

Mitigating data exfiltration risks with Google Cloud

February 2025



Google Cloud

Data is a valuable asset for any organization. In this paper, we discuss the concept of data exfiltration and how data theft is a risk for organizations of all sizes. We then cover some of the most common attack vectors that lead to data exfiltration, as well as Google Cloud security controls that can be used to mitigate these attacks.

Data exfiltration overview

Data exfiltration can be defined as a form of data theft, where a malicious actor gains unauthorized access to corporate data and transfers (or copies) this data to a target only accessible to the malicious actor. The victims of data exfiltration attacks can be a single individual or an enterprise organization. Exfiltration attempts against an organization typically involve large amounts of sensitive data being targeted, such as credential data (i.e. employee/customer authentication information, cryptographic keys, etc), personally identifiable information (i.e. employee/customer emails, home addresses, government identification numbers, etc), company trade secrets, personal financial information (i.e. employee/customer financial records, etc), and other similar sensitive data.

Data exfiltration is the intended outcome of a variety of cyberattacks such as credential theft, phishing, and insider risks to name a few. In general, threats that can lead to data exfiltration can be categorized into external threats and internal threats. External threats are initiated by a malicious actor that does not have any affiliation with the intended target - these sort of attacks are predicated on the attacker first gaining access to this data and only then successfully exfiltrating it. Internal threats are typically initiated by an internal actor (i.e. malicious employee, malicious actor who already infiltrated your environment) who already has access to the data and thus must only circumvent security measures to exfiltrate this data.



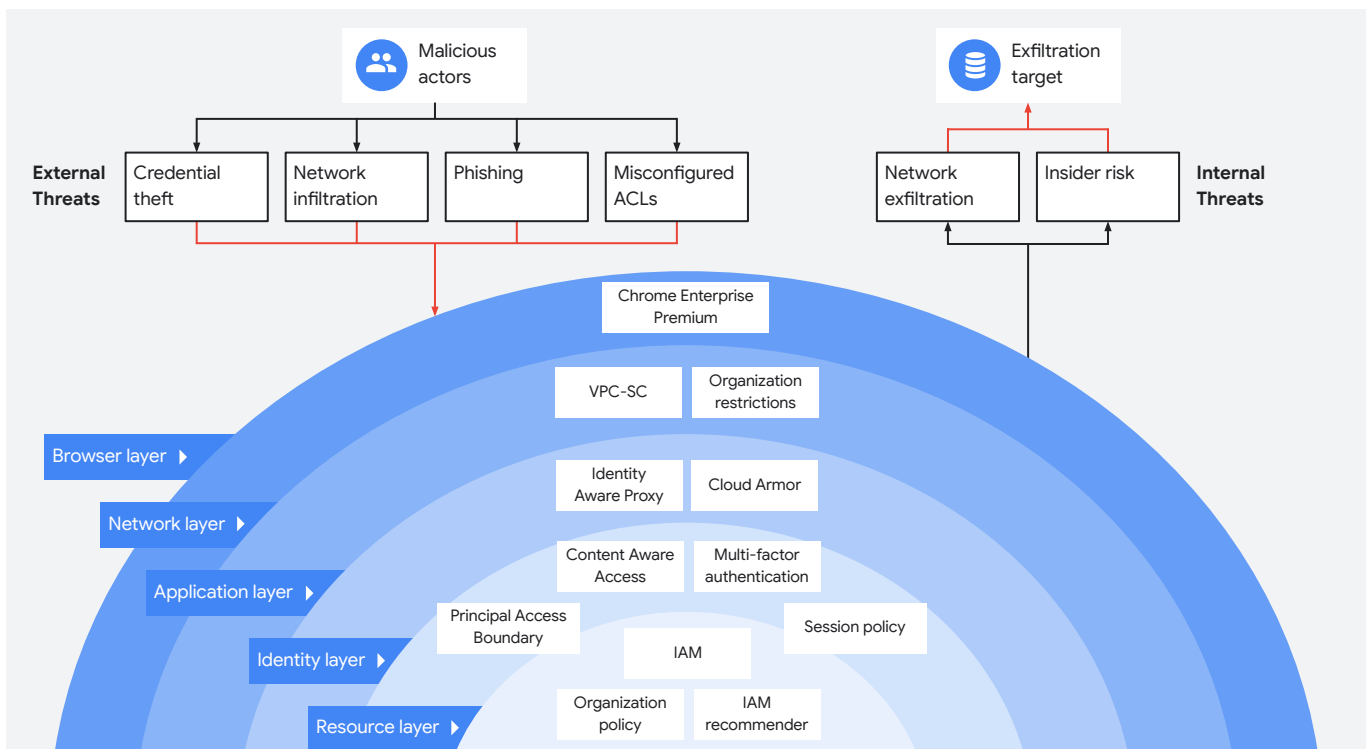
Detecting data exfiltration after a successful attack is difficult and often too late for any meaningful resolution to the attack. As such a defense-in-depth security posture must be established to proactively mitigate data exfiltration risks. Google Cloud offers a variety of cloud-native security controls, and best-practice recommendations to help you establish a strong security posture that proactively mitigates data exfiltration risks.

