

O31

SQL Server High Availability

Mike Shelton

IBM @server xSeries
Technical Conference

Aug. 9 - 13, 2004

Chicago, IL

Overview

- Defining High Availability
- Setting High Availability Goals and Identifying Barriers
- SQL Server 2000 HA Technology
 - Failover Clustering
 - Log Backup Shipping
 - Replication
- High Availability Operations and Support

Defining Five Nines

- What Is 99.999%?
 - Target:
 - Online and available to users 24 hours a day, 365 days a year
 - Total outages less than 5.26 minutes per year
- Is Five Nines the Only Option?
 - Four Nines: 52.6 minutes of downtime a year
 - Perhaps 12 hours a day, 5 days a week
 - Allows for planned downtime: maintenance, etc.

Availability



Un-managed

Well-managed Nodes

Masks Some Hardware Failures

Well-managed Packs and Clones

Masks Hardware Failures

Masks Operations Tasks (e.g. Software Upgrades)

Masks Some Software Failures

Well-managed Geoplex

Masks Site Failures (Power, Network, Fire, Move...)

Masks Some Operations Failures

Business Makes High Availability Necessary

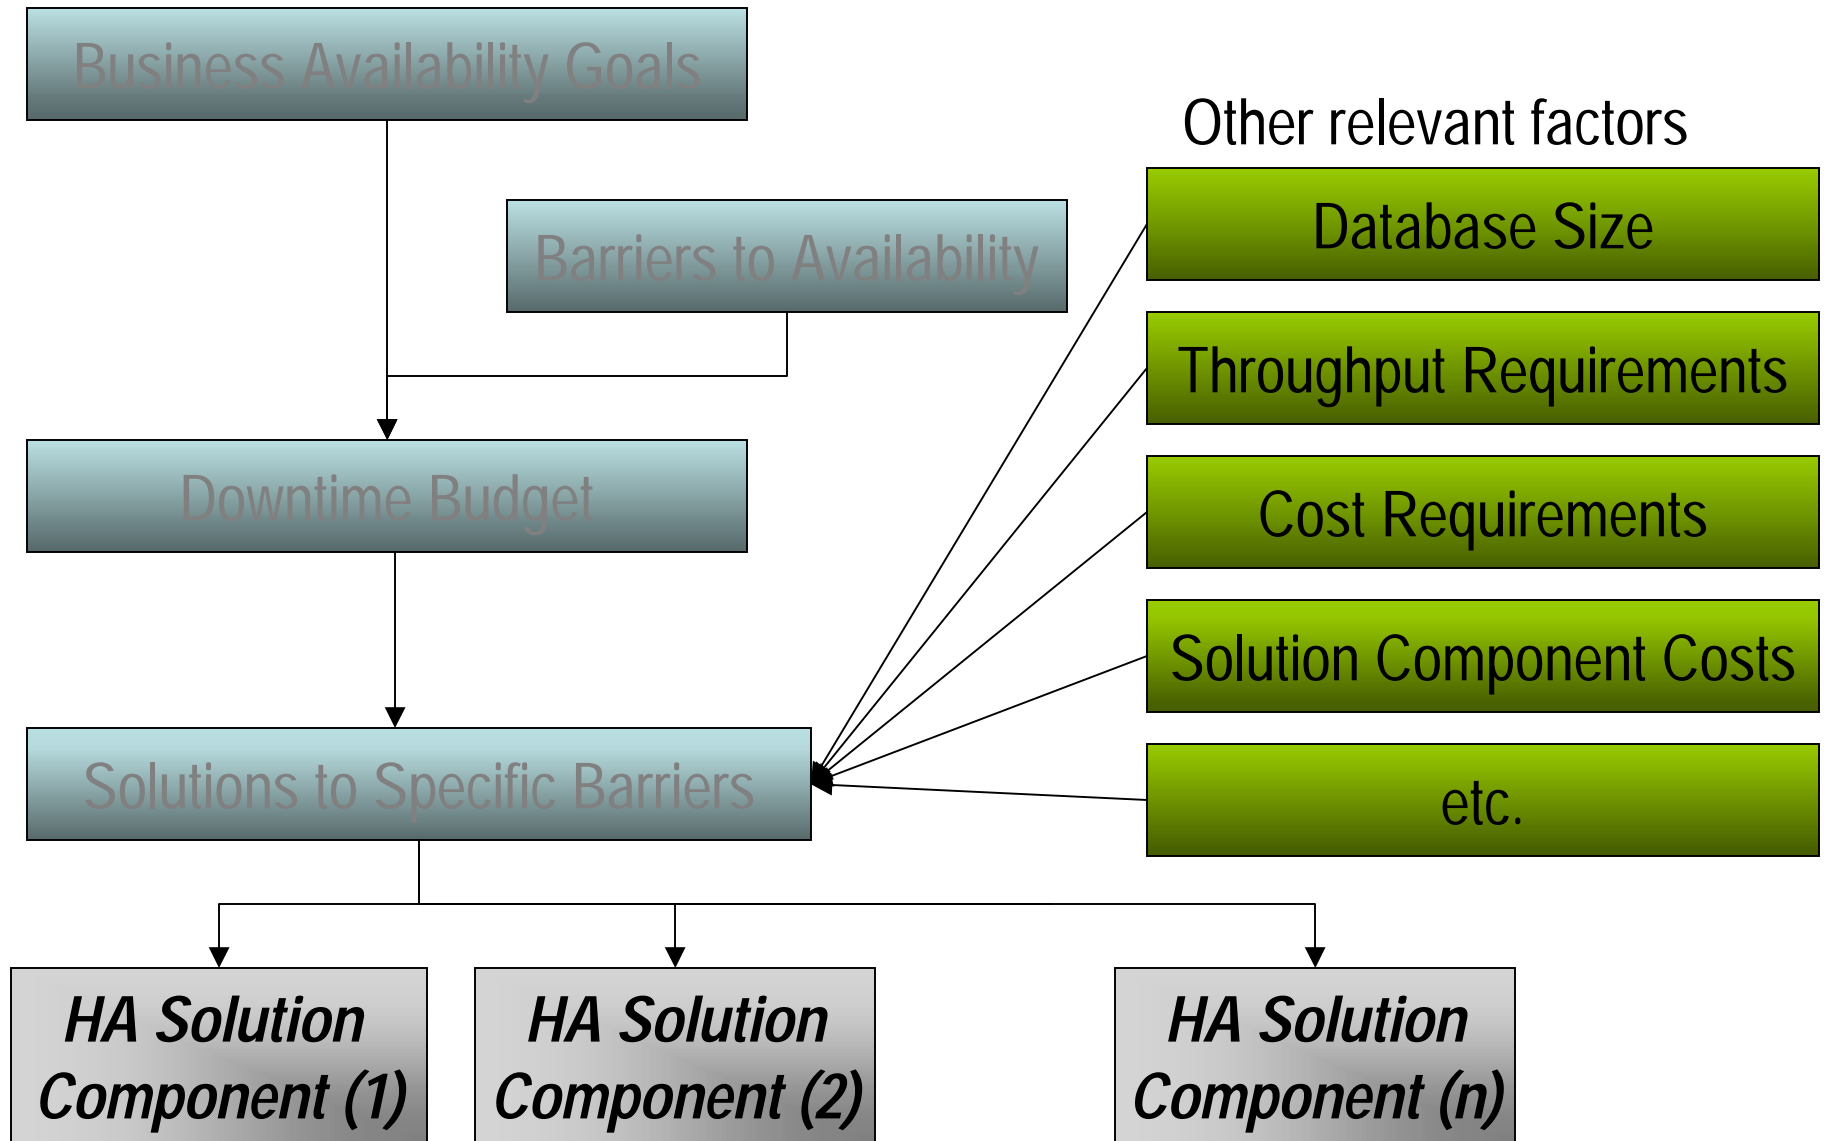
- Reliance on technology
 - Hospital has high cost for downtime (lives!)
- Product or service availability
- Continuous improvement of products, services, and processes
 - Example: It would be a failure if “the person to call” left the company two years ago and nobody can currently offer expertise.
- DBA interest in continuous improvement of products, services, and processes

