



PS7012

High Performance Storage Array User Manual



RASILIENT CONFIDENTIAL PROPRIETARY INFORMATION This document contains confidential information that is proprietary to Rasiliant or confidential information which Rasiliant has legal duty to protect from disclosure. The information contained in this document may not be disclosed without the express written permission of Rasiliant Systems, Inc.

COPYRIGHT NOTICE

All rights, including copyright, in the content of this manual are owned or controlled by RASILIENT SYSTEMS, INC. and protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior consent of RASILIENT.

TRADEMARKS

All product names or brands mentioned herein are the trademarks of RASILIENT SYSTEMS, INC., its subsidiaries or other respective owners.

DISCLAIMER

This manual provides information regarding set-up and installation of the product herein. Nothing herein may be construed as granting any right or license relating to any intellectual property rights of this manual or product. Unless otherwise provided in the Purchase and Sale Agreement for this product, manufacturer and distributor of this product will not be liable whatsoever relating to the distribution and/or use of this product. In addition, manufacturer and distributor of this product hereby specifically disclaim any express or implied warranties of merchantability, fitness for a particular purpose, or non-infringement of third-party rights in connection with this product.

Manufacturer of this product has the right to change specifications and product descriptions at any time without notice.

Table of Contents

About this Manual	10
CONVENTIONS	11
SAFETY SYMBOLS	12
SAFETY PRECAUTIONS	13
Technician Notes.....	13
Electrostatic Discharge Precautions	13
Rack Warnings.....	14
System Warnings.....	15
Battery Replacement Safety	15
REGULATORY AND INTEGRATION INFORMATION.....	16
Regulatory Compliance Identification Numbers	16
Product Safety Compliance.....	16
Product EMC Compliance.....	16
Communications Commission Notice.....	17
Class A Equipment	17
Declaration of Conformity for Products Marked with the FCC Logo - United States Only.....	18
Power Cords	18
Chapter 1: Introduction	19
AUDIENCE ASSUMPTIONS.....	20
ABOUT THIS GUIDE	20
PACKING CHECKLIST	21
SPECIFICATIONS.....	22
SYSTEM OVERVIEW	23
PRODUCT FEATURES	23
Front View.....	27
Back View	27
System LED Information	28
Chapter 2: Hardware Installation	32
HARDWARE INSTALLATION	33
Check the Installation Environment	33
Unpacking.....	33
Rail Kit Assembly	34
Mounting the Chassis in the Rack	38
Hard Drive Assembly and Installation	44

Front Bezel Installation	45
Controller Canister Installation.....	48
Power Supply Unit (PSU) Installation	50
Cabling and Access to electricity	51
SAS Cable Connections	51
iLink Cable	55
Power Cords	56
Chapter 3: System Configuration	58
System Configuration	59
Initialize the System	59
Option 1 – Use Pre-configured IP Address to Connect to System	59
Option 2 – Configure Private IP Address to Connect to System	61
Rules for Adding Storage	65
Disk Group Parameters.....	66
Configure the System.....	67
Storage Manager Overview.....	67
Using the Setup Wizard.....	74
Step 1 – Network Setup	75
Step 2 – System Initialization	76
Step 3 – Disk Group Setup	78
Step 4 – Virtual Disk Setup	94
Step 5 – Host Setup	100
Step 6 – Present Virtual Disks.....	103
Additional Configuration Options	107
Change Storage System Password	109
Change the Display Refresh Options	110
Change System Health Alert Options	111
Configuring the iSNS Client	112
System Upgrade Options.....	113
Inspection Read Options	115
Language Options	119
Drive Write Cache Options.....	120
System Beacon LED Options.....	122
Performance Data	122
Disk Power Management Status	124
Advanced Logging Mode.....	125
Firmware Upgrade Manager.....	127
Hyper Cache Configuration	128
Configuring Notification.....	129

Chapter 4: Use and Maintain	131
Using and Maintaining the System.....	132
System Status	132
Properties Page.....	132
Hardware Status Page	133
System Log	134
Support Menu.....	135
Performance Data	136
System Health Alert	138
Failure Conditions and Corrective Action	140
Degraded System Availability	140
Turning the System Health Alert Off	140
RAID Controller Failure	141
SAS IO Expander Failure (Lost Connection to Expansion Enclosure).....	141
SAS Cable Disconnect (Lost Connection to Expansion Enclosure).....	142
Drive Failure.....	143
Unhealthy RAID Disk Group.....	143
ECC Memory Error.....	143
Faulty Cache Battery/Charger	144
Network Connection Lost.....	144
Temperature Threshold Exceeded.....	144
Voltage Threshold Exceeded	145
Fan Failure	146
List of Health Alert Notifications:	146
Controller Modules	147
Controller Features.....	147
Storage.....	148
Disk Group Roaming:	148
Disk Orphans:	149
Chapter 5: Online Disk Cloning	150
Proactive Disk Cloning	150
Preparation	150
Cloning procedure	150
Fixing the bad stripes	152
Start cloning	153
Cloning status.....	154
Complete cloning	154
Replace the bad hard drive.....	156
An alternative method to display RAID stripe table	156

Q&A.....	158
Appendix A – RAID Concepts & Cache Configurations	161
RAID Levels Defined.....	161
RAID 0.....	161
RAID 1.....	161
RAID 5.....	162
RAID 6.....	162
RAID 10.....	162
RAID 50.....	163
RAID 60.....	163
Cache Settings Defined.....	164
Write-back Cache Option.....	164
Write Through Cache Option.....	164
Auto Cache Setting.....	165
Appendix B – Using the Command Line Interface (CLI).....	166
Connecting a Terminal to Access the CLI.....	166
Accessing the CLI – Serial Port or SSH Session.....	167
dahelp Command.....	167
Local/Remote Option.....	167
Index of CLI Commands.....	168
CLI Commands.....	170

List of Figures

Figure 1: PS7012 Chassis Front View	27
Figure 2: Back View	27
Figure 3: Enclosure Front LEDs	28
Figure 4: PS7012 RBOD Controller Rear LEDs.....	29
Figure 5: HDD LEDs	30
Figure 6: PS7012 Drive Slot Numbers and LEDs	30
Figure 7: PSU LEDs	31
Figure 8: Step 1	34
Figure 9: Step 2	35
Figure 10: Step 3	36
Figure 11: Step 4	37
Figure 12: Step 5-1	38
Figure 13: Step 5-2	39
Figure 14: Step 6	40
Figure 15: Step 7	41
Figure 16: Step 8	42
Figure 17: Step 9	43
Figure 18: The SAS IO cable connection - PS312e2.....	51
Figure 19: SAS cable connection to PS390e	52
Figure 20: Connecting a SAS cable	53
Figure 21: SAS Cable Alignment Key	54
Figure 22: iLink Cable	55
Figure 23: iLink Ethernet Ports expanded view	55
Figure 24: Secure Power Cords.....	57
Figure 25: CLI Setup Wizard	63
Figure 26: Storage Manager Login	67
Figure 27:Management Console Layout	68
Figure 28: Console Banner.....	69
Figure 29: Navigation Tabs	69
Figure 30: Setup Wizard	70
Figure 31: Navigation Tree.....	71
Figure 32: Work Area	72
Figure 33: System Status Icon	73
Figure 34: Network Setup	75

Figure 35: System Initialization--Step 1	77
Figure 36: System Initialization -- Step 2	77
Figure 37: Disk Group Setup	79
Figure 38: Create Disk Group	80
Figure 39: DG Name, Disks, Controller, Protection level.....	81
Figure 40: Create DG -- Advanced Features	84
Figure 41: Creating DG -- in progress	85
Figure 42: Disk Group Properties	86
Figure 43: RAID Build Progress.....	87
Figure 44: Disk Group Discovery	91
Figure 45: Covert Orphan Disk(s) to Spare(s)	92
Figure 46: Expand a Disk Group	93
Figure 47: Virtual Disk Setup	94
Figure 48: Create Virtual Disk	95
Figure 49: Virtual Disk Parameters.....	96
Figure 50: Virtual Disk Properties	97
Figure 51: Extending VD Capacity	99
Figure 52: Host Properties Page	100
Figure 53: Adding a Host -- Step 1 & 2.....	101
Figure 54: Adding a Host - Steps 3 & 4.....	102
Figure 55: Host Properties	102
Figure 56: Chose a VD to present	104
Figure 57: VD Presentation	105
Figure 58: Present Virtual Disk Page	105
Figure 59: Current Presentation.....	106
Figure 60: Options Top part of page.....	108
Figure 61: Options Bottom half of page	109
Figure 62: Change Storage System Password.....	110
Figure 63: Display Refresh Options	111
Figure 64: System Health Alert Options	112
Figure 65: iSNS Configuration	113
Figure 66: Upgrade Manager	115
Figure 67: Options Tab.....	117
Figure 68: Inspection Read Interface.....	118
Figure 69: Inspection Read Progress	119
Figure 70: Language Option	120

Figure 71: Drive Write Cache Options	121
Figure 72: System Beacon LED Options	122
Figure 73: Performance Data	123
Figure 74: Start Recording Data	124
Figure 75: Disk Power Management Status	125
Figure 76: Advanced Logging Mode	126
Figure 77: Firmware Upgrade	127
Figure 78: Hyper Cache Configuration	128
Figure 79: Email Notification	129
Figure 80: Adding Email Addresses	130
Figure 81: Storage System Properties	132
Figure 82: Hardware Status	134
Figure 83: System Log	135
Figure 84: Support Menu	136
Figure 85: Performance Monitor Page	137
Figure 86: VD Data Rate Display	138
Figure 87: System Health Alerts Page	139
Figure 88: Find the DG in Disk Group Properties	153
Figure 89: Start Cloning	153
Figure 90: RAID Cloning is ongoing	154
Figure 91: Disk RAID Cloning in Progress Page	154
Figure 92: Complete cloning	155
Figure 93: New Disk Group	155
Figure 94: The status of the source disk is orphan	155
Figure 95: RAID Stripe Display shown in navigation tree	156
Figure 96: RAID Stripe Display	157
Figure 97: Click Fix button and then click Clone button	157
Figure 98: System is initially building	158
Figure 99: Error information page	158
Figure 100: Source disk turns into a faulty disk during cloning	159
Figure 101: Target disk turns into a faulty disk during cloning	159
Figure 102: No. 1 disk turns into a faulty disk during cloning	160

List of Tables

Table 1: Safety Compliance	16
Table 2: EMC Compliance	16
Table 3: Manual Organization	20
Table 4: Packing Checklist	21
Table 5: System Specifications	22
Table 6: Enclosure front LEDs Functions	28
Table 7: HDD LED DEFINITION.....	30
Table 8: PSU LED DEFINITIONS.....	31
Table 9: Disk Group Parameters	66
Table 10: Disk Group Parameters	82
Table 11: Inspection Read Action.....	116
Table 12: CLI command list	168

About this Manual

Conventions

Safety Symbols

Safety Precautions

Regulatory and Integration Information

ABOUT THIS MANUAL

CONVENTIONS

To make sure that you perform certain tasks properly, take note of the following symbols used throughout this manual.



Warning:

Provides information to prevent injury in the process of completing a task.



Caution:

Provides information to prevent damage to the components in the process of completing a task.



Important:

Provides information required for completing a task.



Note:

Provides tips to aid in completing a task.

SAFETY SYMBOLS

The following symbols are placed on some components of the system to alert the user to potential hazards.



WARNING: Electric Shock Hazard - To reduce risk of injury from electric shock hazards; do not open this component.



WARNING: Contains No User or Field Serviceable Parts - To reduce risk of injury from electric shock hazards; do not open this component.



WARNING: Hot Surface or Component - To reduce risk of injury from a hot component; allow the surface to cool before touching.



WARNING: Insert Network Interface Only - Any receptacle (e.g. RJ45) marked with this symbol indicates a network interface connection. To reduce the risk of electric shock, fire or damage to equipment, do not plug telephone or telecommunications connectors into this receptacle.



WARNING: This symbol, on power supplies or systems, that the equipment is supplied by multiple sources of power. To reduce the risk of injury from electric shock, remove all power cords to completely power down the system.



WARNING: This symbol indicates that the component exceeds the recommended weight for one individual to handle safely. To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manual material handling.

SAFETY PRECAUTIONS



Technician Notes

- Only authorized technicians should attempt to repair this equipment.
- Before installing this system, carefully read all the manuals included with the system.
- All repair procedures allow only module replacement. Because of the complexity of the individual boards and sub-assemblies, no one should attempt to make repairs at the component level or make modifications to any printed wiring board. Improper repairs can create a safety hazard.
- To reduce the risk of personal injury from electric shock and hazardous energy levels, do not exceed the level of repairs specified in these procedures.
- The system is designed to be electrically grounded. To ensure proper operation, plug the AC power cord into a properly grounded AC outlet only.



Electrostatic Discharge Precautions

- Electrostatic discharge (ESD) can damage static sensitive devices or micro circuitry. Proper packaging and grounding techniques are required to prevent damage.
- Keep electrostatic-sensitive parts in their containers until they arrive at a static free work area.
- Use a wrist strap connected to the work surface as well as properly grounded tools and equipment.
- Keep the area free of nonconductive materials such as ordinary plastic tools and foam packaging.
- Avoid touching pins, leads, or circuitry.
- Always place drives with printed circuit board (PCB) assembly-side down.
- Grasp cards and boards by the edges. Hold drives by the frame.

Avoid touching the solder joints or pins.

- If you need to lay the device down while it is out of the antistatic bag, lay it on the antistatic bag. Before picking it up again, touch the antistatic bag and the metal frame of the system unit at the same time.



Rack Warnings

- If you plan to rack mount this product, please follow the rack manufacturer's safety instructions.
- Install the enclosure only in a rack that has been properly secured in an area with suitable environmental conditions.
- Have someone assist you during physical installation.
- To properly ventilate the system, you must provide at least 7.6 cm of clearance at the front and back of the system.
- To reduce the risk of personal injury or damage to equipment, always ensure that the rack is adequately stabilized prior to extending a component outside the rack. A rack may become unstable if more than one component is extended. Extend only one at a time.
- Do not stand or step on any components in the rack.
- If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient. Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (T_{ma}) specified by the manufacturer.
- Do not overload the AC power supply branch circuit that provides power to the rack. Observe extension cable and power strip ratings. Ensure that the total ampere rating of all equipment plugged into the extension cable or power strip does not exceed 80 percent of the ampere ratings limit for the extension cable or power strip.



System Warnings

- Avoid dust, humidity, and extreme temperatures; place the system on a stable surface.
- To reduce the risk of personal injury from hot surfaces, allow the hot-plug disk modules and other system modules to cool before touching them.
- To reduce the risk of electric shock or damage to the equipment, do not disable the power cord grounding plug. The grounding plug is an important safety feature.
- Ensure the power cord is inserted into a grounded electrical outlet that is easily accessible at all times. Unplug the power cord from the power supply module to shut off power to the equipment.
- Protect the storage system from power fluctuations and temporary power interruptions with a regulating uninterruptible power supply (UPS). This device protects the hardware from damage caused by power surges and voltage spikes and keeps the system operational during a power failure.
- The storage system must always be operated with all hot plug modules installed or slot covers in place to ensure proper cooling.
- Route power cords so that they will not be walked on or pinched by items placed upon or against them. Pay particular attention to the plug, electrical outlet, and the point where the cords exit from the product.



Battery Replacement Safety

- This storage is provided with an internal Lithium battery pack. There is a danger of explosion and risk of personal injury if the battery is incorrectly replaced or mistreated.

REGULATORY AND INTEGRATION INFORMATION

Regulatory Compliance Identification Numbers

For the purpose of regulatory compliance certification and identification, this system is assigned a serial number. This system serial number can be found on the product label, along with the required approval markings and information. When requesting certification information for this product, always refer to this serial number. This serial number should not be confused with the marketing name or model number.

Product Regulatory Compliance

Product Safety Compliance

This system complies with the following safety requirements:

Table 1: Safety Compliance

UL/IEC/EN 60950-1	Safety of Information Technology Equipment
-------------------	--

Worldwide Safety approvals can be supplied upon request. Please contact your sales representative for approvals.

Product EMC Compliance

This product has been assembled from components that comply with the following electromagnetic compatibility (EMC) regulations.

Table 2: EMC Compliance

Radiated Emission:	CISPR 22 (FCC, EN 55022, VCCI, MIC, BSMI, C-Tic) Class A
Conducted Emission:	CISPR 22 (FCC, EN 55022, VCCI, MIC, BSMI, C-Tic) Class A EN61000-3-2 Harmonic Emissions EN61000-3-3 Flicker Emissions
Immunity:	EN55024

	<p>IEC 61000-4-3; Radiated Immunity</p> <ul style="list-style-type: none"> • IEC 61000-4-4; EFT • IEC 61000-4-5; Surge • IEC 61000-4-6; Conducted Immunity • IEC 61000-4-11; Voltage Variations and Dips
--	--

Communications Commission Notice

Part 15 of the Federal Communications Commission (FCC) Rules and Regulations has established Radio Frequency (RF) emission limits to provide an interference-free radio frequency spectrum. Many electronic devices, including computers, generate RF energy incidental to their intended function and are, therefore, covered by these rules. These rules place computers and related peripheral devices into two classes, A and B, depending upon their intended installation. Class A devices are those that may reasonably be expected to be installed in a business or commercial environment. Class B devices are those that may reasonably be expected to be installed in a residential environment (for example, personal computers). The FCC requires devices in both classes to bear a label indicating the interference potential of the device, as well as additional operating instruction for the user.

The rating label on the device shows which class (A or B) the equipment falls into. Class A devices do not have an FCC logo or FCC ID on the label. Class B devices have an FCC logo or FCC ID on the label. Once the class of the device is determined, refer to the following corresponding statement.

Class A Equipment

This equipment has been assembled with components that comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

Declaration of Conformity for Products Marked with the FCC Logo - United States Only

This device complies with Part 15 of the FCC Rules Operation and is subject to the following two conditions: (1) this device may not cause harmful interference that may cause undesired operation. For questions regarding your product, please contact your sales representative.

To identify this product, refer to the Part, Series, or Model number found on the product.

Power Cords

The power cord set included in the system meets the requirements for use in the country where the system was purchased. If this system is to be used in another country, contact your sales representative to purchase a power cord that is approved for use in that country.

The power cord must be rated for the product and for the voltage and current marked on the product's electrical ratings label. The voltage and current rating of the cord should be greater than the voltage and current rating marked on the product. In addition, the cross-sectional area of the wires must be a minimum of 1.00mm² or 18AWG, and the length of the cords must be between 1.8m (6 feet) and 3.6m (12 feet). If you have questions about the type of power cord to use, contact your sales representative.

The following statement applies only to rack-installed products that are GS-Marked: This equipment is not intended for use at workplaces with visual display units, in accordance with §2 of the German ordinance for workplaces with visual display units

Chapter 1: Introduction

Audience Assumptions

About This Guide

Packing Checklist

Specifications

Product Features

System Overview

INTRODUCTION

AUDIENCE ASSUMPTIONS

This manual assumes that you are a service technician or network administrator familiar with computer hardware, data storage, and network administration terminology and tasks.

ABOUT THIS GUIDE

This user guide provides step by step instructions on how to install, use and maintain the PS7012 SAN RAID Array. This manual is generally organized as follows:

Table 3: Manual Organization

Introduction	General introduction to the PS7012 and its components.
Preparation	What you need to do before you start.
Installation	Detailed description of each hardware module and instructions for installation and removal.
Configuration	Initializing your system, setting up your system on a network, configuring your storage, and connecting hosts.
Operation	Ongoing usage and maintenance information provided in a reference format and organized by major system component.

PACKING CHECKLIST

Make sure you have all the components that shipped with your system. If any item is damaged or missing, please contact your sales representative for a replacement. The PS7012 is shipped with the following:

Table 4: Packing Checklist

Chassis	2U/12 bay rack-mounted chassis
Controllers Module(s)	Two RAID controller modules
Power Supply Modules	Two power supply modules.
Chassis Brackets	Two chassis brackets for rack installation.
Hard Disk Drives (HDDs)	Up to twelve hot plug 3.5" SAS drive canisters.
Power Cords	Two power cords

SPECIFICATIONS

Table 5: System Specifications

CPU	Intel Xeon Broadwell
Controllers	Dual Redundant Active/Active mode
Memory	Up to 64GB DDR4 ECC (Per Controller)
Storage	Up to 12x 3.5", up to 18TB SAS 3.0 (12 Gbps) drives, (main chassis) Supports up to 108 drives with external JBODs
Cache Type	Mirrored/Persistent
Battery Backup (BBU)	Controllers have protection for memory contents if power loss occurs. Memory data is backed up to an on-board SSD. No time limit for power loss.
Operating Temperature	10°C - 35°C
Networking	2x 1GbE, 2x 10GbE RJ45 per controller Optional 2x SFP+ Fiber ports
Input/Output	2x USB, 1x Serial COM, 2x SAS
Dimensions (WxHxD)	17.6"x 3.5"x 20.9"/ 447 x 89 x 531mm
Weight	62.5 lbs/28.4 kg
Power	89-264VAC, 47Hz-63Hz input, via IEC C-14
Watts	800W, 80plus Gold Level, 1+1 redundant, hot swap PSU

SYSTEM OVERVIEW

PS7012 is a low cost, high bandwidth and high availability RAID array. Unlike most arrays designed to support a single controller with proprietary architectures, the PS7012 can be configured as a single or dual controller system. It is built around a flexible and open architecture utilizing dual active-active storage controllers, and a custom tuned Linux operating system. It is offered in a 2U rack mount form factor and can be populated with 12 hot plug hard disk drives.

The PS7012 offers direct attach iSCSI SAN host connectivity. Each controller module on the PS7012 comes with two Gigabit Ethernet ports and dual 10Gbps ports for higher performance and connectivity. Each PS7012 controller equipped with a battery backed unit (BBU), and redundant power supplies.

The capacity of the PS7012 can be expanded by connecting multiple 2U/12 bay PixelStor312e2, or 90-bay PS390e, or 92-bay PS392e. These can be daisy chained to a PS7012 via fast 48 Gbps SAS connectors for low cost high performance scalability. Further details about PS312e2, PS390e, and PS392e are covered in their respective User Manual.

PRODUCT FEATURES

Unparalleled Flexibility – The PS7012 offers iSCSI SAN host connectivity within the same box. This functionality makes the PS7012 one of the most flexible storage building blocks on the market today. The flexibility inherent in the PS7012 allows customers to address diverse needs in the storage space without the need for multiple specialized platforms.

Zero Down-Time Architecture - In the PS7012, every active component, including drives, controllers, fan trays, and power supplies, are hot swappable without interruption in service or loss of data. Its active-active dual controllers, sophisticated mirrored and battery backed cache and path failover capabilities give the PS7012 best-in-class system fault tolerance.

Active-Active Controllers & Transparent Fail-over & Fail-back - Fast throughput and data availability are paramount for dedicated storage systems. This is why PS7012 is designed with active-active controllers to improve both performance and data availability.

Each Disk Group has an assigned “primary” controller to take advantage of the fastest route to the drives. So, optimal performance is achieved when traffic goes through both controllers to their associated Disk Groups. For availability, both controllers can serve any Disk Groups within the system. When there is a controller failure Disk Groups assigned to the failed controller will automatically “migrate” to the surviving controller without user intervention. When the failed controller is replaced, the systems will automatically fail-back to its original state. In PS7012, the relationship between two controllers is tight. Each controller has the precise status of the other, which results in quick and accurate actions when failure happens.

Mirrored Cache - With Dual controllers, the “auto cache” option puts a mirrored copy of the write cache on the redundant controller. The data is fully protected. If there is a failure in any controller, the system flushes the mirrored copy and automatically switches to the “write-through” setting. This is the best compromise between performance and protection.

Multi-Path IO and Cache Coherency - The PS7012 makes no assumption regarding the behavior of the hosts since they can switch from one path to the other. This is normally referred as “multi-path IO” – a technology to ensure the end-to-end high availability for path failure. RASILIANT Systems patent-pending technology supports path failure with our cache coherency features. In the extreme case, the host can send commands to either controller randomly.

Dedicated SATA Channels - PS7012’s multi-path design provides a dedicated drive channel from each controller. The elimination of inter-drive dependencies enhances performance. The independent controller access provides additional data availability.

Disk Group Roaming - In the PS7012, the concept of Disk Groups is introduced, which represents a set of drives that share similar characteristics. The Disk Group configuration, including RAID, and volume, roams with the drives. Once inserted, the Disk Group can be discovered and the service can start with a mouse click.

Within the enclosure, the drives can also be moved around. The roaming operation can be carried out while the system is on-line.

Users can take advantage of this flexibility to arrange their storage for different on-line and off-line tasks. This feature provides a basic and inexpensive way to upgrade a system and retain your data or move drives to another system for basic system recovery.

Capacity Expansion – The PS7012 has the best capacity expansion capability on the market. The 2U/12 bay controller unit can be expanded with up to eight 2U/12 bay PS312e2 expansion enclosures for a total capacity of 108 drives. The PS390e is 4U/90-bay, and PS392e is 5U/92-bay dual controller expansion enclosure. These enclosures can be hot added to the system via the PS7012 48Gbps SAS connectors. The capacity expansion feature simplifies the task of adding additional storage capacity. You can expand Disk Groups across enclosures. The user can expand the LUN capacity on-line within the Disk Group, or expand the Disk Group by adding disks from the pool of global spare disks.

Software RAID - The PS7012 architecture focuses on the critical performance elements of protocol processing and data movement. The powerful multi-core server processor, DDR4 memory, and server chipset present a superior data mover architecture.

Global Spares - In PS7012, any drives not assigned to a Disk Group are automatically available as global hot spares. If a drive fails, the RAID controller can automatically rebuild its data on the spare drive without requiring intervention by the administrator. The rebuild operation occurs in the background while the controller processes normal operations.

LUN Masking - Authentication and Access Control List (ACL) - Enterprise data must be secured and closely guarded. In PS7012, LUNs are only exposed to the hosts that have authority to access them. This is generally referred to as LUN masking. In addition, CHAP authentication is used for each iSCSI session login. All management access is also protected with the access control list, which filters out the machines that are not on the list.

Robust Enclosure - Both the enclosure and the drive carriers are designed for operation without IO errors in various temperature and vibration conditions. The whole enclosure is built in one piece. Extensive thermal and vibration testing has been conducted on the enclosure by itself and in a densely populated rack. Special attention has been taken to recover the data path failure due to signals stuck at zero or one.

Simple, Intuitive Configuration - The web based management utility called the PixelStor Storage Manager has been designed to be intuitive and extremely easy to use. It requires little to no training. With the PixelStor Storage Manager, a system administrator can set up a system in 6 simple steps. The utility allows you to:

- Assign System IP address(s)
- Select and name multiple drive groups and establish different RAID levels for each group.
- Create and size virtual drives for each drive group. The system supports up to 256 LUNs.
- Create a list of available hosts with WWN's and IP addresses
- Present Volumes to Hosts
- Establish access rights to each virtual drive with the option to enable powerful security policies to each virtual disk.
- Add capacity,
- Check system status,
- Set email alerts in the event of system problems.

System Management and Full Monitoring Capability - RASILIENT Systems is committed to developing storage arrays that reduce administrative complexity and management costs.

PS7012 comes with powerful remote management software that dramatically simplifies setup and monitoring of storage volumes. The robust self-monitoring technology and automatic notification service maximizes storage availability while minimizing the pressure on system administrators. In the event of a hardware failure, PS7012 can automatically notify you via email. The user interface also provides monitoring for each hardware module. The log and all system information are kept in persistent memory, which can be recovered after reboot. For remote diagnostics, the administrator can download the system log to a PC / Laptop.

Front View




Figure 1: PS7012 Chassis Front View

Back View



Figure 2: Back View

System LED Information

PS7012 Enclosure Front LEDs:

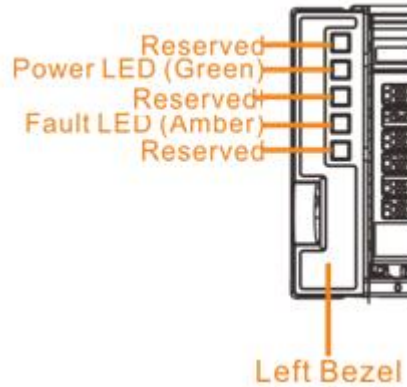


Figure 3: Enclosure Front LEDs

Table 6: Enclosure front LEDs Functions

LED Name	LED Color	LED Status	Description
Power LED	Green	ON	Enclosure power is on
		OFF	Enclosure power is off
Fault LED	Amber	ON	Fault conditions exist in the enclosure
		OFF	Normal operation

RBOD Controller Rear LEDs:




Figure 4: PS7012 RBOD Controller Rear LEDs

Hard Drive LEDs:




Figure 5: HDD LEDs

<div style="display: flex; flex-direction: column; align-items: center;"> ● ● ● </div>	DRIVE 0	●	DRIVE 1	●	DRIVE 2	●	DRIVE 3	●
	DRIVE 4	●	DRIVE 5	●	DRIVE 6	●	DRIVE 7	●
	DRIVE 8	●	DRIVE 9	●	DRIVE 10	●	DRIVE 11	●

Figure 6: PS7012 Drive Slot Numbers and LEDs

Table 7: HDD LED DEFINITION

LED Name	LED Color	LED Status	Description
Activity LED	Green	ON	This LED is controlled by the HDD itself; refer to the HDD spec for the LED definition.
		OFF	
		Blink	
Fault LED	Amber	ON	Fault conditions exist on this HDD
		OFF	Normal operation

PSU LED:

Figure below shows PSU front view. There are three LEDs on the PSU face. two LEDs are for AC IN and DC output status. One is for PSU fault.





Figure 7: PSU LEDs

Table 8: PSU LED DEFINITIONS

LED Name	LED Color	LED Status	Description
AC input status LED	Green	ON	The PSU input voltage is within specification
		OFF	The input voltage is not OK
DC Ready LED	Green	ON	PSU 12V and 5V outputs are within specification
		OFF	PSU output is out of specification
Fault LED	Red	ON	Fault conditions exit on the PSU
		OFF	Normal operation

Chapter 2: Hardware Installation

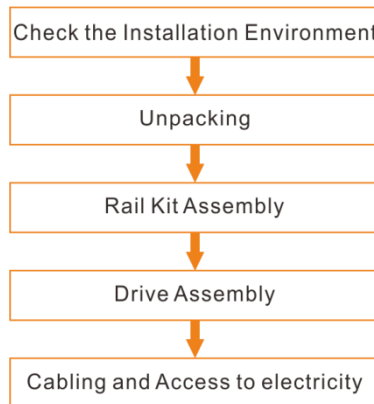


- installation Environment
- Unpacking
- Rail Kit Assembly
- Mounting the System onto a Rack
- Hard Drive Assembly and Installation
- front Bezel Installation
- Controller Canister Installation
- Power Supply Unit (PSU) installation
- Cabling and Access to Electricity
- SAS Cables
- iLink Cables
- Power Cords

HARDWARE INSTALLATION

Check the Installation Environment

- Check the cabinet or mounting bracket: ensure that rack or mounting bracket should be able to hold eight times of the total weight of equipment.
- Check the cooling space: ensure thermal dissipation space has been enough.
- Check another installation environment requirements, please refer to the preface of this user manual.



Unpacking

Follow the following steps to unpack the product:

1. Inspect the package to make sure that it has not been damaged. If you find damage, contact the local authorized Rasilient agent.
2. Open the package and check whether any item on the packing list is missing.
3. Take all items out of the package and check whether any item is in poor condition.
4. After unpacking, please be sure to examine and check the parts in good condition; after the check, suggest to install equipment or put to the cabinet.

Rail Kit Assembly

Step 1: Attach the spring as below picture

Attach one spring to these two metal studs rail as seen below. Redo for the right side rail




Figure 8: Step 1

Step 2: Put the rail into the rack as below picture. (Some compression will be required to overcome the spring force)The Rail with the metal studs (default configuration) will support the round holes on the rack post.

- 1) Open the metal latches.
- 2) Attach the rail ending with inner surface of Rack front and rear post.
- 3) Rotate back the latches.




Figure 9: Step 2

Step 3: Lock the S/S latches to the close position. The shipping bag contains optional hardware for mounting conditions like square holes or threaded holes.


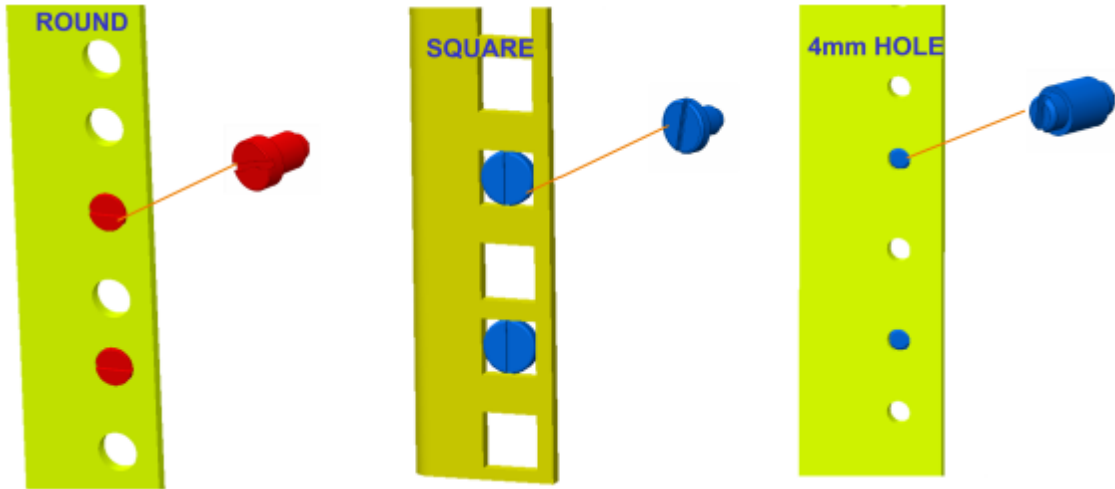


Figure 10: Step 3

Determining the different hole types on the rack post:

- The rail kit will work with round, square and 4mm hole racks.
- The rail kit comes configured with hardware for use in a round hole rack.
- If a square hole rack is being used, remove the metal stud on rail kit, and replace with the larger metal stud supplied in the accessory bag.
- If a 4mm hole rack is being used, remove the metal stud on rail kit, and replace with the smaller stud supplied in the accessory bag.



Step 4: Assemble the other rail with rack post as well as seen on the picture below.




Figure 11: Step 4

Mounting the Chassis in the Rack

Step 5: Put the chassis on the Rail Kit and push it into the rack, and then fix one (or two) black M5 pan-head screw on each rack front post.




Figure 12: Step 5-1




Figure 13: Step 5-2

Step 6: Fix two black M5 pan-head screws on each rack rear post.




Figure 14: Step 6

Step 7: Snap left ear bezel and right ear bezel.




Figure 15: Step 7

Step 8: Open the buckles to buckle the ears of chassis rotation, then use screws to fasten the chassis on the rack.




Figure 16: Step 8

Step 9: Install the bezel on two sides' ear.





Figure 17: Step 9


Hard Drive Assembly and Installation



This section simply introduces the HDD assembly process which may happen on customer side.

Step 1: Install the HDD with HDD carrier.




Step 2: Install the HDD carrier.







Step 3 : Install the front bezel.



Controller Canister Installation

Rotate the PS7012 canister handle with a stop angle and push it into the chassis. Then rotate the canister handle to close.






③ Rotate the Canister handle to close

Power Supply Unit (PSU) Installation

Rotate the PSU handle with a stop angle and push the PSU into the chassis. Then rotate the PSU handle to close.



Cabling and Access to electricity

After the equipment has been installed, please check again that the equipment has been fixed safely. After the check, please wire the equipment and access to electricity.

SAS Cable Connections

This section give a proposal of the SAS IO cable connections to PS312e2 and PS390e JBOD enclosure(s).




