

iSCSI or Fibre Channel



David Pašek | DELL Datacenter Solution Architect
LinuxMeeting Ústí nad Labem
20.11.2009

Storage technologies

What is What?

- internal storage (not subject of this presentation)
 - block level storage
 - BUSes for access IDE, SATA, SCSI, SAS
- DAS (not subject of this presentation)
 - block level storage
 - there is not switch (direct attached)
 - BUSes for access SAS, FC, iSCSI, USB, eSATA
- NAS (not subject of this presentation)
 - **shared** file level storage – **storage consolidation** – **file server consolidation**
 - optimized file server – NFS, CIFS
- **SAN (this is what this presentation covers)**
 - **shared** block level storage – **storage consolidation**
 - BUSes for access (FC, iSCSI, FCoE)



iSCSI vs FibreChannel

What are technological differences?

Fibre Channel

- HBA
- Dedicated high-speed network
- Speed
 - 1, 2, 4, 8 Gbps and 10 Gbps (not backward compatible)
- Multipathing (redundancy, performance)
- Security
 - Zoning, LUN Masking, vSan (cisco)
- Cabling
 - optical fiber, coaxial copper, twisted pair copper
- Topology
 - Point-to-point, loop, switched mode

iSCSI

- standard NIC, optimized NIC (TOE), HBA
- Shared or Dedicated high-speed network
- Speed
 - 1 Gbps or 10 Gbps (is coming)
- Multipathing (redundancy, performance)
- Security
 - bi-directional chap, vlan
- Cabling
 - optical fiber, twisted pair copper
- Topology
 - Point-to-point, switched mode



iSCSI vs FibreChannel

What are technological differences?

Fibre Channel

- special protocol for storage
- loss free protocol
- SCSI encapsulated in FC

iSCSI

- running on top of Ethernet and TCP/IP
- higher overhead
- SCSI encapsulated in standard TCP protocol



iSCSI vs FibreChannel

Performance

Disk subsystem performance metrics

- throughput MBps
- transactions IOPS (tps)
 - host IOPS
 - storage IOPS => host IOPS plus RAID overhead

RAID overhead

- overhead for write transaction
- RAID 1 – storage IO = 2x host write IO
- RAID 10 – storage IO = 2x host write IO
- RAID 3, 5 – storage IO = 4x host write IO (EMC RAID 5 algorithm)



iSCSI vs FibreChannel

Performance

Disk throughput needs bandwidth

- 8Gb FC has benefit against 1Gb iSCSI
- iSCSI MPIO can help – linux openiscsi initiator up to 8 paths
- but MPIO for 4Gb FC can beat 1Gb iSCSI
- 10Gb iSCSI is the iSCSI future for high bandwidth demand

Disk transactions

- Disk spindles are needed
 - even 1Gb bandwidth can handle hundred thousands transactions
- Disk inside any storage is SATA, SCSI or SSD (flash)
- We need disk speed and combine more disk spindles (RAID)
 - SATA 7200 RPM = approx. 70 IOPS
 - SCSI 10k RPM = approx. 150 IOPS
 - SCSI 15k RPM = approx. 180 IOPS
 - EMC EFD = **approx. 5000 IOPS**



iSCSI vs FCoE

Everything over Ethernet

- iSCSI runs over Ethernet and TCP/IP which all admins are familiar with
- that's why **FCoE** (Fibre Channel over Ethernet) is coming (no TCP/IP, only over Ethernet)
- **FCoE** unified and shared network for all datacenter leveraging 10Gb Ethernet (DCE)
- DCE – Data Center Ethernet – Data Loss-free Ethernet
- **FCoE** pretty new - wait some time for production deployment or try to be early adopter with all consequences



iSCSI vs FibreChannel

What is better for me?

