Google Cloud

Google Cloud minimum viable secure platform

Checklist

Authentication and authorization

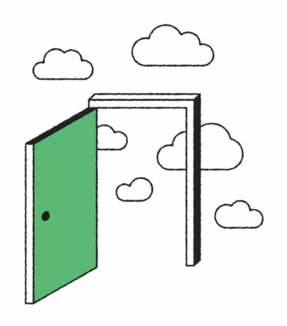

Identity management

	Guidance	Why
Basic		
☑	Decide on your source of truth for provisioning managed user identities. Patterns include creating user identities in Cloud Identity, syncing identities from an existing identity provider, or using Workforce Identity Federation.	This might be more a question of onboarding than a single checklist item, but it's a prerequisite to many of the other identity-related checklist items.
☑	Do not have a single super admin or org admin – select between 2 and 20.	Having a single super or org admin oftentimes results in lockout scenarios, and it carries a higher risk as one person is able to make platform-altering changes, potentially with no oversight.
☑	Enforce strong and unique passwords for all user accounts. Consider using a password manager.	Weak or no credentials are still a common pattern that attackers exploit therefor ensuring that strong and unique passwords are used.
Intermediate		
☑	Set organization-wide policies that prevent adding external members to Google Groups.	By default, external user accounts can be added to groups in Cloud Identity. We recommend that you configure sharing settings so that group owners can't add external members. Note that this restriction doesn't apply to the super admin account or to other delegated administrators with Groups admin permissions. Because federation from your identity provider runs with administrator privileges, the group sharing settings don't apply to this group synchronization. We recommend that you review controls in the identity provider and synchronization mechanism to ensure that non-domain members aren't added to groups, or that you apply group restrictions.
☑	Set session length for Google Cloud services to expire at least every day.	Leaving an account signed in for an extended period is a security risk. Enforcing a maximum session duration automatically ends the session after a set time, forcing a new, secure sign-in. This step significantly limits the window of opportunity for an attacker to use a stolen password and ensures access is regularly reverified.

Identity management

Guidance	Why
Intermediate	
☑ Don’t allow use of unmanaged consumer accounts. Consolidate any unmanaged consumer accounts, and consider a solution to prevent the creation of further unmanaged consumer accounts with your domain.	Unmanaged consumer accounts are not governed by your joiner-mover-leaver (JML) processes, so they introduce the risk that an employee still has access to your resources after they leave their job. These accounts are also treated as external with regard to controls like domain restricted sharing.
☑ Super admin accounts are separate from day-to-day user accounts. They’re dedicated and to be used only when making critical changes. Turn on multiparty approval for admin actions.	Super admin accounts should not be used for day-to-day activities, nor should they be associated with other enterprise activities – like email. Turning on multiparty approval means sensitive actions have to be approved by two admins, thus potentially preventing cases where a malicious user is able to compromise an admin account and lock out other admin users.
☑ Enable multi-factor authentication (MFA) for all Google Accounts and Cloud Identity users, not just admins. Super admins are enabled by default.	Passwords alone are not a strong enough measure therefor MFA is essential to add another layer of defense.
Advanced	
☑ Disable account self-recovery for super admin accounts.	Super admin account self-recovery is off by default for all new customers – existing customers might have this setting on. Turning this setting off helps to mitigate the risk that a compromised phone, compromised email, or social engineering attack could let an attacker gain super admin privileges over your environment. Plan an internal process for a super admin to contact another super admin in your organization if they have lost access to their account, and ensure that all super admins are familiar with the process for support-assisted recovery.
☑ Set idle session timeout to 15 minutes for sensitive use cases.	Idle sessions may be used by attackers for credential theft. Customers like UPS have asked for this feature.
☑ Provide hardware security keys, if possible, to super admins or org admins as a 2nd factor.	While MFA is a critical and essential defense, admin accounts are the highest-value targets for sophisticated attacks. Hardware security keys provide an even higher level of protection because they are phishing-resistant. This makes them the strongest possible defense against account takeover for your most critical admin users, building upon your standard MFA policy.

