



Storage Performance Council (SPC)

www.storageperformance.org
spcadmin@storageperformance.org



Presentations and Speakers

- ***SPC-1 Benchmark***

Mel Boksenbaum (*Hitachi Data Systems Corporation*)

- ***Comparing SPC-1 results***

Walter E. Baker (*Gradient Systems, Inc.*)

- ***SPC-1 Results***

Randy Kerns (*Evaluator Group*)

- ***SPC-2***

Bruce McNutt (*IBM Corporation*)

- ***SPC Future***

Leah Schoeb (*Sun Microsystems, Inc.*)

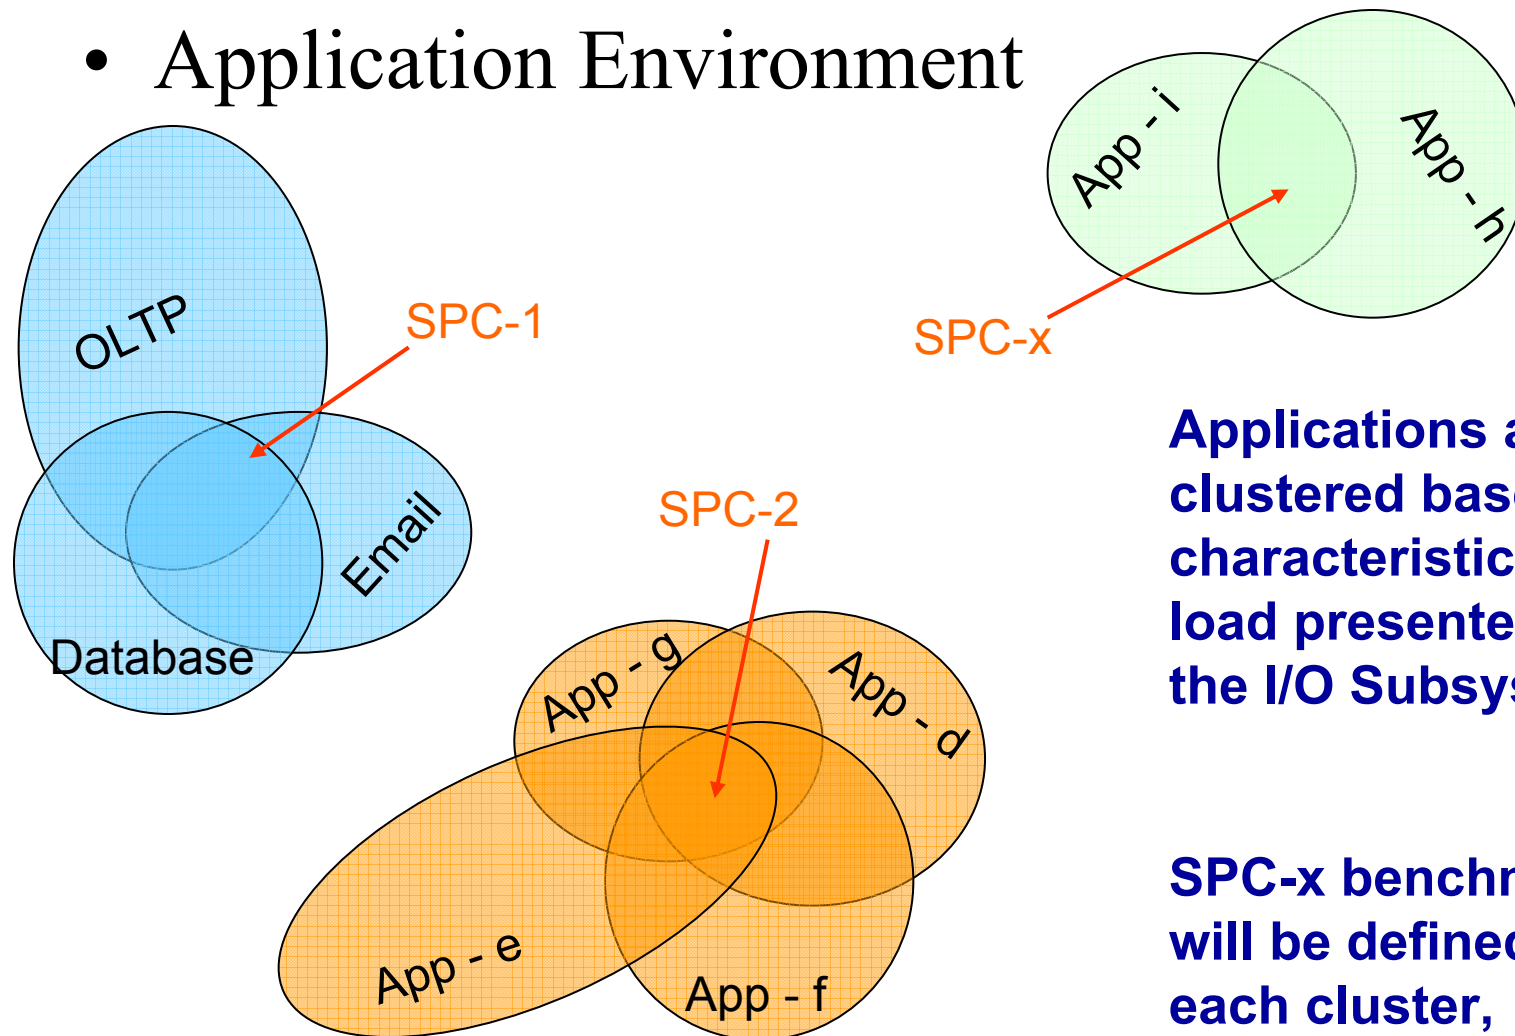


SPC-1 Benchmark

Mel Boksenbaum

SPC Benchmark-1™ Vision

- 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



SPC Benchmark-1™ Vision


- **Important Metrics**

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



SPC Benchmark-1™ Vision

