Your Independent HP Business Technology Community

WI-01-NH

SQL Server on HP Integrity

Failure is not an option

Thomas Grohser
Senior Database Engineer, bwin

SQL Server on HP Integrity Agenda

- The Mission
- The Solution
  - Standardizing
  - Zero data loss
  - High availability
  - Scale up

SQL Server on HP Integrity
The Mission

- VLDB – A database that needs attention it's not size alone
- SLA
  - Zero data loss & 100% transactional consistency on financial transactions
  - 99.99x% availability @ 24 x 7
  - 450.000+ SQL Statements per second
  - Assumed worst case scenario: full datacenter failure with complete data loss within the datacenter
- Budget: unlimited (not kidding)

SQL Server on HP Integrity
The Solution

- Standardize everything
- Work by the book
- Have some clever guys at hand

if the book runs out of pages

SQL Server on HP Integrity Standardizing

- Operating System
  - Version, Edition, Service Pack, Patch Level
- File System and Disks

SQL Server on HP Integrity File System Sample

- C:\
- C:\Windows
- C:\Install
- C:\SQL01
- C:\SQL01\BIN
- C:\SQL01\TEMPLOG01
- C:\SQL01\TEMPDATA01
- C:\SQL01\LOG01
- C:\SQL01\LOG02
- C:\SQL01\DATA01
- C:\SQL01\DATA02
- C:\SQL01\DATA03
- C:\SQL01\DATA04

SQL Server on HP Integrity File System expansion

- C:\
- C:\Windows
- C:\Install
- C:\SQL01
- C:\SQL01\BIN
- C:\SQL01\TEMPLOG01
- C:\SQL01\TEMPDATA01
- C:\SQL01\LOG01
- C:\SQL01\LOG02
- C:\SQL01\DATA01
- C:\SQL01\DATA02
- C:\SQL01\DATA03
- C:\SQL01\DATA04

SQL Server on HP Integrity File System settings

- Stripesize 64/128/256 kB
  depending on storage
- Partition alignment 64/128 kB
  depending on storage
- Cluster size 64 kB
- 100% write cache
- 0% read cache

SQL Server on HP Integrity Standardizing

- Operating System
  - Version, Edition, Service Pack, Patch Level
- File System and Disks
- SQL Server
  - Version, Edition, Service Pack, Patch Level
- Network
  - Separate network for data and backup
  - IP Schema
- Documentation, Documentation, ….

SQL Server on HP Integrity
Zero data loss

- Redundant NIC
- Redundant Power Supply
- Data files on SAN

  (RAID 1/0 Multipath/2 Fabrics)
- Transaction log files on RAID 101

SQL Server on HP Integrity RAID 101

RAID Controller

RAID Controller

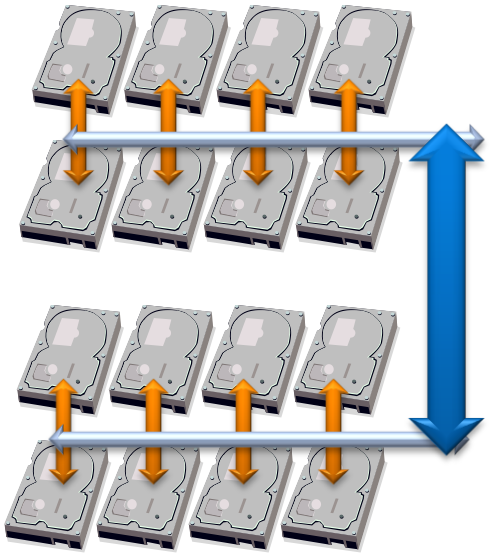

HW RAID 1

HW RAID 0

SW RAID 1

SQL Server on HP Integrity
Zero data loss

- Redundant NIC
- Redundant Power Supply
- Data files on SAN
  - (RAID 1/0 Multipath/2 Fabrics)
- Transaction log files on RAID 101

Availability: 0,00%

Data loss: 100,00%

SQL Server on HP Integrity
Zero data loss

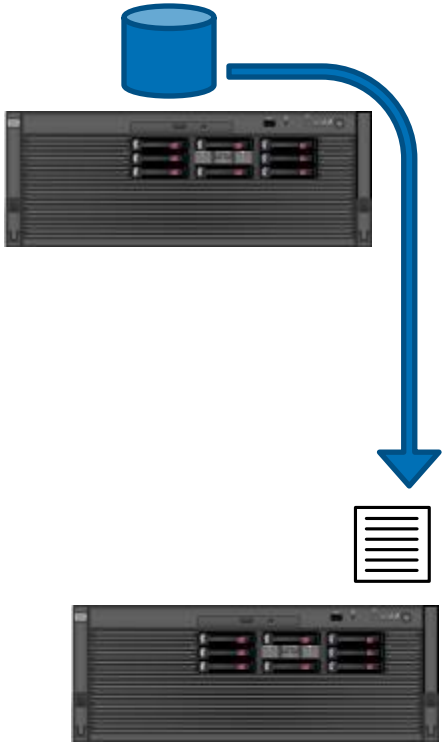

- Full backup every 24 h

Availability: 0,00%

Data loss: 100,00%

SQL Server on HP Integrity
Zero data loss

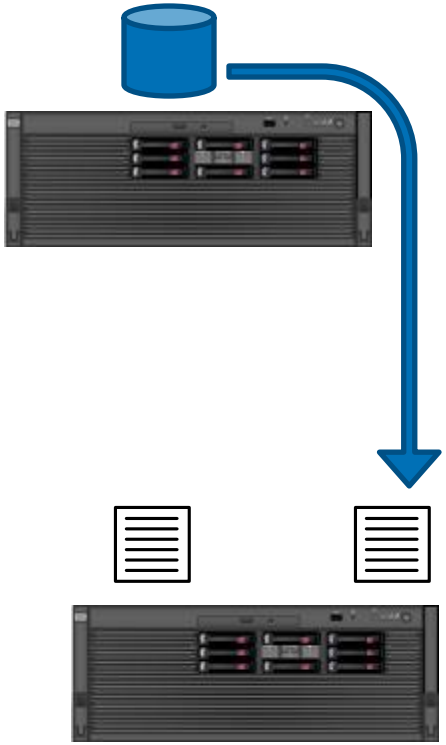

- Full backup every 24 h
- Alternating files

Availability: 0,00%

Data loss: 100,00%

SQL Server on HP Integrity
Zero data loss

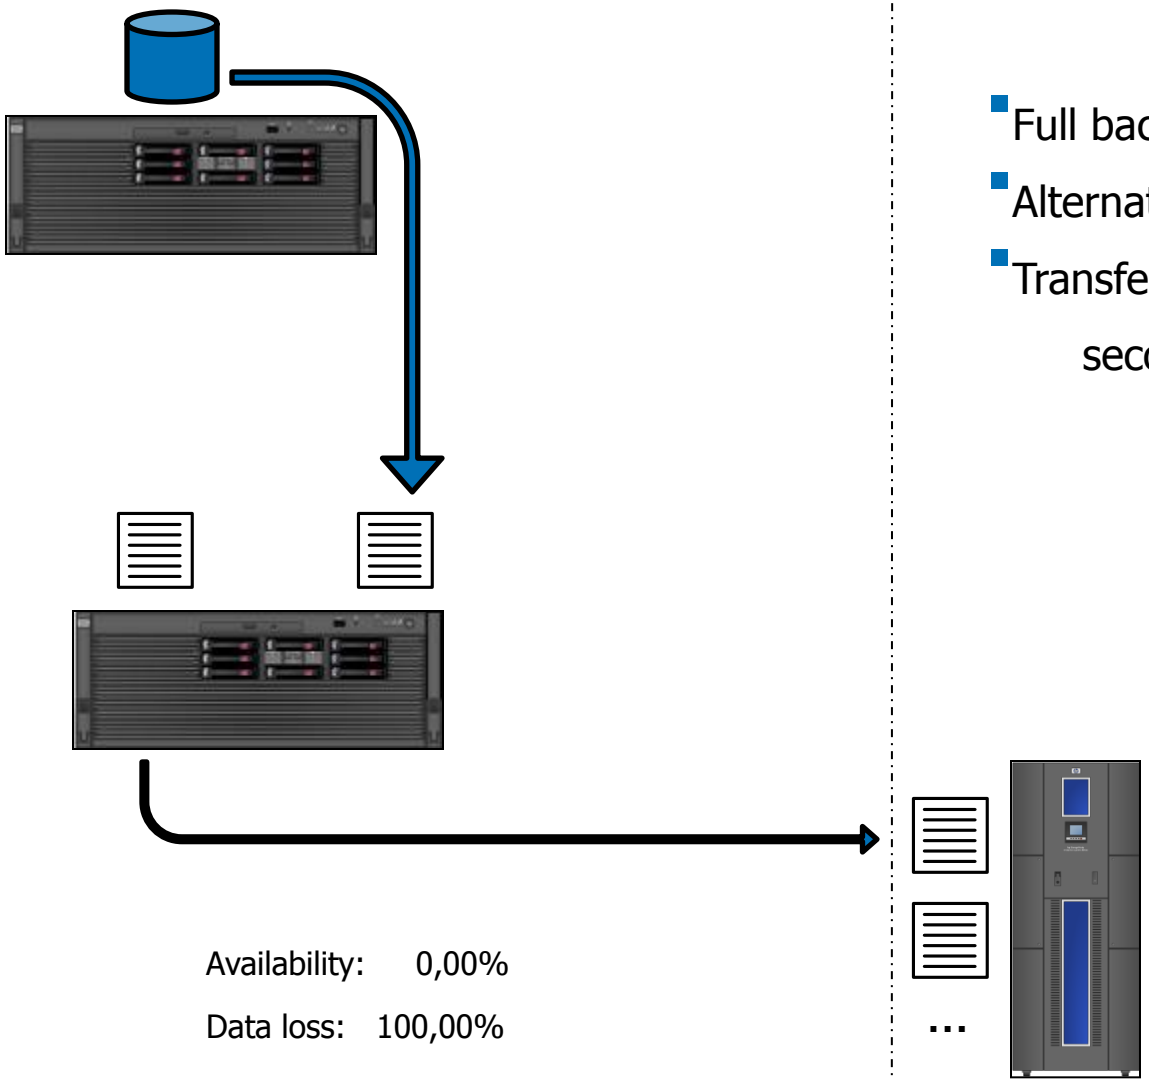

- Full backup every 24 h
- Alternating files
- Transfer of files to tape in second location

