You can use Google Cloud to host disaster recovery solutions for your SAP systems that are running on Google Cloud, on-premises, or on another cloud provider.

Disaster recovery (DR) and *high availability* (HA) are two distinct elements in the larger concept of *business continuity*. This guide focuses on disaster recovery.

A DR solution is designed to restore application processing after a natural or man-made disaster or an infrastructure failure causes a wide-scale outage. Such outages disable not only the primary application processing system, but any standby system that protects against small-scale system or local infrastructure failures.

Other characteristics of a DR solution that distinguish it from a high-availability solution include:

- The recovery system in a DR solution is typically not a hot-standby system.
- A disaster recovery procedure usually requires manual intervention to recover or restore application processing from backups or replication of the data, systems, and infrastructure.
- The solution includes a recovery site that is in a different geographical location than the primary system.
- The recovery time objective (RTO) (https://wikipedia.org/wiki/Disaster_recovery#Recovery_time_objective) for a disaster recovery solution is typically measured in hours, if not days.

The following example architecture shows an SAP system that includes both an HA solution and a DR solution. The DR site, in which the system will be restored after a disaster, is on the right in **Region 2**. The primary system is on the left in **Region 1** and is configured for high availability with failover clusters that span two zones. The backup function of the DR solution is represented by dashed lines that span both regions. For storage, the systems use a multi-regional Cloud Storage bucket, which is shown at the bottom of the diagram.

The example also shows a database. Although application recovery depends on data recovery, DR procedures for database systems are not covered in this guide. Refer to the database documentation and any applicable SAP notes. This guide focuses specifically on the SAP NetWeaver system (SAP ASCS/ERS) and the application servers (SAP PAS and AAS).

