WINCHESTERSYSTEMS®
Purpose-Built Storage

High Availability Configuration

on

FlashNAS ZFS Systems

Table of Contents

FlashNAS ZFS Series

FlashNAS ZFS series systems deliver consolidated storage for application server and file server deployment. They are available in a variety of hardware configurations, including high availability active-active dual controllers for assured redundancy and fast failover. Power supplies are also redundant and energy-efficient. In all situations, FlashNAS ZFS products protect data and your continued ability to work and provide services to your customers. Additionally, controllers, power supplies, and cooling modules use a modular cable-free design that makes installation, maintenance, and upgrades simple and quick.

FlashNAS ZFS products are enhanced by the ZFS file system, which has sophisticated data corruption prevention and healing capabilities built in. You gain access to features such as unlimited snapshot, remote replication, and pool mirror. Powerful computing components and up to 1.5PB in storage via JBOD make FlashNAS ZFS series systems highly scalable and capable solutions for every enterprise and organization.

Active/Active redundant controller configuration:

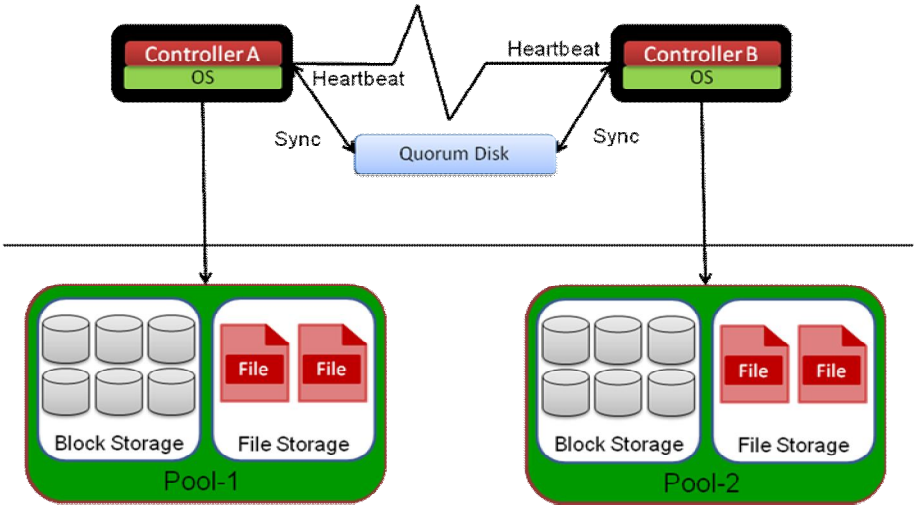

Figure 1: high availability architecture

The FlashNAS ZFS series features active/active controller architecture instead of active/standby, as shown by figure 1. Employing a so-called internal heartbeat mechanism, each controller is aware of the status of its twin controller. Both sync with the quorum disk to ascertain pool ownership.

This architecture nearly doubles performance and offers high availability protection. In real world applications such as iSCSI volume access or share folders, the FlashNAS ZFS series provides better data integrity and higher availability than those delivered by comparable products, as explained below.

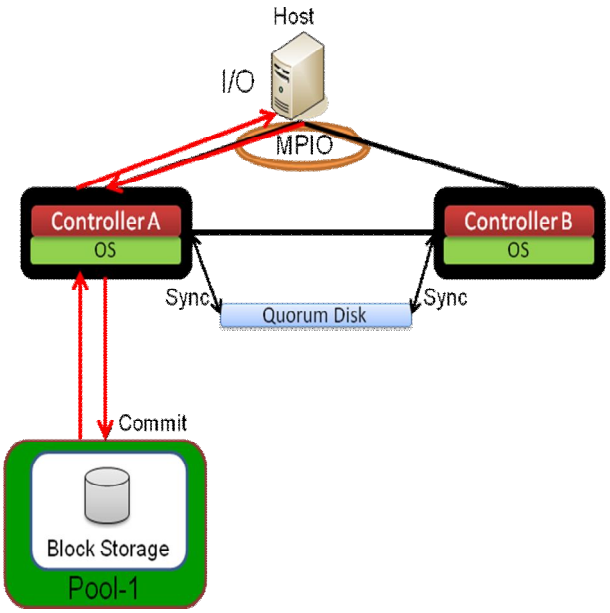

Figure 2: iSCSI volume data integrity

As figure 2 shows, the host receives a commit message when the I/O writes to block storage. The FlashNAS ZFS series uses write-through mode to protect iSCSI volume data.

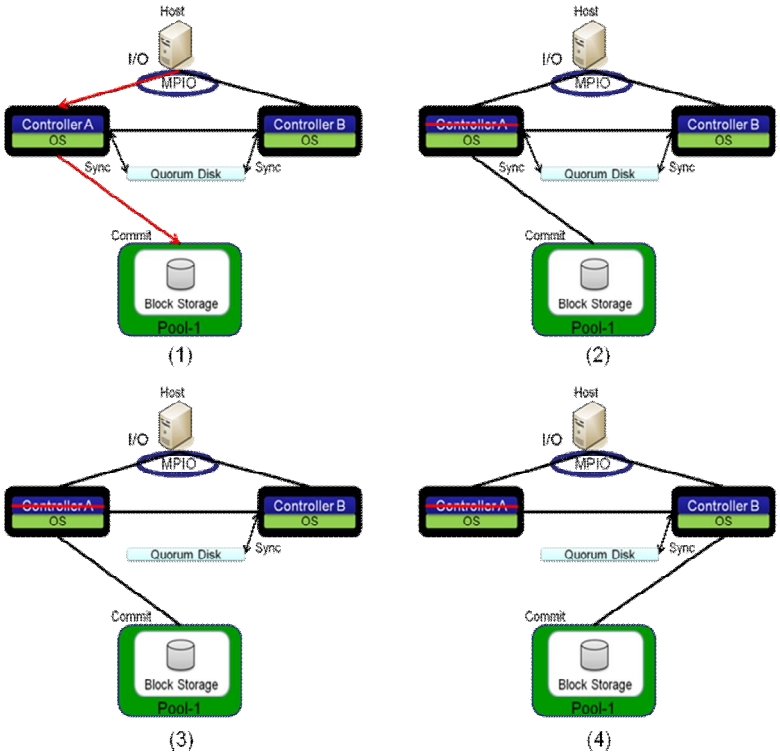

Figure 3 : iSCSI volume high availability

As figure 3-1 shows, the FlashNAS ZFS series leverages MPIO to provide connections from controller A and controller B for to the host. In figure 3-2, we unplug controller A to see how the FlashNAS ZFS series provides high availability. Figure 3-3 shows controller A cannot sync the quorum drive immediately, while figure 3-4 indicates that when controller B notices controller A has failed, controller B becomes the owner of Pool-1. I/O from the host passes through controller B rather than controller A in this case, ensuring continuous availability.

