VTrak E5000 Series

16Gb Fibre Channel to 12 Gb SAS/SATA

High-Availability Storage System

Product Manual

Version 1.0

© 2017 PROMISE Technology, Inc. All Rights Reserved.

About This Manual

This Product Manual describes how to setup, use, and maintain the VTrak E5300f, VTrak E5320f, Vtrak E5600f, and VTrak E5800f external disk array subsystems.

This manual includes a full table of contents, index, chapter task lists and numerous cross-references to help you find the specific information you are looking for.

The terms “VTrak E5000” or “subsystem” are used in examples or descriptions throughout this manual to refer to any of the available VTrak E5000 Series models. The terms “unit” or “device” can refer to any VTrak E5000 Series or VTrak J5000 Series model.

This manual includes a full table of contents, chapter task lists and cross-references to help you find the specific information.

Manual Overview

The manual is organized into chapters as follows:

- "Introduction" on page 1, this chapter provides a general overview of the available devices in the VTrak E5000 Series.
- "Hardware Installation" on page 16 describes the steps necessary for installing subsystem hardware including installing hard disks and placing the device into a rack system.
- "System Setup" on page 51 describes the basics of how to use the Setup Wizard with the embedded Web-based Promise Array Management – Professional (WebPAM PROe) software to setup one or more RAID arrays. It also describes setting up a serial connection and using the the built-in command-line interface (CLI) for basic settings such as changing IP settings and creating RAID arrays.
- "WebPAM PROe - System Configuration" on page 77 provides a more detailed description of the various menus used for managing the VTrak E5000 Series and connected VTrak J5000 Series expansion devices.
- "Managing with the CLI" on page 225 describes using the CLI to manage the VTrak E5000 Series through the network or via serial connection.
- "Maintenance" on page 339 describes how to replace hardware components including RAID controllers, power supplies, and cooling units; how to update firmware for subsystems and physical drives.
- "Technology Background" on page 355 provides a description of the technologies and concepts that underlie networked RAID storage systems generally and the VTrak E5000 Series subsystems in particular.
- "Troubleshooting" on page 397 describes what to do in response to specific problems that might be encountered over the lifetime operation of the VTrak E5000 Series subsystems. Included in the chapter are descriptions of the various types of alerts and notices delivered though the management interfaces (WebPAM PROe, CLI) or hardware (LEDs and audible signals).
- "Contacting Technical Support" on page 460 includes how to contact technical support, how to return a system for repair, and warranty information.

Also included are four levels of notices:

Warning

A Warning notifies you of probable equipment damage or loss of data, or the possibility of physical injury, and how to avoid them.

Caution

A Caution informs you of possible equipment damage or loss of data and how to avoid them.

Important

An Important message calls attention to an essential step or point required to complete a task, including things often missed.

Note

A Note provides helpful information such as hints or alternative ways of doing a task.

INTRODUCTION

This chapter covers the following topics:

- **"VTrak Overview"**
- **"Features"**
- "VTrak E5000 Series Models"
- "VTrak J5000 Series Models"
- **"Specifications"**
- **"Hardware"**
- "Front of VTrak E5800f and E5600f"
- "Front of VTrak E5300f and E5320f"
- "Back of VTrak E5800f and E5600f"
- "Back of VTrak E5300f and E5320f"

VTrak Overview

The PROMISE VTrak E5000 Series supports 12Gb/s and 6 Gb/s SAS and 6 Gb/s and SATA disks and embedded advanced storage I/O processing platform while providing a reliable, flexible and easy to manage RAID storage system.

The VTrak E5000 Series controller features four high speed 16 Gb/s Fibre Channel host ports per controller.

Features

Controller Module Features

Drive Support

- 3.5-inch and 2.5-inch form factor
- Hard disk drives (HDDs) and Solid State drives (SSDs)
- SAS, 6 Gb/s and 12 Gb/s
- SATA, 6 Gb/s
- Supports any mix of SAS and SATA drives simultaneously in the same enclosure

For a list of supported drives, go to PROMISE support: http://www.promise.com/support/

SATA drives require a SAS-to-SATA adapter, available from PROMISE.

External I/O Ports per Controller

- Four 16 Gb/s Fibre Channel ports, compatible with 8 Gb/s and 4 Gb/s.

  (Each Fibre Channel port support 16 Gb/s or 8 Gb/s Fibre Channel SFP only.)
- Two 12G SFF-8644 mini-SAS connectors per controller for JBOD expansion

Data Cache

- 8 GB data cache per controller located in controller.

  A portion of the data cache is shared with the firmware.
- Protected with battery backup unit (BBU) and advanced copy to Flash scheme.

Operational Features

RAID Level Support

- 0, 1, 5, 6, 10, 50, and 60

RAID Stripe Size Support

- 64K, 128K, 256K, 512K, and 1MB

Hot Spare Drives

- Global
- Dedicated
- Revertible option

Maximum LUNs Supported

- LUNs: 1024
- Array: 32

Advanced Storage Features

- Advanced Cache Mirroring over PCIe Gen3
- Simple, drag-and-drop LUN Masking and Mapping
- Asymmetric LUN Unit Access (ALUA)
- Volume Copy
- PerfectFlash - Non-Disruptive Software Update
- I/O performance & power monitoring tools
- Guaranteed Latency Technology (an advanced OEM feature)
- USB Service Log
- LDAP Support for central user management

Background Activities

- Media Patrol
- Background Synchronization
- Foreground Initialization
- Rebuild
- Redundancy Check
- Disk SMART Polling
- Online Capacity Expansion (OCE)
- RAID Level Migration (RLM)
- UPS Monitoring
- SSD Health Monitoring
- Feature rich task scheduler for background activities

PerfectRAID Features

- Predictive Data Migration (PDM)
- Intelligent Bad Sector Remapping
- SMART Error Handling
- NVRAM Error Logging
- Disk Slot Power Control
- Read/Write Check Table
- Write Hole Table

GreenRAID Features

- Four levels of advanced power management disk drive (MAID) support
- Efficient 80 Plus Gold Certified power supplies

VTrak E5000 Series Models

Model	Controller Units	Interface	Number of Drives	Power Supplies
E5800fD	2	Fibre Channel (8 ports up to 16 Gb/s)	24	2
E5600fD	2	Fibre Channel (8 ports up to 16 Gb/s)	16	2
E5300fD	2	Fibre Channel (8 ports up to 16 Gb/s)	12	2
E5320fD	2	Fibre Channel (8 ports up to 16 Gb/s)	24 (2.5”)	2
E5800fS	1	Fibre Channel (4 ports up to 16 Gb/s)	24	2
E5600fS	1	Fibre Channel (4 ports up to 16 Gb/s)	16	2
E5300fS	1	Fibre Channel (4 ports up to 16 Gb/s)	12	2
E5320fS	1	Fibre Channel (4 ports up to 16 Gb/s)	24 (2.5”)	2

VTrak J5000 Series Models

Model	Controller Units	Interface	Number of Drives	Power Supplies
J5800	2	SFF-8644 SAS	24	2
J5600	2	SFF-8644 SAS	16	2
J5300	2	SFF-8644 SAS	12	2
J5320	2	SFF-8644 SAS	24 (2.5”)	2

Note: VTrak J5000 Series also available with single controller.

System Management

Management Interfaces

- Browser-based management with WebPAM PROe over Ethernet
- Command Line Interface (CLI) over Serial Port, Ethernet via Telnet, or SSH
- Third Party Management Support via SNMP and CIM

Supported Browsers

Browsers run on the host PC or server, from which you monitor and manage the VTrak E5000 subsystem using WebPAM PROe. The browsers listed here meet the minimum version requirements for browser compatibility:

- Mozilla Firefox 14.0.1
- Google Chrome 20.0.1132.57 m
- Internet Explorer 7 (Version: 7.0.5730.13)
- Internet Explorer 8 (Version: 8.0.6001.18702)
- Internet Explorer 9 (Version:)
- Safari 5.1.7 for Windows
- Safari 5.1.7 for MAC
- Mozilla Firefox for Linux 3.6.13

For the latest list of supported browsers, go to PROMISE support:

http://www.promise.com/support/

Specifications

Form factor	2U, 3U, and 4U 19” rack mount
Drives supported	12 Gb/s or 6 Gb/s SAS, 6 Gb/s SATA HDD and SSD Supports any mix of SAS and SATA drives simultaneously in the same enclosure.
I/O Ports per E5000f controller	Four 16 Gb/s Fibre Channel ports (compatible with 8 and 4 Gb/s link speeds) Two 12G SFF-8644 mini-SAS connectors per controller for JBOD expansion.
Data Cache per controller	8 GB data cache per controller A portion of the data cache is shared with the controller firmware Protected with battery backup units (BBU).
Storage Expansion	Cascade up to nine VTrak J5000 JBOD expansion units. VTrak J5000 Series supports 12, 16 or 24 drives per device. *See VTrak J5000 Series Product Manual for details.*
Operational	
RAID support	0, 1, 5, 6, 10, 50, 60
RAID stripe size	64K, 128K, 256K, 512K, 1MB
Hot Spare Drives	Global, Dedicated and Revertible option
Maximum LUNs	1024 per system / 32 per array

General	Description			
Power Supplies	Efficient 80 PLUS GOLD certified redundant PSU			
Current (Maximum)	9A @ 100 Vac; 4A @ 240 Vac			
Power Conversion Efficiency	>80% @ 110V (>20% load), >80% @ 240V (>20% load) Temperature Range			
	E5320f	**E5300f**	**E5600f**	**E5800f**
Dimensions (Height, Width, Depth)	88 x 46.7 x 420 mm 3.5 x 17.6 x 16.5 in	88 x 446.7 x 507 mm 3.5 x 17.6 x 19.96 in	131 x 446.7 x 507 mm 5.2 x 17.6 x 19.96 in	174.4 x 446.7 x 507 mm 6.87 x 17.6 x 19.96 in
Weight (w/o drives) (w drives)	 16.2 Kg (35.7 lbs) 20.5 Kg (45.2 lbs)	 18.9 Kg (41.7 lbs) 26.8 Kg (59.1 lbs)	 22.3 Kg (50.7 lbs) 33.4 Kg (73.6 lbs)	 24.8 Kg (54.7 lbs) 40.5 Kg (89.2 lbs)

Safety & Environment	Description
EMI / RFI Statements	EMC Class A: CE, FCC, VCCI, BSMI, RCM Safety: IEEE CB, UL/cUL and TUV
Environmental Standards	RoHS, GreenPC, WEEE
Temperature Range	Operational: 5° to 35°C (41° to 95°F) Non-Operational: -40° to 60°C (-40° to 140°F)
Humidity Range	Operational: 20% to 80% (Non-Condensing) Non-Operational: ~ 95% (Non-Condensing)
Acoustic Noise Levels	< 60dB, 25C
Shock	Operational: 5G, 11 ms duration Non-Operational: 30G, 11ms duration
Vibration	Operational: 0.2G, 5 to 500Hz (sine wave) ; 0.41G, 3-10-200-500Hz (Random) , Non-Operational: 1G, 5 to 500Hz (sine wave) ; 2.256G, 5-80-350-500Hz (Random)

Support & Warranty	Description
Support	• 24 hour, 7 days a week, 365 days a year e-mail and phone support (English only) • 24 hour, 7 days a week, 365 days a year access to PROMISE support site • Firmware and compatibility lists
Warranty	3-year full system limited warranty, optional extended warranty, on site parts replacement program

Hardware

The following section provides a summary of the front and back panel hardware features of the VTrak E5000 Series enclosures.

Front of VTrak E5800f and E5600f

Both VTrak E5000 enclosures feature handles on each side used to secure the enclosure to an equipment rack.

Front view of VTrak E5800f

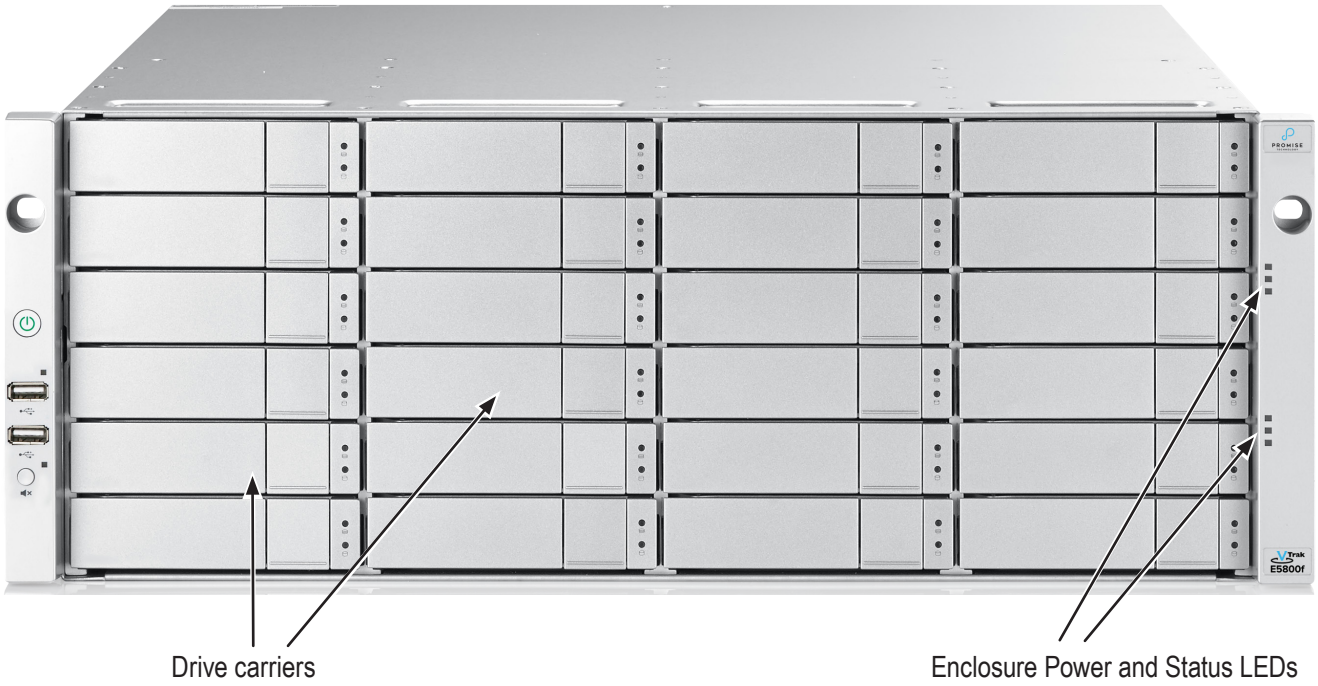

Front view of VTrak E5600f

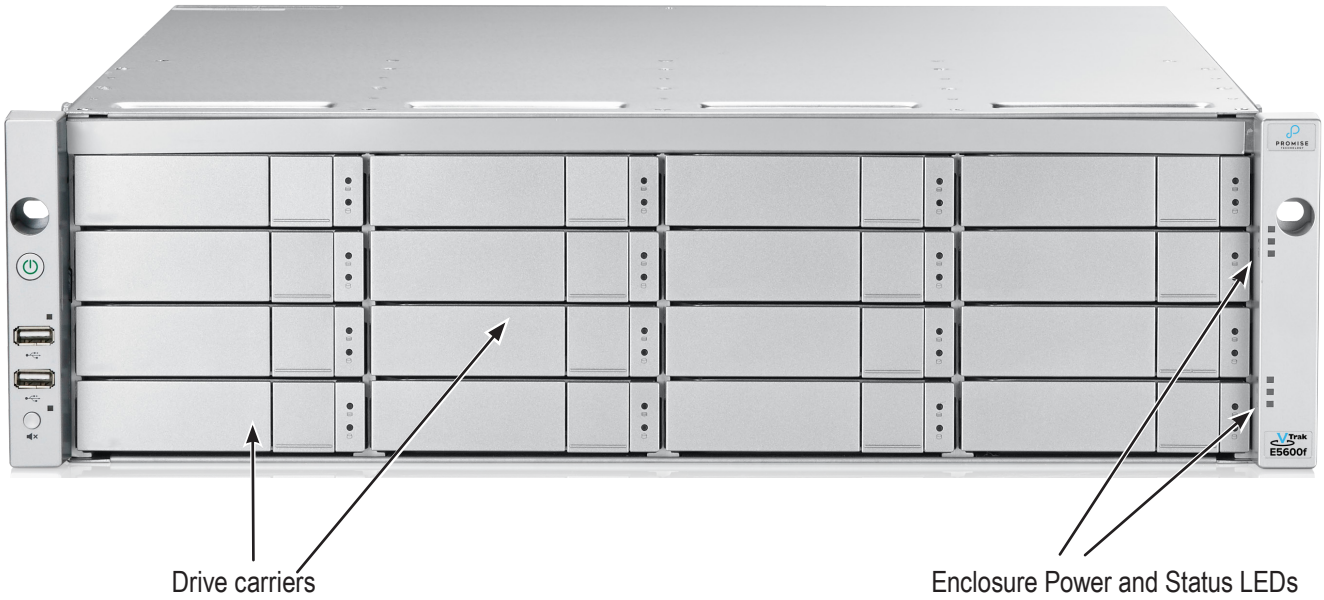

Front of VTrak E5300f and E5320f

The front panel of VTrak E5000 enclosures provide access to storage disk drive carriers, a view of drive status LED indicators located on the front of each drive carrier, the LED indicators for system monitoring.

See the illustrations of the front view of each model below.

VTrak E5300f front view

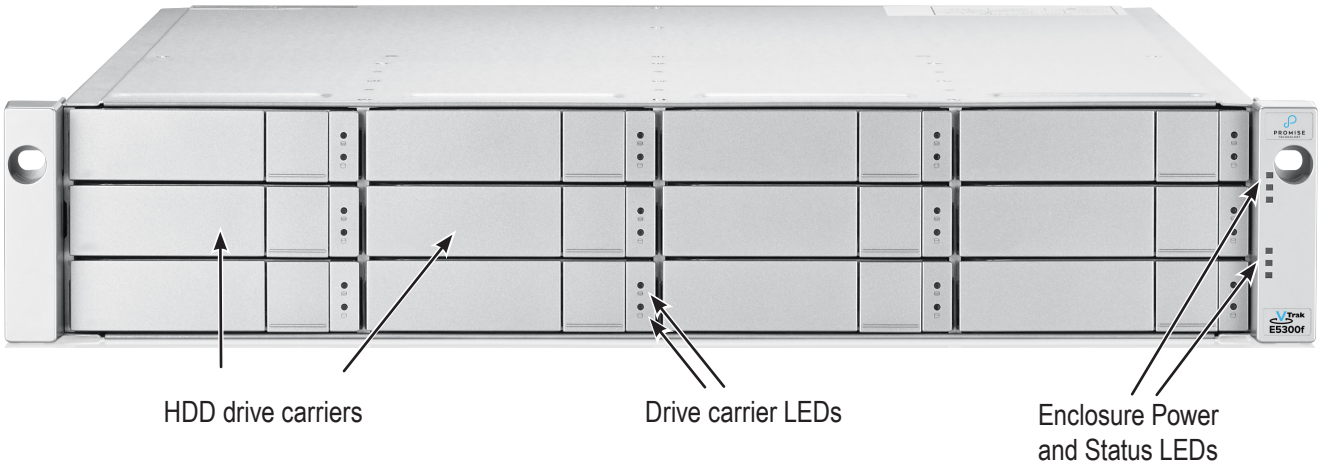

VTrak E5320f front view

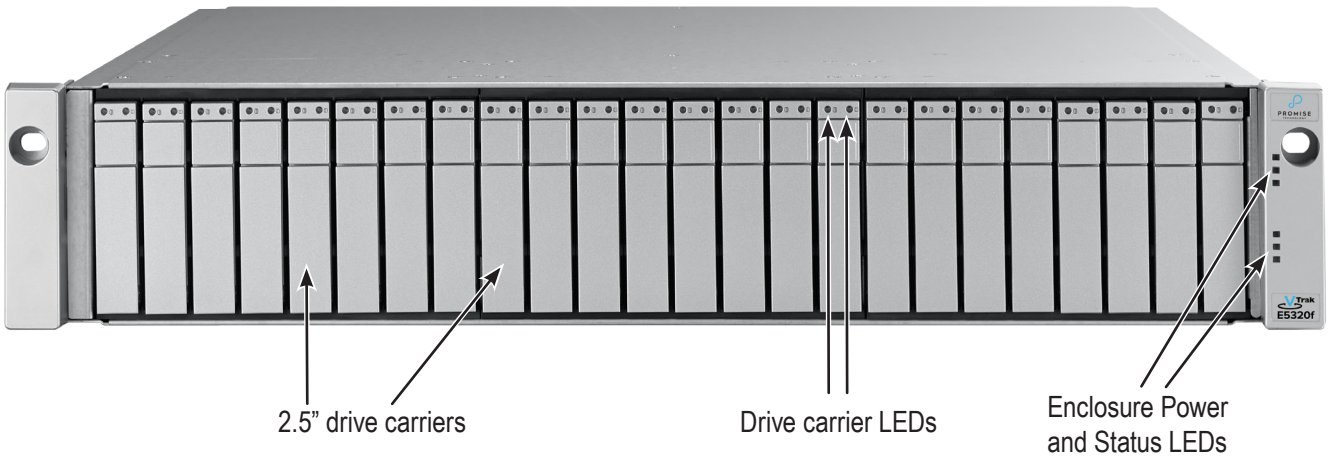

Back of VTrak E5800f and E5600f

The rear of the VTrak E5000 Series enclosure provides access to the power supply units, which include the cooling fans, and the system controller(s).

Each controller has as RJ-45 Gigabit Ethernet port used for management (and for Metadata communications when used as part of the VTrak A Class shared network file system), an RS-232 serial management port using an RJ-11 connector, four Fibre Channel data ports per controller using standard SFP sockets, one SAS Expansion port for additional drive enclosures, and various LED indicators which are described below.

Back view of VTrak E5800fD

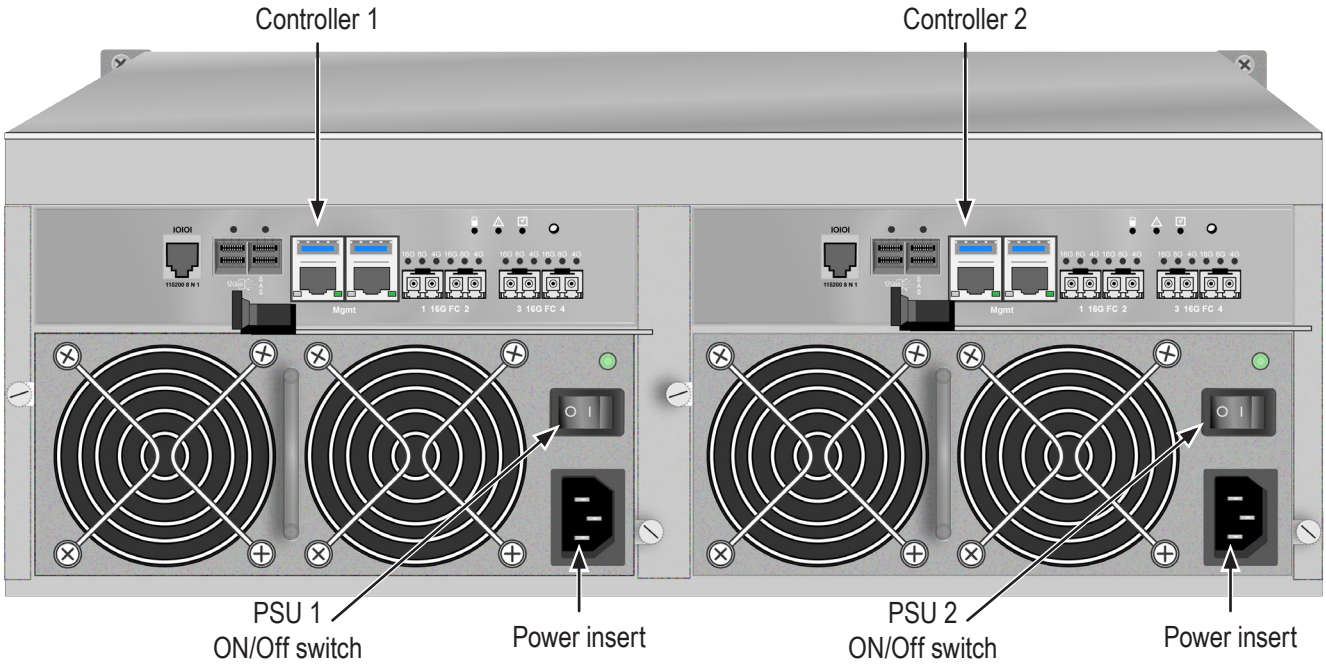

Back of VTrak E5300f and E5320f

The rear of the VTrak E5000 Series enclosure provides access to the power supply units, which include the cooling fans, and the system controller(s).

Each controller has as RJ-45 Gigabit Ethernet port used for management (and for Metadata communications when used as part of the VTrak A Class shared network file system), an RS-232 serial management port using an RJ-11 connector, four Fibre Channel data ports per controller using standard SFP sockets, one SAS Expansion port for additional drive enclosures, and various LED indicators which are described below.

VTrak E5300f and VTrak E5320f rear view

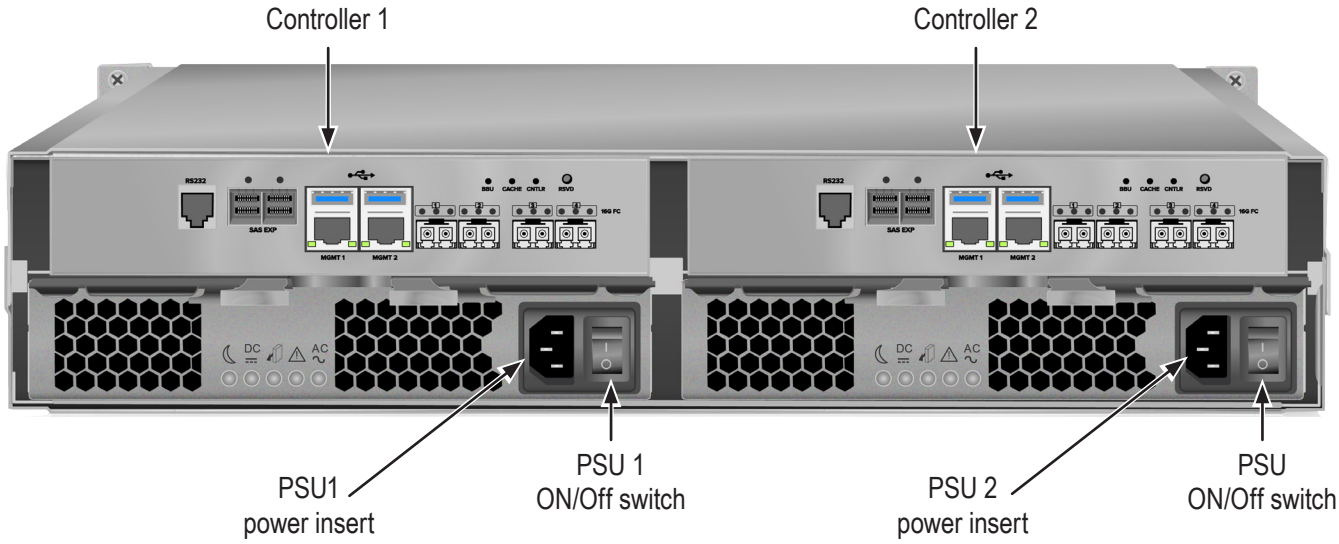

WARRANTY AND SUPPORT

WARRANTY

- Three year complete system limited warranty
- Battery Backup Unit has a one year limited warranty
- Optional 2-year extended warranty
- Optional onsite parts replacement program

Promise Technology, Inc. (“Promise”) warrants that for three (3) years from the time of the delivery of the product to the original end user except for one (1) year warranty on the battery backup unit:

a) the product will conform to Promise’s specifications;

b) the product will be free from defects in material and workmanship under normal use and service.

This warranty:

a) applies only to products which are new and in cartons on the date of purchase;

b) is not transferable;

c) is valid only when accompanied by a copy of the original purchase invoice;

d) is not valid on spare parts.

This warranty shall not apply to defects resulting from:

a) improper or inadequate maintenance, or unauthorized modification(s), performed by the end user;

b) operation outside the environmental specifications for the product;

c) accident, misuse, negligence, misapplication, abuse, natural or personal disaster, or maintenance by anyone other than a Promise or a Promise authorized service center.

Hardware Installation

This chapter presents basic information on unpacking the VTrak E5000 Series enclosure and mounting it in an equipment rack, making the connections for data and management* paths and connecting the power. It also describes how to power on the system and what to look for while it is powering up.

The main sections in Hardware Setup include the following:

- Unpacking
- Mounting the VTrak enclosure in a rack
- Installing Physical Drives
- Making Management and Data Connections
- Connecting the Power
- Power on the VTrak E5600fD/E5800fD

Depending on the details of your order, the VTrak E5000 Series enclosure might be shipped with hard drives installed, or it might require that you install hard drives. The section “Installing Physical Drives” on page 25 provides instruction for installing hard disks.

* A Vtrak E5000 Series unit used as a storage node on a SAN managed by the VTrak E5000 Series uses the “Management” port to connect to the Metadata communications network. See the VTrak E5000 Series Product Manual for more information.

Warning

Turn off the power and disconnect the power cord before servicing this device.

Unpacking

Packing List

The VTrak E5000 Series box contains the following items:

- VTrak E5000 Unit
- One Quick Start Guide printed
- Sliding rail assembly for rack mounting
- DB9-to-RJ11 serial data cable

Warning

The electronic components within the VTrak enclosure are sensitive to damage from Electro-Static Discharge (ESD). Observe appropriate precautions at all times when handling the VTrak or its subassemblies.

Mounting the VTrak enclosure in a rack

This section provides instructions for installing the VTrak E5000 Series enclosure into a rack

Caution

To lighten the enclosure, remove the power supplies, and remove all hard drive carriers. Replace the power supplies and drive carriers after the unit is mounted in your rack.

Cautions

- **Do not populate any unit with hard drives until it has been securely installed in the rack.**
- **At least two persons are required to safely lift, place, and attach the unit into a rack system.**
- **Do not lift or move the unit by the handles, power supplies or the controller units. Hold the system itself.**
- **Do not install the unit into a rack without rails to support the system.**
- **Only a qualified technician who is familiar with the installation procedure should mount and install the unit.**
- **Mount the rails to the rack using the appropriate screws and flange nuts, fully tightened, at each end of the rail.**
- **Do not load the rails unless they are installed with screws as instructed.**
- **The rails available for the PROMISE VTrak unit are designed to safely support that PROMISE VTrak unit when properly installed. Additional loading on the rails is at the customer's risk.**
- **PROMISE Technology, Inc. cannot guarantee that the mounting rails will support your PROMISE VTrak unit unless you install them as instructed.**

To install the VTrak into a rack with the supplied mounting rails:

1. Check the fit of the mounting rails in your rack system.

Installing the rails onto the rack

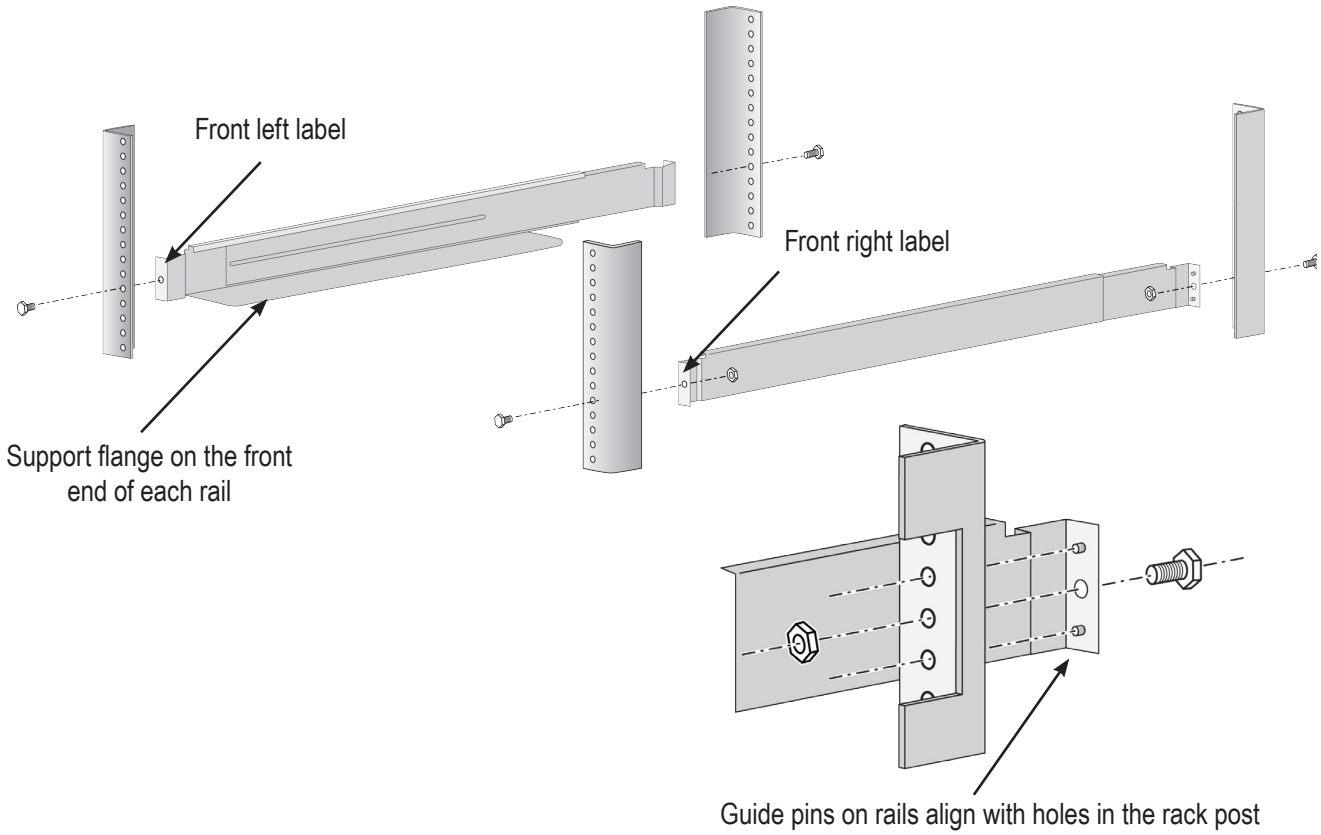

2. Adjust the length of the mounting rails as needed.

- The rear rail slides inside the front rail. The rails are composed of two sliding sections and do not require adjusting screws.
- The front-left and front-right mounting rail ends are labeled.
- Be sure the front rail support is on the bottom facing inward.
- All rail ends, front and rear, attach at the outside of the rack posts.
- The guide pins at the rail ends align with the holes in the rack posts.
- Use the attaching screws and flange nuts from your rack system. Tighten the screws and nuts according to instructions for your rack system.

Rail ends attach to the outside of each post

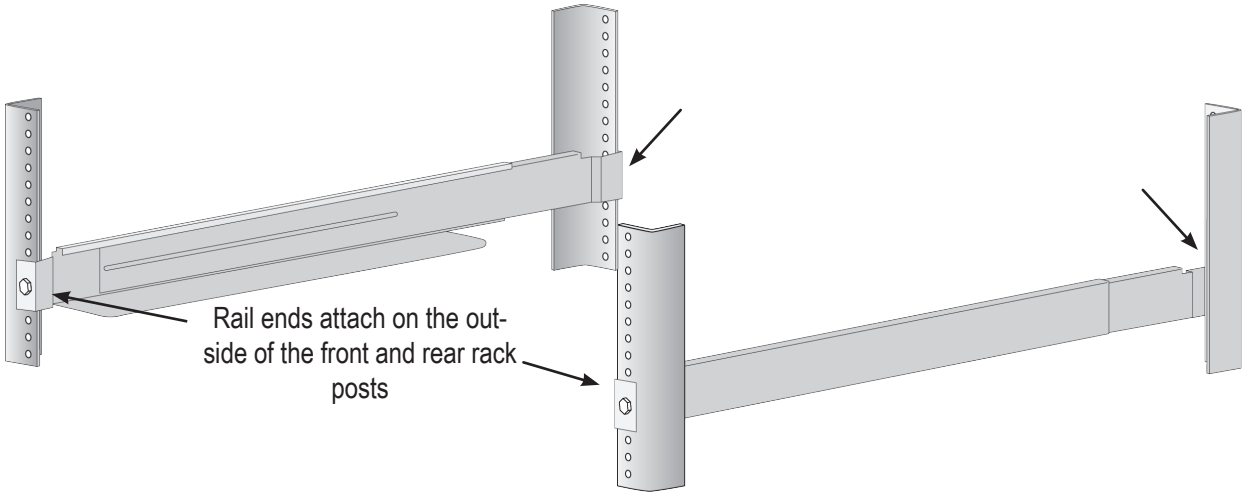

3. Place the VTrak onto the rails.

- At least two persons are required to safely lift the system.
- Lift the VTrak itself. Do not lift the system by its brackets.

Warning

Two persons are needed to safely place the unit onto the rails.
DO NOT lift the unit by the handles

Placing the VTrak system onto the rack rails

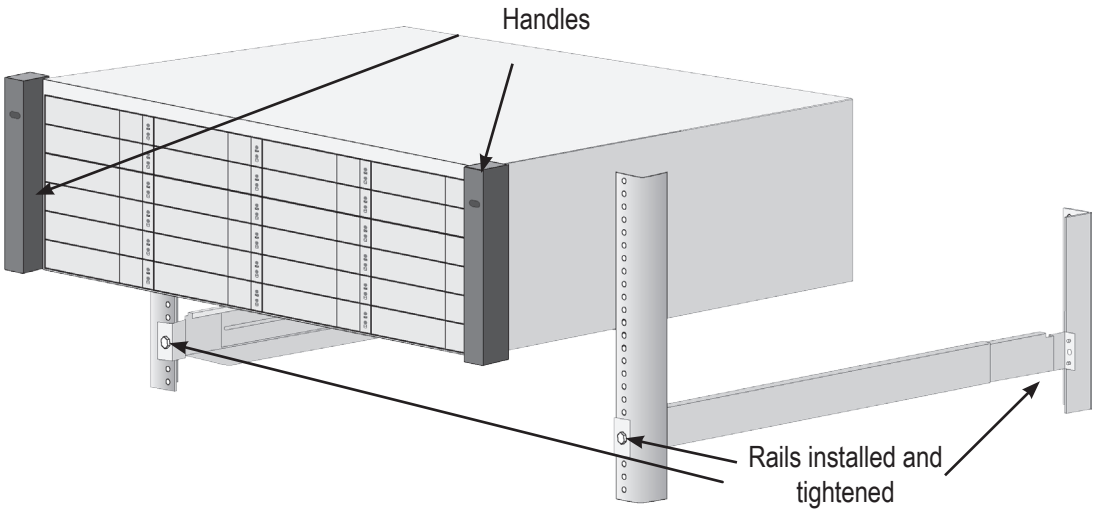

4. Secure the enclosure to the rack.

- Use the included screws and flange nuts to lock the unit in to place in the rack.
- Use the attaching screws and flange nuts that came with the VTrak.

Secure to rack

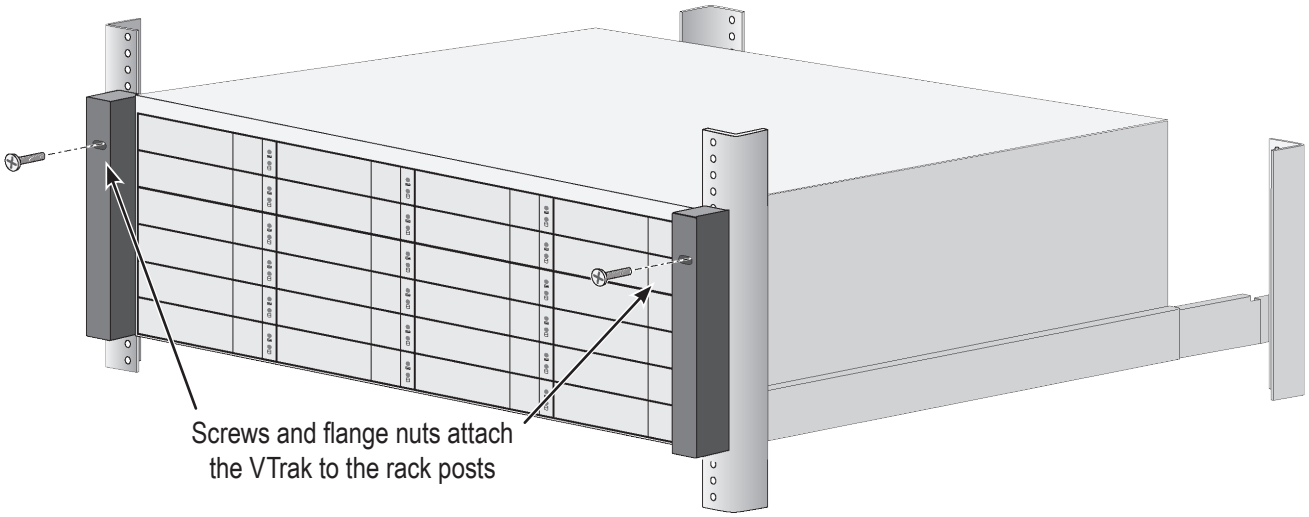

System installed in rack

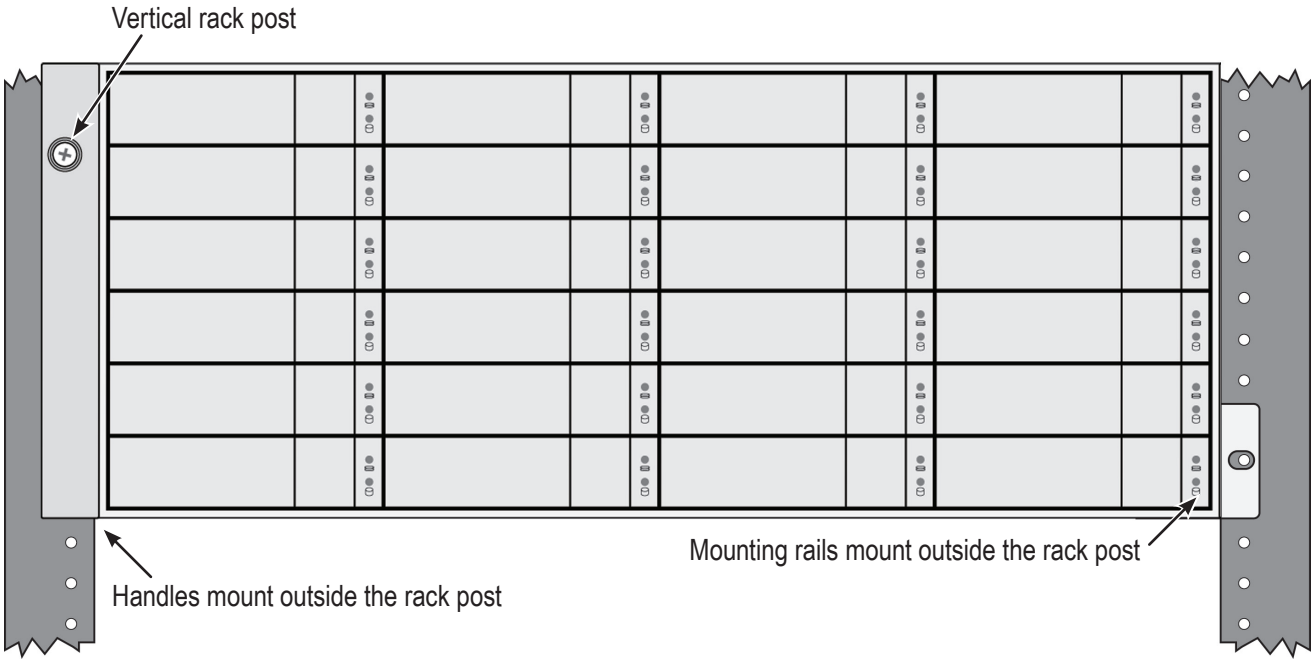

Mounting the VTrak E5300/J5320

To install the 2U VTrak enclosure into a rack with the supplied mounting rails:

1. Determine what height to place the 2U enclosure in the rack, then place the right and left rack rails at the same height on in the right and left rack position. Choose the mounting holes accordingly for your rack system. Consult the documentation for your rack if you are unsure which holes to use. Note that three holes are required on each front post, the uppermost of the three to be used for the flange nuts to anchor the enclosure to the rack posts.

Determine position for rack rails

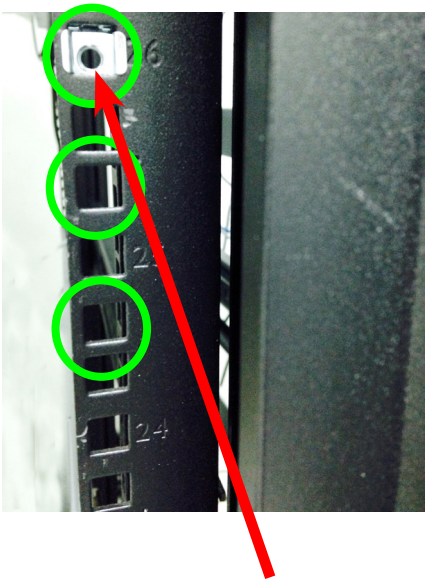

Insert flange nut in each front post

Notice that each end of the sliding rails have a lever to operate the lock mechanism that grips the rack post.

Lock release lever (back left)

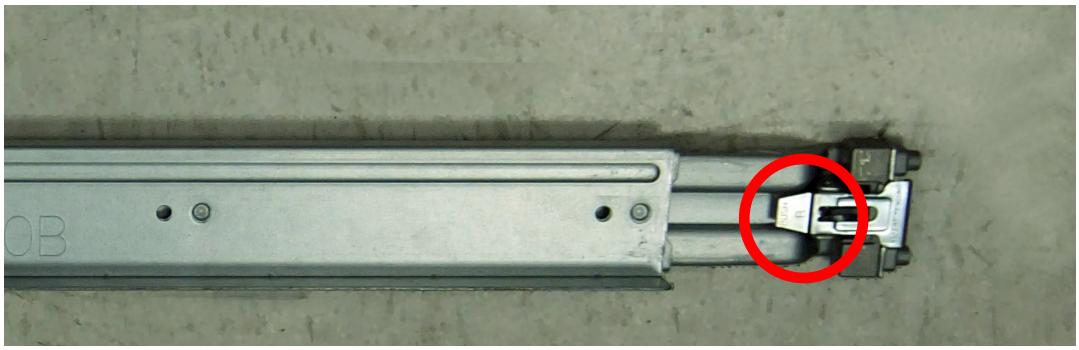

Press lever to release lock (front right)

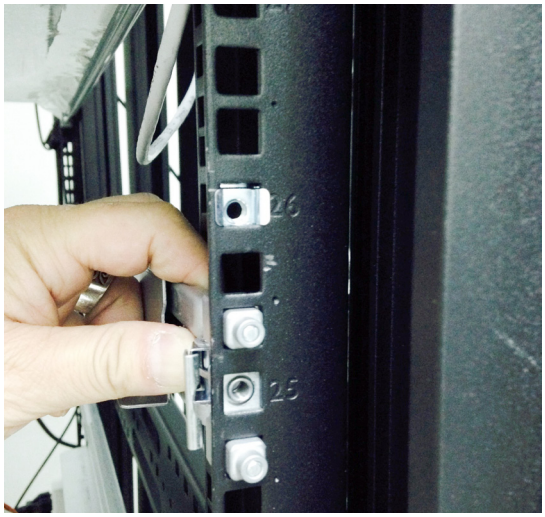

2. Secure the rails to the rack posts. Make sure the rack rails are properly oriented in the rack.

   To set the rails into the rack posts and secure the rails, follow these steps:

   a. Press the spring lock then insert the studs into the selected square holes on the rack post.

   b. Press the spring lock on the other end of the rail and insert the studs into the selected mounting hole on the rack post. If necessary, extend the rail to reach the post.

   c. Use the rail screws to anchor the rack rail to the post.

   d. Make sure the rack rail is aligned, secure, stable and in the correct place.

   e. Perform steps a through c above for the other rail.

   f. Make sure the rack rails are aligned, secure, stable and in place. See figure below.

Back left rail secured to post

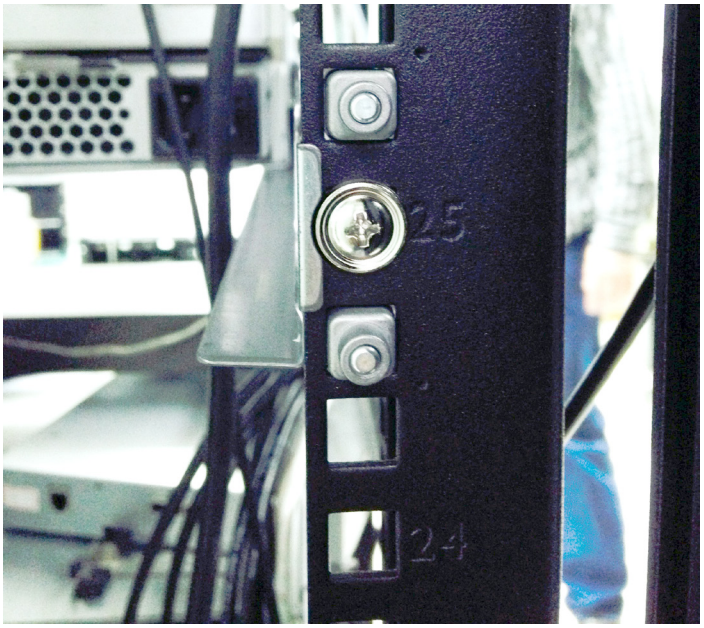

3. Secure the enclosure to the rack.

- Use the included screws and flange nuts to lock the unit in to place in the rack.
- Use the attaching screws and flange nuts that came with the mounting hardware.

Insert screws on each side of the front of the enclosure to secure it to the rack posts

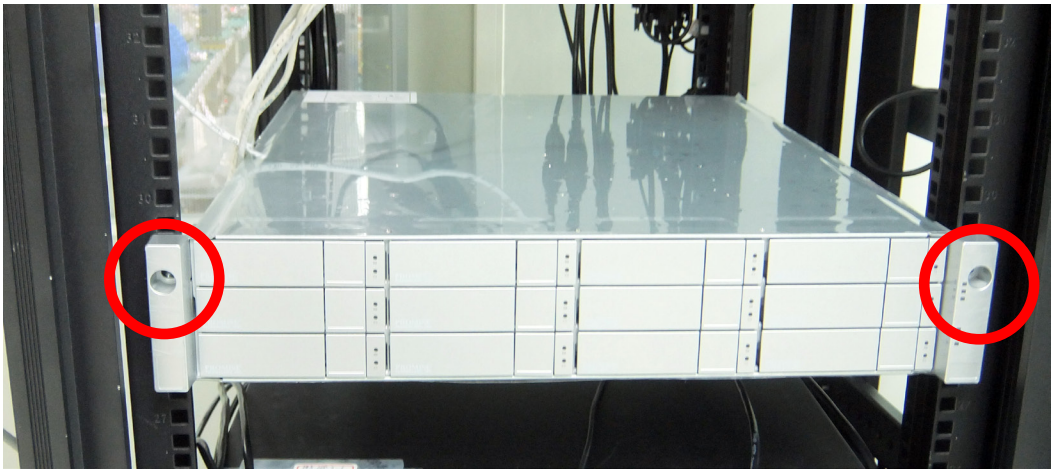

Installing Physical Drives

The VTrak E5000 Series subsystems support:

- SAS hard disks
- 3.5-inch hard disk drives for VTrak E5800fD, VTrak E5600fD, and VTrak E5300fD.
- 2.5-inch disk drives for VTrak E5320fD

For a list of supported physical drives, download the latest compatibility list from the PROMISE support website.

Number of Drives Required

The table below shows the number of drives required for each RAID level

Level	Number of Drives		Level	Number of Drives
RAID 0	1 or more		RAID 6	4 to 32
RAID 1	2 only		RAID 10	4 or more*
RAID 5	3 to 32		RAID 50	6 or more

*Must be an even number of drives.

Caution

The VTrak E5000 Series supports disk drive hot-swapping. To avoid hand contact with an electrical hazard, do not remove more than one drive carrier a time.

Drive Slot Numbering

You can install any suitable disk drive into any slot in the enclosure. The diagrams below shows how drive slots are numbered. Slot numbering is reflected in the WebPAM PROe and CLI user interfaces.

You can install any suitable disk drive into any slot in the enclosure.

Be sure to install all of the drive carriers into the VTrak E5000 enclosure to ensure proper airflow, even if you do not populate all the carriers with physical drives.

Drive slot numbering on VTrak E5800f

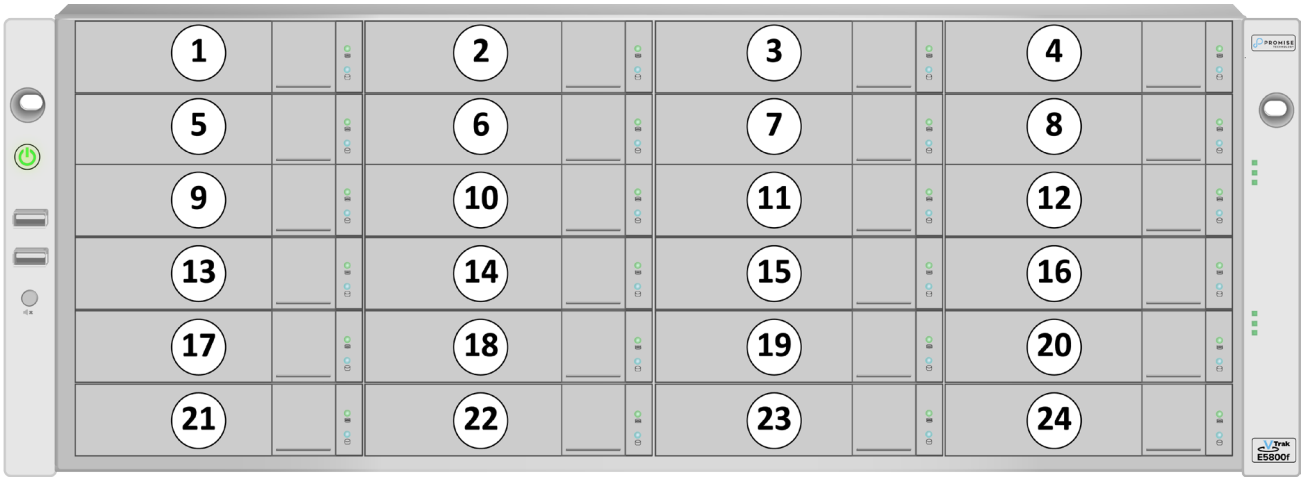

Drive slot numbering on VTrak E5600f

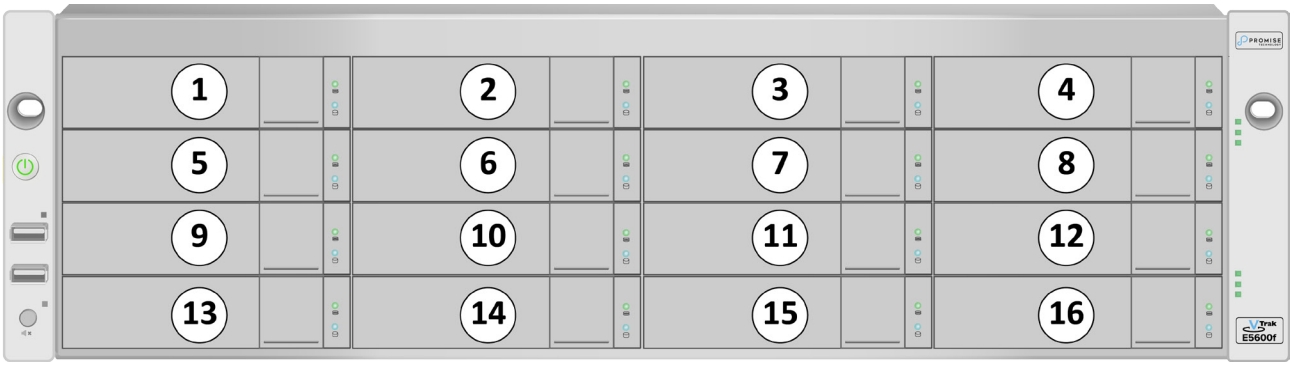

Drive slot numbering on VTrak E5300f

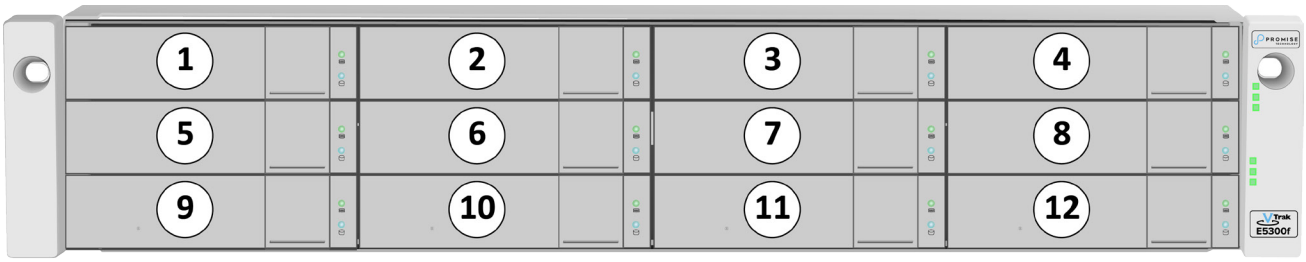

Drive slot numbering on VTrak E5320f

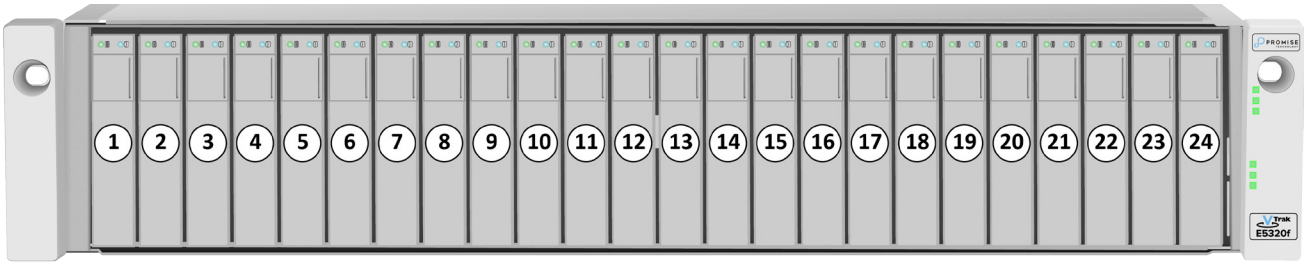

Installing Your Drives

The drive carrier accommodates 2.5-inch and 3.5-inch drives, with or without a SAS-to-SATA adapter.

Cautions

Swing open the drive carrier handle before you insert the drive carrier into the enclosure.

To avoid hand contact with an electrical hazard, remove only one drive carrier a time.

Important

SATA drives require a SAS-to-SATA adapter, available from PROMISE Technology at http://www.promise.com

SAS drives do not require adapters.

1. Press the drive carrier release button. The handle springs open.
2. Grasp the handle and gently pull the empty drive carrier out of the enclosure.

Drive carrier front view

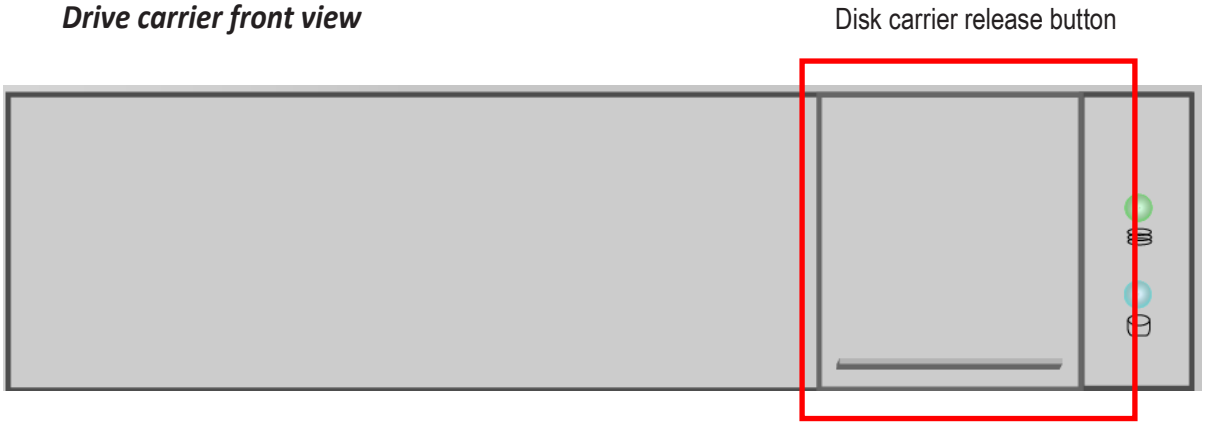

3. If you are installing SATA drives, attach a SAS-to-SATA adapter onto the power and IO connectors of each drive.
4. Carefully lay the drive into the carrier with the power and IO connectors facing away from the carrier handle.

5. Position the drive in the carrier so the mounting holes line up.
    - 2.5-inch drive mounting screws go through the bottom of the carrier.
    - SAS-to-SATA adapter mounting screws go through the bottom of the carrier.
    - 3.5-inch drive mounting screws go through the sides of the carrier.

Drive carrier bottom view

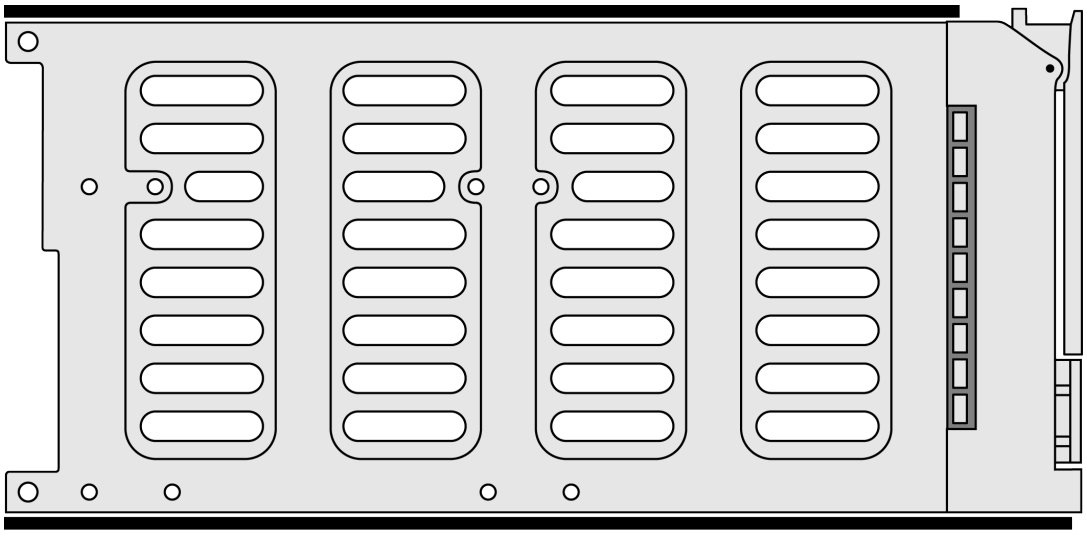

Drive carrier side view

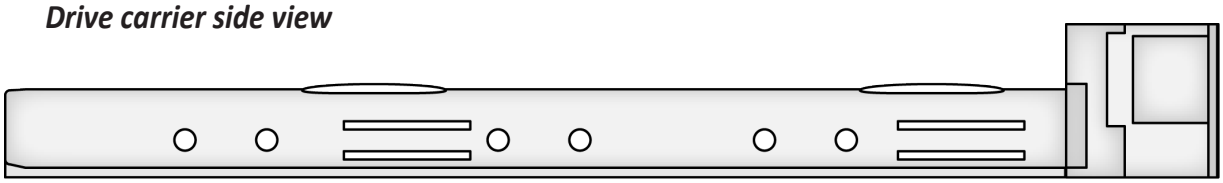

6. Insert the screws through the proper holes in the carrier and into the drive or adapter.
    - Use the screws supplied with the shipment or the SAS-to-SATA adapter.
    - Install four screws per drive.
    - Install two screws per adapter.
    - Snug each screw. Be careful not to over tighten.
7. With the drive carrier handle in open position, gently slide the drive carrier into the enclosure.

Important

Press the release button to push the drive carrier into position.

Proper drive installation ensures adequate grounding and minimizes vibration. Always attach the drive to the carrier with four screws.

2.5-inch Hard Disk Drives

The VTrak E5320fD features 24 drive carriers designed to fit 2.5” drives. Notice that the carriers are oriented vertically, with the carrier release button at the top. The lever mechanism to remove the carrier from the drive bay works exactly the same as the 3.5” carriers, except they are positioned vertically rather than horizontally.

Caution

Swing open the drive carrier handle before you insert the drive carrier into the enclosure.

2.5” drive carrier front view

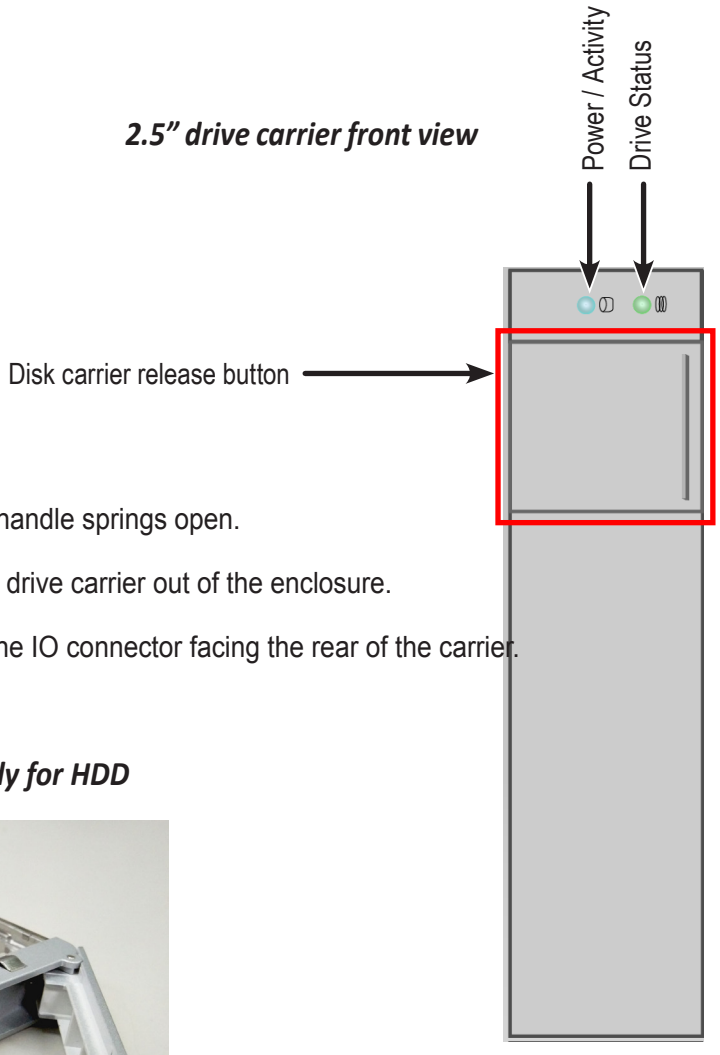

1. Press the drive carrier release button. The handle springs open.
2. Grasp the handle and gently pull the empty drive carrier out of the enclosure.
3. Carefully lay the drive into the carrier with the IO connector facing the rear of the carrier.

