Ask The Perf Guy:
Sizing Guidance for Exchange 2013 SP1

Jeff Mealiffe
Principal Program Manager
Microsoft

Agenda

Impact of 2013 architecture

How to size Exchange 2013 roles

The sizing model: caveats

Trust, but verify

Impact of 2013 architecture

Exchange architecture overview

Exchange 2013 targets balanced use of hardware

Consider hardware platforms that provide the right balance of resources

3 roles for sizing: Mailbox, Client Access Server, Active Directory

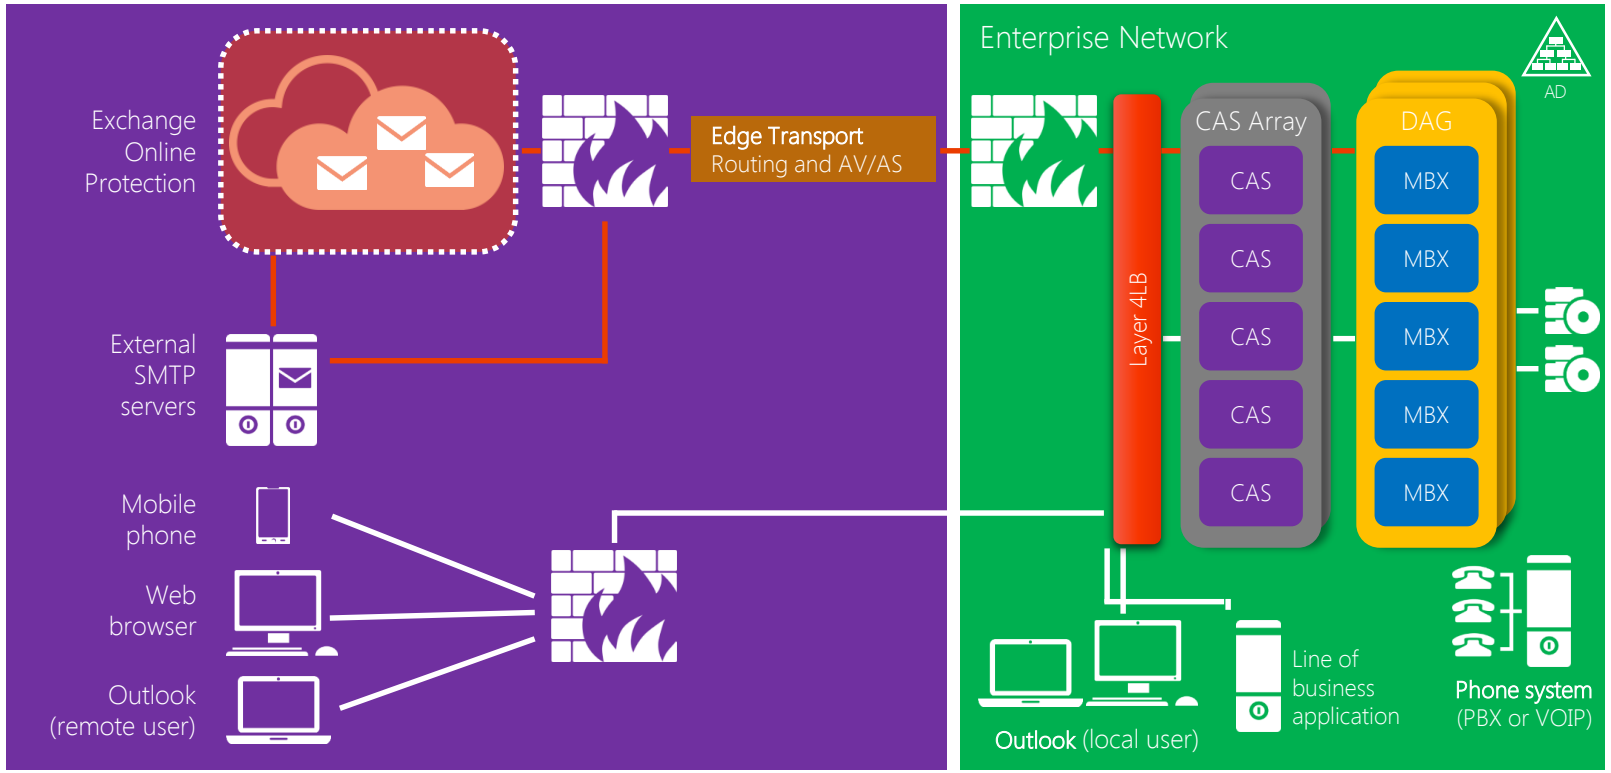

Mailbox role overview

Hub & CAS roles were deprecated*

Mailbox takes over responsibility for the old Hub & CAS components

Benefits	"Considerations"
Simplified deployment & connectivity model	Increased resource usage (see managed availability, content indexing in particular)
Cache efficiencies	Everything interacts (and workload management mediates)
Hardware efficiencies (balanced resource utilization)	Unified Messaging on all Mailbox servers

