Preventing Data Exfiltration

An essential function of computer and network security is keeping sensitive data inaccessible to unauthorized third parties. This document explores characteristics of data exfiltration risks, and discusses industry-wide best practices for securing data. It explains how to use tools and features in Google Cloud Platform to reduce risks, detect data exfiltration, and respond to exfiltration events. When possible, security threats and defense approaches will be described in a cloud-independent context. The evolving regulatory environment, especially the European General Data Protection Regulation (GDPR) that becomes mandatory in 2018, has added new emphasis to the deployment of data exfiltration prevention mechanisms.

Defining data exfiltration

In this document, *data exfiltration* is defined as when an authorized person extracts data from the secured systems where it belongs, and either shares it with unauthorized third parties or moves it to insecure systems. Authorized persons include employees, system administrators, and trusted users. Data exfiltration can occur due to the actions of malicious or compromised actors, or accidentally.

To reduce the risk of data exfiltration, organizations must integrate security awareness and best practices into their culture. They must consistently evaluate the risks of every interaction with computer networks, devices, applications, data, and other users. Organizations may also decide to institute periodic audits to verify that best practices are followed.

Confronting data exfiltration risks in the cloud

Many traditional data security strategies are based on hardening the physical perimeter defenses of private networks. As a cloud consumer, however, you don't control the physical network infrastructure your services use. In public clouds, the network fabric of the hosting provider is shared, and there is no perimeter in the traditional sense. Securing data in the cloud requires new security approaches and methods of auditing data access.

As an analogy, consider that public cloud infrastructures were among the first in the industry to adopt the use of commodity hardware in the data center. While this can result in higher rates of hardware failure, fallout from such failures is minimized through redundancy and intelligent service architecture. In such an environment, services must be able to absorb multiple failures gracefully. Service architectures are designed for hardware and network failures by dividing up processing, storage, authentication, and other tasks across multiple machines and geographies, thereby minimizing the impacts from any one failure event. In securing your data, you should take a similar approach: design an architecture to minimize downtime and limit the effects to the rest of your system in the event of a security compromise.

To satisfy the most security-conscious customers, the public cloud security model incorporates this thinking. Google Cloud offers Shielded VM (https://cloud.google.com/security/shielded-cloud/shielded-vm) to provide verifiable integrity of your Compute Engine virtual machine (VM) instances, so you can be confident your instances haven't been compromised by boot- or kernel-level malware. Shielded VM's verifiable integrity is achieved through the use of Secure Boot, vTPM-enabled Measured Boot, and integrity monitoring.

Additionally, providers can deploy specialized agents to produce telemetry about user and host activity in cloud-based virtual machines (VMs). This provides the Security Operations Center (SOC) with visibility into activities within and around the frequently-morphing security boundary.

Note: By default, Google does not currently install agents to provide this telemetry in VMs. Customers are free to do so in their VMS.

Providers also introduce explicit chokepoints, such as bastion host (https://cloud.google.com/solutions/connecting-securely#bastion) for communication with fleets of VMs, network proxy servers, network egress servers, and cross-project networks. These measures can reduce the risk of data exfiltration, but cannot eliminate it completely.

Your organization must also establish a strong detection and response infrastructure for data exfiltration events. The right cloud infrastructure will give you rapid detection of risky or improper activity, limit the "blast radius" of the activity, and minimize the window of opportunity for the exfiltrating actor.

Data exfiltration event categories

Data exfiltration events can be categorized by common technological, organizational, and physical characteristics. In the sections that follow, we'll look at some of these categories and discuss prevention and mitigation strategies for each.

Outbound mail

In this scenario, actors use authorized telecommunications infrastructure, such as business email or mobile devices, to transmit sensitive data from secure computer systems to untrusted third parties or insecure private systems. The sensitive data can be transmitted as plain text in an email or text message, or attached as a file. This method is often used to exfiltrate the contents of organization emails, calendars, databases, images, planning documents, business forecasts, and source code.

Many email and messaging systems save drafts to the cloud, so it isn't sufficient to check for sensitive data when the message is sent. If a person has outside access to their business email or other messaging service that supports saved drafts, they can use that feature for exfiltration. By saving a draft from devices and networks with access to sensitive data, and then accessing the draft from another client, the actor avoids logging and auditing systems.

Prevention and Mitigation

This scenario involves telecommunications systems selected and authorized by your organization, which gives you more options for securing against data exfiltration than scenarios involving private or third party tools.

