

DATA DRIVEN GLOBAL VISION CLOUD PLATFORM STRATE
ON POWERFUL RELEVANT PERFORMANCE SOLUTION CLO
VIRTUAL BIG DATA SOLUTION ROI FLEXIBLE DATA DRIVEN

WHITE PAPER

Management Abstraction With Hitachi Storage Advisor

What You Don't See Is as Important as What You Do See (WYDS)

By Hitachi Data Systems

October 2016

Contents

Executive Summary 2

Introduction 3

Auto Everything 3

Autobuild: Parity Group Initialization (VSP Gx00 Series Only) 3

WYDS: What’s in the Create Parity Group Logic? 4

Autosize: Pool Creation 5

WYDS: What’s in the Create Pool Logic? 6

Autopath: Port Assignment 7

WYDS: What’s in the Attach Volume Logic? 8

Autozone: Zone Creation 9

WYDS: What’s in the Attach Volume Logic? 10

Complete the Job 10

Autoabstract File-Pool Creation 11

WYDS: What’s in the File-Pool Creation Logic? 12

Summary 13

Executive Summary

Hitachi Storage Advisor (HSA) is a configuration management application for the Hitachi Virtual Storage Platform (VSP) family, including VSP G1500, VSP G1000, VSP Gx00 series, VSP F1500 and VSP Fx00 series. It enables a simplified, smart and extensible approach to managing unified storage infrastructure resources. With intuitive graphical user interfaces (GUIs) and recommended configuration practices, end users can streamline system configurations and storage management operations to easily provision storage for new applications. There is no need for in-depth knowledge of the underlying storage resource details.



With new levels of management abstraction, HSA provides integrated best practices to the configuration management process, which enable administrative complexities to be hidden from the user. The key is intelligent, guided configuration workflows, including recommended system configurations to streamline administrative operations. These guided workflows raise this point: What you don't see within Storage Advisor is as important as what you do see. This white paper helps to explain these recommended configuration workflows and describes the logic behind them, from the creation of storage parity groups and pools to the assignment of ports and zones. This approach makes provisioning storage to servers, LUN to port assignment and SAN fabric zoning much less complex.

Introduction

In this white paper, we focus on Hitachi Storage Advisor (HSA) for the storage component of the Hitachi infrastructure management portfolio and how it addresses industry requirements to simplify the overall unified storage management experience.

Hitachi Storage Advisor is a configuration management application that reduces infrastructure complexities and utilizes a new graphical user interface to simplify the user experience. It also offers a REST-based application-programming interface (API) for management extensibility. HSA abstracts the complexities of traditional Fibre Channel storage and storage area network (SAN) management, which results in lower operational costs and speeds time to deploy new infrastructure resources.

Much has been discussed through the years with respect to Fibre Channel-based storage management being too complex, resulting in rising management costs. What if the management software had built-in intelligence to apply recommended practices for traditional Fibre Channel storage and SAN management, including parity group initialization, pool creation, port assignment and SAN zoning?

In order for Hitachi Storage Advisor to reduce the complexities associated with Fibre Channel and SAN management, new levels of automation had to be implemented. The rest of this white paper describes these new types of automation that Hitachi Storage Advisor employs.

Auto Everything

The key to Storage Advisor's management abstraction is a concept of "auto everything," which covers the following areas of parity groups, storage pools, ports and zoning:

