


PROMISE[®]
TECHNOLOGY, INC.



Best Practices Guide

Promise[®] VTrak E610f External Disk Array Subsystem and
FalconStor[®] VTL

November 2, 2009

Table of Contents

About This Guide	2
Introduction	2
Guidelines	3
Choosing a Suitable Topology	4
DAS.....	4
HA Options for DAS.....	4
Storage Level	4
Appliance Level	4
SAN.....	5
HA Options for SAN.....	5
Access Control in a SAN Topology	5
Configuring the VTrak	6
Single/Dual Controller Module	6
Memory Option	6
Hard Drive.....	6
Recommended Settings	6
Alert Notification.....	7
Configuring the RAID	8
RAID 5 and RAID 6.....	8
Global Revertible Hot Spare	9
Logical Drive Attributes	9
Initializing the Logical Drives	10
Setting up Access Control	10
Assigning Storage to VTL	10
Enabling LUN Mapping.....	10
Configuring LUN Mapping	11
Assign VTL Tape to VTL Clients.....	11
Tuning the Performance	11
Distributing I/O Load among Subsystems	11
Distributing I/O Load on Logical Drives	11
Distributing I/O Load on Fibre Channel Ports.....	11
Identifying and Removing Slow Drives	11
Summary	12
About Promise	13
Contact Us	13

List of Figures

Figure 1. DAS topology	4
Figure 2. SAN topology	5
Figure 3. Recommended RAID configuration	8
Figure 4. How global revertible hot spares work.....	9

List of Tables

Table 1. Recommended settings	7
Table 2. Recommended settings for logical drive attributes	10

About This Guide

This *Best Practices Guide* describes the recommended VTrak E610f configurations and tips for achieving the best performance with the VTrak E610f and the VTL appliance. It is intended for system administrators, storage architects, system integrators, and resellers who are planning to install or qualify FalconStor VTL using Promise VTrak E610f.

Introduction

The explosive growth of data generated by businesses continues to bring challenges in the area of backup and recovery operations. Increasing volumes of data and narrowing backup windows have made traditional tape backup increasingly difficult to maintain. In addition, tape backup has often led to longer recovery periods, leaving more room for media and mechanical failures. Businesses require faster and more reliable ways to back up and recover their data.

Virtual tape appliances, such as FalconStor Virtual Tape Library (VTL) Storage Appliances and Enterprise Edition, address the challenges associated with tape backup. FalconStor VTL solutions offer fast backup, data deduplication, enterprise-wide replication, and tape integration in one solution, without requiring changes to the existing environment. Built-in data deduplication features automatically and transparently eliminate redundant data, extending data retention and allowing users to reduce or even eliminate investments in physical tape media.

However, FalconStor VTL solutions require massive amounts of storage that must scale both in performance and in capacity. In the market, many existing disk array products are designed for traditional enterprise transactional computing. They cannot provide the scalability that VTL applications require, and they are usually too expensive to deploy in large volume.

Promise VTrak E610f external disk array subsystem was designed to work seamlessly with VTL applications, providing exceptional input/output (I/O) performance for concurrent VTL operations like backup, restore, and deduplication with a host of features, including:

- Aggregated 1.6Gb/s bandwidth and low overhead implementation
- Support for RAID levels 0, 1, 10, 1E, 5, 6, 50, and 60
- No single point of failure design
- Dual 4Gb Fibre Channel host interface ports
- Robust error recovery
- Active-active dual controllers

The VTrak E610f has 16 drive bays, and can support up to 80 SAS or SATA drives through Promise VTrak J610S expansion units, or up to 80 terabytes (TB) of raw capacity (based on using a 1TB SATA drive). It offers a highly scalable, reliable, and cost-effective storage solution for FalconStor VTL Appliances and Enterprise Edition, appropriate for small and medium-size businesses (SMBs) to Enterprise customers.

The recommendations in this best practices guide were tested on the following hardware and software versions:

- Promise VTrak E610f external disk array subsystem, firmware version 3.33.0000.00 (SR 2.5)
- Promise VTrak J610s external disk subsystem, firmware version 1.07.0000.12
- FalconStor VTL Server v5.1

Guidelines

This section provides an overview of the topics discussed in this best practices guide for configuring and tuning the VTrak E610f and FalconStor VTL to achieve optimal performance. Detailed instructions are provided in the corresponding sections of this *Guide*.

