Penetration Test Report

*****

*****

ESKA

Hlybochytska Street, 17B, **Kyiv, Ukraine**

office@eska.global

eska.global

+380 (44) 247 10 21

Table of Contents

Executive Summary

ESKA was contracted by ***** to conduct a penetration test in order to determine its exposure to a targeted attack, and Infrastructure Assessment to evaluate configurations regarding security best practices. All activities regarding penetration test were conducted in a manner that simulated a malicious actor engaged in a targeted attack against ***** with the goals of:

- 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

Penetration test was expanded with source code analysis for determination of programming errors and unsecure data flows. Efforts were placed on the identification and exploitation of security weakness 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 provided by ***** credentials and accesses. The assessment was conducted in ***** with the recommendations outlined in 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 Infrastructure Assessment were conducted according to Google Cloud Platform (GCP) security best practices with the goals of:

- Ensure that necessary security controls are integrated into the design and implementation of a project.

Check and evaluate security configurations that should ensure the Confidentiality, Integrity and Availability of ***** sensitive data and other resources.

Summary of Results

Initial reconnaissance of the ***** infrastructure and services of a settings, that need attention. The results provided us with a listing 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 total). After using a custom "Gray Box" technique on the ***** infrastructure we were able to find list of issues according to Google Security Checklist. HIGH-level and some WARNING-level issues was additionally checked with custom scripts and techniques, with set of tools like Burp, MetaSploit, etc. There are not any critical results, but this is need additional attention anyway. Uncovering the passwords via brute-force was not completed with using basic techniques. Cloud penetration testing (uses simulated cyberattacks against target systems to identify vulnerabilities) engages concept, that is performed on cloud-native systems. This form of security testing is used to identify security risks and vulnerabilities, and 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 there were found an IDOR vulnerability, that allows this user to change company name and the avatar of the company, and Stored XSS vulnerability. Additionally, there were found 2 vulnerabilities regarding API with **CRITICAL**-risk and **HIGH**-risk ratings. Other vulnerabilities have **LOW** and Informational risk ratings but still considerable to be remediated.

1. INFRASTRUCTURE ASSESSMENT

1.1. Discovery

For the purposes of this assessment, CLIENT provided cloud account with View permission, suitable for "Gray Box" Pentest. During enumeration stage was founded 2 projects, 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, query outputs etc. For security, it is recommended to always use SSL encryption 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 determine 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 run overs on deadlock timeout can be an indication of an underlying issue. Logging such waits on locks by enabling the log_lock_waits flag 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 that 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 classified as error messages appropriately. 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.

   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 that 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 – It is recommended to use uniform bucket-level access to unify and simplify how you grant access 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 to Cloud Storage resources then is 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.

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

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-meta-data. Project wide SSH keys can be used to login into all the instances within 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.

   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 could be an indicator of sensitive services exposed to 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 could be an indicator of sensitive services exposed to 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:

   *****

*****

*****

*****

*****

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 prevent an attacker from gaining access to all of your project, it is recommended to not use the default Compute Engine service account. Instead, you should create a new service account and assigning only the permissions needed 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 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 risks of misconfigurations and increases the difficulty of audits.

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 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 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.

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 wanted 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 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 practices recommends 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.

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. It should be ensured that keys are 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 assigned application or a VM, Service Account Access holder can 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, am.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 into 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 – It is recommended to prevent use of user-managed service account keys, as anyone who has access to the keys will be able to access resources through the service account. Best practice recommends using 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 on an approximately 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 nodes, 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 ecessary 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 blocks untrusted IP addresses from outside Google Cloud Platform. Addresses from inside GCP can still reach your master 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 any Network Policy in a namespace selecting a particular pod, that pod will reject any connections that are not allowed by any 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 PodSecurityPolicy 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.

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 is ated 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 ensures that security updates are applied and kept up to date.

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 audit-able 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 with the Organization Role Administrator role or the IAM Role Administrator role can create custom roles. Monitoring role creation, deletion and updating activities will help in identifying any over-privileged role at 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 alerted to concerned recipients.

