Cybersecurity Action Team

Applying the NCSC Zero Trust Principles on Google Cloud

Google Cloud Whitepaper
March 2022

Google Cloud

Table of Contents

DISCLAIMER: This whitepaper applies to Google Cloud products described at cloud.google.com. The content contained herein is correct as of December 2021 and represents the status quo as of the time it was written. Google's security policies and systems may change going forward, as we continually improve protection for our customers.

Introduction

Advanced cyber attacks pose an unprecedented threat to critical infrastructure and mission systems for tens of thousands of organizations. Now more than ever, a zero trust approach to security is necessary for countering those attacks, reducing risks, and aligning with guidance from bodies such as NIST and NCSC.

In 2009, Google began an internal initiative - called BeyondCorp - to reimagine their security architecture with regards to how employees and devices access internal applications.

BeyondCorp is Google's implementation of the zero trust model. It builds upon a decade of experience at Google, combined with ideas and best practices from the community. By shifting access controls from the network perimeter to individual users, BeyondCorp enables secure work from virtually any location without the need for a traditional VPN."

Unlike the traditional perimeter security model, BeyondCorp dispelled the notion of network segmentation as the primary mechanism for protecting sensitive resources. Instead, all Applications were made accessible through user and device centric controls in conjunction with contextual based policy engine - that allows for the authorized **user** access an authorized **application** in the approved **context**.

A zero trust model recognises that:

- **No single component within an interdependent network can be trusted implicitly.** A secure-by-design architecture grants specific, limited permissions for authenticated requests.
- **Perimeter-based authentication located at the boundary of the network is no longer sufficient**. Trust cannot be assumed simply because an actor is within the network, but must be continually monitored and authenticated based on behaviour and context.
- **Authentication of more than just "users" is necessary**, since software and services run automated functions in the cloud. Those components also require a form of credentials and continuous monitoring of their behaviour on the network alongside human actors.
- As workloads shift from physically isolated machines to shared/containerised workloads, **strong isolation practices must be reintroduced** so that, if an attack occurs, it can be walled-in to the compromised domain while the remainder of the resources are kept secure.
- The rapid diversification of devices accessing resources owing to **mass working from home and the participation of third parties in the supply chain, multiplies the administrative effort** of device-level security control and requires additional monitoring/analytics tools
- **Enforcement of security policies should no longer be applied on a per-application basis**, but via centralized policy management tools, which can consistently set and verify policies based on the tier of risk/permissions that a given app is assigned.

The criticality of these factors has accelerated in recent years, with increased adoption of Cloud services; a global pandemic driving mass work-from-home and an increasingly complex cyber threat landscape requiring specialist skills that are in high demand.

While the importance of a zero trust strategy is clear, the path to one is not. For many organizations, adopting a zero trust posture means significantly changing how resources and security are managed. The journey to zero trust is complex and lengthy - meaning it is best achieved with a comprehensive roadmap of changes, migrations, and deployments.

This whitepaper outlines how your organization can leverage Google technology to align with the NCSC Zero Trust Architecture Principles. This is a technical guide aimed at Enterprise and Security Architects charged with developing and executing a zero trust strategy.