Choosing a Suitable Topology and HA Option

- DAS
- SAN

Configuring the VTrak

- Single/Dual Controller Module
- Memory Option
- Hard Drive
- Recommended Settings
- Alert Notification

Configuring the RAID

- RAID 5 and RAID 6
- Global Reversible Hot Spare
- Logical Drive Attributes
- Initializing the Logical Drives

Setting up Access Control

- Assigning Storage to VTL
- Assigning VTL Tape to VTL Clients

Tuning the Performance

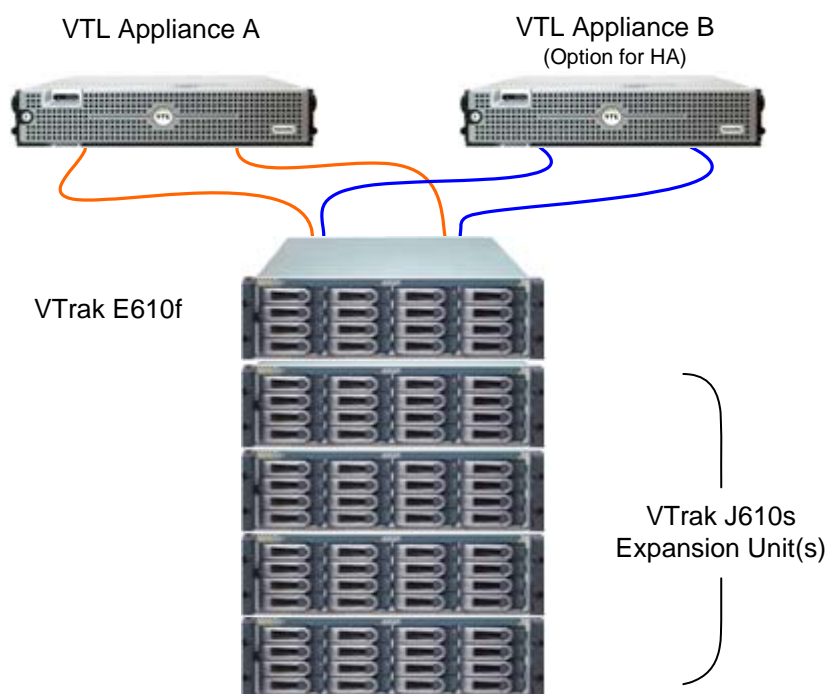
- Distributing I/O Load among Subsystems
- Distributing I/O Load on Logical Drives
- Distributing I/O Load on Fibre Channel Ports
- Identifying and Removing Slow Drives

Choosing a Suitable Topology

To support different VTL market segment requirements, the VTrak E610f offers two topology options for VTL: Direct Attached Storage (DAS) and Storage Area Network (SAN). A DAS topology (shown in Figure 1) is a lower cost approach, suitable for small to medium scale deployments. A SAN topology is massively scalable in terms of both capacity into the petabyte range and of performance up to multiple terabytes per hour, for large-scale backup applications with stringent backup window requirements. With both the DAS and SAN topologies, there are High Availability (HA) options that can be configured at the storage (VTrak) and appliance (VTL) levels for meeting different application requirements.

DAS

Figure 1. DAS topology



In a DAS topology, the VTL appliances are directly connected to the VTrak E610f through their Fibre Channel ports.

The VTrak E610f can be connected to up to four Promise VTrak J610s expansion units.

HA Options for DAS

There are two **levels of high** availability options for meeting different application requirements.

Storage Level


The VTrak E610f supports dual controller modules, which offer fully redundant paths to the VTL appliance with no single point of failure. As shown in Figure 1 above, each VTL appliance has two Fibre Channel ports, each of which can be connected to one Fibre Channel port on each of the storage controllers.

Appliance Level

This HA option offers redundancy at the VTL appliance level for the VTL cluster. If one of the VTL appliances fails, the other VTL appliance takes over the VTL requests from the backup clients without service interruption. In this case, an optional VTL appliance could be installed with two Fibre Channel ports, each connected to the remaining Fibre Channel port on each of the storage controllers.

SAN

Figure 2. SAN topology



In a SAN topology (Figure 2), multiple VTL appliances are connected to multiple VTrak E610f units through one or more Fibre Channel switches (collectively called an FC fabric). Similar to the DAS topology, the capacity of each VTrak E610f could be expanded using VTrak J610s expansion units through SAS connections.

