

Storage Performance Council (SPC)

www.storageperformance.org

spcadmin@storageperformance.org





Presentations and Speakers

- ❑ **SPC-1 Benchmark**
Mel Boksenbaum (*Hitachi Data Systems*)

- ❑ **Comparing SPC-1 Results**
Walter E. Baker (*SPC Administrator & Auditor*)

- ❑ **SPC-2 Benchmark**
Bruce McNutt (*IBM Corporation*)

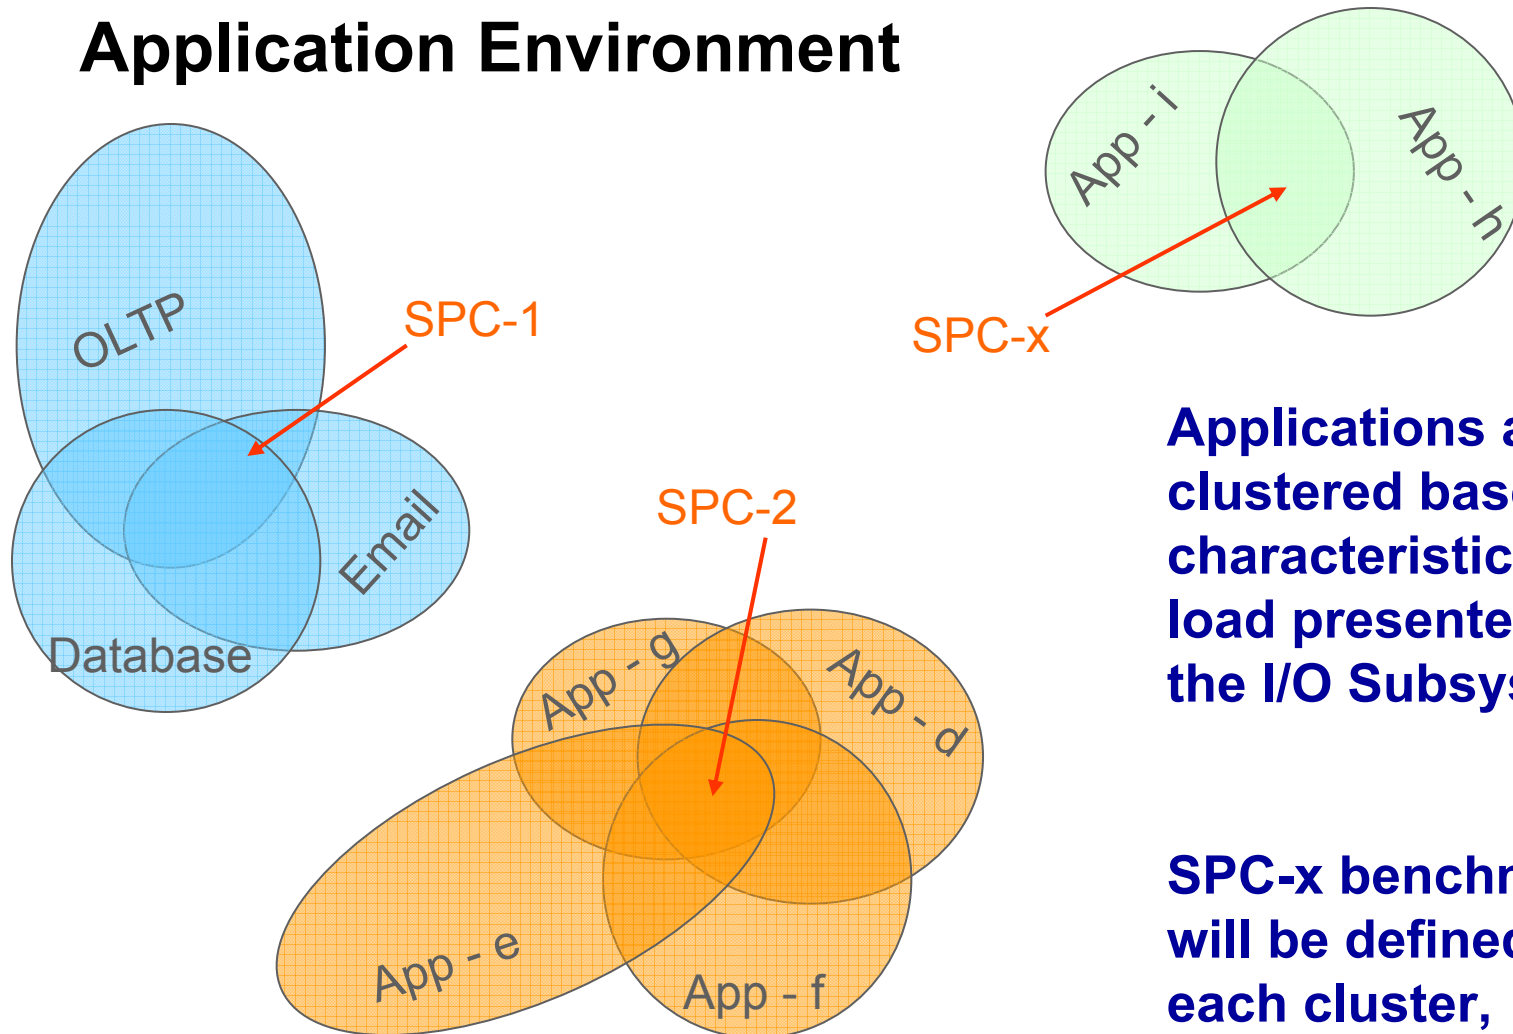
- ❑ **SPC-3 Benchmark**
Leah Schoeb (*Sun Microsystems, Inc.*)