What High Availability Is Not

- A Technology Solution From a Vendor
- A Scalability Solution
- An IT Decision Without Business Knowledge
- A Business Decision Isolated From the Cost of Downtime

- It is:
 - A solution involving people and process, and, very likely, technology

High Availability Framework



Setting High Availability Goals

- Identify Stakeholders
- Value to the Business
- How the System Is Used
- Cost Limitations
- Never Lead With Technology


Identifying Barriers

- A Barrier Impacts a System's Availability
- Use Monitoring to Identify Barriers
- Any High Availability Element Can Become a Barrier
 - Hardware
 - Communication/connectivity
 - Environmental
 - Services
 - Software
 - Process
 - Application or User Error
 - Staffing

Calculating the Time and Cost of High Availability

- Cost of Downtime
- Probability of Occurrence
- Cost to Remove Barriers

Total Cost of Ownership Factors



Microsoft SQL Server 2000 High Availability Technology

- SQL Server 2000 Editions Suitable for High Availability
- Comparison of Standby Options
- High Availability Features
- Failover Clustering
- Log Shipping
- Transactional Replication
- Using Combinations of Technologies

SQL Server 2000 Editions Suitable for High Availability

- Enterprise Edition
 - Most Scalable and Highly Available
 - Includes Failover Clustering
 - Includes Log Shipping Features
 - Suitable for Production
- Developer Edition
 - Full Featured (Same As Enterprise Edition)
 - Suitable for Development and Testing

Comparison of Standby Options

- Hot Standby
- Warm Standby
- Cold Standby

High Availability Features

- Standby Type
- Failure Detection
- Automatic Failover
- Masks Disk Failure
- Masks SQL Process Failure
- Masks Other Process Failure
- Meta Data Support
- Transactionally Consistent
- Transactionally Current
- Perceived Downtime
- Transparent to Client
- Special Hardware Needed
- Distance Limit
- Complexity
- Standby Accessible
- Impact on Performance
- Impact on Backup Strategy


Failover Clustering

- Types of Clusters
- Windows Clustering
- SQL Server 2000 Failover Clustering
- How Failover Clustering Works
- Enhancements to Failover to Clustering
- High Availability Features in Failover Clustering

Types of Clusters

- Windows Cluster
- Failover Cluster
- Federated Cluster
- Network Load Balancing Cluster

