

June 2025
Google Cloud
CMMC Level 2 Implementation Guide
FINAL

Google Cloud

CMMC Level 2 Implementation Guide

Table of Contents

Table of Contents	2
Overview of the CMMC Implementation Guide	6
Purpose	6
Audience	6
Scope	6
Overview of CMMC	7
CMMC Program	7
Cloud Service Provider (CSP) Requirements	8
Key CMMC Requirements	8
Overview of Google Cloud	9
Purpose	9
Assured Workloads	9
Customer Responsibility Matrix (CRM)	9
How to Use this Guide	10
Control Implementation Tables	10
Control Categories	11
1. Controls requiring customer implementation in Google Cloud	11
2. Controls requiring implementation outside of Google Cloud	12
3. Controls natively implemented by Google	13
CMMC L1 and L2 Control Implementation Guidance	13
AC.L1-3.1.1	13
AC.L1-3.1.2	17
AC.L2-3.1.3	20
AC.L2-3.1.4	23
AC.L2-3.1.5	26
AC.L2-3.1.6	30
AC.L2-3.1.7	33
AC.L2-3.1.8	36
AC.L2-3.1.9	39
AC.L2-3.1.10	41
AC.L2-3.1.11	43
AC.L2-3.1.12	45
AC.L2-3.1.13	50
AC.L2-3.1.14	53
AC.L2-3.1.15	56
AC.L2-3.1.16	59
AC.L2-3.1.17	60

AC.L2-3.1.18	60
AC.L2-3.1.19	63
AC.L1-3.1.20	66
AC.L2-3.1.21	69
AC.L1-3.1.22	70
AT.L2-3.2.1	73
AT.L2-3.2.2	74
AT.L2-3.2.3	75
AU.L2-3.3.1	75
AU.L2-3.3.2	81
AU.L2-3.3.3	82
AU.L2-3.3.4	85
AU.L2-3.3.5	86
AU.L2-3.3.6	88
AU.L2-3.3.7	89
AU.L2-3.3.8	90
AU.L2-3.3.9	95
CM.L2-3.4.1	96
CM.L2-3.4.2	98
CM.L2-3.4.3	101
CM.L2-3.4.4	104
CM.L2-3.4.5	105
CM.L2-3.4.6	106
CM.L2-3.4.7	111
CM.L2-3.4.8	116
CM.L2-3.4.9	117
IA.L1-3.5.1	118
IA.L1-3.5.2	121
IA.L2-3.5.3	126
IA.L2-3.5.4	129
IA.L2-3.5.5	131
IA.L2-3.5.6	134
IA.L2-3.5.7	136
IA.L2-3.5.8	139
IA.L2-3.5.9	142
IA.L2-3.5.10	144
IA.L2-3.5.11	146
IR.L2-3.6.1	147
IR.L2-3.6.2	158

IR.L2-3.6.3	158
MA.L2-3.7.1	159
MA.L2-3.7.2	160
MA.L2-3.7.3	160
MA.L2-3.7.4	161
MA.L2-3.7.5	162
MA.L2-3.7.6	162
MP.L2-3.8.1	163
MP.L2-3.8.2	164
MP.L1-3.8.3	167
MP.L2-3.8.4	168
MP.L2-3.8.5	172
MP.L2-3.8.6	172
MP.L2-3.8.7	173
MP.L2-3.8.8	174
MP.L2-3.8.9	174
PS.L2-3.9.1	179
PS.L2-3.9.2	179
PE.L1-3.10.1	183
PE.L2-3.10.2	184
PE.L1-3.10.3	184
PE.L1-3.10.4	185
PE.L1-3.10.5	186
PE.L2-3.10.6	186
RA.L2-3.11.1	187
RA.L2-3.11.2	191
RA.L2-3.11.3	195
CA.L2-3.12.1	199
CA.L2-3.12.2	199
CA.L2-3.12.3	200
CA.L2-3.12.4	204
SC.L1-3.13.1	205
SC.L2-3.13.2	211
SC.L2-3.13.3	214
SC.L2-3.13.4	217
SC.L1-3.13.5	222
SC.L2-3.13.6	225
SC.L2-3.13.7	227
SC.L2-3.13.8	230

SC.L2-3.13.9	232
SC.L2-3.13.10	234
SC.L2-3.13.11	237
SC.L2-3.13.12	240
SC.L2-3.13.13	241
SC.L2-3.13.14	243
SC.L2-3.13.15	245
SC.L2-3.13.16	247
SI.L1-3.14.1	249
SI.L1-3.14.2	254
SI.L2-3.14.3	258
SI.L1-3.14.4	261
SI.L1-3.14.5	265
SI.L2-3.14.6	269
SI.L2-3.14.7	274
Appendix A Key terms, acronyms & definitions	277

Overview of the CMMC Implementation Guide

Purpose

The purpose of this Cybersecurity Maturity Model Certification (CMMC) Implementation Guide is to provide detailed guidance on the use of Google Cloud to support customers' **CMMC Version 2.0 Level 2 ("CMMC")** compliance needs to safeguard **Controlled Unclassified Information (CUI)**.

This implementation guide also demonstrates how Google Cloud can enable customers to meet their cybersecurity and compliance goals.

This document provides guidance on Google Cloud capabilities to support customer compliance with the CMMC. Organizations should seek independent legal advice relating to their responsibilities under CMMC. Nothing in this document is intended to provide or be used as a substitute for legal advice.

Audience

The audience for this CMMC Implementation Guide may include any organization seeking compliance with the Defense Federal Acquisition Regulation Supplement (**DFARS**) Clause 252.204-7012 and subsequent DFARS requirements, including DFARS Clause 252.204.7021, National Institute of Standards and Technology (**NIST**) Special Publication (SP) 800-171, and/or CMMC Level 2, including but not limited to the Defense Contract Management Agency (**DCMA**), members of the Defense Industrial Base (**DIB**), federal contractors and subcontractors, organizations seeking assessment (**OSAs**), or other Google Cloud customers, hereto referred to as a **Customer**. Specifically, this guide should be used by individuals or teams responsible for implementing and managing their Google Cloud environment.

Scope

The scope of the CMMC Implementation Guide is limited to the controls found in [NIST SP 800-171 Revision 2](#) for a CMMC Level 1 or 2 implementation by utilizing the CMMC level identifier (i.e., L1, L2).

This guide is limited to a Customer's [Google Cloud Console](#) configuration, specifically for Customers utilizing Google Cloud **Assured Workloads**. Where applicable, additional Google Cloud services such as **Google Compute Engine (GCE)**, **Google Kubernetes Engine (GKE)**, **Artifact Registry**, **Persistent Disk**, **Google Cloud Storage**, **Key Management Service (KMS)**,

Pub/Sub, and **Cloud SQL** are referenced as it relates to the Customer's implementation of CMMC controls. The scope is limited to these services within the Assured Workloads environment and does not extend to additional services, systems, applications, tools, endpoints, or other processes that may fall within the Customer's **CUI boundary**.

Overview of CMMC

CMMC Program

The CMMC Program was established by the Department of Defense (DoD) to formally verify that DoD contractors and subcontractors within the DIB implement appropriate cyber security measures necessary to safeguard Federal Contract Information (FCI) and Controlled Unclassified Information (CUI). To achieve this, the CMMC program leverages the security controls outlined in NIST SP 800-171 Rev 2 Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations.

After several years of deliberation, the CMMC Rule ([32 CFR Part 170](#)) formally went into effect on December 16, 2024. The table below defines the [requirements for each CMMC Level and assessment type](#):

CMMC Level	Security Requirements	Assessment Requirement	Assessment Frequency
Level 1 (Self)	15 requirements from FAR clause 52.204-21	Self assessment	Annual
Level 2 (Self)	110 requirements from NIST 800-171 Rev 2	Self assessment	Every 3 years
Level 2 (Certification)	110 requirements from NIST SP 800-171 Rev 2	C3PAO assessment	Every 3 years
Level 3 (Certification)	110 requirements from NIST SP 800-171 Rev 2	DCMA assessment	Every 3 years

	24 requirements from NIST SP 800-172		
--	--	--	--

Note: NIST SP 800-171 Rev 2 control requirements are a subset of [NIST SP 800-53](#), which provide FedRAMP control baselines. See Appendix D of NIST SP 800-171 for a complete control mapping of NIST 800-171 controls to the relevant NIST SP 800-53 security controls. These controls have been assessed and authorized as part of the existing [Google Cloud FedRAMP authorization](#).

Cloud Service Provider (CSP) Requirements

In addition to the security requirements listed above, the [DFARS 252.204-7012](#) clause (“DFARS”) also dictates that any contractor intending to leverage an external Cloud Service Provider (CSP) to store, process, or transmit CUI must ensure that the CSP meet a [FedRAMP Moderate Equivalency](#) and complies with the additional requirements in the clause.

Google Cloud infrastructure and specific Google Cloud Services Offerings (CSOs) have obtained a [FedRAMP High Authority to Operate \(ATO\)](#), and therefore can support Customers seeking compliance with DFARS and CMMC.

Key CMMC Requirements

Customers can consider the following during the CMMC journey:

- 1. Understand requirements:** Know the CMMC levels, required levels within contracts, and respective control requirements. Understand the individual objectives of each requirement which must be satisfied to meet the control.
- 2. Define scope and boundary:** Identify CUI and how CUI is received, handled, stored and managed within Google Cloud, as well as other systems, applications, etc. CMMC has provided an official [CMMC Level 2 Scoping Guide](#) to better identify CUI assets and define a boundary.
- 3. Conduct a gap analysis:** Conduct a gap analysis against your defined boundary to identify and address security gaps in the technical controls. Consider using the official [CMMC Level 2 Assessment Guide](#) to mimic how an assessor would approach testing the controls.
- 4. Leverage technologies:** Configuring key security features in Google Cloud to enhance security and implement controls.

Overview of Google Cloud

Purpose

Google Cloud is a public cloud computing platform that offers a wide range of services for compute, storage, machine learning, big data analytics, and more. It allows users to build, deploy, and run applications on Google's infrastructure. When developing applications or running workloads on Google Cloud, customers can mix and match the available Google Cloud [key services](#) into combinations that provide the infrastructure they need.

Assured Workloads

Google Cloud [Assured Workloads](#) is a service designed to help organizations enforce security and compliance controls on their cloud workloads, particularly those subject to specific regulatory, regional, or data sovereignty requirements.

Assured Workloads helps organizations meet stringent compliance standards by applying predefined control packages to specific Google Cloud folders. This ensures that resources within those folders adhere to the necessary regulations.

Google Cloud Assured Workloads also allows Customers to specify the geographic regions where their data can be stored and processed, helping meet data residency laws and requirements. It also offers controls to manage service provider access to data, supporting data sovereignty needs. It creates secure and controlled boundaries within Google Cloud, restricting resource deployment and configuration to comply with chosen frameworks such as CMMC.

Customers of these services must select the appropriate Assured Workloads Data Boundary (FedRAMP Moderate or FedRAMP High), to ensure that the implementation of inheritable CMMC controls aligns with the specific requirements of their environment. Customers should also consider any data residency or citizenship requirements that may exist within their CUI boundary in support of their contractual commitments.

Customer Responsibility Matrix (CRM)

Google maintains a Google Cloud CMMC **Customer Responsibility Matrix (CRM)** to define which controls can be inherited or partially inherited from Google and which are a customer's responsibility. Inheritable controls have been audited by a FedRAMP 3PAO and mapped relevant to NIST SP 800-171 controls.

Google is responsible for the design, development, release and maintenance of the cloud services platform it provides and the performance and availability of the related common infrastructure on which the services are provided. Service management relies on the

underlying infrastructure and related data security mechanisms to provide reliable and secure services to customers.

Once a Google Cloud environment has been deployed, a Customer can use the key services available in the Google Cloud Console to configure and customize services and meet control requirements within domains such as Access Control, Audit and Accountability, etc. APIs are also available to integrate Google Cloud environments with existing client infrastructure or third-party service providers.

For specifics on Google implementation on inherited/partially inherited controls and customer responsibility, reach out to your Google customer representative and request the CRM.

How to Use this Guide

Control Implementation Tables

This section provides an overview of how to interpret the control implementation tables below in order to implement and manage the Google Cloud security controls.

For each control that requires customer implementation, a table will include the following content

- **Control Domain:** The CMMC control domain, based on NIST SP 800-171
- **Control #:** The CMMC control identifier, composed of three distinct elements, the domain, the CMMC Level (e.g., L1, L2), and security requirement number aligned to NIST SP 800-171
- **Control Description:** The CMMC control language descriptor
- **Key Services:** Recommended tools and features within Google Cloud that can be utilized by the Customer as part of the control implementation. Additional or alternative tools and features may also be utilized and are up to the Customer to determine the best approach for their CUI Boundary.
- **Control Responsibility:** Control responsibility as defined in the CRM. Controls identified as “Shared” and “Customer” are controls that require specific implementation by the Customer.
- **Customer Implementation Description:** A step-by-step description for the Customer to follow to utilize the Key Services to support their CMMC control. The guidance takes into consideration the [NIST SP 800-171A](#) control objectives.
- **Supplemental Guidance:** Where additional, more detailed, implementation guidance exists, a reference will be provided for the Customer for further exploration and reading

An example control implementation table is depicted below:

Control Domain	Access Control		
Control #	AC.L1-3.1.1		
Control Description	Limit system access to authorized users, processes acting on behalf of authorized users, or devices (including other systems)		
Key Services	IAM	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
As a Customer of Google Cloud, you should consider - Step 1 - Step 2 - Step 3			
Supplemental Guidance			
Links to additional documentation			

Control Categories

This guide has categorized the CMMC controls into three (3) distinct categories.

1. Controls requiring customer implementation in Google Cloud

Specific to the Customer's Google Cloud environment, the Customer is responsible for leveraging and configuring underlying services to enable compliance with the control. While the Customer Implementation Description provides technical implementation, there may be additional administrative requirements that exist outside of the scope of Google Cloud (e.g., policies and procedures) the Customer may need to deploy to meet the control.

The controls in this category include:

AC.L1-3.1.1	AC.L1-3.1.19	CM.L2-3.4.5	MP.L2-3.8.4	SC.L2-3.13.10
AC.L1-3.1.2	AC.L1-3.1.20	CM.L2-3.4.6	MP.L2-3.8.9	SC.L2-3.13.11
AC.L2-3.1.3	AC.L1-3.1.22	CM.L2-3.4.7	PS.L2-3.9.2	SC.L2-3.13.13

AC.L2-3.1.4	AU.L2-3.3.1	IA.L1-3.5.1	RA.L2-3.11.1	SC.L2-3.13.14
AC.L2-3.1.5	AU.L2-3.3.2	IA.L1-3.5.2	RA.L2-3.11.2	SC.L2-3.13.15
AC.L2-3.1.6	AU.L2-3.3.3	IA.L2-3.5.3	RA.L2-3.11.3	SC.L2-3.13.16
AC.L2-3.1.7	AU.L2-3.3.4	IA.L2-3.5.4	CA.L2-3.12.3	SI.L2-3.14.1
AC.L2-3.1.8	AU.L2-3.3.5	IA.L2-3.5.5	SC.L1-3.13.1	SI.L2-3.14.2
AC.L2-3.1.9	AU.L2-3.3.6	IA.L2-3.5.6	SC.L1-3.13.2	SI.L2-3.14.3
AC.L2-3.1.10	AU.L2-3.3.7	IA.L2-3.5.7	SC.L2-3.13.3	SI.L2-3.14.4
AC.L2-3.1.11	AU.L2-3.3.8	IA.L2-3.5.8	SC.L2-3.13.4	SI.L2-3.14.5
AC.L2-3.1.12	AU.L2-3.3.9	IA.L2-3.5.9	SC.L2-3.13.5	SI.L2-3.14.6
AC.L2-3.1.13	CM.L2-3.4.1	IA.L2-3.5.10	SC.L2-3.13.6	SI.L2-3.14.7
AC.L2-3.1.14	CM.L2-3.4.2	IA.L2-3.5.11	SC.L2-3.13.7	
AC.L2-3.1.15	CM.L2-3.4.3	IR.L2-3.6.1	SC.L2-3.13.8	
AC.L1-3.1.18	CM.L2-3.4.4	MP.L2-3.8.2	SC.L2-3.13.9	

2. Controls requiring implementation outside of Google Cloud

These controls do not apply to a Google Cloud implementation. The Customer's CUI boundary may contain systems, applications, facilities, or tools outside of Google Cloud; it is the sole responsibility of the Customer to decide how to address controls for the organizational systems that fall outside the Google Cloud environment.

The controls in this category include:

AT.L2-3.2.1	CM.L2-3.4.9	MP.L2-3.8.7	CA.L2-3.12.1	
AT.L2-3.2.2	IR.L2-3.6.2	MP.L2-3.8.8	CA.L2-3.12.2	
AT.L2-3.2.3	IR.L2-3.6.3	PS.L2-3.9.1	CA.L2-3.12.4	

CM.L2-3.4.8	MP.L2-3.8.1	PE.L2-3.10.6	SC.L2-3.13.12	
-------------	-------------	--------------	---------------	--

3. Controls natively implemented by Google

Natively implemented controls are those the Customer inherits directly from Google. While Google has implemented the control, there may be opportunities for Customers to deploy security enhancements via Key Services.

The Customer's CUI boundary may contain systems, applications, facilities, or tools outside of Google Cloud; it is the sole responsibility of the Customer to decide how to address controls for the organizational systems that fall outside the Google Cloud environment.

The controls in this category include:

AC.L2-3.1.16	MA.L2-3.7.2	MA.L2-3.7.6	PE.L1-3.10.1	PE.L1-3.10.5
AC.L2-3.1.17	MA.L2-3.7.3	MP.L1-3.8.3	PE.L2-3.10.2	
AC.L2-3.1.21	MA.L2-3.7.4	MP.L2-3.8.6	PE.L1-3.10.3	
MA.L2-3.7.1	MA.L2-3.7.5	MP.L2-3.8.6	PE.L1-3.10.4	

CMMC L1 and L2 Control Implementation Guidance

Control Domain	Access Control		
Control #	AC.L1-3.1.1		
Control Description	Limit system access to authorized users, processes acting on behalf of authorized users, or devices (including other systems).		
Key Services	- Cloud Identity / Google Workspace - IAM - Service Accounts - IAP - OS Login - VPC Firewall Rules - CAA - Google Workspace MDM - Endpoint Verification	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			

Google Cloud Customers are responsible for:

- a. Identifying authorized users of their Google Cloud environment
- b. Identifying processes acting on behalf of authorized users which may access their Google Cloud environment
- c. Identifying devices (and other systems) authorized to connect to their Google Cloud environment
- d. Limiting Google Cloud access to authorized users;
- e. Limiting Google Cloud access to processes acting on behalf of authorized users
- f. Limiting Google Cloud access to authorized devices (including other systems)

When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control:

- Cloud Identity / Google Workspace
- Identity and Access Management (IAM) (including Conditions)
- Service Accounts
- Identity-Aware Proxy (IAP)
- OS Login
- VPC Firewall Rules
- Context-Aware Access (CAA)
- Google Workspace Mobile Device Management (MDM) (for device authorization signals)
- Endpoint Verification

A description of relevant features and implementation guidance is included below.

Cloud Identity / Google Workspace: Use these services to manage the lifecycle of authorized user identities.

- **Manage Users and Groups**
 1. Navigate to the Google Admin console (admin.google.com).
 2. Use **Directory > Users** to create, suspend, or delete user accounts based on authorization status.
 3. Use **Directory > Groups** to create groups for efficient role assignment within IAM. Ensure group memberships reflect authorized access needs.
 4. Implement strong authentication policies (Password Management, 2-Step Verification) under **Security > Authentication**.

Identity and Access Management (IAM): Authorize access for users and processes (service accounts) to specific Google Cloud resources.

- **Grant Least Privilege Roles**

1. Navigate to **IAM & Admin > IAM** in the Google Cloud Console.
2. Select the resource scope (Organization, Folder, Project).
3. Click **Grant Access**.
4. Enter authorized principals (users, groups, service accounts).
5. Assign the most specific predefined roles or custom roles that grant only the necessary permissions. Avoid basic roles (Owner, Editor, Viewer) for routine access.
6. Regularly review grants using IAM or Policy Analyzer (**IAM & Admin > Policy Analyzer**).

Service Accounts: Represent authorized processes (applications, VMs).

- **Manage Service Accounts**

1. Navigate to **IAM & Admin > Service Accounts**.
2. Create service accounts with specific purposes. Avoid using default service accounts where possible.
3. Grant minimal necessary IAM roles directly to the service account via the **IAM & Admin > IAM** page.
4. Securely manage keys or prefer keyless authentication methods (attaching to VMs, Workload Identity Federation).

Identity-Aware Proxy (IAP): Control access to web applications and VMs based on user identity and device context (if CAA is used).

- **Configure IAP**

1. Navigate to **Security > Identity-Aware Proxy**.
2. Enable IAP for target resources (App Engine, Load Balancer backends, Compute Engine instances for TCP).
3. Grant the appropriate IAP-secured... role to authorized users/groups via the IAP resource's IAM policy.
4. Optionally apply Context-Aware Access levels (see below) to restrict access based on device authorization status.

OS Login: Control SSH access to Linux VMs using IAM identities.

- **Configure OS Login**

1. Enable OS Login via metadata (**Compute Engine > Metadata > enable-oslogin = TRUE**).
2. Grant roles/compute.osLogin or roles/compute.osAdminLogin IAM roles to authorized users for SSH access.
3. Disable SSH key access via metadata to ensure IAM/OS Login is the sole path for authorization.

VPC Firewall Rules: Control network-level access, limiting which devices/networks can reach resources.

- **Configure Firewall Rules**

1. Navigate to **VPC network > Firewall**.
2. Implement default-deny ingress/egress rules.
3. Create specific Allow rules permitting traffic only from authorized sources (e.g., specific IP ranges representing authorized networks/devices, IAP IP ranges) to necessary destinations and ports.

Context-Aware Access (CAA): Control access based on user identity and device context (posture).

- **Create Device-Based Access Levels**

1. Ensure devices are managed (e.g., via Workspace MDM) and/or have Endpoint Verification deployed.
2. Navigate to **Security > Access Context Manager**.
3. Create Access Levels based on device attributes (e.g., device.is_compliant_device == true, device.encryption_status == "ENCRYPTED", specific OS requirements) representing authorized devices.

- **Apply Access Levels**

1. Apply these Access Levels to IAP-secured resources or other services supporting CAA to ensure only authorized devices connect.

Google Workspace MDM & Endpoint Verification: Provide signals about device authorization status to CAA. MDM manages mobile devices; Endpoint Verification reports status for desktops/laptops.

- **Configure MDM/Endpoint Verification** (Primarily in Admin Console and via client deployment)

1. Set up MDM policies in the Admin Console (**Devices > Mobile & endpoints**).
2. Deploy Endpoint Verification extension/app to laptops/desktops.

3. Ensure devices are enrolled and reporting status correctly for CAA evaluation.

Additional Considerations

- **Authentication:** Strong authentication (MFA) is crucial for verifying user identity before authorization is checked (covered by controls like 3.5.3).
- **Regular Reviews:** Periodically review user accounts, group memberships, IAM policies, firewall rules, and device compliance rules to ensure only currently authorized entities have access.
- **Least Privilege:** Consistently apply the principle of least privilege across all access control mechanisms.

Supplemental Guidance

- [Overview of Cloud Identity](#)
- [IAM overview | IAM Documentation | Google Cloud](#)
- [Service accounts overview | IAM Documentation | Google Cloud](#)
- [Identity-Aware Proxy overview | Google Cloud](#)
- [About OS Login | Compute Engine Documentation | Google Cloud](#)
- [VPC firewall rules | Cloud NGFW](#)
- [Context-Aware Access Overview](#)
- [Chrome Enterprise Premium overview | BeyondCorp Enterprise | Google Cloud](#)
- [Endpoint Verification overview](#)
- [Mobile device management overview \(Google Workspace\)](#)

Control Domain	Access Control		
Control #	AC.L1-3.1.2		
Control Description	Limit system access to the types of transactions and functions that authorized users are permitted to execute.		
Key Services	• IAM • OS Logins	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for:			

- a. Defining the types of transactions and functions that authorized users are permitted to execute in Google Cloud
- b. Limiting access to the defined types of transactions and functions using roles

When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control:

- Identity and Access Management (IAM) (Permissions, Predefined Roles, Custom Roles, Conditions)
- OS Login

A description of relevant features and implementation guidance is included below.

Identity and Access Management (IAM): Limit access to specific functions and transactions by assigning roles containing granular permissions.

- **Assign Least Privilege Predefined Roles**
 1. Navigate to **IAM & Admin > IAM**.
 2. When granting access (**Grant Access**), select predefined roles that most closely match the user's required functions (e.g., roles/compute.networkViewer, roles/storage.objectCreator, roles/cloudsql.client).
 3. Avoid assigning broad basic roles (roles/owner, roles/editor, roles/viewer) unless absolutely necessary, as they grant extensive permissions beyond specific functions.
- **Create and Assign Custom Roles (Key Implementation)**
 1. Navigate to **IAM & Admin > Roles**.
 2. Click **+ Create Role**.
 3. Provide a Title, Description, and ID (e.g., vmOperatorMinimal, billingDataViewer).
 4. Click **+ Add Permissions**.
 5. In the filter, search for and select *only* the specific permissions corresponding to the authorized functions. Permissions often follow the pattern service.resource.verb (e.g., compute.instances.start, compute.instances.stop, billing.accounts.get). Add only what is needed.
 6. Click **Add**.
 7. Click **Create**.
 8. Assign this custom role to authorized users/groups via the **IAM & Admin > IAM** page instead of broader predefined roles.
- **Use IAM Conditions (Optional Refinement)**

1. When assigning a role (**IAM & Admin > IAM > Edit Principal**), click **Add condition**.
2. Define conditions based on time, source IP, or resource attributes to further restrict *when* or *how* permitted functions can be executed. For example, allow the `compute.instances.delete` function only from a specific admin subnet.

OS Login: Control the level of function permitted within a Linux VM SSH session based on IAM authorization.

- **Assign OS Login Roles**

1. Enable OS Login for the project or specific VMs.
2. Grant the roles/compute.osAdminLogin role *only* to users authorized to perform administrative functions (e.g., run `sudo` commands) within the VM via SSH.
3. Grant the standard roles/compute.osLogin role to users who only need non-administrative functions within their SSH session. This directly limits their ability to execute privileged OS-level functions.

Additional Considerations

- **Least Privilege:** The core principle here is ensuring users/processes have *only* the permissions (functions) they absolutely need. Custom roles are often necessary to achieve this effectively.
- **Permissions Mapping:** Understand that actions in the Cloud Console or via `gcloud` correspond to specific IAM permissions (API calls). Use the documentation or Policy Simulator to identify needed permissions.
- **Regular Audits:** Periodically review assigned roles and the permissions within custom roles using IAM tools and Policy Analyzer (**IAM & Admin > Policy Analyzer**) to ensure they haven't become excessive over time (permission creep).

Supplemental Guidance

- [IAM overview | IAM Documentation | Google Cloud](#)
- [IAM roles and permissions index | IAM Documentation | Google Cloud](#)
- [Create and manage custom roles | IAM Documentation | Google Cloud](#)
- [Overview of IAM Conditions | IAM Documentation | Google Cloud](#)
- [Manage conditional role bindings | IAM Documentation | Google Cloud](#)
- [Introduction to the Organization Policy Service | Resource Manager Documentation | Google Cloud](#)
- [Creating and managing organization policies | Resource Manager Documentation |](#)

[Google Cloud](#)

- [Service accounts overview | IAM Documentation | Google Cloud](#)
- [Best practices for using service accounts | IAM Documentation | Google Cloud](#)
- [Overview of role recommendations | Policy Intelligence | Google Cloud](#)
- [Analyze allow policies | Policy Intelligence | Google Cloud](#)

Control Domain	Access Control								
Control #	AC.L2-3.1.3								
Control Description	Control the flow of CUI in accordance with approved authorizations								
Key Services	• VPC Service Controls • VPC Firewall Rules • IAM • IAP • Cloud Storage & BigQuery Access Controls • Sensitive Data Protection (incl. Cloud DLP)	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Managing and limiting connections between your Google Account and third-parties b. Configuring Data Loss Prevention (DLP) to detect and prevent the loss, leakage, or misuse of CUI in your Google Cloud environment When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • VPC Service Controls • VPC Firewall Rules • Identity and Access Management (IAM) • Identity-Aware Proxy (IAP) • Cloud Storage & BigQuery Access Controls • Sensitive Data Protection (incl. Cloud Data Loss Prevention) A description of relevant features and implementation guidance is included below.									

VPC Service Controls: Create security perimeters around Google managed services to control data flow across boundaries (projects, networks, internet).

- **Configure Perimeters and Rules**

1. Navigate to **Security > VPC Service Controls**.
2. Create a **New Perimeter**.
3. Add the projects containing CUI to the perimeter.
4. Select the Google Cloud services to restrict (e.g., storage.googleapis.com, bigquery.googleapis.com).
5. Define **Ingress Rules**: Specify which sources (identities, networks, access levels) outside the perimeter are authorized to access the restricted services *inside*.
6. Define **Egress Rules**: Specify which destinations or identities outside the perimeter are authorized to receive data *from* the restricted services. Deny unauthorized flows by default.
7. Use **Dry Run** mode initially to test policies without enforcement. Monitor audit logs for violations before enforcing.

VPC Firewall Rules: Control network traffic flow at the IP, port, and protocol level within and between VPCs and external networks.

- **Configure Network Flow Control**

1. Navigate to **VPC network > Firewall**.
2. **Egress Control**: Implement a default-deny egress policy. Create specific Allow rules permitting outbound traffic only to authorized destinations (IPs/ports) needed for CUI flow (e.g., to a specific partner network via VPN). Use network tags or service accounts on VMs to apply rules granularly.
3. **Ingress Control**: Implement a default-deny ingress policy. Allow inbound traffic only from authorized sources (e.g., specific corporate IPs, IAP ranges) to designated resources.
4. **Internal Flow Control**: Create firewall rules between subnets or using tags/service accounts to enforce authorized communication paths *within* the VPC if needed (e.g., restricting database access to specific application tiers).

Identity and Access Management (IAM): Authorize who can access data or configure flow controls.

- **Apply Least Privilege**

1. Grant IAM roles for accessing CUI data stores (Cloud Storage buckets, BigQuery datasets) only to explicitly authorized users, groups, and service accounts.
2. Restrict permissions for configuring flow control mechanisms (e.g., `accesscontextmanager.policyAdmin` for VPC SC, `compute.networkAdmin` for Firewalls) to authorized administrators.

Identity-Aware Proxy (IAP): Control user access flows into web applications and VMs.

- **Authorize Inbound User Flows**

1. Enable IAP (**Security > Identity-Aware Proxy**) for applications or VMs processing CUI accessible from untrusted networks.
2. Grant appropriate IAP-secured... roles only to authorized users/groups, ensuring only they can initiate data flows within the protected resource.

Cloud Storage & BigQuery Access Controls: Provide resource-level control over data access.

- **Configure Data Access**

1. Use fine-grained IAM policies on Cloud Storage buckets and objects. Avoid public access (`allUsers/allAuthenticatedUsers`) for CUI.
2. Use BigQuery Dataset IAM, Authorized Views, Row-Level Security, and Column-Level Security to ensure users can only access/query the specific CUI data they are authorized for, controlling the flow of information presented to them.

Cloud DLP Helps detect CUI in potentially unauthorized flows.

- **Monitor Data Flows**

1. Configure DLP inspection jobs (**Security > Data Loss Prevention**) to scan Cloud Storage buckets or potentially network traffic for CUI patterns.
2. Configure alerts or actions based on findings to identify potential policy violations or misconfigurations allowing unauthorized CUI flow.

Additional Considerations

- **Map Authorized Flows:** Clearly document the intended and authorized paths for CUI data flow before configuring technical controls.
- **Layered Defense:** Combine VPC Service Controls (service-level boundary), VPC Firewalls (network-level boundary), and IAM (identity-level authorization) for comprehensive flow control.

- **Start with Dry Run:** Use VPC Service Controls dry-run mode extensively during initial setup to avoid blocking legitimate business processes. Analyze violation logs carefully.

Supplemental Guidance

- [Overview of VPC Service Controls | Google Cloud](#)
- [Create a service perimeter | VPC Service Controls | Google Cloud](#)
- [VPC firewall rules | Cloud NGFW](#)
- [Identity-Aware Proxy overview | Google Cloud](#)
- [IAM overview | IAM Documentation | Google Cloud](#)
- [Introduction to data governance in BigQuery | Google Cloud](#)
- [Overview of access control | Cloud Storage](#)
- [Sensitive Data Protection Documentation | Google Cloud](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.4		
Control Description	Separate the duties of individuals to reduce the risk of malevolent activity without collusion		
Key Services	• IAM • Google Groups • Resource Manager • Service Accounts • Cloud Build • Cloud Audit Logs	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining the duties of individuals that require separation - Assigning responsibilities to separate individuals - Implementing separation of duties through assigned group and role authorizations When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control: • Identity and Access Management (IAM) (Custom Roles, Predefined Roles, Permissions) • Google Groups (via Google Workspace/Cloud Identity Admin Console)			

- Resource Manager (Folders & Projects)
- Service Accounts
- Cloud Build (Manual Approvals feature)
- Cloud Audit Logs (for verification)

A description of relevant features and implementation guidance is included below.

Identity and Access Management (IAM) - Custom Roles: Enforce Separation of Duties (SoD) in Google Cloud by creating roles with specific, non-overlapping duties.

- **Define and Create Custom Roles for SoD**
 1. Identify the conflicting duties (e.g., Network Configuration vs. Firewall Rule Management).
 2. Navigate to **IAM & Admin > Roles** in the Cloud Console.
 3. Click **+ Create Role** for the first duty (e.g., NetworkConfigurer).
 4. Add *only* the specific permissions needed for that duty (e.g., permissions related to creating VPCs, subnets).
 5. Click **Create**.
 6. Click **+ Create Role** for the second, conflicting duty (e.g., FirewallAdmin).
 7. Add *only* the specific permissions needed for that duty (e.g., permissions related to creating, modifying, deleting firewall rules). Ensure this role *does not* include the permissions granted to NetworkConfigurer.
 8. Click **Create**.
 9. Repeat for all identified duties requiring separation.

Identity and Access Management (IAM) - Role Assignment: Assign the distinct roles created above to different individuals or groups.

- **Assign Separated Roles**
 1. Navigate to **IAM & Admin > IAM**.
 2. Select the appropriate resource scope (Organization, Folder, Project).
 3. Click **Grant Access**.
 4. Enter the principal (user or preferably a Google Group representing a job function, e.g., network-config-team@domain.com).
 5. Assign the specific custom role created for that duty (e.g., NetworkConfigurer).
 6. Click **Save**.
 7. Click **Grant Access** again.

8. Enter a *different* principal (user or group, e.g., firewall-admin-team@domain.com).
9. Assign the specific custom role for the *conflicting* duty (e.g., FirewallAdmin).
10. Click **Save**.
11. **Crucially** Avoid assigning conflicting custom roles (like NetworkConfigurer *and* FirewallAdmin) or conflicting high-privilege predefined roles (like Compute Network Admin *and* Compute Security Admin) to the same individual user or group if SoD is required between those functions.

Google Groups (Managed in Admin Console): Use groups to manage membership for specific duties, making role assignment cleaner.

- **Create Duty-Based Groups**

1. In the Google Admin console (admin.google.com), go to **Directory > Groups**.
2. Create groups corresponding to specific duties (e.g., gcp-billing-managers, gcp-prod-deployers, gcp-security-auditors).
3. Assign authorized users to these groups.
4. In Google Cloud IAM, assign the relevant custom or predefined roles to these groups instead of individual users.

Resource Manager (Folders & Projects): Isolate environments or functions to apply different IAM policies.

- **Structure Resources**

1. Organize projects under Folders representing different environments (e.g., Development, Production) or business units.
2. Apply IAM policies at the Folder level. For example, grant broader roles in the Development folder but highly restricted, separated roles in the Production folder. This prevents developers, for instance, from having conflicting duties in the production environment.

Service Accounts: Use distinct service accounts for different automated tasks.

- **Separate Service Account Duties**

1. Create separate service accounts (**IAM & Admin > Service Accounts**) for distinct stages in automation (e.g., build-robot@..., deploy-robot@...).
2. Grant minimal, specific roles to each service account (e.g., the build robot might only need artifact repository write access, while the deploy robot needs Compute Engine deployment permissions).

Cloud Build (Manual Approvals): Enforce SoD within CI/CD workflows.

- **Configure Deployment Approvals**
 1. In a Cloud Build pipeline YAML file or trigger configuration (**Cloud Build > Triggers > Edit > Approval**), enable the **Require approval** option for critical steps like deploying to production.
 2. Specify authorized approvers (users or groups) who must be different from the user/service account initiating the build/deployment.

Additional Considerations

- **Policy First:** SoD implementation starts with defining organizational policy about which duties must be separated. Google Cloud tools then provide the technical enforcement.
- **Identify Critical Functions:** Focus SoD efforts on high-risk areas like financial controls (billing), security configuration, production deployments, and sensitive data access.
- **Audit Trail:** Use Cloud Audit Logs (**Logging > Logs Explorer**) to review actions and verify that separated duties are being maintained in practice. Look for who performed sensitive actions.
- **Complexity:** Implementing fine-grained SoD can increase operational complexity. Balance the security benefits with operational needs.

Supplemental Guidance

- [IAM overview | IAM Documentation | Google Cloud](#)
- [Create and manage custom roles | IAM Documentation | Google Cloud](#)
- [Resource hierarchy | Resource Manager Documentation | Google Cloud](#)
- [Service accounts overview | IAM Documentation | Google Cloud](#)
- [Cloud Audit Logs Overview](#)

Control Domain	Access Control						
Control #	AC.L2-3.1.5						
Control Description	Employ the principle of least privilege, including for specific security functions and privileged accounts						
Key Services	• IAM • Organization Policies • Service Accounts	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr></table>	☐	Google	☒	Shared
☐	Google						
☒	Shared						

			☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining privileged functions - Identifying non-privileged users/roles - Preventing non-privileged users from executing privileged functions using Google pre-built roles, custom roles or a combination of Google pre-built roles and custom roles When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control: - Identity and Access Management (IAM) - Organizational Policies - Service Accounts A description of relevant features and implementation guidance is included below. Identity and Access Management (IAM): Use IAM to manage who (principals) has what access (roles) to which resources. This is the primary mechanism for authorizing access based on validated requirements. Grant Roles - Navigate to IAM & Admin > IAM in the Google Cloud Console. - Select the appropriate resource scope (Organization, Folder, or Project) from the top dropdown. - Click Grant Access. - In the New principals field, enter the email addresses of users, groups, or service accounts requiring access. Ensure these principals have a validated need for the requested access. - In the Assign roles section, select the predefined or custom role(s) that grant only the necessary permissions for their tasks. Avoid overly broad Basic roles (Owner, Editor, Viewer). (Optional) Click Add condition to create an IAM Condition for time-limited access or resource-specific constraints. - Click Save. Create Custom Roles (if predefined roles are too broad) - Navigate to IAM & Admin > Roles. - Select the resource scope (Organization or Project).			

3. Click **+ Create Role**.
 4. Provide a Title, Description, and unique ID.
 5. Click **+ Add Permissions**.
 6. Filter and select only the specific permissions required for the role, aligning with the validated authorization.
 7. Click **Add**.
 8. Click **Create**.
 9. Grant this custom role following the steps in "Grant Roles".
- **Review Access**
 1. Regularly visit **IAM & Admin > IAM** to review current access grants.
 2. Use **IAM & Admin > Policy Analyzer** to understand who has access to specific resources.
 3. Check **Security > Recommendations Hub** (powered by IAM Recommender) for suggestions to remove unused or excessive permissions.

Organization Policies Use Organization Policies to enforce constraints across your cloud resources, limiting resource configurations and ensuring compliance with authorization boundaries set at a higher level.

- **Set Organization Policies**
 1. Navigate to **IAM & Admin > Organization Policies**.
 2. Select your Organization from the scope selector at the top.
 3. Filter or search for policies relevant to access control (e.g., constraints/iam.allowedPolicyMemberDomains to restrict identities, constraints/compute.vmExternalIpAccess to limit external access, constraints/gcp.resourceLocations to control where resources can be created).
 4. Click on the policy name.
 5. Click **Edit**.
 6. Choose **Customize**.
 7. Configure the policy rules (e.g., Allow or Deny specific values, enforce certain settings). The configuration options depend on the specific policy.
 8. Set **Enforcement** to **Enforce** (or **Replace/Merge** depending on inheritance needs).
 9. Click **Save**.

Service Accounts Use Service Accounts for non-human identities (applications, VMs). Grant them specific, minimal roles based on their function's validated authorization requirements.

- **Create and Assign Roles to Service Accounts**
 1. Navigate to **IAM & Admin > Service Accounts**.
 2. Click **+ Create Service Account**.
 3. Enter a Name, ID, and Description.
 4. Click **Create and Continue**.
 5. In the **Grant this service account access to project** step, select the minimal roles needed for its authorized function. Click **Continue**.
 6. (Optional) Grant user access to the service account if users need to act as the service account.
 7. Click **Done**.
- **Manage Service Account Permissions**
 1. Like user accounts, review the roles granted to service accounts via **IAM & Admin > IAM**. Filter principals by type **Service Account**.
 2. Remove unnecessary roles.
- **Limit Key Usage**
 1. Avoid creating external service account keys whenever possible. Attach service accounts directly to Compute Engine VMs or use Workload Identity Federation.
 2. If keys are necessary, manage them securely and rotate them regularly (**Service Accounts > select account > Keys tab**).

Additional Considerations

- Document the process for requesting, approving, and validating access authorizations before granting permissions in IAM.
- Implement regular access reviews to ensure granted permissions align with current, validated authorizations. Remove access that is no longer required.
- Integrate Cloud Audit Logs with a monitoring system or SIEM to track IAM policy changes and access patterns.

Supplemental Guidance

- [Policy Analyzer for allow policies | Policy Intelligence | Google Cloud](#)
- [Use IAM securely | IAM Documentation | Google Cloud](#)
- [Privileged Access Manager overview | IAM Documentation | Google Cloud](#)
- [Overview of role recommendations | Policy Intelligence | Google Cloud](#)

- [Introduction to the Organization Policy Service | Resource Manager Documentation | Google Cloud](#)
- [Creating and managing organization policies | Resource Manager Documentation | Google Cloud](#)
- [Service accounts overview | IAM Documentation | Google Cloud](#)
- [Best practices for using service accounts | IAM Documentation | Google Cloud](#)
- [IAM overview | IAM Documentation | Google Cloud](#)
- [IAM roles and permissions index | IAM Documentation | Google Cloud](#)
- [Create and manage custom roles | IAM Documentation | Google Cloud](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.6		
Control Description	Use non-privileged accounts or roles when accessing nonsecurity functions		
Key Services	• Cloud Identity / Google Workspace • IAM	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining nonsecurity functions - Requiring users to use non-privileged accounts or roles when accessing nonsecurity functions When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: - Cloud Identity / Google Workspace - Identity and Access Management (IAM) A description of relevant features and implementation guidance is included below. Cloud Identity / Google Workspace: Use these services to manage user identities. To meet this control, ensure users who require privileged access also have a standard, non-privileged account for routine tasks, or clearly defined roles are used to separate functions. • Provision Separate Accounts (if applicable)			

1. Navigate to the Google Admin console (admin.google.com).
2. Go to **Directory > Users**.
3. For each user requiring administrative/privileged access in Google Cloud, ensure they have:
 - One account intended for administrative tasks (e.g., admin-user@domain.com). This account will be granted privileged IAM roles.
 - A separate, standard account for daily non-security functions like email and document collaboration (e.g., user@domain.com). This account should have minimal or no privileged IAM roles in Google Cloud.
4. Instruct users to log in with the appropriate account based on the task they are performing.

- **Manage Groups for Role Assignment**

1. In the Admin console, go to **Directory > Groups**.
2. Create distinct groups for privileged access (e.g., gcp-admins@domain.com) and standard access (e.g., gcp-users@domain.com).
3. Assign the corresponding user accounts (privileged or standard) to these groups. These groups can then be used efficiently within IAM.

Identity and Access Management (IAM): Use IAM to assign roles with specific permissions to the user accounts or groups defined in Cloud Identity/Workspace. This enforces the separation of privileged and non-privileged functions.

- **Assign Roles to Differentiated Accounts/Groups**

1. Navigate to **IAM & Admin > IAM** in the Google Cloud Console.
2. Select the appropriate resource scope (Organization, Folder, Project).
3. Click **Grant Access**.
4. In **New principals**:
 - Enter the email address of the privileged user account or privileged group (e.g., admin-user@domain.com or gcp-admins@domain.com).
 - Assign the necessary privileged roles (e.g., Organization Administrator, Security Admin, custom admin roles). Apply least privilege even for admin accounts.
 - Click **Save**.
5. Click **Grant Access** again.
6. In **New principals**:
 - Enter the email address of the standard user account or standard user group (e.g., user@domain.com or gcp-users@domain.com).

- Assign minimal, non-privileged roles (e.g., Compute Viewer, Project Viewer, or custom roles with only basic access needed for non-security tasks). Avoid granting Editor or Owner roles to standard accounts used for daily tasks.
 - Click **Save**.
- **Utilize Custom Roles for Granularity**
 1. If predefined roles are too broad for either privileged or non-privileged functions, create custom roles (**IAM & Admin > Roles > + Create Role**).
 2. Define roles specifically for non-security tasks (e.g., "Application User Role") with minimal permissions.
 3. Define roles for specific administrative functions (e.g., "Network Admin Role," "Billing Admin Role") granting only the required privileged permissions.
- **Review Permissions Regularly**
 1. Periodically review the roles assigned to both standard and privileged accounts/groups in **IAM & Admin > IAM** to ensure they remain appropriate and adhere to the principle of least privilege and account usage policies.

Additional Considerations

- Clearly document the policy requiring the use of non-privileged accounts for non-security functions.
- Implement training to ensure users understand which account or role to use for different tasks.
- Use Cloud Audit Logs (**Logging > Logs Explorer**, filter by principal email and look for API calls) to monitor activities performed by privileged versus non-privileged accounts. This helps verify that privileged accounts are not being used for routine, non-security related tasks.

Supplemental Guidance

- [Google Cloud Well-Architected Framework | Cloud Architecture Center](#)
- [IAM overview | IAM Documentation | Google Cloud](#)
- [IAM roles and permissions index | IAM Documentation | Google Cloud](#)
- [Create and manage custom roles | IAM Documentation | Google Cloud](#)
- [Privileged Access Manager overview | IAM Documentation | Google Cloud](#)
- [Cloud Audit Logs overview](#)
- [Google Workspace Admin Help](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.7		
Control Description	Prevent non-privileged users from executing privileged functions and audit the execution of such functions		
Key Services	- IAM - Cloud Audit Logs	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining privileged functions - Identifying non-privileged users - Preventing non-privileged users from executing privileged functions by assigning them to the correct role When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: - Identity and Access Management (IAM) - Cloud Audit Logs (part of Cloud Logging) A description of relevant features and implementation guidance is included below. Identity and Access Management (IAM) IAM is the primary mechanism for <i>preventing</i> non-privileged users from executing privileged functions. By assigning roles with specific permissions, you control who can perform which actions. Assign Least Privilege Roles - Follow the guidance provided for control 3.1.5 and 3.1.6 to assign roles. - Critically evaluate the permissions within any role assigned to non-privileged users. Ensure these roles <i>do not</i> contain permissions for actions considered privileged (e.g., *.setIamPolicy, *.delete, *.create, *.update on critical infrastructure or security services like IAM, KMS, Security Command Center, Logging configurations). - Use predefined roles cautiously. Roles like Viewer roles generally prevent privileged actions. Avoid granting Editor or Owner roles to non-privileged users. Use Custom Roles for Precision			

1. If predefined roles grant unnecessary privileged permissions alongside needed non-privileged ones, create custom roles (**IAM & Admin > Roles > + Create Role**).
 2. Build custom roles containing *only* the non-privileged permissions required for a user's job function.
 3. Assign highly privileged permissions (like Organization Administrator, Project Owner, or specific admin roles like Security Admin, Network Admin) *only* to designated administrative accounts.
- **Regularly Review IAM Policies**
 1. Use **IAM & Admin > IAM** and **Policy Analyzer** to periodically review who has which permissions, ensuring non-privileged users have not inadvertently gained access to privileged functions.

Cloud Audit Logs Cloud Audit Logs capture the execution of functions, including privileged ones, providing visibility and accountability. Admin Activity audit logs are crucial for this control.

- **Verify Admin Activity Logging**
 1. Admin Activity logs are enabled by default for most Google Cloud services and cannot be disabled. They record actions that modify resource configuration or metadata. These actions are typically the privileged functions referred to by this control.
 2. There is no cost for Admin Activity logs, and they have a default retention period (consult documentation for current retention).
- **View Execution of Privileged Functions**
 1. Navigate to **Logging > Logs Explorer** in the Google Cloud Console.
 2. Build a query to view relevant activities. To see actions performed by administrative users or involving sensitive permissions
 - Filter by Log Name Select **activity** (Admin Activity).
 - Filter by Resource Specify projects, folders, or the organization if needed.
 - Filter by Principal Enter the email of an administrative user or service account.
 - Filter by Method Name Search for specific API methods associated with privileged functions (e.g., google.iam.admin.v1.SetIamPolicy, google.cloud.resourcemanager.v3.Projects.DeleteProject, google.compute.v1.Instances.Delete).
 3. Example Query

SQL

```
logName="projects/[PROJECT_ID]/logs/cloudaudit.googleapis.com%2Factivity"
protoPayload.methodName "SetIamPolicy" OR protoPayload.methodName
".Delete"
```

(Replace [PROJECT_ID] with your project ID or adjust the logName path for folder/organization logs)

Review the log entries to see which principal performed the action, on which resource, and when.

- **Configure Log Sinks (Optional but Recommended)**
 1. For long-term retention beyond the default period, or for analysis in other tools (like Security Command Center Premium, BigQuery, Splunk), configure log sinks.
 2. Navigate to **Logging > Log Router**.
 3. Click **Create Sink**.
 4. Name the sink, choose the destination (Cloud Storage bucket, BigQuery dataset, Pub/Sub topic, Splunk), and define an inclusion filter (you can filter for only Admin Activity logs if desired, though exporting all audit logs is often recommended).
 5. Click **Create Sink**.

Additional Considerations

- Ensure that permissions required to modify audit logging configurations (e.g., creating sinks, modifying exclusions) are themselves treated as privileged and restricted appropriately using IAM.
- Integrate log reviews into regular operational security procedures. Use alerting capabilities in Cloud Monitoring or Security Command Center to notify on specific high-risk privileged actions.

Supplemental Guidance

- [IAM Overview](#)
- [Understanding Roles](#)
- [Create and manage custom roles](#)
- [Cloud Audit Logs Overview](#)
- [Viewing Audit Logs \(Logs Explorer\)](#)

- [Overview of Log Sinks](#)
- [Querying Admin Activity Logs](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.8		
Control Description	Limit unsuccessful logon attempts		
Key Services	• Google Account Security Features • Cloud Identity / Google Workspace • Cloud Audit Logs • Identity Platform	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining the means of limiting unsuccessful logon attempts - Implementing the means of limiting unsuccessful logon attempts When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control: • Google Account Security Features (inherent) • Cloud Identity / Google Workspace (for monitoring and related security settings like MFA) • Cloud Audit Logs (for monitoring) • Identity Platform (for application-specific configuration) A description of relevant features and implementation guidance is included below. Google Account Security Features: Employ sophisticated, adaptive protections to detect and mitigate brute-force login attempts against Google Accounts used to access the Cloud Console. These include detecting suspicious patterns, requiring CAPTCHAs, sending security alerts, and potentially implementing temporary account lockouts. You should consider 2-step verification (aka, two-factor authentication) when implementing this control. After you turn on 2-step verification, you need to complete a second step to verify it's you if you choose to sign in with a password. To help protect your account, Google will ask you to complete a specific second step. • Implementation			

- These features are inherent to the Google identity platform and are not directly configured with specific thresholds (e.g., "lock after 5 attempts") by administrators in the Cloud Console or Admin Console for standard sign-ins. Reliance is placed on Google's automated protections. Strong password policies and MFA are the primary customer configurations to enhance this protection.

