

# Things are not what they appear to be

Effects of Modern Storage and Virtualisation Technologies

**OSL Data Centre Technology**  
Spring Tour 2011 • Luxembourg, 8/6/2010

**Bert Miemietz**

OSL Gesellschaft für  
offene Systemlösungen mbH

“Konsequent zu Ende denken kann man  
nur mit respektloser Heiterkeit.”

“Resolute logical thinking requires  
disrespectful cheerfulness”

- Gerhard Branstner -

# OSL offers data centre solutions for

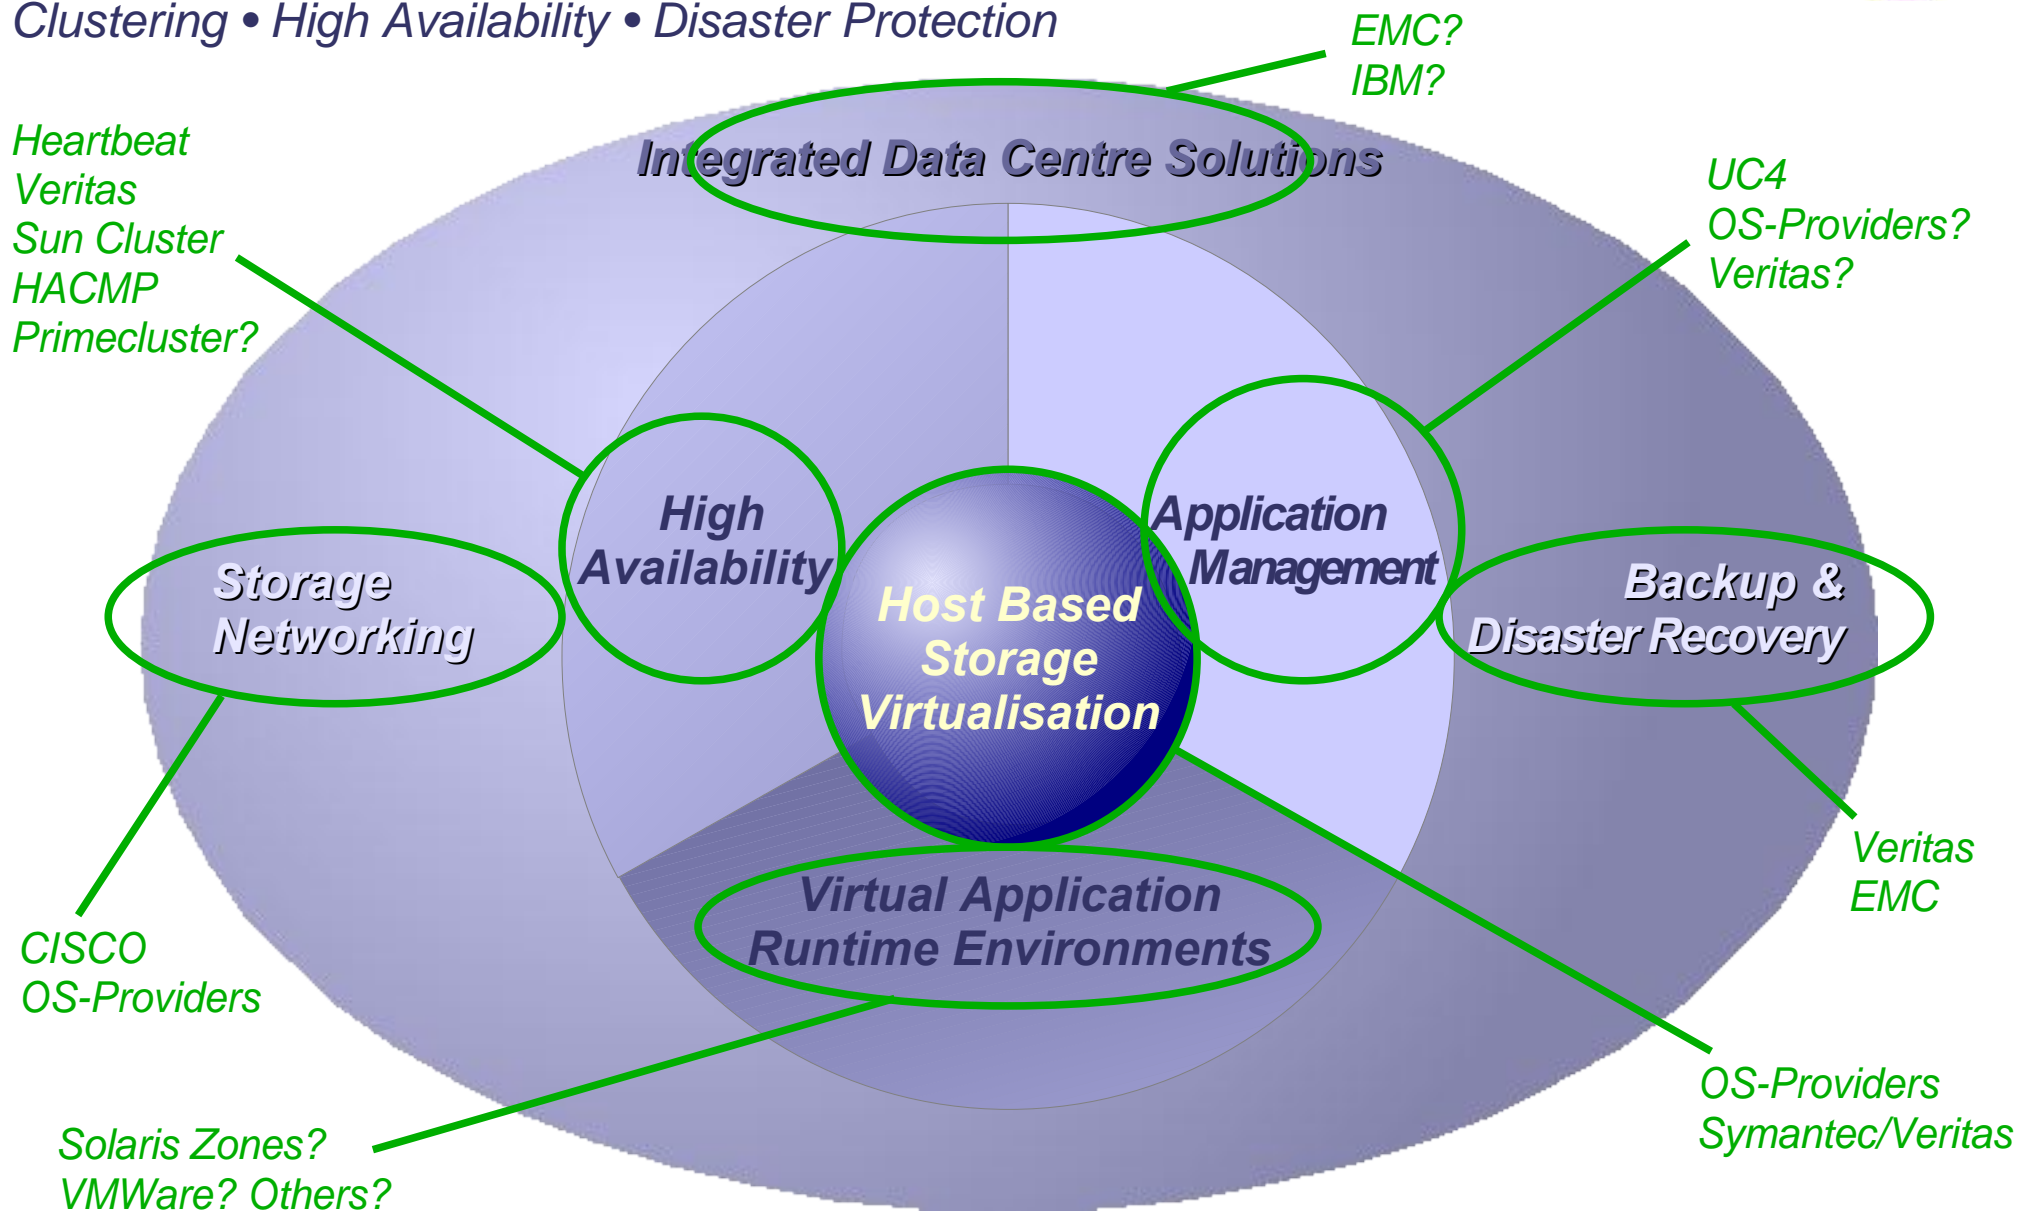
Storage Networking & Virtualization • Volume Management  
Clustering • High Availability • Disaster Protection



