

IBM® Storage

Managing Electronic Health Records with Epic, IBM Storage FlashSystem, and IBM Storage Sentinel

The IBM logo, consisting of the letters 'IBM' in a bold, black, sans-serif font. Each letter is composed of horizontal stripes, with the 'I' having three stripes, the 'B' having six stripes, and the 'M' having four stripes.

© Copyright International Business Machines Corporation 2023.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this document	1
Executive summary	2
Scope	3
Prerequisites	3
Getting started: Epic EHR solution with IBM Storage FlashSystem.	4
Epic	4
InterSystems IRIS	4
Typical Epic configuration	8
Lab setup	15
Setting up an IBM Storage Copy Data Management and IBM Storage Sentinel environment to scan Epic databases	18
Summary	28
Authors	28
Resources	29
Notices	31
Trademarks	32
Terms and conditions for product documentation	33
Applicability	33
Commercial use	33
Rights	33
Privacy policy considerations	33



About this document

The information in this blueprint is intended to facilitate the deployment of IBM® Storage FlashSystem® for the Epic Corporation electronic health record (EHR) solution. It describes the requirements and specifications for configuring IBM Storage FlashSystem and its parameters.

This document also describes the steps that are required to configure the server that hosts the EHR application. To complete these tasks, you must have a working knowledge of IBM Storage FlashSystem and Epic applications.

Also, this document describes the steps that are required to configure IBM Storage Sentinel for cyber resiliency. To complete these tasks, you must have a working knowledge of IBM Storage Copy Data Manager.

Note: The information in this document is distributed on an “as is” basis without any warranty that is either expressed or implied. Support assistance for the use of this material is limited to situations where IBM Storage FlashSystem storage devices are supported and entitled, and where the issues are not specific to a blueprint implementation.

Executive summary

Data is flowing into healthcare systems at an unprecedented rate. Between EHRs, digitized diagnostics, and wearable medical devices, it is estimated that the average person produces more than 1 million gigabytes of health-related data in their lifetime.¹

All this new data creates tremendous potential for healthcare organizations to understand their patients better, and to provide them with more personalized, evidence-based care.

At the same time, managing this immense volume of data can be a serious challenge for organizations that are not adequately prepared. Adopting next-generation healthcare IT solutions today is the first step toward making the most of this new healthcare data while setting the stage to pursue cognitive healthcare in the future.

EHRs are one area where managing exponential data growth is key. EHRs give providers rapid access to patient data when and where they need it, but only if they are supported by an effective data storage solution.

To help healthcare organizations achieve continuous access to EHR data, IBM works closely with [Epic](#), a leading provider of EHR solutions to ensure that IBM Flash Storage offerings can support current and anticipated Epic workload requirements.

By testing IBM Storage Systems with the Epic EHR platform, IBM can offer solutions to joint customers of the two firms, which help them reduce risk, accelerate decisions and actions, improve patient care, and overcome the challenge of massive data growth.

The business and financial effects of cyberattacks continue to rise. Cyberattacks can occur in various ways. They can take many different forms and continue to evolve. Whether the cyberattacker's goal is targeted at stealing confidential customer data, holding valuable information for ransom, or wiping out data, organizations must have an overall cybersecurity strategy in place.

Traditional high availability and disaster recovery (HADR) approaches to data protection work well for their intended purposes, but they are inadequate to protect against cyberattacks. Storage-based remote replication for HADR replicates all changes (malicious or not) to the remote copy.

IBM Storage Sentinel is a cyber resiliency solution that is designed to help organizations enhance ransomware detection and incident recovery. IBM Storage Sentinel automates the creation of immutable backup copies of your data, and then uses machine learning to detect signs of possible corruption and generate forensic reports that help you quickly diagnose and identify the source of the attack. Because IBM Storage Sentinel can intelligently isolate infected backups, your organization can identify the most recent verified and validated backup copies, which accelerate your time to recovery.

¹ Source: IBM Research®

Scope

This blueprint provides the following information:

- ▶ A solutions architecture and the related solution configuration information for the following essential software and hardware components:
 - IBM Storage FlashSystem
 - IBM b-type Storage Networking
 - Dell PowerEdge R750 server
 - InterSystems IRIS
 - Epic Software Suite
 - IBM Storage Copy Data Management
 - IBM Storage Sentinel
 - IBM Safeguarded Copy
- ▶ Detailed technical configuration steps for configuring IBM Storage FlashSystem
- ▶ Server configuration details for IRIS database and Epic applications

This technical report does *not* include the following items:

- ▶ Provide performance analysis from a user perspective.
- ▶ Replace any official manuals and documents that are issued by IBM.
- ▶ Explain the installation and configuration of IRIS and Epic applications.

Prerequisites

For this blueprint, it is assumed that the reader has basic knowledge of the following topics:

- ▶ IRIS
- ▶ Epic applications
- ▶ Intel x86 running VMware ESXi8
- ▶ VMware guest running Red Hat Enterprise Linux (RHEL) Version 8 for IRIS and Epic applications
- ▶ VMware guest running SUSE Linux Enterprise 15 SP4 for IBM Storage Sentinel
- ▶ IBM Storage FlashSystem all-flash storage arrays
- ▶ IBM Storage Copy Data Management
- ▶ IBM Storage Sentinel
- ▶ IBM Safeguarded Copy

Getting started: Epic EHR solution with IBM Storage FlashSystem

This section describes the Epic implementation of EHRs, and the components that are involved in the solution.

Epic

The integrated EHR software from Epic 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 software implementations use the following two main database technologies:

- ▶ Operational database: The online transaction processing (OLTP) database runs IRIS from InterSystems Corporation as the main database engine.
- ▶ Analytical database: Analytical databases typically run on Microsoft SQL Server or Oracle database software.

Although the analytical database component has the highest bandwidth, the IRIS (operational database) performance is by far the most critical aspect to user experience. Therefore, it is here that most of the attention on storage performance is focused.

InterSystems IRIS

To better understand the Epic environment, the underlying database system that is used by Epic must be examined because IRIS interfaces directly with the IBM hardware. Virtually the entire Epic environment relies on IRIS to accomplish its work.

The IRIS database engine is based on a B-Tree data storage structure. Internal to IRIS, the user data is managed as 8 KB blocks.

IRIS maintains a global buffer cache in the computer's real memory. All transactions (read/write) between the user and the database are read in or written from the global buffer. The global buffer acts as a storage cache, which reduces operating system-level input/output (I/O) requests. It also acts as a global data locking communication system that provides mutually exclusive access to data that is referenced or changed by multiple users.

Data that is referenced by one or more users initially is read from the storage device (disk) into the global buffer. The data objects are now accessible for repeated operations, including updates to the contents. Access and occur rapidly because the data is kept in memory.

As data blocks in the global buffer are updated, they are considered "dirty." IRIS flushes the dirty blocks at a regularly scheduled interval of about 80 seconds, or when the percentage of dirty blocks in the total global buffer exceeds the internal threshold, whichever occurs first.

IRIS features dedicated write daemon processes that perform the update operations. The flush of the dirty blocks is typically referred to as a *write daemon cycle*.

IRIS uses a two-phase update technique for database updates:

- ▶ In the first phase during the write daemon cycle, updates are first written to the CACHE.WIJ file (WIJ).
- ▶ The updates to the database file occur only after the WIJ updates complete successfully as part of the second phase.

After all database updates are committed, IRIS marks the WIJ as “clean.” The first WIJ writes are sequential writes of 252 KB blocks.

When the global buffer is not being flushed, the I/O requests that are issued by IRIS are strictly read-only. The I/O requests that are issued do not include the IRIS Journal File, which is a time-sequenced log of database changes that is written out to disk sequentially.

Although writing to the Journal File is a continuous process, it does not require many resources. Therefore, the random read operations can occur with little or no interference.

Every 80 seconds, the flush or write burst is started from the global buffer. While IRIS continues to issue 100% read requests, the database engine also generates many requests in a noticeably short period. Therefore, Epic has strict read latency guidelines to avoid degrading user performance.

Write latencies also become increasingly important for high-end scalability. For large implementations, this write latency can lead to a clear conflict between optimal read performance and optimal write performance during intense write bursts.

IBM Storage FlashSystem

For decades, IBM offered a range of high-performance, ultra-low latency storage solutions. Now, IBM Storage FlashSystem combines the performance of flash and NVM Express (NVMe) with the reliability and innovation of IBM FlashCore® technology. It also includes the rich feature set of IBM Storage Virtualize in one powerful new storage platform that provides the following features:

- ▶ The option to use industry-standard NVMe flash drives or new third-generation IBM FlashCore modules (FCMs) with the line-speed internal performance, multi-dimensional data protection, and innovative flash management features that are provided by IBM FlashCore technology.
- ▶ The software-defined storage functions of IBM Storage Virtualize, which enables a full range of industry-leading data services, such as dynamic tiering, IBM FlashCopy® management, data mobility, and high-performance data encryption.
- ▶ SCSI UNMAP support and all the thin-provisioning, copy management, and efficiency features that customers expect from IBM Storage Virtualize-based storage.

The IBM Storage FlashSystem products deliver new levels of storage density. A key innovation involves the transformation of IBM FlashCore technology into a standard 2.5-inch solid-state drive (SSD) form factor with NVMe interfaces, form the basis of the storage array. *IBM FlashCore technology* refers to the IBM innovations that enable IBM Storage FlashSystem solutions to deliver consistent microsecond latency, extreme reliability, and a wide range of operational and cost efficiencies.

