

# Top Tips and Tricks for SQL Server Failover Clustering

Allan Hirt  
Clustering MVP  
allan@sqlha.com  
<http://www.sqlha.com>



# The Basics Part One – Windows

- Use Windows Server 2008 R2
  - 64-bit only
  - For 32-bit, use 32-bit W2K8 (not R2) ... or W2K3 (out of support)
- Clustering still an Enterprise or Datacenter feature of Windows
- Windows Core variant not supported for SQL Server (standalone or clustered)
- No OS rolling upgrade from W2K3 → W2K8/W2K8 R2
- Up to 16 node support (SQL Server Standard only 2 on all OS editions)

# The Basics Part Two – SQL Server

- SQL 2005/2008/2008 R2 supported in a side-by-side configuration on W2K8/W2K8 R2
  - 2000/2005/2008/2008 R2 under W2K3
- 2000 not supported on W2K8 or W2K8 R2
- Still one default instance per cluster (up to 25 supported per cluster)
- Minimum supported versions:

	Windows Server 2008 RTM/SP2	Windows Server 2008 R2
SQL Server 2005	SP2+	SP3+
SQL Server 2008	RTM	SP1+
SQL Server 2008 R2	RTM	RTM

# What's New In Windows Server 2008/R2?

- Failover clustering rewritten as of Windows Server 2008 RTM
- Administration tool = Failover Cluster Manager
- Cluster validation
  - No more HCL/Windows Server Catalog
  - Must pass to be supported
  - HW must be 2008/2008 R2 logoed
  - SQL Server 2008+ relies on the results
  - Prevent a false positive
- Most tasks online
- PowerShell
- Hyper-V integration (i.e. Live Migration)

# Networking

- Still need redundancy and separate paths
  - Teaming fully supported on *all* NICs if desired
  - Not recommended to disable IPv6 on NICs
  - Minimum of two networks
    - iSCSI will require a dedicated network
    - Live Migration may increase networks (Public, Private, Live Migration traffic, CSV, etc.)
- [http://technet.microsoft.com/en-us/library/ff428137\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/ff428137(WS.10).aspx)
- SQL Server does not support multiple subnets feature of W2K8 failover clustering
    - *Must still use VLAN if configuring a geographically dispersed cluster with SQL Server 2005 and SQL Server 2008 RTM/R2*
    - Denali (SQL Server vNext) supports this

# Security

- Windows/AD
  - New - Cluster Name Object (CNO) and Virtual Cluster Object (VCO)
  - Domain account needed for cluster creation and administration
    - Requires Create Computer Objects right on Computers OU
    - If cannot give CCO, create the CNO and VCO manually
  - Domain account NOT used to run the cluster service – runs under a special context
- SQL Server
  - Still requires service accounts
  - Use a Service SID; do not have to use domain groups unless using W2K3
  - No local Administrator requirements unless using something like xp\_cmdshell (off by default)
- Watch firewall in Windows
- Never use a domain admin account or assign domain admin privileges – THEY ARE NOT NEEDED

# Anti-Virus, SQL Server, and Clustering

- Not recommended if not needed
- If required, set exclusions:
  - .mdf (data), .ldf (log), .ndf (additional data)
  - .bak (default backup extension), .trn (default t-log backup extension)
  - All directories with Analysis Services data, log, temporary files, backups
  - Entire quorum/witness disk
  - Filestream containers
  - \MSDTC directory for MSDTC disks (if used)
  - \Cluster subdirectory under %windir%
- SQL Server & anti-virus KB 309422
- Cluster & anti-virus KB 250355

# Storage Configuration

**Install a SQL Server Failover Cluster**






## Cluster Disk Selection

Select shared cluster disk resources for your SQL Server failover cluster.

Specify the shared disks to be included in the SQL Server resource cluster group. The first drive will be used as the default drive for all databases, but this can be changed on the Database Engine or Analysis Services configuration pages.