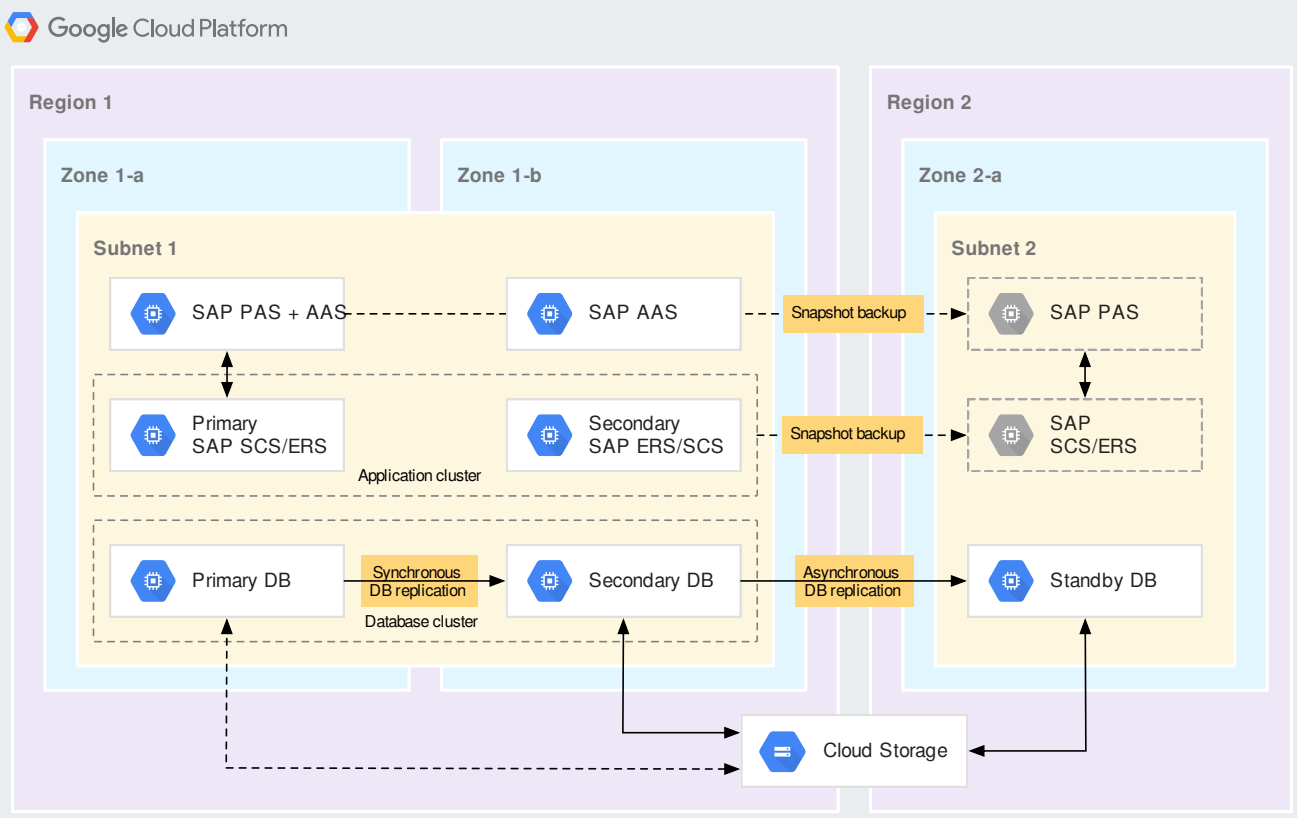

For information about implementing high-availability SAP systems on Google Cloud, see High-availability planning guide for SAP NetWeaver on Google Cloud (/solutions/sap/docs/netweaver-ha-planning-guide).

For information about high-availability and disaster recovery for SAP HANA on Google Cloud, see SAP HANA high availability and disaster recovery planning guide (/solutions/sap/docs/sap-hana-hadr-planning-guide).

For general information about DR on Google Cloud, see Disaster Recovery Planning Guide (/solutions/dr-scenarios-planning-guide).

If your primary SAP system is not running on Google Cloud, you can take a *lift-and-shift* approach to your DR solution, in which the architecture and software of the DR system on Google Cloud is the same as that of the primary system. Or you can take a *cloud-native* approach, in which, as a part of the design of your DR solution, you optimize the recovered system for the cloud, supported by Google Cloud or Google Cloud partner products or services.

If you use a lift-and-shift approach, you do need to confirm that the architecture of your primary system is completely supported on Google Cloud. For either approach, you need to make sure that all of the software that you use is properly licensed for use on Google Cloud.

For more information about licensing, see Google Cloud Platform Terms of Service (/terms/).

The design elements of a DR solution for SAP NetWeaver on Google Cloud can include the following:

- DR site location
- Networking
- Security
- Virtual machine (VM) considerations
- Backup options
- Storage

To implement each of these design elements, you have a number of options to meet your recovery and cost objectives.

When choosing a Google Cloud region for your DR site, you need to consider:

- The potential impact area of any disaster that might occur at your primary site
- The location of the users of your SAP NetWeaver system

- Whether the Google Cloud resources and features that your SAP system uses are available in the region and zone that you choose

Choose a Google Cloud region for your DR site that is far enough away from your primary site so that the DR site will be unaffected by any disaster that might occur at the primary site. A distance of 100 miles or more is usually considered sufficient, but regulations or organizational guidelines might require different minimum distance.

If your primary SAP system is running on Google Cloud, place your DR site in any Google Cloud region that is close to your users, other than the region that your primary system is running in. Google Cloud regions are located far enough away from each other so that no two Google Cloud regions are likely to be impacted by the same disaster.

If your primary SAP system is running outside of Google Cloud, place your DR site in a region that is as close to your users as possible without being in the potential impact zone of any disaster that might occur at the primary site.

In your DR region, place your DR system in a zone that supports the VM instance types and other infrastructure that your SAP and database systems require.

After you select a region for your DR site, you might need to increase your resource quotas (/compute/quotas) for the region to provide enough resources for the DR system prior to an event.

For the locations of all of the Google Cloud regions, see Regions and zones (/compute/docs/regions-zones/).

To see the features that are available in each region, see Available regions & zones (/compute/docs/regions-zones/#available).

To review which Google Cloud resources are global, regional, and zonal, see Global, Regional, and Zonal Resources (/compute/docs/regions-zones/global-regional-zonal-resources).

On Google Cloud, Virtual Private Cloud (VPC) provides networking functionality that can extend around the globe.

You need to create a VPC network for your DR site if one does not already exist. You also need to create a subnetwork and IP range for the DR site.

