HPE GreenLake

Technical white paper

HPE Alletra 9000 Storage configuration guidance for Epic EHR deployments

Contents

Executive summary

Delivery of modern healthcare services relies on always-accessible patient-centric information; therefore, electronic health records (EHR) systems are core to most healthcare providers' operations. The growing volume of patient data means increasingly high-performance infrastructure is required to ensure suitable EHR data access while retaining the required levels of data resilience and security. Storage performance is therefore critical.

HPE Alletra 9000 Storage represents an optimal platform for mission-critical target workloads such as Epic Chronicles and Epic Cogito. Indeed, Epic requires all-flash array (AFA) storage for all production copies of a provider's databases, and it is a specific requirement for the InterSystems IRIS (IRIS) database platforms running Chronicles. The HPE Alletra 9000 AFA delivers performance, resilience, and security and enables healthcare providers to make key decisions based on real-time data-driven intelligence.

This technical paper provides guidance for the design and support of HPE Alletra 9000 arrays to meet Epic's Storage Products Guide requirements for its EHR application. It assumes a basic level of familiarity with Epic solution components and access to customer-specific sizing and configuration information in the Epic Hardware Configuration Guide (supplied directly to customers by Epic). This paper also describes some of the architectural elements of the HPE Alletra 9000 family.

The recommendations for the configuration of HPE Alletra 9000 systems are based on testing carried out by HPE engineers, are reviewed with Epic, and are combined with knowledge gained from past Epic deployments. Hewlett Packard Enterprise also continues to work with Epic on testing to ensure that HPE Alletra 9000 arrays continue to meet the performance, scalability, and resilience requirements for Epic systems.

Note
Epic continues to use the HPE Alletra 9000 Storage array within its test lab to test current and future versions of Epic. It is also listed in the Epic Storage Products and Technology Solutions (SPaTS) guide for Epic's databases. Customers can access the guide by contacting their Epic technical advisors or searching for it in Galaxy. (Epic does not formally certify vendor platforms.)

Target audience: The intended audience for this paper includes HPE customers and the HPE partner community and employees. The paper provides test and validation information for HPE Alletra 9000 arrays to healthcare IT professionals involved in the design and management of infrastructure needed to run Epic EHR.

Solution overview

This HPE Alletra 9000 Storage validation project was conducted using an HPE Alletra 9000 series array, multiple HPE ProLiant Gen10 servers for compute with VMware vSphere®, and an additional physical server as the physical backup server.

HPE Alletra 9000 building blocks

HPE Alletra 9000 arrays are designed for mission-critical workloads with high latency sensitivity and availability requirements. Offering unconstrained scalability across consolidated and next-generation applications, the arrays are backed by a 100% availability guarantee delivered through full hardware redundancy and software fault tolerance.

HPE Alletra 9000 arrays are self-regulating and self-managing. Downtime is reduced and performance is maintained through constant monitoring of the rate at which the storage resources are consumed and by running a separate instance of HPE Alletra 9000 OS on each controller node. HPE Alletra 9000 coordinates the multiple controller nodes through a high-speed, full-mesh passive network topology (the high-performance data movement engines of the HPE Alletra 9000 architecture) to deliver a system-wide unified cache that is coherent and fault-tolerant.

Further operational autonomy is provided by HPE InfoSight, which uses powerful AI combined with cloud-based machine learning to predict and prevent issues across the infrastructure stack.

Compaction technologies such as thin provisioning, thin deduplication, compression, and thin reclamation are incorporated and fully automated in HPE Alletra 9000 systems. Thin provisioning allows a volume to be created and made available as a logical unit number (LUN) to a host without dedicating physical storage until it is needed.

HPE Alletra 9000 consists of the basic building blocks listed in Table 1.

Table 1. HPE Alletra 9000 building blocks

Basic component	Description
 **Four-way storage base**	The storage base building block contains a chassis that can house the controller nodes. This is a 4U chassis that includes four node slots and 48 SFF drive slots.
 **HPE Alletra 9000 controller node**	The controllers compute element of HPE Alletra 9000 series arrays is responsible for the connectivity and functionality of the storage array. HPE Alletra 9000 arrays may be configured with either two or four controller nodes. Each controller SKU includes two nodes and locking power cords.
 **Power supplies**	An HPE Alletra 9000 array requires two 1700W power supplies per node pair and supports AC or DC power. Each power supply includes a power supply with fan and battery.
 **NVMe SSD (SFF) drives**	The supported drives are NVMe SSDs. The drives offered are standard FIPS 140-2 validated self-encrypting drives and FIPS encrypted drives with TAA compliance.
 **Four-port 16/32 Gb/s Fibre Channel HBA** **Four-port 10/25 Gb/s iSCSI HBA** **Four-port 10GbE Base-T iSCSI HBA**	Adapters provide the front-end connectivity options for the array. Supported adapters include a four-port 16/32 Gb/s Fibre Channel adapter, a four-port 10/25 Gb/s iSCSI adapter, and a four-port 10GbE Base-T iSCSI adapter.

Storage environments and database servers

The HPE Alletra 9000 AFA enables healthcare providers to make decisions based on real-time data-driven intelligence. Storage performance is critical and the HPE Alletra 9000 represents an optimal platform for mission-critical target workloads such as a Chronicles database. AFA storage is recommended for the healthcare provider's production databases and is a specific requirement for IRIS.

Healthcare organizations typically manage multiple data environments. In addition to Epic software, examples may include (at a high level):

- Microsoft SQL Server or Oracle® database server hosting business reporting tools, such as Clarity, the Epic analytical reporting module
- Virtual machine (VM) data
- Web binary large object (BLOB) data

In addition to receiving performance specifications required for production databases, customers may also receive guidance on the number of volumes per storage controller. Epic classifies the individual workloads by data characteristics and assigns them to specific disk pools.

Each disk pool in the data environment is configured to meet unique capacity, performance, and availability requirements. For example, Epic Disk Pool 1 supports the production database and has a workload different from the Epic Disk Pool 2 reporting database.

