PROMISE®
TECHNOLOGY, INC.

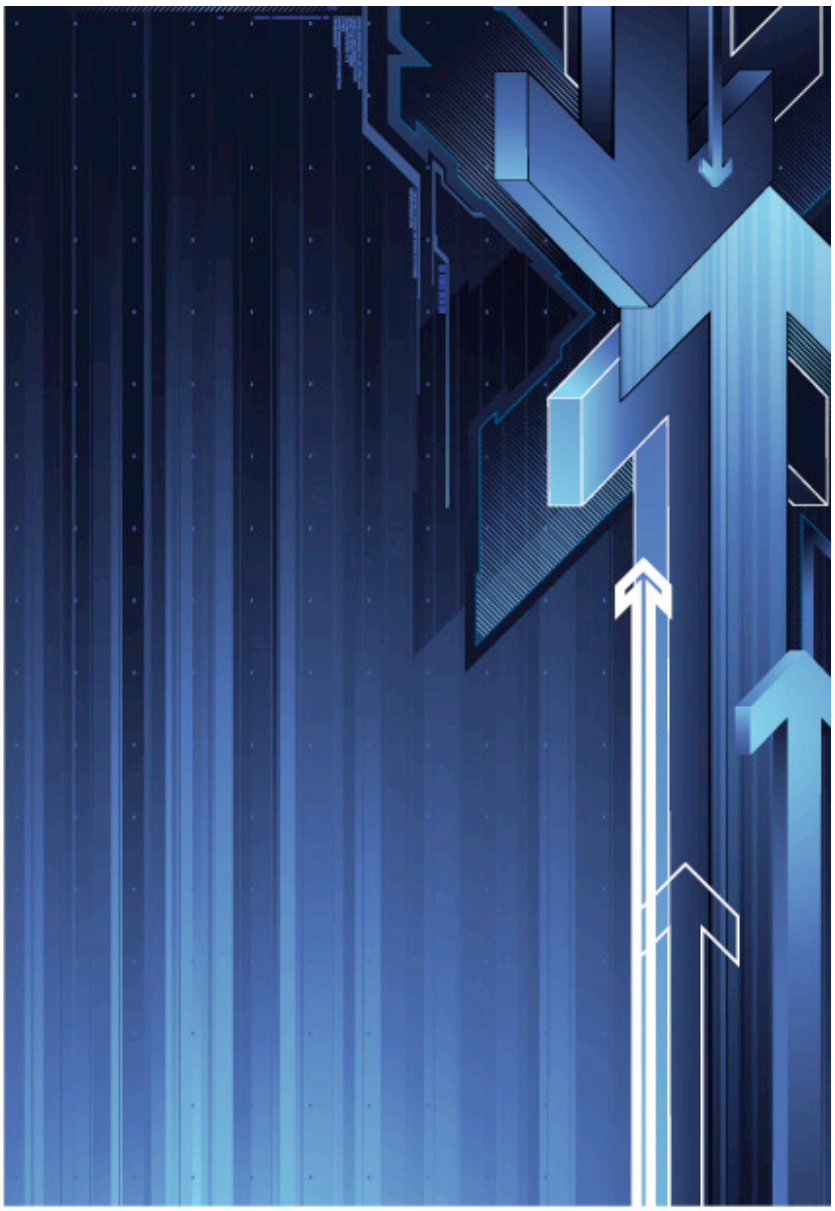

VTrak E-Class for Cloud Applications

Web based application supporting thousands of users.

Date: Q1.2010.008

Introduction

While web-based service businesses have been around for over a decade, the scope of services available has expanded dramatically to include many tasks that were once reserved for local physical hardware. With the popularity of notebooks, netbooks and other thin client wireless computing devices, the movement of applications and related data to a centralized “cloud” is now widespread. These cloud/web-based services have made it easy for the average user to manage and access their data from multiple devices, never having to struggle with the underlying server and storage technology.

Cloud Storage

For the following example deployment we will describe Promise storage technology as the key component to a cloud providers successful endeavor. Customers using the cloud service will be given all the features of corporate network storage including a remote copy of their local “home directory” without having to purchase expensive hardware and software. Other large files like video, audio and still image files will be a part of the content stored and shared using this service. For the corporate user this service could take the worry and expense out of the day-to-day administration associated with managing users laptops and remote office data. For the home user this service provides an enterprise level of reliable storage and offsite vaulting of their large files that was previously unaffordable.

Challenges

- Application services must always be available (99.999% availability).
- Server architecture needs to have enough bandwidth and raw throughput to support an ever-growing number of concurrent users.
- Home directory synchronization must complete in a reasonable amount of time (overnight).
- Service provider administrators need to be able to manage the whole operation remotely.
- All interaction with the service is done through secure connections.
- All aspects of the service must scale transparently without downtime.