Failover Clustering

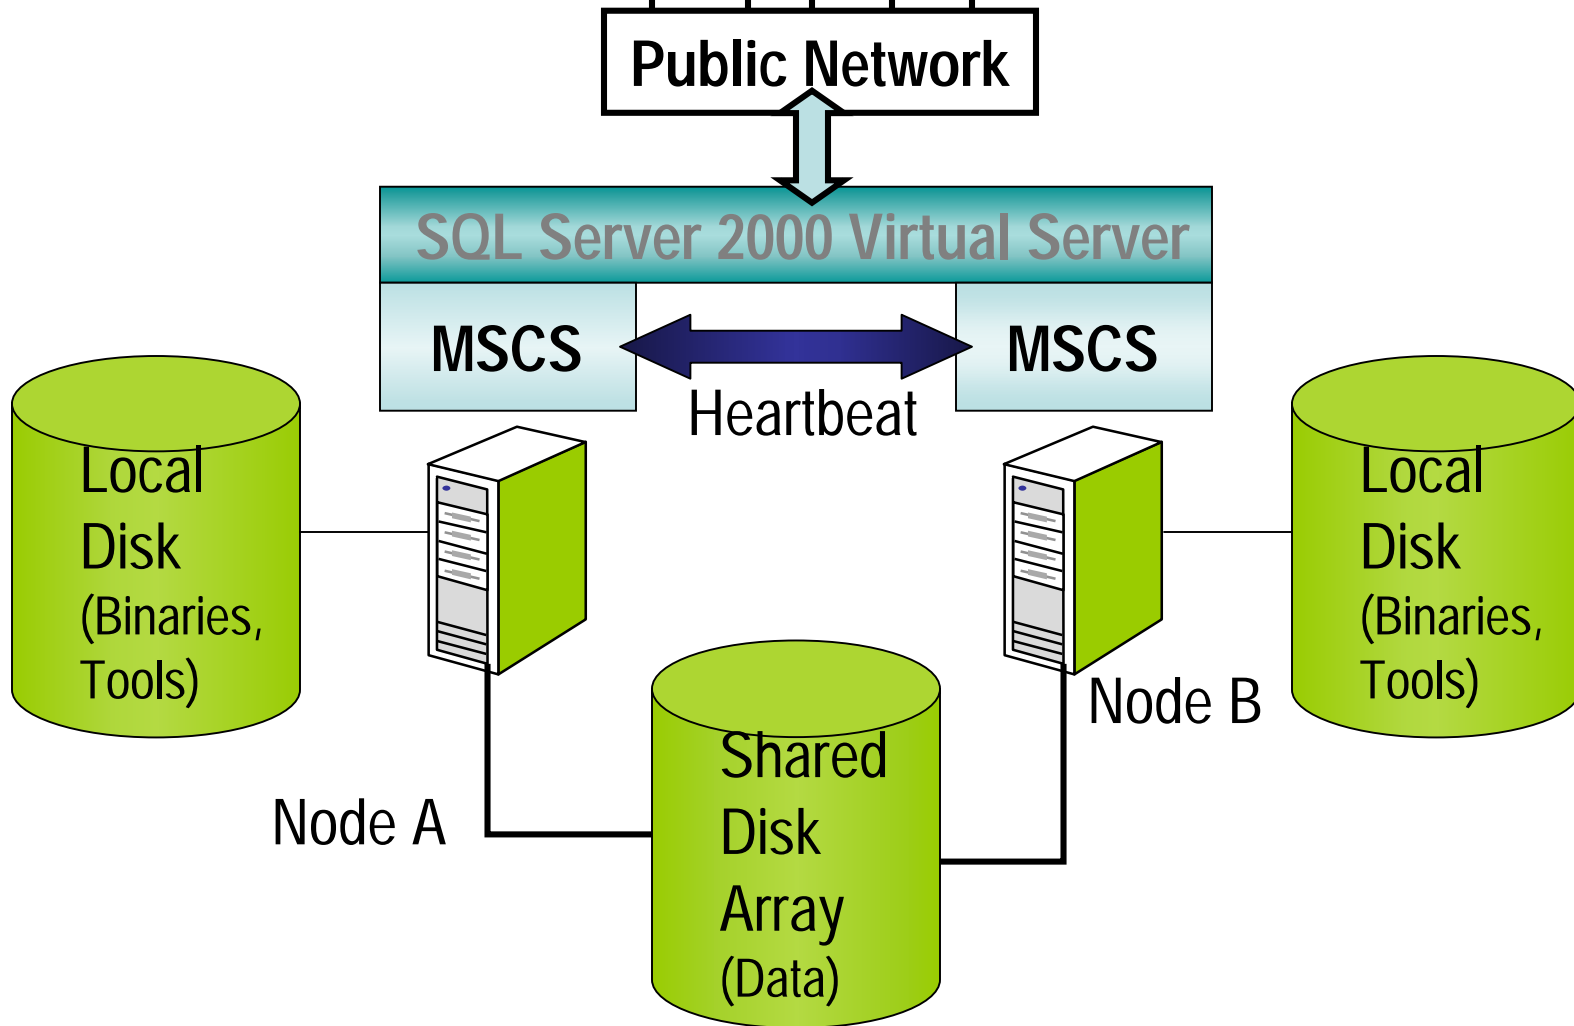


Windows Clustering

- Hardware Components
 - Cluster node
 - Heartbeat
 - External network
 - Shared cluster disk array
 - Quorum drive
- Software Components
 - Cluster name
 - Cluster IP address
 - Cluster administrator account
 - Cluster resource
 - Cluster group
- Virtual Server

SQL Server 2000 Failover

Clustering



How Failover Clustering Works

- Operating System Checks
 - Heartbeat Checks Availability of Nodes and Virtual Servers
- SQL Server Checks
 - Looks-alive Check Every 5 Seconds
 - IsAlive check runs `SELECT @@VERSION` query
- Failover to Another Node
 - Windows Clustering Attempts Restart on Same Node or Fails Over to Another Node
 - SQL Server Service Starts
 - Brings master Online
 - Database Recovery Proceeds
 - End Users and Applications Must Reconnect