Table 2 shows an example layout of an Epic environment utilizing separate storage pools for production and reporting on HPE Alletra 9000 Storage. Data for disaster recovery should always be stored on separate media. Automated backup of the database using freeze and thaw commands can be scheduled with HPE Alletra 9000 Virtual Copy scripts, and hardware snapshots can be used for backups by supported data protection applications with HPE StoreOnce. Zerto, a Hewlett Packard Enterprise Company can be used for data protection in a virtualized environment.

Table 2. Example of Epic disk pools per primary and disaster recovery arrays

Primary array	Data	Disaster recovery array	Data
Pool 1	Production ODB databases	Pool 3	Disaster recovery ODB databases
Pool 1	Production journals	Pool 3	Disaster recovery journals
Pool 1	Production application files	Pool 3	Disaster recovery application files
Pool 1	Production support (SUP)		
Pool 1	Release testing (REL)		
Pool 1	Test, build, and training	Pool 3	Disaster recovery test, build, and training
Pool 1	Nightly backup	Pool 3	Disaster recovery nightly backup
Pool 2	Reporting databases	NAS 2	Disaster recovery archive and journal
Pool 2	Reporting journals	NAS 2	Disaster recovery BLOB data
Pool 2	Cogito SQL Servers		
NAS 1	Archive data and journal		
NAS 1	BLOB data		

Epic Reporting Analytics tier databases use a SQL reporting application. The reporting environment also includes a data warehouse built on a SQL database. Database storage layout recommendations are provided in the Cogito section of the Database Storage Layout Recommendations document provided by Epic on Galaxy.

Storage workload

The workload of each database server varies depending on the role and the specific data environment it is supporting. The Disk Pool 1 database files typically incur continuous random I/O patterns with large bursts of writes every 80 seconds. In contrast, the Disk Pool 2 Clarity database server has a different workload characteristic. The workloads include:

- **Extract, transform, load (ETL) (reporting) workload** — Performs mostly sequential writes with some random reads.
- **Reporting workload** — Generates intensive random read queries at a rate relative to the volume of reporting volume generated by the health organization. Demand can increase at certain times to meet reporting requirements.

Epic performance requirements include criteria such as:

- Read/write performance (IOPS)
- Read/write throughput (MB/sec)
- Write operations per 80-sec write burst
- Write volume per 80-sec burst (MB)
- Minimum SAN write cache requirement (MB)
- Average read/write response (ms)

Note

SQL Server reports are drawn from data extracted from the reporting IRIS database.

Highly available disaster-tolerant solution

It is essential that critical healthcare data be included in business continuity planning. Hewlett Packard Enterprise recommends the use of high availability storage systems to mitigate hardware component failure. This recommendation extends from basic hardware such as redundant power supplies, to networking such as multipath networking. At the storage-node level, nondisruptive upgrades and storage expansion are standard requirements. For this reason, Hewlett Packard Enterprise highly recommends that for medium-to-large mission-critical healthcare environments, the HPE Alletra 9000 system is deployed with four controller nodes. However, two controller nodes are sufficient for small-to-medium sized environments.

Epic customers typically implement mirroring of the IRIS database using built-in InterSystems mirroring technologies. The InterSystems IRIS Data Integrity Guide describes how mirroring on secondary servers maintains a copy of selected databases as they are updated on a primary server. This approach can provide a level of disaster recovery. Array-based replication solutions complement database protection and are a requirement for data not held in the IRIS database.

HPE Alletra 9000 arrays can provide an additional level of disaster-tolerant high availability within a data center and between data centers. Solutions can be designed to meet a defined recovery point objective (RPO) and recovery time objective (RTO). HPE Support Center and HPE Alletra 9000: Configuring and managing data replication using Remote Copy provide details on implementing a highly available disaster-tolerant solution using HPE Alletra 9000 systems and best practices recommendations.

HPE Alletra 9000 configuration guidance

This section provides recommendations that are applicable to HPE Alletra 9000 systems that support Epic databases.

HPE Alletra 9000 can be configured with either two or four controller nodes. The fundamental configuration rule is to have a completely balanced system where each node pair is configured with the same type and number of drives, and where all nodes are configured with the same types and numbers of HBAs.

Host adapters

HPE Alletra 9000 Storage requires at least one host adapter per node because it does not have any built-in host ports. Each node must have the same number and types of HBAs. This means that in a four-node system, all four nodes must have the same HBA configuration.

Drives

HPE Alletra 9000 Storage supports only NVMe SSDs and a maximum of 240 drives per system.

- The minimum drive quantity per node pair is eight.
- The maximum drive quantity per node pair is 120.
- The minimum drive upgrade increment is two drives per node pair.
- A new HPE Alletra 9000 array must be configured with all drives of the same capacity.
- HPE Alletra 9000 Storage supports only RAID 6 configurations (6:2 to 10:2).

Considerations when upgrading drives include the following guidelines:

- The minimum upgrade drive quantity is two drives per node pair.
- When upgrading capacity in a four-node configuration, add equal quantities and types of drives to both node pairs.
- If it is not run automatically, the best practice is to manually run Autonomic Rebalance (also known as “tunesys”) after adding drives.

Drive capacities

Table 3 shows the total usable capacity available per drive. This does not take into consideration the overhead for sparing and RAID. The size of each chunklet is 1 GiB (1024 MiB).

Table 3. Number of chunklets per drive for each drive size

Drive type	Capacity	Chunklets (GiB)
SSD	1920 GB	1788
SSD	3840 GB	3576
SSD	7680 GB	7152
SSD	15360 GB	14305

NVMe SSD drive mixing rules

HPE Alletra 9000 Storage supports NVMe SSDs of different capacities. A new array must be configured with all drives of the same capacity. Mixing drives of different capacities after an upgrade is allowed. However, a system with drives of an identical size will have its capacity utilization and therefore its I/O more evenly distributed among the drives. For this reason, even after an upgrade, mixing drive sizes in a system is not recommended.

