PURESTORAGE®

WHITE PAPER

Simplifying SAP HANA Disaster Recovery with ActiveDR

Extend SAP HANA Resilience with storage-based replication from Pure Storage®

Contents

Executive Summary

SAP HANA enables organizations to run advanced analytics and high-speed transactions in a single system. This seamless, low-latency environment is critical for business operations, making any disruption potentially catastrophic. Ensuring data consistency and minimizing downtime is essential for maintaining business continuity.

Traditional disaster recovery solutions are often complex, costly, and unable to meet strict recovery time objectives and recovery point objectives. Pure Storage FlashArray™ ActiveDR® delivers near-synchronous replication with near-zero recovery point objectives, even across geographically dispersed data centers. Its continuous storage replication provides business continuity, seamless failover testing, and non-disruptive failovers and failbacks.

With ActiveDR, SAP HANA administrators and business leaders can confidently implement disaster recovery strategies, enhance resilience, and protect critical application infrastructure without operational disruption.

Solution Overview

SAP HANA, combined with ActiveDR, provides businesses with a powerful solution for continuous, near-synchronous data replication. In simple terms, this means that a local database site can replicate its data almost simultaneously to a remote site. This setup minimizes the risk of data loss and ensures that, in the event of failure, systems can be quickly recovered with virtually no loss of recent data. The result is enhanced operational continuity and improved customer satisfaction by keeping critical business processes running smoothly.

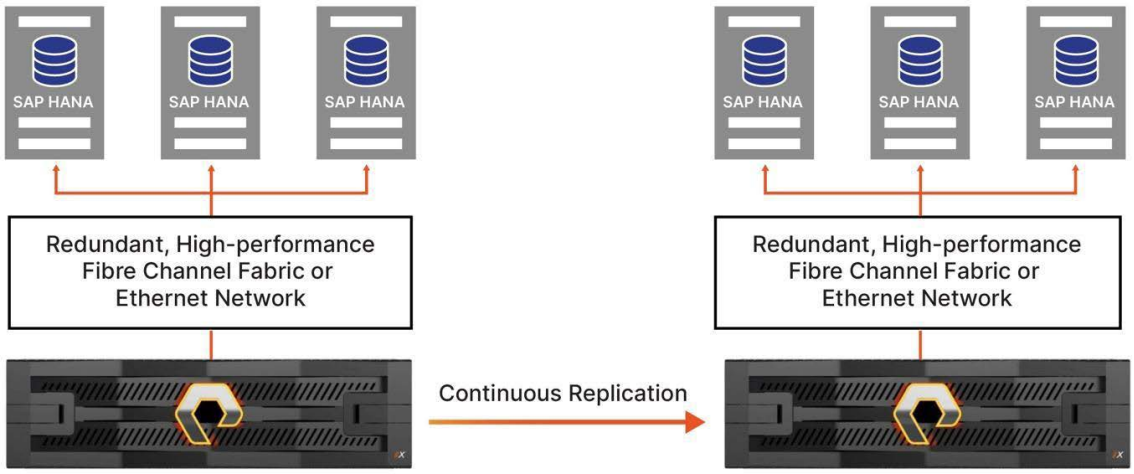

FIGURE 1 High-level overview of the SAP HANA, FlashArray, and ActiveDR solution.

This solution includes the following deployment types and topologies:

- SAP HANA scale-up and scale-out deployment types
- SAP HANA databases running on either bare metal or virtual machines
- FlashArray storage using Fibre Channel or Ethernet connectivity
- A local or wide area network between data centers

Solution Benefits

ActiveDR offers several strategic benefits for businesses using SAP HANA:

- **No additional licensing**: ActiveDR is included with FlashArray at no extra cost. Additionally, with ActiveDR, the secondary SAP HANA instance is on standby, reducing resource costs.
- **Reduced complexity**: ActiveDR can replicate databases at the storage level with varied deployments, including scale-up or scale-out configurations, enabling administrators to focus on database management for their unique business use cases.
- **Simple, non-disruptive test failovers**: Storage and database administrators can test ActiveDR failovers with SAP HANA without disrupting operations at the primary site. This is ideal for any SAP HANA configuration, including test and development workflows, as it doesn't impact recovery point objectives or recovery time objectives.
- **Fast recovery and failovers**: Storage administrators can perform failovers on protected SAP HANA volumes with a single command. Failovers include protected volumes and any protection group snapshots.
- **Near-zero recovery point objectives**: ActiveDR continuously streams writes between the source and target FlashArray systems, which provides near-zero recovery point objectives for protected databases. This has no performance impact on the source FlashArray or SAP HANA, even in third-site system replication scenarios.
- **Near-synchronous data replication between sites**: ActiveDR provides near-synchronous storage replication between primary and secondary sites. Data at a secondary site can be quickly attached and used by SAP HANA in the event of a primary site failure, which can help simplify disaster recovery workflows, especially in scale-up or scale-out configurations.
- **No latency impact**: Primary and secondary arrays can replicate at nearly any distance without affecting the performance of SAP HANA.

Technology Overview

The following sections provide an overview of the technologies that are used in an ActiveDR environment for SAP HANA.

Pure Storage FlashArray

FlashArray is an all-flash storage solution that gives storage and database administrators a fast, scalable, unified block and file storage platform that is ideal for high-performance SAP HANA databases.

By providing a unified interface and simple-to-use tools, FlashArray gives storage administrators the ability to quickly and seamlessly replicate, move, and manage data. FlashArray also deduplicates and compresses all data before it is written, efficiently reducing the size of data without impacting performance. Storage and database administrators can further increase storage capacity by using the FlashArray snapshot capabilities to create snapshots of production databases. They can then use those snapshots in development or testing environments.

The following FlashArray product models are certified for SAP HANA:

- **FlashArray//X™**: Provides high-performance, high-capacity storage that is ideal for performance-oriented workloads
- **FlashArray//XL™**: Provides high-performance storage at scale that helps reduce the number of arrays needed to run large applications

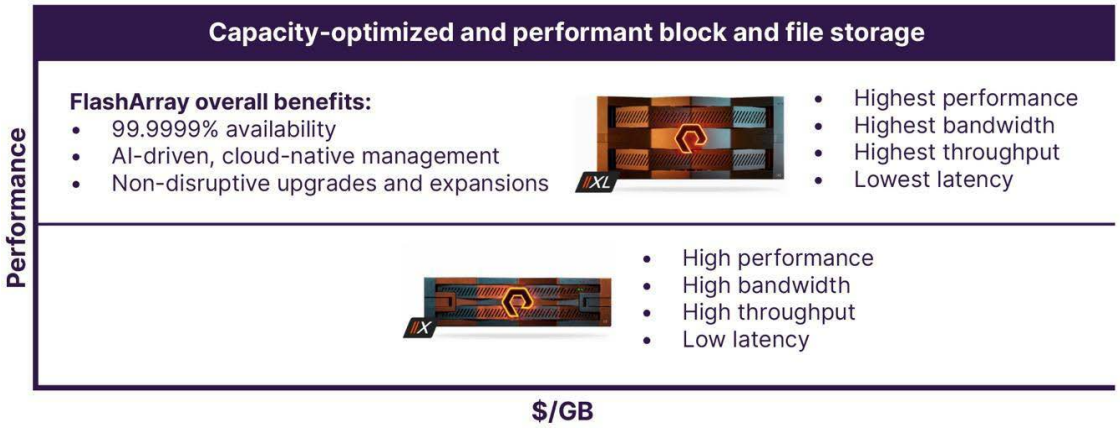

FIGURE 2 The FlashArray product family is certified by SAP HANA.

ActiveDR

ActiveDR provides near-synchronous storage replication between two FlashArray systems within or across disparate data centers to protect against threats such as hardware failures, ransomware attacks, and user error. ActiveDR enables a near-zero recovery point objective that improves business resilience for critical application infrastructure compared to traditional asynchronous replication.

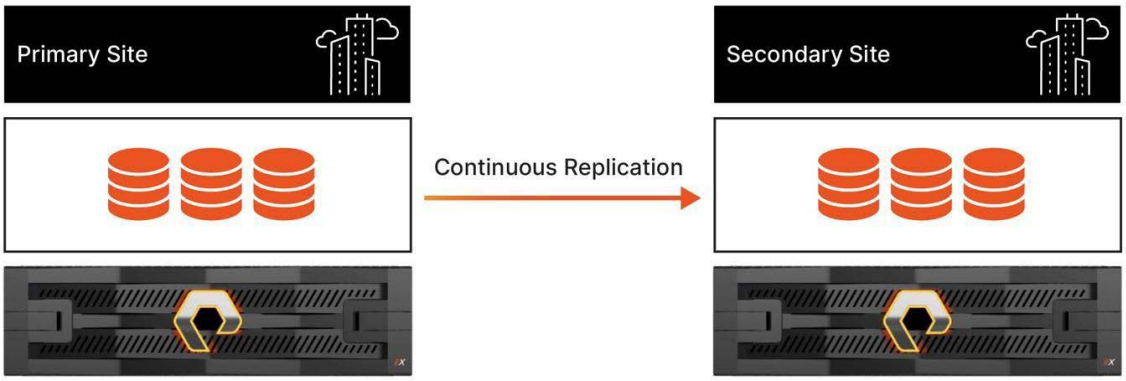

FIGURE 3 Near-synchronous storage replication between two FlashArray systems with ActiveDR.

ActiveDR simplifies disaster recovery by providing continuous data protection and enabling non-disruptive testing of disaster recovery workflows. Organizations can perform test failovers, live failovers, resynchronization, and failbacks without disrupting production or halting replication, ensuring operational continuity while validating recovery readiness.

ActiveDR uses near-synchronous replication, making it ideal for latency-sensitive workloads and geographically distributed environments. Unlike synchronous replication, it doesn't require remote acknowledgment of writes, allowing efficient use of existing wide area networks without performance trade-offs.

ActiveDR consists of three components:

- **Pods**: These are storage-management containers that organize storage objects and configuration settings into groups that are failed over and failed back as a unit. A pod can contain volumes, volume snapshots, and protection groups. Additionally, a pod can contain configuration settings such as protection group snapshot schedules, snapshot retention policies, and quality-of-service volume limits.
- **Replica links**: These provide replication between pods and provide directional and auto-reversing capabilities. Once a replica link is created, ActiveDR is automatically enabled.
- **Connected FlashArray systems**: ActiveDR requires a minimum of two FlashArray systems connected over a network to replicate data between. With no latency requirements, these systems can be within the same data center or as far apart as on different continents.

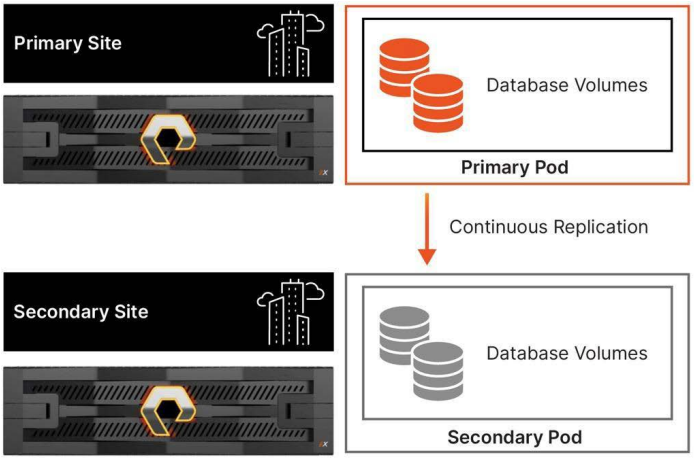

FIGURE 4 A typical ActiveDR deployment that includes pods and replica links between FlashArray systems.

Pod-based replication helps simplify storage management across sites by replicating all configuration changes made on a primary site array to the secondary site array, which helps simplify management and disaster recovery storage failovers. ActiveDR also supports multi-directional replication for different pods. For example, a database administrator might have a pod at their primary site that they want to replicate to their secondary site, while the secondary site might contain a pod that they want to replicate back to their primary site. ActiveDR lets them easily configure the pods to replicate in either direction between sites.

