Disaster Avoidance in Google Cloud Implementations

Rukayat Damilola Alimi
Hood College
401 Rosemont Avenue
Frederick, MD 21701

Olusola Gbenga Olufemi
Hood College
401 Rosemont Avenue
Frederick, MD 21701

ABSTRACT

One of the great benefits of cloud-based services is its ability to provide easy deployment [8]. Then, the right security mindsets and awareness for proper cloud administration will pay off more, to securely utilize these services. However, as all cloud providers will always say – the security of the cloud is their responsibility, the security in the cloud lies with the client users. This saying invariably makes the study of what the users of these cloud services need to know and act on so much important. The focus in this study will be on Google cloud, a user-friendly and cost-saving cloud, which offers good ROI and peace of mind for business users. Google Cloud Platform (GCP) and G Suite (now Google Workspace), are the distinct Google entities that get cloud customers set up completely for the journey to Google Cloud services' adoption [6].

General Terms

Security, Disaster, Implementation, Administrator, Password, Traffic, Permissions, Monitoring, Networks, ROI.

Keywords

GCP, G Suite, Workspace, 2SV, GKE, Kubernetes, Roles, Compute Engine, IAM, Misconfiguration, Cluster, Remote Access, SSH, API.

1. INTRODUCTION

How proficiently can adoption or moving to Google Cloud be done without increasing the risk of a security breach or data loss? In what ways can one best assess Google Cloud environment for misconfigurations? The response to these questions will surely draw attention to necessary cloud security measures. In summary, quickly identifying threats and misconfigurations, and then mitigating them before these cause damage or data loss can be considered a more straightforward answer. Google offers the following major cloud services in Google Cloud Platform (GCP): Compute, Networking, Storage & Databases, Big Data, Machine Learning, Identity & Security Management, and Developer Tools. A key security feature in GCP is Cloud Security Command Center (Cloud SCC), which provides customers with a centralized security tasks dashboard. This offers customers necessary precautionary assistance to prevent, detect, and respond to threats in their GCP cloud environment, once turned on and sets up correctly. All built-in security capabilities and products in SCC also help in increasing customers' visibility into misconfigurations, vulnerabilities, and threats in their cloud environment [8].

On the other hand, G Suite (now Workplace), a bundle of productivity applications, is another Google Cloud product comprised of Gmail, Meet, Drive, Chat, Docs, Sheets, Slides, Forms, Sites, and more. This also provide to customers more valued benefits in Google cloud. Google Cloud offers a large array of services that can be tuned to meet with business specific needs when properly managed. Google Cloud provides modern security infrastructure from the user, to the device, to the application, to the platform [8].

Google have built everything from hardware, networking to the custom Linux software stack with customer security in mind. Homogeneity, combined with ownership of the entire stack, have greatly reduce Google security footprint and allow customers to react to threats faster on their own [8]. In the succeeding sections, it will be provided in details measures cloud admin users need to follow, to keep the cloud environment with enormous sensitive data safe as much as possible, while managing Google cloud services. Cloud advent is accompanied by developmental opportunities, as well as challenges, which are mostly security problems [1].

2. 2SV ENFORCEMENT

To start with, an Identity Access Management (IAM) system in Google Cloud is used to define and control users' access to production services, using security protocols that authenticate users, through the provision of short-lived personal public key certificates. The issuance of these personal certificates is in turn guarded by 2-Step Verification (2SV) [5]. 2SV provides an extra barrier between business and cybercriminals, trying to steal credentials (usernames and passwords) to access business data [5]. Making sure all GCP users belong to the G Suite domains the administrator oversees is very necessary. In so doing, he can effectively enforce 2SV. From G Suite domain, the Cloud admin should follow recommended Google instructions to enable and enforce 2SV. Enforcing 2SV with a short enrollment grace period, while still enforcing secured remote social engineering and access, is a superior approach compared to using passwords only [5]. Turning on 2-Step Verification is a single most important action administrator can take to protect business data. 2SV can ward off bad guys, even if they are in possession of organization's account passwords [5]. Hence, with 2SV, cloud accounts can best be protected with both password and phone alerts [5]. As one can deduce in Figure 1, we can literally say that 2SV resides in-between the user password and the user account to ensure the safety of his sensitive data.