- ❑ **Q&A**
Randy Kerns (*Evaluator Group*)

SPC-1 Benchmark

Mel Boksenbaum

Application Environment

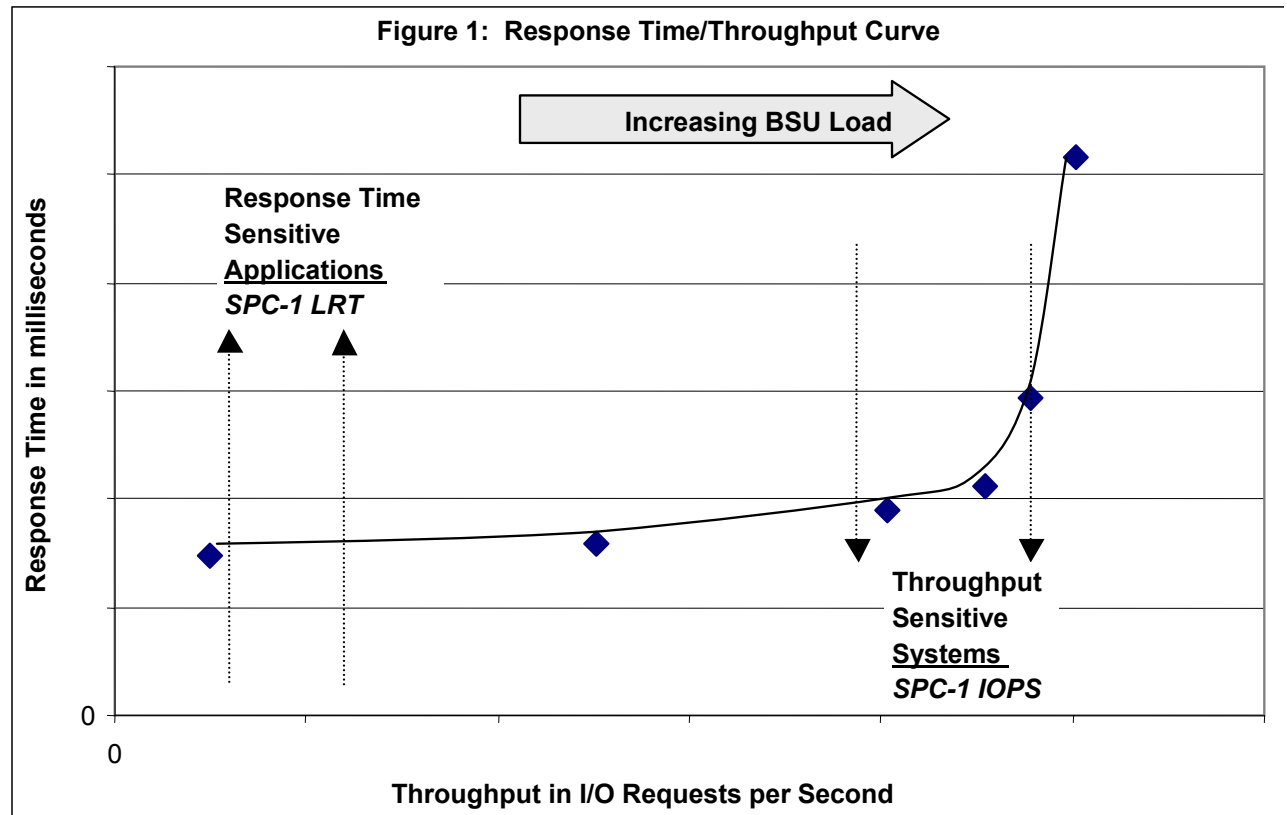


Applications are clustered