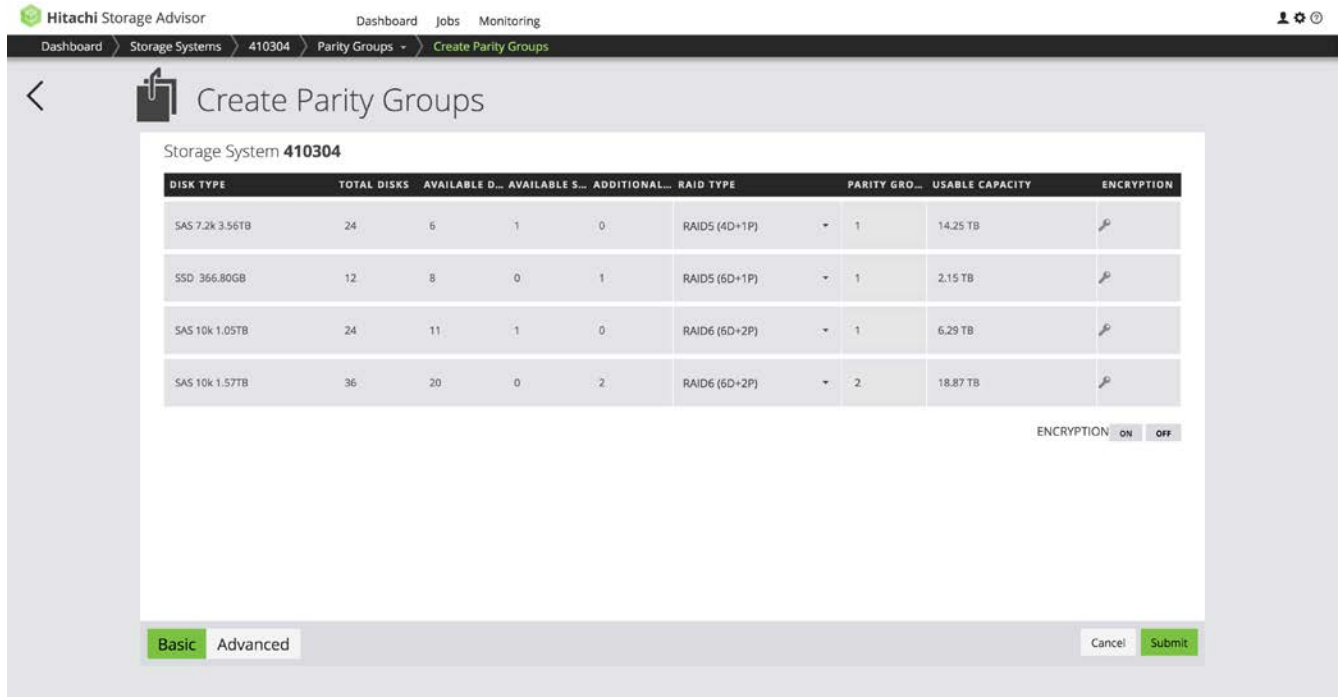
- Autobuild: parity group initialization.
- Autoselect: pool creation.
- Autopath: port assignment.
- Automap: zone creation.

Autobuild: Parity Group Initialization (VSP Gx00 Series Only)

Before a dynamic storage pool (dynamic pool, dynamic tier or thin image) can be created and, subsequently, volumes can be created, parity groups must be built and initialized.

Recommended RAID-6 parity groups are presented to the user. Selecting the defaults will result in all available raw capacity being formatted and initialized based on recommended practices. As shown in Figure 1, with one click, the job is done.

Figure 1. Create Parity Groups



WYDS: What’s in the Create Parity Group Logic?

When presented the default proposal for creating parity groups, Hitachi Storage Advisor takes into account all different media types detected in the storage system. It uses best practices to determine the default RAID type that is recommended per media type and the recommended hot spare ratio per media type.

When proposing the default plan to create parity groups on the storage system, consider three points:

1. Based on the total number of disks detected per media type, number of spare disks that already exists on the array, and the recommended hot spare allocation ratio, HSA calculates and proposes how many new hot spare disks must be assigned per media type.
2. Additionally, HSA also recommends the default RAID configuration that should be used for each media type.
3. Finally, based on the RAID type, layout and number of disks per media type, HSA presents the number of parity groups that can be created and the total usable capacity available upon creating parity groups.

To calculate the number of hot spare disks, HSA uses the best practices shown in Table 1.

Table 1. HSA Best Practices for Hot Spare Disks

Disk Type	Hot Spare Ratio
Solid State Disk (SSD)	1 per 32
Flash Module Drive (FMD)	1 per 24
Serial-Attached SCSI (SAS) 15K, SAS 10K	1 per 32
SAS 7.2K	1 per 16

To propose the default RAID configuration, HSA uses the best practice shown in Table 2.

Table 2. HSA Best Practices for Default RAID

Disk Type and Speed	RAID layout
SSD, FMD	RAID-6: 6D+2
SAS 15K, SAS 10K	RAID-6: 6D+2
SAS 7.2K	RAID-6: 14D+2

If, for any media type, best practices cannot be fulfilled, then Hitachi Storage Advisor will indicate no parity groups can be created. In such cases, end users can select HSA advanced mode to create parity groups without complying with included best practices.

When the end user reviews the available options and proceeds to create parity groups, Hitachi Storage Advisor does the following:

1. For each media type, it calculates the total number of disks and reserves hot spares.
 - a. Disk selection for hot spares:
 - i. HSA first checks to see if sufficient hot spares already exist. If hot spares need to be assigned, then HSA selects a disk from a tray that has the least (or no) assigned hot spares.
2. For the requested RAID layout for each media type and speed, HSA calculates the number of parity groups that can be created and proceeds to select disk for the required number of parity groups.
 - a. Disk selection for creating each parity group:
 - i. Select identical disks in terms of disk type, speed and capacity.
 - ii. Distribute the disks of a parity group across as many trays as possible and as many back-end paths as possible, up to maximum of four.
3. Hitachi Storage Advisor then creates and quickly formats storage logical devices (LDEV) on each parity group created so that the parity groups are ready for pool creation:
 - a. It creates the largest possible equal size volumes and consumes as much space as possible with a maximum size of 2.99TB and minimum size of 8GB.

