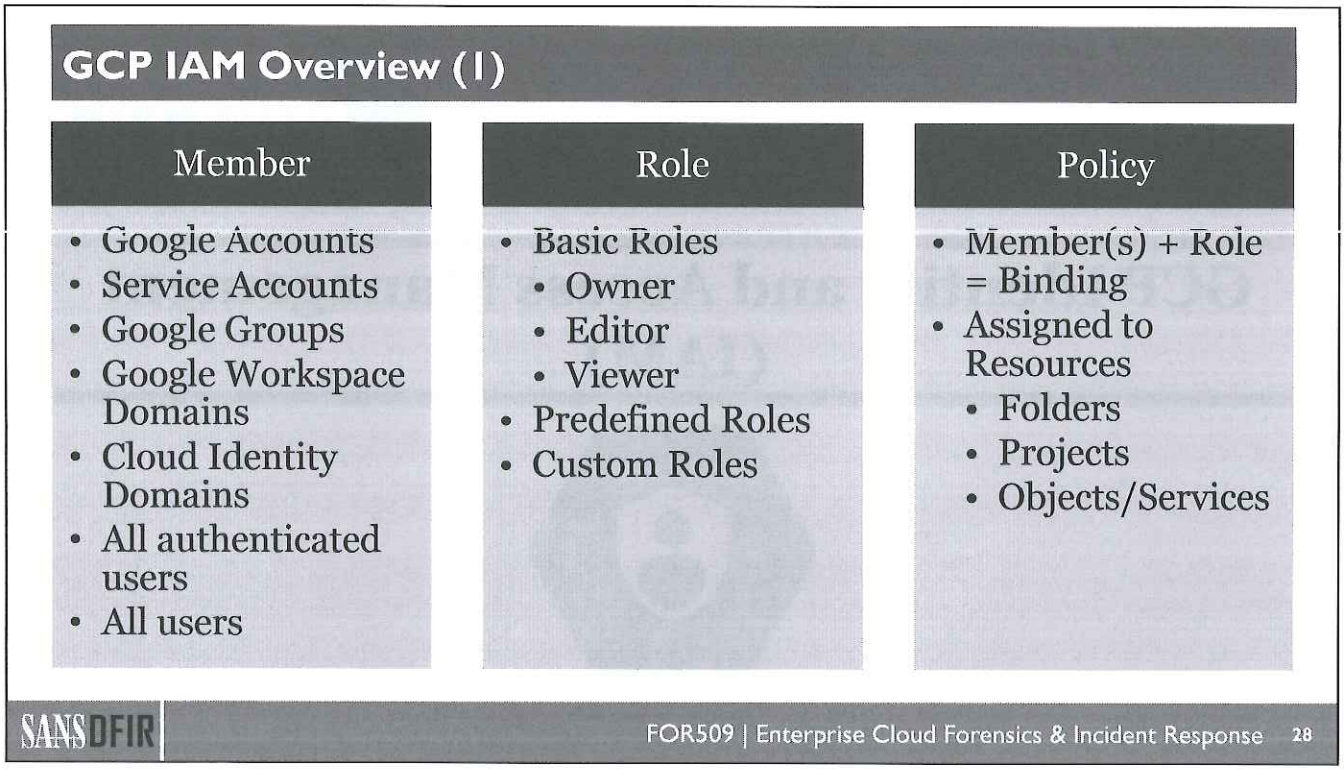

Identity and Access Management (IAM) is a key resource that all investigators will heavily use in GCP[1]. Because of this, it's important that you have a strong understanding of how IAM works and how it's setup within GCP. This section will cover the fundamentals of IAM and provide you with the building blocks to understand IAM in more depth. While IAM is the main backbone to identity and access management within GCP, you should be aware that some large enterprises may use other third-party tools to aid with their IAM management, however, the basis for what resources have access to will still be bound by the concepts we'll present in this section.

GCP defines IAM objects into "Members", "Roles", and "Policies"

Member
A member is an individual user, group, or service account that is treated as an object that can have permissions assigned. Members are not limited to only users within GCP, it can also be users or groups within Google Workspace, or users and service accounts in other GCP organizations. An email address is the primary identifier for members including service accounts and groups, however, this can also be extended to a domain when you are defining a very large group of members together.

Roles
Roles represent a collection of permissions assigned to a common group that grant or restrict access to Members. Within Roles there are three common categories, there are; Basic Roles, Predefined Roles and Custom Roles.

Basic Roles have existed since GCP was first developed by Google and consist of three sub-groups; Owner, Editor, and Viewer. These types of Roles were really developed as a simple and easy way to get started with GCP; however, they can present challenges with securing a GCP instance that we will look at later in this section.

Predefined Roles are similar to Basic Roles; however, they provide a much more fine-grained approach to permissions allow you to be a more specific with the types of permissions you may want to assign a Member. These role types, as the name suggests, are predefined by GCP. In terms of securing your GCP these are the types of Roles you should be assigning Members.

Custom Roles[2] allow you to create your own role types using the thousands of permissions available in GCP. These are ideal when you need to create custom roles that are within the Predefined Roles offered in GCP.

Policy
Policies are a way to tie together a number of Members, assign them a Role and bind it to a resource within GCP. A Policy is assigned to resources within GCP and are used to check permissions when a member attempts to accesses a resource.

References:
[1] https://for509.com/nrdg0
[2] https://for509.com/g0kqa

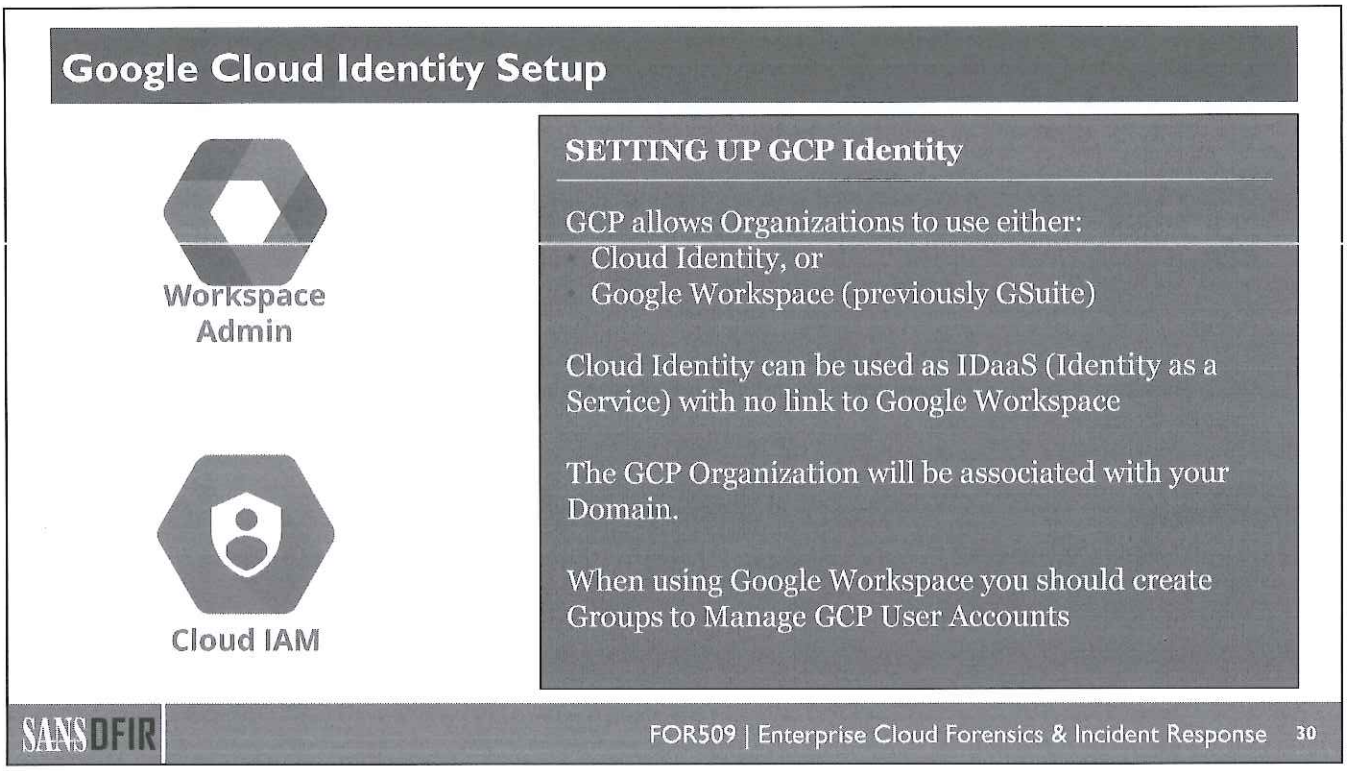

When you setup GCP you will need to determine how authentication and Members are managed. GCP provides you with two standard options being either the IAM service within GCP referred to as "Cloud Identity" or through the use of "Google Workspace"[1]. There are other IAM services, outside of GCP and Google Workspace, that make it possible to create members and authenticate within GCP, however, these usually use a connector into one of the two options mentioned above. Authentication and Member management outside of the two options mentioned here are beyond the details of this course. However, the principals will be similar.

When you link your Google Workspace to GCP you also link your primary domain, within Google Workspace, to GCP. This results in your GCP Organization being your Google Workspace's primary domain.

As you'll see shortly, using Google Workspace can become a simple way to managing users and groups for GCP access.

References:

[1] https://for509.com/o8za4

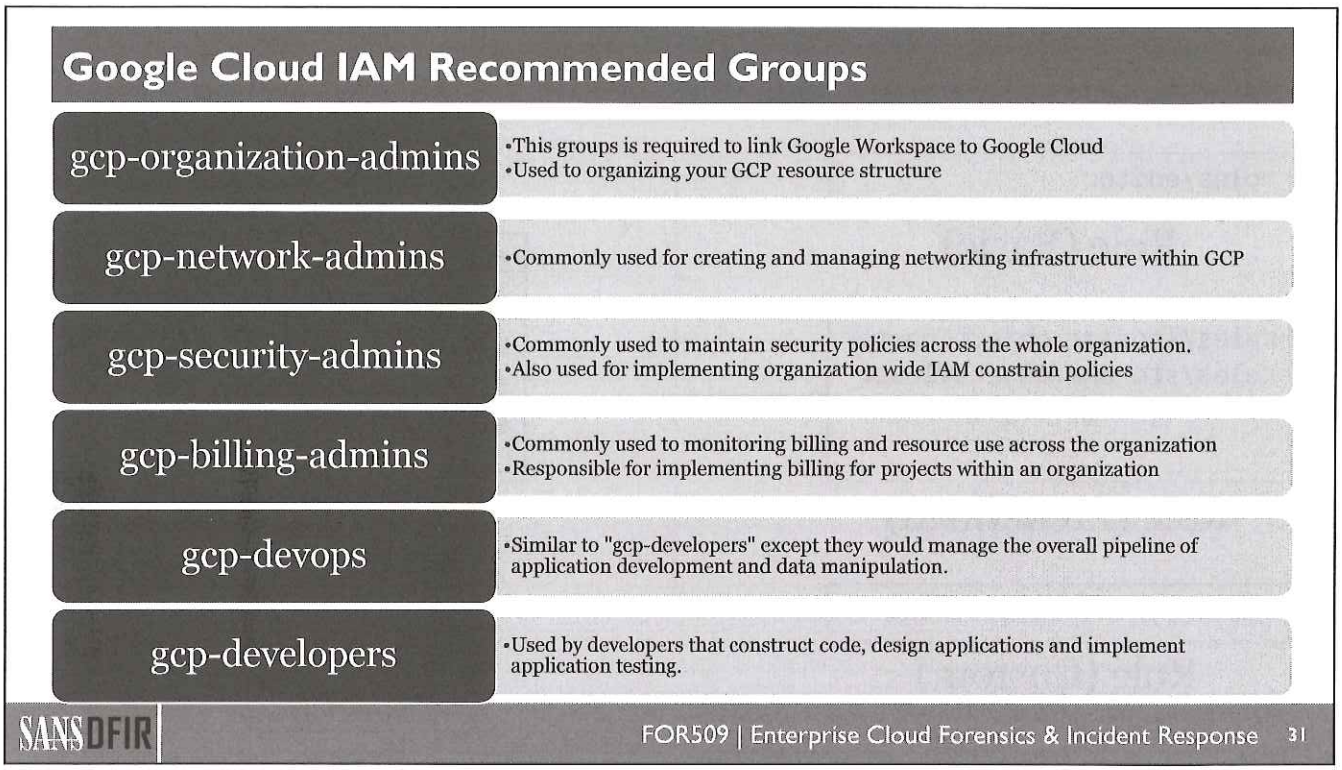

When setting up Google Workspace to work alongside of GCP IAM, you have to create the "gcp-organization-admins" group and put your Organizational admins within that group. Without this group GCP will not link to Google Workspace.

The rest of the groups presented in this page are intended as best practice groups for your Google Workspace and your users[1]. These groups will help you better define the users and allow you to assign Roles to Google Workspace groups for a better structure of permissions. You can of course put a single member/user within multiple groups; however, you should avoid this as part of best practice with separating permissions.

References:

[1] https://for509.com/x6h3g

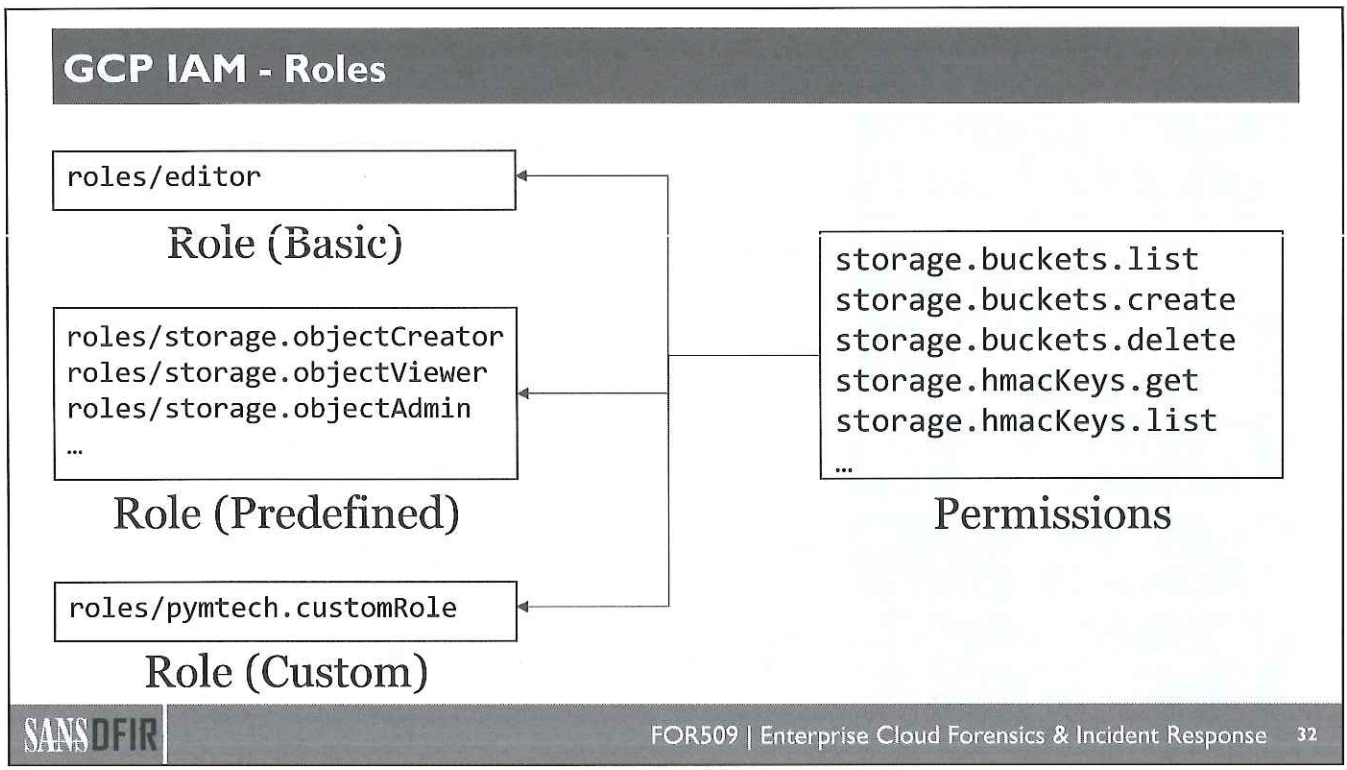

To better understand how permissions are defind and how Roles link with permissions, consider a Role a group of permissions. The permissions within GCP are preset granular permissions for each resource one permissions on its own is not overly useful, however, it's when you cluster or group these permissions together you can define a set of permissions as a Role for a Member.

As mentioned previously, Basic Roles are predefined roles with a predefined set of permissions that cannot be changed. However, they do offer the greatest flexibility for quickly building within GCP. Although, they are very broad sets of permissions that may provide more permissions than you initially need or intend Members to have.

Predefined Roles are pre-made groups of permissions from GCP. They group permissions together for a more narrow defined Role[1]. These are well documented within GCP's documentation to allow Organizations to better define what each Member is allowed to do within GCP[1].

Custom Roles allow Organizations to make their own permissions and assign them to Custom Roles. Along with creating your own permissions, you can use the existing permissions within GCP and mix them together with any custom permissions you decide to make. They allow an organization to specific exactly what permissions are assigned to a Role, however, there are hundreds of permissions, and this can be an extremely long task to get right.

Reference:

[1] https://for509.com/0mors

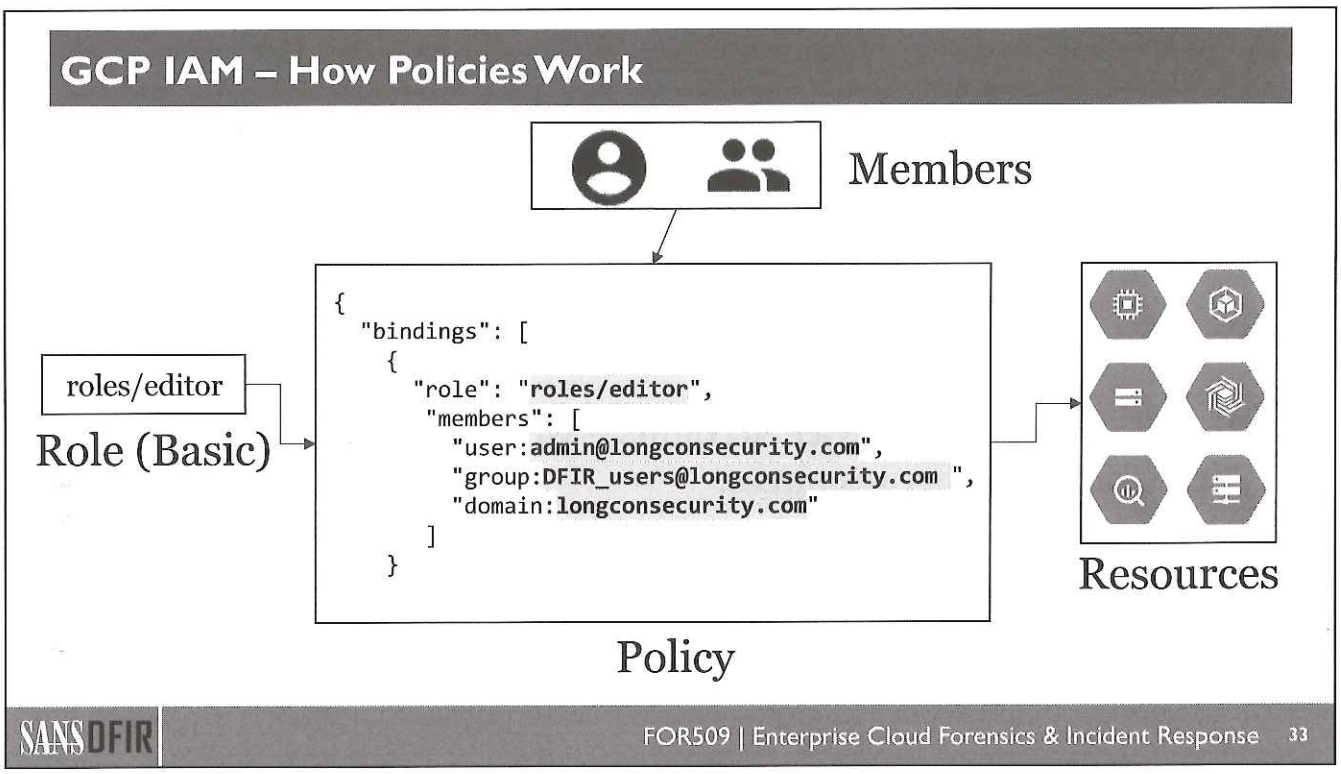

A Policy is made up for the binding together of Roles, which are groupings of permissions, and Members into a single Policy object. This Policy is then applied to one or many resources within GCP.

Policies are not applied to members! This is one of the fundamental differences to GCP compared to what you've learnt so far with AWS and Azure. You cannot query a Member and determine all the permissions that a member may have. You can only query a Resource to determine who has permissions for that Resource, this is because a Resource has the Policy applied to it.

To complicate further an administrator's ability to determine what permissions a Member has, you can have resources in a completely different Organization grant access to a member in another unrelated Organization. This means one of the Members in your Organization may have high level permissions to a Resource in another Organization and you would be unable to determine this from within your own Organization.

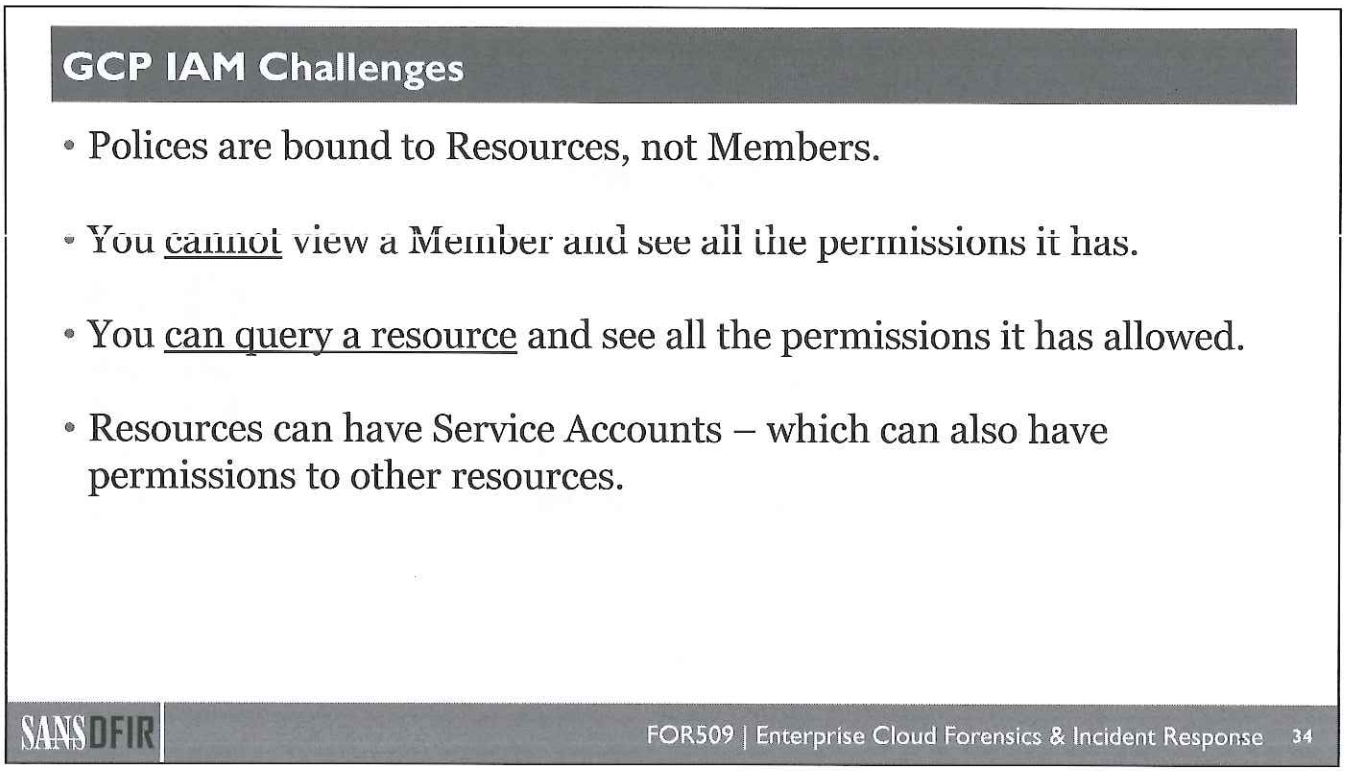

The challenge with how GCP applies permissions to Resources, instead of Members, is a change in the model of how most people expect identity management to work. This means you can't simply query a Member to understand all the permissions they may have, or all the Resources they could possibility have access to. This is often a challenge when it comes to incident response if an account with GCP is being misused. It results in investigators having to rely on logging a lot more than they would to determine a Members permissions. In the world of Microsoft Active Directory, as mentioned previously in the course, you could query a user and see all their groups and each group's permissions. That's no longer the case within GCP.

To understand the permissions a Member may have, you have to query the Resource. This can have some benefits from an investigation perspective. It means if you know a certain Resource was abused you can simply query the permissions that Resource has assigned and see all the Members that can access it and what their permission level is on that Resource.

As mentioned in the previous pages, this can be challenging as you could have Members with access to Resources in another GCP Organization, that you have no control over, but you Members could have access to it. To extend this challenge further, if a Member within your GCP Organization accessed a Resources within another GCP Organization the logging for that access would be in the *other* Organization, not yours.

To extend out the permissions model a little further, Resources can have their own Service Accounts if they need to access other Resources. For example, if you have a VM Resource that needed to access a Storage Bucket for reading and/or writing data, it will require a Service Account for the VM Resource with that Service Account given permissions on the Storage Bucket via a Policy applied to the Storage Bucket. Service Accounts work in the same way we have discussed Member accounts, in fact that are considered Member accounts. This means the same permission challenge mentioned above also applies to Service Accounts.

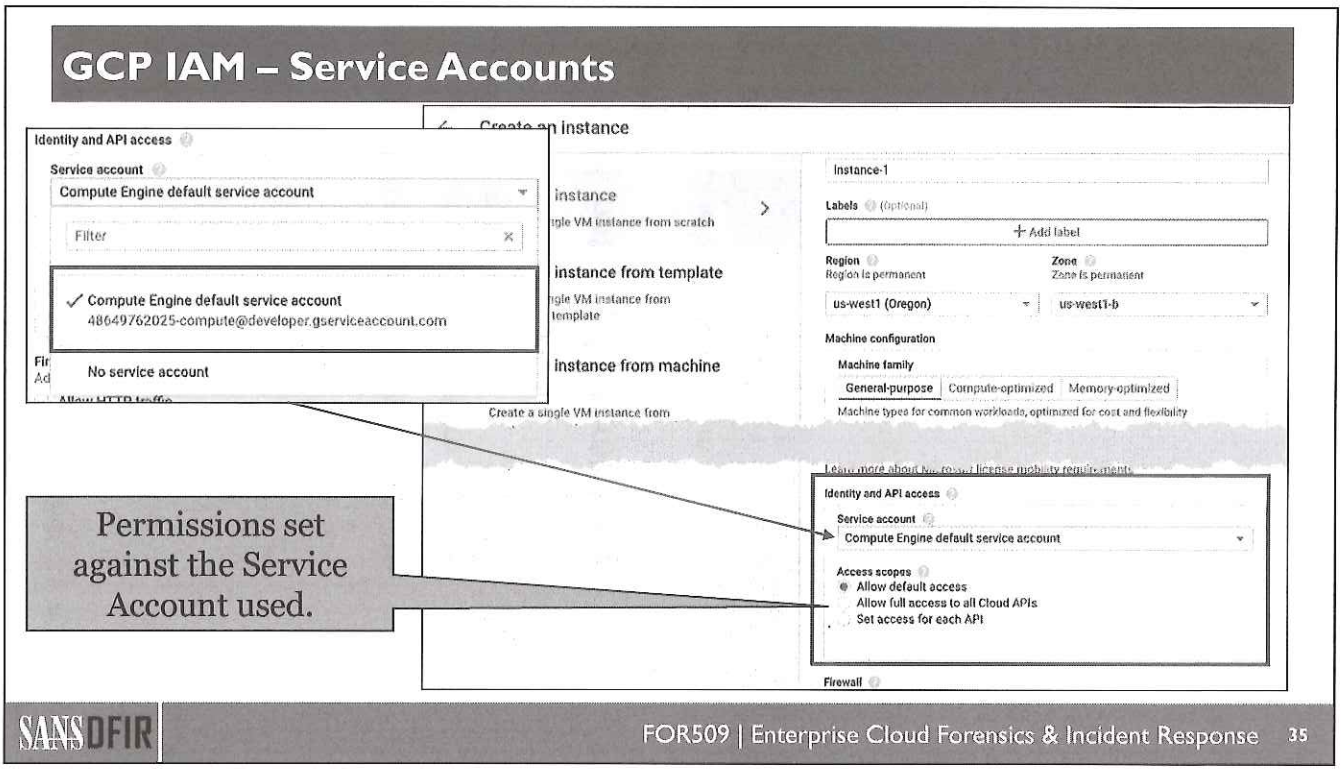

In most cases Service Accounts are created as you create new Resources within GCP that may be used to access other Resources within GCP. When you created a new VM a Service Account is created and assigned to that VM at the time it is created[1]. This is shown in the example above where a new VM is being created and a Service Account for the resource is created at the same time. This Service Account will also appear in IAM and allow you to manage the account the same way you would any other Member in IAM.

You could delete this account in IAM; however, it will likely cause access issues if you're attempting have your VM interact with other GCP Resources. You can also control the permissions this account has on other Resources, which is usually the best way to keep your Service Accounts under control in terms of their level of permissions.

References:

[1] https://for509.com/w5ykm

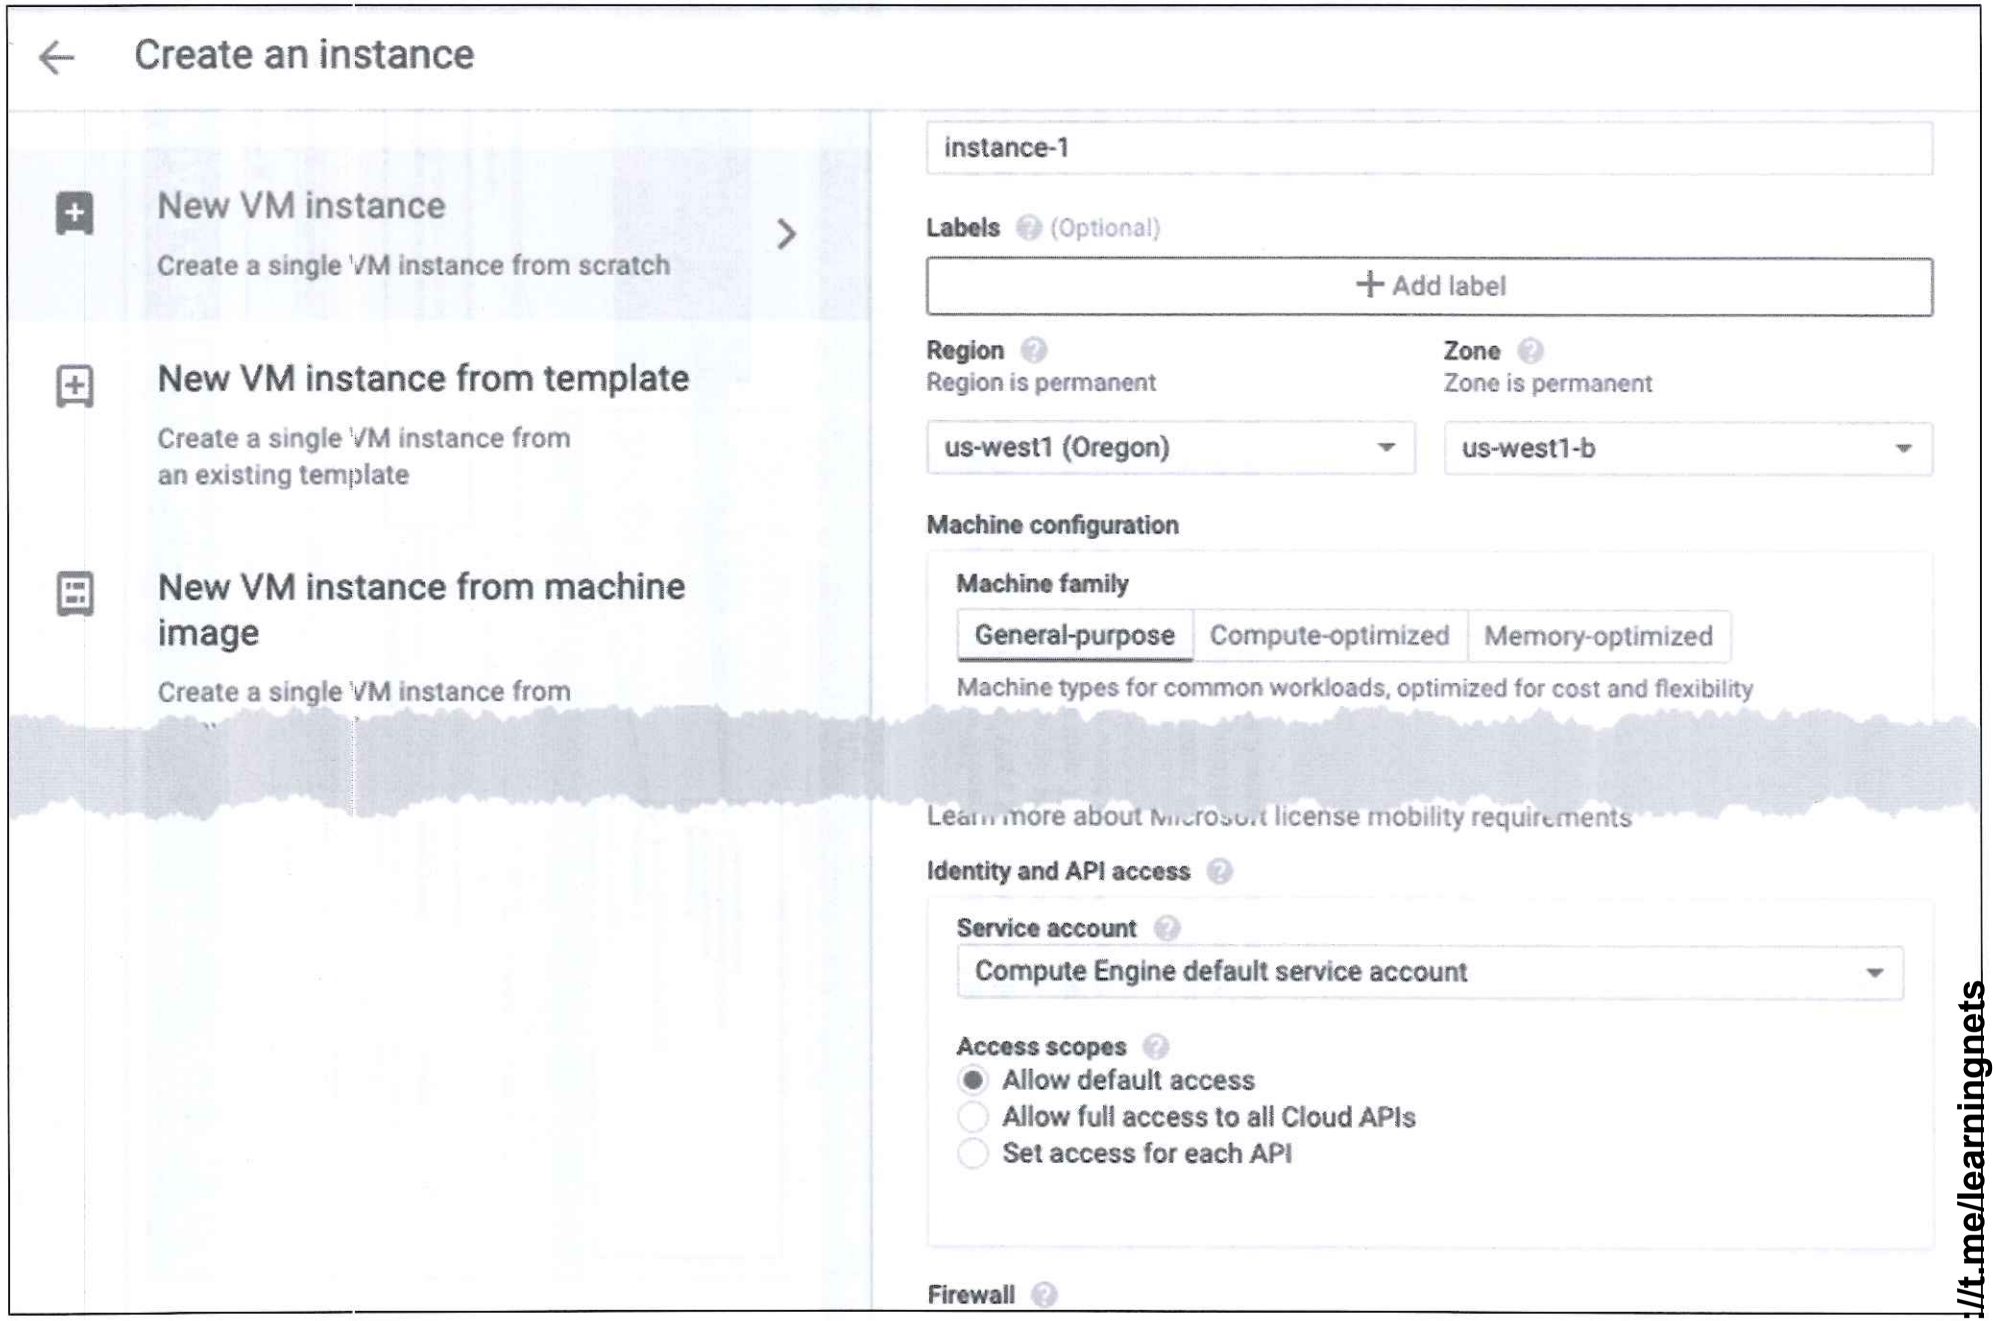

GCP IAM Challenges

Common Lateral Movement Attack

1. Infrastructure Team set up "Development" VM for testing
2. Developer creates code within VM but needs a Storage Container
3. Developer creates Storage Container ~~in their~~ own Project Folder
4. Developer grants VM Service Account access to Storage Container
5. Infrastructure Team set up separate "Production" VM
6. Developer runs code in "Production" VM
7. Developer's code needs access to Storage Container
8. Developer grants access to the same Storage Container with the "Production" VM

This page intentionally left blank.