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

Description - Monitoring changes 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

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 getting impacted.

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

Description - Monitoring for Create or Update Firewall rule events gives insight to 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 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/logging/docs/audit/configure-data-access#getiampolicy-setiampolicy

Affected Logging Configurations:

*****

*****

2. WARNING - Alerts Setup Issues

Alerts 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 audit-able at any point in time.

Alerts 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.

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 with the Organization Role Administrator role or the IAM Role Administrator role can create custom roles. Monitoring role creation, deletion and updating activities will help in identifying any over-privileged role at early stages.

Alerts 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 alerted to concerned recipients.

Alerts Doesn't Exist for SQL Instance Configuration Changes

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

Alerts 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 getting impacted.

Alerts Doesn't Exist for VPC Network Firewall Rule Changes

Description - Monitoring for Create or Update Firewall rule events gives insight to network access changes and may reduce the time it takes to detect suspicious activity.

Alerts Doesn'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 list of issues according to Google Security Checklist. HIGH-level and some WARNING-level issues was additionally checked with custom scripts and techniques, with set of tools like Burp, MetaSploit, etc. There are not any critical results, but this is need additional attention anyway.

2. APPLICATION ASSESSMENT

2.1. Introduction

The cybersecurity team performed Pentest on the ***** application a blackbox and whitebox approach, simulating attack vectors that attackers 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.

