TECHNICAL BRIEF

VSM® and SVA component redundancy

Virtual Storage Manager® (VSM®) system and Shared Virtual Array (SVA)

APRIL 2004

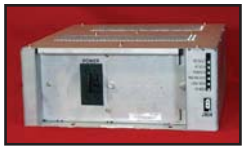

Figure 1.

1 ABSTRACT

In a world of networked computers, workgroups, data sharing and file servers, the very notion that a company's data should be at the mercy of a single component isn't an option. Access to data has to be completely reliable and that doesn't mean a set of mirror-image storage systems within one data complex. That's cost prohibitive. What it does mean is that the data storage system must be designed to be fault tolerant. A component may fail, but data processing MUST continue! The V-Series Shared Virtual Array® (SVA™) disk systems and Virtual Storage Manager® (VSM®) system were engineered with reliability in mind and have built-in redundancy in nearly every major component that is a potential point of failure that could halt processing.

Figure 2.

2 POWER SUPPLY REDUNDANCY

There are redundant power supplies, power distribution units and backup batteries in the VSM and SVA units. A failure of any single power component should not bring the unit down. All of these power components normally operate at less than half of their rated capacity.

Operating these components in such a fashion greatly improves their life expectancy and allows full power to flow to all active components even if half of the power distribution system becomes unavailable. The various power components are discussed in the following sections. All of these components are hot-pluggable — power does not have to be removed from the data storage subsystem to replace any of these parts. If one of these components fails, processing continues.

2.1 THE AC POWER DISTRIBUTION UNIT

There are two AC power distribution units in a VSM or SVA. Each is attached to its own line cord. Typically, each line cord is then attached to a separate customer power circuit so if a circuit loses power, only one AC distribution unit will shut down. Partial loss of power will not bring down a VSM or SVA. (**Figure 1**).

Figure 3.

2.2 THE LOGIC POWER SUPPLIES

There are two logic power supplies in a VSM or SVA. These provide power to the logic cards. If either one of the power supply fails, the load is assumed by the other. Either supply is capable of supplying power to the logic cards for as long as necessary until the failing supply is replaced. (**Figure 2**).

2.3 THE ARRAY POWER SUPPLIES

There are two disk array power supplies for each array, for a total of eight in a fully configured subsystem. In the event of a failure, either supply of a disk array is capable of supplying the necessary power to the disk array for as long as required until the failing supply is replaced. (**Figure 3**).

Figure 4.

2.4 THE BACKUP BATTERIES

In the event of loss of power to both customer power sources, a battery backup system protects NVS (Non-Volatile Storage) for at least 72 hours (assuming the batteries are fully charged). A primary battery provides at least 72 hours of DC power. A second, redundant battery offers at least 72 hours of protection should the primary battery fail. These two batteries operate in parallel for a total of at least 144 hours (assuming both batteries are fully charged). (**Figure 4**).

3 LOGIC REDUNDANCY

The VSM and SVA control unit features a high degree of modularity and component redundancy. By design, if a logic card fails, processing continues.

3.1 STORAGE CLUSTERS

Each of the two storage clusters in the control unit is comprised of eight independent storage paths. Any single failure or service action on a control bus impacts only 25 percent of available path resources.

3.2 CACHE

Cache in the subsystem is modular. If a cache card fails, the card is fenced off and the remainder of cache continues to operate. The failed card can be removed and replaced in the online environment.

3.3 NON-VOLATILE STORAGE (NVS)

NVS is modular; if a card fails, the card is fenced off and the remainder of NVS remains operative. A card failure has minimal impact on system performance, and complete data availability is maintained. The failed NVS card can be removed and replaced in the online environment.

3.4 SHARED MEMORY

The shared memory in the control unit is duplexed; both parts of which are shared by multiple components of the system. The primary shared memory bus of one storage cluster is connected to the secondary shared memory bus of the other storage cluster. Redundancy is achieved by writing and reading simultaneously to the primary and secondary buses of each storage cluster.

3.5 SHARED VIRTUAL ARRAY SUPPORT PROCESSOR

There are two support processors (ISP cards) in a control unit. One ISP card performs the maintenance and support functions while the other is on standby. In the event of a failure in the active ISP card, the standby ISP card assumes control of the maintenance and support functions.

In the event of a component failure, it is the active ISP card that will initiate a "call home" to StorageTek, reporting what failed and making error logs available to support personnel. Service personnel can then be dispatched to the customer's site with a replacement part — possibly before the customer even realizes that anything happened.

Note **That "call home" connection, used for both the trouble report and as a service connection to StorageTek, cannot directly access or transmit usable customer data to ANYONE.**

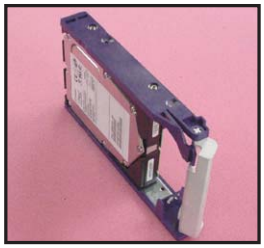

Figure 5.

