
hp
invent

Luc Vogeleer

Global Program Manager

Microsoft Service Operation

Exchange-2000 in the DataCenter

Microsoft Exchange

History

Q2 1996 4.0 Client/Server replacement for MS-Mail

Q1 1997 5.0 Internet protocol

Q4 1997 5.5 Unlimited store, clustering

2000 – Exchange-2000

Exchange-2000

A new technology generation

Active Directory replaces the Directory Services

- The GAL is replaced by the GC

SMTP becomes the routing engine

Partitioning the information store

URL everywhere to access the Web Store

Events everywhere

- Transport
- Store

Exchange-2000 Storage Management

- One Store process per server
- Up to 4 ESE instance per store process
- Shared transaction logs
- Up to 5 MDB per ESE instance
- A MDB consist in a STM and EDB file
  - STM: Streaming Internet Content
  - EDB: MAPI and Properties

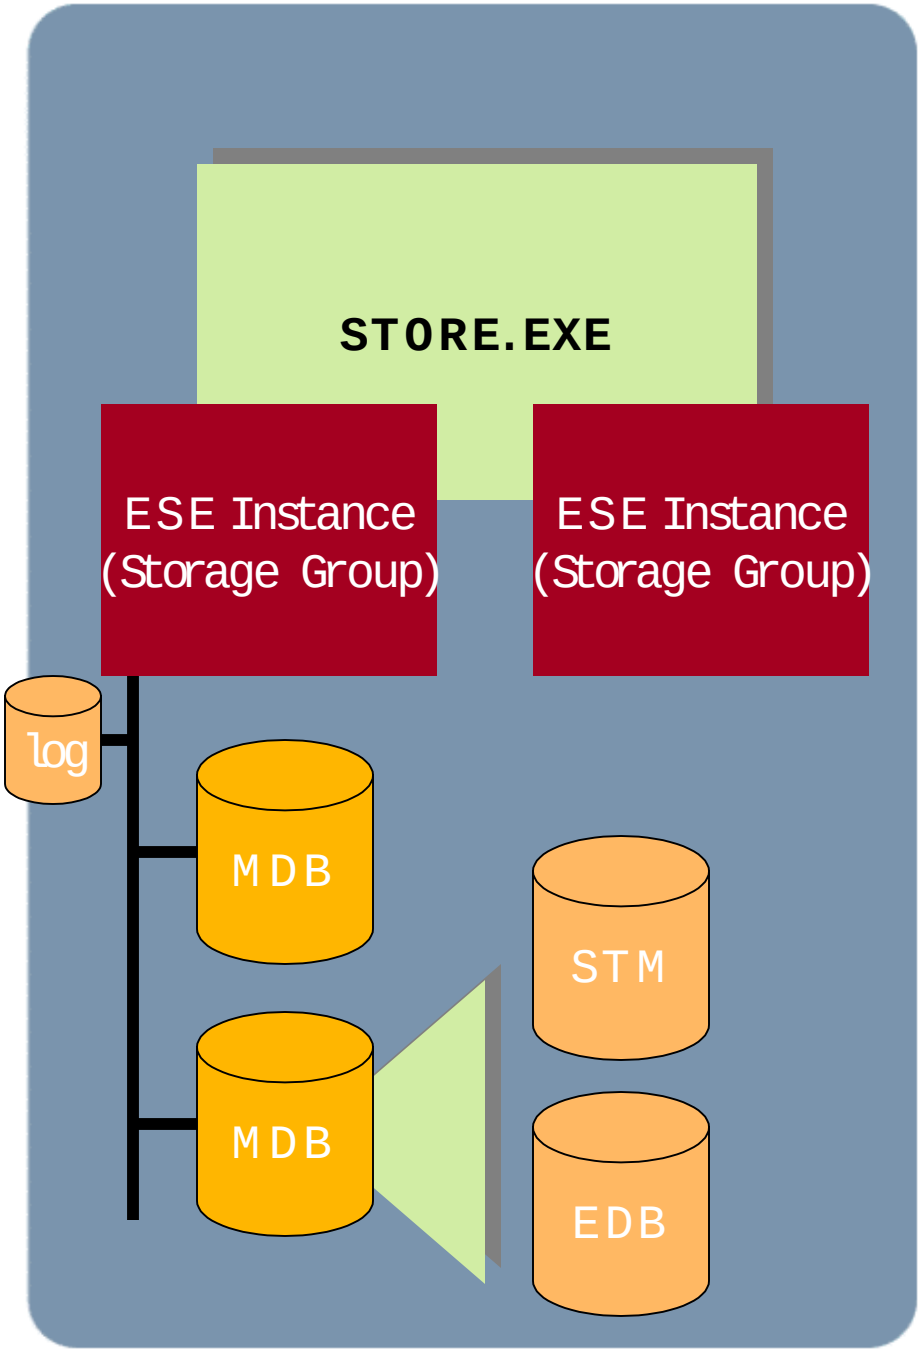

Exchange 2000 Databases

EDB Files

- Properties Database
- MAPI Messages and Attachments
- Headers for STM pages

STM Files

- Raw 'streaming' data (MIME, documents, multimedia, etc)

Current database =
EDB + STM +
Unflushed Log entries

Transaction Logging
(*.LOG)

"Write-ahead" logging

- Cached version of database pages modified but not written to disk immediately
- Appending to the log is fast
- Writing to database files is batched for optimal performance

Checkpoint file (*.chk) tracks location of the last transaction in log files flushed to the database (.edb or .stm)

A database is "consistent" when all transactions have completed and been flushed to the database