For the service provider delivering these services, the IT solution requirements continue to grow but the goals remain the same: customer data must be there when it’s requested, regardless of component failure or even natural disaster. The solution must deliver consistent performance reliability and manageability to support the Service Levels Agreements (SLA) between the service provider and the end customer.

Host-independent high available RAID storage provides the foundation for any mission-critical server-based business today. The benefit of host independence combined with full-component redundancy insures the data from one server is easily accessible to others should the server hardware fail. Servers with internal RAID storage simply cannot provide this level of data availability.

Architecture

For the service provider, the first area of concern on a public network like the Internet is security. We would assume that the storage application was written to prevent unauthorized access to the users data at all stages of the process, as well as resist hackers both inside and outside of the provider's network. Both end-users and remote administrators must access the application and resources through a firewall; typically an application-delivery controller or network accelerator is used to optimize the application being provided. After authentication, the end-user accesses the service through a high performance “front-end” Ethernet network, which in most cases is logically and physically separate from the “back-end” network where the administrators manage the service and physical hardware (figure 1). The servers, VTrak E-Class storage and other managed devices have their Ethernet ports attached to this “back-end” network.

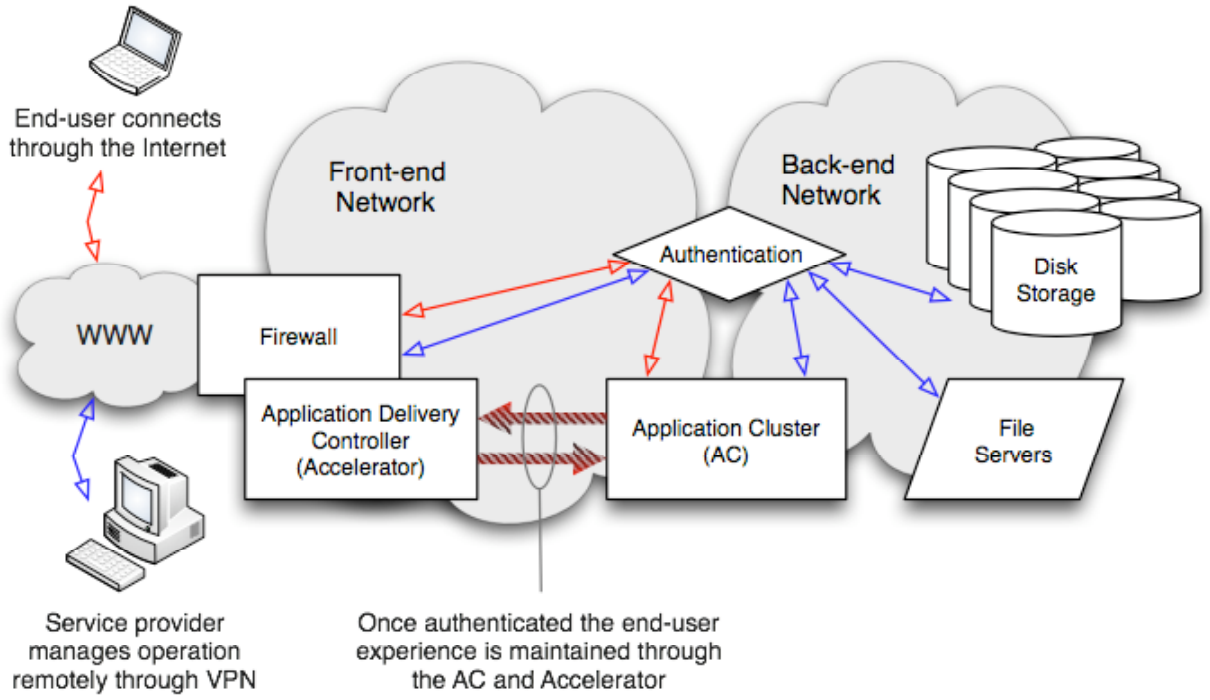

Figure 1: High level topology

With the VTrak management port connected to the back-end network, management interfaces are protected from front-end traffic and malicious attacks are virtually eliminated. VTrak uses secure socket layer (SSL) technology on WebPAM PRO, its browser-based user interface and Secure Shell (SSH) technology on its Command Line Interface (CLI), enabling administrators to manage the storage hardware remotely with full 128-bit encryption end-to-end.

The next step in planning the deployment is to understand the data processing requirements at each functional stage to determine the optimal choice of hardware and software to deliver reliable service, in order to meet customers' needs and keep their business.

Server Profiles

The workload at each step in the cloud storage process has unique processing profiles, so separating these tasks onto dedicated hardware allows the CPU and other physical resources to be scaled as needed, eliminating bottlenecks and maintaining overall resource utilization.

Server virtualization would also be utilized as Promise storage is certified for use on a number of popular platform, but details about server virtualization techniques is beyond the scope of this document.

