



Technical Report

NetApp Sizing Guidelines for MEDITECH Environments

Brahmanna Chowdary Kodavali, NetApp
March 2016 | TR-4190

TABLE OF CONTENTS

1	Introduction	4
1.1	Scope	4
1.2	Audience	5
2	MEDITECH and BridgeHead Workloads	5
2.1	MEDITECH Workload Description	5
2.2	BridgeHead Workload Description	7
3	Sizing NetApp Storage for MEDITECH and BridgeHead Workloads	8
3.1	Sizing Methodology for NetApp Storage Systems	8
3.2	Sizing for MEDITECH Production and BridgeHead Backup Workloads	9
3.3	Sizing Parameters for NetApp SPM Sizing Tool	11
3.4	Using SPM Sizing Results	18
3.5	Additional Sizing Requirements	18
3.6	Additional Storage for MEDITECH Environments	19
4	Sizing Examples	19
4.1	Sizing Example for 20-Host MEDITECH (Category 2) 6.x and BridgeHead Workloads with Clustered Data ONTAP 8.3 and Flash Pool	20
4.2	All Flash FAS Sizing Example for 60-Host MEDITECH (Category 2) 6.x and BridgeHead Workloads with Clustered Data ONTAP 8.3.1	33
5	Sizing and Storage Layout Recommendations to Deploy 7 MEDITECH (Category 1) 6.x Hosts and FAS2240HA Storage System with 24 Internal Disks	44
5.1	MEDITECH, BridgeHead, and NetApp FAS Storage Deployment Requirements	44
5.2	Storage Configuration Recommendations	45
5.3	Storage Layout for FAS2240HA with a Maximum of 24 Internal Disks	45
6	MEDITECH Certified and Validated SPM Storage Configurations	46
6.1	Certified and Validated SPM Storage Configurations	46
6.2	Performance Validation	47
	References	48
	Version History	48

LIST OF TABLES

Table 1)	MEDITECH platform-specific I/O characteristics	6
Table 2)	Summary of MEDITECH workload I/O characteristics and requirements	7
Table 3)	Summary of BridgeHead workload I/O characteristics and requirements.	8
Table 4)	Fixed sizing parameters: assumptions for the MEDITECH production workload.	11

Table 5) Fixed sizing parameters: assumptions for the BridgeHead workload	12
Table 6) Fixed sizing parameters: assumptions for the storage hardware configuration.....	14
Table 7) Variable sizing parameters: assumptions for the MEDITECH production workload.	14
Table 8) Variable sizing parameters: assumptions for the BridgeHead workload.....	16
Table 9) Variable sizing parameters: assumptions for the storage hardware configuration.	16
Table 10) Deployment information for 20-host MEDITECH (category 2) 6.x environment with Flash Pool sizing.	20
Table 11) Example of MEDITECH workload SPM parameter values for 20-host MEDITECH (category 2) 6.x environment.....	23
Table 12) Example of BridgeHead workload SPM parameter values for 20-host MEDITECH environment.	26
Table 13) Example of storage hardware configuration SPM parameter values for 20-host MEDITECH environment with FAS8040 and Flash Pool.	29
Table 14) Drive calculations per aggregate in a 20-host environment (from SPM report).....	30
Table 15) Disk count summary for FAS8040HA storage system configured with Flash Pool for 20-host MEDITECH (category 2) 6.x environment.....	31
Table 16) Adjusted disk count summary for FAS8040HA storage system configured with Flash Pool for 20-host MEDITECH (category 2) 6.x environment.	32
Table 17) Deployment information for 60-host MEDITECH environment with Flash Pool sizing.....	33
Table 18) Example of MEDITECH workload SPM parameter values for 60-host MEDITECH environment.....	36
Table 19) Example of BridgeHead workload SPM parameter values for 60-host MEDITECH environment.	39
Table 20) Example of storage hardware configuration SPM parameter values for 60-host MEDITECH environment with AFF8060.	42
Table 21) Drive calculations per aggregate in a 60-host environment (from SPM report).....	43
Table 22) Deployment requirements for seven MEDITECH (category 1) 6.x hosts with one BridgeHead backup server and an FAS2240HA storage system with internal disks.	44
Table 23) Disk layout for the FAS2240HA storage system configured with 24 internal disks and Flash Pool for 7 MEDITECH (category 1) 6.x hosts.	45
Table 24) Certified storage configurations for MEDITECH environments.	46

LIST OF FIGURES

Figure 1) Screenshot of SPM custom application for MEDITECH (category 2) 6.x workload with 20 hosts.....	22
Figure 2) Screenshot of SPM custom application for BridgeHead workload.	25
Figure 3) Screenshot of SPM storage configuration.....	28
Figure 4) Projected system utilization across two controllers in a 20-host MEDITECH environment.	30
Figure 5) Screenshot of SPM custom application for MEDITECH (category 2) 6.x workload.....	35
Figure 6) Screenshot of SPM custom application for BridgeHead workload.	38
Figure 7) Screenshot of SPM storage configuration.....	41
Figure 8) Projected system utilization across two controllers in a 60-host MEDITECH environment.	43

1 Introduction

This document provides sizing guidelines for NetApp® storage that supports MEDITECH environments.

Section 2, “MEDITECH and BridgeHead Workloads,” describes the I/O characteristics and the performance requirements of the MEDITECH production and BridgeHead backup workloads.

Section 3, “Sizing NetApp Storage for MEDITECH and BridgeHead Workloads,” describes how NetApp technologies are applied to achieve the MEDITECH and BridgeHead workload performance requirements. This section also describes the sizing methodology used to determine NetApp storage sizing according to the types of MEDITECH and BridgeHead workloads.

Section 4, “Sizing Examples,” provides examples of how to size NetApp All Flash FAS (AFF) and hybrid storage with NetApp Flash Pool™ solutions for the MEDITECH and BridgeHead workloads.

Section 5, “Sizing and Storage Layout Recommendations to Deploy 7 MEDITECH (Category 1) 6.x Hosts and FAS2240HA Storage System with 24 Internal Disks,” provides the storage sizing and layout recommendations for small MEDITECH deployments.

Section 6, “MEDITECH Certified and Validated SPM Storage Configurations,” lists the set of validated storage configurations certified by MEDITECH and describes the performance metrics validated.

1.1 Scope

This document describes how to use the [NetApp System Performance Modeler](#) (SPM) tool to size the NetApp FAS storage platform required for the MEDITECH hosts (database servers) in a production environment that uses either NetApp Data ONTAP® 8 operating in 7-Mode or clustered Data ONTAP 8. This document is based on SPM version 2.2.

The document covers using the NetApp SPM sizing tool for MEDITECH hosts (database servers) that use BridgeHead backup software.

This document does not cover the following subjects:

- Use of NetApp sizing tools for non-MEDITECH environments
- Use of NetApp sizing tools for other MEDITECH servers (see the following note)
- Virtualization or Windows Server operating system storage requirements for MEDITECH file servers
- Sizing for nonproduction workloads
- Best practices, deployment methodology, or architecture

Note: This document covers sizing guidelines for only the MEDITECH hosts (also known as database servers, file servers, or MAGIC machines). When sizing your NetApp storage, consider other MEDITECH servers that might exist in the environment. These servers could be the MEDITECH data repository application, the scanning and archiving application, background job clients, connection servers, print servers, and so on. NetApp systems engineers in the field should understand all MEDITECH workloads intended to run on the customer's NetApp storage. These engineers should consult with the NetApp MEDITECH Independent Software Vendor team to determine a proper and complete sizing configuration.

The information provided in this document is based on the following MEDITECH certified and tested storage configurations:

- AFF8060HA on clustered Data ONTAP 8.3.1
- FAS8040HA on clustered Data ONTAP 8.3
- FAS3250HA and FAS3220HA on Data ONTAP 8.1.2 7-Mode
- FAS3250HA on clustered Data ONTAP 8.2
- FAS2552HA on clustered Data ONTAP 8.2.x
- FAS2240HA on clustered Data ONTAP 8.2

For more information, see section 6, “MEDITECH Certified and Validated SPM Storage Configurations.”

1.2 Audience

This document is for NetApp and partner systems engineers and professional services personnel. NetApp assumes that the reader has the following types of knowledge:

- A good understanding of storage performance concepts and I/O characterization
- Technical familiarity with NetApp storage systems and the SPM tool

E-mail questions or comments about this technical report to ng-healthcare-dsg@netapp.com.

2 MEDITECH and BridgeHead Workloads

When you size NetApp storage systems for MEDITECH environments, you must consider both the MEDITECH production workload and the BridgeHead backup workload. BridgeHead software is the certified and validated solution for managing and backing up MEDITECH databases by using NetApp storage.

MEDITECH Host Definition

A MEDITECH host is essentially a database server. Depending on the software platform, the host might also be referred to as a MEDITECH file server or a MAGIC machine. The rest of this document uses the term “MEDITECH host” to refer to the MEDITECH file server and the MAGIC machine.

The following sections describe the I/O characteristics and performance requirements of these two workloads.

2.1 MEDITECH Workload Description

In a MEDITECH environment, multiple servers that run MEDITECH software perform various tasks as an integrated system (referred to as the MEDITECH system). For more information, see the MEDITECH documentation:

- **For production MEDITECH environments**, consult the appropriate MEDITECH documentation to determine the number of MEDITECH hosts and storage capacity that must be included as part of sizing the NetApp storage system.
- **For new MEDITECH environments**, consult the hardware configuration proposal document. For existing MEDITECH environments, consult the hardware evaluation task document. The hardware evaluation task document is associated with a MEDITECH ticket. Customers can request either of these documents from MEDITECH.

You can scale the MEDITECH system to provide increased capacity and performance by adding hosts. Each host requires storage capacity for its database and application files. The storage presented to each MEDITECH host must also be capable of supporting the I/O generated by the host. In a MEDITECH environment, a LUN is presented to each host to support that host’s database and application storage requirements. The type of MEDITECH category and the type of platform that you deploy determine the workload characteristics of each MEDITECH host and, consequently, of the system as a whole.

MEDITECH Categories

MEDITECH associates the deployment size with a category number ranging from 1 to 6. Category 1 represents the smallest MEDITECH deployments; category 6 represents the largest. Examples of the MEDITECH applications specifications associated with a category include metrics such as the number of hospital beds, inpatients per year, outpatients per year, Emergency Room visits per year, exams per

year, inpatient prescriptions per day, outpatient prescriptions per day, and so on. For more information, see the MEDITECH category reference sheet. You can obtain this sheet from MEDITECH through the customer or through the MEDITECH system installer.

Note: MEDITECH category 1 hosts have less demanding random read latency than hosts in MEDITECH categories 2 through 6. Category 1 hosts have average IOPS per host requirements compared to hosts in categories 2 through 6. Hosts in MEDITECH categories 3 through 6 have the same random read and write latency as MEDITECH category 2 hosts. Hosts in categories 3 through 6 have average requirements for IOPS per host compared to MEDITECH category 2 hosts. MEDITECH categories 2 through 6 differ in the number of hosts deployed. For more information, see Table 1 and Table 2. This document focuses on sizing the NetApp FAS storage system for hosts in MEDITECH categories 1 and 2.

MEDITECH Platforms

MEDITECH has three platforms:

- MEDITECH 6.x
- Client/Server 5.x (C/S 5.x)
- MAGIC

For the MEDITECH 6.x and C/S 5.x platforms, the I/O characteristics of each host are defined as 100% random with a request size of 4K. For the MEDITECH MAGIC platform, each host's I/O characteristics are defined as 100% random with a request size of either 8K or 16K. According to MEDITECH, the request size for a typical MAGIC production deployment is either 8K or 16K.

