SDC
STORAGE DEVELOPER CONFERENCE
SNIA ■ SANTA CLARA, 2014

Continuous Availability: A Scenario Validation Approach

Aniket Malatpure

Ningyu He

Microsoft

Agenda

- Continuous Availability
- Cluster-in-a-Box
- Test Scenarios
- Validation Toolkit
- Key Takeaways

Continuous Availability (CA)

- **Continuous Availability**
  - Transparent failover of application data storage
  - Application sees contained IO delay
- **Value propositions**
  - Servicing without downtime
  - Reliable, low-cost, file-based storage
- **Deployment Scenarios**
  - Server application storage platform
  - File Server consolidation
  - Virtual Desktop Infrastructure
- **Deployment Variations**
  - Multiple customer segments (small businesses to enterprises & hosted clouds)
  - Multiple networking configurations (Ethernet, Infiniband, RDMA etc.)
  - Multiple storage options (JBOD, RAID, SAS, SATA, FC, IP SAN etc.)

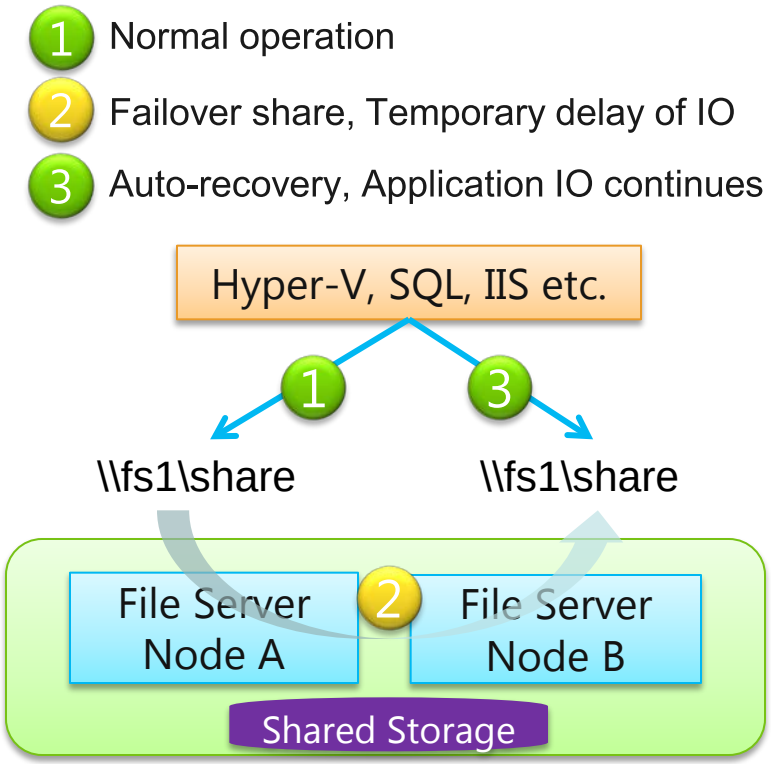

Cluster-in-a-Box (CiB)

- Volume platform for Continuous Availability
- Design Considerations
  - At least one node and storage always available, despite failure or replacement of any component
  - Dual power domains
  - Internal interconnect between nodes, controllers
  - Flexible PCIe slot for LAN options
  - External SAS ports for JBOD expansion
  - Office-level power and acoustics for entry-level NAS

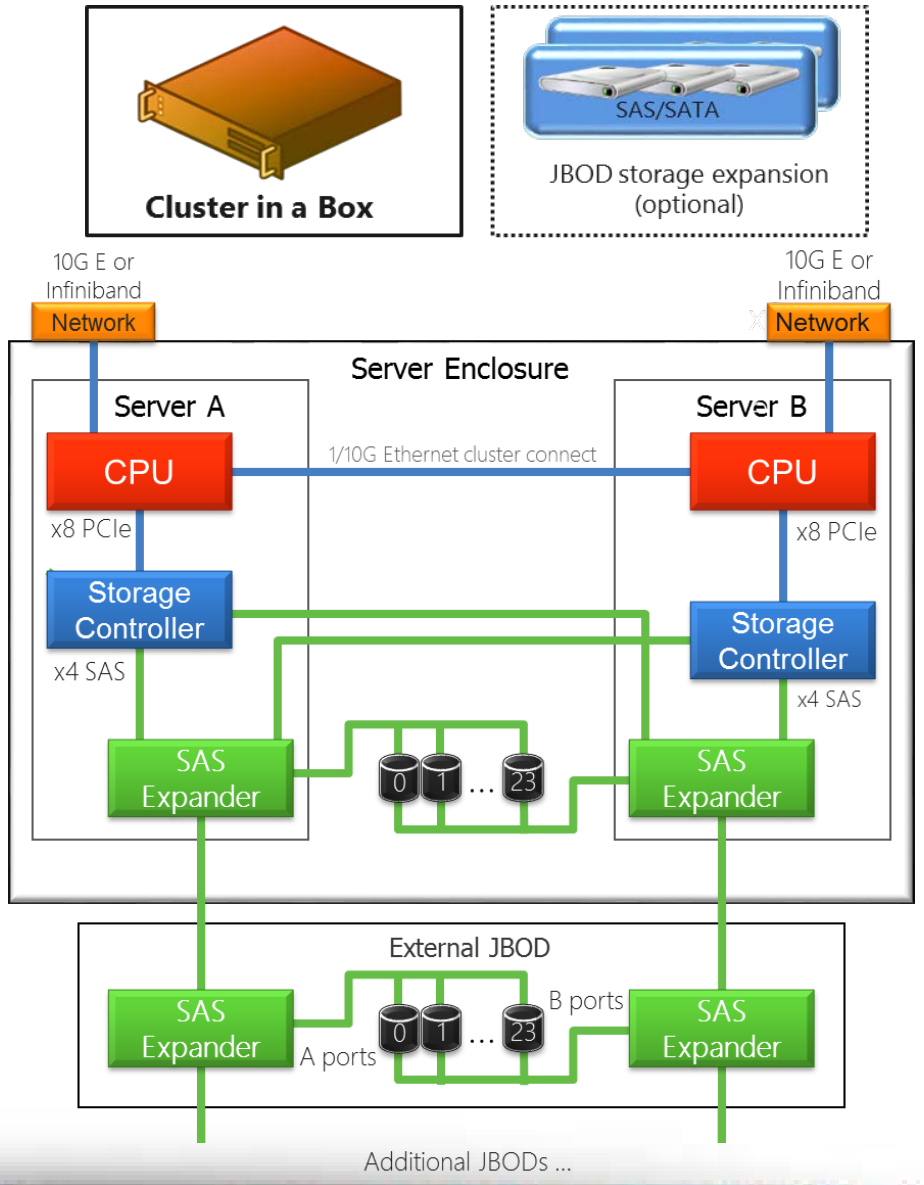

What are the Requirements ?

GOALS ARE FUNCTIONALITY AND CUSTOMER EXPERIENCE, NOT SPECIFIC TECHNOLOGY OR IMPLEMENTATIONS

Requirement Area	Goals	Requirement Summary
RAID	• Survive single-component failures • Reliable handling of node failure by storage systems without data loss	• RAID levels 1, 5, 6 or 10 (or equivalent) • RAID "write hole" solution
Windows Failover Clustering	Needed to survive node failure	• SPC-3 Persistent Reservation support • Shared LUNs accessible from all nodes
LUN access after failover	Transparent failover at the application layer requires • Data Integrity • Prevention of read/write timeouts	• Preserve data for all acknowledged writes • 5 seconds maximum blocking time for up to 100 LUNs • All I/O requests must complete within 25s
Minimum active mode	Customer expectation is active use of ALL nodes	Dual-Active (non-shared LUN access)
Recovery Processing	Transparent failover should not require IT admin attention to continue normal operation	Recovery processing resulting from node failure completes automatically
Performance	Customer performance expectations must be met with Windows in Write-Through cache mode	System builders are permitted to publish performance measurements for the device in clustered configurations