based on characteristics of load presented to the I/O Subsystem

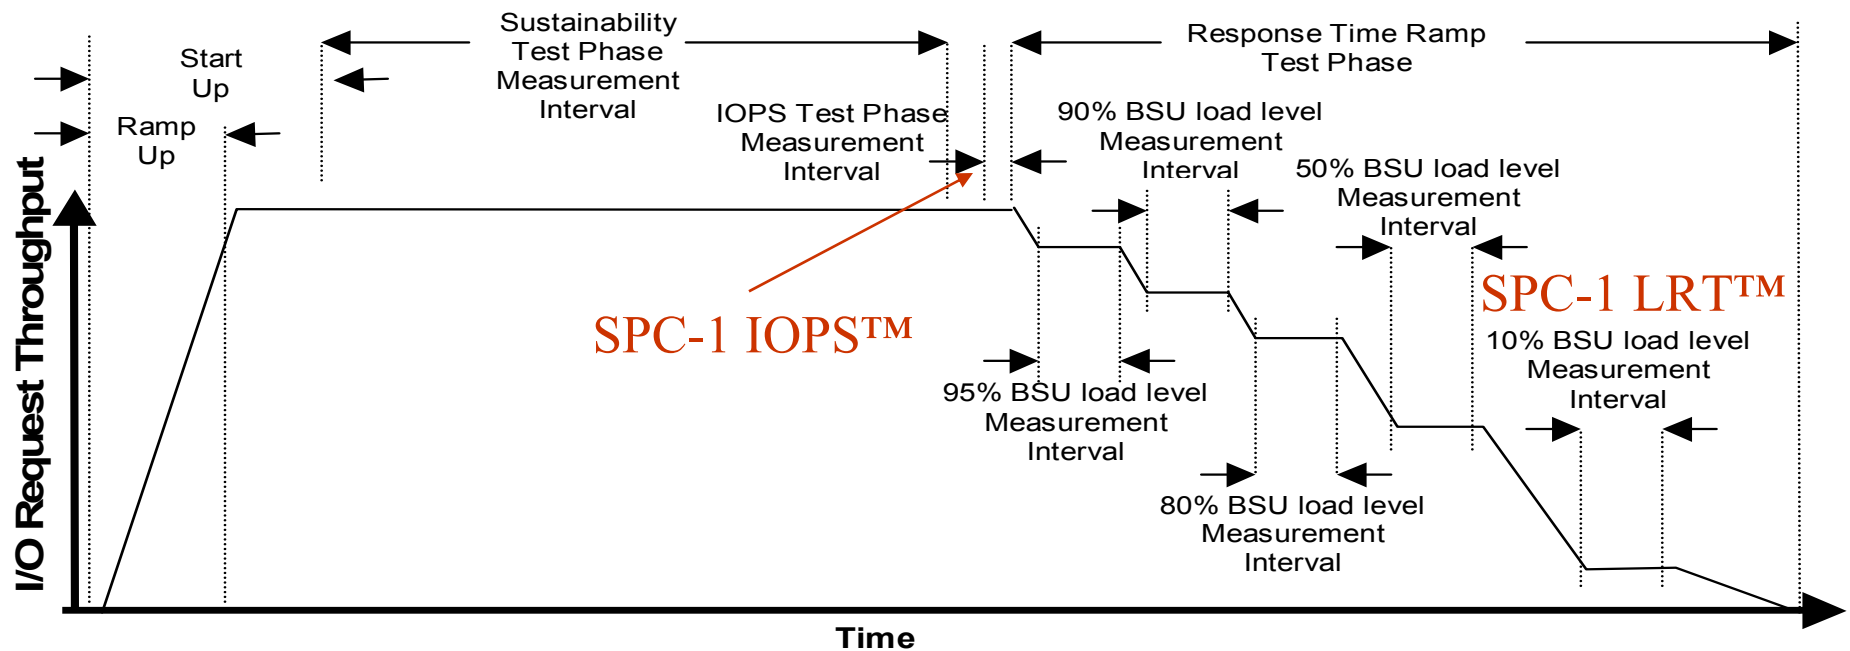
SPC-x benchmark will be defined for each cluster, starting with SPC-1