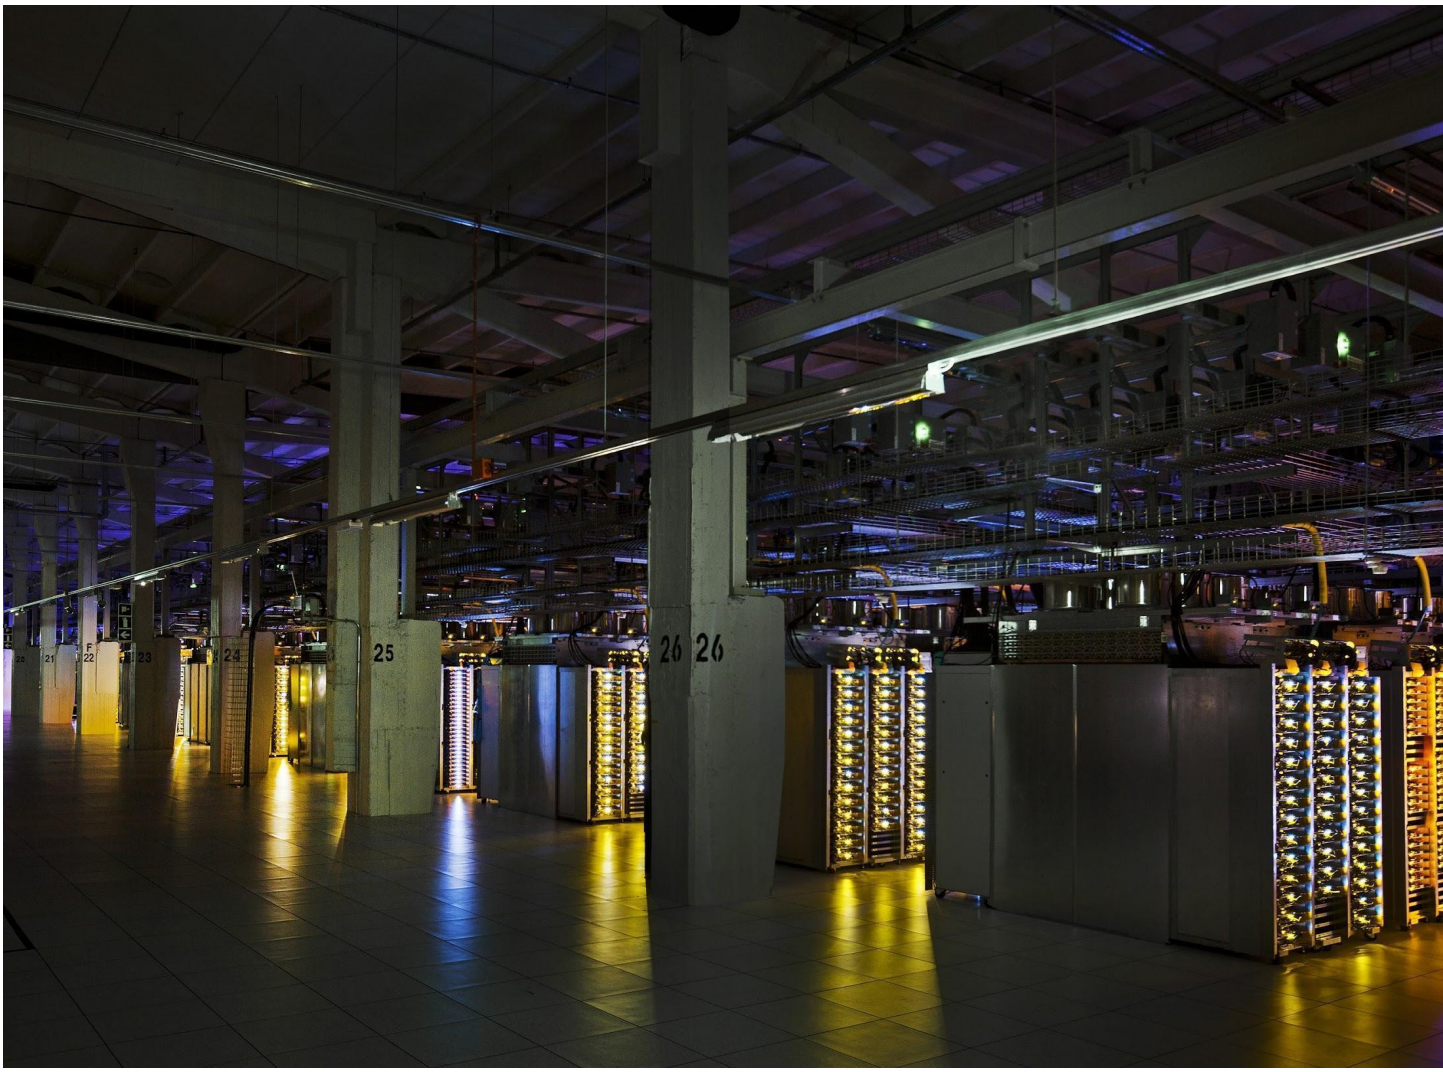

NCSC Zero Trust Architecture Design Principles

For more detailed information from the NCSC see
https://www.ncsc.gov.uk/collection/zero-trust-architecture

1 - Know your architecture, including users, devices, services and data

In order to get the benefits from zero trust, you need to understand each component of your architecture. This will allow you to identify where your key resources are, the main risks to your architecture and also avoid any late-stage pitfalls integrating legacy services that do not support zero trust.

Understanding your current environment is critical to zero trust adoption, just as it applies more broadly to Cloud migrations. Because of this similarity, many of the tools and processes used as part of a Cloud migration can be applied in support of a zero trust migration.

Asset Discovery and Inventory effectiveness

There are different ways to build an inventory. While the quickest way to get started is to proceed manually, this approach can be difficult for a large production environment. Information in manually built inventories can quickly become outdated, and the resulting migration might fail because it was based on incomplete or imperfect visibility into an organization's assets.

Building the inventory is not a one-time exercise. If your current environment is highly dynamic, you should also spend effort in automating the inventory creation and maintenance, so you maintain an up-to-date view of all the items in your environment at any given time.

Google partners with multiple companies to assist you in your migration journey. Refer to the Finding help section of our Cloud Architecture site, for more information.

Data Discovery

Similar techniques can be used to manage discovery and cataloging of data. Automated tools such as Data Catalog can support teams in building this inventory.

Many business leaders also assume that the cloud is inherently secure because there is a general misunderstanding of the shared responsibility model and the difference in preventive security vs detection & response to threats. Of course, they are partially correct - cloud infrastructure is certainly more secure than their data centers. However, their own usage of the cloud often isn't. Analyst firms often remind us that the vast majority of cloud security problems and data breaches occur due to the fault of cloud users and not cloud service providers.

Transition Plan and Risk Assessment

As well as the current environment, you should consider the target architecture. Is the migration to zero trust being undertaken as part of a wider Cloud adoption programme? For example, some services could be migrated to Cloud-hosted equivalents - reducing the risk associated with those services in a zero trust environment.

Google Cloud's Professional Services Organisation (PSO) is able to work directly with customers to help plan a zero trust architecture implementation for first-move systems and workloads.

2 - Know your User, Service and Device identities

An identity can represent a user (a human), service (software process) or device. Each should be uniquely identifiable in a zero trust architecture. This is one of the most important factors in deciding whether someone or something should be given access to data or services.

User Identity

Access to company data is no longer limited to your physical office or your employees. Instead, in today's transformed workforce - increasingly connected, collaborative and in the cloud - the security perimeter has become dispersed and elastic, wrapped around each user and device.

Moreover, 'users' no longer refers to simply employees, but also vendors, partners, contractors and customers. Each of these groups has their own requirements - access to different information and applications, from different locations and different devices.

Google Cloud Identity can help - an identity, access and device management (IAM/EMM) platform that helps organizations maximize user and IT efficiency, protect company data with Google-grade security, and transition to a digital workspace at their own pace. Cloud Identity was designed for zero trust environments from the start.

Architects may wish to review the Reference Architectures provided for Cloud Identity to select a model that best fits their organization.

Requirement	Google Cloud Identity Feature
Create groups	Cloud Identity supports communication and collaboration groups (includes email lists), and configuration groups.
Define roles that have been configured to be 'least privilege'	Permissions are granted to users of Google Cloud Platform via Cloud IAM, and Workspace users via the Administration tools. • For Cloud IAM, Role recommendations help you identify and remove excess permissions from your principals, improving your resources' security configurations. • For Workspace, comprehensive configuration and reporting options allow you to define and monitor the implementation of least privilege.
Support strong, modern authentication methods such as multi-factor or passwordless authentication.	Cloud Identity 2-Step Verification puts an extra barrier between your business and cybercriminals who try to steal usernames and passwords to access business data. **Turning on 2-Step Verification is the single most important action you can take to protect your business**. Cloud Identity supports a range of second factors, including security keys. Because security keys are the strongest 2-Step Verification method, consider using them in your business. • Security keys—The strongest 2-Step Verification method, and they don't require users to enter codes. You can buy compatible security keys from a retailer you trust, such as Titan Security Keys from the Google Store. Or your users can use their phone's built-in security key (available on phones running Android 7+ or iOS 10+). • Alternatives to security keys—If you decide not to use security keys, Google prompt or the Google Authenticator app are good alternatives. Google prompt provides a better user experience because users simply tap their device when prompted instead of entering a verification code. • Text messages are discouraged
Securely provision credentials to users	Cloud Identity can provision credentials in a number of ways, including automatically expiring links, or integrating with existing enterprise services to allow remote onboarding of employees. Administrators can further control how users enroll with the required 2FA factors through policy configuration.

