Education

Storage Performance Management Overview

Brett Allison, IntelliMagic, Inc.

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

- Storage Performance Management (SPM), Storage Architecture and SMI-S
  - This session will appeal to Storage Managers, Performance and Capacity Managers, and those that are seeking a fundamental understanding of storage performance management. This session includes an overview of the processes, technology and skills required to implement SPM, as well as an overview of disk storage system architecture, and the SMI-S specification as it relates to block level performance. The focus is on block level storage systems.

Storage Performance Management (SPM)

How Large is Storage Expenditure?

- Where will you be in 2012?
- Reduce storage costs by implementing SPM

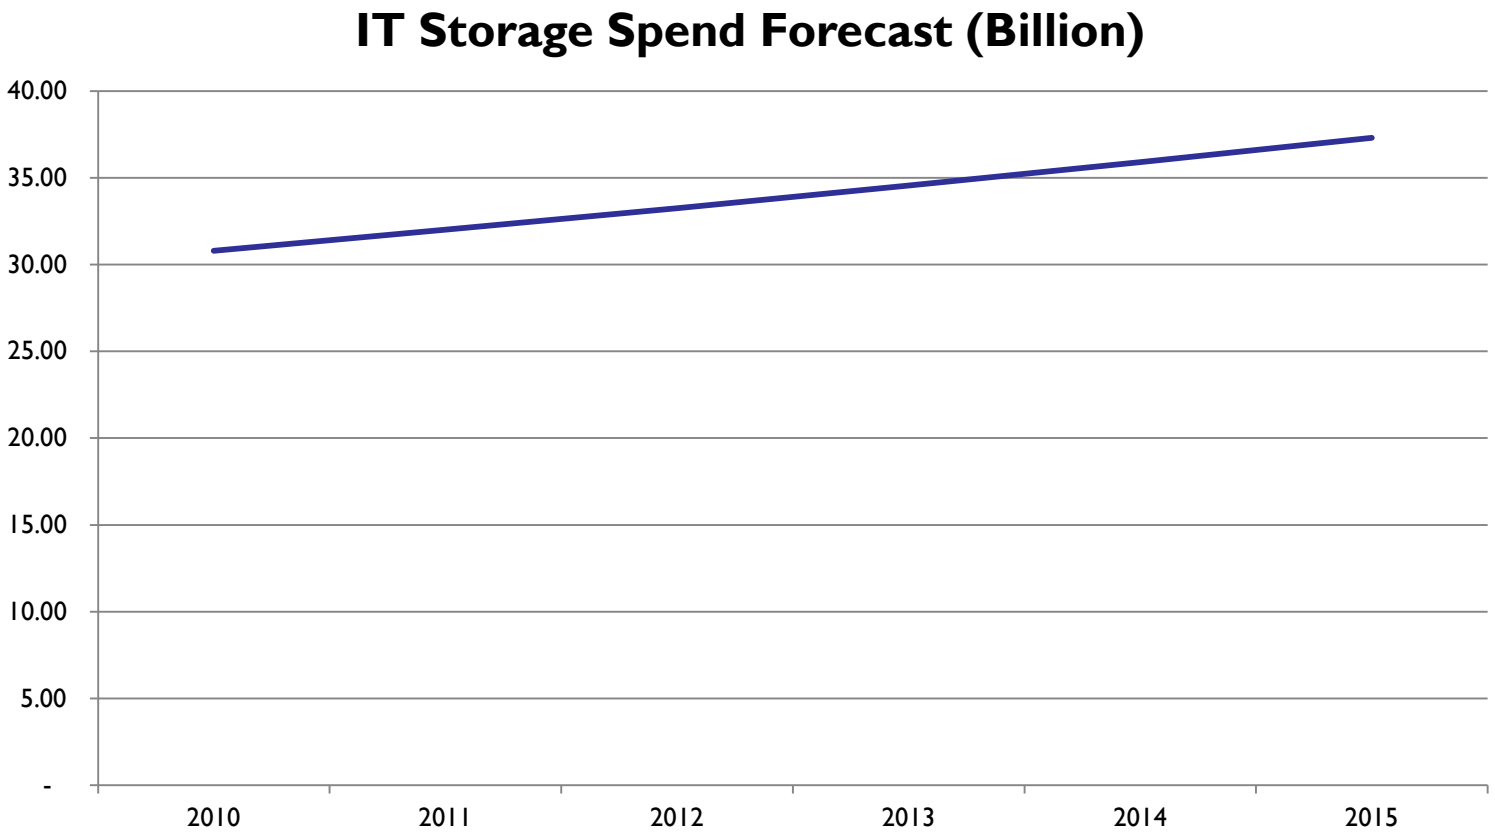

Source: [1]

SPM Defined

- SPM is the process of ensuring
  - that users constantly receive required I/O service levels to avoid performance problems;
  - that storage assets are efficiently configured and used to avoid over spending on hardware.
- Risk avoidance is more important than saving money.

Logical Volume Management (SRM)