Application Cluster Servers

The Application Cluster (AC) servers themselves require very little storage, but must have powerful processors and internal memory to support high performance I/O when serving thousands of unique requests. AC servers manage the customer experience while online, locating and directing the flow of end-customer data.

• High IOP and throughput
• Servers organized in pools to accommodate thousands of users
• CPU and memory intensive processing
• Small internal storage required

Authentication Servers

Closely tied to AC servers, Authentication servers have similar CPU and memory requirements but also need fully redundant, high-performance storage to store user-access and volume-location information requested by the AC servers. For this task, a VTrak E-Class populated with high performance 15K SAS drives is used.

- High IOP
- Servers organized in High Availability clusters
- CPU and memory intensive processing
- High Availability shared disk storage required with good random IOP disk performance

File Server Storage Nodes

File server Storage Nodes (SN) store all user data including a journal with details about changes made to the users data. With throughput as the primary task, SN servers must have massive storage capacity provided by SATA drives, along with fat network connectivity to the AC servers. In our example, the file servers are built on clustered Windows servers and shared to AC using Windows Distributed File Services (DFS). DFS allows a virtualized namespace to be used across multiple servers, simplifying connectivity and replication tasks.

- High Throughput
- Servers organized in failover clustered pairs
- Large memory cache
- High capacity, High Availability shared disk storage required

Storage Profile

With the basic user access profiles, network topology and data processing profiles laid-out, the next step is to logically layout the storage platform. The storage hardware needs to be flexible enough to support full component and host failover. The solution must also deliver performance, capacity and value that enable the service provider to offer the service at a competitive price. VTrak E-Class offers a No-Single-Point-of-Failure (NSPOF) design, and choice of Fibre Channel (FC) or Serial Attached SCSI (SAS) host interfaces, allowing for a number of valid configurations to meet uptime and availability requirements for this application.

Figure 2: Front image of VTrak E-Class RAID Subsystem

Flexible Storage

For the high capacity file servers low cost Direct Attached Storage (DAS) using SAS host connectivity will be used. To support multiple hosts and the small file/transactional data profile needed by the authentication servers SAN attached storage will be used. The promise VTrak E-Class enclosure design supports both high performance SAS and high capacity SATA drives. This allows the same enclosure to be used for different access profiles.

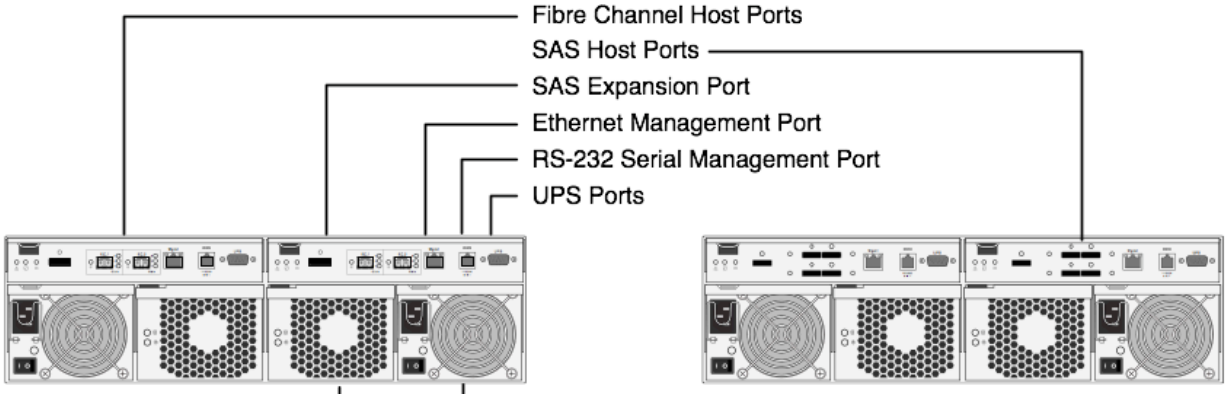

Fibre Channel - VTrak E-Class FC host connected models feature two host ports and one JBOD expansion port per controller.

SAS -VTrak E-Class SAS host connected models feature four host ports and one JBOD expansion port per controller

Figure 3: Rear views of VTrak E-Class in the 3U/16 bay form factor

Flexible Connectivity and Parts Interchange

VTrak E-Class features a high degree of parts interchangeability (figure 3). The authentication storage and the file server storage use the same pool of replacement parts. This arrangement minimizes the provider's investment in spare parts.

RAID Expansion

VTrak E-Class JBOD connectivity enables the user to dynamically expand the RAID subsystem incrementally or deploy fully expanded systems. The VTrak E-Class architecture supports attaching up to four JBOD subsystems.

Each 2U subsystem equals additional 12 drive bays, for a total of 60 bays in 10 units of standard-depth rack space.