The ratio of reads and writes varies depending on the platform deployed. MEDITECH provides an estimate of the average mix of read and write percentages, and the average sustained I/O per second (IOPS) value required for each MEDITECH host on a particular MEDITECH platform. Table 1 summarizes the platform-specific I/O characteristics provided by MEDITECH.

Table 1) MEDITECH platform-specific I/O characteristics.

MEDITECH Category	MEDITECH Platform	Average Random Read %	Average Random Write %	Average Sustained IOPS per MEDITECH Host
1	6x	20	80	250
2–6	6x	20	80	750
	C/S 5.x	40	60	600
	MAGIC	90	10	400

In a MEDITECH system, the average IOPS level of each host should equal the IOPS values defined in Table 1. The IOPS values specified in Table 1 are used as part of the sizing methodology described in section 3.1, "Sizing Methodology for NetApp Storage Systems," to determine the correct storage sizing based on each platform.

MEDITECH requires the average random write latency to stay below 1ms for each host. However, temporary increases of write latency up to 2ms during backup and reallocation jobs are considered acceptable. MEDITECH also requires the average random read latency to stay below 7ms for category 1 hosts and below 5ms for category 2 hosts. These latency requirements apply to every host regardless of the MEDITECH platform.

Table 2 summarizes the I/O characteristics that you must consider when you size NetApp storage for MEDITECH workloads.

Table 2) Summary of MEDITECH workload I/O characteristics and requirements.

MEDITECH Category	Parameter	MEDITECH 6.x Platform	C/S 5.x Platform	MAGIC Platform
1–6	Request size	4K	4K	8K or 16K (See note following table.)
	Random/sequential	100% random	100% random	100% random
1	Average sustained IOPS	250	N/A	N/A
2–6	Average sustained IOPS	750	600	400
1–6	Read/write ratio	20% read, 80% write	40% read, 60% write	90% read, 10% write
	Write latency	<1ms	<1ms	<1ms
	Temporary peak write latency	<2ms	<2ms	<2ms
1	Read latency	<7ms	N/A	N/A
2–6	Read latency	<5ms	<5ms	<5ms

Note: MEDITECH hosts in categories 3 through 6 have the same I/O characteristics as those in category 2. For MEDITECH categories 2 through 6, each category differs in the number of hosts deployed.

Note: According to MEDITECH, the request size for a typical MAGIC production deployment is either 8K or 16K.

The NetApp storage system should be sized to satisfy the performance requirements described in Table 2. In addition to the MEDITECH production workload, the NetApp storage system must be able to maintain these MEDITECH performance targets during backup operations by BridgeHead, as described in section 2.2, “BridgeHead Workload Description.”

2.2 BridgeHead Workload Description

BridgeHead backup software backs up the LUN used by each MEDITECH host in a MEDITECH system. For the backups to be in an application-consistent state, the backup software quiesces the MEDITECH system and suspends I/O requests to disk. While the system is in a quiesced state, the backup software issues a command to the NetApp storage system to create a NetApp Snapshot® copy of the volumes that contain the LUNs. The backup software subsequently unquiesces the MEDITECH system, which allows production I/O requests to continue to the database. The software creates a NetApp FlexClone® volume based on the Snapshot copy. This volume is used as the backup source while production I/O requests continue on the parent volumes that host the LUNs.

The workload generated by the backup software results from the sequential reading of the LUNs that reside in the FlexClone volumes. The workload is defined as a 100% sequential read workload with a request size of 64K. For the MEDITECH production workload, the performance criterion is to maintain the required IOPS and the associated read/write latency levels. For this backup workload, however, the attention is shifted to the overall data throughput (MBps) generated during the BridgeHead backup operation. Specifically, BridgeHead requires the backup of all MEDITECH LUNs to be completed within an eight-hour backup window. NetApp recommends that the backup of all MEDITECH LUNs be completed in six hours or less. Doing so compensates for events such as an unplanned increase in the MEDITECH workload, NetApp Data ONTAP background operations, or data growth over time. Any of these events might incur additional backup time. Regardless of the amount of application data stored, the

BridgeHead backup software performs a full block-level backup of the entire LUN for each MEDITECH host.

Calculate the sequential read throughput that is required to complete the backup within this window as a function of the factors involved:

- The desired backup duration
- The number of LUNs
- The size of each LUN to be backed up

For example, in a 50-host MEDITECH environment in which each host's LUN size is 200GB, the total LUN capacity to back up is 10TB.

To back up 10TB of data in 8 hours, the following throughput is required:

$$\begin{aligned} &= (10 \times 10^6) \text{MB} \div (8 \times 3,600) \text{s} \\ &= 347.2 \text{MBps} \end{aligned}$$

However, to account for unplanned events, a conservative backup window of 5.5 hours is selected to provide headroom beyond the 6 hours recommended.

To back up 10TB of data in 5.5 hours, the following throughput is required:

$$\begin{aligned} &= (10 \times 10^6) \text{MB} \div (5.5 \times 3,600) \text{s} \\ &= 500 \text{MBps} \end{aligned}$$

At the throughput rate of approximately 500MBps, the backup can complete within a 5.5-hour time frame, comfortably within the BridgeHead 8-hour backup requirement.

Table 3 summarizes the I/O characteristics of the BridgeHead workload to use when you size the storage system.

Table 3) Summary of BridgeHead workload I/O characteristics and requirements.

Parameter	All Platforms
Request size	64K
Random/sequential	100% sequential
Read/write ratio	100% read
Average throughput	Depends on the number of MEDITECH hosts and the size of each LUN; backup must complete within 8 hours.
Required backup duration	8 hours

3 Sizing NetApp Storage for MEDITECH and BridgeHead Workloads

3.1 Sizing Methodology for NetApp Storage Systems

NetApp field engineers and partners use the SPM application to determine the storage system specifications required to satisfy workload needs for each customer. The NetApp SPM tool uses storage models that take into account NetApp technologies, statistical field data, and workload details to generate storage sizing recommendations. This is the best practice method for determining the type of storage controller, the amount of NetApp Flash Pool intelligent caching, and the number of disks required for a specific workload or for combined workloads. This method also provides estimates for system utilization metrics such as disk-to-host operations ratio and system utilization level.

For more information about NetApp SPM, see [NetApp TR-4050: System Performance Modeler](#). You can also access the SPM tutorial from SPM > Help.

3.2 Sizing for MEDITECH Production and BridgeHead Backup Workloads

SPM uses input from a variety of workload and storage system parameters to produce storage configuration recommendations that meet the storage capacity and performance requirements of MEDITECH and BridgeHead workloads.

Specified characterizations of the MEDITECH and BridgeHead workloads are entered into the SPM sizing tool as input values. Based on these values, SPM generates storage configuration recommendations that meet the requirements of these workloads. The MEDITECH platform, its category, and the number of hosts to be deployed are key factors in determining inputs to the NetApp SPM for the combined MEDITECH production and BridgeHead backup workloads.

Section 3.3, “Sizing Parameters for NetApp SPM Sizing Tool,” describes the parameters used to size the NetApp storage for these workloads.

NetApp Flash Solutions

The MEDITECH read and write latency requirements listed in Table 2 indicate that a high number of disks will probably be required to provide adequate read and write performance. Given this type of workload, NetApp requires a caching solution to be deployed in all storage systems that support MEDITECH environments. The use of a NetApp caching solution can help significantly reduce the number of spinning disks that are required in MEDITECH production environments.

NetApp offers two caching solutions for FAS storage controllers:

- **NetApp Flash Cache™**, which is used mainly for high random read workloads
- **Flash Pool**, which is used for both read-intensive and write-intensive workloads

Because MEDITECH 6.x environments are highly write intensive, NetApp recommends that you use hybrid storage with Flash Pool or NetApp AFF solutions. The next two sections discuss how Flash Pool capacity is determined on hybrid FAS storage and on NetApp AFF storage for MEDITECH environments.

NetApp recommends that you use the Flash Pool caching solution to meet MEDITECH performance requirements. For more information, see [NetApp TR-4070: Flash Pool Design and Implementation Guide](#).

Flash Pool Capacity

To provide a significant benefit for read and write performance, the Flash Pool capacity should be large enough to accommodate the entire working set of the MEDITECH production workload. The amount of data that is actively accessed at any given time (the working set size) per MEDITECH host is expected to be smaller than the size of the entire LUN.

MEDITECH Working Set Size

MEDITECH does not specify the exact size of the working set. Therefore, for the purpose of sizing the storage, NetApp estimates that it is 10% of each MEDITECH LUN.

To determine Flash Pool capacity, NetApp recommends that you specify the working set size multiplied by two. The extra capacity offers space for additional operational headroom. Furthermore, as the size of the MEDITECH application data grows over time, the amount of Flash Pool capacity can be increased as needed.

For example, in a 50-host MEDITECH environment, with a host LUN size of 200GB, NetApp recommends that you specify a minimum Flash Pool capacity of 2TB (50 MEDITECH hosts × 200GB × 10% × 2).

Flash Pool caching consists of solid-state drives (SSDs) in a NetApp RAID DP® configuration. The proper selection of SSD capacity is used to specify the Flash Pool capacity as close to the target capacity as possible. SPM is used to determine the usable Flash Pool capacity given the number of SSDs used and the SSD per disk capacity specified. The usable Flash Pool capacity is the maximum capacity that the NetApp Data ONTAP operating system uses for data caching. The usable Flash Pool capacity is less than the effective storage capacity of an SSD RAID group.

For example, in a 60-host MEDITECH environment with a LUN size of 200GB, NetApp recommends that you establish a minimum usable Flash Pool capacity of 2.4TB (60 MEDITECH hosts × 200GB × 10% × 2). To size for at least 2.4TB of usable Flash Pool capacity (assuming that 200GB SSD drives are used), SPM determines that 9 × 400GB SSD data drives (3600GB effective RAID capacity) are required to provision 2.4TB of usable Flash Pool capacity.

For more information about NetApp Flash Pool intelligent caching, see [NetApp TR-4070: Flash Pool Design and Implementation Guide](#) and [Flash Pool Technical FAQ](#).

The NetApp Flash Pool disk-partitioning feature is implemented in Data ONTAP 8.3 and later. Taking advantage of Flash Pool disk partitioning helps to reduce the number of SSD parity disks required and to optimize their use.

NetApp All Flash FAS

NetApp offers high-performance AFF arrays to address MEDITECH workloads that demand high throughput and that have random data access patterns and low latency requirements. For MEDITECH workloads, AFF arrays offer performance advantages over systems based on hard disk drives (HDDs). The combination of flash technology and enterprise data management delivers advantages in three major areas: performance, availability, and storage efficiency.

Disk Type

The low read-latency requirement of the MEDITECH production workload dictates the use of a 10K or 15K RPM SAS disk type. Because the 15K RPM SAS disk type reached its end of availability (EOA) in November 2013, new deployments use the 10K RPM SAS disk type.

For AFF systems, NetApp recommends that you use SSDs of 400G or higher because the 200GB SSD disk type has reached its EOA.

High-Availability Controller Pairs

NetApp FAS and AFF storage systems are deployed with two or more storage controllers configured as active-active high-availability (HA) pairs. This configuration enables you to spread the MEDITECH production and BridgeHead backup workloads across all storage controllers. Each FAS storage controller is installed with its own CPU and Flash Pool. The use of one or more storage controller pairs makes possible more total CPU and Flash Pool capacity than a single storage controller can provide.

Furthermore, the HA configuration allows storage service continuity during a storage controller failure. If a storage controller fails, its partner controller takes over the access and management of the disks owned by the failed storage controller. This configuration provides continuous storage access to external applications.