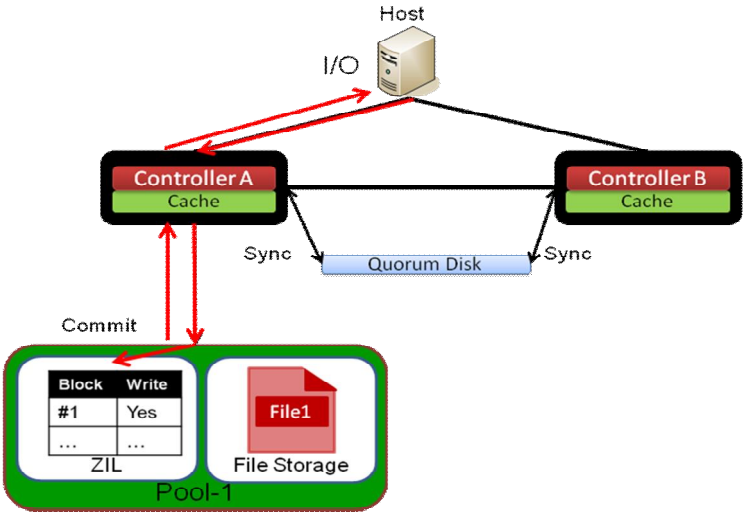

Figure 4: share-folder data integrity

The FlashNAS ZFS series leverages ZFS Intent Log (ZIL) to ensure share folder data integrity. ZIL records data and whether that data has been written onto disk or not.

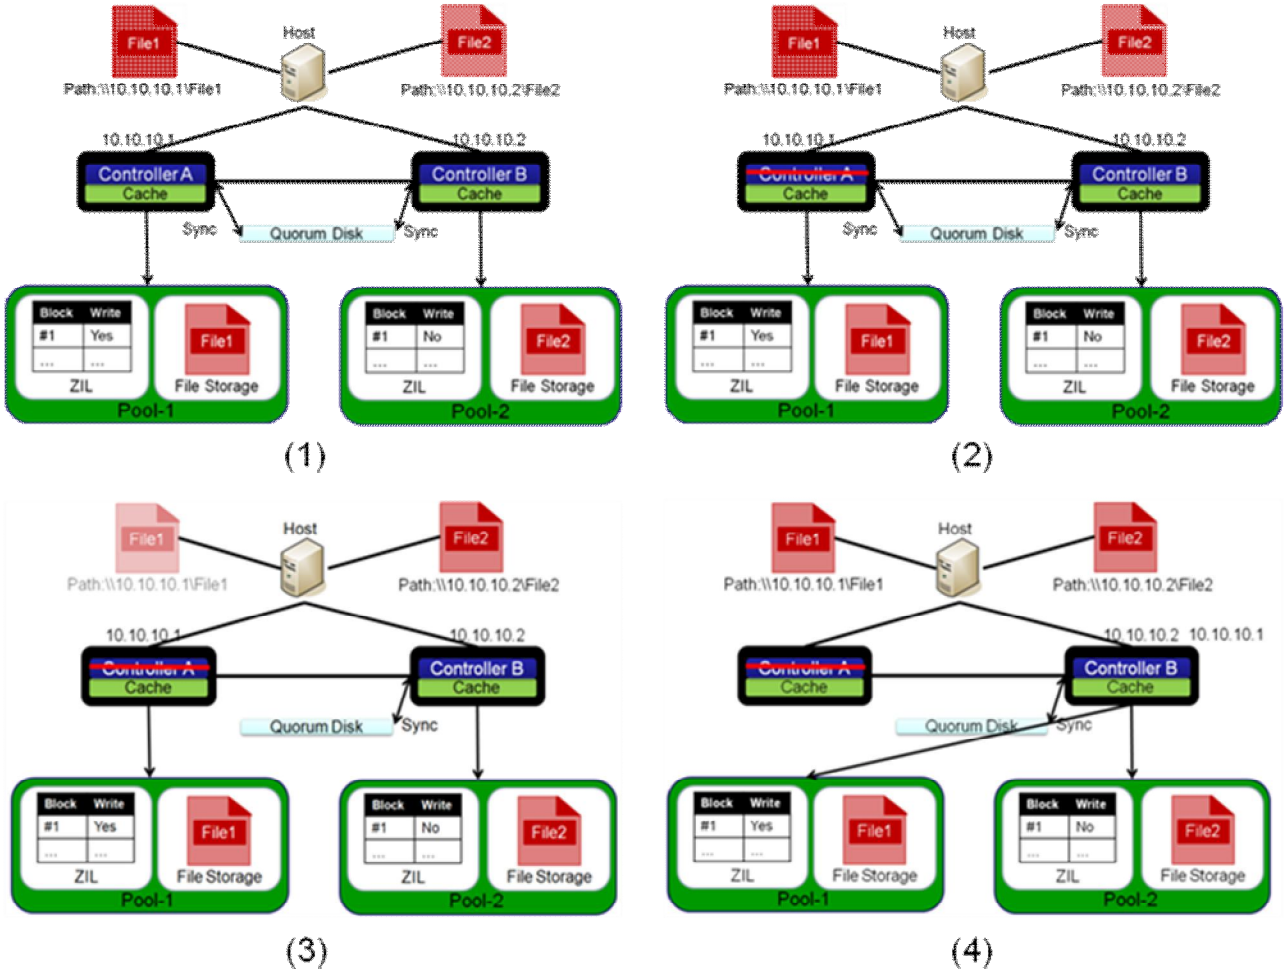

Figure 5: share folder high availability

Figure 5-1 shows how the FlashNAS ZFS series provides two share folders belonging to controller A and controller B, respectively. In figure 5-2, we unplug controller A to see how the FlashNAS ZFS series delivers high availability in more scenarios. As figure 5-3 shows, the first share folder path is \\10.10.10.1\File1. The share folder becomes lost initially once the IP disappears. Also, controller A cannot sync the quorum drive immediately. In figure 5-4, we see an IP that belongs to controller A executing failover to controller B. Once IP failover completes, \\10.10.10.1\File1 becomes accessible again.

