

Penetration Test Report

Hlybochytska Street, 17B, Kyiv, Ukraine
office@eska.global
eska.global
+380 (44) 247 10 21

Table of Contents

Executive Summary	4
Summary of Results	5
1. INFRASTRUCTURE ASSESSMENT	7
1.1. Discovery	7
1.2. Cloud SQL testing results	8
1.3. Cloud Storage testing results	13
1.4. Cloud Compute Engine testing results	15
1.5. IAM testing results	23
1.6. Kubernetes Engine testing results	29
1.7. Stackdriver Logging & Monitoring testing results	37
1.8. Conclusion	42
2. APPLICATION ASSESSMENT	43
2.1. Introduction	43
2.2. User enumeration	44
2.3. Google Captcha bypass	47
2.4. SSRF on JSON API functionality	50
2.5. XSS	54
2.6. IDOR-Privilege Escalation	57
2.7. Security Misconfiguration - Exposed Test environment	60
2.8. Code Review	62
2.9. Advisory information	63
2.10. Risk Rating	65

Appendix A: Infrastructure Assessment Results	66
Cloud SQL test results	66
Cloud Storage test results	70
Cloud Compute Engine test results	72
IAM test results	78
Kubernetes Engine test results	82
Stackdriver Logging & Monitoring test results	86
Appendix B: Vulnerability Detail and Mitigation	89
Risk Rating Scale	89
SSRF on JSON API functionality	89
XSS	90
IDOR-Privilege Escalation	90
Google Captcha bypass	91
User enumeration	91
Security Misconfiguration - Exposed Test environment	92
Code Review: Insecure Randomness	92
Appendix C: About ESKA	94

Executive Summary

ESKA was contracted by ***** to conduct a penetration test in order to determine its exposure to a targeted attack and complete an infrastructure assessment to evaluate configurations regarding security best practices. All activities regarding the penetration test were conducted in a manner that simulated a malicious actor engaged in a targeted attack against ***** with the motivation to:

- Identify if a remote attacker could penetrate *****'s defenses.
- Determine the impact of a security breach on:
 - confidentiality of the company's private data
 - internal infrastructure and availability of ***** information systems.

The penetration test was expanded with source code analysis for the determination of programming errors and insecure data flows. Emphasis was placed on the identification and exploitation of security weaknesses that could allow a remote attacker to gain unauthorized access to organizational data. The attacks were conducted with the level of access that a general application user would have. The source code analysis was conducted with credential and access arrangements provided by ***** . The assessment was conducted in ***** with the recommendations outlined in the NIST SP 800-115 "Technical Guide to Information Security Testing and Assessment" with all tests and actions being conducted under controlled conditions.

All activities regarding the infrastructure assessment were conducted according to Google Cloud Platform (GCP) security best practices to ensure that necessary security controls are integrated into the design and implementation of a project.

Plus, the assessment can check and evaluate security configurations that should ensure the confidentiality, integrity and availability of ***** sensitive data and other resources.

Summary of Results

An initial reconnaissance of the ***** infrastructure and service of settings that need attention was completed. The results provided us with a highlighted list of specific settings in the infrastructure. An examination of the Google Cloud infrastructure revealed 2 **High**-level and 526 **Warning**-level issues within 2 projects (35 in total). After using a custom "gray box" technique on the ***** infrastructure, we were able to find a variety of issues according to Google's security checklist. High-level and some warning-level issues were additionally checked with custom scripts and techniques, using tools such as Burp Suite or Metasploit. There were no critical results reported, but it's good practice to perform these complementary tests. Uncovering the passwords via brute-force was not completed by using basic techniques. Cloud penetration testing (using simulated cyberattacks against targeted systems to identify vulnerabilities) is performed on cloud-native systems. This form of security testing is used to identify security risks and vulnerabilities and to provide actionable remediation advice.

Initial reconnaissance of the ***** network resulted in the discovery of a user enumeration vulnerability that allows an attacker to enumerate registered emails that exist in application. With Google Captcha bypass vulnerability, there is a possibility to brute force users' passwords and get access to users' accounts. While using provided credentials of the user company member, an IDOR vulnerability was discovered that allows this user to change the company name, company avatar and stored XSS vulnerability. Additionally, 2 vulnerabilities were found regarding API with **Critical**-risk and **High**-risk ratings. Other vulnerabilities

had **Low** and informational risk ratings but were still considerable enough to be restored.

CONFIDENTIAL

1. INFRASTRUCTURE ASSESSMENT

1.1. Discovery

For the purposes of this assessment, CLIENT provided their cloud account with view permission, suitable for the "gray box" penetration test. During the enumeration stage 2 projects were identified that need attention (Figure 1).

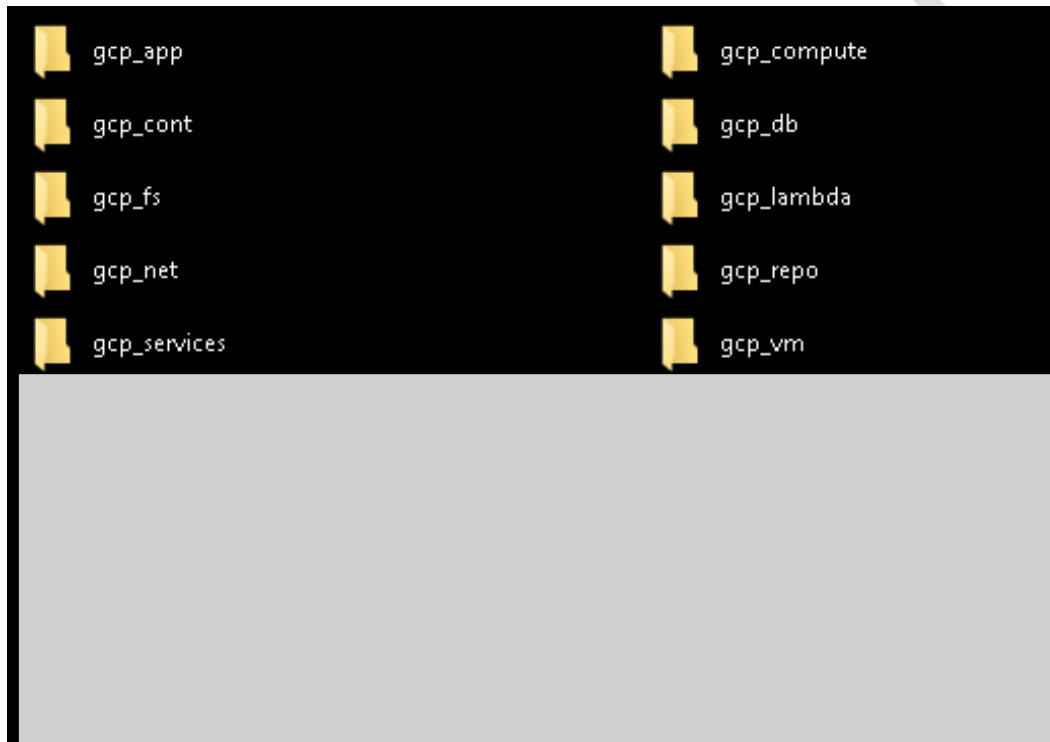


Figure 1 – Discovery process result files

1.2. Cloud SQL testing results

1. **HIGH - Cloud SQL Database Instances Have Public IPs**

Description - To lower the organization's attack surface, Cloud SQL databases should not have public IPs. Private IPs provide improved network security and lower latency for your application.

Remediation - From console:

1. Go to the Cloud SQL Instances page in the Google Cloud Console by visiting <https://console.cloud.google.com/sql/instances>.
2. Click the instance name to open its instance details page.
3. Select the Connections tab.
4. Deselect the public IP checkbox.
5. Click save to update the instance.

Compliance –

CIS Google Cloud Platform Foundations version 1.3.0

Affected Projects –

Databases –

Gathered information sample:

```
Project ID: [REDACTED]
Automatic Backups: Enabled
Last Backup: Invalid date format
Logs: Unknown
SSL Required: Disabled
Public IP Address: [REDACTED]
Private IP Address: None
Local Infile Flag is Off: true
Cross db Ownership Chaining Flag is Off: None
Contained Database Authentication Flag is Off: None
Log Checkpoints Flag is On: false
Log Connections Flag is On: false
Log Disconnections Flag is On: false
Log Lock Waits Flag is On: false
Log Min Messages Flag set Appropriately: false
Log Temp Files Flag set to 0: false
Log Min Duration Statement Flag set to -1: false
Authorized Networks: None
Users:
[REDACTED]
```

2. **WARNING** - Instance Not Requiring SSL for Incoming Connections

Description - SQL database connections if successfully trapped (MITM) can reveal sensitive data like credentials, database queries or query outputs. For security, it is always recommended that SSL encryption is used when connecting to your instance.

Compliance –

CIS Google Cloud Platform Foundations version 1.3.0

References –

<https://cloud.google.com/sql/docs/postgres/configure-ssl-instance>

3. **WARNING - Instance with Binary Logging Disabled**

Description - The benefits of enabling binary logs (replication, scalability, auditability, point-in-time data recovery, etc.) can improve the security posture of the Cloud SQL instance.

References -

<https://cloud.google.com/sql/docs/mysql/instance-settings>

<https://cloud.google.com/sql/docs/mysql/replication/tips>

4. **WARNING - Log Checkpoints Database Flag for PostgreSQL Instance Is Off**

Description - Enabling log_checkpoints cause checkpoints and restart points to be logged in the server log. Some statistics are included in the log messages, including the number of buffers written and the time spent writing them. This parameter can only be set in the postgresql.conf file or on the server command line. This recommendation is applicable to PostgreSQL database instances.