Cluster Disk 4  
 Cluster Disk 5

Available shared disks:

Qualified	Disk	Message
	Cluster Disk 1	The disk resource 'Cluster Disk 1' cannot be used because it is a cluster quoru...
	Cluster Disk 2	The disk resource 'Cluster Disk 2' is already in use by resource 'SQL Server (PI...
	Cluster Disk 3	The disk resource 'Cluster Disk 3' is already in use by resource 'SQL Server'. T...
	Cluster Disk 4	
	Cluster Disk 5	

Refresh

< Back    Next >    Cancel    Help

nt



# MSDTC, SQL Server 2008/R2, and Windows Server 2008/R2

- OS supports multiple DTCs – no longer need to share one per Windows failover cluster

Use MSDTC instance installed to the local group

Else

Use the mapped instance of MSDTC

Else

Use the cluster's default instance of MSDTC

Else

Use the local machine's instance of MSDTC

# SQL Server Installation Gotchas, Tips, and Tricks 1

- Installing is a node by node, instance by instance affair
  - Solves occasional issues selected customers had with previous single-shot installers
  - One main install for each instance, all others add node operations
- Slipstream with SQL Server 2008 SP1 and beyond
  - W2K8 R2 requires SQL Server 2008 to be at SP1 anyway

# SQL Server Installation Gotchas, Tips, and Tricks 2

- Script if doing a lot of installations
  - Repeatable – saves a lot of time
  - Upgrade is scriptable starting with SQL 2K8 RTM
  - SQL 2K8 R2 new parameter:  
IACCEPTSQLSERVERLICENSETERMS
  - Don't put passwords in .ini files; designate at execution time

```
SETUP.EXE /ConfigurationFile="AddNode.ini"  
/SQLSVCPASSWORD="Password"
```

# SQL Server Installation Gotchas, Tips, and Tricks 3

- Instance name
  - Unique in the domain
  - Different than the Windows failover cluster name or the underlying node names
  - Becomes the network name
- Named instance name
  - Unique in the cluster

## SQL Server Installation Gotchas, Tips, and Tricks 4

- Make the instance ID the same as the clustered instance name

Specify a network name for the new SQL Server failover cluster. This will be the name used to identify your failover cluster on the network.

