

# Optimizing IBM Storwize V7000F for Epic EHR applications

---

## Overview

### Challenge

- Address high performance and low latency application and user requirements for Epic EHR platform

### Solution

- Deploy a tested and proven all-flash storage solution for Epic EHR: IBM Storwize V7000F is ideally suited for midsized workload challenges including operational and analytical databases
  - Deliver high efficiency with reduced costs, superior performance, scalability, and high availability.
- 

*A configuration and best practices guide for this flash-based, advanced storage platform for Epic healthcare applications*

## Introduction

Healthcare organizations are experiencing massive data growth and in order to ensure continued and uninterrupted access to their Epic electronic healthcare record (EHR) data, healthcare providers need to consider proven storage solutions such as IBM® Storwize® V7000F to address their high performance, and low latency requirements.

Epic's integrated EHR software covers all aspects of healthcare operations, including patient engagement, clinicals, specialties, distance care, mobile, managed care, billing and revenue, interoperability and government regulations, in addition to patient record access. Epic implementations at client sites are based on the following two inclusive technologies that contribute to the overall workload throughput:

- **Operational database:** The online transaction processing (OLTP) database runs Caché from InterSystems Corporation as the main database engine. Caché database performance is by far the most critical aspect to meeting user requirements.
- **Analytical database:** The analytical databases typically run on Microsoft SQL Server or Oracle database software. This component has the highest bandwidth requirement.

## IBM Storwize V7000F

The IBM Storwize family provides all-flash and hybrid storage solutions with common functionality, management and mobility. These solutions are built with IBM Spectrum Virtualize™ to enable efficient, proven data virtualization across IBM and non-IBM storage systems.



IBM Storwize V7000F, a product within the IBM Storwize family, is a cost-effective flash storage system allowing clients to optimize capital and operational expenses and enable maximum efficiency and flexibility for heterogeneous storage systems. For more details on Storwize V7000F, see: <https://www.ibm.com/us-en/marketplace/storwize-v7000f>.



Figure 1: IBM Storwize V7000F

## Storage requirements for Epic

The Epic application can generate considerable data traffic with its operational and analytic databases. Generally, on any new IBM or Epic related opportunity, a copy of the full hardware configuration guide that has been developed for the customer by Epic is required. In generic terms, the hardware configuration guide is a *requirements document*, and third-party vendors build and evaluate solutions based on the Epic documented requirements.

Epic reinforces this practice to ensure that the solution developed by the hardware vendor (IBM or any other similar vendor) meets the customer's server and storage requirements, as defined by Epic. The hardware configuration guide provides specifications for the complete Epic environment, including processing, storage and network requirements. This allows the IBM server and storage teams to properly develop and review IBM Power Systems™ and IBM storage solutions for inclusion in customer proposals.

This paper acts as a prerequisite document to solution design, planning, and delivery process for server and storage solutions in an IBM and Epic opportunity. This hardware configuration guide creation involves a top-down process, taking clinical inputs to determine server and storage requirements for the solution. It is not a bottom-up process that starts with a certain IBM product selection, and then builds up to an IBM and Epic solution to meet the client's clinical requirements.

## Best practices recommendations with IBM Storwize V7000F

The following best practices are recommended for an Epic configuration, with IBM Storwize V7000F for optimal performance, and operations:

- With Storwize V7000F, the database files need to be isolated from the journal files and cannot co-exist under the same controller. For recoverability purposes, it is important that an up-to-date copy of the production journal files be available on a second controller outside the production array.
- Usage of an even number of VDisks for operational data and similarly an even number of VDisks for analytical data is recommended.
- Distributed RAID6 is the recommended Redundant Array of Independent Disks (RAID) array configuration with all flash drives on Storwize V7000F for the operational and analytical databases.
- While RAID10 was recommended in the yesteryears, with the 15K RPM spinning drives, it is now considered not relevant with all-flash based arrays (such as Storwize V7000F) that can deliver latencies at less than 1 millisecond.
- Distributed RAID allows a RAID6 array to be distributed over a larger set of drives. Previously, if you created a RAID6 array over eight drives, the data was striped across them with each stripe having a data strip on six of the drives and a parity strip on the seventh and eighth drive. Additionally, unique IBM algorithms add the distributed sparing functionality. It is the idea that instead of having a spare sitting idle on the side that is not being used, each drive in the Storwize V7000F array gives up some of its capacity to make a total spare.
- It is recommended to create 16 VDisks to be presented to the server for each Cache instance. This must be verified to ensure that the volumes are equally spread across both the Storwize V7000F node canisters within an I/O group.
- Usage of an even number of VDisks for operational data and similarly an even number of VDisks for analytical data is recommended.
- Validate that the preferred Node ID of the VDisks is alternatively even and odd, as illustrated in Figure 2.