Availability: 0,00%

Data loss: 100,00%

SQL Server on HP Integrity
Zero data loss

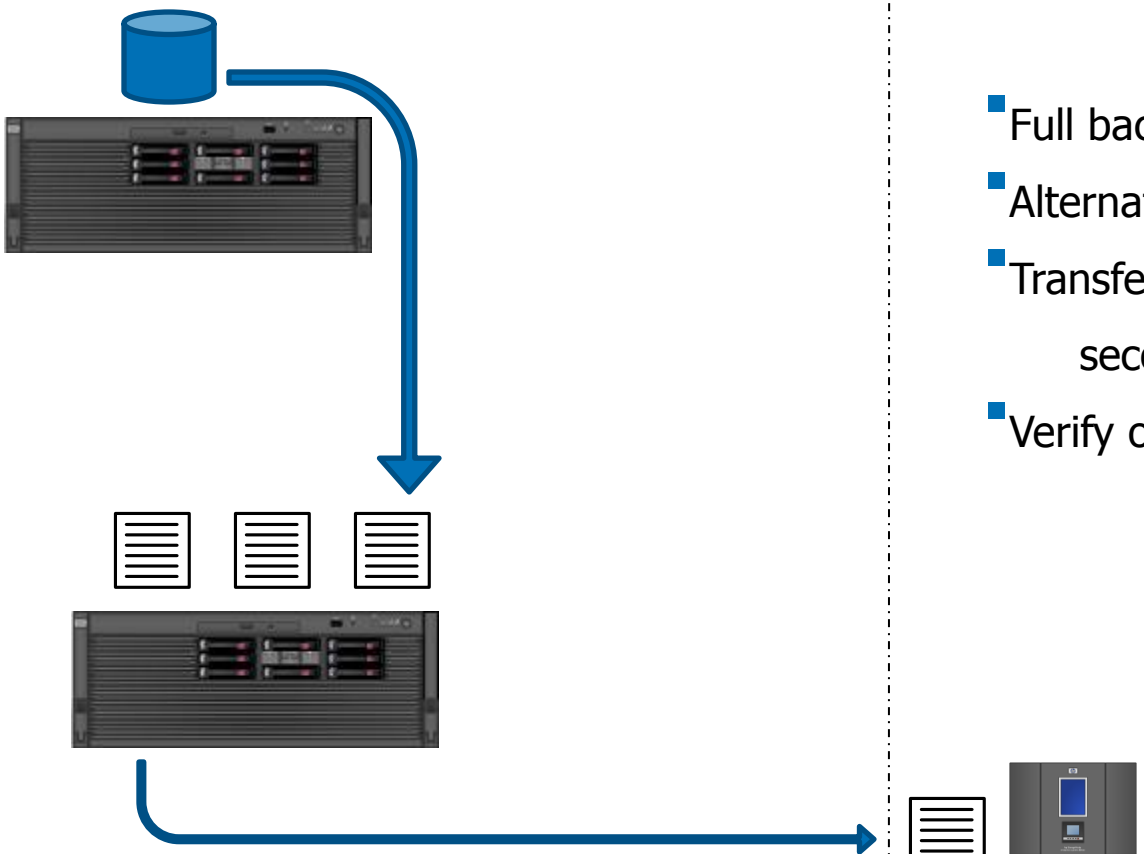

- Full backup every 24 h
- Alternating files
- Transfer of files to tape in second location
- Verify of backups (daily)

Availability: 98,00%

Data loss: 100,00%

SQL Server on HP Integrity
Zero data loss

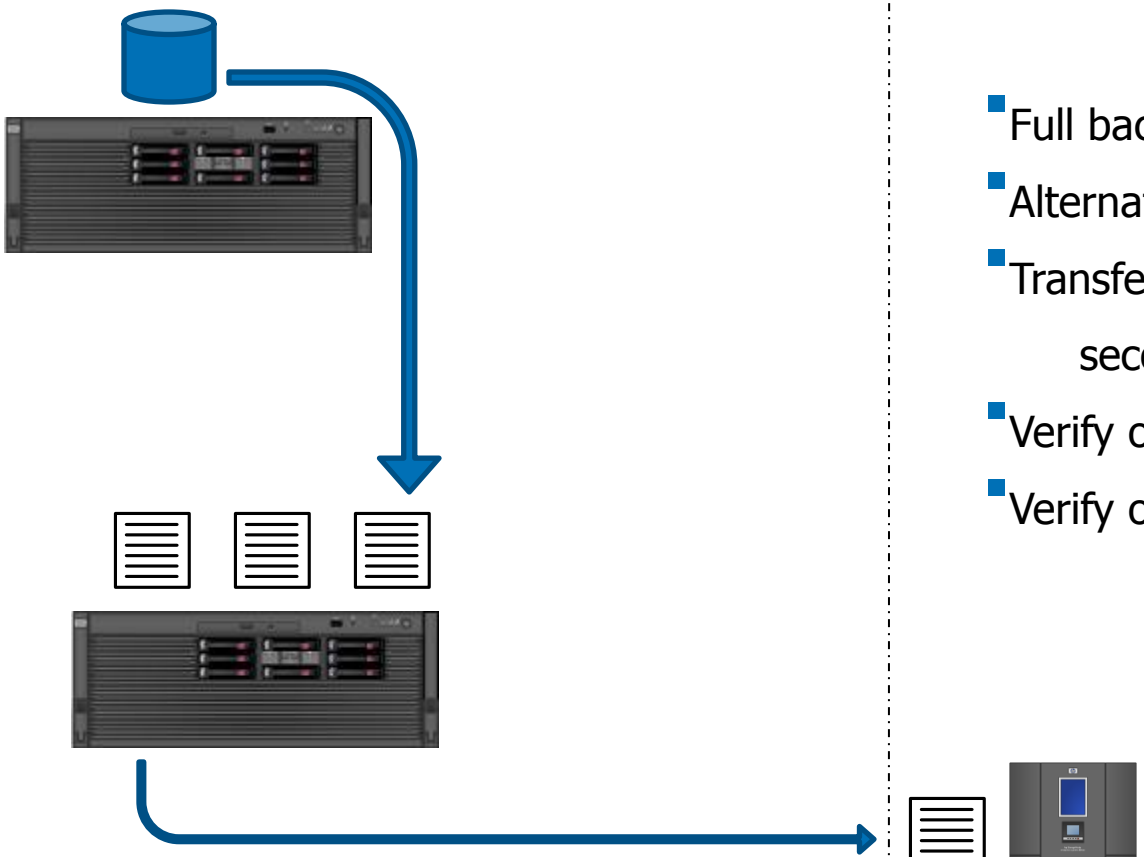

- Full backup every 24 h
- Alternating files
- Transfer of files to tape in second location
- Verify of backups (daily)
- Verify of tapes