Compliance -

CIS Google Cloud Platform Foundations version 1.3.0

References -

<https://www.postgresql.org/docs/13/runtime-config-logging.html>

https://cloud.google.com/sql/docs/postgres/flags#setting_a_database_flag

5. **WARNING - Log Connections Database Flag for PostgreSQL Instance Is Off**

Description - PostgreSQL does not log attempted connections by default. Enabling the log_connections setting will create log entries for each attempted connection as well as successful completion of client authentication which can be useful in troubleshooting issues and to determine any unusual connection attempts to the server. This recommendation is applicable to PostgreSQL database instances.

Compliance -

CIS Google Cloud Platform Foundations version 1.3.0

References -

<https://www.postgresql.org/docs/13/runtime-config-logging.html>

<https://cloud.google.com/sql/docs/postgres/flags>

6. **WARNING - Log Disconnections Database Flag for PostgreSQL Instance Is Off**

Description - PostgreSQL does not log session details such as duration and session end by default. Enabling the log_disconnections setting will create log entries at the end of each session which can be useful in troubleshooting issues and determining any unusual activity across a time period. The log_disconnections and log_connections work hand in hand and generally the pair would be enabled/disabled together. This recommendation is applicable to PostgreSQL database instances.

Compliance –

CIS Google Cloud Platform Foundations version 1.3.0

References –

<https://www.postgresql.org/docs/13/runtime-config-logging.html>

<https://cloud.google.com/sql/docs/postgres/flags>

7. **WARNING - Log Lock Waits Database Flag for PostgreSQL Instance Is Off**

Description - The deadlock timeout defines the time to wait on a lock before checking for any conditions. Frequent runovers on deadlock timeout can be an indication of an underlying issue. This type of logging can enable the log_lock_waits flag. This can be used to identify poor performance due to locking delays or if a specially-crafted SQL is attempting to starve resources through holding locks for excessive amounts of time. This recommendation is applicable to PostgreSQL database instances.

Compliance –

CIS Google Cloud Platform Foundations version 1.3.0

References –

<https://www.postgresql.org/docs/13/runtime-config-logging.html>

<https://cloud.google.com/sql/docs/postgres/flags>

8. **WARNING - Log Min Duration Statement Database Flag for PostgreSQL Instance Is Not Set To - 1**

Description - Logging SQL statements may include sensitive information which should not be recorded in logs. This recommendation is applicable to PostgreSQL database instances.

Compliance – CIS Google Cloud Platform Foundations version 1.3.0

References –

<https://www.postgresql.org/docs/13/runtime-config-logging.html>

<https://cloud.google.com/sql/docs/postgres/flags>

9. **WARNING - Log Min Messages Database Flag for PostgreSQL Instance Is Not Set**

Description - Auditing helps in troubleshooting operational problems and also permits forensic analysis. If log_min_error_statement is not set to the correct value, messages may not be appropriately classified as error messages. Considering general log messages as error messages would make it difficult to find actual errors. Considering only stricter severity levels as error messages may miss actual errors and not log their SQL statements. The log_min_error_statement flag should be set in accordance with the organization's logging policy. This recommendation is applicable to PostgreSQL database instances.

Compliance – CIS Google Cloud Platform Foundations version 1.3.0

References –

<https://www.postgresql.org/docs/13/runtime-config-logging.html>

<https://cloud.google.com/sql/docs/postgres/flags>

10. **WARNING - Log Temp Files Database Flag for PostgreSQL Inst. Is Not Set To 0**

Description – If all temporary files are not logged, it may be more difficult to identify potential performance issues which may be due to either poor application coding or deliberate resource starvation attempts.

Compliance – CIS Google Cloud Platform Foundations version 1.3.0

References –

<https://www.postgresql.org/docs/13/runtime-config-logging.html>

<https://cloud.google.com/sql/docs/postgres/flags>

NOTE: All 3 SQL Instances - ***** have same issues

1.3. Cloud Storage testing results

1. **WARNING - Bucket with Logging Disabled**

Description – Enable access and storage logs in order to capture all events which may affect objects within target buckets.

Compliance – CIS Google Cloud Platform Foundations version 1.0.0, reference 5.3

References –

<https://cloud.google.com/storage/docs/access-logs>

Buckets affected –

2. **WARNING - Bucket with Versioning Disabled**

Description – Enable Object Versioning to protect Cloud Storage data from being overwritten or accidentally deleted.

References –

<https://cloud.google.com/storage/docs/using-object-versioning>

Buckets affected –

3. **WARNING - Uniform Bucket-Level Access Is Disabled**

Description – Uniform bucket-level access is recommended to unify and simplify how access is granted to your Cloud Storage resources. In order to support a uniform permission system, Cloud Storage has uniform bucket-level access. Using this feature disables ACLs for all Cloud Storage resources: access is then granted exclusively through Cloud IAM. Enabling uniform bucket-level access guarantees that if a Cloud Storage bucket is not publicly accessible, then no object in the bucket is publicly accessible.

Compliance – CIS Google Cloud Platform Foundations version 1.1.0, reference 5.2

References –

<https://cloud.google.com/storage/docs/uniform-bucket-level-access>

<https://cloud.google.com/storage/docs/using-uniform-bucket-level-access>

<https://cloud.google.com/storage/docs/org-policy-constraints#uniform-bucket>

[bucket](#)

Buckets affected –

1.4. Cloud Compute Engine testing results

1. **WARNING - Block Project SSH Keys Disabled**

Description – Project-wide SSH keys are stored in Compute/project metadata. Project wide SSH keys can be used to login into all the instances within a project. Using project-wide SSH keys ease the SSH key management; but if compromised, these pose a security risk which can impact all the instances within the project.

Compliance –

CIS Google Cloud Platform Foundations version 1.1.0, reference 4.3

References –

<https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys>

Instances affected – ALL

2. **WARNING - Default Firewall Rule in Use**

Description – Some default firewall rules were in use. This could potentially expose sensitive services or protocols to other networks.

Rules –

Example figure:



3. **WARNING - Default Network should be removed**

Description – The default network has a preconfigured network configuration and automatically generates insecure firewall rules. These automatically created firewall rules do not get audit logged and cannot be configured to enable firewall rule logging.

Compliance –

CIS Google Cloud Platform Foundations version 1.1.0, reference 3.1

References –

https://cloud.google.com/compute/docs/networking#firewall_rules

<https://cloud.google.com/compute/docs/reference/latest/networks/insert>

<https://cloud.google.com/compute/docs/reference/latest/networks/delete>

<https://cloud.google.com/vpc/docs/firewall-rules-logging>

<https://cloud.google.com/vpc/docs/vpc#default-network>

<https://cloud.google.com/sdk/gcloud/reference/compute/networks/delete>

4. **WARNING - Firewall INGRESS Rule Allows Public Access (0.0.0.0/0) to a Sensitive Port**

Description – The firewall rule was found to be exposing a well-known port to all source addresses. Well-known ports are commonly probed by automated scanning tools and may be an indicator of sensitive services exposed to the internet. If such services need to be exposed, a restriction on the source address could help to reduce the attack surface of the infrastructure.

Firewall Elements:

5. **WARNING - Firewall Rule Allows Internal Traffic**

Description – Firewall rule allows ingress connections for all protocols and ports among instances in the network.

Firewall Elements:

6. **WARNING - Firewall Rule Allows Port Range(s)**

Description – It was found that the firewall rule was using port ranges. Sometimes ranges could include unintended ports that should not be exposed. As a result, when possible, explicit port lists should be used instead.

Firewall Elements:

7. **WARNING - Firewall Rule Allows Public Access (0.0.0.0/0)**

Description – The firewall rule was found to be exposing potentially open ports to all source addresses. Ports are commonly probed by automated scanning tools and may be an indicator of sensitive services exposed to the internet. If these services need to be exposed, a restriction on the source address could help to reduce the attack surface of the infrastructure.


Firewall Elements:

8. **WARNING - Firewall Rule Opens All Ports (0-65535)**

Description - The firewall rule allows access to all ports. This widens the attack surface of the infrastructure and makes it easier for an attacker to reach potentially sensitive services over the network.

Firewall Elements:

Example figure:



9. **WARNING - Instance Disk without Snapshots**

Description – You should have snapshots of your in-use or available disks taken on a regular basis to enable disaster recovery efforts.

References –

<https://cloud.google.com/compute/docs/disks/create-snapshots>

<https://cloud.google.com/compute/docs/disks/scheduled-snapshots>

<https://cloud.google.com/compute/docs/disks/snapshot-best-practices>

Affected Instances:

10. WARNING - Instance without Deletion Protection

Description – It is good practice to enable this feature on production instances, to ensure that they may not be deleted by accident.

References –

<https://cloud.google.com/compute/docs/instances/preventing-accidental-vm-deletion>

Affected Instances: ALL

11. WARNING - Instances Configured to Use Default Service Account

Description - The default Compute Engine service account has the editor role on the project which allows read and write access to most Google Cloud services. To defend against privilege escalations if your VM is compromised and to prevent an attacker from gaining access to all of your projects, use of the default Compute Engine service account is not recommended. Instead, you should create a new service account and assign the permissions only required by your instance.

Compliance – CIS Google Cloud Platform Foundations version 1.1.0, reference 4.1

References –

<https://cloud.google.com/compute/docs/access/service-accounts>

<https://cloud.google.com/compute/docs/access/create-enable-service-accounts-for-instances>

<https://cloud.google.com/sdk/gcloud/reference/compute/instances/set-service-account>

12. WARNING Instances Have Public IP Addresses

Description – To reduce your attack surface, Compute Engine instances should not have public IP addresses. Instead, instances should be configured behind load balancers to minimize the instance's exposure to the internet.

Compliance –

CIS Google Cloud Platform Foundations version 1.1.0, reference 4.9

References –

https://cloud.google.com/load-balancing/docs/backend-service#backends_and_external_ip_addresses

<https://cloud.google.com/compute/docs/instances/connecting-advanced#sshbetweeninstances>

<https://cloud.google.com/compute/docs/instances/connecting-to-instance>

https://cloud.google.com/compute/docs/ip-addresses/reserve-static-external-ip-address#unassign_ip

<https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints>

13. **WARNING - Network without Instances**

Description – Maintaining unused resources increases the risk of misconfigurations and audit difficulty.

