

# ► Exchange 2013 High Availability and Site Resilience

Jeff Mealiffe  
Senior Program Manager  
Exchange Product Group

# High Availability

# High availability challenges

High availability focuses on database health

Best copy selection insufficient for new architecture

Management challenges around maintenance and DAG network configuration

# High availability enhancements

Managed Availability

Best Copy and Server Selection


Maintenance Mode

DAG Network Autoconfig


# Managed availability

# Managed availability

If a protocol goes down on a mailbox server, every active database loses access to that protocol



For most protocols, quick correction is provided through restart action



If restart fails, often a failover is triggered

- Protocols control recovery sequence
- Recovery sequence optimized thru Office 365 experience; Service experience accrues to enterprise!

# Managed availability

## Provides

- Reliable and scalable monitoring framework for Exchange components
- Broader perspective across groups of Exchange servers

## Provides

- Sequencing mechanism to control when recovery actions are done vs. when an alert is issued (and a human engaged)
- Common set of recovery actions

## Provides

- Set of enhancements to the best copy selection (BCS) process
- Mechanism for in/out of service for Mailbox and CAS (maintenance mode)

# Managed availability

## Restart

Service - kill and start a service; optional dump

AppPool - restart an app pool; optional dump

Server - bugcheck the machine

## Failover , Offline, Online

Database - failover a single active database

Server - failover all active databases

Protocol off - set health state for protocol to offline

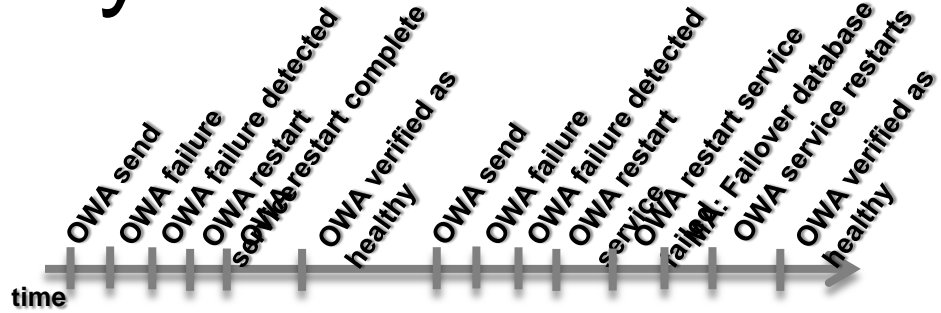
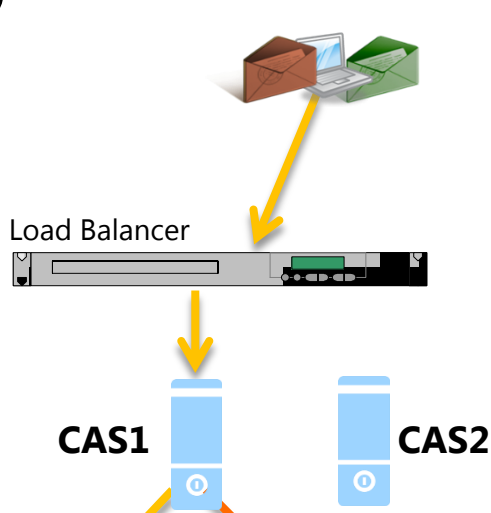
Protocol on - when a health set is green