Physical Disk Choices

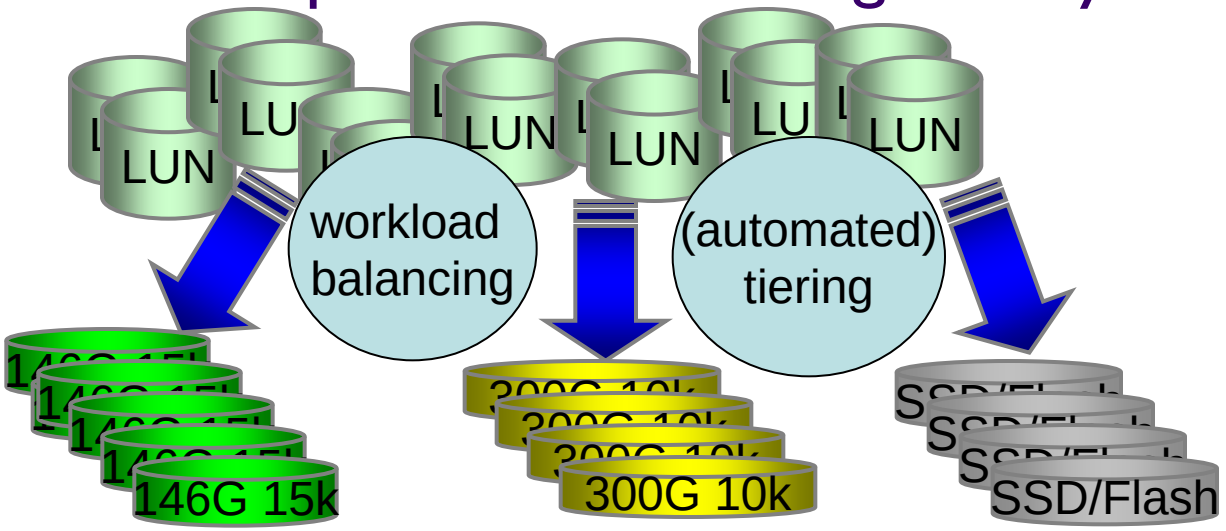

Risks

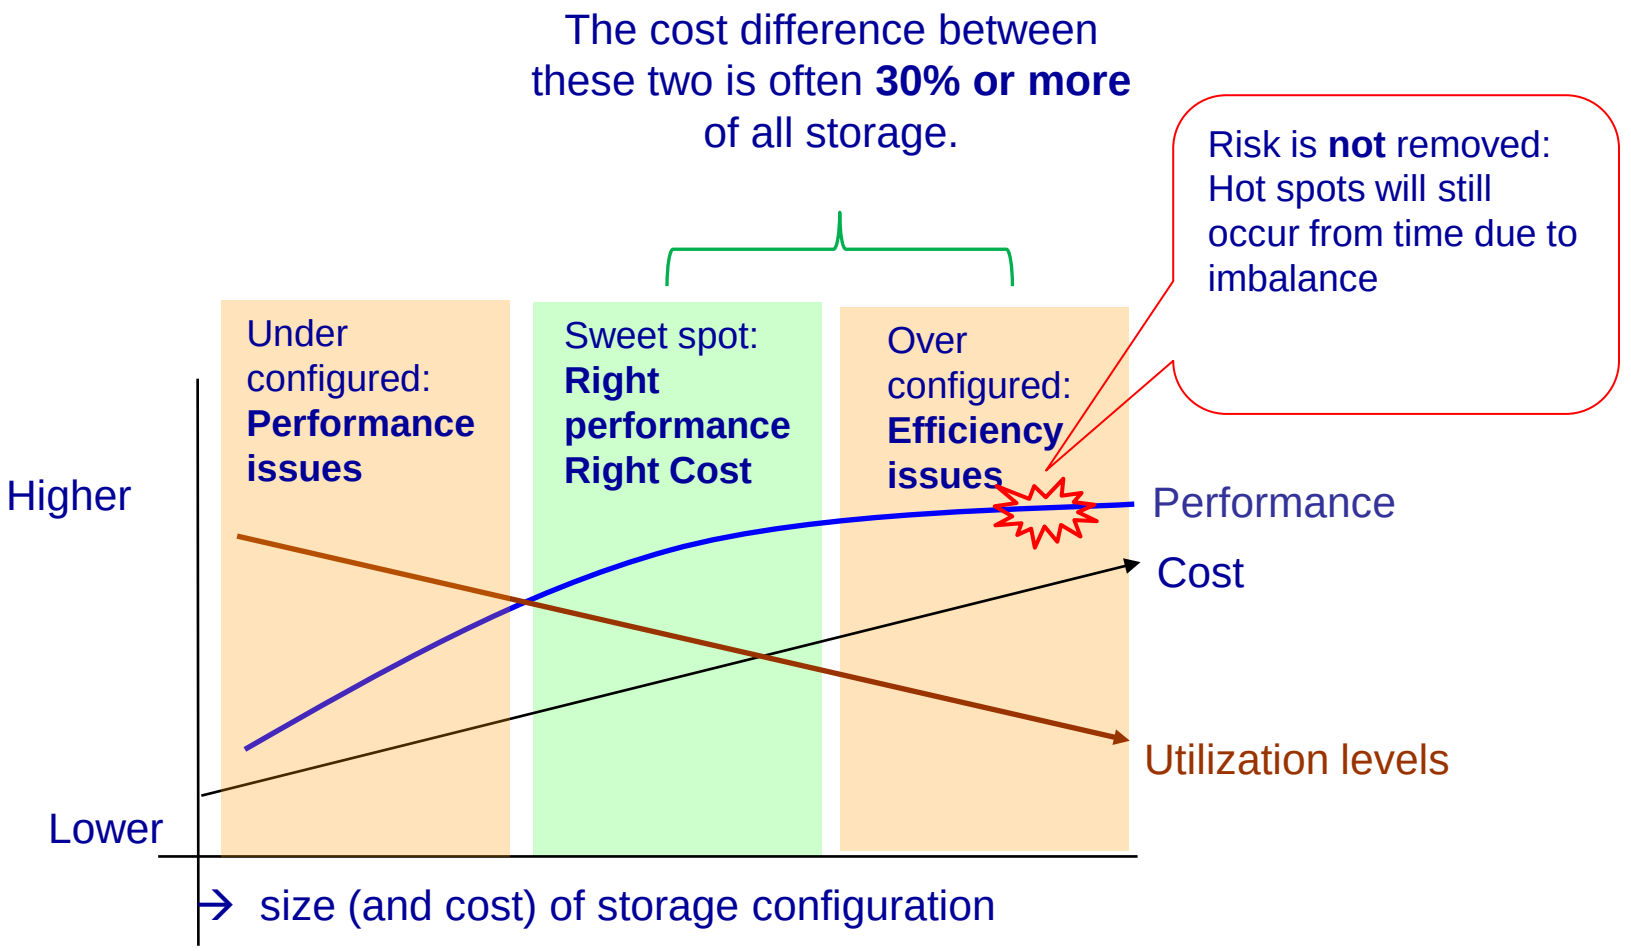

The Four SPM Primary Processes

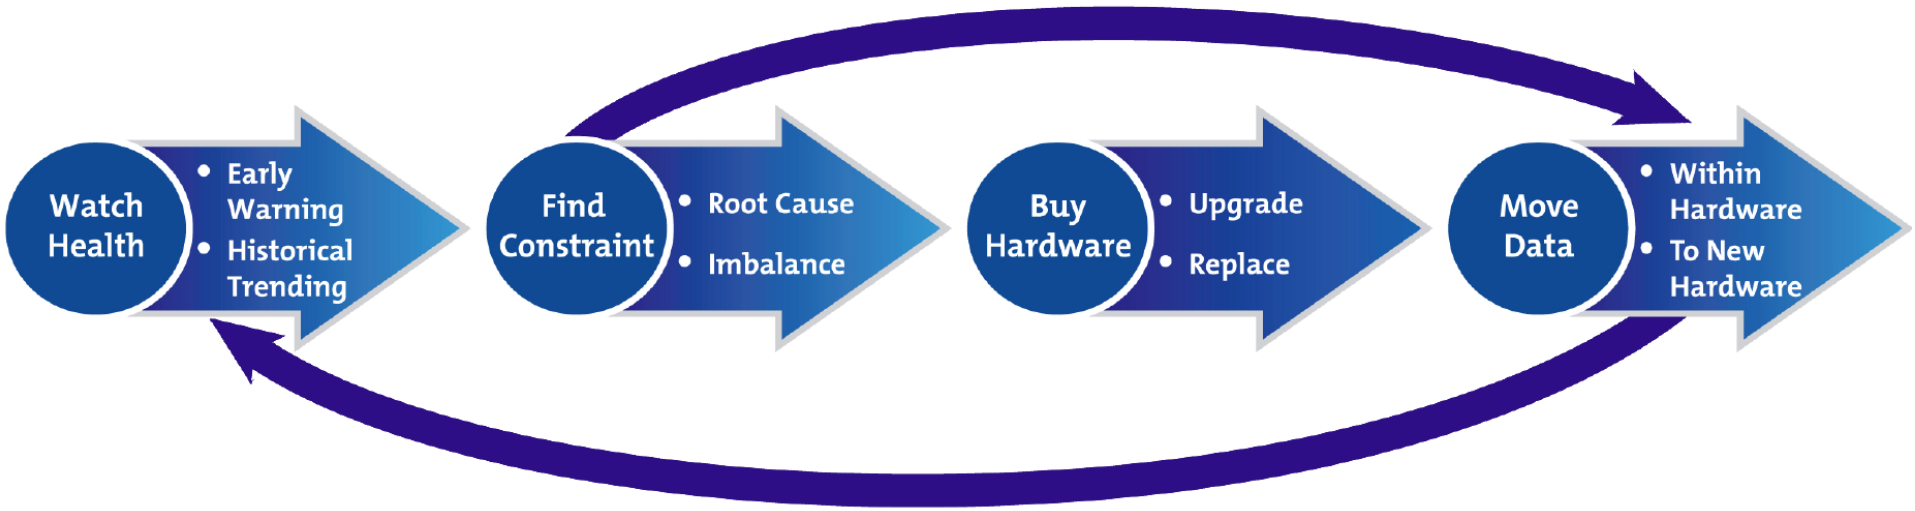

SPM Maturity Stages

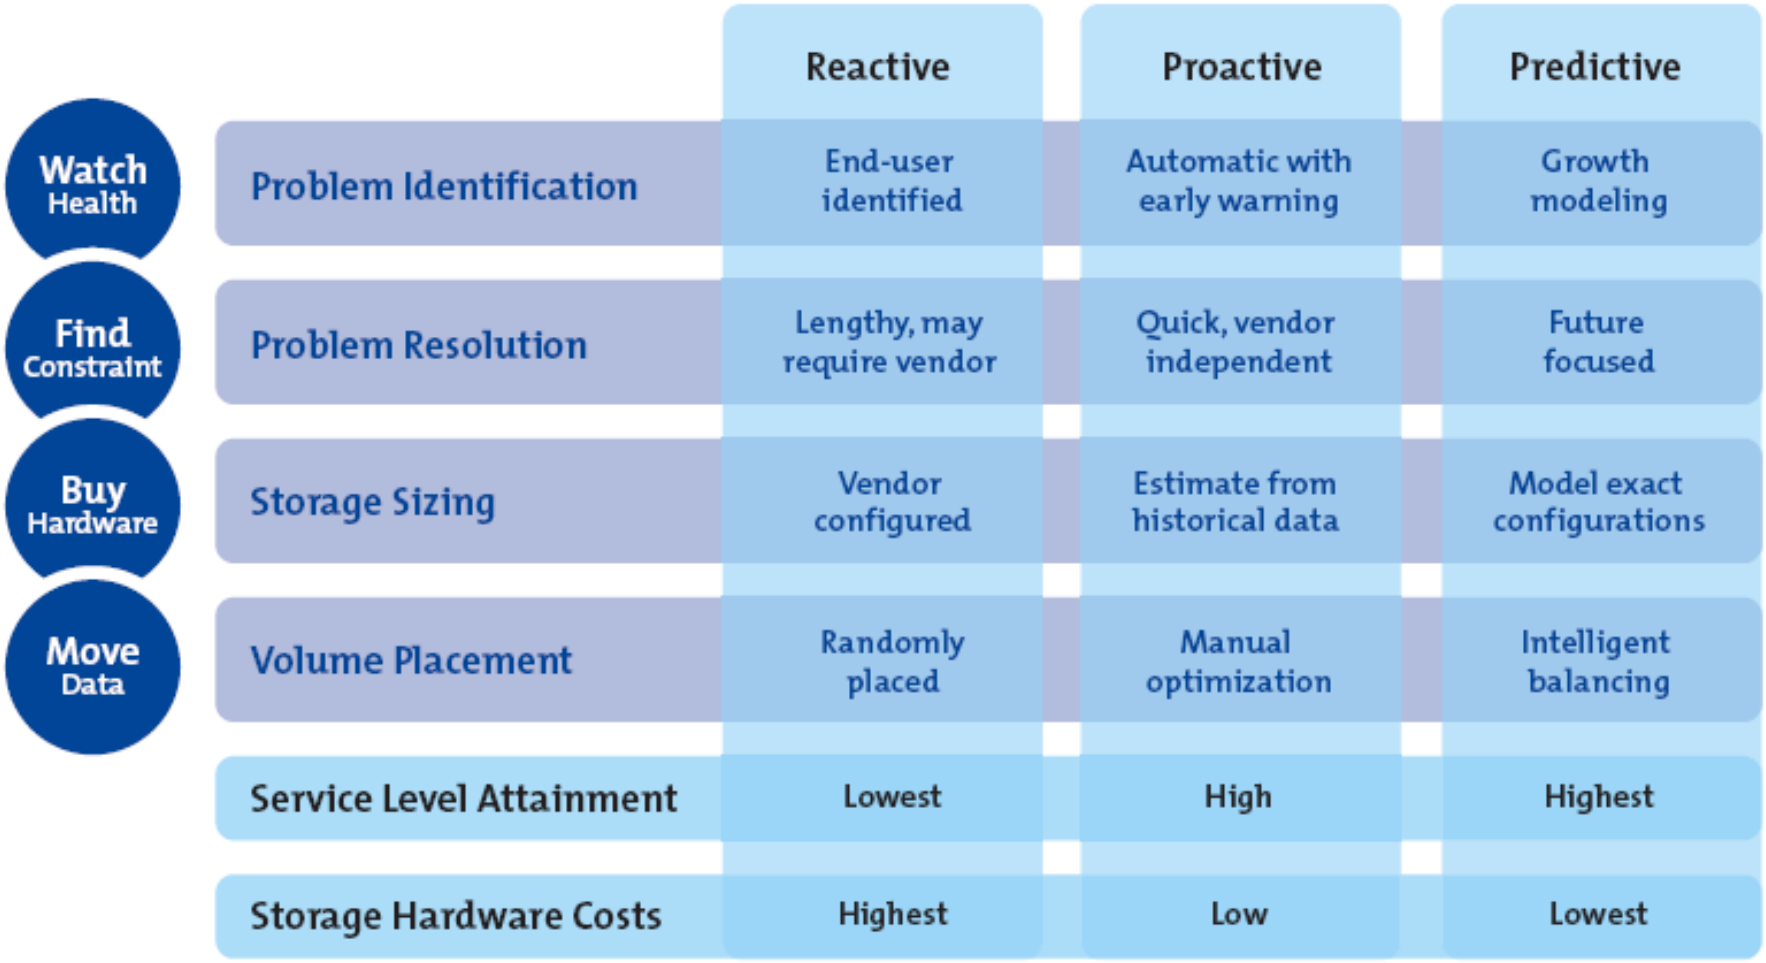

SMI-S

SMI-S

SMI-S is a vendor independent protocol to manage storage subsystems via a HTTP/HTTPS based protocol.