Comprehensive GUI: with the detailed GUI that ships with the FlashNAS ZFS series, customers gain clear insight into all configurations.

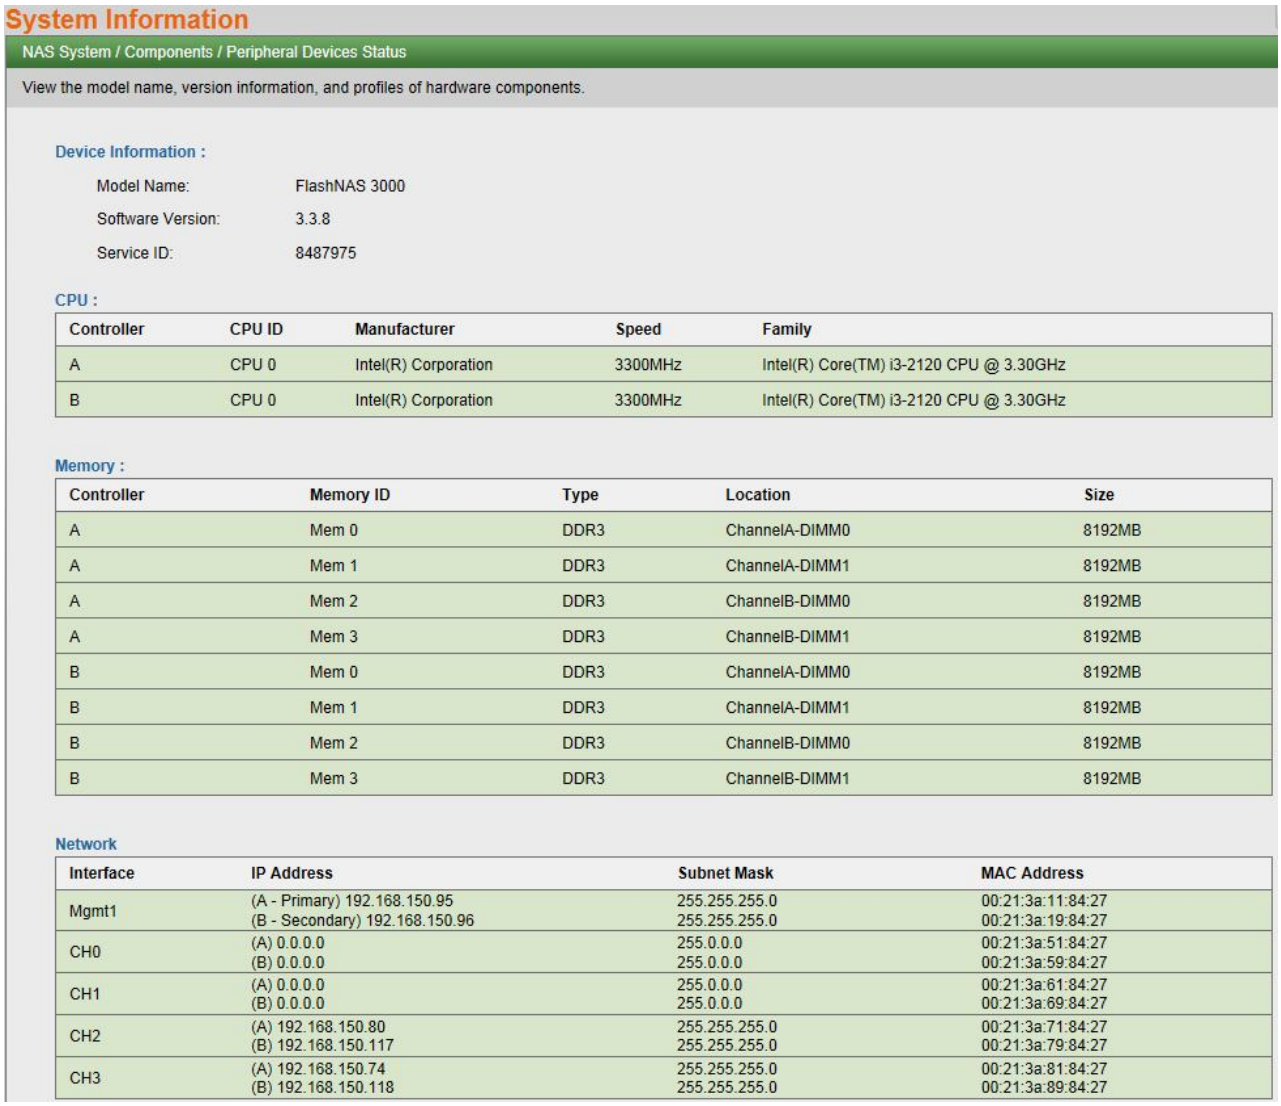

System Information

NAS System / Components / Peripheral Devices Status

View the model name, version information, and profiles of hardware components.

Device Information :

Model Name:	FlashNAS 3000
Software Version:	3.3.8
Service ID:	8487975

CPU :

Controller	CPU ID	Manufacturer	Speed	Family
A	CPU 0	Intel(R) Corporation	3300MHz	Intel(R) Core(TM) i3-2120 CPU @ 3.30GHz
B	CPU 0	Intel(R) Corporation	3300MHz	Intel(R) Core(TM) i3-2120 CPU @ 3.30GHz

Memory :

Controller	Memory ID	Type	Location	Size
A	Mem 0	DDR3	ChannelA-DIMM0	8192MB
A	Mem 1	DDR3	ChannelA-DIMM1	8192MB
A	Mem 2	DDR3	ChannelB-DIMM0	8192MB
A	Mem 3	DDR3	ChannelB-DIMM1	8192MB
B	Mem 0	DDR3	ChannelA-DIMM0	8192MB
B	Mem 1	DDR3	ChannelA-DIMM1	8192MB
B	Mem 2	DDR3	ChannelB-DIMM0	8192MB
B	Mem 3	DDR3	ChannelB-DIMM1	8192MB

