uc4you
unified communication

Exchange 2013
Server Architecture: Part 2

Jeff Mealiffe
Senior Program Manager
Exchange Product Group

Microsoft

Agenda

- Part 1
  - Overview of the new Architecture
  - The Client Access server role

- Part 2
  - The Mailbox server role
  - Transport Architecture
  - Service Availability Improvements

Exchange 2013 server role architecture

2 Building Blocks

Client Access Array

- Evolution of E2010 CAS Array
- SMTP Front-End

Database Availability Group

- Evolution of E2010 DAG
- Includes core server protocols

Loosely coupled

- Functionality
- Versioning
- User partitioning
- Geo affinity

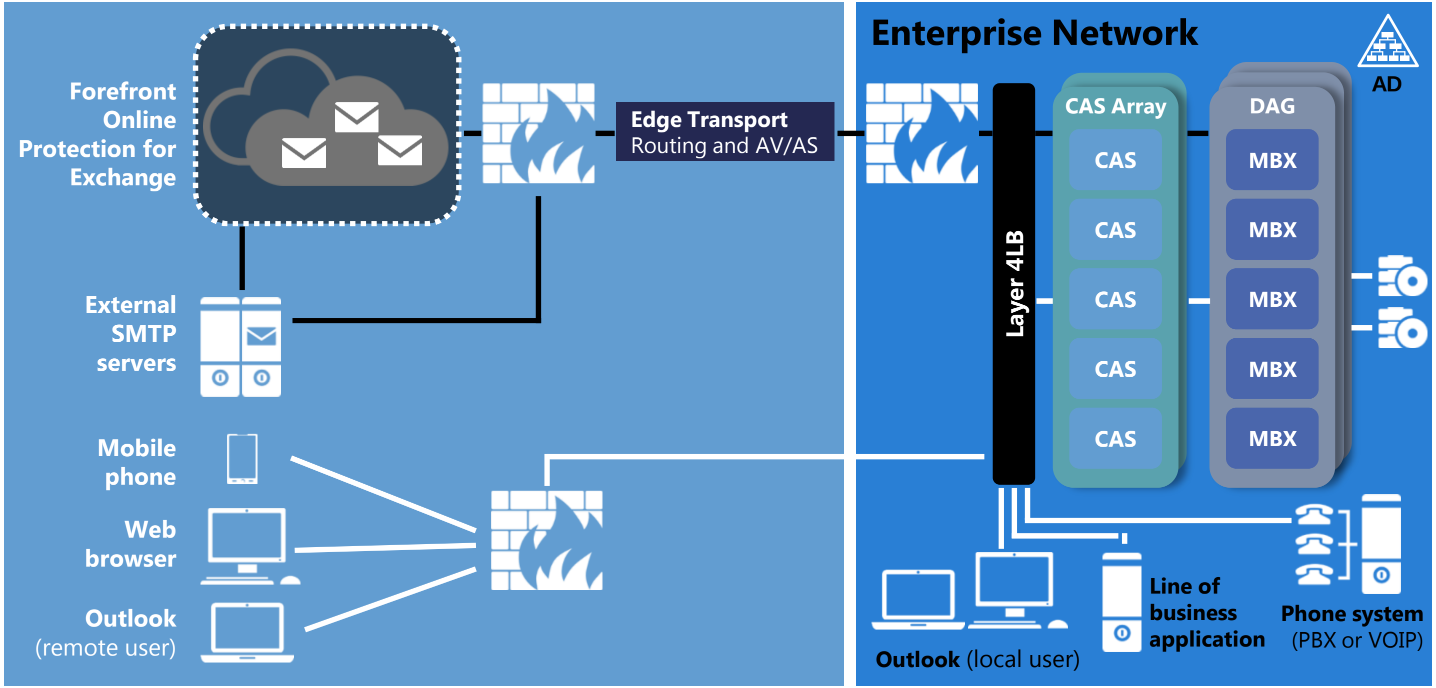

uc4you unified communication

The mailbox server role

The mailbox server role

- A server that hosts all the components that process, render and store the data
- Clients do not connect directly to MBX2013 servers; connectivity is through CAS2013
- Evolution of E2010 DAG
  - Collection of servers that form a HA unit
  - Databases are replicated between servers in a given DAG
  - Servers can be in different locations, for site resiliency
  - Maximum of 16 Mailbox servers
  - 50 database copies / server