- Common standard across all vendors
- Provide both topology and performance information
- Standard scope includes storage, switches, and servers
- Performance measurement is commonly supported

http://www.snia.org/tech_activities/standards/curr_standards/smi

SMI-S Implementation

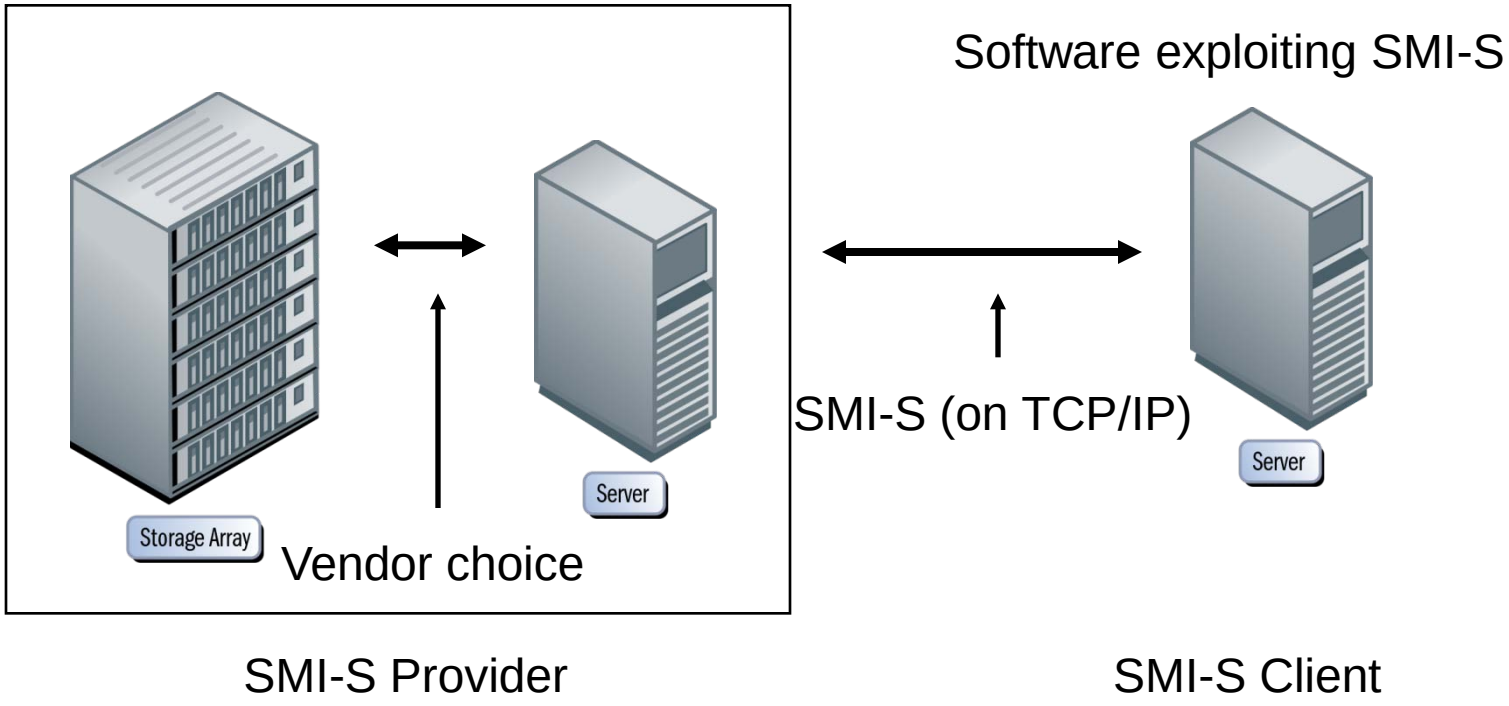

What ElementTypes Does SMI Define for Performance?

ElementType	Component	Vendor Specific
ElementType2	Cumulative statistics for the storage system	
ElementType3	Front-end Controllers	
ElementType4	Peer Storage System (Mirroring)	✓
ElementType5	Back-end Controllers	
ElementType6	Front-end FC ports	
ElementType7	Back-end Ports	
ElementType8	Volumes	
ElementType9	Extent – Intermediate storage	✓
ElementType10	Disk Drive	
ElementType11	Arbitrary Logical Units – Controller commands	✓
ElementType12	Remote Replica Group – Remote Mirror	✓

Statistics Summary by ElementType

Statistic Property	Top Level Computer System	Component Computer System (Front-end)	Component Computer System (Peer)	Component Computer System (Back-end)	Front-end Port	Back-end Port	Volume (LogicalDisk)	Composite Extent	Disk
Statistic Time	R	R	R	R	R	R	R	R	R
TotalIOs	R	R	R	R	R	R	R	R	R
Kbytes Transferred	R	O	O	O	R	O	R	R	R
IOTimeCounter	O	O	O	O	O	O	O	N	O
ReadIOs	O	R	R	N	N	N	R	N	R
ReadHitIOs	O	R	R	N	N	N	R	N	N
ReadIOTimeCounter	O	O	O	N	N	N	R	N	N
ReadHitIOTimeCounter	O	O	O	N	N	N	O	N	N
Kbytes Read	O	O	O	O	N	N	O	N	O
Write Ios	O	R	R	N	N	N	R	N	O
WriteHitIOs	O	R	R	N	N	N	R	N	N
WriteIOTimeCounter	O	O	O	N	N	N	O	N	O
WriteHitIOTimeCounter	O	O	O	N	N	N	O	N	N
KbytesWritten	O	O	O	O	N	N	O	N	O
IdleTimeCounter	N	N	N	O	O	N	O	O	O
MaintOp	N	N	N	N	N	N	N	O	O
MaintTimeCounter	N	N	N	N	N	N	N	O	O

R - Required

O - Optional

N – Not Specified

Relationships – A simple example

- Configuration data is necessary to provide the relationships between the elements:
  - Which LUNs are defined on which extent pool
  - Which physical drives make up an array group
  - Which port (types) are connected to each (host adapter)

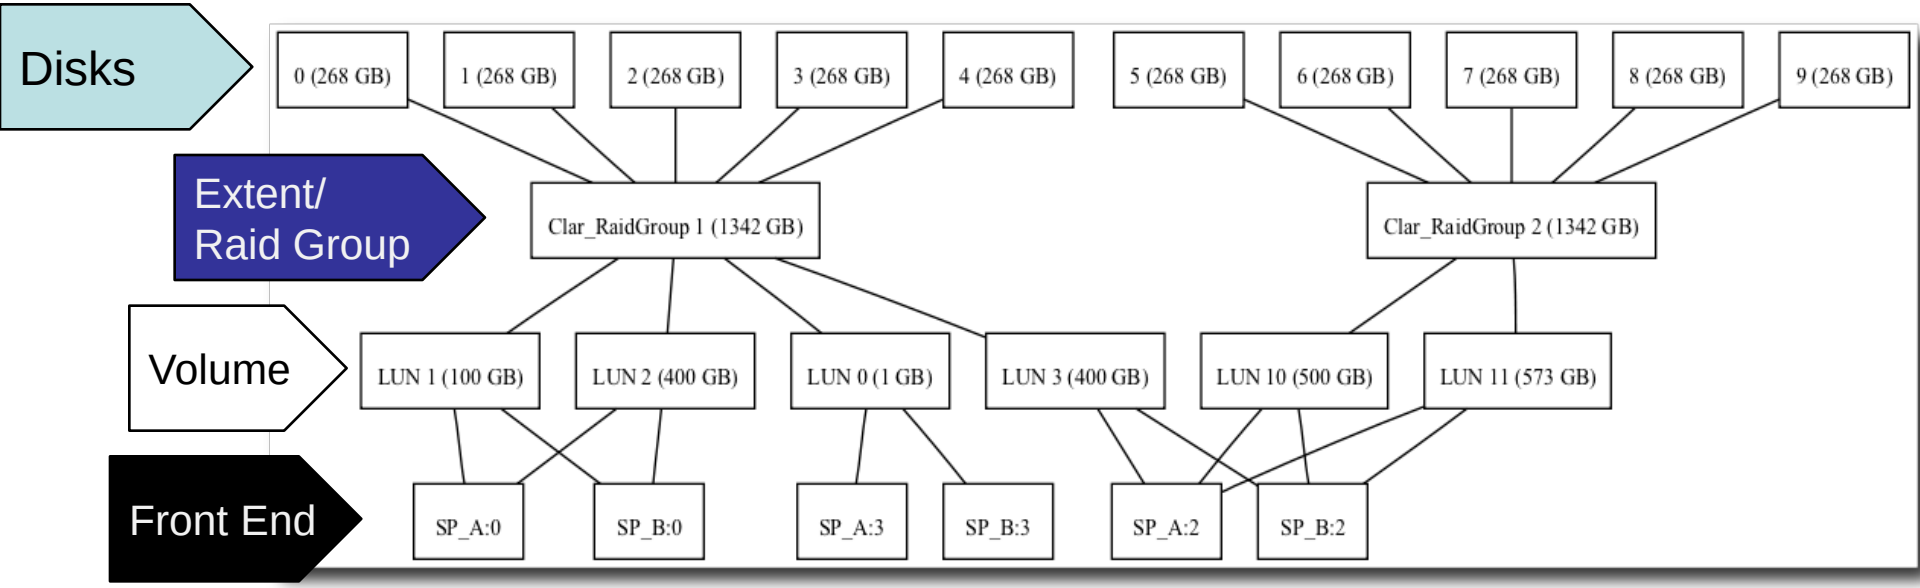

Storage System Architectures and Measurement with SMI-S

Architectural Trends Impacting Performance

- Spinning disks are getting bigger faster than they are getting faster!
- Requests per GB now exceeds access density capability of spinning drives for many workloads
- Trends driving density **increase** include
  - De-duplication
  - Thin provisioning
- SSD could remove disk as bottleneck in the future
- Commodity hardware

Disk Storage System (DSS) Architecture

- All vendors agree:
  - Frontend (host) and
  - Backend (disk) and
  - Cache and
  - Volumes are required
- Do the metrics provided tell us what we need to know?
  - SMI-S has a well defined model for performance metrics and relationships.
- High level storage system metrics are not enough!
  - We can only rely on these if the I/O is evenly spread across components