Access management

Guidance	Why
Basic	
☑ Use Identity and Access Management (IAM) job functions to assign permissions based on known job functions. And if subset is needed, use as reference for custom role.	Job functions are predefined roles that allow admins to easily provide a set of permissions that is limited to a job function, thus improving productivity and reducing the back-and-forth of asking for permissions.
Intermediate	
☑ Remove the domain-wide Project Creator and Billing Account Creator roles that are granted by default to all members in a new organization.	New organizations grant the Project Creator and Billing Account Creator roles to all managed user identities in the domain. While this is useful for getting started, it is not best practice. Allowing the profileration of billing accounts leads to messy overhead and has technical consequences for splitting services across multiple Billing Accounts. Allowing free-form project creation can lead to projects that don't adhere to your governance conventions. Instead, you should remove these roles and establish a project factory process to request new projects and associate them with billing.
☑ Use Privileged Access Manager (PAM) for managing privileged access. • For all other access, use access groups • Let group memberships expire automatically • Implement an approval workflow for group memberships	Using the least privilage model allows customers to only provide access when needed, for the resources needed. Using pre-built roles provides ease of use and reduces sprawl using custom roles customers do not have to worry about managing life cycle.
Advanced	
☑ If using an external identity provider, set up post-SSO verification.	We recommend that you enable an additional layer of control based on Google's sign-in risk analysis. After you apply this setting, users might see additional risk-based login challenges at sign-in if Google deems that a user sign-in is suspicious.
☑ Enable principal access boundaries to limit principal access and protect against phishing and data exfiltration. Our default recommendation is to enable a boundary for the organization to avoid external phishing attacks.	Principal access boundaries improve security greatly by reducing the extent of an attack with any compromised identity, and they also prevent any external phishing attacks and other exfiltration attacks.

Credential management

Guidance	Why
Basic	
☑ Disallow the use of service account keys, except for cases where there are no viable alternatives.	Before deciding to use service account keys, make sure to consider all possible alternatives. Use organizational policy constraints to disable service account key creation and service account key upload by default, and allow the use of service account keys on an exceptional basis only.
Intermediate	
☑ Limit over-permissioning of service accounts by setting the iam.automaticIamGrantsForDefaultServiceAccounts constraint to false.	By default, some systems grant overly broad permissions to automated accounts, which is a potential security risk. If an attacker compromises a single part of the system, they could gain control over the entire project. This policy constraint disables those automatic, high-level permissions, forcing a more secure, deliberate approach where only the minimal necessary permissions are granted.
☑ If service account keys must be used, then ensure keys are rotated at least every 90 days.	A rotation interval limits how long an adversary can have access to the system. Without a rotation interval, the adversary has access forever. Workload Identity Federation is the strong control – the checklist item here is for exceptions when a service account key needs to be used.
☑ Use Workload Identity Federation to let CI/CD systems and workloads running on other clouds authenticate to Google Cloud.	Workload Identity Federation lets workloads that run outside of Google Cloud authenticate without requiring a service account key. By avoiding service account keys and other long-lived credentials, Workload Identity Federation can help you reduce the risk of credential leakage.

Organization

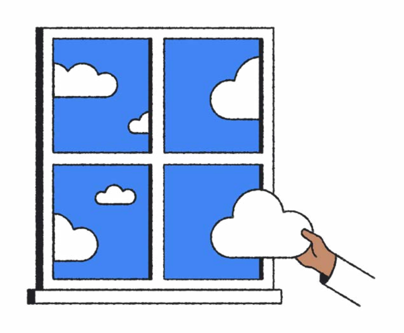

Resource management

Guidance	Why
Intermediate	
☑ Ensure only identities from your organization are allowed in your Google Cloud environment by enabling the constraints/iam.managedallowedPolicyMembers org policy or by creating a custom attribute with it.	This policy prevents employees from granting access to external accounts outside of your organization's control that do not follow your security policies for MFA or password management. This control is critical for preventing unauthorized access, ensuring that only trusted, managed corporate identities can be used.
Advanced	
☑ Consider setting location constraints on resources that can be enforced by the constraints/gcp.resourceLocations organization policy constraint.	This policy allows your organization to enforce that your resources and data are only created and saved in specific, approved geographic regions.
☑ Consider creating a constraint list of allowed API and service that can be enforced by the constraints/gcp.restrictServiceUsage organization policy constraint.	This policy allows your organization to create an allowlist of only the services it has approved. This prevents employees from using unvetted services.