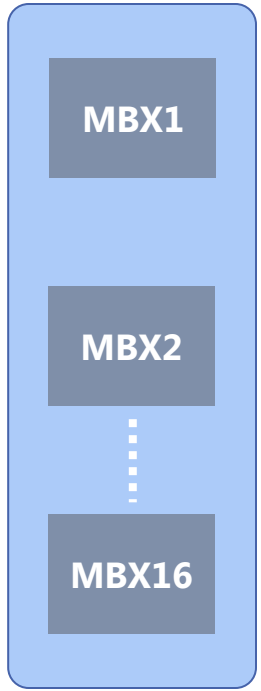

The new store process

- Store is effectively made up of three processes
  - Replication service
  - Store service process/controller
  - Store worker process
- Replication service initiates failovers and is responsible for issuing mount/dismount operations
- Store service process/controller manages the store worker processes
- Each database has its own Store worker process

Exchange IOPS trend

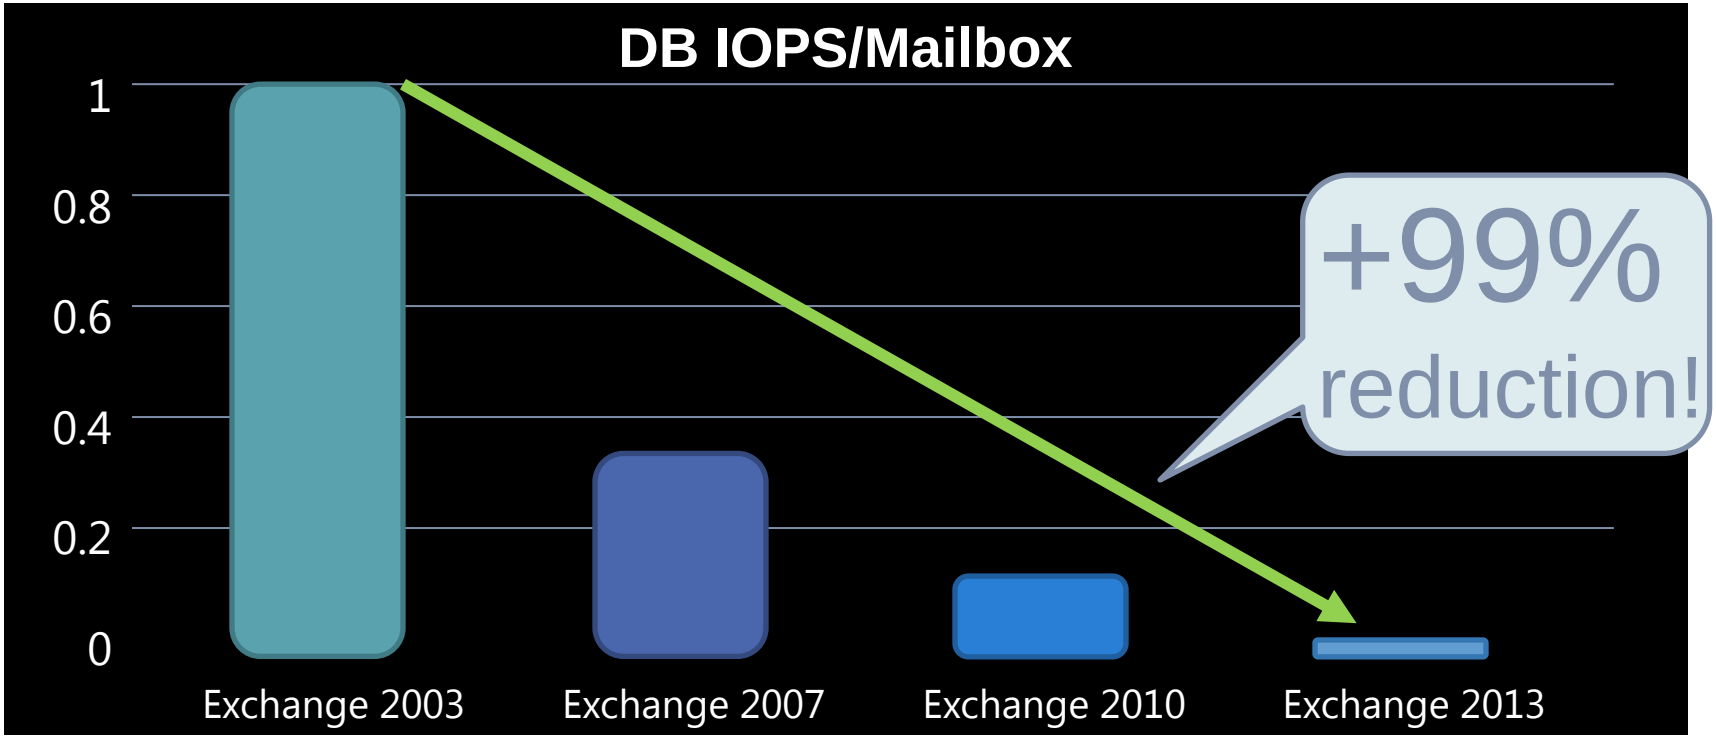

uc4you
unified communication

Large mailboxes for the win!

- Large Mailbox Size 100GB+
  - Aggregate Mailbox = Primary Mailbox + Archive Mailbox + Recoverable Items
  - 1-2 years of mail (minimum)
- Increased knowledge worker productivity
- Eliminate or reduce PST reliance
