

# **HPE Remote Copy Active Sync**

A new synchronous replication mode supported on HPE Alletra 9000 and HPE Primera Storage arrays



# Contents

Executive summary.....	3
Failure domains.....	3
Configuring an Active Sync solution.....	5
Failure recovery with Active Peer Persistence versus Active Sync.....	6
Preparing the environment for Active Sync.....	8
Setting the group for Active Sync.....	9
Host array proximity.....	9
Configuring Active Sync through Data Services Cloud Console.....	9
Configuring Active Sync through the onboard UI.....	11
Configuring Active Sync through the SSMC.....	12
Enabling Active Sync through the CLI.....	14
Best practices for Active Sync.....	14
Summary.....	15
Resources.....	16



## Executive summary

In May 2021, Hewlett Packard Enterprise classic Peer Persistence achieved a new level of increased capability and manageability with the release of Active Peer Persistence for the new HPE Alletra 9000 Storage arrays running HPE Alletra OS 9.3 and for HPE Primera arrays running OS 4.3. In classic Peer Persistence, only the primary array in a Peer Persistence solution services I/O for virtual volumes (VVs) in a Remote Copy (RC) group that is configured for synchronous replication. Active Peer Persistence, in contrast, allows both the primary and the secondary array to service application I/O requests for VVs in a synchronously replicated RC group. Because Active Peer Persistence allows I/O to both the primary and the secondary array for an RC group, it offers an improved solution that enables load balancing of server I/O across both arrays for an application's VVs. It also supports shared VV cluster solutions that need extreme availability by allowing the servers in the cluster to send I/O to both arrays. An Active Peer Persistence configuration requires the servers in the solution to have host connectivity to both arrays. This connectivity enables the automatic transparent failover of RC groups from one array to the other without impact to the applications running on the servers that use the RC groups. (For more information about automatic recovery from failures, refer to the [HPE Active Peer Persistence technical white paper](#).)

Classic Peer Persistence and Active Peer Persistence are outstanding products that align with most customer environments for host connectivity. Some environments, however, might require the ability to read and write both arrays in a synchronous replication configuration, yet not meet the Peer Persistence requirement for servers to have connectivity to both arrays in the replication solution. VMware® refers to this kind of configuration as a “non-uniform” server connect configuration because the VMware ESXi™ servers in a VMware vSphere® Metro Storage Cluster (vMSC) are not connected to both arrays; they are connected non-uniformly to the array in the local data center only. With the release of the 9.4 OS for HPE Alletra 9000 arrays and the 4.3 OS for HPE Primera arrays, HPE supports non-uniform host connectivity for RC groups that are configured for synchronous replication. HPE calls this an **Active Synchronous** or **Active Sync** solution. In an Active Sync solution, data is replicated synchronously between two arrays in a Remote Copy configuration with both the primary and the secondary array able to service I/O for the VVs in an RC group (reads and writes), but with the servers using the VVs in the Active Sync RC group connected only to their local array and hence able to send I/O requests only to their local array (non-uniform server connect).

With Active Sync, VMs running on a vMSC stretched cluster can be migrated seamlessly between the ESXi servers in separate data centers while still maintaining access to all of their data. This ability to migrate VMs between any ESXi servers in the vMSC cluster significantly simplifies administration of the environment by removing the need to fail over data stores from one data center to another before migrating VMs between the data centers. In the event of an array failure or a replication link failure between arrays, the vMSC cluster fails over any affected VMs to an ESXi server that has access to the array servicing the VVs required by the VMs.

### Target audience

The audience for this paper includes HPE and partner presales solution architects, post-sales support engineers, customers interested in Active Sync capability, and engineers who want to understand the concepts involved.

### Document purpose

This paper discusses in detail the deployment of a Remote Copy Active Sync solution, including Active Sync operation, configuration recommendations, and best practices for its use. The integration of a host-level vMSC stretched cluster is integral to an Active Sync solution, and this paper explains the responsibilities of host-level cluster software in the solution.

Although the examples in this paper refer to VMware, and in particular to VMware vMSC clusters, other operating systems and cluster solutions will be supported in the future. For more information about server and cluster support with Remote Copy Active Sync, refer to the [HPE Single Point of Configuration Knowledge \(SPOCK\)](#).

## Failure domains

When designing a solution based on Active Sync replication, you should consider the **failure domain** being protected — that is, the geographic impact of a failure or disaster on the data centers and arrays. Its use of synchronous replication enables Active Sync to provide the foundation for achieving a recovery point objective (RPO) of zero (no data loss), but it is generally used only for failure domains of limited size. Because Active Sync uses synchronous replication, as the failure domain increases in size, so does the impact on host write I/O latency as the distance between the arrays increases (that is, the round-trip latency imposed by distance). Active Sync is appropriate for four failure domain types:

- Data center
- Campus
- Metro
- Geo



Power failures, floods, hurricanes, and tornadoes are examples of disasters that might take down both data centers on a common campus, so a data center or campus-based solution does not protect from these failures. To protect a larger failure domain, you must place the data centers containing the arrays in the solution far enough apart that a failure cannot affect both data centers simultaneously. The larger the failure domain, the less the chance of a disaster affecting both sites simultaneously. However, there is a practical limit to the size of a failure domain that can be protected with synchronous replication because of the effects that distance places on write I/O latency.