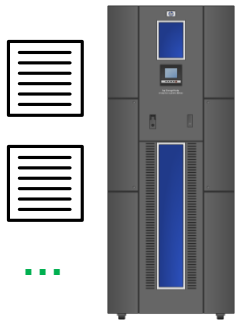

Availability: 98,00%

Data loss: 48 hours

SQL Server on HP Integrity
Zero data loss

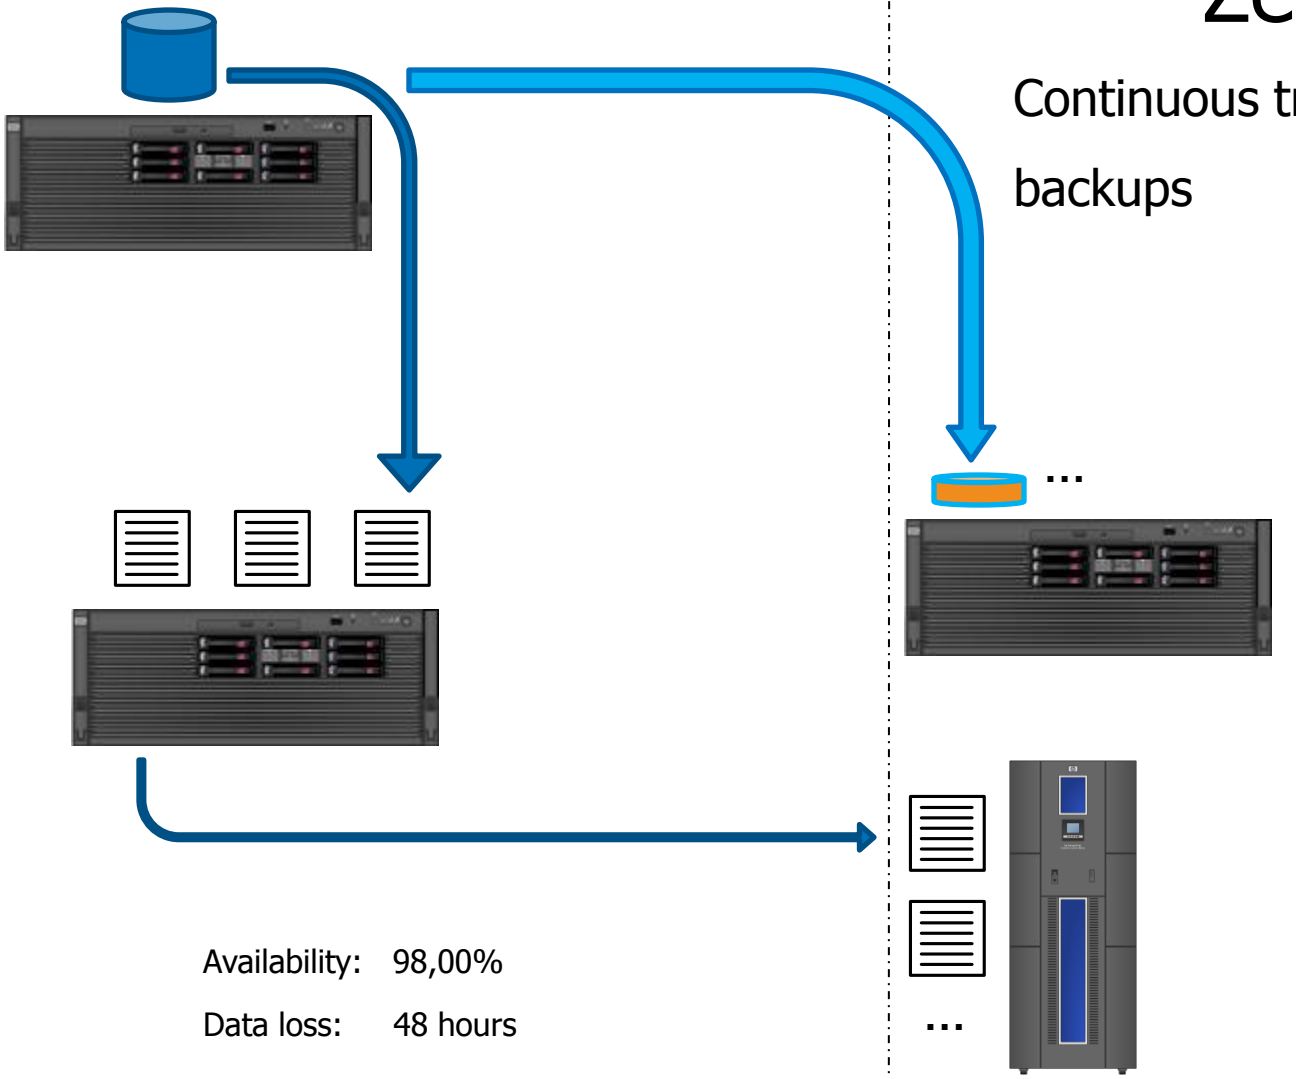

Continuous transaction log backups

Availability: 98,00%

Data loss: 48 hours

SQL Server on HP Integrity
Zero data loss

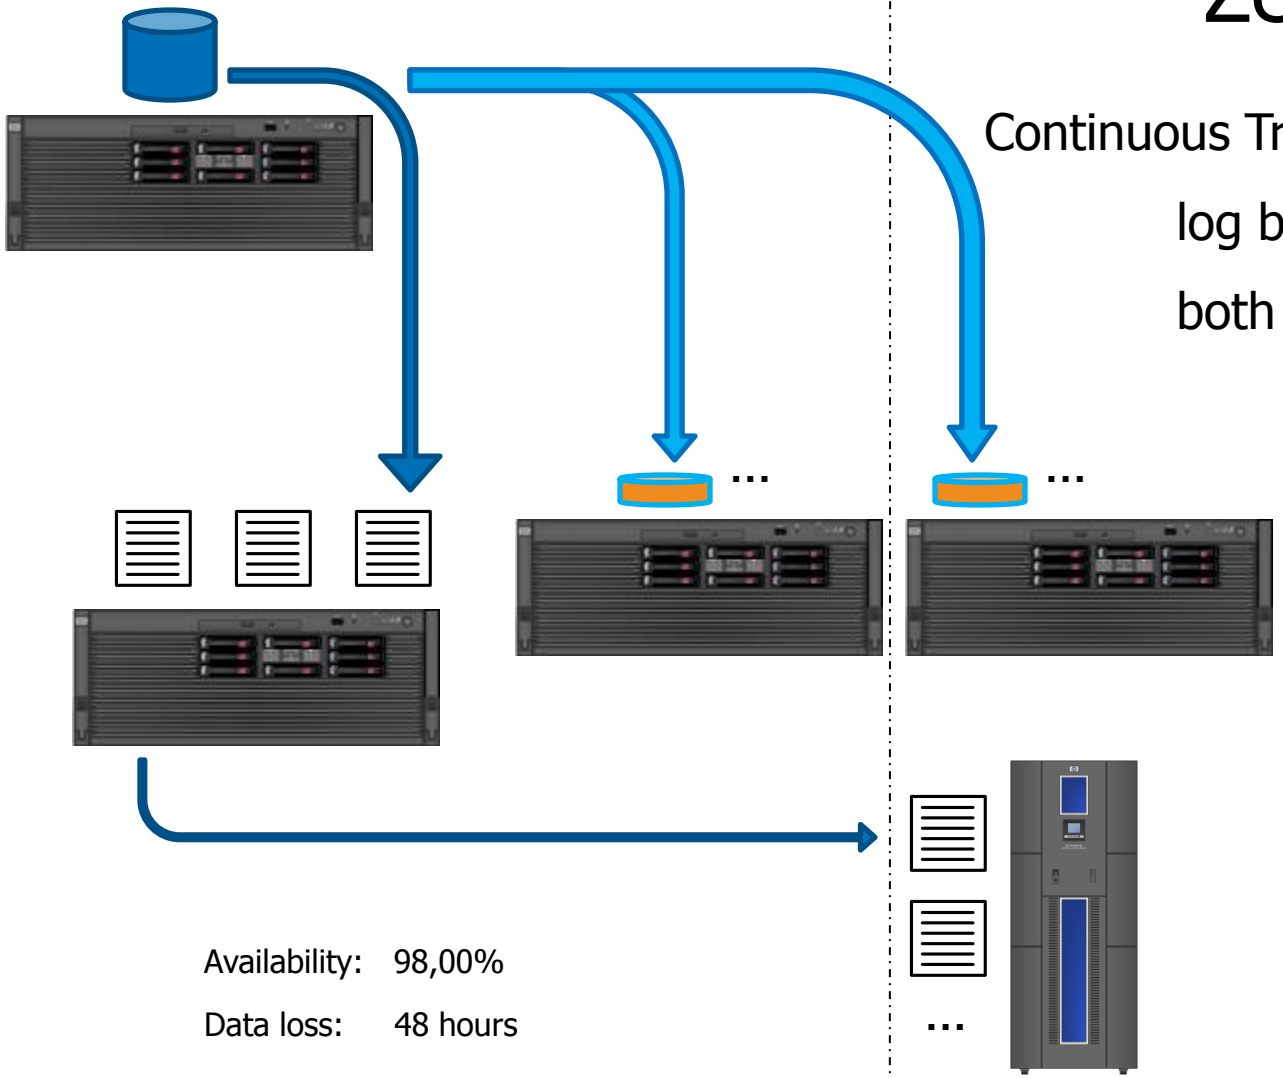

Continuous Transaction log backups to both locations

Availability: 98,00%

Data loss: 48 hours