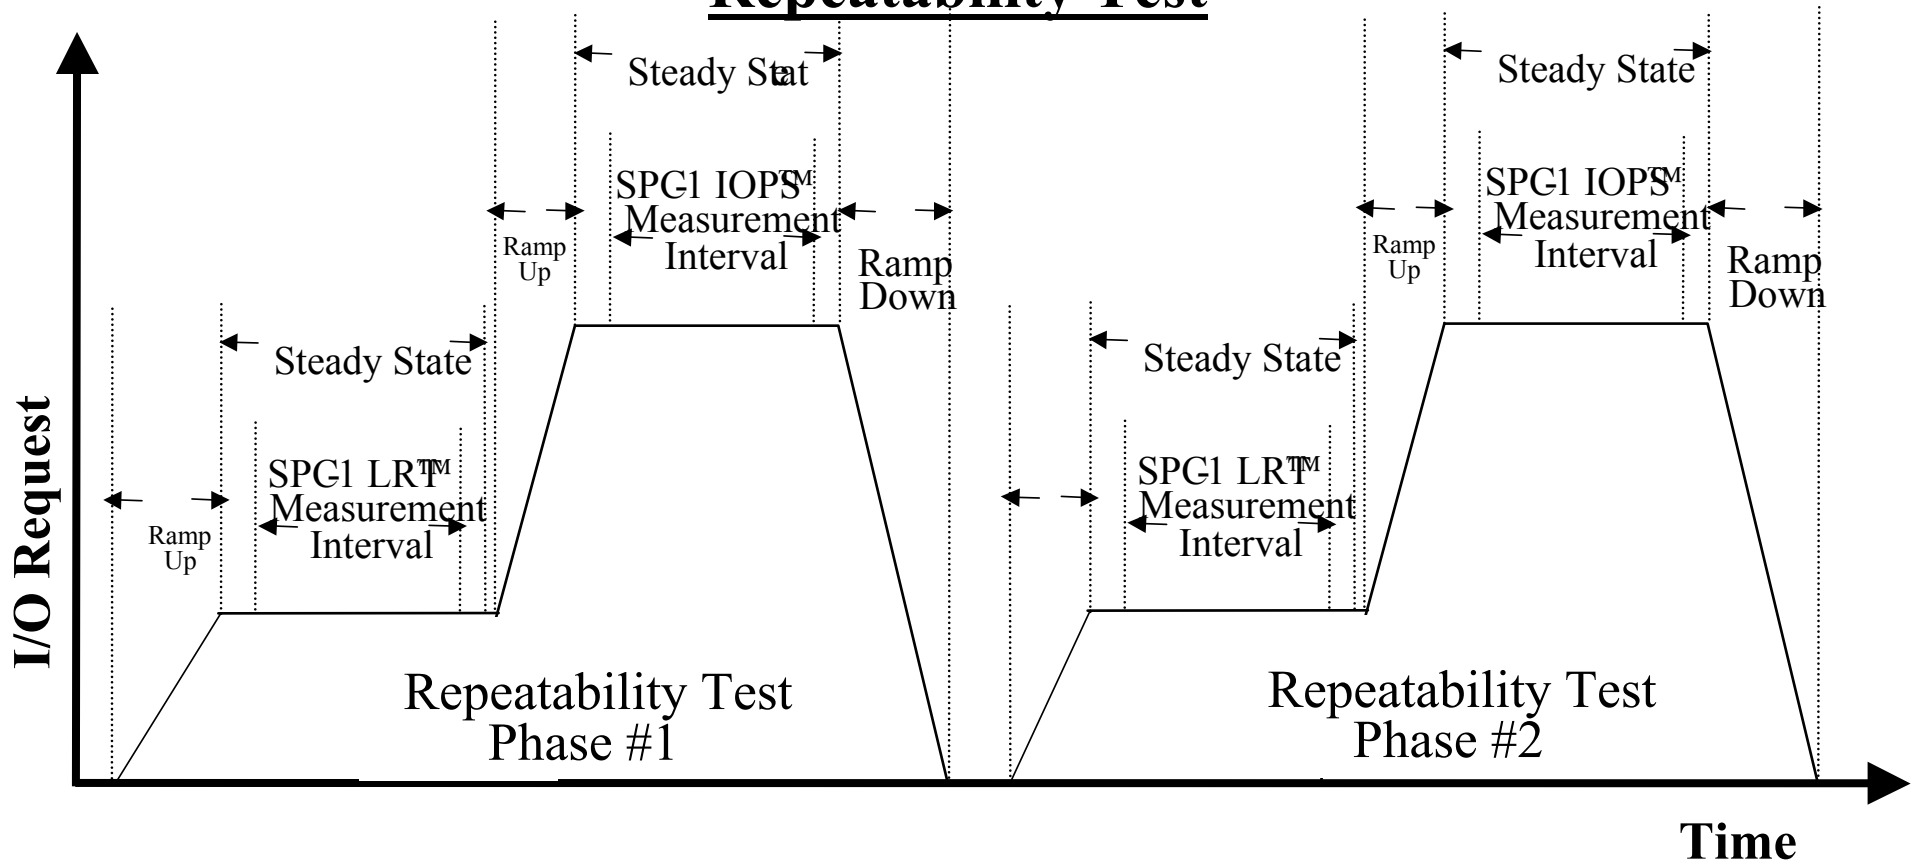
- Metric Test Components



- & Repeatability & Data Persistence

SPC Benchmark-1™ Vision


Repeatability Test






SPC Benchmark-1™ Vision

Data Persistence Test



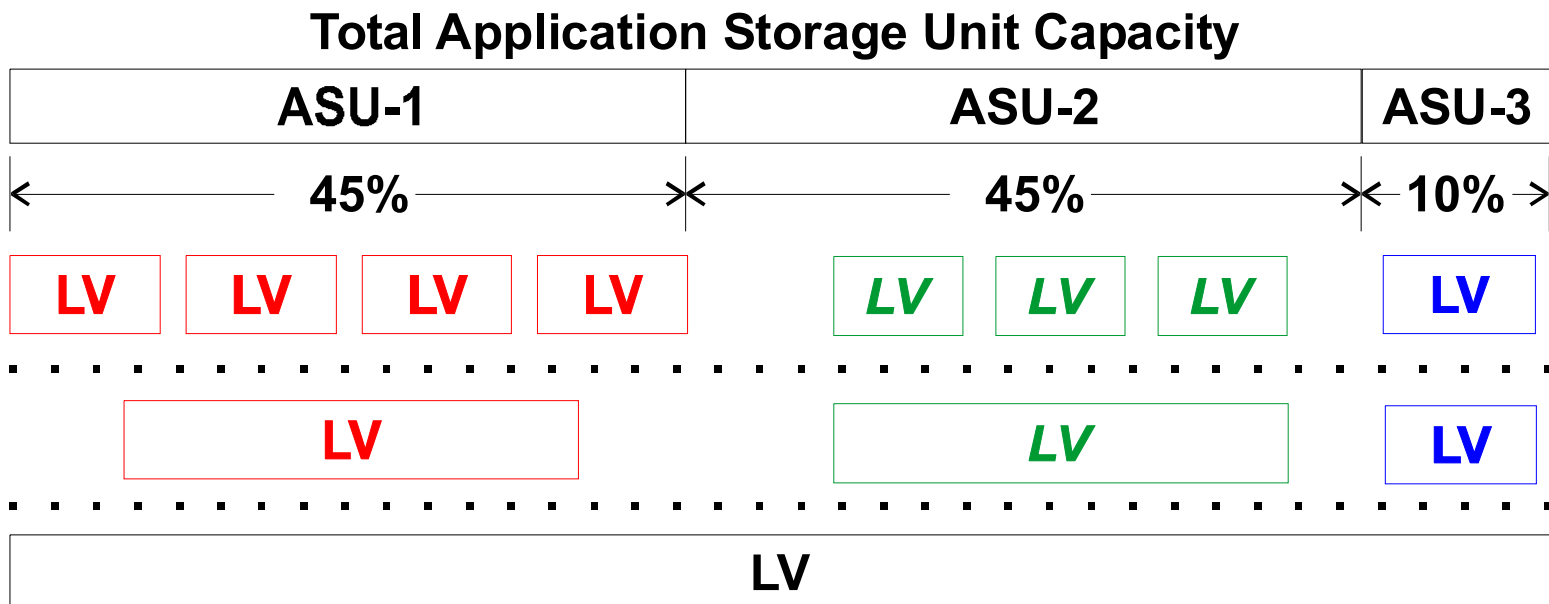
Structure of SPC-1 I/O





Running The SPC-1 Benchmark

- Tested Storage Configuration (TSC)
 - Must conform to the Specification





About The Benchmark Kit

- **SPC-1 benchmark kit includes:**
 - Benchmark executables & scripts
 - Solaris
 - Windows 2000
 - AIX
 - HP/UX
 - Linux (future)
 - Results summary & FDR Production tools
 - Documentation

Comparing SPC-1 Results:

Walter E. Baker



Comparing SPC-1 Results

– Storage Configuration

- Features and Functionality
