

Microsoft® Official Course



Planning and Implementing High Availability

Microsoft®

Module Overview

- High Availability on Exchange Server 2013
- Configuring Highly Available Mailbox Databases
- Configuring Highly Available Client Access Servers

Lesson 1: High Availability on Exchange Server 2013

- Components of High Availability
- What Is a Database Availability Group?
- Understanding How Database Availability Groups Work
- Understanding How High Availability Works with Client Access Servers
- Understanding How Transport High Availability Works
- Understanding How High Availability Works with Edge Transport Servers
- What Is Site Resilience?
- Discussion: Virtualization High Availability Technologies versus Exchange High Availability Technologies for Mailbox Servers

Components of High Availability

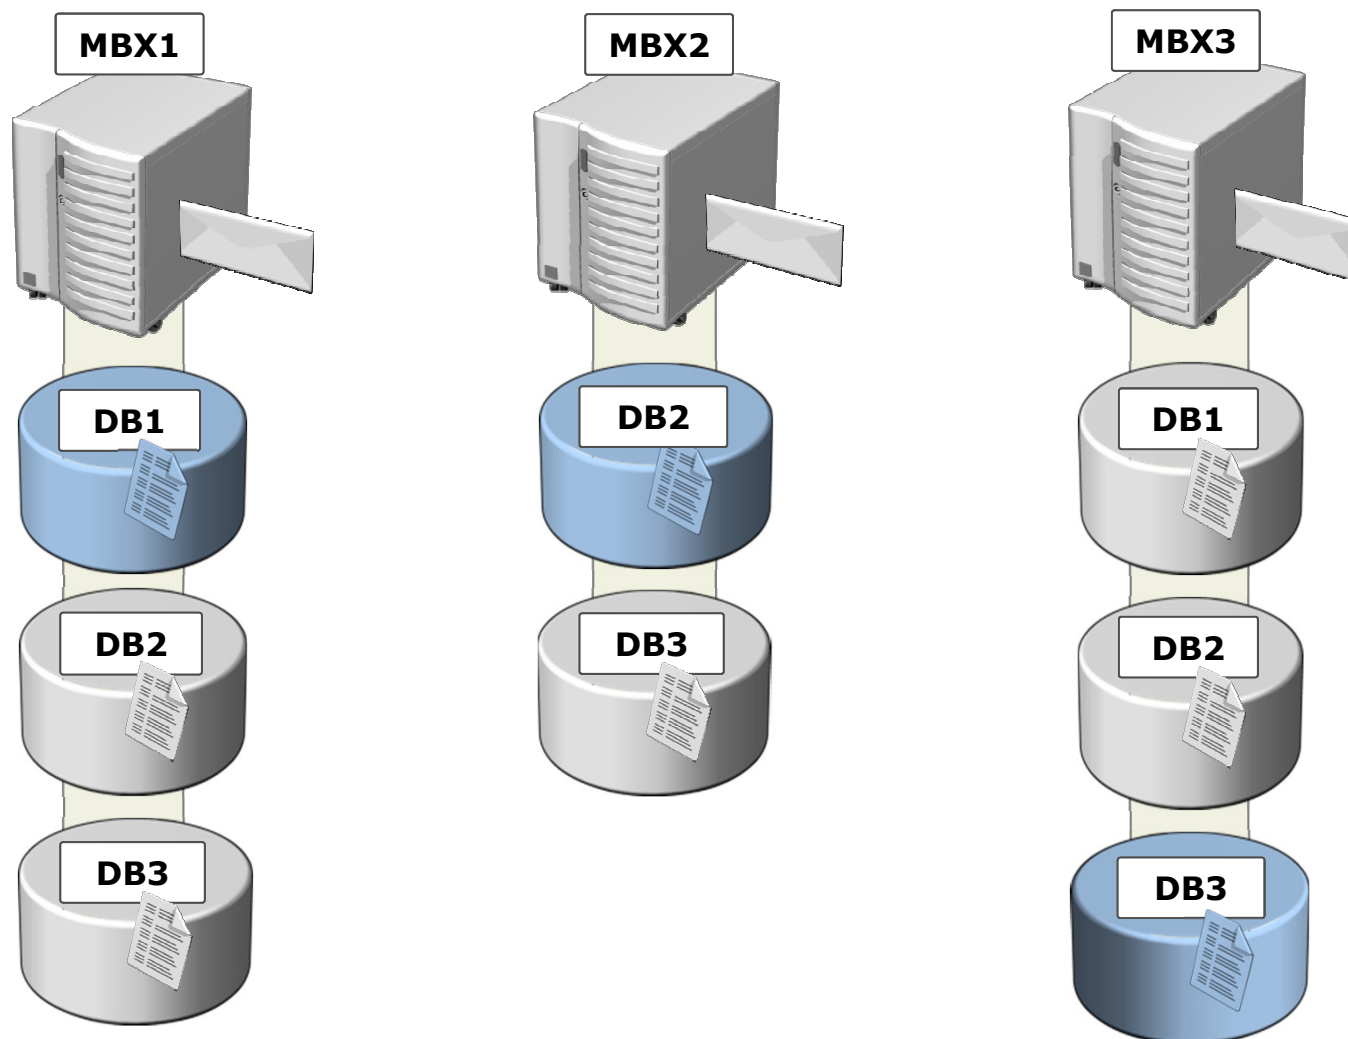
- All parts of the application and infrastructure it relies on must be highly available
 - Data center infrastructure
 - Server hardware
 - Storage
 - Network infrastructure
 - Internet
 - Network services