The diagram below visualizes a variety of external and internal threat vectors which can lead to data exfiltration and the various Google Cloud security controls that can be used to mitigate these attacks.

Subsequent sections of this paper will cover additional details about these attacks as well as how Google Cloud security controls can protect your valuable data from these attacks. While these attacks are covered independently it's important to be aware that often, a successful data exfiltration event will be composed of multiple chained attack types (as referenced in the [cyber kill chain](#) framework).

For example, a malicious actor can employ social engineering techniques to identify enterprise administrators working at the company they are looking to infiltrate. Once identified, the malicious actor can employ phishing techniques to capture the enterprise administrator's credentials. Using these credentials, the malicious actor gains access to the company's cloud environment and infiltrates the production network hosting a highly sensitive application. There, the attacker can set up a long-running job that exfiltrates data from this network to an external target.

As such this paper details a variety of threats, some of which do not individually result in data exfiltration (i.e. network infiltration, credential theft), but are often the first entrypoint in a sequence of attacks that eventually lead to data exfiltration. So while it's important to thoroughly understand each individual threat type and the controls available to mitigate it, it's just as important to consider how a defense in depth strategy can be used to intercept the attack across the various security layers within your infrastructure.



Credential theft



Credential theft is the act of stealing personal information such as usernames, passwords, cookies, SSH keys, etc. The outcome of credential theft is often unauthorized data access and subsequent data exfiltration. Typically credential theft attacks are initiated by external malicious actors in order to gain access to sensitive data. Malicious actors employ a variety of techniques to obtain user credentials such as phishing (i.e. email based phishing with fake login page), cross-site scripting or man-in-the middle attacks for cookie thefts, brute force attacks (i.e. continuous trial and error attempts to guess weak credentials), hardware based attacks such as USB keyloggers (i.e. malicious attacker disposes of infected USB sticks with software that tracks all keyboard strokes when plugged in), and many others.

Password strength requirements or passkeys are now a common security feature across the internet. Combined with the proliferation of user account credentials for website access, the average individual is expected to remember many distinct combinations of usernames and passwords. This can result in the use of a password manager to store and retrieve these credentials. This credential source becomes a single point of failure, and is ripe for credential attack opportunities. Without additional security guardrails on these credentials, a malicious actor that manages to gain access to this credential source can now begin orchestrating an attack using these stolen credentials.

While credential theft is often discussed in the context of user identities, these types of attacks can be just as effective in stealing machine identity credentials. Machine identities are used to authenticate devices within an organization's network including servers, laptops, phones, etc. In the context of public cloud, machine identities are often used to uniquely identify a cloud resource such as a Virtual Machine or containers. The compromise of such a machine identity can be used by a malicious actor to successfully exfiltrate data out of these cloud resources.

