



Epic

Enterprise applications

NetApp
December 17, 2024

Table of Contents

- Epic 1
 - Epic on ONTAP 1
 - Epic on ONTAP 2
 - Epic architecture and design 7
 - Configuration and best practices 15
 - Additional information for Epic on ONTAP 20

Epic

Epic on ONTAP

The key to digital transformation is doing more with your data.



This documentation replaces these previously published technical report *TR-3923: NetApp Best Practices for Epic*.

Hospitals require large amounts of data to start the journey of digital transformation. Part of the process of treating patients, managing staff schedules and medical resources is that information is collected and processed. However, many actions are still performed manually or through outdated systems. The one constant is that the amount of data continues to grow exponentially and thus, becomes increasingly difficult to manage.

The major cause of this problem is that hospital data is often stored in data silos. Too much time is spent on manual entries and updates which lead to burnout and errors. This document is about one part of healthcare data, Epic Electronic Health Records (EHR). However, the data management strategy covered here can and should be applied to all healthcare data. NetApp has a proven history of modernizing and simplifying digital infrastructure. Our intelligent data infrastructure forms the foundation of digital transformation.

NetApp offers a single data management solution for all healthcare needs, and we are able to guide hospitals through their journey toward digital transformation. Building a foundation with structure and smart solutions, healthcare can extract the full value of this precious information. This framework can help medical providers diagnose diseases faster and develop individualized treatment plans to better support decision-making processes in emergency situations. You will also be able to build your own intelligent data infrastructure and enable your hospital to unlock data silos, facilitate data interoperability, and protect sensitive patient information.

Use this document as a guide to successfully construct and implement Epic EHR. Rather than building multiple Epic silos, build a single Epic data infrastructure and transform your hospital.

Purpose

This document describes best practices for integrating NetApp storage into an Epic software environment. It contains the following sections:

- A technical understanding of the Epic software environment and its storage requirements across various configurations.
- Epic storage considerations, describing important decision-making factors for Epic solutions.
- NetApp storage recommendations, describing NetApp storage configuration best practices for satisfying Epic storage requirements.

Scope

This document does not cover the following subjects:

- Quantitative performance requirements and sizing guidance, which are addressed in [TR-3930i: NetApp Sizing Guidelines for Epic](#) (NetApp login required)

Audience

NetApp assumes that the reader has the following background knowledge:

- A solid understanding of SAN and NAS concepts
- Technical familiarity with ONTAP storage systems
- Technical familiarity with the configuration and administration of ONTAP

Epic on ONTAP

EPIC on ONTAP

Epic is easier with ONTAP.

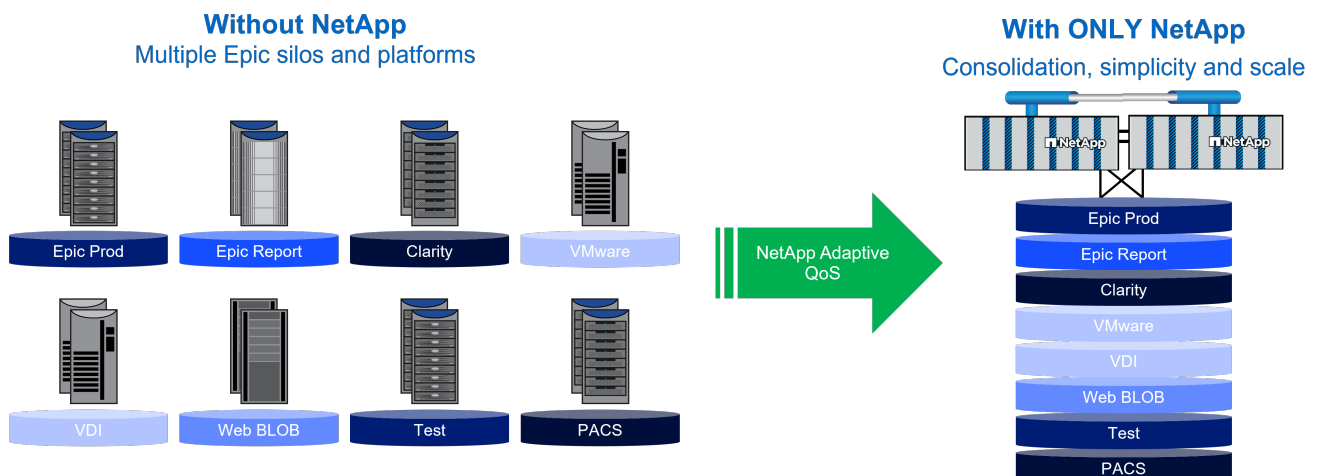
ONTAP is a data management platform that allows you to consolidate Epic workloads while meeting all of your performance, data protection, and data management requirements.

Only on NetApp can you standardize all your healthcare workloads for SAN, NAS, and Object on a single high-availability data management platform. ONTAP is the most widely deployed storage software platform in the world and comes with almost 30 years of constant innovation. You can meet all your Epic challenges with native ONTAP data management tools and application integration. There is no need to purchase a multitude of third-party tools to fill gaps in the solution.

A lot of storage vendors offer traditional, reliable, and fast block storage. They work well but are typically deployed in silos to run a single workload such as production, report, clarity, VDI, VMware, and NAS. Each of these silos has different hardware and different management tools, and they are typically managed by different IT groups. This traditional approach adds to the biggest problem with healthcare today - complexity.

NetApp makes data management easier and more efficient. Instead of throwing money at the problem with oversized silos, ONTAP uses innovation and technology to enable a consistent and guaranteed SLA for each workload on a single platform over any protocol with integrated data protection. These capabilities and tools also extend out to the cloud of your choice as illustrated below.

Scale and simplicity for Healthcare



Epic on ONTAP availability

At the core of ONTAP are nondisruptive operations that enable you to avoid costly interruption to business operations.

NetApp delivers over 99.999999% availability based on production data, which is called home through NetApp Active IQ. Each HA pair in the cluster has no single point of failure. ONTAP dates back to 1992 and is the most widely deployed data management software in the world with an exceptional history of delivering reliable storage. Now, with Active IQ proactively monitoring and automatically resolving 97% of issues, availability is higher with significantly less support cases.

Epic recommends the use of HA storage systems to mitigate hardware component failure. This recommendation extends from basic hardware (such as redundant power supplies) to networking (such as multipath networking).

When you need to upgrade storage, scale up, scale out, or rebalance workloads across the cluster, there is no effect to patient care. You might move data, but you never again need to disrupt patient care with data migrations or forklift upgrades. Move to next-generation technology, future proof, and avoid hardware lock-in. NetApp even offers a 100% written availability guarantee.

More information on NetApp's reliability, availability, serviceability, and security capabilities can be found in the [NetApp ONTAP reliability, availability, serviceability, and security](#) white paper.

Epic on ONTAP consolidation

One of the major challenges in healthcare is the inefficiency of siloed environments.