## Escalate

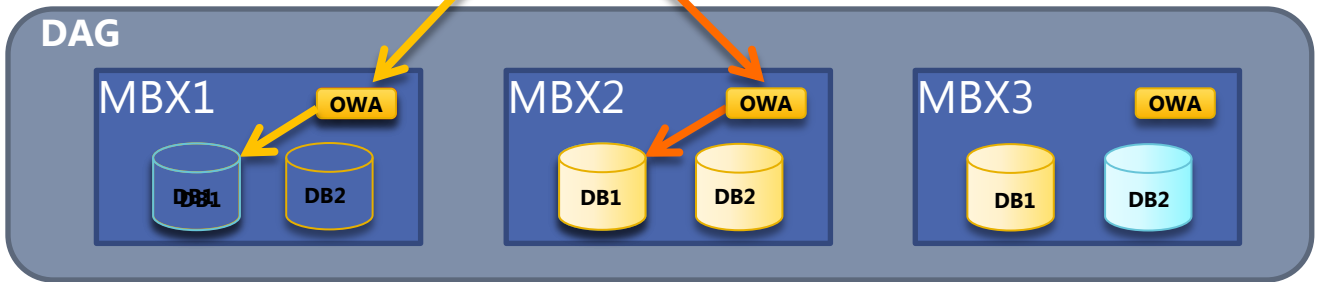
Notify a human of an issue



# Managed availability



Managed Availability = Monitoring + HA  
"Stuff breaks, but the Experience does not"



# Managed availability

- MA failovers are a recovery action from detected failure
  - Detected via a synthetic operation or live data
  - Throttled in time and across the DAG
  - Many failures resolved by a service or app pool restart; sometimes by a forced reboot
- MA failovers come in two forms
  - Server: Protocol failure can trigger server failover
  - Database: Store-detected database failure can trigger database failover

# Single copy alert

- Now runs natively as part of Managed Availability
- Alert is per-server to reduce flow
- Still triggered across all machines with copies
- Monitoring triggered through a notification
- Logs 4138 (red) and 4139 (green) events

# Best copy & server selection

# Best copy selection changes

- Process for finding the “best” copy of a specific database to activate
  - Inputs: list all copies, per-copy health and status
- Exchange 2010 used several criteria
  - Copy queue length
  - Replay queue length
  - Database copy status – including activation blocked
  - Content index status
- Not good enough for Exchange Server 2013, because protocol health is not considered

# Best copy and server selection

- Still an Active Manager algorithm performed at \*over time based on extracted health of the system
- Replication health still determined by same criteria and phases
- Criteria now includes health of the entire protocol stack
  - Considers a prioritized protocol health set in the selection
  - Four priorities – critical, high, medium, low (all health sets have a priority)
  - Failover responders trigger added checks to select a “protocol not worse” target

# Best copy and server selection

1

## All Healthy

Checks for a server hosting a copy that has all health sets in a healthy state

2

## Up to Normal Healthy

Checks for a server hosting a copy that has all health sets Medium and above in a healthy state

3

## All Better than Source

Checks for a server hosting a copy that has health sets in a state that is better than the current server hosting the affected copy

4

## Same as Source

Checks for a server hosting a copy of the affected database that has health sets in a state that is the same as the current server hosting the affected copy

# Maintenance mode



# Maintenance mode

- New functionality to support in/out of service
  - Server switchovers (move active copies; keep them from coming back)
  - Set-ServerComponentState to take CAS or Mailbox offline
- Out-Of-Service
  - Mailbox: no active databases + Transport is offline
  - CAS: per-protocol NLB health check; all proxy services offline
    - If state indicates offline then protocol does not acknowledge NLB health check
- Separate tracking for:
  - Health – MA triggered
  - Sidelined – operator initiated
  - Functional – setup running
  - Deployment – machine being configured

# DAG network autoconfig

# DAG network autoconfig

- Automatic or manual DAG network config
- Default is Automatic
  - Requires specific configuration settings on MAPI and Replication network interfaces
  - Manual edits and EAC controls blocked when automatic networking is enabled
  - Set DAG to manual network setup to edit or change DAG networks
- DAG networks automatically collapsed in multi-subnet environment

ENTERPRISE OFFICE 365

Recipients  
Permissions  
Compliance Management  
Organization  
Protection  
Mail Flow  
Mobile  
Public Folders  
Unified Messaging  
**Servers**  
Hybrid

https://e15cas1/?reqId=17561913&pmwid=2&ReturnObjectType=1&dtm=0&id=DAG%255C...