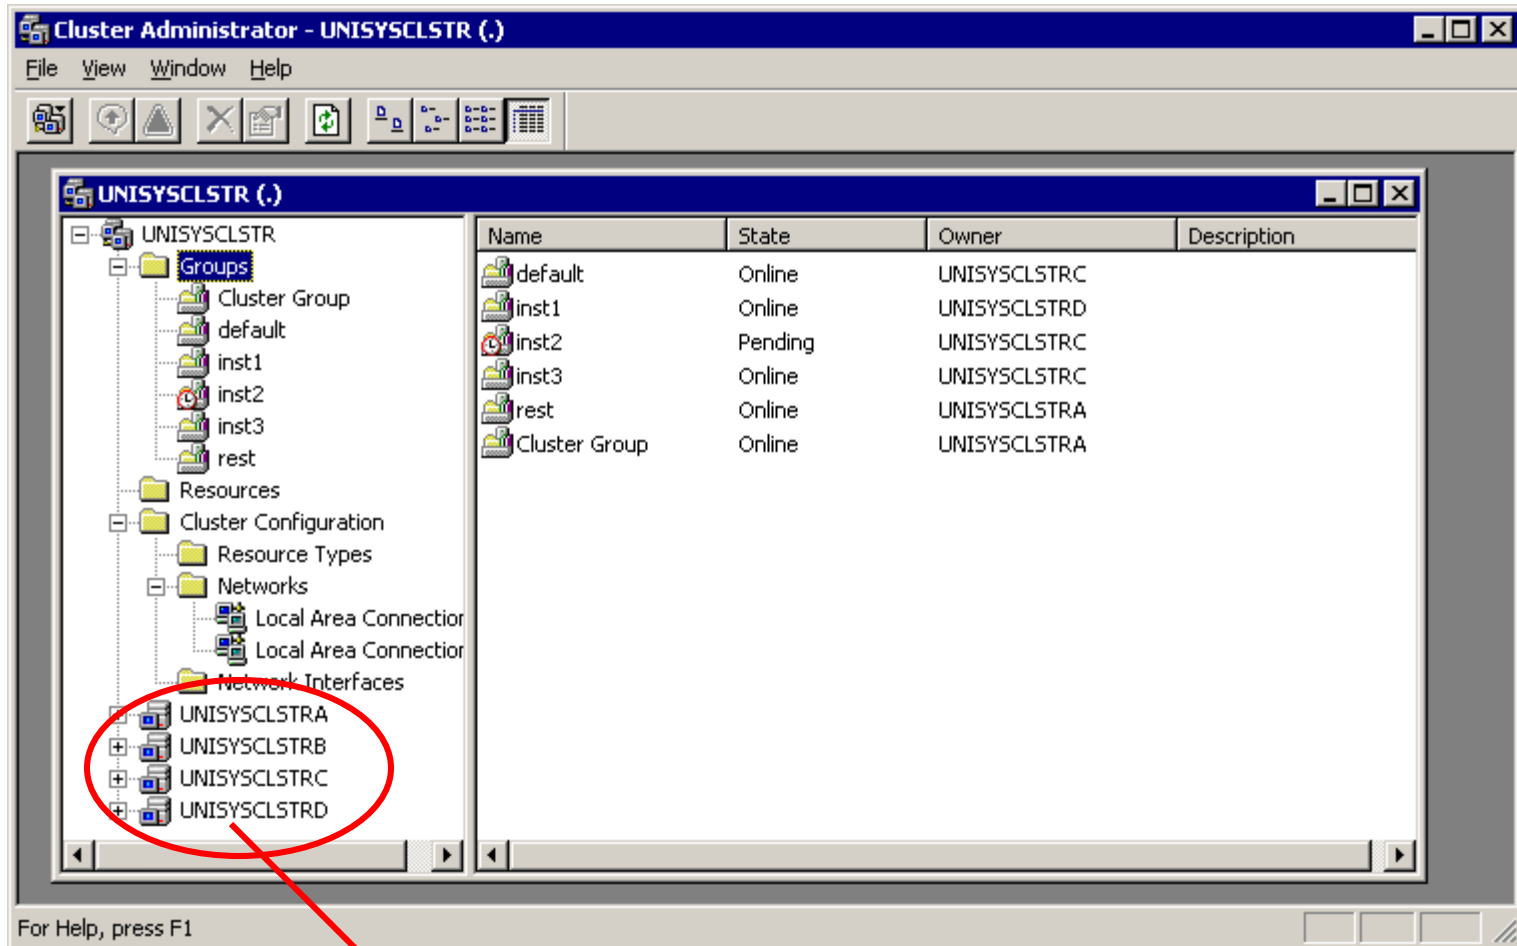
Enhancements to Failover Clustering

- SQL Server Setup Installs/Uninstalls a Cluster
- Service Packs Can be Applied Directly to Virtual Servers
- SQL Server Supports Multiple Instances and Multiple Network Addresses
- Failover and Failback to or From Any Node in a Cluster
- SQL Server 2000 on Windows 2000 Datacenter Server Supports 4 Server Nodes in a Cluster
- All Nodes Have Local Copies of SQL Server Tools and Executables
- Rerunning the Setup Program Updates Failover Cluster Configurations
- SQL Server Service Manager or SQL Server Enterprise Manager Now Start/Stop SQL Server Services

Failover Clustering Terminology

- “Multiple Active Instances” or “Single Active Instance”
- No longer “Active/Active” or “Active/Passive”

Cluster Administrator



4-nodes

High Availability Features in Failover Clustering (1 of 2)

Availability Feature	Failover Clustering
Standby Type	Hot
Failure Detection	Yes
Automatic Failover	Yes
Masks Disk Failure	No; Shared Disk
Masks SQL Process Failure	Yes
Masks Other Process Failure	Yes
Meta Data Support	All System and Database
Transactionally Consistent	Yes
Transactionally Current	Yes, Always Up to Date

High Availability Features in Failover Clustering (2 of 2)

Availability Feature	Failover Clustering
Perceived Downtime	30 Seconds + DB Recovery
Transparent to Client	Yes, Reconnect to Same IP
Special Hardware Needed	Specialized Hardware from Cluster HCL
Distance Limit	100 Miles
Complexity	More
Standby Accessible	Standby never accessible
Impact on Performance	No Impact
Impact on Backup Strategy	Must be able to backup from any node