What Is a Database Availability Group?

- A DAG is a collection of servers that provides the infrastructure for replicating and activating database copies. DAGs:
 - Require the failover clustering feature, although all installation and configuration is done with the Exchange Server management tools
 - Use Active Manager to control failover
 - Use an enhanced version of the continuous replication technology that Exchange Server 2007 introduced
 - Can be created after the Mailbox server is installed
 - Allow a single database to be activated on another server in the group without affecting other databases
 - Allow up to 16 copies of a single database on separate servers
 - Define the boundary for replication


Understanding How Database Availability Groups Work

- Continuous replication protects databases across servers in the DAG

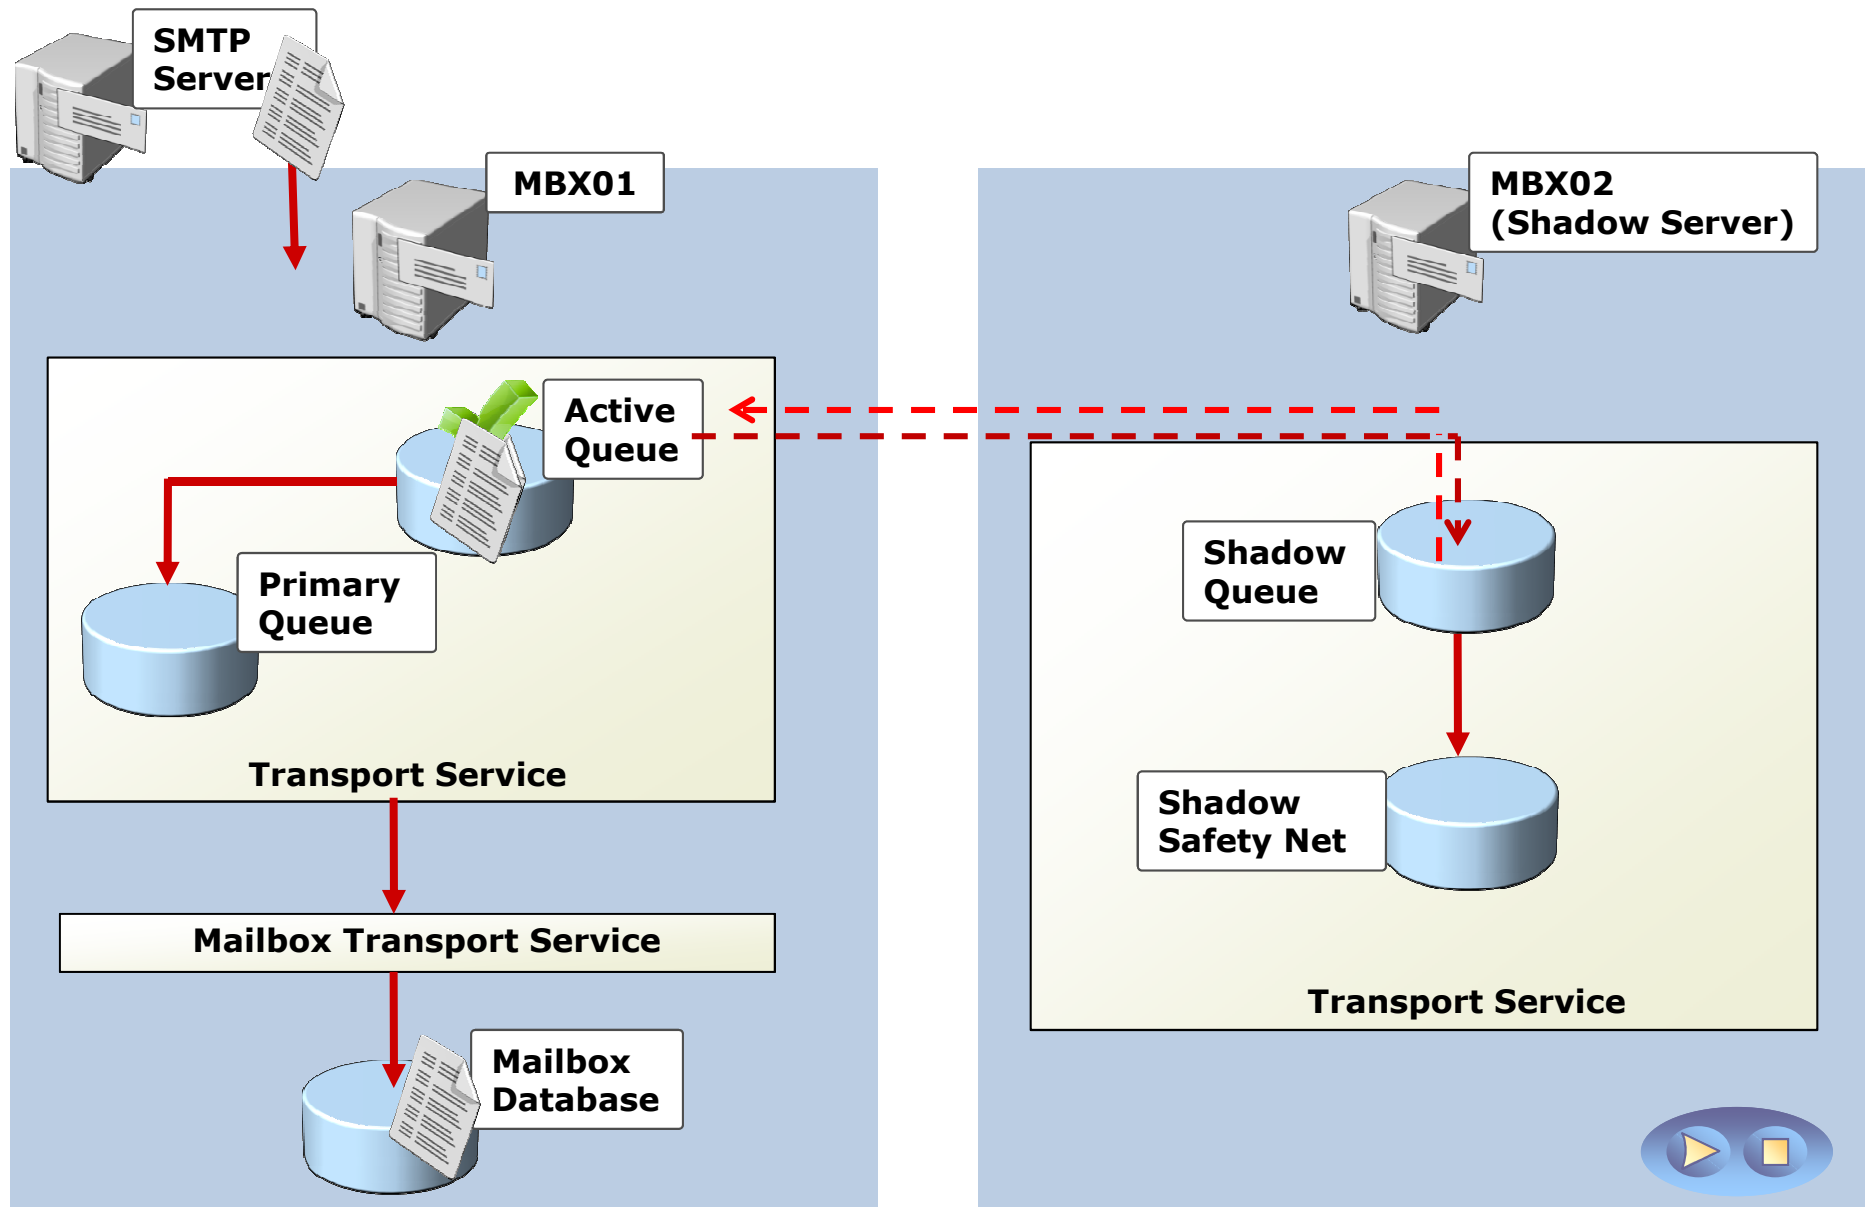


Understanding How High Availability Works with Client Access Servers

