# Microsoft Exchange Server 2010 High Availability

Harold Wong
Blogs.technet.com/b/haroldwong

# Active Manager

- Transition
  HighAvaila

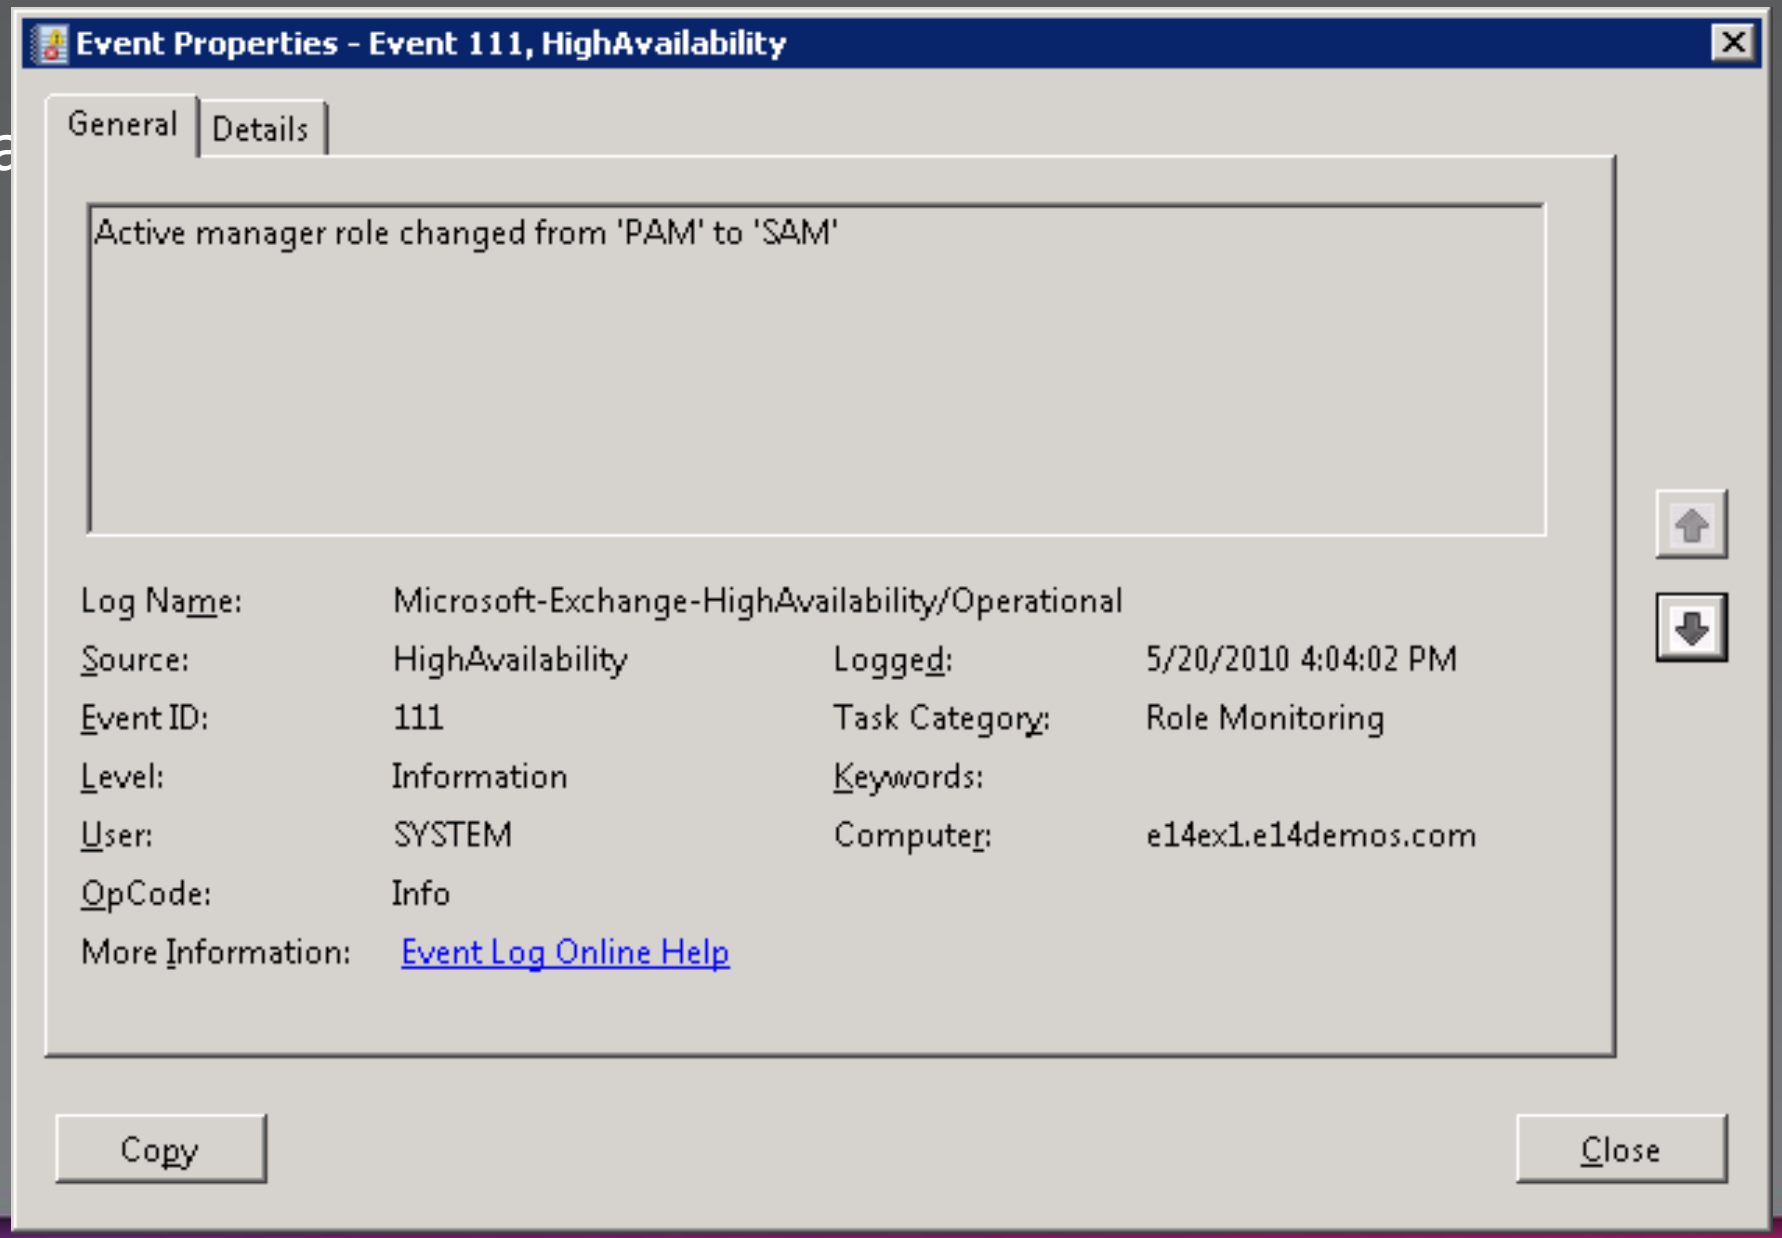

# Active Manager

- Primary Active Manager (PAM)
  - Runs on the node that owns the cluster group
  - Gets topology change notifications
  - Reacts to server failures
  - Selects the best database copy on *overs
  - Detects failures of local Information Store and local databases
- Standby Active Manager (SAM)
  - Runs on every other node in the DAG
  - Detects failures of local Information Store and local databases
    - Reacts to failures by asking PAM to initiate a failover
  - Responds to queries from CAS/Hub about which server hosts the active copy
- Both roles are necessary for automatic recovery
  - If the Replication service is stopped, automatic recovery will not happen

# Active Manager

- Startup process depends on whether AM is Standalone or in a DAG
  - Standalone
    - Replication service starts and reads configuration from Active Directory
    - Sets Active Manager role to Standalone
    - Active Manager queries Active Directory every 30 seconds for changes
      - If it has been added to a DAG, the DAG Active Manager logic is started
  - DAG
    - Replication service starts and reads configuration from Active Directory
    - Replication service assumes SAM role and sets CurrentPAM to *unknown*
    - Replication service determines current PAM holder (who owns cluster group)
      - If local server is PAM, Replication service assumes PAM role
      - If remote server is PAM, Replication service maintains SAM role
    - Replication service sets CurrentPAM to PAM role holder

# Active Manager

- Replication service thread monitors for cluster group changes and reacts as follows:
  - If DAG member owns cluster group and CurrentPAM is set to another member, it will