Network

Interface	IP Address	Subnet Mask	MAC Address
Mgmt1	(A - Primary) 192.168.150.95 (B - Secondary) 192.168.150.96	255.255.255.0 255.255.255.0	00:21:3a:11:84:27 00:21:3a:19:84:27
CH0	(A) 0.0.0.0 (B) 0.0.0.0	255.0.0.0 255.0.0.0	00:21:3a:51:84:27 00:21:3a:59:84:27
CH1	(A) 0.0.0.0 (B) 0.0.0.0	255.0.0.0 255.0.0.0	00:21:3a:61:84:27 00:21:3a:69:84:27
CH2	(A) 192.168.150.80 (B) 192.168.150.117	255.255.255.0 255.255.255.0	00:21:3a:71:84:27 00:21:3a:79:84:27
CH3	(A) 192.168.150.74 (B) 192.168.150.118	255.255.255.0 255.255.255.0	00:21:3a:81:84:27 00:21:3a:89:84:27

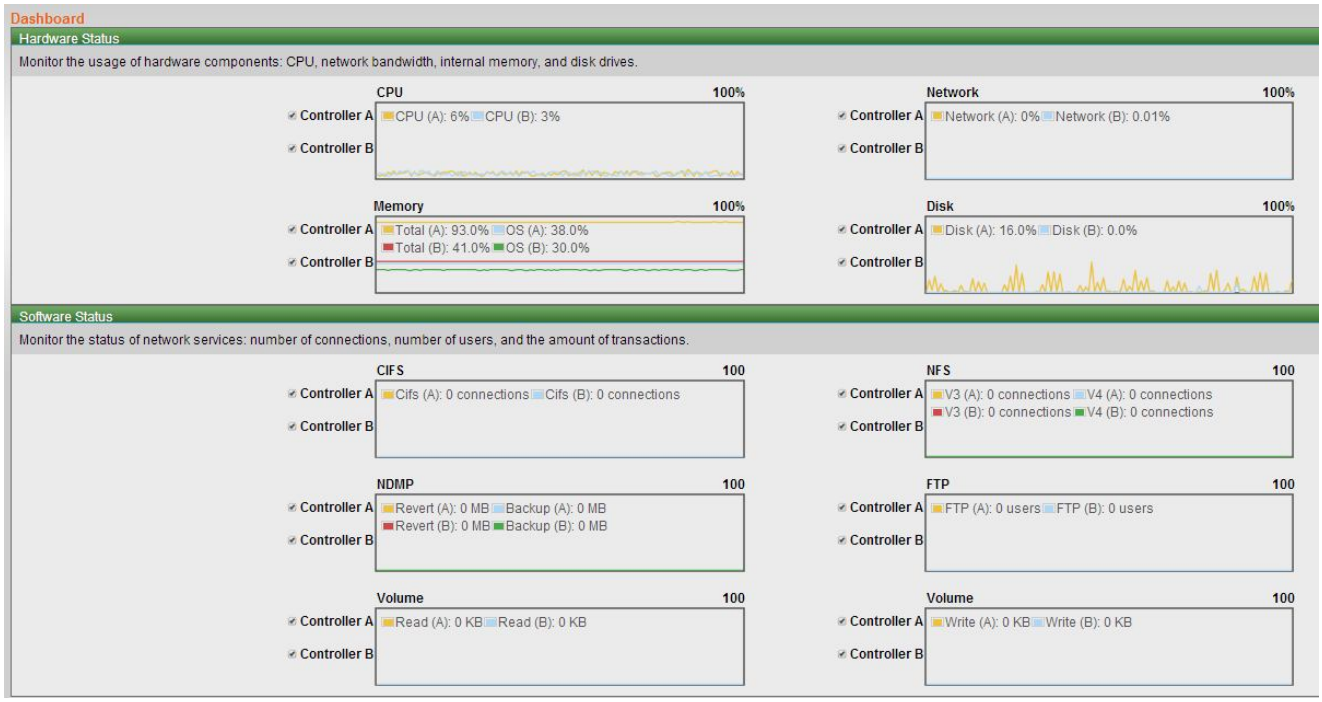
Dashboard
Hardware Status
Monitor the usage of hardware components: CPU, network bandwidth, internal memory, and disk drives.
CPU 100%
Controller A
CPU (A): 6% CPU (B): 3%
Controller B
Network 100%
Controller A
Network (A): 0% Network (B): 0.01%
Controller B
Memory 100%
Controller A
Total (A): 93.0% OS (A): 38.0%
Total (B): 41.0% OS (B): 30.0%
Controller B
Disk 100%
Controller A
Disk (A): 16.0% Disk (B): 0.0%
Controller B
Software Status
Monitor the status of network services: number of connections, number of users, and the amount of transactions.
CIFS 100
Controller A
Cifs (A): 0 connections Cifs (B): 0 connections
Controller B
NFS 100
Controller A
V3 (A): 0 connections V4 (A): 0 connections
V3 (B): 0 connections V4 (B): 0 connections
Controller B
NDMP 100
Controller A
Revert (A): 0 MB Backup (A): 0 MB
Revert (B): 0 MB Backup (B): 0 MB
Controller B
FTP 100
Controller A
FTP (A): 0 users FTP (B): 0 users
Controller B
Volume 100
Controller A
Read (A): 0 KB Read (B): 0 KB
Controller B
Volume 100
Controller A
Write (A): 0 KB Write (B): 0 KB
Controller B