Affected Instances:

14. **WARNING - OS login Disabled**

Description – Enabling OS Login ensures that SSH keys used to connect to instances are mapped with IAM users. Revoking access to an IAM user will revoke all the SSH keys associated with that particular user. It facilitates centralized and automated SSH key pair management which is useful in handling cases like response to compromised SSH key pairs and/or revocation of external/third-party/vendor users.

Compliance –

CIS Google Cloud Platform Foundations version 1.1.0, reference 4.4

References –

<https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys>

Affected Instances: ALL

15. **WARNING - Shielded VM Disabled**

Description – Shielded VM offers verifiable integrity of your Compute Engine VM instances so you can be confident instances haven't been compromised by boot- or kernel-level malware or rootkits. Shielded VM's verifiable integrity is achieved through the use of Secure Boot, a virtual trusted platform module (vTPM) enabled Measured Boot and integrity monitoring.

Compliance –

CIS Google Cloud Platform Foundations version 1.1.0, reference 4.8

References –

<https://cloud.google.com/compute/docs/instances/modifying-shielded-vm>

<https://cloud.google.com/shielded-vm>

<https://cloud.google.com/security/shielded-cloud/shielded-vm#organization-policy-constraint>

Affected Instances: ALL

16. **WARNING - VM Disks Not Customer-Supplied Encryption Keys (CSEK) Encrypted**

Description - By default, Google Compute Engine encrypts all data at rest. Compute Engine handles and manages this encryption for you without any additional actions on your part. However, if you want to control and manage this encryption yourself, you can provide your own encryption keys.

Compliance –

CIS Google Cloud Platform Foundations version 1.1.0, reference 4.7

References –

https://cloud.google.com/compute/docs/disks/customer-supplied-encryption#encrypt_a_new_persistent_disk_with_your_own_keys

<https://cloud.google.com/compute/docs/reference/rest/v1/disks/get>

https://cloud.google.com/compute/docs/disks/customer-supplied-encryption#key_file

Affected Instances: ALL

1.5. IAM testing results

1. **WARNING - Basic Role in Use**

Description – Basic roles grant significant privileges. In most cases, usage of these roles is not recommended and does not follow security best practice.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 1.4

CIS Google Cloud Platform Foundations version 1.1.0, reference 1.5

References –

<https://cloud.google.com/sdk/gcloud/reference/iam/service-accounts/>

<https://cloud.google.com/iam/docs/understanding-roles>

<https://cloud.google.com/iam/docs/understanding-service-accounts>

Affected Roles:

2. **WARNING - Gmail Account in Use**

Description – It is recommended fully-managed corporate Google accounts be used for increased visibility, auditing and for controlling access to Cloud Platform resources. Email accounts based outside of the user's organization, such as personal accounts, should not be used for business purposes.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 1.1

CIS Google Cloud Platform Foundations version 1.1.0, reference 1.1

References –

<https://cloud.google.com/docs/enterprise/best-practices-for-enterprise-organizations#manage-identities>

<https://support.google.com/work/android/answer/6371476>

<https://cloud.google.com/sdk/gcloud/reference/organizations/get-iam-policy>

<https://cloud.google.com/sdk/gcloud/reference/beta/resource-manager/folders/get-iam-policy>

<https://cloud.google.com/sdk/gcloud/reference/projects/get-iam-policy>

<https://cloud.google.com/resource-manager/docs/organization-policy/org-policy-constraints>

<https://cloud.google.com/resource-manager/docs/organization-policy/restricting-domains>

Affected Roles:

3. WARNING - IAM Role Assigned to User

Description - Best practice would be to grant roles to a Google suite group instead of individual users where possible. It is easier to add and remove members from a group instead of updating a Cloud IAM policy to perform these actions.

References –

<https://cloud.google.com/iam/docs/understanding-roles>

<https://cloud.google.com/iam/docs/using-iam-securely>

Bindings affected:

4. **WARNING - Lack of Service Account Key Rotation**

Description – Rotating service account keys will reduce the window of opportunity for an access key that is associated with a compromised or terminated account to be used. Service account keys should be rotated to ensure that data cannot be accessed with an old key which might have been lost, cracked or stolen keys should be rotated every 90 days.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 1.6

CIS Google Cloud Platform Foundations version 1.1.0, reference 1.7

References –

https://cloud.google.com/iam/docs/understanding-service-accounts#managing_service_account_keys

<https://cloud.google.com/sdk/gcloud/reference/iam/service-accounts/keys/list>

<https://cloud.google.com/iam/docs/service-accounts>

Affected Accounts:

5. WARNING - Service Account with Admin Privileges

Description – Service accounts represent service-level security of the resources (application or a VM) which can be determined by the roles assigned to it. Enrolling service accounts with administrative privileges grants full access to an assigned application or a VM. A service account access holder can be a user.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 1.4

CIS Google Cloud Platform Foundations version 1.1.0, reference 1.5

References –

<https://cloud.google.com/sdk/gcloud/reference/iam/service-accounts/>

<https://cloud.google.com/iam/docs/understanding-roles>

<https://cloud.google.com/iam/docs/understanding-service-accounts>

Affected Accounts:

6. WARNING - User with Privileged Service Account Roles at the Project Level

Description – Granting the iam.serviceAccountUser, iam.serviceAccountTokenCreator, or iam.serviceAccountActor role to a user for a project gives the user access to all service accounts in the project, including service accounts that may be created in the future. This can result in an elevation of privileges by using service accounts and corresponding Compute Engine instances.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 1.5

CIS Google Cloud Platform Foundations version 1.1.0, reference 1.6

References –

<https://cloud.google.com/iam/docs/service-accounts>

<https://cloud.google.com/iam/docs/granting-changing-revoking-access>

<https://cloud.google.com/iam/docs/understanding-roles>

<https://cloud.google.com/iam/docs/granting-changing-revoking-access>

<https://console.cloud.google.com/iam-admin/iam>

Affected Bindings:

7. WARNING - User-Managed Service Account Keys

Description – The use of user-managed service account keys is not recommended as anyone who has access to the keys will be able to access resources through the service account. Best practice would be to use GCP-managed keys, which are utilized by Cloud Platform services such as App Engine and Compute Engine. These keys cannot be downloaded. Google will keep the keys and automatically rotate them on an approximate weekly basis.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 1.3

CIS Google Cloud Platform Foundations version 1.1.0, reference 1.4

References –

https://cloud.google.com/iam/docs/understanding-service-accounts#managing_service_account_keys

<https://cloud.google.com/resource-manager/docs/organization-policy/restricting-service-accounts>

Affected Service Accounts -

1.6. Kubernetes Engine testing results

1. **WARNING - Clusters Lacking Labels**

Description – Labels enable users to map their own organizational structures onto system objects in a loosely coupled fashion without requiring clients to store these mappings. Labels can also be used to apply specific security settings and auto configure objects at creation.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 7.5

References –

https://cloud.google.com/kubernetes-engine/docs/how-to/hardening-your-cluster#use_namespaces_and_rbac_to_restrict_access_to_cluster_resources

Affected Clusters:

2. **WARNING - Default Service Account in Use**

Description – You should create and use a minimally privileged service account to run your Kubernetes Engine cluster instead of using the Compute Engine default service account.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 7.17

CIS GKE Benchmark version 1.0.0, reference 6.2.1

References -

<https://www.cisecurity.org/benchmark/kubernetes/>

https://cloud.google.com/kubernetes-engine/docs/how-to/hardening-your-cluster#use_least_privilege_sa

https://cloud.google.com/kubernetes-engine/docs/concepts/cis-benchmarks#default_values_on

Affected Clusters:

3. WARNING - Lack of Access Scope Limitation 2

Description – If you are not creating a separate service account for a node, you should limit the scope of the node service account to reduce the possibility of privilege escalation in an attack. This ensures that your default service account does not have permissions beyond those necessary to run your cluster. While the default scopes are limited, they may include scopes beyond the minimally required scopes needed to run a cluster. If you are accessing private images in Google Container Registry, the minimally required scopes are only logging.write, monitoring and devstorage.read_only.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 7.18

References –

<https://cloud.google.com/kubernetes-engine/docs/how-to/access-scopes>

Affected Clusters:

4. WARNING - Master Authorized Networks Disabled

Description – Master authorized networks block untrusted IP addresses from outside the Google Cloud Platform. Addresses from inside GCP can still reach your master network through HTTPS, provided that they have the necessary Kubernetes credentials.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 7.4

CIS GKE Benchmark version 1.0.0, reference 6.6.3

References –

<https://www.cisecurity.org/benchmark/kubernetes/>

<https://cloud.google.com/kubernetes-engine/docs/how-to/authorized-networks>

https://cloud.google.com/kubernetes-engine/docs/how-to/hardening-your-cluster#restrict_network_access_to_the_control_plane_and_nodes

https://cloud.google.com/kubernetes-engine/docs/concepts/cis-benchmarks#default_values_on

Affected Clusters:

5. WARNING - Network Policy Disabled

Description – By default, pods are non-isolated; they accept traffic from any source. Pods become isolated by having a network policy that selects them. Once there is a network policy in a namespace selecting a particular pod, that pod will reject any connections that are not allowed by the network policy.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 7.11

CIS GKE Benchmark version 1.0.0, reference 6.6.7

References –

<https://www.cisecurity.org/benchmark/kubernetes/>

https://cloud.google.com/kubernetes-engine/docs/how-to/hardening-your-cluster#restrict_with_network_policy

https://cloud.google.com/kubernetes-engine/docs/concepts/security-overview#network_security

https://cloud.google.com/kubernetes-engine/docs/concepts/cis-benchmarks#default_values_on

Affected Clusters:

6. WARNING - Pod Security Policy Disabled

Description – A pod security policy is a cluster-level resource that controls security sensitive aspects of the pod specification. The pod security policy objects define a set of conditions that a pod must run in order to be accepted into the system, as well as defaults for the related fields.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 7.14

CIS GKE Benchmark version 1.0.0, reference 6.10.3

References –

<https://www.cisecurity.org/benchmark/kubernetes/>

<https://cloud.google.com/kubernetes-engine/docs/how-to/pod-security-policies>

<https://kubernetes.io/docs/concepts/policy/pod-security-policy>

https://cloud.google.com/kubernetes-engine/docs/concepts/cis-benchmarks#default_values_on

Affected Clusters:

7. WARNING - Private Cluster Disabled

Description – A private cluster is a cluster that makes your master accessible from the public internet. In a private cluster, nodes do not have public IP addresses, so your workloads run in an environment that is rated from the internet. Nodes have been addressed only in the private RFC address space. Nodes and masters communicate with each other privately using VPC peering.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 7.15

CIS GKE Benchmark version 1.0.0, reference 6.6.4

CIS GKE Benchmark version 1.0.0, reference 6.6.5

References –

<https://www.cisecurity.org/benchmark/kubernetes/>

https://cloud.google.com/kubernetes-engine/docs/how-to/hardening-your-cluster#restrict_network_access_to_the_control_plane_and_nodes

https://cloud.google.com/kubernetes-engine/docs/concepts/cis-benchmarks#default_values_on

Affected Clusters:

8. WARNING - Private Google Access Disabled

Description – Enabling private Google access allows VMs on a subnetwork to use a private IP address to reach Google APIs rather than an external IP address.