Important Metrics

- ❑ SPC-1 IOPS (Maximum Throughput)
- ❑ SPC-1 LRT (Best Response Time @ lite load)
- ❑ Price
- ❑ Capacity
- ❑ Data Protection

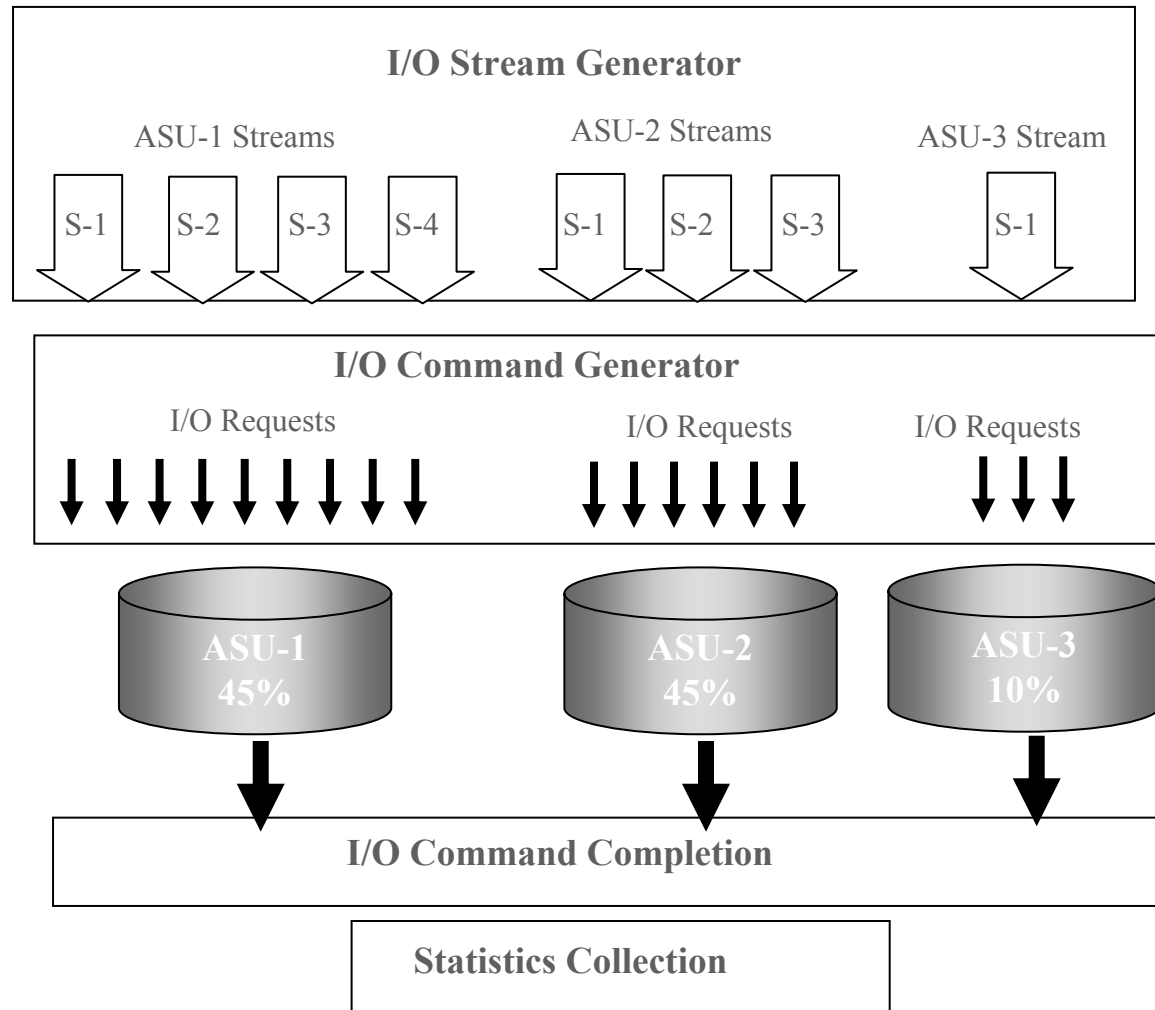


Metric Test Components



& Repeatability & Data Persistence

Structure of SPC-1 I/O





About The Benchmark Kit

SPC-1 benchmark kit includes:

- ❑ Benchmark executables & scripts
 - Solaris
 - Windows 2000,2003
 - AIX
 - HP/UX
 - Linux (future)
- ❑ Results summary & FDR Production tools
- ❑ Documentation



New Variations for SPC-1

- ◆ **SPC1-M: Multi-Address Space Benchmark**
- ◆ **SPC1-F: File System Benchmark**

Other Possibilities...

Comparing SPC-1 Results

Walter E. Baker



Comparing SPC-1 Results

□ Full Disclosure Report (FDR)

- Contains all information necessary to reproduce the reported result.
 - Storage creation and configuration details
 - Test measurement execution details
- Contains the detailed measurement data from the benchmark execution.
- Detailed measurement data provides a means for more in-depth comparison between SPC-1 results.
 - SPC-1 IOPS™ (*I/O Request throughput*)
 - **Response Time**
 - **Storage Capacity**
 - Pricing



Comparing SPC-1 Results


□ SPC-1 IOPS™ and Response Time

- 3 Application Storage Units (ASUs)
 - Each ASU is a set of Logical Volumes that is the target of the SPC-1 I/O requests.
- The FDR contains I/O request throughput and response time at 60-second intervals for each ASU.
- Two SPC-1 results with similar SPC-1 IOPS™ may have significantly different behavior and performance at the ASU level.

Comparing SPC-1 Results

□ SPC-1 IOPS™ Response Time Variability


Reported Average Response Time is approximately the same for both results .



Minimal ramp-up to achieve a uniform, consistent response time across all ASUs.

Longer ramp-up and variable response time across the ASUs.

Superior sequential write performance (ASU-3).






