

Why was the HP StorageWorks XP12000 Disk Array able to take the bullet and keep running? white paper



Executive summary.....	2
Physical separation and redundancy.....	2
Physical design and manufacturing.....	3
Electrical isolation.....	4
Logical isolation.....	4
Full active controller redundant architecture.....	5
Summary.....	5
For more information.....	6

Executive summary

This white paper describes the features of the HP StorageWorks XP12000 Disk Array that enabled it to be shot and keep operating. National Technical System (NTS) administered the “bulletproof” test at its ballistics testing facility in Arkansas. The video of the shooting event is available at www.hp.com/go/storageworks/bulletproofxp. It is strongly recommended that the reader view both the video clip and the documentary.

The reasons that the XP12000 Disk Array continued operating fall into five areas:

- Physical separation and redundancy
- Physical design and manufacturing
- Electrical isolation
- Logical isolation
- Full active controller redundant architecture


Each of these areas is a typical feature of high-availability array design, but they are implemented with extraordinary rigor in the XP12000 Disk Array.

This white paper describes how the XP12000 Disk Array was able to sustain multiple catastrophic failures induced by an external force (the bullet) and continue delivering data services to the streaming video server. In addition, the array was returned to full redundant functionality using standard repair procedures. The array that HP displays today is the array that was shot and has been repaired so it now provides the same redundancy as any similar configured XP12000 Disk Array. In short, it could be shot again and perform as demonstrated in the video clip. The only difference would be a second entry and exit hole in the physical sheet metal.

Physical separation and redundancy

The XP12000 Disk Array has physically redundant hardware components that are located in separate areas in the array. The redundant circuit cards are located in front and back pairs when the array is viewed from the side. Each circuit card plugged into a slot in the front of the array has a redundant “paired” circuit card plugged into the back of the array. The path of the bullet through the array penetrated all circuit cards plugged into the front of the array. None of the circuit cards in the back of the array were physically impacted.

XP12000 Center Cabinet



Physical design and manufacturing

The overall design and manufacturing of the XP12000 Disk Array is sturdy and high quality. The side covers as well as the internal sheet metal of the physical structure are well attached so that the physical structure maintained its integrity. This caused the physical damage to the frame structure to be limited to the holes created by the bullet. The structure was not weakened.

In addition, the physical circuit card assemblies are designed and manufactured so that they are very solid. Physical damage to the assemblies was limited to a hole being punched in the card and components being knocked off the assemblies as the bullet hit them. Heat sinks, surface mount chips, memory modules, and memory module connectors mounted on assemblies were removed from the array separate from the removal of the circuit assemblies they were originally attached to. The damage done by these loose-flying objects was contained in the array due to the physical design. The only impact of the loose components physically was that removal of the circuit board assemblies required some "jiggling."

The result is that after the shot the physical structure was intact and the physical damage to electrical parts was limited to the circuit board assemblies.

Electrical isolation

The electrical power circuits are isolated in the XP12000 Disk Array so that power issues such as short circuits that can be caused when a bullet penetrates all layers of several circuit card assemblies are isolated from impacting other parts of the array. In typical situations, this isolation prevents such issues from interfering with other circuit assemblies in the same physical cluster. In the case of shooting the array, all circuit assemblies in the physical cluster were damaged and this isolation prevented any interference with the other physical cluster, allowing continuing function.

This is internal isolation and different from the multiple AC power inputs that all HP XP Disk Arrays have to provide increased robustness for input power disturbances and outages.

Without such isolation, the power system in the array could have failed in such a way that the power to the other physical cluster could have been impacted resulting in the array not being able to continue uninterrupted operation.

Logical isolation

The distributed microprocessor design provides logical isolation for the logical processes in the HP XP12000 Disk Array. However, there is an additional part of this logical isolation related to how the cache and shared memory (SM) in the array are used.

