This guide provides an overview of how SAP Adaptive Server Enterprise (ASE) works on Google Cloud Platform (GCP), and provides details that you can use when planning the implementation of a new SAP ASE system.

For details about how to deploy SAP ASE on GCP, see:

- SAP ASE Deployment Guide for Linux (/solutions/sap/docs/sap-ase-deployment-linux)

- SAP ASE Deployment Guide for Windows (/solutions/sap/docs/sap-ase-deployment-windows)

GCP consists of many cloud-based services and products. When running SAP products on GCP, you mainly use the IaaS-based services offered through Compute Engine (/compute/docs/) and Cloud Storage (/storage/docs/), as well as some platform-wide features, such as tools (/docs/overview/developer-and-admin-tools).

See the GCP platform overview (/docs/overview/) for important concepts and terminology. This guide duplicates some information from the overview for convenience and context.

For an overview of considerations that enterprise-scale organizations should take into account when running on GCP, see best practices for enterprise organizations (/docs/enterprise/best-practices-for-enterprise-organizations).

GCP offers three main ways to interact with the platform, and your resources, in the cloud:

- The Google Cloud Console, which is a web-based user interface.

- The `gcloud` command-line tool, which provides a superset of the functionality that Cloud Console offers.

- Client libraries (/sdk/cloud-client-libraries), which provide APIs for accessing services and management of resources. Client libraries are useful when building your own tools.

tant: If you are using the `gcloud` command-line tool, always make sure that you are using the latest version of the GC o install the latest version of the GCP SDK, see the gcloud install instructions. (https://cloud.google.com/sdk/downlc

SAP deployments typically utilize some or all of the following GCP services:

Service	Description
VPC Networking (/compute/docs/networking)	Connects your VM instances to each other and to the Internet. Each instance is a member of either a legacy network with a single global IP range, or a recommended subnet network, where the instance is a member of a single subnetwork that is a member of a larger network. Note that a network cannot span GCP projects, but a GCP project can have multiple networks.
Compute Engine (/compute/)	Creates and manages VMs with your choice of operating system and software stack.
Persistent disks (/compute/docs/disks/)	Persistent disks are available as either standard hard disk drives (HDD) or solid-state drives (SSD).
Google Cloud Console (https://console.cloud.google.com/)	Browser-based tool for managing Compute Engine resources. Use a template to describe all of the Compute Engine resources and instances you need. You don't have to individually create and configure the resources or figure out dependencies, because the Cloud Console does that for you.
Cloud Storage (/storage/)	You can back up your SAP database backups into Cloud Storage for added durability and reliability, with replication.
Stackdriver Monitoring (/monitoring/)	Provides visibility into the deployment, performance, uptime, and health of Compute Engine, network, and persistent disks. Stackdriver collects metrics, events, and metadata from GCP and uses these to generate insights through dashboards, charts, and alerts. You can monitor the compute metrics at no cost through Stackdriver Monitoring.
Cloud IAM (/iam/docs/)	Provides unified control over permissions for GCP resources. Control who can perform control-plane operations on your VMs, including creating, modifying, and deleting VMs and persistent disks, and creating and modifying networks.

You can use the pricing calculator (/products/calculator/) to estimate your usage costs. For more pricing information, see Compute Engine pricing (/compute/pricing), Cloud Storage pricing (/storage/pricing), and Stackdriver pricing (/stackdriver/pricing_v2).

GCP resources are subject to quotas. If you plan to use high-CPU or high-memory machines, you might need to request additional quota. For more information, see Compute Engine resource quotas.

(/compute/quotas)

A basic single-node SAP ASE installation on GCP comprises the following components:

- One Compute Engine VM running your SAP ASE database.

- Four or five attached persistent disk drives:

Drive contents	Linux	Windows
Root directory of the database instance	sybase/[DBSID]	ASE (D:)
Database data files	/sybase/[DBSID]/sapdata	ASE Data (E:)
Database transaction logs	/sybase/[DBSID]/logdir	ASE Log (L:)
Database temporary table space	/sybase/[DBSID]/saptmp	ASE Temp (T:)
Diagnostic tablespace for SAPTOOLS	/sybase/[DBSID]/sapdiag	Not applicable

Optionally, you can expand your installation to include the following as well:

- A drive for backups. On Linux the backup drive is `/sybasebackup`. On Windows it is `Backup (X:)`.