- Eliminate or reduce third-party archive solutions
- Outlook 2013 allows you to control OST size!
  - Gives more options around mailbox deployments

Time	Items	Mailbox Size
1 Day	150	11 MB
1 Month	3300	242 MB
1 Year	39000	2.8 GB
2 Years	78000	5.6 GB
4 Years	156000	11.2 GB

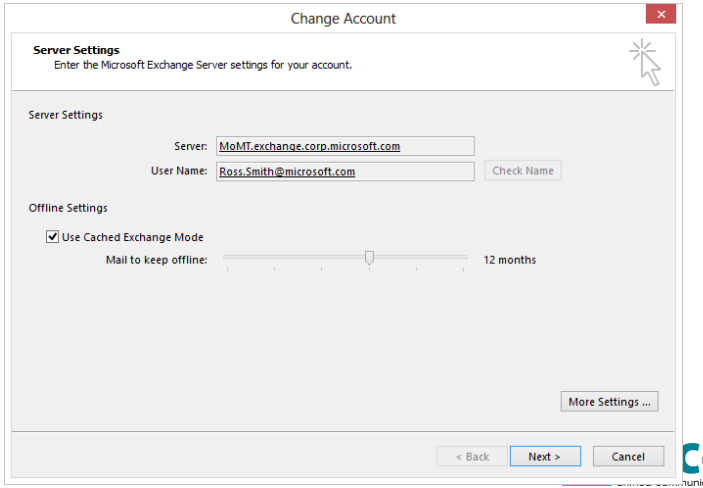

New Exchange search infrastructure

- Leverages Search Foundation
  - Common, actively developed search platform used across Office server products
  - Does consume more memory (1/6 available memory) to improve query performance
- Provides
  - Significantly improved query performance compared to E2010
  - Significantly improved indexing performance compared to E2010
- Feature parity with E2010 search
- Leverages the same cmdlets like Get-MailboxDatabaseCopyStatus