Enable federated authentication to services (e.g. SAML 2.0, OAuth 2.0 or OIDC)	Google provides a large selection of ready-to-use integrations for popular third party applications, and you can use standard protocols such as SAML, OAuth, and OpenID Connect to integrate your custom applications.
Manage user identities in external services, where applicable (e.g. SCIM 2.0)	Using automated user provisioning, you can automatically save any changes to users' identities in the Google Admin console for all supported applications. The setup and configuration of user provisioning vary from app to app. Learn more about the specific setup process by clicking on the app you are interested in. The following list is the complete set of applications preconfigured to support user provisioning: https://support.google.com/cloudidentity/topic/7661972
Support your joiners, movers and leavers processes	Cloud Identity supports a range of integrations for account provisioning (Joiners) and suspension / deletion (Leavers) - including API driven management of user accounts. Movers can also be managed via API or external system, but administrators can also take advantage of Groups, and Dynamic Groups. A dynamic group is a Google Group whose memberships are automatically managed using a membership query or a query on employee attributes, such as job role or building location. For example, a membership query might be "all users whose job role is Technical Writer in my organisation." This approach helps automatically reassign permissions based on multiple user attributes, ensuring that a user does not aggregate permissions from multiple roles over their time with the organisation. **Using an HR Information System or Third Party IdP** Many HR information systems (HRIS), IdPs, and adapters only support one-way user provisioning. This means that changes performed in the HRIS or IdP are propagated to Cloud Identity or Google Workspace, but changes performed in Cloud Identity or Google Workspace are not propagated back. To prevent inconsistencies caused by one-way provisioning, designate your external IdP as the source of truth. Exclusively use your external IdP (or HRIS) to create, modify, or delete users, and rely on automated provisioning to have changes be propagated to Google Workspace and Cloud Identity. By designating your external IdP as the source of truth, you limit the risk of inconsistencies and of having manual modifications overridden by the IdP.

Support third party federated ID (accepting identities from trusted 3rd parties' user directories)	Cloud Identity can be federated with third party identity providers (IdP), such as Active Directory or Azure Active Directory. This enables you to: • Automatically provisioning relevant user accounts from an external authoritative source to Cloud Identity or Google Workspace. • Enabling users to use an external IdP to authenticate to Google services. See the Best Practises for Federating guide for more details.

Migration

Comprehensive Migration support is provided in the Cloud Identity documentation, with the approach contingent on the chosen identity architecture.

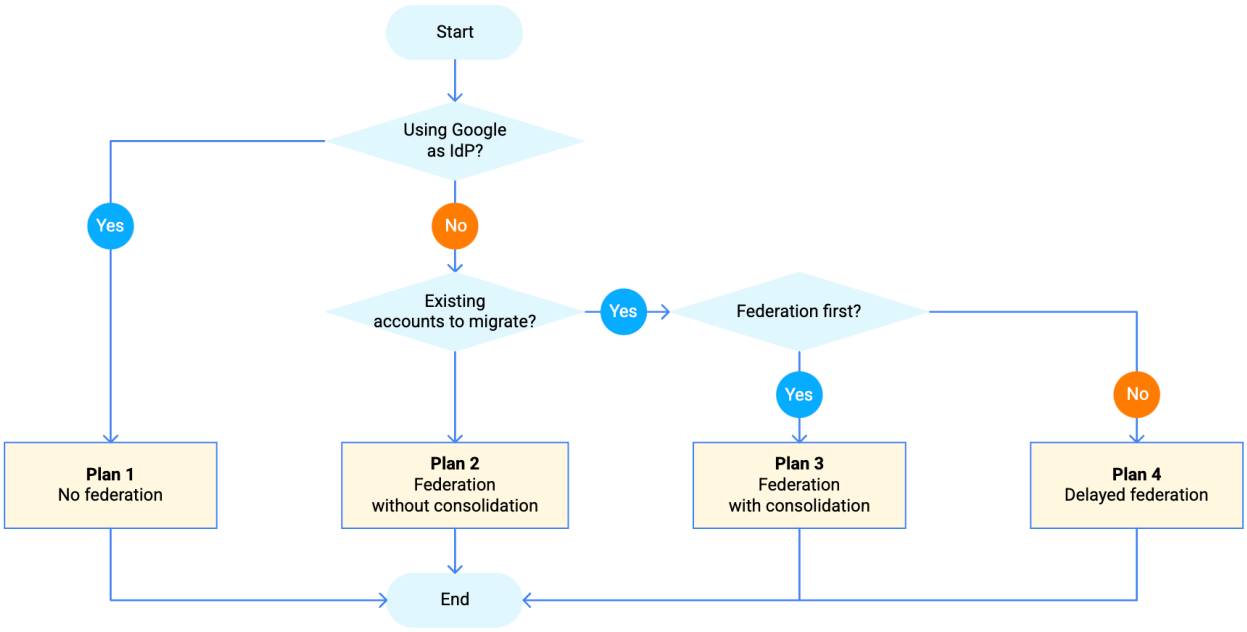

Review the Reference architectures article to select the architecture that most closely matches your requirements.

External Access

In addition to federation, Google Workspace and Cloud Platform allow administrators to restrict which external domains are permitted to access content and projects within the organization.

This approach allows your external users to present their own verified credentials, and allows content and project owners to share directly with those users. Organisational Administrators use Domain Restricted Sharing to set guardrails over which external domains can be permitted by those owners.

Furthermore, Context Aware Access policies can be used to set fine grained controls over what services and data the external users are able to access.

Service Tokens

Google APIs such as the Prediction API and Google Cloud Storage can act on behalf of your application without accessing user information. In these situations your application needs to prove its own identity to the API, but no user consent is necessary. Similarly, in enterprise scenarios, your application can request delegated access to some resources.

For these types of server-to-server interactions you need a service account, which is an account that belongs to your application instead of to an individual end-user. Your application calls Google APIs on behalf of the service account, and user consent is not required. (In non-service-account scenarios, your application calls Google APIs on behalf of end-users, and user consent is sometimes required.)

The OAuth 2.0 tokens issued to the service account are both scoped and time limited.

Service Identity

A service account is a special kind of account used by an application or a virtual machine (VM) instance, not a person. Applications use service accounts to make authorized API calls, authorized as either the service account itself, or as Google Workspace or Cloud Identity users through domain-wide delegation.