    - Verify with all other DAG members
    - Assume PAM role
    - Set CurrentPAM to itself
  - If DAG member does not own cluster group, but it is configured as CurrentPAM,
    - Indicates that the cluster group has been moved to another DAG member
    - All outstanding Active Manager operations are immediately finished
    - CurrentPAM is set to new owner of cluster group
    - DAG member assumes SAM role
  - If DAG member does not own cluster group, and is not configured as CurrentPAM, DAG member maintains SAM role

# Best Copy Selection

- Active Manager selects the “best” copy to become the new active copy when the existing active copy fails
  - Sorts copies by currency (copy queue length) to minimize data loss
    - Breaks ties during sort based on Activation Preference
  - Selects from sorted listed based on which set of criteria met by each copy
  - Attempt Copy Last Logs (ACLL) runs and attempts to copy missing log files from previous active copy
  - Is database mountable? Is copy queue length < AutoDatabaseMountDial?
    - If Yes, database is marked as current active and mount request is issued
    - If not, database that meets next set of criteria tried
- During best copy selection, any servers that are unreachable or “activation blocked” are ignored

# Best Copy Selection

| Criteria | Copy Queue Length | Replay Queue Length | Content Index Status |
|---|---|---|---|
| 1 | < 10 logs | < 50 logs | Healthy |
| 2 | < 10 logs | < 50 logs | Crawling |
| 3 | N / A | < 50 logs | Healthy |
| 4 | N / A | < 50 logs | Crawling |
| 5 | N / A | < 50 logs | N / A |
| 6 | < 10 logs | N / A | Healthy |
| 7 | < 10 logs | N / A | Crawling |
| 8 | N / A | N / A | Healthy |
| 9 | N / A | N / A | Crawling |
| 10 | Any database copy with a status of Healthy, DisconnectedAndHealthy, DisconnectedAndResynchronizing, or SeedingSource | | |