Consider implementing some of the following prevention and mitigation strategies:

- Monitor the volume and frequency of data transmission by your users over email and other organizational messaging tools. If the average user sends 5 megabytes of data on average per day, a user sending 500 megabytes should trigger an alert.
- Retain a log of addresses used to send email, what devices emails are sent from, and the addresses of recipients. These can help you identify the nature and scope of a data exfiltration event. The Administrator security checklist (https://support.google.com/a/answer/2984349?hl=en&ref_topic=2683865) explains how to audit an email account for security risks in Gmail Enterprise (https://gsuite.google.com/intl/en_us/products/gmail/).
- Scan emails sent from systems with access to sensitive data to ensure they don't contain unauthorized content. This can be made easier by tagging sensitive content with markers, like keywords or hashes.
- Prevent sending messages over insecure channels, such as using http instead of https, and alert your IT security staff of attempts.

Downloads to insecure devices

These cases occur when a user accesses sensitive data through authorized channels and then transfers the data to an insecure local device. Actors may use laptops, smartphones, external drives, cameras, or specialized devices to capture sensitive data for exfiltration. The actor can download existing files from your services in the cloud, or copy data into new files. If the files are transferred to unmonitored or insecure devices, they are at high risk for exfiltration.

Prevention and Mitigation

Cloud-based networks have advantages in preventing this kind of event. Many methods for transferring data to a local device require a physical connection to transferrable media. If the data is instead stored in the cloud, it has to be downloaded before it's transferred. These downloads are subject to security and tracking features of the hosting service and clients.

Consider implementing some of these policies and techniques:

- Prohibit downloads of very sensitive data. Depending on how your data is used and processed in the cloud, users may never need to download it to local hardware. If possible, keep all data in the cloud and perform all computation in the cloud. If data is technically downloadable, establish a policy that prohibits downloads, classify and label sensitive data, and keep access logs of data that is requested and served via secured interactions and API calls. See Viewing Activity Logs (https://cloud.google.com/compute/docs/activity-logs) for details.
- Use a Cloud Access Security Broker (CASB) to regulate connections between authorized clients and cloud services according to your organization's security policies.
- Wrap files with Digital Rights Management (DRM) tools. This puts permissions-aware security and encryption on each file.
- Implement dynamic watermarking in your authorized clients to record the user responsible for screenshots or photographs of computer displays containing sensitive information.

Uploads to external services

Similar to the previous category of events, this category often involves downloading sensitive data to local infrastructure. The actor then uploads the data to a third-party through a web browser client or other unmonitored software. Third party services could be innocuous-seeming web sites like a social network, where the actor could accidentally upload the wrong images or paste the wrong text. Sophisticated malicious actors may be able to pass small amounts of sensitive data, like user credentials or encryption keys, as URL parameters to specialized web applications.

Prevention and Mitigation

The risk of this kind of event can be reduced through the same policy restrictions on downloads that protect against local copying of sensitive data. However, a policy does not eliminate the risk of screenshots or copied text being uploaded to social media, file sharing web sites, or other cloud services.

Security practices to consider for combating this kind of risk include:

- Prohibit any data from being downloaded. Keep all data in the cloud and perform all computation in the cloud. Data should be requested and served by secured and logged API interactions. See Viewing Activity Logs (https://cloud.google.com/compute/docs/activity-logs) for details.
- Prevent installation of insecure third party software, such as social media apps or unauthorized browser plugins, on devices with access to sensitive data.
- Use a CASB to regulate traffic from cloud access points and enforce encryption policies for all data transmitted to clients.

Insecure cloud behavior

Using cloud services introduces some new categories of data exfiltration risks that IT security professionals should be conscious of. These include a range of cases where employees, users, or administrators use features of the cloud provider suite in insecure ways. There is data exfiltration potential from any actor who has the ability to requisition or modify virtual machines (VMs), deploy code, or make requests to cloud storage or computation services.

Cloud networks have public frontends and the ability to communicate with the broader internet. Securing and authorizing the behavior of services running in the cloud is essential to providing data security. Actors with sufficient permissions can initiate outbound transmission of sensitive data, move sensitive data from secure containers to less secure ones, or create unauthorized cloud services on behalf of an organization.