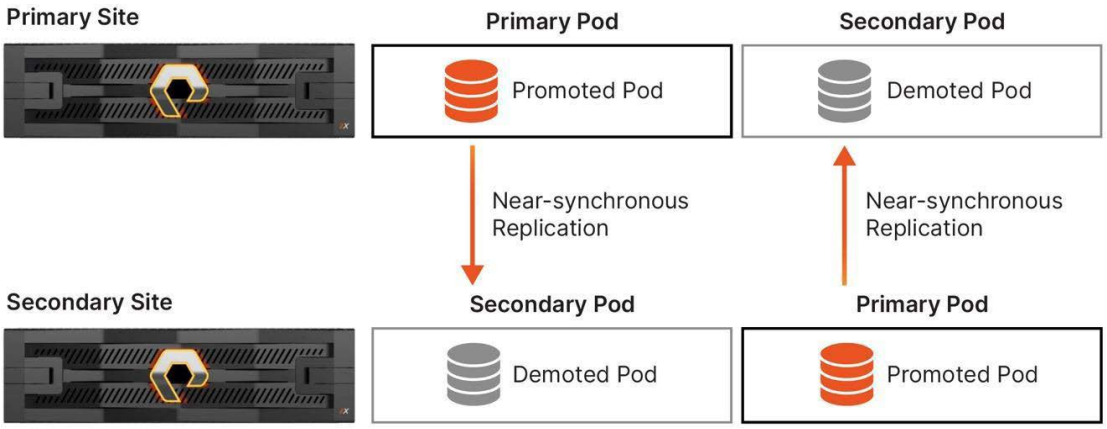

FIGURE 5 Pods in either site can be replicated to either site.

Differences Between ActiveCluster, Asynchronous Replication, and ActiveDR

FlashArray has several storage replication functions that can be used to protect an organization's data. These include synchronous replication, which is used by the ActiveCluster solution; asynchronous replication of storage array volume snapshots; and near-synchronous replication with ActiveDR.

ActiveCluster

ActiveCluster uses synchronous replication to maintain copies of data between two FlashArray systems. When data is written to a primary site FlashArray, it is simultaneously copied to a secondary site FlashArray. Once the data is written to both arrays, the write is acknowledged to the host system. This method of replication is recommended when the latency between arrays is 11ms or less, which means that ActiveCluster should be used between arrays in the same data center, or between data centers that have very-low-latency wide area network capabilities. For more information about ActiveCluster, see ActiveCluster Solution Overview.

Asynchronous Replication

Asynchronous replication is a snapshot-based solution that uses space-efficient snapshots to replicate data between FlashArray storage devices. Asynchronous snapshot replication occurs according to a customer specified schedule, while ActiveDR is a streaming-based solution that continuously replicates volume data between FlashArray systems at different sites. When asynchronous replication is enabled on a volume on the primary site FlashArray, a snapshot of the volume is created on the primary site array and then replicated to the secondary site array. The first snapshot transfer is a baseline, which is a complete copy of the entire contents of the volume. All subsequent transfers are incremental transfers that result by comparing existing data on the storage array with the newly created snapshot to determine what data to send to the secondary site array. For more information about asynchronous replication, see FlashArray Asynchronous Replication Configuration and Best Practices Guide.

ActiveDR

ActiveDR offers near-synchronous, pod-based replication. ActiveDR is ideal for continuously replicating data between high-distance or high-latency sites and is best suited for disaster recovery scenarios. Unlike ActiveCluster and asynchronous snapshot replication, it works best in any latency scenario and does not require continuous copies of data or host confirmation during replication.

For more information about ActiveDR, visit the Pure Storage support site.

SAP HANA

SAP HANA is a column-oriented in-memory database platform. This in-memory architecture provides high-speed processing and real-time analytics for users, and it also uses savepoints to periodically flush changed data to persistent storage (disk) to ensure data durability and prevent data loss in case of failure. It combines online transaction processing and online analytical processing for running transactional and analytical workloads simultaneously. SAP HANA is certified with FlashArray, making it capable of processing massive amounts of data with near-zero latency and protecting that data with ActiveDR storage replication.

SAP HANA can function as a database and development platform, enabling businesses to eliminate redundant data and speed up workloads. Combined with ActiveDR, it offers a unique platform for testing applications and disaster recovery workflows in a variety of scaling and non-scaling deployments. Combined with SAP HANA system replication, ActiveDR offers businesses unique options for improving performance, disaster recovery, and high availability.

ActiveDR Disaster Recovery Architecture

A traditional ActiveDR deployment involves the configuration of a replication link between two or more FlashArray systems at different sites to ensure seamless failover for disaster recovery scenarios. In this architecture, a SAP HANA instance is kept offline at the target site in standby for failover in the event of a disruption in the SAP HANA instance on the primary site. This configuration guarantees minimal downtime and fast recovery time objectives and recovery point objectives for the target system, ensuring business continuity during disruptions.

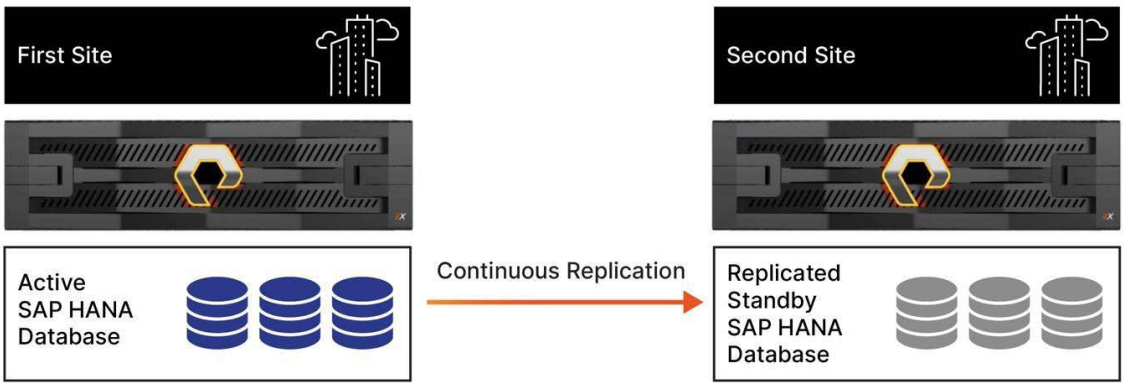

FIGURE 6 An ActiveDR site on standby.

SAP HANA System Replication and ActiveDR

SAP HANA offers a built-in software replication solution known as HANA system replication. SAP HANA System replication is set up so that a secondary system is configured as an exact copy of the active primary system, with the same number of active hosts in each system. Unlike ActiveDR, this replication occurs at the database layer, not the storage layer. System replication replicates the logical state of the database on a configured (synchronous, asynchronous, etc.) basis.

SAP HANA system replication can be used alongside ActiveDR to create a comprehensive disaster recovery strategy by adding a third replication site. While system replication is running, the secondary system is configured identically to the primary system and is on standby until a failover takes place. ActiveDR can be used to replicate either the primary or secondary site at the storage level for an additional layer of protection.

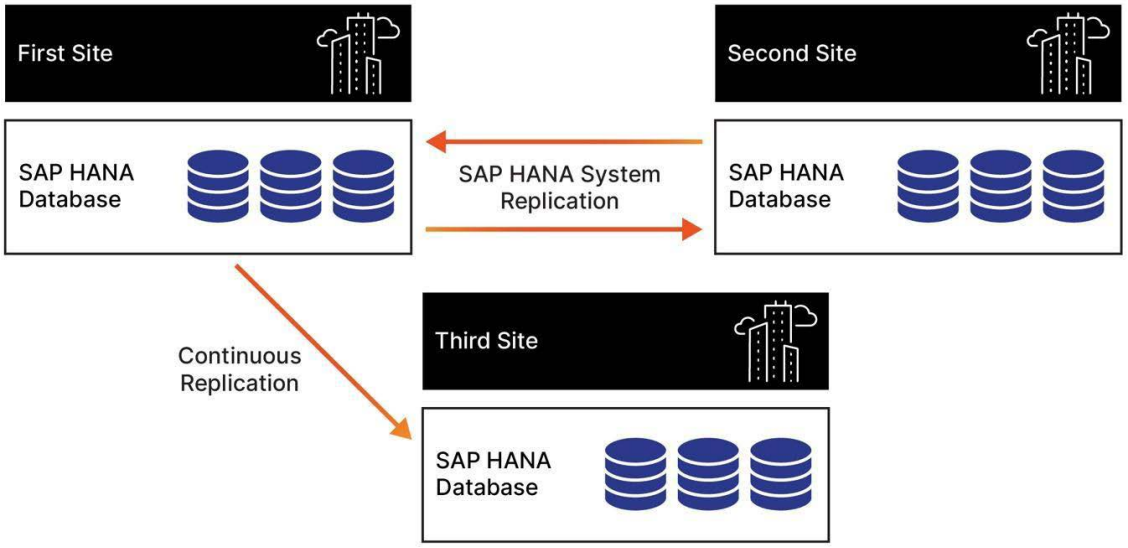

FIGURE 7 SAP HANA system replication and ActiveDR storage replication duplicating data to different sites.

In disaster scenarios, administrators will have two failover sites, one using SAP HANA system replication and another using storage replication, which allows them to prioritize the data most important to their current operations. This flexibility enables businesses to tailor their recovery strategies based on their current needs and the criticality of various business functions. In the event of a primary system failure, the secondary SAP HANA system can take over operations, but if the secondary site also encounters an issue, ActiveDR replication at a tertiary site provides an additional layer of protection.

This dual-layered approach improves resilience and ensures critical data remains accessible and secure while maintaining high availability and comprehensive disaster recovery capabilities.

For more information about SAP HANA system replication, see SAP HANA System Replication.

Considerations for Networking, Failover Logic, and Recovery Workflows

When designing an ActiveDR deployment for SAP HANA, it's important to consider restrictions and response workflows for various situations. This section describes three key areas to consider.

Networking

While ActiveDR functions at nearly any distance, making latency a minimal issue, it is important to consider bandwidth and network redundancy to ensure high availability. Data needs to be protected and available, and any bottlenecks, whether bandwidth or network failure-related, could impact the system. It's ideal to eliminate single points of failure via redundancy to minimize the risk of downtime.

Failover Logic

Failover logic should prioritize recovery for the most critical site first. An effective failover strategy will ensure the system attempts recovery from the closest or least-impacted site before escalating to the next. If using SAP HANA system replication in addition to ActiveDR, automatic failover can be configured, but it could present issues with data inconsistencies or performance impacts. Administrators should consider establishing clear failover priorities to streamline the process during a disaster scenario.

Recovery Workflows

It's critical to align the recovery point objectives and recovery time objectives with business needs. To do this, recovery workflows should include testing, validation, documentation, and automation. An effective failback process should be carefully planned and automated where possible with clear guidelines to prevent data inconsistencies and interruptions to business continuity. Creating a disaster recovery playbook eliminates questions and keeps everyone aligned in emergency situations.

Scale-up vs. Scale-out Deployment Types

SAP HANA offers two deployment architectures for scaling a system: scale-up and scale-out.. Scale-up means increasing the size of one physical machine by increasing the amount of RAM available for processing. Scale-out means combining multiple independent computers into one system to overcome the hardware limitations of a single physical server and to distribute the load between multiple servers.

ActiveDR integrates with these scaling methods to provide options for replication and data protection with minimal performance impact for an additional layer of security in disaster scenarios. However, ActiveDR works differently depending on the scaling configuration:

- In scale-up architectures, ActiveDR replicates both the data and log volumes from a single host to a secondary site.
- In scale-out architectures, ActiveDR replicates the data and log volumes for all worker hosts that are part of the SAP HANA scale-out deployment.

