



Enabling Centricity Clinical Archive with IBM Spectrum Scale

*Leverage an integrated solution to achieve
actionable insights with seamless image
exchange and better access to clinical
information*

*Prashant Avashia, IBM Systems
Gustavo Hernandez, GE Healthcare
Howard Jiang, IBM Systems*

March 2016



Table of contents

| | |
|-------------------------------------------------------------------------|-----------|
| Abstract | 1 |
| Introduction: Centricity Clinical Archive | 2 |
| Introduction: IBM Spectrum Scale | 4 |
| IBM Spectrum Scale: Configurations, tests, and results | 6 |
| Planning information and recommendations | 11 |
| Networking connectivity guidelines | 14 |
| Enabled solution delivery recommendations | 14 |
| File systems layout recommendations..... | 15 |
| Service checks, commands, and procedures to ensure cluster health | 16 |
| Storage estimation and sizing considerations for a VNA solution | 22 |
| Solution benefits – summary | 23 |
| Acknowledgments | 23 |
| Glossary | 24 |
| Resources | 28 |
| About the authors | 28 |
| Trademarks and special notices | 29 |



Abstract

Patients expect their physicians to provide them a higher-quality healthcare with intelligent, real-time insights from their radiological images, clinician notes, and lab results, with easy digital access to their personal medical information, diagnostic guidance, and treatment options. To deliver improved quality of care, physicians and specialists are now demanding high-quality real-time imaging data from the radiologists. The primary driver for the transformation of radiological imaging services from volume-based imaging to patient-centric value-based imaging has enabled a closer dialogue between the radiologists and the referring physicians.

IBM worked with GE Healthcare, a subsidiary of General Electric Company, to combine Centricity Clinical Archive with the IBM Spectrum Scale software, into a proven solution. This solution offers a standard-based, secure, and robust architecture that delivers a scalable, reliable solution for departmental and enterprise imaging environments.

This paper describes the solution architecture for deploying Centricity Clinical Archive on IBM Spectrum Scale software in an active-active synchronous cluster (split-cluster) configuration. This paper provides recommendations, installation steps, and best practices to help ensure an efficient installation of the joint solution with acceptable performance that meets the needs of the organization.



Introduction: Centricity Clinical Archive

Centricity Clinical Archive is a multiology, multisite, vendor-neutral solution that transforms rigid departmental information silos into fluid, information-sharing, and infrastructure-sharing systems to help clinicians make informed decisions with greater efficiency.

By liberating patient data from departmental silos, **Centricity Clinical Archive** relieves clinicians from navigating through multiple disparate systems, putting patient data directly at the point of care, where it is needed the most.

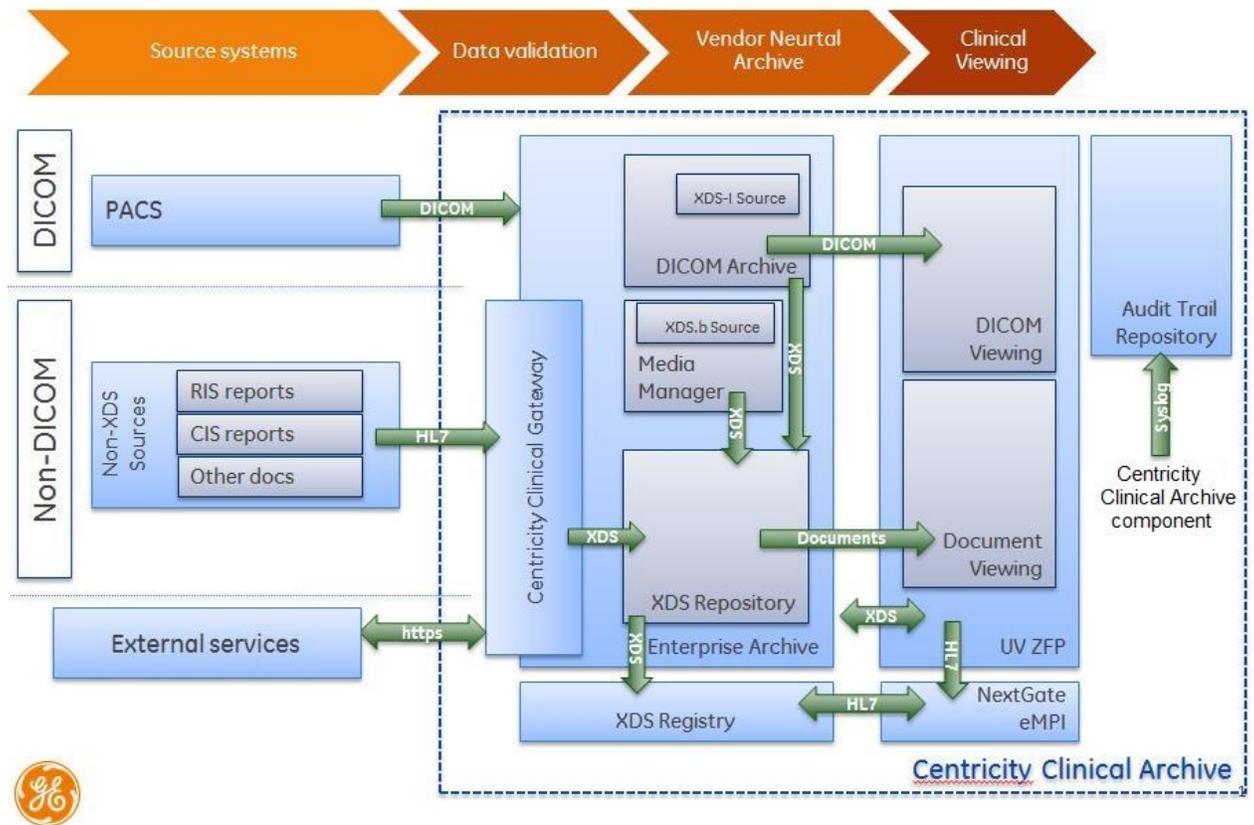


Figure 1: Solution architecture: Centricity Clinical Archive solution

Centricity Clinical Archive provides a powerful vendor-neutral archiving and viewing solution.

Centricity Clinical Archive includes the following product components: Centricity Enterprise Archive, Universal Viewer ZFP, XDS Registry, Centricity Clinical Gateway, NextGate® MatchMetrix® EMPI, and Lexmark PACS Scan™ as described in this section.



Product components

- Centricity Enterprise Archive (EA) provides a multiology and multisite clinical content repository, enabling consolidation of IT infrastructure for archiving and managing unstructured medical content (images, reports, documents and so on) using DICOM and XDS.
- Centricity Universal Viewer Zero Footprint (ZFP) is a zero footprint clinician viewer providing anywhere anytime access (where internet connection is available) to patient clinical records for care providers.
- XDS Registry is an *Integrating the Healthcare Enterprise* (IHE) compliant registry that stores the catalog of patient clinical information.
- NextGate MatchMetrix enterprise master patient index (EMPI) is used to link patient identifiers across multiple clinical information systems and create a commonly known patient identifier for each patient.
- Lexmark PACS Scan is an interface engine to connect nonstandard devices and systems to store data in the content repository (EA) using either DICOM or IHE XDS.
- Centricity Clinical Gateway provides the messaging interface engine which allows clinical information systems such as HIS, RIS, and so on to insert and update medical content in the content repository (EA) and in the EMPI.
- Audit Trail Repository receives audit messages from the other components of the Centricity Clinical Archive solution through IHE-compliant interfaces based on the IHE Audit Trail and Node Authentication (ATNA) profile.
- Media Manager supports the acquisition, classification, and storage of non-DICOM images in Centricity Clinical Archive using mobile devices at the bedside and other patient care locations.
- Patient Consent Manager supports the creation of electronic consent documents, a digital representation of the patient's consent to a specific privacy policy. These consent documents are stored in the XDS Document Repository and can be used to control access to all of the patient's documents known to the XDS document registry.