HA Options for SAN

A SAN topology also provides two levels of high availability options, similar to the DAS topology – at the storage level and at the VTL appliance level. The difference is the appliances and storage are always connected through the Fibre Channel switches instead of being directly connected. As illustrated in Figure 2 above, this topology offers completely redundant paths from all VTL appliances to all storage subsystems, with no single point of failure.

Access Control in a SAN Topology

As the SAN topology places multiple VTL appliances and VTrak E610f units inside the same storage network, it is important to configure the access control among them properly. This can be either achieved through zoning on the Fibre Channel switch or through logical unit number (LUN) mapping on the VTrak E610f units. Refer to the section in this best practices guide entitled “Setting up Access Control” for more details about how to configure the access control between the VTL appliances and VTrak units.

Configuring the VTrak

This section lists best practices for configuring your VTrak E610f with FalconStor VTL for optimal performance. In addition to the best practices specified in this section, we recommend you also follow the general best practices for Promise VTrak E-Class RAID solutions in the following knowledge base article.

10093 – Best practices for E-Class VTrak RAID

<http://kb.promise.com/KnowledgebaseArticle10093.aspx?Keywords=best+practice>

Single/Dual Controller Module

Depending on your application availability and performance requirements, you may choose either the single or dual controller module configuration. To achieve high availability and optimal performance, **Promise recommends using the dual controller module configuration.**

Memory Option

For high performance VTL applications, **Promise recommends using two gigabytes of memory, on the controller, and four-cell battery for ample cache memory backup retention.**

Two gigabytes of memory and four-cell battery are not the default configuration of VTrak E610f. When ordering, ask for the following SKUs:

- Two gigabyte memory, SKU VTEMEM2G
- Four-cell battery, SKU VTBAT4CELL

To place your order, please contact Promise Sales. See “Contact Us” on page 5.

Hard Drive

VTL is generally a high capacity application. For this reason, **Promise recommends using high capacity, low cost enterprise-capable SATA drives.** Please refer to the Promise VTrak E610f compatibility matrix on the Promise website, under Support > Compatibility Matrix, for compatible hard drive models.

Recommended Settings

The following table lists the recommended settings for the VTrak E610f for VTL and Single Instance Repository (SIR) applications:

Table 1. Recommended settings

Setting	Value	Rationale
Subsystem		
Redundancy Type	Active-Active	Supports high availability and load balancing with VTL/SIR. The default setting is Active-Active.
Controller		
LUN Affinity	Disabled	Ensures compatibility with VTL/SIR. The default setting is Disabled.
Force Read Ahead	Enabled	Achieves the optimal performance for multiple stream backup/restore and deduplication. The default setting is Disabled. Please refer to VTrak E610f Product Manual about how to enable it.
Adaptive Write Back Cache	Enabled	Protects the data integrity in case of Battery Backup Unit (BBU) failure or a low battery condition. The default setting is Enabled.
Auto Rebuild	Enabled	Enable Automatic Rebuild to replace bad drive or Predictive Drive Migration (PDM) to replace a drive that is going to fail. These actions improve data reliability by reducing or eliminating periods of lost redundancy. The default setting is Enabled.
SATA Physical Drive Settings		
ReadAheadCache	Enabled	Helps to achieve the optimal performance for multiple stream backup/restore and deduplication. The default setting is Enabled.
Write Cache	Enabled	Helps to achieve the optimal performance for multiple stream backup/restore and deduplication. The default setting is Enabled.
Medium Error Threshold	64	This value is the number of bad blocks tolerated before the controller marks a physical drive as Dead. By enabling this setting, you can identify problem drives before they fail, thereby improving system performance and reliability. The default setting is 0, meaning disabled. Please refer to E610f Product Manual about how to set Medium Error Threshold. Also see the best practice at: http://www.promise.com/support/support_eng1.asp , And search for keyword "Medium Error Threshold."

Promise recommends leaving all other settings at their factory default values.

Alert Notification

It is important to set up the alert notification on the Promise VTrak E610f external disk array subsystem, as this feature notifies the system administrator and allows him to act accordingly in the event of any status change or failure condition.

The administrator should set up event notification either through Email or Simple Network Management Protocol (SNMP) during the initial setup phase.