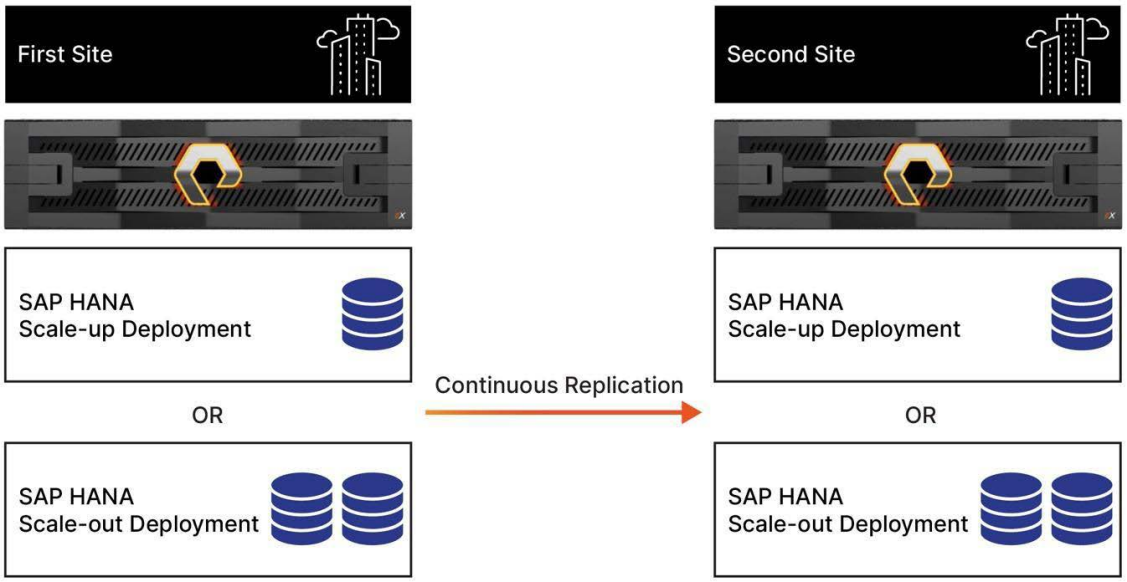

FIGURE 8 ActiveDR replication between two arrays running SAP HANA scale-up or scale-out architectures.

Test and Development Workflows Using ActiveDR

For testing, development and quality assurance cases, ActiveDR offers unique functionality ideal for teams consistently testing application workflows.

Secondary testing environments can easily be enabled without impacting primary site operations. To do this, administrators might temporarily suspend replication and spin up a secondary testing environment for a variety of use cases including workload validation, development, and software testing. This doesn't incur additional infrastructure costs, making it ideal for organizations with robust testing, development, and quality assurance workflows.

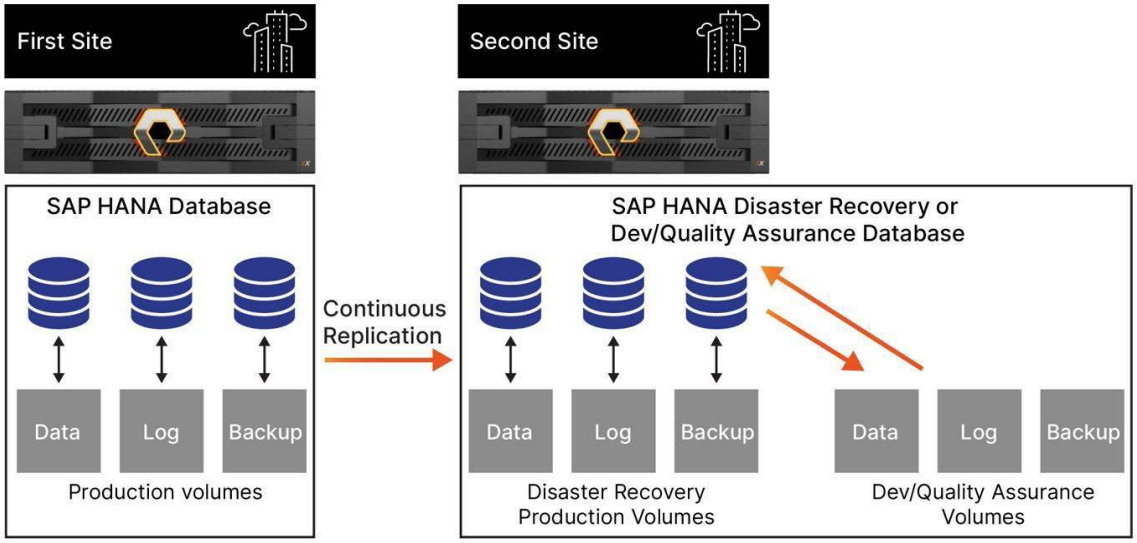

FIGURE 9 A testing and development SAP HANA deployment using ActiveDR.

Once testing is completed, administrators can resume replication from a specific point of suspension, ensuring no impact to the primary site's operations or business continuity.

Using ActiveDR with SAP HANA

The following sections describe how to configure and test ActiveDR with SAP HANA. These sections assume that the primary and secondary sites are already configured to work with FlashArray pods.

Note: All workflow examples are for scale-up deployments.

Here are the minimum requirements to configure ActiveDR with SAP HANA:

- The same SAP HANA version deployed at two sites with the secondary site configured for failover.
- (Optional) System replication for SAP HANA high availability and storage replication for SAP HANA disaster recovery. See ActiveDR and SAP HANA System Replication for more details.

Deployment and Disaster Recovery Configuration

The following steps must be completed using the Pure Storage user interface (GUI or CLI) to initiate ActiveDR replication.

Configuring ActiveDR for SAP HANA involves the following general steps:

1. Create volumes at the primary site.
2. Configure SAP HANA to store files on FlashArray volumes at the primary site.
3. Move the primary site's volumes into a pod.
4. Create an ActiveDR replica link between the source pod and the target replication pod on the secondary FlashArray.

Configure the ActiveDR Primary Site Pod in the Pure Storage User Interface

To configure the primary site pod:

1. From the primary site's Pure Storage user interface, select **Storage** in the navigation pane, and then select the **Pods** tab.

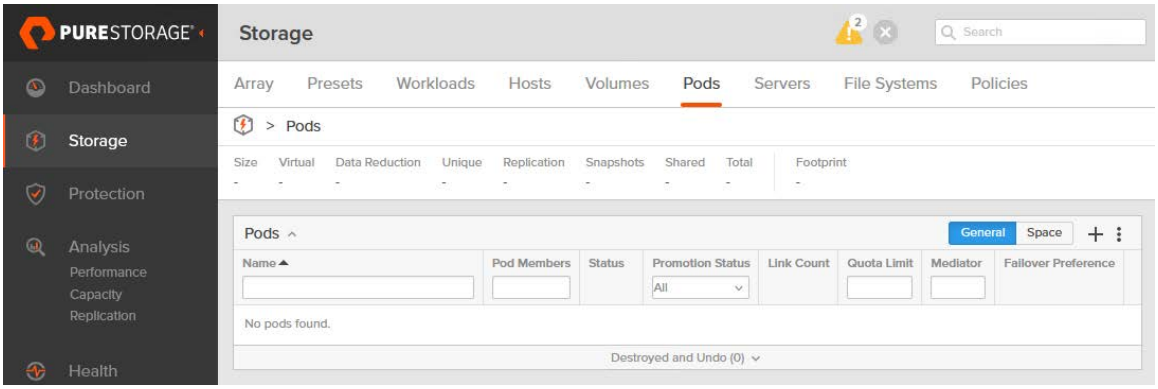

FIGURE 10 The Pods tab in the Pure Storage user interface.

2. Click the **+** icon in the **Pods** group to create a new pod.
3. In the **Name** field in the **Create Pod** dialogue box, enter a name for the pod, and then click **Create**. The pod appears in the **Pods** group.

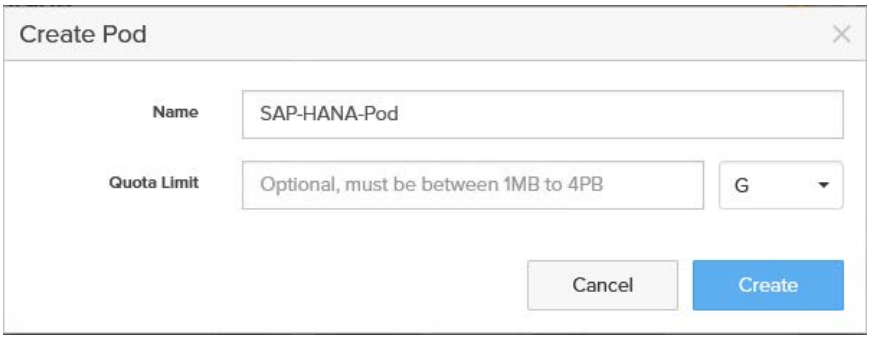

FIGURE 11 Creating the pod.

4. In the **Pods** group, select the name of the pod that was created to display the **Pod management** view.

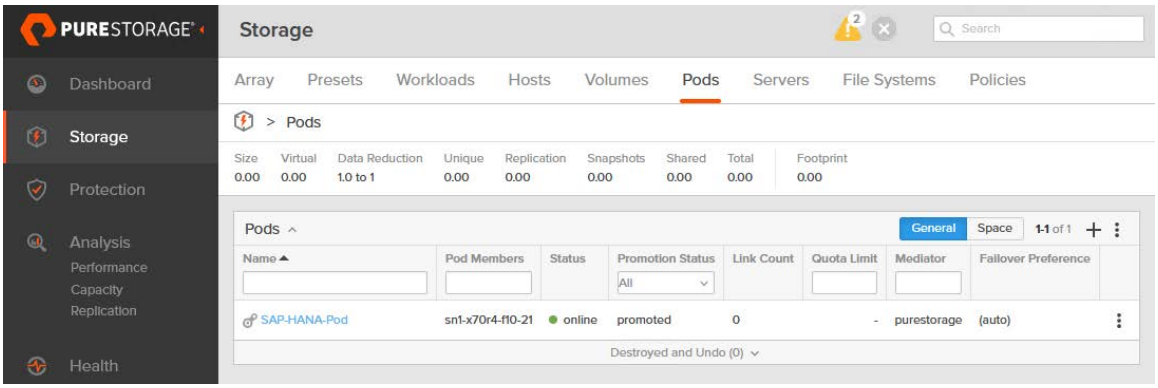

FIGURE 12 The Pod management view.

5. In the **Volumes** group, click the ellipsis, and then click **Move In** to display the **Move Volumes In** dialogue box.

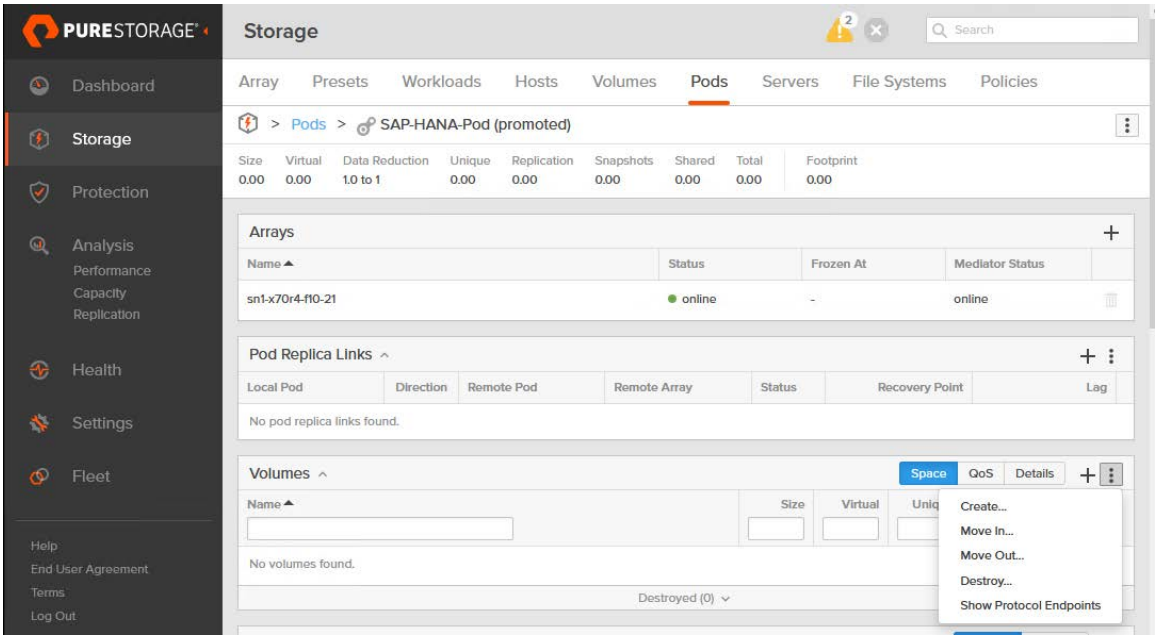

FIGURE 13 Moving the volumes.

