



Epic architecture and design

Enterprise applications

NetApp

December 17, 2024

Table of Contents

- Epic architecture and design 1
 - Epic architecture 1
 - Epic sizing 4
 - Epic storage requirements 4
 - Epic four-node architecture 5
 - Epic six-node architecture 6
 - Epic eight-node architecture 8

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.

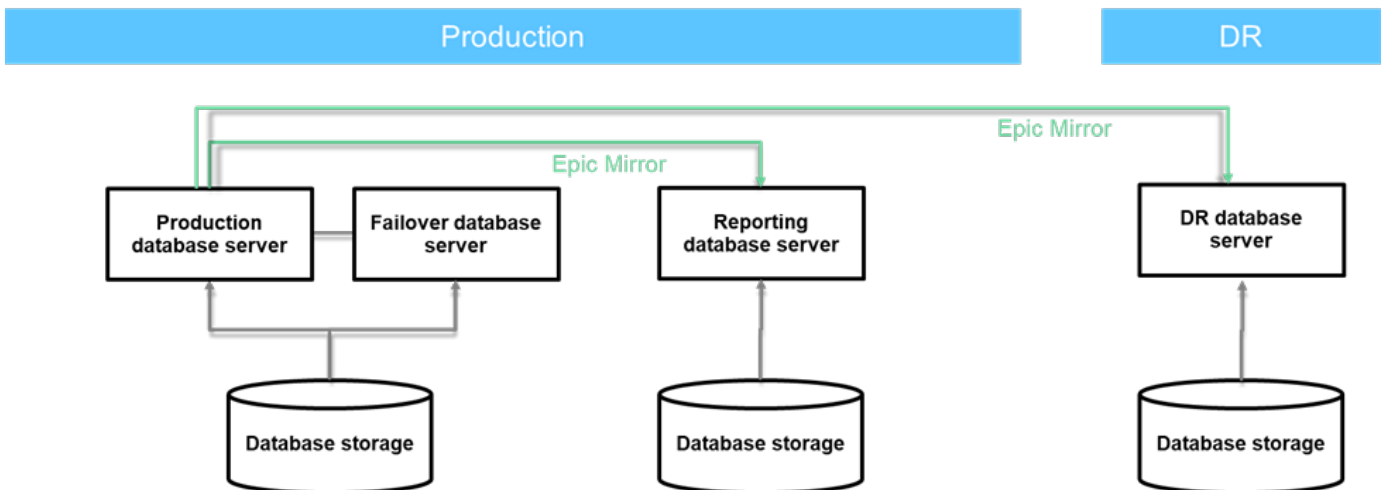
The updated Iris database has many advantages:

- 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 in 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.

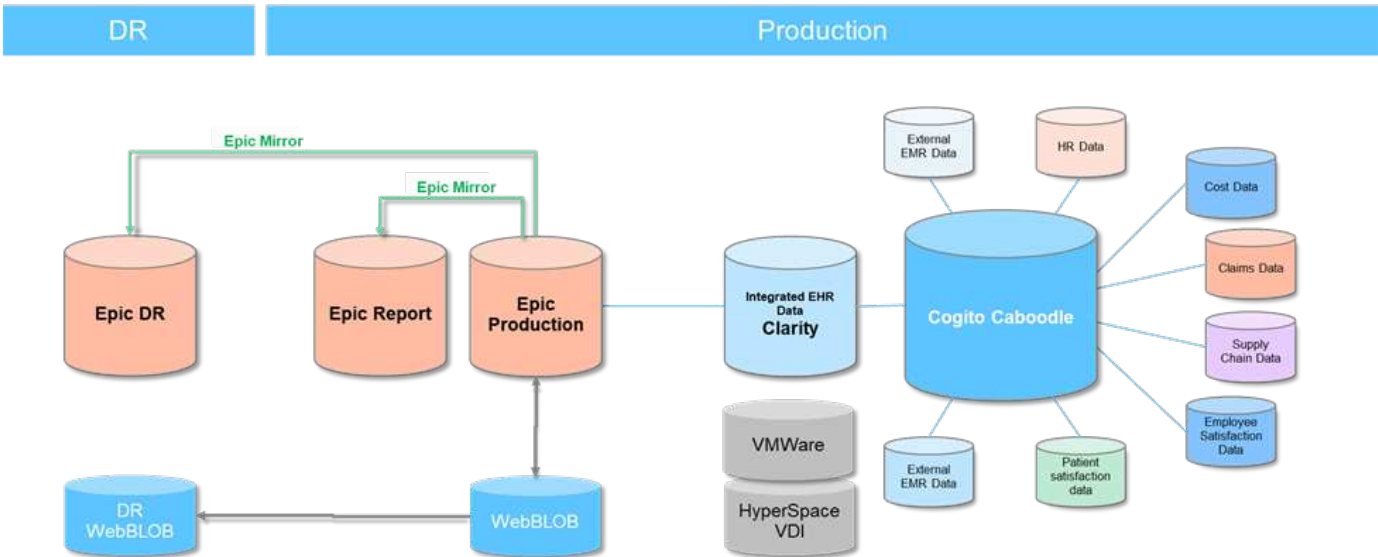
In addition to Iris ODB data servers, Epic software environments commonly include other components such as the following and as shown in the figure below:

