

SQL Server Clustering for Dummies



Mark Broadbent
@retracement

mark.broadbent@sqlcloud.co.uk

tenbulls.co.uk

LinkedIn: <http://linkd.in/gylT4l>



Presentation and materials available from: <http://bit.ly/gAqN32>

A bit more about me!

- More than 20 years in IT and more than 14 years using SQL Server. Have worked my many large global corporations and SMEs such as Microsoft, Nokia, Hewlett Packard and Encyclopaedia Britannica.
- Presented at SQLBits 7 and due to speak at SQLRally 1 in Orlando
- Microsoft Certified IT Professional Database Development SQL 2008
- Microsoft Certified IT Professional Database Administrator SQL 2008
- Microsoft Certified IT Professional Database Administrator SQL 2005
- Microsoft Certified Application Developer (C# .net)
- Microsoft Certified Database Administrator (SQL 2000)
- Microsoft Certified Systems Engineer + Internet
- Participate on #sqlhelp ,MSDN Forums, Stackoverflow & Serverfault.
- Used to be active in the MS newsgroups until their demise :(
- Run the LinkedIn groups
- Linux Mint User Group <http://www.linkedin.com/groups?gid=2989801>
- SQL Server Scripting <http://www.linkedin.com/groups?gid=3033621>

Agenda

- Introduction to Clustering ...What, why and who?
- DEMO
- Installation ...Preparation, Validation and Installation
- Administration ...Problems and benefits
- DEMO
- Houston ...we have a problem

Introduction

- What is a SQL Cluster?
 - One or more “clustered” SQL Instances on one or more physical servers running MS Windows
- Why use one?
 - Single point of failure is the OS
 - Is Virtualization protecting the Virtual OS?
- Who supports?
 - Diverse range of knowledge required, perhaps too much for a specialist DBA?
 - How do we monitor? Not all tools are cluster aware
 - Who controls the Cluster?
- Is Clustering right for us?
 - Management expectation is that Clustering is 100% HA. It is not
 - Are skills in place, if not HA = Hopefully Available
 - Do your applications recover from unavailability?

Terminology

Terms can mean the same thing or sometimes tend to be used interchangeably (often wrongly!) ...including by me. Depending upon who you talk to they can cause confusion. They include:-

- **Cluster**
- **Cluster Node**...Node...Server
- *Service or Application... Virtual Server... **Cluster Group**... Failover Cluster... SQL Instance... SQL Server... and *ahem* Server*
- *Active/Passive... **Single Instance**... **Multi Instance** ... N+n Cluster*
- **Resources** (take your pick)
- **Private Network**...interconnect...internal
- **Public Network**...LAN (or whatever network is the one clients will connect through facing!)
- SAN... iSCSI target... **Shared Storage**... LUN... Partition... Volume... Disk
- **Quorum**... Voting Disk... Majority

A Dummies Dictionary

- Cluster –Connected Windows servers running Cluster service with the ability to own the Cluster Name and IP
- Cluster Node –A Windows server that is Clustered
- Cluster Group –Collection of clustered resources
- Single Instance Cluster –One SQL Instance installed to a Cluster containing one or more Cluster Nodes
- Multi Instance Cluster –Come on people!
- Resources –those things in the Cluster Groups remember!
- Public/ Private/ Storage networks –logical networks across Clustered nodes (each NIC configured in own subnet)
- Quorum –voting mechanism to ensure correct ownership of shared resources. E.g. Magician is “No majority: Disk only” style Quorum.

Failover Cluster Manager

The screenshot shows the Failover Cluster Manager console for a cluster named 'magrathea.hhgg.com'. The left-hand tree view is annotated with callouts: 'Cluster Name' points to the cluster name, 'Cluster Groups' points to the 'Services and applications' group, 'Cluster Nodes' points to the 'babel' node, 'Events' points to 'Cluster Events', and 'Cluster Networks' points to the 'Private Network'. The right-hand pane shows two detailed views: 'Disk Witness in Quorum' and 'Network Connections'.

Cluster Name: magrathea.hhgg.com

Cluster Groups:

- Services and applications
 - marvin (SQL01) instance
 - zaphod (SQL02) instance
- Nodes
 - babel
 - boghog
 - fish
- Storage
- Networks
 - Private Network
 - Public Network
 - Storage Network
- Cluster Events

Disk Witness in Quorum

Disk	Status	Current Owner
Disk Witness in Quorum		
Quorum	Online	fish
marvin (SQL01) instance		
Data Mount Point	Online	fish
Disk D:	Online	fish
MSDTC	Online	fish
zaphod (SQL02) instance		
Disk G:	Online	boghog

Shared Storage

Cluster Networks

Name	Status	Current Owner
Network Connections		
babel - private	Up	babel
boghog - private	Up	boghog
fish - private	Up	fish