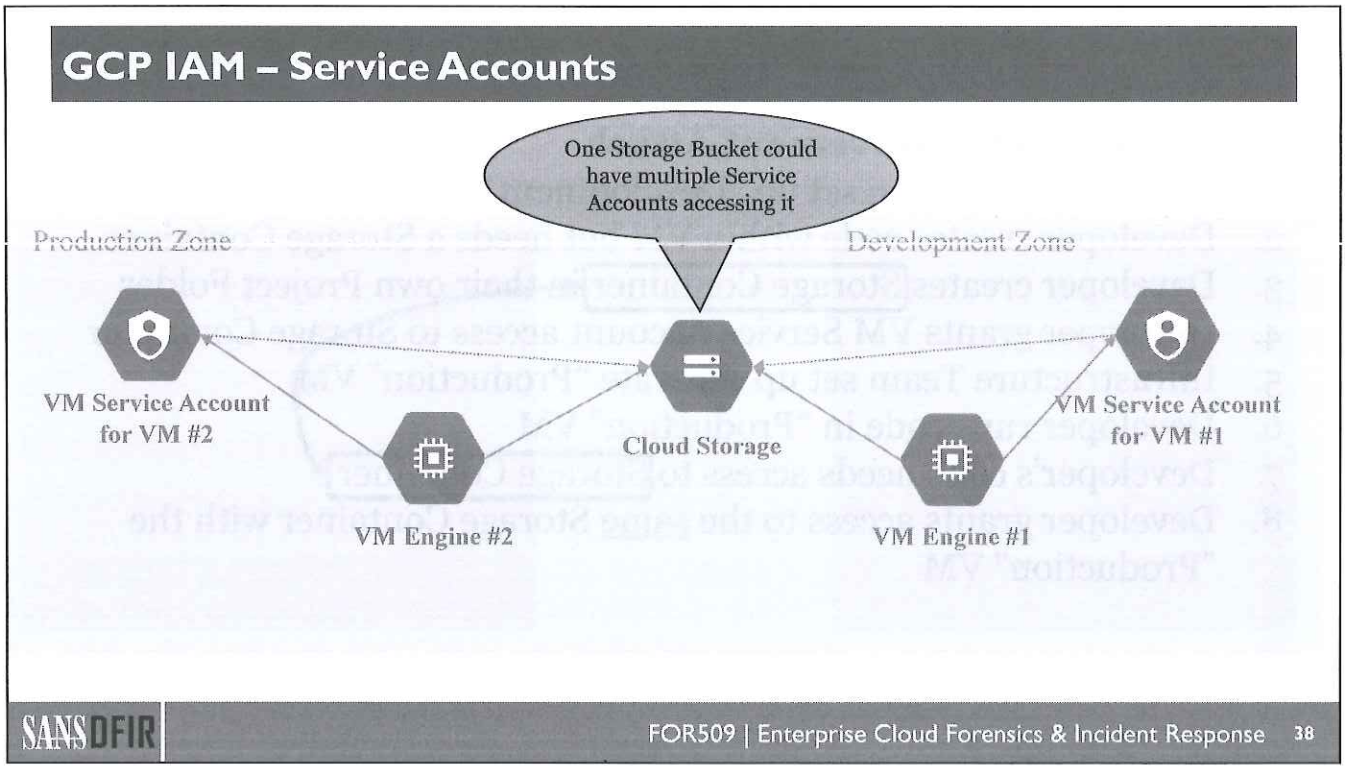

Service accounts are used to access other resources within GCP. The resource you want to access has to allow the service account permission to access it. You would assign a policy to the resource that allows the Service Account permission.

This also means it's hard to control what Service Accounts have access to if a developer grants access to a resource that they control.

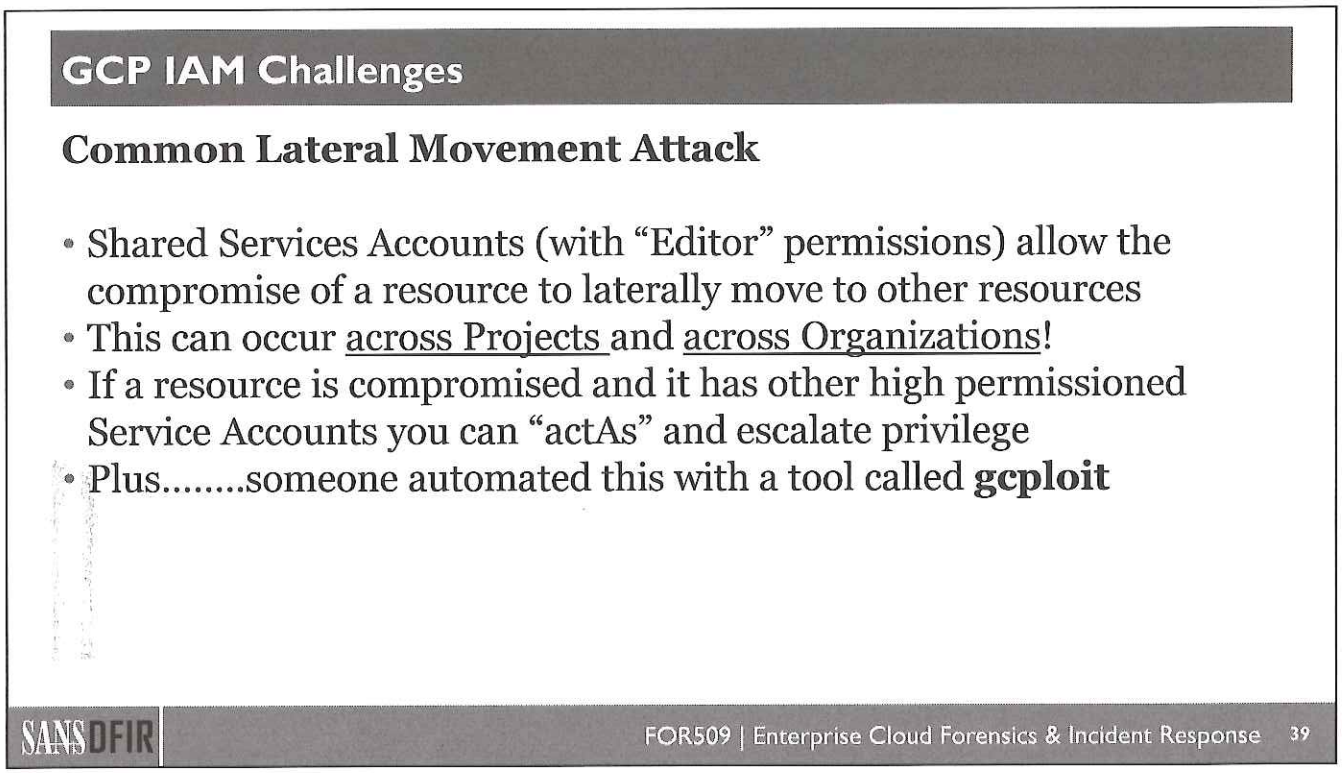

gcploit is written by Dylan Ayrey[1] and then presented by Dylan and Alison D as part of Blackhat/Defcon 2020[2].

References:

[1] https://for509.com/wqfoe

[2] https://for509.com/4quy0

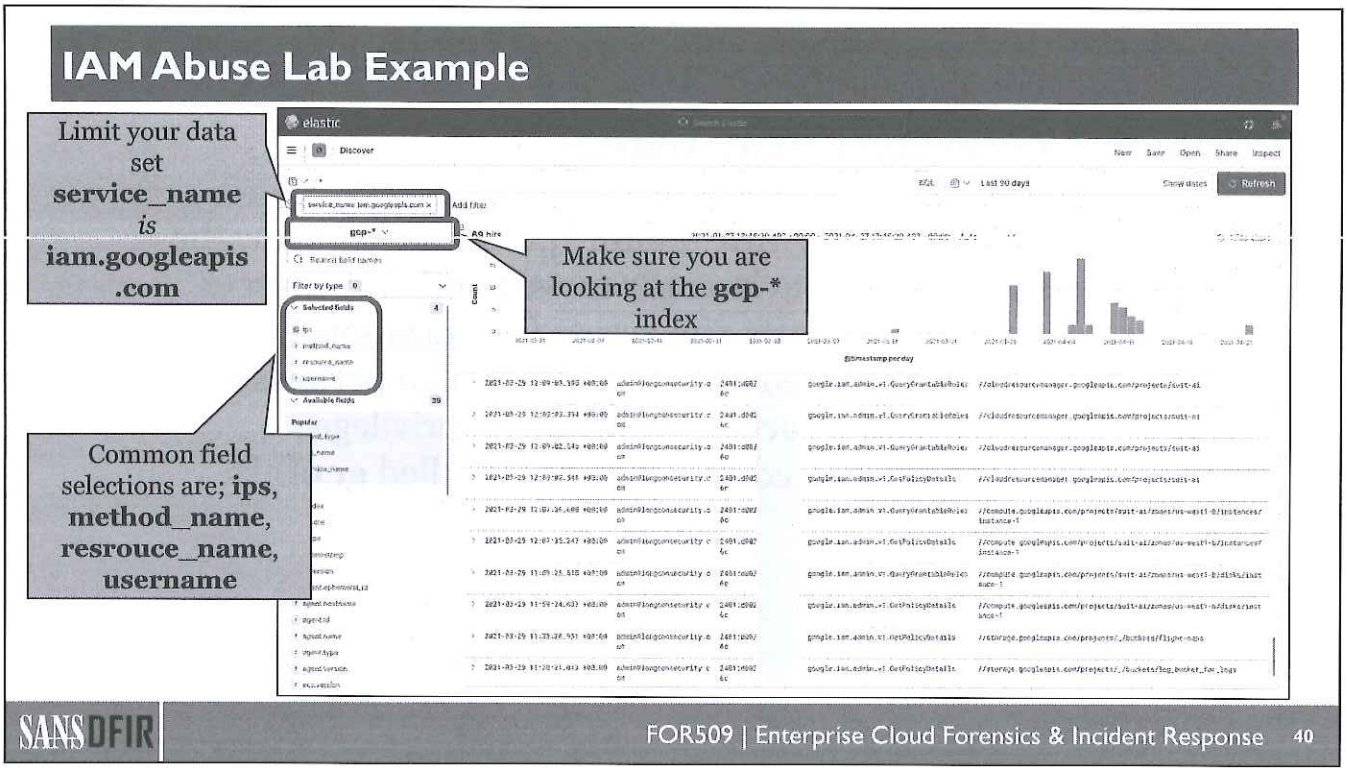

The example above shows you SOF-ELK with the IAM logs extracted from GCP in JSON format and loaded into SOF-ELK. In the Lab, that is coming up, your data set will have a combination of IAM, and other data mixed in, so it will be important to get familiar with the fliting functionality in SOF-ELK.

The data set under the "gcp-*" index will hold all the GCP data, except for Flow Logs. When you're viewing this index, you will see all GCP data loaded in SOF-ELK. Using the "Add filter" button below the search bar, create a new filter that uses the "service_name" field and "is" "iam.google.apis.com". This will limit your data to only IAM logs. When reviewing IAM actions this filter is simply the starting point, however, when you're looking at malicious IAM activity you will also need to pivot to other logs to determine actions that may have been taken by a threat actor.

The filter function at the top will allow you to enable and disable the filter by simply left clicking on it to "Temporarily disable" if you want to see all GCP data related to your search in the search bar at the top.

In addition to using the filter you may also want to add specific fields to see as summary values for the log entries. For IAM events, the common fields that are useful are; @time, username, ips, method_name, and resource_name. These will provide you with a summary view of each log entry to better understand what is occurring in the logs with a quick overview.

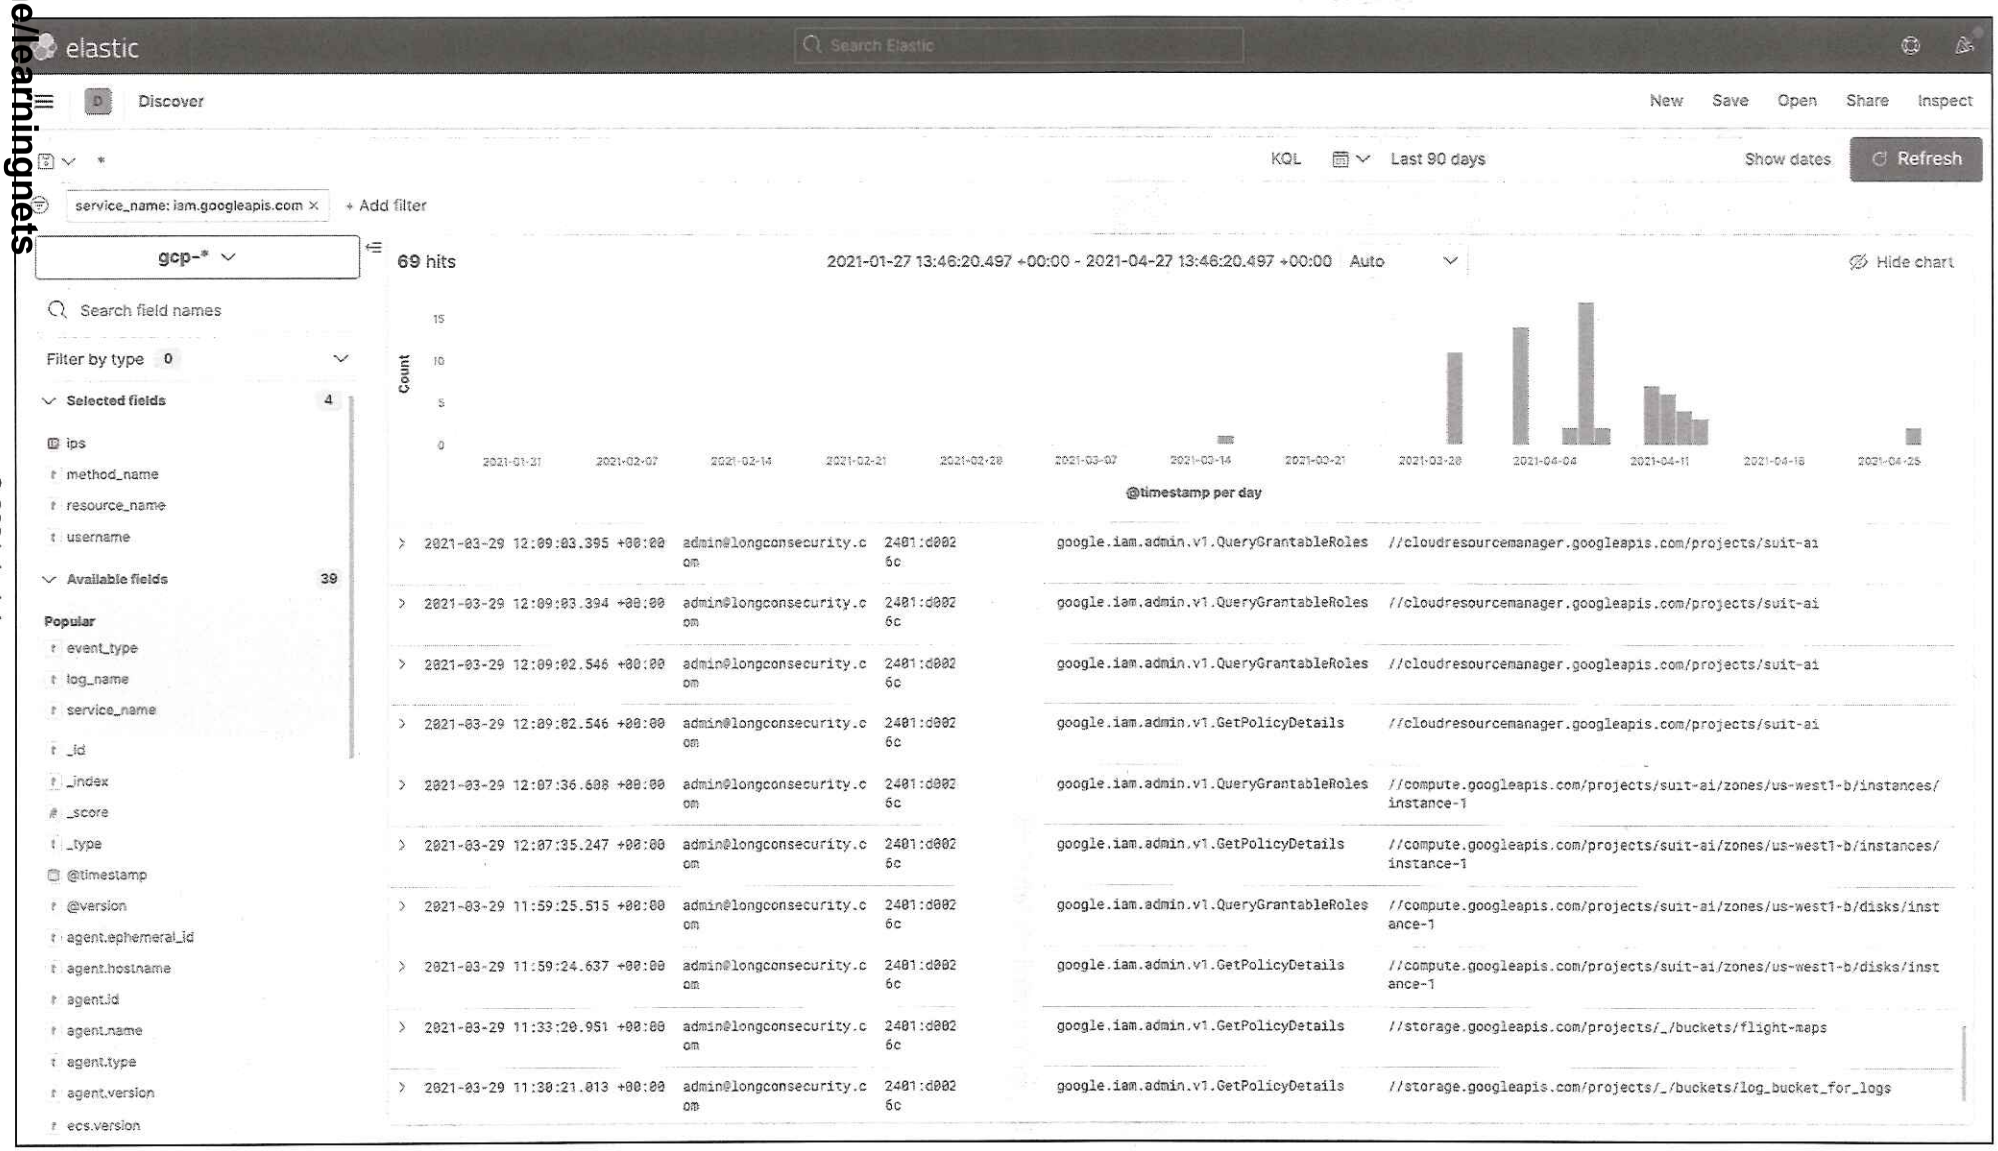

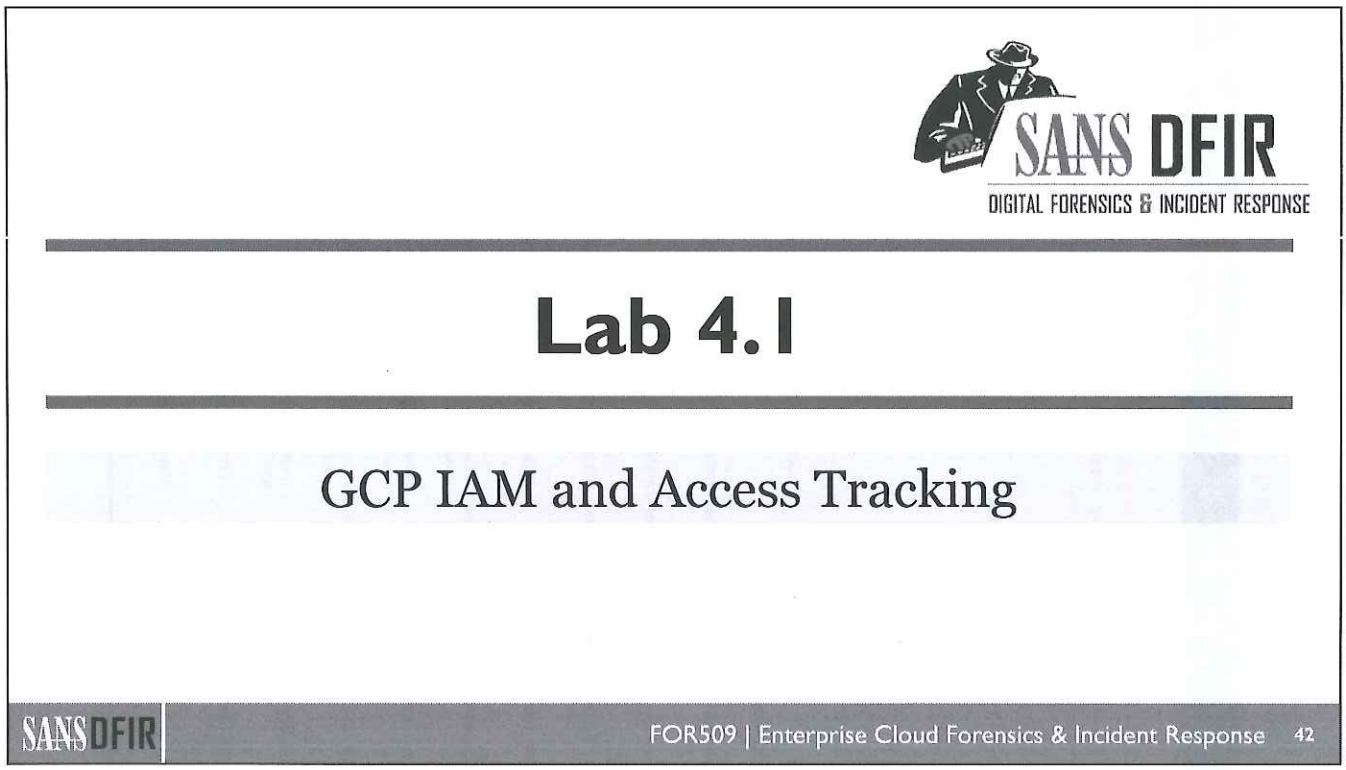

This page intentionally left blank.

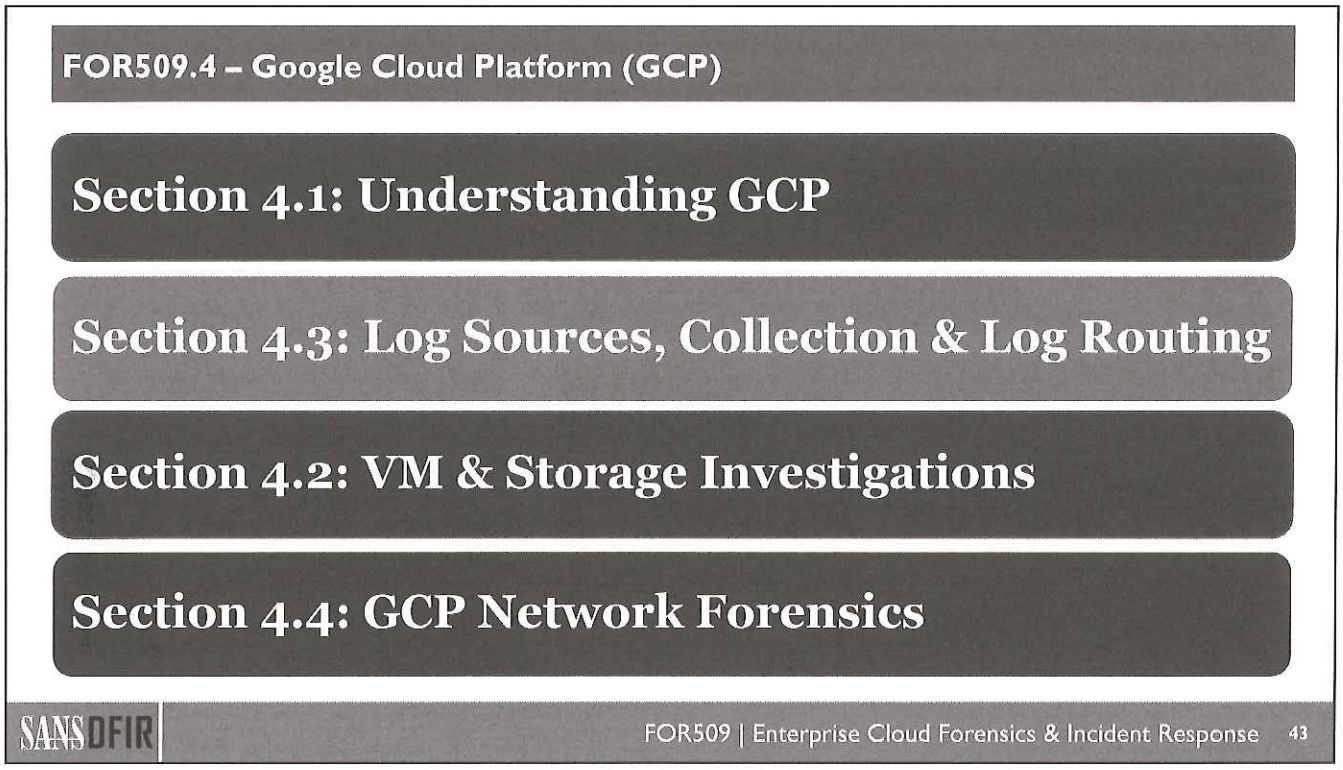

This page intentionally left blank.

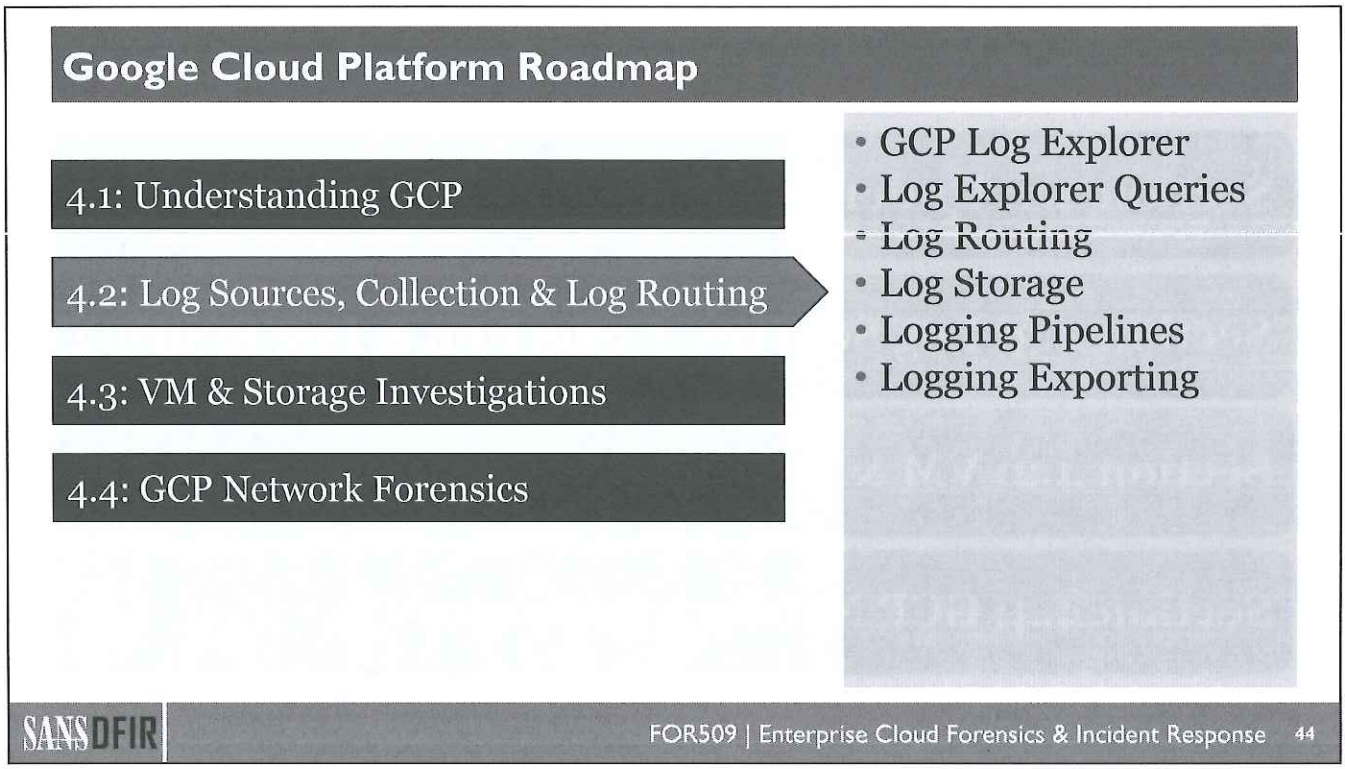

This page intentionally left blank.

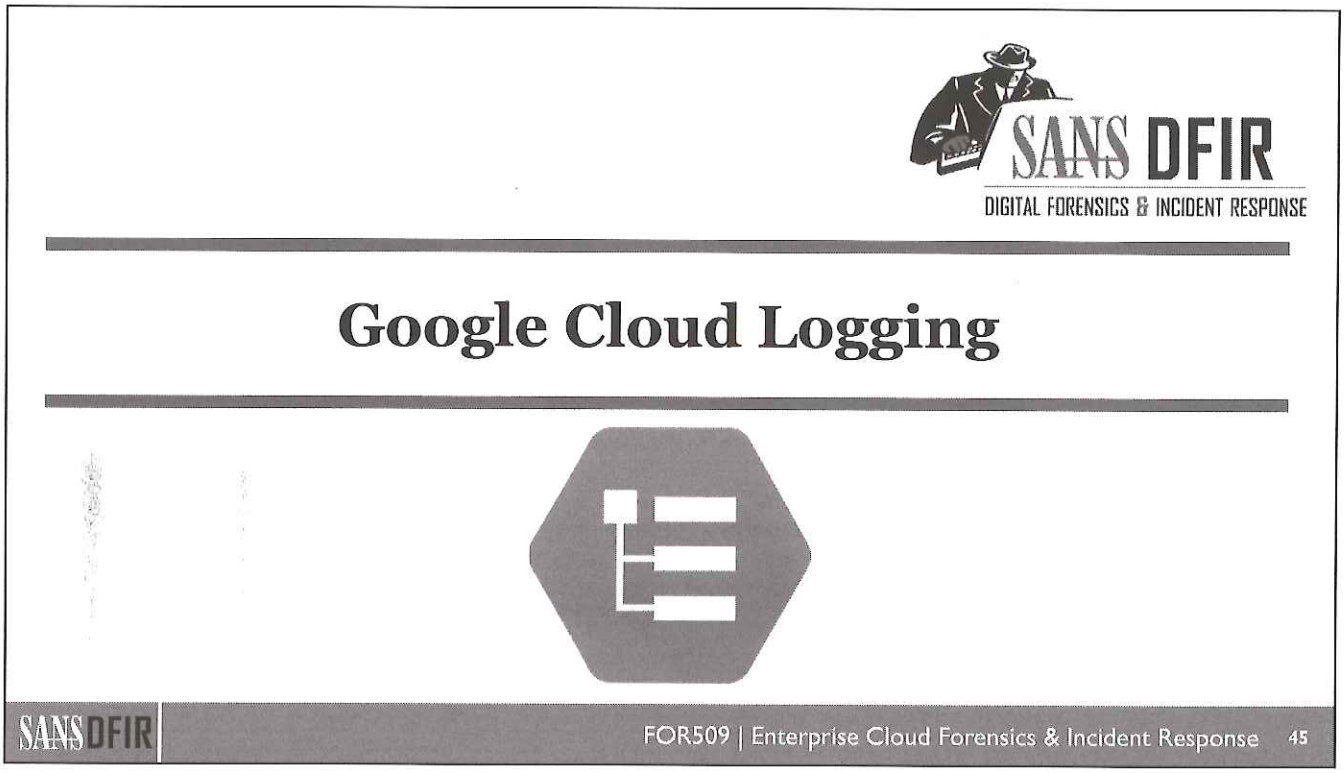

This page intentionally left blank.

Enabling Auditing & Logging (1)

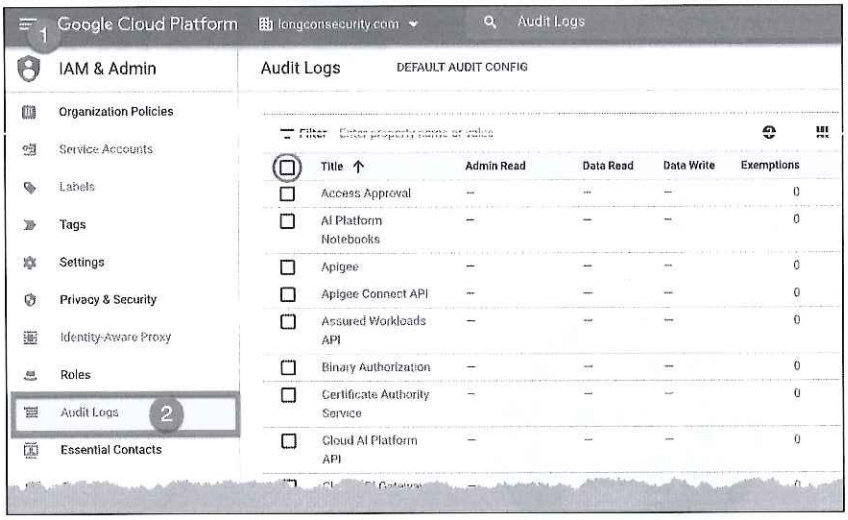

NOTE

Getting Logging inside of GCP works in two ways, either;

- **Audit Logs** from the Platform, or
- You ship logs from another application / system / service

This page intentionally left blank.

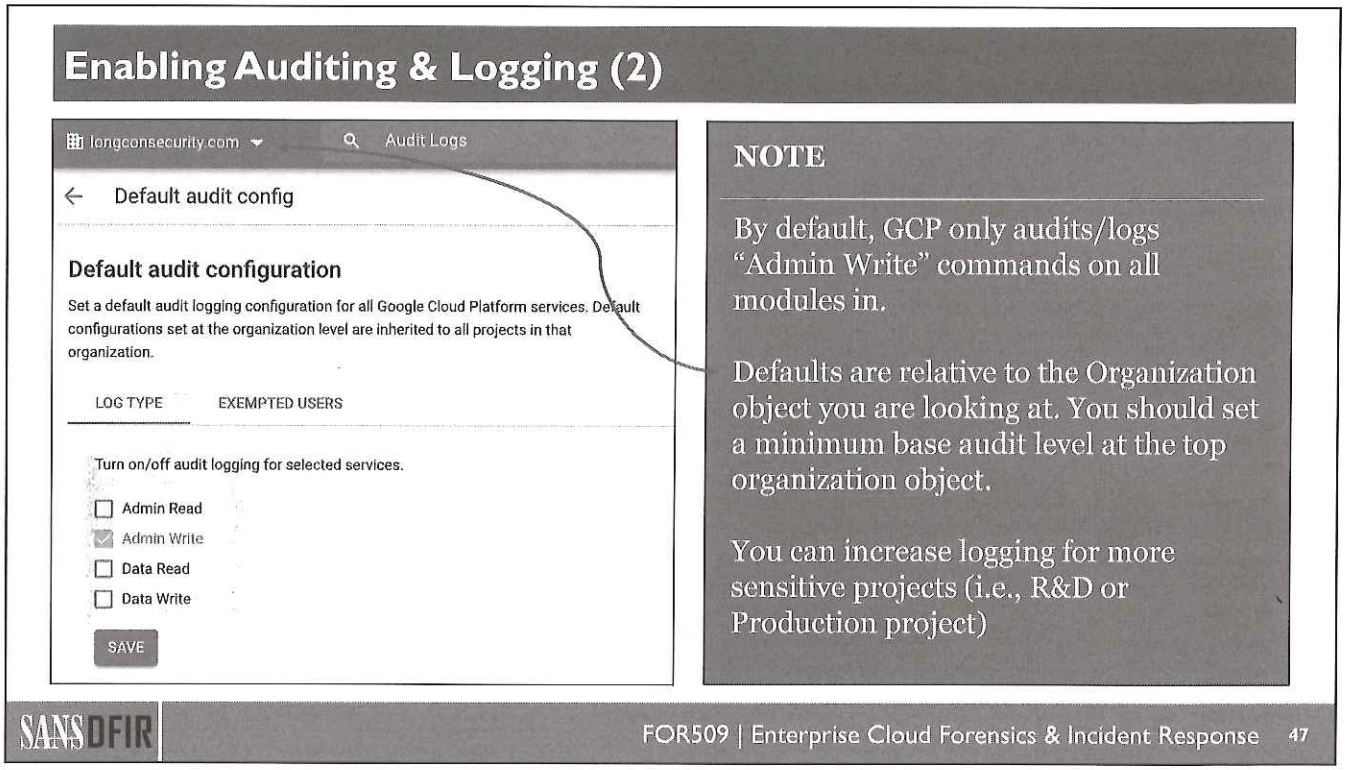

This page intentionally left blank.

Accessing Logging

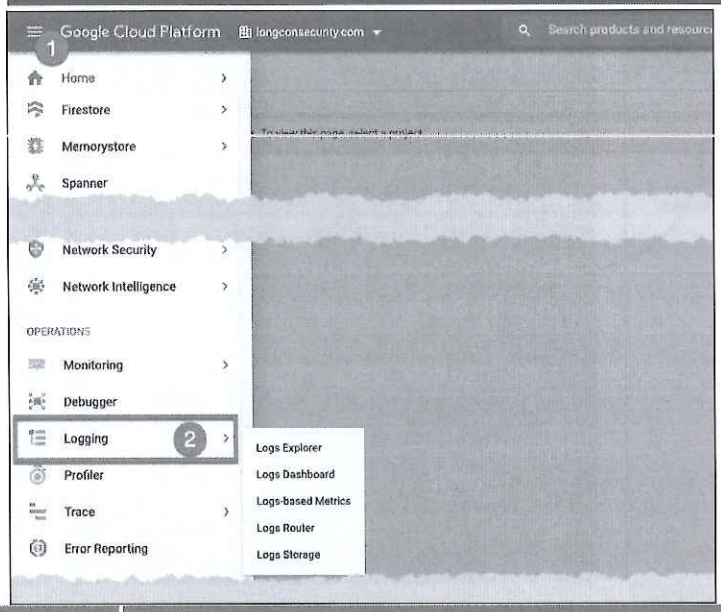

NOTE

GCP has recently taken significant steps to upgrade the production, collection and searching available for logs.