- 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.

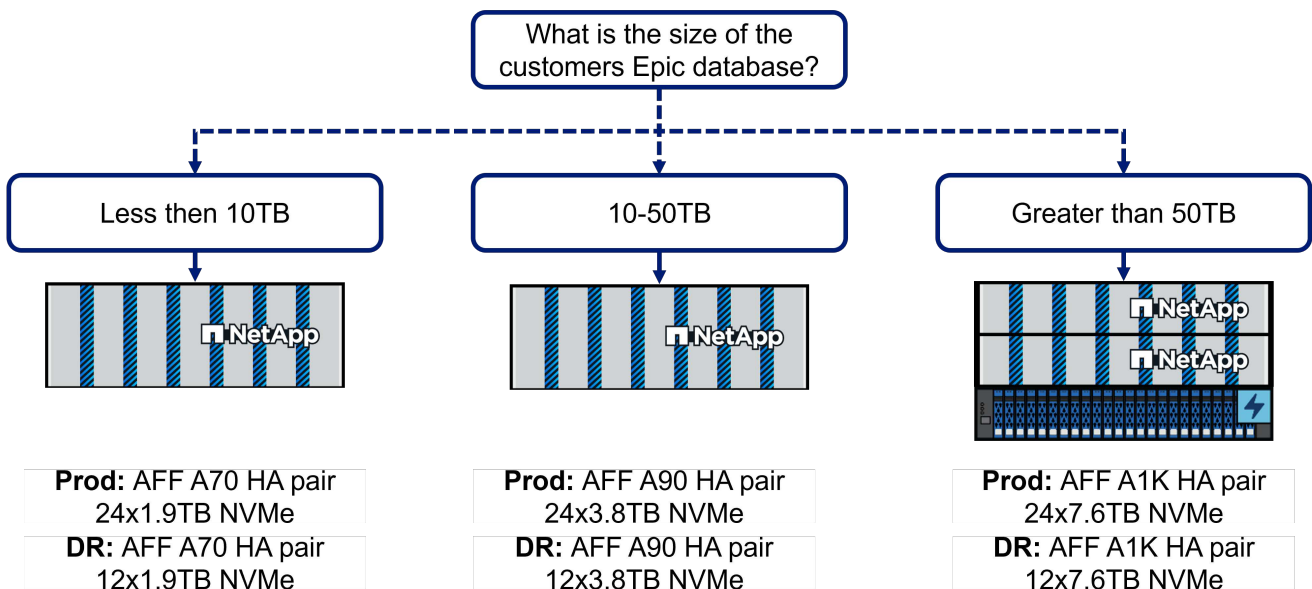
You can use the diagram shown below to select a small-medium-large Epic storage architecture. These designs include running all workloads listed in the Hardware Configuration Guide. The sizing tree is based on data from over 100 hardware configuration guides and should be mostly accurate estimate.

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.