Aggregates

NetApp recommends that you place the combined MEDITECH production and BridgeHead backup workloads on a dedicated aggregate on each storage controller of an HA pair. Place any other workloads on separate aggregates.

Important

BridgeHead does not require additional capacity for its workload.

BridgeHead leverages FlexClone to clone the volumes that host MEDITECH production LUNs and to back up all LUNs from the FlexClone volumes.

For optimal performance, determine which disks the sizing tool recommends for the BridgeHead workload and add these disks to the MEDITECH production aggregates.

3.3 Sizing Parameters for NetApp SPM Sizing Tool

This section describes the SPM parameters that are used for proper sizing of NetApp storage in MEDITECH production environments. When you size for MEDITECH production environments, you must define the following items in the SPM tool:

- MEDITECH workload (defined as an SPM custom application)
- BridgeHead workload (defined as an SPM custom application)
- Storage hardware configuration

Note: This document is based on SPM 2.2. Subsequent SPM releases might change the storage sizing workflow, the GUI, or the SPM parameter terms used. NetApp recommends that you review the latest version of the SPM user tutorial to see the most current storage sizing workflow. Also review the GUI presentation of input parameters to define the workloads and storage hardware configuration. The “SPM User Guide” is available in the Help menu of the SPM application (Help > SPM Help).

The following two sections describe the fixed and variable sizing parameters of these SPM items.

Fixed SPM Sizing Parameters

Table 4 through Table 6 list the fixed NetApp SPM parameters that are used to size NetApp storage in MEDITECH environments. Use these fixed parameters as the basis for all NetApp storage sizing to determine a storage configuration that satisfies the MEDITECH production and BridgeHead backup workloads.

Fixed SPM Sizing Parameters for MEDITECH Production Workload

Table 4 describes the fixed sizing parameters for the MEDITECH production workload.

Table 4) Fixed sizing parameters: assumptions for the MEDITECH production workload.

SPM Parameter	Value	Comments
Basic Inputs		
Protocol	FCP	Integration requirement as provided by MEDITECH.
Throughput type	IOPS	Requirement as provided by MEDITECH.
Latency		
Random read latency (ms)	5	Requirement as provided by MEDITECH.
Random write latency (ms)	1	During the backup window and during reallocations, 2ms write latency is acceptable.

SPM Parameter	Value	Comments
I/O Percent		
Sequential read and write (%)	0	MEDITECH production workloads are 100% random.
Working Set Size		
Active working set (%)	10	Without definitive information from MEDITECH, NetApp estimates the working set size to be 10% of the required capacity.
I/O Sizes		
Sequential read and write size (KB)	IO-Size-64K	Not applicable because MEDITECH production workloads are 100% random.
Layout Hints		
Can this workload be placed on a shared aggregate?	Yes	The MEDITECH production workload is placed on an aggregate that is shared only with the BridgeHead backup workload.
Can the workload be split on different aggregates across the cluster?	Yes	Default value.
Can this workload be placed on a Flash Pool aggregate?	Yes	Applies only if Flash Pool is chosen as the caching solution. Clear the Flash Cache checkbox. (Not applicable for AFF systems.) Note: This option might not be available until the first SPM sizing evaluation is attempted.
Flash Pool		
Random overwrite (%)	100	Applicable only if Flash Pool is chosen as the caching solution. During validation testing in the performance validation lab, a 100% overwrite rate was observed with the MEDITECH workload simulator. Note: This random overwrite value might differ from that of the production workload. Note: This option might not be available until the first SPM sizing evaluation is attempted. Note: This option is not applicable for AFF systems.

Fixed SPM Sizing Parameters for BridgeHead Workload

Table 5 describes the fixed sizing parameters for the BridgeHead workload.

Table 5) Fixed sizing parameters: assumptions for the BridgeHead workload.

SPM Parameter	Value	Comments
Basic Inputs		
Protocol	FCP	Requirement as provided by BridgeHead.

SPM Parameter	Value	Comments
Throughput type	MBps	Requirement as provided by BridgeHead.
Required capacity (TB)	0.001	Because the BridgeHead workload shares the same data aggregates as the MEDITECH workload, the required capacity should be set at 0TB. Because SPM requires the required capacity value to be larger than 0TB, NetApp recommends that you set the value to 0.001TB.
Latency		
Random read latency (ms)	20	Default value.
I/O Percent		
Random read and write (%)	0	Not applicable for backup operations. Note: The BridgeHead backup workload is 100% sequential.
Sequential read (%)	100	The BridgeHead backup workload is 100% sequential read.
Sequential write (%)	0	Not applicable for backup operations. Note: The BridgeHead backup workload is 100% sequential.
Working Set Size		
Active working set (%)	0	The SPM working set size is defined by random read and write operations. This workload is 100% sequential read.
I/O Sizes		
Random read and write size (KB)	IO-Size-4K	Not applicable for backup operations Note: The BridgeHead backup workload is 100% sequential.
Sequential read size (KB)	IO-Size-64K	BridgeHead backup uses a sequential read size of 64K.
Sequential write size (KB)	IO-Size-64K	Not applicable for backup operations Note: The BridgeHead backup workload is 100% sequential.
Layout Hints		
Can this workload be placed on a shared aggregate?	Yes	The MEDITECH production workload is placed on an aggregate that is shared only with the BridgeHead backup workload.
Can the workload be split on different aggregates across the cluster?	Yes	Default value.
Can this workload be placed on a Flash Pool aggregate?	Yes	Applicable only if Flash Pool is chosen as the caching solution Note: This option might not be available until the first SPM

SPM Parameter	Value	Comments
		sizing evaluation is attempted.
How will the workload traffic arrive in the cluster?	On a single node	Default option. Note: This option might not be available until the first SPM sizing evaluation is attempted.
Flash Pool		
Random overwrite (%)	0	Applicable only if Flash Pool is chosen as the caching solution. The BridgeHead backup workload is 100% sequential. Note: This random overwrite value might differ from that of the production workload. Note: This option might not be available until the first SPM sizing evaluation is attempted. Note: Not applicable for AFF systems.

Fixed SPM Sizing Parameters for Storage Hardware Configuration

Table 6 describes the fixed sizing parameters for the storage hardware configuration.

Table 6) Fixed sizing parameters: assumptions for the storage hardware configuration.

SPM Parameter	Value	Comments
HA pair	Yes	Required to enable resiliency and high availability if one storage controller fails
System headroom (%)	30	Default value
Spare disks per node	2	Default value

Variable SPM Sizing Parameters

Table 7 through Table 9 describe the parameters that might vary according to the size and hardware used for a MEDITECH deployment. These include NetApp storage controller types and the total LUN capacity required for the MEDITECH environment.

Variable SPM Sizing Parameters for MEDITECH Production Workload

Table 7 describes the variable NetApp SPM input parameters for the MEDITECH production workload.

Table 7) Variable sizing parameters: assumptions for the MEDITECH production workload.

SPM Parameter	Value	Comments
Basic Inputs		
Protocol	FCP or iSCSI	Select the iSCSI protocol for FAS2240HA. Select FCP for any other FAS models. Note: MEDITECH and BridgeHead specify FCP as the preferred protocol.

SPM Parameter	Value	Comments
Throughput	Depends on the MEDITECH category, platform, and number of MEDITECH hosts deployed	The throughput per MEDITECH host depends on the MEDITECH category and platform. For example, for a 50-host MEDITECH (category 2) 6.x environment, the throughput value is 37,500 IOPS (50 × 750 IOPS). For more information, see Table 1.
Required capacity (TB)	Depends on the number of MEDITECH hosts and the LUN size per host	The required capacity is the sum of the provisioned LUN capacity for each MEDITECH host. For example, for a 50-host MEDITECH environment, given a 200GB LUN size per host, the required capacity value is 10TB (50 × 200GB).
Latency		
Random read latency (ms)	5 or 7	Requirement as provided by MEDITECH. Note: Select 7ms for MEDITECH category 1 deployment and 5ms for MEDITECH category 2 through 6 deployment.
I/O Percent		
Random read and random write	Depends on the MEDITECH platform	As per the workload description provided by MEDITECH: <ul style="list-style-type: none"> MEDITECH 6.x: 80% random write and 20% random read C/S 5.x: 60% random write and 40% random read MAGIC: 10% random write and 90% random read For more information, see Table 2.
I/O Sizes		
Random read size (KB)	IO-SIZE-4K / 8K / 16K	Requirement as provided by MEDITECH: <ul style="list-style-type: none"> MEDITECH 6.x: Random read size is 4K. C/S 5.x: Random read size is 4K. MAGIC: Random read size depends on the MEDITECH deployment, which can be either 8K or 16K. Consult MEDITECH to determine whether the site requires the 8K request size or the 16K request size. For more information, see Table 2.
Random write size (KB)	IO-SIZE-4K / 8K / 16K	Requirement as provided by MEDITECH: <ul style="list-style-type: none"> MEDITECH 6.x: Random write size is 4K. C/S 5.x: Random write size is 4K. MAGIC: Random read size depends on the MEDITECH deployment size, which can be either 8K or 16K. Consult MEDITECH to determine whether the site requires the 8K request size or the 16K request size. For more information, see Table 2.

Note: To determine the LUN sizes required for a MEDITECH production environment, consult the MEDITECH document “Hardware Configuration Proposal” (for a new deployment) or “Hardware Evaluation Task” (for an existing deployment).

Variable SPM Sizing Parameters for BridgeHead Workload

Table 8 describes the variable SPM sizing input parameters for the BridgeHead workload.

Table 8) Variable sizing parameters: assumptions for the BridgeHead workload.

SPM Parameter	Value	Comments
Basic Inputs		
Protocol	FCP or iSCSI	Select the iSCSI protocol for FAS2240HA and FCP for any other FAS models. Note: MEDITECH and BridgeHead specify FCP as the preferred protocol.
Throughput	Depends on the number of MEDITECH hosts and the LUN size per host	BridgeHead requires the backup of all MEDITECH LUNs within an 8-hour backup window. However, NetApp recommends a conservative backup window of less than 6 hours. For example, in a 50-host MEDITECH environment with 200GB LUNs per host, the total data to be backed up is 10TB (50 MEDITECH hosts × 200GB/host). To back up 10TB of data in 5.5 hours (conservatively, in under 6 hours), the following throughput is required: = $(10 \times 10^6) \text{MB} \div (5.5 \times 3600) \text{s}$ = 500MBps
Required capacity (TB)	Depends on the number of MEDITECH hosts and the LUN size per host	See the corresponding parameter in Table 7.

Variable SPM Sizing Parameters for Storage Hardware Configuration

Table 9 describes the variable SPM sizing input parameters for the storage hardware configuration.

Table 9) Variable sizing parameters: assumptions for the storage hardware configuration.

