

IT-Symposium 2005

Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

LRP Landesbank
Rheinland-Pfalz
Ein Unternehmen der LBBW-Gruppe

+ HITACHI
DATA SYSTEMS

Present
for the
DECUS IT - Symposium 2005

Vortrag : 2E06
Erfahrungsbericht einer SAN
Implementierung mit verschiedenen OS

Mittwoch 6. April 2005 15:15
Swissôtel Düsseldorf/Neuss Raum: Mars

Partner Beyond Technology

© 2005 Hitachi Data Systems
www.decus.de

1

IT-Symposium 2005

Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

HITACHI
DATA SYSTEMS

LRP Landesbank
Rheinland-Pfalz
Ein Unternehmen der LBBW-Gruppe

**Hitachi Lightning
9970V
Storage Area Network
Implementation at
Landesbank
Rheinland-Pfalz**
Detlef Buerdorff
(LRP)
Nigel Hone
(HDS)


© 2005 Hitachi Data Systems
www.decus.de

2

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Vortrag : 2E06
Erfahrungsbericht einer SAN Implementierung mit verschiedenen OS


- **Inhalt**
- **Es werden die Erfahrungen bei einer konkreten SAN Installation, die zwei Standorte umfasst, geschildert. Es galt die Betriebssysteme OpenVMS Cluster, Windows 2000 Cluster, AIX Cluster und vor allem zweier Mainframe mit einem gemeinsamen Storage Subsystem, welches Disaster tolerant und ausfallsicher ausgelegt wurde, zu versorgen. Die hierbei gesammelten Erfahrungen und Lösungswege insbesondere die Konfiguration unter OpenVMS und Windows werden hier dargelegt.**

© 2005 Hitachi Data Systems April 6, 2005 
www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Presentation Overview

- Basic SAN construction.
- LRP Datacenter Architecture
- Quick look at Geo-Mirror techniques
- Chase an OpenVMS disk through the SAN

© 2005 Hitachi Data Systems April 6, 2005 
www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

First: The SAN basics

- **Basic SAN construction.**
 - LRP Datacenter Architecture
 - Quick look at Geo-Mirror techniques
 - Chase an OpenVMS disk through the SAN

© 2005 Hitachi Data Systems April 6, 2005 5

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

SCSI basics

- Server
- SCSI Host Bus Adapter (HBA)
 - Normally SCSI TargetID (TiD) 7
- Connects to SCSI Device
 - Has a TargetID
 - Has LUNS

- Daisy Chained to next SCSI Target
 - Up to 15 SCSI Devices
 - Last one terminated
 - remember?

Total addressable devices = 15 TiDs X 8 LUNS = 120 Devices (per HBA)

© 2005 Hitachi Data Systems April 6, 2005 6

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

SAN basics

- Server
- FibreChannel HBA *
- Connects to FC Switch
 - which connects to other switches
- This is called a „FABRIC“
- Switches also connect to Storage

- This has a 3 Byte FibreChannel address
- and up to 256 LUNs (0 to 255)
- Of course we have more storage
- This is a SAN

Of course more servers

* each HBA has unique WWN (like a MAC)

Total „theoretical“ addressable devices = 16Mio TiDs x 256 LUNs = 4096 mio

© 2005 Hitachi Data Systems April 6, 2005

www.decus.de 7

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Always two Fabrics

- We always have dual redundant fabrics
 - Made up of paired switches
- That connect to the storages

And this can be 10Km with LogWave Laser

THESE ARE JBODs

© 2005 Hitachi Data Systems April 6, 2005

www.decus.de 8

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

We use a „Virtualising Array“

- The storages(on previous slide) with 256 maximum LUNs were JBODs
 - Just a Bunch Of Disks
- Our Raid Arrays (Lightning 9970V 9980V, Thunder 9570V 9585V) „virtualise“
- There are „pairs“ of „ports“
- Connected to the two Fabrics
- JBODs are simulated *
- Each one for a different Server