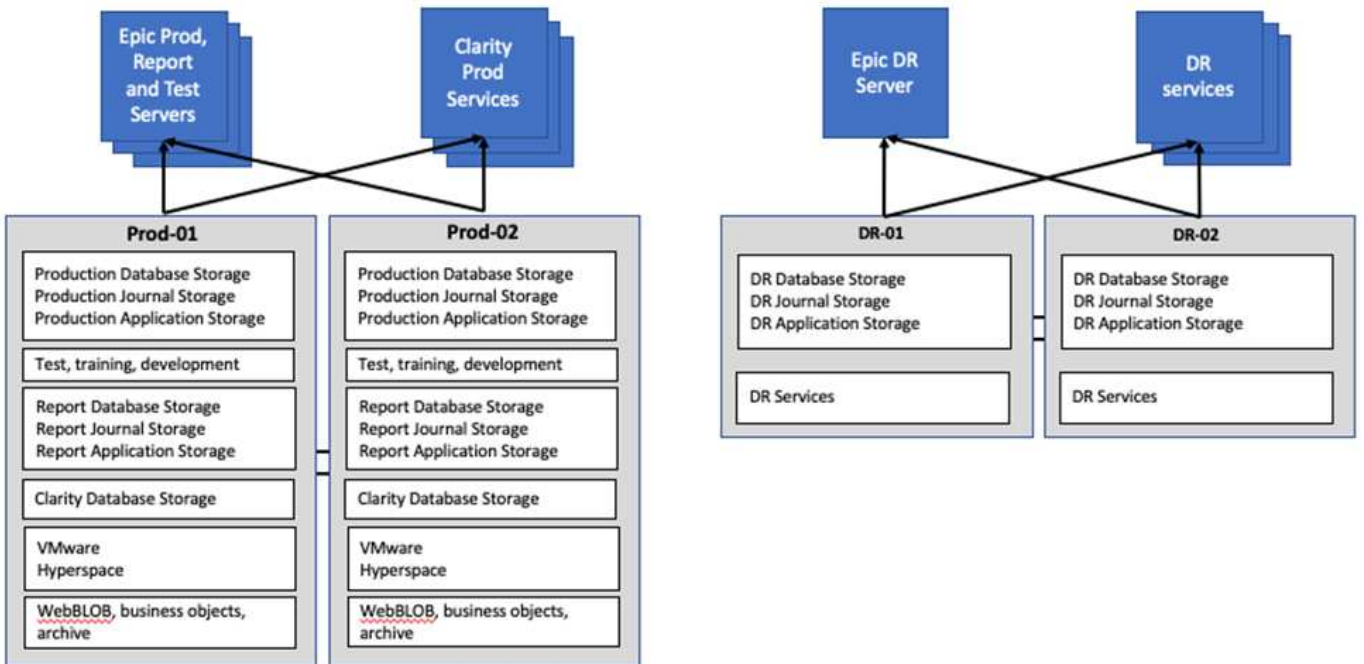
NetApp guarantees a minimum floor level performance by accepting the SLM recommended AQuoS policies. Epic supports consolidating storage pools on ONTAP onto significantly less hardware. For more information, see the Epic quarterly SPATS document. Basically, pool1, pool2, and NAS1 (listed in the Epic 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



