HITACHI
Inspire the Next

Hitachi Virtual Storage Platform Gx00 Models with Epic Systems and Hitachi Thin Image using Flash Module Drives

TECHNICAL PAPER

Reference Architecture Guide

By Diana Milan

July 2016

Hitachi Data Systems

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Contents

Hitachi Virtual Storage Platform Gx00 Models with Epic Systems and Hitachi Thin Image using Flash Module Drives

Reference Architecture Guide

This reference architecture guide shows how the Hitachi Virtual Storage Platform Gx00 storage system environment provides high performance in an integrated solution for an Epic Systems infrastructure.

This document is intended for system engineers, support personnel, and administrators of systems involved in the early stages of planning and designing using Hitachi Virtual Storage Platform Gx00 models in an environment with the Epic online transaction processing database. It documents all applicable configuration options that allow the environment to meet established performance targets. You should have knowledge of the following:

- Hitachi Virtual Storage Platform Gx00 methods
- Hitachi Thin Image methods
- Hitachi Dynamic Provisioning methods

Note — The described architecture in this document was designed specifically for Epic Electronic Record production software and does not reflect other application I/O patterns, nor do they have any influence on the OLTP production by other activities that might run simultaneously on the same array. The guide is strictly for Epic Electronic Record production software type workload and should not be extrapolated to predict performance for other types of workloads on Hitachi Data Systems storage products.

Solution Overview

The Hitachi Virtual Storage Platform Gx00 environment provides high performance in an integrated solution for an Epic Systems infrastructure for typical on-line transaction processing workloads. Tailor your implementation of this solution to meet your specific needs.

Epic makes software for mid-size and large medical groups, hospitals, and integrated healthcare organizations.

This reference architecture guide describes the effects that the amount of cache and number of drives have on deploying this solution with and without Hitachi Thin Image. This shows the performance impact of using Hitachi Thin Image Heterogeneous and different drive counts.

This reference architecture describes the following:

- Hitachi Virtual Storage Platform Gx00 environment baseline with flash module drives (FMDs)
- Hitachi Virtual Storage Platform Gx00 models with FMDs using Hitachi Thin Image
- Hitachi Virtual Storage Platform Gx00 models with FMDs using Hitachi Thin Image, pair operations, and accessing secondary copy
- Hitachi Virtual Storage Platform Gx00 environment baseline with multiple FMD RAID groups

Key Solution Components

Table 1 and Table 2 list the key hardware and software components used in this reference architecture.

Table 1. Reference Architecture Hardware Components

Hardware	Description	Version
Hitachi Virtual Storage Platform Gx00 models (VSP G800)	■ 24 × 1.6 TB FMD devices ■ 512 GB total cache	83-03-20-60/01-08 SVOS6.4.0
Hitachi Virtual Storage Platform Gx00 models (VSP G600)	■ 24 × 1.6 TB FMD devices ■ 256 GB total cache	83-03-20-60/01-08 SVOS6.4.0
Hitachi CB2500 blade server	■ 4 × CB520x B2 ■ Four paths per LUN	A0113/00

Table 2. Reference Architecture Software Components

Software	Version
Hitachi Thin Image (HTI)	83-03-20-60/01-08 or later
Hitachi Dynamic Provisioning (HDP)	83-03-20-60/01-08 or later

Hitachi Virtual Storage Platform Gx00 Models

Hitachi Virtual Storage Platform Gx00 models are based on industry-leading enterprise storage technology. With flash-optimized performance, these systems provide advanced capabilities previously available only in high-end storage arrays. With the Virtual Storage Platform Gx00 models, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Gx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort.

Hitachi Thin Image

Hitachi Thin Image is a storage-based solution that creates RAID-protected snapshots of volumes within Hitachi Virtual Storage Platform family products. The resulting snapshots are cost-effective local by storing only differential data between primary volumes (P-VOLs) and secondary volumes (S-VOLs) of virtual volumes (V-VOLs).

Detailed information on using Thin Image is in *Hitachi Virtual Storage Platform Hitachi Thin Image User Guide* (MK-90RD7179-05).

Hitachi Dynamic Provisioning

On Hitachi storage systems, Hitachi Dynamic Provisioning provides wide striping and thin provisioning functionalities.

Using Dynamic Provisioning is like using a host-based logical volume manager (LVM), but without incurring host processing overhead. It provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) of a logical size you specify of up to 60 TB created against it without initially allocating any physical space.

Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. Dynamic Provisioning distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Dynamic Provisioning has the benefit of thin provisioning. Physical space assignment from the pool to the dynamic provisioning volume happens as needed using 42 MB pages, up to the logical size specified for each dynamic provisioning volume. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. You can rebalance an expanded pool across the current and newly added RAID groups for an even striping of the data and the workload.

Solution Design

This reference architecture implements an Hitachi Virtual Storage Platform (VSP) G800 and VSP G600 to host on OLTP database Epic Electronic Record production software.

The following are the configuration options for this reference architecture:

- **Cache size** — 64, 128, 256, and 512 GB
- **Hitachi Thin Image** — Option of whether or not to use Hitachi Thin Image.

The description of the tested hardware components is in the Key Solution Components section. Specific infrastructure configuration includes the following:

- **Server** — This architecture uses a single Hitachi CB2500 blade server with ESXi 6.0U1 server configuration.
- **Storage System** — This architecture uses LDEVs mapped to each of four ports that are presented to the server as LUNs.
- **SAN Fabric** — This is a direct attached Storage Architecture

This is the storage architecture of this reference architecture. It takes into consideration Hitachi Data Systems and Epic recommended practices for the deployment of the storage design.

Storage Configuration

This reference architecture uses RAID groups and storage pools created with Hitachi Dynamic Provisioning on VSP G800 and VSP G600.

For the 8 drive configuration, one RAID group is created. For the 16 drive configuration, two RAID groups are created. These are RAID-6 (6D+2P) for the following reasons:

- Redundancy (parity is mirrored)
- Performance (blocks are stripped)

Use Hitachi Dynamic Provisioning to eliminate hotspots. In both configurations, create eight LDEVs out of each RAID group. Add all LDEVs to a pool created with Hitachi Dynamic Provisioning.

Create 8 LUNs for production data and map them to four paths to the hosts.

An Epic environment requires at least a 128 GB CLPR size for VSP G800 and at least a 64 GB CLPR size for VSP G600.

Server and Application Architecture

The reference architecture uses one Hitachi Compute Blade 2500. It includes a single Epic OLTP simulator instance. The simulator is a set of read processes doing random read operations and write operations writing to memory and flushing data out to disk every 80 seconds. The read-to-write ratio is 3:1.

There are four CB520x B2 blades with 64 GB RAM. This provides the compute power for Epic Electronic Record production software to handle a large volume of transaction processing in parallel.

SAN Architecture

Map the provisioned LDEVs to multiple ports on VSP Gx00 models. These LDEV port assignments provide multiple paths to the storage system from the host for high availability.

When designing your SAN architecture, follow these recommended practices for a secure, high-performance, and scalable database deployment:

- Use at least two HBAs and place them on different I/O buses within the server. This distributes the workload over the PCI-e bus architecture of the server.
- Use dual SAN fabrics and host-based multipathing software in a business-critical deployment. Connecting two or more paths from the database servers to two independent SAN fabrics provides the redundancy required for critical applications.
- Zone your fabric to handle multiple, unique paths from HBAs to storage ports. Use single initiator zoning where each zone has only a single HBA and one or more storage ports.
- Set MP auto assignment to "enabled".

Network Architecture

This reference architecture has two on-board 1 Gb/sec NIC ports for connectivity and redundancy.

Engineering Validation

This is how the OLTP database used by Epic Electronic Record production software was validated on Hitachi's entry enterprise storage systems, Hitachi Virtual Storage Platform Gx00 models.

Test Methodology

The tests used the RampRun.pl script from Epic. This script calls the Epic simulation tool GenerateIO. While RampRun.pl is responsible for kicking off various monitoring tools to capture metrics throughout the test, GenerateIO is the simulator that generates the I/O workload and patterns similar to an OLTP production environment using Epic Electronic Record production software. GenerateIO ramped up the I/O load to cover a wide range of IOPS levels, as expected in production environments.

