Google Cloud

Professional Cloud Architect

Preparing for Professional Cloud Architect Journey for AWS Professionals

Session 2 topics

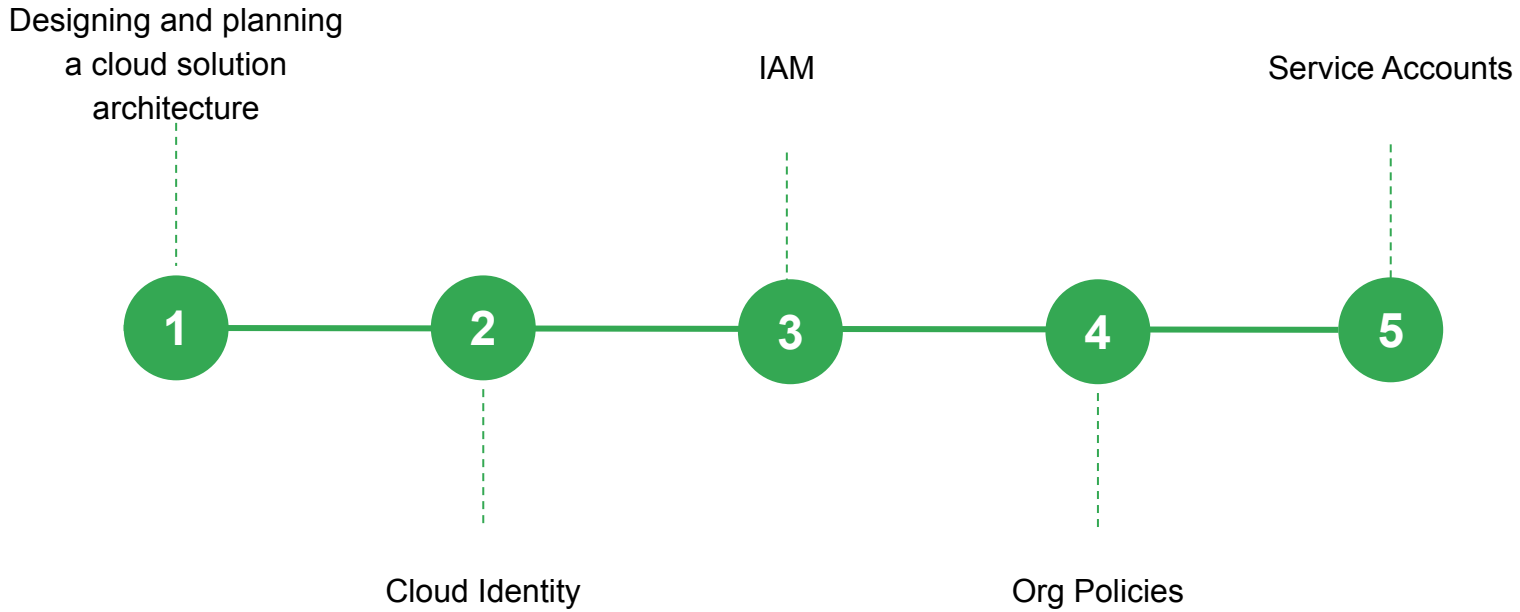

Designing and planning a cloud solution architecture

Define systems in scope for a cloud migration

... and / or decide on “cloud first” approach

Delivery by Drone

- Their website frontend, pilot, and truck management systems run on Kubernetes.
- Positional data for drone and truck location is kept in a MongoDB database clusters
- Drones stream video to virtual machines via stateful connection

Purchase & Product APIs

- APIs are simply built into monolithic apps, and were not designed for partner integration.
- APIs are running on Ubuntu linux VMs

Social Media Highlighting

- Single SuSE linux VM
- MySQL DB
- Redis
- Python

Define business and technical requirements

Business requirements

- Easily scale to handle additional demand when needed and expand to more test markets.
- Streamline development for application modernization and new features/products
- Ensure that developers spend as much time on core business functionality as possible, and not have to worry about scalability wherever possible
- Let partners order directly via API
- Deploy a production version of the social media highlighting service and ensure no inappropriate content

Technical requirements

- Move to serverless services wherever possible
- Ensure that developers can deploy container-based workloads to testing and production environments in a highly scalable environment.
- Standardize on containers where possible, but also allow for existing virtualization infrastructure to run as-is without a re-write, so it can be slowly refactored over time
- Securely allow partner integration
- Stream IoT data from drones