Heartbeat  
Veritas  
Sun Cluster  
HACMP  
Primecluster?

EMC?  
IBM?

UC4  
OS-Providers?  
Veritas?



CISCO  
OS-Providers

Veritas  
EMC

OS-Providers  
Symantec/Veritas

Solaris Zones?  
VMWare? Others?

OSL Gesellschaft für offene Systemlösungen mbH  
[www.osl.eu](http://www.osl.eu)

# *How can that be accomplished by a small company?*

*Just choose a different approach*



- *it is only software*
- *strict focus, e. g. concerning platforms (Solaris, Linux)*
- *sophisticated, modular & open design*
  - > *reduce to standard interfaces*
  - > *build your work on top of others*
- *question common paradigms -> strike new paths*
- *do both: be doubtful and develop long-term strategies*
- *deep technological knowledge, arduous work*
- *no “Enterprise Split Brain”*  
*(planning, engineering, QA, service, marketing – all in a small effective team)*
- *consequent focus on customer needs:*  
*listen carefully, think and make a proper effort!*
- *We do not dominate the market, we just try*  
*to adopt to a changing environment together with our customers*



# How does a powerful Data Centre look like?

15 years ago



*And large  
systems?*



OSL Gesellschaft für offene Systemlösungen mbH  
**[www.osl.eu](http://www.osl.eu)**

# How does a powerful Data Centre look like?

15 years ago



And very large systems?




OSL Gesellschaft für offene Systemlösungen mbH  
**www.osl.eu**

# How does a powerful Data Centre look like?

15 years ago



- *dedicated server*
- *dedicated storage*





OSL Gesellschaft für offene Systemlösungen mbH  
**[www.osl.eu](http://www.osl.eu)**

# How does a powerful Data Centre look like?


Today



*And large systems?*









- *flexible servers*
- *centralized storage*



# How does a powerful Data Centre look like?

Today



OSL Gesellschaft für offene Systemlösungen mbH

[www.osl.eu](http://www.osl.eu)




# How does a powerful Data Centre look like?

Today



Can you see a  
problem?



**SAN**




OSL Gesellschaft für offene Systemlösungen mbH  
**[www.osl.eu](http://www.osl.eu)**

# How does a powerful Data Centre look like?

Today



Can you see a problem?



OSL Gesellschaft für offene Systemlösungen mbH  
[www.osl.eu](http://www.osl.eu)

# Today's Data Centre Infrastructures

*Have a look from a different perspective*



- *the number of (virtual) servers is constantly growing*
- *mass storage today is extremely centralized*
- *servers and storage are connected by often complex storage networks, that are of no interest to the application users*

## **The other side of the medal:**

- *giant mass storage systems that are extremely expensive*
- *new work profiles: storage administrator, SAN administrator*
- *overdimensioned storage networks*
- *increasing interdependencies of load profiles*
- *“availability trap”:* - *better availability of storage infrastructure*  
- *we are increasingly dependent on that availability*
- *Performance? What are the decisive factors for storage performance?*

# Decisive Factors of I/O Performance

The most important performance factors for external RAID systems



- no influence    o little influence    +/- depends on other factors    + strong    ++ very strong


Factor	load profile	
	sequential	random
connection speed (SAN)	+	o
number of disks	++	++
parallelism of RAID system (controller etc.)	++	+
disk attachment type (SATA/SAS)	o	o
disk type (electro-mech. HD / SSD)	+	+
rotation speed of disk (e.-mech.)	o	+/- <i>Zahl der Clients</i>
cache in the RAID system	o	++
OS of the RAID system	o	+/- <i>genaues Lastprofil</i>
other techniques (queue mgmt., policies)	o	+/- <i>Lastprofil / Zahl der Clients</i>

OSL Gesellschaft für offene Systemlösungen mbH  
**www.osl.eu**



# Why not this way?

*OSL software enables a different approach*



OSL Gesellschaft für offene Systemlösungen mbH  
**[www.osl.eu](http://www.osl.eu)**

# Why not this way?

OSL software enables a different approach




<i>number of disks</i>	++	++
<i>parallelism of RAID system (controller etc.)</i>	++	+
<i>disk type (electro-mech. HD / SSD)</i>	+	+
<i>cache in the RAID system</i>	0	++

## Other Effects:

- *different system monitoring / direct administrator responsibility (!)*
- *ease of administration*
- *availability trap is loosing importance*
- *built-in separation of load profiles*
- *often performance gains / better parallelism*
- *enormous cost savings*

# About the Effects of Virtual Machines

*A different pattern of data streams / different priorities*



OSL Gesellschaft für offene Systemlösungen mbH

[www.osl.eu](http://www.osl.eu)