Test Scenario Design

E2E Test Scenarios to Validate CA

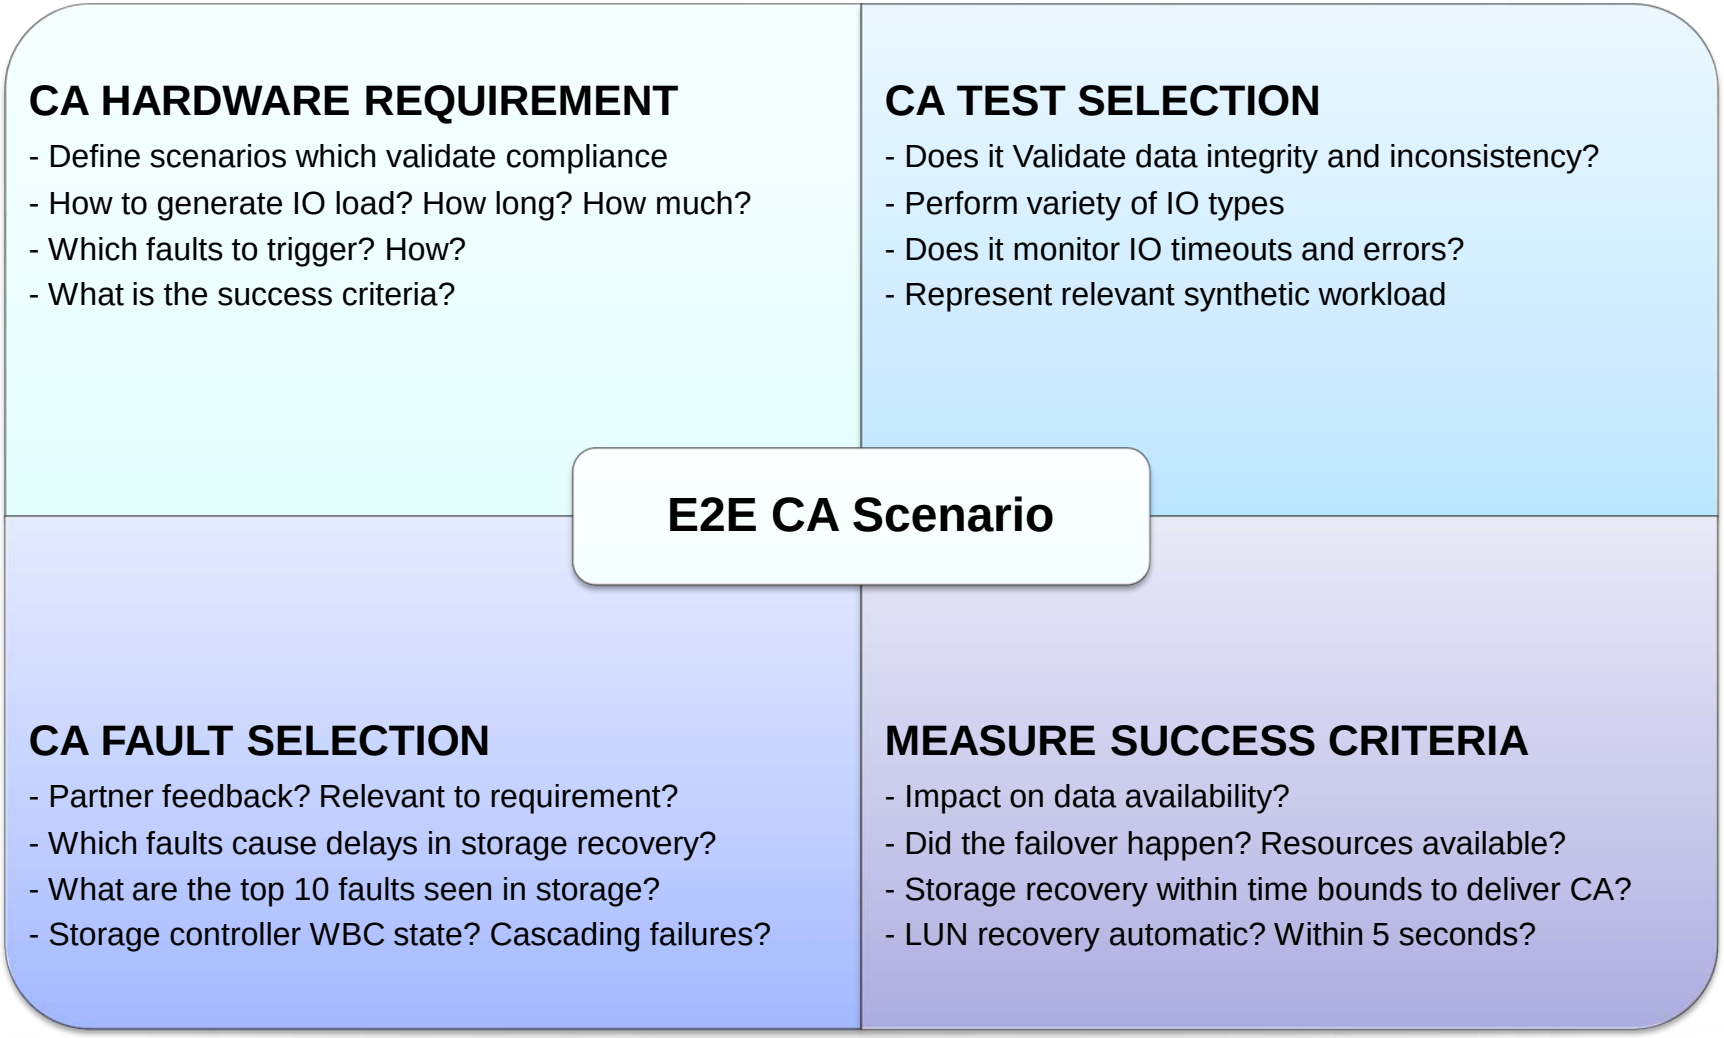

Test Selection for CA Validation

- **CA.Test.SysCache**
  - Data Verification Logo Supplemental Test
  - Variety of IOs (seq., async, mapped, buffered, truncation, file pointer move, scatter-gather, read/write attributes etc.). Ability to detect and halt after detecting corruption

- **CA.Test.RapidFile**
  - Data Verification Logo Supplemental Test
  - Multithreaded async IOs to different regions of file
  - Verify files for data corruption, IO errors or timeouts

- **CA.Test.SQLIOSim**
  - Generates IO Patterns similar to IO activity of SQL Server
  - Load generator stressing the file system, Verification of data integrity
  - Lazywriter, logwriter, checkpoint, grow/shrink, read-ahead, random access
  - Publically available and also ships with SQL Server

Fault Selection for CA Validation

- Goal: Simulate failures that cause resources to move between cluster nodes

'Planned' Faults	'Unplanned' Faults
Cluster Move-Group OS-Reboot	Power Supply Failures OS-Kernel Crash HBA Hang or Reset Cascading HBA Failures
Online RAID Migration	Failures during RAID migration
Dual Active Planned FailOver	Failures during Reconstruction Failures during Capacity Expansion
Physical Drive Hot Swap	Physical Drive Link Failure Physical Drive Failure

How do the tests map to CA Requirements?

Requirement Area	Requirement Summary	Test Scenarios / Workflows
RAID	• RAID levels 1, 5, 6 or 10 (or equivalent) • RAID "write hole" solution	- Physical Drive Failure - Hot swap of Physical Drive - Online RAID Migration - Online Capacity Expansion
Windows Failover Clustering	• SPC-3 Persistent Reservation support • Shared LUNs accessible from all nodes	- FailoverClustering Cluster Validation
LUN access after failover	• Preserve data for all acknowledged writes • 5 seconds maximum blocking time for up to 100 LUNs • All I/O requests must complete within 25 seconds	- Unplanned OS Kernel Crash - Physical Drive to Storage Link Failure - HBA Power Failure - Controller Hang or Reset - Maximum LUNs Failover - Cascading HBA Failures - Planned Cluster Move-Group - Planned OS Reboot
Minimum active mode	Dual-Active (non-shared LUN access)	- Dual-Active Cluster Planned Failovers
Recovery Processing	Recovery processing resulting from node failure completes automatically	- Failover During Reconstruction - Failover During RAID Migration - Failover During Capacity Expansion
Performance	System builders are permitted to publish performance measurements for the device in clustered configurations	- Report generation

Validation Toolkit Overview

Key Features

- Tests performed end to end with clients applying load to a clustered server
  - Runs multiple data integrity tests and synthetic I/O workloads from clients
  - Triggers planned and unplanned cluster failovers and verifies:
    - No loss of data and no data inconsistency observed from the clients
    - No impact on data availability (IO timeouts or failures) is observed
    - Volumes always accessible and LUN downtime is within the required bounds
    - Failover happens within time limit to deliver CA
- Fully automated software faults and extensible plugin framework to enable IHVs simulating hardware faults
- Support for WS2012 and WS2012r2 with iSCSI Target, SMB and NFS CA