© 2005 Hitachi Data Systems * Virtualisation based according to WWN of server

www.decus.de April 6, 2005 9

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

The „Host Storage Domain“

- Here is our management GUI

LUN	CLUDEV	Emulation	Capacity	RAID	Paths
00					
01	01:59	OPEN-E	13.56 GB	5(3D+1P)	2
02	01:55	OPEN-E	13.56 GB	5(3D+1P)	2
03	01:56	OPEN-E	13.56 GB	5(3D+1P)	2
04	01:57	OPEN-E	13.56 GB	5(3D+1P)	2
05	01:58	OPEN-E	13.56 GB	5(3D+1P)	2
06	01:37	OPEN-E	13.56 GB	5(3D+1P)	2
07	01:38	OPEN-E	13.56 GB	5(3D+1P)	2
08	01:39	OPEN-E	13.56 GB	5(3D+1P)	2

Selected LUNs:0 Remaining LUNs(Port):472 Remaining LUNs(GRP):244

HostStorageDomain | Host WWNs | Emulation

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

The point is:

- **We have dual ported hosts for error tolerance (Redundance)**
 - two HBAs per server
 - two fabrics
 - two storage ports defined for each disk
- **We have two FibreConnected datacenters (Desaster Recovery)**
 - Connected via LongWaveLaser

© 2005 Hitachi Data Systems April 6, 2005 **11**

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

The Landesbank Rheinland-Pfalz Data Center

- Basic SAN construction.
- **LRP Datacenter Architecture**
- Quick look at Geo-Mirror techniques
- Chase an OpenVMS disk through the SAN


© 2005 Hitachi Data Systems April 6, 2005 **12**

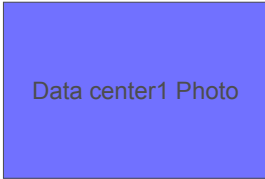
www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

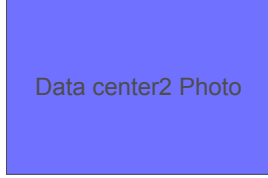
LRP Data Center Buildings

- Two data centers [409m] apart. They are easily reached by building interconnections.





Data center1 Photo



Data center2 Photo

© 2005 Hitachi Data Systems April 6, 2005 **13**

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

SAN Overview

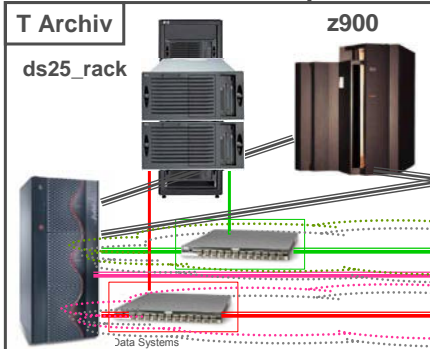
- two data centers
 - 409m apart
- 4 switches
 - in two Fabrics
- 2 x Lightning 9970V
 - Connected to Fabrics
- OpenVMS server farm
 - Also Connected to Fabrics

- Windows Server Farm
 - Also connected to Fabrics
- We also have an AIX server Farm
 - Not shown here
- zOS Mainframe
 - Connected with ESCON
- And a „TrueCopy“ connection

← 409m →

T Archiv


ds25_rack



Data Systems

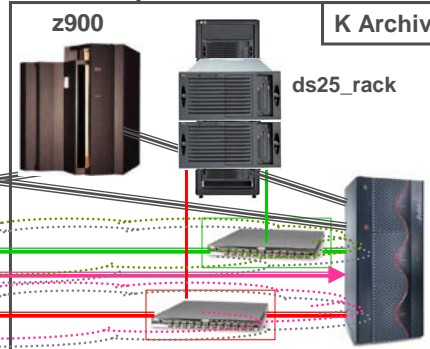
TagesArchive (was for tapes)

z900



K Archiv

ds25_rack



KellerArchive (was for tapes)

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

The Lightning 9970V



Service Processor

Interfaces

Batteries

© 2005 Hitachi Data Systems April 6, 2005 **15**

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

The storage cabling!