Empty 2.5” drive carrier front ready for HDD

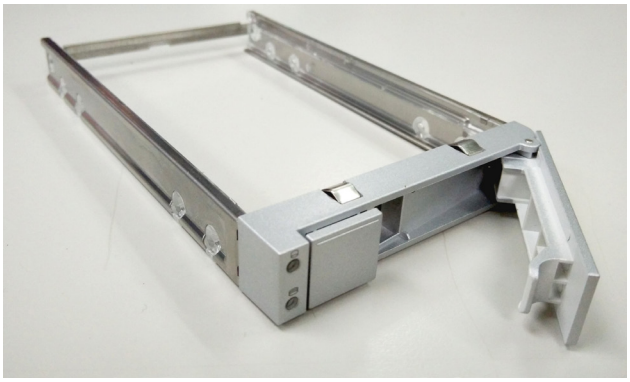

4. Position the drive in the carrier so the mounting holes line up.
5. Insert the screws through the proper holes in the carrier and into the drive or adapter.
   - Install four screws per drive.
   - Snug each screw. Be careful not to over tighten.

2.5" drive carrier with HDD installed, front view

2.5" drive carrier with HDD, 'left' side with HDD undercarriage exposed

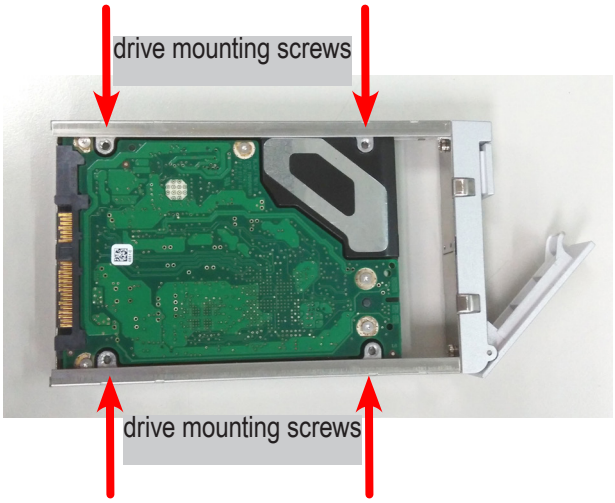

6. With the drive carrier handle in open position, gently slide the drive carrier into the enclosure. The drive carrier should be oriented so the the "top" of the 2.5" HDD is on the right, and the "bottom" of the HDD is on the left. The LED indicators will be at the top.

Making Management and Data Connections

This section describes how to establish a management connection to the VTrak E5000 Series subsystem. There are two methods to establish a management connection, Ethernet and Serial connection. For the initial setup, it is necessary to establish the Ethernet management connection via at least one of the four RJ-45 network ports.

The VTrak E5000 Series also features a Command Line Interface (CLI) for system management via a terminal emulation program (such as Microsoft HyperTerminal).

Examples of VTrak E5000 Series configurations in this section include:

- "Management Path network connection"
- "Fibre Channel SAN connections"
- "Fibre Channel DAS connections"

VTrak E5000 Series controller Network (RJ-45), Serial (RJ-11) and Fibre Channel (FC) ports

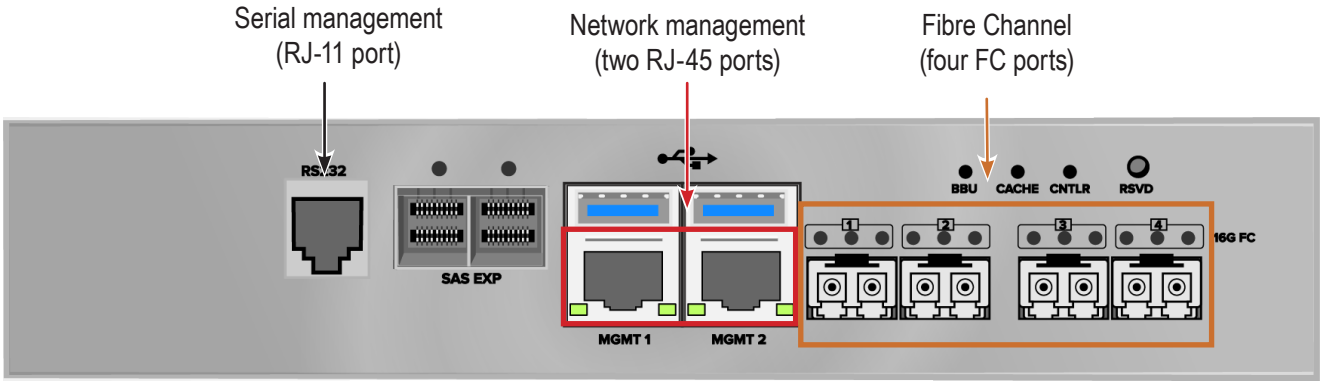

Management Path network connection

Each VTrak E5000 Series controller has two 1000BASE-T Ethernet ports for management over the network. Connect at least one of the management ports for initial setup. See “Management and Fibre Channel SAN data connections” on page 34 for an illustrated example of the network connections.

To establish the management path network connection:

1. Attach one end of an Ethernet cable to the network connector or standard NIC in the Host PC. Attach the other end of the Ethernet cable to a port on a standard network switch.
2. Attach one end of an Ethernet cable to a port on the same network switch. Attach the other end of the Ethernet cable to one of the Management Ports on a controller. Use another Ethernet cable to connect one of the Management Ports on the remaining controller.

If you have multiple VTrak E5000 Series subsystems, repeat steps 1 and 2 as needed.

Note

All RJ-45 network management ports on a VTrak E5000 Series subsystem share the same Virtual IP address. The default Virtual IP address, 10.0.0.1, applies to all RJ-45 network ports. If you change the Virtual IP address, the change applies to all RJ-45 ports.

Fibre Channel SAN connections

The Fibre Channel data network for the VTrak E5000 controllers requires the following items:

- A Fibre Channel connection in each host PC or server
- An SFP transceiver for each connected FC port on the subsystem
- A Fibre Channel switch
- Fiber Optic cabling
- A network switch

Important

For a list of supported HBAs, Switches, and SFP transceivers, download the latest compatibility list from PROMISE support: http://www.promise.com/support/.

Important

If the VTrak E5000 Series is going to be included on a SAN managed by a VTrak A Class unit such as the VTrak A830f, then Management ports are used to connect to the Metadata communications network.

For the Fibre Channel storage area network (SAN):

1. For servers equipped with Fibre Channel HBA cards, connect Fiber Optic cables between the Fibre Channel ports in both host PCs or servers and the ports on a Fibre Channel network switch.
2. Connect Fiber Optic cables between the Fibre Channel port on the VTrak E5000 controllers and a Fibre Channel port on a Fibre Channel switch or Fibre Channel capable switch (SFP).

   If you have multiple VTrak E5000 subsystems, host PCs or servers, repeat the steps as required.

Management and Fibre Channel SAN data connections

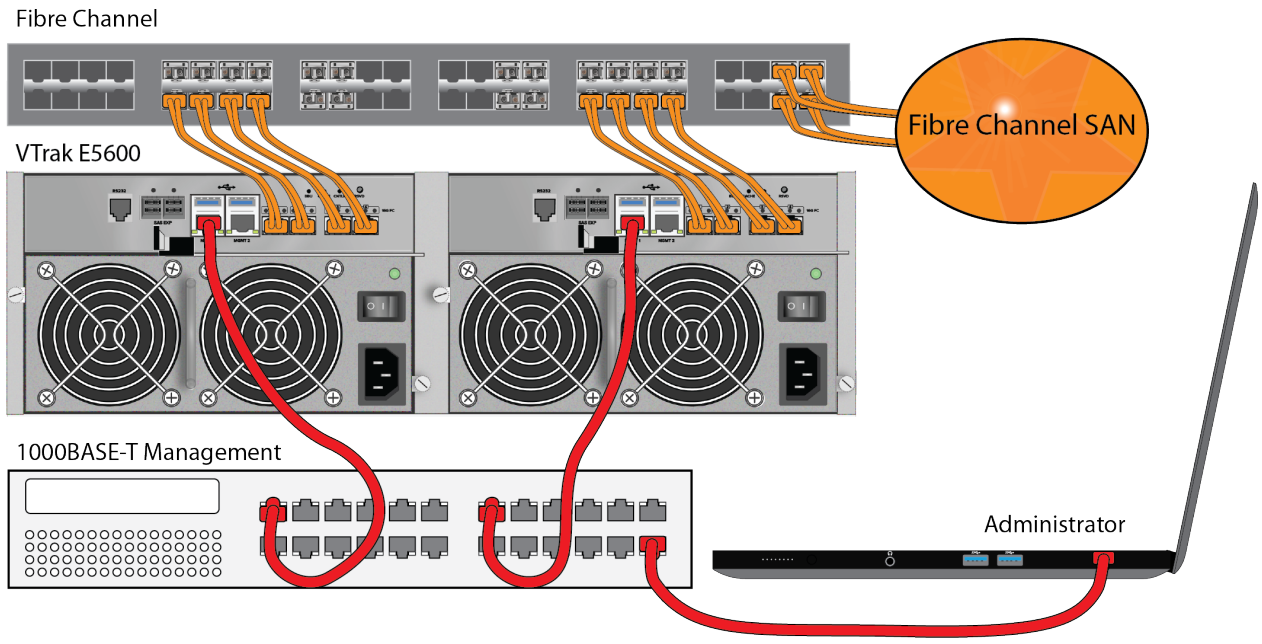

In the illustration above, the network management path is represented by the red cables. The Fibre Channel SAN data path is represented by the orange cables

Making Serial Cable Connections

Serial communication enables any computer that has an available serial port and terminal emulation application to access the VTrak Command Line Interface (CLI) to set up a network connection. The VTrak package includes one RJ11-to-DB9 serial data cable for each controller.

To set up a serial cable connection:

1. Attach the RJ-11 end of the serial data cable to the RJ-11 serial connector on one of the RAID controllers.
2. Attach the DB9 end of the serial data cable to a serial port on the host PC or server.

Controller serial interface

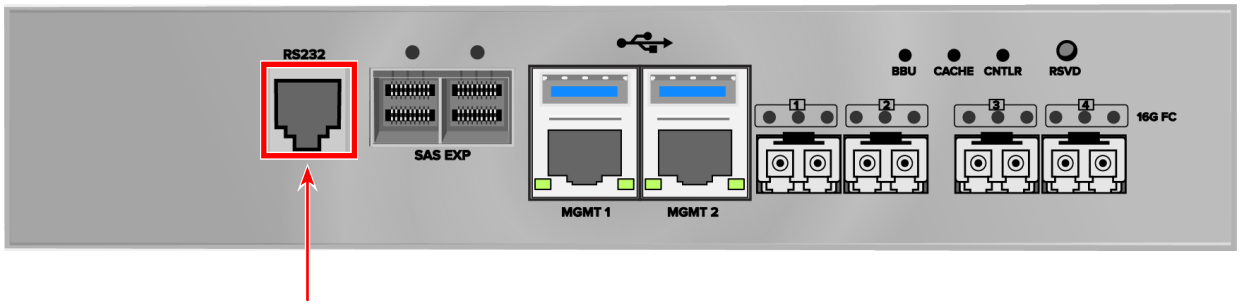

Serial port (RJ-11)
Use the DB9 to RJ-11 adapter to connect

Fibre Channel DAS connections

The following components are required to connect the VTrak E5000 to a host computer using the direct attached (DAS) method.

- At least one available Fibre Channel port - but two are recommended - on each host computer that is to be connected to the VTrak.
- A Fibre Channel transceiver for each port on the VTrak that is to be connected to a host computer.
- Fiber Channel optical cabling (LC/LC 62.5/125µm MMF)

As the name suggests, you are just connecting a host computer directly to the VTrak E5000.

Please see the example in the illustration on the next page.

Direct Attached Storage (DAS) Fibre Channel and network management connections

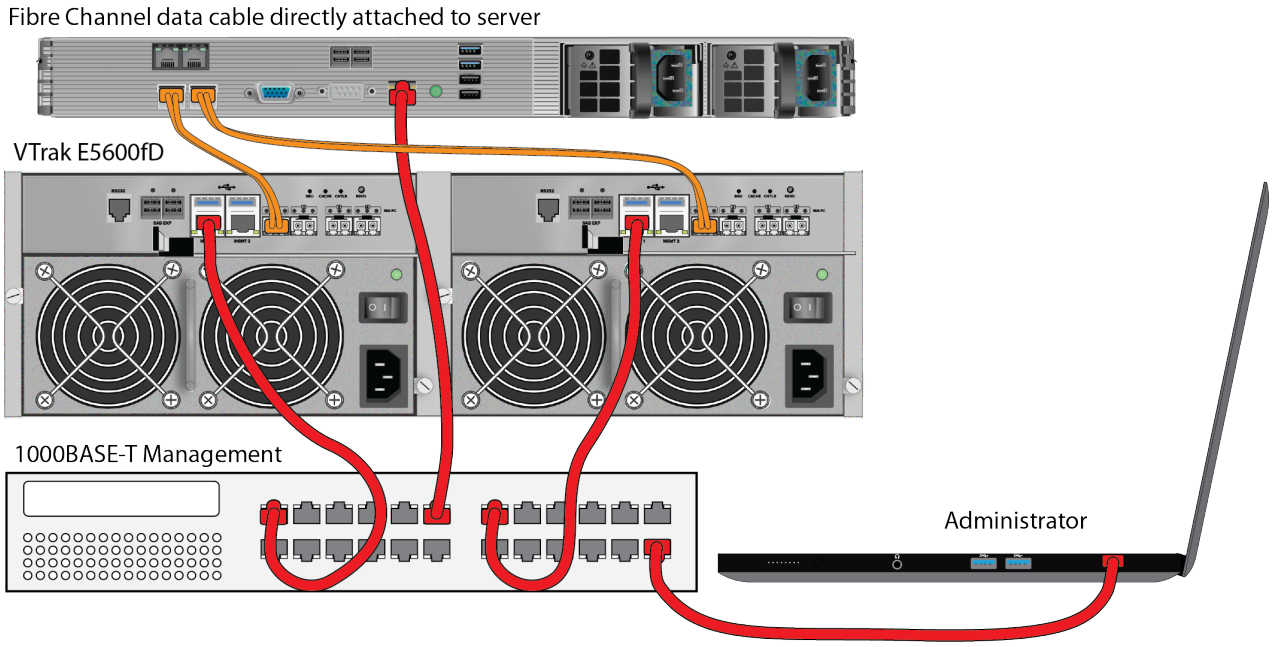

Note that it is recommended you use two Fibre Channel connections from the host computer to the VTrak E5000 so that there is one physical path from the host to each of the VTrak E5000's controller's. This improves performance and provides redundancy.

1. Connect a Fiber Channel optical cable to a Fibre Channel port on the host computer.
2. Connect the other end of the Fiber Channel optical cable to a Fibre Channel port on one of the VTrak E5000 controllers.
3. If you are using a second path, connect another Fiber Channel optical cable to a second Fibre Channel port on the host computer and connect the other end to a Fibre Channel port on the other VTrak E5000 controller.

Fibre Channel with JBOD Expansion

VTrak J5000 Series JBOD expansion requires at least one SFF-8644 to SFF-8644 external SAS cable for each VTrak J5000 unit. The setup description below references the illustration “VTrak E5600fiD with VTrak J5600 JBOD expansion connections” on the next page.

To add VTrak J5000 Series units:

1. Connect SAS expansion port 1 on the left controller of the RAID subsystem to the SAS port 1 on the left I/O module of the first VTrak J5000 unit.
2. Connect the SAS expansion port 1 on the right controller of the RAID subsystem to the SAS data port 1 on the right I/O module of the first VTrak J5000 unit.
3. Connect the SAS data port 2 on left I/O module of the first VTrak J5000 unit to the SAS data port 1 on the left I/O module of the second VTrak J5000 unit.
4. Connect the SAS data port 2 on right I/O module of the first VTrak J5000 unit to the SAS data port 1 on the right I/O module of the second VTrak J5000 unit.
5. Connect the remaining VTrak J5000 units in the same manner.

Keep in mind the following points:

- Keep your data paths organized to ensure redundancy.
- JBOD expansion supports up to nine VTrak J5000 units.

Important

Power on the JBOD units first, when you are ready to power on the enclosures.

Read the VTrak J5000 Series Product Manual for information on the VTrak J5000 Series enclosures.

The VTrak E5000 Series subsystem is shown with SFP transceivers installed.

VTrak E5600fiD with VTrak J5600 JBOD expansion connections

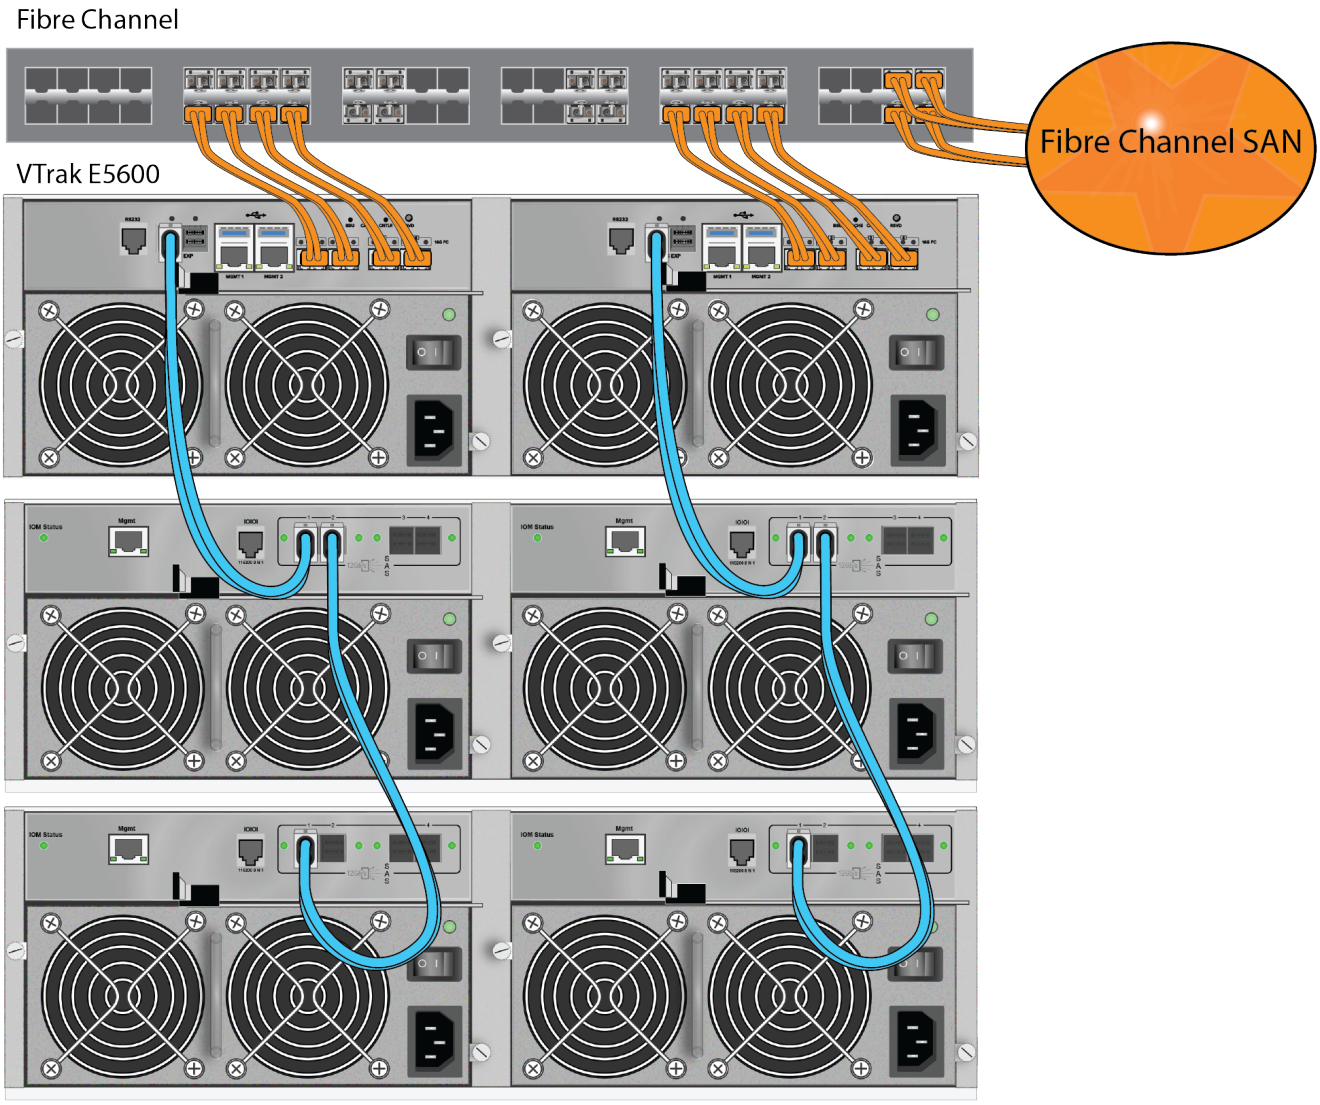

Connecting the Power

VTrak E5000 Series enclosures are equipped with two power supplies for each unit. All VTrak models feature an ON/OFF switch to toggle power located on each power supply unit (PSU).

Important

If you have a SAN, DAS, or Cascade setup with VTrak J5000 Series JBOD units, always power on the JBOD subsystems first.

Connect both power supplies to a suitable power source.

The 2U VTrak E5300fD/E5320fD will power on when the power switch on each power supply is in the *On* position. *See illustration below.*

The 3U VTrak E5600fD and 4U VTrak E5800fD have a power button on the front used to power on the subsystem. *See illustrations on next page.*

Power supplies on VTrak E5320fD and VTrak E5320fD

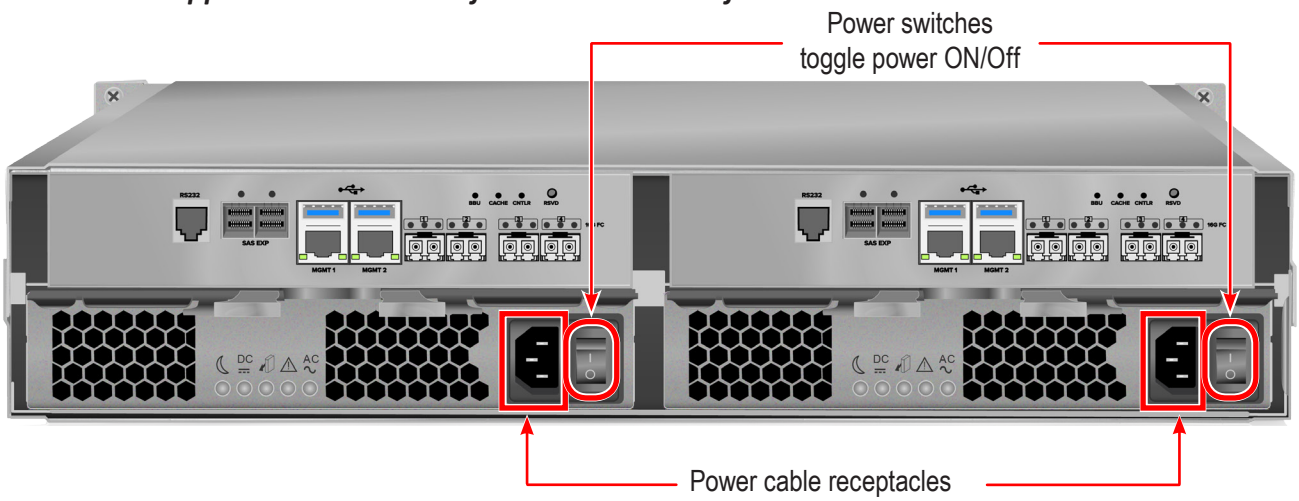

Power supplies on VTrak E5600fD and VTrak E5800fD

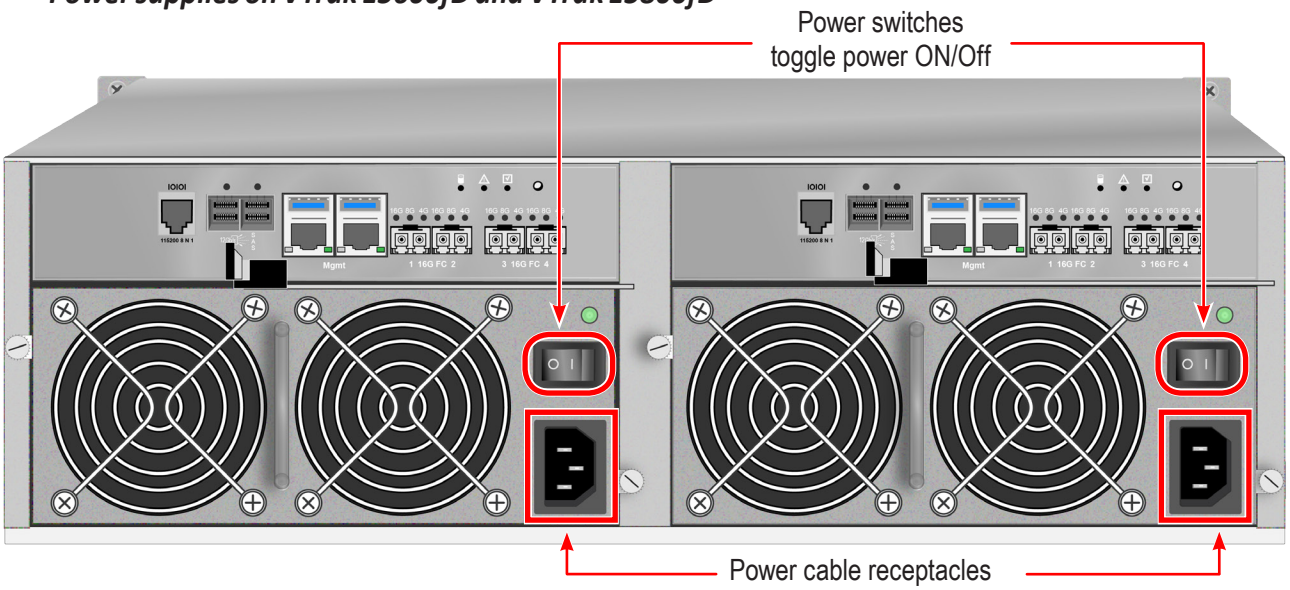

Power button on front left of VTrak E5600fD/E5800fD

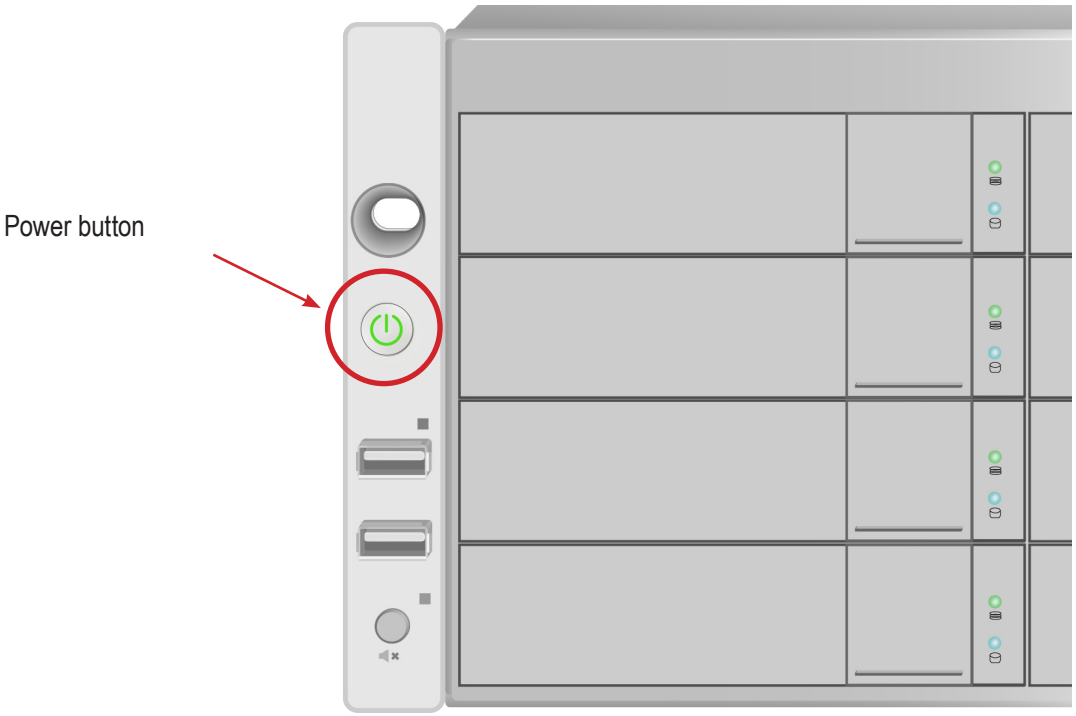

Power on the VTrak E5600fD/E5800fD

With the power supplies connected, turn on both power switches on the back of the enclosure.

Important

If you have a SAN, DAS, or Cascade setup with VTrak J5000 Series JBOD units, always power on the JBOD subsystems first.

Use the Power button on the front left or the VTrak enclosures to first power on any connected VTrak J5000 Series JBOD enclosures, then power on the VTrak E5000 Series enclosures. Observe the LEDs on the right front bracket facing.

Front LED Behavior

When boot-up is finished and the subsystem is functioning normally:

- Power, FRU and Logical Drive LEDs display Green continuously
- Controller Activity LED flashes Green when there is controller activity.
- System Heartbeat LED blinks Green once per second for five seconds, then goes dark for ten secondes, then repeats the same pattern.

Also on the front panel, there are two LEDs on each drive carrier. These report the presence of power and a physical drive, and the current condition of the drive.

LED indicators on front right of VTrak E5600fD/E5800fD

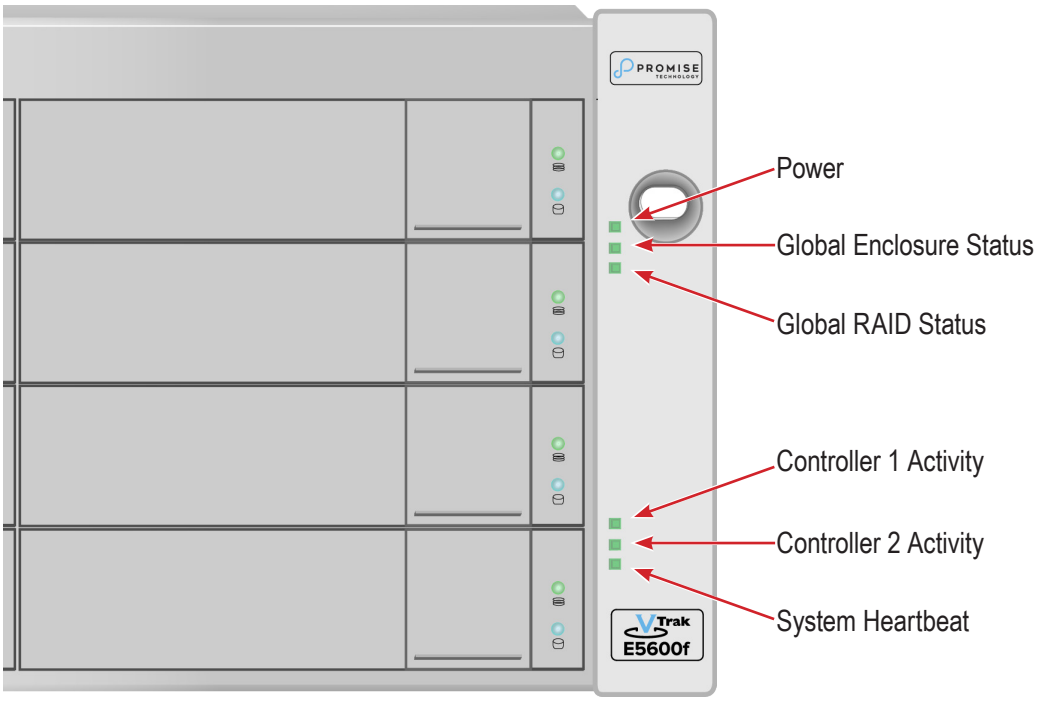

Front right LED Behavior After Boot Up

State	Power	Global Enclosure	Global RAID	Controller Activity	Controller Heartbeat
Dark	No power	No power	—	Unit is off, or controller is not present or FC is not cable connected	—
Steady Green	Normal	Normal	Normal	Unit is up, controller is present and running, at least one FC cable is connected and a link is established.	—
Blinking Green	—	—	—	—	Normal**
Flashing Green	—	—	—	Activity	—
Amber	—	Problem*	Critical	—	—
Red	—	Failure*	Offline	—	—

* Check the LEDs on the back of the VTrak enclosure (controllers and PSUs).

** Blinks blinks green once per second for two seconds for dual controller enclosure; blinks every four seconds for single controller enclosure.

Disk Carrier LEDs - front of every carrier

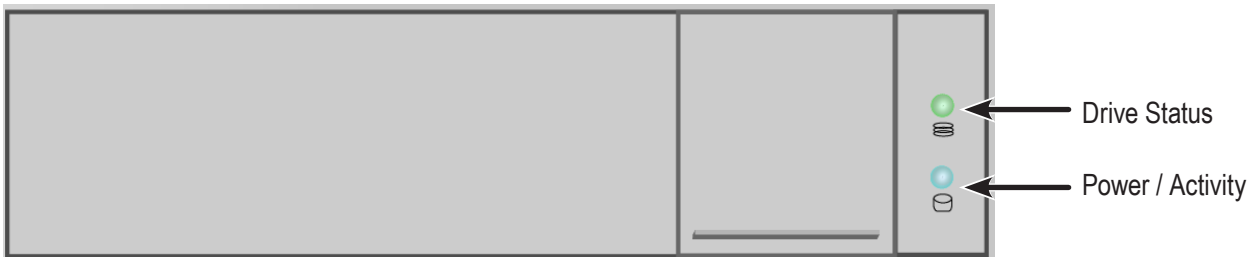

The VTrak spins up the disk drives sequentially to equalize power draw during start-up. After a few moments:

- The Power/Activity LED displays blue when a physical drive is present.
- The Drive Status LED displays green when the physical drive is configured as a member of a disk array or as a spare. When the physical drive is unconfigured, the LED is dark.

Steady means the LED is on.

Blinking means a regular on/off pattern.

Flashing means intermittent and irregular on/off pattern.

Drive Status LED Behavior After Boot Up

State	Power/Activity	Drive Status
Dark	No drive in carrier	Drive is unconfigured
Steady Blue	Drive in carrier	—
Flashing Blue	Activity on drive	—
Steady Green	—	Drive is configured
Blinking Green	—	Locator feature
Amber	—	Drive is rebuilding
Red	—	Drive error of failure

* Configured means the physical drive either belongs to an array or it is assigned as a spare drive.

VTrak E5600f / VTrak E5800f Power Supply LEDs

The LEDs on the rear panel of the VTrak E5600fD and VTrak E5800fD include a single status LED on each power supply. These PSU status LED will light green to indicate normal operation. A red LED indicates a problem or unit failure.

Status LED on Power Supply

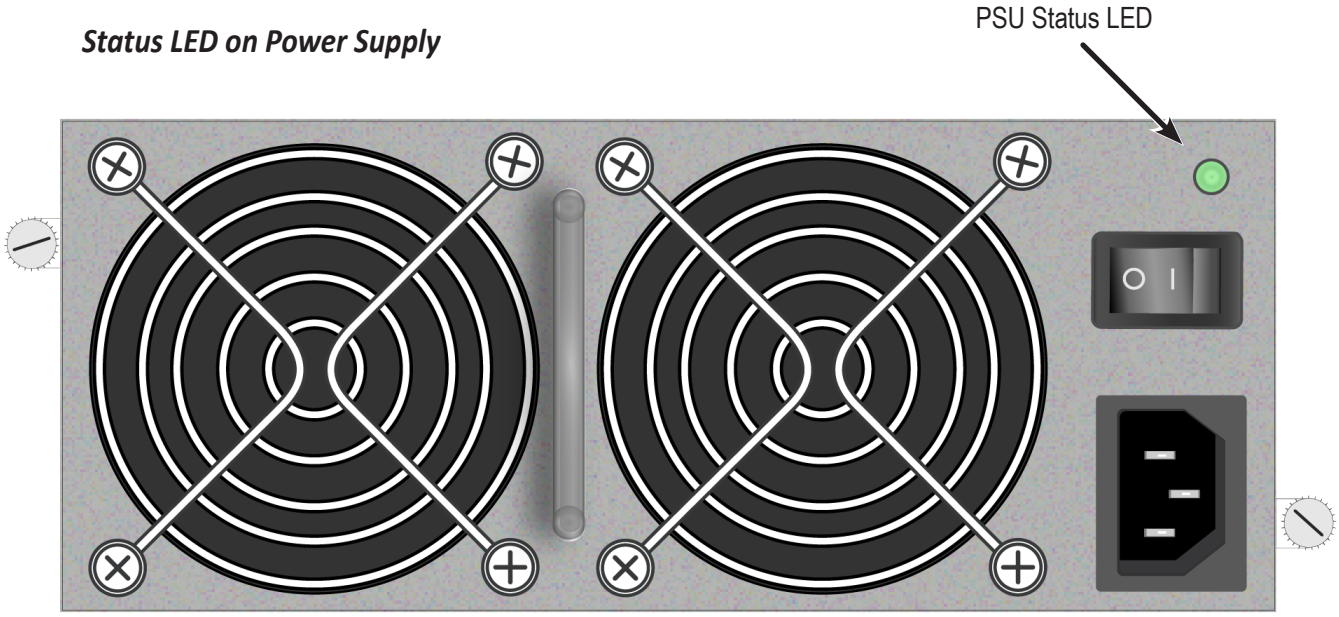

VTrak E5300f / VTrak E5320f Power Supply LEDs

The power supplies for the VTrak E5300 and VTrak E5320 have five LED indicators on each unit. See the table below for PSU LED behavior.

Power supply LED indicators on VTrak E5300

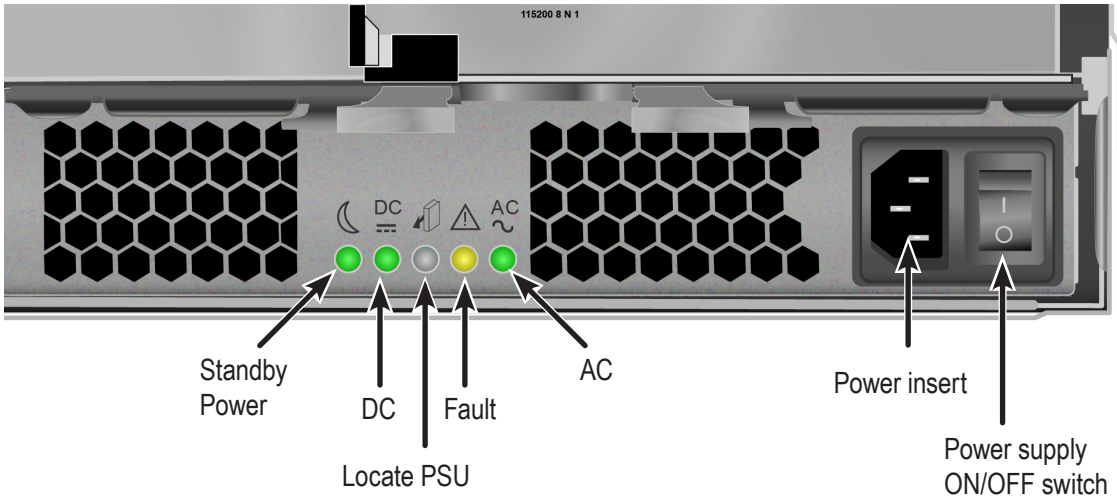

VTrak E5300f/E5320f Power Supply LEDs

LED	Description
Standby Power	This indicates the standby power status. If the standby power is on and the power switch is off, the LED lights green. If the standby power is on and the power switch is on, then the LED will be off.
DC	This indicates if the power supply is properly inserted into the enclosure. The LED lights green when the power supply is properly inserted and the power supply is switch on. It is off if the power supply is not properly inserted, or when the power supply is switched off.
Fault	This indicates the overall health status of the power supply. When the power supply is functioning normally and no problems are detected, it will be off. If a problem with the power supply is detected, it light amber. Note that this LED blinks once when the power supply is switched on, then remains off unless there is a problem.
AC	This indicates that input power is present. When the power supply is switched on, it lights green if input power is available. It is off if power is not present and when the power supply is switched off.
Locate PSU	This flashes blue when using the *Locate PSU* function.

Controller LEDs

When boot-up is finished and the VTrak E5000 subsystem is functioning normally:

- Controller status LEDs display green continuously.
- Ethernet LEDs display green or flash depending on your network connection.
- The FC, SAS, and Expansion LEDs display green or flash during port activity.

See table on next page for complete description of controller LEDs.

Controller LEDs

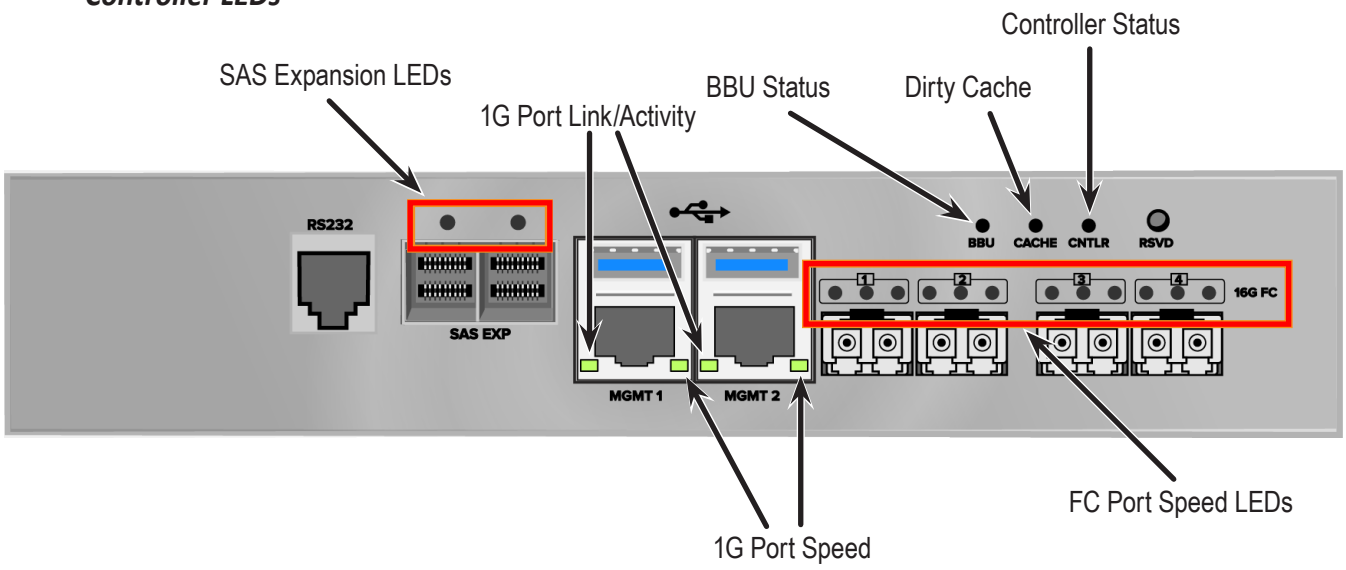

Controller LED Behavior

The table below describes behavior of the LED indicators on the VTrak E5000 Series controller.

LED	Description
SAS Expansion	One LED indicator for each SAS expansion port. These light green when connected, and flash green when there is activity.
Controller Status	This displays the current operational status of the controller. A steady (unblinking) green light indicates the controller is operational. This will blink green during when using the controller locator feature. A blinking amber light indicates a problem. Steady red light indicates controller failure. A flashing red light means the controller is in Maintenance Mode (offline while in Maintenance Mode).
Dirty Cache	Lights steady amber if cache is dirty, meaning that the controller memory cache contains data, otherwise this is dark. This will blink green during when using the controller locator feature.
Battery Status	This lights steady green when the battery status is healthy (normal). Red indicates the battery has failed. A steady amber light indicates there is not enough reserve power in the battery to backup cache memory if the power fails.
FC ports	See next page

LED indicators for Fibre Channel ports, three LEDs for each FC port

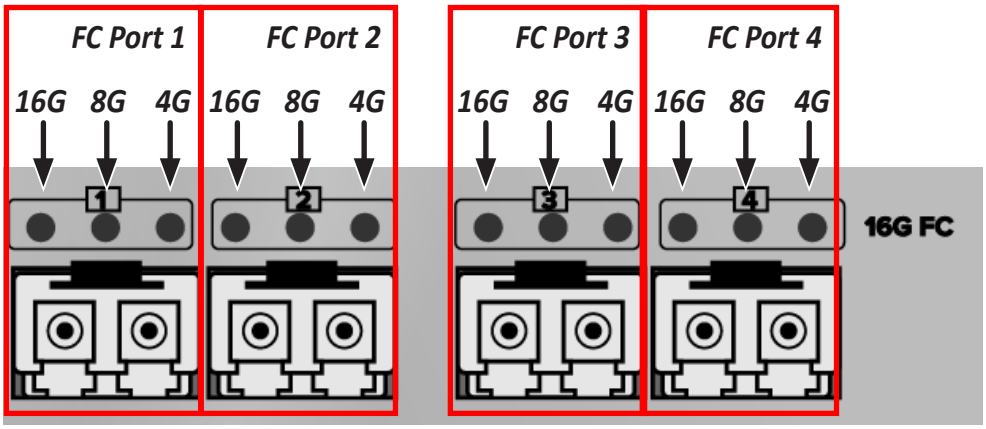

Fiber Channel Port LED behavior

LEDs	Power on (Before Firmware Initialization)	Power on (After Firmware Initialization)	Firmware Fault	4 Gbps Link Up/ ACT	8 Gbps Link Up/ ACT	16 Gbps Link Up/ ACT
Green LED (16 Gbps)	On	Flash	Flash in sequence	Off	Off	On/Flash when active
Green LED (8 Gbps)				Off	On/Flash when active	Off
Green LED (4 Gbps)				On/Flash when active	Off	Off

All Fibre Channel port LED indicators will be dark when the system is powered off. If all three indicators for a port flash simultaneously, then there is no SFP transceiver installed, or the wrong type of transceiver is installed, or the port is not connected.

Resetting the Default Password

This feature resets the Administrator's password to the default factory setting, password. Use this feature when you have forgotten Administrator's password or a new Administrator has been appointed.

The reset applies to the Administrator's login for WebPAM PROe and the CLI. No other user passwords are affected.

To reset the Administrator's default password:

1. Verify that the VTrak has fully booted. See "Front LED Behavior" on page 42.
2. For one of the RAID controllers, locate the password reset switch. *See illustration below*.
3. Insert a pin or a straightened paper clip into the opening and momentarily depress password reset switch.

You only need to press the reset switch on one RAID controller.

The next time the Administrator logs in, use the default password: **password**.

For more information, see "Logging into WebPAM PROe" on page 76 and "Logging Into the CLI" on page 224.

Important

PROMISE recommends that you change the Administrator's default password immediately after reset. See "Changing User Passwords" on page 123.

Password reset switch on controller

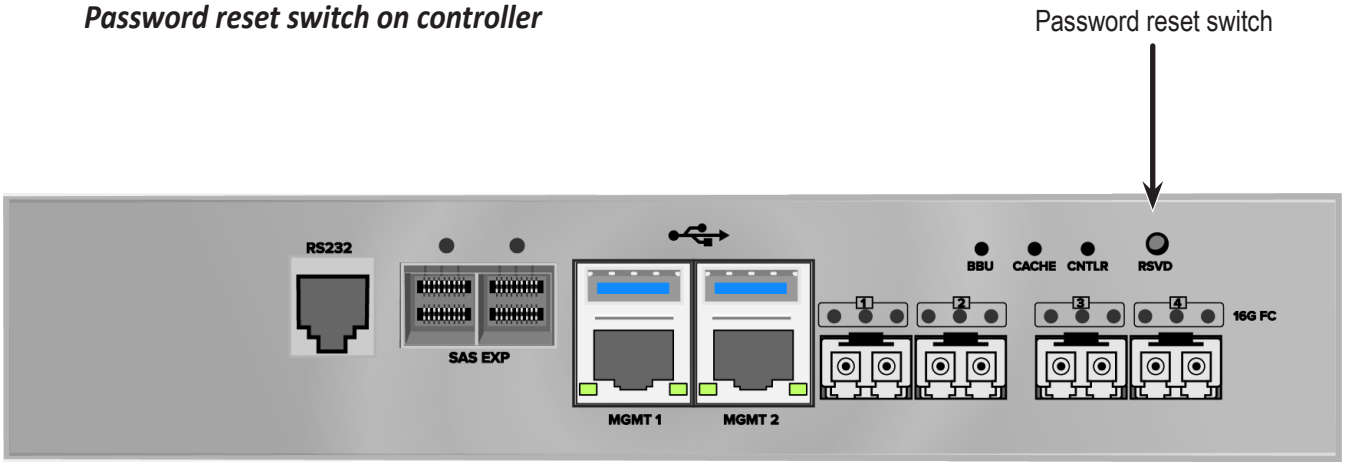

System Setup

Now that the VTrak E5000 Series subsystem is installed and connected, it is time to continue with setting up the storage arrays and perform other administration functions. You have a choice of user interfaces for management and administration of the VTrak E5000 Series. The administrator can choose to use WebPAM PROe, a web-based graphical user interface (GUI), or use command line based interface (CLI) . Administration via CLI and WebPAM PROe are described in detail in separate chapters later in this manual.

This chapter covers the following topics:

- "Setting-up the Serial Connection"
- "About IP Addresses"
- "Setting-up with the CLI"
- "Setting up with WebPAM PROe"
- "Creating Logical Drives"

Setting-up the Serial Connection

The initial connection accesses the serial port using the serial cable connection you made using the RJ-11 to DB9. Use your PC's terminal emulation program, such as Microsoft HyperTerminal, to access the Command Line Interface (CLI).

To make the initial serial connection:

1. Change your terminal emulation application settings to match the following specifications:
   - Bits per second: 115200
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: none
2. Start your PC's terminal VT100 or ANSI emulation program.
3. Press **Enter** once to launch the CLI.
4. At the Login prompt, type **administrator** and press **Enter**.
5. At the Password prompt, type **password** and press **Enter**.

   The screen displays:

```
login as: administrator
administrator@vtak's password:********
------------------------------------------------------------
Promise VTrak Command Line Interface  Utility
Version: X.XX.XXXX.XX Build Date: Xxx X, 2016
------------------------------------------------------------
------------------------------------------------------------
Type help or ? to display all the available commands
Type menu to enter Menu Driven Configuration Utility
------------------------------------------------------------
administrator@cli>
```

To see the full set of CLI commands, at the administrator@cli> prompt, type `help` and press **Enter**.

To see full information about a specific command, at the administrator@cli> prompt, type `help` followed by the command, then press **Enter**.

```
administrator@cli> help net
```

About IP Addresses

- "Default IP Addresses"
- "Choosing DHCP or a Static IP Address"
- "Accessing the MAC Address in the CLI"

Choosing the appropriate IP addresses is essential to manage your VTrak E5000 Series subsystem over a network. You must change the IP addresses of the subsystems as required for your environment.

Default IP Addresses

The default **virtual** management port IP addresses are set to:

- IPv4 – 10.0.0.1
- IPv6 – 2001::1

The virtual management port IP address works with either RAID controller, enabling access to a dual-controller VTrak E5000 Series over your network using a single IP address.

For VTrak E5000 Series systems used as nodes on a VTrak A-Class managed SAN, the IP address of the virtual management port is used to identify the system. Therefore the IP address should be configured as a static IP address.

The default **physical** management port IP addresses are set to:

- Controller 1, IPv4 – 10.0.0.3
- Controller 1, IPv6 – 2001::3
- Controller 2, IPv4 – 10.0.0.5
- Controller 2, IPv6 – 2001::5

The physical management port IP address works with only one RAID controller. The port is used for management of the subsystem and when a controller is in maintenance mode [LINK to troubleshooting]. The management port is also used for Metadata communications when used as a storage node on a SAN managed by a VTrak A-Class.

Choosing DHCP or a Static IP Address

When you setup your VTrak E5000 Series, you have the option of:

- Enabling DHCP and letting your DHCP server assign the IP address to the VTrak E5000 Series's virtual management port.
- Specifying a static IP address for the VTrak E5000 Series's virtual management port.

DHCP is currently supported on IPv4 only. If you use IPv6, you must make your network settings manually.

If you choose to enable DHCP, have your Network Administrator dedicate an IP address for the VTrak E5000 Series, linked to the VTrak E5000 Series's MAC address. This action prevents the DHCP server from assigning a new IP address when the VTrak E5000 Series restarts, with the result that users can no longer log in.

However, *if you choose the DHCP option, it is important to understand the a controller fail-over **might** result in a different IP address being assigned to the subsystem*. With DHCP, there is no guarantee that the IP address will remain the same.

Important

Continuity of the sybsystem IP address is not guaranteed if you choose the DHCP option for IP settings.

Important

Use a static IP address for VTraks that are included as nodes in a VTrak A-Class managed SAN.

Accessing the MAC Address in the CLI

To access the MAC address in the CLI:

At the command prompt, type `net -a list -v` and press **Enter**.

The following information displays:

```
administrator@cli> net -a list -v
------------------------------------------------------------------------------
ActiveCtrlId: 1   Port: 1
MaxSupportedSpeed: 1000Mbps LinkStatus: Up
ProtocolFamily: IPv4(Enabled)    DHCP: Disabled
IP: 10.0.0.1
IPMask: 0.0.0.0
MAC: 00:01:55:61:18:65
DNS: 0.0.0.0
Gateway: 0.0.0.0

ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: ::
IPMask: ::
MAC: 00:01:55:61:18:65
DNS: ::
Gateway: ::
```

Setting-up with the CLI

Setting up the VTrak E5000 Series in the CLI includes these actions:

- "Making Subsystem Date and Time Settings"
- "Virtual Management Port Settings"
    - "Making Virtual Management Port Settings – Automatically"
    - "Making Virtual Management Port Settings – Manually under IPv4"
    - "Making Virtual Management Port Settings – Manually under IPv6"

Making Subsystem Date and Time Settings

To set the subsystem date and time:

1. Type **date -a mod -d** and the date in yyyyy/mm/dd format then press **Enter**.

```
administrator@cli> date -a mod -d 2016/03/25
```

2. Type **date -a mod -t** and the time in hh:mm:ss format, then press **Enter**.

```
administrator@cli> date -a mod -t 14:50:05
```

You can combine date and time settings, such as:

```
administrator@cli> date -a mod -d 2016/03/25 -t 14:50:05
```

Virtual Management Port Settings

Making Virtual Management Port Settings – Automatically

Automatic settings require a DHCP server on your network. DHCP is currently supported on IPv4 only.

To enable automatic management port settings:

1. At the command prompt, type **net -a mod -f ipv4 -s "dhcp=enable"** and press **Enter**.

```
administrator@cli> net -a mod -f ipv4 -s "dhcp=enable"
```

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

2. To verify the setting change, at the command prompt, type **net** and press **Enter**. The following information displays:

```
administrator@cli> net
=========================================
PF   Status    IP              Link
=========================================
IPv4 Enabled 10.0.0.1          Up
IPv6 Disabled ::    Up
```

In the above example:

- PF refers to IP protocol family, v4 or v6
- Status refers to whether the IP protocol is enabled. IPv4 is enabled by default.
- IP is the virtual management port IP address.
- Link indicates whether there is a working network connection.

By default, IPv4 is enabled and IPv6 is disabled. Currently IPv6 does not support DHCP.

Important

Use a static IP address for VTraks that are included as nodes in a VTrak A-Class managed SAN

Making Virtual Management Port Settings – Manually under IPv4

To make IPv4 settings manually on the management port:

1. At the command prompt, type **net -a mod -f ipv4 -s** " followed by:
   - **primaryip=** and the IP address **,**
   - **primaryipmask=** and the subnet mask **,**
   - **primarydns=** and the DNS server IP address **,**
   - **gateway=** and the Gateway server IP address **"** and press **Enter**.

   Example:

   ```
   administrator@cli> net -a mod -f ipv4 -s "primaryip=10.0.0.1,
   primaryipmask=255.255.255.0,primarydns=10.0.0.11,gateway=10.0.0.1"
   ```

After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

2. To verify the settings, at the command prompt, type net -a list -v and press **Enter**.

   The following information displays:

```
administrator@cli> net -a list -v
------------------------------------------------------------------------------
ActiveCtrlId: 1   Port: 1
MaxSupportedSpeed: 1000Mbps LinkStatus: Up
ProtocolFamily: IPv4(Enabled)    DHCP: Disabled
IP: 10.0.0.1
IPMask: 255.255.255.0
MAC: 00:01:55:61:18:65
DNS: 10.0.0.11
Gateway: 10.0.0.1

ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: ::
IPMask: ::
MAC: 00:01:55:61:18:65
DNS: ::
Gateway: ::
```

Making Virtual Management Port Settings – Manually under IPv6

To make IPv6 settings manually on the management port:

1. At the command prompt, type **net -a enable -f ipv6** and press Enter to enable IPv6 on the VTrak E5000 Series.

   After a moment, the command prompt reappears, indicating that your setting was successful.

   ```
   administrator@cli>
   ```

2. At the command prompt, type net -a mod -f ipv6 -s " followed by:
   - **primaryip=** and the IP address **,**
   - **primaryipmask=** and the subnet mask **,**
   - **primarydns=** and the DNS server IP address **,**
   - **gateway=** and the Gateway server IP address **"** and press **Enter**.Example:

   ```
   administrator@cli> net -a mod -f ipv6 -s
   "primaryip=2001:0db8:85a3:0000:0000:8a2e:0370:7334,
   primaryipmask=2001:0db8:fedc:ba98:7654:3210:0246:8acf
   primarydns=2001:0db8:85a3:0000:0000:8a2e:0370:7001,
   gateway=2001:0db8:85a3:0000:0000:8a2e:0370:7002"
   ```

   After a moment, the comand prompt reappears, indicating that your setting was successful.

   ```
   administrator@cli>
   ```

3. To verify the settings, at the command prompt, type net -a list -v and press **Enter**.

   The following information displays:

   ```
   administrator@cli> net -a list -v
   ------------------------------------------------------------------------
   ActiveCtrlId: 1       Port: 1
   MaxSupportedSpeed: 1000Mbps LinkStatus: Up
   ProtocolFamily: IPv4(Enabled)    DHCP: Disabled
   IP: 10.0.0.1
   IPMask: 255.255.255.0
   MAC: 00:01:55:61:18:65
   DNS: 10.0.0.11
   Gateway: 10.0.0.1

   ProtocolFamily: IPv6(Enabled)    DHCP: Disabled
   IP: 2001:0db8:85a3:0000:0000:8a2e:0370:7334
   IPMask: 2001:0db8:fedc:ba98:7654:3210:0246:8acf
   MAC: 00:01:55:61:18:65
   DNS: 2001:0db8:85a3:0000:0000:8a2e:0370:7001
   Gateway: 2001:0db8:85a3:0000:0000:8a2e:0370:7002
   ```

Maintenance Mode Settings

You also have the option to make maintenance mode settings at a later time in WebPRM PROe. The IP address of the management port can be configured to use a different IP address when a controller is in maintenance mode. Maintenance mode is used in the event of a controller failure, or if there is a difference of some kind between two controller on a dual controller subsystem. In maintenance mode, the VTrak goes offline and displays N/A under Readiness Status in the Component List in the Device menu tab. This circumstance requires intervention by the administrator.

Making Maintenance Mode Settings – Automatically

Automatic settings require a DHCP server on your network. DHCP is currently supported on IPv4 only.

You make maintenance mode settings for one controller at a time.

To enable automatic maintenance mode settings:

1. At the command prompt, type `net -a mod -m -c 1 -f ipv4 -s "dhcp=enable"` and press **Enter**.

   ```
   administrator@cli> net -a mod -m -c 1 -f ipv4 -s "dhcp=enable"
   ```

   After a moment, the comand prompt reappears, indicating that your setting was successful.

   ```
   administrator@cli>
   ```

2. To verify the settings changes, at the command prompt, type net -a list -m and press **Enter**.

   The following information displays:

```
administrator@cli> net -a list -m
------------------------------------------------------------------------------
CtrlId: 1    Port: 1
ProtocolFamily: IPv4(Enabled)    DHCP: Enabled
IP: N/A
IPMask: N/A
MAC: 00:01:55:30:65:E9
DNS: N/A
Gateway: N/A

CtrlId: 1    Port: 1
ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: 2001::2
IPMask: ffff::
MAC: 00:01:55:30:65:E9
DNS: ::
Gateway: ::

CtrlId: 2    Port: 1
ProtocolFamily: IPv4(Enabled)    DHCP: Disabled
IP: 10.0.0.3
IPMask: 255.0.0.0
MAC: 00:01:55:30:65:D7
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 2    Port: 1
ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: 2001::3
IPMask: ffff::
MAC: 00:01:55:30:65:D7
DNS: ::
Gateway: ::
```

3. Repeat steps 1 and 2 above but change -c 1 (controller 1) to -c 2 (contoller 2).

Making Maintenance Mode Settings – Manually under IPv4

You make these settings for one controller at a time.

To make maintenance mode settings:

1. At the command prompt, type net -a mod -m -c 1 -s " followed by:
   - **primaryip=** and the IP address **,**
   - **primaryipmask=** and the subnet mask **,**
   - **primarydns=** and the DNS server IP address **,**
   - **gateway=** and the Gateway server IP address **"** and press **Enter**.

   Example:

```
administrator@cli> net -a mod -m -c 1 "primaryip=10.0.0.101, primaryipmask=2
55.255.255.0,primarydns=10.0.0.11,gateway=10.0.0.1"
```

   After a moment, the comand prompt reappears, indicating that your setting was successful.

```
administrator@cli>
```

2. To verify the settings changes, at the command prompt, type `net -a list -m` and press **Enter**. The following information displays:

```
administrator@cli> net -a list -m
------------------------------------------------------------------------------
CtrlId: 1               Port: 1
ProtocolFamily: IPv4(Enabled) DHCP: Disabled
IP: 10.0.0.2
IPMask: 255.0.0.0
MAC: 00:01:55:30:65:E9
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 1               Port: 1
ProtocolFamily: IPv6(Disabled)     DHCP: Disabled
IP: 2001::2
IPMask: ffff::
MAC: 00:01:55:30:65:E9
DNS: ::
Gateway: ::

CtrlId: 2               Port: 1
ProtocolFamily: IPv4(Enabled) DHCP: Disabled
IP: 10.0.0.3
IPMask: 0.0.0.0
MAC: 00:01:55:30:65:E9
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 2               Port: 1
ProtocolFamily: IPv6(Disabled)     DHCP: Disabled
IP: 2001::3
IPMask: ffff::
MAC: 00:01:55:30:65:D7
DNS: ::
Gateway: ::
```

3. Repeat steps 1 and 2 above but change -c 1 (controller 1) to -c 2 (controller 2).

Making Maintenance Mode Settings – Manually under IPv6

You make these settings for one controller at a time.

To make maintenance mode settings:

1. At the command prompt, type net -a enable -f ipv6 -m -c 1 and press **Enter** to enable IPv6.

   After a moment, the comand prompt reappears, indicating that your setting was successful.

   ```
   administrator@cli>
   ```

2. At the command prompt, type net -a mod -m -c 1 -s “ followed by:

   - **primaryip=** and the IP address **,**
   - **primaryipmask=** and the subnet mask **,**
   - **primarydns=** and the DNS server IP address **,**
   - **gateway=** and the Gateway server IP address **“** and press **Enter**.

   Example:

   ```
   administrator@cli> iscsi -a mod -t portal -s
   “primaryip=2001:0db8:85a3:0000:0000:8a2e:0370:7336, primaryipmask=2
   001:0db8:fedc:ba98:7654:3210:0246:8acf,
   primarydns=2001:0db8:85a3:0000:0000:8a2e:0370:7001,
   gateway=2001:0db8:85a3:0000:0000:8a2e:0370:7002”
   ```

   After a moment, the comand prompt reappears, indicating that your setting was successful.

   ```
   administrator@cli>
   ```

1. To verify the settings, at the command prompt, type net -a list -m and press **Enter**.

   The following information displays:

```
administrator@cli> net -a list -m

------------------------------------------------------------------------
CtrlId: 1             Port: 1
ProtocolFamily: IPv4(Enabled)    DHCP: Disabled
IP: 10.0.0.2
IPMask: 255.0.0.0
MAC: 00:01:55:30:65:E9
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 1             Port: 1
ProtocolFamily: IPv6(Enabled)    DHCP: Disabled
IP: 2001:0db8:85a3:0000:0000:8a2e:0370:7336
IPMask: 001:0db8:fedc:ba98:7654:3210:0246:8acf
MAC: 00:01:55:30:65:E9
DNS: 2001:0db8:85a3:0000:0000:8a2e:0370:7001
Gateway: 2001:0db8:85a3:0000:0000:8a2e:0370:7002

CtrlId: 2   Port: 1
ProtocolFamily: IPv4(Enabled)    DHCP: Disabled
IP: 10.0.0.3
IPMask: 0.0.0.0
MAC: 00:01:55:30:65:E9
DNS: 0.0.0.0
Gateway: 0.0.0.0

CtrlId: 2             Port: 1
ProtocolFamily: IPv6(Disabled)   DHCP: Disabled
IP: 2001::3
IPMask: ffff::
MAC: 00:01:55:30:65:D7
DNS: ::
Gateway: ::
```

2. Repeat steps 1, 2, and 3 above but change **-c 1** (controller 1) to **-c 2** (controller 2).

This completes management port and maintenance mode setup.

SETTING UP WITH WEBPAM PROE

Logging into WebPAM PROe

1. Launch your browser.
2. In the browser address field, type in the virtual management port IP address of the VTrak E5000 Series subsystem.

   Use the default virtual management port IP address; or the virtual IP address you set in the CLI. See "Virtual Management Port Settings" on page 57 for IP setting instructions to specify a preferred IP address if you do not intend to use the default IP settings.

   Example:
   - WebPAM PROe uses a secure HTTP connection https://
   - **Enter** the IP address of the VTrak E5000 Series 10.0.0.1

   Together, your entry looks like this: **https://10.0.0.1**
3. When the log-in screen appears:
   - Type **administrator** in the User Name field.
   - Type **password** in the Password field.

   The User Name and Password are case sensitive.
4. Optional. Choose a display language from the drop-down menu.

   WebPAM PROe displays in *English*, *German*, *French*, *Italian*, *Spanish*, *Russian*, *Japanese*, *Traditional Chinese*, *Simplified Chinese*, and *Korean*.
5. Click the **Login** button.