Another reason it is preferable to have a single drive size is that mixing drives of the same type, but different capacities, can lead to extra capacity being used for creating spares. The total spare capacity is spread across all SSDs in proportion to their size. The spare capacity with the minimal setting is equal to the largest capacity drive.

Example: Impact of mixing drive sizes on spare capacity

- Initial system: Eight 3.84 TB SSDs, spare capacity = 3.84 TB
- Upgrade: Add four 7.68 TB SSDs, new spare capacity = 7.68 TB
- The upgraded array has eight 3.84 TB SSDs and four 7.68 TB SSDs, spare capacity = 7.68 TB
- The addition of four 7.68 TB SSDs increases the spare requirement from 3.84 TB to 7.68 TB. Therefore, the addition of four 7.68 TB SSDs does not result in an increase of 30.72 TB (4 x 7.68 TB) but only of 26.88 TB, because 3.84 TB of the new capacity is used as spare space.

By adding similar capacity drives to the array, allocating additional spare space can be avoided. Often, given the cost of NVMe SSDs, such restrictions are difficult to follow. In the cases where mixing NVMe SSD capacities in an HPE Alletra 9000 array is necessary, you must follow these rules:

- When adding more NVMe SSDs, make sure that the configuration remains balanced between node pairs.
- Minimum upgrade drive quantity is two drives per node pair.
- If the same capacity is not available, it is best to use the closest capacity available.

Note

To review all HPE Alletra 9000 best practices recommendations, see the HPE Alletra 9000 technical documentation available in the HPE Support Center.

Common provisioning group

A common provisioning group (CPG) is a virtual pool of logical disks that allows volumes to share resources and to allocate space on demand. It is the layer that defines the logical disk creation characteristics such as RAID type, set size, disk type for chunklet selection, total space warning, and limit points. A thinly provisioned virtual volume (TPVV) created from a CPG automatically allocates space on demand by mapping new regions from the logical disks associated with the CPG. A CPG allows virtual volumes (VVs) to share the resources of the CPG.

Best practices for creating and configuring CPGs include:

- When creating CPGs, accept defaults according to performance and capacity requirements.
- The number of CPGs should be kept to a minimum.
- Do not set growth limits on CPGs.
- If using clones for Epic copies of the ODB, place all VVs in the same CPG. VVs will not dedupe across different CPGs.

Note

For a complete list of precautions regarding CPGs, see the HPE Alletra 9000 Storage Concepts Guide available in the HPE Support Center. For a list of all options and specifiers available for the `createcpg` command, use the HPE Alletra 9000 CLI Help or review the HPE Alletra 9000 Command Line Interface Reference.

RAID level
HPE Alletra 9000 uses RAID 6 as the default RAID level. Depending on the number of drives in the configuration, the array manages the number of drives used for data and parity, for example, 8D+2P or 10D+2P.

Virtual volumes

VVs are exported as LUNs to hosts. VVs are the only data layer visible to the hosts. You can create clones (full copies) or snapshots (virtual copies) of virtual volumes. Clones remain available if the original base volume becomes unavailable. VVs use storage space provided by the CPGs. A clone duplicates all the data from a base volume to a destination volume.

Best practices for creating VVs include:

- Place all Epic VVs and clones in the same CPG to take advantage of deduplication. Enable dedupe and compression (DeCo). All HPE Alletra 9000 performance testing with Epic was conducted with DeCo enabled.
- Epic best practices state that the database volume group should have between 4 and 16 LUNs. All volumes for the ODB should have their own app volume set. Select **InterSystems IRIS** as the application type depending on the version of Epic ODB you are using. All VVs should also be placed in a virtual volume set for clone and snapshot purposes.

HPE Alletra 9000 Adaptive Data Reduction

HPE Alletra 9000 Adaptive Data Reduction (ADR) technologies include HPE Alletra 9000 compression, deduplication, thin provisioning, thin conversion, thin persistence, and thin copy reclamation. These technologies achieve ADR by leveraging built-in hardware capabilities and HPE Alletra 9000 Express Indexing.

Compression is ideal for data that does not have a high level of redundancy. Compression works by looking inside data streams for opportunities to reduce the size of the actual data. The inline lossless compression algorithm of HPE Alletra 9000 is specifically designed to work on a flash-native block size to drive efficiency and performance. It leverages a series of technologies to offer the largest possible space savings. HPE Alletra 9000 compression is based on the LZ4 algorithm. This algorithm is very efficient; it offers not only good compression ratios and compression speed, but also exceptional decompression speed, which is crucial to storage performance.

Data sets that are good candidates for compression include:

- **Databases** — Most databases do not contain redundant data blocks but do have redundant data within blocks so they can benefit from compression.
- **VM images** — VMs where the application data size far exceeds the operating system binaries size might not yield significant deduplication savings but can benefit from compression of the application data.
- **Virtual applications** — Application virtualization environments hosting the Epic Hyperspace client application can achieve excellent compression ratios.
- **Virtual desktop infrastructure (VDI)** — Client virtualization environments with hosted nonpersistent desktops can also achieve excellent compression ratios.

Deduplication

Deduplication works by looking inside data streams for duplicate blocks of information. Deduplication has become standard with disk-based backup because of a high degree of data redundancy and less emphasis on performance.

Data sets that are good candidates for deduplication include clones of databases. Epic requires copies of their ODBs such as support, release, and test. These clones dedupe well.

Data with a low level of redundancy should not be stored on deduplicated volumes. Data sets that are not good candidates for deduplication include:

- **Imaging data** — Most images do not contain redundant data blocks, such as the BLOB data.
- **Deduplicated data** — Data that has already been deduplicated on the host is not compacted further by storage deduplication.
- **Compressed data** — Compression creates a stream of unique data that does not benefit from storage deduplication.
- **Encrypted data** — Host or SAN encryption also results in a stream of unique data that does not benefit from storage deduplication.

Virtual Copy snapshots initially consume no space and increase in size when the parent volume is modified with updated data.

HPE Alletra 9000 Adaptive Data Reduction combines deduplication and compression to help maximize space savings. For data reduction volumes, the data is first checked for duplicate blocks, and the unique blocks are then compressed. Note that the dedupe occurs across all volumes in the CPG.