Failover Time Measurements

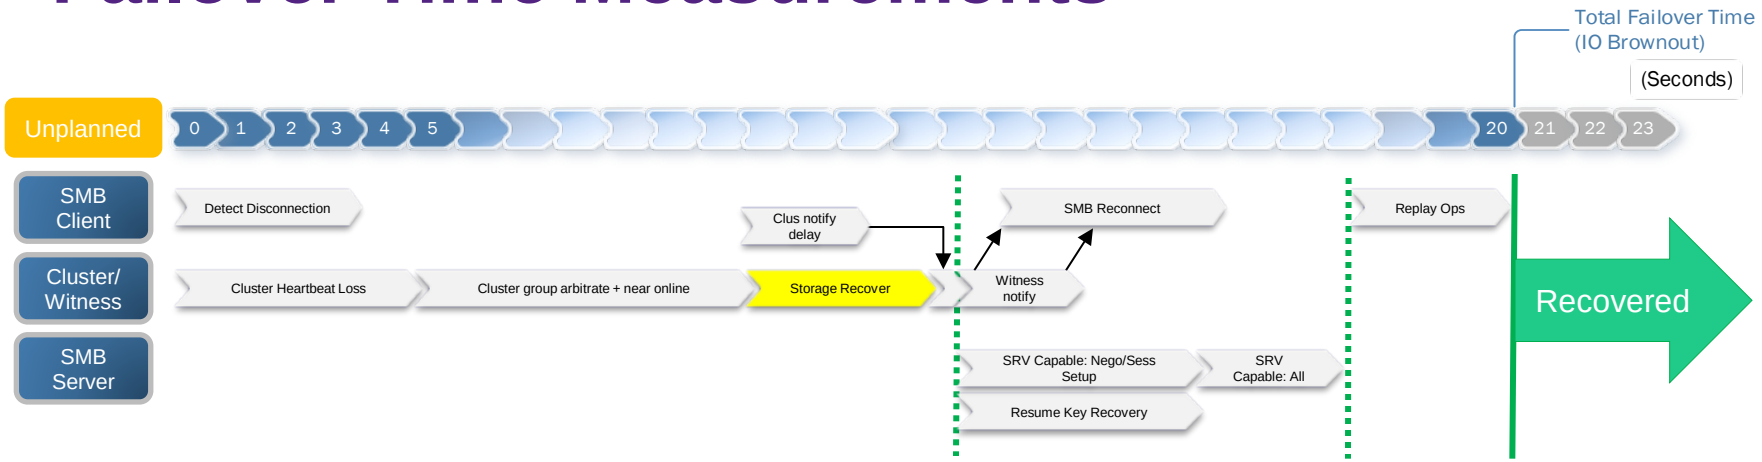

Total Failover Time (IO Brownout)

- Failover time **as seen by the Client** (SMB, NFS, iSCSI) application
- Measured as max IO latency from client during failover
- Tool: Canary
- **COMPONENTS RESPONSIBLE: SMB, NFS, iSCSI, NTFS, Cluster, and Storage Controller**

Storage Recovery

- Time to **recover storage** post cluster failover action
- Measured as time between first disk online call issued on surviving node and last disk coming online on surviving node
- Tool: LUNAccessTime.ps1
- **COMPONENTS RESPONSIBLE: Cluster and Storage Controller**

Disk Online Latency (LUN Recovery)

- **Time to online a volume** (LUN) on surviving node post cluster failover
- Measured as time between PR Arbitrate and volumes arriving on surviving node as seen by MountMgr & Cluster
- Tool: LUNAccessTime.ps1
- **COMPONENTS RESPONSIBLE: Only Storage Controller**