- **Is MORE always better?**
- **You must know real requirements**
- **Each environment is different**
- **Measurement is essential**
- **Price, TCO, ROI**
 - **FC is normally more expensive**
 - **iSCSI can be cheaper**
 - **iSCSI can be sometimes more expensive than FC**
 - **Hardware HBA, dedicated CISCO ethernet switches, enterprise iSCSI storage**

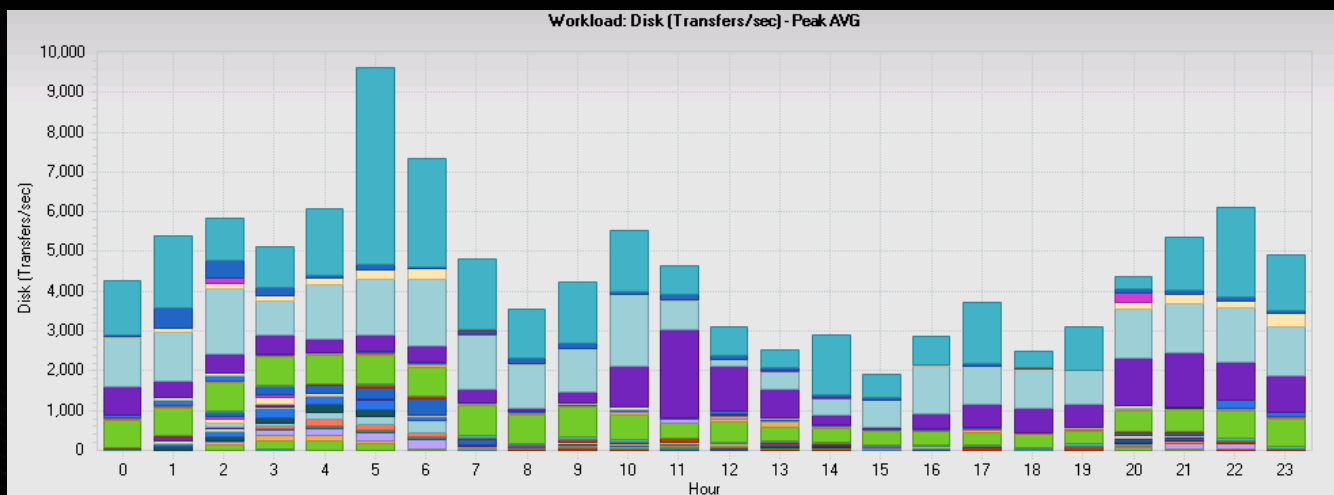
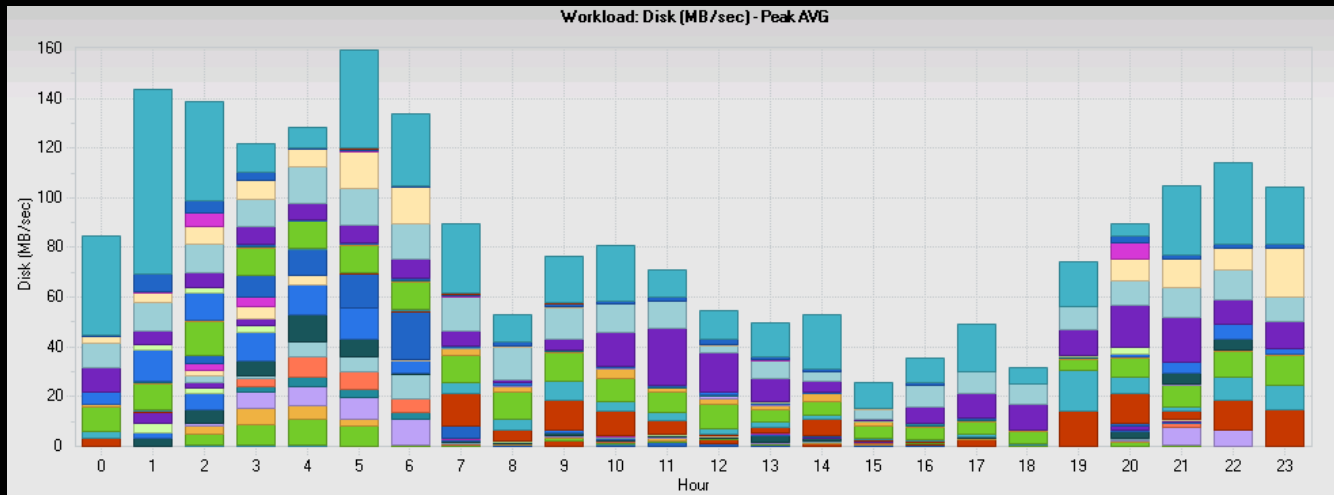