Front-end / Back-end

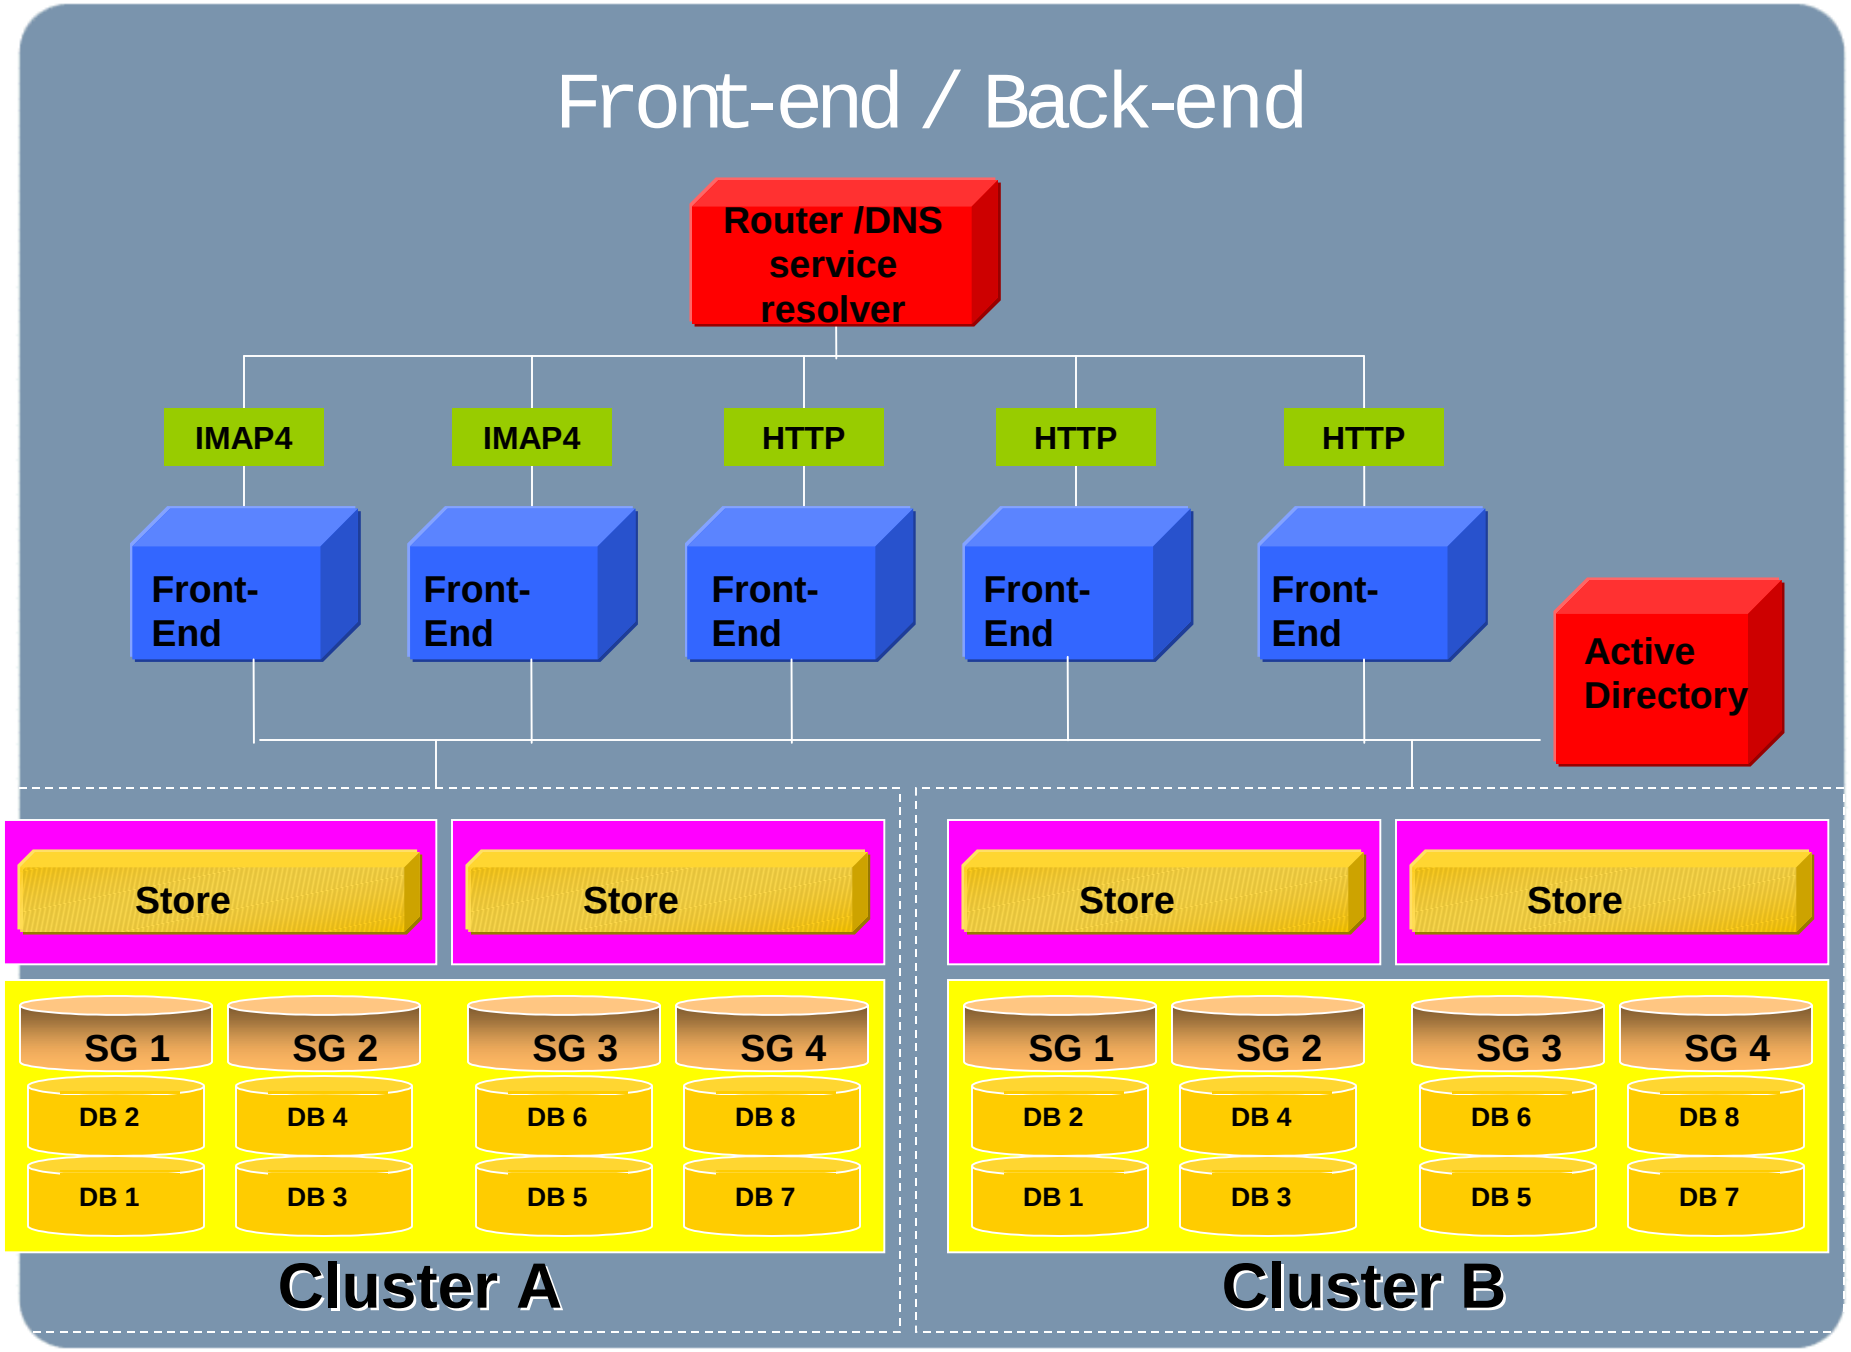

Exchange 2000 Clustering

Designed for *n*-node clustering

Clustering is Active/Active

EVS (Exchange Virtual Server) is the unit of failover

All major components act as cluster resources

- HTTP, IMAP, SMTP, Storage Group, etc.