SM is mirrored in the array (one copy in each physical cluster). Any microprocessor can access either copy of SM for read operations. Both copies of SM must be available for write operations. When one side of SM has failed, such as what happened when the bullet penetrated the SM circuit assembly, the failure is detected by the array and appropriate logical actions are taken. Failed assemblies are blocked logically so that components will not attempt to use them. In the case of SM or cache memory (CM), the blocking of one complete copy results in the array entering a write-through mode to protect the data. In this mode, writes from hosts to the array are responded to as complete when the data has been written to the physical disk drives. In normal operation, the write is acknowledged to the host as soon as it has been written to both sides of the CM. Writes to CM are mirrored in normal operation. This mirroring results in there being two identical (but independent) copies of the written data in the array when the data has not been written to disk. In the case of read operations, the data on the disks matches the data in cache and CM is not mirrored. This is true whether the array is operating normally or in write-through mode. Write-through mode only impacts host write performance (to maintain data integrity). In write-through mode, the data is still written from the Fibre Channel host interface to the remaining good CM unit. Then data is written from that cache unit to the disks. Upon successful completion of the write all the way through to the disk, a write completion message is returned the host. Thus other cache-based write techniques such as local array copies (HP StorageWorks Business Copy) and remote array copies (HP StorageWorks Continuous Access) continue to function correctly, even in the event of a cache module malfunction.

Write-through mode is entered if a CM platform circuit assembly fails completely or if a SM platform assembly fails completely. In the case of shooting the XP12000 Disk Array, both the SM and CM platform circuit assemblies failed completely and the array entered write-through mode.

Since the primary workload on the array during the test was serving streaming video (a read workload), there was no noticeable change in performance when the bullet went through the array and the parts failed.

Full active controller redundant architecture

The XP12000 Disk Array is architected to be fully functional on all paths. For host connectivity, the host manages which paths are used and how. If the same logical device (LDEV) is being accessed by a host over two or more separate paths, the array will perform the I/O operations as directed. It is up to the host to manage accesses for data consistency (such as ensuring that reads by way of one path do not pass writes to the same record by another path). Hence, when there is a failure of one redundant path, the array does not perform a switch action. Rather the host detects the path failure and then uses the redundant paths. The array continues to access the LDEV. For open systems, the redundant paths are the LUN assignments to ports in the array. The same LDEV may have multiple LUN assignments in the array. Any path defined by a LUN assignment to the same LDEV can be used for regular host access. Hence, when a path fails due to a port failure (such as when shot by a bullet in one cluster), other paths configured on ports in the other cluster are ready to go. There is no switch over in the array. Data is in the cache and available to any microprocessor in the array. The control information in SM is also available to any microprocessor. Hence, data service access by way of any operational path is not impeded when another path to the same data fails.

Summary

In the case of the XP12000 Disk Array continuing to function after being shot, it is the combination of all of the preceding features that resulted in continued operation of the streaming video application (video data and the video server program files were on the array) that was running on the connected server. It is also the combination of these features that enabled the array to be repaired to a normal operational state with full redundancy. HP could do it again to the same array and it would continue to operate again.

As a side note, as seen in the video, the .308-caliber bullet shattered into small pieces upon hitting the glass wall of the fish tank (backed by water) and stopped before exiting the tank. Such shattering is typical of a soft lead projectile when it has spread significantly. The spreading was caused by hitting the glass barrier and the fragmenting was caused by the viscosity of the water working over the large cross-section of the fragment caused by the expansion. The combination of the energy reduction due to penetrating the glass wall of the tank and the energy loss in the water resulted in the fragments stopping inside the tank.

For more information

- HP StorageWorks XP12000 Disk Array
<http://h18006.www1.hp.com/products/storageworks/xp12000/index.html>
- Bulletproof Video Page
www.hp.com/go/storageworks/bulletproofxp
- NTS Certificate
<ftp://ftp.compaq.com/pub/products/storageworks/xp12000/NTSImpactTestXP12000.pdf>

To learn more about National Technical Systems, visit the following links:

- NTS Camden Brochure
<http://www.ntscorp.com/pdf/locations/Camden.pdf>
- NTS Website
<http://www.ntscorp.com/>

© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP 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. HP shall not be liable for technical or editorial errors or omissions contained herein.