- A NAT gateway. A NAT gateway allows you to provide Internet connectivity for your VMs while denying direct Internet connectivity to those VMs. You could also configure this VM as a bastion host that allows you to establish SSH connections to the other VMs on your private subnet. See NAT gateways and bastion hosts (#nat_gateways_and_bastion_hosts) for more information.

- NetWeaver directories, including:

  - `/usr/sap` on Linux or `SAP (S:)` on Windows

  - `/sapmnt` on Linux or `Pagefile (P:)` on Windows

- A warm standby server, which you need to set up if you're using SAP ASE's always-on high-availability disaster recovery (HADR) option. The warm standby server requires the same disk configuration as a basic SAP ASE installation.

  For important software requirements for the SAP ASE always-on option, see Supported SAP ASE features (#supported_sap_ase_features).

For more information about setting up HADR for SAP ASE, see the HADR Users Guide (https://help.sap.com/viewer/efe56ad3cad0467d837c8ff1ac6ba75c/latest).

Fault manager servers, such as Pacemaker and other such high-availability resource management software, are not currently supported with SAP ASE on Google Cloud.

Different use cases might require additional devices or databases. For more information, see SAP ASE Configuration Guide for UNIX > Optional Devices and Databases (https://help.sap.com/viewer/ecbccd52e7024feaa12f4e780b43bc3b/latest/en-US/a6a62df6bc2b1014aa43a12fd476bfca.html)
.

You might be able to achieve a sustained use discount (/compute/docs/sustained-use-discounts) of up to 30% for y⌷ ⌷te Engine instances. Use the pricing calculator (/pricing/calculator) to help you estimate your actual costs.

In many ways, running SAP ASE on GCP is similar to running it in your own data center. You still need to think about computing resources, storage, and networking considerations. For more information, see SAP Note 2537664: SAP ASE 16.0 Certification Report for GCP (https://launchpad.support.sap.com/#/notes/2537664).

SAP ASE is certified to run on all Compute Engine machine types, including custom types. For information about different machine types and their use cases, see Machine Types (/compute/docs/machine-types) in the Compute Engine documentation.

The number of vCPUs required varies depending on the application load on SAP ASE. You should allocate a minimum of two vCPUs to your SAP ASE installation. Follow the Performance and Tuning Guides (https://help.sap.com/viewer/91d32d977a174c68829880bc020fc352/latest) for SAP ASE to achieve best use of existing resources, and then increase your computing resources as needed.

Your SAP ASE VM should have at least 4 GB of RAM per vCPU. Of this amount, approximately 80% of your RAM should be allocated to SAP ASE, with the rest allocated to the OS on which SAP ASE is running.

The optimal amount of memory for your use case depends on the complexity of the queries you're running, the size of your data, the amount of parallelism you're using, and the level of performance you're expecting. For further guidance about optimizing your memory configuration, see the SAP ASE Performance and Tuning Series (https://help.sap.com/viewer/91d32d977a174c68829880bc020fc352/latest).

By default, each Compute Engine VM has a small root persistent disk (/compute/docs/disks/create-root-persistent-disks) that contains the operating system. In addition, you should create, attach, format, and mount additional disks for your database, your logs, and your stored procedures.

You can use standard HDD persistent disks or SSD persistent disks as storage for your SAP ASE VMs. The performance of your persistent disks depends on the disk size and the number of vCPUs in the host machine. For a detailed overview of persistent disk performance benchmarks, see Optimizing Persistent Disk and Local SSD Performance (/compute/docs/disks/performance).

Your disk size and performance requirements will depend on your application. Size each device according to your needs. See Configuration Guide for UNIX > Determine the Size of a Database Device (https://help.sap.com/viewer/ecbccd52e7024feaa12f4e780b43bc3b/latest/en-US/a6ff9925bc2b1014a378fef4b820a3a9.html) for guidance.

For a high-level description of persistent disks, see Persistent disks (#persistent_disks) below.

SAP has certified GCP to support the following SAP ASE versions:

- SAP ASE 16.0 SP02 PL06 HF1

- SAP ASE 16.0 SP02 PL07

For more information about the supported SAP ASE versions on GCP, see SAP Note 2537664 (https://launchpad.support.sap.com/#/notes/2537664).