Each 3U subsystem equals additional 16 drive bays, for a total of 80 bays in 15 units of rack space (figure 4).

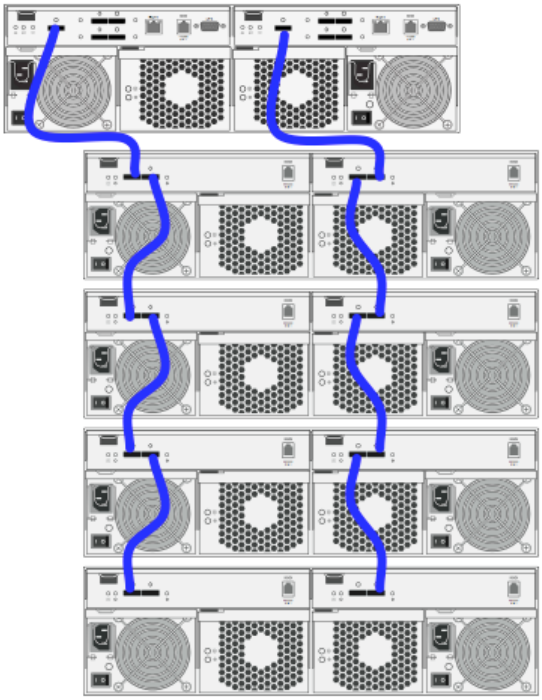

Figure 4: Rear view VTrak E-Class 3U/16-bay with JBOD expansion

SAS Topology

The VTrak E-Class SAS host interface offers a unique feature that allows RAID subsystems to be cascaded, providing massive capacity for directly attached servers. The cascade function works very much like a parallel SCSI daisy chain, where the host bus passes through each discrete RAID subsystem. Each controller features two IN ports and two OUT ports.

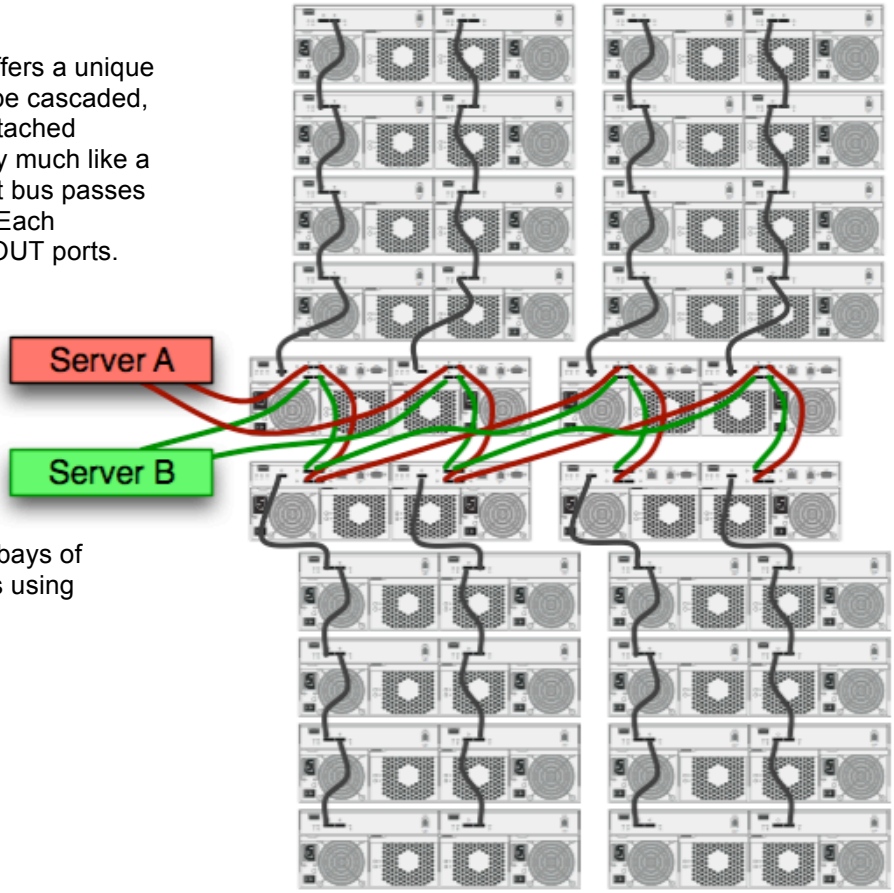

Promise Technology supports up to 230 bays of storage connected to dual or single hosts using cascade (figure 5).

Figure 5: VTrak E-Class with fully redundant SAS host cascade

Fibre Channel Topology

For providers of cloud services, Fibre Channel host connectivity offers the flexibility to go from a small direct-attached solution to a Storage Area Network (SAN) using the same VTrak hardware. Fibre Channel is a mature, high-performance industry standard in enterprise computing. Like Ethernet networking, the greatest benefit of FC is that it can provide diverse connectivity relationships that can improve availability and can effectively scale to very large deployments.