- Accessing and Searching Logs requires Permission
- Logs are Stored and Accessed per Project

This page intentionally left blank.

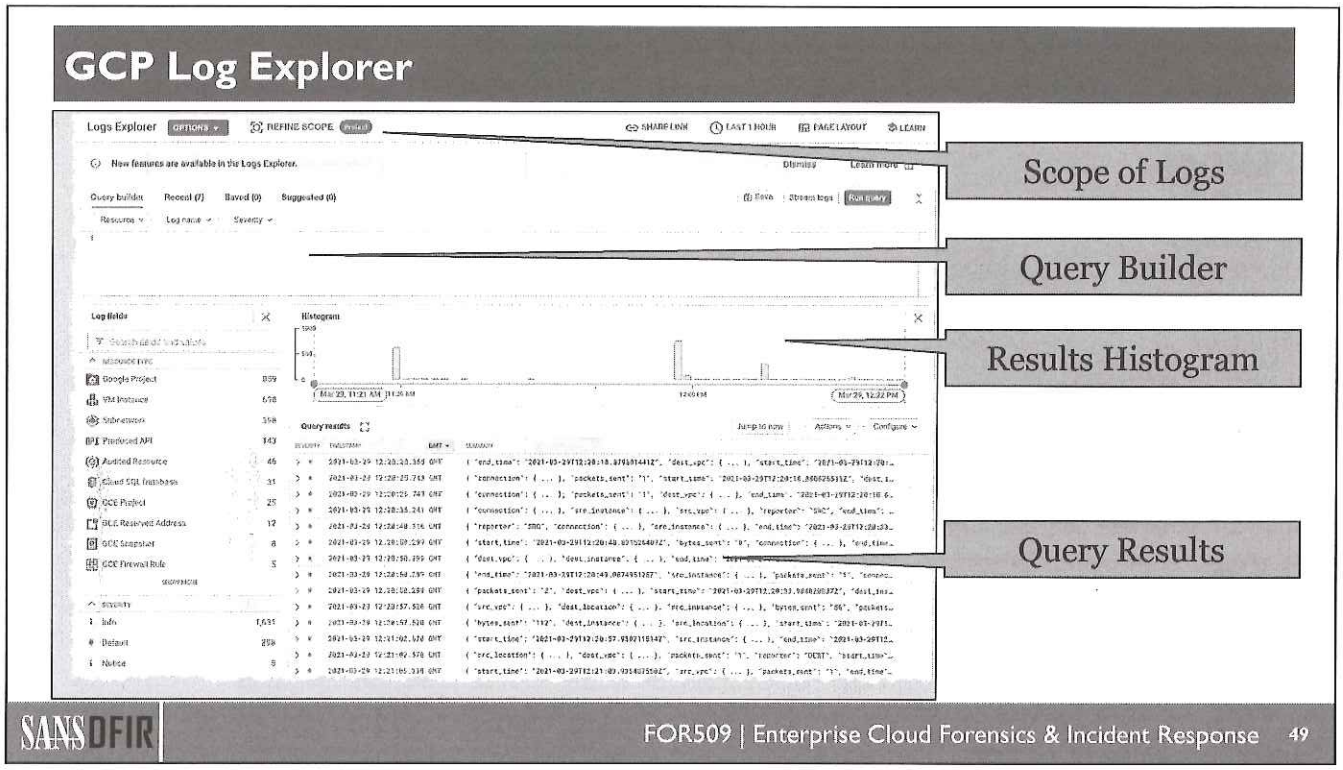

This page intentionally left blank.

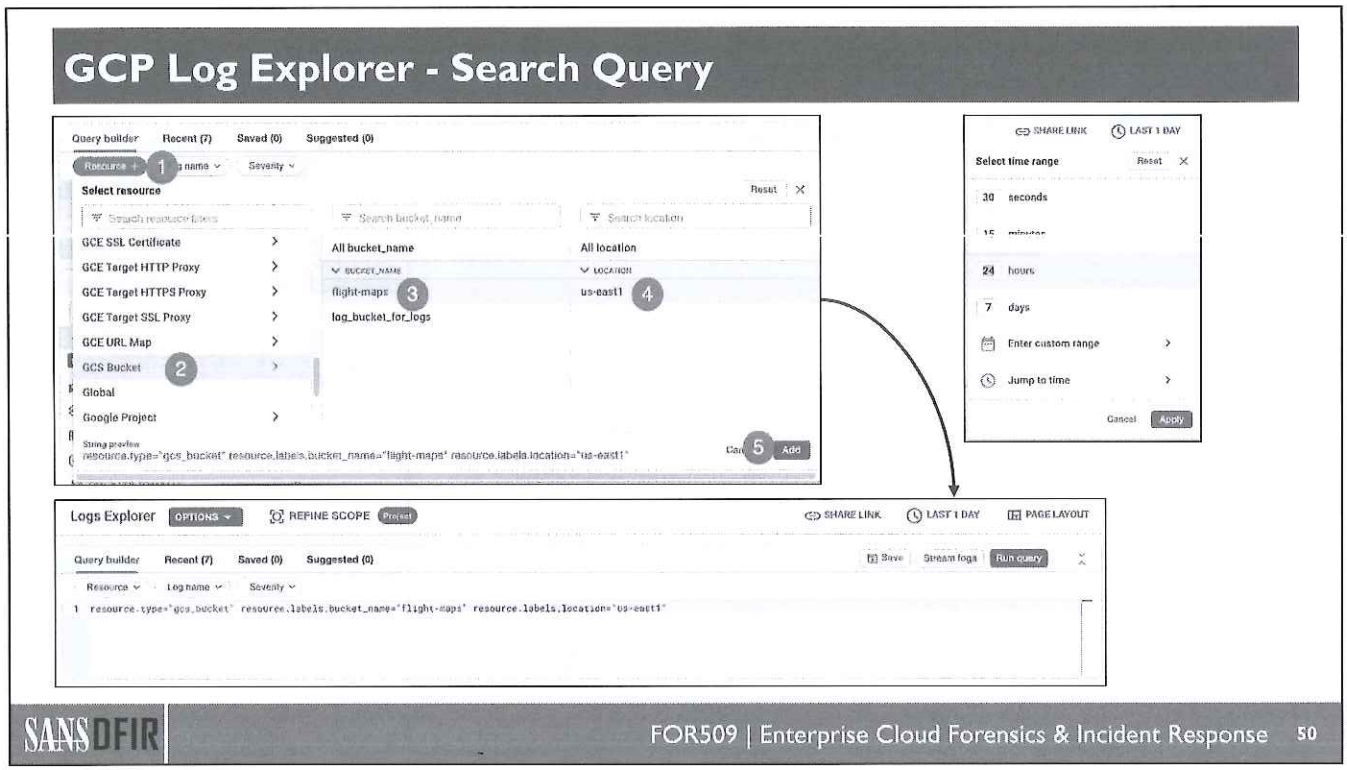

This page intentionally left blank.

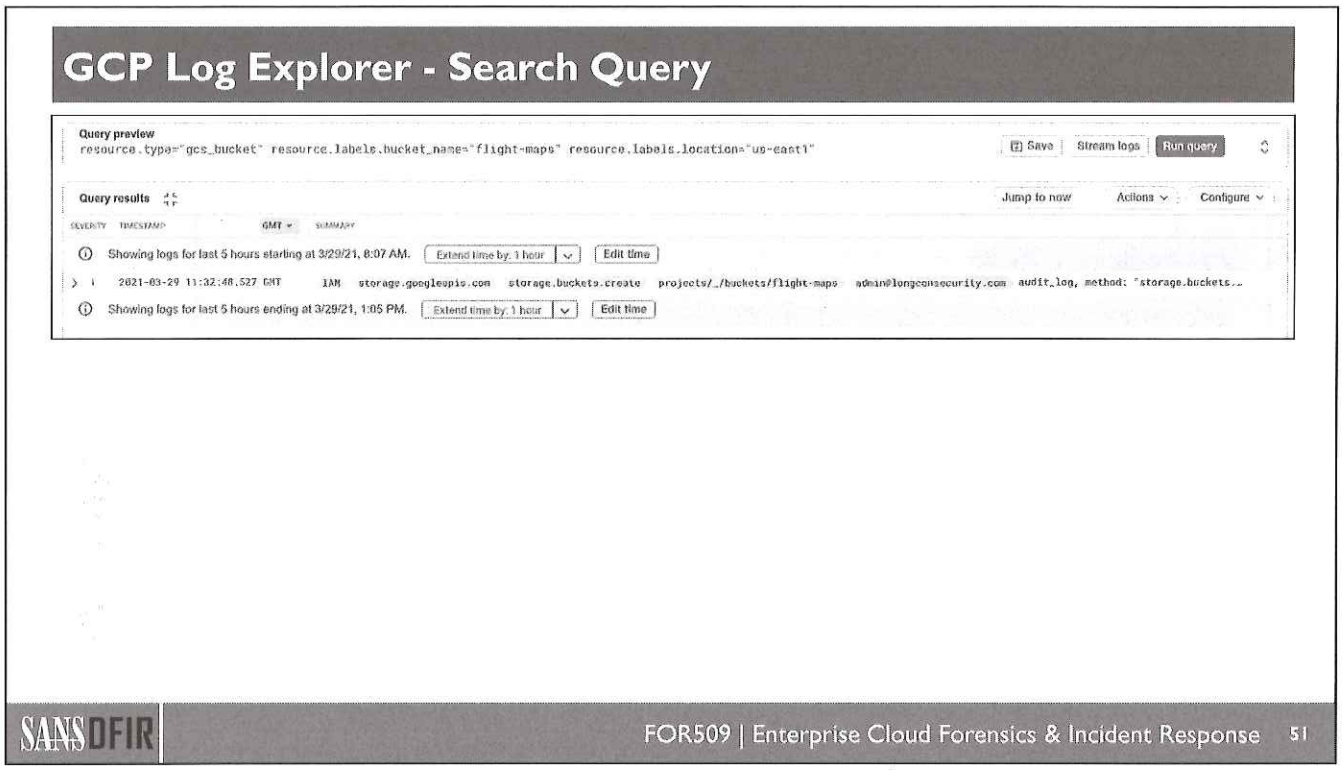

This page intentionally left blank.

GCP Log Explorer - Search Query

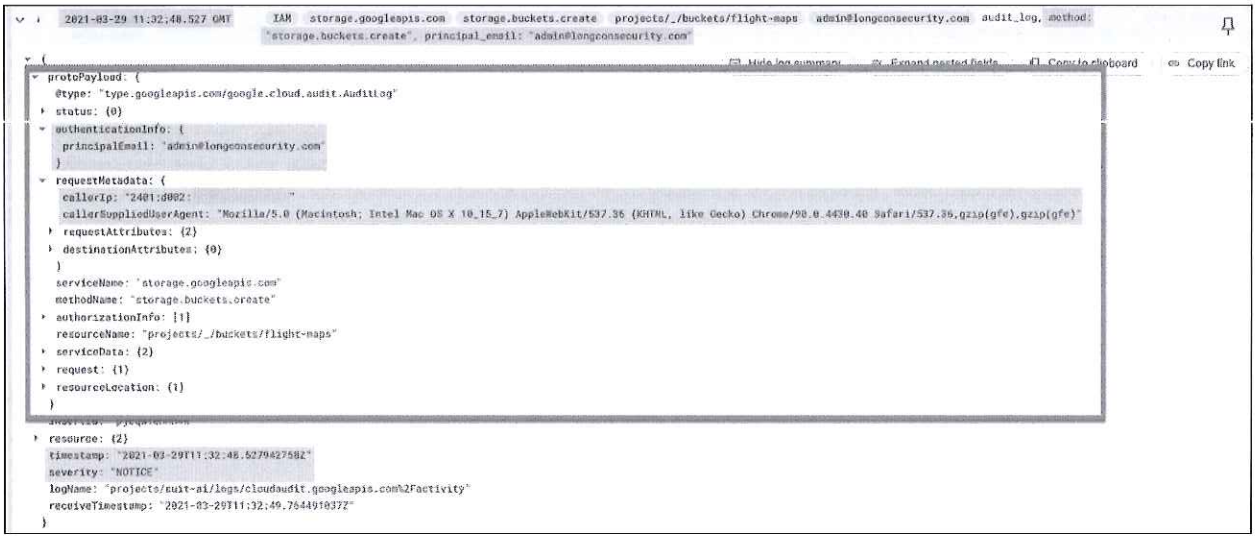

This page intentionally left blank.

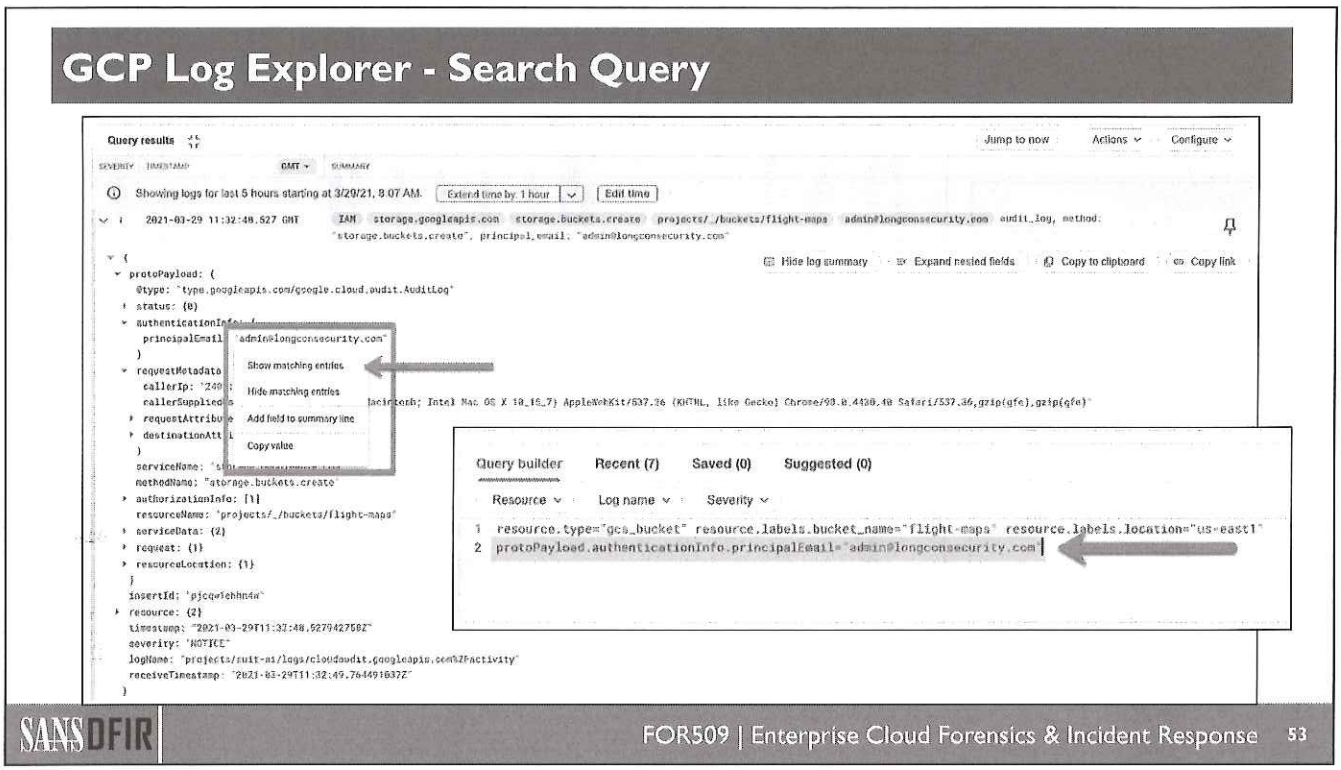

This page intentionally left blank.

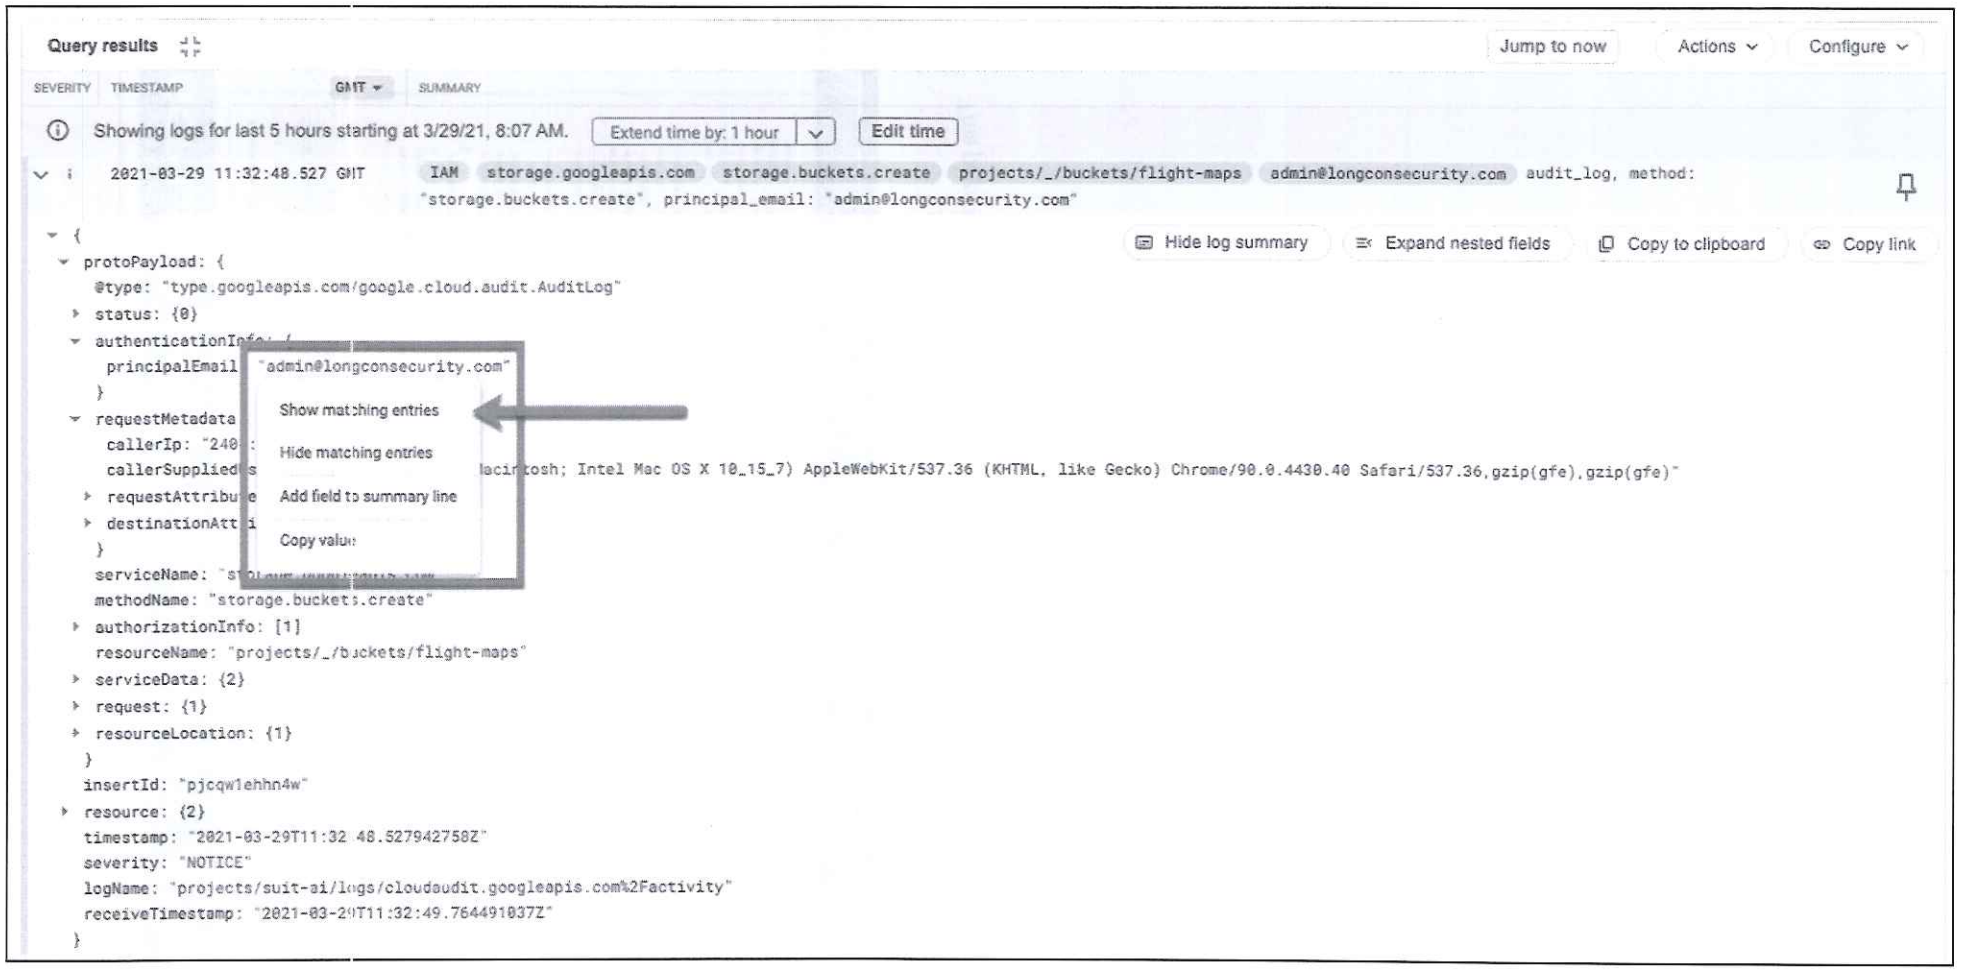

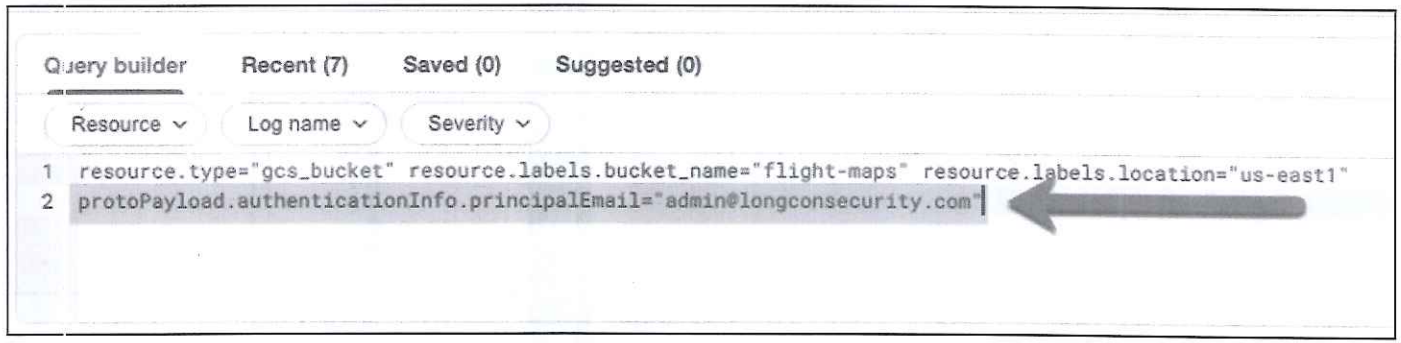

GCP Logging Principals

- All logging is sent to a Google Cloud Logging API
- Google's Cloud Logging API sends all logs to "Log Sinks"
- Log Sinks exist at the Project and Organization level
- Log Sinks are used to determine if a log entry should be routed by the Sink or dropped.
- Log Sink "Exlusions

References:

[1] https://for509.com/tm3b0

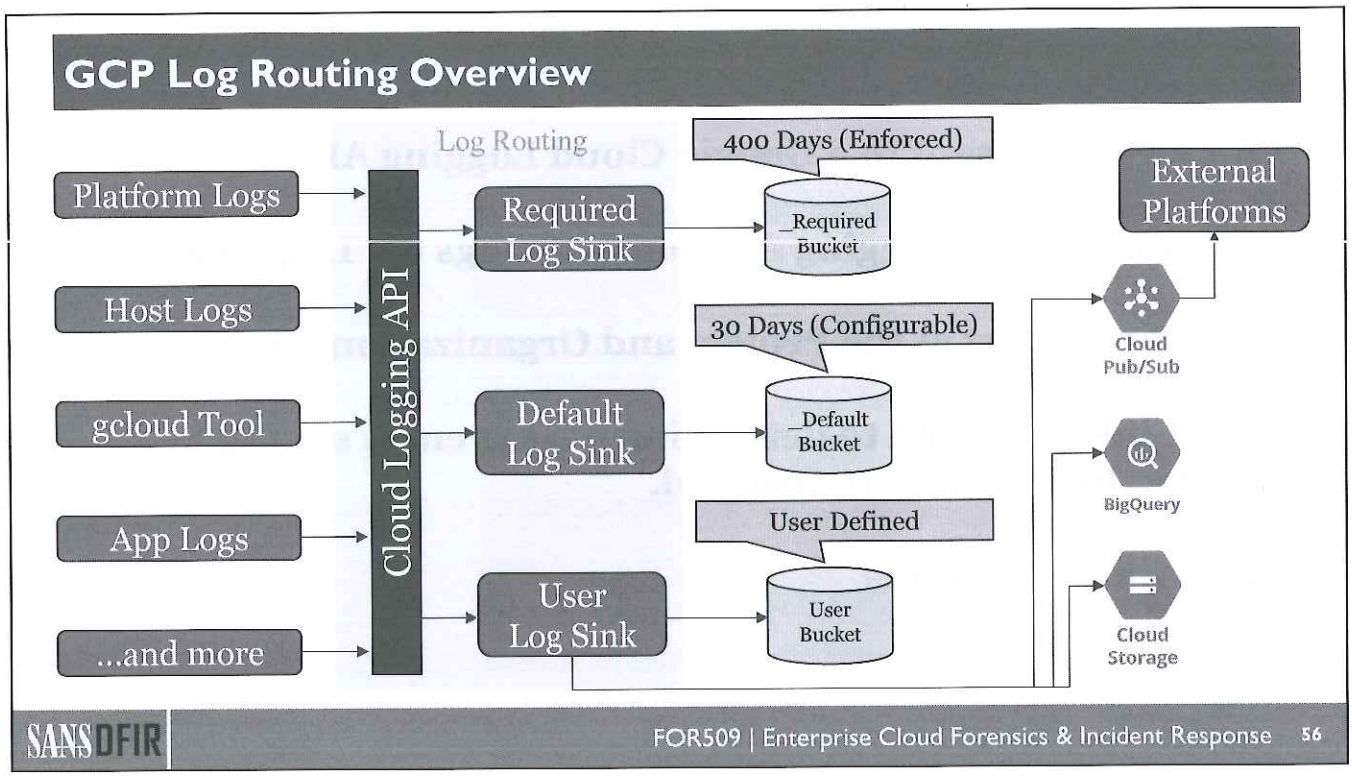

References:

[1] https://for509.com/tm3b0

GCP Logging Storage

- The "*_Required Bucket*" and "*_Required Sink*" cannot be modified.
  - This does not count towards the costs in your GCP Billing
- The "*_Default Bucket*" and "*_Default Sink*" can be modified.
  - This does count towards the costs in your GCP Billing
- Logging Buckets can hold logs between 1 and 3650 days
  - Yes – that's over 9 years!
- Log Routing and Storage can be encrypted using your own keys (CMEK)

References:

[1] https://for509.com/8yt5o

[2] https://for509.com/4unh8

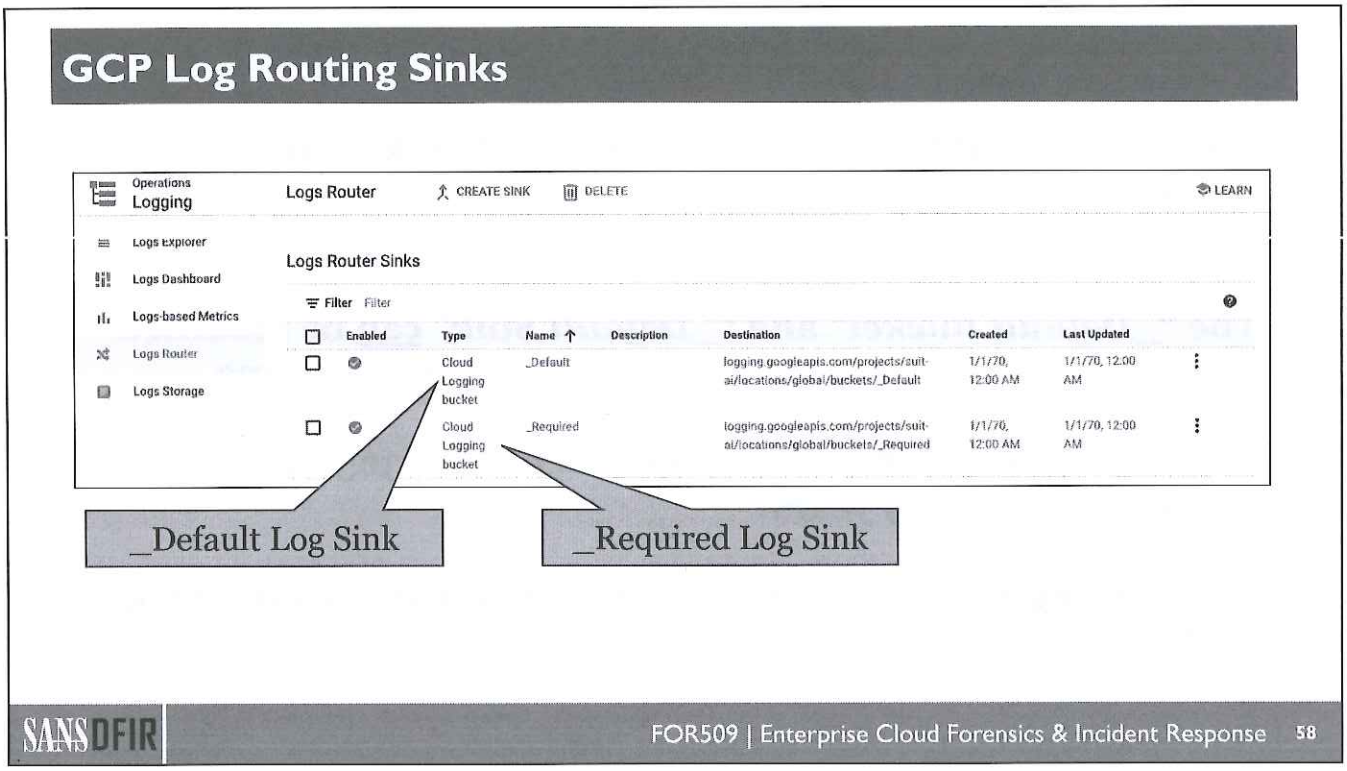

This page intentionally left blank.

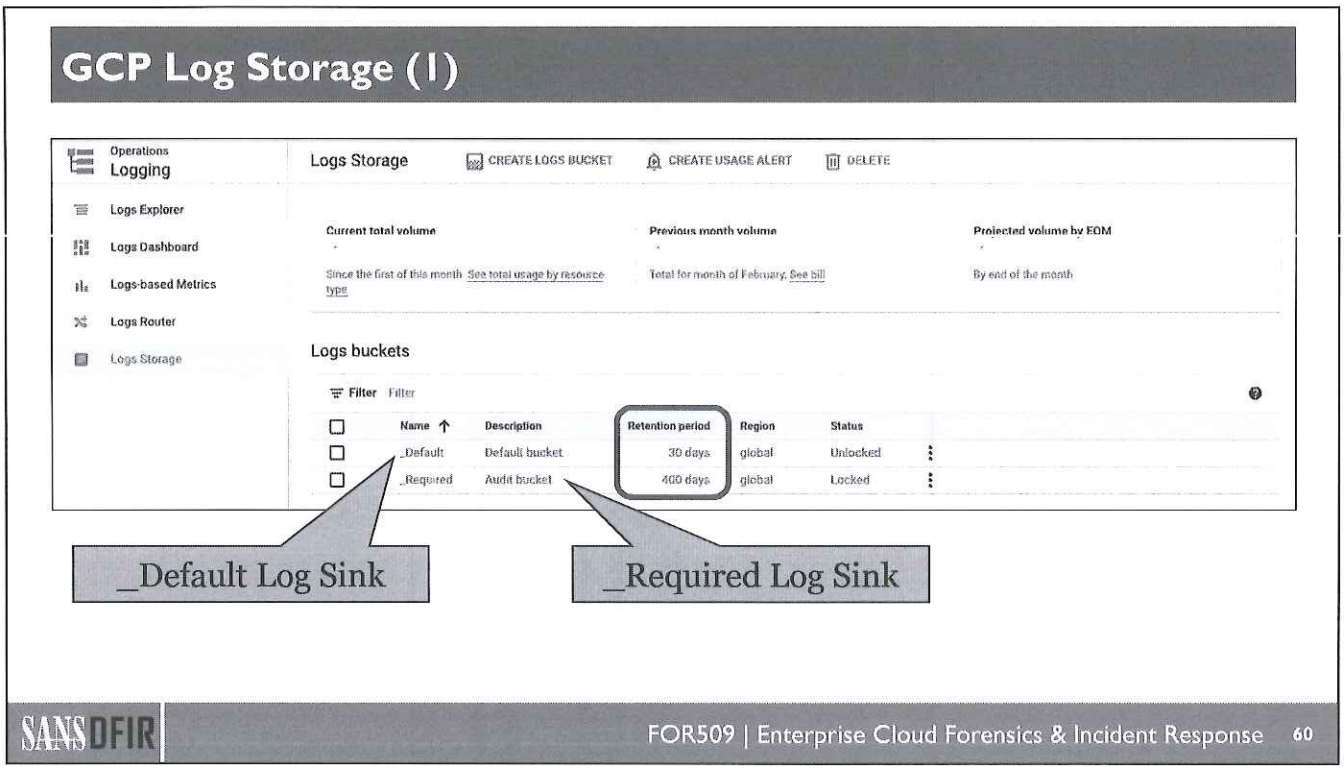

This page intentionally left blank.

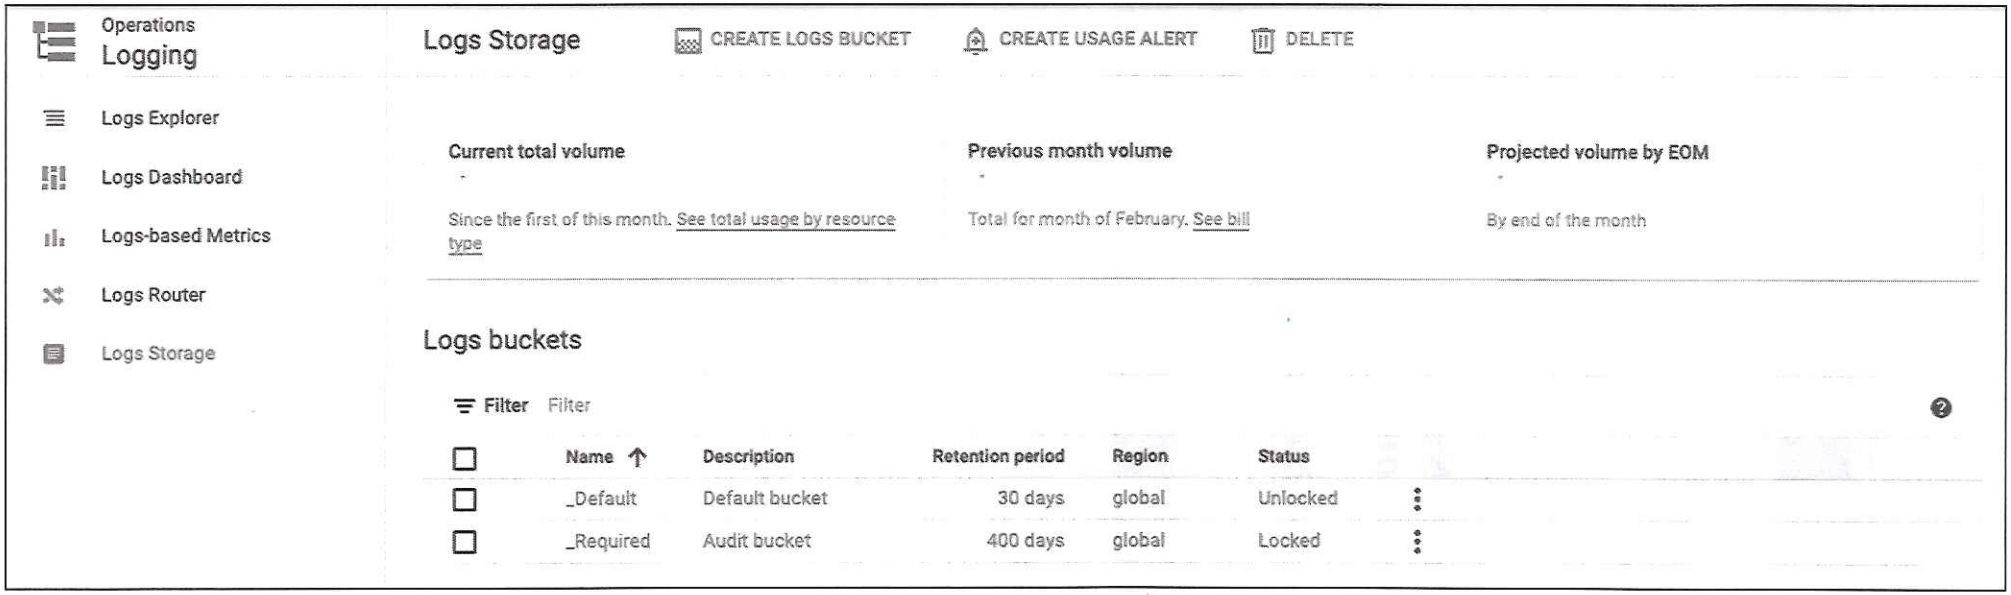

GCP Log Storage (2)

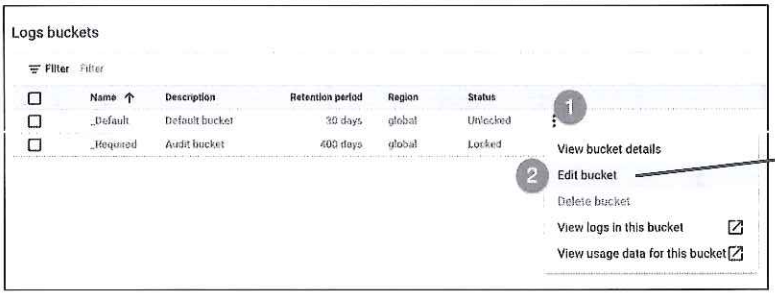

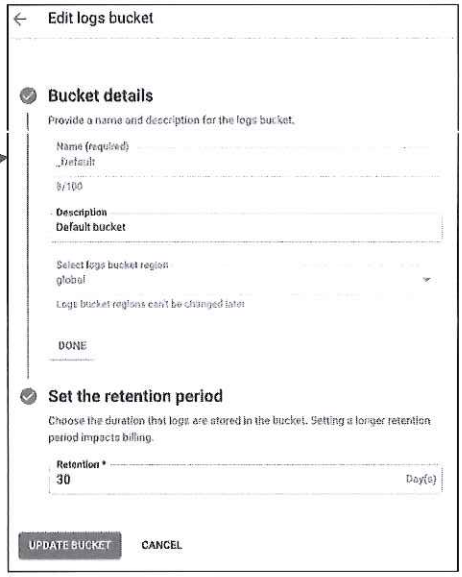

- **Retention time and Name can be changed after creation.**
- **Bucket Location is fixed after initial creation.**

This page intentionally left blank.