Mapping requirements to GCP resources

… through the lens of business & technical requirements

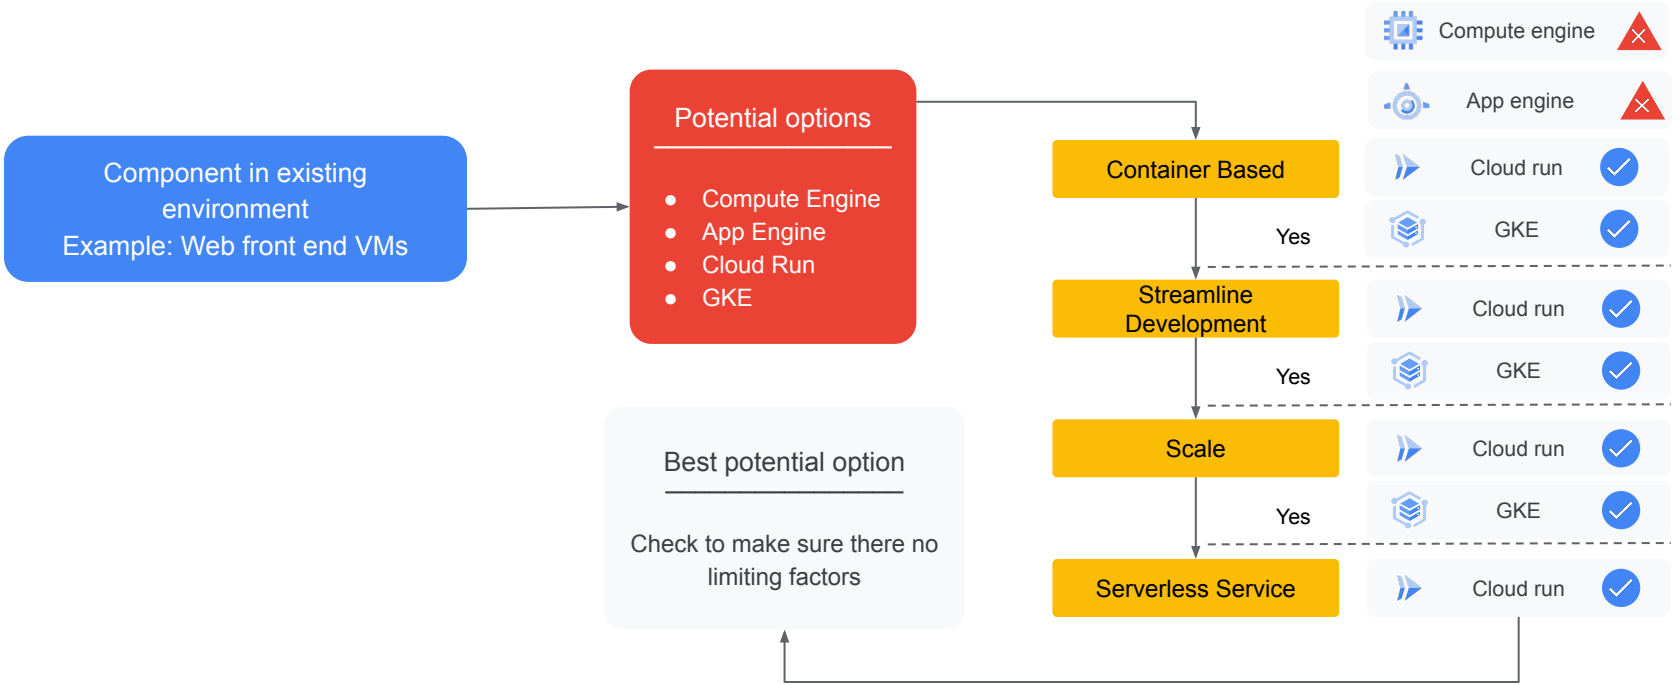

Planning for migration and the future

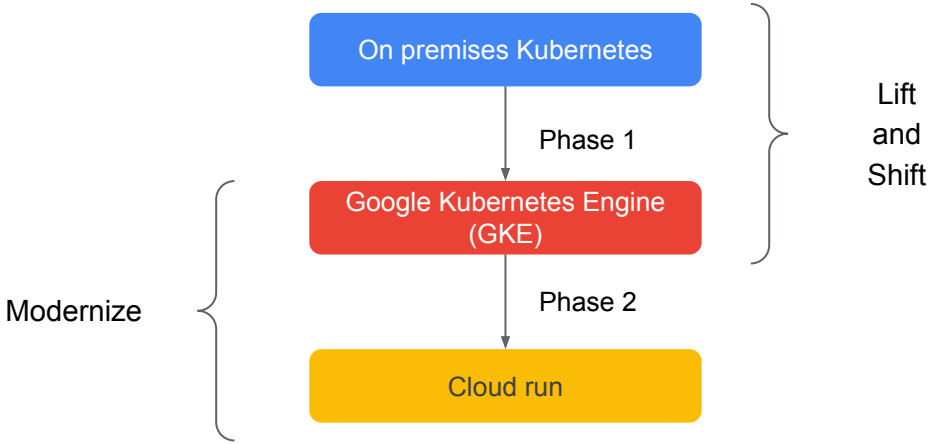

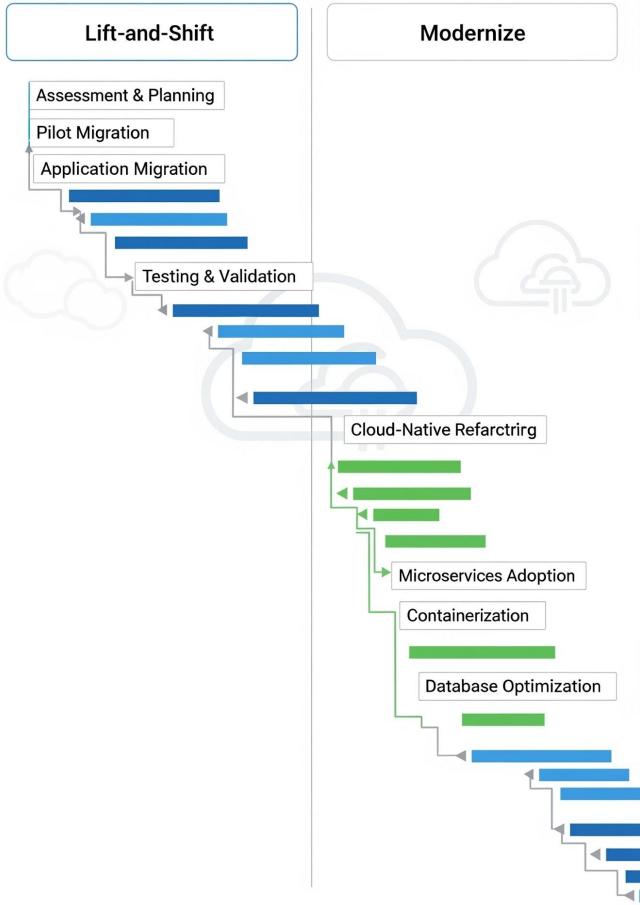

Migration guide & best practices

- Types of migrations and their use-cases
  - For example, “If the current app isn't meeting your goals—for example, you don't want to maintain it, it's too costly to migrate using one of the previously mentioned approaches, or it's not supported on Google Cloud—you can do a rebuild migration.”
- Building inventory of workloads in scope
  - … along with their dependencies!
- Best practices for validating a migration plan.

Diagnostic Question Discussion

You work for a large financial institution that is planning a multi-phase migration of its on-premises workloads to Google Cloud. The migration involves sensitive customer data and requires high availability and disaster recovery capabilities. You want to design a cloud solution architecture that meets these requirements while minimizing costs and ensuring compliance with industry regulations.

What should you do?

A. Migrate all workloads to a single region in Google Cloud to simplify management and reduce latency. Implement basic security measures, such as firewalls and intrusion detection systems, to protect against common threats. Use a combination of managed services and self-managed services to balance cost and flexibility.

B. Design a multi-zone architecture within a single region to reduce costs while maintaining high availability. Implement security measures based on industry best practices, such as using strong passwords and multi-factor authentication. Use primarily self-managed services to have greater control over the environment.

C. Design a multi-region architecture with active-passive failover for high availability and disaster recovery. Implement appropriate security measures, such as data encryption at rest and in transit, to protect sensitive customer data. Use managed services whenever possible to reduce operational overhead and costs.

D. Defer considerations for high availability and disaster recovery until a later phase to minimize initial costs and complexity. Implement security measures as needed based on the sensitivity of the data. Use primarily open-source tools and technologies to minimize licensing costs.

Diagnostic Question Discussion

You work for a large financial institution that is planning a multi-phase migration of its on-premises workloads to Google Cloud. The migration involves sensitive customer data and requires high availability and disaster recovery capabilities. You want to design a cloud solution architecture that meets these requirements while minimizing costs and ensuring compliance with industry regulations.

What should you do?

A. Migrate all workloads to a single region in Google Cloud to simplify management and reduce latency. Implement basic security measures, such as firewalls and intrusion detection systems, to protect against common threats. Use a combination of managed services and self-managed services to balance cost and flexibility.

B. Design a multi-zone architecture within a single region to reduce costs while maintaining high availability. Implement security measures based on industry best practices, such as using strong passwords and multi-factor authentication. Use primarily self-managed services to have greater control over the environment.

C. Design a multi-region architecture with active-passive failover for high availability and disaster recovery. Implement appropriate security measures, such as data encryption at rest and in transit, to protect sensitive customer data. Use managed services whenever possible to reduce operational overhead and costs.

D. Defer considerations for high availability and disaster recovery until a later phase to minimize initial costs and complexity. Implement security measures as needed based on the sensitivity of the data. Use primarily open-source tools and technologies to minimize licensing costs.

Diagnostic Question Discussion

You are a Professional Cloud Architect working with a large retail customer that has a monolithic e-commerce application hosted on-premises. They are experiencing challenges with scalability and performance, especially during peak shopping seasons. They want to migrate this application to Google Cloud and modernize it to be more resilient, scalable, and cost-effective.

Your goal is to design a solution that meets their requirements.

What should you do? (choose two)

A. Migrate the application to Google Cloud using a phased approach, starting with a lift-and-shift migration and gradually modernizing components.

B. Deploy the entire application to a single, large Compute Engine instance to ensure resource availability and minimize management overhead.

C. Continue running the application on-premises and use Cloud CDN and Cloud Load Balancing to enhance performance and scalability.

D. Decompose the monolithic application into microservices and leverage managed services like Google Kubernetes Engine (GKE) and Cloud SQL.

E. Refactor the entire application to serverless architecture using Cloud Functions and Cloud Run to minimize operational overhead and maximize cost savings.

Diagnostic Question Discussion

You are a Professional Cloud Architect working with a large retail customer that has a monolithic e-commerce application hosted on-premises. They are experiencing challenges with scalability and performance, especially during peak shopping seasons. They want to migrate this application to Google Cloud and modernize it to be more resilient, scalable, and cost-effective.

Your goal is to design a solution that meets their requirements.

What should you do? (choose two)

A. **Migrate the application to Google Cloud using a phased approach, starting with a lift-and-shift migration and gradually modernizing components.**

B. Deploy the entire application to a single, large Compute Engine instance to ensure resource availability and minimize management overhead.

C. Continue running the application on-premises and use Cloud CDN and Cloud Load Balancing to enhance performance and scalability.

D. **Decompose the monolithic application into microservices and leverage managed services like Google Kubernetes Engine (GKE) and Cloud SQL.**

E. Refactor the entire application to serverless architecture using Cloud Functions and Cloud Run to minimize operational overhead and maximize cost savings.

Cloud Identity

Managing Cloud Identity: AWS vs Google Cloud

AWS

- Provides Active Directory (AD) integration using the AWS Directory Service. This enables integration with existing AD solutions.
- Existing LDAP solutions can be integrated with AWS IAM using an integration such as AWS Directory Service instead of managing identity separately.

Google Cloud

- User identities are managed outside of Google Cloud.
- The Cloud Identity tool lets organizations define policies and manage their users and groups using the Google Admin console.
- You cannot use IAM to create or manage users or groups - use Cloud Identity or Google Workspace.
- Google Cloud supports AD/LDAP integration via GCDS

How is an Organization Created?

- **Cloud Identity** manages the users and groups that have access to Google Cloud
  - Federated identities from Google Workspace and other identity providers, such as Active Directory and Azure Active Directory
    - Bring existing users/groups into Cloud Identity
    - Use Identity and Access Management (IAM) to manage access to Google Cloud resources

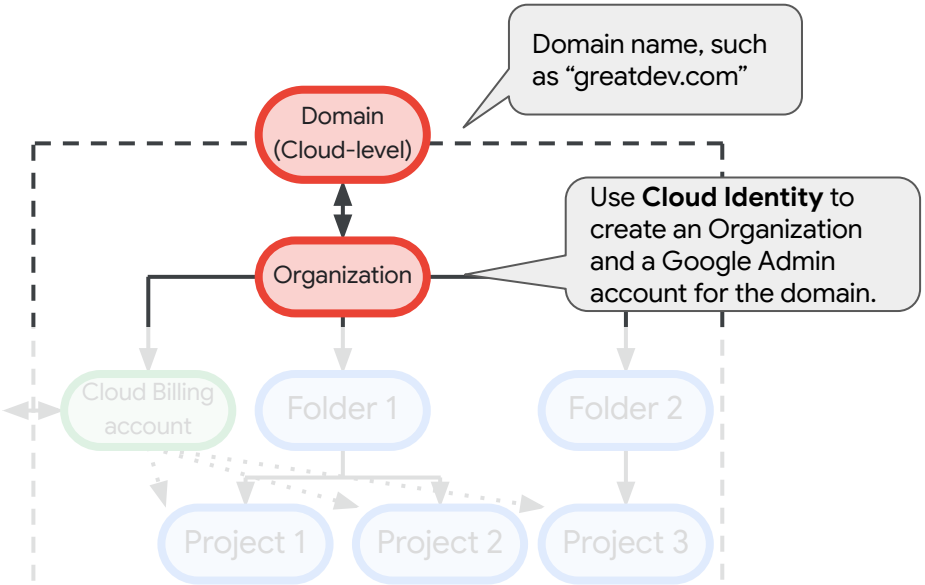

What is Cloud Identity?

- Cloud Identity is an Identity as a Service (IDaaS) solution that **allows you to centrally manage users and groups** who can access GCP and Google Workspace cloud resources
- It is the same identity service that powers Google Workspace and can also be used as IdP for 3rd party applications (supports SAML and LDAP applications)

Members in GCP do not come from GCP!