# Best Copy Selection

- Four copies of DB1
- DB1 currently active on Server1

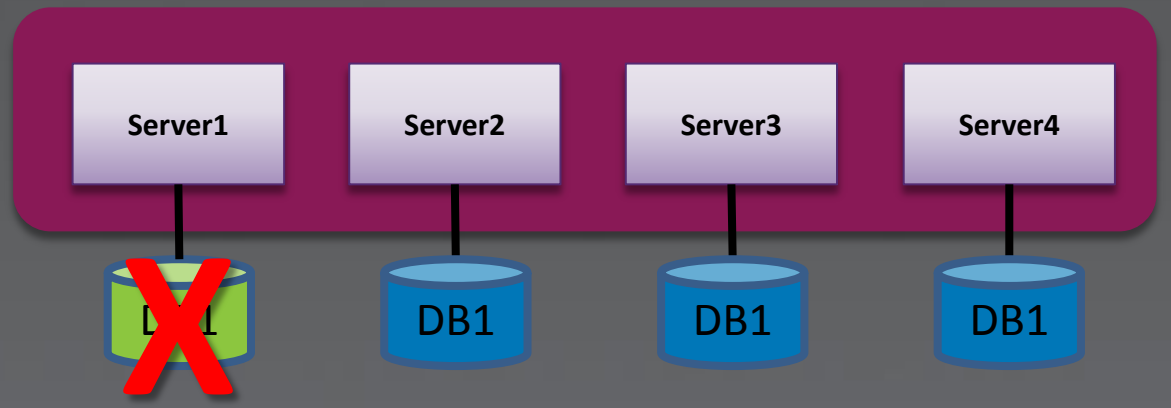


| Database Copy | Activation Preference | Copy Queue Length | Replay Queue Length | CI State | Database State |
|---|---|---|---|---|---|
| Server2\DB1 | 2 | 4 | 0 | Healthy | Healthy |
| Server3\DB1 | 3 | 2 | 2 | Healthy | DiscAndHealthy |
| Server4\DB1 | 4 | 10 | 0 | Crawling | Healthy |

# Best Copy Selection

- Sort list of available copies based by Copy Queue Length (using Activation Preference as secondary sort key if necessary):
  - Server3\DB1
  - Server2\DB1
  - Server4\DB1

| Database Copy | Activation Preference | Copy Queue Length | Replay Queue Length | CI State | Database State |
|---|---|---|---|---|---|
| Server2\DB1 | 2 | 4 | 0 | Healthy | Healthy |
| Server3\DB1 | 3 | 2 | 2 | Healthy | DiscAndHealthy |
| Server4\DB1 | 4 | 10 | 0 | Crawling | Healthy |

# Best Copy Selection

- Only two copies meet first set of criteria for activation (CQL < 10; RQL < 50; CI=Healthy):
  - Server3\DB1 ⟵ Lowest copy queue length – tried first
  - Server2\DB1
  - Server4\DB1

| Database Copy | Activation Preference | Copy Queue Length | Replay Queue Length | CI State | Database State |
|---|---|---|---|---|---|
| Server2\DB1 | 2 | 4 | 0 | Healthy | Healthy |
| Server3\DB1 | 3 | 2 | 2 | Healthy | DiscAndHealthy |
| Server4\DB1 | 4 | 10 | 0 | Crawling | Healthy |

# Best Copy Selection

- After Active Manager determines the best copy to activate
  - The Replication service on the target server attempts to copy missing log files from the source (ACLL)
    - If successful, then the database will mount with zero data loss
    - If unsuccessful (lossy failure), then the database will mount based on the AutoDatabaseMountDial setting
    - If data loss is outside of dial setting, next copy will be tried
  - The mounted database will generate new log files (using the same log generation sequence)
  - Transport Dumpster requests will be initiated for the mounted database to recover lost messages
  - When original server or database recovers, it will run through divergence detection and either perform an incremental resync or require a full reseed

# Datacenter Activation Coordination Mode

# Datacenter Activation Coordination Mode

- DAC mode is a property of a DAG
- Acts as an application-level form of quorum
  - Designed to prevent multiple copies of same database mounting on different members due to loss of network
- RTM: DAC Mode is only for DAGs with three or more members that are extended to two Active Directory sites
  - Should not be enabled for two-member DAGs where each member is in a different Active Directory site or DAGs where all members are in the same Active Directory site
  - In RTM, DAC Mode also enables use of Site Resilience tasks
    - Stop-DatabaseAvailabilityGroup
    - Restore-DatabaseAvailabilityGroup
    - Start-DatabaseAvailabilityGroup
- SP1: DAC Mode can be enabled for all DAGs

# Datacenter Activation Coordination Mode

- Uses Datacenter Activation Coordination Protocol (DACP), which is a bit in memory set to either:
  - 0 = can't mount
  - 1 = can mount
- Active Manager startup sequence
  - DACP is set to 0
  - DAG member communicates with other DAG members it can reach to determine the current value for their DACP bits
    - If the starting DAG member can communicate with all other members, DACP bit switches to 1
    - If other DACP bits are set to 0, starting DAG member DACP bit remains at 0
    - If another DACP bit is set to 1, starting DAG member DACP bit switches to 1