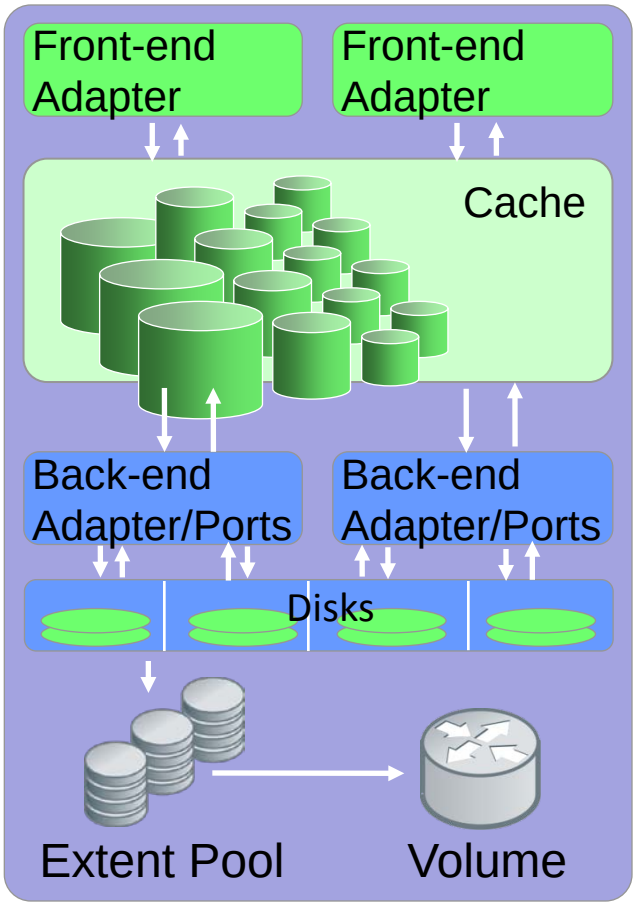

System Level Response Time and Throughput

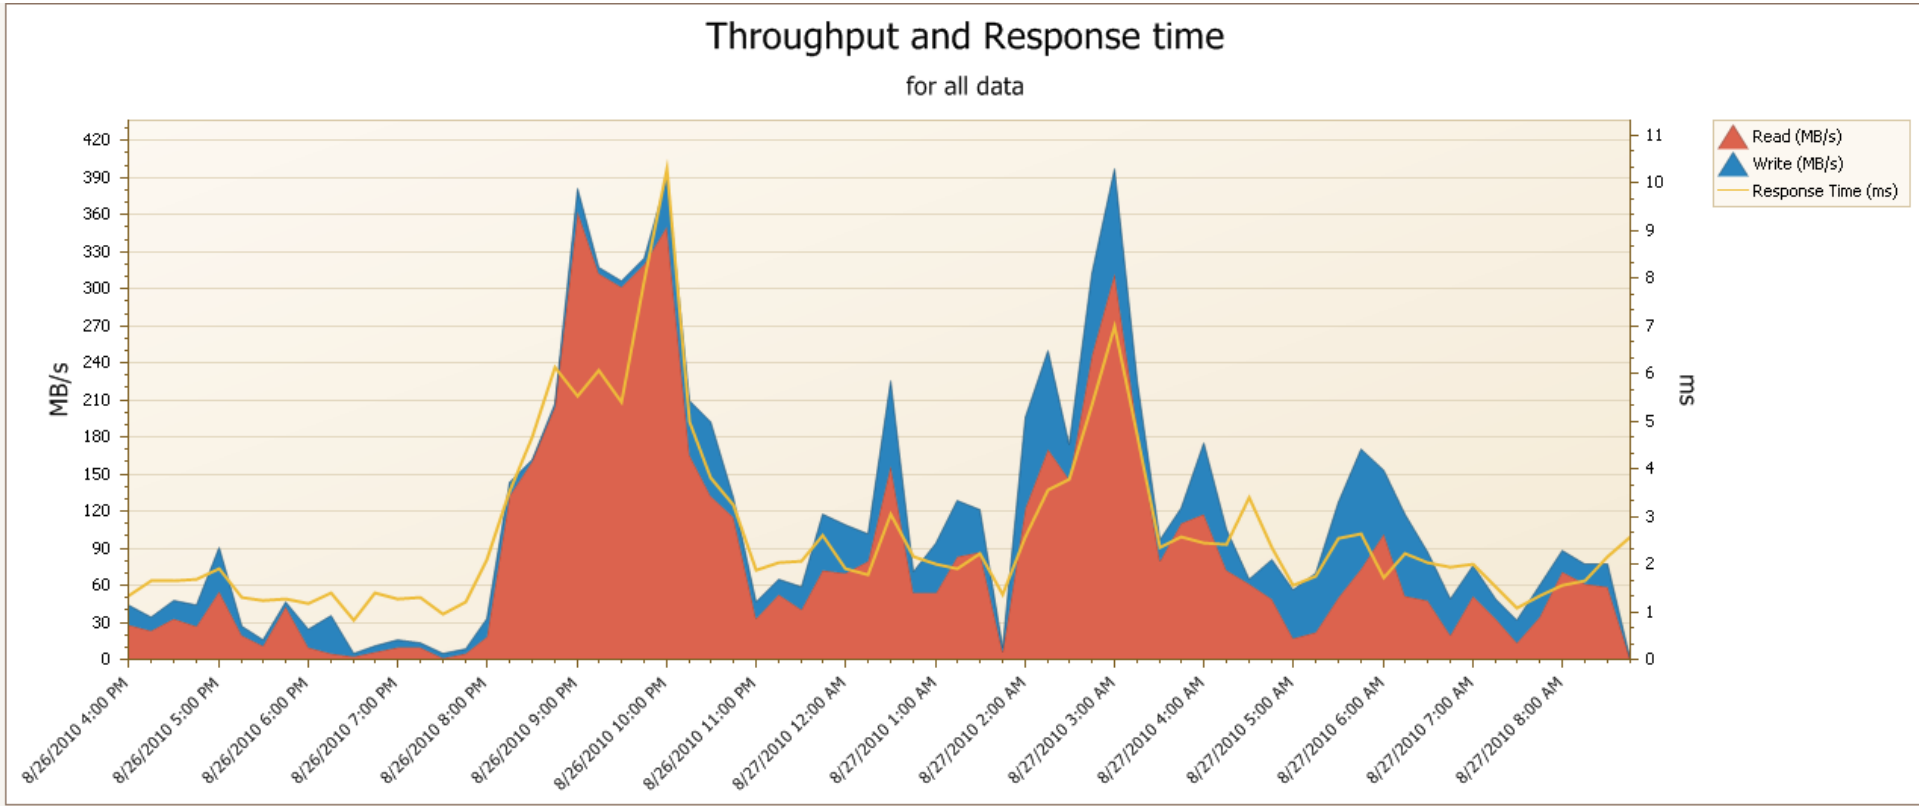

Processors and Cache

- Different implementations use different approaches
- All use cache to store
  - Recently used tracks and records
  - Recently written records
  - Pre-loaded tracks for sequential read
  - Some form of track descriptor tables to facilitate write operations without a disk access
  - Async copy information