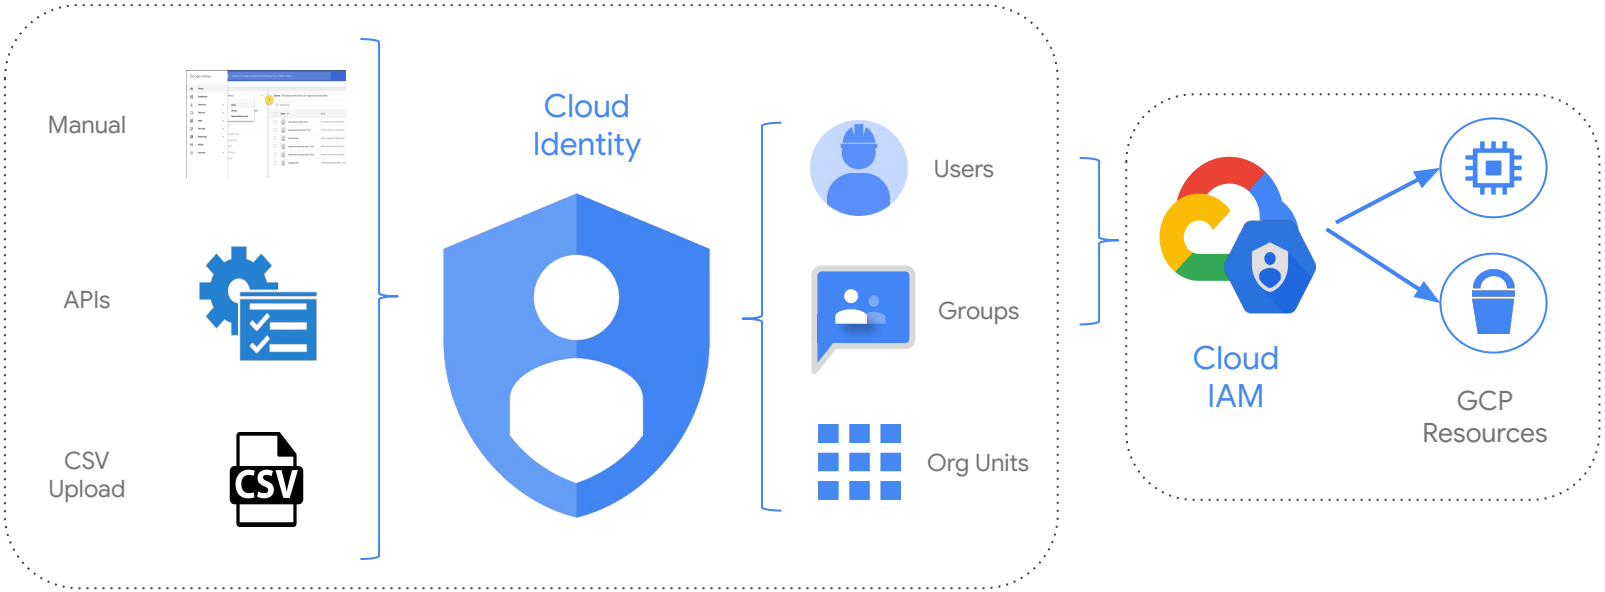

Users and groups created in Cloud Identity are the **Google Identities** that *can* be assigned **IAM roles** in the GCP console

The **Cloud Identity roles** only manage aspects of Cloud Identity such as user/group management, **and are different from GCP roles** which manage permissions to cloud resources

Members in GCP do not come from GCP!

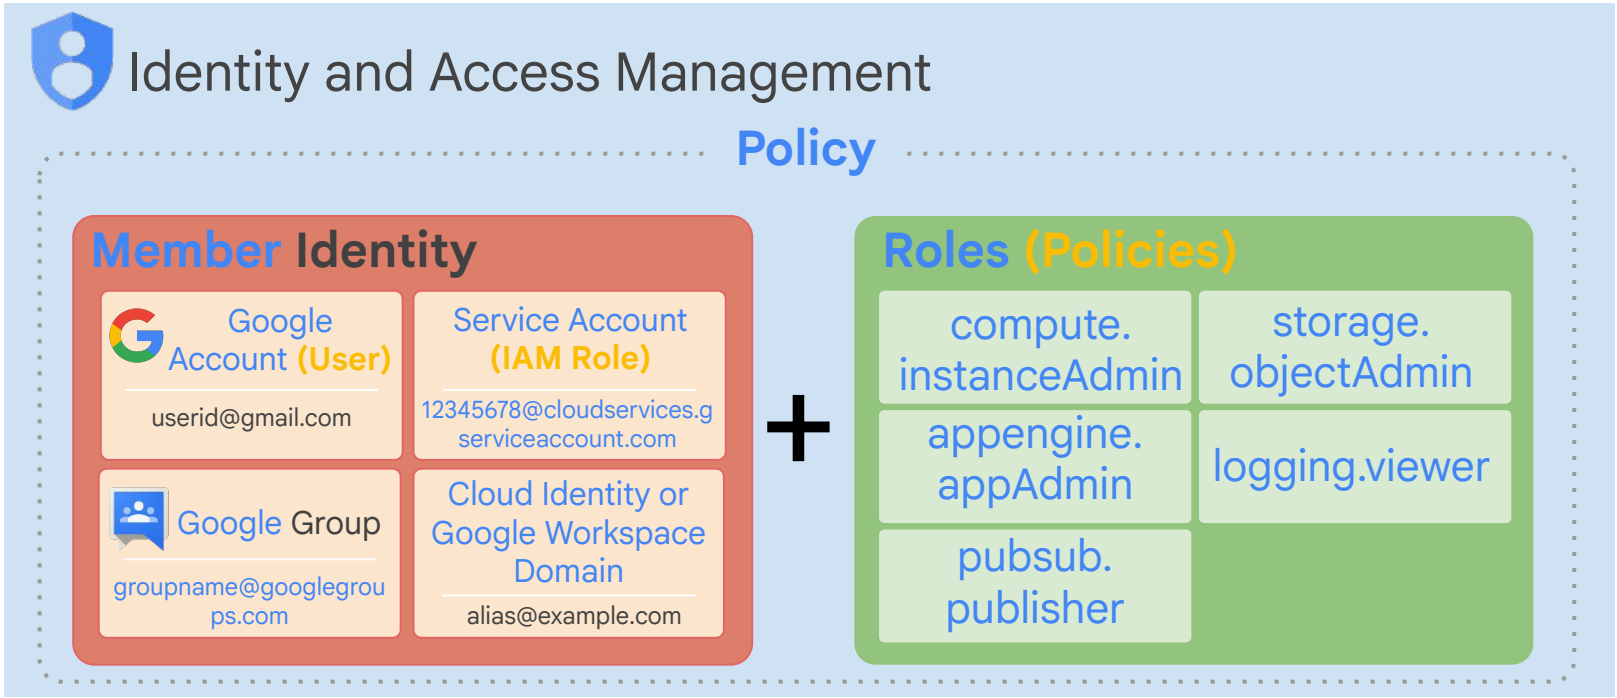

Note: You *cannot* use IAM to create or manage your users or groups.
You can use IAM or IAM Identity Center to manage users and groups.

Cloud Identity: central user and group management

Google's Identity as a Service (IDaaS) solution

- Users and groups that are to be added to Google Cloud need accounts in Cloud Identity

Manually creating user accounts

- Add users individually using the Google Admin console
- Add several users at once by uploading their names in a CSV file

Options for large organizations

- Use Google Cloud Directory Sync to synchronize user data in your existing LDAP directory with your Google account
- Use the Admin SDK Directory API to provision a large number of users with data from your existing LDAP directory, such as Microsoft® Active Directory®
  - Requires programming

*Cloud Identity has advanced management features not covered in this module, e.g., mobile app management, 2-Step verification, etc.

Two consoles for administration

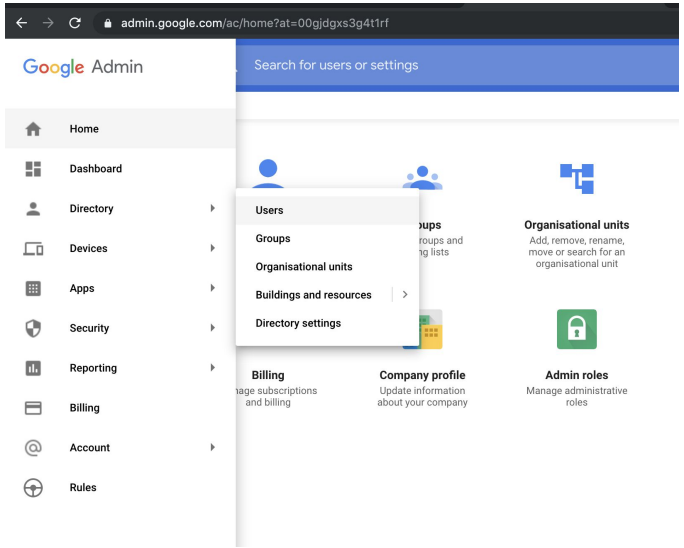

Cloud Identity (`admin.google.com`)