Multiple point solutions are created by various groups that impedes progress. Having a unified strategy to data management brings efficiency to accelerate transformation. Disruptive technology like digitizing patient records, Ransomware, and generative AI all the drive the need for consolidation.

With ONTAP you can consolidate file/block/object and each of your tier 0/1/2/3 workloads, on premises and in the cloud, all running on ONTAP.

Epic on ONTAP efficiency

Epic runs on all-flash arrays where most of the cost is the disk. Therefore, storage efficiency is critical for cost savings.

NetApp inline storage efficiency achieves industry-leading savings on storage with no effects to performance, and we even offer a written efficiency guarantee with the all-flash arrays.


When calculating storage efficiency, it is important to measure raw to usable to effective capacity.

- **Raw capacity** Before any RAID is applied, size of disk by number of disks.
- **Usable capacity** After RAID is applied, how much usable storage is available.
- **Effective capacity** How much storage is provisioned and presented to the host or client.

The figure below is a sample efficiency calculation of a typical Epic deployment including all workloads requiring 852TB of effective storage and with 5.2:1 efficiency delivering 1.32PB of total effective data.



Based on the number of disks, raw-to-usable capacity varies slightly.



NetApp does not use NetApp Snapshot technology or thin provisioning to calculate efficiency in the guarantee program. Doing so would show unrealistic efficiencies of 30-100:1, which do not mean anything when sizing real-world storage capacity.

Epic on ONTAP performance

ONTAP introduced flash technologies in 2009 and has supported SSDs since 2010. This long experience with flash storage allows NetApp to tune ONTAP features to optimize SSD performance and enhance flash media endurance while keeping the feature-rich capabilities of ONTAP.

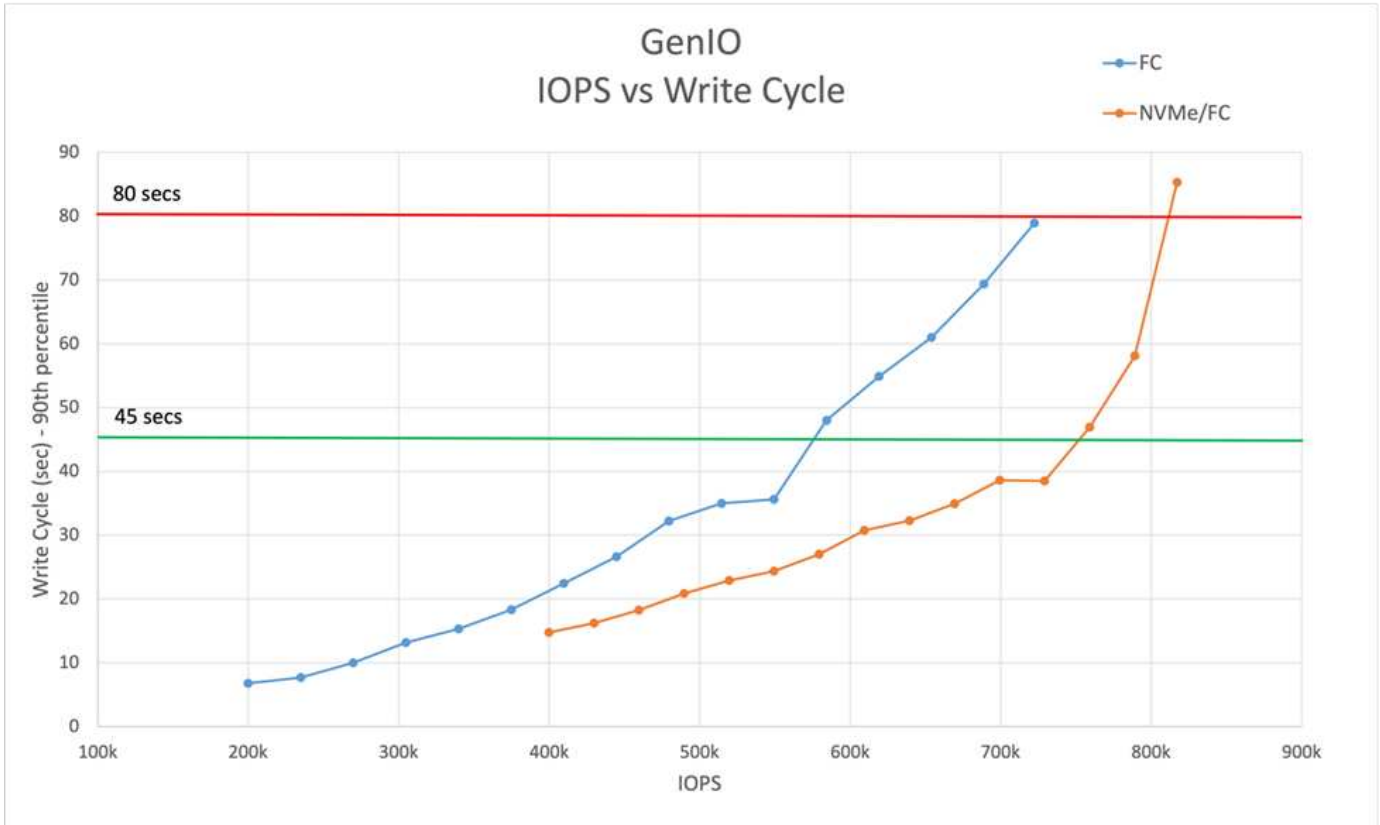
Since year 2020, all Epic ODB workloads are required to be on all-flash storage. Epic workloads typically operate at approximately 1,000-2,000 IOPs per terabyte of storage (8k block, 75%/25% read and write ratio, and 100% random). Epic is very latency-sensitive, and high latency has a visible effect on the end-user experience as well as operational tasks such as running reports, backup, integrity checks, and environment refresh times.

- The limiting factor for all-flash arrays is not the drives but, rather, it is the utilization on the controllers.
- ONTAP uses an active-active architecture. For performance, both nodes in the HA pair write to the drives.
- This result is maximized CPU utilization, which is the single most important factor that allows NetApp to publish the best Epic performance in the industry.
- NetApp RAID DP, Advanced Disk Partitioning (ADP), and WAFL technologies deliver on all Epic requirements. All workloads distribute IO across all the disks. No bottlenecks.
- ONTAP is write-optimized; writes are acknowledged once written to mirrored NVRAM before they are written to disk at inline memory speed.
- WAFL, NVRAM, and the modular architecture enable NetApp to use software to innovate with inline efficiencies, encryption, performance. They also enable NetApp to introduce new features and functionality without impacting performance.
- Historically, with each new version of ONTAP there is an increase in performance and efficiency in the range of 30-50%. Performance is optimal when you stay current with ONTAP.

NVMe

When performance is paramount, NetApp also supports NVMe/FC, the next-generation FC SAN protocol.