Test Framework Architecture

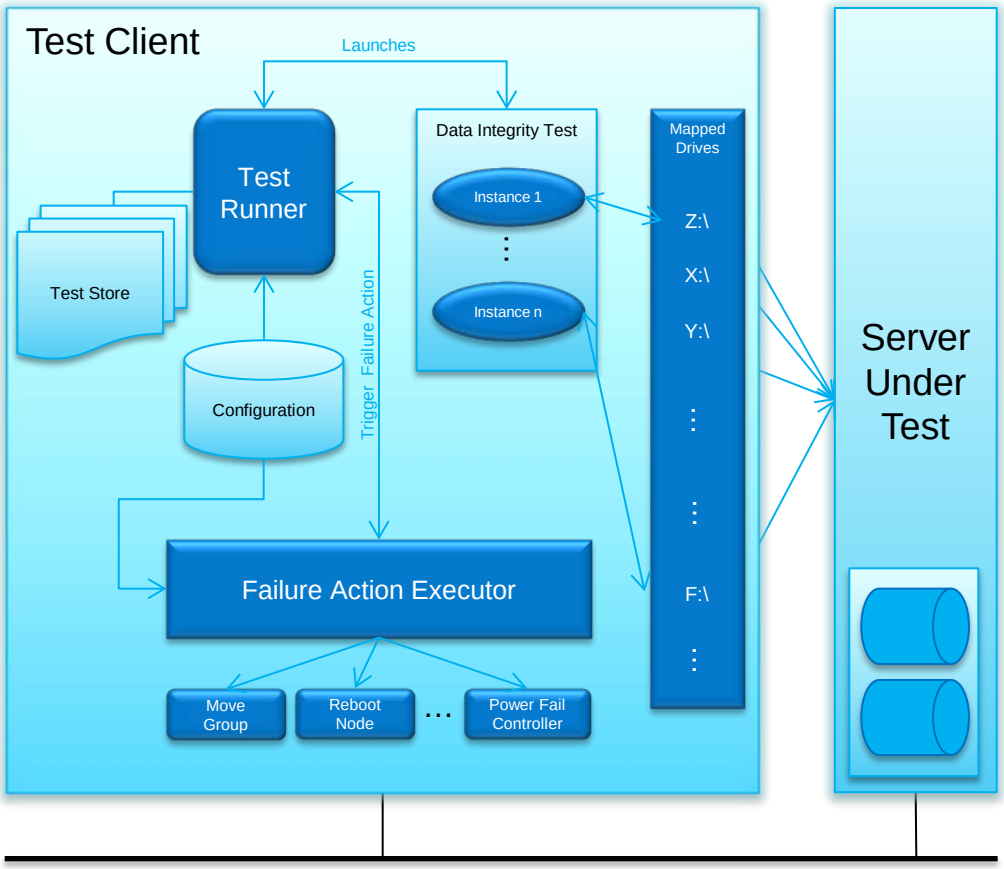

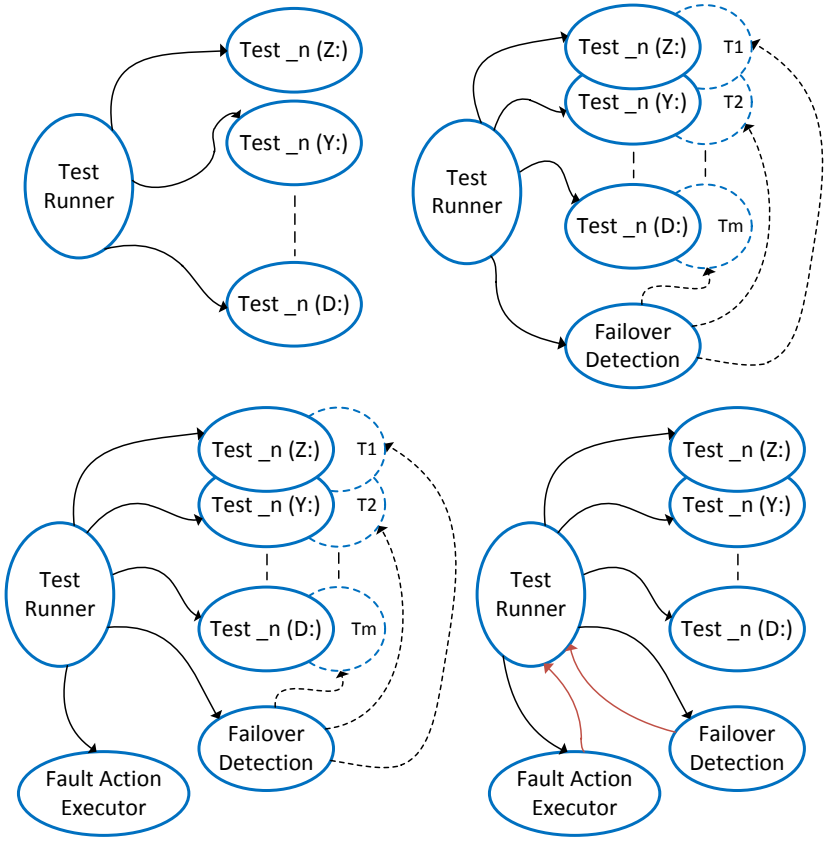

Test Framework - Workflow

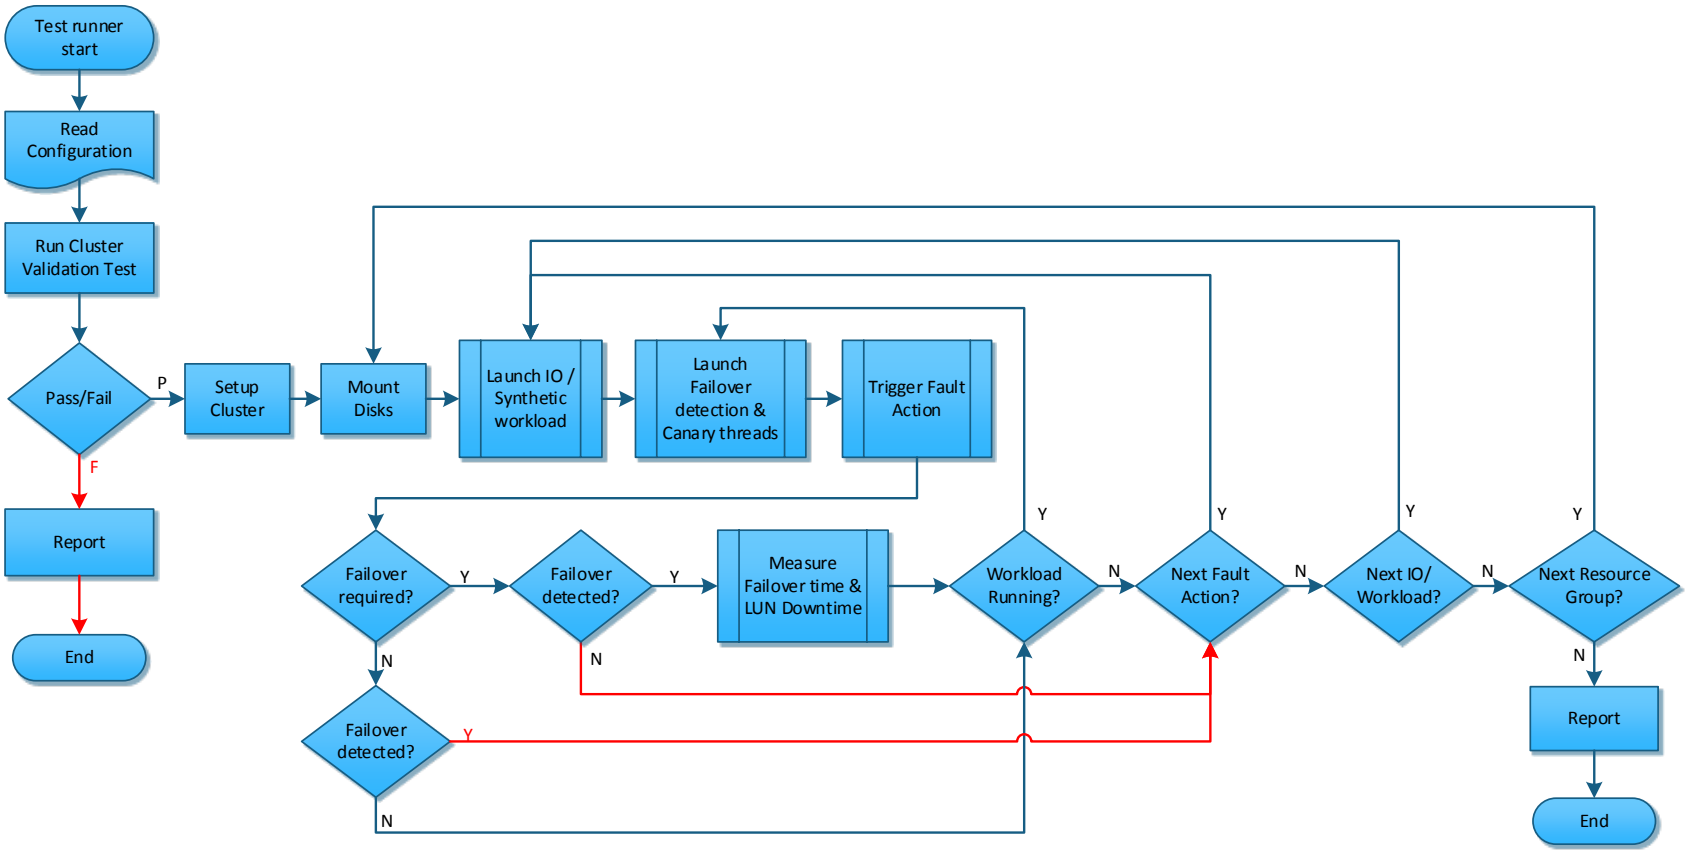

Example Test Scenario – OS Kernel Crash

Review CA Requirement
- LUNs to remain accessible post failover
- Data integrity and prevention of IO timeouts or errors
- 5 seconds LUN online latency for up to 100 LUNs or max supported

CA Test (IO Workload)
- Data integrity test, synthetic I/O workloads
- End to end tests with client applying load
- Tests with ability to detect corruption , inconsistency or timeouts

Test Workflow
CA.LUNAccess.OS-Kernel-Crash

Trigger Fault Action
- Unplanned software failure
- OS kernel crash of active cluster node
- No further interaction from OS to the Storage Controller

Measure Success Criteria
1. Failover of resources is observed
2. LUN recovery within 5 seconds
3. No loss of data, inconsistency, errors or timeouts observed
4. Overall requirement result

```
************************************************************
*** Starting CA Validation Test Scenario Workflow (2  of 4 ) ***
************************************************************
08/26/2014-18:38:23 : Test name:  SysCache Data Integrity Test Suite
08/26/2014-18:38:23 : Failure scenario:  Node software failure via OS kernel crash
08/26/2014-18:38:23 : Starting the IO workload...
08/26/2014-18:38:23 : Workload Duaration: 2 munutes
08/26/2014-18:38:23 : Started CA.Test.SysCache SysCache Data Integrity Test Suite on drive L
08/26/2014-18:38:23 : Started CA.Test.SysCache SysCache Data Integrity Test Suite on drive M
08/26/2014-18:38:23 : Started CA.Test.SysCache SysCache Data Integrity Test Suite on drive N
08/26/2014-18:38:23 : Started CA.Test.SysCache SysCache Data Integrity Test Suite on drive O
08/26/2014-18:38:23 : Started CA.Test.SysCache SysCache Data Integrity Test Suite on drive P
08/26/2014-18:38:23 : Warm-up time: 15 seconds
08/26/2014-18:38:38 : Warm-Up time finished, start the failure actions now
08/26/2014-18:38:38 : New iteration start.
08/26/2014-18:38:38 : Enter PreAction
08/26/2014-18:38:48 : KernelCrash Node software failure via OS kernel crash pre-action done
08/26/2014-18:38:48 : Enter FailureAction
08/26/2014-18:38:48 : Failovergroup is CBCCLU12-iSCSI
08/26/2014-18:38:48 : Started Canary failover time calculator
08/26/2014-18:38:48 : Started LUN downtime calculator
08/26/2014-18:38:48 : Canary::Spawned all the threads, waiting for all to complete within 300 seconds
08/26/2014-18:39:03 : Generating fault KernelCrash Node software failure via OS kernel crash on active node
08/26/2014-18:39:56 : Finished LUN downtime calculator. exit code 0
08/26/2014-18:39:56 : Waiting for canary to exit
08/26/2014-18:39:57 : canary exits
08/26/2014-18:40:33 : All instances of CA.Test.SysCache SysCache Data Integrity Test Suite finished
08/26/2014-18:40:33 : TestInstance::PerformTestInstance - Waiting for end of iteration
08/26/2014-18:40:42 : Finished KernelCrash Node software failure via OS kernel crash on active node. exit code 0
08/26/2014-18:40:42 : Failure scenario ended by Test Runner
08/26/2014-18:40:42 : Enter PostAction
08/26/2014-18:40:42 : Finished iteration index 1. Total to complete 1
08/26/2014-18:40:42 : Signal End of iteration
08/26/2014-18:40:42 : TestInstance::PerformTestInstance - Got the end of iteration
------------------------------------------------------------
Report: Requirement = LUNAccess
Report: Test Scenario = KernelCrash Node software failure via OS kernel crash
Report: Criteria - Failover Observed = Yes, Failover Expected = Yes
Report: Criteria - Max LUN Downtime = 0.581 seconds, pass
Report: Criteria - Data Availability Impact = Pass
Report: Overall Result = Pass
------------------------------------------------------------
************************************************************
*** Completed CA Validation Test Scenario Workflow (2  of 4 ) ***
************************************************************
```

Report

iSCSI Target Resource Group

CA Hardware Requirement	OverallResult (4)	Failure Action	(1) CA-Criteria-01			Failover Time Seen by Client (Worst)	(2) CA-Criteria-02		(3) CA-Criteria-03	
			Result	Failover Expected	Detected		Result	LUN Downtime (Worst)	Result	Impact on Availability
LUNAccess	PASS	Node software failure via OS kernel crash	PASS	true	1 of 1	10.80	PASS	0.58	PASS	5 of 5

How Toolkit Works

- Read CA Requirements
- Read CA Toolkit User Guide
- Hardware Topology
- Cluster Hardware Requirements
- Shared Storage Configuration