- This is what our technician draws!
- Pen colors correspond to cable colors
 - Red
 - ESCON for mainframes
 - Green
 - Shortwave Laser to Switches
 - Yellow
 - Longwave Laser for TrueCopy and ISL (Inter Switch Link)

Lightning 9970Vs


Mainframes

note redundant blade pairs

Switches

ISL

TC



© 2005 Hitachi Data Systems April 6, 2005 **16**

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

The Lightning 9970V interface ports

- Dual redundant interfaces
- Cables

**Yellow(longwave)
TrueCopy**

**Green(shortwave)
to switches**

**Escon
Orange**

www.decus.de April 6, 2005 **17**

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Quick Switch Overview

24 ports

© 2005 Hitachi Data System 05 **18**

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Quick Switch Overview

- ISL: LongWL
- OpenVMS
- Lightning 9970V

Spheroson 4500: S4H1_RZ_SW11 - Microsoft Internet Explorer

Configure: Refresh-11/4/04 at 12:11:52

Ports	Switch	Management	Zoning	Security	Performance
0	ISL 0 zum SW21	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
1	ISL 1 zum SW21	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
2		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
3		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
4	Lightning 1000 CL1A	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
5	Lightning 1000 CL1C	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
6	VK_S01	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
7	HZ_W2V_ADST02	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
8	HZ_VMS_Venus	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	1 Gb/sec
9	RZ_LRPVM1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
10	RZ_Franken	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
11	RZ_Pound	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
12	RZ_Rial	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
13	VK_P01	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
14	HZ_VMS_UF1200	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
15	VK_S01	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
16	HZ_MDST_02 Murex	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate
17		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Gx Port	Negotiate

© 2005 Hitachi Data Systems
www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Recap

- **Lightning 9970V is used for**
 - multiple Operating Systems
 - TrueCopy (Storage based Mirroring)
 - Dual redundant Interface Blades
 - LongWave
 - ShortWave
 - Escon
- **LONG WAVE LASERS USED FOR MIRRORING**
 - Lightning 9970V LongWave port
 - used for TrueCopy
 - storage based mirror
 - Switch LongWave port
 - used for ISL
(Servers can access remote storage)
 - server based mirror (Through a LVM)

© 2005 Hitachi Data Systems April 6, 2005 **20**
www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Geo Mirrors

- Basic SAN construction.
- LRP Datacenter Architecture
- **Quick look at Geo-Mirror techniques**
- Chase an OpenVMS disk through the SAN

© 2005 Hitachi Data Systems April 6, 2005 21

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Three ways to do a Geo-Mirror

1. **Hitachi TrueCopy**
Storage Based Mirror
Asynchronous
Synchronous
2. **Logical Volume Manager mirroring**
Server based
3. **Shadow Database**
after a write...
„redo logs“ sent
Remote server „forward recovers“

HDS TrueCopy

LVM Mirror

Shadow Database

© 2005 Hitachi Data Systems April 6, 2005 22

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Geo-Mirror Recap

- TrueCopy
- Logical Volume Manager
- Shadow DataBase
 - (not used at LRP)

© 2005 Hitachi Data Systems
www.decus.de April 6, 2005 **23**

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

OpenVMS

- Basic SAN construction.
- LRP Datacenter Architecture
- Quick look at Geo-Mirror techniques
- **Chase an OpenVMS disk through the SAN**