SQL Server Network Name: POWER

Default instance

Named instance: WINDOWS

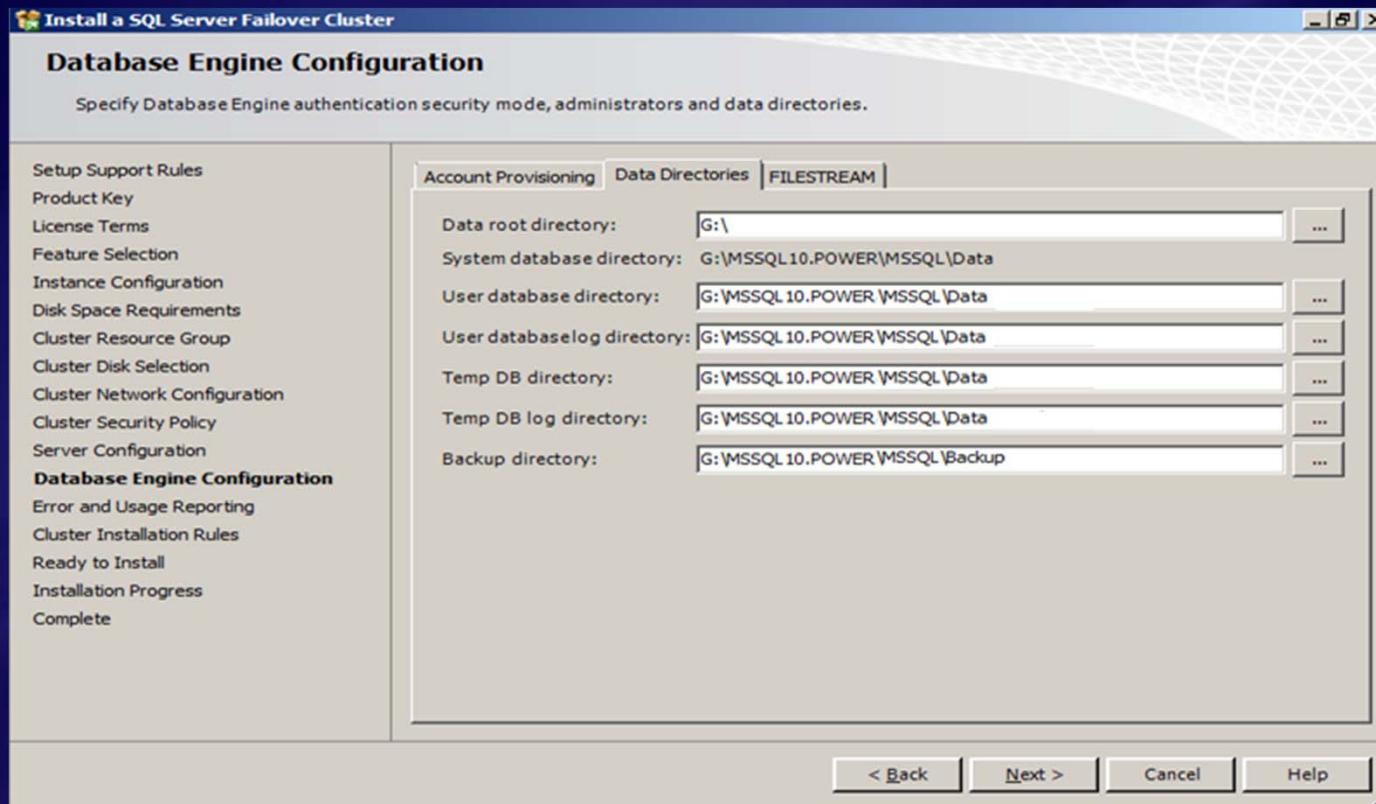
Instance ID: POWER\_WINDOWS

Instance root directory: C:\Program Files\Microsoft SQL Server\ ...

- Slipstream the installation
- All nodes must have the same local configuration
  - Install paths must exist on all nodes

# SQL Server Installation Gotchas, Tips, and Tricks 5

- Watch nested dialogs
- Can customize tempdb location



# Considerations for Multiple Instances 1

- Performance
  - Need to consider the failover condition
  - Memory and CPU
  - Dedicated failover node(s) –  $N+i$  configuration
  - Cannot use Resource Governor to control performance across instances
  - Account for OS overhead

# Considerations for Multiple Instances 2

- Memory
  - Worst case – need enough memory to run all instances on one node
  - Options
    - Hard maximum
    - Minimum
    - Completely dynamic
  - Only adjust with SQL Server
- Processor
  - Similar to memory
  - Can constrain in SQL Server to a degree
  - Use WSRM to constrain % of processor per SQL Server instance – set on each node



# Patching Clustered Instances of SQL Server 1

- Node-by-node process
  - Designed to increase reliability and minimize downtime
- Core goals:
  - When instance not patched, prevent failover to nodes being patched
  - Once instance patched, prevent failover to unpatched nodes
- Can patch multiple instances on a single node at the same time if desired
- Should be considered when planning the initial cluster

# Patching Clustered Instances of SQL Server 2

- Must remove at least half the nodes being patched from possible owners of the network name resource
  - NOT the SQL Server resource
- Add nodes back in after patched
  - Script the removal and addition of nodes
- **IMPORTANT:** maintain quorum
- Can also script the execution of the patch

# Patching Example 1 – One Instance, Two Nodes (1)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP1	No	Yes	SP1

# Patching Example 1 – One Instance, Two Nodes (2)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP1	No	<i>No</i>	SP1

- Remove Node B as a possible owner

# Patching Example 1 – One Instance, Two Nodes (3)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP1	No	No	<i>SP2</i>