Service offerings

- Solution consulting – Assess the current status and future needs including number and type of departmental systems, IT infrastructure, and processes to provide recommendations for long-term sustainable data management and sharing.
- Implementation services – Professional services to help interface existing departmental IT systems from various vendors with the solution, ensuring regular consolidation of information (DICOM and non-DICOM).
- Data migration services – A one-time service to migrate and link historical data from departmental systems.



Introduction: IBM Spectrum Scale

This section highlights the IBM® Spectrum Scale™

IBM Spectrum Scale is a software-defined solution for managing structured and unstructured data with security, reliability, and high performance. It provides system scalability, very high availability, and reliability with no single point of failure in large-scale storage infrastructures.

The IBM Elastic Storage Server is a modern implementation of a software-defined storage system, combining IBM Spectrum Scale™ software with IBM POWER8™ servers and storage enclosures. Unique to IBM Elastic Storage Server is IBM Spectrum Scale RAID software that uses erasure coding instead of traditional Redundant Array of Independent Disks (RAID) for data protection. Because Spectrum Scale is all implemented in software, ESS has all of the storage functionality implemented in software and can use standard servers and storage enclosures as opposed to proprietary storage controllers making ESS a true implementation of software defined storage. The IBM Elastic Storage Server is IBM's building block solution for high performance and shared storage for modern workloads for unstructured data.

IBM Spectrum Storage based solutions deliver the potential of extracting insights from data volumes, and increase business agility by offering its functional capabilities as software, or as a cloud, or as a managed service offering. These all, are complete, ready-to-deploy, proven, high-performance solutions that help to accelerate the time-to-value with reliability, security, and speed.

Part of the IBM Spectrum Storage™ family, Spectrum Scale 4.1.1 introduces new revolutionary features that demonstrate IBM's commitment in providing groundbreaking storage solutions, including:

- File encryption and secure erase
- Transparent flash cache
- Network performance monitoring
- Active File Management (AFM) parallel data transfers
- Network File System (NFS) version 4 support and data migration
- Backup and restore improvements
- File Placement Optimizer (FPO) enhancements

IBM Spectrum Scale allows a group of computers concurrent access to a common set of file data over a storage area network (SAN) infrastructure, a TCP/IP network, InfiniBand® or a combination of other network connection types. The computers can run a mix of IBM AIX®, Linux®, or Microsoft® Windows® operating systems. Spectrum Scale provides storage management, information lifecycle management tools, centralized administration, and shared access to file systems from remote Spectrum Storage clusters providing a global namespace.

A Spectrum Scale cluster can be a single node in a tiered storage solution, two nodes providing a high-availability platform supporting a database application, or thousands of nodes used for applications such as weather-pattern modeling.

A Spectrum Scale file system is built from a collection of storage devices that contain the file system data and metadata. A file system can be built from a single disk or contain thousands of disks storing



petabytes of data. Each file system can be accessible from all nodes within the cluster. There is no practical limit on the size of a file system. The architectural limit for a single file system is more than a yottabyte. Some Spectrum Scale customers use single file systems that are up to 18 PB in size, while others use file systems containing billions of files. Applications access files through standard Portable Operating System Interface (POSIX) file system interfaces. Because all nodes see all file data, any node in the cluster can concurrently read or update a common set of files—enabling applications to scale out easily. Spectrum Scale maintains the coherency and consistency of the file system using sophisticated byte-range locking, token (distributed lock) management, and journaling. This approach means that applications using standard POSIX locking semantics do not need to be modified to run successfully on Spectrum Scale.

Spectrum Scale provides unparalleled I/O performance for unstructured data by:

- Striping data across multiple disks attached to multiple nodes
- Employing high-performance metadata (inode) scans
- Supporting a wide range of file system block sizes to match I/O requirements
- Utilizing advanced algorithms to improve read-ahead and write-behind I/O operations
- Using byte-range locking, based on a sophisticated scalable token management system to provide data consistency while allowing concurrent access of multiple application nodes to files

Spectrum Scale supports POSIX and NFS v4 access control lists (ACLs). NFS v4 ACLs can be used to serve files using NFS v4, but can also be used in other deployments—for example, to provide ACL support to nodes running Windows. To provide concurrent access from multiple operating system types, Spectrum Scale allows you to run mixed POSIX and NFS v4 permissions in a single file system and map user and group IDs between Windows and Linux / UNIX® environments.



IBM Spectrum Scale: Configurations, tests, and results

Configurations and tests

Refer to Figure 2. Through this configuration and test, an IBM Spectrum Scale software set up or installed on IBM Elastic Storage Servers was successfully enabled with Centricity Clinical Archive. The following list of GE Healthcare servers were configured as virtual machines (VMs), for the purposes of this test, as outlined in the following table.

| No. | Application | Software version | Processor (cores) | Memory (Gb) |
|-----|------------------------------------|------------------|-------------------|-------------|
| 1 | Centricity Enterprise Archive (EA) | 4.0.10 | 8 | 24 |
| 2 | XDS Registry | 3.17.1 | 4 | 8 |
| 3 | EMPI | 9.0.1_rev15677 | 4 | 4 |
| 4 | Audit Trail Repository | 1.0 | 4 | 8 |
| 5 | Media Manager | 6.0 SP1 | 4 | 8 |
| 6 | Centricity Universal Viewer ZFP | 6.0 SP1 | 8 | 16 |

Table 1: GEHC applications, configured as virtual machines configured to work with the storage

The VMs were configured with Microsoft Windows 2012 software pre-packaged in self-installable ready to deploy VMWare OVF template formats by GE Healthcare Corporation. In addition, all VMs had SQL Server 2012 Standard Edition installed to host the databases of each of the GE Healthcare products (as embedded software license).

The ippon and nihon systems were configured to be connected to each other as an active-active stretch cluster through a synchronous replication topology. The required file systems residing on the synchronous cluster, were made available to the GE Healthcare virtual machines (VMs) through the protocol nodes configured with Cluster Export Services (CES), and using the Microsoft Server Message Block SMB2 and SMB3 protocols, over the network, as shown in the following figure. In terms of network security, Active Directory (AD) authentication security services were configured in the network.

In a configuration utilizing synchronous replication, a single GPFS cluster is defined over three geographically-separate sites. These sites consist of two production sites (Site A and Site B) and a third tiebreaker site (Site C). The two production sites are connected over 10 GbE links with synchronous replication capabilities. For the purposes of this test, Site A was confined to network 9.11.82.XX, Site B was confined to network 9.11.103.XX and Site C was confined to 9.11.206.XX, as highlighted in Figure 2, of this paper.



The concept of global namespace allows multiple servers that are physically located in different networks or three different sites to share concurrent access to a common set of files and data resident in the GPFS cluster. This concept can be extended to facilitate the ability to share data across multiple GPFS clusters. This allowed customers to share data between organizations in the same building or across different geographic locations.

One or more file systems are created, mounted, and accessed concurrently from the two active production sites (Site A and Site B).

The data and metadata replication features of IBM Spectrum Scale are used to maintain a secondary copy of each file system block, relying on the concept of disk failure groups to control the physical placement of the individual copies:

1. Separate the set of available disk volumes into two failure groups. Define one failure group at each of the active production sites.
2. Create a replicated file system. Specify a replication factor of 2 for both data and metadata.

When allocating new file system blocks, GPFS always assigns replicas of the same block to distinct failure groups. This provides a sufficient level of redundancy allowing each site to continue operating independently when the other site fail.

IBM Spectrum Scale enables the sharing of data across unreliable or high latency networks. This feature is called active file management (AFM). AFM allows you to create associations between GPFS clusters and define the location and flow of file data and helps to automate the management of the data. AFM allows you to now implement a single namespace view across sites, around the world once again, redefining the scope of the term *global namespace*.