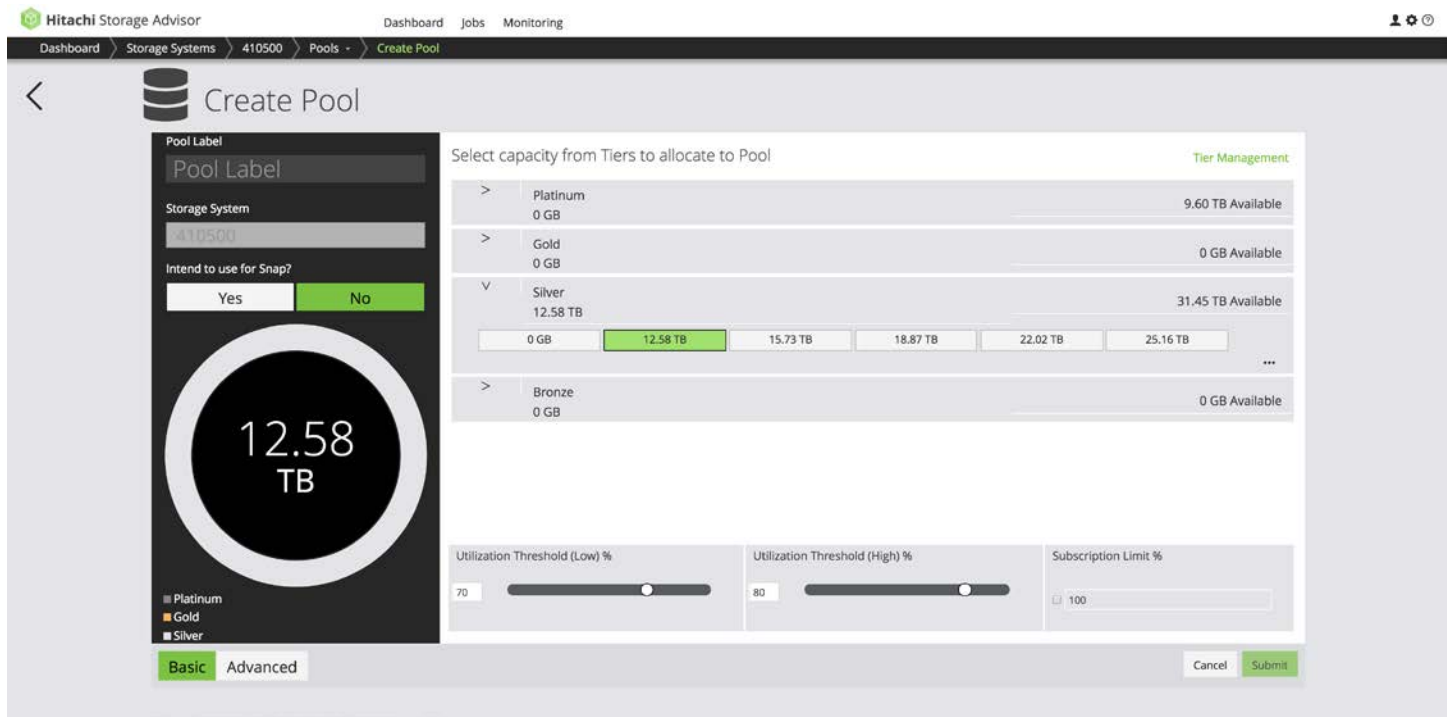
Autosize: Pool Creation

Before volumes can be created and attached to hosts, a pool or pools need to be created (see Figure 2). Three types of pools are supported: Thin (Hitachi Dynamic Provisioning or HDP), Snap (Hitachi Thin Image or HTI) and Tiered (Hitachi Dynamic Tiering or HDT). Choosing between a thin versus tiered pool is very intuitive and is accomplished just by selecting either one or more tiers. To define a snapshot pool, an end user just clicks on the “Yes” button of the “Intend to use for Snap” question.

Note: When building pools, if a single storage tier is chosen, then a thin HDP pool will be built. If multiple storage tiers are chosen (two or three), then a tiered HDT pool will be built. Note: a fourth tier cannot be chosen because HDP and HDT only support a maximum of three tiers.