Usually, Email is a preferred notification method. In certain user environments, there may be an existing alert notification infrastructure that is based on a standard network management protocol such as SNMP. In this case, SNMP notification can be configured to integrate with the existing infrastructure.

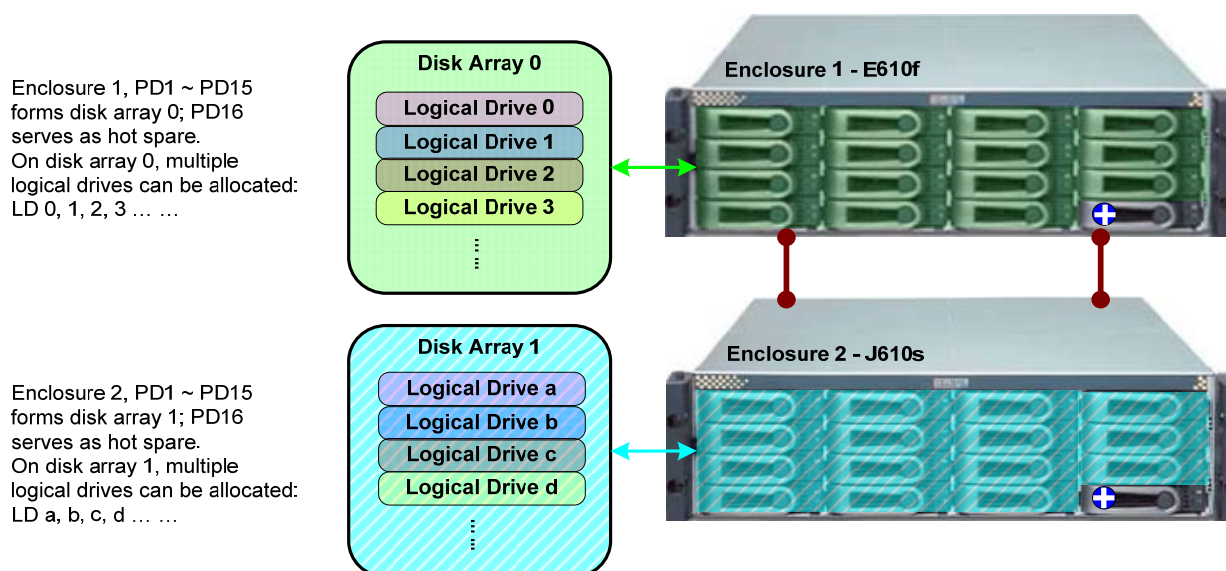
For detailed instructions on how to configure event notification through Email and SNMP, please refer to the VTrak E610f Product Manual. You can find all of the Product Manuals on the Promise website under the Downloads section (Support > Downloads).

Configuring the RAID

Promise VTrak E-Class external disk array subsystems support RAID levels of 0, 1, 1E, 10, 5, 6, 50 and 60. They also support the mixing of different RAID level volumes in the same disk group to meet the requirements of different application performance profiles.

This section provides the proven configurations for typical VTL applications to achieve the best backup and deduplication performance.

Figure 3. Recommended RAID configuration



As illustrated Figure 3, in each enclosure, PD (physical drive) 1 through 15 form a disk array, and PD 16 serves as the hot spare. In each disk array, multiple logical drives (volumes) can be carved out. Depending on the capacity requirement, this configuration can be repeated on up to four VTrak J610s expansion units per VTrak E610f RAID head.

Depending on the application requirement, you may configure more drives as hot spares, or you may also configure fewer hot spares.

RAID 5 and RAID 6

The recommend RAID configuration for the logical drives is RAID 6, as RAID 6 offers excellent reliability due to double spindle failure protection. In the case that logical drives are used to store backup data temporarily for later deduplication processing, the logical drives can be configured as RAID 5.

In either case, you should configure at least one hot spare drive for every 16 physical drives to protect the RAID redundancy in case of drive failure.