Exchange indexing

Reduced processing of body and attachments

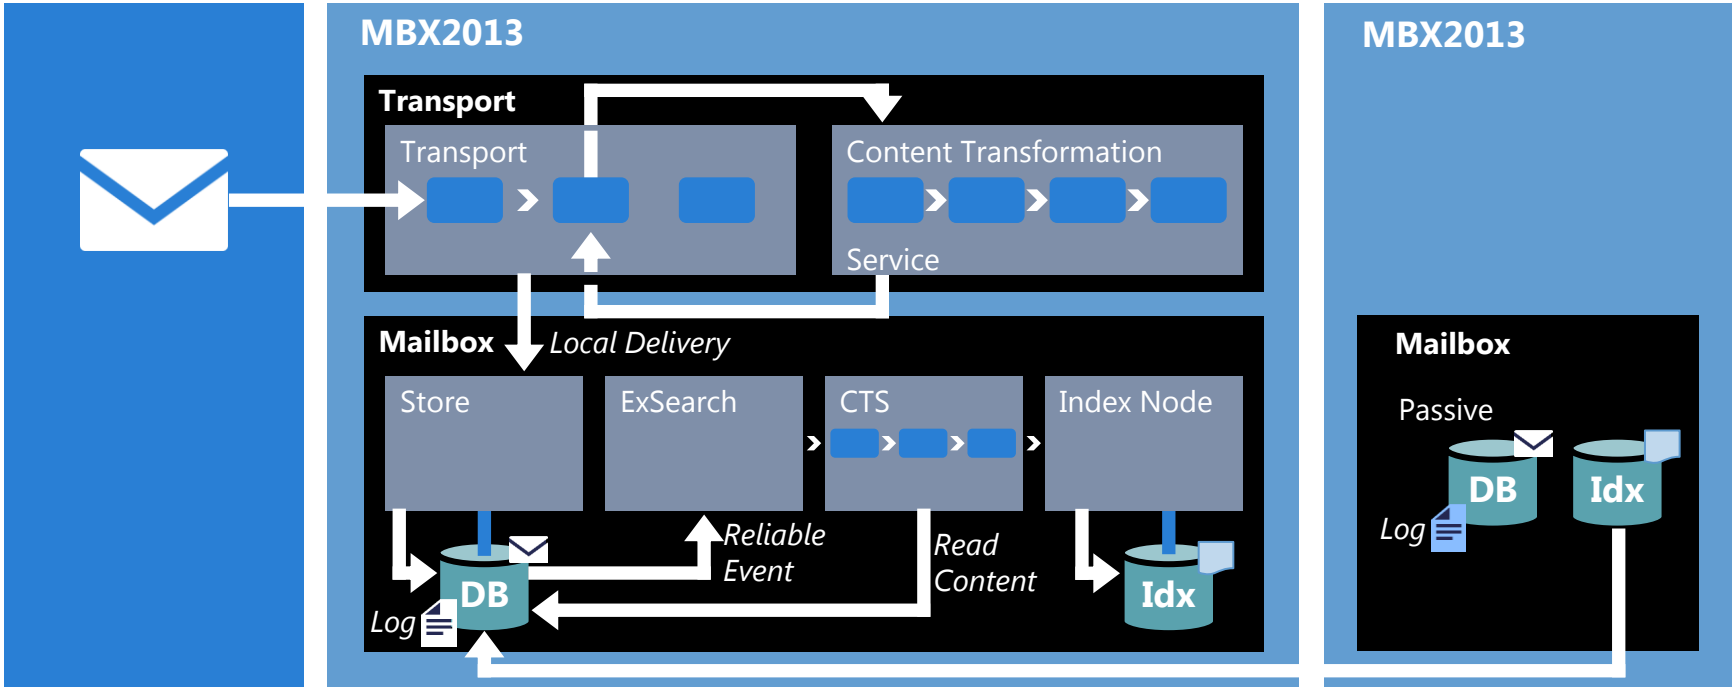

Public folders

Dawn of a new age

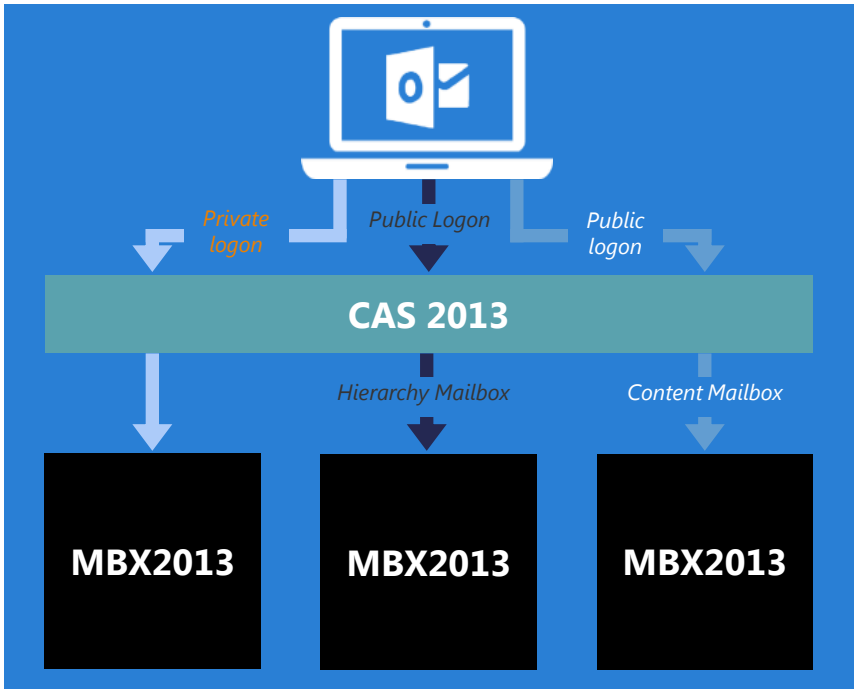

Architectural bet

Public folders are based on the mailbox architecture

Details

- Hierarchy is stored in PF mailboxes (one writeable)
- Content can be broken up and placed in multiple mailboxes
- The hierarchy folder points to the target content mailbox
- Uses same HA mechanism as mailboxes
- No separate replication mechanism
- Single-master model
- Similar administrative features to current PFs (setting quota, expiry, etc.)
- No end-user changes (looks just like today's PFs)

Not all public folder usage scenarios are best served by public folders

Transport architecture

Transport components on client access

Front-end transport service

- Handles all inbound and outbound external SMTP traffic for the organization, as well as client endpoint for SMTP traffic
  - Does not replace the Edge Transport Server Role
- Functions as a layer 7 proxy and has full access to protocol conversation
- Will not queue mail locally and will be completely stateless
- All outbound traffic appears to come from the CAS2013
- Listens on TCP25 and TCP587 (two receive connectors)

Processing Inbound Messages

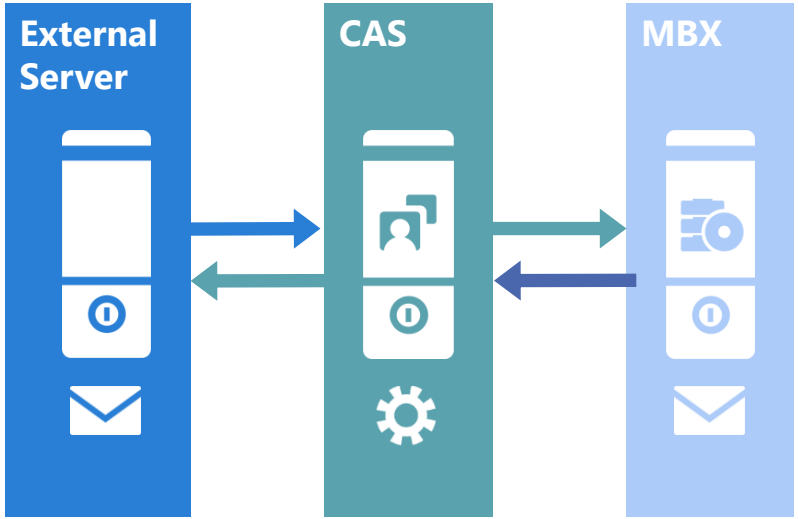

1. New SMTP Connection
2. CAS performs envelope filtering
3. CAS determines route to best MBX server
4. Message delivery begins
    1. If successful, CAS returns 250 OK acknowledgement to external server
    2. If unsuccessful, CAS returns 421 response

Benefits of SMTP front-end service

- The SMTP Front-End Service provides:
  - Protocol level filtering – performs connection, recipient, sender and protocol filtering
  - Network protection – centralized, load balanced egress/ingress point for the organization
  - Mailbox locator – avoids unnecessary hops by determining the best MBX2013 to deliver the message
  - Load balanced solution for client/application SMTP submissions

- Scales based on number of connections – just add more servers

Transport components on mailbox

- Transport in MBX2013 has been broken down into three components
  - Transport Service - Stateful and handles SMTP mail flow for the organization and performs content inspection (Was previously referred to as “Hub Transport”)
  - Mailbox Transport Delivery Service - Receives mail from the Transport service and deliveries to the Mailbox Database
  - Mailbox Transport Submission Service - Takes mail from the Mailbox Databases and submits to the Transport service
- Mailbox Transport performs content conversion