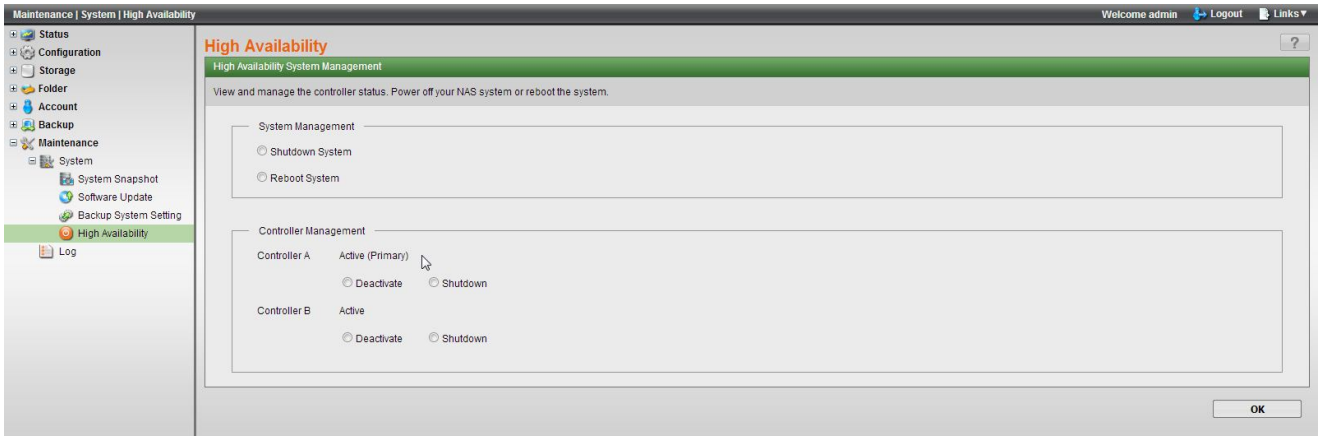
Maintenance | System | High Availability
Welcome admin
Logout
Links
Status
Configuration
Storage
Folder
Account
Backup
Maintenance
System
System Snapshot
Software Update
Backup System Setting
High Availability
Log
High Availability
High Availability System Management
View and manage the controller status. Power off your NAS system or reboot the system.
System Management
Shutdown System
Reboot System
Controller Management
Controller A
Active (Primary)
Deactivate
Shutdown
Controller B
Active
Deactivate
Shutdown
OK

Scenarios without I/O running

We use two methods to illustrate the high availability FlashNAS ZFS series products provide.

Method 1: ping management port IP

Step 1: ping management port IP (controller A): 172.24.110.32

Network

Interface	IP Address	Subnet Mask	MAC Address
Mgmt0	(A - Primary) 172.24.110.32 (B - Secondary) 172.24.110.58	255.255.254.0 255.255.254.0	0:D0:23:06:79:46 0:D0:23:0E:79:46
CH0	(A) 172.24.110.53 (B) 172.24.110.51	255.255.254.0 255.255.254.0	0:D0:23:36:79:46 0:D0:23:3E:79:46
CH1	(A) Offline (B) Offline	--- ---	0:D0:23:46:79:46 0:D0:23:4E:79:46
CH2	(A) Offline (B) Offline	--- ---	0:D0:23:56:79:46 0:D0:23:5E:79:46
CH3	(A) Offline (B) Offline	--- ---	0:D0:23:66:79:46 0:D0:23:6E:79:46

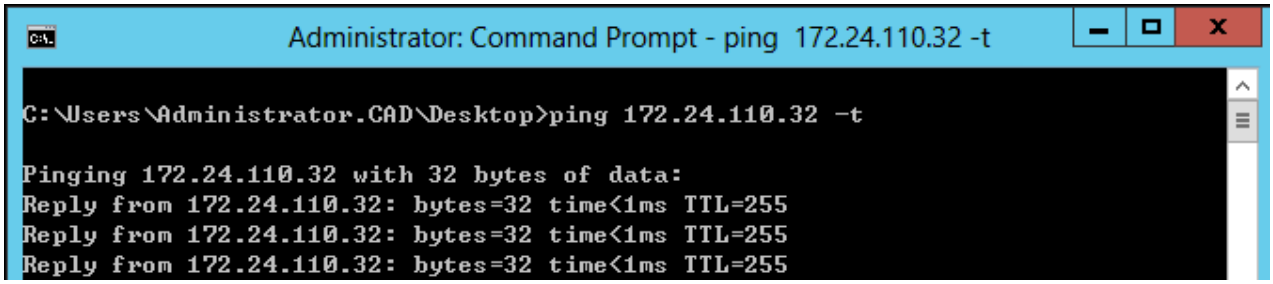

Step 2: unplug controller A

Step 3: check the echo package

```
Reply from 172.24.110.32: bytes=32 time=1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Reply from 172.24.110.86: Destination host unreachable.
Reply from 172.24.110.86: Destination host unreachable.
Reply from 172.24.110.32: bytes=32 time=2279ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
Reply from 172.24.110.32: bytes=32 time<1ms TTL=255
```

Only a few requests timed out, and clients can still receive echo packages from the same management IP.

Step 4: customers can access the management GUI from the same management port IP

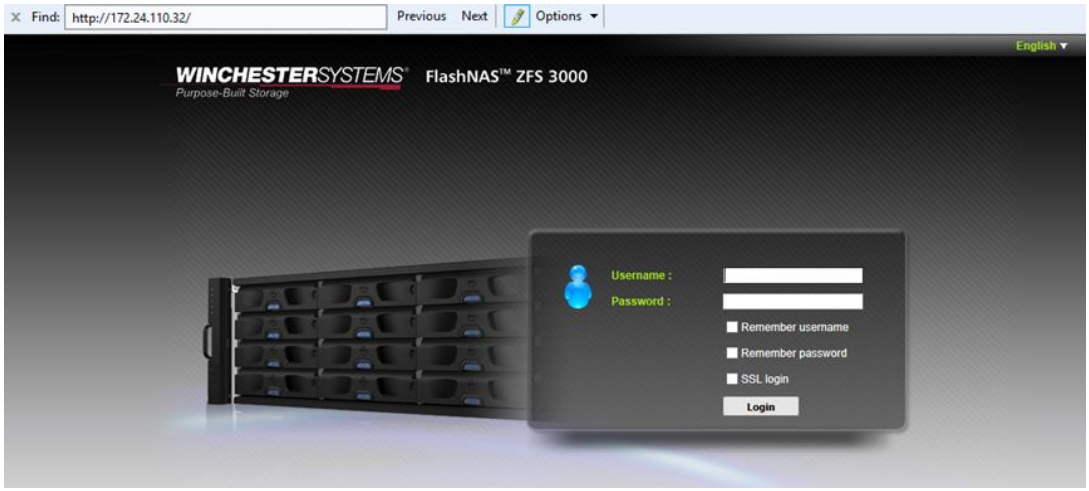

Method 2: FTP

Step 1: check Network Summary

Network

Interface	IP Address	Subnet Mask	MAC Address
Mgmt0	(A - Primary) 172.24.110.32 (B - Secondary) 172.24.110.58	255.255.254.0 255.255.254.0	0:D0:23:06:79:46 0:D0:23:0E:79:46
CH0	(A) 172.24.110.53 (B) 172.24.110.51	255.255.254.0 255.255.254.0	0:D0:23:36:79:46 0:D0:23:3E:79:46
CH1	(A) Offline (B) Offline	--- ---	0:D0:23:46:79:46 0:D0:23:4E:79:46
CH2	(A) Offline (B) Offline	--- ---	0:D0:23:56:79:46 0:D0:23:5E:79:46
CH3	(A) Offline (B) Offline	--- ---	0:D0:23:66:79:46 0:D0:23:6E:79:46