Cloud Identity / Google Workspace (Admin Console): While not directly setting logout thresholds for the Google sign-in, the Admin Console allows monitoring and related security configurations.

- **Monitor Login Activity**

1. Navigate to the Google Admin console (admin.google.com).
2. Go to **Reporting > Reports > Security**. Review reports like "Suspicious login activity," "Account status," and others under "Aggregate reports" for insights into failed attempts or compromised accounts.
3. Go to **Reporting > Audit and investigation > Admin log events** or **User log events** (specifically login events) for detailed logs.

- **Enforce Strong Authentication (Related Controls)**

1. Navigate to **Security > Authentication > Password management** to enforce strong password requirements.
2. Navigate to **Security > Authentication > 2-Step Verification** to enforce MFA, which significantly mitigates the risk associated with compromised passwords or brute-force attempts.

Cloud Audit Logs: Provide detailed records of login attempts, which can be monitored for excessive failures.

- **Monitor Login Events**

1. Navigate to **Logging > Logs Explorer** in the Google Cloud Console.
2. Query for login-related events. Login audit logs are typically found under the Organization or specific projects depending on context. The specific log names and methods might vary, but look for logs related to login or signin. Common patterns include
 - Log Name containing `cloudaudit.googleapis.com/activity` or potentially specific service login logs.
 - Method names like `google.login.LoginService.loginFailure` or similar indicators of failed attempts. Consult the audit logging documentation for the specific services involved (e.g., IAP, specific Google Cloud services).
 - Example filter (may need adjustment based on exact log source)

SQL

logName

"organizations/[YOUR_ORG_ID]/logs/cloudaudit.googleapis.com%2Fdata_access"

protoPayload.methodName = "google.login.LoginService.loginFailure"

3. Configure alerts in Cloud Monitoring based on these logs (e.g., alert if the count of loginFailure events for a single user exceeds a threshold in a short time).

Identity Platform (for Custom Applications): If you use Identity Platform to manage users for your own applications (which might interact with Google Cloud), you can configure limits.

- **Configure Sign-in Limits**
 1. Navigate to **Security > Identity Platform** in the Google Cloud Console.
 2. Go to **Settings > Security**.
 3. Look for options related to "Account lockout" or "Brute force protection." Configuration details may vary based on the sign-in providers enabled (e.g., Email/Password). Adjust settings like the number of failed attempts allowed and lockout duration according to your policy. Refer to Identity Platform documentation for specific provider settings.

Additional Considerations

- While Google's built-in protections are robust, monitoring login failures via Cloud Audit Logs or Admin Console reports provides visibility specific to your organization.
- Ensure a clear process exists for users to regain access if their account is locked due to repeated failed attempts (legitimate or otherwise), often involving an administrator or defined recovery procedures.
- Integrate monitoring data with Security Command Center or an external SIEM for centralized analysis and alerting on potential brute-force attacks.

Supplemental Guidance

- [Make your account more secure - Google Account Help](#)
- [Data retention and lag times - Google Workspace Admin Help](#)
- [Admin audit log events \(Admin Console\)](#)
- [Cloud Audit Logs overview](#)
- [Identity Platform Documentation | Google Cloud](#)
- [Authentication | Identity Platform Documentation | Google Cloud](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.9		
Control Description	Provide privacy and security notices consistent with applicable CUI rules		
Key Services	- IAP - Compute Engine - Google Workspace/Cloud Identity	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Identifying privacy and security notices required by CUI-specified rules consistent with the specific CUI category - Displaying privacy and security notices When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control: - Identity-Aware Proxy (IAP) - Compute Engine (for OS-level banner configuration) - Google Workspace/Cloud Identity (for overall policy communication and limited sign-in customization) A description of relevant features and implementation guidance is included below. Identity-Aware Proxy (IAP): Secure access to web applications, VMs (via SSH/RDP), and GKE applications. You can configure the IAP consent screen, which appears after Google sign-in but before access is granted to the resource, to display your notice. Configure IAP OAuth Consent Screen - Navigate to APIs & Services > OAuth consent screen in the Google Cloud Console. - Select the appropriate User Type (Internal or External). Note Modifying the consent screen often requires domain verification and potentially app verification if external users are involved. - Click Edit App. - Under App information, fill in required details (App name, User support email, Developer contact). - Under App domain, provide links for			

- **Application home page** (Required) Can point to an internal policy page.
 - **Application privacy policy** Link to your organization's privacy policy. The content here supports the notice requirement.
 - **Application terms of service** Link to your organization's terms of service or acceptable use policy. The content here is the primary notice users agree to. Ensure this document contains the required elements (monitoring, authorized use, CUI mention, etc.).
6. Review Scopes (usually automatically configured for IAP).
 7. Review Optional info if needed.
 8. Save and Continue through the steps. The linked Terms of Service and Privacy Policy effectively serve as the notice presented on the consent screen.

Compute Engine (OS-Level Banners): For direct SSH access to Linux VMs (including those managed by OS Login), you can configure a pre-login banner.

- **Configure SSH Banner**

1. Connect to the Linux VM instance using SSH.
2. Edit the SSH daemon configuration file (typically `/etc/ssh/sshd_config`) using a text editor (like nano or vim) with root privileges (sudo).
3. Find the line `#Banner none` or similar. Uncomment it (remove the #) and change none to the path of your banner file, usually `/etc/issue.net`.
4. `Banner /etc/issue.net`
5. Save the configuration file (Ctrl+O, Enter in nano; wq in vim).
6. Create or edit the banner file (`/etc/issue.net`) with root privileges
7. Paste your approved privacy and security notice text into this file.
8. Save the banner file.
9. Restart the SSH service for changes to take effect
10. **Automation** This process can be automated for multiple VMs using startup scripts defined in instance templates or Compute Engine metadata, or through configuration management tools (Puppet, Chef, Ansible, Terraform).

Google Workspace/Cloud Identity: While not providing a direct banner within the Cloud Console itself, these services are central to user identity and policy communication.

- **Policy Communication:** Rely on organizational policies (Acceptable Use Policy, CUI Handling Policy) communicated during employee onboarding and accessible via internal portals. Ensure these policies cover the required notice elements.
- **Sign-in Page Customization (Limited)**
 1. In the Google Admin console (admin.google.com), navigate to **Account > Account settings > Personalization**.

2. You can add a logo and potentially customize some aspects of the sign-in experience, but adding extensive legal text directly to the Google sign-in page is generally not supported. You can link to terms/privacy policies here if desired, reinforcing the notice.

Additional Considerations

- The main Google Cloud Console login page (console.cloud.google.com) uses the standard Google sign-in, which does not support custom banners. The primary technical enforcement points within Google Cloud are IAP-protected resources and direct OS access.
- Ensure the notice text is consistent across different systems (IAP, SSH banners) and accurately reflects organizational policies regarding monitoring, data handling (CUI), and consequences of misuse.
- Keep records of the approved notice text and where it is implemented.

Supplemental Guidance

- [Manage OAuth Clients - Google Cloud Platform Console Help](#)
- [Identity-Aware Proxy overview | Google Cloud](#)
- [About startup scripts | Compute Engine Documentation | Google Cloud](#)
- [About OS Login | Compute Engine Documentation | Google Cloud](#)

Control Domain	Access Control								
Control #	AC.L2-3.1.10								
Control Description	Use session lock with pattern-hiding displays to prevent access and viewing of data after a period of inactivity								
Key Services	• Google Workspace Session Length Controls • Compute Engine	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Defining the period of inactivity after which the system initiates a session lock b. Preventing access to the system and viewing of data by initiating a session lock after the defined period of inactivity c. Concealing previously visible information via a pattern-hiding display after the defined period of inactivity									

While the primary mechanism (OS screen lock) is configured outside the Google Cloud Console, the following feature(s) in Google Workspace and Google Cloud Console are relevant or complementary:

- Google Workspace Session Length Controls
- Compute Engine (for OS-level configuration within VMs and SSH settings)

A description of relevant features and implementation guidance is included below.

Google Workspace Session Length Controls (Complementary Control): Control how long a user's Google Account session remains valid before requiring re-authentication, regardless of activity. It helps manage overall session duration but does not replace the inactivity lock required by this control.

- **Configure Google Session Length**
 1. Navigate to the Google Admin console (admin.google.com).
 2. Go to **Security > Access and data control > Google session control**.
 3. Set the **Web session duration** for Google services (including Cloud Console) to align with organizational policy (e.g., 8 hours, 1 day). Shorter durations increase security but may impact user experience.
 4. Click **Save**.

Compute Engine (VM Access Configuration): Configure screen lock out time at the VM level.

- **OS-Level Screen Lock (VM GUI Access):** If users access graphical desktops on Compute Engine VMs (e.g., via RDP for Windows, VNC for Linux), configure the screen lock timeout within the VM's operating system using standard OS tools (e.g., Group Policy for Windows, dconf/gsettings for Linux desktops). This mirrors the requirement for physical endpoints. Implementation details depend on the specific OS and are managed within the VM, potentially automated via startup scripts or configuration management.
- **SSH Session Timeout (VM CLI Access):** To terminate inactive *SSH connections* (not GUI locks)
 1. Connect to the Linux VM instance using SSH.
 2. Edit the SSH daemon configuration file (`/etc/ssh/sshd_config`) with root privileges.
 3. Add or modify the following lines, adjusting the values based on your policy (e.g., 15 minutes = 900 seconds)
 - `ClientAliveInterval 300 # Send keepalive every 5 minutes`

- ClientAliveCountMax 3# Disconnect after 3 missed keepalives (15 minutes total)
 - Alternatively, configure LoginGraceTime for initial login timeout or TCP KeepAlives at the OS networking level.
4. Save the file and restart the SSH service (sudo systemctl restart sshd).

Additional Considerations

- Use standard operating system settings or endpoint management tools (like Group Policy Objects, MDM profiles) to enforce screen lock timeouts (e.g., 15 minutes) and require password re-authentication to unlock. All modern OS screen locks provide pattern-hiding displays. *Configuration of these endpoint settings occurs outside the Google Cloud Console.*
- This control is distinct from session termination (logging out completely), which is covered by 3.1.11. Session lock is a temporary measure for short absences.
- Ensure the inactivity timeout value defined in policy is consistently applied to endpoint OS configurations and, where applicable, VM OS configurations.
- User training on manually locking screens (Windows Key + L, Ctrl+Cmd+Q on Mac) is a vital supplement to automatic locking.

Supplemental Guidance

- [Set session length for Google services - Google Workspace Admin Help](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.11		
Control Description	Terminate (automatically) a user session after a defined condition		
Key Services	• Google Workspace Session Length Controls • Application/Web Server Configuration	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Defining conditions requiring a user session to terminate b. Terminating a user session automatically after any of the defined conditions occur			

The following feature(s) in Google Cloud Console and related services may be used to support this control:

- Google Workspace Session Length Controls
- Application/Web Server Configuration (within customer deployments)

A description of relevant features and implementation guidance is included below.

Google Workspace Session Length Controls: Automatically terminates the user's Google Account session after a fixed duration, requiring re-authentication. This is the primary control for the overall Google Cloud Console session duration.

- **Configure Google Session Length**

1. Navigate to the Google Admin console (admin.google.com).
2. Go to **Security > Access and data control > Google session control**.
3. Set the **Web session duration** for Google services. Choose a duration consistent with your organization's policy (e.g., 8 hours, 1 day, or a custom time). Setting this to a required duration fulfills the "terminate after a defined condition" (the condition being the elapsed time).
4. Click **Save**.

Compute Engine (SSH Session Timeout): Terminate inactive SSH connections to Linux VMs.

- **Configure SSH Inactivity Timeout**

1. Connect to the Linux VM instance using SSH.
2. Edit the SSH daemon configuration file (`/etc/ssh/sshd_config`) with root privileges (`sudo`).
3. Add or modify the following lines, setting the interval and count to match your inactivity policy (e.g., `ClientAliveInterval 300` and `ClientAliveCountMax 3` for a 15-minute inactivity timeout)
 - `ClientAliveInterval 300` # Seconds between keepalive messages
 - `ClientAliveCountMax 3` # Number of missed keepalives before disconnect
4. Save the configuration file.
5. Restart the SSH service.
6. This configuration can be automated using startup scripts or configuration management tools.

Application/Web Server Configuration (Customer Responsibility): For web applications or other services hosted on Compute Engine, GKE, App Engine, etc., inactivity session

timeouts must be configured within the application itself or the hosting web/application server.

- **Implementation**

- Configuration methods vary greatly depending on the technology stack (e.g., session.gc_maxlifetime in PHP, session-timeout in web.xml for Java applications, ProxyTimeout or application-specific settings in Apache/Nginx, settings within frameworks like Django or Ruby on Rails). Customers must configure these timeouts according to their application documentation and organizational policy. *This configuration occurs within the customer's deployment, not directly through general Cloud Console infrastructure settings.*

Additional Considerations

- **Distinguish from Session Lock (AC.L2-3.1.10):** Session termination logs the user out or disconnects the session entirely, requiring re-authentication or reconnection. Session lock ([AC.L2-3.1.10](#)) only locks the screen, preserving the session state.
- **Defining the "Condition":** Your organization must clearly define the conditions for termination (e.g., maximum duration like 12 hours for the Google session, inactivity period like 15 minutes for SSH).
- **Scope:** Recognize that different mechanisms control different session types (overall Google auth session, SSH connection session, web application session). Apply appropriate controls to each relevant session type.
- **Cloud Console Inactivity/:** There is not a direct setting in Google Cloud or Workspace to terminate the main Cloud Console web session based purely on *inactivity* within the browser tab. Termination is primarily governed by the overall session duration set in Workspace session controls. Application-specific inactivity timeouts must be handled by the applications themselves.

Supplemental Guidance

- [Set session length for Google services - Google Workspace Admin Help](#)

Control Domain	Access Control
Control #	AC.L2-3.1.12
Control Description	Monitor and control remote access sessions

Key Services	• Cloud Audit Logs • VPC Flow Logs • Firewall Rules Logging • SCC • IAM • IAP • OS Login • VPC Firewall Rules • Context-Aware Access • VPC Service Controls	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Permitting remote access sessions permitted - Identifying the types of permitted remote access - Controlling remote access sessions - Monitoring remote access sessions When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Cloud Audit Logs (Admin Activity, System Event, Data Access, Login) • VPC Flow Logs • Firewall Rules Logging • Security Command Center (SCC) • Identity and Access Management (IAM) with Conditions • Identity-Aware Proxy (IAP) • OS Login • VPC Firewall Rules • Context-Aware Access (part of Chrome Enterprise Premium) • VPC Service Controls A description of relevant features and implementation guidance is included below. Cloud Audit Logs: Provide visibility into actions performed within your Google Cloud environment, including remote access events. • Enable & Review Logs - Admin Activity logs are enabled by default. Ensure Data Access logs are enabled for relevant services like IAP (specifically iap.googleapis.com/AccessDenied and			

iap.googleapis.com/ReceivedThroughIAP) and OS Login (compute.googleapis.com/SshLogin). Enable via **IAM & Admin > Audit Logs**.

2. Navigate to **Logging > Logs Explorer**.
3. Query for relevant remote access events
 - Console/API Logins Search for google.login.LoginService events or Admin Activity logs showing actions by users. Filter by IP address if needed.
 - IAP Access Filter for logName "organizations/[ORG_ID]/logs/iap.googleapis.com%2Frequests" or similar project/folder paths. Look at jsonPayload.access_granted status and principal email.
 - OS Login SSH Filter for logName "projects/[PROJECT_ID]/logs/compute.googleapis.com%2Foslogin".
 - SSH Key-based Logins Check standard Linux audit logs (/var/log/auth.log or similar) within the VM, potentially shipping these logs to Cloud Logging using the Ops Agent.
4. Configure sinks (**Logging > Log Router**) to send logs to BigQuery for analysis or SCC for threat detection.

VPC Flow Logs & Firewall Rules Logging: Provide network-level visibility into remote connections.

- **Enable Logging**
 1. For VPC Flow Logs Navigate to **VPC network > VPC networks**, select a network, select a subnet, click **Edit**, turn **Flow logs** On, and configure sampling/aggregation.
 2. For Firewall Rules Logging When creating or editing a firewall rule (**VPC network > Firewall**), set **Firewall logs** to **On**. Log allows and denies relevant to remote access ports (e.g., 22, 3389, 443).
- **Monitor Logs** Analyze these logs in Logs Explorer or BigQuery to identify connection attempts, sources, destinations, and allowed/denied traffic related to remote access.

Security Command Center (SCC): Provide centralized visibility, threat detection, and posture management.

- **Utilize SCC**
 1. Activate SCC Standard or Premium tier (**Security > Security Command Center**).

2. Ensure log sources (Audit Logs, Firewall Logs, VPC Flow Logs) are configured to flow into SCC.
3. Monitor SCC dashboards for findings related to remote access (e.g., Firewall Insights, Event Threat Detection findings like "SSH Brute Force," "Anomalous IAM Grant").
4. Review assets and compliance status related to remote access controls.

Identity-Aware Proxy (IAP): Act as a secure, managed access point for specific resources, enforcing identity and context.

- **Configure IAP**

1. Navigate to **Security > Identity-Aware Proxy**.
2. Enable IAP for desired resources (App Engine apps, backend services for Load Balancers, Compute Engine instances). This requires configuring an OAuth consent screen first.
3. In the IAP settings for the resource, click **Add Principal**.
4. Grant the IAP-secured Web App User role (for web apps) or IAP-secured Tunnel User role (for SSH/RDP) to authorized users or groups.
5. (Optional) Apply Access Levels (Context-Aware Access) to enforce conditions like source IP or device compliance.

OS Login: Manage SSH access to VMs using IAM identities and policies.

- **Enable & Configure OS Login**

1. Enable OS Login for a project or individual VMs via metadata **Compute Engine > Metadata**, add enable-oslogin with value TRUE.
2. Grant necessary OS Login IAM roles to users/groups
 - roles/compute.osLogin Allows login without administrator privileges.
 - roles/compute.osAdminLogin Allows login with administrator privileges.
3. (Optional) Enable OS Login with 2-Step Verification (MFA) via organization policy constraints/compute.requireOsLoginMfa.
4. Remove project-wide or instance metadata SSH keys to ensure access is solely controlled via OS Login and IAM.

VPC Firewall Rules: Control network traffic flow to restrict remote access.

- **Configure Firewall Rules**

1. Navigate to **VPC network > Firewall**.

2. Create ingress rules for necessary remote access protocols (e.g., TCP 22 for SSH, TCP 3389 for RDP, TCP 443 for HTTPS) with the highest priority (lowest number).
3. Crucially, set the **Source filter** to be as restrictive as possible
 - Use Source IP ranges only allowing specific corporate network IPs or bastion host IPs.
 - If using IAP for SSH/RDP, allow only IAP's TCP forwarding IP range (35.235.240.0/20).
4. Apply rules to specific **Targets** (using network tags or service accounts) rather than all instances.
5. Ensure a lower-priority Deny rule exists or rely on the implicit deny-all ingress rule.

Context-Aware Access (Chrome Enterprise Premium): Provide fine-grained, attribute-based access control.

- **Implement Access Levels**
 1. Navigate to **Security > Access Context Manager**.
 2. Create **Access Levels** based on attributes like IP address, device policy compliance (requires endpoint verification), user identity, etc.
 3. Apply these Access Levels as conditions in IAM policies (**IAM & Admin > IAM > Edit Principal > Add Condition**) or within IAP resource configurations.

Additional Considerations

- Define and document your remote access policy clearly.
- Enforce Multi-Factor Authentication (MFA / 2-Step Verification) for all remote access, especially privileged access (see control [IA.L2-3.5.3](#)). Google Workspace 2SV enforcement and OS Login MFA policy help achieve this.
- Regularly review remote access logs and firewall rules for appropriateness and potential anomalies.
- Use bastion hosts or IAP instead of allowing direct external SSH/RDP access to VMs whenever possible.

Supplemental Guidance

- [Cloud Audit Logs overview](#)
- [Configure VPC Flow Logs | Google Cloud](#)
- [Firewall Rules Logging | Cloud NGFW](#)
- [Security Command Center overview | Google Cloud](#)

- [Identity-Aware Proxy overview | Google Cloud](#)
- [Enable IAP for Compute Engine | Identity-Aware Proxy | Google Cloud](#)
- [About OS Login | Compute Engine Documentation | Google Cloud](#)
- [VPC firewall rules | Cloud NGFW](#)
- [Chrome Enterprise Premium overview | BeyondCorp Enterprise | Google Cloud](#)
- [Overview of VPC Service Controls | Google Cloud](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.13		
Control Description	Employ cryptographic mechanisms to protect the confidentiality of remote access sessions		
Key Services	• GFE / Standard TLS Encryption • IAP • Cloud VPN • Cloud Load Balancing • Compute Engine	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Identifying cryptographic mechanisms to protect the confidentiality of remote access sessions b. Implementing cryptographic mechanisms to protect the confidentiality of remote access sessions When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Google Front End (GFE) / Standard TLS Encryption (Inherent) • Identity-Aware Proxy (IAP) (Inherent TLS Encryption) • Cloud VPN (IPsec Configuration) • Cloud Load Balancing (SSL Certificate Management) • Compute Engine (SSH Configuration within VMs) A description of relevant features and implementation guidance is included below. GFE / Standard TLS Encryption (Inherent): Connections to most Google Cloud services, including the Google Cloud Console (console.cloud.google.com) and APIs			

(*googleapis.com), are automatically protected using HTTPS with strong TLS encryption managed by Google.

- **Implementation**

- No specific customer configuration is required to enable this transport encryption; it is enforced by Google. Ensure clients accessing these services support modern TLS versions (1.2+).

Identity-Aware Proxy (IAP) (Inherent TLS Encryption): Secure access using HTTPS for web applications and TLS-encrypted tunnels for SSH/RDP connections initiated via gcloud or the console.

- **Implementation**

- The transport encryption for IAP connections is handled automatically by the service. Customers configure *access policies* via IAM, but the TLS encryption itself is inherent to how IAP operates. Ensure backend web applications behind IAP also support HTTPS if end-to-end encryption is required.

Cloud VPN: Provide encrypted IPsec tunnels between your remote network and your Google Cloud VPC network.

- **Configure IPsec Encryption**

1. When creating a Classic VPN or HA VPN tunnel (**Network connectivity > VPN**), you will configure IKE (Internet Key Exchange) versions and encryption protocols.
2. Choose IKEv2 (preferred over IKEv1).
3. Select strong Phase 1 (IKE) and Phase 2 (IPsec) encryption algorithms, integrity algorithms, and Diffie-Hellman (DH) groups that meet your security requirements (e.g., AES-256-GCM for encryption and integrity). Refer to the supported IKE ciphers documentation.
4. Configure pre-shared keys or certificates for authentication.
5. Ensure the peer VPN gateway on the remote network is configured with matching strong cryptographic parameters.

Cloud Load Balancing: Encrypt client connections to your applications using TLS/SSL certificates.

- **Configure SSL Certificates**

1. Navigate to **Network Security > Load balancing**, select your external HTTPS Load Balancer, and click **Edit**.

2. Go to **Frontend configuration**.
3. Under **Certificates**, click **Add Certificate**.
4. Choose **Google-managed certificate** (recommended, handles automatic provisioning and renewal) or upload a **Self-managed certificate**.
5. If using Google-managed, provide the domain(s).
6. Configure an **SSL Policy (Network Security > SSL policies)** to enforce minimum TLS versions (e.g., TLS 1.2) and specific cipher suites, disabling weak ones. Attach this policy to the load balancer's target HTTPS proxy.
7. Ensure **Backend configuration** uses HTTPS protocol if encryption between the load balancer and backends is required.

Compute Engine (SSH Configuration): The SSH protocol encrypts session data. Ensure strong algorithms are used within the VM's SSH server configuration.

- **Configure SSH Server (within VM)**
 1. Connect to the Linux VM instance using SSH.
 2. Edit the SSH daemon configuration file (/etc/ssh/sshd_config) with root privileges.
 3. Ensure Protocol 2 is specified (disallowing SSHv1).
 4. Review and specify strong KexAlgorithms, Ciphers, and MACs. Remove known weak options. Consult current security best practices for recommended algorithms (e.g., use AES-GCM modes, SHA2-based MACs, secure key exchange methods).
 5. Save the file and restart the SSH service (sudo systemctl restart sshd).

This configuration can be automated using startup scripts or configuration management tools.

Additional Considerations

- **Cryptographic Standard:** Use current industry best practices for cryptographic protocols (TLS 1.2+, IKEv2, SSH Protocol 2) and algorithms (AES-GCM, SHA-2, strong DH groups/elliptic curves). Avoid deprecated or weak options (SSLv3, TLS 1.0/1.1, MD5, SHA-1, DES, 3DES).
- **FIPS 140:** If required, consult Google Cloud's compliance documentation to understand which services and configurations meet FIPS 140 validation requirements. Some services may operate in a FIPS-compliant mode.

- **End-to-End Encryption:** Consider the entire path. While TLS may protect the connection to a load balancer, ensure traffic remains encrypted between the load balancer and backend instances if required by policy.

Supplemental Guidance

- [Encryption in transit for Google Cloud | Security](#)
- [Identity-Aware Proxy documentation | Google Cloud](#)
- [Supported IKE ciphers | Cloud VPN](#)
- [Create an HA VPN gateway to a peer VPN gateway | Google Cloud](#)
- [SSL certificates overview | Load Balancing | Google Cloud](#)
- [Use Google-managed SSL certificates | Load Balancing](#)
- [SSL policies for SSL and TLS protocols | Load Balancing | Google Cloud](#)
- [About FIPS-validated encryption in GKE | Google Kubernetes Engine \(GKE\)](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.14		
Control Description	Route remote access via managed access control points		
Key Services	• IAP • Cloud VPN / Cloud Interconnect • VPC Firewall Rules • External Cloud Load Balancers • Compute Engine	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Identifying and implementing managed access control points - Routing remote access through managed network access control points When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Identity-Aware Proxy (IAP) • Cloud VPN / Cloud Interconnect • VPC Firewall Rules • External Cloud Load Balancers			

- Compute Engine (for Bastion Host pattern)

A description of relevant features and implementation guidance is included below.

Identity-Aware Proxy (IAP): Provide a secure access control point for web applications and SSH/RDP access without requiring VPNs or exposing resources directly.

- **Configure IAP as Access Point:**

1. Enable IAP (**Security > Identity-Aware Proxy**) for target resources (App Engine, GKE/Compute Engine via Load Balancer backend services, Compute Engine instances for TCP forwarding).
2. Ensure resources do NOT have direct external IP access or overly permissive firewall rules bypassing IAP.
3. Configure IAM policies on the IAP-secured resource, granting IAP-secured Web App User or IAP-secured Tunnel User roles only to authorized remote users/groups. Access is granted only after authentication and authorization through IAP.
4. Configure VPC Firewall rules (see below) to allow traffic *only* from IAP's IP range (35.235.240.0/20 for TCP forwarding) to the backend resources on the required ports (e.g., 22, 3389, 80, 443).

Cloud VPN / Cloud Interconnect: Provide a dedicated network gateways connecting remote networks (on-premises data centers, other clouds) to your VPC.

- **Configure Gateways as Access Points:**

1. Set up HA VPN tunnels or Interconnect attachments (**Network connectivity > VPN or Interconnect**). These gateways become the managed entry points for traffic from those specific remote networks.
2. Configure VPC Firewall rules (see below) to allow traffic from the designated remote network IP ranges (defined in VPN/Interconnect configuration) to the necessary internal Google Cloud resources.

VPC Firewall Rules: Enforce the routing of traffic through the chosen managed access control points by blocking direct connections.

- **Enforce Routing via Firewalls:**

1. Navigate to **VPC network > Firewall**.
2. **Deny Direct Access** Create high-priority (low number) Deny rules or ensure no Allow rules permit direct ingress from the internet (0.0.0.0/0) to internal application VMs or sensitive ports (like 22, 3389, database ports).

3. **Allow Access from Managed Points** Create lower-priority (higher number) Allow rules specifically permitting traffic *only* from the IP ranges of your managed access points to the necessary backend resources and ports.
Examples:
 - Allow TCP:22 from 35.235.240.0/20 (IAP) to VMs tagged iap-ssh-target.
 - Allow TCP:443 from 0.0.0.0/0 to External HTTPS Load Balancer frontend IP (if LB is the managed point).
 - Allow required ports from your specific on-premises IP range(s) (defined in VPN/BGP session) to internal resources, only if using Cloud VPN/Interconnect as the access point.
 - Allow TCP:22 from your designated bastion host IP address(es) to internal VMs tagged ssh-target.

External Cloud Load Balancers: Serve as managed ingress points for distributing traffic to backend services and is often used with IAP.

- **Configure Load Balancer as Access Point:**
 1. Create an External HTTP(S) Load Balancer, TCP Proxy Load Balancer, or SSL Proxy Load Balancer (**Network Services > Load balancing**).
 2. Configure backend services pointing to your internal instance groups or network endpoint groups.
 3. Assign a static external IP address to the load balancer's frontend.
 4. Configure firewall rules to allow traffic to the load balancer, but deny direct traffic to the backend instances from external sources.
 5. Optionally, enable IAP on the backend service associated with the HTTPS load balancer.

Compute Engine (Bastion Host Pattern): Create a secure jump server.

- **Implement Bastion Host:**
 1. Create a minimal Compute Engine VM instance to serve as the bastion. Harden the OS image.
 2. Assign a static external IP address *only* to the bastion host (or use IAP to access it without an external IP).
 3. Configure VPC Firewall rules to allow SSH access (TCP:22) *only* to the bastion host, and *only* from authorized source IPs (e.g., corporate network) or via IAP.
 4. Configure VPC Firewall rules to allow SSH access from the bastion host's internal IP address to other internal VMs that require administrative access.

5. Do not allow direct external SSH access to the internal VMs.

Additional Considerations

- **Minimize External IPs:** Avoid assigning external IP addresses directly to VMs that do not explicitly require them. Use internal IPs and route access via managed points like IAP or Load Balancers.
- **Defense in Depth:** Use multiple controls. For example, route traffic through a VPN gateway *and* then require IAP authentication for specific application access.
- **Monitoring:** Ensure your chosen managed access control points (IAP, VPN, Load Balancers) have logging enabled and are actively monitored (as covered in [AC.L2-3.1.12](#)).

Supplemental Guidance

- [Identity-Aware Proxy overview | Google Cloud](#)
- [Enable IAP for Compute Engine | Identity-Aware Proxy | Google Cloud](#)
- [Cloud VPN overview](#)
- [VPC firewall rules | Cloud NGFW](#)
- [Securely connecting to VM instances | Compute Engine Documentation | Google Cloud](#)
- [Using IAP for TCP forwarding | Identity-Aware Proxy | Google Cloud](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.15		
Control Description	Authorize remote execution of privileged commands and remote access to security-relevant information		
Key Services	• IAM • OS Login • Cloud Audit Logs • SCC	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Identifying privileged commands authorized for remote execution via Admin SDK b. Identifying security-relevant information authorized to be accessed remotely via Admin SDK c. Authorizing the execution of the identified privileged commands via remote access			

- d. Authorizing access to the identified security-relevant information via remote access

When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control:

- Identity and Access Management (IAM) (including Roles, Permissions, and Conditions)
- OS Login
- Cloud Audit Logs (for verification)
- Security Command Center (SCC) (access to findings requires authorization)

A description of relevant features and implementation guidance is included below.

Identity and Access Management (IAM): Authorize privileged commands (API calls) and access to security information remotely.

- **Assign Specific Privileged Roles**
 1. Navigate to **IAM & Admin > IAM**.
 2. Identify users or groups requiring remote privileged access.
 3. Grant predefined roles that contain the necessary permissions for their specific remote tasks (e.g., roles/compute.admin for full Compute Engine control, roles/logging.privateLogViewer to view specific logs, roles/securitycenter.findingsEditor to manage SCC findings). Grant these roles *only* to explicitly authorized personnel.
 4. Avoid overly broad roles like roles/owner or roles/editor for remote access.
- **Create Custom Roles for Least Privilege**
 1. If predefined roles are too broad, navigate to **IAM & Admin > Roles > + Create Role**.
 2. Create roles containing *only* the specific permissions required for the remote privileged tasks (e.g., a role with only compute.instances.start, compute.instances.stop permissions for a remote operator).
 3. Assign these granular custom roles instead of broader predefined ones.
- **Use IAM Conditions (Optional Granularity)**
 1. When granting a privileged role (**IAM & Admin > IAM > Edit Principal**), click **Add condition**.
 2. Create conditions based on attributes like:
 - request.time Authorize access only during specific hours.
 - request.origin.ip Authorize access only from specific IP addresses (e.g., corporate network).

- Access Levels (via Access Context Manager) Authorize only if the user meets predefined context requirements (device compliance, geo-location).
- 3. This allows time-bound or context-aware authorization for remote privileged access.

OS Login: Authorize privileged command execution (sudo) within Linux VMs during remote SSH sessions, based on IAM roles.

- **Configure OS Login for Sudo Authorization**

1. Enable OS Login on the project or specific VMs (**Compute Engine > Metadata > enable-oslogin = TRUE**).
2. Grant the IAM role roles/compute.osAdminLogin to users/groups authorized to execute privileged commands (sudo) within the VM during remote SSH sessions.
3. Grant the roles/compute.osLogin role for users needing remote SSH access *without* sudo privileges. OS Login maps these roles to appropriate POSIX permissions, explicitly authorizing the level of access within the remote session.
4. Ensure standard SSH key-based access (project/instance metadata keys) is disabled so OS Login/IAM is the sole authorization path for SSH.

Cloud Audit Logs & Security Command Center (Verification): Access security-relevant information like audit logs or SCC findings requires specific authorization.

- **Authorize Access to Security Information**

1. Grant roles like roles/logging.viewer, roles/logging.privateLogViewer, or specific Logs View Accessor roles (**Logging > Log Router > select sink > Edit Sink** or **Logging > Logs Storage > Log Bucket**) only to users authorized to view audit logs remotely.
2. Grant roles like roles/securitycenter.findingsViewer or roles/securitycenter.securityCenterAdminViewer only to users authorized to view SCC information remotely.

- **Audit Privileged Actions:** Use Cloud Audit Logs (Admin Activity, OS Login logs) to verify that privileged commands executed remotely align with the authorizations granted via IAM and OS Login.

Additional Considerations

- **Remote Access Method:** Ensure the remote connection itself is established via authorized methods (e.g., IAP, authorized VPN) as covered in controls like [AC.L2-3.1.14](#).
- **Least Privilege:** Regularly review IAM roles assigned for remote access to ensure they remain necessary and grant minimal required privileges.
- **Separation of Duties:** Assign distinct roles for different types of remote privileged functions where appropriate.
- **MFA:** Enforce Multi-Factor Authentication (via Google Workspace 2-step verification or OS Login MFA policy) for accounts granted remote privileged authorizations.

Supplemental Guidance

- [IAM overview | IAM Documentation | Google Cloud](#)
- [IAM roles and permissions index - Documentation](#)
- [Create and manage custom roles | IAM Documentation | Google Cloud](#)
- [Overview of IAM Conditions | IAM Documentation | Google Cloud](#)
- [About OS Login | Compute Engine Documentation | Google Cloud](#)
- [Set up OS Login | Compute Engine Documentation | Google Cloud](#)
- [Cloud Audit Logs overview](#)
- [Access control with IAM | Security Command Center | Google Cloud](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.16		
Control Description	Authorize wireless access prior to allowing such connections		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: - Identifying wireless access points - Authorizing wireless access prior to allowing such connections <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems,</i>			

applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.

Supplemental Guidance

- N/A

Control Domain	Access Control		
Control #	AC.L2-3.1.17		
Control Description	Protect wireless access using authentication and encryption		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: - authenticating wireless access to the system - protecting wireless access to the system using encryption <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Access Control		
Control #	AC.L2-3.1.18		
Control Description	Control connection of mobile devices		
Key Services	- Google Workspace MDM - CAA - IAP	Control Responsibility	☐ Google ☐ Shared ☒ Customer

	Cloud Identity / Google Workspace Device Inventory		
Customer Implementation Description			
Google Cloud Customers are responsible for: - Identifying mobile devices that process, store, or transmit CUI - Authorizing mobile device connections - Logging and monitoring mobile device connections Controlling the connection of mobile devices (smartphones, tablets) involves managing which devices can access organizational data or resources and enforcing security configurations on those devices. This is primarily achieved through Mobile Device Management (MDM) policies and context-aware access rules. When configured correctly, the following feature(s) in Google Workspace and Google Cloud Console may be used to support this control: - Google Workspace Mobile Device Management (MDM) (via Admin Console) - Context-Aware Access (CAA) (via Cloud Console - Access Context Manager) - Identity-Aware Proxy (IAP) (via Cloud Console - for enforcing CAA) - Cloud Identity / Google Workspace Device Inventory (via Admin Console) A description of relevant features and implementation guidance is included below. Google Workspace Mobile Device Management (MDM) (Configured in Admin Console): Allow administrators to enforce security policies on mobile devices accessing Google Workspace data. Set Up Mobile Management - Navigate to the Google Admin console (admin.google.com). - Go to Devices > Mobile & endpoints > Settings > Universal settings. - Configure basic or advanced mobile management according to your organization's needs and license level. Enforce Security Policies - Under Universal settings > Security, enable policies like: - Require users to set a passcode. Configure complexity requirements. - Require device encryption (addresses 3.1.19). - Block compromised devices. - Configure platform-specific settings (Android, iOS) under Devices > Mobile & endpoints > Settings for additional controls like minimum OS versions, app management, etc.			

- **Control Device Connections (Approval)**
 1. Go to **Universal settings > General > Account setup**.
 2. Consider enabling **Device approvals**, which requires an administrator to approve each new device before it can sync Google Workspace data, thereby controlling initial connection for Workspace access.
- **View and Manage Devices**
 1. Go to **Devices > Mobile & endpoints > Devices** to view a list of managed devices.
 2. From here, administrators can approve, block, or wipe specific devices.

Context-Aware Access (CAA) (Configured in Cloud Console): Create granular access control policies for Google Cloud resources based on user identity and context, including device compliance managed by Google Workspace MDM (or other integrated partners).

- **Create Device-Based Access Levels**
 1. Ensure devices are enrolled in Google Workspace MDM and reporting compliance status (may require specific Workspace license tiers).
 2. In the Google Cloud Console, navigate to **Security > Access Context Manager**.
 3. Click **New Access Level**.
 4. Give the Access Level a name (e.g., `compliant_mobile_devices`).
 5. Choose **Attribute based** condition mode.
 6. Add conditions based on device attributes reported by MDM:
 - `device.is_managed_device == true`
 - `device.is_compliant_device == true` (Checks overall compliance reported by MDM)
 - `device.os_type == "android" or "ios"`
 - `device.os_version >= "13.0"` (Example OS version check)
 - `device.screen_lock_secured == true`
 - `device.encryption_status == "ENCRYPTED"`
 7. Combine conditions using AND/OR logic as needed.
 8. Click **Create**.

Identity-Aware Proxy (IAP) (Configured in Cloud Console): Apply the CAA Access Levels to control connections to *specific resources* protected by IAP.

- **Apply Access Levels to IAP Resources**
 1. Navigate to **Security > Identity-Aware Proxy**.
 2. Select the IAP-secured resource (e.g., a backend service, App Engine app).

3. Edit the access policy for the resource.
4. In the role bindings (e.g., for IAP-secured Web App User), add a Condition.
5. Set the Condition Type to **Access Level**.
6. Select the device-based Access Level you created (e.g., `compliant_mobile_devices`).
7. Save the changes. Now, only connections from mobile devices meeting the criteria defined in the Access Level will be allowed to access this specific resource via IAP.

Additional Considerations

- **MDM Enrollment:** For these controls to be effective, mobile devices accessing sensitive data or resources generally need to be enrolled in your MDM solution (e.g., Google Workspace MDM).
- **Scope of Control:** MDM controls the *device* configuration and access to Workspace data. CAA + IAP control access to specific Google Cloud/IAP resources based on device state. Define which resources require mobile device connection controls.
- **User Experience:** Balance security requirements with user productivity. Ensure clear communication about MDM enrollment and policies.

Supplemental Guidance

- [Chrome Enterprise Premium overview | BeyondCorp Enterprise | Google Cloud](#)
- [Create an access level for Access Context Manager | Google Cloud](#)
- [Access level attributes | Access Context Manager | Google Cloud](#)
- [Setting up context-aware access with Identity-Aware Proxy | Google Cloud](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.19		
Control Description	Encrypt CUI on mobile devices and mobile computing platforms		
Key Services	• Google Workspace MDM CAA • Endpoint Verification • IAP / VPC Service Controls / IAM Conditions	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			

Google Cloud Customers are responsible for:

- a. Identifying mobile devices and mobile computing platforms that process, store, or transmit CUI; and
- b. Employing encryption to protect CUI on identified mobile devices and mobile computing platforms.

This control requires that CUI stored locally on mobile devices (smartphones, tablets) and mobile computing platforms (typically interpreted as laptops) is protected using encryption. When configured correctly, the following feature(s) in Google Workspace and Google Cloud Console may be used to support this control:

- Google Workspace Mobile Device Management (MDM) (via Admin Console)
- Context-Aware Access (CAA) (via Cloud Console - Access Context Manager)
- Endpoint Verification (helper application/extension, deployment managed by customer)
- Identity-Aware Proxy (IAP) / VPC Service Controls / IAM Conditions (for applying CAA rules)

A description of relevant features and implementation guidance is included below.

Google Workspace Mobile Device Management (MDM) (Configured in Admin Console): Use MDM to enforce encryption on managed Android and iOS devices.

- **Enforce Mobile Device Encryption:**
 1. Navigate to the Google Admin console (admin.google.com).
 2. Go to **Devices > Mobile & endpoints > Settings > Universal settings**.
 3. Ensure appropriate management level (e.g., Advanced) is selected.
 4. Under **Universal settings > Security**, review or enable the policy **Require device encryption**. (Specific wording might vary slightly or be under platform-specific settings like Android / iOS).
 5. Enabling this policy typically prevents unencrypted managed devices from syncing Google Workspace data (like Gmail or Drive files containing CUI).

Context-Aware Access (CAA) & Endpoint Verification (Configured in Cloud Console): Use CAA to verify encryption status on laptops accessing Google Cloud resources. Endpoint Verification must be deployed on the laptops to report status.

- **Deploy Endpoint Verification:** Follow the Google Workspace documentation to deploy the Endpoint Verification Chrome extension and native helper app to managed laptops.

- **Create Encryption Check Access Level:**
 1. In the Google Cloud Console, navigate to **Security > Access Context Manager**.
 2. Click **New Access Level**.
 3. Give it a name (e.g., encrypted_laptops).
 4. Choose **Attribute based** condition mode.
 5. Add the condition `device.encryption_status == "ENCRYPTED"`. This requires the device to report that its disk is encrypted via Endpoint Verification.
 6. Optionally add other conditions (e.g., minimum OS version, screen lock).
 7. Click **Create**.
- **Apply the Access Level:** Use this Access Level (encrypted_laptops) to control access to sensitive Google Cloud resources from laptops:
 - **IAP:** Apply the Access Level to IAP-secured resources (**Security > Identity-Aware Proxy > Select Resource > Edit > Apply Access Levels**).
 - **VPC Service Controls:** Use the Access Level in Ingress Rules for service perimeters to control access to protected APIs.
 - **Other Google Cloud Services (Limited):** Some Google Cloud services might allow applying Access Levels directly or via IAM Conditions.

OS-Level Encryption Configuration (External): The actual act of encrypting the laptop's disk (BitLocker, FileVault, LUKS) is performed using operating system tools or policies managed outside the Google Cloud/Workspace consoles.

Additional Considerations

- **CUI Handling Policy:** Define clearly where CUI is permitted to be stored. If CUI storage on mobile devices/laptops is prohibited by policy and technical controls (e.g., VDI-only access), this control's applicability might focus solely on verifying that prohibition. However, encryption is still a best practice.
- **Enrollment/Deployment:** MDM policies require device enrollment. CAA checks require Endpoint Verification deployment and successful reporting.
- **FIPS 140:** If FIPS 140-validated encryption is required for CUI, ensure the OS-level encryption (BitLocker, FileVault) is configured in its FIPS-compliant mode, where applicable. Google Workspace MDM and CAA primarily enforce/check the *presence* of encryption, not necessarily its FIPS mode configuration, although certain compliance signals might be available depending on MDM integration.

Supplemental Guidance

- [Apply settings for mobile devices \(Google Workspace\)](#)

- [Chrome Enterprise Premium overview | BeyondCorp Enterprise | Google Cloud](#)
- [Create an access level for Access Context Manager | Google Cloud](#)
- [Access level attributes | Access Context Manager | Google Cloud](#)
- [Endpoint Verification overview](#)
- [Turn endpoint verification on or off - Google Workspace Admin Help](#)

Control Domain	Access Control								
Control #	AC.L1-3.1.20								
Control Description	Verify and control/limit connections to and use of external systems								
Key Services	• VPC Firewall Rules • Cloud NAT • Private Google Access • VPC Service Controls • IAP • Organization Policies • Cloud DNS • Cloud Armor	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Identifying connections permitted to external systems from your Google Workspace account; b. Identifying the use of external systems; c. Verifying connections to external systems from your Google Workspace account; d. Verifying the use of external systems; e. Controlling and limiting connections to external systems; and f. Limiting and controlling the use of external systems. When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • VPC Firewall Rules • Cloud NAT • Private Google Access • VPC Service Controls • Identity-Aware Proxy (IAP) • Organization Policies • Cloud DNS (Response Policy Zones) • Cloud Armor									

A description of relevant features and implementation guidance is included below.

VPC Firewall Rules: Control network connections at the IP/port level.

- **Control Egress Connections (To External Systems)**
 1. Navigate to **VPC network > Firewall**.
 2. Create a low-priority (high number, e.g., 65534) egress rule that denies all outbound traffic (Action Deny, Direction Egress, Destinations 0.0.0.0/0, Protocols and ports Deny all). Apply it broadly or use target tags/service accounts.
 3. Create higher-priority (lower number) egress Allow rules only for necessary outbound connections. Be specific:
 - Destinations Specify known external IP ranges required (e.g., partner IPs, specific SaaS IPs, OS update server IPs). Avoid 0.0.0.0/0 in Allow rules unless broadly necessary and other controls are in place.
 - Protocols and ports Specify only required ports (e.g., tcp:443 for HTTPS).
 - Targets Apply rules only to VMs/services that need that specific egress path using network tags or service accounts.
- **Control Ingress Connections (From External Systems)**
 1. Rely on the default-deny ingress rule or create an explicit one.
 2. Create specific ingress Allow rules only for necessary inbound traffic, preferably targeting managed access points like Load Balancers or IAP infrastructure, not directly to internal VMs.
 3. Restrict Source IP ranges to known external IPs whenever possible. Use IAP (see below) for user access instead of opening ports directly.

Cloud NAT: Provide managed outbound connectivity for private instances without external IPs.

- **Configure Controlled Egress**
 1. Navigate to **Network services > Cloud NAT**.
 2. Create a NAT gateway associated with your VPC network and region(s).
 3. Configure it to apply to specific subnets or all subnets in the network.
 4. Assign specific external IP addresses for egress traffic (Static or Automatic).
 5. VMs without external IPs in the configured subnets will use Cloud NAT for outbound internet connections, controlled by egress firewall rules.

Private Google Access: Allow private instances to reach Google APIs and services without internet egress.

- **Enable Private Access**

1. Navigate to **VPC network > VPC networks**, select a network, select a subnet, click **Edit**.
2. Turn **Private Google Access** On.
3. Instances in that subnet without external IPs can now reach most Google APIs (like Cloud Storage, BigQuery) via internal paths, limiting exposure. Requires appropriate DNS configuration (private.googleapis.com or restricted.googleapis.com).

VPC Service Controls: Create security perimeters around Google managed services to control data flow.

- **Limit External Access to Services**

1. Navigate to **Security > VPC Service Controls**.
2. Create a service perimeter, adding the projects and restricting the services (e.g., Cloud Storage, BigQuery) that handle sensitive data.
3. Configure Ingress and Egress rules for the perimeter:
 - Restrict which identities or networks outside the perimeter can access the protected services (Ingress).
 - Restrict or prevent protected services from sending data to destinations outside the perimeter (Egress), controlling connections to potentially external buckets, datasets, etc.

Identity-Aware Proxy (IAP): Control and verify inbound connections from external users.

- **Verify & Control Inbound User Connections**

1. Enable IAP for web applications or VM SSH/RDP access (**Security > Identity-Aware Proxy**).
2. Assign appropriate IAM roles (IAP-secured...) only to authorized external users or groups. IAP verifies the user's identity before allowing the connection to the Google Cloud resource.

Organization Policies: Enforce constraints to limit configurations enabling external connections.

- **Apply Constraints:**

1. Navigate to **IAM & Admin > Organization Policies**.
2. Find and configure policies like:
 - constraints/compute.vmExternalIpAccess Restrict or deny the assignment of external IPs to VMs.
 - constraints/sql.restrictPublicIp Prevent Cloud SQL instances from having public IPs.

Additional Considerations

- **Inventory:** Maintain an inventory of authorized external connections (both directions).
- **Default Deny:** Implement default-deny for both ingress and egress firewall rules. Explicitly allow only necessary traffic.
- **Monitoring:** Regularly review Firewall Rules Logs, VPC Flow Logs, NAT logs, and VPC SC audit logs to verify compliance with connection policies and detect anomalies.
- **Advanced Filtering:** Consider Cloud DNS Response Policy Zones (RPZ) to block name resolution for unauthorized external domains or Cloud Armor (for Load Balancers) / Secure Web Proxy (for egress) for more advanced traffic filtering.

Supplemental Guidance

- [VPC firewall rules | Cloud NGFW](#)
- [Using Egress Firewall Rules](#)
- [Cloud NAT overview](#)
- [Configure Private Google Access | VPC](#)
- [Overview of VPC Service Controls | Google Cloud](#)
- [Identity-Aware Proxy overview | Google Cloud](#)
- [Organization policy constraints | Resource Manager Documentation | Google Cloud](#)

Control Domain	Access Control		
Control #	AC.L2-3.1.21		
Control Description	Limit use of organizational portable storage devices on external systems		
Key Services	N/A	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Customers are responsible for: - Identifying and documenting the use of portable storage devices containing CUI on external systems - Defining the use of portable storage devices containing CUI on external systems - Limiting the use of portable storage devices containing CUI on external systems is limited as defined			

Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.