Comparing SPC-1 Results

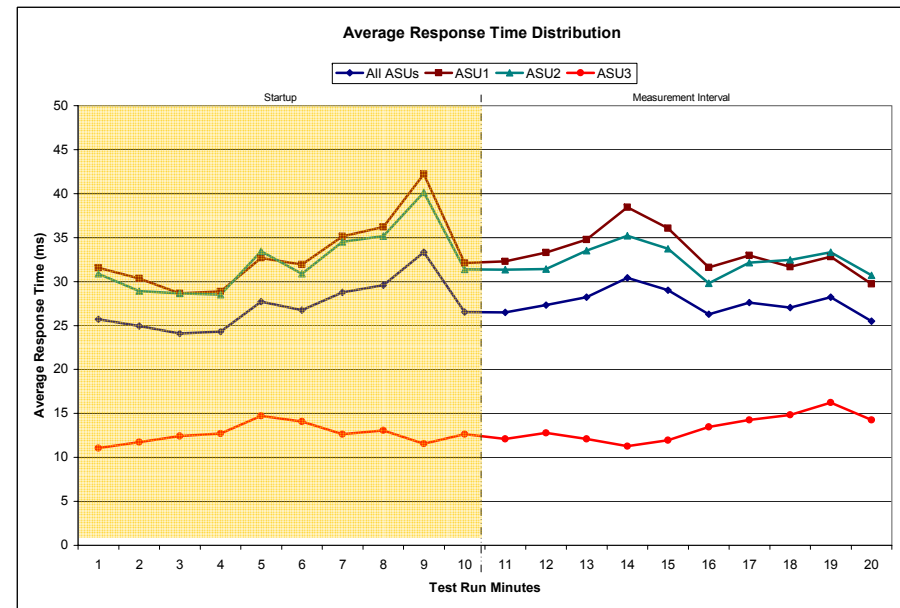
SPC-1 IOPS™ Response Time Comparison

Result 1



Reported Average Response Time for Result 1 is ~30% less than Result 2.

Result 2



Result 1 Sequential write (ASU-3) Response Time is ~130% greater than Result 2.



Comparing SPC-1 Results

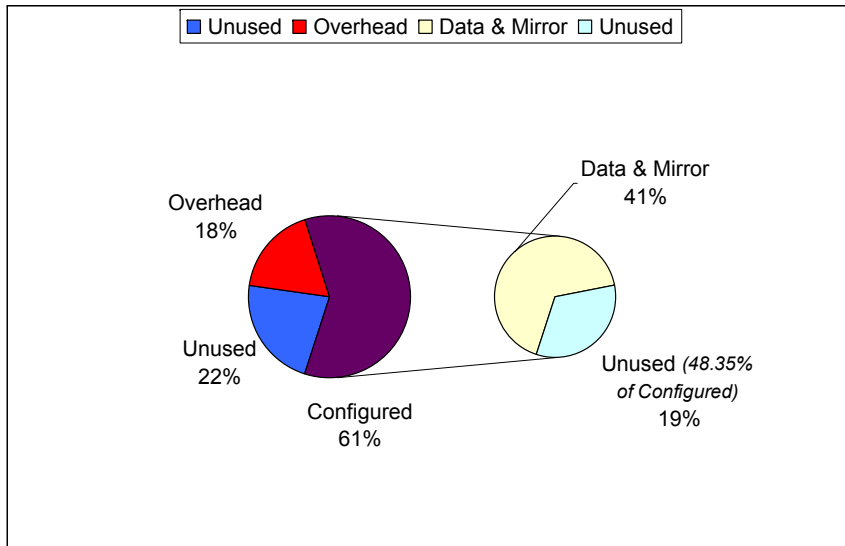
□ SPC-1 Storage Capacity Utilization

- Unused Storage
 - What does the vendor recommend for unused storage capacity (“headroom”)?
 - Are their SPC-1 configurations consistent with those recommendations?
 - What is the performance impact of any “short stroking”?
- Required Overhead: A “cost” of usable storage
 - Significant differences in required overhead between various SPC-1 configurations.
 - With a greater percentage of storage required for overhead, what is the real cost per usable megabyte?