Prevention and Mitigation

Maintaining secure behavior for your cloud services requires precise, narrowly scoped permissions, and comprehensive logging. Wherever possible, prohibit actors from accessing the backends of your services. For most tasks that an employee or administrator needs to complete on a VM, there are automated agents and frontend clients that are secure and monitorable. Use these where it is possible to limit the number of persons with direct SSH access to your cloud machines. When feasible, scan all data sent to the broader Internet to identify sensitive information. For applications that process information from external users or systems, consider scanning that information to prevent inadvertent collection, storage, or sharing of sensitive data like Personally Identifiable Information (PII).

For VMs in the cloud, consider these security principles:

- Set up IP tables on your VMs that prohibit outgoing connections to unknown addresses. This can reduce the risk that an actor successfully transmits sensitive data out of your network.
- Avoid giving your VMs public IP addresses, and use a Network Address Translation (NAT) service to process ingoing and outgoing connections. Read this guide on setting up a NAT gateway (https://cloud.google.com/compute/docs/networking#natgateway) for Compute Engine to learn more.
- Consider using a bastion host (https://cloud.google.com/solutions/connecting-securely#bastion) in the cloud to mediate and monitor connections to other hosts.
- Disable remote management software like Remote Desktop Protocol (RDP) or Windows Remote Management (WinRM) agents on machines that do not need them.
- Use Private Google Access (https://cloud.google.com/compute/docs/private-google-access) to enable virtual machine (VM) instances on a subnetwork to reach Google APIs and Services using an internal IP address rather than an external IP address.
- Consider Cross-Project Networking (XPN) (https://cloud.google.com/compute/docs/xpn/) to share Google Cloud Platform (GCP) virtual networks across projects in your Cloud Organization.
- Limit direct SSH access to VMs to those persons with critical and unavoidable need. Google Compute Engine provides comprehensive SSH key management tools (https://cloud.google.com/compute/docs/instances/ssh-keys) for controlling access to VMs.

For cloud storage services like Cloud Storage (https://cloud.google.com/storage) or Cloud Bigtable (https://cloud.google.com/bigtable), the following practices can reduce exfiltration risks:

- Use Identity and Access Management (Cloud IAM) (https://cloud.google.com/iam) to provide users and applications with the narrowest set of permissions necessary for accessing data. Store data with different sensitivity and access requirements in different containers, to allow for permissions to be as granular as possible.
- Monitor and limit the rate at which data can be read from your storage resources. Use monitoring agents to alert your security team if there are attempts to move much more data than expected in the normal use case.
- Make permissions to very sensitive data temporary and subject to frequent review and revocation. For example, App Engine (https://cloud.google.com/appengine/docs/quotas) quotas can be useful here.
- Have a human team regularly audit the set of persons with access to very sensitive containers.
- Keep thorough logs of all access to your storage services. Ideally, the set of persons granted access to the storage services will be separate from the set of persons with access to the logs. This reduces the risk tampering with the logs by malicious actors. Consider using Cloud Storage access log utilities (https://cloud.google.com/storage/docs/access-logs) to write log data to a separate storage bucket.
- More comprehensive security best practices (https://cloud.google.com/storage/docs/best-practices#security) for Cloud Storage are available.

Enforcing compliance with security policies

The rich infrastructure provided by the Google Cloud Platform (GCP) creates multiple opportunities for customers to develop solutions that target their needs. At the same time, the rich infrastructure also brings new challenges such as enforcing desired security policies across different projects in an organization. To simplify security management, GCP introduced a hierarchy of entities in which all resources reside. This hierarchy is rooted in the concept of Organizations. Organizations may optionally contain folders or projects. Folders may optionally contain sub-folders or projects. All GCP Service resources belong to a project.

Using this hierarchy, Organization -> Folder -> Project -> GCP Service -> Resource, security policies can be set at any level of the hierarchy and are inherited down the lower levels of the hierarchy. Security policies are evaluated from top-to-bottom in the resource hierarchy and, as soon as an “allow” answer is obtained, the access to the resource is granted.

Enforcing compliance

Using the resource hierarchy, and the inheritance of security policies simplifies auditing to ensure that desired security policies are uniformly followed. Because of the inheritance property, administrators can demonstrate that, for example, all projects allow the same set of auditors to inspect their data. They achieve this by having such a security policy at the Organization level and never overriding it at a lower level. These security policies are specified in software audit activities and their verification can be automated.