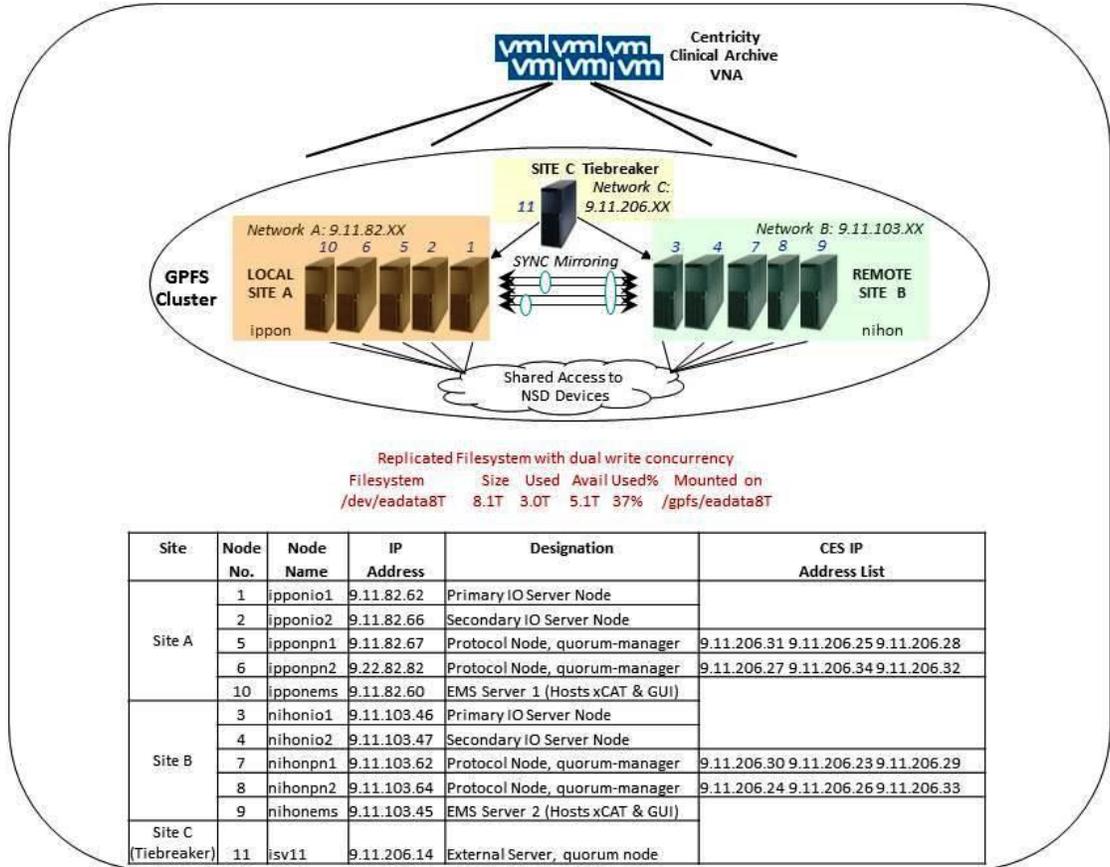


Figure 2: IBM GPFS stretch cluster servicing GE clinical applications through SMB2/3 protocols over a sync-replication environment

Test observation and results:

- The GE Healthcare Centricity Enterprise Archive software was successfully installed, configured, and tested, with IBM Spectrum Scale software, over Common Internet File System (CIFS) protocols, with successful results.
- Tests included sending and receiving medical images through DICOM and XDS capabilities, using Windows SMB 2/3 protocols.
- No critical issues or showstoppers were reported with installation, configuration, and tests, conducted as mentioned.



Test results

Centricity Enterprise Archive can successfully read and write medical images from and to the IBM ESS synchronous cluster preconfigured with IBM Spectrum Scale software. The following table highlights the tests jointly conducted by IBM and GE and the successful results.

| No. | Test | Result |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|
| 1 | Configure storage: Configure local and remote storage locations. | Configuration only. No result at this step. |
| 2 | <p>Store anonymized studies with different modalities (CT, CR, MR) simultaneously for a period of 1 hour.</p> <p>Confirm that all images are successfully archived. For this purpose, a Hossim (hospital simulation system) script has been designed by GE Healthcare. It tries to store about 5000 studies in an hour with a throughput of about 500 GB per hour.</p> <p>This is a stress test to determine the maximum throughput that EA can sustain on the IBM test lab.</p> <p>During this test, the following information will be collected:</p> <ul style="list-style-type: none"> • Hossim log with the achieved throughput in GBph and the number of studies • - Performance monitor report for processor, memory, and I/O from the virtual machine instance running EA. | PASS |
| 3 | <p>Store anonymized studies with different modalities (CT, CR, MR) simultaneously for a period of 3 hours.</p> <p>Confirm that all images are successfully archived. For this purpose, a Hossim script has been designed to subject EA to a constant load for a relatively prolonged period of time.</p> <p>During this test the following information will be collected:</p> <ul style="list-style-type: none"> • Hossim log with the achieved throughput in GBph and number of studies • Performance monitor report for processor, memory, and I/O from the box running EA. | PASS |

| | | |
|---|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| 4 | <p>Clinical viewing: View 10 studies on the DICOM viewer. Determine:</p> <ul style="list-style-type: none"> • First image display (how many seconds until the first image appears on the viewer)- • that all chosen studies are shown correctly, that is, all series and images are streamed correctly from EA to the viewer | PASS |
| 5 | <p>Store a large study Capture performance counters: Memory, processor, network I/O</p> | PASS |
| 6 | <p>Route large study Capture performance counters: Memory, processor, network I/O</p> | PASS |
| 7 | <p>Create an information lifecycle management (ILM) DICOM move rule to migrate one entire library to another on the same archive. Retrieve an existing study to check that the move has been done correctly and images are accessible. Measure the time taken to move the contents of an entire library to another. View the exam on the Universal Viewer ZFP patient timeline and high-fidelity DICOM viewer Capture performance counters: Memory, processor, network I/O</p> | PASS |
| 8 | <p>Register two patients on the XDS Registry and MPI. Store DICOM studies and XDS documents for each patient using Hossim and Media Manager. Perform the following steps:</p> <ol style="list-style-type: none"> 1. Register the patients in XDS Registry and MPI using the 7Edit tool (available at the URL: www.7edit.com). Send ADT messages to both systems with the patient's demographics. 2. Make sure that an XDS submission rule has been defined and enabled in EA to provide and register DICOM studies on the XDS infrastructure. 3. Store studies for each of the patients. 4. Using the Media Manager Client Functional Check application, provide and register images for both patients. | PASS |
| 9 | <p>Open Universal Viewer ZFP for one of the patients and view the patient timeline. Determine the time it takes to display the patient timeline. View one of the studies using the DICOM viewer measuring the time it takes to do first image display.</p> | PASS |

Table 2: IBM and GE Healthcare tests run jointly, with successful results



Planning information and recommendations

This section provides specific instructions, planning information, and useful documentation that can help in planning and successfully installing Centricity Clinical Archive with IBM ESS synchronous cluster preconfigured with IBM Spectrum Scale.

Recommendations with IBM Spectrum Scale software:

The following information can help in successful planning and implementing **IBM Spectrum Scale** software in your environment. This location offers excellent recommendations on how to install, administer, maintain, and use IBM Spectrum Scale.

IBM Spectrum Scale Knowledge Center:

- IBM Spectrum Scale Knowledge Center (at: ibm.com/support/knowledgecenter/STXKQY_4.1.1/ibmspectrumscale411_welcome.html?cp=STXKQY_4.1.1%2F0) contains the information required to set up the **IBM Spectrum Scale software** and manage the system using its web-based management GUI and its command-line interface (CLI).
- At this time, the latest version of **IBM Spectrum Scale** version is marked as 4.1.1 in the Knowledge Center and is dated September 21, 2015.
- When you begin your planning, make sure that you are using the most recent documentation, and the latest version of the software.