iSCSI vs FibreChannel

Real measurements from Czech datacenters

Small environment 12 servers – 160MBps and 10.000 IOPS

10.000 host IOPS = 15.000 IOPS in R10 (R:W 50:50) = 88 disks 15k RPM




iSCSI vs FibreChannel

Real measurements from Czech datacenters

Small environment 14 servers – 30MBps and 2.000 IOPS

2.000 host IOPS = 3.000 IOPS in R10 (R:W 50:50) = 18 disks 15k RPM





iSCSI vs FibreChannel

Real measurements from Czech datacenters

Small environment 52 servers – 240MBps and 8.000 IOPS

8.000 host IOPS = 12.000 IOPS in R10 (R:W 50:50) = 68 disks 15k RPM





iSCSI vs FibreChannel

Real measurements from Czech datacenters

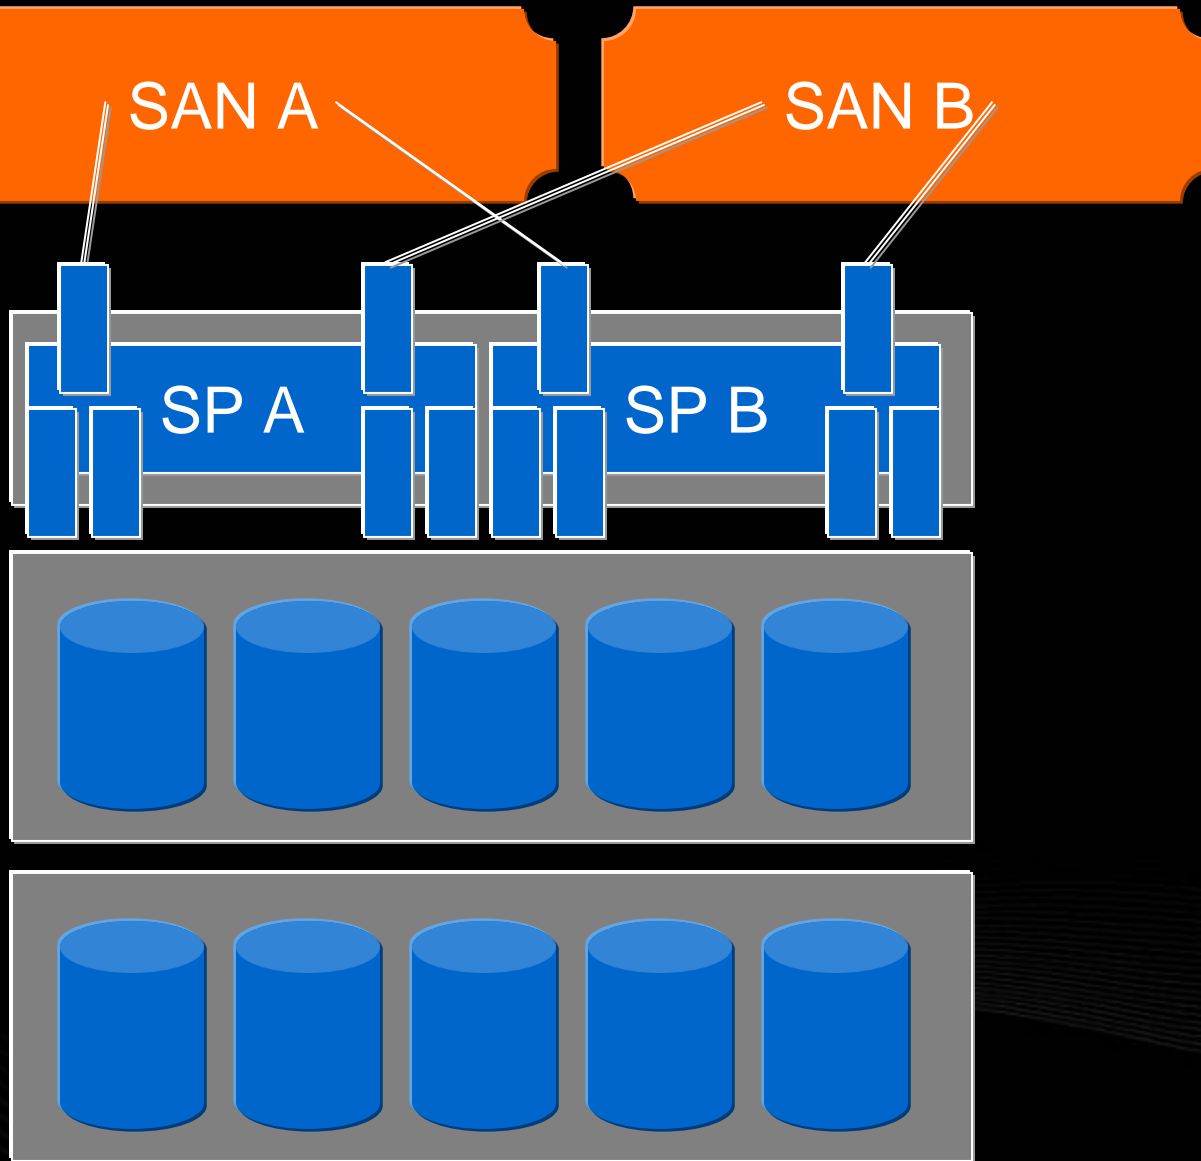
Small environment 78 servers – 300 MBps and 16.000 IOPS