If your primary system is on Google Cloud, configuring your network is easier if both the primary and DR sites are in the same VPC network. However, if necessary, you can instead place the primary and DR sites in different VPC networks or even different projects.

When designing your DR solution, you need to consider the following communication paths:

- The connection between the primary site and the recovery site in Google Cloud
- The internal communication between the applications, databases, and servers that make up your SAP system
- The connection between your users and the SAP system

For connections from sites that are external to Google Cloud, Google Cloud provides a variety of networking products (/products/networking) to support each of these connection points.

The connection between the primary site and the DR site is required to store backups or provide a replication path between the two systems so that the recovery resources are immediately available in the event of a disaster and for testing your disaster procedures.

If your primary system is running on Google Cloud, making your backups available at the DR site is almost automatic. Compute Engine snapshots can be designated as multi-regional. Other backups can be stored directly from the primary system into multi-regional Cloud Storage buckets for instant availability at the DR site.

If your primary system is not running on Google Cloud, you can connect your primary site to your DR site on Google Cloud by using Cloud Interconnect (/interconnect/docs/concepts/overview) or Cloud VPN (/vpn/docs/concepts/overview).

For an example of a highly available Cloud Interconnect topology that, with a few changes, could be adapted to a DR scenario, see Establishing 99.99% Availability for Dedicated Interconnect (/interconnect/docs/tutorials/dedicated-creating-9999-availability).

For some examples of highly available, multi-region VPN gateway topologies that can also be adapted to a DR scenario, see Cloud VPN topologies (/vpn/docs/concepts/topologies).

A key consideration when setting up your connection to Google Cloud from an off-platform site is bandwidth. The connection needs to adequately support the regular transfer of data and backups to Google Cloud.

For more information about your options for connecting to Google Cloud from off-platform sites, see Hybrid connectivity products (/hybrid-connectivity/).

If your primary system is running on Google Cloud, maintaining the connectivity between the SAP applications, databases, and servers at the DR site is a relatively simple matter of modeling the host names, subnets, firewalls, and so forth on the primary site.

If your primary system is not running on Google Cloud, more effort might be required in the design phase to translate the networking architecture of the primary site to the DR site on Google Cloud.

Testing your DR procedures is critical to identifyand address any connectivity issues before your business needs to depend on the recovered system.

In a recovery, after the SAP system is restored to the DR site, the traffic from your users needs to be rerouted to the recovered system. Typically, you do this by updating the network aliases in the DNS entries to the new IP addresses, which are regional.

If your networking architecture uses VPC routes, you will need to update the routes during recovery.

On Google Cloud, you can use Cloud DNS (/dns/docs/overview), or you can use another DNS solution.

If your primary system is running on Google Cloud and you are using regional networking resources like Cloud NAT or a regional Cloud Load Balancing, you need an instance of the resource in each region.

For more information:

- Load balancing overview (/load-balancing/docs/load-balancing-overview)
- Cloud NAT (/nat/docs/overview)

Make sure that the same ports are open at the DR site as are open at primary site.

You can define firewall rules (/vpc/docs/firewalls) in your VPC network to control traffic to and from your VMs.

If you have an Active Directory service on Windows Server, you need to set it up before recovery and keep it in sync with the Active Directory instance at the primary site.

You need to implement the same security controls and permissions that you have at you primary site at your DR site. The same compliance regulations apply to your recovered environment, as well.

Any user or role that requires access at the primary site requires access at the DR site as well.

For general information about designing security for a DR solution on Google Cloud, see Implementing security and compliance controls (/solutions/dr-scenarios-planning-guide#implementing_security_and_compliance_controls).

You can speed the deployment of your Compute Engine virtual machine (/compute/docs/instances/) (VM) instances and avoid configuration errors at your DR site by using Google Cloud products and features like Cloud Deployment Manager, instance templates, and custom images.

Google Cloud provides Deployment Manager (/deployment-manager/) templates for SAP on GCP that you can use to pre-define and deploy your SAP system at the DR site. Using the Deployment Manager templates speeds deployment, reduces configuration errors, and ensures that your SAP systems meet SAP support requirements.