For example, a Compute Engine VM can run as a service account, and that account can be given permissions to access the resources it needs. This way the service account is the identity of the service, and the service account's permissions control which resources the service can access.

A service account is identified by its email address, which is unique to the account.

Service accounts differ from user accounts in a few key ways:

- Service accounts do not have passwords, and cannot log in via browsers or cookies.
- Service accounts are associated with private/public RSA key-pairs that are used for authentication to Google.
- You can let other users or service accounts impersonate a service account.
- Service accounts do not belong to your Google Workspace domain, unlike user accounts. If you share Google Workspace assets, like docs or events, with your entire Google Workspace domain, they are not shared with service accounts. Similarly, Google Workspace assets created by a service account are not created in your Google Workspace domain. As a result, your Google Workspace and Cloud Identity admins can't own or manage these assets.

https://cloud.google.com/iam/docs/service-accounts

At the application level, Anthos Service Mesh allows you to configure a layer of service context-aware and request context-aware network security that is independent of the security of the underlying network. Because of this, Anthos Service Mesh lets you adopt a defense-in-depth posture that is consistent with zero trust security principles. It lets you achieve this posture through declarative policies and without modifying any application code.

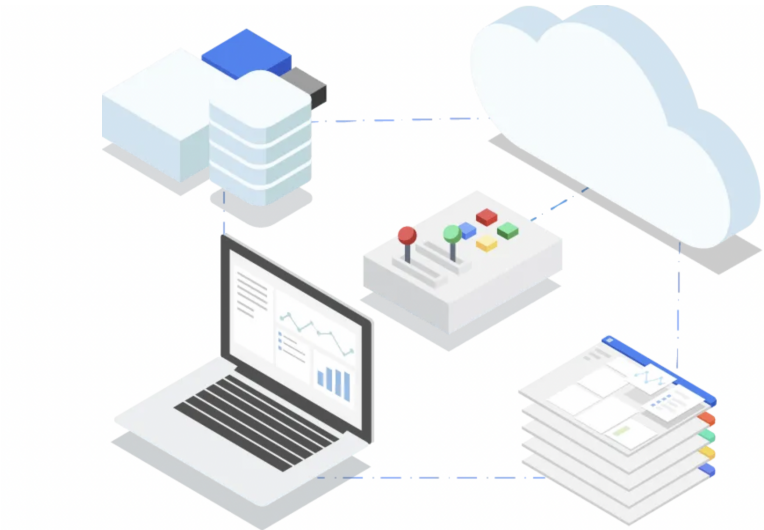

Device Identity

As noted in the NCSC guidance, “the strength of the device’s identity depends on the device type, hardware and platform”. ChromeOS devices are equipped with a Trusted Platform Module (TPM), regardless of price point. The presence of a TPM allows administrators to enable Verified Access.

Verified Access ensures that a device connecting to your network has been unmodified and is policy-compliant. Verified Access serves as an access point for a network service (such as a VPN gateway, a sensitive server, an enterprise Certificate Authority (CA), or an enterprise Wi-Fi access point) to get a hardware-backed cryptographic guarantee of the identity of the device and user that’s trying to access it. Learn more about how Verified Access works.

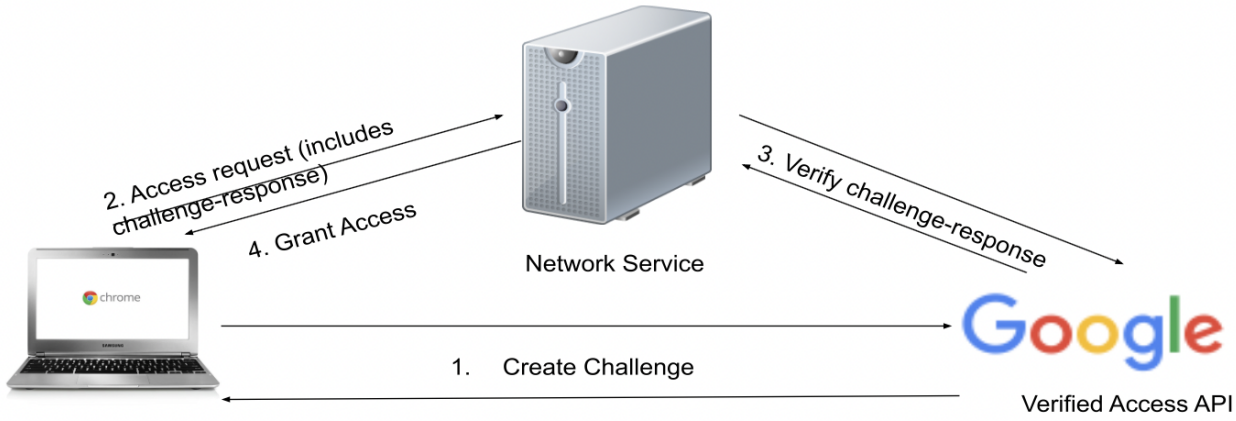

Verified Access uses the Trusted Platform Module (TPM) - present in every Chrome OS device - to enable enterprise network services to cryptographically confirm the identity and status of verified boot and enterprise policy using a Google server-side Application Programming Interface (API).

You need to enable the Verified Access feature in the Google Admin console and force-install a Chrome extension on your users’ devices. Once you’ve done this, your network service talks to the Verified Access API to determine the policy compliance and talks to Google to (optionally) determine the identity of the client device.

3 - Assess your user behaviour, devices and services health

User behaviour, and service or device health, are important indicators when looking to establish confidence in the security of your systems, making them important signals for policy engines. Therefore, having the ability to measure user behaviour, device and service health is key in a zero trust architecture.

The Cloud Identity security center provides advanced security information and analytics, and added visibility and control into security issues affecting your domain.

The security center expands on advanced settings in the Google Admin console to surface your security data through insightful, customizable reports that you can share with colleagues in your organisation. Administrators can also monitor the configuration of Google Admin console settings from the security health page. Additionally, admins can use the investigation tool to identify, triage, and take action on security and privacy issues in your domain.

Devices

The reporting on device health varies depending on the device type. For example, mobile devices running Android or iOS will report whether they have been rooted or jailbroken in the Compromised Device reporting.

Device security health events are presented via dashboards, but can also be used as part of device management rules.

A device management rule is triggered by an event on a managed device. When the event is detected, the rule checks for any conditions you specify. If the conditions are met, an action is carried out.