Transport components on mailbox

Responsibilities

- Receives all inbound mail to the organization (Proxied through CAS or direct)
- Submits all outbound mail from the organization (Proxied through CAS or direct)
- Handles all internal message processing such as Transport Rules, Content Filtering, and Anti-Virus
- Performs mail flow routing
- Queue messages
- Supports SMTP extensibility

Routing optimizations

- Next hop selection is broken down into distinct delivery groups:
  - Routable DAG
  - Mailbox Delivery Group
  - Connector Source Servers
  - AD Site (Hub Sites; Edge Subscriptions)
  - Server list (DG expansion servers)
- Queuing is per delivery group, connector, or mailbox
- Once message is received at final destination, Transport will deliver the message via SMTP to Mailbox Transport on the server hosting the active database copy
- Send/Delivery-Agent Connectors can have source servers from multiple DAGs or AD Sites, and can be proxied through CAS

Mail delivery

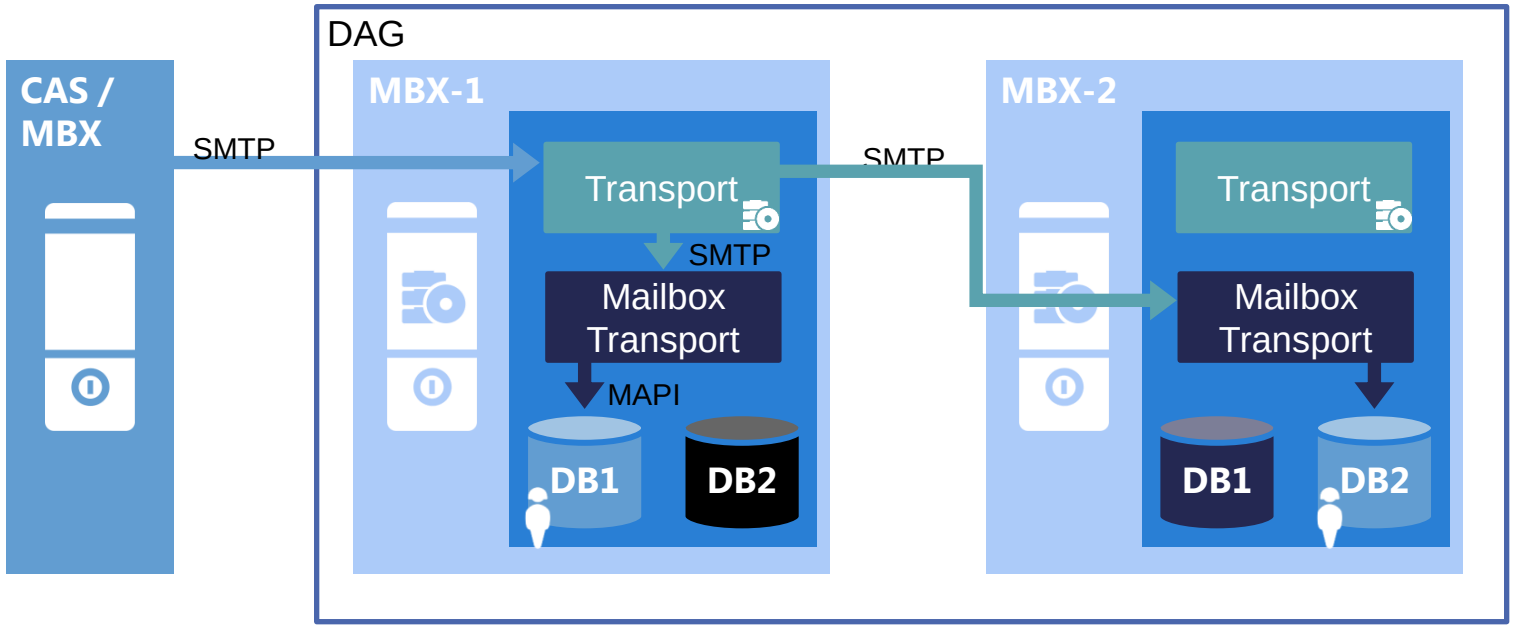

Service availability improvements

uc4you
unified communication

Managed availability

- Monitoring and recovery infrastructure is integrated with Exchange's high availability solution
- Detects and recovers from problems as they occur and are discovered
- Is user focused – if you can't measure it, you cannot monitor it