 - Data Protection, Storage Device Capacity, Storage Device Performance (10K vs. 15K RPM), etc.

– SPC-1 Reported Data

- SPC-1 IOPS™
- SPC-1 Price-Performance
- Total ASU Capacity
- Total Three-Year Price

– Additional Full Disclosure Report Data

- More insight into the SPC-1 reported data



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 IOPSTM (*I/O Request throughput*)
 - Response Time
 - Storage Capacity
 - Pricing



Comparing SPC-1 Results

– **BSUs, ASUs, and I/O Requests**

- **Business Scaling Unit (BSU)** – scaling factor for SPC-1
 - 50 I/O requests generated per BSU
 - I/O requests presented via 5 random I/O streams and 3 sequential I/O streams per BSU
- **3 Application Storage Units (ASUs)**
 - Each ASU is a set of Logical Volumes that is the target of the SPC-1 I/O requests (50 x BSUs specified) via the 8 I/O streams per BSU.



Comparing SPC-1 Results

– SPC-1 I/O Request Distribution

- The distribution of I/O requests across streams/ASUs is fixed by the benchmark specification.

	Type of I/O	% of Total	Read %	Write %
ASU-1/Stream 1 :	Random	03.5	50	50
ASU-1/Stream 2:	Random	28.1	50	50
ASU-1/Stream 3:	Sequential	07.0	100	0
ASU-1/Stream 4:	Random	21.0	50	50
ASU-2/Stream 1:	Random	01.8	30	70
ASU-2/Stream 2:	Random	07.0	30	70
ASU-2/Stream 3:	Sequential	03.5	100	0
ASU-3/Stream 1:	Sequential	28.1	0	100

- *ASU-1/Stream 1 and ASU-2/Stream 1 I/O requests are uniform across the storage.*
- *ASU-1/Stream 2, ASU-1/Stream-4, and ASU-2/Stream-2 I/O requests are targeted for “hot spots” across the storage.*