Managing Users, Groups, and Authentication settings

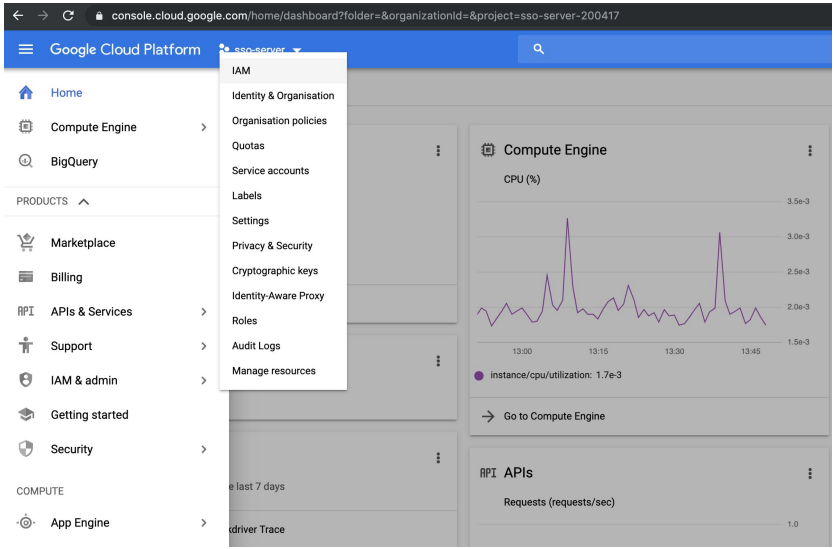

GCP (`console.cloud.google.com`)

Roles & Authorization for GCP

Two key admin functions

	(CI) Super Admin	(Cloud IAM) Org. Admin
Role	**GCP Org. Admin by default**	It can add/assume any other IAM roles
Manages	**User/group account lifecycle** and Org's security settings	**IAM policies** and Resource Manager hierarchy
Delegates	**GCP Org. Admin role** and CI admin roles	**GCP IAM roles** to users and groups
Managed in	**Admin console**	**GCP console**
Visibility	**Cloud Identity and GCP environments**	**GCP environment**