Selection of array groups is totally automated. Default parity groups capacities are displayed in increments of usable capacity.

Figure 2. Pool Creation: Create Pool Dialog



WYDS: What’s in the Create Pool Logic?

When presenting options for pool creation, Hitachi Storage Advisor takes into account different storage tiers of media that are available on a given storage system. These tiers map to the underlying media type in the storage system, as shown in Table 3.

Table 3. Storage Tiers and Media Types

Tier	Media Type
Platinum	Solid-State Disk (SSD), Flash Module Drive (FMD)
Gold	Serial-Attached SCSI (SAS) 15K
Silver	SAS 10K
Bronze	SAS 7.2K

Based on the number of identical parity groups available for each tier, HSA calculates various possible pool sizes for each tier. While doing so, HSA checks the following:

1. Selects identical parity groups that use the same disk type, disk capacity, disk speed, RAID type and RAID layout.
2. Uses the entire capacity of the parity group (that is, all the LDEVs on a given parity group) so that parity groups are not split across multiple storage pools.
 - a. The only exception will be a parity group on which a command device is created. In this case, the remaining capacity of the parity group can be allocated to a pool.
3. Uses minimum of four parity groups to create a pool.
4. Uses default subscription to 100% for caution.

Based on the above rules, HSA determines all possibilities of combining parity groups to calculate different possible pool sizes for each tier. It also ensures that the maximum pool size is within the supported limits of the hardware.

If insufficient parity groups are available to meet the above-listed best practices, HSA will show no options to create a pool. In such cases, the end user can go to HSA advanced mode to create pools without complying with all the best practices.

When the end user selects a size of each tier required for the pool, HSA does the following:

1. Identifies the parity groups needed to create the pool using the same rules used to calculate pool sizes.
2. Creates the pool using selected parity groups:
 - a. If more than one tier was selected, HSA created a HDT pool and enables continuous monitoring and tiering (24 hour period).
 - b. If no end user inputs are provided for setting thresholds, it sets default pool thresholds to 70% and 80% and oversubscription to 100%.
 - c. For snapshot pools, HSA ensures that single tier pools are created.

When expanding a pool, HSA uses the same best practices. When a single tier pool is expanded, one or more parity groups can be added to the pool. If a new tier is being added to an existing pool:

1. HSA ensures that the new tier uses at least four or more identical parity groups.
2. HDT monitoring and tiering are enabled if an HDP pool is being converted to an HDT pool.