6. Select the SAP HANA volumes that will be moved into the pod, and then click **Move**. The volumes appear in the **Volumes** group using the naming convention, *<pod name>::<volume name>*.

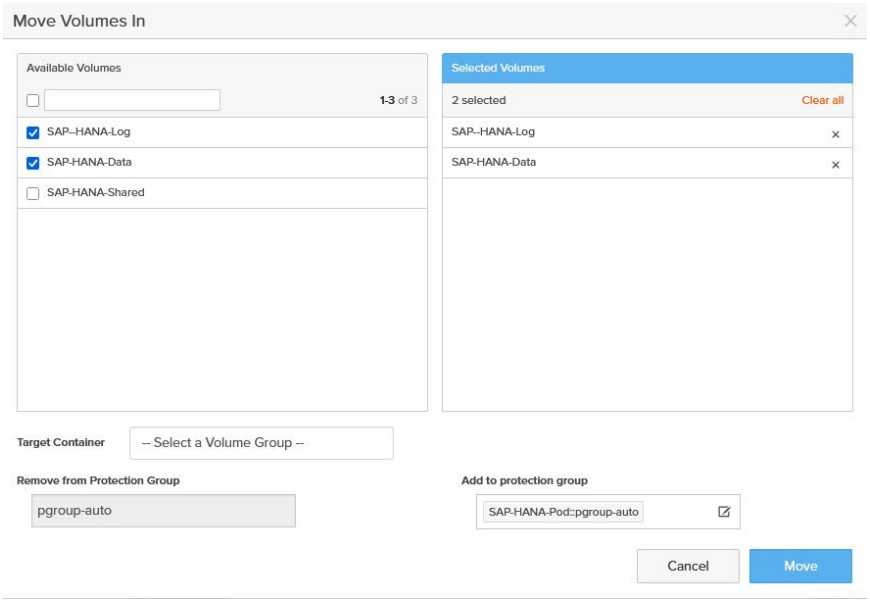

FIGURE 14 Moving volumes into a pod.

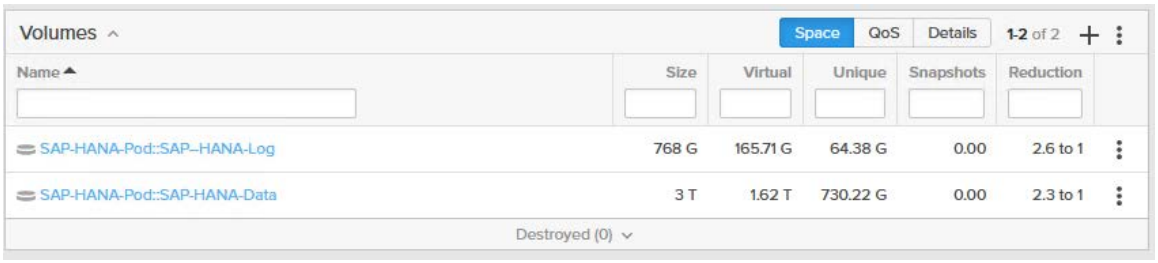

FIGURE 15 Displaying the volumes in a pod.

The primary site pod is now configured with the database data and log volumes. The next step is to create a replica link to the remote site's FlashArray, which enables ActiveDR replication.

Create a Replica Link in the Pure Storage User Interface

1. From the remote site's Pure Storage user interface, select **Protection** in the navigation pane.
2. In the **Array Connections** group, click the ellipsis, and then select **Get Connection Key**.

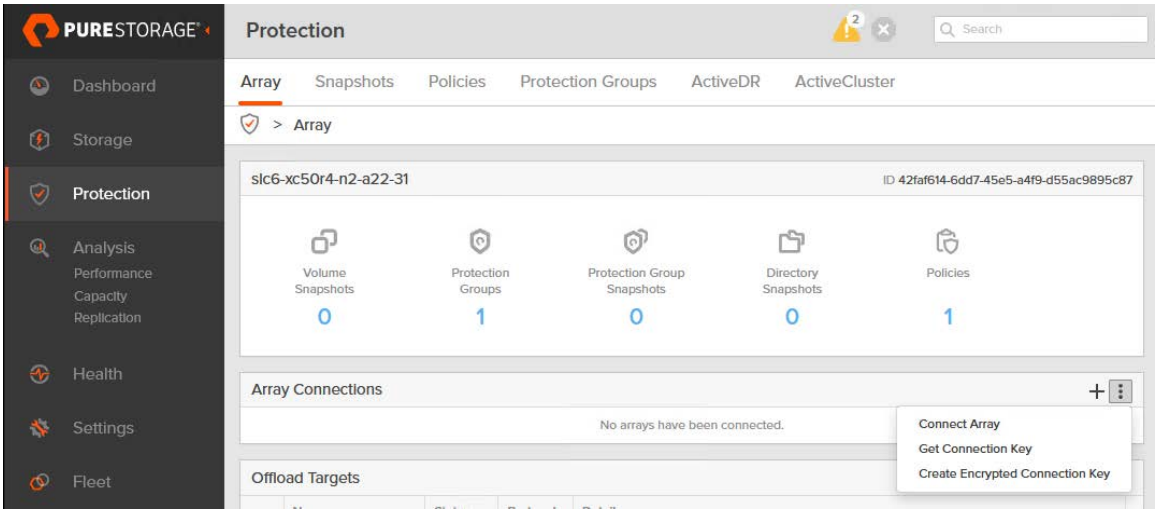

FIGURE 16 Getting a connection key on the **Array** tab.

3. In the **Connection Key** dialog box, click **Copy** to copy the connection key to the clipboard.

FIGURE 17 The connection key dialogue box.

4. Return to the primary site's Pure Storage FlashArray user interface and navigate to the **Protection** tab.
5. In the **Array Connections** group, click the ellipsis, and then select **Connect Array**.

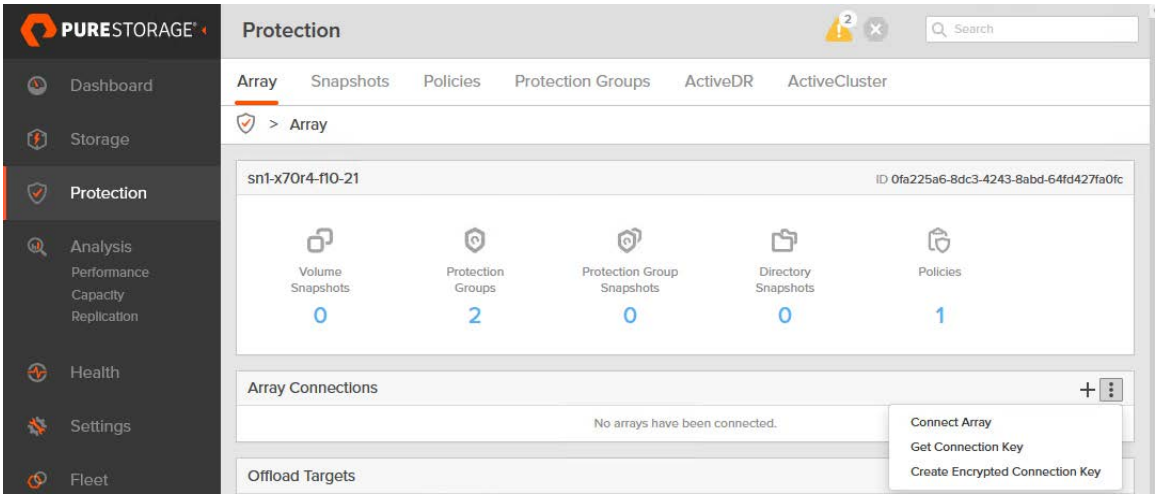

FIGURE 18 Protection view for the primary site FlashArray.

6. In the **Connect Array** dialogue box, enter or select the following information:
   - **Management Address**: The IP address of the remote site's FlashArray.
   - **Type**: Select **Async Replication** from the drop-down list.
   - **Connection Key**: Paste the connection key that you copied from step 3 into the remote site's Pure Storage user interface.
   - **Replication Transport**: Select **Ethernet (IP)** from the drop-down list.
   - **Encrypted**: Select the button to toggle replication encryption on or off.

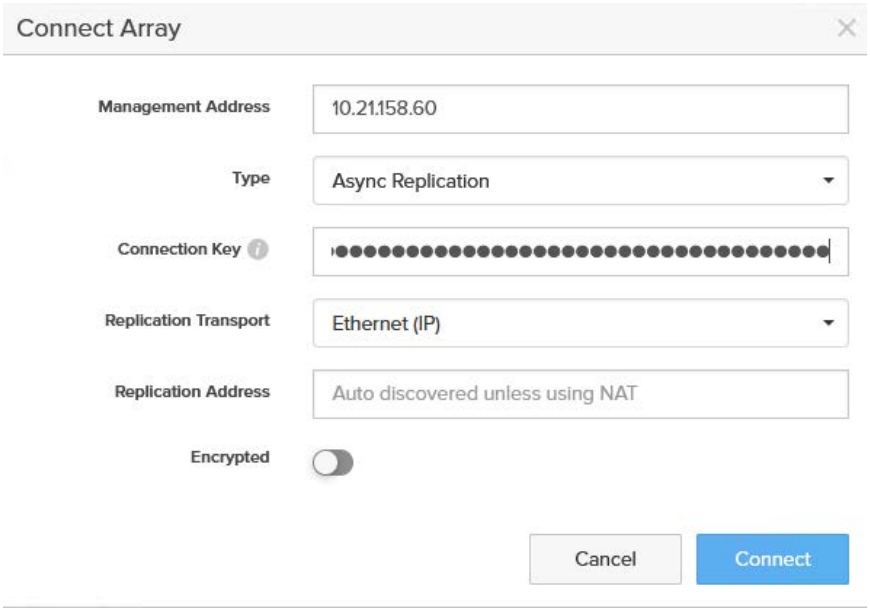

FIGURE 19 Connecting the primary and secondary arrays.

7. Click **Connect**.

The secondary site array now appears in the **Array Connections** group.

FIGURE 20 Displaying the secondary site in the **Array Connection**s group.

8. In the **Pod Replica Links** group for the pod that was just created, click the ellipsis, and then select **Create**.

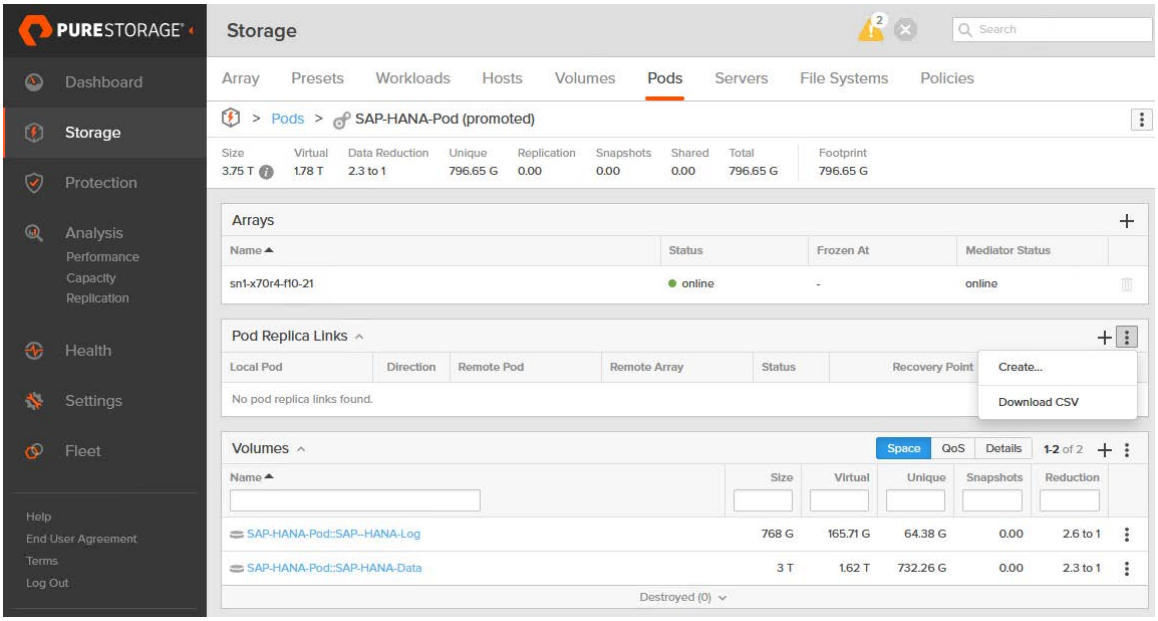

FIGURE 21 Creating a replica link for the previously created pod.