WebPAM PROe log-in screen and language options

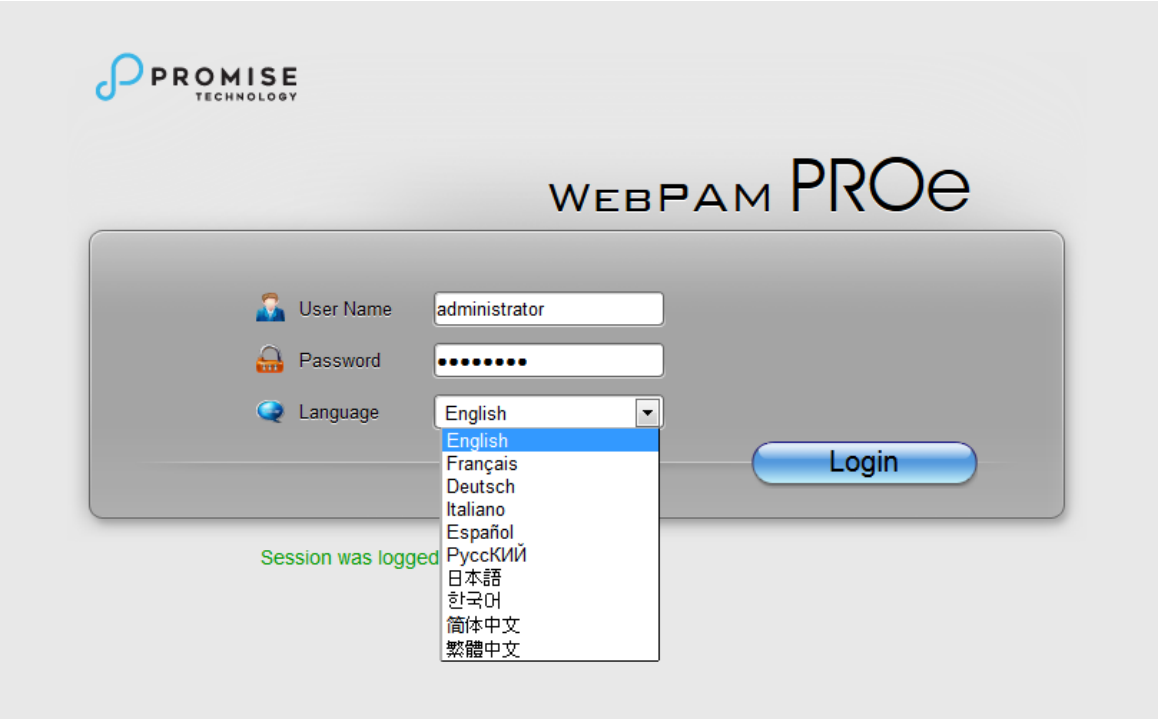

Important

PROMISE recommends that you change the Administrator's default password immediately after setup is completed.

See "Changing User Passwords".

Note

All RJ-45 network management ports on a VTrak E5000 Series subsystem share the same Virtual IP address. The default Virtual IP address, 10.0.0.1, applies to all RJ-45 network ports. If you change the Virtual IP address, the change applies to all RJ-45 ports.

After log-in, the WebPAM PROe opens with the Dashboard tab.

WebPAM PROe Dashboard tab

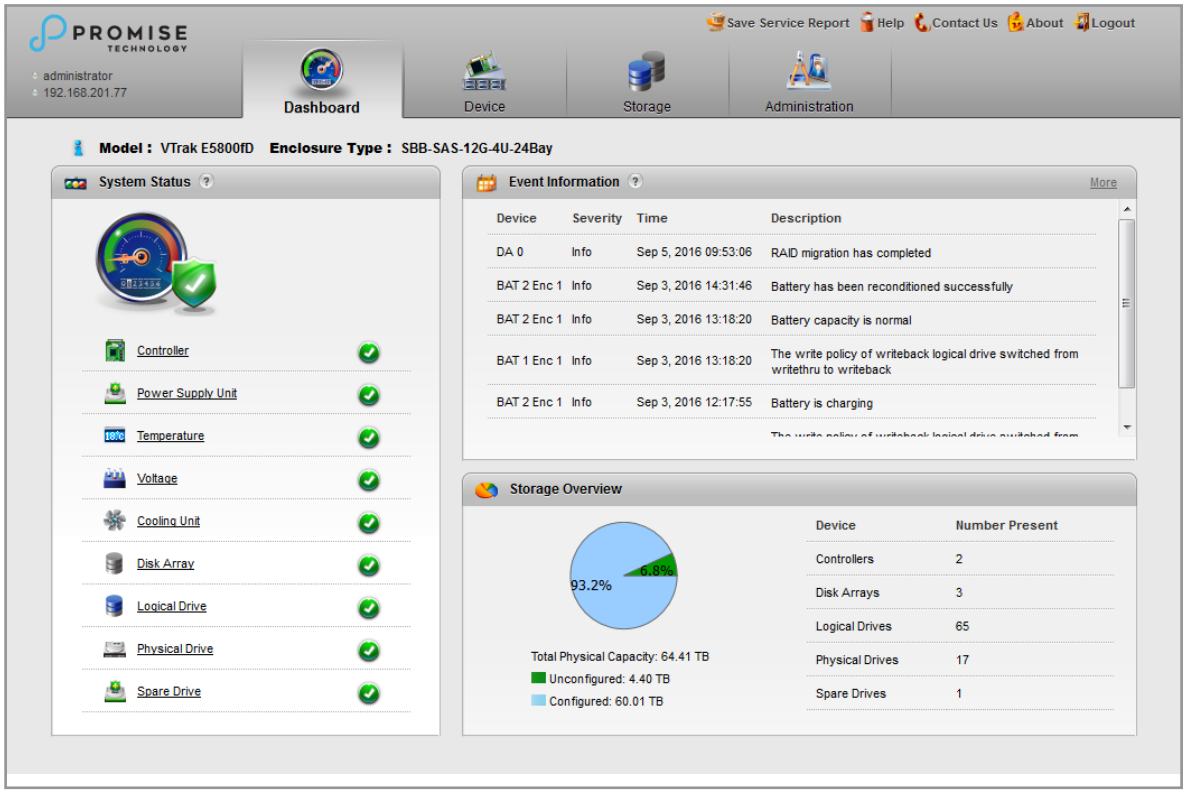

Creating Logical Drives

On a newly activated RAID system, there are no disk arrays or logical drives. The term “disk array” includes arrays composed of hard disk drives or solid state drives.

To create your disk arrays and logical drives:

1. Click the Storage tab, then click the *Wizard* option.

   Or, click Disk Array under System Status.

   The Wizard screen appears with three creation alternatives:

   - Optimal
   - Advanced

2. Click one of these configuration option buttons to continue.

The Disk Configuration Wizard main menu

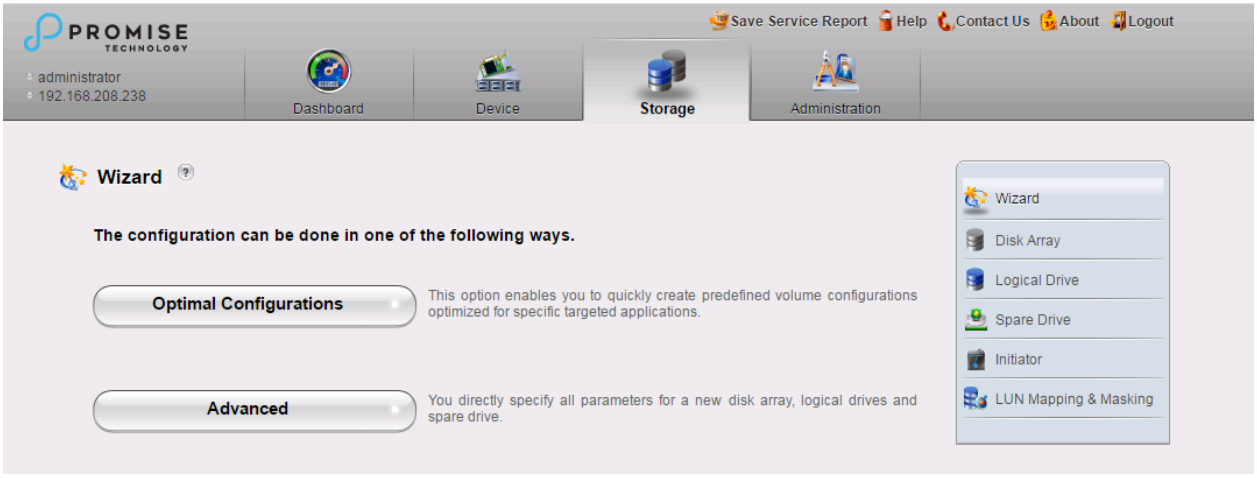

Optimal Configuration

When you choose the Optimal Configurations option, you choose a script de-signed to set up your disk arrays, logical drives, and spare drives for a specific target application.

Each script requires a specific model of RAID subsystem. And most scripts re-quire a specific model and number of JBOD expansion units. You cannot modify these scripts.

Choose an optimal configuration script

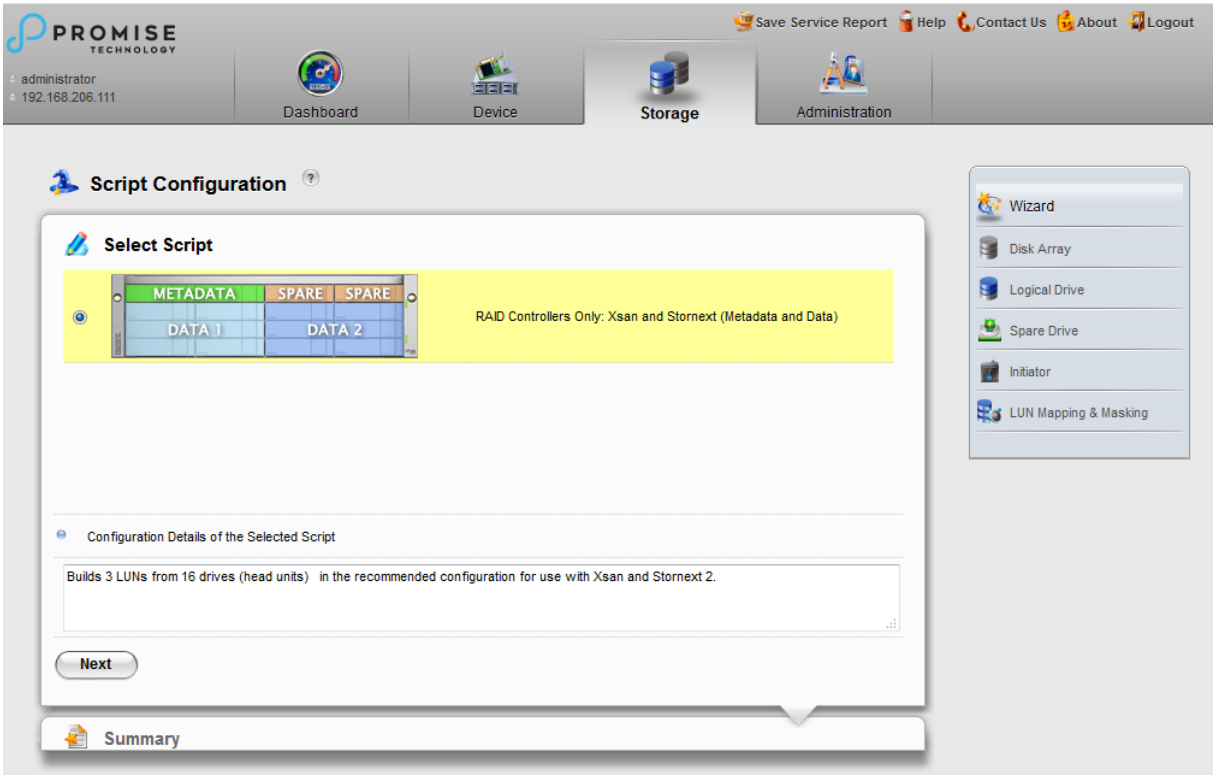

Summary of optimal configuration

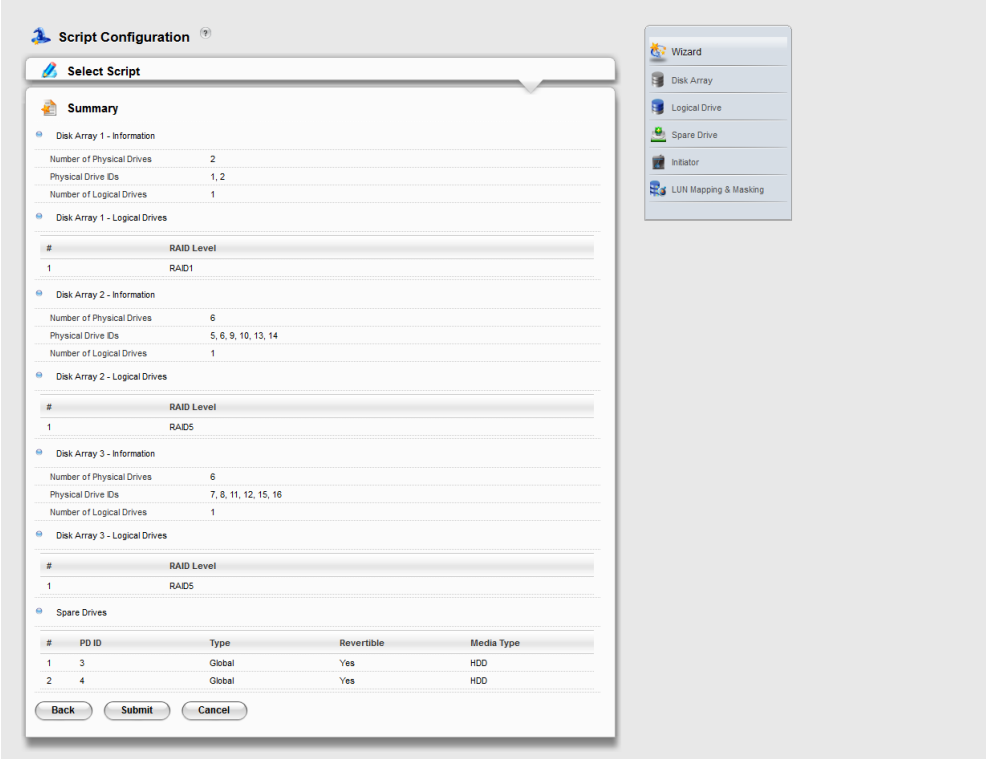

Type 'confirm' to verify and implement optimal configuration

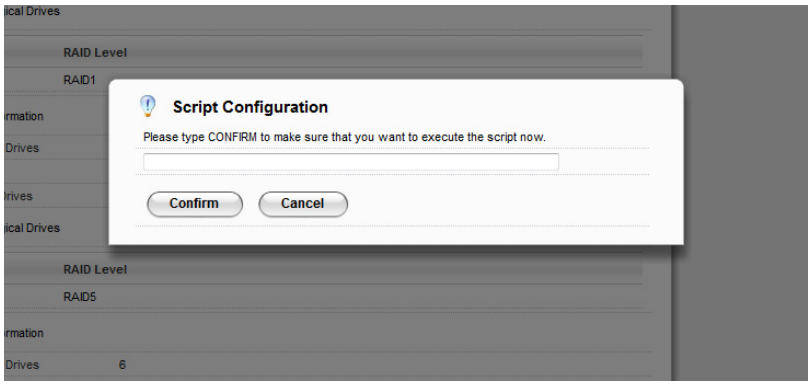

Advanced Configuration

When you choose the **Advanced** option, the **Create Disk Array** menu appears.

Step 1 – Disk Array Creation

1. **Enter** your information and choose your options.
   - **Enter** a disk array alias in the field provided.
   - Check the box to enable Media Patrol
   - Check the box to enable Predictive Data Migration (PDM)
   - Check the box to enable Power Management
   - Choose a media type – Hard disk drive (HDD) or solid state drive (SSD)
2. Click the enclosure graphic to view information about physical drives.

   Look for drives with a green LED dark, a blue LED lit, and no crosshatching over the carrier.
3. Click a physical drive to select it for your array.

   The physical drive's ID number is added to the Selected list.
4. Click the **Next** button to continue.

   The **Create Logical Drive** screen appears.

Step 2 – Logical Drive Creation

1. **Enter** your information and choose your options.
   - **Enter** a logical drive alias in the field provided.
   - Choose a RAID level from the drop-down menu.
     The choice of RAID levels depends on the number of physical drives in your array.
   - Note the **Max:** capacity value. Then enter a capacity value the field provided and choose a unit of measure from the drop-down menu.
   - Choose a stripe size from the drop-down menu.
     The choices are 64 KB, 128 KB, 256 KB, 512 KB, and 1 MB.
   - Choose a sector size from the drop-down menu.
     The choices are 512 B, 1 KB, 2 KB, and 4 KB.
   - Choose the Read Cache Policy from the drop-down menu
     The choices are Read Cache, Read Ahead (cache), and None.
   - Choose the Write Cache Policy from the drop-down menu - The choices are WriteThru (write through) and WriteBack. Write back requires a Read Cache or Read Ahead Read Cache Policy.
   - Uncheck the Perfect Rebuild check box if do not need perfect rebuild for this LD.
   - Click the **Add** button to continue.

   The logical drive you just created appears in the **New Logical Drives** list.

2. Click the **Next** button to continue.

   The **Create Spare Drive** screen appears.

Step 3 – Spare Drive Creation

Creating a spare drive is optional but highly recommended.

1. **Enter** your information and choose your options.
   - Check the **Revertible** box if you want this spare drive to be revertible.
   - Choose the option for the type spare drive you want.
     Global – Replaces a failed drive in any disk ar ray.
     Dedicated – Replaces the failed drive only in the assigned disk array.
2. Click the enclosure graphic to view information about physical drives.
3. Click a physical drive to select it for your spare drive.

   The physical drive's ID number is added to the Selected list.
4. Click the **Next** button to continue.

Step 4 – Summary

The Summary screen lists the disk arrays, logical drives, and spare drives that you specified.

If you accept these parameters, click the **Submit** button.

If you do NOT accept these parameters, review and modify your selections in the previous steps.

Enabling LUN Mapping and Masking

These features are optional for each logical drive. The Enable LUN Mapping dialog box appears after you create a logical drive.

To enable LUN Mapping:

1. Click the **OK** button in the **Enable LUN Mapping** dialog box.

   The **LUN Mapping & Masking** screen appears.
2. Check the **Enable LUN Masking** box to enable LUN Masking.
3. Click the **LUN Mapping** button to continue.

   The initiator list screen displays.
4. Choose the initiators you want to use from the drop-down menu and click the **Next** button.

   The screen displays a list of initiators and a list of logical drives.
5. Click and drag a logical drive from the logical drives list to the initiators list.
6. Click the **Next** button when you are done.

   The screen displays a list of initiator IDs and corresponding LUN maps that you specified.
7. Click the **Submit** button to create the LUN map.

   The screen displays a list of initiator IDs and corresponding LUN maps.

You can also set LUN mapping and masking at a later time. Click the **Administration** tab, then click the **LUN Mapping & Masking** option.

Logging out of WebPAM PROe

There are two ways to log out of WebPAM PROe:

- Close your browser window
- Click **Logout** on the WebPAM PROe banner

Clicking **Logout** brings you back to the Login Screen. See WebPAM PROe log-in screen with display language options.

After logging out, you must enter your user name and password in order to log in again.

Using WebPAM PROe over the Internet

The above instructions cover connections between VTrak E5000 Series and your company network. It is also possible to connect to a VTrak E5000 Series from the Internet.

Your MIS Administrator can tell you how to access your network from outside the firewall. Once you are logged onto the network, you can access the VTrak E5000 Series using its IP address.

WebPAM PROe - System Configuration

This chapter describes system configuration using WebPAM PROe. It contains the following topics:

Logging into WebPAM PROe

1. Launch your browser.
2. In the browser address field, type in the virtual management port IP address of the VTrak E5000 subsystem.

   Use the IP address you set in the CLI (page 37).

   Example:

   - WebPAM PROe uses a secure HTTP connection.https://
   - Enter the IP address of the VTrak E5000 .....

   For example, if your VTrak E5000 has an IP address: 10.0.0.1 your entry looks like this:

   https://10.0.0.1
3. When the login screen appears:
   - Type **administrator** in the User Name field.
   - Type **password** in the Password field.
   - Click the **Login** button.

   The User Name and Password are case sensitive.
4. Optional. Choose a display language from the drop-down menu.

   WebPAM PROe displays in English, German, French, Italian, Spanish, Russian, Japanese, Traditional Chinese, Simplified Chinese, and Korean.
5. Click the **Login** button.

   After login, the WebPAM PROe main menu appears.

Choosing the Display Language

WebPAM PROe displays in multiple languages. You choose the display language when you log in.

If you are already logged in and you want to change the display language:

1. Click **Logout** at the top right corner of the screen.

   The Login screen appears.

2. Click the Language drop-down menu and highlight the language you prefer.

Login language selection menu

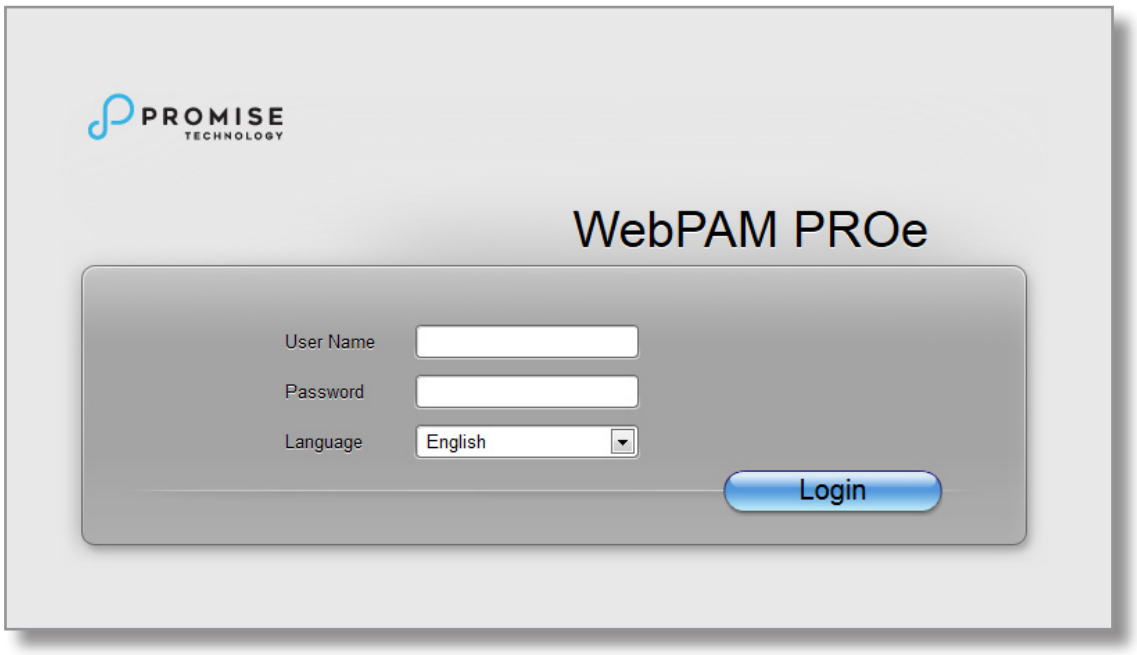

3. Reenter your user name and password.
4. Click the **Login** button.

   WebPAM PROe opens in the language you chose.

Perusing the Interface

The WebPAM PROe interface consists of a header and four tabs, each with specific functions.

- Header

  Top left corner of the window:

  - Name of logged-in user
  - IP address – Virtual IP address of the RAID subsystem

- Top right corner of the window
  - Save Service Report – Saves a detailed report to your Host PC
  - Help – Accesses the Help Welcome screen
  - Contact Us – Technical support contact information
  - About – Information about WebPAM PROe
  - Logout – Exits WebPAM PROe

- Discovery tab
  - Displays other PROMISE RAID systems on your network
  - Enables direct login to other PROMISE RAID systems

List continues on next page

- Dashboard tab
  - RAID subsystem model and type of enclosure
  - System status
  - Event information – Most recent NVRAM events
  - Storage overview – Capacities, number of devices
- Device tab
  - Enclosure front and back views
  - Topology
  - Enclosure component list and settings
  - Physical drive management
  - UPS (Uninterruptable Power Supply) management
  - Fibre Channel management
- Storage tab
  - Wizard – Automatic or Advanced configuration
  - Disk array management
  - Logical drive management
  - Initiator management
  - LUN mapping and masking
- Administration tab
  - Subsystem settings, clearing statistics, NTP, and controller lock
  - User management, including LDAP and role mapping
  - Software services
  - Runtime and NVRAM event logs
  - Background activity, settings and schedules
  - Firmware updates
  - Image version
  - Performance monitor
  - PSU wattage monitor
  - Restore factory default settings
  - Import/Export user database and configuration script
  - Network management

Web PAM PROe Main menu/Dashboard

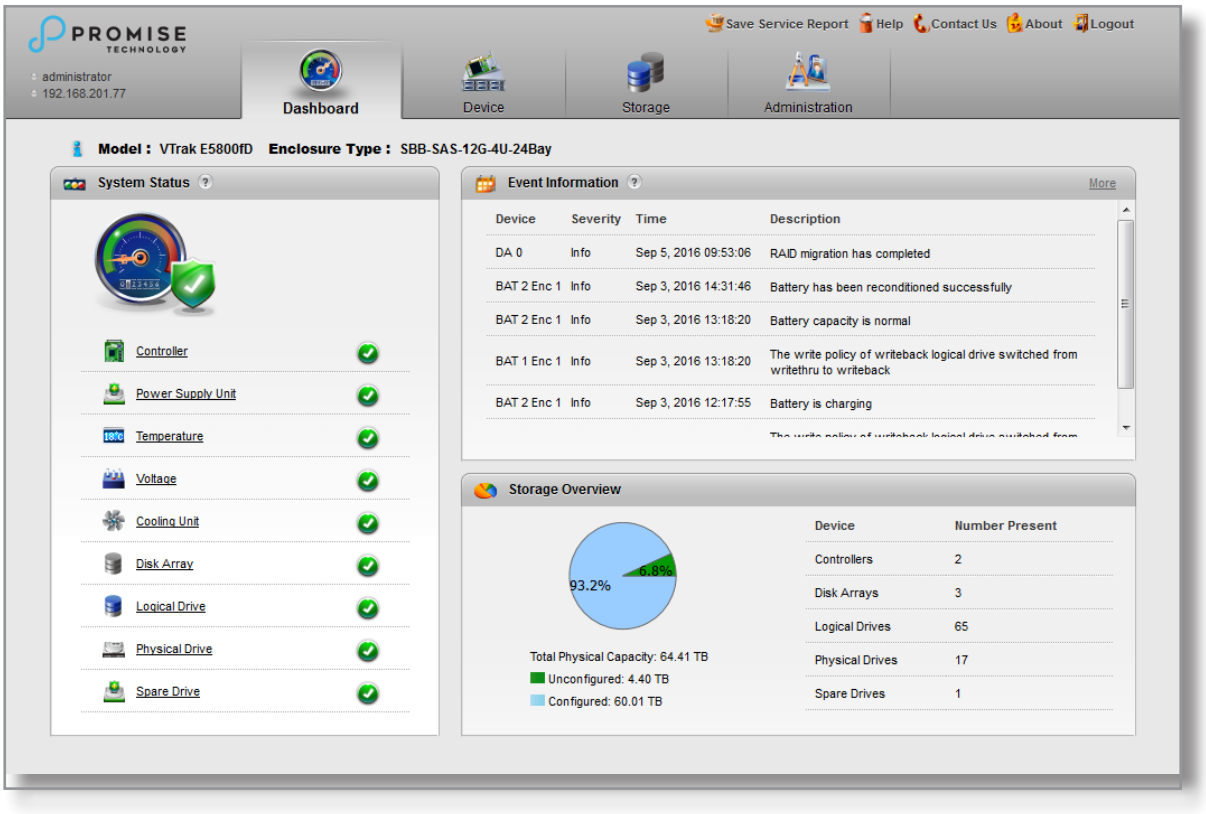

Logging out of WebPAM PROe

There are two ways to log out of WebPAM PROe:

- Close your browser window
- Click **Logout** on the WebPAM PROe banner

Clicking Logout brings you back to the Login Screen.

After logging out, you must enter your user name and password in order to log in again.

Viewing the Storage Network

To view the other subsystems on your Storage Network, click the **Discovery** tab at the left edge of the WebPAM PROe window.

Discovery tab in Main menu

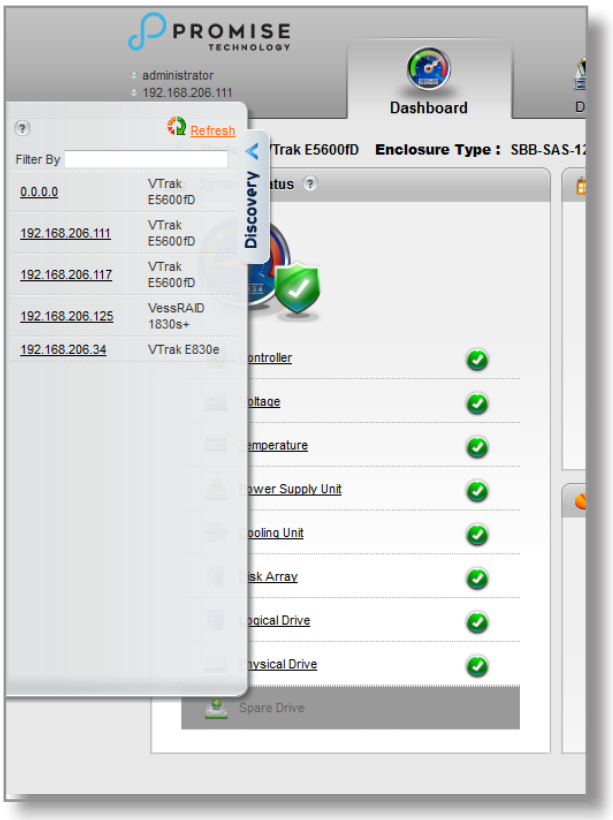

Logging onto a Subsystem

To log onto a subsystem in the list, double-click the subsystem.

Caution

The new subsystem displays in the same browser tab. Click your browser's back button to return to the original subsystem.

Filtering the Subsystem List

To filter the list, so it shows only specific subsystems, enter a characteristic into the Filter By field and press Enter. To filter by IP address, enter the IP routing prefix for the range you want to display. For example, typing “10.0” in the entry field reveals all subsystems with IP address beginning with the “10.0” prefix.

Refreshing the List

To refresh the list, click the **Refresh** link.

Managing Subsystems

Subsystem management includes:

- “Viewing Subsystem Information”
- “Making Subsystem Settings”
- “Locking or Unlocking the Subsystem”
- “Restoring Factory Default Settings”
- “Clearing Statistics”
- “Saving a Service Report”
- “Importing a Configuration Script”
- “Exporting a Configuration Script”
- “Restarting the Subsystem”
- “Shutting Down the Subsystem”
- “Restarting the Subsystem after a Shutdown”

Viewing Subsystem Information

To view subsystem information, click the **Administration** tab.

The list of subsystems and host controllers displays.

Subsystem information includes:

• Alias, if assigned	• Vendor
• Model	• WWN – World Wide Name
• Serial number	• Part number
• Revision number	• Cache Mirroring (*Status*)
• Number of connected JBOD enclosures	• Maximum number of JBOD expansion units supported
• Number of controllers present	• Maximum number of controllers supported
• Redundancy status	• Redundancy type
• System date and time	

Making Subsystem Settings

To make subsystem settings:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Settings** button.
4. Make changes as required:
   - Enter an alias or change the existing alias in the field provided.
   - Choose a redundancy type from the drop-down menu.
     The choices are **Active-Active** and **Active-Standby**
   - Check the box to enable **Cache Mirroring** (or uncheck to disable)
5. Click the **Save** button.

Locking or Unlocking the Subsystem

The lock prevents other sessions (including sessions with the same user) from making a configuration change to the controller until the lock expires or a forced unlock is done. When the user who locked the controller logs out, the lock is automatically released.

Setting the Lock

To set the lock:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Lock / Unlock** button.
4. In the Lock Time field, type a lock time in minutes.

   1440 minutes = 24 hours
5. Click the **Lock** button.

Resetting the Lock

To reset the lock with a new time:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Lock / Unlock** button.
4. In the Lock Time field, type a new lock time in minutes.

   1440 minutes = 24 hours

5. Click the Lock button.

Releasing the Lock

To release a lock that you set:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Lock / Unlock** button.
4. Click the **Unlock** button.

Releasing a Lock set by another user

To release somebody else's lock:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Lock / Unlock** button.
4. Check the **Force Unlock** box.
5. Click the **Unlock** button.

Restoring Factory Default Settings

This feature restores settings to their default values.

Caution

Use this feature only when required and only on the settings that you must reset to default in order to set them correctly.

To restore all settings to their default values:

1. Click the **Administration** tab.
2. Click the **Restore Factory Default** icon.
3. In the Restore factory default settings screen, check the boxes beside the settings you want to reset to default value (see **Factory Default Settings (by type)** table below).
4. Click the **Submit** button.
5. In the Confirmation box, type the word “confirm” in the field provided and click the **Confirm** button.

Factory Default Settings (by type)

Firmware

- Background activity settings
- Controller settings
- Enclosure settings
- Fibre Channel port settings
- Management network settings
- Physical drive settings
- Subsystem settings

Software

- BGA scheduler settings
- Service settings
- Web Server Settings
- SNMP settings
- Telnet settings
- SSH settings
- Email settings
- Netsend settings
- CIM Settings
- NTP settings
- User settings
- UPS settings
- LDAP Settings

Clearing Statistics

This function clears statistical data on the RAID controllers, Fibre Channel ports, physical drives, and logical drives.

To clear subsystem statistics:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Click the **Clear Statistic**s button.
4. Type the word "**confirm**" in the field provided.
5. Click the **Confirm** button.

Saving a Service Report

A Service Report is a detailed report covering the configuration and status of all components in your RAID system. A support technician or field engineer might request a service report for the purpose of diagnosis and troubleshooting.

To save a Service Report file:

1. Click **Save Service Report** in the Header (very top of the web interface, next to the **Help** link).

   Information for the report is gathered and compiled. This action takes up to a few minutes, depending on the size of your RAID system

2. Click the **Save File** option, then click the **Save** button.

   The report saves to your Host PC as a compressed HTML file.

3. Double-click the downloaded file to decompress it.

4. Double-click the report to open it in your default browser.

   The Service Report topics are listed on the next page.

The Service Report includes the following topics:

- About – Report utility
- Battery Info – Cache backup batteries
- BBM Info – Bad Block Manager
- BGA Summary – Status and settings
- BGA Schedules – Scheduled activities
- Buzzer Info
- Controller Info
- Debug Syslog – Diagnostic information
- Disk Array Info – ID, alias, and capacities only
- Disk Array Dump Info – Diagnostic information
- Disk Array Verbose Info – All disk array information
- Enclosure Info
- Error Table Info – Read check, write check, and inconsistent blocks
- Event Info – NVRAM – List of NVRAM events
- Event Info – Runtime – List of Runtime events
- Fibre Channel Node Info
- Fibre Channel Device Info
- Fibre Channel Initiator Info
- Fibre Channel Port Info
- Fibre Channel SFP Info
- Fibre Channel Stats Info
- Flash Image Version Info
- LDAP Info
- LogDrive Info – Basic logical drive information
- LogDrive Dump Info – Diagnostic information
- Logical Drive Verbose Info – Full logical drive information
- Lunmap Info – LUN map type, LUN masking status, and LUN entries
- Network Info – Virtual port
- Network Maintenance Info – Maintenance mode ports
- Phydriv Info – Basic physical drive information
- Phydriv Verbose Info – Full physical drive information
- PD SMART Info – Physical drive ID, model, type, and SMART status
- PSU Wattage Info – Enclosure power consumption, power supply input and output, and power on time
- SWMGT Info – Software management
- Service Setting – Email
- Service Setting – Netsend
- Service Setting – NTP
- Service Setting – SLP
- Service Setting – SNMP
- Service Setting – SSH
- Service Setting – Telnet
- Sessions Info
- Spare Info – Basic spare drive information
- Spare Dump Info – Diagnostic information
- Spare Verbose Info – Full spare drive information
- Statistic Info
- Subsystem info
- UPS Info
- User Info

Importing a Configuration Script

You can write a CLI configuration script to automatically configure your VTrak E5000 subsystem. The script must be a plain, non-encrypted text file. From there, you can import the script from the Host PC and perform the configuration automatically.

Cautions

Do NOT attempt to write or modify a configuration script until you receive guidance from Technical Support.

Importing a configuration script overwrites the current settings on your VTrak E5000 subsystem.

Or you can save the configuration from one VTrak E5000 RAID subsystem, export it, and then import it to automatically configure your other VTrak E5000 RAID subsystems. To import a configuration script:

1. Click the **Administration** tab.
2. Click the **Import/Export** icon.
3. Click the **Import** option.
4. Choose **Configuration Script** from the **Type** drop-down menu.
5. Click the **Browse** button and navigate to the configuration script and click the **OK** button.
6. Click the **Next** button.

   The system verifies that the file is a valid configuration script and displays any errors or warnings.
7. Click the **Submit** button to continue.
8. In the **Confirmation** box, type the word "**confirm**" in the field provided and click the **Confirm** button.

   The configuration script is imported and applied automatically.

Exporting a Configuration Script

You can save the configuration from one VTrak E5000 RAID subsystem, export it, and then import it to automatically configure your other VTrak E5000 RAID subsystems.

To export a configuration script:

1. Click the **Administration** tab.
2. Click the **Import/Export** icon.
3. Click the **Export** option.
4. Choose **Configuration Script** from the **Type** drop-down menu.
5. Click the **Submit** button.
6. In the Open dialog box, click the **Save File** option, then click the **OK** button. Alternatively, if your browser settings are set to begin downloading immediately, there will be no 'Open' dialog box.

   The file is saved to your PC as "Configscript.txt".

Cautions

Do NOT attempt to write or modify a configuration script until you receive guidance from Technical Support.

Restarting the Subsystem

This function shuts down the subsystem and then restarts it.

To restart the subsystem:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Choose the option to apply the restart to the **Subsystem**, **Controller 1** *only* or **Controller 2** *only.*
4. Click the **Shutdown/Restart** button.
5. Click the **Restart** button.
6. Type the word “confirm” in the field provided.
7. Click the **Confirm** button.

   When the controller shuts down, your WebPAM PROe connection is lost.
8. Wait at least two minutes.
9. In your browser, click **Logout** in the WebPAM PROe Header, then log in again.

   If you cannot log in immediately, wait 30 seconds and try again.

Shutting Down the Subsystem

This function shuts down the RAID subsystem without restarting it.

To shutdown the subsystem:

1. Click the **Administration** tab.
2. Click the **Subsystem Information** icon.
3. Choose the option to apply the shutdown to the **Subsystem**, **Controller 1** *only* or **Controller 2** *only*
4. Click the **Shutdown/Restart** button.
5. Click the **Shutdown** button.
6. Type the word "**confirm**" in the field provided.
7. Click the **Confirm** button.

   When the controller shuts down, your WebPAM PROe connection is lost.

Important

If your RAID subsystem manages JBOD expansion units, you must follow the proper startup procedure.

Restarting the Subsystem after a Shutdown

To start the RAID subsystem:

1. Press the Power button on the front left side of the device being restarted.
2. Wait at least two minutes.
3. Open your browser and log into WebPAM PROe.

   If you cannot log in immediately, wait 30 seconds and try again.

MANAGING RAID CONTROLLERS

RAID controller management includes:

- "Viewing Controller Information"
- "Making Controller Settings"
- "Viewing Controller Statistics"
- "Locating a Controller"
- "Viewing the Flash Image Information"
- "Updating Firmware on a RAID Subsystem"
- "Viewing Battery Information"
- "Reconditioning a Battery"
- "Buzzer Settings"
- "Silencing the Buzzer"

Viewing Controller Information

To view controller information:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the controller you want, then click the **View** button.

   Controller information includes:

   - Controller ID
   - Alias – If assigned
   - Readiness Status
   - Power On Time
   - LUN Mapping method
   - Serial Number
   - WWN – World Wide Name
   - Dirty Cache Usage – Percentage
   - Boot Loader Build Date
   - Firmware Build Date
   - Software Build Date
   - Operational Status
   - SCSI Protocol Supported
   - Part Number
   - Hardware Revision
   - Cache Usage – Percentage
   - Boot Loader Version
   - Firmware Version
   - Software Version

4. Click the **Advanced Information** tab.

Advanced controller information includes:

- Slot 1 Memory Type
- Slot 2 Memory Type
- Slot 3 Memory Type
- Slot 4 Memory Type
- LUN Affinity *
- Controller Role
- Flash Size
- NVRAM Size
- Coercion *
- SMART *
- Write Back Cache Flush Interval *
- Adaptive Writeback Cache *
- Forced Read Ahead (cache) *
- Power Saving Standby Time *
- Cache Line Size
- Slot 1 Memory Size
- Slot 2 Memory Size
- Slot 3 Memory Size
- Slot 4 Memory Size
- ALUA *
- Flash Type
- NVRAM Type
- Preferred Cache Line Size
- Coercion Method *
- SMART Polling Interval *
- Enclosure Polling Interval *
- Host Cache Flushing *
- Power Saving Idle Time *
- Power Saving Stopped Time *
- SSD Trim Support *
- SSD SMART Polling Interval*
- SSD SMART Life Threshold*
- SSD SMART Remaining Block Threshold*

Items with an asterisk (*) are adjustable under Controller Settings.

Making Controller Settings

In a dual-controller RAID subsystem, settings made to one controller are applied to both controllers.

To make controller settings:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the controller you want, then click the **Settings** button.
4. Make settings changes as required:
    - Enter, change or delete the alias in the **Alias** field.
    - LUN Affinity – Check the box to enable or uncheck to disable.
      RAID controllers must be set to Active-Active.
    - **ALUA** – Check the box to enable or uncheck to disable.
      RAID controllers must be set to Active-Active. See “Making Subsystem Settings” on page 87 and “ALUA” .
    - **SMART Log** – Check the box to enable or uncheck to disable.
    - **SMART Polling Interval** – Enter a value into the field, 1 to 1440 minutes
    - **HDD Power Saving** – Choose time periods from the drop-down menus.
      After an HDD has been idle for the set period of time:
        * **Power Saving Idle Time** – Parks the read/write heads.
        * **Power Saving Standby Time** – Lowers disk rotation speed.
        * **Power Saving Stopped Time** – Spins down the disk (stops rotation).

Notes

Power Management must be enabled on the disk array for the HDD Power Saving settings to be effective. See “Making Disk Array Settings”

Power Management functions are limited to the features your HDDs actually support.

- **Coercion** – Check the box to enable or uncheck to disable.
- **Coercion Method** – Choose a method from the drop-down menu:
    * *GBTruncate*
    * *10GBTruncate*
    * *GrpRounding*
    * *TableRounding*
- **Write Back Cache Flush Interval** – Enter a value into the field, 1 to 12 seconds.
- **Enclosure Polling Interval** – Enter a value into the field, 15 to 255 seconds.
- **Adaptive Writeback Cache** – Check the box to enable or uncheck to disable. See “Adaptive Writeback Cache” on page 402.
- **Host Cache Flushing** – Check the box to enable or uncheck to disable. See “Host Cache Flushing” on page 403.
- **Forced Read Ahead** (cache) – Check the box to enable or uncheck to disable. See Forced Read-Ahead Cache.
- **SSD Trim Support** – Check the box to enable or uncheck to disable.
- **SSD SMART Polling Interval** – Used to set the time interval in number of hours to poll the drive SMART status. Enter a value into the field 0 to 255 seconds.
- **SSD SMART Life Threshold** – Used to set the life threshold in percentage. Enter a value into the field 5 to 25 seconds. Default is 10.
- **SSD SMART Remaining Block Threshold** – Used to set the remaining reserved block threshold in percentage. Enter a value into the field 5 to 25 seconds. Default is 10.

5. Click the **Save** button.

Viewing Controller Statistics

To view controller statistics:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the controller you want, then click the **View** button.
4. Click the **Statistics** tab.

Controller statistics include:

- Data Transferred
- Read Data Transferred
- Errors
- Read Errors
- I/O Requests
- Read IO Requests
- Statistics Start Time
- Write Data Transferred
- Non-Read/Write Errors
- Write Errors
- Non-Read/Write Requests
- Write I/O Requests
- Statistics Collection Time

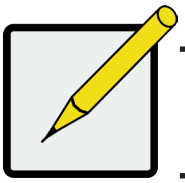

Note

To clear controller statistics, see "Clearing Statistics" on page 90.

Locating a Controller

This feature causes the controller LEDs to blink for one minute to assist you in locating the controller on a RAID subsystem or JBOD expansion unit.

To locate a controller:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the controller you want, then click the **Locate** button.

   The controller status LEDs blink for one minute.

Viewing the Flash Image Information

To view the flash image information for the RAID subsystem enclosure:

1. Click the **Administration** tab.
2. Click the **Image Version** icon.
3. Click the Enclosure you want to See and click the triangular button.

   RAID subsystems have the following components in their flash image:

   - Kernel
   - Firmware
   - Software
   - Ramdisk
   - OEM Customization
   - BIOS
   - 12G Expander
   - System Libraries
   - Applications
   - Mount Scripts
   - SAS Expander

   - **Running** – The version that is currently running on the subsystem or expansion unit.
   - **Flashed** – This version was updated but does not run until the subsystem restarts.

   See”Updating Firmware on a RAID Subsystem” on page 104

JBOD expansion units have only one component in their flash image, SEP firmware. It only appears as running.

Updating Firmware on a RAID Subsystem

Use this function to flash (update) the firmware on the VTrak E5000.

Download the latest firmware image file from PROMISE support:

http://www.promise.com/support/ and save it to your Host PC or TFTP server.

Important

Verify that no background activities are running on the RAID subsystem.

To update the firmware on the RAID subsystem and JBOD expansion units:

1. Click the **Administration** tab.
2. Click the Firmware **Update** icon.
3. Click the **Controller Firmware Update** tab.

   The Controller Firmware Update screen appears showing the current Image Version Number and Build Date.
4. Choose a download option:
   - **Local File through HTTP** – Click the **Browse** button, locate the firmware image file, click the file to choose it, then click the **Open** button.
   - **TFTP Server** – Enter the TFTP Server host name or IP address, port number and file name.
5. Optional. Check the Non-disruptive Image Update (NDIU) box.

   NDIU updates the RAID controllers and I/O modules one at a time, enabling I/O operations continue during the firmware update. Updates with this option take a longer period of time to complete. Only dual controller models support this feature.
6. Click the **Next** button.

   The next screen shows the Flash Image (firmware image file) Version Number and Build Date.
7. Click the **Submit** button.

   The progress of the update displays.

Warning

Do NOT power off the RAID subsystem during the update!

Do NOT move to any other screen until the firmware update operation is completed!

When the update is completed a message tells you to reboot the subsystem,

8. Click the **OK** button.

   - If you chose the Disruptive Flash Method, the RAID subsystem and JBOD expansion units automatically restart.
   - If you chose the Non-Disruptive Flash Method, the system automatically flashes and restarts the RAID controllers one at a time.

Automatic Restart

If you did NOT check the NDIU box, the RAID subsystem and JBOD expansion units automatically restart. That action temporarily disrupts I/O operations and drops your WebPAM PROe connection.

To reestablish your WebPAM PROe connection:

1. Wait no less than two minutes.
2. Click **Logout** in the WebPAM PROe Header, then log in again.

   If you cannot log in, wait 30 seconds and try again.
3. In your browser, click Logout in the WebPAM PROe Header, then log in again.

   If you cannot log in immediately, wait 30 seconds and try again.

Viewing Battery Information

Batteries maintain power to the controller cache in the event of a power failure, thus protecting any data that has not been written to a physical drive.

To view battery information:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the battery you want, then click the **View** button.

Battery information includes:

- (Battery) ID
- Operational status – Fully charged, recondition means a reconditioning is in process
- Battery chemistry – LiON, etc.
- Remaining capacity – Battery capacity as a percentage
- Battery cell type – Number of cells
- Estimated hold time – Time in hours that the battery can power the cache
- Temperature threshold discharge – Maximum temperature allowed when the battery is discharging
- Temperature threshold charge – Maximum temperature allowed when the battery is charging
- Battery temperature – Actual battery temperature
- Cycle count – Number of times the battery was reconditioned
- Voltage (millivolts)
- Current (milliamps)
- Total Power On Hours

Reconditioning a Battery

Batteries maintain power to the controller cache in the event of a power failure, thus protecting any data that has not been written to a physical drive. Reconditioning is the action of discharging and recharging a battery to preserve its capacity and performance.

Reconditioning is a background activity, it might affect I/O performance. When the recondition is completed, the battery's cycle count increments by one.

Battery reconditioning is disabled by default. You can change the reconditioning status and schedule.

To recondition a battery immediately:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the battery you want, then click the **Recondition** button.
4. If a battery reconditioning is already scheduled, a pop-up dialog menu appears explaining this. If you wish to recondition now, type "confirm" in the entry field and click on the **Confirm** button.

Battery operations status changes to "Recondition" and the battery's remaining capacity and estimated hold time fall and rise reflecting the discharge and recharge cycles of the reconditioning. That behavior is normal.

Making Schedule Changes

To make changes the scheduled battery reconditioning:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

The list of Background Activities displays.

3. Click the **Scheduler** button.
4. Mouse-over **Battery Reconditioning** and click the Settings button.
5. Make settings changes as required:
   - Start Time
   - Uncheck the Enable This Schedule box to disable this activity.
   - Recurrence Pattern
   - Start From
   - End On

Click the **Save** button to apply the new settings.

Buzzer Settings

To make buzzer settings:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Buzzer and click the **Settings** button.
4. Check the E**nable Buzzer** box to enable the buzzer, or uncheck the box to disable.
5. Click the **Save** button.

Caution

If you disable the buzzer, it is disabled for all events. To stop the buzzer from sounding, you can use either the Mute Alarm button located on the front of the subsystem hardware, or follow the instructions below.

Silencing the Buzzer

To silence the buzzer, you can press the Mute Alarm button on the front of the VTrak enclosure hardware, located on the left side under the power button and USB ports. Or, click on the virtual Mute Buzzer button at the top of the user interface in WebPAM PROe. *This only appears when the buzzer is sounding.*

Managing Enclosures

Enclosure management includes the following functions:

- "Viewing Enclosure Topology"
- "Viewing the Enclosures Summary"
- "Locating an Enclosure"
- "Viewing Enclosure Information"
- "Making Enclosure Settings"
- "Viewing FRU VPD Information"
- "Viewing Power Supply Status"
- "Viewing Cooling Unit Status"
- "Viewing Temperature Sensor Status"
- "Viewing Voltage Sensor Status"

Viewing Enclosure Topology

This feature displays the connection topology of the VTrak E5000 subsystem. Topology refers to the manner in which the data paths among the enclosures are connected. There are three methods:

- **Individual Subsystem** – A single subsystem
- **JBOD Expansion** – Managed through one subsystem or head unit
- **RAID Subsystem Cascading** – Managed through one subsystem or head unit

To view enclosure topology:

1. Click the **Device** tab.
2. Click the **Topology** icon.

The topology or data connections of your system displays.

Topology display

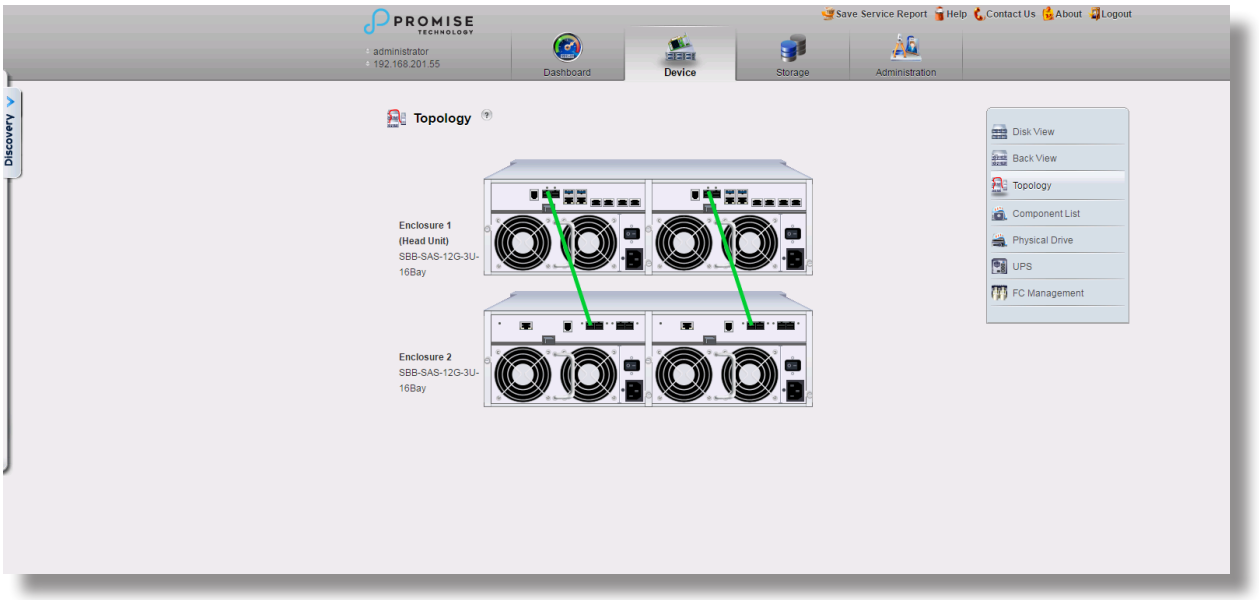

Viewing the Enclosures Summary

Enclosure Management includes information, status, settings and location. To access Enclosure Management:

1. Click the **Device** tab.
2. Click the **Component List** icon.

   The following information is shown:

   - Enclosure ID number
   - Status
   - Enclosure Type
   - Status Description (specific components in need of attention, if any)

Locating an Enclosure

To locate an enclosure:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure you want, then click the **Locate** button.

   The enclosure LEDs blink for one minute.

Viewing Enclosure Information

To view enclosure information:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **View** button.

   Enclosure information includes:

   - Enclosure ID
   - Enclosure Type
   - Enclosure Warning Temperature Threshold
   - Controller Warning Temperature Threshold
   - Max Number of Controllers
   - Max Number of Fans
   - Max Number of Temperature Sensors
   - Max Number of Batteries
   - Enclosure Critical Temperature Threshold
   - Controller Critical Temperature Threshold
   - Max Number of Physical Drive Slots
   - Max Number of Blowers
   - Max Number of Power Supply Units
   - Max Number of Voltage Sensors

   For information on Enclosure problems, See “Enclosure Problems” on page 426.

Making Enclosure Settings

To make Enclosure settings:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **Settings** button.

   Enclosure settings include:

   - Enclosure Warning Temperature Threshold [47-51]°C
   - Enclosure Critical Temperature Threshold [57-61]°C
   - Controller Warning Temperature Threshold [61-65] [66-70] [71-78]°C
   - Controller Critical Temperature Threshold [68-72] [73-77] [81-88]°C

Note: there are three Warning, and three Critical threshold segments for Controller sensors.

4. In the field provided, type the temperature in degrees C for each threshold value.
5. Click the Save button.

Viewing FRU VPD Information

FRU VPD refers to Vital Product Data (VPD) information about Field Replaceable Units (FRU) in the enclosure. The number and type of FRU depends on the subsystem model.

To view FRU VPD information:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **FRU VPD** button.

   Use this information when communicating with Technical Support and when ordering replacement units.

   For contact information, See "Contacting Technical Support" on page 58.

Viewing Power Supply Status

To view the status of the power supplies and the fans that cool those power supplies:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **View** button.
4. Scroll down to view the power supplies.

   The screen displays the operational and fan status of the power supplies. If any status differs from normal or the fan speed is below the Healthy Threshold value, a malfunction is indicated in the Status column.

   See "Replacing a Power Supply" on page 58.

Viewing Cooling Unit Status

To view the status of the cooling units:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **View** button.
4. Scroll down to view the Fans.

The screen displays the status and speed of the cooling units. If blower speed is below the Healthy Threshold, a malfunction is indicated in the Status column. See “Enclosure Problems” on page 426.

Viewing Temperature Sensor Status

To view the status of the temperature sensors:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **View** button.
4. Scroll down to view the Temperature Sensors.

If any temperature exceeds the Healthy Threshold value, an overheat condition exists in the enclosure. See ""Making Enclosure Settings” on page 113" and “Enclosure Problems” on page 426.

Viewing Voltage Sensor Status

To view the status of the voltage sensors:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the Enclosure and click the **View** button.
4. Scroll down to view the Voltage Sensors.

If any voltage is outside the Healthy Threshold values, a voltage malfunction in the enclosure is indicated in the Status column. See "Enclosure Problems" on page 426.

Managing UPS Units

Uninterruptible Power Supply (UPS) Management includes the following functions:

- "Viewing UPS Units"
- "Making UPS Settings"
- "Viewing UPS Information"

Viewing UPS Units

To view a list of UPS units supporting the VTrak E5000:

1. Click the Device tab.
2. Click the UPS icon.

   Information in the UPS List includes:

   - **ID** – The ID number of the UPS
   - **Status** – OK means Normal.
     On AC means the UPS is connected to a viable external AC power source.

     On Battery means the external AC power source is offline and the UPS is running on battery power.
   - **Model** – Model name of the UPS
   - **Battery Capacity** – Backup capacity expressed as a percentage.
   - **Remaining Minutes** – Number of minutes the UPS is expected to power your system in the event of a power failure.

Making UPS Settings

These settings control how the VTrak E5000 subsystem detects the UPS unit and responds to data reported by the UPS unit.

To make UPS settings:

1. Click the **Device** tab.
2. Click the **UPS** icon.
3. Click the **UPS Settings** button.
4. Perform the following actions as required:
   - Verify the Current UPS Communication method. *See Note 1*:
     * **SNMP** – Network connection.
     * **USB**
     * **Unknown** – No connection.
   - Choose a Detection Setting from the drop-down menu:
     * **Automatic** – Default. If a UPS is detected when the subsystem boots, the settings changes to Enable.
     * **Enable** – Monitors UPS. Settings changes, reports warnings, and logs events.
     * **Disable** – Does not monitor UPS.
   - Type values into the Threshold fields:
     * **Running Time Remaining Threshold** – Actual time below this value resets adaptive writeback cache to writethrough.
   - For UPS units with network cards, type the IP addresses or DNS names in fields UPS 1 and UPS 2. *See Note 2*.
5. Press **Submit** to save your settings.

Note 1: VTrak E5000 supports multiple UPS units using network or USB connections, but not a combination of both methods.

Note 2: To specify UPS units by DNS names, ask your IT administrator to add the DNS names to the DNS server, before you make UPS settings.

Viewing UPS Information

To view information about a specific UPS unit:

1. Click the **Device** tab.
2. Click the **UPS** icon.
3. Mouse-over UPS and click the **View** button.

UPS information includes:

- **UPS ID**
- **Model Name**
- **Serial Number**
- **Firmware Version**
- **Manufacture Date**
- **Battery Capacity** – Backup capacity expressed as a percentage.
- **Remaining Backup Time** – Number of minutes the UPS is expected to power your system in the event of a power failure.

Managing Network Connections

Network Connections Management includes the following functions:

- "Making Virtual Management Port Settings"
- "Making Maintenance Mode Settings"

Making Virtual Management Port Settings

The VTrak E5000 subsystem has a virtual management port, enabling you to log into a VTrak E5000 with dual controllers using one IP address.

Before you change settings, please See "About IP Addresses" on page 36.

You initially made these settings during subsystem setup. You can change them later as required.

Caution

Changing virtual management port settings can interrupt your WebPAM PROe connection and require you to log in again.

To make virtual management port settings:

1. Click the **Administration** tab.
2. Click the **Network Management** icon.
3. Click the **Virtual Management Port** tab.
4. Click the protocol family whose settings you want to change and click the **Configuration** button.
5. Make the following settings are needed:
   - Check the E**nable** box to enable this protocol family.
   - Check the E**nable DHCP** box to enable a DHCP server to make your network settings. DHCP is currently supported in IPv4 only.
   - For manual network settings, type the RAID subsystem's IP address, subnet mask, gateway IP address, and DNS server IP address into the fields provided.
6. Click the **Submit** button.

Making Maintenance Mode Settings

Each controller has its own IP addresses for access when the controller goes into maintenance mode.

Before you change settings, please See “About IP Addresses” on page 68.

To make maintenance mode settings:

1. Click the **Administration** tab.
2. Click the **Network Management** icon.
3. Click the **Maintenance Mode** tab.
4. Click the controller and protocol family whose settings you want to change and click the **Configuration** button.
5. Make the following settings are needed:
   - Check the **Enable** box to enable this protocol family.
   - Check the **Enable DHCP** box to enable a DHCP server to make your network settings. DHCP is currently supported in IPv4 only.
   - For manual network settings, type the **IP address, subnet mask, gateway IP address**, and **DNS server IP address** into the fields provided.
6. Click the **Submit** button.

Managing Users

User management includes:

- "Viewing User Information"
- "Creating a User"
- "Making User Settings"
- "Changing User Passwords"
- "Deleting a User"
- "Setting User Event Subscriptions"
- "Importing a User Database"
- "Exporting a User Database"

The **Administrator** or a **Super User** can perform these tasks.

Viewing User Information

To view user information:

1. Click the **Administration** tab.
2. Click the **User Management** icon.

   The list of users displays. User information includes:

   - User name
   - Status
   - Privilege level
   - Display Name
   - Email address
   - User Type

Important

Click the new user and click the Subscription button to set up email for a user to receive event notification. See "Setting User Event Subscriptions" on page 126.

Creating a User

This action requires **Administrator** or **Super User** privileges.

To create a user:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. Click the **Add User** button.
4. In the **Add User** dialog box, enter the information in the fields provided:
   - Name – This is the user's login name
   - Display Name
   - Password
   - Retype Password
   - User Email – Required for event notification
5. Choose a privilege level from the drop-down menu. See the table below for a description of the privilege types.
6. (Optional) Uncheck the Enable box to disable this User account.
7. Click the **Save** button. The user is added to the list.

User Privileges

Level	Meaning
View	Allows the user to See all status and settings but not to make any changes
Maintenance	Allows the user to perform maintenance tasks including Rebuilding, PDM, Media Patrol, and Redundancy Check
Power	Allows the user to create (but not delete) disk arrays and logical drives, change RAID levels, change stripe size; change settings of components such as disk arrays, logical drives, physical drives, and the controller
Super	Allows the user full access to all functions including create and delete users and changing the settings of other users, and delete disk arrays and logical drives. The default "administrator" account is a **Super User**

Making User Settings

This action requires **Administrator** or a **Super User** privileges.

To make user settings:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. In the User list, click the user you want, then click **Settings**.
4. Make settings changes as required:
   - For the E**nable** box, check to enable this user account, uncheck to disable this user account
   - In the User Settings dialog box, enter a new **Display Name** or **User** E**mail** address
   - Choose a new **Privilege** level from the drop-down menu. See the table on the next page.
5. Click the **Save** button.

User Privileges

Level	Meaning
View	Allows the user to See all status and settings but not to make any changes
Maintenance	Allows the user to perform maintenance tasks including Rebuilding, PDM, Media Patrol, and Redundancy Check
Power	Allows the user to create (but not delete) disk arrays and logical drives, change RAID levels, change stripe size; change settings of components such as disk arrays, logical drives, physical drives, and the controller
Super	Allows the user full access to all functions including create and delete users and changing the settings of other users, and delete disk arrays and logical drives. The default “administrator” account is a **Super User**

Changing User Passwords

This action requires **Administrator** or **Super User** privileges.

To change a user's password:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. In the User list, click the user you want, then click **Change Password**.
4. In the Change Password dialog box, enter the information in the fields provided:
   - New Password
   - Retype Password
5. Click the **Save** button.

Deleting a User

This action requires **Administrator** or **Super User** privileges

To delete a user:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. In the User list, click the user you want, then click the **Delete** button.
4. In the **Confirmation** box, type the word "**confirm**" in the field provided and click the **Confirm** button.

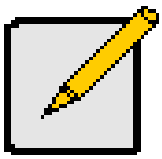

Note

The Administrator account cannot be deleted.

Setting User Event Subscriptions

By default, all users have event notification:

- Enabled
- Set to the Major (severity) level for all events

Subscribing users receive notification of events at the chosen severity level and all higher levels.

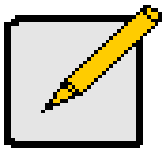

Note

Each user must have a valid Email address to receive events.

See "Making User Settings" on page 124.

Changing a user subscription requires **Administrator** or **Super User** privileges.

To set a user event subscription:

1. Click the **Administration** tab.
2. Click the **User Management** icon.
3. In the User list, click the user you want, then click the **Subscription** button.
4. Make settings changes as required:
   - For the E**nable** Event **Notification** box, check to enable for this user, uncheck to disable.
   - Click to change the priority options for each category of event.
5. Click the **Save** button.

Importing a User Database

You can save the user information and settings from one VTrak E5000 RAID subsystem, export it, and then import it to automatically configure your other VTrak E5000 RAID subsystems.

Caution

Importing a user database overwrites the current users and user settings on your VTrak E5000 subsystem.

To import a user database:

1. Click the **Administration** tab.
2. Click the **Import/Export** icon.
3. Click the **Import** option.
4. Choose **User Database** from the **Type** drop-down menu.
5. Click the **Browse** button and navigate to the user database file and click the **OK** button.
6. Click the **Next** button.

   The system verifies that the file is a valid user database and displays any errors or warnings.
7. Click the **Submit** button to continue.
8. In the **Confirmation** box, type the word "**confirm**" in the field provided and click the Confirm button.

   The user database is imported and applied automatically.

Exporting a User Database

You can save the user information and settings from one VTrak E5000 RAID subsystem, export it, and then import it to automatically configure your other VTrak E5000 RAID subsystems.

To export a user database:

1. Click the **Administration** tab.
2. Click the **Import/**E**xport** icon.
3. Click the E**xport** option.
4. Choose **User Da**tab**ase** from the **Type** drop-down menu.
5. Click the **Submit** button.
6. In the **Open** dialog box, click the **Save File** option, then click the **OK** button.

The file is saved to your PC as “User.dat”.

Note

The user database file is not designed to be opened or edited in the field.

Managing Background Activities

Background activity management includes:

- “Viewing Current Background Activities”
- “Viewing Scheduled Background Activities”
- “Adding a Scheduled Background Activity”
- “Changing a Background Activity Schedule”
- “Enabling/Disabling Scheduled Background Activity”
- “Media Patrol”
- “Redundancy Check”
- “Initialization”
- “Rebuild”
- “Migration”
- “PDM”
- “Transition”
- “Synchronization”
- “Battery Reconditioning”

Background activities perform a variety of preventive and remedial functions on your physical drives, disk arrays, logical drives, and other components.

You can run a background activity immediately or schedule it to run at a later time. Scheduling options are described below.

Setting options for each activity are listed after the scheduling options. These settings determine how the background activity affects I/O performance.

Viewing Current Background Activities

To view a list of current background activities:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background appears.

   Currently running activities show a progress bar.

Viewing Scheduled Background Activities

To view a list of scheduled background activities:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background appears.
3. Click the **Scheduler** button.

   The list of currently scheduled background activities appears.

Adding a Scheduled Background Activity

To add a new scheduled background activity:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background appears.
3. Click the **Scheduler** button.

   The list of currently scheduled background activities appears.
4. Click the **Add Schedule** button.
5. Check the E**nable Media Patrol** box to enable, uncheck to disable.

This settings enables or disables Media Patrol for your entire RAID system.

6. Click the **Confirm** button.
7. Choose the option for the activity you want:
   - Media Patrol
   - Redundancy Check
   - Spare Check
   - Battery Recondition
8. Choose a **Start Time** from the drop-down menus.

   The menus have a 24-hour clock.
9. Choose a **Recurrence Pattern** option, daily, weekly, or monthly.
   - For the Daily option, enter an interval in the Every field.
   - For the Weekly option, enter an interval in the Every field and choose one or more days of the week.
   - For the Monthly option, choose, Day of the Month option then choose a number from the drop-down menu.
     The day of the week option then choose the day of the month from the drop-down menus.
10. Choose a **Start From** date from the drop-down menus.
11. Choose an E**nd On** option,
    - No end date or perpetual.
    - End after a specific number of activity actions.
    - Until date from the drop-down menus.
12. For **Redundancy Check**, choose,
    - **Auto Fix** option – Attempts to repair the problem when it finds an error. Check to enable
    - **Pause on Error** option – The process stops when it finds a non-repairable error. Check to enable
    - **Select LD** – Check the boxes for the logical drives to run Redundancy Check. Check at least one logical drive
13. Click the **Save** button.