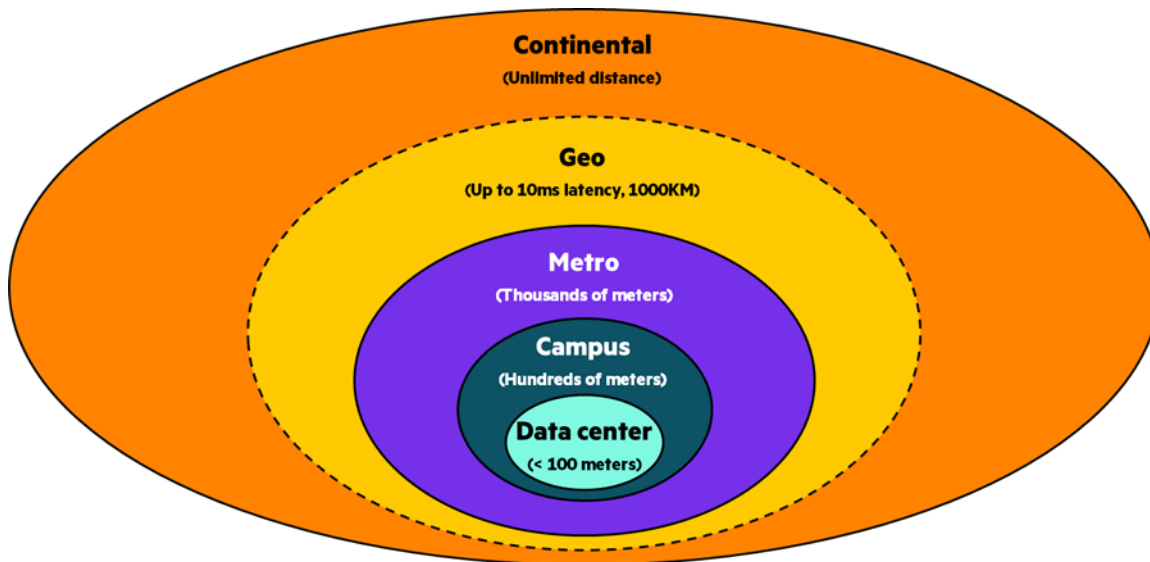


Figure 1. Failure domain size comparison

Figure 1 shows the failure domains that a Remote Copy Active Sync solution can protect. In this graphic, the larger the oval, the larger the failure domain. The maximum failure domain that can be protected by Active Sync replication on HPE Alletra 9000 and HPE Primera Storage arrays is a geo domain (the dotted oval), which encompasses data center, campus, and metro failure domains.

The maximum practical size of a failure domain is dictated by the need to replicate write data synchronously between the arrays configured in the solution and by the effect that the replication latency has on server write I/O latency. The absolute maximum latency supported by Remote Copy for synchronous replication is 10 ms round trip. That is 1000 km (or approximately 625 miles) over a dark fiber infrastructure and much less over a switched network. Most customer solutions cannot tolerate application write I/O latency greater than a millisecond or two. This requirement forces most customers to deploy a solution for a failure domain no larger than about 100 km to 200 km in size — significantly smaller than the 1000 km maximum that is supported.

A good rule of thumb for determining the maximum failure domain is 1 ms round trip of link latency per 100 km of dark fiber. For replication over a switched network (WAN), there is no correlation between distance and latency, so you must measure the specific latency over the switched network.



### Configuring an Active Sync solution

The process of deploying and configuring an Active Sync solution on HPE Primera or HPE Alletra 9000 Storage arrays is almost identical to that of deploying and configuring an Active Peer Persistence solution — with a few minor differences. The primary difference between the two solutions regards server cabling. With an Active Peer Persistence solution, the servers in the vMSC stretched cluster have I/O connections to both the primary and the secondary array in the Remote Copy configuration (Figure 3). With an Active Sync solution, the servers making up the stretched cluster are connected only to the array in their local data center (Figure 2).

These are key differences between Active Peer Persistence and Active Sync configurations:

- With Active Sync, the ESXi servers in the solution are cabled only to their local array, not to the other array in the configuration.
- With Active Sync, the ESXi servers in the solution should boot locally or, if they boot from SAN, their boot disks should not be replicated.

Hewlett Packard Enterprise highly recommends that you deploy an Active Sync configuration only when it is impossible to deploy a full Active Peer Persistence solution. If the servers can be connected to both arrays in the Remote Copy configuration, then a full Active Peer Persistence solution complete with a quorum witness should be deployed because an Active Peer Persistence solution provides far better availability than an Active Sync solution can for all array failures.

Figures 2 and 3 show the difference in server cabling between an Active Sync solution with non-uniform host connectivity compared to an Active Peer Persistence solution with uniform host connectivity. Notice that for an Active Sync solution, the servers in the stretched cluster are cabled only to the array in the local data center, whereas for the Active Peer Persistence solution the servers in the stretched cluster are connected to the arrays in both data centers.