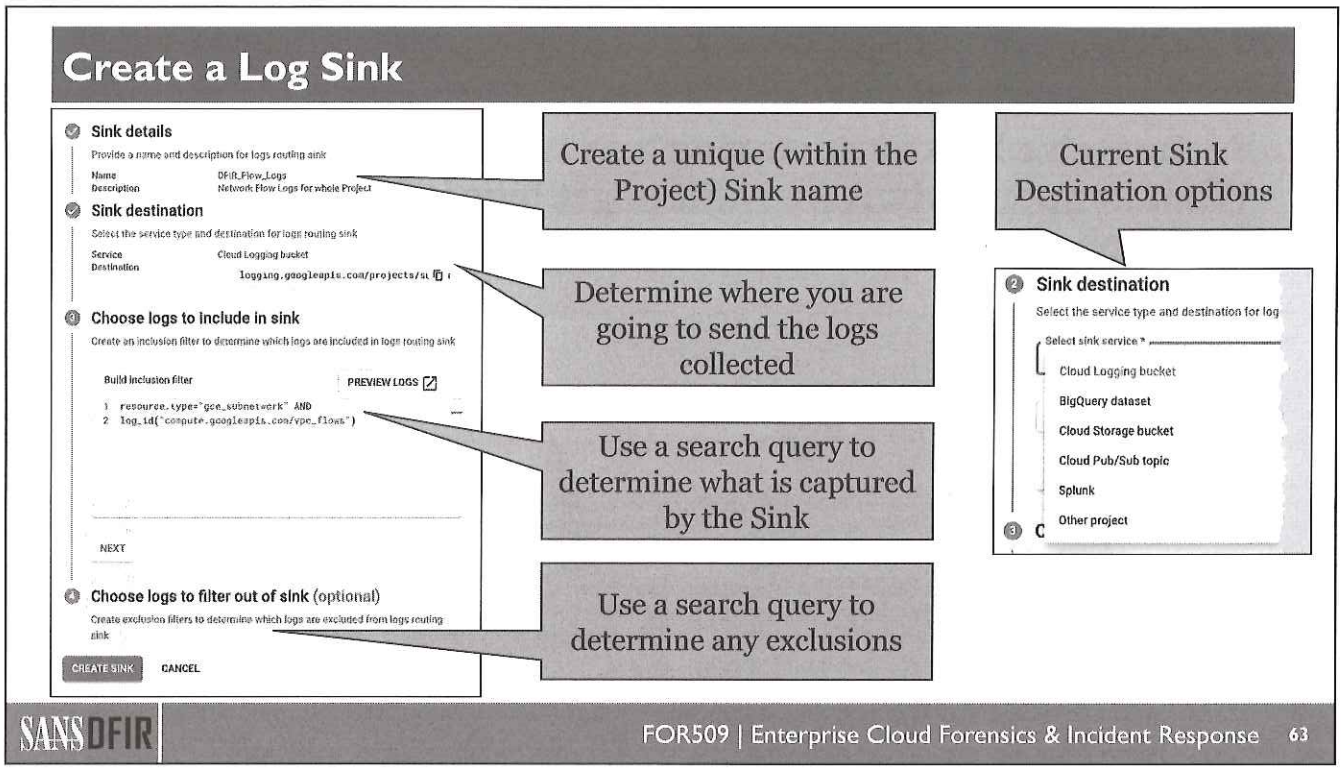

Query Used: *resource.type="gce_subnetwork" AND log_id("compute.googleapis.com/vpc_flows")*

Reference:

[1] https://for509.com/gsuzk

Exporting Logs Outside of GCP

- **So far we've looked at storing logs within GCP**
- **What happens when we want to get logs into SOF-ELK/Splunk**

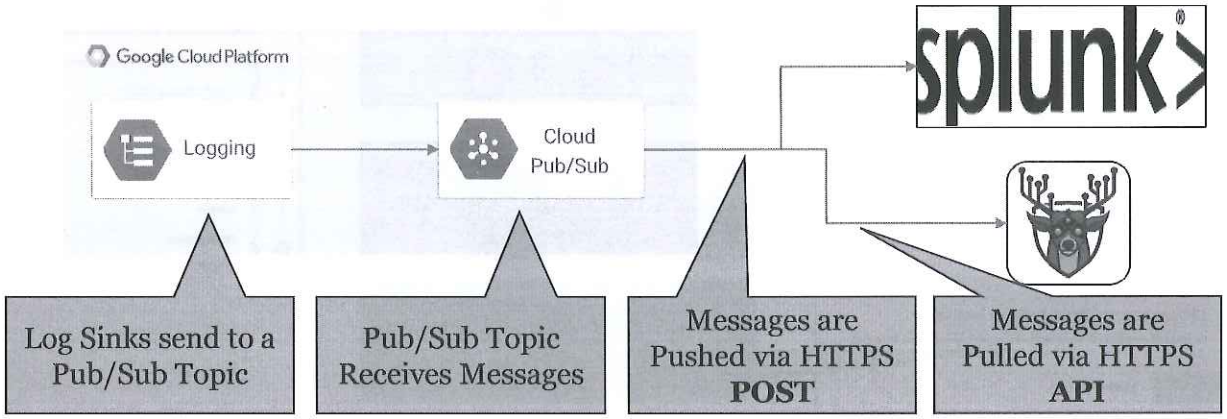

References:

[1] https://for509.com/bp54a

[2] https://for509.com/r12cx

[3] https://for509.com/knxcj

[4] https://for509.com/1slyv

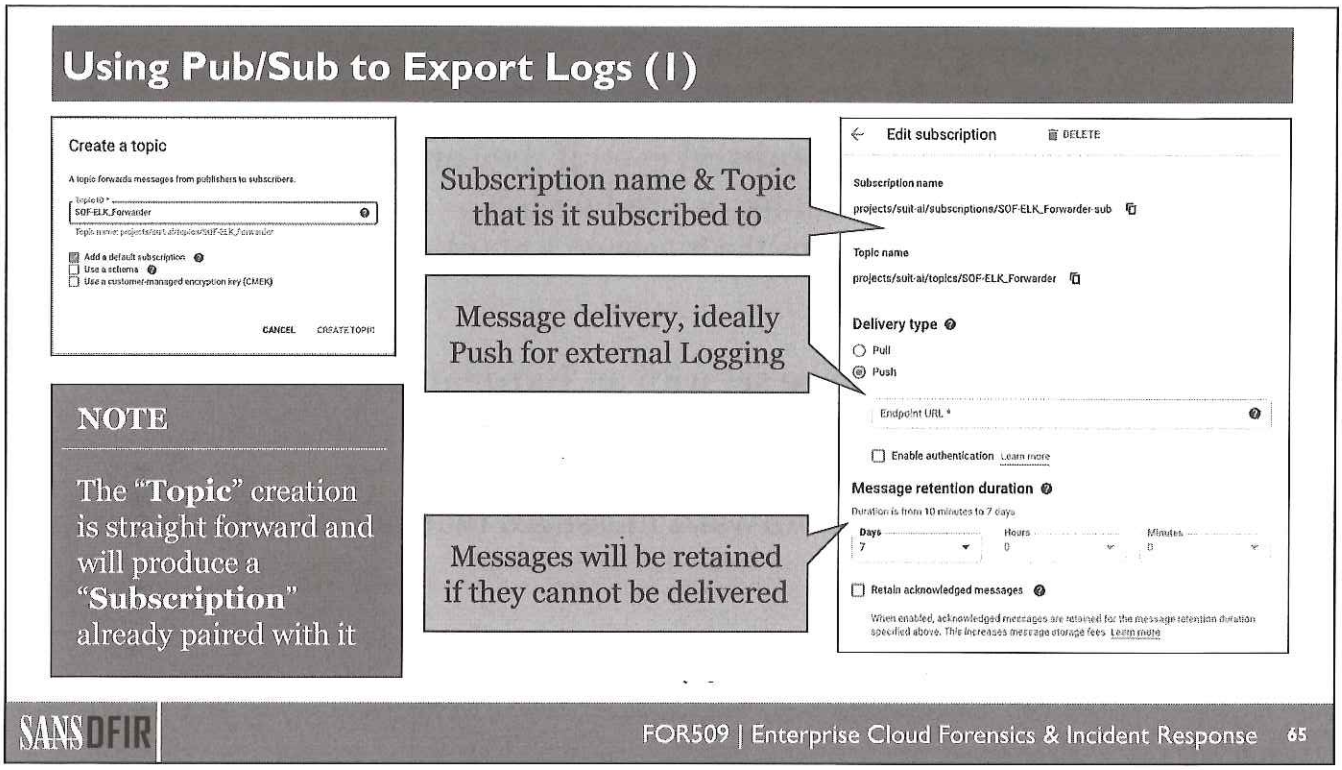

This page intentionally left blank.

Using Pub/Sub to Export Logs (2)

- **Messages** can also be **Retained** even if they are delivered successfully, this has a cost associated with it.
  - This will follow the time set by the **Message Retention Duration** for undelivered messages
- **Subscriptions** can also be set with an **Expiration Period** if you want them removed after receiving no messages.
- **Subscriptions** can auto-forward messages that cannot be delivered using **Dead Lettering**

This page intentionally left blank.

Pub/Sub Push Message Type - Example

- **Pushed** messages over HTTP use the below data structure in a POST request from Pub/Sub.

```
{
    "message": {
        "attributes": {
            "key": "value"
        },
        "data": "SGVsbG8gQ2xvdWQgUHViL1N1YiEgSGVyZSBpcyBteSBtZXNzYWdlIQ==",
        "messageId": "2070443601311540",
        "message_id": "2070443601311540",
        "publishTime": "2021-02-26T19:13:55.749Z",
        "publish_time": "2021-02-26T19:13:55.749Z",
    },
   "subscription": "projects/myproject/subscriptions/mysubscription"
}
```


References:

[1] https://for509.com/r12cx

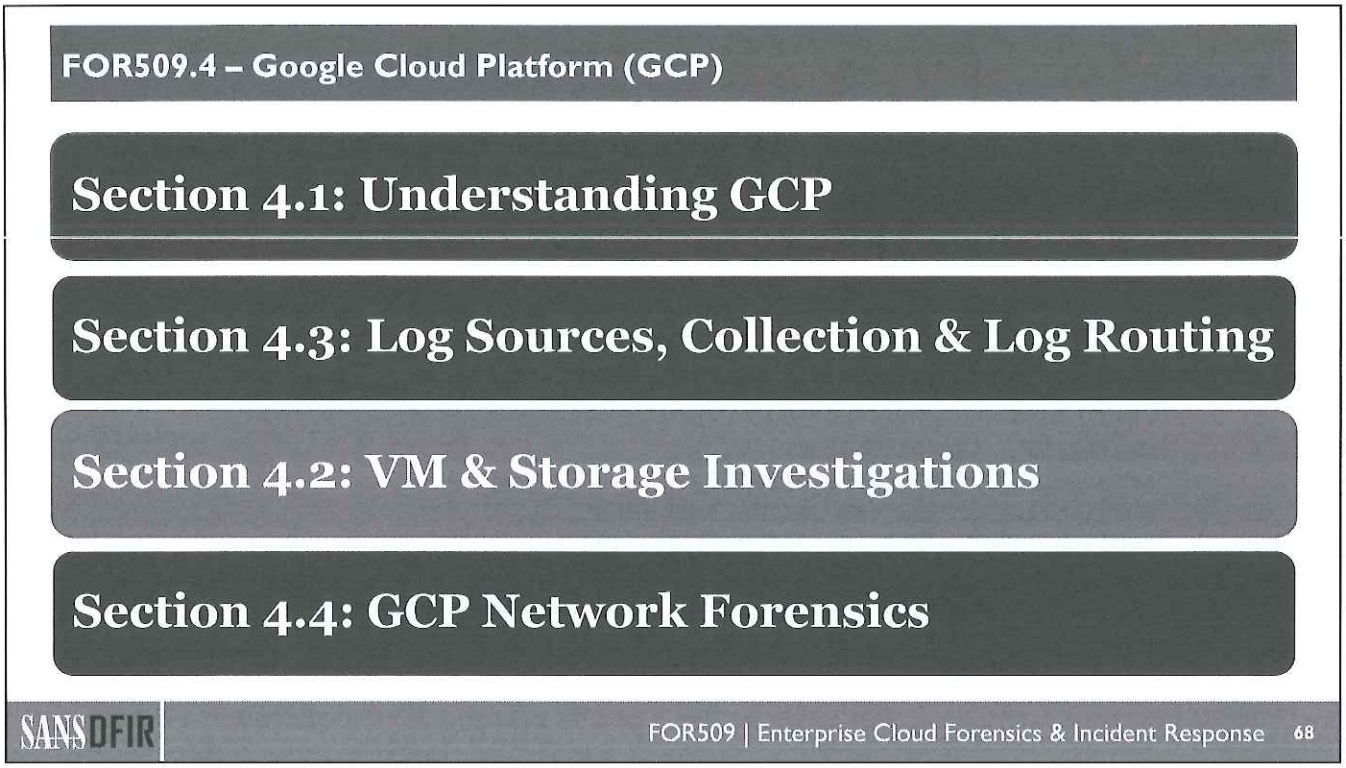

This page intentionally left blank.

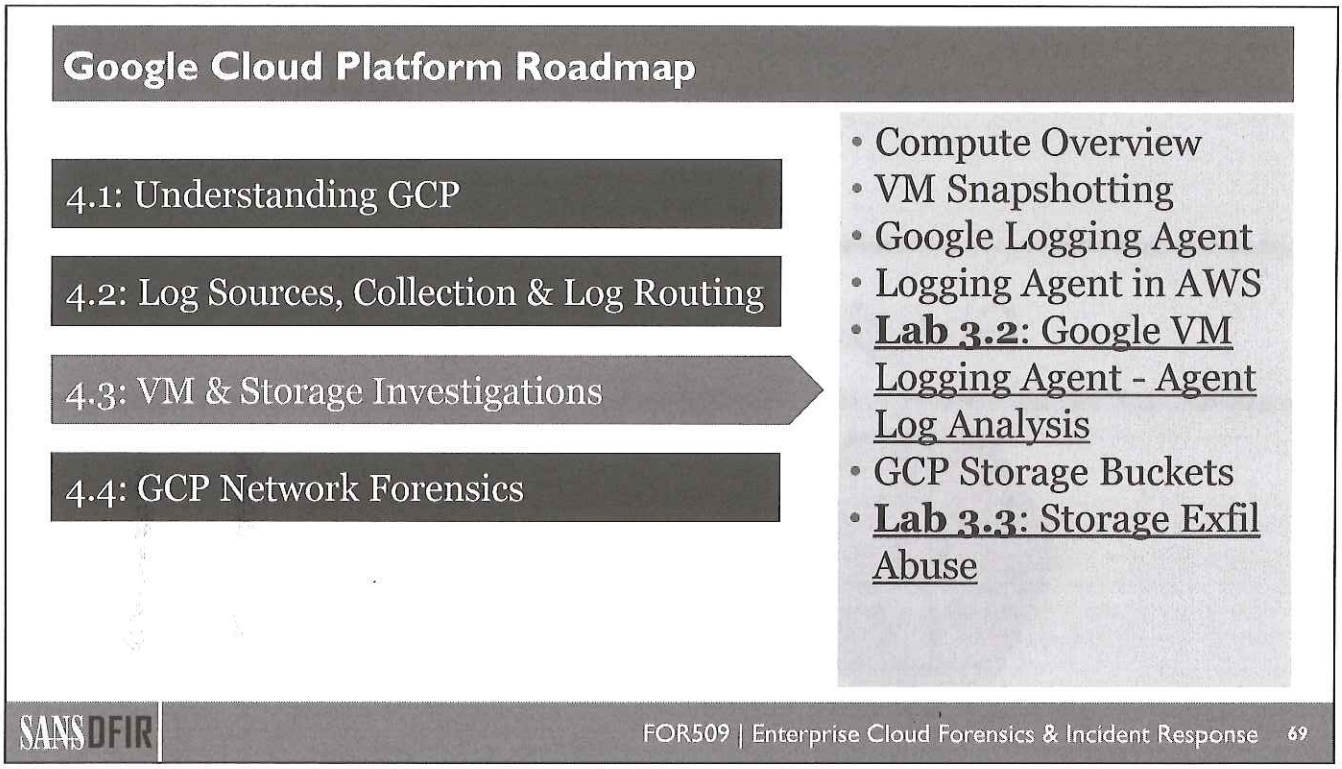

This page intentionally left blank.

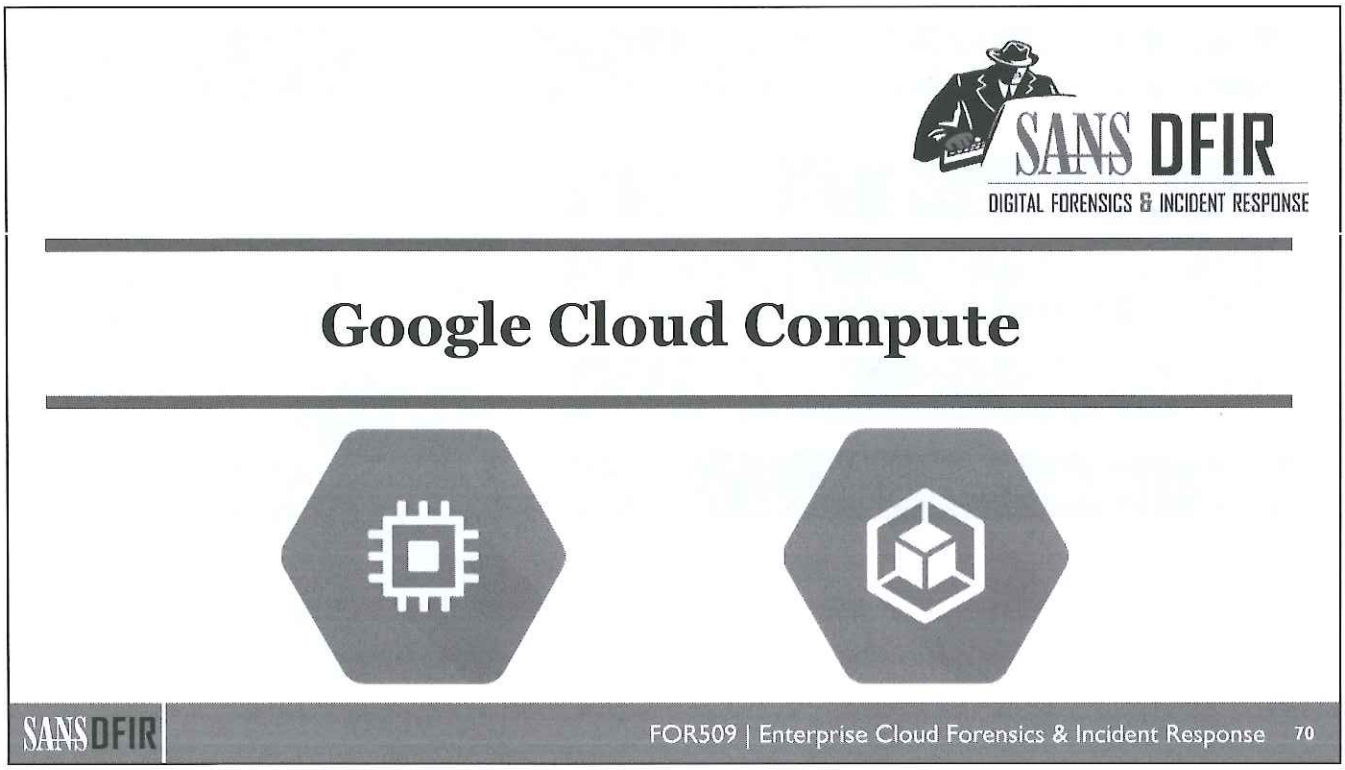

This page intentionally left blank.

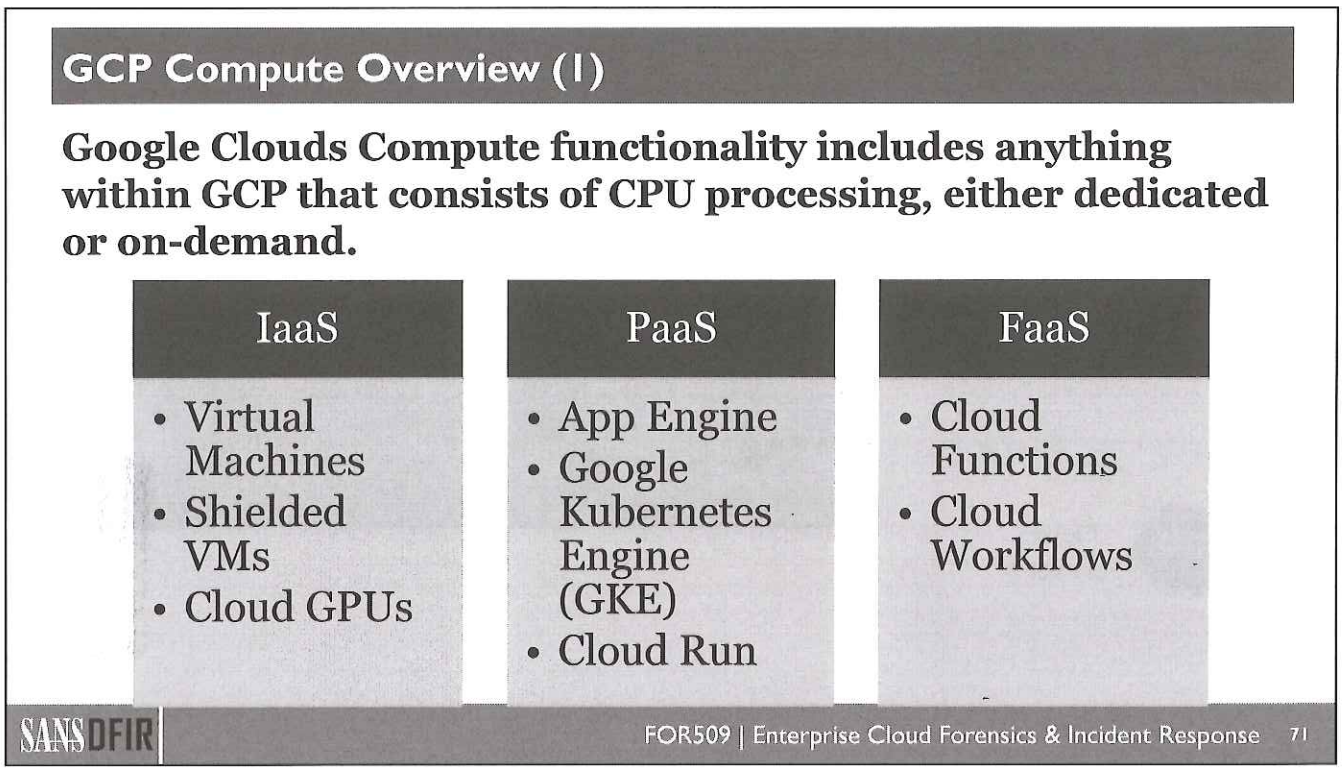

This page intentionally left blank.

GCP Compute Overview (2)

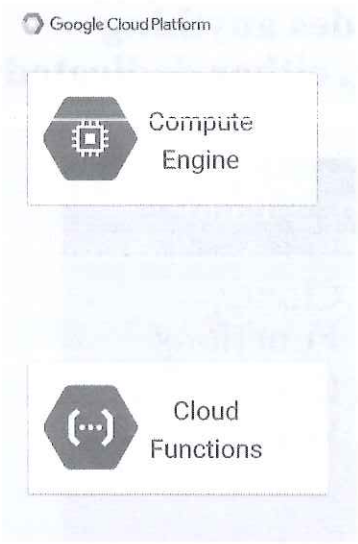

NOTE

- There is an ever-growing number of **compute type resources** in GCP
- For the purpose of this class, we're going to **only focus on two.**
- The principals for these **will apply to many of the others.**

This page intentionally left blank.

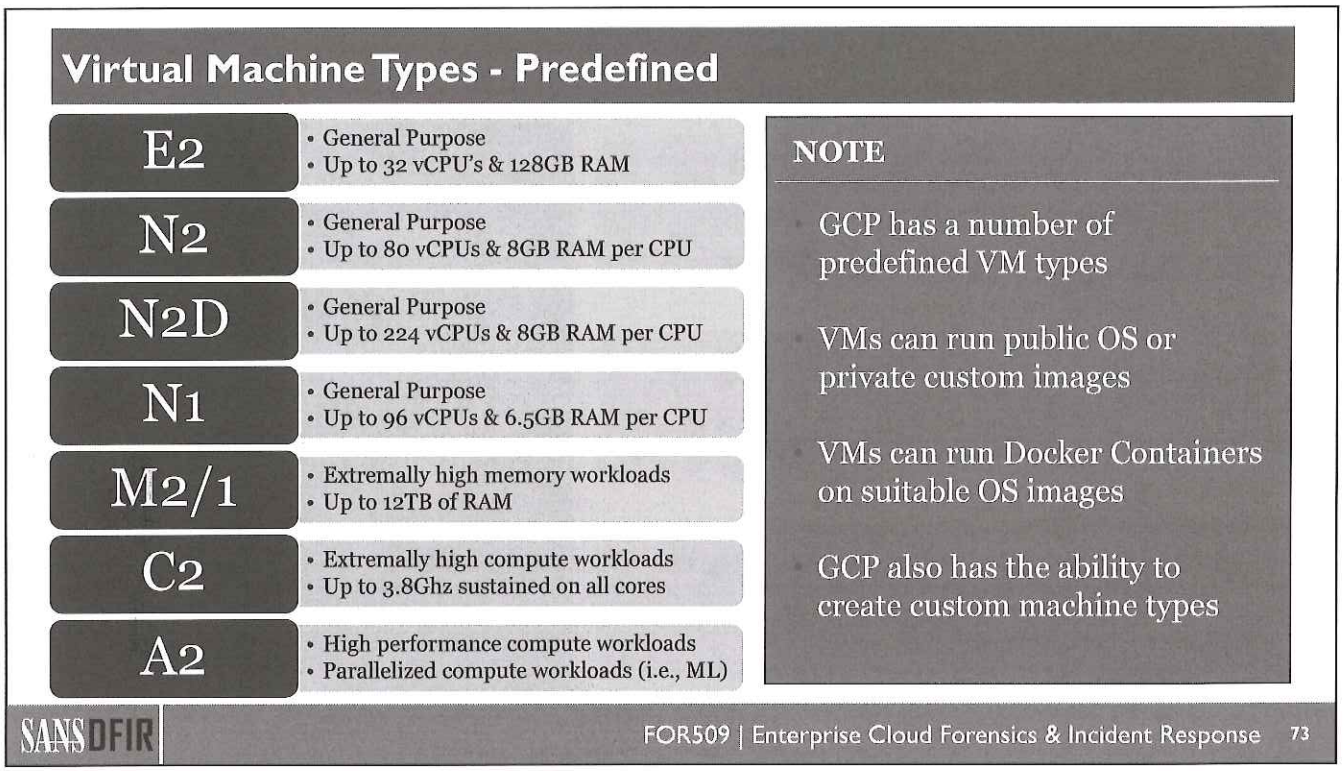

References:

[1] https://for509.com/r9pb8

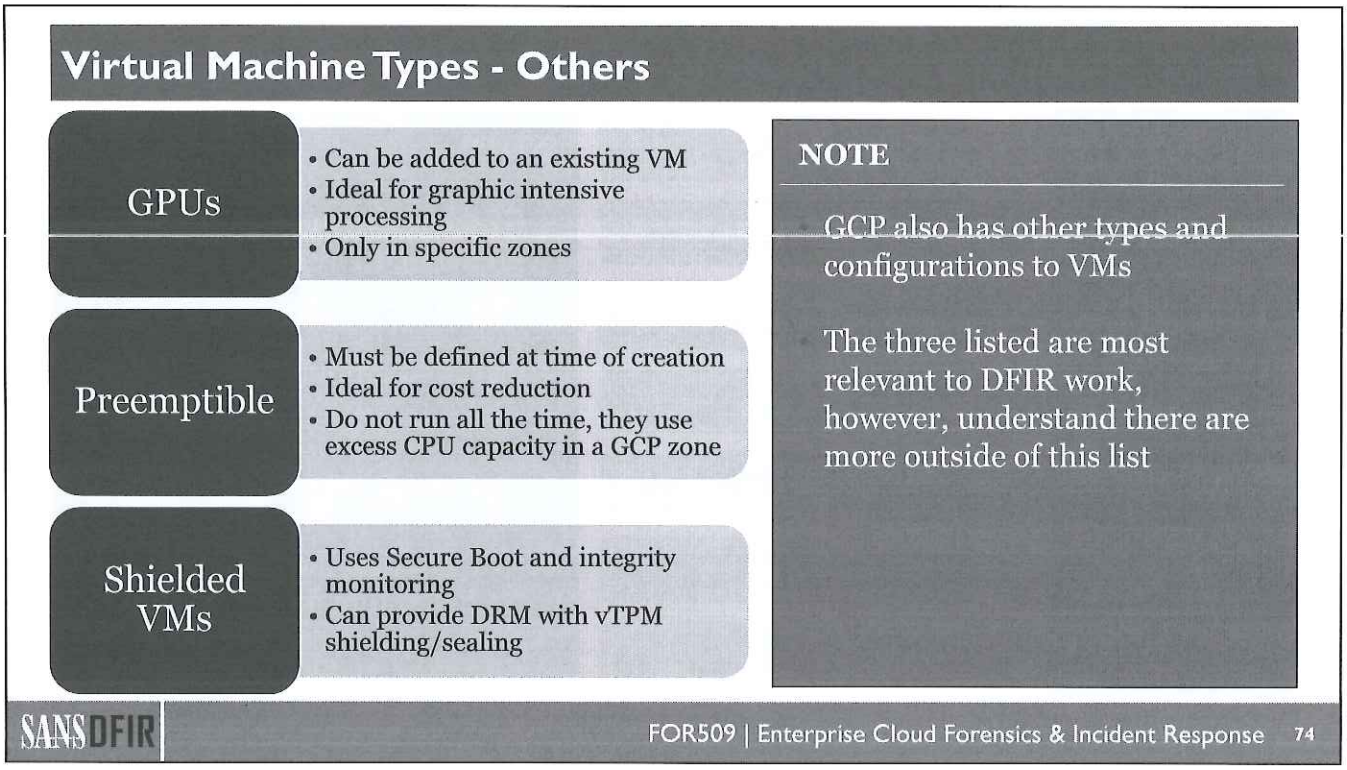

References:

[1] https://for509.com/w9prz

[2] https://for509.com/5xbh3

[3] https://for509.com/oarxb

Virtual Machine – Storage & Snapshots (1)

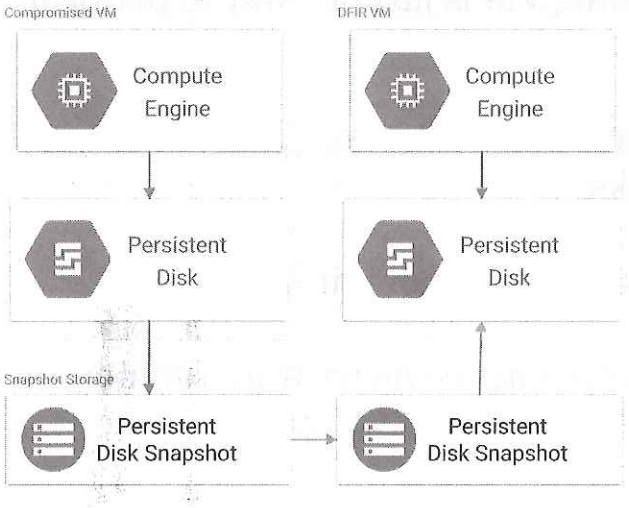

Capturing a VM Disk Image

- Requires the Compromised VM to have its Disk "Snapshot"
- The snapshot needs to be shared with the user/project/org that has your DFIR VM
- Once shared it can be copied
- The copy can be converted back to a Disk
- The Disk can be mounted as a "read only" drive to your DFIR analysis VM
- From here, normal Disk Forensics would occur (i.e., FOR508)

This page intentionally left blank.

Virtual Machine – Storage & Snapshots (2)

- The method of snapshotting a running VM is just one way to perform forensics on a running VM.
- This method has been available since a lot of the IaaS platforms first enabled the ability to snapshot disks.
- It also has one large draw back – it's a single point in time.
- GCP has built out the concept of collecting telemetry from running VMs and applications on those VMs, known as "**Cloud Logging Agent**"

This page intentionally left blank.