Another option for configuring your VMs in advance are Compute Engine instance templates. Using instance templates is a way to speed deployment and reduce manual configuration of your VMs during recovery. However, because they do require some reconfiguration for the DR

site, recovering your VMs from a recovery boot disk, as described in the next section, might be easier.

For more information about instance templates, see, Instance templates (/compute/docs/instance-templates/).

You can also use deployment orchestration tools, such as Terraform (https://www.hashicorp.com/integrations/google-cloud), to manage you infrastructure deployments on Google Cloud.

Depending on your RTO, you can pre-deploy your Compute Engine instances or, because deploying a VM takes only minutes, you can deploy at the time of recovery.

If you pre-deploy your VMs, you can stop them to save costs or you can use them for other non-essential purposes until they are needed for recovery.

You can also minimize cost by consolidating a distributed system on fewer VMs at the recovery site. For example, if the application servers at the primary site are on dedicated hosts at the primary site, at the DR site, you might put the application servers on the same host as the SAP Central Services. However, you need to weigh the cost savings against the increased complexity of having different configurations at each site.

You can create a custom image (/compute/docs/images#custom_images) from the boot disk of the host at your primary system and then use the custom image to create the recovery instance at the DR site.

If your system is running on Google Cloud, create a custom image if you created and modified a Compute Engine persistent boot disk for your primary system. If you are using an unmodified Google Cloud public image, you can use the Google Cloud public image at the recovery site also. For more information, see Creating, deleting, and deprecating custom images (/compute/docs/images/create-delete-deprecate-private-images).

If your system is not running on Google Cloud, you can import a boot disk image (/compute/docs/images/import-existing-image) to Google Cloud from your on-prem environment, or import virtual disks (/compute/docs/images/importing-virtual-disks) from VMs that are running on your local workstation or on another cloud platform.

Google Cloud provides a number of different backup options that you can choose from when you design a DR solution:

- Compute Engine custom images
- Compute Engine persistent disk snapshots
- Replication

To back up the boot disk of the primary system for use in recovery, you can store it on Google Cloud as a Compute Engine custom image. For more information, see Recovery boot disk (#recovery_boot_disk).

To back up SAP or other file systems that are on a Compute Engine persistent disk, you can use persistent disk snapshots (/compute/docs/disks/create-snapshots).

You can also define a snapshot schedule to take snapshots automatically at regular intervals. See Creating scheduled snapshots for persistent disk (/compute/docs/disks/scheduled-snapshots).

Snapshots provide only block-level consistency. To ensure file-level consistency, you might need to implement mechanisms to suspend write activity on the target file systems during these snapshots.

Depending on your shared storage solution and recovery point objectives, you can replicate your file systems. Replication can be used for databases, block-level storage, or files.

Designing a DR solution that uses replication is best for business critical applications that have a very low tolerance for data loss.

If your primary system is running on Google Cloud, multi-region buckets and multi-region snapshots are replicated for you between the selected regions.

You can also use replication provided by third storage solutions.

When you design a DR solution on Google Cloud, you are likely to use multiple types of storage, depending on where your primary system is running and what you are storing.

For backups other than disk snapshots, such as files that you upload from an SAP system that is not running on Google Cloud, create a bucket in Cloud Storage (/storage/) that you can access from both the primary and DR sites. When you create a bucket, you choose a default storage class and a location.

Select a default storage class based on the service level agreement (SLA) you require, your expected usage of the storage, and your cost constraints. For DR, the coldline storage class is often a good option.

For you bucket location, choose a location that includes the region of your DR site and, if your primary system is running on Google Cloud, the region of your primary site.

If your primary system is on Google Cloud, choose a multi-region location that includes the regions of both the primary and DR sites so that you can access the bucket from both sites.

If your primary system is not on Google Cloud, to save costs you can select a single-region location in the region that includes your DR site.

If your primary system is on Google Cloud and you are using a third-party solution for shared storage, your storage options might be determined by the solution. Refer to the solution documentation or your Google Cloud support representative.

For information about pricing, see Cloud Storage pricing (/storage/pricing).

When you create a snapshot, you can specify a storage location. The location of a snapshot affects its availability (/storage/sla) and can incur networking costs (/storage/pricing#network-pricing) when you create the snapshot or restore it to a new disk.