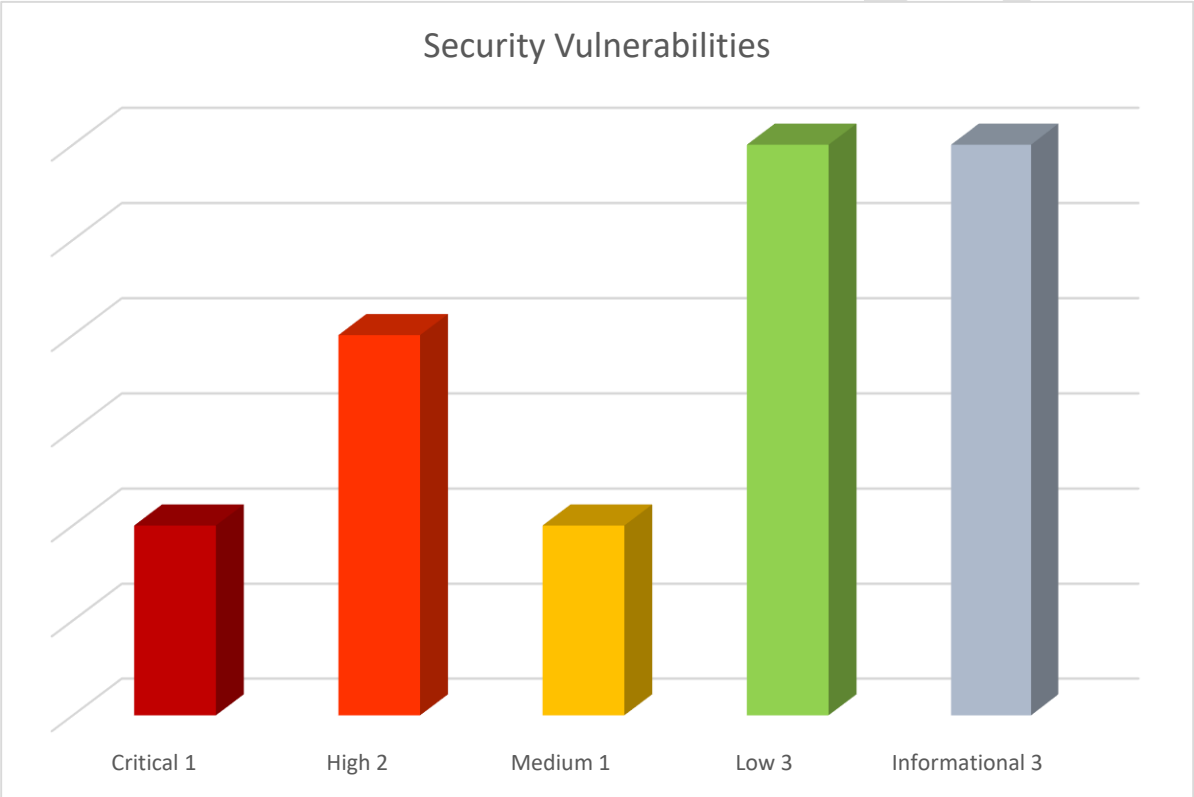

2.2. User enumeration

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 email or password is wrong, in both cases "Wrong email or password." message is displayed to the user. Thus, it is not indicated that which(email or password) is incorrect one.

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 whether any 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 mail address, the following message is displayed.

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

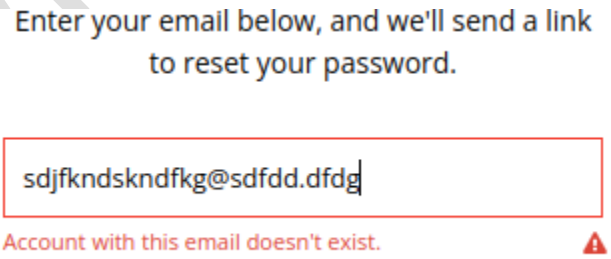

2. In the registration page, if the user tries to register with an existing email address, the following error message will be displayed.

But, with non-existing 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, 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

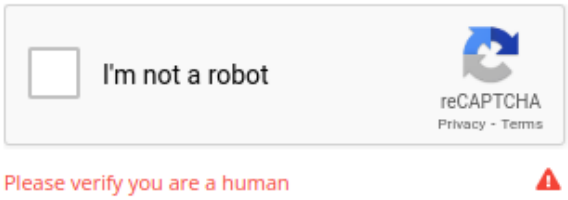

- When the user confirms the captcha and clicks to “Login” button, a request is sent as follows:

- The problem here is, if we re-use 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 he wants until he finds the right one.
- The screenshot below proves that despite checking 50 passwords, no protection mechanism is triggered

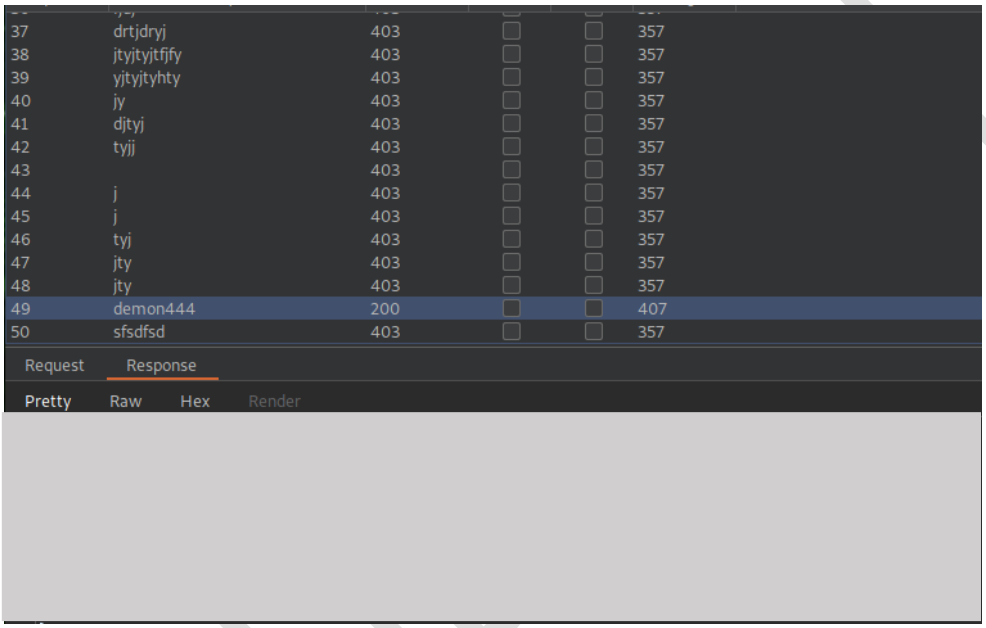

37	drtjdryj	403			357
38	jtyjtyjtfjfy	403			357
39	yjtyjtyhty	403			357
40	jy	403			357
41	djtyj	403			357
42	tyjj	403			357
43		403			357
44	j	403			357
45	j	403			357
46	tyj	403			357
47	jty	403			357
48	jty	403			357
49	demon444	200			407
50	sfsdfsd	403			357

