

Applying Zero Trust on Google Cloud

Matching Google Cloud services with NIST 800-207

Brian Schmult, Hani Raouda, John Hurring, Mikesheh Khanal, Muhuji Mshana, and Alida Wilms



Table of contents

Introduction	4
Operation Aurora	5
Understanding zero trust	6
What is zero trust?	6
Why is zero trust becoming popular?	6
How is zero trust different?	6
What are the benefits of zero trust?	7
Zero trust solutioning	8
“North star”	8
Components	8
Core components	8
Functional components	9
Operational principles	9
Project planning	11
Approaching zero trust	11
Understand the why	11
Training	12
Assemble a team	12
Zero trust change management process	14
Critical phase 1 steps	15
Perform a zero trust assessment	15
Create the DAAS Inventory	15
Create the signal inventory	16
Inventory prioritization	17
Complete the zero trust capability gap analysis	17
Migrating	19
Building the plan	19
Implementing zero trust	20
Step 1: Define the protect surface	20
Step 2: Map the transaction flows	21
Step 3: Build a zero trust architecture	22
Step 4: Create a zero trust policy	23
Step 5: Monitor and maintain the network	23
Architecture	24
Design concepts	24

Design patterns	28
1. Enterprise with remote employees	28
2. Enterprise with a multi-cloud application	30
3. Cross-enterprise collaboration	31
4. Enterprise with non-employee access	32
5. Enterprise with public-facing services	33
Zero trust capability matrix	35
Mapping Google Cloud services to the NIST 800-207 pillars	42
BeyondCorp Enterprise	45

Introduction

This guide is intended to provide readers with an understanding of the following:

- What is zero trust and why it matters
- How to build a project plan for a zero trust migration
- What Google Cloud services align to [NIST 800-207](#) pillars

Implementing zero trust is not something that can be done overnight, in a silo, with a sole vendor, or by one team. A successful journey is driven by significant amounts of detailed planning, cross-business unit collaboration, organizational buy-in, and stakeholder support; all accompanied by the right selection of vendors and capabilities. The end state of this journey is a paradigm shift that will fundamentally alter current approaches to securing an enterprise, as achieving zero trust impacts every user, device, workload, data source, asset, and service within an organization.

Operation Aurora

As technology continues to rapidly evolve and continues to redefine *cutting edge*, threat vectors are also evolving at a near 1:1, if not 2:1 rate. Each new cutting-edge capability creates a new potential vulnerability or threat. Google recognized the need for a new cybersecurity approach in 2009 after [Operation Aurora](#), and embarked on a decade-long journey into zero trust. When Google set out on this journey, the problem statement encompassed objectives such as the following:

- Design and implement a modern enterprise network
- Enable the ability to make intelligent centralized access decisions in real time
- Use the same model of access from any geographic location
- Decouple access logic from identity system of record
- Create common access policy definitions

While Google might have been among the first major adopters of zero trust, much has changed since 2009. Today, zero trust principles are embedded in many of our products, and help ensure the integrity, safety, and security of our customer's workloads. To complement our product integrations, Google has also created a [Zero Trust Foundations](#) workshop with the [Google Cybersecurity Action Team \(GCAT\)](#), the Office of the CISO, and the [Google Cloud Consulting \(GCC\)](#) organization to help customers adopt zero trust architectures.

Understanding zero trust

It's time to challenge the established approaches to data security

What is zero trust?

[NIST](#), to [Forrester](#), to [Gartner](#) provide various definitions for zero trust. *Zero trust* is an operational security framework that is designed to help adopters plan, build, and deploy dynamic context-based security controls within workloads.

Zero trust is not one vendor or an off-the-shelf product to buy or lease. You can't migrate to zero trust overnight.

Why is zero trust becoming popular?

There has been a notable push to adopt zero trust, with arguably the most prominent being the [U.S. Department of Defense's](#) mandate that was released on May 12th, 2021. The mandate says that "incremental improvements will not give us the security we need; instead, the Federal Government needs to make bold changes and significant investments in order to defend the vital institutions that underpin the American way of life."

This message convinced many organizations to modernize their security approach. Like the DoD, many aspirational adopters have a desire to take a modern approach to cybersecurity. For some, the strongest driver is rooted in enabling a secure work-from-anywhere and bring-your-own-device (BYOD) policy. Others want to adopt zero trust to create cost savings through VPN elimination. Finally, for some; it's the [law](#). Whatever the reason, migrating to zero trust takes time and effort. Google's journey took nearly a decade to complete, and is still evolving as new technologies and threats emerge.

How is zero trust different?

Conceptually, one can think of zero trust as a castle. A traditional castle has a wall to guard its citizens and buildings. While that wall is immensely fortified and can provide a sense of safety, it has one major flaw. After you're behind the wall, you can roam about the kingdom freely, gaining access to anything that may be of interest to you, with little resistance along the way. Now, imagine if instead of relying on one giant wall to protect the castle; there were walls, barriers, or locked doors around every building and every room of that castle. In this scenario, if someone were to be inside your castle, the damage potential is significantly reduced as everything valuable has been secured, locked away, and is otherwise inaccessible.

While the castle comparison may provide a conceptual understanding, it is worthwhile to explain how it is technically different. In a traditional perimeter-based model, trusted zones are established that let a user access assets and resources within that zone. While there are controls that can be implemented to increase the security posture of those zones (such as stronger authentication, principle of least privilege, more sophisticated firewalls, logging, and monitoring), it is still a less than optimal approach as there is a wide-blast radius. Zero trust focuses on protecting resources, not protecting network segments, because network location is

no longer a core component of enterprise security posture. Zero trust assumes there is no implicit trust granted to resources based on physical location, network location, or asset ownership. Instead, access relies on dynamic context such as user behavior, device posture, operating system, patch-level, geo-location, and date and time constraints in addition to the conventionally used controls like user identity, role and RBAC-based or ABAC-based permissions. Additionally, resource access is granted on a per session basis, which means that all the previously mentioned context-based verification checks run the next time a request is submitted for that resource after the previous session has expired or terminated.

What are the benefits of zero trust?

There are significant benefits to making the transition to zero trust, and not all of them are rooted in technology. While we won't cover all of the benefits in this document, at a high-level, some of the benefits are the following:

- Reduction in costs
- Better user experience while increasing security posture
- Compliance
- Marketability and competitive differentiation in the marketplace

Zero trust solutioning

“North star”

[NIST 800-207](#) is the primary guiding document for the Google implementation of zero trust architecture. This section describes several concepts, terms, and principles that are critical to understanding some nuances of designing and implementing a zero trust solution.

There are five pillars of NIST’s zero trust architecture that are meant to provide a means of identifying and grouping components of a solution architecture. It is important to underscore that the strength of the pillars is underpinned by a need for a well-defined strategy around governance, automation, visibility, and analytics. NIST defines the pillars as:

- **Identity:** The globally or locally unique identity of users, devices, and resources that interact with the system.
- **Device:** IP address-enabled components that can access secure networks and resources.
- **Application:** Workloads that process, store, and transmit data; and whether or not the workloads are part of the network.
- **Network:** The perimeter that contains the data, assets, applications, and services (DAAS) components that must be protected.
- **Data:** The most important part of the solution that all pillars will process, store, and transmit.

Components

There are two general classifications of components: core and functional. The core components are the MVP of zero trust. Without the core components, the architecture isn’t complete and might not function as desired. The functional components are items that can help increase the efficiency and effectiveness of your solution architecture.

Core components

The core components are the following:

- **Policy engine (PE):** The ultimate decision point that grants access to resources. The PE establishes a trust algorithm to grant, deny, or revoke access to resources. In Google Cloud, the policy engine can be implemented using services such as Identity and Access Management (IAM), Cloud Asset Inventory, and Security Command Center.
- **Policy administrator (PA):** Responsible for establishing and shutting down communication paths between entities and resources. PAs facilitate session-specific authentication, which is used to access resources within an environment. In Google Cloud, the policy administrator can be implemented using services such as IAM, Cloud Asset Inventory, and Security Command Center.
- **Policy enforcement point (PEP):** Responsible for enabling, monitoring, and closing connections between entities and enterprise resources. The PEP forwards requests and

receives policy updates from the PA. The trusted zone is just beyond the PEP. In Google Cloud, the PEP can be implemented using services such as Identity-Aware Proxy (IAP), Cloud Load Balancing, and Cloud Armor.

For more information about Google Cloud services mapping to zero trust components, see [Zero trust capability matrix](#).

Functional components

Functional components might be different, depending on the requirements and position of the enterprise. Typical functional components include the following:

- **Monitoring:** Continuous monitoring helps with early detection for issues such as service issues, potential breaches, and unauthorized requests. These issues can impact production systems and potentially affect downstream revenue.
- **Automation:** Automation improves the efficiency and effectiveness of the trust algorithm and policy engine, all while reducing time spent on manual intervention.
- **Governance:** Governance means having the proper security tooling, identity and access management (IAM) configurations, analytics, incident management playbooks, and operational playbooks.

Operational principles

Zero trust is not only a technical shift, it is an operational shift as well. At the core of the concepts that are outlined below is the phrase “never trust, always verify”.