Supplemental Guidance

- N/A

Control Domain	Access Control		
Control #	AC.L1-3.1.22		
Control Description	Control information posted or processed on publicly accessible information systems		
Key Services	• Cloud Storage Public Access Prevention • IAM • IAP • VPC Service Controls • Cloud DLP (now part of Sensitive Data Protection) • Organization Policies	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Identifying individuals authorized to post or process information on publicly accessible systems - Identifying procedures to ensure CUI is not posted or processed on publicly accessible systems - Establishing a review process prior to posting of any content to publicly accessible systems - Ensuring content on publicly accessible systems is reviewed to ensure that it does not include CUI - Implementing mechanisms to remove and address improper posting of CUI When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Cloud Storage Public Access Prevention • Identity and Access Management (IAM)			

- Identity-Aware Proxy (IAP)
- VPC Service Controls
- Cloud DLP (Data Loss Prevention) (now part of Sensitive Data Protection)
- Organization Policies

A description of relevant features and implementation guidance is included below.

Cloud Storage Public Access Prevention: Provide a strong, centralized way to prevent public access to Cloud Storage buckets, overriding individual bucket/object permissions that might allow public access.

- **Enable Public Access Prevention:**
 1. Navigate to **IAM & Admin > Organization Policies**.
 2. Search for the policy `storage.publicAccessPrevention`.
 3. Click **Edit**.
 4. Select **Customize**.
 5. Choose **Replace** under **Applies to**.
 6. Under **Policy values**, select **Custom**.
 7. Select **Allow** or **Deny**. Choose **Deny** all to enforce public access prevention.
 8. Click **Save**. This can be applied at the Organization, Folder, or Project level.
 9. Alternatively, for per-bucket enforcement Navigate to **Cloud Storage > Buckets**, select a bucket, go to the **Permissions** tab, and under **Public access**, ensure it states "Not public" and that "Public access prevention" is set to "Enforced".

Identity and Access Management (IAM): Control who has permission to configure resources, including potentially making them public.

- **Restrict Public Access Permissions**
 1. Navigate to **IAM & Admin > IAM**.
 2. Review roles assigned to users and service accounts.
 3. Avoid granting roles with permissions that allow making resources public (e.g., `storage.buckets.setIamPolicy`, `storage.objects.setIamPolicy` with `allUsers` or `allAuthenticatedUsers`, permissions to deploy public App Engine/Cloud Run services) unless strictly necessary and only to authorized personnel.
 4. Use least privilege; create custom roles if needed.

Identity-Aware Proxy (IAP): Secure web applications and other resources, ensuring they are not publicly accessible without authentication.

- **Protect Web-Accessible Systems**

1. For any web application (on Compute Engine, GKE, App Engine) that processes CUI but needs to be reachable from the internet, enable IAP (**Security > Identity-Aware Proxy**).
2. Configure IAP access policies, granting the IAP-secured Web App User role only to authorized users/groups.
3. This places an authentication and authorization layer in front of the application, preventing unauthenticated public access even if the underlying compute resource has network paths from the internet.

VPC Service Controls: Create perimeters around sensitive data stores to prevent data exposure.

- **Create Perimeters for CUI Data Stores**

1. Navigate to **Security > VPC Service Controls**.
2. Create a service perimeter, adding projects containing CUI stored in supported services (like Cloud Storage, BigQuery).
3. Restrict the services within the perimeter.
4. Configure Ingress/Egress rules to control access:
 - Prevent access from public IP ranges (Ingress).
 - Prevent data from being copied or accessed from resources outside the perimeter, especially potentially public ones (Egress).

Cloud DLP (Data Loss Prevention): Scan storage and data streams to detect potential CUI exposure.

- **Scan for Exposed CUI**

1. Navigate to **Security > Data Loss Prevention**.
2. Create DLP inspection jobs to scan Cloud Storage buckets (even those intended to be private) for CUI patterns (using predefined or custom infoTypes).
3. Configure actions, such as logging findings to Security Command Center or Cloud Logging, sending Pub/Sub notifications for remediation workflows, or potentially de-identifying data (use with caution).
4. Regularly review findings to identify and remove any CUI found in improperly configured locations.

Organization Policies: Help prevent configurations that lead to public accessibility.

- **Apply Restrictive Policies**

1. Navigate to **IAM & Admin > Organization Policies**.
2. Enforce policies like constraints/compute.vmExternallpAccess (Deny) or constraints/sql.restrictPublicIp (Enforce) to limit direct public exposure of compute and database resources.

Additional Considerations

- **Policy and Procedures:** Technical controls should support administrative policies. Have clear policies designating who can authorize public content, review processes before publication, and procedures for removing CUI if found on public systems.
- **Training:** Ensure personnel involved in managing web content or publicly accessible systems are trained on CUI handling and identification.
- **Definition of Public:** Clearly define what "publicly accessible" means in your environment (e.g., requires no authentication, accessible from any IP address).

Supplemental Guidance

- [Public access prevention | Cloud Storage](#)
- [IAM permissions for Cloud Storage](#)
- [Identity-Aware Proxy overview | Google Cloud](#)
- [Overview of VPC Service Controls | Google Cloud](#)
- [Sensitive Data Protection Documentation | Google Cloud](#)
- [Creating and managing organization policies | Resource Manager Documentation | Google Cloud](#)

Control Domain	Awareness and Training		
Control #	AT.L2-3.2.1		
Control Description	Ensure that managers, systems administrators, and users of organizational systems are made aware of the security risks associated with their activities and of the applicable policies, standards, and procedures related to the security of those systems		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for:			

- a. Identifying security risks associated with organizational activities involving CUI
- b. Identifying policies, standards, and procedures related to the security of the system
- c. Informing managers, systems administrators, and users of the system of the security risks associated with their activities
- d. Informing managers, systems administrators, and users of the system of the applicable policies, standards, and procedures related to the security of the system

Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.

Supplemental Guidance

- N/A

Control Domain	Awareness and Training		
Control #	AT.L2-3.2.2		
Control Description	Ensure that organizational personnel are adequately trained to carry out their assigned information security-related duties and responsibilities		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Defining information security-related duties, roles, and responsibilities b. Assigning information security-related duties, roles, and responsibilities to designated personnel c. Adequately training personnel to carry out their assigned information security related duties, roles, and responsibilities <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			

- N/A

Control Domain	Awareness and Training		
Control #	AT.L2-3.2.3		
Control Description	Provide security awareness training on recognizing and reporting potential indicators of insider threat		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Identifying potential indicators associated with insider threats - Providing security awareness training on recognizing and reporting potential indicators of insider threat to managers and employees <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Audit and Accountability
Control #	AU.L2-3.3.1
Control Description	Create and retain system audit logs and records to the extent needed to enable the monitoring, analysis, investigation, and reporting of unlawful or unauthorized system activity

Key Services	• Cloud Logging • Logs Explorer • Cloud Audit Logs • Cloud Logging Storage • Cloud Storage • BigQuery • Pub/Sub	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for - Specifying audit logs needed (i.e., event types to be logged) to enable the monitoring, analysis, investigation, and reporting of unlawful or unauthorized system activity - Defining the content of audit records needed to support monitoring, analysis, investigation, and reporting of unlawful or unauthorized system activity - Defining retention requirements for audit records - Retaining audit records, as defined Google is responsible for - Generating audit records - Ensuring audit records, once created, contain the defined content When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control: • Cloud Logging • Logs Explorer • Cloud Audit Logs • Cloud Logging Storage (Log Buckets) • Cloud Storage • BigQuery • Pub/Sub A description of relevant features and implementation guidance is included below. Cloud Logging: Collect log data from Google Cloud resources. Your applications, on-premise resources, and resources from other cloud providers can send log data to Cloud Logging. You can also configure alerting policies so that Cloud Monitoring notifies you when certain kinds of events are reported in your log data. For regulatory or security reasons, you can determine where your log data is stored.			

In the Google Cloud console, you can view and analyze your log data either with **Logs Explorer** or the **Logs Analytics** pages.

Logs Explorer: Troubleshoot and analyze the performance of your services and applications. You can use Logs Explorer to retrieve, view, and analyze log entries that are stored in log buckets. Viewing and analyzing individual log entries and a sequence of log entries can help you troubleshoot problems.

- **To begin using the Logs Explorer:**
 1. In the Google Cloud console, go to the **Logs Explorer** page
 2. Select a Google Cloud project, folder, or organization
 3. The log entries displayed by the Logs Explorer page depend on the following
 - The resources searched for log entries.
 - The time-range setting.
 - Your Identity and Access Management (IAM) roles on the searched resources.
 - Your query filters the search results. For example, adding the query `severity>=ERROR` results in the display listing only those log entries with a severity level of at least ERROR.

By default, the Logs Explorer page searches the resources listed in the default log scope for log entries. When the default log scope isn't accessible, the page searches for the log entries that originate in your selected project, folder, or organization. For projects, the search results include the log entries that are routed to the project by a sink in another project, and then stored in a log bucket.

Cloud Audit Logs: Write audit logs that record administrative activities and accesses within your Google Cloud resources. Cloud Audit Logs provides the following audit logs for each Google Cloud project, folder, and organization:

- Admin Activity audit logs
- Data Access audit logs
- System Event audit logs
- Policy Denied audit logs
- **Cloud Logging**
 1. In the Google Cloud console, go to the **Logs Explorer** page.
 2. Select an existing Google Cloud project, folder, or organization.
 3. To display all audit logs, enter either of the following queries into the query-editor field, and then click **Run query**
 - `logName "cloudaudit.googleapis.com"`

- `protoPayload.@type="type.googleapis.com/google.cloud.audit.AuditLog"`
- 4. To display the audit logs for a specific resource and audit log type, in the **Query builder** pane, do the following
 - In **Resource type**, select the Google Cloud resource whose audit logs you want to see.
 - In **Log name**, select the audit log type that you want to see
 - i. For Admin Activity audit logs, select **activity**.
 - ii. For Data Access audit logs, select **data_access**.
 - iii. For System Event audit logs, select **system_event**.
 - iv. For Policy Denied audit logs, select **policy**.
- 5. Click **Run query**.

Cloud Logging Storage (Log Buckets): Use log buckets as containers in your Google Cloud projects, billing accounts, folders, and organizations to store and organize your logs data. The log entries that you store in Cloud Logging are indexed, optimized, and delivered to let you analyze your logs in real time. Cloud Logging buckets are different storage entities than the similarly named Cloud Storage buckets.

“_Required” and “_Default” Log Buckets

For each Google Cloud project, billing account, folder, and organization, Logging automatically creates two log buckets “_Required” and “_Default”.

Cloud Logging automatically routes the following types of log entries to the “_Required” bucket:

- Admin Activity audit logs
- System Event audit logs
- Google Workspace Admin Audit logs
- Enterprise Groups Audit logs
- Login Audit logs
- Access Transparency logs.

Cloud Logging retains the log entries in the “_Required” bucket for 400 days; you can't change this retention period.

Any log entry that isn't stored in the “_Required bucket” is routed to the “_Default” bucket, unless you disable or otherwise edit the “_Default” sink.

For example, Cloud Logging automatically routes the following types of log entries to the “_Default bucket”:

- Data Access audit logs
- Policy Denied audit logs

Cloud Logging retains the log entries in the “**_Default bucket**” for 30 days, unless you configure custom retention for the bucket.

User-Defined Log Buckets: You can also create user-defined log buckets in any Google Cloud project. By applying sinks to your user-defined log buckets, you can route any subset of your log entries to any log bucket, letting you choose the Google Cloud project in which your log entries are stored, and lets you store log entries from multiple resources in one location. You can configure custom retention for the bucket.

- To create a user-defined log bucket:
 1. In the Google Cloud console, go to the Logs Storage page.
 2. Click Create log bucket.
 3. Enter a Name and Description for your bucket.
 4. Optional Upgrade your bucket to use Log Analytics.
 - Select **Upgrade to use Log Analytics**.
 - When you upgrade a bucket to use Log Analytics, you can query your logs in the **Log Analytics** page by using SQL queries. You can also continue to view your logs by using the Logs Explorer.
 5. Optional To select the storage region for your logs, click the **Select log bucket region** menu and select a region. If you don't select a region, then the *global* region is used, which means that the logs could be physically located in any supported region.
 6. Optional To set a custom retention period for the logs in the bucket, click **Next**.
 In the **Retention period** field, enter the number of days, between 1 day and 3650 days, that you want Cloud Logging to retain your logs. If you don't customize the retention period, the default is 30 days.
 You can also update your bucket to apply custom retention after you create it.
 7. Click **Create bucket**.
 After the log bucket is created, Logging upgrades the bucket and creates the dataset link, if these options were selected.

Cloud Storage: Customers that have longer audit log retention requirements can export logs to Cloud Storage. Cloud Storage is an option to provide long-term storage of log entries. The Cloud Storage bucket can be in the same project in which log entries originate, or in a different project.

BigQuery: Customers can export logs to BigQuery and use the capabilities of BigQuery to analyze log data. To do so, you will have to upgrade a log bucket to use Log Analytics, and

then create a linked dataset. With this configuration, Logging stores your log data but BigQuery can read the log data.

- **BigQuery**

1. In the Google Cloud console, go to the **Logs Storage** page.
2. Click **Create log bucket**.
3. Enter a **Name** and **Description** for your bucket.
4. Select **Upgrade your bucket to use Log Analytics**.
5. To view your logs in BigQuery, select **Create a new BigQuery dataset that links to this bucket** and enter a unique dataset name.
6. To select the storage region for your logs, click the **Select log bucket region** menu and select a region. If you don't select a region, then the global region is used, which means that the logs could be physically located in any supported region.
7. To set a custom retention period for the logs in the bucket, click **Next**.
 - In the **Retention period** field, enter the number of days, between 1 day and 3650 days, that you want Cloud Logging to retain your logs. If you don't customize the retention period, the default is 30 days.
 - You can also update your bucket to apply custom retention after you create it.
8. Click **Create bucket**.
 - After the log bucket is created, Logging upgrades the bucket and creates the dataset link, if these options were selected.

When you select this option, BigQuery can read the data stored in your log bucket. You can now query in the BigQuery interface where you can join your log data.

Pub/Sub: Customers can route their Cloud Logging logs to Pub/Sub. Google recommends using Pub/Sub for integrating Cloud Logging logs with third-party software.

Additional Considerations

- You can also collect logs from applications that you write or from your third party applications by using either a client library or a logging agent. Further information is provided in the Supplemental Guidance.

Supplemental Guidance

- [Cloud Logging overview](#)
- [View logs by using the Logs Explorer | Cloud Logging | Google Cloud](#)
- [Cloud Audit Logs overview](#)
- [Configure log buckets | Cloud Logging](#)

- [View logs routed to Cloud Storage](#)
- [Configure log buckets | Cloud Logging](#)
- [View logs routed to Pub/Sub | Cloud Logging | Google Cloud](#)
- [Cloud Logging overview](#)

Control Domain	Audit and Accountability		
Control #	AU.L2-3.3.2		
Control Description	Ensure that the actions of individual system users can be uniquely traced to those users so they can be held accountable for their actions.		
Key Services	• Cloud Audit Logs	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for - Defining the content of the audit records needed to support the ability to uniquely trace users to their actions is defined - Ensuring audit records, once created, contain the defined content Google is responsible for - Ensuring audit records, once created, contain the defined content When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control • Cloud Audit Logs A description of relevant features and implementation guidance is included below. Cloud Audit Logs: Write audit logs that record administrative activities and accesses within your Google Cloud resources. Audit logs record the identity that performed the logged operations on the Google Cloud resource. The caller's identity is held in the "AuthenticationInfo" field of "AuditLog" objects. Cloud Audit Logs provides the following audit logs for each Google Cloud project, folder, and organization: • Admin Activity audit logs			

- Data Access audit logs
- System Event audit logs
- Policy Denied audit logs

Most audit logs can be viewed in **Cloud Logging** by using the Google Cloud console.

1. In the Google Cloud console, go to the **Logs Explorer** page.
2. Select an existing Google Cloud project, folder, or organization.
3. To display all audit logs, enter either of the following queries into the query-editor field, and then click **Run query**
 - **logName "cloudaudit.googleapis.com"**
 - **protoPayload."@type"="type.googleapis.com/google.cloud.audit.AuditLog"**
4. To display the audit logs for a specific resource and audit log type, in the **Query builder** pane, do the following
 - In **Resource type**, select the Google Cloud resource whose audit logs you want to see.
 - In **Log name**, select the audit log type that you want to see
 - i. For Admin Activity audit logs, select **activity**.
 - ii. For Data Access audit logs, select **data_access**.
 - iii. For System Event audit logs, select **system_event**.
 - iv. For Policy Denied audit logs, select **policy**.
5. Click **Run query**.

Additional Considerations

- If additional logging is needed besides the log types listed above, consider routing logs from those applications or services to a central repository. Ensure that the necessary content is captured and forwarded to the central repository.

Supplemental Guidance

- [Cloud Audit Logs overview](#)

Control Domain	Audit and Accountability		
Control #	AU.L2-3.3.3		
Control Description	Review and update logged events		
Key Services	• Logs Explorer • Cloud Monitoring	Control Responsibility	☐ Google ☒ Shared ☐ Customer

Customer Implementation Description

Google Cloud Customers are responsible for

- a. Defining a process for determining when to review logged events
- b. Review event types being logged in accordance with the defined review process
- c. Updating event types being logged based on the review

When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control

- Logs Explorer - Log-Based Alerting Policy
- Cloud Monitoring - Log-Based Alerting Policy

A description of relevant features and implementation guidance is included below.

Logs-Based Alerting Policies are a capability within **Cloud Logging**. Refer to [AU.L2-3.3.1](#) for a description of **Cloud Logging** and general implementation guidance.

Log-Based Alerting Policies: When you want to be notified anytime a specific message occurs in a log entry, use log-based alerting policies. Log-based alerting policies are useful for catching security-related events in log entries. Log-based alerting policies can be configured from Logs Explorer and/or Cloud Monitoring.

Logs Explorer - Log-Based Alerting Policy: You can configure an alerting policy to notify you whenever a specific message appears in your included logs. For example, if you want to know when an audit log records a particular data-access message, you can get notified when the message appears.

- **Implementation**

1. In the Google Cloud console, go to the **Logs Explorer** page.
2. Use the **Query** pane to build a query that matches the message you want to use in your log-based alerting policy.
Click **Run query** to validate the query.
3. In the **Query results** toolbar, expand the **Actions** menu and select **Create log alert**.
4. In the **Alert details** pane, do the following
 - Enter a name for your alerting policy in the **Alert Policy Name** field. For example "Network address invalid IPv4 value".
 - Select an option from the **Policy severity level** menu. Incidents and notifications display the severity level.

- Optional Add documentation for your alerting policy. You can include information that might help the recipient of a notification diagnose the problem.
- 5. To advance to the next step, click **Next**.
- 6. In the **Choose logs to include in the alert pane**, do the following
 - Enter a query and or edit the existing query to filter available logs. Log entries that match the query cause the alerting policy to trigger.
 - To verify the query, click **Preview logs**.
 - Optional Extract log labels. You can create labels from log fields to be displayed in any incidents and notifications created by the alert.
- 7. Click **Next**.
- 8. Select the minimum time between notifications. This value lets you control the number of notifications you get from Monitoring if this condition is met multiple times. For this example, select **5 min** from the options.
- 9. Optional Select the incident autoclose duration. By default, the incident autoclose duration is set to 7 days.
- 10. Click **Next**.
- 11. Select one or more notification channels for your alerting policy. For this example, select an email notification channel.
If you already have an email notification channel configured, then you can select it from the list. If not, click **Manage notification channels** and add an email channel. For information about creating notification channels, see [Create and manage notification channels](#).
- 12. Click **Save**.

Cloud Monitoring - Log-Based Alerting Policy: Cloud Monitoring can be used to observe trends in your logs and to notify you when conditions you describe occur. You can use alerting policies to monitor, in near real time, when a message appears in your log entries. These alerting policies are called log-based alerting policies. You can view, edit, and delete log-based alerting policies by using the Google Cloud console for Monitoring.

- **Implementation**

1. In the Google Cloud console, go to the **Alerting** page
2. To see all policies and to enable filtering, in the **Policies** pane, click **See all policies**.
3. From the **Policies** page you can edit, delete, copy, enable, or disable an alerting policy
 - To edit or copy a policy, click **More options**, and select an option. Editing and copying a policy is similar to the procedure described in [Create a log-based alerting policy](#). You can change and, in some cases, delete the values in the fields. When done, click **Save**.

- You can also edit a log-based alerting policy by clicking its name in the list of policies.
- To delete a policy, click **More options** and select **Delete**. In the confirmation dialog, select **Delete**.
- To enable or disable the alerting policy, click the toggle located under the heading **Enabled**.

Additional Considerations

- There are additional alerting options that can be configured from log entries. For further information on the alerting options and a comparison of those options, refer to the link in the supplemental guidance.

Supplemental Guidance

- [Configure log-based alerting policies | Cloud Logging | Google Cloud](#)
- [Monitor your logs | Cloud Logging](#)

Control Domain	Audit and Accountability								
Control #	AU.L2-3.3.4								
Control Description	Alert in the event of an audit logging process failure								
Key Services	• Logs Explorer • Cloud Monitoring	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for a. Identifying personnel or roles to be alerted in the event of an audit logging process failure b. Defining types of audit logging process failures for which alert will be generated c. Alerting identified personnel or roles in the event of an audit logging process failure When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control • Logs Explorer • Cloud Monitoring									

To address this requirement, establish log-based alerting policies within **Logs Explorer** or **Cloud Monitoring**. These policies should be configured to notify specified personnel or roles upon detection of audit logging process failures or the potential for exceeding maximum storage capacity. Refer to [AU.L2-3.3.3](#) for log-based alerting policy configuration details.

Supplemental Guidance

- [Configure log-based alerting policies | Cloud Logging | Google Cloud](#)

Control Domain	Audit and Accountability		
Control #	AU.L2-3.3.5		
Control Description	Correlate audit record review, analysis, and reporting processes for investigation and response to indications of unlawful, unauthorized, suspicious, or unusual activity		
Key Services	Google SecOps	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for a. Defining audit record review, analysis, and reporting processes for investigation and response to indications of unlawful, unauthorized, suspicious, or unusual activity b. Correlating the defined audit record review, analysis, and reporting processes When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control - Google Security Operations (SecOps) A description of relevant features and implementation guidance is included below. Google Security Operations: Ingests logs from customers, normalizes the data, and detects security alerts. Google Security Operations SIEM provides self-service features around data ingestion, threat detection, alerts, and case management. Google Security Operations can also ingest alerts from other SIEM systems. These alerts are ingested into your Google Security Operations SIEM account, where they can be analyzed.			

The Google Security Operations ingestion mode includes the following types of data ingestion:

- **Ingestion of raw logs into Google Security Operations:** Raw logs are ingested using the Google Security Operations SIEM forwarders, ingestion API, directly from Google Cloud, or using a data feed.
- **Ingestion of alerts generated by other SIEMs:** Alerts generated in other SIEMs are ingested as follows:
 - Google Security Operations ingests alerts from the other SIEM systems, EDRs, or ticketing systems using Google Security Operations SOAR connectors or Google Security Operations SOAR webhooks.
 - Google Security Operations SOAR ingests the events associated with the alerts and creates a corresponding detection.
 - Google Security Operations SOAR processes the alerts and the ingested events.
 - Customers can create detection engine rules to identify patterns in ingested events and generate additional detections.

To use Google SecOps, you must enable Google SecOps audit logging. Google SecOps will write Data Access audit logs and Admin Activity audit logs to the project. You cannot disable Data Access logging using Google Cloud console.

To access audit log configuration options in the Google Cloud console, follow these steps:

1. In the Google Cloud console, go to the **Audit Logs** page
2. In the **Data Access audit logs configuration** table, select one or more Google Cloud services from the **Service** column.
3. In the **Log Types** tab, select the Data Access audit log types that you want to enable for your selected services.
4. Click **Save**.

Supplemental Guidance

- [Google SecOps overview | Google Security Operations](#)
- [Configure a Google Cloud project for Google SecOps | Google Security Operations](#)
- [Google SecOps data ingestion | Google Security Operations](#)
- [Google SecOps audit logging information | Google Security Operations](#)
- [Enable Data Access audit logs | Cloud Logging | Google Cloud](#)

Control #	AU.L2-3.3.6								
Control Description	Provide audit record reduction and report generation to support on-demand analysis and reporting								
Key Services	- Google SecOps SIEM Dashboards - Google SecOps Threat Detection Alerts	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for a. Providing an audit record reduction capability that supports on-demand analysis b. Providing a report generation capability that supports on-demand reporting. When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control: - Google SecOps SIEM Dashboards - Google SecOps Threat Detection Alerts A description of relevant features and implementation guidance is included below. Google SecOps SIEM Dashboards and Threat Detection Alerts are capabilities within Google SecOps. Refer to AU.L2-3.3.5 for a description of Google SecOps and general implementation guidance. Google SecOps SIEM Dashboards: Google Security Operations SIEM dashboards can be used to view and analyze the data in Google Security Operations SIEM, including security telemetry, ingestion metrics, detections, alerts, and IOCs. Google Security Operations SIEM provides you with multiple default dashboards. You can also create custom dashboards. The default dashboards are listed below: - Main - Cloud Detection and Response - Context Aware Detections - Risk - Data Ingestion and Health - IOC Matches - Rule Detections - User Sign In Overview Google SecOps Threat Detection Alerts									

- The **Alerts and IOCs** page displays all the alerts and indicators of compromise (IOC) currently impacting your enterprise. Alerts are tied to data identified as a threat by your security systems. Investigating alerts gives you context about the alert and related entities.
 1. In the Google Cloud Console, go to the **Google SecOps page**
 2. Click **Detection > Alerts and IOCs**
 3. The Alerts tab displays a list of all of the current alerts in your enterprise. Click an alert name in the list to pivot to Alert view. Alert view displays additional information about the alert and its status.
You can view the severity, priority, risk score, and verdict of each alert at a glance. The color-coded icons and symbols help you to quickly identify which alerts need your attention.

Supplemental Guidance

- [Dashboards overview | Google Security Operations](#)
- [Create a custom dashboard | Google Security Operations](#)
- [Understand the Google SecOps platform | Google Security Operations](#)
- [View Alerts and IOCs | Google Security Operations](#)
- [Investigate an alert | Google Security Operations](#)

Control Domain	Audit and Accountability								
Control #	AU.L2-3.3.7								
Control Description	Provide a system capability that compares and synchronizes internal system clocks with an authoritative source to generate time stamps for audit records								
Key Services	● Google Public NTP ● Private Cloud Servers	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for a. internal system clocks are used to generate time stamps for audit records b. an authoritative source with which to compare and synchronize internal system clocks is specified c. internal system clocks used to generate time stamps for audit records are compared to and synchronized with the specified authoritative time source									

Google is responsible for:

- d. internal system clocks are used to generate time stamps for audit records

When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control

- Google Public NTP
- Private Cloud Servers

Note: These features are not navigable from the Google Cloud Console, but they are services made available by Google to configure time configuration and synchronization services for customers to configure their network settings to use Google's public NTP server. Google Cloud logs inherit time configuration settings from Google's internal NTP infrastructure. To address the customer portion of this requirement, the customer is responsible for configuring the servers of any product deployed on Google Cloud to point to the correct NTP servers (with the features listed above).

A description of relevant features and implementation guidance is included below.

Google Public NTP

Google Public NTP is a free, global time service that you can use to synchronize to Google's atomic clocks. It implements the leap smear to smoothly handle leap seconds without disruptions.

Private Cloud Servers

Google recommends that Google Compute Engine VMs are configured to use Google's internal private cloud servers. Compute Engine VMs

Additional Considerations

- Google recommends that customers do not configure Google Public NTP together with non-leap-smearing NTP servers.

Supplemental Guidance

- [Configuring Clients | Public NTP | Google for Developers](#)
- [Configure NTP on a VM | Compute Engine Documentation | Google Cloud](#)

Control #	AU.L2-3.3.8		
Control Description	Protect audit information and audit logging tools from unauthorized access, modification, and deletion		
Key Services	• IAM • Cloud Logging • Google SecOps	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for - Protecting audit information from unauthorized access - Protecting audit logging tools from unauthorized access Google is responsible for: - Protecting audit information from unauthorized modification - Protecting audit information from unauthorized deletion - Protecting audit logging tools from unauthorized modification - Protecting audit logging tools from unauthorized deletion When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control • IAM Roles and Permissions • Cloud Logging • Google SecOps A description of relevant features and implementation guidance is included below. IAM Roles and Permissions are a capability within IAM. Refer to AC.L1-3.1.1 for a description of IAM and implementation guidance. IAM Roles and Permissions - Cloud Logging: To address this requirement, access to sensitive audit data and logging tools should be strictly limited to individuals who have a need-to-know based on their specific roles and responsibilities. IAM provides predefined roles to grant granular access to specific Google Cloud resources such as Cloud Logging. Refer to AU.L2-3.3.1 for a description of Cloud Logging and general implementation guidance. The table in the Supplemental Guidance link lists the predefined roles for Cloud Logging. You can grant the predefined roles at the Google Cloud project level or, in most cases, any type			

higher in the resource hierarchy. For this control, it is important to consider which users are given roles with elevated privileges such as the Logging Admin role, Logs Configuration Writer role, and other IAM roles specific to Cloud Logging.

It is also worth considering that certain Google Cloud Project-level roles will provide permissions to users within Cloud Logging due to the hierarchical structure of IAM. Some of the downstream permissions granted are listed below:

- To give view access to most Google Cloud services, grant the Viewer (**“roles/viewer”**) role.
This role includes all permissions granted by the Logs Viewer (**“roles/logging.viewer”**) role.
- To give editor access to most Google Cloud services, grant the Editor (**“roles/editor”**) role.
This role includes all permissions granted by the Logs Viewer (**“roles/logging.viewer”**) role, and the permissions to write log entries, delete logs, and create log-based metrics. However, this role doesn't let users create sinks, read Data Access audit logs that are in the “**_Default**” bucket, or read logs that are in user-defined log buckets.

For additional details on configuring IAM, refer to [AC.L1-3.1.1](#).

IAM Roles and Permissions - Google SecOps

Refer to [AU.L2-3.3.5](#) for a description of **Google SecOps** and general implementation guidance.

If Google SecOps is implemented in a customer's environment, then access to Google SecOps should be limited to protect audit information and audit logging tools. This can be done with configuring roles and permissions within IAM through IAM policies that define which users and groups have access to Google SecOps features. These IAM policies are defined using predefined roles and permissions provided by Google SecOps or custom roles that you create.

The pre-defined IAM roles and permissions for Google SecOps (**“Chronicle API roles”**) are listed below:

- **Chronicle API Admin** (**“roles/chronicle.admin”**)
 - Full access to the Chronicle API services, including global settings.
- **Chronicle API Data Governor** (**“roles/chronicle.dataGovernor”**)
 - Grants elevated access to control the lifecycle of the Chronicle instance and its data.
- **Chronicle API Editor** (**“roles/chronicle.editor”**)
 - Modify Access to Chronicle API resources.

- **Chronicle API Federation Admin** (“roles/chronicle.federationAdmin”)
 - Full access to Chronicle Federation features.
- **Chronicle API Federation Viewer** (“roles/chronicle.federationViewer”)
 - Readonly access to Chronicle Federation Features.
- **Chronicle API Global Data Access** (“roles/chronicle.globalDataAccess”)
 - Grants global access to data i.e. all data can be accessed.
- **Chronicle API Limited Viewer** (“roles/chronicle.limitedViewer”)
 - Grants read-only access to Chronicle API resources, excluding Rules and Retrohunts.
- **Chronicle API Restricted Data Access** (“roles/chronicle.restrictedDataAccess”)
 - Grants access to data controlled by Data Access Scopes. Intended to be refined by IAM Conditions.
- **Chronicle API Restricted Data Access Viewer** (“roles/chronicle.restrictedDataAccessViewer”)
 - Grants read only access to Chronicle API resources without global data access scope.
- **Chronicle SOAR Admin** (“roles/chronicle.soarAdmin”)
 - Grants admin access to Chronicle SOAR.
- **Chronicle SOAR Threat Manager** (“roles/chronicle.soarThreatManager”)
 - Grants threat manager access to Chronicle SOAR.
- **Chronicle SOAR Vulnerability Manager** (“roles/chronicle.soarVulnerabilityManager”)
 - Grants vulnerability manager access to Chronicle SOAR.
- **Chronicle API Viewer** (“roles/chronicle.viewer”)
 - Read-only access to the Chronicle API resources.

For additional details on configuring IAM, refer to [AC.L1-3.1.1](#).

Google SecOps User Management per Module: Provide an additional layer of access control for the various modules of the platform. Access is based on a combination of SOC roles, environments, and permission groups. Within your organization, you can experiment with the permission levels, environments, or environment groups to determine the full scope of responsibility for each user group in the Google Security Operations platform.

SOC Roles

The Google SecOps platform comes with predefined SOC roles but customized roles can be added. The predefined SOC roles are defined as follows:

- Tier 1 performs basic triage on the alerts
- Tier 2 reviews high priority security threats
- Tier 3 handles major incidents
- SocManager manages the SOC team

- CISO top level manager within your organization
- Administrator has access to the entire Google SecOps platform

One of these SOC roles is used as a default role to be automatically assigned to incoming cases, and cases can be reassigned to other SOC roles.

Environments

You can define different environments and environment groups to create data segregation. This ensures a logical separation between the majority of the platform modules such as cases, playbooks, ingestion, and dashboards. You can experiment with the platform configuration settings so that only analysts who are associated with a specific environment or environment group can see cases from that environment or environment group.

Permission Groups

The permission groups decide the level of access each group has to different modules and settings in the platform. The settings are on a granular level. The Google Security Operations platform comes with predefined permission groups but more can always be added. The predefined groups are:

- Admin
- Basic
- Readers
- View Only
- Collaborators
- Managed
- Managed-Plus

Additional Considerations

- Google's Workforce Identity Federation lets you grant on-premises or multi cloud workloads access to Google Cloud resources, without having to use a service account key. You can use Workforce Identity Federation with any IdP that supports OpenID Connect (OIDC), or SAML 2.0, such as Microsoft Entra ID, Active Directory Federation Services (AD FS), Okta, and others.
- Google SecOps provides the ability to grant permissions to specific users to use only the SIEM features in Google SecOps (such as investigating raw data) or only the SOAR features of Google SecOps (such as managing cases).

Supplemental Guidance

- [Access control with IAM | Cloud Logging](#)
- [Control access to platform with SOC roles, environments, and groups | Google Security Operations](#)

- [Configure a third-party identity provider | Google Security Operations](#)
- [Add SIEM or SOAR users to Google SecOps](#)

Control Domain	Audit and Accountability								
Control #	AU.L2-3.3.9								
Control Description	Limit management of audit logging functionality to a subset of privileged users								
Key Services	• IAM • Cloud Logging • Google SecOps	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for a. Defining a subset of privileged users granted access to manage audit logging functionality b. Limiting management of audit logging functionality to the defined subset of privileged users When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control • IAM Roles and Permissions • Cloud Logging • Google SecOps To address this requirement, limit management of audit logging functionality to the pre-defined subset of privileged users within Cloud IAM for Cloud Logging and Google SecOps. These users should have the appropriate roles and responsibilities that warrant being granted the permissions to manage audit logging functionality. Refer to AU.L2-3.3.8 for configuration guidance for the roles that can be assigned to the privileged users who manage the system’s audit logging functionality.									
Supplemental Guidance									
• Access control with IAM Cloud Logging • Control access to platform with SOC roles, environments, and groups Google Security Operations • Configure a third-party identity provider Google Security Operations									

- [Add SIEM or SOAR users to Google SecOps](#)

Control Domain	Configuration Management								
Control #	CM.L2-3.4.1								
Control Description	Establish and maintain baseline configurations and inventories of organizational systems (including hardware, software, firmware, and documentation) throughout the respective system development life cycles								
Key Services	Cloud Asset Inventory	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for a. Establishing a system inventory b. Ensuring the system inventory includes hardware, software, firmware, and documentation c. Maintaining, reviewing, and updating the inventory throughout the system development life cycle When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control: - Cloud Asset Inventory <u>Establish a System Inventory</u> Cloud Asset Inventory: Cloud Asset Inventory is a feature available through the Google Cloud Console that would support the documentation and maintenance of an inventory for assets deployed in Google Cloud. Cloud Asset Inventory is a global metadata inventory service that lets you view, search, export, monitor, and analyze your Google Cloud asset metadata, with up to 35 days of create, update, and delete history. Assets that haven't changed in the past 35 days report their latest status. Asset metadata can come from the following places: Google Cloud resources, such as Compute Engine VM instances, Cloud Storage buckets, and App Engine instances.									

- **Policies** set on Google Cloud resources, such as IAM policies, organization policies, and Access Context Manager policies.
- **Runtime information** from OS inventory management.

To view your list of assets in the Google Cloud Console, follow these steps

1. In the Google Cloud console, click **Activate Cloud Shell**.
2. In Cloud Shell, run the command that corresponds with the asset type that you would like to view - [gcloud CLI reference for listing all options of Cloud Assets](#)
 - To list all the VM instances in a specific project, run the following command:

```
"gcloud asset list \  
  --project=PROJECT_ID \  
  --asset-types=compute.googleapis.com/Instance \  
  --content-type=resource"
```

Replace PROJECT_ID with the ID of the project whose assets you want to list. In the previous code sample, an asset type of **"compute.googleapis.com/Instance"** is used to only list Compute Engine VM instances.

A content type of **"resource"** has also been set. This specifies that resource metadata should also be returned in the response. If no content type is set, then only basic information about each asset is returned, such as the asset name, the last time it was updated, and what project it's in.

Baseline Configurations for Google Cloud

When establishing a baseline for Google Cloud, it is best to implement secure configuration settings. You can leverage CIS benchmarks, DISA STIGs, and/or configurations in this Google Cloud CMMC Implementation Guide to create a documented, maintained, reviewed, and updated baseline. Although not a feature that is accessed via the Google Cloud console, the CIS Google Cloud Computing Platform Foundation Benchmark offers specific technical configuration guidelines and is publicly available. Refer to the [supplemental guidance](#) for the CIS Google Cloud Computing Platform Benchmark.

Baseline Configurations for Products Deployed on Google Cloud

Depending on the Google Cloud product feature deployed, the approach to satisfying this control will differ. It is important to understand that for each product deployed, you will need to inventory the necessary components, establish baseline configurations and implement a process to document, review and update those baseline configurations.

Implementation guidance for each Google Cloud product feature is outside of the scope of this guide; however, see below for a few examples to demonstrate where and how baselines can be established and configured at the Google Cloud product level:

- *Google Compute Engine* lets you create and run instances on Google infrastructure.
 - Machine, OS, networking and security configurations can be customized in a multitude of ways. Refer to the [supplemental guidance](#) for more information on creating instances.
- *Google Kubernetes Engine*
 - *Artifact Registry* lets you store your approved container images, OS packages (if using private package repositories), and other software artifacts that form part of your system baselines.
- *Cloud Deployment Manager* allows you to specify all the resources needed and create reusable templates for your applications in a declarative format using YAML
 - Store your Deployment Manager configurations and templates in a version control system.

Supplemental Guidance

- [Cloud Asset Inventory - View your assets](#)
- [gcloud CLI reference for listing all options of Cloud Assets](#)
- [CIS Google Cloud Platform Foundation Benchmark](#)
- [DISA STIGS](#)
- [Google Compute Engine - Create an instance](#)
- [Google Cloud Deployment Manager documentation](#)
- [Artifact Registry documentation | Google Cloud](#)

Control Domain	Configuration Management		
Control #	CM.L2-3.4.2		
Control Description	Establish and enforce security configuration settings for information technology products employed in organizational systems		
Key Services	• Organization Policy Service	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			

Google Cloud Customers are responsible for:

- a. Establishing security configuration settings for Google Cloud and including them in the baseline configuration; and
- b. Enforcing security configuration settings for Google Cloud

When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control:

- Organization Policy Service

A description of relevant features and implementation guidance is included below.

To address this control, Customers should establish baseline configurations that incorporate security configuration settings. Refer to [CM.L2-3.4.1](#) for details on establishing **Baseline Configurations for Google Cloud** and references for incorporating CIS benchmarks, DISA STIGs and/or the configurations provided in this Google Cloud CMMC Implementation Guide.

Organization Policy Service: Configure a single constraint that restricts one or more Google Cloud services. The organization policy is set on an organization, folder, or project resource to enforce the constraint on that resource and any child resources. An organization policy contains one or more rules that specify how, and whether, to enforce the constraint.

Benefits of Organization Policies:

- Centralize control to configure restrictions on how your organization's resources can be used.
- Define and establish guardrails for your development teams to stay within compliance boundaries.
- Help project owners and their teams move quickly without worry of breaking compliance.

Additionally, when a Google Cloud service acts or is in a state that is counter to the organization policy restriction configuration within the scope of its resource hierarchy then it is considered a violation. Google Cloud services will enforce constraints to prevent violations, but the application of new organization policies is usually not retroactive. If an organization policy constraint is retroactively enforced, it will be labeled as such on the organization policy constraints page.

If a new organization policy sets a restriction on an action or state that a service is already in, the policy is considered to be in violation, but the service won't stop its original behavior. You

will need to address this violation manually. This prevents the risk of a new organization policy completely shutting down your business continuity.

- **Creating and editing a boolean policy:**

1. In the Google Cloud console, go to the **Organization policies** page.
2. From the project picker, select the project, folder, or organization for which you want to edit organization policies.
3. The **Organization policies** page displays a filterable list of organization policy constraints that are available.
4. Select a constraint from the list on the **Organization policies** page. The **Policy details** page that appears describes the constraint and provides information about how the constraint is applied.
5. To update the organization policy for this resource, click **Manage policy**.
6. On the **Edit policy** page, select **Override parent's policy**.
7. Select **Add a rule**.
8. Under **Enforcement**, select whether enforcement of this organization policy should be on or off.
9. To enforce the policy, click **Set policy**.

- **Creating and editing a list constraint:**

1. In the Google Cloud console, go to the **Organization policies** page.
2. From the project picker, select the project, folder, or organization for which you want to edit organization policies.
3. The **Organization policies** page displays a filterable list of organization policy constraints that are available.
4. Select a constraint from the list on the **Organization policies** page. The **Policy details** page that appears describes the constraint and provides information about how the constraint is applied.
5. To update the organization policy for this resource, click **Manage policy**.
6. On the **Edit policy** page, select **Override parent's policy**.
7. Under **Policy enforcement**, select an enforcement option:
8. To merge and evaluate the organization policies together, select **Merge with parent**. For more information about inheritance and the resource hierarchy, see Understanding Hierarchy Evaluation.
9. To override the inherited policies completely, select **Replace**.
10. Select **Add a rule**.
11. Under **Policy values**, select whether this organization policy allows all values, denies all values, or specifies a custom list.
12. If you specify a custom list of values, then under **Policy type**, select whether the given values should be accepted or denied by the organization policy.

Control Domain	Configuration Management		
Control #	CM.L2-3.4.3		
Control Description	Track, review, approve or disapprove, and log changes to organizational systems		
Key Services	- Cloud Audit Logs - Organization Policy Service - Security Command Center	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Tracking changes to the system - Reviewing changes to the system - Approving or disapproving changes to the system - Logging changes to the system Google is responsible for: - Logging changes to the system When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control: - Cloud Audit Logs - Organization Policy Service - Security Command Center - Security Posture (<i>Enterprise or Premium Tier feature</i>)			

A description of relevant features and implementation guidance is included below.

To address this control, Customers should implement a configuration change control process that tracks, reviews, approves/disapproves and logs changes. There are no singular key services within the Google Cloud console that establishes a holistic configuration change control process; however, Cloud Logging, Organizational Policy Service, and Security Command Center can support the implementation of this control.

Cloud Audit Logs: Writes Admin Activity audit logs, which record operations that modify the configuration or metadata of a resource. You can't disable Admin Activity audit logs. Cloud Audit Logs would provide tracking capabilities for changes that occur within supported Google Cloud resources.

Also refer to [AU.L2-3.3.2](#) for a description of **Cloud Audit Logs** and general implementation guidance.

Organization Policy Service: Use an organization policy in dry-run mode to monitor how a policy change would impact your workflows before it is enforced. The dry-run functionality provides customers with the capability to review and evaluate the effects of an organization policy change on their Google Cloud organization before it is implemented and enforced.

Refer to [CM.L2-3.4.2](#) for a description of **Organization Policy Service** and general implementation guidance.

- **Create an organization policy in dry-run mode for a Boolean constraint:**
 1. In the Google Cloud console, go to the **Organization policies** page.
 2. From the project picker, select the resource for which you want to set the organization policy.
 3. Select the **Restrict Resource Service Usage** constraint from the list on the **Organization policies** page.
 4. Select the **Dry run** tab.
 5. Click **Manage dry run policy**.
 6. On the **Edit dry run policy** page, select **Override parent's policy**.
 7. Click **Add rule**.
 8. Under **Enforcement**, select **On**, and then click **Done**.
 9. To enforce the organization policy in dry-run mode, click **Set dry run policy**. Once you verify that the organization policy in dry-run mode works as intended, you can set the live policy by clicking **Set policy**.

You can verify the status of your organization policy in dry-run mode by going to the **Dry run** tab of an organization policy constraint.

For projects that have an organization policy in dry-run mode applied to them, you can see the audit logs by clicking **View rejection logs**. For this organization policy, the audit logs display violations as if the custom organization policy is enforced.

- **Create an organization policy in dry-run mode for a list constraint:**
 1. In the Google Cloud console, go to the **Organization policies** page.
 2. From the project picker, select the resource for which you want to set the organization policy.
 3. Select the **Restrict Resource Service Usage** constraint from the list on the **Organization policies** page.
 4. Select the **Dry run** tab.
 5. Click **Manage dry run policy**.
 6. On the **Edit dry run policy** page, select **Override parent's policy**.
 7. Under **Policy enforcement**, click **Replace**.
 8. Click **Add rule**.
 9. From **Policy values**, select **Custom**.
 10. From **Policy type**, select **Deny**.
 11. In the **Custom values** box, enter “**compute.googleapis.com**”, and then click **Done**.
 12. If this is a custom constraint, you can click **Test changes** to simulate the effect of this organization policy. For more information, see [Test organization policy changes with Policy Simulator](#).
 13. To enforce the organization policy in dry-run mode, click **Set dry run policy**. You can also set the live policy by clicking **Set policy**.

You can verify the status of your organization policy in dry-run mode by going to the **Dry run** tab of an organization policy constraint.

For projects that have an organization policy in dry-run mode applied to them, you can see the audit logs by clicking **View rejection logs**.

Security Command Center - Security Posture

The security posture service is only available to you if you purchase a subscription of the Security Command Center Premium tier or the Enterprise tier and activate Security Command Center at the organization level.

The security posture service is a built-in service for the Security Command Center that lets you define, assess, and monitor the overall status of your security in Google Cloud. A security posture helps you detect and mitigate any drift from your defined benchmark.

To enforce a posture with all its policies on a Google Cloud resource, you deploy the posture. After you deploy your posture, you can monitor your environment for any drift from your defined posture. Security Command Center reports instances of drift as findings that you can review, filter, and resolve. In addition, you can export these findings in the same way that you export any other findings from the Security Command Center.

Supplemental Guidance

- [Google Cloud services with audit logs](#)
- [Create a dry-run organization policy](#)
- [Security posture service](#)
- [Manage a security posture](#)

Control Domain	Configuration Management		
Control #	CM.L2-3.4.4		
Control Description	Analyze the security impact of changes prior to implementation		
Key Services	• Organization Policy Service	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Analyzing the security impact of changes prior to implementation When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control: • Organization Policy Service To address this control, Customers should implement a configuration change control process that includes an analysis of the security impact and ramifications of proposed changes prior to the change being implemented. There are no singular key services within the Google Cloud console that establishes a holistic configuration change control process;			

however, Organizational Policy Service can support the implementation of this control. The Organizational Policy Service offers a dry-run mode which allows Customers to assess and determine the potential impacts of changes to an organizational policy.

Refer to [CM.L2-3.4.3](#) for a description of **Organization Policy Service - dry-run mode** and general implementation guidance.

Supplemental Guidance

- [Create a dry-run organization policy](#)

Control Domain	Configuration Management		
Control #	CM.L2-3.4.5		
Control Description	Define, document, approve, and enforce physical and logical access restrictions associated with changes to organizational systems		
Key Services	• IAM	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining logical access restrictions associated with changes to the system - Documenting logical access restrictions associated with changes to the system - Approving logical access restrictions associated with changes to the system - Enforcing logical access restrictions associated with changes to the system Google is responsible for: - Defining physical access restrictions associated with changes to the system - Documenting physical access restrictions associated with changes to the system - Approving physical access restrictions associated with changes to the system - Enforcing physical access restrictions associated with changes to the system When configured correctly, the following key service(s) in the Google Cloud Console may be used to support this control: • Identity and Access Management (IAM)			

To address this control, define and document the type of changes that require logical access restrictions. Once that is determined, Customers can leverage the functionality of IAM to approve and enforce logical access restrictions associated with changes to the system. IAM is a tool to manage fine-grained authorization for Google Cloud, letting you control who can do what on which resources. By assigning roles in IAM according to job functions, organizations can ensure that only authorized individuals are permitted to make system changes. You can leverage pre-defined roles in IAM that are created and maintained by Google.

Refer to [AC.L1-3.1.1](#) for a description of **IAM** and general implementation guidance.

Supplemental Guidance

- [Grant an IAM role by using the Google Cloud console](#)
- [Predefined roles and permissions](#)
- [IAM permissions reference](#)

Control Domain	Configuration Management								
Control #	CM.L2-3.4.6								
Control Description	Employ the principle of least functionality by configuring organizational systems to provide only essential capabilities								
Key Services	• Custom VM Images / Hardened OS Images • Shielded VMs • Serverless Platforms • Organization Policy Service • VPC Firewall Rules • GKE features • IAM • API Gateway / Cloud Endpoints • Enabled APIs management per project	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Defining essential system capabilities based on the principle of least functionality									

b. Configuring the system to provide only the defined essential capabilities

When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control:

- Custom VM Images / Hardened OS Images (Compute Engine)
- Shielded VMs (Compute Engine)
- Serverless Platforms (Cloud Run, Cloud Functions, App Engine)
- Organization Policy Service
- VPC Firewall Rules
- Google Kubernetes Engine (GKE) features (e.g., Autopilot, Policy Controller, Security Posture Dashboard)
- IAM (Identity and Access Management)
- API Gateway / Cloud Endpoints
- Enabled APIs management per project

A description of relevant features and implementation guidance is included below.

Custom VM Images / Hardened OS Images (Compute Engine)” Start with a minimal operating system and remove or disable unnecessary software, services, libraries, and ports to reduce the attack surface.

- **Use Google-Provided Hardened Images**
 - When creating a VM in **Compute Engine > VM instances > Create Instance**, under Boot disk, click **Change**.
 - Select an operating system. Google provides various OS images, including security-hardened versions of Container-Optimized OS or Shielded VM compatible images.
- **Create Custom/Hardened Images**
 - Launch a VM from a standard Google-provided OS image.
 - Connect to the VM and manually remove unnecessary packages, disable unneeded services (e.g., `systemctl disable <service-name>`), close unused ports (via OS firewall), and apply other security hardening configurations based on security benchmarks (e.g., CIS Benchmarks).
 - After hardening, create a custom image from this VM's boot disk. Stop the VM instance. Go to **Compute Engine > Storage > Images > Create Image**.
 - Name the image, select "Disk" as the Source, and choose the hardened VM's boot disk.
 - Use this custom hardened image as the source for future VM deployments.
- **Regularly Update and Patch Images** Ensure that custom images are regularly updated with security patches and that configurations are periodically reviewed for continued adherence to least functionality.

Shielded VMs (Compute Engine): Shielded VMs offer verifiable integrity of your Compute Engine VM instances, ensuring they haven't been compromised by boot-level or kernel-level malware or rootkits.