As a best practice for working with deduplication and compression, use dedupe and compression for IRIS databases (Disk Pools 1, 2, and 3). Part of pools NAS 1 and 2 contain Epic BLOB data, which is made up of images that do not compress or dedupe. Do not use DeCo for Epic BLOB data.

Sizing for performance, capacity, and availability

Use the Epic Hardware Configuration Guide to help define the storage requirements. The unique workloads of each database server must be accounted for when sizing for performance, capacity, and availability. Additional back-end workloads such as Remote Copy, Virtual Copy, RAID overhead, and ADR could impact the front-end performance IOPS and latency. Healthcare organizations might also have unique requirements that could increase the amount of required storage and performance from the array.

With the workload and capacity data, a baseline sizing exercise should include the use of HPE sizing tools for the HPE Alletra 9000 system. These tools are available for HPE partners to help size and configure cost-effective and performance-optimized HPE Alletra 9000 solutions. The tool converts a set of basic capacity and performance metrics (requirements) into a ready-to-order bill of materials (BOM) list.

The sizing tool has several modifiable inputs for performance and sizing:

- Drive count
- Model of drive
- RAID type and set size
- High availability
- Workload
- Block size
- Read/write ratio

HPE Alletra 9000 clones and snapshots

The integrity and recoverability of clones and snapshots are critical. These copies are typically restored to an alternate server and then the integrity of the copy database is validated. Clones or snapshots used for backups are also presented to a backup server to offload the integrity and copy operations.

HPE Alletra 9000 supports two technologies to provide users with copies of HPE Alletra 9000 base VVs:

- Physical copies (clones)
- Virtual copies (snapshots)

These technologies allow you to access a point-in-time (PIT) copy of a base VV and to read data from or write data to that PIT copy. You can meet RPO and RTO requirements by using these technologies.

HPE Alletra 9000 clone and snapshot technologies enable you to use scripts to call freeze/thaw operations and automate the process. By using these technologies, you can create re-syncable clones/snaps for production database data volumes. As an example, a typical set of create, refresh, and update steps follow this order:

1. Create re-syncable clones/snaps for production ODB and journal volumes.
2. Create re-syncable snaps for SUP.
3. Refresh the clone/snap copy of the production ODB.
4. Refresh the SUP copy.
5. Create a virtual copy of the clones for backup purposes.
6. Update the virtual copy of the clones after refreshing the clone copy.

Figure 1 shows a data protection strategy based on HPE Storage clone and snapshot technologies. It is typical to have 5 to 20 copies (clones or snapshots) of a database. This strategy includes copies for the production databases, support databases, release environment, test, and training environments.

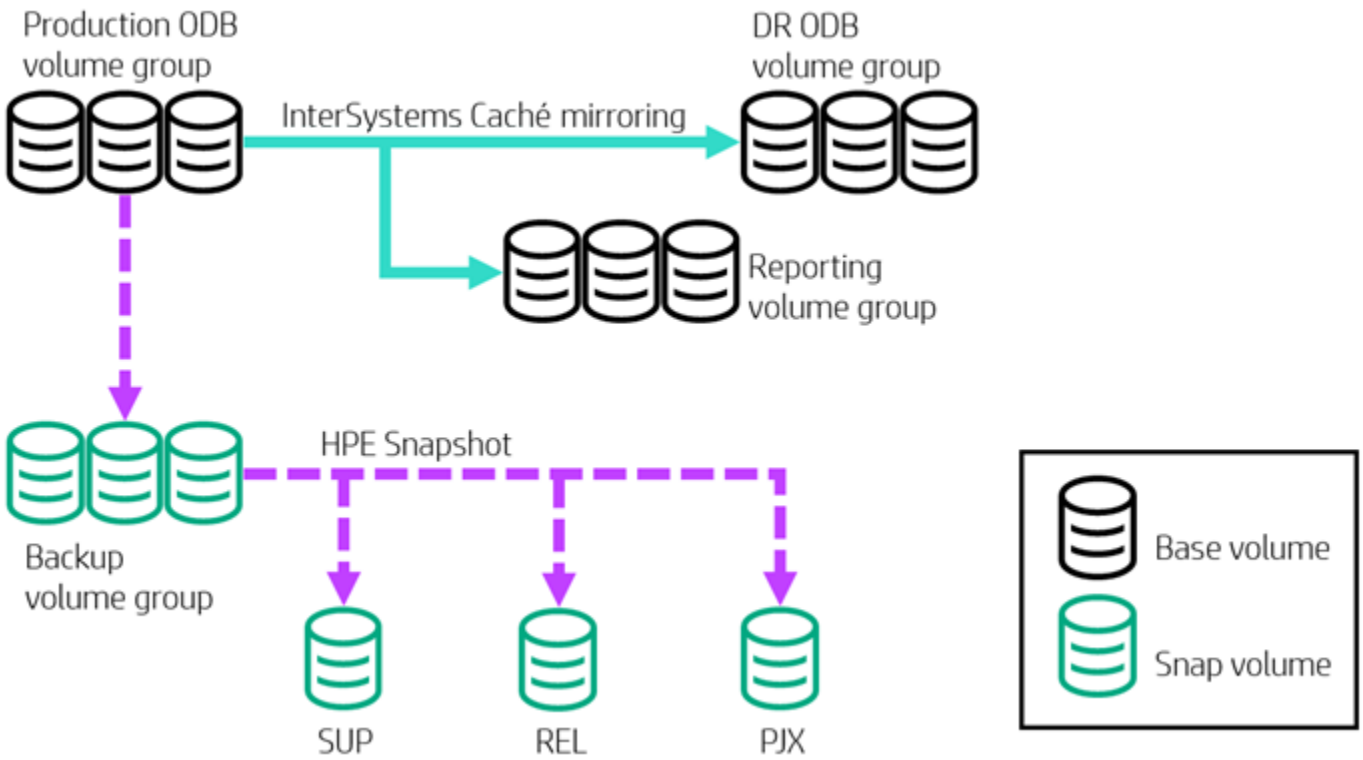

Figure 1. Example of physical and virtual copies used for database protection

HPE Alletra 9000 physical copies

A physical copy is a full copy of a volume that duplicates all the data from one original base volume to a destination volume. Any change to either volume causes it to lose synchronization with the other volume, which can be corrected by resynchronizing the two volumes. The destination volume can be resynchronized with its parent volume by using a saved snapshot so that only the changes since the last copy or resynchronization need to be copied. You must specify a snapshot of the parent saved during the creation of the physical copy because it is required for the resynchronization operation. After the resynchronization is complete, a new snapshot for the next resynchronization operation replaces the old snapshot.