## ReplicationDagNetwork01

Help

\*Database availability group network name:  
ReplicationDagNetwork01

Description:

Subnets:  
+ ✎ -

SUBNET	STATUS
10.0.0/8	Up

Network interfaces:

NETWORK INTERFACE	STATUS
10.0.0.1	Up
10.0.0.2	Up

Enable replication

save cancel

100%

# Site resilience

# Site resilience challenges

Operationally complex

Mailbox and Client Access recovery connected

Namespace is a SPOF

# Site resilience enhancements

Operationally simplified

Mailbox and Client Access  
recovery independent

Namespace provides redundancy

# Site resilience

- Previously loss of CAS, CAS array, VIP, LB, some portion of the DAG required admin to perform a datacenter switchover
- In Exchange Server 2013, recovery happens automatically
  - The admin focuses on fixing the issue, instead of restoring service



# Site resilience

- Previously, CAS and Mailbox server recovery were tied together in site recoveries
- In Exchange Server 2013, recovery is independent, and may come automatically in the form of failover

# Site resilience

- DNS resolves to multiple IP addresses
- Almost all protocol access in Exchange 2013 is HTTP
- HTTP clients have built-in IP failover capabilities
- Clients skip past IPs that produce hard TCP failures
- Admins can switchover by removing VIP from DNS
- Namespace no longer a SPOF
- No dealing with DNS latency

# Site resilience


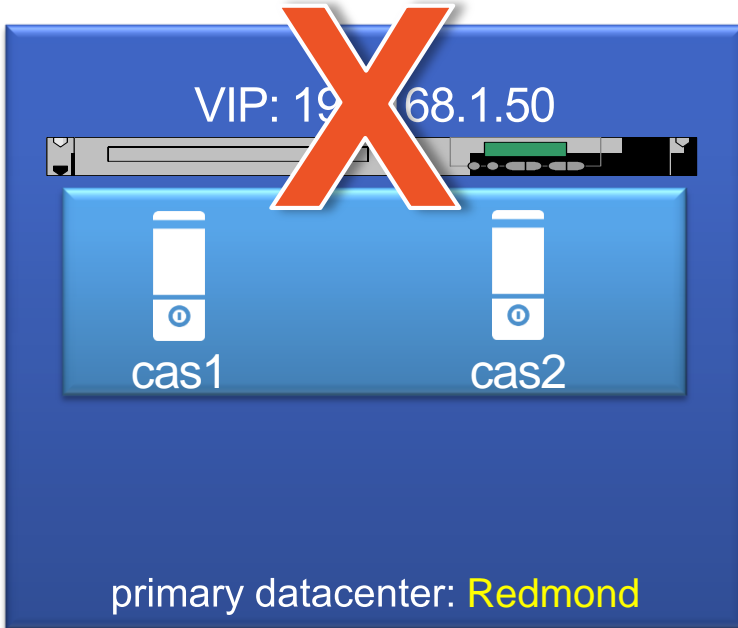
- With the namespace simplification, consolidation of server roles, separation of CAS array and DAG recovery, de-coupling of CAS and Mailbox by AD site, and load balancing changes...
- if available, three locations can simplify mailbox recovery in Exchange Server 2013 and provide datacenter failovers

# Site resilience

- You must have at least three locations
  - Two locations with Exchange; one with witness server
- Exchange sites must be well-connected
- Witness server site must be isolated from network failures affecting Exchange sites