Request Response

Pretty Raw Hex Render

In this simulation, attacker tried 50 passwords, and was able to find the correct one in the 49th request. In the response, 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 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.

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

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

While analyzing the application, it turned out that, 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), attacker can access to metadata server API and retrieve **confidential** data from VM metadata server.

Exploitation process:

As Proof of Concept, we'll retrieve VM metadata from Google Cloud VM metadata server

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

4. Attacker saves the block and click to "Test Chat" button.

5. When the chat starts, the HTTP request is sent. It's possible to view webhook history in ***** page.

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

In this response, confidential information such as Kubernetes environment variables, Network interfaces, service accounts, 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

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

.etc

Remediation:

We discussed this vulnerability and your developer team needs a unique solution for this case. Because the best practice of remediation can bad reflect your business process.

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 HTML file to the server, then executed it on the victim's browser.

1. Attacker creates a new "JSON API" bot block.

[redacted]

2. Uploads the file to the server, using /case/attachment endpoint.

Figure 1

3. Delivers the malicious URL to victim user. The URL is value of "*****" parameter in the response. (Above screenshot)

   URL: *****

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

4. With the help of 2nd line (Figure 1), the path section after the domain is deleted using the "pushState" method. This makes the URL even more realistic.

5. This behavior can lead the following security issue:

An attacker copies ***** login page, and modify it to send user submitted data to his own server. 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

Note: Since the session token is stored under ***** domain's Local Storage, this vulnerable endpoint has not 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:

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. 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" (*****), he cannot modify this company's settings.

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

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

4. After sending 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 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 turned out to be available to 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 could "practice" their exploits on exposed staging environments until they are ready and able to take down the live(prod.) application in one shot.

Remediation:

To remediate this issue, some Access Controls should be implemented. There are some choices, like implementing VPN, adding 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 to 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

- *****, line 119

- *****, line 44

Remediation:

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

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

2.9. Informational: Advices

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 an up-to-date version.

2. ***** : Ngrok service

   In the future, please notice this service. Because developers can publish sensitive services or API's with Ngrok application.

3. Dangerous allowlist policy

   The ***** application uses the Flask Jinja template engine. We notice when we started to build a bot, we can inject mathematical operations on the template engine but developers use the ***** library for security. Therefore we couldn't inject a malicious payload(for example: A remote code execution payload) into the system.

   [redacted]

   But we could bypass some restrictions like a “double bracket”:

   [redacted]

Application deleted our double bracket payload

An attacker can bypass that restriction using a “triple bracket”.

In the 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. "*****" application has passed penetration testing check with a 5/10 score. The overall risk identified to ***** as a result of the penetration test is **Medium**.

Appendix A: Infrastructure Assessment Results

Cloud SQL testing 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	Projects: • ***** • ***** Databases: • ***** • ***** • *****
Recommendation	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.

Name	**Instance Not Requiring SSL for Incoming Connections**
Risk	**WARNING**
Description	SQL database connections if 1successfully trapped (MITM); can reveal sensitive data like credentials, database queries, query outputs etc. For security, it is recommended to always use SSL encryption 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 data recovery, etc.) 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 determine 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 such waits on locks by enabling the log_lock_waits flag 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.

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 log_min_error_statement is not set to the correct value, messages may not be classified as error messages appropriately. 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 may be due to either poor application coding or deliberate resource starvation attempts.