Compliance – CIS Google Cloud Platform Foundations version 1.0.0, reference 7.16

References –

https://cloud.google.com/kubernetes-engine/docs/how-to/hardening-your-cluster#restrict_network_access_to_the_control_plane_and_nodes

Affected Clusters:

9. WARNING - Nodes Auto-Upgrade Disabled

Description – Auto-upgrades automatically ensure that security updates are applied and are current.

Compliance –

CIS Google Cloud Platform Foundations version 1.0.0, reference 7.8

CIS GKE Benchmark version 1.0.0, reference 6.5.3

References –

<https://www.cisecurity.org/benchmark/kubernetes/>

<https://cloud.google.com/kubernetes-engine/docs/how-to/node-auto-upgrades>

https://cloud.google.com/kubernetes-engine/docs/concepts/cis-benchmarks#default_values_on

Affected Clusters:

1.7. Stackdriver Logging & Monitoring testing results

1. **WARNING - Log Metric Filter Issues**

Log Metric Filter Doesn't Exist for Audit Configuration Changes

Description - Configuring the metric filter and alerts for audit configuration changes ensures the recommended state of audit configuration is maintained so that all activities in the project are auditable at any point in time.

Log Metric Filter Doesn't Exist for Cloud Storage IAM Permission Changes

Description - Monitoring changes to Cloud Storage bucket permissions may reduce the time needed to detect and correct permissions on sensitive Cloud Storage buckets and objects inside the bucket.

Log Metric Filter Doesn't Exist for Custom Role Changes

Description - Google Cloud IAM provides predefined roles that give granular access to specific Google Cloud Platform resources and prevent unwanted access to other resources. However, to cater to organization-specific needs, Cloud IAM also provides the ability to create custom roles. Project owners and administrators who are an organization role administrator or IAM role administrator can create custom roles. Monitoring role creation, deletion and updating activities will help in identifying any over-privileged role in the early stages.

Log Metric Filter Doesn't Exist for Project Ownership Assignments/Changes

Description - Project ownership has the highest level of privileges on a project. To avoid misuse of project resources, the project ownership assignment/change actions mentioned above should be monitored and highlighted for concerned recipients.

Log Metric Filter Doesn't Exist for SQL Instance Configuration Changes

Description - Monitoring changes to SQL instance configuration may reduce the time needed to detect and correct misconfigurations done on the SQL server.

Log Metric Filter Doesn't Exist for VPC Network Changes

Description - It is possible to have more than one VPC within a project. In addition, it is also possible to create a peer connection between two VPCs enabling network traffic to route between VPCs. Monitoring changes to a VPC will help ensure VPC traffic flow is not being impacted.

Log Metric Filter Doesn't Exist for VPC Network Firewall Rule Changes

Description - Monitoring for create or update firewall rule events gives insight into network access changes and may reduce the time it takes to detect suspicious activity.

Log Metric Filter Doesn't Exist for VPC Network Route Changes

Description - Google Cloud Platform (GCP) routes define the path that network traffic takes from a VM instance to another destination. The other destination can be inside the organization VPC network (such as another VM) or outside of it. Every route consists of a destination and the next hop. Traffic whose destination IP is within the destination range is sent to the next hop for delivery. Monitoring changes to route tables will help ensure that all VPC traffic flows through an expected path.

Compliance -

CIS Google Cloud Platform Foundations version 1.1.0

References -

<https://cloud.google.com/logging/docs/logs-based-metrics/>

<https://cloud.google.com/monitoring/custom-metrics/>

<https://cloud.google.com/monitoring/alerts/>

<https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

<https://cloud.google.com/logging/docs/audit/configure-data-access#getiampolicy-setiampolicy>

Affected Logging Configurations:

2. WARNING - Alerts Setup Issues

Alerts Don't Exist for Audit Configuration Changes

Description - Configuring the metric filter and alerts for audit configuration changes ensures the recommended state of audit configuration is maintained so that all activities in the project are auditable at any point in time.

Alerts Don't Exist for Cloud Storage IAM Permission Changes

Description - Monitoring changes to Cloud Storage bucket permissions may reduce the time needed to detect and correct permissions on sensitive Cloud Storage buckets and objects inside the bucket.

Alerts Doesn't Exist for Custom Role Changes

Description - Google Cloud IAM provides predefined roles that give granular access to specific Google Cloud Platform resources and prevent unwanted access to other resources. However, to cater to organization-specific needs, Cloud IAM also provides the ability to create custom roles. Project owners and administrators who are an organization role administrator IAM role administrator can create custom roles. Monitoring role creation, deletion and updating activities will help in identifying an over-privileged role at early stages.

Alerts Don't Exist for Project Ownership Assignments/Changes

Description - Project ownership has the highest level of privileges on a project. To avoid misuse of project resources, the project ownership assignment/change actions mentioned above should be monitored and alerted to concerned recipients.

Alerts Don't Exist for SQL Instance Configuration Changes

Description - Monitoring changes to SQL instance configuration changes may reduce the time needed to detect and correct configurations on the SQL server.

Alerts Don't Exist for VPC Network Changes

Description - It is possible to have more than one VPC within a project. In addition, it is also possible to create a peer connection between two VPCs enabling network traffic to route between VPCs. Monitoring changes to a VPC will help ensure VPC traffic flow is not being impacted.

Alerts Don't Exist for VPC Network Firewall Rule Changes

Description - Monitoring for create or update firewall rule events gives insight into network access changes and may reduce the time it takes to detect suspicious activity

Alerts Don't Exist for VPC Network Route Changes

Description - Google Cloud Platform routes define the paths network traffic takes from a VM instance to another destination. The other destination can be inside the organization VPC network (such as another VM) or outside of it. Every route consists of a destination and a next hop. Traffic whose destination IP is within the destination range is sent to the next hop for delivery. Monitoring changes to route tables will help ensure that all VPC traffic flows through an expected path.

Compliance –

CIS Google Cloud Platform Foundations version 1.1.0

References –

<https://cloud.google.com/logging/docs/logs-based-metrics/>

<https://cloud.google.com/monitoring/custom-metrics/>

<https://cloud.google.com/monitoring/alerts/>

<https://cloud.google.com/logging/docs/reference/tools/gcloud-logging>

<https://cloud.google.com/storage/docs/access-control/iam>

Affected Logging Configurations:

1.8. Conclusion


An examination of the Google Cloud infrastructure revealed 2 **high**-level and 526 **warning**-level issues within 2 projects (35 total). After using a custom "Gray Box" technique on the ***** infrastructure we were able to find a list of issues according to Google's security checklist. High-level and some warning-level issues were additionally checked with custom scripts and techniques and with a set of tools like Burp Suite or Metasploit. An essential test, but with no critical results at this point.

CONFIDENTIAL

2. APPLICATION ASSESSMENT

2.1. Introduction

The cybersecurity team performed a penetration test on the ***** application using a black box and white box approach, simulating attack vectors that intruders could perform in real life. When the web application was tested, 7 security vulnerabilities were found and rated for their critical, high, medium and low level.



2.2. User enumeration


Category: OWASP Top 10 (A07:2021-Identification and Authentication Failures)

Severity: Low

Vulnerability explanation:

User enumeration is when a malicious actor can use brute-force techniques to either guess or confirm valid users in a system. In the login page of the ***** platform, if one of the emails or passwords is wrong, in both cases the "wrong email or password" message is displayed to the user. Thus, it is not clear which(email or password) is incorrect.

For example, in the screenshot below, although the email address cybersec@example.com is registered, the response message does not show any information about it.



However, this mechanism was not implemented on the "password reset" and "registration" pages. Therefore, an attacker can find out if an email address is registered on the Proto platform.

Exploitation process:

To verify the vulnerability, we will use one existing(cybersec@example.com) and one non-existing(not-registered@gmail.com) email address.

Steps to Reproduce:

1. In the password reset page, when the user enters an existing email address, the following message is displayed:



But, with the non-existent email, the user will see the following message:

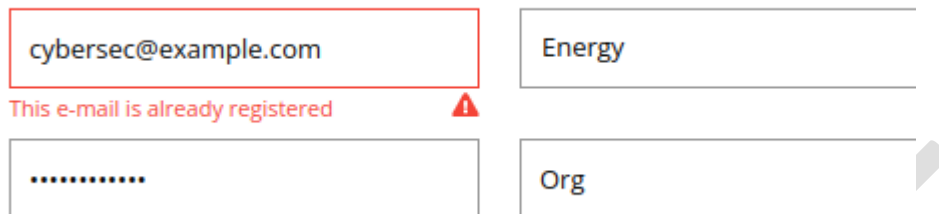
Enter your email below, and we'll send a link
to reset your password.

sdjfkndskndfkg@sdfdd.dfdg

Account with this email doesn't exist.