# Monitoring Exchange 2010 High Availability

Built-in Tools

# Monitoring Best Practices

- Ensure that your servers are operating reliably and that your database copies are healthy are key objectives for daily messaging operations
- Actively monitor hardware, the Windows operating system, Exchange 2010 services, and database and database copy health
- Monitoring actively and daily enables you to:
  - Meet service level agreements (SLAs)
  - Ensure regular administrative tasks have completed (e.g., backups)
  - Detect and address issues that might affect service or data availability
- Exchange 2010 includes several built-in tools for monitoring high availability
- System Center Operations Manager automates and enhances these tools with Exchange 2010 Management Pack

# Get-MailboxDatabaseCopyStatus

- Used to view information about copies of a particular database, a specific copy of a database on a specific server, or about all database copies on a server
- Examples
  - Get status for all copies of a database

    ```
    Get-MailboxDatabaseCopyStatus -Identity DB2 | FL
    ```

  - Get status for all copies on the local server

    ```
    Get-MailboxDatabaseCopyStatus -Local | FL
    ```

  - Get status for all copies on a remote server

    ```
    Get-MailboxDatabaseCopyStatus -Server MBX2 | FL
    ```

  - Get status, log shipping and seeding network information

    ```
    Get-MailboxDatabaseCopyStatus -Identity DB3\MBX1 -ConnectionStatus | FL
    ```

# Test-ReplicationHealth

- Designed for proactive monitoring of continuous replication, the availability of Active Manager, and the health and status of the underlying cluster service, quorum, and network components
- Can be run locally on or remotely against any Mailbox server in a DAG
- Example
  - Test the health of a DAG member

```
Test-ReplicationHealth -Identity MBX1
```

# Crimson Channel Event Logging

- Windows Server 2008 includes two categories of event logs
  - Windows logs (includes legacy Application, Security and System event logs, as well as new Setup and ForwardedEvent logs)
  - Applications and Services logs
    - New category of events logs used for storing events from a single application or component, rather than events that might have system-wide impact
    - This new category is referred to as an application's 'crimson channel'.
    - Includes four general subtypes (there can be custom ones, too)
      - Admin (useful for troubleshooting; contain guidance for problem resolution)
      - Operational (somewhat useful; require a bit more interpretation)
      - Analytic (hidden and disabled by default)
      - Debug (used by developers when debugging applications)