IBM Spectrum Scale product documentation and related publications:

For IBM Spectrum Scale product documentation and related publications, refer to: ibm.com/support/knowledgecenter/STXKQY_4.1.1/ibmspectrumscale411_welcome.html

In this reference:

- Make sure to review the applicable publications and documents identified in IBM Spectrum Scale V4.1.1 product documentation.
- Review the IBM Spectrum Scale FAQs.
- Pay particular attention to understand the product. There are a number of learning modules and tutorials.

Additional references

1. Elastic Storage Server (ESS), see IBM Knowledge Center: ibm.com/support/knowledgecenter/SSYSP8/sts_welcome.html
2. IBM General Parallel File System (IBM GPFS™), see IBM Knowledge Center: ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html
3. IBM POWER8™ processor-based servers, see IBM Knowledge Center: ibm.com/support/knowledgecenter/POWER8/p8hdx/POWER8welcome.htm



4. The DCS3700 storage enclosure, see:
 - a. System Storage DCS3700 Quick Start Guide, GA32-0960-03:
<http://www-01.ibm.com/support/docview.wss?uid=ssg1S7005178&aid=1>
 - b. IBM System Storage DCS3700 Storage Subsystem and DCS3700 Storage Subsystem with Performance Module Controllers: Installation, User's, and Maintenance Guide, GA32-0959-07:
[ibm.com/support/docview.wss?uid=ssg1S7004920](http://www.ibm.com/support/docview.wss?uid=ssg1S7004920)
5. The IBM Power Systems™ EXP24S I/O Drawer (FC 5887), see IBM Knowledge Center:
[ibm.com/support/knowledgecenter/8247-22L/p8ham/p8ham_5887_kickoff.htm](http://www.ibm.com/support/knowledgecenter/8247-22L/p8ham/p8ham_5887_kickoff.htm)
6. Configuring Protocols Quick Overview for Spectrum Scale, see:
<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20%28GPFS%29/page/Protocols%20Quick%20Overview%20for%20IBM%20Spectrum%20Scale>
7. Extreme cluster/cloud administration toolkit (xCAT), go to the xCAT website:
http://sourceforge.net/p/xcat/wiki/Main_Page/
8. Validate that you have all required planning information and perform a Technical Delivery Assessment (TDA) installation of your Elastic Storage Server or IBM Spectrum Scale or both. The Solution Assurance Pre-sale and Pre-install TDA checklists are to be used by IBM and IBM Business Partner teams conducting required TDAs on IBM Elastic Storage Server opportunities. The purpose of these checklists is to help account teams and Business Partners ensure that they are proposing a supported configuration that meets customer requirements and expectations and ensures that customer environments are properly prepared for deployment and installation. Complete the checklists as this can help to validate the planned setup and environment. You can find the documentation and checklists at the following URL:
[ibm.com/support/assure/assur30i.nsf/WebIndex/SA1053](http://www.ibm.com/support/assure/assur30i.nsf/WebIndex/SA1053)
9. Ensure that this information is validated by IBM SMEs and architects before starting the installation.
10. Perform hardware installation planning: Validate required cables, connectors, network cards, networking connectivity, location, layout, floor space, rack layout, power, cooling, and so on.
11. Define and validate the networking environment, switches, and communication services: IP addresses, existing Network Time Protocol (NTP) servers and IP addresses, existing Domain Name System (DNS) servers and IP addresses, domains, authentication servers, net group ID and mapping support, backup servers, method chosen, antivirus scan engines and so on. These servers should be available and pre-exist in the common network.
12. Plan for system implementation:
 - Define the local and remote Ethernet networking requirements for connectivity, security, DNS services, and sub-netting considerations.
 - Define the network requirements for management and data access
 - Define the network interfaces of I/O nodes, protocol nodes, event management services (EMS) servers, xCAT servers, Hardware Management Console (HMC) systems, including subnets, routing rules and virtual LANs (VLANs)
 - Define the logical configuration of the system.
 - Define the pools, exports / shares, file systems, file sets, and directory structures.

13. Define the users required for the management and monitoring roles of the Spectrum Scale itself and for file-based access requiring authentication. Configure them within the authentication service / directory server.
14. Define the authorizations required for every user accessing every file, file system, file set, and directory structures.

Checkpoints for security authentication

- Decide which implementation of security authentication service will be used.
- The Spectrum Scale 4.1.1 supports only *one* authentication method at a time and changing it later is *not recommended*. Therefore, it is important to carefully decide and select the method up-front.

Implementation

The installation, implementation, and configuration tasks are grouped into major steps of the installation process. Each step is performed sequentially and in most cases, must be completed before the next step can begin.

- Follow the implementation task checklist.
- Unpack, rack, and cable hardware.
- Conduct power-on self-test (POST) and self-check.
- Validate license keys, licenses, and install the latest software.
- Initialize the system.
- Perform base configuration.
- Perform manual setup and configuration changes.
- Plan and configure networking.
- Set up alerting and call-home capabilities.
- Set up directory services and security authentication services.
- Run health check of system.
- Define user security functions and access privileges.
- Configure storage controller and logical disks.
- Configure antivirus capabilities.
- Configure and use file services.

Tips from prior implementations: Some tasks might not apply to a system or environment, and it would be beneficial to have tasks corroborated and validated with IBM teams, prior to the start of any implementation project:

1. Make sure to download and use the correct version of the Spectrum Scale 4.1.1 software configuration guide before you start your Spectrum Scale 4.1.1 planning and installation efforts.
2. Save all your TDA documents and the completed software configuration guide for easy reference at a later date.
3. Save a physical copy of these completed documents in the rack / cabinet of the Spectrum Scale installed system.
4. It is absolutely necessary to have a completed, verified, and validated TDA worksheet before starting any Spectrum Scale installation. More eyes during the TDA and planning processes are



definitely helpful, and this effort catches more errors in the planning stage before the actual installation. This diligence prevents wrong assumptions, incorrect planning, configuration errors, and other issues.

5. IBM Spectrum Scale is a complicated system with its built-in intelligence. Any pre-existing inefficiencies or errors in a client network need to be corrected, and addressed prior to its use.

Training services: In-house instructors from IBM authorized partners such as Arrow, Avnet, Ingram Micro, LearnQuest, and Global Knowledge companies periodically deliver relevant education workshops worldwide for understanding, recommending, sizing, and implementing IBM Spectrum Scale based solutions. These workshops are typically four days long. Specific information on workshop details, timings, locations, and more is available at the following URL:

ibm.com/services/learning/ites.wss/zz/en?pageType=course_description&courseCode=H005G&cc=us

The following workshop for the IBM Spectrum Scale systems is recommended.

- H0050G: IBM Spectrum Scale 4.1 Basic Administration for Linux (Updated)

IBM Spectrum Scale and ESS shipments

Make sure that all required cables, cards, disks, and components are available and have been validated, cross checked, and inventoried with the packing lists enclosed with the shipment.

Networking connectivity guidelines

It is recommended to run all Centricity Clinical Archive application servers (or VMs) or the physical hosts to be set up with MTU=1500. Also, make sure to configure the virtual switch (vSwitch) ports to MTU=1500, if configuring with VMWare.

The networking services are unconstrained to the physical medium over which TCP/IP message traffic is carried. Various network interfaces are supported, including but not limited to: 10Base-T Ethernet, 100Base-T Ethernet, 1000Base-T Ethernet, and asynchronous transfer mode (ATM) OC-3.

The physical media supported depends on network cabling and interfaces equipment available at the Centricity picture archiving and communication system (PACS) installation site and interface equipment commercially available.