Operational principles include the following:

- **Never trust:** Eliminating the idea of a trusted zone is often tough to conceptualize. However, with zero trust, there isn't a trusted zone because every user, device, machine, and component is required to authenticate before they are granted access. With zero trust, the Policy Engine checks for items like the latest patch version and [impossible travel](#). Every user, device, workload, application, and data flow is treated as untrusted.
- **Always verify:** Authentication validates the subject and authorization validates that the subject has the right permissions to access the object. Access to any resource must be authenticated and authorized continuously following a least privilege approach using dynamic security policies. There is no such thing as implicit trust in zero trust.
- **Assume breach:** Consciously operate, secure, and defend assets with the assumption that an adversary already has presence within the environment. This principle includes the following:
 - Deny by default and thoroughly scrutinize all users, devices, data flows, and requests for access.
 - Log, inspect, and continuously monitor all configuration changes, resource accesses, and network traffic for suspicious activity.

- Widespread adoption of granular segmentation.
- Limit access to all resources on a network.
- Treat data as an enterprise resource and encrypt it at all points within its lifecycle.
- **Verify explicitly:** Access to all resources is conducted in a consistent and secure manner using dynamic and static attributes. Contextual access decisions are made to resources.
- **Enforce least privilege access:** Access creates risk from a security perspective. A subject needs access to the object but too much access or the presence of that access presents risks to misuse or compromise of the privilege. The principle of least privilege calls for limiting access to the object to the minimum permissions to get the job done.

Access control practices like risk-based (or adaptive) access control, role-based access control (RBAC), and just-in-time (JIT) access have been adopted to implement least privilege.

Everything has an identity and the evaluation of identity and context are what grants access to services.
- **Comprehensive security monitoring:** Zero trust on your network must be upheld at all times. Continuous monitoring of control effectiveness and signal analytics is a critical operational principle.

Project planning

Approaching zero trust

Migrating to zero trust requires an understanding of NIST 800-207 and your organization's posture with respect to the NIST 800-207 framework. You must consider your employees' skill sets, available capacity, existing tooling, and alignment with enterprise commitments when building the path for this journey.

Understand the why

When preparing to start your journey, spend time to thoroughly understand why your organization can benefit from adopting zero trust. Understanding your organization's reasons will enable the teams to clearly define success criteria, metrics, and key requirements.

Additionally, understanding your organization's reasons will guide the priorities within the project plan, demonstrate success, and track improvement. Over a long period of time it is easy to lose sight of the reasons, and it can be difficult to demonstrate success or improvements along the way. Stakeholder support underscores the need for clearly defined metrics and requirements, because much of the work likely will not be highly visible or immediately quantifiable. Metrics are also likely to be different for each company. Some common metrics include items such as:

- Total number of breaches that occur
 - This metric requires that you have accurate data on the number of breaches before you start this journey.
- Total cost of breaches which occur
 - By reducing the affected area, zero trust can reduce impact and costs related to a breach.
- Number of user complaints about security
 - Zero trust can improve user experience.
- Number of users who are able to access applications and data without any issues
- Number or percentage of applications that are protected by zero trust
- Percentage of employees who are aware of zero trust and how it works
- Number or percentage of incidents that are prevented because of zero trust
- Satisfaction survey for employees related to the implementation and use of zero trust

Training

To date, most organizations have become accustomed to operating with implicit trust and, on occasion, transitive trust. Operating in this way builds muscle memory when handling tasks like access requests, creating or removing certain permissions, and monitoring traffic. Zero trust introduces a radically different mode of operation and will challenge existing operational muscle memory.

Having a plan to carefully address this risk and upskill teams will ensure the long term sustainability and success of zero trust for your organization. While every organization will have different priorities, here are some items to consider when building your training plans:

- What new tools have you adopted, what did they replace?
- How familiar are you with the UI?
- Have your workflows been altered due to new tooling, were any operational playbooks or processes for threat mitigation impacted?
- With access control being rooted in the device rather than the firewall, how does this change your operational posture?

Assemble a team

As mentioned, to deliver a successful zero trust journey you must coordinate across business units and multi-functional teams. The following table describes some examples of role profiles and job titles that may be beneficial to include in your journey.

Team	Level	Possible titles
Management team	Executive Provides the driving force that brings key stakeholders to the table and imposes accountability. Most importantly, they make decisions that have a lasting impact.	CxO, SVP, VP, or MD
	Technical manager Works alongside the program and project team to steer the direction of the journey, provide opinions on impactful technical decisions, and ensure metrics have been met.	Head of, manager, or director
	Program and project Works in collaboration with the executive sponsorship. They synchronize project	Program manager, technical program manager, project manager, scrum master

Team	Level	Possible titles
	<p>items like schedule tracking, sprints, migration wave formation, migration playbook readiness, and general reporting.</p>	
<p>Technical teams</p>	<p>Security engineer Unless the corporate structure involves multiple security teams, handles change requests that impact security such as changing firewall rules and changing permissions.</p>	<p>Security architect, security engineer, security analyst</p>
	<p>Networking engineer Handles network-related requests and activities such as setting up isolated networks, assigning IP addresses, changing topology, and dealing with ports, IP addresses, and protocols. The networking engineer works closely with the security engineers and in some cases the two levels are combined.</p>	<p>Systems engineer, systems architect, network administrator</p>
	<p>Infrastructure-as-code (IaC) engineer Writes infrastructure code to create infrastructure based on approved change requests.</p>	<p>Systems engineer, site reliability engineer, DevOps architect, platform engineer, operations engineer</p>

Team	Level	Possible titles
DAAS owners	<p>Workload and application owners</p> <p>Owns the assets and workloads that are migrating to zero trust. Unlike the other teams, these individuals are not stateful, and rotate as their workload is being migrated. The owners are responsible for producing architectural diagrams, playbooks, test plans, and post go-live approvals.</p>	Product owner, product lead, lead developer, product architect

Table 1 - Example of a zero trust team structure

Zero trust change management process

Because many management and technical teams are involved in a typical zero trust migration project, a sound change management process is important to set in place at the very start of the migration. Change management encapsulates all changes that impact production systems through impact analysis by each affected party (networking team issuing and allowlisting a new IP address range or the security team relaxing a security policy for new contractors). A good change process creates accountability, achieves faster consensus, and can harmonize deployments to support alignment between management and technical teams.

Critical phase 1 steps

High priority items to complete at the start of a zero trust journey

Perform a zero trust assessment

The first critical step is conducting a current state assessment, often known as a maturity assessment. The critical state assessment identifies a few key items, such as your posture relative to NIST standards. Before you can design a migration plan, you must understand what your business is doing well from a technical and process perspective and where your business needs to improve.

Second, the critical state assessment lets you understand which services or software within the current environment might have zero trust based capabilities (for example, the ability to enable dynamic context). Again, identifying what services are strong or weak relative to zero trust capabilities lets you understand what to change to strengthen your support for zero trust. Identifying a service's zero trust capabilities also provides insight for required software purchases and the associated costs.

Lastly, but definitely not least; the critical state assessment provides visibility and documentation of your assets, users, traffic flows, dependency mappings, and business logic of workloads.

Conduct the critical state assessment using the NIST 800-207 pillars as your guide. NIST 800-207 outlines five pillars (Identity, Device, Application, Networking, and Data). You can determine your maturity score for each capability in each pillar by checking whether the pillar's capability is not available, but requires manual effort, semi-automated, fully automated, or automated and AI-based or ML-based. For more help, see Google's [Zero Trust Foundations](#) offering, which provides guidance for conducting an assessment and generating key outputs such as maturity scores.

Create the DAAS Inventory

You must identify your data, assets, applications, and services (DAAS) to understand your protect surface. The DAAS inventory is also important when building the migration plan. The following are some examples of what to include in the DAAS inventory:

- **Data:** Databases, data warehouses, or datalakes
- **Applications:** Any web app or non-web app
- **Assets:** Devices, servers, laptops, cameras, resources, and anything that can store, process, or transmit data
- **Services:** APIs or endpoints which can be hosted internally or externally

Within zero trust, DAAS underscores the fact that the scope of work goes beyond application workloads, because building strategies for assets such as mobile devices are core requirements. In some cases, an organization can have tens of thousands of DAAS elements, which can also have interdependencies on one another.

The output of this inventory provides you with the total numbers of the following:

- Data sources, data source categories, and associated classifications
- Applications, application categories, and associated classifications
- Assets, asset categories, and associated classifications
- Services, service categories, and associated classifications
- Users, user categories, and associated classifications
- Devices, device categories, and associated classifications

In addition to the DAAS inventory, it is recommended to also create the inventory of end users and their devices, which could range from company assets to BYODs.