For example, Google Cloud uses Service Accounts to assign cloud resources a machine identity. Google Service Accounts employ a variety of credentials, one of which are [service account keys](#). Typically [Google discourages the use of service account keys](#), however these are widely used by Google Cloud customers as they are simple to use. These keys can be used to authenticate and make requests against Google Cloud resources. These keys are often hardcoded within an application's code during the development process, and sometimes developers upload their application's code to public code repositories without removing the Service Account Keys embedded within the code. A malicious actor can stumble upon or search for these types of credentials across public code repositories in an attempt to use these credentials to orchestrate a successful data exfiltration attack. Google Cloud now [actively searches through popular code repositories for instances of these credentials](#) and automatically disables services accounts linked to these credentials in order to mitigate risks associated with credential leaks.



1

Mitigating credential theft with multi-factor authentication

Multi-factor authentication is strong protection against credential theft. Malicious actors that manage to steal passwords through the various techniques described above will not be able to get access using stolen credentials as they will lack the additional authentication requirements enabled on the respective identity.

Google Cloud customers who use Google as their primary Identity Provider should [enable multi-factor authentication directly through Cloud Identity](#).

Google Cloud customers who employ identity federation with a third party identity provider (through [identity synchronization and single sign-on](#)) should also [enable multi-factor authentication on these federated identities](#) directly through Cloud Identity. When Google Cloud customers use Cloud Identity with their own Identity Provider, through SAML or OIDC, Cloud Identity queries the provider for an attestation when the session expires or when Google requires a reauthentication. In the default configuration, Identity Providers silently approve all these attestations. However, Identity Providers can generally be configured to always require re-entering credentials, or always require multi-factor authentication whenever Google requests an attestation. This configuration can be set up to only apply to the application representing Google, and not for all applications that the Identity Provider federates.

2

Mitigating credential theft with Context Aware Access

Context-Aware Access is a feature of Google Cloud which allows for attribute based access control to Google Cloud resources, APIs, and the Cloud Console.

[Access Context Manager](#) enables Google Cloud organization administrators to define fine-grained, attribute-based access control rules through Access Levels. Access levels describe additional [device, user, or request attributes](#) to be met in order for a resource request to be successful.

An example of such a device attribute is a verified X.509 certificate in order to enforce [certificate-based access restrictions](#). Certificated based access provides strong protection against cookie theft, requiring users to present an X.509 certificate in addition to existing credentials such as cookies. These certificates are typically stored in the Trusted Platform Module (TPM) of the user's device, making it extremely difficult to be stolen by an attacker. Many enterprises already deploy certificates to their user's devices, and Google Cloud allows customers to either reuse their existing mTLS certificates, or use new certificates just for Google Cloud.

It's important to note that Access Context Manager is not responsible for policy enforcement. Instead Access Levels can be used in conjunction with other Google Cloud security policies in order to establish attribute based control (a.k.a Context Aware Access) for those security policies.

Using Access Levels, administrators should [limit what IPs can access Cloud Console and APIs](#) for their Google Cloud organization, making stolen cookies useless unless the attacker is using allowlisted IPs (such as the corporate network or VPN IPs), or depending on the Access Level definition, using a corporate-managed device to orchestrate their attack.

Context-Aware Access is also supported through [VPC-SC](#) in order to establish attribute based access control to a VPC-SC perimeter. Through [context-aware ingress rules](#), Google Cloud recommends that administrators restrict inbound access to a sensitive resource perimeter based on device attributes (including X.509 certificates), or the network origin of the request to prevent perimeter access by a malicious actor who successfully compromised an identity within your organization.