The use of SAN file systems can further expand Fibre Channel's networking capability, by enabling the storage to be pooled across physical hardware, providing both aggregated performance and transparent expansion to multiple hosts simultaneously. With a SAN file system, each server connected to the SAN has software installed that communicates with a "Master" or "Metadata server" over a dedicated Ethernet network, out-of-band from storage I/O. The Metadata server controls all access to the storage, virtualizing the physical hardware presented to the servers. Storage can then be provisioned and shared efficiently.

VTrak features initiator based LUN masking at no additional charge (figure 6). LUN masking allows shared storage targets and private storage targets to coexist securely across large SANs or even on a single VTrak subsystem.

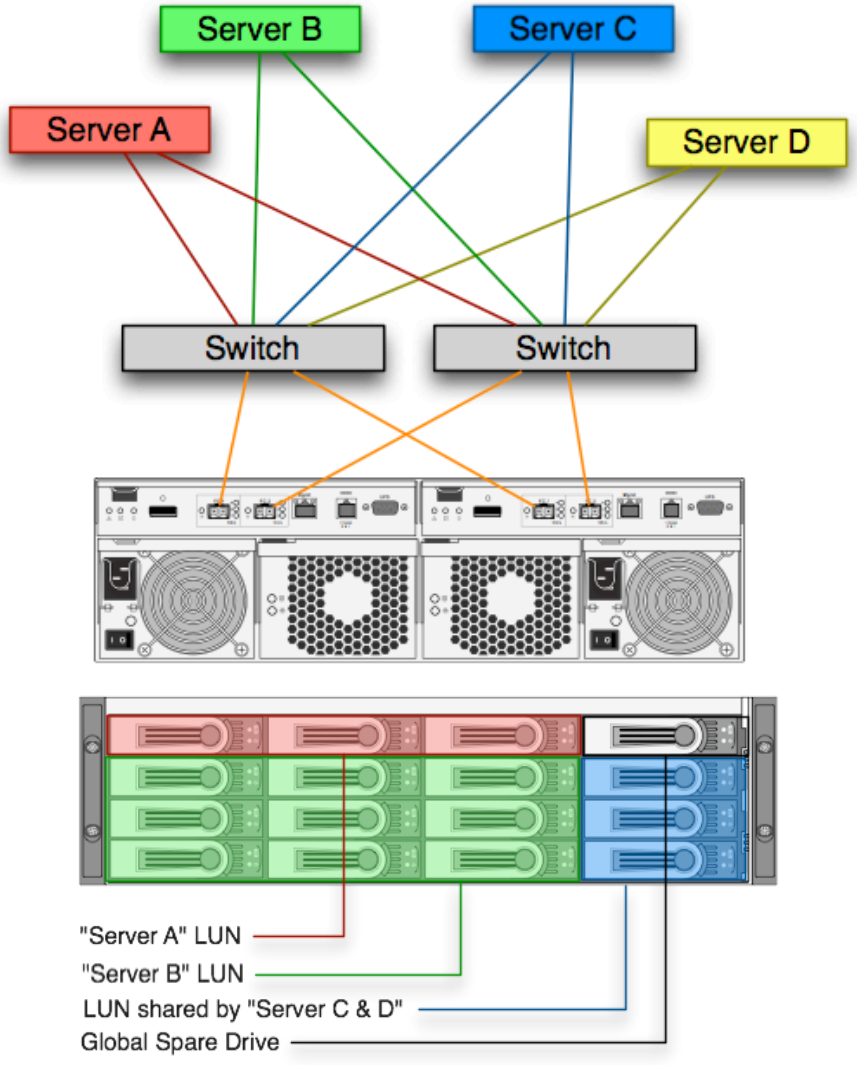

Figure 6: Fibre channel redundant switched topology

Data Storage Processing

End User Application

After purchasing the service, the end user downloads and installs a lightweight applet on his computer. The applet provides a secure remote login to the AC servers. The AC works with the Authentication servers that store user account information and file locations. The Authentication server then directs the AC to file locations on a storage volume shared by “SN servers”.

Authentication Database Cluster

To accommodate the high transaction rates needed to support thousands of connected users as well as the need to share the same data across a growing pool of servers, a VTrak with Fibre Channel SAN connectivity and 15K SAS drives offers the best fit.

Starting small, a single VTrak will be configured to deliver maximum performance to a pair of physical servers (figure 7). An Enterprise Data Center Edition of the database software supports both Scaling-up: Adding newer more powerful hardware, and Scaling-out: Adding more servers.