(Cloud Identity) Super Admin ≠ **(Cloud IAM) Organization Admin**

Cloud Identity as an identity provider: Typical architecture

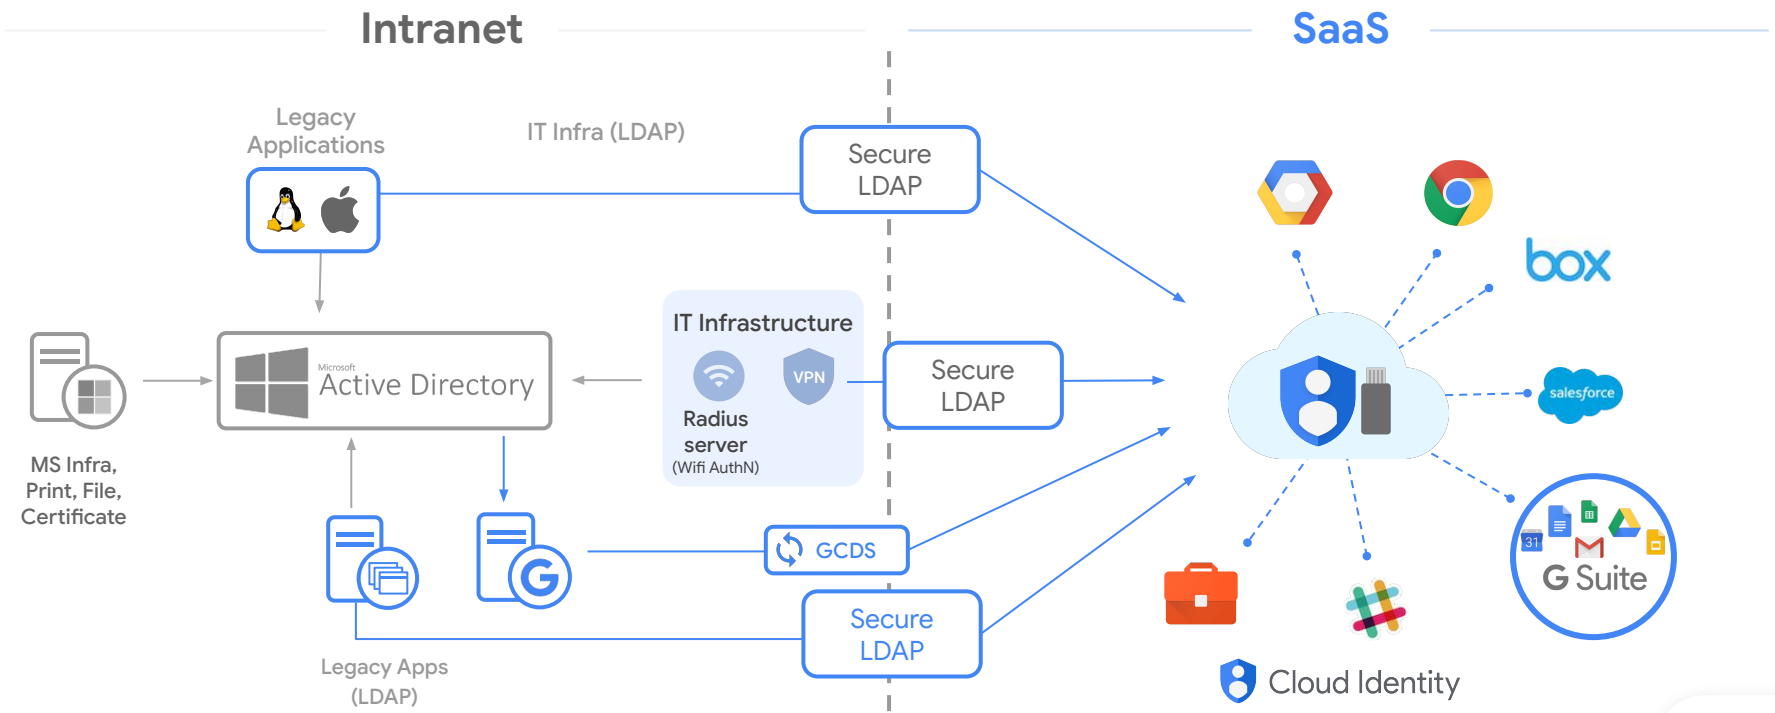

3rd party as an identity provider: Typical architecture

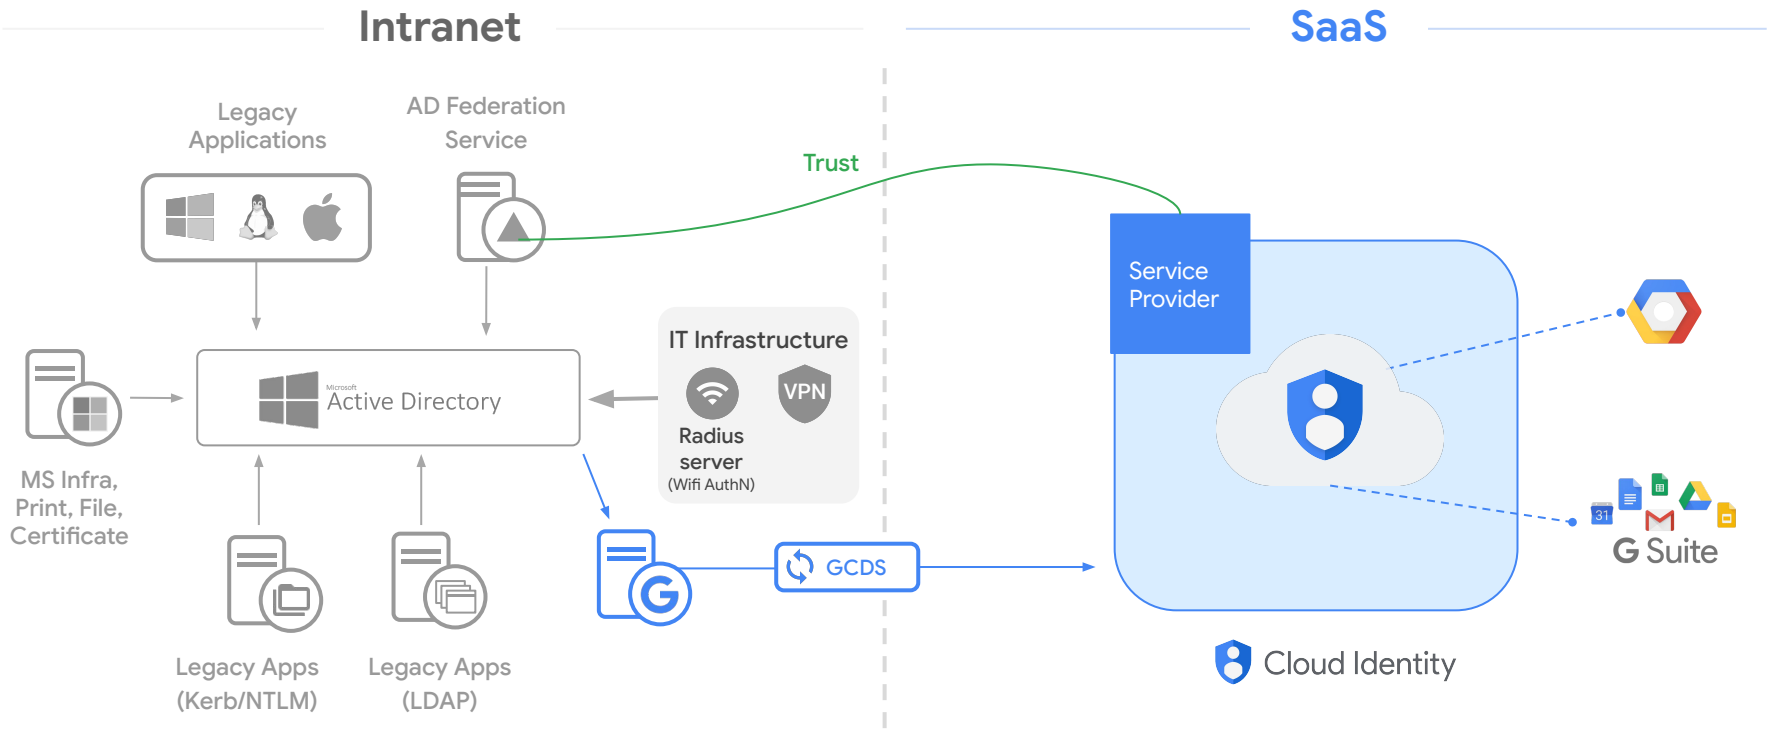

User provisioning options

Method	Effort	Staff involved	Notes
Manual provisioning	High	Google Workspace admin	Easiest method, but not scalable
CSV upload via Admin Console	Medium	Google Workspace admin	More flexibility, but not scalable
Google Cloud Directory Sync	**Medium**	**LDAP Admin**	**Integrates with LDAP, scalable, requires no programming**
3rd party tools (Okta, Ping, Azure AD, ...)	Medium	LDAP admin	Scalable, may incur additional cost
Admin SDK Directory API	High	LDAP Admin Development staff	Scalable, flexible, requires in-depth programming

Best Practice

Diagnostic Question Discussion

Your company wants to start using Google Cloud resources but wants to retain their on-premises Active Directory domain controller for identity management.

What should you do?

A. Use the Admin Directory API to authenticate against the Active Directory domain controller.

B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.

C. Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.

D. Use Compute Engine to create an Active Directory (AD) domain controller that is a replica of the on-premises AD domain controller using Google Cloud Directory Sync.

Diagnostic Question Discussion

Your company wants to start using Google Cloud resources but wants to retain their on-premises Active Directory domain controller for identity management.

What should you do?

A. Use the Admin Directory API to authenticate against the Active Directory domain controller.

B. Use Google Cloud Directory Sync to synchronize Active Directory usernames with cloud identities and configure SAML SSO.

C. Use Cloud Identity-Aware Proxy configured to use the on-premises Active Directory domain controller as an identity provider.

D. Use Compute Engine to create an Active Directory (AD) domain controller that is a replica of the on-premises AD domain controller using Google Cloud Directory Sync.

https://cloud.google.com/architecture/identity/federating-gcp-with-active-directory-introduction

IAM

IAM: GCP vs AWS

Feature	Google Cloud IAM	AWS IAM
Permission Model	**Role-Based.** Bind a *Member* to a ***Role*** **on a Resource**.	**Policy-Based.** Attach a ***Policy*** **to a Principal** (User/Role).
Hierarchy	**Full Inheritance.** Permissions flow down from Org > Folder > Project.	**No Inheritance.** Permissions are scoped to an Account. (OUs provide guardrails, not grants).
"Role" means...	A **set of permissions** (like compute.viewer).	An **identity** that can be assumed (like EC2-S3-Access-Role).
"Policy" means...	The *entire collection of bindings* on one resource.	A JSON document listing Allow/Deny statements.
Day-to-Day Task	"I will grant user@a.com the Storage Viewer role **on Project B**."	"I will attach the S3ReadOnly policy **to the 'dev-group'**."

Comparing IAM: AWS vs GCP

AWS

- Offers built-in and custom **policies**
- The collection of permissions and resources those permissions apply is a policy.
- The account root can fully administer roles and policies unless otherwise restricted by a Service Control Policy.
- Policies can be attached to users, groups of users, or roles.

Google Cloud

- Offers built-in and custom **roles**
- It has sets of predefined roles, and with definitions of where those roles can be applied.
- It grants granular access to specific Google Cloud resources and prevents unwanted access to other resources.
- Roles are collections of permissions.

Identity and Access Management

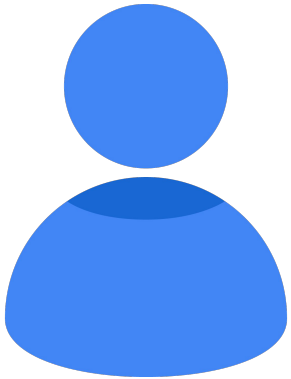

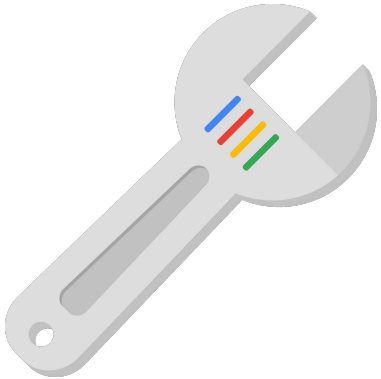

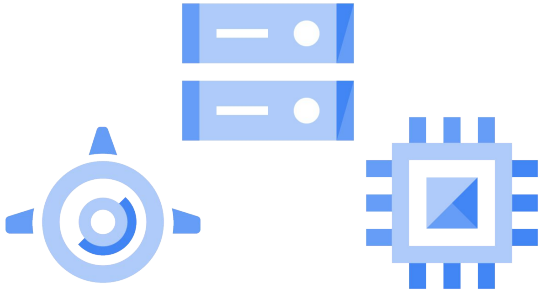

Who

can do what

on which resource

IAM inheritance

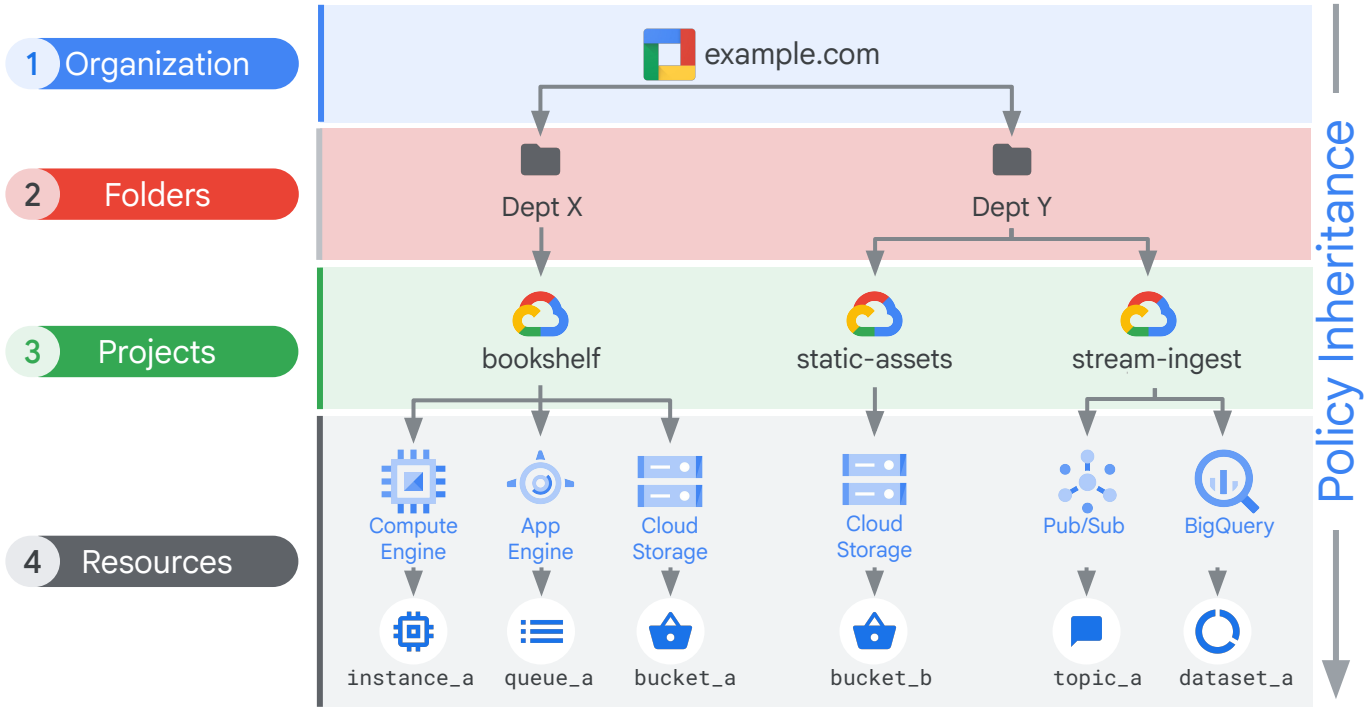

Assign permissions to multiple accounts with IAM Identity Center and Permission Sets.

AWS does not support policy inheritance, but you can use Service Control Policies (SCPs) to limit granted permissions and AWS supports SCP inheritance within an organization.

IAM objects

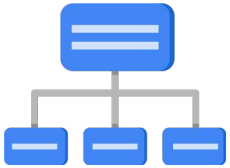

Organization

AWS Organizations

Folders

OUs in Organizations

Projects

Accounts

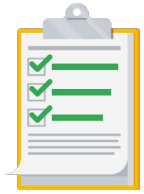

Resources

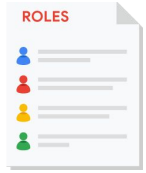

Roles

Policies (intra-account) Policy Sets (across accounts)

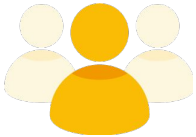

Members

Users, Groups, Roles

IAM policies

- A policy consists of a list of bindings.
- A binding binds a list of members to a role.