Infrastructure

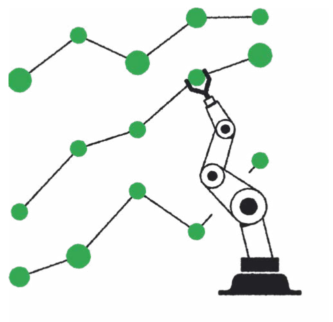

Compute resource management

Guidance	Why
Basic	
☑ Turn off serial port access by setting the compute.disableSerialPortAccess organization policy constraint.	You should disable serial port access on your Compute Engine VMs to eliminate a powerful access channel that completely bypasses your firewall rules and other network security controls. The interactive serial console is primarily intended for emergency troubleshooting, but leaving it enabled creates a persistent potential backdoor that can be targeted by attackers. Disabling it enforces a defense-in-depth security posture by forcing all administrative access through standard, audited pathways like SSH, which are protected by IAM and Identity-Aware Proxy (IAP).
☑ Disable IPv6 external subnet creation unless specifically required.	You should consider disabling IPv6 on systems and networks where it's not actively managed or required, as this reduces your overall attack surface. Many organizations have mature security controls and monitoring for IPv4, but their tools and policies may not fully extend to IPv6, creating a significant blind spot for threats. Running a dual-stack network also introduces operational complexity, requiring specific configurations and expertise to manage and troubleshoot effectively. Therefore, if you don't have a clear business driver for IPv6, disabling it can simplify your environment and ensure all traffic is consistently filtered through your established IPv4 security posture.
☑ Turn on vTPM and integrity monitoring attributes of Shielded VMs for your instances – now preselected defaults.	You should use the vTPM (virtual trusted platform module) and integrity monitoring attributes of Shielded VMs to ensure your virtual machines boot only with trusted, unmodified code. The vTPM provides a secure, virtual cryptoprocessor that generates and stores cryptographic measurements of the entire boot sequence, from the UEFI firmware to the kernel drivers. Integrity monitoring then continuously compares these runtime measurements against a known-good baseline established when the VM was first created. This provides a verifiable chain of trust, automatically alerting you or taking action if it detects any malicious modifications, like those from a bootkit or rootkit, ensuring your workload's integrity from the moment it powers on.

Compute resource management

	Guidance	Why
Intermediate		
☑	Use OS Login to easily manage SSH keys with IAM policies by setting the compute.requireOsLogin organization policy constraint.	OS Login centralizes VM access by tying SSH permissions to a user's Google identity and IAM roles, eliminating the need to manage individual SSH keys on each machine. This is crucial for security because removing a user's IAM role instantly revokes their access across all instances, protecting against unauthorized entry from stale accounts. The system simplifies key management to prevent dangerous "key sprawl" and provides a clear, centralized audit trail for all login events in Cloud Audit Logs. Furthermore, you can enforce two-factor authentication, adding a critical layer of protection against stolen SSH keys and credentials.
☑	Unless needed, prevent the creation of Compute Engine instances with public IP addresses by default by setting the compute.vmExternalIpAccess organization policy constraint.	You should prevent Compute Engine instances from having public IP addresses to drastically reduce their exposure to the public internet. Any instance with a public IP is immediately discoverable and becomes a direct target for automated scans, brute-force attacks, and attempts to exploit vulnerabilities. Instead, you should require instances to use private IPs and manage access through controlled, authenticated, and logged pathways like the Identity-Aware Proxy (IAP) tunnel or a bastion host. Adopting this "deny by default" posture is a foundational security best practice that minimizes your attack surface and enforces a zero-trust approach to your network.

Container resource management

	Guidance	Why
Basic		
☑	Use GKE Autopilot clusters.	Autopilot clusters offer robust security measures, with many container or Google Kubernetes Engine (GKE) security best practices enabled by default.
☑	Use least privilege IAM service accounts for nodes.	Access to the GKE control plane is restricted to a single DNS-based endpoint, significantly reducing the attack surface without the need for additional firewall rules or bastion hosts.
☑	Restrict network access to the control plane using a DNS-based endpoint.	The control plane is the management center for a Kubernetes cluster, and exposing it to the public internet makes it a prime target for attackers. This setting makes it private, removing it from the public internet entirely. This ensures that only trusted devices within your organization's private network can manage the cluster, drastically reducing the risk of an external attack.
☑	Prefer the usage of Container-Optimized OS for a hardened and managed container OS.	General-purpose operating systems include many extra programs that are not needed to run containers, creating a larger, unnecessary target for attackers. Container-Optimized OS is a minimal, locked-down operating system that significantly reduces this attack surface by including only what is necessary. As managed OS, Container-Optimized OS also has security patches automatically applied by Google, which ensures critical vulnerabilities are fixed and reduces your operational workload.
Intermediate		
☑	Use Workload Identity Federation for GKE to authenticate to Google Cloud APIs from GKE workloads.	Simplest and safest way to obtain identities to call Google Cloud APIs.
☑	Create private nodes to reduce internet exposure.	Reduces internet exposure by ensuring GKE nodes do not have a public IP.
☑	Use Google Groups for role-based access control (RBAC), which also lets you integrate with your existing user account management practices, such as revoking access when someone leaves your organization.	Facilitates efficient management of cluster access using IAM and Google Groups, ideal for medium to large organizations that use Google Groups.