9. In the **Create Replica Link** dialogue box, select the remote site FlashArray from the **Remote Array** drop-down list, and then select a remote pod.

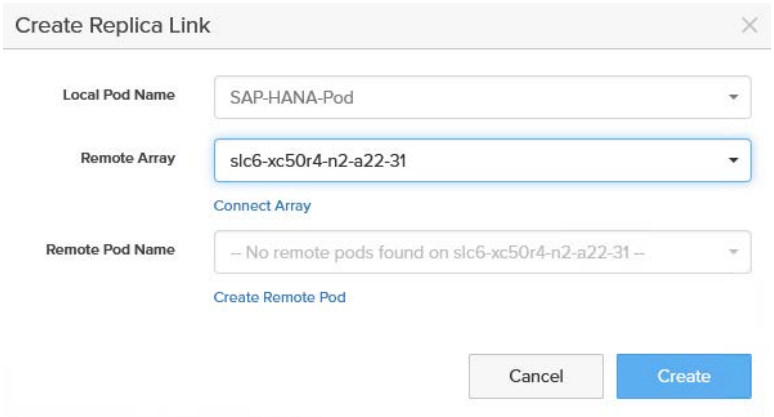

FIGURE 22 Creating the replica link.

Note: If a pod hasn't been created on the remote site's FlashArray, click **Create Remote Pod**, enter a pod name in the **Name** field, and then click **OK**.

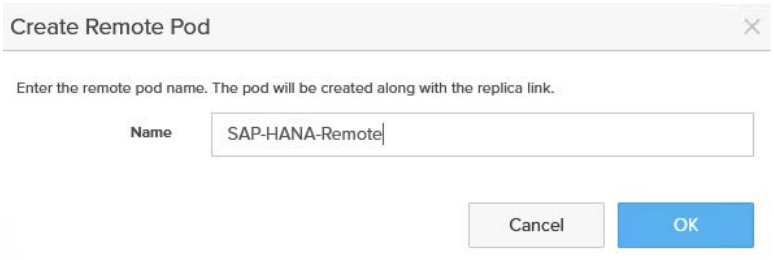

FIGURE 23 Naming the pod.

10. Select **Create**.

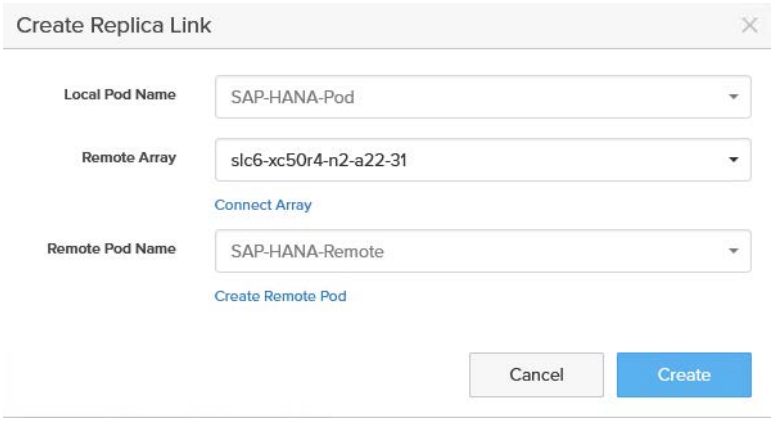

FIGURE 24 Finalizing the replica link.

Once a replica link has been created from the primary site's pod to the secondary site's pod, the primary site's pod begins a baseline replication of the volumes to the secondary site's pod. During a baseline replication operation, the primary site FlashArray replicates the volumes to the secondary site, which, depending on the size of the data, can take a significant amount of time. The status of the baselining progress can be monitored in the **Status** field in the **Pod Replica Links** group. When the baseline replication completes, the **Status** field changes from "baselining" to "replicating."

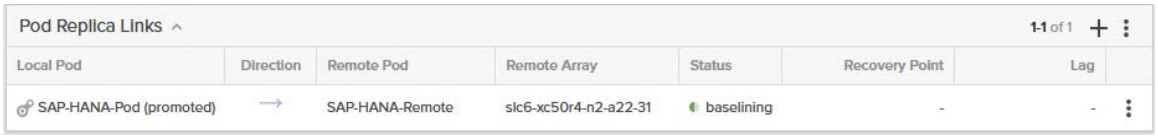

FIGURE 25 The pod status shows baselining.

FIGURE 26 The pod status shows replicating.

In addition to monitoring the pod replica link status at the primary site, the replica link status can also be monitored in the secondary site's Pure Storage user interface in the **Pod Replica Links** group.

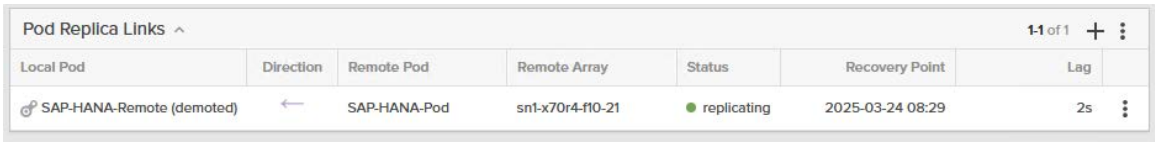

FIGURE 27 Monitoring the pod replica link status on the secondary site's FlashArray.

11. (Optional) Connect your DR hosts to the target pod volumes on the replication target FlashArray. Although the volumes within the target pod are write-disabled, you can pre-connect them to hosts at the target FlashArray system to reduce the number of steps required during a failover.

ActiveDR Disaster Recovery Workflows for SAP HANA

For disaster recovery, there are a few options for testing failovers and workflows for real scenarios. Here's how to perform each one.

Non-disruptive Failover Testing

Non-disruptive failover with ActiveDR allows for any one of the disaster recovery/target sites to bring a pod online without changes being replicated back to the production/primary pod. During a non-disruptive test, the databases in the primary pod remain online and continue to replicate transactions to the FlashArray at the secondary site. This ensures uninterrupted production operations while enabling failover testing and allowing the target database to be used for testing and development purposes once the disaster recovery testing is complete.

Note: When using this process, the primary site will still queue changes, which will be applied to the target site once the process is undone automatically, with no manual intervention required.

During this process, while the target pod is promoted, ActiveDR will maintain replication from the source pod in the background, continuing to stream new content to the target array. The target FlashArray system will store the replicated content in a separate accounting bucket using metadata and pointers for space efficiency. The replicated content will not be shown in the target pod while the pod is in a promoted state.

When failover testing is complete, simply demote the target pod. Demoting a target pod will cause any test data written into that pod to be discarded from the pod. The content that was replicated from the source will then be attributed to the target pod.

To perform a non-disruptive failover, the state of the environment must be as follows:

- There are two FlashArray storage platforms with an ActiveDR replica link.
- There are two separate SAP HANA systems running the same version. These can be either scale-up or scale-out hosts.
- There is a demoted remote pod on a separate array in the same active pod replica link.

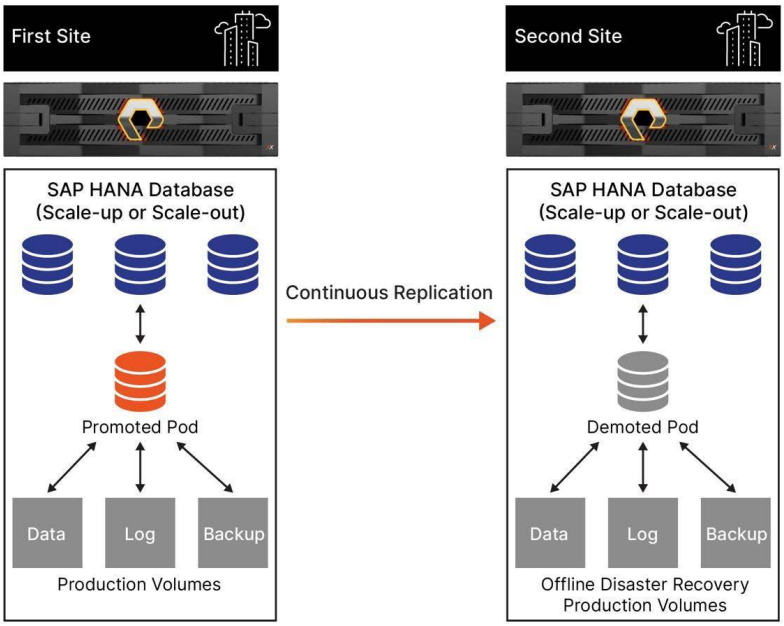

FIGURE 28 The state of the SAP HANA database before performing failover testing.

Performing the Failover

A non-disruptive failover consists of the following steps:

1. Promote the ActiveDR pod on the target array.
2. Connect the log and data volumes to the target systems.
3. Mount them to the relevant SAP HANA volume locations.

A failback after a non-disruptive failover only requires that the pod be demoted on the target system. Replication will automatically resume from the source pod to the target.

To complete the non-disruptive failover and failback between two arrays and different SAP HANA host systems:

1. In the **Storage** section on the production/primary array, click the **Pods** tab, and then observe the state of the **Pod Replica Links** in a specific pod's details, which will show as promoted.

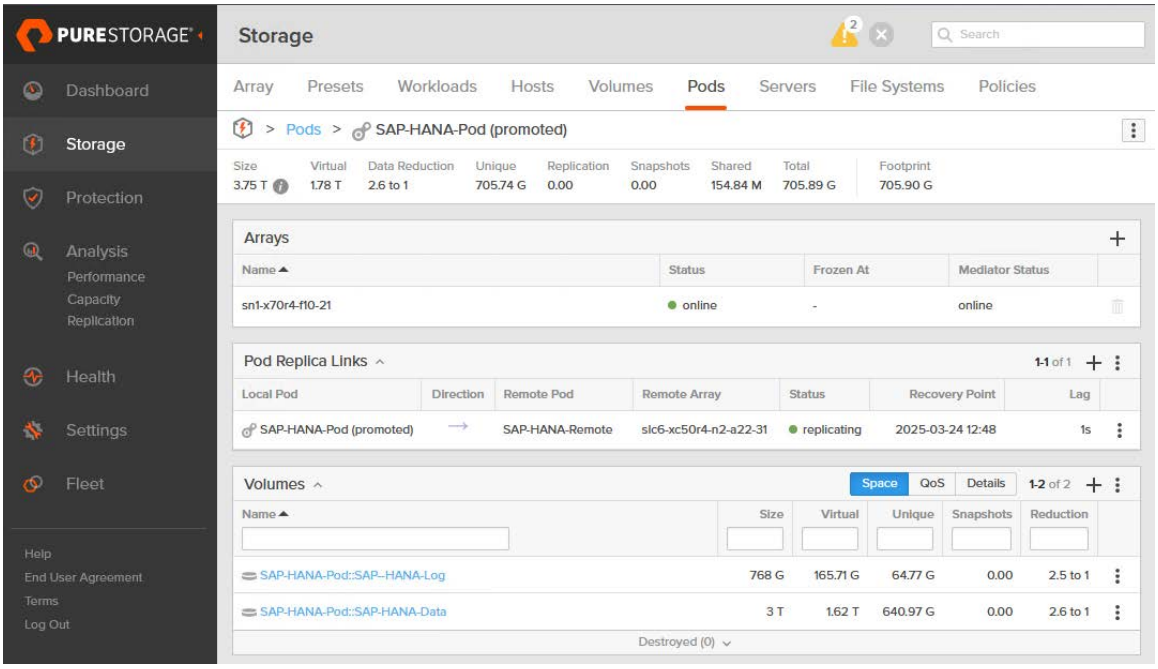

FIGURE 29 Displaying the state of the production/primary array pod.