Lastly, Context-Aware Access is also supported through [Identity Aware Proxy](#) which secures applications hosted on Google Cloud. Identity Aware Proxy provides authentication capabilities to your cloud applications as well as authorization control via IAM Policies. When authoring IAM Policies against applications protected by Identity Aware Proxy, administrators should enforce [context aware access through Access Levels](#). These policies will help mitigate the risk of a malicious actor using compromised identities to access cloud applications protected by Identity Aware Proxy.

3

Mitigating cookie theft with session length control

To mitigate scenarios where a malicious actor successfully executes a cookie theft attack, [authentication session length](#) should be reduced to minimize the risk exposure window associated with this threat. Through [Cloud Identity Reauthentication Policies](#), administrators have control over the user authentication session length (min 1 hour, max 24 hours) and the re-authentication mechanism (i.e. password, security key) for the user to successfully re-authenticate. The session length control applies to: access via Google Cloud Console, the gcloud command-line tool, and any other app that requires Google Cloud scopes.

Phishing



Phishing is a form of social engineering where an external actor attempts to deceive your company's employees into revealing sensitive information. This information can be company data or an individual's credential information. Phishing attacks often rely on human error and pressure tactics for a successful attack. These types of attacks can be highly targeted attacks (where a specific individual within an organization is targeted) and thus highly sophisticated as the malicious actor only has a small attack surface. Alternatively, the malicious actor can target a high volume of individuals (i.e. all employees working for a company) with a less sophisticated approach in hopes that a single, distracted employee falls victim to the attack.

Traditional phishing techniques typically resort to email based phishing, fake websites, social media based phishing attacks and other similar techniques. These traditional phishing techniques typically target a large volume of individuals. For example, a malicious attacker can leverage large databases of corporate email addresses available on the dark web to send emails that masquerade as a legitimate email from a trusted source in order to trick recipients into revealing their credentials, or uploading sensitive corporate data.

As many enterprises have made the switch to public cloud, these traditional phishing techniques have also evolved to target public cloud infrastructure.

For example, imagine a scenario in which ACME (a fictional enterprise company) hosts a large majority of their core business data on Google Cloud. The ACME Security team has configured a security posture that ensures only ACME employees can access ACME data within this organization, thus the malicious actor can not get direct access to this data. Instead the malicious actor sets up a fake ACME Google Cloud organization (i.e. names the fake organization ACM3) and grants all ACME users access to this organization. To successfully execute their attack, the malicious actor is relying on human error on the part of an ACME user who mistakes the fake organization for the legitimate ACME organization and unintentionally uploads ACME data to a non-ACME environment.



1

Mitigating phishing attacks with Chrome Enterprise Premium

[Chrome Enterprise Premium](#) enables secure enterprise browsing thus mitigating traditional phishing techniques. Through frontline intelligence and AI, dynamic URL filtering, and site categorization, Chrome Enterprise Premium blocks access to malicious websites that employ phishing techniques.

2

Mitigating phishing attacks principal access boundaries

[Principal access boundary policies](#) enable security administrators to enforce identity centric policies on all enterprise identities they manage in order to restrict what Google Cloud resources these identities are allowed to access (regardless of where they are granted access via IAM policies). Google Cloud recommends enforcing a policy such as “identities associated with my Google Cloud organization can only interact with resources in my organization” to mitigate the cloud based phishing attacks as your identities will no longer be able to access the malicious resource shared with them.

3

Mitigating phishing attacks with organization restrictions

[Organization restrictions](#) enable security administrators to enforce that all Google Cloud bound network traffic originating from managed devices (under their control) is restricted to a specific set of target Google Cloud organizations. Enforcing these constraints on managed devices mitigates cloud based phishing attacks as those devices are restricted to only accessing resources in authorized organizations as defined by your security team.



Misconfigured access control



A misconfiguration in access control policies can result in unauthorized access and subsequent data exfiltration. It is not uncommon for large organizations to host hundreds of thousands of digital assets within public cloud infrastructure. With such a large asset inventory, a centralized security team can not reasonably manage access control for each individual asset. Instead these centralized security teams typically delegate access control management to individual teams or departments that own those assets.