| Name      | Preferred Node ID | State    |
|-----------|-------------------|----------|
| EpicCmp1  | 2                 | ✓ Online |
| EpicCmp2  | 1                 | ✓ Online |
| EpicCmp11 | 2                 | ✓ Online |
| EpicCmp12 | 1                 | ✓ Online |
| EpicCmp13 | 2                 | ✓ Online |
| EpicCmp14 | 1                 | ✓ Online |
| EpicCmp15 | 2                 | ✓ Online |
| EpicCmp16 | 1                 | ✓ Online |
| EpicCmp3  | 2                 | ✓ Online |
| EpicCmp4  | 1                 | ✓ Online |
| EpicCmp5  | 2                 | ✓ Online |
| EpicCmp6  | 1                 | ✓ Online |
| EpicCmp7  | 2                 | ✓ Online |
| EpicCmp8  | 1                 | ✓ Online |
| EpicCmp9  | 2                 | ✓ Online |
| EpicCmp10 | 1                 | ✓ Online |

Figure 2: Node ownership of volumes

- Data compression is not recommended for Epic applications.
- Proper sizing of the overall cache requirements should include related copy services requirements. At the time of IBM FlashCopy® resynchronization, FlashCopy will use Storwize V7000F write cache as well. Therefore, it is important to take it into account for the overall cache sizing (to include cache for the front-end application I/O, **plus** additional cache for the back-end FlashCopy services).
- The incremental FlashCopy target should follow the same rules as the production database to use the maximum available cache.
- The FlashCopy target volumes should be created with the same size of production volumes.
- The ownership of FlashCopy target volumes should preferably be the same as that of the FlashCopy source volumes.
- The incremental FlashCopy target should follow the same rules as the production database to use the maximum available cache.

| Name        | Preferred Node ID | State    |
|-------------|-------------------|----------|
| EpicCmp1    | 2                 | ✓ Online |
| EpicCmp1_01 | 2                 | ✓ Online |
| EpicCmp2    | 1                 | ✓ Online |
| EpicCmp2_01 | 1                 | ✓ Online |
| EpicCmp3    | 2                 | ✓ Online |
| EpicCmp3_01 | 2                 | ✓ Online |
| EpicCmp4    | 1                 | ✓ Online |
| EpicCmp4_01 | 1                 | ✓ Online |

Figure 3: Node ownership of FlashCopy source and target volumes

- The FlashCopy mappings or relationships should be configured with the *incremental* option set to *on* to allow the FlashCopy mapping to record the changes from the previous FlashCopy on the same mapping. This allows the FlashCopy background copy function to copy only the changes from the previous FlashCopy operation.
- It is recommended to set the FlashCopy grain size to 64 KB.

- In a FlashCopy relationship, with the incremental option switched on, it is important to note the background copy rate along with the timing when you issue a FlashCopy operation. The higher background copy rate increases the speed of synchronization of FlashCopy volumes.
- At the time of the FlashCopy resynchronization, FlashCopy uses the write cache as well. Therefore, it is important to take it into account the overall Storwize V7000F cache sizing to include cache for the front-end Epic application I/O, plus additional cache for the back-end FlashCopy services.
- When using FlashCopy to create backup, it is essential to configure FlashCopy to be incremental and to tune the background copy rate from a default value of 50% to a value that will allow the background copy to finish in time without impacting production performance during the backup window.
- The recommendation is to set the background copy rate between 70% to 80%. This assumes, it is a small environment, for a small environment, that is configured with a typical database size of about 20TB, where the FlashCopy operation is executed about every eight hours. This setting usually gives the optimal balance between FlashCopy copy time and production I/O time when running with incremental FlashCopy operations.
- It is important to understand the relationship between the input/output operations per second (IOPS) performance, background copy rate, and the frequency of FlashCopy operations. The more frequent the FlashCopy operation the lesser the amount of data changes to be applied, and hence quicker the synchronization with an average background copy rate. The lesser the amount of changes, the lesser the task for background copy and lesser the latency for production I/O.
- Based on the above observation, the three important variables to consider for configuring FlashCopy services:
  - The number of IOPS
  - Background copy rate
  - The time between two FlashCopy entities.
- It is important to balance all the three variables to arrive at the best possible configuration for a balanced Cache workload, depending on the IOPS requirement.

## General Server recommendations for optimal configuration of databases

The IBM Storwize V7000F system has been independently tested successfully on Intel x86 platforms running RHEL, IBM Power Systems running AIX, and IBM/Cisco VersaStack converged architectures, running RHEL operating systems. In all these different tests with different environments and operating systems, Storwize V7000F has consistently delivered acceptable numbers.

The following server recommendations for optimal configuration of databases are general and might vary from platform to platform and OS to OS.

- It is recommended to configure native or IBM *pcmpath* multipathing driver on the host.
- The queue depth of each physical disk must be configured to 1024 or to a maximum configurable value.
- A single logical volume group should be created using all the 16 logical unit numbers (LUNs) mapped on the host.
- Using the logical volume, eight logical volumes should be created, each using 12% of the space of the volume group.
- The remaining 4% can be used for the ninth logical volume that will be used for write image journaling.
- The logical volumes must be configured in a striped mode to enable uniform I/O distribution.
- Using the first eight logical volumes, eight file systems for production data should be created.
- Using the ninth logical volume, a file system for write image journaling should be created.