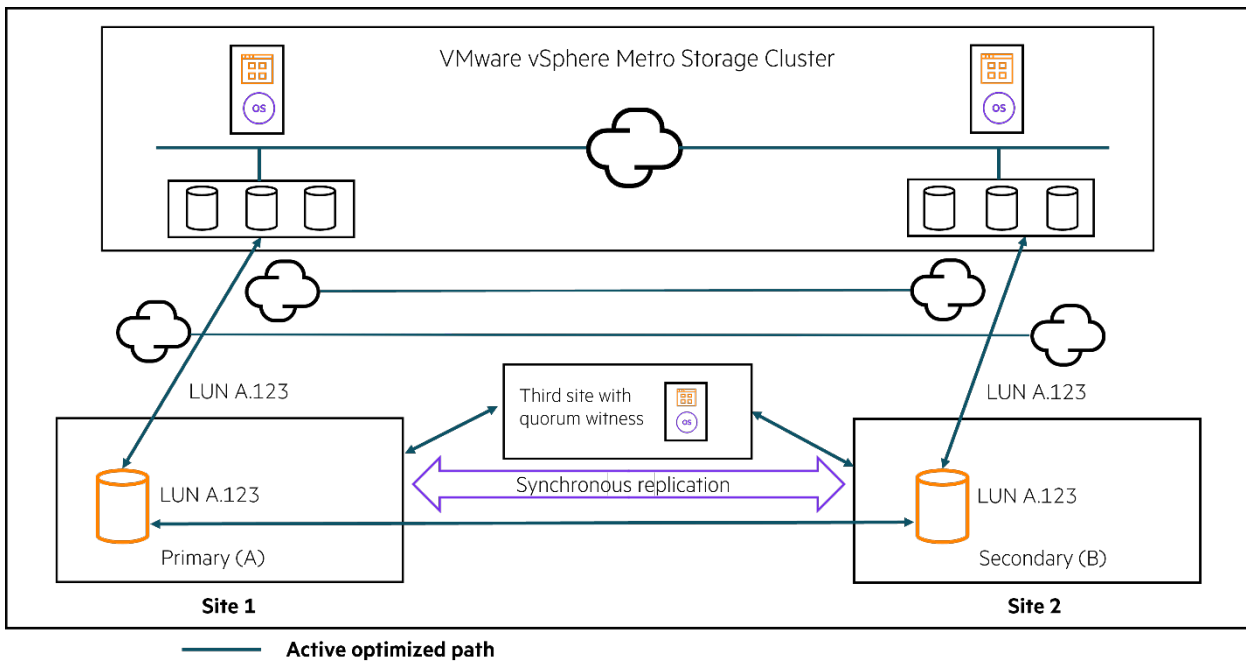


Figure 2. vMSC stretched cluster with Active Sync showing non-uniform server cabling



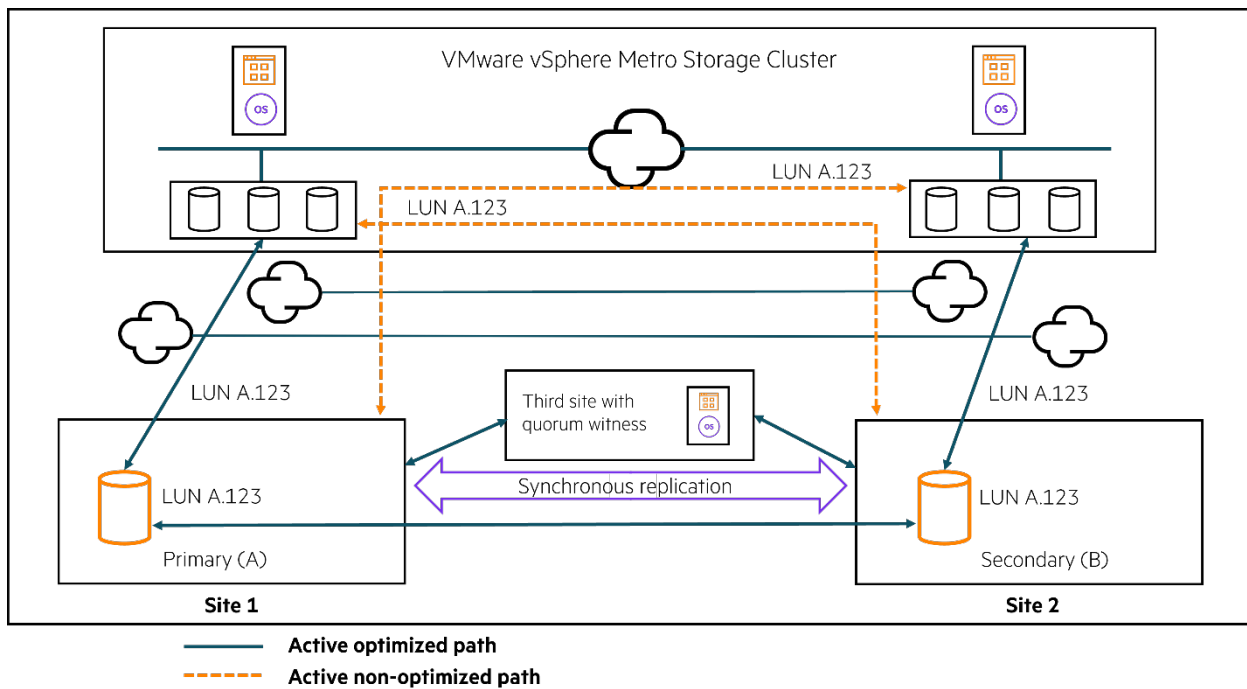


Figure 3. vMSC stretched cluster with Active Peer Persistence showing uniform server cabling

### Failure recovery with Active Peer Persistence versus Active Sync

Its use of the non-uniform host connectivity configuration makes the Active Sync solution much more susceptible to an application outage than a full Active Peer Persistence solution in many failure scenarios. Having the servers in the stretched cluster connected only to the array in the local data center results in the cluster servers losing access to their data for any situation that prevents the local array from being able to access data, including administrative tasks such as stopping the RC group. In these situations, the VMs and applications running on the ESXi servers that lose data access must be restarted by the vMSC cluster on ESXi servers that continue to have storage access to the other array in the solution. This means that an RPO of zero is achieved for all application data in an Active Sync solution because synchronous replication is used, but the recovery time objective (RTO) is not zero for any VMs that must be restarted. In all of the following failure scenarios, the RTO and RPO would be zero if an Active Peer Persistence solution had been deployed.