Clustering and Storage
Active/Active (2 node only)

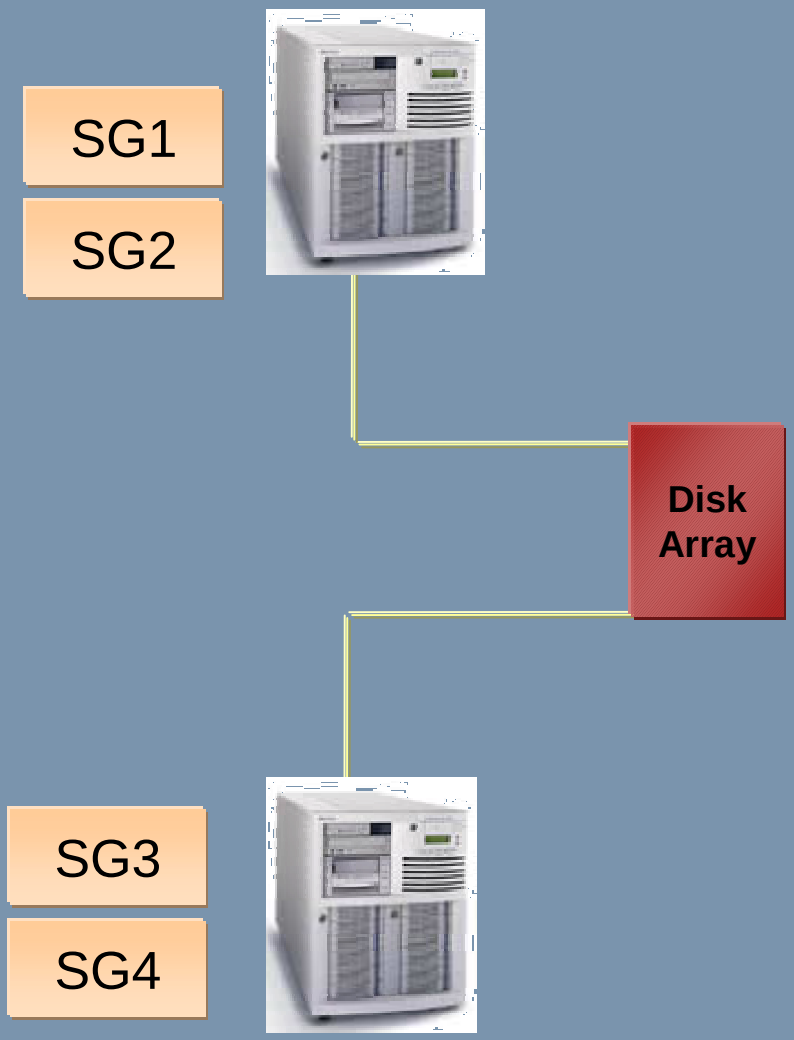

Clustering and Storage
Active/Active (2 node only)

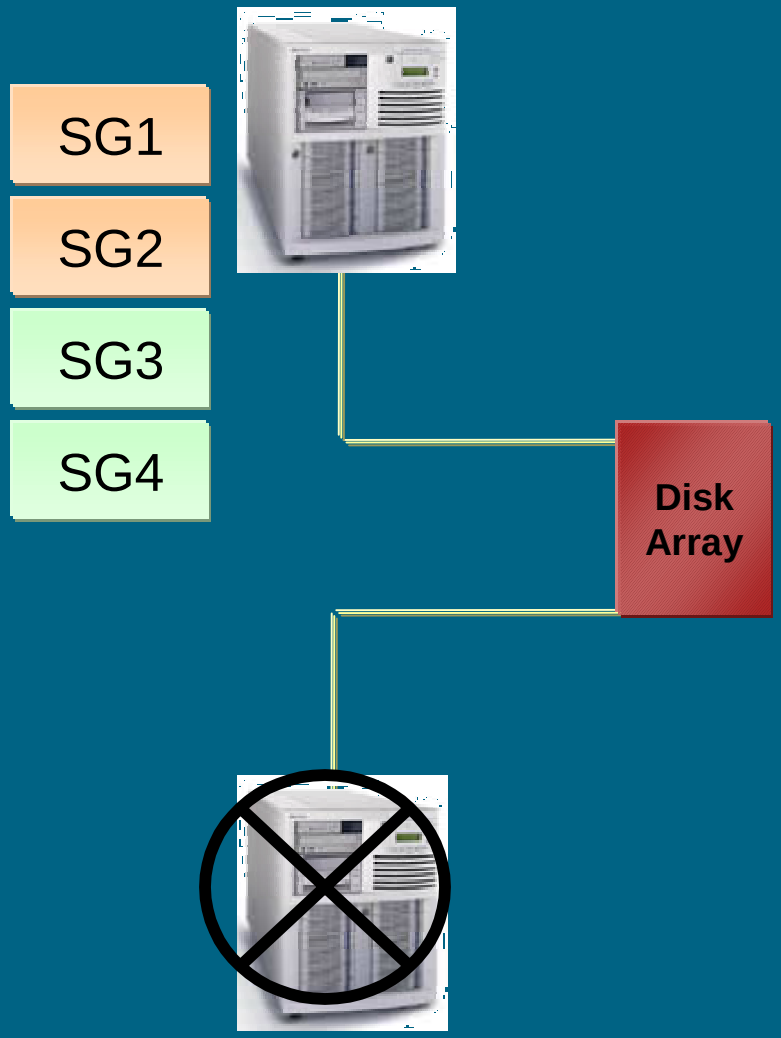

Clustering and Storage active/passive (n+1)