SAP supports most SAP ASE 16 features on GCP. However, the following features are unsupported:

- Real Time Data Services

- Web Services

- Kerberos

- IPv6

- SAP ASE High Availability option with native cluster managers

- Raw disk as database devices

- Kernel Process Mode

- Tivoli Storage Management

The SAP ASE always-on option is supported on the following operating systems in GCP:

- Red Hat Enterprise Linux (RHEL) Server 7.4

- SUSE Linux Enterprise Server (SLES) 12 SP3

- Windows Server 2012 R2

On Linux, the SAP ASE alway-on option has the following requirements:

- ASE 16.0 SP02 PL07

- The `libgcc` system library

    - On RHEL 7.4, install `libgcc.i686`

    - On SLES 12 SP3, install `libgcc_s1`

- The `glibc` system library

    - On RHEL 7.4, install `glibc.i686`

    - On SLES 12 SP3, install `glibc`

For more information about the supported SAP ASE features on GCP, see SAP Note 2537664 (https://launchpad.support.sap.com/#/notes/2537664).
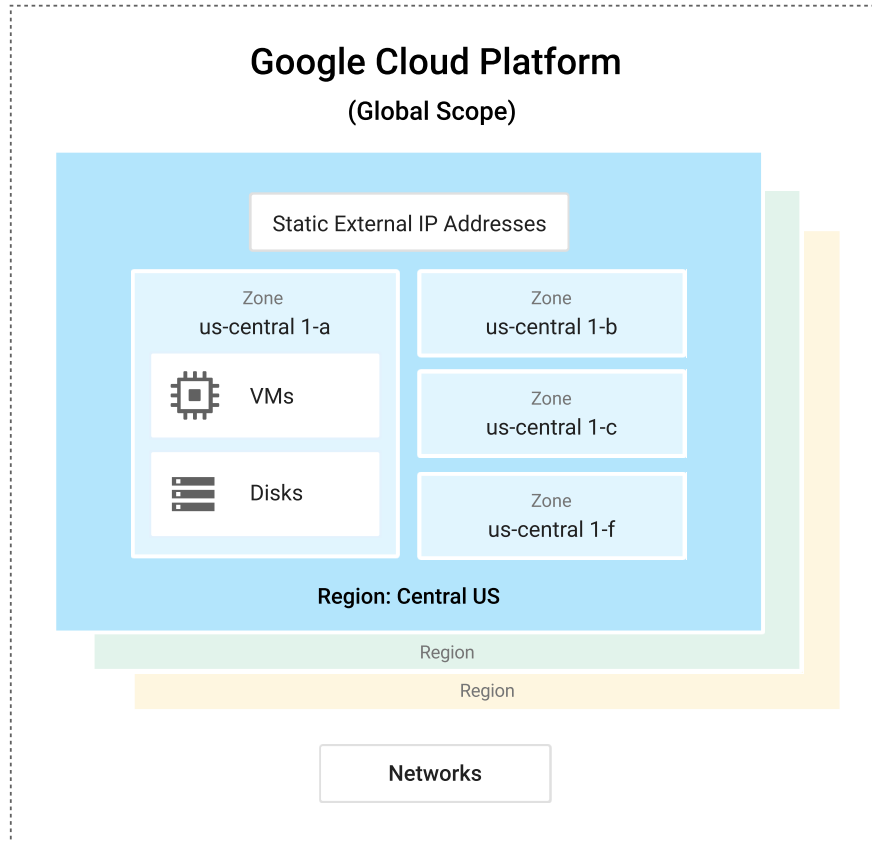
SAP has certified GCP to run SAP ASE on the following SUSE Linux Enterprise Server (SLES), Red Hat Enterprise Linux (RHEL), and Windows Server operating system images:

- RHEL 7.3 and 7.4

- SLES 12 SP2 and SP3

- Windows Server 2012 R2

For more information about Compute Engine images, see Images (/compute/docs/images).

When you deploy a VM, you must choose a region and zone. A region is a specific geographical location where you can run your resources, and corresponds to a data center location. Each region has one or more zones.

Global resources, such as preconfigured disk images and disk snapshots, can be accessed across regions and zones. Regional resources, such as static external IP addresses, can be accessed only by resources that are in the same region. Zonal resources, such as VMs and disks, can be accessed only by resources that are located in the same zone.