Create the signal inventory

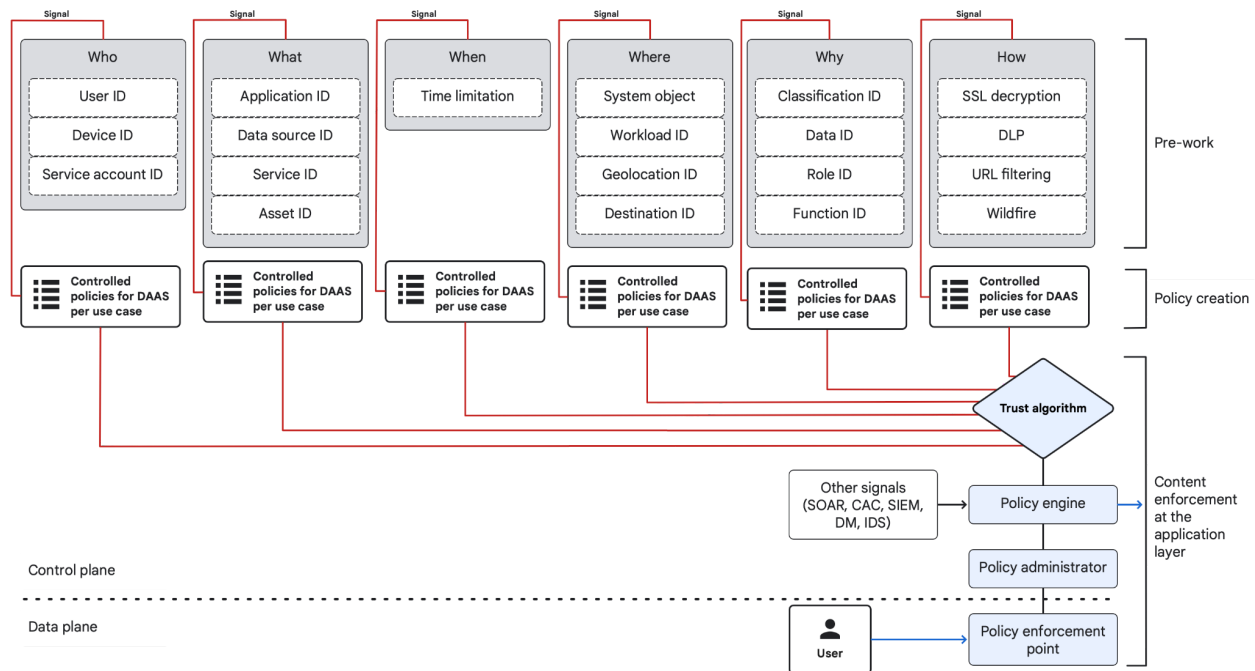
Signals are indicators that can assess the security posture of a device or user and can come from a variety of sources, such as device logs and network traffic. You can also use signals to identify potential threats such as data exfiltration, unauthorized access attempts, and malware infections. Additionally, with regards to zero trust, signals are used by the trust algorithm to make access control decisions. Some examples of how the trust algorithm can use signals for access decisions are the following:

- **Device compliance:** Determine whether a device is compliant with security policies.
- **User behavior:** Track and learn user behavior to identify potential unusual activity.

Identifying adequate security signals (which are layer 7 context-aware) and building the inventory is a critical component on the journey to starting your migration, as it outlines the sources of the signals that enable contextual awareness. As part of the authorization process with zero trust, signals pass through the policy engine to the trust algorithm. Creating a well-built, accurate, and agreed upon signal inventory helps ensure that there is a reliable and effective means of consuming signals on a session basis.

One of the best ways to approach defining your signal inventory is to use the Kipling method. The Kipling method applies to each signal, which translates to “*Who* can do *What* actions on *Which* resources, *When*, *Where*, *Why*, *How*.”

The following diagram visually describes how the NIST 800-207 core components (the policy engine, policy administrator, and policy enforcement point) integrate with an organization's zero trust security signals (that is, who needs access to what resource, where, why, when, and how?).



Raouda, Hani. (2023). Zero Trust context aware authorization . Retrieved June 10, 2023

Figure 2.1 - ZT Policy Signals

Inventory prioritization

After you create the DAAS and signal inventories, the next step in the migration path is to categorize and prioritize the inventories. Inventories can be quite expansive and this step helps organize applications, users, devices and services into smaller categories that have commonalities such as security or regulatory requirements. During this prioritization process, you must [understand the why](#) and what priorities the business is trying to achieve. Aligning the inventories with business priorities helps to create data-driven migration plans.

Complete the zero trust capability gap analysis

The [critical state assessment](#) plays a significant role in being able to complete the capability gap analysis. The capability gap analysis is meant to display how an organization currently compares to a zero trust best practice. For example, if assessing the Device pillar, an example of a capability evaluation could be the following:

- Identify where you align using current tooling and processes

- **Traditional:** Organization has limited visibility into device compliance.
- **Advanced:** Organization employs compliance enforcement mechanisms for most devices.
- **Optimal:** Organization constantly monitors and validates all devices' security posture.

This type of assessment must be conducted across all five pillars of the NIST 800-207 framework. The output of this analysis should identify what products' capabilities are helping or hurting maturity scores, what the easy changes would be, and should provide the business with a view of what potential new products or changes are needed. While there are many tools to help with this analysis, the [CISA maturity model](#) has proven to be a very effective tool to assess and define posture.

Migrating

Tactical

From a process perspective, a zero trust migration is similar to cloud migration projects or general projects, in that standard program management cadences and processes are used (backlog creation, backlog refinement, sprints, retrospectives, and so on). The intricacies of each step are where the previous design concepts, principles, and decisions have the most impact. Use the inventories that you created to determine what the path of least resistance is for your scenario; the path that best aligns with the objectives that are set by the business.

Building the plan

Step 1: Objective-based roadmapping: The most important step in the process is building the roadmap. Having a roadmap that's grounded in reality and mapped to business objectives sets attainable goals and establishes tangible returns on business objectives. After you identify the objectives, you can prioritize them. Ordering objectives in a way that meets the needs of the business and is attainable with the current team helps set clear direction while managing expectations as well as capacity. Setting clear cross-team expectations at this phase will help reduce turmoil and identify wins.

Step 2: Identify first-mover criteria: Starting with simple workloads like SAML-based and modern web applications and SaaS products is highly recommended to allow for quick learning, creation of reusable assets, and quick wins for demonstrated ROI. These migrations are not something that is completed in a silo, overnight, or with a sole vendor. Zero trust is not a product or a service, but a combination of people, processes, and vendor-agnostic technologies. The roadmap to zero trust can be measured in weeks, months, or even years, depending on the complexity of the use cases and the scale of your digital landscape (for example, application complexity, regulatory requirements, and downtime windows are all things that can impact the duration). Also, an executive management's appetite and commitment level play a significant role to the timeline and the outcome of the zero trust journey.

Step 3: Define outcomes: As mentioned, embarking on this journey is no small task as it spans across business units, and will take time. Desired outcomes in a project plan are specific to your organization and the reasons why you started in the first place. Outcomes should be the long pole in a project plan, for instance if the business wants to begin enabling dynamic context first to train a model, that should be one of the major milestones in the project plan. Additionally, having clearly defined success criteria enables the teams to identify tangible wins and also enables the business to know what is coming and when.

Implementing zero trust

Analyst John Kindervag of [Forrester Research](#) has outlined a five-step process that can be applied to any workload type. The following diagram shows these steps, which are further described in the sections below.




Figure 3.1 - Zero Trust 5-step process ([John Kindervag](#))

These implementation steps are designed to be flexible, repeatable, and technology-agnostic. This process lets an organization start with a small, bounded initial protect surface (or set of DAAS elements), work through the rest of the steps with that initial protect surface to establish their approach, and then add additional protect surfaces as their zero trust strategy matures and expands.

Step 1: Define the protect surface

When you define the protect service, you identify the DAAS elements that you want to protect. A protect surface should have only one DAAS component and a zero trust environment can contain one or more protect surfaces. DAAS elements include organizational units, groups, services, service accounts, and more.

The following diagram shows some of the elements that you must consider in this step.




Figure 4.1 - Zero Trust Step 1 - Define the protect surface

Step 2: Map the transaction flows

In this step, you map the transaction flows to and from the protect surface (IP addresses, ports, protocols, and so on.) Transaction flows let you answer how your elements communicate. The following diagram shows an example of transaction flows.

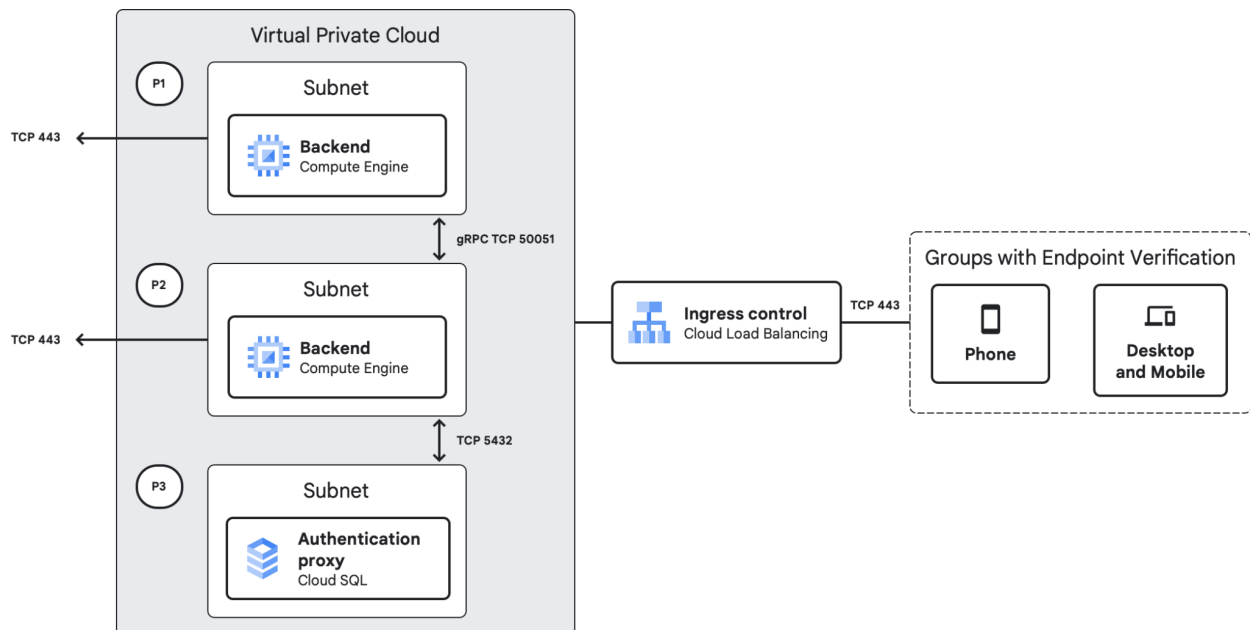


Figure 4.2 - Zero Trust Step 2 - Map Transaction Flows

Step 3: Build a zero trust architecture

Each zero trust architecture is tailor-made for every different protect surface. As you design the zero trust environment, consider how you will restrict access for unauthenticated users and add Layer 7 context awareness. The following diagram shows some Google Cloud services that you can consider as part of your design.