1. In the registration page, if the user tries to register with an existing email address, the following error message will be displayed:



The screenshot shows a registration form with four input fields arranged in a 2x2 grid. The top-left field contains the email address 'cybersec@example.com' and is highlighted with a red border. Below this field, the error message 'This e-mail is already registered' is displayed in red text, accompanied by a small red warning triangle icon. The top-right field contains the text 'Energy'. The bottom-left field contains a series of dots, indicating a masked password. The bottom-right field contains the text 'Org'.

But, with non-existent one, a new account will be created.

Based on these two response messages, it's possible to determine whether an email address is registered.

Remediation:

The same response should be returned whether the email address entered by the user exists or not.

2.3. Google Captcha bypass

Category: A2:2017 – Broken Authentication

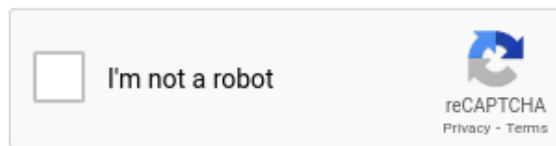
Severity: Medium

Vulnerability explanation:

There are 3 pages that uses reCAPTCHA; login, password reset and registration. While testing these functionalities, it turned out that the captcha provided by reCAPTCHA is not validated.

Exploitation process:

- In the login page, captcha is required:



Please verify you are a human



- When the user confirms the captcha and clicks to "login", a request is sent as follows:



- The problem here is, if we re-use a captcha token, or even remove it completely, the server doesn't show any error message and the request is accepted.



- Further analysis revealed that there is no rate-limiting here either. This means, an attacker can try as many passwords as they want until the correct one is found.
- The screenshot below proves that despite checking 50 passwords, no protection mechanism is triggered:

37	drtjdryj	403	<input type="checkbox"/>	<input type="checkbox"/>	357
38	jtyjtyjtjfy	403	<input type="checkbox"/>	<input type="checkbox"/>	357
39	yjtyjtyhty	403	<input type="checkbox"/>	<input type="checkbox"/>	357
40	jy	403	<input type="checkbox"/>	<input type="checkbox"/>	357
41	djtyj	403	<input type="checkbox"/>	<input type="checkbox"/>	357
42	tyjj	403	<input type="checkbox"/>	<input type="checkbox"/>	357
43		403	<input type="checkbox"/>	<input type="checkbox"/>	357
44	j	403	<input type="checkbox"/>	<input type="checkbox"/>	357
45	j	403	<input type="checkbox"/>	<input type="checkbox"/>	357
46	tyj	403	<input type="checkbox"/>	<input type="checkbox"/>	357
47	jty	403	<input type="checkbox"/>	<input type="checkbox"/>	357
48	jty	403	<input type="checkbox"/>	<input type="checkbox"/>	357
49	demon444	200	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	407
50	sfsdfsd	403	<input type="checkbox"/>	<input type="checkbox"/>	357

Request
Response

Pretty
Raw
Hex
Render

In this simulation, the attacker tried 50 passwords and was able to find the correct one on the 49th request. In the response, the victim user's session token is sent to the attacker.

Note: The captcha doesn't work on any of the 3 pages named above. Each of them has its own impact. These are:

- Login – brute force
- Password reset – sending large volumes of email from proto's mail address.
- Registration – creating a large number of fake accounts.

Remediation:

There is a logical flaw in captcha implementation. Make sure that every request is checked for the correct captcha and is then processed.

CONFIDENTIAL

2.4. SSRF on JSON API functionality

Category: A10:2021 – SSRF

Severity: **Critical**

Vulnerability explanation:

SSRF flaws occur whenever a web application is fetching a remote resource without validating the user-supplied URL. It allows an attacker to coerce the application to send a crafted request to an unexpected destination, even when protected by a firewall, VPN or another type of network access control list (ACL).

According to the documentation, the "JSON:API" bot block provides the user with the ability to send a HTTP request to another (bot owner's) server.

But it's also possible to send HTTP requests to internal addresses that belong to the ***** environment and are not accessible by users.

While analyzing the application, it appears the environment is deployed on Google Cloud Platform. So, it automatically has access to the metadata server API **without any additional authorization**. As the requests are issued from the server (VM), an attacker can access the metadata server API and retrieve **confidential** data from VM metadata server.

Exploitation process:

As proof of the concept, we retrieve VM metadata from the Google Cloud VM metadata server.

1. An attacker creates a new "JSON:API" bot block.
2. As a URL, the address of the metadata server of the Google Cloud system is entered.

(<http://metadata.google.internal/computeMetadata/v1/?recursive=true>)

3. Regarding Google Cloud's documentation, in order to query metadata information, the Metadata-Flavor header must be in all requests. In traditional SSRF exploitation, it's not possible to add the HTTP header to requests. But, as the "JSON:API" block provides users with the ability to add headers to requests, it's not a problem for an attacker.



4. The Attacker saves the block and click the "test chat" button.
5. When the chat starts, the HTTP request is sent. It's possible to view webhook history in the ***** page.

In this page, we can see requests and responses that are sent via the JSON:API block.



In this response, confidential information such as Kubernetes environment variables, network interfaces, service accounts and SSH keys are leaked.

Example of leaked data:



As staging and development environments are accessible, we've tested this vulnerability on those environments. There are more local users that uses SSH, such as:

- *****
- *****
- *****

Remediation:

With this vulnerability, a developer team would need a unique solution. This is because the best practice of remediation can sometimes negatively impact the business process.

CONFIDENTIAL

2.5. XSS

Category: A03:2021 – Injection

Severity: High

Vulnerability explanation:

Unrestricted file upload leads stored-XSS.

Endpoint: *****

In the “create case” page, no validation is performed on the uploaded files. In that case, a user can upload an arbitrary file to the server. Then, in the preview page of the attachment, this file will be served.

Exploitation process:

As a PoC, we have uploaded the HTML file to the server, then executed it onto the victim's browser.

1. The attacker creates a new “JSON:API” bot block.



2. The file is uploaded to the server, using /case/attachment endpoint.




Figure 1

1. The malicious URL is delivered to the victim user. The URL is the value of "*****" parameter in the response (above screenshot).

URL: *****

When the victim clicks on the link, the following page will be displayed.



Figure 1

1. The malicious URL is delivered to the victim user. The URL is the value of "*****" parameter in the response (above screenshot).

URL: *****

When the victim clicks on the link, the following page will be displayed.

2. With the help of the 2nd line (Figure 1), the path section after the domain is deleted using the "pushState" method. This makes the URL even more realistic.
3. This behaviour can lead the following security issue:

An attacker copies the ***** login page, and modifies it to send user submitted data to his own server. Then, the HTML file is uploaded via the vulnerable endpoint and sends it to the user, convincing them that it is the login page. As the domain(*****) belongs to ***** , the victim has a high chance of being deceived.

Note: Since the session token is stored under ***** domain's local storage, this vulnerable endpoint has no access to it, so, it's not possible to escalate this to "account takeover".

Remediation:

https://cheatsheetseries.owasp.org/cheatsheets/Cross_Site_Scripting_Prevention_Cheat_Sheet.html

2.6. IDOR-Privilege Escalation

Category: A01:2021 – BROKEN ACCESS CONTROL

Severity: High



Vulnerability explanation:

An unprivileged user(Company.MEMBER) can change the name and avatar of the company (even if it doesn't belong to the company).

Exploitation process:

1. The attacker gets ***** of the target company. It can be retrieved with the help of "*****" endpoint.


Note: As the current user is "Company.MEMBER" on the company called "SecureComp" (*****), they cannot modify this company's settings.

- 
2. Then, attacker modifies settings of the company where he has "Company.ADMIN" permission(*****), and intercepts the request with proxy software.
- 

3. As seen in the above screenshot, company ID is sent via request. However, it is not checked if the user who sent the request has the correct privilege for that company. So, user can replace his own company's ID (*****) with the target company's ID(*****). A modified request is shown in the following image.



4. After sending the request, the target company's name is changed. To confirm this, we can use "*****" endpoint.




Source code analysis:

Source code of the vulnerable endpoint is shown in the following screenshot:



Privilege check process is implemented in the line 101 (*****).

The CREATE_COMPANY method is defined in the ***** file.



Unlike other methods, no verification is performed here. This method will return a true value in all cases.

Remediation:

The only real solution to this issue is to implement access control. The user needs to be authorized for the requested information before the server provides it.

2.7. Security misconfiguration – exposed test environment

Category: A6:2017 – Security Misconfiguration

Severity: Low

Vulnerability explanation:

When analyzing public resources belonging to *****, some resources that shouldn't be public were actually accessible by everyone.

Some of them are as follows:

- *****
- *****
- *****
- *****

Such test environments may contain source code of future features that are not yet meant to be publicly available. Such exposed test environments pose weak entry points into internal networks and can lead to data exposure and leaks. In addition to potential leaks, since most test environments are not regularly monitored, attackers apply their exploits on exposed staging environments until they are ready and able to take down the live(prod.) application in one attempt.

Remediation:

To remediate this issue, some access controls should be implemented. There are some choices, such as: implementing VPN, adding an additional security layer (login page, MFA) or any security measure to confirm whether the person who wants to access one of those resources has permission to enter the test environment.

2.8. Code Review

Category: Insecure Randomness

Severity: Low

Vulnerability explanation:

Standard pseudorandom number generators cannot withstand cryptographic attacks. A PRNG is an algorithm used to produce random-looking numbers with certain desirable statistical properties. In order for a PRNG to be cryptographically secure, it must be resistant to prediction.