#### Active Peer Persistence failure scenarios and recovery

The following list explains how an Active Peer Persistence solution responds to and recovers from different failure scenarios.

- Primary array failure:** In the event of a complete failure of the primary array, the Active Peer Persistence solution automatically and transparently fails over all RC groups from the primary array to the secondary array. Host I/O processing and applications continue uninterrupted, using the ESXi server's paths to the secondary array following the automatic transparent failover of the RC groups. RPO=0 and RTO=0.
- Secondary array failure:** In the event of a complete failure of the secondary array, the ESXi servers' paths to the secondary array are lost and host-based multi-pathing software automatically starts to use the ALUA "active non-optimized" paths to the primary array. Host I/O processing and applications continue uninterrupted. RPO=0 and RTO=0.
- Replication link failure:** If all replication links between the primary and secondary arrays fail, the host paths to the secondary array transition to ALUA "unavailable," and hosts can no longer send I/O to the secondary array. The host starts to use its ALUA active non-optimized path to the primary array, and application processing continues uninterrupted. RPO=0 and RTO=0.
- RAID set failure or metadata failure:** The paths to an array that suffers a RAID set or metadata failure will transition to ALUA "unavailable." If the failure happens on the primary array, an Automatic Transparent Failover to the secondary array occurs. The servers in the solution then use their paths to the healthy array in the Active Peer Persistence configuration, and application processing continues uninterrupted. RPO=0 and RTO=0.
- Host disconnect (host loses all its active optimized paths to an array):** The servers use their active non-optimized paths to the other array in the Active Peer Persistence configuration, and application processing continues uninterrupted. RPO=0 and RTO=0.



- **Primary array reboot/shutdown:** In the event of a primary array reboot or shutdown, the secondary array in the Active Peer Persistence solution automatically and transparently fails over all primary RC groups. Host I/O processing and applications continue uninterrupted following the automatic transparent failover of the RC groups. RPO=0 and RTO=0.
- **Cluster server failure:** If an ESXi server in the vMSC stretched cluster fails, the virtual machines running on that cluster server and the associated applications also fail. Recovery occurs at the cluster layer with the vMSC cluster restarting the failed virtual machines on other ESXi servers in the cluster. RPO=0, but RTO <> 0 for the failed VMs.
- **Data center failure:** If the entire data center fails, both the array in that data center and all ESXi clustered servers in that data center also fail. Active Peer Persistence performs an automatic failover of the RC groups for the failed array, but all of the VMs and associated applications running on the ESXi servers in the failed data center must be recovered by the stretched cluster with the associated impact on application RTO. RPO=0 but RTO <> 0 for failed the VMs.

As this list shows, an Active Peer Persistence solution recovers transparently (RPO=0 and RTO=0) for any failure that does not affect the servers in the stretched cluster solution.

### Active Sync failure scenarios and recovery

Failure scenario recovery differs significantly between an Active Sync solution and an Active Peer Persistence solution. Active Sync, like Active Peer Persistence, is synchronous replication in which both the primary and secondary array for an RC group can accept read and write requests for VVs in the RC group. It can perform an Automatic Failover of the RC group from the primary array to the secondary array if the primary array fails, a quorum witness is configured for the group. But in the case of an array failure (either the primary or the secondary array), all VMs running on ESXi servers connected to the failed array will suffer an outage and must be recovered by the vMSC cluster. Recovery of the failed VMs by the vMSC cluster is not transparent.

The following list explains how an Active Sync-based solution responds to different failure scenarios:

- **Primary array failure:** If the primary array fails, Active Sync executes an automatic failover of the RC group if the `auto_failover` policy is set on the group and there is an operating quorum witness. This causes the secondary array to become primary for the RC groups. The servers in the vMSC cluster connected to the secondary array at the time of the failure have no interruption to their storage, and the VMs and applications running on those servers continue to run. The VMs running on the servers in vMSC cluster connected to the failed primary array are failed over to the secondary site by the vMSC cluster and are restarted — with a resulting impact on application RTO.
- **Secondary array failure:** The servers in the ESXi vMSC cluster connected to the primary array are not interrupted, and the VMs and applications running on those servers continue to run. The VMs running on the servers in the vMSC cluster connected to the failed secondary array fail, and the vMSC cluster automatically fails them over to cluster servers in the primary array site — with a resulting impact on application RTO.
- **Replication link failure:** If all replication links between the arrays fail, the ALUA paths to the secondary array transition from ALUA “active optimized” to ALUA “unavailable.” The VMs running on the servers in the vMSC cluster connected to the secondary array fail, and the vMSC cluster does not automatically fail them over to cluster servers in the primary array site. Recovering these VMs and their associated applications is a manual process.
- **RAID set failure or metadata failure:** Paths to the array suffering the RAID set failure or metadata failure transition to ALUA “unavailable”; the RC group stops; and the cluster servers connected to that array lose access to their data. If the primary array suffers the failure, an automatic failover from the primary to the secondary array occurs. The VMs and applications running on the cluster servers that lose storage connectivity fail and are recovered at the stretched cluster layer.
- **Host disconnect:** If a server loses access to its VVs because it is disconnected from the array, all host I/O stops, and the associated VMs and applications running on the server fail. These failed VMs must be restarted on a connected host in the cluster.
- **Primary array reboot/shutdown:** If the primary array is shut down or rebooted, an automatic failover occurs, and the secondary array becomes primary for the RC groups. The servers in the ESXi vMSC cluster connected to the secondary array are not interrupted, and the VMs and applications running on those servers continue to run. The VMs running on the servers in the ESXi vMSC cluster connected to the rebooted/shut down primary array are failed over to the secondary site by VMware HA clustering and restarted with a resulting impact on application RTO.
- **Secondary array reboot/shutdown:** The servers in the ESXi vMSC cluster connected to the primary array have no interruption, and the VMs and applications running on those servers continue to run. The VMs running on the servers in the ESXi vMSC cluster connected to the secondary array fail, and the vMSC cluster does not automatically fail them over to cluster servers in the primary array site. Recovering these VMs and their associated applications is a manual process.