IBM FlashCore innovations include a hardware-accelerated nonvolatile memory (NVM) architecture and advanced flash management features, such as IBM Variable Stripe RAID technology, IBM engineered error-correction codes, and proprietary garbage-collection algorithms that increase flash endurance and accelerate performance while reducing latency.

IBM Storage FlashSystem use the powerful inline, hardware-accelerated data compression technology that provides consistent, high-performance data reduction across the full range of workloads.

Flexibility is built into the IBM Storage FlashSystem architecture. Customers can choose IBM FCMs in multiple capacities or opt for industry-standard NVMe-enabled flash drives. IBM FCMs can be complemented with Storage Class Memory (SCM) NVMe drive technology. SCM technology offers even lower latency, and when combined with IBM FCM drives, it can be used for even more demanding workloads.

Distributed RAID 6 (DRAID 6) helps rebuild data faster if a disk failure occurs by striping the data, parity, and spare capacity across the disks.

IBM Storage FlashSystem 9500 delivers 48 drives in a single 4U form factor, which increases the performance, capacity, and connectivity of the high-end IBM FlashSystem offering. It combines the NVMe storage of IBM FCMs with IBM Storage Virtualize software, and it is supported by IBM Storage Insights with IBM Storage Expert Care service and support options.

IBM Storage Networking b-type (Brocade) Gen 7 SAN infrastructure

IBM Storage Networking b-type SAN (Brocade) Gen 7 FC SAN platforms deliver far more than only speed and latency improvements. It can make protecting and managing your data center more efficient. Integrated security and autonomous SAN technology enables a cyber resilient network that safeguards your SAN against cyberattacks, IT disruptions, and disasters, while learning, optimizing, and healing on its own. These capabilities automate processes to ensure optimal performance and enable nonstop operations. Concurrently, they strengthen the level of security in your network to protect continuously against security threats.

With Gen 7, you can get the most out of the infrastructure you have while easing your migration to new storage technologies. Delivering 64G speeds and ultra-low latency, Gen 7 technology provides an instant performance boost for data-intensive applications and can handle SCM, all-flash arrays (AFAs), NVMe, and more.

IBM b-type (Brocade) Gen 7 platforms raise the bar for network availability and performance through its autonomous SAN capabilities. The Gen 7 SAN uses built-in telemetry to track, “learn”, and “baseline” traffic patterns so that the SAN to make accurate decisions about traffic and congestion management and deliver high availability with consistent performance and reliability. If an event occurs that might degrade application performance, the Gen 7 SAN works with end devices to pinpoint and resolve issues, often without human intervention and faster than a human might respond. For example, if a credit-stall results in slowed traffic between devices, the Gen 7 SAN can autonomously take immediate action to identify and quarantine the misbehaving device while redirecting traffic flows to resume normal operation.

Fibre Channel fabrics are secure by design because of controlled access between servers and storage, and isolation within the data center. IBM b-type (Brocade) Gen 7 technology provides integrated security that further reduces vulnerabilities from malware and hijacking attacks by automatically validating the integrity of the switch operating system, security settings, and hardware. IBM b-type (Brocade) Fabric OS (FOS) automatically provides the added level of security that is needed to ensure that your operations are protected, and extra security enhancements to validate the integrity and security of your IBM b-type (Brocade) hardware and software. These features are integrated and turned on automatically.

Features include Secure Boot, Brocade Trusted FOS (TruFOS) Certificates, FOS hardening with removal of root access, and automated distribution of SSL certificates through the IBM b-type (Brocade) SANnav Management Portal. IBM b-type (Brocade) TruFOS Certificates ensure that enterprises running IBM b-type directors and switches are covered by support; validated; and securely enabled to perform critical operations without having to worry about whether the operating system has been tampered with.

The IBM b-type (Brocade) SANnav Management Portal provides a complete view of the network, including traffic flows and health conditions across the SAN. Its intuitive, drill-down interface leverages IBM b-type (Brocade) Autonomous SAN technology and presents data in clear dashboards that enable a faster, more comprehensive SAN management experience. It also streamlines management workflows to accelerate the deployment of new applications, switches, hosts, and targets. In addition, security features are built into the IBM b-type (Brocade) SANnav Management Portal to help administrators protect their network.

Figure 1 shows the IBM b-type (Brocade) Gen 7 FC SAN family.




Figure 1 IBM b-type (Brocade) Gen 7 FC SAN family

Broadcom Emulex GEN 7 Fibre Channel host bus adapters

The Broadcom Emulex® Gen 7 Fibre Channel (FC) host bus adapters (HBAs) are designed to meet the requirements for enterprise applications and demanding data center workloads. As server designs expand their performance capability with increased processor core counts, improved memory density, and faster PCIe bandwidth the I/O, performance requirements of applications drive the storage performance bottleneck to the HBA and SAN. Emulex Fibre Channel HBAs are purpose-built for storage networks to meet the requirements for high availability, scalability, predictable performance, and low latency.

Compared to previous-generation HBAs, Emulex Gen 7 FC HBAs offer up to 2x higher bandwidth, 3x better latency, enhanced security, and operational efficiency for 32 GB Fibre Channel (GFC) and 64 GFC SANs. Gen 7 32GFC provides seamless compatibility to 16 GFC and 8 GFC networks. Emulex LPe36000-series HBAs are available with single and dual 64 GFC optics and provide seamless compatibility to 32 GFC and 16 GFC networks.

The Emulex Gen 7 HBA is based on a modern dynamic multi-core architecture that delivers the industry's highest reliability and performance. Its design uses fully redundant paths inside the ASIC connecting fast path data engines to reroute transactions to other fast path data engines if a failure of one of the engines occurs. Each HBA port has its own firmware image that can be updated without a restart to maximize uptime. This design provides the highest levels of dynamic performance and overall reliability compared to other obsolete HBA designs. The LPe35002 can deliver 12,800 MBps (two 32 GFC ports) or 25,600 MBps (two 64 GFC ports) full duplex. For the LPe36002 model, they deliver 3x better hardware latency, and the 64 GFC LPe36000-series HBAs deliver industry-leading performance up to 10 million IOPS.

Emulex Gen 7 HBAs support NVMe over Fibre Channel (NVMe/FC), which provides lower latency versus the traditional Fibre Channel SCSI protocol (SCSI FCP). Emulex Gen 7 HBAs can support NVMe/FC and SCSI FCP concurrently, which protects investment and enables data centers to move to end-to-end NVMe over FC SANs at their own pace.

IBM is uniquely positioned to offer Emulex Gen 7 end-to-end from host server to storage device. Emulex Gen 7 HBAs are available for storage connectivity on IBM Storage FlashSystem storage solutions. For server connectivity, IBM Power servers also have options for both 32 GFC and 64 GFC Emulex Gen 7 HBAs.

Typical Epic configuration

Epic uses the following fundamental architecture models:

- ▶ Single Symmetric Multiprocessing (SMP)

The single database server architecture provides the greatest ease of administration.

- ▶ Enterprise Cache Protocol (ECP)

This tiered architecture has a central database server with a single data storage repository. Unlike the SMP architecture, most processing needs are offloaded to application servers.

The application servers contain no permanent data. This architecture offers increased scaling over the SMP architecture.

Many Epic customers use the SMP architecture. Each architecture includes a production database server that is clustered in an active-passive configuration to a failover server.

Figure 2 shows a high-level configuration of various components of the Epic solution that is built with IBM Storage FlashSystem.