Snapshots can be stored in either one Cloud Storage multi-regional location (/storage/docs/bucket-locations#location-mr), such as `asia`, or one Cloud Storage regional location (/storage/docs/bucket-locations#location-r), such as `asia-south1`.

A multi-regional storage location provides higher availability and might reduce network costs when creating or restoring a snapshot. A regional storage location gives you more control over the physical location of your data.

Regardless of the location options you choose, the snapshots must be stored in a location that is accessible from your DR site.

For more information about snapshot locations, see Selecting the storage location for a snapshot (/compute/docs/disks/create-snapshots#selecting_a_storage_location).

For pricing information for snapshot storage, see Persistent disk pricing (/compute/pricing#persistentdisk).

After you add custom image files to your custom images list, Compute Engine manages the storage for the images. Images in a custom images list are global resources, so are available in any region.

For information about image storage pricing, see Image storage (/compute/pricing#imagestorage).

After your DR plan is complete, test it regularly, noting any issues that come up and adjusting your plan accordingly.

Be sure to test all aspects of your DR plan, including things like:

- Backups and backup schedules
- Data transfer from off-platform sites
- Recovery from the stored backups
- Security controls and system access

- Network routing

When you test your DR systems, your primary systems will continue to run. To prevent conflicts or split brain scenarios, you need to isolate the test system from the primary system and its users.

For general information about DR testing on Google Cloud, see Disaster Recovery Planning Guide (/solutions/dr-scenarios-planning-guide).

Certain Google Cloud products, functions, and services are key to designing a DR solution that meets your RPOs and RTOs.

When designing a DR solution on Google Cloud to meet a particular RPO, there are two variables that control the point in time you can recover to:

- Backup frequency
- Backup retention

Backup frequency determines the maximum amount of time that can elapse between your last backup and a disaster. If you create your DR backups every 24 hours, you could potentially lose almost 24 hours worth of data updates if disaster strikes 23 hours and 59 minutes after your last backup was taken.

Replication can reduce the maximum amount of elapsed time since your last backup to near zero; however, replication is expensive, so you might use it only for databases and business critical application files.

In a DR solution, point-in-time copies or snapshots are commonly used to backup the SAP application file systems that are required for recovery.

On Google Cloud, you can automate Compute Engine persistent disk snapshots by creating an hourly, daily, or weekly snapshot schedule. However, because the Compute Engine snapshots control for consistency only at the block level, consider suspending write activity on the target file systems during snapshots to ensure file-level consistency.

The primary cost to consider when choosing a backup frequency is the cost of data transfer. Storage cost is also a consideration, but your backup retention policy might have a greater impact on storage costs than your backup frequency.

For more information about snapshot schedules, see Creating scheduled snapshots for persistent disk (/compute/docs/disks/scheduled-snapshots).

Backup retention determines how far back in time you can move your recovery point. Retaining backup copies helps to protect against logical errors by allowing you to recover to a point in time before the logical error exists.

You can set retention policies for Compute Engine snapshots and Cloud Storage buckets that automatically delete your backup files after a specified amount of time.

The primary cost to consider when choosing a retention policy is the cost of storage. Compute Engine snapshots reduce the amount of storage that is required for multiple snapshots by storing only incremental block-level changes after the first complete snapshot is stored.

For more information about defining retention policies for snapshots, see Snapshot retention policy (/compute/docs/disks/scheduled-snapshots#retention_policy).

For information about setting retention policies on Cloud Storage buckets, see Retention policies using Bucket Lock (/storage/docs/bucket-lock).

When designing your Google Cloud DR solution to meet a particular RTO, the readiness of the infrastructure, software, file systems, and data at the DR site is the primary controlling variable.

For example, to achieve a very low RTO, you might maintain a *hot standby* system at the DR site, with pre-deployed infrastructure, active SAP systems, and replicated data. However, the low RTO comes at a higher cost.

You can balance costs and recovery times by setting up some low- or no-cost infrastructure in advance, and deferring some set up steps to the time of recovery.

For example, you can configure and deploy a VM in advance, but then stop the VM. The resources attached to the VM, such as any external static IPs or persistent disks, might still incur charges, but the stopped VM itself does not.