- **Cluster server failure:** If a server in the stretch vMSC cluster fails, the VMs and associated applications running on that cluster server also fail. Recovery occurs at the stretched cluster layer with an RPO of zero but with an impact on application RTO.
- **Data center failure:** If an entire data center fails, both the array in that data center and all cluster servers in that data center fail. Active Sync performs an automatic failover for the RC groups that are primary on the failed array, but all of the VMs and associated applications running on them must be recovered by the stretched cluster, with the associated impact on application RTO.

In all of these scenarios, except for the cases in which servers in the stretched cluster failed, if an Active Peer Persistence solution had been deployed, the failures would have been recovered from transparently with no application outage. This is why Hewlett Packard Enterprise recommends using Active Peer Persistence with uniform host connectivity configurations over an Active Sync solution when possible.

---

**Note**

Best practice is to always deploy a uniform host connectivity configuration with an Active Peer Persistence-based solution when possible.

---

## Preparing the environment for Active Sync

As explained previously, an Active Sync solution has the same basic configuration requirements as an Active Peer Persistence solution, with a few minor differences. With Active Peer Persistence, all servers in the solution must be cabled to both arrays, but with an Active Sync solution, the servers in the stretched cluster are cabled only to the array in the local data center and are not cross-connected to the other array in the solution. Also, with an Active Sync solution, you must ensure that the servers in the stretched cluster either boot locally or boot from SAN using a non-replicated boot volume.

---

**Note**

For more information about how to configure your environment for Active Peer Persistence/Active Sync, refer to the [HPE Active Peer Persistence technical white paper](#).

---

To prepare the environment for Active Sync, complete the following tasks:

- Check [HPE SPOCK](#) to confirm that the hosts, cluster, and operating system versions intended to participate in the Active Sync solution are supported.
- Configure a stretched cluster across the hosts in the two data centers.
- Make sure that hosts in the stretched cluster are connected only to the array in the local data center.
- Follow all vendor best practice recommendations for a stretched cluster.
- Host OS (physical server/hypervisor) should be configured either for local boot or boot from SAN. If boot from SAN is chosen, be sure to boot from a non-replicated boot volume.
- Ensure that you have the same WWN on both arrays for the VVs in the Active Sync RC group.
- The individual host definitions and host sets must exist on both systems. For the hosts that are not attached to a given site, the host definition can be empty (no initiators).
- Export the volumes that will be replicated to all hosts in the stretched cluster (including empty definitions for non-connected hosts, as noted earlier).
- Ensure that hosts and arrays are zoned properly.
- Set up and configure a quorum witness (necessary for automatic failover in the event of an array failure).
- Ensure proper replication link sizing.

---

**Note**

For the latest Active Sync host and operating system support information, refer to [HPE SPOCK](#).

---



## Setting the group for Active Sync

After you have confirmed that your configuration is supported and properly configured for Active Sync, you are ready to enable the Active Sync group. Enabling a group for Active Sync is the same as enabling the group for Active Peer Persistence. You enable Active Sync by applying the `active_active` group policy to the synchronous RC group through the CLI or by selecting the **Active Peer Persistence** option from the High Availability Configuration screen in the HPE StoreServ Management Console (SSMC).

---

### Note

The SSMC cannot be used to manage HPE Alletra 9000 arrays. Data Services Cloud Console, the CLI, or the onboard UI must be used to manage HPE Alletra 9000 arrays.

---

When the Active Peer Persistence policy (SSMC) or the `active_active` policy (CLI) is applied to the synchronous RC group, it immediately becomes an Active Sync RC group. Setting the Active Sync policy on a RC group also enables the following RC group policies by default:

- `Auto_recover`
- `Auto_synchronize`

The `auto_failover` and `path_management` group policies are set for an Active Sync group by default and cannot be modified. The `auto_recover` and `auto_synchronize` options are also set by default for an Active Sync group. These options can be modified (turned off) through the CLI, but it is a best practice to have these two options set. Hewlett Packard Enterprise recommends leaving them set.

---

### Note

Best practice is to leave the `auto_recover` and `auto_synchronize` policies set for an Active Sync RC group.

---

## Host array proximity

When the `active_active` policy is first applied, and the group becomes an Active Sync group, all host I/O paths to the primary array are set to ALUA “active optimized,” and all host I/O paths to the secondary array are set to ALUA “active non-optimized.” After applying the `active_active` policy, you must adjust the host proximity settings for the group to proximity “all” for all the servers using the Active Sync group. DSCC, the onboard UI, SSMC, or the CLI can be used to adjust the host proximity for the servers in the stretched cluster, so they have an ALUA “active optimized” path to the array in their local data center.

For more information about host proximity settings, refer to the [HPE Active Peer Persistence technical white paper](#).

---

### Note

For an Active Sync RC group, the host proximity setting for all servers in the stretched cluster using the RC group must be set to “all.”

---

## Configuring Active Sync through Data Services Cloud Console

Data Services Cloud Console (DSCC) is the comprehensive fleet management hub for end-to-end data and storage lifecycle management. Composed of various integrated applications that encompass storage array setup and management, block and file storage, data protection and more, DSCC offers the most intelligent and comprehensive solution to managing your entire HPE Storage ecosystem.