SPM Parameter	Value	Comments
Deployment mode	Clustered Data ONTAP or 7-Mode	Select the target Data ONTAP mode.
Data ONTAP version	Data ONTAP 7-Mode (version 8.1.x or later) or clustered Data ONTAP (version 8.2.x or later)	Versions validated with MEDITECH: <ul style="list-style-type: none"> Data ONTAP 7-Mode, version 8.1.2 Clustered Data ONTAP, versions 8.2.x and 8.3.x For more information, see Table 2.
Controller platform	<ul style="list-style-type: none"> FAS2240HA FAS2552HA FAS32XXHA FAS62XXHA FAS80XXHA AFF80XXHA 	Specify the type of storage controller. For a list of platforms that have been tested with the MEDITECH and BridgeHead workloads, see section 6, "MEDITECH Certified and Validated SPM Storage Configurations." The products listed there meet the MEDITECH performance validation requirements. The AFF80XXHA, FAS80XXHA, and FAS62XX storage controllers can be used for large MEDITECH environments. The controllers can also be used for shared workload environments that require the greater capacity

SPM Parameter	Value	Comments
		and/or performance of the FAS80XX, AFF80XX, and FAS62XX families. For more information, see section 6.1, “Certified and Validated SPM Storage Configurations.”
Degraded performance OK on HA takeover event	Depends on the customer’s need	This is a checkbox option. The default position is for the checkbox to be selected, which means that some degree of performance degradation is acceptable if a storage controller in an HA pair is down for a short time. The storage HA pair operates with degraded performance until the failed HA partner is brought back up. If this checkbox is deselected, storage performance is maintained while a storage controller in an HA pair is down. SPM reserves 50% of system resources to guarantee the performance. SPM might recommend additional storage controller HA pairs because of the 50% system resource reservation. Note: When you configure this option, it is important to consider the customer’s requirements.
Disk type	SAS-10K, SSD	Required to provide support for environments with high IOPS levels and low read latencies.
Flash acceleration options	Flash Pool	Flash Pool intelligent caching is required if you deploy MEDITECH environments with challenging read and write performance targets. For more information, see NetApp TR-4070: Flash Pool Deployment and Integration Guide . If you select Flash Pool as the caching solution, select the Flash Pool option and deselect the Flash Cache and None options.
HA pair	Selected	Default option. Note: Storage controllers are configured in an HA configuration.
Default disk shelf	DS2246 or DS4243	Select the type of disk shelf. Currently, only the DS2246 and DS4243 disk shelves have been tested with the MEDITECH and BridgeHead workloads. These shelves meet the MEDITECH performance validation requirements. For more information, see section 6.1, “Certified and Validated SPM Storage Configurations.”
Default drive type	SAS-10K, SSD	Select the type of disk. Currently, the SAS-10K and SSD disk types have been tested with the MEDITECH and BridgeHead workloads. These disk types meet the MEDITECH performance validation requirements. For more information, see section 6.1, “Certified and Validated SPM Storage Configurations.”
Flash Pool		
SSD disk capacity	Select a disk capacity.	Applies only if Flash Pool is chosen as the caching solution.

SPM Parameter	Value	Comments
		Because the required Flash Pool capacity can be calculated, NetApp does not recommend that you use the <code>AUTO_SUGGEST</code> option.
Flash Pool capacity per controller (GB)	Depends on the MEDITECH working set size	<p>Applies only if Flash Pool is chosen as the caching solution.</p> <p>It is important to allocate sufficient usable Flash Pool capacity for any MEDITECH environment. The MEDITECH working set size is estimated to be 10% of the total LUN size used by the MEDITECH hosts. NetApp recommends that you allocate an amount of Flash Pool capacity that is greater than or equal to the working set size multiplied by 2.</p> <p>For example, for a 60-host MEDITECH environment, given 200GB of LUN storage per host, NetApp recommends that you allocate a minimum cache capacity of 60 MEDITECH hosts × 200GB/host × 10% × 2 = 2.4TB.</p> <p>You can size a total of 2.9TB of usable Flash Pool capacity by using 12 × 400GB SSD drives (1488GB usable Flash Pool capacity with 2400GB SSD capacity [RAID DP] per controller).</p>

Note: To determine the LUN sizes required for your specific MEDITECH production environment, consult the MEDITECH document “Hardware Configuration Proposal” (for a new deployment) or “Hardware Evaluation Task” (for an existing deployment).

3.4 Using SPM Sizing Results

SPM might divide the MEDITECH production and BridgeHead backup workloads by placing the MEDITECH workload on one storage controller and the BridgeHead backup workload on another. This division has the potential to create an uneven distribution of data disks and an uneven use of resources on the storage controllers.

The MEDITECH and BridgeHead workloads share the same aggregates across all storage controllers (random read and write by MEDITECH and sequential read by BridgeHead). Therefore, the combined MEDITECH and BridgeHead workloads can be spread evenly across all storage controllers. To do so, make sure the capacity of the aggregates for the MEDITECH hosts is distributed evenly across all storage controllers. The following assumptions apply:

- The data disk/aggregate ratio is the same across all data aggregates dedicated to the MEDITECH hosts.
- The IOPS per MEDITECH hosts adhere to the MEDITECH platform specifications listed in Table 2.

3.5 Additional Sizing Requirements

In addition to following the storage configuration recommendations from SPM, you must also satisfy the following sizing requirements to properly size the NetApp FAS storage system for the combined MEDITECH host and BridgeHead backup workloads.

Storage System with Flash Pool

When you size a system that has Flash Pool capabilities, make sure to meet the requirements stated in the “Flash Pool Guidelines” section of the NetApp SPM sizing report. The SPM-recommended number of SAS drives might not meet these requirements. One requirement states:

“For Flash Pool sizings resulting in greater than 75% reduction in HDD count compared to the configuration without Flash Pool, it is recommended that the HDD count is retained at 75% reduction when quoting to customers. This limit provides a buffer in situations where all the IOPS need to be served from HDDs while the SSD cache is still warming up, during boot up.” (SPM 1.4.2(P1) Sizing Report).

Note: To satisfy these requirements, it might be necessary to add more SAS drives to the number of data drives recommended by SPM. To understand how adding another SAS drive might help you to size properly for NetApp storage with Flash Pool, see section 4.2, “All Flash FAS Sizing Example for 60-Host MEDITECH (Category 2) 6.x and BridgeHead Workloads with Clustered Data ONTAP 8.3.1.”

Minimum Required Storage Capacity

The minimum recommended storage capacity for the MEDITECH hosts must be equal to or larger than the sum of storage required by each MEDITECH host multiplied by 1.5. The additional 50% of storage is provisioned to anticipate the additional storage consumption of the Snapshot copies created by the BridgeHead backup application and the change in Snapshot copy storage consumptions that results from traditional volume reallocate operations.

For example, a sizing of 4TB required capacity for the MEDITECH workload requires a minimum capacity of 6TB. If the number of data disks specified in the SPM report creates a capacity of less than 6TB, additional data disks must be added so that the final storage capacity is at least 6TB.

Note: The total capacity of LUNs for the MEDITECH hosts remains the same at 4TB. However, the total capacity of the NetApp volume that contains the LUNs is 6TB.

For more information, see [NetApp TR-4300: NetApp FAS Storage Systems for MEDITECH Environments Best Practices Guide](#).

Minimum and Maximum RAID Group Sizes

The minimum and maximum RAID group sizes depend on the type of data disk and data network used. For recommendations about minimum and maximum RAID group sizes, see [NetApp TR-3838: Storage Subsystem Configuration Guide](#).

Storage System Root Volumes and Spare Disks

The storage controller root volumes require a small amount of additional storage for Data ONTAP. Also, some disks should be set aside in the NetApp storage shelves as spare disks. For the best practices to determine the number of spare disks to use, see [NetApp TR-3838: Storage Subsystem Configuration Guide](#).

3.6 Additional Storage for MEDITECH Environments

A system in a MEDITECH environment might also contain the following types of storage:

- VMware vSphere datastore for MEDITECH host system disks (VMDKs)
- Other MEDITECH application workloads (for example, data repository, scanning, and archiving)

This storage should be sized on a separate aggregate in keeping with the workload I/O characteristics and storage capacity required for the customer's environment.

4 Sizing Examples

In this section, two examples show how to size storage with Flash Pool and AFF systems by following the SPM guidelines presented in section 3, “Sizing NetApp Storage for MEDITECH and BridgeHead Workloads.”

These examples are representative of the following two use cases:

- Section 4.1 describes a use case for customers who deploy new MEDITECH systems with NetApp hybrid storage. NetApp recommends that you deploy new MEDITECH systems with clustered Data ONTAP 8.2.1 or later and that you use Flash Pool caching.
- Section 4.2 describes a use case for customers who deploy new MEDITECH systems with NetApp AFF storage. To get the full benefit of AFF, NetApp recommends that you deploy new MEDITECH systems with clustered Data ONTAP version 8.3 or later.

4.1 Sizing Example for 20-Host MEDITECH (Category 2) 6.x and BridgeHead Workloads with Clustered Data ONTAP 8.3 and Flash Pool

The following sizing example shows how the NetApp storage configuration is determined for a BridgeHead and MEDITECH (category 2) 6.x environment that uses Flash Pool intelligent caching. You can use Flash Pool caching with clustered Data ONTAP 8.3 to satisfy the MEDITECH (category 2) 6.x platform's requirement of 5ms random read latency. It is important to provision sufficient Flash Pool capacity to support the MEDITECH working set size. In addition to the storage provisioned for the MEDITECH hosts, more disks are also allocated for the Data ONTAP root aggregate and spare disk (per controller).

The sizing process involves the following tasks:

1. Collect deployment information.
2. Use the NetApp SPM tool to determine the number of data disks required to satisfy the capacity and performance needs of the combined MEDITECH production and BridgeHead backup workloads.
3. Select the RAID group size that is appropriate for ease of deployment and data growth. For recommended RAID group sizes, see [NetApp TR-3838: Storage Subsystem Configuration Guide](#).
4. Determine the total number of disks required on each storage controller, taking into consideration the root aggregate and the number of spare disks.
5. Verify that the number of data disks meets the sizing requirements listed in section 3.5, "Additional Sizing Requirements."

Note: Storage for other workloads such as virtual machine infrastructure is considered separately from the MEDITECH production and BridgeHead backup workloads.

Step 1: Collect Deployment Information

This example uses FAS8040 storage controllers configured as an HA pair and Flash Pool aggregates. In addition, 10K RPM 600GB SAS disks are chosen as the data drives with the DS2246 disk shelves. Different disk types might require different disk shelf types.

Table 10 lists the deployment information for this example.

Table 10) Deployment information for 20-host MEDITECH (category 2) 6.x environment with Flash Pool sizing.

Item	Comments
NetApp storage system	FAS8040 in HA configuration
Data ONTAP version	Clustered Data ONTAP 8.3
Caching solution	Flash Pool is selected as the caching solution
Disk type	NetApp 10K RPM SAS 600GB drives
Disk shelf	NetApp DS2246
Data network protocol	FCP

Item	Comments
Performance expectation if one storage controller is down	As stated in the deployment requirement, some degree of performance degradation is acceptable if a storage controller is down for a short time
MEDITECH category	2
MEDITECH platform	6.x
MEDITECH host count	20
MEDITECH host LUN size	200GB
Backup software	BridgeHead
Backup time window	Less than 8 hours
Data growth in the next 2 years	50%
Legacy data	None; this is a new deployment

Step 2: Use SPM to Determine Number of Disks Required for MEDITECH and BridgeHead Backup Workloads

The next step is to translate the deployment information from Table 10 into input for the SPM sizing tool. The tables in the following sections list the SPM parameters for the MEDITECH and BridgeHead backup workload and the FAS8040 hardware configuration used in this example.

To create the new workload, log in to the [SPM](#) tool and click the Workload tab, which is shown in Figure 1.

Figure 1) Screenshot of SPM custom application for MEDITECH (category 2) 6.x workload with 20 hosts.

Workload Title: MEDITECH 6.x with 20 hosts

Workload Type: Custom

▼ Basic Inputs (Required) ?

Capacity (TiB): 4

Throughput: 15000 IOPS

▼ Detailed Inputs ?