2. In the **Storage** section on the target/disaster recovery array, click the **Pods** tab, and then observe the state of the **Pod Replica Links**, which will show as demoted.

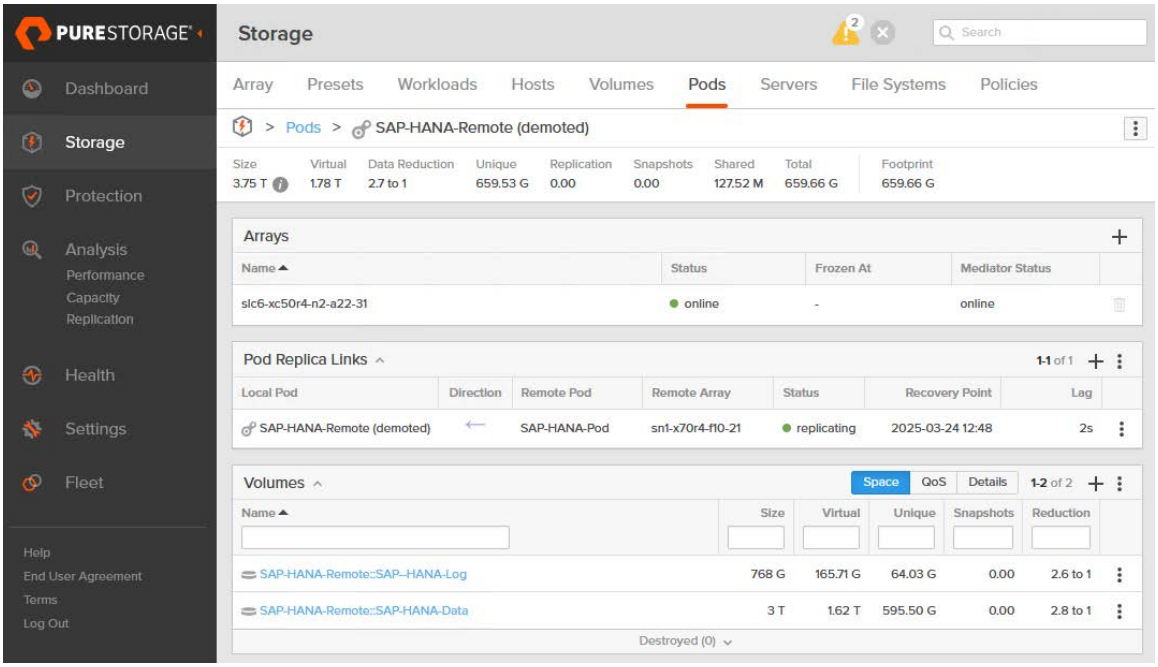

FIGURE 30 Displaying the state of the target/disaster recovery array pod.

3. When ready to perform the failover, navigate to the specific pod in **Pods** under **Storage** on the target/disaster recovery array, click the ellipsis, and then select **Promote**.

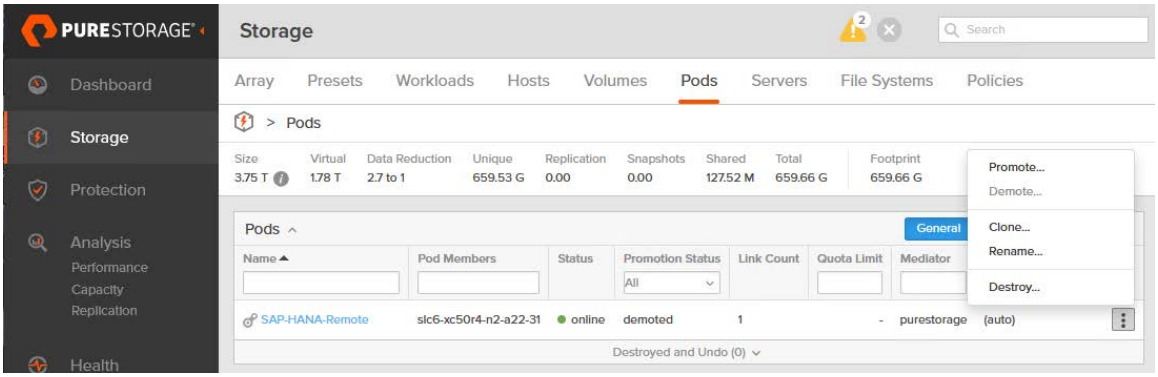

FIGURE 31 Promoting the target/disaster recovery array pod.

4. In the dialogue box, click **Promote**.

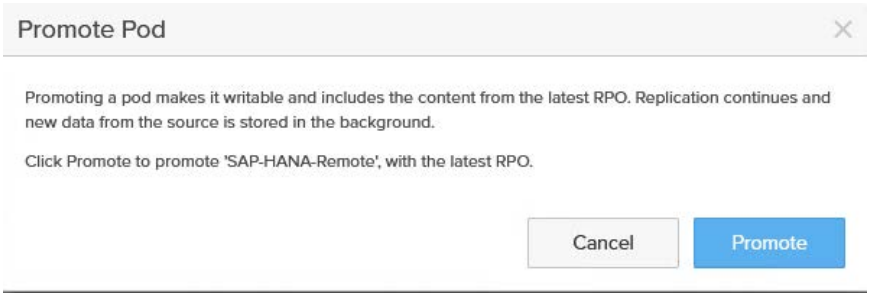

FIGURE 32 The dialogue box to confirm promotion.

5. When the pod has been promoted, it will have a status of **promoted**.

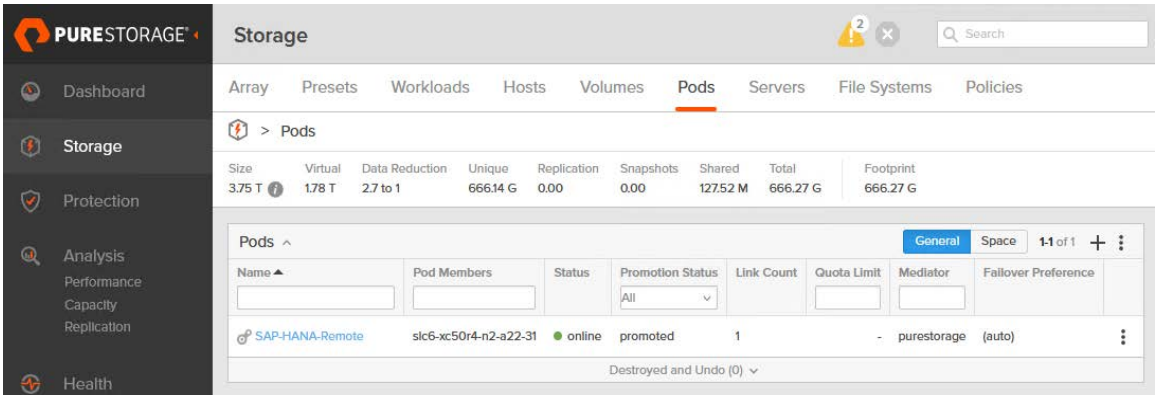

FIGURE 33 The promoted pod on the target/disaster recovery array.

6. On the array to be failed over to connect the data and log volumes to the target systems. attach the pod's volumes to a specific host. This can be done on the **Hosts** tab for a specific host. (If not already connected during the initial implementation).

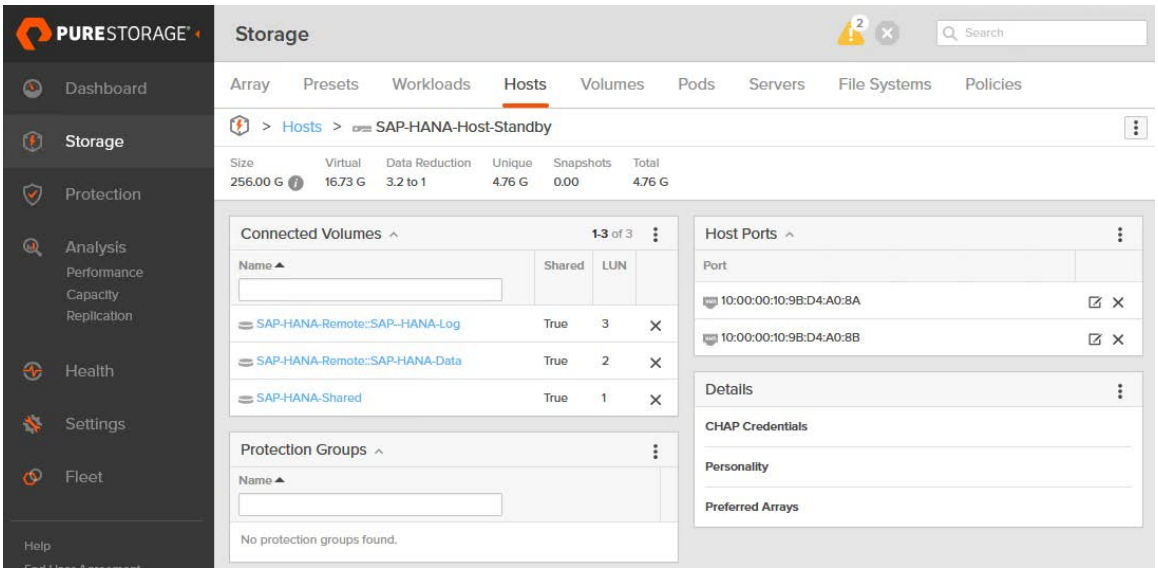

FIGURE 34 Attaching the pod's volumes to a host.

7. Once complete, mount the data and log file systems on the HANA server (scale-up) and then start the instance. To learn more, read SAP HANA Host Failover with ActiveCluster and ActiveDR.

Once the secondary site is operational, monitor the status of the SAP HANA database to ensure there are no issues with data loss or integrity. Ensure that business applications can connect and remain operational with the secondary site online. If there are no issues, resume replication from the primary site by demoting the target array pod.

Resuming Replication to DR site

1. Stop the running SAP HANA database at the DR site

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function StopSystem HDB
```

2. Unmount the SAP HANA volumes