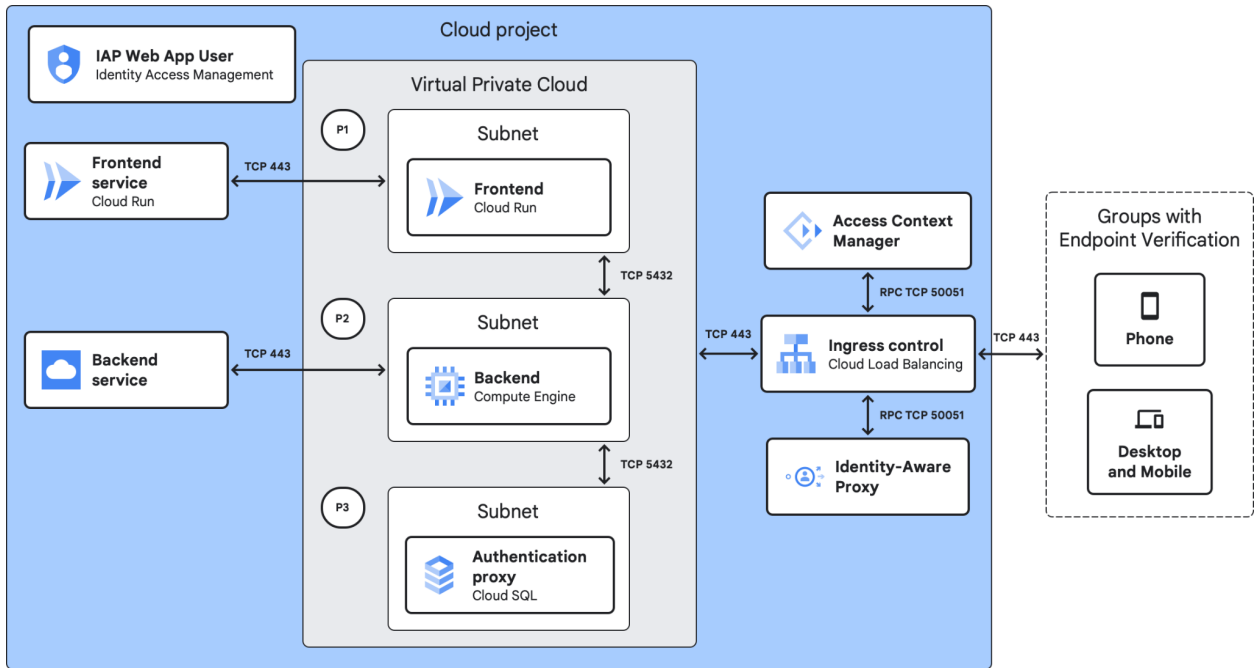


Figure 4.3 - Zero Trust Step 3 - Build Zero Trust Architecture

Step 4: Create a zero trust policy

During this step, you define and map the zero trust Layer 7 dynamic context signals to each protect surface using the Kipling method (who, what, where, when, why, and how). The following diagram describes signals that you can consider, various inputs to a policy, and how to build context enforcement.

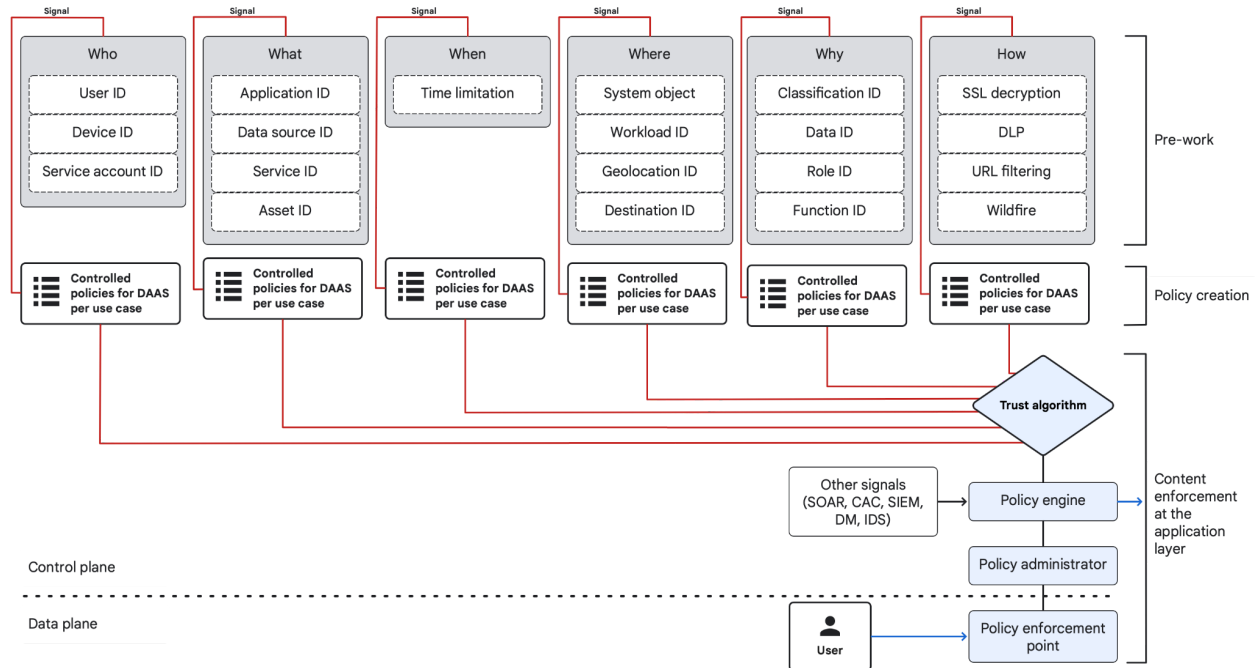


Figure 4.4 - Zero Trust Step 4 - Create Zero Trust Policy

Step 5: Monitor and maintain the network

You must set up the ability to inspect logs, monitor violations, and which cases require appropriate action. Depending on the services used, telemetry from cloud endpoints, on-premises endpoints, and other endpoints are analyzed using dynamic attributes.

Services that can help you include the following:

- Security Command Center for risk management
- Chronicle SIEM
- Chronicle SOAR
- Mandiant incident response services
- Automation of infrastructure as code
- Cloud Logging and Cloud Monitoring

Architecture

Tactical steps

Implementing zero trust properly requires a well thought-out design that aligns business objectives, industry standards, and a best-in-class technology stack. Meeting these objectives is where zero trust architecture comes into play. The NIST 800 207 publication says that “Zero trust architecture is an enterprise cybersecurity architecture that is based on zero trust principles and designed to prevent data breaches and limit internal lateral movement.”

Design concepts

The components and operational principles play a factor in developing the architecture. However, there are also design concepts and patterns to help in designing your architecture on Google Cloud. The following concepts guide decision making along the journey:

- **Define mission outcomes:** Derive the zero trust architecture from organization-specific mission requirements that identify the critical DAAS.
- **Architect from the inside out:** First focus on protecting critical DAAS, then focus on security for all access paths to the critical DAAS.
- **Create access control policies:** Create security policies and apply them consistently across all environments (for example, LAN, WAN, endpoint, perimeter, mobile, and so on).
- **Inspect and log all traffic:** Establish full visibility of all activity across all layers from endpoints and the network to enable analytics that can detect suspicious activity.

Using these concepts, you can start a zero trust journey at any point of maturity and transform an enterprise application or infrastructure so that they follow Zero trust principles.

The following diagram shows how trust is propagated in a traditional perimeter-based architecture.




Figure 5.1 - Traditional perimeter-based architecture (Castle walls)

You must examine traditional architectures of implicit trust in your organization. Using a phased approach, migrate to granular policy-driven access controls based on zero trust components.

The following diagram shows how zero trust architecture determines access using policy-based authorization. The connection flows from the user to the data using the different zero trust components on an on-premises network. In this architecture, a central policy administrator (with a trust algorithm) is used to evaluate trust between identity, device, network, application, and data components. PEPs are deployed between each component.