For example, you can block a device when the account registration state changes on Android devices because a user unregisters their corporate account from the device. In this example:

- The event is an account registration state change on a device.
- The first condition is that the device type is Android.
- The second condition is that a user unregisters their account from the device (Account state is Unregistered from).
- The action is blocking the device.

You can create your own rule or work with a predefined template. For the scope, you can assign a rule to your whole organisation, an organisational unit, or a group in Google Groups. You can also exclude a group.

Note: Device management rules let you approve, block, or wipe a device in response to a specific event. To control access to Google apps for devices based on device attributes such as OS version, security status, IP address, geographic location, or ownership, you can use Context-Aware Access levels. Context-Aware Access is described in Section 4.

Services

Services deployed in Google Cloud can leverage the Security Command Centre (SCC). In addition to the asset discovery and inventory capabilities described in Section 1, SCC offers:

Threat prevention: Understand the security state of your Google Cloud assets. Uncover common web application vulnerabilities such as cross-site scripting or outdated libraries in your web applications running on App Engine, GKE, and Compute Engine. Quickly resolve misconfigurations by clicking directly on the impacted resource and following the proscribed steps on how to fix it.

Threat detection: Detect threats using logs running in Google Cloud at scale. Detect some of the most common container attacks, including suspicious binary, suspicious library, and reverse shell.

The zero trust Infrastructure itself (including Context Aware Access and Identity Aware Proxy) are battle-tested components managed by Google on your behalf - based on BeyondCorp. BeyondCorp is Google's implementation of the zero trust model. It builds upon a decade of experience at Google, combined with ideas and best practices from the community.

Users

Cloud Identity tracks and logs user behaviours, providing administrators with reporting and alerts.

You can use the user login attempts report to identify spikes in the amount of failed and suspicious logins in your domain. You can also view statistics about the challenge methods that have been used. This chart enables you to identify and investigate attempts to hijack user accounts in your organisation.

As an administrator, you can use email alerts to notify you if there's suspicious sign-in activity for your users. For example, Google might notice a sign-in attempt that doesn't match a user's normal behaviour. Usually, before Google sends you an alert, Google presents the user with an extra security question or challenge. If the user fails or abandons the challenge, the alert is sent.

Infrastructure

Cloud Identity tracks and logs user behaviours, providing administrators with reporting and alerts.

Depending on the architectural choices made under the Shared Responsibility Model[1] (IaaS, PaaS, SaaS, Hybrid on-prem), Google offers a range of telemetry options (for example, VPC Flow Logs), which can be combined with existing Cloud Logging sources to provide a common view of security events.

If required, Google Professional Services can assist with migration planning of existing infrastructure and services, taking advantage of Cloud native and managed services to reduce the burden on customer teams.

[1] In addition to the video content from Google Cloud Next, the Shared Responsibility Matrix for PCI-DSS is available via our Security and Compliance pages: https://cloud.google.com/security/compliance/pci-dss

Deeper analysis of telemetry can be conducted in Chronicle. Chronicle enables you to examine the aggregated security information for your enterprise going back for months or longer. Use Chronicle to search across all of the domains accessed within your enterprise. You can narrow your search to any specific asset, domain, or IP address to determine if any compromise has taken place.

Chronicle is a cloud service, built as a specialized layer on top of core Google infrastructure, designed for enterprises to privately retain, analyze, and search the massive amounts of security and network telemetry they generate. Chronicle normalizes, indexes, correlates, and analyzes the data to provide instant analysis and context on risky activity.

Chronicle can ingest numerous security telemetry types through a variety of methods, including:

- Forwarder: A lightweight software component, deployed in on-premise networks, that supports syslog, packet capture, and existing log management or security information and event management (SIEM) data repositories.
- Ingestion APIs: APIs that enable logs to be sent directly to the Chronicle platform, eliminating the need for additional hardware or software in customer environments.
- Third party integrations: Integration with third party cloud APIs to facilitate ingestion of logs, including sources like Office 365 and Azure AD.

4 - Use policies to authorize requests

Each request for data or services should be authorized against a policy. The power of a zero trust architecture comes from the access policies you define. Policies can also help to facilitate risk managed sharing of data or services with guest users or partner organisations.

The policy engine is a key component of the zero trust architecture, it uses multiple signals and provides a flexible and secure access control mechanism that adapts to the resources being requested.

The BeyondCorp Enterprise solution provides the components required to make policy based decisions.

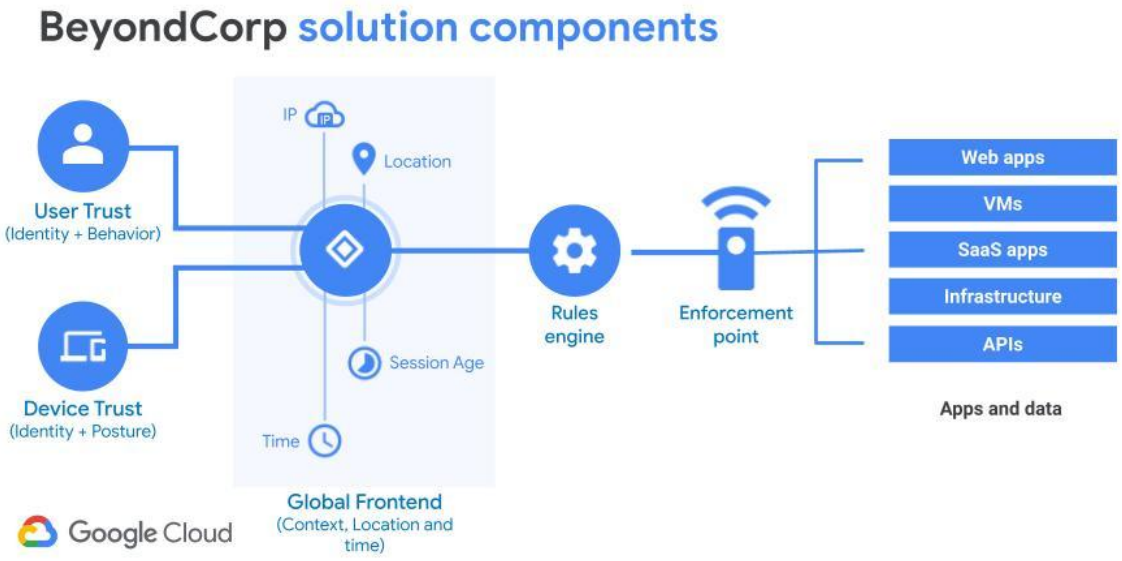

The architecture maps to the NCSC principles as follows:

01

The Policy Enforcement Point is provided by Identity Aware Proxy, Identity and Access Management, Cloud Identity or VPC Service Controls, depending on the request type.