CAS role reintroduction

We brought back the Client Access Server as an intelligent proxy

Benefits	"Considerations"
Stateless	Net-new role in 2013, adds performance "cost" – can potentially mitigate by shifting some resources from LB to CAS
Load balancing optimizations	
Low CPU & memory footprint	
Connection scalability	

How to size Exchange 2013 roles

First, use the calculator

http://aka.ms/E2013Calc

Calculator turns published sizing guidance into a modeling tool

Try out various failure scenarios

Understand the impact of different hardware & storage choices

Provides scripts for DAG, database & copy creation

Many new features

CAS sizing
Transport storage sizing
Multiple databases per-volume (JBOD) support
High availability architecture improvements

Note: Baseline platform for CPU guidance changed in 2013. Don't directly compare results from 2010 & 2013 calculators.

Multi-role: just do it

Very few reasons not to consider multi-role (Mailbox+CAS) deployment

Multi-role simplifies deployment, can reduce server count

Benefit of increased availability at the CAS layer

Issues remain with Windows NLB + DAG (WSFC)

Certificate management may be a concern

Storage capacity requirements

Size for mailbox size on disk, content indexes, log space

Method for computing space requirement similar to Exchange 2010, with some important changes

20% database overhead is now 0%

CI size is now 20% of EDB

Plus space for additional index set per volume (master merge)

Note that impact of space for master merge is reduced with multiple DBs per-volume

May find that .edb is smaller than sum of mailbox size

IOPS Requirements

As in previous releases, Exchange 2013 reduced IOPS requirements (~33% reduction compared to 2010)

We have seen higher reduction in various tests, guidance is conservative and based on production observations

No separate guidance for HA vs. non-HA databases

Checkpoint depth is now consistent for all scenarios, so IOPS requirements are the same