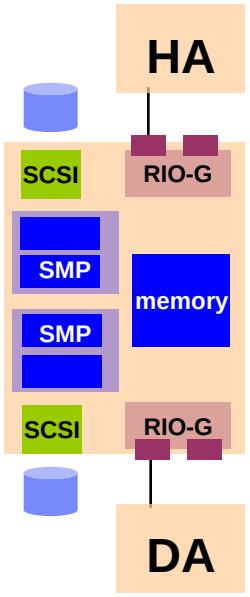

Midrange: centralized cache management

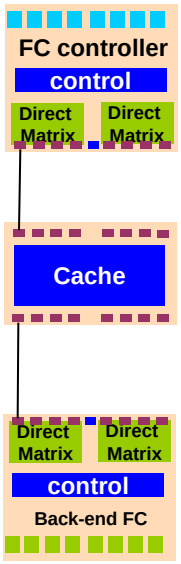

Cache shared between engines

Fixed cache assignment

Read Cache Hit % by DSS

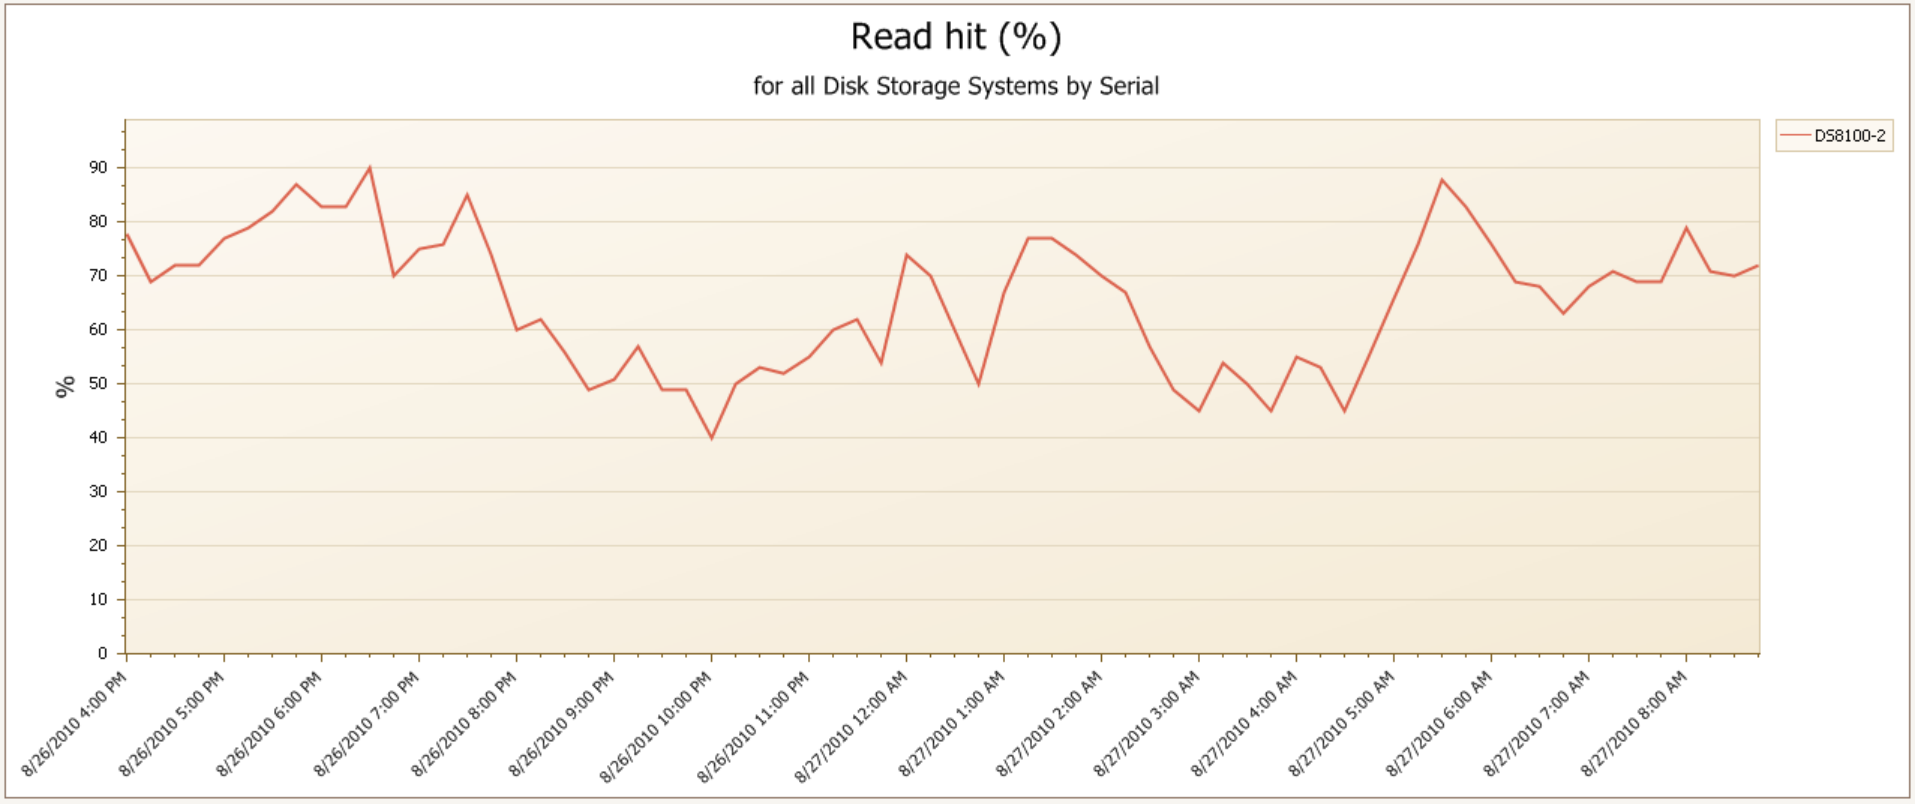

Front-end Adapter

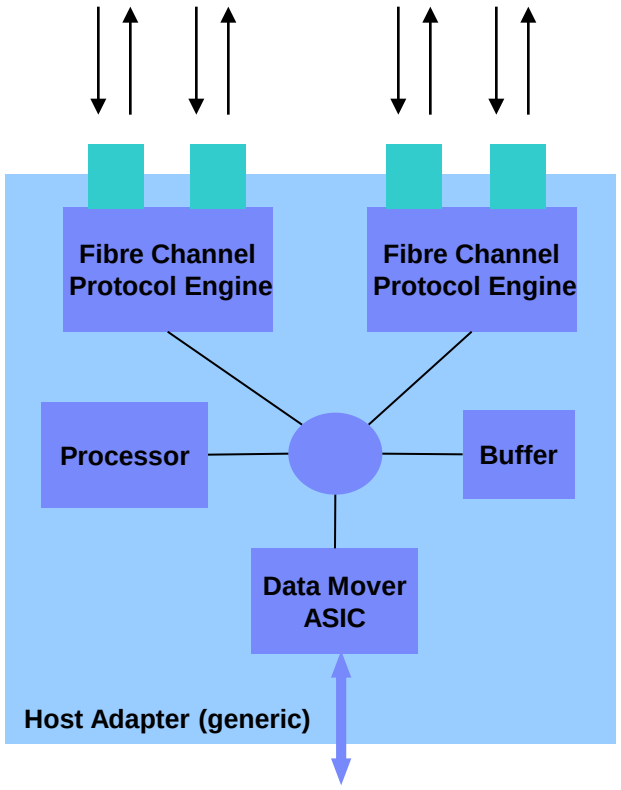

- Provides connectivity between disk subsystem and hosts
- Cards support ESCON, SCSI, FICON Fibre, SAS and/or iSCSI sometimes FICON and Fibre with one card
- Implementations differ greatly in maximum data handling capability, especially for FICON and Fibre
- Even though ports are rated as (e.g.) 8 Gbit/s, no implementation achieves this speed due to overhead.