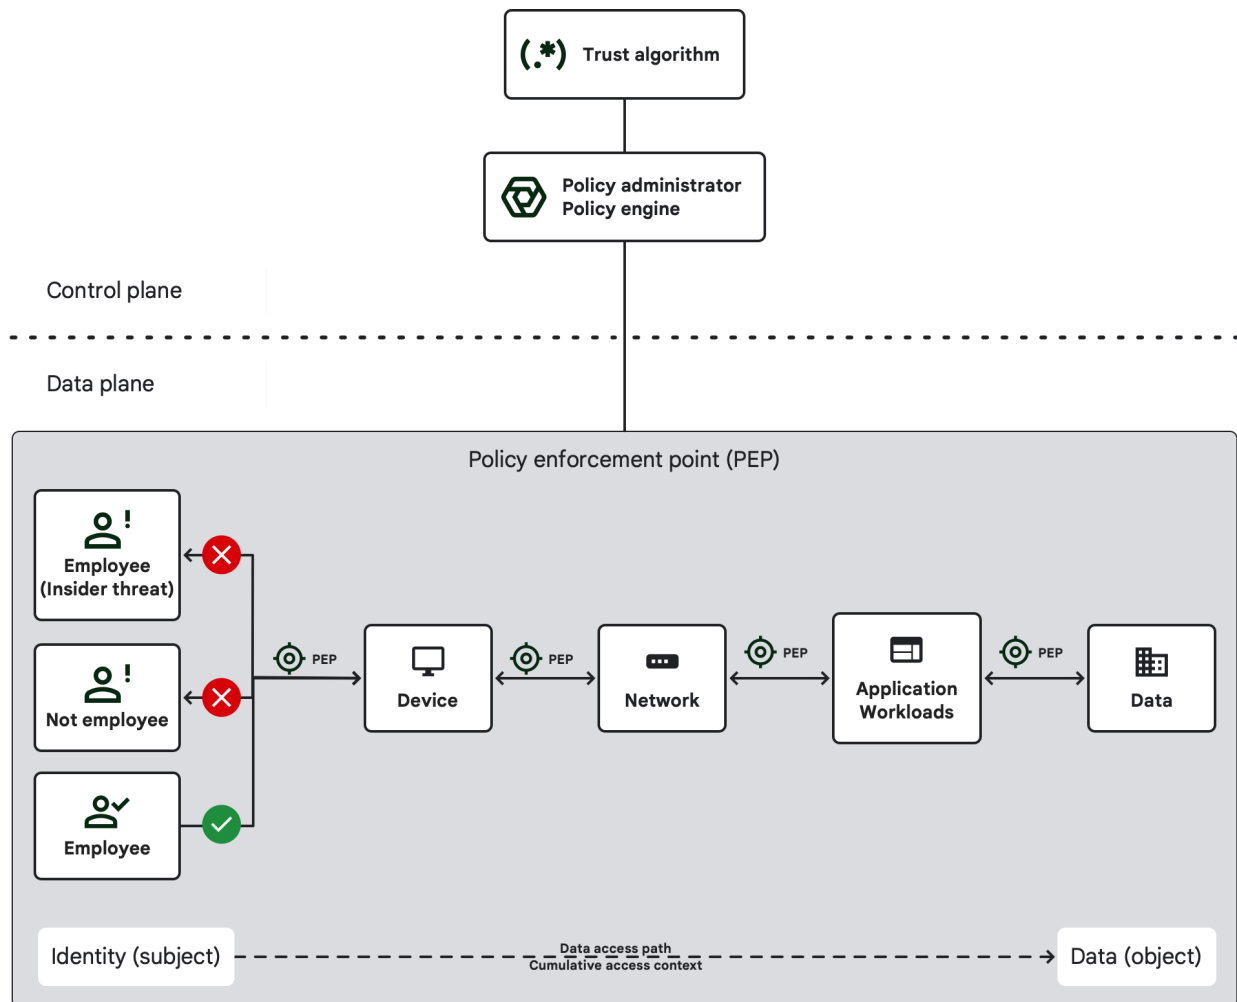


Figure 5.2 - Zero Trust architecture - With zero trust, enterprise trust is decoupled and access (propagation) is controlled by policy-based authorization

Implementing zero trust architecture as you migrate your workloads to the cloud makes migration more seamless because most cloud services support zero trust by default. Cloud services have components like PEP for authentication and authorization that enable continuous policy evaluations with overarching policy administrators and policy engines.

Additionally, foundational components like logging and monitoring, governance, and automation are default services on Google Cloud.

The following diagram shows how zero trust can be applied on a cloud environment by using cloud services.

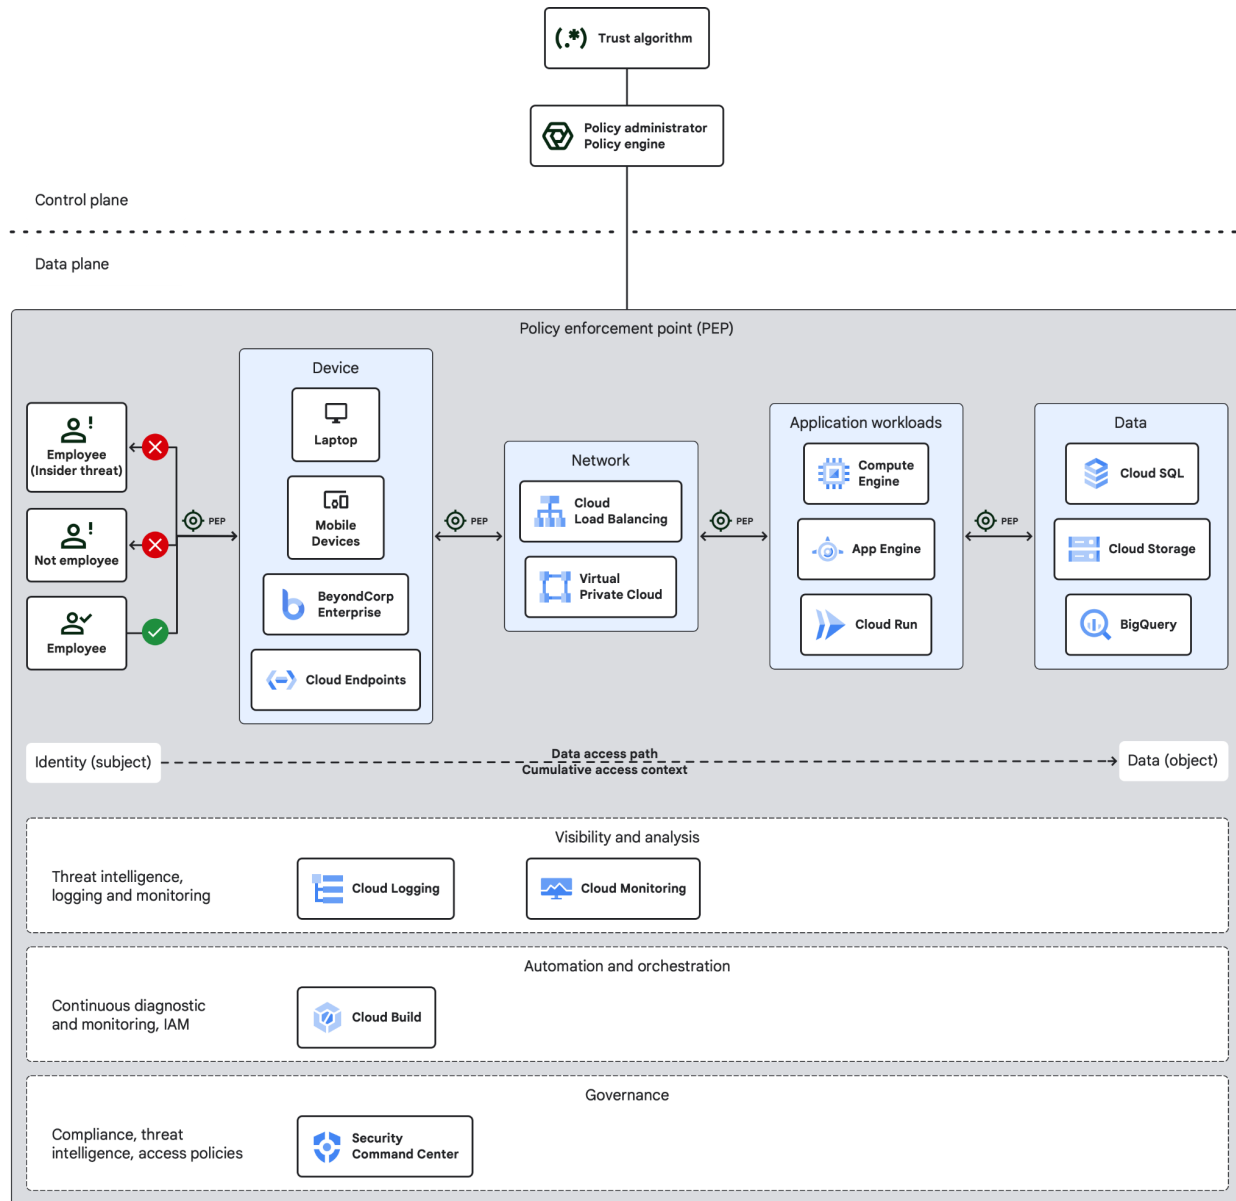


Figure 5.3 - Cloud-based zero trust environment

With cloud services such as Google Cloud services, determining access using policy-based authorization is a default as authorization is waived in all Google Cloud services. In this diagram, Cloud IAM and Cloud Identity are the main policy administrator components that run the trust algorithm for the IAM policies. You can use Access Context Manager as one of the main PEPs to evaluate trust as connection flows from the user's device with BeyondCorp to the network (Virtual Private Cloud (VPC), or load balancers), application (App Engine, Cloud Run, or Compute Engine) to data services (Cloud SQL, Cloud Storage, or BigQuery).

Design patterns

The following patterns illustrate how design concepts can be used to implement zero trust Architecture in common modern cloud architecture patterns.

1. Enterprise with remote employees

This design pattern includes a single headquarter and one or more geographically dispersed locations that are not joined by an enterprise-owned physical network connection. Employees may be teleworking or in a remote location and using enterprise-owned or personally-owned devices. This architecture leverages policy enforcement and policy administration points in the cloud. The following diagram shows an example of this design pattern.