# About the Effects of Virtual Machines

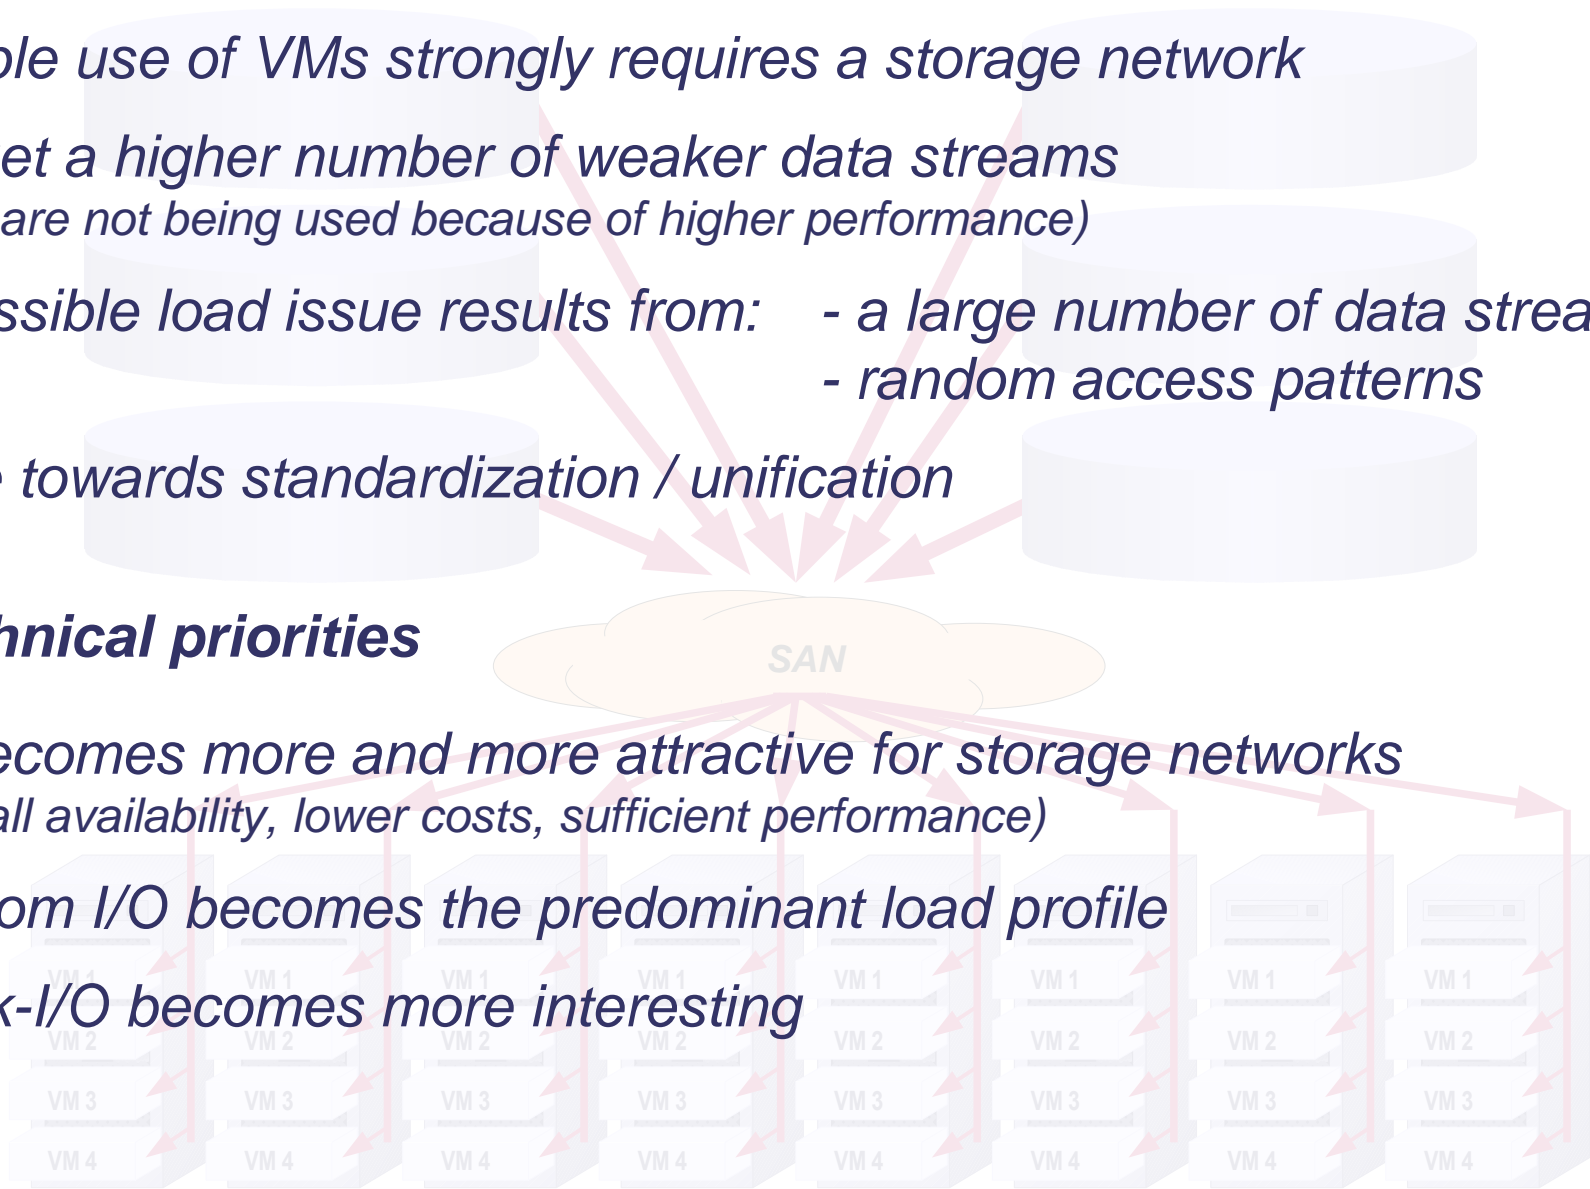
*A different pattern of data streams / different priorities*



- *flexible use of VMs strongly requires a storage network*
- *we get a higher number of weaker data streams  
(VMs are not being used because of higher performance)*
- *a possible load issue results from:*
  - *a large number of data streams*
  - *random access patterns*
- *drive towards standardization / unification*

## New technical priorities

- *IP becomes more and more attractive for storage networks  
(overall availability, lower costs, sufficient performance)*
- *random I/O becomes the predominant load profile*
- *block-I/O becomes more interesting*