An equipment list and configuration information for the physical media supported is available on request.

Enabled solution delivery recommendations

To ensure a valid solution, this section offers the following recommendations:

- File systems layout recommendations
- Service checks, commands, and procedures to ensure cluster health.
- Recommended IBM storage sizing guidance for a PACS solution.

File systems layout recommendations

In strict terms of storage, a distinction is made between short-term online storage and long-term archive storage as follows:

- Short-term storage (STS): New images sent from the modality to the Centricity Enterprise Archive server (version 4.0.10) are registered in the Microsoft SQL database and then stored in the Spectrum Scale 4.1.1 file system, directly.
- The Centricity Enterprise Archive server communicates with the Spectrum Scale 4.1.1 file system, using the jCIFS (Samba version of java enabled CIFS, equivalent to the CIFS protocol from Microsoft), and then writes the images into the CIFS share.
- The architecture does not allow for any caching of data or images, locally on the Enterprise Archive VM.
- Long-term storage (LTS): Images are copied from the short-term storage to the long-term storage based on specific rules. This step is the actual archiving. The long-term storage provides a very large storage capacity. Images can be archived there for many years.

IBM Spectrum Scale 4.1.1 systems fulfill the following requirements of electronic archiving systems.

- Long-term storage of the images
- Access to short-term storage with CIFS and SMB3 protocols
- Access to long-term storage with CIFS and SMB3 protocols
- Quick access to all archived images (independent of the time of archiving)

For improved performance in a normal and a typical production environment, lay out the file systems for Centricity Enterprise Archive (version 4.0.10) as per the following guidelines and recommendations:

- Recommend using the block size of 256 KB for both, short-term storage and long-term storage. With a 256 KB block size file system, it should also be noted, very clearly, that it also creates smaller sub-blocks below the file system that are 1/32 of the block size of the file system. This means that it creates 8 KB sub-blocks, and there are 32 sub-blocks underneath a single 256 KB file system block. So, if you have small 8 KB files, it would write 32 files into a single 256 KB file system block. Contrary to a popular misconception, it does not write a single 8 KB block, and jumps to the next file system block. The default 256 KB block size supports many general file services, and most applications, and more. During many years of testing, it has been found that very few workloads benefit from a file system block size of 8 KB, 16 KB, or 64 KB, specifically.



Service checks, commands, and procedures to ensure cluster health

This section provides an introduction on various execution steps to ensure a correct setup of a synchronous cluster and maintain good cluster health. Proper commands with correct syntax, and relevant options are documented here, at each step. See man pages for a brief description of all commands and relevant options.

1. **Cluster creation:** Run the `gssgenc1uster` command on the EMS to create the cluster. This command creates a GPFS cluster using all the nodes in the node group if the `-G` option is used. The command assigns server licenses to each I/O server node. So, it prompts for license acceptance (or use the `-accept-license` flag). It applies the best practice GPFS configuration parameters for a GNR-based NSD server. At the end of cluster creation, the SAS HBA firmware, storage enclosure firmware and drive firmware are upgraded, if needed. The firmware update can be skipped using the `--no-fw-update` option. This command can take a long time to run.
2. **Verify that the GPFS cluster is active:** Run the `mmgetstate` command from an I/O server node to verify that the cluster is active and operational. Use the `-a` option to include all the nodes in the GPFS cluster. For example:

```
[root@ipponio1 ~]# mmgetstate -L -a -s
```

| Node number | Node name | Quorum | Nodes up | Total nodes | GPFS state | Remarks |
|-------------|-------------|--------|----------|-------------|------------|-------------|
| 1 | ipponio1-hs | 3 | 5 | 13 | active | |
| 2 | ipponio2-hs | 3 | 5 | 13 | active | |
| 3 | nihonio1-hs | 3 | 5 | 13 | active | |
| 4 | nihonio2-hs | 3 | 5 | 13 | active | |
| 5 | ipponpn1-hs | 3 | 5 | 13 | active | quorum node |
| 6 | ipponpn2-hs | 3 | 5 | 13 | active | quorum node |
| 7 | nihonpn1-hs | 3 | 5 | 13 | active | quorum node |
| 8 | nihonpn2-hs | 3 | 5 | 13 | active | quorum node |
| 11 | nihonems-hs | 3 | 5 | 13 | active | |
| 12 | ipponems-hs | 3 | 5 | 13 | active | |
| 13 | isv11-hs | 3 | 5 | 13 | active | quorum node |
| 14 | isv22-hs | 3 | 5 | 13 | active | |
| 15 | isv28-hs | 3 | 5 | 13 | active | |

Summary information

```
-----
Number of nodes defined in the cluster:      13
Number of local nodes active in the cluster: 13
Number of remote nodes joined in this cluster: 0
Number of quorum nodes defined in the cluster: 5
Number of quorum nodes active in the cluster: 5
Quorum = 3, Quorum achieved
```



3. **Create the recovery groups:** The `gssgenclusterrgs` command creates the recovery groups and the associated vdisks for logtip, logbackup and loghome. For data and metadata vdisk, NSD, and file system creation, use the `gssgenvdisks` (preferred method) command.
4. **Verify recovery group creation:** To look at the detail for one of the recovery groups, log into one of the I/O server nodes and run the `mmlsrecoverygroup` command, as documented in the IBM Publication Center at URL:

https://www01.ibm.com/support/knowledgecenter/SSFKCN_4.1.0/com.ibm.cluster.gpfs.v4r1.gpfs100.doc/bl1adm_mmlsrecoverygroup.htm

As example:

```
[root@ipponio1 ~]# mmlsrecoverygroup
```

| Recovery group | VDisks | Declustered | |
|----------------|--------|-------------|------------------------------------------------------------------|
| | | Arrays with | Servers |
| rg_nihonio1_hs | 5 | 39 | nihonio1-hs.tuc.stglabs.ibm.com, nihonio2-hs.tuc.stglabs.ibm.com |
| rg_nihonio2_hs | 5 | 39 | nihonio2-hs.tuc.stglabs.ibm.com, nihonio1-hs.tuc.stglabs.ibm.com |
| rg_ipponio1-hs | 4 | 31 | ipponio1-hs.tuc.stglabs.ibm.com, ipponio2-hs.tuc.stglabs.ibm.com |
| rg_ipponio2-hs | 4 | 31 | ipponio2-hs.tuc.stglabs.ibm.com, ipponio1-hs.tuc.stglabs.ibm.com |

5. **Create VDisks, NSD, and file systems:** The command `gssgenvdisks` can be used to create VDisks, NSDs, and filesystem on existing recovery groups. If no recovery groups are provided, all available recovery groups are used. If the command is run on the EMS node (or any other node) that is not part of the cluster, a contact node which is part of cluster must be provided. The contact node must be reachable from the node (for example, EMS node) where the command is run. This command also allows to add suffix to the VDisk names which can be useful when creating multiple file systems. A unique suffix can be used with the VDisk names to associate them with a different file system (we will see some example shortly). Default reserve capacity is set to 1%. If the VDisk data block size is less than 8 MB, the reserve capacity should be increased to 2%. For example, the following command is used to create an 8 TB file system for testing Centricity Clinical Archive:

```
gssgenvdisks --vdisk-suffix _eadt8 --create-vdisk --create-filesystem --
filesystem-name eadata8T --data-vdisk-size 750 --metadata-blocksize 512k --
data-blocksize 512k --filesystem-options "--version 4.1.0.8"
```

6. The command in step 5 creates the VDisks, NSDs, and the file system by using all available declustered arrays from Ippon and the Nihon ESS servers, as depicted in Figure 1, of this paper. To review the VDisks that are created using the command in step 5, run the following command:

```
mmlsvdisk | grep eadt8
```