Figure 18: The SAS IO cable connection - PS312e2

Connecting to PS390e:




Figure 19: SAS cable connection to PS390e



1. Connect SAS cable from first PS312e2 JBOD controller to PS7012's **right hand side** SAS connector only as shown above. Similarly connect the SAS cable from the right most SAS port of PS390e canister to the right most SAS connector of PS7012 as shown in Figure 19.
2. Connectors on both ends of the SAS cable are universally keyed. You can connect either end of the cable to the Canister or the RAID controller. Orient the SAS cable connector to align with the notch (key) on the controller receptacle. See figures below.
3. To remove the SAS cable, pull the pull-tab to release the cable from the connector on the Canister and the host system.




Figure 20: Connecting a SAS cable




Figure 21: SAS Cable Alignment Key

iLink Cable



The following Ethernet cables must be connected between the two Ethernet ports labeled 6 and 7. Green cable must be connected to top port **6** and red cable must be connected to bottom port **7** as shown below:




Figure 22: iLink Cable





Figure 23: iLink Ethernet Ports expanded view

Power Cords

There is no power button for the system, when there is specific input power is applied to the PSU, the system will power up.



There is no power button for the system, when the input power is cut off or both power cables are plugged out, the system will power off.



When power is on, check all the LEDs are lighting normally. If not, check the cables or

reference [Table 8: PSU LED DEFINITIONS](#)

After connecting the power cords attach these to the cable straps as shown in the illustration below:




Figure 24: Secure Power Cords

Chapter 3: System Configuration

System Configuration

Initialize the System

Configure the System

Rules for Adding Storage

Disk Group Parameters

System Configuration

Once the physical hardware components are installed, you are ready to start configuring your storage. First you will need to set up the PS7012 on your network.



It is always recommended that you set up the PS7012 on protected storage network separate from your data network for security and performance reasons.

Initialize the System

The PS7012 ships with 4 Gigabit Ethernet (GE) ports per system. These are bonded together using RASILIANT's IP redirect technology. RASILIANT offers additional flexibility by offering the ability to set one of these GE link links as a separate out-of-band management port. In addition the system ships with four 10 Gigabit ports per system. These are also bonded together. These can be bonded to the same system IP address as GE ports or a separate IP address (recommended).

The bonding of IP address using Rasiliant's IP redirect technology provides the best performance because the system intelligently balances IO requests across the multiple bonded links. This approach also offers the best fault tolerance because if any link goes down, hosts can continue to access the PS7012 via the surviving links.

With the physical hardware components installed, you are now ready to initialize the system and start configuring your storage. The first step in setting up the PS7012 is to get it ready to connect to your network. There are two options for setting up the PS7012 on your network.

Option 1 – Use Pre-configured IP Address to Connect to System

Each PS7012 comes pre-configured with an IP address. Once you have connected your PS7012 to your network and started the system, you can use the web based

PixelStor Storage Manager included with your system to configure your storage. Using the PixelStor Storage Manager and default IP addresses instead of the CLI is recommended, especially if this is your first time setting up a PS7012 . The PixelStor Storage Manager is designed to get your system up and running in six simple steps.



There will be only one base System IP address regardless of number of physical GE ports or number of controllers.

1. Power-on PS7012, by plugging in AC power cords to both power supplies. The green AC and DC power LEDs on the power supplies will turn on and stay on. The green activity light on each installed disk drive should be on. The PS7012 comes with its own pre-installed system software. When the system is powered on, the system software boots automatically.



When PS7012 is started for the first time, the system may take 3-4 minutes. If the system is restarted during the initialization sequence, the process starts over from the beginning. Please wait for the system to finish its initialization sequence to avoid start-up delays.

2. Move to a workstation with a web browser such as Microsoft Internet Explorer on the same subnet as PS7012 . Type <https://192.168.0.1> into the web browser address window to access the PixelStor Storage Manager.



The workstation that will be used for system configuration should be on the same subnet as PS7012 . Set the Workstation IP address to 192.168.0.x and the net mask to 255.255.255.0 before trying to access the PixelStor Storage Manager.

3. You will see the PS7012 login. Enter the following login information when prompted.
Password: root
User Name:root

4. This will take you to Systems Properties page. Use the Setup Wizard on the far right side of the browser window (Steps 1-6) to guide you through the setup process.

Option 2 – Configure Private IP Address to Connect to System

Before connecting the PS7012 to your network, you will first need to access the built-in Command Line Interface (CLI) that is included with your system. The PS7012 is designed with one serial port on the back of each controller module. This can be used to connect a serial console to the PS7012 to manage the system. The CLI can be accessed via a terminal emulator running on a PC or laptop connected directly to the mini Serial port on the back of the PS7012 . On PC or laptop end the interface is USB.



See Appendix B – Using the Command Line Interface, for more information on the CLI and a full list of commands. The serial port is configured for 115200 Baud, N81.

1. To minimize any IP address conflicts, make sure the PS7012 is not connected to your network.
2. Connect the supplied serial cable to the COM port on either controller.
3. Connect the opposite end of the cable to a USB port on workstation or laptop with a terminal emulation program (e.g. TeraTerm or Putty).
4. You must use the following terminal emulation settings on the workstation/laptop to access the CLI:
 - Baud Rate – 115,200
 - Data Bits – 8
 - Parity - None
 - Stop Bits – 1
 - Flow Control Off
5. Power-on the PS7012, by plugging in AC power cords to both power supplies. The green DC power LED on the power supplies should turn on and stay on. The green LED on each installed disk drive should be on. The PS7012 comes with its own pre-installed CLI and system software. When the system is powered on, the software boots automatically.



When the PS7012 is started for the first time, the system may take 3-4 minutes. If the system is restarted during the initialization sequence, the process starts over from the beginning. Please wait for the system to finish its initialization sequence to avoid start-up delays. The mini-serial connector can be connected to monitor information sent to the console during the boot process.

6. Once connected and running the PS7012 will prompt you for the following login information on the workstation/laptop terminal screen:
 - Login – root
 - Password – root



If you have not done so already, consult with your network administrator for IP addresses that you can use to set up PS7012 on your network

7. This will start the CLI setup wizard.

```

RAStor - HyperTerminal
File Edit View Call Transfer Help

Setting raid5_force_reconstruct to 0
itoring

Starting UPS monitoring:[ OK ]
After daemoinze
Starting snmpd
Starting snmpd: [ OK ]
Starting cifsys
da_host register to the MEM successfullylogger: dalun -R succeeded
<iSCSI>: Set portal group 0 ip 192.168.0.2

Portal 192.168.0.2 with tag 0 is set
Starting httpd: [ OK ]
Starting crond: [ OK ]

Battery present and fully charged
Battery charger is good
Main Chassis: Power Supply 0 present
Main Chassis: Power Supply 0 FAN is OK
Main Chassis: Power Supply 0 DC is OK
Main Chassis: Power Supply 1 is present
Main Chassis: Power Supply 1 FAN is OK
Main Chassis: Power Supply 1 DC is OK
Fan board is present.
expander 0, Power Supply 0 present
expander 0, Power Supply 0 DC is OK
expander 0, Power Supply 0 FAN is OK
expander 0, Power Supply 1 is present
expander 0, Power Supply 1 DC is OK
expander 0, Power Supply 1 FAN is OK

login: root
Password:

This wizard will walk you through the network configuration.
You can configure the rest of the system using the web GUI.
The default values are in the bracket, you can press "Enter" to
select the default value. If you make any mistake, you can
press "ctrl-c", then logout and login to run this wizard again.

Enter IP address for system [192.168.0.1]:192.168.123.60
Enter netmask [255.255.255.0]:
Enter mtu [1500 - Jumbo frame disabled]:
Enter NIC1 IP [1]:192.168.123.61
Enter NIC2 IP [1]:192.168.123.62
Enter NIC3 IP [1]:192.168.123.63
Enter NIC4 IP [1]:192.168.123.64
Enter NIC5 IP [1]:192.168.123.65
Enter NIC6 IP [1]:192.168.123.66
Enter default gateway for system [1]:255.255.255.0
Invalid IP address 255.255.255.0
Enter default gateway for system [1]:192.168.123.254
Enter DNS address for system [1]:192.168.123.4
Portal 192.168.123.60 with tag 0 is set
[root@(none) ~]#
[root@(none) ~]#
[root@(none) ~]#

```

Figure 25: CLI Setup Wizard



This wizard will only run automatically when you log in from the serial terminal as "root". It will NOT run when you login using SSH or log in as a non-root user. Once the network is configured, the wizard will not run automatically again. However, you can run the wizard manually by issuing the command 'dasetup' from the CLI (See Appendix B – Using the Command Line Interface for more information)

8. Once launched, the wizard will walk you through the network configuration. The default values are in the brackets. You can hit your enter key to accept the default value or type in the private network information you received from your network administrator to change the default value. If you make any mistake, you can press "ctrl-c", then logout and login to run this wizard again. Here is a sample of the what you will be asked to provide:

- dedicated management port: [yes,no]
 - mtu: [1500, 9000]
 - NIC IP: [a.b.c.d, dhcp, disable]
 - Configure NIC1 as dedicated management port [no]:
 - Enter system IP address for 1-Gb ports []:
 - Enter system IP address for 10-Gb ports []:
 - Enter netmask [255.255.255.0]:
 - iSCSI discovery response with discovery IP [no]:
 - Enter mtu [1500 - Jumbo frame disabled]:
 - Enter C0 NIC1 IP [dhcp]:
 - Enter C0 NIC2 IP [dhcp]:
 - Enter C0 NIC8 IP 10G [dhcp]:
 - Enter C0 NIC9 IP 10G [dhcp]:
 - Enter C1 NIC1 IP [dhcp]:
 - Enter C1 NIC2 IP [dhcp]:
 - Enter C1 NIC8 IP 10G [dhcp]:
 - Enter C1 NIC9 IP 10G [dhcp]:
 - Enter default gateway IP []:
 - Enter DNS address for system []:
9. The PS7012 comes standard with dual 10Gbps RJ45 ports. Optionally it can have a quad port 10Gbps adapter. In this case then you will be asked to enter IP addresses for these additional ports (NIC10, NIC11). NIC 1, 2 and NIC8 thru 11 either use DHCP or static addresses as required. Once wizard is complete, connect the PS7012 to your network (RJ45 network cables not provided). You are now ready to launch the PixelStor web based management utility to install and configure your PS7012.

Rules for Adding Storage

PS7012 and PixelStor 312e2/390e/392e expansion enclosures follow the concept of “disk groups”. A “disk group” (DG) is composed of an associated group of disks with similar RAID properties. All DGs must be assigned a RAID level. The system supports RAID 0, 1, 10, 5, 6, 50, 60

The following rules must be followed when creating DGs for a PixelStor system:



To improve the utilization of the available storage capacity and ensure reliability, it is recommended that all disk drive modules within a DG be the same size and from the same manufacturer. This is because for each RAID Group, the system scales down to the size of the smallest drive (except for RAID 0).

- A DG can span across interconnected PS7012 and PixelStor expansion systems.
- Drives within a DG can reside in any order across any of the enclosures in a system.
- Capacity of a DG can be expanded by adding disks from the available pool of global spare disks. (Adding disks to a disk group may require multiple disks.)
- DGs can be accessed by either controller, but must be assigned a primary controller.
- Drives not assigned to a disk group will be classified as “Global Spares”
- Global spare(s) can reside on any of the interconnected chassis to protect any DG. A Global Spare cannot be added to RAID 0 DG.
- Failed disks or disks disassociated from their disk group will be classified as “Orphans” by the system.

Disk Group Parameters

Table 9: Disk Group Parameters

Disk Group (DG) Parameters	RAID 0	RAID5	RAID50	RAID6	RAID60	RAID1	RAID10
Minimum # of Drives per DG	2	3	6	4	8	2	4
Maximum # of Drives	108	108	108	108	108	108	108
Maximum # of Drives per System	108	108	108	108	108	108	108
DGs Allowed per System	1-16	1-16	1-16	1-16	1-16	1-16	1-16
RAID Levels Allowed per DG	1	1	1	1	1	1	1
Drives Must be Added in Groups of	Singles	Three	Six	Four	Eight	Pairs	Pairs
Global Spare Capability	N/A	Yes	Yes	Yes	Yes	Yes	Yes
Access to DG by both Controllers	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Configure the System

Storage Manager Overview

PS7012 ships with a comprehensive web based management console, “PixelStor Storage Manager”. The web based management console enables users to install, configure and maintain the PS7012 from an intuitive easy to use web interface. To access the PixelStor Storage Manager:

1. Login to the PS7012 from any laptop or workstation on the same subnet as the one you assigned to your PS7012 in step 8 of Initialize System section. You do this by opening a browser window and typing in any of the IP addresses you setup on the PS7012 in the previous section.
2. The workstation that will be used for system configuration should be on the same subnet as PS7012 .
3. Be sure to add the IP address to your Local Intranet access list in Internet Explorer otherwise your browser may not generate the correct output..

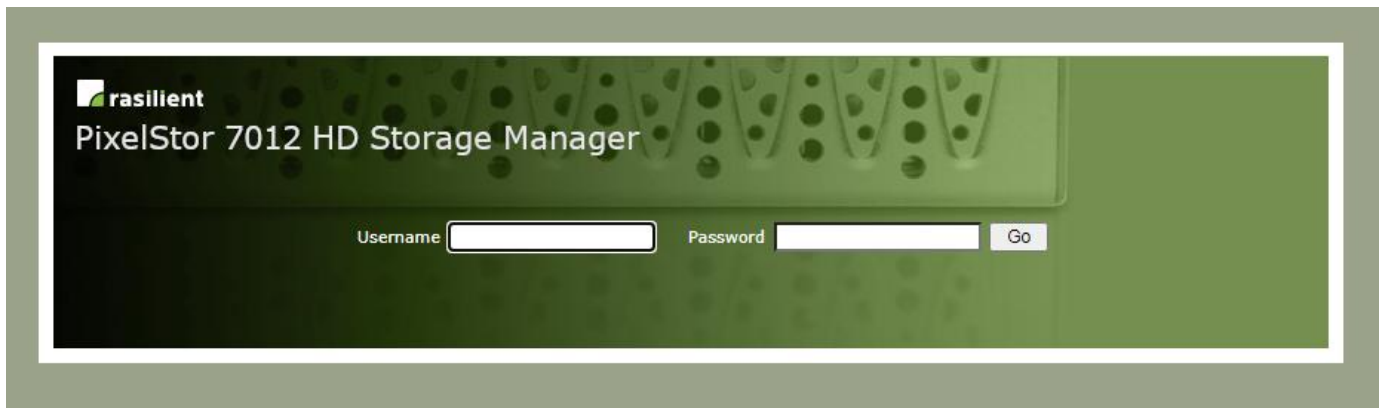


Figure 26: Storage Manager Login

4. Enter the following login information when prompted.
User Name:root
Password: root

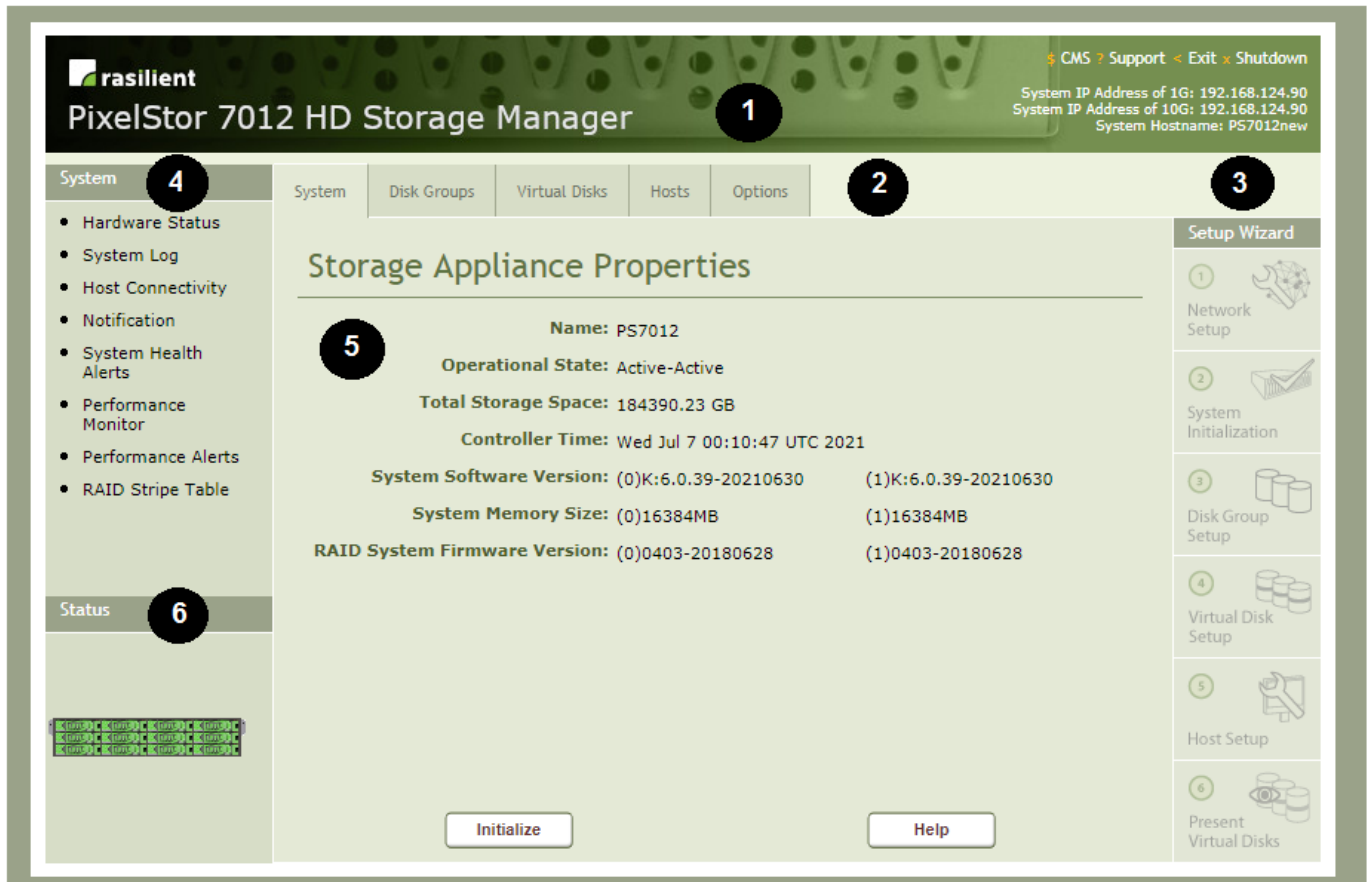


Figure 27:Management Console Layout

This will take you to Systems Properties page of the management console. There are 6 primary sections within the Storage Manager (see Figure 27:Management Console Layout)

- ❶ Console Banner
- ❷ Navigation Tabs
- ❸ Setup Wizard
- ❹ Navigation Tree
- ❺ Work Area
- ❻ System Status Icon

The Console Banner provides links to the support page, an exit link and a link to the Shutdown page. For easy reference, the System IP address and the Storage System name are displayed as shown in Figure 28:



Figure 28: Console Banner

The Navigation Tabs provide quick access to utilities that allow you to manage the PS7012 as shown in Figure 29. Each is discussed in greater detail in subsequent sections of this document.



Figure 29: Navigation Tabs

The Setup Wizard shown serves as a setup guide and progress bar that takes you through the six basic steps needed to configure and use your PS7012 .

These steps include:

1. Network Setup
2. System Initialization
3. Disk Group Setup
4. Virtual Disk Setup
5. Host Setup
6. Present Virtual Disks



Figure 30: Setup Wizard

The Navigation tree shown in Figure 31 provides a second level of navigation associated with the tab that is currently selected. For example, when System tab is selected, you can go to 5 additional sub categories. These include:

1. Hardware Status
2. System Log
3. iSCSI Statistics
4. Notification
5. System Health Alert
6. Performance Monitor
7. Performance Alerts
8. RAID Stripe Table

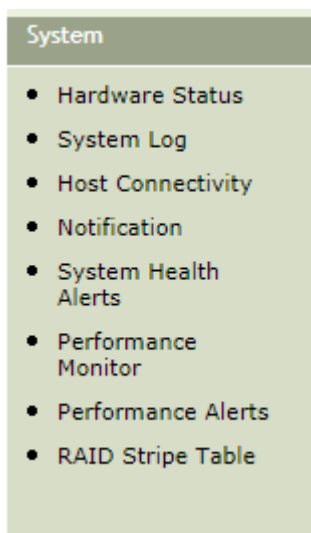
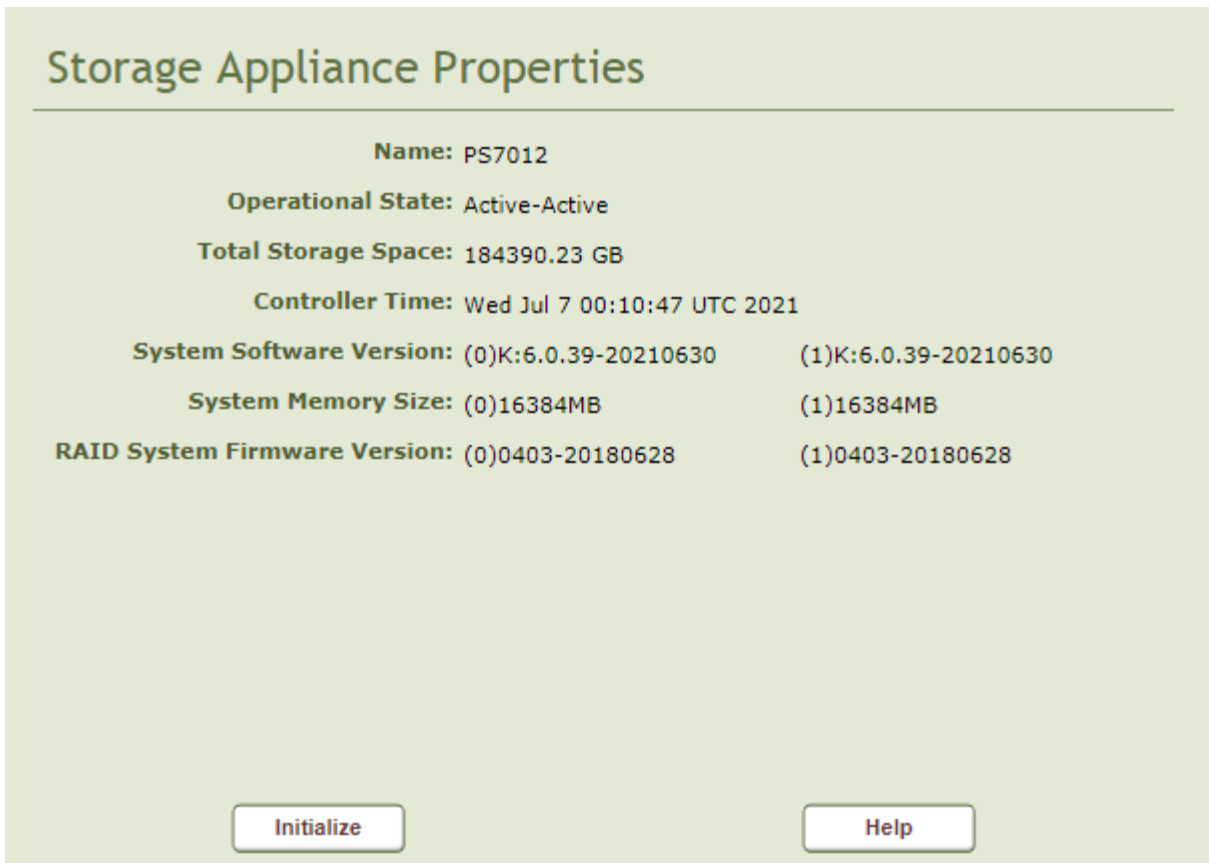


Figure 31: Navigation Tree

The Work Area, shown in Figure 32, changes depending on the Navigation tab you have selected. For each Navigation tab you will find information, input fields and action buttons associated with the specific Navigation tab/Navigation tree you have selected.



The screenshot displays the 'Storage Appliance Properties' window. It features a title bar at the top, followed by a horizontal line. Below the line, the following information is presented:

- Name:** PS7012
- Operational State:** Active-Active
- Total Storage Space:** 184390.23 GB
- Controller Time:** Wed Jul 7 00:10:47 UTC 2021
- System Software Version:** (0)K:6.0.39-20210630 (1)K:6.0.39-20210630
- System Memory Size:** (0)16384MB (1)16384MB
- RAID System Firmware Version:** (0)0403-20180628 (1)0403-20180628

At the bottom of the window, there are two buttons: 'Initialize' on the left and 'Help' on the right.

Figure 32: Work Area

The System Status Icons shown in Figure 33 provides a visual representation of the hardware status for the PS7012 .Healthy components are presented in green and faulty components are represented in Red. Figure 33 shows how a drive color changes when a drive has a fault condition

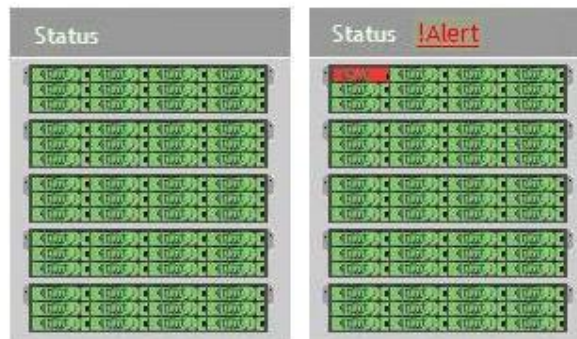
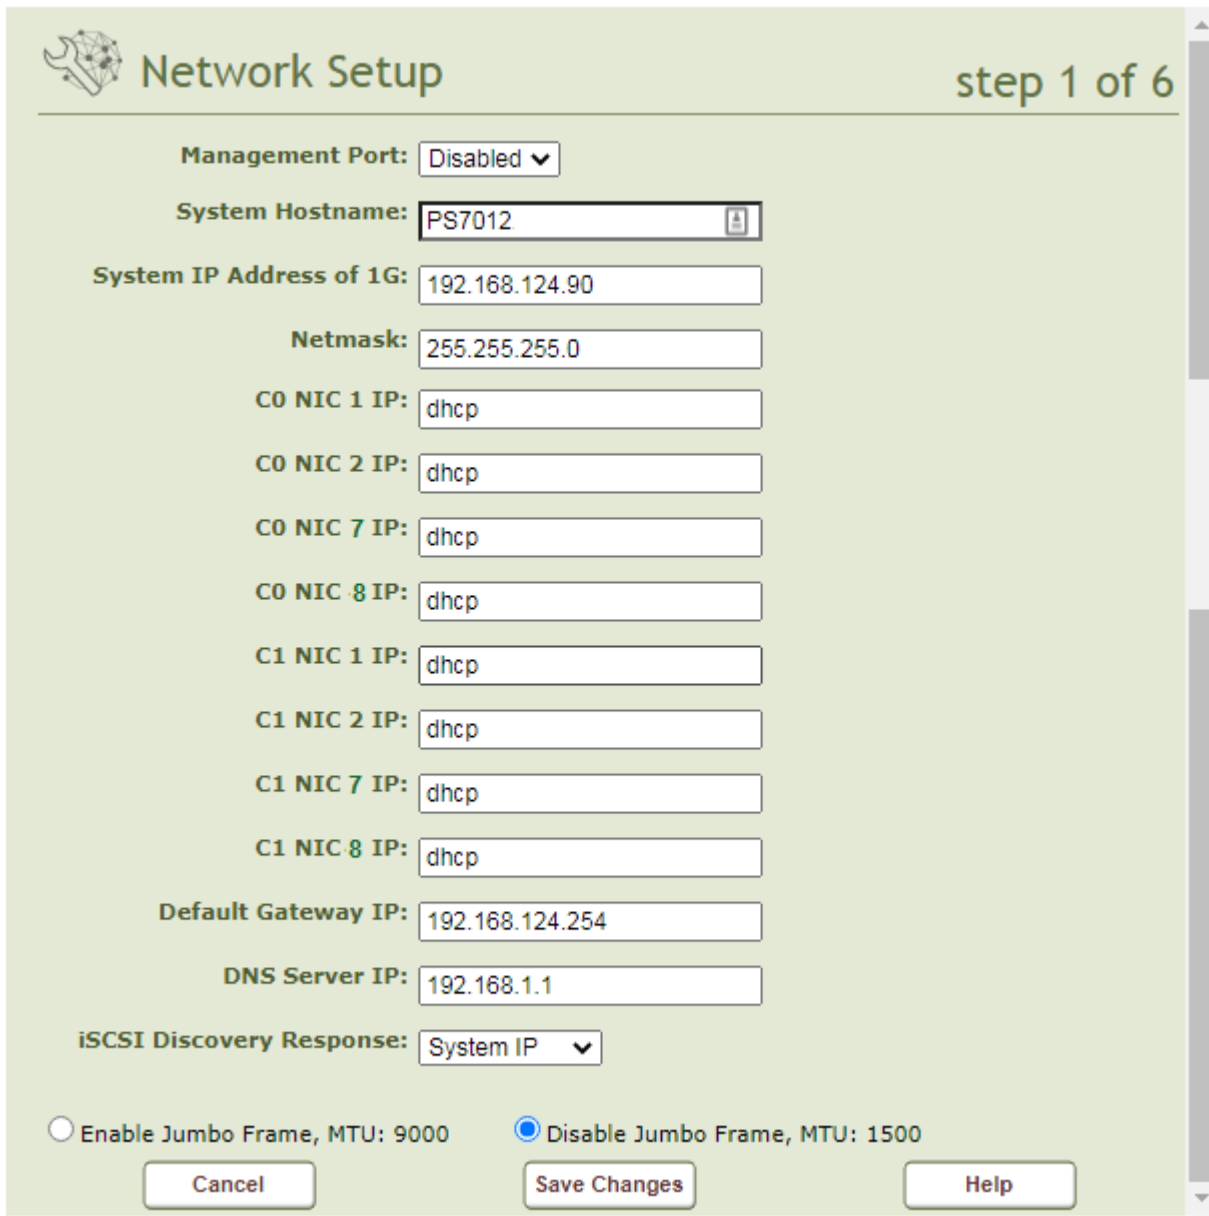


Figure 33: System Status Icon

Using the Setup Wizard

The Setup Wizard which is located on the right hand side of the Storage Manager serves as a setup guide and a progress bar for the six basic steps for setting-up a PS7012 . When you click on any of the steps in the Setup Wizard, that step is highlighted and the work area opens the associated pages for that step. Upon completion of a step, that step in the setup wizard is shaded to show that you completed that step in the setup procedure.

Once you have completed the initial setup, you can use the Setup Wizard to jump to any of the steps you care to repeat. Those familiar with the system may choose to use the Navigation tabs to get to the same pages.



The screenshot shows a web-based configuration interface for network settings. At the top left is a gear icon and the title "Network Setup". At the top right, it says "step 1 of 6". The form contains several fields: "Management Port" is a dropdown menu set to "Disabled"; "System Hostname" is a text box containing "PS7012"; "System IP Address of 1G" is a text box with "192.168.124.90"; "Netmask" is a text box with "255.255.255.0"; there are eight "NIC IP" fields (C0 NIC 1-8 and C1 NIC 1-8), all set to "dhcp"; "Default Gateway IP" is a text box with "192.168.124.254"; "DNS Server IP" is a text box with "192.168.1.1"; "iSCSI Discovery Response" is a dropdown menu set to "System IP". At the bottom, there are two radio buttons: "Enable Jumbo Frame, MTU: 9000" (unselected) and "Disable Jumbo Frame, MTU: 1500" (selected). At the very bottom are three buttons: "Cancel", "Save Changes", and "Help".

Figure 34: Network Setup

Step 1 – Network Setup

Use Network Setup (see Figure 34) to configure network connectivity to the system. If you logged-in using the default IP address or if you changed the default IP

address using the CLI wizard, you can make adjustments using this step. This step allows you to make the following changes:

1. Set System Name
2. Enter a system IP address. This bonds all the installed 1-GbE ports. Then, assign individual IP addresses to each GbE port.
3. For 10GbE ports enter a separate or same system IP address as 1-GbE ports



All IP addresses must be assigned in order for port bonding to operate properly.

4. Enter Net-mask value
5. Enter Default Gateway (GW) address
6. Enter Domain Name Server (DNS) IP address
7. Check the box to activate Jumbo Frames
8. Or Check the box to de-select Jumbo Frames
9. Click on Save Changes for your settings take effect.



If you change the IP address or netmask value you used to log into the controller, you will have to login again after clicking on Save Changes with your new IP Address.

Step 2 – System Initialization

Step 2 System Initialization, shown in Figure 35 and Figure 36, is used to define a system name and to set the date and time on the system. Once you have entered the system name, you can click on Finish to save these settings or click on Next to set the date and time.

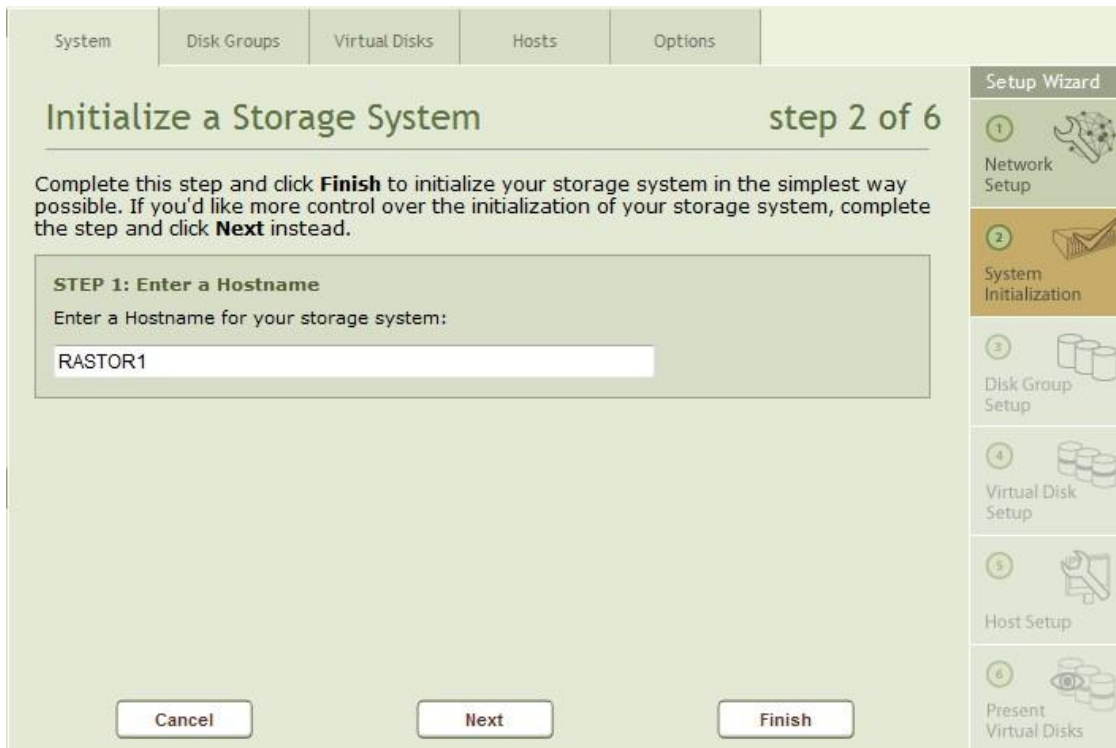


Figure 35: System Initialization--Step 1

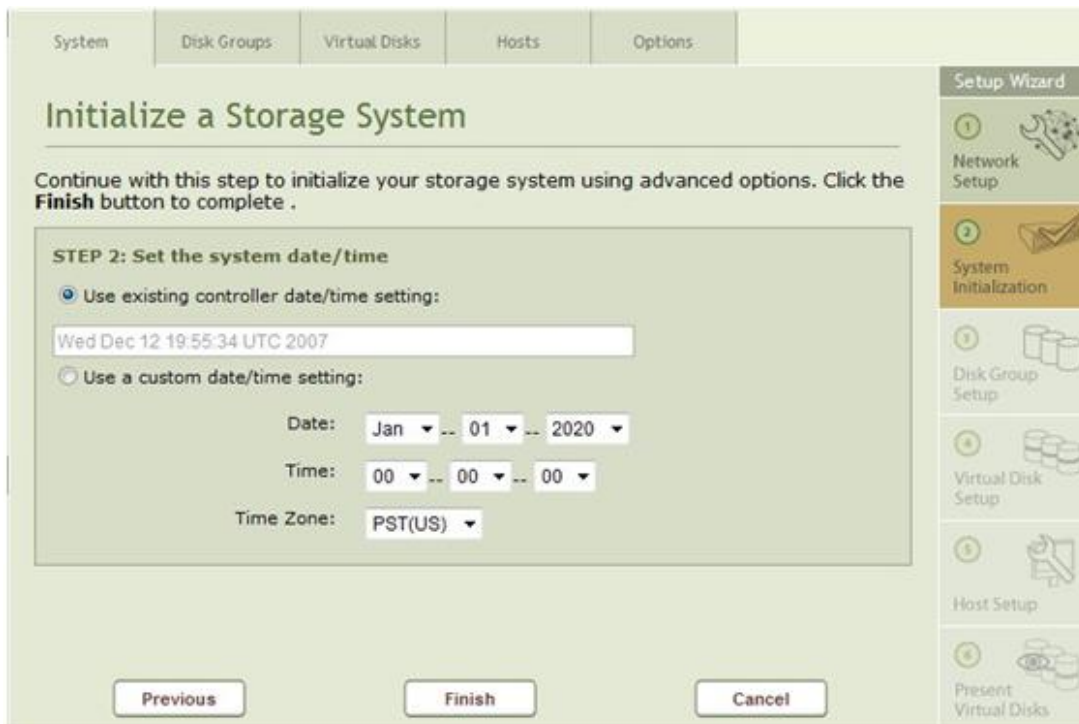


Figure 36: System Initialization -- Step 2

Step 3 – Disk Group Setup

A Disk Group is a collection of disk drives that is created by selecting disks to become members of the Disk Group and building a RAID device with the selected disks. Clicking on Step 3, Disk Group Setup, shown in Figure 37, opens the Disk Group Properties page where you see an overview of the Disks in the system.