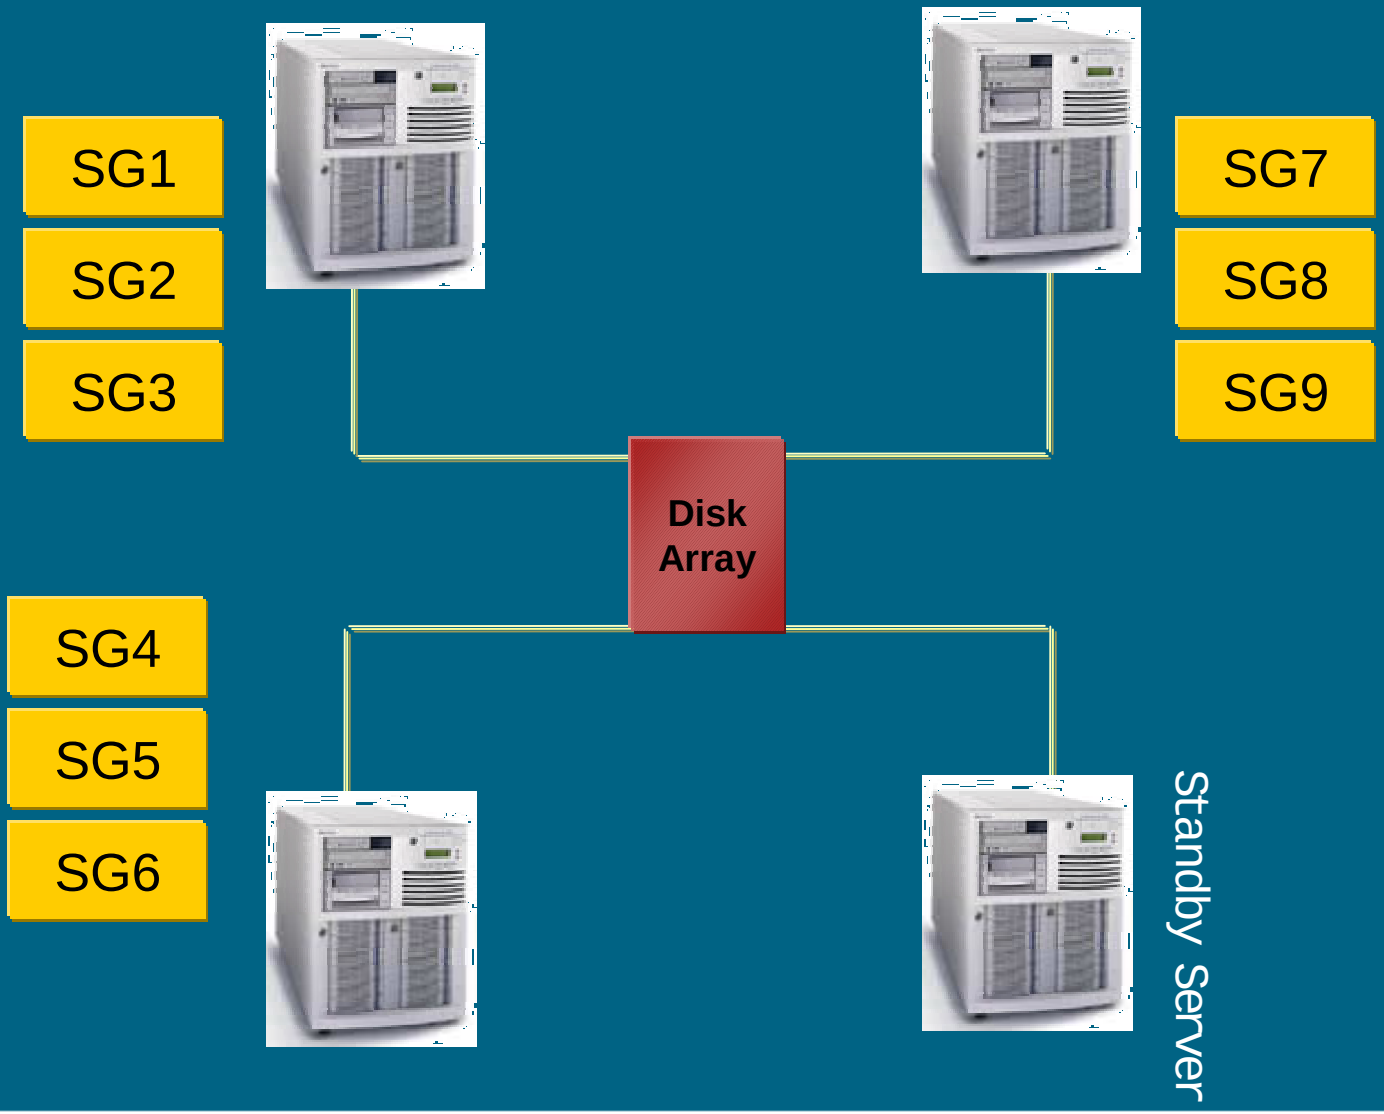

The Benefits Of Centralization in a corporate datacenter

Reduced Complexity

- Operating Procedures

Reduce Personnel Costs

- Help desk, administrators, engineers

Reduce Hardware/Software Costs

- Increased resource utilization
- Centralized/shared storage
- Reduced software licenses and distribution costs
- Volume purchasing

Increased Customer Service Levels and User Satisfaction

Large, highly distributed enterprises

- Densely populated regions/geographies
  - Corporate headquarters or campuses
  - State or local government
- Widely dispersed offices/locations
  - Retail or grocery chains
  - Banks, financial institutions, insurance agencies
  - State or local government

Candidates For Centralized Exchange Environments in a corporate datacenter

Exchange-2000
In the
Datacenter

Cracking the myth

IS NOT

Exchange-2000
On
Window-2000
Datacenter

Windows 2000 server family

deploy the right product for your needs

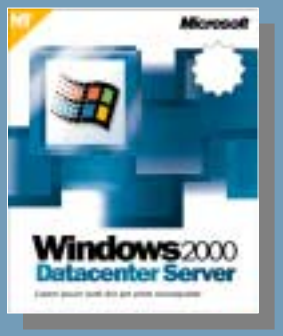

Up to 32-way SMP
Up to 32GB memory
4 node clustering
RAS HCL Testing
Process Control Manager
WinSock Direct

8-way SMP
8 GB RAM
2 node clustering
Load balancing

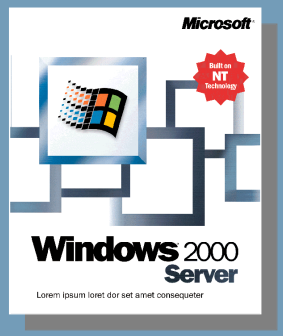

4-way SMP
4 GB RAM
IIS CPU & bandwidth throttling

Value

Functionality and Capacity

- 4 node clustering
  - Requires Exchange-2000 SP1
- 32 GB of RAM
  - Best use with 4GB or less
- 32-way CPU
  - Doesn't scale above 8 CPU
  - Scale out
- Process control
  - Tuning & Application Stacking
- Winsock Direct
  - Benefits in a front-end/Back-end architecture

How does Exchange-2000 take advantage of those features?

Any usage model for 32 way systems?

Partition in 4 8-way?

What is it?

- Big scalability, clean modularity
- Expensive, skills not common
- A lot of proprietary technology

Where best?

- Big non-distributed databases outpacing 8 way platforms

Why pass?

- Itanium and Windows 64 are the technologies for scale up due to higher speed and memory allocation efficiency
- ROI is shaky

… Exchange-2000 In the Datacenter

What is …

Always-On Infrastructure

Centralized Management

Agreed SLA

… Exchange-2000 In the Datacenter

What is …

Always-On Infrastructure

Centralized Management

Agreed SLA

Windows-2000 datacenter program

Best practices

HP Windows 2000 Datacenter Solution

Microsoft's Windows 2000 Datacenter Server O/S

HP Solution Manager

Integrated services and hardware

HP Datacenter Program: What does it bring to you?

- Stress Test
- HP MSO Consulting

- Readiness assessment
- Process design & management
- Education
- Remote monitoring

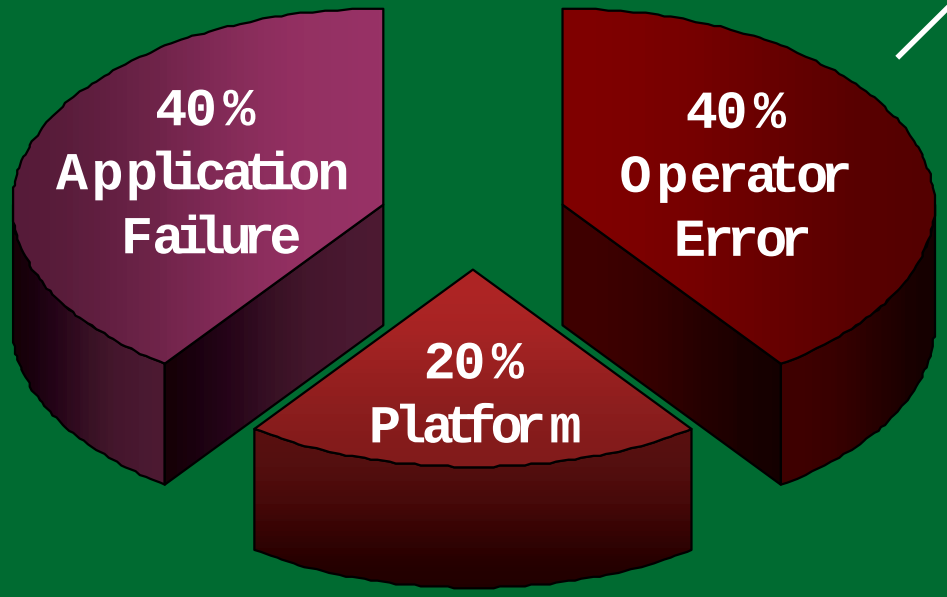

- Datacenter HCL
- Stress test
- HP reliability
- Certification of kernel intrusive applications