- **Enable Shielded VM Features**

- When creating a VM, in the "Shielded VM" section, ensure "Enable Secure Boot," "Enable vTPM," and "Enable Integrity Monitoring" are checked.
- Secure Boot helps ensure that the system only runs authentic software.
- vTPM enables Measured Boot by validating boot components.
- Integrity Monitoring helps you understand and make decisions about the integrity of your VM instances by comparing baseline boot measurements against subsequent boot measurements. Findings are available in Security Command Center.

Serverless Platforms (Cloud Run, Cloud Functions, App Engine): These platforms abstract away the underlying OS and infrastructure management, inherently reducing the configurable surface area. You deploy code or containers, and Google manages the infrastructure with a focus on security and essential capabilities.

- **Cloud Run / Cloud Functions**

- Deploy your application or function. The environment is managed by Google.
- Configure specific settings such as environment variables, memory allocation, and concurrency as needed for the function, but OS-level hardening is largely handled by Google.
- Use IAM permissions to control who can invoke the function/service and what Google Cloud APIs the function/service itself can access via its runtime service account.

- **App Engine (Standard Environment)**

- Deploy applications in a sandboxed environment. Runtimes are curated and secured by Google.
- Configuration is primarily at the application level (app.yaml).

Organization Policy Service: Enforce constraints across your Google Cloud resource hierarchy, which can be used to restrict available functionality.

- **Disable Unnecessary Service Usage**

- Navigate to **IAM & Admin > Organization Policies**.
- Use the constraints/gcp.restrictServiceUsage constraint to define which Google Cloud APIs and services can be used within an organization, folder, or project. Create an allowlist or denylist of services.
- For example, if your organization does not use Bigtable, you can disable the bigtable.googleapis.com service for all projects.

- **Restrict VM Capabilities**
 - Use constraints like `constraints/compute.disableSerialPortAccess`, `constraints/compute.vmExternallpAccess` (to disallow external IPs on VMs), or `constraints/compute.disableNestedVirtualization`.
- **Restrict Service Account Key Creation**
 - Enforce `constraints/iam.disableServiceAccountKeyCreation` to prevent the creation of external service account keys, promoting more secure authentication methods.
 - Changes to organization policies should be documented and approved.

VPC Firewall Rules: Restrict network traffic to and from your VMs to only what is essential.

- **Implement Default Deny**
 - Google Cloud VPC firewalls are stateful and have an implied deny all ingress and allow all egress rule. Explicitly create rules to allow only necessary inbound traffic on specific ports and protocols to your VMs from specific sources.
 - Navigate to **VPC network > Firewall > Create Firewall Rule**.
 - Specify Target tags or Target service accounts to apply the rule to specific VMs.
 - For Source filter, use IP ranges or Source tags/service accounts, preferring specific sources over 0.0.0.0/0.
 - Under Protocols and ports, specify only the essential ones (e.g., tcp 443 for HTTPS).
- **Restrict Egress Traffic**
 - While the default is “allow all egress,” consider creating more restrictive egress rules if VMs only need to connect to specific internal or external destinations.

Google Kubernetes Engine (GKE): Configure GKE clusters and workloads for least functionality.

- **Use GKE Autopilot**
 - Autopilot clusters are managed by Google with pre-configured security best practices, reducing the operational burden and attack surface related to node configuration and cluster setup.
- **GKE Security Posture Dashboard**
 - Navigate to **Security > Security Posture** in the Google Cloud Console.
 - Review findings for GKE clusters, which can highlight workload configurations that deviate from security best practices (e.g., privileged containers, hostPID/hostIPC usage).

- **GKE Policy Controller (based on Open Policy Agent Gatekeeper)**
 - If not using Autopilot, or for finer-grained control, install Policy Controller on your clusters.
 - Define and enforce policies (constraints) that restrict pod capabilities, such as disallowing privileged containers, host network access, or specific volume types.
 - Navigate to **Kubernetes Engine > Security & Policy > Policy** to manage Policy Controller.
- **Role-Based Access Control (RBAC)**
 - Configure Kubernetes RBAC to grant users and service accounts within the cluster only the permissions necessary for their tasks.
- **Minimal Base Images for Containers** Use minimal, distroless, or scratch base images for your containers to reduce the included OS packages and libraries.

IAM (Identity and Access Management): Apply the principle of least privilege to identities, limiting the capabilities they can invoke on Google Cloud resources.

- **Grant Minimal Roles**
 - Assign users, groups, and service accounts predefined or custom IAM roles that grant only the permissions essential for their tasks. Avoid using basic roles like Owner (`roles/owner`) broadly.
- **Service Account Permissions**
 - When a service (e.g., Compute Engine VM, Cloud Function) needs to access other Google Cloud APIs, ensure its associated service account has only the permissions required for its intended functionality. For example, if a VM only needs to read from a Cloud Storage bucket, grant its service account the Storage Object Viewer role, not Storage Admin.

API Gateway / Cloud Endpoints: When exposing APIs, ensure only intended functionality is accessible.

- **Define Specific Methods/Paths**
 - Configure API Gateway or Cloud Endpoints to expose only specific HTTP methods (GET, POST, etc.) and resource paths of your backend services.
 - Implement authentication and authorization at the gateway level to restrict who can access the exposed functionality.

Enabled APIs Management Per Project: Only enable the Google Cloud APIs that are strictly necessary for the applications and services within a specific project.

- **Review and Disable Unused APIs**
 - Navigate to **APIs & Services > Library** or **APIs & Services > Enabled APIs**.

- Review the list of enabled APIs for each project.
- If an API is enabled but not used, disable it to reduce the project's potential attack surface and functionality. Click on the API and then **Disable API**. This requires the serviceusage.services.disable permission.

Additional Considerations

- **Regular Audits:** Periodically audit system configurations, firewall rules, IAM policies, and enabled APIs to ensure they continue to adhere to the principle of least functionality. Remove or disable any capabilities that are no longer essential.
- **Change Management:** Integrate checks for least functionality into your change management process. When a new system is deployed or an existing one is modified, verify that only essential capabilities are enabled.
- **Documentation:** Document the rationale for enabled capabilities and configurations, especially where they deviate from a more restrictive default.

Supplemental Guidance

- [Grant an IAM role by using the Google Cloud console](#)
- [Predefined roles and permissions](#)
- [IAM permissions reference](#)
- [IAM – Understanding roles](#)
- [API Gateway overview](#)
- [Enabling and disabling services](#)
- [GKE Autopilot overview](#)
- [GKE Security Posture dashboard](#)
- [GKE Policy Controller](#)
- [Creating and managing custom images](#)
- [Shielded VM](#)
- [Serverless on Google Cloud](#)
- [Organization Policy Service constraints](#)
- [VPC firewall rules overview](#)

Control Domain	Configuration Management		
Control #	CM.L2-3.4.7		
Control Description	Restrict, disable, or prevent the use of nonessential programs, functions, ports, protocols, and services.		
Key Services	● OS Hardening ● VPC Firewall Rules	Control Responsibility	☐ Google ☒ Shared

	• Organization Policy Service • GKE Policy Controller • Binary Authorization • Serverless Platforms • Enabled APIs management per project • IAM		☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining essential programs - Defining the use of nonessential programs - Restricting, disabling or preventing the use of nonessential programs as defined - Defining essential functions - Defining the use of nonessential functions - Restricting, disabling or preventing the use of nonessential functions as defined - Defining essential ports - Defining the use of nonessential ports - Restricting, disabling, or preventing the use of nonessential ports as defined - Defining essential protocols - Defining the use of nonessential protocols - Restricting, disabling or preventing the use of nonessential protocols as defined - Defining essential services - Defining the use of nonessential services - Restricting, disabling, or preventing the use of nonessential services as defined When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control: • OS Hardening within Custom VM Images (Compute Engine) • VPC Firewall Rules • Organization Policy Service • Google Kubernetes Engine (GKE) Policy Controller (and built-in security contexts) • Binary Authorization • Serverless Platforms (Cloud Run, Cloud Functions, App Engine) • Enabled APIs management per project			

- Identity and Access Management (IAM) (indirectly, by limiting permissions to use/configure services)

A description of relevant features and implementation guidance is included below.

OS Hardening within Custom VM Images (Compute Engine): Configure the operating system on your VMs to ensure nonessential software and services are removed or disabled.

- **Create and Maintain Hardened Images**
 - Start with a minimal base OS image provided by Google Cloud.
 - Connect to a template VM instance and
 - Uninstall unnecessary software packages and applications.
 - Disable system services that are not required for the VM's specific role (e.g., using `systemctl disable <service-name>` on Linux).
 - Configure host-based firewalls (like `ufw` or `firewalld` on Linux, Windows Firewall) to block all nonessential inbound and outbound ports and protocols, complementing VPC firewall rules.
 - After hardening, create a custom image from this VM's boot disk (**Compute Engine > Storage > Images > Create Image**).
 - Use this custom hardened image for deploying new VMs. This directly prevents nonessential programs and services from being present or active.

VPC Firewall Rules VPC firewall rules are critical for preventing network traffic on nonessential ports and protocols.

- **Implement Deny-by-Default Ingress Rules**
 - Create explicit allow rules for only essential inbound ports, protocols, and source IP ranges/tags/service accounts.
 - The implicit deny-all ingress rule will then block all other nonessential traffic.
 - Navigate to **VPC network > Firewall > Create Firewall Rule**.
 - Example If a web server only needs HTTPS, create an allow rule for tcp 443 from required sources. All other ports (e.g., tcp 80 if not redirecting, tcp 22 unless bastion access is specifically configured and restricted) would be blocked by the implicit deny.
- **Implement Egress Restrictions**
 - Create explicit allow rules for essential outbound traffic on specific ports and protocols to specific destinations.
 - If default allow-all egress is too permissive, create a lower-priority deny-all egress rule and then higher-priority allow rules for specific necessary outbound connections. This prevents VMs from initiating connections on nonessential ports/protocols.

Organization Policy Service This service can enforce constraints that disable or restrict nonessential Google Cloud services, APIs, or features across your organization, folders, or projects.

- **Prevent Use of Nonessential Cloud Services/APIs**
 - Navigate to **IAM & Admin > Organization Policies**.
 - Utilize the constraints/gcp.restrictServiceUsage constraint. Configure it to create an "allowlist" of only essential Google Cloud services and APIs, effectively preventing the use of all others. Alternatively, use a "denylist" for specific nonessential services.
- **Disable Specific Service Functions**
 - Use constraints like constraints/compute.disableSerialPortAccess to prevent serial port access if not needed.
 - Enforce constraints/iam.disableServiceAccountKeyCreation to prevent the generation of static service account keys, which can be a nonessential risk if alternatives exist.
 - Prevent public IP addresses on VMs using constraints/compute.vmExternallpAccess if internal access is sufficient.
 - These policies actively prevent the enabling or use of these functions.

Google Kubernetes Engine (GKE) Policy Controller For GKE clusters, Policy Controller (using OPA Gatekeeper) allows you to define and enforce policies that restrict or prevent nonessential behaviors and capabilities within your pods and containers.

- **Define and Apply Restrictive Policies**
 - Install Policy Controller in your GKE clusters (**Kubernetes Engine > Security & Policy > Policy**).
 - Implement policies from the provided library or create custom constraints to
 - Prevent privileged containers (k8spspprivilegedcontainer).
 - Disallow use of host networking or host ports (k8spsphostnetworkingports).
 - Restrict use of specific volume types (e.g., hostPath volumes).
 - Forbid running containers as root user or with specific Linux capabilities.
 - These policies actively prevent pods with nonessential (and potentially risky) configurations from being scheduled.

Binary Authorization This service ensures that only trusted and verified container images are deployed on GKE or Cloud Run, effectively preventing the deployment and execution of containers that might contain nonessential or unauthorized programs.

- **Create and Enforce Attestation Policies**

- Navigate to **Security > Binary Authorization**.
- Define a policy that specifies attestors (e.g., security scanners, QA sign-off) that must verify an image before it can be deployed.
- Images without the required attestations (meaning they haven't been vetted to contain only essential programs) will be blocked from deployment.
- This prevents nonessential programs that might be bundled in unverified container images.

Serverless Platforms (Cloud Run, Cloud Functions, App Engine) These platforms inherently restrict the execution environment, preventing the use of arbitrary nonessential programs, ports, or services beyond what is defined for the deployed application or function.

- **Leverage Managed Runtimes**

- By deploying to these platforms, you rely on Google to manage the underlying OS and infrastructure, which is configured for essential capabilities only.
- You configure your application's specific dependencies and minimal required settings, not the broad OS environment.

Enabled APIs Management Per Project Each Google Cloud project can have various APIs enabled. Disabling those not essential for the project's purpose restricts available functionality.

- **Audit and Disable Nonessential APIs**

- Go to **APIs & Services > Enabled APIs** for each project.
- Review the list of APIs. If an API is not actively used or required for the project's workload, click on the API and then **Disable API**.
- This prevents accidental or unauthorized use of these APIs and their associated functions and services.

Identity and Access Management (IAM) While primarily for authorization, IAM contributes by preventing identities from using or configuring services if they lack permissions. If a service's API cannot be called, its functions are effectively disabled for that identity.

- **Apply Least Privilege to Service Accounts**

- Ensure that service accounts used by applications or VMs have only the IAM roles and permissions necessary for their specific tasks. If a service account for a VM does not have permissions for `pubsub.googleapis.com`, that VM cannot use Pub/Sub services, effectively disabling that service for the VM's application logic.

Additional Considerations

- **Continuous Monitoring and Review** Regularly review system configurations, firewall logs (for denied traffic), Organization Policy compliance, and API usage to identify and disable any nonessential elements that may have been overlooked or become unnecessary over time.
- **Change Control** Incorporate a step in your change management process to verify that new or modified systems do not introduce nonessential programs, functions, ports, protocols, or services.
- **Application Whitelisting (OS Level)** For VMs, consider implementing application whitelisting technologies (e.g., AppLocker on Windows, or third-party tools) to ensure only explicitly approved programs can execute. This is an advanced form of preventing nonessential programs.

Supplemental Guidance

- [Grant an IAM role by using the Google Cloud console](#)
- [Predefined roles and permissions](#)
- [IAM permissions reference](#)
- [Creating and managing custom images](#)
- [VPC firewall rules](#)
- [Organization Policy Service constraints](#)
- [GKE Policy Controller](#)
- [Binary Authorization overview](#)
- [Enabling and disabling services \(APIs\)](#)
- [Google Cloud Serverless options](#)
- [IAM Predefined Roles](#)

Control Domain	Configuration Management		
Control #	CM.L2-3.4.8		
Control Description	Apply deny-by-exception (blacklisting) policy to prevent the use of unauthorized software or deny-all, permit-by-exception (whitelisting) policy to allow the execution of authorized software		
Key Services	N/A	Control Responsibility	☐ Google ☐ Shared ☒ Customer

Customer Implementation Description
Google Cloud customers are responsible for: - Specifying a policy specifying whether whitelisting or blacklisting is to be implemented - Specifying the software allowed to execute under whitelisting or denied use under blacklisting - Implementing whitelisting to allow the execution of authorized software or blacklisting to prevent the use of unauthorized software <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>
Supplemental Guidance
N/A

Control Domain	Configuration Management		
Control #	CM.L2-3.4.9		
Control Description	Control and monitor user-installed software.		
Key Services	N/A	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description	Google Cloud Customers are responsible for: - Establishing a policy for controlling the installation of software by users - Controlling the installation of software by users - Monitoring the installation of software by users <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>		
Supplemental Guidance			

- N/A

Control Domain	Identification and Authentication								
Control #	IA.L1-3.5.1								
Control Description	Identify system users, processes acting on behalf of users, and devices								
Key Services	• Cloud Identity / Google Workspace • IAM • Endpoint Verification • Identity Platform	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Identifying system users b. Identifying processes acting on behalf of users c. Identifying devices accessing the system When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Cloud Identity / Google Workspace • Identity and Access Management (IAM) - Service Accounts • Endpoint Verification • Identity Platform A description of relevant features and implementation guidance is included below. Cloud Identity / Google Workspace Use these services to manage user identities. Cloud Identity provides core identity services, while Google Workspace includes those plus collaboration tools like Gmail and Calendar. Both allow administrators to create, manage, and delete user accounts, establishing the primary identifier (usually email address) for human users accessing Google Cloud resources or other integrated applications. • Create User Accounts (Admin Console) 1. Sign in to the Google Admin console (admin.google.com) with an administrator account. 2. Navigate to Directory > Users. 3. Click Add new user.									

4. Fill in the user's details (First name, Last name, Primary email). Assign the correct domain if multiple exist.
5. Configure password options (auto-generate or create) and whether the user must change it at next sign-in.
6. (Optional) Assign the user to an Organizational Unit.
7. Click **Add New User**.
8. Provide the new user with their account details and initial password.

- **Manage User Lifecycle**

1. Regularly review users in **Directory > Users** to suspend or delete accounts for individuals who no longer require access.
2. Use Organizational Units to apply different policies or settings to groups of users.

Identity and Access Management (IAM) - Service Accounts Use IAM Service Accounts to create and manage identities for non-human entities like applications, scripts, or Compute Engine instances. These accounts allow processes to authenticate and be authorized to use Google Cloud APIs and resources.

- **Create a Service Account**

1. In the Google Cloud Console, navigate to **IAM & Admin > Service Accounts**.
2. Click **+ Create Service Account**.
3. Enter a **Service account name**, which helps generate the ID. Add a description.
4. Click **Create and Continue**.
5. (Optional) Grant IAM roles to the service account in the next step. Click **Continue**.
6. (Optional) Grant users access to manage or use this service account. Click **Done**.

- **Identify Service Account Usage**

1. Associate service accounts with resources like Compute Engine VMs during creation or by editing the instance settings (**Identity and API access** section). This identifies the process running on the VM.
2. Use Cloud Audit Logs to trace actions performed by service accounts. Filter logs in Logs Explorer for the service account's email address as the principal.

Endpoint Verification Use Endpoint Verification to collect information about desktop and laptop devices (OS, encryption status, etc.) accessing organizational data. This information helps identify devices and can be used with Access Context Manager to enforce device-based access controls.

- **Enable Endpoint Verification Monitoring**

1. Sign in to the Google Admin console (admin.google.com).
 2. Navigate to **Devices > Mobile & endpoints > Settings > Universal settings**.
 3. Click **Data Access > Endpoint Verification**.
 4. Select the target Organizational Unit.
 5. Check the box for **Monitor which devices access organization data**.
 6. Click **Save**.
- **Deploy Endpoint Verification Extension**
 1. In the Admin console, navigate to **Devices > Chrome > Apps & extensions > Users & browsers**.
 2. Select the target Organizational Unit.
 3. Click the **+** button and select **Add Chrome app or extension by ID**.
 4. Enter the extension ID `callobklhcbilhphincmohgkigmfcg`
 5. Choose **From the Chrome Web Store**. Click **Save**.
 6. Set the **Installation policy** to **Force install**.
 7. Enable **Allow access to keys** and **Allow enterprise challenge** in the extension settings panel.
 8. Click **Save**.
 9. Users may need to install a native helper application depending on the OS.
 - **View Device Inventory**
 1. In the Admin console, navigate to **Devices > Overview**.
 2. Click **Endpoints** to see the list of devices reporting information via Endpoint Verification.

Identity Platform Use Identity Platform to add identity management to your custom applications. It supports various authentication methods (email/password, social, SAML, OIDC, phone) for identifying end-users of your applications, distinct from the organizational users managed in Cloud Identity/Workspace.

- **Enable Identity Platform**
 1. In the Google Cloud Console, navigate to **Security > Identity Platform**.
 2. Click **Enable Identity Platform**.
- **Configure Providers**
 1. Navigate to **Identity Platform > Providers**.
 2. Click **Add a Provider**.
 3. Select and configure the desired sign-in methods (e.g., Email/Password, Google, SAML). Follow the on-screen instructions for each provider.
- **Manage Application Users**
 1. Users who authenticate via configured providers are managed within Identity Platform. Navigate to **Identity Platform > Users** to view and manage these user accounts specific to your application.

Additional Considerations

- Establish clear naming conventions for users and service accounts to maintain uniqueness and clarity.
- Implement processes for timely removal or suspension of identifiers when access is no longer required.
- Use Cloud Audit Logs to monitor authentication events and actions taken by identified users and processes.

Supplemental Guidance

- [Add an account for a new user \(Cloud Identity\)](#)
- [Create Cloud Identity user accounts](#)
- [Manage a user's security settings \(Workspace\)](#)
- [Manage user accounts \(API\)](#)
- [IAM Service accounts overview](#)
- [Types of service accounts](#)
- [Endpoint Verification overview](#)
- [Set up Endpoint Verification on your devices \(Quickstart\)](#)
- [Turn endpoint verification on or off \(Admin Console\)](#)
- [Identity Platform Documentation Overview](#)
- [Identity Platform Product Page](#)

Control Domain	Identification and Authentication								
Control #	IA.L1-3.5.2								
Control Description	Authenticate (or verify) the identities of users, processes, or devices, as a prerequisite to allowing access to organizational systems								
Key Services	• Cloud Identity / Google Workspace • IAM • IAP • Access Context Manager • Identity Platform	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Authenticating or verifying the identity of each user as a prerequisite to system access b. Authenticating or verifying the identity of each process acting on behalf of a user is authenticated or verified as a prerequisite to system access									

- c. Authenticating the identity of each device accessing or connecting to the system is authenticated or verified as a prerequisite to system access

When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control:

- Cloud Identity / Google Workspace (with 2-Step Verification/MFA)
- Identity and Access Management (IAM) - Service Accounts & Workload Identity Federation
- Identity-Aware Proxy (IAP)
- Access Context Manager
- Identity Platform

A description of relevant features and implementation guidance is included below.

Cloud Identity / Google Workspace (with 2-Step Verification/MFA): Use these services to manage user identities and enforce strong authentication. Configure password policies and mandate the use of 2-Step Verification (Multi-Factor Authentication - MFA) to verify user identities beyond just a password.

- **Configure Password Policies (Workspace/Cloud Identity)**
 1. In the Google Admin console (admin.google.com), navigate to **Security > Authentication > Password management**.
 2. Set password strength requirements (length, complexity).
 3. Configure password reuse policies.
 4. Set password expiration frequency.
 5. Click **Save**.
- **Enforce 2-Step Verification (MFA)**
 1. In the Admin console, navigate to **Security > Authentication > 2-step verification**.
 2. Select the top Organizational Unit or a specific one.
 3. Check **Allow users to turn on 2-Step Verification**.
 4. Set **Enforcement** to **On** or **Enforced from date**. *Recommendation Start with "Turned on" allowing users to enroll, communicate the requirement, track enrollment, and then move to "Enforced".*
 5. Select allowed **Methods**. *Recommendation Prioritize Security Keys and Google Prompts over SMS/Voice calls.*
 6. Configure the **New user enrollment period**.
 7. Click **Save** (or **Override** if configuring an OU).

Identity and Access Management (IAM) - Service Accounts & Workload Identity Federation: Use IAM to manage how service accounts (representing processes/applications) authenticate. Avoid exporting keys; instead, attach service accounts

directly to resources like Compute Engine VMs or use Workload Identity Federation for external workloads.

- **Attach Service Account to a VM (Preferred Authentication)**
 1. During VM creation (**Compute Engine > VM instances > Create Instance**), under **Identity and API access**, select the desired service account. The VM uses the Google Cloud metadata service to get tokens automatically.
 2. For existing VMs Stop the VM, click **Edit**, change the **Service account** under **Identity and API access**, and restart the VM.
- **Manage Service Account Keys (Use only if necessary)**
 1. Minimize key usage. If needed Navigate to **IAM & Admin > Service Accounts**, select the account, go to the **Keys** tab, click **Add Key > Create new key**. Secure the downloaded file.
 2. Implement regular key rotation.
 3. Consider disabling key creation via Organization Policies (**iam.disableServiceAccountKeyCreation**).
- **Use Workload Identity Federation (For external workloads)**
 1. Configure a Workload Identity Pool (**IAM & Admin > Workload Identity Federation**).
 2. Add an identity provider (e.g., AWS, Azure AD, OIDC, SAML 2.0). Configure provider details and attribute mapping.
 3. Grant external identities IAM roles, often the Service Account Token Creator role (**roles/iam.serviceAccountTokenCreator**) to allow them to impersonate a specific Google Cloud service account. External workloads exchange their native credentials for short-lived Google Cloud credentials.

Identity-Aware Proxy (IAP): Use IAP to establish a central authorization layer for applications accessed via HTTPS (App Engine, Compute Engine backends, GKE) and for TCP connections (SSH/RDP to VMs). IAP verifies user identity and context *before* allowing access to the application or resource.

- **Enable IAP for HTTPS Resources**
 1. Navigate to **Security > Identity-Aware Proxy**.
 2. If prompted, configure the **OAuth consent screen**. Provide an application name and support email.
 3. Find the backend service or App Engine app under **HTTPS Resources**.
 4. Toggle the switch in the **IAP** column to **On**.
 5. In the right-side panel (**Info Panel**), click **Add Principal**.
 6. Enter the email addresses of users or groups who need access.
 7. Assign the **IAP-Secured Web App User** role (**roles/iap.httpsResourceAccessor**).

8. Click **Save**.
- **Enable IAP for TCP Resources (e.g., SSH/RDP)**
 1. Navigate to **Security > Identity-Aware Proxy**.
 2. Go to the **SSH and TCP Resources** tab.
 3. Find the VM instances you want to protect. (Requires appropriate firewall rules allowing traffic from IAP's IP range 35.235.240.0/20).
 4. Select the instances.
 5. In the right-side panel (**Info Panel**), click **Add Principal**.
 6. Enter users/groups.
 7. Assign the **IAP-Secured Tunnel User** role (**roles/iap.tunnelResourceAccessor**).
 8. Click **Save**. Users will then use `gcloud compute ssh --tunnel-through-iap` or similar methods.

Access Context Manager: Use Access Context Manager to define fine-grained access control policies based on attributes like user identity, device security status (requires Endpoint Verification), IP address, and location. This ensures users are authenticated *within the required context*.

- **Create an Access Level**
 1. Navigate to **Security > Access Context Manager**.
 2. If prompted, select your Organization.
 3. Click **New Access Level**.
 4. Provide a **Name** for the level (e.g., `authenticated_corp_device`).
 5. Choose **Basic** mode.
 6. Define **Conditions**. Click **Add Attribute**.
 7. Select attributes like
 - **IP Subnetworks** Specify allowed source IP ranges.
 - **Device Policy** Require attributes like OS version, screen lock, disk encryption (relies on Endpoint Verification data).
 - **Require Verified Chrome OS** For managed ChromeOS devices.
 8. Set conditions to **AND** or **OR**.
 9. Click **Save**.
- **Apply Access Level (Example VPC Service Controls)**
 1. When creating or editing a VPC Service Control perimeter (**Security > VPC Service Controls**), you can add the created Access Level under **Ingress Policies** or **Egress Policies** to allow authenticated access based on the defined context.

Identity Platform: Use Identity Platform to manage authentication for users of your custom-built applications, supporting various methods.

- **Configure Authentication Methods**
 1. In the Google Cloud Console, navigate to **Security > Identity Platform > Providers**.
 2. Click **Add a Provider**.
 3. Choose and configure desired methods
 - **Email/Password** Basic username/password authentication.
 - **Phone** SMS-based verification.
 - **Google, Facebook, Twitter, GitHub, etc.** Federated sign-in.
 - **SAML / OpenID Connect** Integrate with enterprise identity providers.
 4. Follow configuration steps for each selected provider (e.g., obtaining client IDs/secrets, configuring callback URLs).
- **Enable Multi-Factor Authentication (MFA)**
 1. Navigate to **Identity Platform > Providers > Multi-factor authentication**.
 2. Configure MFA options (e.g., SMS-based second factor).
 3. Implement MFA requirements in your application flow using the Identity Platform SDKs.

Additional Considerations

- Prioritize MFA (2-Step Verification) for all user accounts, especially administrators.
- Avoid using static, long-lived credentials like service account keys whenever possible. Prefer attached service accounts or Workload Identity Federation.
- Configure appropriate session timeouts for web consoles and applications.
- Regularly audit authentication logs (Cloud Audit Logs, IAP logs) for suspicious activity.

Supplemental Guidance

- [Deploy 2-Step Verification \(Cloud Identity\)](#)
- [Protect your business with 2-Step Verification \(Workspace\)](#)
- [Multi-factor authentication requirement for Google Cloud](#)
- [Authentication methods at Google](#)
- [Workload Identity Federation](#)
- [Best practices for working with service accounts](#)
- [IAP Overview](#)
- [Setting up IAP for Compute Engine](#)
- [IAP How-to Guides](#)
- [Access Context Manager Overview](#)
- [Creating a basic access level](#)

- [Identity Platform Authentication Concepts](#)
- [Identity Platform Product Page](#)

Control Domain	Identification and Authentication								
Control #	IA.L2-3.5.3								
Control Description	Use multifactor authentication for local and network access to privileged accounts and for network access to non-privileged accounts								
Key Services	• Cloud Identity / Google Workspace • IAM • IAP	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Identifying privileged accounts b. Implementing multifactor authentication for network access to privileged accounts c. Implementing multifactor authentication for network access to non-privileged accounts Google is responsible for: d. Implementing multifactor authentication for local access to privileged accounts When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Cloud Identity / Google Workspace (with 2-Step Verification) • Identity and Access Management (IAM) • Identity-Aware Proxy (IAP) A description of relevant features and implementation guidance is included below. Cloud Identity / Google Workspace (with 2-Step Verification): Use the Google Admin console to mandate MFA for all users signing into their Google account, which is used to access Google Cloud Console, APIs, and IAP-protected resources. This addresses the requirement for network access MFA for both privileged and non-privileged accounts. • Enforce 2-Step Verification (MFA) for All Users 1. In the Google Admin console (admin.google.com), navigate to Security > Authentication > 2-step verification.									

2. Select the top Organizational Unit to apply the policy to everyone.
 3. Check **Allow users to turn on 2-Step Verification**.
 4. Set **Enforcement** to **On** or **Enforced from date**. *Google requires MFA for all admin accounts and recommends it for all users.*
 5. Under **Methods**, select the allowed verification methods. *Recommendation Select "Any except verification codes via text, phone call" or "Only security key" for higher security, especially for privileged users.*
 6. Configure the **New user enrollment period** (grace period for new accounts). Setting to **None** requires enrollment before first sign-in after policy application.
 7. Click **Save**.
- **(Optional) Stricter MFA for Privileged Users**
 1. Create a Group or Organizational Unit specifically for users with privileged IAM roles.
 2. Navigate to **Security > Authentication > 2-step verification**.
 3. Select the privileged Group or OU.
 4. Configure stricter **Enforcement** settings (e.g., shorter or no enrollment period) or restrict **Methods** (e.g., "Only security key").
 5. Click **Save** or **Override**.

Identity and Access Management (IAM): Use IAM to define which accounts are considered privileged by assigning specific roles. While IAM doesn't enforce MFA itself, it's critical for identifying the users and groups that require privileged access and should potentially have stricter MFA policies applied in the Admin Console.

- **Identify Privileged Roles**
 - Review predefined roles like roles/owner, roles/editor (avoid if possible), roles/organizationAdmin, roles/iam.securityAdmin, etc.
 - Analyze custom roles to identify those granting sensitive permissions (e.g., modifying IAM policies, deleting resources, accessing sensitive data).
- **Assign Privileged Roles Appropriately**
 - Grant privileged roles only to users who require them, following the principle of least privilege.
 - Use Groups in the Admin Console to manage membership for privileged roles, simplifying policy application (including MFA policies).

Identity-Aware Proxy (IAP): Use IAP to secure network access (HTTPS, SSH, RDP) to applications and Compute Engine instances. IAP enforces authentication based on the user's Google identity *before* allowing traffic to the resource. This means that if MFA is enforced for the user's account in Cloud Identity/Workspace, IAP effectively requires MFA for network access to the protected resource. This addresses the "network access" component for both

privileged and non-privileged accounts accessing specific systems, and the "local access" component for privileged accounts needing SSH/RDP access to VMs.

- **Secure Network Access to VMs (SSH/RDP) with IAP**
 1. Ensure appropriate firewall rules allow TCP traffic from IAP's IP range (35.235.240.0/20) to your target VMs on ports 22 (SSH) or 3389 (RDP). Block direct access from other sources.
 2. Navigate to **Security > Identity-Aware Proxy**.
 3. Go to the **SSH and TCP Resources** tab.
 4. Select the VM instances.
 5. In the right-side panel (**Info Panel**), click **Add Principal**.
 6. Enter the users/groups needing access.
 7. Assign the **IAP-Secured Tunnel User** role (**roles/iap.tunnelResourceAccessor**).
 8. Click **Save**.
 9. Users will connect using `gcloud compute ssh --tunnel-through-iap` (or RDP equivalent), authenticating with their Google identity (including MFA if enforced).
- **Secure Network Access to Web Applications (HTTPS)**
 1. Follow the steps outlined in IA.3.5.2 for enabling IAP for HTTPS resources. Access will require authentication with the user's Google identity, including MFA if enforced.

Additional Considerations

- **Phishing-Resistant MFA:** Strongly encourage or require the use of phishing-resistant MFA methods like Security Keys (FIDO hardware tokens or phone-based keys) for all users, especially administrators.
- **Separate Admin Accounts:** Administrators should use dedicated accounts with privileged roles, separate from their standard user accounts, and these admin accounts must have MFA enforced.
- **Recovery:** Establish secure procedures for users who lose access to their MFA methods (e.g., backup codes, administrator recovery).
- **Auditing:** Regularly audit MFA enrollment status (**Admin Console > Reporting > User Reports > Security**) and monitor sign-in logs (Cloud Audit Logs) for MFA-related events or failures.

Supplemental Guidance

- [Protect your business with 2-Step Verification](#)
- [Deploy 2-Step Verification](#)
- [Use IAM securely | IAM Documentation | Google Cloud](#)
- [Best practices for using service accounts | IAM Documentation | Google Cloud](#)

- [Privileged Access Manager overview | IAM Documentation | Google Cloud](#)
- [Identity-Aware Proxy overview | Google Cloud](#)
- [Overview of TCP forwarding | Identity-Aware Proxy | Google Cloud](#)

Control Domain	Identification and Authentication								
Control #	IA.L2-3.5.4								
Control Description	Employ replay-resistant authentication mechanisms for network access to privileged and non-privileged accounts								
Key Services	• Cloud Identity / Google Workspace • IAM	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Implementing replay-resistant authentication mechanisms for network account access to privileged and non-privileged accounts This control is primarily met through the inherent behavior of Google's identity systems and customer processes, rather than specific console configurations to set a reuse period. However, the following feature(s) in Google Cloud Console may be used to support this control: • Cloud Identity / Google Workspace (User Identifiers) • Identity and Access Management (IAM) (Service Accounts) A description of relevant features and implementation guidance is included below. Cloud Identity / Google Workspace (User Identifiers): Google Workspace and Cloud Identity enforce a waiting period before a deleted user's primary email address can be reused as a primary address for a new user or as an alias. • Platform Behavior: When a user account is deleted, Google prevents that primary email address identifier from being immediately reassigned. This waiting period allows for potential account recovery and prevents impersonation or misdirected communication. The typical period is around 24 hours, but can extend up to 20 days if data recovery is possible, or potentially longer for accounts subject to legal holds or recently deleted admin accounts. • Customer Process									

1. Do not attempt to create a new user or alias with the exact primary email address of a recently deleted user immediately. Plan for a waiting period of at least 24 hours, potentially longer.
2. Consider suspending user accounts instead of deleting them if access needs to be temporarily revoked but the identifier or data retained.
3. If an identifier (email address) needs to be available quickly for another purpose (like a group), consider renaming the user account *before* deletion. Note that renaming retains the old address as an alias by default, which would also need to be removed.

Identity and Access Management (IAM) - Service Accounts: IAM ensures that service account email addresses, which serve as their unique identifiers within a project (service-account-name@project-id.iam.gserviceaccount.com), cannot be reused within that same project after the service account is deleted.

- **Platform Behavior:** Once an IAM service account is deleted from a project, its unique email identifier cannot be assigned to a new service account created within that same project. While you can undelete a service account within 30 days, if you create a new one with the same name during that period, the new one is treated as a completely separate identity and prevents undeletion of the original.
- **Customer Process**
 1. Use distinct and meaningful naming conventions for service accounts to minimize the chance of needing to reuse an identifier.
 2. If a service account is deleted accidentally, attempt to undelete it within the 30-day window before creating a new one with the same identifier if the original permissions are needed.
 3. Disable service accounts before deleting them to verify they are no longer in use and reduce operational impact.

Additional Considerations

- **Suspension vs. Deletion:** Suspending a user account blocks access but preserves the account, data, and identifier. Deletion initiates a removal process with a limited recovery window (20 days for users, 30 days for service accounts) and prevents identifier reuse as described.
- **Aliases:** User email aliases might have different, potentially shorter, reuse restrictions compared to primary email addresses after being removed from an account, but relying on this is not recommended for meeting the control's intent regarding primary identifiers.

- **Workforce Identity Federation:** For external identities managed via Workforce Identity Federation, user identifiers enter a 30-day soft-deletion state before becoming potentially reusable.

Supplemental Guidance

- [Delete or remove a user from your organization](#)
- [Delete and undelete service accounts](#)
- [Delete Workforce Identity Federation users and their data](#)
- [Create, shut down, and restore projects](#)

Control Domain	Identification and Authentication		
Control #	IA.L2-3.5.5		
Control Description	Prevent reuse of identifiers for a defined period		
Key Services	• Cloud Identity / Google Workspace • IAM	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining a period within which identifiers cannot be reused Google is responsible for: - Preventing the reuse of identifiers within the defined period Note: Google Cloud Console and Google Workspace Admin Console do not offer a built-in, configurable feature to automatically disable user or service accounts based <i>solely</i> on a defined period of inactivity. Meeting this control typically requires manual processes or custom automation built by the customer. However, the following key service(s) in Google Cloud Console may be used to support this control: • Cloud Identity / Google Workspace (User Accounts) • Identity and Access Management (IAM) - Service Accounts A description of relevant features and implementation guidance is included below.			

Cloud Identity / Google Workspace (User Accounts): Use Admin Console reports for manual review or leverage APIs for custom automation to disable inactive user accounts.

- **Manual Review and Disablement**

1. Define the inactivity period (e.g., 90 days).
2. Periodically review user activity in the Google Admin console
 - Navigate to **Directory > Users**. The **Last sign-in** column provides an approximate time.
 - Navigate to **Reporting > Reports > User Reports > Accounts** for more detailed login activity or **Security** for security-related events. Use filters to identify users with no recent activity.
3. For users exceeding the defined inactivity period, suspend their accounts
 - Go to **Directory > Users**.
 - Find the inactive user.
 - Click **More options > Suspend user**.
 - Confirm by clicking **Suspend**.

- **Custom Automation (Conceptual Steps)**

1. Develop a script or application (e.g., using Cloud Functions scheduled with Cloud Scheduler).
2. Use the Reports API (users.usage.get) to fetch the accountslast_login_time for users. (See Result 1.3)
3. Compare the last login time against your defined inactivity threshold.
4. For users exceeding the threshold, use the Directory API (users.update) to set the suspended field to true. (See Results 2.1, 2.4, 2.5)

Identity and Access Management (IAM) - Service Accounts: Use IAM console information, gcloud commands, or Security Command Center Premium to identify inactive service accounts, then disable them manually or via automation.

- **Manual Review and Disablement**

1. Define the inactivity period (e.g., 90 days).
2. Periodically review service account activity
 - In the Google Cloud Console, navigate to **IAM & Admin > Service Accounts**. Check the **Last authentication** information if available (may not be present for all accounts or recent activity).
 - Use `gcloud iam service-accounts list --format='value(email, activity.lastAuthTime)'` (Note activity.lastAuthTime is a newer field replacing older methods, check documentation for current command syntax and availability).
 - Use Activity Analyzer (**IAM & Admin > Activity Analyzer**) to check last usage dates.

3. For service accounts exceeding the defined inactivity period, disable them
 - Navigate to **IAM & Admin > Service Accounts**.
 - Select the inactive service account.
 - Click **Disable Service Account**. (See Result 3.1)
- **Using Security Command Center (SCC) Premium Detector**
 1. Ensure SCC Premium is active and configured.
 2. Monitor for findings from the **Dormant service account** detector (or similarly named inactivity detector). This typically flags accounts inactive for 90+ days by default.
 3. Review findings in the SCC console (**Security > Security Command Center > Findings**).
 4. Based on findings, manually disable the identified service accounts via the IAM console or gcloud.
- **Custom Automation (Conceptual Steps)**
 1. Develop automation (e.g., Cloud Function).
 2. Trigger the automation based on SCC findings pushed to Pub/Sub or on a schedule (Cloud Scheduler).
 3. If triggered by schedule, use the IAM API or gcloud to list service accounts and check last authentication times.
 4. For service accounts exceeding the threshold, use the IAM API (projects.serviceAccounts.disable) or gcloud (gcloud iam service-accounts disable) to disable the account. (See Result 3.1)

Additional Considerations

- **Defining Inactivity:** Carefully define the inactivity period based on organizational risk tolerance and operational needs. A period like 90 days is common but should be justified.
- **Suspension/Disabling vs. Deletion:** Disabling (suspending for users) is reversible and generally preferred over immediate deletion, allowing for investigation or potential reactivation if needed. Implement separate processes for eventual deletion of long-disabled accounts.
- **Impact:** Before disabling, consider the potential impact on applications or resources that might rely on the identifier, especially for service accounts. Disabling is safer than deleting but can still cause outages if the account is unexpectedly needed.
- **Automation Complexity:** Building reliable automation requires careful development, testing, error handling, and monitoring.

Supplemental Guidance

- [View user's last sign-in](#)
- [User log events](#)

- [Reports API Users Usage Report](#)
- [Suspend a user temporarily](#)
- [Directory API Manage user accounts](#)
- [Disable and enable service accounts](#)
- [Service accounts overview](#)
- [Overview of Event Threat Detection](#)
- [Policy Intelligence](#)

Control Domain	Identification and Authentication								
Control #	IA.L2-3.5.6								
Control Description	Disable identifiers after a defined period of inactivity								
Key Services	• Cloud Identity / Google Workspace • Identity Platform • Password Sync • IAM	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Defining a period of inactivity after which an identifier is disabled b. Disabling identifiers after the defined period of inactivity Note: Google Workspace/Cloud Identity does not provide a direct setting in the Admin Console for administrators to define the exact number of failed attempts or lockout duration for standard Google user accounts. Google's internal security mechanisms handle this automatically. However, the following key service(s) in Google Cloud Console may be used to support this control: • Cloud Identity / Google Workspace (User Accounts) • Identity Platform • Password Sync • Identity and Access Management (IAM) - Service Accounts A description of relevant features and implementation guidance is included below. Cloud Identity / Google Workspace (User Accounts): Rely on Google's built-in protections and monitor activity.									

- **Understand Built-in Protections:** Google automatically detects patterns indicative of brute-force attacks (like multiple consecutive failed logins) and may respond by presenting challenges (e.g., CAPTCHA) or temporarily locking the account to protect it. The exact thresholds are not customer-configurable.
- **Monitor Activity and Alerts**
 1. Configure Admin alerts for security events Navigate to **Security > Alert center > Settings (gear icon)** and configure email notifications for relevant alerts like "Suspicious login," "User suspended due to suspicious activity," or "Leaked password detected."
 2. Regularly review the **Alert center (Security > Alert center)** for triggered alerts.
 3. Review login activity logs Navigate to **Reporting > Audit and investigation > User log events**. Filter for login events (Event name is Login) and look for patterns of failures (Login Failure Type). (See Result 2.2)

Identity Platform (Custom Applications): Use Blocking Functions to implement custom logout logic for users authenticating to applications managed by Identity Platform.

- **Configure Blocking Functions**
 1. Develop a Cloud Function triggered beforeSignIn.
 2. Within the function, implement logic to
 - Track the number of recent failed login attempts for the user (this typically requires using an external store like Firestore or Cloud Memorystore, as the function itself is stateless).
 - Check if the number of failures within a defined time window exceeds the organizational threshold.
 - If the threshold is exceeded, throw an `HttpsError` from the function to block the sign-in attempt. (See Results 3.2, 3.4)
 - If the login is successful, reset the failure counter for that user.
 3. Deploy the function and register it as a blocking function trigger in the Identity Platform settings (**Security > Identity Platform > Settings > Triggers**).

Password Sync (for Synced AD Accounts): If using Google Cloud Directory Sync (GCDS) with the Password Sync feature to synchronize Active Directory passwords to Google, configure the account lockout policy.

- **Configure Account Lockout Policy**
 - Within the Password Sync configuration settings on the server where it's installed, adjust the parameters related to account lockout thresholds (number of attempts) and lockout duration according to your organizational

policy. Refer to the specific Password Sync documentation for configuration details as this is set within the tool, not the Cloud Console.

Identity and Access Management (IAM) - Service Accounts: Service account authentication failures do not typically result in account lockout but generate API errors.

- **Handle API Errors:** Applications using service accounts should implement robust error handling to detect authentication failures (e.g., HTTP 401/403 errors). Log these errors and potentially implement alerting or circuit-breaker patterns if failures are persistent, indicating a possible configuration issue or compromised key (if keys are used).

Additional Considerations

- **MFA:** Enforcing Multi-Factor Authentication ([IA.L2-3.5.3](#)) significantly reduces the risk associated with compromised passwords and the reliance on simple password failure lockouts.
- **Monitoring:** Actively monitoring Admin alerts and audit logs is crucial for identifying and responding to potentially malicious login attempts, even with automatic protections in place.
- **User Experience:** Balance strict lockout policies with user experience. Overly aggressive lockouts can hinder legitimate users. Ensure clear procedures exist for users to regain access after a lockout.

Supplemental Guidance

- [About admin alerts for suspicious login activity](#)
- [Alert Center overview](#)
- [User log events](#)
- [Login Audit Activity Events \(API Details\)](#)
- [Customizing authentication flow using blocking functions](#)
- [Troubleshoot Password Sync](#)

Control Domain	Identification and Authentication		
Control #	IA.L2-3.5.7		
Control Description	Enforce a minimum password complexity and change of characters when new passwords are created		
Key Services	• Cloud Identity / Google Workspace • Identity Platform	Control Responsibility	☐ Google ☒ Shared

			☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining password complexity requirements - Defining password change of character requirements Google is responsible for: - Enforcing minimum password complexity requirements when new passwords are created - Enforcing minimum password change of character requirements as defined when new passwords are created Note: Password complexity and history controls do not apply to IAM Service Accounts, as they use cryptographic keys or other authentication methods instead of passwords. However, the following key service(s) in Google Cloud Console may be used to support this control: - Cloud Identity / Google Workspace (User Accounts) - Identity Platform A description of relevant features and implementation guidance is included below. Cloud Identity / Google Workspace (User Accounts): Use the Google Admin Console to configure password complexity and reuse policies for users managed directly by Google Workspace or Cloud Identity. Configure Password Requirements - Sign in to the Google Admin console (admin.google.com). - Navigate to Security > Authentication > Password management. - Select the Organizational Unit to apply the policy to (select the top level for all users). - In the Strength section - Check Enforce strong password. This requires users to use a mix of character types (upper, lower, number, symbol) and avoids common dictionary words or contextual information. - In the Length section - Set the Minimum length (between 8 and 100 characters). (Optional) Set a Maximum length. - In the Reuse section			

- Uncheck **Allow password reuse**. This prevents users from reusing a certain number of their previous passwords based on Google's internal history tracking (the exact number of generations is not customer-configurable but reuse is prevented).
- 7. (Optional) Configure **Password Expiration**.
- 8. (Optional) Check **Enforce password policy at next sign-in** to force users with non-compliant passwords to change them immediately.
- 9. Click **Save** (or **Override** if configuring an OU).

Identity Platform (Custom Applications using Email/Password Provider): Use the Identity Platform API or Admin SDK to enforce password complexity. Password history/reuse prevention requires custom implementation.

- **Configure Password Complexity (via API/SDK)**
 1. Use the Identity Platform Admin SDK (e.g., Node.js) or the REST API (`projects.updateConfig`) to set a `passwordPolicyConfig`. (See Result 3.1, 3.5)
 2. Within the policy configuration, specify constraints such as
 - `requireUppercase` (true/false)
 - `requireLowercase` (true/false)
 - `requireNumeric` (true/false)
 - `requireNonAlphanumeric` (true/false)
 - `minLength` (integer 6-30)
 - `maxLength` (integer up to 4096)
 3. Set the `enforcementState` to `ENFORCE`.
- **Implement Password History/Reuse Prevention (Custom Logic)**
 1. This feature is not directly configurable via the Identity Platform password policy API.
 2. Custom logic is needed, typically involving
 - Securely storing a hash of previous passwords (e.g., in Firestore) associated with the user during password change events. Only store hashes, never plaintext passwords.
 - Using an Identity Platform Blocking Function triggered before `Create` (for initial password) or during password reset flows.
 - Within the function, retrieve the user's password history hashes.
 - Hash the proposed new password using the same algorithm and salt parameters.
 - Compare the new hash against the stored history hashes.
 - If a match is found within the defined number of prohibited generations, throw an `HttpError` from the blocking function to reject the password change.

- If the password change is successful, add the new hash to the user's history (and potentially prune old history).

Additional Considerations

- **Define Policies:** Establish clear, documented password complexity and history requirements based on organizational risk assessment.
- **User Education:** Inform users about the password policy requirements and provide guidance on creating strong, unique passwords. Encourage the use of password managers.
- **Alternatives:** Consider promoting phishing-resistant authentication methods like Security Keys (using WebAuthn/FIDO2) or federated sign-in (SAML/OIDC) with external identity providers, which can reduce reliance on user-managed passwords.
- **Identity Platform Integration:** If using Identity Platform, consider configuring Cloud Identity/Google Workspace as a SAML or OIDC identity provider instead of relying solely on the native email/password provider. This allows users to authenticate with their managed Google account, inheriting the password policies set in the Admin Console.

Supplemental Guidance

- [Enforce and monitor password requirements for users](#)
- [Manage a user's security settings \(includes password reset\)](#)
- [Enable, disable, and use password policies \(API/SDK\)](#)
- [Method getPasswordPolicy \(REST API details\)](#)
- [Customizing authentication flow using blocking functions \(for history logic\)](#)
- [Google Password Manager](#)

Control Domain	Identification and Authentication		
Control #	IA.L2-3.5.8		
Control Description	Prohibit password reuse for a specified number of generations		
Key Services	• Cloud Identity / Google Workspace • Identity Platform • IAP	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for:			

- a. Specifying the number of generations during which a password cannot be reused is specified

Google is responsible for:

- b. Prohibiting reuse of passwords during the specified number of generations

This control is primarily met by the standard design of Google's authentication and session management systems. The main configuration action for customers is setting appropriate session lengths. However, the following key service(s) in Google Cloud Console may be used to support this control:

- Cloud Identity / Google Workspace (Session Management)
- Identity Platform (Token-based Session Management)
- Identity-Aware Proxy (IAP) (Session Management)

A description of relevant features and implementation guidance is included below.

Cloud Identity / Google Workspace (Session Management): Google's authentication system establishes a session upon successful login (which may involve a password and MFA initially). Subsequent access to Google Workspace applications and the Google Cloud Console within that session relies on secure session cookies/tokens, not repeated password entry.

- **Configure Session Duration:** Control how long these sessions last before re-authentication is required
 1. Sign in to the Google Admin console (admin.google.com).
 2. Navigate to **Security > Access and data control > Google Session control**.
 3. Select the Organizational Unit to apply the policy to.
 4. Under **Web session duration**, choose the desired length (e.g., 1 hour, 8 hours, 24 hours, 14 days). Setting a duration ensures users must re-authenticate after that period, though the system doesn't ask for the password for every action *within* the session.
 5. (Optional) Configure Google Cloud specific session controls under **Security > Access and data control > Google Cloud session control** to set re-authentication frequency specifically for Cloud Console and gcloud CLI access.
 6. Click **Save** (or **Override**).