3. INHERITED IAM PERMISSIONS

3.1 IAM Policies

Roles are granted to users by creating an IAM policy, which is a collection of statements that defines who has what type of access. A policy is attached to a resource and is used to enforce access control whenever such resource is accessed by different users in an organization [11].

For shrewd prevention of G Suite and GCP Organization takeover by unscrupulous users, create admin only user account(s) and remove admin privileges from all "normal" users. In GCP, the all-embracing structure is based on an Organization, which usually corresponds to a domain name, like *adapt.com*. Within the Organization, there are Folders, which contain either other Folders or Projects. Folders stand as Organizational Units (OUs), to represent certain business units. Projects contain resources, such as databases and virtual machines. The Figure 2 absolutely pictures how Organization, Folders, Projects and Resources are all interwoven and hierarchical. Users can be made to have access to the Organization, or just some parts of it. Misconfiguration here, such as giving a normal user 'Project Owner' permissions, could cause an administrator to lose control of Projects in Organization, as this permission will be automatically inherited throughout the Organization. As a reminder, Cloud administrator should endeavor to architect permissions perfectly well. It should be well-noted that child resources inherit policies from the parent resources' policies. Let's take for instance a policy for a project granting a user the ability to administer (VM) instances, then a normal user can administer any VM in that project, regardless of the restriction policy that is set on each VM in the project [12]. Figure 3 shows how parent's policies are being inherited by a child.

3.2 Title and Authors

Damilola studied Information Technology, Hood College - Maryland. She's a wife, a mother and an ardent researcher, with specialization in IT experimentation and innovations in digital device profound values and precautionary utilization.

Olusola is a Cloud and Information Security Specialist. He studied Information Technology. Olusola is a husband, a father, and has been a life-long learner.

4. ADMINISTRATIVE ROLES

It is always a best practice to avoid basic roles such as roles/owner, roles/editor, and roles/viewer, for security-critical resources. Rather, approve of predefined roles to allow the least-permissive access necessary [12]. Let the smallest scope needed be the only reason for granting roles. For example, if a user only needs access to publish a specific topic, grant a Publisher role to the user for that specific topic [12]. One may like to use Google-managed role, then choose the predefined role provided by Google in the platform. Google also provides the option to create fresh or new custom roles that include only the recommended permissions the administrator wants to apply or effect [13]. This custom role can be modified at will by adding or removing permissions.

Putting more emphasis on this, as a best practice for Cloud administrator, use more granular permissions for your users instead of primitive roles, so that users are given the least privileges necessary to complete their tasks. Also, watch and alert on primitive creation and existence in your GCP environment, especially at the Organization root. Ensure users are not assigned roles directly. Instead, groups are the primary method of assigning roles and permissions in Cloud IAM. A provided on-premises or Google managed Active Directory may be used to assign users group membership, hence group creation and membership are strictly controlled [9]. Google has always recommended adopting predefined roles wherever possible. Hence, ensure that IAM permissions that are granted to users and groups are not overly permissive. The administrator should eliminate the use of primitive roles in Google Cloud Organization, since the wide scope of permissions inherent in these roles goes against the principles of least privilege [9]. Hence predefined roles, which are designed with inherent separation of duties should be used always, then add custom roles as needed in the cloud security architecture [9].