Because you can configure and deploy a Compute Engine VM on Google Cloud relatively quickly, you might be able to do it at the time of recovery and still meet your RTO, especially if you use Deployment Manager to deploy your DR system or create the VM from a template and a custom image you create in advance.

If you do not pre-deploy your infrastructure and systems, check your resource quotas in the region of the DR site periodically to ensure that the quotas are sufficient to deploy the DR system.

Pre-deploying as much of the DR infrastructure and systems as your budget allows can help ensure that your DR system fits within your quotas and that the GCP resources that your system needs are available when your system needs them.

The following architecture diagrams show examples of DR designs for different RTOs.

Each of the diagrams show multi-region backup options on the left and a corresponding simplified SAP configuration at the DR site on the right.

Each example shows one possible combination of DR design elements. To adjust your DR design to your objectives and circumstances, you might combine elements from any or all of the examples.

The following architecture diagram shows a low-RTO DR design example.

In this design, you maintain a nearly functional SAP system at the DR site. The VMs are deployed and active. The SAP application servers and the database server are active, but the

application service is stopped.

The backup options you are likely to use in this scenario include OS images stored in and persistent disk snapshots stored in Compute Engine, and data replication between the primary and DR sites.

Because data replication is used, the risk of data loss is limited to the last database synchronization.

To extend your RPO back in time, you can add data backups to Cloud Storage, which is not shown in the diagram. For the persistent disk snapshots, you can use scheduled snapshots and adjust the retention policy to your RPO.

The actions you need to take to recover a low-RTO system like this include:

- If required, performing a final synch of the files or recovering the files from a persistent disk snapshot.
- Switching the primary database to the DR site.
- Starting the application at the DR site.

Of the three example architectures shown, the low-RTO example is the most costly.

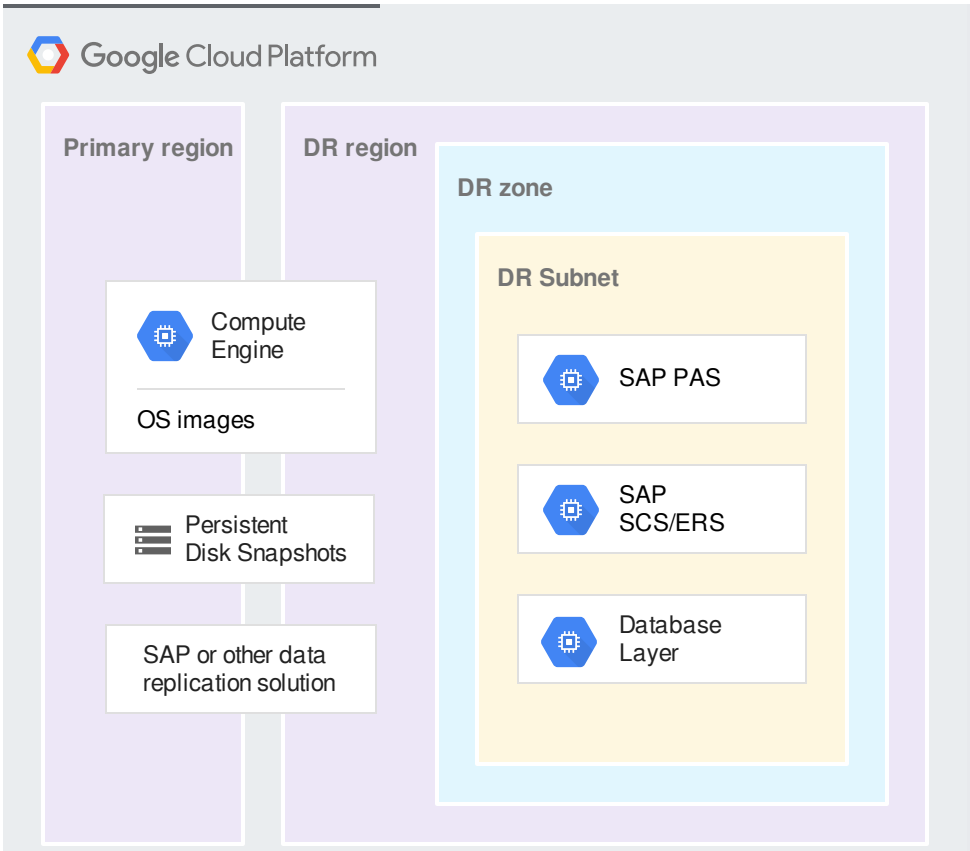

The following example DR architecture has a higher RTO than the preceding example, but a lower cost.