You can create and manage physical copies in groups to reduce the number of management tasks. You can create a consistent group of physical copies from a list of VVs and then group the physical copies into autonomic groups that are managed as a single physical copy.

Physical copies can be online or offline:

- With online physical copies, the destination volume is created automatically and can be exported immediately to the host because cloning happens in the background at a lower priority, so the performance of active I/O is not affected. Because the copy is a background process, the duration of the cloning task depends on how busy the array is.
- Offline physical copies require a destination volume with a user space size at least as large as the user space of the base volume being copied. Offline physical copies cannot be exported until the data copying is complete.

HPE Alletra 9000 physical copy resynchronization

Another efficient way of managing the data in an HPE Alletra 9000 environment with Virtual Copy is by updating the clones. The `createvvcopy -r` command updates a VV with new data. By using this command, you avoid having to remove an exported VLUN, performing a disk rescan, and exporting it to mount the updated disk. Although you can use the command successfully in many environments, review HPE guidance and limitations:

- Use the `setuseracl` command for users that require permission to run the `createvvcopy -r` command on specific volumes.
- Although the impact is minimal, Hewlett Packard Enterprise recommends that you schedule the `createvvcopy -r` command during low I/O periods.
- After creating the copy, you can save or delete the PIT snapshot used to create the clone. If you save the snapshot, the target volume becomes a clone that can be resynchronized later with the parent. If you do not save the snapshot, the data is copied from the parent to the target volume, but the target does not become a clone. Rather, the target remains a base volume and cannot be resynchronized with the parent.
- Copying data from a large parent volume can take several minutes to hours. You cannot export the clone volume to a host until copying is complete.
- If necessary, you can stop the cloning process while it is running. If you stop the cloning process, the target volume reverts to being a base volume and will contain incomplete data.

The `createvvcopy -r` command:

1. First verifies if the user has permission to run the `createvvcopy -r` command on the specified VVs.
2. Checks if there is a snapshot associated with the parent volume.
   - If a snapshot was not specified when the clones were created, the resynchronization command will fail.
   - If a snapshot is available, the command will progress.

HPE Alletra 9000 Virtual Copy snapshots

HPE Alletra 9000 Virtual Copy is the snapshot implementation used to provide a PIT virtual copy of a virtual volume. This process ensures that you can always obtain the original data if a problem occurs when data is updated on a VV. HPE Alletra 9000 Virtual Copy snapshots are always thin, nonduplicative (only one copy of the change blocks is kept and then referenced), and without reservation.

A Virtual Copy snapshot initially consumes no space and increases in size as it records changes to the original volume. This allows an earlier state of the original volume to be re-created by starting with the current state and rolling back all the changes that have been made since the Virtual Copy snapshot was created.

You can create and manage Virtual Copy snapshots in groups to reduce the number of management tasks. You can create a consistent group of Virtual Copies from a list of VVs and group Virtual Copies into autonomic groups that are managed as a single Virtual Copy. A Virtual Copy snapshot can belong to any CPG and any RAID level, giving you flexibility and efficiency in terms of storage consumption.

When snapshot backups are intended for use as a recovery source for the production database, the backup method must provide database consistency. To ensure the integrity of the snapshot, IRIS provides methods to freeze writes to databases while the snapshot is created. Only physical writes to the database files are frozen during the creation of the snapshot. This allows user processes to continue performing updates. The snapshot is typically a PIT capture of all file systems in use by the system. After writes are thawed, the snapshot may be used as a backup source and then either rejoined to production or left online as a warm backup (depending on the specific technology used).

HPE Alletra 9000 Virtual Copy snapshot restore and update

You can use HPE Alletra 9000 Virtual Copy snapshots to restore a parent VV. The VV is restored to match the PIT data when the snapshot is taken. Such an operation gives you the option to recover data if a file is deleted or if data is inconsistent.

You can restore the data by using the HPE Alletra 9000 CLI. You can perform this operation using the online option.

Another efficient way of managing the data in an HPE Alletra 9000 environment with Virtual Copy is by updating the snapshots. The `updatevv` command updates a snapshot VV with a new snapshot. By using this command, you avoid having to remove an exported VLUN, performing a disk rescan, and exporting it to mount the updated disk. Although you can use the `updatevv` command successfully in many environments, review the HPE guidance and limitations:

- Use the `setuseracl` command for users that require permission to run the `updatevv` command on specific volumes.
- Although the impact is minimal, Hewlett Packard Enterprise recommends that you schedule the `updatevv` command during low I/O periods.
- The `updatevv` command is:
  - Not used with a physical copy
  - Not used on Remote Copy snapshots
  - Not used on snapshots that have child snapshots
  - Not used on snapshots that have a retention time

The following steps provide a high-level summary of the `updatevv` command:

1. The command first verifies if the users have permission to run the `updatevv` command on the specified VVs.
2. If the snapshot is read/write (RW):
   - The snapshot's read-only (RO) parent is replaced by a new RO snapshot of the same name.
   - A new RW snapshot is created of the new RO snapshot.
3. If the specified snapshot is an RO snapshot, the command only updates the requested snapshots. Best practices for working with snapshots and clones include:
   - Leverage physical copies (clones) when the use of the clone requires high random read and write I/O and will be retained for several months, such as in a release testing and project environment.
   - Leverage virtual copies (snapshots) for copies of the ODB that are refreshed frequently such as a production support environment or for nightly backups.
   - Epic typically requires eight virtual volumes for the production ODB. The virtual copies (snapshots) and physical copies (clones) of the ODB should be placed in their own virtual volume set.
   - When issuing the snapshot or clone refresh command, issue it against the virtual copy/physical copy virtual volume set and not each VV individually. This reduces the time and the number of commands issued to the array. For example, place the SUP snapshots in a virtual volume set named "EpicSUPsnapset" or place the REL clones in a virtual volume set named "EpicRELcloneset" and issue the `updatevv` command against those sets.