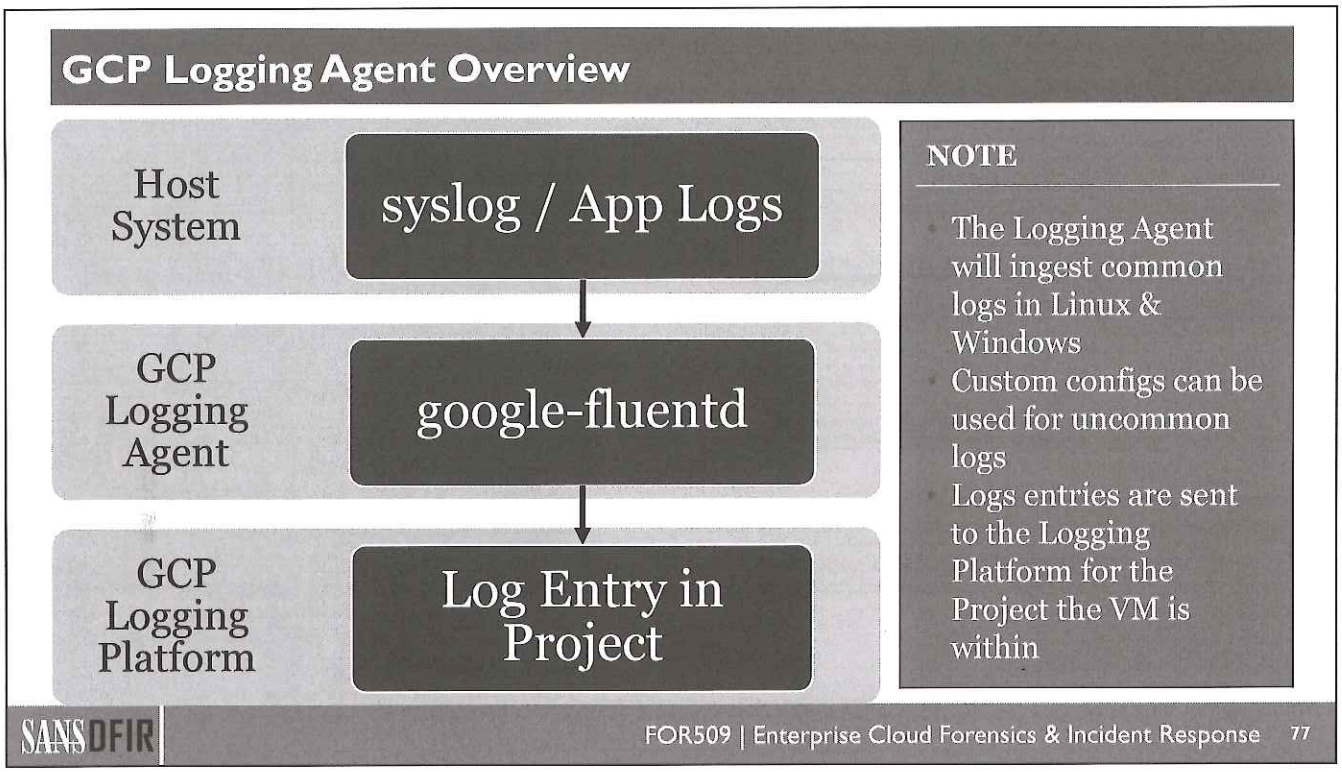

References:

[1] https://for509.com/st7qb

[2] https://for509.com/kvzui

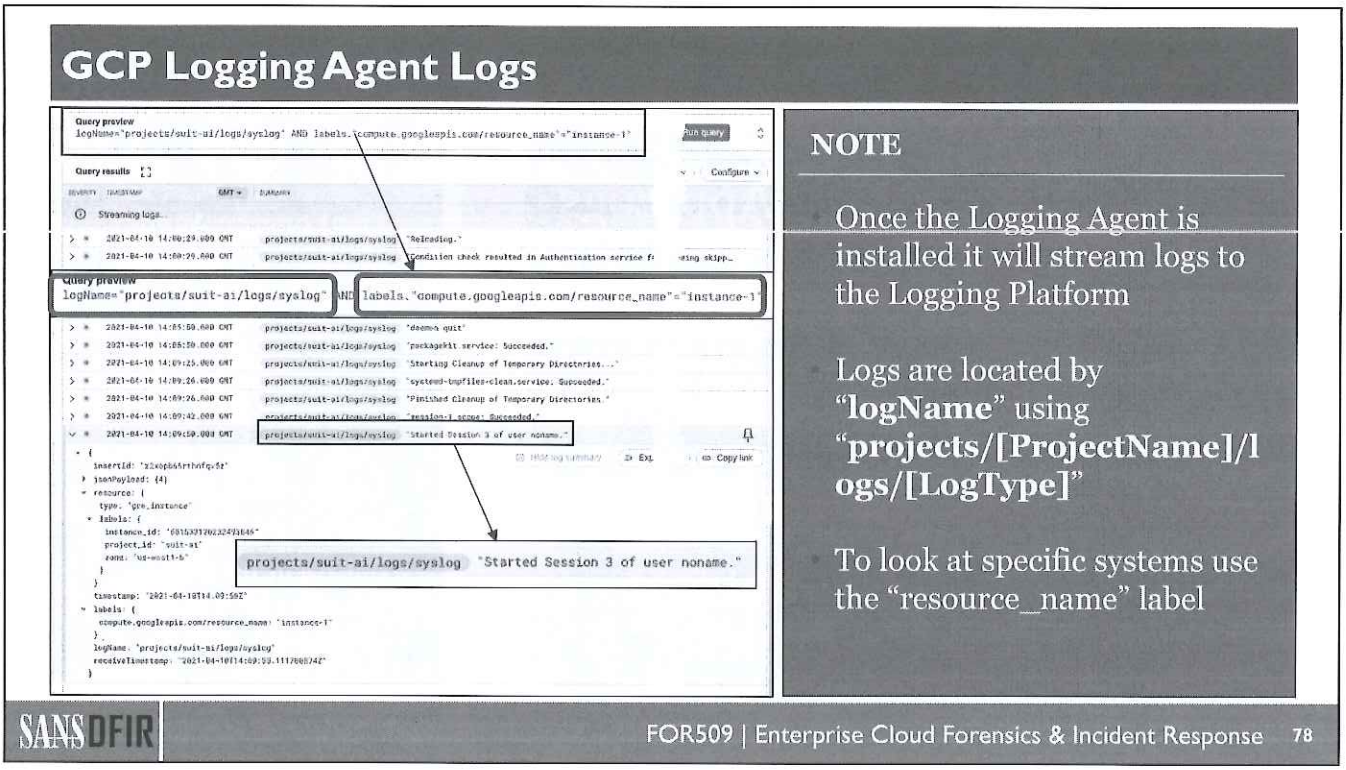

References:

[1] https://for509.com/m5x2c

[2] https://for509.com/v39uh

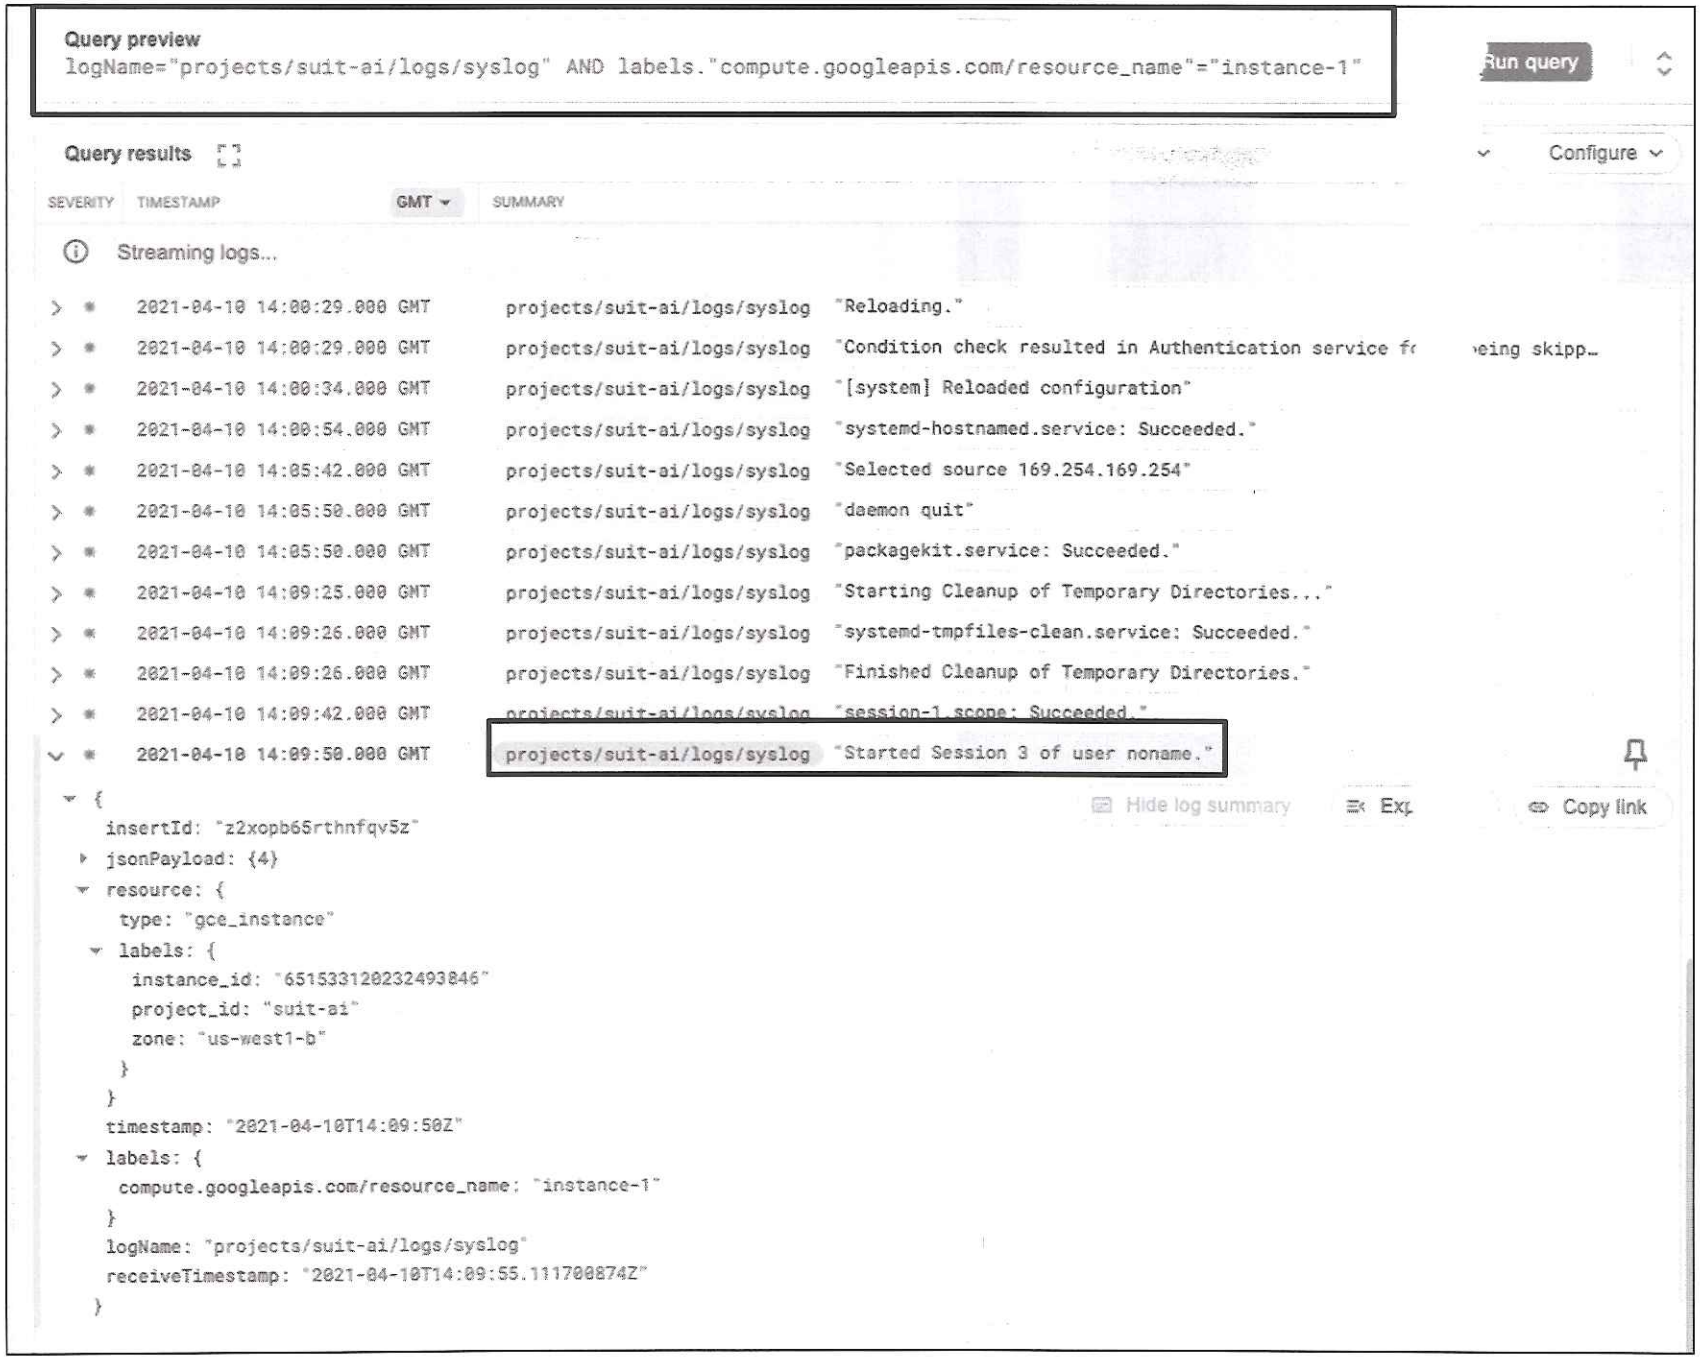

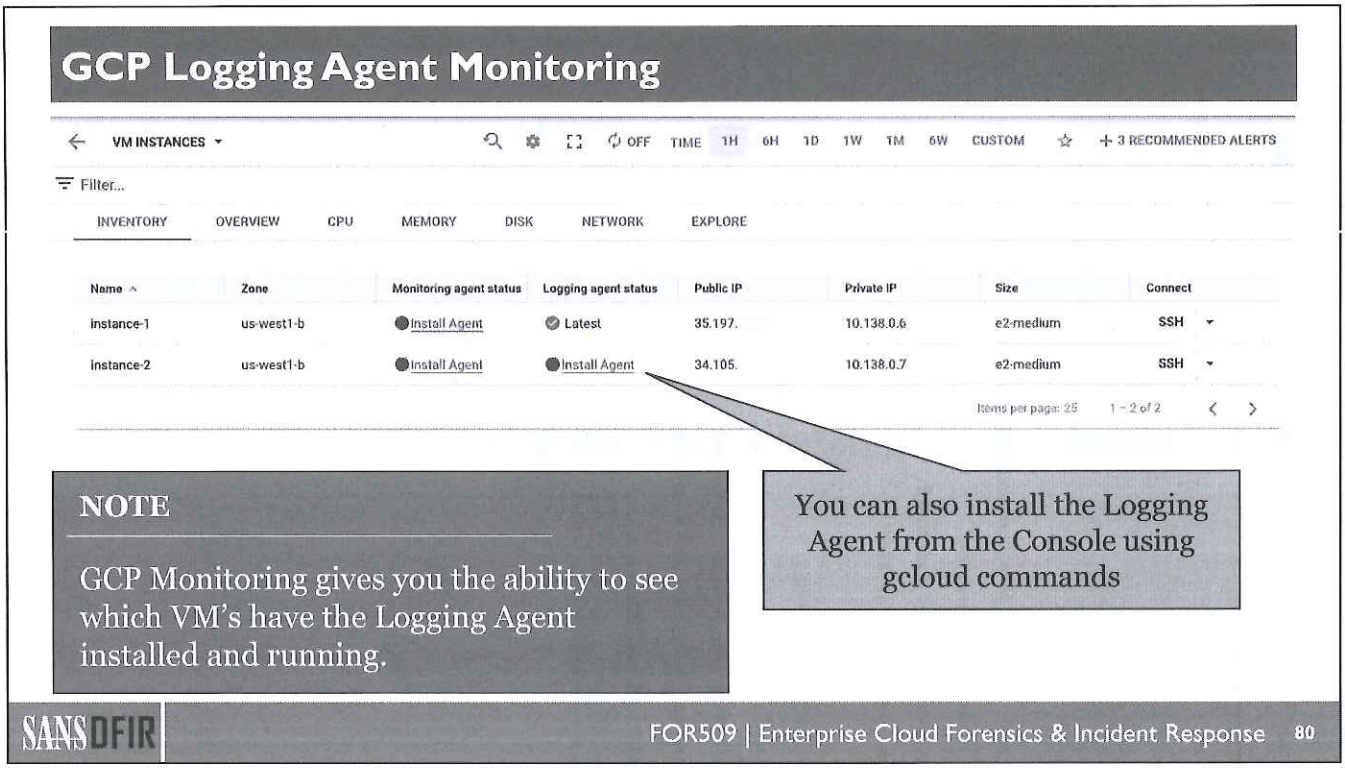

References:

[1] https://for509.com/3sjlx

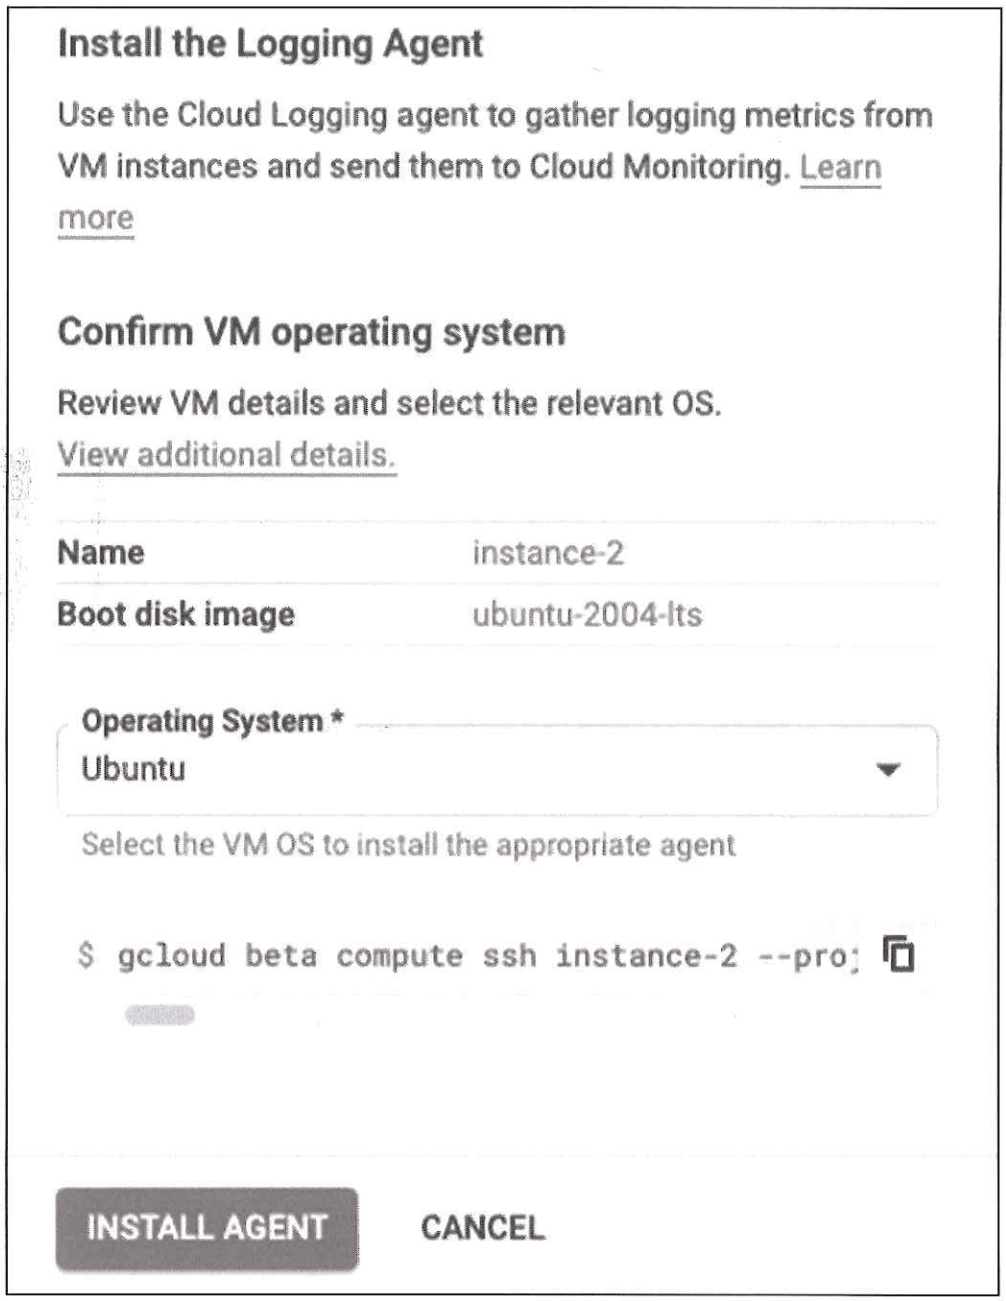

GCP Logging Agent – Structured and Unstructured (1)

- Logging Agent is based on "**fluentd**"
- **Fluentd** has a number of support parsers that **google-fluentd** can use to transform data to JSON when configured

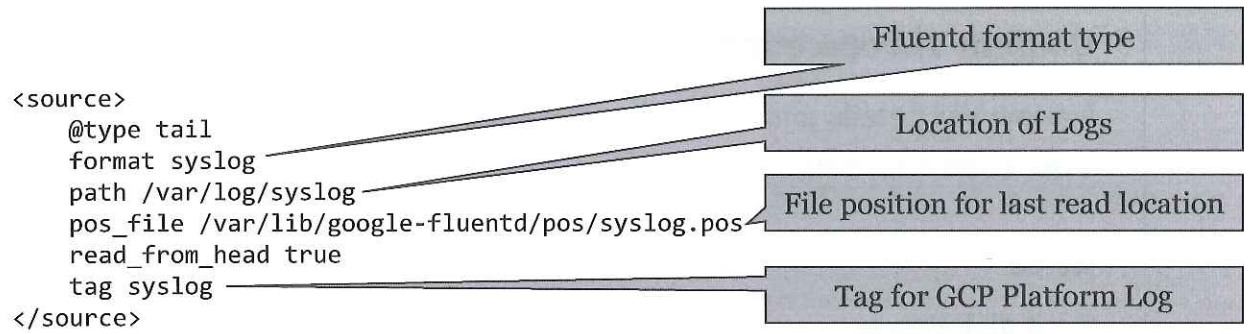


```
<source>
    @type tail
    format syslog
    path /var/log/syslog
    pos_file /var/lib/google-fluentd/pos/syslog.pos
    read_from_head true
    tag syslog
</source>
```


References:

[1] https://for509.com/0g215

[2] https://for509.com/5pftw

GCP Logging Agent – Structured and Unstructured (2)

- **Fluentd** existing parsers can transform syslog data into JSON

```
<1> Apr 10 14:09:50 instance-1 system: Started Session 3 of user noname.
```

```
jsonPayload: {
        pid: "1"
        message: "Started Session 3 of user noname."
        ident: "systemd"
        host: "instance-1"
}
timestamp: "2021-04-10T14:09:50Z"
```


References:

[1] https://for509.com/0g215

[2] https://for509.com/5pftw

GCP Logging Agent with AWS and Ops Agent

- **The GCP Logging Agent is also compatible with AWS Linux and Windows VM's**
  - **This requires an AWS Connector – usually a separate Project within your Organization**
- **There is also an "Ops Agent" in GCP**
  - **This is used for VMs with very high log throughput or resource utilization**
  - **This agent is also used for performance monitoring of VMs (i.e., CPU, Memory, Disk, etc.)**

References:

[1] https://for509.com/5k1zv

[2] https://for509.com/dwntc

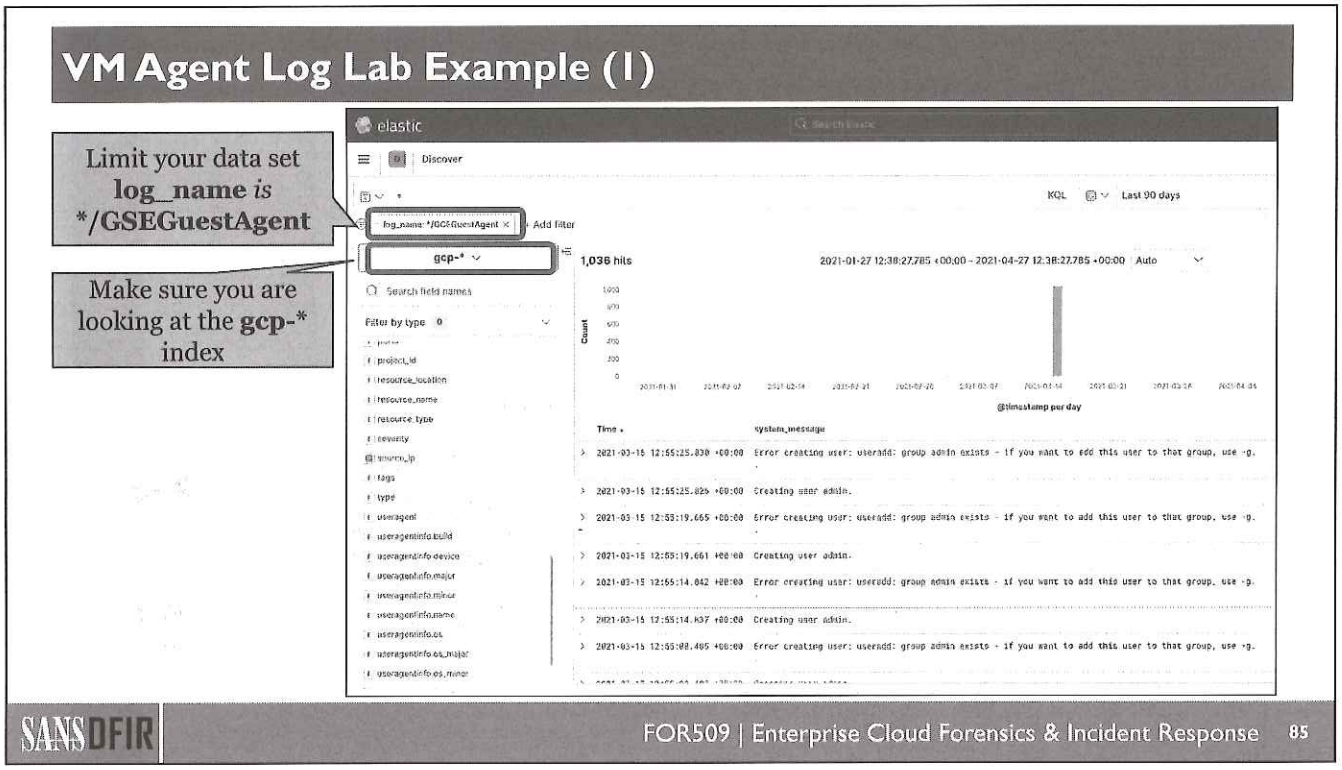

The example above shows you SOF-ELK with the VM Agent logs extracted from GCP in JSON format and loaded into SOF-ELK. In the Lab, that is coming up, your data set will already be limited to only the GCP VM Agent logs, however, if you have loaded other data into SOF-ELK you will need to limit your data set.

The data set under the "gcp-*" index will hold all the GCP data, except for Flow Logs. When you're viewing this index, you will see all GCP data loaded in SOF-ELK. Using the "Available fields" section locate the "log_name" field, once you click on this field it will show you the other services that are also loaded into SOF-ELK from previous GCP data you may have been working with. Select the one of the entries that end with "\GCEGuestAgent", once selected you will need to modify this search slightly, click on the filter you just created at the top and alter the Value to be "*/GCEGuestAgent". For this example in the course, we are going to expose all the VM Agent logs, however, in practice you may want to limit this to specific projects or folders. You can still use this technique to limit your logs to a specific project or folder.

Remember, you can always click on the filter you've just put in place and select "Temporarily disable" if you want to see all GCP data related to your search in the search bar at the top.

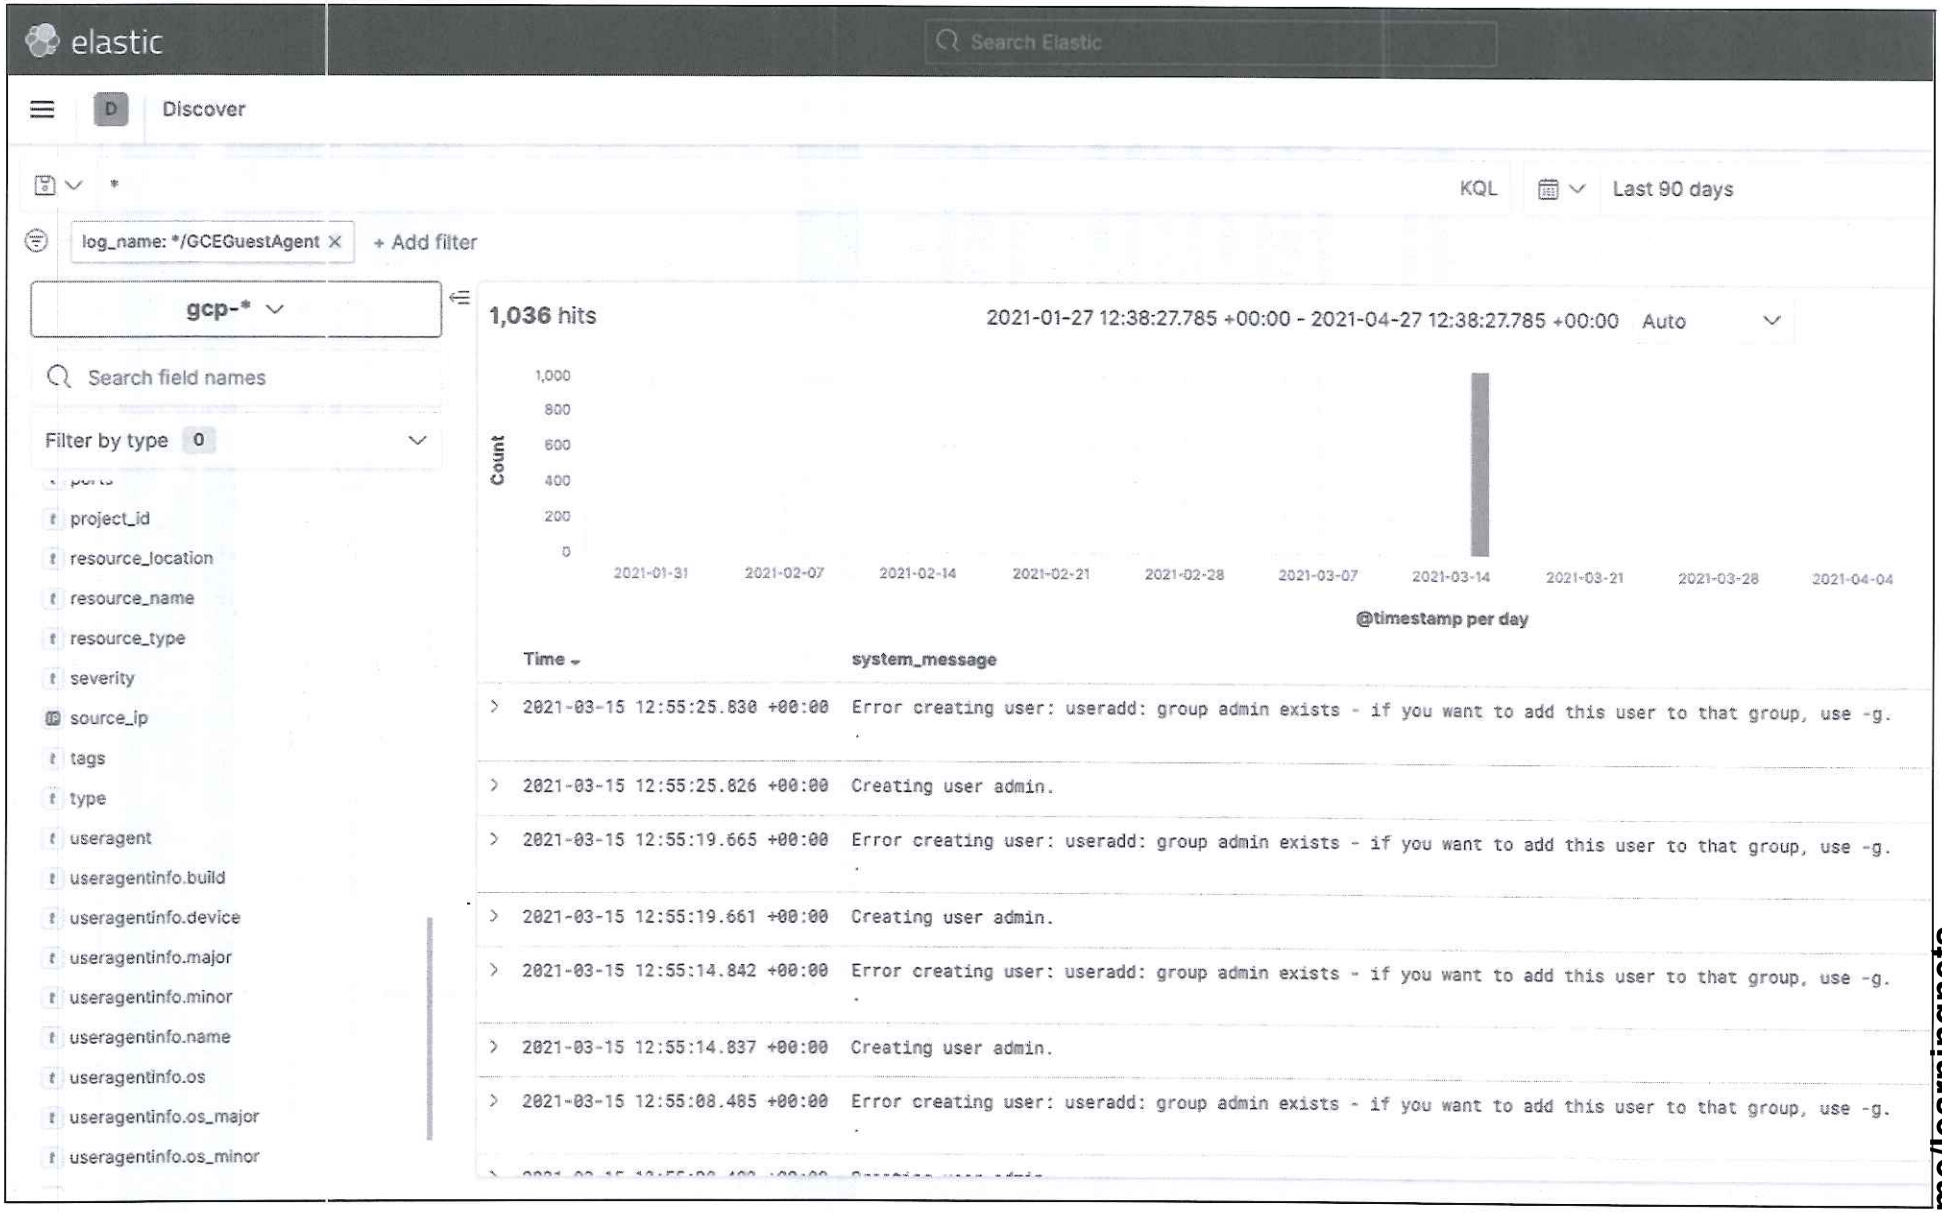

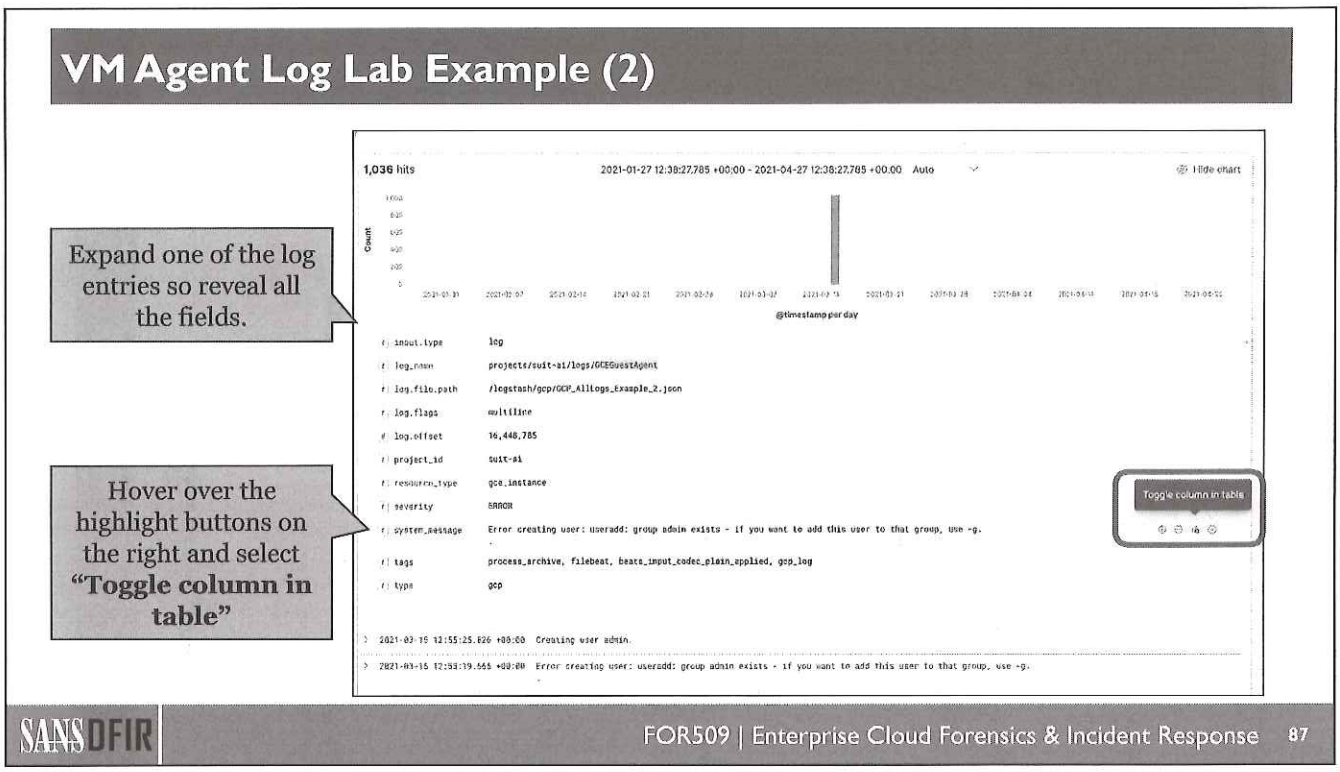

You may want to get a better view of the data that has been capture by the VM Agent. You can do this by adding in a field of data to the non-expended data view. If you want to see System Messages in you unexpended explorer view, you can expend one of the records and find the "system_message" field. If you hover of the right-hand side of the field, you will see an icon that represents "Toggle column in table", if you select this it will put this field in the unexpanded view of the record making it easier to see data in the Explorer view.