### A typical Epic configuration with IBM Storwize V7000F

Figure 4 depicts the high-level configuration of various components of the Epic solution built with IBM Storwize V7000F:

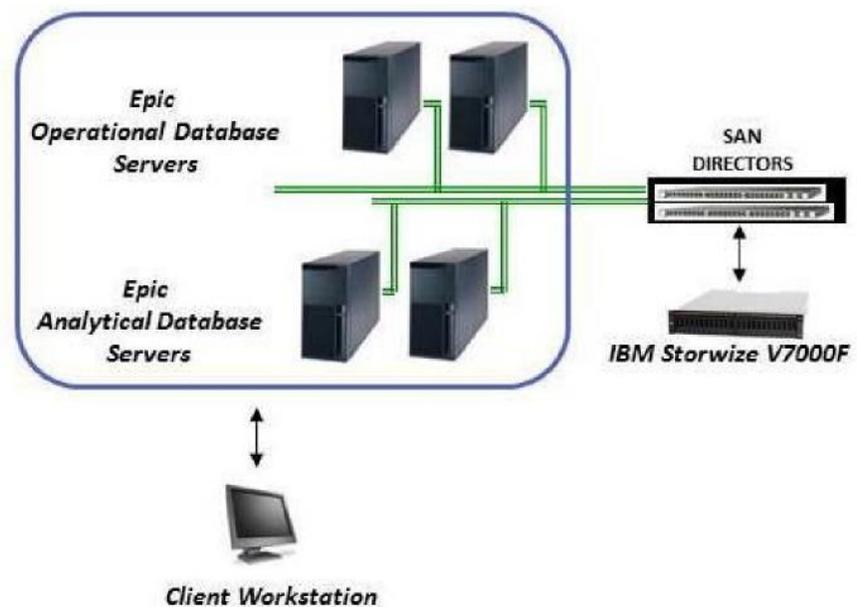


Figure 4: Simple Symmetric Multiprocessing (SMP) architecture where the end users access the main Epic production database servers directly, supported by IBM Storwize V7000F

### Test setup

This section provides details about the simulated lab environment for the Epic with Storwize V7000F configuration. A simulated lab environment was used to validate the information pertaining to the best practices recommendations detailed in the previous sections.

**Test system includes:**

- IBM Storwize V7000F storage system
- 40 TB flash disks

**Operational consideration and full disclosure information:**

- The tests were run with an installed patch, as Storwize V7000F was tested with code version 7.8.1.0.
- This patch takes advantage of some newer code capabilities (mostly internal features).
- The patch is not required in Storwize code version 8.1 or later because version 8.1 code is expected to be in general availability (GA) with this patch already included in the code on September 22, 2017.
- To ensure backwards compatibility, this patch will be required for code 7.8.1.0 or earlier.

**Server:**

- One Lenovo M4 (32 processors, 180 GB, four 16 Gbps FC ports, RHEL 7.3) server

## **Extending the solution using additional features for greater benefit**

Additional benefits such as improved data protection, on IBM Storwize V7000F can be achieved easily with the use of other IBM products and features:

- IBM FlashCopy along with IBM Spectrum Protect™ can provide a valuable backup and archive solution for the Epic database.
- Management, Operation, and Automation of Epic backups can be enabled using IBM Spectrum Copy Data Management™.

## **Acknowledgments**

Many thanks to the IBM offering management team, IBM Storwize development team, IBM client executives, IBM Systems members, and other associates worldwide, who contributed with their recommendations and excellent guidance, during the test run and review process and enabled successful completion and validation of this qualification project.

## Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld  
[ibm.com/partnerworld/systems](http://ibm.com/partnerworld/systems)
- IBM Redbooks  
[ibm.com/redbooks](http://ibm.com/redbooks)
- IBM Publications Center  
[www.elink.ibm.link.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US](http://www.elink.ibm.link.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US)
- IBM Storwize V7000 Knowledge Center  
[https://www.ibm.com/support/knowledgecenter/ST3FR7\\_7.8.1/com.ibm.storwize.v7000.781.doc/v7000\\_ichome\\_781.html](https://www.ibm.com/support/knowledgecenter/ST3FR7_7.8.1/com.ibm.storwize.v7000.781.doc/v7000_ichome_781.html)
- About Epic  
<http://www.epic.com/software>

## About the authors

**Prashant Avashia** is a senior architect for storage and software-defined infrastructure in IBM Systems Group. He has architected and delivered cloud, analytics, and software-defined solutions for global healthcare clients. You can reach Prashant at [pavashia@us.ibm.com](mailto:pavashia@us.ibm.com).

**Sandeep Zende** is an architect for healthcare solutions in IBM Systems Group. He has architected and delivered storage migrations, disaster recovery solutions, and software-defined solutions for several clients, including healthcare clients. You can reach Sandeep at [sandeep.zende@in.ibm.com](mailto:sandeep.zende@in.ibm.com).



---

© Copyright IBM Corporation 2017  
IBM Systems  
3039 Cornwallis Road  
RTP, NC 27709

Produced in the United States of America

IBM, the IBM logo and [ibm.com](http://ibm.com) are trademarks or registered trademarks of the International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked items are marked on their first occurrence in the information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other product, company or service names may be trademarks or service marks of others.

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in the IBM operates.