```
umount /hana/data; umount/hana/log
```

3. In the secondary site's Pure Storage array user interface, navigate to the specific pod in Pods under the Storage tab and then select the POD used during the testing. Select the ellipsis next to the pod name, and then select **Demote**.

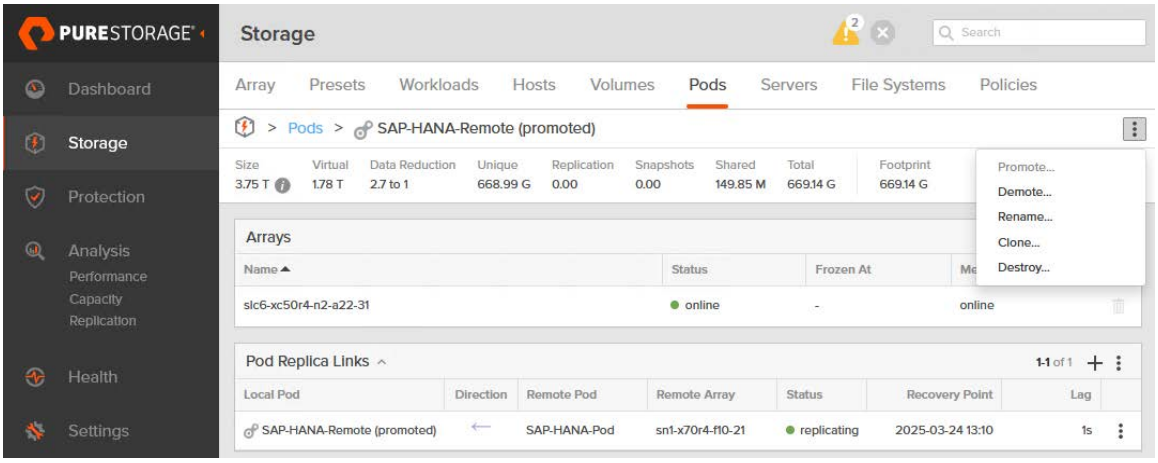

FIGURE 35 Demoting the pod on the secondary array.

4. Check the status of the pod on the primary array, where it will show as **replicating**.

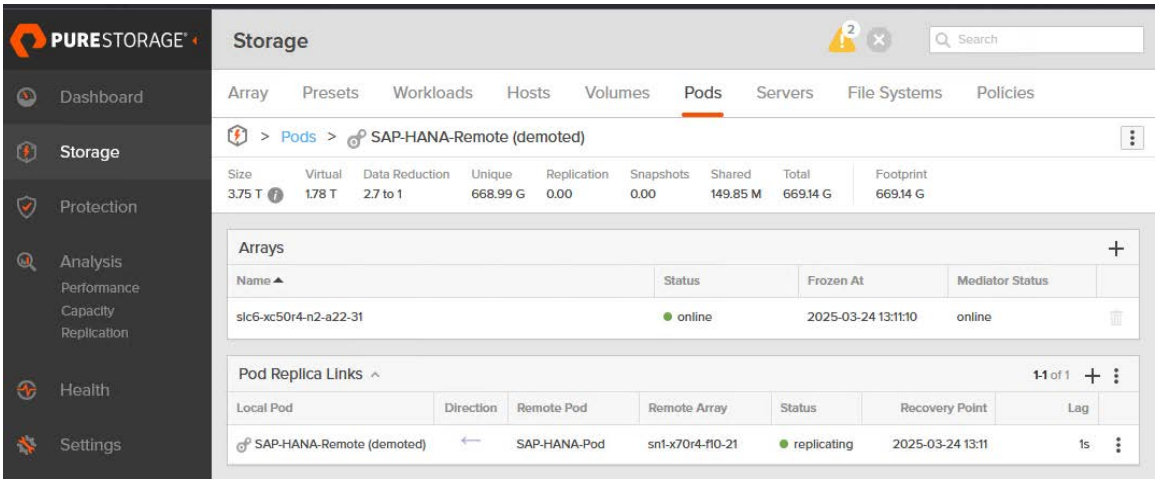

FIGURE 36 Pod at the primary site's status showing as replicating.

Unplanned Failover (Disaster Recovery Scenario)

In the event of an unplanned failover, the secondary site must be brought online.

1. Navigate to the **Protection** tab in the disaster recovery array's Pure Storage user interface to view the status of all arrays. The status will show as **unhealthy**.

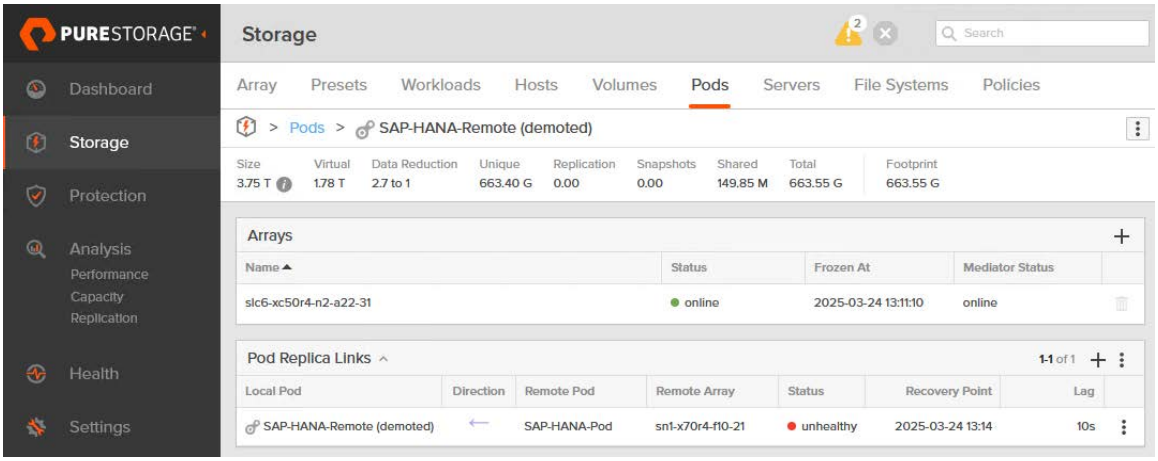

FIGURE 37 Status showing unhealthy.

2. In the **Storage** section on the disaster recovery array, click the **Pods** tab, and then observe the state of the **Pod Replica Links** in a specific pod's details, which will show as **demoted**.
3. Navigate to the specific pod in **Pods** under **Storage** on the disaster recovery array, click the ellipsis, and then select **Promote**.

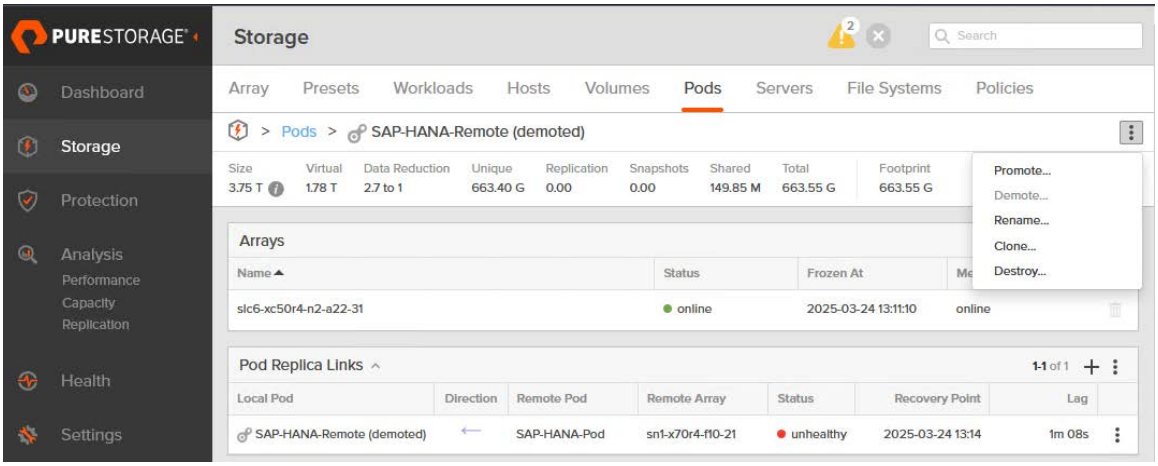

FIGURE 38 Promoting the pod on the disaster recovery array.

4. In the dialogue box that opens, click **Promote**.

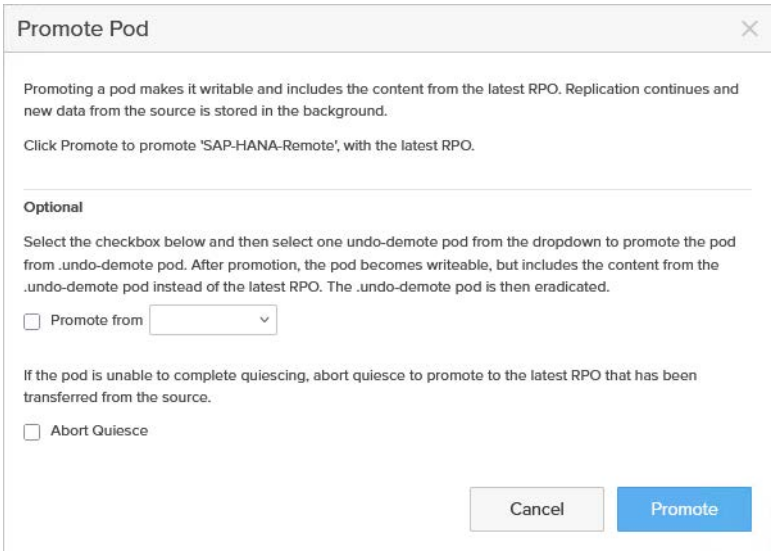

FIGURE 39 Confirming promotion of the pod.

5. When the pod has been promoted, it will have a status of **promoted**.

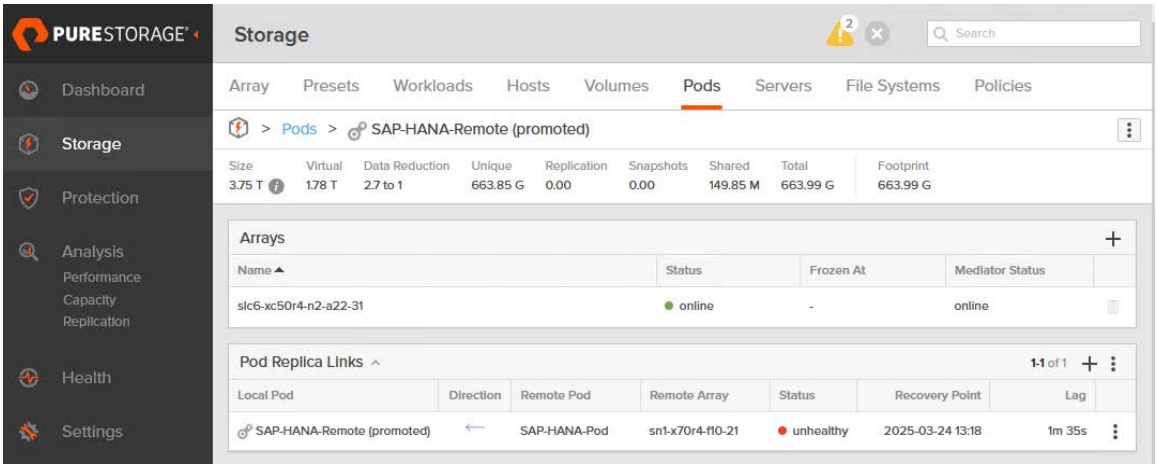

FIGURE 40 The promoted pod on the disaster recovery array.

6. Connect the log and data volumes to the target system(s) (If not already done during the original implementation). Mount the file systems on the HANA server (scale-up).

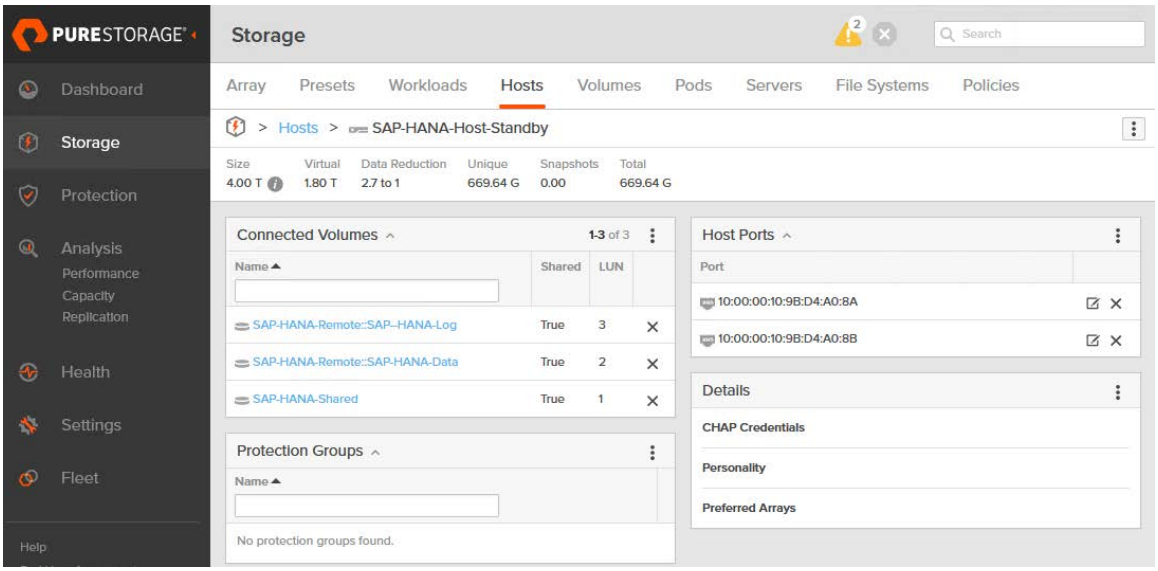

FIGURE 41 Attaching the pod's volumes to a host.

7. Start SAP HANA database

Failback after Disaster Recovery

After verifying that the primary FlashArray has fully recovered and is operational, storage replication can be brought back online to replicate data from the secondary FlashArray back to the primary FlashArray.

1. Stop the primary site's SAP HANA system.

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function StopSystem HDB
```

2. Wait until the instances have stopped on each system by using the following command:

```
/usr/sap/hostctrl/exe/sapcontrol -nr <instance number> -function GetProcessList
```

3. Unmount the HANA data and log volumes on the HANA server

```
umount <path to volume>
```

4. On the primary site's Pure Storage user interface, select **Storage** in the navigation pane, and then select the **Pods** tab.
5. Click the ellipsis for the original source pod, and then select **Demote**.