Best Practices for Always-On Exchange-2000

Data protection

Backup/Restore procedures

Clustering

Management and monitoring

Virus protection

Separate Exchange servers per functions:

- Mailboxes, Public folders, Bridgehead servers, front-end/back-end

Best Practices for Always-On Exchange-2000 (2)

Testing environment part of your Exchange deployment

- Fully part of your Exchange topology
- Used for evaluating new software or Service Packs

Drill Exercise

World Wide deployment topology

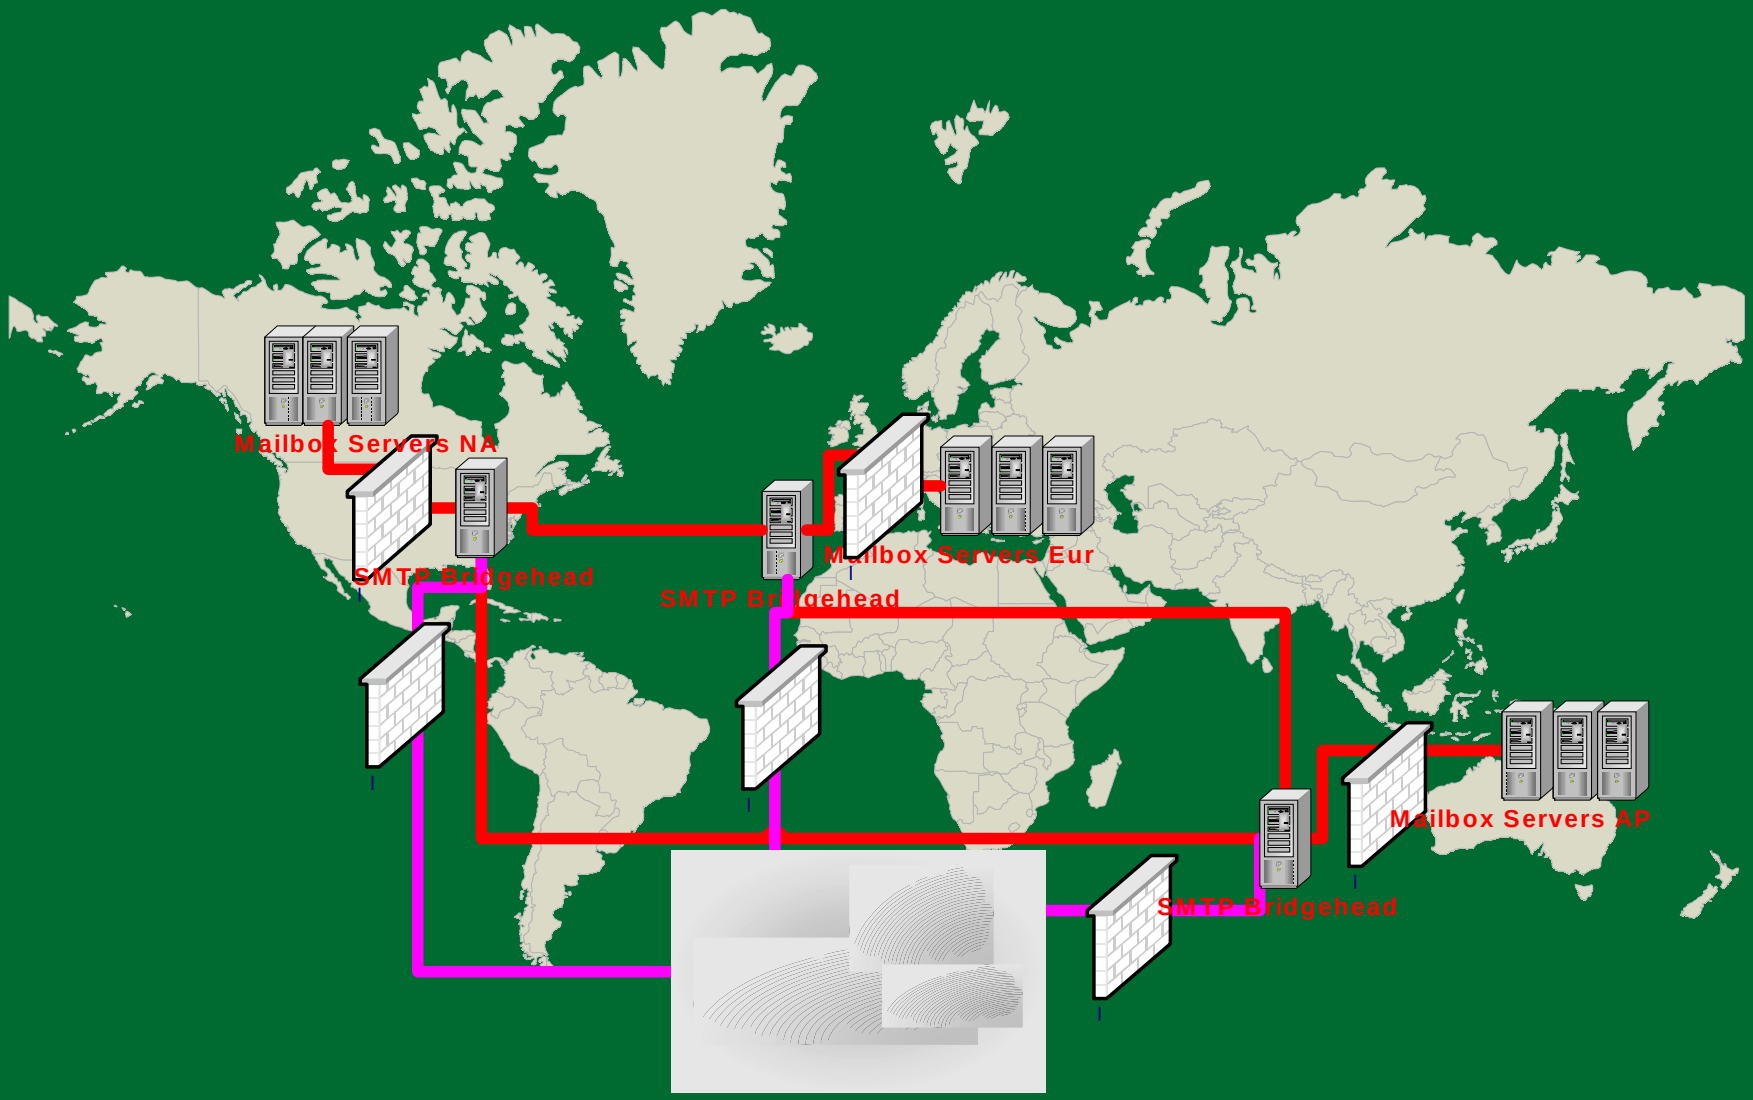

… Exchange-2000 In the Datacenter

What is …

Always-On Infrastructure

Centralized Management

Agreed SLA

Hosting Model

System Management and Monitoring

Storage Management

Backup and restore

Exchange in an hosted Environment

All software components, hardware live at the Enterprise datacenter

Centralized management and backup operation

Front-end/Back-end

- Up to 32 nodes with unified namespace across the farm
- HTTP-WebDAV, POP3/IMAP4/SMTP
- Dynamic adjusts, distribute requests, without re- authentication
- Offload SSL processing
- Additional security layer
- Seamless user distribution across multiple servers
  - Allows for seamless server consolidation