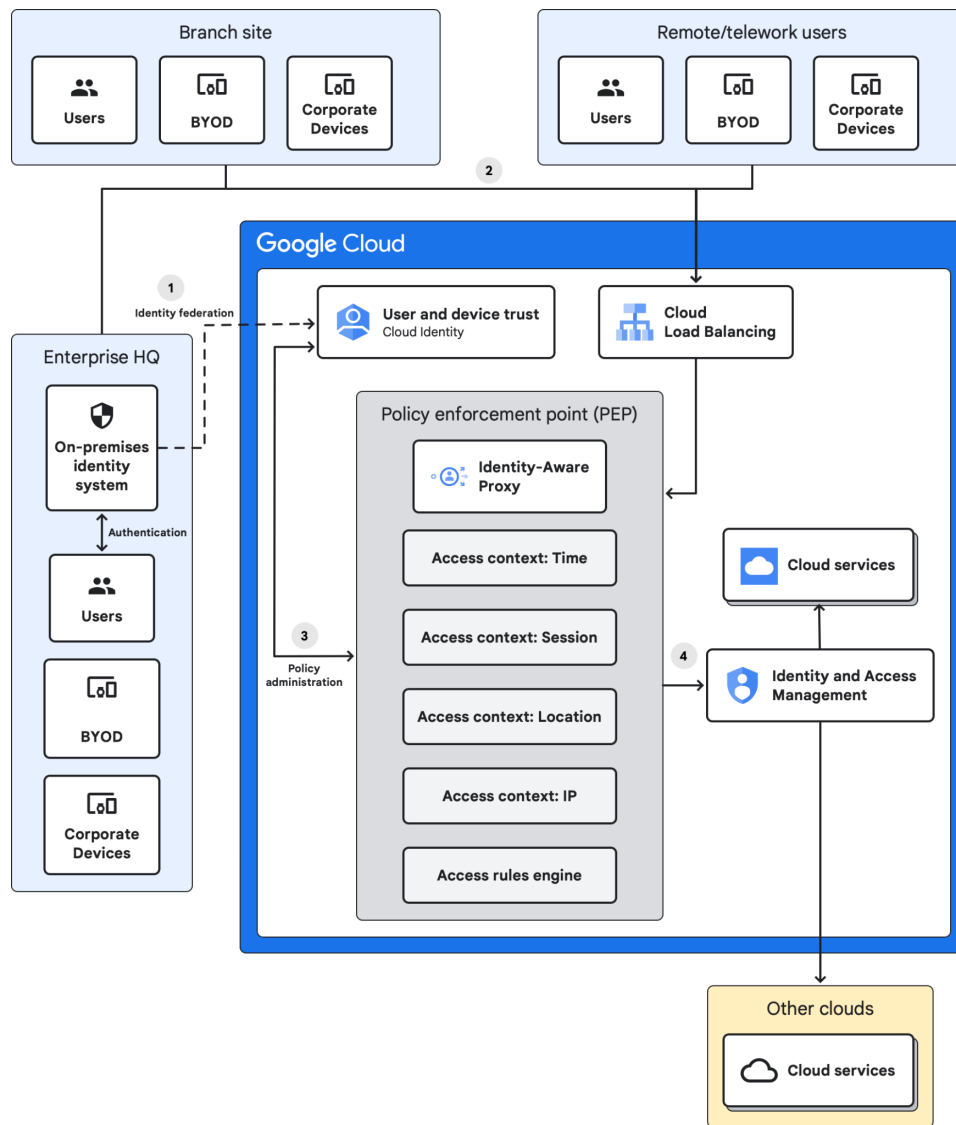


Figure 6.1 - Zero Trust Pattern: Enterprise with remote employees

In the previous diagram, the numbers describe the following:

1. Corporate identity systems are federated with a cloud identity as a service (IDaaS) so that the business can keep one central identity system rather than maintain independent systems.
2. Enterprise users from the enterprise headquarters, branch sites, or remote networks use corporate devices or managed BYOD to request services access. Cloud Identity checks user and device trust.
3. The Identity Aware Proxy (IAP) PEP performs policy administration that is based on user and device context.
4. After user and device policies are checked, users are granted access to cloud services based on their roles and permissions.

2. Enterprise with a multi-cloud application

This design pattern includes an enterprise that uses multiple cloud service providers (CSPs) to host an application, services, or data. The following diagram shows an application that is hosted in CSP-A and that can connect directly to a data source in CSP-B without tunneling back to the enterprise network.

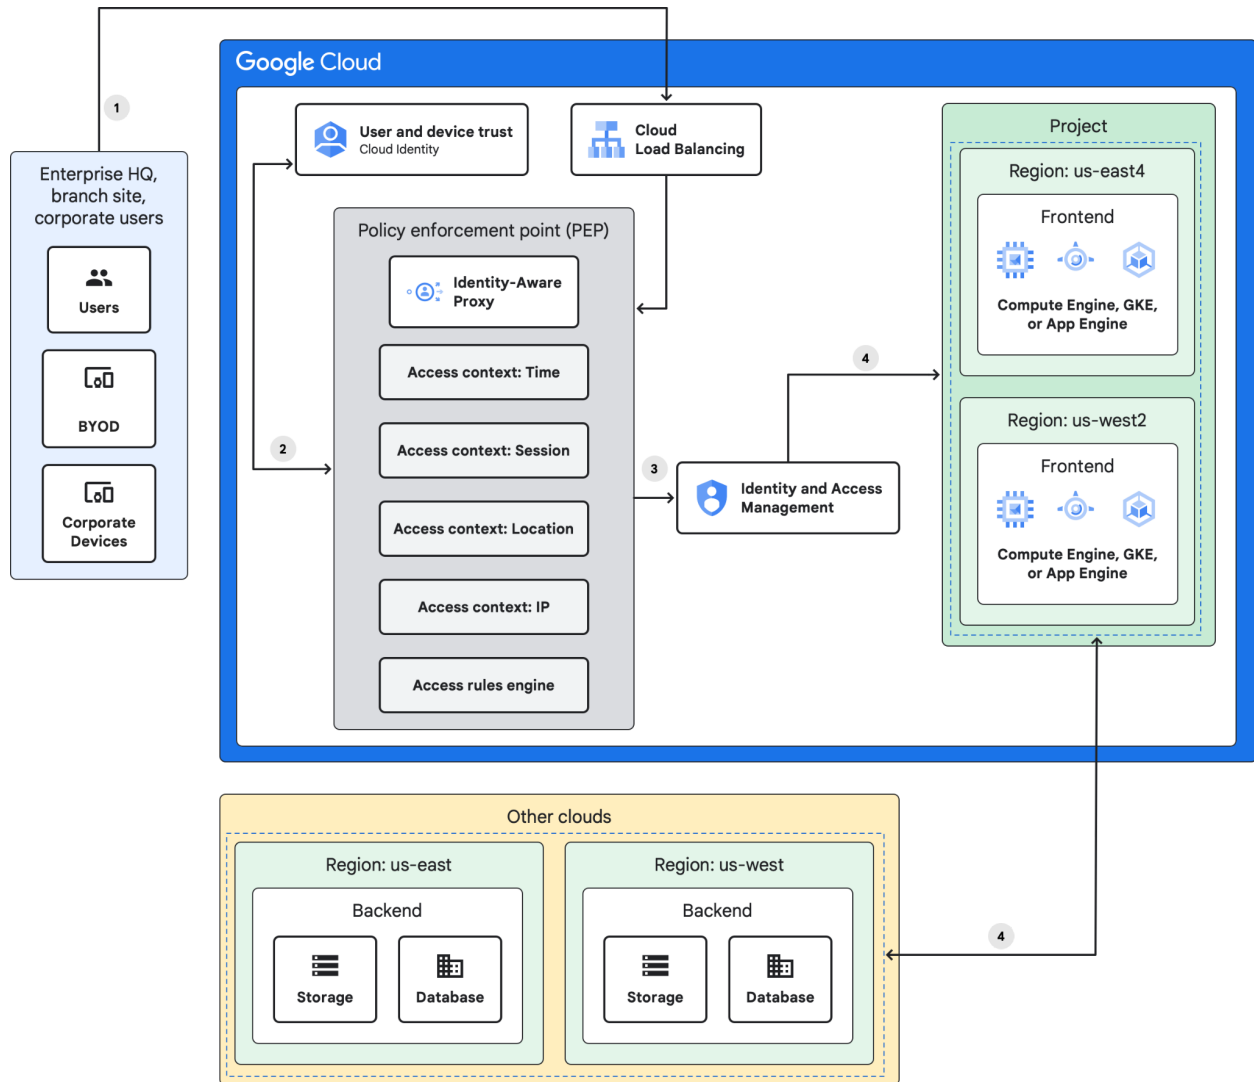


Figure 6.2 - Zero Trust Pattern: Enterprise with a multi-cloud application

In the previous diagram, the numbers describe the following:

1. Enterprise users from the enterprise network, branch sites, or remote networks use corporate devices or managed BYOD to access a multi-cloud application.
2. The IAP PEP performs policy administration that is based on user and device context.
3. After user and device policies are checked, users are granted access to the application based on their roles and permissions.

- Frontend services for the cloud application can be hosted in one cloud, with direct peering, VPN connection, or private connection to another cloud where the application hosts the backend services.

3. Cross-enterprise collaboration

This design pattern includes two enterprises that are working together and have a federated identity management system. Policy enforcement and policy administration provides access to hosted cloud services. The following diagram shows Enterprise-A and Enterprise-B as separate organizations working on a joint project.




