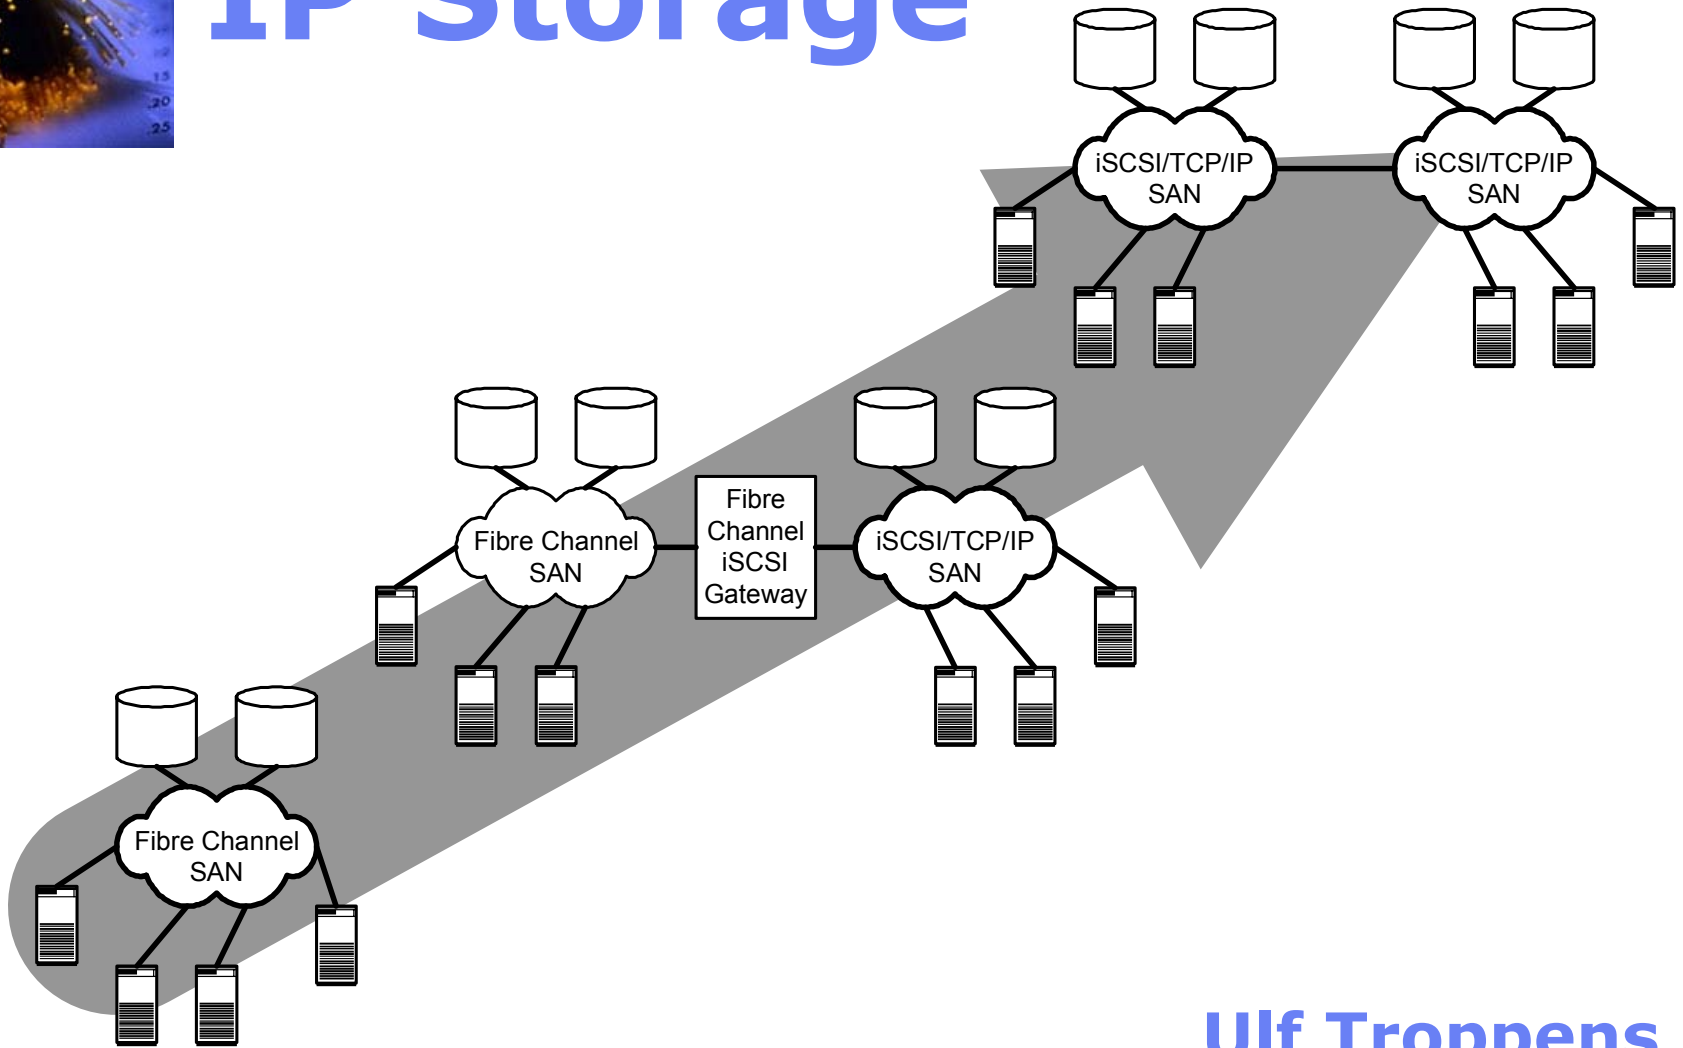


Cost Reduction with IP Storage



Ulf Troppens



Agenda

- What is IP Storage?
- Why IP Storage?
- Customers are using IP storage today!
- Is IP Storage suitable for high-end SANs?
- Conclusion