As can be seen in the figure below, our GenIO testing achieved a much greater number of IOPS using NVMe/FC protocol versus the FC protocol. The NVMe/FC connected solution achieved over 700k IOPS before surpassing the 45-second write cycle threshold. By replacing SCSI commands with NVMe, you also significantly reduce the utilization on the host.



Epic on ONTAP scalability

The Epic Hardware Configuration Guide accounts for ~20% growth per year for 3 years. However, environments can also grow unexpectedly.

NetApp can seamlessly scale performance and capacity up to 12 nodes for NAS, SAN, and Object clusters. As a result, you can nondisruptively scale up and out as your business grows.

Epic Iris brings added scaling capabilities. It enables larger customers that have multiple Epic instances to consolidate to a single instance. The [NetApp Verified Architecture Epic on Modern SAN](#) document demonstrates that Epic can seamlessly scale consolidated workloads to 720K IOPS on an single HA and scale out to over 4M IOPS in a cluster. You can scale up nondisruptively by upgrading controllers or adding disks to existing clusters.

NAS, SAN, and Object data also has the ability to be moved nondisruptively between nodes in the cluster. Each HA pair in the cluster can be any combination of ONTAP FAS and AFF system types and sizes. You can balance your workloads across a single cluster to maximize your storage investment.

ONTAP also provides the option to use object storage on StorageGRID or the cloud as a backup target and/or automatic cold storage tiering target. This capability enables you to free up expensive, all-flash disks, tier snapshots, and cold data automatically to Object.

The result is that Epic simply runs better with the NetApp product portfolio, leveraging ONTAP, multiple protocols, StorageGRID, and the cloud of your choice. These products provide options for disaster recovery,

archiving, analytics, tiering, and more.

Epic storage efficiency configuration

A snapshot is a point-in-time copy of a volume that is read-only.

A snapshot puts a logical lock on all the blocks in the active file system. NetApp ONTAP Snapshot copies are near instant, and they do not use any additional storage.

Write Anywhere File Layout, or WAFL, is a write-only file system; it does not perform additional IO, such as copying the data in a snapshot-protected block before being overwritten. No data is ever moved; therefore, snapshots have no effect on storage capacity or performance. Snapshots provide tremendous savings in storage while augmenting the backup solution.

FlexClone

A NetApp ONTAP FlexClone volume is a clone of an existing volume or a snapshot of an existing volume. It is otherwise an ONTAP volume like any other, and can itself be cloned, protected with snapshots, and configured with a QoS policy.

As with snapshots, a FlexClone volume does not require any additional space at creation time. Only changes to the clone require additional capacity.

Epic requires 10 to 30 copies of the production databases for various operational requirements such as streaming backups, integrity checks, and staging upgrade environments. The need for a solution built on FlexClone volumes has increased with the move to more frequent upgrades.



A fully automated Epic backup solution and Epic refresh solution are provided by NetApp as part of the solution using Ansible and native NetApp tools.

Epic on ONTAP security

Security is the number one concern for organizations and healthcare executives today. It has never been more difficult to manage, and organizations are challenged with compliance, data governance, antivirus protection, and ransomware.

A complete guide to Epic and storage security is beyond the scope of this document; however, [Security Hardening Guide for ONTAP](#) details all the extensive and advanced security features available with ONTAP.

NetApp Active IQ Unified Manager monitors for security violations based on the information included in [TR-4569](#) and reports them in the dashboard to simplify security management. These tools can help your organization meet your security goals to protect, detect, and remediate against attacks.

NetApp has also partnered with security vendors to provide integration through [NetApp FPolicy](#) software to enhance your security offering. Furthermore, [multifactor \(MFA\) authentication](#) can be added to secure your Epic environment against unauthorized access with leaked credentials.

Finally, ONTAP native Snapshot copies and immutable SnapLock technologies with [ONTAP cyber vault](#), offer a unique air gap capability to protect your patient records against ransomware. See NetApp documentation on [The NetApp Solution for Ransomware](#). For a more strategic approach to security, see [NetApp and Zero Trust](#).

Epic architecture and design

Epic architecture

This section describes the Epic software environment and the key components that require storage. It provides key considerations to help guide storage design.

Epic, headquartered in Verona, Wisconsin, makes software for medium to large medical groups, hospitals, and integrated healthcare organizations. Customers also include community hospitals, academic facilities, childrens' organizations, safety-net providers, and multi-hospital systems. Epic-integrated software spans clinical, access, and revenue functions and extends into the home.

It is beyond the scope of this document to cover the wide span of functions supported by Epic software. From the storage system point of view, however, all Epic software shares a single patient-centric database for each deployment. Epic is transitioning from the InterSystems Caché database to the new InterSystems Iris database. Because the storage requirements are the same for Caché and Iris, we will refer to the database as Iris throughout the rest of this document. Iris is available for the AIX and Linux operating systems.

InterSystems Iris

InterSystems Iris is the database used by the Epic application. In this database, the data server is the access point for persistently stored data. The application server manages database queries and makes data requests to the data server. For most Epic software environments, the use of the symmetric multiprocessor (SMP) architecture in a single database server suffices to service Epic applications' database requests. In large deployments, a distributed model can be supported by using InterSystems' Enterprise Caché Protocol (ECP).

The use of failover-enabled clustered hardware enables a standby data server to access the same storage as the primary data server. It also enables the standby data server to take over processing responsibilities during a hardware failure.

InterSystems also provides technologies to satisfy data replication, disaster recovery, and high-availability (HA) requirements. InterSystems' replication technology is used to replicate an Iris database synchronously or asynchronously from a primary data server to one or more secondary data servers. NetApp SnapMirror is used to replicate WebBLOB storage or for backup and disaster recovery.

The updated Iris database has many advantages:

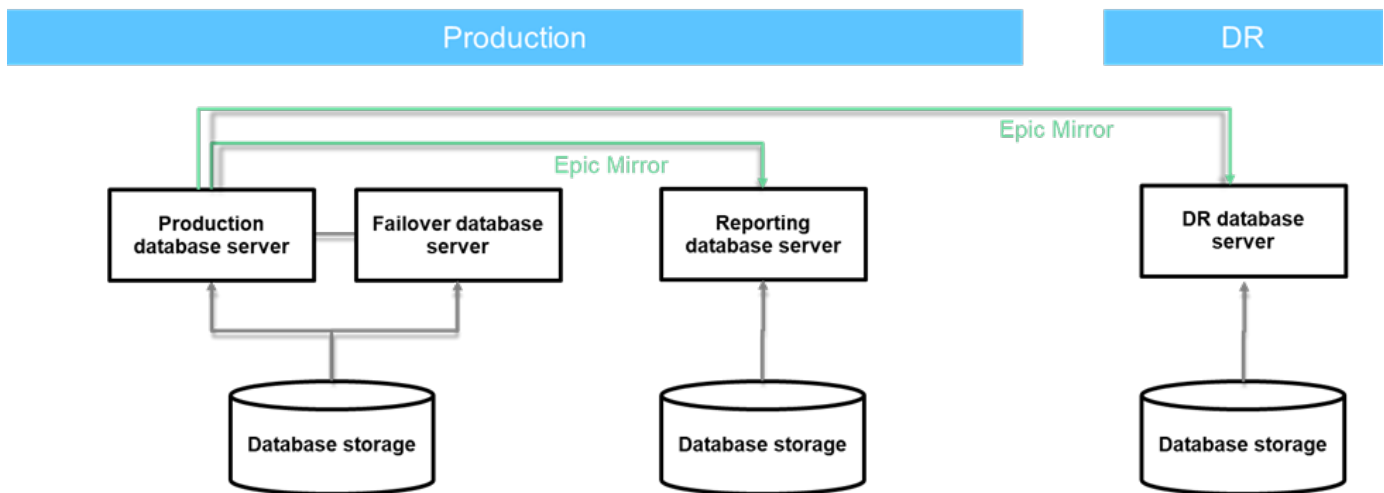
- Increased scale and enables larger organizations with multiple Epic instances to consolidate into one larger instance.
- A licensing holiday where customers can now move between AIX and Red Hat Enterprise Linux (RHEL) without paying for a new platform license.