AWS does not have an equivalent, but you can use Service or Resource Summary views to see what permissions have been assigned

Three types of IAM roles (IAM policies)

Basic

Simply do NOT use those!

Predefined (AWS managed policies)

Custom (customer managed policies)

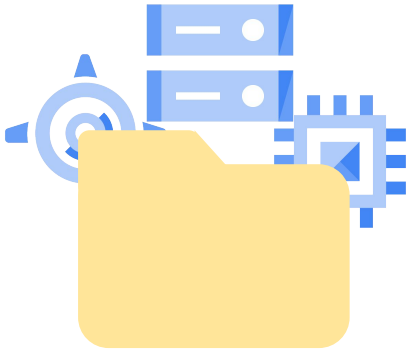

IAM basic roles apply across all Google Cloud services in a project

can do what

on all resources

IAM basic roles offer fixed, coarse-grained levels of access for Google Cloud

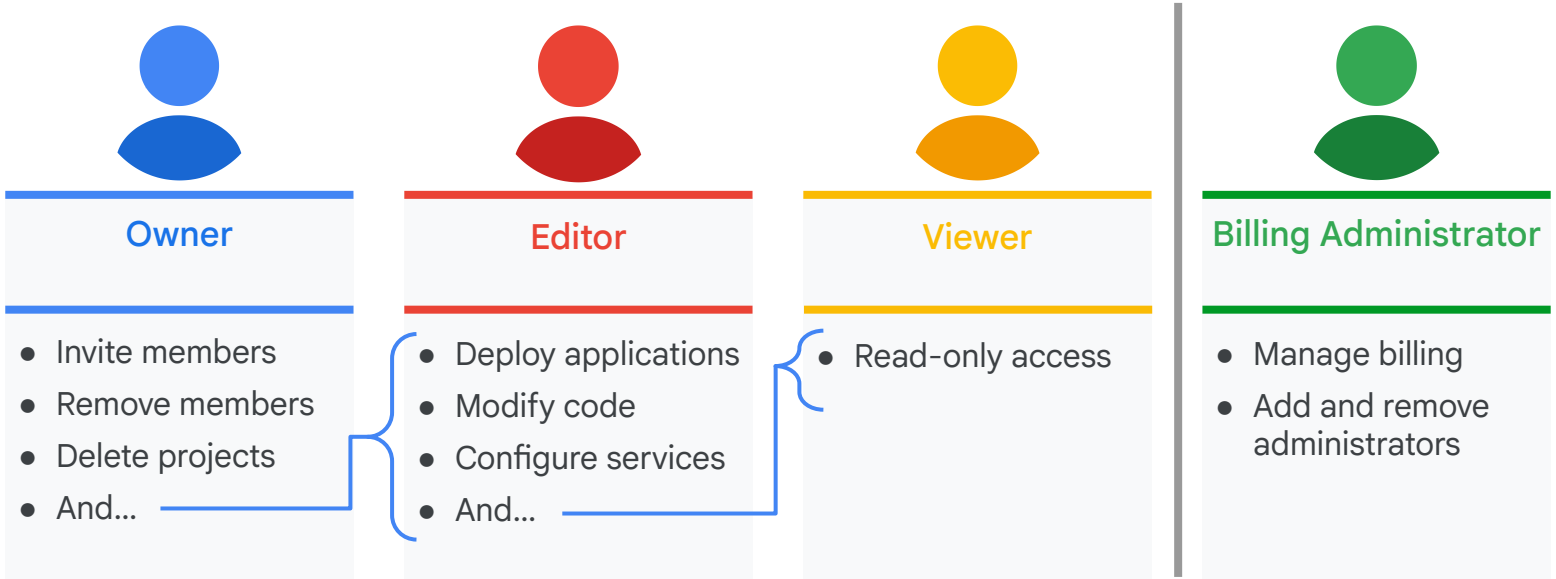

IAM predefined roles apply to a particular Google Cloud service in a project

Google IAM predefined roles = AWS managed policies

can do what

on specified resources in this project (account), folder (OU), or organization

IAM predefined roles offer more fine-grained permissions on particular services

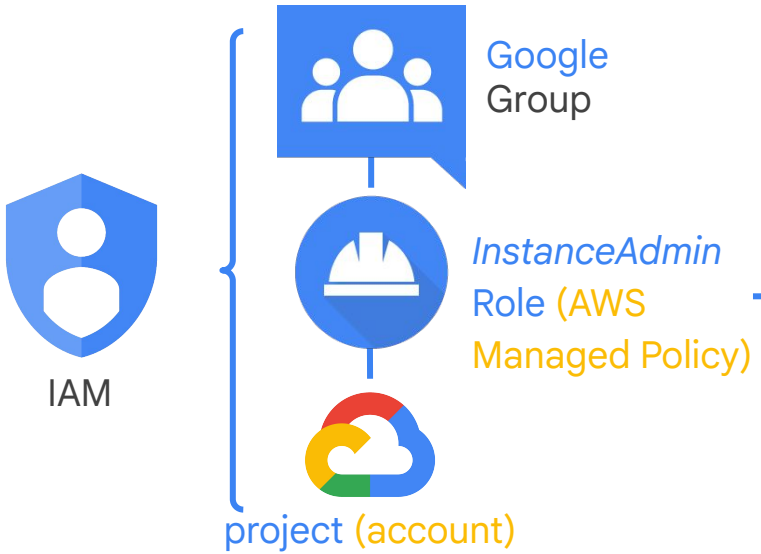

List of Permissions

- ✔ `compute.instances.delete`
- ✔ `compute.instances.get`
- ✔ `compute.instances.list`
- ✔ `compute.instances.setMachineType`
- ✔ `compute.instances.start`
- ✔ `compute.instances.stop`

...

IAM custom roles let you define a precise set of permissions

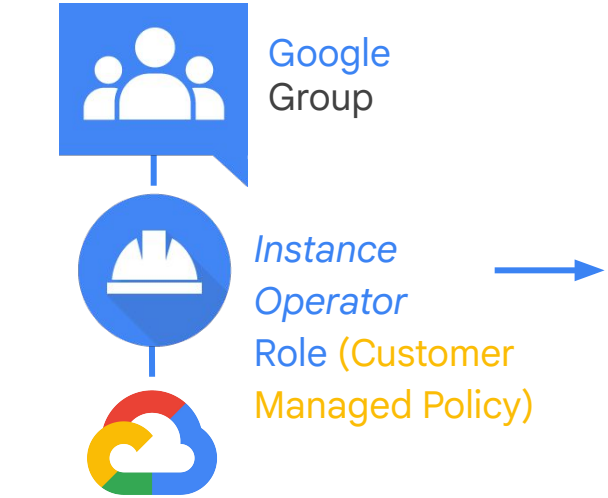

✔ `compute.instances.get`
✔ `compute.instances.list`
✔ `compute.instances.start`
✔ `compute.instances.stop`

IAM conditions

- IAM Conditions are used to enforce conditional ABAC for cloud resources
- With IAM Conditions, you can choose to grant resource access to identities (members) only if configured conditions are met
- Conditions are specified in the role bindings of a resource's IAM policy. When a condition exists, the access request is only granted if the condition expression evaluates to true.

Deny IAM Policies

- Inherited through the resource hierarchy just like IAM allow policies
- Attached to project, folder or organization
- Denies override grants further down the hierarchy
- Currently, must be created via command line

First create a deny policy and store it in a file

People in the dev@example.com group are not allowed to create or delete service account keys

```
{
  "deniedPrincipals": [
    "principalSet://goog/group/dev@example.com"
  ],
  "deniedPermissions": [
    "iam.googleapis.com/serviceAccountKeys.create",
    "Iam.googleapis.com/serviceAccountKeys.delete"
  ]
}
```

Next apply the deny policy

```
gcloud iam policies create POLICY_ID \
    --attachment-point=[proj-id|folder-id|org-id] \
    --kind=denypolicies \
    --policy-file=POLICY_FILE
```

IAM policy pattern example

inheritance

`roles/browser` → `domain:example.org`
all domain users should be able to see the hierarchy

`roles/viewer` → `group:first-lvl-support@`
support users should be able to view logs and VMs

`roles/compute.admin` → `group:linux-os@`
instance admins should manage resources

`roles/logging.logViewer` → `group:app-team-1@`
app admins should view logs in dev

`roles/storage.admin` → `serviceAccount:app1@`
ad-hoc permissions on project or single resource

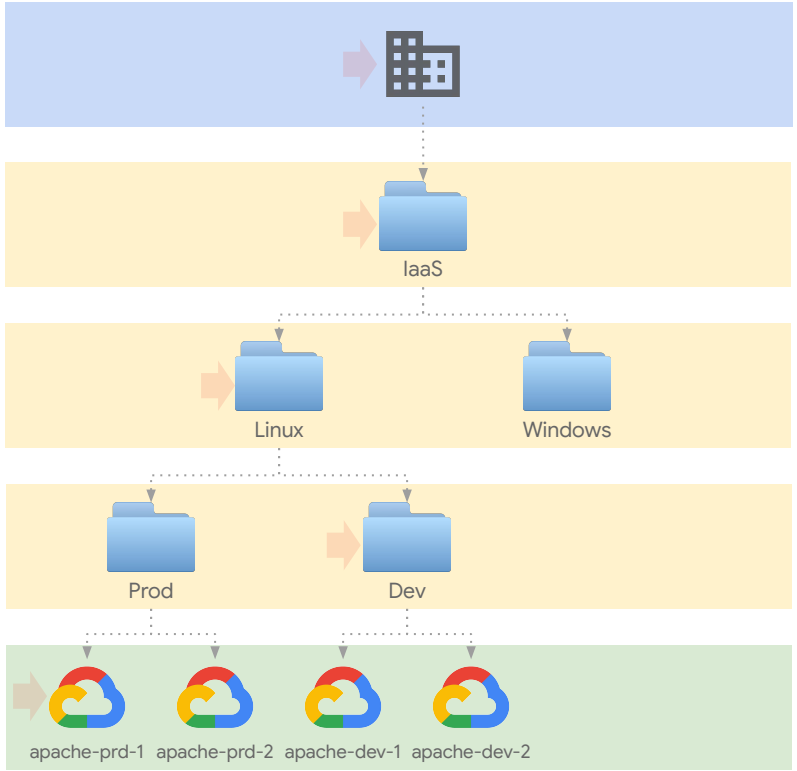

Diagnostic Question Discussion

You are responsible for the Google Cloud environment in your company. Multiple departments need access to their own projects, and the members within each department will have the same project responsibilities. You want to structure your environment for minimal maintenance and follow Google-recommended practices.