Global Revertible Hot Spare

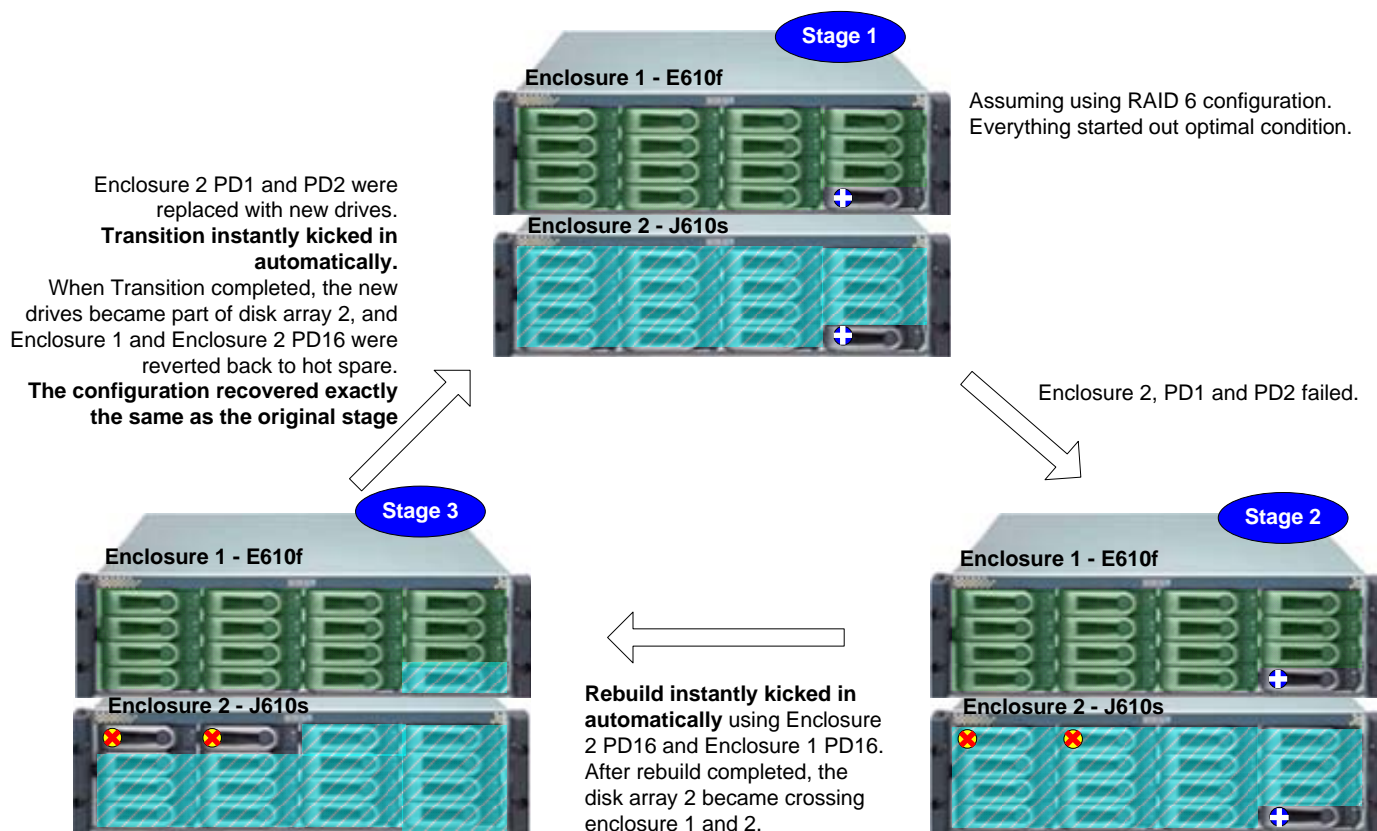
To achieve the best reliability and manageability, **Promise recommends configuring all of the spares as global revertible hot spares**. The advantage of global revertible hot spares is that each of the hot spares can be used for rebuilding the same logical drive or different logical drives, in the same enclosure or in different enclosures. Therefore, when there are unevenly distributed drive issues, there are still enough hot spares to provide coverage and to allow RAID redundancy to be restored in the shortest time frame.

As revertible spares, when the rebuild completes and the failed drive is replaced with a new drive, the used spare will automatically be freed up and become a revertible spare again after the reversion process, called Transition, initiates and completes in the background.

Due to reliability considerations, **Promise does not recommend mixing drives from different enclosures in the same disk array**. Therefore, when a rebuild completes through using a hot spare from a different enclosure, Promise recommends replacing the failed drive as soon as possible. When the new drive is inserted, Transition will start automatically in the background. Eventually, the disk array will consist of drives in the same enclosure again.

The follow diagram illustrates how global revertible hot spares work.