Cloud Storage testing 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: • ***** • ***** • ***** • ***** • ***** • ***** • ***** • ***** • *****

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: • ***** • ***** • ***** • ***** • ***** • ***** • ***** • ***** • ***** • *****

Name	**Uniform Bucket-Level Access Is Disabled**
Risk	**WARNING**
Description	It is recommended to use uniform bucket-level access 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 then is 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: • ***** • ***** • ***** • ***** • ***** • *****

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 into all the instances within 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: • ***** • ***** • ***** • *****

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 do not get 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 be exposing 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 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: • ***** • *****

Name	**Firewall Rule Allows Internal Traffic**
Risk	WARNING
Description	Firewall rule allows ingress connections for all protocols and ports among instances in the network.
Affected Assets	Firewall Rules: • *****

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: • *****

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 exposed to 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: • ***** • ***** • ***** • ***** • *****

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: • ***** • *****

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: 1. ***** 1.1. ***** 1.2. ***** 1.3. ***** 1.4. ***** 1.5. ***** 1.6. *****

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

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 prevent an attacker from gaining access to all of your project, it is recommended to not use the default Compute Engine

	service account. 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 risks of misconfigurations and increases the difficulty of audits.
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 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.
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 wanted 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: • ***** • *****

Name	**Gmail Account in Use**
Risk	WARNING
Description	It is recommended fully-managed corporate Google accounts be used for increased visibility, auditing, and 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: ***** • 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**
	***** • ***** • ***** • ***** • ***** • ***** • ***** • ***** ***** • ***** • ***** • *****

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	Accounts: ***** • ***** • ***** • ***** • ***** • ***** ***** • ***** • *****

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 the roles assigned to it. Enrolling Service Accounts with administrative privileges grants full access to assigned application or a VM, Service Account Access holder can user.
Affected Assets	Accounts: ***** • ***** • ***** • ***** • ***** • ***** ***** • ***** • ***** • *****

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 in the project, including service accounts that may be created in the future. This can result into elevation of privileges by using service accounts and corresponding Compute Engine instances.
Affected Assets	Accounts: ***** • *****

Name	**User-Managed Service Account Keys**
Risk	WARNING
Description	It is recommended to prevent use of user-managed service account keys, as anyone who has access to the keys will be able to access resources through the service account. Best practice recommends using 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 on an approximately weekly basis.
Affected Assets	Accounts: ***** • ***** • *****

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	Clusters: ***** • *****

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	Clusters: ***** • ***** ***** • ***** • *****

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 only logging.write, monitoring, and devstorage.read_only.
Affected Assets	Clusters: ***** • ***** ***** • ***** • *****

Name	**Master Authorized Networks Disabled**
Risk	WARNING
Description	Master authorized networks blocks untrusted IP addresses from outside GoogleCloud Platform. Addresses from inside GCP can still reach your master through HTTPS provided that they have the necessary Kubernetes credentials.
Affected Assets	Clusters: • ***** ***** • ***** • *****

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 NetworkPolicy that selects them. Once there is any NetworkPolicy in a namespace selecting a particular pod, that pod will reject any connections that are not allowed by any NetworkPolicy.
Affected Assets	Clusters:

Name	**Network Policy Disabled**
	***** • ***** ***** • ***** • *****

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 PodSecurityPolicy 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	Clusters: ***** • ***** ***** • ***** • *****

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 PodSecurityPolicy 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	Clusters:

Name	**Pod Security Policy Disabled**
	***** • ***** ***** • ***** • *****

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 addressed only in the private RFC address space. Nodes and masters communicate with each other privately using VPC peering.
Affected Assets	Clusters: ***** • ***** ***** • ***** • *****

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	Clusters: ***** • *****

Name	**Private Google Access Disabled**
	***** • ***** ***** • ***** • *****

Name	**Nodes Auto-Upgrade Disabled**
Risk	WARNING
Description	Auto-upgrades automatically ensures that security updates are applied and kept up to date.
Affected Assets	Clusters: ***** • ***** • *****