# Crimson Channel Event Logging

- Active

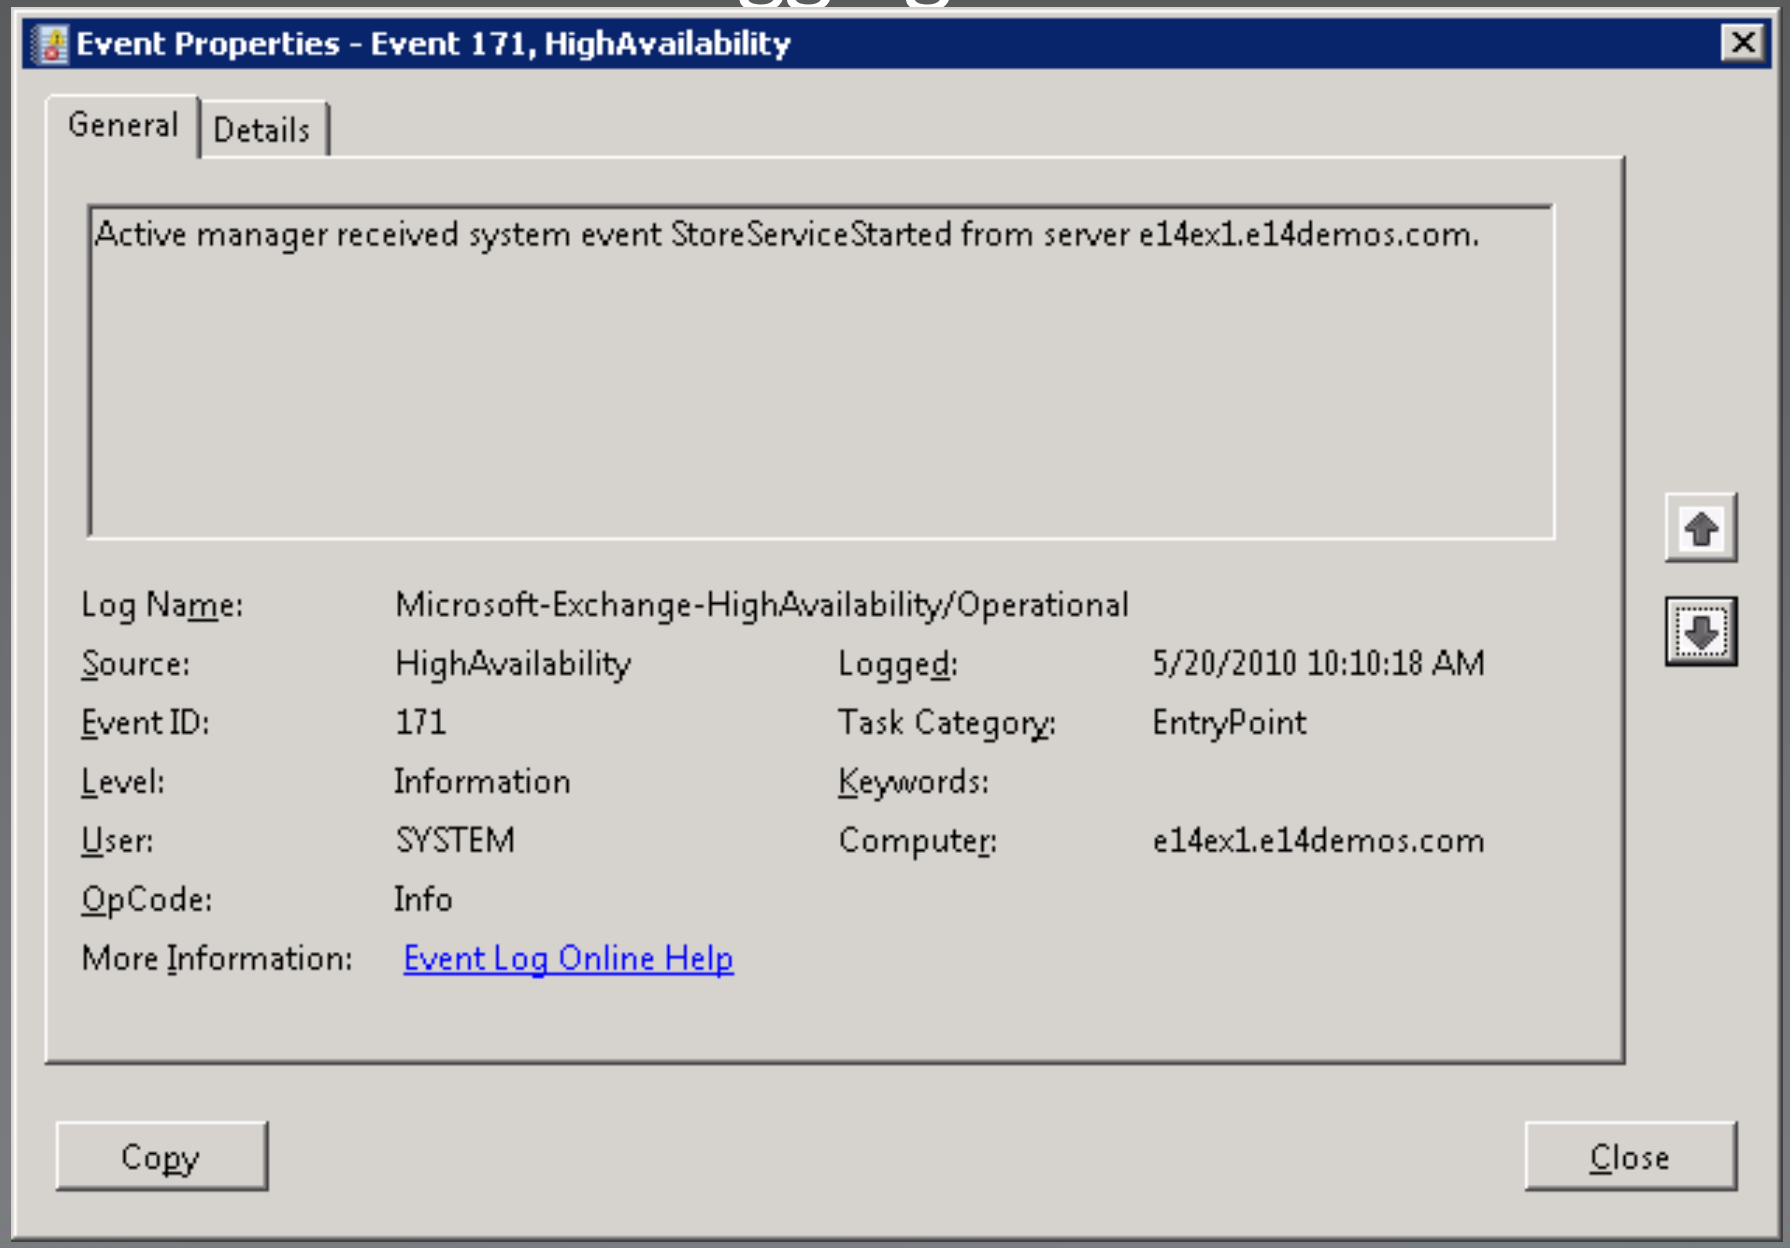

# CollectReplicationMetrics.ps1

- Collects replication performance data metrics for a DAG in real-time
- Script represents an active form of monitoring
  - Collects metrics in real-time while running
    - Hours as active copy, Hours as passive copy
    - Minutes unavailable, Minutes Resynchronizing, Minutes Failed, Minutes Suspended, Minutes FailedAndSuspended, Minutes Disconnected
    - Average log generation rate, Peak log generation rate
    - Average log copy rate, Peak log copy rate
    - Average log replay rate, Peak log replay rate
    - Percentage of time log copying used replication network *N* (for each *N*)
    - Percentage of the time log copying was using MAPI network

# CollectOverMetrics.ps1

- Workflow script that collects information about various switchover and failover-related statistics
- Script represents a passive form of monitoring
  - Collects and analyzes events that have already been recorded
- Script settings and output are customizable
- Examples

```
CollectOverMetrics.ps1 -DatabaseAvailabilityGroup DAG1 -Database:"DB*"
-GenerateHTMLReport –ShowHTMLReport
```

```
CollectOverMetrics.ps1 -DatabaseAvailabilityGroup DAG1
-GenerateHTMLReport -ShowHTMLReport
```