02

The Enforcement Point queries the Rules Engine - provided by Access Context Manager on Google Cloud

03

Requests that do not match the required policy are dropped by the Enforcement Point

04

Each request in a given session is evaluated by the Rules engine, allowing for real time continuous evaluation. For example, if an element of context changes, such as geolocation, the request will be dropped or re-authenticated.

Continuous Evaluation

Protect the policy engine

Under the shared responsibility model, customers are required to define appropriate access policies, but are **not** responsible for the security of Access Context Manager itself. Cloud IAM permissions to configure and monitor the various components are defined by the customer, and so should be configured according to standard principles of least privilege.

Use multiple signals to make access decisions

Multiple signals are fed into the access decision, including User and Device posture, as well as signals from the Global Frontend - such as IP address, geolocation, session age and time of day. Credential Strength (e.g. hardware second factor) can also be used as a signal.

Risk-based engines

Access levels are used for permitting access to resources based on contextual information about the request. Using access levels, you can start to organize tiers of trust. For example, you might create an access level called High_Level that will permit requests from a small group of highly-privileged individuals. You might also identify a more general group to trust, such as an IP range that you want to permit requests from. In that case, you might create an access level called Medium_Level to permit those requests.

Once you have defined access levels, enforcement services can use them to determine whether to honor a request. For example, you might specify that while many resources are available to "Medium_Trust," certain more sensitive resources require the "High_Trust" level. These checks are applied in addition to standard Identity and Access Management policy.

https://cloud.google.com/access-context-manager/docs/overview#access-levels

Other considerations

Follow these best practices to help ensure a smooth rollout of Context-Aware Access policies in your company. These best practices are based on customer feedback.

Avoid locking out employees, partners, or external collaborators

Don't block access to Google Workspace services, such as Gmail, that you use to share communications with your users (and that they also need to communicate with you).

Identify IP ranges that partners, external collaborators, and clients need.

Keep in mind that some Google Workspace services, such as Forms and Sites, don't have a mobile app and will be blocked on phones.

Roll out device policies in phases

Discover — Enforce the use of Endpoint verification so you know which devices are accessing (or will be accessing) Google Workspace data. Find out information about each device, such as if it's encrypted, running an up-to-date operating system, and if it's a company-owned or personal device.

Note that if you enforce a Context-Aware device policy before the user can sign in to Endpoint verification, the user may get access denied even if their device meets the enforced Context-Aware policy. This is because syncing the device attributes through Endpoint verification may take a few seconds. To avoid this, be sure to have users sign into Endpoint verification before you enforce a Context-Aware device policy.

Remediate — Get your devices under IT management and in compliance with company standards in preparation for device policy enforcement. This should help reduce help desk tickets and support calls.

Enforce — Enforce policies to restrict access to apps based on device context. Identify the organisations, sub-organisations, and groups, and then apply device policies in a phased rollout. Base your rollout plan on the device composition of each organisation or group, and plan for sufficient help desk support.

See this support article for more details on implementing Context-Aware Access
https://support.google.com/a/answer/9275380

5 - Authenticate & Authorise everywhere

Authentication and authorisation decisions should consider multiple signals, such as device location, device health, user identity and status to evaluate the risk associated with the access request. We do this as we assume the network is hostile and want to ensure all connections that access your data or services are authenticated and authorized.

Multi-factor

You can make 2-Step Verification (Multi-factor authentication) optional or required for your users. We recommend enforcing 2-Step Verification for your administrator account and users who work with your most important business information.

- Forwarder: A lightweight software component, deployed in on-premise networks, that supports syslog, packet capture, and existing log management or security information and event management (SIEM) data repositories.
- Ingestion APIs: APIs that enable logs to be sent directly to the Chronicle platform, eliminating the need for additional hardware or software in customer environments.

While Google supports a wide variety of second factor authentication methods, you should consider using Security keys, as they provide the strongest protection.

- Security keys—The strongest 2-Step Verification method, and they don't require users to enter codes. You can buy compatible security keys from a retailer you trust, such as Titan Security Keys from the Google Store. Or your users can use their phone's built-in security key (available on phones running Android 7+ or iOS 10+).
- Alternatives to security keys—If you decide not to use security keys, Google prompt or the Google Authenticator app are good alternatives. Google prompt provides a better user experience because users simply tap their device when prompted instead of entering a verification code.
- Text messages are discouraged—They rely on external carrier networks and might be intercepted.

Usability

Rolling out 2-Step verification should be carefully planned to avoid locking out legitimate users of your service. See the deployment planning support article for further details.

To balance usability and security, Cloud Identity will look for suspicious activity on accounts and prompt for 2-Step verification if it sees a suspicious event. We determine whether a sign-in is suspicious when our risk-analysis system identifies an attempt that's outside the normal pattern of user behaviour. For example, a user might try to sign in from an unusual location or in a manner associated with abuse.

If you are not using 2-Step verification, a range of alternative login challenges can be used - for example, prompting the user to enter their Employee ID.

Administrators can also use Context-Aware Access policies to tailor the security requirements to best balance usability and security - for example, applying extra constraints if a user is accessing services from a particular geography.

Service to service

In 2019, Google published a whitepaper on our BeyondProd model to explain how we protect our cloud-native architecture and to help organisations learn to apply the security principles that we established internally.

We developed and optimized for the following security principles:

- **Protection of the network at the edge**, so that workloads are isolated from network attacks and unauthorized traffic from the Internet.
- **No inherent mutual trust between services**, so that only known, trusted, and specifically authorized callers can utilize a service. This stops attackers from using untrusted code to access a service. If a service does get compromised, it prevents the attacker from performing actions that allow them to expand their reach. This mutual distrust helps to limit the blast radius of a compromise.
- **Trusted machines** — designed with Titan to be secure from boot on up — running code with known provenance, so that service identities are constrained to use only authorized code and configurations, and run only in authorized, verified environments.
- **Choke points for consistent policy enforcement across services**. For example, a choke point to verify requests for access to user data, such that a service's access is derived from a validated request from an authorized end user, and an administrator's access requires business justification.
- **Simple, automated, and standardized change rollout**, so that infrastructure changes can be easily reviewed for their impact on security, and security patches can be rolled out with little impact on production.
- **Isolation between workloads sharing an operating system**, so that if a service is compromised, it can't affect the security of another workload running on the same host. This limits the "blast radius" of a potential compromise.