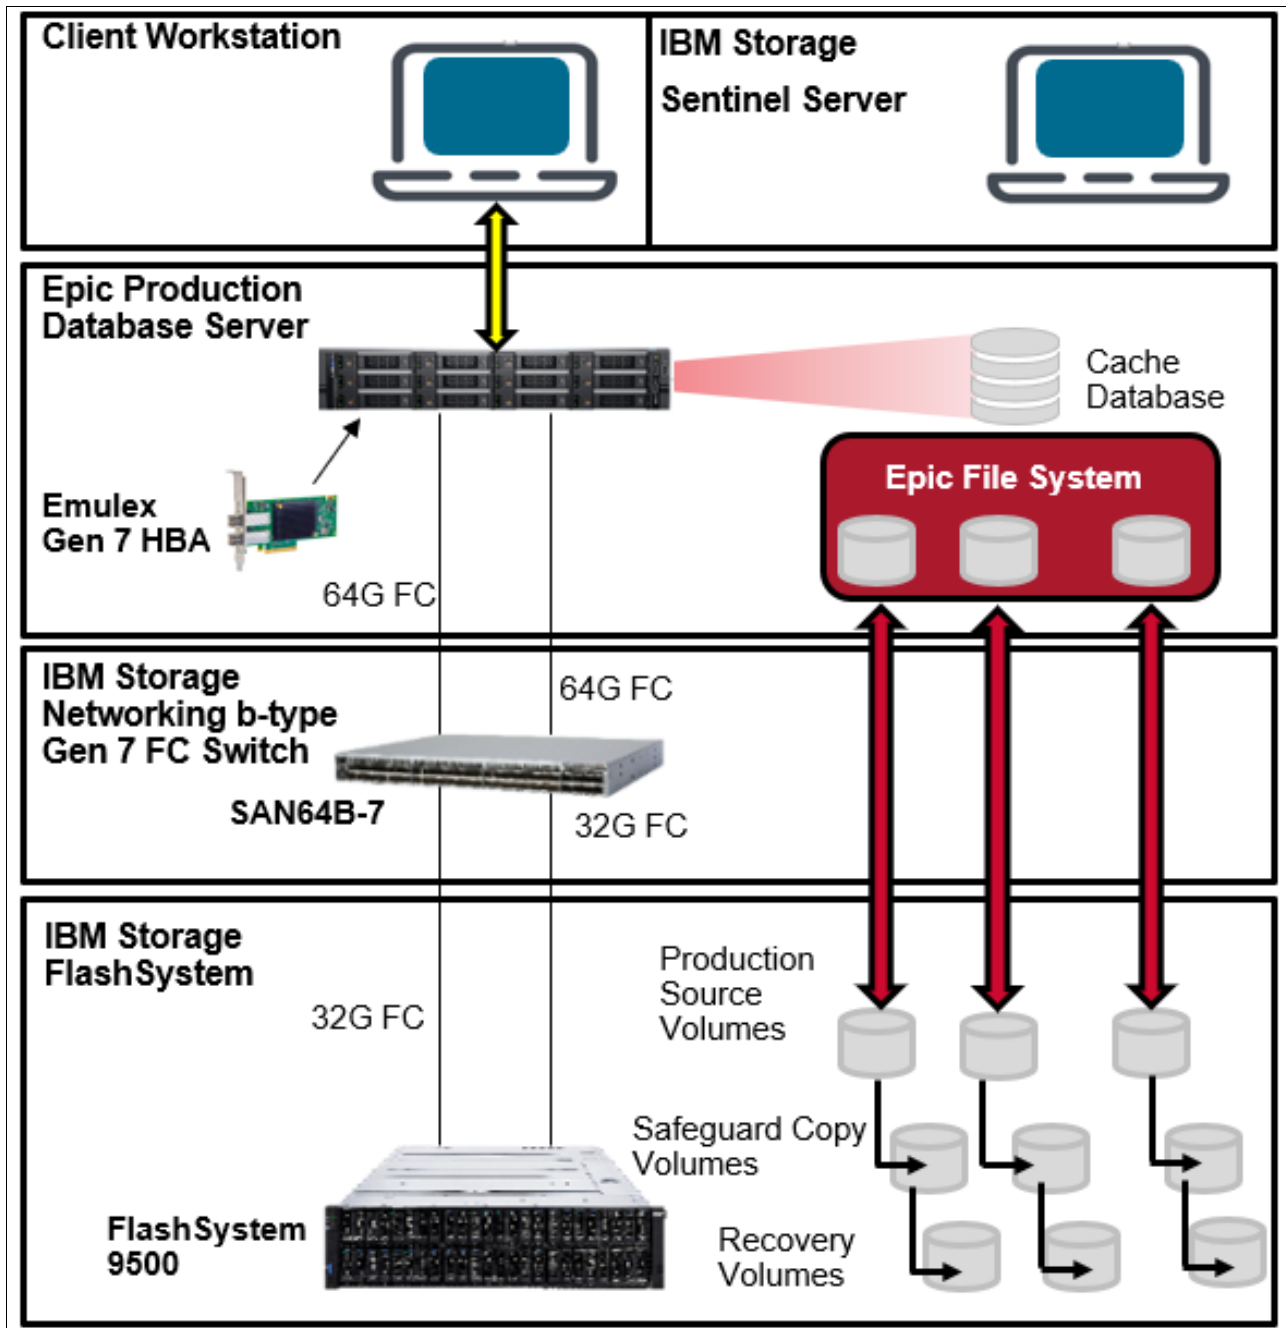


Figure 2 High-level configuration of various components of the Epic solution

Storage configuration for Epic

For more information about Epic storage configuration, see the latest [Epic Blueprint](#) for IBM FlashSystem or contact epicibm@us.ibm.com.

FlashCopy configuration recommendations

To accommodate SafeGuarded Copy, a FlashCopy configuration should be done as FlashCopy NoCopy relationships with the smaller, 64-KB grain sizes so that the SafeGuarded Copy relationships and the ad hoc or scheduled FlashCopy relationships maintain a contiguous inheritance scheme. Before implementing SafeGuarded Copy, see the latest [Epic Blueprint](#) for IBM FlashSystem or contact epicibm@us.ibm.com for this IBM FlashSystem Guide.

Note: If the storage configuration does not allow FlashCopy NoCopy, contact epicibm@us.ibm.com for a separate recommendation about how to implement SafeGuarded Copy.

Configuring SafeGuarded Copy

Before configuring SafeGuarded Copy, configure the FlashCopy relationships by using best practices and the smaller, 64-KB grain size. This action defines the grain size for the SafeGuarded Copy relationships.

The first stage of configuring SafeGuarded Copy for IBM Storage Sentinel is to create a SafeGuarded Child Pool for the protected volumes:

1. Select **Pools** → **Pools** menu or right-click and highlight **Actions**, and then select **Create Child Pool**, as shown in Figure 3.

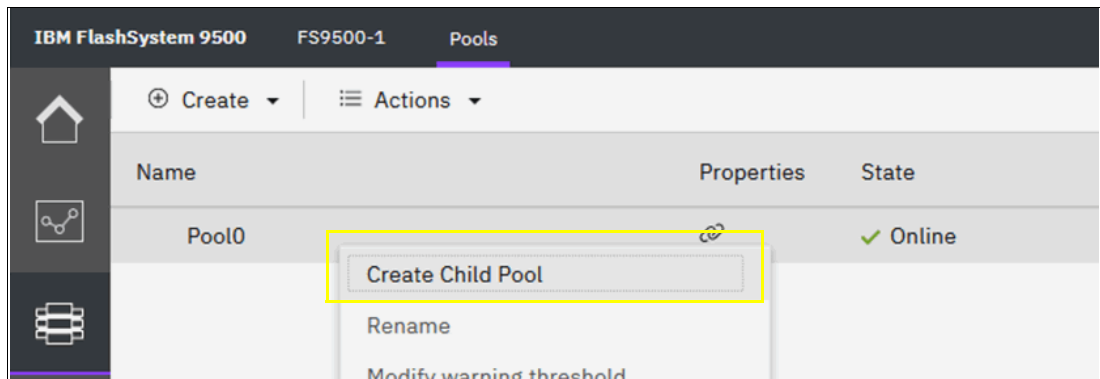


Figure 3 Create a SafeGuard Child Pool

2. Select **Child pool** and click **Next**, as shown in Figure 4.

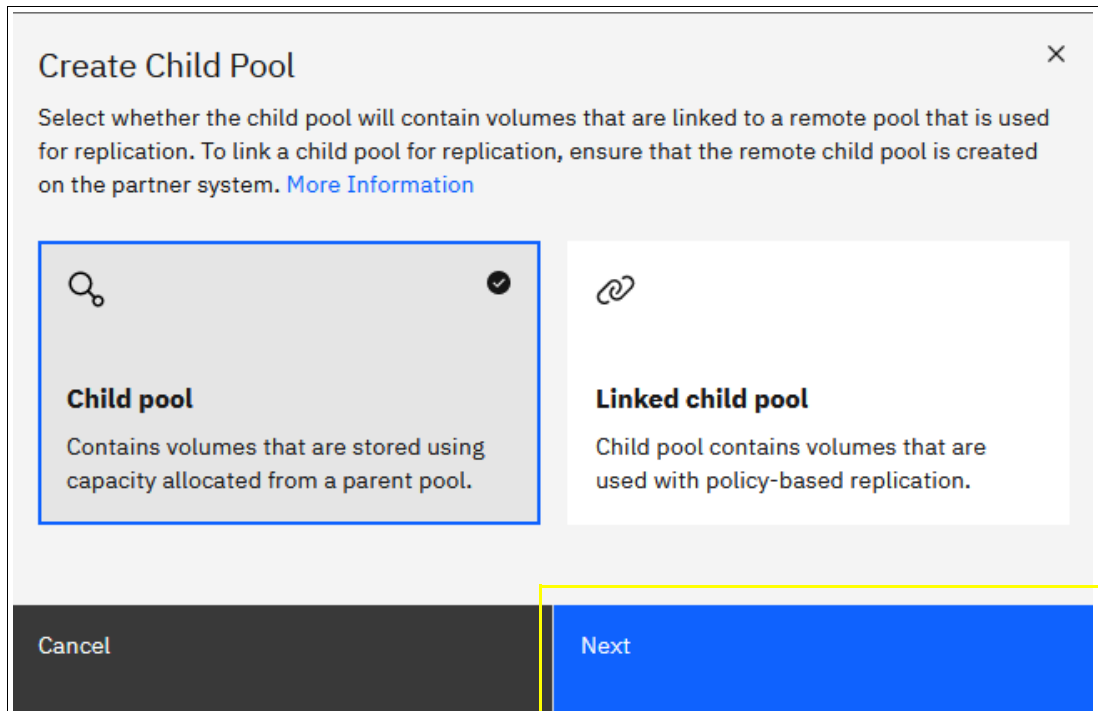


Figure 4 Selecting Child pool

- After you name the new child pool, define the capacity as at least the size of the effective capacity of the protected volumes to enable growth of the preserved data in a disaster, as shown in Figure 5. This effective capacity limit can be increased as necessary.

Figure 5 Child Pool parameters

- Click **Create** to create the Child pool. Figure 6 shows the Child pool that was created.