Note

For detailed information on using the HPE Alletra 9000 CLI `updatevv` and `createvvcopy` commands, see the HPE Alletra 9000 Command Line Interface Reference.

HPE Alletra 9000 Virtual Lock — immutable snapshots

HPE Virtual Lock software provides ongoing protection for mission-critical data by safeguarding individual virtual volumes and virtual copies (snapshots) from encryption and accidental or unauthorized deletion. HPE Virtual Lock is an integral part of the HPE Alletra operating system and does not require any software or agent installation at the host level.

There are two ways to create an immutable snapshot with HPE Virtual Lock through the Data Services Cloud Console (DSCC): manually and with Protection Policies.

Creating an HPE Virtual Lock immutable snapshot manually

To create a manual HPE Virtual Lock immutable snapshot, log in to the DSCC and go to **Block Storage**. Select **Storage** from the menu, and then select the Volume Set. Click **Take Snapshot**. The following window will appear, as shown in Figure 2.

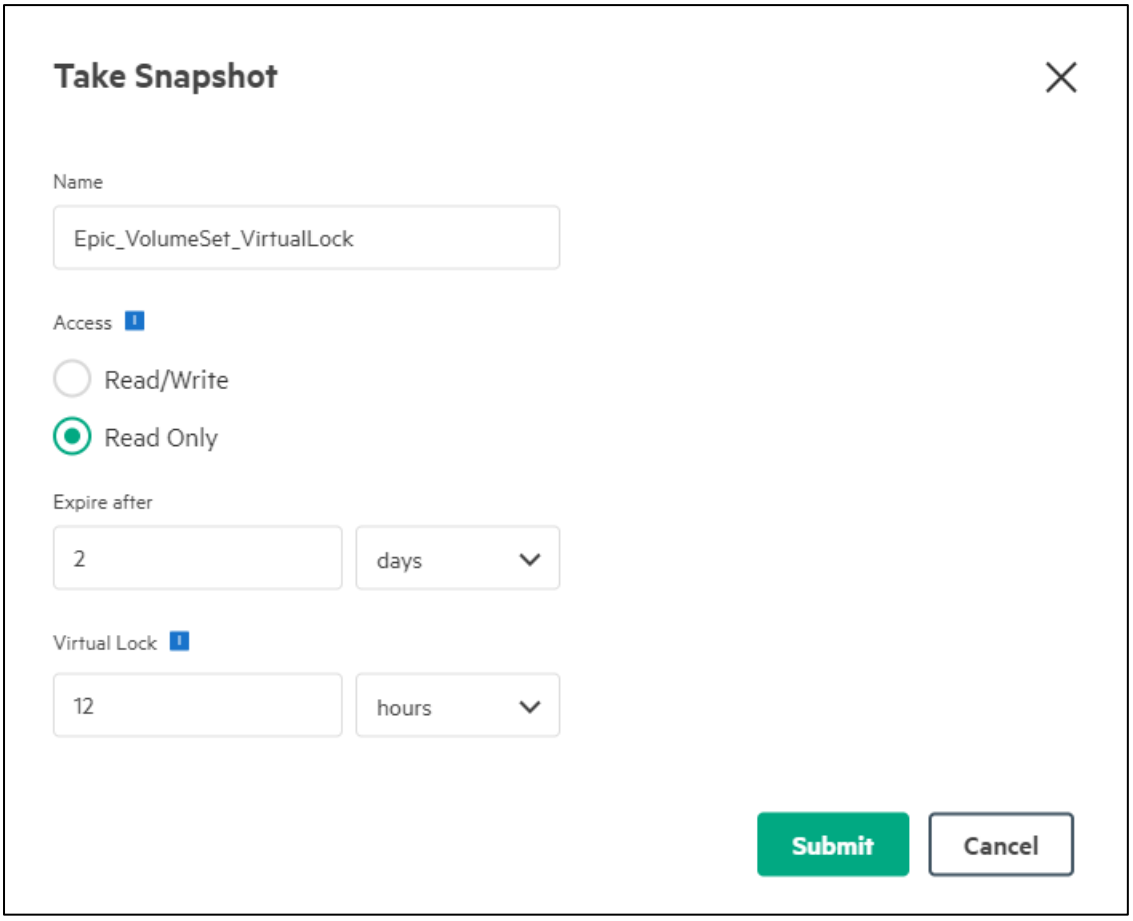

Figure 2. The "Take Snapshot" window in the DSCC

Give the HPE Virtual Lock snapshot a name, set it to **Read Only**, and fill out the **Expire after** and **Virtual Lock** fields. The **Expire after** fields are used to assign a retention period for the snapshot. The **Virtual Lock** fields are used to assign an immutable period. The **Virtual Lock** setting must be shorter than the **Expire after** setting because the immutable snapshot cannot be deleted.

Creating an HPE Virtual Lock immutable snapshot Protection Policy

To create a Protection Policy that includes immutable snapshots using HPE Virtual Lock, log in to the DSCC and go to **Block Storage**. Select **Storage** from the menu, and then click the Volume Set name. Next, click **Policies** and then the "**+**" symbol to add a new Protection Policy. In the pop-up window, select **Scheduled Snapshot**, and then click **Next**. The following window will appear, as shown in Figure 3.