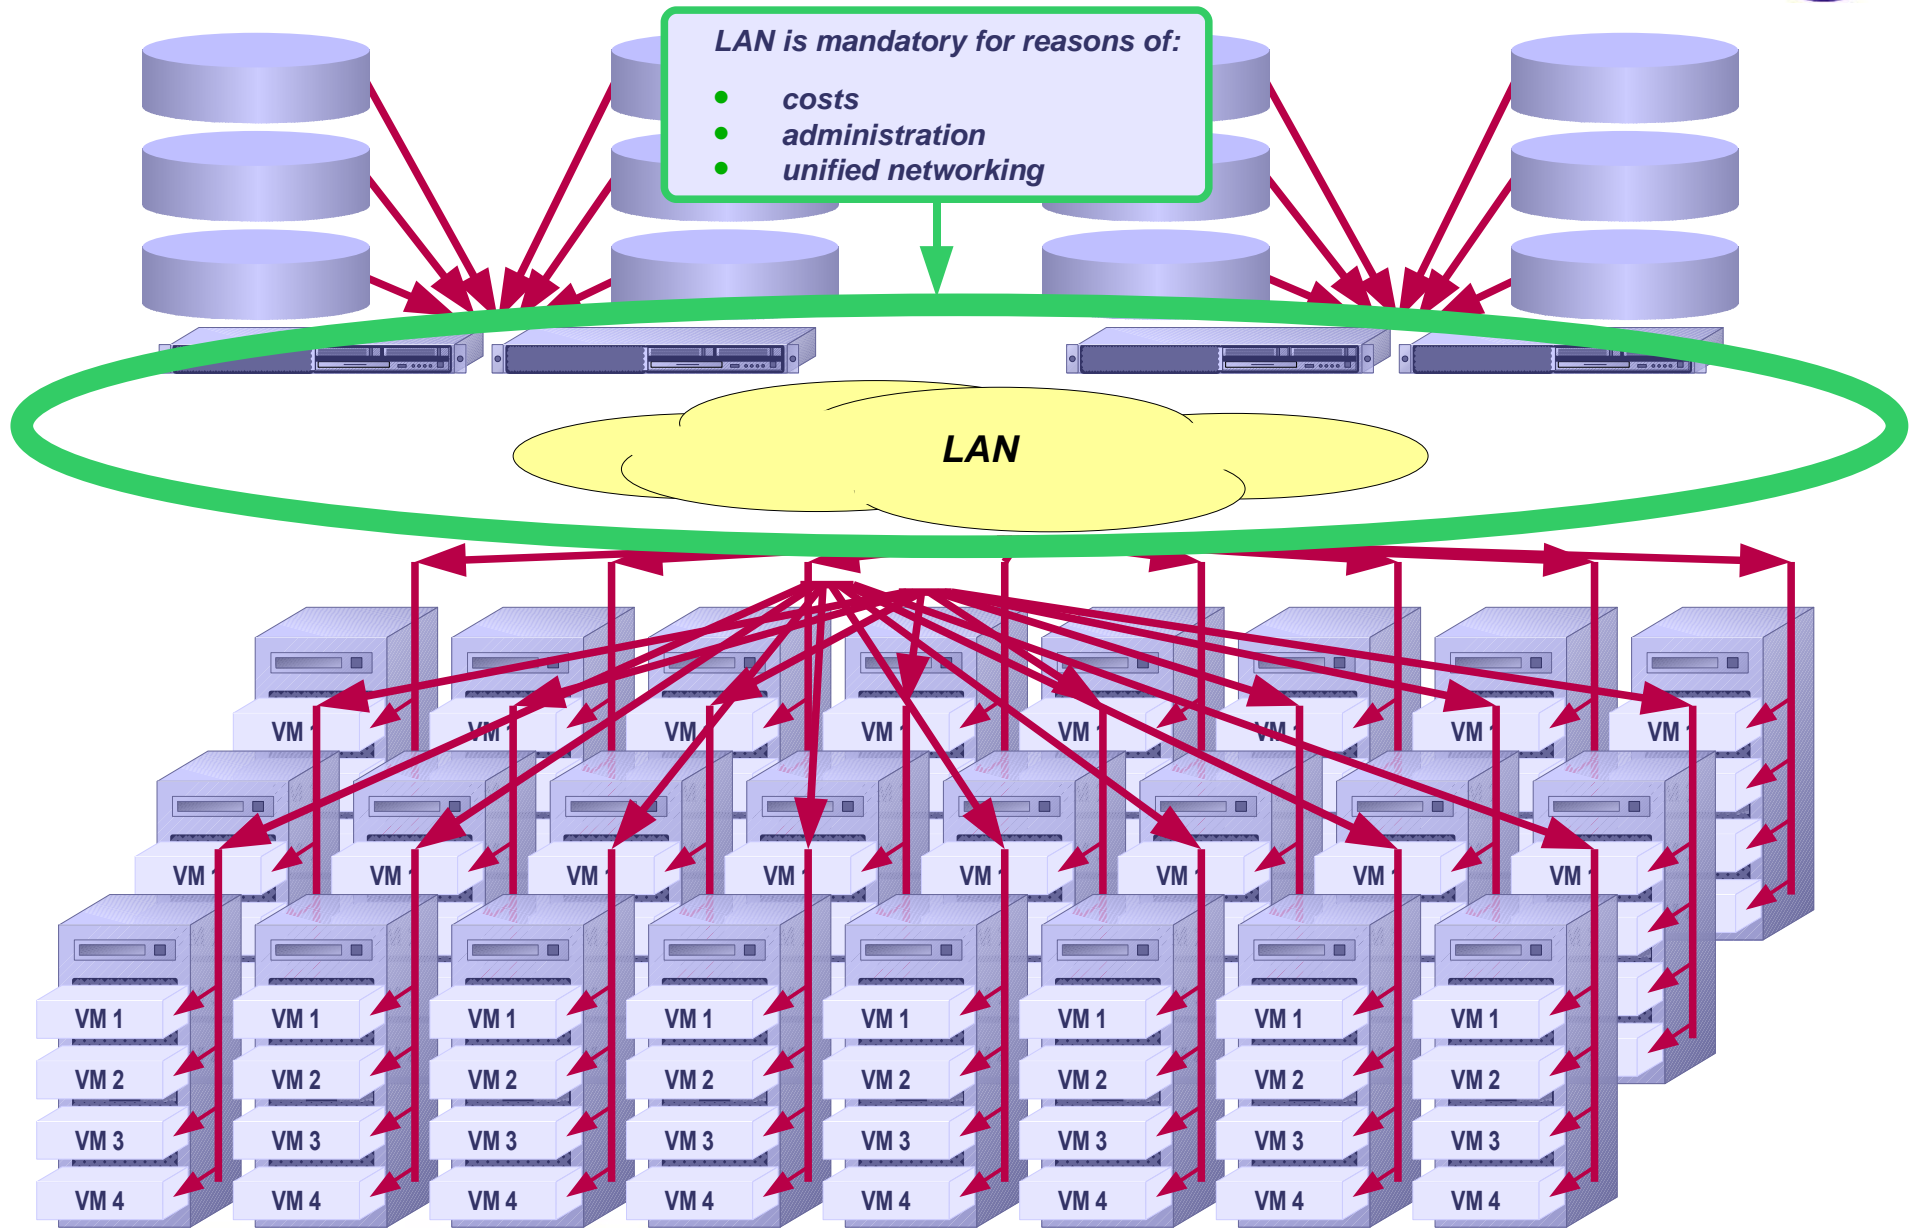


OSL Gesellschaft für offene Systemlösungen mbH  
[www.osl.eu](http://www.osl.eu)



# The Escalation: Cloud Infrastructures

*At least cloud infrastructures bring about a new quality*



OSL Gesellschaft für offene Systemlösungen mbH  
[www.osl.eu](http://www.osl.eu)

# ***“One For All” or: “It Could Be Done So Easily ...”***


## ***Advantages and Disadvantages of Unified Storage and Highly Integrated Systems***

# Unified Storage

*The solution to all problems?*



- *the promise:*
  - NFS / CIFS / block-I/O all from a single machine
  - integration of different storage platforms
  - solution of all storage problems
- *the method:*
  - filesystem-based data organisation
  - integrated RAID features
  - cute details (over-provisioning, snapshots ...)