What should you do?

A. Create a Google Group per department and add all department members to their respective groups. Create a folder per department and grant the respective group the required IAM permissions at the folder level. Add the projects under the respective folders.

B. Grant all department members the required IAM permissions for their respective projects.

C. Create the folder per department and grant the respective members of the department the required IAM permissions at the folder level. Structure all projects for each department under the respective folders.

D. Create a Google Group per department and add all department members to their respective groups. Grant each group the required IAM permissions for their respective projects.

Diagnostic Question Discussion

You are responsible for the Google Cloud environment in your company. Multiple departments need access to their own projects, and the members within each department will have the same project responsibilities. You want to structure your environment for minimal maintenance and follow Google-recommended practices.

What should you do?

A. **Create a Google Group per department and add all department members to their respective groups. Create a folder per department and grant the respective group the required IAM permissions at the folder level. Add the projects under the respective folders.**

B. Grant all department members the required IAM permissions for their respective projects.

C. Create the folder per department and grant the respective members of the department the required IAM permissions at the folder level. Structure all projects for each department under the respective folders.

D. Create a Google Group per department and add all department members to their respective groups. Grant each group the required IAM permissions for their respective projects.

Diagnostic Question Discussion

You have a user in your environment that has read-only access to four Cloud Storage buckets through the use of IAM. An Organization Administrator mistakenly adds that user to the Project Admin group. The Project Admin group has the Project Editor role associated with it for the project where the four Cloud Storage buckets reside

What are the consequences of adding the user to the Project Admin group?

A. The user will now have full access to every item in the bucket except for the four buckets, which the user will have read-only access to.

B. The user will now have full access to every item in the project.

C. Groups have no association with privileges, therefore, nothing will change..

D. The user will have read-only access to the entire project.

Diagnostic Question Discussion

You have a user in your environment that has read-only access to four Cloud Storage buckets through the use of IAM. An Organization Administrator mistakenly adds that user to the Project Admin group. The Project Admin group has the Project Editor role associated with it for the project where the four Cloud Storage buckets reside

What are the consequences of adding the user to the Project Admin group?

A. The user will now have full access to every item in the bucket except for the four buckets, which the user will have read-only access to.

B. The user will now have full access to every item in the project.

C. Groups have no association with privileges, therefore, nothing will change..

D. The user will have read-only access to the entire project.

Organization Policy

Organization policies (Service Control Policies)
(≠ IAM policies)

An organization policy is:

- A configuration of restrictions
- Defined by configuring a constraint with desired restrictions.
- Applied to the organization node, folders or projects.

In AWS Organizations, Service Control Policies (SCPs) can do some of the same things, but are mostly restrictions on granted permissions that apply to everyone

Most common Organization Policies

Policy Constraint	Description
compute.vmExternalIpAccess	A list of project/zone/instance names that are allowed to have external IP addresses and deny all others. Attempts to create any other VMs with an external IP address will fail.
compute.trustedImageProjects	A list of projects that contain trusted images that can be used as the basis for a VM and deny all others. Attempting to instantiate a VM with an image from another project is denied.
compute.skipDefaultNetworkCreation	Disables the creation of default VPC when creating a project. The default VPC uses auto mode subnetworks and includes default firewall rules which are often incompatible with production deployments.
iam.disableServiceAccountKeyCreation	This boolean constraint disables the creation of service account external keys where this constraint is set to `True`.
gcp.resourceLocations	This list constraint defines the set of locations where location-based GCP resources can be created. Policies for this constraint can specify multi-regions such as asia and europe, regions such as us-east1 or europe-west1, or individual zones such as europe-west1-b as allowed or denied locations.
sql.restrictPublicIp	This boolean constraint restricts configuring Public IP on Cloud SQL instances where this constraint is set to True. This constraint is not retroactive, Cloud SQL instances with existing Public IP access will still work even after this constraint is enforced. By default, Public IP access is allowed to Cloud SQL instances.
sql.disableDefaultEncryptionCreation	Restrict default Google-managed encryption on Cloud SQL instances
compute.requireShieldedVm	This boolean constraint, when set to True, requires that all new Compute Engine VM instances use Shielded disk images with Secure Boot, vTPM, and Integrity Monitoring options enabled. Secure Boot can be disabled after creation, if desired. Shielded VM features add verifiable integrity and exfiltration resistance to your VMs.

Organization policy pattern example

`constraints/compute.vmExternalIpAccess` → `false`

no VMs should be able to use external IPs

`constraints/compute.vmExternalIpAccess` → `true`

DMZ appliances should be able to use external IPs

inheritance

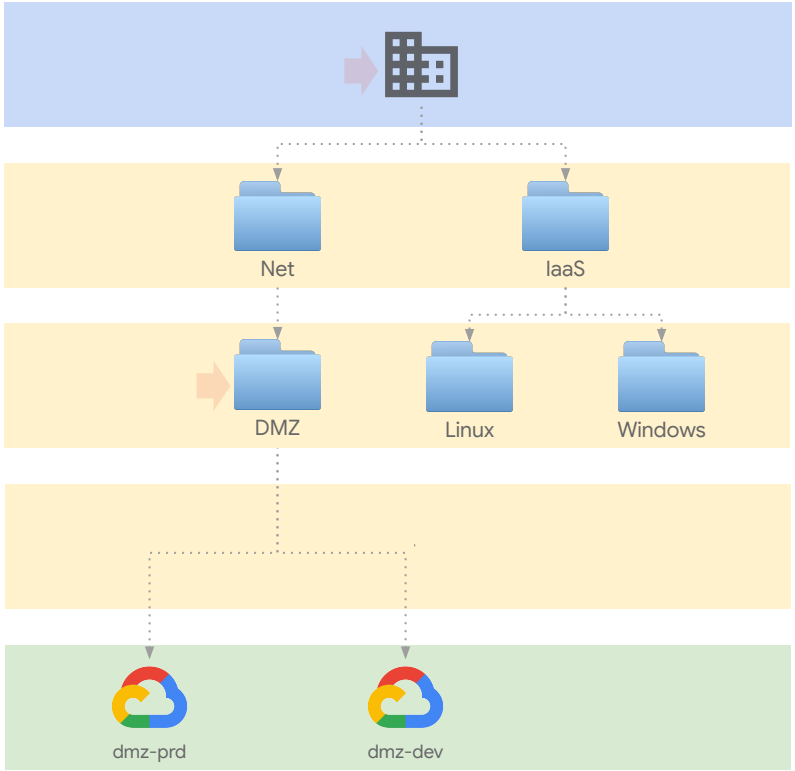

Organization Policies vs IAM Policies

Organization Policies	IAM Policies
Constraints that allow you to: ● Limit resource sharing based on domain. ● Limit the usage of Identity and Access Management service accounts. ● Restrict the physical location of newly created resources.	Effectively they're bindings which specify what access should be granted to principal on resources.
Focuses on "what". Allows to set restrictions on specific resources to determine how they can be configured	Focuses on "who". Let's you authorize who can take action on specific resources based on permissions
Can be set on different levels (org, folder, project), propagate down but lower-level policy overwrites a higher-level one.	Effective IAM Policy on each level is a SUM of all privileges (* with an exception of "deny policies", which are not covered on the exam as of Q1 '23)
Both should be used as part of a security posture! It's NOT one or the other.	

Service Accounts

Differences in service-to-service authentication (1/3)

AWS

- Applications use IAM roles and **instance profiles** for permissions management on applications.
- Instance profiles are a container for an IAM role that can be attached to an application running in an AWS EC2 container, providing the permissions granted by the named role.

Google Cloud

- Service accounts are a type of member in Google Cloud IAM that let you give access permissions to virtual machines and other services
- Google Cloud uses **service accounts** to control service-to-service authentication using Google Cloud IAM.

Differences in service-to-service authentication (2/3)

Topic	AWS	Google Cloud
Access management	AWS IAM roles	Google Cloud IAM roles
Principle type	A role with the appropriate permissions is created and configured in the instance profile associated with a service.	Service accounts are assigned to Google Cloud instances and can be used to access Google Cloud services, resources, and application code.

Differences in service-to-service authentication (3/3)

Topic	AWS	Google Cloud
Authentication	Instance profiles contain a specified role that can be used to access specified AWS resources without the need for credential management. The instance profile handles temporary credentialing and rotation of those credentials.	Service accounts are named with an email address and use cryptographic keys to access resources.