Step 2: use FTP client software to connect a share folder

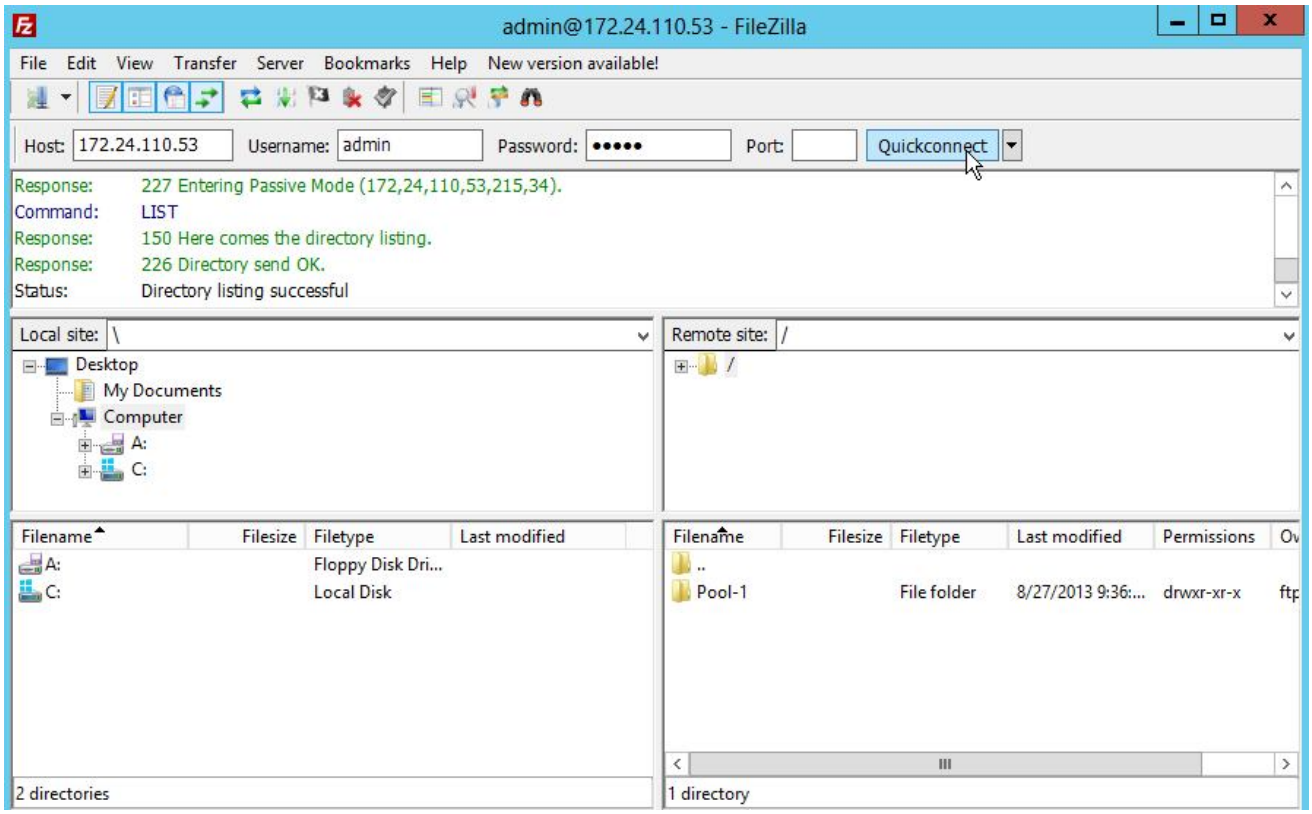

Step 3: prepare a file and upload to the share folder

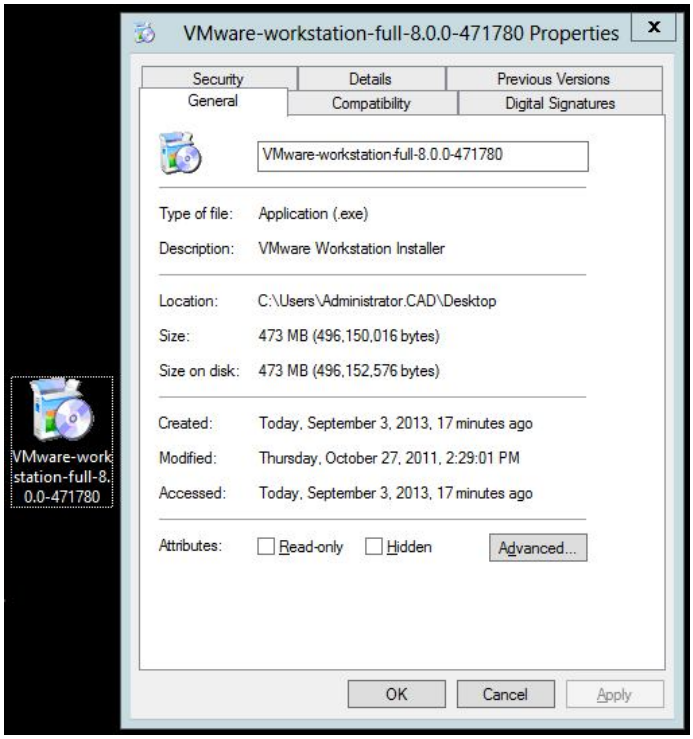