Identity Platform (Token-based Session Management): Applications using Identity Platform authenticate users initially (potentially with a password). After that, the SDKs

manage the user's session using short-lived ID tokens and longer-lived refresh tokens. API calls or access checks rely on validating the ID token, not the user's password.

- **Platform Behavior:** The Identity Platform SDKs handle the process of obtaining, caching, and refreshing ID tokens automatically. Applications typically use listeners to detect authentication state changes and use the current ID token for authenticated requests. This avoids needing the user's password after the initial sign-in.
- **Customer Process:** Implement standard authentication state management as described in Identity Platform documentation. Ensure application logic relies on verifying ID tokens passed from the client SDK for authorization decisions, rather than re-collecting passwords.

Identity-Aware Proxy (IAP) (Session Management): Protect access to applications by verifying the user's Google-authenticated session on each request. It relies on the session established with Google (as configured by Google Session Control) and does not require the user to re-enter their password for each access attempt through the proxy.

- **Platform Behavior:** IAP intercepts requests, checks for a valid Google session cookie, verifies the user's identity and authorization (IAM permissions), and then passes the request to the backend application, often adding a signed header containing the authenticated user's identity. Password re-entry is not part of this per-request check.

Additional Considerations

- **Session Security:** Google employs security measures like HttpOnly and Secure flags for its session cookies. For custom applications, follow secure session management best practices.
- **Re-authentication for Sensitive Actions:** While the control prohibits *routine* password re-entry, Google services or well-designed applications might trigger step-up authentication (often involving MFA) for particularly sensitive actions (e.g., changing security settings, deleting critical resources). This is a security best practice and generally compatible with the control's intent.
- **Service Accounts:** This control is not applicable to service accounts as they do not use password-based authentication.

Supplemental Guidance

- [Set session length for Google services](#)
- [Set session length for Google Cloud services](#)
- [Google Cloud Blog on Session Length Defaults](#)
- [Identity Platform users in projects](#)
- [Managing sessions with external identities](#)

- [Session Management Best Practices](#)
- [Identity-Aware Proxy overview](#)

Control Domain	Identification and Authentication								
Control #	IA.L2-3.5.9								
Control Description	Allow temporary password use for system logons with an immediate change to a permanent password								
Key Services	Cloud Identity / Google Workspace	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☐</td><td>Shared</td></tr><tr><td>☒</td><td>Customer</td></tr></table>	☐	Google	☐	Shared	☒	Customer
☐	Google								
☐	Shared								
☒	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Requiring an immediate change to a permanent password when a temporary password is used for system logon When configured correctly, the following feature(s) in Google Cloud Console (or associated admin consoles) may be used to support this control: - Google Workspace Admin Console (for Google Workspace and Cloud Identity users) A description of relevant features and implementation guidance is included below. Google Workspace Admin Console (for Google Workspace and Cloud Identity users): Manage users, including their passwords and sign-in policies for accounts that access Google Cloud resources. Setting a Temporary Password and Forcing Change for a New User: - Navigate to the Google Admin console (admin.google.com). - Go to Directory > Users. - Click Add new user. - Fill in the user's details (First name, Last name, Primary email). - Under the Password section, select Create password. - Enter a temporary password for the user or allow Google to generate one. It is recommended to use an automatically generated strong password. - Crucially, ensure the option Ask for a password change at the next sign-in is turned ON (this is typically the default setting).									

8. Click **Add New User**.
 9. Securely communicate the temporary password to the user, along with instructions to log in immediately and change it.
- **Resetting a Password and Forcing Change for an Existing User:**
 1. Navigate to the Google Admin console (admin.google.com).
 2. Go to **Directory > Users**.
 3. Hover over the user whose password you want to reset and click **Reset password**.
 4. In the "Reset password" dialog, choose to **Create password** (either automatically generate or manually create a temporary one).
 5. Ensure the option **Ask for a password change at the next sign-in** is **turned ON**.
 6. Click **Reset**.
 7. Securely communicate the new temporary password to the user, along with instructions to log in immediately and change it.

Upon their next sign-in to their Google account (which could be for accessing Google Cloud Console, gcloud CLI, or other Google services), the user will be immediately prompted to create a new, permanent password before they can proceed.

Additional Considerations

- **External Identity Providers (IdP):** If you are using an external IdP (e.g., Azure AD, Okta) to federate identities into Google Cloud, the management of temporary passwords and the enforcement of password changes on initial logon are handled by that external IdP. You will need to configure these policies within your IdP's administration interface according to its specific capabilities. Google Cloud will honor the authentication assertion from the federated IdP.
- **Service Accounts:** This control primarily applies to human user accounts. Service accounts in Google Cloud use cryptographic keys rather than passwords for authentication and are not subject to interactive logon password change policies.
- **Password Policies:** Beyond temporary passwords, configure strong password policies in the Google Admin console (Security > Authentication > Password management) to enforce complexity, expiration, and reuse prevention for permanent passwords.
- **User Training:** Inform users about the process and the importance of choosing a strong, unique permanent password once they log in with their temporary password.
- **Monitoring:** While not a direct configuration for this control, monitor sign-in audit logs (available in Google Workspace Admin console under Reporting > Audit and

investigation > Admin log events or User log events) for unusual password reset activities if needed.

Supplemental Guidance

- [Add an account for a new user - Domain verified - Google Workspace Admin Help](#)
- [Reset a user's password - Google Workspace Admin Help](#)
- [Manage a user's security settings - Google Workspace Admin Help](#)
- [Best practices for federating Google Cloud with an external identity provider | Cloud Architecture Center](#)

Control Domain	Identification and Authentication								
Control #	IA.L2-3.5.10								
Control Description	Store and transmit only cryptographically-protected passwords								
Key Services	• Cloud Identity / Google Workspace • Identity Platform • Google Cloud Infrastructure	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for (related to hosted applications and systems): a. Cryptographically protecting passwords in storage b. Cryptographically protecting passwords in transit Google is responsible for (related to Google infrastructure): a. Cryptographically protecting passwords in storage b. Cryptographically protecting passwords in transit When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Cloud Identity / Google Workspace (Session Management) • Identity Platform • Google Cloud Infrastructure (Transport Encryption) A description of relevant features and implementation guidance is included below.									

Identity Platform: For applications using Identity Platform's built-in email/password authentication provider, Google handles password transmission and storage securely, similar to standard Google accounts.

- **Platform Behavior (Native Provider):** When users sign up or sign in using the email/password provider, passwords are transmitted over HTTPS and stored by Google using secure hashing techniques.
- **Custom Authentication / Federation**
 - If integrating with external identity providers (SAML/OIDC), the responsibility for secure password storage lies with the external provider. Ensure the chosen provider follows industry best practices.
 - If building a custom authentication system that integrates with Identity Platform (e.g., exchanging custom tokens for Identity Platform tokens), the customer's backend system is responsible for securely handling password verification. This *must* involve storing only secure hashes (e.g., bcrypt, scrypt, Argon2 with salts) of user passwords and performing comparisons against hashes derived from user input during login attempts. Plaintext password storage is insecure and should never be implemented.

Google Cloud Infrastructure (Transport Encryption): All communication with Google Cloud services, including authentication portals and APIs, occurs over encrypted channels (HTTPS/TLS), protecting data, including passwords entered by users, during transmission.

- **Platform Behavior:** Google enforces HTTPS for its login pages and APIs, ensuring passwords are encrypted between the user's browser/client and Google's servers.

Additional Considerations

- **Reliance on Google's Security:** For accounts managed by Google, meeting this control depends on the security practices implemented within Google's infrastructure, as detailed in their security whitepapers and compliance certifications.
- **Custom Implementation Risks:** If implementing custom password handling, using established, well-vetted cryptographic libraries and following best practices for hashing and salting is critical to avoid vulnerabilities.

Supplemental Guidance

- [Google Security Overview](#)
- [Google Workspace Security Whitepaper](#)
- [Google Security Whitepaper](#)
- [Authentication Concepts](#)
- [Encryption in transit for Google Cloud](#)

Control Domain	Identification and Authentication		
Control #	IA.L2-3.5.11		
Control Description	Obscure feedback of authentication information		
Key Services	- Google Sign In UI - Google Cloud CLI - Identity Platform	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for (inside hosted applications and systems): - Obscuring authentication information during the authentication process Google is responsible for (using Google Features): - Obscuring authentication information during the authentication process This control is met by default behavior in Google's standard interfaces and requires adherence to standard UI development practices for custom applications. However, the following key service(s)/interface(s) in Google Cloud Console may be used to support this control: - Google Sign-in User Interface - Custom Applications (including Identity Platform) - Google Cloud CLI (gcloud) A description of relevant features and implementation guidance is included below. Google Sign-in User Interface: When users authenticate to Google services (including Google Cloud Console and Google Workspace) via the standard web-based sign-in page, the password field automatically masks the characters being typed. Platform Behavior: This is standard, built-in behavior for password input fields on the web to prevent onlookers ("shoulder surfing") from seeing the password as it is typed. No customer configuration is required for this. Custom Applications (including Identity Platform): For custom web applications, including those using Identity Platform for authentication, the application's user interface must obscure password fields. Implementation			

- Developers building the application's frontend should use the standard HTML element `<input type="password">` for password entry fields. This element instructs the browser to mask the input automatically. Similar standard components should be used in mobile application development frameworks. Relying on standard components ensures compliance with this requirement.

Google Cloud CLI (gcloud): When the gcloud CLI prompts for sensitive information, such as a password during certain authentication flows, it typically does not echo the characters to the screen, providing obscurity.

- **Platform Behavior:** This is a common security practice for command-line tools handling sensitive input. No customer configuration is required for this behavior.

Additional Considerations

- **Standard Practice:** Obscuring password input is a fundamental security practice implemented by default in nearly all web browsers and operating systems for fields designated as password inputs.
- **Service Accounts:** This control is generally not applicable to service account authentication using keys, as the keys are typically handled programmatically or via file uploads, not typed interactively where feedback obscuring is relevant.

Supplemental Guidance

- [Google Security Overview](#)
- [Google Workspace Security Whitepaper](#)
- [Google Security Whitepaper](#)

Control Domain	Incident Response
Control #	IR.L2-3.6.1
Control Description	Establish an operational incident-handling capability for organizational systems that includes preparation, detection, analysis, containment, recovery, and user response activities

Key Services	• SCC • Cloud Logging • Cloud Monitoring • Organizational Policies • IAM • Chronicle Security Operations • VPC Flow Logs • Google Cloud Armor • Snapshots and Forensics • VPC Firewall Rules • Compute Engine • VPC Service Controls	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Establishing an operational incident handling capability; that - includes preparation - includes detection - includes analysis - includes containment - includes recovery - includes user response activities. When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control: • Security Command Center (SCC) • Cloud Logging • Cloud Monitoring • Organizational Policies • Identity and Access Management (IAM) • Chronicle Security Operations • VPC Flow Logs • Google Cloud Armor • Snapshots and Forensics • VPC Firewall Rules • Compute Engine • VPC Service Controls Preparation			

Security Command Center: Use Security Command Center to gain visibility into your security posture, identify misconfigurations, detect threats, and manage vulnerabilities. This is a foundational step in preparing for incidents.

- **Enable Security Command Center:**
 - Navigate to **Security > Security Command Center** in the Google Cloud Console.
 - Select your organization and choose a tier (Standard or Premium). Premium offers more advanced threat detection.
 - Activate desired services like Security Health Analytics, Web Security Scanner, Event Threat Detection, Container Threat Detection, and VM Threat Detection.
- **Configure Sources:**
 - Go to **Settings > Sources** within Security Command Center to ensure all relevant services are providing findings.
- **Set up Notifications:**
 - Go to **Settings > Notifications**.
 - Configure Pub/Sub topics to receive findings and integrate with other tools (e.g., SIEM, ticketing systems).
 - Set up email notifications for critical findings.

Cloud Logging: Centralized logging is critical for preparing for incident analysis. Ensure all relevant logs are being collected and stored appropriately.

- **Enable Data Access audit logs:**
 - Navigate to **IAM & Admin > Audit Logs**.
 - Select services and log types (Admin Read, Data Read, Data Write) to be logged. Be mindful of log volume and cost.
- **Create Log Sinks:**
 - Navigate to **Logging > Log Router**.
 - Click **Create Sink**.
 - Name the sink and choose a destination (Cloud Storage bucket, BigQuery dataset, Pub/Sub topic, or another Google Cloud project). This is crucial for log retention and analysis.
 - Define an inclusion filter to specify which logs to export (e.g., `resource.type="gce_instance"` or `severity>=ERROR`).
 - For long-term retention for compliance, sink logs to a Cloud Storage bucket with appropriate lifecycle policies or to BigQuery.

Cloud Monitoring: Set up alerts for suspicious activities or critical system states to enable early detection.

- **Create Log-based Metrics and Alerts:**
 - Navigate to **Logging > Logs Explorer**.
 - Enter a query for conditions you want to monitor (e.g., failed SSH attempts, IAM policy changes).
 - Click **Actions > Create metric**. Configure the metric.
 - Navigate to **Monitoring > Alerting**.
 - Click **+ Create Policy**.
 - Click **Add Condition**. Select the log-based metric you created.
 - Configure the trigger (e.g., threshold, duration).
 - Configure notification channels (email, SMS, PagerDuty, Slack, Pub/Sub).
- **Create Metric-based Alerts:**
 - Navigate to **Monitoring > Alerting**.
 - Click **+ Create Policy**.
 - Click **Add Condition**. Select a metric (e.g., CPU utilization, disk I/O, network traffic).
 - Configure the trigger and notification channels.

Organization Policies: Use Organization Policies to enforce security configurations across your Google Cloud environment, reducing the likelihood of incidents caused by misconfigurations.

- **Define and Enforce Policies:**
 - Navigate to **IAM & Admin > Organization Policies**.
 - Review the list of available policy constraints.
 - For each relevant constraint (e.g., "Disable Serial Port Access", "Restrict Public IP access for VMs"), click on it.
 - Click **Edit Policy**.
 - Choose **Customize**.
 - Select **Replace** or **Merge with parent**.
 - Set **Enforcement** to **Enforce** or **Allow / Deny** based on the policy logic.
 - Specify the desired values or rules (e.g., deny all for "Define allowed external IP addresses for VM instances" to prevent public IPs).
 - Apply the policy at the organization, folder, or project level.

IAM (Identity and Access Management): Properly configured IAM ensures that users and services have only the necessary permissions (principle of least privilege), reducing the potential impact of a compromised account.

- **Regularly Review IAM Policies:**

- Navigate to **IAM & Admin > IAM**.
- Review principals and their assigned roles at the organization, folder, and project levels.
- Use **IAM Recommender** to identify overly permissive roles.
- Remove or replace basic roles (Owner, Editor, Viewer) with more granular predefined or custom roles.
- **Implement IAM Conditions:**
 - When granting roles, use conditions to restrict access based on date/time, resource attributes, or request attributes for more fine-grained control.

Detection

Security Command Center: Use Security Command Center's integrated services for threat detection.

- **Monitor Findings:**
 - Regularly review the **Findings** tab in Security Command Center.
 - Filter by severity, category (e.g., "Threat"), or specific detector (e.g., Event Threat Detection, VM Threat Detection).
- **Event Threat Detection:** This service uses Google's intelligence to detect threats like malware, cryptomining, and outgoing DDoS attacks based on log analysis. Ensure it's enabled in Security Command Center.
- **VM Threat Detection:** This service detects potentially malicious applications running inside your virtual machines by analyzing signals from the hypervisor and guest OS.
- **Container Threat Detection:** Detects common container runtime attacks.

Cloud Logging & Cloud Monitoring (Alerts): As configured in the Preparation phase, alerts from Cloud Monitoring (based on logs or metrics) will serve as a primary detection mechanism.

- **Respond to Alerts:**
 - Ensure a defined process for acknowledging and investigating alerts received through configured notification channels.

Chronicle Security Operations: For organizations requiring advanced threat detection and hunting capabilities, Chronicle ingests security telemetry at scale and applies Google's threat intelligence.

- **Forward Logs to Chronicle:**
 - Configure log forwarding from Google Cloud (and other sources) to your Chronicle instance. This typically involves setting up a feed in Chronicle and

configuring log sinks in Cloud Logging to publish to a Pub/Sub topic that Chronicle ingests.

- **Use Detection Engine:**
 - Leverage Chronicle's rules engine (YARA-L) to create custom detection rules based on observed TTPs (Tactics, Techniques, and Procedures).
 - Utilize out-of-the-box threat detections provided by Chronicle.

VPC Flow Logs: Enable VPC Flow Logs to capture information about IP traffic going to and from network interfaces in your VPC network.

- **Enable VPC Flow Logs:**
 - Navigate to **VPC network > Subnets**.
 - Select a subnet. Click **Edit**.
 - Turn **Flow logs** to **On**. Configure sampling rate and aggregation interval as needed. Logs are sent to Cloud Logging.
- **Analyze Flow Logs:**
 - Use Logs Explorer to query VPC Flow Logs for suspicious traffic patterns (e.g., connections to known malicious IPs, unexpected data exfiltration).
 - Consider exporting flow logs to BigQuery for more complex analysis or to Chronicle.

Google Cloud Armor: Detects and mitigates web-based attacks and DDoS.

- **Review Security Policy Logs:**
 - If using Google Cloud Armor, ensure logging is enabled for security policies.
 - Navigate to **Logging > Logs Explorer** and filter for `resource.type="http_load_balancer"` and `jsonPayload.enforcedSecurityPolicy`.
 - Look for logs indicating blocked requests or preconfigured WAF rule triggers.

Analysis

Cloud Logging (Logs Explorer): Use Logs Explorer as the primary tool for analyzing logs related to a detected incident.

- **Query Logs**
 - Navigate to **Logging > Logs Explorer**.
 - Use the query builder or write custom queries to filter logs based on time range, resource, severity, log name, and keywords relevant to the incident.
 - Example queries:
 - `resource.type="gce_instance" AND jsonPayload.event_subtype="compute.instances.insert"` (VM creation)
 - `protoPayload.methodName="SetIamPolicy"` (IAM policy changes)

- `resource.labels.instance_id="your-vm-id"` AND `severity=ERROR` (Errors on a specific VM)

- **Correlate Events:**

- Use timestamps and request IDs (if available) to correlate events across different log sources.

Chronicle Security Operations: Chronicle excels at large-scale analysis and threat hunting.

- **Investigate Alerts:**

- Use Chronicle's UI to investigate alerts, view associated entities (users, assets, domains, IPs), and analyze timelines of activity.

- **Threat Hunting:**

- Perform IOC (Indicator of Compromise) searches across all ingested telemetry.
- Use Chronicle's UDM Search to look for patterns of behavior.

Security Command Center: Review findings related to the incident for context and impacted assets.

- **Examine Finding Details:**

- Click on a finding in Security Command Center to see details, affected resources, and sometimes recommended remediation steps.

Snapshots and Forensics: In some cases, you may need to perform deeper forensic analysis.

- **Create Disk Snapshots:**

- For a potentially compromised VM, navigate to **Compute Engine > Snapshots**.
- Click **Create Snapshot**. Select the source disk of the affected VM.
- This snapshot can be used to create a new disk for forensic analysis in an isolated environment.

- **Isolate for Analysis:**

- Create a dedicated, isolated VPC network and project for forensic analysis.
- Create a new VM in this isolated environment and attach a disk created from the snapshot of the potentially compromised system.

Containment

IAM: Limit or revoke access for compromised accounts.

- **Revoke Roles:**

- Navigate to **IAM & Admin > IAM**.
- Identify the compromised user or service account.
- Select the account and click **Edit principal**.
- Delete relevant roles or, for immediate broad containment, remove all roles.
- **Disable Service Accounts:**
 - Navigate to **IAM & Admin > Service Accounts**.
 - Select the compromised service account and click **Disable**.
- **Disable User Accounts (Google Workspace/Cloud Identity):**
 - If using Google Workspace or Cloud Identity, an administrator can suspend the user's account via the Admin console (admin.google.com).

VPC Firewall Rules: Isolate compromised resources at the network level.

- **Create Deny Rules**
 - Navigate to **VPC network > Firewall**.
 - Create new firewall rules with a high priority (lower number) to deny traffic to/from the compromised VM(s) or specific services.
 - Target rules using network tags or service accounts associated with the compromised resources.
 - For example, create an ingress rule denying all traffic to a tagged VM and an egress rule denying all traffic from that VM.
- **Modify Existing Rules**
 - Temporarily modify existing "allow" rules to be more restrictive or disable them if they are contributing to the incident.

Compute Engine: Stop or isolate compromised VMs.

- **Stop or Suspend a VM**
 - Navigate to **Compute Engine > VM instances**.
 - Select the compromised VM and click **Stop** or **Suspend**. Stopping is generally preferred for containment.
- **Remove from Load Balancer**
 - If the VM is part of an instance group managed by a load balancer, remove it from the instance group (**Compute Engine > Instance groups**).

Google Cloud Armor: Block malicious traffic.

- **Update Security Policies**
 - Navigate to **Network Security > Cloud Armor**.
 - Select the relevant security policy.

- Add new rules to deny traffic based on IP address, IP range, region code, or other request attributes identified during the analysis phase.

VPC Service Controls (if applicable): If data exfiltration is suspected and VPC Service Controls are in use, review and tighten perimeter configurations.

- **Review Perimeter Violations**
 - Check Cloud Audit Logs for VpcServiceControlsAudit entries to identify any unauthorized attempts to access services protected by perimeters.
- **Modify Perimeters (use with caution)**
 - If necessary, and after careful consideration of impact, modify service perimeters (**Security > VPC Service Controls**) to further restrict services or data access. This is a powerful control and changes should be well-understood.

Eradication & Recovery

Remove Malicious Code/Artifacts:

- Based on analysis, identify and remove any malicious files, processes, or configurations from affected systems. This might involve:
 - Rebuilding VMs from a known good image/template.
 - Restoring specific files from backups after verifying their integrity.

Restore from Backups/Snapshots

- **Restore VM from Snapshot**
 - Navigate to **Compute Engine > Snapshots**.
 - Select the snapshot taken before the incident (or a clean snapshot).
 - Click **Create instance** to create a new VM from the snapshot, or create a disk from the snapshot to attach to an existing or new VM.
- **Restore Data from Cloud Storage Backups**
 - If data was stored in Cloud Storage and backed up (e.g., using versioning or by copying to another bucket), retrieve the last known good version.
- **Restore Persistent Disks**
 - If using Persistent Disk snapshots, create a new disk from a clean snapshot and attach it to a VM.

Rebuild Systems: In many cases, rebuilding from a known-good state (Infrastructure as Code, golden images) is preferable to cleaning a compromised system.

- **Use Deployment Manager or Terraform:** If infrastructure is defined as code, redeploy affected resources.
- **Use Instance Templates:** Create new VMs from trusted instance templates.

Patch and Harden

- Ensure all systems are patched to the latest security levels.
- Review configurations and apply hardening guides based on lessons learned.
- Update vulnerabilities identified by Security Command Center.

Post-Incident Activities (User Response / Lessons Learned)

Communication

- Develop a communication plan for notifying internal stakeholders, customers (if applicable), and regulatory bodies as required. Google Cloud does not directly provide tools for this external communication, but your internal processes should cover it.

Lessons Learned Report

- Document the incident: how it was detected, analyzed, contained, and remediated.
- Identify the root cause.
- Evaluate the effectiveness of the incident response plan and tools.
- Propose improvements to security controls, monitoring, and response procedures.

Update Incident Response Plan

- Based on lessons learned, update the incident response plan, playbooks, and contact lists.

Security Command Center & Monitoring Adjustments

- Fine-tune Security Command Center settings. Mute findings that are not relevant or create custom modules for specific organizational risks.
- Adjust Cloud Monitoring alert thresholds and create new alerts based on the nature of the incident to improve future detection.

Access Transparency and Access Approval: While not directly for customer incident handling, these services provide visibility and control over Google administrator access, which can be relevant in certain scenarios or for forensic investigation if Google support involvement was part of the incident.

- **Review Access Transparency Logs:**
 - Navigate to **IAM & Admin > Access Transparency**.
 - Review logs if there's a need to understand Google administrator actions on your resources (usually related to support tickets).
- **Configure Access Approval:**
 - Navigate to **IAM & Admin > Access Approval**.
 - Configure policies to require your explicit approval before Google personnel can access your data or configurations for support purposes.

Additional Considerations

- **Incident Response Playbooks:** Develop specific playbooks for common incident types (e.g., malware infection, compromised credentials, data exfiltration).
- **Regular Testing:** Conduct tabletop exercises and simulations to test your incident response plan and team readiness.
- **External SIEM/SOAR Integration:** If using a third-party SIEM or SOAR, ensure logs and alerts from Google Cloud (via Pub/Sub sinks) are correctly ingested and parsed.
- **Forensic Preparedness:** Understand how to preserve evidence (e.g., disk snapshots, log exports) in a forensically sound manner if legal or further investigation is required.
- **Shared Responsibility Model:** Clearly understand the shared responsibility model for security in the cloud. Google secures the underlying infrastructure, while you are responsible for securing what you put in the cloud.

Supplemental Guidance

- [Security Command Center documentation | Google Cloud](#)
- [Cloud Logging documentation](#)
- [Cloud Monitoring documentation](#)
- [Google Security Operations documentation](#)
- [Identity and Access Management documentation - IAM](#)
- [VPC firewall rules | Cloud NGFW](#)
- [Google Cloud Armor documentation](#)
- [Create archive and standard disk snapshots | Compute Engine Documentation | Google Cloud](#)
- [Overview of Access Transparency | Google Cloud](#)
- [Overview of Access Approval | Google Cloud](#)
- [Introduction to the Organization Policy Service | Resource Manager Documentation | Google Cloud](#)
- [Observability: cloud monitoring and logging | Google Cloud](#)
- [Data incident response process | Security | Google Cloud](#)
- [Google security overview](#)

- [Route log entries | Cloud Logging](#)
- [Alerting overview | Cloud Monitoring](#)

Control Domain	Incident Response		
Control #	IR.L2-3.6.2		
Control Description	Track, document, and report incidents to designated officials and/or authorities both internal and external to the organization		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Tracking incidents - Documenting incidents - Identifying authorities to whom incidents are to be reported - Identifying organizational officials to whom incidents are to be reported - Notifying the identified authorities of incidents - Notifying the identified organizational officials of incidents <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			
N/A			

Control Domain	Incident Response		
Control #	IR.L2-3.6.3		
Control Description	Test the organizational incident response capability		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared

			☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Testing the incident response capability <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Maintenance		
Control #	MA.L2-3.7.1		
Control Description	Perform maintenance on organizational systems		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: a. Performing system maintenance <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Maintenance		
Control #	MA.L2-3.7.2		
Control Description	Provide controls on the tools, techniques, mechanisms, and personnel used to conduct system maintenance		
Key Services	• IAM • IAP for TCP Forwarding • OS Login • VPC Firewall Rules • Cloud Audit Logs • Artifact Registry • Cloud Shell	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: - Controlling tools used to conduct system maintenance - Controlling techniques used to conduct system maintenance - Controlling mechanisms used to conduct system maintenance - Controlling personnel used to conduct system maintenance <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Maintenance		
Control #	MA.L2-3.7.3		
Control Description	Ensure equipment removed for off-site maintenance is sanitized of any CUI		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			

Google is responsible for:

- a. Sanitizing equipment to be removed from organizational spaces for off-site maintenance of any CUI

Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.

Supplemental Guidance

- N/A

Control Domain	Maintenance		
Control #	MA.L2-3.7.4		
Control Description	Check media containing diagnostic and test programs for malicious code before the media are used in organizational systems		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: a. Checking media containing diagnostic and test programs for malicious code before being used in organizational systems that process, store, or transmit CUI <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Maintenance		
Control #	MA.L2-3.7.5		
Control Description	Require multifactor authentication to establish nonlocal maintenance sessions via external network connections and terminate such connections when nonlocal maintenance is complete		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: - Using multifactor authentication to establish nonlocal maintenance sessions via external network connections - Terminating nonlocal maintenance sessions established via external network connections when nonlocal maintenance is complete <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			
Supplemental Guidance			
N/A			

Control Domain	Maintenance		
Control #	MA.L2-3.7.6		
Control Description	Supervise the maintenance activities of maintenance personnel without required access authorization.		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			

Google is responsible for:

- a. Supervising maintenance personnel without required access authorization during maintenance activities.

Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.

Supplemental Guidance

- N/A

Control Domain	Media Protection		
Control #	MP.L2-3.8.1		
Control Description	Protect (i.e., physically control and securely store) system media containing CUI, both paper and digital		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for a. Physically controlling paper media containing CUI b. Physically controlling digital media containing CUI c. Securely storing paper media containing CUI d. Securely storing digital media containing CUI <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Media Protection		
Control #	MP.L2-3.8.2		
Control Description	Limit access to CUI on system media to authorized users		
Key Services	• IAM • Cloud Identity / Google Workspace • Cloud Storage Access Controls • Persistent Disk Access Controls • Database Access Controls • Service Accounts • Cloud Audit Logs	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Limiting access to CUI on system media to authorized users. When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control for <i>digital</i> media: • Identify and Access Management (IAM) • Cloud Identity / Google Workspace • Cloud Storage Access Controls (IAM) • Persistent Disk Access Controls (IAM + OS Controls) • Database Access Controls (Example BigQuery) • Service Accounts • Cloud Audit Logs (Data Access) A description of relevant features and implementation guidance is included below. Identity and Access Management (IAM): The primary tool for defining authorized access. Grant roles to specific identities (users, groups, service accounts) at the necessary resource scope (organization, folder, project, or resource level like a Cloud Storage bucket). • Implementation 1. Navigate to IAM & Admin > IAM. 2. Identify the resource containing CUI (e.g., project, specific storage bucket). 3. Click Grant Access.			

4. In **New principals**, enter the email addresses of authorized users, Google Groups (recommended for user management), or service accounts.
5. In **Assign roles**, select the most specific predefined role(s) that grant the necessary permissions (e.g., roles/storage.objectViewer for read-only access to Cloud Storage, roles/compute.viewer for viewing disk info, roles/bigquery.dataViewer for reading BigQuery data). Avoid basic roles (Owner, Editor, Viewer) whenever possible. Create custom roles if predefined roles are too broad.
6. **Save** the policy. Regularly review bindings under IAM.

Cloud Identity / Google Workspace: Manage user accounts and groups used in IAM policies.

- **Implementation**

1. Use the Google Admin console (admin.google.com) for Google Workspace or Cloud Identity (Free/Premium).
2. Create Google Groups (e.g., cui-project-alpha-readers@yourdomain.com, cui-storage-bucket-editors@yourdomain.com).
3. Add authorized user accounts to these groups.
4. Grant IAM roles to these Google Groups instead of individual users for easier management.

Cloud Storage Access Controls (IAM): Control access to buckets and objects. Uniform bucket-level access is recommended for CUI.

- **Implementation**

1. Ensure the target bucket uses **Uniform** access control mode (Bucket details > Permissions tab).
2. Use IAM (as described in #1) applied at the project level (if access applies to all buckets) or directly on the specific bucket (Permissions tab > Grant Access) to assign roles like roles/storage.objectViewer, roles/storage.objectCreator, roles/storage.objectAdmin to authorized principals (users, groups, service accounts).

Persistent Disk Access Controls (IAM + OS Controls): IAM controls who can manage the disk resource itself (attach, detach, snapshot). OS controls manage access to the file system *on* the disk once attached.

- **Implementation**

1. Use IAM (as described in #1) to grant Compute Engine roles (e.g., roles/compute.instanceAdmin) only to users/services authorized to manage VM/disk lifecycle.
2. Within the VM's operating system (Linux/Windows), use standard OS tools (e.g., chmod, chown, NTFS permissions) to restrict file/directory access on the mounted Persistent Disk file system to only authorized OS users. OS users should correspond to the organization's authorized personnel. OS Login can help manage Linux user access based on IAM.

Database Access Controls (Example BigQuery): Combine IAM for API-level actions with dataset/table level permissions.

- **Implementation (BigQuery)**

1. Use IAM (as described in #1) to grant project/dataset level roles like roles/bigquery.dataViewer, roles/bigquery.dataEditor, roles/bigquery.user to authorized principals.
2. Navigate to **BigQuery**, select the dataset containing CUI.
3. Click **Sharing > Permissions > Add Principal**.
4. Add authorized principals (users, groups, service accounts) and assign dataset-level IAM roles (e.g., BigQuery Data Viewer).
5. For table-level control (less common, more granular), use Authorized Views or manage table ACLs via the API/bq command-line tool.

Service Accounts: Grant identities to applications needing access to media containing CUI.

- **Implementation**

1. Navigate to **IAM & Admin > Service Accounts**.
2. Create dedicated service accounts for specific applications/tasks needing access to CUI.
3. Grant minimal necessary IAM roles (as described in #1) directly to these service accounts.
4. Securely manage keys or preferably use Workload Identity Federation or attach service accounts directly to Compute Engine instances.

Cloud Audit Logs (Data Access): Log read/write operations on resources holding CUI.

- **Implementation**

1. Navigate to **IAM & Admin > Audit Logs**.

2. For services storing CUI (e.g., Cloud Storage, BigQuery, Compute Engine), enable Data Access audit logs (DATA_READ, DATA_WRITE). Note potential cost implications.
3. Monitor these logs in **Logging > Logs Explorer** to verify access patterns align with authorized usage. Filter logs by principal, resource, and permissions checked.

Additional Considerations

- **Principle of Least Privilege:** Always assign the minimum permissions required for a user or service to perform its authorized function. Avoid broad roles.
- **Regular Access Reviews:** Periodically review who has access to digital media containing CUI and remove permissions that are no longer necessary. Use tools like IAM Recommender or Policy Analyzer.
- **Group Management:** Use groups in Cloud Identity/Workspace to manage user access. Adding/removing users from groups is simpler than modifying numerous IAM policies.
- **Separation of Duties:** Where appropriate, structure roles and permissions to separate sensitive actions (e.g., data reading vs. data deletion vs. access administration).

Supplemental Guidance

- [IAM Overview](#)
- [Understanding IAM Roles](#)
- [Using IAM with Cloud Storage](#)
- [Access control for Compute Engine resources](#)
- [Introduction to BigQuery access control](#)
- [Cloud Audit Logs Overview](#)
- [Enabling Data Access audit logs](#)

Control Domain	Media Protection
Control #	MP.L1-3.8.3
Control Description	Sanitize or destroy information system media containing CUI before disposal or release for reuse

Key Services	• Data Deletion Mechanisms • Cloud KMS • Cloud Audit Logs	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: - Sanitizing or destroying customer-controlled information system media containing CUI before disposal - Sanitizing system media containing CUI before it is released for reuse.			
Supplemental Guidance			
• N/A			

Control Domain	Media Protection		
Control #	MP.L2-3.8.4		
Control Description	Mark media containing CUI indicating the distribution limitations		
Key Services	• Resource Manager Labels • Resource Manager Tags • Cloud DLP • Data Catalog • Naming Conventions • IAM Conditions • Security Command Center	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Marking media containing CUI with applicable CUI markings - Marking media containing CUI with distribution limitations When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control for <i>digital</i> media: • Resource Manager Labels			

- Resource Manager Tags
- Cloud DLP
- Data Catalog
- Naming Conventions
- IAM Conditions
- Security Command Center

A description of relevant features and implementation guidance is included below.

Resource Manager Labels: Apply key-value pairs to resources for identification, categorization, and automation. Labels are flexible and good for human-readable markings.

- **Implementation**

1. Define a consistent labeling scheme (e.g., cui = true, data_sensitivity = high, distribution = no_foreign, handling = proprietary).
2. Navigate to the resource you want to mark (e.g., **Compute Engine > VM instances**, select instance > **Edit**; **Cloud Storage > Buckets**, select bucket > **Edit labels**; Project settings > **Labels**).
3. Under the **Labels** section, click **Add Label**.
4. Enter the **Key** and **Value** according to your defined scheme.
5. Add multiple labels as needed.
6. **Save** the changes.
7. Use labels for filtering resources, cost allocation, or in custom scripts/automation.

Resource Manager Tags: Apply key-value pairs, often used for programmatic control and hierarchical policy inheritance. Tags are distinct from labels and network tags.

- **Implementation**

1. Define Tag Keys (e.g., dataClassification, accessLevel) and allowed Tag Values (e.g., public, internal, cui, level1, level2, level3) at the Organization or Project level via **IAM & Admin > Tags**. Requires specific permissions (roles/resourcemanager.tagAdmin).
2. Bind specific Tag Values to resources (Projects, Folders, Organization). Navigate to the resource, look for a **Tags** section or option, and select the appropriate pre-defined Tag Key and Value to attach. Requires roles/resourcemanager.tagUser.
3. Use tags in IAM Conditions or Organization Policies for enforcement based on CUI status.

Cloud Data Loss Prevention (DLP): Discover potential CUI within data stores and report findings, which can inform manual marking or trigger automated labeling actions.

- **Implementation**

1. Configure DLP Inspection Jobs or Discovery Scans (as described for MP.3.8.1) targeting resources expected to hold CUI.
2. Review DLP findings in the console or Security Command Center. Findings indicate the presence and type of potential CUI.
3. Based on findings, manually apply appropriate Labels or Tags to the resources.
4. (Advanced) Configure DLP job actions to publish findings to Pub/Sub, triggering a Cloud Function that automatically applies predefined labels/tags to the resource where CUI was found.

Data Catalog: Create a metadata inventory of data assets, allowing for richer classification, including applying tags for CUI status, sensitivity, and handling rules.

- **Implementation**

1. Navigate to **Dataplex > Catalog**.
2. Enable Data Catalog API if needed. Search for data assets (e.g., BigQuery tables, Pub/Sub topics).
3. Select an asset entry. Use the **Attach tags** option.
4. Create or use existing Tag Templates (defining fields like cui_present, distribution_code, handling_instructions).
5. Fill in the values for the specific asset according to your CUI policy.
6. Use Data Catalog for searching and understanding data assets based on these markings.

Naming Conventions (Procedural): Establish and enforce clear naming standards for resources that indicate the presence and sensitivity of data.

- **Implementation**

1. Define a standard (e.g., [prj|bkt|vm|disk]-[env]-[cui_indicator]-[description]). Example prj-prod-cui-financials, bkt-dev-noncui-logs.
2. Document the standard and train developers and operators.
3. Use linters or policy checks in CI/CD pipelines to enforce naming conventions.

IAM Conditions: Use applied labels or tags to create conditional role bindings, limiting access based on the resource's marking.

- **Implementation**

1. When granting access in **IAM & Admin > IAM**, click **Add IAM Condition**.
2. Select Condition Type **Resource > Tag** or **Resource > Label**.
3. Specify the Tag Key/Value or Label Key/Value (e.g., `resource.matchTag('your_org_id/dataClassification', 'cui')` or `resource.labels.cui == 'true'`).
4. Build the condition expression to grant access only if the resource has the appropriate CUI marking.

Security Command Center: Display assets and their associated labels, tags, and DLP findings, providing a central place to view marked resources and identified CUI.

- **Implementation**

1. Navigate to **Security > Security Command Center > Assets**.
2. View assets by type or project. Columns for Labels and Tags can often be displayed.
3. Use filters based on labels or tags to quickly identify CUI-marked resources.
4. Review DLP findings related to CUI under the **Findings** tab.

Additional Considerations

- **Consistency is Key:** Apply markings consistently across all resources according to the defined policy. Inconsistent marking reduces effectiveness. Use guidance from the [CUI Marking Handbook](#) published in accordance with Executive Order 13556.
- **Policy Definition:** A clear, documented policy defining the marking scheme (keys, values, what they mean) and when/how to apply them is essential.
- **Physical Media:** Remember that this control also applies to physical media (printouts, external hard drives, tapes). Procedures for physically marking, handling, and storing these items according to CUI guidelines are required outside of Google Cloud configuration.
- **Automation:** Leverage DLP findings, Cloud Functions, and potentially policy checks in deployment pipelines (e.g., Terraform using `google_tags_tag_binding`) to automate marking where possible.

Supplemental Guidance

- [Create and update labels](#)
- [Resource Manager Tags Overview](#)
- [Creating and managing tags](#)
- [Attaching tags to resources](#)

- [Cloud Data Loss Prevention \(DLP\) Documentation](#)
- [Data Catalog Overview](#)
- [IAM Conditions Overview](#)
- [Security Command Center Assets display](#)

Control Domain	Media Protection		
Control #	MP.L2-3.8.5		
Control Description	Control access to media containing CUI and maintain accountability for media during transport outside of controlled areas		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: - Controlling access to media containing CUI - Maintaining accountability for media containing CUI during transport outside of controlled areas <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			
Supplemental Guidance			
N/A			

Control Domain	Media Protection
Control #	MP.L2-3.8.7
Control Description	Control the use of removable media on system components. Items handling CUI.

Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Controlling the use of removable media on system components. <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Media Protection		
Control #	MP.L2-3.8.8		
Control Description	Prohibit the use of portable storage devices when such devices have no identifiable owner.		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Prohibiting the use of portable storage devices when such devices have no identifiable owner <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			

- N/A

Control Domain	Media Protection								
Control #	MP.L2-3.8.9								
Control Description	Protect the confidentiality of backup CUI at storage locations								
Key Services	• IAM • Cloud Storage • Persistent Disk Snapshots • Cloud SQL Backups • Google Cloud Backup and DR Service • Cloud KMS • VPC Service Controls	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Protecting the confidentiality of backup CUI at storage locations. When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control: • Identity and Access Management (IAM) • Cloud Storage (as a backup target) • Persistent Disk Snapshots • Cloud SQL Backups • Google Cloud Backup and DR Service • Cloud Key Management Service (Cloud KMS) • VPC Service Controls A description of relevant features and implementation guidance is included below. Identity and Access Management (IAM): Use IAM to restrict access to backup CUI and the systems used to manage backups. • Least Privilege for Backup Access: ○ Grant specific IAM roles to users, groups, or service accounts that require access to backup data or backup administration functions. Avoid broad roles like Owner or Editor.									

- Examples:
 - For Cloud Storage buckets holding backups roles/storage.objectViewer for read-only access to backups, roles/storage.objectAdmin for managing backups (use with caution).
 - For Persistent Disk snapshots roles/compute.storageAdmin allows managing snapshots; more granular roles might be needed if separating creation from restoration.
 - For Cloud SQL roles/cloudsql.editor or roles/cloudsql.admin for managing instances and backups; roles/cloudsql.viewer for viewing.
- Navigate to **IAM & Admin > IAM**, select the relevant project, and carefully assign roles.
- **Service Account for Backup Operations:**
 - If using automated backup scripts or third-party tools, use dedicated service accounts with narrowly scoped permissions necessary for backup and restoration tasks.

Cloud Storage (as a backup target): Often used for storing database dumps, file backups, and other backup data.

- **Encryption at Rest:** Data in Cloud Storage is encrypted at rest by default using Google-managed encryption keys.
 - **CMEK for Backup Buckets:** For enhanced control over backup confidentiality:
 1. Create a key in **Cloud Key Management Service (KMS)** (Security > Key Management).
 2. When creating the Cloud Storage bucket intended for backups, under **Encryption type**, select "Customer-managed encryption key (CMEK)" and specify the KMS key.
 3. Alternatively, you can set a default KMS key on an existing bucket for new backup objects.
- **Access Control:**
 - Apply granular IAM permissions at the bucket level (and object level if necessary) as described in the IAM section.
 - Enforce **Public Access Prevention** on buckets storing backup CUI (default for new buckets, can be set via Bucket settings > Permissions > Public access prevention).
- **Object Versioning and Lifecycle Rules:**
 - Enable Object Versioning on backup buckets to protect against accidental deletion or overwriting of backup files.

- Use Object Lifecycle Management (Bucket > Lifecycle tab) to manage retention and deletion of old backups securely.

Persistent Disk Snapshots (Compute Engine VM Backups): Use to back up the state of VM disks.

- **Encryption at Rest:**
 - Persistent Disk snapshots are encrypted at rest by default.
 - If the source Persistent Disk is encrypted with CMEK, the snapshot will also be encrypted with the same CMEK. The KMS key is required to restore from the snapshot or create a disk from it.
- **Access Control:**
 - IAM roles control who can create, delete, list, and use snapshots (e.g., roles/compute.storageAdmin, roles/compute.instanceAdmin).
 - Snapshots can be stored in custom locations (regions or multi-regions) for CUI storage requirements.

Cloud SQL Backups (Managed Database Backups): Cloud SQL provides automated and on-demand backups for database instances.

- **Encryption at Rest:** Cloud SQL data, including backups, is encrypted at rest by default.
 - **CMEK for Cloud SQL Backups:** Some database types in Cloud SQL support CMEK for the database instance, which extends to its backups.
 1. Enable CMEK when creating or editing a Cloud SQL instance (Instance details > Edit > Data protection > Encryption). Select a key from Cloud KMS.
- **Access Control:**
 - IAM roles for Cloud SQL (e.g., roles/cloudsql.editor, roles/cloudsql.admin) control who can manage instances, including backup and restore operations.
 - Access to exported backups (e.g., to Cloud Storage) is then controlled by Cloud Storage IAM permissions.

Google Cloud Backup and DR Service: A managed backup and disaster recovery solution for various workloads.

- **Encryption:**
 - This service encrypts backup data both in transit and at rest. Review the service's specific documentation for details on its encryption mechanisms and key management options (often leveraging Google-managed keys or allowing CMEK integration for the storage targets it uses).
- **Access Control:**

- Access to the Backup and DR console and its operations is controlled by IAM permissions specific to the backupdr.googleapis.com service. Assign roles like roles/backupdr.admin or roles/backupdr.user based on responsibilities.

Cloud Key Management Service (Cloud KMS): Manage encryption keys for CMEK used to protect backup CUI.

- **Key Management for Backups:** When using CMEK for Cloud Storage buckets, Persistent Disk snapshots, or Cloud SQL instances, Cloud KMS is used to create and manage these keys.
 1. Navigate to **Security > Key Management**.
 2. Create key rings and keys in appropriate regions, following best practices for key rotation and permissions.
 3. Restrict IAM permissions on KMS keys (roles/cloudkms.cryptoKeyEncrypterDecrypter, roles/cloudkms.viewer) to control who can use them for encrypting/decrypting backup data.

VPC Service Controls: Help prevent exfiltration of backup CUI from storage locations.

- **Perimeter Protection for Backup Storage:**
 1. Navigate to **Security > VPC Service Controls**.
 2. Create or amend service perimeters to include projects where backup CUI is stored and the relevant storage services (e.g., storage.googleapis.com, cloudsql.googleapis.com, compute.googleapis.com).
 3. This restricts data movement, ensuring backups cannot be copied to unauthorized locations outside the perimeter.

Additional Considerations

- **Data Classification:** Clearly identify which backups contain CUI to apply appropriate protection levels.
- **Backup Data Lifecycle Management:** Define and implement policies for how long backup CUI should be retained and how it should be securely disposed of (e.g., using Cloud Storage lifecycle rules).
- **Geographic Location of Backups:** Store backups in geographic regions that comply with any data residency requirements for the CUI.
- **Regular Testing of Backups:** Regularly test the restoration process from backups containing CUI to ensure their integrity and availability, performing these tests in a secure manner.
- **Least Privilege for Restoration:** Ensure that permissions to restore backups are also tightly controlled and only granted when necessary.

- **Immutable Backups:** For critical CUI backups, consider using features like Cloud Storage Bucket Lock in "Compliance" mode to make backup data immutable for a defined retention period.

Supplemental Guidance

- [IAM overview | IAM Documentation | Google Cloud](#)
- [Cloud Storage documentation](#)
- [Create archive and standard disk snapshots | Compute Engine Documentation | Google Cloud](#)
- [Use customer-managed encryption keys \(CMEK\) | Cloud SQL for MySQL Backup and DR Service documentation | Google Cloud](#)
- [Cloud Key Management Service overview | Cloud KMS | Google Cloud](#)
- [Overview of VPC Service Controls | Google Cloud](#)
- [Google Cloud security best practices center](#)

Domain	Personnel Security		
Control #	PS.L2-3.9.1		
Control Description	Screen individuals prior to authorizing access to organizational systems containing CUI		
Key Services	N/A	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: Screening individuals prior to authorizing access to organizational systems containing CUI <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			

- N/A

Control Domain	Personnel Security								
Control #	PS.L2-3.9.2								
Control Description	Ensure that organizational systems containing CUI are protected during and after personnel actions such as terminations and transfers								
Key Services	• Cloud Identity / Google Workspace • IAM • Google Groups or Cloud Identity Groups • Cloud Audit Logs	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Terminating system access and credentials consistent with personnel actions such as termination or transfer b. Ensuring system access and credentials are terminated consistent with personnel actions such as termination or transfer c. Ensuring the system is protected during and after personnel transfer actions When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control: • Cloud Identity / Google Workspace (Identity Lifecycle Management) • Identity and Access Management (IAM) • Google Groups or Cloud Identity Groups • Cloud Audit Logs A description of relevant features and implementation guidance is included below. Cloud Identity / Google Workspace (Identity Lifecycle Management): Manage the lifecycle of user identities that access Google Cloud. Prompt action here is critical upon personnel changes. • Guidance for Terminations ○ Immediately suspend or delete the user's Google account (Cloud Identity or Google Workspace) upon notification of termination. Suspension is often preferred initially to preserve data and allow for investigation if needed, followed by deletion according to organizational policy.									

- Reset the user's password and revoke all active login sessions and application-specific passwords.
- Manage data owned by the user (e.g., transfer Google Drive ownership, export email if needed) according to your offboarding policy.
- **Guidance for Transfers/Role Changes**
 - While the identity itself may not change, ensure HR processes trigger a review of the user's group memberships and direct role assignments.
- **Implementation (Google Admin Console - admin.google.com)**
 - **Suspend a User**
 1. Navigate to **Directory > Users**.
 2. Hover over the user and click **More options > Suspend user**. Confirm the suspension.
 - **Delete a User**
 1. Navigate to **Directory > Users**.
 2. Hover over the user and click **More options > Delete user**.
 3. You will be prompted to transfer ownership of their data (e.g., Drive files, Calendar primary events). Assign a new owner or choose to not transfer. You may also have options to export user data.
 4. Confirm the deletion.
 - **Reset Password**
 1. Navigate to **Directory > Users**.
 2. Click on the user's name.
 3. Click **Reset password**. You can auto-generate a password or create one. You can also choose to require the user to change their password at next sign-in (less relevant for terminations).
 - **Sign Out / Revoke Sessions**
 1. Navigate to **Directory > Users**.
 2. Click on the user's name.
 3. Go to the "Security" section.
 4. Click on "Signed in devices" or "Login cookies" (options may vary slightly) and look for options to sign out or revoke sessions. You can also revoke 2-Step Verification backup codes and security keys.
 - **Manage App Passwords** If app passwords were used, these should be revoked from the user's security settings (often also managed by the admin via the user's security panel).

Identity and Access Management (IAM): Control authorization to Google Cloud resources. Permissions must be promptly updated based on personnel actions.

- **Guidance for Terminations:** Remove all direct IAM role assignments for the terminated user from all Google Cloud resources (organization, folders, projects).

- **Guidance for Transfers/Role Changes:** Review all existing IAM role assignments for the user. Remove roles that are no longer necessary for their new position and add any new roles required, adhering to the principle of least privilege.
- **Implementation (Google Cloud Console)**
 - Navigate to **IAM & Admin > IAM**.
 - **For Terminations**
 1. Filter by the terminated user's email address in the "Filter" bar across your organization, all folders, and all projects.
 2. For each role binding found for the user, select the checkbox next to their name and click **Remove Access** or the delete/trash can icon.
 - **For Transfers/Role Changes**
 1. Filter by the user's email address.
 2. Review their current roles. Edit the member's roles by clicking the pencil icon next to their entry, remove unnecessary roles, and add new ones as required for their new responsibilities. Click **Save**.