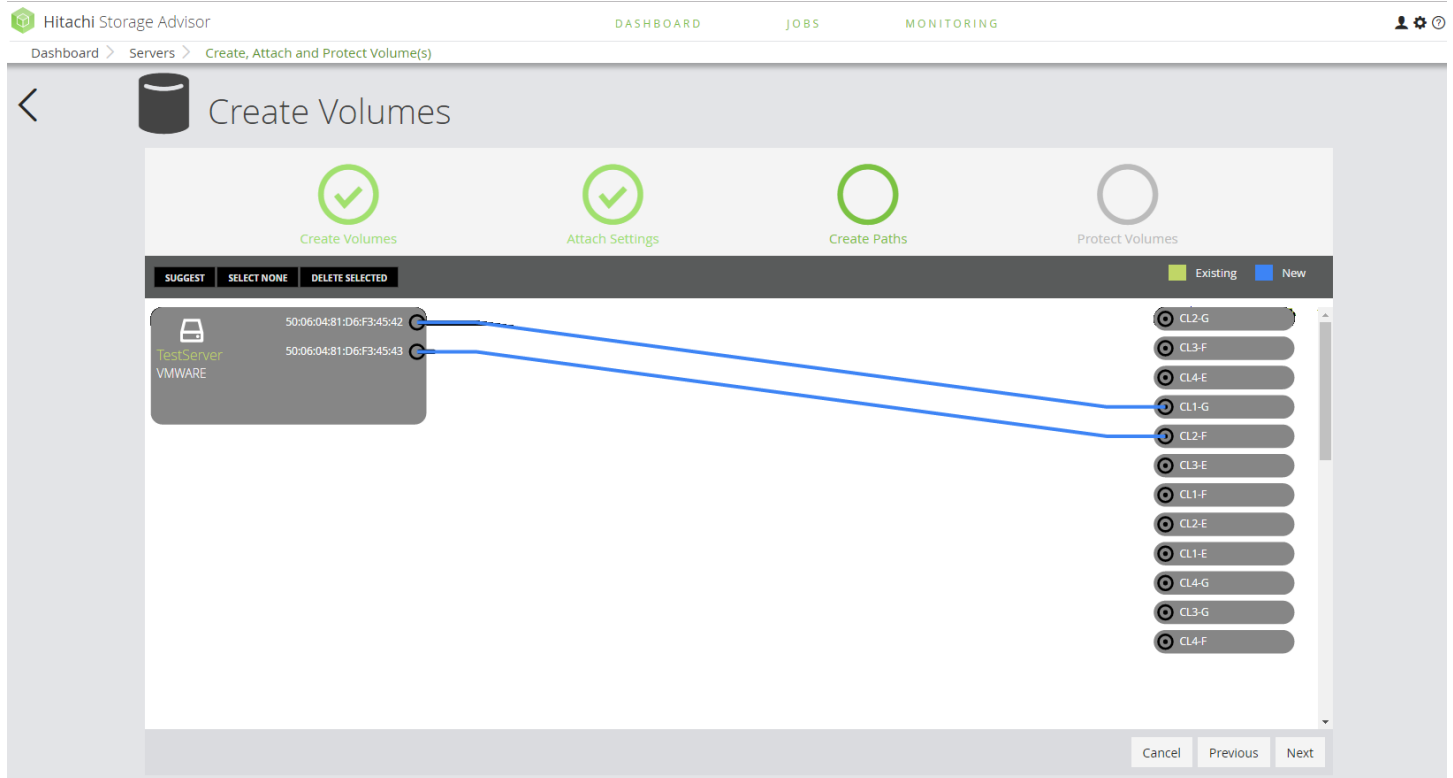
Autopath: Port Assignment

With pools having been created, we are ready to create volumes, attach these volumes to servers, and protect these volumes. The Hitachi Storage Advisor workflow to accomplish these end-to-end provisioning and protection operations can be accessed from the host-managed object on the HSA dashboard.

The complexity of SAN fabric management and server-based provisioning falls into this area of port assignment and zoning. It is imperative that the complexity be masked here as much as possible in order to deliver on simplifying the overall user experience.

Clicking next on the Create Volume dialog as shown in Figure 3 will have the port assignment auto selected based on port count utilization.

Figure 3. Port Assignment: Create Volume Dialog



WYDS: What's in the Attach Volume Logic?

When attaching one or more volumes to one or more servers, the user can either specify the ports to be selected or let Hitachi Storage Advisor autoselect ports and worldwide names (WWNs) for attach. When autoselecting ports, a port selection algorithm is used to decide which storage ports are to be used to present a volume to a target server or server cluster. There are several criteria for selecting ports for attach:

- Port type: Fibre Channel ports are used for attaching a volume.
- Port utilization.
- Number of paths.
- Reusing existing ports.
- Redundancy via fabrics, even and odd storage ports.

HSA will attempt to select multiple ports for redundancy. The storage array itself is split into an odd side (CL1-A, CL3-A, and so forth) and an even side (CL2-A, CL4-A, and so forth), where each side is on its own controller. When selecting ports, the number of ports HSA will attempt to select is twice the number of fabrics. In an isolated storage environment, the best effort would be to use two storage ports. The ports are duplicated to take advantage of the redundancy offered by the storage controllers.

For each fabric, HSA uses the following algorithm:

- If a path(s) for the server already exists, HSA will reuse it by ensuring that the intended host mode and host mode option match the existing path.
- If existing paths cannot be reused, HSA proceeds to select the ports.
 - HSA will try to select one odd and one even storage port. Otherwise, HSA will select the available port that has the least LUNs attached. In the future, HSA will select ports by I/O throughput, instead of by the least number of LUNs attached.

If no server host bus adapter (HBA) or WWNs are specified, HSA uses all WWNs for a server to create a path.