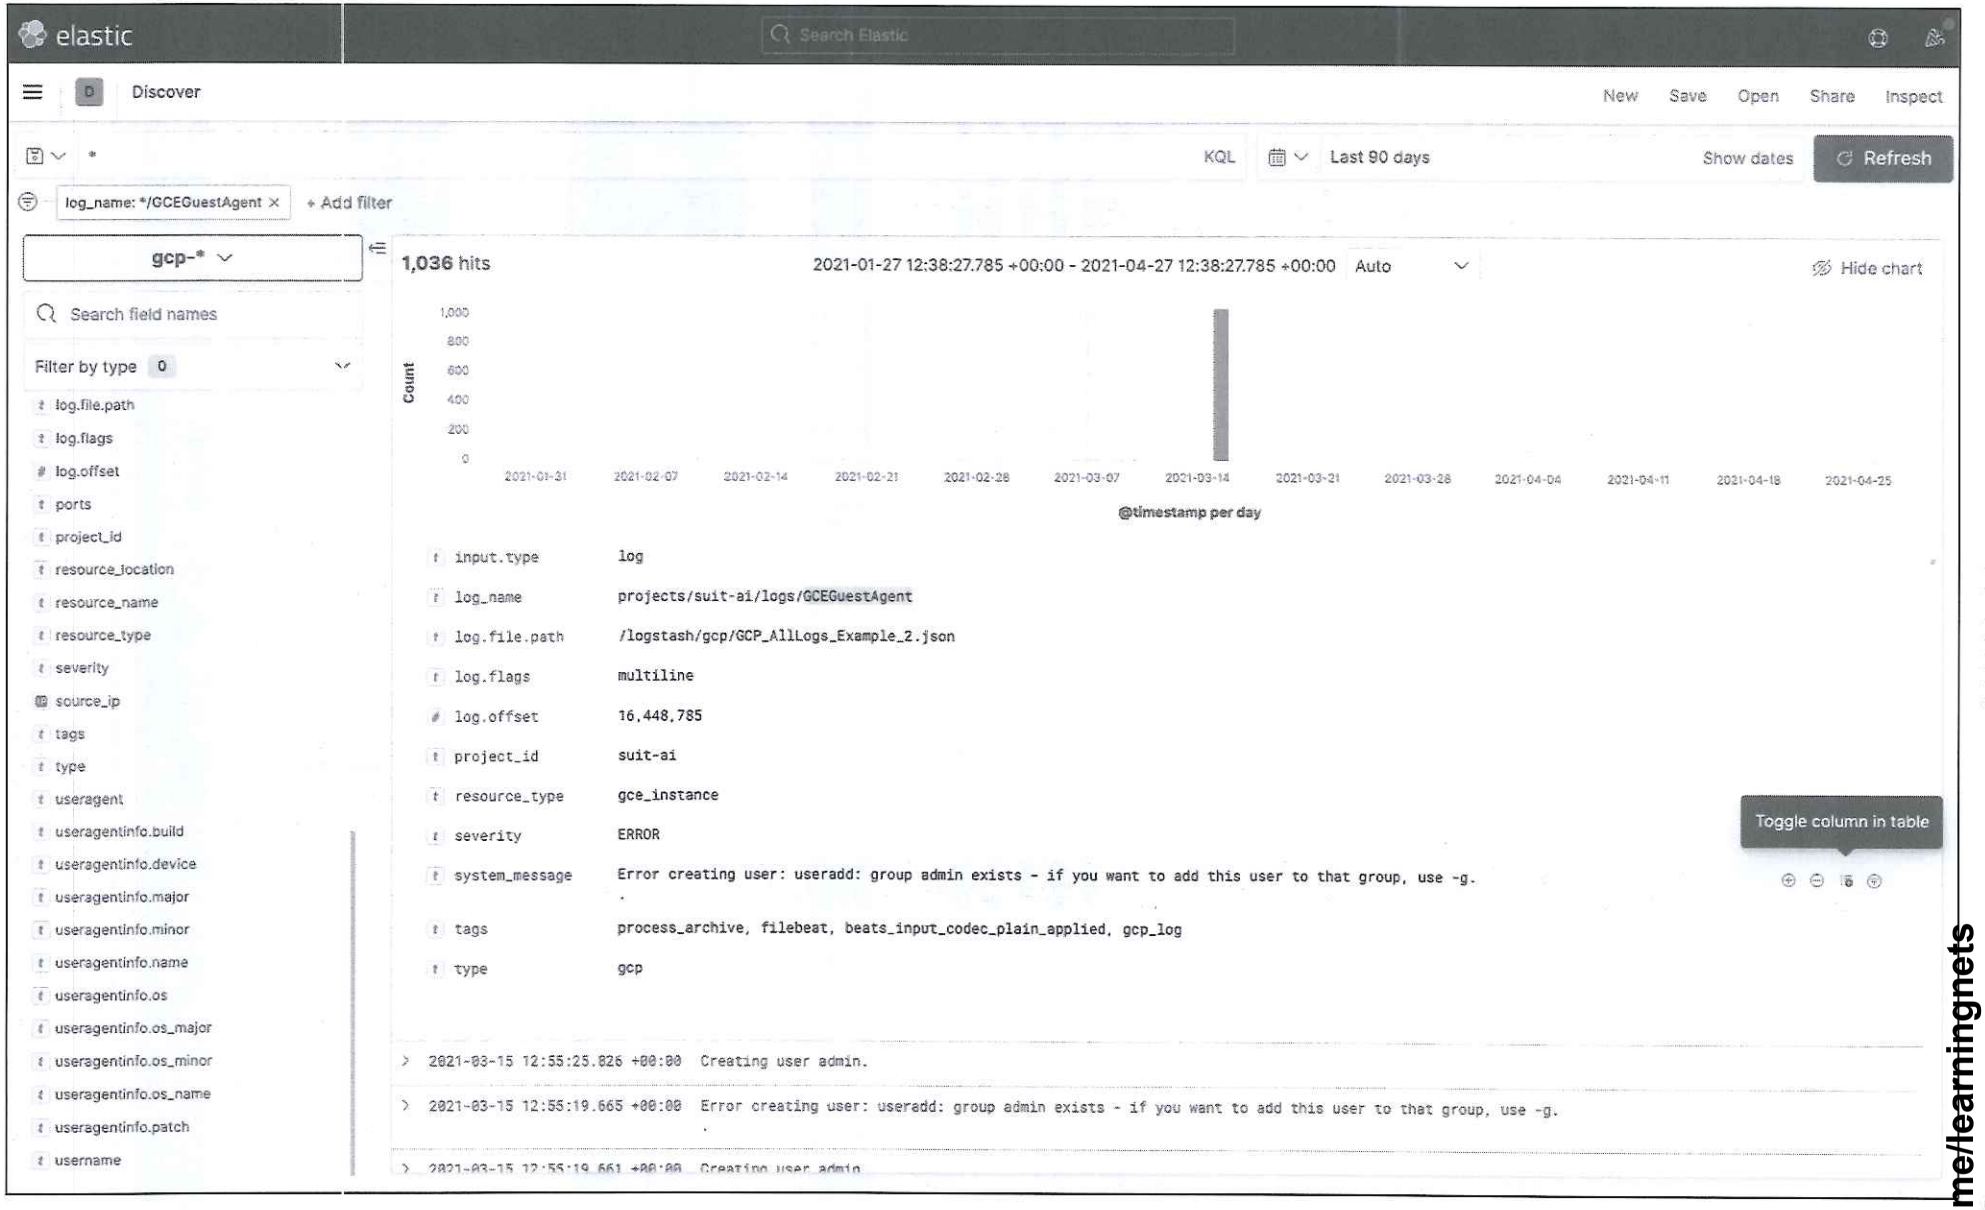

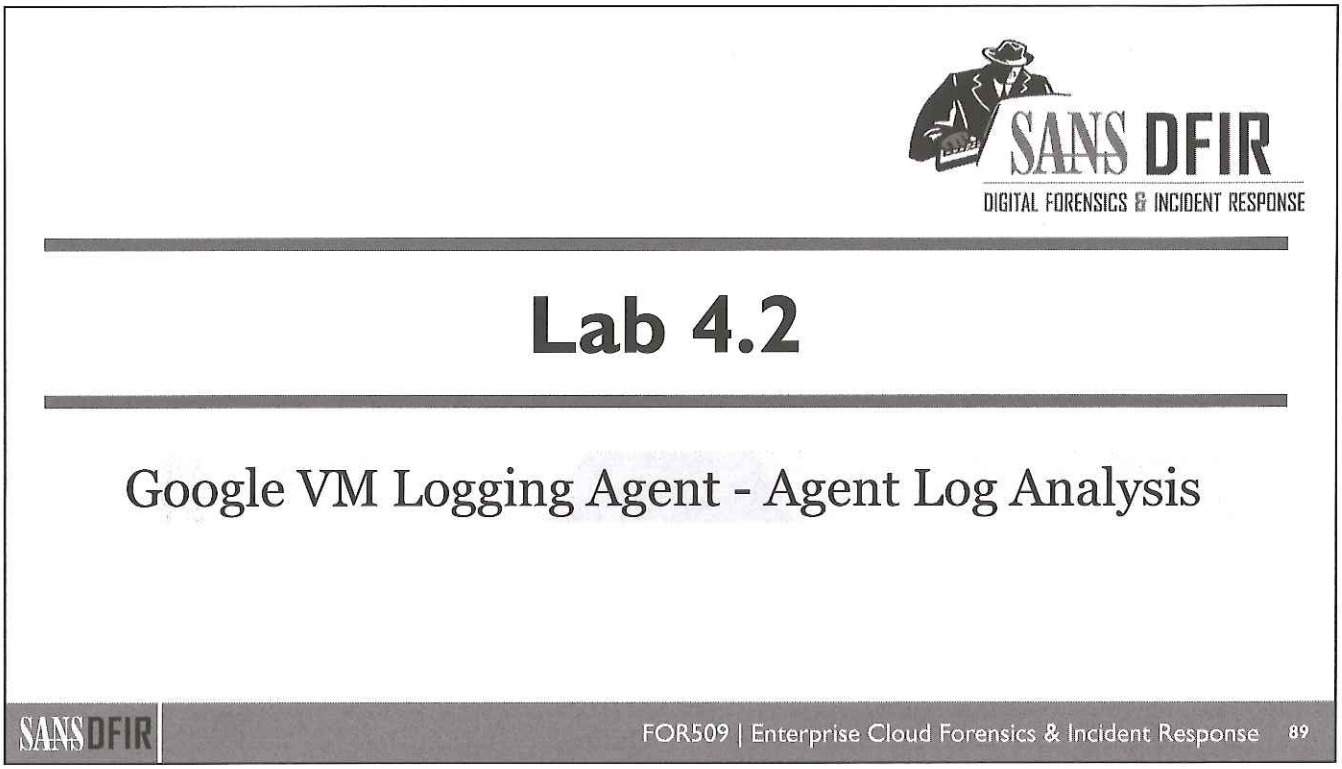

This page intentionally left blank.

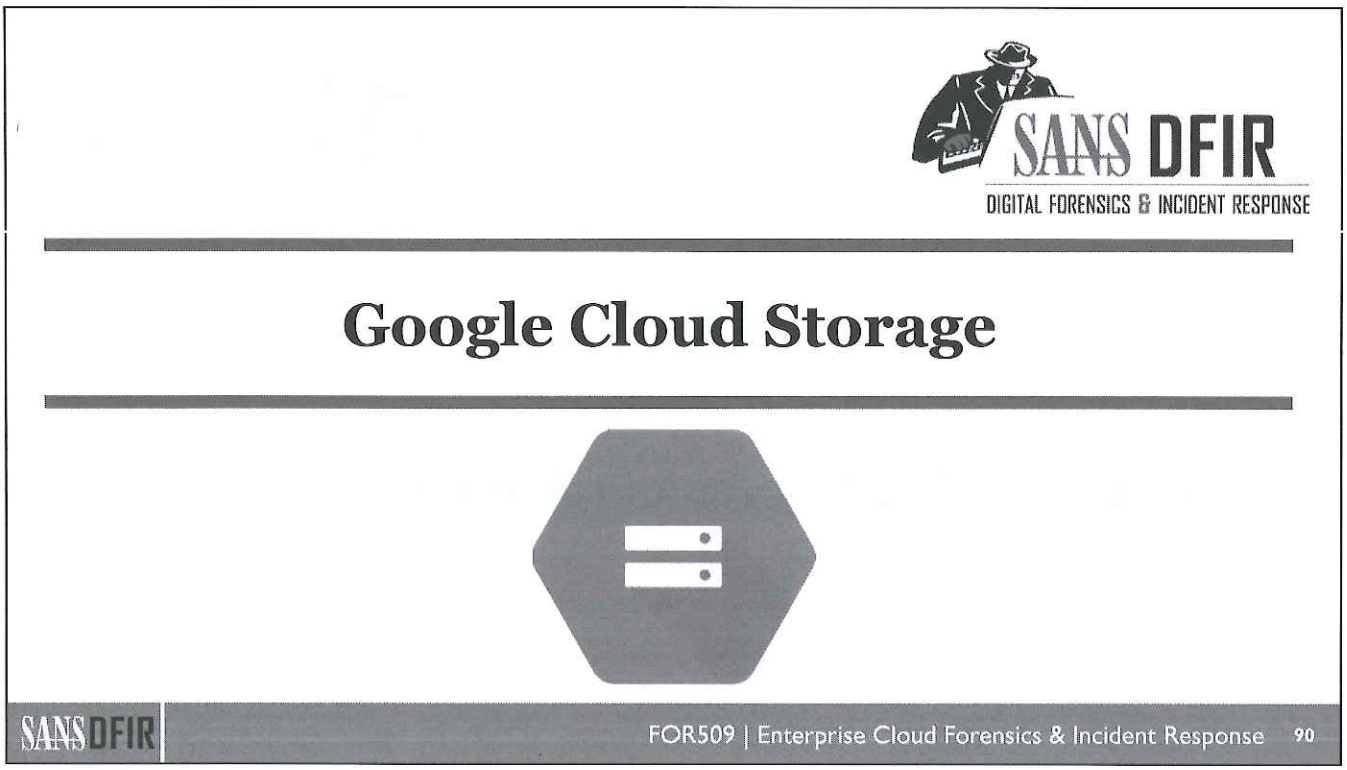

This page intentionally left blank.

GCP Storage Overview

There are several different Storage types in GCP along with features that can be used on the different storage types. For DFIR, we're going to focus no "Cloud Storage" (Buckets)

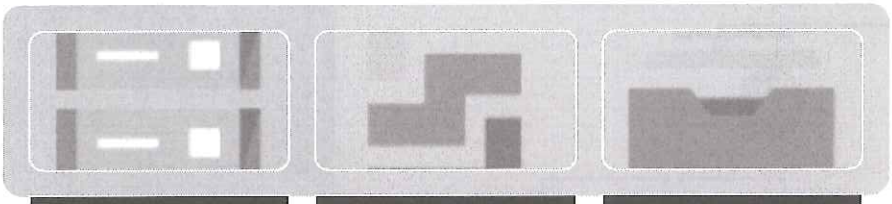

Cloud Storage	Block Storage (Persistent / SSD)	File Storage

This page intentionally left blank.

GCP Buckets Overview

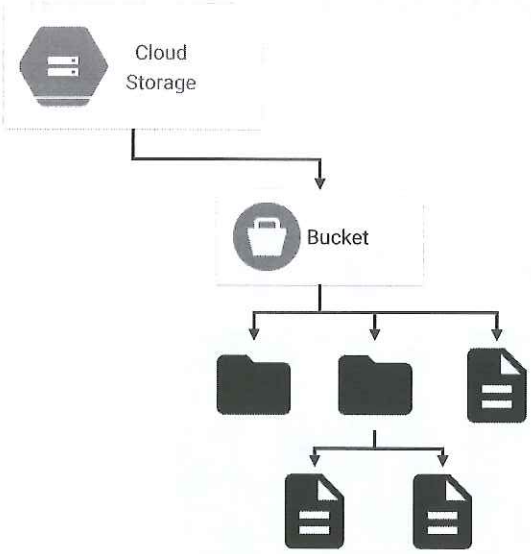

NOTE

- Cloud Storage allows you to create "Buckets" within "Projects" to store data.
- Within in a Bucket you can have folders, sub-folders and files.
- All the contents within the Bucket follow the permissions and polices of the parent Bucket.

This page intentionally left blank.

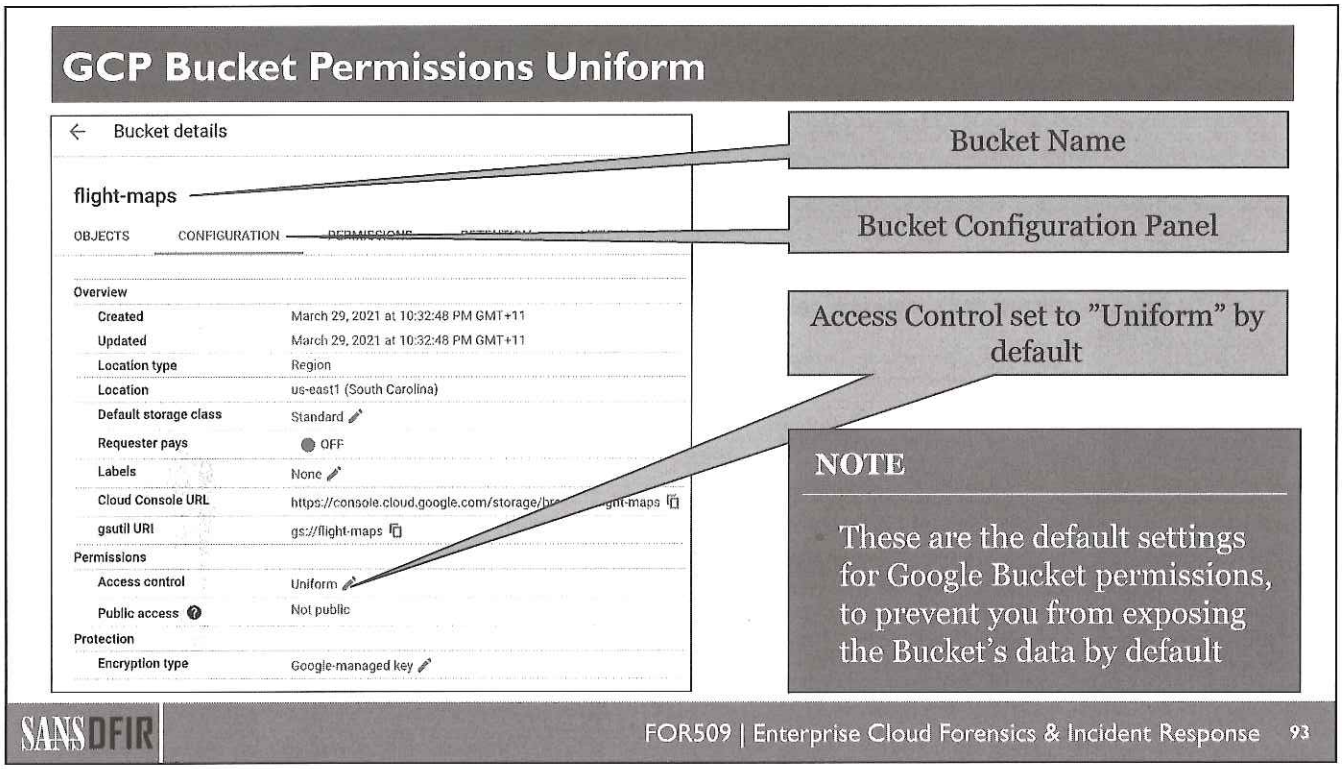

References:

[1] https://for509.com/phmxg

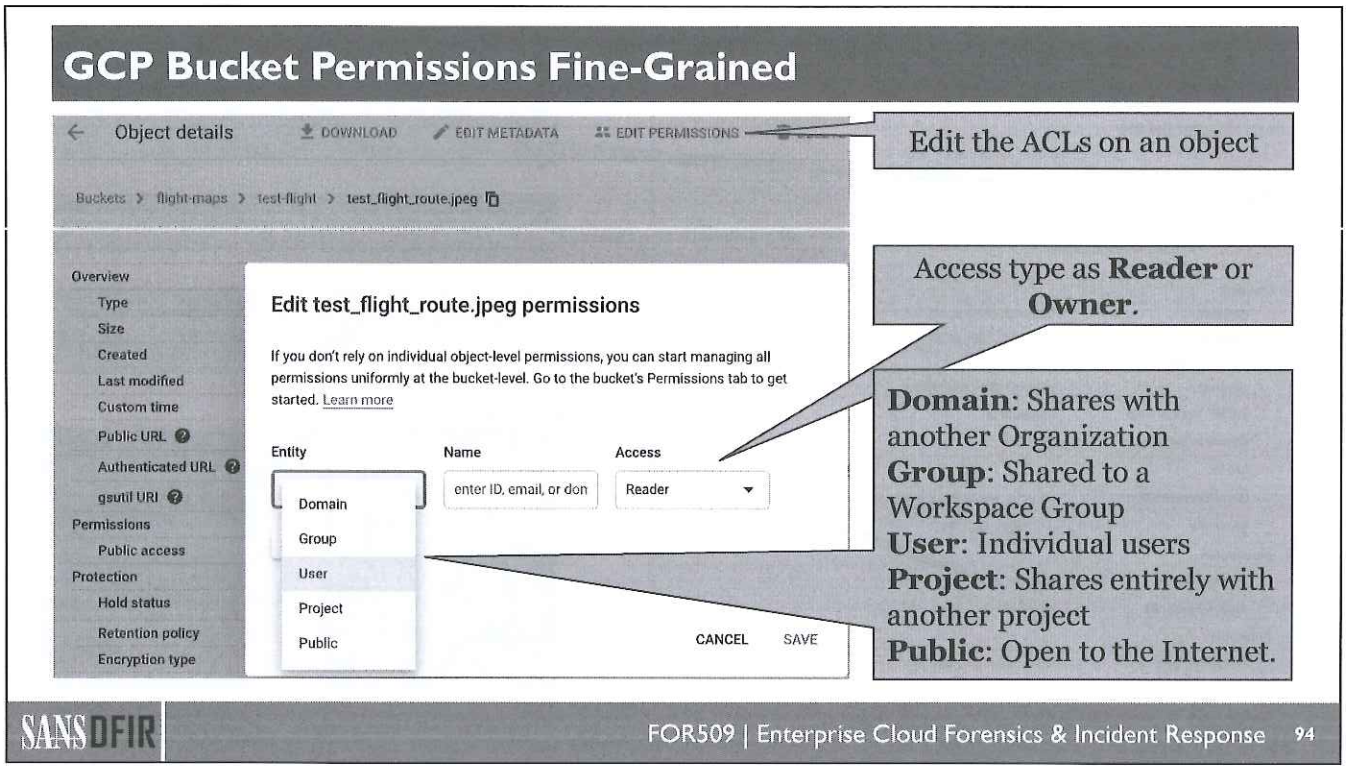

References:

[1] https://for509.com/48hrz

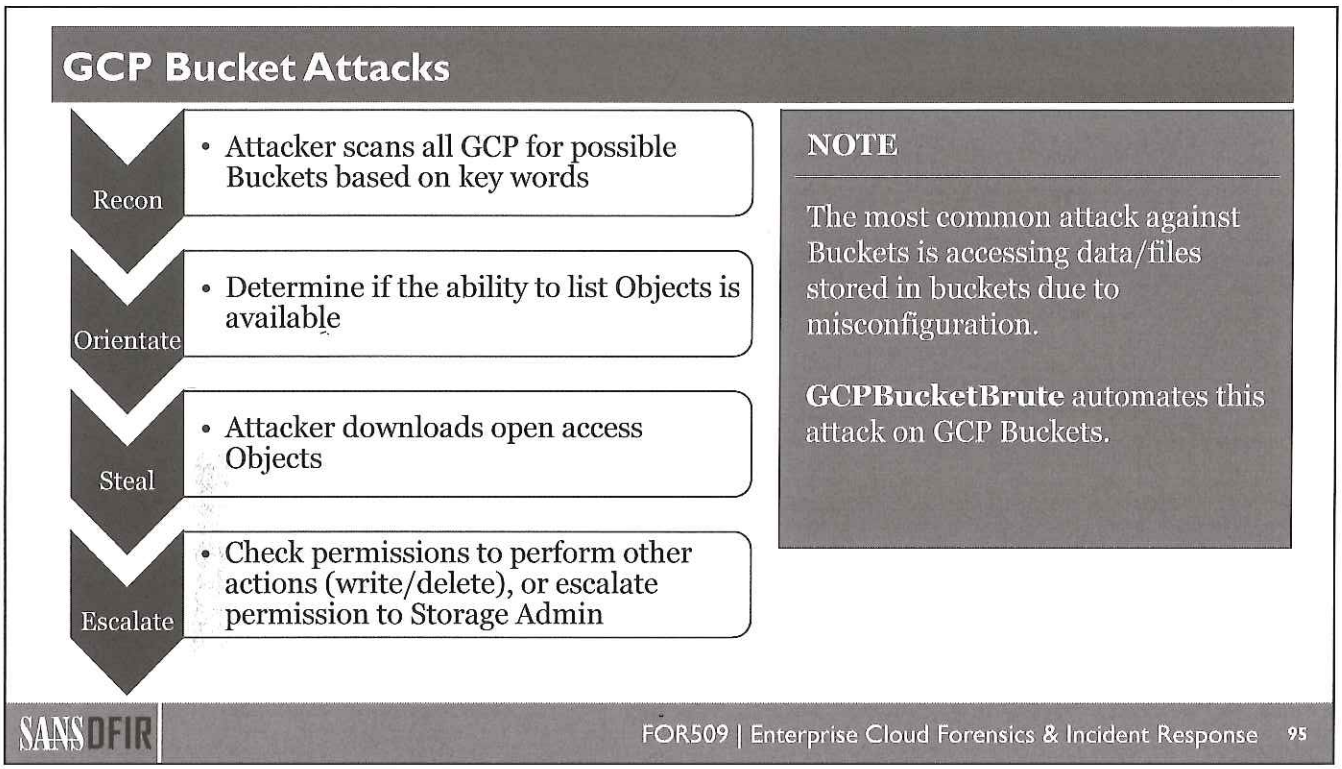

Resources:

[1] https://for509.com/afj2q

[2] https://for509.com/95h84

Cloud Storage – Enabling Object Logging

WARNING: Object logging is not turned on by default
This mean no object access logging to detect DLP

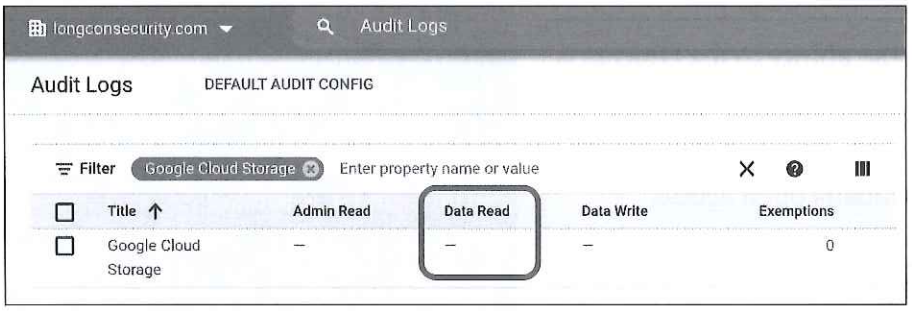

NOTE

You need to enable "**Data Read**" if you want to see access to Objects in Buckets.

Only do this on sensitive data – it will produce a *LOT* of logs.

This page intentionally left blank.

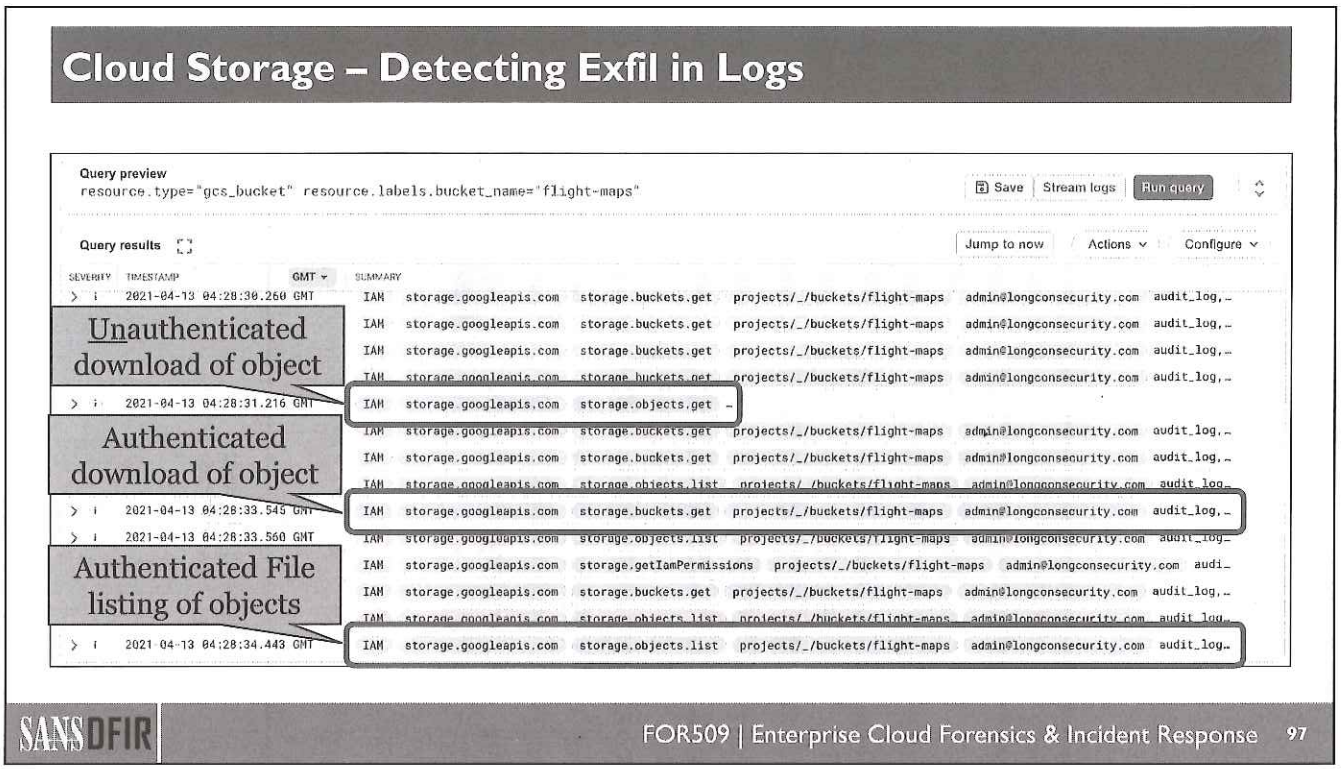

This page intentionally left blank.

Query preview

```
resource.type="gcs_bucket" resource.labels.bucket_name="flight-maps"
```

Save | Stream logs | Run query

Query results

Jump to now | Actions | Configure

SEVERITY	TIMESTAMP (GMT)	SUMMARY
i	2021-04-13 04:28:30.260 GMT	IAM storage.googleapis.com storage.buckets.get projects/_/buckets/flight-maps admin@longconsecurity.com audit_log,…
i	2021-04-13 04:28:30.291 GMT	IAM storage.googleapis.com storage.buckets.get projects/_/buckets/flight-maps admin@longconsecurity.com audit_log,…
i	2021-04-13 04:28:30.295 GMT	IAM storage.googleapis.com storage.buckets.get projects/_/buckets/flight-maps admin@longconsecurity.com audit_log,…
i	2021-04-13 04:28:31.371 GMT	IAM storage.googleapis.com storage.buckets.get projects/_/buckets/flight-maps admin@longconsecurity.com audit_log,…
i	2021-04-13 04:28:31.216 GMT	IAM storage.googleapis.com storage.objects.get …
i	2021-04-13 04:28:31.268 GMT	IAM storage.googleapis.com storage.buckets.get projects/_/buckets/flight-maps admin@longconsecurity.com audit_log,…
i	2021-04-13 04:28:33.400 GMT	IAM storage.googleapis.com storage.buckets.get projects/_/buckets/flight-maps admin@longconsecurity.com audit_log,…
i	2021-04-13 04:28:33.541 GMT	IAM storage.googleapis.com storage.objects.list projects/_/buckets/flight-maps admin@longconsecurity.com audit_log…
i	2021-04-13 04:28:33.545 GMT	IAM storage.googleapis.com storage.buckets.get projects/_/buckets/flight-maps admin@longconsecurity.com audit_log,…
i	2021-04-13 04:28:33.560 GMT	IAM storage.googleapis.com storage.objects.list projects/_/buckets/flight-maps admin@longconsecurity.com audit_log…
i	2021-04-13 04:28:33.600 GMT	IAM storage.googleapis.com storage.getIamPermissions projects/_/buckets/flight-maps admin@longconsecurity.com audi…
i	2021-04-13 04:28:33.618 GMT	IAM storage.googleapis.com storage.buckets.get projects/_/buckets/flight-maps admin@longconsecurity.com audit_log,…
i	2021-04-13 04:28:34.423 GMT	IAM storage.googleapis.com storage.objects.list projects/_/buckets/flight-maps admin@longconsecurity.com audit_log…
i	2021-04-13 04:28:34.443 GMT	IAM storage.googleapis.com storage.objects.list projects/_/buckets/flight-maps admin@longconsecurity.com audit_log…

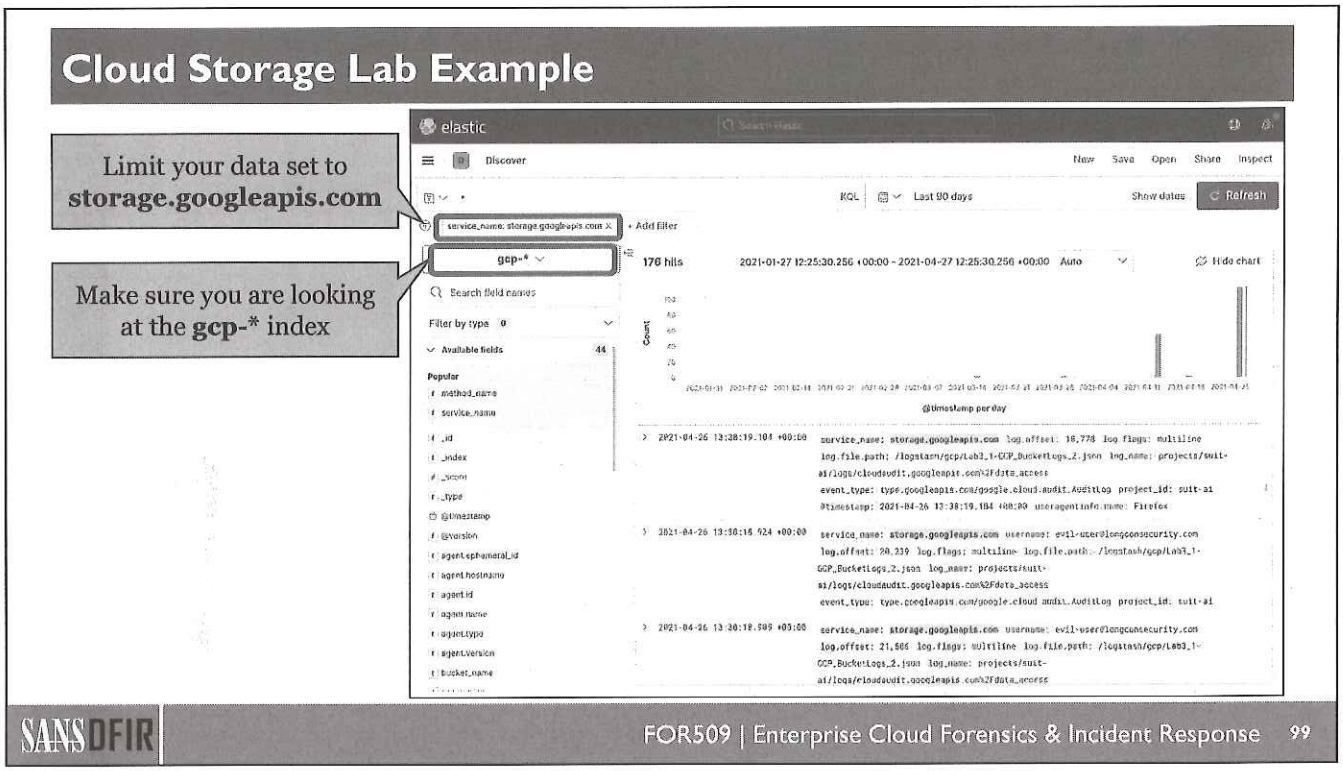

The example above shows you SOF-ELK with the Bucket Storage logs extracted from GCP in JSON format and loaded into SOF-ELK. In the Lab this is coming up your data set will already be limited to only the Bucket Storage logs, however, if you have loaded other data into SOF-ELK you will need to limit your data set.

The data set under the "gcp-*" index will hold all of the GCP data, except for Flow Logs. So, when you're viewing this index you will see all GCP data loaded in SOF-ELK. Using the "Available fields" section locate the "service_name" field, once you click on this field it will show you the other services that are also loaded into SOF-ELK from previous GCP data you may have been working with. Select the "storage.googleapis.com" "service_name" to limit your data to only logs related to GCP Storage data. This will aid in ensuring searches you run in the search bar are limited to GCP Storage data. Can you always click on the limit you've just put in place and select "Temporarily disable" if you want to see all GCP data related to your search in the search bar at the top.