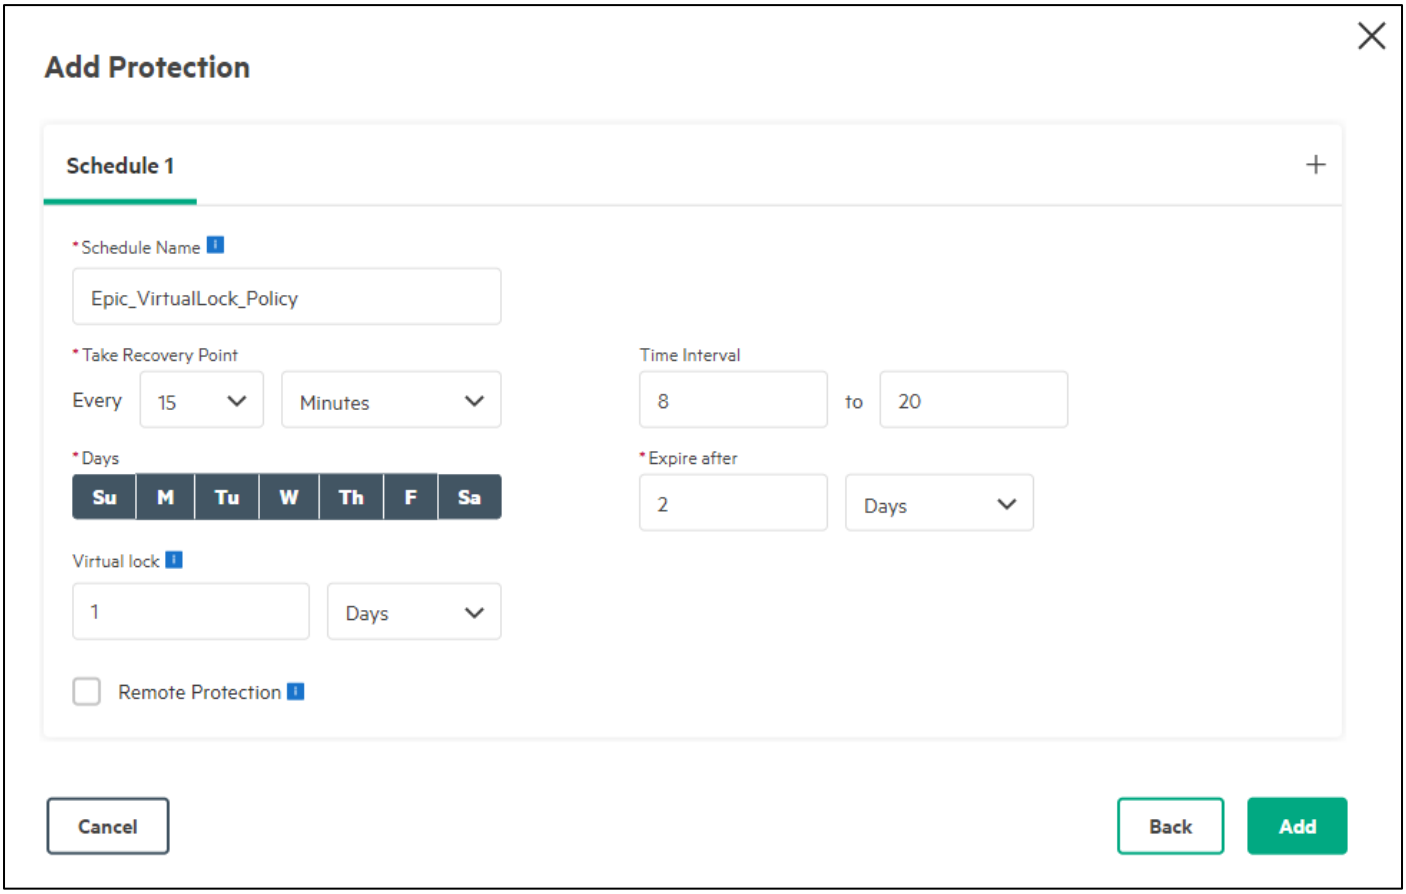

Figure 3. The "Add Protection" window in the DSCC

Give the Protection Policy a name. The **Take Recovery Point** setting determines how often snapshots are taken, the **Time Interval** setting defines the time window during the selected days to take snapshots, the **Expire after** setting defines how long the snapshot is kept before it is deleted, and the **Virtual lock** setting defines how long the snapshot is immutable. **Remote Protection** can be used to create coordinated snapshots on both the local and remote volume sets on the local and remote storage arrays.

Note
If additional Protection Policies are needed, clicking the "**+**" at the top-right of the window will add a new schedule that can be configured independently of other schedules.

For more information about HPE Virtual Lock immutable snapshots, see the HPE Virtual Lock—Build a Resilient Defense against Ransomware technical white paper.

HPE Alletra 9000 management and monitoring

Effective monitoring tools are required to manage storage in mission-critical environments such as healthcare. The HPE Alletra 9000 OS helps simplify, automate, and expedite storage management by handling provisioning and change management automatically and intelligently — at a subsystem level — without administrator intervention. These tools allow the rapid identification or prediction of storage system bottlenecks.

HPE Alletra 9000 management console

Beginning with HPE Alletra 9000 OS 9.4, Data Services Cloud Console (DSCC) simplifies the initial deployment of the storage array. Previously, the array was initialized on-premises through the onboard UI. With DSCC, a short pre-initialization is done on the array using the HPE Cloud Connectivity Wizard. This will connect the array to the HPE GreenLake service, allowing the start of the Setup Service in DSCC where blueprints reduce the amount of information needed to be entered manually when initializing the storage array. Furthermore, this

allows the subscription key to be presented on screen, rather than being sent through email. See the HPE Alletra 9000 Simplified Setup with Data Services Cloud Console technical white paper for more details.

HPE Alletra 9000 command line interface

In addition to the DSCC, HPE Alletra 9000 Storage administrators can also use the HPE Alletra 9000 CLI to provide rich instrumentation of all physical and logical objects for one or more storage systems, thus eliminating the need for the extra tools and consulting often required for diagnosis and troubleshooting.