Figure 6.3 - Zero Trust Pattern: Cross-enterprise collaboration

In the previous diagram, the numbers describe the following:

- Enterprise identity systems for both organizations are federated with Cloud Identity so that users can access services from anywhere.
- Enterprise users from both organizations use corporate devices or managed BYOD to access cloud services from the corporate network, branch sites, or remotely.
- The IAP PEP performs policy administration that is based on user and device context.

4. After user and device policies are checked, users are granted access to cloud services based on their roles and permissions.
5. Users from Enterprise-B only have access to shared cloud services. Enterprise-A users can access shared and private cloud resources.

4. Enterprise with non-employee access

This design pattern includes an enterprise that has contracted service providers who require limited access to corporate resources to do their work. Visiting providers get access to the Internet and limited resources. The following diagram shows how corporate staff and contractors can get access to resources.




Figure 6.4 - Zero Trust Pattern: Enterprise with non-employee access

In the previous diagram, the numbers describe the following:

1. Enterprise users are grouped by personnel type: staff or contractor. Contractors get Internet access using a guest network. The corporate identity system is federated with Cloud Identity, which groups the users and devices for staff and contractor.
2. Enterprise users and contractors use corporate devices or managed BYOD to access the Internet and cloud services from the corporate network.

3. The IAP PEP performs policy administration that is based on group association, and user and device context.
4. After user and device policies are checked, users are granted access to cloud services based on their group, roles, and permissions.
5. Contracted users only have access to limited, shared cloud resources. Enterprise users can access shared and private cloud resources.

5. Enterprise with public-facing services

This design pattern includes an enterprise with a service that is available to the public. Requesting assets aren't enterprise-owned so the tenets of zero trust do not directly apply. The following diagram shows some of the steps that an enterprise can take to limit attacks to the public service.




Figure 6.5 - Zero Trust Pattern: Enterprise with public-facing services

In the previous diagram, the numbers describe the following:

1. The enterprise identity for the corporate system is federated with Cloud Identity so that users can access services from anywhere.

2. Enterprise users use corporate devices or managed BYOD to access cloud services.
3. Public users connect to the public-facing web app through a cloud load balancer with DDoS protection.
4. For enterprise users, an IAP PEP performs policy administration.
5. Enterprise resources and public resources are separate and are regulated by access control policies.
6. Enterprise users are granted access to enterprise cloud services, and public users are only granted access to the web application services.

Zero trust capability matrix

The following table describes the components of a zero trust architecture and how Google Cloud services map to the various components.

It is important to note that Google started our zero trust journey before NIST defined the zero trust core components (that is, the policy engine, policy administrator, and policy enforcement point). As a result, the table below contains multiple Google Service services that can be mapped to a combination of NIST 800-207 Zero Trust components instead of just a single mapping.

Component	Google Cloud capabilities		
	Good	Better	Best
<p>Policy engine (PE)</p> <p>Core component</p> <p>This component is responsible for orchestrating context-aware signals to provide the ultimate decision of granting or denying access to a resource for a given subject.</p>	<p>Access Context Manager</p> <p>Standard policy attributes:</p> <ul style="list-style-type: none"> • Identity • Device 	<p>Access Context Manager</p> <p>Standard policy attributes:</p> <ul style="list-style-type: none"> • Identity • Device 	<p>Access Context Manager</p> <p>Advanced policy attributes:</p> <ul style="list-style-type: none"> • Identity • Device • Browser • Third-party signals
<p>Policy administrator (PA)</p> <p>Core component</p> <p>This component is responsible for establishing and shutting down the communication path between a subject and a resource (by using commands to relevant PEPs).</p>	<p>IAM</p> <p>Cloud Armor</p> <p>Cloud Asset Inventory</p> <p>Cloud Load Balancing</p> <p>IAP for on-premises</p>	<p>IAM</p> <p>Cloud Armor</p> <p>Cloud Asset Inventory</p> <p>Cloud Load Balancing</p> <p>IAP for on-premises</p>	<p>IAM</p> <p>Cloud Armor</p> <p>Cloud Asset Inventory</p> <p>Cloud Load Balancing</p> <p>IAP for on-premises</p>

Component	Google Cloud capabilities		
	Good	Better	Best
<p>Policy enforcement point (PEP)</p> <p>Core component</p> <p>This system is responsible for enabling, monitoring, and eventually terminating connections between a subject and an enterprise resource. The PEP communicates with the PA to forward requests and receive policy updates from the PA.</p>	<p>Gateway:</p> <p>GFE</p> <p>Cloud Identity</p> <p>Client-side:</p> <p>Chrome Enterprise</p> <ul style="list-style-type: none"> - Endpoint verification extension - Helper (for non-browser, client-side applications) 	<p>Gateway:</p> <p>GFE</p> <p>Cloud Identity</p> <p>IAP</p> <p>Client-side:</p> <p>Chrome Enterprise</p> <ul style="list-style-type: none"> - Endpoint verification extension - Helper (for non-browser, client-side applications) - Context-aware data protection 	<p>Gateway:</p> <p>GFE</p> <p>Cloud Identity</p> <p>IAP</p> <p>Client-side:</p> <p>Chrome Enterprise</p> <ul style="list-style-type: none"> - Endpoint verification extension - Helper (for non-browser, client-side applications) - Context-aware data protection - MDM
<p>Continuous diagnostics and mitigation (CDM) system - Endpoint protection</p> <p>Indirect component</p> <p>This system gathers information about the enterprise asset's current state and applies updates to configuration and software components.</p>	<p>Calling endpoint status and configuration:</p> <p>Chrome Enterprise</p> <ul style="list-style-type: none"> -Endpoint verification extension - Helper (for non-browser, client side applications) <p>Managed device inventory and configuration:</p> <p>Google Endpoint Management</p>	<p>Calling endpoint status and configuration:</p> <p>Chrome Enterprise</p> <ul style="list-style-type: none"> -Endpoint verification extension - Helper (for non-browser, client side applications) <p>Managed device inventory and configuration:</p> <p>Google Endpoint Management</p>	<p>Calling endpoint status and configuration:</p> <p>Chrome Enterprise</p> <ul style="list-style-type: none"> -Endpoint verification extension - Helper (for non-browser, client side applications) <p>Managed device inventory and configuration:</p> <p>Google Endpoint Management</p>

Component	Google Cloud capabilities		
	Good	Better	Best
	Cloud Identity Premium	Cloud Identity Premium Security Command Center Premium	Cloud Identity Premium Security Command Center Premium Google Cloud asset inventory: Cloud Asset Inventory VM Manager Artifact Registry
<p>Industry compliance system</p> <p>Indirect component</p> <p>This system ensures that the enterprise remains compliant with any regulatory regime that it may fall under (such as FISMA, or any healthcare or financial industry information security requirements).</p>	<p>Google Cloud asset compliance: Security Command Center Premium</p>	<p>Google Cloud asset compliance: Security Command Center Premium</p> <p>Google Cloud compliance configuration: Risk and compliance as code</p> <p>Security foundations blueprint</p>	<p>Google Cloud asset compliance: Security Command Center Premium</p> <p>Google Cloud compliance configuration: Risk and compliance as code</p> <p>Google Cloud compliance architecture: Assured Workloads</p>

Component	Google Cloud capabilities		
	Good	Better	Best
<p>Threat intelligence feeds</p> <p>Indirect component</p> <p>These feeds provide information from internal or external sources that help the policy engine make access decisions.</p>	<p>Service delivering it:</p> <p>SafeBrowsing Anti-phishing or malicious URLs malware protection or VirusTotal (Not dependant on product, but uses the feeds)</p> <p>Cloud Identity</p>	<p>Service delivering it:</p> <p>SafeBrowsing Anti-phishing or malicious URLs malware protection or VirusTotal (Not dependant on product, but uses the feeds)</p> <p>Cloud Identity</p>	<p>Service delivering it:</p> <p>SafeBrowsing Anti-phishing or malicious URLs malware protection or VirusTotal (Not dependant on product, but uses the feeds)</p> <p>Cloud Identity</p> <p>Source: Google Cloud Threat Intelligence</p>
<p>Network and system activity logs</p> <p>Indirect component</p> <p>This enterprise system aggregates asset logs, network traffic, resource access actions, and other events that provide real-time (or near real-time) feedback on the security posture of enterprise information systems.</p>	<p>Endpoint logging:</p> <p>Cloud Audit Logs Chrome browser logs</p> <p>Google Cloud asset logging:</p> <p>Cloud Logging Policy Intelligence</p> <p>Network telemetry:</p> <p>VPC Flow Logs Cloud DNS logs Cloud NAT logs Firewall logs</p>	<p>Endpoint logging:</p> <p>Cloud Audit Logs Chrome browser logs</p> <p>Google Cloud asset logging:</p> <p>Cloud Logging Policy Intelligence</p> <p>Network telemetry:</p> <p>VPC Flow Logs Cloud DNS logs Cloud NAT logs Firewall logs</p>	<p>Endpoint logging:</p> <p>Cloud Audit Logs Chrome browser logs</p> <p>Google Cloud asset logging:</p> <p>Cloud Logging Policy Intelligence</p> <p>Network telemetry:</p> <p>VPC Flow Logs Cloud DNS logs Cloud NAT logs Firewall logs</p>

Component	Google Cloud capabilities		
	Good	Better	Best
	HTTPS Load Balancing logs VM logs Security Command Center findings	HTTPS Load Balancing logs VM logs Security Command Center findings Remediation Cloud Functions Pub/Sub Chronicle	HTTPS Load Balancing logs VM logs Security Command Center findings Remediation Cloud Functions Pub/Sub Chronicle Siemplify (SOAR)
<p>Data access policies</p> <p>Indirect component</p> <p>These policies are the attributes and rules that define access to enterprise resources. This set of rules could be encoded in (using a management interface) or dynamically generated by the PE.</p>	BeyondCorp Threat and Data Protection Access Context Manager	BeyondCorp Threat and Data Protection Access Context Manager Advanced policies Workload or resource: IAM (Conditions for just-in-time for Google Cloud workloads) VPC Service Controls	BeyondCorp Threat and Data Protection Access Context Manager Workload or resource: IAM (Conditions for just-in-time for Google Cloud workloads) VPC Service Controls Data: Data Catalog and Cloud Data Loss Prevention Encryption: Cloud Key Management Service (KMS) .