Caché database servers and storage usage

- **Production** In Epic software environments, a single patient-centric database is deployed. In Epic's hardware requirements, the physical server hosting the primary read/write Iris data server is called the production database server. This server requires high performance all-flash storage for files belonging to the primary database instance. For high availability, Epic supports the use of a failover database server that has access to the same files. Iris uses Epic Mirror to replicate to read-only Report, disaster recovery and support read-only copies. Each type of database server can be switched to read/write mode for business continuity reasons.
- **Report** A reporting mirror database server provides read-only access to production data. It hosts an Iris data server configured as a backup mirror of the production Iris data server. The reporting database server

has the same storage capacity requirements as the production database server. Reporting write performance is the same as production but read workload characteristics are different and sized differently.

- **Supports read-only** This database server is optional and not shown the figure below. A mirror database server can also be deployed to support Epic supports read-only functionality, in which access is provided to a copy of production in read-only mode. This type of database server can be switched to read/write mode for business continuity reasons.
- **Disaster recovery** To meet business continuity and disaster recovery objectives, a disaster recovery mirror database server is commonly deployed at a site geographically separate from the production and/or reporting mirror database servers. A disaster recovery mirror database server also hosts an Iris data server configured as a backup mirror of the production Iris data server. If the production site becomes unavailable for an extended time, this backup mirror database server can be configured to act as a mirror read/write instance (SRW). The backup mirror database server has the same file storage requirements as the production database server. In contrast, the backup mirror database storage is sized the same as the production storage from a performance perspective for business continuity.



- **Test** Healthcare organizations often deploy development, testing, and staging environments. Additional Iris data servers for these environments also require storage, which can be accommodated by the same storage system. Epic has specific requirements and constraints for providing additional storage from a shared storage system. These specific requirements are addressed generically by the best practices in this document.

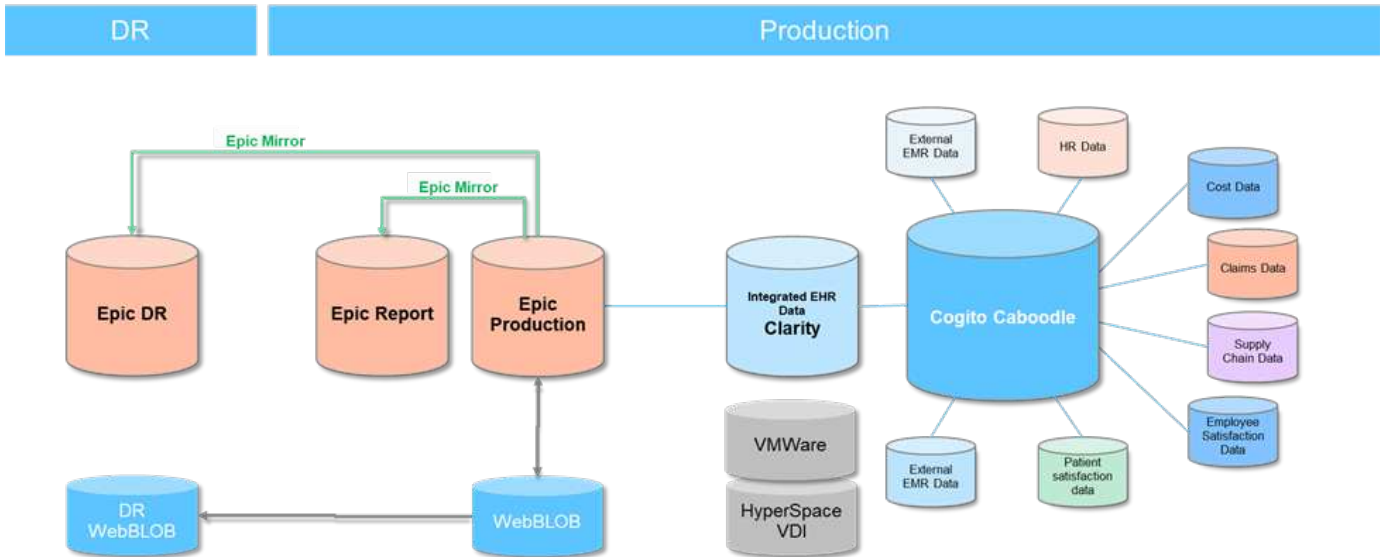
In addition to Iris ODB data servers, Epic software environments commonly include other components such as the following and as shown in the figure below:

- An Oracle or Microsoft SQL Server database server as a back end to Epic’s Clarity business-reporting tools



Clarity is used to report on data extracted daily from the reporting Iris database.

- WebBLOB server (SMB)
- Multipurpose database server
- Multipurpose virtual machines (VMs)
- Hyperspace for client access



The storage requirements of all these multiple workloads, pools, NAS and SAN protocols can be consolidated and hosted by a single ONTAP cluster. This consolidation enables healthcare organizations to have a single data management strategy for all Epic, and Non-Epic, workloads.

Operational database workloads

Each Epic database server performs I/O on the following types of files:

- Database files
- Journal files
- Application files

The workload of an individual database server depends on its role in the Epic software environment. For example, production database files typically incur the most demanding workload, consisting of 100% random I/O requests. The workload of any mirror database is typically less demanding and has fewer read requests. Journal file workloads are mainly sequential.

Epic maintains a workload model for storage performance benchmarking and customer workload. For more information about the Epic workload model, benchmark results, and guidance on using NetApp sizing tools to correctly size storage for Epic environments, see [TR-3930i: NetApp Sizing Guidelines for Epic](#) (NetApp login required).

Epic also provides each customer with a customized hardware configuration guide containing I/O projections and storage capacity requirements. The final storage requirements might include development, testing, and/or staging environments, and any other ancillary workloads which may be consolidated. Customers can use the hardware configuration guide to communicate the total storage requirements to NetApp. This guide contains all the data needed to size an Epic deployment.