Using the Block Storage application within DSCC, Active Sync can be applied to existing volume sets on eligible arrays.

To enable Active Sync on an existing volume set, complete the following steps:

1. Navigate to the volume set and then to the Protection Policies tab.
2. Select **Add protection policy** and add a remote protection policy.
3. Select **Synchronous replication** and then select the **Active Sync** checkbox, as shown in Figure 4.



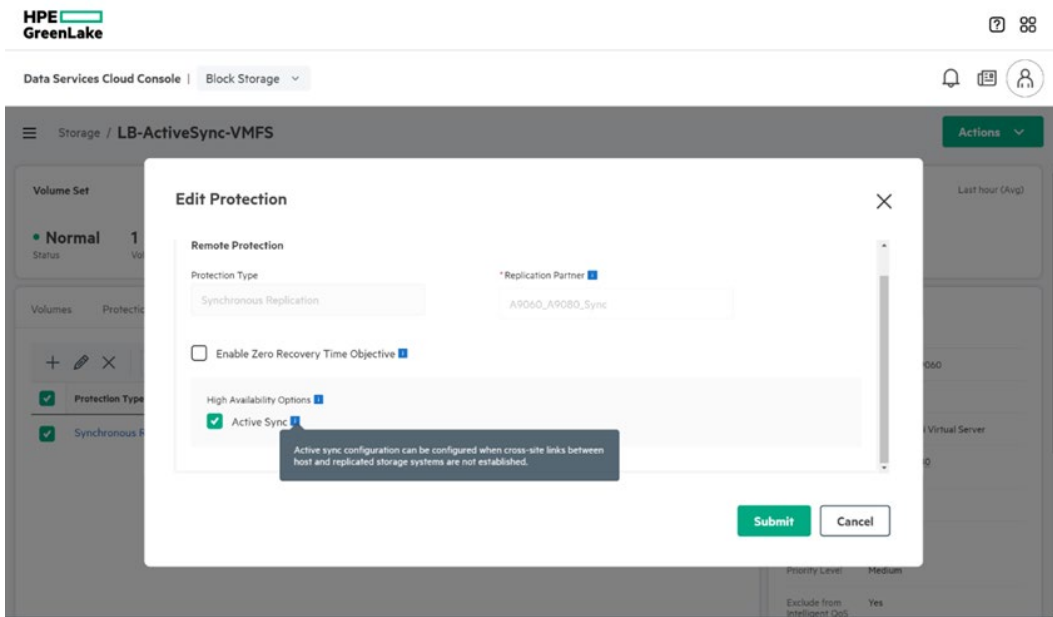


Figure 4. Enabling Active Sync on a volume set in DSCC

The RC group is created automatically, the volumes are admitted, the target volumes are created, and replication starts. The proximity is automatically set to "All," as shown in Figure 5, and it does not need to be modified.