As the cluster grows, new servers are added dynamically to the SAN. When new physical servers are added, their host initiators appear in the VTrak management interface, where they are mapped or masked to the storage. VTrak has embedded performance monitoring that gives administrators real-time performance statistics, so they can see the current load on the subsystem. Once maximum storage performance thresholds are reached, more VTrak E-Class units should be added.

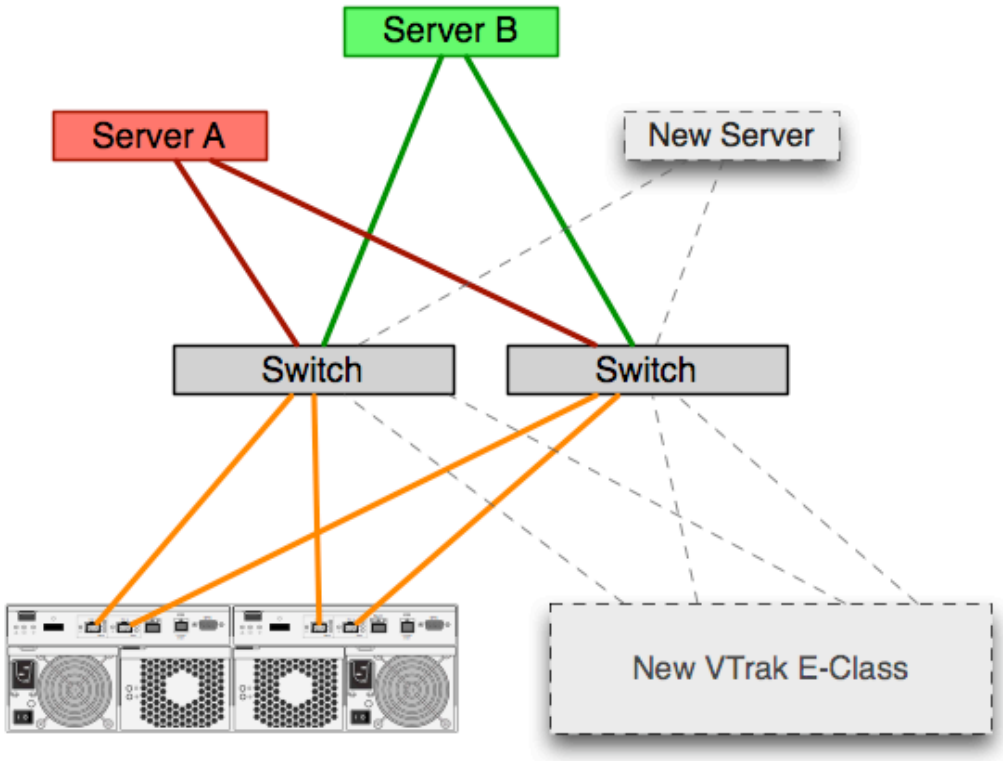

Figure 7: Database cluster with plan for future scale out

File Server Storage Nodes

To offer mission critical services at competitive prices, service providers must take every possible cost advantage when selecting hardware for the SNs. The use of directly attached fail-over clusters offers a mature server topology that is both easy to troubleshoot and familiar to system administrators. Widely distributed third party software, allows for multiple RAID volumes to be presented as a single volume to the ACs (figure 8).

When the customer logs in to the service for the first time, the AC creates the customer's folder on the file server SN, and then creates two files. The first file is a small catalog file that tracks changes to the customer's online data volume. The second file is a data blob that contains all of the user's data. When changes are made or new files are added, the AC updates the customer record indicating success or failure of the transaction. When the customer exceeds their initial provisioned space upgrades can be offered utilizing any available SN .

To retrieve a file, the user starts the applet. Once the session has been authenticated, the user can access the online data as if it was a local hard drive

Because the customer's transactions are stored in a catalog file, the cloud service provider can offer additional services like point-in-time snapshots to the end user data at any time.