During the deployment phase, Epic provides a Database Storage Layout Guide, which provides more granular LUN-level details that can be used for an advanced storage design. Note that the Database Storage Layout Guide is general storage recommendations and not specific to NetApp. Use this guide to determine the best storage layout on NetApp.

Epic sizing

One of the key architecture considerations when sizing an Epic storage environment is the ODB database size.

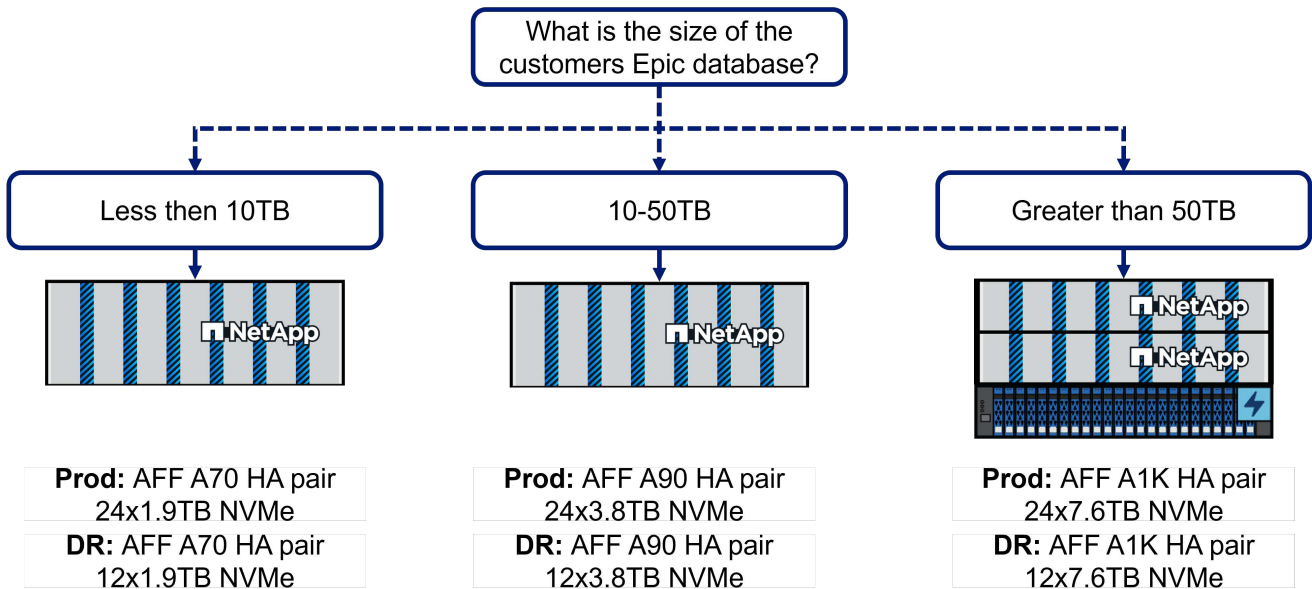
You can use the diagram shown below to select a small-medium-large Epic storage architecture. These designs include running all workloads listed in the Hardware Configuration Guide. The sizing tree is based on data from over 100 hardware configuration guides and should be mostly accurate estimate.

It is important to note that this is only a starting point only. You should work with our Epic alliance team to confirm any Epic designs. The team can be reached at epic@netapp.com. Every deployment needs to accommodate customer requests while adhering to Epic and NetApp recommended best practices.

- Small Epic architecture with an Epic database less than 10TB
- Medium Epic architecture with an Epic database from 10TB to 50TB
- Large Epic architecture with an Epic database from greater than 50TB

Epic sizing decision tree

Work with the NetApp Epic alliance team to validate designs



Epic storage requirements

Dedicated storage resources are generally provided for the production database, whereas mirror database instances share secondary storage resources with other Epic software-related components, such as the Clarity reporting tools.

Other software storage environments, such as those used for application and system files, are also provided by the secondary storage resources.

Beyond sizing considerations, Epic has the following additional storage layout rules and key considerations:

- Since 2020, all operational database (ODB) workloads must be on all-flash arrays.
- Epic recommends each pool of storage to be on separate physical hardware, including pool1, pool2, pool3,

NAS1, and NAS2.



A node in a cluster can be considered as a pool of storage. With ONTAP 9.4 or later and AQuoS, you can create protected pools using policies.

- New Epic 3-2-1 backup recommendation.
 1. Copy located in remote site (disaster recovery)
 2. One of the copies must be on a different storage platform than the primary copy
 3. Copies of the data



Customers who use NetApp SnapMirror to back up NetApp do not meet the 3-2-1 recommendations. The reason is ONTAP to ONTAP does not satisfy the second requirement listed above. You can use SnapMirror directly from ONTAP to object storage on-premises (through StorageGRID, for example) or to the cloud to meet Epic requirements.

For more information about storage mandates, see the following Epic guides available in Galaxy:

- SAN Considerations
- Storage Products and Technology Status (SPaTS)
- Hardware Configuration Guide