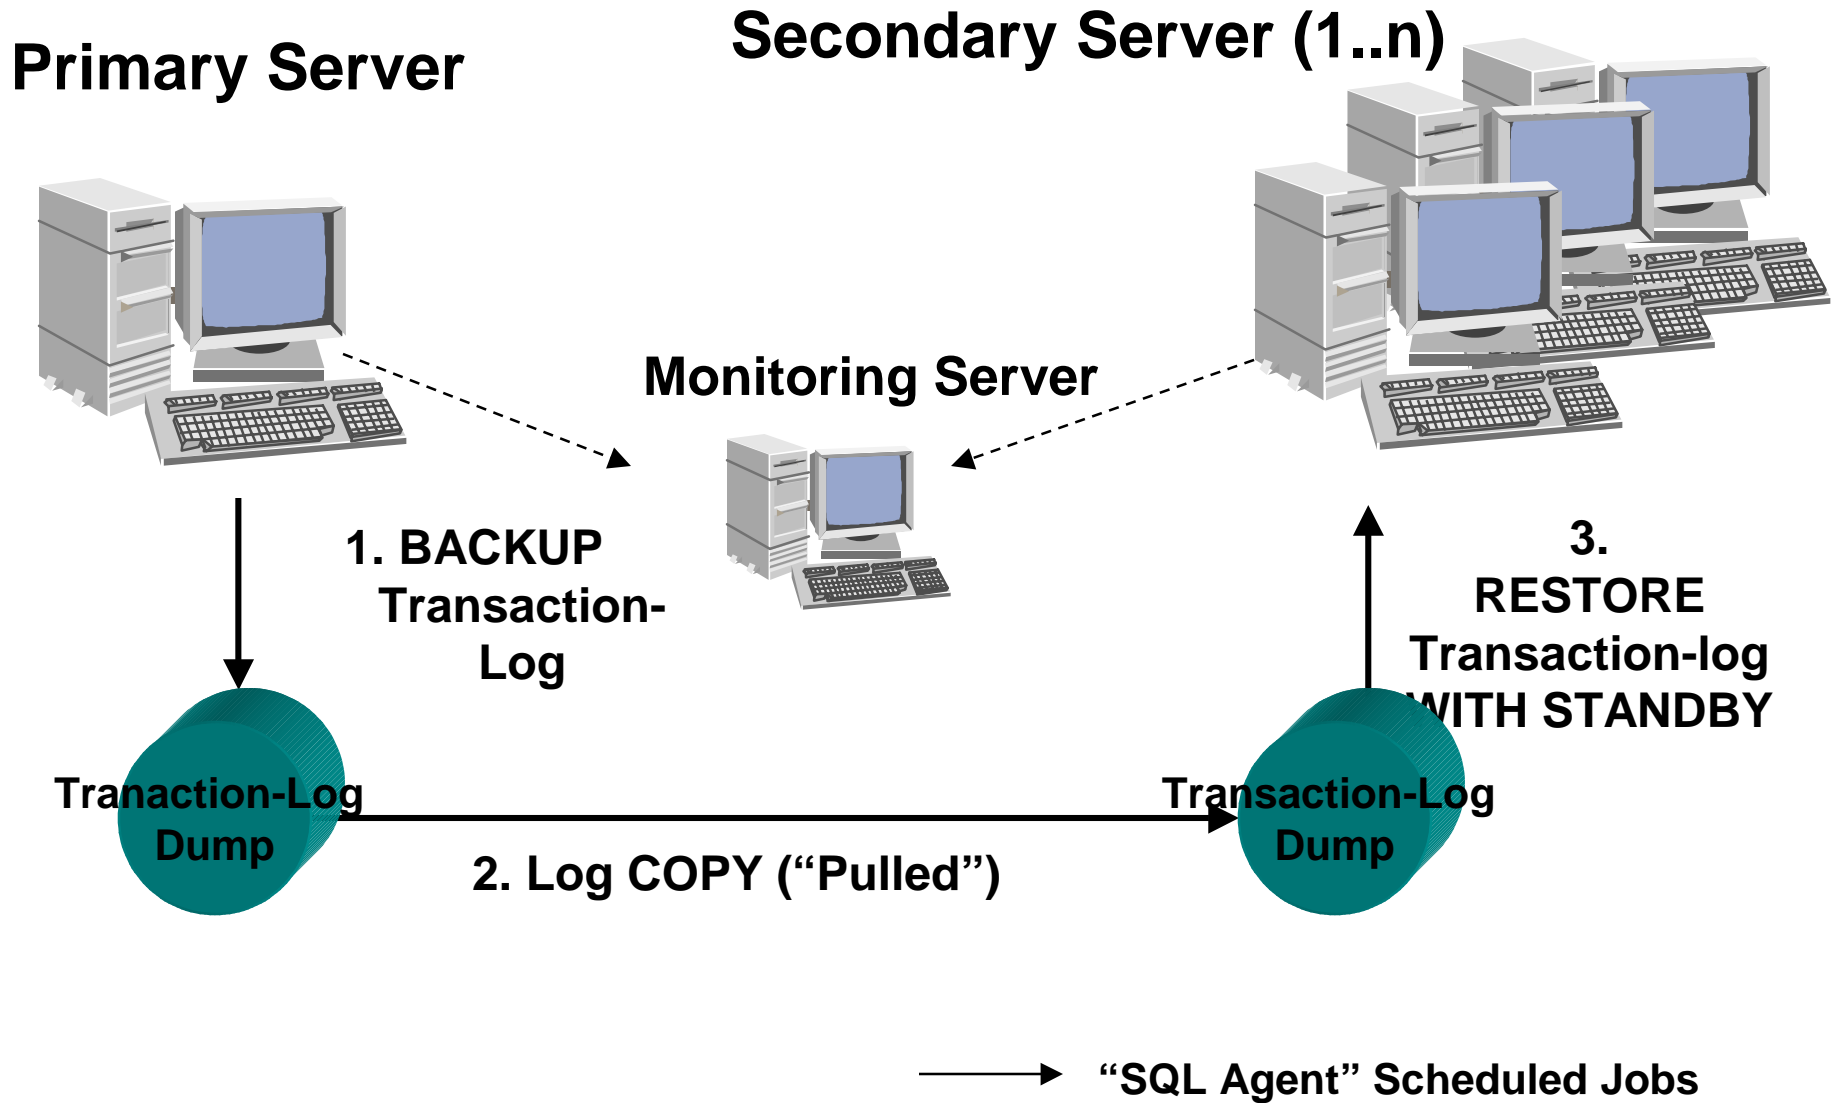
SQL Server Failover Clustering

- Hot Standby Solution
- Best High Availability Configuration
 - Redundant System
 - Shared Access to the Database Files
 - Recovery in Seconds
 - Automatic Failure Detection
 - Automatic Failover
 - Minimal Client Application Awareness
- Built on Microsoft Cluster Server

Log Shipping

- Warm Standby Solution
- Applies Transaction Log From Primary Server (Primary) to Warm Standby (Secondary)
- Attributes of Log Shipping
 - Warm Standby Available for Limited Read-Only Use
 - All Logged Schema and Data Changes Applied
 - Cannot Filter Changes for Partitioning or Subsets
- Manual Failure Detection; Manual Failover
- For Yukon we call this “Log Backup Shipping”
 - Differentiate it from Real-time Log Shipping

Log Shipping Architecture



Log Shipping Monitor

The screenshot displays the SQL Server Enterprise Manager interface. The left pane shows the server tree with 'Log Shipping Monitor' selected under the 'Management' folder of the 'Test' server. The right pane lists 37 items, including various databases and their log shipping pairs. A 'Log Shipping Pair Properties' dialog box is open, showing details for a pair between 'LOGSHIPPING1.Finance' and 'LOGSHIPPING2.Finance'.

Log Shipping Pair Properties

Status | Source | Destination

LOGSHIPPING1.Finance → LOGSHIPPING2.Finance

Status: Secondary out-of-sync

Current date / time on Monitor server: Wednesday, February 09, 2000