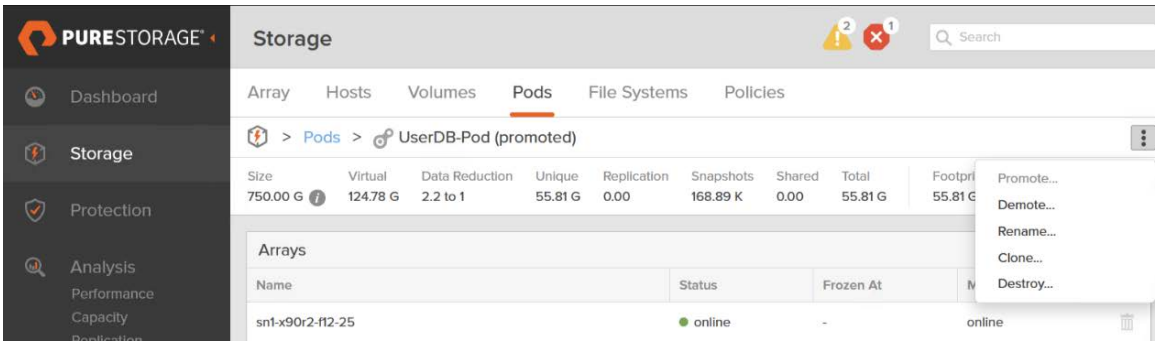

FIGURE 42 Demoting the primary site's pod.

6. From the **Demote** dialog, select the **Skip Quiesce** option and then click **Demote**. **This will make the original source pod volumes read-only and will reverse the direction of replication.**

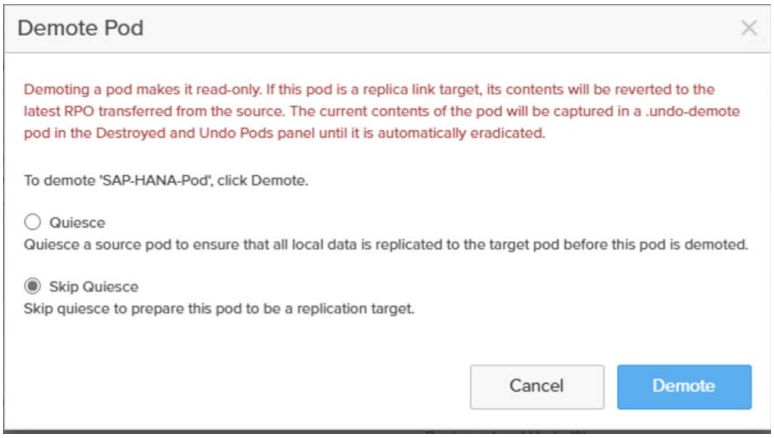

FIGURE 43 Confirming the pod's demotion.

7. Move the SAP instance back to the Primary Site.
8. Monitor the **Status** field in the **Pod Replica Links** group until the **Status** field changes from "baselining" to "replicating."
9. Stop the secondary site's SAP HANA system.
10. Verify the HANA instance has halted.
11. On the primary site's Pure Storage user interface, select **Storage** in the navigation pane, and then select the **Pods** tab.
12. Navigate to the primary site's pod, click the ellipsis, and then select **Promote**.

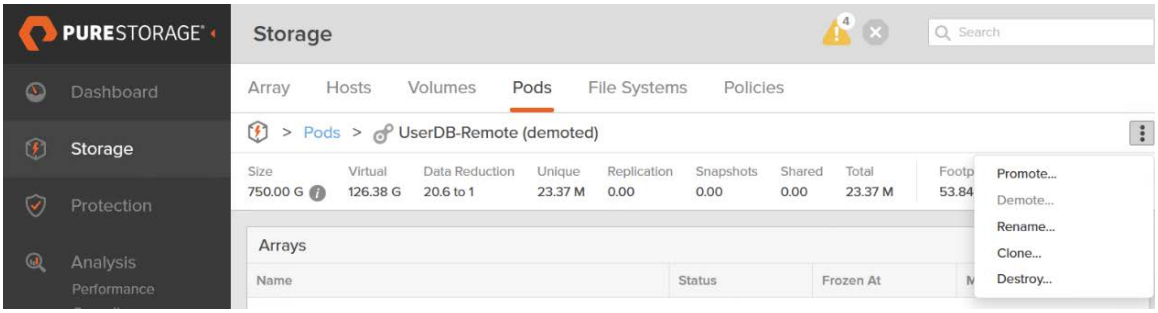

FIGURE 44 Promoting the primary site's pod.

13. From the **Promote** dialogue, select **Promote**. The primary site's pod is promoted and put into a read/write state.

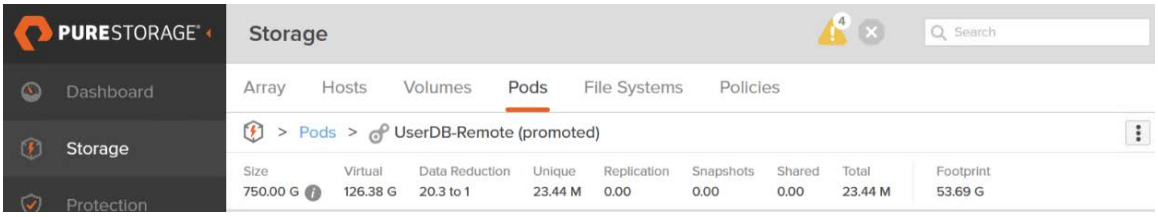

FIGURE 45 Confirming the pod's promotion.

With the primary site's pod promoted and in a read/write state, the volumes can be attached to the primary site's SAP HANA database. The following steps highlight how to do this:

1. If necessary, adjust the applications' database connection settings to point back to the primary SAP HANA system's IP address or hostname.
2. Perform connectivity tests to confirm that applications can access the database without issues.
3. Verify that business applications can connect to the SAP HANA database without issues.
4. On the secondary site's Pure Storage user interface, select Storage in the navigation pane, and then select the Pods tab.
5. Click the ellipsis for the original target pod, and then select **Demote**.
6. From the Demote dialog, select the **Quiesce** option and then click **Demote**. This will make the original target pod volumes read-only and will reverse the direction of replication from the primary site to the secondary site.

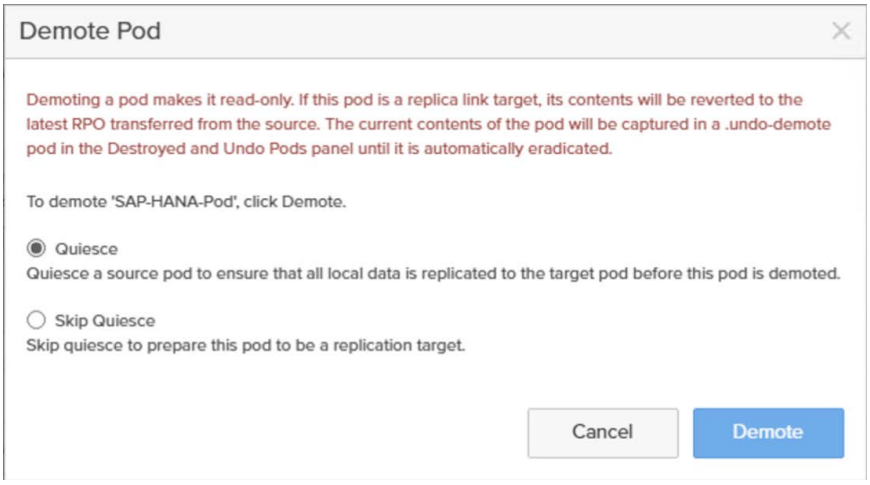

FIGURE 46 Confirming the pod's demotion.

ActiveDR and SAP HANA System Replication

To use ActiveDR alongside SAP HANA system replication, the following must be in place:

- The same SAP HANA version deployed on multiple sites with sites configured for failover
- FlashArray configured at a site for storage replication

System replication is ideal for customers who need high availability for business operations and who want swift recovery from planned maintenance, faults, and disasters. Besides the standard setup, in which a primary system ships all the data to the secondary system, a multitier or a multitarget system replication can be configured, which replicates data changes to more than one failover system.

To create a stronger disaster recovery posture, ActiveDR can be used alongside system replication. The ActiveDR site can be configured to replicate any of the secondary sites configured for system replication, even in a multitarget system. Because the ActiveDR site will be using storage replication instead of system replication, the data backed up to the different sites is protected in a multi-layer disaster recovery strategy.

The ActiveDR deployment is flexible and can be utilized by administrators to replicate the primary SAP HANA system or any of the secondary sites as desired. In failover scenarios, administrators will have two options for recovery available as restore points: a high availability site in one region between the availability zones and another using ActiveDR that helps in recovering the system from catastrophic regional failures.

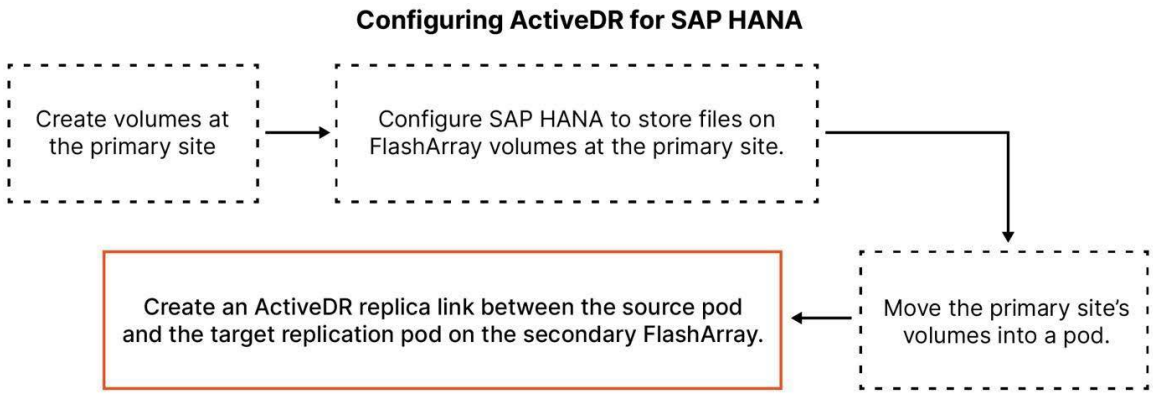

FIGURE 47 Flowchart showing the steps for implementing ActiveDR alongside system replication.

To configure a site for ActiveDR alongside system replication, follow the steps in Deployment and Disaster Recovery Configuration to configure pods for replicating the SAP HANA sites.

Conclusion

ActiveDR provides database and storage administrators with a robust disaster recovery tool for SAP HANA environments. The near-synchronous replication of ActiveDR works together with the powerful in-memory computing of SAP HANA to create a seamless environment protected against data loss. A near-zero recovery point objective is invaluable for businesses relying on SAP HANA; ActiveDR functionality allowing for live failovers and failback testing makes it a valuable combination for protecting mission-critical data with minimal complexity, cost, and operational impact.

In combination with SAP HANA features, ActiveDR provides database and storage administrators with a wide range of choices for performance or failover protection to maintain recovery point objectives, ensuring teams can meet service-level agreements and refine disaster procedures. Other integrations can provide features like virtualization, improved processing, and high availability, making SAP HANA a strong choice for businesses needing a flexible, scalable database for their business.

For more information, visit the Pure Storage website or try the functionality in a test drive.

CM444-01-en 05/25