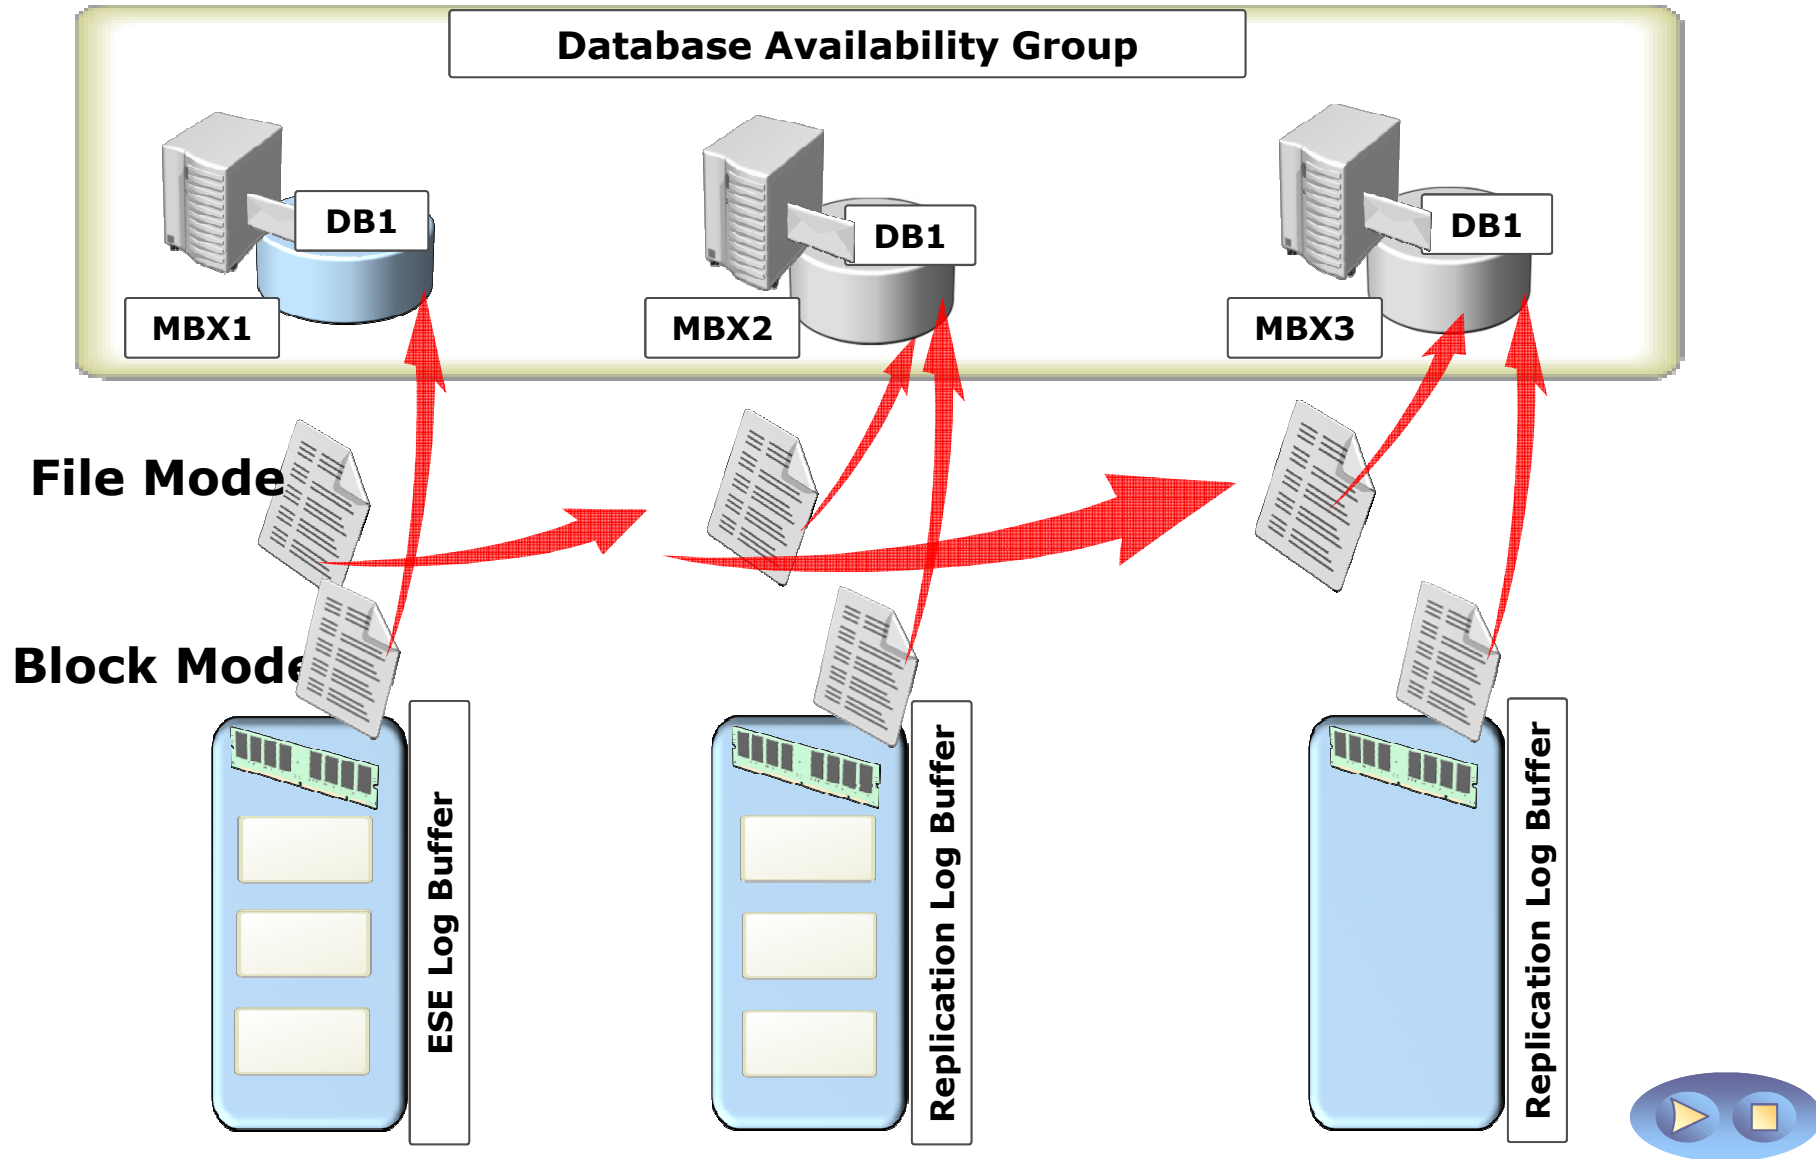
- You make Client Access servers highly available with multiple Client Access servers and by configuring load balancing
- You can achieve high availability and load balancing by using:
 - DNS round-robin
 - Network Load Balancing
 - Hardware-based load balancing