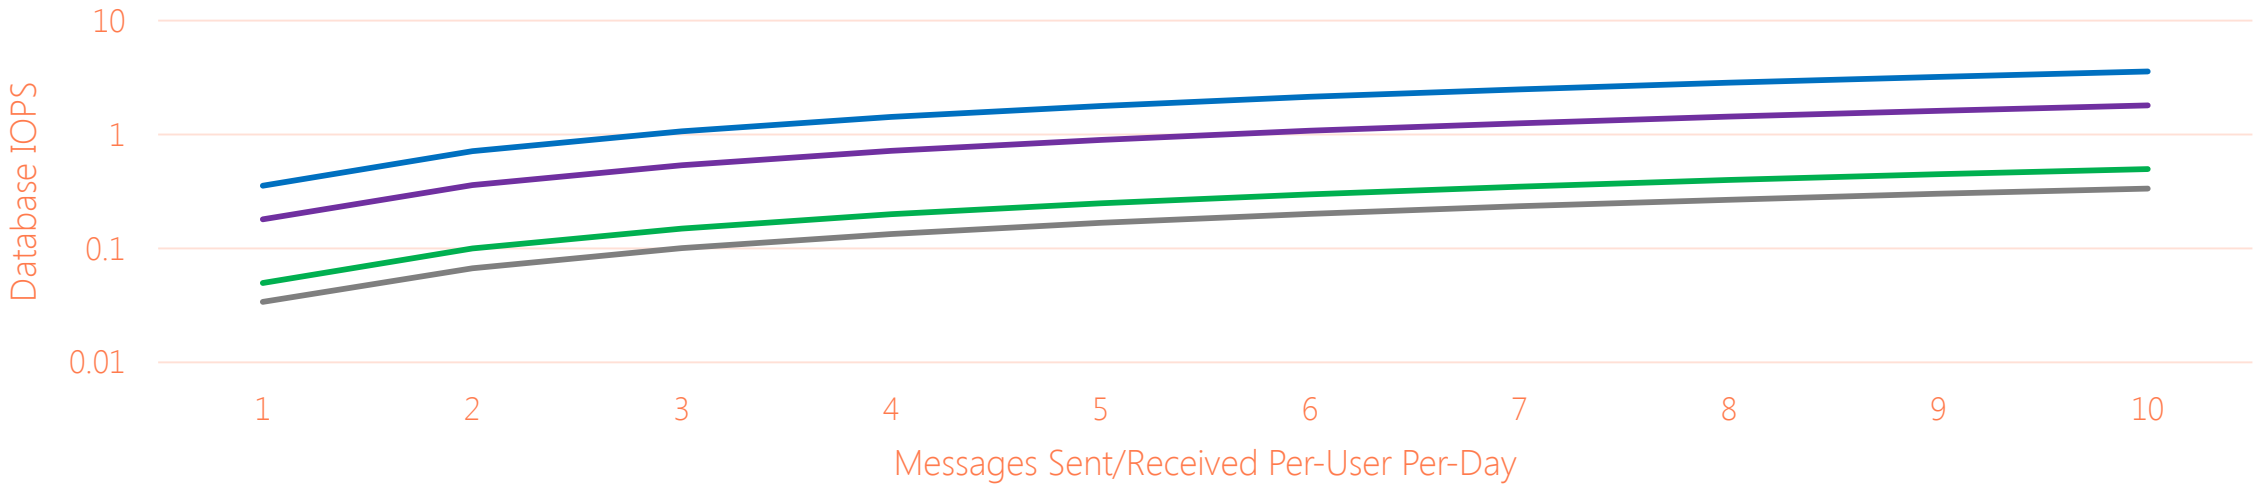

Transport storage requirements

Transport capacity requirements include queue and Safety Net

Guidance demonstrates method for calculating capacity requirements

Transport queue database takes advantage of ESE IO improvements to reduce IOPS

Microsoft production observations show ~1 DB IO per 75KB message

Low IOPS suggest that placing transport queue on system/install volume is now feasible in many scenarios

Significant transport throughput benefit seen from a protected write cache disk controller, set to 100% write cache

Processor requirements

As in Exchange 2010, mcycle requirements are per-user for active & passive copies

Per-passive multiplier on the active has been removed in 2013

Guidance includes a multi-role mcycle requirement for the active copy – simplifies sizing

Messages sent or received per mailbox per day	Mcycles per User, Active DB Copy or Standalone (MBX only)	Mcycles per User, Active DB Copy or Standalone (Multi-Role)	Mcycles per User, Passive DB Copy
50	2.13	2.93	0.69
100	4.25	5.84	1.37
150	6.38	8.77	2.06
200	8.50	11.69	2.74
250	10.63	14.62	3.43
300	12.75	17.53	4.11
350	14.88	20.46	4.80
400	17.00	23.38	5.48
450	19.13	26.30	6.17
500	21.25	29.22	6.85

Note: Baseline platform for CPU guidance changed in 2013. Mcycle requirements in 2010 and 2013 cannot be directly compared.

Hyperthreading & Exchange 2013

Turn off hyperthreading (SMT)!

SMT provides gain in processor throughput, but overall the gain is not worth the "cost" based on our lab measurements

Significant impact to some Exchange service memory footprints

Tuning .NET for store

Best practice to install .NET 4.5.2 (or later)

Contains many fixes that benefit Exchange

Reduces memory consumption in each store worker (note no impact to sizing guidance)
Decreases CPU spent in .NET garbage collector
Benefits Mailbox & multi-role

Extremely important on CAS role for MAPI/HTTP

Memory requirements

Memory on the Mailbox role sized based on ESE cache requirements

Cache requirements have remained constant from 2010

Overall cache sized to 25% of RAM, so guidance (based on total system memory) is 4x of 2010 cache sizing recommendation

Messages sent or received per mailbox per day	Mailbox role memory per active mailbox (MB)
50	12
100	24
150	36
200	48
250	60
300	72
350	84
400	96
450	108
500	120

Memory requirements

Multi-role servers require additional memory for CAS based on user concurrency during worst-case failure

$$2\text{GB} + \left(2\text{GB} \times \frac{(\text{worst-case active DBs per-server} \times \text{users per-DB} \times \text{mbx mcycles per-user}) \times 0.375}{\text{per-core mcycles}}\right)$$

Minimum memory requirements based on database count must be observed to ensure optimal ESE cache utilization

Per-server DB copies	Minimum physical memory (GB)
1-10	8
11-20	10
21-30	12
31-40	14
41-50	16

Pagefile guidance

SP1

Old guidance was “size of RAM + 10MB”

Exchange servers typically have *lots* of RAM, resulting in ridiculous pagefile sizes

New guidance for Exchange 2013 is to continue to use a fixed pagefile size, but cap pagefile at 32778 MB if using > 32GB RAM

```
MIN(RAM+10MB,32778MB)
```

CAS processor requirements

CAS CPU is sized using a percentage of Mailbox CPU active user requirements

2013 CAS requires 37.5% of Mailbox active-user mcycles, down from 75% in 2010

Given significant reduction, ensure that enough CAS servers are deployed to handle failures and provide high availability (particularly in small sites)

CAS memory requirements

CAS memory is sized using a simple formula of 2GB + 2GB per-CPU core.

The per-core value assumes utilized CPU cores at peak (worst case failure), so this can get a little complicated

$$\text{Per-server CAS memory} = 2\text{GB} + 2\text{GB} \times \left(\frac{\frac{\text{total user count}}{\text{CAS server count in worst case}} \times \text{Mailbox mcycles per-user} \times 0.375}{\text{mcycles per-core}} \right)$$

Note no CAS memory reduction from 2010, but decreased CAS server count should result in overall memory reduction

Active Directory requirements

Recommend deploying 1 AD GC core for every 8 Mailbox cores handling active load (assuming 64-bit GCs

Size memory such that the entire NTDS.DIT can be contained within RAM for optimal query performance

Size does matter

How big is too big?