SQL Server on HP Integrity
Zero data loss

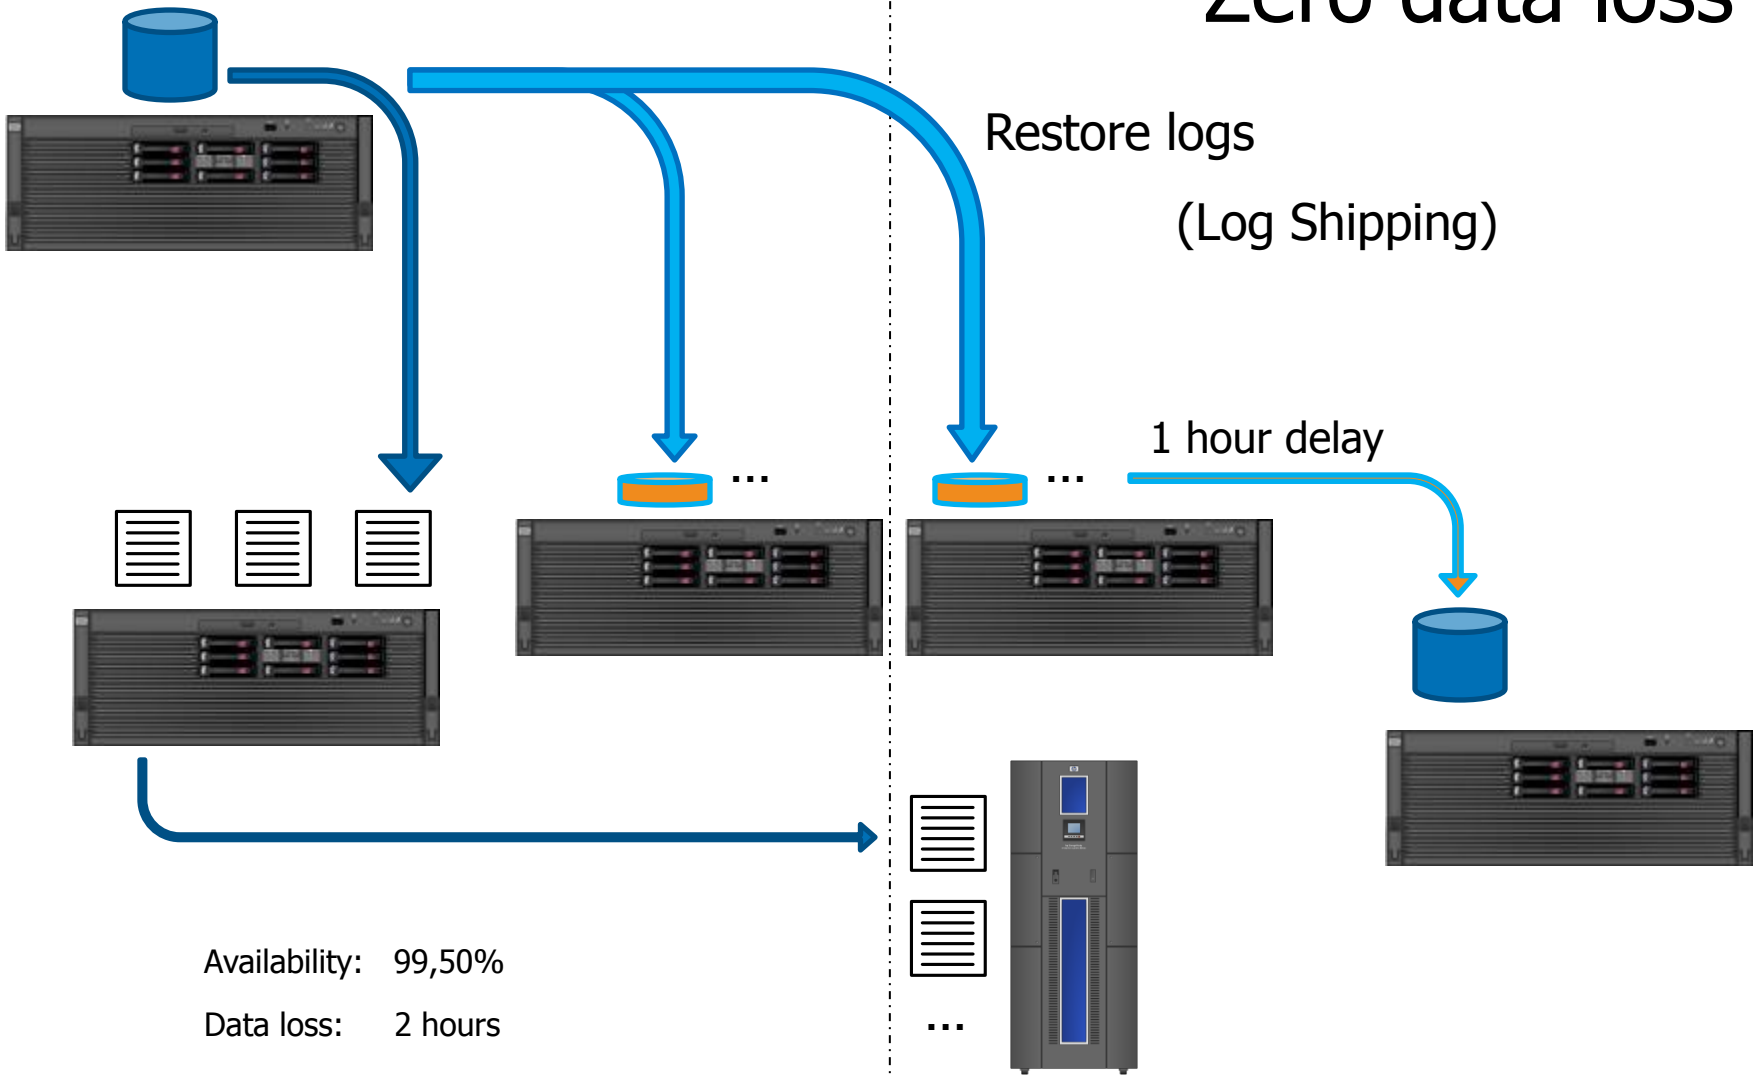

Availability: 99,50%

Data loss: 2 hours

SQL Server on HP Integrity
Zero data loss

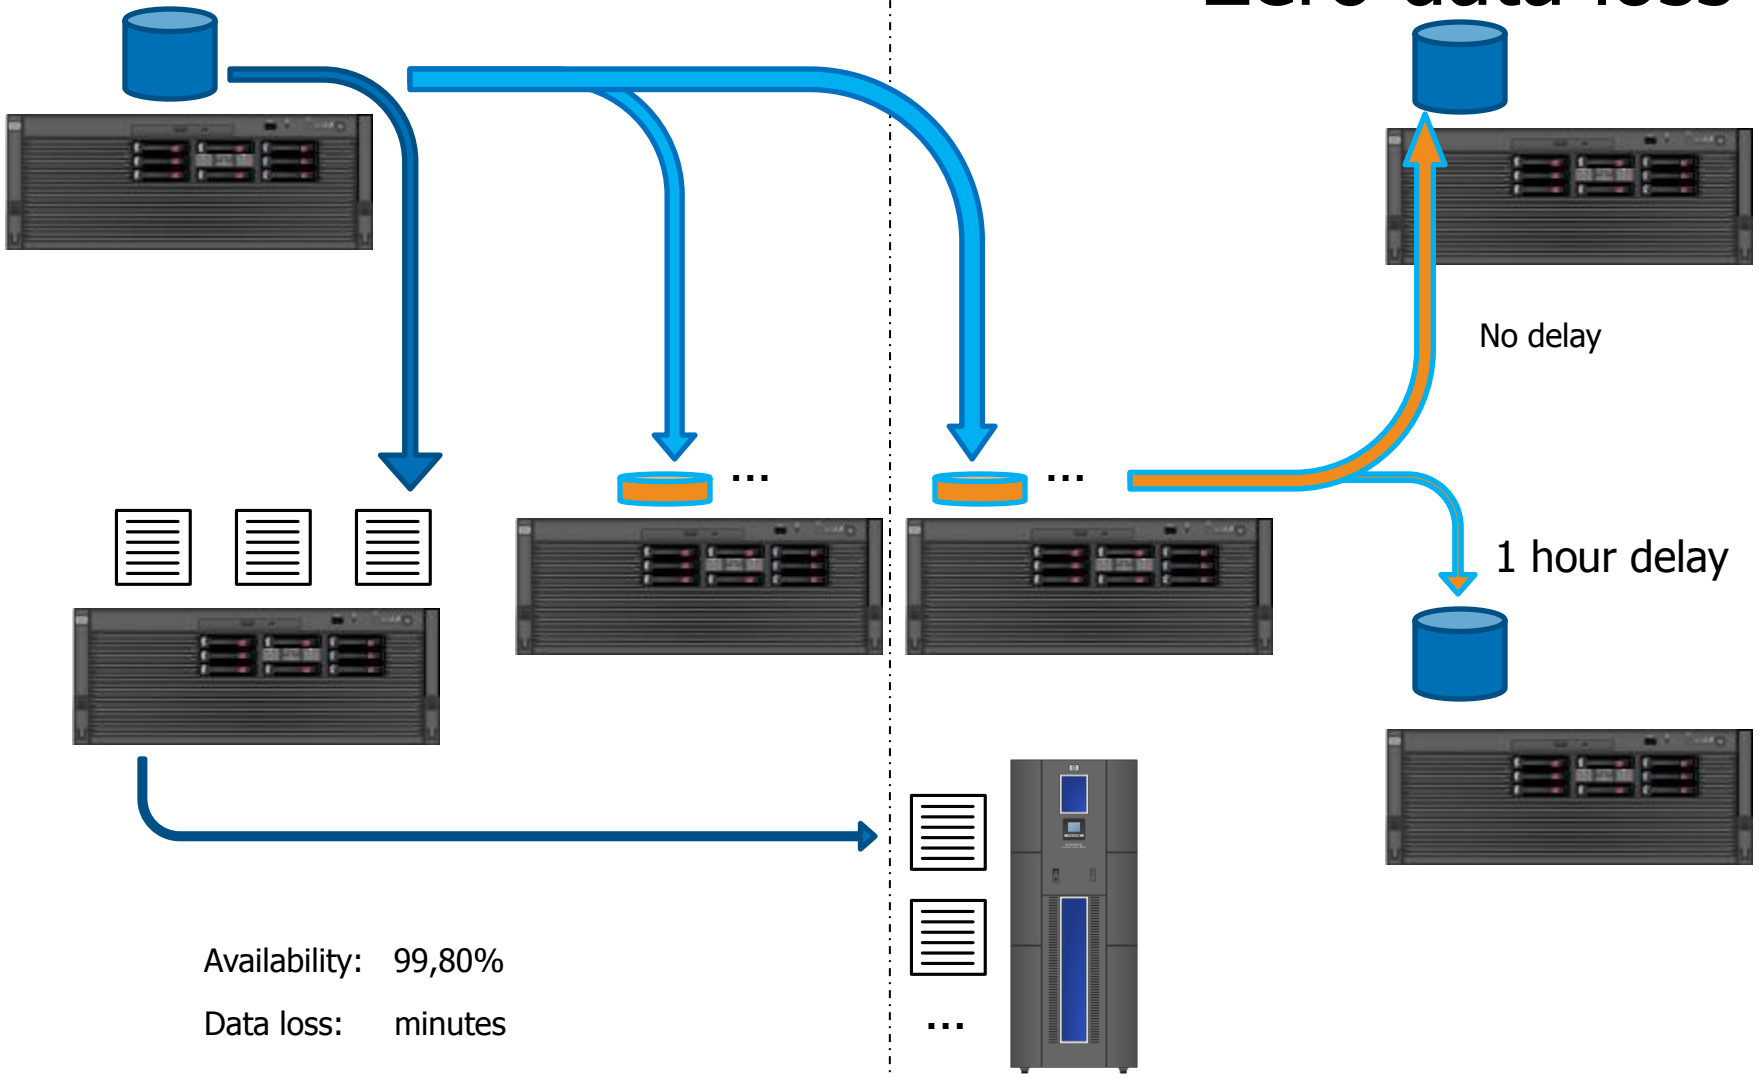

Availability: 99,80%

Data loss: minutes

SQL Server on HP Integrity
Zero data loss

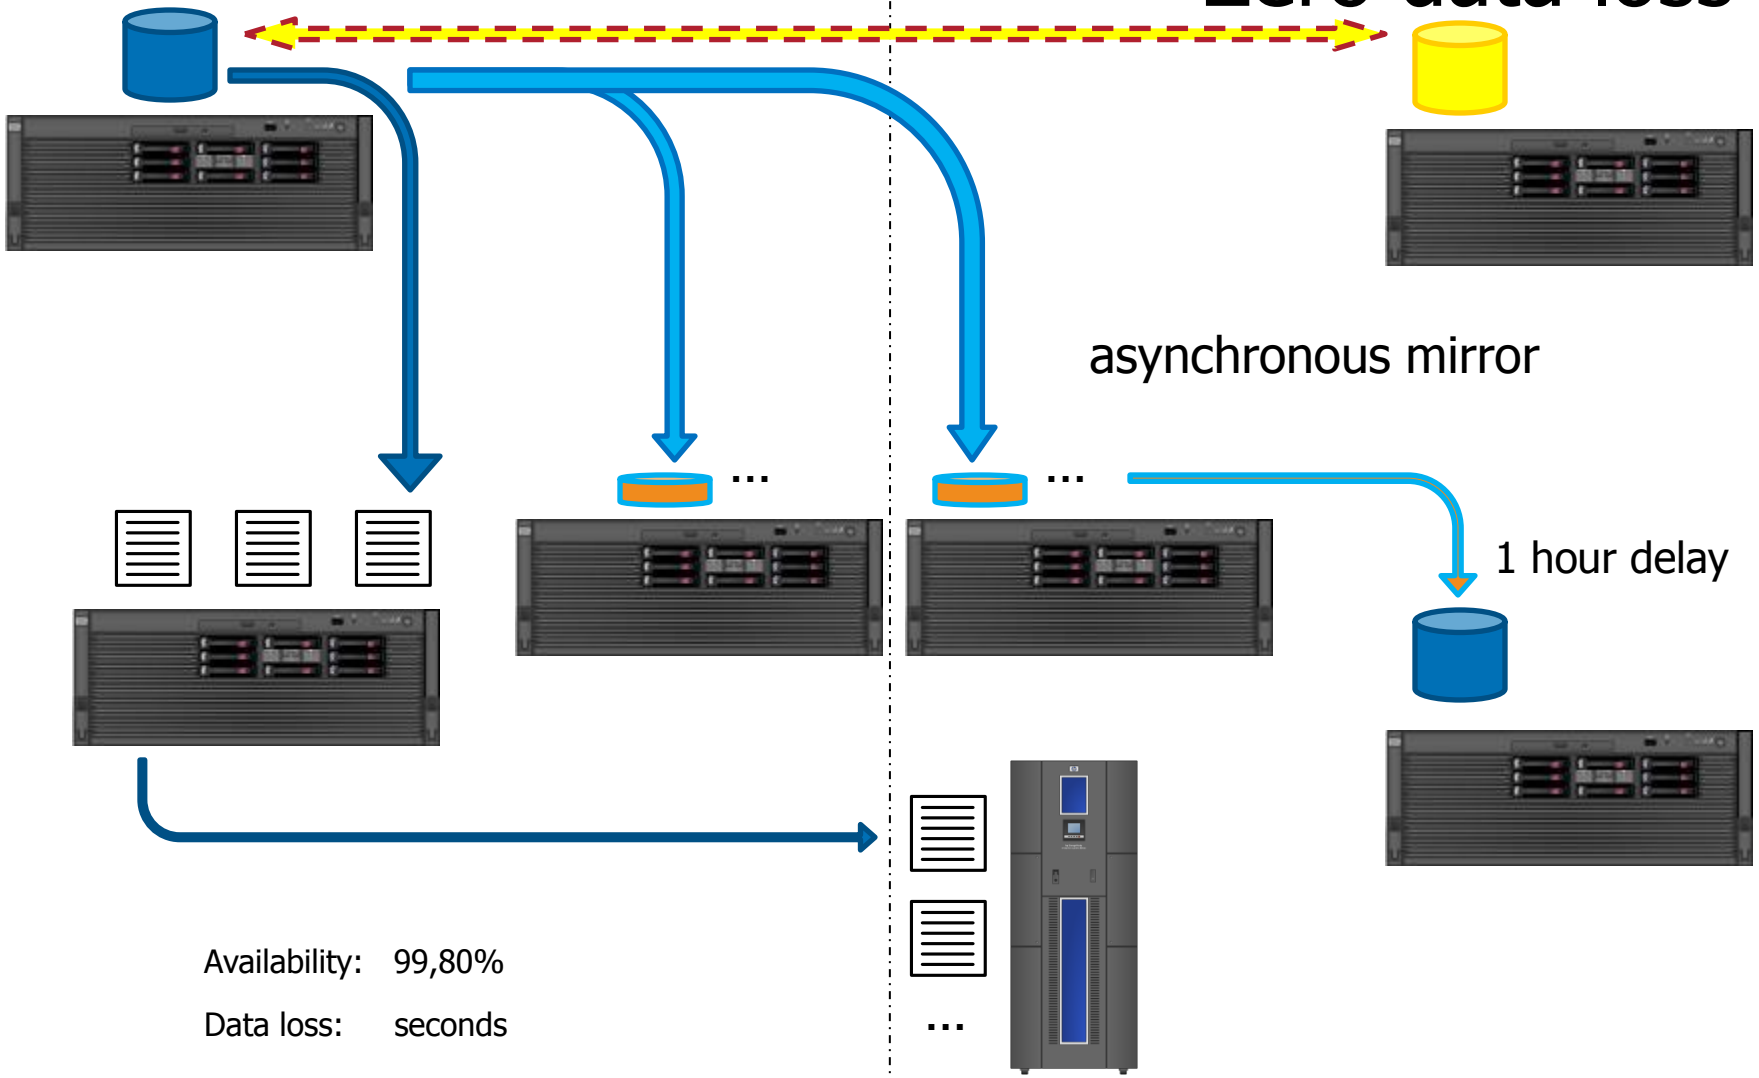

SQL Server on HP Integrity
Zero data loss

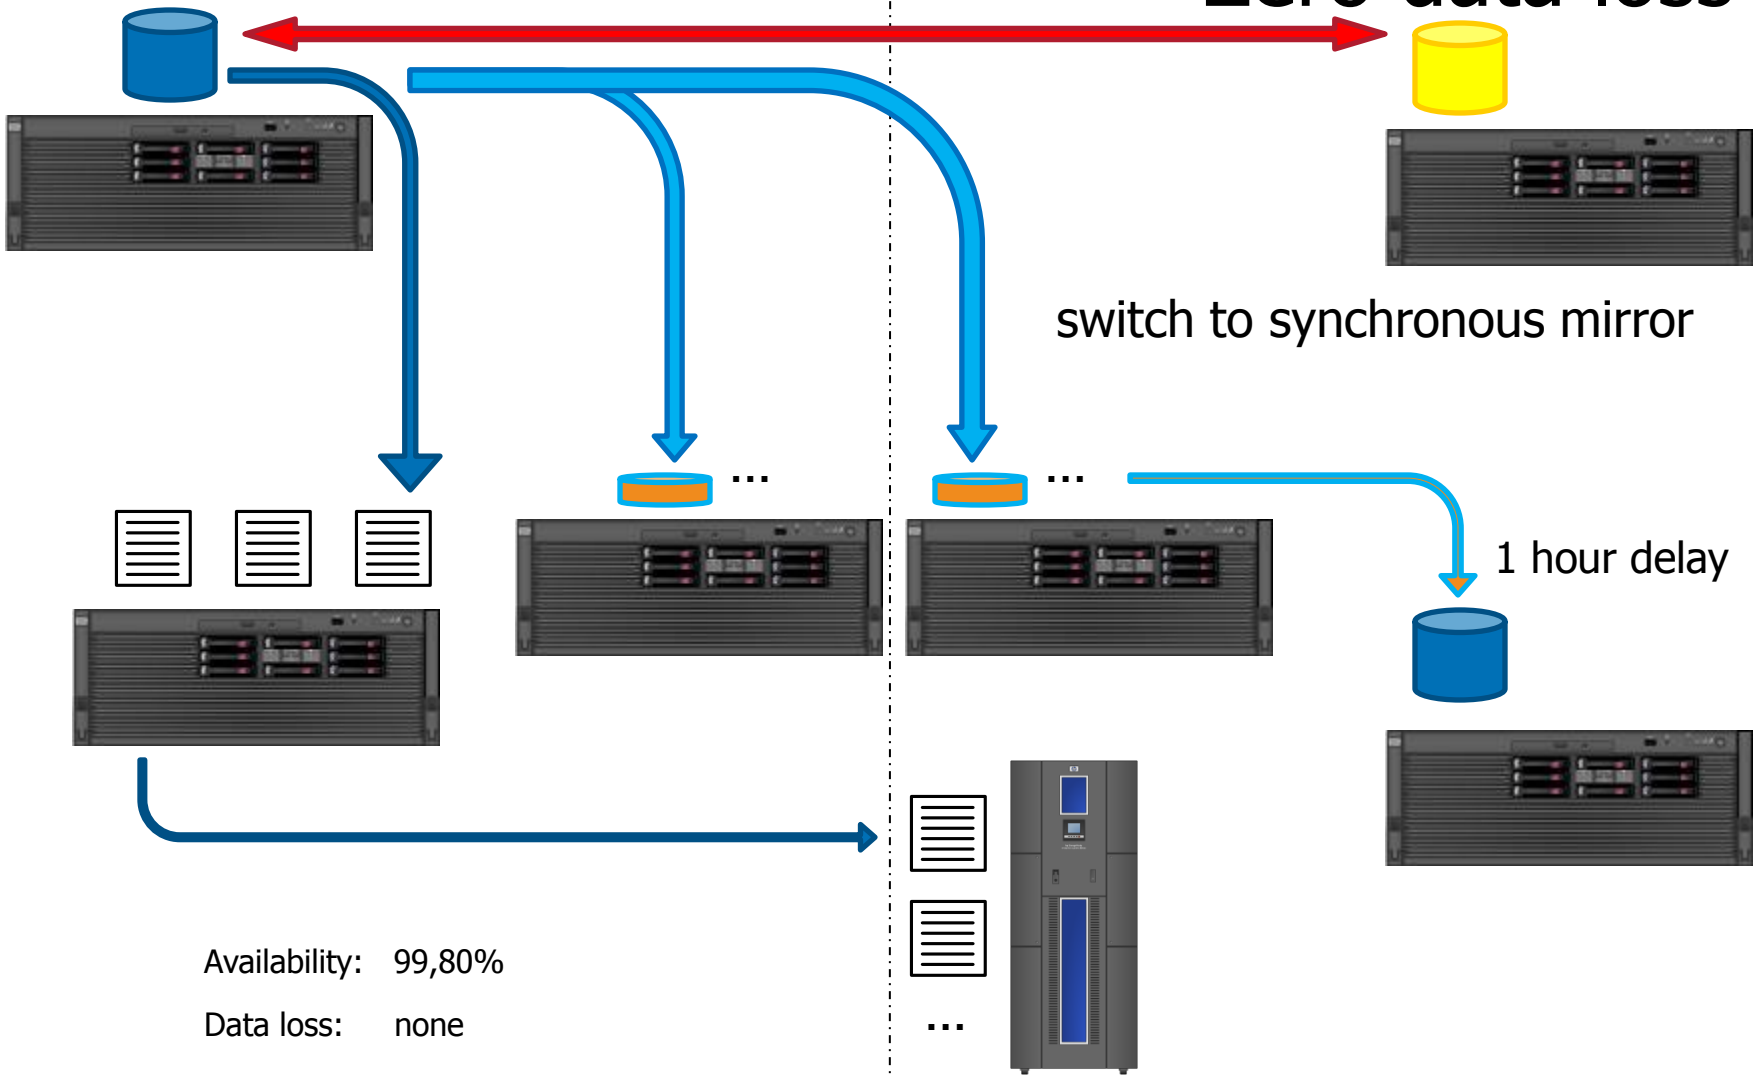

SQL Server on HP Integrity