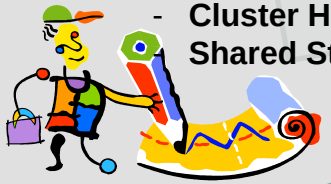

Planning

Setup

Running

Faults

Result Report

Submission

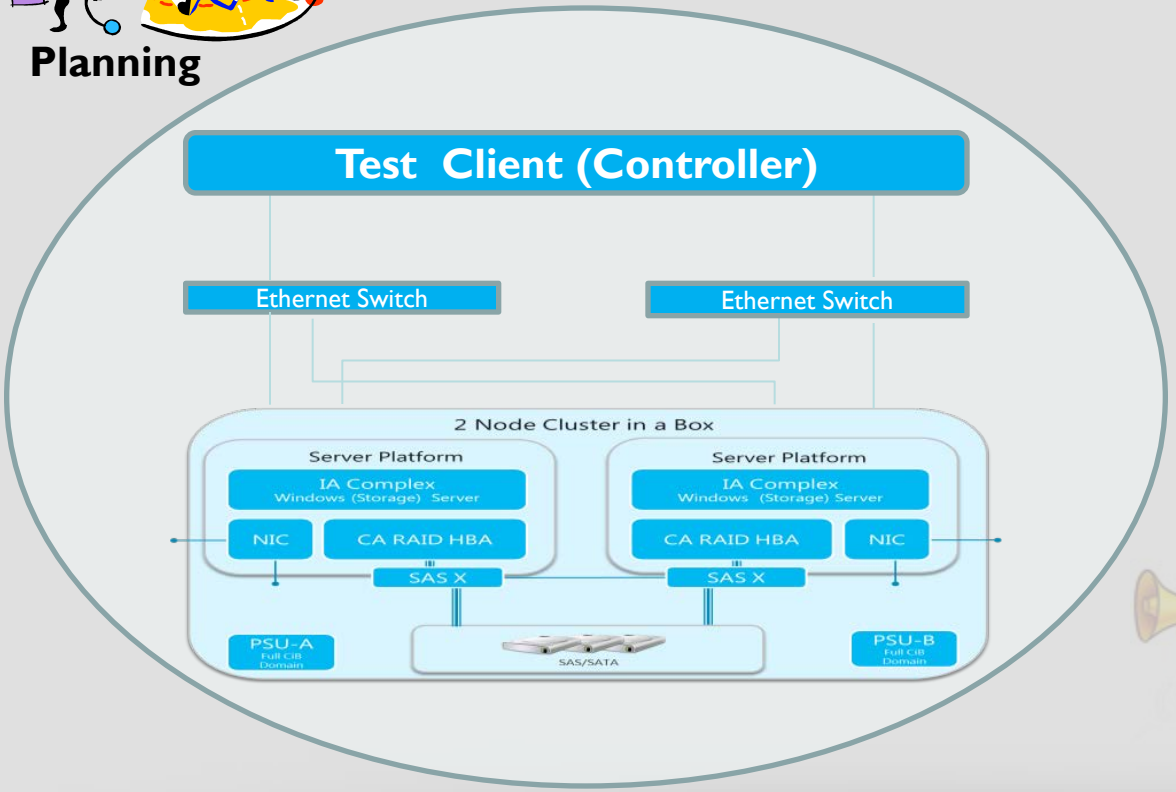

How Toolkit Works

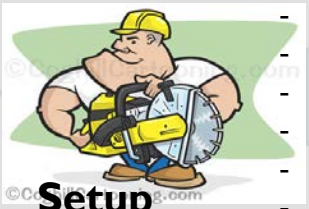

Setup

- Toolkit Configuration
- Launch Toolkit
- Run Cluster Validation Test
- Create Cluster
- Create Resource Groups
- Create Shares and VHDs
- Mount Shares/Disks from Client