Container resource management

	Guidance	Why
	Advanced	
☑	Use GKE Sandbox to provide an extra layer of security to prevent untrusted code from affecting the host kernel on your cluster nodes.	Enhances workload isolation for untrusted or sensitive workloads, providing an additional layer of protection against container escape attacks.
☑	Use Binary Authorization to make sure trusted images are deployed to Google Kubernetes Engine.	Ensures that only verified and trusted container images are allowed to be deployed in your clusters, strengthening software supply chain security.
☑	Use Confidential GKE Nodes to enforce encryption of data in-use in your nodes and workloads.	Secures highly sensitive workloads by encrypting data in use through confidential computing.
☑	Run your own certificate authorities and keys in GKE.	Users can manage their own certificate authorities and keys within GKE, offering greater control over cryptographic operations. To request access to this feature, contact your Google Cloud account team.
☑	Encrypt Kubernetes Secrets at rest using Cloud Key Management Service (Cloud KMS) managed keys.	Provides an additional layer of security for etcd data by allowing you to encrypt Kubernetes Secrets with a key you control.
☑	Use customer-managed encryption keys (CMEK) for node boot disk encryption.	Allows you to encrypt a Kubernetes node’s boot disk with a key that you manage.

Data protection

Storage protections

Guidance	Why
Basic	
☑ Ensure that all of your storage buckets have consistency with regard to access policy by enabling the storage.uniformBucketLevelAccess organization policy constraint.	Using two different and conflicting systems to manage permissions on storage buckets is complex and a common cause of accidental data leaks. This setting turns off the legacy system (access control lists, or ACLs) and makes the modern, centralized system (IAM) the single source of truth for all permissions.
Intermediate	
☑ Prevent storage buckets from being accessed from public sources without authentication by enabling the storage.publicAccessPrevention policy.	Accidentally making private data public in a storage bucket is one of the most common and severe causes of data breaches. This policy acts as an organization-wide safety net that actively blocks any setting that would make a bucket publicly accessible.

Encryption

Guidance	Why
Advanced	
☑ Consider your encryption management strategy and use either Cloud KMS with Autokey, Cloud External Key Manager (Cloud EKM), or a combination of these that is suitable for your strategy.	This control allows your organization to use and manage its own keys to meet your specific requirements. This provides granular, auditable control over data access, including the ability to immediately block access to data by disabling the key.

Database security

	Guidance	Why
	Basic	
☑	Prevent Cloud SQL from having a public IP address and being directly exposed to the internet by setting the sql.restrictPublicIp organization policy constraint.	Databases normally should not be directly exposed to the public internet. This policy prevents your databases from getting public IPs, ensuring they are private and only accessible from trusted, internal applications.
	Intermediate	
☑	Ensure that BigQuery does not have datasets open to public access unless intended.	Datasets in BigQuery often contain sensitive data. This control helps you ensure that you do not accidentally or unitentionaly expose data to the public internet.

Network security

Routing and networks

Guidance	Why
Basic	
☑ Skip default network creation by enabling the compute.skipDefaultNetworkCreation organization policy value.	The default network is an auto mode Virtual Private Cloud (VPC) network with pre-populated IPv4 firewall rules to allow internal communication paths. Generally, this is not a great security posture for production.
☑ Enable Private Google Access on all subnets.	Enabling Private Google Access is a necessary step for accessing Google Cloud services without a public IP. By default, it is not enabled on new resources and requires additional steps to explicitly enable it.
Intermediate	
☑ For applications exposed behind Cloud Load Balancing, use Google Cloud Armor default policies or configure your own policies to add Layer 3 to Layer 7 network protection for externally facing applications or services.	Google Cloud Armor security policies protect your application by providing Layer 7 filtering and by scrubbing incoming requests for common web attacks or other Layer 7 attributes to potentially block traffic before it reaches your load-balanced backend services or backend buckets. Each security policy is made up of a set of rules that can be configured on attributes from Layer 3 through Layer 7.
☑ Configure your firewalls and DNS settings to route egress traffic to Google Cloud APIs using either the private VIP or restricted VIP.	By default, traffic bound for Google services uses the default domains, which include all Google services. This introduces a few problems: • Using default domains requires a very broad egress firewall rule, and it is difficult to create other firewall rules to deny egress traffic without understanding the full range of Google default domains • A Google Cloud-only customer who thinks they have blocked internet egress to web services might be concerned when they realize their settings allow egress to non-Google Cloud exfiltration paths like Google Drive
Advanced	
☑ Enable VPC Service Controls as an additional layer of protection to prevent potential data loss.	VPC Service Controls can help prevent data exfiltration by creating isolation perimeters around your cloud resources, sensitive data, and networks.