The fields displayed on the Disk Group Properties page are:

- Operational State - Shows the status of the Disk Groups
- Total Disk Groups - Shows the number of Disk Groups created
- Total Grouped Disks - Shows the number of disk drives that are members of a Disk Group
- Total Spare Disks - Shows the number of spare disks that are eligible to become members of a Disk Group
- Total Orphan Disks - Shows the number of disks that have Disk Group information, but are not currently a member of a Disk Group
- Total Not-Owner Disks - Shows the number of disks that belong to a Disk Group owned by the remote controller. (This state is possible for dual configurations only). Disks are shown as Not-Owner disks when the remote controller is not available. This condition may occur when the remote controller is in the process of booting. You may wait for the remote controller to boot or use the Takeover DG to reassign these disks to the local controller.
- Total Faulty Disks - Shows the number of disks that have been declared faulty in the storage system.
- Total Unqualified Disks - Shows the number of disks that have been produced by Rasilient in the storage system and are not usable.
- Total Hyper Cache - Total number of SSD drives used as Hyper Cache in the system.
- Total Physical Disks - Shows the number of disks installed in the storage system.

Buttons

- Cancel - Return to the Storage Systems Properties page.
- Create DG - Create a new Disk Group.
- Discover DG - Manually discover Disk Groups.
- Orphan to Spare - Convert Orphan disks into Spare Disks. This removes all Disk Group information from the selected disks.
- Takeover DG - Reassign Not-Owner disks to the local controller
- Help - Display the help page.



Figure 37: Disk Group Setup

To create a Disk Group, click on the Create DG button of the Disk Group Properties page. This will open the Create a Disk Group page shown in Figure 38.

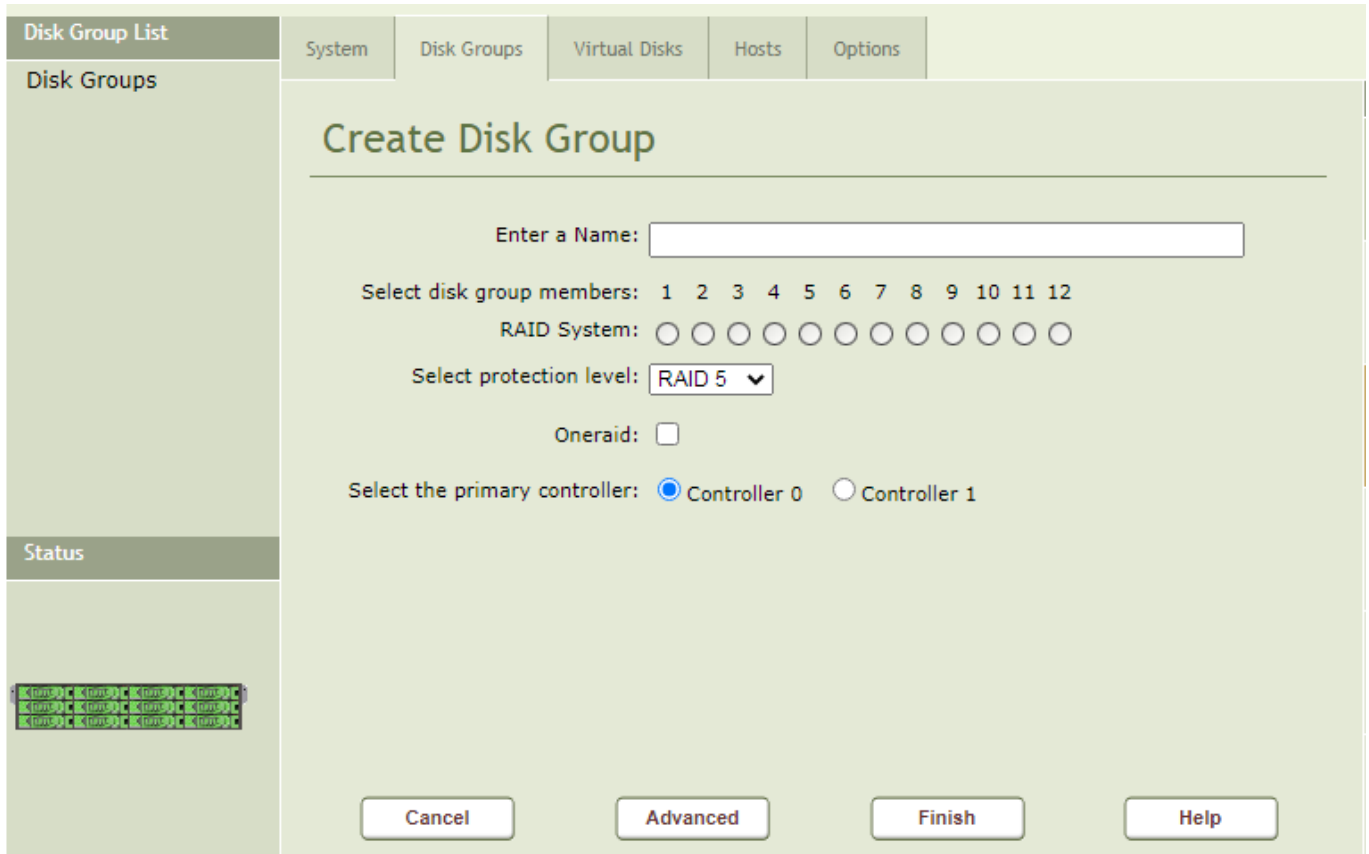


Figure 38: Create Disk Group

Once you open the Create a Disk Group page, you need to enter a name for the Disk Group (see Figure 39). The Disk Group name should not contain spaces or special characters. You will also need to select the drives to be used as members of the Disk Group. Then select a RAID level from the pull down menu by clicking on the down arrow of "Select protection level:". The RAID level selected for the Disk Group further imposes requirements associated with that RAID level.

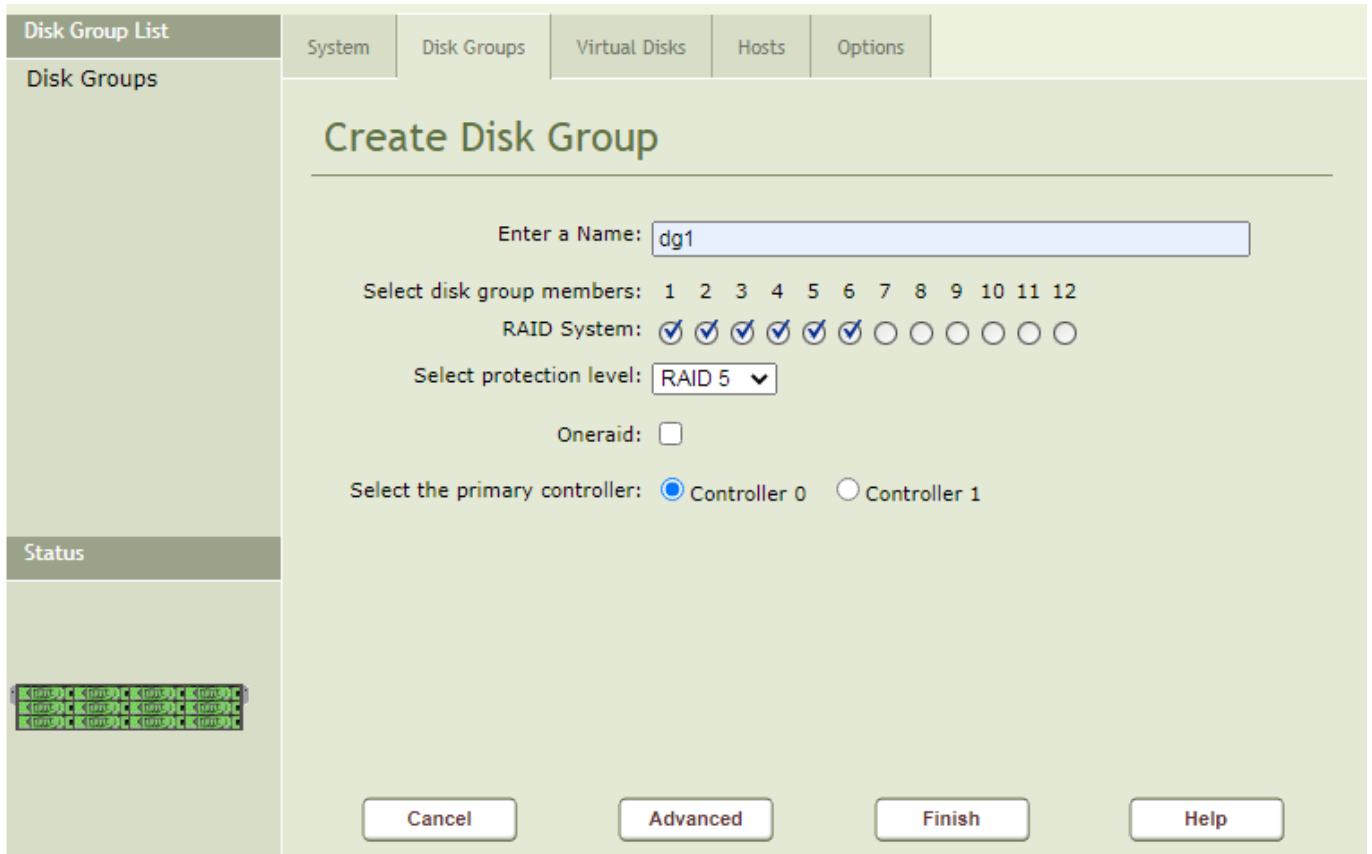


Figure 39: DG Name, Disks, Controller, Protection level

The following table describes the requirements for each RAID level:

Table 10: Disk Group Parameters

Disk Group (DG) Parameters	RAID 0	RAID5	RAID50	RAID6	RAID60	RAID1	RAID10
Minimum # of Drives per DG	2	3	6	4	8	2	4
Maximum # of Drives	108	108	108	108	108	108	108
Maximum # of Drives per System	108	108	108	108	108	108	108
DGs Allowed per System	1-16	1-16	1-16	1-16	1-16	1-16	1-16
RAID Levels Allowed per DG	1	1	1	1	1	1	1
Drives Must be Added in Groups of	Singles	Three	Six	Four	Eight	Pairs	Pairs
Global Spare Capability	N/A	Yes	Yes	Yes	Yes	Yes	Yes
Access to DG by both Controllers	Yes	Yes	Yes	Yes	Yes	Yes	Yes

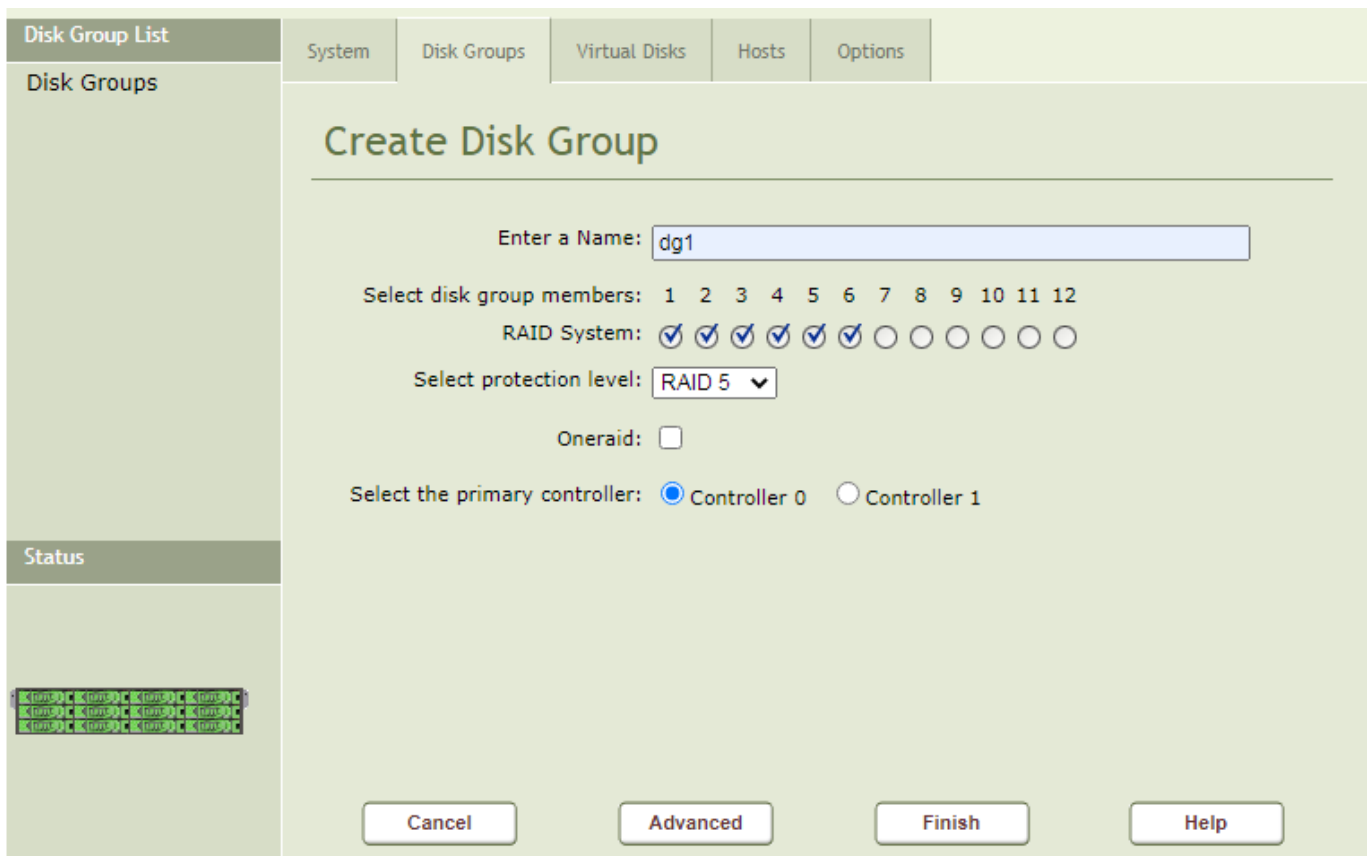
Drives that are not assigned to a Disk Group are available as Global Spares unless they are in an Orphan state or Failed Disk state. A Disk Group built as RAID 0 cannot use a global spare.

It is recommended that for a RAID 1, RAID 5, RAID 6, RAID 50, RAID 60 or RAID 10 Disk Groups, the drives the same size. The usable space on any given drive in the

Disk Group will be equal to the size of the smallest drive in the Disk Group (except for RAID 0).

Next, you select a primary controller for the new Disk Group on the Create a Disk Group page. When using an iSCSI initiator, this setting will associate the Disk Group with the IP address of the controller you select as defined in the network setup. Simply click on the button for the controller you wish to associate with the Disk Group.

Figure 39 shows the selection of Controller 0 for the new Disk Group. Clicking on the Finish button will create the Disk Group. Clicking on the advanced button will open the Advanced Features page shown in Figure 40, where you can set an Occupancy Warning level and add any comments to the Disk Group attributes.



Advanced Features in Create Disk Group

Continue with these steps and click on Finish to create your Disk Group.

Enter the occupancy alarm level

%

Enter your Comments

Figure 40: Create DG -- Advanced Features

Once you click on Finish, a progress page will open, shown in Figure 41.



Figure 41: Creating DG -- in progress



You can continue to create additional Disk Groups or start using the Disk Group while it is being created.

When the Disk Group has been created, a properties page for the newly created Disk Group is opened. You can get to any specific Disk Group Properties page by selecting the Disk Group you are interested in from the Disk Group List on the left hand side of the Management Console. The currently selected Disk Group is highlighted in the Disk Group List as shown in Figure 42.

Disk Group Properties

Name:

Operational State: Active

Total Disks: 12

Total Capacity: 167627.50 GB

Occupancy: 0.00 GB

Select a Primary Controller: ▼

Protection Level: RAID6

Disk Group Members: 1 2 3 4 5 6 7 8 9 10 11 12

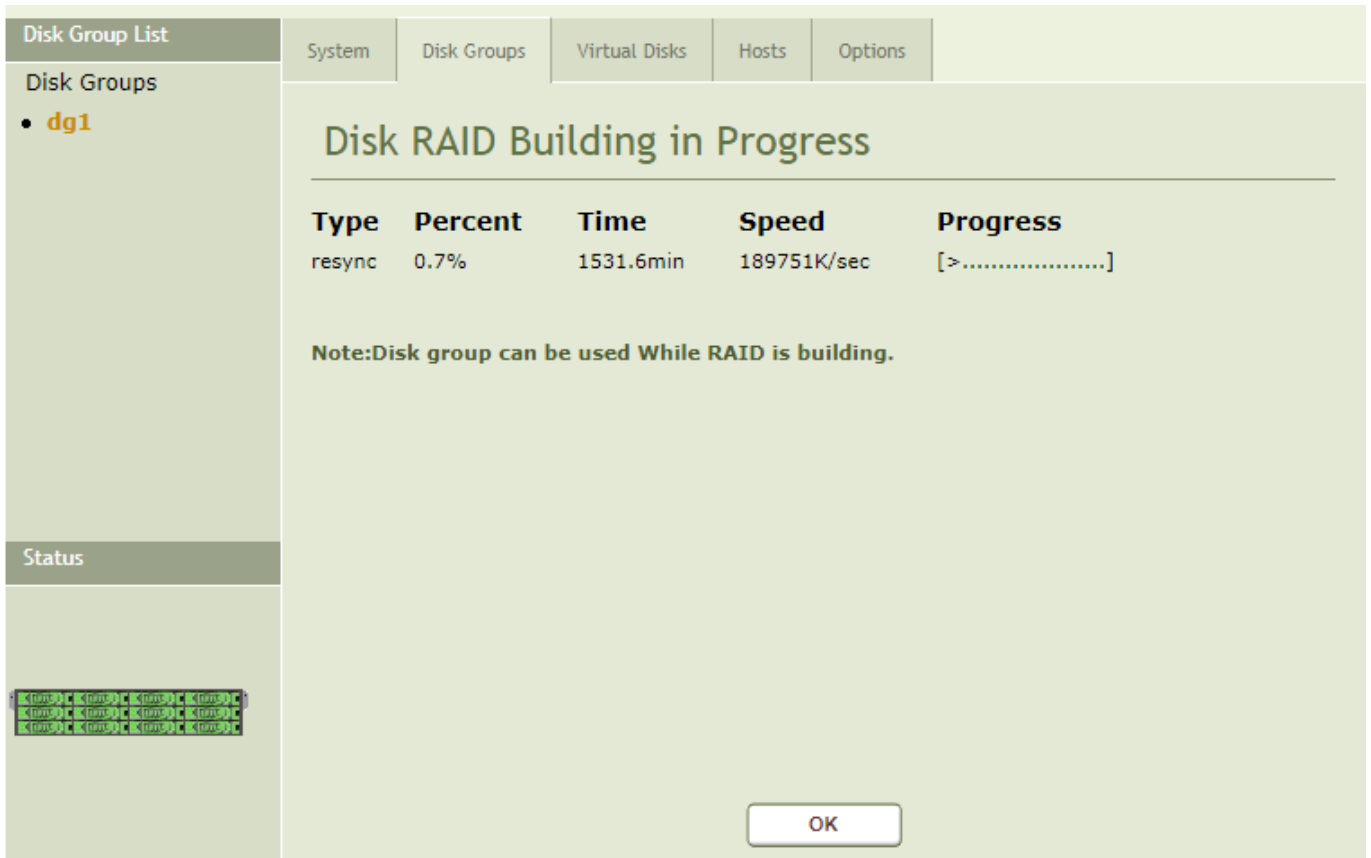
RAID System:

Comments:

Note:Disk group can be used While RAID is building.

Figure 42: Disk Group Properties

A RAID build is automatically started on the Disk Group when it is created according to the RAID level specified. The RAID is Building button is highlighted while the RAID is building. You can use the Disk Group while the RAID is building. You may also click on the "RAID Building" button to open the RAID Build Progress page shown in Figure 43.



The screenshot shows a software interface with a sidebar on the left and a main content area. The sidebar has a 'Disk Group List' section with 'dg1' selected, and a 'Status' section with a small RAID array icon. The main content area has tabs for 'System', 'Disk Groups', 'Virtual Disks', 'Hosts', and 'Options'. The 'Disk Groups' tab is active, displaying a 'Disk RAID Building in Progress' dialog box. The dialog box contains a table with the following data:

Type	Percent	Time	Speed	Progress
resync	0.7%	1531.6min	189751K/sec	[>.....]

Below the table is a note: "Note:Disk group can be used While RAID is building." At the bottom of the dialog box is an 'OK' button.

Figure 43: RAID Build Progress

The Disk Group Properties page for a specific Disk Group shows:

- Name - The name of the Disk Group
- Operational State - A Disk Group may be in one of the following states:
 - Active is where the Disk Group is up and running with all member disks with no faults on any of the disks in the Disk Group.
 - Inactive is where the Disk Group has been deactivated and is not running or unavailable to any host.
 - RAID Degraded is where the Disk Group is up and running but is either missing one or more Disk Group Members or where one or more of the Disk Group members has been declared faulty. In this state, the Disk Group is usable but without RAID protection.
 - RAID Unhealthy is where the Disk Group is missing two or more disk members. An Unhealthy Disk Group is not usable.
- Total Disks - Shows the total number of disks in the Disk Group
- Total Capacity - Shows the Disk Group capacity in Gigabytes
- Occupancy - Shows the percentage of the capacity of the Disk Group that has been allocated to Virtual Disks
- Select Primary Controller - Shows the current Primary controller selection or allows you to change the primary controller selection through the pull down menu. Selecting an alternative primary controller will not take effect until you click on the Save Changes Button.

- Protection Level - Shows the RAID protection level of the Disk Group
- Disk Group Members - Shows the disk drives that are members of the Disk Group
- Comments - Shows any comments that were entered from the advance features of the Create Disk Group page or allows you to enter comments for the specific Disk Group. Comments will not be associated with the Disk Group until you click on the Save Changes button.

Buttons

- RAID Building - The RAID Building button is shown while a RAID is building. This occurs during the initial creation of a Disk Group or when the RAID is being reconstructed. A RAID device is reconstructed when a member has been declared faulty and has been replaced by a global spare drive or by physically replacing the faulty drive.
- RAID Cloning - The "RAID is Cloning" button will only appear, when the disks are being cloned in this disk group. Clicking on this button will show the disk RAID cloning status and progress.
- Deactivate/Activate - Clicking on the Deactivate button will take a Disk Group off line, making all of the Virtual Disks unavailable to host on the network. For Disk Groups that have been deactivated, clicking on Active will make the Disk Group and all the Virtual Disks contained in the Disk Group available to hosts.
- Expand - Click on Expand to increase the capacity of a Disk Group. You add members to the Disk Group in sets that are consistent with the RAID type.

- Delete - Click on Delete to remove a Disk Group from the system. You must remove any Virtual Disk from the Disk Group before deleting the Disk Group.
- Save Changes - Click on save changes to apply any changes to the Disk Group.
- Help - Displays the Help page

When the RAID build is done the RAID is Building button is no longer visible.

Disk Group Discovery

Disk Groups may be moved between chassis. When disks belonging to a Disk Group are installed in a new chassis, they may appear as Orphaned Disks. The Discovery process, shown in Figure 44: Disk Group Discovery, will examine disks to determine if they belong to a Disk Group, if all the Disk Group members are present, the Disk Group will be reactivated.



You should de-activate a Disk Group before moving the disks to a new or different chassis. Removing multiple drives from an active Disk Group without deactivation can create a failure scenario that makes the Disk Group unhealthy and unusable

You may select to Force Disk Group(s) online to bring a Disk Group online even when it is in a degraded state. You cannot use Force Disk Group(s) online to bring an unhealthy Disk Group online.

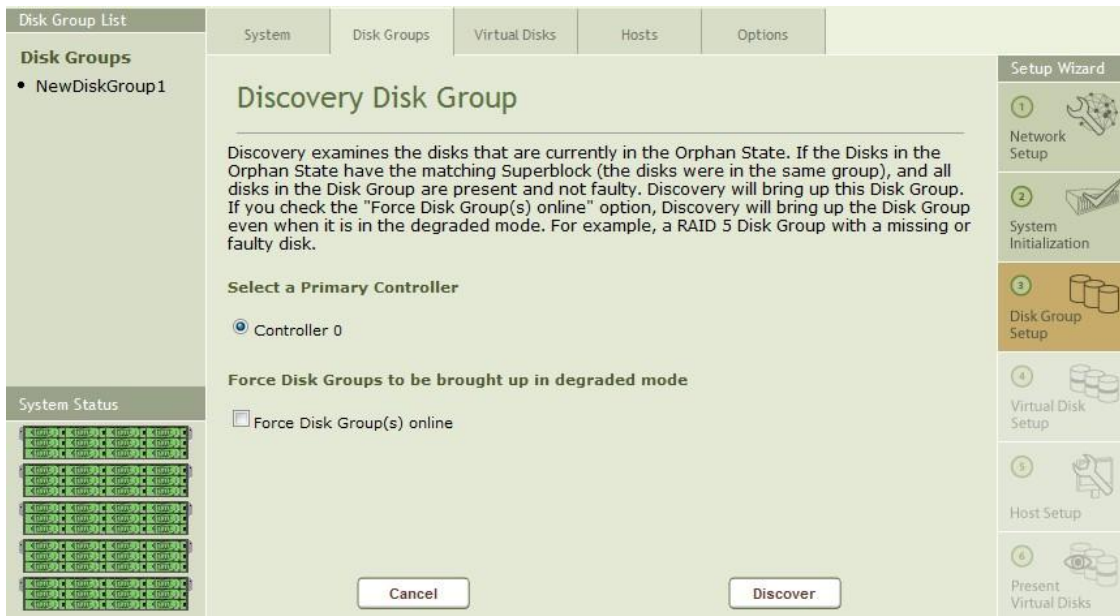


Figure 44: Disk Group Discovery

Orphan Disks

An Orphan Disk is a disk that contains information about membership in a Disk Group that has either been removed from the Disk Group or is designated as an Orphan because there are not enough of the Disk Group members available in the system to build the Disk Group.

This condition can occur when a disk has been improperly removed from the system and a hot spare disk has replaced the Orphan Disk in the Disk Group. Since the Disk Group is rebuilt using the Spare Disk, the Orphan cannot be used in the Disk Group.

This condition may also occur when some of the disks from a Disk Group are moved to a new chassis. Without enough Disk Group members present to reassemble the Disk Group, the system labels the disks as orphans. In this case, you can either locate and add the missing Disk Group members or convert the Orphan Disks to Spare Disks and rebuild the Disk Group from available Spare Disks. In this case all information on the disks will be lost.

Converting an Orphan Disk to a Spare will remove all data from the disks. Only those disks that are in an Orphan state may be selected to be converted. Click on

the Convert button in the Orphan Disk to Spare disk page, shown in Figure 45, to convert the disks.



Figure 45: Covert Orphan Disk(s) to Spare(s)

Expanding a Disk Group

The ability to expand a Disk Group depends upon the number of drives and the type of RAID used in the Disk Group and having enough spare drives to complete the expansion. Expanding Disk Groups follows the same rules for each RAID type as when you created the Disk Group (See Table 10: Disk Group Parameters .



When expanding a RAID 5 DG, the expansion set will be appended to the original set and the total size adjusted accordingly. The striping of the original RAID 5 set will not be re-arranged.

To expand a Disk Group, select the Disk Group you wish to expand from the Disk Group List on the left hand side of the Management Console and click on the Expand button at the bottom of the properties page for that Disk Group shown in Figure 46: Expand a Disk Group.

This will open the Expand a Disk Group window shown in Figure 46: Expand a Disk Group, where you select the additional drives for the Disk Group.

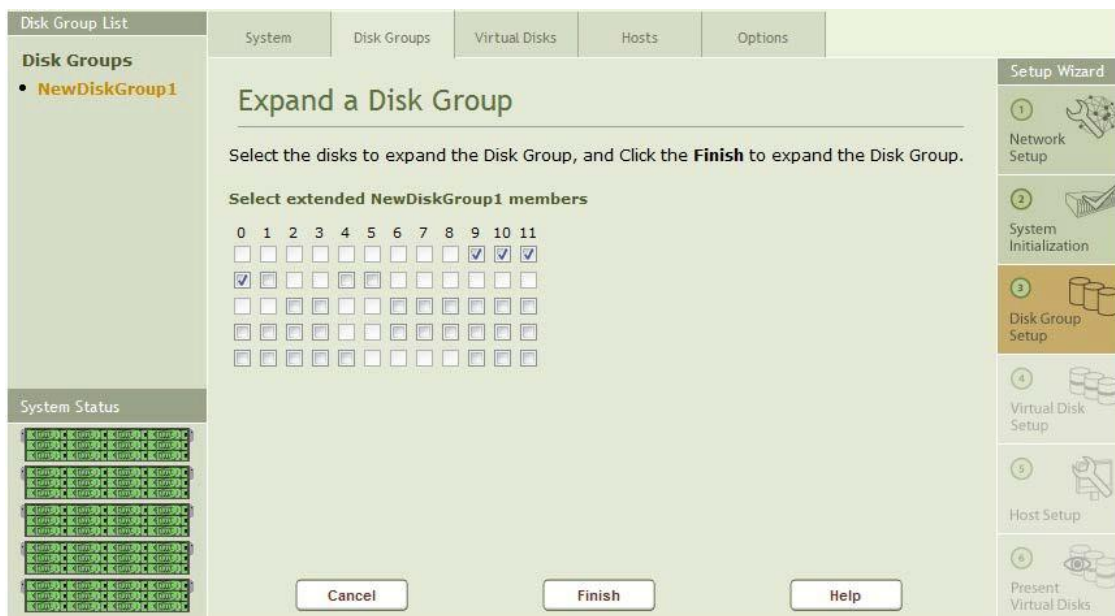


Figure 46: Expand a Disk Group

Once you have selected the drives to add to the Disk Group, click the Finish button to start the expansion.

Step 4 – Virtual Disk Setup

Once you have built at least one Disk Group, clicking on Step 4, Virtual Disk Setup will open the Virtual Disk Properties pages shown in Figure 47: Virtual Disk Setup.



The number of Virtual Disks (also known as logical volumes or LUNs) in all the Disk Groups in the system must not exceed 256.




Figure 47: Virtual Disk Setup

The Virtual Disk Properties page in, shows the Total number of Virtual Disk and Total active Disk Groups in the work area.

Buttons on the page include:

- Cancel - Returns the user to the System Properties Page
- Create VD - Opens the page to create a Virtual Disk

- [Help](#) - Displays help information.

Clicking on Create VD will open the page shown in Figure 48: Create Virtual Disk.