Changing a Background Activity Schedule

To change an existing scheduled background activity:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background appears.
3. Click the **Scheduler** button.

   The list of currently scheduled background activities appears.
4. Click the background activity and click the **Settings** button.
5. Make settings changes as required:
   - Choose a **Start Time** from the drop-down menus.
     The menus have a 24-hour clock.
   - Choose a **Recurrence Pattern** option, daily, weekly, or monthly.
     For the Daily option, enter an interval in the Every field.
     For the Weekly option, enter an interval in the Every field and choose one or more days of the week.
     For the Monthly option, choose the Day of the Month option or the day of the week option, and choose the day from the drop-down menu.
   - Choose a **Start From** date from the drop-down menus.
   - Choose an E**nd On** option,
     No end date or perpetual.
     End after a specific number of activity actions.
     Until date from the drop-down menus.
   - For **Redundancy Check**, choose,
     Auto Fix option – Attempts to repair the problem when it finds an error. Check to enable
     Pause on Error option – The process stops when it finds a non-repairable error. Check to enable
     Select LD – Check the boxes for the logical drives to run Redundancy Check. Check at least one logical drive
6. Click the **Save** button.

Enabling/Disabling Scheduled Background Activity

Background activity schedules are enabled by default when you create the schedule. If you want to stop a background activity now but plan to use it again in the future, disable the scheduled activity rather than deleting it.

To enable or disable change an existing scheduled background activity:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background appears.
3. Click the **Scheduler** button.

   The list of currently scheduled background activities appears.
4. Click the background activity and click the **Settings** button.
5. Uncheck the E**nable This Schedule** box to disable this schedule.

   Check the box to enable this schedule.
6. Click the **Save** button.

Deleting a Scheduled Background Activity

To change an existing scheduled background activity:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background appears.
3. Click the **Scheduler** button.

   The list of currently scheduled background activities appears.
4. Click the background activity and click the **Delete** button.
5. In the confirmation box, click the confirm button.

Media Patrol

Media Patrol is a routine maintenance procedure that checks the magnetic media on each disk drive. Media Patrol checks are enabled by default on all disk arrays and spare drives. Media Patrol is concerned with the media itself, not the data recorded on the media. If Media Patrol encounters a critical error, it triggers a PDM if PDM is enabled on the disk array.

See “Making Disk Array Settings” on page 186.

Making Media Patrol Settings

To make Media Patrol settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background appears.
3. Click the **Settings** button.
4. Check the E**nable Media Patrol** box to enable, uncheck to disable.

   This settings enables or disables **Media Patrol** for your entire RAID system.
5. Click the **Confirm** button.

You can also enable or disable **Media Patrol** on individual disk arrays.

Redundancy Check

Redundancy Check is a routine maintenance procedure for fault-tolerant disk arrays (those with redundancy) that ensures all the data matches exactly. Redundancy Check can also correct inconsistencies. See “Redundancy Check on a Logical Drive” on page 201.

Making Redundancy Check Settings

To make Redundancy Check settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Click the **Settings** button.
4. Click the **Redundancy Check Rate** drop-down menu and choose a rate:
   - **Low** – Fewer system resources to Redundancy Check, more to data read/write operations.
   - **Medium** – Balances system resources between Redundancy Check and data read/write operations.
   - **High** – More system resources to Redundancy Check, fewer to data read/write operations.
5. Click the **Confirm** button.

Initialization

Technically speaking, **Initialization** is a foreground activity, as you cannot access a logical drive while it is initiating.

Initialization is normally done to logical drives after they are created from a disk array. Initialization sets all data bits in the logical drive to zero. The action is useful because there may be residual data on the logical drives left behind from earlier configurations. For this reason, Initialization is recommended whenever you create a logical drive. See “Initializing a Logical Drive” on page 200.

Making Initialization Settings

To make initialization settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.
3. Click the **Settings** button.
4. Click the Logical Drive Initialization Rate drop-down menu and choose a rate:
   - **Low** – Fewer system resources to Initialization, more to data read/write operations.
   - **Medium** – Balances system resources between Initialization and data read/write operations.
   - **High** – More system resources to Initialization, fewer to data read/write operations.
5. Click the **Confirm** button.

Rebuild

When you rebuild a disk array, you are actually rebuilding the data on one physical drive.

- When a physical drive in a disk array fails and a spare drive of adequate capacity is available, the disk array begins to rebuild automatically using the spare drive.
- If there is no spare drive of adequate capacity, but the Auto Rebuild function is ENABLED, the disk array begins to rebuild automatically as soon as you remove the failed physical drive and install an unconfigured physical drive in the same slot. See “Making Rebuild Settings” below.
- If there is no spare drive of adequate capacity and the Auto Rebuild function is DISABLED, you must replace the failed drive with an unconfigured physical drive, then perform a **Manual Rebuild**.

See “Rebuilding a Disk Array” on page 189 and “Managing Spare Drives” on page 206.

Also see “Disk Array Degraded / Logical Drive Critical” on page 436 and “Disk Array Offline / Logical Drive Offline” on page 437.

Making Rebuild Settings

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Click the **Settings** button.
4. Click the **Rebuild Rate** drop-down menu and choose a rate:
   - **Low** – Fewer system resources to the Rebuild, more to data read/write operations.
   - **Medium** – Balances system resources between the Rebuild and data read/write operations.
   - **High** – More system resources to the Rebuild, fewer to data read/write operations.
5. Check the E**nable Auto Rebuild** box to enable Auto Rebuild (rebuilds when you swap out the failed drive with a new one).
6. Click the **Confirm** button.

Migration

The term “Migration” means either or both of the following:

- Change the RAID level of a logical drive.
- Expand the storage capacity of a logical drive.

See “Migrating a Logical Drive’s RAID Level” on page 202 and “RAID Level Migration” on page 382.

Making Migration Settings

To make migration settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Click the **Settings** button.
4. Click the **Migration Rate** drop-down menu and choose a rate:
   - **Low** – Fewer system resources to Migration, more to data read/write operations.
   - **Medium** – Balances system resources between Migration and data read/write operations.
   - **High** – More system resources to Migration, fewer to data read/write operations.
5. Click the **Confirm** button.

PDM

Predictive Data Migration (PDM) is the migration of data from the suspect physical drive to a spare drive, similar to rebuilding a logical drive. But unlike Rebuilding, PDM constantly monitors your physical drives and automatically copies your data to a spare drive before the physical drive fails and your logical drive goes Critical.

See “Running PDM on a Disk Array” on page 188 and “PDM” on page 365.

Making PDM Settings

To make PDM settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Click the **Settings** button.
4. Make the following settings are required:
   - Click the **PDM Rate** drop-down menu and choose a rate:
     * **Low** – Fewer system resources to PDM, more to data read/write operations.
     * **Medium** – Balances system resources between PDM and data read/write operations.
     * **High** – More system resources to PDM, fewer to data read/write operations.
   - Highlight the current values in the block threshold fields and input new values.
     Reassigned block threshold range is 1 to 512 blocks.

     Error block threshold range is 1 to 2048 blocks.
5. Click the **Confirm** button.

Transition

Transition is the process of replacing a revertible spare drive that is currently part of a disk array with an unconfigured physical drive or a non-revertible spare drive.

See “Running a Transition on a Spare Drive” on page 211.

Making Transition Settings

To make Transition settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Click the **Settings** button.
4. Click the **Transition Rate** drop-down menu and choose a rate:
   - **Low** – Fewer system resources to Transition, more to data read/write operations.
   - **Medium** – Balances system resources between Transition and data read/write operations.
   - **High** – More system resources to Transition, fewer to data read/write operations.
5. Click the **Confirm** button.

Synchronization

Synchronization is automatically applied to redundant logical drives when they are created. Synchronization recalculates the redundancy data to ensure that the working data on the physical drives is properly in sync.

Mouse-over on the logical drive, click the View button, and look under Logical Drive Information beside the line that says Synchronized. A ***Yes*** means the logical drive was synchronized. See “Viewing Logical Drive Information” on page 192.

Making Synchronization Settings

To make Synchronization settings:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Click the **Settings** button.
4. Click the Synchronization Rate drop-down menu and choose a rate:
   - **Low** – Fewer system resources to Synchronization, more to data read/write operations.
   - **Medium** – Balances system resources between Synchronization and data read/write operations.
   - **High** – More system resources to Synchronization, fewer to data read/write operations.
5. Click the **Confirm** button.

Battery Reconditioning

Batteries maintain power to the controller cache in the event of a power failure, thus protecting any data that has not been written to a physical drive. Reconditioning is the action of discharging and recharging a battery to preserve its capacity and performance.

Battery Reconditioning is a Background Activity that must be scheduled to employ. See “Adding a Scheduled Background Activity” on page 130. Also see “Viewing Battery Information” on page 106 and “Battery Reconditioning” on page 141.

Managing Storage Services

Storage service management includes:

- "Viewing Services"
- "Email Service"
- "SLP Service"
- "Webserver Service"
- "Telnet Service"
- "SSH Service"
- "SNMP Service"
- "CIM Service"
- "Netsend Service"

Viewing Services

This feature displays all software services running on the RAID subsystem.

To view the list of software services:

1. Click the **Administration** tab.
2. Click the Services icon.

The Services list displays the Status and Start Type of the services available. These services are described in the sections that follow. To view the actions available for the various services in the list, click in any row to see the **Start** button and **Settings** button for the respective service, as well as the **Stop** button and **Restart** button for services that are currently running.

Email Service

Email service enables the RAID subsystem to send you Email messages about events and status changes. By default, Email service is set to Automatic.

Stopping Email Service

To stop the Email service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Email service and click the **Stop** button.
4. Click the **Confirm** button.

To start the Email service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Email service and click the **Start** button.

Restarting Email Service

To restart the Email service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Email service and click the **Restart** button.

Making Email Settings

To change Email service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Email service and click the **Settings** button.
4. Make settings changes as required:
    - Choose a startup type,
        * Automatic – (default) Starts and runs with the subsystem.
        * Manual – You start the service when you need it.
    - SMTP Server IP address.
    - SMTP Authentication – The Yes option enables authentication. The No option disables.
    - SMTP Authentication Username – Required if SMTP authentication is enabled.
    - SMTP Authentication Password – Required if SMTP authentication is enabled.
    - Email Sender (From) Address – The sender's name shown on notification messages.
    - Email Subject – The subject line of the notification message.
5. To enable SSL for the connection to the Email server, check the **Enable SSL** box.
6. To send a test email, one time, check the **Send a Test Email** box.
7. Click the **Save** button.
8. Click the **Confirm** button.

Note

To verify your settings, send a test message.

SLP Service

Service Location Protocol (SLP) discovers services over the Internet. SLP applies to IPv4 protocol only.

Stopping SLP Service

To stop the SLP service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SLP service and click the **Stop** button.
4. Click the **Confirm** button.

To start the SLP service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SLP service and click the **Start** button.

Restarting SLP Service

To restart the SLP service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SLP service and click the **Restart** button.

Making SLP Settings

To change SLP service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SLP service and click the **Settings** button.
4. Choose a startup type:
   - Automatic – (default) Starts and runs with the subsystem.
   - Manual – You start the service when you need it.
5. Click the **Save** button.
6. Click the **Confirm** button.

Webserver Service

Webserver service connects the WebPAM PROe interface to the RAID subsystem though your browser.

Stopping Webserver Service

To stop the Webserver service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Webserver service and click the **Stop** button.
4. Click the **Confirm** button.

To start the Webserver service after stopping it:

1. Click the **Administration** tab.
2. Click the Services icon.
3. Click the Webserver service and click the **Start** button.

Restarting Webserver Service

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Webserver service and click the **Restart** button.

Making Webserver Settings

To change Webserver service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Webserver service and click the **Settings** button.
4. Make settings changes as required:
    - Choose a startup type,
        * Automatic – (default) Starts and runs with the subsystem.
        * Manual – You start the service when you need it.
    - Session Time Out – Default is 24 minutes.
5. Click the **Save** button.
6. Click the **Confirm** button.

Telnet Service

Telnet service enables you to access the RAID subsystem's Command Line Interface (CLI) through a network connection.

Stopping Telnet Service

To stop the Telnet service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **Telnet** service and click the **Stop** button.
4. Click the **Confirm** button.

To start the Telnet service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **Telnet** service and click the **Start** button.

Restarting Telnet Service

To restart the Telnet service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click Telnet service and click the **Restart** button.

Making Telnet Settings

To change Telnet service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **Telnet** service and click the **Settings** button.
4. Make settings changes as required:
   - Choose a startup type,
     * Automatic – Starts and runs with the subsystem.
     * Manual – (default) You start the service when you need it.
   - Port number – Default is 2300.
   - Max Number of Concurrent Connections – Default is 4.
     Maximum number is 4.
   - Session Time Out – Default is 24 minutes.
5. Click the **Save** button.
6. Click the **Confirm** button.

SSH Service

Secure Shell (SSH) service enables you to access the subsystem's Command Line Interface (CLI) through a network connection.

Stopping SSH Service

To stop SSH service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SSH** service and click the **Stop** button.
4. Click the **Confirm** button.

To start SSH service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click SSH service and click the **Start** button.

Restarting SSH Service

To restart SSH service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SSH** service and click the **Restart** button.

Making SSH Settings

To change SSH service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SSH service and click the **Settings** button.
4. In the SSH Settings tab, make settings changes as required:
    - Choose a startup type,
        * Automatic – (default) Starts and runs with the subsystem.
        * Manual – You start the service when you need it.
    - Port number - Default is 22.
    - Max Number of Concurrent Connections – Default is 4.
      Maximum number is 4.
    - Session Time Out - Default is 24 minutes.
5. Click the **Save** button.
6. Click the **Confirm** button.

SSH Public Key Management

To upload a public key for SSH service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SSH service and click the **Settings** button.
4. Choose the **SSH Public Key Management** tab.
5. Click the **Browse** button and locate the file with the public key.
6. Click the **Upload** button. The new SSH public key configuration appears listed in the bottom of the menu.

   Up to sixteen public keys can be stored.

To remove a public key configuration:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the SSH service and click the **Settings** button.
4. Choose the **SSH Public Key Management** tab.
5. Click the selection box for the public key configuration to be removed, and click on the **Delete** button.

SNMP Service

Simple Network Management Protocol (SNMP) service enables the SNMP browser to obtain information from the RAID subsystem. The Trap Sink is where SNMP events are sent and can be viewed.

Stopping SNMP Service

To stop the SNMP service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SNMP** service and click the **Stop** button.
4. Click the **Confirm** button.

To start the SNMP service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SNMP** service and click the **Start** button.

Restarting SNMP Service

To restart the SNMP service:

1. Click the Administration tab.
2. Click the Services icon.
3. Click the SNMP service and click the Restart button.

Making SNMP Settings

To change SNMP service settings:

1. Click the Administration tab.
2. Click the Services icon.
3. Click the SNMP service and click the Settings button.
4. Make settings changes as required:
    - Choose a startup type,
        * Automatic – (default) Starts and runs with the subsystem.
        * Manual – You start the service when you need it.
    - Port Number – Default is 161.
    - System Name – No default.
    - System Location – Default is USA.
    - System Contact – Default is admin@yourcompany.com.
    - Read Community – Default is public.
    - Write Community – not currently supported
5. Click the **Save** button.
6. Click the **Confirm** button.

See next section for adding an SNMP Trap Sink.

Adding an SNMP Trap Sink

To add a trap sink:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SNMP** service and click the **Settings** button.
4. Enter a trap sink server IP address in the field provided.
5. Choose a trap filter (event severity level). See “Event Severity Levels” on page 160.
6. Click the **Add** button.
7. Click the **Confirm** button.

Deleting an SNMP Trap Sink

To delete a trap sink:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SNMP** service and click the **Settings** button.
4. In the **Trap Sink** list and click the Trap Sink you want to delete.
5. Click the **Trash** icon. The trap sink is deleted.
6. Click the **Save** button.
7. Click the **Confirm** button.

CIM Service

Stopping CIM Service

To stop the CIM service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **CIM** service and click the **Stop** button.
4. Click the **Confirm** button.

To start the CIM service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **CIM** service and click the **Start** button.

Restarting CIM Service

To restart the CIM service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **CIM** service and click the **Restart** button.

Making CIM Settings

The Common Information Model (CIM) service provides a database for information about computer systems and network devices.

To make CIM service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Mouse-over CIM and click the **Settings** button.
4. Make settings changes as required.

   Choose a startup type.

   Automatic - (default) Starts and runs with the subsystem

   Manual - You start the service when you need it

   - **CIM HTTP Enabled** - Default is Yes
   - **CIM HTTP Port** number - Default is 5988
   - **CIM HTTPS Enabled** - Default is No
   - **CIM HTTPS Port** number - Default is 5989
   - **CIM Authentication** - Default is No
   - When **CIM Authentication** is Yes, these fields appear.
     * **Change Password** - Default is No
     * **CIM User Name** - Default is cim. No changes are possible
   - If **Change Password** is *Yes*, these fields appear.
     * **Old User Password**
     * **New User Password**
     * **Retype Password** (new user)
5. Click the **Save** button.
6. Click the **Confirm** button.

Netsend Service

Netsend service sends RAID subsystem events in the form of text messages to the Host PC and other networked PCs configured to receive Netsend event messages by setting up Netsend server accounts.

This service is set to Manual startup by default. It does not run unless you start it manually or change the startup type to Automatic.

Starting Netsend Service

To restart the Netsend service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the Start button.

Stopping Netsend

To stop the Netsend service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the Netsend service and click the Stop button.
4. Click the **Confirm** button.

Restarting Netsend Service

To start the Netsend service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the **Start** button.

Making Netsend Settings

To change Netsend service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the **Settings** button.
4. Choose a startup type,
   - Automatic – Starts and runs with the subsystem.
   - Manual – (default) You start the service when you need it.
5. Click the **Save** button.
6. Click the **Confirm** button.

Adding Netsend Server Accounts

To add a Netsend server account:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the **Settings** button.
4. Enter the recipient server IP address in the field provided.
5. Choose a recipient filter (event severity level). See “Event Severity Levels” on page 160.
6. Click the **Add** button.

   The recipient server is added to the list.
7. Click the **Save** button.
8. Click the **Confirm** button.

Event Severity Levels	
Level	**Description**
Fatal	Non-recoverable error or failure has occurred.
Critical	Action is needed now and the implications of the condition are serious.
Major	Action is needed now.
Minor	Action is needed but the condition is not a serious at this time.
Warning	User can decide whether or not action is required.
Information	Information only, no action is required.

Deleting Netsend Server Accounts

To delete a Netsend server account:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **Netsend** service and click the **Settings** button.
4. In the **Message Event Severity Filter** list, click the recipient server you want to delete.
5. Click the **Trash** icon.

   The recipient server is deleted.
6. Click the **Save** button.
7. Click the **Confirm** button.

SSH Service

Stopping SSH Service

To stop the SSH service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click **SSH** service and click the **Stop** button.
4. Click the **Confirm** button.

To start the SSH service after stopping it:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **SSH** service and click the **Start** button.

Restarting

To restart the SSH service:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Click the **SSH** service and click the **Restart** button.

Making SSH Settings

Secure Shell (SSH) service enables you to access the subsystem's Command Line Interface (CLI) through a network connection.

To make SSH service settings:

1. Click the **Administration** tab.
2. Click the **Services** icon.
3. Mouse-over **SSH** and click the **Settings** button.
4. Make settings changes as required.
   - Choose a startup type.
     * **Automatic** - (default) Starts and runs with the subsystem
     * **Manual** - You start the service when you need it
   - **Port** number - Default is 22
   - **Max Number of Concurrent Connections** - Default is 4. Maximum number is 4.
   - **Session Time Out** - Default is 24 minutes
5. Click the **Save** button.
6. Click the **Confirm** button.

Working with the Event Viewer

Working with the Event Viewer includes the following functions:

- "Viewing Runtime Events"
- "Saving Runtime Events"
- "Clearing Runtime Events"
- "Viewing NVRAM Events"
- "Saving NVRAM Events"
- "Clearing NVRAM Events"

The Event Viewer displays log of subsystem events. Events are classified as:

- Runtime Events – A list of and information about the 1023 most recent runtime events recorded since the subsystem was started.
- **NVRAM Events** – A list of and information about the most important events over multiple subsystem startups. NVRAM events are stored in non-volatile memory.

Event Severity Levels

Level	Description
Fatal	Non-recoverable error or failure has occurred.
Critical	Action is needed now and the implications of the condition are serious.
Major	Action is needed now.
Minor	Action is needed but the condition is not a serious at this time.
Warning	User can decide whether or not action is required.
Information	Information only, no action is required.

Viewing Runtime Events

To display Runtime Events:

1. Click the **Administration** tab.
2. Click the **Events** icon.

   The log of Runtime Events appears. Events are added to the top of the list. Each item includes:

   - **Index number** – Begins with 0 at system startup.
   - **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
   - **Event ID** – Hexadecimal code for the specific event
   - **Severity** – see table on previous page
   - **Time** – Date and time the event happened.
   - **Description** – A description of the event in plain language.
3. Press the up and down arrow keys to scroll through the log.

Saving Runtime Events

This feature saves a plain text file of runtime events to your host PC or server using your browser.

To save the Runtime Events log:

1. Click the **Administration** tab.
2. Click the **Events** icon.
3. Click the **Save** button.
4. Follow your browser's procedure to save the event file to the desired location.

Clearing Runtime Events

To clear the Runtime Events log:

1. Click the **Administration** tab.
2. Click the **Events** icon.
3. Click the **Clear** button.
4. In the Confirmation box, type the word "**confirm**" in the field provided and click the Confirm button.

Viewing NVRAM Events

This screen displays a list of and information about the most important events over multiple subsystem startups.

To display NVRAM events:

1. Click the **Administration** tab.
2. Click the **Events** icon.
3. Click the **NVRAM** Event**s** button.

   The log of NVRAM Events appears. Events are added to the top of the list. Each item includes:

   - **Index number** – Begins with 0 at system startup.
   - **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
   - **Event ID** – Hexadecimal code for the specific event
   - **Severity** – See ."Event Severity Levels" on page 163.
   - **Timestamp** – Date and time the event happened.
   - **Description** – A description of the event in plain language.
4. Press the up and down arrow keys to scroll through the log.

Saving NVRAM Events

This feature saves a plain text file of NVRAM events to your host PC or server using your browser.

To save NVRAM Events:

1. Click the **Administration** tab.
2. Click the Event**s** icon.
3. Click the **NVRAM** Event**s** button.
4. Click the **Save** button.
5. Follow your browser's procedure to save the event file to the desired location.

Clearing NVRAM Events

To clear the Runtime Events log:

1. Click the **Administration** tab.
2. Click the Event**s** icon.
3. Click the **Clear** button.
4. In the Confirmation box, type the word "confirm" in the field provided and click the **Confirm** button.

MONITORING PERFORMANCE

Performance monitoring includes:

- "Monitoring I/O Performance"
- "Monitoring PSU Wattage"

Monitoring I/O Performance

The Performance Monitor displays real-time performance statistics for logical drives, physical drives, and Fibre Channel data ports. The vertical scale adjusts dynamically to accommodate the statistical data.

Because it reports performance in real-time, to see data in the monitor, there must be I/O data activity taking place between the VTrak E5000 subsystem and the Host.

To monitor performance:

1. Click the **Administration** tab.
2. Click the **Performance Monitor** icon. Follow the instructions below for the menu category you want to view.

Controller

1. Under **Controller**, choose the metric you want to see from the **Measurement** drop-down menu.
   - Cache usage by %
   - Dirty cache usage by %
2. Click the **Select Controllers** button and check the boxes for the one or both controllers to view the cache usage data.

Logical Drive

1. Under **Logical Drive**, choose the metric you want to see from the **Measurement** drop-down menu.
   - Bandwidth in MB/s
   - Cache usage by %
   - Dirty cache usage by %
   - Maximum latency in ms
   - Average latency in ms

- Minimum latency in ms
- I/Os per second

2. Click the **Select Logical Drives** button and check the boxes for the logical drives you want to see.
    - Total of all logical drives
    - Up to 8 individual logical drives

Physical Drive

1. Under **Physical Drive**, choose the metric you want to see from the **Measurement** drop-down menu.
    - Bandwidth in MB/s
    - Maximum latency in ms
    - Average latency in ms
    - Minimum latency in ms
    - I/Os per second
2. Click the Select **Physical Drives** button and check the boxes for the physical drives you want to see.
    - Total of all physical drives
    - Up to 8 individual physical drives

Port

1. Under **Port**, choose the metric you want to see from the **Measurement** drop-down menu.
    - Bandwidth in MB/s
    - Maximum latency in ms
    - Average latency in ms
    - Minimum latency in ms
    - I/Os per second
2. Click the **Select Ports** button and check the boxes for the ports you want to see:
    - Total of all ports
    - Up to 8 individual ports

Since the **Performance Monitor** is a real-time display, it does not accumulate information and there is no clear or save function.

To save performance statistics for analysis or troubleshooting, save a **Service Report**, open the report, and look under **Statistic Info**.

Monitoring PSU Wattage

The PSU Wattage Monitor displays real-time performance statistics for logical drives, the input power of all enclosures and the input power of an individual. The vertical scale adjusts dynamically to accommodate the statistical data.

Because it reports performance in real-time, to see data in the monitor, there must be I/O data activity taking place between the VTrak E5000 subsystem and the Host.

To monitor performance and power use:

1. Click the **Administration** tab.
2. Click the **PSU Wattage Monitor** icon.
3. Under **Input Power of an individual** Enclosure, click the **Select** Enclosure**s** button and check the boxes for the enclosures you want to see.

Since the PSU Wattage Monitor is a real-time display, it does not accumulate information and there is no clear or save function.

To save performance and power statistics for analysis or troubleshooting, save a Service Report, open the report, and look under PSU Wattage Info.

Managing Physical Drives

Physical drive management includes:

- "Viewing Physical Drives"
- "Viewing Physical Drive Information"
- "Making Global Physical Drive Settings"
- "Making Individual Physical Drive Settings"
- "Viewing Physical Drive Statistics"
- "Viewing Physical Drive SMART Log Information"
- "Saving the Physical Drive SMART Log"
- "Locating a Physical Drive"
- "Forcing a Physical Drive Offline"
- "Clearing a Stale or a PFA Condition"
- "Updating Firmware on a Physical Drive"

For physical disk troubleshooting, see "Physical Drive Problems" on page 516.

Viewing Physical Drives

To view a list of physical drives in the RAID system:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.

The list of enclosures and the physical drives inside them displays.

Physical drive information includes:

- ID – ID number of the physical drive
- Status – Green check , yellow ! , and red X icons
- Model – Make and model of the drive
- Type – SAS or SATA, HDD or SSD
- Location – Enclosure number and slot number
- Configuration – Array number and sequence number, spare number, unconfigured, or stale configuration
- Capacity – In GB

Viewing Physical Drive Information

To view physical drive information:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **View** button.

Physical drive information includes:

- Physical Drive ID – ID number of the physical drive
- Alias – If assigned
- Configurable Capacity – Usable capacity in GB
- Block Size – Typically 512 Bytes
- Configuration – Array number and sequence number, spare number,
- Drive Interface – SATA 1.5Gb/s or 3Gb/s, SAS 3Gb/s or 6Gb/s
- Firmware Version – Firmware version on the drive
- Visible To – Controllers that can access this physical drive
- Location – Enclosure number and slot number
- Physical Capacity – Total capacity in GB
- Used Capacity – Capacity actually used in GB
- Operational Status – OK is normal, Stale, PFA, Dead
- Model Number– Make and model of the drive
- Serial Number – Serial number of the drive
- Protocol Version – ATA/ATAPI or SCSI protocol version

Advanced information for SAS physical drives includes:

- Read Cache – Enabled or disabled
- Read Cache Support – Yes or No
- Write Cache – Enabled or disabled
- Write Cache Support – Yes or No
- Enable Read Look Ahead Support – Yes or No
- Read Look Ahead Cache – Enabled or disabled
- Command Queuing – Enabled or disabled
- Command Queuing Support – Yes or No
- WWN – World Wide Name
- Port 1 Negotiated Physical Drive Speed
- Port 1 SAS Address
- Port 2 Negotiated Physical Drive Speed
- Port 2 SAS Address
- Drive Temperature in ºC
- Drive Reference Temperature in ºC
- Power Saving Level – Enabled or disabled
- SSD SMART Life Percentage
- SSD SSD Remaining Reserved Blocks Percentage
- SSD Trim Type
- Medium Error Threshold
- SAS SATA Bridge Firmware Version
- SAS SATA Bridge Boot Loader Version

Advanced information for SATA physical drives includes:

- Read Look Ahead Cache – Enabled or disabled
- Read Cache Support – Yes or No
- SMART Feature Set – Yes or No
- SMART Self Test – Yes or No
- SMART Error Logging – Yes or No
- Queue Depth - Number of commands
- Maximum Multiple DMA Mode Supported
- Maximum Ultra DMA Mode Supported
- DMA Mode
- Power Saving Level – Enabled or disabled
- APM Support – Standby or Active
- Medium Error Threshold
- Drive Temperature
- Drive Reference Temperature
- SSD SMART Life Percentage
- SSD SSD Remaining Reserved Blocks Percentage
- SSD Trim Type

Making Global Physical Drive Settings

To make global physical drive settings:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the **Global Physical Drive Settings** button.
4. Check the boxes to enable, uncheck to disable.

   For **SATA** drives:

   - Enable Write Cache
   - Enable Read Look Ahead Cache
   - DMA Mode
   - Medium Error Threshold

   For **SAS** drives:

   - Enable Write Cache
   - Enable Read Look Ahead Cache
   - Enable Read Cache
   - Medium Error Threshold
5. Click the **Save** button.

Making Individual Physical Drive Settings

To make individual physical drive settings:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **Settings** button.
4. On the **Settings** tab:
   - Enter, change, or delete the alias in the **Alias** field.
5. On the **SMART Log Settings** tab:
   - Check the box to enable the SMART log.
6. Click the **Save** button.

Viewing Physical Drive Statistics

To view physical drive statistics:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **View** button.
4. Click the **Statistics** tab.

   Physical drive statistics include

- Data Transferred
- Read Data Transferred
- Write Data Transferred
- Errors - Number of errors
- Non Read/Write Errors
- Read Errors
- Write Errors
- I/O Request – Number of requests
- Non Read/Write Request – Number of requests
- Read I/O Request – Number of requests
- Write I/O Request – Number of requests
- Statistics Start Time – Time and date
- Statistics Collection Time – Time and date
- Avg Response Time Ctrl 1 – Controller 1 average response time
- Avg Response Time Ctrl 2 – Controller 2 average response time
- Max Response Time Ctrl 1 – Controller 1 maximum response time
- Max Response Time Ctrl 2 – Controller 2 maximum response time

Viewing Physical Drive SMART Log Information

SMART Log is disabled by default. SMART Log information is only visible if it is first enabled. Keep in mind that enabling SMART Log can degrade performance on some physical drives, and it is might not be supported on some drives.

To view physical drive SMART Log information when it is enabled:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **View** button. The Physical Drive Information display appears.
4. Click the **SMART Log** tab.

When enabled, SMART Log information includes:

- Physical Drive ID
- SMART Status
- Current Drive Temperature
- Maximum Start Stop Count
- Manufactured Year
- SMART Support
- SMART Health Status
- Drive Trip Temperature
- Current Start Stop Count
- Manufactured Ween In Year

Saving the Physical Drive SMART Log

To save the physical drive SMART Log:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **View** button.
4. Click the **SMART Log** tab.
5. Click the **Save Advanced SMART Log** button.

Your browser saves a text file containing the SMART Log to its designated download folder.

Locating a Physical Drive

This feature causes the drive carrier LEDs to blink for one minute to assist you in locating the physical drive, and is supported by RAID subsystems and JBOD expansion units.

To locate a physical drive:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **Locate** button.

The drive carrier status LED flashes for one minute.

Forcing a Physical Drive Offline

This feature applies only to physical drives assigned to disk arrays.

Caution

Forcing a physical drive offline is likely to cause data loss. Back up your data before you proceed. Use this function only when required.

Important

Forcing a physical drive offline causes your logical drives to become degraded. If Auto Rebuild is enabled and a spare drive is available, the disk array begins rebuilding itself automatically.

To force a physical drive offline:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the **down arrow** button to list the physical drives in the enclosure.
4. Mouse over the physical drive you want to force offline.
5. Click the **Force Offline** button.
6. In the **Confirmation** box, type the word "**confirm**" in the field provided and click the **Confirm** button.

Clearing a Stale or a PFA Condition

This procedure is used to clear configuration data on a physical drive; or if the physical drive is stale or the has errors putting it in PFA status.

Stale – The physical drive contains obsolete disk array information.

PFA – The physical drive has errors resulting in a prediction of failure.

To clear a **Stale** or a **PFA** condition:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the **Clear** button.

   If the physical drive has both a Stale condition and a PFA condition, the first click removes the Stale condition. Click the **Clear** button a second time to remove the PFA condition.

   If a physical drive is still online and shows a PFA error but “Clear PFA” does

   not appear, use PDM to copy the data to a new physical drive. See “Running PDM on a Disk Array” on page 188.

After you clear a PFA error, watch for another PFA error to appear. If it does, replace the physical drive.

Updating Firmware on a Physical Drive

This feature applies only to PROMISE-supported physical drives. For a list of supported drives, go to http://www.promise.com/support/.

If you have physical drives in your RAID system that are not PROMISE-supported, follow the firmware update procedure from the drive manufacturer.

To update physical drive firmware, first get the firmware file from the manufacturer website and save it to your Host PC or TFTP server. Then follow these steps:

1. Click the **Administration** tab.
2. Click the Firmware **Update** icon.
3. Click on **PD Firmware Update**.
4. Choose a download option:
   - **Local File through HTTP** – Click the **Browse** button, locate the firmware image file, click the file to choose it, then click the **Open** button.
   - **TFTP Server** – Enter the TFTP Server host name or IP address, port number and file name.
5. Click the **Next** button.

   The next screen shows the Flash Image (firmware image file).
6. Click the **Submit** button.

   The progress of the update displays.

Warning

Do NOT power off the RAID subsystem during the update!

Do NOT move to any other screen until the firmware update operation is completed!

7. Click the **OK** button.

Managing Disk Arrays

For disk array troubleshooting, see “Disk Array and Logical Drive Problems” on page 517.

Disk array management includes:

- “Viewing Disk Arrays”
- “Viewing Disk Array Information”
- “Creating a Disk Array Manually”
- “Creating a Disk Array with the Wizard”
- “Deleting a Disk Array”
- “Locating a Disk Array”
- “Making Disk Array Settings”
- “Running Media Patrol on a Disk Array”
- “Running PDM on a Disk Array”
- “Preparing a Disk Array for Transport”
- “Rebuilding a Disk Array”

Viewing Disk Arrays

To view a list of disk arrays:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

   The list of disk arrays appears.

Disk array information includes:

- **ID** – DA0, DA1, DA2, etc.
- **Alias** – If assigned
- **Status** – A green check icon means OK
- **Capacity** – Data capacity of the array
- **Free Capacity** – Unconfigured or unused capacity on the physical drives
- **Media Patrol** – Enabled or disabled on this array
- **No. of Logical Drives** – The number of logical drives on this array

Viewing Disk Array Information

To view disk array information:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

   The list of disk arrays appears.
3. Click the disk array you want, then click the **View** button.

   Array information displays, including:

   - **Disk Array ID** – DA0, DA1, DA2, etc.
   - **Alias** – If assigned
   - **Operational Status** – OK is normal
   - **Media Patrol** – Enabled or disabled on this array
   - **PDM** – Enabled or disabled on this array
   - **Power Management** - Enabled or disabled on this array
   - **Total Physical Capacity** – Data capacity of the array
   - **Configurable Capacity** – Maximum usable capacity of the array
   - **Free Capacity** –Unconfigured or unused capacity on the physical drives
   - **Max Contiguous Free Capacity -** The largest contiguous free capacity available.
   - **Current Power Saving Level** - Default is disabled
   - **User Set Power Saving Level** - Enabled or disabled on this array
   - **Number of Physical Drives** – The number of physical drives in this array
   - **Number of Logical Drives** – The number of logical drives on this array
   - **Number of Dedicated Spares** - Number of spare drives dedicated for this array
   - **Available RAID Levels** – RAID levels you can specify on this array

Disk Array Operational Status

- ***OK*** – This is the normal state of a logical drive. When a logical drive is OK, it is ready for immediate use. For RAID Levels other than RAID 0 (Striping), the logical drive has full redundancy.
- ***Synchronizing*** – Synchronizing is a temporary condition. It is a function of initialization that occurs when a logical drive is created. The purpose is to verify the integrity of data and redundancy in the logical drive. When a logical drive is Synchronizing, it functions and your data is available. However, access is slower due to the synchronizing operation.
- ***Critical/Degraded*** – This indicates a physical drive failure. A degraded logical drive still functions and your data is still available. However, the logical drive has lost redundancy (fault tolerance). You should determine which drive (or drives) are at fault and replace it as soon as possible.
- ***Rebuilding*** – This condition is temporary. When a physical drive has been replaced, the logical drive automatically begins rebuilding in order to restore redundancy (fault tolerance). When a logical drive is rebuilding, it functions and your data is available. However, access is slower due to the rebuilding operation.
- ***Transport Ready*** – After you perform a successful Prepare for Transport operation, this condition means you can remove the physical drives of this disk array and move them to another enclosure or different drive slots. After you relocate the physical drives, the disk array status shows OK.

Creating a Disk Array Manually

This feature creates a disk array only. You can also use the Wizard to create a disk array with logical drives and spare drives at the same time.

This action requires **Super User** or **Power User** privileges.

To create a disk array:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.
3. Click the **Create Disk Array** button.
4. Accept the defaults or make changes:
   - Enter an alias in the **Alias** field
     Maximum of 32 characters; letters, numbers, space between characters, and underline.
   - **Media Patrol** – Uncheck to disable on this array.
   - **PDM** – Uncheck to disable on this array.
   - **Power Management** – Uncheck to disable on this array.
   - **Choose a media type** – Hard disk drive (HDD) or solid state drive (SSD)
5. In the **Select Physical Drives** diagram, click the drives to add them to your array. Look for drives with a green LED dark, a blue LED lit, and no crosshatching over the carrier.

   The ID numbers of the chosen drives appear in the field below the diagram.
6. When you have finished your settings and choices, click the **Submit** button.

   The new array appears in the list.

   If you are done creating disk arrays, click the **Finish** button.

   To create additional disk arrays, click the **Create More** button.

After you create a disk array, create a logical drive on it.

Creating a Disk Array with the Wizard

The Wizard creates disk arrays and logical drives automatically. It has four options.

- ***Automatic*** – Creates a new disk array following a default set of parameters. Creates a hot spare drive for all RAID levels except RAID 0, when five or more unconfigured physical drives are available. You can accept or reject the proposed arrangement but you cannot modify it. See instructions in “Automatic Configuration” on page 68.
- ***Advanced*** – Enables you to specify all parameters for a new disk array, logical drives and spare drives. See instructions in “Advanced Configuration” on page 69.

Deleting a Disk Array

Caution

If you delete a disk array, you also delete any logical drives that belong to it, along with the data in those logical drives. Back up any important data before deleting a disk array.

This action requires **Administrator** or **Super User** privileges.

To delete a disk array:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.
3. Click the disk array you want, then click the **Delete** button.
4. In the **Confirmation** box, type the word "**confirm**" in the field provided and click the **Confirm** button.

Locating a Disk Array

This feature causes the drive carrier LEDs to flash for one minute to assist you in locating the physical drives that make up this disk array.

To locate a disk array:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

   The list of disk arrays appears.
3. Click the disk array you want, then click the **Locate** button.

   The drive carrier status LEDs for the array flash for one minute.

Drive carrier status LED

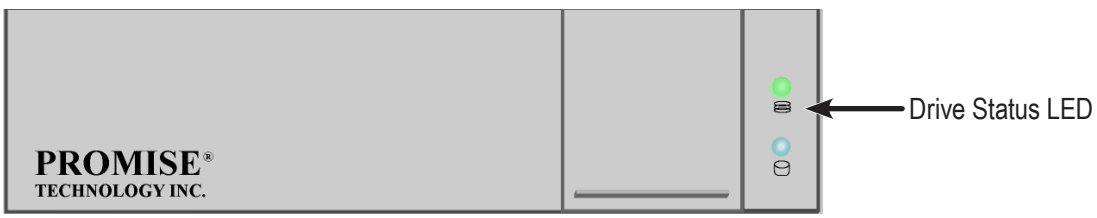

Making Disk Array Settings

To make disk array settings:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

   The list of disk arrays appears.
3. Click the disk array you want, then click the **Settings** button.
4. Make settings changes as required:
   - Enter, change or delete the alias in the **Alias** field
     Maximum of 32 characters; letters, numbers, space between characters, and underline.
   - **Media Patrol** – Check to enable, uncheck to disable on this array.
   - **PDM** – Check to enable, uncheck to disable on this array.
   - **Power Management** – Check to enable, uncheck to disable on this array.
5. Click the **Save** button.

Notes

You can also enable or disable Media Patrol for the entire RAID system. See "Making Media Patrol Settings" on page 134.

HDD Power Saving must be enabled on the RAID controller for the Power Management settings to be effective. See "Making Controller Settings" on page 100.

Power Management functions are limited to the features your HDDs actually support.

Running Media Patrol on a Disk Array

Media Patrol is a routine maintenance procedure that checks the magnetic media on each disk drive. If Media Patrol encounters a critical error, it triggers PDM if PDM is enabled on the disk array.

For more information, see “Media Patrol” and “PDM”.

Running Media Patrol

To run **Media Patrol:**

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Mouse-over Media Patrol and click the **Start** button.

Stopping, Pausing or Resuming Media Patrol

To stop, pause or resume **Media Patrol**:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background appears.
3. The progress of the running Media Patrol is displayed. Notice there are buttons on which you can click to **Stop**, **Pause** or **Resume** the Media Patrol.
4. If you click the **Stop** button, you will need to confirm your action in a confirmation dialog menu that appears. Type “confirm” and click the **Confirm** button to stop running Media Patrol.

Running PDM on a Disk Array

Predictive Data Migration (PDM) is the migration of data from the suspect disk drive to a spare disk drive.

For more information, see “PDM” on page 435.

Running PDM

To run PDM on a disk array:

1. Click the **Administration** tab.
2. Click the Background **Activities** icon.

   The list of background activities appears.
3. Mouse-over PDM and click the **Start** button.
4. From the **Source Physical Drive** drop-down menu, choose a Source disk array and physical drive.
5. From the **Target Physical Drive** drop-down menu, choose a Target physical drive.
6. Click the **Confirm** button.

Stopping, Pausing or Resuming PDM

To stop, pause or resume PDM:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. The progress of the running PDMis displayed. Notice there are buttons on which you can click to **Stop**, **Pause** or **Resume** the Media Patrol.
4. If you click the **Stop** button, you will need to confirm your action in a confirmation dialog menu that appears. Click the **Confirm** button to stop running PDM.

You can also enable or disable PDM on individual disk arrays. See “Making Disk Array Settings” on page 186.

Preparing a Disk Array for Transport

This feature prepares the physical drives that make up the disk array to be removed from the enclosure and installed in a different location.

To prepare a disk array for transport:

1. Click the **Storage** tab.
2. Click the **Disk Array** icon.

   The list of disk arrays appears.
3. Click the disk array you want, then click the **Transport** button.
4. You need to confirm your action in a confirmation dialog menu that appears. Type "confirm" and click the **Confirm** button.

   The status changes to ***Transport Ready***.
5. Remove the physical drives and install them in their new location.

   For more information, see "Installing Physical Drives" on page 35.

Rebuilding a Disk Array

When you rebuild a disk array, you are actually rebuilding the data on one physical drive.

If there is no spare drive of adequate capacity and the **Auto Rebuild** function is *DISAB*LED, you must replace the failed drive with an unconfigured physical drive, then perform a Manual Rebuild. See "Making Rebuild Settings" on page 137.

Important

If your replacement disk drive was formerly part of a different disk array or logical drive, you must clear the configuration data on the replacement drive before you use it. See "Clearing a Stale or a PFA Condition" on page 179.

Performing a Manual Rebuild

To perform a manual rebuild:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Mouse-over Rebuild and click the **Start** button.
4. From the Source Physical Drive drop-down menu, choose a Source disk array and physical drive.

   Arrays have an ID No. Physical drives have a Seq. No. (sequence number)
5. From the Target Physical Drive drop-down menu, choose a Target physical drive.
6. Click the **Confirm** button.

   When the disk array is rebuilding:

   - The disk array shows a green check icon and ***Rebuilding*** status.
   - Logical drives under the disk array continue to show a yellow ! icon and ***Critical*** status.

Stopping, Pausing or Resuming a Rebuild

To stop, pause or resume a Rebuild:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background appears.
3. The progress of the rebuild is displayed. Notice there are buttons on which you can click to **Stop**, **Pause** or **Resume** rebuilding.
4. If you click the **Stop** button, you will need to confirm your action in a confirmation dialog menu that appears. Click the **Confirm** button to stop rebuilding.

Managing Logical Drives

Logical drive management includes:

- “Viewing Logical Drives”
- “Viewing Logical Drive Information”
- “Viewing Logical Drive Statistics”
- “Viewing Logical Drive Check Tables”
- “Creating a Logical Drive Manually”
- “Deleting a Logical Drive”
- “Making Logical Drive Settings”
- “Locating a Logical Drive”
- “Initializing a Logical Drive”
- “Redundancy Check on a Logical Drive”
- “Migrating a Logical Drive’s RAID Level”
- “Creating a LUN Clone”

Viewing Logical Drives

To view a list of logical drives:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

   The list of logical drives appears.

   Logical Drive information includes:

   - **ID** – LD0, LD1, LD2, etc.
   - **Alias** – If assigned.
   - **Status** – A green check
   - icon means OK.
   - **Capacity** – Data capacity of the logical drive.
   - **RAID Level** – Set when the logical drive was created.
   - **Stripe** – Set when the logical drive was created.
   - **Cache Policy** – Read cache and Write cache settings.
   - **Array ID** – ID number of the disk array where this logical drive was created.

Viewing Logical Drive Information

To view logical drive information:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

   The list of logical drives appears.

   Click the logical drive you want, then click the **View** button.

Logical Drive information displays, including:

- **Logical Drive ID** – LD0, LD1, LD2, etc.
- **Alias** – If assigned
- **Array ID** –- ID number of the disk array where this logical drive was created
- **RAID Level** – Set when the logical drive was created
- **RAID 5/6 Algorithm** - Algorithm used to create RAID 5 or RAID 6 (if applicable)
- **Operational Status** – OK means normal
- **Capacity** – Data capacity of the logical drive
- **Number of Axles** – For RAID 10, 2 axles. For RAID 50 and 60, 2 or more axles
- **Physical Capacity** – Data capacity of the physical drives
- **Number of Physical Drives** – The number of physical drives in the disk array
- **Stripe size** – Set at logical drive creation
- **Read Policy** – Adjustable
- **Sector size** – Set at logical drive creation
- **Write Policy** – Adjustable
- **Current Write Policy** – Adjustable
- **Preferred Controller ID** – For RAID subsystems with dual controllers
- **Tolerable Number of Dead Drives** – Number of physical drives that can fail without the logical drive going offline
- **Host Accessibility** - Normal, read-only, write-only, or not visible to host
- **Synchronized** – A new logical drive shows **"**No**"** until synchronizing is completed. See "Synchronization" on page 141.
- **Parity Pace** – Pertains to some RAID levels
- **WWN** – Worldwide Name, a unique identifier assigned to this logical drive
- **RAID6 Scheme** – Pertains to RAID 6 only
- **Serial Number** – Assigned to this logical drive
- **ALUA Access State for Ctr**l 1 - Active, optimized or standby
- **ALUA Access State for Ctr**l 2 - Active, optimized or standby
- **Association State** - for LUN clone
- **Storage Service Status** - for LUN clone
- **Perfect Rebuild** - Enable or disable

Viewing Logical Drive Statistics

To view logical drive statistics:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

   The list of logical drives appears.
3. Click the logical drive you want, then click the **View** button.
4. Click the **Statistics** tab.

   Logical Drive statistics display, including:

   - **Data Transferred** – In Megabytes
   - **Read Data Transferre**d – In Megabytes
   - **Write Data Transferred** – In Megabytes
   - **Errors**
   - **Non-Read/Write Errors**
   - **Read** E**rrors**
   - **Write** E**rrors**
   - **I/O Requests**
   - **Non-Read/Write I/O Requests**
   - **Read I/O Requests**
   - **Write I/O Requests**
   - **Statistics Start Time**

   Statistics Collection Time

   To clear physical drive statistics, see "Clearing Statistics" on page 90.

Viewing Logical Drive Check Tables

This feature enables you to view error tables. Use this information to evaluate the integrity of the logical drive and to determine whether corrective action is needed.

To view logical drive check tables:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

   The list of logical drives appears.
3. Click the logical drive you want, then click the **Check Table** button.
4. Choose an option:
   - **All** – All errors. The default choice.
   - **Read Check** – The Read Check Table is used to hold LBA entries with Invalid Data Blocks on a per LD (logical disk) basis. Invalid Data Blocks are the data in a sector which is physically good but the data may be invalid. Read Check Table is used to record these LBAs which have invalid data. A read to this LBA by the host will be returned as Medium Error. A write to that LBA will clear the entry.
   - **Write Check** – Write Check errors or this logical drive.
   - **Inconsistent Block** – Inconsistent blocks for this logical drive. Mirror data for RAID Levels 1 and 10 or Parity data for RAID Levels 3, 5, 6, 30, 50, and 60. Identified by the Redundancy Check.

   The Check Table lists:
   - **Entry Number** – A number assigned to each block of entry.
   - **Table Type** – Read Check, Write Check or Inconsistent Block.
   - **Start Logical Block Address** – LBA of the first block for this entry.
   - **Count** – Number of errors or continuous blocks starting from this LBA.

   To clear the check tables, use the CLI command **checktable**.

   For example, to clear a read checktable entry on logical drive 0

   ```
   checktable -a clear -t rct -l 0
   ```

Note: It is not necessary to reboot after clearing the checktable, a rescan from the host side should be sufficient.

Creating a Logical Drive Manually

This feature creates a logical drive only. You can also use the Wizard to create a disk array with logical drives and spare drives at the same time.

This action requires **Super User** or **Power User** privileges.

To create a logical drive manually:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.
3. Click the **Create Logical Drive** button.
4. Click the option button of the disk array you want to use and click the **Next** button.
5. Optional. Enter an alias in the **Alias** field.

   Maximum of 32 characters; letters, numbers, space between characters, and underline.
6. Choose a RAID level.

   The choice of RAID levels depends the number of physical drives in the disk array.
7. RAID 50 and 60 only. Specify the number of axles for your array.
8. In the **Capacity** field, accept the default maximum capacity or enter a lesser capacity and size in MB, GB or TB.

   Any remaining capacity is available for an additional logical drive.
9. For each of the following items, accept the default or change the settings as required:
   - Choose a Stripe size.
     64 KB, 128 KB, 256 KB, 512 KB, and 1 MB are available.
   - Choose a Sector size.
     512 B, 1 KB, 2 KB, and 4 KB are available. Read Cache, Read Ahead, and No Cache are available.
   - Choose a Write (cache) Policy.
     Write Back and Write Through (Thru) are available.
   - Choose a Preferred Controller ID

- Check box of Perfect Rebuild Enable / Disable Perfect Rebuild

10. Click the **Add** button.

    The new logical drive appears on the list at the right.

    If there is capacity remaining, you can create an additional logical drive.

11. When you are finished, click the **Submit** button.

    The new logical drive or drives appear in the logical drive list.

    New redundant logical drives are automatically synchronized. You can access the logical drive during synchronization.

Deleting a Logical Drive

Caution

If you delete a logical drive, you also delete all the data in the logical drive. Back up any important data before deleting the logical drive.

This action requires **Administrator** or **Super User** privileges.

To delete a logical drive:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.
3. Click the logical drive you want, then click the **Delete** button.
4. In the Confirmation box, type the word "**confirm**" in the field provided and click the **Confirm** button.

Making Logical Drive Settings

To make logical drive settings:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

   The list of logical drives appears.
3. Click the logical drive you want, then click the **Settings** button.

Note

NAS Disk Pools and Share Disks do not have a feature corresponding to Logical Drive settings

4. Make settings changes as required:
   - Enter, change, or delete the alias in the Alias field.
     Maximum of 32 characters; letters, numbers, space between characters, and underline.
   - Choose a Read (cache) Policy.
     Read Cache, Read Ahead, and No Cache are available.
   - Choose a Write (cache) Policy.
     Write Back and Write Through (Thru) are available.
   - Choose a Preferred Controller ID
   - Check box of Perfect Rebuild Enable / Disable Perfect Rebuild Note that if Perfect Rebuild is disabled, it cannot be enabled again on the LD.
5. Click the **Save** button.

   For more information, see “Cache Policy” on page 473.

Note

The Write Cache is always set to WriteThru when Read Cache is set to NoCache.

Locating a Logical Drive

This feature causes the drive carrier LEDs to flash for one minute to assist you in locating the physical drives that make up this logical drive.

To locate a logical drive:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

   The list of logical drives appears.
3. Click the logical drive you want, then click the **Locate** button.

   The drive carrier status LEDs of the disk carriers making up the logical drive flash for one minute.

Drive carrier status LED

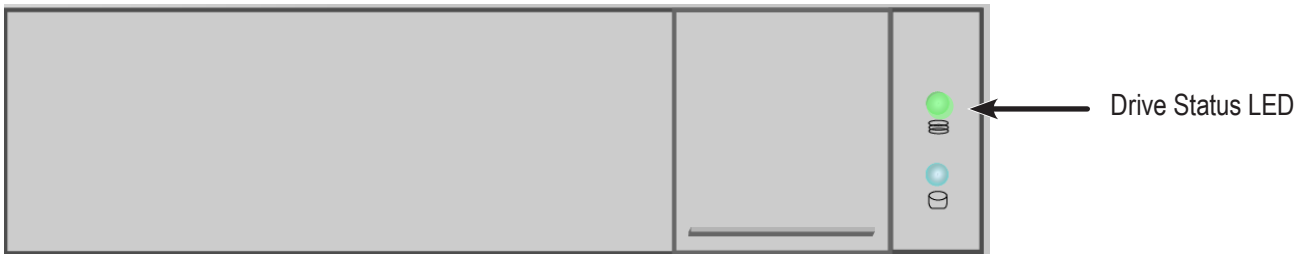

Initializing a Logical Drive

Initialization is normally done to logical drives after they are created from a disk array.

Warning

When you initialize a logical drive, all the data on the logical drive is lost. Backup any important data before you initialize a logical drive.

To initialize a logical drive:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Click the **Start** button.
4. Check the box to the left of the logical drive you want to initialize.
5. Choose the initialization option you want.

   Enter a hexadecimal value in the Initialization Pattern in Hex field or use the default 00000000 value.

   For ***Quick Initialization***, check the box and enter a value in the Quick Initialization Size field. This value is the size of the initialization blocks in MB.
6. Click the **Confirm** button, then type the word "**confirm**" in the field provided, and click **Confirm** button in the pop-up menu.

Stopping, Pausing or Resuming an Initialization

To stop, pause or resume Initialization:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Click in the line of the LD you want to effect, the **Stop**, **Pause**, or **Resume** button appears for a few seconds. Click the button for the action you want to take for the ongoing initialization.
4. If you want to stop initialization, click the **Confirm** button that appears to verify that you want to stop. Then type the word "**confirm**" in the field provided, and click **Confirm** button in the pop-up menu.

Redundancy Check on a Logical Drive

Redundancy Check is a routine maintenance procedure for fault-tolerant disk arrays (those with redundancy) that ensures all the data matches exactly. Redundancy Check can also correct inconsistencies.

To run **Redundancy Check** on a logical drive:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Mouse-over **Redundancy Check** and click the **Start** button.
4. Check the boxes to the left of the logical drives you want to run.
5. Check the options you want:
   - **Auto Fix** – Attempts to repair the problem when it finds an error
   - **Pause on Error** – The process stops when it finds a non-repairable error
6. Click the **Confirm** button, then type the word "**confirm**" in the field provided, and click **Confirm** button in the pop-up menu.

Stopping, Pausing or Resuming a Redundancy Check

To stop, pause or resume **Redundancy Check**:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Click in the line of the LD you want to effect, the **Stop**, **Pause**, or **Resume** button appears for a few seconds. Click the button for the action you want to take for the ongoing redundancy check.
4. If you want to stop a redundancy check, click the **Confirm** button that appears to verify that you want to stop. Then type the word "**confirm**" in the field provided, and click **Confirm** button in the pop-up menu.

Migrating a Logical Drive's RAID Level

The term "Migration" means either or both of the following:

- Change the RAID level of a logical drive.
- Expand the storage capacity of a logical drive.

Before you begin a migration, examine your current disk array to determine whether:

- The physical drives in your array can support the target RAID level.
- There is sufficient capacity to accommodate the target logical drive size.

If you need to add physical drives to your array, be sure there are unassigned physical drives installed in your RAID system before you begin migration.

See "Migration" on page 138, and "RAID Level Migration" on page 382.

Migrating a Logical Drive

To migrate a logical drive:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Mouse-over **Migrate** and click the **Start** button.
4. In the **Select Disk Array** drop-down menu, choose the source disk array.
5. In the **Select Physical Drives** diagram, click the drives to add them to your array.

   The ID numbers of the chosen drives appear in the field below the diagram.
6. Click the **Next** button.
7. Check the box next to the logical drive you want to modify.
8. From the drop-down menu, choose a target RAID level.

   The choice of RAID levels depends the number of physical drives in the disk array. See the Note below.
9. In the **Capacity** field, accept the current capacity.

   Or check the **Expand Capacity** box and enter a greater capacity and size in MB, GB or TB.

   If there is capacity remaining, you can create an additional logical drive.
10. Click the **Next** button.

    The logical drive ID numbers, with the original and target RAID levels and capacities are shown
11. To start migration, click the **Confirm** button then type the word "**confirm**" in the field provided, and click **Confirm** button in the pop-up menu.

Note

If you are adding an even number of physical drives to a RAID 10 array and you want the target array to be RAID 10, you must specify RAID 10 under RAID level.

Creating a LUN Clone

A LUN clone is an exact copy of the original LUN or logical drive, including all the data it contains, at one point in time. Use a LUN clone as a backup or to copy a LUN from one system to another.

A LUN clone has the same capacity, stripe size, read and write policies as the original LUN. However, the LUN clone can be a different RAID level. The choice of RAID levels depends on the disk array. And if you have multiple disk arrays, you can create the LUN clone on a different disk array than the original LUN.

This action requires **Super User** or **Power User** privileges.

LUN Clone options

The VTrak E5000 includes a new LUN Clone option, the Online LUN Clone. This is used to create a copy of a LUN without stopping I/O on the source LUN. All data on the source LUN is copied and synchronized in a background operation. The cloning process runs in the background and continues until it is explicitly stopped by the administrator. This is in contrast to the Offline LUN Clone which requires that the source LUN go offline during the process.

First decide if the LUN is to be cloned to a Disk Array or a Logical Drive and whether to do an offline or online clone.

Important

The action of creating an Offline LUN momentarily takes the original source LUN or logical drive offline, meaning nobody can read or write to it.

To create a LUN clone of a logical drive:

1. Click the **Storage** tab.
2. Click the **Logical Drive** icon.

   The Logical Drive list appears.
3. Click the logical drive you want, then click the **LUN Clone to DA** button to clone the LUN to a Disk Array or **LUN Clone to LD** button to clone the LUN to a Logical Drive.
4. Make settings as required:
   - For **LUN Clone to DA** check available Disk Array in the list as the target Disk Array, Choose a RAID level for the copy of the LD, then choose *Online* or *Offline* for the type of cloning process.
   - For **LUN Clone to LD** check available Logical Drives in the list for the target LD, and choose *Online* or *Offline* for the type of cloning process.
   - For **Online Clone** enter a time in minutes for **Mirror Write End Time**.

     Up to 8 clones of a LUN can be created at a time if there are enough Disk Arrays or LDs available.
5. Click the **Next** button and review your choices.
6. Click the **Start** button to begin the cloning process. You need to **Confirm** the Lun Clone start in a pop-up menu.

   The cloning progress bar displays.

Note the **Target Logical Drive ID**. Use this number to identify the LUN clone in the Logical Drive list.

If you chose a redundant RAID level, the LUN clone is automatically synchronized after creation.

After the LUN clone is created, you can manage it like any other logical drive. See “Making Logical Drive Settings” on page 198, “Locating a Logical Drive” on page 199, and “Deleting a Logical Drive” on page 197.

If LUN mapping is enabled, for users to access the LUN clone, you must map it to an initiator. See “Managing LUNs” on page 215.

Managing Spare Drives

Spare drive management includes:

- "Viewing Spare Drives"
- "Viewing Spare Drive Information"
- "Creating a Spare Drive Manually"
- "Deleting a Spare Drive"
- "Making Spare Drive Settings"
- "Locating a Spare Drive"
- "Running Spare Check"
- "Running a Transition on a Spare Drive"

Viewing Spare Drives

To view a list of spare drives:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.

   Spare Drive information displays, including:

   - **ID** – Spare0, Spare1, etc.
   - **Operational Status** – OK means normal
   - **Configurable Capacity** – Usable capacity of the spare drive
   - **Physical Drive ID** – ID number of the physical drive chosen for this spare
   - **Revertible** – Yes or No
   - **Spare Type** – Global or Dedicated
   - **Dedicated to Array** – ID number of the disk array to which the spare is dedicated

Viewing Spare Drive Information

To view spare drive information:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.

   The list of spare drives appears.
3. Click the spare drive you want, then click the **View** button.

   Spare Drive information displays, including:

   - **Spare Drive ID** – Spare0, Spare1, etc.
   - **Physical Drive ID** – ID number of the physical drive chosen for this spare
   - **Location** - Enclosure and slot location
   - **Model Number** - Physical drive manufacturer and model number
   - **Operational Status** – OK means normal
   - **Spare Type** – Global or Dedicated
   - **Physical Capacity** – Total data capacity of the spare drive
   - **Revertible** – Yes or No
   - **Configurable Capacity** – Usable capacity of the spare drive
   - **Spare Check Status** – Not Checked or Healthy
   - **Media Patrol** – Enabled or Not Enabled
   - **Dedicated to Array** – ID number of the disk array to which the spare is dedicated

Creating a Spare Drive Manually

This feature creates a spare drive only. You can also use the Wizard to create a disk array with logical drives and spare drives at the same time.

This action requires **Super User** or **Power User** privileges.

To create a spare drive:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.
3. Click the **Create Spare Drive** button.
4. For each of the following items, accept the default or change the settings as required:
   - Check the **Revertible** box if you want a revertible spare drive.
     A revertible spare drive returns to its spare drive assignment after you replace the failed physical drive in the disk array and run the Transition function.
   - **Global** – Can be used by any disk array
   - **Dedicated to newly created disk array** – The disk array you are now creating.
5. In the **Select Physical Drives** diagram, click a drive to choose it for your spare.

   The ID number for chosen drive appears in the field below the diagram.
6. Click the **Submit** button.

   If you are done creating spare drives, click the **Finish** button.

   To create another spare drive, click the **Create More** button.

Deleting a Spare Drive

This action requires Administrator or a **Super User** privileges.

To delete a spare drive:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.
3. Click the spare drive you want, then click the **Delete** button.
4. In the Confirmation box, type the word "**confirm**" in the field provided and click the **Confirm** button.

Making Spare Drive Settings

To make spare drive settings:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.
3. Click the spare drive you want, then click the **Settings** button.
4. For each of the following items, accept the default or change the settings as required:
   - In the **Revertible** drop-down menu, choose Yes or No.
   - In the **Spare Type** drop-down menu, choose **Global** or **Dedicated**.
   - If you use chose a Dedicated spare, check the box beside the disk array to which this spare drive is assigned.
5. Click the **Save** button.

Locating a Spare Drive

Spare drives are located in the same way as individual physical drives.

To locate a spare drive:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.

   The list of spare drives appears.
3. In the spare drive list, identify the physical drive ID number.
4. Click the **Device** tab.
5. Click the **Physical Drive** icon.

   The list of physical drives appears.
6. Click the physical drive with the matching ID number and click the **Locate** button.

   The drive carrier LED blinks for one minute.

Running Spare Check

Spare Check verifies the status of your spare drives.

To run spare check:

1. Click the **Storage** tab.
2. Click the **Spare Drive** icon.

   The list of spare drives appears.
3. Click the spare drive you want, then click the **Spare Check** button.
4. Click the **Confirm** button.

   After the “Spare Check completed” message appears, click the **View** button to see *Spare Check Status*.

Running a Transition on a Spare Drive

Transition is the process of replacing a revertible spare drive that is currently part of a disk array with an unconfigured physical drive or a non-revertible spare. You must specify an unconfigured physical drive of the same or larger capacity and same media type as the revertible spare drive.

See “Transition” on page 466.

Running a Transition

To run a transition on a revertible spare drive:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Mouse-over **Transition** and click the **Start** button.
4. From the **Source Physical Drive** drop-down menu, choose a Source disk array and the revertible spare drive.

   Arrays have an ID No. The revertible spare has a Seq. No. (sequence number).
5. From the **Target Physical Drive** drop-down menu, choose a Target unconfigured drive.
6. Click the **Confirm** button.

Stopping, Pausing or Resuming a Transition

To stop, pause or resume Transition:

1. Click the **Administration** tab.
2. Click the **Background Activities** icon.

   The list of background activities appears.
3. Mouse-over **Transition** and click the **Stop**, **Pause**, or **Resume** button.

MANAGING INITIATORS

Initiator management includes:

- "Viewing Initiators"
- "Adding a Fibre Channel Initiator"
- "Deleting a Fibre Channel Initiator"

Viewing Initiators

The VTrak E5000's initiator list displays initiators available for mapping to a LUN or logical drive. You must add initiators to the VTrak E5000's initiator list to make them available for mapping to a LUN.

To view a list of initiators:

1. Click the **Storage** tab.
2. Click the **Initiator** icon.

   The list of initiators appears. Initiator information includes:

   - **Index** – Initiator 0, Initiator 1, Initiator 2, etc.
   - **Initiator Name**

     Fibre Channel – The World Wide Port Name of the initiator, composed of a series of eight, two-digit hexadecimal numbers.

Adding a Fibre Channel Initiator

You must add an initiator to the VTrak E5000's initiator list in order to map your LUN or logical drive to the initiator.

Method 1: Inputting the Initiator Name

This action requires **Administrator** or **Super User** privileges.

To add a Fibre Channel initiator to the list:

1. Click the **Storage** tab.
2. Click the **Initiator** icon.
3. Click the **Add Initiator** button.
4. Input the initiator name in the fields provided.

   A Fibre Channel initiator name is the World Wide Port Name of the initiator, composed of a series of eight, two-digit hexadecimal numbers.
5. Click the **Submit** button.

   The initiator is added.

Method 2: Adding from a List

This action requires **Administrator** or **Super User** privileges.

To add a Fibre Channel initiator to the list:

1. Click the **Device** tab.
2. Click the **Fibre Channel Management** icon.
3. Click the Logged In **Device** tab.
4. Check the box next to the initiator you want to add.
5. Click the **Add to Initiator List** button.

   The initiator is added, and its check box grays out.

Deleting a Fibre Channel Initiator

Caution

If you delete an initiator, you delete the LUN map associated with that initiator. Verify that the LUN map is no longer needed before deleting the initiator

This action requires **Administrator** or **Super User** privileges.