- Install SP2 on Node B

# Patching Example 1 – One Instance, Two Nodes (5)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP1	No	<i>Yes</i>	SP2

- Add Node B back as a possible owner

# Patching Example 1 – One Instance, Two Nodes (6)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	<i>No</i>	Yes	SP1	<i>Yes</i>	Yes	SP2

- Fail instance to Node B to upgrade instance

# Patching Example 1 – One Instance, Two Nodes (7)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	No	<i>No</i>	SP1	Yes	Yes	SP2

- Remove Node A as a possible owner



# Patching Example 1 – One Instance, Two Nodes (8)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	No	No	<i>SP2</i>	Yes	Yes	SP2

- Install SP2 on Node A

# Patching Example 1 – One Instance, Two Nodes (9)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	No	<i>Yes</i>	SP2	Yes	Yes	SP2

- Add Node A back as a possible owner

# Patching Example 1 – One Instance, Two Nodes (10)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	<i>Yes</i>	Yes	SP2	<i>No</i>	Yes	SP2

- Fail instance to Node A
- Process complete

# Patching Example 2 – Two Instances, Two Nodes (1)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP1	No	Yes	SP1
Instance 2	No	Yes	SP1	Yes	Yes	SP1

# Patching Example 2 – Two Instances, Two Nodes (2)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP1	No	<i>No</i>	SP1
Instance 2	No	Yes	SP1	Yes	Yes	SP1

- Remove Node B as a possible owner from Instance 1

# Patching Example 2 – Two Instances, Two Nodes (3)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP1	No	No	<i>SP2</i>
Instance 2	No	Yes	SP1	Yes	Yes	<i>SP1*</i>

- Install SP2 for shared files and Instance 1 on Node B

# Patching Example 2 – Two Instances, Two Nodes (4)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP1	No	No	SP2
Instance 2	<i>Yes</i>	Yes	SP1	<i>No</i>	Yes	SP1*

- Fail Instance 2 over to Node A

# Patching Example 2 – Two Instances, Two Nodes (5)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP1	No	<i>Yes</i>	SP2
Instance 2	Yes	Yes	SP1	No	Yes	SP1*

- Add Node B back as a possible owner for Instance 1
- Reboot Node B



# Patching Example 2 – Two Instances, Two Nodes (6)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	<i>No</i>	Yes	SP1	<i>Yes</i>	Yes	SP2
Instance 2	Yes	Yes	SP1	No	Yes	SP1*

# Patching Example 2 – Two Instances, Two Nodes (7)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	No	<i>No</i>	SP1	Yes	Yes	SP2
Instance 2	Yes	Yes	SP1	No	Yes	SP1*

- Remove Node A as a possible owner of Instance 1

# Patching Example 2 – Two Instances, Two Nodes (8)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	No	No	<i>SP2</i>	Yes	Yes	SP2
Instance 2	Yes	Yes	<i>SP1*</i>	No	Yes	SP1*

- Install SP1 on Node A against Instance 1 and the shared files

# Patching Example 2 – Two Instances, Two Nodes (9)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	No	<i>Yes</i>	SP2	Yes	Yes	SP2
Instance 2	Yes	Yes	SP1*	No	Yes	SP1*

- Add Node A back as a possible owner of Instance 1

# Patching Example 2 – Two Instances, Two Nodes (10)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	<i>Yes</i>	Yes	SP2	<i>No</i>	Yes	SP2
Instance 2	<i>No</i>	Yes	SP1*	<i>Yes</i>	Yes	SP1*

- Fail Instance 2 to Node B to facilitate a reboot
- Reboot Node A
- Fail Instance 1 back to Node A to complete its update

# Patching Example 2 – Two Instances, Two Nodes (11)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP2	No	Yes	SP2
Instance 2	No	<i>No</i>	SP1*	Yes	Yes	SP1*

- Remove Node A as a possible owner of Instance 2

# Patching Example 2 – Two Instances, Two Nodes (12)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP2	No	Yes	SP2
Instance 2	No	No	<i>SP2</i>	Yes	Yes	SP1*