Google Groups or Cloud Identity Groups: If access is primarily managed via groups (a recommended practice), modifying group membership is an efficient way to update access.

- **Guidance for Terminations:** Remove the terminated user from all Google Groups or Cloud Identity Groups that grant access to Google Cloud resources.
- **Guidance for Transfers/Role Changes:** Review the user's group memberships. Remove them from groups associated with their old role and add them to groups relevant to their new role.
- **Implementation (Google Admin Console - admin.google.com or Cloud Console for Cloud Identity Groups)**
 1. Navigate to **Directory > Groups** (in Admin Console) or manage Cloud Identity Groups via API/gcloud/Console if used primarily for IAM.
 2. Select the relevant group.
 3. Go to **Members**.
 4. Find the user and remove them from the group.
 5. For transfers, add them to new, relevant groups.

Cloud Audit Logs: Use Cloud Audit Logs to verify that user accounts were suspended/deleted and IAM roles/group memberships were changed in accordance with personnel action timelines.

- **Implementation**
 - **Reviewing Logs for User Account Changes (Admin Activity Audit Log)**
 1. In the Google Admin console, navigate to **Reporting > Audit > Admin**. Filter by events like "Suspend user," "Delete user," "Change password."

- **Reviewing Logs for IAM Changes (Cloud Audit Logs)**
 1. In Google Cloud Console, navigate to **Logging > Logs Explorer**.
 2. Filter for `protoPayload.methodName="SetIamPolicy"` to review IAM policy changes.
 3. Filter for group membership changes if using Cloud Identity Groups managed via API (logs `google.identity.groups.v1.GroupService.UpdateMembership` or similar).

Additional Considerations

- **Timeliness:** Automate de-provisioning actions where possible by integrating HR systems with Google Cloud identity tools (e.g., using Google Cloud Directory Sync with triggers, or custom SCIM solutions). Speed is critical, especially for terminations.
- **Checklists:** Use detailed offboarding and transfer checklists to ensure all necessary steps are completed, including access revocation from Google Cloud and other enterprise systems.
- **Asset Retrieval:** Ensure all physical assets (laptops, phones, security tokens) are retrieved from terminated employees.
- **Data Ownership and Access:** Clearly define procedures for transferring ownership of data (e.g., in Google Drive, project resources) and managing access to shared mailboxes or resources previously managed by the departed or transferred individual.
- **Review Shared Resources:** Audit access to resources shared directly by the individual (e.g., Google Drive files shared publicly or with specific external users) and adjust as necessary.
- **External Collaborators:** Apply similar principles if the personnel action involves an external collaborator who had access to your Google Cloud resources.

Supplemental Guidance

- [Manage access to projects, folders, and organizations | IAM Documentation | Google Cloud](#)

Control Domain	Physical Protection
Control #	PE.L1-3.10.1
Control Description	Limit physical access to organizational systems, equipment, and the respective operating environments to authorized individuals

Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: a. Identifying authorized individuals allowed physical access b. Limiting physical access to organizational systems to authorized individuals c. Limiting physical access to equipment to authorized individuals d. Limiting physical access to operating environments to authorized individuals <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Physical Protection		
Control #	PE.L2-3.10.2		
Control Description	Protect and monitor the physical facility and support infrastructure for organizational systems		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: a. Protecting the physical facility where organizational systems reside b. Protecting the support infrastructure for organizational systems c. Monitoring the physical facility where organizational systems reside d. Monitoring the support infrastructure for organizational systems			

Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.

Supplemental Guidance

- N/A

Control Domain	Physical Protection		
Control #	PE.L1-3.10.3		
Control Description	Escort visitors and monitor visitor activity		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: - Escorting visitors - Monitoring visitor activity <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Physical Protection
Control #	PE.L1-3.10.4
Control Description	Maintain audit logs of physical access

Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: a. Maintaining audit logs of physical access <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			
Supplemental Guidance			
• N/A			

Control Domain	Physical Protection		
Control #	PE.L1-3.10.5		
Control Description	Control and manage physical access devices		
Key Services	N/A	Control Responsibility	☒ Google ☐ Shared ☐ Customer
Customer Implementation Description			
Google is responsible for: a. Identifying physical access devices b. Controlling physical access devices c. Managing physical access devices <i>Based on the scope of this implementation guide, Google is responsible for the implementation of this control. However, your CUI Boundary may include systems, applications, facilities, or tools outside of Google Cloud, therefore, additional control implementation responsibility may be required.</i>			

Supplemental Guidance
N/A

Control Domain	Physical Protection		
Control #	PE.L2-3.10.6		
Control Description	Enforce safeguarding measures for CUI at alternate work sites		
Key Services	N/A	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining safeguarding measures for CUI for alternate work sites - Enforcing safeguarding measures for CUI for alternate work sites <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			
N/A			

Control Domain	Risk Assessment		
Control #	RA.L2-3.11.1		
Control Description	Periodically assess the risk to organizational operations (including mission, functions, image, or reputation), organizational assets, and individuals, resulting from the operation of organizational systems and the associated processing, storage, or transmission of CUI		
Key Services	- SCC - VM Manager - Cloud DLP - IAM	Control Responsibility	☐ Google ☐ Shared

	• Network Analyzer and Firewall Insights • Cloud Logging and Cloud Audit Logs • Cloud Monitoring		☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining the frequency to assess risk to organizational operations, organizational assets, and individuals; and - Assessing risk to organizational operations, organizational assets, and individuals resulting from the operation of an organizational system that processes, stores, or transmits CUI with the defined frequency. When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control: • Security Command Center (SCC) (including Security Health Analytics, Event Threat Detection, VM Threat Detection, Container Threat Detection, Web Security Scanner, Asset Inventory, Risk Manager (Premium), and Compliance Reporting (Premium)) • VM Manager (Vulnerability Reporting and Patch Management) • Cloud Data Loss Prevention (DLP) • Identity and Access Management (IAM) (including Policy Analyzer, Policy Simulator, and IAM Recommender) • Network Analyzer and Firewall Insights (within Network Intelligence Center) • Cloud Logging and Cloud Audit Logs • Cloud Monitoring A description of relevant features and implementation guidance is included below. Security Command Center (SCC): Activate and utilize SCC to gain comprehensive visibility into your Google Cloud assets, vulnerabilities, threats, misconfigurations, and compliance status. These findings are direct inputs for identifying risks. SCC Premium's Risk Manager feature directly supports risk quantification and prioritization. • Implementation ○ Activation: Navigate to Security > Security Command Center and activate the desired tier (Standard or Premium). ○ Enable Security Sources: Ensure all relevant built-in services are enabled (Security Health Analytics, Event Threat Detection, VM Threat Detection, Container Threat Detection, Web Security Scanner for your web applications). ○ Asset Inventory: Regularly review the discovered assets (Security Command Center > Assets) to understand the scope of systems that need			

to be included in your risk assessment. Note assets that process, store, or transmit CUI.

- **Vulnerability & Misconfiguration Findings (Security Health Analytics, VM Threat Detection, etc.):** Review findings in **Security Command Center > Findings**. These directly identify vulnerabilities (e.g., open firewalls, public buckets, unpatched software, insecure IAM configurations) that must be assessed for risk. Filter by severity and asset type.
- **Threat Detection Findings (Event Threat Detection, etc.):** Findings indicating active threats (e.g., malware, brute force, data exfiltration attempts) help in assessing the likelihood and impact of specific threat events.
- **Risk Manager (SCC Premium):** If using SCC Premium, configure Risk Manager by defining business criticality for your projects. Risk Manager then uses SCC findings, asset information, and your business context to generate risk scores and reports, helping you prioritize which vulnerabilities and threats pose the greatest risk to your CUI and operations.
- **Compliance Reporting (SCC Premium):** Use compliance reports to identify gaps against specific benchmarks (e.g., CIS, NIST). These gaps represent areas of increased risk.

VM Manager (Vulnerability Reporting and Patch Management): Use VM Manager to identify known vulnerabilities (CVEs) in the operating systems of your VMs. This data is a critical input for assessing risks to these systems, especially those handling CUI.

- **Implementation**

- **Enable VM Manager:** Ensure VM Manager is set up for your projects (**Compute Engine > VM Manager**). This typically involves enabling the OS Config API and installing/activating the OS Config agent on VMs.
- **Review Vulnerability Reports:** Regularly review vulnerability reports (**Compute Engine > VM Manager > Vulnerability reports**) which list CVEs affecting your VMs, along with severity levels. This information directly feeds into your risk assessment's vulnerability identification phase.
- **Patch Compliance:** Use patch compliance data (**Compute Engine > VM Manager > Patch compliance**) to understand the patch status of your VMs. Unpatched systems represent a higher risk.

Cloud Data Loss Prevention (DLP): Use Cloud DLP to understand where CUI is located (stored, processed). Knowing the location and type of CUI is essential for assessing the potential impact if those systems or data stores are compromised.

- **Implementation**

- Navigate to **Security > Data Loss Prevention**.

- Create inspection jobs or job triggers to scan Cloud Storage buckets, BigQuery tables, and Datastore for CUI using predefined or custom infoType detectors.
- The results of these scans will highlight datasets that require heightened risk assessment due to the presence of CUI.

Identity and Access Management (IAM): Use IAM tools to identify overly permissive access rights (a common vulnerability) and to understand the potential blast radius of compromised accounts.

- **Implementation**

- **IAM Recommender:** Review recommendations in **IAM & Admin > IAM or Security > Recommendations Hub** to identify and remove excessive permissions granted to users or service accounts.
- **Policy Analyzer:** Use **IAM & Admin > Policy Analyzer** to understand who has what access to which resources. This helps identify potential access paths that could be exploited.
- **Policy Simulator:** Before making IAM changes as part of risk mitigation, use **Policy Simulator** to test the impact of those changes.

Network Analyzer and Firewall Insights (within Network Intelligence Center): Use these tools to analyze your VPC network configurations and firewall rules for potential security risks.

- **Implementation**

- **Network Analyzer:** Enable the Network Analyzer API and review insights in **Network Intelligence > Network Analyzer**. It identifies issues like IP utilization insights, firewall rule shadowing, and connectivity issues.
- **Firewall Insights:** Review **VPC network > Firewall Insights** to identify overly permissive rules (e.g., rules allowing access from any IP address to sensitive ports), shadowed rules, or highly utilized rules. These insights directly inform network-related risk assessments.

Cloud Logging and Cloud Audit Logs: Analyze logs for trends, anomalies, or specific events (like repeated failed logins, unauthorized access attempts, or policy violations) that might indicate an increased likelihood of certain threats materializing.

- **Implementation**

- Ensure comprehensive logging is configured.

- Periodically review logs or set up alerts for security-relevant events in **Logging > Logs Explorer**. While primarily for incident detection, patterns observed can inform the likelihood component of risk assessments.

Cloud Monitoring: System outages or performance degradation can pose operational risks. Alerts from Cloud Monitoring on critical system health or security metrics (e.g., from log-based alerts) can highlight areas of concern that need to be factored into risk assessments.

- **Implementation**

- Configure uptime checks and alert policies for critical systems handling CUI via **Monitoring > Alerting**.
- The history of alerts and system performance can provide input on the stability and reliability of systems, which is relevant to operational risk assessment.

Additional Considerations

- **Periodicity:** Establish a defined frequency for conducting risk assessments (e.g., annually) and ensure they are also triggered by significant changes (new systems, major architectural changes, new CUI types, emerging threats).
- **Contextualization:** The information from Google Cloud tools needs to be contextualized with your organization's specific mission, operations, asset valuation (especially for CUI), and threat intelligence to accurately assess risk.
- **Documentation:** Thoroughly document your risk assessment process, findings, and resulting risk treatment plans. This documentation is essential for CMMC compliance.
- **Shared Responsibility:** Clearly understand the shared responsibility model. Focus your risk assessment on your responsibilities within Google Cloud and the interfaces with Google's responsibilities.

Supplemental Guidance

- [Security Command Center Overview](#)
- [About VM Manager | Compute Engine Documentation | Google Cloud](#)
- [Cloud Data Loss Prevention \(DLP\) Documentation](#)
- [IAM Policy Analyzer Overview](#)
- [Network Analyzer Overview](#)
- [Firewall Insights Overview](#)
- [NIST SP 800-30 Rev. 1, Guide for Conducting Risk Assessments | CSRC](#)

Control Domain	Risk Assessment		
Control #	RA.L2-3.11.2		
Control Description	Scan for vulnerabilities in organizational systems and applications periodically and when new vulnerabilities affecting those systems and applications are identified.		
Key Services	• SCC • VM Manager • Artifact Registry	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Defining the frequency to scan for vulnerabilities in organizational systems and applications - Performing vulnerability scans on organizational systems with the defined frequency - Performing vulnerability scans on applications with the defined frequency - Performing vulnerability scans on organizational systems when new vulnerabilities are identified - Performing vulnerability scans on applications when new vulnerabilities are identified When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control: • Security Command Center (SCC) (including Security Health Analytics, Web Security Scanner) • VM Manager (for OS and software package vulnerabilities on Compute Engine VMs) • Artifact Registry (with integrated Container Analysis for container image scanning) A description of relevant features and implementation guidance is included below. Security Command Center (SCC): Identify vulnerabilities at the cloud configuration level and for web applications. • Security Health Analytics (Periodic and Continuous Scanning for System Configurations): Enable Security Health Analytics to automatically scan your Google Cloud resources (projects, folders, organization) for misconfigurations and common vulnerabilities. Many detectors run continuously or on a frequent periodic basis. ○ Implementation ■ Navigate to Security > Security Command Center. Ensure it is activated (Standard or Premium tier).			

- Security Health Analytics is a built-in service. Its findings will appear in the **Findings** tab.
- Review findings regularly. These can include issues like overly permissive IAM roles, public Cloud Storage buckets, exposed database instances, disabled logging for critical services, and other configuration vulnerabilities that represent risks to your systems and CUI.
- New detectors are added by Google, and existing ones are updated, effectively scanning for new types of configuration vulnerabilities as they become relevant.
- **Web Security Scanner (Periodic and On-Demand Scanning for Applications):**
Use Web Security Scanner to find vulnerabilities in your App Engine, Google Kubernetes Engine (GKE), and Compute Engine web applications (must be publicly accessible and not behind IAP-only authentication for unauthenticated scans, or support Google accounts for authenticated scans).
 - **Implementation**
 1. Navigate to **Security > Security Command Center > Configuration > Web Security Scanner**.
 2. Click **Create scan**.
 3. Configure the scan
 - Provide a display name, starting URLs for your application.
 - Schedule Set to run periodically (e.g., weekly, daily) or run on-demand.
 - Authentication Choose "None" or "Google Account" (if your app uses Google Sign-In for authentication).
 - Export to Security Command Center Ensure this is checked so findings appear in SCC.
 4. Review scan results in **Security Command Center > Findings**, filtering by "Category" for "Web Security Scanner." Common findings include cross-site scripting (XSS), mixed content, and outdated libraries.
 5. When new web vulnerability types become widely known, you can trigger on-demand scans or rely on Google updating the scanner's capabilities for periodic scans.

VM Manager (Periodic and Event-Driven OS Vulnerability Scanning for Systems):

Enable VM Manager for your projects to get continuous vulnerability reporting for your VMs. VM Manager uses OS inventory data and correlates it with CVE databases. Scans are performed periodically (typically daily for active VMs after initial setup), and the vulnerability information is updated as new CVEs are published.

- **Implementation**

- **Enable API and Agent**

- Ensure the OS Config API is enabled for your projects.
 - Ensure the OS Config agent is installed and running on your Compute Engine VMs. Most Google-provided OS images include this agent. Verify its status.

- **Setup VM Manager** Navigate to **Compute Engine > VM Manager**. If it's the first time, you may need to click "Enable."

- **Review Vulnerability Reports**

- Go to **Compute Engine > VM Manager > Vulnerability reports**.
 - This dashboard shows an overview of vulnerabilities across your VMs, filterable by severity, CVE ID, OS, etc. It lists affected VMs and the specific CVEs.
 - Because VM Manager's vulnerability data sources are regularly updated, it effectively scans "when new vulnerabilities affecting those systems are identified" by re-evaluating the OS inventory against the latest CVE information.

Artifact Registry (with Container Analysis - Periodic and Event-Driven Scanning for Applications/Systems): Store your container images in Artifact Registry. Vulnerability scanning can be enabled to automatically scan images when they are pushed and to continually monitor them for newly discovered vulnerabilities.

- **Implementation**

- **Enable API:** Ensure the Container Analysis API and Artifact Registry API are enabled.

- **Store Images:** Push your container images to an Artifact Registry repository.

- **Enable Vulnerability Scanning**

- Vulnerability scanning is automatically enabled for images pushed to Artifact Registry in regions where Container Analysis is supported. You can check status and findings.
 - Navigate to **Artifact Registry > Repositories**, select your repository, then select an image. Vulnerability information will be displayed if available.
 - Alternatively, view findings directly in **Security > Container Analysis > Vulnerabilities**.

- **Scan on Push & Continuous Analysis**

- Images are typically scanned upon push.
 - Container Analysis also supports continuous analysis, where already-pushed images are re-scanned as new vulnerability

information is added to its database. This addresses the requirement to scan when new vulnerabilities are identified.

Additional Considerations

- **Defining "Periodically":** Your organization must define what "periodically" means for different types of scans based on risk, compliance requirements, and system criticality (e.g., daily for OS vulnerabilities via VM Manager, weekly for web applications via Web Security Scanner, continuous for cloud configurations via Security Health Analytics).
- **Authenticated Scans:** For Web Security Scanner, authenticated scans can provide deeper coverage. Configure this if your application supports Google Account authentication for test users.
- **Remediation Workflow:** Vulnerability scanning is only the first step. Establish a documented process for prioritizing, assigning, tracking, and verifying the remediation of identified vulnerabilities.
- **Third-Party Tools:** While this guidance focuses on Google Cloud resources, you may also use third-party vulnerability scanners. Security Command Center can ingest findings from some integrated third-party tools, providing a consolidated view.
- **Scope:** Ensure your scanning strategy covers all in-scope organizational systems and applications that process, store, or transmit CUI.
- **Documentation:** Maintain records of scan configurations, schedules, scan reports, identified vulnerabilities, and remediation activities for compliance and audit purposes.

Supplemental Guidance

- [Vulnerability findings | Security Command Center | Google Cloud](#)
- [Overview of Web Security Scanner | Security Command Center | Google Cloud](#)
- [Using Web Security Scanner | Security Command Center | Google Cloud](#)
- [About VM Manager | Compute Engine Documentation | Google Cloud](#)
- [View vulnerability reports | VM Manager | Google Cloud](#)
- [Artifact analysis and vulnerability scanning | Artifact Registry documentation | Google Cloud](#)

Control Domain	Risk Assessment
Control #	RA.L2-3.11.3
Control Description	Remediate vulnerabilities in accordance with risk assessments

Key Services	• SCC • VM Manager • Artifact Registry • VPC Firewall • IaC Tools	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Identifying vulnerabilities - Remediating vulnerabilities in accordance with risk assessments When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control: • Security Command Center (SCC) (for remediation recommendations and tracking finding states) • VM Manager (especially Patch Management for OS vulnerabilities) • Artifact Registry (for identifying fixed versions for container vulnerabilities) • VPC Firewall Rules / Network Analyzer / Firewall Insights (for remediating network configuration vulnerabilities) • Infrastructure as Code (IaC) Tools (e.g., Cloud Deployment Manager, Terraform, for applying configuration changes) A description of relevant features and implementation guidance is included below. Security Command Center (SCC): Use the detailed information and remediation steps provided within SCC findings to guide your remediation efforts. Monitor the status of findings in SCC to confirm if your actions have resolved the underlying issue. • Implementation ○ Review Remediation Steps 1. Navigate to Security > Security Command Center > Findings. 2. Click on a specific finding to view its details. Many findings (especially from Security Health Analytics) include a "Next steps" or "Remediation" section with specific guidance or links to documentation on how to fix the issue. ○ Track Finding State: After attempting remediation, observe if the finding's state changes to "INACTIVE" or is no longer reported by SCC in subsequent scans. This helps verify remediation. ○ Mute Findings (for Accepted Risks/False Positives): If a finding cannot be remediated immediately or is deemed an accepted risk (with proper			

authorization) or a false positive, use the "Mute finding" option. Provide a clear justification for muting. This is part of managing the remediation lifecycle.

- **Integrate with Ticketing Systems:** SCC can publish finding notifications to Pub/Sub. Configure a Pub/Sub subscription to forward these notifications to your ticketing system (e.g., Jira, ServiceNow) to create, assign, and track remediation tasks based on risk.

VM Manager (Patch Management for OS Vulnerabilities): Use VM Manager's Patch Management capabilities to apply OS patches to your Compute Engine instances in a controlled manner. This directly remediates many CVEs reported by the Vulnerability Reporting feature.

- **Implementation**

- **Review Patch Compliance and Available Patches**

1. Navigate to **Compute Engine > VM Manager > Patch compliance**. This dashboard shows which VMs have available patches.
2. Correlate this with findings from **Vulnerability reports** to identify which patches address specific CVEs.

- **Deploy Patches**

1. Go to **Compute Engine > VM Manager > Patch deployments**.
2. Click **Create patch deployment**.
3. Configure the deployment
 - Name and description.
 - **Target VMs** Select VMs by instance name, group, or zone, using instance filters (e.g., by OS, labels).
 - **Patch config** Specify patch types (e.g., security, critical). For Windows, select KB IDs if needed. For Linux, options include dist-upgrade.
 - **Schedule** Run on-demand ("Run now") or schedule for specific times/recurrence (e.g., during maintenance windows). Configure rollout options (zonal, max unavailable).
 - (Optional) Pre-patch and post-patch scripts.

- **Verify Patching:** After a patch deployment job completes, review its status and re-check Vulnerability Reports or Patch Compliance dashboards to confirm that relevant vulnerabilities are no longer reported or that systems are compliant.

Artifact Registry (Remediating Container Vulnerabilities): Review vulnerability findings for container images. The findings often indicate the vulnerable package and sometimes the version in which the vulnerability is fixed. Rebuild your container images using updated base

images or patched libraries, then push the new image to Artifact Registry, which will trigger a new scan.

- **Implementation**

- **Identify Fixed Versions:** Review vulnerability details in **Security > Container Analysis > Vulnerabilities** or within the Artifact Registry UI for a specific image. Look for information on patched versions of OS packages or application libraries.
- **Update Dockerfiles/Source Code:** Modify your Dockerfile to use a newer, patched base image, or update your application dependencies (e.g., requirements.txt, pom.xml) to fixed versions.
- **Rebuild and Re-Push Image:** Rebuild your container image and push it to Artifact Registry.
- **Verify Remediation:** Check the vulnerability scan results for the new image version in Artifact Registry or Container Analysis to confirm the vulnerability is no longer present. Update your deployments (e.g., GKE workloads) to use the newly remediated image.

VPC Firewall Rules / Network Analyzer / Firewall Insights: Remediating network vulnerabilities involves modifying firewall rules or other network configurations.

- **Guidance:** Based on findings from Firewall Insights (e.g., overly permissive rules) or Network Analyzer, update your VPC firewall rules or network configurations.

- **Implementation**

1. Navigate to **VPC network > Firewall**.
2. To remediate an overly permissive rule (e.g., 0.0.0.0/0 source range)
 - Select the rule. Click **Edit**.
 - Change the "Source filter" and "Source IPv4 ranges" to more specific, authorized IP ranges.
 - Adjust "Protocols and ports" to allow only necessary services. Click **Save**.

Infrastructure as Code (IaC) Tools: If your Google Cloud resources are managed by IaC (e.g., Terraform, Cloud Deployment Manager), remediate configuration-based vulnerabilities by updating the code and re-applying it.

- **Guidance:** Modify your IaC templates to reflect the secure configuration (e.g., correct IAM bindings, secure firewall rules, non-public bucket settings).
- **Implementation**
 1. Edit the relevant IaC configuration files.
 2. Use the IaC tool's "plan" or "preview" feature to review the proposed changes.

3. Apply the changes (e.g., terraform apply or gcloud deployment-manager deployments update). This ensures remediations are codified and consistently applied.

Additional Considerations

- **Risk-Based Prioritization:** Always use your risk assessment (Control [AC.L2-3.11.1](#)) to prioritize which vulnerabilities to remediate first. Focus on those that pose the greatest risk to CUI and organizational operations.
- **Change Management:** Integrate vulnerability remediation activities into your formal change management process, especially for production systems.
- **Testing Remediations:** Before applying remediations in production, test them in a non-production environment if possible to ensure they do not cause unintended operational impacts.
- **Compensating Controls:** If a vulnerability cannot be immediately remediated (e.g., no patch available, operational constraints), document this, implement compensating controls to reduce risk, and create a plan for eventual remediation. This should be part of a formal risk acceptance process.
- **Continuous Monitoring:** Remediation is not a one-time fix. Continuously monitor for new vulnerabilities and re-assess previously remediated ones to ensure they remain effective.

Supplemental Guidance

- [Security Command Center - Managing findings](#)
- [VM Manager - Applying OS patches](#)
- [Artifact Registry - Viewing and resolving vulnerabilities](#)
- [Modifying VPC firewall rules](#)
- [Infrastructure as Code with Google Cloud](#)

Control Domain	Security Assessment		
Control #	CA.L2-3.12.1		
Control Description	Periodically assess the security controls in organizational systems to determine if the controls are effective in their application		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer

Customer Implementation Description

Google Cloud Customers are responsible for:

- a. Defining the frequency of security control assessments
- b. Assessing security controls with the defined frequency to determine if the controls are effective in their application

Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.

Supplemental Guidance

- N/A

Control Domain	Security Assessment		
Control #	CA.L2-3.12.2		
Control Description	Develop and implement plans of action designed to correct deficiencies and reduce or eliminate vulnerabilities in organizational systems		
Key Services	N/A	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Identifying deficiencies and vulnerabilities to be addressed by the plan of action b. Developing a plan of action to correct the identified deficiencies and reduce or eliminate identified vulnerabilities c. Implementing the plan of action to correct the identified deficiencies and reduce or eliminate identified vulnerabilities <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			

- N/A

Control Domain	Security Assessment		
Control #	CA.L2-3.12.3		
Control Description	Monitor security controls on an ongoing basis to ensure the continued effectiveness of the controls		
Key Services	• SCC • Cloud Logging • Cloud Monitoring • IdAMRecommender • VM Manager • OS Configuration management • Network Analyzer and Firewall Insights	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Monitoring security controls on an ongoing basis to ensure the continued effectiveness of those controls. When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control: • Security Command Center (SCC) (especially Security Health Analytics, Event Threat Detection, VM/Container Threat Detection, Compliance Reporting, and Posture Management) • Cloud Logging (with Log-based Alerts via Cloud Monitoring) • Cloud Monitoring (Dashboards, Metric-based Alerts, Uptime Checks) • Identity and Access Management (IAM) Recommender • VM Manager (continuous Patch Compliance and Vulnerability Reporting) • OS Configuration management (continuous compliance monitoring for VMs) • Network Analyzer and Firewall Insights (periodic analysis providing ongoing awareness) A description of relevant features and implementation guidance on how it supports ongoing monitoring of security controls is included below.			

Security Command Center (SCC): Leverage SCC's various detectors and features to maintain continuous visibility into your security posture and the state of key security controls. Respond promptly to new findings.

- **Implementation**

- **Security Health Analytics:** This service continuously scans your Google Cloud environment for misconfigurations and vulnerabilities related to IAM, networking, storage, API keys, and more. These findings directly reflect the effectiveness of your configured controls. Regularly review new findings via the **Security Command Center > Findings** dashboard or set up notifications.
- **Event Threat Detection, VM Threat Detection, Container Threat Detection (SCC Premium/Enterprise):** These services provide near real-time detection of active threats. An alert from these services often indicates that a preventative control may have been bypassed or is ineffective, requiring immediate attention.
- **Compliance Reporting (SCC Premium):** For supported benchmarks (e.g., CIS, NIST), SCC continuously evaluates your resources against control objectives. Review these reports via **Security Command Center > Compliance** to monitor your ongoing compliance status, which reflects control effectiveness.
- **Posture Management (SCC Enterprise):** Allows you to define custom security postures (desired state of controls) and continuously monitor your environment for deviations. Create posture deployments to monitor specific controls critical for CUI protection.
- **Notifications:** Configure real-time notifications for new SCC findings by setting up continuous exports to Pub/Sub (**Security Command Center > Settings > Continuous Exports**) and then integrating Pub/Sub with your alerting or ticketing systems.

Cloud Logging (with Log-based Alerts via Cloud Monitoring) Create specific log-based alerts for activities that would indicate a security control failure, circumvention, or unexpected changes to critical security configurations.

- **Implementation**

- **Ensure Comprehensive Logging:** Verify that Admin Activity, System Event, Data Access (for CUI-related services), Policy Denied, VPC Flow Logs, and Firewall Rules Logging are enabled and being ingested into Cloud Logging.
- **Develop Log-Based Alerts:** Navigate to **Logging > Log-based Alerts** (or via Cloud Monitoring > Alerting). Create alerts for conditions such as
 - Critical IAM policy changes
protoPayload.methodName="SetIamPolicy" AND

protoPayload.serviceName="cloudresourcemanager.googleapis.com" on production CUI projects.

- Firewall rule modifications protoPayload.methodName "compute.firewalls.patch" OR protoPayload.methodName "compute.firewalls.insert"
- VPC Service Controls violations logName "cloudaudit.googleapis.com%2Fpolicy"
- Disabling of essential logging (e.g., audit log sink failures).
- Attempts to access sensitive data that are denied by IAM.
- **Configure Notification Channels:** Ensure these alerts notify the appropriate security personnel or teams immediately.

Cloud Monitoring (Dashboards, Metric-based Alerts, Uptime Checks): Create dashboards to visualize key security metrics and the status of critical controls. Set up alerts for deviations in these metrics.

- **Implementation**

- **Security Dashboards** Go to **Monitoring > Dashboards**. Create custom dashboards displaying
 - Number of critical/high SCC findings over time.
 - IAM role change frequency.
 - Firewall rule modification counts.
 - Rates of denied network connections to sensitive systems.
 - Health and performance of security infrastructure (e.g., custom security agents, Cloud IDS instances).
- **Metric-Based Alerts** Create alerts (**Monitoring > Alerting**) for
 - Unusual spikes in Data Access logs for sensitive CUI buckets/datasets.
 - Failure of critical security-related compute instances (e.g., a bastion host, a custom security logging VM).
 - Uptime check failures for critical applications handling CUI, which might indicate a broader issue affecting security.

Identity and Access Management (IAM) Recommender: Regularly review and act upon recommendations from IAM Recommender to ensure access controls (a key security control) remain effective and adhere to least privilege.

- **Implementation**

- Check for recommendations directly on the **IAM & Admin > IAM** page (recommendations often appear at the top or as icons next to principals).
- Also review **Security > Recommendations Hub**, filtering for IAM-related recommendations.

- The ongoing nature of these recommendations helps in continuously monitoring and improving IAM control effectiveness.

VM Manager (Continuous Patch Compliance and Vulnerability Reporting): Use VM Manager's dashboards as a continuous monitoring tool for the effectiveness of your vulnerability and patch management controls.

- **Implementation**

- Regularly (e.g., daily or weekly as part of operational checks) review **Compute Engine > VM Manager > Patch compliance and Vulnerability reports**.
- The data is updated frequently (typically daily for active VMs), providing an ongoing view of the control status. A decline in patch compliance or a rise in unaddressed critical vulnerabilities indicates a lapse in control effectiveness.

OS Configuration management: Define desired OS configurations (e.g., password complexity, software restrictions, security settings) as OS policies. Monitor the compliance status to ensure these controls remain effective.

- **Implementation**

- Create OS policy assignments via **Compute Engine > OS Configuration management** that define your security baselines for VMs handling CUI.
- The service continuously monitors and reports compliance. Regularly review the compliance dashboard to detect and address any drift from the desired secure state, thus monitoring the continued effectiveness of these OS-level controls.

Network Analyzer and Firewall Insights: Periodically run analyses or review insights, but consider their findings as part of an ongoing awareness of network control effectiveness. Firewall Insights, for example, can show how rules are being hit over time.

- **Implementation**

- Review **Network Intelligence > Network Analyzer** and **VPC network > Firewall Insights** reports. New insights may appear as your network or traffic patterns change, providing ongoing data points for control effectiveness.

Additional Considerations

- **Define Baselines:** For effective ongoing monitoring, establish clear baselines for what constitutes "normal" or "secure" for your controls and configurations.
- **Automation of Response:** Where feasible, automate responses to certain alerts that indicate control failures (e.g., automatically isolating a VM that triggers a critical threat detection finding, with appropriate safeguards).

- **Review and Tune Monitoring:** Regularly review the effectiveness of your monitoring setup itself. Are you getting too many false positives? Are there gaps in what you're monitoring? Adjust configurations, queries, and alert thresholds as needed.
- **Documentation:** Document your ongoing monitoring procedures, including tools used, alert configurations, review frequencies, and response actions for different types of control deviations.

Supplemental Guidance

- [Security Command Center - Overview](#)
- [Security Command Center - Continuous Exports](#)
- [Configure log-based alerting policies | Cloud Logging](#)
- [Create and manage custom dashboards | Cloud Monitoring](#)
- [Overview of role recommendations | Policy Intelligence | Google Cloud](#)
- [About VM Manager | Compute Engine Documentation | Google Cloud](#)
- [OS Configuration management Overview](#)

Control Domain	Security Assessment		
Control #	CA.L2-3.12.4		
Control Description	Develop, document, and periodically update system security plans that describe system boundaries, system environments of operation, how security requirements are implemented, and the relationships with or connections to other systems		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Developing a system security plan - Describing and documenting the system boundary in the system security plan - Describing and documenting the system environment of operation in the system security plan - Identifying the security requirements identified and approved by the designated authority as non-applicable			

- e. Describing and documenting the method of security requirement implementation in the system security plan
- f. Describing and documenting the relationship with or connection to other systems in the system security plan
- g. Defining the frequency to update the system security plan
- h. Updating the system security plan with the defined frequency

Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.

Supplemental Guidance

- N/A

Control Domain	System and Communications Protection		
Control #	SC.L1-3.13.1		
Control Description	Monitor, control, and protect organizational communications (i.e., information transmitted or received by organizational systems) at the external boundaries and key internal boundaries of the systems		
Key Services	• VPC Firewall Rules • Google Cloud Armor • Cloud NAT • VPC Service Controls • Cloud VPN / Cloud Interconnect • VPC Flow Logs • Firewall Rules Logging • Cloud IDS • SCC	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Defining the external system boundary b. Defining key internal system boundaries c. Monitoring communications at the external system boundary d. Monitoring communications at key internal boundaries e. Controlling communications at the external system boundary			

- f. Controlling communications at key internal boundaries
- g. Protecting communications at the external system boundary
- h. Protecting communications at key internal boundaries

When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control:

- VPC Firewall Rules
- Google Cloud Armor
- Cloud NAT
- VPC Service Controls
- Cloud VPN / Cloud Interconnect
- VPC Flow Logs
- Firewall Rules Logging
- Cloud IDS
- Security Command Center (SCC)

A description of relevant features and implementation guidance is included below.

VPC Firewall Rules: Control network traffic flow to and from VM instances and other resources within a VPC network based on IP addresses, protocols, ports, network tags, and service accounts. They are fundamental for defining and enforcing network boundaries.

- **Implementation**

1. Navigate to **VPC network > Firewall** in the Google Cloud Console.
2. Create rules following a default-deny principle, explicitly allowing only necessary traffic.
3. Define ingress rules to control incoming traffic (e.g., allow HTTPS from specific sources to web servers). Use specific source IP ranges, tags, or service accounts instead of 0.0.0.0/0 where possible.
4. Define egress rules to control outgoing traffic (e.g., allow specific VMs to reach external update servers on specific ports).
5. Use network tags or service accounts to apply rules to specific groups of VMs for microsegmentation (controlling internal boundaries).
6. Assign priorities to rules to manage precedence (lower numbers have higher priority).
7. Enable Firewall Rules Logging on critical allow and deny rules for monitoring.

Google Cloud Armor: Provide defense against DDoS attacks and application-layer attacks (like XSS, SQLi) for applications and services behind supported Google Cloud external load balancers. Acts as a web application firewall (WAF) at the network edge.

- **Implementation Guidance**

1. Navigate to **Network Security > Cloud Armor**.
2. Create a Security Policy.
3. Configure rules within the policy
 - Use preconfigured WAF rules (e.g., OWASP Top 10) to block common web attacks.
 - Define custom rules based on IP addresses/ranges (allow/deny lists), geo-location, request headers, etc.
 - Configure rate limiting rules to mitigate volumetric attacks.
 - Enable Adaptive Protection (requires Cloud Armor Enterprise) for automated DDoS detection and mitigation suggestions.
 - Set rule actions (Allow, Deny, Preview). Use Preview mode to test rules before enforcement.
4. Attach the security policy to one or more backend services associated with your external HTTP(S) load balancer.

Cloud NAT (Network Address Translation): Allow VM instances without external IP addresses to initiate outbound connections to the internet or other networks, while controlling and masking their private source IPs. Useful for managing egress traffic from private subnets.

- **Implementation**

1. Navigate to **Network services > Cloud NAT**.
2. Click **Create NAT gateway**.
3. Provide a name, select the VPC network and region.
4. Select a Cloud Router (or create one).
5. Choose **Source IP address translation options** Automatic (Google-managed IPs) or Manual (static IPs you reserve). Manual is useful for allow-listing source IPs with external services.
6. Choose **NAT mapping** Specify which subnet IP ranges should use the NAT gateway (e.g., all ranges, specific primary/secondary ranges).
7. Configure logging options as needed.
8. Click **Create**.

VPC Service Controls: Create security perimeters around Google-managed services (like Cloud Storage, BigQuery) to prevent data exfiltration by controlling which identities and networks can access services within the perimeter and how data can move across the perimeter boundary.

- **Implementation**

1. Navigate to **Security > VPC Service Controls**.
2. Select or create an Access Policy for your Organization.

3. Click **New Perimeter**.
4. Give the perimeter a name and choose a configuration type (**Enforced** or **Dry Run** - start with Dry Run).
5. Add the Google Cloud **Projects** to be included within the perimeter.
6. Select the **Restricted Services** (e.g., storage.googleapis.com, bigquery.googleapis.com) that will be protected by the perimeter.
7. Configure **VPC Accessible Services** to control which APIs can be called from *within* the perimeter networks.
8. (Optional) Configure **Ingress/Egress Policies** and **Access Levels** (defined in Access Context Manager) for fine-grained control over traffic crossing the perimeter boundary based on identity, IP, or device context.
9. Click **Create Perimeter**. Monitor Dry Run logs before moving to Enforced mode.

Cloud VPN / Cloud Interconnect: Establish secure, private connections between your on-premises network or other clouds and your Google Cloud VPC network. These act as controlled external boundaries.

- **Implementation (Cloud VPN)**

1. Navigate to **Hybrid Connectivity > VPN**.
2. Create a HA (High Availability) VPN Gateway in Google Cloud.
3. Configure tunnel(s) connecting to your peer (on-premises) VPN gateway, including IKE versions, pre-shared keys, and routing options (dynamic via BGP with Cloud Router, or static/policy-based).
4. Configure corresponding firewall rules on both Google Cloud (VPC Firewall) and your on-premises firewall to control traffic flow across the tunnel.

- **Implementation (Cloud Interconnect)**

1. Requires physical connections (Dedicated Interconnect) or partner connections (Partner Interconnect).
2. Order the connection via the Google Cloud Console (**Hybrid Connectivity > Interconnect**).
3. Configure VLAN attachments and BGP sessions (via Cloud Router) to establish routing between your VPC and the external network.
4. Apply VPC Firewall rules to control traffic entering/leaving your VPC via the Interconnect attachment.

VPC Flow Logs: Capture samples of network IP traffic flows sent from and received by VM instances. Useful for network monitoring, forensics, and understanding traffic patterns.

- **Implementation**

1. Navigate to **VPC network > Subnets**.

2. Select the subnet(s) you want to monitor.
3. Click **Edit**.
4. Set **Flow logs** to **On**.
5. Configure **Aggregation interval**, **Sample rate**, and **Metadata** inclusion. *Note Higher sampling rates and metadata inclusion generate more logs.*
6. Click **Save**. Logs will be sent to Cloud Logging.

Firewall Rules Logging: Record when specific firewall rules allow or deny traffic, including connection details (IPs, ports, protocol). Helps audit rule effectiveness and troubleshoot connectivity.

- **Implementation**

1. Navigate to **VPC network > Firewall**.
2. Select the firewall rule(s) you want to enable logging for.
3. Click **Edit**.
4. Set **Logs** to **On**.
5. Click **Save**. Logs are sent to Cloud Logging.

Cloud IDS (Intrusion Detection System): Managed network-based threat detection service that monitors ingress, egress, and intra-VPC traffic for malicious activity using signatures and threat intelligence.

- **Implementation**

1. Navigate to **Network Security > Cloud IDS**.
2. Click **Create Endpoint**.
3. Select the network, region, zone, and severity level for alerts.
4. Configure Private Service Access if not already done.
5. After endpoint creation, configure **Packet Mirroring** policies (**Compute Engine > Packet Mirroring**) to send traffic from desired subnets, tags, or instances to the created Cloud IDS endpoint for inspection.
6. Monitor findings in Cloud Logging or Security Command Center.

Security Command Center (SCC): Provide a centralized platform for security monitoring, threat detection, and posture management across Google Cloud. Aggregates findings from various sources, including Cloud IDS, Firewall Rules Logging insights, VPC Service Controls violations, etc.

- **Implementation**

1. Navigate to **Security > Security Command Center**.
2. Activate SCC for your organization or project (Standard or Premium tier).
3. Ensure relevant services (Cloud IDS, Firewall logging, etc.) are enabled and configured to send findings/logs.

4. Use the SCC dashboard to review findings, assets, vulnerabilities, and compliance status related to network security.
5. Configure notifications for critical findings.

Additional Considerations

- **Defense-in-Depth:** Employ multiple layers of controls (e.g., Cloud Armor + VPC Firewall + VPC Service Controls).
- **Least Privilege Network Access:** Configure firewall rules to allow only the minimum necessary communication between resources.
- **Regular Review:** Periodically review firewall rules, VPC SC policies, and logs to ensure they remain effective and aligned with current requirements.
- **Automation:** Consider automating responses to certain alerts or findings using tools like Cloud Functions or Security Orchestration, Automation, and Response (SOAR) platforms integrated with SCC.

Supplemental Guidance

- [VPC Firewall Overview](#)
- [Using VPC Firewall Rules](#)
- [Google Cloud Armor Overview](#)
- [Google Cloud Armor Product Details](#)
- [Cloud NAT Overview](#)
- [VPC Service Controls Overview](#)
- [VPC Service Controls Detailed Overview](#)
- [VPC Flow Logs Overview](#)
- [Firewall Rules Logging Overview](#)
- [Using Firewall Rules Logging](#)
- [Cloud IDS Overview](#)
- [Cloud IDS Product Page](#)
- [Security Command Center Overview](#)
- [Security Command Center Documentation](#)

Control Domain	System and Communications Protection
Control #	SC.L2-3.13.2
Control Description	Employ architectural designs, software development techniques, and systems engineering principles that promote information security in system development

Key Services	• VPC Networking Features • IAM • Organization Policy Service • Secure Software Development and Deployment • Defense in depth approach • Managed Services	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Identifying architectural designs that promote effective information security - Identifying software development techniques that promote effective information security - Identifying systems engineering principles that promote effective information security - Employing identified architectural designs that promote effective information security - Employing identified software development techniques that promote effective information security - Employing identified systems engineering principles that promote effective information security When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Secure Network Architecture (VPC) Networking Features • Identity and Access Management (IAM) • Organization Policy Service • Secure Software Development and Deployment • Defense in depth approach • Managed Services A description of relevant features and implementation guidance is included below. Secure Network Architecture (VPC): Design VPC networks with security in mind from the start. • Implementation			

1. Avoid using the default VPC network for production workloads; create custom VPCs.
2. Segment networks using VPCs and subnets (e.g., separate environments like development, staging, production; separate application tiers).
3. Use private IP addresses (RFC 1918) for internal resources and control internet access via Cloud NAT or secure proxies.
4. Implement fine-grained VPC Firewall Rules based on the principle of least privilege (default deny, allow specific protocols/ports/sources/destinations using tags or service accounts).
5. Utilize VPC Service Controls to create perimeters around sensitive data services.
6. Follow guidance from **Best practices and reference architectures for VPC design**.

Identity and Access Management (IAM): Implement the principle of least privilege rigorously.

- **Implementation**

1. Avoid using primitive roles (Owner, Editor, Viewer) in production. Grant predefined roles that meet specific needs.
2. Create and use custom IAM roles when predefined roles are too permissive.
3. Grant roles to Google Groups instead of individual users where practical for easier management.
4. Use IAM Conditions for attribute-based or temporary access.
5. Regularly review IAM policies using tools like Policy Analyzer and apply IAM Recommender suggestions.
6. Follow **IAM best practices**.

Organization Policy Service: Enforce preventative guardrails across the organization, folders, or projects.

- **Implementation**

1. Navigate to **IAM & Admin > Organization Policies**.
2. Review available constraints relevant to security architecture (e.g., `compute.vmExternalIpAccess` to restrict external IPs, `iam.disableServiceAccountKeyCreation` to prevent key downloads, `compute.trustedImageProjects` to restrict VM images, `constraints/gcp.resourceLocations` to restrict physical locations).
3. Define and apply policies to enforce desired constraints at the appropriate level of the resource hierarchy. Test policies before enforcing them broadly.

Secure Software Development and Deployment: Integrate security into the development pipeline.

- **Implementation**

1. Use Cloud Build to create automated, repeatable build processes defined in configuration files.
2. Enable Cloud Build to generate build provenance (SLSA level 3 evidence) to verify build integrity.
3. Store container images and packages in Artifact Registry.
4. Enable vulnerability scanning in Artifact Analysis (via the Container Scanning API) to automatically scan artifacts stored in Artifact Registry.
5. Implement Binary Authorization policies to ensure that only containers signed by trusted attestors (which can be linked to build provenance or vulnerability scan results) can be deployed to GKE or Cloud Run.

Defense-in-Depth Approach: Layer multiple security controls.

- **Implementation**

1. Combine network controls (VPC Firewall, Cloud Armor), identity controls (IAM, IAP), data controls (VPC Service Controls, Encryption), and monitoring/detection (Logging, Cloud IDS, SCC) rather than relying on a single defense mechanism.

Managed Services: Leverage Google-managed services where suitable.

- **Implementation**

1. Using services like Cloud SQL, GKE Autopilot, or Cloud Run shifts responsibility for securing and patching the underlying infrastructure and parts of the operating system/runtime to Google, benefiting from their secure engineering practices.

Additional Considerations

- **Security Reviews:** Incorporate security architecture reviews into the system design process.
- **Threat Modeling:** Identify potential threats and design controls accordingly.
- **Secure Coding:** Train developers on secure coding practices (e.g., OWASP Top 10 awareness).
- **Infrastructure as Code (IaC):** Use tools like Cloud Deployment Manager or Terraform to define infrastructure, enabling version control, review, and automated deployment, which supports secure engineering principles.

Supplemental Guidance

- [Google Cloud Security Foundations Guide](#)
- [Best practices and reference architectures for VPC design](#)
- [General Hybrid/Multicloud Networking Best Practices \(includes VPC design\)](#)
- [Use IAM securely](#)
- [Best practices for using service accounts](#)
- [Organization Policy Service Introduction](#)
- [Secure Cloud Build Security Overview](#)
- [Artifact Analysis Vulnerability Scanning](#)
- [Binary Authorization Overview](#)
- [Google Security Whitepaper](#)
- [Defense in Depth eBook](#)

Control Domain	System and Communications Protection								
Control #	SC.L2-3.13.3								
Control Description	Separate user functionality from system management functionality								
Key Services	• IAM • Cloud Identity / Google Workspace	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Identifying user functionality b. Identifying system management functionality c. Separating user functionality from system management functionality When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Identity and Access Management (IAM) • Cloud Identity / Google Workspace (Account Separation & OUs) A description of relevant features and implementation guidance is included below. Identity and Access Management (IAM): Use IAM roles to grant permissions, strictly adhering to the principle of least privilege. • Assign Least Privilege Roles									

1. In the Google Cloud Console, navigate to **IAM & Admin > IAM**.
2. Select the resource level (Organization, Folder, or Project).
3. Identify the principal (user account, group, or service account) to grant access. Use dedicated administrator accounts for management roles and standard accounts for user roles.
4. Click **Grant Access**.
5. In the **Assign roles** section, select the most specific predefined role(s) required for the function. Avoid basic roles like Owner or Editor for routine tasks. Examples
 - *User Functionality* Assign roles like roles/compute.viewer, roles/storage.objectViewer, or custom roles with application-specific, non-administrative permissions.
 - *System Management* Assign roles like roles/compute.admin, roles/iam.organizationAdmin, roles/resourcemanager.projectIamAdmin, etc., *only* to dedicated administrator accounts.
6. Create custom IAM roles (**IAM & Admin > Roles**) if predefined roles are too broad, bundling only the necessary permissions.
7. Click **Save**.
- **Regularly Audit Roles**
 1. Use **IAM & Admin > Policy Analyzer** to understand who has access to what.
 2. Review **IAM Recommender** suggestions (**IAM & Admin > IAM** or **Security > Recommendations Hub**) to identify and remove excessive permissions.

Cloud Identity / Google Workspace (Account Separation & OUs): Manage user accounts and organizational structure to support the separation of functions.

- **Use Separate Administrative Accounts**
 1. Create distinct Google accounts for users who perform administrative duties. One account (user@example.com) should be used for daily non-privileged tasks (email, docs), and a separate account (admin-user@example.com) should be used exclusively for administrative access to Google Cloud or Workspace.
 2. Ensure these admin accounts are clearly identifiable.
- **Utilize Organizational Units (OUs) for Policy Separation (Recommended)**
 1. In the Google Admin console (admin.google.com), navigate to **Directory > Organizational units**.
 2. Create separate OUs for standard users and administrators (e.g., Users and Administrators).
 3. Move the respective user accounts into the appropriate OUs.

4. Navigate to security settings (e.g., **Security > Authentication > 2-step verification**, **Security > Access and data control > Google Session control**) and apply stricter policies to the Administrators OU compared to the Users OU (e.g., mandatory security keys for MFA, shorter session timeouts).

Additional Considerations

- **Operational Discipline:** Administrators must avoid using their privileged accounts for routine, non-administrative tasks like checking email or browsing the web. Privileged accounts should only be used when performing administrative actions.
- **Role Documentation:** Clearly document the purpose and permissions associated with custom IAM roles.
- **Just-in-Time Access:** Explore using IAM Conditions for time-bound access or investigate future Google Cloud Privileged Access Manager capabilities to grant temporary elevated privileges instead of persistent administrative roles where feasible.

Supplemental Guidance

- [Use IAM securely \(includes least privilege\)](#)
- [Understanding Roles \(Basic, Predefined, Custom\)](#)
- [IAM Roles Overview](#)
- [Security best practices for administrator accounts \(Workspace\)](#)
- [Super administrator account best practices \(Resource Manager\)](#)
- [How the organizational structure works](#)

Control Domain	System and Communications Protection		
Control #	SC.L2-3.13.4		
Control Description	Prevent unauthorized and unintended information transfer via shared system resources.		
Key Services	• IAM • VPC Service Controls • Organization Policy Service • VPC Firewall Rules • Cloud DLP	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for:			

- a. Preventing unauthorized and unintended information transfer via shared system resources

When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control:

- Identity and Access Management (IAM)
- VPC Service Controls
- Organization Policy Service
- VPC Firewall Rules
- Cloud Data Loss Prevention (DLP)

A description of relevant features and implementation guidance is included below.