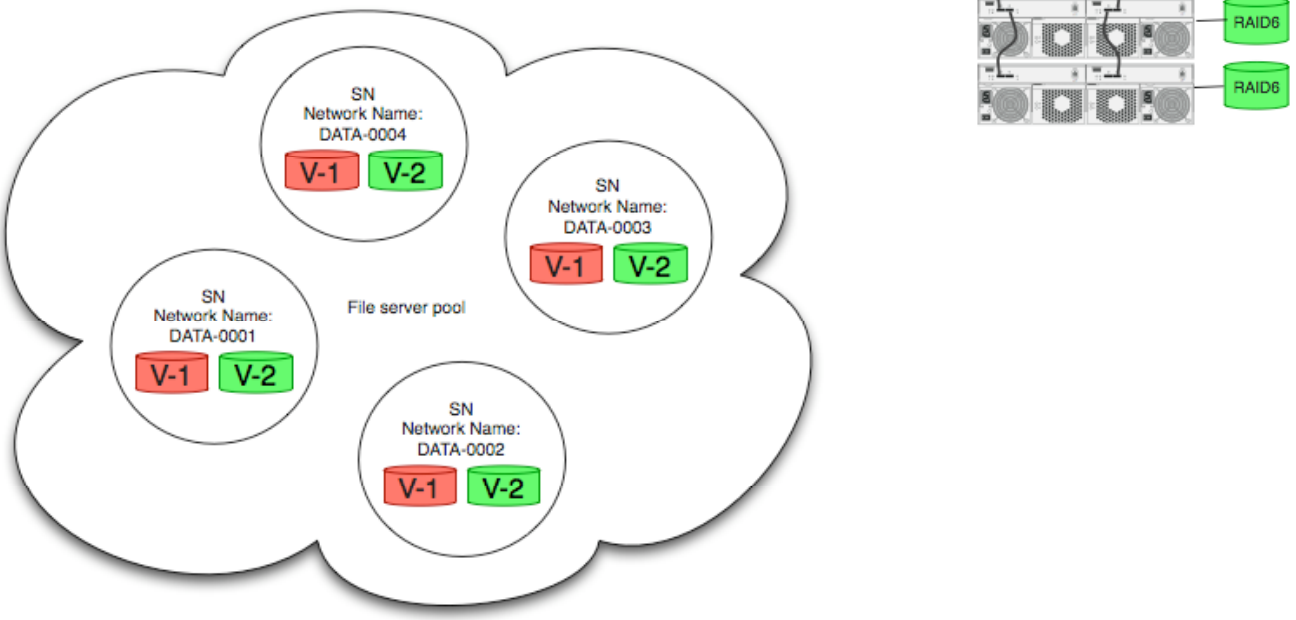

Figure 8: File Server Storage Node detail, storage shares as they appear to ACs

Mirrored Hardware

To the cloud storage provider, customer data is primary data regardless of whether it is the customer's second or third copy. The data must always be available in an instant. This means fully redundant mission-critical hardware must be specified across the whole data center (figure 9).

As in any mission critical application, a second backup copy of the data should be stored on separate hardware. In our example, SN servers C and D continuously mirror SN servers A and B. This provides both stripe based RAID and complete physical hardware duplication. When a customer wants to retrieve a file, both locations of the files are available to be read through the same namespace. An additional remote, asynchronous copy could also be made once the service grows to multiple geographic locations