These individual teams often do not have a security specialist to enact access control management and instead rely on engineers/developers for these tasks. When considering the sheer number of access control changes across all teams over the course of a year, it's highly likely that mistakes will happen. Mistakes in access control policies such as typos, unawareness of company security standards, or simply choosing convenience over security can result in access control vulnerabilities which can be easily exploited by an external malicious actor.

While security teams can configure a security configuration auditing pipeline with alerts to surface misconfigurations and playbooks for resolutions, a detective approach can often be too late to prevent data exfiltration. Instead, these security teams must employ a preventative approach.

1

Mitigating misconfigured access with Organization Policy Service

[Organization policies](#) enable security administrators to establish constraints on resource configurations within their organization. As you plan for establishing your access control posture, Google Cloud recommends enforcing an organization-wide [Domain Restricted Sharing \(DRS\) Organization Policy](#).

DRS policies should be used by security administrators to enforce that principals which are added to IAM Policies within their organization belong to a set of authorized Google Cloud organizations or specific Workspace domains. Domain Restricted Sharing policies mitigate risks associated with misconfigured IAM Policies on resources within your organization. Given the vast amount of IAM policies in a large organization, there is a high probability that a misconfigured access control can occur which can result in a sensitive resource being exposed to the internet - this can be accidental, such a typo in a principal reference, or intentional such as a malicious insider looking exfiltrate data.

As such centralized security teams must rely on security guardrail controls such as DRS policies to enforce restrictions on IAM Policy configurations to ensure that only authorized principals can be added in IAM policies against resources belonging to their enterprise organization. With such restrictions in place, a mistake or explicit attempt to add a non-authorized user to IAM policies will be denied.

2

Mitigating misconfigured access controls **with custom organization policies**

While Domain Restricted Sharing is specifically focused on restricting which principals are allowed to be added to IAM Policies within their organization, the integration of [IAM with custom organization policies](#) unlocks a number of new use-cases.

Through this integration, administrators can enforce that IAM policies within a given resource context (i.e. project, organization, etc) must match a set of criteria such as:

- Only allow specific roles to be granted against resources in this project.
- Deny specific roles to be granted against resources in this project.
- Only allow specific members to be granted access via policies against this folder.
- Deny “allUsers” grant for any resources in this organization.

This integration provides significantly more granularity in the type of restrictions which can be enforced against IAM policies within your organization. By applying these restrictions, administrators can then delegate the ability to manage policies directly to developers while knowing that these developers can only create, update, or delete policies according to the exact rules they've enforced. For a full list of supported use-cases, please visit this [section of our documentation](#).

3

Mitigating misconfigured access controls **with IAM Recommender**

In addition to the type of restrictions which can be enforced with Organization Policy constraints, [IAM Recommender](#), available through the [Cloud Infrastructure Entitlement Management \(CIEM\)](#) suite in Security Command Center, should be used to identify and remediate misconfigured access controls which [grant too much access](#) or are subject to [lateral movement vulnerabilities](#).

Using ML-based policy findings, IAM recommender generates role recommendations which help you identify and remove excess permissions from your principals, improving your resources' security configurations. These recommendations help you enforce the principle of least privilege by ensuring that principals have only the permissions that they actually need, thus helping mitigate data exfiltration risks due to ambient, over-permissioned IAM grants.



Network infiltration



As cloud resources are deployed and managed on the cloud provider's infrastructure, all cloud resource access must be performed over a network connecting the originating device to the cloud provider network. This network can take many forms. For example the network can be composed entirely of private network links (i.e. Direct Interconnect between customer data center and Google Cloud network), the network can be entirely public (i.e. cloud resource access via API calls over the public internet), or a combination of both public and private network segments (i.e. hybrid networks). As such, ensuring that traffic traversing this network is secure is critical to ensuring your cloud resources are secure.