In this design, the database server is active to support replication from the primary site. The VMs for the application servers are not active.

The backup options you are likely to use in this scenario include OS images and persistent disk snapshots stored in Compute Engine, and data replication between the primary and DR sites.

Because data replication is used, the risk of data loss is limited to the last database synchronization.

To extend your RPO back in time, you can store data backups in a Cloud Storage bucket, which is not shown in the diagram. For the persistent disk snapshots, you can use scheduled snapshots and adjust the retention policy to your RPO.

Depending on the requirements of your database system, you might be able to deploy your database on a smaller VM at the DR site than is required at the primary site, which can help

reduce costs until the DR system is activated. In this case, you need to resize the VM to the required size during recovery.

The actions you need to take to recover a system like this include:

- If necessary, deploying the VM instances for SAP NetWeaver and the application servers from persistent disk snapshots or images.
- Synchronizing the file systems from persistent disk snapshots or other storage.
- If necessary, resizing the database VM instance.
- Switching the primary database to the DR site.
- Starting the application service at the DR site.

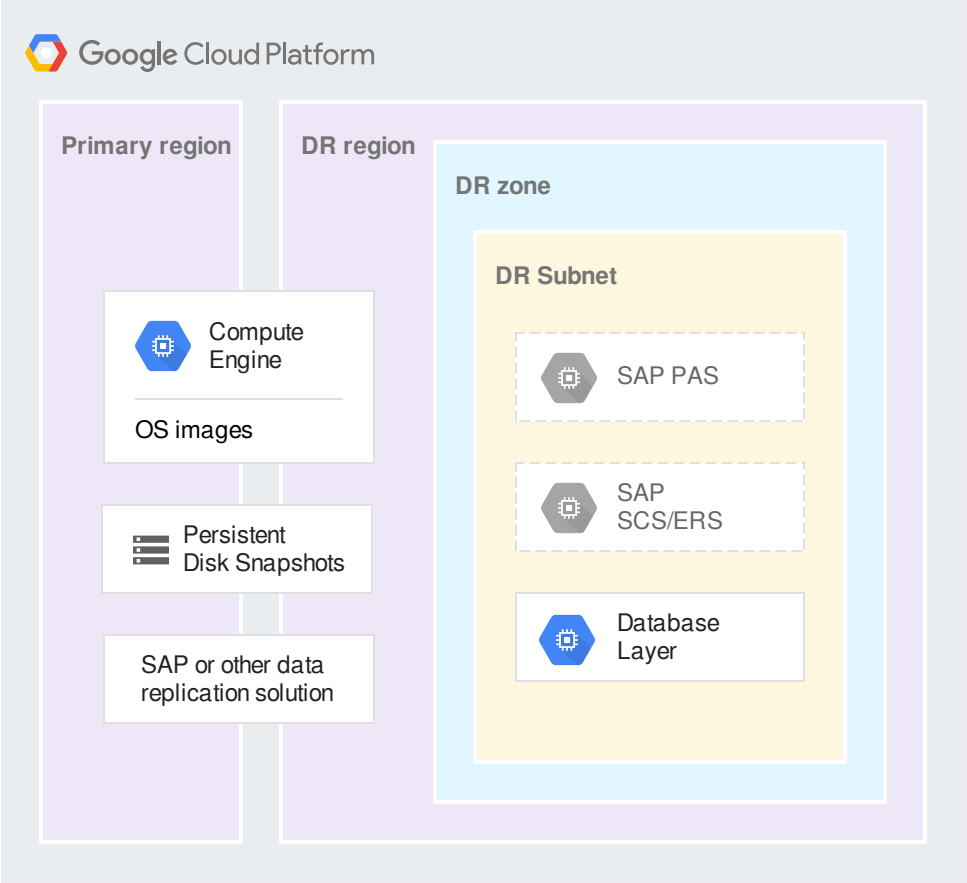

The following architecture diagram has the highest RTO of the examples shown and is the lowest-cost option.