Understanding How Transport High Availability Works



Understanding How High Availability Works with Edge Transport Servers



What Is Site Resilience?

- Site resilience is the ability of the messaging system to survive the failure of a site
 - An alternate data center can be located at another company location that has sufficient capacity
 - A DAG can span Active Directory sites
 - Other roles and services must already exist in the site

Lesson 2: Configuring Highly Available Mailbox Databases

- What Is a Quorum?
- Planning Software and Hardware Components for Database Availability Groups
- What Is Active Manager?
- What Is Continuous Replication?
- Configuring a Database Availability Group
- Configuring Databases for High Availability
- What are Lagged Mailbox Database Copies?
- Demonstration: How to Create and Configure a Database Availability Group
- Understanding the Failover Process
- Planning, Monitoring and Managing a Database Availability Group
- Demonstration: How to Monitor Replication Health

What Is a Quorum?

- A quorum uses a consensus of voters to ensure that enough cluster members are available to provide services
- Exchange Server 2013 DAG quorums:
 - Are based on votes in Windows Server 2012
 - Allow nodes, file shares, and shared disks to have votes, depending on the quorum mode
 - Use node majority with a witness server for quorum:
 - DAGs with an even number of mailbox servers use the witness server
 - DAGs with an odd number of mailbox servers use node majority
- You can have non-voting nodes, if needed