Agenda

- ***What is IP Storage?***
- Why IP Storage?
- Customers are using IP storage today!
- Is IP Storage suitable for high-end SANs?
- Conclusion



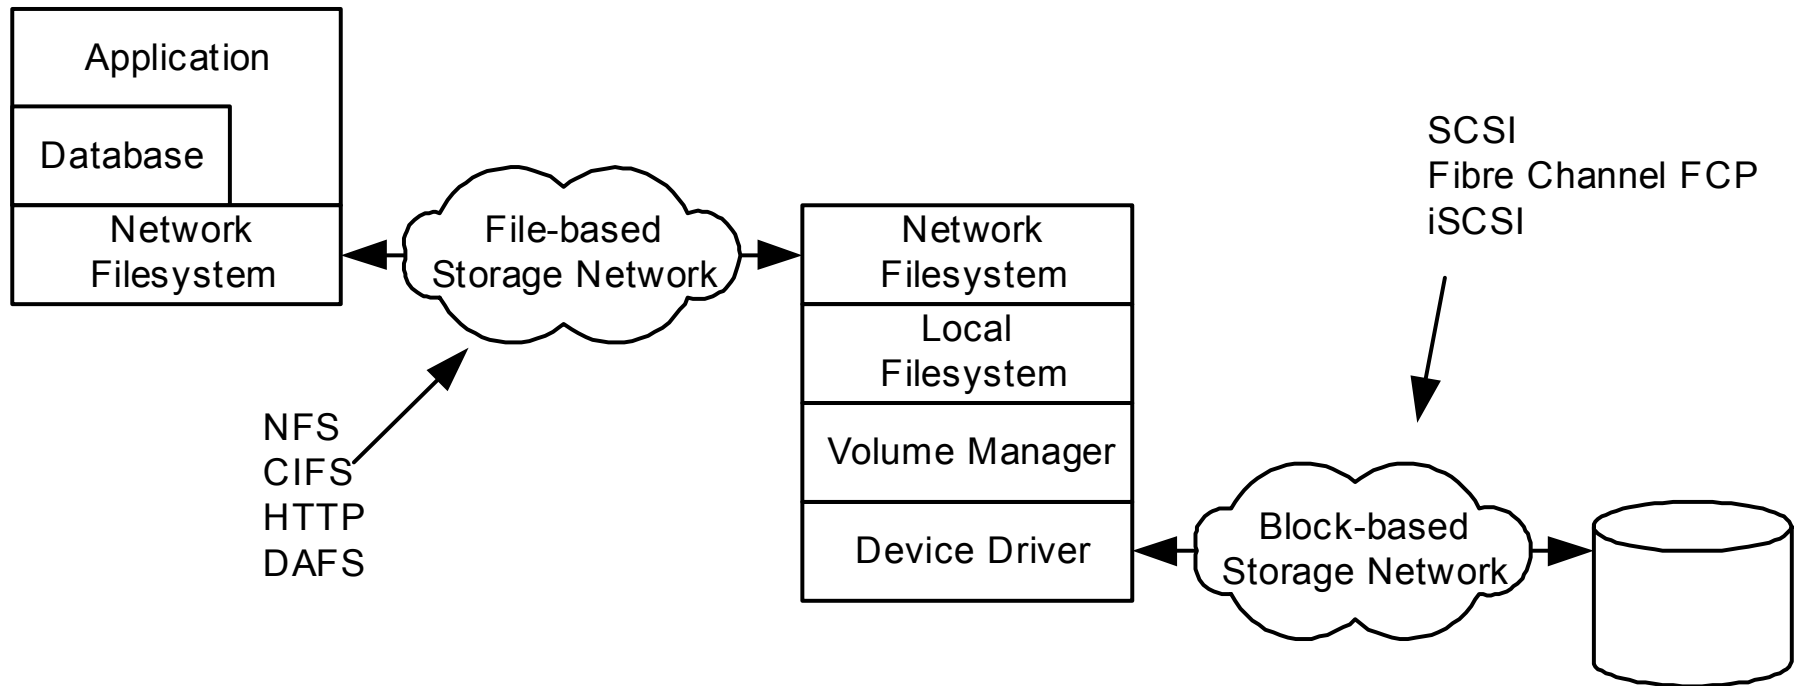
Key Concept "Storage Network"