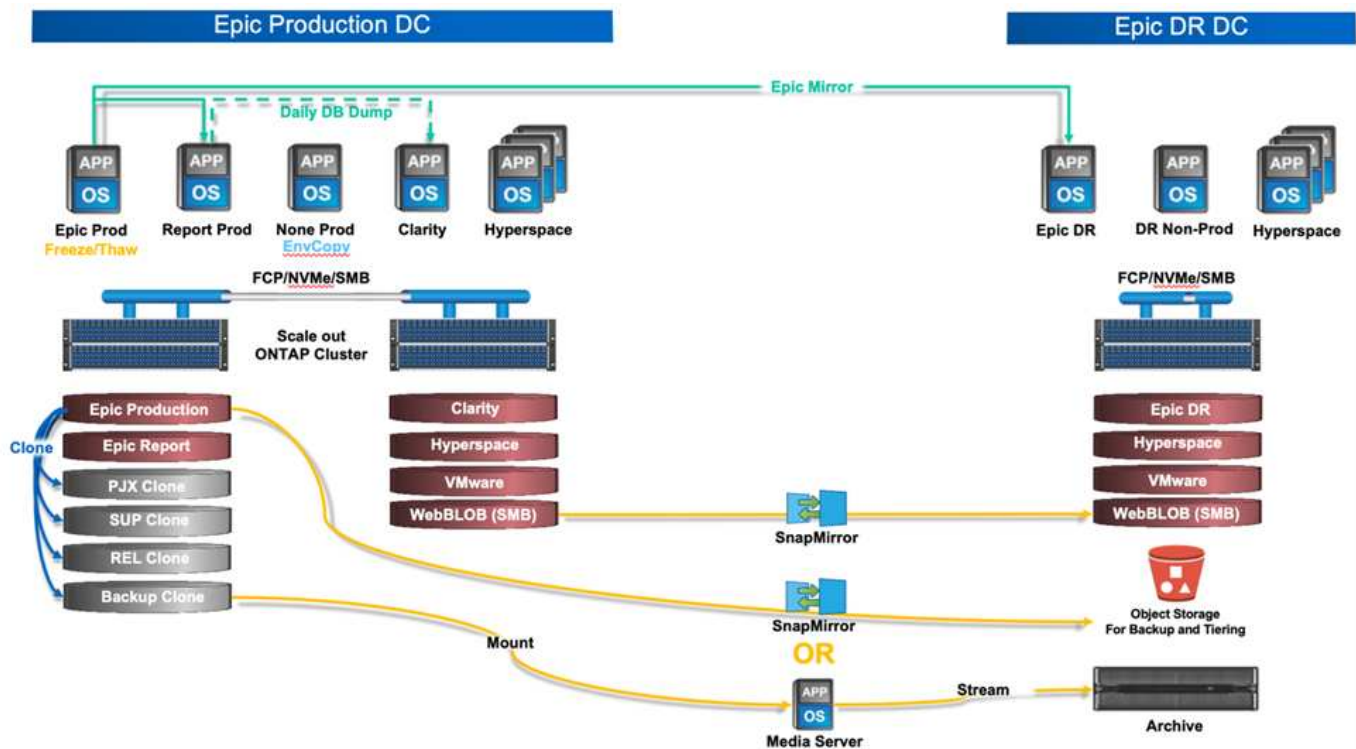
Epic six-node architecture

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.