Front-end Controller/Adapter Read Response Time

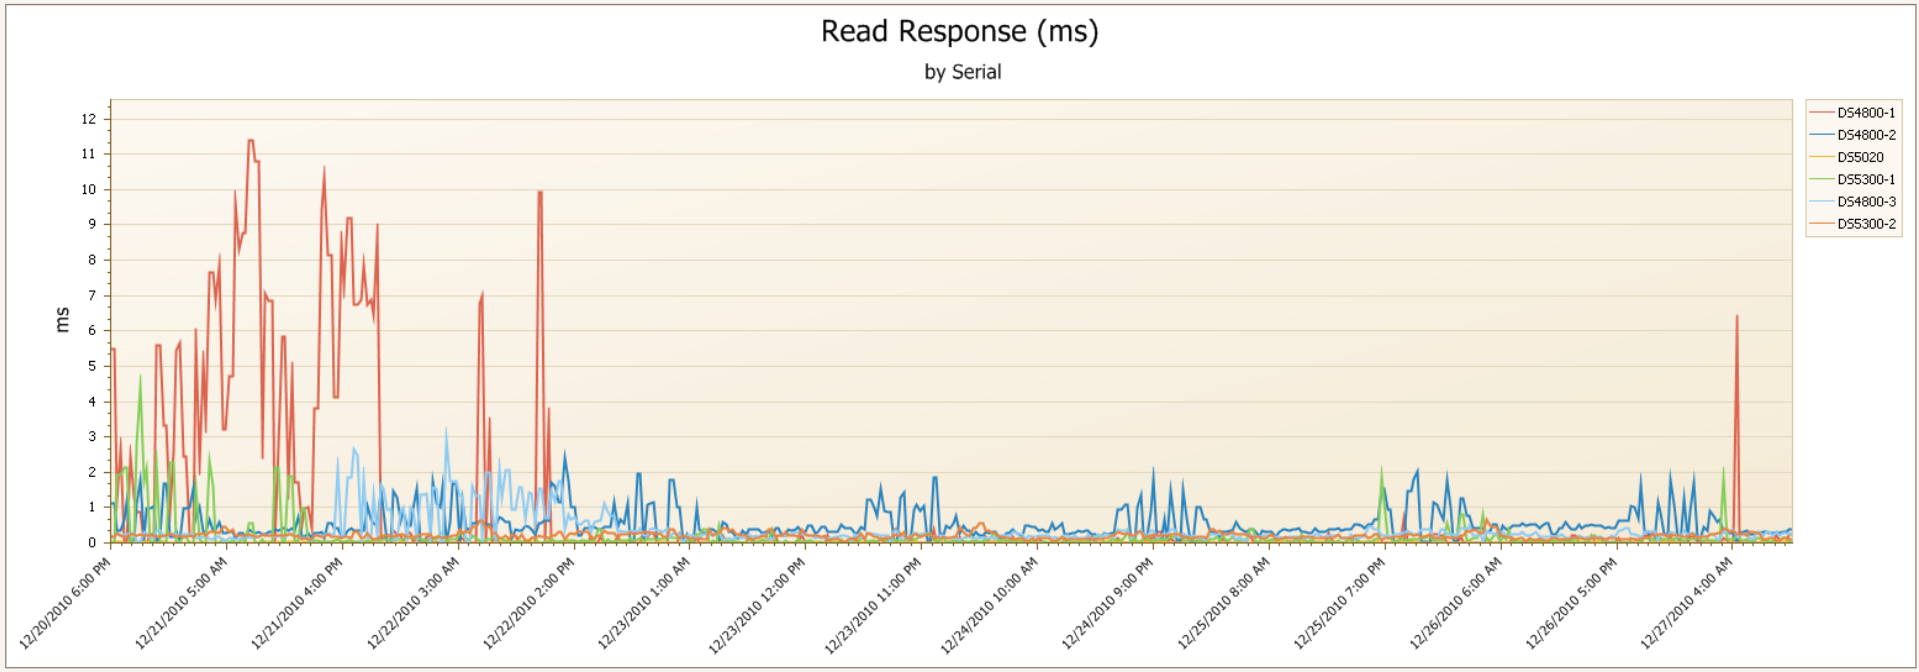

Front-end Ports

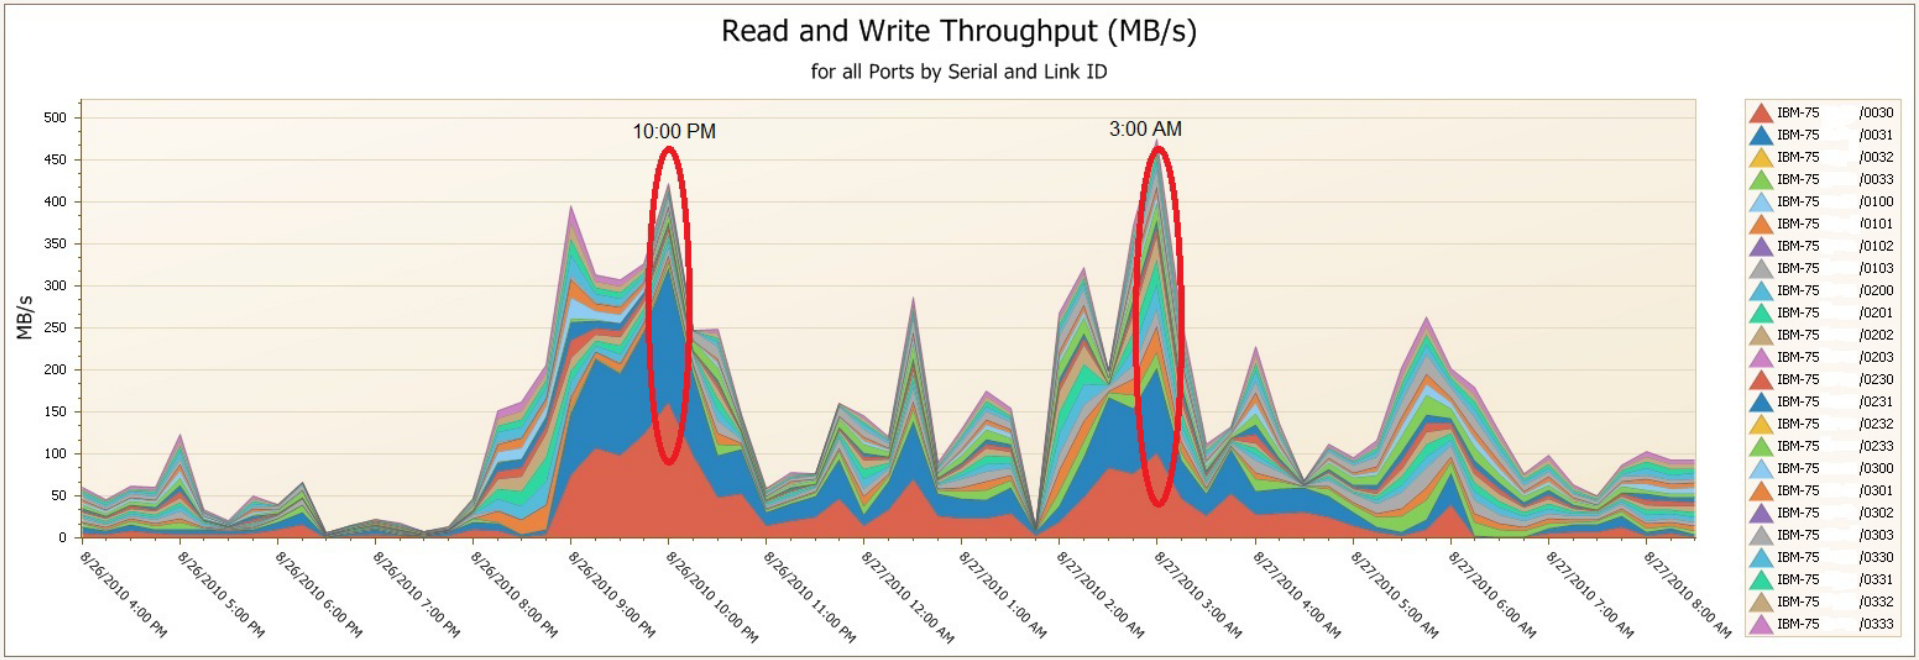

Device Adapters

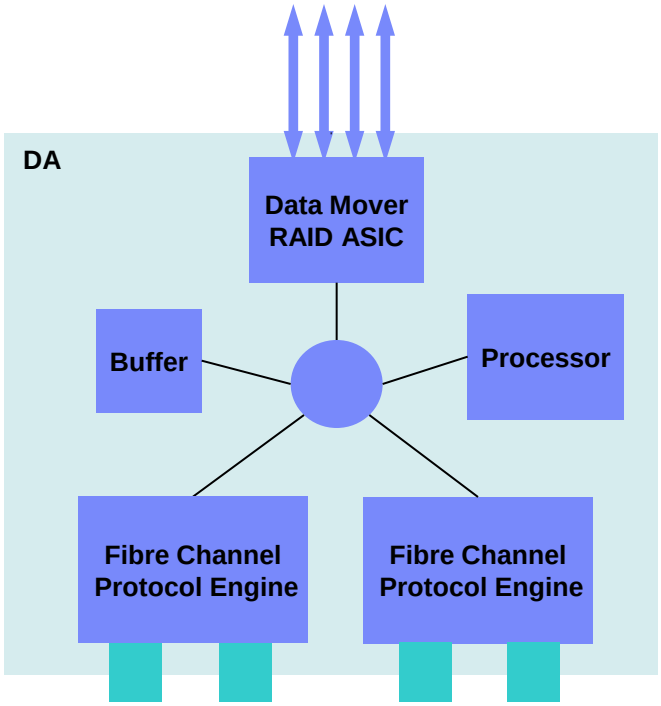

- Connect HDDs to internal disk system resources
- Manage RAID operations, sometimes using cache memory for RAID computations
- Configured in pairs to provide redundancy if one adapter fails
- HDD interfaces include various generations of SCSI, SSA, FC-AL, SATA and SSD
- FC-AL switched back-end are gradually being replaced by SAS back-ends

Backend Adapter Throughput

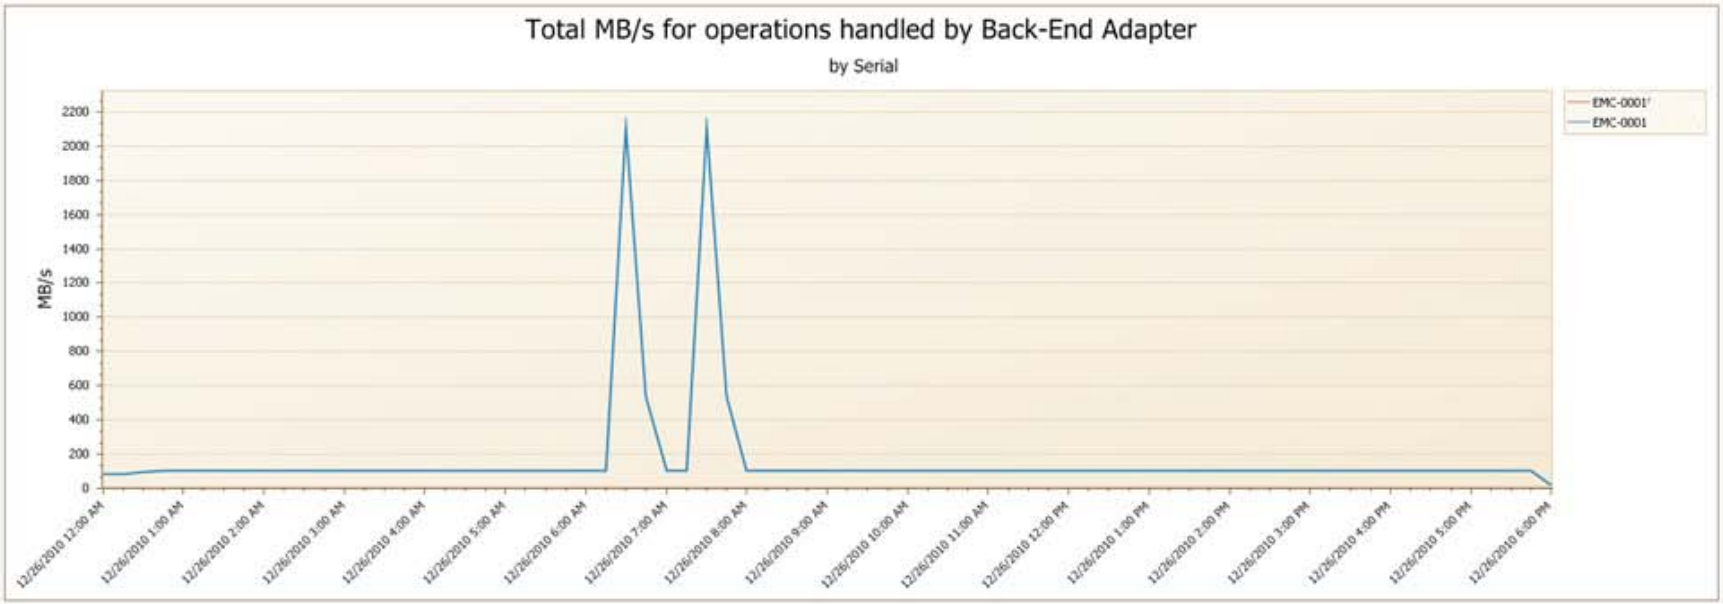

Logical Components: RAID Group Sets/Storage Pools

- RAID Groups are created from Physical Disks
- RAID Group Sets created from RAID Groups
- Volumes created from RAID Group Sets