Code Block:

- *****, line 480

```
[REDACTED]
```

- *****, line 119

```
[REDACTED]
```

- *****, line 44

```
[REDACTED]
```

Remediation:

We recommend using the secrets module's PRNG as follows:

<https://docs.python.org/dev/library/secrets.html#secrets.SystemRandom>

2.9. Advisory information

1. The old version of the software: Grafana v7.1.1


The older version of Grafana has multiple security vulnerabilities. We advise updating your system to a more up-to-date version.

2. ***** : Ngrok service

In the future, it may be beneficial to appraise this service. This is because developers can publish sensitive services or APIs with an Ngrok application.

3. Dangerous allowlist policy

The ***** application uses the Flask Jinja template engine. We noticed that when we started to build a bot, we were able to inject mathematical operations onto the template engine; however, developers use the ***** library for security. Therefore, we couldn't inject a malicious payload (such as a remote code execution payload) into the system.




Nevertheless, we could bypass some restrictions like a "double bracket":



The application deleted our double bracket payload.



An attacker can bypass that restriction using a "triple bracket".



In future, you can use complex restrictions on important functionality.

2.10. Risk Rating

During the test, our team found high and critical-level security vulnerabilities. The "*****" application has passed the penetration testing check with a 5/10 score. The overall risk identified to ***** as a result of the penetration test is **Medium**.

CONFIDENTIAL

Appendix A: Infrastructure Assessment Results

Cloud SQL test results

Name	Cloud SQL Database Instances Have Public IPs
Risk	HIGH
Description	To lower the organization's attack surface, cloud SQL databases should not have public IPs. Private IPs provide improved network security and lower latency for your application.
Affected Assets	<p>Projects:</p> <ul style="list-style-type: none"> ***** ***** <p>Databases:</p> <ul style="list-style-type: none"> ***** ***** *****
Recommendation	<ol style="list-style-type: none"> 1. Go to the Cloud SQL instances page in the Google Cloud Console. https://console.cloud.google.com/sql/instances. 2. Click the instance name to open its instance details page. 3. Select the connections tab. 4. Deselect the public IP checkbox. 5. Click save to update the instance.

Name	Instance Not Requiring SSL for Incoming Connections
Risk	WARNING
Description	SQL database connections if successfully trapped (MITM) can reveal sensitive data like credentials, database queries or query outputs. For security, it is recommended that SSL encryption is always used when connecting to your instance.

Name	Instance with Binary Logging Disabled
Risk	WARNING
Description	The benefits of enabling binary logs (replication, scalability, auditability, point-in-time and data recovery) can improve the security posture of the cloud SQL instance.

Name	Log Checkpoints Database Flag for PostgreSQL Instance Is Off
Risk	WARNING
Description	Enabling log_checkpoints causes checkpoints and restart points to be logged in the server log. Some statistics are included in the log messages, including the number of buffers written and the time spent writing them. This parameter can only be set in the postgresql.conf file or on the server command line. This recommendation is applicable to PostgreSQL database instances.

Name	Log Connections Database Flag for PostgreSQL Instance Is Off
Risk	WARNING
Description	PostgreSQL does not log attempted connections by default. Enabling the log_connections setting will create log entries for each attempted connection as well as successful completion of client authentication which can be useful in troubleshooting issues and to determine any unusual connection attempts to the server. This recommendation is applicable to PostgreSQL database instances.

Name	Log Disconnections Database Flag for PostgreSQL Instance Is Off
Risk	WARNING
Description	PostgreSQL does not log session details such as duration and session end by default. Enabling the log_disconnections setting will create log entries at the end of each session which can be useful in troubleshooting issues and determining any unusual activity across a time period. The log_disconnections and log_connections work hand in hand and generally the pair would be enabled/disabled together. This recommendation is applicable to PostgreSQL database instances.

Name	Log Lock Waits Database Flag for PostgreSQL Instance Is Off
Risk	WARNING
Description	The deadlock timeout defines the time to wait on a lock before checking for any conditions. Frequent run overs on deadlock timeout can be an indication of an underlying issue. Logging incidences by enabling the log_lock_waits flag can identify poor performance. These may be due to locking delays or if a specially-crafted SQL is attempting to starve resources through holding locks for excessive amounts of time. This recommendation is applicable to PostgreSQL database instances.

Name	Log Min Duration Statement Database Flag for PostgreSQL Instance Is Not Set To -1
Risk	WARNING
Description	Logging SQL statements may include sensitive information that should not be recorded in logs. This recommendation is applicable to PostgreSQL database instances.

Name	Log Min Messages Database Flag for PostgreSQL Instance Is Not Set
Risk	WARNING
Description	Auditing helps in troubleshooting operational problems and also permits forensic analysis. If the log_min_error_statement is not set to the correct value, messages may not be appropriately classified as error messages. Considering general log messages as error messages would make it difficult to find actual errors, while considering only stricter severity levels as error messages may skip actual errors to log their SQL statements. The log_min_error_statement flag should be set in accordance with the organization's logging policy. This recommendation is applicable to PostgreSQL database instances.

Name	Log Temp Files Database Flag for PostgreSQL Inst. Is Not Set To 0
Risk	WARNING
Description	If all temporary files are not logged, it may be more difficult to identify potential performance issues that are either due to poor application coding or deliberate resource starvation attempts.

Cloud Storage test results

Name	Bucket with Logging Disabled
Risk	WARNING
Description	Enable access and storage logs in order to capture all events which may affect objects within target buckets.
Affected Assets	Buckets: <ul style="list-style-type: none"> • ***** • ***** • ***** • ***** • ***** • ***** • ***** • ***** • *****

Name	Bucket with Versioning Disabled
Risk	WARNING
Description	Enable object versioning to protect Cloud Storage data from being overwritten or accidentally deleted.
Affected Assets	Buckets: <ul style="list-style-type: none"> • ***** • ***** • ***** • ***** • ***** • ***** • ***** • ***** • ***** • ***** • ***** • *****

Name	Uniform Bucket-Level Access Is Disabled
Risk	WARNING
Description	Uniform bucket-level access is recommended to unify and simplify how you grant access to your Cloud Storage resources. In order to support a uniform permissioning system, cloud storage has uniform bucket-level access. Using this feature disables ACLs for all cloud storage resources: access to Cloud Storage resources is then granted exclusively through Cloud IAM. Enabling uniform bucket-level access guarantees that if a storage bucket is not publicly accessible, no object in the bucket is publicly accessible either.
Affected Assets	Buckets: <ul style="list-style-type: none"> • ***** • ***** • ***** • ***** • ***** • *****

Cloud Compute Engine testing results

Name	Block Project SSH Keys Disabled
Risk	WARNING
Description	Project-wide SSH keys are stored in *****. Project wide SSH keys can be used to login to all the instances within the project. Using project-wide SSH keys eases the SSH key management but if compromised, poses the security risk which can impact all the instances within project.
Affected Assets	All

Name	Default Firewall Rule in Use
Risk	WARNING
Description	Some default firewall rules were in use. This could potentially expose sensitive services or protocols to other networks.
Affected Assets	Firewall Rules: <ul style="list-style-type: none"> • ***** • ***** • ***** • *****

Name	Default Network should be removed
Risk	WARNING
Description	The default network has a preconfigured network configuration and automatically generates insecure firewall rules. These automatically created firewall rules are not audit logged and cannot be configured to enable firewall rule logging.

Name	Firewall INGRESS Rule Allows Public Access (0.0.0.0/0) to a Sensitive Port
Risk	WARNING
Description	The firewall rule was found to expose a well-known port to all source addresses. Well-known ports are commonly probed by automated scanning tools and could be an indicator of sensitive services exposed to the internet. If such services need to be

Name	Firewall INGRESS Rule Allows Public Access (0.0.0.0/0) to a Sensitive Port
	exposed, a restriction on the source address could help to reduce the attack surface of the infrastructure.
Affected Assets	Firewall Rules: <ul style="list-style-type: none"> ***** *****

Name	Firewall Rule Allows Internal Traffic
Risk	WARNING
Description	The firewall rule allows ingress connections for all protocols and ports among instances in the network.
Affected Assets	Firewall Rules: <ul style="list-style-type: none"> *****

Name	Firewall Rule Allows Port Range(s)
Risk	WARNING
Description	It was found that the firewall rule was using port ranges. Sometimes ranges could include unintended ports that should not be exposed. As a result, when possible, explicit port lists should be used instead.
Affected Assets	Firewall Rules: <ul style="list-style-type: none"> *****

Name	Firewall Rule Allows Public Access (0.0.0.0/0)
Risk	WARNING
Description	The firewall rule was found to be exposing potentially open ports to all source addresses. Ports are commonly probed by automated scanning tools and could be an indicator of sensitive services being exposed to the internet. If such services need to be exposed, a restriction on the source address could help to reduce the attack surface of the infrastructure.

Name	Firewall Rule Allows Public Access (0.0.0.0/0)
Affected Assets	Firewall Rules: <ul style="list-style-type: none"> • ***** • ***** • ***** • ***** • *****

Name	Firewall Rule Opens All Ports (0-65535)
Risk	WARNING
Description	The firewall rule allows access to all ports. This widens the attack surface of the infrastructure and makes it easier for an attacker to reach potentially sensitive services over the network.
Affected Assets	Firewall Rules: <ul style="list-style-type: none"> • ***** • *****

Name	Instance Disk without Snapshots
Risk	WARNING
Description	You should have snapshots of your in-use or available disks taken on a regular basis to enable disaster recovery efforts.
Affected Assets	Instances: <ol style="list-style-type: none"> 1. ***** <ol style="list-style-type: none"> 1.1. ***** 1.2. ***** 1.3. ***** 1.4. ***** 1.5. ***** 1.6. *****

Name	Instance Disk without Snapshots
	<p>1.7. *****</p> <p>*****</p> <p>2. *****</p> <p>2.1. *****</p> <p>2.2. *****</p> <p>2.3. *****</p> <p>2.4. *****</p> <p>2.5. *****</p> <p>2.6. *****</p> <p>3. *****</p> <p>3.1. *****</p> <p>3.2. *****</p> <p>3.3. *****</p> <p>3.4. *****</p> <p>3.5. *****</p>

Name	Instance without Deletion Protection
Risk	WARNING
Description	It is good practice to enable this feature on production instances to ensure that they may not be deleted by accident.
Affected Assets	All

Name	Instances Configured to Use Default Service Account
Risk	WARNING
Description	The default Compute Engine service account has the editor role on the project, which allows read and write access to most Google Cloud services. To defend against privilege escalations if your VM is compromised and to prevent an attacker from gaining access to all of the project, it is recommended the default Compute Engine service account is

	not used. Instead, you should create a new service account and assigning only the permissions needed by your instance.
--	--

Name	Instances Have Public IP Addresses
Risk	WARNING
Description	To reduce your attack surface, compute instances should not have public IP addresses. Instead, instances should be configured behind load balancers to minimize the instance's exposure to the internet.