Last backup file: Finance_tlog_200002091600.TRN

Updated: 2/9/2000 4:00:14 PM

Backup delta: 1 minutes

Last file copied: Finance_tlog_200002091600.TRN

Updated: 2/9/2000 4:00:19 PM

Copy delta: 1 minutes

Last file loaded: Finance_tlog_200002091600.TRN

Updated: 2/9/2000 4:00:26 PM

Load delta: 0 minutes

OK Cancel Help

Log Shipping HA Features (1 of 2)

Availability Feature	Failover Clustering	Log Backup Shipping
Standby Type	Hot	Warm
Failure Detection	Yes	No, NLB helps
Automatic Failover	Yes	No, NLB helps
Masks Disk Failure	No; Shared Disk	Yes
Masks SQL Process Failure	Yes	Yes
Masks Other Process Failure	Yes	No
Meta Data Support	All System and Database	Database Only
Transactionally Consistent	Yes	Yes
Transactionally Current	Yes, Always Up to Date	No, Since Last Log Backup
Perceived Downtime	30 Seconds + DB Recovery	Seconds + DB Recovery Time

Log Shipping HA Features (2 of 2)

Availability Feature	Failover Clustering	Log Backup Shipping
Transparent to Client	Yes, Reconnect to Same IP	No, App must know standby
Special Hardware Needed	Specialized Hardware from Cluster HCL	No; Duplicate system needed
Distance Limit	100 Miles	Dispersed
Complexity	More	Some
Standby Accessible	Standby never accessible	Yes, Multiple Copies, Read-only; % depends on update frequency
Impact on Performance	No Impact	Minimal – File Copy on Primary
Impact on Backup Strategy	Must be able to backup from any node	Minimal – many small backups