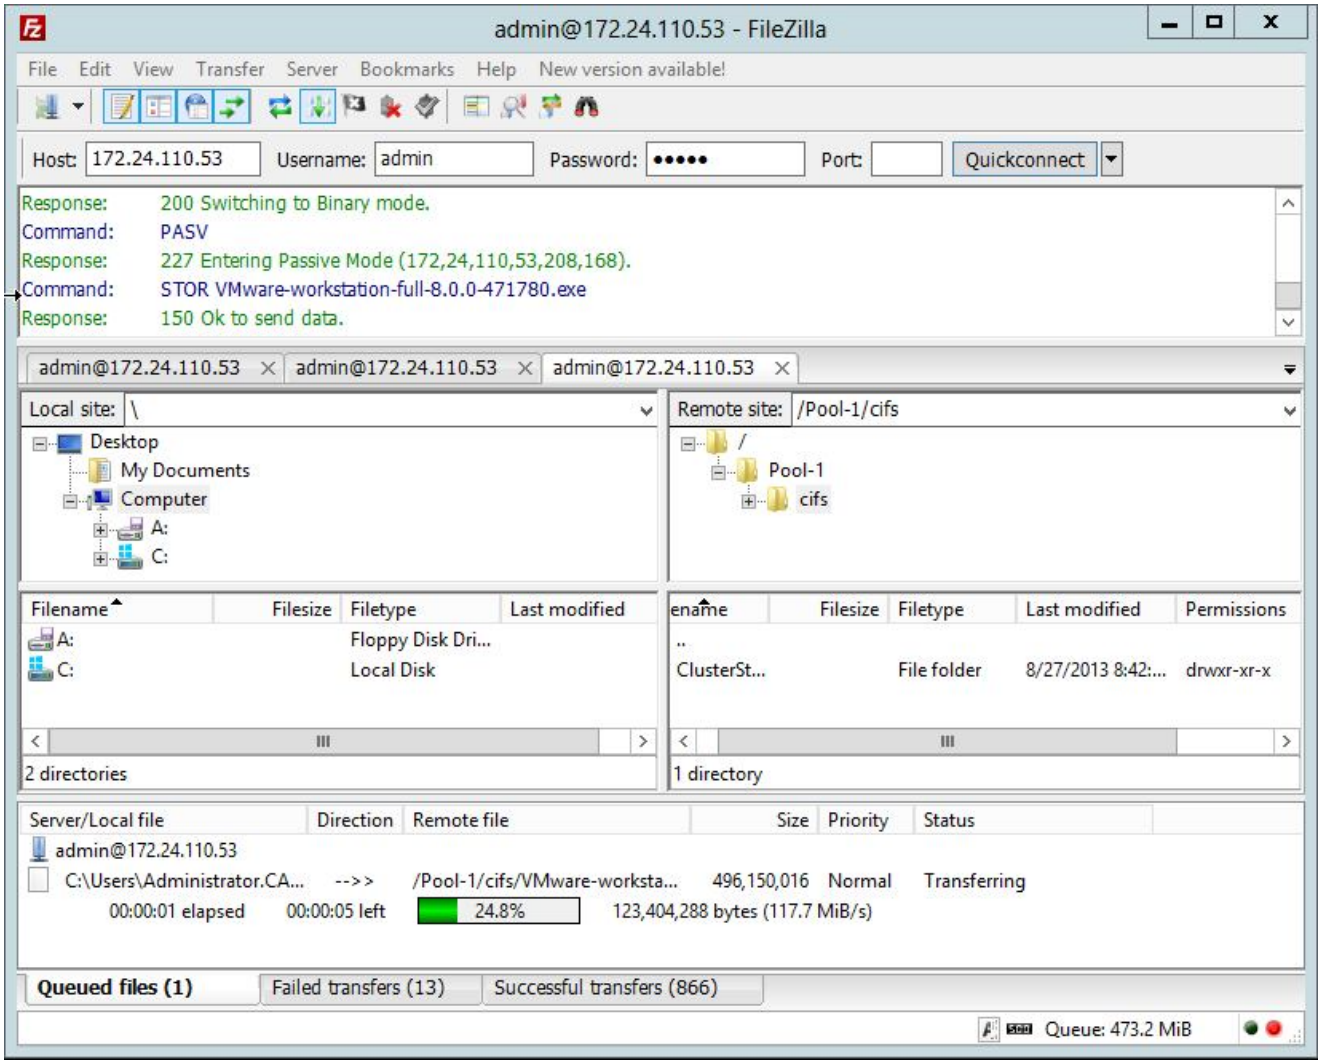

Step 4: unplug controller A

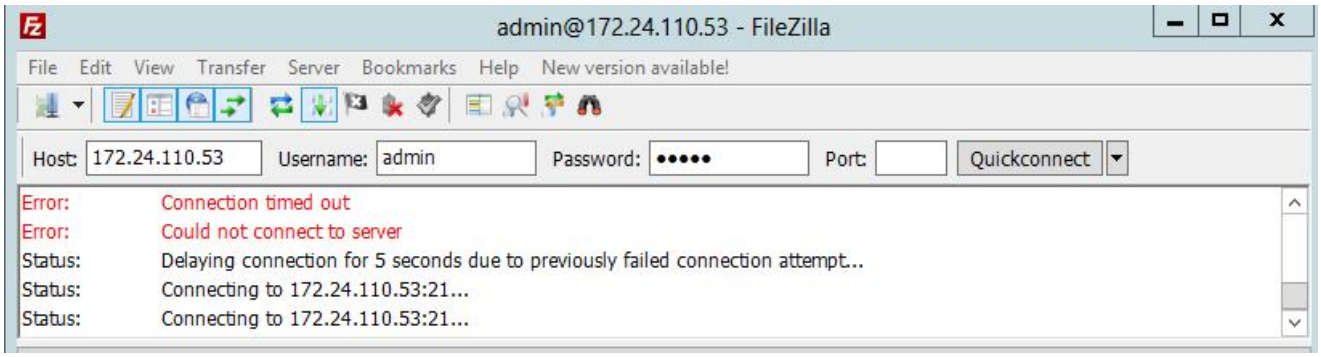

The FTP client software automatically reconnects to the share folder and continues previously started file transmission.