Zero data loss

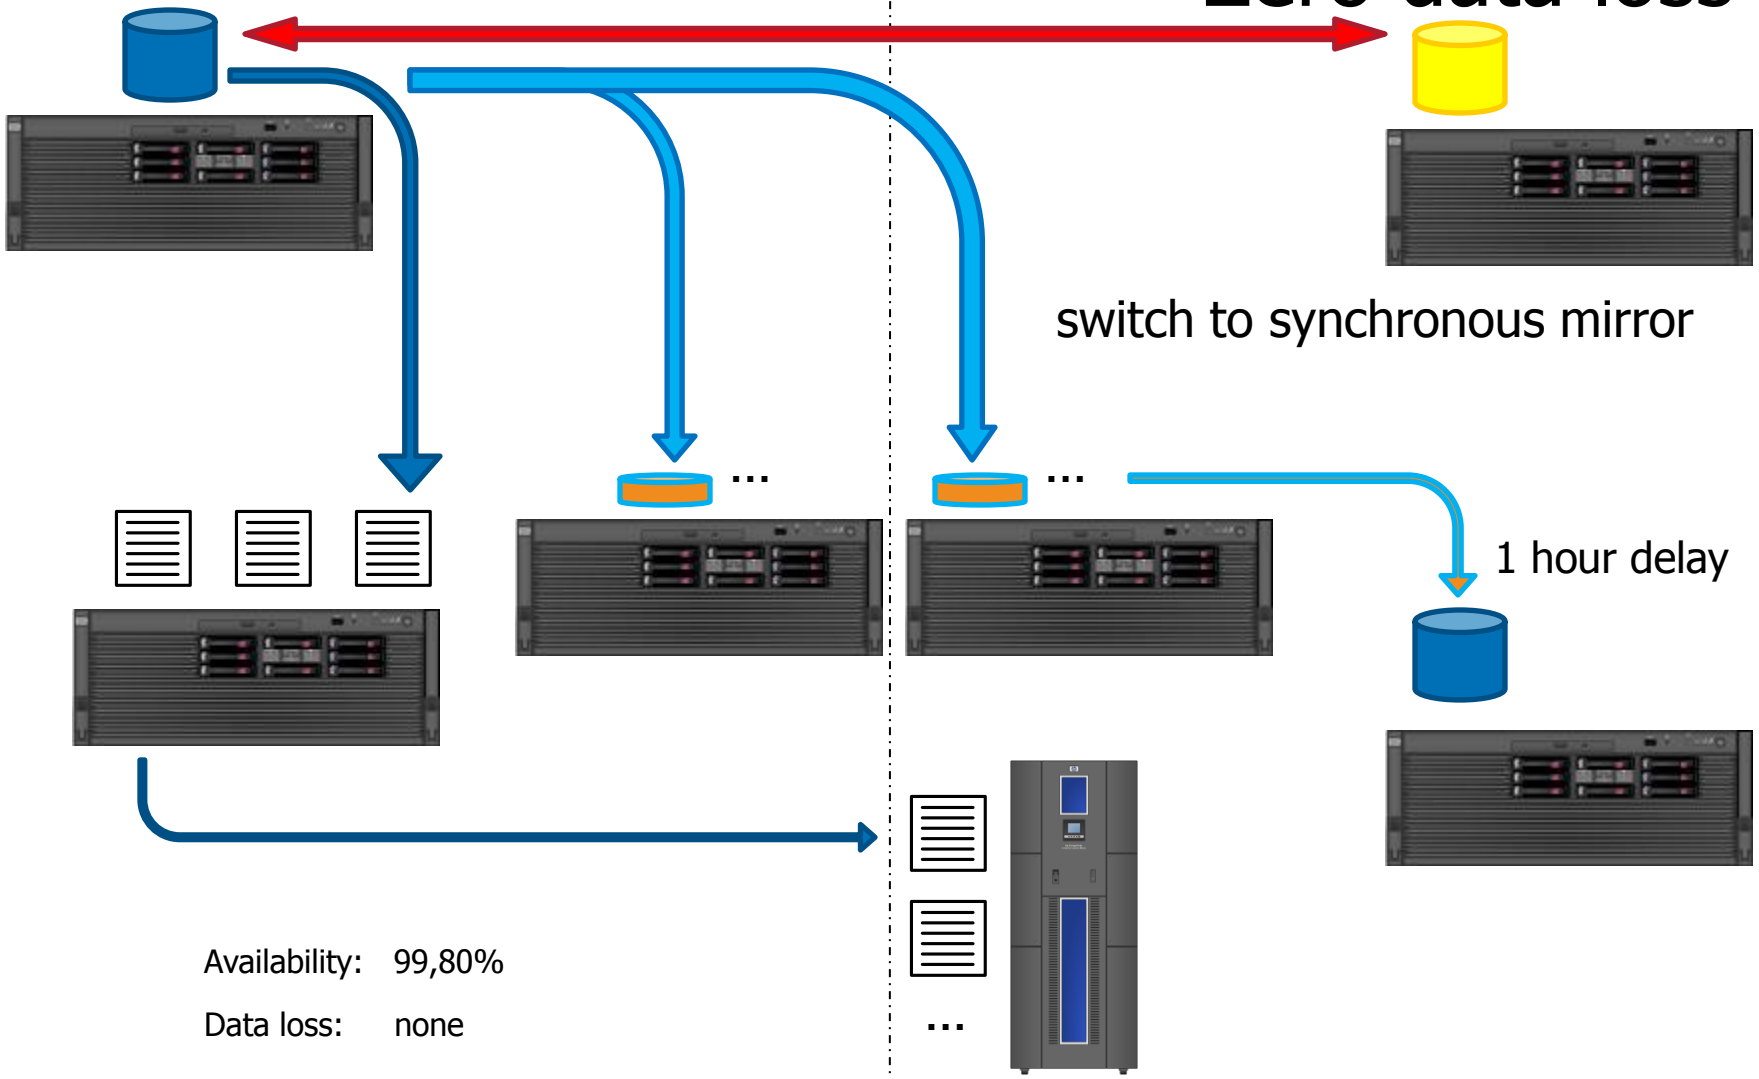

SQL Server on HP Integrity Availability

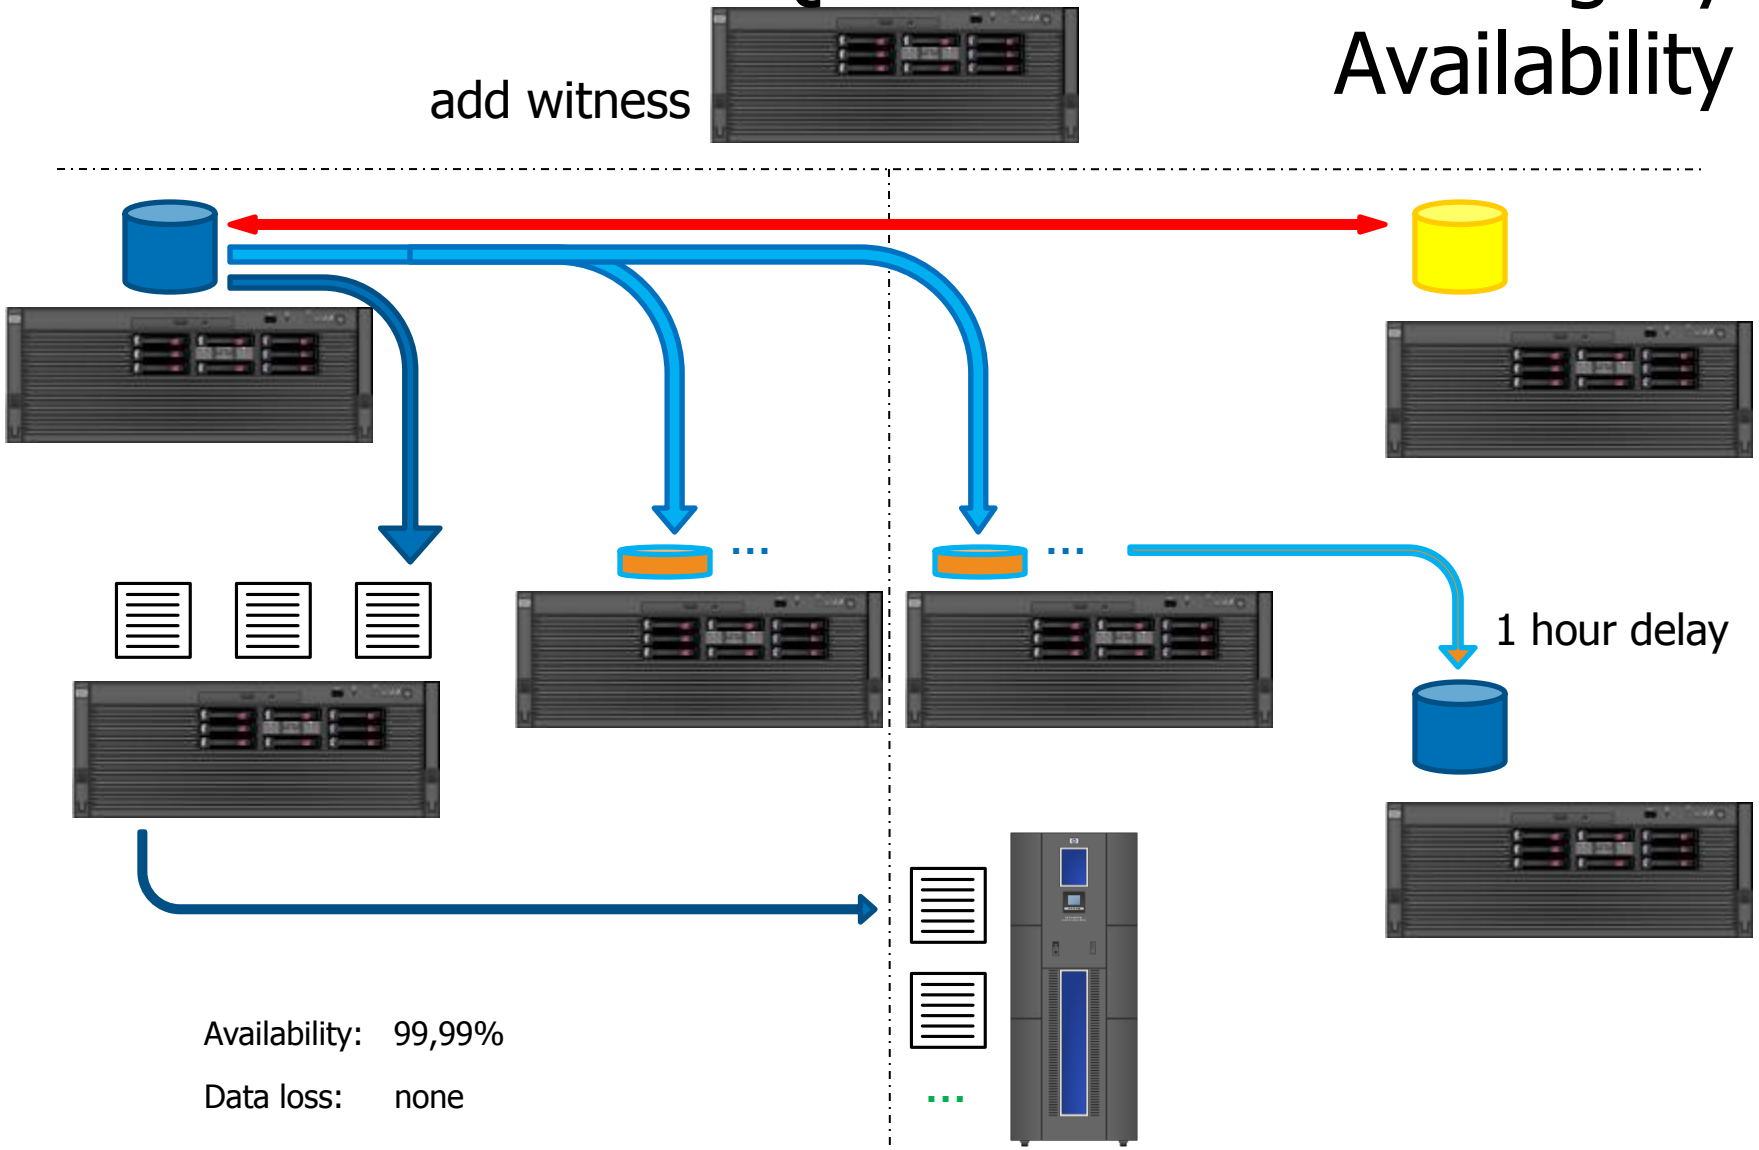

SQL Server on HP Integrity Scale Up

- Selected CPU: IA64 / Itanium 2
- Selected server/memory architecture: SMP / NUMA
- Windows Server 2008 for Itanium-Based Systems
- SQL Server 2008 Enterprise Edition

SQL Server on HP Integrity Scale Up – Single NUMA node

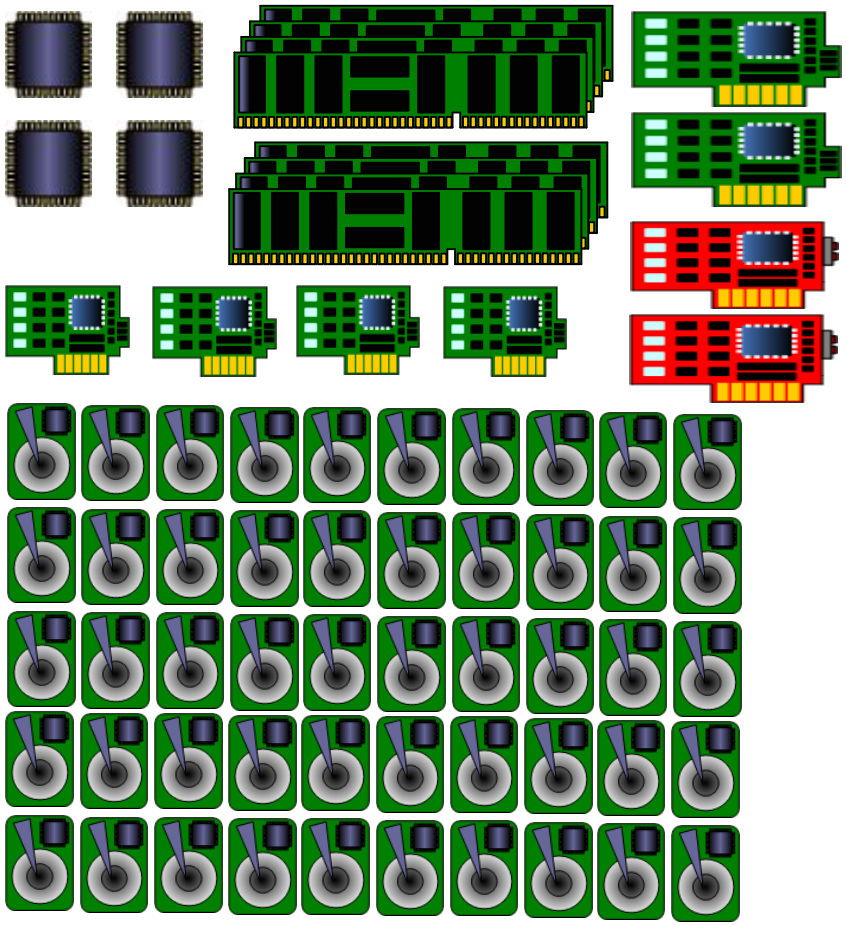