In this design, you can pre-deploy the servers and then stop them to avoid incurring charges for the VMs or, to avoid costs associated with persistent disks and other VM-related resources, you can deploy the VMs as part of your recovery process,.

The backup options you are likely to use in this scenario include OS images stored and persistent disk snapshots that are stored in Compute Engine, and data backups that are stored in Cloud Storage.

To meet your RPO, you can adjust both your backup frequency and the retention policies of your snapshot schedules and of the backups in the Cloud Storage bucket.

The actions you need to take to recover a system like this include:

- If necessary, deploying the VM instances for SAP NetWeaver, the application servers, and the database server from multi-regional persistent disk snapshots or images that are stored in Compute Engine.
- Synchronizing the file systems from persistent disk snapshots or other storage.
- Recovering the database from backup files in multi-regional Cloud Storage bucket or elsewhere.
- Switching the primary database to the DR site.
- Starting the application service at the DR site.

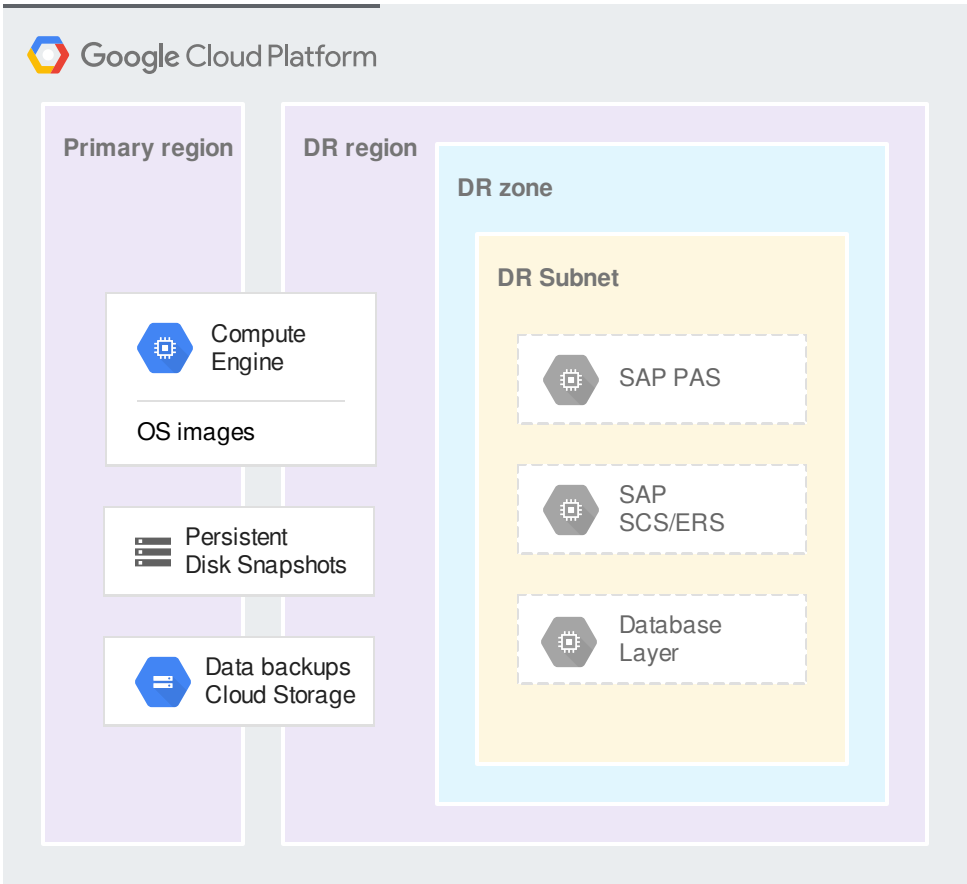

You can find partners to help with your DR solution in the Google Cloud Partner Directory (https://cloud.withgoogle.com/partners/).

You can also find various products and partners to support your DR solution on Google Cloud Marketplace (https://console.cloud.google.com/marketplace/browse?q=%22Disaster%20recovery%22).