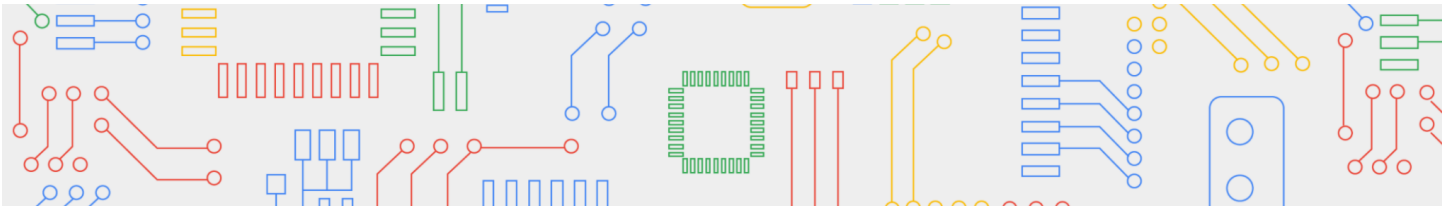

Many of the capabilities of BeyondProd are embedded in Anthos, Google Cloud's managed application platform, in features like Binary Authorization and Anthos Service Mesh.

Anthos

Binary Authorization

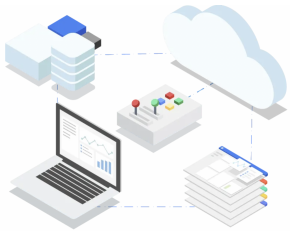

Anthos Service Mesh

By applying the security principles in the BeyondProd model to your own cloud-native infrastructure, you can benefit from our experience to strengthen the deployment of your workloads, including how your communications are secured and how they affect other workloads.

If you are looking to apply the principles of BeyondProd in your own environment, there are many components through Anthos, Google Kubernetes Engine (GKE) and open source that you can leverage to achieve a similar architecture:

- Envoy or Traffic Director, for managing TLS termination and policies for incoming traffic;
- Anthos Service Mesh for a zero-trust security model toolset to automatically and declaratively secure services and their communication;
- Mutual TLS, as part of ASM or Istio for RPC authentication, integrity, encryption, and service identities;
- Anthos Identity Services to support identity federation across environments;
- Binary Authorization for deploy-time enforcement checks such as code provenance;
- Anthos Config Management Policy Controller, to enforce programmable policies for clusters and prevent configuration changes from violating security, operational, or compliance controls;
- Shielded GKE Nodes, for secure boot and integrity verification; and
- gVisor or GKE Sandbox, for workload isolation.

6 - Focus your monitoring on users, devices and services

In a zero trust architecture, it is highly likely that your monitoring strategy will change to focus on users, devices and services. Monitoring of these devices, services and users behaviours will help you establish their health. Monitoring should link back to the policies you have set to gain assurance in their configuration.

Protective Monitoring

Cloud native monitoring solutions provide a richer set of protective monitoring capabilities than traditional network boundary logging - e.g. at a VPN chokepoint. Comprehensive protective monitoring in a zero trust environment will likely involve a range of teams - from those who are supporting users and devices through to service and product owners.

Google collates protective monitoring and investigative tools in two primary locations:

Cloud Identity — Security Center, which primarily relates to device and user configurations and behaviour.

Google Cloud Platform — Security Command Center, allows you to manage the security posture of services deployed in GCP and beyond.

For those operating in hybrid or multi-cloud environments. Cloud Logging can be extended to cover both on-premise and other cloud vendors, ensuring that security teams have a consistent view across their environment.

BYOD and Guest devices

For devices the Organisation does not own and control - BYOD and Guest devices - administrators can configure options that balance the level of control the Organisation has over a personal device with the security requirements of that organisation.

A work profile can be set up on an Android device to separate work apps and data from personal apps and data. With a work profile you can securely and privately use the same device for work and personal purposes—your organisation manages your work apps and data while your personal apps, data, and usage remain private. Mobile application management options exist for iOS devices, via the Device Policy App.

More generally, Context-Aware access can be used to define multiple Access, or trust, levels. For example, you could define policies that permit access to certain services only to company-owned devices, while Guest and BYOD devices could access lower sensitivity applications.

Network Monitoring

Google Cloud cybersecurity product management and engineering teams have observed the maturation of the industry from access control to the addition of intrusion prevention, and more recently, analytics-based detection and automated response. As such, we provide a range of tools across different network threat signal types - summarized in the table below.

Tool	Contents	Method	Use
IPFIX (Netflow)	Simple flow descriptions, optionally sampled	Streamed from networking devices across their hybrid, multi-cloud environment into a datalake, e.g. Chronicle or Elastic	Mostly NOC, some SOC, for connectivity patterns, help determine access control. All connections that pass the network instance (N-S & E-W).
VPC Flow Logs	Detailed flow metadata, sampled, not all flows recorded	Enabled on Google Cloud workloads	NOC & SOC (such as for DDoS detection), connectivity patterns, help determine access control. VPC-Internet, VPC-VPC.
Packet Mirroring	Full packets copied	Enable on Google Cloud VPCs, networks, tags, workloads, filtered traffic	NOC & SOC, all analytics, detections, threat hunting, etc. All traffic, N-S & E-W.
Cloud IDS	Detected threat logs	Enable on Google Cloud VPCs, networks, tags, workloads, filtered traffic; built with Packet Mirroring and Palo Alto Networks industry-leading threat detection technology & research	SOC, full detections, pre-written. All traffic, N-S & E-W, including intra-VPC and intra-container-pod.
Network Forensics & Telemetry Blueprint	Full packets and full flow metadata (uses Packet Mirroring)	DIY assembly of Packet Mirroring, Zeek, Pub/Sub, datalake and analytics tool(s), e.g. Chronicle	NOC & SOC, full analytics, predictive alerts, custom detections, threat hunting, detection mechanism development and tuning.

For more detail, see this blog post on Network security threat detection - Comparison of analytics methods

7 - Don't trust any network, including your own

Don't trust any network between the device and the service it's accessing, including the local network. Communications over a network, to access data or services, should use a secure transport protocol to gain assurance that your traffic is protected in transit and less susceptible to threats.