- Detach the storage from the host and substitute the direct connection with a new network
- Handle storage and servers as separate entities
- Move from a server centric IT-architecture to a storage centric IT-architecture



Storage Networks in the I/O Path





Techniques for Storage Networks

File-based

- NFS
- CIFS
- HTTP
- DAFS

Block-based

- SCSI
- Fibre Channel (FCP)
- *Internet SCSI (iSCSI)*
- *Internet FCP (iFCP)*

RAM-based

- Remote Direct Memory Access (RDMA)

Transport protocols

- Fibre Channel (FC-2, FC-3)
- TCP/IP
- UDP/IP
- Virtual Interface Architecture (VIA) / Remote Direct Memory Access (RDMA)

Gateway and Tunneling protocols


- IP over Fibre Channel (IPFC)
- *Fibre Channel over IP (FCIP)*

Transmission Technics

- Fibre Channel (FC-0, FC-1, FC-2)
- Fast Ethernet
- GigabitEthernet
- SCSI Cable
- InfiniBand




Fibre Channel FCP



- Cabling: Fibre Channel
- Protocol: SCSI




Internet SCSI (iSCSI)



- Cabling: TCP/IP/Ethernet
- Protocol: SCSI




Fibre Channel over IP (FCIP)





Internet FCP (iFCP) Metronet FCP (mFCP)





Internet Storage Name Service (iSNS)

- ... is just a name service
 - ... used by iSCSI and iFCP/mFCP
- Note: FCIP does not need a name service!



Comparison

	Storage / Server (HBA)	Fabric Services (Routing)	Transport (Wire)
Fibre Channel FCP	Fibre Channel	Fibre Channel	Fibre Channel
FCIP	Fibre Channel	Fibre Channel	IP / Ethernet
iFCP	Fibre Channel	IP / Ethernet	IP / Ethernet
iSCSI	IP / Ethernet	IP / Ethernet	IP / Ethernet



Agenda

- What is IP Storage?
- ***Why IP Storage?***
- Customers are using IP storage today!
- Is IP Storage suitable for high-end SANs?
- Conclusion



Why IP Storage?