Identity and Access Management (IAM): Use IAM to define granular access control over who (users, groups, service accounts) can perform what actions on which Google Cloud resources. This is foundational to preventing unauthorized access that could lead to unintended data transfers.

- **Grant Roles**
 1. Navigate to **IAM & Admin > IAM** in the Google Cloud Console.
 2. Select the appropriate resource scope (Organization, Folder, or Project).
 3. Click **Grant Access**.
 4. In the **New principals** field, add the user, group, or service account.
 5. From the **Assign roles** dropdown, select the most restrictive predefined or custom role that grants only the necessary permissions (principle of least privilege). Avoid overly broad basic roles like Owner or Editor for routine tasks.
 6. (Optional) Add an IAM Condition for context-aware access.
 7. Click **Save**.
- **Audit IAM Policies**
 1. Regularly review IAM policies using **IAM & Admin > IAM**.
 2. Utilize the **Policy Analyzer** (IAM & Admin > Policy Analyzer) to understand who has what access to which resources.
 3. Review IAM recommendations in **Security Command Center > Recommendations** or directly in IAM for over-privileged accounts.

VPC Service Controls: Use VPC Service Controls to create security perimeters around Google-managed services to protect against data exfiltration. Data within a perimeter cannot leave the perimeter except through explicitly configured channels.

- **Create an Access Policy (if one doesn't exist)**
 1. Navigate to **Security > VPC Service Controls**.

2. An organization access policy is typically created by default.
- **Create a Service Perimeter**
 1. Click **New Perimeter**.
 2. Enter a **Perimeter Name** and optional description.
 3. Choose **Regular** perimeter type for most use cases.
 4. Select **Enforced** mode to actively block violations, or **Dry Run** mode to log violations without blocking (recommended for initial setup and testing).
 5. **Add Projects** Click **Add Projects** and select the projects containing sensitive data and services you want to protect.
 6. **Restrict Service** Click **Add Services** and select the Google APIs/services (e.g., storage.googleapis.com, bigquery.googleapis.com) to be secured within the perimeter.
 7. (Optional) **VPC Accessible Services** Configure what services can be accessed from within the perimeter's VPC networks. Restrict this to only necessary services.
 8. (Optional) **Ingress/Egress Policies & Access Levels** Define fine-grained rules for allowing specific access to or from the perimeter based on identity, IP, or device context (using Access Context Manager). This is crucial for allowing legitimate data transfers while blocking others.
 9. Click **Create Perimeter**.
- **Monitor** If using Dry Run mode, monitor Cloud Audit Logs (filter for VpcServiceControlsAuditData) to identify potential violations before switching to Enforced mode.

Organization Policy Service: Use Organization Policies to enforce organization-wide constraints on how resources can be configured and used, helping prevent configurations that could lead to unintended data transfer.

- **View/Edit Policies**
 1. Navigate to **IAM & Admin > Organization Policies**.
 2. Select your Organization, Folder, or Project.
 3. Find a relevant policy constraint from the list, for example:
 - **Domain Restricted Sharing**
(constraints/iam.allowedPolicyMemberDomains) Restricts which customer identity domains can be added to IAM policies.
 - **Define allowed external IPs for VM instances**
(constraints/compute.vmExternallpAccess) Prevents or limits the assignment of public IPs to VMs.
 - **Restrict Public Access for Cloud Storage**
(constraints/storage.publicAccessPrevention) Enforces that new and existing buckets prevent public access.

- **Disable Cloud DNS Peering, Cross-Project Binding, Service Directory registration for private zones**

(constraints/dns.restrictPrivateZoneSharing) Limits how private DNS zones can be shared.

- **Customize a Policy**

1. Select the policy constraint.
2. Click **Edit Policy**.
3. Choose **Customize**.
4. Select the **Enforcement** type (e.g., Replace, Merge with parent).
5. Configure the specific rules for the constraint (e.g., for iam.allowedPolicyMemberDomains, add the Google Cloud Customer IDs of allowed domains).
6. Click **Save**.

VPC Firewall Rules: Use VPC firewall rules to control network traffic to and from your VM instances and other resources within a VPC network. This helps prevent unauthorized network connections that could be used for data transfer.

- **Create a Firewall Rule**

1. Navigate to **VPC network > Firewall**.
2. Click **Create Firewall Rule**.
3. **Name** Provide a descriptive name.
4. **Network** Select the VPC network.
5. **Priority** Set a priority (lower numbers are higher priority).
6. **Direction of traffic** Choose **Ingress** (incoming) or **Egress** (outgoing).
7. **Action on match** Select **Allow** or **Deny**. Default-deny is a best practice for ingress traffic.
8. **Targets** Specify which instances the rule applies to (e.g., all instances, specific target tags, or service accounts).
9. **Source/Destination filter** For Ingress, define allowed **Source IPv4 ranges** (e.g., internal subnets, specific IP addresses). For Egress, define **Destination IPv4 ranges**. Be as specific as possible. Avoid 0.0.0.0/0 unless absolutely necessary and carefully considered.
10. **Protocols and ports** Specify protocols (e.g., tcp, udp) and port numbers (e.g., tcp:443).
11. (Optional) **Enable logging** for the firewall rule to audit connections.
12. Click **Create**.

- **Review Implied Rules:** Remember that each VPC network has two implied firewall rules (allow all egress, deny all ingress) at the lowest priority. Your rules override these.

Cloud Data Loss Prevention (DLP): Use Cloud DLP to discover, classify, and protect sensitive data within your Google Cloud environment. While not a direct transfer prevention tool, it identifies where sensitive data resides, enabling better application of other controls like IAM or VPC Service Controls.

- **Create a DLP Job or Template:**
 1. Navigate to **Security > Data Loss Prevention**.
 2. Go to **Create > Job or Job trigger** to scan existing data or **Configuration > Templates** to define reusable scan configurations.
 3. **Choose Input Data** Specify the Cloud Storage bucket, BigQuery table, or Datastore kind to scan.
 4. **Configure Detection**
 - Select or create **InfoType detectors** (e.g., PII, financial data, credentials).
 - Set **Likelihood levels** to adjust sensitivity.
 5. **Add Actions**
 - **Publish to Security Command Center** Send findings to Security Command Center.
 - **Publish to Pub/Sub** Trigger automated workflows.
 - **Save to BigQuery** Store findings for analysis.
 - **De-identify data** Create a de-identified copy (for specific use cases).
 6. **Schedule (for Job Triggers)** Configure how often the scan should run.
 7. Click **Create**.
- **Review Findings:** Analyze DLP findings to understand where sensitive data is located and to inform your data protection strategies.

Additional Considerations

- **Confidential Computing:** For highly sensitive workloads running on Compute Engine, use Confidential VMs. Confidential Computing encrypts data in-use (in memory) while it is being processed, providing an additional layer of protection on shared compute resources. This is enabled via a checkbox during VM creation.
- **Principle of Least Privilege:** Consistently apply the principle of least privilege for all IAM roles and firewall rules. Only grant permissions and network access that are strictly necessary.
- **Regular Audits:** Regularly audit IAM policies, firewall rules, Organization Policies, and VPC Service Controls configurations to ensure they remain aligned with security requirements and haven't been inadvertently changed.
- **Logging and Monitoring:** Ensure comprehensive logging is enabled (Cloud Audit Logs, VPC Flow Logs, Firewall Rules Logging) and integrate with Security Command Center or a SIEM to detect and alert on suspicious activities or policy violations that might indicate unauthorized data transfer attempts.

- **Data Classification:** Implement a data classification scheme to identify the sensitivity of your data. This will help determine the appropriate level of protection needed for different datasets and guide the configuration of controls like VPC Service Controls and DLP.
- **Secure Data Deletion:** When resources are de-provisioned, ensure that data is securely deleted according to your organization's policies to prevent unintended information recovery or transfer from decommissioned shared resources. Google Cloud provides mechanisms for data deletion.

Supplemental Guidance

- [IAM overview | IAM Documentation | Google Cloud](#)
- [Overview of VPC Service Controls | Google Cloud](#)
- [Create a service perimeter | VPC Service Controls | Google Cloud](#)
- [Introduction to the Organization Policy Service | Resource Manager Documentation | Google Cloud](#)
- [Using constraints | Resource Manager Documentation | Google Cloud](#)
- [VPC firewall rules | Cloud NGFW](#)
- [Sensitive Data Protection Documentation | Google Cloud](#)
- [Confidential Computing | Google Cloud](#)
- [Well-Architected Framework Security, privacy, and compliance pillar | Cloud Architecture Center](#)

Control Domain	System and Communications Protection		
Control #	SC.L1-3.13.5		
Control Description	Implement subnetworks for publicly accessible system components that are physically or logically separated from internal networks		
Key Services	• Google Cloud Default Encryption • Enforcing HTTPS • Securing Hybrid Connectivity	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Identifying publicly accessible system components - Physically or logically separating subnetworks for publicly accessible system components from internal networks			

When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control:

- Default Encryption within Google Cloud
- Enforcing HTTPS for Applications (Cloud Load Balancing)
- Enforcing HTTPS for App Engine
- Securing Hybrid Connectivity (Cloud VPN / Interconnect)

A description of relevant features and implementation guidance is included below.

Default Encryption within Google Cloud: Google Cloud automatically encrypts traffic at multiple layers when it moves outside physical boundaries controlled by Google (e.g., between data centers). Communication with most Google Cloud APIs and services via the console or standard endpoints defaults to using HTTPS/TLS.

- **Platform Behavior:** Customers benefit from Google's default encryption-in-transit measures for traffic flowing between Google data centers and for standard API/service access. No specific customer configuration is typically required for this default protection.

Enforcing HTTPS for Applications (Cloud Load Balancing): Ensure traffic between end-users and your applications hosted on Google Cloud is encrypted using HTTPS/TLS.

- **Implementation**
 1. Use an External HTTP(S) Load Balancer, Internal HTTP(S) Load Balancer, SSL Proxy Load Balancer, or TCP Proxy Load Balancer.
 2. Navigate to **Network Services > Load balancing**. Select or create a load balancer.
 3. In the **Frontend configuration**
 - Select **HTTPS** (or SSL/TLS depending on LB type) as the protocol.
 - Under **Certificates**, click **Create a new certificate** or select existing ones.
 - Choose **Google-managed certificate** (recommended for public-facing domains) or **Upload my own certificate** (self-managed).
 - If Google-managed, provide the domain name(s). Google will provision and renew the certificate automatically.
 - If self-managed, upload your certificate file, private key file, and optionally the certificate chain.
 - (Optional but recommended) Configure an SSL Policy to control the TLS versions and cipher suites allowed.

4. Configure backend services. Consider enabling encryption between the load balancer and backends if required.
5. Ensure applications are configured to handle HTTPS traffic and potentially redirect HTTP requests to HTTPS.

Enforcing HTTPS for App Engine: App Engine provides managed SSL certificates for custom domains.

- **Implementation**

1. Map your custom domain to your App Engine application (**App Engine > Settings > Custom Domains**).
2. Update your DNS records as instructed.
3. App Engine automatically provisions and renews Google-managed SSL certificates for your mapped custom domains. HTTPS will be enabled by default.
4. Optionally, configure secure always in your app.yaml (standard environment) or application code (flexible environment) to redirect all traffic to HTTPS.

Securing Hybrid Connectivity (Cloud VPN / Interconnect): Encrypt traffic flowing between your Google Cloud VPC and external networks.

- **Cloud VPN**

1. Navigate to **Hybrid Connectivity > VPN**.
2. Create an HA VPN gateway and configure tunnels. Cloud VPN uses the IPsec protocol suite to encrypt tunnel traffic by default.
3. Configure corresponding firewall rules to allow traffic through the tunnel.

- **Cloud Interconnect**

1. Cloud Interconnect provides a private connection, but traffic is not encrypted by default over the attachment.
2. To encrypt traffic over Interconnect, deploy **HA VPN over Cloud Interconnect**. This configures IPsec tunnels that run over the private Interconnect VLAN attachments. Follow the steps for creating HA VPN, but specify the Interconnect attachments as the target interfaces.
3. Alternatively, for Dedicated Interconnect, configure **MACsec for Cloud Interconnect** for link-layer encryption if supported by your on-premises equipment and connection speed.

Additional Considerations

- **Certificate Management:** If using self-managed certificates, ensure processes are in place for timely renewal and secure key management. Google-managed certificates simplify this.

- **Protocol and Cipher Strength:** Use strong, current TLS versions (e.g., TLS 1.2, 1.3) and cipher suites. Configure SSL policies on load balancers to enforce minimum standards.
- **End-to-End Encryption:** While TLS to the load balancer or VPN encryption protects traffic up to the Google Cloud boundary, consider application-level encryption if end-to-end protection (client-to-application-logic) is required for highly sensitive data.

Supplemental Guidance

- [Encryption in transit for Google Cloud](#)
- [Encryption For Cloud Security](#)
- [SSL certificates overview](#)
- [Use Google-managed SSL certificates](#)
- [Use self-managed SSL certificates](#)
- [Securing custom domains with App Engine HTTPS/SSL](#)
- [Cloud VPN overview](#)
- [Cloud Interconnect overview](#)
- [VPN over Cloud Interconnect](#)
- [MACsec for Cloud Interconnect](#)

Control Domain	System and Communications Protection								
Control #	SC.L2-3.13.6								
Control Description	Deny network communications traffic by default and allow network communications traffic by exception (i.e., deny all, permit by exception)								
Key Services	• Google Cloud Networking Services • IaC • Organization Policy Service • SCC	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Denying network communications traffic by default b. Allowing network communications traffic by exception When configured correctly, the following feature(s) in Google Cloud Console may be used to support this control:									

- Networking Services
- Infrastructure as Code (IaC)
- Organizational Policy Service
- Security Command Center (SCC)

A description of relevant features and implementation guidance is included below.

Networking Services: Define the required security settings for network components based on risk and compliance needs.

- **Implementation**

1. Reference established benchmarks like the CIS Google Cloud Platform Foundation Benchmark.
2. Consult Google Cloud security best practices documentation (e.g., Security Foundations Guide, VPC design best practices).
3. Define settings for
 - **VPC Networks** Use custom mode VPCs, avoid default networks, plan IP ranges carefully.
 - **Firewall Rules** Default-deny ingress/egress, allow specific needed traffic based on tags/service accounts, enable logging for critical rules.
 - **Load Balancers** Require HTTPS, use strong SSL policies (TLS 1.2+), enable access logging, integrate with Cloud Armor.
 - **Cloud VPN/Interconnect** Use strong authentication, enable logging, configure appropriate routing, encrypt traffic (IPsec for VPN, HA VPN/MACsec for Interconnect).
 - **Cloud NAT** Use manual IP allocation if needed for allow-listing, enable logging.
 - **DNS** Use Cloud DNS with DNSSEC enabled for public zones.
4. Document these baseline configurations.

Infrastructure as Code (IaC): Use IaC tools to automate the deployment and management of network resources according to the established baselines.

- **Implementation**

1. Use Cloud Deployment Manager (Google Cloud native) or Terraform (popular third-party) to write configuration templates (YAML for Deployment Manager, HCL for Terraform).
2. Define VPCs, subnets, firewall rules, load balancers, NAT gateways, etc., in code, reflecting the baseline settings.

3. Store these configuration files in a version control system (like Cloud Source Repositories, GitHub, GitLab) to track changes, facilitate review, and enable rollbacks.
4. Integrate IaC deployments into CI/CD pipelines (using Cloud Build or other tools) for automated and consistent application of configurations.

Organization Policy Service: Use Organization Policies to prevent configurations that deviate from critical security requirements.

- **Implementation**

1. Navigate to **IAM & Admin > Organization Policies** in the Google Cloud Console.
2. Identify constraints relevant to communication system configurations (e.g., `compute.vmExternalIpAccess`, `compute.restrictXpnProjectLienRemoval`, `compute.restrictSharedVpcSubnetworks`, `compute.skipDefaultNetworkCreation`, `compute.restrictFirewallCreationForTypes`, `iam.disableServiceAccountKeyCreation`).
3. Configure policies at the organization, folder, or project level to enforce these constraints (e.g., Deny external IPs on most projects, Allow only specific types of firewall rules).
4. Use policy conditions (e.g., based on tags) for more granular enforcement if needed.

Security Command Center (SCC): Use Security Health Analytics within SCC Premium to detect misconfigurations related to network security.

- **Implementation**

1. Ensure SCC Premium is activated and Security Health Analytics is enabled.
2. Navigate to **Security > Security Command Center > Findings**.
3. Filter findings by category or detector related to networking. Security Health Analytics includes detectors for common network misconfigurations, such as
 - Open firewall ports (SSH, RDP, databases, etc.)
 - Open RDP/SSH access from the internet
 - Default network usage
 - Disabled VPC Flow Logs or Firewall Rules Logging
 - Legacy network usage
 - Public IP addresses on instances
4. Review findings regularly and remediate misconfigurations to align with established baselines.
5. Configure notifications for high-severity or critical network-related findings.

Supplemental Guidance

- [Cloud Deployment Manager Overview](#)
- [Infrastructure as Code on Google Cloud \(mentions Terraform, Infra Manager\)](#)
- [Organizational Policy Service Introduction](#)
- [Creating and managing policies](#)
- [Overview of Security Health Analytics](#)
- [Security Health Analytics Findings \(includes network detectors\)](#)
- [Google Cloud Security Foundations Guide](#)

Control Domain	System and Communications Protection								
Control #	SC.L2-3.13.7								
Control Description	Prevent remote devices from simultaneously establishing non-remote connections with organizational systems and communicating via some other connection to resources in external networks (i.e., split tunneling)								
Key Services	- IAP / Chrome Enterprise Premium - VPN Policies	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☐</td><td>Shared</td></tr><tr><td>☒</td><td>Customer</td></tr></table>	☐	Google	☐	Shared	☒	Customer
☐	Google								
☐	Shared								
☒	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Preventing remote devices from simultaneously establishing non-remote connections with the system and communicating via some other connection to resources in external networks (i.e., split tunneling). Note: Preventing split tunneling is typically configured on the client-side VPN software or the VPN gateway that remote users connect to, or managed via endpoint policies. Google Cloud VPN itself primarily handles the site-to-site tunnel connection. A recommended Google Cloud approach to mitigate risks associated with VPNs and split tunneling is to adopt a Zero Trust model using IAP. When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control: - Identity-Aware Proxy (IAP) / Chrome Enterprise Premium (Alternative Approach) - Google Admin Console VPN Policies (Managed Endpoints)									

A description of relevant features and implementation guidance is included below.

Identity-Aware Proxy (IAP) / Chrome Enterprise Premium (Alternative Approach):

Implement IAP to provide secure, context-aware access to specific web applications, VMs (via SSH/RDP), and GKE applications without needing a traditional VPN, thus avoiding split tunneling issues for those resources.

- **Implementation**

1. Identify applications and resources (Compute Engine VMs, App Engine apps, GKE services via Load Balancer) that require secure remote access.
2. Navigate to **Security > Identity-Aware Proxy** in the Google Cloud Console.
3. Enable IAP for the identified resources (requires appropriate load balancer setup for web apps, or specific firewall rules for TCP forwarding).
4. Configure IAM permissions by granting the appropriate IAP roles (e.g., roles/iap.httpsResourceAccessor for web apps, roles/iap.tunnelResourceAccessor for SSH/RDP) to authorized users or groups for each resource.
5. (Optional, for enhanced security) Use Access Context Manager (**Security > Access Context Manager**) to create access levels based on IP address, device posture (requires Endpoint Verification), identity attributes, etc., and apply these access levels to IAP policies to enforce context-aware access.
6. Users access protected resources via their standard browser (for web apps) or specific gcloud commands (gcloud compute ssh --tunnel-through-iap) which are proxied securely through IAP based on their identity and context, without a full VPN tunnel.

Google Admin Console VPN Policies (Managed Endpoints): For ChromeOS devices managed through Google Workspace, administrators can configure specific VPN settings, potentially including split tunnel behavior depending on the VPN type.

- **Implementation**

1. Sign in to the Google Admin console (admin.google.com).
2. Navigate to **Devices > Networks > VPN**.
3. Select the Organizational Unit (OU) containing the target ChromeOS devices.
4. Add or configure a VPN connection (e.g., L2TP/IPsec, OpenVPN).
5. Review the available configuration options for the specific VPN type. Look for settings related to routing or split tunneling (e.g., options to route all traffic or only specific subnets). Configure these settings to route all traffic through the VPN to prevent split tunneling, if the option is available for the chosen VPN type. *Note Explicit "disable split tunnel" checkboxes may not exist for all VPN types; routing all traffic achieves the same effect.*

6. Click **Save**.

Additional Considerations

- **Security vs. Performance:** Disabling split tunneling enhances security by routing all traffic through monitored and controlled channels, but it can consume more bandwidth on the corporate network and potentially introduce latency for non-corporate traffic.
- **Zero Trust Model:** Adopting IAP aligns with a Zero Trust security model, which is increasingly considered a more secure approach than traditional perimeter-based VPN access.

Supplemental Guidance

- [Cloud VPN Overview](#)
- [IAP Overview](#)
- [IAP Product Page](#)
- [Chrome Enterprise Premium Overview](#)

Control Domain	System and Communications Protection								
Control #	SC.L2-3.13.8								
Control Description	Implement cryptographic mechanisms to prevent unauthorized disclosure of CUI during transmission unless otherwise protected by alternative physical safeguards								
Key Services	• Google Cloud Default Encryption at Rest • CMEK • Cloud KMS • CSEK • Client-Side Encryption	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Identifying cryptographic mechanisms intended to prevent unauthorized disclosure of CUI b. Implementing either cryptographic mechanisms to prevent unauthorized disclosure of CUI during transmission.									

When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control:

- Google Cloud Default Encryption at Rest
- Customer-Managed Encryption Keys (CMEK)
- Customer-Supplied Encryption Keys (CSEK)

A description of relevant features and implementation guidance is included below.

Customer-Managed Encryption Keys (CMEK): Provide greater control by allowing you to use keys managed in Cloud KMS to protect data in supported Google Cloud services.

- **Implementation**

1. Navigate to **Security > Key Management** in the Google Cloud Console.
2. Create a Key Ring in a desired location (should be co-located with the resources it will protect).
3. Create a cryptographic Key within the Key Ring (typically Symmetric, AES-256-GCM). Note the Key Resource Name.
4. Grant the appropriate service agent (e.g., the Cloud Storage service agent, Compute Engine service agent) for the service needing encryption the roles/cloudkms.cryptoKeyEncrypterDecrypter IAM role on the specific KMS key.
5. When creating a resource in a CMEK-supported service (e.g., a Cloud Storage bucket, Persistent Disk, BigQuery table, Cloud SQL instance)
 - Look for the "Encryption" or "Encryption key management" section in the creation options.
 - Select the "Customer-managed key" option.
 - Provide the Key Resource Name of the KMS key created in step 3.
6. Complete the resource creation. The service will now use your KMS key to wrap/unwrap the data encryption keys used for that resource.

Customer-Supplied Encryption Keys (CSEK): Provide your own AES-256 key for encrypting specific Cloud Storage objects or GCE Persistent Disks. Google does not store your key.

- **Implementation**

1. **Cloud Storage:** CSEK must be provided via the API, client libraries (like Python, Java, Go), or Google Cloud storage commands during object upload, download, copy, or rewrite operations using specific headers (x-goog-encryption-algorithm, x-goog-encryption-key, x-goog-encryption-key-sha256). It cannot be set via the Cloud Console web UI.

2. **Compute Engine:** CSEK can be provided when creating or attaching a Persistent Disk using the gcloud CLI or API. The key must be supplied whenever the disk is attached to a VM.
3. **Key Management:** You are entirely responsible for generating, securing, backing up, and providing the key whenever needed. Loss of the key means permanent loss of access to the data.

Additional Considerations

- **Key Management:** Using CMEK or CSEK shifts key management responsibilities to you. Securely managing key access (IAM for KMS keys), rotation, backup, and disaster recovery is critical. CSEK requires the most customer effort.
- **FIPS 140:** Google Cloud's default encryption and Cloud KMS software keys utilize cryptographic modules validated against FIPS 140 Level 1. Cloud HSM offers FIPS 140 Level 3 validation for CMEK keys.
- **Performance:** Default encryption and CMEK have minimal performance impact. CSEK involves key transfer with each operation, which can add latency. Client-side encryption impact depends on the application implementation.
- **Service Support:** Check documentation for which specific services support CMEK and CSEK.

Supplemental Guidance

- [Google Cloud default encryption at rest](#)
- [Encryption For Cloud Security \(Overview\)](#)
- [CMEK Overview](#)
- [Cloud Storage CMEK](#)
- [Compute Engine CMEK](#)
- [BigQuery CMEK](#)
- [CSEK Overview](#)
- [Cloud Storage CSEK](#)
- [Compute Engine CSEK](#)
- [KMS Overview](#)
- [KMS Product Page](#)
- [FIPS 140 Validation](#)

Control Domain	System and Communications Protection
Control #	SC.L2-3.13.9
Control Description	Terminate network connections associated with communications sessions at the end of the sessions or after a defined period of inactivity

Key Services	• Web Sessions • Compute Engine Instance Sessions (SSH/RDP) • Customer Application Sessions	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Protecting the authenticity of communications sessions When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Web Sessions (Google Cloud Console, Google Workspace Apps) • Compute Engine Instance Sessions (SSH/RDP) • Customer Application Sessions A description of relevant features and implementation guidance is included below. Web Sessions (Google Cloud Console, Google Workspace Apps): Use Google Admin Console session controls to define the maximum time a user can remain signed in before needing to re-authenticate. This terminates the session based on total duration, not necessarily inactivity within the session. • Implementation 1. Sign in to the Google Admin console (admin.google.com). 2. For general Workspace app sessions Navigate to Security > Access and data control > Google Session control. 3. For Google Cloud Console and gcloud CLI sessions Navigate to Security > Access and data control > Google Cloud session control. 4. Select the Organizational Unit (OU) to apply the policy to. 5. Configure the Web session duration (for Workspace) or Reauthentication frequency (for Google Cloud). Choose a duration (e.g., 1 hour, 8 hours, 12 hours, 1 day) after which the session associated with the network connection will expire, forcing re-authentication. 6. For Google Cloud sessions, also select the Reauthentication method (Password or Security key). 7. Click Save (or Override).			

Compute Engine Instance Sessions (SSH/RDP): Inactivity timeouts for interactive SSH or RDP sessions must be configured at the operating system level within the VM instance itself. This is outside the standard Google Cloud Console configuration for the VM resource.

- **Implementation (Conceptual - Requires VM Access)**

1. **Linux (SSH)** Edit the SSH daemon configuration file (/etc/ssh/sshd_config). Set ClientAliveInterval to the desired inactivity timeout in seconds (e.g., 900 for 15 minutes) and ClientAliveCountMax to 0. The server will send keep-alive messages; if ClientAliveCountMax is 0, the connection will terminate after ClientAliveInterval seconds of inactivity with no client response. Restart the SSH service after making changes.
2. **Windows (RDP)** Configure idle session timeout limits using Group Policy (Computer Configuration > Administrative Templates > Windows Components > Remote Desktop Services > Remote Desktop Session Host > Session Time Limits) or local security policy.
3. **Note** While IAP can control access to SSH/RDP based on the Google session duration, it does not inherently manage the *inactivity timeout* within the established SSH or RDP session itself.

Custom Application Sessions: Applications deployed on GKE, App Engine, Compute Engine, etc., require their own session management logic.

- **Implementation**

- Application developers must implement server-side or client-side mechanisms to track session activity and terminate sessions after a defined period of inactivity or upon explicit logout. This often involves session cookies with expiration times and server-side session state tracking.

Additional Considerations

- **Policy Definition:** Clearly define both maximum session duration and inactivity timeout periods in organizational policy.
- **Balance Security and Usability:** Shorter timeouts improve security but can impact user productivity if set too aggressively.
- **Session Termination vs. Inactivity:** Be aware that the Google Admin Console settings primarily control maximum session *duration*, forcing re-authentication regardless of activity. Inactivity timeouts usually require OS or application-level configuration.

Supplemental Guidance

- [Set session length for Google services](#)
- [Set session length for Google Cloud services](#)

- [Quotas and limits \(includes inactivity/duration limits\)](#)

Control Domain	System and Communications Protection								
Control #	SC.L2-3.13.10								
Control Description	Establish and manage cryptographic keys for cryptography employed in organizational systems.								
Key Services	• Cloud KMS • IAM • Cloud Audit Logs • Secret Manager	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Establishing cryptographic keys whenever cryptography is employed b. Managing cryptographic keys whenever cryptography is employed When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Cloud Key Management Service (Cloud KMS) • Identity and Access Management (IAM) for KMS • Cloud Audit Logs for KMS • Secret Manager (for CSEK / Application-Level Keys) A description of relevant features and implementation guidance is included below. Cloud Key Management Service (Cloud KMS): Use Cloud KMS as the central service for managing cryptographic keys required for CMEK or application-level cryptography. • Establish Keys (Generation/Import) 1. Navigate to Security > Key Management in the Cloud Console. 2. Click Create Key Ring and specify a name and location. 3. Within the key ring, click Create Key. 4. Specify a key name, protection level (Software, HSM, or External via EKM connection), purpose (e.g., Symmetric encrypt/decrypt, Asymmetric sign), and algorithm. 5. (Optional) Configure rotation schedule for symmetric keys. 6. Click Create.									

7. Alternatively, import existing key material following the key import procedures.
- **Manage Access (IAM)**
 1. Navigate to **Security > Key Management**.
 2. Select the Key Ring or specific Key.
 3. Use the **Permissions** tab in the info panel (or navigate to **IAM & Admin > IAM**) to grant roles.
 4. Assign administrative roles (e.g., roles/cloudkms.admin, roles/cloudkms.importer) only to personnel responsible for key management.
 5. Assign usage roles (e.g., roles/cloudkms.cryptoKeyEncrypterDecrypter, roles/cloudkms.signerVerifier, roles/cloudkms.publicKeyViewer) to users or service accounts that need to perform cryptographic operations, following least privilege.
 - **Manage Lifecycle (Rotation/Destruction)**
 1. For symmetric keys, configure an automatic **Rotation period** and **Next rotation** time during key creation or by editing the key later.
 2. Rotate symmetric keys manually if needed by selecting the key and clicking **Rotate**.
 3. Rotate asymmetric or imported keys manually by creating a new key version with new material (imported or generated) and updating applications to use the new version.
 4. To retire a key version Select the key, go to **Versions**, select the version, and click **Disable**.
 5. To destroy a key version After disabling, select the disabled version and click **Schedule Destruction**. There is a default recovery period (typically 24 hours) before permanent destruction.

Cloud Audit Logs for KMS: Enable and monitor audit logs for key management activities and key usage.

- **Implementation**
 1. Navigate to **IAM & Admin > Audit Logs**.
 2. Find "Cloud Key Management Service (KMS) API" in the list.
 3. Ensure **Admin Read**, **Admin Write** log types are enabled (enabled by default).
 4. Enable **Data Read** (for decryption/verification operations) and **Data Write** (for encryption/signing operations) if required for detailed auditing of key usage. *Note Data Access logs can generate significant volume and cost.*
 5. Monitor logs in Cloud Logging (**Logging > Logs Explorer**) or export them to a SIEM using log sinks.

Secret Manager (for CSEK / Application-Level Keys): Use Secret Manager to store and manage the actual key material for keys not directly managed within KMS (like CSEK keys).

- **Implementation**

1. Navigate to **Security > Secret Manager**.
2. Click **Create Secret**.
3. Provide a name for the secret (e.g., csek-key-for-bucket-abc).
4. Enter the key material (e.g., the base64-encoded AES-256 key) in the **Secret value** field.
5. Configure replication policy (Automatic or User-managed).
6. Click **Create Secret**.
7. Use IAM permissions on the secret resource to grant access only to the specific service accounts or users that need to retrieve the key material for cryptographic operations.
8. Use Secret Manager's versioning feature when rotating these keys.

Additional Considerations

- **Key Management Policy:** Maintain a documented policy outlining key lifecycle procedures, roles, responsibilities, and usage guidelines.
- **Separation of Duties:** Use distinct IAM roles to separate key administration tasks from key usage tasks.
- **Key Backup (External Keys):** For keys managed outside KMS (CSEK keys stored in Secret Manager, application-level keys), ensure robust backup and recovery procedures are in place. Loss of these keys means loss of data.
- **Compliance:** Select appropriate KMS protection levels (Software, HSM, EKM) based on compliance requirements (e.g., FIPS 140).

Supplemental Guidance

- [Cloud KMS Overview](#)
- [Cloud KMS Product Page](#)
- [Cloud KMS access control with IAM](#)
- [Cloud KMS Permissions and roles](#)
- [Key rotation](#)
- [Destroying and restoring key versions](#)
- [Cloud HSM Overview](#)
- [Cloud HSM Architecture](#)
- [Cloud External Key Manager \(EKM\) Overview](#)
- [Secret Manager Overview](#)
- [Secret Manager Product Page](#)

- [Cloud KMS audit logging](#)

Control Domain	System and Communications Protection								
Control #	SC.L2-3.13.11								
Control Description	Employ FIPS-validated cryptography when used to protect the confidentiality of CUI								
Key Services	• Google Cloud Default Encryption Mechanisms • Cloud KMS • Cloud VPN • Cloud Load Balancing	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Employing FIPS-validated cryptography to protect the confidentiality of CUI When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • Google Cloud Default Encryption (At Rest & In Transit) • Cloud Key Management Service (Cloud KMS) • Cloud VPN • Cloud Load Balancing • Secret Manager (for CSEK / Application-Level Keys) A description of relevant features and implementation guidance is included below. Google Cloud Default Encryption (At Rest & In Transit): Google Cloud's standard encryption mechanisms meet the FIPS 140 validation requirement. • Platform Behavior: As documented by Google, default encryption for data at rest (e.g., in Cloud Storage, Persistent Disk, BigQuery) and data in transit (e.g., TLS to Google Front Ends, traffic between data centers) uses the BoringCrypto module, which is FIPS 140 validated (Level 1). • Configuration: No specific configuration is needed to leverage this default FIPS-validated protection. Customers can rely on this for CUI protection unless higher levels of validation (e.g., FIPS 140 Level 3) or specific algorithms are mandated.									

Cloud Key Management Service (Cloud KMS): If choosing to use customer managed keys for encryption, ensure that the appropriate key protection level is based on FIPS requirements when using customer-managed keys.

- **Implementation**

1. Navigate to **Security > Key Management**.
2. When creating a new key (**Create Key**)
 - Under **Protection level**, choose
 - **Software** Uses cryptographic modules validated at FIPS 140 Level 1. Suitable for many CUI scenarios.
 - **HSM** Uses hardware security modules validated at FIPS 140 Level 3. Select this if Level 3 validation is required by policy or regulation for the keys protecting CUI.
3. Use keys with the selected protection level for CMEK or application-level cryptography involving CUI.

Cloud VPN: The Cloud VPN service enables VPN tunnels to use FIPS-approved cryptographic algorithms.

- **Implementation**

1. Navigate to **Hybrid Connectivity > VPN**.
2. When creating or editing a Cloud VPN tunnel, configure the IKE and IPsec settings.
3. In the **IKE version** section, select IKEv2 (preferred).
4. In the **IKEv2 encryption algorithms**, **Integrity algorithms**, and **DH groups** sections (Phase 1 & Phase 2), select ciphers approved by FIPS. Generally, this includes
 - Encryption AES-128-GCM, AES-256-GCM, AES-128-CBC, AES-256-CBC.
 - Integrity/PRF SHA-256, SHA-384, SHA-512 based HMACs or PRFs.
 - DH Groups Groups 14 (2048-bit), 15 (3072-bit), 16 (4096-bit), 19 (ECP-256), 20 (ECP-384), 21 (ECP-521) are generally acceptable. Avoid weaker groups like 2.
5. Consult the Cloud VPN supported ciphers list and FIPS standards for specific approved algorithm identifiers. Ensure the peer VPN gateway is also configured with compatible, FIPS-approved algorithms.

Cloud Load Balancing (SSL Policies): Load balancing SSL policies for HTTPS and SSL load balancers can be configured to use FIPS-approved TLS versions and cipher suites.

- **Implementation**

1. Navigate to **Network Security > SSL policies**.

2. Click **Create Policy**.
3. Set the **Minimum TLS version** to TLS 1.2 or TLS 1.3.
4. For the **Profile**, select **Custom**.
5. Click **Configure Features**.
6. Deselect any cipher suites that do not use FIPS-approved algorithms. Generally, prioritize suites using AES-GCM or CHACHA20-POLY1305 for encryption and SHA2 (SHA-256, SHA-384) for integrity. Examples of strong, commonly FIPS-approved ciphers available in Google Cloud include
 - TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
 - TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
 - TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256
 - TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
 - TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
 - TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256
 - (For TLS 1.3) TLS_AES_128_GCM_SHA256, TLS_AES_256_GCM_SHA384, TLS_CHACHA20_POLY1305_SHA256
7. Avoid older algorithms like RC4, DES, 3DES, MD5, SHA1.
8. Click **Save**.
9. Attach this custom SSL policy to the target HTTPS or SSL proxy associated with your load balancer (**Network Services > Load balancing > Select LB > Edit > Frontend configuration**).

Additional Considerations

- **FIPS 140-3:** Note that FIPS 140-3 is gradually replacing FIPS 140-2. Check Google Cloud compliance documentation for updates on FIPS 140-3 validations. FIPS 140-2 validated modules remain acceptable for federal agencies until 21-Sep 2026.
- **Module vs. Algorithm:** The requirement is for a FIPS-validated *module*. Using FIPS-approved *algorithms* (like AES) within a non-validated module does not meet the requirement. Relying on Google's default encryption or Cloud KMS ensures use of a validated module.
- **Client-Side Implementation:** If implementing encryption client-side or within custom applications handling CUI, customers are responsible for selecting and using libraries or modules that are FIPS 140-2/3 validated.

Supplemental Guidance

- [FIPS 140 Validated](#)
- [Cloud KMS Protection levels](#)
- [Cloud KMS Overview](#)
- [Cloud VPN supported IKE ciphers](#)
- [Cloud Load Balancing SSL policies overview](#)

- [Configuring SSL policies](#)
- [NIST Cryptographic Module Validation Program](#)

Control Domain	System and Communications Protection		
Control #	SC.L2-3.13.12		
Control Description	Prohibit remote activation of collaborative computing devices and provide indication of devices in use to users present at the device		
Key Services	N/A	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for - Identifying collaborative computing devices - Ensuring collaborative computing devices provide indication to users of devices in use - Prohibiting remote activation of collaborative computing devices <i>Google Cloud does not natively offer any features that, if configured correctly, will address this control. You should implement this control outside of Google Cloud, in accordance with your CUI boundary.</i>			
Supplemental Guidance			
N/A			

Control Domain	System and Communications Protection		
Control #	SC.L2-3.13.13		
Control Description	Control and monitor the use of mobile code		
Key Services	- Google Cloud Armor - SCC	Control Responsibility	☐ Google ☐ Shared ☒ Customer

Customer Implementation Description

Google Cloud Customers are responsible for:

- a. Controlling the use of mobile code
- b. Monitoring the use of mobile code

When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control:

- Google Cloud Armor
- Security Command Center (SCC)

A description of relevant features and implementation guidance is included below.

Google Cloud Armor: Implement WAF rules to filter malicious requests, including common script injection attacks.

- **Implementation**

1. Ensure your web application is fronted by a supported Google Cloud external load balancer.
2. Navigate to **Network Security > Cloud Armor**.
3. Create or edit a Security Policy attached to the load balancer's backend service(s).
4. Add rules using preconfigured WAF rulesets
 - Click **Add Rule**.
 - Choose **Apply preconfigured WAF rules**.
 - Select rules relevant to mobile code, specifically xss-stable or xss-canary (Cross-Site Scripting).
 - Choose an appropriate **Sensitivity level** (lower levels have fewer false positives but might miss some attacks; higher levels are more protective but riskier for false positives). Start with a lower level (e.g., 1) in Preview mode.
 - Set the **Action** to **Deny** (or **Preview** initially for testing).
5. Configure other rules as needed (e.g., SQLi protection).
6. Monitor logs (ensure Firewall Rules Logging is enabled for the VPC or specific rules are logged via Cloud Armor policy) to tune rules and identify blocked threats or false positives.

Security Command Center (SCC): Use Web Security Scanner to identify vulnerabilities that could be exploited by mobile code.

- **Implementation**

1. Ensure SCC Premium is activated.

2. Navigate to **Security > Security Command Center > Web Security Scanner**.
3. Create **Custom scans** targeting your public-facing web applications hosted on App Engine, GKE, or Compute Engine.
 - Define **Starting URLs**.
 - Configure authentication if needed (using a test account).
 - Schedule the scans.
4. Monitor findings in **Security > Security Command Center > Vulnerabilities** (filtering for XSS category, among others).
5. Remediate identified vulnerabilities in the application code.

Additional Considerations

- **Balancing Act:** Completely blocking JavaScript will break most modern web functionality. Control focuses on limiting execution from untrusted sources and managing extensions.
- **Browser Updates:** Keeping browsers updated (managed via Admin Console auto-update policies) is crucial as updates often contain security fixes related to mobile code execution.
- **Extension Vetting:** Thoroughly vet any extensions before allowlisting them, paying close attention to requested permissions.

Supplemental Guidance

- [Set Chrome policies for users or browsers](#)
- [Allow or block apps and extensions](#)
- [Google Cloud App security policy overview](#)
- [Preconfigured WAF rules overview \(includes XSS\)](#)
- [Tune preconfigured WAF rules](#)
- [Web Security Scanner Overview](#)
- [Using Web Security Scanner](#)
- [MDN Web Docs – Content Security Policy](#)
- [Apigee - Configure a content security policy](#)

Control Domain	System and Communications Protection		
Control #	SC.L2-3.13.14		
Control Description	Control and monitor the use of Voice over IP (VoIP) technologies		
Key Services	● Compute Engine/GKE	Control Responsibility	☐ Google ☐ Shared ☒ Customer

Customer Implementation Description

Google Cloud Customers are responsible for:

- a. Controlling use of Voice over Internet Protocol (VoIP) technologies
- b. Monitoring use of Voice over Internet Protocol (VoIP) technologies

When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control:

- Compute Engine/GKE

A description of relevant features and implementation guidance is included below.

Controlling/Monitoring VoIP Hosted on Google Cloud (Compute Engine/GKE)

- **Configure VPC Firewall Rules**
 1. Navigate to **VPC network > Firewall** in the Google Cloud Console.
 2. Create specific allow ingress rules for necessary VoIP ports (e.g., UDP/TCP 5060/5061 for SIP, UDP 10000-20000 for RTP/SRTP, or the specific range used by your application) targeting your VoIP server instances (using tags or service accounts).
 3. Restrict the **Source IP ranges** in these rules to only trusted locations (e.g., office IPs, specific remote user VPN ranges, trusted peering IPs). Avoid allowing 0.0.0.0/0.
 4. Configure appropriate egress rules if servers need outbound access.
 5. Enable **Logs** for critical allow/deny rules related to VoIP traffic.
- **Enable VPC Flow Logs**
 1. Navigate to **VPC network > Subnets**.
 2. Select the subnet(s) hosting VoIP servers.
 3. Click **Edit**.
 4. Set **Flow logs** to **On**. Configure aggregation interval and sample rate as needed (higher sampling captures more detail but increases log volume).
 5. Click **Save**.
 6. Monitor logs in **Cloud Logging** for unusual traffic patterns, large data transfers, or connections to unexpected IPs.
- **Deploy Cloud IDS (Optional)**
 1. Navigate to **Network Security > Cloud IDS**.
 2. Create an IDS endpoint in the relevant region(s).
 3. Configure **Packet Mirroring (Compute Engine > Packet Mirroring)** to mirror traffic from your VoIP server instances/subnets to the Cloud IDS endpoint.
 4. Monitor Cloud IDS findings in **Cloud Logging** or **Security Command Center** for potential threats targeting VoIP protocols or infrastructure.

- **VoIP Application Security**

1. Ensure the VoIP software itself is configured securely (strong passwords, user authentication, TLS for signaling SIP TLS, SRTP for media encryption). This configuration happens within the application, not the Google Cloud Console.

Additional Considerations

- **Media Encryption (SRTP):** For hosted VoIP, strongly recommend configuring and enforcing Secure Real-time Transport Protocol (SRTP) within the VoIP application to encrypt voice media streams.
- **Signaling Encryption (SIP TLS):** For hosted VoIP, use TLS to encrypt SIP signaling traffic where possible to protect call setup information.
- **Network Quality:** Monitor network latency and jitter (using Cloud Monitoring metrics for VMs/LBs and potentially VoIP application metrics) as these significantly impact VoIP quality.
- **Compliance:** Ensure any call recording features (Google Voice or hosted solutions) comply with relevant legal and regulatory requirements regarding consent and notification.

Supplemental Guidance

- [VPC Firewall Rules Overview](#)
- [Using VPC Firewall Rules](#)
- [Using VPC Flow Logs](#)
- [Cloud IDS Overview](#)

Control Domain	System and Communications Protection		
Control #	SC.L2-3.13.15		
Control Description	Protect the authenticity of communications sessions		
Key Services	• TLS / HTTPS • Cloud VPN • Cloud Interconnect IAP • Anthos Service Mesh	Control Responsibility	☐ Google ☐ Shared ☒ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Protecting the authenticity of communications sessions			

When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control:

- TLS/HTTPS
- Authenticated VPN/Interconnect
- Identity Aware Proxy (IAP) with Header Validation
- Mutual TLS (mTLS) with Anthos Service Mesh (Optional/Advanced)

A description of relevant features and implementation guidance is included below.

TLS/HTTPS for Web Communications: Use TLS to authenticate servers and encrypt web traffic for applications hosted on Google Cloud.

- **Implementation**

1. Configure External or Internal HTTP(S) Load Balancers with SSL/TLS certificates (Google-managed or self-managed) via **Network Services > Load balancing**. Ensure frontend protocols are set to HTTPS.
2. Configure SSL policies (**Network Security > SSL policies**) to enforce strong TLS versions (1.2+) and appropriate cipher suites.
3. For App Engine or Cloud Run, map custom domains and utilize the automatically provisioned Google-managed certificates or upload custom certificates to enable HTTPS.
4. Configure applications to redirect HTTP traffic to HTTPS.

Authenticated VPN/Interconnect: Ensure hybrid connections use protocols that authenticate endpoints.

- **Implementation**

1. When configuring Cloud VPN (**Hybrid Connectivity > VPN**), use strong, unique pre-shared keys for IKE authentication. Where possible and supported by the peer device, consider certificate-based authentication for IPsec.
2. When using Cloud Interconnect (**Hybrid Connectivity > Interconnect**), enhance authenticity by deploying HA VPN over Cloud Interconnect (uses IPsec authentication) or configuring MACsec (link-layer authentication and encryption) on supported Dedicated Interconnect connections.

Identity-Aware Proxy (IAP) with Header Validation: Use IAP to verify user identity before allowing access and enable backend validation for higher assurance.

- **Implementation**

1. Enable IAP for relevant resources (Web applications, VMs via TCP Forwarding) as described under SC.3.13.7 (**Security > Identity-Aware Proxy**).

2. Grant appropriate IAM roles (IAP-Secured Web App User or IAP-Secured Tunnel User) to authorized principals.
3. **Backend Validation (Crucial for authenticity assurance)** Modify backend application code to
 - Retrieve the JWT assertion from the X-Goog-IAP-JWT-Assertion HTTP header.
 - Fetch Google's public keys for IAP (available at a public URL).
 - Use a standard JWT library to validate the token's signature against Google's public keys.
 - Verify the aud (audience) claim in the JWT matches the expected audience value for your backend service (obtainable via gcloud or console).
 - Verify the iss (issuer) claim is <https://cloud.google.com/iap>.
 - Ensure the token is not expired (exp) and was issued in the past (iat).
 - If validation succeeds, trust the email and sub (subject ID) claims within the JWT as the authenticated user's identity. Deny requests that fail validation or lack the header.

Mutual TLS (mTLS) with Anthos Service Mesh (Optional/Advanced): For service-to-service communication within a mesh (e.g., GKE), enforce mTLS where services mutually authenticate each other using certificates.

- **Implementation**

1. Ensure Anthos Service Mesh (managed Istio) is installed on your GKE cluster(s).
2. Configure PeerAuthentication policies at the mesh, namespace, or workload level. Set the mtls.mode to STRICT to require mTLS for incoming connections.
3. Configure DestinationRule resources with trafficPolicy.tls.mode set to ISTIO_MUTUAL for services making outbound calls to ensure they initiate mTLS connections.
4. Anthos Service Mesh manages the certificate issuance and rotation for service identities within the mesh.

Additional Considerations

- **Certificate Management:** Securely manage private keys for self-managed SSL certificates and certificates used for Cloud VPN authentication. Use Certificate Manager or rely on Google-managed certificates where possible.
- **IAP Header Validation:** Implementing backend validation of the IAP JWT header is critical to fully protect against scenarios where IAP might be bypassed or misconfigured.

- **Layered Security:** Combine these session authenticity controls with strong user authentication (MFA) and network segmentation for a defense-in-depth posture.

Supplemental Guidance

- [SSL certificates overview \(Load Balancing\)](#)
- [Use Google-managed SSL certificates](#)
- [Cloud VPN overview](#)
- [Best practices for Cloud VPN \(includes using strong PSKs\)](#)
- [IAP Product Page](#)
- [Securing your app with signed headers \(JWT validation\)](#)
- [Configure transport security \(mTLS\)](#)
- [ASM Security Overview](#)

Control Domain	System and Communications Protection								
Control #	SC.L2-3.13.16								
Control Description	Protect the confidentiality of CUI at rest								
Key Services	• Google Cloud Default Encryption at Rest • CMEK • CSEK	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Protecting the confidentiality of CUI at rest When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control: • Google Cloud Default Encryption at Rest • Customer-Managed Encryption Keys (CMEK) • Customer-Supplied Encryption Keys (CSEK) A description of relevant features and implementation guidance is included below. Google Cloud Default Encryption at Rest: This is the baseline protection applied automatically by Google Cloud to protect CUI confidentiality at rest.									

- **Platform Behavior:** All customer data written to persistent storage in services like Cloud Storage, Persistent Disk, Cloud SQL, BigQuery, etc., is automatically encrypted at the storage layer using AES-256 (or sometimes AES-128 for older disks) within FIPS 140 validated modules. Google manages the keys used for this encryption transparently.
- **Configuration:** No specific customer configuration is required to enable default encryption; it is always active. This default protection addresses the core requirement for encrypting CUI at rest.

Customer-Managed Encryption Keys (CMEK): Provides enhanced control by allowing you to use keys managed in Cloud KMS to protect data in supported Google Cloud services.

- **Implementation**
 1. Navigate to **Security > Key Management** in the Cloud Console.
 2. Create a Key Ring and a cryptographic Key (e.g., Symmetric AES-256-GCM). Choose the appropriate **Protection level** (Software or HSM) based on compliance needs. Note the Key Resource Name.
 3. Grant the appropriate service agent (e.g., Cloud Storage service agent) the roles/cloudkms.cryptoKeyEncrypterDecrypter IAM role on the specific KMS key.
 4. When creating a resource in a CMEK-supported service (e.g., Cloud Storage bucket, Persistent Disk, BigQuery table)
 - Select the "Customer-managed key" option in the encryption settings.
 - Provide the Key Resource Name of the KMS key.
 5. Complete resource creation.