A zero trust architecture changes the way traditional user protections such as malicious website filtering and phishing protection are implemented, these may need to be provided by different solutions in your zero trust architecture.

Services deployed on Google Cloud take advantage of the same secure-by-design infrastructure, built-in protection, and global network that Google uses to protect your information, identities, applications, and devices. Our stack builds security through progressive layers that deliver true defense in depth at scale.

We encrypt data in transit between our facilities and at rest, ensuring that it can only be accessed by authorized roles and services with audited access to the encryption keys. Learn more about how we encrypt data at rest and how we encrypt data in transit.

Enforcing device usage policy

Safe Browsing launched in 2007 to protect users across the web from phishing attacks, and has evolved to give users tools to help protect themselves from web-based threats like malware, unwanted software, and social engineering across desktop and mobile platforms.

Our Transparency Report includes details on the threats that Safe Browsing identifies. The Transparency Report includes our Site Status diagnostic tool that you can use to see whether a site currently contains content that Safe Browsing has determined to be dangerous.

Safe Browsing can be enforced via Chrome Browser policies in Cloud Identity and Workspace.

Additional protections against hostile local networks can be enforced in the same way - including configuring HSTS preloading, DNS over HTTPS, restricting user ability to bypass SSL errors, and DNS interception checks.

8 - Choose services designed for zero trust

Services may not support zero trust and thus may require additional resources to integrate and increase support overhead. In these scenarios it may be prudent to consider alternative products and services that have been designed with zero trust in mind.

Using products that utilize standards-based technologies allows for easier integration and interoperability between services and identity providers.

BeyondCorp is Google's implementation of the zero trust model. It builds upon a decade of experience at Google, combined with ideas and best practices from the community. By shifting access controls from the network perimeter to individual users, BeyondCorp enables secure work from virtually any location without the need for a traditional VPN.

BeyondCorp began as an internal Google initiative to enable every employee to work from untrusted networks without the use of a VPN. Now, BeyondCorp is used by most Googlers every day to provide user- and device-based authentication and authorization for Google's core infrastructure and corporate resources.

https://cloud.google.com/beyondcorp

While BeyondCorp focuses on the security of users and devices, securing our Cloud native services is equally important - implemented at Google as BeyondProd, and detailed in our whitepaper.

It is important to note that the majority of business application services (for example, a web based staff timesheet) will have not been built explicitly as “designed for zero trust”. For those applications, architects should focus on the standards that those applications support - ensuring they can be successfully integrated with the chosen zero trust architecture.

Legacy services

BeyondCorp Enterprise customers can secure HTTP or HTTPS based on-premises applications (outside of Google Cloud) with Identity-Aware Proxy (IAP) by deploying a connector. When a request is made for an on-premises app, IAP authenticates and authorizes the user request and then routes the request to the connector. For more details, see the step-by-step guidance on the Identity-Aware Proxy documentation page.

Look for standards

Google Cloud undertakes independent verification of our security, privacy, and compliance controls to help you meet your regulatory and policy objectives. Find details on our full set of compliance offerings, like ISO/IEC 27001/27017/27018/27701, SOC 1/2/3, PCI DSS, and FedRAMP certifications, and alignment with GDPR, Cyber Essentials and NCSC Cloud Security Principles, among others, in our compliance resource centre.

Managed services in the cloud

As NCSC note in their guidance: *"There are a number of cloud hosted services that have been designed for zero trust. It's important that you're confident you can trust the vendors running these services."*

Protecting the privacy of Google Cloud Platform and Google Workspace customers is a priority and codified in our Enterprise Privacy Commitments, which guide our security and privacy practices. Learn more about how we protect privacy and keep you in control in our privacy resource centre.

We work to earn your trust through transparency. We state and adhere to a concrete set of trust principles that govern our approach to security. We only process data in accordance with our terms and data protection agreements and clearly outline our policies on responding to government requests.

Trust Principles

Security Approach

Data Process

Policies

Path to alignment with NCSC Zero Trust Architecture Design Principles

While this paper highlights the technologies and services that can form the foundation of a zero trust architecture, aligning these to a given organization can be a complex task. This is especially true in heterogeneous environments, employing technologies from multiple vendors.

Open Cloud

Google's Open Cloud approach embraces its partner ecosystem rather than competing against it. With that in mind, we offer managed open source services operated by our partners that are tightly integrated into Google Cloud, providing a seamless user experience across management, billing, and support. This makes it easier for our enterprise customers to build on open source technologies.

On Premises

Applications and services do not need to be migrated to the Cloud to take advantage of Google zero trust solutions. BeyondCorp Enterprise supports a hybrid model, providing zero trust controls and protection for on-premises resources.

Professional Services

For bespoke and in-depth guidance, Google Cloud Professional Services Organisation (PSO) have a range of custom offerings that can assist organizations on their journey to Zero Trust, including:

- **Zero Trust Foundations** —This zero trust PSO Engagement will assist customers with developing an executable zero trust strategy. Engagement will consist of a deep-dive assessment, comprehensive zero trust workshop, and a customised zero trust strategy document. Follow-on PSO engagements can be used to execute zero trust gaps identified during the Foundations engagement.
- **Cloud Deploy: Zero Trust** — If customers require additional support to implement zero trust elements for enterprise workloads, they can leverage follow-on PSO engagements for implementation.

Combined, these PSO offerings:

01	02	03	04
Help customers identify gaps in their zero trust posture	Help customers create a zero trust plan that aligns with NCSC zero trust Architecture Principles	Help customers adopt an executable zero trust strategy for their organisation, based on their zero trust plan	Assist customers' zero trust migration and implementation with PSO support

Conclusion

While this paper highlights the technologies and services that can form the foundation of a zero trust architecture, aligning these to a given organization can be a complex task. This is especially true in heterogeneous environments, employing technologies from multiple vendors.

Defining and delivering against a zero trust strategy has never been more important for organisations - who are adapting to working patterns radically altered by the pandemic, as well as increased threats from advanced cyber threats.

This paper has highlighted how Google Cloud solutions can be used to implement such a strategy, aligned with the architecture principles set out by the NCSC.

Organisations managing complex environments may require more support; which is available from Google Professional Services as a series of defined deliverables, helping you build both strategy and a plan for implementation.

If you would like to discuss further, please contact your Google Cloud account team, or reach us via uki-pubsec-cloud@google.com.