A malicious actor can exploit weaknesses in the network security configuration to penetrate the network or a single device within this network. Once within the network, this malicious actor can use techniques such as lateral movement to gain access to other systems, and reach their ultimate target. For example, a misconfigured cloud firewall policy could enable a malicious actor to gain access to an application running in the cloud that contains sensitive company information. Without enabling layered security controls, once within a network, a malicious actor can begin exfiltrating data out of this network to an external target.

1

Mitigating network infiltration risks with firewall policies

Google Cloud Compute resources such as virtual machines require attachment to [Virtual Private Cloud \(VPC\) networks](#). [Firewall policies](#) are a collection of firewall rules which enable network administrators to establish ingress and egress rules that controls inbound and outbound traffic within the VPC network.

Through [firewall ingress rules](#), a network administrator can lock-down inbound access to a given VPC network by ensuring that all unauthorized ingress connections are blocked by default. A security administrator can leverage [hierarchical firewall policies](#) to enforce a similar inbound network constraint across their entire organization. This allows organization level administrators to enforce ingress guardrails on all VPC networks within their organization while delegating the administration of other network security settings to lower-level network administrators.

2

Mitigating network infiltration risks with VPC-SC

In addition to Firewall Policies, [VPC-SC](#) can be used to provide more granular ingress control within a [VPC-SC perimeter](#). Once a perimeter is established, [ingress rules](#) can be configured to restrict inbound network traffic based on the source or identity of the request as well as the target of the request (i.e. API, resource, etc). By default, a service perimeter will only allow free communication within the perimeter and deny traffic from outside of the perimeter.

3

Mitigating network infiltration risks with organization policies

As many Google Cloud customers expose cloud resource on the public internet (i.e. VMs running web applications with public IP, Cloud SQL with public IP, etc), these resources become a prime target for network infiltration attacks.

[Organization policies](#) enable security administrators to establish constraints on resource configurations within their organization. In order to restrict cloud resources exposure to the public internet, [there are a number of prebuilt constraints](#) which you can enable to protectively mitigate network infiltration risks. The following are examples of such constraints, [Enforce Public Access Prevention on Cloud Storage Buckets](#), [Restrict Public IP access on Cloud SQL instances](#), [Disabling external IP access for VMs](#).

4

Mitigating network infiltration risks with Cloud Armor

While the previous section described scenarios where cloud resources are exposed to the public internet, many Google Cloud customers deploy internet facing applications hosted on IaaS or PaaS Google Cloud services. [Cloud Armor](#) will protect your Google Cloud deployments from multiple types of threats, including attacks such as cross-site scripting (XSS) and SQL injection attacks which can otherwise be used to steal credentials or gain access to the underlying infrastructure. This additional layer of defense will help mitigate risks associated with vulnerable cloud applications which a malicious actor can target as the infiltration point.



Other network security best practices

In addition to the recommendations above, there are other [network security best practices](#) to consider when architecting your network to further reduce data exfiltration risks associated with network infiltration and exfiltration threats. Some of these recommendations are product specific while others are general recommendations for securely architecting your cloud networks.

Network exfiltration




Network exfiltration refers to risks associated with data exfiltration originating from within your network. How a malicious actor can infiltrate your network is discussed in the network infiltration section. This section focuses on security risks after a malicious actor gains access to your network. The main intention of a malicious actor after they gain access to a network is to exfiltrate assets that are freely accessible on this network.

Techniques such as packet sniffing, man-in-the-middle, etc can be used to extract data directly from network traffic. However access to a network, especially a private cloud network, can also result in the malicious actor gaining direct access to unprotected assets on the network (i.e. databases, web servers, etc) since, often, organizations will

employ a “castle-and-moat” network security model - effectively employing a model in which network access is tightly controlled, however once inside the network, there are no additional security checkpoints. Instead of this castle-and-moat architecture, a [Zero-Trust security model](#) should be employed to mitigate these risks.