Customer-Supplied Encryption Keys (CSEK): Allows you to provide your own AES-256 key for encrypting specific Cloud Storage objects or GCE Persistent Disks. Google does not store your key.

- **Implementation**
 1. CSEK must be provided via API, client libraries, or Google Cloud CLI commands using specific headers/flags for each operation involving the data (writes, reads). This is not configured via a persistent setting in the Cloud Console web UI for the data itself.
 2. The customer bears full responsibility for key generation, security, backup, and availability. Loss of the key results in permanent data loss. Due to the operational overhead and risk, CMEK is often preferred over CSEK.

Additional Considerations

- **Key Management:** Implementing CMEK or CSEK requires robust key management processes, including secure storage, access control (IAM for KMS keys), rotation, and backup/recovery (especially critical for CSEK).
- **FIPS 140 Validation:** Google's default encryption and Cloud KMS (Software and HSM levels) utilize FIPS 140 validated modules, directly supporting this CMMC requirement when encryption is employed.
- **Audit Logging:** Enable Cloud Audit Logs for KMS (including Data Access logs if needed) and relevant storage services to track key usage and data access.

Supplemental Guidance

- [Default encryption at rest](#)
- [Encryption For Cloud Security \(Overview\)](#)
- [CMEK Overview](#)
- [Cloud Storage CMEK](#)
- [Compute Engine CMEK](#)
- [CSEK Overview](#)
- [Cloud Storage CSEK](#)
- [Compute Engine CSEK](#)
- [KMS Overview](#)
- [FIPS 140 Validation](#)

Control Domain	System and Information Integrity		
Control #	SI.L1-3.14.1		
Control Description	Identify, report, and correct system flaws in a timely manner		
Key Services	• SCC • VM Manager • Artifact Registry • Network Analyzer • IAM • Cloud Logging & Cloud Monitoring • OS Configuration management • IaC	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Specifying the time within which to identify system flaws - Identifying system flaws within the specified time frame			

- c. Specifying the time within which to report system flaws
- d. Reporting system flaws within the specified time frame
- e. Specifying the time within which to correct system flaws
- f. Correcting system flaws within the specified time frame

When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control:

- Security Command Center (SCC) (Security Health Analytics, Web Security Scanner, Event/VM/Container Threat Detection)
- VM Manager (Vulnerability Reporting, Patch Management)
- Artifact Registry with Container Analysis
- Network Analyzer & Firewall Insights
- IAM Recommender
- Cloud Logging & Cloud Monitoring (Alerts on new critical findings)
- Direct Console Configuration for services like IAM, VPC Firewall, Cloud Storage based on identified misconfigurations.
- Infrastructure as Code (IaC) tools (e.g., Cloud Deployment Manager, Terraform)
- OS Configuration management

A description of relevant features and implementation guidance is included below.

Identify System Flaws

- **Security Command Center (SCC):** Use SCC as a primary tool to identify a broad range of flaws, from cloud resource misconfigurations to web application vulnerabilities and potential runtime threats that might indicate underlying flaws.
 - **Implementation**
 - Ensure SCC is activated (**Security > Security Command Center**) with all relevant services enabled (Security Health Analytics, Web Security Scanner, Event/VM/Container Threat Detection for premium tiers).
 - **Security Health Analytics:** Continuously identifies misconfigurations in your Google Cloud resources (e.g., open firewalls, public storage, insecure API key settings, weak IAM policies). These are system flaws.
 - **Web Security Scanner:** Periodically scan your web applications (**SCC > Configuration > Web Security Scanner**) for common vulnerabilities like XSS, mixed content, etc.
 - Review findings regularly in **SCC > Findings**.
- **VM Manager (Vulnerability Reporting):** Use VM Manager to identify known software vulnerabilities (CVEs) in the OS and installed packages on your Compute Engine VMs.
 - **Implementation**

- Enable VM Manager for your projects (**Compute Engine > VM Manager**). Ensure OS Config agents are active on VMs.
 - Regularly review **Vulnerability reports** which list identified CVEs, their severity, and affected VMs.
- **Artifact Registry with Container Analysis:** Scan container images stored in Artifact Registry for known vulnerabilities in OS packages and application libraries.
 - **Implementation**
 - Push container images to Artifact Registry. Scanning is typically automatic.
 - Review vulnerability findings in **Security > Container Analysis > Vulnerabilities** or within the Artifact Registry UI for specific images.
- **Network Analyzer & Firewall Insights:** Identify flaws in network configurations and firewall rules.
 - **Implementation**
 - Review insights from **Network Intelligence > Network Analyzer** and **VPC network > Firewall Insights** for issues like overly permissive rules or suboptimal network configurations.
- **IAM Recommender:** Identify flaws related to excessive IAM permissions.
 - **Implementation**
 - Review recommendations in **IAM & Admin > IAM** or **Security > Recommendations Hub**.

Report System Flaws (Facilitating Internal Reporting): Identified flaws need to be reported to internal teams responsible for remediation. Google Cloud tools can provide the necessary information and automation hooks.

- **Security Command Center (SCC):** Use SCC's notification capabilities to ensure that when flaws are identified, the relevant internal teams are promptly informed. Finding details from SCC serve as the basis for internal flaw reports.
 - **Implementation**
 - Configure continuous exports of SCC findings to Pub/Sub (**SCC > Settings > Continuous Exports**).
 - Subscribe your internal ticketing system (e.g., Jira, ServiceNow) or alerting channels to this Pub/Sub topic to automatically create tasks or notifications for flaw remediation based on severity or type.
- **Cloud Logging & Cloud Monitoring:** Create alerts for new, high-severity flaws detected by SCC, VM Manager (if findings are logged appropriately), or other integrated tools.
 - **Implementation**

- Develop log-based alerts (**Logging > Log-based Alerts**) that trigger when specific high-severity flaw patterns appear in logs (e.g., logs from the Pub/Sub topic fed by SCC).
- Ensure notifications from these alerts are routed to the teams responsible for flaw management and reporting.

Correct System Flaws (in a Timely Manner): Flaw correction involves implementing changes to address the identified weaknesses. "Timely manner" is defined by the customer's risk assessment and SLAs.

- **VM Manager (Patch Management):** For OS-level CVEs on Compute Engine VMs identified by Vulnerability Reporting, use VM Manager's Patch Management to apply necessary patches.
 - **Implementation**
 - Based on internal flaw reports and remediation priorities, create patch deployments in **Compute Engine > VM Manager > Patch deployments**.
 - Schedule these deployments during appropriate maintenance windows to correct identified vulnerabilities in a timely fashion.
- **Security Command Center (Remediation Guidance):** Many SCC findings include "Next steps" or remediation guidance. Follow this guidance to correct misconfigurations.
 - **Implementation** For example, if SCC reports a publicly accessible Cloud Storage bucket, the correction involves navigating to **Cloud Storage > Buckets**, selecting the bucket, editing permissions, and removing public access.
- **Direct Console Configuration / IaC:** For flaws related to misconfigurations in IAM, firewalls, network settings, or other service configurations, use the respective service consoles or, preferably, update your Infrastructure as Code (IaC) templates and re-deploy.
 - **Implementation (Examples)**
 - **IAM Flaws:** Modify policies in **IAM & Admin > IAM**.
 - **Firewall Flaws:** Edit rules in **VPC network > Firewall**.
 - **IaC:** Update Terraform (.tf) or Deployment Manager (.yaml) files with the correct configurations and apply the changes. This ensures corrections are documented as code and consistently applied.
- **Artifact Registry / Container Rebuilds:** For container image vulnerabilities, the correction involves updating base images or vulnerable libraries, rebuilding the image, and re-deploying applications.
 - **Implementation**

- Modify Dockerfiles, update dependencies, rebuild images, push to Artifact Registry, and update GKE or Cloud Run deployments to use the new, corrected image.
- **OS Configuration management:** If a flaw is due to an OS configuration drift from a secure baseline, use OS Configuration management to re-apply or enforce the correct settings.
 - **Implementation**
 - Create or update OS policy assignments in **Compute Engine > OS Configuration management** to push the corrected configuration to affected VMs.

Additional Considerations

- **Flaw Prioritization:** Use a risk-based approach (informed by your risk assessments - [RA.L2-3.11.1](#)) to prioritize the order and timeliness of flaw correction. Address flaws affecting CUI systems or high-impact vulnerabilities first.
- **Change Management:** Integrate flaw correction activities into your organization's change management process to ensure changes are tested, approved, and documented.
- **Verification of Correction:** After applying corrective actions, re-scan or re-assess (using tools like SCC, VM Manager, Web Security Scanner) to verify that the flaw has been successfully eliminated or mitigated to an acceptable level.
- **Documentation:** Maintain comprehensive records of identified flaws, the reporting process, corrective actions taken, responsible parties, dates, and verification results. This is crucial for audits and continuous improvement.
- **"Timely Manner":** Your organization must define what "timely" means for different types of flaws based on their severity and potential impact. This should be documented in your vulnerability management policy or SLAs.

Supplemental Guidance

- [Security Command Center Overview](#)
- [About VM Manager | Compute Engine Documentation | Google Cloud](#)
- [Artifact analysis and vulnerability scanning | Artifact Registry documentation | Google Cloud](#)
- [Configure Security Command Center services | Google Cloud](#)
- [OS images | Compute Engine Documentation | Google Cloud](#)

Control Domain	System and Information Integrity
Control #	SI.L1-3.14.2

Control Description	Provide protection from malicious code at designated locations within organizational systems		
Key Services	• VM Manager • Artifact Registry • Binary Authorization • VPC Firewall Rules • Google Cloud Armor • Shielded VMs • Organization Policies • OS images • SCC • Cloud IDS • Cloud Logging • Cloud Monitoring	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: - Identifying designated locations for malicious code protection - Providing protection from malicious code at designated locations When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control: • VM Manager (Patch Management) • Artifact Registry with Container Analysis (Vulnerability and potentially basic malware scanning for images) • Binary Authorization (for container deployments) • VPC Firewall Rules • Google Cloud Armor (Web Application Firewall) • Shielded VMs (for boot integrity on Compute Engine) • Organization Policies (for restricting risky configurations) • Use of Google-provided hardened OS images • Security Command Center (SCC) (especially Event Threat Detection, VM Threat Detection, Container Threat Detection - primarily Premium/Enterprise tiers) • Cloud IDS (Intrusion Detection System) • Cloud Logging and Cloud Monitoring (for custom alerts and ingesting logs from host-based protection tools) A description of relevant features and implementation guidance is included below. Preventative Measures for Malicious Code Protection • VM Manager (Patch Management): Keeping operating systems patched is a fundamental step in preventing malware that exploits known software vulnerabilities.			

- **Implementation**
 - Use **Compute Engine > VM Manager > Patch deployments** to schedule and apply OS patches to your Compute Engine VMs regularly. This reduces the attack surface for many types of malicious code.
- **Artifact Registry with Container Analysis:** Scan container images for known vulnerabilities before deployment. Some scanners may also detect common malware packaged in images.
 - **Implementation**
 - Store container images in **Artifact Registry**. Vulnerability scanning is typically enabled by default and will scan images upon push and continuously monitor for new vulnerabilities. Review findings in **Security > Container Analysis > Vulnerabilities**.
- **Binary Authorization:** Enforce that only trusted and attested container images (which should have been scanned and vetted for malicious code) can be deployed to Google Kubernetes Engine (GKE) or Cloud Run.
 - **Implementation**
 - Configure Binary Authorization policies in **Security > Binary Authorization**, requiring attestations from trusted sources (e.g., successful vulnerability scans, internal reviews) before an image can be deployed.
- **VPC Firewall Rules:** Implement strict ingress and egress firewall rules to limit network attack vectors and prevent communication with known malicious command-and-control (C2) servers or malware distribution sites.
 - **Implementation**
 1. Navigate to **VPC network > Firewall**.
 2. Apply a default-deny ingress policy. Only allow traffic on necessary ports from trusted sources.
 3. Implement egress filtering to restrict outbound connections from your VMs to only known-good destinations. Consider using threat intelligence feeds for IP/domain blocklists (can be implemented with higher priority deny rules or via Cloud NGFW).
- **Google Cloud Armor (Web Application Firewall):** Protect web applications and APIs from common web attacks, some of which can be used to inject or deliver malicious code.
 - **Implementation**
 1. Navigate to **Network Security > Cloud Armor**.
 2. Create security policies and attach them to your HTTP(S) Load Balancers.
 3. Utilize preconfigured WAF rules (e.g., for SQL injection, XSS, common attacks) and consider custom rules.

4. Use IP allow/deny lists and geo-based blocking to restrict access from known malicious sources.
- **Shielded VMs (Compute Engine):** Enhance VM security by ensuring boot integrity, protecting against boot-level or kernel-level malware and rootkits.
 - **Implementation**
 - When creating VMs (**Compute Engine > VM instances > Create Instance**), in the "Shielded VM" section, enable "Secure Boot," "vTPM," and "Integrity Monitoring."
- **Organization Policies:** Enforce policies that reduce the risk of malware, such as restricting the creation of external IPs on VMs or disabling risky APIs.
 - **Implementation**
 - Configure relevant constraints in **IAM & Admin > Organization Policies** (e.g., constraints/compute.vmExternallpAccess).
- **Use of Google-provided Hardened OS Images:** Start with Google's hardened OS images (e.g., Container-Optimized OS, Shielded VM images) which have a reduced attack surface.
 - **Implementation**
 - Select these images when creating new Compute Engine VMs.

Detective Measures for Malicious Code

- **Security Command Center (SCC) - Threat Detection (Primarily Premium/Enterprise Tiers):** Enable and monitor SCC's threat detection services for identifying active malicious code behavior in your Google Cloud environment.
 - **Implementation**
 - **Event Threat Detection:** Detects threats like malware, crypto mining, C2 communication, and data exfiltration based on log analysis. (Findings appear in **SCC > Findings**).
 - **VM Threat Detection:** Detects runtime threats within VMs, such as suspicious executables, kernel modifications, and connections to malicious domains.
 - **Container Threat Detection:** Detects runtime threats in GKE clusters like unexpected binary execution, reverse shells, or unexpected library loads.
 - Configure notifications for these findings via Pub/Sub to your security operations team.
- **Cloud IDS (Intrusion Detection System):** Deploy Cloud IDS to monitor ingress and egress traffic within your VPC networks for known malicious signatures and anomalous activity.
 - **Implementation**
 1. Navigate to **Network Security > Cloud IDS**.

2. Create an IDS endpoint for each network you want to monitor.
 3. Configure IDS policies to specify threat severity levels for alerting.
 4. Review IDS threat logs and alerts in Security Command Center and Cloud Logging.
- **Cloud Logging and Cloud Monitoring (for Custom/Host-Based Detections):** Centralize logs from host-based anti-malware/EDR solutions running on your Compute Engine VMs into Cloud Logging. Create alerts in Cloud Monitoring for detected malicious code events from these tools.
 - **Implementation**
 1. Install and configure the Ops Agent on VMs to forward security tool logs (e.g., from `/var/log/` or Windows Event Logs where AV/EDR logs are written) to Cloud Logging.
 2. Create log-based alerts (**Logging > Log-based Alerts**) for specific log entries that indicate malware detection by your host-based tools.

Additional Considerations

- **Defense-in-Depth:** No single tool is a silver bullet. Employ multiple layers of malicious code protection (network, host, application, container).
- **Regular Updates:** Keep all software, including OS, applications, AV/EDR signatures, and WAF rules, up to date.
- **User Awareness Training:** Train users to identify and avoid phishing, malicious websites, and unsafe downloads, as users are often the initial entry point for malicious code.
- **Principle of Least Functionality:** Configure systems to provide only essential capabilities, reducing the attack surface available to malicious code (see [CM.L2-3.4.6](#)).
- **Incident Response Plan:** Have a plan in place to respond to malicious code incidents, including isolating affected systems, eradicating the malware, and recovering.

Supplemental Guidance

- [Overview of Event Threat Detection | Security Command Center | Google Cloud](#)
- [Virtual Machine Threat Detection overview | Security Command Center | Google Cloud](#)
- [Container Threat Detection overview | Security Command Center | Google Cloud](#)
- [About Patch | VM Manager | Google Cloud](#)
- [Artifact analysis and vulnerability scanning | Artifact Registry documentation | Google Cloud](#)
- [Binary Authorization overview | Google Cloud](#)
- [Product overview | Google Cloud Armor](#)

- [What is Shielded VM? | Google Cloud](#)
- [Cloud IDS Overview](#)

Control Domain	System and Information Integrity								
Control #	SI.L2-3.14.3								
Control Description	Monitor system security alerts and advisories and take action in response								
Key Services	• SCC • VM Manager • Artifact Registry with Container Analysis • Recommendations Hub • Google Cloud Security Bulletins	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Identifying response actions to system security alerts and advisories b. Monitoring system security alerts and advisories c. Taking action in response to system security alerts and advisories When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control: • Security Command Center (SCC) (Security Health Analytics, Event/VM/Container Threat Detection, Web Security Scanner) • VM Manager (Vulnerability Reporting as advisories for VMs) • Artifact Registry with Container Analysis (Vulnerability findings as advisories for containers) • Recommendations Hub (Aggregated recommendations as advisories) • Google Cloud Security Bulletin A description of relevant features and implementation guidance is included below. Monitor System Security Alerts and Advisories • Security Command Center (SCC): SCC is the primary Google Cloud service for centralizing security alerts (findings) from various built-in detectors. Configure SCC and continuously monitor its findings. ○ Implementation									

- Ensure SCC is activated (**Security > Security Command Center**) with relevant services (Security Health Analytics, Event/VM/Container Threat Detection, Web Security Scanner) enabled for your projects/organization.
 - **Monitor Findings:** Regularly review the **Findings** dashboard in SCC. Filter by severity, category, or asset to focus on the most critical alerts.
 - **Notifications:** Set up real-time notifications for new SCC findings. Configure continuous exports to Pub/Sub (**SCC > Settings > Continuous Exports**) and integrate this with email, chat ops (e.g., Slack), or your ticketing system so that relevant teams are immediately alerted to new findings.
- **VM Manager (Vulnerability Reporting):** Treat vulnerability reports from VM Manager as security advisories specific to your Compute Engine VMs. Monitor these for new, critical vulnerabilities affecting your systems.
 - **Implementation**
 - Regularly review **Compute Engine > VM Manager > Vulnerability reports**. Note that VM Manager is updated with new CVE information, so it effectively provides ongoing advisories as new vulnerabilities are discovered that affect your scanned OS packages.
- **Artifact Registry with Container Analysis:** Monitor vulnerability scan results for your container images in Artifact Registry. These findings act as advisories for vulnerabilities within your containerized applications.
 - **Implementation**
 - Regularly review scan results in **Artifact Registry** (select image) or **Security > Container Analysis > Vulnerabilities**. New vulnerabilities discovered by Google will trigger re-analysis of existing images.
- **Recommendations Hub:** Monitor the Recommendations Hub for proactive advisories from Google Cloud services (e.g., IAM Recommender, cost optimization, security hardening).
 - **Implementation**
 - Navigate to **Home > Recommendations Hub**. Review and assess recommendations, particularly those related to security.
- **Google Cloud Security Bulletins:** Actively monitor official Google Cloud Security Bulletins for advisories that may affect the Google Cloud services you use or require customer action.
 - **Implementation**
 - Subscribe to or regularly visit the Google Cloud Security Bulletins page <https://cloud.google.com/support/bulletins>.

Additional Considerations

- **Take Appropriate Action in Response:** Once an alert or advisory is received and assessed, timely and appropriate action is required.
- **Timeliness of Action:** Establish internal SLAs or guidelines for how quickly different types of alerts and advisories must be addressed, based on risk and potential impact (especially to CUI).
- **Prioritization:** Implement a process to prioritize alerts and advisories based on severity, exploitability, system criticality, and CUI impact.
- **Incident Response Integration:** Some alerts may indicate an active security incident. Ensure your monitoring processes are integrated with your incident response plan (see [IR.L2-3.6.1](#)).
- **Documentation:** Document all alerts/advisories received, the assessment performed, actions taken, timeliness of response, and outcomes. This is crucial for demonstrating compliance.
- **Verification:** After taking corrective action, verify that the action was effective in addressing the alert or advisory (e.g., the SCC finding becomes inactive, the vulnerability is no longer detected by VM Manager).

Supplemental Guidance

- [Enable finding notifications for Pub/Sub | Security Command Center | Google Cloud](#)
- [Security Command Center - Continuous Exports to Pub/Sub](#)
- [VM Manager - Viewing vulnerability reports](#)
- [Artifact Registry - Viewing and resolving vulnerabilities](#)
- [Security Bulletins | Customer Care | Google Cloud](#)

Control Domain	System and Information Integrity		
Control #	SI.L1-3.14.4		
Control Description	Update malicious code protection mechanisms when new releases are available		
Key Services	• SCC • Cloud IDS • Google Cloud Armor • VM Manager • Artifact Registry with Container Analysis • VM Manager • Custom Google Cloud Armor rules	Control Responsibility	☐ Google ☒ Shared ☐ Customer

Customer Implementation Description

Google Cloud Customers are responsible for:

- a. Updating malicious code protection mechanisms when new releases are available

When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control

- Security Command Center (SCC) (Event/VM/Container Threat Detection, Security Health Analytics)
- Cloud IDS (Threat signature sets)
- Google Cloud Armor (Managed Protection Plus / preconfigured WAF rule sets)
- VM Manager (Vulnerability Reporting - CVE database updates)
- Artifact Registry with Container Analysis (Vulnerability database updates)
- Host-based protection on Compute Engine VMs (Updates managed within the VM OS by customer; Google Cloud tools like OS Configuration management and Cloud Logging can assist).
- VM Manager (Patch Management) (For OS updates that may include security tool updates or support their effectiveness).
- Custom Google Cloud Armor rules (Customer reviews and updates these via the console).

A description of relevant features and implementation guidance is included below.

Leveraging Google-Managed Updates for Malicious Code Protection: For several Google Cloud-native security services, Google automatically handles the updates to the underlying threat intelligence and detection mechanisms. Customers ensure these services are enabled and properly configured to benefit from these updates.

- **Security Command Center (SCC) - Threat Detection & Security Health Analytics:** The detection logic, threat signatures, and vulnerability information used by SCC's services (Event Threat Detection, VM Threat Detection, Container Threat Detection, Security Health Analytics) are continuously updated by Google based on the latest threat intelligence.
 - **Implementation**
 - Ensure these services are enabled within SCC for your projects/organization (**Security > Security Command Center > Settings** and ensure relevant detectors are on).
 - By using these services, you are inherently benefiting from Google's ongoing updates to their detection capabilities. No direct customer action is needed in the console to "update" the threat intelligence itself.

- **Cloud IDS:** Cloud IDS uses threat signature sets managed and updated by Google (derived from Palo Alto Networks' threat intelligence).
 - **Implementation**
 - When creating or managing a Cloud IDS endpoint and its associated IDS policy (**Network Security > Cloud IDS**), you select a threat severity profile. Google automatically updates the signatures within these profiles.
 - Ensure your Cloud IDS endpoints are active and policies are appropriately configured to use these Google-managed signature sets.
- **Google Cloud Armor (Managed Protection Plus / Preconfigured WAF Rules):** If using Google Cloud Armor's Managed Protection Plus subscription or its preconfigured WAF rules (e.g., OWASP ModSecurity Core Rule Set), Google manages and updates these rule sets.
 - **Implementation**
 - When configuring a Cloud Armor security policy (**Network Security > Cloud Armor**), select and apply the relevant Google-managed preconfigured WAF rules.
 - Ensure your subscription (if applicable for advanced managed rules) is active. Google handles the updates to these rules.
- **VM Manager (Vulnerability Reporting) & Artifact Registry (Container Analysis):** The vulnerability databases (CVE information, etc.) used by VM Manager for OS scanning and by Container Analysis for image scanning are regularly updated by Google.
 - **Implementation**
 - Ensure VM Manager is active for your VMs and Container Analysis is enabled for Artifact Registry.
 - These services will automatically use the latest vulnerability information when performing their scans or re-evaluating existing resources/images.

Managing Updates for Customer-Deployed Malicious Code Protection Mechanisms:

For software or configurations customers deploy and manage themselves (e.g., AV/EDR on VMs, custom WAF rules).

- **Host-based Protection on Compute Engine VMs (e.g., AV/EDR - Customer Responsibility within OS):** This is the most significant area of customer responsibility for updates. Ensure any AV/EDR software installed on your Compute Engine VMs is configured to receive and apply regular updates for its software engine, signature files, and threat definitions from the software vendor.
 - **Implementation (Supported by Google Cloud tools)**

- **OS Configuration management**
 - Use **Compute Engine > OS Configuration management** to create OS policy assignments that can help ensure AV/EDR services are running and that update mechanisms (e.g., scheduled tasks for updates within the OS) are configured.
 - You could potentially use guest policies to verify the version of the AV/EDR agent if it can be queried via a script, and then report on outdated versions.
- **Cloud Logging & Monitoring**
 - Configure the Ops Agent on your VMs to collect logs from your AV/EDR software (if it logs to standard locations like syslog or Windows Event Log).
 - Create log-based alerts in **Logging > Log-based Alerts** to notify you if AV/EDR logs indicate update failures or significantly outdated definitions.
- **Startup/Shutdown Scripts (Less Ideal)** While possible to add update commands to VM startup scripts (**Compute Engine > VM instances > Edit VM > Automation**), relying on the AV/EDR's built-in, robust update scheduler is preferred.
- **VM Manager (Patch Management - Supporting Role):** While primarily for OS and common application patches, keeping the underlying OS patched via VM Manager ensures that the environment where your malicious code protection tools run is secure and that any dependencies they have are met. Some AV/EDR updates might also be bundled by OS vendors.
 - **Implementation**
 - Regularly apply OS patches using **Compute Engine > VM Manager > Patch deployments**.
- **Custom Google Cloud Armor Rules:** If you have created custom Cloud Armor rules, you are responsible for periodically reviewing and updating them based on new threat intelligence or changes in your application's risk profile.
 - **Implementation**
 - Navigate to **Network Security > Cloud Armor**.
 - Select your security policy and review/edit your custom rules.
 - This is a manual or process-driven update by the customer, using the console to implement the changes.

Additional Considerations

- **Define "New Releases":** Your policy should define what triggers an update (e.g., vendor notification of a new signature file, critical software update for AV/EDR, new WAF rule recommendation).

- **Automated Updates:** Prioritize configuring malicious code protection tools (especially host-based AV/EDR) for fully automated updates of signatures and engines whenever possible. Manual updates are prone to delays.
- **Testing Updates:** For significant software version updates to malicious code protection tools, consider testing them in a non-production environment first to ensure compatibility and no adverse operational impact.
- **Verification of Updates:** Implement a process to periodically verify that updates are being successfully applied across your systems. This could involve checking agent dashboards, reviewing logs, or using OS Configuration management to check agent versions.
- **Disconnected Systems:** If you have systems that are not always connected to the internet (less common in cloud but possible for certain architectures), ensure they have a mechanism to receive updates when they do connect, or a process for manual updates.

Supplemental Guidance

- [Security Command Center Overview](#)
- [Cloud IDS Overview](#)
- [Google Cloud Armor preconfigured WAF rules overview](#)
- [OS Configuration management](#)
- [Ops Agent - Logging](#)
- [VM Manager - Patch Management](#)

Control Domain	System and Information Integrity		
Control #	SI.L1-3.14.5		
Control Description	Perform periodic scans of organizational systems and real-time scans of files from external sources as files are downloaded, opened, or executed		
Key Services	• SCC Premium • VM Manager • Artifact Registry with Artifact Analysis • Google Cloud Armor • Cloud Storage	Control Responsibility	☐ Google ☒ Shared ☐ Customer
Customer Implementation Description			
Google Cloud Customers are responsible for: Defining the frequency for malicious code scans			

- b. Performing malicious code scans with the defined frequency
- c. Performing real-time malicious code scans of files from external sources as files are downloaded, opened, or executed

When configured or utilized correctly, the following feature(s) in Google Cloud Console may be used to support this control:

- Security Command Center (SCC) Premium
- VM Manager
- Artifact Registry with Artifact Analysis (Container Analysis)
- Web Security Scanner (part of SCC)
- Google Cloud Armor
- Cloud Storage with Cloud Functions and VirusTotal API integration (customer-built integration)
- Container Threat Detection (part of SCC Premium)

A description of relevant features and implementation guidance is included below.

Security Command Center (SCC) Premium: Provide centralized visibility and control over security findings, including vulnerabilities and misconfigurations.

- **Implementation (Periodic Scans):**
 1. **Activate SCC Premium:** Navigate to **Security > Security Command Center** in the Google Cloud Console. Select your organization and activate the Premium tier to access advanced features like Security Health Analytics, Web Security Scanner, and Container Threat Detection.
 2. **Security Health Analytics:** This service automatically and periodically scans your Google Cloud resources for common misconfigurations and vulnerabilities across services like Compute Engine, Cloud Storage, IAM, networking, etc. Findings are displayed in the SCC dashboard.
 - Review findings under **Security Command Center > Vulnerabilities** and **Security Command Center > Misconfigurations**.
 - Prioritize and remediate findings based on severity.
 3. Ensure necessary APIs are enabled for comprehensive scanning (e.g., Compute Engine API for VM-related checks).

VM Manager: Manage operating systems for large VM fleets running on Compute Engine.

- **Implementation (Periodic Patch/Configuration Scans)**
 1. **Enable VM Manager**
 - Navigate to **Compute Engine > VM Manager**.

- Enable the service for your project. This requires the OS Config API to be enabled and the OS Config agent to be installed and running on your VMs.
- 2. **Patch Scanning**
 - Access **VM Manager > Patch**. View patch compliance dashboards to see which VMs have missing patches.
 - Patch data is updated daily for connected VMs. This constitutes a periodic scan for patch vulnerabilities.
 - (Optional) Configure patch deployments to automate patching.
- 3. **OS Configuration Management**
 - Use OS policies (under **VM Manager > OS policies**) to define and enforce desired OS configurations (e.g., software installed, agent versions, specific file states). While not a direct vulnerability scanner, it helps maintain system integrity by periodically checking and enforcing desired states.

Artifact Registry with Artifact Analysis (Container Analysis): Scan container images for known vulnerabilities.

- **Implementation (Periodic Container Image Scans):**
 1. **Store Images in Artifact Registry** Push your container images to **Artifact Registry**.
 2. **Enable Vulnerability Scanning**
 - When creating an Artifact Registry repository or by editing an existing one, ensure that vulnerability scanning (powered by Artifact Analysis) is enabled. This is often on by default for new repositories.
 - Artifact Analysis automatically scans images upon push and can re-scan them periodically as new vulnerability information becomes available.
 3. **Review Findings**
 - View vulnerability findings in **Artifact Registry > Images > [Your Image] > Vulnerabilities**, or centrally in **Security Command Center > Vulnerabilities**.

Web Security Scanner (part of SCC Premium): Identify common web application vulnerabilities.

- **Implementation (Periodic Web App Scans):**
 1. Navigate to **Security Command Center > Web Security Scanner**.
 2. Click **Create scan**.
 3. Configure the scan:

- Specify starting URLs for your application (must be publicly accessible or accessible from Google's scanning infrastructure).
 - Choose authentication methods if required.
 - Set a schedule for periodic scans (e.g., weekly).
4. Review scan results in SCC under **Vulnerabilities**.

Google Cloud Armor: Provide DDoS protection and Web Application Firewall (WAF) capabilities at the edge.

- **Implementation (Filtering File Uploads - Edge Protection)**

1. Navigate to **Network Security > Cloud Armor**.
2. Create or configure a security policy.
3. Add rules to the policy:
 - Use preconfigured WAF rules (e.g., to block known malicious patterns or restrict file types).
 - Create custom rules based on request headers (e.g., Content-Type, Content-Disposition), request size, or other attributes to filter or block potentially malicious file uploads before they reach your backend services.
4. Attach the security policy to a backend service associated with your load balancer.

Cloud Storage with Cloud Functions and VirusTotal API (Event-driven File Scanning):

This is a customer-implemented solution using Google Cloud components for scanning files upon upload to Cloud Storage.

- **Implementation**

- **Develop a Cloud Function**
 - Write a Cloud Function (e.g., in Python, Node.js) that is triggered by object creation events in a specific Cloud Storage bucket (`google.cloud.storage.object.finalize`).
- **Integrate with VirusTotal API**
 - Within the Cloud Function, obtain a VirusTotal API key (VirusTotal is a Google service).
 - Use the API to submit the newly uploaded file (or its hash) to VirusTotal for analysis.
 - Process the scan report from VirusTotal.
- **Take Action** Based on the scan report, the Cloud Function can:
 - Move malicious files to a quarantined bucket.
 - Add metadata to the file indicating its scan status.

- Send notifications (e.g., via Pub/Sub to another alerting system or email).
- **IAM Permissions**
 - Ensure the Cloud Function's service account has permissions to read from the source bucket, write to a quarantine bucket (if used), and any other necessary services. <!-- end list -->
- *Note This provides a scan upon upload, not necessarily as files are "opened or executed" on an arbitrary VM, but it addresses "files from external sources as files are downloaded [into Cloud Storage]"*.

Container Threat Detection (part of SCC Premium): Detect common container runtime threats.

- **Implementation (Runtime Analysis):**
 1. Ensure SCC Premium is active and Container Threat Detection is enabled (typically enabled by default when SCC Premium is activated for a GKE environment).
 2. Ensure your GKE clusters are running a supported version and that telemetry is being sent.
 3. Monitors for activities such as unexpected binary execution, reverse shell, unexpected library load, or new privilege escalation.
 4. Review findings in **Security Command Center > Threats**.

Additional Considerations

- **Defining "Real-time":** For files on VMs, "real-time" typically implies an agent on the OS kernel hooks. Google Cloud does not offer a native Google agent for this. The event-driven Cloud Storage/VirusTotal approach is near real-time for new files arriving in a bucket.
- **Third-Party Solutions:** For comprehensive real-time endpoint protection (AV/EDR) on Compute Engine VMs, customers almost always deploy third-party security solutions. These can often integrate their findings into Security Command Center.
- **Defense in Depth:** No single tool is a silver bullet. Use a combination of preventative controls (e.g., IAM, firewalls, secure configurations) and detective controls (scanning, logging, monitoring).
- **Incident Response Plan:** Have a plan in place to respond to vulnerabilities or malicious files detected by these scanning mechanisms.
- **Logging:** Ensure appropriate audit logging (Admin Activity, Data Access, System Event) and VPC Flow Logs are enabled to support investigations. Chronicle Security Operations can be used for advanced threat detection and investigation using this telemetry.

Supplemental Guidance

- [Security Command Center overview | Google Cloud](#)
- [About VM Manager | Compute Engine Documentation | Google Cloud](#)
- [Artifact Analysis overview | Documentation | Google Cloud](#)
- [Overview of Web Security Scanner | Security Command Center | Google Cloud](#)
- [Product overview | Google Cloud Armor](#)
- [Trigger functions from Cloud Storage using Eventarc | Cloud Run Documentation | Google Cloud](#)
- [VirusTotal API v3 Overview](#)
- [Container Threat Detection overview | Security Command Center | Google Cloud](#)

Control Domain	System and Information Integrity								
Control #	SI.L2-3.14.6								
Control Description	Monitor organizational systems, including inbound and outbound communications traffic, to detect attacks and indicators of potential attacks								
Key Services	• VPC Flow Logs • Firewall Rules Logging • Cloud IDS • Google Cloud Armor • Packet Mirroring • SCC • Cloud Logging • Cloud Monitoring	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								
Customer Implementation Description									
Google Cloud Customers are responsible for: a. Monitoring the system to detect attacks and indicators of potential attacks b. Monitoring inbound communications traffic to detect attacks and indicators of potential attacks c. Monitoring outbound communications traffic to detect attacks and indicators of potential attacks When configured correctly, the following key service(s) in Google Cloud Console may be used to support this control: • VPC Flow Logs • Firewall Rules Logging • Cloud IDS (Intrusion Detection System) • Google Cloud Armor (for web traffic)									

- Packet Mirroring (for deep packet inspection with third-party tools)
- Security Command Center (SCC) (especially Event Threat Detection, VM Threat Detection, Container Threat Detection)
- Cloud Logging (Cloud Audit Logs, Ops Agent logs from VMs, application logs)
- Cloud Monitoring (Alerting on logs and metrics)

A description of relevant features and implementation guidance is included below.

Monitor Network Communications Traffic (Inbound and Outbound)

- **VPC Flow Logs:** Enable VPC Flow Logs for all subnets, especially those containing systems that handle CUI. These logs provide records of IP traffic to and from VM instances, which can be analyzed for anomalous patterns, connections to suspicious IPs, unexpected data transfers, or internal network reconnaissance.
 - **Implementation**
 1. Navigate to **VPC network > VPC Flow Logs**.
 2. For each subnet, click **Turn on** or edit settings to enable flow logs.
 3. Configure sampling rate, aggregation interval, and metadata as needed.
 4. Consider exporting flow logs to BigQuery for advanced analysis and long-term storage via a log sink in **Logging > Log Router**.
 5. Analyze logs for
 - Connections to/from known malicious IP addresses or unusual geographic locations.
 - Large, unexpected data transfers (especially outbound).
 - Internal port scanning or unusual east-west traffic patterns.
- **Firewall Rules Logging:** Enable logging for VPC firewall rules to record when rules are hit (both allow and deny actions). This helps detect attempts to access unauthorized ports, successful connections on allowed ports that might be part of an attack, or reconnaissance activity.
 - **Implementation**
 1. When creating or editing a firewall rule in **VPC network > Firewall**, select **On** for "Firewall rules logging."
 2. Review these logs in Cloud Logging (filter by logName "projects/YOUR_PROJECT_ID/logs/compute.googleapis.com%2Ffirewall").
 3. Look for
 - High volumes of denied traffic to specific ports (indicating scanning or attack attempts).
 - Unexpected allowed connections to sensitive services.

- **Cloud IDS (Intrusion Detection System):** Deploy Cloud IDS to monitor mirrored VPC traffic for known malicious signatures and network anomalies, providing direct detection of network-based attacks and indicators.
 - **Implementation**
 1. Navigate to **Network Security > Cloud IDS**.
 2. Create an IDS endpoint for each VPC network you wish to monitor, specifying the region and severity profile for threats.
 3. Configure Packet Mirroring policies (**Compute Engine > Packet Mirroring**) to forward traffic from selected subnets or instances to your Cloud IDS endpoint.
 4. Monitor Cloud IDS threat alerts, which are surfaced in Security Command Center and Cloud Logging.
- **Google Cloud Armor:** For web applications and services exposed via Google Cloud HTTP(S) Load Balancers, use Cloud Armor to detect and block web attacks (e.g., DDoS, SQL injection, XSS) and monitor inbound web traffic.
 - **Implementation**
 1. Navigate to **Network Security > Cloud Armor**.
 2. Create security policies and attach them to your load balancer backend services.
 3. Implement preconfigured WAF rules, custom rules, IP allow/deny lists, and rate limiting.
 4. Monitor Cloud Armor logs (available via Cloud Logging when enabled for the backend service) for blocked requests, WAF rule hits, and DDoS mitigation activity.
- **Packet Mirroring (with Third-Party Tools):** For more in-depth analysis than Cloud IDS provides, or to use existing third-party NIDS/NIPS or security analytics tools, use Packet Mirroring to send a copy of network traffic from your VMs to a collector instance running your chosen tools.
 - **Implementation**
 1. Set up collector instances (VMs running your NIDS/security tools).
 2. Create Packet Mirroring policies in **Compute Engine > Packet Mirroring**, specifying mirrored sources (subnets, instances, or network tags) and the collector ILB forwarding rule.

Monitor Organizational Systems (for Attacks and Indicators)

- **Security Command Center (SCC) - Threat Detection (Primarily Premium/Enterprise Tiers):** Enable and actively monitor SCC's threat detection services for near real-time identification of active threats and compromises within your Google Cloud environment.
 - **Implementation**

- Ensure SCC is activated (**Security > Security Command Center**) with Event Threat Detection, VM Threat Detection, and Container Threat Detection enabled.
 - **Event Threat Detection** Analyzes logs (Cloud Audit Logs, DNS, etc.) to find threats like malware C2, phishing, brute force, data exfiltration, IAM abuse.
 - **VM Threat Detection** Monitors VMs for runtime indicators like unexpected processes, kernel rootkits, and connections to malicious domains.
 - **Container Threat Detection** Detects runtime threats within GKE clusters.
 - Review findings in **SCC > Findings** and configure notifications (e.g., via Pub/Sub to SIEM/SOAR or ticketing systems) for immediate action on critical/high severity findings.
- **Cloud Logging (Audit Logs, Ops Agent Logs, Application Logs):** Collect and analyze logs from all system components to detect anomalous or malicious activity.
 - **Implementation**
 1. Ensure Admin Activity, Data Access (for CUI services), System Event, and Policy Denied audit logs are enabled and reviewed/alerted upon.
 2. Install and configure the Ops Agent on all Compute Engine VMs to collect OS-level logs (syslog, Windows Event Logs, security logs) and application logs. Forward these to Cloud Logging.
 3. Analyze these logs in **Logs Explorer** for
 - Repeated failed login attempts.
 - Unauthorized privilege escalation attempts.
 - Execution of suspicious commands or processes.
 - Unexpected system reboots or service failures.
 - Application-specific error patterns indicative of attacks (e.g., SQL injection patterns in web server logs).
- **Cloud Monitoring (Alerting):** Create alerts based on logs and metrics that indicate an ongoing attack or IoAs/IoCs.
 - **Implementation**
 1. In **Monitoring > Alerting**, create log-based alerts for specific patterns identified above from Cloud Logging.
 2. Create metric-based alerts for unusual system behavior, such as
 - Sustained high CPU/memory/network usage on VMs that doesn't correlate with expected load.
 - Anomalous number of active connections.
 - Deviations from baseline API call rates.

3. Ensure alerts are routed to your security operations team for investigation.

Additional Considerations

- **Define "Indicators of Potential Attacks":** Your organization should develop a list of IoAs/IoCs relevant to your environment, based on threat intelligence, system types, and CUI sensitivity. This will guide your monitoring queries and alert configurations.
- **Baselining:** Establish normal behavior baselines for your systems and network traffic. This makes it easier to identify anomalies that could indicate an attack.
- **Threat Intelligence Integration:** Consider integrating external threat intelligence feeds (e.g., lists of known malicious IPs or domains) into your monitoring processes (e.g., by creating firewall rules or Cloud Armor policies to block them, and alerting on any hits).
- **Incident Response Playbooks:** Have playbooks ready for responding to different types of detected attacks or indicators.
- **Regular Review and Tuning:** Continuously review and tune your monitoring configurations, alert thresholds, and detection rules to improve accuracy, reduce false positives, and adapt to evolving threats and changes in your environment.

Supplemental Guidance

- [Configure VPC Flow Logs | Google Cloud](#)
- [Firewall Rules Logging | Cloud NGFW](#)
- [Cloud IDS Overview](#)
- [Product overview | Google Cloud Armor](#)
- [Packet Mirroring | VPC | Google Cloud](#)
- [Overview of Event Threat Detection | Security Command Center | Google Cloud](#)
- [Security Command Center - VM Threat Detection](#)
- [Logging query language | Google Cloud](#)
- [Alerting overview | Cloud Monitoring](#)

Control Domain	System and Information Integrity								
Control #	SI.L2-3.14.7								
Control Description	Identify unauthorized use of organizational systems								
Key Services	• SCC • Cloud Logging • Cloud Monitoring • Cloud IDS	Control Responsibility	<table><tr><td>☐</td><td>Google</td></tr><tr><td>☒</td><td>Shared</td></tr><tr><td>☐</td><td>Customer</td></tr></table>	☐	Google	☒	Shared	☐	Customer
☐	Google								
☒	Shared								
☐	Customer								

	• IAM Policy Analyzer & Recommender • Google SecOps • BigQuery		
Customer Implementation Description			
Google Cloud Customers are responsible for: a. Defining authorized use of the system b. Identifying unauthorized use of the system When configured correctly, the following key service(s) in Google Cloud Console and related services may be used to support this control: • Security Command Center (SCC) Event Threat Detection, VM Threat Detection, Container Threat Detection, Security Health Analytics. • Cloud Logging Cloud Audit Logs (Admin Activity, Data Access, Policy Denied), VPC Flow Logs, Firewall Rules Logging, Ops Agent logs. • Cloud Monitoring Log-based and metric-based alerting. • Cloud IDS • IAM Policy Analyzer & Recommender • Chronicle Security Operations (Case Management, if used). • BigQuery (for analysis and creating records from exported logs/findings) A description of relevant features and implementation guidance is included below:. Identifying Unauthorized Use: Refer to the detailed guidance for SI.L1-3.14.2 for comprehensive information on configuring Cloud Logging, Cloud Monitoring, Security Command Center (including its various detectors like Event Threat Detection, VM Threat Detection), Cloud IDS, and IAM tools to <i>identify</i> unauthorized use. The core idea is to collect rich telemetry and use automated detection and alerting to flag suspicious activities and deviations from authorized behavior. Recording Identified Unauthorized Use: Once unauthorized use is identified, it must be thoroughly recorded. Google Cloud services provide the raw data and means to preserve this information. • Cloud Logging (Log Sinks and Retention for Record Keeping): Ensure that all relevant logs (Audit Logs, VPC Flow Logs, Firewall Logs, application logs, etc.) that provide evidence of or context to the unauthorized use are retained securely and for a sufficient period to support investigation, legal requirements, and potential reporting. ○ Implementation ■ Navigate to Logging > Log Router.			

- Create log sinks to export logs to
 - **Cloud Storage** For long-term, immutable archival of logs. Configure appropriate storage classes (e.g., Archive for cost-effectiveness) and bucket policies (e.g., retention policies, Bucket Lock if extreme immutability is needed).
 - **BigQuery** For detailed analysis, querying, and creating structured records of unauthorized events. This allows for easier aggregation and reporting.
 - Configure log bucket retention rules (**Logging > Logs Storage**) to meet your organization's record-keeping policies.
- **Security Command Center (SCC) - Findings as Records:** SCC findings themselves act as initial records of potential or confirmed unauthorized use detected by Google Cloud's native tools. Details within the finding are crucial for documentation.
 - **Implementation**
 - When a finding from services like Event Threat Detection or VM Threat Detection is confirmed as unauthorized use, ensure its details (JSON payload, affected assets, timeline) are captured or referenced in your primary incident tracking/case management system.
 - If a finding is muted (e.g., as part of a risk acceptance for a low-impact issue that is still technically "unauthorized" but accepted), the justification for muting serves as a record of that decision.
- **Google Security Operations (Case Management - if used):** If your organization uses Google SecOps, its case management capabilities are designed for recording details of security investigations, including those into unauthorized use.
 - **Implementation**
 - Create cases in SecOps for confirmed instances of unauthorized use. Link relevant UDM events, IoCs, investigation notes, and timelines within the case to build a comprehensive record.
- **BigQuery (for Aggregated Records and Analysis):** By exporting logs and SCC findings to BigQuery, you can create powerful queries to aggregate data related to an instance of unauthorized use, generate timelines, and produce structured data for incident records.
 - **Implementation**
 - Develop SQL queries in BigQuery to extract and format data relevant to specific incidents of unauthorized use. The results can be exported or used to populate incident report templates.

Reporting Identified Unauthorized Use (Facilitating Customer's Internal and External Reporting): Google Cloud tools can facilitate the customer's reporting process by providing

data and enabling automated notifications. The actual reporting to officials is a customer process.

- **Security Command Center (SCC) - Pub/Sub Notifications for Reporting Workflows:** Configure SCC to send finding notifications to Pub/Sub. This can trigger automated workflows to create tickets in your incident management system, notify relevant internal teams (security operations, legal, compliance), and start the internal reporting process for unauthorized use.
 - **Implementation**
 1. Navigate to **Security Command Center > Settings > Continuous Exports**.
 2. Create an export to a Pub/Sub topic for all findings or filtered findings (e.g., high/critical severity, specific categories like "Threat").
 3. Develop a subscriber application (e.g., a Cloud Function) or use a SIEM/SOAR integration to process messages from this Pub/Sub topic and initiate internal reporting procedures.
- **Cloud Monitoring Alerts - Direct Notifications:** Alerts configured in Cloud Monitoring for indicators of unauthorized use should be routed directly to personnel responsible for initiating the internal reporting and incident response process.
 - **Implementation**
 1. When creating alert policies in **Monitoring > Alerting**, configure notification channels (e.g., email distribution lists for security teams, PagerDuty, Slack) that ensure responsible parties are immediately notified.
 2. The alert notification itself serves as an initial internal report.

Additional Considerations

- **Define "Unauthorized Use":** Ensure your organization has clear, documented definitions of what constitutes unauthorized use of systems and data, as this underpins all identification, recording, and reporting efforts.
- **Incident Response Plan Integration:** The processes for identifying, recording, and reporting unauthorized use must be integral components of your overall incident response plan (see [IR.L2-3.6.1](#) and [IR.L2-3.6.2](#)).
- **Timeliness:** Establish SLAs for how quickly identified unauthorized use must be recorded and reported internally. External reporting timelines are often dictated by regulations or contractual agreements.
- **Chain of Custody for Evidence:** If digital evidence from Google Cloud logs or findings is required for legal or disciplinary action, ensure processes are in place to maintain its integrity (e.g., using immutable storage for exported logs, documenting an Ediscovery hold if necessary via Google Vault for Workspace data).

- **Legal and Regulatory Obligations:** Be aware of all legal, regulatory (e.g., data breach notification laws), and contractual (e.g., DoD reporting requirements for incidents involving CUI) obligations for reporting unauthorized use.

Supplemental Guidance

- [Route logs to supported destinations | Cloud Logging](#)
- [Enable finding notifications for Pub/Sub | Security Command Center | Google Cloud](#)
- [Google SecOps Search API | Google Security Operations](#)
- [BigQuery documentation | Google Cloud](#)
- [Cloud Security and Data Protection Services | Google Workspace](#)

Appendix A Key terms, acronyms & definitions

Term	Definition
Assured Workloads	Assured Workloads provides Google Cloud users with the ability to apply controls to a folder in support of regulatory, regional, or sovereign requirements.
CMMC	The Cybersecurity Maturity Model Certification (CMMC) is a program which verifies that companies in the Defense Industrial Base (DIB) have adequate cybersecurity safeguards in place
CUI	Controlled Unclassified Information (CUI) is a category of federal information that is not classified but still requires protection to ensure national security
CUI Boundary	The Controlled Unclassified Information (CUI) boundary refers to the defined organizational systems and facilities in which CUI is stored, processed, and transmitted

Customer	Users of Google Cloud, including but not limited to the DCMA, members of the DIB, federal contractors and subcontractors, OSAs, etc.
DCMA	The Defense Contract Management Agency (DCMA) is a federal agency that manages contracts for the Department of Defense and other federal agencies
DFARS	The Defense Federal Acquisition Regulation Supplement (DFARS) is a set of regulations that apply to a United States Department of Defense (DoD) contracts
DIB	The Defense Industrial Base (DIB) is a network of organizations, facilities, and resources that support the United States military with services, resources, and equipment
FAR	The Federal Acquisition Regulation (FAR) is a set of regulations that govern how the federal government provisions goods and services
NIST	The National Institute of Standards and Technology (NIST) is a United States federal agency that develops measurement standards, including several cybersecurity frameworks such as NIST SP 800-171 and NIST SP 800-53.
OSA	Organizations seeking assessment (OSAs) are organizations aiming to comply with CMMC and/or organizations considering bidding on future Department of Defense (DoD) contracts with CMMC obligations.