The I/O latency data gathered during these tests include the following:

- Read response time measured by the Ranread utility from Epic
  - The Ranread utility is used in production environments to measure read response times.
- Write response time and write cycle length logged by the GenerateIO simulator

Epic determined a list of I/O performance targets to ensure that end-user experience is not affected by any bottlenecks at the storage level. These are the I/O performance targets:

- Average read latencies must be 12 m/sec or less
- 99% of read latencies must be below 60 m/sec
- 99.9% of read latencies must be below 200 m/sec
- 99.99% of read latencies must be below 600 m/sec
- The write cycle must be completed within 45 seconds

Workload

Testing included scenarios to ensure the following:

- The write response time was below 1 m/sec, preferably in the 0.4 m/sec to 0.6 m/sec range
- The write daemon cycle was in the 45 second range

Testing Procedure

When testing the environment, run the tests for each of these server cache size partitions:

- On VSP G800 use 128, 256, and 512 GB
- On VSP G600 use 64, 128, and 256 GB

To validate the solution, this testing was done on the reference architecture:

1. Ran the testing procedure on the environment for each of the server cache sizes to establish the performance baseline data without using the following:
    - Hitachi Thin Image
2. Ran the testing procedure on the environment adding Hitachi Thin Image to the base configuration.
3. Ran the testing procedure on the environment adding Hitachi Thin Image to the base configuration and executed pair operations as well as secondary copy read access.
4. Ran the testing procedure on the environment adding a second RAID group to the base configuration.

The **GenerateIO Test** testing procedure measured the following in the environment with the different cache sizes and number of drives:

- Front and back-end IOPS
- Read and write latency
- "iostat"

Data Gathering

Performance statistics were collected at the following levels:

- **Storage**
    - Hitachi Storage Navigator collected storage performance data
- **Operating System**
    - GenerateIO collected operating system statistics

Test Results

This summarizes the key observations from the test results for using Epic Electronic Record production software with Hitachi Virtual Storage Platform Gx00 models.

Baseline Configuration Test

This summarizes the results from the baseline configuration test.

- Using a single FMD RAID group in Hitachi Virtual Storage Platform Gx00 models exhibits the following performance:
    - A single RAID group configuration can drive up to 55,000 host IOPS on the VSP G800
    - A single RAID group configuration can drive up to 45,000 host IOPS on the VSP G600
- There is a no significant difference for VSP G800 in terms of read and write performance between 128, 256, and 512 GB cache under an Epic workload.
- There is a no significant difference for VSP G600 in terms of read and write performance between 64, 128, and 256 GB cache under an Epic workload.
- These cache configuration options free up cache in the storage system to be shared with other applications.

Hitachi Thin Image Configuration Test

In this test we set up the target disks or secondary volume (s-VOL) with RAID-6 (6D+2P) disk groups.

Comparing the results to the baseline configuration, the pairing operation did not affect performance, but the split pair operation increased the write cycle time in milliseconds by 400%. This is the operation to perform backups so the IOPS activity during this period should be lower than during regular operation hours.

Hitachi Thin Image Configuration Test with Secandary Volume Read Operation

In this test the target disks or secondary volume (s-VOL) were split from the source disks or primary volume (p-VOL) and a read operation was executed on the secondary volume. This scenario mimics the activity from a backup server in the event of a backup.

Comparing the results to the baseline configuration, the read operation executed on the secondary volume did not affect the read or write performance.

Multiple RAID Group Configuration Test

These are the results of the multiple RAID group configuration test.

- Using multiple FMD RAID groups in Hitachi Virtual Storage Platform Gx00 models exhibits a performance improvement of 10% from the baseline configuration.
  - A two RAID group configuration can drive up to 68,000 host IOPS on the VSP G800
  - A two RAID group configuration can drive up to 57,000 host IOPS on the VSP G600

Conclusion

The major finding is that write cycle time is not dependent on the cache partition size. With one of the I/O performance targets being 45 seconds the cache size must be set to at least 128 GB for VSP G800 for a small implementation and 256 GB or 512 GB for a large implementation and at least 64 GB for VSP G600 for a small implementation and 128 GB or 256 GB for a large implementation regardless of whether or not Hitachi Thin Image is present.

For cache optimization, a cache partition dedicated to the production database for Epic Electronic Medical Record production software must be used and LUNs must be distributed across all MP blades.

A multiple RAID group configuration is recommended in order to take advantage of the performance improvement and space.

If Hitachi Thin Image is chosen as the replication tool for backup, a low I/O window must be established during the pair split operation in order to maintain expected performance for Epic Electronic Medical Record production software

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems Global Services website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate Resources website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services Education website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the Hitachi Data Systems website.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

AS-518-00 July 2016.