Easy to use

- provides tight Windows integration
- offers customizable policies
- product's exceptional report generator can publish reports directly to IIS

- HP OpenView ManageX

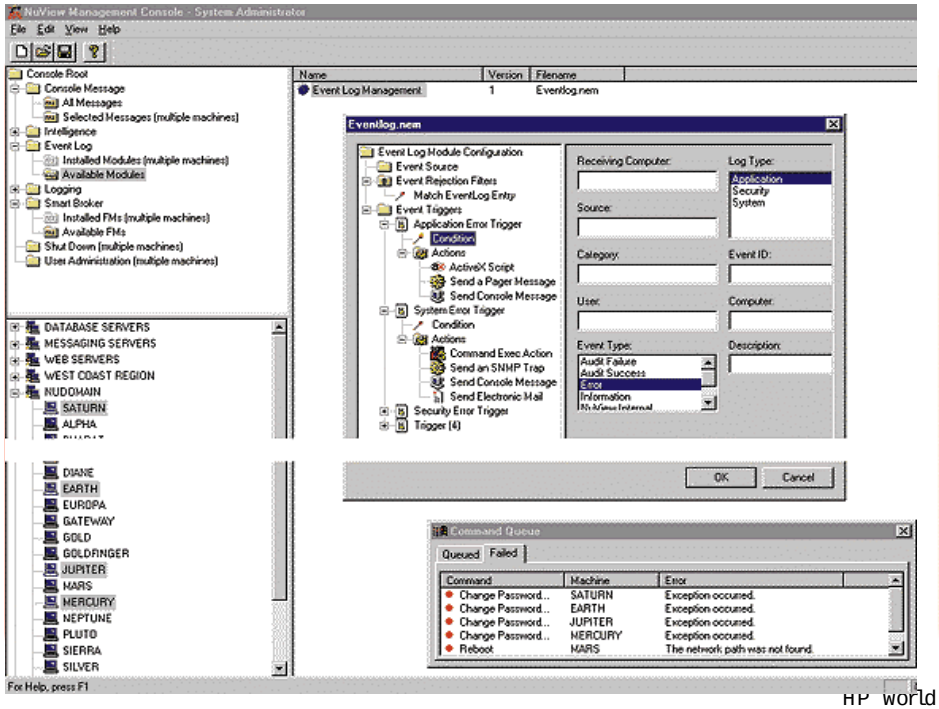

System management and monitoring solutions

Storage Management

LOG, STM and EDB files
- Different I/O characteristics
- LOG – Sequential, small synchronous I/O
- EDB - 4K random I/O – Sync Read – Async Write
- STM – 32KB I/O or larger

Controller level caching

Up to 512 disks (24TB)

32 GB cache

928 host connections

Characterization work with 7030 Mailboxes (4 servers)

Snapshot/Split Mirror Storage

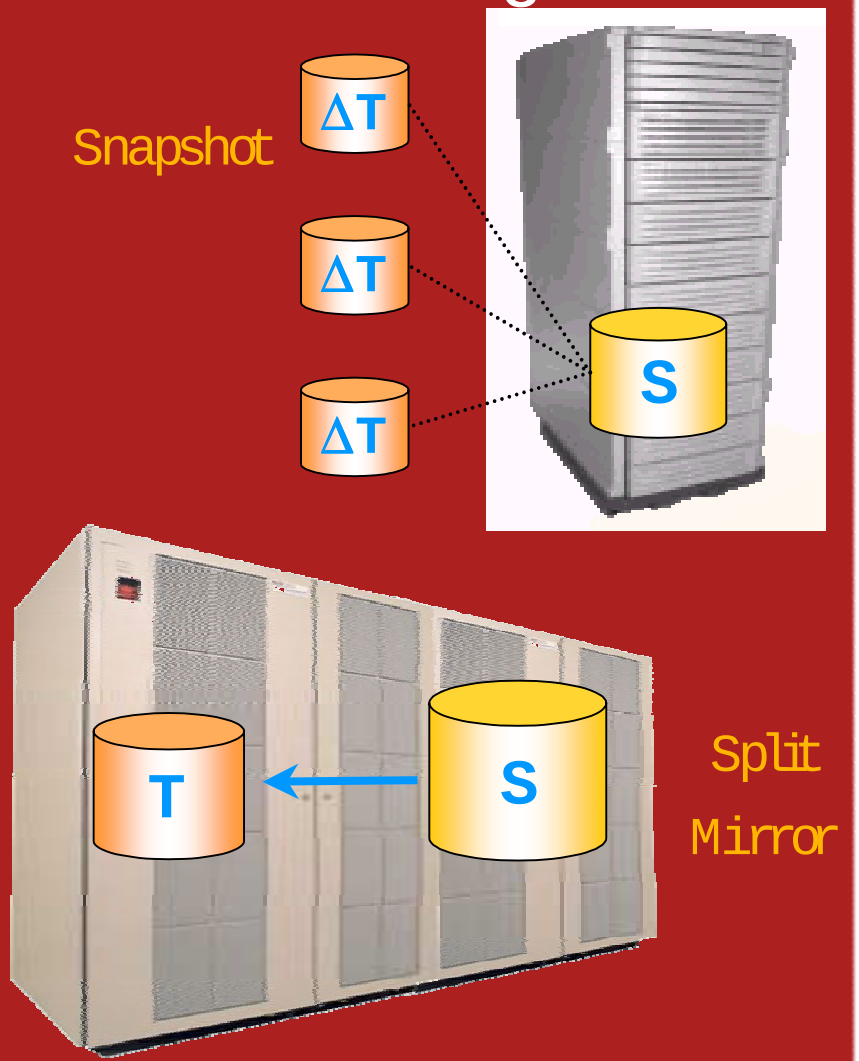

- Methods for establishing point-in-time images of data
- Snapshot maintains point-in-time copies using delta images
- Split Mirror maintains separate complete disk images for point-in-time copies
- Either method can be managed in software or in array firmware

- Exchange services must be stopped when taking the Snapshot
- Offline backup procedure (Check related KB articles)
- Log files
- Page level integrity check

… Exchange-2000 In the Datacenter

What is …

Always-On Infrastructure

Centralized Management

Agreed SLA

What level of services

How many users per servers

What is in your SLA?

All Exchange architecture are driven by the Service Level Agreement

How many physical locations?

Downtime per year

- 24/7 or 8/5 or in between
- Balance between
  - Investment Cost
  - Downtime cost
  - Risk

What risk is protected?