Name	Properties	State
Pool0		Online
SafeGuard		Online

Figure 6 Child pool created

- When the SafeGuarded child pool is created, create a Volume Group to associate with all volumes that are appropriate for each application instance. You can create the Volume Group selecting **Volumes** → **Volume Groups**, as shown in Figure 7 on page 13.

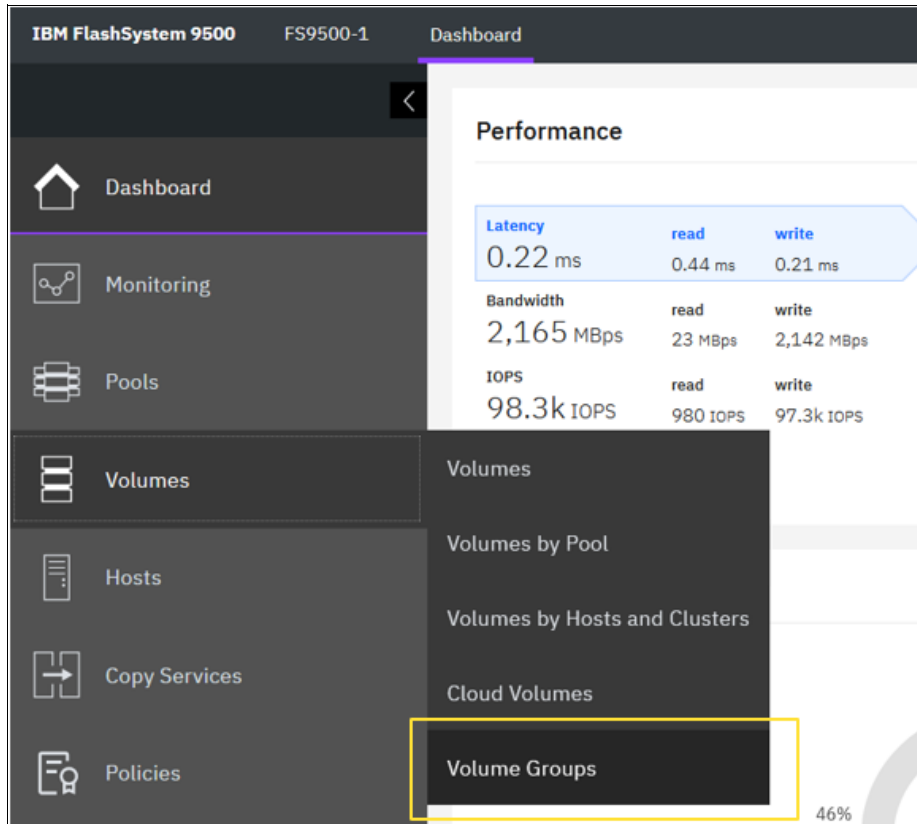


Figure 7 Selecting Volume Groups

6. Click **Create Volume Group**, as shown in Figure 8.



Figure 8 Clicking Create Volume Group

- Enter a name for the Volume Group, select **Choose existing volumes**, and click **Next**, as shown in Figure 9.

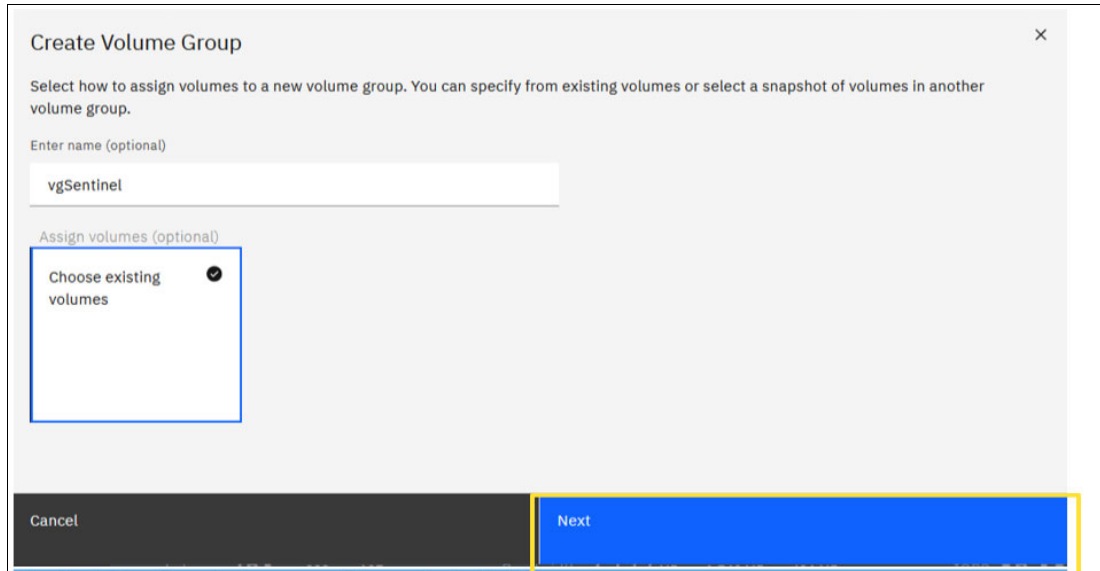


Figure 9 Create Volume Group

For IBM Storage Sentinel, all volumes that are associated with the Epic IRIS instance must be in the same Volume Group, which includes the OS instance that might be in a VMware data store. Figure 10 shows the volumes in the newly created Volume Group.

Name	State	Pool	Protocol Type	Capacity	Used Capacity
EPICDB0	Online	Pool0	SCSI	4.00 TiB	4.00 TiB
EPICDB1	Online	Pool0	SCSI	4.00 TiB	4.00 TiB
EPICDB2	Online	Pool0	SCSI	4.00 TiB	4.00 TiB
EPICDB3	Online	Pool0	SCSI	4.00 TiB	4.00 TiB
EPICDB4	Online	Pool0	SCSI	4.00 TiB	4.00 TiB
EPICDB5	Online	Pool0	SCSI	4.00 TiB	4.00 TiB
EPICDB6	Online	Pool0	SCSI	4.00 TiB	4.00 TiB
EPICDB7	Online	Pool0	SCSI	4.00 TiB	4.00 TiB
ESX_Datastore_0	Online	Pool0	SCSI	20.00 TiB	108.27 GiB

Figure 10 Volumes in the newly created Volume Group

Lab setup

This section provides more information about the simulated lab environment for configuring Epic with IBM FlashSystem.

The test system included the following components:

- ▶ IBM Storage FlashSystem 9500 Storage System (8.5.2.2 code level)
- ▶ 2 TB of cache (system total)
- ▶ 24 4.8 TB of IBM FCMs
- ▶ Eight 32-Gb Fibre Channel Ports
- ▶ One Dell PowerEdge R750 Server
- ▶ Dual Intel 40-core processors (80 cores total):
 - 256 GB of memory
 - Two Emulex LPe36002-M64 64 GFC HBAs
 - VMware ESXi 8.0) server
 - IBM b-type Gen 7 Fibre Channel switch (IBM SAN64B-7 / Brocade 720)

Figure 11 shows our lab environment.




Figure 11 Our lab environment

Next, we describe the process that was used to configure the storage, and the server that was used by the IRIS database and Epic applications.

SAN zoning was done by using a single initiator, dual-target method. Each of the host interface ports was zoned to a pair of ports on the IBM Storage FlashSystem 9500. For example, Host port 1 was zoned to the IBM FlashSystem Node 1 Port 1 and Node 2 Port 1. This method helps ensure balance across nodes for data access and provides redundancy if there is a port, adapter, or node failure.

To aid in this zoning configuration, multiple aliases were created. Each host port had an alias that was defined, and then each pair of IBM FlashSystem target ports (Node 1 Port 1 and Node 2 Port 1) were put into their own alias. Then, the zones were created by using a pair of aliases, one host port, and one of the target pair aliases. IBM FlashSystem uses N-Port ID Virtualization (NPIV) and presents multiple worldwide port names (WWPNs) to the SAN fabric. Use the WWPNs that permit host access for alias and zone creation. In this setup, the NPIV WWPNs that enabled host access through the SCSI protocol were used.

Example 1 shows an example of the configuration that was used in the lab setup.