Figure 4. How global revertible hot spares work



As illustrated in Figure 4, global revertible hot spares ensure overall system reliability in case of drive failure. In addition, they also facilitate management by enabling the entire rebuild and Transition process to be automated. Once the process completes, the system is restored to its original state.

Logical Drive Attributes

To achieve the best performance, Promise recommends configuring each of the logical drives with the attribute values as listed in the table below:

Table 2. Recommended settings for logical drive attributes

Attribute Name	Value	Rationale
Stripe Size	64KB	Backup, restore, and deduplication I/O requests are usually 128KB, 256KB, or larger. A stripe size of 64KB can ensure the VTL or SIR I/O request will stripe across as many drive spindles as possible, hence achieving the best I/O performance.
Read Cache Policy	Read Ahead	VTL restore or SIR read requests are sequential for single streams. The Read Ahead setting boosts performance for sequential I/O. Even in the case of multiple stream reads, Promise lab tests indicated it boosted the performance significantly, combined with Force Read Ahead enabled at the controller level.
Write Cache Policy	Write Back	VTL backup requests are sequential write operations for single stream. However, a VTL appliance can usually serve multiple backup clients simultaneously. With multiple stream VTL backup, the write access becomes random. By enabling the Write Back cache, it avoids performance degradation due to random drive access, therefore improving the performance dramatically.
Other Attributes	Leave as default	

Initializing the Logical Drives

After creating the logical drives, make sure to initialize all of them. This action ensures optimal performance.

Depending on the capacity of the logical drives, it may take 40 minutes to a few hours to initialize all of the logical drives. Therefore, you are wise to plan ahead before running a performance test or placing a solution into production.

For detailed instructions about how to initialize logical drives, please refer to VTrak E610f Product Manual. You can find all of the Product Manuals on the Promise website under the Downloads section (Support > Downloads).

Setting up Access Control

This section offers instructions and recommendations for how to assign storage to VTL as well as how to assign VTL tape to VTL clients in a SAN topology. *These recommendations apply to SAN topologies only.*

Assigning Storage to VTL

Assigning storage to VTL can be achieved either through configuring zoning on the Fibre Channel switch or through configuring LUN Mapping on the VTrak E610f. This best practices guide provides instructions on configuring LUN Mapping on the VTrak E610f. For instructions on configuring zoning on the Fibre Channel switch, please refer to the Fibre Channel switch User Manual.

Enabling LUN Mapping

By default, LUN Mapping is disabled on the VTrak E610f. When LUN Mapping is disabled, all of the logical drives on the external array storage subsystem are accessible to all of the Fibre Channel initiator hosts. For a DAS topology or for a SAN topology using zoning access control, LUN Mapping should be disabled.

However, for a SAN topology with access control set through LUN Mapping, LUN Mapping must first be enabled on the subsystem before access control can be configured. Please refer to the VTrak E610f Product Manual for detailed instructions on how to enable LUN Mapping.

Configuring LUN Mapping

Once LUN Mapping is enabled, it can be configured per logical drive and per initiator, in other words, on a per VTL appliance basis. To assign a logical drive to a VTL appliance, create a LUN map between the logical drive and the Fibre Channel initiator ports on the VTL appliance.

For detailed instructions on how to configure LUN Mapping, please refer to the VTrak E610f Product Manual.

Assign VTL Tape to VTL Clients

A SAN topology is typically used to connect VTL appliances and VTL clients, where backup software has been installed. Access control is required to configure which client has access to which VTL appliance and, more specifically, to which VTL tape on the appliance.

Similar to assigning storage to VTL, the FalconStor VTL solution also offers a way to assign VTL tape devices to VTL clients, similar to LUN Mapping on the VTrak E610f. For detailed instructions on how to assign VTL Tape to VTL clients, please refer to the FalconStor VTL User Manual.

Tuning the Performance

This section lists tips for fine-tuning the VTrak E610f for optimal performance with the FalconStor VTL solution.

Distributing I/O Load among Subsystems

You can determine the I/O load on each subsystem through the real-time performance monitoring function in Promise WebPAM PROe software. Promise WebPAM PROe software is a web-based graphical user interface embedded inside each VTrak E610f. To access the graphical user interface, launch your Internet browser and type in the IP address of the Management port; you will be connected to the Promise WebPAM PROe software login screen.