© 2005 Hitachi Data Systems
www.decus.de April 6, 2005 **24**

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

OpenVMS disk configuration

- Virtual Unit
 - DSA132
- 2 ShadowSet members
 - \$1\$DGA4441
 - \$1\$DGA345
- multiple paths (2 of them)
 - \$1\$DGA4441
 - PGB0.5006-0E80-0359-FC00
 - PGA0.5006-0E80-0359-FC10
 - \$1\$DGA345
 - PGB0.5006-0E80-035A-0100
 - PGA0.5006-0E80-035A-0110

0359FC = machine serial number in hex 035A01 = machine serial number in hex

www.decus.de April 6, 2005 25

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

The OpenVMS commands

- VENUS_SYSTEM >> show dev dg


```

$1$DGA345: (VENUS) ShadowSetMember 0 (member of DSA132:)
$1$DGA4441: (VENUS) ShadowSetMember 0 (member of DSA132:)

Disk $1$DGA345:, device type HITACHI OPEN-E, is online, device has multiple I/O
paths, member of shadow set DSA132:, served to cluster via MSCP Server,
error logging is enabled.

Host name          "VENUS"   Host type, avail AlphaServer 4100 5/466
Alternate host name "UF1281"  Alt. type, avail AlphaServer 4100 5/466

I/O paths to device      3
Path PGA0.5006-0E80-035A-0110 (VENUS), primary path, current path.
Path PGB0.5006-0E80-035A-0100 (VENUS).
Path MSCP (UF1281).
```

© 2005 Hitachi Data Systems April 6, 2005 26

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

A quick word about the WWN

- 5006-0E80-035A-0110 is an Hitachi WWN
 - 0060E8 = OUI (Organizationally Unique Identifier)
 - HITACHI COMPUTER PRODUCTS (AMERICA), INC.
 - See <http://standards.ieee.org>
 - 0035A01= Machine serial number in HEX
 - this is 219649 In Decimal (for the bookkeepers)
 - 10 is the port address
 - 00 = CL1-A , 01 = CL1-B etc
 - 10 = CL2-A , 11 = CL2-B etc

© 2005 Hitachi Data Systems April 6, 2005 27

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Word about the Ldev

- Remember this?:
 - here are the Ldevs
 - OpenVMS uses them
 - 0159 for instance
 - keep them different on different systems
 - we had issues with this

```

P00_VENUS>>> wwidmgr -show wwid
polling kgpsa0 (KGPSA-C) slot 4, bus 0 PCI, hose 1
kgpsaa0.0.0.4.1  PGA0  WWN 2000-0000-c926-270c
polling kgpsa1 (KGPSA-C) slot 5, bus 0 PCI, hose 1
kgpsab0.0.0.5.1  PGB0  WWN 2000-0000-c926-265c
[0] UDID:4441 WWID:5006-0e80-0359-fc10-0000-59fc-0000-1159
[12] UDID:345 WWID:5006-0e80-035a-0110-0000-5a01-0000-0159
            
```

LUN	CULDEV
00	
01	01:59
02	01:55
03	01:56
04	01:57
05	01:58
06	01:37
07	01:38
08	01:39

© 2005 Hitachi Data Systems April 6, 2005 28

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Recap

- Use Volume Shadowing and a LongWave ISL to create a GeoCluster.
- Use Dual pathing for resilience
- Understand the HBA and WWN addressing
- Be careful with Internal LDEV addresses.

© 2005 Hitachi Data Systems April 6, 2005 **29**

www.decus.de

IT-Symposium 2005 Vortrag 2E06 06/04/2005
HITACHI
Inspire the Next

Questions

- **Basic SAN construction.**
- LRP Datacenter Architecture
- Quick look at Geo-Mirror techniques
- Chase an OpenVMS disk through the SAN

Hmm..... do we have time for questions?

© 2005 Hitachi Data Systems April 6, 2005 **30**

www.decus.de

IT-Symposium 2005 Vortrag 2E06 08/04/2005
HITACHI
Inspire the Next

That's it

DONE

Thanks for your
attention

© 2005 Hitachi Data Systems
www.decus.de April 6, 2005 **31**