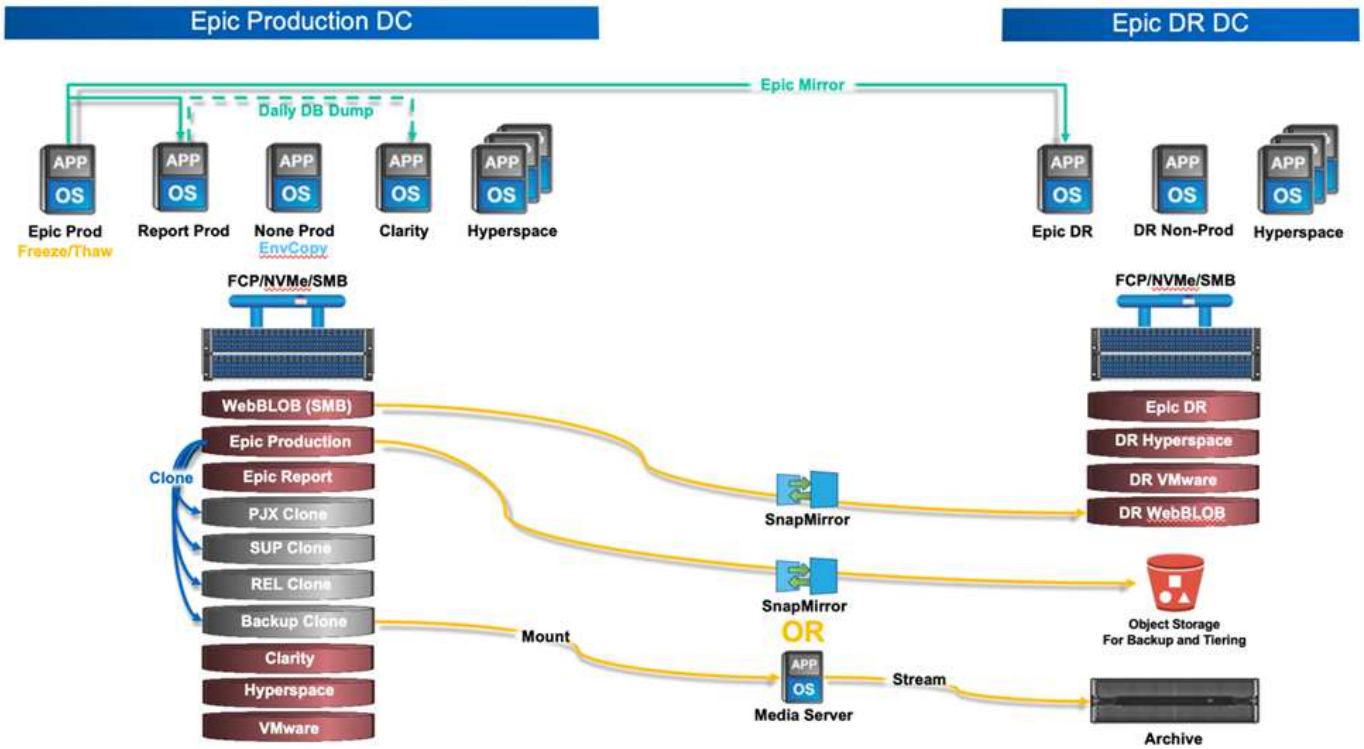
Epic four-node architecture

The figures below show the storage layout for a four-node architecture: an HA pair in production and an HA pair in disaster recovery. The size of the controllers and number of disks are based on the latter sizing image.


NetApp guarantees a minimum floor level performance by accepting the SLM recommended AQuoS policies. Epic supports consolidating storage pools on ONTAP onto significantly less hardware. For more information, see the Epic quarterly SPATS document. Basically, pool1, pool2, and NAS1 (listed in the Epic Hardware Configuration Guide) can all be run on a single HA pair with the workloads spread evenly across the two controllers. In disaster recovery, Epic pool 3 and NAS 3 are also split between the two controllers in the HA pair.

Test full copy environments (such as SUP, REL, and PJX) are cloned from either Epic Production, Epic Report, or Epic Disaster Recovery. For information about Epic backup and refresh, see the section titled, "Data management".

Four-node architecture



Four-node workload placement



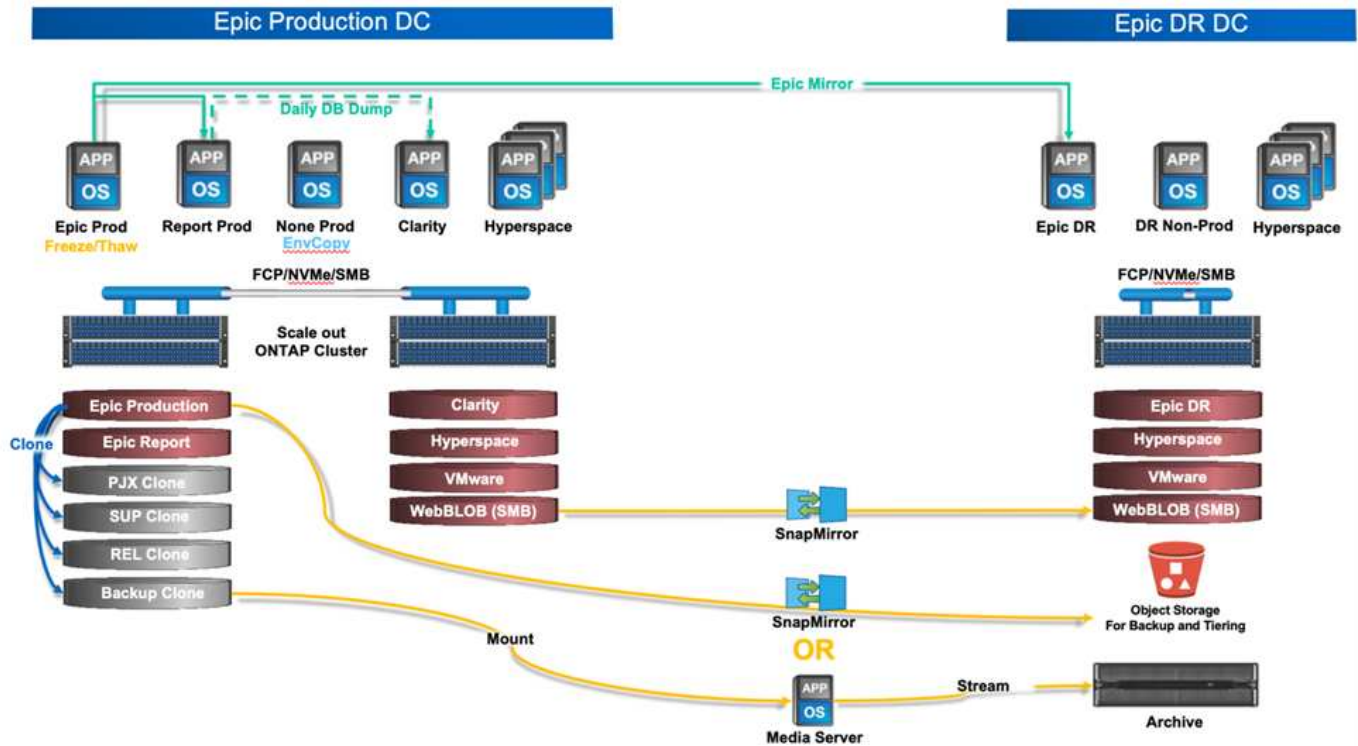
Epic six-node architecture

Customers might want to start with a six-node design or seamlessly scale out from four to six nodes with growing demand. With scale-out, you can nondisruptively move workloads between nodes and rebalance across the cluster.


This architecture offers the best performance and capacity balance on the cluster. Epic Production, Epic

Report, and Epic Test all run on the first HA pair. The second HA pair is used for Clarity, Hyperspace, VMware, NAS1, and the remaining Epic workloads. Disaster recovery is the same as the four-node architecture in the previous section.