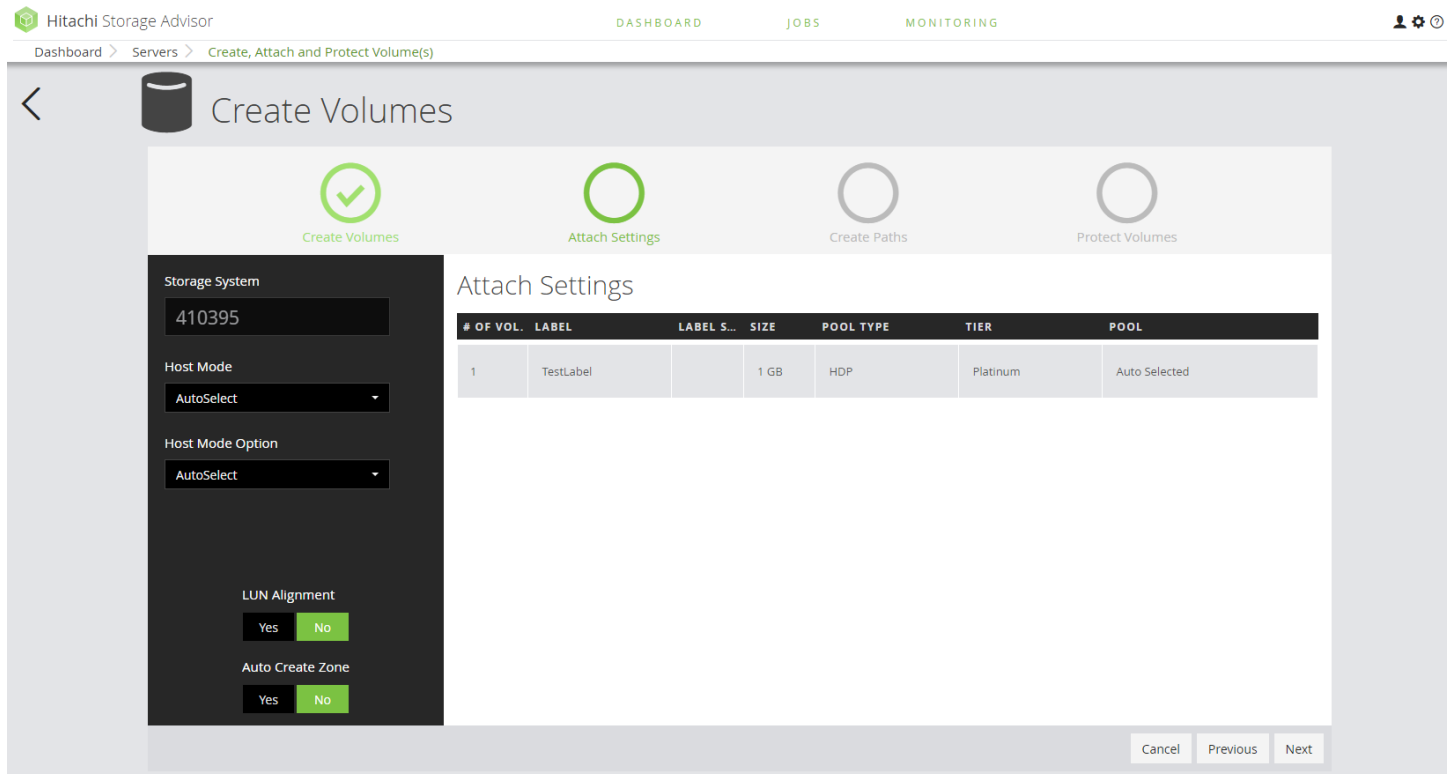
Additionally, logical unit number (LUN) IDs are autoassigned when attaching a volume to a server. HSA scans all available LUN IDs and picks the first available LUN ID that can be used. If an end user requests LUN alignment to be true, HSA ensures that the same LUN ID is used for all the paths that are created for the volume on one or more servers.

Autozone: Zone Creation

To complete the exercise of attaching the volumes to a host server, a SAN fabric zone has to be created and mapped. Clicking the “Yes” button for “Auto Create Zone” provides a much-needed point of automation. (See Figure 4.)

A prerequisite to using this feature is to have already brought a fabric switch on board, allowing it to be accessed from the Hitachi Storage Advisor dashboard. The OpenStack Cinder driver API is used to set the zone between the server and auto-port assignment. HSA internally hosts the Open Stack Cinder driver developed by Brocade and Cisco internally in the HSA virtual appliance. The end user does not have to provide an OpenStack environment. What could be simpler?

Figure 4. Zone Creation: Create Volume(s) Dialog



WYDS: What's in the Attach Volume Logic?

After selecting storage ports and server HBA ports to use to attach the LUN, HSA can also create Fibre Channel zones, provided that the principal Fibre Channel switches are added to the HSA inventory. For each server-storage port pair, HSA will ensure:

- Both ports are visible on a fabric.
- Will reuse an existing, matching zone or create a new zone if no such zone exists.