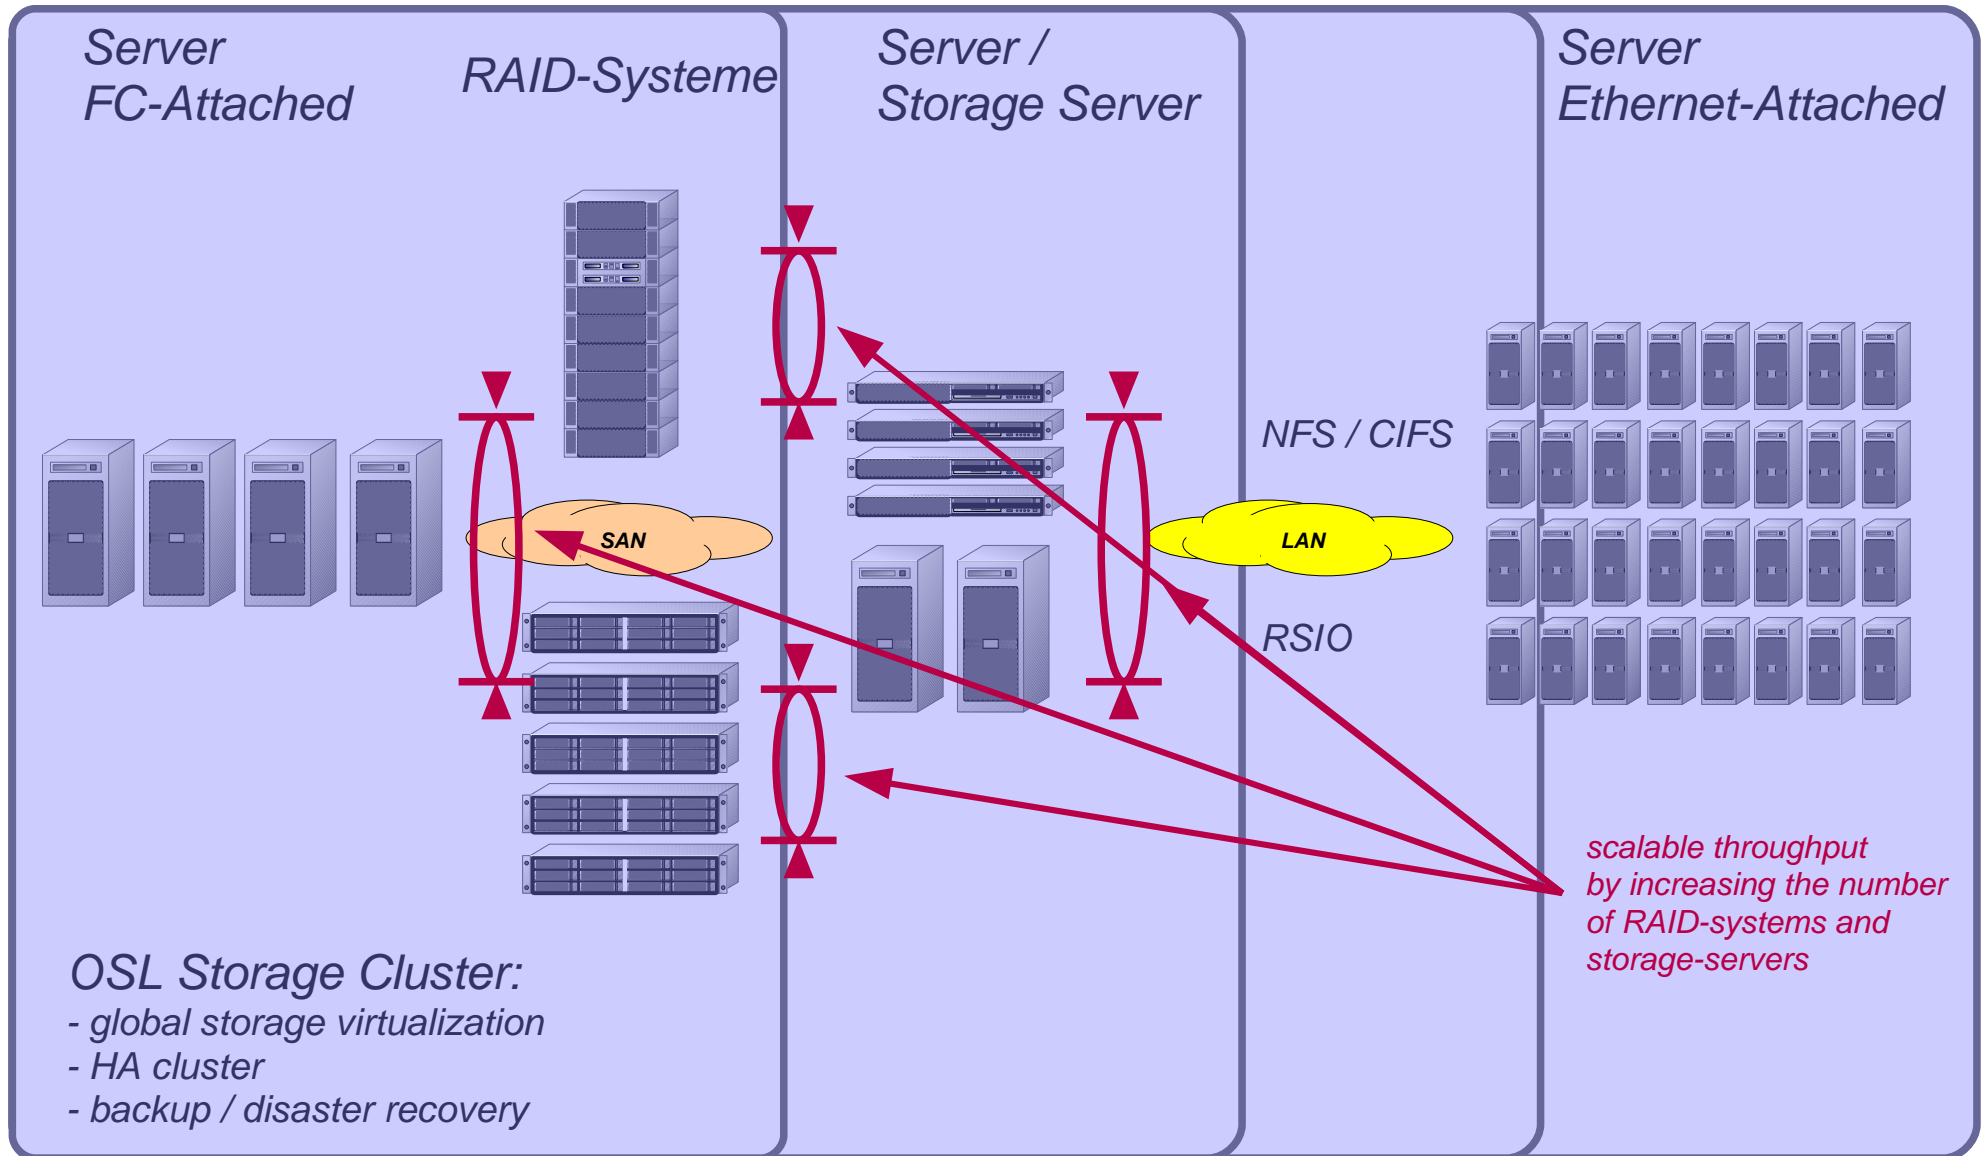


- *the reality:*
  - often excellent results for NFS / CIFS
  - weaker performance with block-I/O (filesystem layout!)
  - enormous interdependencies of different client load profiles
  - complex skill requirements for the administrator
  - poor scalability
  - almost no deterministic response times in block-I/O
  - in most cases mind-blowing prices, loss of liberty

OSL Gesellschaft für offene Systemlösungen mbH  
**www.osl.eu**

# Our alternative draft

Get more by use of standard building blocks



OSL Gesellschaft für offene Systemlösungen mbH  
**www.osl.eu**



# ***Our alternative draft***

*Where are the differences*



- *Unified Integrated Storage is not a product but a concept*
- *the concept is supported by host-based software*
- *use of standard components*  
*(good performance, low price, no vendor lock-in)*
- *almost at-will scalability*
- *maximum throughput only limited by the network*
