



Configuration and best practices

Enterprise applications

NetApp
December 17, 2024

Table of Contents

- Configuration and best practices 1
 - Epic on ONTAP - Host Utilities 1
 - Epic LUN and volume configuration 1
 - Epic and file protocols 3
 - Epic performance management 3
 - Epic on ONTAP - protocols 3
 - Epic storage efficiency configuration 4
 - Epic storage efficiency configuration 4
 - Storage sizing for Epic 5

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.

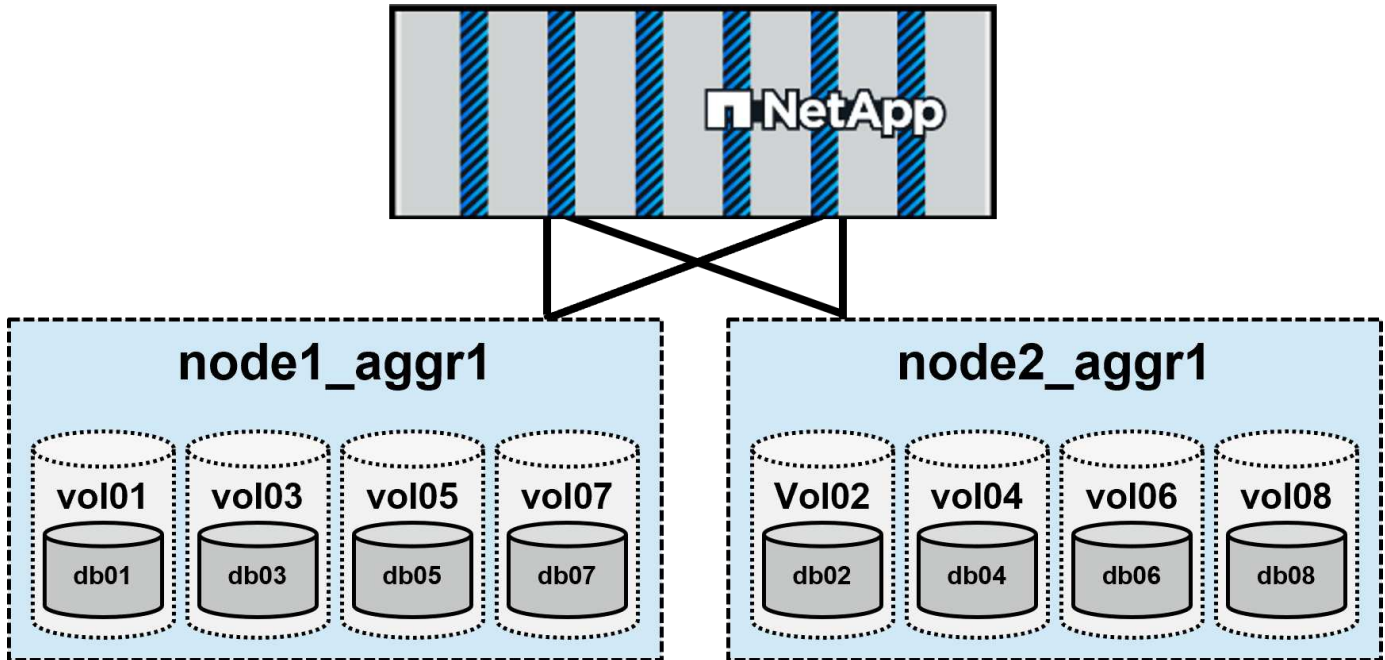
To maximize performance for a workload, such as Epic ODB or Clarity, each layout works best also for NetApp storage. By using eight volumes, write IO is spread evenly across controllers, maximizing CPU utilization. For replication and backup, it's best to limit the number of volumes to eight to simplify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

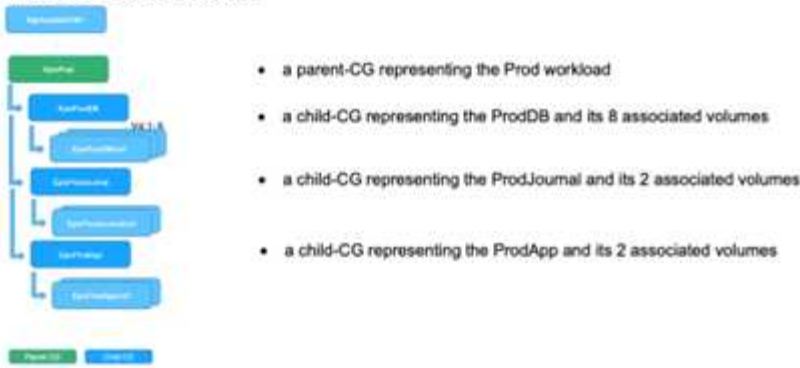
Applications with storage spread across more than one volume with one or more LUNs of appropriate quantities for the workload need the contents to be backed up together ensuring consistent data protection require CGs.

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:

**CG/EPIC EHR
Prod: CG Storage Layout**



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:

**CG/EPIC EHR
CG Storage Layout**



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.

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.