Name	Network without Instances
Risk	WARNING
Description	Maintaining unused resources increases the risk of misconfiguration and audit difficulty.
Affected Assets	Network instances: 1. ***** 1.1. ***** 1.2. *****

Name	OS login Disabled
Risk	WARNING
Description	Enabling OS login ensures that SSH keys used to connect to instances are mapped with IAM users. Revoking access to an IAM user will revoke all the SSH keys associated with that particular user. There is the ability to facilitate centralized and automated SSH key pair management which is useful in handling cases like responding to compromised SSH key pairs and/or the revocation of external/third-party/vendor users.
Affected Assets	All

Name	Shielded VM Disabled
Risk	WARNING
Description	Shielded VM offers verifiable integrity of your Compute Engine VM instances, so you can be confident your instances haven't been compromised by boot or kernel-level malware or rootkits. Shielded VM's verifiable integrity is achieved through the use of Secure Boot, virtual trusted platform module (vTPM) enabled Measured Boot and integrity monitoring.
Affected Assets	All

Name	VM Disks Not Customer-Supplied Encryption Keys (CSEK) Encrypted
Risk	WARNING
Description	By default, Google Compute Engine encrypts all data at rest. Compute Engine handles and manages this encryption for you without any additional actions on your part. However, if you want to control and manage this encryption yourself, you can provide your own encryption keys.
Affected Assets	All

IAM testing results

Name	Basic Role in Use
Risk	WARNING
Description	Basic roles grant significant privileges. In most cases, usage of these roles is not recommended and does not follow security best practice.
Affected Assets	Roles: <ul style="list-style-type: none"> • ***** • *****

Name	Gmail Account in Use
Risk	WARNING
Description	It is recommended that fully-managed corporate Google accounts be used for increased visibility, auditing and for controlling access to Cloud Platform resources. Email accounts based outside of the user's organization, such as personal accounts, should not be used for business purposes.
Affected Assets	Roles: <ul style="list-style-type: none"> • ***** • User: ***** • Project ID: ***** • Bindings: ***** • ***** • User: ***** • Project ID: ***** • Bindings: *****

Name	IAM Role Assigned to User
Risk	WARNING
Description	Best practices recommend granting roles to a Google Suite group instead of to individual users when possible. It is easier to add members to and remove members from a group instead of updating a Cloud IAM policy to add or remove users.
Affected Assets	Roles:

Name	IAM Role Assigned to User
	<p>*****</p> <ul style="list-style-type: none"> • ***** • ***** • ***** • ***** • ***** • ***** • ***** <p>*****</p> <ul style="list-style-type: none"> • ***** • ***** • *****

Name	Lack of Service Account Key Rotation
Risk	WARNING
Description	Rotating service account keys will reduce the window of opportunity for an access key that is associated with a compromised or terminated account to be used. Service account keys should be rotated to ensure that data cannot be accessed with an old key which might have been lost, cracked or stolen. It should be ensured that keys are rotated every 90 days.
Affected Assets	<p>Accounts:</p> <p>*****</p> <ul style="list-style-type: none"> • ***** • ***** • ***** • ***** • ***** <p>*****</p> <ul style="list-style-type: none"> • ***** • *****

Name	Service Account with Admin Privileges
Risk	WARNING
Description	Service accounts represent service-level security of the resources (application or a VM) which can be determined by roles that are assigned. Enrolling service accounts with administrative privileges grants full access to an assigned application or a VM. A service account access holder can be a user.
Affected Assets	<p>Accounts:</p> <p>*****</p> <ul style="list-style-type: none"> ***** ***** ***** ***** ***** <p>*****</p> <ul style="list-style-type: none"> ***** ***** *****

Name	User with Privileged Service Account Roles at the Project Level
Risk	WARNING
Description	Granting the *****, *****, or ***** role to a user for a project gives the user access to all service accounts within the project, including service accounts that may be created in the future. This can result in an elevation of privileges by using service accounts and corresponding Compute Engine instances.
Affected Assets	<p>Accounts:</p> <p>*****</p> <ul style="list-style-type: none"> *****

Name	User-Managed Service Account Keys
Risk	WARNING
Description	It is recommended that user-managed service account keys are kept confidential as anyone who has access to the keys will be able to view resources through the service account. Best practice would be to use GCP-managed keys which are used by Cloud Platform services, such as App Engine and Compute Engine. These keys cannot be downloaded. Google will keep the keys and automatically rotate them approximately on a weekly basis.
Affected Assets	Accounts: ***** <ul style="list-style-type: none"> • ***** • *****

CONFIDENTIAL

Kubernetes Engine testing results

Name	Clusters Lacking Labels
Risk	WARNING
Description	Labels enable users to map their own organizational structures onto system objects in a loosely coupled fashion, without requiring clients to store these mappings. Labels can also be used to apply specific security settings and auto-configure objects at creation.
Affected Assets	<p>Clusters:</p> <p>*****</p> <ul style="list-style-type: none"> *****

Name	Default Service Account in Use
Risk	WARNING
Description	You should create and use a minimally privileged service account to run your Kubernetes Engine cluster instead of using the Compute Engine default service account.
Affected Assets	<p>Clusters:</p> <p>*****</p> <ul style="list-style-type: none"> ***** <p>*****</p> <ul style="list-style-type: none"> ***** *****

Name	Lack of Access Scope Limitation 2
Risk	WARNING
Description	If you are not creating a separate service account for a node, you should limit the scopes of the node service account to reduce the possibility of a privilege escalation in an attack. This ensures that your default service account does not have permissions beyond those necessary to run your cluster. While the default scopes are limited, they may include scopes beyond the minimally required scopes needed to run a cluster. If you are accessing private images in Google Container

Name	Lack of Access Scope Limitation 2
	Registry, the minimally required scopes are logging.write, monitoring and devstorage.read_only.
Affected Assets	<p>Clusters:</p> <p>*****</p> <ul style="list-style-type: none"> • ***** <p>*****</p> <ul style="list-style-type: none"> • ***** • *****

Name	Master Authorized Networks Disabled
Risk	WARNING
Description	Master authorized networks block untrusted IP addresses from outside the Google Cloud Platform. Addresses from inside GCP can still reach your master through HTTPS provided they have the necessary Kubernetes credentials.
Affected Assets	<p>Clusters:</p> <ul style="list-style-type: none"> • ***** <p>*****</p> <ul style="list-style-type: none"> • ***** • *****

Name	Network Policy Disabled
Risk	WARNING
Description	By default, pods are non-isolated; they accept traffic from any source. Pods become isolated by having a network policy that selects them. Once there is a network policy in a namespace selecting a particular pod, that pod will reject any connections that are not allowed by the network policy.
Affected Assets	Clusters:

Name	Network Policy Disabled
	<p>*****</p> <ul style="list-style-type: none"> • ***** <p>*****</p> <ul style="list-style-type: none"> • ***** • *****

Name	Pod Security Policy Disabled
Risk	WARNING
Description	A Pod Security Policy is a cluster-level resource that controls security sensitive aspects of the pod specification. The pod security policy objects define a set of conditions that a pod must run with in order to be accepted into the system, as well as defaults for the related fields.
Affected Assets	<p>Clusters:</p> <p>*****</p> <ul style="list-style-type: none"> • ***** <p>*****</p> <ul style="list-style-type: none"> • ***** • *****

Name	Private Cluster Disabled
Risk	WARNING
Description	A private cluster is a cluster that makes your master inaccessible from the public internet. In a private cluster, nodes do not have public IP addresses, so your workloads run in an environment that is isolated from the internet. Nodes have addresses only in the private RFC address space. Nodes and masters communicate with each other privately using VPC peering.
Affected Assets	Clusters:

	<p>*****</p> <ul style="list-style-type: none"> ● ***** <p>*****</p> <ul style="list-style-type: none"> ● ***** ● *****
--	--

Name	Private Google Access Disabled
Risk	WARNING
Description	Enabling private Google access allows VMs on a subnetwork to use a private IP address to reach Google APIs rather than an external IP address.
Affected Assets	<p>Clusters:</p> <p>*****</p> <ul style="list-style-type: none"> ● ***** <p>*****</p> <ul style="list-style-type: none"> ● ***** <p>*****</p> <ul style="list-style-type: none"> ● ***** ● *****

Name	Nodes Auto-Upgrade Disabled
Risk	WARNING
Description	Auto-upgrades automatically ensure that security updates are applied and kept up to date.
Affected Assets	<p>Clusters:</p> <p>*****</p> <ul style="list-style-type: none"> ● ***** ● *****