# CheckDatabaseRedundancy.ps1

- Recently released script
  - http://msexchangeteam.com/archive/2010/05/20/454976.aspx
  - Already shipped as updated to SCOM Management Pack
  - Expected in Update Rollup 4, and will ship in SP1 (might be modified)
- Script is an active form of monitoring
  - Single Copy Alert - monitors redundancy of replicated mailbox databases
    - Validates that there are at least two configured and healthy and current copies
  - Fires red alert when down to a single copy for 20 minutes or more; and fires green alert when database has been green for 10 minutes
- Best Practice: Run this script regularly, part of normal monitoring options
  - Run script every 60 minutes
  - Without SCOM, run as Scheduled Task in Windows 2008 R2
    - Windows 2008 SP2 has issue with long-running tasks

# CheckDatabaseRedundancy.ps1

- Examples

```
CheckDatabaseRedundancy.ps1 -MailboxDatabaseName DB1

CheckDatabaseRedundancy.ps1 -MailboxServerName EX1
```

- Scheduled Task Examples
  - Windows Server 2008 SP2

```
CheckDatabaseRedundancy.ps1 -MonitoringContext -
SleepDurationBetweenIterationsSecs:0 -TerminateAfterDurationSecs:1 -
SuppressGreenEventForSecs:0 -ReportRedEventAfterDurationSecs:0 -
ReportRedEventIntervalSecs:0 -ShowDetailedErrors
```

  - Windows Server 2008 R2

```
schtasks /create /TN "Check Database Redundancy" /TR "Powershell.exe -
NonInteractive -WindowStyle Hidden -command 'C:\Program
Files\Microsoft\Exchange Server\V14\bin\RemoteExchange.ps1'; Connect-
ExchangeServer -auto; C:\Ops\CheckDatabaseRedundancy.ps1 -MonitoringContext -
ShowDetailedErrors -SummaryMailFrom:'SMTPFromAddress@contoso.com' -
SendSummaryMailTos:@('SMTPToAddress@contoso.com') -ErrorAction:Continue" /RU
SYSTEM /SC HOURLY
```

# Improvements in Service Pack 1

Replication and Copy Management enhancements in SP1

# Improvements in Service Pack 1

- Continuous replication changes
  - Enhanced to reduce data loss
  - Eliminates log drive as single point of failure
- Automatically switches between modes:
  - File mode (original, log file shipping)
  - Block mode (enhanced log block shipping)
- Switching process:
  - Initial mode is file mode
  - Block mode triggered when target needs Exx.log file
  - All healthy passives processed in parallel
  - File mode triggered when block mode falls too far behind