16000 host IOPS = 24000 IOPS in R10 (R:W 50:50) = 133 disks 15k RPM



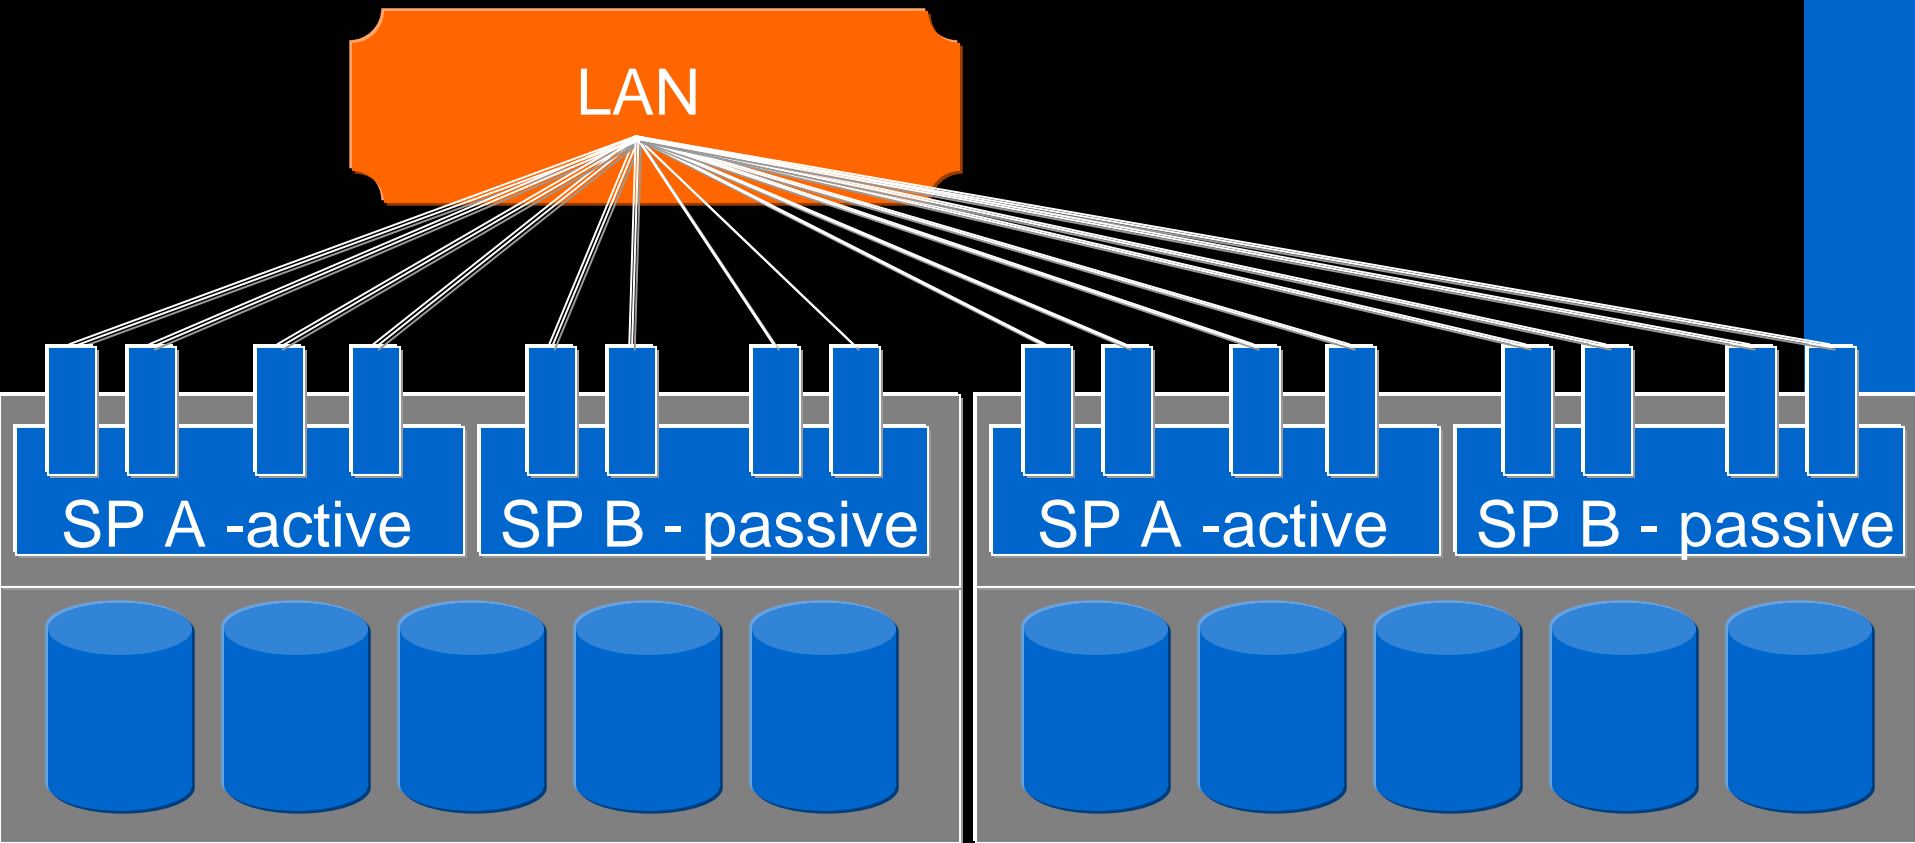
Storage Architecture

FC disk arrays – Scale Up architecture



Storage Architecture

iSCSI disk arrays – DELL EqualLogic – Scale out architecture



iSCSI vs FibreChannel

Recommendations for iSCSI

- **Two or more enterprise Ethernet switches**
- **Dedicate switches just for iSCSI SAN – don't mix LAN and SAN even it's technically possible**
- **If more than two – use two Ethernet clouds and use stacking for switches in the cloud**
- **Enable Jumbo Frame support – MTU 9000 – end-to-end**
 - **NICs, Switch, Storage ports, vSwitches if VMware ESX is in the path**



iSCSI vs FibreChannel

Recommendations for iSCSI on linux

- **Use latest stable openiscsi initiator (or iSCSI in your enterprise distro RHEL 5.3 and upper, SLES 10 and upper)**
- **Linux software iSCSI initiator is more performing then hardware iSCSI HBA but it consumes CPU**
- **Hardware iSCSI HBA allow boot from SAN and saves CPU**
- **Use multipathing – masks multiple paths and maps devices persistently**
- **Measure and monitor – Capacity Planning**



iSCSI vs FibreChannel vs FCoE

FUTURE

- **Large environment stay with Fibre Channel**
- **Large environment are testing FCoE and planning integration and deployment of FCoE**
- **Small and mid environment (up to 100 virtual or physical workloads) can use iSCSI**
 - **Customers will require TCO model for each solution**
 - **Hardware investment**
 - **Cabling investment**
 - **Supportability**
 - **Staff Training and Education**



Q & A