Google Cloud Platform

(Global Scope)

Static External IP Addresses

Zone us-central 1-a	Zone us-central 1-b
VMs	Zone us-central 1-c
Disks	Zone us-central 1-f

Region: Central US

Region

Region

Networks

When choosing regions and zones for your VMs, keep the following in mind:

- **The location of your users and your internal resources**, such as your data center or corporate network. To decrease latency, select a location that is in close proximity to your users and resources.

- **The location of your other SAP resources**. Your SAP application and your database must be in the same zone.

Persistent disks are durable storage devices that function similarly to the physical disks in a desktop or a server. Google manages the hardware behind these devices to ensure data redundancy and to optimize performance. Persistent disks are available as either standard hard disk drives (HDD) or solid-state drives (SSD). Standard HDD persistent disks are efficient and economical for handling sequential read-write operations, but are not optimized to handle high rates of random input-output operations per second (IOPS).

Persistent disks are located independently from your VMs, so you can detach or move persistent disks to keep your data, even after you delete your VMs. Persistent disk performance scales automatically with disk size, so you can resize your existing persistent disks or add more persistent disks to a VM to meet your performance and storage space requirements.

GCP also offers local SSD (/compute/docs/disks/#localssds) disk drives. Although local SSDs can offer some advantages over persistent disks, don't use them as part of an SAP ASE system. VM instances with local SSDs attached cannot be stopped and then restarted.

If your security policy requires truly internal VMs, you need to set up a NAT proxy manually on your network and a corresponding route so that VMs can reach the Internet. It is important to note that you cannot connect to a fully internal VM instance directly by using SSH. To connect to such internal machines, you must set up a bastion instance that has an external IP address and then tunnel through it. When VMs do not have external IP addresses, they can be reached only by other VMs on the network, or through a managed VPN gateway. You can provision VMs in your network to act as trusted relays for inbound connections, called bastion hosts, or network egress, called NAT gateways. For more transparent connectivity without setting up such connections, you can use a managed VPN gateway resource.

Bastion hosts (https://wikipedia.org/wiki/Bastion_host) provide an external facing point of entry into a network containing private-network VMs. This host can provide a single point of fortification or audit and can be started and stopped to enable or disable inbound SSH communication from the Internet.
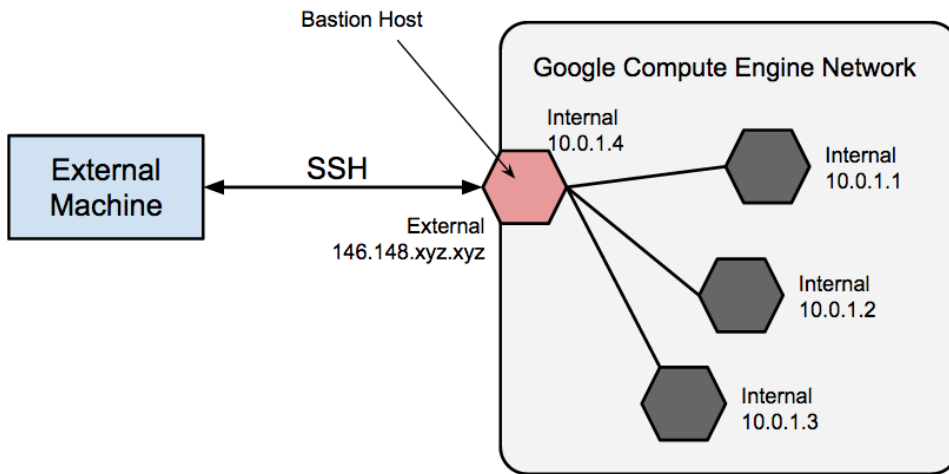
You can achieve SSH access to VMs that do not have an external IP address by first connecting to a bastion host. A complete hardening of a bastion host is outside the scope of this guide, but you can take some initial steps, including:

- Limit the CIDR range of source IPs that can communicate with the bastion.

- Configure firewall rules to allow SSH traffic to private VMs from only the bastion host.

By default, SSH on VMs is configured to use private keys for authentication. When using a bastion host, you log into the bastion host first, and then into your target private VM. Due to this two-step login, you should use SSH-agent forwarding to reach the target VM instead of storing the target VM's private key on the bastion host. You must do this even if you are using the same key-pair for both bastion and target VMs, as the bastion has direct access only to the public half of the key-pair.