# Improvements in Service Pack 1

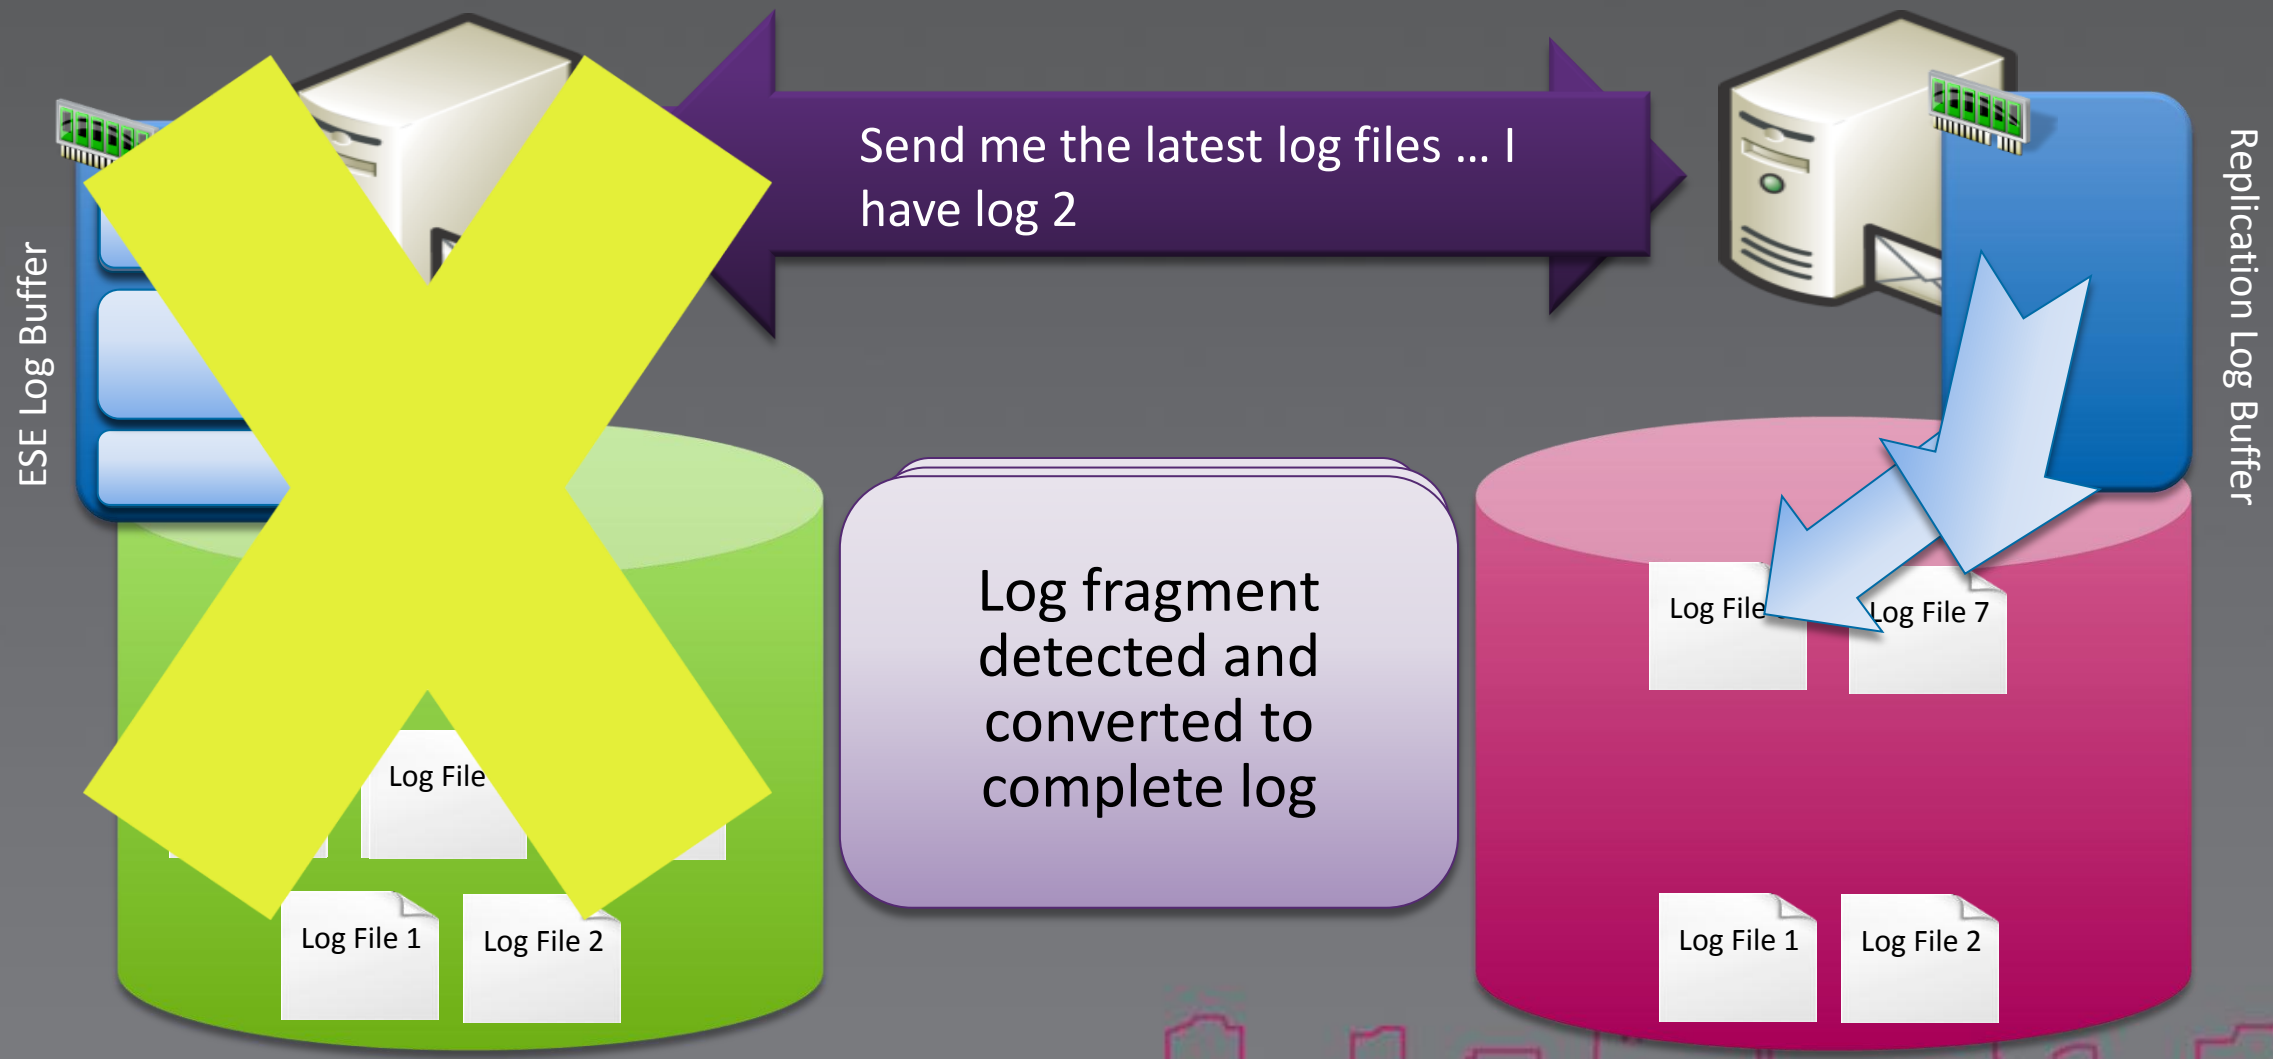


Continuous Replication – Block Mode

# Improvements in Service Pack 1

- SP1 introduces RedistributeActiveDatabases.ps1 script (keep database copies balanced across DAG members)
  - Moves databases to the most preferred copy
  - If cross-site, tries to balance between sites
- Targetless admin switchover altered for stronger activation preference affinity
  - First pass of best copy selection sorted by activation preference; not copy queue length
  - This basically trades off even distribution of copies for a longer activation time.  So you might pick a copy with more logs to play, but it will provide you with better distribution of databases

# Improvements in Service Pack 1

- *over Performance Improvements
  - In RTM, a *over immediately terminated replay on copy that was becoming active, and mount operation did necessary log recovery
  - In SP1, a *over drives database to clean shutdown by playing all logs, and mount brings database up-to-date (no recovery required)