High Impact

Random Read Latency (ms): 5

IO Percent

Workload IO Type: ☒ Only Random ☐ Mixed ☐ Only Sequential

Random Read (%): 20

Random Write (%): 80

Sequential Read (%): 0

Sequential Write (%): 0

Working Set Size

Active Working Set (%): 10

Low Impact

Protocol: FCP

IO Sizes

Random Read Size (KB): 4KB

Random Write Size (KB): 4KB

Sequential Read Size (KB): 64KB

Sequential Write Size (KB): 64KB

Flash Pool

Random Overwrite (%): 100


Random Overwrite is a sub component of Random Write(%) for workloads where Flash Pool is used.

Layout Hints

Can this workload be placed on a shared aggregate? ☒

Can this workload be split on different aggregates across the cluster? ☒

How will the workload traffic arrive into the cluster? ☒ On a single node ☐ All nodes in the cluster



Data Protection

Enable SnapMirror: ☒ Yes ☐ No

The SPM tool parameters listed in Table 11 define the MEDITECH workload for a custom application.

Table 11) Example of MEDITECH workload SPM parameter values for 20-host MEDITECH (category 2) 6.x environment.

SPM Parameter	Value	Comments
Identifiers		
Workload title	MEDITECH (category 2) 6.x with 20 hosts	Workload name.
Workload type	Custom	Select the custom workload type for the MEDITECH workload.
Basic Inputs (Required)		
Capacity (TB)	4	In this example, the LUN size per MEDITECH host is 200GB, and 20 hosts must be deployed. Therefore, the total storage capacity is 20 × 200GB = 4TB.
Throughput	15,000	Because there are 20 MEDITECH (category 2) 6.x hosts and each host averages 750 IOPS, the total IOPS is 21,000. For more information, see section 2.1, “MEDITECH Workload Description.”
Throughput type	IOPS	Throughput in IOPS instead of MBps.
Detailed Inputs		
Random read latency (ms)	5	MEDITECH workload requirement as described in section 2.1, “MEDITECH Workload Description.”
I/O Percent		
Workload I/O type	Only random	MEDITECH workloads are 100% random.
Random read (%)	20	MEDITECH workload requirement as described in section 2.1, “MEDITECH Workload Description.”
Random write (%)	80	MEDITECH workload requirement as described in section 2.1, “MEDITECH Workload Description.”
Working Set Size		
Active working set (%)	10	MEDITECH workload requirement as described in section 2.1, “MEDITECH Workload Description.”
Protocol	FCP	Deployment requirement.
I/O Sizes		
Random read and write size (KB)	IO-SIZE-4K	MEDITECH workload requirement as described in section 2.1, “MEDITECH Workload Description.”

SPM Parameter	Value	Comments
Sequential read and write size (KB)	N/A	Default value. Note: Sequential read and write size does not apply to the MEDITECH workload because this workload is 100% random.
Flash Pool		
Random overwrite (%)	100	Flash Pool overwrite is always 100% for MEDITECH workloads. Note: The Flash Pool random overwrite option does not appear when a MEDITECH workload is defined. The option appears only after Flash Pool is selected under Hardware.
Layout Hints		
Can this workload be placed on a shared aggregate?	Selected	Default value. Note: The MEDITECH workload shares its aggregate with the BridgeHead backup software.
Can the workload be split on different aggregates across the cluster?	Selected	Default value. Note: The MEDITECH workload is split evenly between two storage controllers as specified in the deployment specification in Table 10.
How will the workload traffic arrive into the cluster?	On a single node	Default option. Note: This option might not be available until the first SPM sizing evaluation is attempted.
Data Protection		
Enable NetApp SnapMirror® software	Selected	Click Yes to enable NetApp SnapMirror replication technology.

Next, the BridgeHead workload is defined through SPM. BridgeHead software is used to manage and perform nightly backup on site. The backup window is set to a maximum of 8 hours. Conservatively, a maximum backup time of 6 hours is used to calculate the amount of throughput required. For the calculation, see the corresponding SPM parameter for throughput.

Figure 2 shows the parameters for the BridgeHead workload.

Figure 2) Screenshot of SPM custom application for BridgeHead workload.

Workload Title

Workload Type

▼ Basic Inputs (Required) ?

Capacity (TiB)

Throughput

▼ Detailed Inputs ?

High Impact

Random Read Latency (ms)

IO Percent

Workload IO Type ☐ Only Random ☐ Mixed ☒ Only Sequential

Random Read (%)

Random Write (%)

Sequential Read (%)

Sequential Write (%)

Working Set Size

Active Working Set (%)

Low Impact

Protocol

IO Sizes

Random Read Size (KB)

Random Write Size (KB)

Sequential Read Size (KB)

Sequential Write Size (KB)

Flash Pool

Random Overwrite (%)

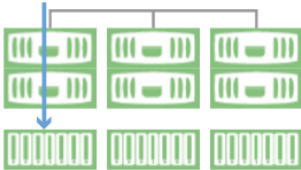
Random Overwrite is a sub component of Random Write(%) for workloads where Flash Pool is used.

Layout Hints

Can this workload be placed on a shared aggregate? ☒

Can this workload be split on different aggregates across the cluster? ☒

How will the workload traffic arrive into the cluster? ☒ On a single node ☐ All nodes in the cluster



Data Protection

Enable SnapMirror ☐ Yes ☒ No

When you use the SPM tool to create a custom application, use Table 12 to define the BridgeHead workload.

Table 12) Example of BridgeHead workload SPM parameter values for 20-host MEDITECH environment.

SPM Parameter	Value	Comments
Identifiers		
Workload title	BridgeHead workload for 20-host MEDITECH (category 2) 6.x environment	Workload name.
Workload type	Custom	Select the custom workload type for the BridgeHead workload.
Basic Inputs (Required)		
Capacity (TB)	4	Same as MEDITECH workload capacity.
Throughput	234	Throughput is specified in MBps. Total data to back up is 4TB (20 MEDITECH hosts × 200GB per host). Total time allocated for backup is 5 hours or 18,000 seconds (5 hours × 3,600 seconds/hour). Throughput required is 234 (4 × 1024 ² MB ÷ 18,000 seconds).
Throughput type	MBps	Throughput in MBps instead of IOPS.
Detailed Inputs		
Random read latency (ms)	20	Default value. Note: Random read latency does not apply to the BridgeHead workload because this workload is 100% sequential read.
I/O Percent		
Workload I/O type	Only sequential	BridgeHead workload is always 100% sequential read.
Sequential read (%)	100	BridgeHead workload is 100% sequential read.
Sequential write (%)	0	BridgeHead workload is 100% sequential read.
Working Set Size		
Active working set (%)	10	Same as MEDITECH working set size.
Protocol	FCP	Deployment requirement.
I/O Sizes		
Random read and write size (KB)	Not applicable (N/A)	N/A

SPM Parameter	Value	Comments
Sequential read and write size (KB)	64KB	BridgeHead sequential read size is 64KB.
Flash Pool		
Random overwrite (%)	0	There are no writes for the BridgeHead workload. It is a backup-only workload.
Layout Hints		
Can this workload be placed on a shared aggregate?	Selected	Default value. Note: The MEDITECH workload shares its aggregate with the BridgeHead backup software.
Can the workload be split on different aggregates across the cluster?	Selected	Default value. Note: The BridgeHead workload is split evenly between two storage controllers.
How will the workload traffic arrive into the cluster?	On a single node	Default selection. Note: This option might not be available until the first SPM sizing evaluation is attempted.
Data Protection		
Enable SnapMirror	No	Click No to deselect SnapMirror.

Figure 3 shows the SPM parameters for the storage configuration.

Figure 3) Screenshot of SPM storage configuration.

☐ Let SPM autosuggest hardware for me.

Basic Inputs (Required) ⓘ

[Note: For sizing EOA hardware, use Performance Estimation (Reverse Sizing)]

Deployment Options ☐ All-Flash FAS (AFF) ☐ MetroCluster

Data ONTAP ☒ Clustered ONTAP ☐ 7-Mode

8.2.x

Controller Platform ☐ Entry ☒ Mid ☐ High

Hold Ctrl to select multiple Controller Platforms.

FAS2554 FAS2552 FAS2550	FAS8040 FAS8020	FAS8080 FAS8060
-------------------------------	--------------------	--------------------

Degraded Performance OK on HA takeover event ☒

Disk Type ☐ SAS ☐ SATA ☐ SSD

Flash Acceleration Options ☐ None ☐ FlashCache ☒ FlashPool

☐ Enable Storage Pool Sharing

System Headroom (%) 30

Capacity Reserve (%) 50

Workload to Shelf/Drive Type Mapping ⓘ

☐ Enable Drive Type Selection per workload

Disk Shelf for all workload(s) DS2246

Drive Type for all workload(s)

SAS-10K-900GB
SAS-10K-600GB
SAS-10K-1.8TB
SAS-10K-1.2TB
NETAPP-SSD-800GB

Maximum 2 drive types are allowed to size per sizing. Please select any 2 drive types from the list.

Detailed Inputs ⓘ

Is 2-Node MetroCluster? ☐

HA Pair ☒

FlashCache

Card Type AUTO-SUGGEST

Cards per Controller 1

Card Capacity Select flash card capacity...

Flash Pool

Flash Pool Disk Capacity AUTO-SUGGEST

Flash Pool Capacity per controller (GB) 800

[Flash Pool limits](#)

Other Inputs

Map to Full Shelves ☐

Spare Disks per Drive Type 2

System Age Empty System

Use Table 13 to define the “Storage Hardware Configuration” section in the SPM tool.

Table 13) Example of storage hardware configuration SPM parameter values for 20-host MEDITECH environment with FAS8040 and Flash Pool.


SPM Parameter	Value	Comments
Manual Selection Option		
Let SPM autosuggest hardware for me.	Deselect the checkbox	Deselecting this option enables you to select hardware manually.
Basic Inputs (Required)		
Data ONTAP	Clustered Data ONTAP	Deployment requirement.
Data ONTAP version	8.2.x	Deployment requirement.
Controller platform	Mid FAS8040	Deployment requirement.
Degraded performance OK on HA takeover event	Selected	Default value. Note: As stated in the deployment requirement, some degree of performance degradation is acceptable if a storage controller is down for a short time.
Flash acceleration options	Flash Pool	Deployment requirement.
System headroom (%)	30	Default value.
Capacity reserve	50%	Anticipated data growth in the next 2 years.
Workload to shelf/drive type mapping	Selected	To continue disk shelf selection, select the Workload to Shelf/Drive Type Mapping drop-down menu.
Disk shelf for all workloads	DS2246	Deployment requirement. Note: 10K RPM SAS disks require disk shelf type DS2246.
Drive type for all workloads	SAS-10K-600 GB	Required to provide support for environments with high IOPS levels and low read latencies.
Flash Pool		
Flash Pool disk capacity	400GB	Select 400GB disks or later versions because the 200GB disk is at end of support.
Flash Pool capacity per controller (GB)	800GB	800GB is minimum for FAS8040 versions.
Other Inputs		
Map to full shelves	Deselected	Default value.
Spare disks per drive type	2	Default value.
System age	Empty system	Default value for a new deployment.

After the SPM input parameter values are defined for the storage controller and for the MEDITECH and BridgeHead workloads, you can generate the sizing report:

- To generate the report, click Calculate at the bottom of the SPM tool UI.
- To export the report to MS Word, click Export in the lower-right corner.

Figure 4 shows the projected system utilization across the two controllers as displayed in the SPM report.

Figure 4) Projected system utilization across two controllers in a 20-host MEDITECH environment.