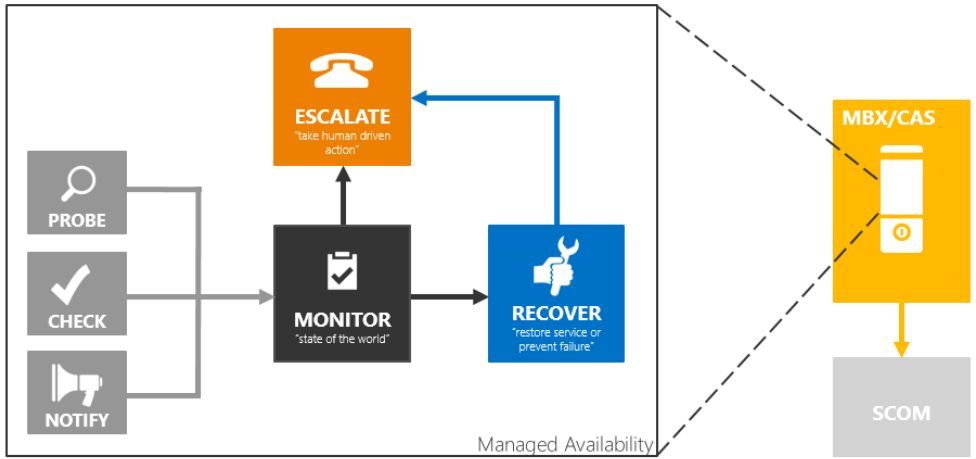

Transport High Availability Improvements

- Every message is redundantly persisted before its receipt is acknowledged to the sender
- Delivered messages are kept redundant in transport similar to active messages
- Every DAG represents a transport HA boundary and owns its HA implementation
  - If you have a stretched DAG, you also have transport site resilience
- Resubmits due to transport DB loss or MDB *over are fully automatic and do not require any manual involvement

Transport HA

1. Maintain a copy of the message in the queue database but don't acknowledge the DATA verb
2. Generate a shadow copy on another MBX2013 server in the DAG (remote site preferred)
3. Wait for acknowledgement from the shadow server
4. Send acknowledgement to SMTP client
5. Delete message from queue after SafetyNet threshold has expired