To delete a Fibre Channel initiator:

1. Click the **Storage** tab.
2. Click the **Initiator** icon.
3. Click the initiator you want, then click the **Delete** button.
4. In the Confirmation box, type the word "**confirm**" in the field provided and click the **Confirm** button

   The initiator is removed from VTrak E5000's initiator list.

MANAGING LUNS

LUN management includes:

- "Viewing LUN Maps"
- "LUN Mapping and Masking"
- "Adding a LUN Map"
- "Editing a LUN Map"
- "Deleting a LUN Map"
- "Enabling and Disabling LUN Masking"

Viewing LUN Maps

To view a list of LUN maps:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon.

   The list of LUN maps appears.

LUN Mapping and Masking

This feature applies to Fibre Channel subsystems and controls user access to storage resources.

- LUN Mapping – Maps LUNs to an initiator, so that the initiator can only access only the specified LUNs.
- LUN Masking – The process of applying a LUN Map.

To access LUN mapping:

1. Click the **Storage** tab.
2. Click the **LUN Mapping** button.

To enable LUN Masking:

1. Click the **Storage** tab.
2. Click to check the **LUN Masking** option box.

Adding a LUN Map

For Fibre Channel systems, you can set up an Initiator LUN map.

A maximum of 256 logical drives can be mapped to a Fibre Channel initiator.

To assign a LUN to an initiator, add the initiator first. See “Adding a Fibre Channel Initiator” on page 213.

To add a LUN map:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon.
3. Click the **LUN Mapping** button.
4. Click the drop-down menu to choose the initiators for the LUN map and choose Fibre Channel initiators.
5. Click the **Next** button.

   The second LUN Mapping screen appears.
6. Click a logical drive to highlight it. Then click the < button to assign the logical drive to an initiator.

   Or click the << button to assign all logical drives to an initiator.

   Each logical drive can have only one unique. If you map several LUNs to the same initiator, they all need to have different LUN IDs, beginning with LUN 0, and proceed numbering in consecutive order.
7. Click the **Next** button.

   The final LUN Mapping screen appears showing the initiator or port and LUN map.
8. Click the **Submit** button.

   The new LUN map is created.

Editing a LUN Map

Editing a LUN map is the action of assigning a logical drive or LUN to an initiator. By changing the assignment, you change the initiator's access.

To edit a LUN map:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon. The list of LUN maps appears.
3. Click the LUN map you want to change, then click the **Setting** button.
4. Click the drop-down menu to choose the initiators for the LUN map and choose Fibre Channel initiators.
5. Click a logical drive to highlight it. Then click the < button to assign the logical drive to an initiator.

   Or click the << button to assign all logical drives to an initiator.

   Each logical drive can have only one unique. If you map several LUNs to the same initiator, they all need to have different LUN IDs, beginning with LUN 0, and proceed numbering in consecutive order.
6. Click the **Next** button.

   The LUN Mapping screen shows the edited LUN map.
7. Click the **Submit** button.

Deleting a LUN Map

Deleting a LUN map prevents the initiator from accessing the LUN while LUN masking is enabled.

To delete a LUN map:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon.

   The list of LUN maps appears.
3. Click the LUN map you want, then click the **Delete** button.
4. In the Confirmation box, type the word “confirm” in the field provided and click the **Confirm** button.

Enabling and Disabling LUN Masking

Disabling LUN masking allows all initiators to access all LUNs in your data storage. However, disabling LUN masking does not delete existing LUN maps.

These actions require **Administrator** or **Super User** privileges.

To enable or disable LUN masking:

1. Click the **Storage** tab.
2. Click the **LUN Mapping & Masking** icon.
3. Check the box to enable LUN Masking.

   Or uncheck the box to disable LUN Masking.

   LUN Masking starts or stops as soon as you make your setting.

Managing Fibre Channel Connections

Fibre Channel management includes:

- "Viewing Fibre Channel Node Information"
- "Viewing Fibre Channel Port Information"
- "Making Fibre Channel Port Settings"
- "Viewing Fibre Channel Initiators on the Fabric"
- "Viewing Fibre Channel Logged-in Devices"
- "Viewing Fibre Channel SFPs"

Also see "Adding a Fibre Channel Initiator" on page 213 and "Deleting a Fibre Channel Initiator" on page 214.

Viewing Fibre Channel Node Information

To view Fibre Channel node information:

1. Click the **System** button to go to System mode. *Fibre Channel nodes are managed in System mode.*
2. Click the **Device** tab.
3. Click the **Fibre Channel Management** icon.
4. Click the **Node** tab.

   Node information includes:

   - **Worldwide Node Name (WWNN)**
   - **Maximum Frame Size**
   - **Supported Fibre Channel Class**
   - **Supported Speed**

Viewing Fibre Channel Port Information

To view Fibre Channel port information:

1. Click the **System** button to go to System mode. *Fibre Channel nodes are managed in System mode.*
2. Click the **Device** tab.
3. Click the **Fibre Channel Management** icon.
4. Click the **Port** tab.
5. Mouse-over a Fibre Channel port to access and click the **View** button.

   Port information includes:

   - **Fibre Channel Port ID - Fibre Channel port number**
   - **State - Link status**
   - **Topology Attached**
   - **WWNN - World Wide Node Name**
   - **Fabric WWPN**
   - **Current Speed**
   - **Configured Link Speed**
   - **Hard ALPA**
   - **Location - Controller ID number**
   - **Identifier - (hexadecimal)**
   - **Alias WWPN -**
   - **WWPN - Worldwide Port Name**
   - **Fabric WWNN**
   - **Link Type**
   - **Configured Topology**

Making Fibre Channel Port Settings

To make Fibre Channel port settings:

1. Click the **System** button to go to System mode. *Fibre Channel nodes are managed in System mode.*
2. Click the **Device** tab.
3. Click the **Fibre Channel Management** icon.
4. Click the **Port** tab.
5. Click the Fibre Channel port you want to access and click the **Settings** button.
6. Make these changes as required:
   - Choose a configured link speed from the drop-down menu.
     The choices are Auto (default), 4 Gb/s, 8 Gb/s and 16 Gb/s.
   - Choose a topology from the drop-down menu.
   - Enter a Hard ALPA in the field provided.
     Enter 255 to disable Hard ALPA.
7. Click the **Save** button.

Port Setting Information

The table below shows the type of attached topology you achieve based on your connection type and the configured topology you select.

Example 1: If you connect the VTrak E5000 to a Fibre Channel switch and choose NL-Port topology, you create a Public Loop attached topology.

Example 2: If you have a Point-to-Point attached topology, you made a direct connection (no Fibre Channel switch) and selected N-port topology.

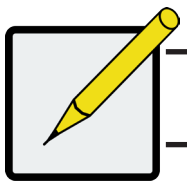

Note

In some cases, HBA settings to N-Port only work if connected to the switch. Refer to your HBA manual for more information.

Viewing Fibre Channel Port Statistics

To view Fibre Channel port statistics:

1. Click the **System** button to go to System mode. *Fibre Channel nodes are managed in System mode.*
2. Click the **Device** tab.
3. Click the **Fibre Channel Management** icon.
4. Click the **Statistics** tab.
5. Mouse over the Fibre Channel port you want to access and click the **View** button.

To clear Fibre Channel port statistics, see “Clearing Statistics” on page 90.

Viewing Fibre Channel Initiators on the Fabric

To view a list Fibre Channel initiators on the fabric:

1. Click the **System** button to go to System mode. *Fibre Channel nodes are managed in System mode.*
2. Click the **Device** tab.
3. Click the **Fibre Channel Management** icon.
4. Click the **Initiators on Fabric** tab.

Also see “Viewing Initiators” on page 212.

To add a Fibre Channel initiator in the list, select it and click on the **Add to Initiator List** button. This is the method described in “Adding a Fibre Channel Initiator” on page 213.

Viewing Fibre Channel Logged-in Devices

Logged-in devices refers to all Fibre Channel devices currently logged into the VTrak E5000. The device list includes:

- **Fibre Channel ports**
- **Fibre Channel switches, if attached**
- **Fibre Channel initiators**

To view a list Fibre Channel logged-in devices:

1. Click the **Device** tab.
2. Click the **Fibre Channel Management** icon.
3. Click the **Logged In Device** tab.

Also see “Viewing Initiators” on page 212.

To add a Fibre Channel initiator in the list, select it and click on the **Add to Initiator List** button. This is the method described in “Adding a Fibre Channel Initiator” on page 213.

Viewing Fibre Channel SFPs

The term SFP refers to Small Form Pluggable transceivers used in Fibre Channel ports. The SFPs convert electrical signals to optical signals and send them over the Fibre Channel fabric, where another transceiver converts the optical signal back to an electrical signal again.

To view a list Fibre Channel SFPs:

1. Click the **Device** tab.
2. Click the **Fibre Channel Management** icon.
3. Click the **SFP** tab.

SFP information includes:

- **Fibre Channel port ID**
- **Controller ID**
- **Connector type**
- **Transceiver type**
- **Transceiver code**
- **Vendor name**

MANAGING WITH THE CLI

Making a Serial Connection

Before you begin, be sure the RJ11-to-DB9 serial data cable is connected between the Host PC and the VTrak E5000 enclosure, and that both machines are booted and running.

Serial port on the VTrak E5000f controller

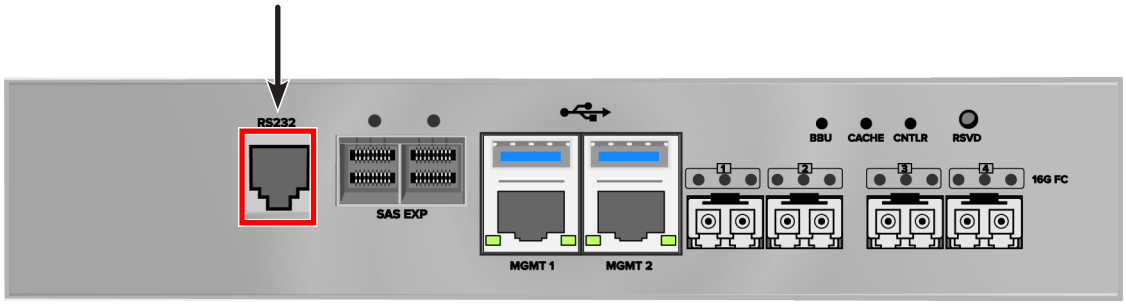

Then do the following actions:

1. Change your terminal emulation program settings to match the following specifications:
    - Bits per second: 115200
    - Data bits: 8
    - Parity: None
    - Stop bits: 1
    - Flow control: none
2. Start your PC's terminal VT100 or ANSI emulation program.
3. Press Enter once to launch the CLI.

Logging Into the CLI

1. At the Login prompt, type the user name and press Enter.

   The default user name is ***administrator***.

2. At the Password prompt, type the password and press Enter.

   The default password is ***password***.

   The CLI screen appears.

   CLI administrator prompt

Table of Supported Commands

The table below and on the following pages provides a brief description of the CLI commands available on the VTrak E5000 Series.

Command	Action
help	When used alone will display this menu. When used in conjunction with a command (example: help array) it will display help information for that particular command.
?	This can be used in place of the help command or optionally can be used as a switch for a command (example: array -?) to provide command usage.
about	View utility information.
array	View or edit array information. Create, edit, or delete logical drives in an existing array. To physically locate an array in an enclosure. Accept an incomplete array condition.
assn	View, create, or delete associations between logical drives.
battery	View battery information or to recondition a battery.
bbm	View or clear the BBM defect list of the specified configured physical drive.
bga	View status of all current background activities. Enable or disable relevant background activities. Modify the background task rate for each of the background tasks.
bgasched	View, add, modify or delete bga scheduled background activities.
buzz	View buzzer status, enable/disable and turn on/off buzzer.
checktable	View logical drive error tables.
clone	View logical drive clone status and progress. Start, stop a clone.
config	For express or automatic configuration. For advanced configuration please see the ‘array’ command.
ctrl	View or edit controller information and settings.
Note: Commands are NOT case sensitive.	

Table of Supported Commands (Continued)

Command	Action
date	View or edit system time.
enclosure	View or edit enclosure and SEP information and settings. Locate an enclosure via LEDs.
event	View or clear events logs.
export	Subsystems only. Export files to remote TFTP host.
factorydefaults	Restore settings to factory defaults.
fc	View or edit fc information and settings. Fibre Channel host interface product only.
import	Import files or license from remote TFTP host
init	View logical drive initialization status and progress. Start, stop, pause, or resume an initialization or a quick initialization.
initiator	View initiator list, add or delete initiator entry.
Note: Commands are NOT case sensitive.	

Table of Supported Commands (Continued)

Command	Action
ldap	View or edit LDAP information and settings.
lunmap	View the LUN mapping and masking table. Enable or disable LUN mapping and masking on Fibre Channel host interface product. Add, delete or modify an LMM entry.
logdrv	View or edit logical drive information and settings. Locate a logical drive via LEDs.
logout	Logout session for the current user.
maintenance	Enter or exit maintenance mode.
migrate	Start and monitor disk array migration process.
mp	View media patrol status and progress. Start, stop, pause, or resume media patrol.
net	View or edit Ethernet network information and settings.
ntp	View or edit NTP status and settings.
password	Modify a user's password.
pdm	View PDM status and progress. Start, stop, pause, or resume PDM process.
perfstats	Start and view performance statistics for controllers, logical drives, physical drives or ports.
phydrv	View or edit physical drive information and settings. Locate a physical drive via LEDs.
ping	Ping another system through management port.
ptiflash	Update system software and firmware through tftp server.
rc	View redundancy check status and progress. Start, stop, pause or resume redundancy check.
rb	View rebuild status and progress. Start, stop, pause, or resume a rebuild process.
Note: Commands are NOT case sensitive.	

Table of Supported Commands (Continued)

Command	Action
san	These are the primary SAN management function commands. Use to get information and status of the SAN and individual nodes and clients, modify the SAN, and to add or delete nodes and clients.
sasdiag	SAS diagnostic command.
sas	View or edit SAS host interface port information and settings. SAS host interface product only.
sc	View spare check status. Start spare check.
scsi	View or edit parallel SCSI information and settings. Parallel SCSI host interface product only.
session	View the list of active sessions.
shutdown	Shutdown or restart system.
smart	S.M.A.R.T diagnostic for physical drives.
spare	Create or modify hot spare drives.
stats	View or reset statistics.
subscription	View, modify, enable or disable event notification.
subsys	View or edit subsystem information and settings.
swmgt	View, start or stop software component.
sync	View logical drive synchronization status and progress.
syslog	View, add, delete or modify syslog remote servers.
topology	View SAS topology, the physical connections and device information. For products that support multiple enclosures only.
transit	View transition status and progress. Start, stop, pause, or resume a transition process.
ups	View or modify UPS information and status.
user	List, modify, create and delete user accounts on subsystem.
zoning	List, modify SAS zoning on subsystem.
Note: Commands are NOT case sensitive.	

Notes and Conventions

Commands and options are NOT case sensitive.

Not all extended keys are supported. However, you can use the backspace and the left and right arrow keys for command line editing. In addition, the up and down arrow keys allow scrolling through the command history buffer.

If you need context-sensitive help, type one of the following commands:

- <command> -h
- <command> -?
- help <command>

That action will display full context-sensitive help for the specific command. Each command when used alone, such as “array” will display a summary of relevant information. If more information is desired, the -v verbose mode can be used. This will provide information for all relevant aspects of that command.

Usage terminology is as follows:

- [square braces] depict an optional switch
- <arrow braces> depict user input

  Type “ | more” at the end of each command, to display info page by page

Note

CLI commands and parameters pertaining to iSCSI data networks are not applicable to the VTrak enclosures equipped with the VTrak E5000f controller. Future controller models available for VTrak E5000 Series might include a controller equipped with iSCSI ports.

about

Usage

about

Summary

Displays utility information.

array

Usage

array [-a <action>] [-d <DaId>] [-c <array count>] [-v]

array -a add [-s “<list of array params>”] [-d <DaId>] -p <PdId list>

[-c <Ld count>] [-l “<list of Ld params>”]

array -a mod -d <DaId> -s “<list of array settings>”

array -a del -d <DaId list>

array -a locate -d <DaId>

array -a accept -d <DaId> [-t <condition type>]

array -a addld -d <DaId> [-c <Ld count>] -l “<list of ld settings>”

array -a delld -l <LdId list>

array -a transport -d <DaId>

Summary

The array command is the main command for performing advanced configuration and maintenance tasks on disk arrays.

This command is used to list, create, modify, delete, and locate disk arrays. Also to add and delete logical drives.

Note that you cannot mix Hard Disk Drives (HDD) and Solid State Drives (SSD) in the same disk array.

Options

-a <action>		Specifies the action to perform.
	list	(Default) Displays a summary of all or a specific or a specified number of arrays.
	add	Add/create an array and possibly some logical drives at the same time.
	addld	Add/create a logical drive to an array that already exists.
	delld	Delete an existing logical drive from an array.
	mod	Modify settings of an existing array.
	del	Delete an existing array and all its associated logical drives.
	locate	Locate an array.
accept		Accepts the condition of an incomplete array. There are two conditions that will cause the array to report as incomplete:

- **Missing Drive**, i.e. when one or more drives are missing in the array.
- **Missing NVRAM Watermark**, the NVRAM Watermark of migration not found on the controller, however the DDF on the array indicates the migration is in progress.

  When the either of the above condition happens, all the logical drives on the array will become offline.

	The user can do an accept to acknowledge the incomplete condition of the array, and try to recover the access to the array and logical drives on the array, in other words, to bring the logical drives online. However, it is a high-risk and non-revertible operation, and it may result in data loss. Therefore, it is recommended to try to clear the condition first, e.g., to put the missing drives back; or to roam the array back to the original controller and wait until the migration completes.
transport	To gracefully make the array and the logical drives on the array offline to get ready for transport the array to another subsystem.
untransport	Cancels the transport action on the array and the logical drives. Brings the array back online in the original subsystem. Not supported for HBA products.
online	To set all the dead physical drives of array online at a time.
-d <DA ID>	The disk array ID. Valid values are 0-255. Used to specify the desired array ID when creating (add) an array. Used to specify the array ID when listing array information modifying, deleting, locating, accepting, adding or deleting a logical drive. Only one array may be specified.
-p <PD ID list>	Used to specify which physical drives are to be used in an array. Used in conjunction with -a add. PD IDs can be used singly or separated by comma. Additionally a sequential group of physical drives can be specified by placing a ~ between numbers such as 1~6. This will include physical drives 1,2,3,4,5,6.
-s “<option>=<value>”	Used to specify settings for an array. This is used when creating (add) or modifying (mod) an array. These options are comma separated.

alias=	A user specified name used to identify an array. It can be up to 32 characters long, containing alpha-numeric characters, blank spaces and underscores. The beginning and ending blank spaces will be discarded.
mediapatrol=	Used to enable or disable media patrol for this array.
enable	The default is enable.
disable	
powermanagement=	Apply controller power management setting.
enable	(Default) Controller power management setting will be applied to this disk array.
disable	Controller power management setting will not be applied to this disk array.
pdm=	Used to enable or disable PDM for this array.
enable	The default is enable.
disable	
-l “<option>=<value>”	Used to specify the settings for a logical drive when creating a logical drive during logical drive addition to an existing array (addld) or during array creation (add). These options are comma separated.
<LD ID list>	used to specify a list of Logical Drive IDs when used with the -a delld switch
ID=	Used to assign a specific ID to the logical drive if the user does not want one auto assigned. Valid values are 0-255.
Alias=	A user specified name used to identify the logical drive.

Raid=	Used to specify the RAID level of the logical drive.
0	Striping.
1	Mirroring on two drives.
5	Parity, requiring 3 or more drives.
10	Mirroring on even number of drives.
50	Striping on multiple RAID 5, requiring 6 or more drives.
6	Allow two drive failure, requiring 4 or more drives.
60	Striping on multiple RAID 6, requiring 8 or more drives.
Capacity=	Used to specify the desired capacity of the logical drive. It can be specified in megabytes (mb), gigabytes (gb) or terabytes (tb). Up to 2 decimal places are allowed to be specified. If not specified, all free capacities will be used for this logical drive.
CapacityRounding=	Enable or disable capacity rounding for logical drive creation.
enable	(Default) Enable capacity rounding.
disable	Disable capacity rounding.
Stripe=	Used to specify the stripe size of the logical drive. The possible parameters are 64KB, 128KB, 256KB, 512KB, and 1024KB. If not specified, the default 64KB will be used.
Sector=	Used to specify the desirable sector size of the logical drive. The possible parameters are 512B, 1KB, 2KB, and 4KB. It must not be greater than the Stripe size. It will be auto-adjusted not to exceed the max supported sector size of the controller, please see controller info If not specified, the default 512B will be used.
WritePolicy=	Used to specify the write policy for the logical drive.
writethru	Writes are not cached
writeback	Writes are cached

ReadPolicy=	Used to specify read policy for the logical drive.
readahead	Reads extra data to reduce read times for sequential data.
readcache	Caches the reads in case the same request is made again.
nocache	No read cache.
parity=	Used to specify the method of parity distribution for RAID 5, 50, 6, and 60.
left	Left asymmetric.
right	Right asymmetric.
Axle=	RAID 50, 3 to 32 drives per axle. RAID 60, 4 to 32 drives per axle. Range is 2 to 16 axles.
codec=	Used to specify the codec scheme for RAID 6 and 60.
p+q	
q+q	
PreferredCtrlId=	Used to specify which controller the LD is preferred for LUN affinity. Valid value is 1 or 2. If value is not specified, LUN affinity will be auto balanced.
PerfectRebuild=	Used to specify which logical drives supply to Perfect Rebuild.
enable	(Default)The maximum amount of logical drive which support perfect rebuild is 30. If you select perfectrebuild=enable, only sectors where write changes occurred are rebuilt.
disable	Perfect rebuild will not supply to this logical drive.
-c <array count>	Specifies the number of arrays to give a summary of when used with the -a list option. For example ‘array -a list -c3’ will give a summary for the first 3 arrays on that controller.

<Ld count>	Also specifies the number of logical drives to be created when used with the -a add option. If this -c option is used, all the logical drives will be created with the same settings and only one -l can be specified.
-t <condition type>	Specify the type of incomplete condition to accept. If not specified, it will accept the current incomplete condition by default.
missingdrive	The condition of missing drive in the array.
missingwatermark	The condition of missing NVRAM watermark of the array.
-v	Verbose mode. Used with -a list to display all properties of the given array.

Examples

```
array -v -c 1
 array -a add -s "alias=MyArray,mediapatrol=enable" -p 1,3,5~9
   -l "raid=5,capacity=50gb,stripe=256kb,sector=1kb"
 array -a add -p 1,3,5~9 -l "raid=5,capacity=50gb,stripe=256kb"
   -l "raid=0,capacity=100gb"
 array -a mod -d 1 -s "alias=YourArray,mediapatrol=disable"
 array -a del -d 3
 array -a locate -d 0
 array -a accept -d 2
 array -a delld -l 1
```

assn

Usage

assn [-a <action>][-t <type>][-l <LdId>][-d <TargetLdId>][-r][-c<count>][-v]

assn -a add -t <type> -l <SourceLdId> -d <TargetLdId(1,2,3...)> [-r]

assn -a del -l <SourceLdId> -d <TargetLdId>

assn -a list [-t <type>] [-l <LdId>] [-c count] [-v]

Summary

Assn is used to manage association between two logical drives, including list and delete existing associations, and create new associations.

Options

-a <action>	Which action to perform.
list	Display a list of existing associations for all or specified logical drives.
add	Create association between specified source logical drive and destination logical drive.
del	Delete existing association between specified source logical drive and destination logical drive.
-t <assn type>	What kind of association to be created.
clone	Clone association.
-l <source ld>	Source logical drive Id.
-d <destination ld>	Destination logical drive Id.
-r	Instructs to retain this association after corresponding background operation done. For clone, by default the association will not be retained.
-c <count>	Specifies the number of associations to give a summary of when used with the -a list option.
-v	Verbose mode. Used with -a list.

Examples

```
assn *shows a list of association of specified logical drive*

assn       -a add -t clone -l 0 -d 1 -r

assn       -a del -l 0 -d 1
```

battery

Usage

battery [-a <action>] [-b <batId>]

battery -a recondition -b <batId>

Summary

Battery is used to display the current status of a battery indicating the percentage of charge left.

This command is also used to recondition a battery. Reconditioning of a battery attempts to fully discharge, and then recharge it. In addition the battery will be reconditioned automatically once per month.

-a <action>	Which action to perform.
list	(Default) List information for all batteries or a specific battery unit.
recondition	Recondition a specific battery.
-b <battery ID>	Used to specify which battery in a given enclosure.

Examples

```
battery

battery -a recondition -b 1
```

bbm

Usage

bbm [-a <action>] [-p <PdId>]

bbm -a clear -p <PdId>

Summary

The bbm command displays or clears the Bad Block Map (BBM) defect list for all configured physical drives.

Options

-a <action>		Specifies the action to perform.
	list	(Default) List the BBM information.
	clear	Clears the BBM list. For configured SATA drives only.
	-p <PdId>	Specifies the physical drive id. For the -a list option, the default is all physical drives. For the -a clear option, you must specify a physical drive id.

Examples

```
bbm -p 1

bbm -a clear -p 3
```

bga

Usage

bga [-a <action>]

bga -a mod -s “<list of settings>”

Summary

The bga command displays all current background activities and makes settings for each background activity.

Options

-a <action>		Specifies the action to perform.
	list	(Default) Lists current background activities.
	mod	Makes changes to one of the settings.
-s “<option>=<value>”		Specifies which background activity settings to change.
	autorebuild=	Enable or disable auto-rebuild and auto-transition.
	enable	Auto-rebuild initiates a rebuild of an array when an unconfigured drive is inserted into the slot of a dead drive.
	disable	Auto-transition means transitioning is initiated on a used revertible spare in the following condition: 1. When the rebuild has been completed using the revertible spare, and 2. When an unconfigured drive is inserted into the slot of the dead drive which the was part of the array. or When a non-revertible spare has been inserted or created, and is applicable to the array This option affects all arrays on the subsystem.
	enable	
	disable	
mediapatrol=		Verifies the media of the array and/or spares to find bad blocks on

	physical disks before you use that block. This feature is enabled and disabled for individual arrays on a per array basis.
enable	
disable	
BBMThreshold=	(1-2048) Threshold value to trigger PDM based on the Bad Block Monitor count of reassigned blocks on the PD.
MediaPatrolThreshold=	(1-2048) Threshold to trigger PDM based on the Media Patrol count of error blocks on the PD.
<bg task>=<rate>	Background task rates determine what percentage of the IO load on the controller will be dedicated to the background task. A lower number means the task takes longer to complete, a higher number will cause the task to complete faster, all other things being equal.
rebuildrate= (low=25, medium=50, high=75)	Rebuild rate determines the rate at which rebuild will run.
low	
medium	
high	
pdmrate=	PDM rate determines the rate at which PDM will run. (low=25, medium=50, high=75)
low	
medium	
high	

transitionrate=	Transition rate determines the rate at which transition will run. (low=25, medium=50, high=75)
low	
medium	
high	
syncrate=	Synchronization rate determines the rate at which synchronization will run. (low=25, medium=50, high=75)
low	
medium	
high	
initrate=	Initialization rate determines the rate at which initialization will run. (low=25, medium=50, high=75)
low	
medium	
high	
rcrate=	Redundancy check rate determines the rate at which redundancy check will run. (low=25, medium=50, high=75)
low	
medium	
high	
migrationrate=	Migration rate determines the rate at which migration (low=25, medium=50, high=75)
low	
medium	
high	

Examples

```
bga

bga -a mod -s “autorebuild=enable,rebuildrate=high,syncrate=low”
```

bgasched

Usage

bgasched -a <action> -t <type> -s <list of settings>

bgasched -a add -t <type> -s <list of settings>

bgasched -a mod -t rc -i <RC scheduler id> -s <list of settings>

bgasched -a mod -t <type> -s <list of settings>

bgasched -a del -t <type>

bgasched -a del -t rc -i <RC scheduler id>

Summary

bgasched is used to display all scheduled background activities as well as to allow the user to add, modify or delete date and time of the scheduled activities.

-a <action>		Which action to perform.
	list	(Default) Displays information of BGA scheduler.
	add	Create a new BGA scheduler. If exists RC scheduler, cannot add RC scheduler with all LDs. The max number of RC scheduler is 4. Only 1 for other schedulers.
	mod	Modify a exist scheduler. Can not change ldid parameter of an exist RC scheduler to all LDs.
	del	Delete a exist scheduler.

-t <type>	Specifies what type of scheduler.
mp	Media Patrol Schedule.
rc	Redundancy Check Schedule.
br	Battery Reconditioning Schedule.
sc	Spare Drive Check Schedule.
-i <RC scheduler id>	Specifies the RC scheduler ID.
	It's used for list/modify/delete RC scheduler. If the option is not specified, assumed to all.
-s "<option>=<value>"	Used to specify which BGA scheduler settings to change.
status=	Specifies status type of scheduler.
enable	Enable a scheduler.
disable	Disable a scheduler. The default is disable.
starttime=	Used to specify start time of scheduler in the following format
	hh:mm where hour's range is 0-23, minute's range are 0-59.
	The default is 20:00 for MP, 22:00 for RC and SC, 02:00 for others.
recurtype=	Specifies recurrence type of scheduler.
daily	
weekly	The default is weekly.
monthly	
recurInterval=	Specifies recurrence Interval.
	This option is for Daily and Weekly recurrence type.
	For Daily type, the range is 1-255. (default is 1)
	For Weekly type, the range is 1-52.
	For Weekly type, the default is 4 for MP, 2 for RC,1 for others.
dow=	Day of Week. This is for Weekly or Monthly recurrence type scheduler.

	Regarding Monthly type, if daypattern (see below) is day of week, it will be used.
	For Weekly, the range is [Sun|Mon|Tues|Wed|Thur|Fri|Sat]. For multiple values, divide with spaces. The default is 'Fri' for MP, 'Wed' for RC, 'Tues' for SC,'Sun Mon Tues Wed Thur Fri Sat' for others.
	For Monthly, the range is [Sun|Mon|Tues|Wed|Thur|Fri|Sat].The default is Sat.
daypattern =	Specifies the daypattern type for Monthly recurrence type scheduler.
dom	Day of month.
dow	Specific day of week.
dom=	Day of Month, for 'dom' daypattern type. The range is 1~31. The default is 1.
wom=	Week ordinal, for 'dow' daypattern type. The range is (1st| 2nd |3rd |4th |Last). The default is 1st.
month=	Months. The range are 1~12 divided by space or '~'. The default is 1~12.
startfrom=	Start day of range of occurrence in the following format

		mm/dd/yyyy where month's range is 1-12, day's range is 1-31.
		The default is current date of system.
	endon=	Used to specify end time of scheduler.
	0	(Default) No end time.
	n	An integer N indicates after N times.
	mm/dd/yyyy	End date, month's range is 1-12 and day's range is 1-31.
	autofix=	Fix inconsistent data.
	enable	
	disable	The default is disable.
	pause=	Pause on error.
	enable	
	disable	The default is disable.
	ldid=	The list of LDID.
		For add action, if the option is not specified, assumed to for all LDs.
		For multiple value, divided by space or '~'.
	-v	Verbose mode. Used with -a list.

Examples

```
 bgasched
 bgasched -a mod -t rc -i 1 -s "status=disable,ldid=1 3~5 7"
 bgasched -a add -t mp -s "recurtype=monthly,daypattern=dow,wom=2nd,dow=Su
n,month= 1 3~6, endon=10"
 bgasched -a add -t sc -s "recurtype=weekly,dow= Mon Wed Fri,starttime=12:0
0,endon=1/1/2010"
```

buzz

Usage

buzz [-a <action>]

buzz -a list buzz -a enable buzz -a disable buzz -a on

buzz -a off

Summary

The buzz command displays the status of the buzzer, and enables, disables, turns on or turns off the buzzer.

Options

-a <action> Specifies the action to perform.

list	(Default) List the status of the buzzer.
enable	Enable the buzzer.
disable	Disable the buzzer.
on	Turn on the buzzer.
off	Turn off the buzzer.

checktable

Usage

checktable [-t <tableType>] -l <LdId>

Summary

The checktable command displays the error check tables of a logical drive.

Options

-t <tableType>	Specifies which error table to display. The default displays all tables.
rct	Displays the read check table.
wct	Displays the write check table.
ibt	Displays the inconsistent block table.
-l <LdId>	Specifies the logical drive ID.

Examples

```
checktable -l 10 -t rct

checktable -l 10
```

clone

Usage

clone [-a <action>] [-l <SourceLdId>] [-d <TargetLdId>] [-r]

clone -a start -l <SourceLdId> [-d <TargetLdId(1,2,3...)>] [-r]

clone -a stop -l <SourceLdId> [-d <TargetLdId>]

clone -a list [-l <LdId>]

Summary

This command allows the user to start or stop a Clone as well as to check on the progress of a running Clone.

There are two methods to start a Clone. One is specify the destination logical drive to perform clone, another is specify an existing array.

Options

-a <action>	Which action to perform.
list	(Default) Displays the current active Clone(s) and their status(es).
start	Start a Clone.
stop	Stop a Clone.
-l <source ld>	Specifies which source logical drive to perform clone action on.
-d <destination ld>	Specifies which destination logical drive to perform a clone action on. For start, if multiple destinations are specified, a maximum of 8 allowed. If not specified, all existing associations on the source logical drive (specified by -l) will be started or stopped.
-r	Instructs to retain this association after corresponding background operation done. For clone, by default the association is not retained.

The following are used to specify a existing array to perform clone

-s “<option>=<value>”

id=<array id>	Specifies an array id.
Raid=	Used to specify the RAID level of the logical drive.
0	Striping.
1	Mirroring on two drives.
5	Parity, requiring 3 or more drives.
10	Mirroring on even number of drives.
50	Striping on multiple RAID 5, requiring 6 or more drives.
6	Allow two drive failure, requiring 4 or more drives.
60	Striping on multiple RAID 6, requiring 8 or more drives.
Axle=	Used to specify the number of axles for RAID50 and RAID60.
-c <Ld count>	Specifies the number of logical drives to be created.

Examples

```
clone

clone -a start -l0 -d1

clone -a stop -l0 -d1

clone -a start -l 0 -s “id=1,raid=5” -c 2
```

config

Usage

config -a auto

config -a expr [-r y|n] [-c y|n] [-p y|n] [-m y|n] [-s y|n] [-t <AppType>] [-l <NumLd>]

Summary

The config command has two options, Automatic (auto) and Express (expr).

Automatic configuration takes all available unconfigured physical drives to create an optimized disk array following a default set of parameters. There are no options.

Express configuration takes your input, creates one or two arrays, and spreads their capacity evenly over all of the logical drives that you specify.

The redundancy option creates redundant logical drives (RAID 1, 10, 5, 50, 6, or 60).

The capacity option enables optimizes the logical drives for capacity. The performance option optimizes the logical drives for performance.

If you choose all three options, redundancy gets highest priority and capacity gets lowest priority.

Note that you cannot combine HDDs and SSDs in the same disk array. If your system has both type of drives, it will create separate disk array/logical drive sets for each type of physical drive.

Options

-a <action>	Specifies the action to perform.
auto	Automatic configuration with no options. Creates an optimized disk array. One or more logical drives are created automatically.
expr	Express configuration. RAID level is dependant on the options chosen.
-r <y|n>	Selects the redundancy option.
-p <y|n>	Selects the performance option.
-c <y|n>	Selects the capacity option.
-s <y|n>	Includes a spare drive in the array. Note: Requires 5 or more unconfigured physical drives.
-t <AppType>	Specifies the intended application for this array.
video	Sequential large block reads.
data	Random read/write mix, small to medium sized IO.
log	Sequential small block write.
other	Random read/write mix, small to medium sized IO.
fileserver	Random read/write mix, small to medium sized IO.
-l <num of LDs>	Specifies how many logical drives to include in the configuration. Array capacity is divided evenly among the logical drives.

Examples

```
config -a auto

config -a expr -ry -p y -c n -sy -t data -l2
```

ctrl

Usage

ctrl [-a <action>] [-i <ctrlId>] [-c <ctrl count>] [-v]

ctrl -a mod [-i <ctrlId>] -s “<list of settings>”

ctrl -a clear [-i <ctrlId>] [-t <condition type>]

Summary

The ctrl command displays controller information and changes controller settings.

Options

-a <action>	Specifies the action to perform.
list	(Default) Lists controller information.
mod	Changes controller settings.
clear	Clears controller conditions.
-i <ctrl ID>	Specifies the controller ID. For high availability products, controller ID is required when setting alias of controller.
-c <ctrl count>	Controller count. Required for information on multiple controllers.
-s “<option>=<value>”	Specifies which settings to change.
alias=	A user-specified name for the controller. Up to 48 characters long, alpha- numeric characters, blank spaces and underscores The beginning and ending blank spaces are discarded.

Following setting is shared if there are dual controllers:

coercion=	Enables or disables disk coercion. Disk coercion will truncate the size of the physical drives. Makes different size drives appear to be the same size. For example, a 90.1 GB drive would appear as the same size as an 89.8 GB drive. Important when using drives of different manufacturers for rebuilds or as hot spares. Coercion settings are shared if there are dual controllers:
enable	
disable	
coercionmethod=	The method of coercion.
GBTruncate	Truncates the drive to the nearest 1-billion byte boundary.
10GBTruncate	Truncates the drive to the nearest 10-billion byte boundary.
GrpRounding	Truncates the drive using an intelligent algorithm. This allows the maximum amount of usable space while at the same time attempting to keep drives in the same size group the same size. For example a 253 GB drive would appear the same size as a 248 GB drive.
TableRounding	This uses a pre-defined coercion table to determine how much will be truncated.
smart=	Enables or disables polling drive SMART status.
enable	
disable	
smartpollinginterval=	(1 - 1440) Sets the time interval in number of minutes to poll the drive SMART status.

cacheflushinterval=	(1-12) Sets the time interval in seconds to flush the controller writeback cache.
migrationstorage=	To set which place to store the migration watermark.
ddf	Uses the DDF area on the physical drives of the disk array.
nvram	Uses the NVRAM on the controller.
lunaffinity=	To enable or disable LUN affinity, allowing LD access only to certain controller. For products that have high availability only.
enable	
disable	
alua=	To enable or disable asymmetric logical unit access.
enable	
disable	
pollinterval=	(15 - 255) Sets interval in seconds to poll enclosure SEP information.
adaptivewbcache=	Enables or disables adaptive writeback cache.
enable	Writeback logical drives will change the write policy based on the availability of protection. If BBU or UPS is available, the write policy is retained as Writeback, otherwise the policy is switched to Writethru.
disable	The write policy of the writeback logical drives are not changed irrespective of the availability of BBU or UPS.

hostcacheflushing=	Subsystems only. To enable or disable host cache flushing. When enabled, **`synchronize cache scsi`** command from host is supported. *Note that this is for high availability products only.*
enable	
disable	
forcedreadahead=	Enables or disables forced read ahead caching. For high availability products only.
enable	
disable	
powersavingidletime=	After an HDD has been idle for the set period of time, parks the read/write heads. Set the time interval in number of minutes. Valid values are 0(never), 15, 30, 60(= 1 hour)..1440(=24 hours).
powersavingstoppedtime=	After an HDD has been idle for the set period of time, Spins down the disk (stops rotation). Set the time interval in number of minutes. Valid values are 0(never), 15, 30, 60 (= 1 hour)..1440 (=24 hours).
VAAIsupport=	To enable or disable VAAI support.
enable	
disable	
ssdsmartpollinterval=	(0...255) To set the time interval in number of hours to poll the drive SMART status. Default value is 24 hours, disabled polling when the value is 0.
ssdlifethreshold=	(5...25) To set the life threshold in percentage. Default is 10.

ssdremainreserveblockthreshold=	(5...25) To set the remaining reserved block threshold in percentage. Default is 10.
-t <condition type>	Used to specify the type of condition to clear. It is valid only when the command action is “clear”.
watermark	Watermark, the only supported condition for now. It is used together with -a clear to clear the orphan migration watermark in the controller NVRAM. This will work only when migration storage is set to NVRAM prior to starting migration.
-l	Display local controller’s id that CLI runs through its serial port.
-v	Verbose mode. Used with -a list.

Examples

```
ctrl

ctrl -v

ctrl -l

ctrl -a mod -i 1 -s “alias=ctrl1, coercion=enable”

ctrl -a mod -s “powersavingstoppedtime=180”
```

date

Usage

date

date -a mod [-d <date>] [-t <time>] [-z <timezone>]

Summary

The date command allows the user to view and modify the system time.

Options

-a <action>	Which action to perform.
list	(Default) Displays the current system time.
mod	Modifies the current system time.
-d <date>	Used to specifies date in the following format: yyyy/mm/dd where month's range is 1-12 and day's range is 1-31.
-t <time>	Used to specifies time in the following format: hh:mm:ss where hour's range is 0-23, minute's and seconds' range are 0-59.
-z <timezone>	Specify the time zone. The time zone range is GMT-12 ~ GMT12.

Examples

```
date

date -a mod -d 2004/02/25 -t 14:50:05

date -a mod -z GMT-8
```

enclosure

Usage

enclosure [-a <action>] -v

enclosure -a mod -s <list of settings>

enclosure -a locate [-t <FRU type> -f <FRU id>]

Summary

The enclosure command provides status and information about the various components of the enclosure unit. It is also used to set thresholds for temperature and polling. In addition when using the -v option all VPD (Vendor Provided Data) will be displayed.

Options

-a <action>	Which action to perform.
list	(Default) Displays information and status of the enclosure.
mod	Allows the user to modify settings when coupled with the -s switch.
locate	Allows the user to locate an enclosure by flashing LEDs
-e <encl id>	Enclosure ID. The default value is 1 if unspecified. For list action, the default is for all enclosures if unspecified.
-i <sensor id>	IOM Sensor ID. For controller temperature warning (1..6) and critical threshold. If unsepecified, only the IOM sensor 1,4 will be set, the IOM sensor 2,3,5,6 will not changed.
-s “<option>=<value>”	Used to specify which Enclosure settings to change.
tempwarning=	(47 - 51) Temperature, displayed in Celsius, that the SEP will consider as a warning threshold.
tempcritical=	(57 - 61) Temperature, in Celsius, that the SEP will consider as a critical threshold.

ctrltempwarning=	Controller temperature, displayed in Celsius, that the controller will consider as a warning threshold. (61-65) For IOM sensor 1,4. (66-70) For IOM sensor 2,5. (71-78) For IOM sensor 3,6.
ctrltempcritical=	Controller temperature, displayed in Celsius, that the controller will consider as a critical threshold. (68-72) For IOM sensor 1,4. (73-77) For IOM sensor 2,5. (81-88) For IOM sensor 3,6.
-t <FRU type>	Used with action locate to indicate which type of FRU to locate. If -t is not specified, it indicates to locate the enclosure.
ctrl	To locate controller.
cooling	To locate cooling unit. It only works with SAS type enclosure.
psu	To locate power supply unit. It only works with SAS type enclosure.
-f <FRU id>	Used with action locate and -t <FRU type> option to indicate which FRU to locate. The valid values for FRU id are 1,2,3 and 4.
-v	Verbose mode. Used with -a list. VPD information will also be displayed when using this switch.

Examples

```
enclosure

enclosure -v

enclosure -a mod -s “tempwarning=40,tempcritical=70”
```

event

Usage

event [-a <action>] [-l <location>] [-i <SeqNo>] [-c <event count>] [-v]

event -a clear [-l <location>]

Summary

The event command displays and clears the RAM and NVRAM event logs.

Options

-a <action>		Specified the action to perform.
	list	(Default) Displays the events for the specified location. RAM events are displayed if no location is specified.
	clear	Clear events for a specified location.
-l <location>		Specifies the location from which to display or clear events.
	ram	(Default) All events are stored in RAM. These events are lost after rebooting.
	nvram	Some events are also stored in NVRAM. These events remain after rebooting and are a subset of the RAM events.
	bbu	These events are stored in the Battery backed area of the RAM.
-i <sequence ID>		Specifies a specific event by its sequence number. This is a starting point. Requires the -a list option. You can use the -c option.
-c <event count>		Specifies the number of events to retrieve when displaying events.
-v		Verbose mode. Requires the -a list option.

Examples

```
event

event -v

event -l nvram

event -a clear -l nvram

event -c 200

event -a list -i 852 -c 200
```

export

Usage

export -t <fileType> [-s <tftpServer>] [-p <port>] [-x <fileExt>] -f <fileName>

Summary

The export command exports certain types of configuration files to a TFTP server.

Options

-t <file type>	Specifies the type of file to export.
userdb	User database file.
configscript	Configuration script.
servicereport	System service report file in compressed HTML format.
-f <file name>	Specifies the name of the file to be exported.
-x <file ext>	Specifies the type of the file.
txt	Saves service report as a text file.
html	Saves service report as a compressed HTML file (default).
-s <TFTP server>	Specifies tftp server's IP or host name.
-p <port num>	The port number of the TFTP server. Default is 69.

Examples

```
export -t userdb -s 192.168.1.1 -f userdb.bin

export -t servicereport -s 192.168.1.1 -f servicereport

export -t servicereport -s 192.168.1.1 -f servicereport -x txt
```

Note: Make sure that you have a file named <fileName>.<html|txt>.gz (e.g. servicereport.txt.gz or servicereport.html.gz) created on the specified TFTP server, with write permissions.

factorydefaults

Usage

factorydefaults -a <action> -t <type>

Summary

The factorydefaults command restores specified settings to the factory default values.

Options

-a <action>	Specifies the action to perform.
restore	Restore the factory default settings.
-t <type>	Specifies the type of settings to restore.
all	All settings.
allfw	All firmware settings.
allsw	Subsystems only. All software settings.

The following are individual Firmware settings:

bga	Background activity settings.
ctrl	Controller settings.
encl	Enclosure settings, including temperature thresholds, buzzer, etc.
fc	fc port settings. Fibre Channel host interface product only.
netmgmt	Subsystems only. Network settings of management ports.
phydrv	Physical drive settings.
sas	SAS host interface port setting. SAS host interface port product only.
scsi	Parallel SCSI channel settings. Parallel SCSI host interface product only.

subsys	Subsystem settings.

The following are individual Software settings:

bgasched	bga scheduler settings.
service	service startup type settings.
snmp	snmp settings.
telnet	telnet settings.
ssh	ssh settings.
email	email settings.
netsend	netsend settings.
ntp	ntp settings.
user	user settings.
ups	ups manager configuration settings.

Examples

```
factorydefaults -a restore -t phydrv

factorydefaults -a restore -t all
```

fc

Usage

fc [-a <action>] [-t <Type>] [-i <CtrlId>] [-p <PortId>] [-v]

fc -a mod -t <Type> -i <CtrlId> -p <PortId> -s “<list of settings>”

Summary

The fc command is used to view and modify Fibre Channel information and settings.

Options

-a <action>	Which action to perform.
list	(Default) Gives summary information about Fibre Channel status.
mod	Modify Fibre Channel settings.
reset	Reset Fibre Channel port(s)
-t <type>	Specifies what type of information to display or modify.
node	Display Fibre Channel node information.
port	(Default) Specifies Fibre Channel port as the device type to display or modify information.
SFP	Display port SFP (Small Form Factor Pluggable) information.
stats	Display port statistics information.
device	Display port logged in devices information.
fabricdevices	Display devices list from fabric switch.
-i <ctrlId>	Controller Id. Default to be all available controllers for listing if -i is not specified. Default to be controller 1 for modifying if -i is not specified.
-p <port id>	Port Id. Default to be all ports for listing if -p is not specified. Default to be port 1 for modifying if -p is not specified.
-s “<option>=<value>”	Specifies Fibre Channel settings to change.

linkspeed=	Fibre Channel link speed.
4gb	4 GB/s
8gb	8 GB/s
16gb	16 GB/s
auto	Automatic
topology=	Fibre Channel topology method.
nlport	NL-Port
nport	N-Port
auto	Automatic
hardalpa=	Hard Arbitrated Loop Physical Address (ALPA)
0..255	Value 255 will disable hard ALPA.
-v	Verbose mode. Used with -a list.

Examples

```
fc

fc -t port -v

fc -a mod -t port -p 1 -s “linkspeed=8gb”
```

import

Usage

import -t <file type> -s <TFTP server> -f <file name> -p <port num> -i

Summary

The import command is used to import files from a remoter TFTP host.

Options

-t <file type>

userdb	User database file.
configscript	Configuration script.
-s <TFTP server>	Specifies tftp server's IP or host name.
-f <file name>	Specifies the name of the file to import.
-p <port num>	The port number of the TFTP server. Default is 69.
-i	Get format validation information about imported file only. File is not really applied to subsystem yet.

Examples

```
import -t userdb -s 192.168.10.168 -f userdb.xml
```

init

Usage

init [-a <action>] [-l <LdId>]

init -a start -l <LdId> [-q <size>] [-p <pattern>]

init -a stop -l <LdId>

init -a pause -l <LdId>

init -a resume -l <LdId>

Summary

The init command starts, stops, pauses, and resumes a logical drive initialization. A full initialization writes to the entire logical drive space and can takes several minutes, depending on the size of the logical drive.

A quick initialization writes to the first and last few megabytes of the logical drive. Typically, a quick initialization is completed in a few minutes.

Options

-a <action>		Specifies the action to perform.
	list	Displays a list of the initialization processes in progress or paused and their status. The default action.
	start	Start an initialization.
	stop	Stop an initialization.
	pause	Pause an initialization.
	resume	Resume an initialization.
-l <LD ID>		Specifies the logical drive to be initialized.
-q <size>		(1-1024) Specifies the amount of data in megabytes (MB) for a quick initialization.
-p <pattern>		Specifies the pattern for a full initialization. The pattern can range

from 1 to 128 bytes (HEX string), and is padded to even number of bytes, such as, fff padded to 0fff.

Pattern is not supported for quick initialization.

Examples

```
init

init -a stop -l0

init -a start -l0 -p5a5a0101
```

initiator

Usage

initiator [-a <action>] [-i <Index>] [-c <Count>]

initiator -a add [-i <Index>] -n <Name>

initiator -a del -i <Index>

Summary

Use this to display information about the current initiator list as well as to add or delete an initiator.

Options

-a <action>	Which action to perform.
list	(Default) Displays the current initiator list.
add	Add an initiator to the list.
del	Delete an initiator from the list.
-i <Index>	(0-2047) Used to specify the index of the initiator. For -a list option, it is the starting index and may be used with -c option. For other options, it is the specific index.

Caution

For -a add option, if the index specified is already in use, the existing initiator name is overwritten with new name.

-c <Count>	Used to specify the number of initiators to be listed. Only used with -a list option.
-n <Name>	Used to specify the name of the initiator. For a Fibre Channel host interface product, the name should be the initiator's WWPN in hex format, e.g. aa-bb-cc-dd-ee-ff-11-22 For a SAS host interface product, the name should be the initiator's SAS address in hex format, e.g. aa-bb-cc-dd-ee-ff-11-22 For slot based lun mapping product, the first byte is slot id. For example, for slot 2, the name is 02-00-00-00-00-00-00-00

Examples

```
initiator -i 1 -c 2

initiator -a add -n iqn.vendorcompany.com
```

ldap

Usage

ldap [-a <action>] [-t <rolemap>] [-v]

ldap -a add -t rolemap [-s “<externalrole=..., privilege=...>”]

ldap -a mod [-s “<list of settings>”]

ldap -a del -t rolemap [-s “<externalrole=...>”]

Summary

The ldap command is used to view and modify LDAP configurations.

It is used to list, add, modify, and delete role mapping rules as well as to enable or disable the external LDAP.

The default is disabled external LDAP.

Options

-a <action>		Which action to perform.
	list	List LDAP settings.
	add	Add role mapping.
	mod	Modify LDAP settings or role mapping.
	del	Delete role mapping.
	test	Test LDAP bind operation to verify current configuration. Note this operation only validates server IP and port attributes.
-t <type>		The type of external LDAP setting to configure.
	auth	LDAP authentication configuration. Default value.
	rolemap	Role mapping rules. Maps external roles, such as an LDAP group, to internal privileges.

-s “<option>=<value>”	Specify the external LDAP settings or role mapping.
auth settings	
ldap=	Enables or disables LDAP authentication. The default value is disable.
enable	
disable	
emailforevent=	Specify to enable or disable event subscription for LDAP authenticated user. The default value is disable.
enable	
disable	
server=	LDAP server hostname or IP address. The default server is 127.0.0.1.
servertype=	LDAP server type. The default type is Open Directory.
ad	Windows Active Directory.
od	Open Directory.
ol	Linux Open LDAP.
port=	LDAP server port number. The default value is 389.
timeout=	The timeout allowed for communication with LDAP server.
1..30	The default value is 10 seconds.
binddn=	User on LDAP server that is permitted to search the LDAP directory within the defined search base.
bindpw=	Password of binddn.
uidattribute=	Attribute name that stores user’s UID property in LDAP server.
basedn=	Base distinguished name (DN) to use for common LDAP search base.

anonymousbind=	Allows the system to bind to an LDAP server without providing the Bind DN and password.
enable	The default value is enable.
disable	
objectclass=	Object class used to filter information from LDAP search result for a specific user. Can be configured only when event subscription is enabled. See emailforevent above.
attributeoffullname=	Attribute name that stores user's full name property in LDAP server. Can be configured only when event subscription is enabled. See emailforevent above.
attributeofemail=	Attribute name that stores user's email address propertyin LDAP server. Can be configured only when event subscription is enabled. See emailforevent above.
rolepolicy=	Specifies role policy. The default value is "default".
default	All LDAP users have privileges specified by default privilege parameter.
explicit	LDAP users have privileges according to role mapping configurations.
defaultprivilege=	Specifies the privilege level for creating LDAP user. Only takes effect when rolepolicy is set to default.
super	Super user has full control.
power	Power user cannot modify users nor delete configurations.
maintenance	Maintenance user can only perform background tasks.
view	View user can only view.

basednofgroup=	Similar to base DN above. It defines the base distinguished name to use for common LDAP lookup for group information. Only needed when rolepolicy is set to explicit.
objectclassofgroup=	Defines object class to store group entry in LDAP server. Only needed when rolepolicy is set to explicit.
attributeofgroupid=	Attribute name which stores group ID. If the value of this field5s is mapped to an internal privilege, members of this group have the corresponding privilege when they log in. Only needed when rolepolicy is set to explicit.
rolemap settings	
externalrole=	External Role name. For LDAP, usually is LDAP the user group id.
privilege=	Specifies the privilege to grant for external roles.
super	Super user has full control.
power	Power user cannot modify users nor delete configurations
maintenance	Maintenance user can only perform background tasks.
view	View user can only view.

Examples

```
ldap

ldap -v

ldap -a mod -s”server= xxx.xxx.xxx.xxx”

ldap -a mod -t auth -s “server=xxx.xxx.xxx.xxx,

binddn=username, bindpw=password, basedn=’dc=xxx,dc=domain,dc=com’,

objectclass=person, attributeoffullname=displayName,

attributeofemail=mail”

ldap -a add -t rolemap -s “externalrole=rolename, privilege=view”

ldap -a del -t rolemap -s “externalrole=rolename”

ldap -a disable -l external
```

lunmap

Usage

lunmap [-a <action>] [-i <InitiatorId>] [-r <CtrlId>] [-p <PortId>] [-c <Count>]

lunmap -a addld -i <InitiatorId> [-l <LdIdList>] [-m <LunMap>]

lunmap -a delld -i <InitiatorId> [-l <LdIdList>]

lunmap -a add [-i <InitiatorId>] -n <Name> [-l <LdIdList>] [-m <LunMap>]

lunmap -a del -i <InitiatorId>

lunmap -a enable

lunmap -a disable

Summary

The lunmap command displays information about the current LUN mapping and masking (LMM) table information and enables you to add, modify, and delete LMM entries. LMM can be enabled or disabled.

Options

-a <action>	Which action to perform.
list	(Default) Displays the current LMM table.
enable	Enables LMM.
disable	Disables LMM.
add	Adds an LMM entry and its LUN maps to the table.
del	Deletes an LMM entry from the table.
addld	Adds or modifies an LUN map for an existing LMM entry.
delld	Deletes a LUN map for an existing LMM entry.
mod	Specifies LUN mapping type.

-r <CtrlId>	Specifies the Ctrl ID for a port-based LMM entry. Valid only for Fibre Channel host interface.
-p <PortId>	Specifies the Port ID for a port-based LMM entry. Valid only for Fibre Channel host interface.
-i <InitiatorId>	(0-2047) Specifies the initiator ID for an initiator based LMM entry. For -a list option, it is the starting index. May be used with -c option.
-c <Count>	Specifies the number of LMM entries to be listed. Only used with -a list option.
-n <Name>	Specifies the initiator name. For the Fibre Channel host interface, name is the initiator's WWPN in hex format, such as aa-bb-cc-dd-ee-ff-11-22. For the SAS host interface, the name is the SAS address, such as aa-bb-cc-dd-ee-ff-11-22. For slot-based LUN mapping, the first byte is the slot ID. For example, slot 2, the name is 02-00-00-00-00-00-00-00.
-l <Ld ID list>	(0-1023) Specifies the logical drive IDs.
-m <LUN map list>	(0-1023) Specifies the LUN mapping values. Please check the maximum number of LUNs supported by host OS.

-s “<option>=<value>” Specifies settings for LMM entry. Modifies an LMM entry.

type=

initiator For initiator-based LUN mapping.

port For port-based LUN mapping.

Examples

```
lunmap -i 1 -c 2

lunmap -a addld -i 1 -l 2 -m 2

lunmap -a delld -i 1 -l 2

lunmap -a enable

lunmap -a add -n iqn.promise.com -l 0,1 -m 0,1
```

logdrv

Usage

logdrv [-a <action>] [-l <LdId>] [-c <Ld count>] [-v]

logdrv -a locate -l <LdID>

logdrv -a mod -l <LdId> -s “<list of ld settings>”

Summary

The logdrv command displays information about the logical drives and is used to make changes on logical drive settings.

To create a logical drive please see the array command.

Options

-a <action>		Specifies the action to perform.
	list	(Default) Displays a summary of one or more logical drives.
	mod	Changes logical drive settings.
	locate	Locates a logical drive within the enclosure by flashing drive carrier LEDs.
-l [<LD ID>]		Logical drive ID.
-c [<LD count>]		Logical drive count. Requires the -a list option.
-s [“<option>=<value>”]		Specifies the logical drive settings to change.
alias=		A user-specified name for the logical drive. Up to 32 characters, containing alpha-numeric characters, blank spaces and underscores. Beginning and ending blank spaces are discarded.

WritePolicy=	Specifies logical drive write policy.
writethru	Writes are not cached.
writeback	Writes are cached.
	Note: Cannot be set if ReadPolicy is set to "nocache."
ReadPolicy=	Specifies logical drive read policy.
readahead	Reads extra data to help reduce read times of sequential data.
readcache	Caches reads in the case the same request is made again.
nocache	No caching algorithm.
PreferredCtrlId=	Specifies which controller the LD is prefers for LUN affinity.
	Valid value is 1 or 2.
PerfectRebuild=	Used to specify which logical drives supply to Perfect Rebuild.
disable	Perfect rebuild will not supply to this logical drive.
-v	Verbose mode. Used with -a list.

Examples

```
logdrv

logdrv -v

logdrv -a mod -l0 -s"readpolicy=readahead"

logdrv -a locate -l2
```

logout

Usage

logout

Summary

The logout command is used to logout the current user from the session.

Examples

```
logout
```

migrate

Usage

migrate [-a <action>] [-d <DaId>]

migrate -a start -d <DaId> -p <PdIds> -l <LdSettings>

Summary

The migrate command allows the user to migrate logical drives inside a particular disk array. The supported migrations are online capacity expansion, RAID level migration and stripe size migration.

Options

-a <action>	Which action to perform.
list	(Default) Displays the migration status of specified disk array. If no array ID specified, all migration status will be displayed.
start	start a specific migration progress.
-d <DA ID>	Used to specify the array ID for migration.
-p <PD ID list>	Used to specify which physical drives are to be added in an array.
-l “<option>=<value>”	Used to specify settings for logical drive migration.
id=	(Required) Specifies the logical drive ID.
capacity=	Specifies the new logical drive capacity. Not to specify it unless intending to expand the capacity
capacityrounding=	Enable or disable capacity rounding for logical drive
migration	
enable	(Default) Enable capacity rounding.
disable	Disable capacity rounding.
raid=	Specifies the new logical drive RAID level.
axle=	Specifies the axle number for hybrid RAID Levels when RAID Level is changed.
stripe=	Specifies the new logical drive stripe size. This is currently not supported and is ignored.

Examples

```
migrate -d 1

migrate -a start -d 1 -p 10 -l “id=0,capacity=10gb”
```

mp

Usage

mp -a <action>

Summary

The mp command activates Media Patrol. Media Patrol searches the physical drives for media errors. When an error is found, Media Patrol attempts to repair the error. If it fails to correct the error, Media Patrol attempts to remap the sector. Note: Sector remapping is not currently supported.

You can start, stop, pause, or resume Media Patrol and monitor its progress and status.

Options

-a <action>	Specifies the action to perform.
list	(Default) Displays the status and progress of Media Patrol.
start	Starts Media Patrol.
stop	Stops Media Patrol.
pause	Pauses Media Patrol.
resume	Resumes a paused Media Patrol.

Examples

```
mp

mp -a stop

mp -a resume
```

net

Usage

net [-a <action>] [-f <protocol family>] [-m] [-v]

net -a mod [-f <protocol family>] [-m]

-s “<list of settings>”

Summary

Net is used to display the TCP/IP specific information for the management port.

In addition to displaying IP address and subnet mask, changes to DHCP and DNS settings can be changed. Most often this command will be used during initial setup to either setup a static IP address or to display what DHCP assigned IP address the enclosure is using.

Options

-a <action>		Which action to perform.
	list	(Default) Displays a list of IP configurations.
	mod	To modify current network settings.
	enable	To enable IPv4/IPv6.
	disable	To disable IPv4/IPv6.
	-m	Maintenance mode.
-c <ctrl ID>		Specifies the controller ID. When the action is to modify setting and -c is not specified, the value is default to be the current controller id. Used with -m maintenance mode.
-f <protocol family>		To specify which protocol family will be modified, enabled or disabled.
	ipv4	(Default)IPv4.
	ipv6	IPv6.

-s “<option>=<value>”	List the various settings for the MGMT ports. These options are comma separated. Works only with modify command.
primaryip=	Specify the primary IP address.
primaryipmask=	Specify the primary subnet mask.
ipmasklen=	Specify the primary subnet mask length.
gateway=	Specify the gateway.
dhcp=	Enable or disable DHCP support. Currently only supported for ipv4.
enable	
disable	
primarydns=	Set an IP address of the primary DNS server.
wol=	Enable or disable Wake On Lan support. This option is valid only for management port.
enable	
disable	
-v	Verbose mode. Used with -a list.

Examples

```
net               *shows a list of ip info for all network ports*
net -a enable -f ipv4
net -a mod -m -c 1 -s “primaryip=10.0.0.2”
net -a mod -f ipv4 -s “primaryip=192.168.1.10, primaryipmask=255.255.255.0”
```

ntp

Usage

ntp [-a <action>]

ntp -a list

ntp -a mod -s “<list of settings>”

ntp -a test -t <time server>

ntp -a sync

Summary

The ntp command enables a user to view NTP status, add an NTP server, modify NTP settings, test the NTP server connection, and synchronize subsystem time with the NTP server.

Options

-a <action>	Which action to perform.
list	(Default) Displays NTP information.
mod	Change the settings for NTP.
test	Test time server.
sync	Sync time with time server.
-s “<option>=<value>”	Used to specify what options to change.
ntp=	Enable and disable ntp service.
enable	
disable	
server1=	Specific to the time servers.
.....	(max of 3 servers)

dst=	Enable and disable Daylight Saving Time.
enable	
disable	
dststarttime=	Used to specify the DST start time.
	The format is Month-WeekOfMonth-DayOfWeek.
	Month range is [Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec].
	WeekOfMonth range is [1st, 2nd, 3rd, 4th, Last].
	DayOfWeek range is [Sun, Mon, Tues, Wed, Thur, Fri, Sat].
dstendtime=	Used to specify the DST end time.
	The format is Month-WeekOfMonth-DayOfWeek.
	Month range is [Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec].
	WeekOfMonth range is [1st, 2nd, 3rd, 4th, Last].
	DayOfWeek range is [Sun, Mon, Tues, Wed, Thur, Fri, Sat].
-t <time server>	Specifies the time server to test.
	Used with -a test. Returns only failure reports.

Examples

```
ntp -a list

ntp -a mod -s “server1=ABC.123.XYZ” (adds a NTP server)

ntp -a mod -s “ntp=enable, timezone=-8, server1=ABC.123.XYZ,

dst=enable, dststarttime=Mar-2nd-Sun, dstendtime=Nov-1st-Sun”

ntp -a test -t ABC.123.XYZ

ntp -a sync
```

password

Usage

password [-u <username>] [-t<type>] [-p<protocol>]

Summary

Allows a user to change their password. A normal (non super user) user will never use the -u option, as they are allowed only to change their password. For any user who wants to change its own password, it will be first prompted for their old password before inputting their new password.

For a super user, the -u option can be used to change the password of other users. When changing the password another management user, the old password is not required.

Maximum password length is 31 characters, no spaces.

Options

-u <username>	Whose password to change. Current management user is the default.
-t <type>	Specified the user type.
mgmt	(Default) Change management user password.
snmp	Change snmp user password.
-p<protocol>	Specifies the secure protocol type. For snmp user only.
auth	Change authentication protocol password.
priv	Change privacy protocol password.

Examples

```
password
old password:******
new password:******
Retry password:******
password -u snmpuser -t snmp -p auth
Input auth password:******
Retype auth password:******
```

pdm

Usage

pdm [-a <action>] [-d <DaId>] [-s <SeqNo>]

pdm -a start -d <DaId> -s <SeqNo> -p <PdId>

pdm -a stop -d <DaId> -s <SeqNo>

pdm -a pause -d <DaId> -s <SeqNo>

pdm -a resume -d <DaId> -s <SeqNo>

Summary

This command allows the user to start, stop, pause or resume a PDM as well as to check on the progress of a running or paused PDM.

PDM (Predictive Data Migration) is an operation to replace a drive in the disk array, which has a PFA condition, with a destination physical drive. The destination physical drive can be an unconfigured drive, a global spare, or a dedicated spare to this disk array.

During PDM, the data on the PFA drive will be transferred to the destination drive while the IO remains going on. After PDM, the destination drive becomes part of the disk array; the PFA drive will become unconfigured and PFA condition will remain on.

The PFA drive cannot be used for further configuration until the PFA condition is cleared by the user.

To clear the PFA condition of a physical drive, please refer to **phydrv** *command with option -a clear.*

Options

-a <action>	Which action to perform.
list	(Default) Displays the current active or paused PDM(s) and their status(es).
start	Starts a manual PDM.
stop	Stops a PDM.
pause	Pauses a PDM.
resume	Resumes a paused PDM.
-d <DA ID>	Specifies which disk array to perform PDM action on.
-s <sequence Num>	Specifies the sequence number of the physical drive that has a PFA condition.
-p <PD ID>	Specifies physical drive ID of the destination drive.

Examples

```
pdm

pdm -a start -d0 -s2 -p10

pdm -a stop -d0 -s2
```

phydrv

Usage

phydrv [-a <action>] [-p <PdId>] [-c <Pd count>] [-v]

phydrv -a mod -p <PdId> -s “<list of settings>”

phydrv -a locate -p <PdId>

phydrv -a online -p <PdId>

phydrv -a offline -p <PdId>

phydrv -a clear -t <condition type> -p <PdId>

Summary

The phydrv command displays physical drive information, changes physical drive settings, locates individual drives, and forces a drive to an online or offline state.