Through the performance monitoring function, you can see the I/O load of each subsystem and can determine whether the loads are evenly distributed. If the loads are not evenly distributed, you can adjust the access control between the VTL clients and the access control between the VTL and the VTrak E610f to balance the I/O load.

Please refer to the section in this best practices guide on Setting up Access Control for more details about how to configure access control.

Distributing I/O Load on Logical Drives

Through the WebPAM PROe performance monitoring function, you can also find out whether the I/O load is evenly distributed among the logical drives. If the load is not evenly distributed, you may adjust the access control between the VTL and the logical drives of the VTrak E610f to balance the I/O load.

Please refer to the section in this best practices guide on Setting up Access Control for more details about how to configure access control.

Distributing I/O Load on Fibre Channel Ports

You can use the performance monitoring function to find out whether the I/O load on the Fibre Channel ports connected to the VTL is evenly distributed. If the load is not distributed evenly, you may change the preferred path to the logical drives through the FalconStor VTL management console.

Please refer to the FalconStor VTL User Manual for detailed instructions on how to change the preferred path to the logical drives from the storage subsystem.

Identifying and Removing Slow Drives

Not all SATA drives are equal. Some drives perform significantly slower than others, in term of response time and latency. Any slow drive brings down the overall VTL solution performance. Therefore, **Promise recommends stress testing the drives during a pre-production test, then identifying and removing the slow drives before you deploy the solution to the production environment.**

You may also find out whether there are any slow drives using the WebPAM PROe performance monitoring function. The WebPAM PROe real-time performance monitoring function is available in VTrak E610f firmware version 3.33.0000.00 and

later. Please refer to the VTrak E610f Product Manual for more detailed instructions on how to use the WebPAM PROe software and the real-time performance monitoring function.

Summary

The FalconStor VTL solution addresses the business challenge of making backup and recovery fast and reliable. Storage plays a significant role in the reliability and performance of the VTL solution. Through efficient use of the Promise VTrak E610f external disk array subsystem, the VTL backup and recovery performance becomes exceptional, the backup and recovery windows are significantly shortened, and the total cost of the VTL solution is dramatically reduced.

Promise VTrak E610f external disk array subsystem is a versatile storage solution that can support many different application requirements. However, for different applications, best practices may vary. In order to achieve the best performance when planning FalconStor VTL installation with Promise VTrak E610f external disk array subsystem, please follow the guidelines in this *Best Practices Guide* and refer to the additional resources specified in this paper.

About Promise

With a long history of innovation, Promise Technology develops and manufactures sophisticated RAID solutions recognized worldwide, ranging from a complete line of RAID controller cards to SAS/SATA RAID subsystems. Meeting enterprise, mid-range, and entry-level data protection needs, Promise products are distinguished by their common RAID code, unified management interface, and unparalleled support. Known as the originator of SATA/ATA RAID products, Promise's comprehensive product base includes High Availability (HA) standalone RAID subsystems with standards-based management interfaces, host-based (internal) RAID controllers, NAS appliances for SOHO, and SATA ASICs integrated into an extensive list of industry-leading motherboards. Headquartered in Milpitas, California, Promise has operations throughout Asia and Europe.

Contact Us

:: Promise USA

Milpitas, CA
1-408-228-1400
sales@promise.com

:: Promise EMEA

El Son, The Netherlands
31-40-235-2600
sales@eu.promise.com

:: Promise Taiwan

Hsinchu, Taiwan
886-3-578-2395
sales@tw.promise.com

:: Promise Germany

Dortmund, Germany
49-02-31-56-76-48-0
sales-de@eu.promise.com

:: Promise China

Beijing, China
86-10-8857-8085 or 8095
sales@cn.promise.com

:: Promise Italy

Rome, Italy
39-06-3671-2626
sales-it@eu.promise.com

:: Promise Japan

Tokyo, Japan
81-03-5333-3631
sales@jp.promise.com

:: Promise UK

Wokingham, England
44-0-870-112-59-77
sales@eu.promise.com

:: Promise Australia

Sydney, Australia
61-2-8065-2055
sales@au.promise.com

To learn more, visit www.promise.com

© 2009 Promise Technology, Inc. The information contained herein is subject to change without notice. The only warranties for Promise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Promise shall not be liable for technical or editorial errors or omissions contained herein.

Promise[®] is a trademark of Promise Technology, Inc. All other brand and product names are registered marks of their respective proprietors.

Issue v1.1, Nov 2009