- Technically, Fibre Channel is better than IP
 - Fibre Channel has an integrated protocol stack
 - Fibre Channel is well established in production environments

- IP Storage (iSCSI, FCIP, iFCP/mFCP)
 - IP Storage is in the early market
 - IP Storage has an higher protocol overhead

- **BUT**
 - TCP/IP/Ethernet has utterly wiped out other protocols
 - IP Storage will be sufficient (most likely)
 - IP Storage will be cheaper (most likely)



TCP/IP/Ethernet has utterly wiped out other protocols

- Data is already IP
 - ▶ LAN (NFS, CIFS, application specific protocols, ...)
 - ▶ Internet (HTTP, FTP, ...)
 - ▶ MAN (Pilots with Gigabit Ethernet)
- File level storage is IP
 - ▶ NFS, CIFS, FTP, HTTP (see above)
- Block level storage starts using IP
 - ▶ FCIP, iSCSI
- Voice starts moving to IP (VoIP)
 - ▶ New office buildings !!
 - ▶ Cisco IBM Agreement
 - Gigabit Ethernet switch for the desktop

→Everything will be TCP/IP/Ethernet !!!



IP Storage will be cheaper than Fibre Channel

- One single IP based network is easier to manage than three
 - ▶ Only IP skills required
 - ▶ Only one physical network in office buildings
 - ▶ More flexible assignment of bandwidth

- IP Storage components will be cheaper than FC components
 - ▶ IP market volume >> Fibre Channel market volume
 - ➔ economy of scale is better for IP components (mass production)
 - ▶ More competition
 - ➔ Cisco "meets" Brocade, Inrange, and McData
 - ➔ Ethernet NIC vendors "meet" Emulex, JNI and QLogic



Why IP Storage?

- **BUT**

- TCP/IP/Ethernet has utterly wiped out other protocols ➤
- IP Storage will be sufficient (most likely) ???
- IP Storage will be cheaper (most likely) ➤




Agenda

- What is IP Storage?
- Why IP Storage?
- ***Customers are using IP storage today!***
- Is IP Storage suitable for high-end SANs?
- Conclusion



Case Study 1: Automated Tape Duplication


- Goal: Automate tape duplication for disaster protection
- Two locations (10km distance)
- Gigabit Ethernet connection used during office hours
→ **Idle over night**






Case Study 1: Automated Tape Duplication

- Goal: Automate tape duplication for disaster protection
- Two locations (10km distance)
- Gigabit Ethernet connection used during office hours
→ **Idle over night**
- Tape applications
 - ▶ iSeries (AS/400)
→ No iSCSI drivers available !
 - ▶ Legato Networker running on Sun Solaris