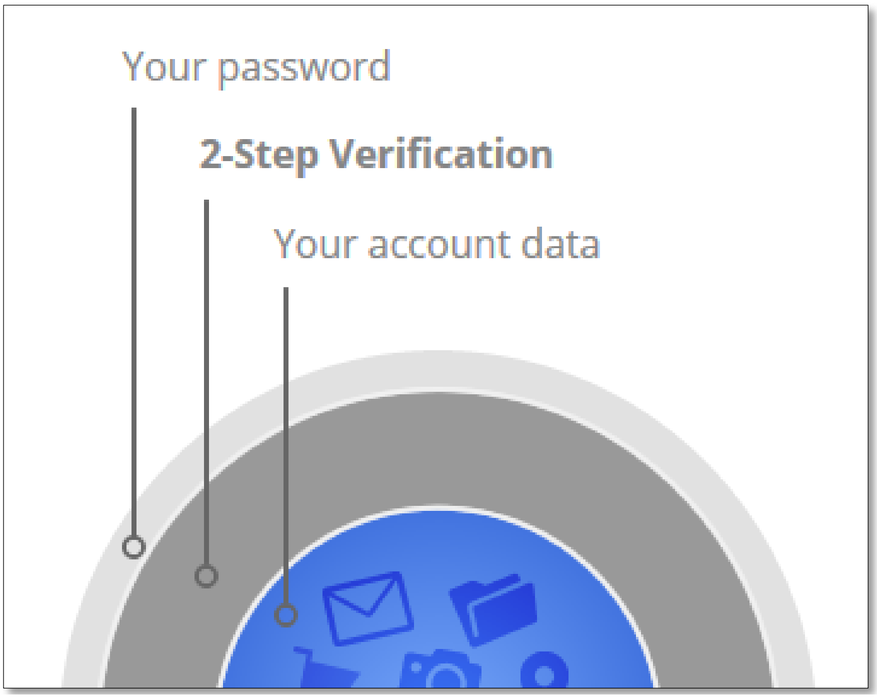

Figure 1: 2FV as a defensive intermediary - Google

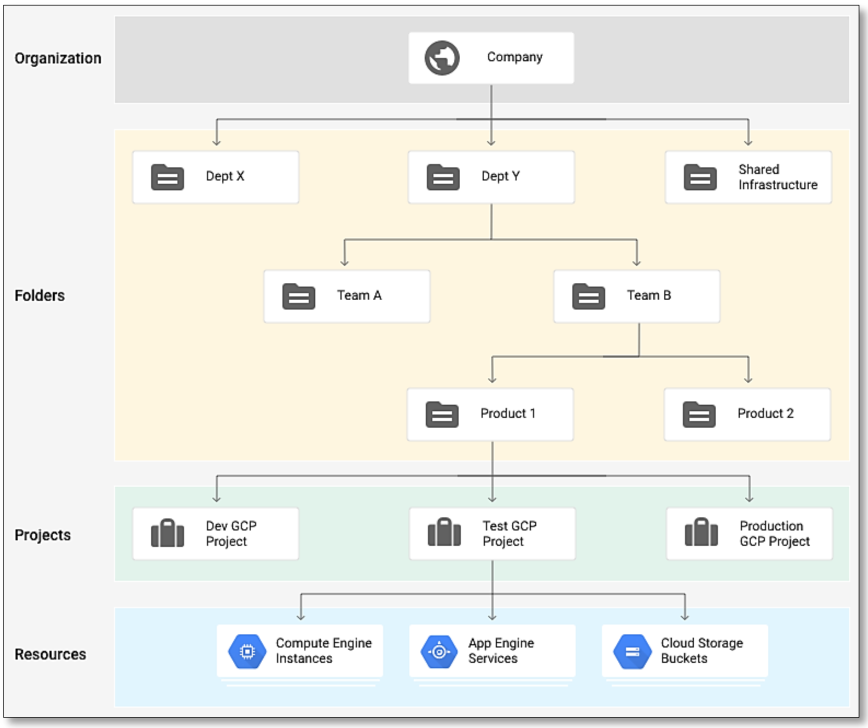

Figure 2: Organization, Folders, Projects and Resources hierarchy – Google

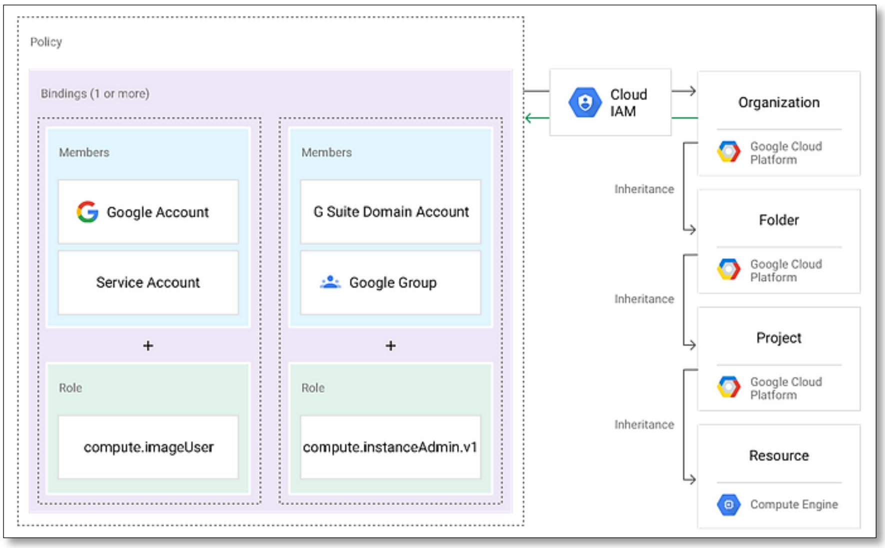

Figure 3: Inherited IAM Permissions - Google

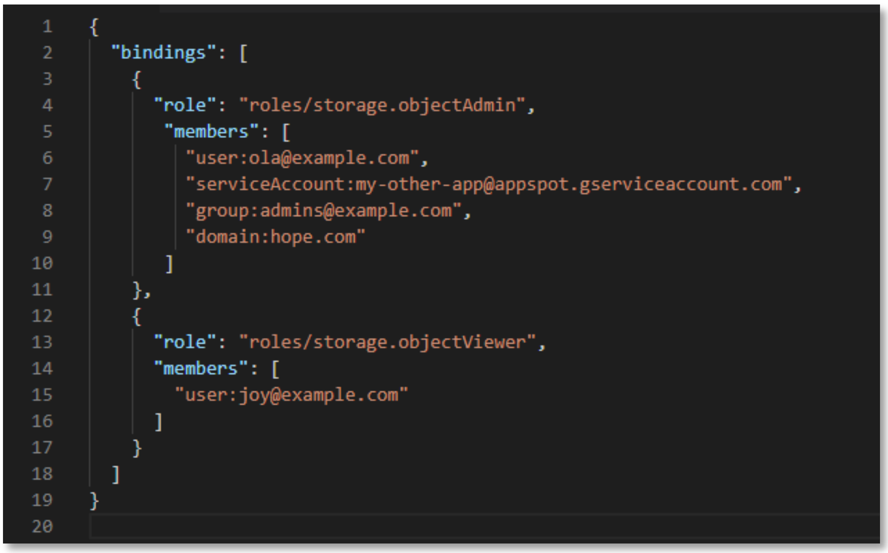

```
{
  "bindings": [
    {
      "role": "roles/storage.objectAdmin",
       "members": [
         "user:ola@example.com",
         "serviceAccount:my-other-app@appspot.gserviceaccount.com",
         "group:admins@example.com",
         "domain:hope.com"
       ]
    },
    {
      "role": "roles/storage.objectViewer",
      "members": [
        "user:joy@example.com"
      ]
    }
  ]
}

```

Figure 4: Code snippet showing the structure of an IAM policy

5. NETWORKS ACCESS IN GKE

Kubernetes is known to be an open-source container-orchestration system for automating computer application deployment, scaling and management. This is one of the wonders of Google security mechanisms. Google Kubernetes (GKE) is a great solution to quickly utilize Kubernetes for cluster management and orchestration. GKE master nodes are not in an administrator's GCP Compute Network, as configured by Google. GKE master nodes are hosted in a Google owned project network. Then, this are connected to an organization's worker nodes via VPC Network Peering [7]. Hence, the master nodes are publicly accessible. This means by default that the control plane of GKE cluster and nodes have addresses that are internet routable, therefore can be accessed from any IP address [9]. Best practices to ward off intruders include creating an Authorized Network for GKE Master access, in accordance to Google recommendations in documentation. This entails adding authorized networks for cluster master access. To disable direct internet access to nodes, specify the gCloud tool option *--enable-private-nodes* at cluster creation [9]. Keeping version of Kubernetes in use up to date is one of the simplest things you can do to improve security in clusters. Kubernetes frequently introduces new security features and provides security patches, which must be taken into consideration and be installed to fortify the KGE [9].

6. DEFAULT NETWORKS AND RULES

Google cloud is pre-populated default network with firewall rules which can be deleted or modified [3]. Cloud Interconnect and managed VPN allow you to create encrypted channels between your private IP environment on-premise and Google's network [15]. This allows you to completely disconnect instances from the public internet, while the instances are still being accessible from one's own private infrastructure [15]. All traffic coming from outside a network is blocked and no packet is allowed into an instance, except with clear firewall rules by default. To allow incoming network traffic, you need to set up firewalls to permit these connections [3].