Backup duration & frequency

Maximum time to recover

Single message or mailbox recover service

Lost messages

Response time

Mailbox quotas

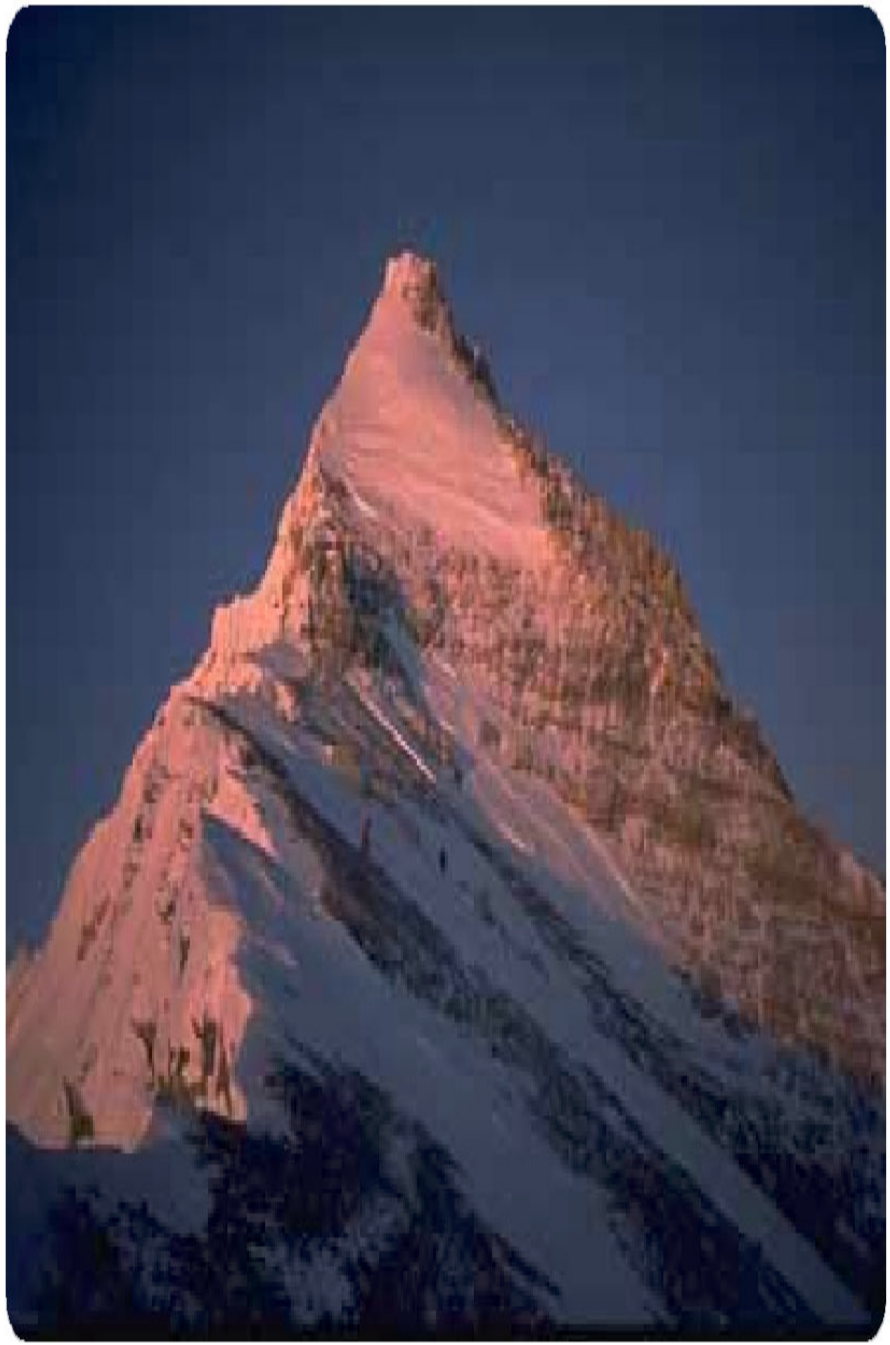

5 nines = 5 minutes

99.99% = 1 hour

99.9% = 9 hours

99% = 3.5 days

Exchange 2000

Scheduled downtime

- Server Maintenance
- Service Pack Installation
- Firmware Upgrade

Up and Running in under 3 minutes

Unscheduled downtime

- Hardware Problem
- Software Problem

Up and Running in

Failover

Backup Types

Type	Copies DB	Copy Logs	Truncates Logs
Full(Normal)	X	X	X
Copy/Daily	X	X	
Incremental		X	X
Differential		X	

Backup performance

Speed to backup multiple storage group simultaneously

	DLT 8000		Ultrium	
Number of Storage Groups Simultaneously Backed Up	Total MB/s	Total GB/hr	Total MB/s	Total GB/hr
1	**10**	**36**	**34**	**124**
2	**20**	**72**	**66**	**240**
3	**30**	**108**	**87**	**315**
4	**40**	**144**	**106**	**380**

Restore performance performance

under optimal conditions

	DLT 8000		Ultrium	
Number of Storage Groups Simultaneously restored	Total MB/s	Total GB/hr	Total MB/s	Total GB/hr
1	**9**	**34**	**41**	**148**
2	**18**	**68**	**68**	**246**
3	**27**	**102**	**86**	**308**
4	**36**	**136**	**84**	**302**

- Maximum of 4 Storage Groups per Server
- Maximum of 5 MDB Stores per Storage Group
- Each Storage Group contains log files for all of its Mailbox Stores
- Backup Unit is the Storage Group
- How long to restore a MDB?

Full server

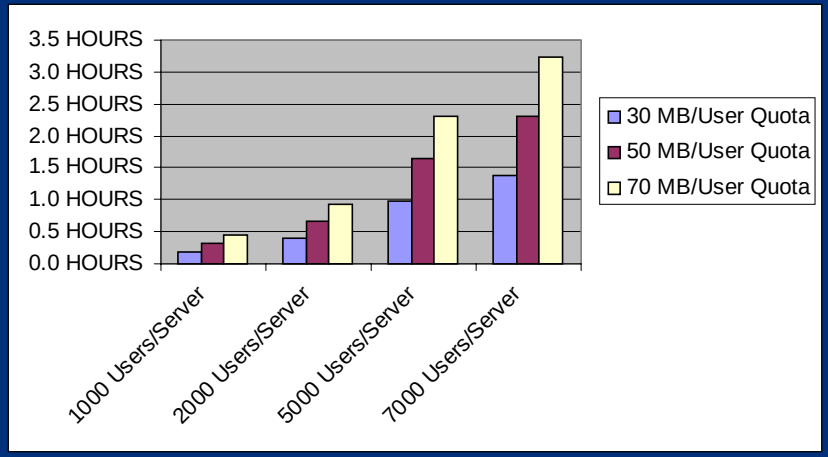

Single MDB

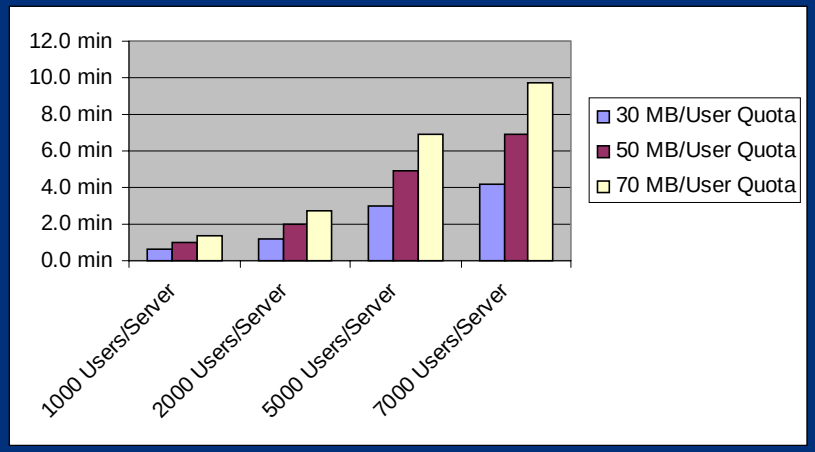

What about ?

