netskope

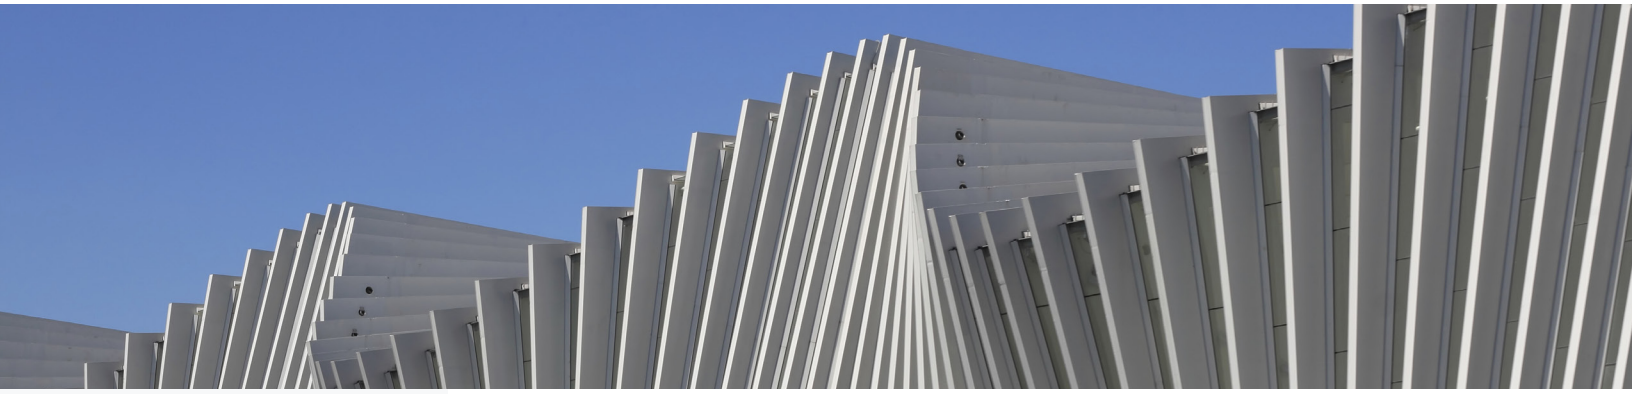

Cloud Data Loss Prevention Reference Architecture

A framework designed to help organizations more accurately find sensitive content in the cloud and integrate with on-premises DLP solutions.

Introduction

Enterprise assets are increasingly moving from applications deployed in private data centers to public cloud services. Matched with this trend are data access patterns evolving from campus locations to almost everywhere with the proliferation of smartphones and tablets, further increasing the risk of enterprise data loss. The latest Netskope Cloud Report shows that 17.9% of all files in enterprise-sanctioned cloud apps constitute a data policy violation – which doesn't include data in unsanctioned cloud apps that employees use to get their jobs done. Enterprises experiencing heavy cloud use need a way to extend their data loss prevention (DLP) frameworks to include cloud services and accommodate these new access patterns to effectively protect corporate IP and sensitive data such as personally identifiable information (PII), protected health information (PHI), and/or payment card industry information (PCI).

Reference Architecture Goals

This reference architecture has three primary goals for implementation of a cloud DLP solution to protect data assets in an enterprise:

- Maintain employee user experience. Enterprises that don't want to backhaul cloud traffic via VPN need efficient data access that preserves the user experience.
- Protect existing on-premises DLP investments. Enterprises have significant investments in highly tuned, effective DLP solutions, and want to protect those investments.
- Preserve existing incident response processes. Enterprises need to keep existing incident response processes that have been developed with existing tools, plans, and investments as to prevent any disruption.

Key Architectural Tenets

1. **Derive context from cloud service transactions and set policy based on it before moving to the next stage of data identification:** Reduce the scope of content you inspect en route to or from, or at rest in, the cloud by using cloud service transaction context such as identity, location, device, or activity. Specifically, stop activities (regardless of data type) that are non-conforming (e.g., you may have a policy that states that unmanaged devices downloading any content from a sanctioned app are disallowed).

2. **Use a classification framework to identify or categorize sensitive content:** Use a classification framework to identify or categorize sensitive content. Apply persistent metadata to maintain that data identity as content moves to and from cloud services.

3. **Apply data classification to discover sensitive content in the cloud:** Inspect content en route to or from, or at rest in, cloud services in order to determine if that content is sensitive. Allow conforming content to remain or pass by unencumbered.

4. **Quarantine and redirect potentially sensitive content to an on-premises DLP solution:** To ensure a high quality user experience, quarantine remaining content and notify the user. Redirect content to an on-premises DLP solution via secure Internet Content Adaptation Protocol (ICAP) to evaluate against highly-tuned policies in order to determine violations.

5. **Enforce policies and initiate incident response:** Enforce policy actions and initiate incident response workflows in applicable solutions, whether they are on-premises and/or in the cloud, based on ICAP responses from on-premises DLP solutions, together with the context derived from the cloud service transactions.

6. **Ensure user accountability:** Ensure user accountability by coaching users on potential violations and allowing them to provide feedback (e.g., report a false positive or provide a justification).

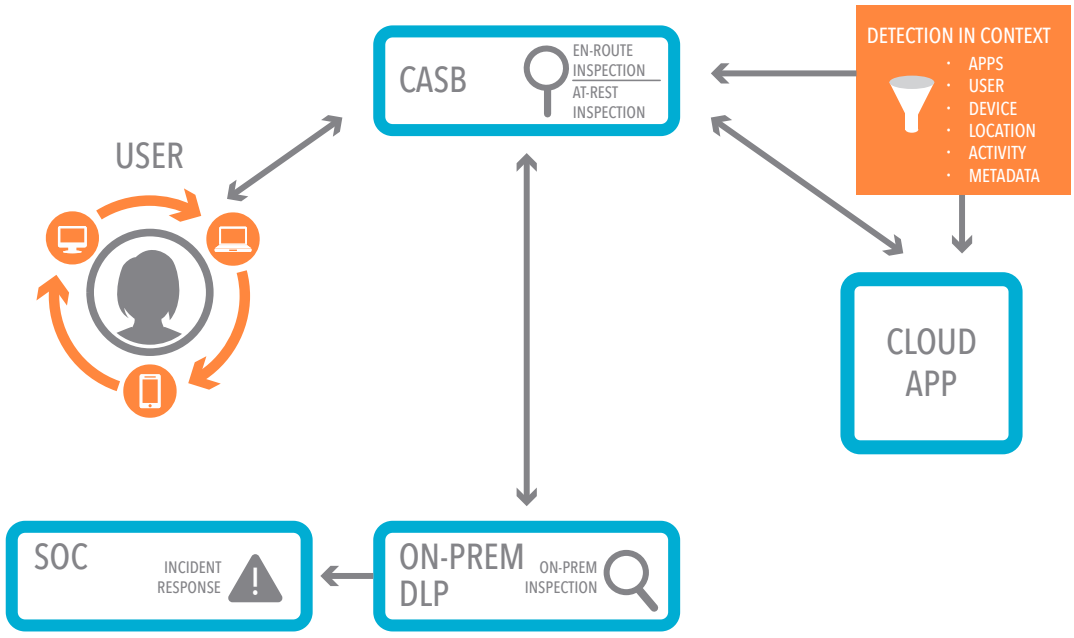

Figure 1. This diagram shows a CASB deployment in context of a typical enterprise environment and shows CASB directing potential violations to on-premises DLP solutions for follow-up and incident response.