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 audit-able 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 with the Organization Role Administrator role or the IAM Role Administrator role can create custom roles. Monitoring role creation, deletion and updating activities will help in identifying any over-privileged role at 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 alerted to concerned recipients.
Log Metric Filter Doesn't Exist for SQL Instance Configuration Changes	Monitoring changes 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 getting impacted.
Log Metric Filter Doesn't Exist for VPC Network Firewall Rule Changes	Monitoring for Create or Update Firewall rule events gives insight to 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 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

Name	Description
	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 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 audit-able at any point in time.
Alerts 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.
Alerts 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 with the Organization Role Administrator role or the IAM Role Administrator role can create custom roles. Monitoring role creation, deletion and updating activities will help in identifying any over-privileged role at early stages.
Alerts 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 alerted to concerned recipients.
Alerts Doesn't Exist for SQL Instance Configuration Changes	Monitoring changes to SQL instance configuration changes may reduce the time needed to detect and correct is configurations done on the SQL server.
Alerts 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

Name	Description
	enabling network traffic to route between VPCs. Monitoring changes to a VPC will help ensure VPC traffic flow is not getting impacted.
Alerts Doesn't Exist for VPC Network Firewall Rule Changes	Monitoring for Create or Update Firewall rule events gives insight to network access changes and may reduce the time it takes to detect suspicious activity.
Alerts Doesn't Exist for VPC Network Route Changes	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.
Affected Logging Configurations:	
***** *****	

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, SSH keys are leaked.
Recommendation	We discussed this vulnerability and your developer team needs a unique solution for this case. Because the best practice of remediation can bad reflect your business process.

XSS

Risk	HIGH
Category	A03:2021 – Injection
Description	Unrestricted file upload leads Stored-XSS.
Impact	An attacker copies ***** login page, and modify it to send user submitted data to his own server. 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	Unprivileged user (Company.MEMBER) can change the name and avatar of the company (even if it doesn't belong to the company).
Impact	1. Attacker gets ***** of the target company. It can be retrieved with the help of ***** endpoint. 2. Then, attacker modifies settings of the company where he has "Company.ADMIN" permission (*****), and intercept the request with proxy software. 3. Company ID is sent via request. However, it is not checked whether the user who sent the request has enough privilege on that company. So, user can replace his own company's ID (*****) with the target company's ID (*****). Modified request is shown in the following image. 4. After sending 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, Registration. While testing these functionalities, it turned out that, 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, • Login – Brute-Force • Password reset – Sending large volumes of email from ***** mail address. • Registration – Creating large number of fake accounts
Recommendation	There is a logical flaw in captcha implementation. Make sure that 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	In the registration page, if the user tries to register with an existing email address, the following error message will be displayed.  But, with non-existing one, a new account will be created. Based on these two response messages, it's possible to determine whether an email address is registered.
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 turned out to be available to everyone. 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 could "practice" their exploits on exposed staging environments until they are ready and able to take down the live(prod.) application in one shot.
Impact	Some of them are as follows: • ***** • ***** • ***** • *****

Recommendation	To remediate this issue, some Access Controls should be implemented. There are some choices, like implementing VPN, adding 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 to 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	• *****, line 480 [redacted] • *****, line [redacted] • *****, line 44 [redacted]
Recommendation	We recommend using the secrets module's PRNG as follows: https://docs.python.org/dev/library/secrets.html#secrets.SystemRandom

Appendix C: About ESKA

We are the providers of external and internal network penetration services, which could help reveal vulnerabilities before "real" hackers do. All this in a controlled and secure framework and without exploiting the security gaps found, so you could see the holes in your cybersecurity and fill them with the modern cybersecurity tools – no one unwanted could ever get in.

A week rarely goes by without reports of attacks on sensitive systems. It results in financial damage, and the reputation and trust of customers and partners' crumble.

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, above all, the security check through a penetration test by an independent third party is an effective means.

So, what is exactly a penetration test? A penetration test is an authorized, planned, and a simulated cyber-attack 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 or other sensitive data or otherwise damage an organization.

During the penetration test, trained 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 hacker attacks, penetration tests can give a clear picture of the system's security situation.

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