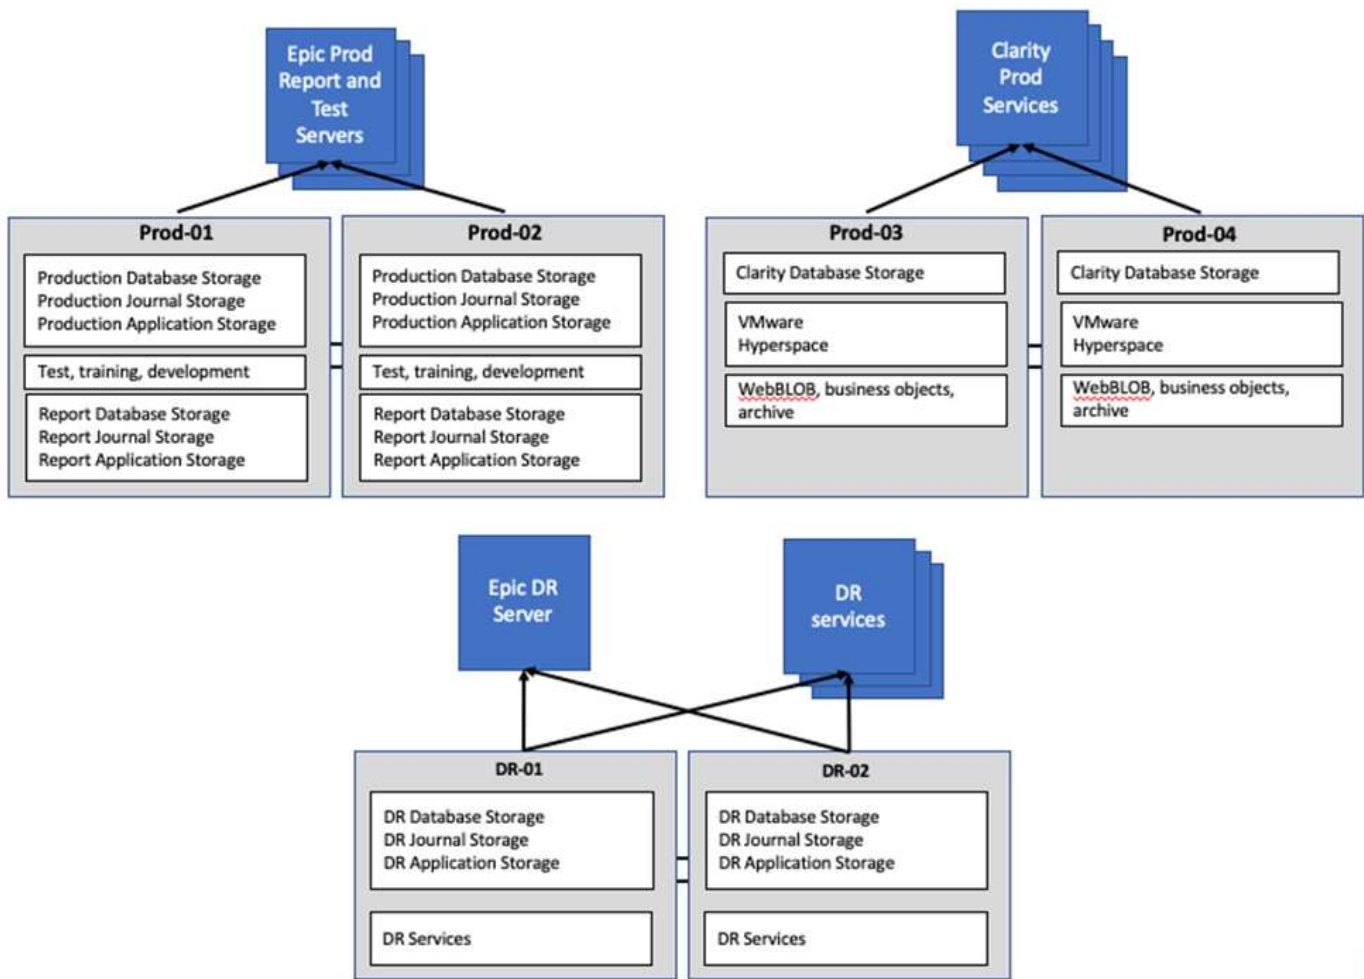
This architecture offers the best performance and capacity balance on the cluster. Epic Production, Epic