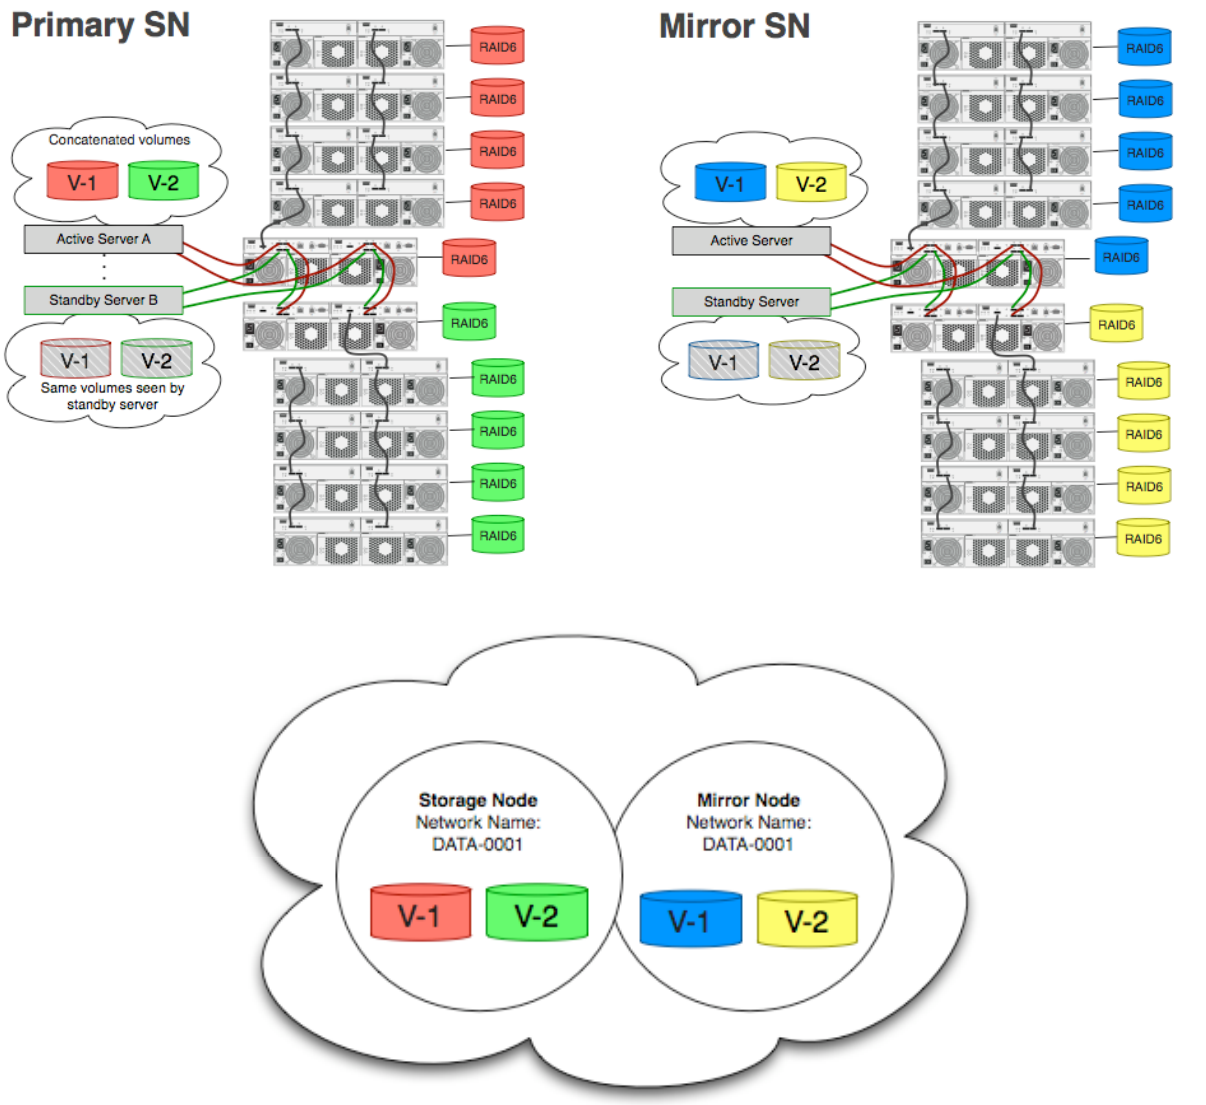

Figure 9: Mirrored Storage Nodes, redundant data under single namespace

Storage in the Data Center

Now that we have an overview of our planned cloud architecture, let us review several questions that must be considered before deployment:

Q: Is there adequate fault tolerance and redundancy implemented for the service levels that are offered?

A: VTrak E-Class features a fully redundant design to provide maximum uptime. That said, data availability is driven by many factors that can be external to the intrinsic reliability of RAID storage. Multiple availability strategies should be employed.

Q: Are planned operating systems and applications fully tested with the hardware?

A: Promise has an extensive Approved Vendors List (AVL). Promise works with our partners to make sure our products interoperate with all the leading industry suppliers. Promise has even made major hardware investments to duplicate key customer configurations in our lab.

Q: Does the supplier have technical training to assure a solid deployment?

A: Promise currently offers two levels of training for the VTrak E-Class. Level-1 training is general product knowledge about the VTrak E-Class RAID subsystem, provided free to the public as a video download. Level-2 training is provided quarterly for a fee and provides a deeper study into VTrak deployment.

Q: Do the planned standard operating procedures match those recommended by the suppliers?

A: Promise offers best-practice implementation information through the support organization. Before a major deployment of VTrak storage, it is highly recommended that the user attend a VTrak Level-2 training class as well as consult our support resources.

Q: Do the specified technology suppliers have similar goals?

A: Promise is focused on high-availability storage. We are here to keep your data online longer.

Q: Are reliable sales and support contacts in place with all of the technology suppliers?

A: Promise offers 24/7 phone and email technical support as well as pre- and post-sales contacts to help our customers succeed. Give us a call!

Conclusion

VTrak E-Class offers a single flexible storage platform upon which to build a cloud-based business. The ability to configure VTrak E-Class for high performance with random read and write applications using SAS drives as well as scale the system to massive capacities using SATA drives enables companies to use the same platform across functional tiers with compelling operational cost savings over competing solutions. VTrak E-Class can even provide a tiered solution under a single RAID head allowing providers to scale from the smallest footprint to extremely large installation.

Promise service and support staff is available 24/7 to assist customers with anything from component replacement to best-practice techniques to keep your data available.

PROMISE TECHNOLOGY, INC.

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

:: Promise Italy
Rome, Italy
0039-06-3671-2626
sales-it@eu.promise.com

:: Promise UK
Workingham, England
44-0-870-112-59-77
sales@eu.promise.co)

:: Promise Taiwan
Hsin-Chu, Taiwan
886-3-578-2395
sales@tw.promise.com

:: Promise China
Beijing, China
8610-8857-8085 or 8095
sales-china@promise.com

:: Promise Japan
Tokyo, Japan
81-03-6801-8063
sales@jp.promise.com

:: Promise ANZ
Sydney, Australia
61-2-8065-2055
sales@anz.promise.com

To learn more, visit www.promise.com

I