Identification and redaction of sensitive data

One of the first steps in managing sensitive data is knowing where it is. Once identified, you are better equipped to set access control to ensure proper access/handling and use techniques to reduce the sensitivity through redaction, masking, or de-identification of the data. Once the data is in redacted form it ceases to convey the sensitive nature of it, like being a specific social security number, a valid credit card number, or Personally Identifiable Information (PII).

The traditional challenge in redacting large amounts of diverse data is the need to automate recognition, classification, and appropriate redaction. An advancement is to have system support to reason about the content in data fields in an automated manner. This level of automated visibility into arbitrary data streams allows applications to decide what data to transmit to which end-points, which systems to use to store the different kinds of data that are being managed, and when to alert about specific kinds of data being transmitted.

Prevention and mitigation

In Google Cloud, the Data Loss Prevention (DLP) lets you understand and manage sensitive data. It provides fast, scalable classification and optional redaction for sensitive data elements like credit card numbers, names, Social Security numbers, passport numbers, US and selected international driver's license numbers, and phone numbers. Cloud DLP supports text, structured data, and images – just submit data to Cloud DLP or specify data stored on your Google Cloud Storage, BigQuery, and Cloud Datastore instances. The findings from Cloud DLP can be used to automatically monitor or inform configuration of Identity and Access Management (Cloud IAM) settings, data residency, or other policies. Cloud DLP can also help you redact or mask certain parts of this data in order to reduce the sensitivity or help with data minimization as part of a least-privileged or need-to-know policy. Techniques available are masking, format-preserving encryption, tokenization, and bucketing across structured or free text data.

Rogue administrators

By design, most computer systems grant unchecked power to designated administrators. Malicious or compromised administrators will have sufficient permissions to perpetrate any of the scenarios discussed in this document, and furthermore have the greatest ability to obliterate logs and evidence of their actions. Reducing these risks requires separation of powers and authority over parts of your network, and enabling administrators to monitor each other.

Prevention and Mitigation

Limiting the authority of any single actor is essential to reducing the risks presented by a rogue administrator.

In order to accomplish their assigned tasks, administrators will have the ability to exfiltrate data, however the following principles can be applied to reduce the scope and magnitude of such events should they occur:

- Secure a large network against administrator malfeasance by logging the actions of administrators in a place they cannot access. Use a separate security team to manage monitoring and logging services.
- Consider making all administrator access temporary with a short expiration period. In many networks, administrators do not need persistent access.
- Require multiple actors to approve administrative actions. By treating all administrative actions like source code that requires approval, you can reduce the risks presented by a single actor.

Employee terminations

The Computer Emergency Response Team (CERT) at the Software Engineering Institute (https://www.sei.cmu.edu/) of Carnegie Mellon University produced a 2011 paper (https://resources.sei.cmu.edu/library/asset-view.cfm?assetID=9875) showing that employees were more likely to engage in data exfiltration when they anticipated imminent termination. A pending employee termination is a period of increased risk which demands additional attention from IT security teams.

Prevention and Mitigation

For networks with very sensitive data, consider connecting logging and monitoring systems to HR software that records an upcoming termination and set more conservative thresholds for alerting security teams to abnormal behavior by these users.

Conclusion

The flexibility, cost-savings, and power of using public cloud infrastructure requires increased vigilance and new approaches to securing data from exfiltration. You can architect your organization's policies and implementations to account for the environment by using techniques described in this paper:

- Minimize the "blast radius" of data exfiltration events through compartmentalization of data.
- Create redundancy and approvals in system administrator workflows to increase accountability.
- Use granular permissions and grant access to sensitive data only to those whose job function requires it.
- Use logging to increase the transparency into the access and movement of data in your organization.
- Restrict and monitor ingress and egress to machines in your organization using networking rules, identity and access management (Cloud IAM), and bastion hosts.
- Create a baseline of normal data flows, such as amounts of data accessed or transferred, and geographical locations of access against which to compare abnormal behaviors.

Next steps

- Read about the Organization Policy Service (https://cloud.google.com/resource-manager/docs/organization-policy/overview).
- Read about Data Loss Prevention (DLP) (https://cloud.google.com/dlp/).
- Read the Google Cloud Security Overview (https://cloud.google.com/security/).
- Visit the Security Partner Ecosystem (https://cloud.google.com/security/partners/) page to learn about security-centric GCP partners.
- Visit the Google Cloud & the GDPR (https://cloud.google.com/security/gdpr/) page.

Last updated December 4, 2019.