Options

-a <action>	Specifies the action to perform.
list	(Default) Displays all physical drives, their make, model number, and
array	they belong to. Their status is also shown.
mod	Modifiies physical drive settings.
locate	Flashes the physical drive’s LED so you can location it.
online	Forces a drive from an Offline to an Online state.
offline	Forces a drive from an Onine to an Offline state.
clear	Clears a drive’s condition.
-p <PD ID>	Specifies the physical drive ID.
-c <count>	Specifies number of drives when their ID numbers are sequential.
-t <condition type>	Specifies type of condition to clear. Requires the -a clear option.

pfa	Clears a PFA condition on the drive.
staleconfig	Clears a stale configuration on the drive.
-d <drive type>	Specifies type of settings to modify. Requires the -a mod option. Defaults to be all if -d is not specified.
sata cmdqueuing.	SATA related setting(s): writecache, rlacache, and The SATA settings apply to all SATA physical drives.
all	All drives where is applicable.
-s “<option>=<value>”	Specifies which settings to change.
alias=	User-specified name, only for configured physical drives. Up to 32 characters, containing alpha-numeric characters, blank spaces and underscores. Beginning and ending blank spaces are discarded. The following global settings are for physical drives that support these features:
writecache=	Enables or disables write cache on the physical drive(s).
enable disable	
rlacache=	Enables or disables read look ahead cache on the physical drive(s).
enable disable	

cmdqueuing=	Enables or disables command queuing on the physical drive(s).
enable	
disable	
temppollint=	(15-255) Drive temperature polling interval in seconds. If value is 0, polling is disabled. For high availability products only.
mediumerrorthreshold=	(0-4294967294) Medium error threshold. If the threshold is reached, the physical drive is marked as dead. The default value is 0, indicating that physical drive is not marked as dead for medium errors. For high availability products only.
-v	Verbose mode. Used with -a list.

Examples

```
phydrv phydrv -v

phydrv -a locate -p 9

phydrv -a mod -s "writecache=enable,rlacache=enable"

phydrv -a offline -p 8

phydrv -a online -p 8
```

ping

Usage

ping -t <PortType> [-l <CtrlId>] [-p <PortId>] -i <ipAddr | DestinationId>

[-d <DestinationType>] [-c <packetCount>]

Summary

Allows the user to ping another network device from the management port or FC port to verify that the device is able to be “seen” by the enclosure.

Options

-t <port type>	The type of port to ping through. If -t is not specified, the default value is mgmt port.
mgmt	Management port. For embedded only.
fc	FC port. FC host interface product only.
-l <CtrlId>	Controller id. It is required when port type is fc.
-p <port ID>	Physical port id. Port id is required when port type is fc.
-i <IP address | DestinationId>	IP address to ping destination identifier for fc port.
-d <Destination Type>	The type of destination for the ping. It is required when port type is fc.
FCID	
WWPN	
-c <packet count>	Number of packets to ping. Range from 1 to 65535. It is optional for mgmt type.

Examples

```
ping -t mgmt -i 192.168.1.1   # for embedded

ping -t iscsi -l 1 -p 1 -i 192.168.1.1 -c 5

ping -t fc -l 1 -p 1 -d WWPN -i 26-03-00-01-55-60-01-0e
```

ptiflash

Usage

ptiflash [-a <action>] [-t] [-s <ServerIP>] -f <FileName> [-p <PortNum>]

[-e <encl id>] [-i <ctrl id>] [-n] [-d <pd id>] [-l] [-y]

Summary

This is the flash utility for the controller and physical drives. It is used to flash images such as firmware or software for controllers and drive firmware image for physical drives. For embedded, in order to update the flash image, the user must have a TFTP server setup that is accessible from the enclosure's management port and the flash image located on the TFTP server. For in-band, the flash image located on the local host must be accessible.

Please note that only one flash process should be running at one time.

Options

-a <action>	Which action to perform.
start	(Default) To start the flash process.
versioninfo	To get the version and build information of running images of all controllers or the specified controller.
-t	Indicates that TFTP get method is to be used to obtain the file from a TFTP server.
-s <servername|ipaddress>	Specifies the hostname or IP address of the TFTP server which contains the image file.
-f <filename>	Specifies the filename of the flash image. Include the folder name A flash image could be either a controller flash image or a physical drive firmware update image.
-p <port number>	Specifies the port number of the TFTP server. If no port number is given, the default value that will be used is 69.
-e <encl id>	Specifies the Enclosure ID. Only used with -a versioninfo option. If not specified, default value is all enclosures.
-i <ctrl id>	Specifies the Controller ID. Only used with -a versioninfo option. Enclosure id is required when controller id is specified. If not specified, default value is all controllers.

-v	Verbose mode. Only used with -a versioninfo option. To show the version and build information of all the flash images of all controllers or the specified controller.
-n	Start the flash process/image update in NDIU mode. This mode is applicable only if the system is in redundant state. Default mode of flash is DIU (disruptive) mode.
-d <device id>	Specifies the physical drive IDs. Only for physical drive firmware update. If not specified, all the physical drives, which are supported by the specified physical drive firmware, are selected.
-l	Display the status of currently running flash process.
-y	Enable non-interactive mode.

Examples

```
ptiflash -t -s 192.168.1.1 -f fw_multi.ptif -p 69  # for embedded

ptiflash -f fw_multi_20031010.ptif   # for in-band

ptiflash -l # list currently running flash process

ptiflash -t -s 192.168.1.1 -f fw_multi.ptif -n  # for NDIU mode

ptiflash -t -s 192.168.1.1 -f pd_fw.ptif -d 1,2

# update the pd firmware for pd 1 and 2 using the pd_fw.ptif image.
```

rc

Usage

rc [-a <action>] [-l <LdId>]

rc -a start -l <LdId> [-n] [-p]

rc -a stop -l <LdId>

rc -a pause -l <LdId>

rc -a resume -l <LdId>

Summary

The rc command starts, stops, pauses and resumes a Redundancy Check and monitors the progress of a running Redundancy Check.

Options

-a <action>	Specifies action to perform.
list	(Default) Displays active and paused Redundancy Checks and their status.
start	Starts a Redundancy Check.
stop	Stops a Redundancy Check.
pause	Pauses a Redundancy Check.
resume	Resumes a paused Redundancy Check.
-l <Ld ID>	Specifies the logical drive ID on which to run the Redundancy Check.
-n	Do not fix inconsistent data. This option causes Redundancy Check to run without correcting inconsistent data. All inconsistency errors are reported.
-p	Pause on error. This option causes Redundancy Check to pause when it encounters inconsistent data. The default is to continue on error.

Examples

```
rc

rc -a start -l3 -n -p

rc -a start -l3

rc -a stop -l2
```

rb

Usage

rb [-a <action>] [-d <DaId>] [-s <SeqNo>]

rb -a start -d <DaId> -s <SeqNo> -p <PdId>

rb -a stop -d <DaId> -s <SeqNo>

rb -a pause -d <DaId> -s <SeqNo>

rb -a resume -d <DaId> -s <SeqNo>

Summary

This command allows the user to start, stop, pause or resume a Rebuild as well as to check on the progress of a running or paused Rebuild.

Options

-a <action>		Which action to perform.
	list	(Default) Displays the current active or paused rebuild(s) and their status(es).
	start	Starts a manual rebuild.
	stop	Stops a rebuild.
	pause	Pauses a rebuild.
	resume	Resumes a paused rebuild.

-d <DA ID>	Specifies which disk array to perform rebuild action on.
-s <sequence Num>	Specifies the sequence number of the physical drive that was marked offline and will used for the rebuild.
-p <PD ID>	Identifies the physical drive ID that will be used in the rebuild process.

Examples

```
rb

rb -a start -d0 -s2 -p10

rb -a stop -d0 -s2
```

san

Usage

san -a <action> [-t <type>] [-i <device id list>] [-s <list of settings>] [-d <list of device setting>] [-f] [-v] [-y]

Summary

This command allows the user to list, modify a SAN, as well as add devices into or remove devices from the SAN. For any VTrak A-Class, there is always a SAN running on it.

Options

-a <action>	Which action to perform.
list	(Default) List SAN information.
mod	Modify an existing SAN, VTrak A-Class only.
add	Add devices into the SAN. Used with options -t and -d, VTrak A-Class only.
del	Remove devices from the SAN. Used with options -t and -i.
-t <type>	Device type. It's required when add or delete device(s). For list action, if not specified, both node and client information will be displayed.
node	Storage Node.
client	VTrak FS Client.
-i <device id list>	Used to specify which storage nodes or clients are to be deleted or listed, for devices in the SAN only.
-s "<option>=<value>"	
name=	Name of the SAN.

-d “<option>=<value>”

ip=	IP address of the device, used while adding a device into the SAN. For IPv4 protocol only.
-f	Force delete, used with -a del.
-y	Enable non-interactive mode. Used with -a del.
-v	Verbose mode, used with -a list on VTrak A-Class to show current status of the SAN.

Examples

```
san -a mod -s “name=new_san”
san -a add  -t node -d “ip=192.168.1.100”
-d “ip=192.168.1.102”
san -a list -t node
san -a list -v
san -a del -t node -i 1,4~7
san -a del -t node -i 1 -f
```

sas

Usage

sas [-a <action>] [-t <Type>] [-i <CtrlId>] [-p <PortId>] [-v]

sas -a mod -t <Type> -i <CtrlId> -p <PortId> -s “<list of settings>”

Summary

The sas command is used to view and modify SAS host port info and settings on SAS host interface product only.

Options

-a <action>	Which action to perform.
list	(Default) Gives summary information about SAS host port status.
mod	Modify SAS host port settings.
clear	Clear port statistics information.
-t <type>	Specifies what type of information to display or modify.
port	(Default) Specifies SAS host port as the device type to display or modify information.
stats	Display or clear port statistics information.
initiator	Display initiator list connected to subystem.
phystats	Display PHY level statistics information.
-i <ctrlId>	Controller Id. Default to be all available controllers for listing if -i is not specified. Default to be controller 1 for modifying if -i is not specified.
-p <port id>	Port number. Default to be all ports if -p is not specified when listing.
-s “<option>=<value>”	Specifies SAS host port settings to change.
cablesignalstrength	Adjust link cable signal strength. The value is from 1 to 8. Use cable

	length in meters as a guideline to select. For example, if cable length is 2 meter, the cable signal strength should be the value around 2. If 2 is not a good value, select the value such as 1 or 3.
-v	Verbose mode. Used with -a list.

Examples

```
sas

sas -t port -i 1 -p 1 -v

sas -a mod -t port -i 1 -p 1 -s “cablesignalstrength=1”
```

sasdiag

Usage

sasdiag -a <action> -e <EnclosureId> -i <expanderId> [-p <PHYId>]

Summary

Diagnostic command for getting SMP discovery info, getting PHY error log, or clear the error log. For products that support multiple enclosures only.

Options

-a <action>	Which action to perform.
discover	Display SMP general discovery information.
errorlog	Display error log on a certain expander.
clearerrlog	Clear error log on a certain PHY.
-l <PHY Location>	The location where PHY resides. If -l is not specified, the default value is expander.
expander	
c2cport	
-e <Enclosure ID>	Used to specify which enclosure ID.
-i <Expander ID>	Used to specify which expander ID.
-p <PHY ID>	Used to specify which PHY ID you wish to issue clear errorlog. Only used with -a clearerrlog option.

Examples

```
sasdiag -a discover -l expander -e 1 -i 1
sasdiag -a errorlog -l expander -e 1 -i 1
sasdiag -a clearerrlog -l expander -e 1 -i 1 -p 1
sasdiag -a errorlog -l c2cport
```

sc

Usage

sc [-a <action>] [-i <SpareId>]

sc -a start [-i <SpareId>]

Summary

The sc command starts a Spare Check and monitors the status of a running Spare Check.

Options

-a <action>		Specifies the action to perform.
	list	(Default) Displays Spare Check status.
	start	Starts the Spare Check.
-i <Spare ID>		Specifies the spare ID on which to run Spare Check. Valid value range is 0~255.

Examples

```
sc

sc -a start -i 3
```

scsi

Usage

scsi [-a <action>] [-c <ChannelId>] [-i <TargetId>] [-v]

scsi -a list -c <ChannelId>

scsi -a list -c <ChannelId> -i <TargetId>

scsi -a list -t target

scsi -a list -c <ChannelId> -t target

scsi -a mod -c <ChannelId> -s “<List of Settings>”

scsi -a enable -c <ChannelId> -i <Target Id List>

scsi -a disable -c <ChannelId> -i <Target Id List>

Summary

The parallel SCSI command is used to view and modify parallel SCSI info and settings. These include things like parallel SCSI termination and targetlist.

Options

-a <action>	Which action to perform.	
list	(Default) Gives summary information about parallel SCSI status.	
enable	To enable the specified target IDs of the specified channel.	
disable	To disable the specified target IDs of the specified channel.	
mod	To modify the specified channel termination setting.	
-t target	List all targets infomation one or all channel(s).	
-c <Channel ID>	Channel number.	
-i <Target ID>	0..15	Used to specify the target ID. Used with -a list option to display the target information and statistics.

-i <Target ID list>	0..15	Used to specify which targets are to be used in the list. Used in conjunction with -a enable or -a disable. Target IDs can be used singly or separated by comma. Additionally a sequential group of targets can be specified by placing a ~ between numbers such as 1~6. This will include targets 1,2,3,4,5,6.
-s “<option>=<value>”		Specifies which Parallel SCSI settings to change.
termination=		Parallel SCSI termination configuration
auto		
on		
off		
-v		Verbose mode, display statistics information.

Examples

```
scsi

scsi -a list  -c 1

scsi -a list  -c 1 -i 1

scsi -a list -t target

scsi -a list -c 1 -t target

scsi -a mod -c 1 -s “termination=on”

scsi -a enable  -c 1 -i 1,3,5,7~15

scsi -a disable -c 1 -i 1,3,5,7~15
```

session

Usage

session

session -h (this command)

Summary

This command lists the current active sessions.

Examples

```
session
```

shutdown

Usage

shutdown -a <action> [-i <ctrlId>

shutdown -a shutdown

shutdown -a restart

shutdown -a restart -i 2

shutdown -a shutdown -i 1

Summary

Shutdown is the command used to shutdown or restart a controller or subsystem.

Options

-a <action>	Which action to perform.
shutdown	To shutdown the controller or subsystem.
restart	To restart the controller or subsystem.
-i <ctrlId>	Controller ID or subsystem. If -i is not specified, the default value is subsystem.
1	Controller 1.
2	Controller 2.
subsys	Subystem.

smart

Usage

smart [-a <action>] [-p <PdId>]

Options

-a <action>	Which action to perform.
list	(Default) Displays the status of S.M.A.R.T. diagnostic for phydrv drive(s).
enable	Enable S.M.A.R.T.
disable	Disable S.M.A.R.T.
-p <PdId>	Specifies physical drive ID of the destination drive.
	If not specified, the destination drive will be all physical drives.
-v	Verbose mode. Used with -a list.

Summary

S.M.A.R.T diagnostic for physical drives.

Examples

```
smart

smart -v

smart -a list -p 1

smart -a enable -p 1
```

spare

Usage

spare [-a <action>]

spare -a list [-i <SpareId>] [-d <DaId>] [-v]

spare -a add [-i <SpareId>] -p <PdId> [-t g|d] [-r y|n] [-d <DaId list>] [-s “<list of settings>”]

spare -a mod -i <SpareId> [-t g|d] [-r y|n] [-d <DaId list>] [-s “<list of settings>”]

spare -a del -i <SpareId>

Summary

The spare command displays a list of hot spare drives and creates, modifies, and deletes hot spare drives.

A global hot spare can replace a failed drive from any redundant disk array.

A dedicated hot spare is assigned to one or more redundant disk arrays, and can only replace a drive that belongs to one of the assigned arrays.

A revertible hot spare can transition back to spare status after it replaces a failed drive in a disk array. See the transit command.

The hot spare drive must be of equal or greater size than the drive it replaces. The spare drive must be the same media type, HDD or SSD, as the other physical drives in the disk array.

Options

-a <action>	Specifies the action to perform.
list	(Default) Displays a list of hot spare drives.
add	Adds new hot spare drives.
mod	Changes hot spare drive settings.
del	Deletes a hot spare drive.
-i <Spare Id>	Specifies the ID of the spare drive.
-p <PD ID>	Specifies the ID of the physical drive. Requires the -a add option to configure a drive as a spare.
-d <DA ID or DA ID List>	

	Specifies the disk array ID. Requires the -a list option. Displays a list of global spares and spares dedicated to this disk array. When used with other actions, it specifies the disk array IDs to which this spare is dedicated.
-t <type>	Specifies the type of hot spare drive.
g	A global hot spare drive.
d	A dedicated hot spare drive.
-r <revertible>	Specifies whether the spare drive is revertible.
y	Yes.
n	No.
-s “<option>=<value>”	Specifies options for the spare drive.
mediapatrol=	Enables or disables Media Patrol.
enable	
disable	

Examples

```
spare

spare -a add -p 14 -t g -r y

spare -a mod -i 1 -t d -d 0,1 -s “mediapatrol=disable”

spare -a del -i 0
```

stats

Usage

stats [-t <type>] [-i <devId>] [-c <Count>]

stats -a clear

Summary

The stats command displays statistics of subsytem, controller, enclosure, physical drives, and logical drives; and resets the statistics count to zero.

Options

-a <action>		Specifies the action to perform.
	list	(Default) Displays the statistics.
	clear	Resets the statistics count to zero.
-t <type>		Specifies the device type.
	ctrl	Controller.
	logdrv	Logical drive.
	phydrv	Physical drive.
	all	All the above options.
-i <devId>		Specifies the device ID. Default is the first available device ID.
-c <Count>		Specifies the device count. Default is all devices.

Examples

```
stats -t logdrv -i 0 -c 5

stats -a list -t all

stats -a clear
```

subsys

Usage

subsys [-a <action>] [-v]

subsys -a mod -s “<list of settings>”

subsys -a lock [-r] [-t <number of minutes>]

subsys -a unlock [-f]

subsys -a chklock

Summary

The subsys command is used to display and make changes to subsystem settings. This is also used to lock the subsystem so that only the current administrator can make modifications.

Options

-a <action>	Specifies the action to perform.
list	(Default) Displays information for the specified subsystem.
mod	Modifies subsystem settings.
lock	Locks the subsystem so other users cannot apply changes. No changes can be made to subsystem settings by other users until the lock expires or the system is unlocked.
unlock	Clears a subsystem lock.
chklock	Checks the status of the lock.
-s “<option>=<value>”	Specifies which subsystem settings to change.
alias=	A user specified name to identify the subsystem. It can be up to 48 characters long, containing alpha-numeric characters, blank spaces and underscores. The beginning and ending blank spaces will be discarded.

redundancytype=	Redundancy type in high availability set up. SAS host interface product doesn't support active-standby. The default value is active-active if not specified.
active-active	Active-Active.
active-standby	Active-Standby.
cachemirroring=	Enable and disable cache mirroring. Cache mirroring will only be available when redundancy type is active-active. The default value is enable if not specified.
enable	
disable	
-t <number of mins>	Used with -a lock. Number of minutes to lock the subsystem. Default is 30 minutes.
-r	Renew the lock timer. Used with -a lock and -t
-f	Force unlock. Only super user has the privilege to do it.
-v	Verbose mode. Used with -a list.

Examples

```
subsys

subsys -v

subsys -a mod -s "alias=MySubsystem"

subsys -a lock -t 60

subsys -a lock -r -t 35

subsys -a unlock

subsys -a chklock
```

swmgt

Usage

swmgt [-a <action>]

swmgt -a mod -n <component name> [-t <startup type>] [-s “<list of settings>]

swmgt -a start -n <component name>

swmgt -a stop -n <component name>

swmgt -a restart -n <component name>

Summary

The swmgt command allows a user to view and modify setting of software components.

Options

-a <action>	Which action to perform.
list	(Default) Displays all software components.
start	Start a software component.
stop	Stop a software component.
restart	Restart a software component.
mod	Change a component’s startup type when system boots.
add	Add trap sink for snmp, public key for ssh or recipient for netsend.
del	Delete trap sinks for snmp, public key for ssh or recipients for netsend.

-n <component name>	Specifies the component name to view setting, modify, start or stop.
email	Email notification.
slp	Service location protocol service agent. SLP service is supported for IPv4 protocol only.
webserver	Web server.
telnet	Telnet.
ssh	SSH.
snmp	SNMP.
cim	CIM.
netsend	Netsend. Netsend service is supported for IPv4 protocol only.
-t <startup type>	Specifies the startup type.
automatic	Component is automatically started when system boots.
manual	Component has to be manually started by issuing command.
-s “<option>=<value>”	Used to specify settings for this component. This is used when modifying (mod). These options are comma separated.

email settings

smtpserver=	SMTP server IP address or SMTP server name.
serverport=	SMTP server port number.
authentication=	SMTP server authentication.
	no
	yes
username=	Username if using SMTP authentication.
senderaddr=	Sender’s email address.
subject=	Email subject.

telnet settings

port=	Port number for telnet daemon.
sessiontimeout=	Session time out in minutes. Maximum 1440.
maxconnection=	Max number of telnet client connection .

ssh settings

port=	Port number for ssh daemon.
sessiontimeout=	Session time out in minutes. Maximum 1440.
maxconnection=	Max number of ssh client connection .

snmp settings

port=	Port number.
sysname=	System name string.
syslocation=	System location string.
syscontact=	System contact information string.
readcommunity=	Read community name.
-i <Index>	Used to specify trap sink index for snmp, public key index for ssh or recipient index for netsend. Only valid for modify or delete trap sink or recipient, delete public key.

-p “<option>=<value>”	Used to specify trap sinks for snmp, public key for ssh or recipients for netsend. Multiple -p option can be entered with -a add option for trap sink or recipient.
trapsinkserver=	Trap sink IP address or trap sink server name. For snmp only.
trapfilter=	Trap filter level. It implies the level and above. For snmp only.
info	
warning	
minor	
major	
critical	
fatal	
recipientserver=	Recipient IP address or recipient server name. For netsend only.
messagefilter=	Message filter level. It implies the level and above. For netsend only.
info	
warning	
minor	
major	
critical	
fatal	
filename=	Ssh public key file name. For ssh only.
server=	TFTP server IP address or server name. For ssh only.
comment=	Ssh public key comment. For ssh only.

Examples

```
swmgt

swmgt -n snmp

swmgt -a start -n snmp

swmgt -a stop -n snmp

swmgt -a mod -n snmp -t automatic

swmgt -a mod -n netsend -i 1 -p "recipientserver=192.168.1.1,messagefilter=info"

swmgt -a add -n netsend -p "recipientserver=192.168.1.1,messagefilter=info"

swmgt -a del -n netsend -i 1

swmgt -a add -n ssh -p "filename=key.pub, server=192.168.1.1,

comment=root@server"

swmgt -a del -n ssh -i 1
```

For adding multiple trapsinkserver (SNMP):

```
swmgt -a add -n snmp -p "trapsinkserver=192.168.1.1,trapfilter=info"

-p "trapsinkserver=192.168.2.1,trapfilter=critical"
```

For adding multiple reciprentserver (Netsend):

```
swmgt -a add -n netsend -p "recipientserver=192.168.1.1,messagefilter=info"

-p "recipientserver=192.168.2.1,messagefilter=critical"
```

sync

Usage

sync [-a <action>] [-l <LdId>]

Summary

The sync command is used for background synchronization, the process of enforcing consistency in logical drives. This is an optional replacement for LDI (logical drive initialization).

Background Synchronization starts automatically when a redundant logical drive is created while still allowing I/O to be performed to the logical drive (unlike LDI). If LDI is started then background synchronization will halt and LDI will run.

This command may also be used to allow the user to check the status of background synchronization. Since background synchronization is started automatically and yields automatically there is no need to explicitly start, stop, pause or resume a background synchronization.

Options

-a <action>	Specifies the action to perform.
list	(Default) Displays the current background synchronization activities and their status.
-l	Specifies the logical drive ID on which background synchronization is running.

Examples

```
sync

sync -l3

sync -a list -l3 **same as example above
```

syslog

Usage

syslog -a <action> [-i <index>] [-s <settings>]

Summary

To configure the settings for syslog remote servers.

Options

-a <action>	Which action to perform.
listserver	(Default) Displays the current settings of the remote syslog servers.
addserver	Add a new remote syslog server.
delserver	Delete an existing remote syslog server.
modserver	Modify the settings of an existing remote syslog server.
-i index	The index of remote syslog servers, ranging from 1 to 4. This option is valid when modifying an existing remote syslog server settings or when deleting a server.
-s “<option>=<value>”	The settings of one remote syslog server. It’s valid only when adding a server or modifying the settings.
server=	Server IP or domain name of the remote syslog server. This setting is valid only when adding the server.
port=	The UDP or TCP port number. If not specified, 514 is used by default.
protocol=	Networking protocol. When adding a server, if this setting is not specified, “udp” is used by default. When modifying a server settings, if this setting is not specified, it means no change to the protocol.

udp	
udp6	
tcp	
tcp6	
severity=	Syslog severity level according to the standard syslog specification in RFC5424
debug	Send the logs with debug and more severe severity.
info	Send the logs with info and more severe severity.
notice	Send the logs with notice and more severe severity.
warning	Send the logs with warning and more severe severity.
error	Send the logs with error and more severe severity.
critical	Send the logs with critical and more severe severity.
alert	Send the logs with alert and more severe severity.
emerg	Only send the logs with emergency severity.

Examples

```
syslog

syslog -a addserver -s “server=192.168.252.252,port=514,severity=error”

syslog -a modserver -i 2 -s “severity=info”

syslog -a delserver -i 4
```

topology

Usage

topology [-a <action>] [-v]

Summary

View enclosures topology, the physical connections and devices. For products that support multiple enclosures only.

Options

-a <action>	Which action to perform.
list	(Default) Displays topology information.
-v	View complete information about topology.

Examples

```
topology
```

transit

Usage

transit [-a <action>] [-d <DaId>] [-s <SeqNo>]

transit -a start -d <DaId> -s <SeqNo> -p <PdId>

transit -a stop -d <DaId> -s <SeqNo>

transit -a pause -d <DaId> -s <SeqNo>

transit -a resume -d <DaId> -s <SeqNo>

Summary

The transit command starts, stops, pauses, and resumes a transition and monitors the progress of a running transition.

Transition is an operation to replace a revertible spare drive currently used in a disk array with an new physical drive, so the reveritible spare can be restored to spare drive status. The destination drive can be an unconfigured drive, a non- revertible global spare, or a non-revertible spare dedicated to the array.

During transition, the data on the revertible spare is transferred to the destination drive while the disk array remains online. After transition, the destination drive becomes the part of the array and the revertible spare is a spared drive once again.

Note that the destination drive must be the same media type, HDD or SSD, as the other physical drives in the disk array.

Options

-a <action>	Specifies the action to perform.
list	(Default) Displays the running and paused transitions and their status.
start	Starts a manual transition.
stop	Stops a transition.
pause	Pauses a transition.
resume	Resumes a paused transition.
-d <DA ID>	Specifies the id of disk array which contains the revertible spare drive.
-s <sequence Num>	Specifies the sequence number of the revertible spare drive in the array.
-p <PD ID>	Specifies the physical drive ID of the destination drive.

Examples

```
transit

transit -a start -d 0 -s 2 -p 10

transit -a stop -d 0 -s 2
```

ups

Usage

ups [-a <action>]

ups -a list [-v]

ups -a mod -s “<list of settings>”

Summary

The ups command allows a user to view and modify ups status and settings. Network UPS is supported for IPv4 protocol only.

Options

-a <action>	Which action to perform.
list	(Default) Displays all current UPS status.
mod	Change the settings for UPS.
-s “<option>=<value>”	Used to specify what options to change.
detection=	Detection mode setting
auto	(Default. Whenever a UPS is detected, it changes the detection mode to “enable”.)
enable	(Monitors UPS, UPS Settings changes, reports warnings and logs events.)
disable	(Monitors Serial UPS only.)
ups1=	UPS1 IP address or Domain Name.
ups2=	UPS2 IP address or Domain Name.

rtr=	Running time remaining threshold in minute. The valid value range is 3~20.
lr=	Critical loading ratio threshold in percentage The valid value range is 1~100.
wt=	Warning temperature threshold in Celsius. The valid value range is 32~42.
-v	Verbose mode. Used with -a list.

Examples

```
ups -v

ups -a mod -s “ups1=192.168.1.1, rtr=5”
```

user

Usage

user [-a <action>] [-u <username>]

user -a add -u <username> -p <privilege> [-s “<list of settings>”]

user -a mod -u <username> [-p <privilege>] [-s “<list of settings>”]

user -a del -u <username>

Summary

The user command allows a user to view and modify an existing user account.

Only a Superuser can create, modify, or delete a user account.

User access levels are: Superuser, Poweruser, Maintenance, and View.

If a password is not specified when the user account is created, there will be no password when you log in.

Use the password command to change a password.

Maximum password length is 31 characters, no spaces.

Options

-a <action>		Which action to perform.
	list	(Default) Displays the current users.
	add	Create a new user.
	mod	Modify an existing user.
	del	Delete a user.
-u <username>		Specifies the username to display, edit or delete.

-p <privilege>	Specifies the privilege level to set for the user.
super	Superuser has max control
power	Poweruser cannot modify users nor delete configs
maintenance	Maintenance user can only perform background tasks
view	View user can only view.
-f	Force delete a user.

-s “<option>=<value>”

status=	Enable/disable this user’s account. Default is enable. Only for local user.
name=	Specifies the user’s display full name.
email=	Specifies an email address for the user.

Examples

```
user -a add -u newuser -p view -s”name=NewUser,
email=MyEmail@yourcompany.com”
Input password: ******
Retype password: ******
user -a mod -u olduser -p super -s”status=disable,name=OldUser”
user -a del -u baduser
```

zoning

Usage

zoning [-a <action>] [-g <group id>] [-i <ctrl id>]

zoning -a mod -g <froup id> -i <ctrl id>] [-s “<list of settings>”]

Summary

The zoning command allows a user to view and modify zoning membership table and permission table.

Options

-a <action>	Which action to perform.
list	(Default) Displays memembership table and permission table.
mod	Modify permission table.
-g <group id>	Specifies first group id.
-i <ctrl id>	Specifies controller id for permission table.
-s “<option>=<value>”	
group=	Specifies second group id for permission table.
permission=	Enable/disable this permission table.

Examples

```
zoning

zoning -a mod -i 1 -g 17 -s”group=10, permission=enable”
```

help

Usage

-a <action> -u <username> -p <privilege> -s “<option>=<value>”

Summary

The user command is used to list, modify, create and delete user accounts on the subsystem.

?

Usage

-a <action> -u <username> -p <privilege> -s “<option>=<value>”

Summary

The user command is used to list, modify, create and delete user accounts on the subsystem.

MAINTENANCE

This chapter covers the following topics:

Updating the Subsystem Firmware

This procedure applies to VTrak E5000 RAID subsystems and VTrak E5000 JBOD expansion units managed by a VTrak E5000 RAID subsystem. There are two methods:

- "Updating with WebPAM PROe" on page 340
- "Updating Vtrak E5000 with USB Support" on page 342

Updating with WebPAM PROe

Download the latest firmware image file from PROMISE support: http://www.promise.com/support/ and save it to your Host PC or TFTP server.

Important

Verify that no background activities are running on the RAID subsystem.

To update the firmware on the RAID subsystem and JBOD expansion units:

1. Click the **Administration** tab.
2. Click the Firmware **Update** icon.
3. Click the **Controller Firmware Update** tab.

   The Controller Firmware Update screen appears showing the current Image Version Number and Build Date.
4. Choose a download option:
   - **Local File through HTTP** – Click the **Browse** button, locate the firmware image file, click the file to choose it, then click the **Open** button.
   - **TFTP Server** – **Enter** the TFTP Server host name or IP address, port number and file name.

5. Optional. Check the Non-disruptive Image Update (NDIU) box.

   NDIU updates the RAID controllers and I/O modules one at a time, enabling I/O operations continue during the firmware update. Updates with this option take a longer period of time to complete. All VTrak E5000 models support this feature.

6. Click the **Next** button.

   The next screen shows the Flash Image (firmware image file) Version Number and Build Date.

Warning

Do NOT power off the RAID subsystem during the update!

Do NOT move to any other screen until the firmware update operation is completed!

7. Click the **Submit** button.
8. In the **Confirmation** box, type the word "**confirm**" in the field provided and click the **Confirm** button.

   The progress of the update displays.

   When the update is completed a message tells you to reboot the subsystem,

9. Click the **OK** button.
   - If you chose the Disruptive Flash Method, the RAID subsystem and JBOD expansion units automatically restart.
   - If you chose the Non-Disruptive Flash Method, the system automatically flashes and restarts the RAID controllers one at a time.

Automatic Restart

If you did NOT check the NDIU box, the RAID subsystem and JBOD expansion units automatically restart. That action temporarily disrupts I/O operations and drops your WebPAM PROe connection.

To reestablish your WebPAM PROe connection:

1. Wait no less than two minutes.
2. Click **Logout** in the WebPAM PROe Header, then log in again.

   If you cannot log in, wait 30 seconds and try again.

3. In your browser, click Logout in the WebPAM PROe Header, then log in again.

   If you cannot log in immediately, wait 30 seconds and try again.

Updating Vtrak E5000 with USB Support

USB support uses the disruptive flash method only. Both RAID controllers and all JBOD I/O modules are updated at the same time and momentarily go offline when the RAID subsystem and JBOD unit reboot.

This procedure requires a USB flash device:

- Formatted to FAT 32
- At least 300 MB of free space

Download the latest OPAS_xxxxx.zip firmware image file from PROMISE support: http://www.promise.com/support/ and save it the root folder of the USB flash device.

Important

Verify that no background activities are running on the RAID subsystem.

To update the subsystem firmware using VTrak E5000's USB Support feature:

1. For VTrak E5600 and VTrak E5800 you have the option to insert the USB flash device into the bottom USB port on the front panel which is for the RAID controller; or on a USB port on the controller.

VTrak E5600 front left side

VTrak E5600 front right side

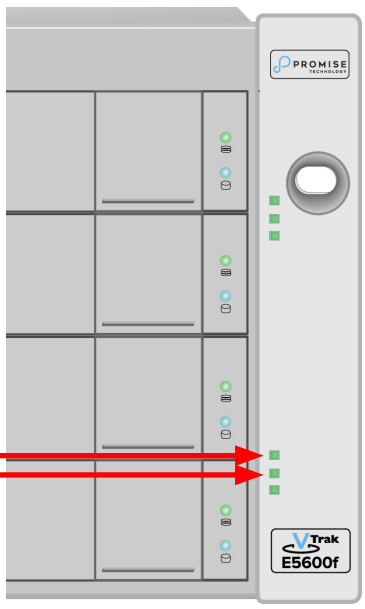

Warning

Do NOT power off the RAID subsystem during the update!

Do NOT move to any other screen until the firmware update operation is completed!

USB ports on VTrak E5000f controller

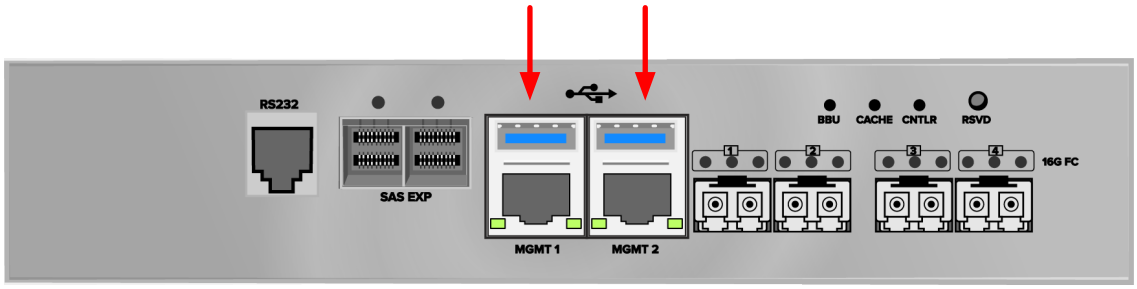

For VTrak E5300 and VTrak E5320, use one of the USB ports on the controller.

2. Wait until the controller activity LED stops blinking green and starts blinking amber.
3. Within 30 seconds, remove the USB flash device, then insert the USB flash device back into the same RAID controller.

   The remove and insert action confirms that you want to update the firmware.
4. Wait until the controller activity LED displays steady green.
5. Remove the USB flash device.

Automatic Restart

After you remove the USB flash device from the RAID controller, the RAID subsystem and any JBOD expansion units automatically restart. That action temporarily disrupts I/O operations and drops your WebPAM PROe connection.

To reestablish your WebPAM PROe connection:

1. Wait no less than two minutes.
2. Click **Logout** in the WebPAM PROe Header, then log in again.

   If you cannot log in, wait 30 seconds and try again.

Failed Update

If the firmware update fails, the controller status LED displays red. See “VTrak E5600 front right side” on page 342 for location of controller status LEDs. Remove the USB flash device.

3. Insert the USB flash device into a USB port on your PC.
4. Go to the **OPAX_xxxxxx** folder to obtain the report and log.

   Possible causes for an update failure include:

   - Less than 300 MB free space on the USB flash device.
   - The VTrak E5000 firmware image is invalid.
   - A background activity is running.

   See “Contacting Technical Support” on page 469.

Updating Physical Drive Firmware

This feature applies only to PROMISE-supported physical drives. For a list of supported drives, go to PROMISE support: http://www.promise.com/support/.

If you have physical drives in your RAID system that are not PROMISE-supported, follow the firmware update procedure from the drive manufacturer.

WebPAM PROe

Download the latest firmware image file from PROMISE support: http://www.promise.com/support/ and save it to your Host PC or TFTP server.

To update the firmware on PROMISE-supported physical drives:

1. Click the **Administration** tab.
2. Click the **Firmware Update** icon.
3. Click the **PD Firmware Update** tab.
4. Choose a download option:
   - **Local File through HTTP** – Click the **Browse** button, locate the firmware image file, click the file to choose it, then click the **Open** button.
   - **TFTP Server** – **Enter** the TFTP Server host name or IP address, port number and file name.
5. Click the **Next** button.
6. Click the **Submit** button.

   The progress of the update displays.

Warning

Do NOT power off the RAID subsystem during the update!

Do NOT move to any other screen until the firmware update operation is completed!

When the update is completed a message tells you to reboot the subsystem.

7. Click the **OK** button.

   Restart the RAID subsystem. See “Restarting a Subsystem” on the next page.

Restarting a Subsystem

This function shuts down the subsystem and then restarts it.

Important

Do NOT turn off the power supply switches on the RAID subsystem or JBOD expansion units.

To restart the subsystem:

1. Click the **Administration** tab.
2. Click the S**ubsystem Information** icon.
3. Click the **Shutdown/Restart** button.
4. Click the **Restart** button.
5. Type the word “**confirm**” in the field provided.
6. Click the **Confirm** button.

   When the controller shuts down, your WebPAM PROe connection is lost.

7. Wait no less than two minutes.
8. In your browser, click Logout in the WebPAM PROe Header, then log in again.

   If you cannot log in immediately, wait 30 seconds and try again.

REPLACING A POWER SUPPLY

The power supplies for the VTrak E5000 models are hot-swappable. Note that the VTrak E5300/E5320 uses a different PSU than the VTrak E5600/E5800.

VTrak E5300 and VTrak E5320 PSU

Before removing a bad power supply, first verify there is a problem. See “VTrak E5300 / VTrak E5320 Power Supply LEDs” on page 416 for a description of power supply LED behavior.

Removing the power supply

To remove the power supply:

Verify that the power supply *AC* LED is amber or red. Proceed with the following steps if the power supply needs replacement:

1. Switch off the power to the power supply you plan to replace.
2. Unplug the power cord.
3. Use thumbs and forefingers to squeeze the latch release tabs on the two levers used to remove the power supply. Do not grab the controller lever by mistake.
4. Pull the levers toward you and out to each side, then slide power supply module out of the enclosure.

Push both gray colored lever release tabs outward

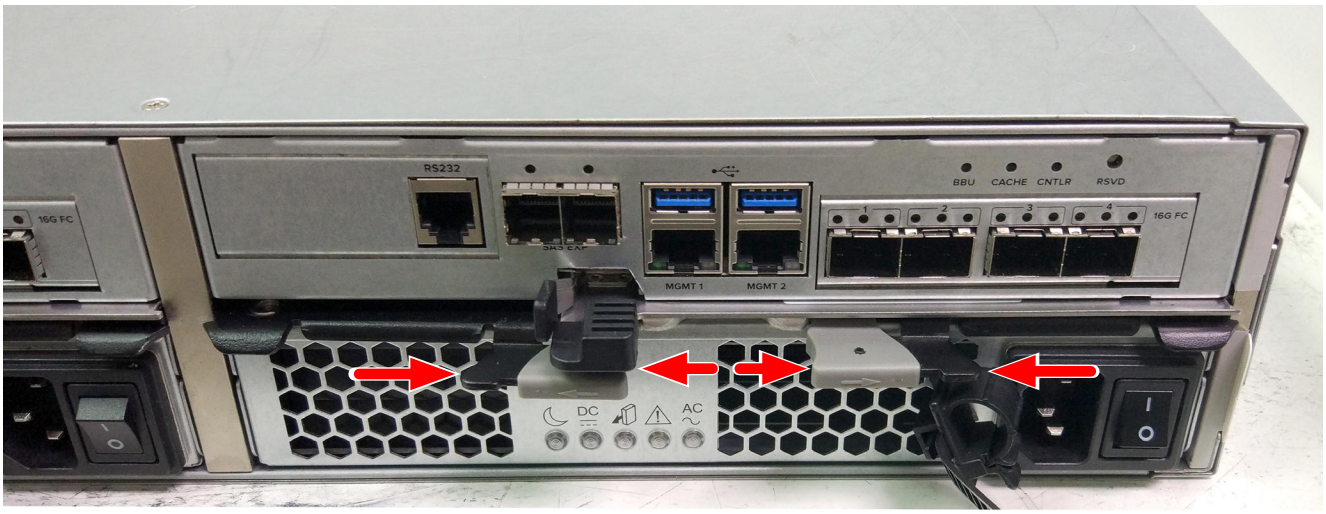

Pull both levers backward and outward, the module should slide out easily

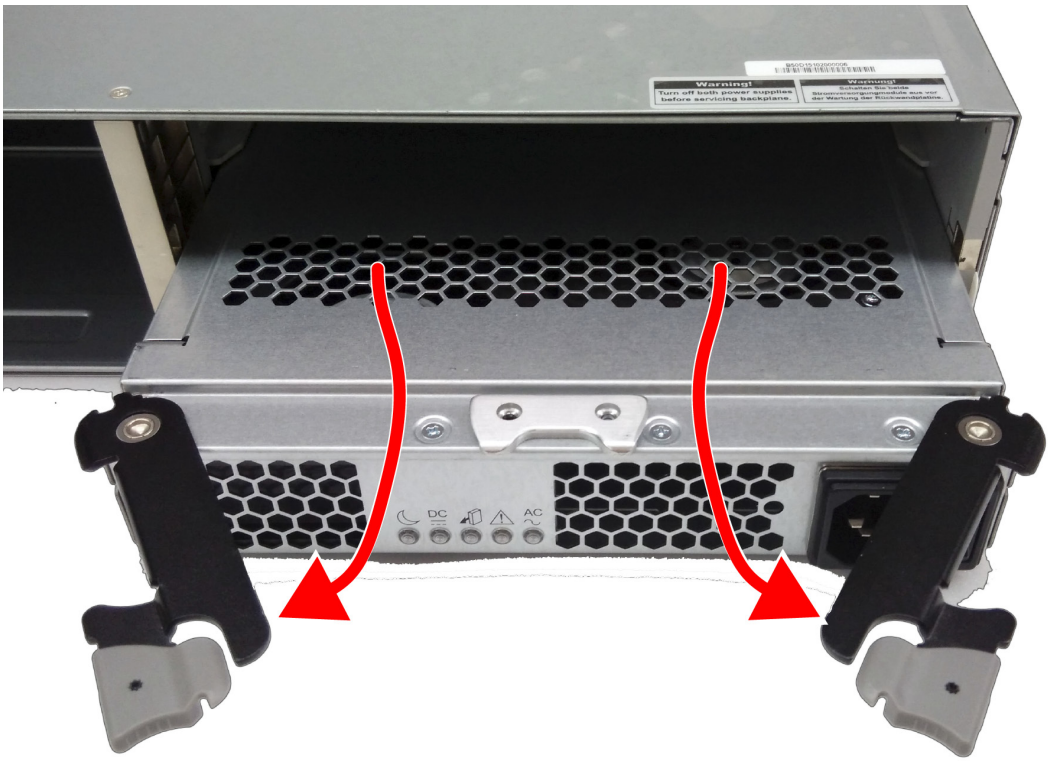

Installing the new power supply

To install the power supply:

1. Carefully slide the power supply into the enclosure until the levers on each side catch on each side of the power supply.
2. Push both levers in to firmly seat the power supply into the enclosure. Both levers will latch to the enclosure so that the levers cannot be pulled out unless you release the latch.
3. Plug in the power cord.
4. Switch on the power supply.

   Verify that the *DC* and *AC* LEDs are green.

This completes the power supply replacement procedure.

Replacing a Power Supply - VTrak E5600 / E5800

The power supply and its fans are replaced as one unit. There are no individually serviceable parts. No tools are required for this procedure.

Removing the Old Power Supply

To remove the power supply:

1. Verify that the PSU status LED is amber or red.
2. Unplug the power cord.
3. Turn the set screw counter-clockwise to loosen it. The screw is retained on the power supply housing.
4. Grasp the handle and pull the power supply straight out of the enclosure.

Power supply for VTrak E5600 and VTrak E5800

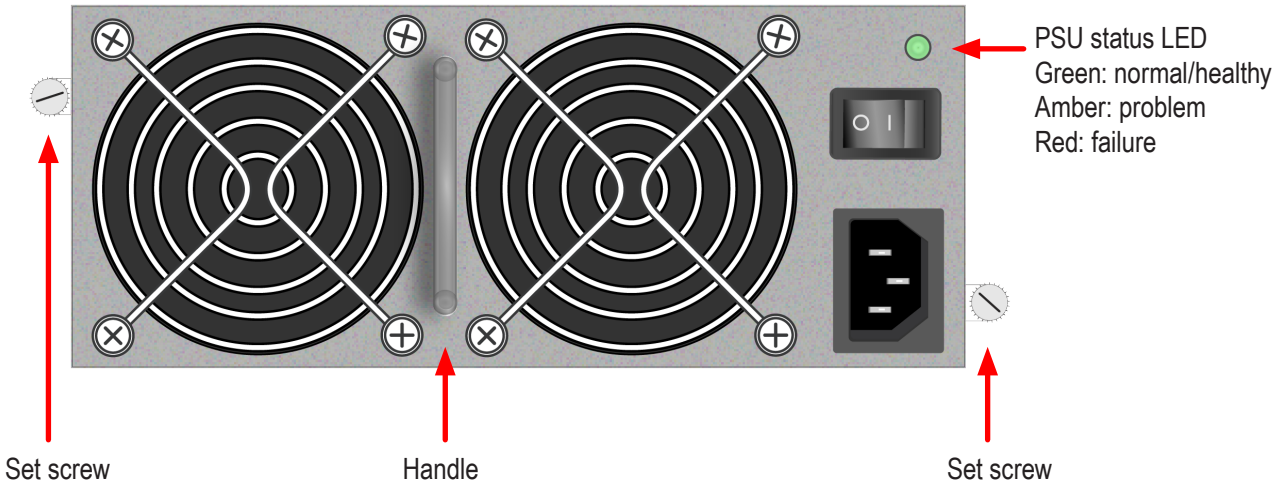

Installing a New Power Supplyt

To install the power supply:

1. Carefully slide the power supply into the enclosure.
2. Turn the set screw clockwise to tighten, DO NOT over tighten.
3. Plug in the power cord.
4. Switch on the power supply.
5. Verify that the new power supply status LED is green. (See illustration on previous page).

Replacing a RAID Controller: Dual Controllers

The RAID controller monitors and manages the logical drives. When the RAID controller is replaced, all of your logical drive data and configurations remain intact because logical drive information is stored on the physical drives.

Important

Do not replace the RAID controller based on LED colors alone. Only replace the RAID controller when directed to do so by PROMISE Technical Support.

The firmware version and amount of SDRAM must be the same on the replacement RAID controller and the other RAID controller in the subsystem.

To obtain firmware and SDRAM information for an installed RAID controller, in WebPAM PROe, click the Administration button then click the Image Version icon.

Removing the Old Controller

To remove a RAID controller:

1. Disconnect all attached cables from the RAID controller,
   - Fibre Channel cables
   - SAS expansion cables
   - Management port cables
   - Serial cable
   - UPS control cable
2. On the controller handle, squeeze the release tab and pull the handle outward. See illustrated guide in “Remove controller” on page 354.
3. Pull the RAID controller out of the subsystem enclosure.

Note

On subsystems with dual RAID controllers, you can hot-swap a controller while the subsystem is running.

Installing the New Controller

To install the new RAID controller:

1. Carefully slide the RAID controller into the enclosure.
2. Gently swing the handle in and press the handle until it locks.
3. Reconnect all cables that were attached to the RAID controller.

- Fibre Channel cables
- SAS expansion cables
- Management port cables
- Serial cable
- UPS control cable

If one of the controllers goes into maintenance mode, see “RAID Controller Problems” on page 430 and “Maintenance Mode” on page 430.

Replacing a RAID Controller: Single Controller

The RAID controller monitors and manages the logical drives. When the RAID controller is replaced, all of your logical drive data and configurations remain intact because logical drive information is stored on the physical drives.

Cautions

The RAID controller is NOT hot-swappable if your VTrak E5000 has only one controller. Power-down the VTrak E5000 before removing it.

Important

Do not replace the RAID controller based on LED colors alone. Only replace the RAID controller when directed to do so by PROMISE Technical Support. See page 435.

Important

The firmware on the replacement RAID controller must be the same version as the original RAID controller or a later version.

The amount of SDRAM in the replacement RAID controller must be the same as the original RAID controller or greater.

To obtain firmware and SDRAM information for the currently installed RAID controller, click the Administration button then click the Image Version icon.

Remove controller

The mechanism used to remove the controller is that same on all VTrak E5000 models. Follow the steps below to remove the old controller module:

1. Squeeze controller handle release between thumb and finger.

Release lever on controller

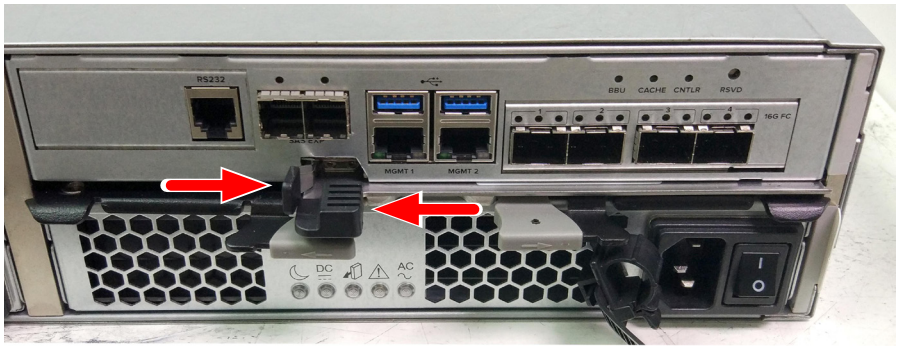

2. Pull handle out, then slide controller unit straight out of the controller.

Pull out lever on controller

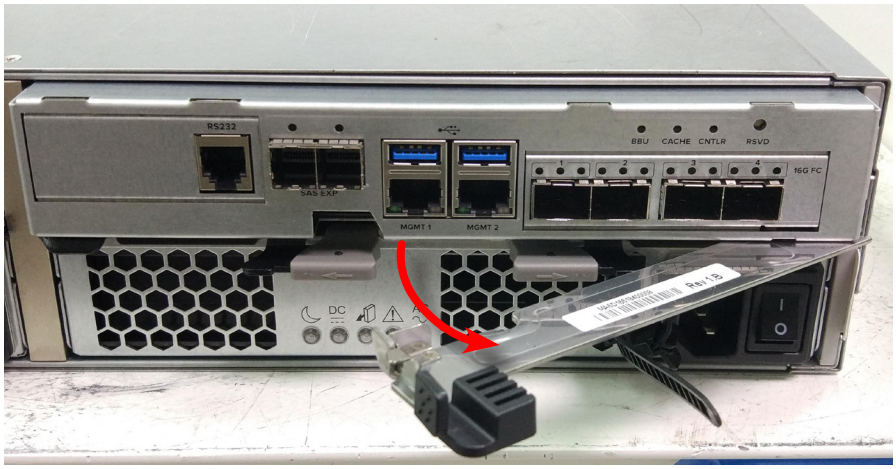

TECHNOLOGY BACKGROUND

This chapter covers the following topics:

Disk Arrays

Disk array technology includes:

- “Media Patrol”
- “PDM”

Media Patrol

Media Patrol is a routine maintenance procedure that checks the magnetic media on each disk drive. Media Patrol checks all physical drives assigned to disk arrays and spare drives. Media Patrol does not check unconfigured drives.

Media Patrol checks are enabled by default on all disk arrays and spare drives. You can disable Media Patrol in the disk array and spare drive settings, however that action is not recommended.

Unlike Synchronization and Redundancy Check, Media Patrol is concerned with the condition of the media itself, not the data recorded on the media. If Media Patrol encounters a critical error, it triggers PDM, if PDM is enabled on the disk array.

Media Patrol has three status conditions:

- ***Running*** – Normal. You can access your logical drives at any time.
- ***Yield*** – Temporary pause while a read/write operation takes place.
- ***Paused*** – Temporary pause while another background runs. Or a pause initiated by the user.

PDM

Predictive Data Migration (PDM) is the migration of data from the suspect phyi-cal drive to a spare drive, similar to rebuilding a logical drive. But unlike Re-building, PDM constantly monitors your physical drives and automatically copies your data to a spare drive before the physical drive fails and your logical drive goes Critical.

The following actions trigger PDM:

- A phyical drive with unhealthy status (see below)
- Media Patrol finds a critical error
- You initiate PDM manually

PDM also counts the number of media errors reported by Media Patrol. A disk drive becomes unhealthy when:

- A SMART error is reported
- The bad sector remapping table fills to the specified level.

Because data would be lost if written to a bad sector, when a bad sector is detected, the physical drive creates a map around it. These maps are saved in the bad sector remapping table, which has a capacity of 512 reassigned blocks and 2048 error blocks. See “Making PDM Settings” on page 154 or “Managing Background Activities” on page 144.

You can specify the maximum levels for the reassigned and error blocks in PDM settings. When the table fills to a specified value, PDM triggers a migration of data from the suspect drive (the disk drive with the bad sectors) to a replacement physical drive.

During data migration, you have access to your logical drives but they respond more slowly to read/write tasks because of the additional operation. The time required for data migration depends on the size of the physical drives.

PDM is enabled on all disk arrays by default. You can disable PDM in the disk array settings, however that action is not recommended. See “Running PDM on a Disk Array” on page 195.

Logical Drives

Logical drive technology includes:

RAID Levels

RAID (Redundant Array of Independent Disks) allows multiple physical drives to be combined together in a disk array. Then all or a portion of the disk array is formed into a logical drive. The operating system sees the logical drive as a single storage device, and treats it as such.

RAID 0 – Stripe

When a logical drive is striped, the read and write blocks of data are interleaved between the sectors of multiple physical drives. Performance is increased, since the workload is balanced between drives or "members" that form the logical drive. Identical drives are recommended for performance as well as data storage efficiency.

RAID 0 Striping interleaves data across multiple drives

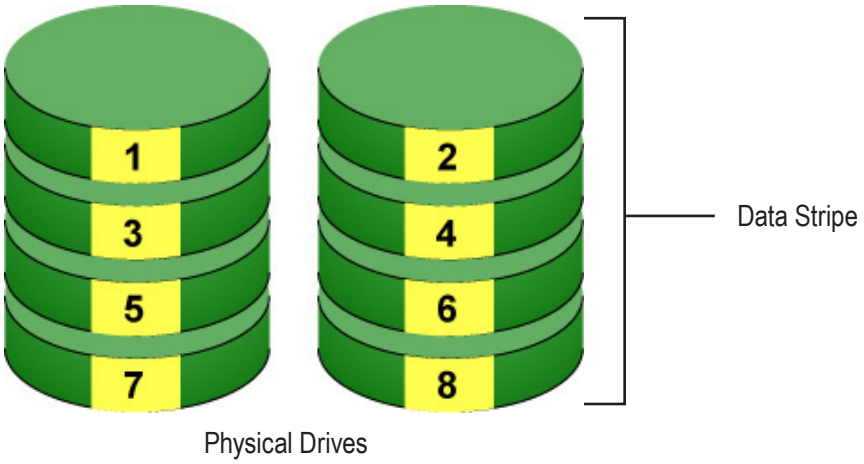

The disk array's data capacity is equal to the number of disk drive members multiplied by the smallest drive's capacity. For example, one 100 GB and three 120 GB drives form a 400 GB (4 x 100 GB) disk array instead of 460 GB.

If physical drives of different capacities are used, there is unused capacity on the larger drives.

RAID 0 logical drives on VTrak consist of one or more physical drives.

Advantages	Disadvantages
• Implements a striped disk array, the data is broken down into blocks and each block is written to a separate disk drive • I/O performance is greatly improved by spreading the I/O load across many channels and drives • No parity calculation overhead is involved	• Not a true RAID because it is not fault-tolerant • The failure of just one drive results in all data in an disk array being lost • Should not be used in mission critical environments

Recommended Applications for RAID 0:

- Image Editing
- Pre-Press Applications
- Any application requiring high

RAID 1 – Mirror

When a logical drive is mirrored, identical data is written to a pair of physical drives, while reads are performed in parallel. The reads are performed using elevator seek and load balancing techniques where the workload is distributed in the most efficient manner. Whichever drive is not busy and is positioned closer to the data is accessed first.

With RAID 1, if one physical drive fails or has errors, the other mirrored physical drive continues to function. Moreover, if a spare physical drive is present, the spare drive is used as the replacement drive and data begins to mirrored to it from the remaining good drive.

RAID 1 Mirrors identical data to two drives

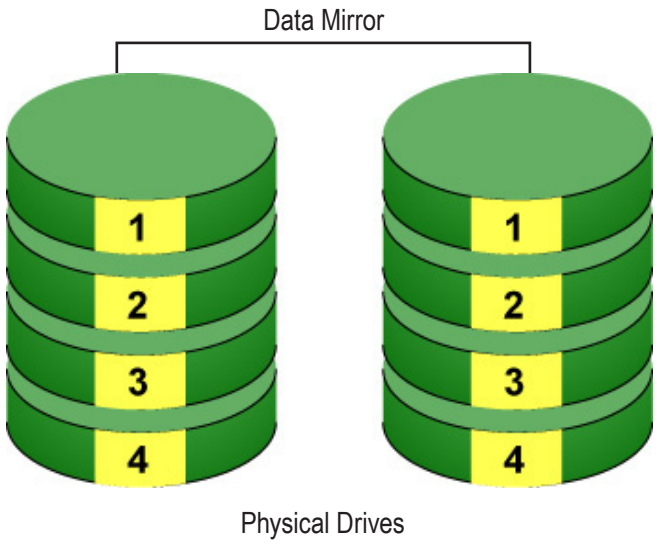

The logical drive's data capacity equals the smaller physical drive. For example, a 100 GB physical drive and a 120 GB physical drive have a combined capacity of 100 GB in a mirrored logical drive.

If physical drives of different capacities are used, there is unused capacity on the larger drive.

RAID 1 logical drives on VTrak consist of two physical drives.

Advantages	Disadvantages
• Simplest RAID storage subsystem design • Can increase read performance by processing data requests in parallel since the same data re-sides on two different drives	• Very high disk overhead – uses only 50% of total capacity

Recommended Applications for RAID 1:

- Accounting
- Payroll
- Financial
- Any application requiring very high availability

RAID 5 – Block and Parity Stripe

RAID 5 organizes block data and parity data across the physical drives. Gener-ally, RAID Level 5 tends to exhibit lower random write performance due to the heavy workload of parity recalculation for each I/O. RAID 5 is generally consi-dered to be the most versatile RAID level. It works well for file, database, appli-cation and web servers.

RAID 5 stripes all drives with data and parity information

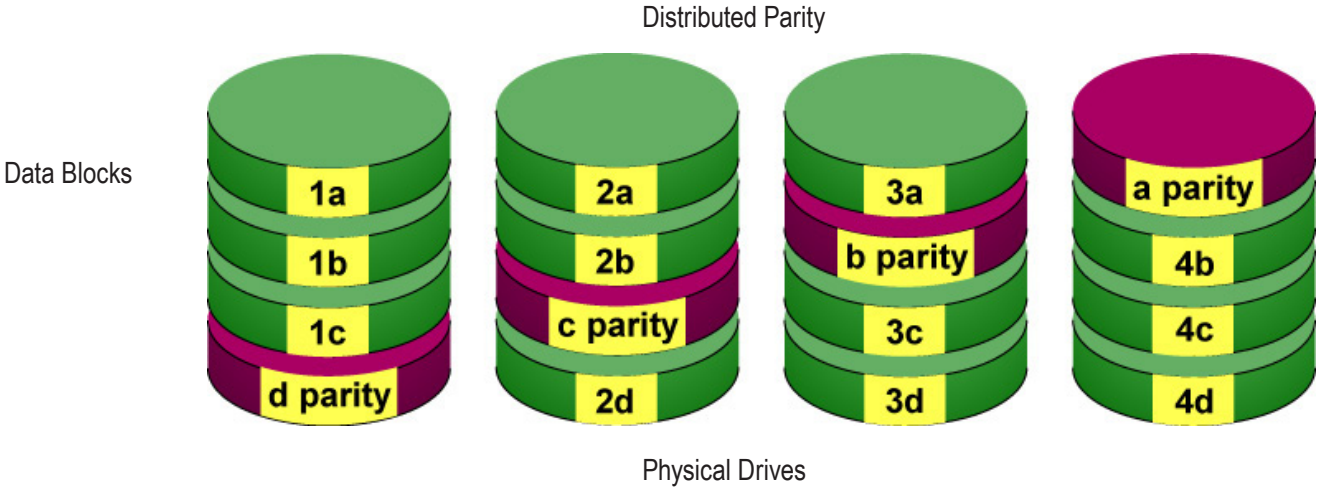

The capacity of a RAID 5 logical drive equals the smallest physical drive times the number of physical drives, minus one. Hence, a RAID 5 logical drive with four 100 GB physical drives has a capacity of 300 GB. A RAID 5 logical drive with two 120 GB physical drives and one 100 GB physical drive has a capacity of 200 GB.

RAID 5 is generally considered to be the most versatile RAID level.

A RAID 5 on VTrak consists of 3 to 32 physical drives.

Advantages	Disadvantages
• High Read data transaction rate • Medium Write data transaction rate • Good aggregate transfer rate • Most versatile RAID level	• Disk failure has a medium impact on throughput

Recommended Applications for RAID 5:

- File and Application servers
- WWW, E-mail, and News servers
- Intranet servers

RAID 6 – Block and Double Parity Stripe

RAID level 6 stores dual parity data is rotated across the physical drives along with the block data. A RAID 6 logical drive can continue to accept I/O requests when any two physical drives fail.

RAID 6 stripes all drives with data and dual parity

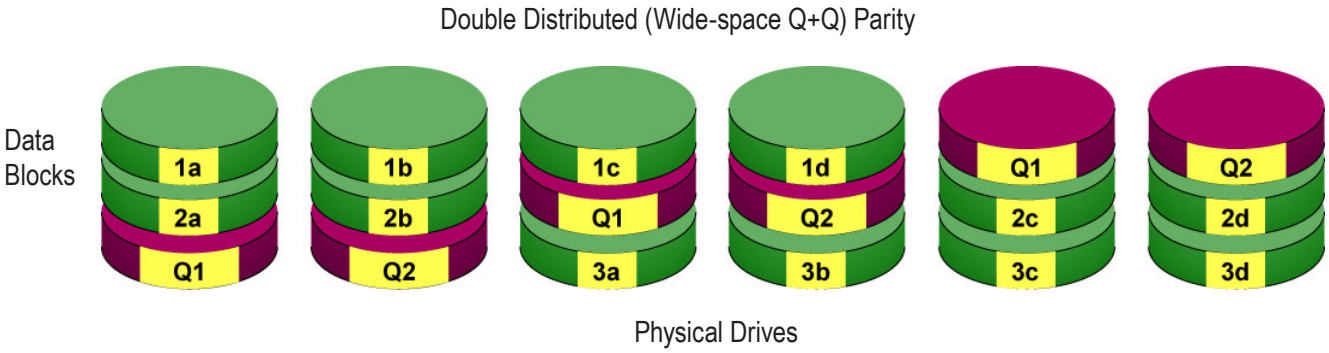

Hence, a RAID 6 logical drive with (7) 100 GB physical drives has a capacity of 500 GB. A RAID 6 logical drive with (4) 100 GB physical drives has a capacity of 200 GB.

RAID 6 becomes more capacity efficient in terms of physical drives as the number of physical drives increases.

RAID 6 provides double fault tolerance. Your logical drive remains available when up to two physical drives fail.

RAID 6 is generally considered to be the safest RAID level.

A RAID 6 on VTrak consists of 4 to 32 physical drives.

Advantages	Disadvantages
• High Read data transaction rate • Medium Write data transaction rate • Good aggregate transfer rate • Safest RAID level, except for RAID 60	• High disk overhead – equivalent of two drives used for parity • Slightly lower performance than RAID 5

Recommended Applications for RAID 6:

- Accounting and Financial
- Database servers
- Any application requiring very high availability

RAID 10 – Mirror + Stripe

Mirror + Stripe combines both of the RAID 1 and RAID 0 logical drive types. RAID 10 can increase performance by reading and writing data in parallel or striping, and duplicating the data, or mirroring.

PROMISE implements RAID 10 by creating a data stripe over one pair of disk drives, then mirroring the stripe over a second pair of disk drives. Some applications refer to this method as RAID 0+1.

PROMISE RAID 10 starts with a data stripe, then mirrors it

The data capacity RAID 10 logical drive equals the capacity of the smallest physical drive times the number of physical drives, divided by two.

In some cases, RAID 10 offers double fault tolerance, depending on which physical drives fail.

RAID 10 arrays require an even number of physical drives and a minimum of four.

Advantages	Disadvantages
• Implemented as a mirrored disk array whose segments are RAID 0 disk arrays • High I/O rates are achieved thanks to multiple stripe segments	• Very high disk overhead – uses only 50% of total capacity

Recommended Applications for RAID 10:

- Imaging applications
- Database servers
- General fileserver

RAID 50 – Striping of Distributed Parity

RAID 50 combines both RAID 5 and RAID 0 features. Data is striped across physical drives as in RAID 0, and it uses distributed parity as in RAID 5. RAID 50 provides data reliability, good overall performance, and supports larger volume sizes.