A representative output:

| VDisk name | code | Recovery group | Declustered array | block -size (KB) |
|-----------------------------------------|-----------------|-----------------------|--------------------------|-------------------------|
| rg_nihonio1_hs_Data_512k_2p_1_eadt8 | 8+2p | rg_nihonio1_hs | DA3 | 512 |
| rg_nihonio1_hs_Data_512k_2p_2_eadt8 | 8+2p | rg_nihonio1_hs | DA2 | 512 |
| rg_nihonio1_hs_Data_512k_2p_3_eadt8 | 8+2p | rg_nihonio1_hs | DA1 | 512 |
| rg_nihonio1_hs_MetaData_512k_2p_1_eadt8 | 3WayReplication | rg_nihonio1_hs | DA3 | 512 |
| rg_nihonio1_hs_MetaData_512k_2p_2_eadt8 | 3WayReplication | rg_nihonio1_hs | DA2 | 512 |
| rg_nihonio1_hs_MetaData_512k_2p_3_eadt8 | 3WayReplication | rg_nihonio1_hs | DA1 | 512 |
| rg_nihonio2_hs_Data_512k_2p_1_eadt8 | 8+2p | rg_nihonio2_hs | DA1 | 512 |
| rg_nihonio2_hs_Data_512k_2p_2_eadt8 | 8+2p | rg_nihonio2_hs | DA2 | 512 |
| rg_nihonio2_hs_Data_512k_2p_3_eadt8 | 8+2p | rg_nihonio2_hs | DA3 | 512 |
| rg_nihonio2_hs_MetaData_512k_2p_1_eadt8 | 3WayReplication | rg_nihonio2_hs | DA1 | 512 |
| rg_nihonio2_hs_MetaData_512k_2p_2_eadt8 | 3WayReplication | rg_nihonio2_hs | DA2 | 512 |
| rg_nihonio2_hs_MetaData_512k_2p_3_eadt8 | 3WayReplication | rg_nihonio2_hs | DA3 | 512 |
| rg_ipponio1_hs_Data_512k_2p_1_eadt8 | 8+2p | rg_ipponio1-hs | DA1 | 512 |
| rg_ipponio1_hs_Data_512k_2p_2_eadt8 | 8+2p | rg_ipponio1-hs | DA2 | 512 |
| rg_ipponio1_hs_MetaData_512k_2p_1_eadt8 | 3WayReplication | rg_ipponio1-hs | DA1 | 512 |
| rg_ipponio1_hs_MetaData_512k_2p_2_eadt8 | 3WayReplication | rg_ipponio1-hs | DA2 | 512 |
| rg_ipponio2_hs_Data_512k_2p_1_eadt8 | 8+2p | rg_ipponio2-hs | DA1 | 512 |
| rg_ipponio2_hs_Data_512k_2p_2_eadt8 | 8+2p | rg_ipponio2-hs | DA2 | 512 |
| rg_ipponio2_hs_MetaData_512k_2p_1_eadt8 | 3WayReplication | rg_ipponio2-hs | DA1 | 512 |
| rg_ipponio2_hs_MetaData_512k_2p_2_eadt8 | 3WayReplication | rg_ipponio2-hs | DA2 | 512 |



Along the same lines, the following command offers a representation of the various NSD devices that are created with these VDisks, for the /gpfs/eadata8T file system:

```
mmlsnsd | grep eadt8
```

| File system | Disk name | NSD servers |
|-------------|-----------------------------------------|-----------------------------------------------------------------|
| eadata8T | rg_nihonio1_hs_Data_512k_2p_1_eadt8 | nihonio1-hs.tuc.stglabs.ibm.com,nihonio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio1_hs_Data_512k_2p_2_eadt8 | nihonio1-hs.tuc.stglabs.ibm.com,nihonio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio1_hs_Data_512k_2p_3_eadt8 | nihonio1-hs.tuc.stglabs.ibm.com,nihonio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio1_hs_MetaData_512k_2p_1_eadt8 | nihonio1-hs.tuc.stglabs.ibm.com,nihonio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio1_hs_MetaData_512k_2p_2_eadt8 | nihonio1-hs.tuc.stglabs.ibm.com,nihonio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio1_hs_MetaData_512k_2p_3_eadt8 | nihonio1-hs.tuc.stglabs.ibm.com,nihonio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio2_hs_Data_512k_2p_1_eadt8 | nihonio2-hs.tuc.stglabs.ibm.com,nihonio1-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio2_hs_Data_512k_2p_2_eadt8 | nihonio2-hs.tuc.stglabs.ibm.com,nihonio1-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio2_hs_Data_512k_2p_3_eadt8 | nihonio2-hs.tuc.stglabs.ibm.com,nihonio1-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio2_hs_MetaData_512k_2p_1_eadt8 | nihonio2-hs.tuc.stglabs.ibm.com,nihonio1-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio2_hs_MetaData_512k_2p_2_eadt8 | nihonio2-hs.tuc.stglabs.ibm.com,nihonio1-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_nihonio2_hs_MetaData_512k_2p_3_eadt8 | nihonio2-hs.tuc.stglabs.ibm.com,nihonio1-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_ipponio1_hs_Data_512k_2p_1_eadt8 | ipponio1-hs.tuc.stglabs.ibm.com,ipponio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_ipponio1_hs_Data_512k_2p_2_eadt8 | ipponio1-hs.tuc.stglabs.ibm.com,ipponio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_ipponio1_hs_MetaData_512k_2p_1_eadt8 | ipponio1-hs.tuc.stglabs.ibm.com,ipponio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_ipponio1_hs_MetaData_512k_2p_2_eadt8 | ipponio1-hs.tuc.stglabs.ibm.com,ipponio2-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_ipponio2_hs_Data_512k_2p_1_eadt8 | ipponio2-hs.tuc.stglabs.ibm.com,ipponio1-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_ipponio2_hs_Data_512k_2p_2_eadt8 | ipponio2-hs.tuc.stglabs.ibm.com,ipponio1-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_ipponio2_hs_MetaData_512k_2p_1_eadt8 | ipponio2-hs.tuc.stglabs.ibm.com,ipponio1-hs.tuc.stglabs.ibm.com |
| eadata8T | rg_ipponio2_hs_MetaData_512k_2p_2_eadt8 | ipponio2-hs.tuc.stglabs.ibm.com,ipponio1-hs.tuc.stglabs.ibm.com |

7. The file system, is then mounted using the following command.

```
mount /gpfs/eadata8T -a
```

8. To make this file system a replicated file system, with two replicas for *data* and *metadata*, spread over the ippon and nihon I/O nodes, run the following command on any I/O node.

```
mmchfs -m 2 -M 2 -r 2
```

NSDs that are in the same building block are given same failure group by default. If file system replication is set to 2 (m=2 or r=2), then there should be more than one building block or the failure group of the NSDs must be adjusted accordingly.