DEMO: Intro to the Failover Cluster

- Network Name/ IP
- SQL Services
- Dependency
- Possible Owner and Preferred Owner
- Remote Desktop

Installation

- Preparation

- Remove Disabled Adapters to avoid "ghosted NICs"
- Ensure .NET Framework 3.5 SP1 is installed
- Disable Firewall on private network
- "Before Installing Failover Clustering" <http://bit.ly/dEZc6H>
- "Windows Server 2008 Failover Cluster Configuration Program" <http://bit.ly/fldSrK>

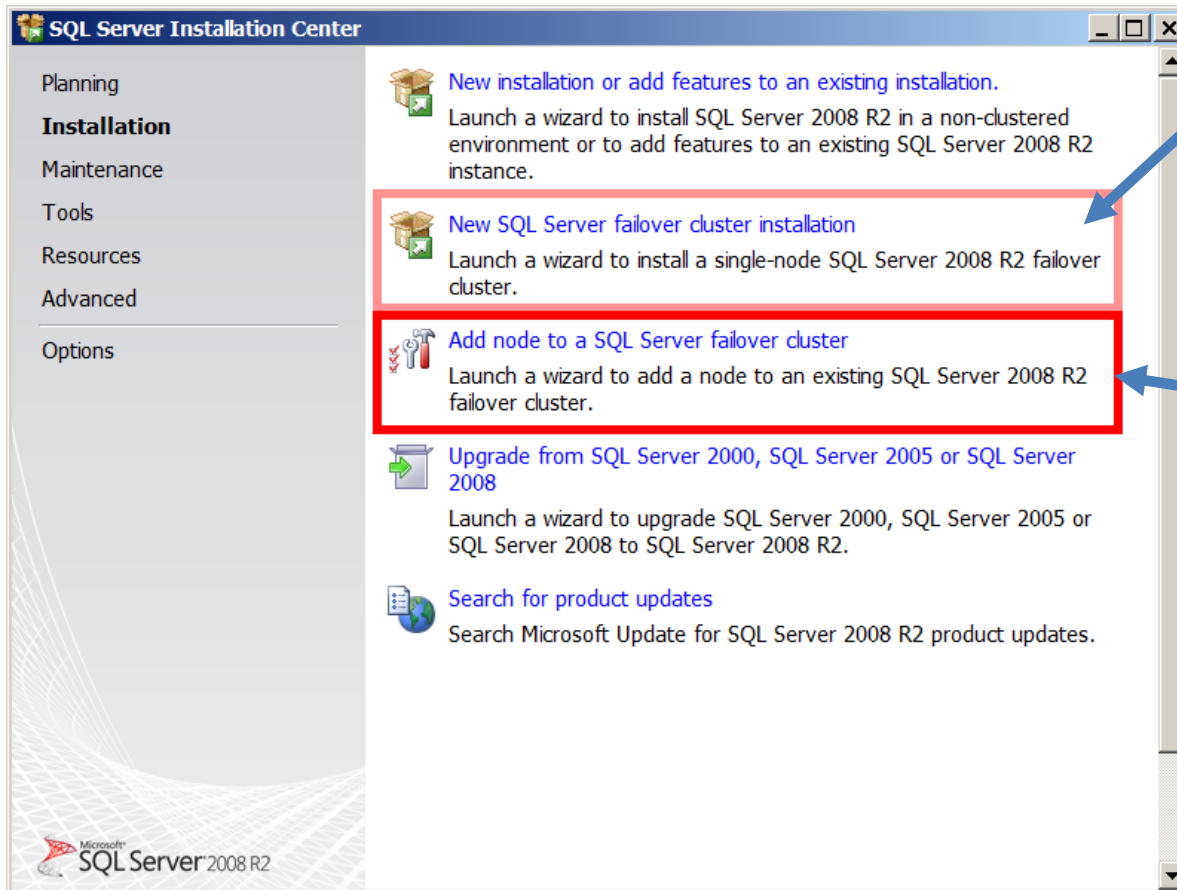
- Validation

- HCL no longer exists, although there are pre-validated vendor solutions
- Successful validation is required
- "Microsoft SQL Server support policy for Microsoft Clustering" <http://bit.ly/a7yDok>
- "Microsoft Support Policy for Windows Server 2008 or Windows Server 2008 R2 Failover Clusters" <http://bit.ly/f17wew>
- "Understanding Cluster Validation Tests" <http://bit.ly/efBIHp>

- Installation

- One SQL installation per instance per node
- "Installing a SQL Server 2008 R2 Failover Cluster" <http://bit.ly/gmXyip>
- Download my Installation Videos from <http://bit.ly/gAqN32>