By default, when you create a Project, a default network and some default network rules are created for all regions and availability zones [3]. Best practice includes removing unused region compute networks from all Projects. All the default network rules allowing SSH, RDP, and ICMP from the Internet should also be removed [3]. This enforces explicit network security rules. However, avoid deleting the rule *gke-<cluster_name>-<random-characters>*, since this is needed for the GKE master to node SSH communication [7]. This firewall rule and network metadata is leveraged by GKE to perform SSH access. In addition, for internal compute communications to further enforce least permissive access, use compute labels with network firewall rules [2]. Use CIDR whitelisting instead of labels, if network communication is from or to outside project's networks [3]. Set alert to monitor any default network configurations. All default rule existence or overly permissive broadly scoped access rules should also be removed promptly.

7. INSTANCE REMOTE ACCESS

VMs and APIs can be configured so that they're not connected to the public internet, but still accessible to system administrators in Google Cloud. To protect VMs, the first thing that needs to be done is to configure a policy that does not allow VMs from obtaining an external IP. As already known, default VPC has firewall rules which allow SSH / RDP, even without an external IP. This means access is only possible from the internal network. In Google Cloud, Cloud Identity-Aware Proxy (IAP) provides a better way to allow secured connection to the machines [14].

Best practices include allowing instances be configured to ignore Project-wide public SSH keys, except for GKE compute nodes. Google documentation can give more light on "Adding or removing project-wide public SSH keys." Google documentation also recommends adding public keys at the instance metadata level. GKE registers its public keys at the Project level, hence GKE compute nodes should be excluded from the project-wide blocking [7]. GKE master can be prevented from being able to connect to your nodes via SSH for cluster communication if it is blocked. This blockage may be avoided if admin writes own automation script to replicate GKE SSH public keys at the Project level to the instance level on GKE cluster nodes instead [7]. Moreover, activating the virtual serial console at the Project level should be avoided. There be in place a way to detect this, get alerts, and auto-remediate if that can be done. Activating virtual serial console allows public access to a compute instance, once SSH keys and Project are known, with no actual VPC network communication. Hence, firewall rules will not block this access.

8. CLUSTERS OPEN TRAFFIC COMMUNICATION

Several methods can be used to control what traffic should flow through clusters and their Pods [10]. All ingress and egress traffic have free flow into and out of all Pods by default, when there are no network policies defined. Tags are used to define the traffic flowing through Pods with network policies [10]. The pods are not isolated and will accept network traffic from any source by default, when a cluster created in GKE [7]. To use a network policy, either create a cluster with a network policy or update your clusters. Enabling network policies will allow you configure these clusters to your specifications.

9. INSTANCES PUBLIC IP ADDRESSES

As much as possible, as a best practice, run services on private IPs, and only expose services via secured public gateways, firewalls, load balancers, Web Application Firewalls (WAFs) or Content Delivery Networks (CDNs), which can log and restrict access where needed [7]. Alerts for validation should be received for any instances created with public IPs, since this configuration increases external exposure and risks.

10. GCP ENABLED SERVICES APIs

The management of all services in GCP are done through a global API gateway infrastructure which are secure [16]. This API-serving infrastructure is only accessible over encrypted SSL/TLS channels, and every request must include a time-limited authentication token, generated via human login or private key-based secrets through the authentication system described above [17]. Hence, it is vital to monitor and alert on whatever key services one is using. This is because malicious use of cloud resources can cost organization to get a very large bill from Google. Hence, disable the GCP APIs for unused services until when needed. Watch and alert/auto-remediate if they are enabled. Do overall monitoring of the GCP space.

11. DEFAULT ENCRYPTION

Google Cloud services always encrypt customer stored data at rest, except in very few cases. Encryption at rest is automatic, and no customer action is required. Any new data stored in persistent disks is encrypted with 256-bit Advanced Encryption Standard (AES-256), and each encryption key is itself encrypted with a regularly rotated set of master keys [4]. It is indeed a great delight GCP provides encryption by default for any data at rest. Still, Google recommends maintaining control over secrets by leveraging CMEKs (or customer supplied keys), wherever in configuration it's being supported [7]. Key Management System (KMS) allows admin to easily assign fine-grained permissions to anyone using the keys. This can help data labeling efforts, and can offer additional outlier visibility through data key usage monitoring/alerting [7].