Six-node architecture



Six-node workload placement

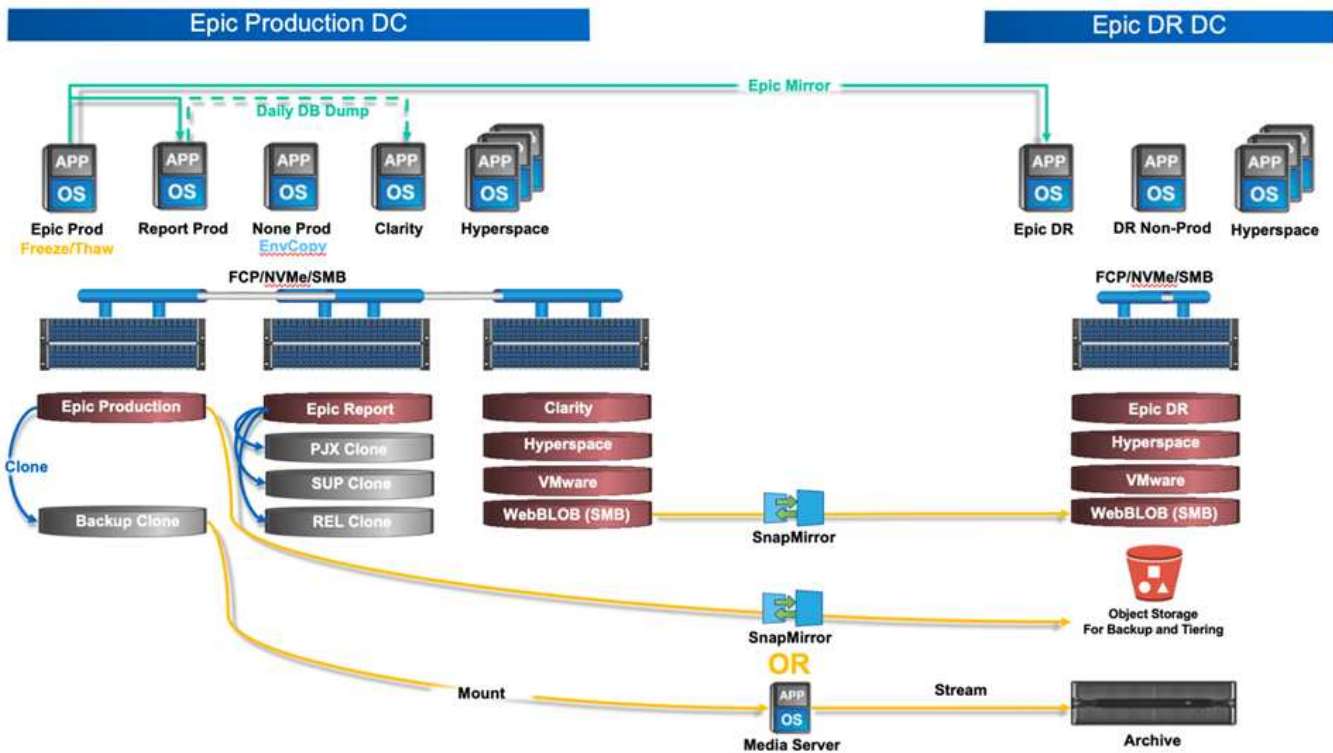


Epic eight-node architecture


The figures below show the scale-out eight-node architecture. Again, you can start with four nodes and scale to six nodes and continue to scale to eight nodes and beyond. This architecture offers the best balance of performance and capacity across the six nodes in production.

The test environments are cloned from Report instead of Production in this design. This offloads test environments and integrity checks from production.

Eight-node architecture



Eight-node workload placement



Configuration and best practices

Epic on ONTAP - Host Utilities

The NetApp Host Utilities are software packages for various operating systems that contain management utilities such as the `sanlun` CLI binary, multipath drivers, and other important files required for proper SAN operations.



NetApp recommends installing the NetApp Host Utilities on hosts that are connected to and accessing NetApp storage systems. For more information, see [Interoperability Matrix Tool](#) and [SAN Hosts](#) documentation.



With AIX, it is especially critical that the Host Utilities are installed before discovering LUNs. This ensures that the LUN multipathing behavior is configured correctly. If discovery was performed without the Host Utilities, the LUNs will need to be deconfigured from the system using the `rmdev -dl` command and then rediscovered via `cfgmgr` or a reboot.

Epic LUN and volume configuration

The Epic Database Storage Layout Recommendations document provides guidance on the size and number of LUNs for each database.

It is important to review this document with the Epic DBA and Epic support as well as finalize the number of LUNs and LUN sizes as they might need to be adjusted. These storage recommendations are important for HBA queue depth, storage performance, ease of operations and ease of expansion.

For the server OS queue depth consideration, use a minimum of eight LUNs (one LUN per volume) for a database. Increase the number of LUNs by the number of nodes in the ONTAP cluster. For example, add 4 LUNs when using a 4 node (2 HA pair) cluster. For larger environments, more LUNs might be required; use the same number of volumes (eight total, distributed across storage node) and add LUNs in multiples of two across the cluster nodes and volumes. This approach enables you to easily scale your Epic environment.

Example 1: 2 Node ONTAP cluster

2 Node, 1 HA pair

8 Volumes, 4 volumes per node

8 LUNs, one LUN per volume

Adding an additional 2 LUNs, one on node01 in volume01, one on node02 in volume02.

Example 2: 4 Node ONTAP cluster

4 Node, 2 HA pair

8 Volumes, 2 volumes per node

8 LUNs, one LUN per volume

Adding an additional 4 LUNs, one on node01 in volume01, one on node02 in volume02, one on node03 in volume03, one on node04 in volume04.


To maximize performance for a workload, such as Epic ODB or Clarity, each layout works best also for NetApp storage. By using eight volumes, write IO is spread evenly across controllers, maximizing CPU utilization. For replication and backup, it's best to limit the number of volumes to eight to simplify operations.

Scaling options

If more storage is required by the server, the easiest option is to grow the LUNs containing volumes. The second option is to add LUNs to the volume groups in multiples of two at a time (one per volume per node).

Example:

Volume and 8-LUN layout



If in a large environment requiring more than 4 nodes or 8 LUNs please consult our Epic alliance team to confirm LUN designs. The team can be reached at epic@netapp.com.

Best practices

- Use 8 LUNs in 8 volumes to start, adding 2 LUNs at a time, across all nodes of the cluster.
- Balance the workloads across the HA pair to maximize performance and efficiency.
- Create LUNs at the size expected for 3 years of growth. (Consult the [ONTAP documentation](#) for maximum LUN sizes.)
- Use thin provisioned volumes and LUNs.
- Use a minimum of eight DB LUNs, two journal LUNs, and two app LUNs. This configuration maximizes storage performance and OS queue depth. More can be used if needed for capacity or other reasons.
- If you do need to add LUNs to volume groups, add eight LUNs at a time.
- Consistency Groups (CGs) are required for the group of volumes and LUNs to be backed up together.
- Do not use QoS during the GenIO or any I/O performance.
- After GenIO or Clarity testing, NetApp recommends deleting the storage and reprovisioning before loading production data.
- It's important that `-space-allocation enabled` is set on the LUNs. If not, any deleted data on the LUNs will not be seen by ONTAP and might cause capacity issues. For more information, see the Epic Storage Configuration Quick Reference Guide.

Epic and file protocols

Combining NAS and SAN on the same all-flash array is supported.



NetApp recommends using FlexGroup volumes for NAS shares, such as WebBLOB (when available).