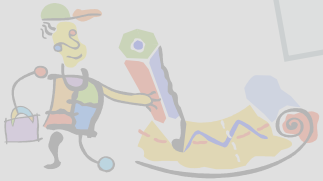

Planning

Faults

Result Report

Submission

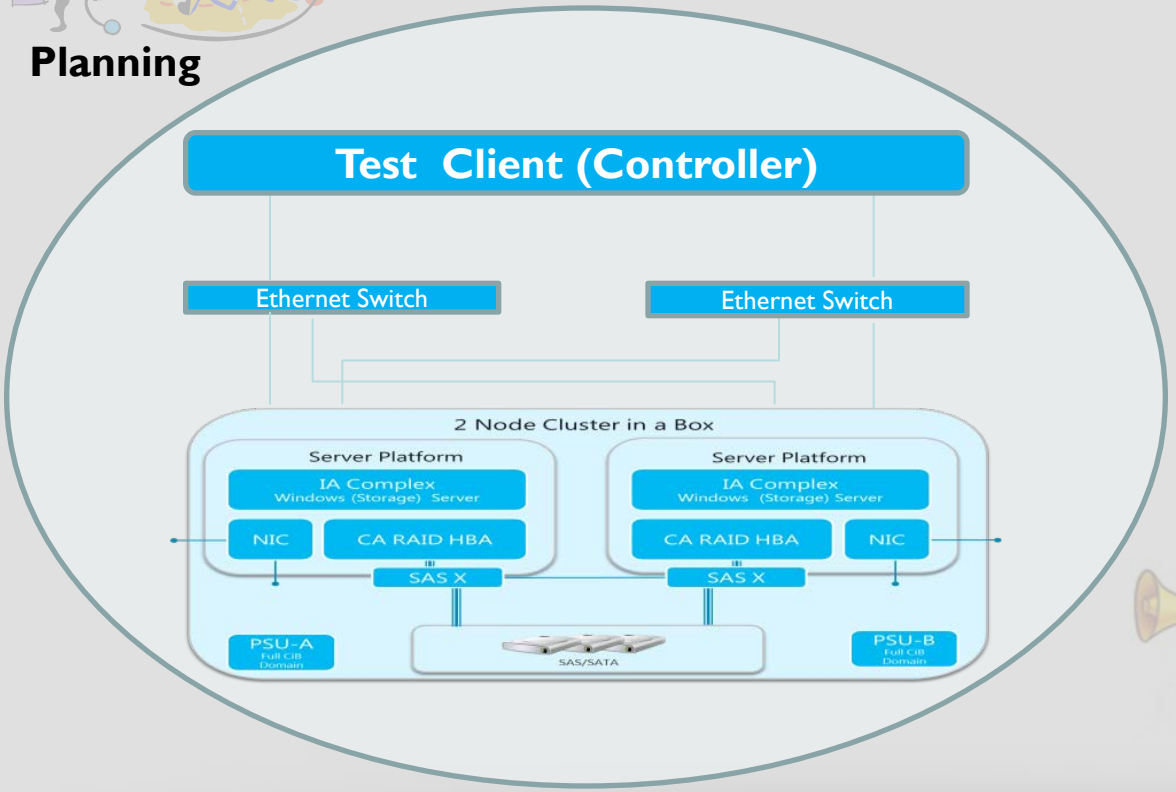

How Toolkit Works

- Runs Data Integrity and IO workload tests from client
- Monitors Cluster

Running

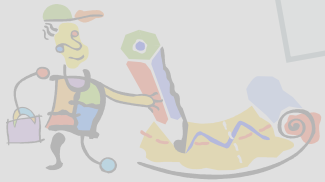

Planning

Setup

Faults

Result Report

Submission

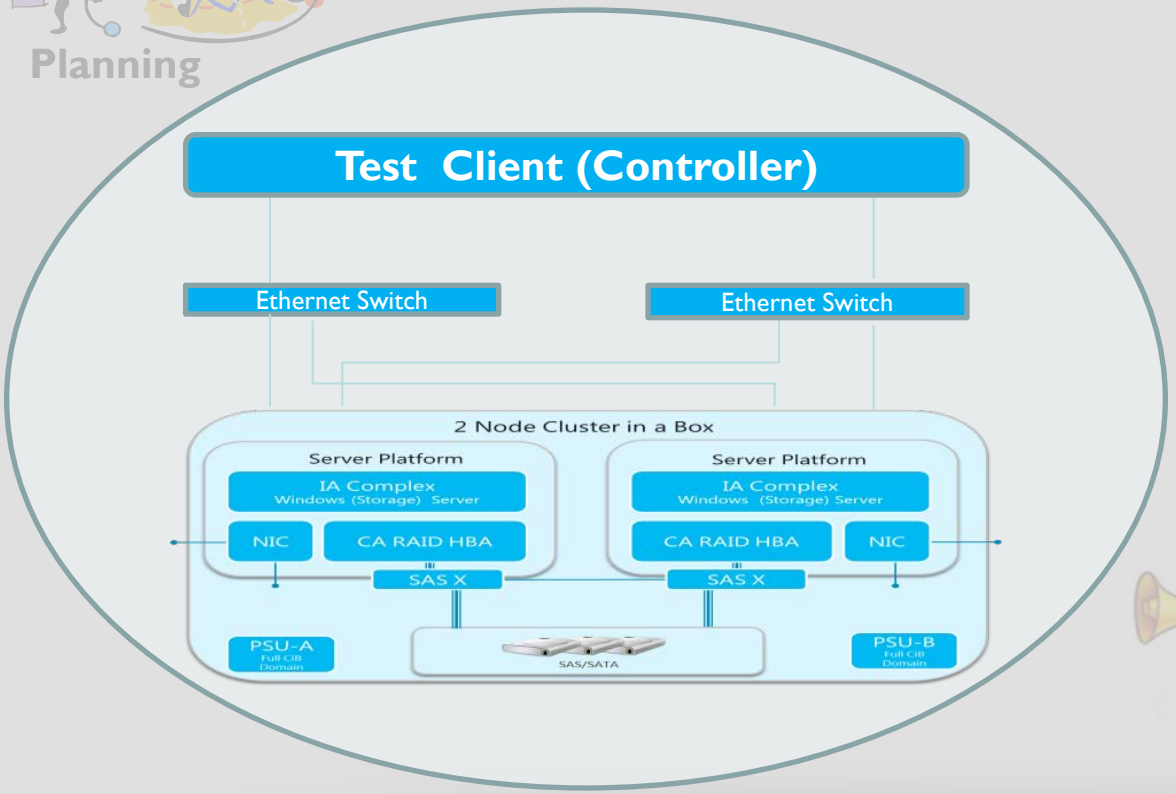

How Toolkit Works

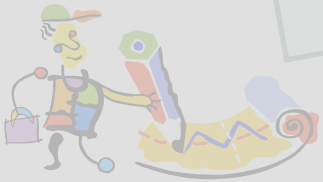

Planning

Setup

Running

Faults

- **Auto trigger software faults**
- **Prompts for hardware faults**
- **Plugins to automate**
- **Detects failover**
- **Measure LUN recovery time**
- **Measure failover time**