HPE InfoSight for HPE Alletra 9000

HPE Alletra 9000 systems collect detailed telemetry information continuously at predefined time frames and send this data to HPE InfoSight for analysis. The collected data is routed into analytics pipelines and then cataloged and displayed in the portal. This telemetry provides HPE InfoSight with information, including:

- System details
- Hardware and software configuration history
- Capacity trending
- Performance trending and key performance indicators (KPIs)
- Health and wellness summaries and status
- Best practices and supportability
- Software and patch recommendations
- Alerts and support cases

How does HPE InfoSight work?

HPE InfoSight performs five main functions:

- **Observe —** Through simultaneous monitoring of all systems in the installed base, HPE InfoSight develops a steady-state understanding of the ideal operating environment for every workload and application. In this way, abnormal behavior can be identified through recognition of the underlining I/O patterns and configurations of each environment.
- **Learn —** Deep system telemetry coupled with global connectivity creates a foundation of data that exploits the experiences of every system connected. Machine learning in the cloud rapidly accelerates the knowledge and global learning of HPE InfoSight.
- **Predict —** For any new problem detected, HPE InfoSight can learn to predict the issue and use pattern-matching algorithms to determine if any other system in the installed base will be susceptible. Additionally, application performance can be modeled and tuned for new infrastructure based on historical configurations and workload patterns.
- **Recommend —** Based on the predictive analytics, HPE InfoSight determines the appropriate recommendation needed to improve and ensure the ideal environment. Recommendations are system operational decisions that free IT and eliminate the guesswork in managing infrastructures.
- **Act —** Through mutual trust between the infrastructure and HPE InfoSight, recommendations can be applied automatically on behalf of IT administrators. When automation is not available, specific recommendations can be delivered through support case automation.

As a best practice, Hewlett Packard Enterprise recommends that HPE Alletra 9000 systems supporting critical data use HPE InfoSight to add notes to their systems. Adding the user-defined metadata to tag mission-critical systems (such as Epic IRIS or Epic ODB) can help provide support.

Note

To learn more, download the HPE InfoSight for HPE Storage User Guide white paper.

HPE Support recommendations

To access documentation and support services, visit the HPE Support Center website. For assistance, visit the Contact HPE website.

Implementing a proof-of-concept

Epic requires each customer to run a benchmark on the storage array before the array is used to support production data. As a matter of best practice for all deployments, Hewlett Packard Enterprise recommends implementing a proof-of-concept using a test environment that matches the planned production environment as closely as possible. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HPE Services representative or your HPE partner.

Hewlett Packard Enterprise recommendations for Epic support

The following third-party support level requirements are component-specific and subject to change:

- **Analytical reporting database** — 24x7 hardware and operating system support
- **ODB for production purposes** — 24x7 hardware and operating system support with 60-minute response time and on-site availability for hardware replacement and diagnostics
- **24x7 high availability support with 30-minute response time** — HPE Proactive Care, HPE Proactive Care Advanced, or HPE Complete Care Service (hp4epic@hpe.com)

Note
See the current Epic Hardware Configuration Guide for support recommendations (supplied directly to customers by Epic).

Information to collect

Before contacting HPE Support, be sure to have the following information at hand:

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on and third-party products or components

Customer self-repair

HPE customer self-repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your HPE authorized service provider will determine whether a repair can be accomplished by CSR.

Note
For more information about CSR, contact your local service provider.

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis and automatic, secure submission of hardware event notifications to HPE Support to initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

Summary

This paper provides information for configuring HPE Alletra 9000 arrays in mission-critical healthcare environments, with recommendations based on testing by HPE engineers and knowledge gained from similar deployments. The Epic Hardware Configuration Guide is used to help define the storage requirements. Additional storage array front-end and back-end workloads and overhead must be accounted for.

Considerations for Epic EHR deployments on HPE Alletra 9000 arrays include:

- Hewlett Packard Enterprise recommends the use of high availability storage systems to mitigate hardware component failure.

- HPE Alletra 9000 systems are recommended to include a minimum of four controller nodes for Epic environments.
- Verify with benchmarking the performance requirements of the different workloads.
- Enable HPE Alletra 9000 deduplication and compression for Epic Chronicles database.
- The HPE Alletra 9000 array can support isolation of the production database on a single array using multiple CPGs and disk filtering. When considering the use of CPG disk filters, avoid using the specifier for Physical Disk 1.
- Array-based replication solutions complement database protection and is a requirement for non-IRIS data.
- Additional future growth and performance requirements and recommendations should be included when planning the system.

Resources

Hewlett Packard Enterprise Support Center
hpe.com/support/hpesc

Single Point of Connectivity Knowledge (SPOCK) Storage compatibility matrix (login required)
hpe.com/storage/spock

Assistance in system capacity and performance sizing (HPE NinjaSTARS)
saf.hpe.com/

HPE GreenLake Uptime Guarantee
hpe.seismic.com/Link/Content/DC_mnb8nj9bU6WY_48sPF7Ww

HPE Alletra 9000
hpe.com/us/en/hpe-alletra.html

HPE Alletra 9000 QuickSpecs
h20195.www2.hpe.com/v2/getdocument.aspx?docname=a50002571enw

HPE Virtual Lock—Build a Resilient Defense against Ransomware
hpe.com/psnow/doc/a00127331enw

HPE ProLiant DL servers
hpe.com/us/en/servers/proliant-dl-servers.html

HPE Storage white papers and analyst reports
hpe.com/storage/whitepapers

HPE GreenLake Data Services Cloud Console
hpe.com/us/en/storage/data-services-cloud-console.html

HPE Storage Satisfaction Guarantees for HPE GreenLake for Block Storage, HPE Alletra, HPE Primera, and HPE Nimble Storage
hpe.seismic.com/Link/Content/DCehmzjHsxh0G81WVuZ2drhg

HPE Advisory and Professional Services
hpe.com/us/en/services/consulting.html

Learn more at

hpe.com/us/en/solutions/healthcare.html

Explore HPE GreenLake

Chat now (sales)

a00118521ENW, Rev. 2

Hewlett Packard Enterprise