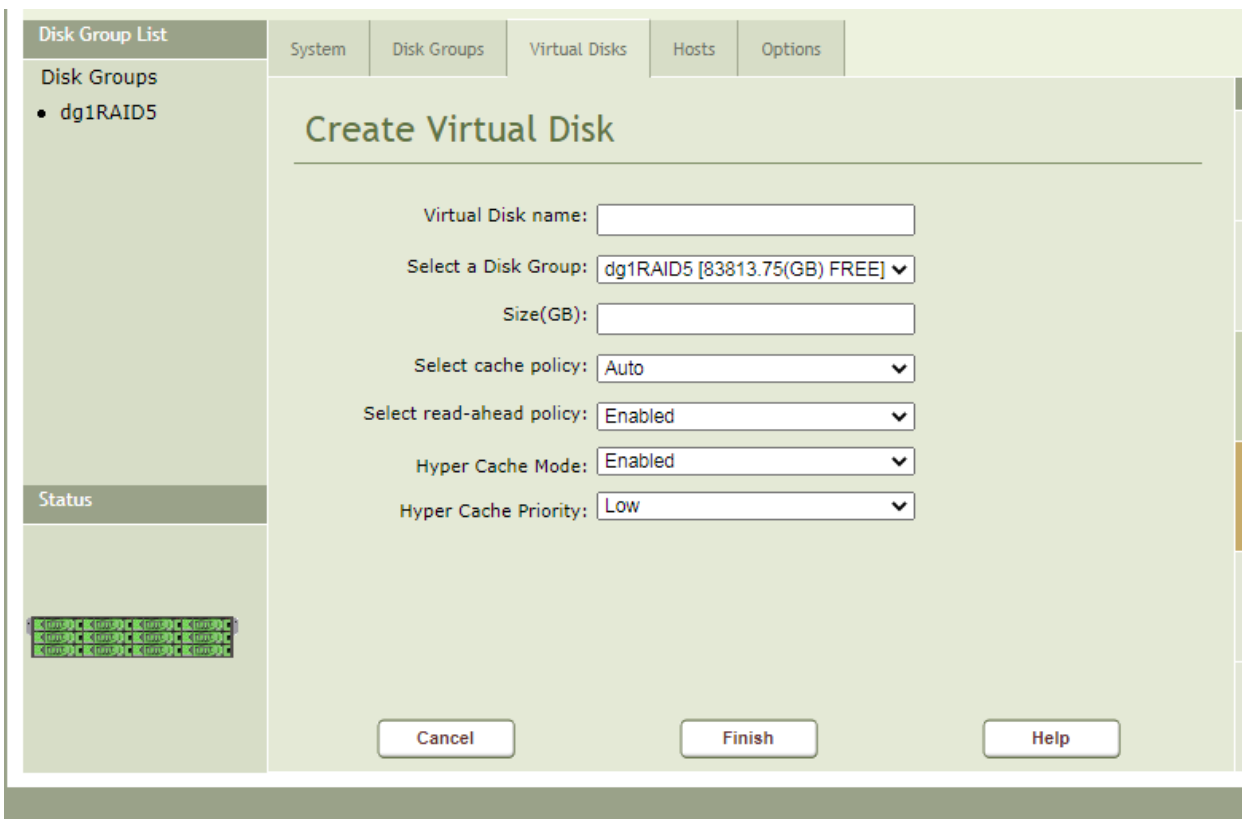


Figure 48: Create Virtual Disk

To create a Virtual Disk, do the following as shown in Figure 49: Virtual Disk Parameters:

1. Enter a name for the Virtual Disk.
2. Select the source Disk Group for the Virtual Disk as shown.

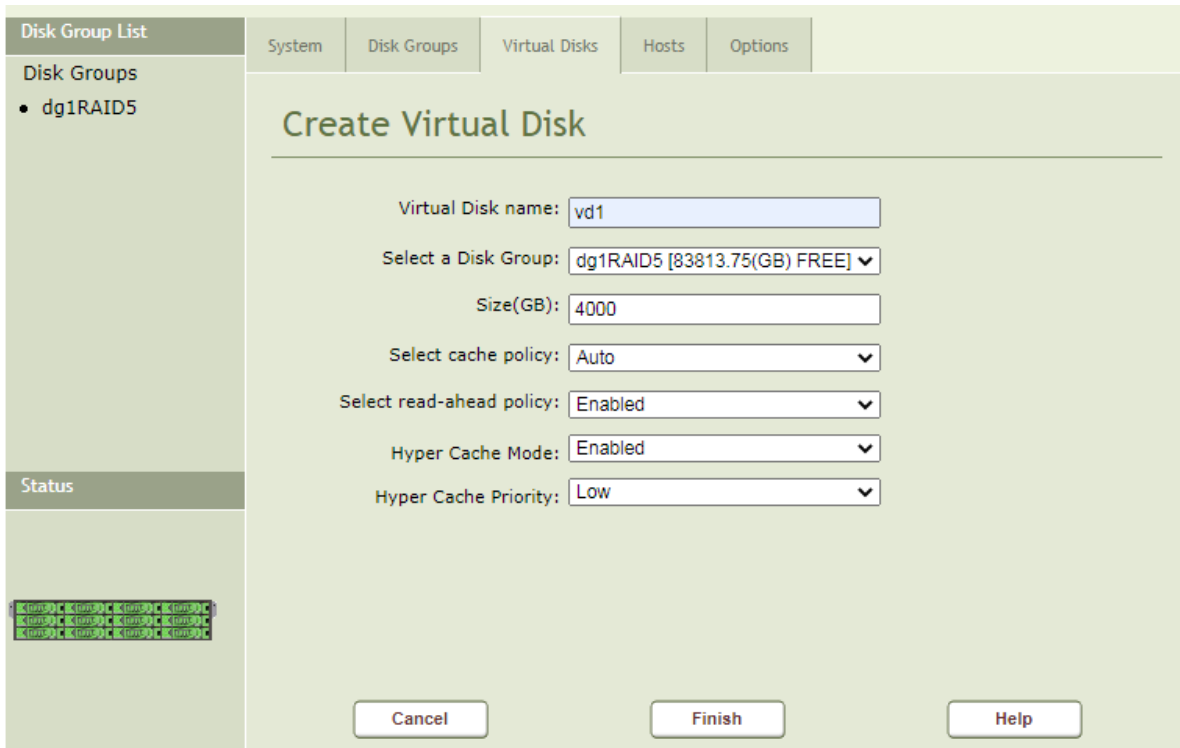


Figure 49: Virtual Disk Parameters

3. Enter the size in Gigabytes (GB) from the available free space in the Disk Group
4. Select the cache policy for the Virtual Disk
 - Auto – provides mirrored cache and write back policies while two controllers are present and operational. Should one controller fail, or be removed the cache policy is switched to write through. Default setting
 - Write Through – passes data to the disk and cache simultaneously.
 - Write Back – does not provide mirrored cache. Write Back stores data in battery protected cache and writes to disks as cycles are available.
5. Select the read-ahead policy

- Enable – allows read ahead operations. The read-ahead logic adjusts for traffic patterns to improve bandwidth. Default setting
 - Disable – turns off read ahead operations. It may be desirable to disable read-ahead in environments with highly random small block read and write operations.
6. Hyper Cache Mode - Enable or disable Hyper Cache
 7. Hyper Cache Priority - Selects priority of VD Hyper Cache



The screenshot displays the 'Virtual Disk Family Properties' page. On the left, a 'Disk Group List' sidebar shows 'dg1RAID5' expanded with 'vd1' selected. The main content area has tabs for 'System', 'Disk Groups', 'Virtual Disks', 'Hosts', and 'Options'. The 'Virtual Disks' tab is active, showing the following properties:

- Name: vd1
- Date/Time Created: Mon Aug 2 15:13:06 2021
- Operational State: Available
- Disk Groups: dg1RAID5
- Cache_Policy: Auto
- Capacity: 4000.00 GB
- Read Ahead Policy: Enabled
- Hyper Cache Mode: Enabled
- Hyper Cache Priority: Low
- Extend to new capacity: [input field] GB

At the bottom, there are five buttons: Cancel, Save Changes, Presentations, Delete, and Help.

Figure 50: Virtual Disk Properties

The Virtual Disk Properties page shown in Figure 50: Virtual Disk Properties displays:

- Name - The name of the VD
- Date/Time Created - Displays the date and time the Virtual Disk was created
- Operational State
 - Active when available

- Inactive when the Disk Group is deactivated or doesn't have an owner
- Disk Group - The name of the Disk Group in which the Virtual Disk resides
- Cache Policy - Auto, Write Through or Write Back.
- Capacity - The size of the Virtual Disk in GB
- Read Ahead Policy - Enabled or Disabled
- Hyper Cache Mode - Selects enable or disable Hyper Cache
- Hyper Cache Priority- priority of VD Hyper Cache
- Extend to new capacity - For expanding the capacity of the VD. You need to press the "Save Changes button" to take effect.

Buttons

- Cancel - Returns you to the Storage System Properties Page
- Save Changes - Works with Name, Cache Policy, Read Ahead policy and Extend Capacity field to make changes effective
- Presentation -Present or unpresent the Virtual Disk to Host(s).
- Delete - Remove the Virtual Disk from the system.
- Help - Displays the help page

Once you have created a Virtual Disk you can change several operational characteristics by entering a new value in a field on the Virtual Disk properties page and then clicking on the Save Changes Button. Figure 51: Extending VD Capacity shows how to expand the capacity for the Virtual disk. The additional space must be available in the Disk Group to allow the Virtual Disk to expand.



Figure 51: Extending VD Capacity

Step 5 – Host Setup

Clicking on Step 5 in the Setup Wizards opens the Host Properties page that is used to describe a Host to the Storage array as shown in Figure 52: Host Properties Page.



On PS7012 there is no need to create a Host definition because the VD is presented to all hosts by default. Only if presentation is intentionally limited to one host will a host need to be defined.

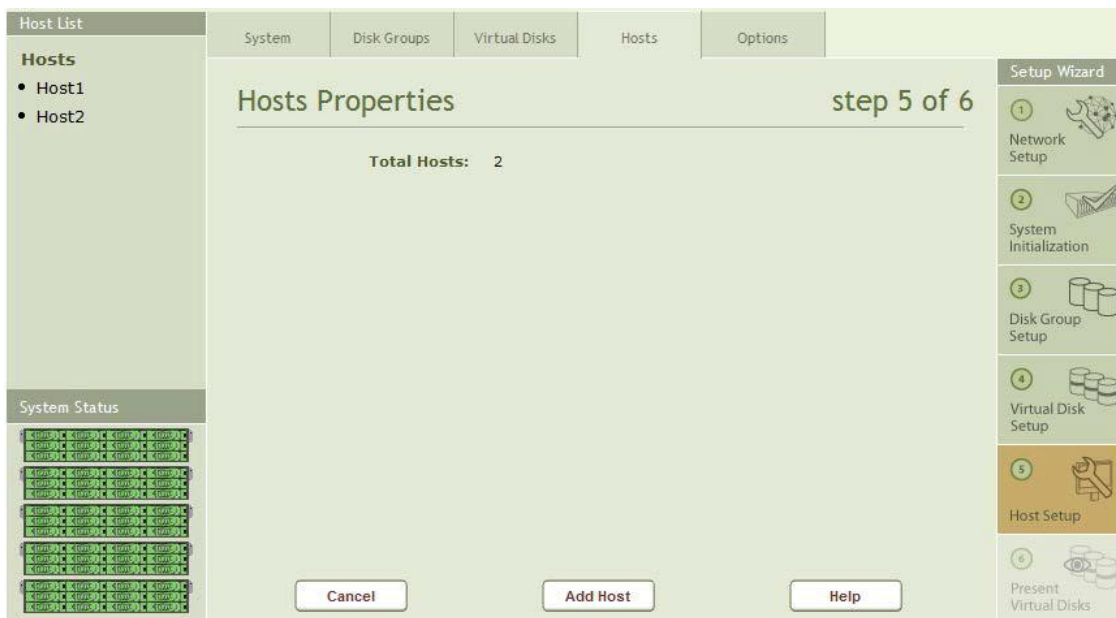


Figure 52: Host Properties Page

The left hand side of the screen shows any Hosts that have been defined in the system. Clicking on the Add Host button will launch the Add Host wizard to guide you through the 4 simple steps to describe a Host to the system. Figure 53: Adding a Host -- Step 1 & 2 shows the first 2 steps to add a host. Enter a Name for the new host in Step 1. Then in Step 2, select the interface to the host and enter the iqn name of the iSCSI initiator.

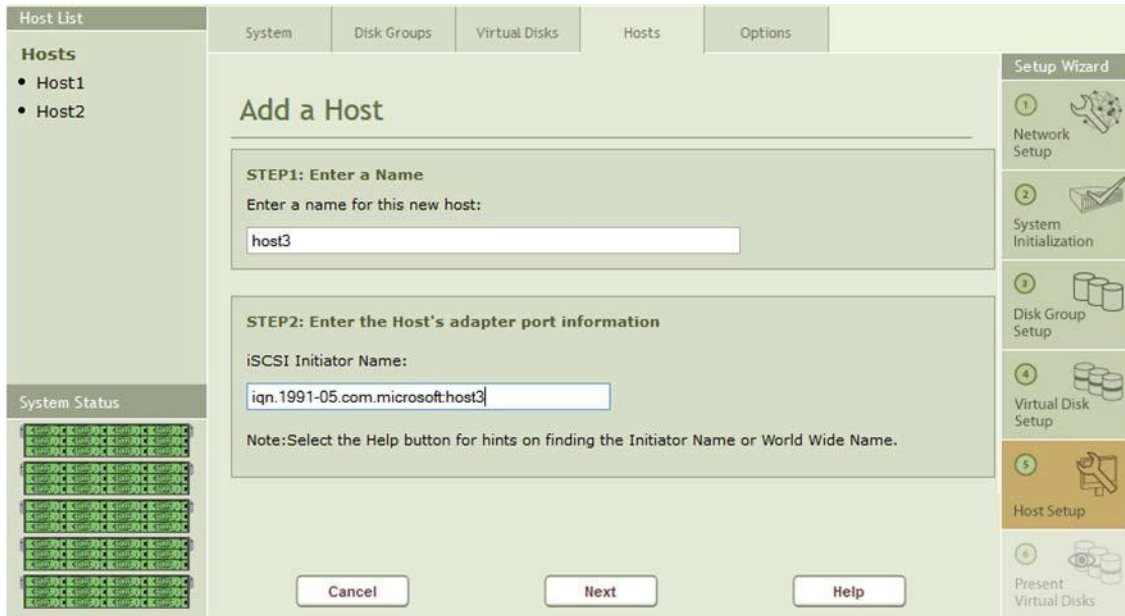


Figure 53: Adding a Host -- Step 1 & 2

Figure 54: Adding a Host - Steps 3 & 4 shows Steps 3 & 4 of adding a host. Use Step 3 to set up CHAP security. Enter the Target Secret and the Initiator Secret in the associated fields. Use Step 4 to add any comments about the Host that will be displayed when you select the host from the Host List.

Click on Finish to complete the process of describing a Host to the system.

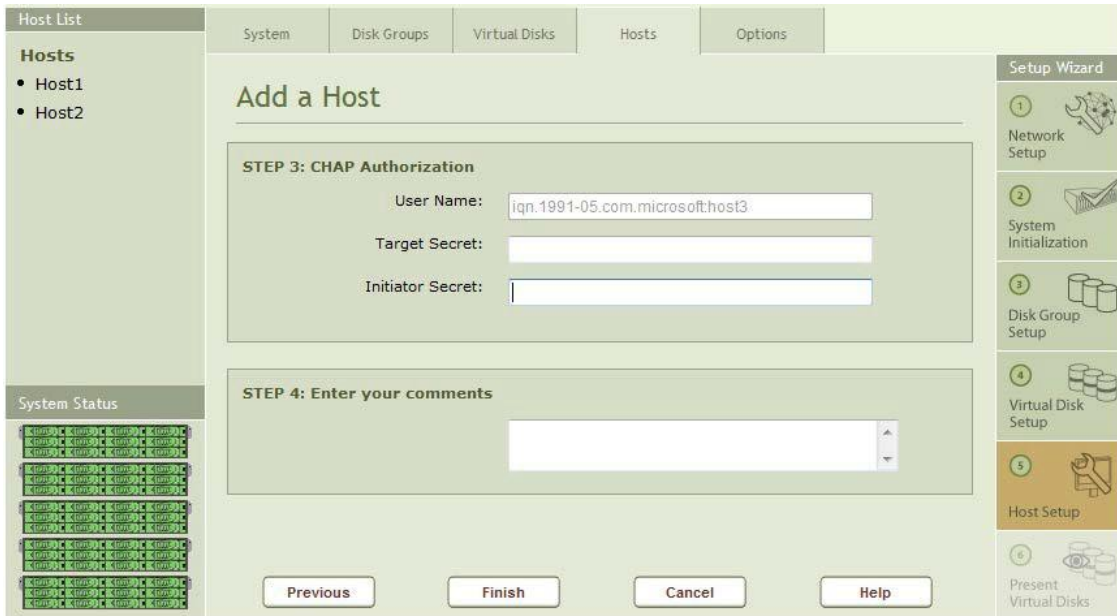


Figure 54: Adding a Host - Steps 3 & 4

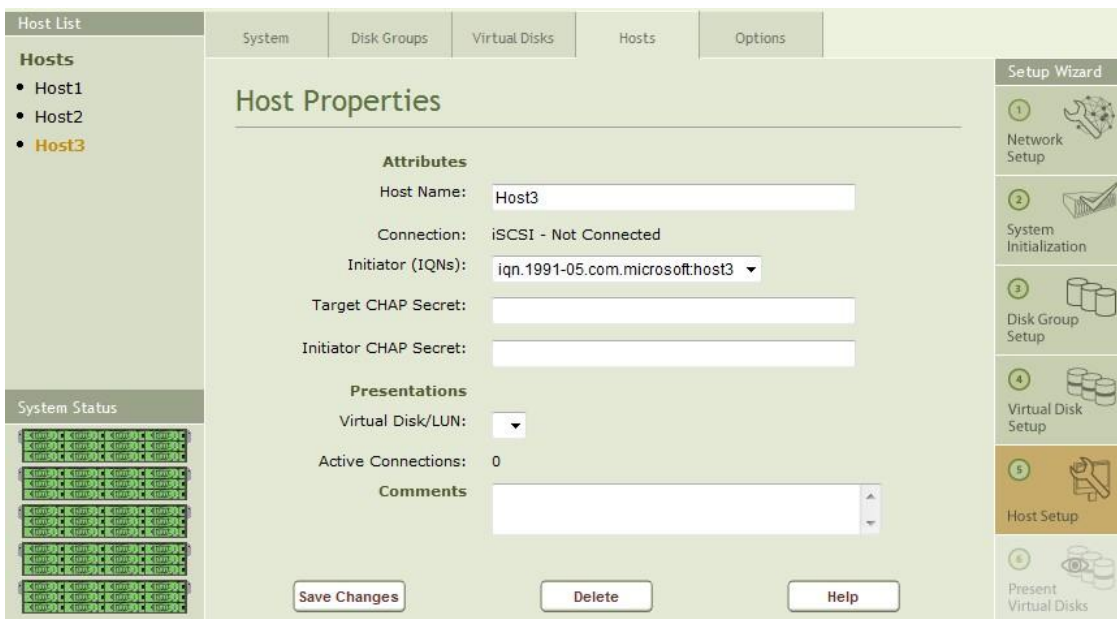


Figure 55: Host Properties

The Host Properties page for each host, shown in Figure 55: Host Properties, contains the following information about the host selected from the Host List:

- Name - Displays the current Host Name, or may be used with the Save Changes button to change the Host Name.
- Connection - Show the connection type and status
- Initiator (IQNs/WWNs) - Shows a list of iSCSI or Fibre channel ports that have been defined for the Host (if installed)
- Target CHAP Secret - Shows the target CHAP secret or works with Save Changes button to add a target CHAP secret.
- Initiator CHAP Secret - Shows the Initiator CHAP secret or works with Save Changes button to add an initiator CHAP secret.
- Presentations - Shows a list of LUNS that have been presented to the host
- Active Connections - Shows the number of active connections to the Host
- Comments - Displays any comments that were added when the Host was defined or works with the Save Changes button to add comments to a Host

Buttons

- Save Changes - Works with Name, Target CHAP Secret, Initiator Secret and comment fields to make changes take effect
- Delete - Host will be deleted.
- Help - Displays the Help page

Step 6 – Present Virtual Disks

The final step in setup is to present the Virtual Disk to the Host(s). Clicking on Step 6 Present Virtual Disks opens the page shown in Figure 56: Chose a VD to present. On this page, you select a VD to present from the pull down menu that shows available VDs and click on the Ok button.

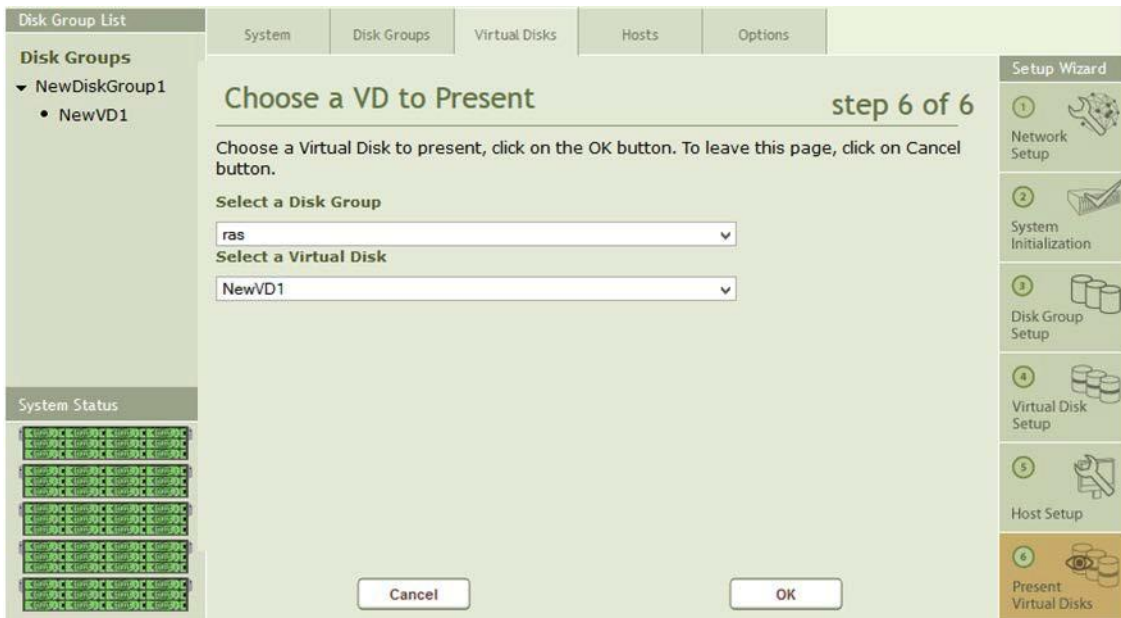


Figure 56: Chose a VD to present

The VD presentation page, shown in Figure 57: VD Presentation, provides details about the presentations of the VD selected in the Disk Group List. Any host that the VD is currently presented to will show in the list box. Since the current VD is new, the default setting is to present the VD to all hosts. If you wish to present to a specific host, select the “present VD to a specific host” option and hit save changes.



Figure 57: VD Presentation

To Present the Virtual Disk to a specific host, select the Host from the pull down list as shown in Figure 58: Present Virtual Disk Page and click on the “Finish” button.

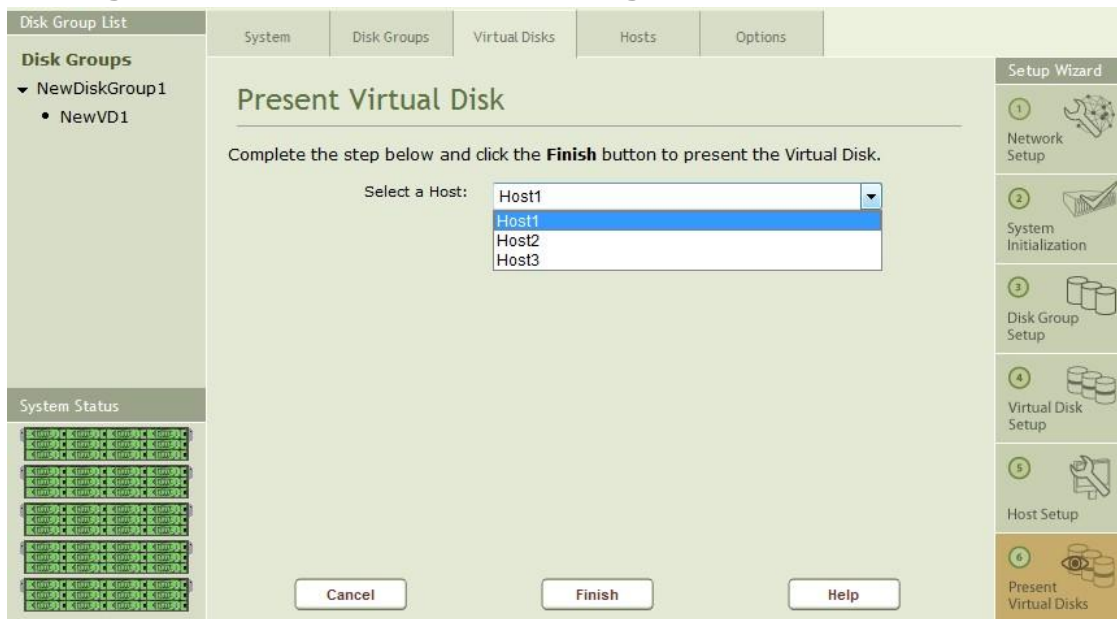


Figure 58: Present Virtual Disk Page

After you hit the finish button, the current presentations for the selected VD are shown (see Figure 59: Current Presentation).

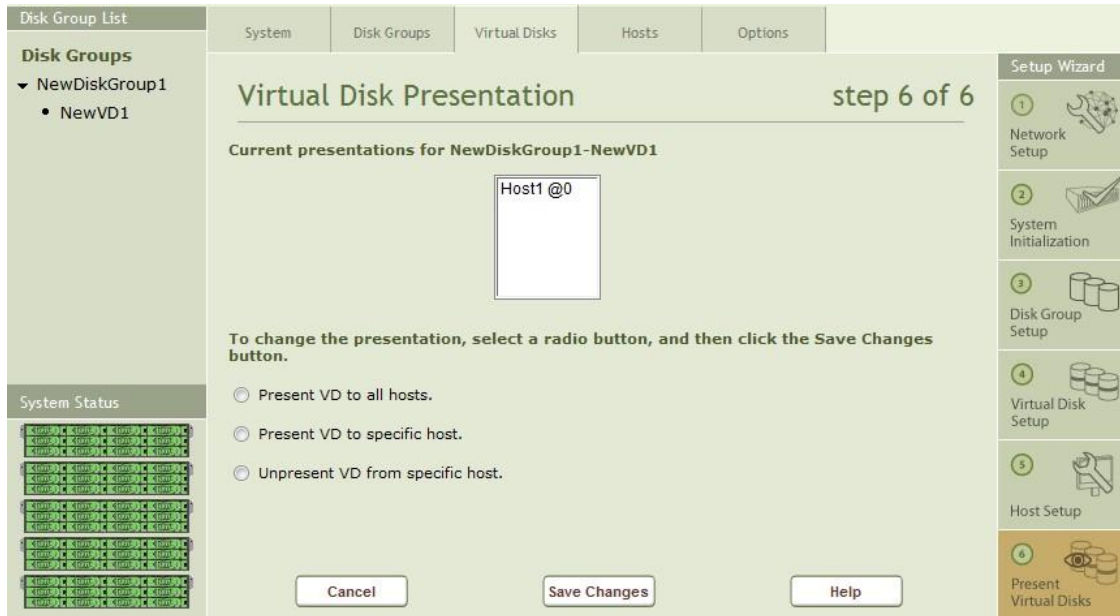


Figure 59: Current Presentation

Additional Configuration Options

The PS7012 allows you to perform additional tasks through the Options Tab. These include:

- Changing Password
- Language Option
- PixelStorOS Upgrade
- System Health Alerts Options
- Performance Monitor
- Display Refresh Options
- System Beacon/Identity LED
- Drive Write Cache Options
- SNMP Configurations
- Chassis Firmware Upgrade
- Hyper Cache Configuration
- Disk Power Management Status
- iSNS Options
- Advanced Logging Mode
- Inspection Read
- NTP Configuration
- UPS Configuration

You can access the options by clicking on the Options tab of the system toolbar in the Management Console. From the Options window you can select the area of interest by clicking on "Go Now" for the option you desire or by clicking on the item in the Option List on the left hand side of the Management Console.

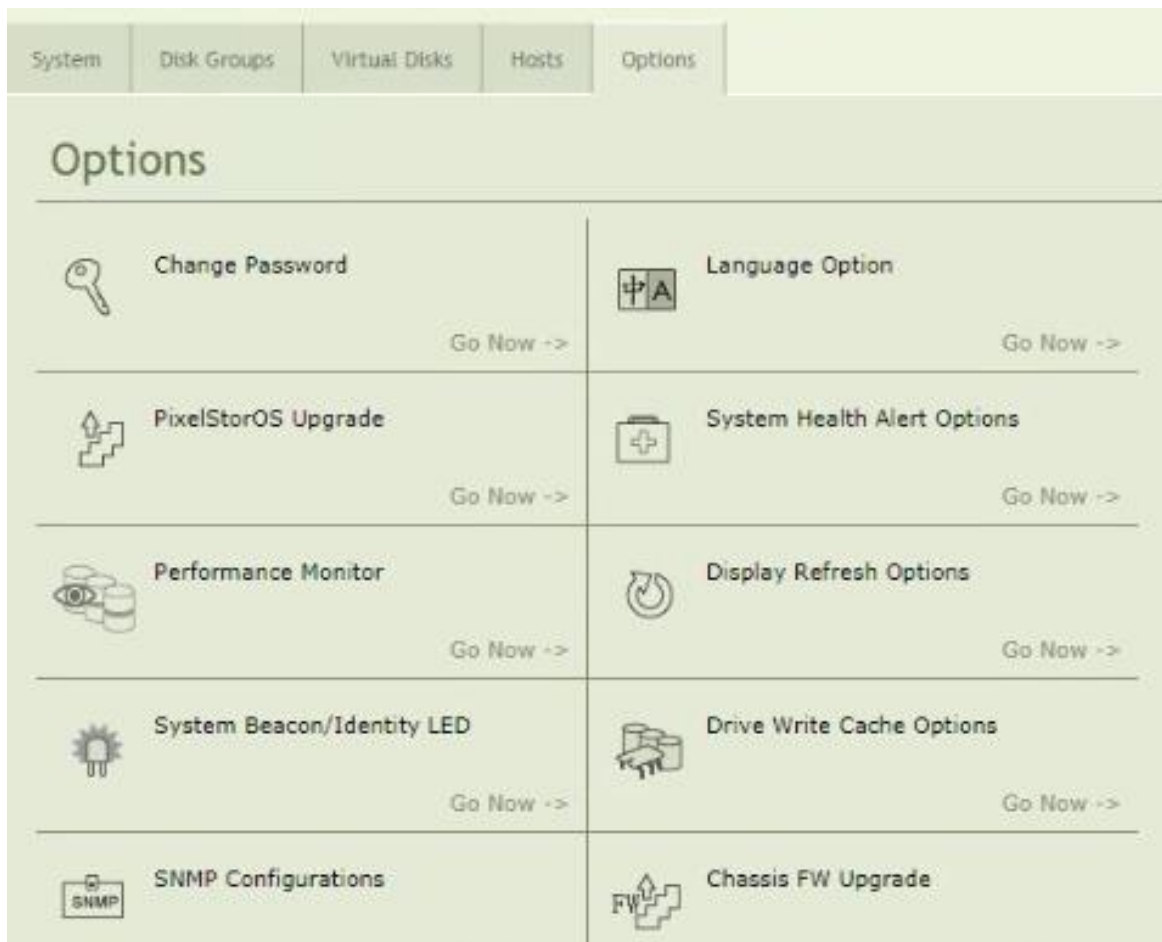


Figure 60: Options Top part of page

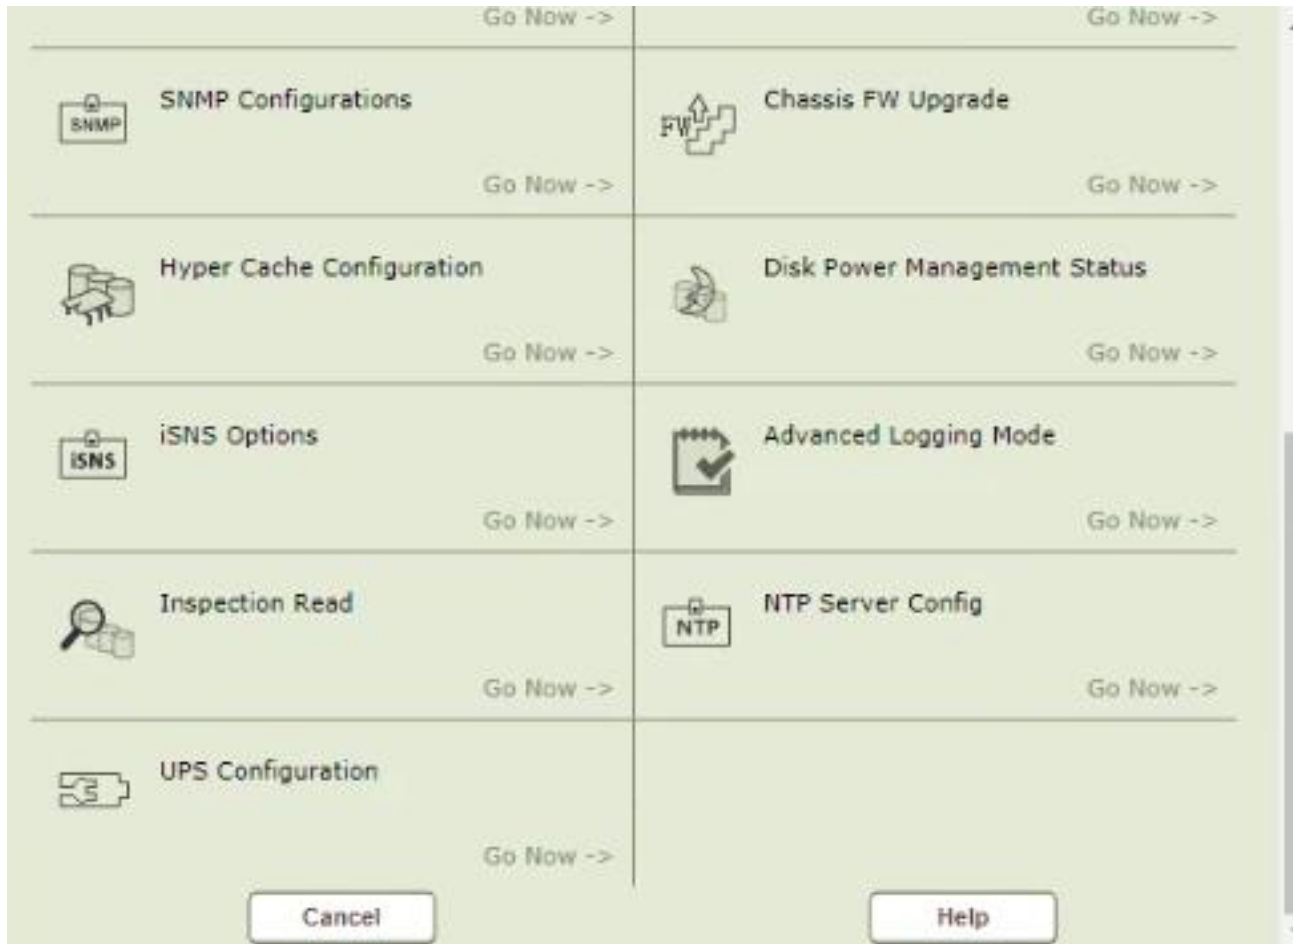


Figure 61: Options Bottom half of page

Change Storage System Password

To change passwords on the system, click on the Change storage system password icon in the options window to bring up the table shown in Figure 62: Change Storage System Password. You may change the password for root or guest through this menu.

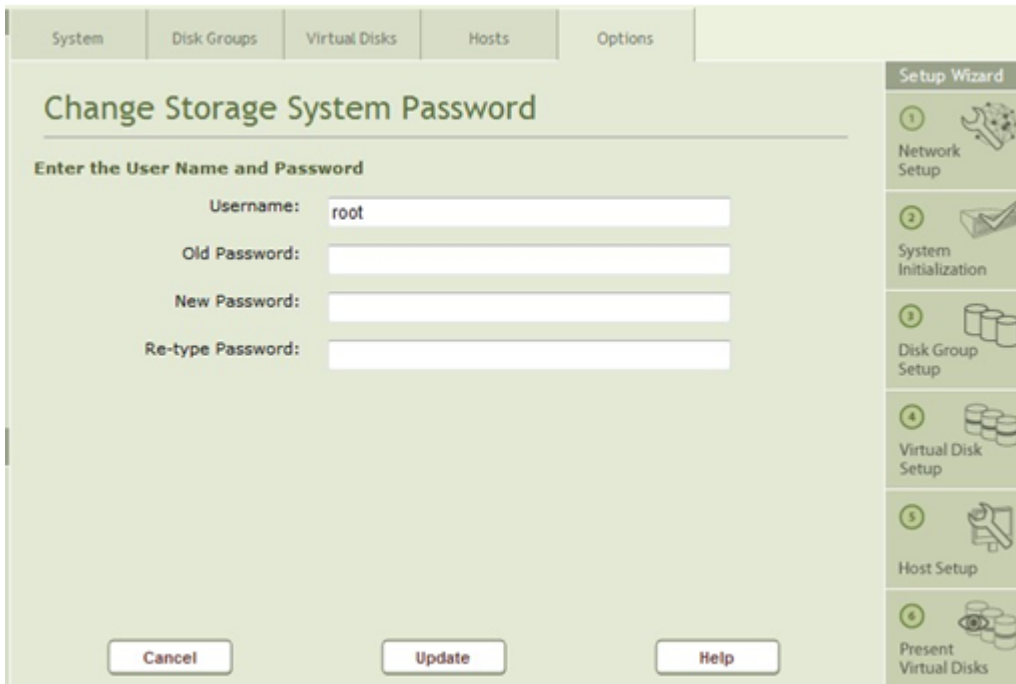


Figure 62: Change Storage System Password

To change the password:

1. Enter the user name
2. Enter the old password
3. Enter the new password
4. Re-enter the new password
5. Click on the update button

Change the Display Refresh Options

By default, the Management Console is refreshed every 5 seconds. You can change the refresh interval by clicking on the Display refresh options icon in the Options window. This will bring up the Display Refresh Options as shown in Figure 63: Display Refresh Options.