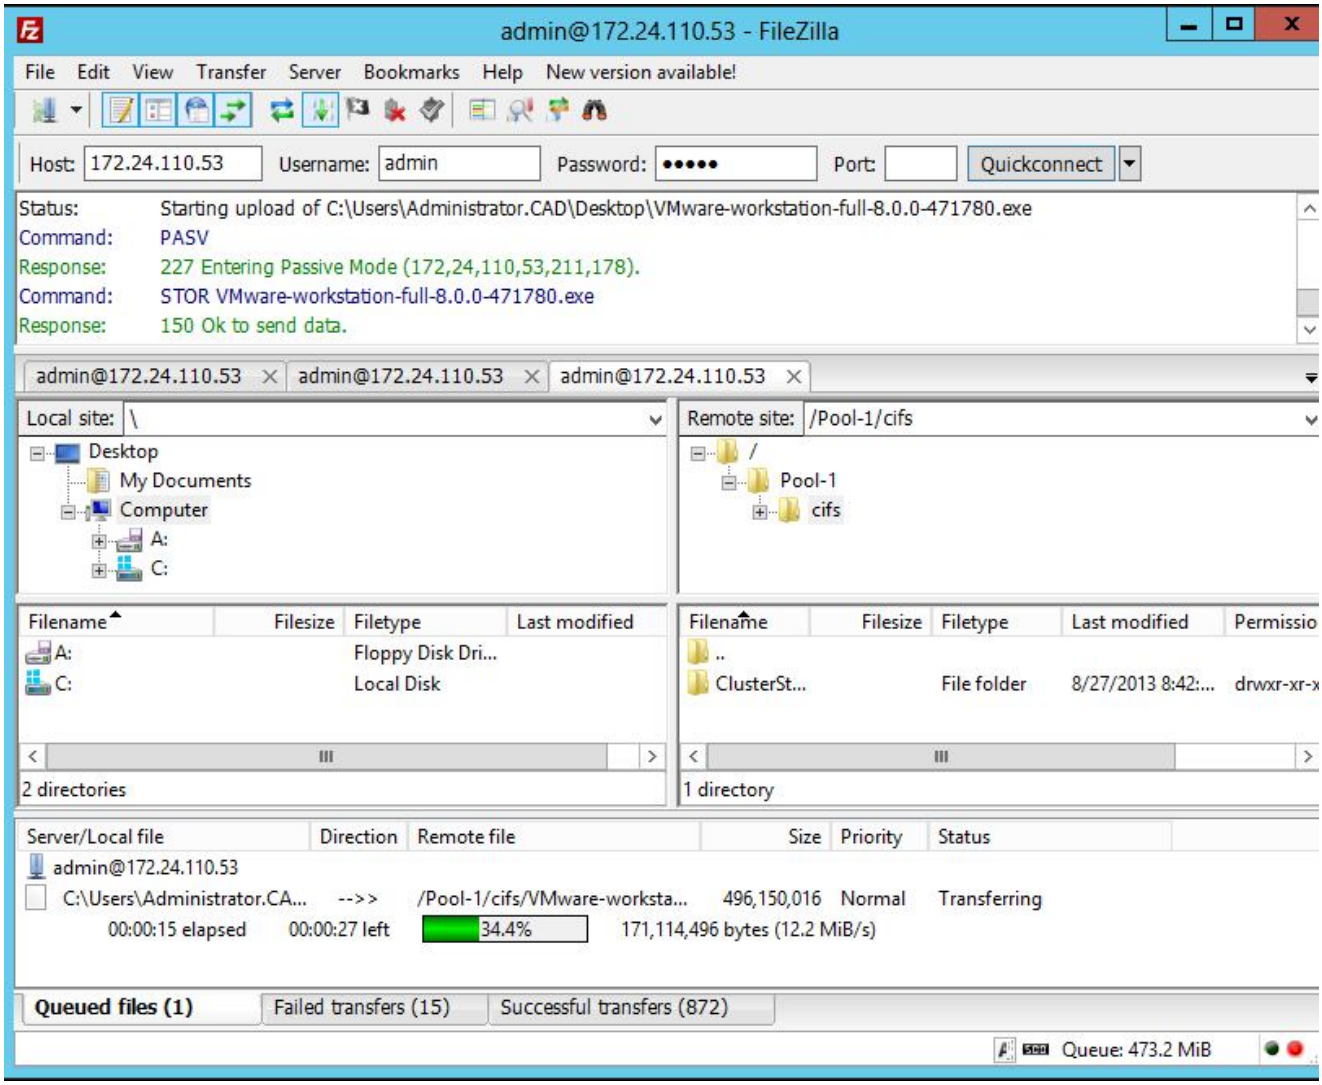

Scenarios with I/O running

We also test high availability under I/O running conditions to establish the length of time required to complete the failover process. Based on this test, failover takes less than 100 seconds even with I/O running. In other words, the FlashNAS ZFS series offers high availability during “live” circumstances, with data transiting across the network as failover occurs. High availability is not limited to “off” or passive conditions (no I/O running).

Test Item	Configuration	Time (sec)	Remark
CIFS Share Folder	A: 4(CIFS) B: 4(CIFS)	35	1. Create 4 CIFS share folders for controller A and create another 4 CIFS share folders for controller B. 2. Running IOMeter on all CIFS share folders. 3. During the I/O, unplug controller A directly.
	A: 8(CIFS) B: 8(CIFS)	35	1. Create 8 CIFS share folders for controller A and create another 8 CIFS share folders for controller B. 2. Running IOMeter on all CIFS share folders. 3. During the I/O, unplug controller A directly.
iSCSI Volume	A: 4(iSCSI) B: 4(iSCSI)	92	1. Create 4 iSCSI volumes for controller A and create another 4 iSCSI volumes for controller B. 2. Running IOMeter on all iSCSI volumes. 3. During the I/O, unplug controller A directly.
	A: 8(iSCSI) B: 8(iSCSI)	97	1. Create 8 iSCSI volumes for controller A and create another 8 iSCSI volumes for controller B. 2. Running IOMeter on all iSCSI volumes. 3. During the I/O, unplug controller A directly.
CIFS Share Folder + iSCSI Volume	A: 4(iSCSI), 4(CIFS) B: 4(iSCSI), 4(CIFS)	92	1. Create 4 CIFS share folders and 4 iSCSI volume for controller A and create another 4 CIFS share folders and another 4 iSCSI volumes for controller B. 2. Running IOMeter on all CIFS share folders and all iSCSI volumes. 3. During the I/O, unplug controller A directly.
	A: 8(iSCSI), 8(CIFS) B: 8(iSCSI), 8(CIFS)	97	1. Create 8 CIFS share folders and 8 iSCSI volume for controller A and create another 8 CIFS share folders and another 8 iSCSI volumes for controller B. 2. Running IOMeter on all CIFS share folders and all iSCSI volumes. 3. During the I/O, unplug controller A directly.

Conclusions

In the above scenarios, the FlashNAS ZFS series shows that it provides customers the benefits of a redundant design. When encountering single controller failure, users do not need to worry about data loss or lack of access. FlashNAS ZFS series systems carry out failover of all IP and pool ownership information to the other controller in order to ensure business continuity and uninterrupted productivity. In terms of data safety and service integrity, the FlashNAS ZFS series provides a greatly enhanced environment to customers.