Design for scale-out, not scale-up

Better alignment with intentions & design points of PG

Ideally focus on "commodity" 2U servers as a platform to help minimize deployment risk

We don't push the "top end" today – and don't want you to either

The sizing model: caveats

We don’t cover everything

Sizing data is limited to the deployments we use to build our models

Not all client types or versions are covered

3rd party solutions generally not included

LOB applications

Hardware variations

Ongoing product changes

Feature enablement/usage

Trust, but verify

Jetstress & Exchange Solution Reviewed Program

Jetstress 2013 released March 2013

http://aka.ms/jetstress2013

Event log captured

Errors associated with specific volumes

Threads controlled globally instead of per-DB, better automatic tuning

Use Jetstress to validate all Exchange capacity before service ready

Validates storage performance & reliability

ESRP Storage v4.0 released last May to storage partners

http://aka.ms/esrp2013

~30 solutions available on ESRP site for Exchange 2013

Summary

Exchange 2013 requires more HW resources than prior releases

Plan to deploy multi-role

Use the calculator

Be careful w/3rd party products & "exceptional" user profiles

Scale out, not up

Run Jetstress

More details available on the blog

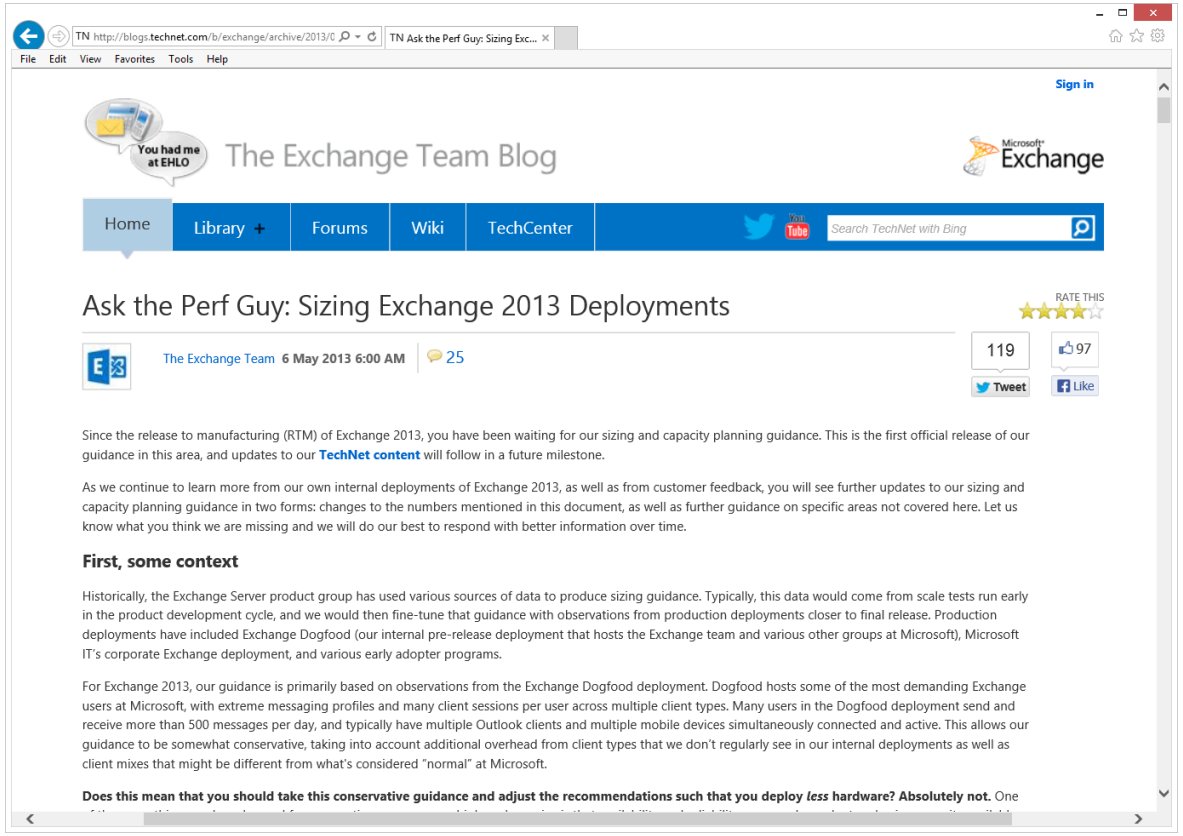

http://aka.ms/Exchange2013SizingGuidanceBlog