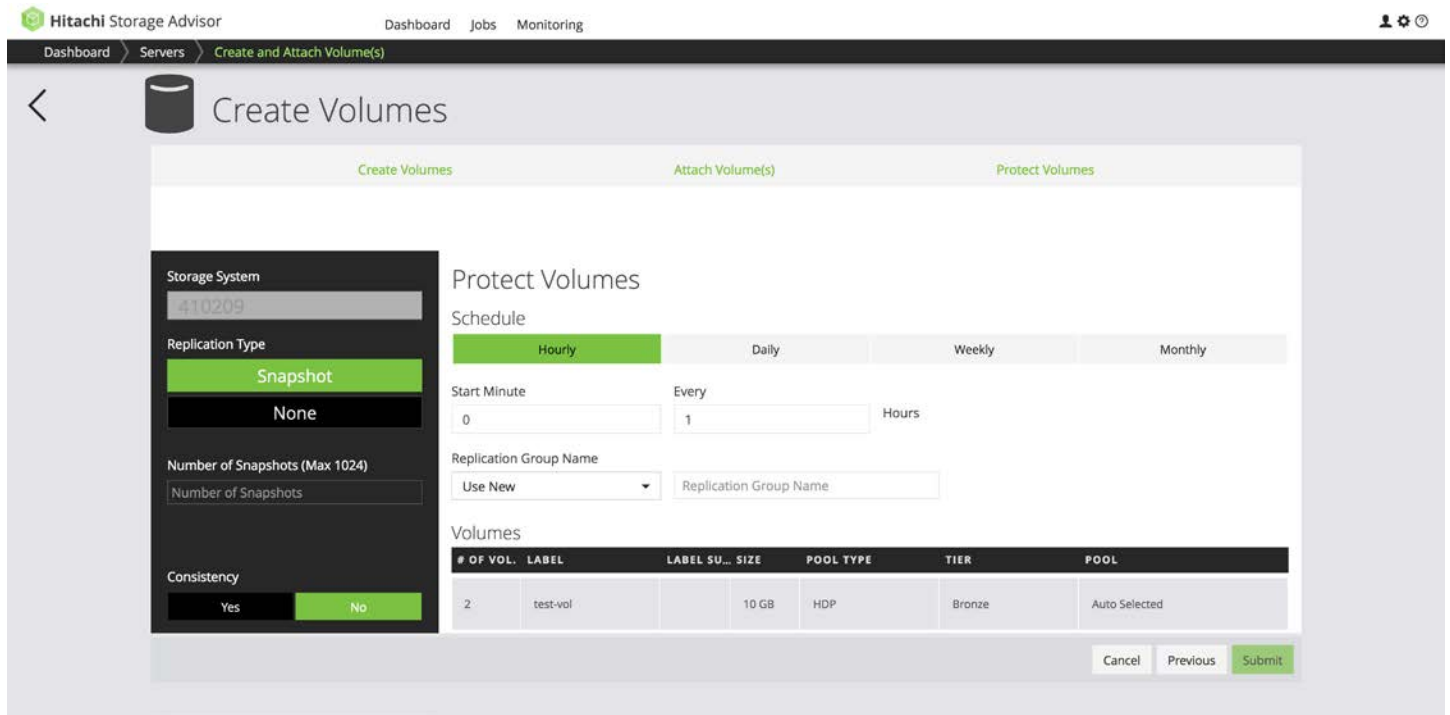
HSA internally calls the OpenStack Cinder driver API using provided credentials to manage zones for any principal switch added to the inventory against the active zone set.

Complete the Job

And finally, protecting these newly created and attached volumes can be accomplished in the same end-to-end workflow. No other Hitachi storage configuration management interface has the ability to create, attach and protect volumes in a single workflow.

The option to set up snapshot schedule is available in create, attach and protect workflow actions, as well as when protecting an existing volume (see Figure 5). The option to clone now, on the other hand, is only available when protecting an existing attached volume.

Figure 5. Protection: Create Volume(s) Dialog



Every time a snapshot schedule runs, Hitachi Storage Advisor completes the following:

1. One secondary snapshot volume is created for each protected volume.
2. New snapshot group is created at every scheduled run and each protected volume and corresponding snapshot (secondary) volume is added to the snapshot group.
3. If consistency flag was checked, then the flag is set at the snapshot group.

Autoabstract File-Pool Creation

With support for the Hitachi Virtual Storage Platform Gx00 series NAS modules, HSA can create file pools while provisioning or abstracting all the underlying block storage resources using field-recommended practices. When using the Hitachi NAS system management unit (SMU) before creating a file pool, the user would be expected to have:

1. Created parity groups.
2. Created block pools.
3. Carved volumes.
4. Attached ports to volumes.
5. Enabled system drives.

With the HSA v2.x unified management release, the user will need to simply select the storage tier and parity group capacity from the options available, to create a file pool (as shown in Figure 6).