Stackdriver logging & monitoring testing results

1. **WARNING** – Log Metric Filter Issues

Name	Description
Log metric filter doesn't exist for audit configuration changes	Configuring the metric filter and alerts for audit configuration changes ensures the recommended state of audit configuration is maintained so that all activities in the project are auditable at any point in time.
Log metric filter doesn't exist for Cloud Storage IAM permission changes	Monitoring changes to Cloud Storage bucket permissions may reduce the time needed to detect and correct permissions on sensitive Cloud Storage buckets and objects inside the bucket.
Log metric filter doesn't exist for custom role changes	Google Cloud IAM provides predefined roles that give granular access to specific Google Cloud Platform resources and prevent unwanted access to other resources. However, to cater to organization-specific needs, Cloud IAM also provides the ability to create custom roles. Project owners and administrators who are an organization role administrator or IAM role administrator can create custom roles. Monitoring role creation, deletion and updating activities will help in identifying any over-privileged role during the early stages.
Log metric filter doesn't exist for project ownership assignments/changes	Project ownership has the highest level of privileges on a project. To avoid misuse of project resources, the project ownership assignment/change actions mentioned above should be monitored and concerned recipients alerted.
Log metric filter doesn't exist for SQL instance configuration changes	Monitoring alterations to SQL instance configuration changes may reduce the time needed to detect and correct misconfigurations done on the SQL server.
Log metric filter doesn't exist for VPC network changes	It is possible to have more than one VPC within a project. In addition, it is also possible to create a peer connection between two VPCs enabling network traffic to route between VPCs. Monitoring changes to a VPC will help ensure VPC traffic flow is not being impacted.
Log metric filter doesn't exist for VPC network firewall rule changes	Monitoring for create or update firewall rule events gives insight into network access changes and may reduce the time it takes to detect suspicious activity.
Log metric filter doesn't exist for VPC network route changes	Google Cloud Platform (GCP) routes define the paths network traffic take from a VM instance to another destination. The other destination can be inside the organization's VPC network (such as

Name	Description
	another VM) or outside of it. Every route consists of a destination and a next hop. Traffic whose destination IP is within the destination range is sent to the next hop for delivery. Monitoring changes to route tables will help ensure that all VPC traffic flows through an expected path.
Affected logging configurations:	

2. **WARNING** – Alerts Setup Issues

Name	Description
Alerts don't exist for audit configuration changes	Configuring the metric filter and alerts for audit configuration changes ensures the recommended state of audit configuration is maintained so that all activities in the project are auditable at any point in time.
Alerts don't exist for Cloud Storage IAM permission changes	Monitoring changes to Cloud Storage bucket permissions may reduce the time needed to detect and correct permissions on sensitive Cloud Storage buckets and objects inside the bucket.
Alerts don't exist for custom role changes	Google Cloud IAM provides predefined roles that give granular access to specific Google Cloud Platform resources and prevent unwanted access to other resources. However, to cater to organization-specific needs, Cloud IAM also provides the ability to create custom roles. Project owners and administrators who are an organization role administrator or IAM role administrator can create custom roles. Monitoring role creation, deletion and updating activities will help in identifying any over-privileged role during the early stages.
Alerts don't exist for project ownership assignments/changes	Project ownership has the highest level of privileges on a project. To avoid misuse of project resources, the project ownership assignment/change actions mentioned above should be monitored and concerned recipients alerted.
Alerts don't exist for SQL instance configuration changes	Monitoring alterations to SQL instance configuration changes may reduce the time needed to detect and correct configurations done on the SQL server.


Name	Description
Alerts don't exist for VPC network changes	It is possible to have more than one VPC within a project. In addition, it is also possible to create a peer connection between two VPCs enabling network traffic to route between VPCs. Monitoring changes to a VPC will help ensure VPC traffic flow is not being impacted.
Alerts don't exist for VPC network firewall rule changes	Monitoring for create or update firewall rule events gives insight into network access changes and may reduce the time it takes to detect suspicious activity.
Alerts don't exist for VPC network route changes	Google Cloud Platform routes define the paths network traffic take from a VM instance to another destination. The other destination can be inside the organization's VPC network (such as another VM) or outside of it. Every route consists of a destination and a next hop. Traffic whose destination IP is within the destination range is sent to the next hop for delivery. Monitoring changes to route tables will help ensure that all VPC traffic flows through an expected path.
Affected logging configurations:	
<p>*****</p> <p>*****</p>	

Appendix B: Vulnerability Detail and Mitigation

Risk Rating Scale

In accordance with NIST SP 800-30, exploited vulnerabilities are ranked based upon likelihood and impact to determine overall risk.

SSRF on JSON API functionality

Risk	CRITICAL
Category	A10:2021 – SSRF
Description	SSRF flaws occur whenever a web application is fetching a remote resource without validating the user-supplied URL.
Impact	Confidential information such as Kubernetes environment variables, network interfaces, service accounts or SSH keys are leaked. 
Recommendation	The developer team may need a unique solution in this event as this process can negatively impact businesses.

XSS

Risk	HIGH
Category	A03:2021 – Injection
Description	Unrestricted file upload leads Stored-XSS.
Impact	An attacker copies the ***** login page and modifies it to send the user submitted data to their own server. The attacker then uploads the HTML file via the vulnerable endpoint and sends it to the user, convincing them that it is the login page. As the domain(*****) belongs to ***** , the victim has a high chance of being deceived.
Recommendation	https://cheatsheetseries.owasp.org/cheatsheets/Cross_Site_Scripting_Prevention_Cheat_Sheet.html


IDOR-Privilege Escalation

Risk	HIGH
Category	A01:2021- BROKEN ACCESS CONTROL
Description	An unprivileged user (Company.MEMBER) can change the name and avatar of the company (even if it doesn't belong to the company).
Impact	<ol style="list-style-type: none"> 1.Attacker gets ***** of the target company. It can be retrieved with the help of ***** endpoint. 2.Then, the attacker modifies settings of the company where they have "Company.ADMIN" permission (*****) and intercepts the request with proxy software. 3.A company ID is sent via a request. However, it is unlikely that the user who sent the request will checked to confirm they have enough privilege on that company. So, the user can replace their own company's ID (*****) with the target company's ID (*****). The modified request is shown in the following image. 4.After sending the request, the target company's name is changed. To confirm this, we can use ***** endpoint.
Recommendation	DNS zone transfers should be restricted only to pre-approved servers.

Google Captcha bypass

Risk	MEDIUM
Category	A2:2017 – Broken Authentication
Description	There are 3 pages that uses reCAPTCHA; login, password reset and registration. While testing these functionalities, it appears the captcha provided by reCAPTCHA is not validated.
Impact	The captcha doesn't work on any of the 3 pages named above. Each of them has its own impact. These are: <ul style="list-style-type: none"> • Login – brute force • Password reset – Sending large volumes of emails from ***** mail address. • Registration – Creating a large number of fake accounts
Recommendation	There is a logical flaw in captcha implementation. Make sure every request is checked for the correct captcha and is then processed.


User enumeration

Risk	LOW
Category	A2:2017 – Broken Authentication
Description	User enumeration is when a malicious actor can use brute-force techniques to either guess or confirm valid users in a system.
Impact	<p>In the registration page, if the user tries to register with an existing email address, the following error message will be displayed:</p>  <p>But, with a non-existent one, a new account will be created. Based on these two response messages, it's possible to determine whether an email address is registered.</p>
Recommendation	The same response should be returned whether the email address entered by the user exists or not.

Security Misconfiguration - Exposed Test environment

Risk	LOW
Category	A6:2017 – Security Misconfiguration
Description	When analyzing public resources belonging to *****, some resources that shouldn't be public were available to everyone. Such test environments may contain source code of future features that are not yet meant to be publicly available. Exposed test environments pose weak entry points into internal networks and can lead to data exposure and leaks. In addition to potential leaks, since most test environments are not regularly monitored, attackers could apply their exploits on exposed staging environments until they are ready and able to take down the live(prod.) application in one attempt.
Impact	Some of them are as follows: <ul style="list-style-type: none"> • ***** • ***** • ***** • *****
Recommendation	To remediate this issue, some access controls should be implemented. There are some choices, like implementing VPN, adding an additional security layer (login page, MFA) or any security measure to confirm whether the person who wants to access one of those resources has permission to access the test environment.

Code Review: Insecure Randomness

Risk	LOW
Category	Insecure Randomness
Description	Standard pseudorandom number generators cannot withstand cryptographic attacks. A PRNG is an algorithm used to produce random-looking numbers with certain desirable statistical properties. In order for a PRNG to be cryptographically secure, it must be resistant to prediction.
Impact	<ul style="list-style-type: none"> • *****, line 480 

	<ul style="list-style-type: none"> • ***** , line <div style="background-color: #cccccc; height: 100px; width: 100%;"></div> <ul style="list-style-type: none"> • ***** , line 44 <div style="background-color: #cccccc; height: 80px; width: 100%;"></div>
Recommendation	<p>We recommend using the secrets module's PRNG as follows: https://docs.python.org/dev/library/secrets.html#secrets.SystemRandom</p>

CONFIDENTIAL

Appendix C: About ESKA

We are the providers of external and internal network penetration services, which help reveal vulnerabilities before “real” hackers do. Our work is completed in a controlled and secure framework which does not exploit any security gaps that are discovered, so clients can see the holes in their cybersecurity and fill them with modern cybersecurity tools for ultimate safety. A week rarely goes by without reports of attacks on sensitive systems. These often result in financial damage and can negatively impact the reputation and trust of customers and partners.

To protect yourself against attacks, adequate countermeasures must be taken at different levels. Well-trained employees and processes that also take IT security into account are essential for effective protection. However, complete and comprehensive security checks that are delivered through a penetration test by an independent third party are the only way to absolutely assure your organization is safe.

So, what exactly is a penetration test? It's an authorized, planned and simulated cyberattack on a company or a public sector institution. The aim is to identify and eliminate previously unknown points of attack before hackers can use them to steal intellectual property, sensitive data or otherwise damage an organization.

During the penetration test, specialist testers attempt to attack your IT systems using the methods of criminal hackers to determine the vulnerability of systems, after which appropriate protective measures can be taken.

There are two types of businesses:

- those that have been already hacked
- those that will be hacked once.

To effectively protect yourself against these attacks, penetration tests can give a clear picture of your system's security situation.

If you would like to discuss your penetration testing needs, please contact us at office@eska.global