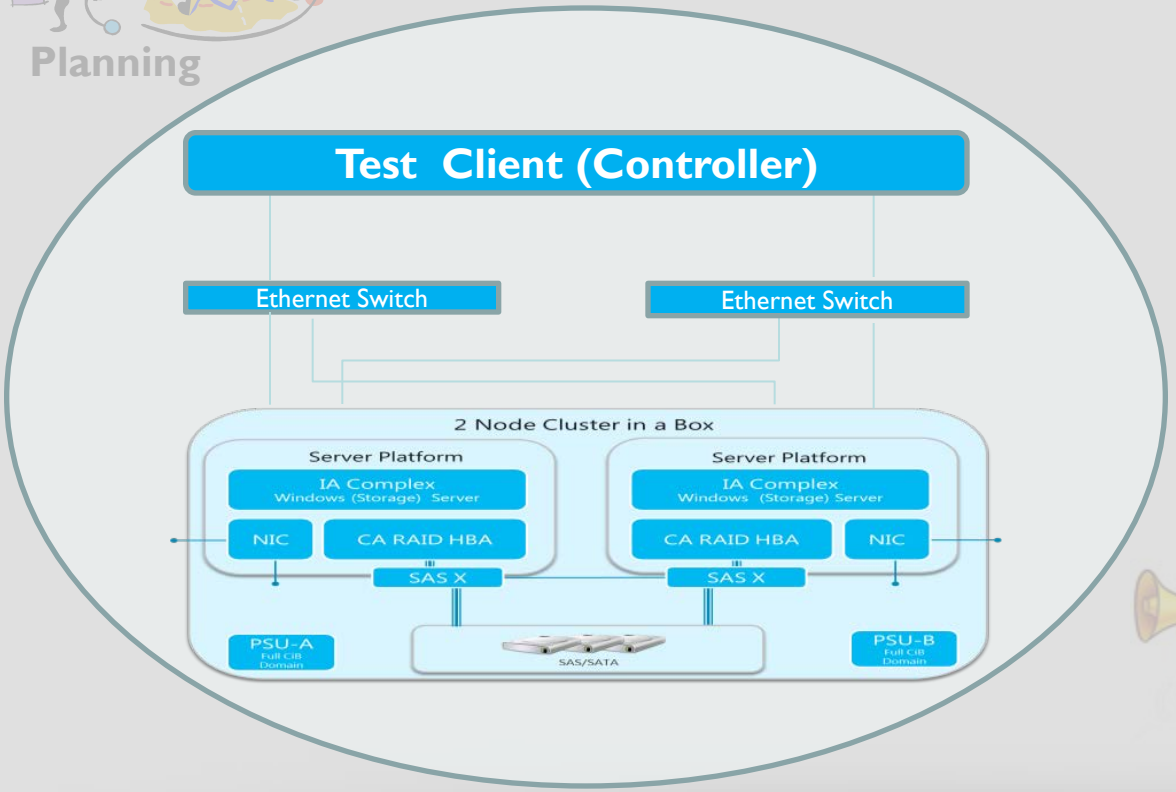

Result Report

Submission

How Toolkit Works

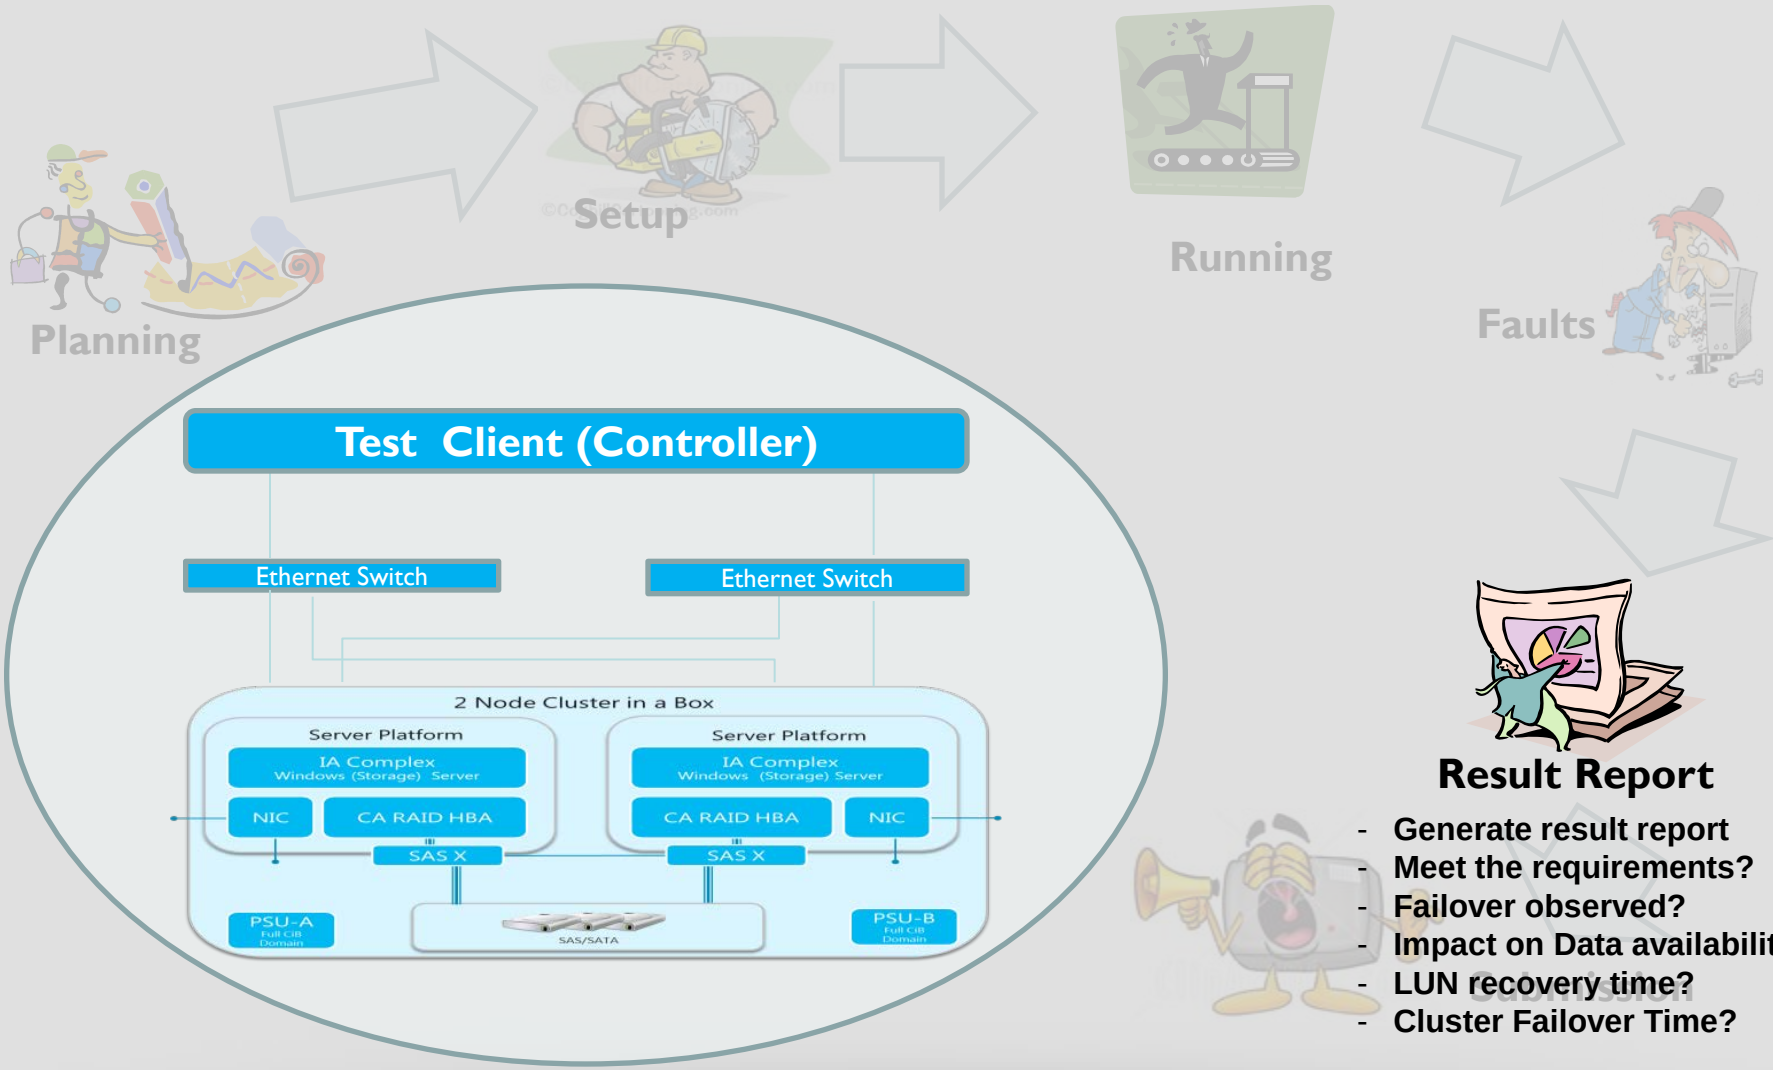

Result Report

- Generate result report
- Meet the requirements?
- Failover observed?
- Impact on Data availability?
- LUN recovery time?
- Cluster Failover Time?

How Toolkit Works

Planning

Setup

Running

Faults

Result Report

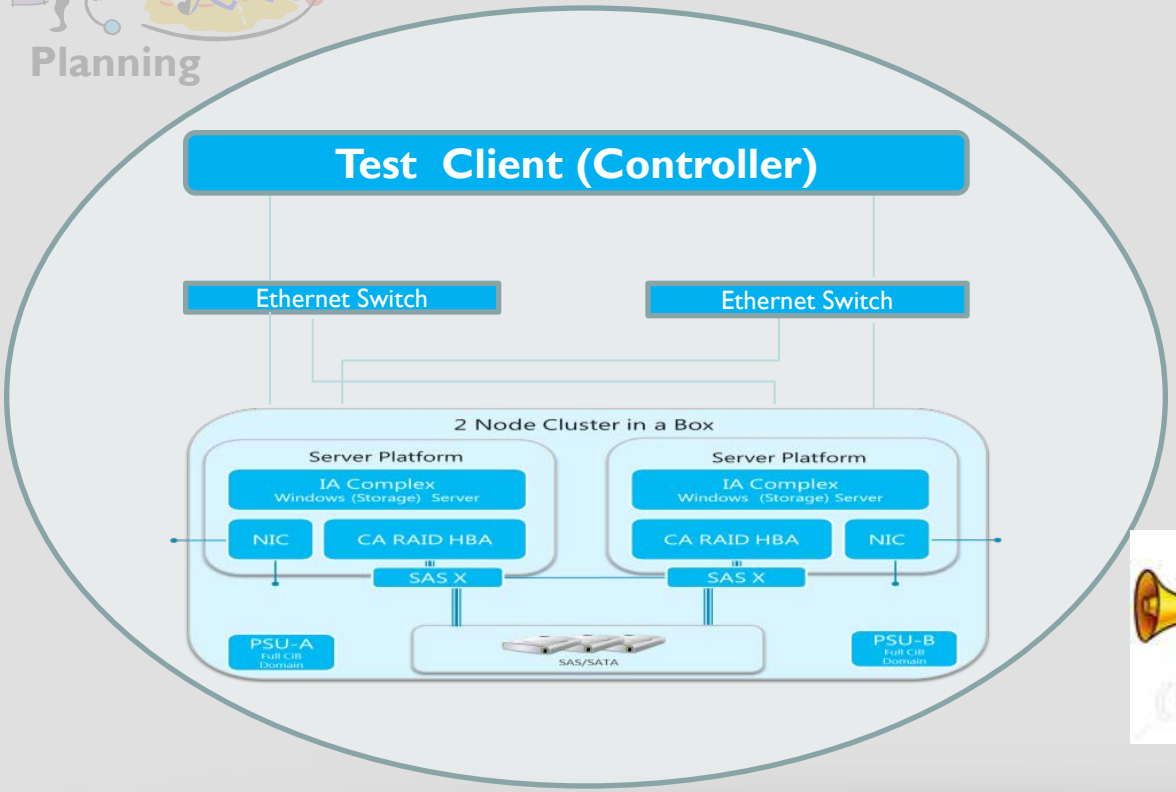

- Analyze the report
- Troubleshooting
- Submit results

Submission

Key Takeaways

- Validate Continuous Availability by using customer scenario-focused testing
- Create validation requirements based on application expectations in customer deployments
- Simulate both ‘Planned’ and ‘Unplanned’ faults during typical application usage patterns
- Measure success based on transparent failure recovery during application usage (rather than implementation specific details)

Q & A

Appendix

Abstract

CONTINUOUS AVAILABILITY: A SCENARIO VALIDATION APPROACH

- **Abstract:**
- Systems designed for ‘Continuous Availability’ functionality need to satisfy strict failure resiliency requirements from scenario, performance and reliability perspective. Such systems normally incorporate a wide variety of hardware-software combinations to perform transparent failover and accomplish continuous availability for end applications. Building a common validation strategy for diversified software and hardware solution mix needs focus on the end-user scenarios for which customers would deploy these systems. We developed the ‘Cluster In a Box’ toolkit to validate such ‘Continuous Availability’ compliant systems. In this presentation, we examine the test strategy behind this validation. We focus on end-to-end scenarios, discuss different user workloads, potential fault inducers and the resiliency criteria that has to be met in the above deployment environment.
- **Learning Objectives**
- 1. Continuous Availability and Transparent Failover
- 2. End-to-end scenario testing strategy
- 3. User workload simulation
- 4. Environment fault injection
- 5. System resiliency SLA/criteria measurement

Test Configuration File Example

```
<?xml version="1.0" encoding="utf-8"?>
<IHVSuite Name="CA Validation Suite">
  <ClusterConfig>
    <Config Name="ClusterName"      Value="CBCCLU12" />
    <Config Name="NodeNames"        Value="28-1303C090808,28-1303C090809" />
    <Config Name="ResourceGroup"    Value="CBCCLU12-iSCSI" />
    <Config Name="ResourceGroup"    Value="CBCCLU12-FS" />
    <Config Name="ResourceGroup"    Value="CBCCLU12-SOFS" />
  </ClusterConfig>
  <TestConfig>
    <Test ID="CA.Test.SysCache"
          Name="SysCache Data Integrity Test Suite"
          Command="readwrit.exe -S -N -B -c -v 0"
          Active="Yes"
          WarmUpTimeInSeconds="15"
          RunTimeInMinutes="30" />
    <Test ID="CA.Test.RapidFile"
          Name="Rapid File Data Availability Test Suite"
          Command="rapidfile.exe"
          Active="Yes"
          WarmUpTimeInSeconds="15"
          RunTimeInMinutes="30" />
    <Test ID="CA.Test.SQLIOSim"
          Name="SQL IO Simulation Test Suite"
          Command="sqliosim.com -cfg sqliosim.default.cfg.ini -size 1"
          Active="No"
          WarmUpTimeInSeconds="15"
          RunTimeInMinutes="30" />
  </TestConfig>
  <FailureScenarios FailureActionStore="FailureActions">
    <FailureScenario ID="KernelCrash"
                     Requirement="LUNAccess"
                     Type="Software"
                     Name="Node software failure via OS kernel crash"
                     FailoverExpected="Yes"
                     Iterations="1">
      <PreAction Command="Identifies and performs kernel crash failure on the active cluster node"
                 WaitPeriodInSeconds="10" />
      <FailureAction Command="FailOverSimulator.ps1 -FailType CrashMachine -ResourceGroup $ResourceGroup$ "
                     TimeoutInSeconds="300" />
      <PostAction Command=""
                  WaitPeriodInSeconds="10" />
    </FailureScenario>
  </FailureScenarios>
</IHVSuite>
```