Firewall management

Guidance	Why
Intermediate	
☑ Limit access to external sources, as by default all access is allowed out. Set specific firewall rules for intended patterns of traffic needing to egress.	By default, systems are often allowed to make outbound connections to the internet, which can be deemed security risk. A “deny by default” policy blocks all outbound traffic and requires specific rules to be made for only the known, necessary destinations.
☑ Limit inbound access to resources and, where possible, restrict to specific resource ranges. If Identity-Aware Proxy (IAP) is configured, inbound SSH and Remote Desktop Protocol (RDP) firewall rules should be set for IAP IP ranges as sources.	Permissive SSH and RDP firewall rules allow for brute force attacks. Instead, Google Cloud identity-aware proxies (such as IAP) for SSH and RDP should be used.

Monitoring, logging, and alerting

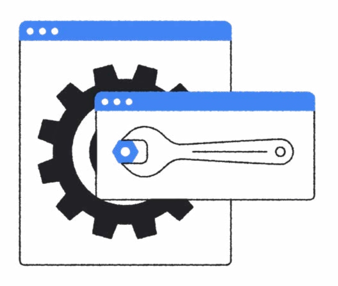

Setting alerts

	Guidance	Why
	Basic	
☑	Subscribe to security bulletin notifications.	Customers can subscribe to security bulletins related to Google Cloud products so they are notified of vulnerabilities and mitigation measures.
☑	Configure essential contacts to ensure you receive important notifications by using a monitored group alias or mailing list.	Google sends critical security alerts (like a potential account compromise) to the email addresses listed as "Essential Contacts." If an individual's email is used for this purpose, the alert will be missed if that person is unavailable or has left the company. Using a monitored group email address ensures these time-sensitive alerts are delivered to an active team that can respond quickly.
☑	Use the billings anomaly feature in Cloud Billing to track any spikes or deviations in expected spend.	A sudden, unexpected spike in a cloud bill is a primary indicator of a security compromise. This is sometimes caused by an attacker who has gained access and is using resources for unauthorized activities. Enabling billing anomaly detection provides an essential early warning system by automatically flagging this suspicious activity.

Logging

	Guidance	Why
	Intermediate	
☑	Turn on Firewall Rules Logging.	By default firewall rules do not automatically log. Firewall Rules Logging allows you to audit, verify, and analyze the effects of your firewall rules. For example, you can determine if a firewall rule designed to deny traffic is functioning as intended. Logging is also useful if you need to determine how many connections are affected by a given firewall rule.
☑	If using Cloud Identity, share audit logs from Cloud Identity to Google Cloud.	Admin Activity audit logs from Google Workspace or Cloud Identity are ordinarily managed and viewed in the Google Admin console, separately from your logs in your Google Cloud environment. These logs contain information that is relevant for your Google Cloud environment, such as user login events. We recommend that you share Cloud Identity audit logs to your Google Cloud environment to centrally manage logs from all sources.

Logging

Guidance	Why
Advanced	
☑ Ensure Access Transparency is turned on.	Standard logs show you what your organization’s own users are doing, but Access Transparency logs show what Google support staff do when they access the account. This access typically only happens in response to a support request. This provides a complete and verifiable audit trail of all access, which is essential for meeting strict compliance and data governance requirements.
☑ Create a log sink to export logs for your security monitoring solution and set the retention period to meet your requirements.	The default log retention periods are often not long enough to meet the 1–7 year requirements mandated by compliance rules (like PCI or HIPAA). Creating a log sink to export logs to a long-term storage location is essential for meeting these legal and regulatory obligations. This also allows you to send logs to a centralized security monitoring system for advanced threat detection.