Case Study 1: Automated Tape Duplication





Case Study 1: Automated Tape Duplication


- Legato Networker:
Remote tape access using application specific protocol





Case Study 1: Automated Tape Duplication

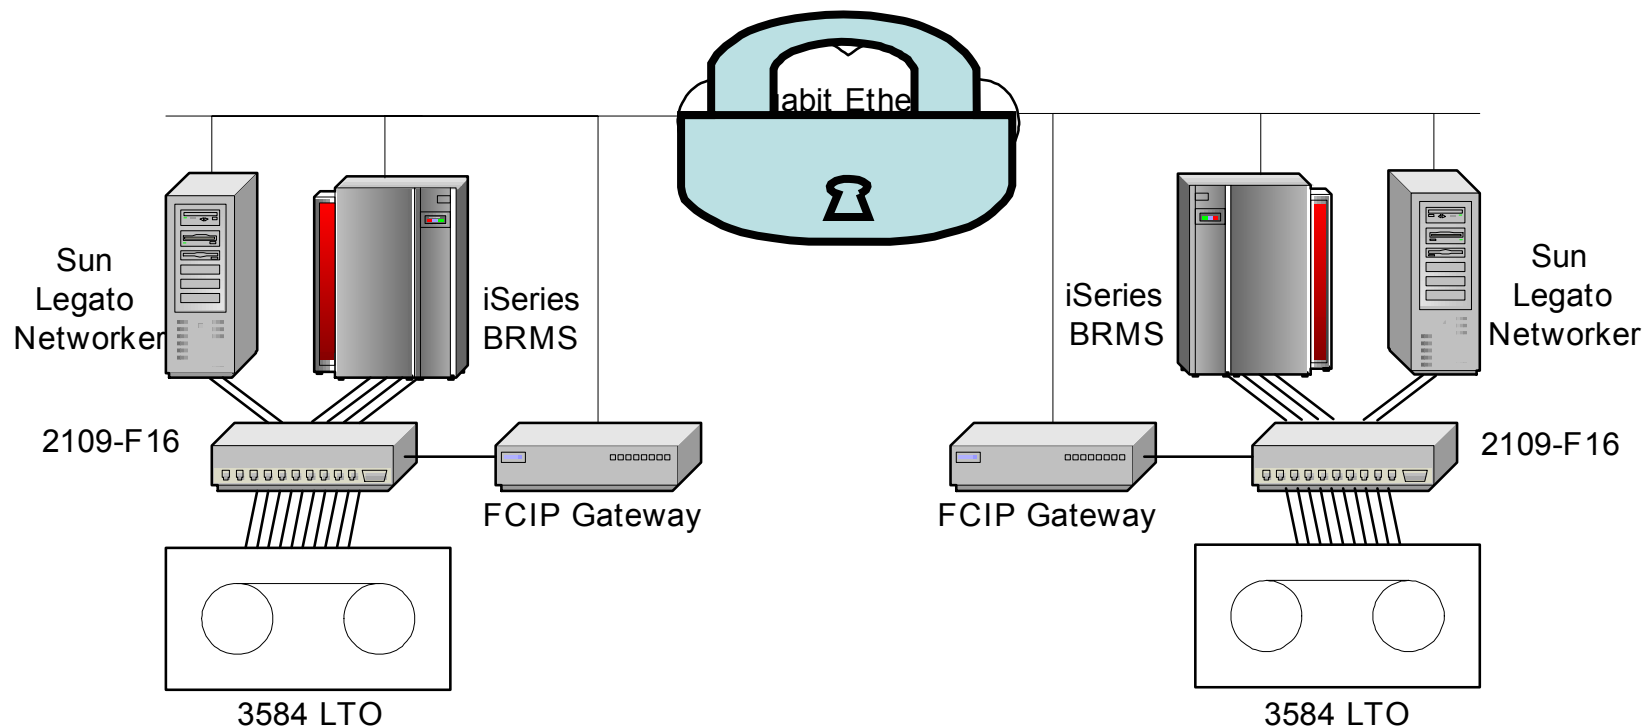
- Legato Networker:
Remote tape access using application specific protocol
- iSeries: Remote tape access using FCIP
→ 30-40 MByte/s sequential write on LTO Generation 1 drive





Case Study 1: Automated Tape Duplication


- Legato Networker: Remote tape access using application specific protocol
- iSeries: Remote tape access using FCIP
→ 30-40 MByte/s sequential write on LTO Generation 1 drive
- Future: Encrypt WAN IP traffic using standard IP components





Case Study 2: Booting via iSCSI

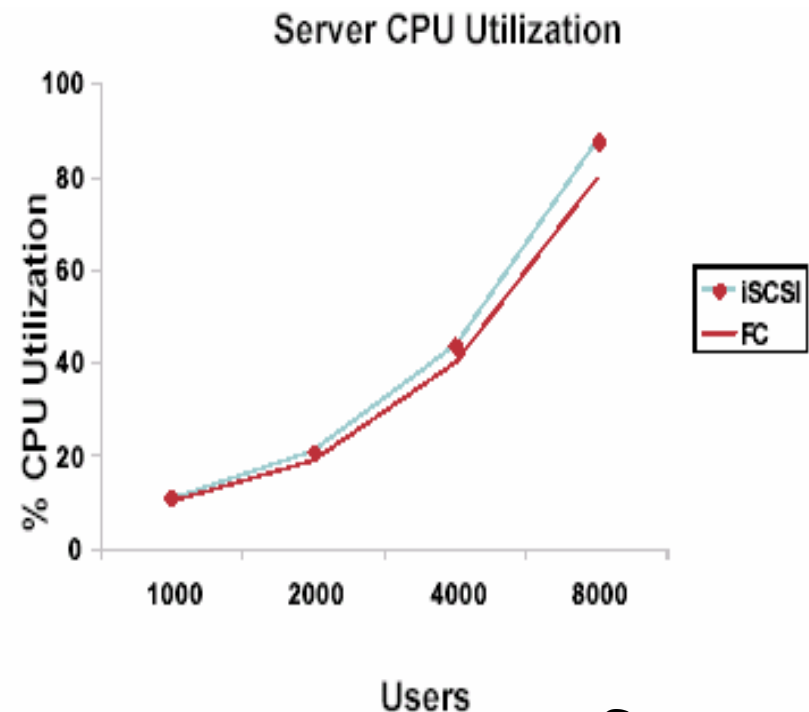

- Boot servers from external Fibre Channel Disksubsystem
→ iSCSI Disksubsystems are not (yet) very common
- Save internal Fibre Channel HBA
→ May be cost saving if one Gigabit Ethernet NIC is sufficient for all traffic
- Use external disks instead of internal disk
→ Lower TCO