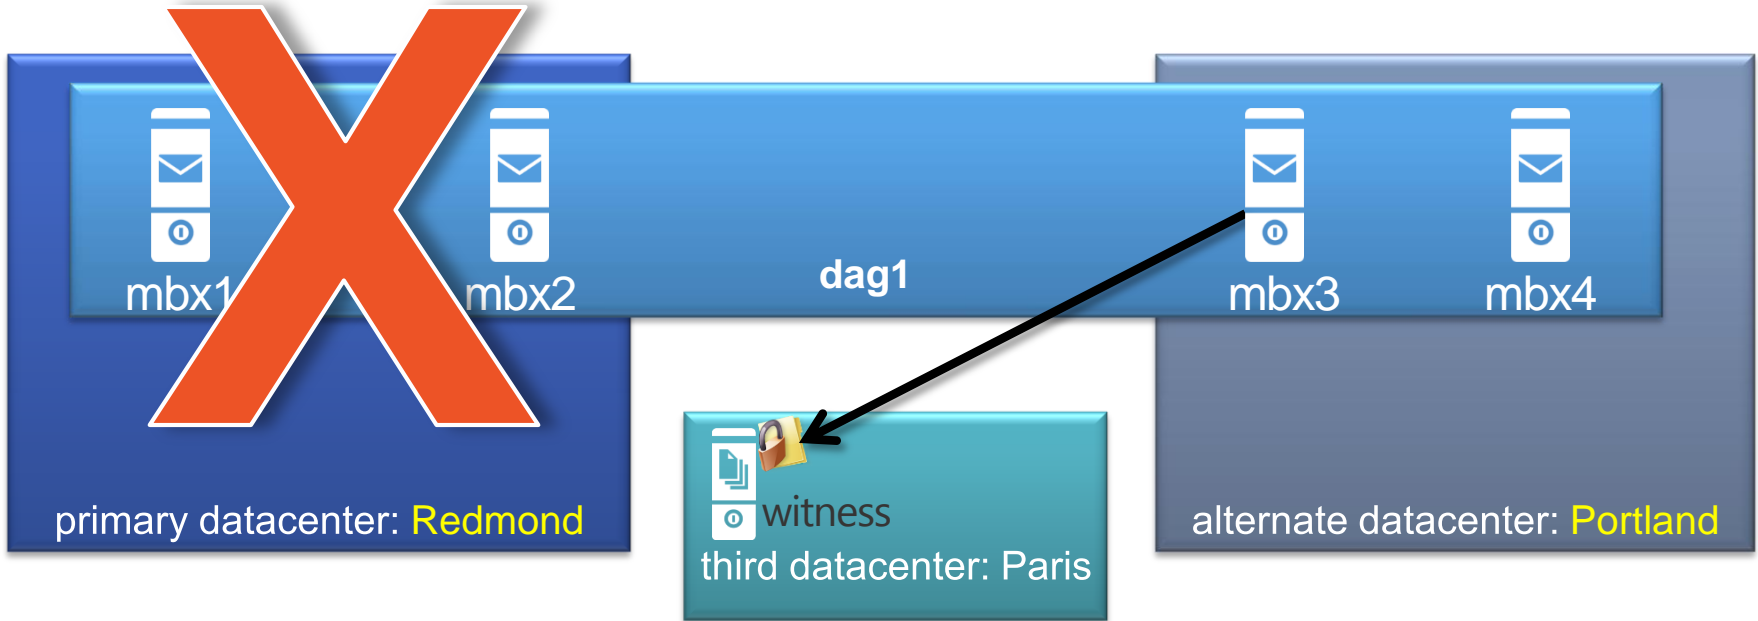
# Site resilience

With Redundancy, failover from one data center to the same center of service VIP fails, VIPs automatically failover to alternate VIP and just work!  
mail.com@comcast.com 192.168.0.50, 50.0.1.50



# Site resilience

Assuming MBX3 and MBX4 are operating and one of them can lock the witness.log file, *automatic failover should occur*



# Site resilience



# Site resilience

1. Mark the failed servers/site as down: `Stop-DatabaseAvailabilityGroup DAG1 -ActiveDirectorySite:Redmond`
2. Stop the Cluster Service on Remaining DAG members: `Stop-Clussvc`
3. Activate DAG members in 2<sup>nd</sup> datacenter: `Restore-DatabaseAvailabilityGroup DAG1 -ActiveDirectorySite:Portland`





# Questions?

# Disclaimer

© 2012 Microsoft Corporation. All rights reserved. Microsoft, Office 365, and other product and service names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries.

The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.