9. **Export File system over SMB2 and SMB3 protocols:** Run the following command on the protocol server node (any) to export the eadata8T file system over SMB2 and SMB3 protocols.

```
mmsmb export add eadata8T /gpfs/eadata8T --option "browseable=yes" --option "guest ok=yes"
```



- On the Centricity Enterprise Archive VM, the exported file system can be mounted as a share, as shown in this step.

```
C:\>net use

Status Local Remote Network
-----
OK G: \\9.11.206.24\leadata8T Microsoft Windows Network
```

- Check the health of the entire cluster: Run the `gnrhealthcheck` command on an I/O node of the cluster to validate the health of the cluster. For example:

```
[root@ipponio1 ~]# gnrhealthcheck

#####
# Beginning topology checks.
#####
Topology checks successful.

#####
# Beginning enclosure checks.
#####
Enclosure checks successful.

#####
# Beginning recovery group checks.
#####
Recovery group checks successful.

#####
# Beginning pdisk checks.
#####
Pdisk checks successful.
```

- File system replication:** To check the status of the replicated file system `/gpfs/eadata8T`, run the `lsattr` command. For example:

```
[root@ipponio1 ~]# mmlsattr /gpfs/eadata8T
replication factors
metadata(max) data(max) file [flags]
-----
2 (2) 2 (2) /gpfs/eadata8T [balanced]
```



13. **File system rebalance:** To rebalance the file system, run the `mmrestripefs` command.

```
[root@ipponi01 ~]# mmrestripefs eadata8T -R
Scanning file system metadata, phase 1 ...
Scan completed successfully.
Scanning file system metadata, phase 2 ...
Scanning file system metadata for data storage pool
Scan completed successfully.
Scanning file system metadata, phase 3 ...
Scan completed successfully.
Scanning file system metadata, phase 4 ...
Scan completed successfully.
Scanning user file metadata ...
100.00 % complete on Thu Nov  5 16:54:43 2015 (2396160 inodes with
total 9632 MB data processed)
Scan completed successfully.
```

14. To validate that the security and authentication services are running successfully, run the following command.

```
mmuserauth service check --server-reachability --data-access-method file
Userauth file check on node: ipponpn2-hs.tuc.stglabs.ibm.com
AD servers status
NETLOGON connection: OK
Domain join status: OK
Machine password status: OK
Service 'gpfs-winbind' status: OK
```

15. To validate protocol services, run the following command.

```
[root@nihonpn1 ~]# mmces service list -a
Enabled services: SMB NFS
ipponpn2-hs.tuc.stglabs.ibm.com: SMB is running, NFS is running
ipponpn1-hs.tuc.stglabs.ibm.com: SMB is running, NFS is running
nihonpn1-hs.tuc.stglabs.ibm.com: SMB is running, NFS is running
nihonpn2-hs.tuc.stglabs.ibm.com: SMB is running, NFS is running
```

16. To validate that the CES IP addresses are externally available for NFS, CIFS, object, and other mounts, run the following command.

```
[root@ipponi01 ~]# mmlscluster --ces

GPFS cluster information
=====
GPFS cluster name:          ippon.tuc.stglabs.ibm.com
GPFS cluster id:           11058106192254440982
```



Cluster Export Services global parameters

```
-----
Shared root directory:           /gpfs/cesSharedRoot/ces
Enabled Services:                SMB NFS
Log level:                       0
Address distribution policy:      even-coverage
```

| Node | Daemon node name | IP address | CES IP address list |
|------|---------------------------------|-------------|-------------------------------------|
| 5 | ipponpn1-hs.tuc.stglabs.ibm.com | 9.11.82.67 | 9.11.206.31 9.11.206.25 9.11.206.28 |
| 6 | ipponpn2-hs.tuc.stglabs.ibm.com | 9.11.82.82 | 9.11.206.27 9.11.206.34 9.11.206.32 |
| 7 | nihonpn1-hs.tuc.stglabs.ibm.com | 9.11.103.62 | 9.11.206.30 9.11.206.23 9.11.206.29 |
| 8 | nihonpn2-hs.tuc.stglabs.ibm.com | 9.11.103.64 | 9.11.206.24 9.11.206.26 9.11.206.33 |

17. After any disk errors and after disk replacement procedures, run the following command to ensure that the new disks are being repaired and rebuilt, and run the command for the specific recovery group, as noted.

```
mmlsrecoverygroup rg_ipponio1-hs -L -pdisk (or)
mmlsrecoverygroup rg_ipponio2-hs -L -pdisk (or)
mmlsrecoverygroup rg_nihonio1-hs -L -pdisk (or)
mmlsrecoverygroup rg_nihonio2-hs -L -pdisk
```

Storage estimation and sizing considerations for a VNA solution

GE Healthcare offers and uses estimation tools that facilitate a granular way of planning, estimating, and sizing the storage requirements to accommodate data for annual patient studies of a client healthcare institution for several years. GE Healthcare and IBM architects work together to quantitatively understand the imaging requirements of a client healthcare institution and can ensure accurate sizing of the vendor neutral archive (VNA) solution by considering the following criteria:

- Estimation of annual storage requirement for 1, 3, or 5 years.
- Compression capabilities
- Workflow considerations
- Prevalence of pre-existing images – legacy data migration considerations
- Specific RAID protection levels
- Choice of disk sizes required for short-term storage and long-term storage archives
- Age of data resident on short-term storage and long-term storage archives
- Policy-based data movement between short-term storage and long-term storage
- GE Healthcare image database size and number of extents included to expand the database to accommodate an increased annual volume of studies after three or more years
- Selection of different modalities available for imaging purposes in each hospital facility



While general estimations are offered here, your mileage might vary, based on your unique requirements, different environments, and architectural considerations

Solution benefits – summary

IBM Spectrum Scale software offers the following significant benefits for healthcare clients running Centricity Clinical Archive in their environments.

- It simplifies administration by consolidating short-term storage and long-term storage into a single system
- It includes integrated management tools to help manage petabytes of data and billions of files. This is accomplished within a single namespace that is easy to administer and can be scaled quickly, as required, by simply adding more scale-out resources—helping eliminate *filer sprawl*.
- The new graphical user interface for common administration tasks can speed provisioning, configuration, and monitoring of an IBM Spectrum Scale cluster. It enhances administrator productivity and simplifies new deployments.
- It provides system scalability, very high availability, and reliability with no single point of failure in large-scale storage infrastructures. Administrators can configure the file system so it remains available automatically if a disk or server fails. IBM Spectrum Scale can be configured to automatically recover from node, storage, and other infrastructure failures.
- With support of synchronous and asynchronous features, IBM Spectrum Scale ensures the survivability of data in any failure situation. With Synchronous replication, it also ensures a concurrent data-write feature. With these built-in capabilities, it easily supports data migration, data archiving, data prefetching, and data compression, with VNA, PACS, and medical imaging software applications.

Acknowledgments

Special thanks to the GE Healthcare product management and engineering teams in the United States and Europe, for loaning the software and licenses that enabled the IBM engineering teams to successfully create an operational test environment and run tests to document real-life results.

Many thanks to the IBM client executives, IBM Systems members, and other associates worldwide, who contributed with their recommendations during the test run and review process and enabled successful completion and validation so that Centricity Clinical Archive can run successfully on IBM Spectrum Scale software.

Glossary

ADT: Healthcare Industry term that stands for admission, discharge, and transfer.

Angiography: A procedure used by clinicians to view blood vessels by injecting a radiopaque dye that outlines these vessels on an X-ray image.

CR: Computed radiography system for imaging.

CT: Computed tomography system for medical imaging. CT was earlier called computerized axial tomography. It is a radiographic technique that produces an image of a detailed cross section of tissue. CT uses a narrowly collimated beam of X-rays that rotates in a full arc around the patient to image the body in cross-sectional slices. An array of detectors, positioned at several angles, records those X-rays that pass through the body. The image is created by a computer that uses multiple attenuation readings taken around the periphery of the body part. The computer calculates tissue absorption and produces a representation of the tissues that demonstrates the densities of the various structures. Tumor masses, infarctions, bone displacement, and accumulations of fluid can be detected. For cardiologic examination, ultrafast CT is electrocardiogram-triggered and this allows visualization of cardiac function and blood flow.

Digital Imaging and Communications in Medicine (DICOM): Is a worldwide standard for exchanging patient images and data. It enables the exchange of images and data between heterogeneous information systems and various imaging and image producing and processing devices (modalities).

Enterprise Archive (EA): GE Healthcare Centricity Enterprise Archive application.

Enterprise master patient index (EMPI): A database index that enables the merging of different medical record numbers (MRN) into a single patient database record.

Electronic medical record (EMR): The electronic record of health information on an individual that is created, gathered, managed, and consulted by licensed clinicians and staff from a single organization who are involved in the individual's health and care.