- *high-performance block I/O can be run over Ethernet*
- *arbitrary isolation of load profiles*
- *special functions, e. g. bandwidth control*
- *RAID administration is reduced to simple operations*  
*(e.g. disk replacement)*
- *more advanced functions come in a unified form from the software*  
*(e. g. backups / DR / system copies)*

# ***Finding the Right Perspective ...***

***On the Eve of New Storage Concepts***

# The Influence of New Storage Technologies

The transition from HD to SSD is just the beginning



	<b>HD</b>	<b>SSD</b>	<b>FeRAM</b>	<b>DDR SDRAM</b>	<b>MRAM</b>
<i>Capacity</i>	<i>&gt; 2 TB</i>	<i>&gt; 512 GB</i>	<i>16 MB (Modul)</i>	<i>&gt; 4 GB (Modul)</i>	<i>Universal RAM ?</i>
<i>Access</i>	<i>4 ms</i>	<i>0,2 ms</i>	<i>0,05 µs</i>	<i>10 ns</i>	
<i>Transfer</i>	<i>140 MB/s</i>	<i>280 MB/s</i>	<i>1,6 GB/s (DDR2)</i>	<i>40 GB/s</i>	

# The Influence of New Storage Technologies

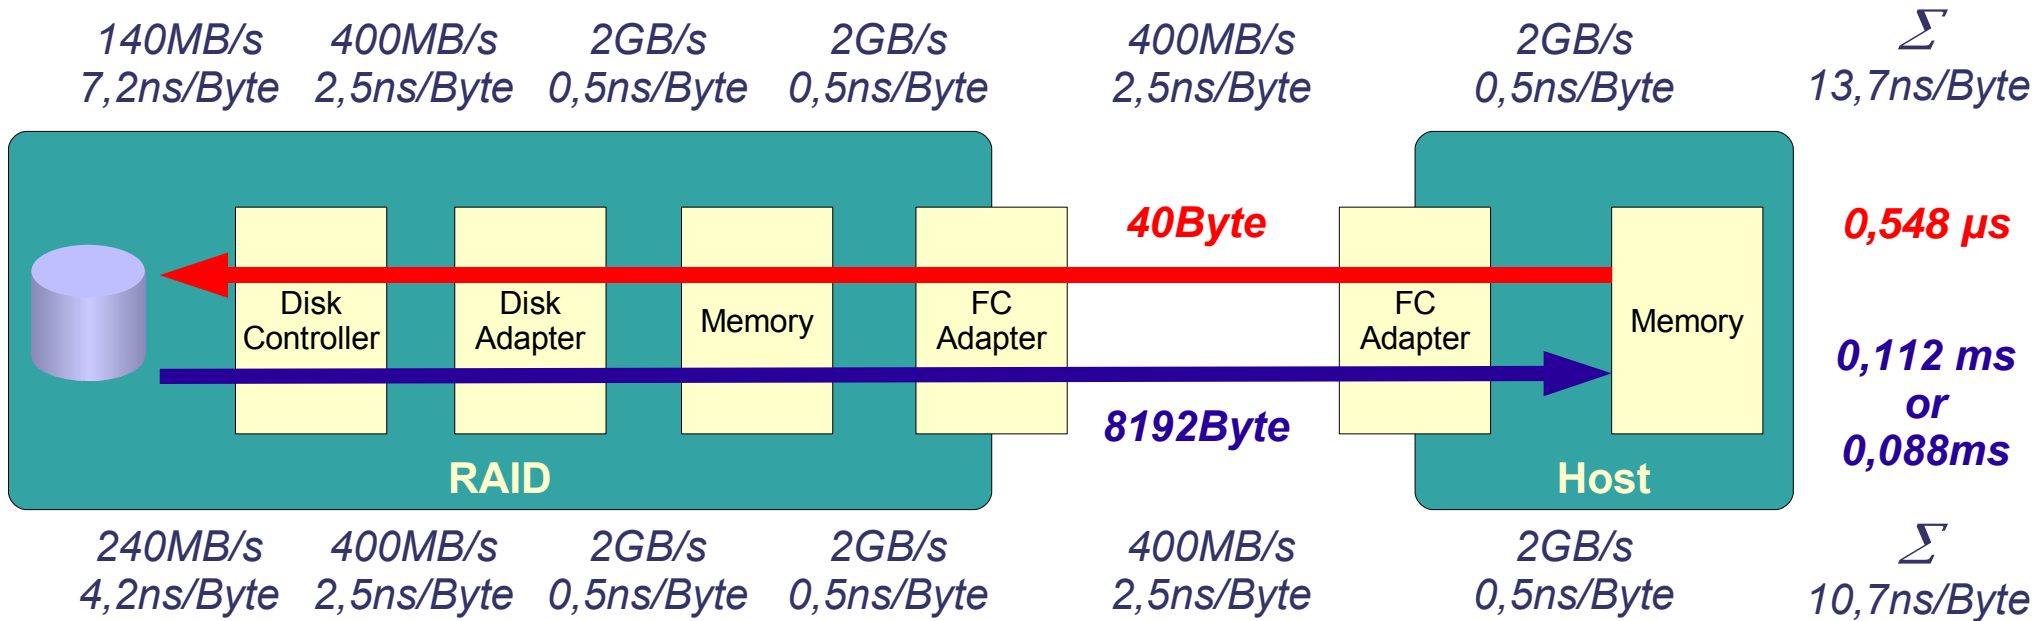
The transition from HD to SSD is just the beginning