- 4 x Dual Core ITANIUM 2 CPUs
  24 MB cache each
- 64 GB memory
- 4 x dual port 1 Gb/s network card
- 2 x dual port HBA (4Gb/s)
- 2 x P800 RAID controller
- 50 x 72 GB 15kRPM SAS disks
- SAN storage as needed
  n x 512GB (on 64 spindles each)

SQL Server on HP Integrity Scale Up

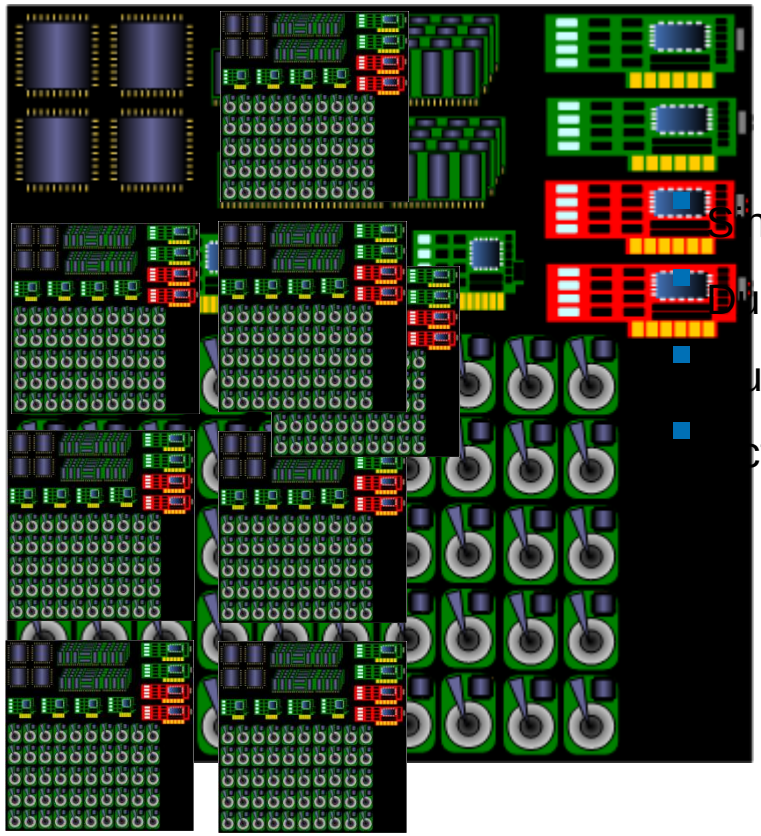

	cores	GB	disks	NIC	HBA
Single	8	64	50	8	4
Dual	16	128	100	16	8
Quad	32	256	200	32	16
Octal	64	512	400	64	32