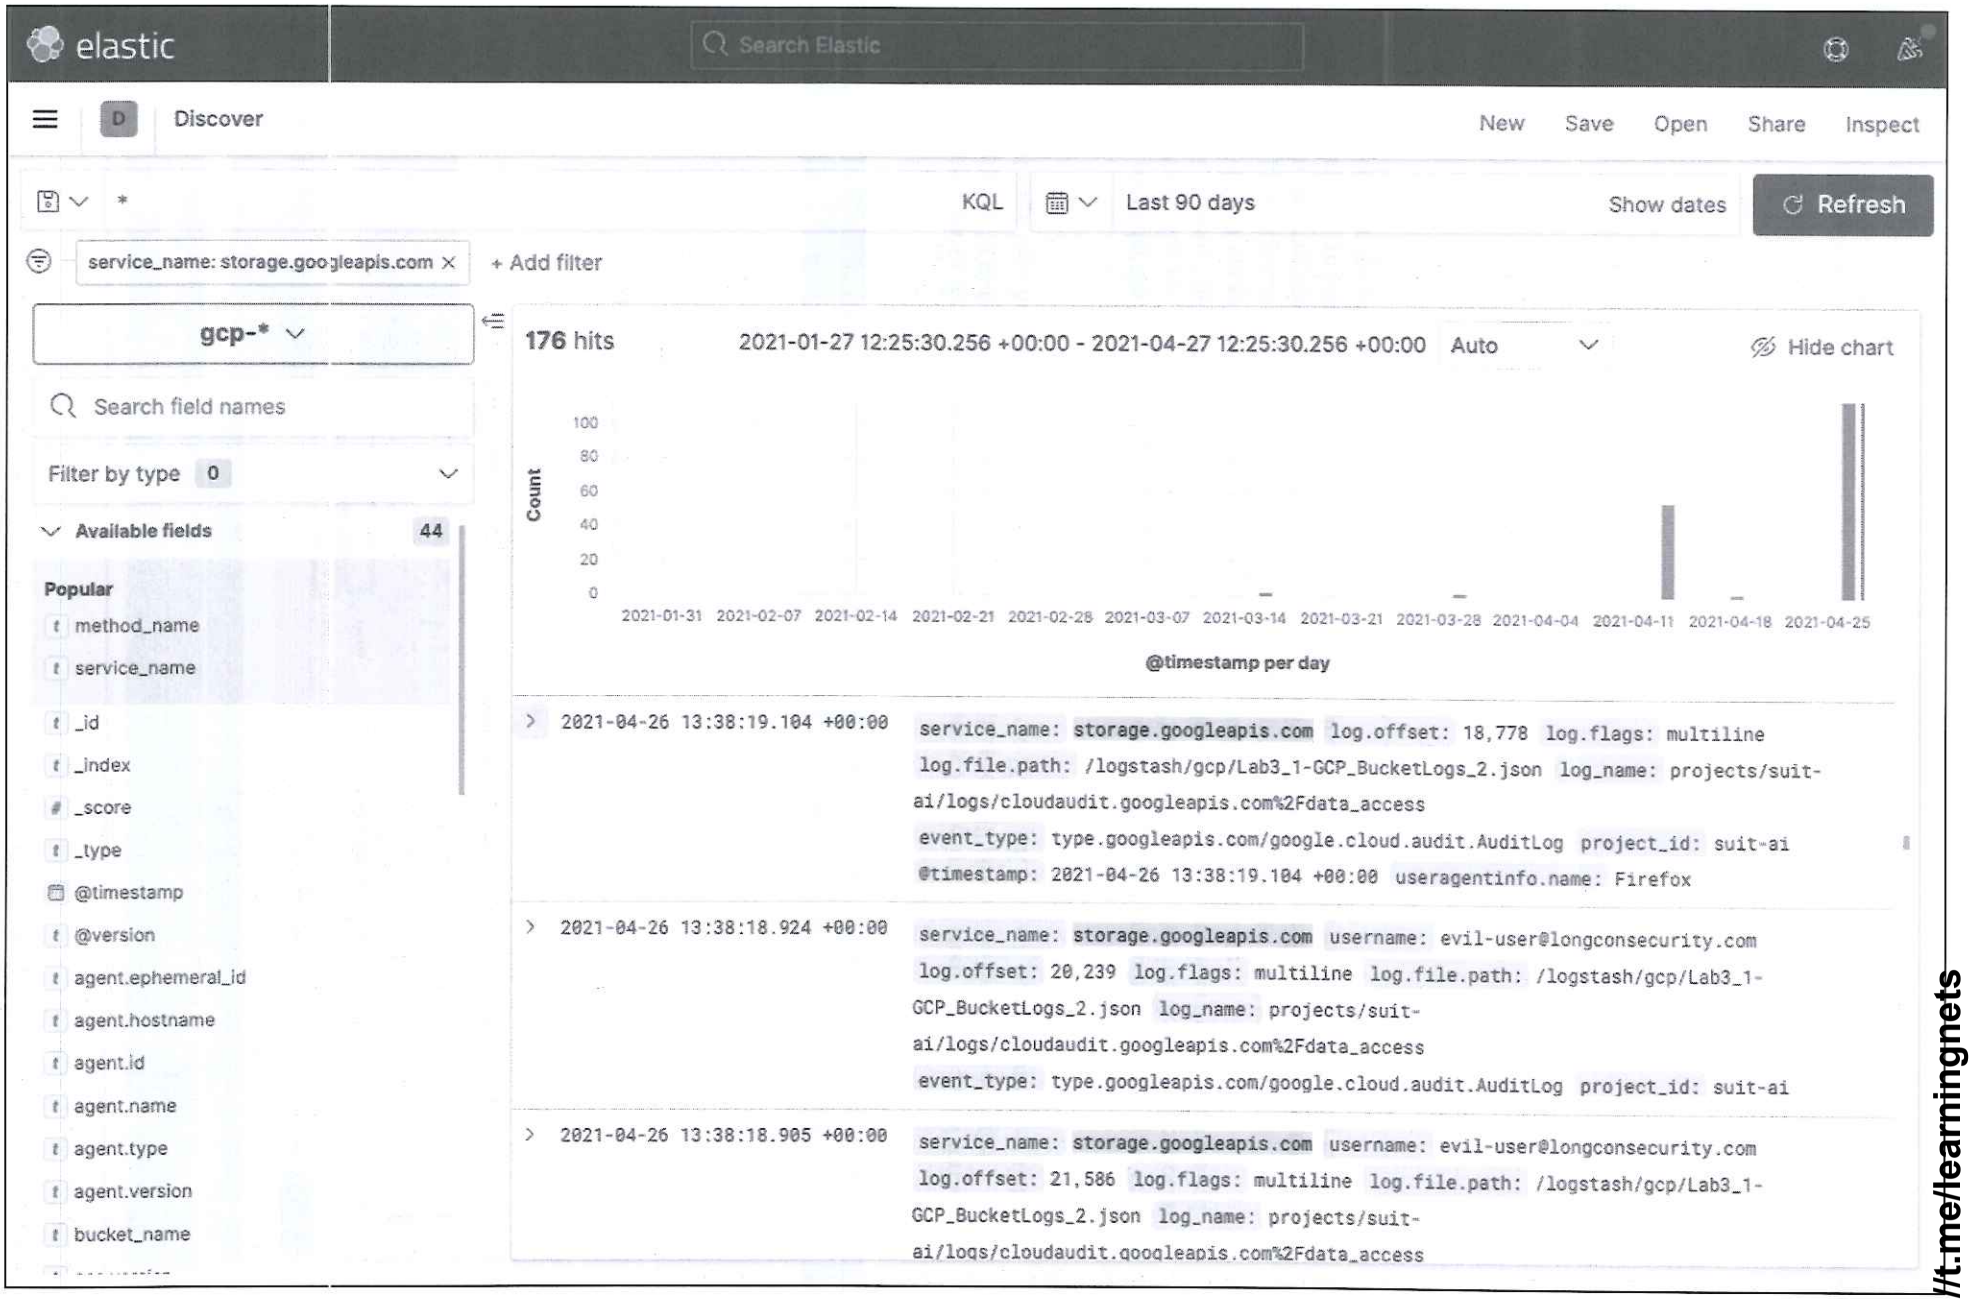

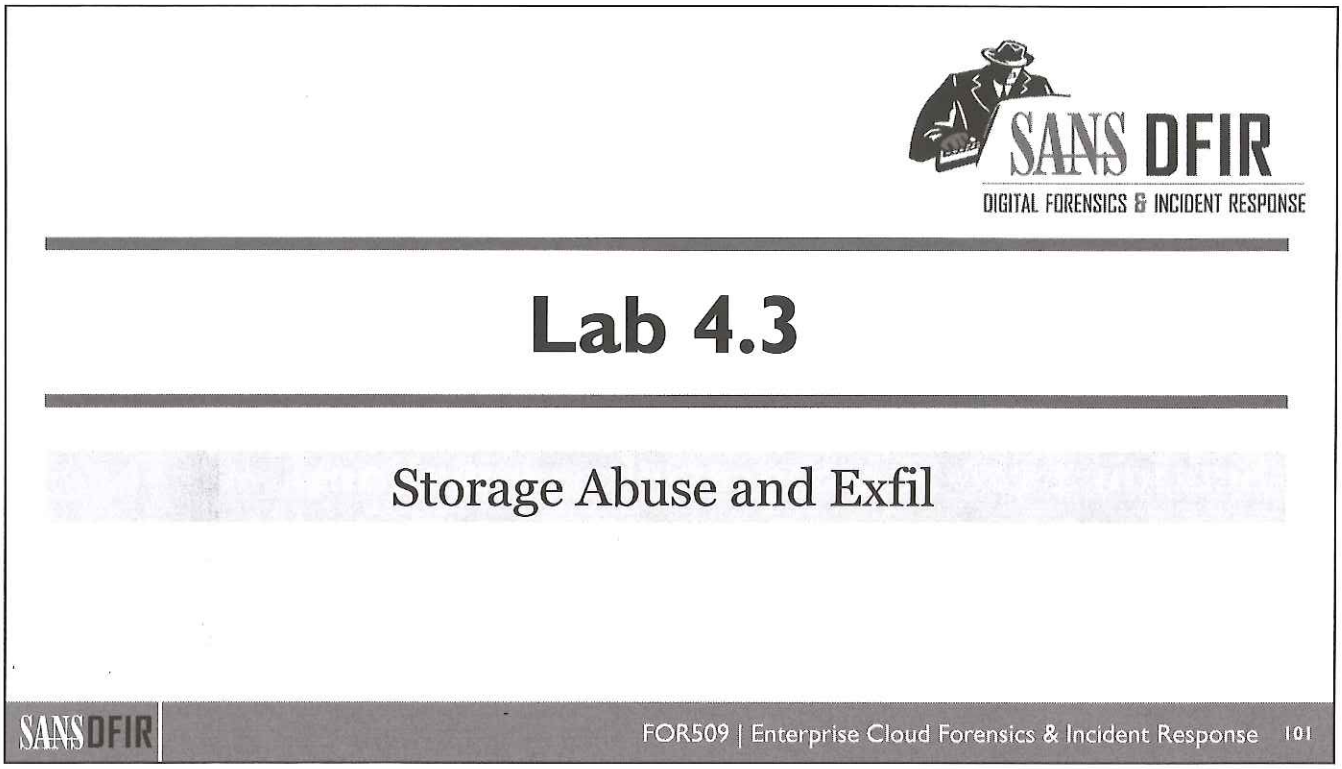

This page intentionally left blank.

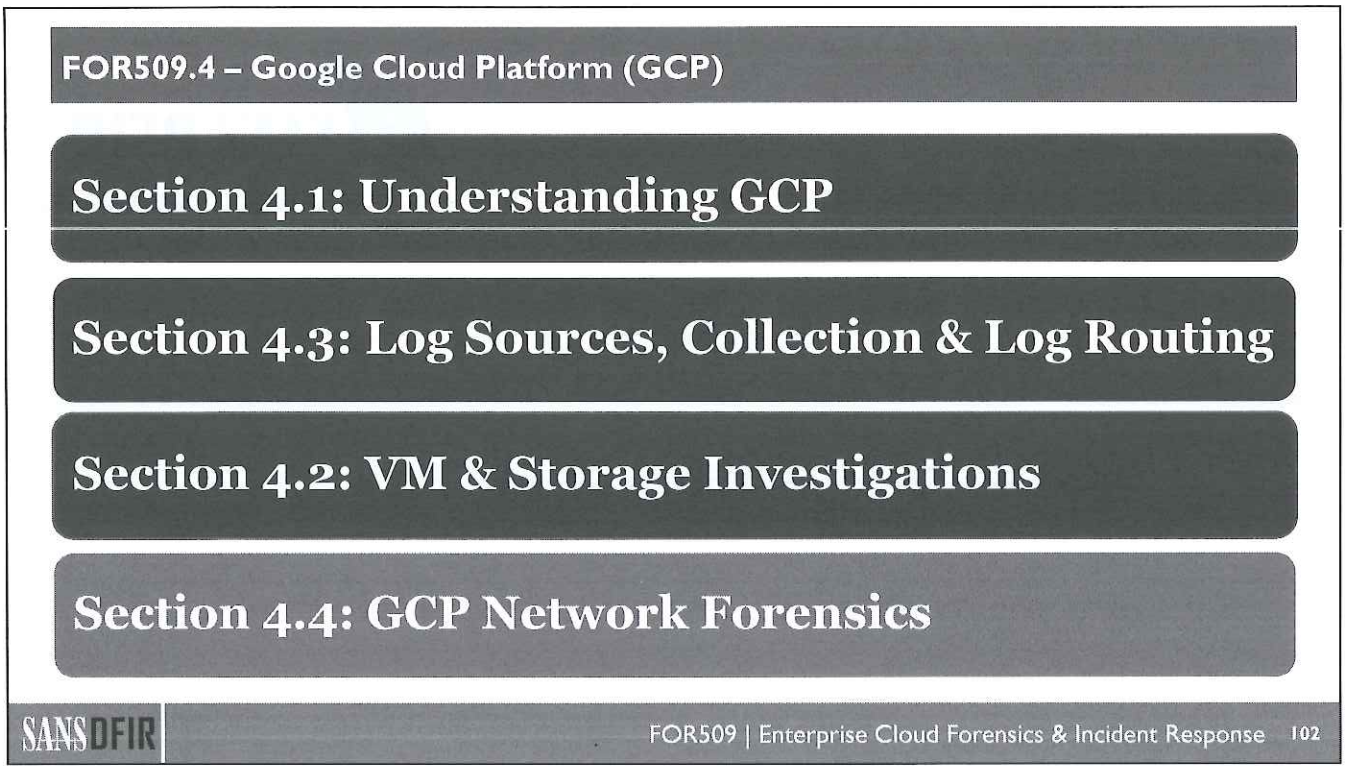

This page intentionally left blank.

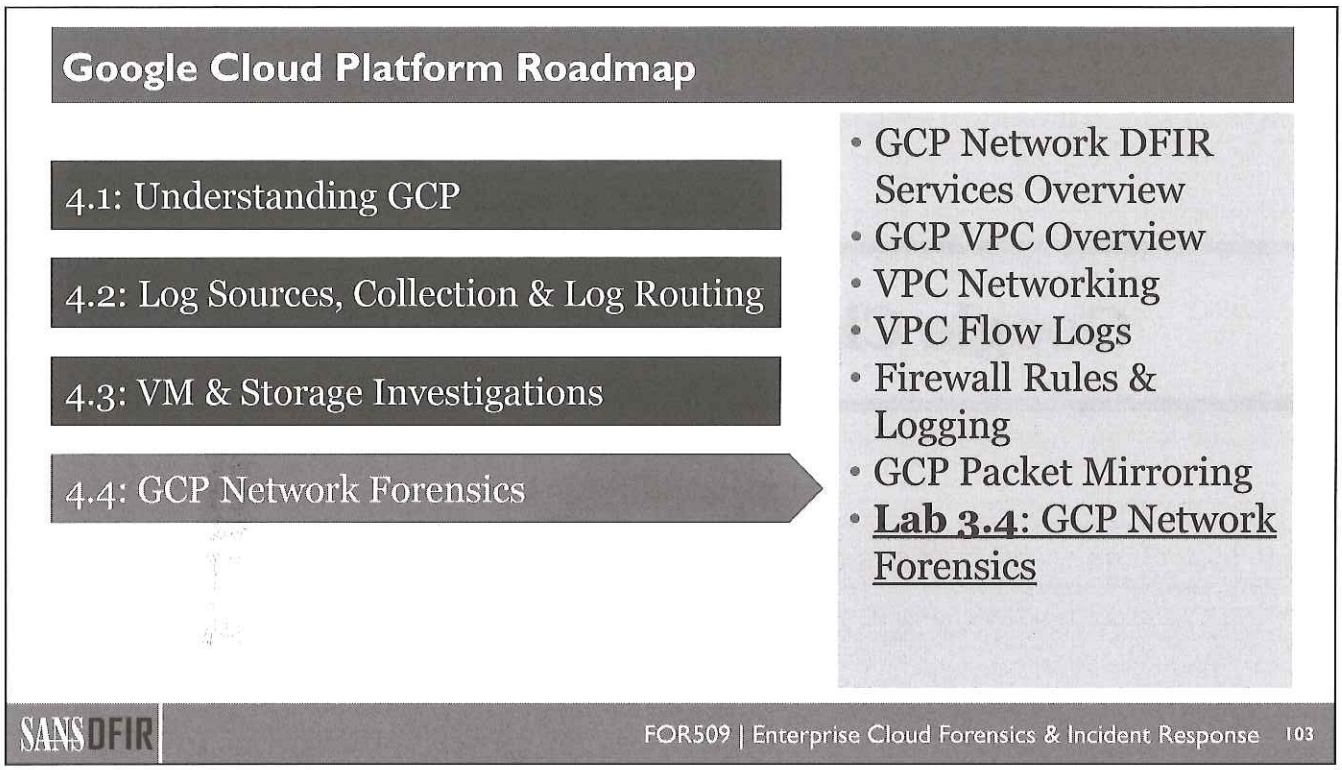

This page intentionally left blank.

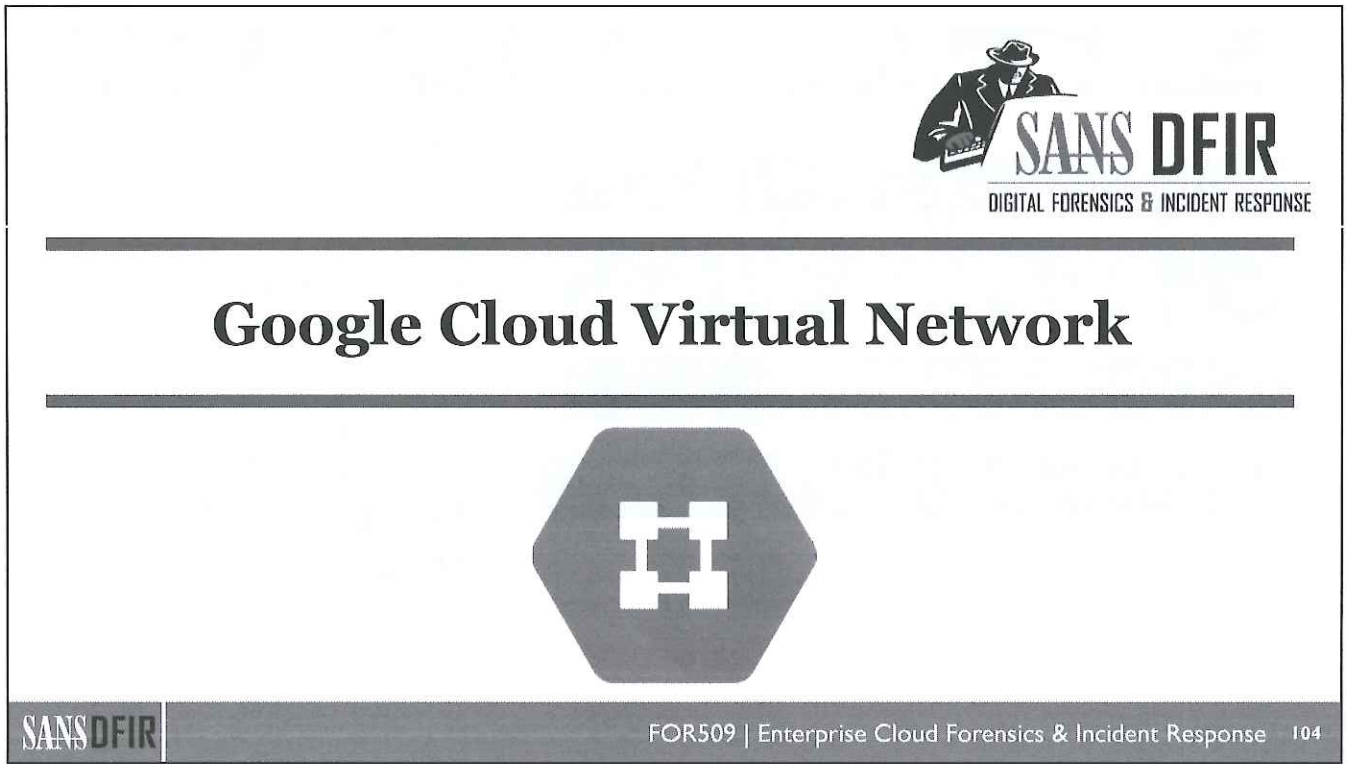

This page intentionally left blank.

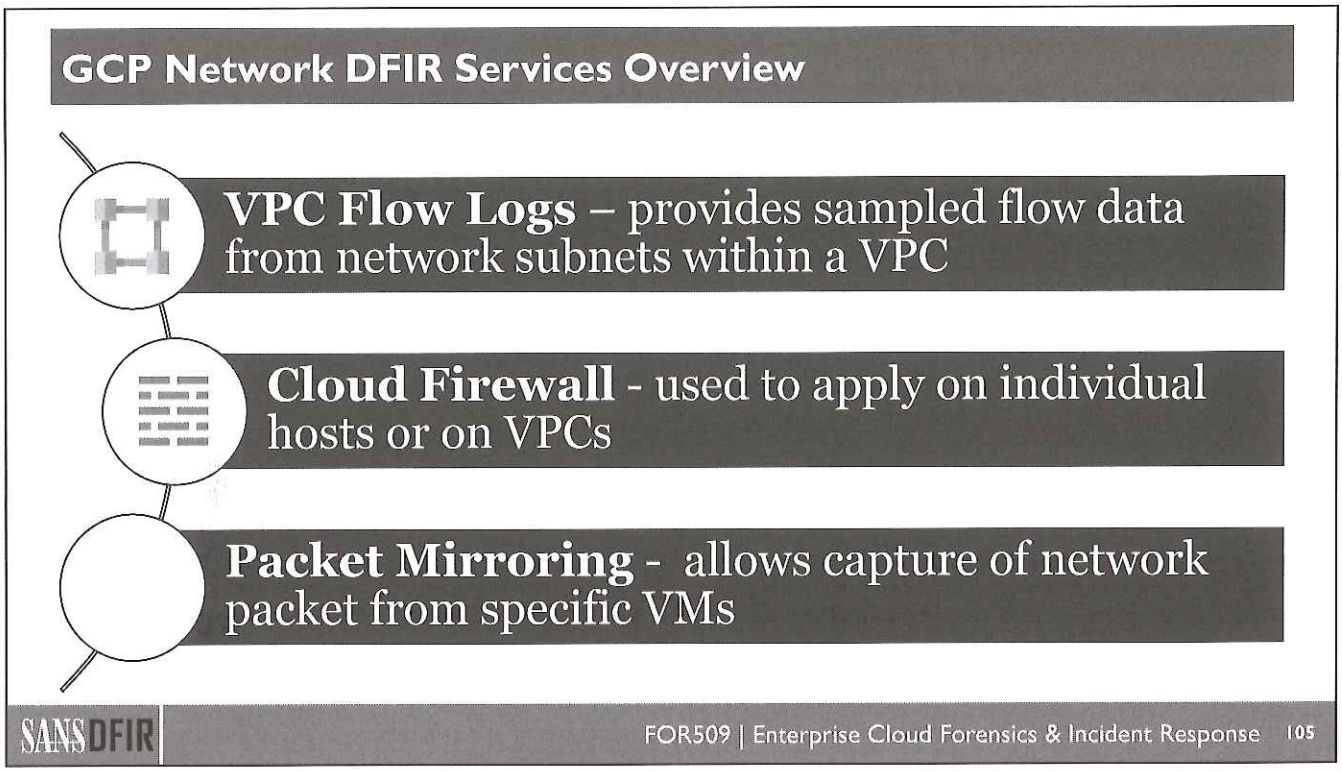

Cloud Firewalls

- Can be defined as a policy at the Folder and Organization level.

GCP VPC Overview

- **Google Cloud's VPC provides network connectivity between compute instances, other VPCs, and VPN tunnels or cloud interconnects.**
- **Andromeda is the technology that support traffic switching and routing for VPCs**
- **Subnets are considered regional, all other VPC resources are considered global (i.e., Firewalls, NAT, Routing)**
- **VPCs only support IPv4 unicast. There is no IPv6, broadcast or multicast traffic within a VPC.**

IPv6 can be achieved by using a global load balancer for public facing services.

References:

[1] https://for509.com/h2klg

[2] https://for509.com/6ctfx

GCP Networking – VPC Network (1)

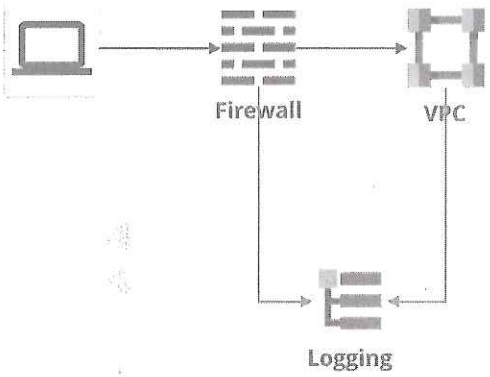

VPC LOGGING

VPC's usually have Firewall rules applied/attached to a VPC – this occurs across all subnets within a VPC.

Firewalls can produce logs per rule, this applies to success or failure. This logging is disabled by default.

VPC's allow provide "sampled" Flow Logging per subnet within a VPC. This is disabled by default.

This page intentionally left blank.

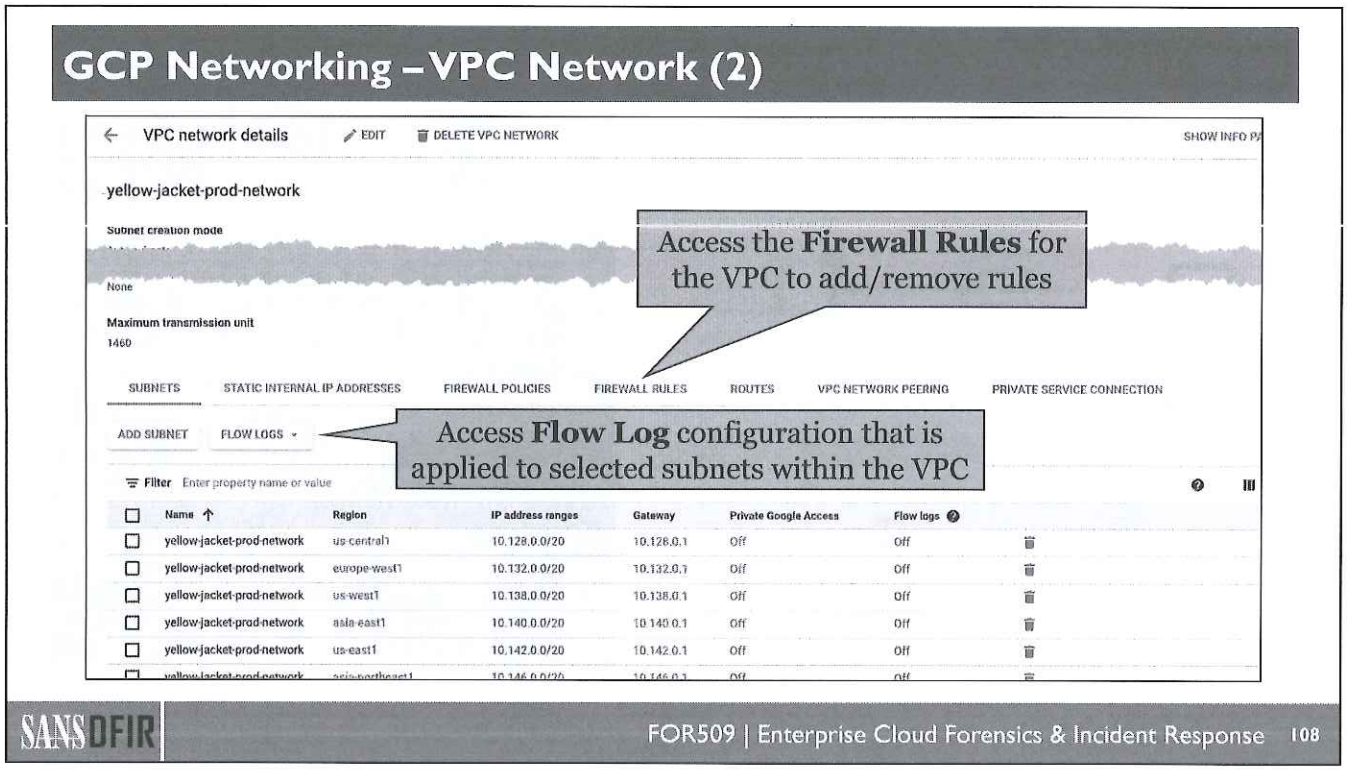

This page intentionally left blank.

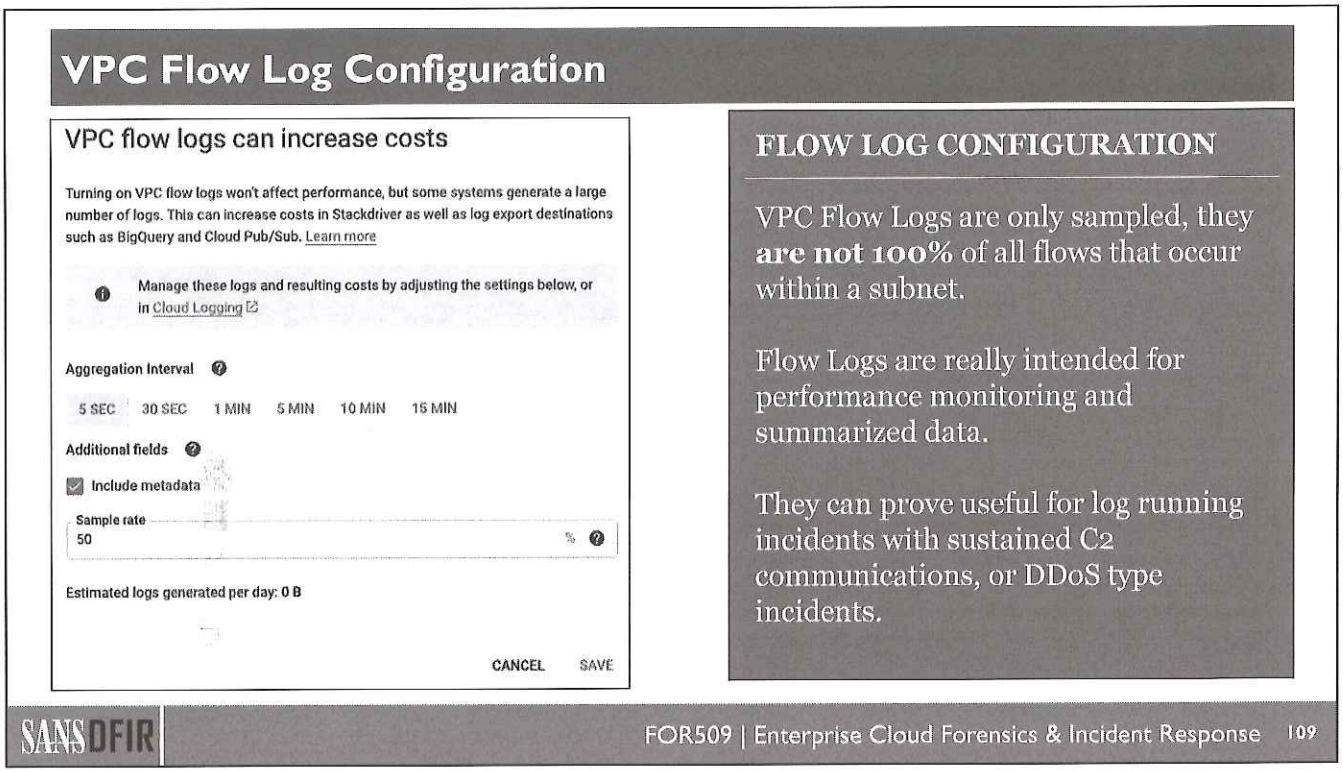

Disabling "Includes metadata" will reduce the overall log storage size.

References:

[1] https://for509.com/knzmq

[2] https://for509.com/aoprn

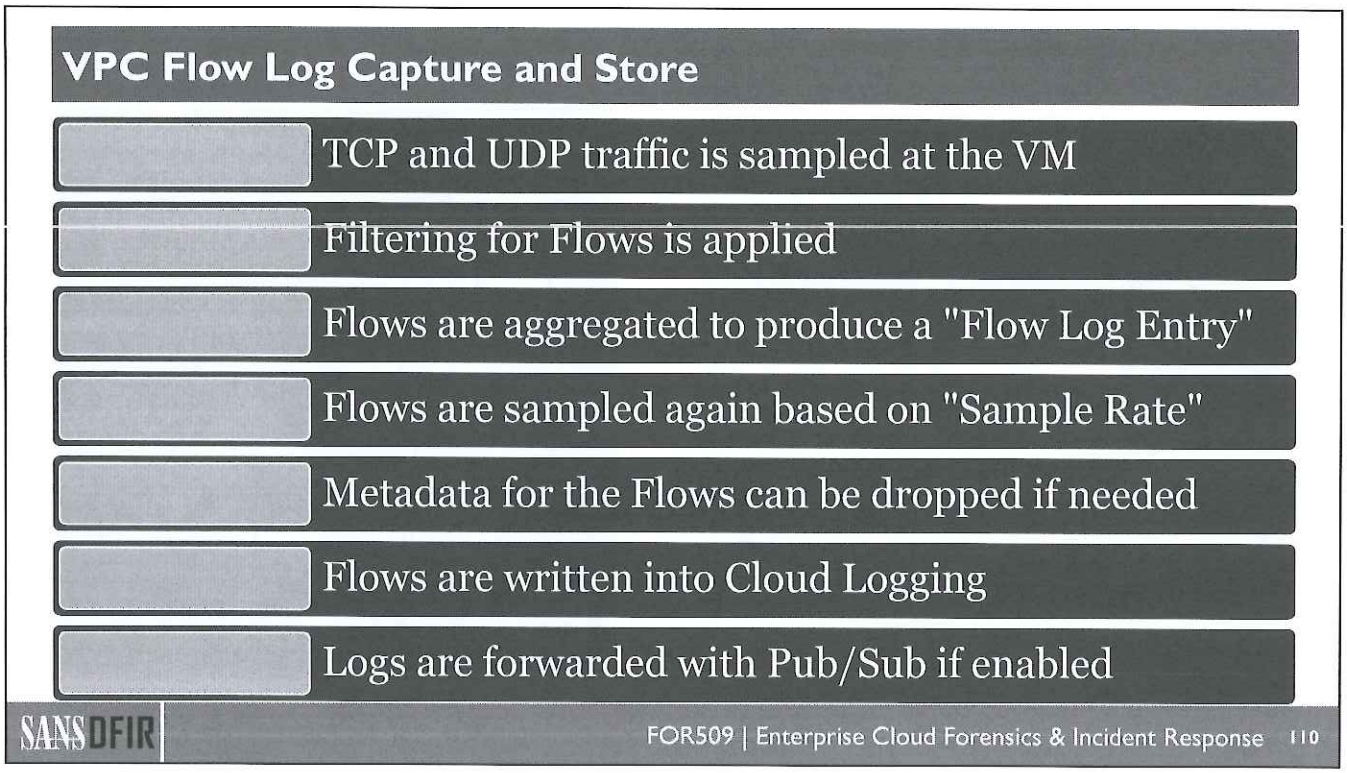

References:
[1] https://for509.com/knzmq

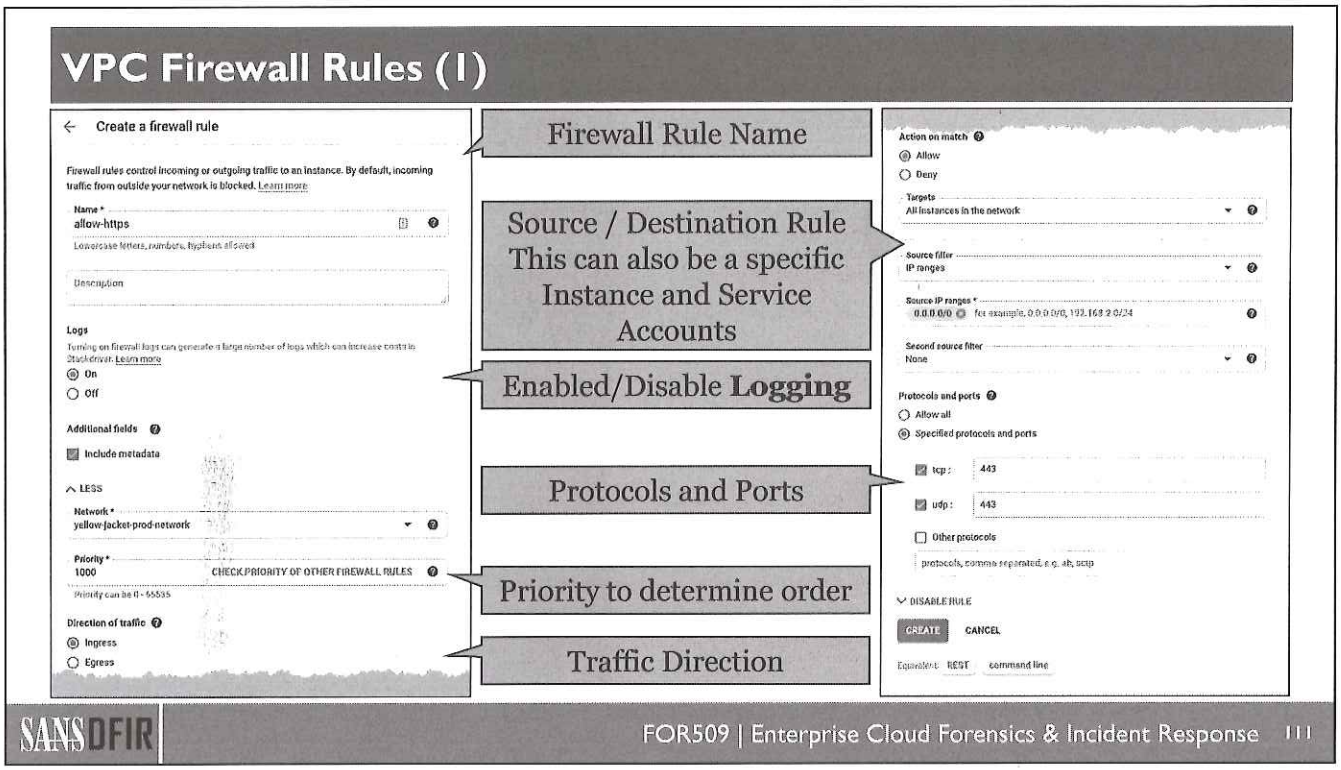

Disabling "Includes metadata" will reduce the overall log storage size.

References:

[1] https://for509.com/798su

VPC Firewall Rules (2)

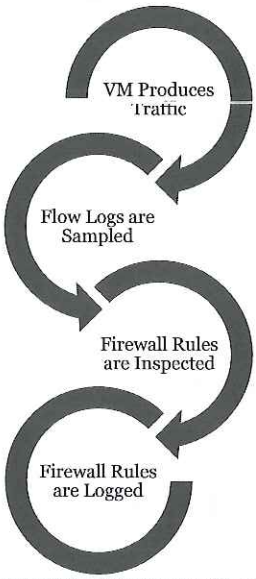

LOGGING ANALAYSIS

Firewall rules are assessed after Flow Logs are generated.

You could see Flow Logs indicating traffic occurred, however, Firewall Rules may prevent the traffic ever leaving the VPC.

Review Firewall Logs and Flow Logs together to determine if traffic ever made it out of your VPC.