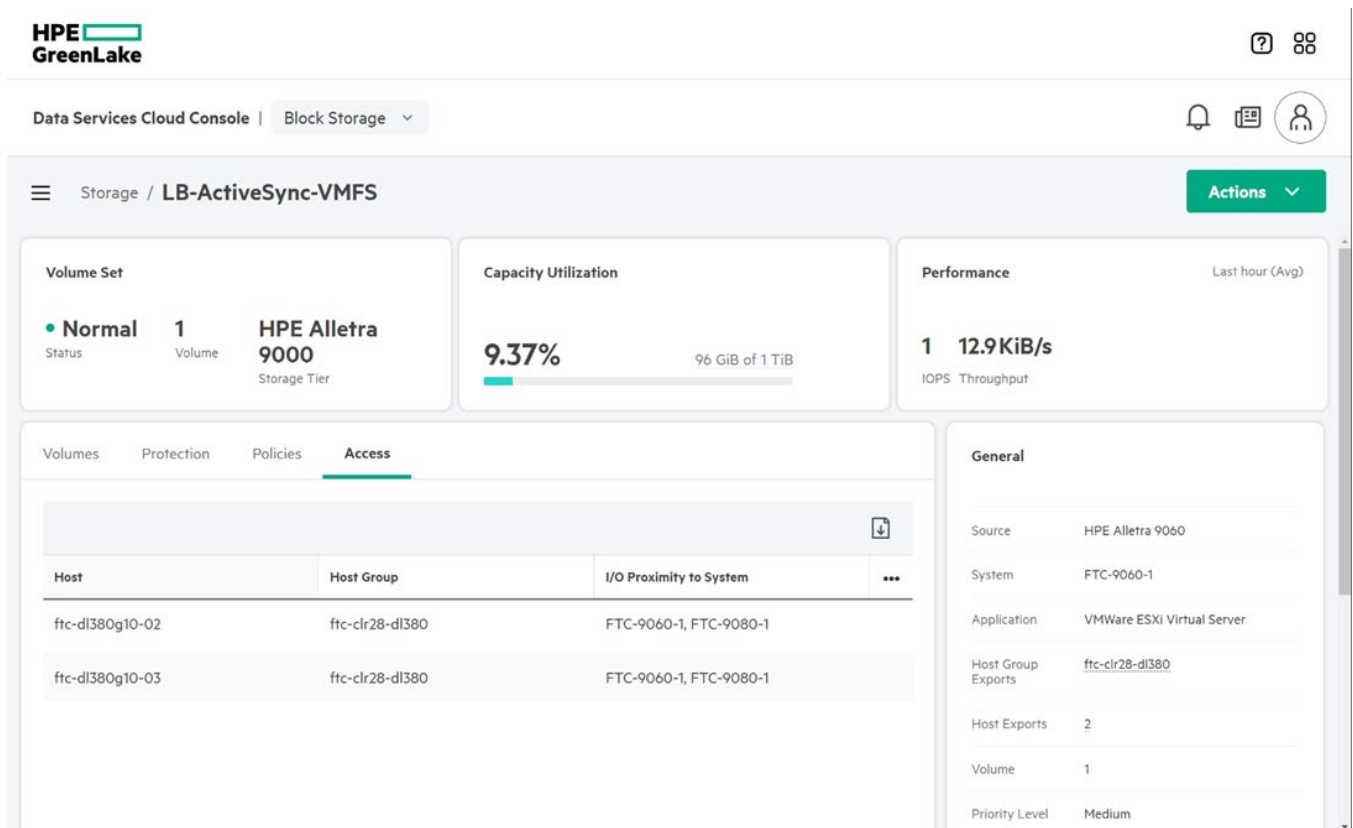


Figure 5. Host proximity setting in DSCC for an Active Sync configuration

### Configuring Active Sync through the onboard UI

For HPE Alletra 9000 and HPE Primera, the onboard UI can be used to configure replication and create Active Sync RC groups, starting with OS 9.5.15 for HPE Alletra 9000 and OS 4.5.15 for HPE Primera.

To configure Active Sync while creating an application set, add a replication policy in the creation wizard. To apply it to an existing application set, navigate to the set, select **Edit**, and add a replication policy. In either case, after applying Active Sync, complete the following steps to configure it:

1. Select a suitable replication partner.
2. Select **Synchronous mode**.
3. Select the **High Availability** checkbox.

Hosts or host sets to which the application set is already exported will populate with the option to select the proximity setting. Select **All** for all hosts that populate.

4. Save the application set with the new replication policy.



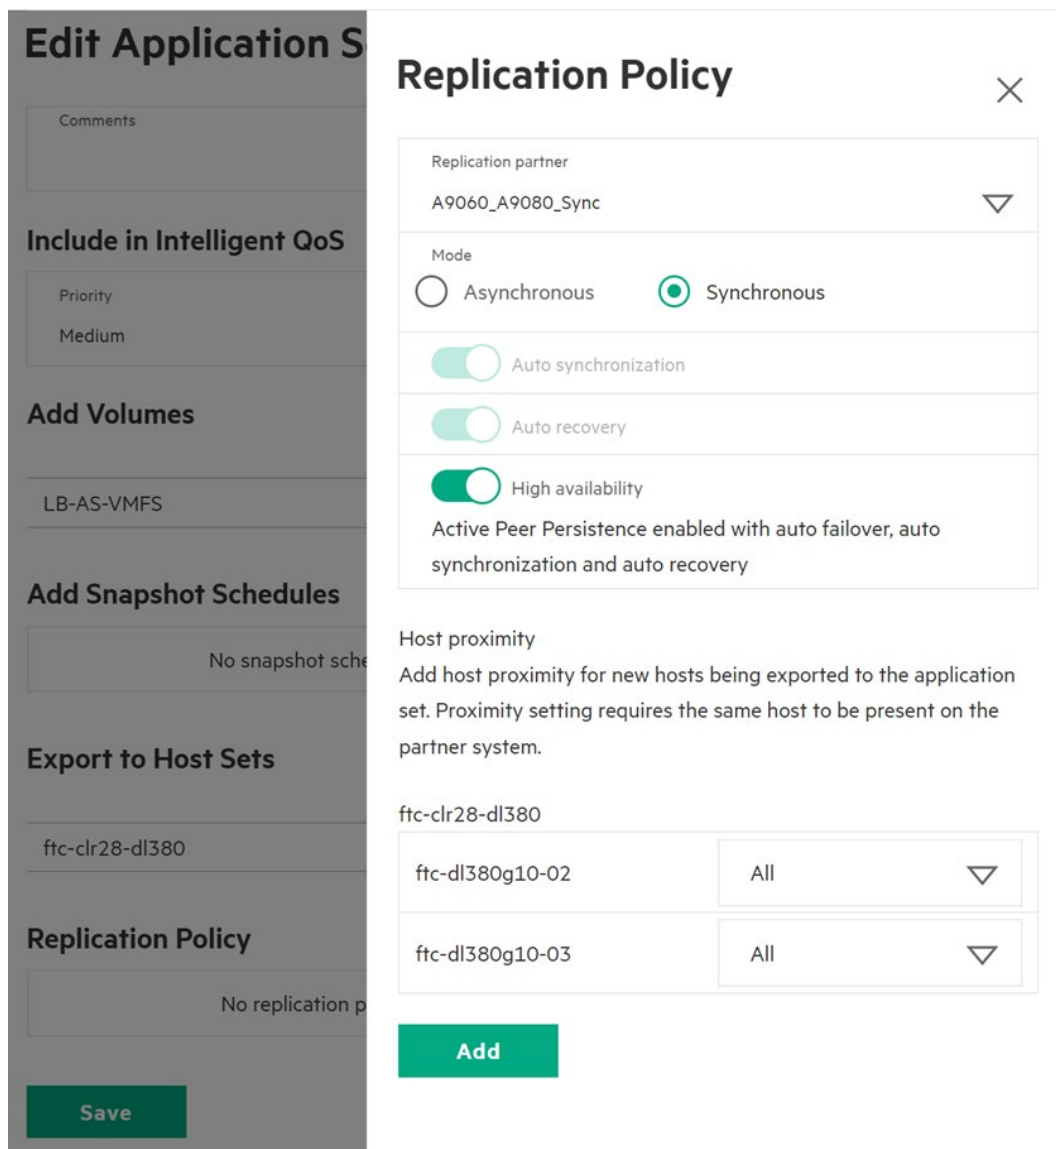


Figure 6. Creating an Active Sync replication policy in the onboard UI

### Configuring Active Sync through the SSMC

Starting with SSMC 3.8, the Create Group menu has new options that enable you to select a high-availability configuration, which includes both Active Peer Persistence and classic Peer Persistence.