- Install SP2 on Node A against Instance 2

# Patching Example 2 – Two Instances, Two Nodes (13)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP2	No	Yes	SP2
Instance 2	No	<i>Yes</i>	SP2	Yes	Yes	SP1*

- Add Node A back as a possible owner of Instance 2



# Patching Example 2 – Two Instances, Two Nodes (14)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP2	No	Yes	SP2
Instance 2	<i>Yes</i>	Yes	SP2	<i>No</i>	Yes	SP1*

- Fail Instance 2 to Node A to upgrade the instance

# Patching Example 2 – Two Instances, Two Nodes (15)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP2	No	Yes	SP2
Instance 2	Yes	Yes	SP2	No	<i>No</i>	SP1*

- Remove Node B as a possible owner of Instance 2

# Patching Example 2 – Two Instances, Two Nodes (16)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP2	No	Yes	SP2
Instance 2	Yes	Yes	SP2	No	No	<i>SP2</i>

- Install SP2 against Instance 2 on Node B

# Patching Example 2 – Two Instances, Two Nodes (17)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP2	No	Yes	SP2
Instance 2	Yes	Yes	SP2	No	<i>Yes</i>	SP2

- Add Node B back as a possible owner of Instance 2

# Patching Example 2 – Two Instances, Two Nodes (18)

	Node A			Node B		
	Owner	Possible Owner	Version	Owner	Possible Owner	Version
Instance 1	Yes	Yes	SP2	No	Yes	SP2
Instance 2	<i>No</i>	Yes	SP2	<i>Yes</i>	Yes	SP2

- Fail Instance 2 over to Node B

# Summary

- Everything is different ... and the same for clusters both in Windows and SQL Server
- Cluster validation is the proverbial king
- The failover condition is key to success
- Use scripts where possible
- Patching cluster nodes must be thought about
- Learn PowerShell

# Resources 1

- [Pro SQL Server 2008 Failover Clustering \(Apress, 2009\)](#)
- [My blog](#) – updated regularly with new stuff!
- [SQL Server 2008 failover clustering whitepaper](#)
- [Creating a Clustered DTC with PowerShell](#)
- [Server Virtualization Validation Program](#)
- [Blog post – Slipstreaming SQL Server 2008 + SP](#)
- [Blog post – Slipstreaming SQL Server 2008 + SP + CU](#)
- [Windows Clustering Team Blog](#)

# Resources 2

- KB Articles:

- [956893 “Support policy for Microsoft SQL Server products that are running in a hardware virtualization environment”](#)
- [897615 Support policy for Microsoft software running in non-Microsoft hardware virtualization software](#)
- [943984 The Microsoft Support Policy for Windows Server 2008 Failover Clusters](#)
- [327518 The Microsoft SQL Server support policy for Microsoft Clustering](#)
- [958734 SQL Server 2008 failover clustering rolling patch and service pack process](#)



## Conditions and Terms of Use

This training package content is proprietary and confidential, and is intended only for users described in the training materials. Content and information designated for limited distribution is provided to you under a Non-Disclosure Agreement and cannot be distributed. Copying or disclosing all or any portion of the content and/or information included in such packages is strictly prohibited.

The contents of this package are for informational and training purposes only and are provided "as is" without warranty of any kind, whether express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, and non-infringement.

Training package content, including URLs and other Internet Web site references, is subject to change without notice. Because Megahirtz LLC must respond to changing market conditions, the content should not be interpreted to be a commitment on the part of Megahirtz LLC, and Megahirtz LLC cannot guarantee the accuracy of any information presented after the date of publication. Unless otherwise noted, the companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred.

## Copyright and Trademarks

© 2011 Megahirtz LLC All rights reserved.

Megahirtz LLC may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in written license agreement from Megahirtz LLC, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Megahirtz LLC.

All products mentioned are either registered trademarks or trademarks for their respective manufacturers and owners in the United States and/or other countries.