Example 1 Example of the configuration that was used in the lab setup

Defined configuration:

```
cfg:EPIC_FS9500_SCSI
    Server0_64Gb_FS9500_1_0; Server0_64Gb_FS9500_1_1;
    Server0_64Gb_FS9500_1_2; Server0_64Gb_FS9500_1_3

zone:Server0_64Gb_FS9500_1_0
    Server0_64Gb_S1P0; FS9500_1_n1p5; FS9500_1_n2p5
zone:Server0_64Gb_FS9500_1_1
    Server0_64Gb_S1P1; FS9500_1_n1p6; FS9500_1_n2p6
zone:Server0_64Gb_FS9500_1_2
    Server0_64Gb_S2P0; FS9500_1_n1p7; FS9500_1_n2p7
zone:Server0_64Gb_FS9500_1_3
    Server0_64Gb_S2P1; FS9500_1_n1p8; FS9500_1_n2p8

alias:Server0_64Gb_S1P0
    10:00:00:10:9b:ce:62:45
alias:Server0_64Gb_S1P1
    10:00:00:10:9b:ce:62:46
alias:Server0_64Gb_S2P0
    10:00:00:10:9b:ce:65:e1
alias:Server0_64Gb_S2P1
    10:00:00:10:9b:ce:65:e2

alias:FS9500_1_n1p5
    50:05:07:68:13:25:00:d7
alias:FS9500_1_n1p6
    50:05:07:68:13:26:00:d7
alias:FS9500_1_n1p7
    50:05:07:68:13:27:00:d7
alias:FS9500_1_n1p8
    50:05:07:68:13:28:00:d7
alias:FS9500_1_n2p5
    50:05:07:68:13:25:00:d8
alias:FS9500_1_n2p6
    50:05:07:68:13:26:00:d8
alias:FS9500_1_n2p7
    50:05:07:68:13:27:00:d8
alias:FS9500_1_n2p8
    50:05:07:68:13:28:00:d8
```

Effective configuration:

```
zone:Server0_64Gb_FS9500_1_0
    10:00:00:10:9b:ce:62:45
```

```
50:05:07:68:13:25:00:d7
50:05:07:68:13:25:00:d8
zone:Server0_64Gb_FS9500_1_1
10:00:00:10:9b:ce:62:46
50:05:07:68:13:26:00:d7
50:05:07:68:13:26:00:d8
zone:Server0_64Gb_FS9500_1_2
10:00:00:10:9b:ce:65:e1
50:05:07:68:13:27:00:d7
50:05:07:68:13:27:00:d8
zone:Server0_64Gb_FS9500_1_3
10:00:00:10:9b:ce:65:e2
50:05:07:68:13:28:00:d7
50:05:07:68:13:28:00:d8
```

Setting up an IBM Storage Copy Data Management and IBM Storage Sentinel environment to scan Epic databases

This section describes the high-level steps that are needed to design and deploy a Copy Data Management and Sentinel environment to protect and scan your Epic databases:

1. Plan and implement your supported server and storage deployment, which includes predefining Safeguarded Copy volume groups.
2. Plan and implement the security settings and user accounts that are needed for creating Safeguarded Copy copies; integrating with vSphere; logging in to Copy Data Management and Sentinel; and so on. Decide whether you will use local accounts or LDAP or Microsoft Active Directory (AD); what the scope of authority will be for each account; and so on.
3. Plan and implement your Sentinel farm at the scale that is needed to perform the scanning of your Epic databases (including any other workload that you plan to support with the Sentinel farm). Plan how you will distribute your scanning workloads across the farm, and how often you will be able to scan your application servers with this farm.

4. Deploy a new Copy Data Management virtual appliance, if needed. If you use LDAP or AD for authentication, configure the security directory. To configure Copy Data Management to use LDAP or AD accounts, click the Configure tab, and then under the Provider Browser pane, click LDAP to register the LDAP server (Figure 12). Then, import the LDAP group by clicking **Access Control**, add a user; and then click **Import LDAP Group** (Figure 13).

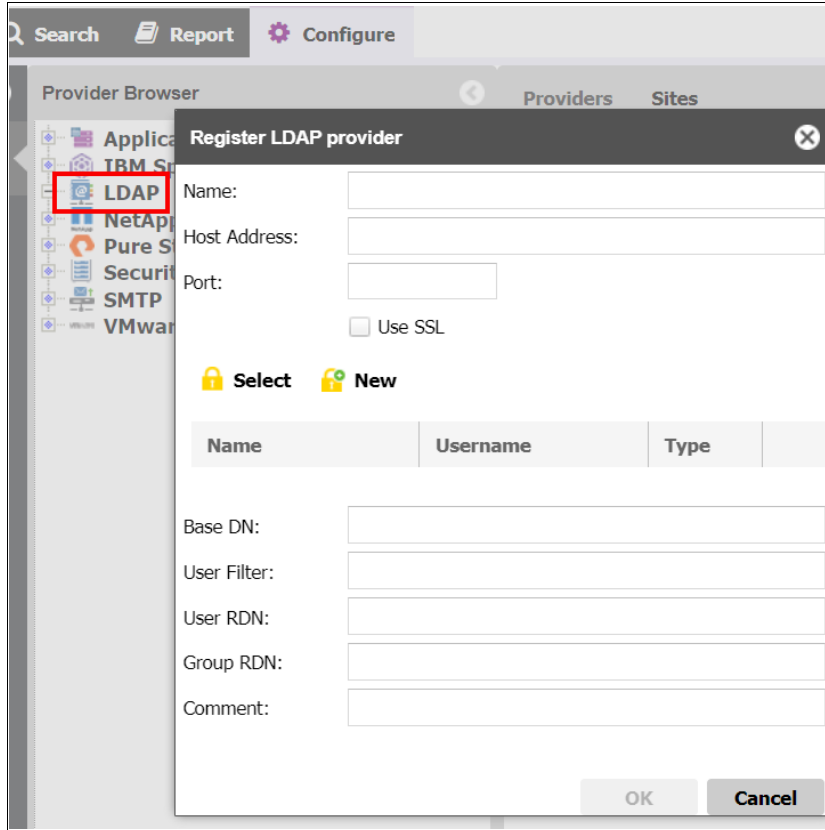


Figure 12 Registering the LDAP server

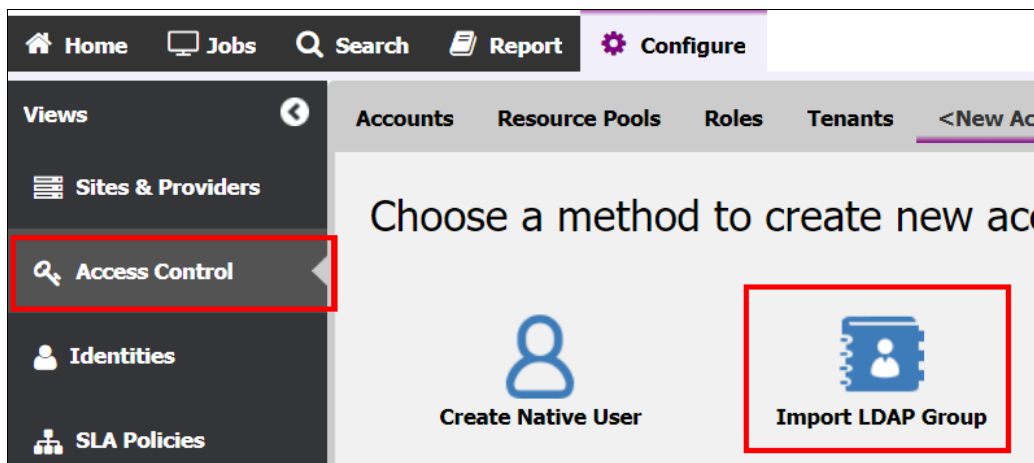


Figure 13 Import LDAP Group pane

5. Define a Copy Data Management site that contains your storage, vSphere, application server, and Sentinel components by clicking **Sites & Providers**, and then selecting **Sites** → **New**, as shown in Figure 14.




Figure 14 Configure Sites

6. Register your storage components, by using the appropriate Storage Virtualize user accounts, which automatically add the storage to the daily Storage Virtualize inventory job and starts an inventory of the newly registered storage. To do this task, click **Sites & Providers**, and then under the Provider Browser pane, select **IBM Spectrum Virtualize** → **Register**. You register the storage under the IBM Spectrum Virtualize node, as shown in Figure 15.

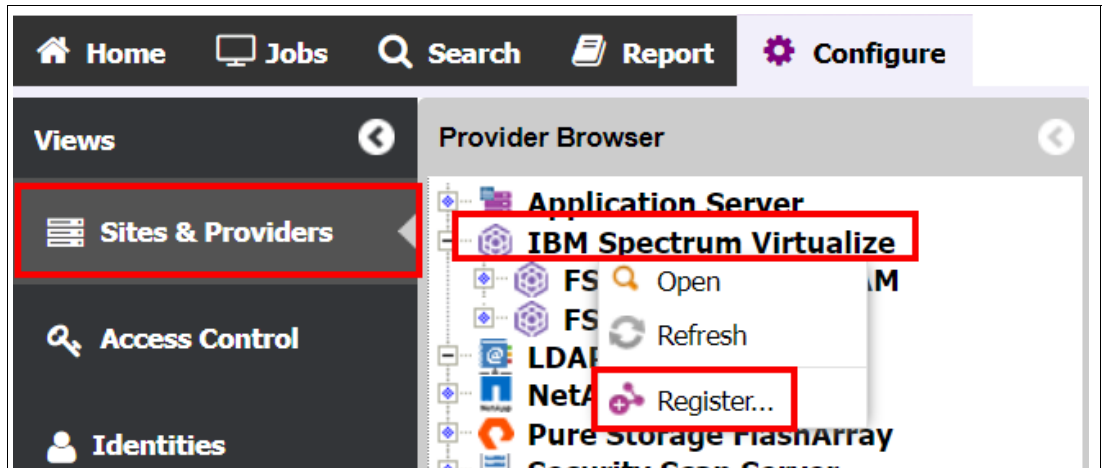


Figure 15 Register your storage components

7. If your application server or Sentinel servers are VMs, register your vCenter servers to Copy Data Management by using the appropriate account or certificate, which automatically adds the vSphere environment to the daily vSphere inventory job and starts an inventory of the newly registered components. Register your vCenter servers by selecting **Sites & Providers**, and then under the Provider Browser pane, select **VMware** → **Register**, as shown in Figure 16.