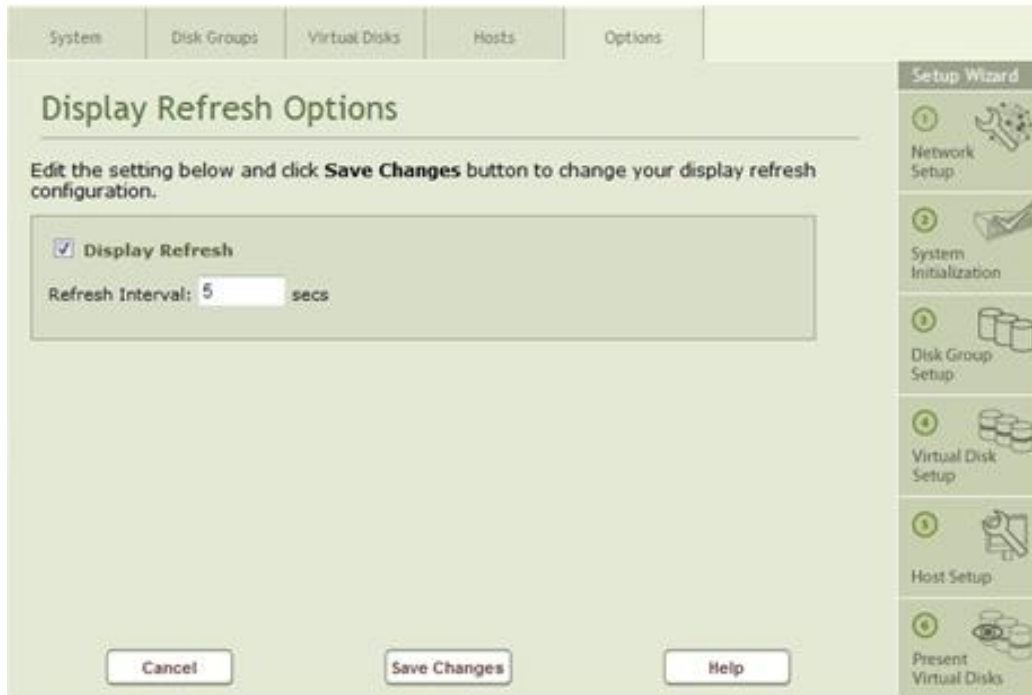


Figure 63: Display Refresh Options

To change the refresh interval:

1. Enter a new value in seconds
2. Click on the Save Changes button

Change System Health Alert Options

System Health Alert Options may be enabled or disabled through the System Health Alert Options page. You can reset the alerts status by temporarily disabling them and enabling the System Health Alerts. Click on the System Health Alert Option icon on the Options page to bring up the System Health Alert page shown in Figure 64: System Health Alert Options.

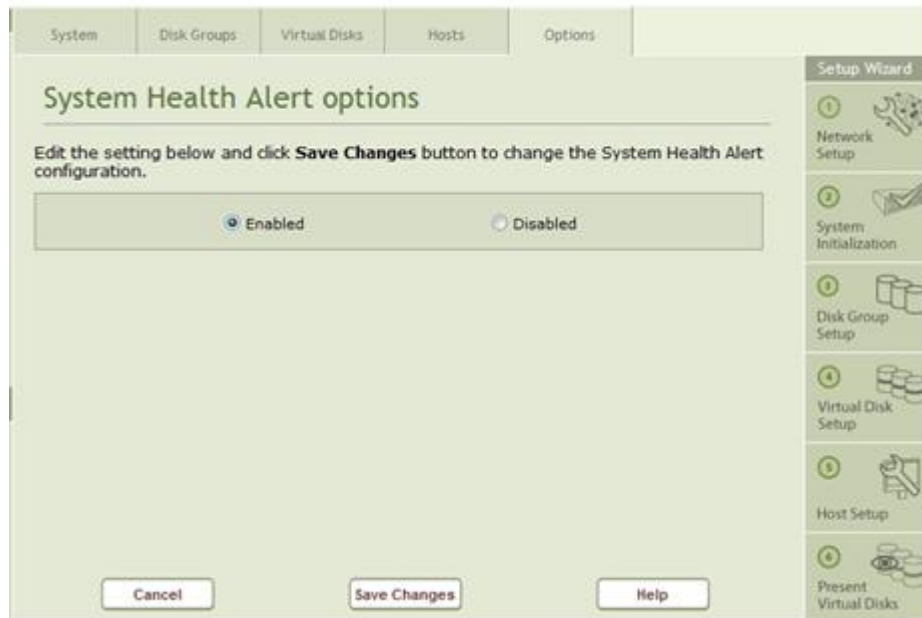


Figure 64: System Health Alert Options

To change health alert options

1. Select Enable or Disable on the System Health Alert Options page
2. Click on the Save Changes button

Configuring the iSNS Client

The PS7012 provides an iSNS client for registration with an iSNS server. To fully utilize the iSNS service you must configure an iSNS server and configure your iSCSI initiator to use the iSNS server.

Once you have defined the IP address and alias for the iSNS Server, you must start the iSNS client on the target. Starting the iSNS client and saving changes will enable automatic iSNS client start up on any subsequent reboot.

To configure and use the iSNS target client, click on the iSNS Options icon on the Options page to bring up the iSNS Configuration page shown in Figure 65: iSNS Configuration.

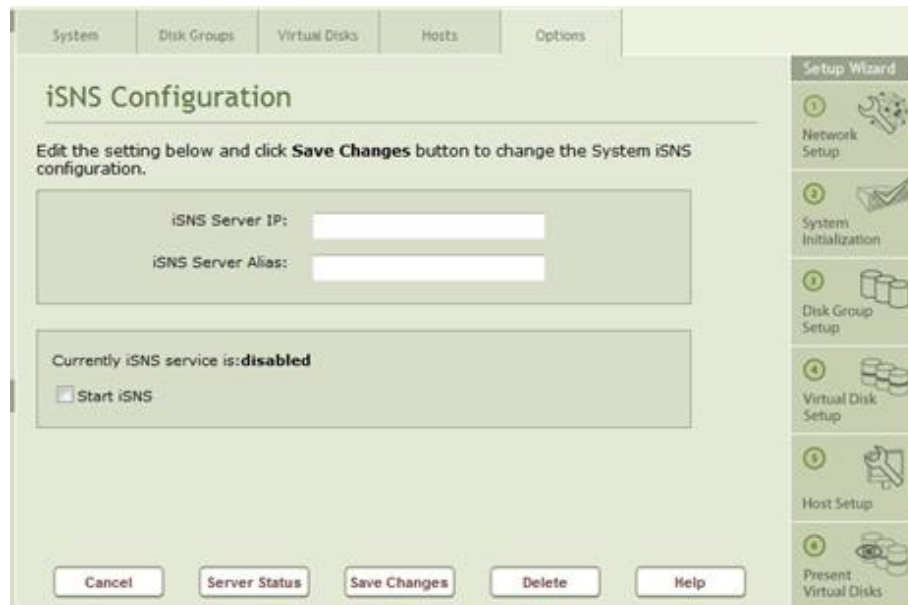


Figure 65: iSNS Configuration

To configure the iSNS client:

1. Click on the iSNS Options icon on the Options page
2. Enter the IP Address of the iSNS server
3. Enter the Server Alias of the iSNS Server
4. Select Start iSNS
5. Click on the Save Changes button

You can also determine the iSNS Server Status by Clicking on the Server Status button of the iSNS Configuration page. This will open a pop up window that will indicate whether the iSNS server you have defined is online or offline.

System Upgrade Options

The PS7012 supports a system upgrade from the Storage Manager. This feature allows you to select the new KDI from your local machine or another system on the same subnet as the system.

The upgrade process is initiated through the Upgrade Manager shown in Figure 66: Upgrade Manager .The upgrade process executes the following actions.

- Load a copy of the specified KDI into memory
- Copy the KDI to the remote controller backup location (typically takes 3-4 minutes).
- Reboot the remote controller with the backup copy moving IP address and Disk Groups to the local controller
- Load the KDI into the backup location on the local controller (typically takes 3-4 minutes).
- Reboot the local controller moving IP address and Disk Groups to the remote controller
- Restore IP addresses
- Restore Disk Groups to primary controller assignments (assuming minimal or no traffic)
- Synchronize backup copy with new KDI (typically takes 3-4 minutes).

Messages are presented in the Upgrade Manger page Upgrade Status area showing the progress throughout the upgrade process. Command response and Storage Manager Refresh will slow during the final step that synchronizes the copies of the KDI.

To upgrade the system software:

1. Click on the System Upgrade icon in the Options page
2. Enter the software filename or use the Browse button to select a new software image
3. Start the upgrade by clicking on the GO button

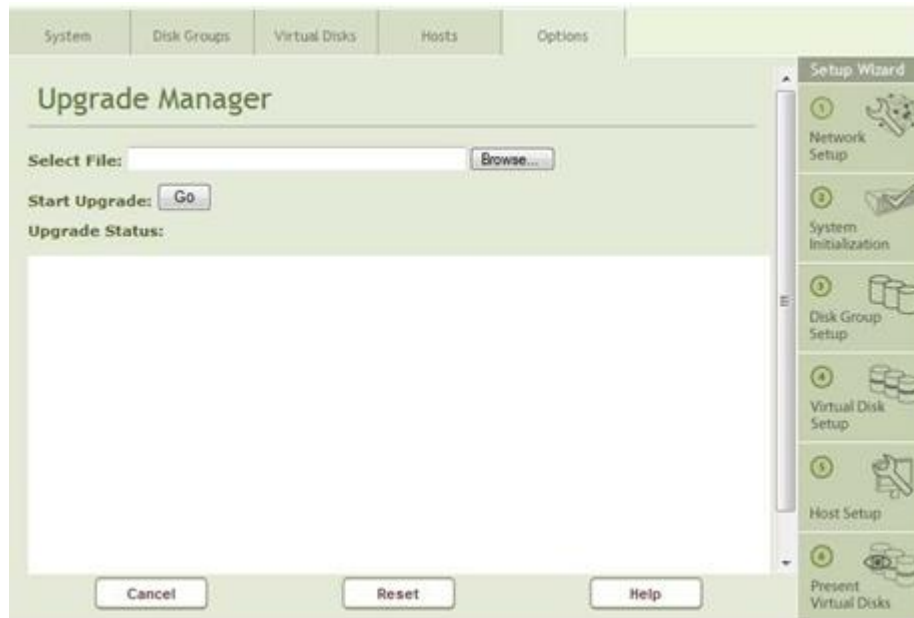


Figure 66: Upgrade Manager

Inspection Read Options

Inspection Read is a utility designed to proactively scan the disks in the storage array to identify any faulty disks before the user attempts to access them.

Priority

Inspection Read may be run in the background at one of three different priority levels. This determines how much time is allocated to Inspection Read by the system. Inspection Read will automatically relinquish the CPU during heavy CPU activity.

Logging

The Inspection Read process reports errors only on the local controller where the process was initiated. Errors and process output is recorded in /var/log/messages and KEM log files.

Inspection Read creates a log message after each device has been successfully inspected.

Devices to Read

Inspection Read will access all sectors within each healthy, non-building; non RAID-0 Disk Group and spare disk drives.

Table 11: Inspection Read Action

Device Type	Action
RAID 0 Device	Do Not Read
Healthy RAID Device (not RAID 0, Not during initial build)	Read
Healthy RAID Device (not RAID 0, During initial build)	Do Not Read
Degraded RAID Device (not rebuilding on spare disk)	Do Not Read
Degraded RAID Device (During rebuild on spare disk)	Do Not Read
Unhealthy RAID Device	Do Not Read
Orphan disk	Do Not Read
Spare disk	Read

Inspection Read monitors the status of Disk Groups and will stop an inspection if the status of the Disk Group changes to a “Do Not Read” state listed in the previous table.

Starting Inspection Read

Inspection Read may be started from the Storage Manager or from the CLI. This section describes using Inspection Read through the Console Manager. To learn how to use Inspection read from the command line interface, issue the command `dainspread -h` from the CLI console.

To start Inspection Read from the PixelStor Storage Manager, select the Options Tab and click on Go Now under Inspection Read as shown in Figure 68: Inspection Read Interface.

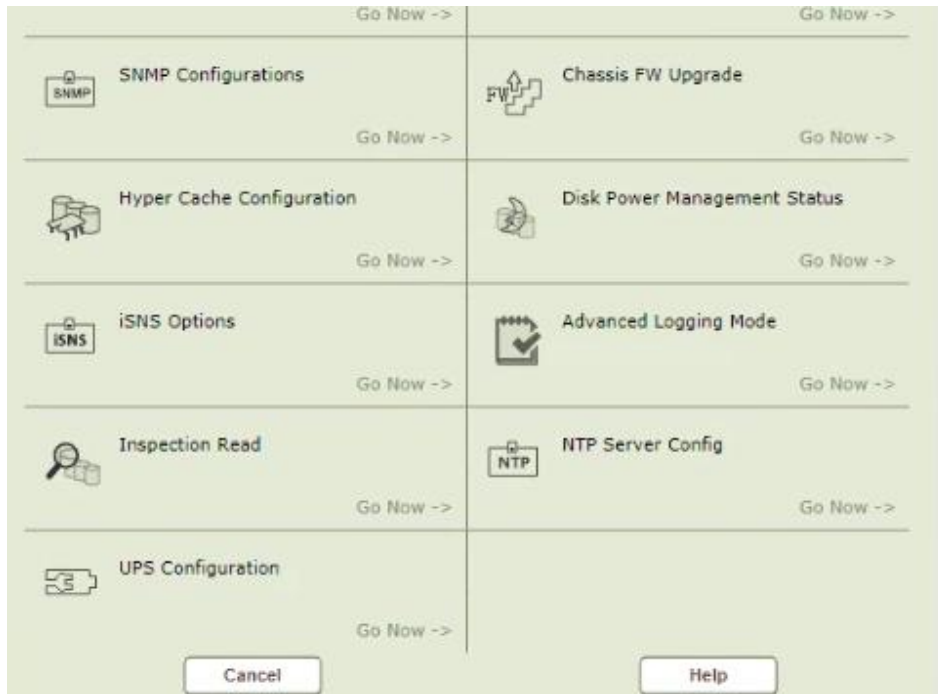


Figure 67: Options Tab

Inspection Read Interface

To run Inspection Read from the Management Console shown in Figure 68: Inspection Read Interface, you must select the priority and any other options you desire then either click on the Start button or the Schedule button to start Inspection Read at a later time.

- Once Inspection Read has started you can use the Stop and Resume buttons to pause or restart Inspection Read from the point it was stopped.
- You can use the Progress button to see information about the progress of Inspection Read on the system as shown in.
- Checking the Repeat box will cause Inspection Read to run continuously until the user clicks on the Stop button.

- Checking the Set Maximum Duration button allows the user to specify a maximum number of hours to run Inspection Read in the Maximum Duration field.

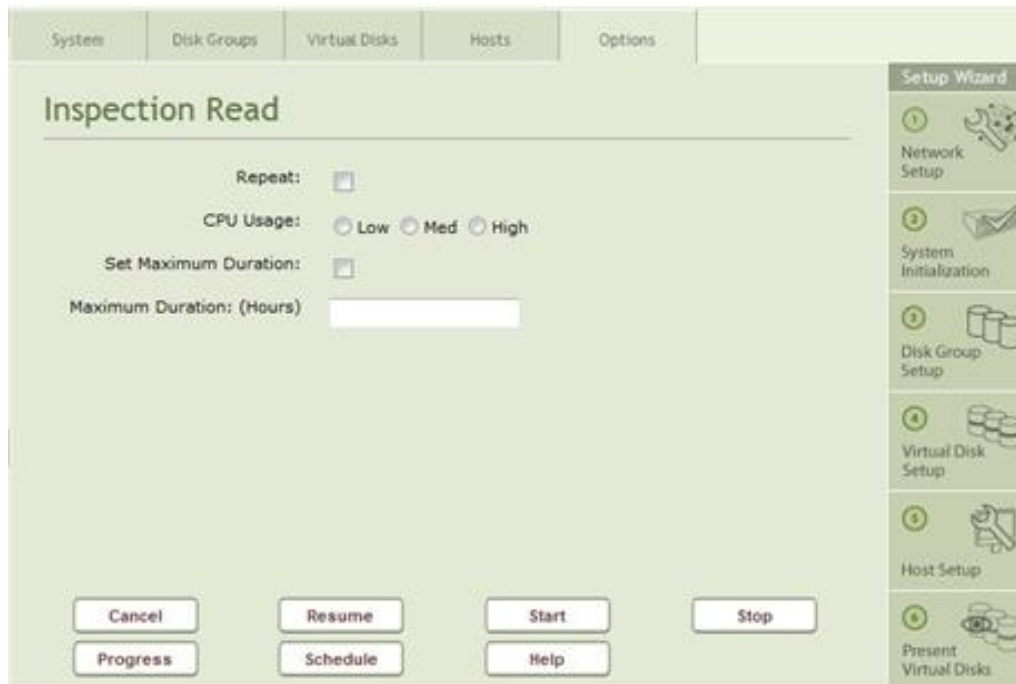


Figure 68: Inspection Read Interface

Inspection Read Progress

Figure 69: Inspection Read Progress shows the Inspection Read progress pages. The display fields include:

- Running – Shows yes if Inspection Read is currently running
- Execution Time – Shows the amount of time that Inspection Read has been running
- Maximum Duration – Shows the total time allocated to run Inspection Read when Maximum duration has been set.
- Device list – Shows the devices and slots for each device under inspection. The Number of blocks inspected, total blocks to be inspected and percentage completed are listed for each device.

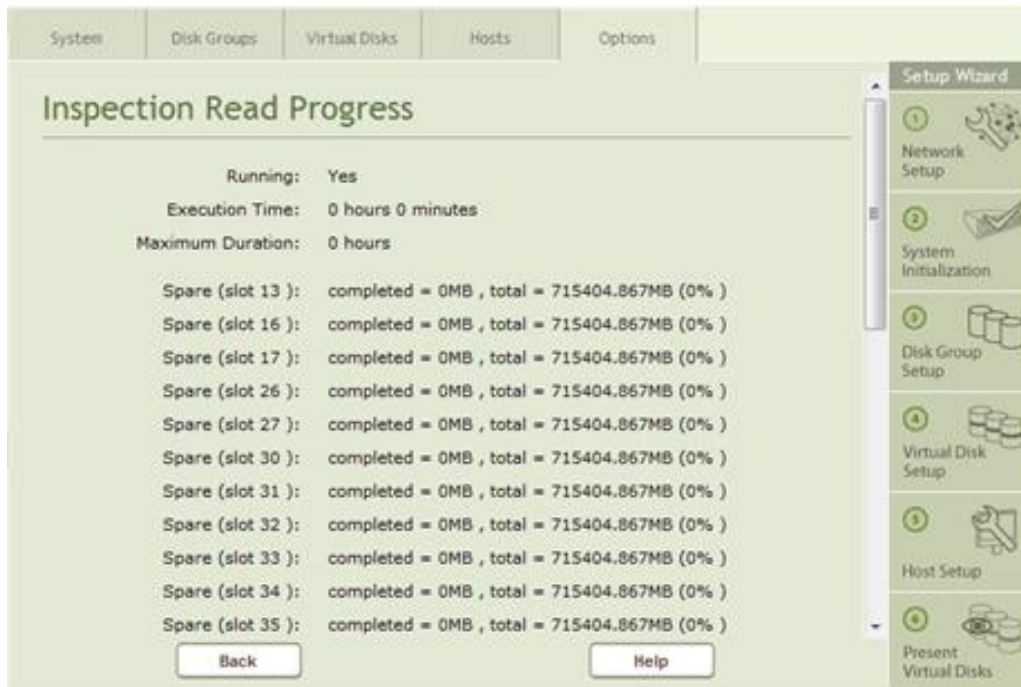


Figure 69: Inspection Read Progress

Language Options

You can select English or Chinese languages as shown in Figure 70: Language Option.

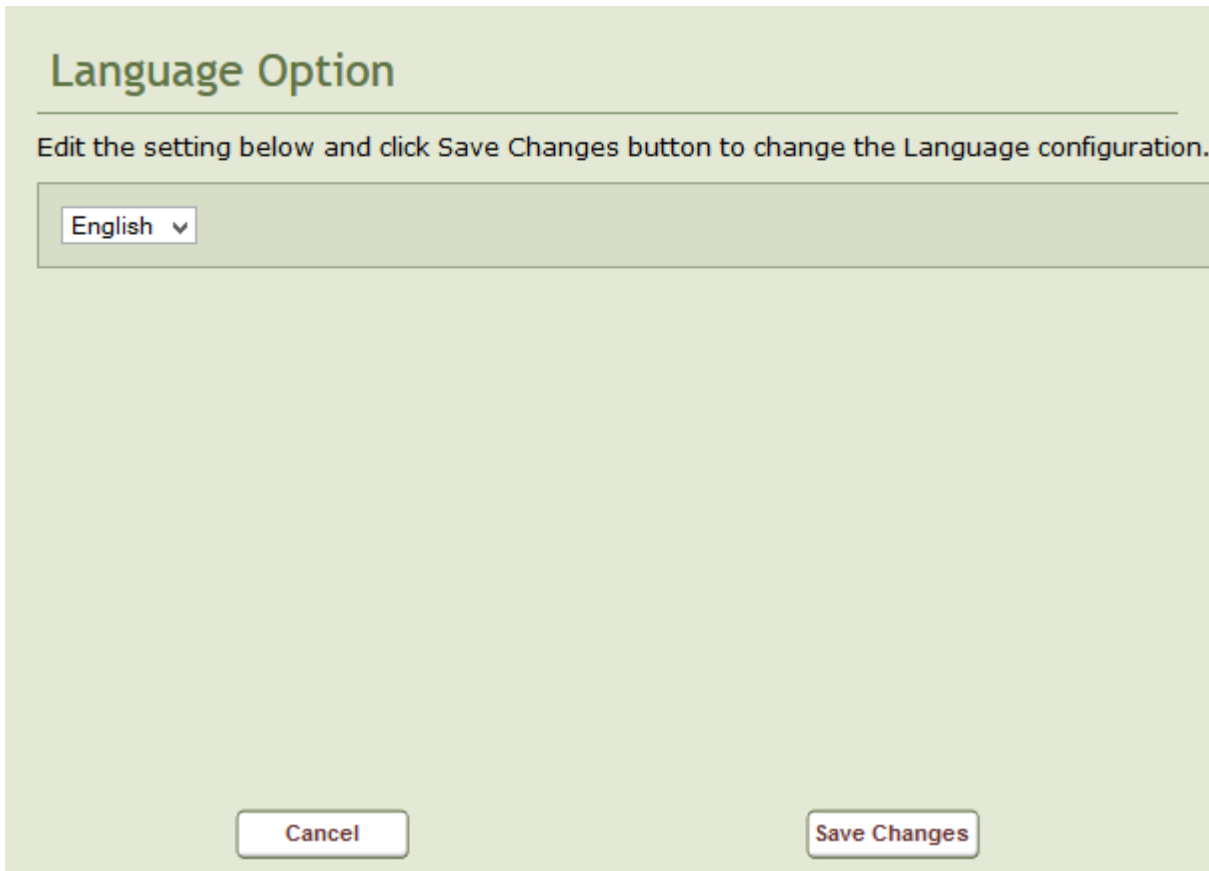


Figure 70: Language Option

Drive Write Cache Options

You can enable or disable the cache on the installed disk drives as shown in Figure 71: Drive Write Cache Options. Enabling the cache increases performance. It is recommended that you have a UPS installed to protect the data that may be in the drive cache in the event of a power failure. If you do not have a UPS installed, it is recommended that you disable the drive cache to maintain data protection.

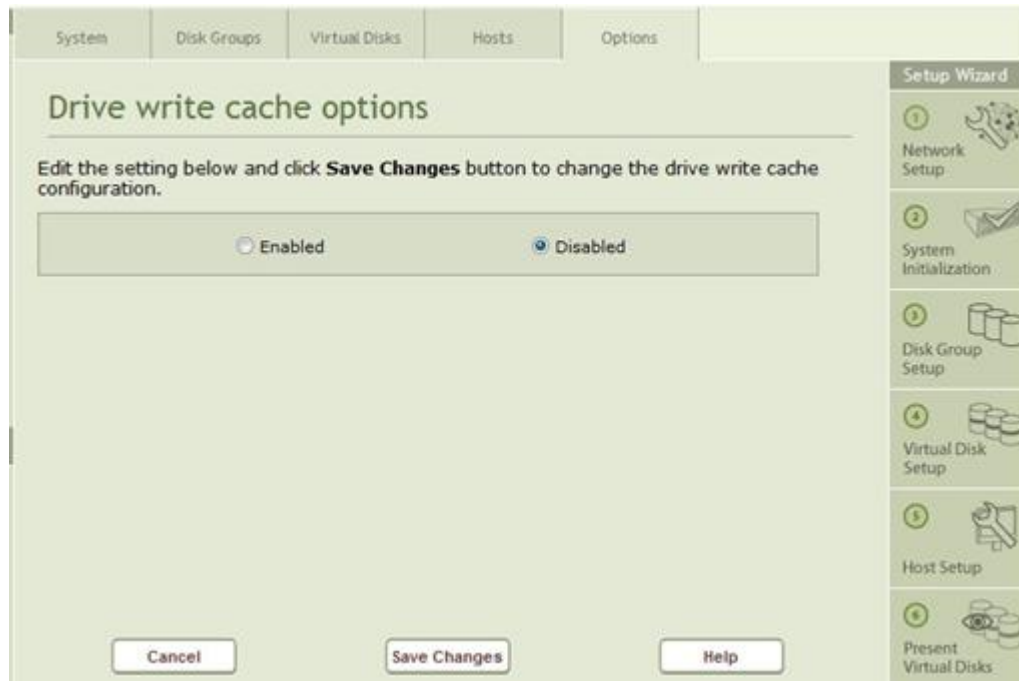


Figure 71: Drive Write Cache Options

To enable/disable drive cache:

1. Click on the Drive Write Cache Options icon in the Options page
2. Select the Enabled or Disabled button
3. Select the Save Changes button

System Beacon LED Options

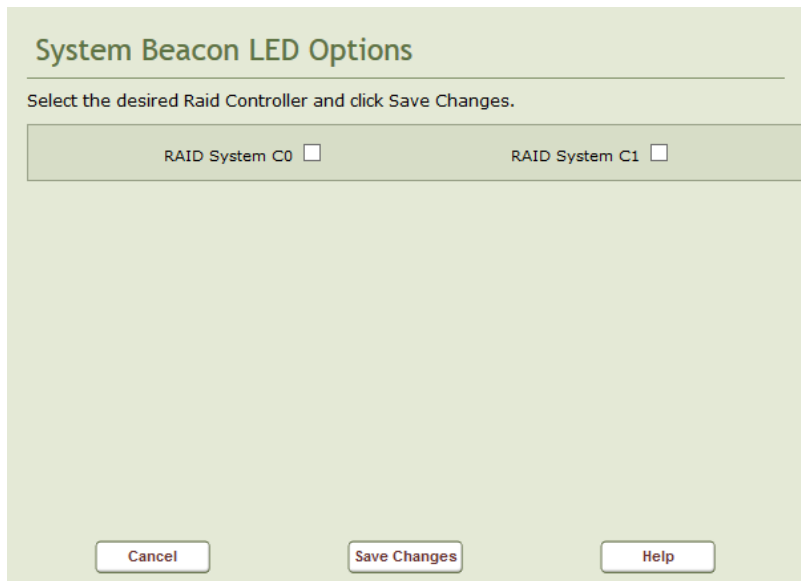


Figure 72: System Beacon LED Options

Performance Data

The PS7012 supports diagnostic data. This data is used to diagnose and monitor the performance of the Storage system and its interaction with the host computers.



In order to view diagnostic information recording must be started.

1. Select the Performance Data page from the 'Options' menu.

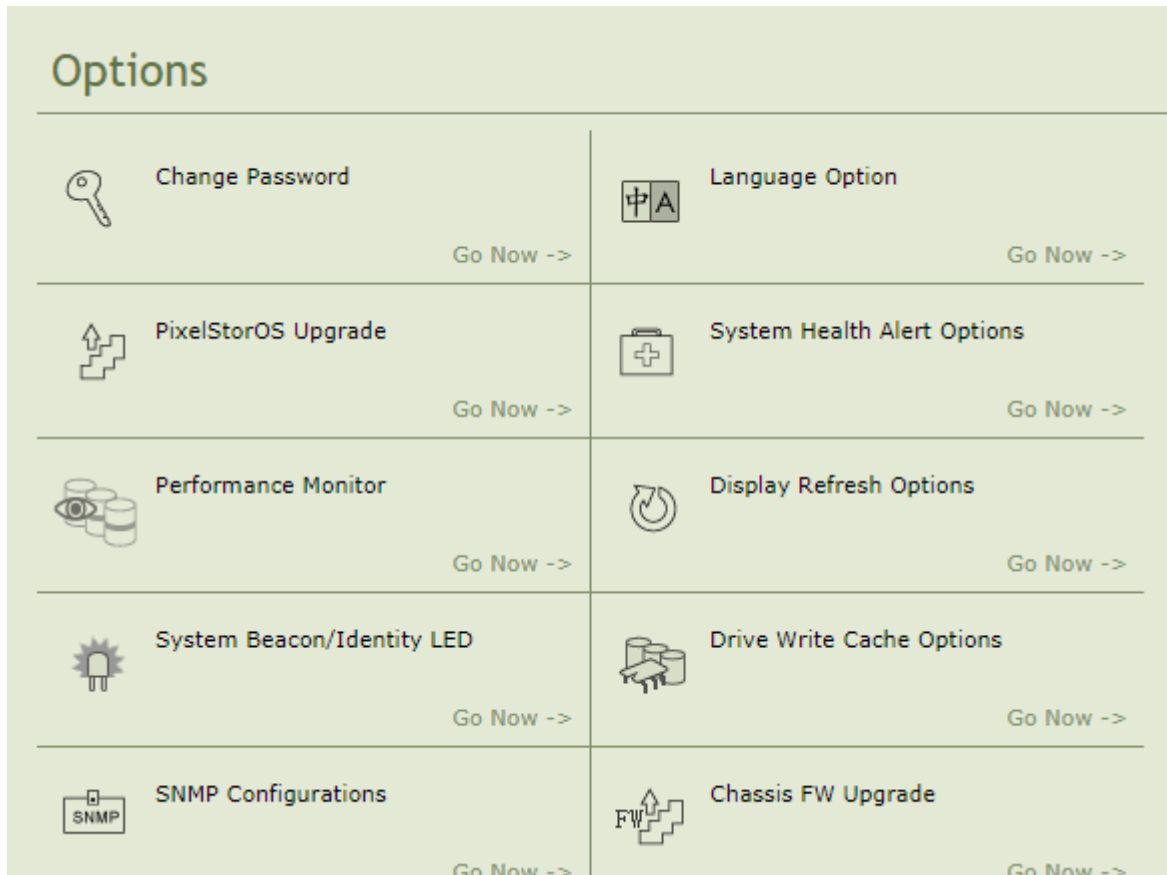


Figure 73: Performance Data

2. Select on "Start" and select sample time as shown in Figure 74: Start Recording Data. The default of 0.5 seconds will normally be used.

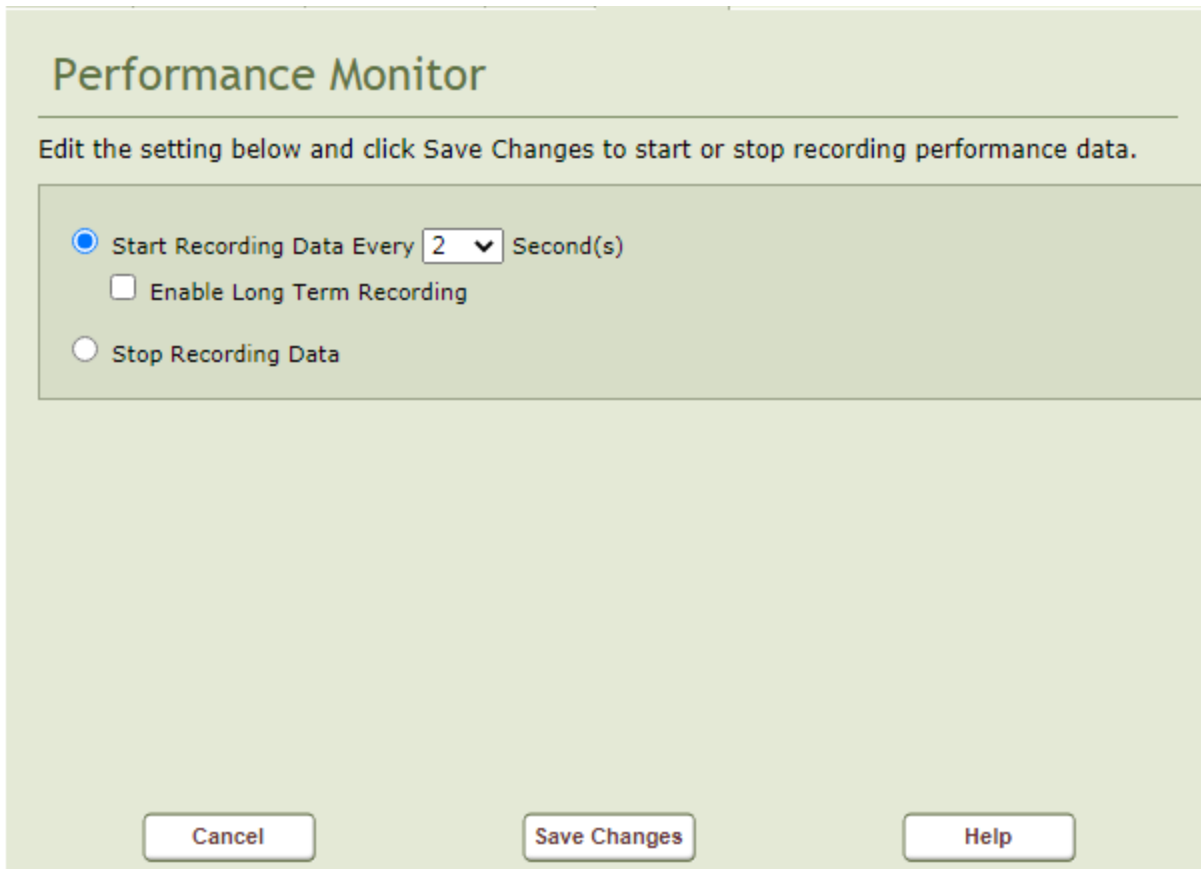


Figure 74: Start Recording Data

Now Performance Data under the system tab may be used.

Disk Power Management Status

Disk Groups

Disk Group Name or Spare/Orphan disks.

Management Mode

This field indicates the current power management status for the disks belonging to this group. If the status is "off", then the system will not manage the power status of the drives in this group. The drives in this group remain in the active state all the time. If the status is "on", then the system will manage the power status of the drives in this group. The drives will go into standby mode after the specified inactivity period.

Inactivity Timer

This field specifies how long the drive should wait before it transits from active mode to standby mode if there is no activity. Valid values: 0, 10, 20, 30, 60, 90, 120, 150, 180, 210, 240, 270, 300 and 330 minutes. A value of 0 disables spin down.

Scheduling Information

This field specifies whether power management scheduling is enabled for this group. If power management scheduling is enabled, it displays the time interval that the drives remain in the active mode. Outside of this time interval, the drives will go to standby mode after the specified period of inactivity.