Comparing SPC-1 Results

SPC-1 Storage Capacity Comparison

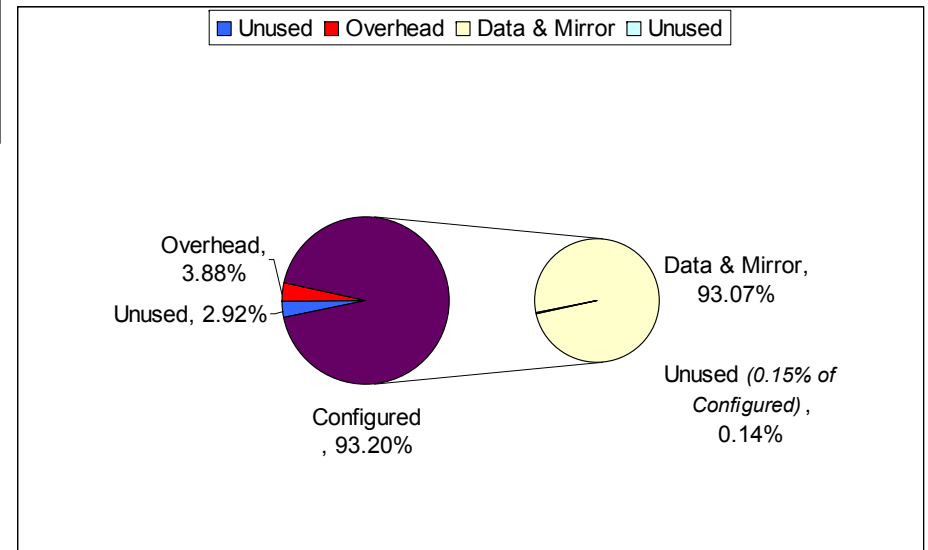
Configuration 1



Configuration 1:

- 22% unused before volume creation**
- 19% additional unused after creation**
- 48% of configured capacity is unused**
- 18% required for overhead**

Configuration 2



Configuration 2:

- 2.92% unused before volume creation**
- 0.14% additional unused after creation**
- 0.15% of configured capacity is unused**
- 3.88% required for overhead**



Comparing SPC-1 Results

- **The FDR provides additional comparisons.**
 - Additional response time/latency comparisons
 - At various throughput levels: 95%, 90%, 80%, 50%, and 10%
 - Over an extended measurement interval (3 hour Sustainability)
 - By ASU/stream type: random, sequential, uniform, etc.
 - By reads/writes within ASU/stream type.
 - Pricing
 - By various storage capacities

SPC-2 Benchmark

Bruce McNutt



SPC-1: Predominately random access

Represents environments such as OLTP, database, and mail server:

	Read (pct)	Write (pct)	All (pct)
Random	29	32	61
Sequential	11	28	39
All	40	60	100



Sequential tests: valuable to many customers

Customers with dedicated servers.

- ❑ Data warehouse & Data mining
- ❑ Scientific Processing
- ❑ Media

Customers with a mix of applications including sequential.

- ❑ Banking & Financial, Insurance, Manufacturing, Medical, Government, many others...
- ❑ To such customers, the sequential component of performance is often the most visible.



SPC-2: A *collection* of sequential tests

Why a collection?

- Types of sequential processing can be *very* different.
- It is easy to include a variety of tests.
- Since the hardware imposes a ceiling on all data rate measurements, a variety of measurements will tend to reinforce each other.



Sequential processing tested in SPC-2

Large file processing: **simple sequential processing of one or more large files.**

Large database queries: **scans or joins of large relational tables.**

Video on demand: **individualized playback from a digital film library.**



SPC-2 Status

A prototype implementation has been tested.

Based on test results, we decided upon significant revisions to the specification.

Please help us review our new specification.

Goal in 2005:

- An approved test available to SPC members.
- An initial set of published SPC-2 results.

SPC-3 Benchmark

Leah Schoeb

A protocol independent benchmark that measures storage performance as seen by applications

- SPC-1 and SPC-2 characterize lower level components in the storage hierarchy
 - Storage from the bottom up
- SPC-3 will take the end-user view of the file system storage environment.
 - Storage from the top down



SPC-3 Environments/workloads

Home Directories

- ❑ Business Applications
- ❑ Backup

Database

- ❑ OLTP
- ❑ Decision Support

Video Post-Processing

- ❑ Digital Film Editing



SPC-3 Critical Questions

Is a measurement at the file system API useful?

What factors make it interesting?

- Scale? Volume, users, bandwidth
- Constraints? Price, latency, protocol
- Topology? NAS, SAN, DAS

Your input is invited and eagerly awaited!

spc3-input@storageperformance.org

Q&A Session

Randy Kerns



Questions for Panel

1. **How can an IT manager use the SPC information?**
2. **Is the SPC-1 benchmark representative of a real environment?**
3. **Why doesn't the SPC have results for all vendor storage systems?**
4. **When will I have a CIFS benchmark?**
5. **Can a vendor "cook" the benchmark results?**
6. **What about results where more drives than were needed for the capacity were used?**
7. **Are the results repeatable?**