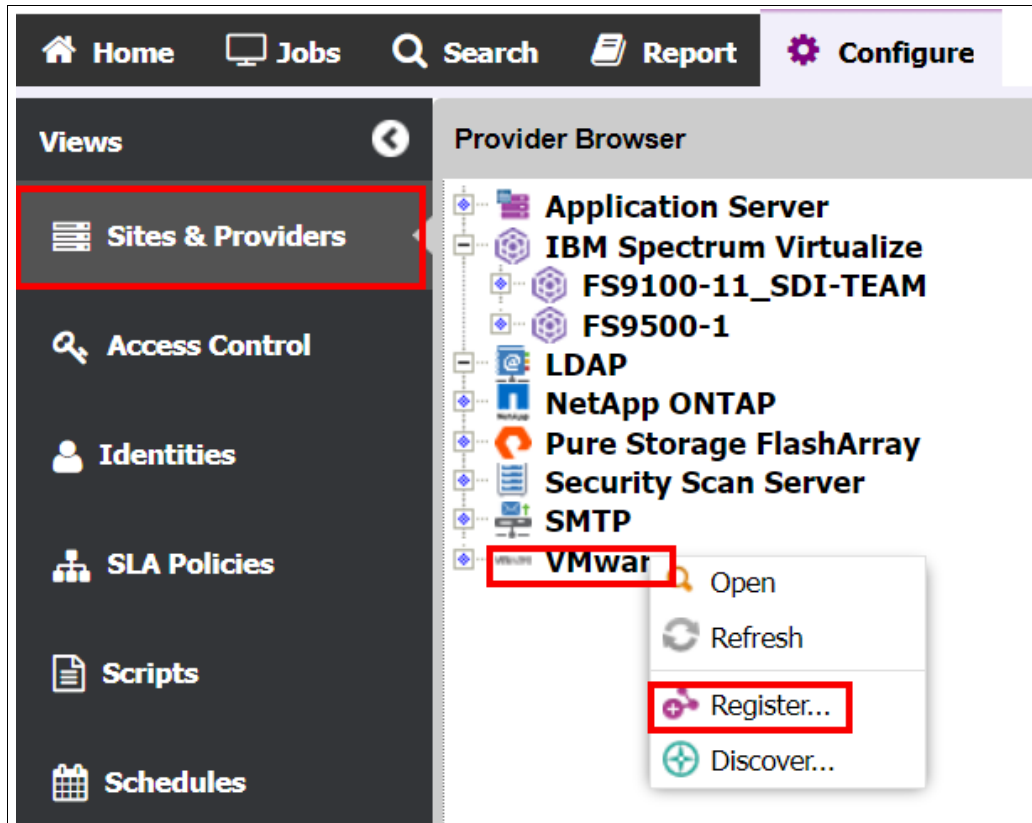


Figure 16 Registering your vCenter servers

8. Register your Epic DB application servers. If your DB servers are VMs, wait until the vCenter inventory completes and you validate that the Epic VMs were found during the inventory by expanding the navigation tree in the **Configuration** tab where you registered the vCenter server. You should see the ESXi hosts and the VMs that are defined to vCenter in that tree, including the application servers that you must register. To register the Epic DB server, click **Sites & Providers**, and then under the Provider Browser pane, select **Application Server** → **Register**. Click **InterSystems Database**, as shown in Figure 17.

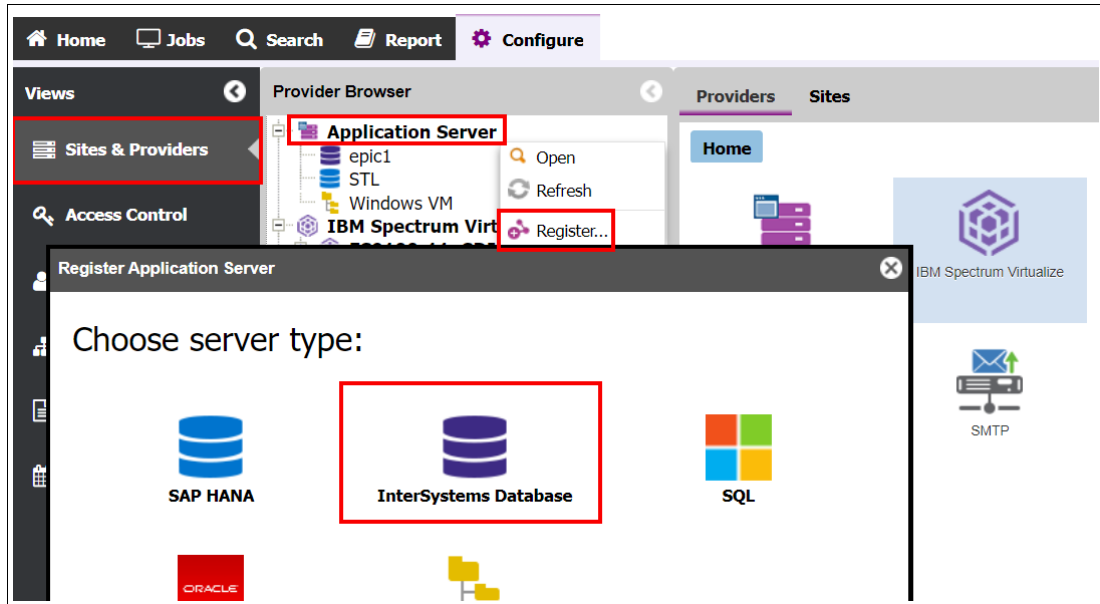


Figure 17 Registering your Epic DB application servers

9. Register your Sentinel servers by using the Security Scanner node in the Configuration tree. If your Sentinel servers are VMs, wait until the vCenter inventory completes and you validate that the Sentinel VMs were found during the inventory by expanding the navigation tree in the **Configuration** tab where you registered the vCenter server. You should see the ESXi hosts and the VMs that are defined to vCenter in that tree, including the Sentinel servers must register. To register the Sentinel server, click **Sites & Providers**, and then under the Provider Browser pane, select **Security Scan Server** → **Register**, as shown in Figure 18.

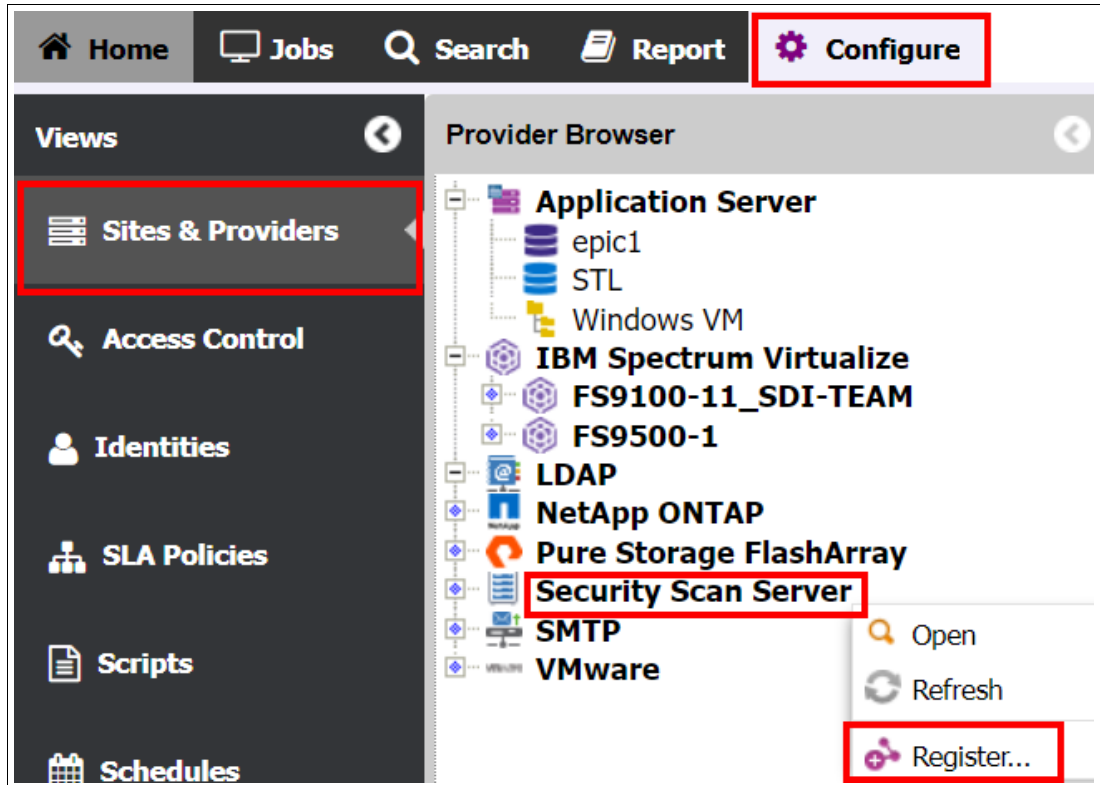


Figure 18 Registering your Sentinel servers

10. Validate each of your components registered correctly. View the results of the Epic inventory job to validate that the Cache or IRIS instances are protected. You see information about the application that is detected and the number of databases that are cataloged. To do this task, select the **Job** tab, expand the application servers, and select **InterSystems Database**. Click the **Default Inventory Job**. Click **History**, and then click the job log's hyperlink to open the job log page.