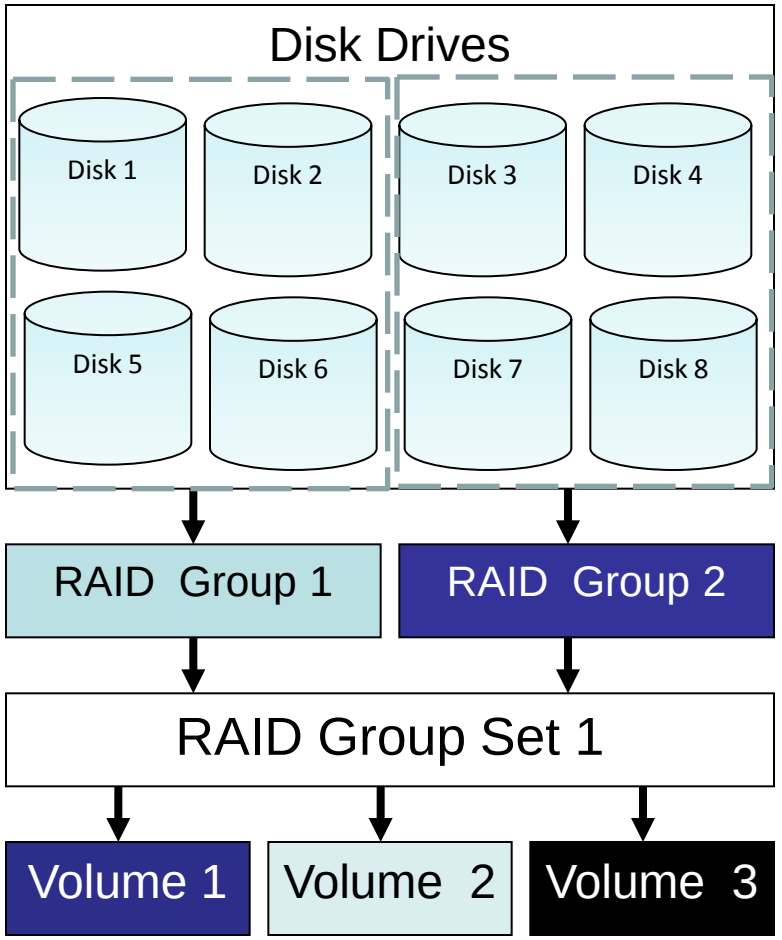

Extent Pool/RAID Group Sets:
Merging Topology and Performance

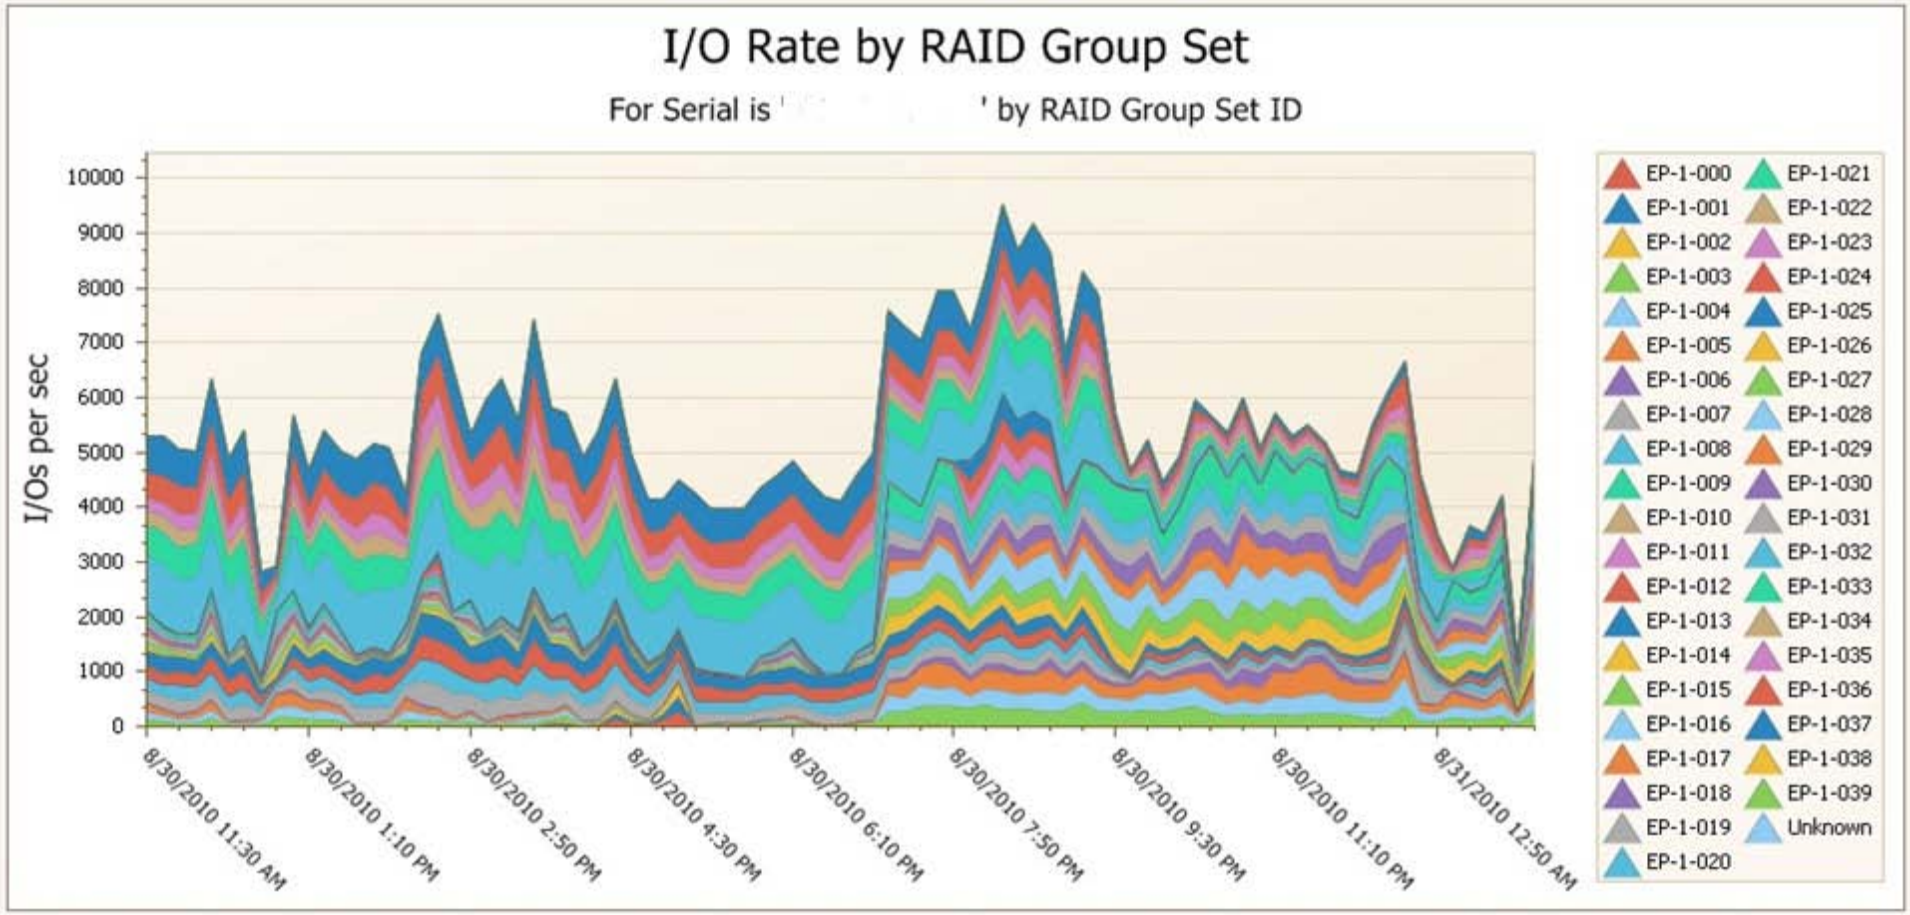

Interpreting
Disk Response Service Levels

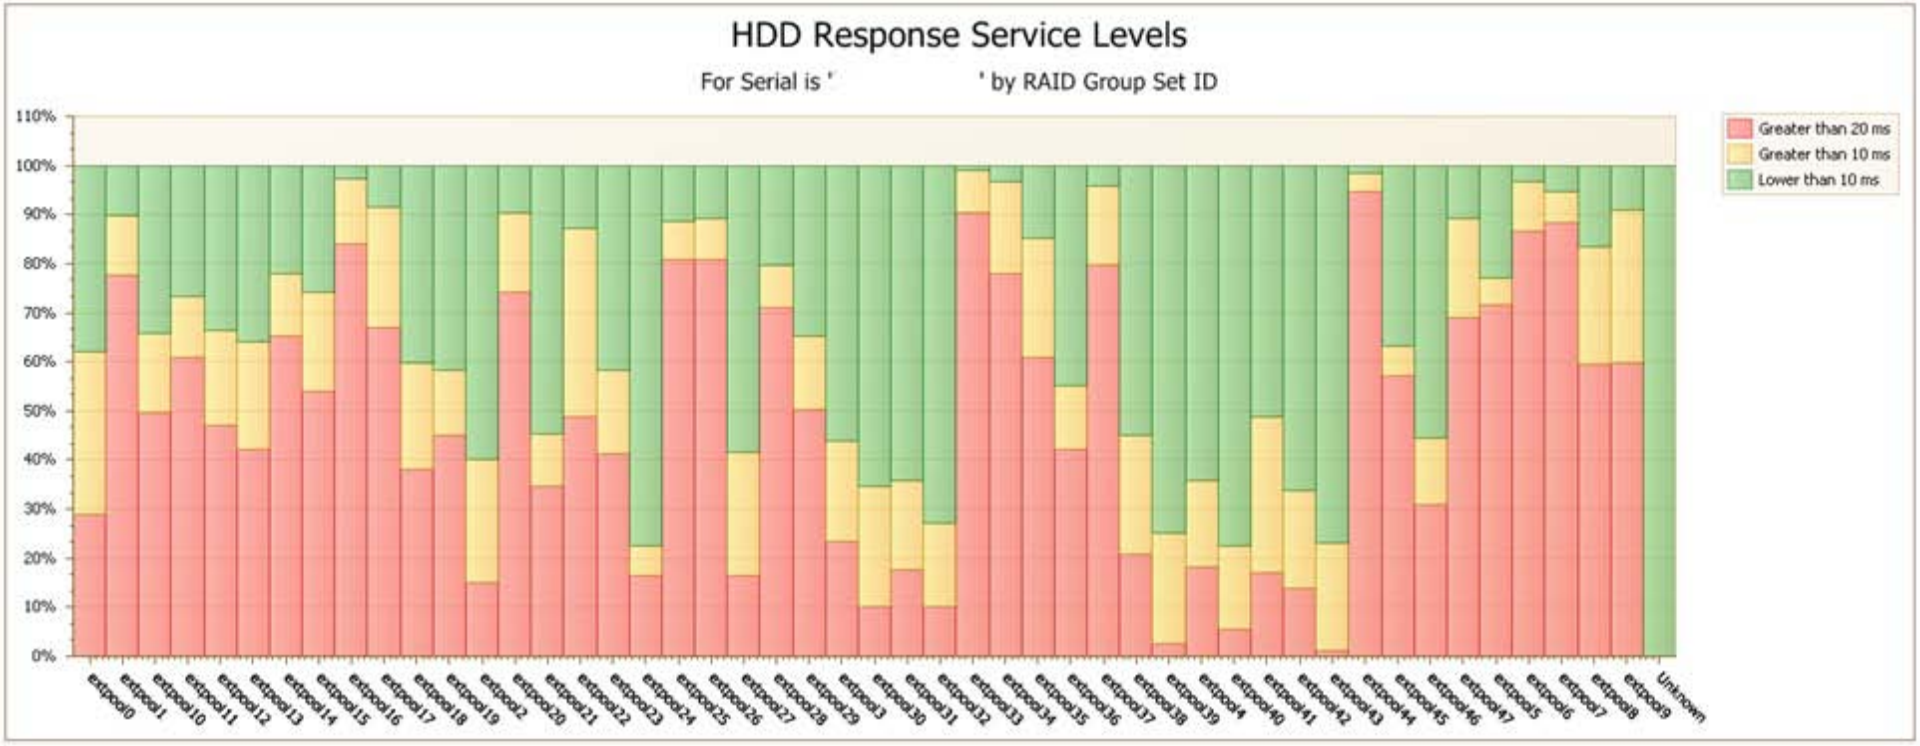

Logical Volumes

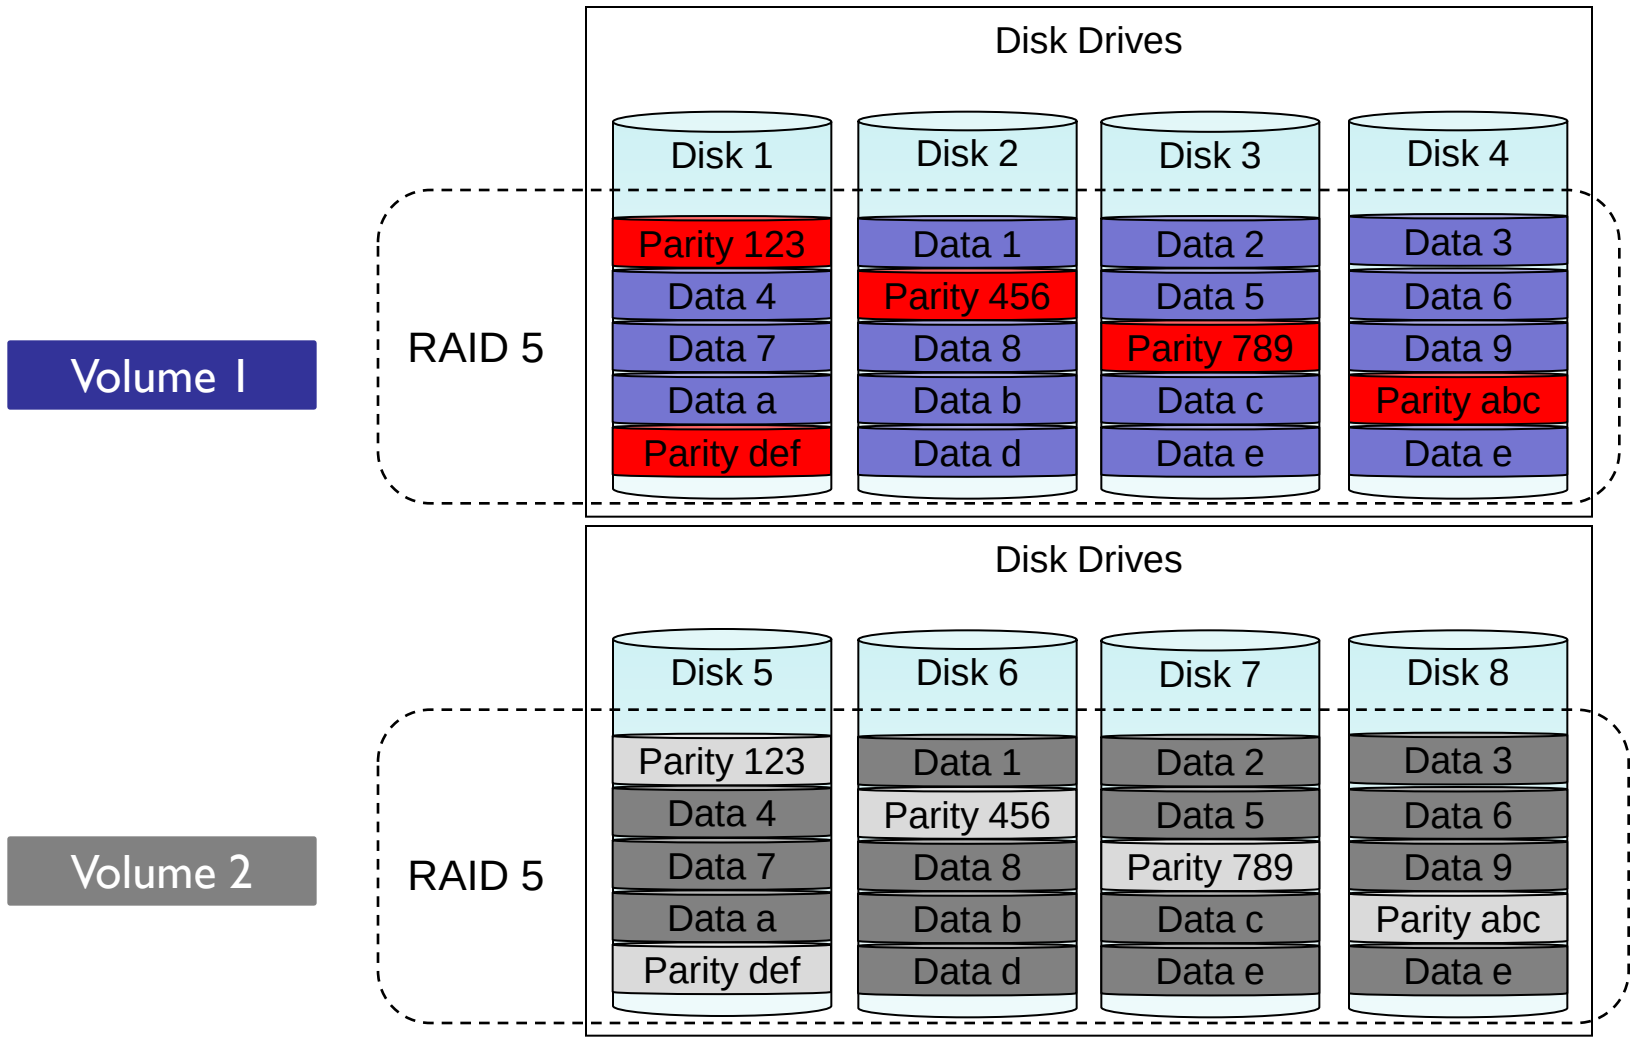

Volume Performance Metrics

Volume – Fast Write Bypass %

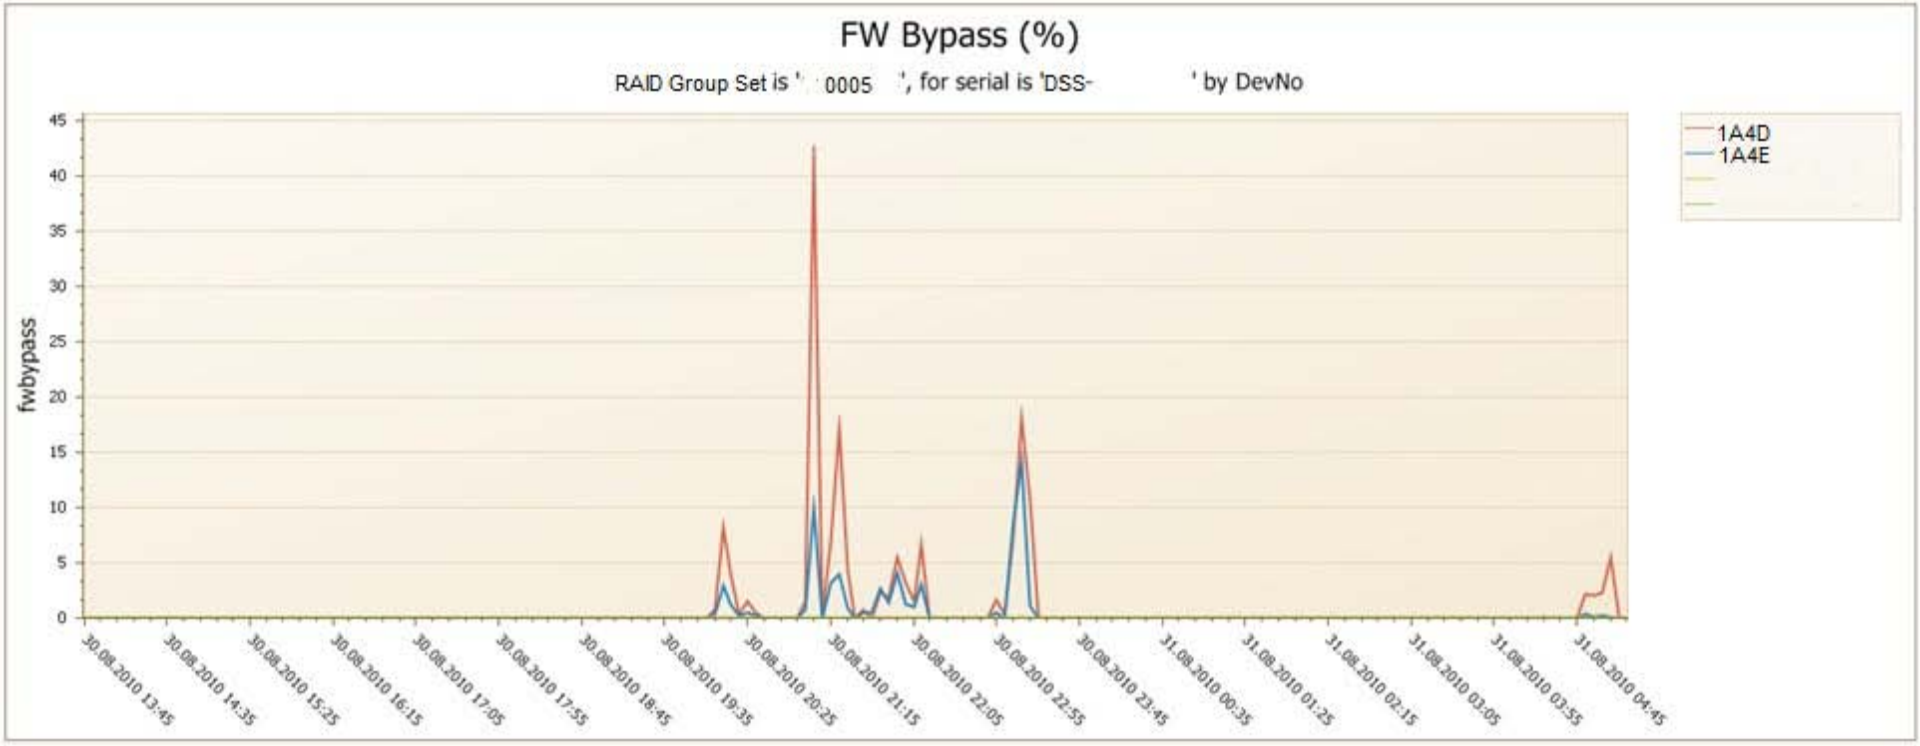

Summary

- The goal of storage performance management (SPM) is to reduce storage costs while maintaining performance SLAs.
- SPM consists of:
  - Processes
  - Measurement
  - Skills
- SMI-S provides a solid foundation for obtaining the necessary measurements to implement SPM

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: tracktutorials@snia.org

Many thanks to the following individuals for their contributions to this tutorial.

Brett Allison **Gilbert Hautekamer**

References

[1]	"User Spending on Enterprise Storage Systems Reached $31 Billion in 2010"	http://storagenewsletter.com/news/marketreport/user-idc--enterprise-storage-systems-2010

Bridging the Visibility Gap

- An effective Storage Performance Management (SPM) solution must:
  - Provide visibility inside the storage system where 70% of the bottlenecks occur, and
  - Automatically correlate your workload metrics with the specific hardware component capabilities.