RAID 50 is a combination of RAID 5 and RAID 0

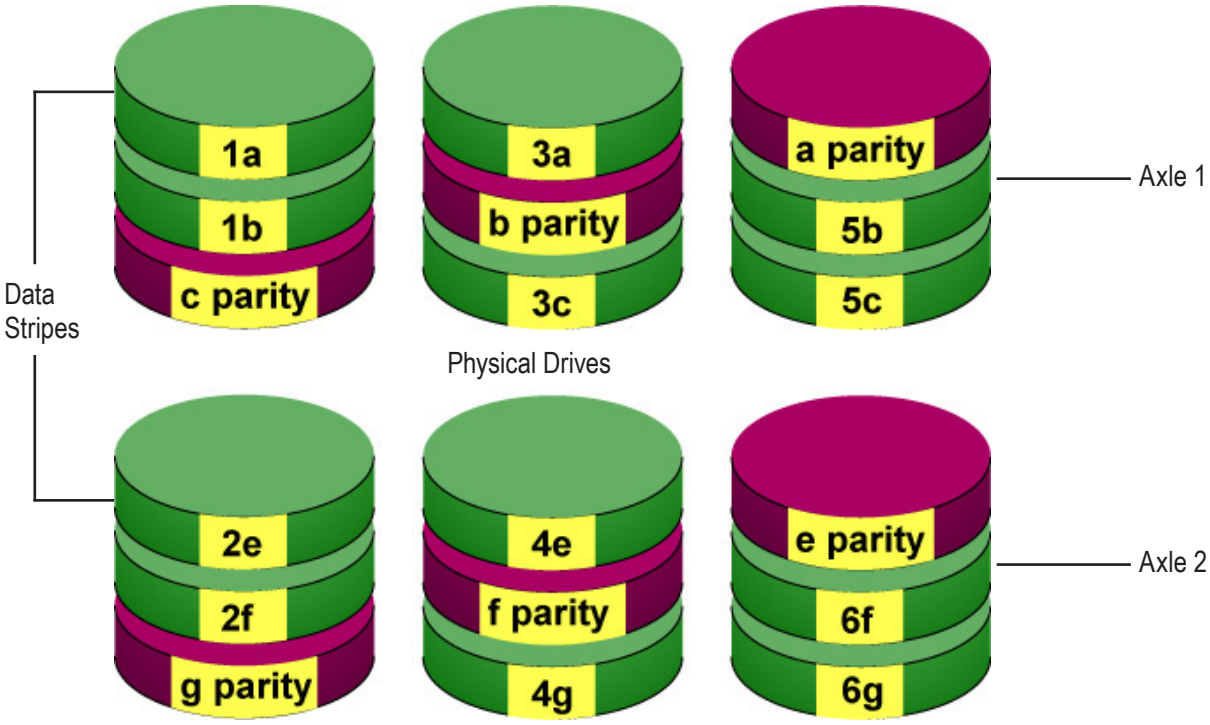

Advantages	Disadvantages
• High Read data transaction rate • Medium Write data transaction rate • Good aggregate transfer rate • High reliability • Supports large volume sizes	• Higher disk overhead than RAID 5

Recommended Applications for RAID 50:

- File and Application servers
- Transaction processing
- Office application with many users accessing small files

The data capacity RAID 50 logical drive equals the capacity of the smallest physical drive times the number of physical drives, minus two.

RAID 50 also provides very high reliability because data is still available even if multiple physical drives fail (one in each axle). The greater the number of axles, the greater the number of physical drives that can fail without the RAID 50 logical drive going offline.

Components	Minimum	Maximum
Number of Axles	2	16
Physical Drives per Axle	3	32
Physical Drives per Logical Drive	6	256

RAID 50 Axles

When you create a RAID 50, you must specify the number of axles. An axle re-fers to a single RAID 5 logical drive that is striped with other RAID 5 logical drives to make RAID 50. An axle can have from 3 to 32 physical drives, de-pending on the number of physical drives in the logical drive.

The chart below shows RAID 50 logical drives with 6 to 32 physical drives, the available number of axles, and the resulting distribution of physical drives on each axle.

RAID 50 Logical Drive

Drives	Axles	Drives/Axle
6	2	3,3
7	2	3,4
8	2	4,4
9	2	4,5
	3	3,3,3
10	2	5,5
	3	3,3,4
11	2	5,5
	3	3,3,4
12	2	6,6
	3	4,4,4
	4	3,3,3,3
13	2	6,7
	3	4,4,5
	4	3,3,3,4
14	2	7,7
	3	4,5,5
	4	3,3,4,4
15	2	7,8
	3	5,5,5
	4	3,4,4,4
	5	3,3,3,3,3
16	2	8,8
	3	5,5,6
	4	4,4,4,4
	5	3,3,3,3,4

RAID 60 – Striping of Double Parity

RAID 60 combines both RAID 6 and RAID 0 features. Data is striped across disks as in RAID 0, and it uses double distributed parity as in RAID 6. RAID 60 provides data reliability, good overall performance and supports larger volume sizes.

RAID 60 is a combination of RAID 6 and RAID 0

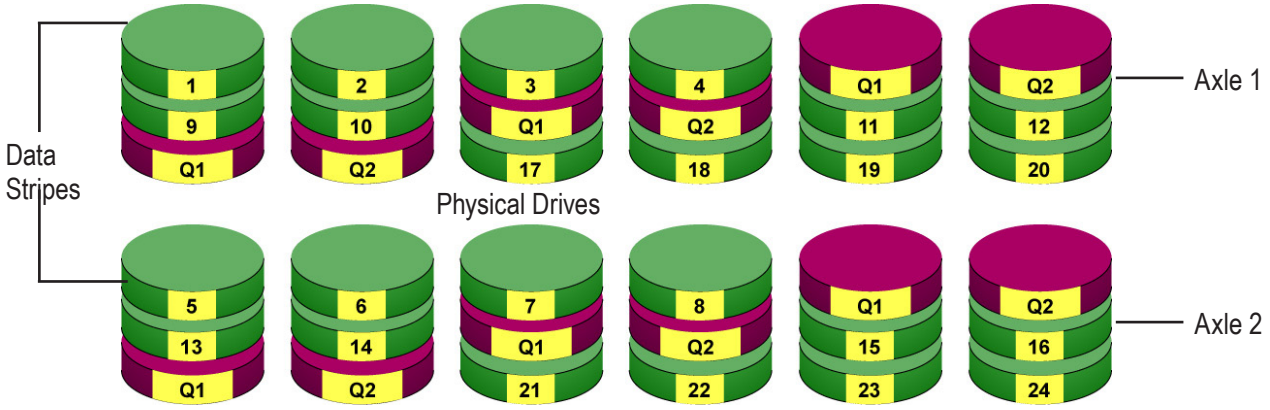

The total capacity of a RAID 60 logical drive is the smallest physical drive times the number of physical drives, minus four.

RAID 60 also provides very high reliability because data is still available even if multiple physical drives fail (two in each axle). The greater the number of axles, the greater the number of physical drives that can fail without the RAID 60 logical drive going offline.

Advantages	Disadvantages
• High Read data transaction rate • Medium Write data transaction rate • Good aggregate transfer rate • Safest RAID level	• High disk overhead – equivalent of two drives used for parity • Slightly lower performance than RAID 50

Recommended Applications for RAID 60:

- Accounting and Financial
- Database servers
- Any application requiring very high availability

Components	Minimum	Maximum
Number of Axles	2	16
Physical Drives per Axle	4	32
Physical Drives per Logical Drive	8	256

RAID 60 Axles

When you create a RAID 60, you must specify the number of axles. An axle refers to a single RAID 6 logical drive that is striped with other RAID 6 logical drives to make RAID 60. An axle can have from 4 to 32 physical drives, de-pending on the number of physical drives in the logical drive.

RAID 60 Logical Drive

Drives	Axles	Drives/Axle
8	2	4,4
9	2	4,5
10	2	4,5
11	2	5,6
12	2	6,6
	3	4,4,4
13	2	6,7
	3	4,4,5
14	2	7,7
	3	4,5,5
15	2	7,8
	3	5,5,5
16	2	8,8
	3	5,5,6
	4	4,4,4,4
17	2	8,9
	3	5,6,6
	4	4,4,4,5
18	2	9,9
	3	6,6,6
	4	4,4,5,5
19	2	9,10
	3	6,6,7
	4	4,5,5,5
20	2	10,10
	3	6,6,7
	4	5,5,5,5
	5	4,4,4,4,4

RAID Level Migration

The term “Migration” means either or both of the following:

- Change the RAID level of a logical drive.
- Expand the storage capacity of a logical drive.

On VTrak, RAID level migration is performed on the disk array but it applies to the logical drives. Migration does not disturb your data. You can access the data while the migration is in progress. When migration is done, your disk array has a different RAID level and/or a larger capacity.

Migration Requirements

The following conditions affect RAID level migration:

- The disk array and logical drive must show a green check icon.
- The Target disk array may require more physical drives than the Source disk array.
- If the Target disk array requires an EVEN number of physical drives but the Source disk array has an ODD number, ADD a physical drive as part of the migration process.
- You cannot reduce the number of physical drives in your disk array, even if the Target disk array requires fewer physical drives than the Source disk array.
- RAID 1 (mirroring) works with two drives only. Only a single-drive RAID 0 disk array can migrate to RAID 1. Other RAID Levels use too many drives to migrate.
- You cannot migrate a disk array when it is Critical or performing activities such as Synchronizing, Rebuilding, and PDM.
- For RAID 6 or RAID 60, you can only migrate between these two RAID le-vels. Destination RAID 60 axles can have up to 16 physical drives. Other limitations might apply.

Source and Target RAID Levels

The tables on the following pages show the migration options for each source logical drive by its RAID level. The available target RAID levels are shown with their requirements.

RAID 0

A RAID 0 source logical drive can migrate to the following target logical drives:

Target	Requirements
RAID 0	Add physical drives.
RAID 1	2 physical drives only. Only a single-drive RAID 0 can migrate to RAID 1 by adding 1 physical drive.
RAID 5	3 physical drives minimum, 32 maximum. RAID 0 must have less than 16 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 6	4 physical drives minimum, 32 maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	4 physical drives minimum. Even number of physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 50	6 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

.

RAID 1

A RAID 1 source logical drive can migrate to the following target logical drives:

Target	Requirements
RAID 0	None.
RAID 5	3 physical drives minimum, 32 maximum. RAID 1 must have less than 32 physical drives. Add 1 or more physical drives.
RAID 10	4 physical drives minimum. Even number of physical drives. Add 2 or more physical drives.
RAID 50	6 physical drives minimum, 32 per axle maximum. Add 4 or more physical drives.

RAID 5

A RAID 5 source logical drive can migrate to the following target logical drives:

Target	Requirements
RAID 0	None.
RAID 5	Add physical drives. 32 maximum.
RAID 6	4 physical drives minimum, 32 maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	4 physical drives minimum. Even number of physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 50	6 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

.

RAID 6

A RAID 6 source logical drive can migrate to the following target logical drives:

Target	Requirements
RAID 6	Add physical drives. 32 maximum.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

RAID 10

A RAID 10 source logical drive can migrate to the following target logical drives:

Target	Requirements
RAID 0	None.
RAID 1	None.
RAID 5	3 physical drives minimum, 32 maximum. RAID 10 must have less than 16 physical drives.
RAID 6	4 physical drives minimum, 32 maximum. RAID 10 must have less than 32 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	Add physical drives. Even number of physical drives.
RAID 50	6 physical drives minimum, 32 per axle maximum.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

If you want a RAID 10 logical drive, there must be an even number of physical drives and you must specify RAID 10 for the target logical drive.

RAID 50

A RAID 50 source logical drive can migrate to the following target logical drives:

Target	Requirements
RAID 0	None.
RAID 5	32 physical drives maximum.
RAID 6	32 physical drives maximum. RAID 50 must have less than 32 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 10	Even number of physical drives.
RAID 50	Add physical drives. 32 per axle maximum.
RAID 60	8 physical drives minimum, 32 per axle maximum. If existing physical drives have no unused space, add 1 or more physical drives.

You can add physical drives to a RAID 50 array but you cannot change the number of axles.

RAID 60

A RAID 60 source logical drive can migrate to the following target logical drives:

Target	Requirements
RAID 6	32 physical drives maximum. RAID 60 must have less than 32 physical drives. If existing physical drives have no unused space, add 1 or more physical drives.
RAID 60	Add physical drives. 32 per axle maximum.

You can add physical drives to a RAID 60 array but you cannot change the number of axles.

Stripe Size

Stripe Size, also called “Stripe Block Size,” refers to the size of the data blocks written to, and read from, the physical drives. Stripe Size is specified when you create a logical drive. You can choose Stripe Size directly when you use the Wizard Advanced Configuration function to create a logical drive.

You cannot change the Stripe Size of an existing logical drive. You must delete the logical drive and create a new one.

The available Stripe Sizes are 64 KB, 128 KB, 256 KB, 512 KB, and 1 MB. 64 KB is the default. There are two issues to consider when choosing the Stripe Size:

- You should choose a Stripe Size equal to, or smaller than, the smallest cache buffer found on any physical drive in the disk array. Selecting a larger value slows read/write performance because physical drives with smaller cache buffers need more time for multiple accesses to fill their buffers.
- If your data retrieval consists of fixed data blocks, such as with some data-base or video applications, then you should choose that size as your Stripe Size.
- If you do not know the cache buffer or fixed data block sizes, choose 64 KB as your Stripe Size. Generally speaking,
- Email, POS, and web servers prefer smaller stripe sizes.
- Video and database applications prefer larger stripe sizes.

Sector Size

A sector is the smallest addressable area on a physical drive. Sector size refers to the number of data bytes a sector can hold. A smaller sector size is a more efficient use of a physical drive’s capacity. 512 bytes (512 B) is the most common sector size, and the default in WebPAM PROe.

Preferred Controller ID

When you create a logical drive using the Advanced method of disk array creation, you can specify the Preferred Controller ID:

- Controller 1 – Assign all logical drives to Controller 1
- Controller 2 – Assign all logical drives to Controller 2.
- Automatic – Alternate logical drive assignments between Controllers 1 and 2.

Automatic is the default and preferred setting because it balances the logical drive assignments for you.

See “Creating a Disk Array Manually” on page 190, “Creating a Disk Array with the Wizard” on page 191, and “Advanced Configuration” on page 86.

Initialization

Initialization is done to logical drives after they are created from a disk array. Full initialization sets all data bits in the logical drive to a specified pattern, such as all zeros. The action is useful because there may be residual data on the logical drives left behind from earlier configurations. For this reason, Initialization is recommended for all new logical drives. See “Initializing a Logical Drive” on page 207.

Caution

When you initialize a logical drive, all the data on the logical drive is lost. Backup any important data before you initialize a logical drive.

Partition and Format

Like any other type of fixed disk media in your system, a RAID logical drive must also be partitioned and formatted before use. Use the same method of partitioning and formatting on an logical drive as you would any other fixed disk.

Depending on the operating system you use, there may or may not be various capacity limitations applicable for the different types of partitions.

Spare Drives

Spare drive technology includes:

- “Definition” (below)
- “Options” (below)
- “Requirements” on page 384
- “Transition” on page 384

Definition

A spare drive is a physical drive that you designate to automatically replace the failed physical drive in a disk array. See “Creating a Spare Drive Manually” on page 215.

The general recommendation is to:

- Provide at least one spare drive for every 16 physical drives in the RAID system
- Configure the spares as global revertible spare drives

Options

There are several options you can specify for a spare drive:

- System Options
- Revertible – Returns to its spare drive assignment after you replace the failed physical drive in the disk array and run the Transition function.
- Media Patrol – By default, Media Patrol runs on spare drives unless you disable it.
- Spare Type
- Global – Can be used by any disk array
- Dedicated – Can be used only by the assigned disk array
- Media Type (type of physical drive)
- Hard Disk Drive (HDD)
- Solid State Drive (SSD)

Requirements

The spare drive must:

- Have adequate capacity to replace the largest physical drive in your disk arrays.
- Be the same media type as the physical drives in your disk arrays.
- A revertible spare drive requires:
- You to replace the failed physical drive in the disk array
- You to run the Transition function

Transition

Transition is the process of replacing a revertible spare drive that is currently part of a disk array with an unconfigured physical drive or a non-revertible spare. The revertible spare drive returns to its original status. In order to run the Transition function, the spare drive must be revertible.

In addition, you must specify an unconfigured physical drive of the same or larger capacity and same media type as the revertible spare drive.

Running a Transition

The Transition feature enables you to specify “permanent” spare drives for your VTrak subsystem. Transition is the process of replacing a revertible spare drive that is currently part of a disk array with an unconfigured physical drive or a non-revertible spare. The revertible spare drive returns to its original status.

Transition happens automatically when the following sequence of events takes place:

- You create a revertible spare drive. See “Creating a Spare Drive Manually” on page 215.
- A physical drive assigned to your disk array fails and the array goes critical or degraded.
- VTrak automatically rebuilds your array to the revertible spare drive and the array becomes functional again.
- You replace the failed physical drive with a new physical drive of equal or greater capacity.
- VTrak automatically transitions (moves) the data from the revertible spare to the new physical drive.

The new physical drive becomes part of the array and the revertible spare drive returns to its original spare status.

Transition happens manually when you specify a different unconfigured physical drive to transition (move) the data from the revertible spare drive.

See the example on the following pages.

Example

Following is an example to explain the Transition function.

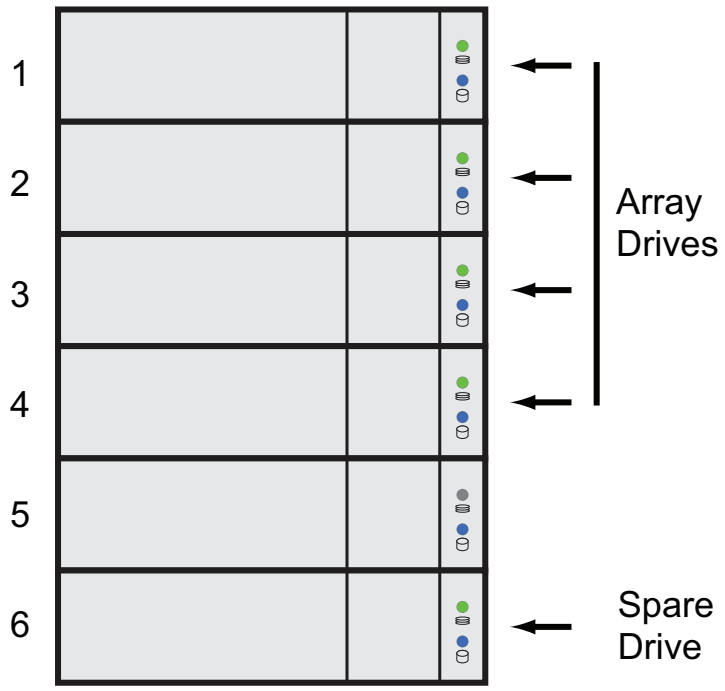

In the example above, there is a four-drive RAID 5 disk array and a global spare drive. Physical drives 1, 2, 3, and 4 belong to the disk array. Physical drive 5 remains unconfigured. Physical drive 6 is a revertible spare drive.

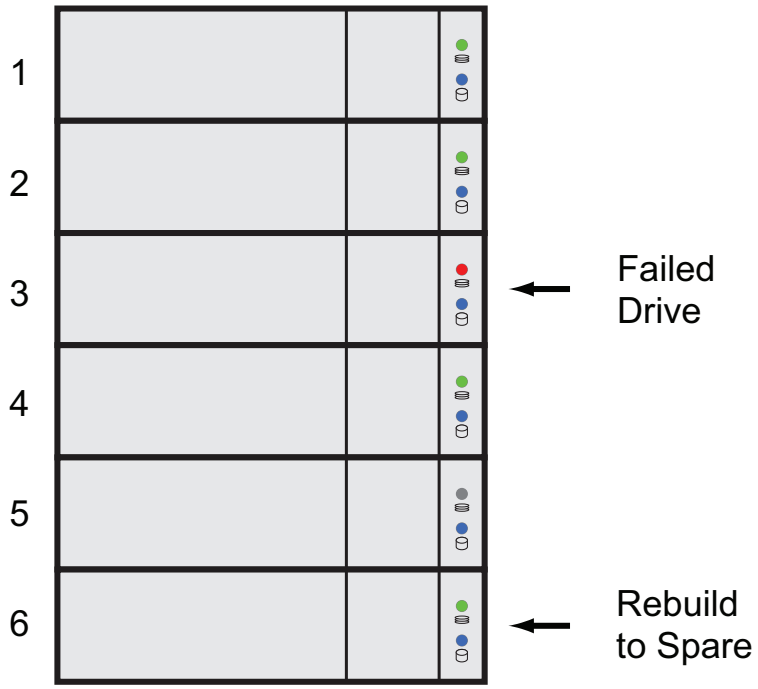

If a physical drive fails in a disk array and there is a spare drive of adequate capacity available, the controller automatically rebuilds the array using the spare drive. In this example, physical drive 3 failed and the array is rebuilt using physical drive 6, the revertible spare drive.

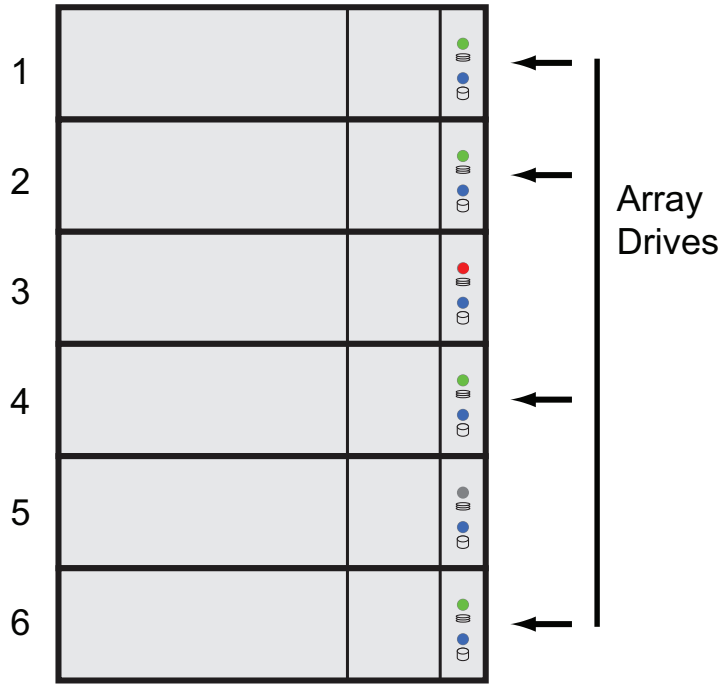

When the rebuild is complete, the spare drive has replaced the failed drive. In this example, failed drive 3 was replaced by spare drive 6. The disk array now consists of physical drives 1, 2, 4, and 6.

There is no spare drive at this moment. Even if physical drive 5 is of adequate capacity, it has not been designated as a spare, therefore the controller cannot use it as a spare.

Automatic Transition

At this juncture, you would replace the failed drive in slot 3 with a new one of the same or greater capacity.

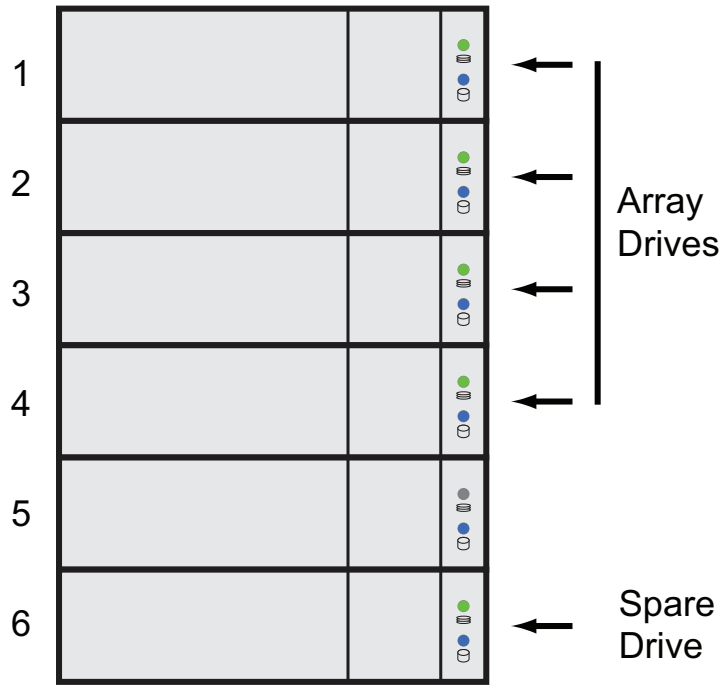

When the VTrak controller detects the new drive in slot 3, the controller:

- Automatically transitions the data on drive 6 to drive 3
- Returns drive 6 to spare status

When the Automatic Transition is finished, physical drives 1, 2, 3, and 4 belong to the disk array and physical drive 6 is a revertible spare drive. The original configuration is restored.

Manual Transition

If you wanted to use the drive in slot 5 as a member of the disk array, rather than the drive in slot 3, you would run the Transition function manually. See “Running a Transition on a Spare Drive” on page 218.

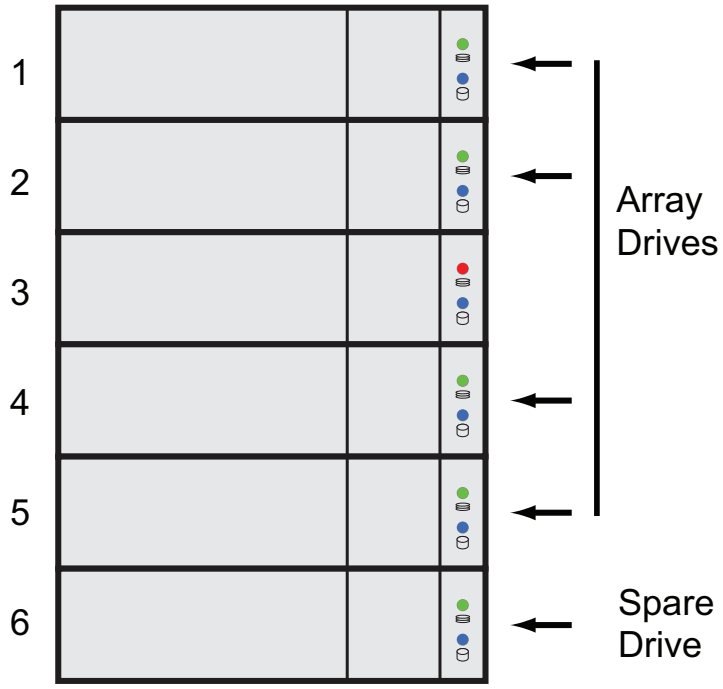

When the Manual Transition is finished, physical drives 1, 2, 4, and 5 belong to the disk array and physical drive 6 is a revertible spare drive.

At this point, you would replace the drive in slot 3. The new drive in slot 3 re-mains unconfigured until you assign it to a disk array or as a spare.

RAID Controllers

RAID controller technology includes;

- "LUN Affinity" (below)
- "ALUA" (below)
- "Cache Policy" on page 391
- "Read Cache Policy" on page 391
- "Write Cache Policy" on page 392
- "Power Saving" on page 394
- "Capacity Coercion" on page 394

LUN Affinity

VTrak subsystems with dual RAID controllers include a LUN Affinity feature. Normally, either controller can access all logical drives. LUN Affinity enables you to specify which controller can access each logical drive. Use this feature to balance the load of your logical drives between the two controllers.

To use LUN Affinity you must:

- Have two RAID controllers in the subsystem.
- Set the redundancy type to Active-Active. See "Making Subsystem Settings" on page 101.
- Enable LUN Affinity. See "Making Controller Settings" on page 113.

On subsystems with two RAID controllers, when Cache Mirroring is disabled, LUN Affinity is enabled automatically.

ALUA

VTrak supports Asymmetric Logical Unit Access (ALUA) on Linux OSes. ALUA is a multipathing tool. It enables an initiator (your host PC or server) to discover target port groups that provide a common failover / failback behavior for your LUNs. ALUA enables the host to see which paths are in an optimal state and which are not.

To use ALUA you must:

- Have two RAID controllers in the subsystem.
- Set the redundancy type to Active-Active. See “Making Subsystem Settings” on page 101.
- Enable LUN Affinity and ALUA. See “Making Controller Settings” on page 113.

Cache Policy

As it is used with VTrak, the term cache refers to any of several kinds of high-speed, volatile memory that hold data moving from your computer to the physical drives or vice-versa. Cache is important because it can read and write data much faster than a physical drive. There are read caches, which hold data as it is read from a physical drive; and write caches, which hold data as it is written to a physical drive.

In order to tune the cache for best performance in different applications, user-adjustable settings are provided. Cache settings are made on the RAID controller. See “Making Controller Settings” on page 113.

Read Cache Policy

- Read Cache – The read cache is enabled but no pre-fetch action.
- Read Ahead – The read cache and predictive pre-fetch feature are enabled. Read-ahead anticipates the next read and performs it before the request is made. Can increase read performance.
- Forced Read Ahead – The read cache and aggressive pre-fetch feature are enabled. See “Forced Read-Ahead Cache” below.
- No Cache – The read cache is disabled.

Write Cache Policy

- Write Back – Data is written first to the cache, then to the logical drive. Better performance. VTrak has a cache backup battery to protect data in the cache from a sudden power failure.
- Adaptive Writeback – See “Adaptive Writeback Cache” below.
- Write Thru – Also “Write Through.” Data is written to the cache and the logical drive at the same time. Safer.

If your write cache policy is set to Write Back, the write policy automatically changes to Write Thru when all of the following conditions occur:

- The logical drive write policy is set to Write Back
- The Adaptive Writeback Cache feature is enabled
- The cache backup battery goes offline

When the battery comes back online, the write policy automatically changes back to Write Back.

Also see “Viewing Battery Information” on page 119.

Forced Read-Ahead Cache

On the VTrak subsystem, you can set the logical drive read cache policy to Forced Read Ahead an enable the aggressive pre-fetch feature.

The Forced Read-Ahead cache policy setting provides predictive pre-fetching of data requests, allowing the controller to aggressively buffer large chunks of data in cache memory to prevent frame drops on high-bandwidth video playback. Not normally enabled for non-video applications.

Adaptive Writeback Cache

On the VTrak subsystem, you can set the logical drive write cache policy to Write Thru or Write Back.

If you set the write cache policy to Write Back, your data is first written to the controller cache, and later to the logical drive. This action is conducted to improve performance. In order to preserve the data in the cache in the event of a power failure, the subsystem has a backup battery that provides continuous power to maintain the cache. To see an estimate of how long the battery can power the cache, see “Viewing Battery Information” on page 119.

The Adaptive Writeback Cache feature protects your data by changing the write cache settings while the cache backup battery is offline. When all of the following conditions occur:

- The logical drive write policy is set to Write Back.
- The Adaptive Writeback Cache feature is enabled.
- The cache backup battery goes offline. (See definition in Note below)

The write policy automatically changes to Write Thru. When the battery comes back online, the write policy automatically changes back to Write Back.

To enable the Adaptive Writeback Cache option, see “Making Controller Settings” on page 113.

Notes

The condition “cache battery goes offline” can be due to one of the following circumstances:

- Battery not installed.
- Battery remaining capacity is not enough to keep data.

Host Cache Flushing

On the VTrak subsystem, you can enable or disable host cache flushing.

When enabled, host cache flushing guards against data loss in the event of a power failure. However RAID performance is slightly reduced.

When disabled, the VTrak subsystem has greater sustained bandwidth and lower latency, which are helpful for real-time video capture.

When you operate the VTrak with host cache flushing disabled, use a UPS to protect against data loss.

Preferred Controller ID

See “Preferred Controller ID” on page 381.

Power Saving

Power saving is a method of conserving energy by applying specific actions to hard disk drives (HDD). After an HDD has been idle for the set period of time, you can elect to:

- Park the read/write heads – Referred to as Power Saving Idle Time on VTrak.
- Reduce disk rotation speed – Referred to as Power Saving Standby Time on VTrak.
- Spin down the disk (stop rotation) – Referred to as Power Saving Stopped Time on VTrak.

Power management must be:

- Set on the RAID controller. See “Making Controller Settings” on page 113.
- Enabled on each HDD. See “Making Disk Array Settings” on page 193.

Capacity Coercion

This feature is designed for fault-tolerant logical drives (RAID 1, 5, 10, 50, and 60). It is generally recommended to use physical drives of the same size in your disk arrays. When this is not possible, the system adjusts for the size differences by reducing or coercing the capacity of the larger drives to match the smaller ones. With VTrak, you can choose to enable capacity coercion and any one of four methods.

Enable capacity coercion and choose a method, see See “Making Controller Settings” on page 113.

- GB Truncate – (Default) Reduces the useful capacity to the nearest 1,000,000,000 byte boundary.
- 10GB Truncate – Reduces the useful capacity to the nearest 10,000,000,000 byte boundary.
- Group Rounding – Uses an algorithm to determine how much to truncate. Results in the maximum amount of usable drive capacity.
- Table Rounding – Applies a predefined table to determine how much to truncate.

Capacity coercion also affects a replacement drive used in a disk array. Normally, when an physical drive fails, the replacement drive must be the same capacity or larger. However, the capacity coercion feature permits the installation of a replacement drive that is slightly smaller (within 1 gigabyte) than the remaining working drive. For example, the remaining working drives can be 80.5 GB and the replacement drive can be 80.3, since all are rounded down to 80 GB. This permits the smaller drive to be used.

Without capacity coercion, the controller does not permit the use of a replacement physical drive that is slightly smaller than the remaining working drives.

Internet Protocols

VTrak supports the IPv4 and IPv6 protocols.

Protocol	Addresses		Example
IPv4	32-bits	4.3 x 109	192.168.10.85
IPv6	128-bits	3.4 x 1038	2001:0000:0000:0000:0000:0000:e2a8:4337 *Abbreviated* 2001:0:0:0:0:0:e2a8:4337

TROUBLESHOOTING

This chapter contains the following topics:

VTrak E5000 is Beeping

VTrak's alarm has five different patterns, as shown below.

Audible alarm sound patterns

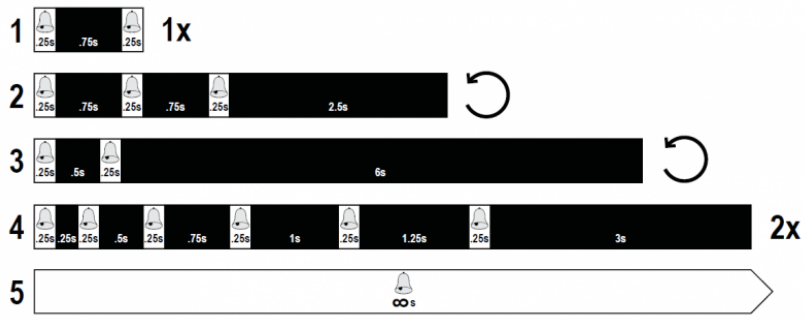

When you first power-up the VTrak, it beeps twice to show normal operation.

See pattern 1, in the figure above.

The audible alarm sounds at other times to inform you that the VTrak needs attention. But the alarm does not specify the condition.

When the alarm sounds:

- Check the front and back of VTrak enclosure for red or amber LEDs.
- If email notification is enabled, check for new messages.
- Check for yellow **!** red **X** icons.
- Check the event log.
- See "Viewing Runtime Events" on page 162 and "Viewing NVRAM Events" on page 163.

When a continuous tone sounds, there are multiple alarm patterns sounding at the same time.

Silencing the Buzzer

Caution

This action disables the buzzer for all events.

To silence the buzzer:

1. Click the **Device** tab.
2. Click the **Component List** icon.
3. Click the **Buzzer** and click the **Settings** button.
4. Uncheck the **Enable Buzzer** box.
5. Click the **Save** button.

Muting the alarm buzzer

To silence the audible alarm, press the Mute button of the front of the enclosure. This button is used to toggle the audible alarm buzzer on (enabled) and off (disabled). The Mute Alarm LED lights when the buzzer has been disabled by the Mute button.

VTrak E5600fD front left side

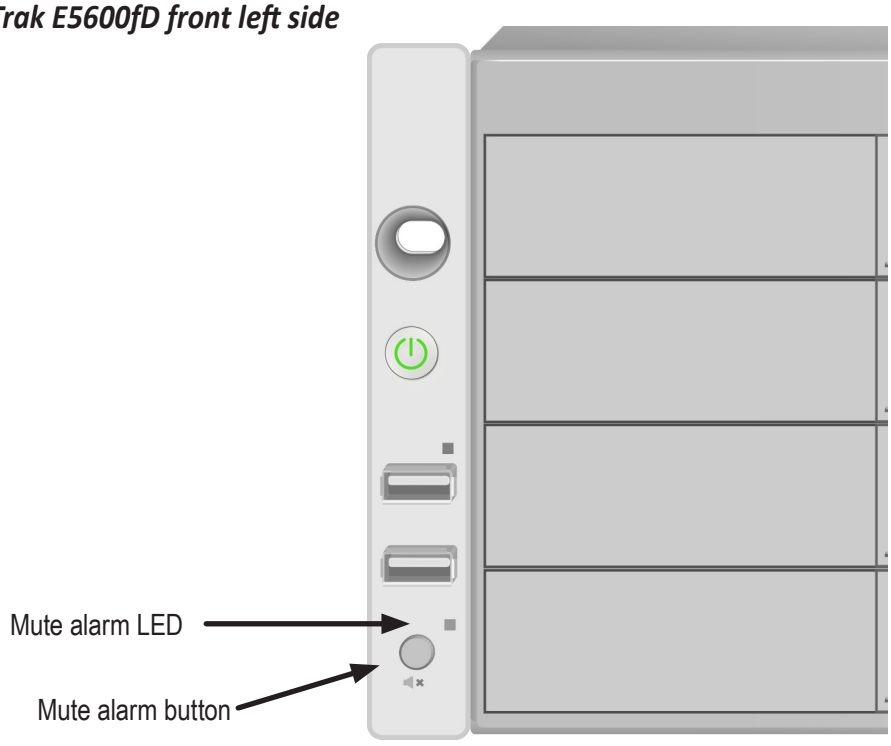

LEDs on the Front of the VTrak E5000

When the power is on, the LEDs on the front of the VTrak E5000 light up.

LED indicators on front right of VTrak E5600fD/E5800fD

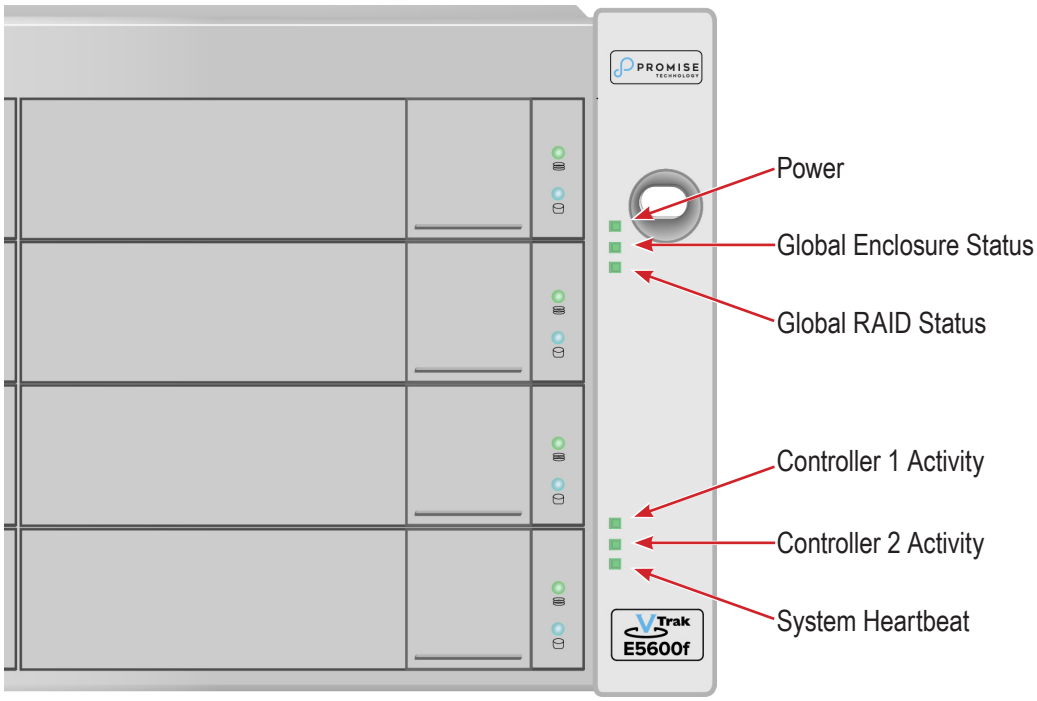

When boot-up is finished and the VTrak is functioning normally:

- **Power**, **Global Enclosure Status**, and **Global RAID Status** LEDs display green continuously.
- **Controller Activity** LED flashes green when there is controller activity.
- **System Heartbeat** LED blinks blue (once a second), and repeats the pattern.

Steady means the LED is on.

Blinking means a regular on/off pattern.

Flashing means an intermittent and irregular on/off pattern.

Dark means the LED is off.

See the table below.

Front right LED Behavior After Boot Up

State	Power	Global Enclosure	Global RAID	Controller Activity	Controller Heartbeat
Dark	No power	No power	—	Unit is off, or controller is not present or FC is not cable connected	—
Steady Green	Normal	Normal	Normal	Unit is up, controller is present and running, at least one FC cable is connected and a link is established.	—
Blinking Green	—	—	—	—	Normal**
Flashing Green	—	—	—	Activity	—
Amber	—	Problem*	Critical	—	—
Red	—	Failure*	Offline	—	—

* Check the LEDs on the back of the VTrak enclosure (controllers and PSUs).

** Blinks blinks green once per second for two seconds for dual controller enclosure; blinks every four seconds for single controller enclosure.

Steady means the LED is on.

Blinking means a regular on/off pattern.

Flashing means an intermittent and irregular on/off pattern.

Dark means the LED is off.

“Enclosure Problems” on page 417, “RAID Controller Problems” on page 421, and “Physical Drive Problems” on page 426.

The Locator feature triggered from WebPAM PROe causes the LEDs to blink on and off for one minute. That action helps you find the physical component.

Drive Carrier LEDs

The VTrak spins up the disk drives sequentially to equalize power draw during start-up. After a few moments:

- The Power/Activity LED displays blue when a physical drive is present.
- The Drive Status LED displays green when the physical drive is configured as a member of a disk array or as a spare. When the physical drive is unconfigured, the LED is dark.

Drive carrier LEDs

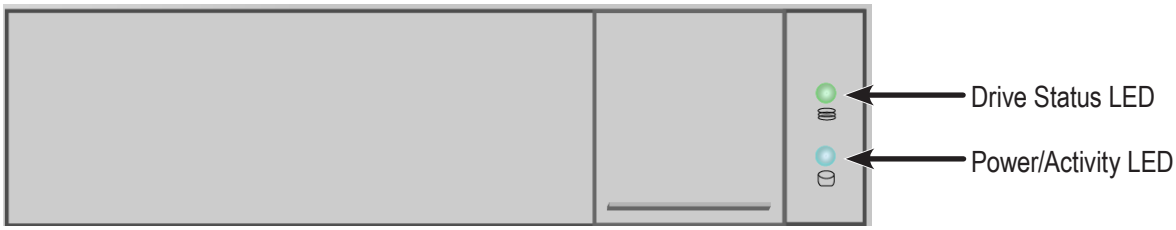

Enclosure Front LEDs on drive carriers

State	Power / Activity	Drive Status
Dark	No drive in carrier	Drive is unconfigured
Steady Blue	Drive is present	—
Flashing Blue	Activity on drive	—
Steady green	—	Drive is configured
Blinking green	—	Locator feature
Amber	—	Drive is rebuilding
Red	—	Drive error or failure

- *Configured* means the physical drive either belongs to an array or it is assigned as a spare drive.
- *Steady* means the LED is on.
- *Blinking* means a regular on/off pattern.
- *Flashing* means intermittent and irregular on/off pattern.

"Enclosure Problems" on page 417, "RAID Controller Problems" on page 421, and "Physical Drive Problems" on page 426.

The Locator feature triggered from WebPAM PROe causes the LEDs to blink on and off for one minute. That action helps you find the physical component.

Controller LED Behavior

The table below describes behavior of the LED indicators on the VTrak E5000 Series controller.

Controller LEDs

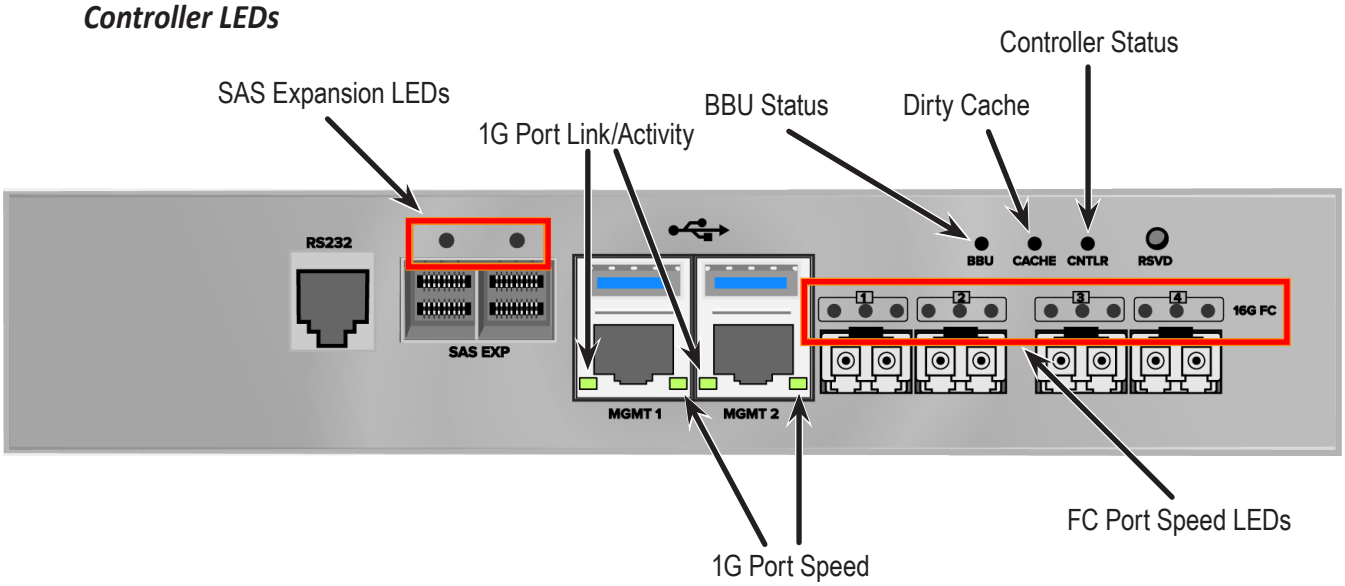

Controller LED Behavior

When boot-up is finished and the VTrak E5000 subsystem is functioning normally:

- Controller status LEDs display green continuously.
- Ethernet LEDs display green or flash depending on your network connection.
- The FC, SAS, and Expansion LEDs display green or flash during port activity.

The table below describes behavior of the LED indicators on the VTrak E5000 Series controller.

LED	Description
SAS Expansion	One LED indicator for each SAS expansion port. These light green when connected, and flash green when there is activity.
Controller Status	This displays the current operational status of the controller. A steady (unblinking) green light indicates the controller is operational. This will blink green during when using the controller locator feature. A blinking amber light indicates a problem. Steady red light indicates controller failure. A flashing red light means the controller is in Maintenance Mode (offline while in Maintenance Mode).
Dirty Cache	Lights steady amber if cache is dirty, meaning that the controller memory cache contains data, otherwise this is dark. This will blink green during when using the controller locator feature.
Battery Status	This lights steady green when the battery status is healthy (normal). Red indicates the battery has failed. A steady amber light indicates there is not enough reserve power in the battery to backup cache memory if the power fails.
FC ports	See next page

LED indicators for Fibre Channel ports, three LEDs for each FC port

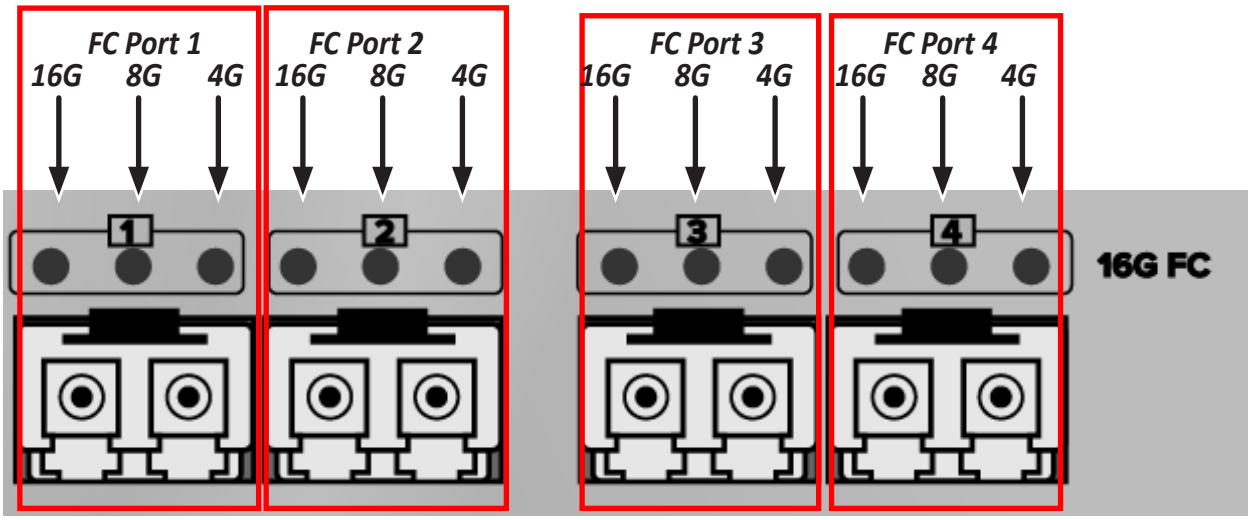

Fiber Channel Port LED behavior

LEDs	Power on (Before Firmware Initialization)	Power on (After Firmware Initialization)	Firmware Fault	4 Gbps Link Up/ ACT	8 Gbps Link Up/ ACT	16 Gbps Link Up/ ACT
Green LED (16 Gbps)	On	Flash	Flash in sequence	Off	Off	On/Flash when active
Green LED (8 Gbps)				Off	On/Flash when active	Off
Green LED (4 Gbps)				On/Flash when active	Off	Off

All Fibre Channel port LED indicators will be dark when the system is powered off. If all three indicators for a port flash simultaneously, then there is no SFP transceiver installed, or the wrong type of transceiver is installed, or the port is not connected.

VTrak E5600 / VTrak E5800 Power Supply LEDs

The LEDs on the rear panel of the VTrak E5600 and VTrak E5800 include a single status LED on each power supply. These PSU status LED will light green to indicate normal operation. A red LED indicates a problem or unit failure.

Status LED on Power Supply

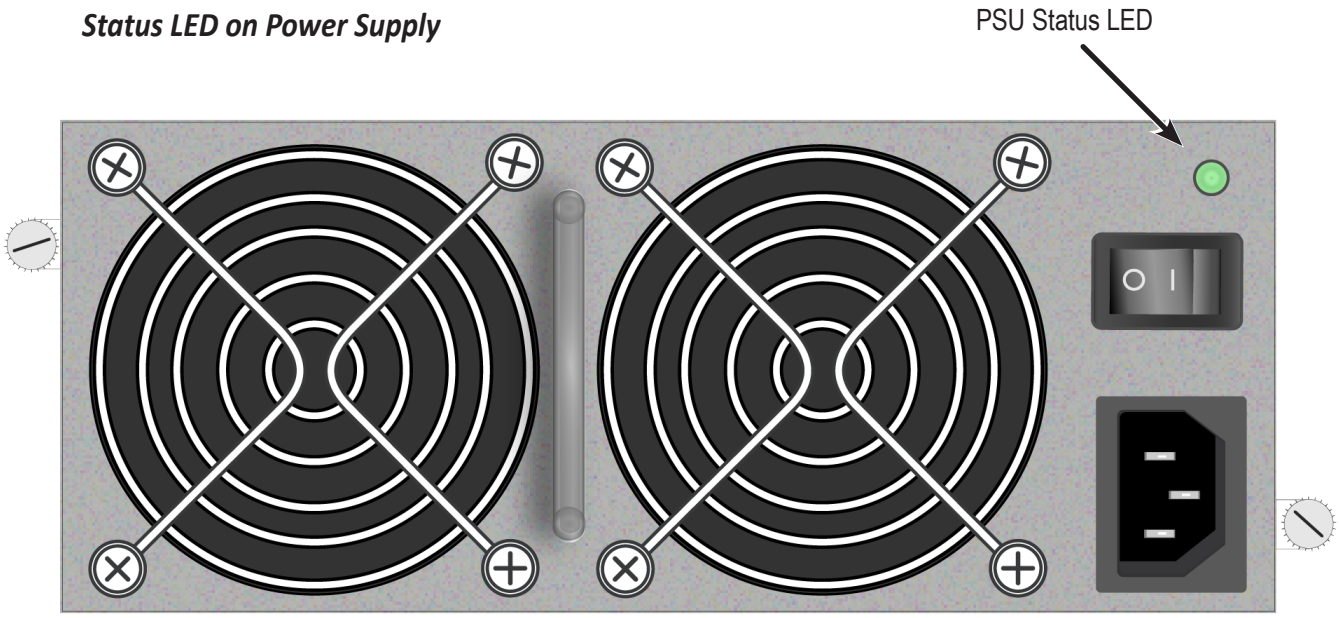

VTrak E5300 / VTrak E5320 Power Supply LEDs

The power supplies for the VTrak E5300 and VTrak E5320 have five LED indicators on each unit. See the table below for PSU LED behavior.

Power supply LED indicators on VTrak E5300

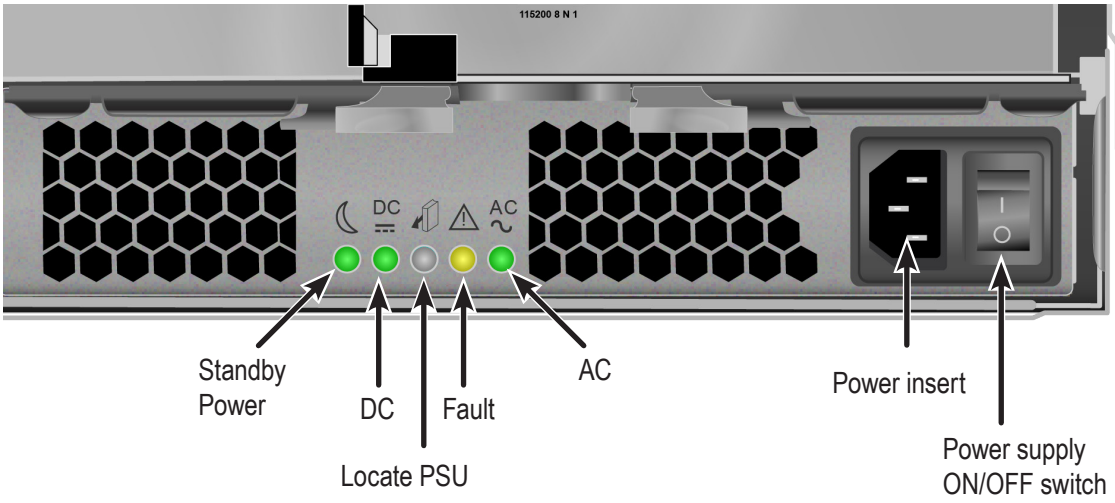

VTrak E5300fD/E5800fD Power Supply LEDs

LED	Description
Standby Power	This indicates the standby power status. If the standby power is on and the power switch is off, the LED lights green. If the standby power is on and the power switch is on, then the LED will be off.
DC	This indicates if the power supply is properly inserted into the enclosure. The LED lights green when the power supply is properly inserted and the power supply is switch on. It is off if the power supply is not properly inserted, or when the power supply is switched off.
Fault	This indicates the overall health status of the power supply. When the power supply is functioning normally and no problems are detected, it will be off. If a problem with the power supply is detected, it light amber. Note that this LED blinks once when the power supply is switched on, then remains off unless there is a problem.
AC	This indicates that input power is present. When the power supply is switched on, it lights green if input power is available. It is off if power is not present and when the power supply is switched off.
Locate PSU	This flashes blue when using the *Locate PSU* function.

Checking Component Installation

To check a component's installation, remove the component, then reinstall the component in its original location. In most cases, this action fixes a bad connection and allows VTrak to detect the component. If this action does not correct the problem, replace the unit.

On VTrak E5000 systems with dual controllers, when one controller's Status LED is amber and the other controller's Status LED is flashing red, it means that the controller with the flashing red LED has entered maintenance mode. See "RAID Controller Problems" on page 421.

If the Controller Status LED continues to display amber after startup, contact PROMISE Technical Support. See "Contacting Technical Support" on page 469.

The Dirty Cache LED flashes during input/output operation. If the LED shines amber and the power is off, there is unsaved data in the cache. Do NOT power down the VTrak while this LED is on.

Viewing Runtime Events

To display Runtime Events:

1. From the **Main Menu**, highlight Event Viewer and press **Enter**.

The log of Runtime Events appears. Events are added to the top of the list. Each item includes:

- **Sequence number** – Begins with 0 at system startup.
- **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
- **Severity** See "Event Severity Levels" on page 415.
- **Timestamp** – Date and time the event happened.
- **Description** – A description of the event in plain language.

2. Press the up and down arrow keys to scroll through the log.

Viewing NVRAM Events

This screen displays a list of and information about 63 most important events over multiple subsystem startups.

To display NVRAM events:

1. From the **Main Menu**, highlight Event Viewer and press **Enter**.
2. Highlight **NVRAM Events** and press **Enter**.

   The log of Runtime Events appears. Events are added to the top of the list. Each item includes:

   - **Sequence number** – Begins with 0 at system startup.
   - **Device** – Disk Array, Logical Drive, Physical Drive by its ID number.
   - **Severity** – See the Table below.
   - **Timestamp** – Date and time the event happened.
   - **Description** – A description of the event in plain language.

3. Press the up and down arrow keys to scroll through the log.

WebPAM PROe Reports a Problem

WebPAM PROe reports these conditions in the header and all four tabs.

Header

The Header displays popup messages, per your configuration.

Popup messages

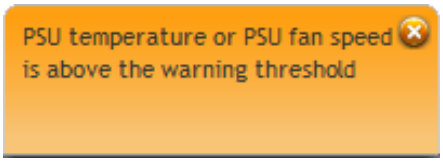

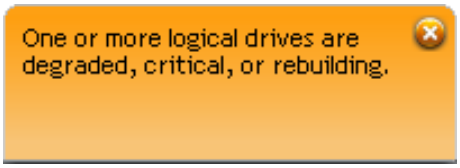

Dashboard Tab

System Status

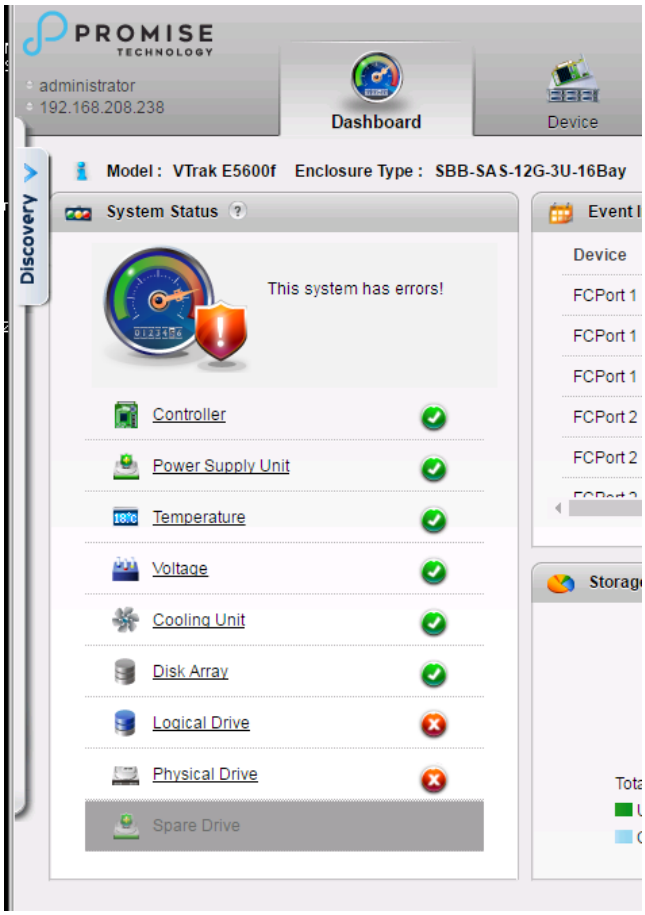

Red **X** and Yellow **!** icons identify components that need attention

Event Information

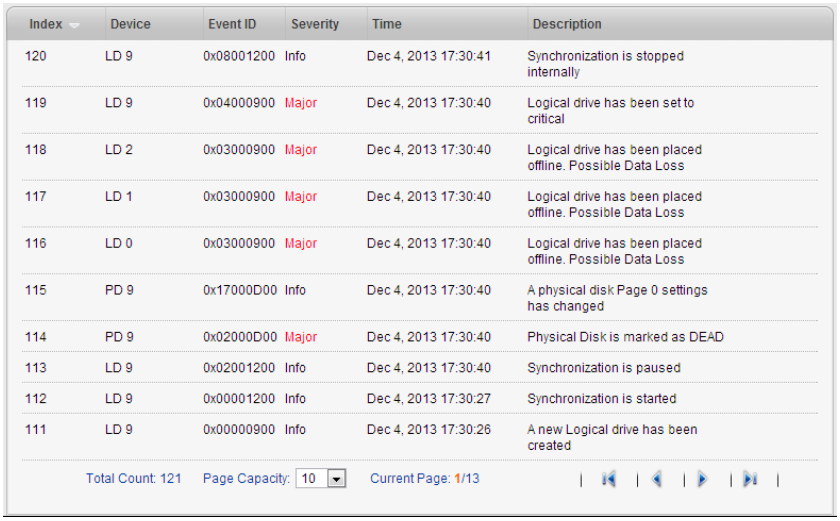

Index	Device	Event ID	Severity	Time	Description
120	LD 9	0x08001200	Info	Dec 4, 2013 17:30:41	Synchronization is stopped internally
119	LD 9	0x04000900	Major	Dec 4, 2013 17:30:40	Logical drive has been set to critical
118	LD 2	0x03000900	Major	Dec 4, 2013 17:30:40	Logical drive has been placed offline. Possible Data Loss
117	LD 1	0x03000900	Major	Dec 4, 2013 17:30:40	Logical drive has been placed offline. Possible Data Loss
116	LD 0	0x03000900	Major	Dec 4, 2013 17:30:40	Logical drive has been placed offline. Possible Data Loss
115	PD 9	0x17000D00	Info	Dec 4, 2013 17:30:40	A physical disk Page 0 settings has changed
114	PD 9	0x02000D00	Major	Dec 4, 2013 17:30:40	Physical Disk is marked as DEAD
113	LD 9	0x02001200	Info	Dec 4, 2013 17:30:40	Synchronization is paused
112	LD 9	0x00001200	Info	Dec 4, 2013 17:30:27	Synchronization is started
111	LD 9	0x00000900	Info	Dec 4, 2013 17:30:26	A new Logical drive has been created

Total Count: 121 Page Capacity: 10 Current Page: 1/13

Event Severity Levels

Level	Description
Fatal	Non-recoverable error or failure has occurred.
Critical	Action is needed now and the implications of the condition are serious.
Major	Action is needed now.
Minor	Action is needed but the condition is not a serious at this time.
Warning	User can decide whether or not action is required.
Information	Information only, no action is required.