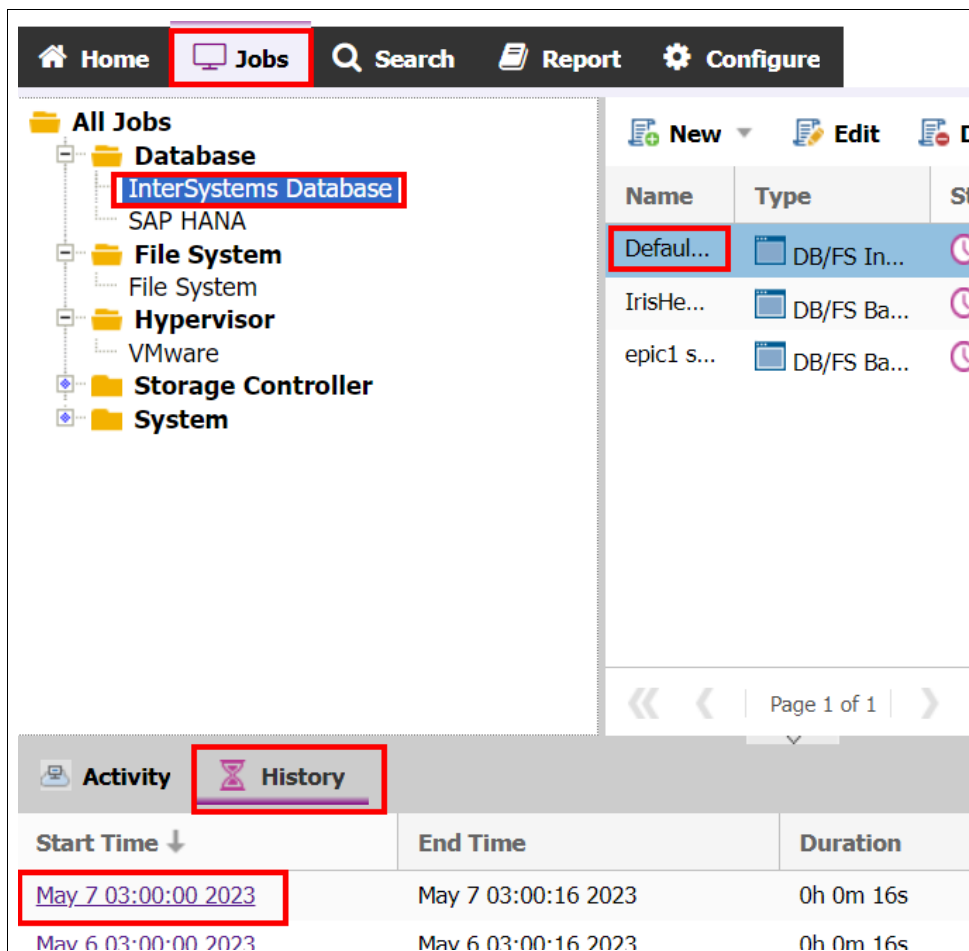


Figure 19 Click the job log hyperlink to open the job log page

Figure 20 shows the job log page.




Figure 20 Job log page

11. Define a service-level agreement (SLA) for your Epic DB data protection by clicking the **Configure** tab, and then clicking **SLA Policies**. Click **New**, and then click **IBM Spectrum Virtualize**. as shown in Figure 21 on page 25.




Figure 21 Define an SLA for your Epic DB data protection

- Give the SLA a unique name (ideally one that identifies the type of protection that the SLA manages). Add a meaningful comment.
- Select the **Source** icon and enter the Frequency and Interval for this SLA.
- Right-click **Source** and select **Add Safeguarded Copy**, as shown in Figure 22.

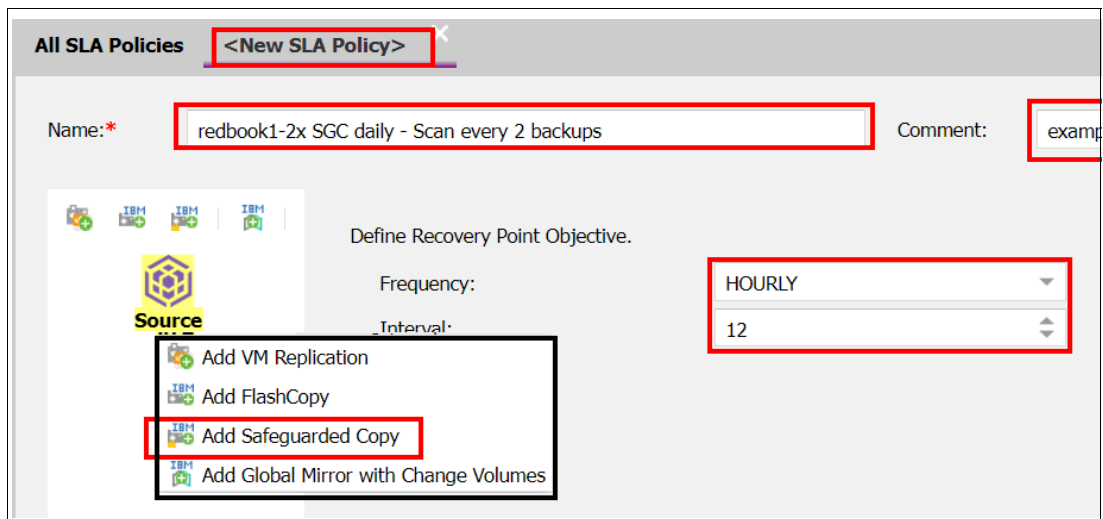


Figure 22 Add Safeguarded Copy

- Select the Safeguarded Copy Volume Group to associate with this SLA.
- Identify the number of days to retain the copies.
- Give this set of Safeguarded Copies a unique and meaningful name.
- Define a volume prefix for the copies. As a best practice, use a unique and meaningful prefix to identify the Copy Data Management instance and SLA name to help correlate copies back to what created those copies.

- h. Select the **Perform Security Scan every:** checkbox to enable security scanning, identify how many backups there should be between scans, and select the Sentinel instance that you previously registered for this workload, as shown in Figure 23.

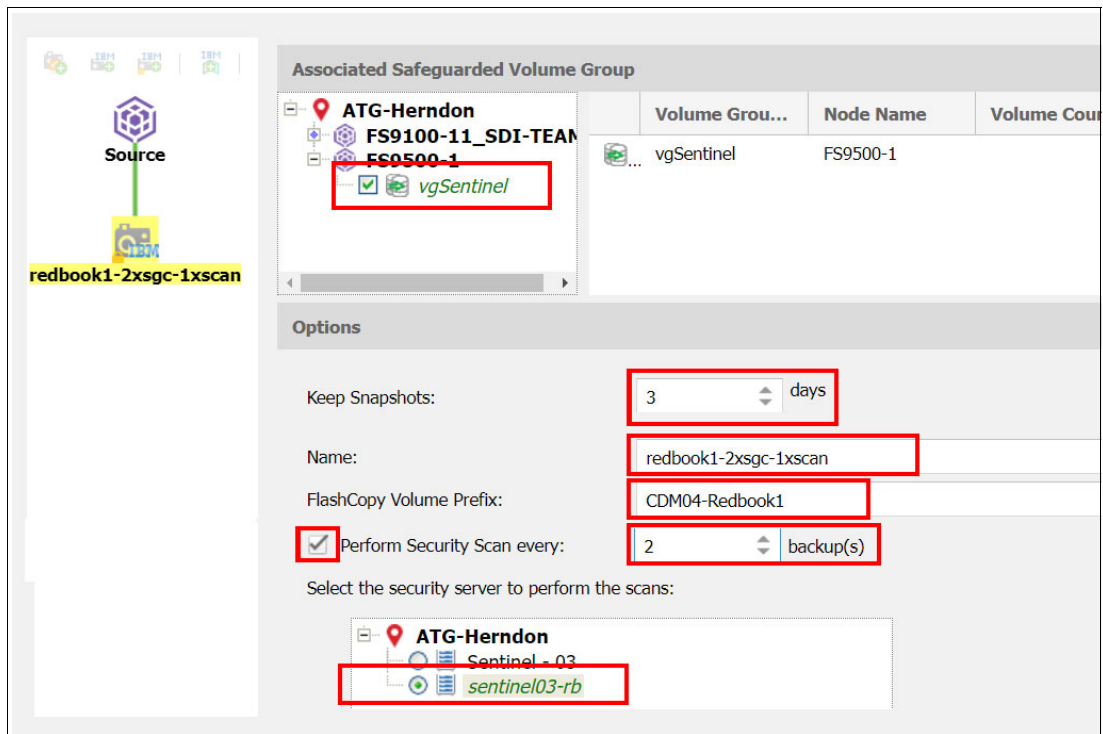


Figure 23 Completing the SLA configuration

- i. Save your SLA.
12. Define a backup job for your Epic application servers by completing the following steps:
- a. Click the **Jobs** tab, expand **Database**, and click **InterSystems Database**.
 - b. Select **New** → **Backup** to add a backup job, as shown in Figure 24.

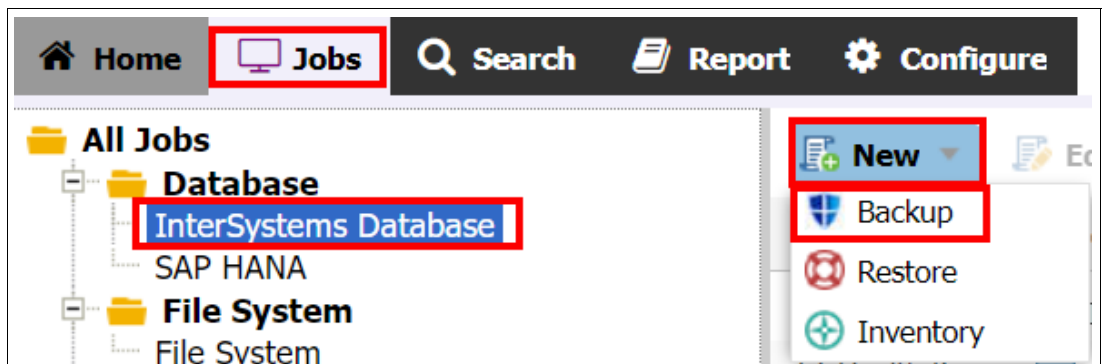


Figure 24 Adding a backup job

- c. Give this job a meaningful name and comment.
- d. Expand the node for the site where your Epic instance is, and expand the Epic server's node. Select the checkbox next to the instance or database.
- e. Select your SLA, as shown in Figure 25 on page 27.