Component	Google Cloud capabilities		
	Good	Better	Best
			Cloud HSM , Cloud External Key Manager Secret Manager
Enterprise public key infrastructure (PKI) Indirect component This system is responsible for generating and logging certificates issued by the enterprise to resources, subjects, services, and applications.	Certificate Authority Service Audit logging	Certificate Authority Service Audit logging	Certificate Authority Service Audit logging Cloud KMS , Cloud HSM , Cloud EKM
ID management system (IAM) Indirect component This system is responsible for creating, storing, and managing enterprise user accounts and identity records (for example, lightweight directory access protocol (LDAP) server).	Cloud Identity IAM	Cloud Identity IAM Certificate Authority Service Google Cloud Directory Sync	Cloud Identity IAM Certificate Authority Service Google Cloud Directory Sync
Security information and event management (SIEM) system Indirect component This system collects security-centric information for later analysis. This data is then used to refine policies and warn of possible attacks against enterprise assets.	Chronicle	Chronicle	Chronicle Chronicle SecOps
Network segmentation gateway Indirect component A Layer 7 control that is designed to segment network traffic based on users, applications, and data.	VPC Service Controls, Cloud Firewall		

Component	Google Cloud capabilities		
	Good	Better	Best
<p>ID assertion service</p> <p>Indirect component</p> <p>This service authenticates users with the environment that includes the protect surfaces, before any Layer 7 inspections take place.</p>	Workload Identity federation		

Table 2 - Zero trust capability matrix

Mapping Google Cloud services to the NIST 800-207 pillars

The following table describes the various Google Cloud services that meet the requirements for the NIST 800-207 pillars.

NIST 800-207 pillar	Capability and service	Policy, service, or function to enable for zero trust
Identity	Identity awareness IAP	Minimum capability: Backend of an external load balancer Ideal capability: End-to-end authentication and authorization of application (apply IAP on frontend, backend, and networks)
	Context management	Minimum capability: Access Context Manager with basic signals (for example, geolocation and IP address on DAAS) Ideal capability: Access Context Manager with advanced signals (for example, device, browser, OS, identity, geolocation, IP address, screen lock, or encrypted disk)
	Access analysis	Minimum capability: IAM recommender and insights, BeyondCorp Enterprise Policy Troubleshooter. Ideal capability: Security analytics (for example, logs, Recommender and insights, Policy Analyzer, and Cloud Asset Inventory)
Networking	Intrusion detection	Minimum capability: Cloud IDS
	SSH or RDP access	Minimum capability: IAP with TCP forwarding and OS Login policy Ideal capability: Force the user to SSH or RDP through Cloud IAM on private IP address resources
	On-premises support	Minimum capability: Hybrid connectivity (Cloud Interconnect or Cloud VPN) Ideal capability: Hybrid network endpoint groups (NEGS)
	SSL offloading	Minimum capability: External proxy Network Load Balancer (can help in performing deep packet inspection (DPI) after the data reaches the web server level)

NIST 800-207 pillar	Capability and service	Policy, service, or function to enable for zero trust
		Ideal capability: Secure Web Proxy (can help in performing DPI before data reaches the web server level)
	Intelligence	Minimum capability: Network Intelligence Center and Security Command Center Ideal capability: Network Intelligence Center, Chronicle, and Security Command Center Premium
	Private connection	Minimum capability: Based on use case and need: <ul style="list-style-type: none"> ● Google Private Access ● Private services access ● Private Service Connect ● Serverless VPC Access
Data	Packet mirroring	Minimum capability: VPC packet mirroring (100% packet capture)
	Deep Packet Inspection (DPI)	Minimum capability: SSL proxy load balancer (can help in performing DPI after the data reaches the web server level) Ideal capability: Secure Web Proxy (can help in performing DPI before data reaches the the web server level)
	Data Loss Prevention (DLP)	Capabilities: Cloud DLP API, Healthcare API, Contact Center AI, Gmail DLP, Workspace DLP (such as chat or docs), trust rules, and so on Ideal capability: Based on use case and need
	Exfiltration prevention	Capabilities: VPC Service Controls with ingress and egress policies and allowed services, organization policies, network firewall policies, and Cloud Firewall
Device	Endpoint security	Minimum capability: BeyondCorp Enterprise with Endpoint Verification, and Google Workspace Enterprise edition Ideal capability: Based on use case and need: Endpoint verification with BeyondCorp Enterprise, and Google Workspace Enterprise edition
	Certificate-based authentication Certificate-based access	Capabilities: Access Context Manager and IAP Ideal capability: Access Context Manager with

NIST 800-207 pillar	Capability and service	Policy, service, or function to enable for zero trust
		IAP, BeyondCorp Enterprise, and internal and external Network Load Balancers
	Agentless BYOD	Capabilities: BeyondCorp Enterprise with Endpoint Verification, protected profiles, and Chrome browser
	Agent BYOD	Capabilities: Not applicable Ideal capability: Third party
Cutting edge	SIEM integration	Capabilities: Chronicle
	SOAR/XDR	Capabilities: Chronicle SOAR (formerly Siemplify)
	AI/ML-Based threat detection	Capabilities: Cloud Identity, Cloud IDS
	Bot traffic management	Capabilities: Global external Application Load Balancer with Cloud Armor
	Compliance as Code Assured Workloads OSCAL	Capabilities: Terraform Vet, Policy Controller with OPA Ideal capability: RCaC using Open Security Controls Assessment Language (OSCAL), OPA, and Common Expression Language (CEL)
	Policy as Code	Capabilities: KRM , Terraform Vet

Table 3 - NIST 800-207 pillars to GCP service mapping

BeyondCorp Enterprise

The following table provides information about BeyondCorp Enterprise capabilities. The tiers are the following:

- **Baseline:** A free tier that is offered as part of Google Cloud and comes with basic conditional controls (IP address and geographic location) on Google Cloud resources.
- **Enterprise Essential:** A paid tier that uses user, device, Chrome and third-party context signals for Google Workspace, Google Cloud resources, and SAML applications.
- **Enterprise Plus:** A paid tier that uses user, device, Chrome and third-party context signals for Google Workspace, Google Cloud resources, SAML applications, on-premises resources, and resources on other clouds.

Category	Feature	Baseline	Enterprise Essential	Enterprise Plus
Application and resource access	Access control to web applications on Google Cloud	Yes (user context)	No	Yes (user, device, Chrome context)
	Access control to SSH, RDP, TCP ports for VMs on Google Cloud	Yes (user context)	No	Yes (user, device, Chrome context)
	Access control to Google Cloud APIs	Yes (user context)	No	Yes (user, device, Chrome context)
	Access control to Google Cloud console	Yes (user context)	No	Yes (user, device, Chrome context)
	Access control to on-premises web applications	No	No	Yes (user, device, Chrome context)
	Access control to thick client or client-server applications	No	No	Yes (user, device, Chrome context)
	Access control to web applications on Amazon Web Services and Microsoft Azure	No	No	Yes (user, device, Chrome context)

Category	Feature	Baseline	Enterprise Essential	Enterprise Plus
	Access control to SAML-based applications where Google is IdP (login time only and not dynamic context awareness)	No	Yes (user, device, Chrome context)	Yes (user, device, Chrome context)
	Access control to Google Workspace Admin console	No	Yes (user, device, Chrome context)	Yes (user, device, Chrome context)
Access policies and advanced settings	Access levels using users and groups	Yes	Yes	Yes
	Access levels using IP addresses and geoLocation	Yes	Yes	Yes
	Access levels using date and time restrictions	No	Yes	Yes
	Access levels using login credential strength	No	Yes	Yes
	Access levels using enterprise certificates	No	Yes	Yes
	Access levels using device security posture	No	Yes	Yes
	Access levels using Google Chrome security posture (CBCM and protected profiles)	No	Yes	Yes

Category	Feature	Baseline	Enterprise Essential	Enterprise Plus
	Access levels using third-party signals	No	Yes	Yes
	Access levels using advanced expression language (not just basic)	No	Yes	Yes
	Same origin policy configuration in HTTP options	Yes	Not applicable (Enterprise Essential applies only to Google Workspace)	Yes
	Custom authentication and 403 pages	No	Not applicable	Yes
User, threat, and data protection category	DLP (Chrome)	No	Yes	Yes
	Malware protection and advanced sandboxing (Chrome)	No	Yes	Yes
	Phishing and malicious URL protection (Chrome)	No	Yes	Yes
	Threat and data protection alerting and reporting (Chrome)	No	Yes	Yes
	Password exfiltration (Chrome)	No	Yes	Yes

Table 4 - BeyondCorp Enterprise