12. ACKNOWLEDGMENTS

Appreciations first and foremost to the Almighty God, the giver of wisdom, knowledge and understanding. Also appreciated are institutions, families, friends and loved ones who have been sources of encouragement in our research journey so far.

13. CONCLUSION

A comprehensive security for Google Cloud requires continuous visibility and control of the many services within Google customer environment. This can be an overwhelming task for customer organizations that are new to Cloud adoption, and the experienced in Cloud usage. More and more organizations and individuals in the U.S, and in many other countries are falling victims because of the diverse weaknesses made opened to ruinous men on the internet [18]. This invariably means that the dividends of security in the Cloud far outweigh its cost, as more threats are being discovered and defeated continuously. Then, any organization putting in more resources in ensuring security while in the cloud gains more, since security effort is a profitable investment in the long run. Cloud problems are not just technical, as these also encompass standardization, the mode of Cloud supervision, laws and regulations, and a host of others [1]. It is also encouraged to consult Google Cloud security documentations regularly, in order to keep abreast with Google Cloud latest updates in use.

14. REFERENCES

[1] Xiaowei Yan et al. 2012. Research and Design of Cloud Computing Security Framework. Available from: https://www.researchgate.net/publication/259764167. [accessed Oct 06, 2020].

[2] Chinta ChandraSekhar et al. 2014. Secure Network Connectivity in the Cloud Computing Environment. International Journal of Innovative Research in Computer and Communication Engineering. Vol. 2, Issue 3, March 2014 [accessed Oct 01, 2020]

[3] VPC Firewalls Rules Overview, https://cloud.google.com/vpc/docs/firewalls [accessed Oct 02, 2020]

[4] M Husni et al. 2020. Security Audit In Cloud-Based Server by Using Encrypted Data AES-256 and SHA-256, IOP Conference. Series: Materials Science and Engineering 830 (2020) 032015 [accessed Oct 03, 2020]

[5] How it Protects, https://www.google.com/landing/2step/index.html#tab=how-it-protects, [accessed Oct 07 2020]

[6] https://support.google.com/a/answer/7284269?hl=en, [accessed Oct 06, 2020]

[7] Securing GCP: Top ten Mistakes to avoid, https://resources.netskope.com/cloud-security-solution-white-papers/securing-gcp-top-ten-mistakes-to-avoid, [accessed Oct 09, 2020]

[8] Google Cloud Security Foundations Guide, https://services.google.com/fh/files/misc/wp_take_command_of_your_security_in_the_cloud_rgb_v15c.pdfhttps://services.google.com/fh/files/misc/google-cloud-security-foundations-guide.pdf, [accessed Oct 01, 2020]

[9] Hardening You Cluster, https://cloud.google.com/kubernetes-engine/docs/how-to/hardening-your-cluster. [accessed Oct 08, 2020]

[10] https://cloud.google.com/kubernetes-engine/docs/concepts/security-overview#node_security

[11] https://cloud.google.com/iam/docs/overview, [accessed Oct 11, 2020]

[12] Using IAM Securely, https://cloud.google.com/iam/docs/using-iam-securely, [accessed Oct 09, 2020]

[13] https://cloud.google.com/iam/docs/recommender-overview, [accessed Oct 07, 2020]

[14] Configuring Secure Remote Access for Compute Engine VMs, https://cloud.google.com/blog/products/identity-security/configuring-secure-remote-access-for-compute-engine-vms, [accessed Oct 09, 2020]

[15] https://ldapwiki.com/wiki/Google%20Cloud%20Securit, [accessed Oct 07, 2020]

[16] Enabling and Disabling Services, https://cloud.google.com/service-usage/docs/enable-disable, [accessed Oct 02 2020]

[17] Improving the security of Google APIs with SSL, http://googlecode.blogspot.com/2011/03/improving-security-of-google-apis-with.html [accessed Oct 01, 2020]

[18] O. G. Olufemi, R. D. Alimi. 2020. Authenticating Device Users via Keyboard Strokes, IJCA - International Journal of Computer Applications, Foundation of Computer Science (FCS), NY, USA, Vol. 175, Issue 17, https://www.ijcaonline.org/archives/volume175/number17/31542-2020920673.