WebBLOB is up to 95% cold data. You can optionally free space on your all-flash array and tier backups and cold data to object storage on-premises or in the cloud using the [FabricPool](#) feature of ONTAP. All of which can be accomplished without any noticeable performance effect. FabricPool is an included feature of ONTAP. Customers can generate a cold (or inactive) data report to review how much benefit could be realized by enabling FabricPool. You can set the age of the data to tier through policy. Epic customers have realized significant savings with this feature.

Epic performance management

Most all-flash arrays can deliver the performance required for Epic workloads. The NetApp differentiator is its ability to set floor level performance policies and guarantee a consistent performance for each application.

Quality of Service (QoS)

NetApp recommends using QoS. The benefit of QoS is the ability to consolidate all Epic workloads. All protocols and pools of storage can reside on less hardware. You do not need to separate pools of storage.

- NetApp recommends having all workloads in the cluster assigned to a QoS policy to better manage headroom on the cluster.
- NetApp recommends balancing all workloads evenly across the HA pair.
- Do not use QoS policies when performing any I/O testing; otherwise, GenIO testing will fail. Analyze the different production workloads for 2-4 weeks before assigning any QoS policies.

Epic on ONTAP - protocols

FCP is the preferred protocol for presenting LUNs.



NetApp recommends single initiator zoning: one initiator per zone with all the required target ports on storage using worldwide port names (WWPNs). The presence of more than one initiator in a single zone is likely to lead to intermittent HBA crosstalk, which causes significant disruption.

After the LUN is created, map the LUN to the initiator group (igroup) containing the WWPNs of the host to enable access.

NetApp also supports using NVMe/FC (if you have versions of AIX and RHEL operating systems that are capable) and enhances the performance. FCP and NVMe/FC can coexist on the same fabric.

Epic storage efficiency configuration

ONTAP inline efficiencies are on by default and work regardless of storage protocol, application, or storage tier.

Efficiencies reduce the amount of data written to expensive flash storage and reduce the number of drives required. ONTAP preserves efficiency with replication. Each of the efficiencies has little to no effect on performance, even for a latency-sensitive application like Epic.



NetApp recommends turning on all efficiency settings to maximize disk utilization. These settings are on by default on AFF and ASA based systems.

The following features make this storage efficiency possible:

- Deduplication saves space on primary storage by removing redundant copies of blocks in a volume that hosts LUNs. This recommended option is on by default.
- Inline compression reduces the amount of data to be written to disk, and a considerable savings in space is realized with Epic workloads. This recommended option is on by default.
- Inline compaction takes 4k blocks that are less than half full and combines them into a single block. This recommended option is on by default.
- Thin replication is at the center of the NetApp data protection software portfolio, which includes NetApp SnapMirror software. SnapMirror thin replication protects business-critical data while minimizing storage capacity requirements. **NetApp recommends** turning on this option.
- Aggregate deduplication. Deduplication has always been at the volume level. With ONTAP 9.2, aggregate deduplication became available, providing additional disk reduction savings. Postprocess aggregate deduplication was added with ONTAP 9.3. **NetApp recommends** turning on this option.

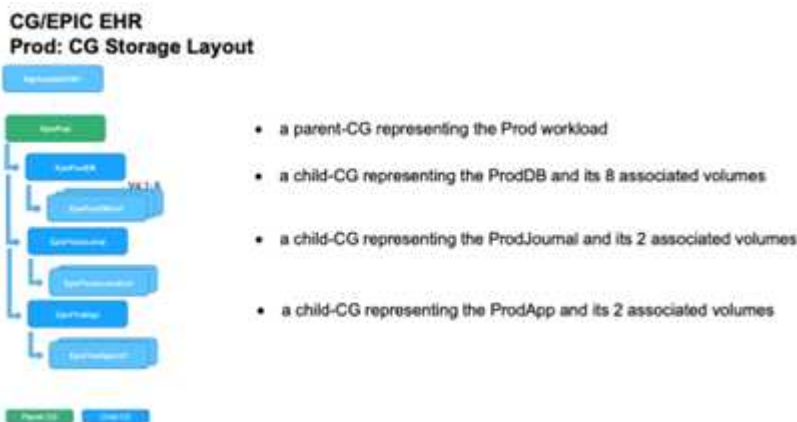
Epic storage efficiency configuration

Applications with storage spread across more than one volume with one or more LUNs of appropriate quantities for the workload need the contents to be backed up together ensuring consistent data protection require CGs.

Consistency groups (CGs for short) provide this capability and more. They can be used nightly to create on-demand or scheduled consistent snapshots using a policy. You can use this to restore, clone and even replicate data.

For additional information on CGs please refer to the [Consistency groups overview](#)

Once the volumes and LUNs are provisioned as detailed in the previous sections of this document, they can then be configured into a set of CGs. The recommended best practice is to set them up as depicted in the picture below:



Consistency group snapshots

A nightly CG snapshot schedule should be set on each of the child-CGs associated with the volumes providing storage for the production database. This will result in a fresh set of consistent backups of these CGs every

night. These can then be used for cloning the production database for use in non-production environments such as development and test. NetApp has developed proprietary CG based automated Ansible workflows for Epic to automate the backup of production databases, the refresh and test environments too.

CG snapshots can be used to support the restore operations of Epic's production database.

For SAN volumes, disable the default snapshot policy on each volume being used for CGs. These snapshots are typically managed by the backup application being used or NetApp's Epic Ansible automation service.

For SAN volumes, disable the default snapshot policy on each volume. These snapshots are typically managed by a backup application or by Epic Ansible automation.[NS2]

WebBLOB and VMware datasets should be configured as just volumes, not associated with CGs. You can use SnapMirror to maintain snapshots on storage systems separate from production.

When complete, the configuration would look as follows:



Storage sizing for Epic

You should work with our Epic alliance team to confirm any Epic designs. The team can be reached at epic@netapp.com. Every deployment needs to accommodate customer requests while adhering to Epic and NetApp recommended best practices.

For information about how to use NetApp sizing tools to determine the correct RAID group size and number of RAID groups for Epic software environment storage needs, see [TR-3930i: NetApp Sizing Guidelines for Epic](#) (NetApp login required).



NetApp Field Portal access is required.

Additional information for Epic on ONTAP

To learn more about the information described in this document, refer to the following documents and/or websites:

- [NetApp Product Documentation](#)
- [ONTAP 9 Documentation](#)
- [Consistency Groups](#)
- [ONTAP and ONTAP System Manager Documentation Resources](#)

- [TR-3930i: NetApp Sizing Guidelines for Epic](#) (NetApp login required)

Epic customer guidance documents

Epic provides customers the following documents for guidance on server, storage, and network. These documents are referenced in this technical report.

- Storage Area Network Considerations
- Business Continuity Technical Solutions Guide
- All-Flash Reference Architecture Strategy Handbook
- Storage Products and Technology Status
- Epic Cloud Considerations
- Hardware Configuration Guide (customer-specific)
- Database Storage Layout Recommendations (customer-specific)

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.