When a VM does not have an assigned, external IP address, it cannot make direct connections to external services, including other GCP services. To allow these VMs to reach services on the Internet, you can set up and configure a NAT gateway (/compute/docs/vpc/special-configurations#natgateway). The NAT gateway is a VM that can route traffic on behalf of any other VM on the network. You should have one NAT gateway per network. Be aware that a single-VM NAT gateway should not be considered highly available, and cannot support high traffic throughput for multiple VMs. For instructions on how to set up a VM to act as a NAT gateway, see either the SAP ASE Deployment Guide for Linux (/solutions/sap/docs/sap-ase-deployment-linux#setting-up-a-nat-gateway) or the SAP ASE Deployment Guide for Windows (/solutions/sap/docs/sap-ase-deployment-windows#setting-up-a-nat-gateway).

After your system is up and running, you can <u>create custom images</u> (/compute/docs/images/create-delete-deprecate-private-images). You should create these images when you modify the state of your root persistent disk and want to be able to easily restore the new state. You should have a plan for how to manage the custom images you create. For more information, see <u>Image Management Best Practices</u> (/solutions/image-management-best-practices).

When planning security for an SAP deployment on Google Cloud, you must identify:

- The user accounts and applications that need access to the Google Cloud resources in your Google Cloud project

- The specific Google Cloud resources in your project that each user needs to access

You must add each user to your project by adding their Google account ID to the project as a member. For an application program that uses Google Cloud resources, you create a *service account*, which provides a user identity for the program within your project.

Compute Engine VMs have their own service account. Any programs that that run on a VM can use the VM service account, as long as the VM service account has the resource permissions that the program needs.

After you identify the Google Cloud resources that each user needs to use, you grant each user permission to use each resource by assigning resource-specific roles to the user. Review the predefined roles that Cloud IAM provides for each resource, and assign roles to each user that provide just enough permissions to complete the user's tasks or functions and no more.

If you need more granular or restrictive control over permissions than the predefined Cloud IAM roles provide, you can create custom roles.

For more information about the Cloud IAM roles that SAP programs need on Google Cloud, see <u>Identity and access management for SAP programs on Google Cloud</u> (/solutions/sap/docs/security-for-sap/iam-for-sap-programs).

For an overview of identity and access management for SAP on Google Cloud, see <u>Identity and access management overview for SAP on Google Cloud</u> (/solutions/sap/docs/security-for-sap/iam-for-sap).

Consider the information in the following sections when planning networking and security.

One of your first lines of defense is to restrict who can reach your network and your VMs by using firewalls (/compute/docs/vpc/firewalls). By default, all traffic to VMs, even from other VMs, is blocked by the firewall unless you create rules to allow access. The exception is the default network that is created automatically with each project and has default firewall rules (/compute/docs/networking#firewall_rules).

By creating firewall rules, you can restrict all traffic on a given set of ports to specific source IP addresses. You should follow the minimum privilege model to restrict access to the specific IP addresses, protocols, and ports that need access. For example, you should always set up a bastion host (/solutions/sap/docs/netweaver-planning-guide#bastion_hosts_nat) and allow SSH into your SAP NetWeaver system only from that host.

You can use a network to define a gateway IP and the network range for the VMs attached to that network. All Compute Engine networks use the IPv4 (http://tools.ietf.org/html/rfc791) protocol. Every GCP project is provided with a default network with preset configurations and firewall rules, but you should add a custom subnetwork and firewall rules based on a minimum privilege model. By default, a newly created network has no firewall rules and hence no network access.

Depending on your requirements, you might want to add additional subnetworks to isolate parts of your network. For more information, see Subnetworks (/compute/docs/vpc/#vpc_networks_and_subnets).

The firewall rules apply to the entire network and all the VMs in the network. You can add a firewall rule that allows traffic between VMs in the same network and across subnetworks. You can also configure firewalls to apply to specific target VMs by using the tagging mechanism (/compute/docs/vpc/add-remove-network-tags).

Some SAP products, such as SAP NetWeaver, require access to certain ports. Be sure to add firewall rules to allow access to the ports outlined by SAP (https://help.sap.com/viewer/575a9f0e56f34c6e8138439eefc32b16/2.0/en-US/616a3c0b1cc748238de9c0341b15c63c.html) .

Routes (/compute/docs/vpc/routes) are global resources attached to a single network. User-created routes apply to all VMs in a network. This means you can add a route that forwards traffic from VM to VM within the same network and across subnetworks without requiring external IP addresses.