Select **Active Peer Persistence**, as shown in Figure 4, to enable the group for Active Sync. The **auto failover**, **auto recover**, and **auto synchronize** options are enabled on the group by default when the Active Peer Persistence high-availability configuration is selected, as shown in Figures 7 and 8.



**Create Group** High Availability Configuration ?

**High Availability Configuration**  Advanced options

Select Active Peer Persistence to ensure that the target port group state of volumes in the specified group will be in Active/Optimized or Active/Non-Optimized state based on Host Proximity setting.

Active Peer Persistence  Peer Persistence  None

Auto failover

Figure 7. Select Active Peer Persistence — the Auto failover option is automatically enabled

**Create Group** Source ?

**Source**  Advanced options

Auto synchronize **Enabled**

Auto recover **Enabled**

Figure 8. The auto synchronize and auto recover policies are set automatically for an Active Sync group

The last step in creating a group is specifying host proximity. All hosts should be set with proximity **All**, as shown in Figure 9.

**Create Group** Host ?

**Host**  Advanced options

Choose Proximity setting for hosts for Primera-630-A. Click on 'Apply to all' to set the same Proximity as chosen for the first host in the table.

Host	Host OS	Proximity	Apply to all
esx3-8s201	VMware (ESXi)	Primary	<input type="checkbox"/>

**Add Hosts**

Checked: Auto failover

**Create** **Create +** **Cancel**

Figure 9. Setting host proximity in the SSMC



## Enabling Active Sync through the CLI

To enable the `active_active` policy for a group from the CLI, run the following CLI command from the primary array:

```
cli% setrcopygroup pol active_active <group_name>
```

To admit the hosts with the correct proximity, the command for setting host proximity to “all” is:

```
admitrcopyhost -proximity all {<host_name>|<host_set>} <group_name>
```

---

### Note

Host proximity must be set to `-proximity all` for all hosts in the stretched cluster that uses the Active Sync group.

---

## Best practices for Active Sync

Consider the following best practices when configuring and operating an Active Sync environment:

- Configure the Active Sync solution by following all the rules for an Active Peer Persistence solution, with the exception of the following:
  - The stretched cluster servers in an Active Sync solution are connected only to the array in the local data center; they are not cross-connected to the array in the distant data center.
  - The ESXi servers in the cluster should boot locally or, if they boot from SAN, they should boot from a non-replicated boot disk.
  - Set the host proximity for all of the servers in the stretched cluster to `-proximity all`.
- Follow all vendor best practice recommendations for their stretched cluster deployment.
- Ensure that the replication links are sized properly for the replication workload.
- Leave the `auto_recover` and `auto_synchronize` options set on all Active Sync RC groups (they are set by default).
- Configure a stretched cluster across the hosts in the two data centers.
- The servers in the stretched cluster should connect only to the array in the local data center.

If you can cross-connect all the servers in the stretched cluster so that they are connected to both arrays, then you should deploy an Active Peer Persistence configuration.

- Follow all vendor best practice recommendations for a stretched cluster.
- The host OS (physical server/hypervisor) should be configured either for local boot or boot from SAN. If boot from SAN is chosen, be sure to boot from a nonreplicated boot volume.
- Ensure that you have the same WWN on both arrays for the VVs in the Active Sync RC group.
- The individual host definitions and host sets must exist on both systems. For the hosts that are not attached to a given site, the host definition can be empty (no initiators).
- Ensure proper server/array zoning.
- Set up and configure a quorum witness (necessary for automatic failover in the event of an array failure).

Without a quorum witness in the configuration, if the primary array fails, the secondary array will not perform an automatic failover. All of the servers in the stretched cluster will lose access to their storage, and all VMs and associated applications will stop.

- Ensure proper replication link sizing.
- If you plan to administratively stop any RC groups in the Active Sync configuration, migrate all VMs that use the secondary array to ESXi cluster servers that use the primary array before stopping the group.



## Summary

Remote Copy Active Sync mode on HPE Alletra 9000 and HPE Primera arrays is a significant improvement over traditional synchronous replication. With the ability to service host I/O for the volumes in an RC group from both arrays in the Active Sync solution, the administrator can now deploy and migrate VMs seamlessly between ESXi servers deployed in a vMSC cluster. This makes the management of a VMware environment much easier, allowing load balancing of I/O for clustered applications between arrays and supporting seamless migration and manual failover of array RC groups to aid in maintenance and upgrade tasks. Hewlett Packard Enterprise recommends deploying and using Active Peer Persistence in environments in which a uniform host connection configuration can be supported for its ability to recover seamlessly from most failure scenarios. Confirm that Active Sync is supported in your environment (see [HPE SPOCK](#) for details on host and operating systems supported with Active Peer Persistence).



## Resources

[HPE Active Peer Persistence technical white paper](#)


[HPE SPOCK](#)

[HPE Alletra 9000: Configuring and managing data replication using Remote Copy](#)

## Learn more at

[hpe.com/us/en/storage.html](https://hpe.com/us/en/storage.html)

Explore **HPE GreenLake** 



© Copyright 2023 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

VMware ESXi, VMware vSphere, and VMware are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All third-party marks are property of their respective owners.

a00135531ENW