References:
[1] https://for509.com/knzmq

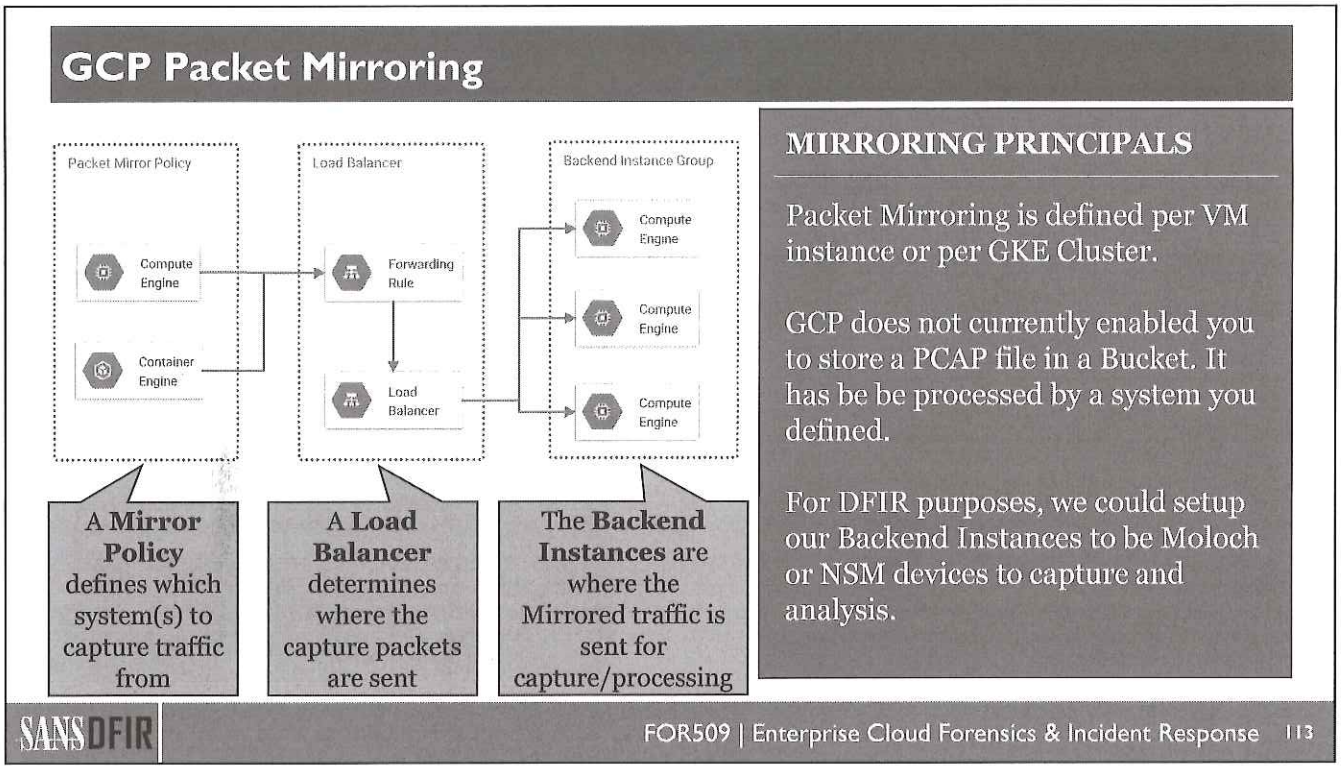

References:

[1] https://for509.com/fp614

[2] https://for509.com/jaq42

GCP Networking – VPC Network (3)

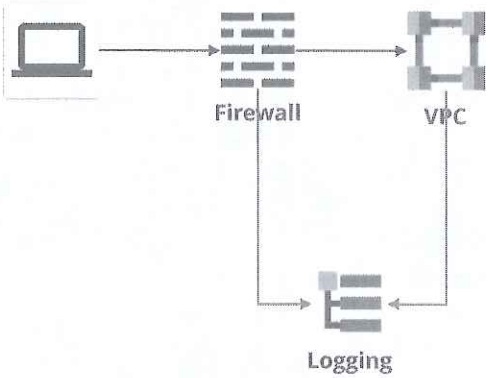

VPC LOGGING

VPC's usually have Firewall rules applied/attached to a VPC – this occurs across all subnets within a VPC.

Firewalls can produce logs per rule, this applies to success or failure. This logging is disabled by default.

VPC's allow provide "sampled" Flow Logging per subnet within a VPC. This is disabled by default.

This page intentionally left blank.

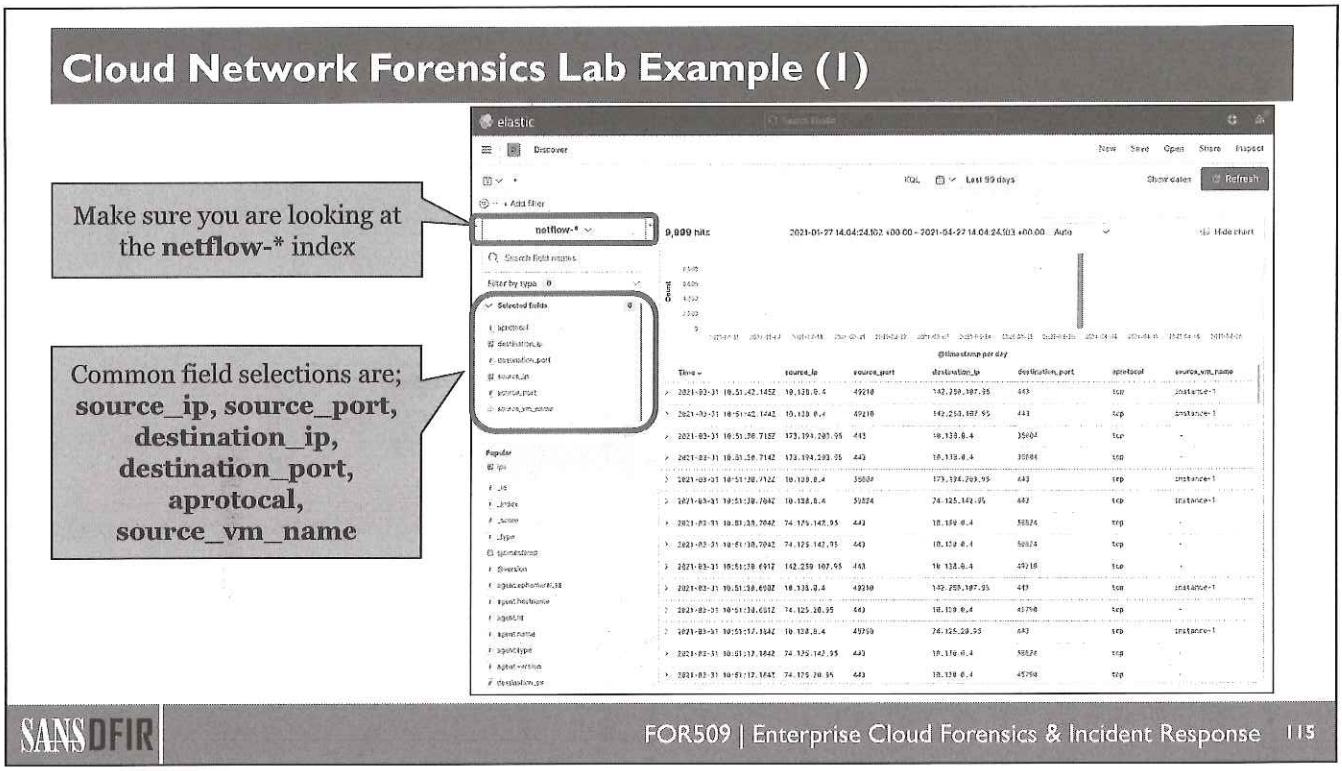

The example above shows you SOF-ELK with the GCP Flow logs extracted from GCP in JSON format and loaded into SOF-ELK.

Unlike the data you have seen so far, which was under the "gcp-*" index, the flow log data is handled in a sperate index called "netflow-*". This is because SOF-ELK is already configured with pre-made dashboards designed to handle flow data from various source, like the ones you've already in the course.

As we've show in previous SOF-ELK configurations, you may want to customize the data you are viewing in the summary view. For flow data there are some fairly common fields; @time, source_ip, source_port, destination_ip, destination_port, aprotocal, along with any additional traffic quantity information you feel is relevant. You should also include the source_vm_name so you understand which system produce or received the flow data you are analyzing.

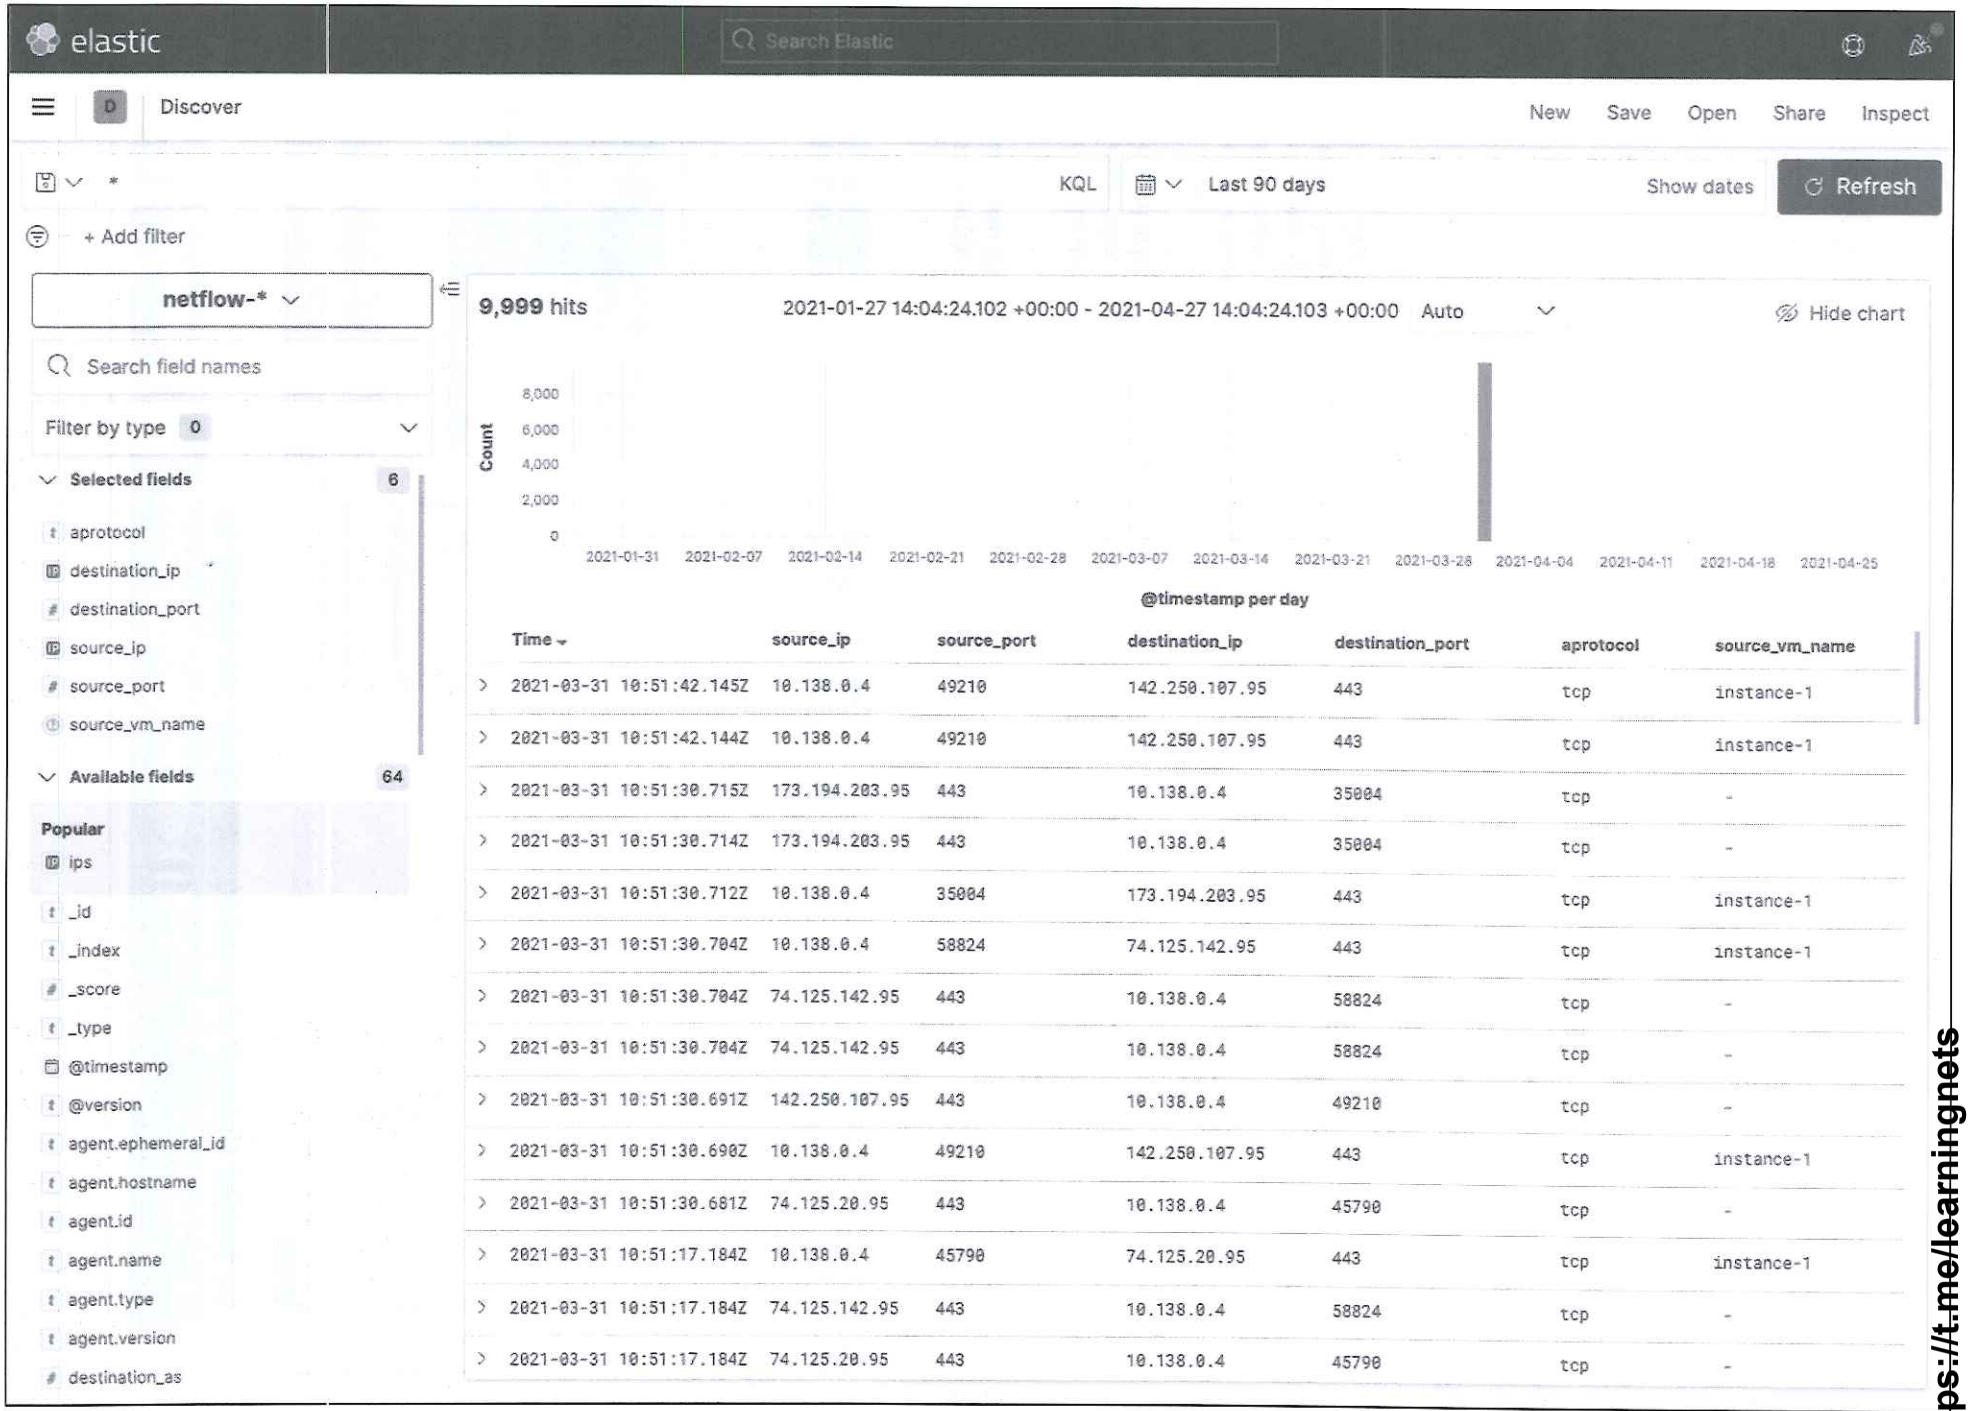

Time	source_ip	source_port	destination_ip	destination_port	aprotocol	source_vm_name
2021-03-31 10:51:42.145Z	10.138.0.4	49210	142.250.107.95	443	tcp	instance-1
2021-03-31 10:51:42.144Z	10.138.0.4	49210	142.250.107.95	443	tcp	instance-1
2021-03-31 10:51:30.715Z	173.194.203.95	443	10.138.0.4	35004	tcp	-
2021-03-31 10:51:30.714Z	173.194.203.95	443	10.138.0.4	35004	tcp	-
2021-03-31 10:51:30.712Z	10.138.0.4	35004	173.194.203.95	443	tcp	instance-1
2021-03-31 10:51:30.704Z	10.138.0.4	58824	74.125.142.95	443	tcp	instance-1
2021-03-31 10:51:30.704Z	74.125.142.95	443	10.138.0.4	58824	tcp	-
2021-03-31 10:51:30.704Z	74.125.142.95	443	10.138.0.4	58824	tcp	-
2021-03-31 10:51:30.691Z	142.250.107.95	443	10.138.0.4	49210	tcp	-
2021-03-31 10:51:30.690Z	10.138.0.4	49210	142.250.107.95	443	tcp	instance-1
2021-03-31 10:51:30.681Z	74.125.20.95	443	10.138.0.4	45790	tcp	-
2021-03-31 10:51:17.184Z	10.138.0.4	45790	74.125.20.95	443	tcp	instance-1
2021-03-31 10:51:17.184Z	74.125.142.95	443	10.138.0.4	58824	tcp	-
2021-03-31 10:51:17.184Z	74.125.20.95	443	10.138.0.4	45790	tcp	-

Cloud Network Forensics Lab Example (2)

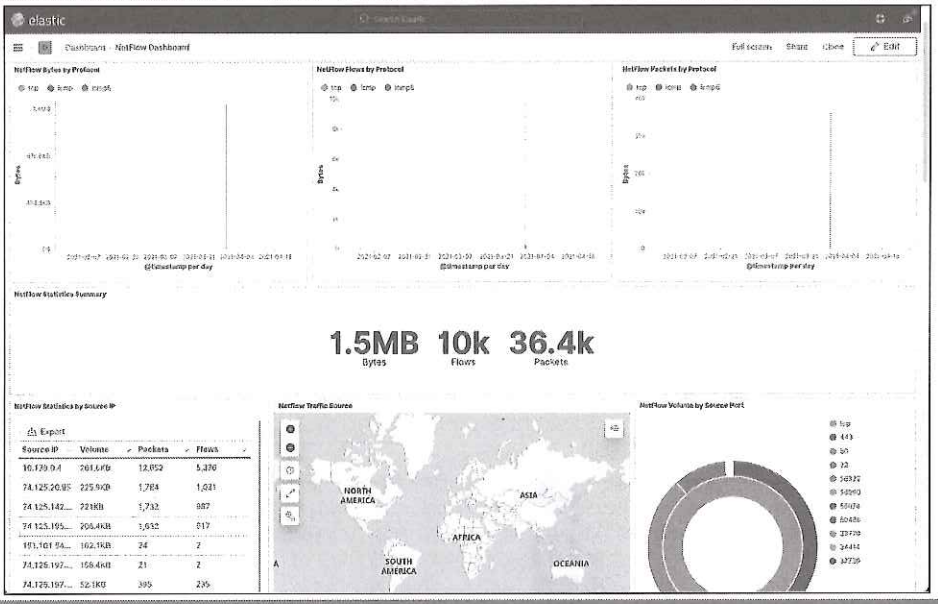

The example shows the pre-made flow data dashboard with the same data you were viewing on the previous page. This uses the power of SOF-ELK to arrange and group flow data. Additionally, this dashboard would also allow you to see flow data information from any other flow data source you loaded into SOF-ELK at the same time. This is simply an example of the same data presented graphically.

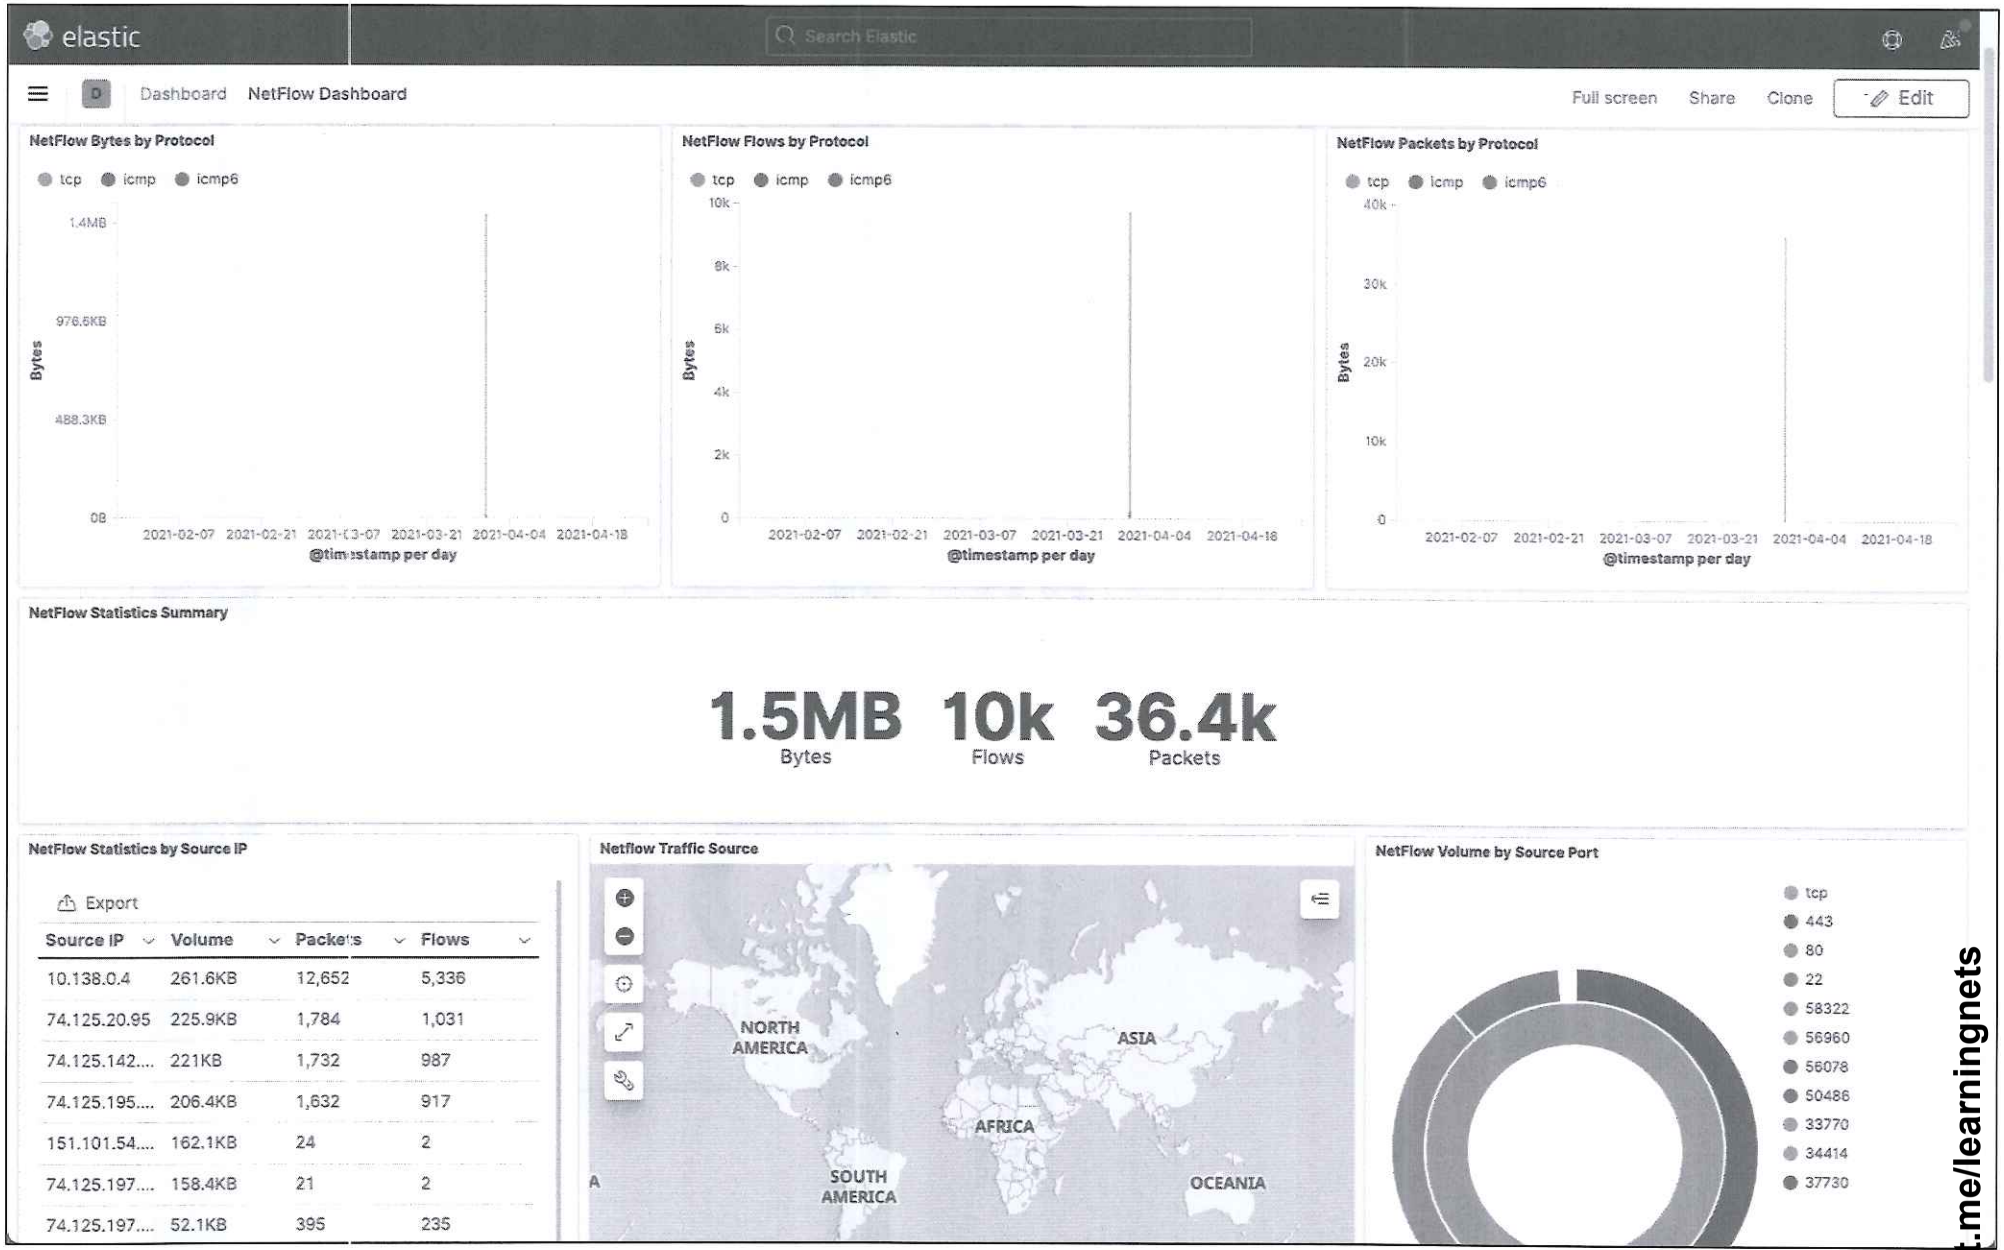

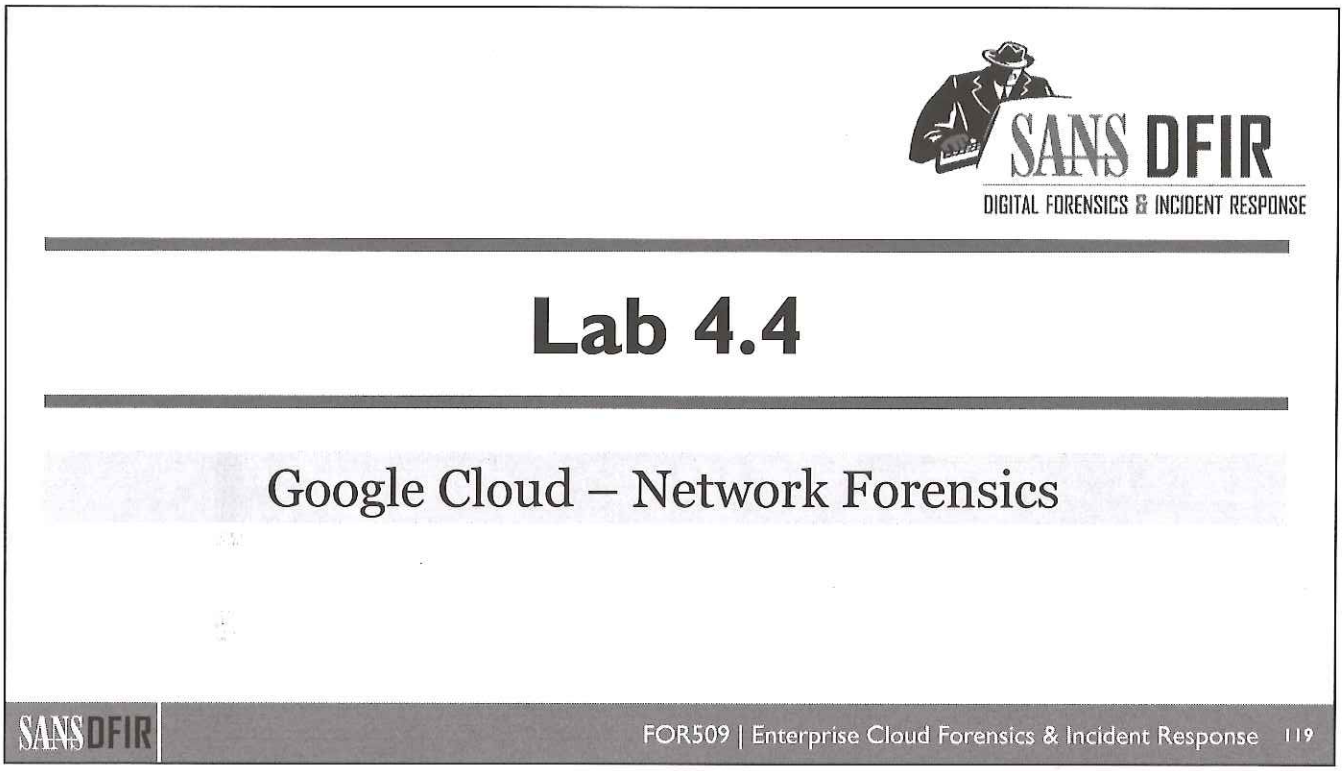

This page intentionally left blank.

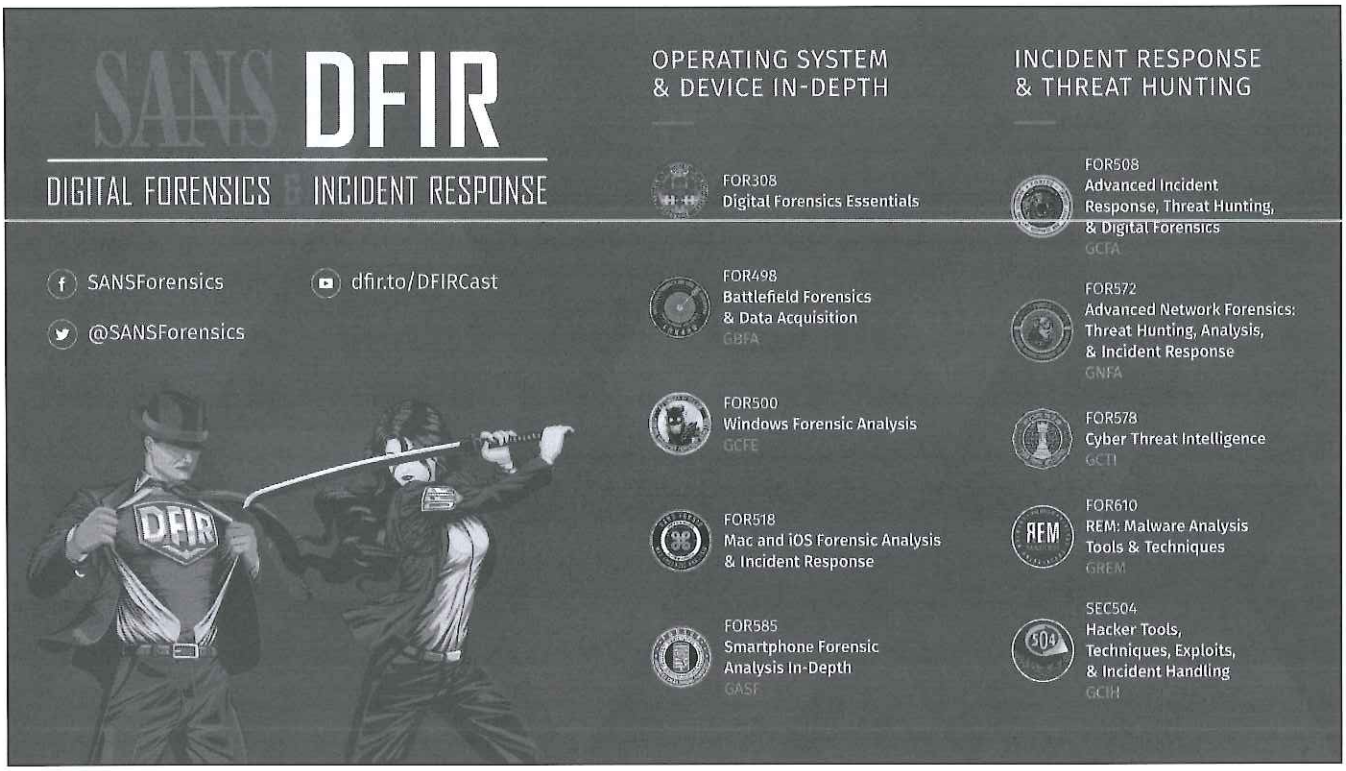

This page intentionally left blank.

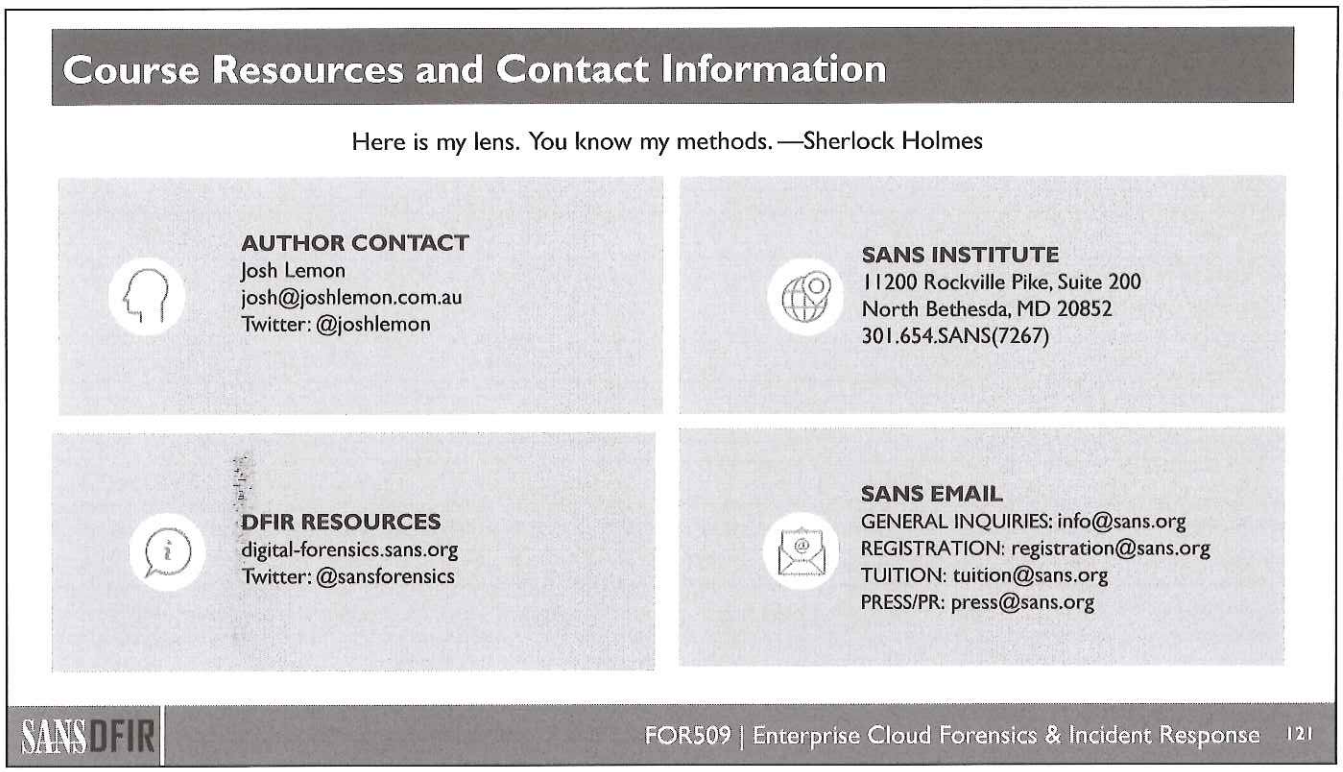

Author: Josh Lemon
Email: josh@joshlemon.com.au
Twitter: @joshlemon

"As usual, SANS courses pay for themselves by Day 2. By Day 3, you are itching to get back to the office to use what you've learned."
Ken Evans, Hewlett Packard Enterprise - Digital Investigation Services

SANS Programs
sans.org/programs

GIAC Certifications
Graduate Degree Programs
NetWars & CyberCity Ranges
Cyber Guardian
Security Awareness Training
CyberTalent Management
Group/Enterprise Purchase Arrangements
DoDD 8140
Community of Interest for NetSec
Cybersecurity Innovation Awards

Search SANSInstitute

SANS Free Resources
sans.org/security-resources

- E-Newsletters
  NewsBites: Bi-weekly digest of top news
  OUCH!: Monthly security awareness newsletter
  @RISK: Weekly summary of threats & mitigations
- Internet Storm Center
- CIS Critical Security Controls
- Blogs
- Security Posters
- Webcasts
- InfoSec Reading Room
- Top 25 Software Errors
- Security Policies
- Intrusion Detection FAQ
- Tip of the Day
- 20 Coolest Careers
- Security Glossary

SANS Institute
8120 Woodmont Avenue | Suite 310
Bethesda, MD 20814
301.654.SANS(7267)
info@sans.org