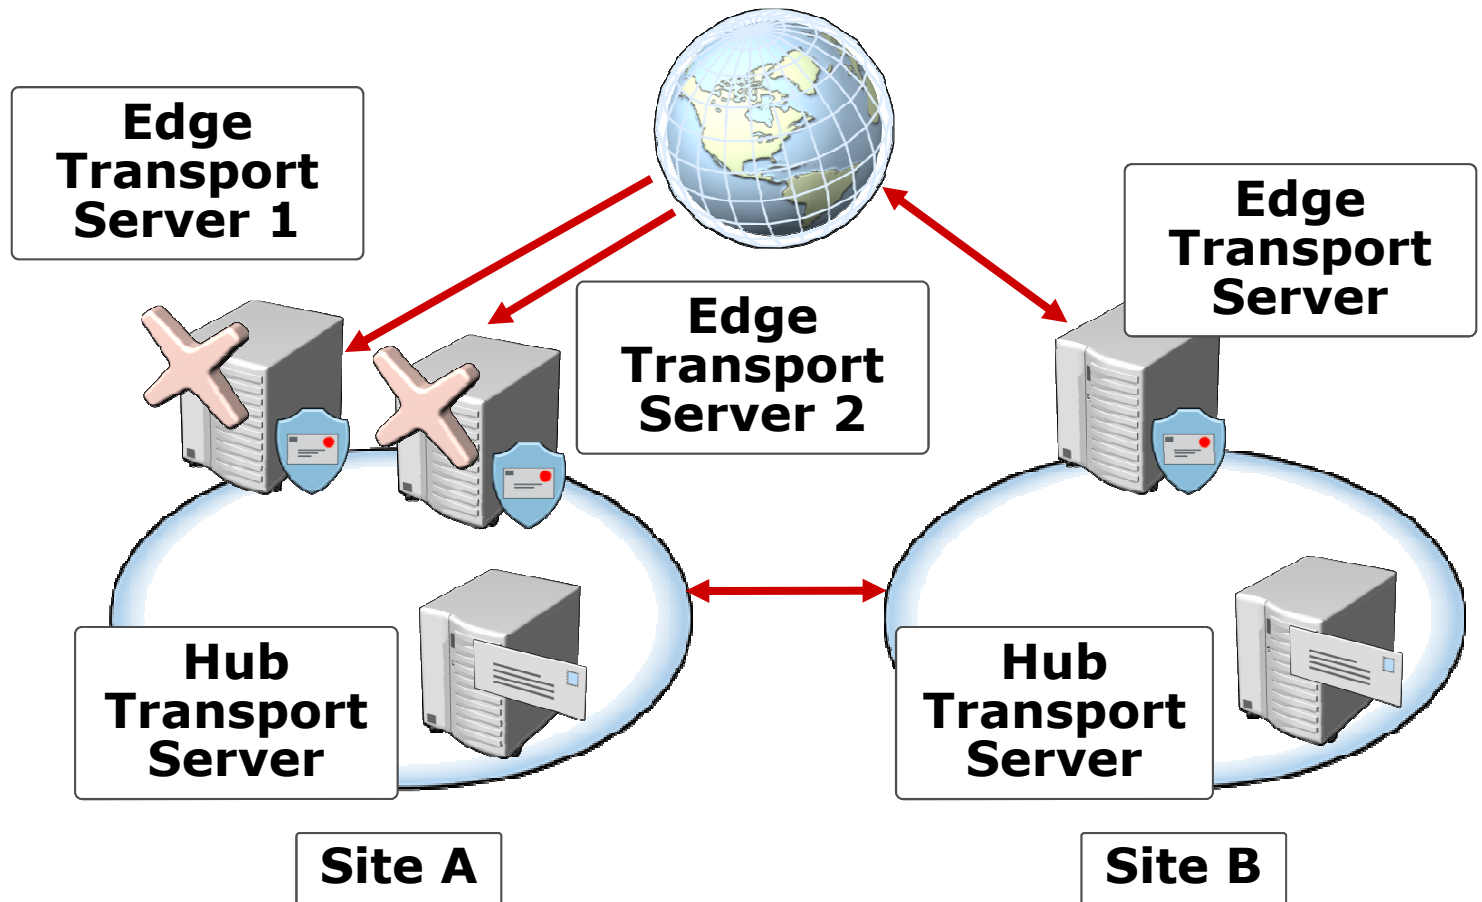
Planning Software and Hardware Components for Database Availability Groups

- General configuration
 - Nodes in the same domain
 - Not domain controllers
- Windows Server 2012 Datacenter or Standard, or Windows Server 2008 R2 Datacenter or Enterprise
- Network configuration
 - Two adapters
 - Less than 250 ms latency
- DAG must have one IP on the MAPI network
- Witness server

What Is Active Manager?

- Active Manager:
 - Runs a process on each server in the DAG
 - The PAM is one node
 - The remaining nodes are SAM
 - Manages which database copies are active and which are passive
 - Stores database state information
 - Manages database switchover and failover processes
 - Does not require direct administration configuration

What Is Continuous Replication?



Configuring a Database Availability Group

- To configure DAGs, you must define the following:
 - Witness Server – server used to store witness information
 - Witness Directory – directory used on the witness server to store witness information
 - Database availability group IP addresses – IP address(es) used by DAG
- Also consider these settings for larger or multi-site implementations:
 - DAG networks including replication
 - DAG network compression
 - DAG network encryption
 - Third-party replication mode
 - Alternative witness server
 - Alternative Witness Directory

Configuring Databases for High Availability

- After creating a DAG, adding Mailbox servers to the DAG, and configuring the DAG, you must still complete the following steps:
 - Create database copies
 - Set activation preference number
 - Set replay lag time
 - Set truncation lag time

What are Lagged Mailbox Database Copies?