The combined MEDITECH production and BridgeHead backup workloads should be located on a shared aggregate on each storage controller. However, NetApp SPM might divide the workloads and allocate them separately, potentially creating an unbalanced load across the two storage controllers. To balance workloads across the controllers, combine the total number of recommended data disks for the BridgeHead workload and distribute the disks equally to MEDITECH data aggregates on both storage controllers. The BridgeHead workload requires no capacity. For more information, see section 3.4, “Using SPM Sizing Results.”

Table 14 lists the suggested drive calculations per aggregate. These calculations are taken from the table on page 5 of the SPM report.

Table 14) Drive calculations per aggregate in a 20-host environment (from SPM report).

Aggregate	Data Drives Calculations					Flash Pool Drives Calculations			Total	RAID Group Sizes		Aggregate Free Space
	Total Drives for Capacity (A)	Total Drives for Performance (B)	Total Drives Max. (A, B)	RAID Drives	Total Data Drives (C)	Data Drives	RAID Drives	Total Flash Pool Drives (D)	Drives (C+D)	Data Drives RAID Group Size	Flash Pool Drives RAID Group Size	
N1_A2	13	9	13	2	15	2	2	4	19	16	16	0%
N2_A1	13	0	13	2	15	2	2	4	19	16	16	0%

According to the SPM report, a minimum number of 26 data disks (sized by capacity) is required to support a 20-host MEDITECH (category 2) 6.x environment. The recommended SSD RAID DP group for Flash Pool per controller is 4 × 400GB SSD drives (2 data + 2 parity drives).

Step 3: Select RAID Group Size for Ease of Deployment and Data Growth

Because an HA configuration contains two storage controllers, each storage controller is required to have a minimum of 13 ($26 \div 2$) data disks.

After you determine the number of disks required to handle the IOPS and latency targets for each storage controller, you can determine a RAID group size that enables expansion of storage in sensible increments for the capacity and performance required. For any environment, NetApp recommends adding disks to aggregates in disk quantities equal to the aggregate's RAID group size.

For this example, a RAID group size of 15 was chosen. Each RAID group consists of 13 data disks and 2 parity disks for a RAID DP group type. Therefore, a storage configuration with one RAID DP group (13 data disks) satisfies the minimum data disk count of 13 per storage controller recommended by the NetApp SPM to support the MEDITECH production and BridgeHead backup workloads.

Step 4: Determine Total Number of Disks Required on Each Storage Controller, Taking into Consideration Root Aggregate and Number of Spare Disks

When you size a MEDITECH environment, you must also consider the number of spare and root volume disks required for each storage controller. This task is explained in the section 3.5 subsection titled "Storage System Root Volumes and Spare Disks." Table 15 summarizes the disk requirements for the FAS8040HA storage system.

Table 15) Disk count summary for FAS8040HA storage system configured with Flash Pool for 20-host MEDITECH (category 2) 6.x environment.

Item	SAS Disk Count	SSD Count
Chosen RAID DP group size	15	4
Number of RAID DP groups per controller	1	1
Total data disks per controller	13	2
Total parity disks per controller	2	2
Spare disk count per controller	2	1
Root volume SAS disk count per controller	5	N/A
Total disks required per controller	22	5
Total disks required for storage system	44	10

Note: A total of 44 x SAS disks and 10 x SSDs is required to deploy a 20-host MEDITECH 6.x environment. The storage of system files for MEDITECH hosts is not included in Table 15.

Step 5: Verify That the Number of Data Disks Meets Sizing Requirements

Satisfying Flash Pool Guidelines in SPM Sizing Report

Regardless of the number of data disks recommended by SPM, the minimum SAS data disks might have to be adjusted (increased) to satisfy the first condition outlined in the "Flash Pool Guidelines" section of the SPM sizing report:

"For Flash Pool sizings resulting in greater than 75% reduction in HDD count compared to the configuration without Flash Pool, it is recommended that the HDD count is retained at 75% reduction when quoting to customers. This limit provides a buffer in situations where all the IOPS need to be served from HDDs while the SSD cache is still warming up, during boot up."

To determine the HDD count without Flash Pool, an SPM sizing is performed for the MEDITECH and BridgeHead workloads without Flash Pool configured.

The following steps make up the sizing process for the similar workload in this example without Flash Pool:

1. Click Perform Sizing to access the Prefilter Hardware Configuration dialog box.
2. In the Basic Inputs > Flash Acceleration Options section:
 - a. Disable Flash Cache by deselecting the Flash Cache checkbox.
 - b. Disable Flash Pool by deselecting the Flash Pool checkbox.
 - c. Enable the no-cache option by selecting the None checkbox.

With caching options disabled, the current storage configuration cannot meet the 5ms random read latency of the MEDITECH workload. SPM responds: `Validation Error: Please provide latency value greater than or equal to 8ms for the workload Id: 1 (i.e. MEDITECH).`

3. Set Custom App (MEDITECH) > Detailed Inputs > Random Read Latency (ms) to 9. Because the 10K SPM SAS disk type is used, the minimum achievable latency is approximately 9ms.

With the preceding modified SPM sizing inputs, SPM recommends 130 × 10K RPM SAS drives to support the 20-host MEDITECH and BridgeHead workloads. Note that the random read latency is increased from 5ms to 9ms to make this sizing possible.

Based on the Flash Pool guidelines referred to at the beginning of this section, 25% of the 130 disk count is made up of 32.5 data disks. Because the number of data disks recommended with Flash Pool is 26, the final number of data disks for deployment with Flash Pool must be increased to a minimum of 32 (32 ÷ 2 disks per controller). This number equals 16 data disks per controller.

Table 16 summarizes the disks required for the FAS8040 storage system after the disk count adjustment.

Table 16) Adjusted disk count summary for FAS8040HA storage system configured with Flash Pool for 20-host MEDITECH (category 2) 6.x environment.

Item	SAS Disk Count	SSD Count
Chosen RAID DP group size	15	4
Number of RAID DP groups per controller	2 (including the root aggregate)	1
Total data disks per controller	16	2
Total parity disks per controller	4	2
Spare disk count per controller	2	1
Root volume SAS disk count per controller	5	N/A
Total disks required per controller	27	5
Total disks required for storage system	54	10

Minimum Required Storage Capacity

The minimum required storage capacity specified in section 3.5, “Additional Sizing Requirements,” refers to the storage capacity needed to accommodate the storage capacity required by the MEDITECH hosts, the anticipated storage consumed by the Snapshot copies created by the BridgeHead backup operation, and the impact on storage consumption from traditional volume reallocate operations. For more information, see [NetApp TR-4300: NetApp FAS Storage Systems for MEDITECH Environments Best Practices Guide](#).

According to the formula specified in section 3.5, the minimum required storage capacity is (20 hosts × 200GB/host) × 1.5 = 6TB.

The total number of 15K RPM 600GB SAS data disks determined in step 4 is 44 (22 data disks per controller × 2 controllers). The storage capacity provisioned is 44 disks × 600GB/disk = 26.4TB.

The storage layout specified in Table 16 meets the storage requirement described in section 3.5 because 26.4TB is larger than 6TB.

Minimum and Maximum RAID Group Sizes

As of the writing of this document, based on the SAS drives and FC data network connectivity used in this example, NetApp recommends a minimum RAID group size of 12 and a maximum RAID group size of 24. For more information, see [NetApp TR-3838: Storage Subsystem Configuration Guide](#).

The RAID group size specified in Table 16 is 15. This size is within the recommended range for the RAID group size.

4.2 All Flash FAS Sizing Example for 60-Host MEDITECH (Category 2) 6.x and BridgeHead Workloads with Clustered Data ONTAP 8.3.1

This section describes an example to determine the NetApp AFF storage configuration for a BridgeHead and MEDITECH (category 2) 6.x environment with 60 hosts. In addition to the storage provisioned for MEDITECH hosts, disks are also allocated for the Data ONTAP root aggregate and a spare disk (per controller).

Complete the sizing process tasks in the following order:

1. Collect deployment information.
2. Use the NetApp SPM tool to determine the number of data disks required to satisfy the capacity and performance needs of the combined MEDITECH production and BridgeHead backup workloads.
3. Select the RAID group size for ease of deployment and data growth. For recommended RAID group sizes, see [NetApp TR-3838: Storage Subsystem Configuration Guide](#).
4. Determine the total number of disks required on each storage controller, taking into consideration the root aggregate and the number of spare disks.
5. Verify that the number of data disks meets the sizing requirements listed in section 3.5, "Additional Sizing Requirements."

Note: Storage for other workloads such as virtual machine infrastructure is considered separately from the MEDITECH production and BridgeHead backup workloads.

Step 1: Collect Deployment Information

This example uses the AFF8060 HA pair. 400GB SSDs are chosen as the data drives with the DS2246 disk shelves.

Table 17 lists the deployment information for this example.

Table 17) Deployment information for 60-host MEDITECH environment with Flash Pool sizing.

Item	Comments
NetApp storage system	AFF8060 in HA configuration
Data ONTAP version	Clustered Data ONTAP 8.3.x
Caching solution	None
Disk type	NetApp 400GB SSDs
Disk shelf	NetApp DS2246
Data network protocol	FCP
Performance expectation if one	As stated in the deployment requirement, some degree of performance

Item	Comments
storage controller is down	degradation is acceptable if a storage controller is down for a short time
MEDITECH category	2
MEDITECH platform	6.x
MEDITECH host count	60
MEDITECH host LUN size	200GB
Backup software	BridgeHead
Backup time window	Less than 8 hours
Data growth in the next 2 years	0%
Legacy data	None; this is a new deployment

Step 2: Use SPM to Determine the Number of Disks Required for MEDITECH and BridgeHead Backup Workloads

Based on the information collected and determined in Table 17, the next step is to translate the information as input to the SPM sizing tool. The MEDITECH and BridgeHead backup workload and AFF8060 hardware configuration SPM parameters used in this example are shown in the following tables.

Log in to the [SPM](#) tool and click the Workload tab to create the new workload, as shown in Figure 5.

Figure 5) Screenshot of SPM custom application for MEDITECH (category 2) 6.x workload.

Workload Title: MEDITECH 6.x with 60 hosts

Workload Type: Custom

Basic Inputs (Required)

Capacity (TiB): 12

Throughput: 45000 IOPS

Detailed Inputs

High Impact

Random Read Latency (ms): 5

IO Percent

Workload IO Type: ☒ Only Random ☐ Mixed ☐ Only Sequential

Random Read (%): 20

Random Write (%): 80

Sequential Read (%): 0

Sequential Write (%): 0

Working Set Size

Active Working Set (%): 10

Low Impact

Protocol: FCP

IO Sizes

Random Read Size (KB): 4KB

Random Write Size (KB): 4KB

Sequential Read Size (KB): 64KB

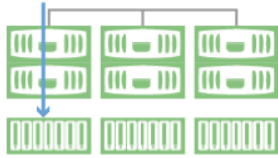
Sequential Write Size (KB): 64KB

Layout Hints

Can this workload be placed on a shared aggregate? ☒

Can this workload be split on different aggregates across the cluster? ☒

How will the workload traffic arrive into the cluster? ☒ On a single node ☐ All nodes in the cluster



Data Protection

Enable SnapMirror: ☒ Yes ☐ No

SnapMirror Role: Source

SnapMirror Type: SnapMirror

Daily Data Change Rate (%): 10

RPO Required (mins): 10

Number of Mirrors: 1

Table 18 lists the parameters that define the MEDITECH workload if the SPM tool is used to create a custom application.

Table 18) Example of MEDITECH workload SPM parameter values for 60-host MEDITECH environment.