Figure 6. Select Capacity for New File Pool

The screenshot displays the 'Create File Pool' configuration page in Hitachi Storage Advisor. On the left sidebar, the storage system is identified as 410500. A circular gauge indicates a total capacity of 37.74 TB. Below the gauge, capacity details are listed: 13 TB Physical Block Pool Capacity, 25 TB Estimated File Over Commit Capacity, File Over Commit Capacity at 200%, and Estimated Max Potential at 0. The main content area, titled 'Select Capacity for New File Pool', offers five storage tier options. The 'Silver' tier (12.58 TB Available) is currently selected, with a green bar and a slider set to 12.58 TB. A warning message states, 'Platinum must be selected in tiered file pool creation'. At the bottom of the tier selection area, there are two sliders for 'Utilization Threshold (Low) %70' and 'Utilization Threshold (High) %80'. 'Cancel' and 'Select' buttons are located at the bottom right of the interface.

WYDS: What's in the File-Pool Creation Logic?

Automating the multi-step file pool creation workflow, HSA has simplified the process by adopting the following practices to create the block pools and system drives:

1. Initial Block Pool Creation

HSA followed recommended block storage based best practices for creating block pools with RAID-6, enforcing one of the following:

- a. Parity group restriction: having a minimum of four parity groups when using SAS or nearline-SAS (NL-SAS) disks or one parity group when using SSD/FMD-DC2 disks, provided all disk types are homogeneous.
- b. Disk count restriction: using a minimum of 48 disks for bronze tier, 24 disks for silver and gold tiers, 8 disks for platinum tier.
- c. If SSD/FMD-DC2 users wish to use RAID-5, HSA provides a warning indicating the possibility of a double disk failure during parity group creation.
- d. When the user chooses more than one tier and one of those contains SSD/FMD-DC2 drives, HSA automatically creates two Hitachi Dynamic Provisioning (HDP) pools with the faster media as Tier 0 (platinum); Hitachi Dynamic Tiering (HDT) pools are not supported at this time.

2. System Drive Sizing and Cardinality

HSA computes the number of system drives and their sizes based on the disk types and their corresponding estimated IOPS (input/output per sec) potential. The number of system drives is computed based on the below guidelines:

- a. The number of system drives (DP-Vols) carved should lie between 4 and 32 and always be an even number.
- b. HSA uses the rule of 8TB per system drive to begin calculations while maintaining a default queue depth of 32 per SAS/NL-SAS system drive and 256 per SSD/FMD-DC2 system drive.
- c. All system drives are thin provisioned, targeting a 200% over-commit value, but the actual over-commit value will be determined during system drive sizing through the following calculations below.
- d. Number of system drive calculation:
 - a. $\langle \text{physical capacity} \rangle = \langle \text{number of parity groups} \rangle * \langle \text{total capacity of one parity group} \rangle$
 - b. $\langle \text{physical capacity} \rangle * 200\% = \langle \text{over-commit capacity target} \rangle$
 - c. $\langle \text{over-commit capacity target} \rangle / \langle \text{capacity per system drive} \rangle = \langle \text{number of system drives} \rangle$
 - i. $\langle \text{cap per system drive} \rangle$: 8TB to start with and goes up by multiples of eight up to 64TB
 - ii. $\langle \text{number of system drives} \rangle$: rounded up to power of two
- e. Tags per drive to estimate maximum I/O potential
 - a. 10K/15K SAS: maximum eight tags
 - b. NL-SAS, SATA: maximum four tags
 - c. SSD: maximum 16 tags
 - d. FMD: maximum 64 tags

- f. Size of system drive calculation:
 - a. $\langle \text{total number of disks} \rangle * \langle \text{max tags per device} \rangle = \langle \text{total tag capability} \rangle$
 - b. $\langle \text{number of system drives} \rangle * \langle \text{queue depth per system drive (32 or 256)} \rangle = \langle \text{IOP potential} \rangle$
 - c. Adjust the $\langle \text{capacity per system drive} \rangle$ to ensure that $\langle \text{total tag capability} \rangle$ is greater than or equal to $\langle \text{IOP potential} \rangle$

Summary

In summary, what Hitachi Data Systems has done begins to address what our users and the IT industry has been asking for: We are reducing the complexity and enabling new levels of ease of use for IT infrastructure management, starting with automating storage and SAN switch fabric operations. The next step will be to apply these similar principles to other aspects of the IT infrastructure, such as compute management.

“What you don't see” (WYDS) in Hitachi Storage Advisor is as important as “what you do see.” Hitachi Storage Advisor abstracts the complexities by hiding and automating management details from end users, resulting in a level of ease of use expected by the industry today.

 **Hitachi Data Systems**

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