- A lagged database copy is a database that uses a delayed replay lag time to commit the log files to the database
- You create a lagged database to prevent:
 - Database Logical Corruption
 - Store Logical Corruption
 - Rogue Admin Protection
- Lagged database copy enhancements in Exchange Server 2013:
 - Automatic log play down
 - Simpler activation with Safety Net
 - Lagged copies can now be configured in EAC

Demonstration: How to Create and Configure a Database Availability Group

- In this demonstration, you will see how to create a new database availability group





Understanding the Failover Process

- If a failure occurs, the following steps occur for the failed database:
 - Active Manager determines the best copy to activate
 - The replication service on the target server attempts to copy missing log files from the best "source":
 - If successful, the database mounts with zero data loss
 - If unsuccessful (failover), the database mounts based on the AutoDatabaseMountDial setting
 - The mounted database generates new log files (using the same log generation sequence)
 - Transport dumpster (safety net) requests are initiated for the mounted database to recover lost messages
 - When the original server or database recovers, it determines if any logs are missing or corrupt, and fixes them if possible

Planning, Monitoring and Managing a Database Availability Group

- Allocate the necessary permissions for managing a DAG
 - Organization Management
 - DAGs
 - Database copies
- Failure may not be noticed
- Exchange Server 2013 includes several scripts and commands for DAG monitoring and management
- Consider using System Center Operations Manager 2012

Demonstration: How to Monitor Replication Health

- In the demonstration, you will see how to view the health status of database copies in the Exchange Administration Center or Exchange Management Shell.

Lesson 3: Configuring Highly Available Client Access Servers

- Planning Software and Hardware Components for Highly Available Client Access Servers
- What Is Network Load Balancing?
- Considerations for Implementing Highly Available Client Access Servers
- Demonstration: Configuring Options for Highly Available Client Access Servers

Planning Software and Hardware Components for Highly Available Client Access Servers

- Client Access servers used by all clients
- Implement load balancing for Client Access servers
- Sites performing a proxy are most important
- Same SSL certificate is used on all nodes
- Consider redundant Internet connections for external users

What Is Network Load Balancing?

- NLB enables high availability and scalability for server applications, typically web based
- NLB enables access to two or more servers through a single IP address
- NLB requires all servers in the NLB cluster to be running the same server applications with the same configuration
- Client requests are distributed across available servers in the NLB cluster, based on a common algorithm
- All servers monitor each other through heartbeat messages
- If a server fails to send heartbeat messages for five seconds, the other servers automatically converge and redistribute the client connections across the available servers

Considerations for Implementing Highly Available Client Access Servers

- When implementing highly available Client Access servers, consider the following:
 - All CAS servers should have a certificate to match their address space
 - All CAS servers should have the same protocols enabled
 - Use a hardware network load balancer
 - Deploy servers with similar hardware, memory and performance

Opgave

- **Konfiguration af Database Availability Groups (DAG)**
- Opret en Database Availability Group.
- [http://technet.microsoft.com/en-us/library/dd298065\(v=exchg.150\).aspx](http://technet.microsoft.com/en-us/library/dd298065(v=exchg.150).aspx)
- Tilføj en mailbox database copy og test failover.
- [http://technet.microsoft.com/en-us/library/dd335158\(v=exchg.150\).aspx](http://technet.microsoft.com/en-us/library/dd335158(v=exchg.150).aspx)

Lab: Implementing High Availability

- Exercise 1: Creating and Configuring a Database Availability Group
- Exercise 2: Deploying Highly Available Client Access Servers
- Exercise 3: Testing the High-Availability Configuration

Logon Information

20341A-LON-DC1

20341A-LON-CAS1

20341A-LON-CAS2

20341A-LON-MBX1

20341A-LON-MBX2

User name: Adatum\Administrator

Password: Pa\$\$w0rd

Estimated Time: 60 minutes

Lab Scenario

You are the messaging administrator for A. Datum Corporation. You have completed the basic installation for four Exchange Server 2013 servers. Now you must complete the configuration so that they are highly available. This basically requires you to configure your mailbox databases as well as your CAS servers to be highly available, and to test if an automatic failover works.

Lab Review

- When do you need to pre-stage the cluster network object for a database availability group?
- In the lab, one mailbox server failed. How did the other mailbox server achieve a quorum?

Module Review and Takeaways

- Review Question
- Best Practice
- Common Issues and Troubleshooting Tips