**Calculation example: What can be gained by the transition from HD to SSD?\***

\*simplified model for sequential IO, access time not considered

	HD	SSD	FeRAM	DDR SDRAM	MRAM
Capacity	1 TB	32 GB	1 MB (module)	4 GB (module)	Universal RAM ?
Access	4 ms	0,2 ms	0,05 $\mu$ s	10 ns	
Transfer	140 MB/s	280 MB/s	1,6 GB/s (DDR2)	40 GB/s	



at 140 MB/s  
at 280 MB/s

about 8900 transfers/s  
about 11350 transfers/s

ca. 70MB/s  
ca. 88MB/s



# The Influence of New Storage Technologies

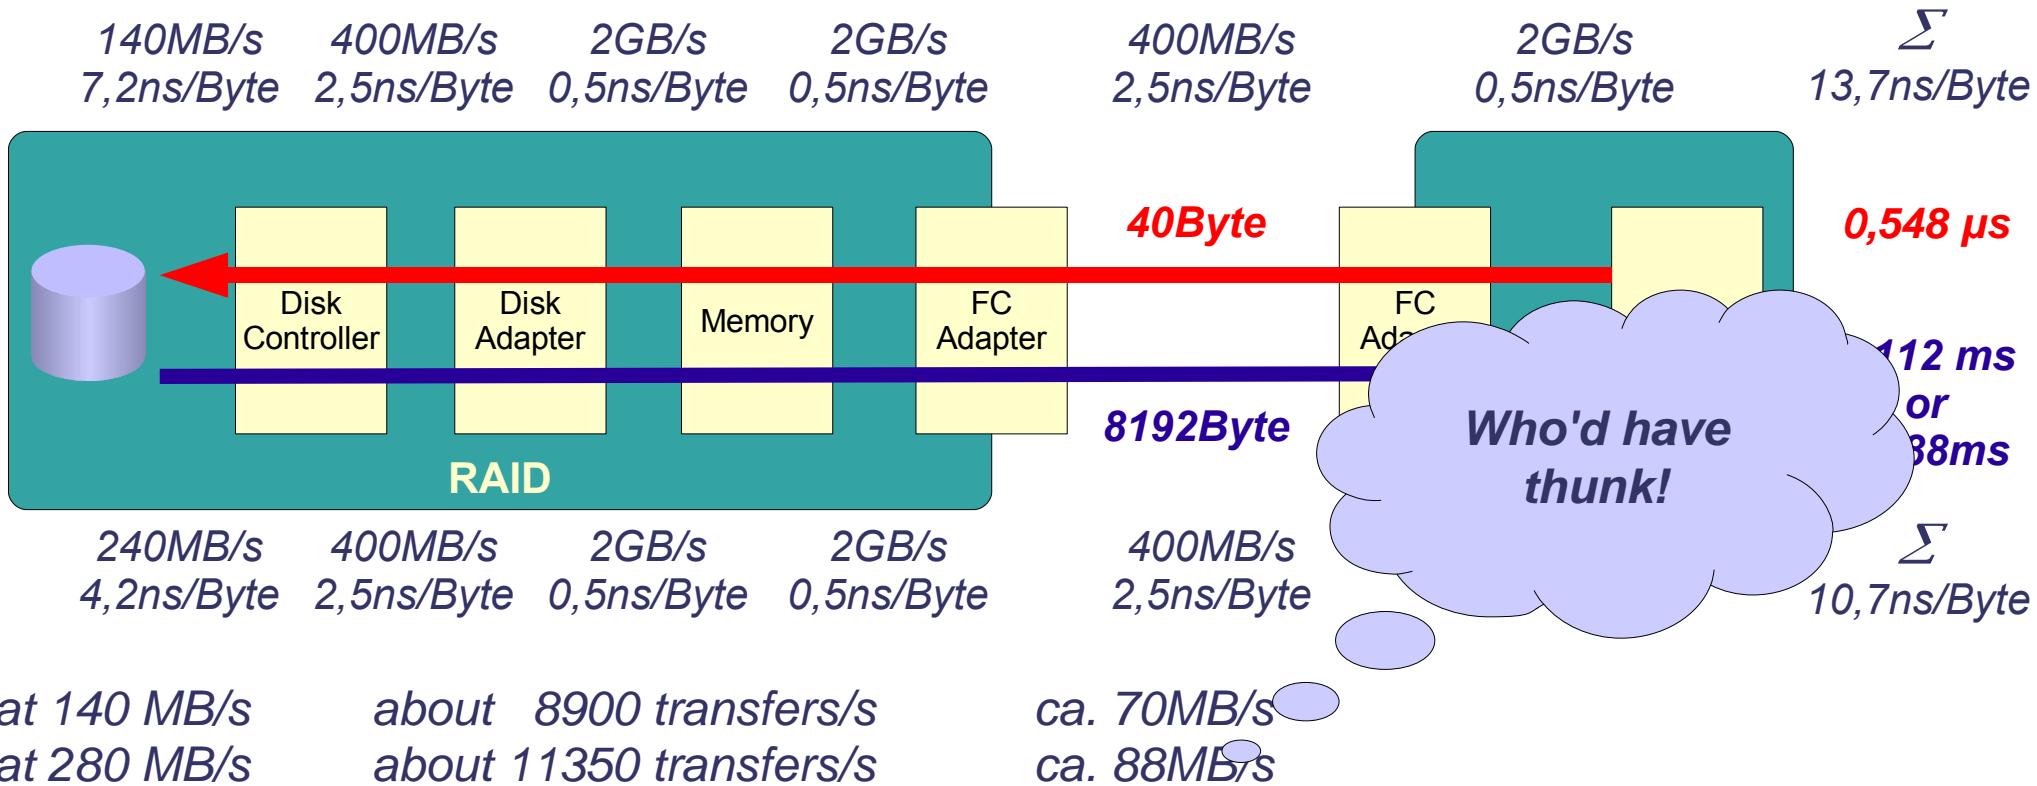
The transition from HD to SSD is just the beginning