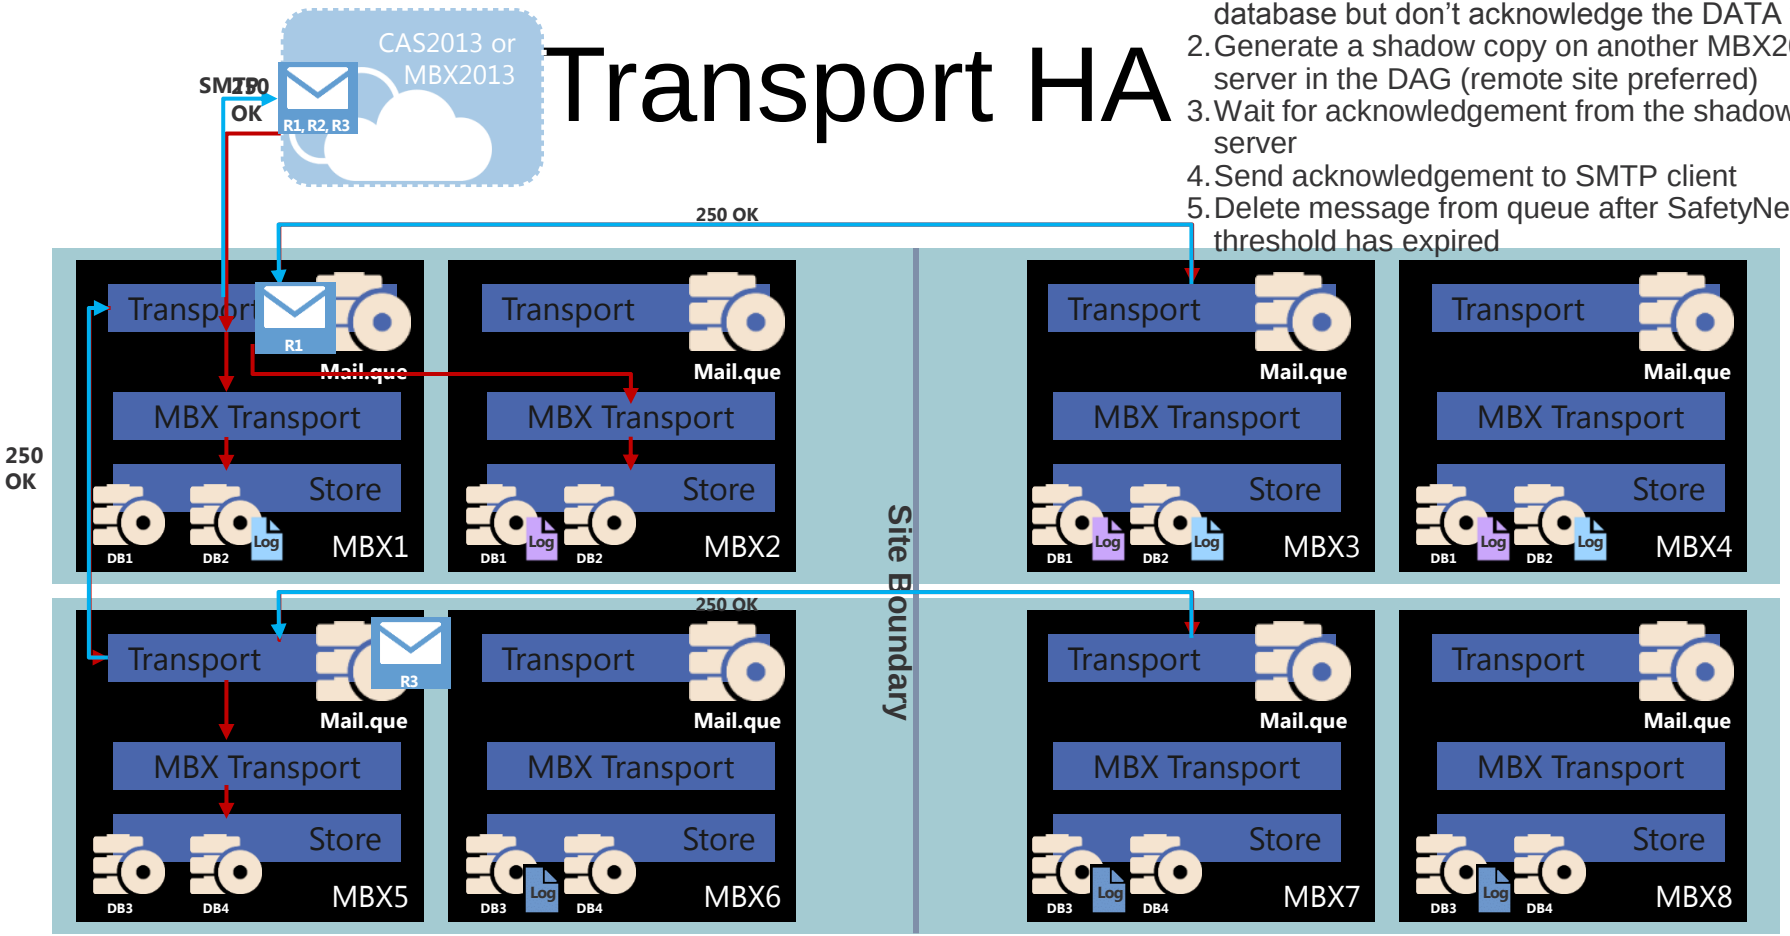

uc4you unified communication

Summary

New Building Blocks

- Facilitates deployments at all scales – from self-hosted small organizations to Office 365
- Provides more flexibility in namespace management

Simplified HA

- All core Exchange functionality for a given mailbox is served by the MBX2013 server where that mailbox's database is currently activated
- Simplifies the network layer
- Transport protection is built-in

Simplified upgrade and inter-op

- All components in a given server upgraded together
- No need to juggle with CAS <-> MBX versions separately

Aligned with hardware trends

- Utilize CPU core increase, cheaper RAM
- Utilize capacity effectively
- Fewer disks/server => simpler server SKUs

Questions?

uc4you
unified communication

Disclaimer

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.

uc4you unified communication