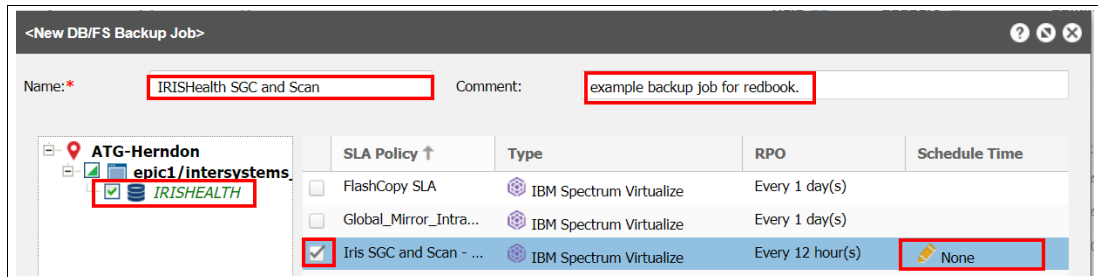


Figure 25 Selecting your SLA

- f. Click the **Schedule Time** column so you can enable a schedule and set a start time. How often the job runs is defined in the SLA. For example, if you defined the SLA to run every 12 hours and schedule the backup to start at midnight, it runs at midnight and at noon.
 - g. If you want to define pre-job, post-job, pre-snapshot, or post-snapshot scripts to run as part of this job, click **Advanced** and identify the scripts that you want to run.
13. Monitor your data protection as shown in Figure 26. You can manually start the backup job at any time by right-clicking the job definition, but you might impact your production Epic instance if you run the job at a time of high activity.

If you run the job manually or wait for its scheduled start time, review the job log for the backup job. You should see activity indicating the following items:

- The instance was quiesced.
- A snapshot was created.
- The instance was unquiesced.
- Whether a security scan should be performed before the job runs.
- Activity showing copies made of the copies and then assigning and mounting those copies to the security scanner.

When the job shows as finished, view the job log in the **History** tab.

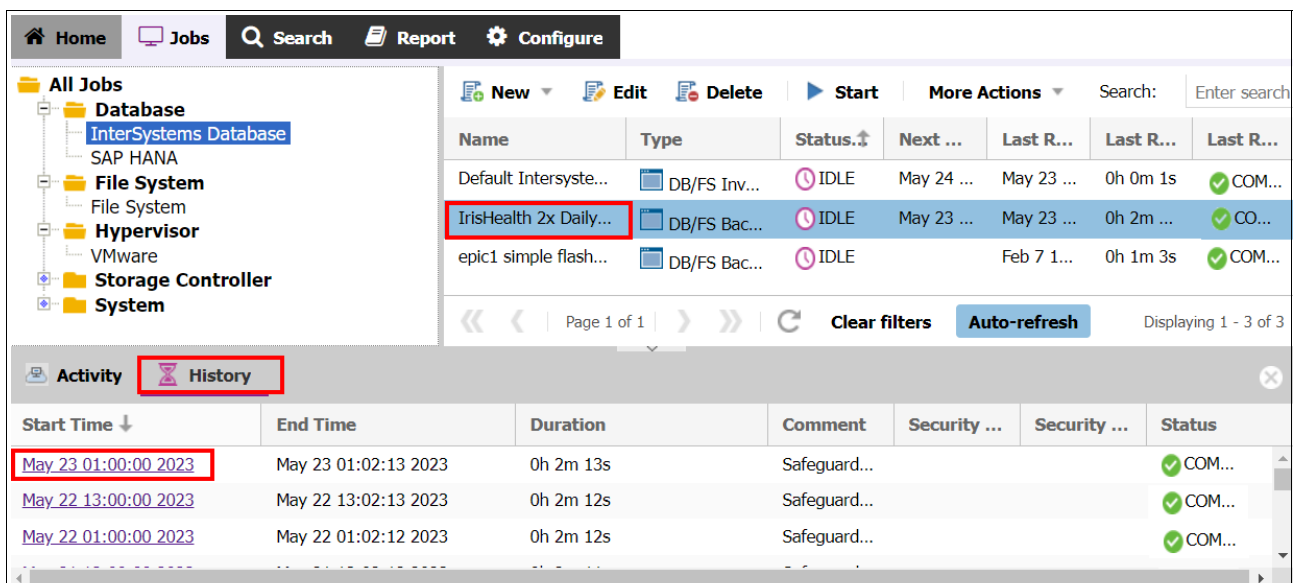


Figure 26 Monitoring your data protection

Summary

For over 15 years, IBM has collaborated with Epic to help healthcare organizations optimize their IT environments.

IBM Flash Storage solutions deliver the performance, scalability, simplicity, and flexibility that is needed to leverage how hospitals and other clinical institutes store, access, secure, manage, share, and analyze data.

With IBM Storage Virtualize for enterprise functions and flash-based NVMe drives for security and performance, IBM Storage FlashSystem is a powerful platform for demanding EHRs.

Authors

This blueprint guide was produced by a team of specialists from around the world.

John A Bernatz is a Senior Storage Technical Specialist with the IBM Advanced Technology Group (ATG) serving as a subject matter expert (SME) for IBM Storage FlashSystem and IBM Storage DS8900F systems. He has a deep systems background that includes x86, IBM Power, and LinuxONE, in addition to designing and implementing storage area networks (SANs). With more than 20 years at IBM, he started in the IBM Global Services Solution Validation Lab as a technician, and he has spent the last 13 years in the ATG as a storage SME.

Dave Gimpl is a Senior Technical Staff Member within the IBM FlashSystem development team as an Integration Architect. David has over 33 years with IBM, spanning file protocol, Fix Central, IBM Power HMC development, cloud computing, and storage resource management software development. He is a Development Lab Advocate for several clients. He provides education about IBM storage products, and leverages contact with those clients to bring feedback directly into product development.

Tim Jeka leads the technical and positioning effort of the Brocade storage networking business with IBM as an IBM North America Systems Engineer of Storage Networking. Tim is responsible for driving technical awareness and positioning of Brocade storage networking products to the IBM field across the Americas. He joined Brocade in 2008, bringing over 34 years of experience in networking, storage, and SAN technology with IBM. Before joining Brocade, Tim was a Networking and Storage Area Networking RDS at IBM. Tim also held Technical Team leadership roles in the IBM Storage and Networking product divisions for over 34 years.

Mark Jones has worked in the enterprise computing industry since 1984. He is the Director of Technical Marketing for Broadcom Corporation, Emulex Connectivity Division. He came to Broadcom through the acquisition of Emulex Corporation in 2015, for which he worked since 2002. He worked as a Strategic Solutions Manager for Burroughs/Unisys for 18 years. He has a Computer Science degree from the University of Redland, and serves as the President Emeritus of the Fibre Channel Industry Association (FCIA).

Matt Key graduated from Texas A&M University in 2006 and took a role at Texas Memory Systems (TMS) to work as an Applications Engineer. In 2012, IBM acquired TMS to provide a strategic flash offering (IBM FlashSystem). After the integration of TMS into IBM, Matt took the role of technical sales lead for flash technology and focused on broadening the message and technical validation of flash into use cases beyond high-performance computing. Matt is part of IBM ATG and focuses on flash who drives success in client education and experience.

Brian Larsen joined Broadcom in July 1991. He has more than 35 years of professional experience in high-end processing, storage, disaster recovery, cloud, virtualization, and networking environments. Larsen is the Director of Partner Business Development, and is responsible for solution and business development within all IBM divisions. In addition to his 14 years in Business Development, Larsen has held positions in sales and consulting, product management, and solutions marketing.

Daniel Thompson has been working in IT for more than 40 years. His specialty is data protection (Backup and Restore, Disaster Recovery, Business Continuity, and Cyber Resiliency). He works at the ATG, IBM Technology, Americas.

Dan Zehnpfennig joined the ATG as a Senior Storage Technical Specialist who concentrates in Health Care ISVs. Dan comes from the IBM Business Partner community, where he was a technical storage seller for more than 20 years. He has been involved with IBM Storage Virtualize products since 2003, and the DS8000 and its precursors. He enjoys helping customers get the most out of their storage.

Resources

For more information, see the following resources:

- ▶ IBM Data Storage for Healthcare:
<https://www.ibm.com/it-infrastructure/storage/healthcare-industry>
- ▶ IBM Epic Alliance Team:
<https://www.epicibm@us.ibm.com>
- ▶ Epic:
<https://www.epic.com/>

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.


COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo) ®
FlashCopy®
IBM®

IBM FlashCore®
IBM FlashSystem®
IBM Research®

IBM Spectrum®
Redbooks®

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.



© Copyright IBM Corporation

July 2023

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.



Please recycle

ISBN 073846127x

REDP-5692-00