4 DISK DRIVE REDUNDANCY

Hardware redundancy is built into the disk arrays as well. Each array has a "hot spare" that is automatically put into service if a drive should fail. The subsystem will logically replace the failing drive with the spare and begin rebuilding the data that was on the failing drive. One drive failing is transparent to the end user. Should a second drive fail before the first one can be replaced or the data is completely rebuilt, the subsystem is capable of handling double bit error correction as well. Once again, should a drive fail, processing continues. (**Figure 5**).

Note **Spare drives are "global" and can be used in any array.**

5 PATH REDUNDANCY

5.1 THE FRONT END PATHS

The architecture of the VSM and SVA provides multiple paths for a host to connect to the data storage system. In the event of an interface card failure, there are other paths available. The alternate path should be connected to the other cluster so, in the event of a cluster failure, there are still paths to the back end data storage.

5.2 THE BACK END PATHS

There is a multipath arrangement for access to the disk drives of the VSM and SVA. In the event of a disruption to one path, there is always another path available into the back end.

6 CONFIGURATION REDUNDANCY

The subsystem configuration is stored on two redundant disk drives. If one fails, the other has a complete picture of the configuration. These drives are not array drives and never contain customer data.

7 COOLING SYSTEM

There are cooling fan assemblies located thoughout the VSM and SVA. All fans have sensors, which tell the ISP card if a fan slows down or stops turning. If any fan fails, the rest maintain sufficient cooling.

8 THOSE NONREDUNDANT PARTS

A quick inspection of the subsystem reveals a few components for which there is no obvious duplicate part — power control panel, the local operator panel and the motherboards.

8.1 THE POWER CONTROL PANEL

There is only one power control panel, and it could be a single point of failure in the subsystem. However, the power control panel consists only of switches and indicator lights (light emitting diodes with an extremely long life expectancy). The failure of an indicator light will not impact data availability. The failure of a switch could bring the subsystem down, but it should be noted that switches typically only fail when they are being moved from one state to another. While the subsystem is operating, it is highly unlikely that a switch is going to fail without some human intervention. (**Figure 6**).

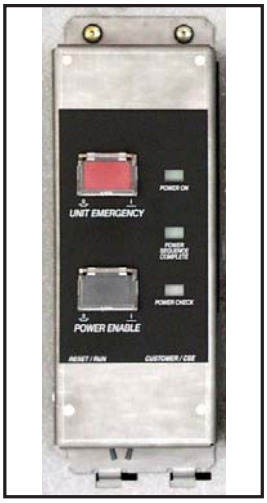

Figure 6.

8.2 THE LOCAL OPERATOR PANEL

There is only one local operator panel. However, if a customer installs the SVA™ Administrator software when from the system console instead of the subsystem's local operator panel. So, in effect, there is a second operator panel — it's just not on the VSM or SVA unit! An operator panel failure during operation will not affect data availability. If the local operator panel fails when the subsystem is powered up, just use the SVA Administrator software to see if the devices are available, and . . . processing continues. (**Figure 7**).

Figure 7.

8.3 THE MOTHERBOARDS

The motherboards of the VSM and the SVA do represent a potential single point of failure. Although without some sort of human intervention, failure of these components is extremely rare and unlikely.

8.4 THE BUSES

The SVA and VSM products use bus technology in many areas. Buses are used for power distribution as well as for communication between many of the processing components. While it's extremely remote, the possibility does exist for a component to fail in such a manner that it disables the entire bus. If such a rare event were to happen, then it is possible that such a failure could cause a disruptive outage.

9 CONCLUSION

In an age when data accessibility is an absolute must, StorageTek has engineered the Virtual Storage Manager system and the V-Series Shared Virtual Array disk system so that while an individual component may fail, the failure is transparent to the end user. Depending on the nature of the failure and the subsystem configuration, the customer's system operator may be notified of a failure by a Service Information Message (SIM), but that SIM is never seen by the end user. Redundant components allow the failing component to be replaced without the end user ever knowing anything happened.

The antecedent of the current model of the VSM and SVA was StorageTek's original Iceberg® disk storage subsystem — a product known for its reliability. The various evolutions of that product, up to the current ones, have steadily reduced the number of parts, further reducing the probability of failure and thus increasing the mean time between failures of the storage system's component modules.

A reduction in the number of parts and the continued use of redundant components means that the end user of the data storage system is even less likely than ever to experience a data access interruption and…processing continues.

ABOUT STORAGETEK®

Storage Technology Corporation (NYSE: STK), a $2 billion worldwide company with headquarters in Louisville, CO, has been delivering a broad range of storage management solutions designed for IT professionals for over 30 years. StorageTek offers solutions that are easy to manage, integrate well with existing infrastructures and allow universal access to data across servers, media types and storage networks. StorageTek's practical and safe storage solutions for tape automation, disk storage systems and storage integration, coupled with a global services network, provide IT professionals with confidence and know-how to manage their entire storage management ecosystem today and in the future.

StorageTek products are available through a worldwide network. For more information, visit www.storagetek.com, or call 1.800.275.4785 or 01.303.673.2800.

WORLD HEADQUARTERS

Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
1.800.877.9220 or 01.303.673.5151

MS 9155 B 04/04