# Improvements in Service Pack 1

- Exchange Manag
  - Manage DAG IP
  - Manage witnes erver/directory

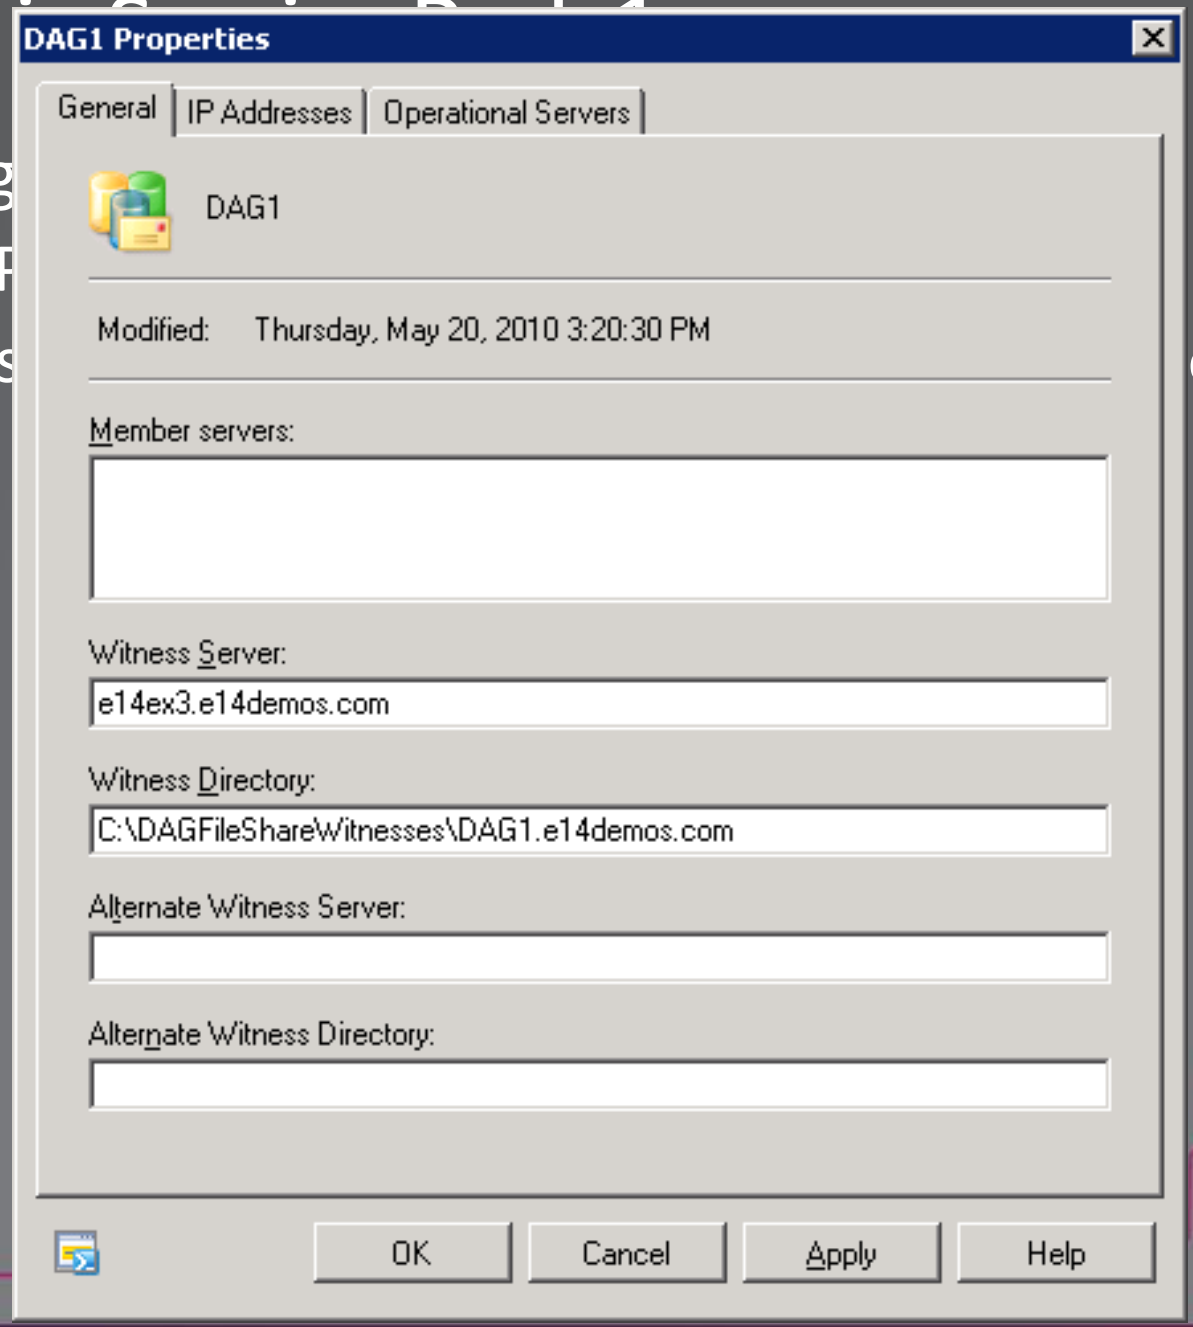

Microsoft®