SPM Parameter	Value	Comments
Identifiers		
Workload title	MEDITECH (category 2) 6.x with 60 hosts	Workload name.
Workload type	Custom	Select the custom workload type for the MEDITECH workload.
Basic Inputs (Required)		
Capacity (TB)	12	In this example, the LUN size per MEDITECH host is 200GB, and there are 20 hosts to be deployed. Therefore, the total storage capacity is $60 \times 200\text{GB} = 12\text{TB}$.
Throughput	45,000	Because there are 60 MEDITECH hosts and each host averages 750 IOPS, the total IOPS is 45,000. For more information, see section 2.1, "MEDITECH Workload Description."
Throughput type	IOPS	Throughput in IOPS instead of MBps.
Detailed Inputs		
Random read latency (ms)	2	MEDITECH workload requirement as described in section 2.1, "MEDITECH Workload Description," is 5. Note: For AFF platforms, the maximum value for this option is 3.
I/O Percent		
Workload I/O type	Only random	MEDITECH workloads are 100% random.
Random read (%)	20	MEDITECH workload requirement as described in section 2.1, "MEDITECH Workload Description."
Random write (%)	80	MEDITECH workload requirement as described in section 2.1, "MEDITECH Workload Description."
Working Set Size		
Active working set (%)	10	MEDITECH workload requirement as described in section 2.1, "MEDITECH Workload Description."
Protocol	FCP	Deployment requirement.
I/O Sizes		
Random read and write size (KB)	IO-SIZE-4K	MEDITECH workload requirement as described in section 2.1, "MEDITECH

SPM Parameter	Value	Comments
		Workload Description."
Sequential read and write size (KB)	N/A	Default value. Note: Sequential read and write size does not apply to the MEDITECH workload because this workload is 100% random.
Flash Pool: Random overwrite (%)	100	Flash Pool overwrite is always 100% for MEDITECH workloads. Note: The Flash Pool random overwrite option does not appear when the MEDITECH workload is defined. The option appears only after Flash Pool is selected under Hardware.
Layout Hints		
Can this workload be placed on a shared aggregate?	Selected	The default value is for the MEDITECH workload to share its aggregate with the BridgeHead backup software.
Can the workload be split on different aggregates across the cluster?	Selected	The default value is for the MEDITECH workload to be split evenly between two storage controllers as specified in the deployment specification in Table 17.
How will the workload traffic arrive into the cluster?	On a single node	Default option. Note: This option might not be available until the first SPM sizing evaluation is attempted.
Data Protection		
Enable SnapMirror	Yes	Click Yes to enable SnapMirror.
SnapMirror role	Source	Select <code>Source</code> as the SnapMirror role and keep the default values in the other fields.

Next, use SPM to define the BridgeHead workload. BridgeHead is used to manage and perform nightly backup on site. The backup window is set to a maximum of 8 hours. Conservatively, a maximum backup time of 5 hours is used to calculate the required throughput. For the calculation, see the corresponding SPM parameter for throughput.

Figure 6 shows the SPM parameters for the BridgeHead workload.

Figure 6) Screenshot of SPM custom application for BridgeHead workload.

Workload Title

Workload Type

▼ Basic Inputs (Required) ?

Capacity (TiB)

Throughput

▼ Detailed Inputs ?

High Impact ☐

Random Read Latency (ms)

IO Percent

Workload IO Type ☐ Only Random ☐ Mixed ☒ Only Sequential

Random Read (%)

Random Write (%)

Sequential Read (%)

Sequential Write (%)

Working Set Size

Active Working Set (%)

Low Impact ☐

Protocol

IO Sizes

Random Read Size (KB)

Random Write Size (KB)

Sequential Read Size (KB)

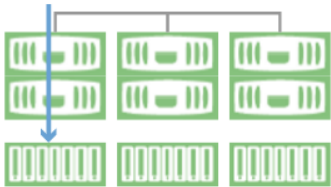
Sequential Write Size (KB)

Layout Hints

Can this workload be placed on a shared aggregate? ☒

Can this workload be split on different aggregates across the cluster? ☒

How will the workload traffic arrive into the cluster? ☒ On a single node ☐ All nodes in the cluster



Data Protection

Enable SnapMirror ☐ Yes ☒ No

Table 19 lists the parameters used to create a custom application in the SPM tool to define the BridgeHead workload.

Table 19) Example of BridgeHead workload SPM parameter values for 60-host MEDITECH environment.

SPM Parameter	Value	Comments
Identifiers		
Workload title	BridgeHead workload for 60-host MEDITECH environment	Workload name.
Workload type	Custom	Select the custom workload type for the BridgeHead workload.
Basic Inputs (Required)		
Capacity (TB)	12	Same as MEDITECH workload capacity.
Throughput	667	Throughput is specified in MBps. Total data to back up is 4TB (20 MEDITECH hosts × 200GB per host). Total time allocated for backup is 5 hours or 18,000 seconds (5 hours × 3,600 seconds/hour). Throughput required is 234 ($4 \times 1024^2 \text{MB} \div 18,000 \text{ seconds}$).
Throughput type	MBps	Throughput is specified in MBps instead of IOPS.
Detailed Inputs		
Random read latency (ms)	20	Default value. Note: Random read latency does not apply to the BridgeHead workload because this workload is 100% sequential read.
Workload I/O type	Only sequential	BridgeHead workload is always 100% sequential read.
Sequential read (%)	100	BridgeHead workload is 100% sequential read.
Sequential write (%)	0	BridgeHead workload is 100% sequential read.
Working Set Size		
Active working set (%)	10	Same as MEDITECH working set size.
Protocol	FCP	Deployment requirement.
I/O Sizes		
Random read and write size (KB)	Not applicable (NA)	N/A
Sequential read and write size (KB)	64KB	BridgeHead sequential read size is 64KB.

SPM Parameter	Value	Comments
Layout Hints		
Can this workload be placed on a shared aggregate?	Selected	Default value. Note: The MEDITECH workload shares its aggregate with the BridgeHead backup software.
Can the workload be split on different aggregates across the cluster?	Selected	Default value. Note: The BridgeHead workload is split evenly between two storage controllers as described in the deployment specification in Table 17.
How will the workload traffic arrive into the cluster?	On a single node	Default option. Note: This option might not be available until the first SPM sizing evaluation is attempted.
Data Protection		
Enable SnapMirror	No	Click No to deselect SnapMirror.

Figure 7 shows the SPM parameters used for storage configuration.

Figure 7) Screenshot of SPM storage configuration.

☐ Let SPM autosuggest hardware for me.

▼ Basic Inputs (Required) ?

[Note: For sizing EOA hardware, use Performance Estimation (Reverse Sizing)]

Deployment Options ☐ All-Flash FAS (AFF) ☐ MetroCluster

Data ONTAP ☒ Clustered ONTAP ☐ 7-Mode

8.3.x

Controller Platform ☐ Entry ☐ Mid ☒ High

Hold Ctrl to select multiple Controller Platforms.

FAS2554 FAS2552 FAS2530	FAS8040 FAS8030	FAS8080 FAS8060
-------------------------------	--------------------	--------------------

Degraded Performance OK on HA takeover event ☒

Disk Type ☐ SAS ☐ SATA ☐ SSD

Flash Acceleration Options ☐ None ☐ FlashCache ☐ FlashPool

☐ Enable Storage Pool Sharing

System Headroom (%) 30

Capacity Reserve (%) 0

▼ Workload to Shelf/Drive Type Mapping ?

☐ Enable Drive Type Selection per workload

Disk Shelf for all workload(s) DS2246

Drive Type for all workload(s)

SAS-10K-900GB
SAS-10K-600GB
SAS-10K-1.8TB
SAS-10K-1.2TB
NETAPP-SSD-800GB

Maximum 2 drive types are allowed to size per sizing. Please select any 2 drive types from the list.

Detailed Inputs ?

Is 2-Node MetroCluster ? ☐

HA Pair ☒

FlashCache

Card Type AUTO-SUGGEST

Cards per Controller 1

Card Capacity 0

Flash Pool

Flash Pool Disk Capacity Select SSD Disk Capacity

Flash Pool Capacity per controller (GB) 800

[Flash Pool limits](#)

Other Inputs

Map to Full Shelves ☐

Spare Disks per Drive Type 2

System Age Empty System

Table 20 lists the SPM parameters used to configure the storage hardware in the example of a 60-host MEDITECH environment.

Table 20) Example of storage hardware configuration SPM parameter values for 60-host MEDITECH environment with AFF8060.


SPM Parameter	Value	Comments
Manual Selection Option		
Let SPM autosuggest hardware for me	Deselect the checkbox	Deselecting this option enables you to select hardware manually.
Basic Inputs (Required)		
Data ONTAP	Clustered Data ONTAP	Deployment requirement.
Deployment options	Select All Flash FAS	Deployment requirement.
Data ONTAP version	8.3.x	Deployment requirement.
Controller platform	High AFF8060	Deployment requirement.
Degraded performance OK on HA takeover event	Selected	Default value. Note: As stated in the deployment requirement, some degree of performance degradation is acceptable while a storage controller is down for a short time.
Flash acceleration options	None	Deployment requirement.
System headroom (%)	30	Default value.
Capacity reserve	0%	Leave capacity reserve at 0%.
Workload to shelf/drive type mapping	Selected	Use the Workload to Shelf/Drive Type Mapping drop-down menu to continue disk shelf selection.
Disk shelf for all workloads	DS2246	Deployment requirement; 400GB SSDs require disk shelf type DS2246.
Drive type for all workloads	SSD	Required to provide support for environments with high IOPS levels and low read latencies.
Other Inputs		
Map to full shelves	Deselected	Default value.
Spare disks per drive type	2	Default value.
System age	Empty system	Default value (for a new deployment).

After you define the SPM input parameter values for the storage controller and for the MEDITECH and BridgeHead workloads, you can generate the sizing report:

- To generate the report, click Calculate at the bottom of the SPM tool UI.
- To export the report to MS Word, click Export in the lower-right corner.

In Figure 8, a graph from the SPM report shows the projected system utilization across the two controllers.

Figure 8) Projected system utilization across two controllers in a 60-host MEDITECH environment.



The combined MEDITECH production and BridgeHead backup workloads should be placed on a shared aggregate on each storage controller. However, NetApp SPM might divide the workloads and allocate the two workloads separately, potentially creating an unbalanced load across the two storage controllers. To balance workloads across the controllers, combine the total number of recommended data disks for the BridgeHead workload and distribute the disks equally to MEDITECH data aggregates on both storage controllers. The BridgeHead workload requires no capacity. For more information, see section 3.4, “Using SPM Sizing Results.”

Table 21 lists the suggested drive calculations per aggregate. These calculations are taken from the table on page 5 of the SPM report.

Table 21) Drive calculations per aggregate in a 60-host environment (from SPM report).

Aggregate	Data Drives Calculations					Total	Size	Aggregate		
	Total Drives for Capacity (A)	Total Drives for Performance (B)	Total Drives Max. (A, B)	RAID Drives	Total Data Drives (C)	Drives (C +D)	Data Drives RAID Group Size	Free Space	Capacity Savings (in GB) from Inline Compression	Capacity Savings (in GB) from Inline Deduplication
N1_A2	5	8	8	2	10	10	16	40.28%	0	0
N2_A1	38	8	38	4	42	42	21	0%	0	0

According to the SPM report, 52 data disks at minimum (sized by capacity) are required to support a 60-host MEDITECH (category 2) 6.x environment.

To balance the load between two controllers, redistribute the disks equally between all MEDITECH data aggregates across the cluster.

Based on the drive calculations report, SPM suggested a total of 77 data disks and 8 parity disks.