As the focus of this paper is data exfiltration, let us assume that a malicious actor was able to gain access to sensitive data by infiltrating your network. At this point, they are looking to exfiltrate this data to a target outside of your network and thus outside of your control. To mitigate this risk, the security controls of this network should prevent outbound traffic destined to unauthorized external targets. Traditionally this is done through a firewall between your internal network and the internet, however in cloud environments, there are additional, more sophisticated controls that should be employed to protect against data exfiltration risks.



1

Mitigating network exfiltration risks with firewall policies

The first line of the defense of your network security posture is [firewall policies](#). Firewall policies are a collection of firewall rules which enable network administrators to establish ingress and egress rules that controls inbound and outbound traffic within the [VPC network](#).

Through [firewall egress rules](#), a network administrator should lock-down outbound access from a given VPC network by ensuring that all unauthorized egress connections are blocked by default. A security administrator should leverage [hierarchical firewall policies](#) to enforce a similar outbound network constraint across their entire organization.

2

Mitigating network exfiltration risks with VPC-SC

In addition to firewall policies, administrators should employ [VPC-SC](#) to enforce granular egress control within a [VPC-SC perimeter](#). Once a perimeter is defined, [egress rules](#) should be configured to restrict outbound network traffic based on the target destination of the request as well as the source identity of the request. By default, a service perimeter will only allow free communication within the perimeter and deny traffic from outside of the perimeter.





3

Mitigating network exfiltration risks with Secure Web Proxy

Network and security administrators should also employ Secure Web Proxy to [secure egress web traffic \(HTTP/S\) within a network](#) that connects a variety of different resource types (i.e. virtual machines, containers, workloads outside of Google Cloud connected by Cloud VPN or Cloud Interconnect, etc). Alternatively a single Secure Web Proxy deployment can also be used to secure multiple networks through [Private Service Connect attachment mode](#).

[Secure Web Proxy](#) should be used to apply granular access policies to egress web traffic. Secure Web Proxy policies are applied to the traffic sent from cloud workloads or applications to the internet. It also monitors outbound access to untrusted web services that don't conform to your policy and logs it to Cloud Logging. Using this tool, administrators can monitor internet usage, discover threats to their network, and respond to threats.

This additional layer of security should be used in conjunction with Firewall policies and VPC-SC to provide comprehensive protection from the network exfiltration risks described in this section.

Other network security best practices

In addition to the recommendations above, there are other [network security best practices](#) to consider when architecting your network to further reduce data exfiltration risks associated with network infiltration and exfiltration threats.

Insider risks



Insider risks refers to threats associated with insiders (i.e. company employees, contractors, partners, etc) intentionally looking to exfiltrate company data. There are a variety of potential motivations for these insiders. For example, a company employee might look to exfiltrate company data if they are about to switch employers to a competitor and would like access to past project data they worked on. Alternatively, a malicious actor could apply pressure (via blackmail, personal threats, etc) to an employee such that they exfiltrate data on behalf of the malicious actor.

The main challenge associated with this type of security threat is that these insiders already have access to the data that you are looking to protect. Unlike with phishing, credential theft, and other threats covered so far where the main focus is on preventing external threats, insider risks must be mitigated by protecting against outbound threats - effectively preventing data exfiltration through security controls on outbound traffic.



1

Mitigating insider risks with VPC-SC

Since insiders already have access to the sensitive data you are looking to protect, it's imperative that you establish absolute control over data movement for resources which host this sensitive data. Through [VPC-SC](#), you should establish a [security perimeter](#) around the resources which host sensitive data and leveraging [perimeter ingress and egress rules](#) you should enforce a security posture which prevents copying data to a resource outside the perimeter.