Calculation example: What can be gained by the transition from HD to SSD?\*


\*simplified model for sequential IO, access time not considered

	HD	SSD	FeRAM	DDR SDRAM	MRAM
Capacity	1 TB	12 TB	1 MB (module)	4 GB (module)	Universal RAM ?
Access	4 ms	0,2 ms	0,05 $\mu$ s	10 ns	
Transfer	140 MB/s	280 MB/s	1,6 GB/s (DDR2)	40 GB/s	



# The Influence of New Storage Technologies

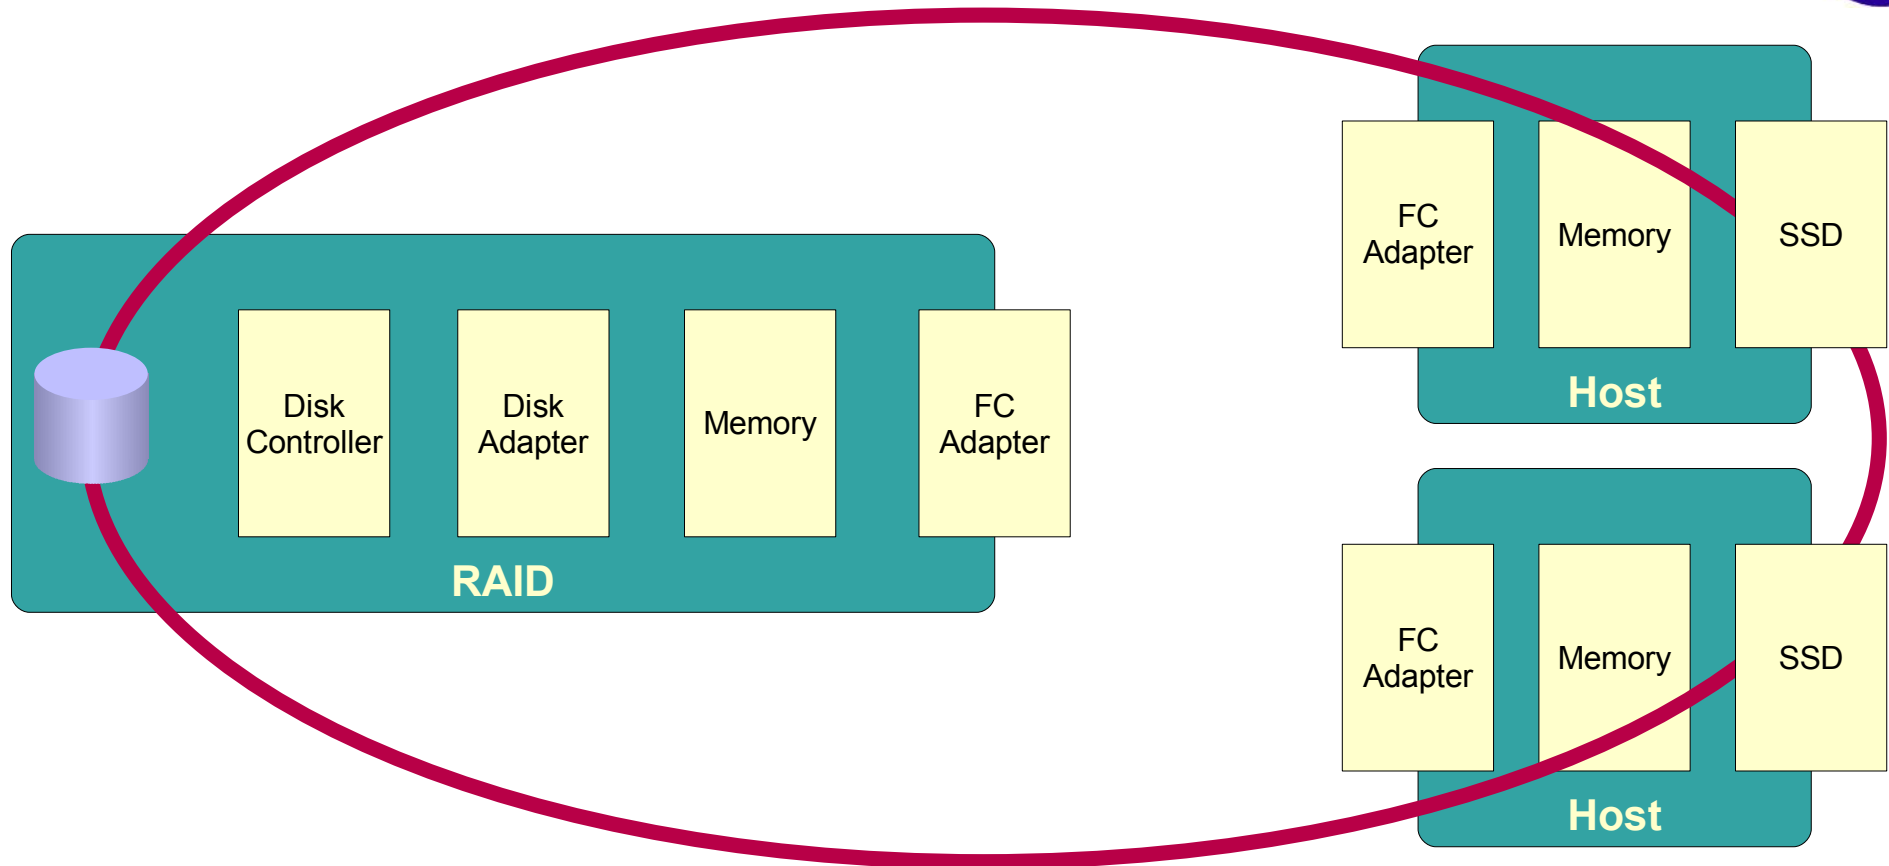
*What is making more sense?*



- *revolutionary new system designs cannot be seen for now*
- *faster serial interconnects?*
- *proprietary system architectures?*
- *combination of “faster and proprietary”?*

# The Influence of New Storage Technologies

What is getting shape today – what OSL is working at



- *only locally integrated mass storage can deliver break-through performance gains*
- *challenges:*
  - *integration with external storage*
  - *full control in clustered environments*
  - *make the performance gains useable in “real life” data centre processes*

OSL Gesellschaft für offene Systemlösungen mbH  
**www.osl.eu**

# *Effects of Modern Storage and Virtualisation Technologies - Summary*



- *Build modern architectures from standard components!*
- *Speed of the network is not enough!*
- *Be ready for hybrid architectures!*
- *Focus on concepts – not on hardware!*