5 Sizing and Storage Layout Recommendations to Deploy 7 MEDITECH (Category 1) 6.x Hosts and FAS2240HA Storage System with 24 Internal Disks

This section applies only to using a FAS2240 storage system in a 2HA configuration with 24 internal disks to deploy 7 MEDITECH (category 1) 6.x hosts with 1 BridgeHead backup server. This storage configuration is certified and tested specifically for small MEDITECH category 1 deployments.

5.1 MEDITECH, BridgeHead, and NetApp FAS Storage Deployment Requirements

The sizing recommendations in this section apply only to the deployment of a MEDITECH system with a FAS2240HA system that meets the requirements listed in Table 22.

Table 22) Deployment requirements for seven MEDITECH (category 1) 6.x hosts with one BridgeHead backup server and an FAS2240HA storage system with internal disks.

Requirement	Value	Notes
MEDITECH category	1	N/A
MEDITECH platform	6.x	N/A
Maximum number of MEDITECH hosts	7 or fewer MEDITECH hosts	The MEDITECH hosts are virtualized by using VMware ESXi hosts.
Total IOPS from the MEDITECH hosts	2,000 IOPS or less	N/A
Read and write latency	Refer to Table 1	N/A
Total storage capacity for MEDITECH hosts	Approximately 350GB	50GB LUN size per MEDITECH host
BridgeHead backup server	One physical or virtual BridgeHead backup server	A minimum of 1Gbps Ethernet bandwidth is required for the iSCSI traffic. If the server is virtualized by using VMware, the Ethernet port of the ESXi host dedicated for iSCSI traffic must be configured in the DirectPath I/O mode.
NetApp FAS storage system	FAS2240HA	N/A
Data ONTAP version	Clustered Data ONTAP 8.2.1 or later	N/A
Flash Pool	Required	N/A
Maximum number of spinning disks and SSDs for the Data ONTAP root, storage for MEDITECH LUNs, VMDK datastore and spare disks	The combined total number of spinning disks and SSDs for Flash Pool should not exceed 24 disks for the entire storage system	No external disk shelf is required because the FAS2240HA has 24 internal disk slots.
FAS spinning disk type	10K RPM SAS drives	N/A
Data protocol	iSCSI	iSCSI is the only option for data network protocol on the FAS2240HA configuration because clustered Data ONTAP uses the

Requirement	Value	Notes
		one PCIe expansion slot on the FAS2240HA for the cluster network.
Site architecture	Single site	No disaster recovery site.
SnapMirror	No SnapMirror setup	N/A

5.2 Storage Configuration Recommendations

The storage configuration recommendations are derived from a tested configuration using the 24 internal disks of the FAS2240HA storage system to support seven MEDITECH (category 1) 6.x hosts. These recommendations are specific to the requirements listed in Table 22. These recommendations should not be used as general guidelines for sizing and storage configuration recommendations unless all conditions in Table 22 are met.

For specific storage recommendations, see [NetApp TR-4300: NetApp FAS Storage Systems for MEDITECH Environments Best Practices Guide](#).

5.3 Storage Layout for FAS2240HA with a Maximum of 24 Internal Disks

Table 23 presents the storage layout recommended for the FAS2240HA with 24 disks for deploying 7 MEDITECH category 1 hosts virtualized through VMware ESXi hosts.

The sizing is based on the following factors:

- A MEDITECH category 1 6.x workload of 2,000 IOPS, 7ms random read latency, and 1ms random write latency
Note: The total storage capacity required is 350GB.
- A BridgeHead workload of 70MBps
- A Flash Pool configuration that uses 200GB SSDs
- FAS2240HA storage controllers with 10K RPM SAS disks
- The NetApp best practice of having a minimum RAID group size of 12

Table 23) Disk layout for the FAS2240HA storage system configured with 24 internal disks and Flash Pool for 7 MEDITECH (category 1) 6.x hosts.

Item	Storage Controller A		Storage Controller B	
	SAS Disk Count	SSD Count	SAS Disk Count	SSD Count
Total number of RAID DP groups for data aggregate	1	1	0	0
RAID DP group size chosen for data aggregate	12	3	0	0
Total data disks	10	1	0	0
Total parity disks	2	2	0	0
Spare disk count	1	1	1	0
Root volume SAS disk count	3	N/A	3	N/A
Total disks required	16	4	4	0

Item	Storage Controller A		Storage Controller B	
Total disks required for storage system	20	4		

Note: A total of 20 SAS disks and 4 SSDs are required to deploy 7 MEDITECH (category 1) 6.x hosts. The data aggregate on storage controller A provides storage to both the LUNs used by the MEDITECH hosts and the VMDK NFS datastore used by the VMware ESXi for the Windows system files of the MEDITECH hosts.

Note: The storage controller pair is configured in the active-passive mode. The active-active mode configuration is required except for MEDITECH deployments that meet the conditions and requirements listed in Table 22.

Note: The VMware ESXi host local storage provides storage for the VMware vCenter server and the virtual BridgeHead backup server.

Data Aggregate Recommendation

The recommendation in Table 23 to use one aggregate for MEDITECH LUNs and VMDK files applies only to the deployment requirements specified in Table 22. The rest of this document provides sizing guidance about other deployments.

6 MEDITECH Certified and Validated SPM Storage Configurations

The sizing methodology described in this document was used to build the storage configurations for the certification of NetApp FAS storage systems with the MEDITECH and BridgeHead workloads.

All certified storage configurations used two storage controllers in an HA configuration. Storage configurations with clustered Data ONTAP were configured with two nodes for the primary cluster.

6.1 Certified and Validated SPM Storage Configurations

Table 24 lists the storage configurations that are tested and certified with MEDITECH.

Table 24) Certified storage configurations for MEDITECH environments.

MEDITECH Platform	Number of MEDITECH Hosts	Storage Controller	Data ONTAP Version	Disk Shelf and Disk Type	Caching Technology and Total Capacity (TB)
6.x (category 2)	116	AFF8060	Clustered Data ONTAP 8.3.1	Disk shelf: DS2246 Disk type: 10K RPM SAS 600GB	AFF
	58	FAS8040	Clustered Data ONTAP 8.3	Disk shelf: DS2246 Disk type: 10K RPM SAS 600GB	Flash Pool: • 2.4TB: 60 hosts
	18	FAS2552	Clustered Data ONTAP 8.2.3	Disk shelf: DS2246 Disk type: 10K RPM SAS 600GB	Flash Pool: • 800GB: 18 hosts
	40	FAS3250HA	Clustered Data ONTAP 8.2	Disk shelf: DS2246 Disk type: 10K RPM SAS 600GB	Flash Pool: • 1.7TB: 40 hosts • 1.1TB: 25 hosts

MEDITECH Platform	Number of MEDITECH Hosts	Storage Controller	Data ONTAP Version	Disk Shelf and Disk Type	Caching Technology and Total Capacity (TB)
	50	FAS3250HA	Data ONTAP 8.1.2 7-Mode	Disk shelf: DS4243 Disk type: 15K RPM SAS 450GB	Flash Cache: • 2TB: 50 hosts • 1TB: 25 hosts
	25	FAS3220HA			
6.x (category 1)	7	FAS2240HA	Clustered Data ONTAP 8.2.1	Disk shelf: Internal Disk type: 10K RPM SAS 900GB	Flash Pool: • 142.5GB
C/S 5.x (category 2)	50	FAS3250HA	Data ONTAP 8.1.2 7-Mode	Disk shelf: DS4243 Disk type: 15K RPM SAS 450GB	Flash Cache: • 2TB: 50 hosts • 1TB: 25 hosts
	25	FAS3220HA			
	40	FAS3250HA	Clustered Data ONTAP 8.2	Disk shelf: DS2246 Disk type: 10K RPM SAS 600GB	Flash Pool: • 1.7TB: 40 hosts • 1.1TB: 25 hosts
MAGIC (category 2)	50	FAS3250HA	Data ONTAP 8.1.2 7-Mode	Disk shelf: DS4243 Disk type: 15K RPM SAS 450GB	Flash Cache: • 2TB: 50 hosts • 1TB: 25 hosts
	25	FAS3220HA			
	40	FAS3250HA	Clustered Data ONTAP 8.2	Disk shelf: DS2246 Disk type: 10K RPM SAS 600GB	Flash Cache: • 2TB: 40 hosts Flash Pool: • 1.7TB: 40 hosts • 1.1TB: 25 hosts

Note: The VMware ESXi hypervisor is used to virtualize the MEDITECH hosts.

Note: MEDITECH has added the FAS6000 and FAS8000 series storage controllers to its certified storage hardware list. Because no test was performed on these FAS controller platforms, the number of supported MEDITECH hosts per controller platform is not specified in the MEDITECH certified hardware list.

SPM System Headroom Percentage Recommendation

In the performance validation lab, the System Headroom Percent value is set to 0% to size the storage for the maximum number of MEDITECH hosts with the minimum number of storage controller HA pairs.

When you size for production environments, NetApp recommends that you use the default value for system headroom percentage (30% in SPM 1.4.2 [P1]).

6.2 Performance Validation

Each storage configuration in Table 24 was tested against its corresponding MEDITECH and BridgeHead workloads. The MEDITECH workload was generated by the MEDITECH application simulator operated by a MEDITECH test engineer.

Key observations recorded during the performance validation include the following:

- The MEDITECH system's overall IOPS level was sustained throughout the duration of the tests.

- Read latency was below the 5ms requirement most of the time for all platforms with both MEDITECH production and BridgeHead backup workloads running concurrently. The read latency reached beyond 5ms for short durations, which MEDITECH deemed acceptable.
- Write latency was below the 1ms requirement most of the time for all platforms with both MEDITECH production and BridgeHead backup workloads running concurrently. The write latency reached beyond 1ms for short durations, which MEDITECH deemed acceptable.
- Average BridgeHead backup times were well within the required eight-hour window.

References

This document references the following NetApp technical reports and other resources:

- Flash Pool Technical FAQ
<https://fieldportal.netapp.com/content/195703?assetComponentId=195801>
- TR-3838: Storage Subsystem Configuration Guide
<http://www.netapp.com/us/media/tr-3838.pdf>
- TR-4050: System Performance Modeler
<https://fieldportal.netapp.com/content/200086?assetComponentId=200184>
- TR-4070: Flash Pool Design and Implementation Guide
<http://www.netapp.com/us/media/tr-4070.pdf>
- TR-4300: NetApp FAS Storage Systems for MEDITECH Environments Best Practices Guide
<https://fieldportal.netapp.com/content/201762?assetComponentId=201860>

In addition, the following resources are relevant to this document:

- Advanced Disk Partitioning Video
<https://fieldportal.netapp.com/content/228948?assetComponentId=229061>
- FlashSelect Tool
<https://one.netapp.com/flashselect/>
- Sample AFF sizing video demonstration
<https://fieldportal.netapp.com/content/228883?assetComponentId=228996>
- Sizing concepts video demonstration
<https://fieldportal.netapp.com/content/228855?assetComponentId=228968>

Note: Some documents in this list are available only through the NetApp Field Portal. For information about how to access the Field Portal, contact your NetApp field representative.

Version History

Version	Date	Document Version History
Version 1.3	March 2016	Updated entire document with the latest information.
Version 1.2	May 2014	Added the concept of MEDITECH category Added sizing with iSCSI protocol Added sizing recommendations for deploying MEDITECH (category 1) 6.x hosts with FAS2240HA using 24 internal disks
Version 1.1	February 2014	Added sizing information for 10K RPM SAS disks with Flash Pool (all platforms) and Flash Cache (MAGIC)
Version 1.0	May 2013	Initial release (15K RPM SAS disk with Flash Cache)

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-4190-0316