- **Supported Resource Group**
  - Singleton File Server + SMB/NFS shares
  - Scale-out File Server + CSV + SMB shares
  - iSCSI Target
- **Multiple IO Workloads**
  - SysCache
  - RapidFile
  - SQLIOSim
- **Pluggable Fault Generator**
  - Automated fault action
  - Manual fault action

Fault Action Configuration

```
<FailureScenarios FailureActionStore="FailureActions">
  <FailureScenario ID="KernelCrash"
                   Requirement="LUNAccess"
                   Type="Software"
                   Name="Node software failure via OS kernel crash"
                   FailoverExpected="Yes"
                   Iterations="1">
    <PreAction Command="Identifies and performs kernel crash failure on the active cluster node"
               WaitPeriodInSeconds="10" />
    <FailureAction Command="FailOverSimulator.ps1 -FailType CrashMachine -ResourceGroup $ResourceGroup$ "
                   TimeoutInSeconds="300" />
    <PostAction Command=""
                WaitPeriodInSeconds="10" />
  </FailureScenario>
  <FailureScenario ID="HBA-Reset"
                   Requirement="LUNAccess"
                   Type="Hardware"
                   Name="HBA controller reset or hung"
                   FailoverExpected="Yes"
                   Iterations="1">
    <PreAction Command=""
               WaitPeriodInSeconds="10" />
    <FailureAction Command="ccu.exe cli controller local lockup this_controller "
                   TimeoutInSeconds="1000" />
    <PostAction Command=""
                WaitPeriodInSeconds="300" />
  </FailureScenario>
  <FailureScenario ID="DriveHotSwap"
                   Requirement="RAID"
                   Type="Hardware"
                   Name="Single physical drive hot swap in RAID set"
                   FailoverExpected="No"
                   Iterations="2">
    <PreAction Command=""
               WaitPeriodInSeconds="0" />
    <FailureAction Command="Please perform single physical drive Hot Swap in the RAID set\n
                            Refer to user guide for more detailed steps"
                   TimeoutInSeconds="-1" />
    <PostAction Command=""
                WaitPeriodInSeconds="300" />
  </FailureScenario>
</FailureScenarios>
```

- **PreAction, FailureAction & PostAction**
  - Generic Commands
  - Plain text, displays message box
  - Can execute .cmd, .ps1, .exe, .*
- **Inbuilt Software Fault Simulator:** FailOverSimulator.ps1
  - Software failure simulator capable of triggering move resource group, reboot node, crash node
- **IHV Example:**
  - Model enables IHVs to integrate their automation to trigger h/w faults
  - **ccu.exe cli controller local lockup this_controller**
- **Manual Example:**
  - Plain text for the FailureAction Command
  - Manually trigger the failure action
  - Close the message box

Cluster Log - LUN Recovery Time

```
2014/08/26-21:02:29.963 INFO  [RCM] rcm::RcmResource::Online: bringing File Server (\\CiB007-FS)'s provider resource 'Cluster Disk 10' online.
2014/08/26-21:02:29.966 INFO  [RCM] Res Cluster Disk 10: Offline -> OnlineCallIssued( StateUnknown )
2014/08/26-21:02:29.966 INFO  [RCM] TransitionToState(Cluster Disk 10) Offline-->OnlineCallIssued.
2014/08/26-21:02:29.980 INFO  [RCM] rcm::RcmResource::Online: NFS-CiB007-FS's provider 'Cluster Disk 10' is already coming online (state is OnlineCallIssued).
2014/08/26-21:02:29.984 INFO  [RCM] STM pre online check for Cluster Disk 10 returned 8
2014/08/26-21:02:29.984 INFO  [RCM] Issuing Online(Cluster Disk 10) to RHS.
2014/08/26-21:02:29.985 INFO  [RES] Physical Disk <Cluster Disk 10>: ResHardDiskOnlineV2: Online request.
2014/08/26-21:02:29.991 INFO  [RCM] HandleMonitorReply: ONLINERESOURCE for 'Cluster Disk 10', gen(0) result 997/0.
2014/08/26-21:02:29.991 INFO  [RCM] Res Cluster Disk 10: OnlineCallIssued -> OnlinePending( StateUnknown )
2014/08/26-21:02:29.991 INFO  [RCM] TransitionToState(Cluster Disk 10) OnlineCallIssued-->OnlinePending.
2014/08/26-21:02:29.992 INFO  [RES] Physical Disk <Cluster Disk 10>: ResHardDiskArbitrateInternal request Not a Space: Uses FastPath
2014/08/26-21:02:29.992 INFO  [RES] Physical Disk <Cluster Disk 10>: ResHardDiskArbitrateInternal: Clusdisk driver handle or event handle is NULL.
2014/08/26-21:02:30.008 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpQueryDiskFromStm: ClusterStmFindDisk returned device='\\?\scsi#disk&ven_lsi&prod_
2014/08/26-21:02:30.067 INFO  [RES] Physical Disk <Cluster Disk 10>: Arbitrate - Node using PR key 774a8dc00001734d
2014/08/26-21:02:30.146 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpPRArbitrate: Fast Path arbitration...
2014/08/26-21:02:30.226 INFO  [RES] Physical Disk <Cluster Disk 10>: Successful reserve, key 774a8dc00001734d
2014/08/26-21:02:30.340 INFO  [RES] Physical Disk <Cluster Disk 10>: Disk is offline
2014/08/26-21:02:30.340 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpSetUnsetDiskFlags(mask=0x00000007, SetCluster=1, SetCsv=0, SetMaintenanceMode=0
2014/08/26-21:02:30.346 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpGetDiskHandle: EXIT, status 0
2014/08/26-21:02:30.348 INFO  [RES] Physical Disk <Cluster Disk 10>: OnlineThread: Successfully cleared CSV state with partmgr.
2014/08/26-21:02:30.352 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpWaitForPartitionsToArrive: Begin wait for \\?\GLOBALROOT\Device\Harddisk16\Partitio
2014/08/26-21:02:30.358 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpWaitForPartitionsToArrive: Status ERROR_IO_PENDING from IOCTL_DISK_ARE_VOLUME
2014/08/26-21:02:30.630 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpWaitForPartitionsToArrive: Wait success and IOCTL_DISK_ARE_VOLUMES_READY compl
2014/08/26-21:02:30.630 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpWaitForPartitionsToArrive: wait for volumes
2014/08/26-21:02:31.520 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpWaitForPartitionsToArrive: wait for volumes completed
2014/08/26-21:02:31.520 INFO  [RES] Physical Disk <Cluster Disk 10>: UnLockVolumesIfEncryptionEnabled
2014/08/26-21:02:31.526 INFO  [RES] Physical Disk <Cluster Disk 10>: ResHardDiskVolumeGuidPathnameAndDriveLetterChecks. Disk {16}.
2014/08/26-21:02:31.527 INFO  [RES] Physical Disk <Cluster Disk 10>: ResHardDiskGetWin32Pathnames: Found 2 mount points for disk {16}, partition {1}.
2014/08/26-21:02:31.527 INFO  [RES] Physical Disk <Cluster Disk 10>: HardDiskpDriveLetterAndVolumeGuidReset: Volume {\\?\Volume{ba4af359-0c43-11e4-80ba-806e
2014/08/26-21:02:31.528 INFO  [RES] Physical Disk <Cluster Disk 10>: VolumeIsNtfs: Volume \\?\GLOBALROOT\Device\Harddisk16\ClusterPartition1\ has FS type NTFS
2014/08/26-21:02:31.529 INFO  [RES] Physical Disk <Cluster Disk 10>: OnlineThread: HardDiskpVerifyVolume part=1 returned 0, state=0x00000000
2014/08/26-21:02:31.529 INFO  [RES] Physical Disk <Cluster Disk 10>: ResHardDiskVerifyMountFolderTargetVolumesAreClustered: Volume \\?\GLOBALROOT\Device\H
2014/08/26-21:02:31.547 INFO  [RES] Physical Disk <Cluster Disk 10>: MountPoint S:\ points to volume \\?\Volume{ba4af359-0c43-11e4-80ba-806e6f6e6963}\
2014/08/26-21:02:31.583 INFO  [RHS] Resource Cluster Disk 10 has come online. RHS is about to report status change to RCM
2014/08/26-21:02:31.584 INFO  [RCM] HandleMonitorReply: ONLINERESOURCE for 'Cluster Disk 10', gen(0) result 0/0.
2014/08/26-21:02:31.584 INFO  [RCM] Res Cluster Disk 10: OnlinePending -> Online( StateUnknown )
2014/08/26-21:02:31.584 INFO  [RCM] TransitionToState(Cluster Disk 10) OnlinePending-->Online.
2014/08/26-21:02:34.642 INFO  [RES] Physical Disk <Cluster Disk 10>: Path S:\nfsShare0 can be on the disk
2014/08/26-21:03:31.583 INFO  [RES] Physical Disk <Cluster Disk 10>: VolumeIsNtfs: Volume \\?\GLOBALROOT\Device\Harddisk16\ClusterPartition1\ has FS type NTFS
```