For external access to Internet resources, launch a VM with no external IP address and configure another virtual machine as a NAT gateway. This configuration requires you to add your NAT gateway as a route for your SAP instance. For more information, see NAT gateways and bastion hosts (#nat_gateways_and_bastion_hosts).

You can securely connect your existing network to GCP through a VPN (https://wikipedia.org/wiki/Virtual_private_network) connection using IPsec (https://wikipedia.org/wiki/IPsec) by using Google Cloud VPN. Traffic traveling between the two networks is encrypted by one VPN gateway, then decrypted by the other VPN gateway. This protects your data as it travels over the Internet. You can dynamically control which VMs can send traffic down the VPN using instance tags on routes (/compute/docs/vpc/using-routes#addingroute). Cloud VPN tunnels are billed at a static monthly rate plus standard egress charges. Note that connecting two networks in the same project still incurs standard egress charges. For more information, see the VPN Overview (/vpn/docs/concepts/overview) and Choosing a VPN Routing Option (/vpn/docs/how-to/choosing-a-vpn).

If you use Cloud Storage to host your backups for your data and log, make sure you use TLS (HTTPS) while sending data to Cloud Storage from your VMs to protect data in transit. Cloud Storage automatically encrypts data at rest. You can specify your own encryption keys if you have your own key-management system.

For security best practices, see Cloud Storage Security (/storage/docs/best-practices#security).

Refer to the following additional security resources for your SAP environment on GCP:

- Securely Connecting to VM Instances (/solutions/connecting-securely)

- Security Center (/security/)

- Compliance in the Google Cloud (/security/compliance)

- Google Cloud security whitepaper (/security/whitepaper)

- Google Infrastructure security design (/security/security-design/resources/google_infrastructure_whitepaper_fa.pdf)

You must have a plan for how to restore your system to operating condition if the worst happens. For general guidance about how to plan for disaster recovery using GCP, see:

- How to Design a Disaster Recovery Plan (/solutions/designing-a-disaster-recovery-plan)

- Disaster Recovery Cookbook (/solutions/disaster-recovery-cookbook)

This section provides information about licensing requirements.

Running SAP ASE on GCP requires you to bring your own license (BYOL). For more information, see:

- SAP Note 2446441 - Linux on Google Cloud Platform (IaaS): Adaption of your SAP License (https://launchpad.support.sap.com/#/notes/2446441)

- SAP Note 2456953 - Windows on Google Cloud Platform (IaaS): Adaption of your SAP License (https://launchpad.support.sap.com/#/notes/2456953)

For more information about SAP licensing, contact SAP.

In Compute Engine, there are two ways to license SLES, RHEL, and Windows Server:

- With pay-as-you-go licensing, your Compute Engine VM hourly cost includes licensing. Google manages the licensing logistics. Your hourly costs are higher, but you have complete flexibility to increase and decrease your costs, as needed. This is the licensing model used for GCP public images that include SLES, RHEL, and Windows Server.

⭐ **Important:** RHEL, SLES, and Windows Server are premium images. Premium images charge for OS licenses in addition to standard VM charges. For details about pricing, refer to the pricing for RHEL images (/compute/disks-image-pricing#rhel_images), pricing for SUSE images (/compute/disks-image-pricing#suse_images), and pricing for Windows Server images (/compute/disks-image-pricing#windows_server_pricing).

- With BYOL, your Compute Engine VM costs are lower because the licensing isn't included. You must migrate an existing license or purchase your own license, which means paying up front, and you have less flexibility.

Google Cloud customers either with a Production Support Role or with Enterprise Support can request assistance with the provisioning and configuration of the Google Cloud resources that are required for SAP systems. Google Cloud Production-level support or Enterprise support is required for support of SAP systems in production environments.

For more information about Google Cloud support options, see Google Cloud Support (/support/).

For SAP product-related issues, log your support request with SAP support (https://support.sap.com/support-programs-services/about/getting-started.html). SAP evaluates the support ticket and, if it appears to be a Google Cloud infrastructure issue, transfers the ticket to the Google Cloud queue.

To deploy SAP ASE on Linux, see the SAP ASE Deployment Guide for Linux (/solutions/sap/docs/sap-ase-deployment-linux).

To deploy SAP ASE on Windows, see the SAP ASE Deployment Guide for Windows (/solutions/sap/docs/sap-ase-deployment-windows).