Device Tab

Front View, showing the drive carrier icons

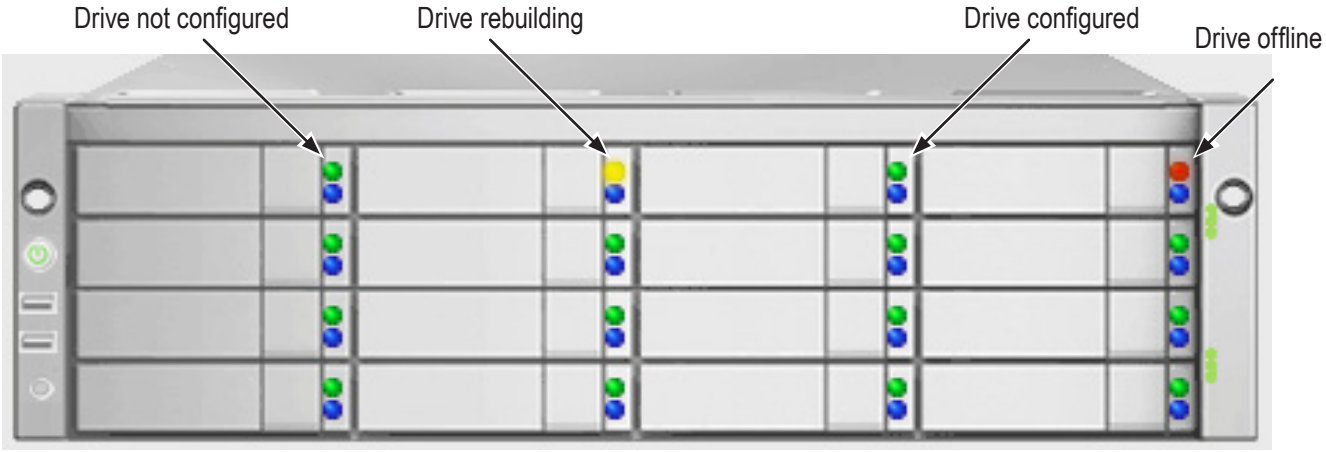

Rear View, with Show Internal Components option

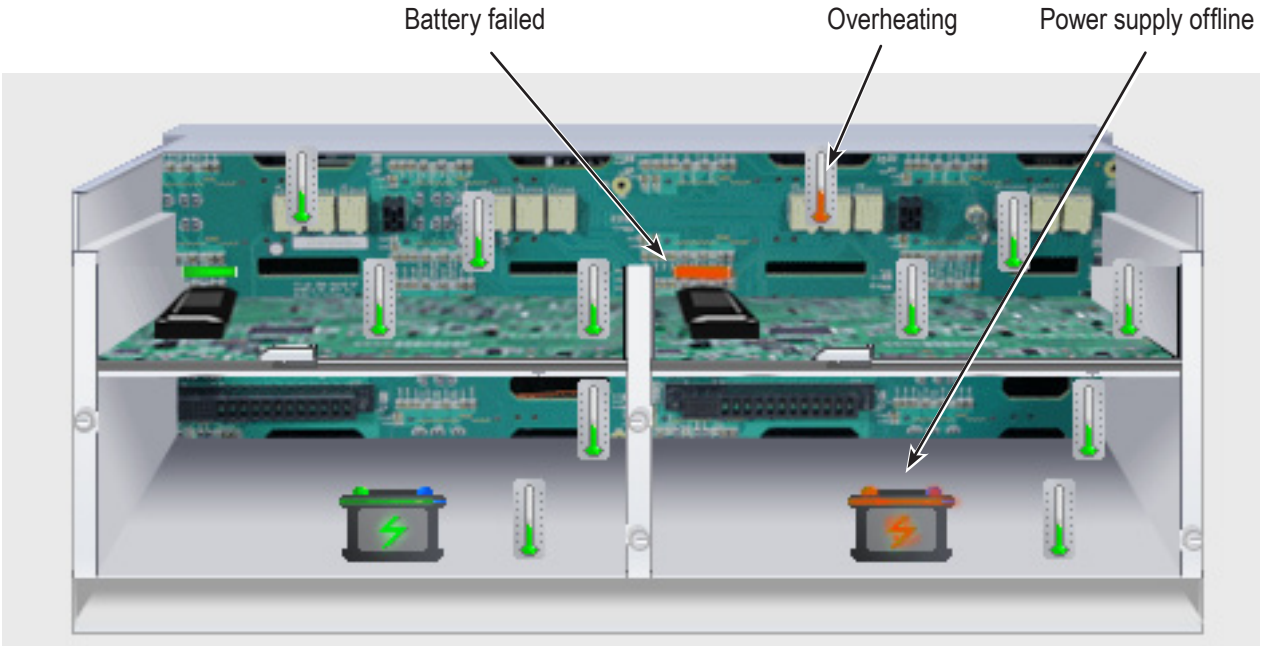

Physical Drive View, physical drive shown dead or offline and marked with a red X icon

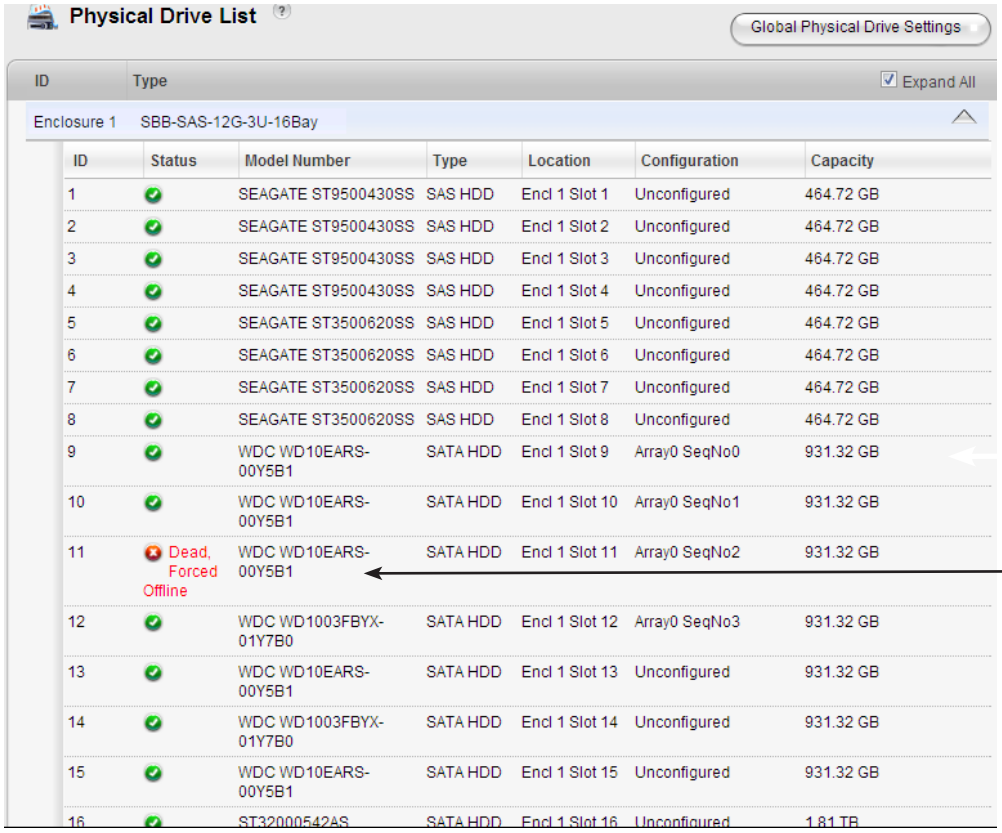

Storage Tab

Disk Arrays

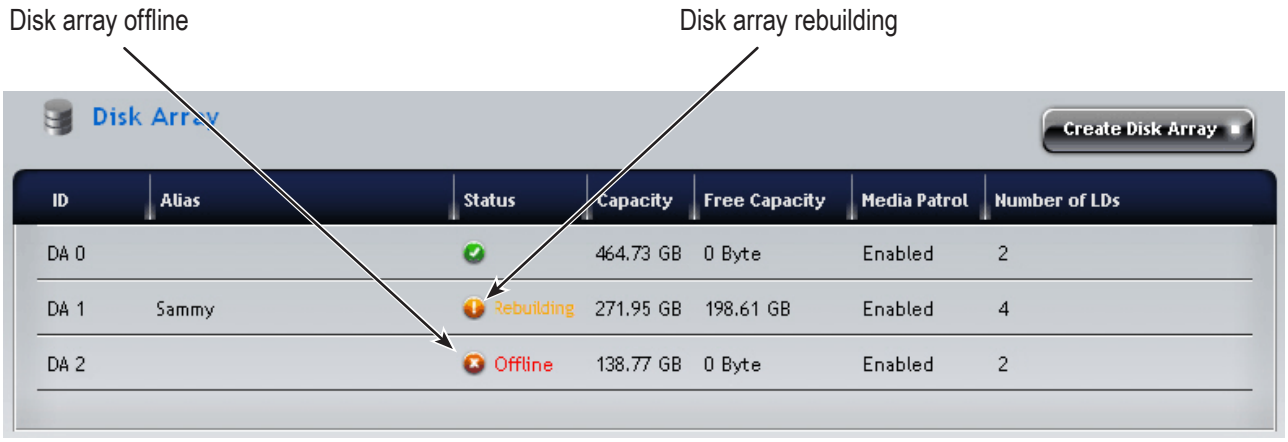

Logical Drives

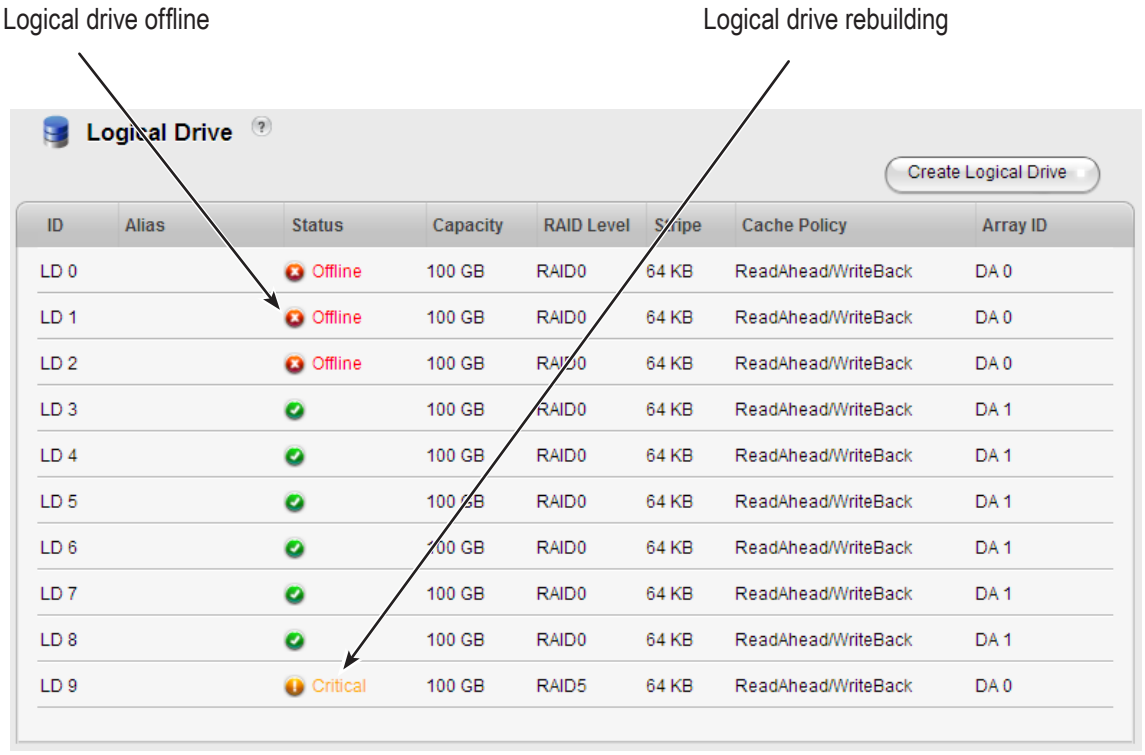

Administration Tab

Events icon

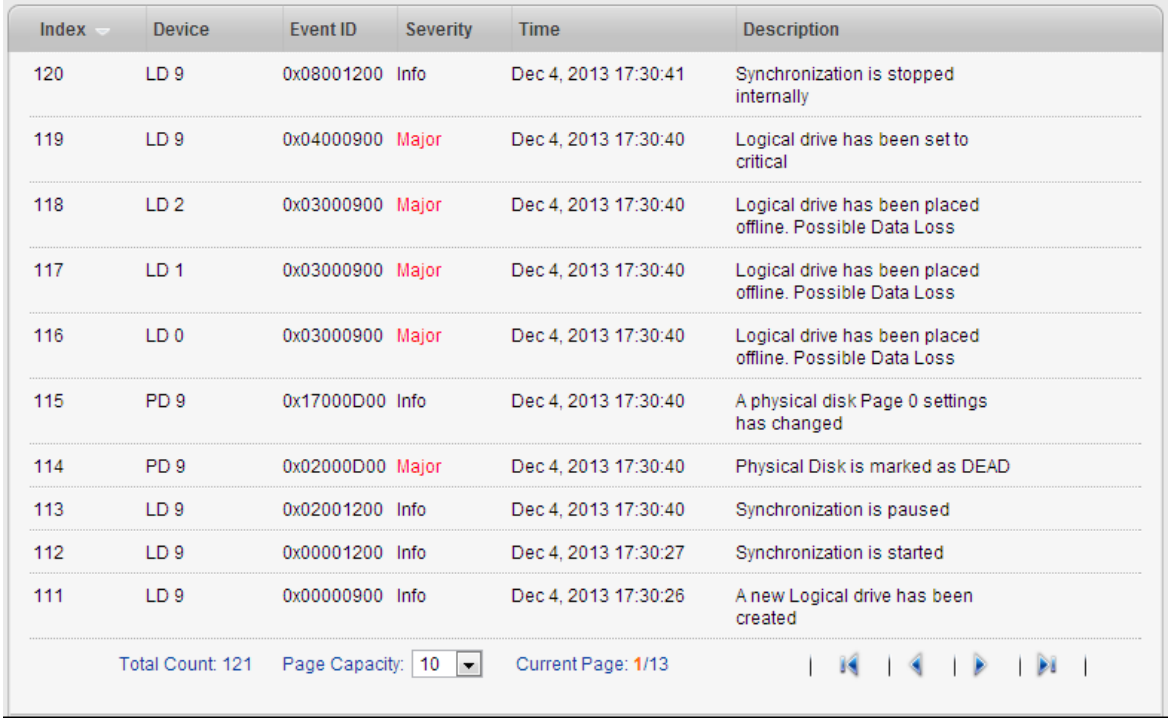

Index	Device	Event ID	Severity	Time	Description
120	LD 9	0x08001200	Info	Dec 4, 2013 17:30:41	Synchronization is stopped internally
119	LD 9	0x04000900	Major	Dec 4, 2013 17:30:40	Logical drive has been set to critical
118	LD 2	0x03000900	Major	Dec 4, 2013 17:30:40	Logical drive has been placed offline. Possible Data Loss
117	LD 1	0x03000900	Major	Dec 4, 2013 17:30:40	Logical drive has been placed offline. Possible Data Loss
116	LD 0	0x03000900	Major	Dec 4, 2013 17:30:40	Logical drive has been placed offline. Possible Data Loss
115	PD 9	0x17000D00	Info	Dec 4, 2013 17:30:40	A physical disk Page 0 settings has changed
114	PD 9	0x02000D00	Major	Dec 4, 2013 17:30:40	Physical Disk is marked as DEAD
113	LD 9	0x02001200	Info	Dec 4, 2013 17:30:40	Synchronization is paused
112	LD 9	0x00001200	Info	Dec 4, 2013 17:30:27	Synchronization is started
111	LD 9	0x00000900	Info	Dec 4, 2013 17:30:26	A new Logical drive has been created

Total Count: 121 Page Capacity: 10 Current Page: 1/13

Event Severity Levels

Level	Description
Fatal	Non-recoverable error or failure has occurred.
Critical	Action is needed now and the implications of the condition are serious.
Major	Action is needed now.
Minor	Action is needed but the condition is not a serious at this time.
Warning	User can decide whether or not action is required.
Information	Information only, no action is required.

Also see these troubleshooting topics:

- “Event Notification Response” on page 438
- “Enclosure Problems” on page 417

USB Support Reports a Problem

This procedure requires a USB flash device:

- Formatted to FAT 32
- At least 300 MB of free space

Caution

Verify that there is no firmware image file on the USB flash device.

If a firmware image file is present, the RAID controller might attempt a firmware update.

To collect a service report using the USB Support feature:

1. Insert the USB flash device into one of the USB ports on the front left of the VTrak E5600fD/VTrak E5800fD; or into a USB port on the controller on VTrak E5300fD/VTrak E5320fD.

   The controller status LED blinks green in half-second intervals.
2. Wait until the controller activity LED stops blinking green and displays steady green.
3. Remove the USB flash device.
4. Insert the USB flash device into a USB port on your PC.
5. On the USB flash device, open the OPA**X**_xxxxxx folder to obtain the report and log.

Enclosure Problems

Enclosure Problems include:

- “Diagnosing an Enclosure Problem” (below)
- “Overheating”
- “Power Supplies”
- “Batteries”

Diagnosing an Enclosure Problem

Check System Status on the Dashboard tab. If a yellow **!** or red **X** appears in the **System Status** box:

1. Click the name link of the component with the red **X** icon.

The Components List of the **Device** tab displays.

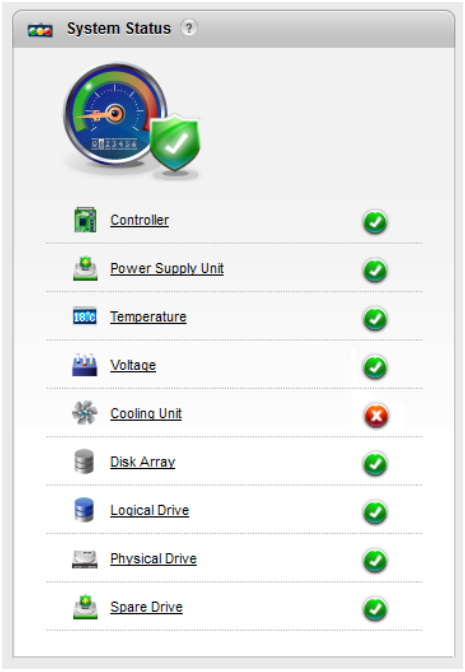

2. Mouse-over Enclosure with the red **X** icon and click the **View** button.

The components list expands and shows the power supplies (PSU) and Cooling Units of the VTrak E5000 enclosure.

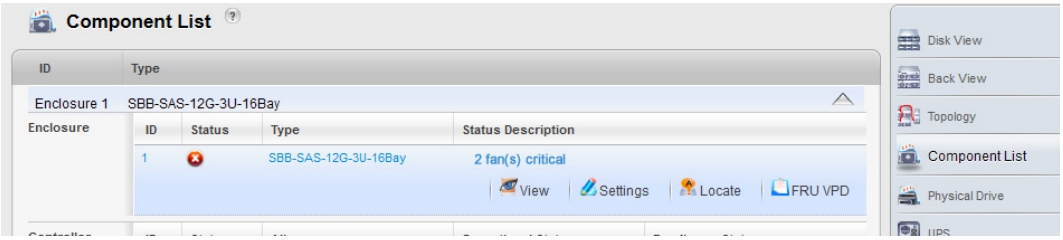

3. Click the **Back View** icon on the Device tab.
4. Click the picture of the enclosure.

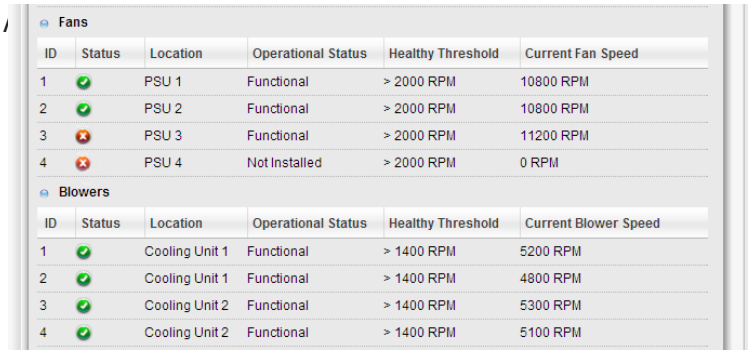

When a power supply fan fails, you must replace the power supply. See “Replacing a Power Supply” on page 424.

If the system reports a fan malfunction, contact Technical Support. See “Contacting Technical Support” on page 469 immediately to schedule replacement of the suspect power supply as soon as possible. Running the unit in this condition for more than three weeks may shorten subsystem life and void your warranty.

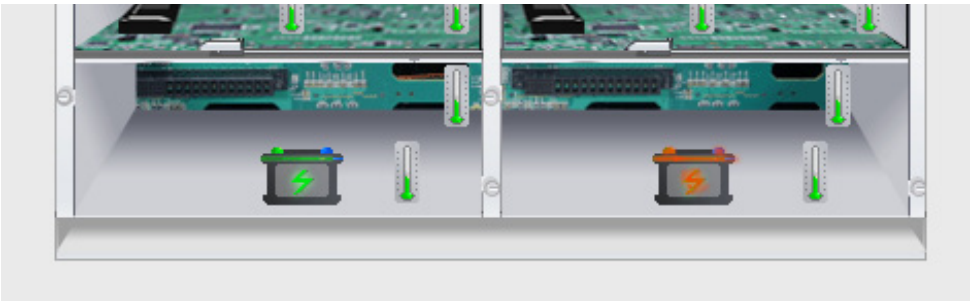

OVERHEATING

Overheating is a potentially serious condition because the excessively high temperatures can lead to physical drive failure and controller malfunction.

Overheating usually results from:

- Fan failure
- Inadequate air circulation around the enclosure

Fan Failure

In the VTrak E5600fD and VTrak E5800fD subsystems, the power supply fans are the Cooling Units for the enclosure.

When a power supply fan fails, you must replace the power supply. See “Replacing a Power Supply - VTrak E5600 / E5800” on page 358 for more information.

Inadequate Air Circulation

Air circulation around the VTrak enclosure might be a more complex problem. Use the thermometer icons to help you locate the specific hot spot. Check for these conditions:

- Accumulated dust or objects blocking the fans
- Less than a minimum of 13 cm (5 inches) space between the back of the enclosure and the wall or other object

Ambient temperature above 35°C (95°F) where the subsystem is operating

To cool down an enclosure:

- Correct any problems identified above.
- Power it down and let it sit for an hour or longer.
- See “Shutting Down the Subsystem” on page 94.

Power Supplies

VTrak subsystems are equipped with redundant power supplies. The advantage of dual power supplies is that, should one fail, the other continues to power the subsystem until the faulty one can be replaced. The subsystem is capable of operating on a single power supply.

The power supplies are hot-swappable, meaning you can leave the subsystem running when you replace the bad one. Be careful, however, to remove the faulty power supply and not the good one, or the subsystem comes to an immediate stop and your data is unavailable until the subsystem is powered and booted again.

See “Replacing a Power Supply” on page 323 for more information.

Batteries

The RAID controllers in the VTrak subsystem use a battery for backup power to protect data in the cache. Should a power failure occur, the battery enables backup of the cache to NAND flash (non-volatile memory). The battery recharges during normal subsystem operation.

In most cases, installing a replacement battery corrects a marginal or failed condition. The battery is located inside the RAID controller housing.

Also see “Battery Reconditioning” on page 141.

RAID Controller Problems

RAID controller problems include:

- “Maintenance Mode” (below)
- “Storage Tab”
- “Taking a RAID Controller out of Maintenance Mode”
- “Unsaved Data in the Controller Cache”

Controller problems occur when one of the controllers goes into maintenance mode.

Maintenance Mode

For VTrak with two RAID controllers, one of them enters maintenance mode in the event of:

- A difference of some kind between the two controllers (described below)
- An internal controller failure

When a controller enters maintenance mode, it goes offline and it displays N/A (not accessible) under Readiness Status.

You must find and correct the cause of the problem and then take the controller out of maintenance mode (see page 423).

Finding and Correcting the Cause of the Problem

External Checks

Make the following external checks to your VTrak subsystem. Be sure that:

- Both RAID controllers are present, fully inserted into their slots, and locked into place.
- The RAID controllers match, meaning both are the same model
- All SAS expansion cables from the RAID controllers to external JBOD units in good condition and are securely connected.

Important

A disconnected SAS expansion cable causes the two RAID con-trollers to see a different set of configured drives. This condition is the most common cause of a controller entering maintenance mode.

Internal Checks

If all external checks are OK, take the following actions:

1. Shut down the VTrak.

See “Shutting Down the Subsystem” on page 94.

2. Remove one of the RAID controllers.

See “Replacing a RAID Controller: Dual Controllers” on page 429.

3. Restart the VTrak.
4. After the VTrak is fully booted, view the controller information.

“Viewing Controller Information” on page 96.

5. Observe and record the following information about the first controller:
   - SDRAM memory size
   - Hardware version
   - Firmware version
6. Shut down the VTrak.
7. Remove the first controller and install the second controller.
8. Repeat steps 3 through 6. Then compare your records.
9. Correct any differences between the two controllers. See “Updating Firmware on a RAID Subsystem” on page 102.

Taking a RAID Controller out of Maintenance Mode

If you shut down the VTrak subsystem in the process of correcting the mainten-ance mode problem, the affected RAID controller boots into normal mode when the VTrak restarts. No further action is required.

If you corrected the problem without shutting down the VTrak subsystem, choose one of the following methods to take the controller out of maintenance mode:

- Restart the VTrak subsystem. See "Restarting the Subsystem" on page 93.
- Establish a serial connection, then use the CLI (see below) or
- Establish a Telnet connection, then use the CLI (see page 424)

Serial Connection

To clear maintenance mode using a serial connection:

1. Change your terminal emulation program settings to match the following specifications:
   - Bits per second: 115200
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: none
2. Start your PC's terminal VT100 or ANSI emulation program.
3. Press **Enter** once to launch the CLI.

The login screen appears.

The following steps show the default Administrator user name and pass-word. Use your own user name and password if you have changed these.

4. At the Login prompt, type **administrator** and press **Enter**.
5. At the Password prompt, type **password** and press **Enter**.

The CLI screen appears.

The prompt should display **MAINTENANCE MODE@cli>**.

If the prompt dis-plays your login name, such as **administrator@cli>**, log into the other controller.

6. At the **MAINTENANCE MODE@cli>** prompt, type maintenance -a exit and press **Enter**.

The controller reboots. The login screen again appears.

7. Close the Serial connection.

Telnet Connection

This procedure requires you to know the IP address of the controller.

To clear maintenance mode using a Telnet connection:

1. Go to the command line prompt (Windows) or click the terminal icon (Linux), then run:

telnet 192.168.1.56 2300

The IP address above is only an example. 2300 is the default Telnet port for VTrak.

The login screen appears.

The following steps show the default Administrator user name and pass-word. Use your own user name and password if you have changed these.

2. At the Login prompt, type **administrator** and press **Enter**.
3. At the Password prompt, type **password** and press **Enter**.

The CLI screen appears.

The prompt should display **MAINTENANCE MODE@cli>**. If the prompt dis-plays your login name, such as **administrator@cli>**, log into the other controller.

4. At the **MAINTENANCE MODE@cli>** prompt, type maintenance -a exit and press **Enter**.

The controller reboots. The Telnet session ends.

Unsaved Data in the Controller Cache

The dirty cache LED (marked with the icon) informs you that there is data in the cache that has not been saved to non-volatile memory. Such data is sometimes called “dirty,” not to suggest it is corrupted in some way but because it has not been saved to a physical drive.

Dirty Cache LED

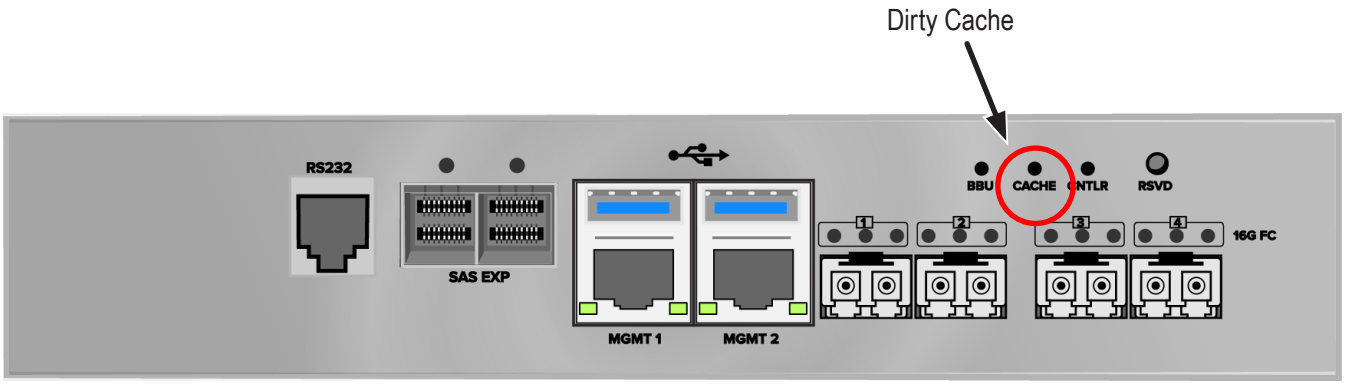

Caution

If there is unsaved data in the controller’s cache, the dirty cache LED shines amber. During this time, do NOT power down the VTrak. Wait until the LED goes dark.

Physical Drive Problems

Physical drives are the foundation of data storage. A physical drive problem can affect your entire RAID system.

When a yellow **!** icon or a red **X** icon appears beside a physical drive, check the drive's operational status:

1. Click the **Device** tab.
2. Click the **Physical Drive** icon.
3. Click the physical drive you want, then click the View button.

Look under Operational Status for the condition of the physical drive.

- Offline – Check the drive for:
- PFA Condition – Caused by a bad block or sector. See Note 1 below.
- Stale Condition – Caused by obsolete array information on the physical drive. See Note 2 below.
- Not Usable – This condition occurs when you have:
- Two controllers in your RAID subsystem and a SATA drive without a SAS-to-SATA adapter. See Note 3 below.
- A missing or defective SAS cable between the RAID subsystem and a JBOD expansion unit.
- Drive Failed or Dead – The physical drive cannot be repaired. You must replace the failed drive. See Note 4 below.

Note 1: Clear the error condition. Then the physical drive is available. See "Clearing a Stale or a PFA Condition" on page 177.

Note 2: Identify the disk array to which the physical drive belongs. Then delete the disk array. If the error condition remains on the physical drive, clear the error condition.

Note 3: Obtain SAS-to-SATA adapters though PROMISE Technology, at

http://www.promise.com. See "Installing Physical Drives" on page 22 for installation instructions.

Note 4: You can set the number of bad blocks tolerated before the controller marks a physical drive as Dead. See "Managing Background Activities" on page 129, "Making PDM Settings" on page 139. See also "Running Media Patrol" on page 185 and "Disk Array Degraded / Logical Drive Critical" on page 427.

Disk Array and Logical Drive Problems

Disk array and logical drive problems include:

- "Disk Array Degraded / Logical Drive Critical" (below)
- "Disk Array Offline / Logical Drive Offline"
- "Repairing an Offline Disk Array or Logical Drive"
- "Rebuilding a Disk Array"
- "Incomplete Array"

Disk array problems typically result from a physical drive failure. The most common problem is a degraded disk array. The RAID controller can rebuild a degraded disk array. See "Rebuilding a Disk Array" on page 430.

Disk Array Degraded / Logical Drive Critical

Disk arrays are made up of physical drives. Logical drives are created on the disk array.

When one of the physical drives in a disk array fails:

- The operational status of the disk array becomes Critical.
- The operational status of the logical drives becomes Critical or Degraded.
- The operational status of the physical drive becomes Dead or Offline.

WebPAM PROe reports these conditions in the following places:

- **Dashboard tab**

A yellow **!** icon beside the disk arrays, logical drives, and physical drives under **System Status**.

Major event for the logical drive under Event **Information**.

Warning event for the physical drive under Event **Information**.

- **Device tab**

Front View – Physical drives are shown Dead or Offline and marked with a red **X** icon, or Missing.

Physical Drive View – Physical drives are shown Dead or Offline and marked with a red **X** icon, or Missing.

- **Storage tab**

Disk Array and Logical Drive are marked Critical with a yellow **!** icon.

RAID 6 and 60 logical drives are marked:

- Degraded with a yellow **!** icon when ONE physical drive is offline.
- Critical with a yellow **!** icon when TWO physical drives are offline.

RAID 0 logical drives show Offline status and a red **X** icon.

If there is no spare drive or unconfigured drive in the RAID system, you must provide the replacement drive. See “Installing Physical Drives” on page 22.

- **Administration tab**

Depending on your settings and availability of a replacement drive, your system automatically rebuilds the degraded disk array. See “Rebuilding a Disk Array” on page 430.

The system sends an Email message about the incident to subscribing users, depending on user settings. See “Setting User Event Subscriptions” on page 126.

Disk Array Offline / Logical Drive Offline

Disk arrays are made up of physical drives. Logical drives are created on the disk array. When a disk array and its logical drives go Offline, the data stored in the logical drives is no longer accessible.

RAID 0 logical drives go Offline when ONE physical drive is removed or fails.

RAID 5, 10, and 50 logical drives go Offline when TWO physical drives are removed or fail.

RAID 6 and 60 logical drives go Offline when THREE physical drives are re-moved or fail.

WebPAM PROe reports these conditions in the following places:

- **Dashboard tab**

A red **X** icon appears beside the disk arrays, logical drives, and physical drives under **System Status**.

Major event for the logical drive under Event **Information**

Warning event for the physical drive under Event **Information**.

- **Device tab**

On Front View and Physical Drive View, physical drives are shown Dead, Offline, or Missing.

- **Storage tab**

Disk array and logical drives are marked with a red **X** icon.

- **Administration tab**

Under **Background Activities**, no Rebuild takes place. See “Repairing an Offline Disk Array or Logical Drive”, below.

The system sends an Email message about the incident to subscribing users, depending on user settings. See “Setting User Event Subscriptions” on page 126.

Repairing an Offline Disk Array or Logical Drive

RAID 1, 5, 6, 10, 50, and 60 Logical Drives

If a fault-tolerant logical drive, RAID 5, 6, 10, 50, and 60, goes Offline, it may be possible to recover your data.

Warning

Take no further corrective action until you have consulted with Technical Support!

RAID 0 Logical Drives

If a logical drive based on a non-fault-tolerant disk array, RAID 0, goes offline, all of the data on the logical drive is lost.

To recreate your logical drive:

1. Identify the failed physical drive. See “Locating a Physical Drive” on page 175.
2. Replace the failed drive. See “Installing Physical Drives” on page 22.
3. If the disk array had more than one physical drive, delete the disk array and re-create it.

   See “Deleting a Disk Array” on page 183 and “Creating a Disk Array Manually” on page 181.
4. Restore the data from your backup source.

Rebuilding a Disk Array

When you rebuild a disk array, you are actually rebuilding the data on one physical drive.

- When a physical drive in a disk array fails and a spare drive of adequate capacity is available, the disk array begins to rebuild automatically using the spare drive.
- If there is no spare drive of adequate capacity, but the Auto Rebuild function is ENABLED, the disk array begins to rebuild automatically as soon as you remove the failed physical drive and install an unconfigured physical drive in the same slot. See “Making Rebuild Settings” on page 137.
- If there is no spare drive of adequate capacity and the Auto Rebuild function is *DISABLED*, you must replace the failed drive with an unconfigured physical drive, then perform a Manual Rebuild. See “Rebuilding a Disk Array” on page 187.

Important

If your replacement disk drive was formerly part of a different disk array or logical drive, you must clear the configuration data on the replacement drive before you use it. See “Clearing a Stale or a PFA Condition” on page 177.

Incomplete Array

A more serious, but far less common problem is an Incomplete Array. An in-complete array results from a physical drive that fails or becomes missing during:

- RAID level migration
- Disk array transport

Migration

Normally, if a physical drive or the controller fails during migration, the disk array goes critical, and you can rebuild it.

Transport

Transport is the action of moving the physical drives of a disk array:

- To different slots in the same enclosure
- From one enclosure to another

If a physical drive fails during a transport, or you do not move all of the physical drives to their new locations, WebPAM PROe displays an incomplete array. When WebPAM PROe discovers an incomplete array, it displays a dialog box asking you to:

- Click the OK button to accept the incomplete array.
- Click the Cancel button to reject the incomplete array.

Accepting an Incomplete Array

Before you accept the incomplete array, be sure all of the physical drives are present and that their drive carriers are properly installed into the enclosure. See “Installing Physical Drives” on page 22.

If you choose to accept the incomplete array:

1. Click **OK** in the incomplete array dialog box.
2. Check the operational status of the logical drives in the array.
    - If the logical drives are Critical, proceed with a rebuild.
    - If the logical drives are Offline, contact Technical Support.
      See “Contacting Technical Support” on page 469.
3. Restore your data from a backup source.

If you choose NOT to accept the incomplete array:

1. Click Cancel in the incomplete array dialog box.
2. Do one of the following:
    - Delete the array. This action deletes all logical drives on the array.
    - Replace the missing physical drive.

Connection Problems

Connection problems include:

- "Serial Connections" (below)
- "Network Connections"
- "Fibre Channel Connections"
- "SAS Connections"
- "Browser Does Not Connect to WebPAM PROe"

Connection problems cause a majority of failures in almost any electrical system. While the installation of the cables and components was correct, they don't function properly, or at all, because:

- A connector is dirty or corroded
- A connector is loose or damaged
- A cable looks OK outside but has an open circuit inside
- The wrong cable was used

The VTrak E5000 ships with a full set of new cables, as required for each specific model. Be sure to use these components because: 1.) They are the proper ones for your RAID subsystem, 2.) They are in brand-new condition, and 3.) You paid for them with the purchase of your subsystem.

Serial Connections

VTrak uses a serial connection for the command line interface (CLI). After you set the IP address, you can access the CLI through a network connection, also. Normally, users prefer WebPAM PROe because of its graphic user interface. But the CLI can do the same jobs. And it works when your network connection is down.

For VTrak, you must use the CLI to set the Management Port IP address in order for WebPAM PROe to connect with it. See "Making a Serial Connection" on page 226 for more information on making the connection. This issue is discussed further under Network Connections, below.

The CLI controls and manages but does not move data. The CLI communicates through a RJ11-to-DB9 serial data cable, supplied with the VTrak. You may choose not use the CLI often and want to disconnect and store the

cable. Consider leaving it connected, so you know where it is the next time you need it.

Network Connections

The VTrak E5000 Series RAID controller has two Ethernet (RJ45) management ports. These are 1 Gbps Ethernet connections designed to connect to your network and be used for management of the subsystem. The VTrak becomes a node on your network like any other PC, server or other component with an IP address.

VTrak E5000 models ship from the factory virtual IP addresses of 10.0.0.1, this applies to both ports on the controller, and all four ports for a duel controller subsystem. You can change IP address settings using the CLI, or with WebPAM PROe. See “Virtual Management Port Settings” on page 53 for instructions on using the CLI to change IP settings for the network ports.

VTrak E5000 Series controller Network (RJ-45) and Serial (RJ-11) ports

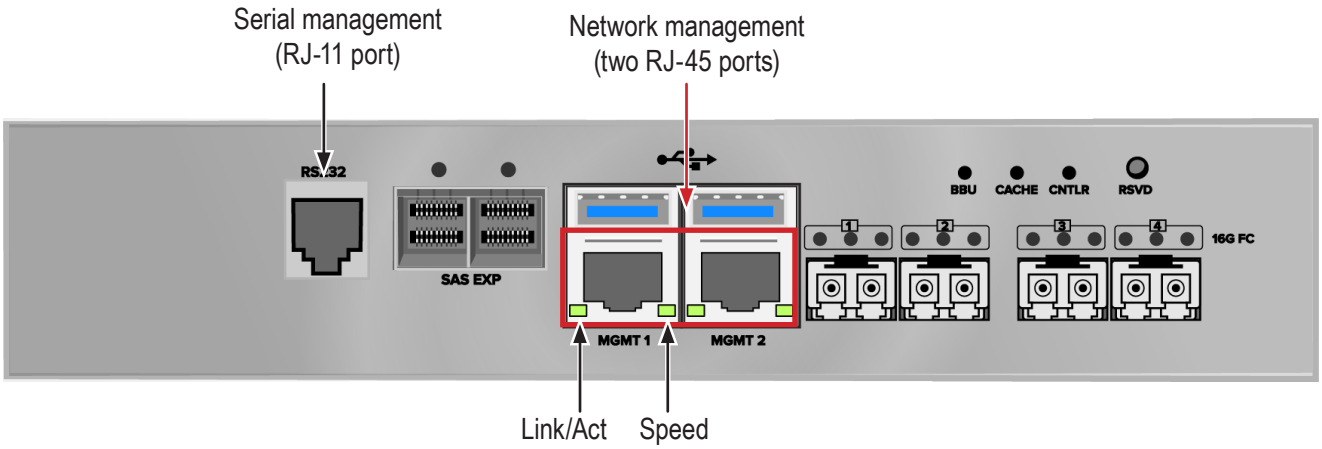

Management Port LEDs

State	Link/Activity	Speed
No link	Dark	Dark
10 Mbps	Steady green	Dark
100 Mbps	Steady green	Flashing orange
1 Gbps	Steady green	Flashing green

Note that the virtual and maintenance ports can accept IP address assignments from a DHCP server.

If you manually assigned an IP address to the VTrak E5000 but there is a DHCP server on your network, there is a chance that the server might assign the IP address to another node. You might see a warning to this effect on your PC's monitor. If this happens, WebPAM PROe may not be able to connect. See your network administrator to work out a suitable arrangement.

Fibre Channel Connections

When there is a connection failure, use WebPAM PROe to verify that VTrak sees the initiators. See "Redundancy Check" on page 135.

If VTrak sees some initiators but not the one you want, the problem is most likely elsewhere in the loop or fabric. If VTrak does not see any initiators:

- Check all of the Fibre Channel connections
- Verify that all nodes are properly connected and powered
- Verify that the fabric router or switch is properly connected powered

For more information, see "Managing Fibre Channel Connections" on page 217.

SAS Connections

Faulty SAS expansion connections are suspected when the link port counter reports a large number of bad link errors.

Link errors can be caused by:

- Debris blocking the SAS cable connector
- A faulty SAS cable
- A faulty controller or I/O module SAS connector

Blocked Cable Connectors

To check for debris blocking the SAS cable connector:

1. Power down the RAID subsystem and JBOD units.
2. Remove the SAS cable and check all SAS connectors for debris.
3. Clean the connectors as required and reconnect the SAS cable.
4. Power up the subsystems and monitor the link port counter for changes in the rate of link error accumulation.

Faulty Cable

To check for a faulty SAS cable:

1. Power down the RAID subsystem and JBOD units.
2. Replace the SAS cable with a new one.
3. Power up the subsystems and monitor the link port counter for changes in the rate of link error accumulation.

Faulty Controller or I/O Module Connector

To check for a bad controller or I/O module SAS connector:

1. With the subsystems online and I/Os running, access the CLI via serial or Telnet connection.
2. At the command prompt, type the following command and press **Enter**.

administrator@cli> sasdiag -a errorlog -l expander -e 1 -i 1

3. At the command prompt, type the following command and press **Enter**.

administrator@cli> sasdiag -a errorlog -l c2cport

By interpreting the two error logs, you can verify which controller or I/O module SAS port is accumulating link errors.

Browser Does Not Connect to WebPAM PROe

If you successfully setup and connected to WebPAM PROe, then suddenly you can no longer connect, it might be the result of the following three conditions:

- DHCP is enabled on your VTrak's virtual management port
- The DHCP server does not have a dedicated IP address for the VTrak
- The VTrak restarted and your DHCP server assigned a new IP address

You must obtain the new IP Address for the virtual management port in order to direct your browser to the VTrak and start WebPAM PROe.

To access the new IP address:

1. Start your PC's terminal VT100 or ANSI emulation program.
2. Press **Enter** once to launch the CLI.
3. At the Login prompt, type **administrator** and press **Enter**.
4. At the Password prompt, type **password** and press **Enter**.
5. Type **net** and press **Enter**.

```
administrator@cli> net

================================================================================

CId         Port Type     IP              Mask           Gateway       Link

================================================================================

Virtual     Mgmt      192.168.10.85  255.255.255.0    192.168.10.1    Up
```

The new virtual management port IP address and other network settings display.

6. Enter the new IP address into your browser to log into WebPAM PROe.

For more information, see "Setting-up the Serial Connection" on page 48, also, "Virtual Management Port Settings" on page 53, and "Logging into WebPAM PROe" on page 62.

Power Cycling the Subsystem

To power cycle a RAID subsystem means to:

- Shut down
- Turn off the power
- Turn on the power
- Restart

Power cycling is sometimes required as a remedial action but only when prompted by a message from software or when directed by Technical Support.

To power cycle the RAID subsystem:

1. Shut down the subsystem.

   When the controllers shut down, your network connection is lost.

2. Manually turn off the system power by pressing the power button of the RAID unit for over five seconds.
3. Wait at least 10 seconds.
4. Manually press the power button of the JBOD units.
5. Manually press the power button of the RAID subsystem.
6. Wait no less than two minutes.
7. Open your browser and log into WebPAM PROe.

   If you cannot log in immediately, wait 30 seconds and try again.

Important

If your RAID subsystem manages JBOD expansion units, always power on the JBOD expansion units first. Then power on the RAID subsystem.

Event Notification Response

When you choose Event Notification, WebPAM PROe sends popup and/or email messages regarding its status. The messages you see depend on your notification selection and what is currently happening in the VTrak. See "Setting User Event Subscriptions" on page 126.

The table below cites:

- Reported Events – Events that require you to take action
- Corrective Actions – The action you should take in response to the event

A list of event categories is shown below.

Reported Event	**Corrective Action**
Battery	
Battery is inserted	No action is required.
Battery charging has failed	Replace the battery.
Battery reconditioning has started	No action is required.
Battery reconditioning has been terminated	Replace the battery.
The write policy of writeback logical drive switched from writeback to writethru	Check the event log to see whether battery is re-conditioning.
The write policy of writeback logical drive switched from writethru to writeback	No action is required.
Battery is charging in high temperature	Monitor the condition. Contact Tech Support if the problem persists.
Battery cannot function with the enclosure or with the attached battery board	Wrong battery installed. Contact Tech Support for assistance.
Logical drive writeback cache maybe enabled without battery support	No action required.
Battery is fully charged	
Battery is not present	Install a battery or verify that the battery is properly connected.
Battery is not accessible	Connect the battery properly or replace the battery.
BBU	
BBU flushing has started	No action is required.
BBU flushing has ended	
BBU flushing has failed	Contact Tech Support if the condition persists.

Reported Event	**Corrective Action**
Blade Server	
Blade Server Inserted	No action is required.
Blade Server Removed	
Cache	
Not available	Contact Tech Support.
Controller	
The controller parameter(s) are changed by user	No action is required.
The controller is reset by Watch Dog timer	Result of a firmware update. If the condition persists, replace the controller.
The controller has new crash information	Contact Tech Support.
The controller's heart beat has started	No action is required.
The controller's heart beat has stopped	
The partner controller's heart beat has started	
The partner controller's heart beat has stopped	
The partner controller's heart beat has skipped	
The controller's main scheduler has frozen	Contact Tech Support if the condition persists.
Controller has entered maintenance mode since configured physical disk seen by partner controller is not seen here	Verify that all SATA drives have an SAS-to-SATA adapter installed.
Controller has entered maintenance mode due to mismatch of physical disks types	Check and correct SAS cabling and connections as needed.
Controller has entered maintenance mode due to mismatch of physical disk WWN	Update to the latest firmware. If the condition persists, replace the controller.
Controller has entered maintenance mode due to mismatch of SATA Disks	Check and correct data cabling and connections as needed.
Controller has entered maintenance mode due to mismatch of Disk IDs	
Controller has entered maintenance mode since no physical disks are seen as seen by Partner controller	

Reported Event	**Corrective Action**
Controller	
Controller is started	No action is required.
Controller is set to Active Mode	
Controller is set to Standby Mode	
Controller Failed Over as partner is removed	Verify that the partner controller is properly installed and all cables are connected.
Controller Failed Over as heart beat stopped	
Controller Firmware mismatch with that of the partner controller	Auto Firmware synchronization upgrades or downgrades the firmware.
Controller set to Maintenance Mode because of hardware mismatch with partner (controller)	Compare controller types and amount of memory installed. Correct or update as needed.
Controller set to Maintenance Mode because of firmware mismatch with partner controller	Update this controller to the same firmware version as the partner controller.
Controller set to Maintenance Mode because Firmware is flashing in the partner controller	Exit out of Maintenance mode after firmware flashing is complete.
Controller set to Maintenance Mode because of flash image version mismatch with partner (controller)	Update this controller to the same flash image version as the partner controller.
Controller has been set to Maintenance mode because there is a mismatch in the Controller Model or Hardware version with that of the partner controller	Replace this controller with the same Model and Hardware version as the partner controller.

Reported Event	**Corrective Action**
Controller	
Controller has been set to Maintenance mode because there is a mismatch in the memory size with that of the partner controller	Replace this controller's memory with the same memory size as the partner controller
Partner Controller has entered maintenance mode to protect user data since one of the configured physical drives was disconnected in the partner controller	Check and correct cable connections to external JBOD enclosures. Rebuild any critical logical drives. Back up array data. Replace the physical drive. Bring controller out of maintenance mode.
Controller was placed on reset during Fail Over processing	No action is required.
Partner Controller was placed on reset during Fail Over processing	
Controller was reset as it was not able to join the running partner controller	Verify that the controller is running. If the condition persists, replace the controller.
The controller has reset because it encountered a firmware problem	If resets happen frequently, update to new firmware or replace the controller.
Controller temperature is above the warning threshold	Check airflow around the VTrak. Check blowers and fans.
The controller temperature is above controller critical threshold	No action is required.
Controller temperature is within the normal range	

Reported Event	**Corrective Action**
CRC	
CRC error is detected while receiving CMD information unit	If this message appears repeatedly, contact Tech Support.
CRC error is detected during Data Out phase	
Disk Array	
New disk array has been created	No action is required.
Disk array has been deleted	
Disk array has been added	
Disk array has been removed	
Disk array settings have been changed	
Disk array is transport ready	Remove physical drives in disk array and insert them into a different subsystem. To cancel Transport Ready Status, remove and reinsert the drives in their original slots.

Reported Event	**Corrective Action**
Drive Interface	
Drive-interface controller is found	No action is required.
Drive-interface controller is NOT found	Restart the VTrak. If this message appears repeatedly, contact Tech Support.
Drive-interface diagnostics has passed	No action is required.
Drive-interface diagnostics has failed	Restart the VTrak. If this message appears repeatedly, contact Tech Support.
Drive-interface controller has generated a general parity error	If this message appears repeatedly, contact Tech Support.
Drive-interface controller has generated a data parity error	
Enclosure	
Enclosure ***temperature is above the threshold***	Check blowers and fans.
Enclosure ***temperature is above the warning threshold***	Check airflow around the VTrak. Check blowers and fans.
Enclosure ***temperature is above the critical threshold***	
Enclosure ***temperature is within the normal range***	No action is required.
Shut down PSUs due to enclosure or controller temperature over threshold	Shut down the VTrak
Event Log	
Event ***logging is enabled***	No action is required.
Event ***logging is disabled***	
Event ***log buffer is cleared in RAM***	
Event ***log buffer is cleared in NVRAM***	
Event ***log buffer is cleared in MDD***	

Reported Event	**Corrective Action**
Fibre Channel	
Fibre Channel controller has detected bus reset	If this message appears repeatedly, contact Tech Support.
Fibre Channel controller has received a "LUN reset" command.	No action is required.
Fibre Channel controller has encountered a fatal error	Restart the VTrak. If this message appears repeatedly, contact Tech Support.
Fibre Channel link is up	No action is required.
Fibre Channel link is down	
Fibre Channel controller settings have changed	
Firmware Update	
Firmware update is started	No action is required.
Firmware update is complete	
Firmware update is fail	Try the update again. If this message repeats, contact Tech Support.
Back-end expander firmware upgrade is started	No action is required.
Back-end expander firmware upgrade is completed	
Back-end expander firmware upgrade failed	Try the update again. If this message repeats, contact Tech Support.
Front-end expander firmware upgrade is started	No action is required.
Front-end expander firmware upgrade is completed	
Front-end expander firmware upgrade failed	Try the update again. If this message repeats, contact Tech Support.

Reported Event	**Corrective Action**
Host Interface	
Host interface controller has detected bus reset	If this message appears repeatedly, contact Tech Support.
Host interface controller has encountered an unrecoverable error	Restart the VTrak. If this message appears repeatedly, contact Tech Support.
Host interface controller has received an "abort task" command.	No action is required.
Host interface controller has received an "abort task set" command.	
Host interface controller has received a "clear ACA" command.	If this message appears repeatedly, contact Tech Support.
Host interface controller has received a "clear task set" command.	No action is required.
Host interface controller has received a "LUN reset" command.	

Reported Event	**Corrective Action**
Host Interface	
Host interface controller is informed that the initiator has detected an error	If this message appears repeatedly, contact Tech Support.
Host interface controller has received illegal secondary identification	
Host interface controller has received a message parity error	
Host interface controller has received a bus reboot	
Host interface link is up	No action is required.
Host interface link is down	Check connections.
Host interface controller has encountered an unknown error	If this message appears repeatedly, contact Tech Support.
Host interface controller has encountered a system error	
Host interface controller has encountered a fatal error	Restart the VTrak. If this message appears repeatedly, contact Tech Support.
Host interface controller settings have changed	No action is required.
Host interface controller has received a 'WARM reset' command	If this message appears repeatedly, contact Tech Support.
Host interface controller has received a "COLD reset" command	
Host Interface controller, MU handshake failed	
Host Interface controller, HMU has stopped	
Host Interface controller, FMU has unloaded	

Reported Event	**Corrective Action**
Initiator	
Initiator sent message for detecting an error	If this message appears repeatedly, contact Tech Support.
JBOD	
JBOD system connected	No action is required.
JBOD system either is removed or malfunctioned	Check Expander firmware and SAS connections.
Logical Drive	
Logical drive initialization has started	No action is required.
Logical drive Initialization is in progress	
Logical drive initialization has completed	
Logical drive initialization has paused	Resume the initialization when ready.
Logical drive initialization has resumed	No action is required.
Logical drive initialization has stopped	If this action was not intentional, check the logical drive's status.
Logical drive initialization marks the logical drive offline	Replace the failed physical drive. Delete and recreate the logical drive.
Logical drive initialization is aborted due to an internal error.	Reduce system load on the VTrak.
Logical drive initialization is queued	No action is required.
Quick logical drive initialization has started	
Quick logical drive initialization has completed	
Quick logical drive initialization has paused	Resume the initialization when ready.
Quick logical drive initialization has resumed	No action is required.
Quick logical drive initialization has stopped	If this action was not intentional, check the logical drive's status.
Quick logical drive initialization marks the logical drive offline	Replace the failed physical drive. Delete and recreate the logical drive.
Quick logical drive Initialization is aborted due to an internal error	Reduce system load on the VTrak.

Reported Event	**Corrective Action**
Logical Drive	
Quick logical drive initialization is queued	No action is required.
A new logical drive has been created	
Logical drive has been deleted	
Logical drive has been placed online	
Logical drive has been placed online. Possible data loss	Check the state of the physical drives, replace any bad drives. Rebuild logical drive.
Logical drive has been set to critical.	
Logical drive has been set to degrade	
Rebuild marks the logical drive synchronized upon rebuild completion	No action is required.
Logical drive settings has been changed through a user command	
One of the error tab***les of a logical drive has been cleared by the user***	
Logical drive axle has been placed online	

Reported Event	**Corrective Action**
Media Patrol	
Media patrol is started	No action is required.
Media patrol is in progress	
Media patrol is completed	
Media patrol is paused	Resume Media Patrol when ready.
Media patrol is resumed	No action is required.
Media patrol is stopped	If this action was not intentional, check the logical drive's status.
Media patrol is aborted due to an internal error.	Reduce system load on the VTrak.
Media patrol is queued	No action is required.
Media patrol is stopped internally	
Online Capacity Expansion	
Online capacity expansion has started	No action is required.
Online capacity expansion has completed	
Online capacity expansion has paused	Resume OCE when ready.
Online capacity expansion has resumed	No action is required.
Online capacity expansion has stopped	If this action was not intentional, check the logical drive's status.
Online capacity expansion has encountered a physical disk error	Check the physical drive check table after OCE is finished.
Online capacity expansion is aborted due to an internal error.	Reduce system load on the VTrak.
Online capacity expansion is queued	No action is required.
Parity	
Parity error is detected during Data Out phase	If this message appears repeatedly, contact Tech Support.

Reported Event	**Corrective Action**
PDM	
PDM is started	No action is required.
PDM is in progress	
PDM is completed	
PDM is paused	Resume PDM when ready.
PDM is resumed	No action is required.
PDM is stopped	If this action was not intentional, check the disk array's status.
PDM is switched to rebuild.	Replace the dead physical drive or reinstall the missing drive.
PDM is stopped internally	The destination drive was removed or used for a rebuild.

Reported Event	**Corrective Action**
Physical Disk	
Physical disk is marked online	No action is required.
Physical disk is marked offline	Replace the physical drive.
Physical disk is marked as DEAD.	
Physical disk has been reset	
Physical disk assigned as global spare	No action is required.
Global Spare has been deleted	
Physical Disk is no longer assigned as a global spare	
Physical disk assigned as dedicated spare	
Dedicated Spare has been deleted	
Physical Disk is no longer assigned as a dedicated spare	
Physical disk has been inserted	
Physical disk has been removed	Insert the physical drive back into the system.
Command on physical disk has been re-tried	If this message appears repeatedly, replace the physical drive
Physical disk ECC error is detected	Replace the physical drive.
Physical disk CRC error is detected	
Bad sector is found on physical disk	If this message appears repeatedly, replace the physical drive.
Error is detected in remap sectors	
Command times out on physical drive	
Physical disk negotiation speed is decreased.	
Previously configured disk is no longer found	Insert the physical drive back into the system.
A physical disk has encountered an unknown (non-ECC) media error.	If this message appears repeatedly, replace the physical drive.
A physical disk has encountered PFA condition	Clear the PFA condition. If this message appears repeatedly, replace the physical drive.
A configured dead physical drive has been inserted	Replace the physical drive.
A physical drive page 0 settings have been changed	No action is required.
A physical drive page 1 settings have been changed (SATA drives)	
A physical drive page 3 settings have been changed (SAS drives)	

Reported Event	**Corrective Action**
Physical Disk	
Physical disk is marked as DEAD due to removal	Replace the physical drive.
Physical disk is marked as DEAD due to failure of reassign sectors command	
Physical disk is marked as DEAD due to PFA condition	
Physical disk is marked as DEAD due to forced offline state	
Physical disk seen by partner controller not seen here	Check and correct SAS connections. Verify that SAS-to-SATA adapters are installed on all SATA drives.
Single ported physical disk seen by Partner controller not seen here	Install an SAS-to-SATA adapter on the SATA drive.
Physical disk reported not ready	Replace the physical drive.

Reported Event	**Corrective Action**
PSU (Power Supply Units)	
PSU is not inserted	Reinstall the power supply unit.
PSU is off	Turn on the power supply or plug in the power cable.
PSU is on	No action is required.
PSU is installed and turned on	
PSU is functional and turned on	
PSU is installed and turned off	Turn on the power supply or plug in the power cable.
PSU is functional and turned off	
PSU is malfunctioning and turned on	Replace the power supply unit.
PSU is malfunctioning and turned off	
PSU has been removed	
PSU 12V/5V/3.3V power is out of the threshold range	
PSU 12V/5V/3.3V power is within the normal range	No action is required.
PSU is critical. This may cause instab***ility of the system***	Check the power to the PSU. Verify that the correct PSU is installed.
PSU Fans	
PSU fan or blower has turned on	No action is required.
PSU fan or blower has turned off	
PSU fan or blower speed is increased	
PSU fan or blower speed is decreased	
PSU fan or blower is malfunctioning	Replace the power supply.
PSU fan or blower is inserted	No action is required.
PSU fan or blower is functioning normally	
PSU fan or blower is NOT installed	Check fans or blowers.
PSU fan status is unknown.	Check for proper installation and turn on the power supply. If the condition persists, replace the power supply.

Reported Event	**Corrective Action**
RAID Level Migration	
RAID level migration is started	No action is required.
RAID migration is in progress	
RAID level migration is completed	
RAID level migration is paused	Resume migration when ready.
RAID level migration is resumed	No action is required.
RAID level migration is stopped	If this action was not intentional, check the logical drive's status.
RAID level migration has encountered a physical disk error	Check the disk drive check table after migration and replace disk drive as needed.
RAID level migration is aborted due to an internal error.	Reduce system load on the VTrak.
RAID level migration is queued	No action is required.
Migration has detected stale NV Watermark	Wait to see if the watermark clears.
Migration has cleared stale NV Watermark	No action is required.
Array was made incomplete due to missing NV Watermark	If the array is online, try migration again. If the array is offline, delete and recreate the array.
User has accepted Incomplete Array. (Caused by a missing NV Watermark)	Rebuild the disk array.

Reported Event	**Corrective Action**
Rebuild	
Rebuild is started	No action is required.
Rebuild is in progress	
Rebuild is completed	
Rebuild is paused	Resume rebuild when ready.
Rebuild is resumed	No action is required.
Rebuild is stopped	If this action was not intentional, check the logical drive's status.
Rebuild stopped internally	Contact Tech Support.
Rebuild is aborted	Reduce system load on the VTrak.
Rebuild is queued	No action is required.
Auto rebuild cannot start	Install a target physical drive of adequate capacity.
Redundancy Check	
Redundancy Check is started	No action is required.
Redundancy Check is completed	
Redundancy Check is paused	Resume Redundancy Check when ready.
Redundancy Check is resumed	No action is required.
Redundancy Check is stopped	
Redundancy Check is aborted due to internal error	Reduce system load on the VTrak.
Redundancy Check encountered inconsistent block(s)	Check the disk drive check table after RC and replace disk drive as needed.
Redundancy Check task is queued	No action is required.
Redundancy check is in progress	
Redundancy Check task is stopped internally	Restore the disk array to functional status.
Redundancy check is started on unsynchronized logical drive	No action is required.

Reported Event	**Corrective Action**
Resource	
Resource is NOT available	Reduce system load on the VTrak.
SCSI	
SCSI host interface controller settings have changed	No action is required.
SEP	
SEP is found	No action is required.
SEP is NOT found	Insert or replace SEP hardware.
SEP I2C device access failure	f this message appears repeatedly, contact Tech Support.
SEP I2C device access recovered from failure	
Spare Check	
Spare check started on the given spare drive	No action is required.
Spare check completed successfully on the given spare drive	
Spare Drives	
Physical disk assigned as global spare	No action is required.
Physical disk is no longer assigned as global spare	
Global Spare has been deleted	
Physical disk assigned as dedicated spare	
Physical disk is no longer assigned as dedicated spare	
Dedicated Spare has been deleted	

Reported Event	**Corrective Action**
SMART	
SMART error is received	If this message appears repeatedly, replace the physical drive.
Stripe Level Migration	
Stripe Level migration is started	No action is required.
Stripe Level migration is completed	
Stripe Level migration is paused	Resume SLM when ready.
Stripe Level migration is resumed	No action is required.
Stripe Level migration is stopped	If this action was not intentional, check the logical drive's status.
Stripe Level migration has encountered a physical disk error	Check the physical drive check table after OCE is finished.
Stripe Level migration is aborted due to an internal error.	Reduce system load on the VTrak.
Stripe Level migration is queued	No action is required.

Reported Event	**Corrective Action**
Synchronization	
Synchronization is started	No action is required.
Synchronization is completed	No action is required.
Synchronization is paused	Resume synchronization when ready.
Synchronization is resumed	No action is required.
Synchronization is stopped	
Synchronization is aborted due to an internal error.	Reduce system load on the VTrak.
Synchronization is queued	No action is required.
Synchronization is stopped internally	
Subsystem (VTrak)	
The Subsystem is started	No action is required.
The Subsystem is stopped	
Subsystem parameter(s) are changed by user	
System is set to Redundant mode	
System is set to Critical mode	Check controller operation.
System is set to Non-Redundant mode	If your system has two controllers, check controller operation.
Transition	
Transition is started	No action is required.
Transition is completed	
Transition is paused	Resume transition when ready.
Transition is resumed	No action is required.
Transition is stopped	If this action was not intentional, check the disk array's status.
Transition was switched to rebuild	Replace the dead physical drive or reinstall the missing drive.
Unknown	
Unknown priority reason is detected	If this message appears repeatedly, contact Tech Support.
Zoning	
Zoning permission settings with the expander has been reset to defaults	No action is required.
Zoning expander has been rebooted.	
Zoning permission settings with the expander different than expected	Settings have been updated correctly. No action is required.

Contacting Technical Support

PROMISE Technical Support provides several support options for PROMISE users to access information and updates. We encourage you to use one of our electronic services, which provide product information updates for the most efficient service and support.

PROMISE E-Support: https://support.promise.com

PROMISE web site: http://www.promise.com//

When you contact us, please have the following information available:

- Product model and serial number
- BIOS, firmware, and driver version numbers
- A description of the problem / situation
- System configuration information, including: motherboard and CPU type, hard drive models, SAS/SATA/ATA/ATAPI drives & devices, and other controllers.

United States

580 Cottonwood Drive

Milpitas, Ca 95035, USA

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Australia

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

EMEA

Netherlands

Science Park Eindhoven 5228

5692 EG Son, The Netherlands

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Austria

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

France

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Germany

Europaplatz 9

44269 Dortmund, Germany

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Sweden

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Switzerland ITF

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Norway ITF

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Belgium

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Luxembourg

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

United Kingdom

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Taiwan

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

China

Room 1108, West Wing, Shi Chuang Plaza, 22 Information Road

Shangdi IT Park, Haidian District, Beijing 100085

Fax: 86-10-8857-8015

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Korea

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Hong Kong

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Singapore

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Japan

3F, Mura Matsu Bldg, 3-8-5, Hongo Bunkyo-ku

Tokyo 113-0033, Japan

Technical Support (E-Support): https://support.promise.com

Web site: http://www.promise.com//

Limited Warranty

PROMISE Technology, Inc. (“PROMISE”) warrants that this product, from the time of the delivery of the product to the original end user:

a) all components, except the cache backup battery, for a period of three (3) years;

b) the cache backup battery, for a period of one (1) year;

c) will conform to PROMISE’s specifications;

d) will be free from defects in material and workmanship under normal use and service.

This warranty:

a) applies only to products which are new and in cartons on the date of purchase;

b) is not transferable;

c) is valid only when accompanied by a copy of the original purchase invoice.

d) Is not valid on spare parts.

This warranty shall not apply to defects resulting from:

a) improper or inadequate maintenance, or unauthorized modification(s), performed by the end user;

b) operation outside the environmental specifications for the product;

c) accident, misuse, negligence, misapplication, abuse, natural or personal disaster, or maintenance by anyone other than a PROMISE or a PROMISE-authorized service center.

Disclaimer of other warranties

This warranty covers only parts and labor, and excludes coverage on software items as expressly set above.

Except as expressly set forth above, PROMISE disclaims any warranties, expressed or implied, by statute or otherwise, regarding the product, including, without limitation, any warranties for fitness for any purpose, quality, merchantability, non-infringement, or otherwise. PROMISE makes no warranty or representation concerning the suitability of any product for use with any other item. You assume full responsibility for selecting products and for ensuring that the products selected are compatible and appropriate for use with other goods with which they will be used.

PROMISE does not warrant that any product is free from errors or that it will interface without problems with your computer system. It is your responsibility to back up or otherwise save important data before installing any product and continue to back up your important data regularly.

No other document, statement or representation may be relied on to vary the terms of this limited warranty.

PROMISE's sole responsibility with respect to any product is to do one of the following:

a) replace the product with a conforming unit of the same or superior product;

b) repair the product.

PROMISE shall not be liable for the cost of procuring substitute goods, services, lost profits, unrealized savings, equipment damage, costs of recovering, reprogramming, or reproducing of programs or data stored in or used with the products, or for any other general, special, consequential, indirect, incidental, or punitive damages, whether in contract, tort, or otherwise, notwithstanding the failure of the essential purpose of the foregoing remedy and regardless of whether PROMISE has been advised of the possibility of such damages. PROMISE is not an insurer. If you desire insurance against such damage, you must obtain insurance from another party.

Some states do not allow the exclusion or limitation of incidental or consequential damages for consumer products, so the above limitation may not apply to you.

This warranty gives specific legal rights, and you may also have other rights that vary from state to state. This limited warranty is governed by the State of California.

Your Responsibilities

You are responsible for determining whether the product is appropriate for your use and will interface with your equipment without malfunction or damage. You are also responsible for backing up your data before installing any product and for regularly backing up your data after installing the product. PROMISE is not liable for any damage to equipment or data loss resulting from the use of any product.

Returning the Product For Repair

If you suspect a product is not working properly, or if you have any questions about your product, contact our Technical Support staff, and be ready to provide the following information:

- Product model and serial number (required)
- Return shipping address
- Daytime phone number
- Description of the problem
- Copy of the original purchase invoice

The technician helps you determine whether the product requires repair. If the product needs repair, the technician issues an RMA (Return Merchandise Authorization) number.

Important

Obtain an RMA number from Technical Support **before** you return the product and write the RMA number on the label. The RMA number is essential for tracking your product and providing the proper service.

Return ONLY the specific product covered by the warranty. Do not ship cables, manuals, CDs, etc.

USA and Canada:	PROMISE Technology, Inc. Customer Service Dept. Attn.: RMA # _____ 47654 Kato Road Fremont, CA 94538
Other Countries:	Return the product to your dealer or retailer. Contact them for instructions before shipping the product.

You must follow the packaging guidelines for returning products:

- Use the original shipping carton and packaging
- Include a summary of the product's problem(s)
- Write an attention line on the box with the RMA number
- Include a copy of your proof of purchase

You are responsible for the cost of insurance and shipment of the product to PROMISE. Note that damage incurred due to improper transport or packaging is not covered under the Limited Warranty.

When repairing returned product(s), PROMISE may replace defective parts with new or reconditioned parts, or replace the entire unit with a new or reconditioned unit. In the event of a replacement, the replacement unit is under warranty for the remainder of the original warranty term from purchase date, or 30 days, whichever is longer.

PROMISE pays for standard return shipping charges only. You must pay for any additional shipping options, such as express shipping.

Information for China RoHS

部件名稱 Products Description	Toxic or Hazardous Substances					
	鉛 (Pb)	汞 (Hg)	鎘 (Cd)	六價鉻 (Cr6+)	多溴联苯 (PBB)	多溴二苯醚 (PBDE)
PCBA	X	○	○	○	○	○
Metal parts	X	○	○	○	○	○
Plastic parts	○	○	○	○	○	○
Cable	○	○	○	○	○	○
Power Supply	X	○	○	○	○	○
Battery	X	○	○	○	○	○
Package	○	○	○	○	○	○

本表格依據 SJ/T 11364 的規定编制。
This table is prepared in accordance with the provisions of SJ/T 11364.
O：表示该有害物质在该部件所有均質材料中的含量均在 GB/T 26572 规定的限量要求以下。
O: Indicates that said hazardous substance contained in all of the homogeneous materials for this part is under the limitation requirement of GB/T 26572.
X：表示该有害物质至少在该部件的某一均質材料中的含量超出 GB/T 26572 规定的限量要求。
X: Indicates that said hazardous substance contained in at least one of the homogeneous materials used for this part is above the limit requirement of GB/T 26572.
此产品符合 EU RoHS 指令 2011/65/EU
The product complies with EU RoHS Directive 2011/65/EU

生產日期代碼參考路徑：*http://www.promise.com/Manufacturing-Dates-of-Products*