Installation



Create Cluster Group and Install SQL instance to it
This is the “SQL Failover Cluster”

Install SQL instance onto Cluster node and join into existing SQL Failover Cluster

Administration for the DBA

- Storage
 - Must be available otherwise instance will fail
 - Filesystem corruption repaired through maintenance mode
- Co-existence
 - Battle for server resources will occur. Provision dedicated node or avoid co-existence
 - Memory. What is your Min and Max Memory?
 - Processor. Solutions could be Resource Governor, WRSM, Processor Affinity and Dedicated failover node/s
- Dependency, Current owner, Preferred Owners and Possible Owners
 - Do not change the default dependencies of the Clustered SQL instance services
 - Ticked Preferred Ownership sets order, un-ticked still potentially possible owner
 - Only Possible Owner (set on resource) will be allowed to own resource
- SQL Patching
 - Rolling updates is possible reducing downtime to potentially a single failover
- Monitoring
 - Is cluster service running on each node? If it is stopped then node is offline in the cluster
 - Can you ping Cluster Group network name and IP. If it's inaccessible then Cluster Group & SQL is be down
 - Monitor the SQL Services within the Cluster group. If they are stopped then so is SQL!
- IP Address, port, network name and instance.
 - Ensure you set a fixed port and avoid using DHCP
 - Can have one default instance per cluster. Cluster Group needs unique network name
 - Instance names must be unique across the Cluster.

DEMO: More SQL Failover Cluster

- Automatic Failover
 - Stopping, starting and pausing nodes
 - Network Failure
- Adding Disks
- PowerShell
- Querying events

Danger Will Robinson!

- SSRS and SSIS are not Cluster aware
 - Put SSRS dbs on Clustered SQL Instance. Use Standalone SSRS web servers and Scale-out using NLB
 - I prefer to install SSIS ONLY if you can police, control and regulate the package deployments
- SSAS should be installed to it's own Cluster Group
 - SSAS is Cluster aware
 - Resource hungry, allocate to dedicated Cluster Nodes
- Be careful with Mirrored Databases in High Safety with Automatic Failover
 - A Cluster Group has failover could cause Database Mirror failover. Is this what you want?
 - Can change partner timeout by
ALTER DATABASE *dbname* SET PARTNER TIMEOUT x
however Cluster Group failover time can vary
- Failover time can be substantially affected by
 - Database Recovery time and performance of servers
- MSDTC configuration is different to a standalone
 - Ensure it is configured correctly
 - Have a proven tested way to test it
- Keep Dummies, developers (in fact anyone but you) away from production!



Disaster Planning

“Hang on lads, I've got a great idea.”



- Is the Shared Storage accessible?
- Can you make storage available elsewhere?
- Do you have Backups?
- What is the restore time?
- Have you available instances?
- What about name resolution?
- Can you redirect your apps?
- Will you maintain same performance level?
- What are our SLAs?

Shouldn't you have thought about this already?

In summary

- We discussed
 - What is a SQL Failover Cluster
 - Overview of the Cluster Group
 - Installation, preparation, and validation
 - Administration of the SQL Failover Cluster
 - Recovery from failure

Further references

- Books
 - Apress - Pro SQL Server 2008 Failover Clustering – Allan Hirt
- Blogs/ Websites
 - MSDN Clustering and High-Availability <http://bit.ly/czSTPz>
 - Allan Hirt blog <http://bit.ly/epAoPj>
 - Mark Broadbent blog/ clustering <http://bit.ly/i8knch>
 - Geoff Hiten blog <http://bit.ly/cd8aCA>
 - David Bermingham blog <http://bit.ly/hYmwXU>
 - Symon Perriman blog <http://bit.ly/eyGa3U>
- Video/ Webcasts
 - SQLBits 8: Multi-site Failover Cluster With SQL 2008 and Denali – Allan Hirt
 - Thinking outside the Box <http://bit.ly/eUe3v5> - Mark Broadbent
 - PASS Summit 2010: DBA324 Clustering for mere mortals - Geoff Hiten
 - TechNet Webcast: Failover Cluster Validation and Troubleshooting with Windows Server 2008 (Level 300) <http://bit.ly/hx8xCl>
 - TechNet Webcast: Windows Failover Clustering PowerShell Cmdlets: Overview and Examples (Level 400) <http://bit.ly/fNBF4O>
 - TechNet Webcast: Failover Clustering 101 (Level 300) <http://bit.ly/fmCOqc>
 - TechNet Webcast: High Availability Server Apps with Windows Server 2008 R2 Failover Clusters (Level 300) <https://msevents.microsoft.com/CUI/WebCastEventDetails.aspx?culture=en-US&EventID=1032448854&CountryCode=US>

SQL Server Clustering for Dummies

THANKS FOR COMING!