2

Mitigating insider risks with principal access boundaries

Malicious insiders looking to exfiltrate data will often attempt to leverage self-owned Google Cloud organizations or resources (outside of your corporate organization) as the target for data exfiltration. For example they might utilize the [Cloud Storage bucket copy API call](#) in order to copy the data stored in a corporate managed Cloud Storage bucket to a bucket they created outside of your organization's control.

Through [principal access boundary policies](#), security administrators should enforce identity-centric policies on all enterprise identities they manage in order to restrict what Google Cloud resources these identities are allowed to access (regardless of where they are granted access via IAM policies). Enforcing a policy such as "identities associated with my Google Cloud organization can only interact with resources in my organization" mitigates the risk of a malicious insider attempting to move or copy corporate data to a

Google Cloud target outside of your control as this policy will prevent their identity from accessing that external data exfiltration target. Since Principal Access Boundary policies support enforcement on Service Accounts, this also mitigates the risk of a malicious insider impersonating a corporate service account and attempting to exfiltrate data through this separate identity.

3

Mitigating insider risks with organization restrictions

In addition to the Principal Access Boundary policy mitigation described above, [organization restrictions](#) should be used as an additional layer of defense which is enforced on the traffic origination from the managed devices that company employees use to access Google Cloud services.

Organization Restrictions enable security administrators to enforce that all Google Cloud bound network traffic originating from managed devices (under their control) is restricted to a specific set of target Google Cloud organizations. Enforcing these constraints on managed devices mitigates risks associated with insiders who attempt to circumvent Principal Access Boundary policies on their identity by utilizing a different non corporate-managed identity when attempting to exfiltrate data (i.e. use corporate identity to download data on their laptop and use a self-owned identity to upload this data to external data exfiltration target).

This security control should be paired with a [Context Aware Access policy](#) that ensures authentication to enterprise Google accounts can only occur from corporate managed devices. Thus employees must use managed devices to authenticate to their enterprise account and their Google Cloud bound traffic from these managed devices will be subject to Organization Restrictions.

Tying it all together

In this paper, we discussed the concept of data exfiltration and how it can be a cyber risk for organizations of all sizes. We also covered some of the most common attack vectors that can lead to data exfiltration and the Google Cloud security controls that can be used to mitigate these attacks.

It is important to note that these controls must be used together to establish a defense-in-depth security model. No single control is sufficient to protect against all data exfiltration attacks. In addition to establishing a security posture to mitigate data exfiltration risks using these controls, it is imperative that you continuously audit and monitor your security posture to validate correctness and identify required changes to keep up with business needs.

As these controls address various security domains, we highly recommend employing the use of [Security Command Center](#) (SCC) to help aggregate findings and recommendations across many of these controls. In addition, SCC can also be used to identify software vulnerabilities or compliance violations, simulate sophisticated attacks across your cloud resources, identify indicators of compromise based on latest threat intelligence findings, and much more.

We understand that getting security right in the cloud can be challenging, and an expert partner committed to your success can make all the difference. At Google Cloud, we are active partners committed to helping you achieve your desired risk and security outcomes. We are not delineators of where our responsibility ends and where yours begins.

Instead, we stand with you from day one, helping you implement best practices for safely migrating to and operating in our Trusted Cloud. We call this operating model shared fate.

With shared fate, we help you operationalize a best-practices approach for securing your cloud footprint.

- A secure-by-design, [secure-by-default infrastructure foundation](#) augmented by layers of security controls that you can configure according to your risk profile.
- [Security foundations](#) that address top security concerns and provide our top recommendations.
- [Secure blueprints](#) that let you deploy and maintain secure solutions using infrastructure as code (IaC).
- [Architecture Framework](#) best practices that address the top recommendations for building security into your cloud infrastructure patterns.
- [Landing zone navigation guides](#) to help you build a secure foundation for your workloads, including resource hierarchy, identity onboarding, security and key management, and network structure.

By following Google Cloud best practices for each of these controls, organizations can significantly reduce the risk of data exfiltration and protect their sensitive data from malicious actors.





Google Cloud