Case Study 3: Microsoft Exchange

- Requires a lot of storage but limited bandwidth
- Minimal CPU utilisation differences between iSCSI software and drivers and Fibre Channel HBAs
- iSCSI meets the performance requirements for many applications



Source: Cisco



Why IP Storage?

- **BUT**

- TCP/IP/Ethernet has utterly wiped out other protocols ➤
- IP Storage will be sufficient (most likely) ➤?
- IP Storage will be cheaper (most likely) ➤



Agenda

- What is IP Storage?
- Why IP Storage?
- Customers are using IP storage today!
- ***Is IP Storage suitable for high-end SANs?***
- Conclusion



Classification of I/O Requirements

	Low latency (local)	High latency (remote)
Average performance	<ul style="list-style-type: none">■ Most applications■ Booting servers	<ul style="list-style-type: none">■ Remote backup■ Replication■ Asynchronous Mirroring■ Data Movers
Very high performance	<ul style="list-style-type: none">■ Some applications (e.g., central databases)	<ul style="list-style-type: none">■ N/A



Classification of I/O Requirements

	Low latency (local)	High latency (remote)
Average performance		
Very high performance		



Technical requirements for very high local performance

	IP Storage
Very high throughput	<ul style="list-style-type: none">▪ Dedicated Networks for IP Storage▪ 10 GBit/s IP SANs in 2H03▪ Trunking
Low CPU load for I/O	<ul style="list-style-type: none">▪ TCP/IP Offload Engines (TOEs)▪ iSCSI HBAs▪ Future: Socked Direct Protocol (SDP)▪ Future: iSCSI RDMA extension
Very low latency for fast application response time	<ul style="list-style-type: none">▪ Still open:<ul style="list-style-type: none">→ FC protocol is still lighter than iSCSI→ Latency in server >> Latency in SAN→ Future: Socked Direct Protocol (SDP)→ Future: iSCSI RDMA extension→ Must be measured in real environments



Why IP Storage?

- **BUT**

- TCP/IP/Ethernet has utterly wiped out other protocols ➤
- IP Storage will be sufficient (most likely) ➤(?)
- IP Storage will be cheaper (most likely) ➤



Agenda

- What is IP Storage?
- Why IP Storage?
- Customers are using IP storage today!
- Is IP Storage suitable for high-end SANs?
- ***Conclusion***



Why IP Storage?

- Technically, Fibre Channel is better than IP
 - Fibre Channel has an integrated protocol stack
 - Fibre Channel is well established in production environments
- IP Storage (iSCSI, FCIP, iFCP/mFCP)
 - IP Storage is in the early market
 - IP Storage has an higher protocol overhead
- **BUT**
 - TCP/IP/Ethernet has utterly wiped out other protocols ➤
 - IP Storage will be sufficient (most likely) ➤(?)
 - IP Storage will be cheaper (most likely) ➤



Conclusion

- FCIP is a key technology for connecting remote Fibre Channel SANs
- Many FCIP products are becoming available in 1H03
- iSCSI is in the early market today
→ Interoperability is still an issue
- iSCSI will meet the performance needs for most applications
- It is still open whether iSCSI is sufficient for very high performance requirements
- Fibre Channel is the best option for most local SANs today
- iSCSI can be a nice add-on for already established Fibre Channel SANs
- iSCSI may replace Fibre Channel in the long term



Quellenangabe

- Inhalte und Abbildungen der Präsentation
 - ▶ Speichernetze – Grundlagen und Einsatz von Fibre Channel SAN, NAS, iSCSI und InfiniBand
 - ▶ Ulf Troppens, Rainer Erkens
 - ▶ 1. Auflage 2003
dpunkt.verlag, Heidelberg
ISBN 3-89864-135-X
- Weiterführende Links
 - ▶ www.speichernetze.com