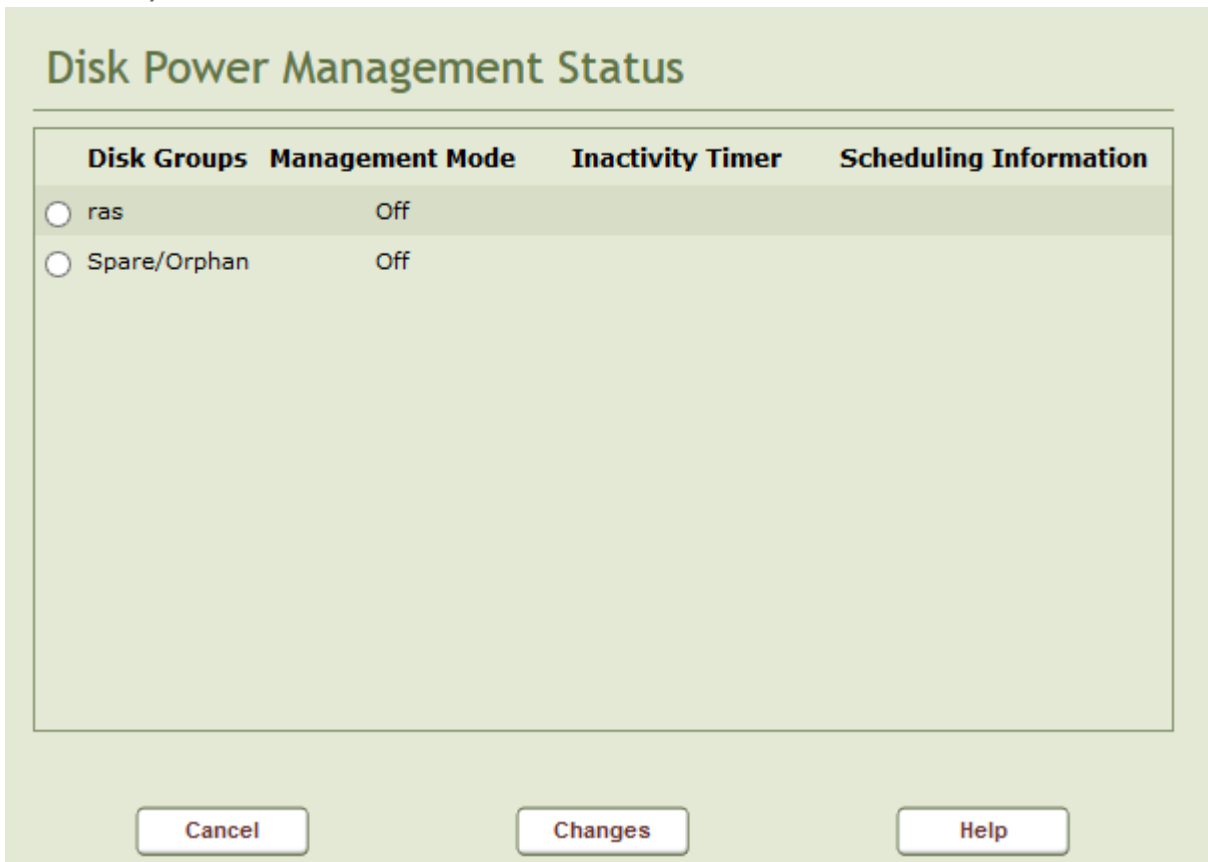


Figure 75:Disk Power Managemet Status

Advanced Logging Mode

- Disabled (default) - This is the default logging mode. We recommend to use it when the system is normally running.

- Enabled - More detailed log information will be recorded. Please use it only when debugging problems.

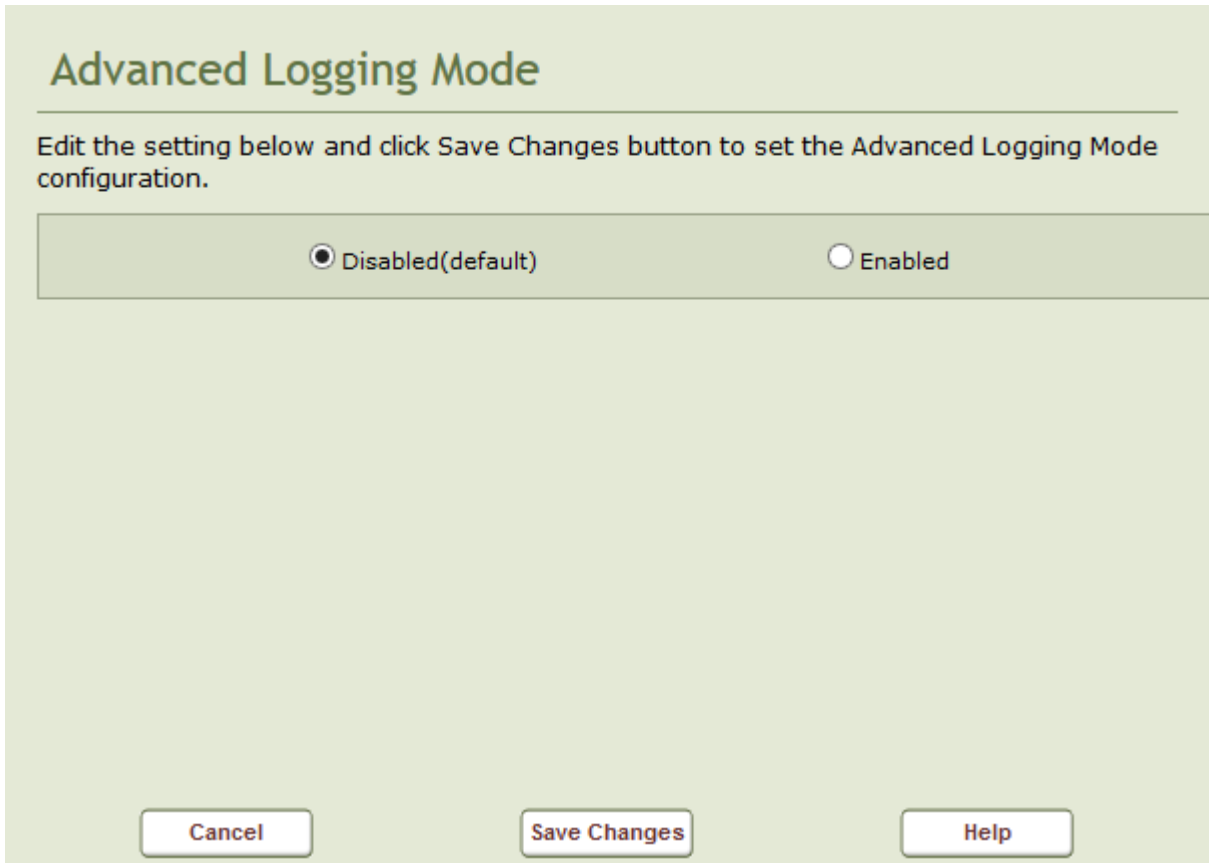


Figure 76: Advanced Logging Mode

Firmware Upgrade Manager

Select File

Enter the absolute path to the firmware file or use the "Browse" button to locate the file. You may obtain the file from another server by entering the IP address of the server and the path to the file in the following format:

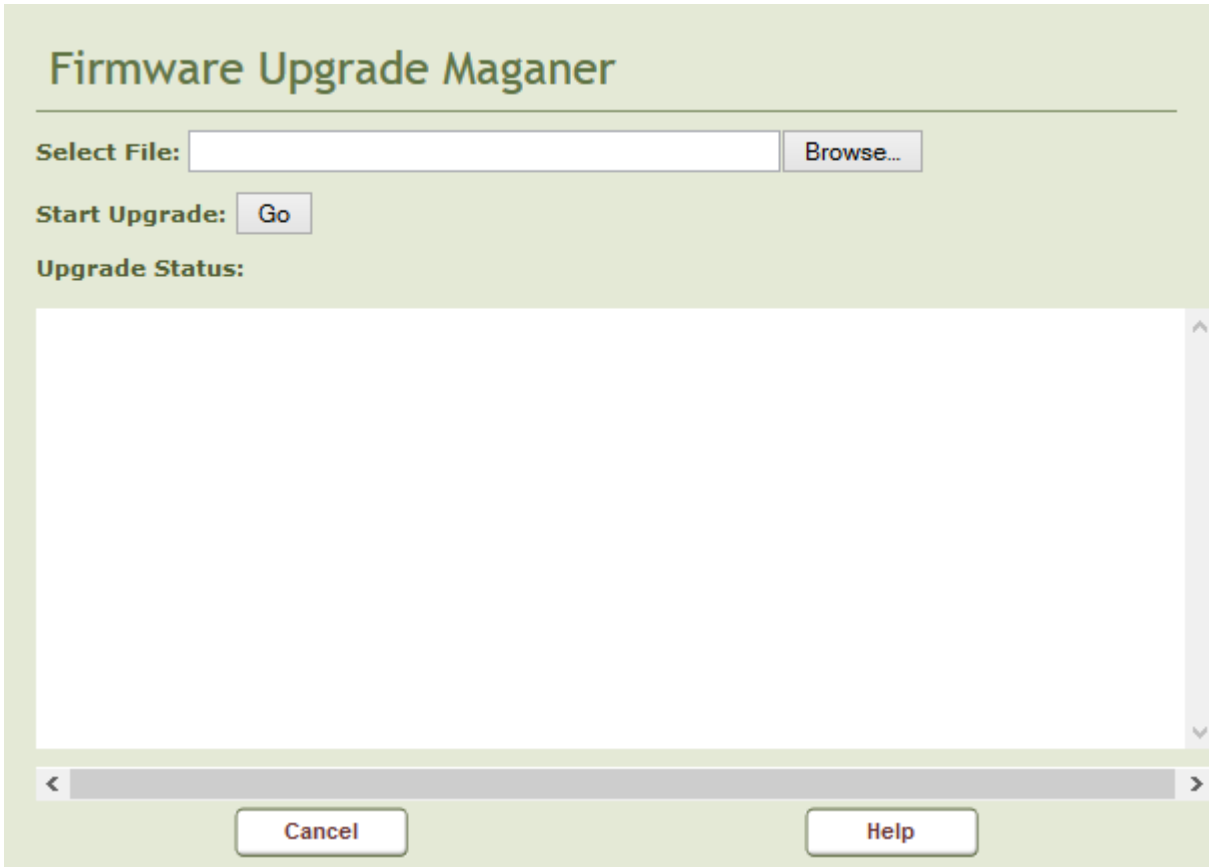
\\192.168.0.250\home\firmware.fuf

Start Upgrade

Click on the "Go" button to begin the upgrade

Upgrade Status

Display status messages and relative progress



The screenshot shows a web-based interface titled "Firmware Upgrade Maganer". It features a "Select File:" label followed by a text input field and a "Browse..." button. Below this is a "Start Upgrade:" label with a "Go" button. Underneath is an "Upgrade Status:" label followed by a large, empty text area with a vertical scrollbar on the right. At the bottom of the interface are two buttons: "Cancel" and "Help".

Figure 77: Firmware Upgrade

Hyper Cache Configuration

Select a Primary Controller

Select the owner of this disk group.

Total Capacity

Total Capacity of the SSD Cache Drive

Hit Rate

Last measured Hit Rate of SSD Cache Drive

Select Hyper Cache members

Select among all the spare disks to be the Hyper Cache members. Circle is SAS drive. Drive in blue is SSD.

VD Hyper Cache Configuration

Selects high or low priority of VD Hyper Cache. Selects enable or disable VD Hyper Cache.

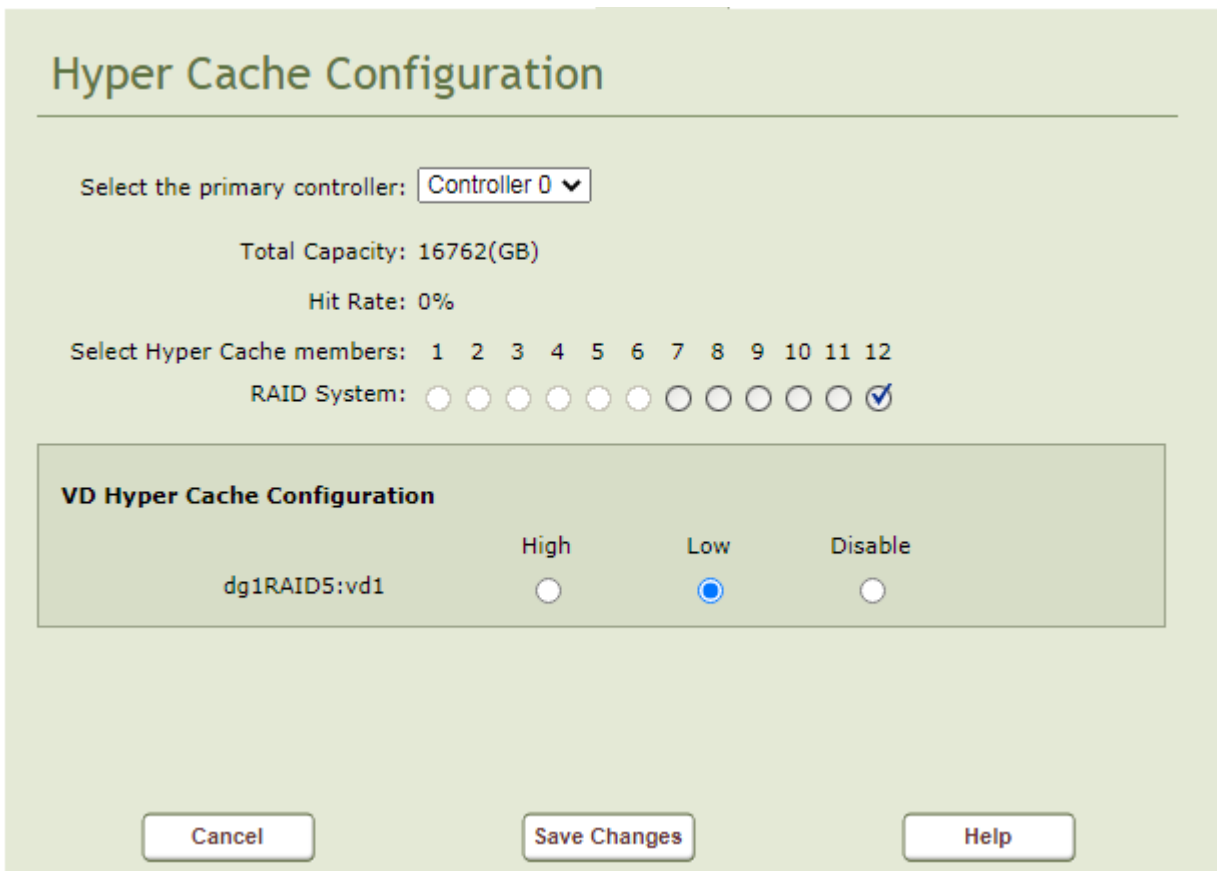


Figure 78: Hyper Cache Configuration

Configuring Notification

The PS7012 supports email notification for alerts and major/critical events. You specify email subscribers to be notified through the notification page of the Management Console.



In order to receive email, you need to make sure you have correctly set up the default Gateway and DNS server in the network setup.

1. Select the Email Notification page from the 'System Status' menu.

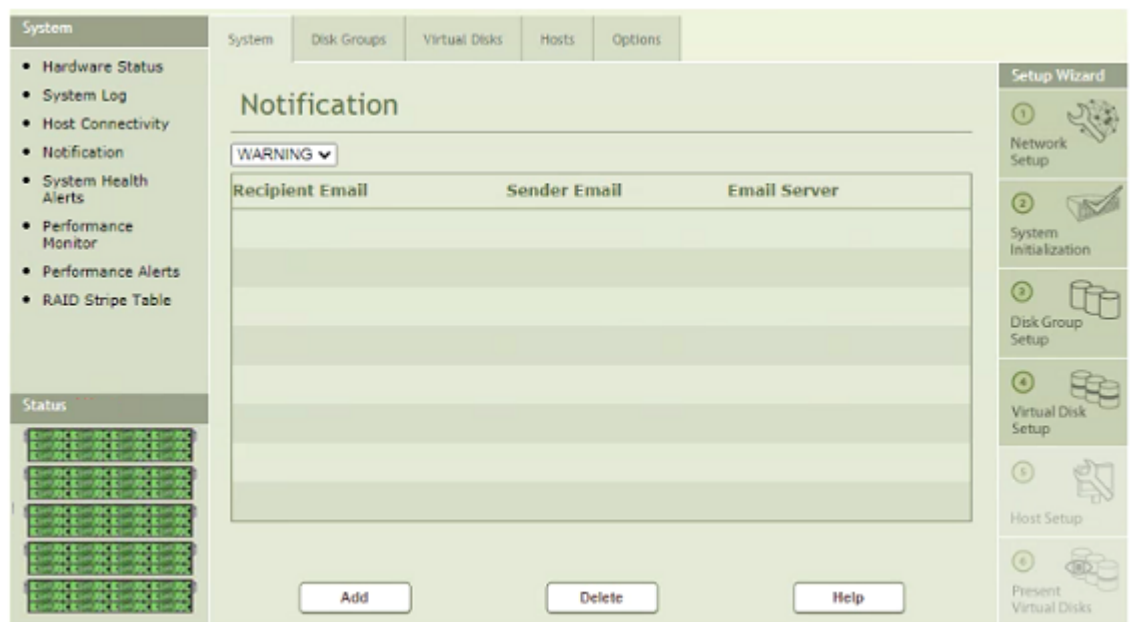


Figure 79: Email Notification

2. Click on "Add" to bring up the add user page as shown in Figure 80: Adding Email Addresses.

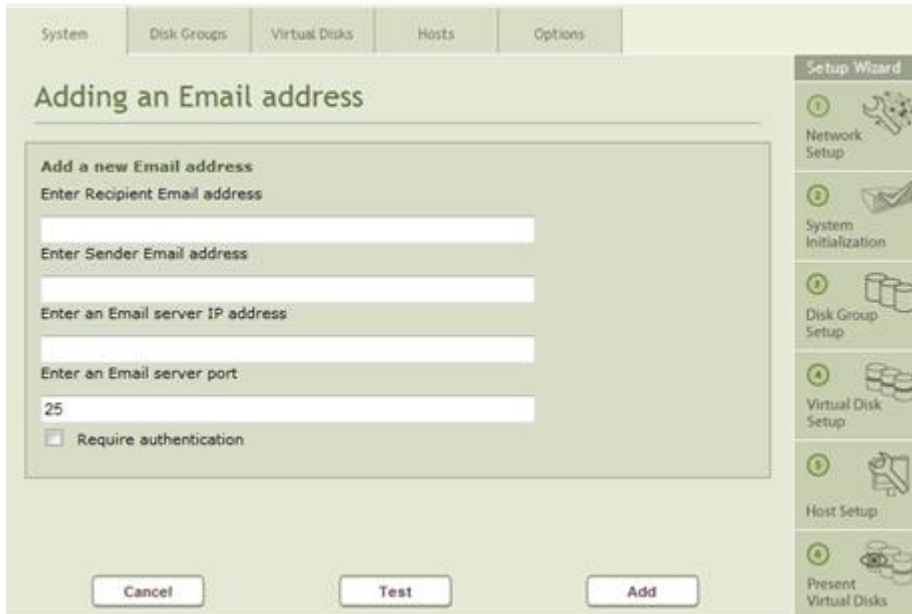


Figure 80: Adding Email Addresses

Enter

- Recipient's Email Address
- Sender's Email Address
- IP address of the mail server
- Port number for the email service on the email server
 - Click on the "Add" button as shown in Figure 80: Adding Email Addresses

Chapter 4: Use and Maintain

System Status
Failure Conditions
Controller Modules
Storage
System Memory
System Battery

Using and Maintaining the System

System Status

This section includes instructions for using, maintaining and servicing your PS7012 . It is organized from a reference perspective and gives details on each major configuration screen and hardware module and its associated configuration features. This section also reviews placement of LEDs and what they indicate.

Properties Page

The System Properties Page is displayed when you log onto the system. This page provides an overview of the system as shown in Figure 81: Storage System Properties.

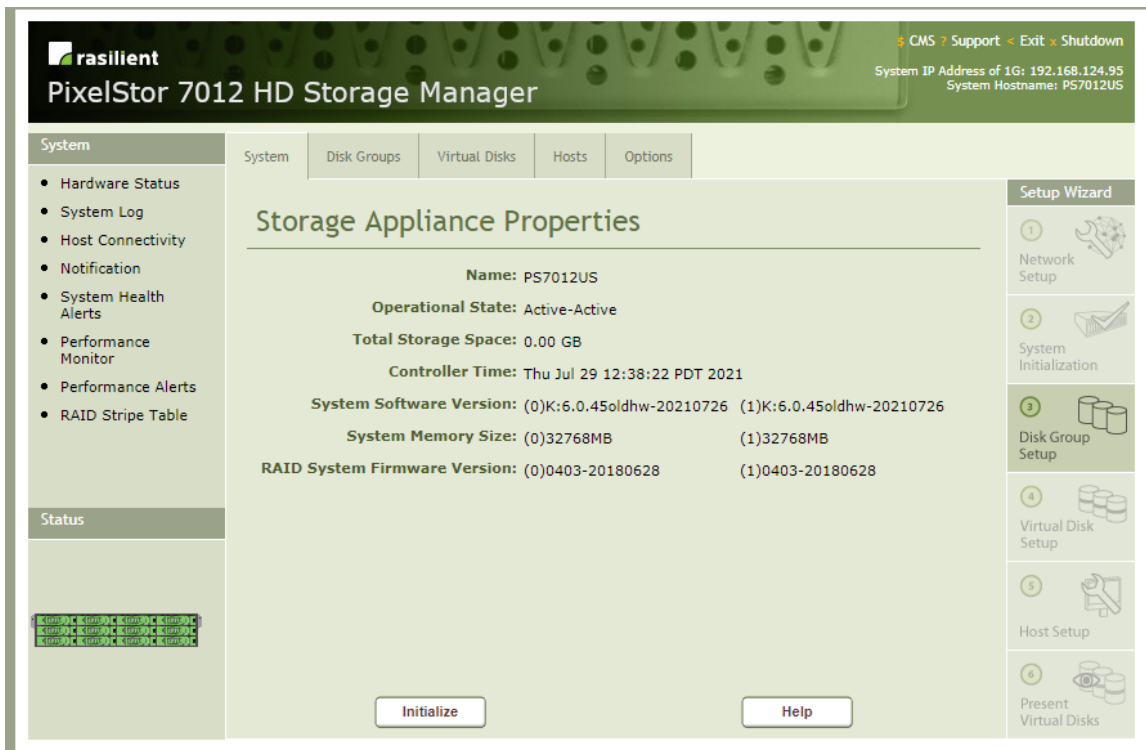


Figure 81: Storage System Properties

The fields in the Storage System Properties page are:

- Name - The hostname of the system

- Operational State - Describes the operational state of each controller in the system
 - Active/Active – Both controllers active running normally.
 - Active/Takeover – One Controller in takeover state. Its IP addresses and disk groups have been taken over by the other controller.
 - Active/- - Single controller active
- Total Storage Space - Displays the total capacity of the system in Gigabytes
- System Software Version - Describes the software version of the local (L) and remote (R) controllers (if two controllers installed)
- Expansion Firmware Version - Describes the F/W version of each enclosure installed in the system.

Hardware Status Page

The hardware status page provides a visual representation of the system status. Moving the mouse over individual components provides additional details about the component. For example moving the mouse over a disk drive will display the manufacturer, model, size and firmware revisions as shown in Figure 82: Hardware Status



Figure 82: Hardware Status

Additionally a failed component will display in red.

System Log

The System Log is accessible through the Storage Manager. To get to the System Log display, select System Log in the System list on the left hand side of the Storage Manager. This will bring up the Log viewing screen as shown in Figure 83: System Log. You can select the level of detail through the pull down menu. Your choices include, info, warning or critical messages.

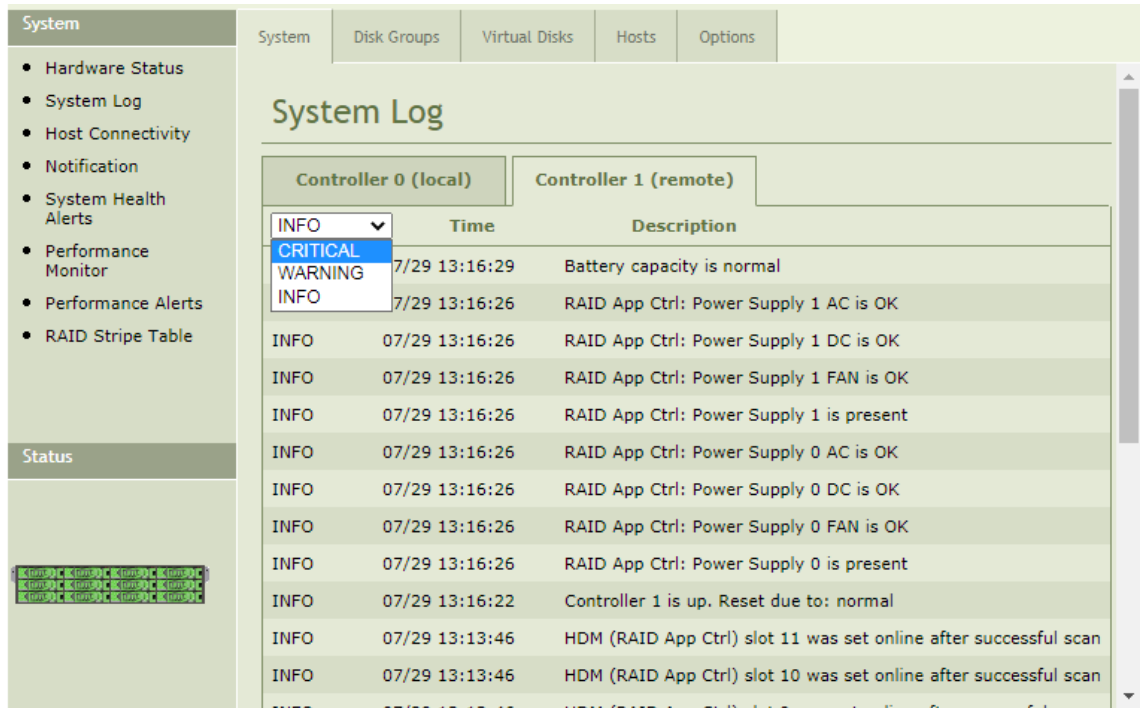


Figure 83: System Log

Support Menu

A support menu is offered for the PS7012 . The support menu can be viewed by selecting the support link at the top right had corner of the web management console. The support window includes information on how to contact technical support. In addition, there are two buttons at the bottom of the menu for saving and sending logs to support contacts.

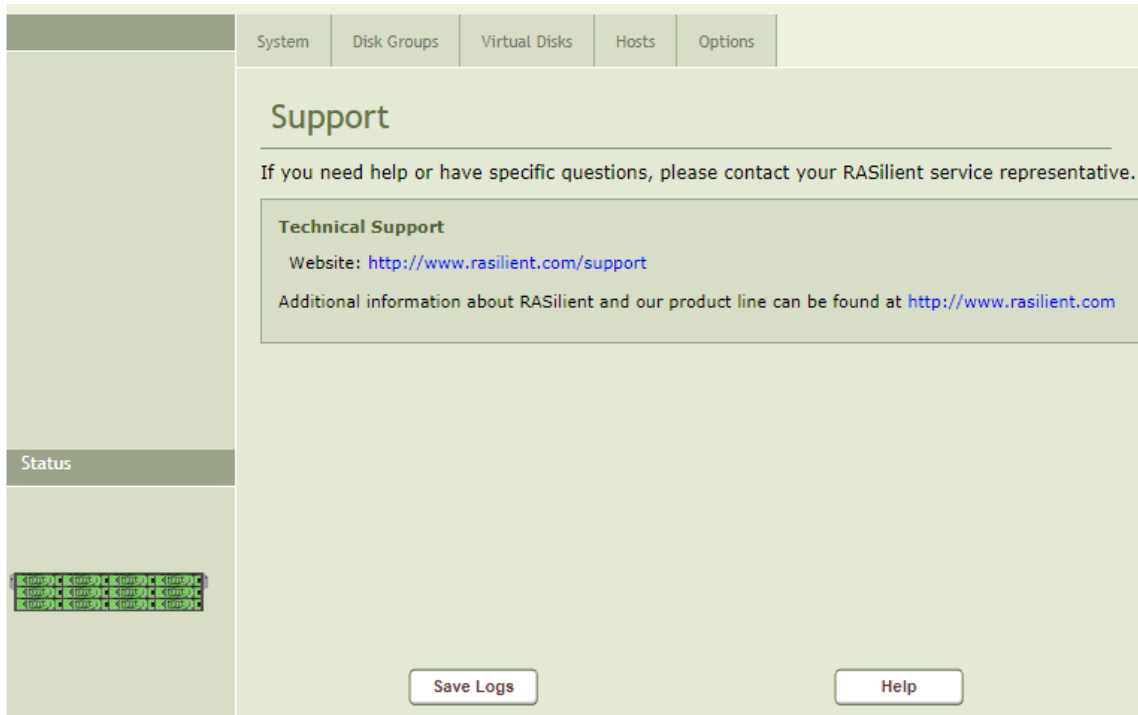


Figure 84: Support Menu

- Save Logs - To save logs to a local or network drive, just select the "Save Logs" button at the bottom of the support menu. Once selected a window will pop up asking if you wish to open or save the file. Select save and choose a location to save the file. The saved file contains logs from one or both controllers (depending on your configuration) and an inventory of the system. This can be used as a record of the status of the system or emailed to a support representative.

Performance Data

Figure 85: Performance Monitor Page shows the Performance Data page. A wide variety of data can be monitored for the storage system. There are 4 selections that can be made Virtual Disk Performance, Cache Count, iSCSI Busy Percentage and CPU Idle. Within the Virtual Disk Performance selection each individual VD can be

select to see how it is performing. This performance is measure with Write Data Rate, Read Data Rate, Command Completion Time and Command LBA.

Cache Count monitors the system cache usage. iSCSI Busy monitors busy messages sent to the host computers. CPU Idle monitors CPU usage.



The information displayed in the graphs is actual data derived from the PS5012.

In Figure 85: Performance Monitor Page is an example the selects the DG1 and vd1. It also selects the write data rate for that VD. Figure 86: VD Data Rate Display shows the output of this selection.

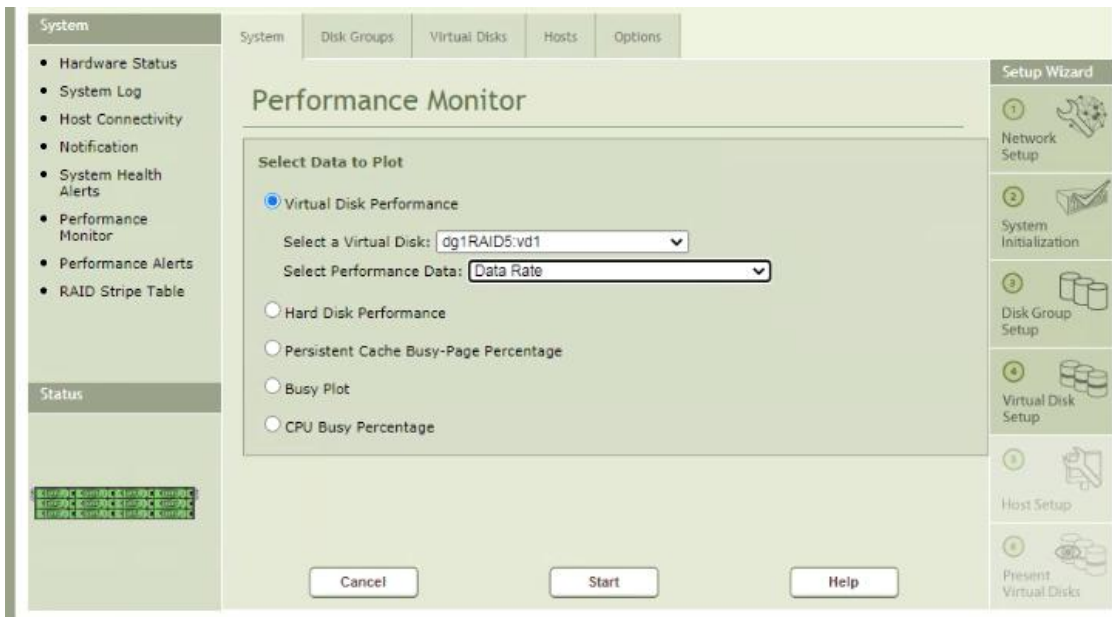


Figure 85: Performance Monitor Page

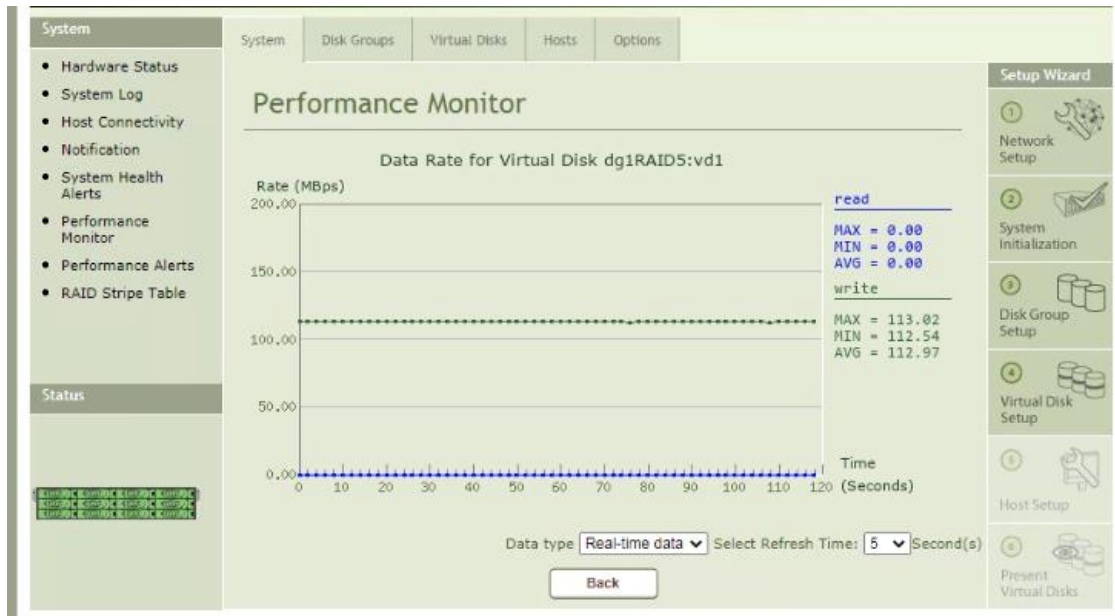


Figure 86: VD Data Rate Display

System Health Alert

RASILIENT Systems also equips PS7012 with a System Health Alert which is a flashing alert visible from the front of the system.

Description – The System Health Alert lets administrators know that the system needs attention. The System Health Alert utilizes all the disk drive LEDs and uses them to flash a warning signal. All the LEDs on the disk drives blink approximately once per second. These lights continue to flash until redundancy is restored or the unhealthy condition is corrected. This visible alert is especially useful in dense rack implementations or situations where it is not easy to see or access the back of the system. Once alerted, you can log-on to the web-based PixelStor Storage Manager to ascertain the cause of the problem and use the system’s Status Indicators (LEDs) to pinpoint the failed component and take corrective action.

Policy – The System Health Alert is triggered whenever system availability degrades. For example, the Health Alert lights flash when component failure or removal causes system availability to drop below its original startup condition. It is also triggered by an unhealthy system state such as an exceeded temperature threshold.



PS7012 does not require redundant components to operate. If the system boots up without network links or controllers, the System Health Alert is NOT triggered. If one of the two power supplies has no AC power, alert will be posted.

To View the System Health Alerts, click on System Health Alerts in the navigation area of the System tab.