High Availability Uses of Log Shipping

- Shorter failover time
- If there is a high incidence of user error and a need to recover data frequently without recovering the whole database
 - Allows you time-delay possibilities
 - 5 hours behind
 - 8 hours behind
- Increase data redundancy
- Less complex hardware – no HCL

When to Consider Using Replication for HA

- After Considering Failover Clustering
- After Considering Log Shipping
- System and Some User Metadata is Not Replicated
- Failure Detection and Failover is Not Automatic
 - Standby Server is Not Identical to the Primary
- Not Guaranteed to be Transactionally Current
 - Merge Replication is not Transactionally Consistent
- Replication Uniquely Allows:
 - Partitioning of Data on the Standby Server
However, standby server is not identical to primary server
 - Offline Access to the Data without Periodic Termination

Transactional Replication

- Warm Standby Solution
- Propagates Transactions From Primary Server (Publisher) to Warm Spare (Subscriber)
- Use Replication to Create
 - A Read-Only Spare
 - A Scale Out Solution
 - A Partitioned Solution
- Manual Failure Detection; Manual Failover

Comparing Clustering, Log Shipping, and Transactional Replication (1 of 2)

Availability Feature	Failover Clustering	Log Backup Shipping	Transactional Replication
Standby Type	Hot	Warm	Warm
Failure Detection	Yes	No, NLB helps	No
Automatic Failover	Yes	No, NLB helps	No, NLB helps
Masks Disk Failure	No; Shared Disk	Yes	Yes
Masks SQL Process Failure	Yes	Yes	Yes
Masks Other Process Failure	Yes	No	No
Meta Data Support	All System and Database	Database Only	Object(s)
Transactionally Consistent	Yes	Yes	Transactional: Yes; Merge: No
Transactionally Current	Yes, Always Up to Date	No, Since Last Log Backup	No, Since Last Distributed Op
Perceived Downtime	30 Seconds + DB Recovery	Seconds + DB Recovery Time	Detect, then manual fail over

Comparing Clustering, Log Shipping, and Transactional Replication (2 of 2)

Availability Feature	Failover Clustering	Log Backup Shipping	Transactional Replication
Transparent to Client	Yes, Reconnect to Same IP	No, App must know standby	No, App must know standby
Special Hardware Needed	Specialized Hardware from Cluster HCL	No; Duplicate system needed	No; Duplicate system needed
Distance Limit	100 Miles	Dispersed	Dispersed
Complexity	More	Some	More
Standby Accessible	Standby never accessible	Yes, Multiple Copies, Read-only; % depends on update frequency	Yes, Multiple Copies, Read-only; perhaps updateable
Impact on Performance	No Impact	Minimal – File Copy on Primary	Minimal Impact
Impact on Backup Strategy	Must be able to backup from any node	Minimal – many small backups	No Impact; but must backup distribution database

Choosing Replication to Implement a Warm Standby

- Data Partitioning
- Offline Access to the Data
- Geographically Separated Warm Standby
- Costs
 - Complexity
 - Process
 - Resources

Using Combinations of Technologies

- Log Shipping With Replication
- SQL Server Failover Clustering With Log Shipping
- Log Shipping With Network Load Balancing

Establish Operational Excellence

- Data center principles
- Change Control
- Staffing
- Disaster Recovery Plan
- Run Book

Monitoring for HA

- Two theories:
 - Every counter 100% of the time
 - Just what you need
- Don't forget Profiler
- Coordinate with Event Logs, SQL Logs, IIS Logs, etc.
 - Time difference between servers
 - HA is a *total* solution ... not just SQL

Backup And Restore

- Last resort for HA!
- Develop a backup strategy
 - Full database backups
 - File/filegroup backups
 - Transaction log backups
 - Image of operating system disk
- Test your backups on another server
- Rotate your tapes off-site

Backup And Restore (cont'd)

- Test your recovery plans
- Pick a point in time to recover to
 - Locate the tapes
 - Test using the graphical interface
 - Test using script only
 - Test with different people on all teams
- Time the drill: how long will it take?

Designing A Disaster Recovery Plan

- One of the keys to HA
 - Without this, you might as well not do HA
- Different plans:
 - Site down
 - Server down
 - Data gone
- Document plan (keep updated) – test, test TEST! Store results/learnings
- Off-site backup storage, including the operational manual (run book)

SQL Server 2000 High Availability Resources

- MSPress title: *SQL Server 2000 High Availability*
Authors: Allan Hirt with Cathan Cook, Kimberly L. Tripp,
Frank McBath ISBN: 0-7356-1920-4
- Microsoft SQL Server 2000 High Availability Series
<http://www.microsoft.com/technet/prodtechnol/sql/2000/deploy/sqlhalp.msp>

High Availability In SQL Server

Questions??