Almost linear scaling

SQL Server on HP Integrity Scale Up

1 NUMA Node Server (rx6600)

- 1 x NUMA node basic configuration

plus

- 2 x P600 (512MB cache)
- 16 x 72 GB 15kRPM SAS disks

SQL Server on HP Integrity Scale Up

- 2 NUMA Node Server (rx7640)
  - 2 x NUMA node basic configuration

SQL Server on HP Integrity Scale Up

4 NUMA Node Server (rx8640)

- 4 x NUMA node basic configuration

plus

- 2 x single port 10 GE NIC

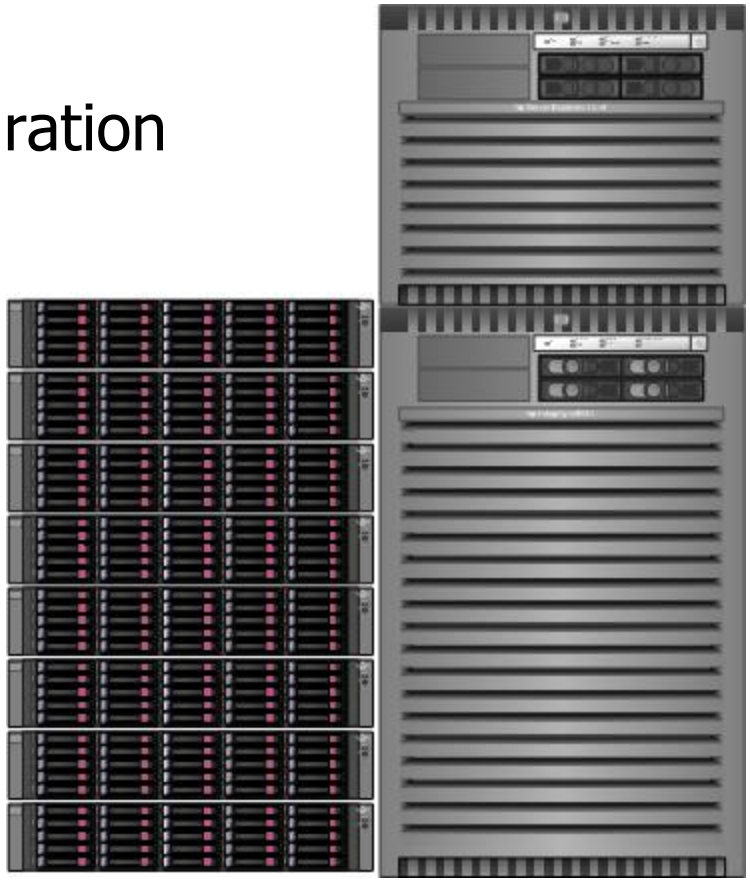

SQL Server on HP Integrity Scale Up

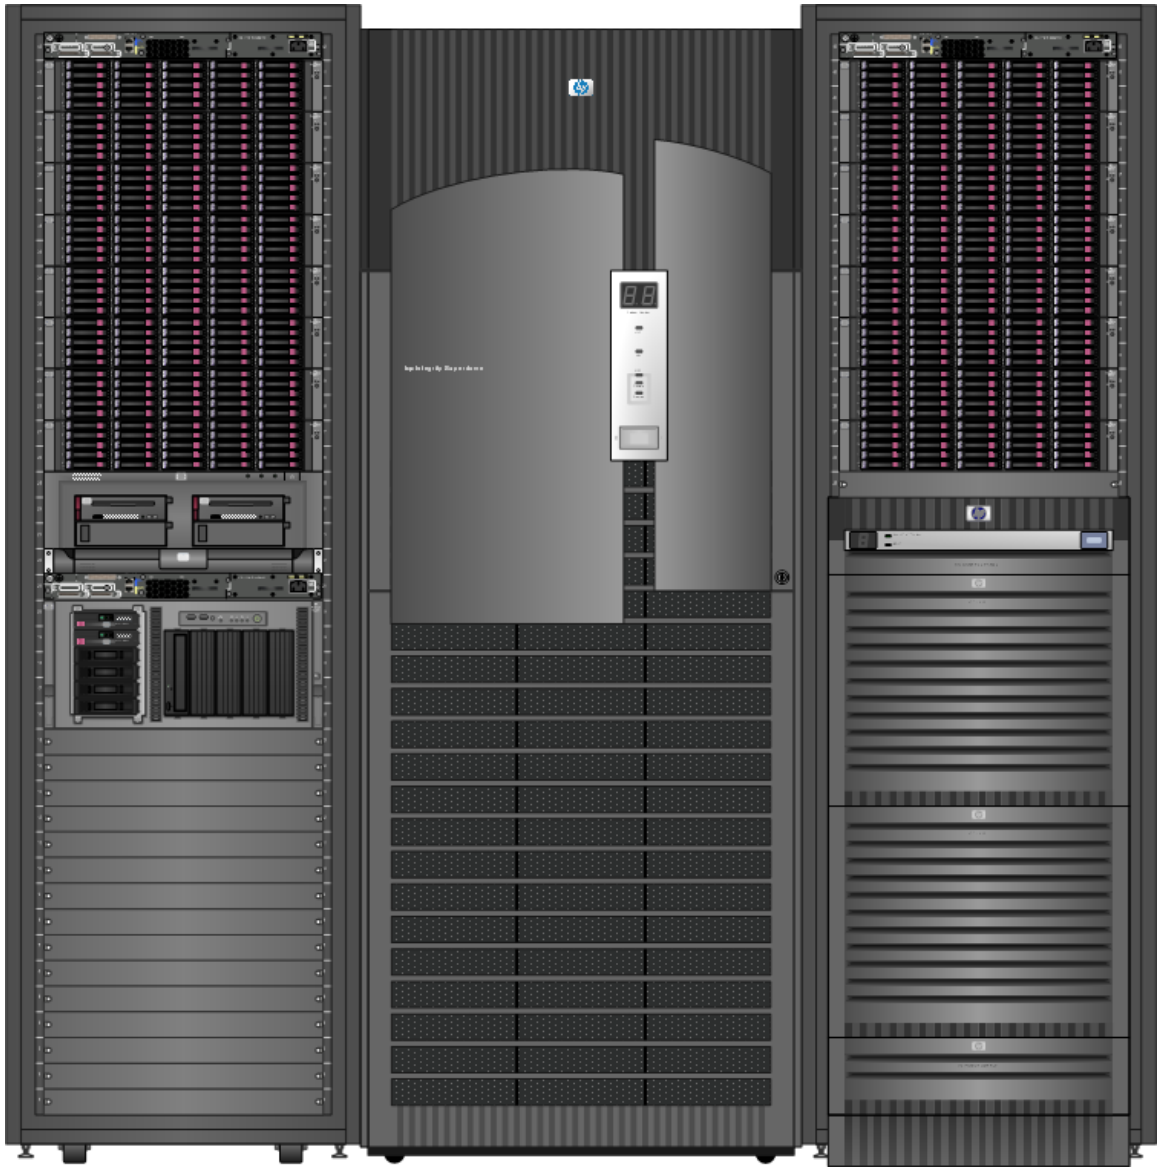

- 8/16 NUMA Node Server
- HP Superdome
  - 8/16 x NUMA node basic configuration

  plus
  - 2 x single port 10 GE network card

SQL Server on HP Integrity Scale Up

HP Superdome as principal and mirror servers

8 NUMA Nodes	TOTAL
4 dual corer IA64 CPU's / 24MB cache	64 Cores / 768 MB cache
– 64 GB Memory	512 GB
– 2 HBA / 2 x 4 Gb/sec	16 HBA / 16 GB/sec
– 2 RAID Controller / 1 GB/sec	16 RAID Controller / 16 GB/sec
– 4 x 1GB NIC	32 NIC / 8 GB/sec

Additional

NODE 1 and 3 (SQL Core Node)	
– 2 x 10GE NIC	2 NIC / 2 GB/sec

SQL Server on HP Integrity Scale Up

SAN configuration

- 2 Fabrics (4Gb/s based)
- 16 LUNs
- 16 exclusive 4Gb/s ports on the SAN BOX

  each LUN mapped via 8 paths to SAN BOX

  (4 required for performance)
- each LUN (512 GB) is build from 64 exclusive spindles (146GB 15kRPM) with short stroking
- 256 GB cache on SAN BOX

SQL Server on HP Integrity Scale Up

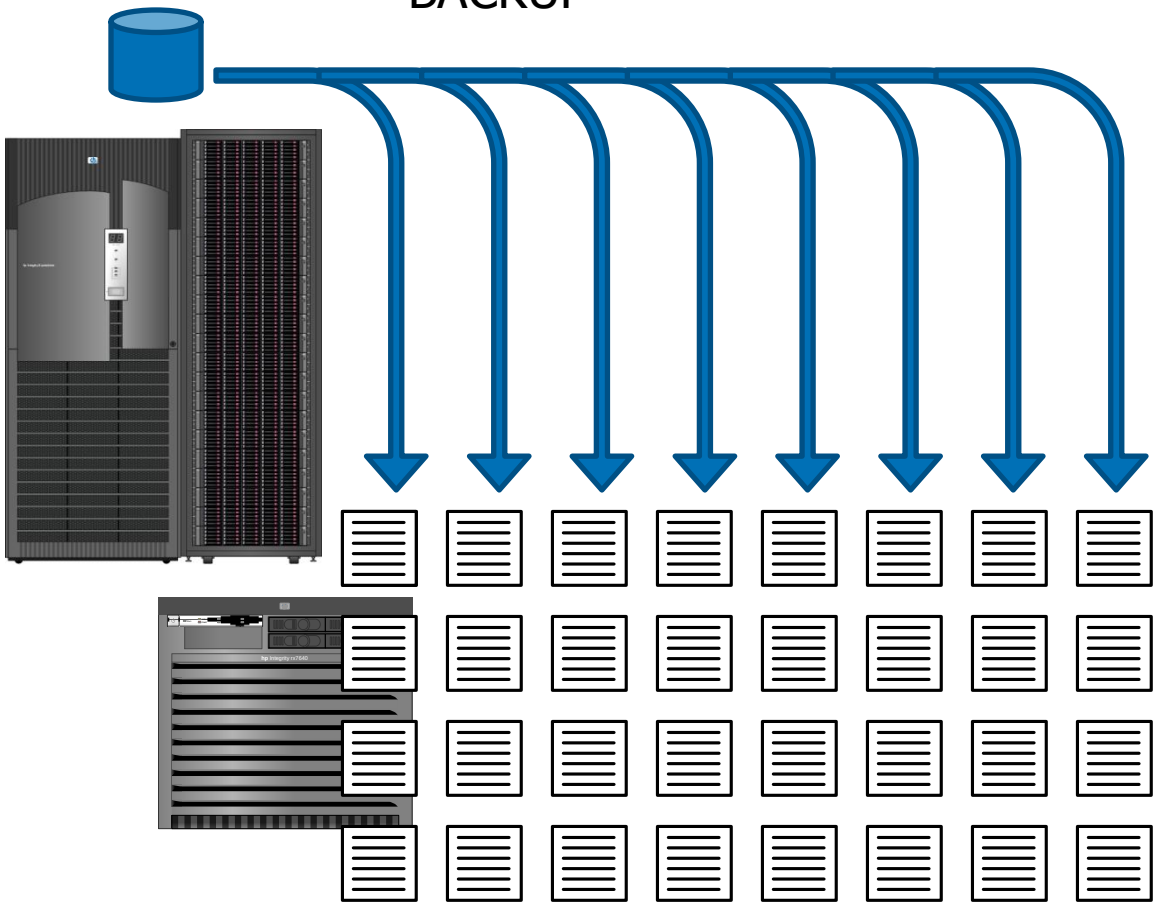

- Use eight parallel one GB/s sec network interface cards (one physical network, eight subnets)
- Use 32 parallel backup files each on a separate set of spindles with aligned partitions
- Transfer four files per network interface card