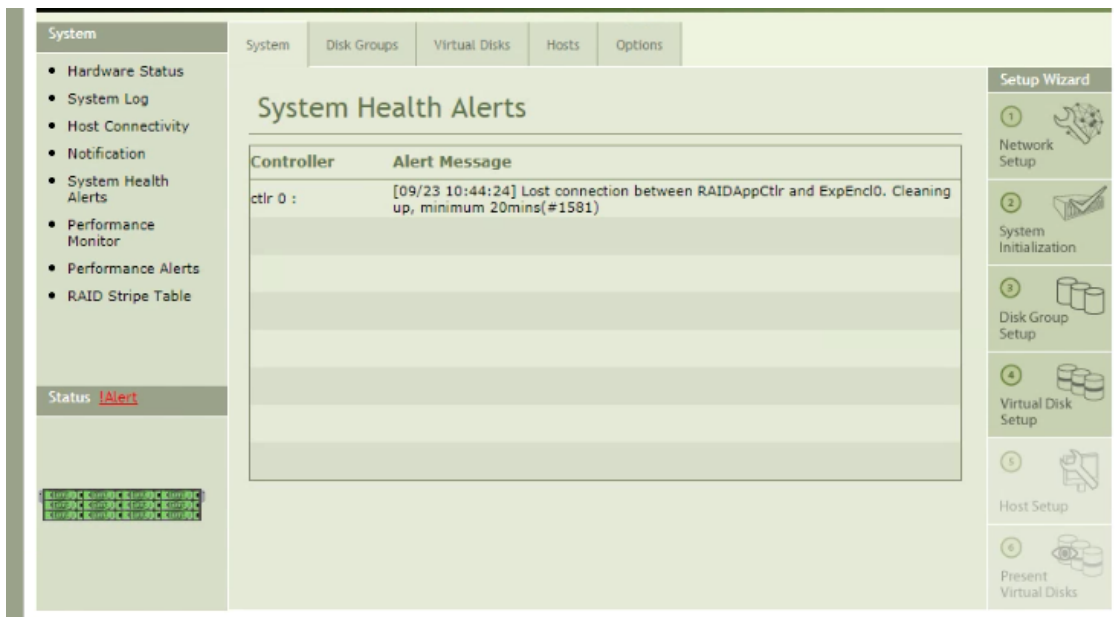


Figure 87: System Health Alerts Page

Failure Conditions and Corrective Action

This section describes specific events in the system that requires corrective action and the corresponding event or alarm.

Degraded System Availability

When the PS7012 is turned on, it captures an inventory of installed components and connected interfaces. Once the user configures Disk Groups, the system also tracks the RAID configuration information. If at any time, the system falls to lower system availability than at startup, the System Health Alert is triggered. This can happen if a redundant component fails or a RAID 1, 5, 6 or 10 disk(s) fail or if an interface fails to respond.

Corrective action:

- Check the PixelStor Storage Manager for the source of the problem.
- Restore redundancy by replacing the failed controller, power supply or failed drive. Once system availability is restored the System Health Alert will stop flashing.
- If the system has a global spare, the System Health Alert will stop flashing once the RAID group is automatically rebuilt by the system.

Turning the System Health Alert Off

Once a system administrator identifies the problem and schedules corrective action, he may want to turn off the System Health Alert. To shut off the flashing lights, the administrator needs to:

- Log-on to the PixelStor Storage Manager.
- Once logged-in, click the "Options" menu.
- In the options menu, click "System Health Alert Options"

- There are two option buttons “enable and “disable”. Hit the disable button to shut off the System Health Alert.



Although the System Health Alert is disabled, the system will continue to log error notifications in PS7012 's system log. Administrators can view the notifications in PixelStor Storage Manager's "System Log" or "System Health Alert" sub-menus.

RAID Controller Failure

When a RAID controller fails in the PS7012 with redundant controllers, all drives including those on interconnected PixelStor enclosures will fail-over to the surviving controller without interruption to data access. If the PS7012 is in a single controller configuration the system will shut down and all data will not be available

Corrective action:

- Dual Controllers - replace the failed controller and the system will automatically fail-back to original state.
- Single controller - replace the failed controller, reconfigure the host access and rediscover the drives groups.

SAS IO Expander Failure (Lost Connection to Expansion Enclosure)

PS7012 capacity can be expanded by adding up to 8 additional capacity expansion enclosures for a maximum capacity of 108 SAS drives. These enclosures are daisy chained to the PS7012 SAS IO expanders and 12Gbps SAS connections.

In the event a SAS IO expander fails, all expansion enclosures and corresponding drives connected to the failed IO expander will become unavailable. Under this scenario, the system will treat all drives associated with that SAS expander as failed. Since disk groups can span across the PS7012 and interconnected expansion enclosures, a loss of one or more expansion enclosures can affect multiple disk

groups. If one drive for RAID 1/5/10 or two drives for RAID 6 are missing from the associated Disk Groups, the system will rebuild them using any available global spare. If multiple drives (2+ for RAID 1/5/10 or 3+ drives for RAID 6) are missing then all writes to the DG will be suspended and the DG will be made available to the system in a read-only state.

Corrective action:

- Dual Controller - the system will fail over and the controller with healthy IO expanders and cable links will take over all the disks without disruption to data access. The administrator can replace the failed expander and the system will automatically fail-back to regain redundancy.
- Single Controller - all connected expansion enclosures (and drives) will become unavailable. The administrator will need to replace the IO expander and do a Disk Group Restore to reactivate all the disk groups brought off-line by the failure.

SAS Cable Disconnect (Lost Connection to Expansion Enclosure)

If a cable fails or is accidentally disconnected, the system will react as if the SAS IO Module failed. (See 5.2.4 "SAS IO Expander Failure" in previous section)

Corrective action:

- Dual Controller - the system will fail over and the controller with healthy a SAS IO connections will take over all the disks without disruption to data access. The administrator can reconnect the SAS cable and the system will automatically fail-back to regain redundancy.
- Single Controller - all connected expansion enclosures (and drives) will become unavailable. The administrator will need to reconnect the SAS cable and do a Disk Group Restore to reactivate all the disk groups brought off line by the failure.

Drive Failure

If a one drive fails in a RAID 1/5/10 DG, or two drives in a RAID 6 DG, the DG will go into degraded mode. If a global spare is present in the system, the spare will be automatically be added to the DG and a rebuild of the DG will be initiated. (Note: RAID 0 by definition has no protection and therefore cannot be rebuilt.)

Corrective action:

- No Spare – Replace the failed drive and the system will automatically start a rebuild of the DG. (It is recommended that a global spare is always installed in the system to insure data availability.
- Spare Present - Replace failed drive. This will automatically be added to the system as a global spare.

Unhealthy RAID Disk Group

If a RAID Group experiences a drive fault that is unrecoverable, the System Health Alert is triggered. For example, this can happen if a RAID 0 Disk Group experiences a drive failure or a RAID 5 Disk Group experiences two consecutive drive failures.

Corrective action:

- Recover the Disk Group from a recent back-up.

ECC Memory Error

The system uses ECC memory on each controller. If an ECC memory error is detected by the system, the System Health Alert is triggered.

Corrective action:

- Check the Web-based console for the source of the problem.
- Replace the controller if necessary.

Faulty Cache Battery/Charger

Each controller comes with a battery backup unit to protect the cache. If the battery fails or stops holding a charge, the System Health Alert is triggered. The backup battery subsystem includes a charger for the backup battery. If the charger becomes defective, the System Health Alert is triggered.

Corrective action:

- Check the Web-based console for the source of the problem.
- Replace the controller if necessary.

Network Connection Lost

If a network link is lost, the System Health Alert is triggered.

Corrective action:

- Check the Web-based console for the source of the problem.
- If there are no system problems, check your network devices or initiators for the source of the problem.

Temperature Threshold Exceeded

The system has multiple temperature sensors with factory thresholds. If any temperature sensor indicates an abnormal value beyond the threshold, the System Health Alert is triggered. The Health Alert signals the fault by flashing all the red drive LEDs every 3 seconds.

Corrective action:

- Log-on to PixelStor Storage Manager. Follow the menu tree on the left side of the screen to “System Status” and click on it. Check the graphical representation of the system for errors or click on the “System Health Alert” sub-menu for details on the alert.
- If the PixelStor Storage Manager indicates a faulty module, go to the system and look for a module with a red Status Indicator. Replace the component (drive, fan, controller, power supply). Once the faulty module is replaced, the System Health Alert will stop flashing.
- Check for missing slot covers. Module slot covers are required to ensure proper system airflow. Replace missing slot covers to ensure temperature stays within specified system limits.
- If the PixelStor Storage Manager indicates a temperature fault and there are no faulty system modules, check for a sudden rise or drop in the room temperature. If the room temperature is beyond those listed in the system specifications section of this manual (See System Specifications) shut down the system until you are able to correct the problem.

Voltage Threshold Exceeded

The system is also equipped with voltage sensors. If any voltage sensor reading indicates an abnormal value beyond the factory threshold, the System Health Alert is triggered.

Corrective action:

- Check the Web-based console for the source of the problem.
- Check Status Indicator on each power supply. Replace if faulty.
- Ensure that the system is connected to an appropriate power source such as the rack’s power distribution unit (PDU) or an uninterruptible power supply (UPS).

Fan Failure

Fan state and fan speed are also monitored. If a fan fails or if the fan speed falls below the factory threshold, the System Health Alert is triggered.

Corrective Action:

- Check the Web-based console for the source of the problem.
- If the Power Supply fan has dropped below the recommended RPM that Power Supply should be replaced.

List of Health Alert Notifications:

- The other controller is down
- RAM ECC error is detected
- Battery not present or not in normal charging state
- Battery is discharging unexpectedly
- Temp sensor reading Tx of xx.xx C > alarm threshold
- Voltage sensor reading Vx of xx.xx out of range
- A Disk Group becomes degraded or unhealthy
- Power supply 0 is not operational
- Power supply 1 is not operational
- Ethernet link is down
- Incompatible KDI – The other controller is upgrading this controller
- Incompatible KDI – The other controller is syncing itself
- Disk Group ([name]) ([raid x/y]) has degraded raid devices, cannot be discovered if rebooted

Controller Modules

The controller modules are designed to hot plug into the chassis and they fit into one of the two controller slots of the PS7012 . The controller modules monitor and control every aspect of the storage subsystem. They control chassis management, control path and storage data path.

Controller Features

CPU	Intel Xeon Broadwell
Controllers	Dual Redundant Active/Active mode
Memory	Up to 64GB DDR4 ECC (Per Controller)
Storage	Up to 12x 3.5", up to 18TB SAS 3.0 (12 Gbps) drives, (main chassis) Supports up to 108 drives with external JBODs
Cache Type	Mirrored/Persistent
Battery Backup (BBU)	Controllers have protection for memory contents if power loss occurs. Memory data is backed up to an on-board SSD. No time limit for power loss.
Operating Temperature	10°C - 35°C
Networking	2x 1GbE, 2x 10GbE RJ45 per controller Optional 2x SFP+ Fiber ports
Input/Output	2x USB, 1x Serial COM, 1x SAS (for expansion)
Dimensions (WxHxD)	17.6"x 3.5"x 20.9"/ 447 x 89 x 531mm
Weight	62.5 lbs/28.4 kg
Power	89-264VAC, 47Hz-63Hz input, via IEC C-14
Watts	800W, 80plus Gold Level, 1+1 redundant, hot swap PSU

This chapter provides the detailed information and removal steps for controller module, including processor, system memory and battery.

Storage

Disk Group Roaming:

PS7012 and PS312e2/390e/392e have the unique capability of allowing disk groups and their properties to roam with the drives. The user can physically remove the drives that make up a disk group and install in different locations within the same array or in a different array and the system will recognize the disk group configuration including the RAID volume.

LUN Masking and host information are stored on the disks of the roamed disk group and are maintained when the disk group is roamed into the new system.

Disk Group roaming is possible because when a disk group is created, the system automatically creates a small 1GB partition on each drive. This partition contains the disk group configuration. In this way, all the information regarding a disk group can be recognized by any PS7012 controller during the automatic boot up discovery or during a manual discovery using the PixelStor Storage Manager or CLI.

The drives can be reinserted in different slots and they can be rediscovered while the system is on-line. You can take advantage of this flexibility to arrange your storage for different on-line and off-line tasks. This feature provides a basic and inexpensive way to upgrade a system and retain your data or move drives to another system for basic system recovery.



All drives in a DG need to be present for the disk group discovery to take place. A disk group with one missing member can be forced online using the PixelStor Storage Manager or CLI discovery command.

Disk Orphans:

If some of the drives from a disk group are not re-inserted or a drive from the disk group goes bad and the disk group is not reconstructed before it roams to a new system, the new system will not be able to discover that disk group. It will know that one or more drives are missing and will designate the drives in the incomplete disk group as orphans. The data in the incomplete disk group is maintained, however user intervention is required to recover the data on the drives. This ensures that the data on the remaining drives is not overwritten. If the drives are designated as orphans, the user has the following options:

Complete the Set – You can track down the remaining drives from the disk group, insert them in any order into the same PS7012 /312e2/390e/392e enclosure and use the PixelStor Storage Manager to do a discovery. Once the whole set is discovered, the disk group will be restored.

Rebuild a Disk Group – If the original group is part of a RAID 1, 5, 6, or 10 disk group, a bad drive maybe the reason for an orphan classification by the system. If this is the case, make sure you have a global spare in the system and then use the PixelStor Storage Manager to do forced discovery of the disk group. Once the disk group is discovered and the bad drive is identified, the system will start the RAID rebuild using the global spare. This will restore the disk group.

Re-configure the Drives – If you cannot complete the disk group or rebuild it, then you will not be able to recreate the disk group (RASILIENT recommends you always back up your data to recover from such a scenario). Alternatively you may just want to reuse drives from an old disk group where the data on the drives is no longer needed. In either case you can create a new disk group or turn the drives into global spares by using the PixelStor Storage Manager.

Chapter 5: Online Disk Cloning

Proactive Disk Cloning

Disk cloning is a user-initiated, online, copy-and-replace of a problematic disk-group member with another disk. Cloning allows the user to proactively replace a disk while still maintaining the disk-group status as healthy. In the coming new release, the system will suggest cloning when a disk's average command completion time or its media errors, managed by the Rasiliant RAID Stripe Table, have increased to unacceptable levels.

Preparation

Spare check

- Please make sure there is at least ONE spare disk in your system.

Cloning procedure

Check RAID Stripe Table

SSH to PS7012 and login with root/root. And then run:

```
"daraidst-s -g <DG Name>"
```

```
[root@Storage-0 ~]# daraidst -s -g Raid5
```

```
RAID Stripes Summary:
```

Chunk Breakdown:

[00] xda	(Degraded/Restricted/Unhealthy =	0/	0/	0)
[01] xdb	(Degraded/Restricted/Unhealthy =	0/	1/	0)
[02] xdc	(Degraded/Restricted/Unhealthy =	6/	0/	7)
[03] xdd	(Degraded/Restricted/Unhealthy =	0/	1/	0)

```

[04] xde      (Degraded/Restricted/Unhealthy = 0/ 0/ 0)
[05] xdf      (Degraded/Restricted/Unhealthy = 0/ 0/ 0)
[06] xdg      (Degraded/Restricted/Unhealthy = 0/ 0/ 0)
[07] xdh      (Degraded/Restricted/Unhealthy = 0/ 0/ 0)
[08] xdi      (Degraded/Restricted/Unhealthy = 0/ 1/ 0)
[09] xdj      (Degraded/Restricted/Unhealthy = 0/ 2/ 0)
[10] xdk      (Degraded/Restricted/Unhealthy = 0/ 2/ 0)
    
```

RAID Stripes: 13 total

6 stripe(s) with 1	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
7 stripe(s) with 2	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 3	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 4	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 5	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 6	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 7	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 8	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 9	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 10	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 11	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)
0 stripe(s) with 12	Degraded/Restricted/Unhealthy/ Unavailable chunk(s)

RAID Stripe Table:

H-Healthy, D-Degraded, U-Unhealthy, R-Restricted, V-unavailable
 (Issue "daraidst -h" to read the detail of the keys.)

Sector:	xda	xdb	xdc	xdd	xde	xdf	xdg	xdh	xdi	xdj	xdk
191360 :	H	H	D	H	H	H	H	H	H	H	H
1809783168 :	H	H	D	H	H	H	H	H	H	H	H
1952775168 :	H	H	U	H	H	H	H	H	H	H	R
2147791744 :	H	H	D	H	H	H	H	H	H	H	H
2395665280 :	H	H	D	H	H	H	H	H	H	H	H
2733225984 :	H	H	U	H	H	H	H	H	H	H	R
2928417408 :	H	H	D	H	H	H	H	H	H	H	H
3123643648 :	H	H	U	H	H	H	H	H	H	R	H
3123645184 :	H	H	U	H	H	H	H	H	H	R	H
3123645312 :	H	H	U	H	H	H	H	H	R	H	H
3123649024 :	H	H	U	R	H	H	H	H	H	H	H
3514106624 :	H	R	U	H	H	H	H	H	H	H	H
3709362816 :	H	H	D	H	H	H	H	H	H	H	H

Fixing the bad stripes

Sum the bad stripes up

RAID stripe table shows 5 statuses: H – Healthy, D – Degraded, U – Unhealthy, R – Restricted and V – unavailable. Sum up the statuses of "D" and "U" of each hard drive. Find out the biggest number and locate that disk. That disk needs to do "cloning" at first. In this example, xdc disk has 13 bad stripes, so it is the first disk to be "cloned".

Find the DG in Disk Group Properties

Go to PS7012 web page and login with root/root. Click Disk Groups Tab and then click that DG name in Disk Group List in the left Navigation Tree. And then click Clone button shown in Figure 88: Find the DG in Disk Group Properties at the bottom.

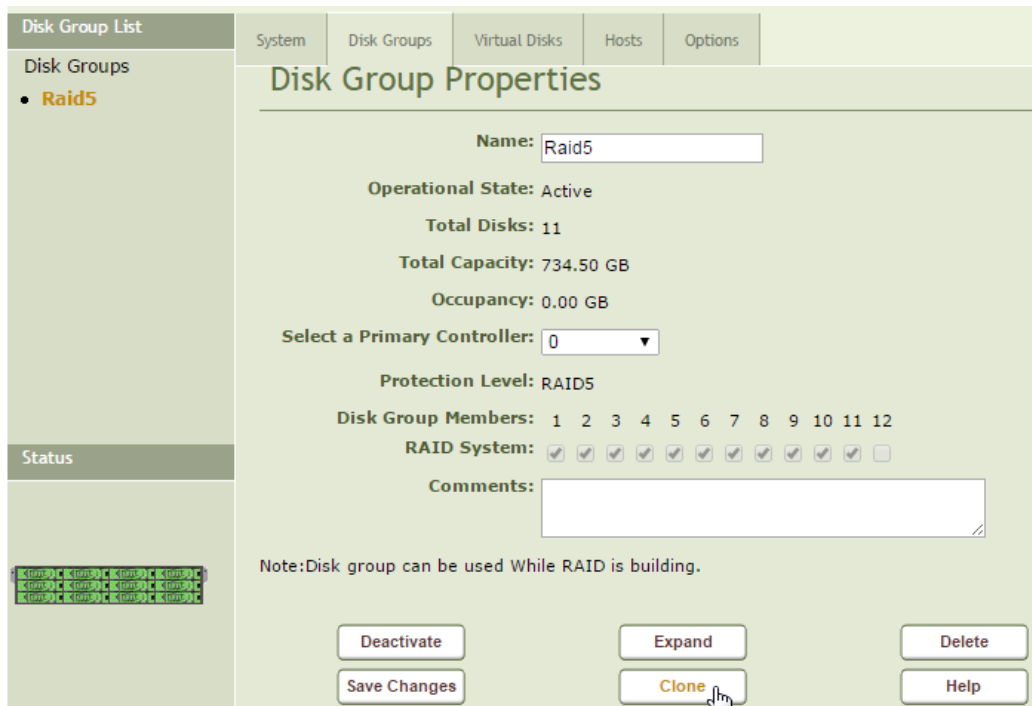


Figure 88: Find the DG in Disk Group Properties

Start cloning

The disks from xda to xdl are represented by number 1 to 12 in Disk Group Numbers and Spares. Check xdc disk-- number 3 and the spare disk-- number 12 and click Finish button shown in Figure 89: Start Cloning. In this example, No. 3 disk is the source disk and No. 12 disk is the target disk.




Figure 89: Start Cloning

Cloning status

The cloning is started automatically. The RAID Cloning button is highlighted while the RAID is cloning shown in Figure 90: RAID Cloning is ongoing. You may also click on the RAID Cloning button to open the RAID Cloning Progress page shown in Figure 90: RAID Cloning is ongoing.

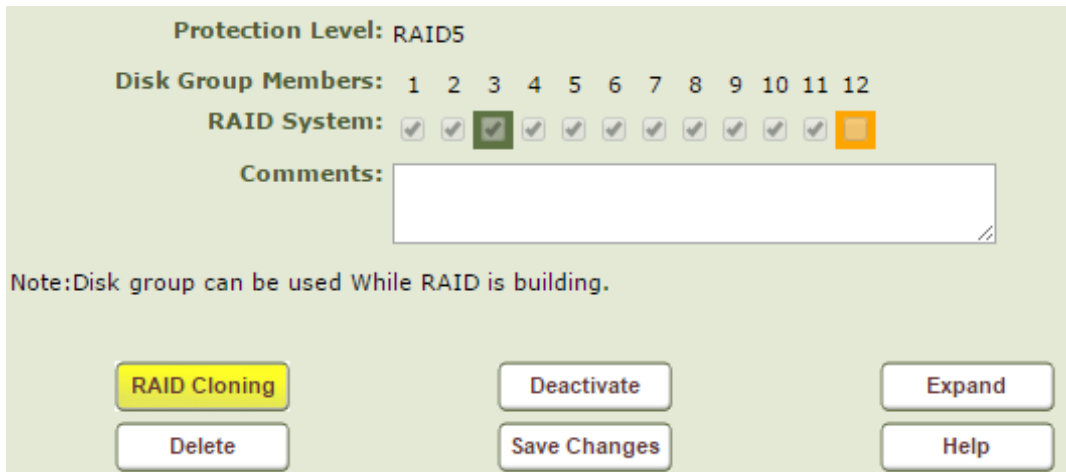


Figure 90: RAID Cloning is ongoing




Figure 91: Disk RAID Cloning in Progress Page

Complete cloning

When system completes cloning, the progress shows "Clone is done" shown in Figure 92: Complete cloning. Click OK button to return Disk Group Properties page.

Disk RAID Cloning in Progress				
Type	Percent	Time	Speed	Progress
		clone is done		

Figure 92: Complete cloning

The target disk –No. 12 disk has been added into the DG shown in Figure 93: New Disk Group. The source xdc disk –number 3 disk has been removed from the DG.

Disk Group Members:	1	2	3	4	5	6	7	8	9	10	11	12
RAID System:	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Figure 93: New Disk Group

The status of number 3 disk -- xdc disk is orphan shown in Figure 94: The status of the source disk is orpha.

Operational State: Initialized - Good
Total Disk Groups: 1
Total Grouped Disks: 11
Total Spare Disks: 0
Total Orphan Disks: 1
Total Not-Owner Disks: 0
Total Faulty Disks: 0
Total Unqualified Disks: 0
Total Hyper Cache: 0
Total Physical Disks: 12

Figure 94: The status of the source disk is orphan

Replace the bad hard drive

Replace No. 3 disk -- xdc disk with a good hard drive.

An alternative method to display RAID stripe table

We provide an alternative method to visually display the errors in RAID stripes table.

SSH to PS7012 and login with root/root. And then run the below command.

```
[root@Storage-0 ~]# echo 1 > /mnt/flash/etc/show_raid_stripe_display
```

Return to the web page and press "F5" button to refresh. RAID Stripe Display will appear in the left navigation tree shown in Figure 95: RAID Stripe Display shown in navigation tree.



Figure 95: RAID Stripe Display shown in navigation tree

Click RAID Stripe Display and the below stripe table is shown in Figure 96: RAID Stripe Display.

Disk Group: **Raid5** ▼

M:Main chassis **E#**:Expander#

LBA:Logical Block Address in 512-byte sectors **Chunk**:64KB(128 sectors) of disk data

Chassis	M	M	M	M	M	M	M	M	M	M	M
Disk Slot	0	1	11	3	4	5	6	7	8	9	10
Error Chunk Count (1203848 Chunks per disk)	3	8	2	5	1	10	0	1	2	6	6
Error Percentage (x10 ⁻⁶ %)	24.9	66.5	16.6	41.5	8.3	83.1	0.0	8.3	16.6	49.8	49.8
LBA											
2000128		■									
2007552								■			
2009344									■		
2095488				■							
2098304						■					
2242688											■
2253952		■									
2260096		■									
2276992				■							
2277888				■							
2280704										■	

Figure 96: RAID Stripe Display

Scroll to the bottom of the page, there are 3 buttons shown in Figure 97: Click Fix button and then click Clone button.

Click Fix button to fix the degraded stripes. After that, if there are still some errors in the system, click Clone button to start cloning. The system starts cloning automatically. It will choose the disk representing the biggest error percentage shown in Figure 97: Click Fix button and then click Clone button to clone.

3639040						■						■
3647104									■			
3854592							■					
3857280				■								
3857408											■	
3864064							■					

Figure 97: Click Fix button and then click Clone button

Q&A

Why I cannot find “Clone” button in Disk Group Properties page?

Answer: Please make sure that your system finishes RAID building. If your system is building RAID, Clone button replaces by Save Changes button shown in Figure 98: System is initially building.




Figure 98: System is initially building

Why does the system pop up the error No. 3196 after starting cloning?

Answer: If the error shown in Figure 99: Error information page appears, please check the capacity of the source disk and target disk. Make sure that the capacity of the source disk is equal or less than the target disk. For further investigation, please send system log to Rasilient support to analyze.

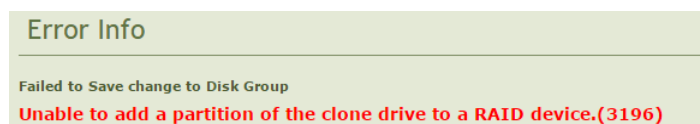


Figure 99: Error information page

If the source disk turns into a faulty disk during cloning, what can I do?

Answer: During cloning, the source disk – No. 3 disk shown in Figure 100: Source disk turns into a faulty disk during cloning turns into a faulty disk. The system will stop cloning and start rebuilding automatically. After the system completes rebuilding, you need to remove the faulty disk -- No. 3 disk shown in Figure 101: Target disk turns into a faulty disk during cloning and replace with a good hard drive. For further investigation please send system log to Rasilient support to analyze.

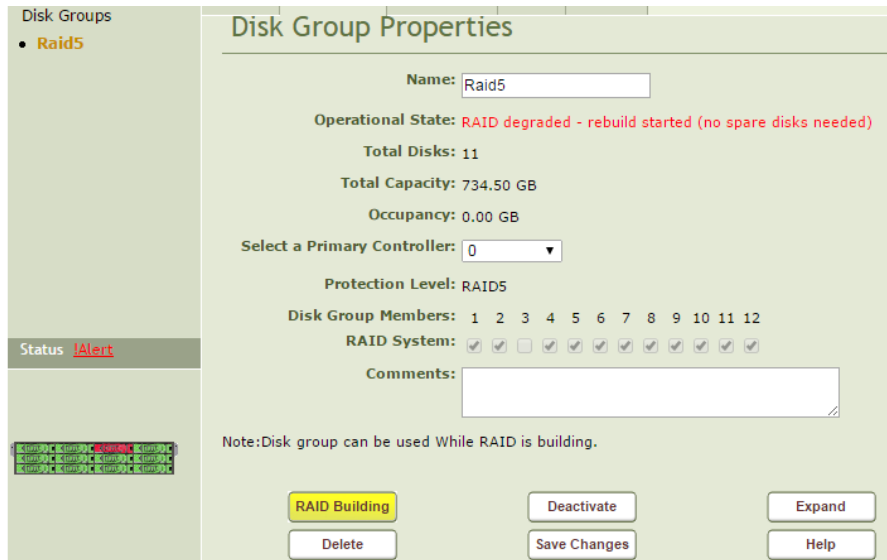


Figure 100: Source disk turns into a faulty disk during cloning

If the target disk turns into a faulty disk during cloning, what can I do?

Answer: During cloning, the target disk – No. 12 disk shown in Figure 101: Target disk turns into a faulty disk during cloning turns into a faulty disk. The system will stop cloning. You need to remove this faulty disk and replace with a good hard drive to restart cloning. For further investigation please send system log to Rasilient support to analyze.

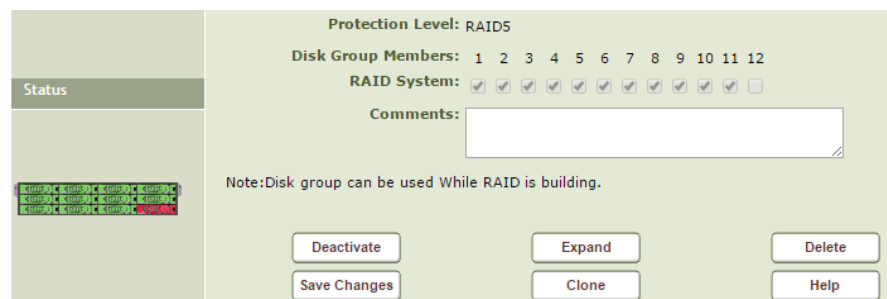


Figure 101: Target disk turns into a faulty disk during cloning

If another disk turns into a faulty disk during cloning, what can I do?

Answer: During cloning, another disk, for example, No. 1 disk shown in Figure 102: No. 1 disk turns into a faulty disk during cloning, turns into a faulty disk. The Disk Groups turns into degraded and cloning is still ongoing. If having another spare drive, the system starts rebuilding automatically. If not, you need to wait for

completing cloning firstly. And then remove the fault disk-- No. 1 disk shown in Figure 102: No. 1 disk turns into a faulty disk during cloning and replace with a good hard drive. The system will start rebuilding automatically. For further investigation please send system log to Rasilient support to analyze.

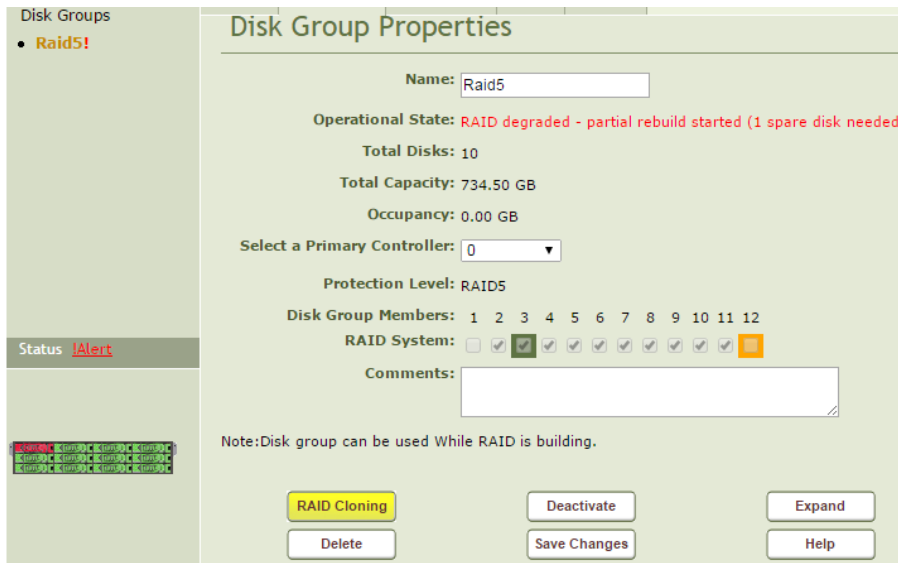


Figure 102: No. 1 disk turns into a faulty disk during cloning

Appendix A – RAID Concepts & Cache Configurations

RAID Levels Defined

RAID 0

Data Striping without Protection - This technique has striping but no redundancy of data. The data is broken down into blocks and each block is written to a separate disk.

- Pros - I/O performance is improved by spreading the I/O across many channels/drives.
- Cons- Not "true" RAID because it is NOT fault-tolerant. The failure of just one drive will result in all data in an array being lost.
- Ideal applications - those that require high bandwidth without need for data protection.

RAID 1

Disk Mirroring - Consists of at least two drives that duplicate the storage of data. There is no striping.

- Pros- Read performance is improved since either disk can be read at the same time. Write performance is the same as for single disk storage. It provides the best performance and the best fault-tolerance in a multi-user system. 100% redundancy of data means no rebuild is necessary in case of a disk failure, just a copy to the replacement disk. Also, under certain circumstances, RAID 1 can sustain multiple simultaneous drive failures.
- Cons - highest disk overhead of all RAID types (100%).
- Ideal applications - those that require very high availability.

RAID 5

Distributed Parity - RAID-5 stores parity information but not redundant data. The parity information can be used to reconstruct data. RAID-5 requires at least three disks for the array.

- Pros - This is the best combination of performance and cost.
- Cons - Disk failure has some impact on throughput.
- Ideal Applications - this is the most versatile RAID level and is applicable to a broad range of applications. It is best for multi-user systems in which performance is not critical or applications that require few write operations.

RAID 6

Block-level striping with *dual* distributed parity - Like RAID 5, RAID 6 stripes data and parity across an array of drives with the additional capability to calculate two sets of data parity for each volume of data. Requires at least 4 disks for the array.

- Pros - Extremely high data fault tolerance. Provides protection against double disk failures and failures while a single disk is rebuilding while other RAID levels can handle at most one fault.
- Cons - RAID 6 performance can be slightly worse than RAID 5 in terms of writes due to the added overhead of more parity calculations. RAID 6 is inefficient when used with a small number of drives due to need for 2 parity drives.
- Ideal Applications - Complex, higher capacity SATA arrays (12-24 drives) where higher availability is critical.

RAID 10

Combination of mirroring and striping - This type offers an array of stripes in which each stripe is a RAID-1 array of drives.

- Pros - This offers higher performance than RAID-1
- Cons - but at much higher cost.
- Ideal Applications - those requiring high performance and fault tolerance.

RAID 50

RAID 50, also called RAID 5+0, combines the straight block-level striping of RAID 0 with the distributed parity of RAID 5. As a RAID 0 array striped across RAID 5 elements, minimal RAID 50 configuration requires six drives. One drive from each of the RAID 5 sets could fail without loss of data; for example, a RAID 50 configuration including three RAID 5 sets can tolerate three maximum potential simultaneous drive failures (but only one per RAID 5 set). Because the reliability of the system depends on quick replacement of the bad drive so the array can rebuild, it is common to include hot spares that can immediately start rebuilding the array upon failure. However, this does not address the issue that the array is put under maximum strain reading every bit to rebuild the array at the time when it is most vulnerable.

RAID 50 improves upon the performance of RAID 5 particularly during writes, and provides better fault tolerance than a single RAID level does. This level is recommended for applications that require high fault tolerance, capacity and random access performance. As the number of drives in a RAID set increases, and the capacity of the drives increase, this impacts the fault-recovery time correspondingly as the interval for rebuilding the RAID set increases.

- Pros - Better write performance and more fault tolerant than RAID 5.
- Cons - Twice the parity drives overhead.
- Ideal Applications - Requiring high fault tolerance, capacity and random access.

RAID 60

RAID 60 (sometimes referred to as RAID 6+0) combines multiple RAID 6 sets (striping with dual parity) with RAID 0 (striping). Dual parity allows the failure of two drives in each RAID 6 array while striping increases capacity and performance without adding drives to each RAID 6 array.

Like RAID 50, a RAID 60 configuration can accommodate 8 or more drives, but should only be used with configurations of more than 16 drives. The usable capacity of RAID 60 is between 50%-88%, depending on the number of data drives in the

RAID set. RAID 60 is similar to RAID 50 but offers more redundancy, making it good for very large capacity servers, especially those that will not be backed up (i.e. video surveillance servers handling large numbers of cameras).

Pros:

- Can sustain two drive failures per RAID 6 array within the set, so it is very safe.
- Very large and reasonable value for money, considering this RAID level won't be used unless there are a large number of drives.

Cons:

- Requires a lot of drives.
- Slightly more expensive than RAID 50 due to losing more drives to parity calculations.

Cache Settings Defined

Write-back Cache Option

Data normally written to disk is first written into the cache. During idle machine cycles, the data is written from the cache into disk.

- Pros- Improves performance because a write to the cache is faster than to disk. Reduces the number of write operations to disk.
- Cons- With this performance improvement comes a risk that data may be lost if the system crashes. The write back cache configuration should only be selected if you have the battery-backed option installed.

Write Through Cache Option

Data is written to disk and also written into the cache.

- Pros – This method ensures complete data integrity and good performance for sequential access, which relies on less cache. If a subsequent read operations need that same data, read performance is improved, because the data is already in the cache.
- Cons- Write performance is not improved with this method.