HL-7: Health Level-7 or **HL7** refers to a set of international standards for transfer of clinical and administrative data between software applications used by various healthcare providers. These standards focus on the application layer, which is **layer 7** in the OSI model.

HIPAA: The Health Insurance Portability and Accountability Act of 1996 was enacted by the United States Congress and signed by President Bill Clinton in 1996. It has been known as the Kennedy-Kassebaum Act or Kassebaum-Kennedy Act after two of its leading sponsors.

HIS: Hospital information system is a comprehensive, integrated information system designed to manage all the aspects of a hospital operation, such as medical, administrative, financial, legal, and the corresponding service processing.

Hossim: GE Healthcare hospital simulation system is a test tool and programming language engineered by GE Healthcare. It is used to produce DICOM traffic such as C-STORE, C-FIND and C-MOVE requests. Hossim is used to produce a load on a DICOM archive such as EA.

Integrating the Healthcare Enterprise (IHE): An initiative by users and companies to achieve maximum interoperability of the digital information systems used in health services. The basis for this is the

integration of the profiles that are based on the working procedures actually followed in radiology. Further information about IHE available at: www.ihe.net.

ILM: Image lifecycle management is a feature in EA to manipulate the DICOM images and XDS documents stored in the archives. The most important operations in EA ILM are as follows:

- Move files from one storage library to another
- Compress DICOM images
- Delete studies or XDS documents

Joint Photographic Experts Group (JPEG): A standardization group for the coding and compression of gray scale and color images. JPEG is a graphics format that enables up to 20-fold compression of individual images. The names of files in the JPEG format usually have the extension *.jpg*.

JPEG 2000: A wavelet-based image compression standard. It was created by the JPEG committee in the year 2000 with the intention of superseding their original discrete cosine transform-based JPEG standard. Compared to JPEG, there is an increase in compression performance of JPEG 2000. The main advantage offered by JPEG 2000 is the significant flexibility of the code stream.

The code stream of a JPEG 2000 compressed image is scalable in nature. It can be decoded in a number of ways. For example, by truncating the code stream at any point, you can obtain a representation of the image at a lower resolution.

Lightweight Directory Access Protocol (LDAP): An application protocol for querying and modifying directory services running over TCP/IP.

Long-term storage: A storage medium for long-term storage or archiving of images.

Mammography: An X-ray imaging procedure of human breasts, to detect tumors, cancers or other abnormalities in breast tissues.

Metadata: *Data about the data*. In DICOM, each image is marked with properties, or attributes, such as patient name, study date, and so on. These attributes constitute the metadata.

Megapixel (MP): Abbreviation for the resolution of monitors. Additionally,

1.3 MP = 1280 × 1024 pixels

2.0 MP = 1600 × 1200 pixels

3.0 MP = 1536 × 2048 pixels

5.0 MP = 2048 × 2560 pixels

Modalities: Different types of radiography equipment and imaging scanners that capture, and transmit medical images. These could be MRI, CT, CR, X-ray, ultrasound, and so on.

MR / MRI: Magnetic resonance imaging is a non-invasive **diagnostic** technique that produces computerized images of internal body tissues and is based on nuclear **magnetic resonance** of atoms within the body induced by the application of radio waves

Medical record number (MRN): An index number for patient identification that is assigned by a HIS or RIS.

National Electrical Manufacturers Association (NEMA): An industry consortium, which in 1983 began a collaboration with the American College of Radiology (ACR) to develop the imaging standard now known as DICOM.

Nuclear medicine: *Nuclear medicine* is a *medical* specialty involving the application of radioactive substances in the diagnosis and treatment of disease.

Picture archiving and communication system (PACS): It facilitates the capture, collection, transport and archival of medical images of patients in digital formats.

PET: Positron emission tomography is a computerized radiographic technique that uses radioactive substances to examine the metabolic activity of various body structures. The patient ingests a substance similar to radioactive glucose, before the imaging procedure. During the tomography imaging procedure, the positrons from the ingested radioactive glucose combine with electrons normally found in the cells of the body, and emit gamma rays. The electronic circuitry and computers of the PET device detect the gamma rays and construct color-coded images of organs and malignancies. The radioactive isotopes used in PET are very short-lived, so that patients undergoing a PET scan are exposed to very small amounts of radiation.

PET-CT: Positron emission tomography-computed tomography is a medical technique that uses two different scanners (or different modalities) to capture images from a PET scanner, and a CT scanner, in the same session, for more accurate clinical visualization and rendering of patient anatomies and images.

POST: A power-on self-test (POST) is a process performed by firmware or software routines immediately after a computer, system, device or other digital electronic appliance is powered on.

Radiology information system (RIS): A system for the radiology department, which takes care of patient administration, documentation, acknowledgment of services, and writing the reports.

Solid-state drive (SSD): In systems and applications, refers to electronic memory storage.

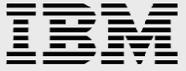
Short-term storage: Is based on storage technologies and connections that allow fast access to image data. Typically, a RAID system that is connected directly to the server is used as short-term storage.

Ultrasound: A method of diagnosing illness and viewing internal body structures in which sound waves of high frequency are bounced off internal organs and tissues from outside the body.

VNA: A Vendor Neutral Archive (**VNA**) is a medical imaging technology in which images and documents (and potentially any file of clinical relevance) are stored (archived) in a standard format with a standard interface, such that they can be accessed in a vendor-neutral manner by other systems. The Vendor Neutral Archive system must meet following requirements.

- It must interface with other clinical systems and disparate PACS imaging equipment for the purpose of communicating imaging data.
- It must interface with other clinical information systems for communication of reports, results, workflow, and so on.
- It must store all objects in a non-proprietary format understood by the community at large.

XDS: XDS stands for cross-enterprise document sharing. Healthcare organizations can use XDS for sharing non-image data with each other. Examples of non-image data include radiology reports, lab reports, primary care physician's diagnoses and emergency-room care records.



X-ray / X-ray RF: Typically, an X-ray beam is focused on a body / body part. While some of the rays are absorbed, any remaining rays are captured on a detector plate, as computerized image. This is X-ray radiography. In fluoroscopy, the image is viewed real time, on a video-screen, while the X-ray beam is focused on the body. It allows the clinicians to visualize the body part and its movements, through the ingestion of contrast agents, such as dyes. Typically, bone fractures are identified through X-ray radiography, while gastroenterological problems are identified through X-ray fluoroscopy.

Resources

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

- GE Healthcare Common Documentation Library
<http://apps.gehealthcare.com/servlet/ClientServlet?REQ=Enter+Documentation+Library>
- IBM Systems on PartnerWorld
ibm.com/partnerworld/systems
- IBM Redbooks
ibm.com/redbooks
- IBM Publications Center
www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US

About the authors

Prashant Avashia is a Software Engineer 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

Gustavo Hernandez is a Software Architect in GE Healthcare, currently involved in the architecture and development of Centricity Clinical Archive, a solution based on IHE standards such as DICOM and XDS and adopted by hospitals and data centers across the world. You can reach Gustavo at GustavoHernandez@ge.com

Howard Jiang is a Network Systems Engineer in IBM Systems Group, currently focused on implementing and integrating PACS, VNA, and RIS imaging solutions in the IBM Tucson Laboratory. You can reach Howard at hjiang@us.ibm.com



Trademarks and special notices

© Copyright IBM Corporation 2016.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this 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 www.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.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

INFINIBAND, InfiniBand Trade Association and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association.

GE and Centricity are trademarks of General Electric Company. Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.



Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

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.

This configuration was installed and tests were performed at an IBM facility. The results expressed in this document may not be applicable to a particular site or installation and individual results may vary.

This document and its contents are provided to you for informational purposes only and do not constitute a representation, warranty or performance guarantee. GE and IBM disclaim liability for any loss, which may arise from reliance on or use of information, contained in this document.