Service accounts (IAM Roles)

- Provide an identity for carrying out server-to-server interactions
- Programs running within Compute Engine (EC2) instances can automatically acquire access tokens with credentials.
- Tokens are used to access to any service API in your project (account) and any other services that granted access to that service account (role).
- Service accounts (roles) are convenient when you're not accessing user data.

Service account use cases

- Typically, service accounts are used in scenarios such as:
  - Running workloads on virtual machines (VMs)
    - E.g., Create a service account with permissions to query BigQuery
      - Attach it to a VM
      - Deploy an application onto the VM that submits SQL commands to BigQuery
  - Running workloads on on-premises workstations or data centers that call Google APIs.
    - E.g., Same example as above, but now the application is running on-premise
  - Running workloads which are not tied to the lifecycle of a human user.
    - E.g., Batch jobs that are scheduled to run periodically

Proprietary + Confidential

Service account keys pose a security risk to your resources

- SA keys are similar to a **password without an expiration date**.
- SA Keys can be leaked accidently and attackers can use it to access your (GCP project or org admin) sensitive GCP resources.
- Usage cannot be audited → compounding the risk

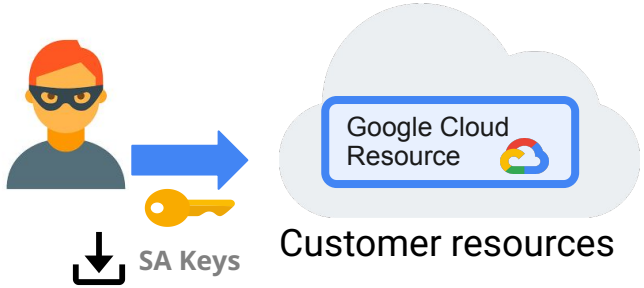

Customers have downloaded > 48 Million Service Account Keys!!

So what's the solution? Ditch the keys and use Workload Identity for GKE & Workload Identity Federation!

Workload Identity Federation: Keyless Access

User Story: As an App Developer, I want to **securely connect my services (= NOT users, like in Workforce Identity Federation!)** to GCP resources without downloading access keys.

Benefits

Keyless access to GCP APIs

Auditability through Cloud logs

Attribute-based access control

Exam Tips:

- Have a look at a great explanation of WIF
- Have a look at Workflow Identity Federation documentation

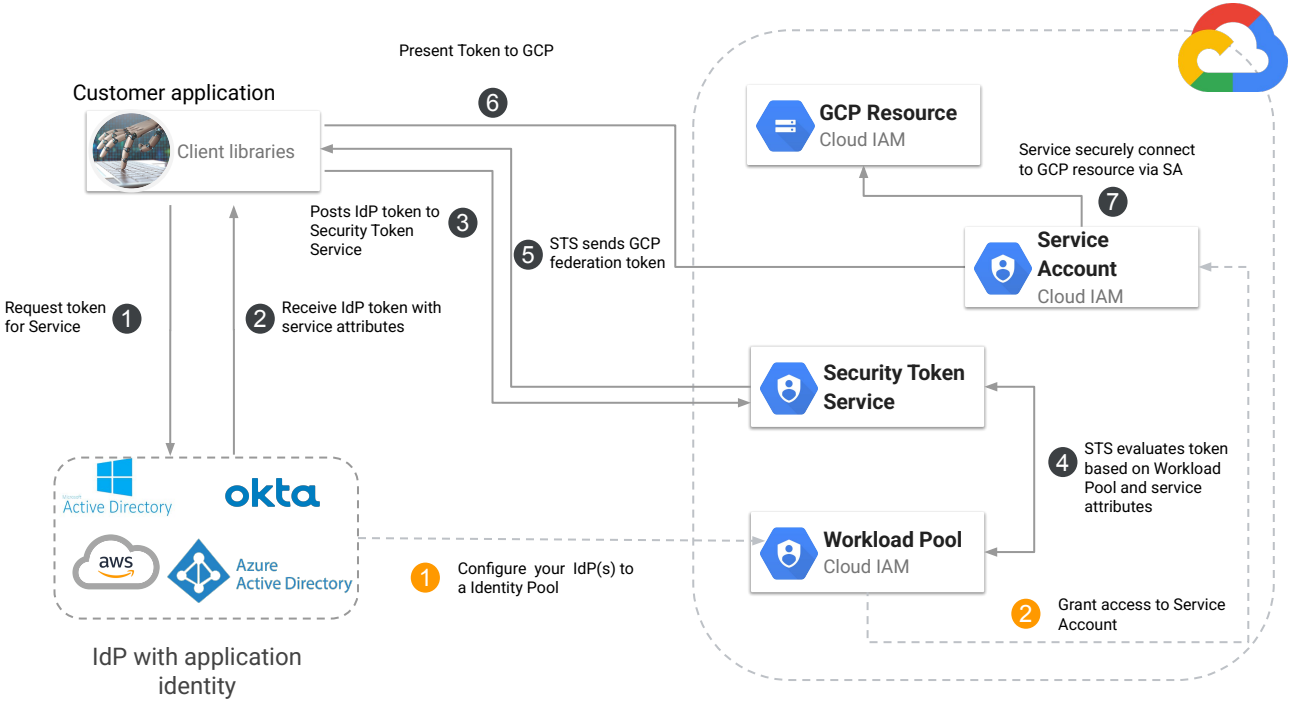

Best practice

Service accounts: Best practices

Workflow

- **Avoid** using the **Default Compute Engine service account**
- Use **dedicated custom service-accounts** for running VM's with **minimal required permissions**
- Consider using service accounts to apply **firewalls** (more details under *Networking* topic)
- For easier **visibility and auditing**, **centrally create** service accounts in **dedicated projects**

Security

- **Don't embed** service account keys (or any authentication secrets) in your code
- **Prevent committing** keys to external source repositories
- **Rotate** user-managed service account **keys frequently**
- **Utilize VPC Service Controls** to create a perimeter around who can authenticate with **Google services** (more details under *Networking* topic)
- Use Forseti to detect old service account keys

Diagnostic Question Discussion

Cymbal Direct's social media app must run in a **separate project** from its APIs and web store. You want to use **Identity and Access Management (IAM)** to ensure a **secure environment.**

How should you set up IAM?

A. Use **separate** service accounts for each component (social media app, APIs, and web store) with **basic** roles to grant access.

B. Use **one** service account for all components (social media app, APIs, and web store) with **basic** roles to grant access.

C. Use separate service accounts for each component (social media app, APIs, and web store) with **predefined or custom** roles to grant access.

D. Use **one** service account for all components (social media app, APIs, and web store) with **predefined or custom** roles to grant access.

Diagnostic Question Discussion

Cymbal Direct's social media app must run in a **separate project** from its APIs and web store. You want to use **Identity and Access Management (IAM)** to ensure a **secure environment.**

How should you set up IAM?

A. Use **separate** service accounts for each component (social media app, APIs, and web store) with **basic** roles to grant access.

B. Use **one** service account for all components (social media app, APIs, and web store) with **basic** roles to grant access.

C. **Use separate service accounts for each component (social media app, APIs, and web store) with predefined or custom roles to grant access.**

D. Use **one** service account for all components (social media app, APIs, and web store) with **predefined or custom** roles to grant access.

Diagnostic Question Discussion

You are writing an application that will be required to read and write from Cloud Storage. You want to ensure that you do not cause any security issues in your configuration.

What would be the best option for granting the application access to Cloud Storage?

A. Create a user account and apply a curated role that gives the user read/write permissions to the Cloud Storage bucket. Then, in the application configuration file, enter the username and password of the account in the file.

B. Open all access to the Cloud Storage bucket so it is accessible to the public, then use the bucket for your application.

C. Create a service account that only has permissions to read and write from the Cloud Storage bucket you are using for the application. Then, apply the service account to the virtual machine that is running your application.

D. Ensure the scope of the instance has full privileges to Cloud Storage.

Diagnostic Question Discussion

You are writing an application that will be required to read and write from Cloud Storage. You want to ensure that you do not cause any security issues in your configuration.

What would be the best option for granting the application access to Cloud Storage?

A. Create a user account and apply a curated role that gives the user read/write permissions to the Cloud Storage bucket. Then, in the application configuration file, enter the username and password of the account in the file.

B. Open all access to the Cloud Storage bucket so it is accessible to the public, then use the bucket for your application.

C. Create a service account that only has permissions to read and write from the Cloud Storage bucket you are using for the application. Then, apply the service account to the virtual machine that is running your application.

D. Ensure the scope of the instance has full privileges to Cloud Storage.

Diagnostic Question Discussion

Which of the following are the best practices recommended by Google Cloud when dealing with Service Accounts? (choose 3):

A. Grant the Service Accounts full set of permissions

B. Do not delete service accounts that are in use by running instances on App Engine or Compute Engine

C. Grant “Service Account User” role to all users in the organization

D. When you create a service account, give it a meaningful name to indicate the purpose of the service account

E. Create separate service accounts for different purposes with minimal set of privileges required for that service.

Diagnostic Question Discussion

Which of the following are the best practices recommended by Google Cloud when dealing with Service Accounts? (choose 3):

A. Grant the Service Accounts full set of permissions

B. Do not delete service accounts that are in use by running instances on App Engine or Compute Engine

C. Grant “Service Account User” role to all users in the organization

D. When you create a service account, give it a meaningful name to indicate the purpose of the service account

E. Create separate service accounts for different purposes with minimal set of privileges required for that service.

Make sure to...

Enjoy the journey as much as the destination!