Auto Cache Setting

This RASILIENT setting automatically turns on write back cache and enables the mirrored cache feature with two controllers installed. If the second controller is removed or fails, the system automatically switches to a write through cache configuration.

- Pros - Good cache performance and complete data protection.
- Cons - Less optimized performance for sequential access due to mirrored cache system overhead.

Appendix B – Using the Command Line Interface (CLI)

The following information is for experienced users that wish to manage the PS7012 with the built in Command Line Interface using:

- A terminal connected to the Serial port of the PS7012 or
- A SSH session to the PS7012

This section describes how to access the CLI and provides a list of commands that can be used for system configuration, array management and administrative functions.

Connecting a Terminal to Access the CLI

Connect the supplied Serial cable to the Serial management port on either controller. Connect the opposite end of the cable to a workstation or laptop with a terminal emulation program. A DB9 connector is provided for the workstation or laptop.

Use the following terminal emulation settings to access the CLI:

- Baud Rate – 115,200
- Data Bits – 8
- Parity None
- Stop Bits – 1
- Flow Control Off

Accessing the CLI – Serial Port or SSH Session

Once connected to the Serial port with the PS7012 is running, you will be prompted for the following information:

- Login – root
- Password – root

When connecting through a SSH session connect to the PS7012 with " <IP Address>". You will be prompted for the following information:

- Login – root
- Password – root

Enter SuperUser mode by executing the "su" command (pwd is "root"). CLI commands must be executed as a SuperUser.

dahelp Command

The "dahelp" command lists all available CLI commands for the PS7012 . Detailed information for each CLI command is listed on the screen by simply executing the specific command without parameters or by specifying the "?" parameter.

Local/Remote Option

Some commands have a local or remote option. This enables the local controller to invoke an action onto the remote controller. The function of the local or the remote operation is specified in the command description. Command option *local* means you are invoking an action on the local controller, and command option *remote* means you are invoking an action on the remote controller.

Some commands don't have local or remote options. Setting the configuration on the local controller with these commands will set the same configuration on the both controllers.

Index of CLI Commands

The following table lists the PS7012 commands available through the command line interface.

Table 12: CLI command list

command	description
daacl	used to manage access control list
daaddrbk	used to add email address to notification list
dacomment	used to add comments to configuration
daconnstat	used to display connection statistics
dadate	used to display or set time and date
dadg	used to manage Disk Groups
dadiskload	download firmware for Seagate hard drives
dadsktst	measure disk baseline IO performance
dafwdwnld	used to download expander firmware
dahelp	used to display information about commands
dahost	used to define and manage host
dahostname	allows user to set the name of the storage array
dahwstat	used to display hardware status
dainspread	used to manage Inspection Read
daisns	used to manage isns registration
dakem	used to display the event log
dalun	used to manage the lun masking table

dantpdate	used to configure an NTP server
dapasswd	used to change the user password
daphyinfo	get expander phy status information
dapm	configure the drive power management
dareboot	used to reboot the local or remote controller
daroute	used to display or set the gateway address
dasavelog	saves the log file to the /pub/ directory
daselblk	select a chassis or a hard drive to blink
dasetup	used to set IP addresses, netmask, gateway and name server
dashutdown	used to shutdown the local or remote controller
dasmartinfo	used to view disk smart information
dasnmp	used to manage snmp groups
dastat	used to display system information
dastonith	used to forcibly reboot the remote controller
dasysalert	used to enable or disable health alerts
daum	used to upgrade or manage kernel images
davd	used to manage virtual disks
daversion	used to display version information of kernel images or boot images

CLI Commands

The CLI Commands for the PS7012 have the prefix “da” as part of the command name. The commands and their descriptions are listed below.

daacl

Use daacl to manage the access control list on at a system level.

daacl - manage access control list.

Usage:

daacl [Options]

Options:

-h/--help: help

-L/--list: list all allowed addresses

-a [address/prefixlen]: add an address into the allowed list

-d [address/prefixlen]: delete an address from the allowed list

Command Examples:

daacl -a 0.0.0.0/0: allow access from all

daacl -a 192.168.1.0/24: allow access from one subnet
daacl -a 192.168.1.1/32:
allow access from one IP

daaddrbk

Use daaddrbk to manage the email addresses for notification.

daaddrbk - add email address to notification list.

Usage:

daaddrbk [-hL] [-ad <email_address>] Options:

- a [recipient_email] : add an recipient email to the address book
- d [recipient_email] : delete an recipient email from the address book
- o [sender_email] : set the sender email
- s [smtp_server_ip] : set the smtp server ip address
- S [1/0] : set whether the smtp server support SSL
- p [smtp_server_port]: set the smtp server port, default is 25
- u [user_name] : set the user name
- w [password] : set the password
- h/help: help
- L: list all the current email address

Command Examples:

daaddrbk -a steve@foo.com -o test@test.com -s 1.1.1.1 -p 25 add a notification email address.

daaddrbk -d steve@foo.com
delete a notification email address.

Note:

To receive email notification, you need to make sure the default gateway/router and name server are setup properly. The system needs to be able to talk to the mail server for the email address you entered.

dacomment

Use dacomment to add or display comments on a system.

dacomment- add comments to configuration.

Usage:

dacomment [Options]

-h: help

-g: get the comment

-s ["comments"]: set the comment

Examples:

dacomment -s "this is storage system for \"server A\"" dacomment -g

daconnstat

Use daconnstat to display the iSCSI connection statistics.

daconnstat - display connection statistics.

Usage:

daconnstat [Options]

Options:

?/-help: help

-show: show statistics

-reset sid: reset rx/tx counters of session "sid" to zero

Command Examples: daconnstat -help daconnstat daconnstat -show
daconnstat -reset 0xffff81004e409000

Notes:

The "daconnstat" without any option has the same effect as "daconnstat -show".

dadate

Use `dadate` to display or set the system date and time on one or two controllers. Without any option specified, the current date is printed out.

`dadate` - Use `dadate` to display or set the system date and time on one or two controllers

Usage:

`dadate` [Options]

Options:

`?/-help:` help information

`-s [MMDDhhmmYY.ss]:` set date and time specified

`-z [GMT|gmt|UTC|utc] offset:` set time zone by offset to GMT/UTC

Command Examples:

`dadate:` prints out the current date and time
`dadate -s 1216183005.32:` sets the date and time

`dadate -z GMT:` set as Greenwich Mean Time
`dadate -z GMT +8:` set as 8th zone east to GMT
`dadate -z GMT -8:` set as 8th zone west to GMT

Notes:

For offset value for your location, please check out the following sites

<http://www.greenwichmeantime.com>

<http://www.worldtimezone.com>

`dadg`

Use `dadg` to manage the Disk Groups.

`dadg` - manage disk groups.

Usage:

`dadg` [Options] [Parameter Options] [Device List]

Options:

- h: help
- help: help
- Ldisk: list disk information
- Ldg [disk-group name]: list disk-group information
- Lall: list generic information for all disk groups
- a [disk-group name]: create a disk group
- d [disk-group name]: delete a disk group
- D [disk-group name]: deactivate a disk group
- A [disk-group name]: activate a disk group
- s [disk-group name]: restore unhealthy disk group to degraded/healthy
- r: discover new disk groups
- e [disk-group name]: expand a disk group
- c [disk-group name]: change a disk-group parameter
- y [single device name]: recycle an orphan disk to a spare disk
- EC [disk-group name]: check and correct RAID consistency Parameter

Options:

- o [controller owner]: [0, 1] for controller 0, controller 1
- l [RAID level]: [0, 1, 5, 6, 10] for RAID-0,1,5,6,10
- v [alarm level]: [1-100] for disk-space alarm threshold
- n [new disk-group name]: new disk-group name
- m [comment string]: explanatory comment
- f: force deletion of disk group with virtual disks, force activation of degraded disk group, force restoration of unhealthy disk group, force discovery to find degraded disk groups, force name change of disk group

Option Requirements:

- option a requires: o, l, v, and device list. m is optional option d requires: f is optional
- option A requires: o. f is optional option s requires: f
- option r requires: o. f is optional option e requires: device list
- option c requires: at least one of v, n, or m. f is required with n option y requires: one device only

Device List Mapping:

/dev/xda: slot 0 (the leftmost slot of head chassis)

/dev/xdb - /dev/xdk: slot 1 - slot 10 of head chassis
 /dev/xdl: slot 11 (the rightmost slot of head chassis)
 /dev/xdaa - /dev/xdal: slot 0 - slot 11 of expansion enclosure 0
 /dev/xdba - /dev/xdbl: slot 0 - slot 11 of expansion enclosure 1
 /dev/xdca - /dev/xdcl: slot 0 - slot 11 of expansion enclosure 2 Device List

Requirements:

RAID-0 requires: 2 to 24 devices in device list RAID-5 requires: 3 to 24 devices in device list RAID-6 requires: 4 to 24 devices in device list
 RAID-1 requires: 2, 4, 6, ..., 24 devices in device list
 RAID-10 requires: 4, 6, 8, ..., 24 devices in device list Device List Examples:

/dev/xda /dev/xdb /dev/xdc /dev/xdg /dev/xdh /dev/xdo
 /dev/xd[a-c,g-h,o]
 /dev/xd[a,b,c,g-h,o]
 /dev/xdaa /dev/xdab /dev/xdac /dev/xdad
 /dev/xda[c-f]

/dev/xdb[a-d,i-l] /dev/xdc[a-d] Command Examples:

dadg -help dadg -Ldisk dadg -Ldg dg1 dadg -Lall
 dadg -a dg1 -o 0 -l 5 -v 100 -m "comment" /dev/xd[a-e] dadg -d dg1
 dadg -d dg1 -f dadg -D dg1
 dadg -A dg1 -o 1 dadg -A dg1 -o 1 -f dadg -s dg1 -f
 dadg -r -o 0 dadg -r -o 0 -f
 dadg -e dg1 /dev/xd[f-i]
 dadg -c dg1 -m "comment new" -v 90 dadg -c dg1 -n dg1new -f
 dadg -y /dev/xdo dadg -EC dg1

dadiskcache

Use dadiskcache to enable or disable the use of the write cache on the drives.

dadiskcache - turns on/off all the disk write caches.

Usage:

dadiskcache [-h | -e | -d | -s]

Options:

- h : Help
- e : Enable all the disk write caches.
- d : Disable all the disk write caches.
- s : Display the current disk write cache state.

dadiskload

Use dadiskload to download the f/w for Seagate drives.

dadiskload - download firmware for Seagate SATA hard drive

Usage:

dadiskload [Options]

Options:

- n [URL or file]: URL or file name for the firmware, the firmware file must be in the format: drivemodel_firmwareversion
- d [devices/all]: specify the hard drives for download, for example /dev/xd[a-h] or specify 'all' to download for all available hard drives
- i: display drive information
- s: only display Seagate drive information (-i required)
- h: display this help information
- u: skip firmware version check for download(in downgrade)
- q: hide download progress information on console

Required options: -d, -n

Command Examples

```
dadiskload -d /dev/xda /dev/xdi /dev/xdd[a-h] -n joe:passwd@192.168.1.1/  
images/ST3400633AS_3.AAD
```

```
dadiskload -d /dev/xda /dev/xdi /dev/xdd[a-h] -n joe:passwd@192.168.1.1/  
images/ST3400633AS_0.AAD -u
```

```
dadiskload -d all -n joe:passwd@192.168.1.1/images/ST3750640NS_3.AEK
```

```
dadiskload -i
```

`dadiskload -i -s`

`dadsktst`

Use `dadsktst` to measure the baseline data transfer performance for disk drives.

`dadsktst` - measure hard drive baseline performance

Usage:

`dadsktst [Options]`

Options:

- f [filename]: store test results in a file
- d [devices/all]: specify the hard drives to test, for example `xda` or specify 'all' to test all available hard drives
- t [r/w/rw]: specify 'r', 'w' or 'rw' for read only, write only or both tests write test can only be performed for 'spare' drives
- v: hide test information on console
- h: help information

Required options: `-d, -t`

Command Examples

`dadsktst -d xda xdi xdg -t rw dadsktst -d all -f results.dat -t w`

Disk List Mapping:

- `/dev/xda:` slot 0 (the top-left slot of RAID System (main chassis))
- `/dev/xdb - /dev/xdk:` slot 1 - slot 10 of RAID System (main chassis))
- `/dev/xdl:` slot 11 (the bottom-right slot of RAID System (main chassis))
- `/dev/xdaa - /dev/xdal:` slot 0 - slot 11 of expansion enclosure 0
- `/dev/xdba - /dev/xdbl:` slot 0 - slot 11 of expansion enclosure 1
- `/dev/xdca - /dev/xdcl:` slot 0 - slot 11 of expansion enclosure 2

dafwdwnld

Use dafwdwnld to download the f/w into the Expanders. This command downloads the f/w automatically into both controllers and all the Expanders.

dafwdwnld - download target system's main chassis and expander firmware

Usage:

dafwdwnld [Options]

Options:

-g/--gui: this tool is called from GUI

-h/--help: display help information

-i/--info: display information of the system's main chassis and expander firmware

-n [URL]/--new [URL]: specify the firmware file name and its location

-u download the firmware, even if it is older than the firmware on the main chassis or expander

-q hide download progress information

Option Requirements:

URL is in the following format: `username:passwd@ip_address/path_name/file_name`

The host specified with 'ip_address' must be a FTP site. If the FTP server accepts 'anonymous' user, the 'username:passwd' can be omitted.

Command Examples

```
dafwdwnld -n joe:passwd@192.168.1.1/images/vsc7513.fuf dafwdwnld -n
```

```
joe:passwd@192.168.1.1/images/vsc7513.fuf -u dafwdwnld -n
```

```
192.168.1.1/pub/vsc7514.fuf
```

```
dafwdwnld -I
```

dapm

Use dapm to display the drive power management state and to set the inactivity period for drive spin-down.

dapm - Use dapm to display drive power management state.
- Use dapm to set inactivity period for drive to spin down.

Usage:

dapm [Options]

Options:

-help: Display help information.

-s Display drive power management state.

-l [inactivityPeriod] Set inactivity period for drive to spin down.

Valid values: 0, 10, 20, 30, 60, 90, 120, 150,
180, 210, 240, 270, 300 or 330 minutes.

A value of 0 disables spin down.

-i Display current inactivity timeout value.

-z schedule time of day to spin up and spin down disks in a disk group.

-y stop current schedule in a disk group.

Parameter options:

-g [disk group name]: disk group name.

-n : spare or orphan drives.

-b [time of day]: time of day [0-23] to spin up the drives.

-e [time of day]: time of day [0-23] to spin down the drives.

Option requirement:

option l requires: g or n. option i requires: g,n optional. option y requires:
g.

option z requires: g,b,e.

Command Examples:

dapm -s Display power management state for all drives. dapm -l 30 -n
 Set inactivity period for spare/orphan drives
to 30 minutes. Drive has no activities for 30 minutes will spin down to enter standby
state.

dapm -l 30 -g dg1 Set inactivity period for disk group dg1 to 30
minutes. Drives in this disk group has no activities for 30 minutes will spin down and
enter standby states.

dapm -i Display current inactivity timeout value for all disk groups and
spare orphan drives.

dapm -i -n Display current inactivity timeout value for spare or orphan drives.

dapm -i -g dg1 Display current inactivity timeout value for
disk group dg1.

dapm -z -g dg1 -b 8 -e 20 Set disk group dg1 to spin up drives at 8AM and spin
down its drives at 8PM.

dapm -y -g dg1 Stop current scheduling for disk group dg1.

dahelp

Use dahelp to display a list of CLI commands

dahelp - display command usage page

Usage:

dahelp [Options]

Options:

-h/help: display this help page

[command]: display usage for the specified command

-s [pattern]: display usage for all the commands that match the pattern

Command List:

daacl - manage access control list

daaddrbk - add email address to notification list dacomment - add comments to
configuration daconnstat - display connection statistics

dadate - display or set the date and time dadg - manage disk groups

dadiskcache - Turn on/off disk write caches for all drives
dadiskload - download firmware for Seagate hard drives
dadsktst - measure disk baseline IO performance
dafwdwnld - download expander firmware
dahelp - display help information about commands
dahost - define or manage host
dahostname - display or set hostname for the storage array
dahwstat - display hardware status
dainspread - manage Inspection Read
daisns - manage iSNS client
dakem - display event log
dalun - manage the lun masking table
dantpdate - configure NTP server
dapasswd - set the password for the user
daphyinfo - get expander phy status information
dapm - configure the drive power management
dareboot - reboot local or remote controller
daroute - display or set gateway address
dasavelog - save the complete log file from the target to the local PC.
daselblk - select a chassis or a hard drive to blink
dasetup - set IP addresses, netmask, gateway and name server
dashutdown - shutdown local or remote controller
dasmartinfo - used to view disk smart information
dasnmp - Add/Delete SNMP managers to send traps to
dastat - display the system information
dastonith - force reboot of remote controller
dasysalert - enable or disable blinking lights
daum - upgrade or manage kernel images
david - manage virtual disks
daversion - display versions of kernel images and boot image

dahost

Use dahost to add or delete host entries from the system.

dahost - define or manage host. Usage:
dahost [Options] [Parameter Options]

Options:

-a: add a new entry
-m: modify an entry

-d: delete an entry
-l: print out all the entries

Parameter Options:

-b [type]: 0 is iSCSI, 1 is FC, 2 is SAS
-w [host's WWN]: World Wide Name
-u [user name]: empty by default
-o [host OS type]: 0 is Windows by default, 1 is Linux, 2 is Others
-n [host name]: host name
-i [host ip address]: empty by default
-p [initiator secret]: empty by default
-s [target secret]: empty by default
-t ["comments"]: empty by default

Option Requirements:

option a requires: w and n are mandatory. i,u,p,s and t is optional option d requires:
w or n

option m requires: w are mandatory, u,o,n,i,p,s and t is optional WWN (-w) can not be changed

Command Examples:

```
dahost -a -w 0x0210000E0802DE2f -b 1 -n server2
dahost -a -w iqn.test3 -o 0 -n test3 :          add a host
dahost -d -w iqn.test3 :          delete a host
dahost -d -n test3 :          delete a host
dahost -a -w iqn.test2 -o 0 -n server2 -t "Server 2": add a host with comments
dahost -m -w iqn.test2 -n newhost :          modify a hostname
dahost -m -w iqn.test2 -t "" :          clear host comments
dahost -m -w iqn.test2 -s "" :          clear target secret
```

dahostname

Use the `dahostname` to change the system name for the PS7012 . For those systems that build an iSCSI target name with this name, you must use the `-f` (force) flag to change the name because target names may be affected.

`dahostname` - display or set hostname for the storage array.

Usage:

`dahostname [-h | -f] [hostname]`

Options:

`-h/help`: display syntax

`-f`: force hostname change. This flag must be specified to change host name for system supporting multiple iSCSI targets.

Command Examples:

`dahostname`: Display current hostname `dahostname -f newhost`: Change hostname to newhost

`dahwstat`

Use `dahwstat` to display the hardware status for the local controller.

`dahwstat` - show hardware status.

Usage:

`dahwstat [Options]`

Options:

`-help`: display help information

`-drives | -drive`: show properties of all the drives

`-drive [no.]`: show properties of a specified drive [0-23]

`-chassis`: show information of chassis components, including controller ID, second controller installed,

Second controller presence, fan readings, temperature readings, voltage readings and battery information

-alert show alert messages if any

Command Examples: dahwstat -alert dahwstat -drive 0 dahwstat -drives dahwstat - chassis

dainspread

Use dainspread to manage the Inspection Read operation.

dainspread - manage Inspection Read.

Usage:

dainspread [Options] [Option Parameters]

Options:

-h/-help: Show help information

-start: Immediately start Inspection Read

-stop: Immediately stop Inspection Read

-resume: Immediately start Inspection Read from last inspected block

-schedule: Add, delete or show Inspection Read schedule. When no parameters are provided, -schedule shows the current schedule

-progress: Show the progress of Inspection Read

Option Parameters:

-repeat: Run continuously until dainspread is issued with the stop option or the specified -duration is satisfied.

-cpu [priority]: Set the Inspection Read priority based upon cpu load [1-3]. Where High is 1, Medium is 2, Low is 3.

The default priority is low unless otherwise specified.

- duration [hours]: The maximum of hours to run [1 - 10000]. Inspection Read will stop after running for the specified number of hours.
- add: Add a new schedule or change the current schedule.
- del: Delete the current schedule.
- sch_resume: Schedule Inspection Read to resume from last inspected block.
- day [Day]: Used to specify the day of week to start Inspection Read [0 - 6], where 0 specifies Sunday.
- hour [hour]: Specifies the hour to start Inspection Read [0 - 23], where 0 specifies midnight.

Option Requirements:

start: repeat, cpu and duration are optional.

stop: no parameters should be specified with stop. resume: repeat, cpu and duration are optional.

schedule: add or del required. If add, then day and hour are required.

sch_resume, repeat, cpu and duration are optional. If del, then no other parameters should be specified.

progress: no parameters should be specified with progress.

Examples:

dainspread -start Start Inspection Read.

dainspread -start -repeat

Start Inspection Read and run continuously until dainspread is issued with the -stop option.

dainspread -start -repeat -cpu 2 -duration 6

Start Inspection Read at a medium priority and repeat until 6 hours have passed.

dainspread -stop Stop Inspection Read

dainspread -resume

Resume Inspection Read from the last inspected block.

dainspread -resume -cpu 1

Resume Inspection Read from the last inspected block at a high priority.

dainspread -resume -repeat -cpu 2 -duration 6

Resume Inspection Read at a medium priority and repeat until Inspection Read has run for 6 hours.

dainspread -schedule

Show the current Inspection Read schedule.

```
dainspread -schedule -add -day 0 -day 1 -hour 23
```

Schedule Inspection Read to start every Sunday and Monday at 11 PM.

```
dainspread -schedule -add -day 0 -hour 23 -duration 6
```

Schedule Inspection Read to start every Sunday at 11 PM and run for 6 hours.

```
dainspread -schedule -add -day 0 -hour 23 -duration 6 -sch_resume
```

Schedule Inspection Read to resume from the last inspected block every Sunday at 11 PM and run for 6 hours.

```
dainspread -schedule -add -day 0 -hour 23 -repeat -cpu 2 -duration 6
```

Schedule Inspection Read to start every Sunday at 11 PM at a medium priority and repeat until Inspection Read has run for 6 hours.

```
dainspread -schedule -del
```

Delete the current Inspection Read schedule.

```
dainspread -progress
```

Show the progress of the current Inspection Read.

daisns

Use daisns to manage isns registration.

daisns - manage iSNS client service

Usage:

```
daisns [Options] [Option Parameters]
```

Options:

-h/-help: help page

-l/-list: list iSNS server IP and alias name

-start: start iSNS service

-stop: stop iSNS service

-setserver: save iSNS server ip and alias name

-delserver: delete iSNS server ip and alias name

-display: display the iSNS server online or offline
-show: display whether the iSNS service is started

Parameters:

-addr: iSNS server IP address
-alias: iSNS server alias name

Option Requirements:

option setserver: addr and alias are mandatory.

Comand Examples: daisns -start daisns -stop daisns -l daisns -display
daisns -setserver -alias test3 -addr 10.4.1.3 daisns -delservice

dakem

Use dakem to display the event log.

dakem - display event log.

Usage:

dakem [Options] [Option parameters]

Options:[-hvEi] [-s severity] [-C n]

-h : help
-v : version
-E : display KEM messages
-i : suppress the special char display
-T : set the alert level for snmptrap/email notification. valid option
parameters are info,warning,critical
-G : get the alert level for snmptrap/email notification

Option Parameters:

-s [severity]: [info][warning][critical][all(default)]
-C [n]: number of messages to be displayed, maximum of 6000

Command Examples

dakem -E: Display the last 256 KEM messages.
dakem -E -C 6000: Display all the KEM messages.
dakem -E -C 6000 -s warning: Display all the KEM message of severity warning and above.
dakem -T warning Set alert level to warning.

dalun

Use dalun to manage the lun masking table. This pertains to Virtual Drive presentation.

dalun - manage the lun masking table.

Usage:

dalun [options] [fields...]

Options:

-h: help
-a: add a new lunmask entry
-d: delete a lunmask entry. You should log off at the host initiator before performing a delete.
This ensures the host has sent all data to the storage array.
-p: print out all the existing lunmask entries

Required fields:

-n [host]
-g [disk_group_name]
-k [virtual_disk_name]
-P [parent_virtual_disk_name]

Notes:

The -P parent_virtual_disk_name is only required if the virtual disk presented is a snapshot virtual disk.

Command Examples:

dalun -a -n host1 -g dg1 -k vd1:
presents virtual disk vd1 in disk group dg1 to host1
dalun -d -n host1 -g dg1 -k vd1:
un-present virtual disk vd1 in disk group dg1 from host1.

dalun -d -n host2 -g dg1 -k snap1 -P vd1:
un-present snapshot snap1 of virtual disk vd1 in disk group dg1 from host2

dantpdate

Use dantpdate to configure an NTP server.

dantpdate - Configure NTP servers

Usage:

dantpdate [Options]

Options:

-help help
-list list current NTP servers
-add hostName/IPAddress add a NTP server to list

-del hostName/IPAddress delete a NTP server from list
-sync synchronize the system clock on both controllers with the NTP server

Command Examples:

dantpdate -help display this help page
dantpdate -list list current list of NTP servers
dantpdate -sync synchronize the system clock on both controllers with the NTP server
dantpdate -add myntpserver add myntpserver to the list of NTP servers
dantpdate -del 65.78.1.1 delete 65.78.1.1 from the list of NTP servers

dapasswd

Use dapasswd to change the passwords on the system.

dapasswd - set the password for the user. Usage:
dapasswd [Options] [Option Parameters]

Options:

-help: display help page
-p: new passwd

Option Parameters:

username: Specify username. It is mandatory.

Command Examples:

dapasswd -p passwd John: sets password for a user John

daphyinfo

Use `daphyinfo` to get Expander phy status information for the local controller.

`daphyinfo` - get expander phy information

Usage:

`daphyinfo` [Options]

Options:

`-s:` get expander phy status information

`-l:` get error log information for phy

`-c:` clear error log information for phy

`-e:` specify the expander to get the information: `rs` -- RAID System (main chassis)

`eeN` -- expander enclosure, N is the enclosure number starting from 0 (default is all expanders)

`-p:` specify the phy number to get the information (default is all phys)

`-h:` display this help information

Required options: `-s` or `-l`

Command Examples

`daphyinfo -s -e rs -p 2 daphyinfo -l -e ee3 daphyinfo -s`

`dareboot`

Use `dareboot` to reboot the local or remote controller. If no parameters are specified, `dareboot` will reboot the local controller.

`dareboot` - reboot local or remote controller.

Usage:

dareboot
dareboot [Options]

Options:

-help: help
-local: reboot local controller
-remote: reboot remote controller

Command Examples: dareboot -help dareboot dareboot -local dareboot -remote

Notes:

The "dareboot" without any option has the same effect as "dareboot -local".

dashutdown

Use dashutdown to shutdown the local or remote controller. If local or remote are not specified, dashutdown operates on the local controller.

dashutdown - shutdown local or remote controller.

Usage:

dashutdown [Options] [Option parameters]

Options:

local: shutdown local controller.

Default is local controller

remote: shutdown remote controller

Option Parameters:

-t [seconds]: execute at now plus seconds specified now: shutdown
now

Notes:

dashutdown requires a time be specified, either "now" or -t <seconds>.

dasmartinfo

Use dasmartinfo to save the smart data from all of the non-faulty disks to the file /tmp/smartinfo.log. This command requires no options.

Usage:

dasmartinfo - used to view disk smart information

daroute

Use daroute to display or set the gateway IP address.

daroute - display or set the gateway address.

Usage:

daroute [Options] [{add | del} default gw <gwIPAddr>]

Options:

-h: help

Option Parameters:

add: add a default gateway

del: delete default gateway

<gwIPAddr>: gateway IP address to add or delete Command Examples:

daroute add default gw 192.168.123.254

dasavelog

Use dasavelog to save the controller logs to the /pub/ directory. managers

dasavelog - Save complete logs from the target to the local PC.

Usage:

dasavelog [Options]

Options:

-s: Creat and save the logyymmddhmmss.zip file to /pub/
User can use ftp from a client machine to retrieve this file. Please log in as an anonymous user, and get this file directly.
-h/-help: help page

daselblk

Use daselblk to select a chassis or hard drive to blink its LED(s). The selected item will flash its Red LED(s). If a single drive is selected, only one LED will flash. If a chassis is selected, all 12 Red LEDs will flash. Select the "-i" option to list the names of the selectable items.

daselblk - select a hardware item to blink

Usage:

daselblk [Options]

Options:

-s [item]: specify the hardware item to blink
-o [blink/cancel]: blink or cancel the blink
-i: list all available hardware items, which can be selected
-f: find the hardware item selected to blink
-h: help information

Required options: -s, -o

Command Examples

To blink a hard drive in slot 4 on RAID System (main chassis): `daselblk -s hdm4 -o blink`

To blink the Expansion Enclosure-0:

`daselblk -s ee0 -o blink`

dasetup

Use `dasetup` to assign the IP addresses, netmask, gateway and dns servers for the PixelStor system. This command may automatically start when you log into the Serial Port for the first time. Invoke this command whenever a NIC port IP address needs to be changed. The output of the `dasetup` command is shown below. Refer to the PS7012 Quick Start Guide for additional setup information.

`dasetup V2.7` is running for a dual controller system.

This wizard will walk you through the network configuration. You can configure the rest of the system using the web GUI.

The default values are in the bracket, you can press "Enter" to select the default value. If you make any mistake, you can press "ctrl-c", then logout and login to run this wizard again. Items with optional values are:

dedicated management port: [yes,no] mtu: [1500, 9000]

NIC IP: [a.b.c.d, dhcp, disable]

Configure NIC1 as dedicated management port [no]: Enter IP address for system [192.168.0.1]:

Enter netmask [255.255.255.0]:

Enter mtu [1500, 9000]: Enter C0 NIC1 IP []: Enter C0 NIC2 IP []: Enter C0 NIC3 IP []:
Enter C0 NIC4 IP []:
Enter C0 NIC5 IP []:
Enter C1 NIC1 IP []:
Enter C1 NIC2 IP []:
Enter C1 NIC3 IP []:
Enter C1 NIC4 IP []:
Enter C1 NIC5 IP []:
Enter default gateway IP []:
Enter DNS address for system []:

dashshutdown

Use dashshutdown to shutdown the local or remote controller.

dashshutdown - shutdown local or remote controller.

usage:

dashshutdown [Options] time

Options:

- help: help
- local: shutdown local controller. This is by default.
- remote: shutdown remote controller

The time argument is mandatory. It specifies the number of minutes to wait before shutdown happens. "now" indicates immediate shutdown.

Command Examples:

dashshutdown -help: display usage page

dashutdown now: shutdown the local controller now dashutdown -remote 10:
shutdown the remote one after 10 minutes

dasnmp

Use dasnmp list, add or delete SNMP managers

dasnmp - Add/Delete SNMP managers to send traps to. Usage:
dasnmp [Options]

Options:

- h/--help: help
- L/--list: list all SNMP managers
- a/--add [address]: add a SNMP manager with/without community
- d/--delete [address]: delete a SNMP manager
- c [community string]: community string of a SNMP manager
- i/--inform [0, 1]: set to informed(1) or non-informed(0) SNMP trap

Option Requirements:

option c requires: a

Command Examples:

dasnmp -a 192.168.1.1: add a SNMP manager without community
dasnmp -a 10.0.0.1 -c pub: add a SNMP manager with community "pub"
dasnmp -d 10.0.0.1: delete a SNMP manager
dasnmp -i 1: set to informed SNMP trap

dastat

Use dastat to display status information for the local controller.

dastat - display the system information.

Usage:

dastat dastat -help

dastonith

Use dastonith to force a reboot on the remote controller. Make sure a "dastonith 0" command is executed soon after the "dastonith 1" command, otherwise the remote controller will not reboot.

dastonith - force reboot of remote controller

Usage:

dastonith [-h/help] [0/1] Options:

-h/help: print this help page.

1: set the stonith condition to the remote controller.

0: release the stonith condition to the remote controller.

Command Examples: dastonith 1

dastonith 0

Notes:

To stonith the remote controller, issue the dastonith commands in the following sequence.

dastonith 1

dastonith 0

The remote controller will reboot after you issue dastonith 0.

dasysalert

Use `dasysalert` to display system Health Alert setting or to change setting - enable or disable. There are two components in the system Health Alert mechanism. Red LEDs on the hard drives blink 3 times when there is a system Health Alert in the specific chassis. Disabling system Health Alerts only turns off the blinking red LEDs and the red Alert! Indication in the GUI. This command does not affect the health messages viewable on the system health page or in the system logs.

`dasysalert` - turn on/off system alert

Usage:

`dasysalert [-h] [on | off]`

Options:

`-h`: help and print this page

`-s/-show`: the current status, 'on' or 'off'

`On`: turn on system alert

`Off`: turn off system alert

Command Examples: `dasysalert on dasysalert off dasysalert -show`

daum

Use `daum` to upgrade or manage kernel images.

`daum` - upgrade or manage kernel images.

Usage:

daum [Options]

Options:

- h/--help: help information
- i/--info: display information about the system images,
- s/--switch: switch the image pointer for next boot operation
- y/--sync: synchronize current and backup images
- n URL/--new URL: upgrade with a new system image
- v show version information of daum

Option Requirements:

URL is in the following format: `username:passwd@ip_address/path_name/file_name`

The host specified with 'ip_address' must be an FTP site. If the FTP server accepts an 'anonymous' user, the 'username:passwd' can be omitted.

Command Examples

```
daum -n joe:passwd@192.168.1.1/images/newkdi daum -n  
192.168.1.1/pub/kdi.0.1
```

davd

Use davd to manage the Virtual Disks of each Disk Group.

davd - manage virtual disks.

Usage:

davd [Options] [Parameter Options]

Options:

- h: help

-help: help
 -L [virtual-disk name]: list virtual-disk information
 -a [virtual-disk name]: create a data virtual disk
 -d [virtual-disk name]: delete a virtual disk
 -e [virtual-disk name]: expand a virtual disk
 -c [virtual-disk name]: change a virtual-disk parameter Parameter Options:
 -g [disk-group name]: disk group containing the virtual disk
 -s [size]: size in MB. Append G or g, T or t for GB, TB
 -m [cache policy]: [0, 1, 2] for auto, write-through, write-back
 -r [read ahead policy]: [0, 1] for enable, disable
 -n [virtual-disk name]: new virtual-disk name
 -f: force deletion of virtual disk force name change of virtual disk
 -io_stat: extra information on IO stat
 -io_stat_perf: extra information on IO performance(default interval is 5 seconds)
 -interval [seconds]: interval seconds with option io_stat_perf

Option Requirements:

option L requires: g

option a requires: g, s, m, r option d requires: g, f

option e requires: g, s

option c requires: g, one of (n,f), m, r Option Notes:

option e: -s [size] specifies the capacity increment

option c: Name changes (-n option) require the -f option,

Command Examples: `davd -help`

`davd -L lv1 -g dg1`

`davd -L lv1 -g dg1 -io_stat davd -L lv1 -g dg1 -io_perf`

`davd -L lv1 -g dg1 -io_perf -interval 5 davd -a lv1 -g dg1 -s 10240 -m 2 -r 0 davd -a`

`lv1 -g dg1 -s 10G -m 2 -r 1 davd -d lv1 -g dg1 -f`

`davd -e lv1 -g dg1 -s 8192 davd -e lv1 -g dg1 -s 8G davd -c lv1 -g dg1 -m 1 davd -c`

`lv1 -g dg1 -r 1`

`davd -c lv1 -g dg1 -n lv1new -f`

daversion

Use `daversion` to display version information of the kernel images (KDI) and other components of the PS7012 system.

`daversion` - displays the boot image and kdi versions.

Usage:

`daversion [Options]`

Options

`-h/--help`: help information

`local`: print out version information of local controller `remote`: print out version information of remote controller

Command Examples: `daversion -help` `daversion` `daversion local` `daversion remote`