- Problem notification and identification
- Analysis and decision
- Loading tapes
- Replaying transaction logs

Single message or mailbox restore

- Check ISV partner software
  - CommVault
    - Backup and Restore Software
    - Incremental backup at the message level
  - KVS
    - Archiving software
      - May archive everything to a secondary media such as optical storage

How many users per servers?

- Driven by the SLA
- Depends on the users profile
  - When do they start working?
  - Same time zone?

- Eg 99.99 availability
- No performance degradation in failover
- Recover a corrupted DB in server in less than 1 hour
  - A good limit will be 2000 Users per Storage Group (with 4 MDB)

How many users?
4 node datacenter cluster-18000 users
(requires Exchange-2000 SP1)

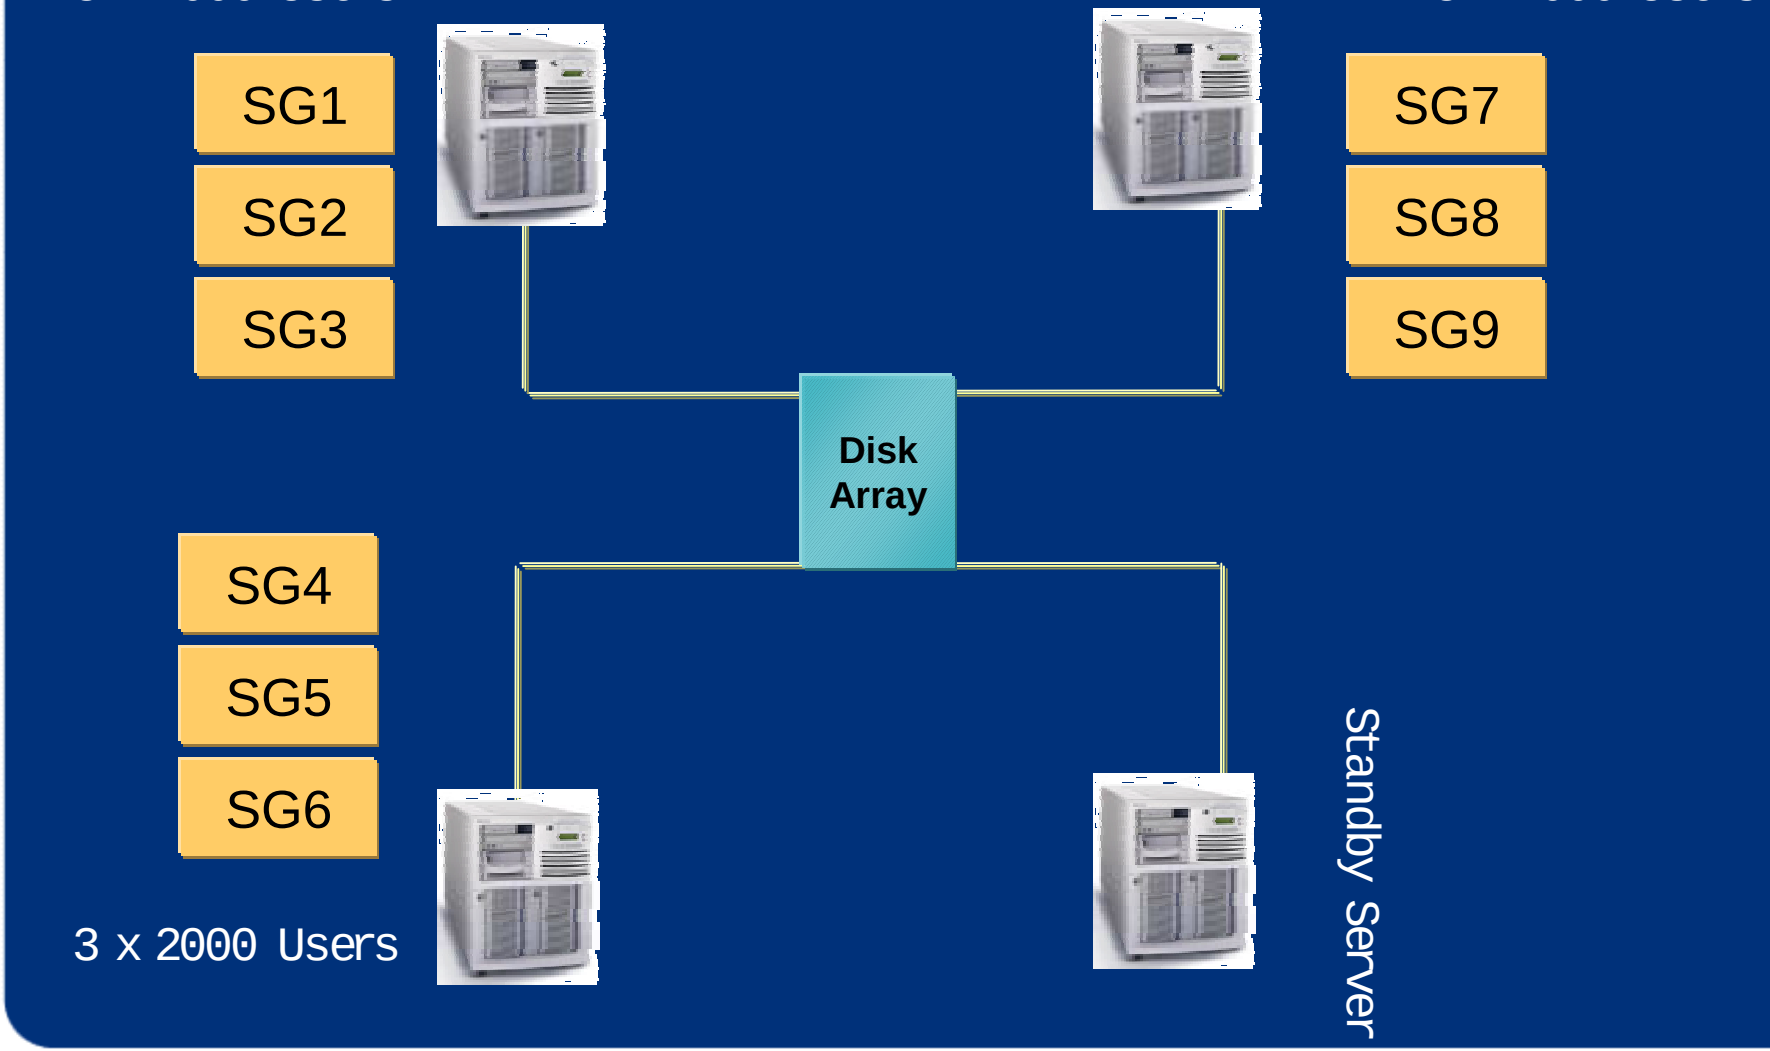

… Exchange-2000 SP1 & 2…

What is next…

… Exchange Version next …

…64 bits…

Assess | Plan & Design | Implement | Maintain

How can hp help you?

Exchange 2000 services from hp

- discovery workshop
- planning and design
- implementation & migration
- support
- licensing
- training
- outsourcing


hp
invent