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

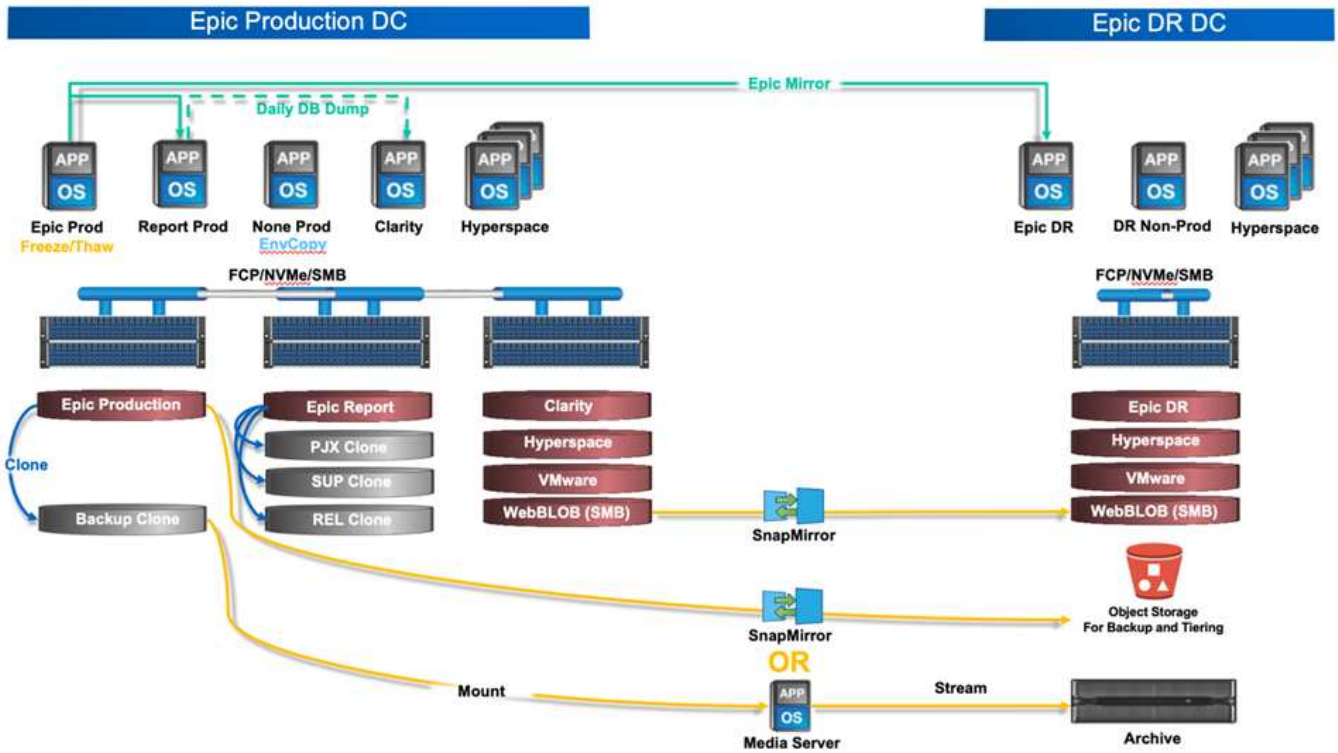


Epic eight-node architecture

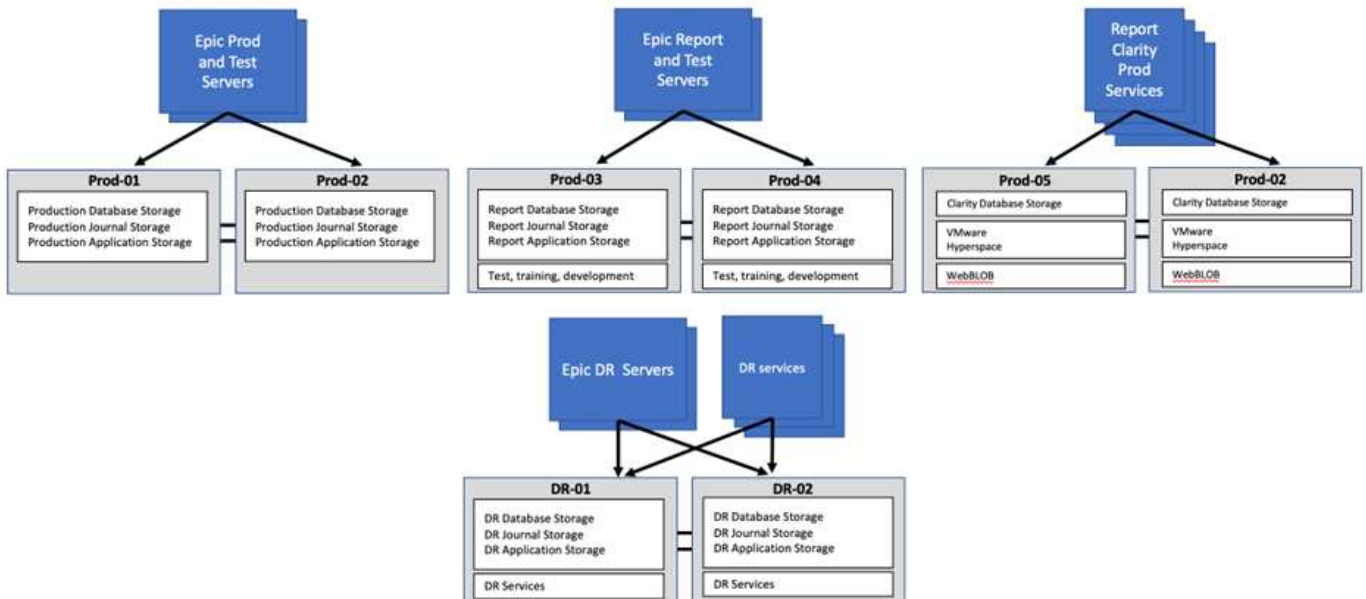
The figures below show the scale-out eight-node architecture. Again, you can start with four nodes and scale to six nodes and continue to scale to eight nodes and beyond. This architecture offers the best balance of performance and capacity across the six nodes in production.

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



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.