SQL Server on HP Integrity Scale Up

SQL Server	IP Address	Network Mask
Network Card 1	192.168.1.2	255.255.255.0
Network Card 2	192.168.2.2	255.255.255.0

File Server	IP Address	Network Mask
Network Card 1	192.168.1.1	255.255.255.0
Network Card 2	192.168.2.1	255.255.255.0

SQL Server on HP Integrity Scale Up

```
BACKUP DATABASE  MyVLDB
  TO
    DISK='\\192.168.1.1\backup\MyVLDB_1.bak',
    DISK='\\192.168.2.1\backup\MyVLDB_2.bak'
WITH
    BLOCKSIZE = 65536
```

- Use Jumbo Frames if you can (+100%) with 9014 bytes frame size

SQL Server on HP Integrity Maintenance

- Index defrag and index rebuilds: NONE
- Rollouts one per day
- Planned downtimes two per year
  each lasting about one hour

SQL Server on HP Integrity Reporting, DWH and BI

- Requirement:
  - have all data from previous day ready by 8 am

SQL Server on HP Integrity Reporting, DWH and BI

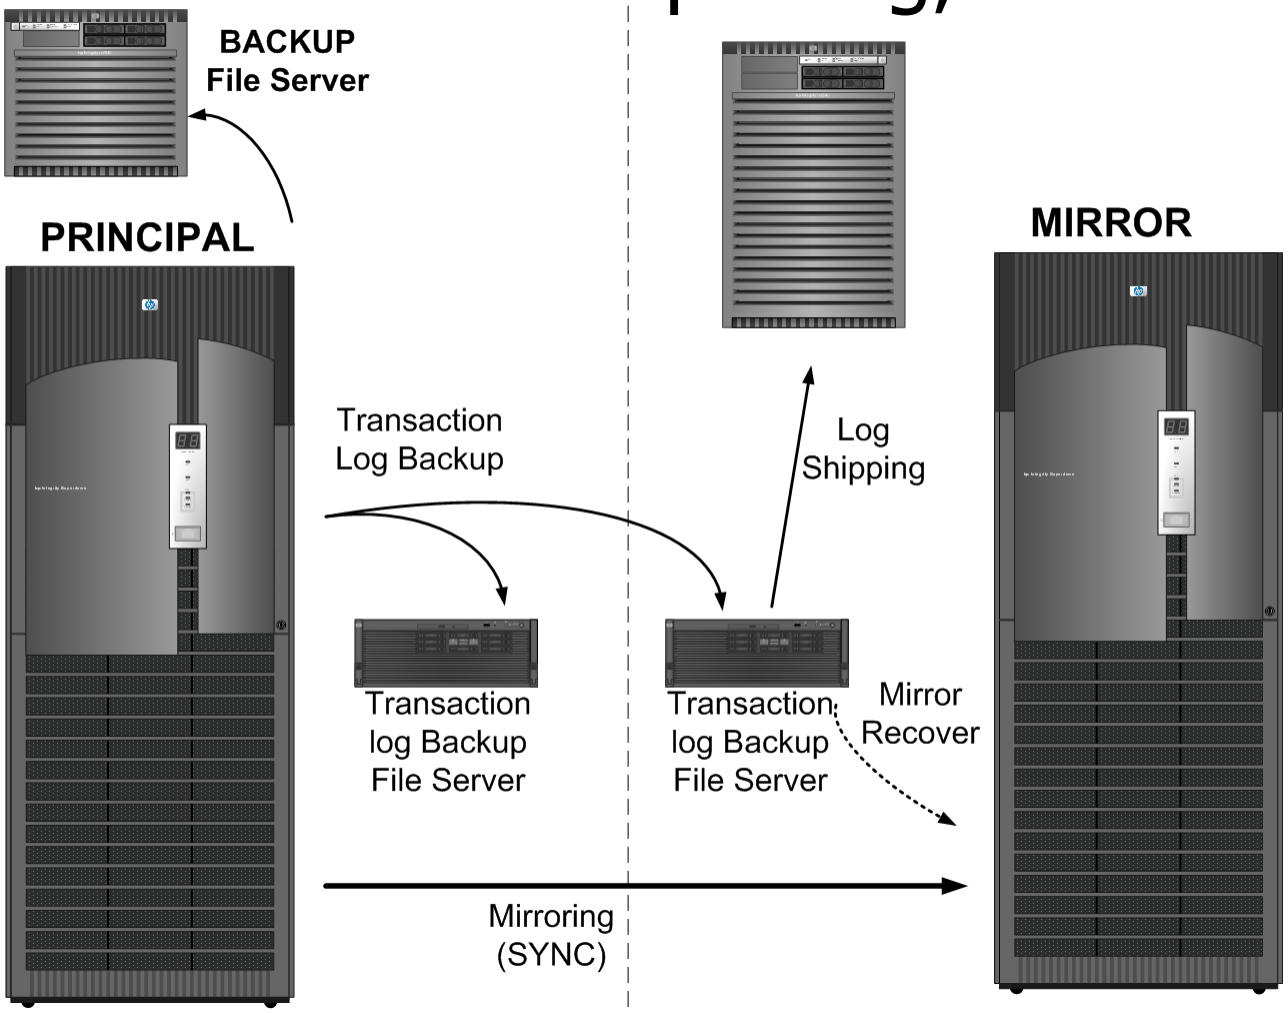

SQL Server on HP Integrity Reporting, DWH and BI

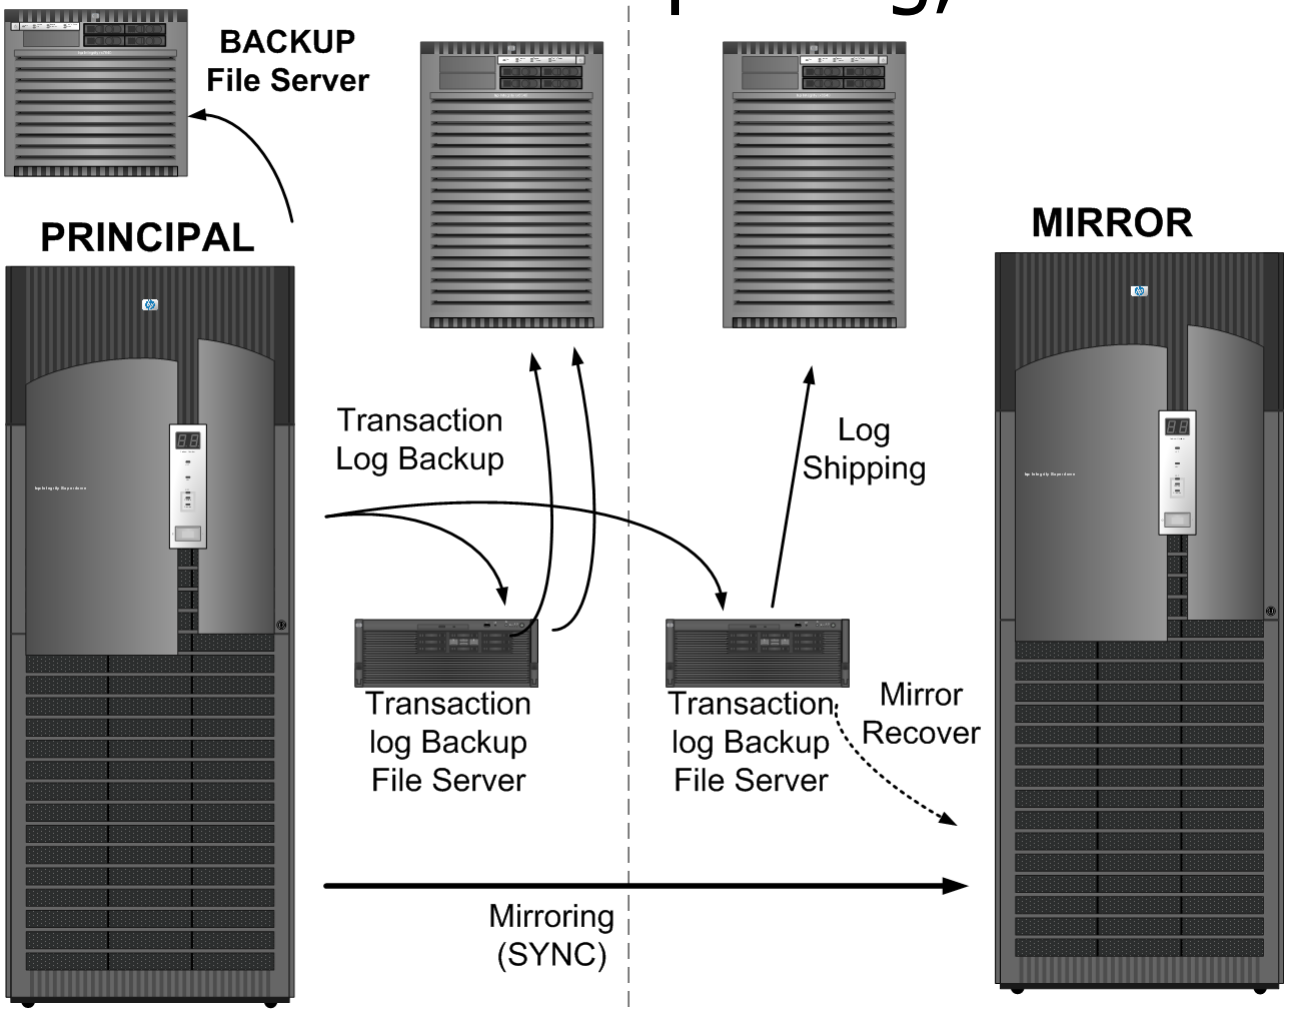

Your Independent HP Business Technology Community

SQL Server on HP Integrity
Questions

?

tg@grohser.com

Thank you!

And remember:

Failure is not an option