Comparing SPC-1 Results

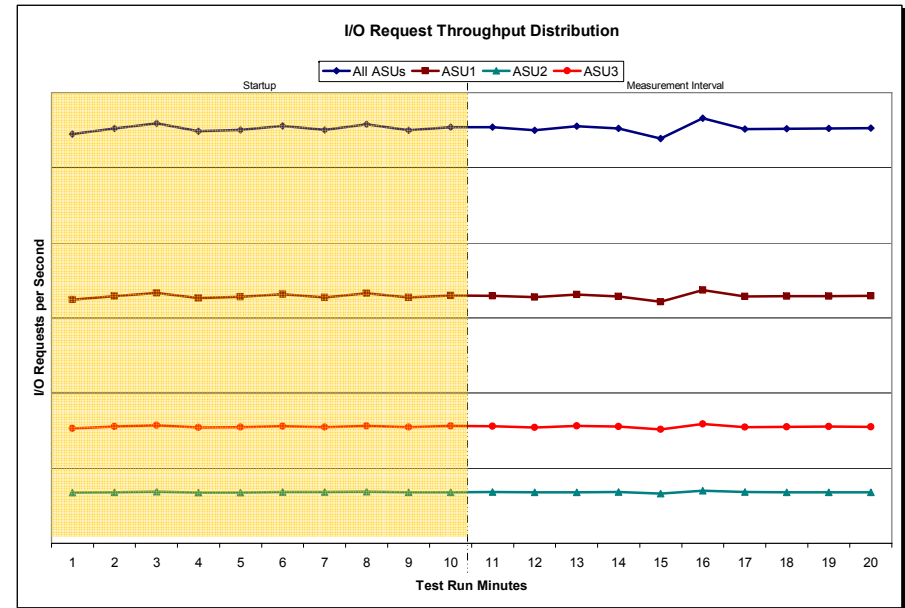
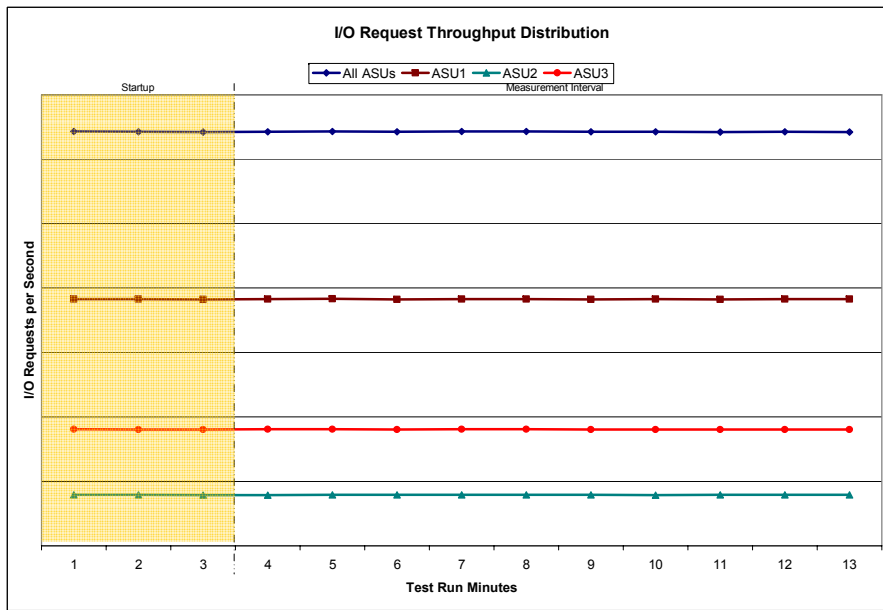
– SPC-1 IOPS™ and Response Time

- 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 IOPSTM - Ramp-Up Duration and I/O Request Throughput Variability






Comparing SPC-1 Results

– SPC-1 IOPS™ Response Time Variability

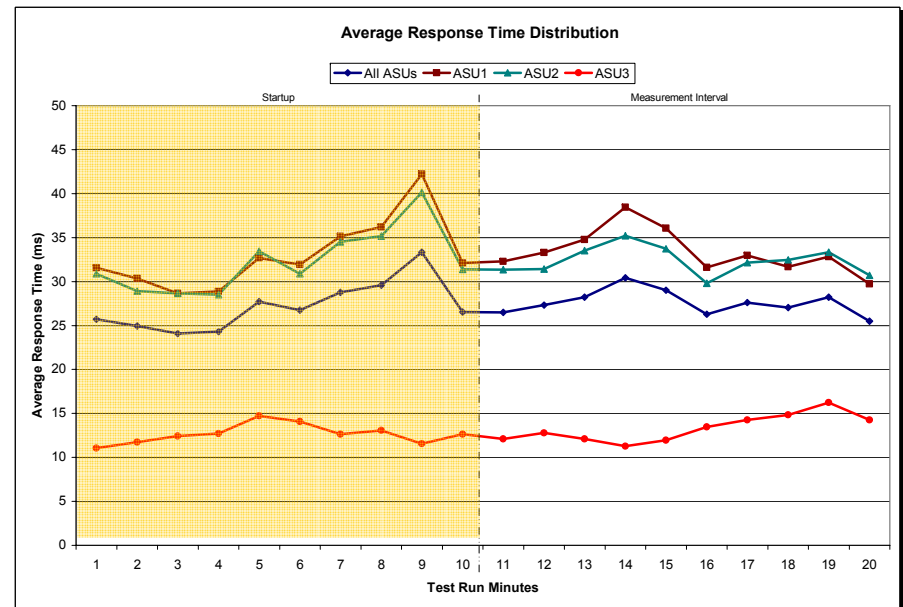
Reported Average Response Time is approximately the same for both results .



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.

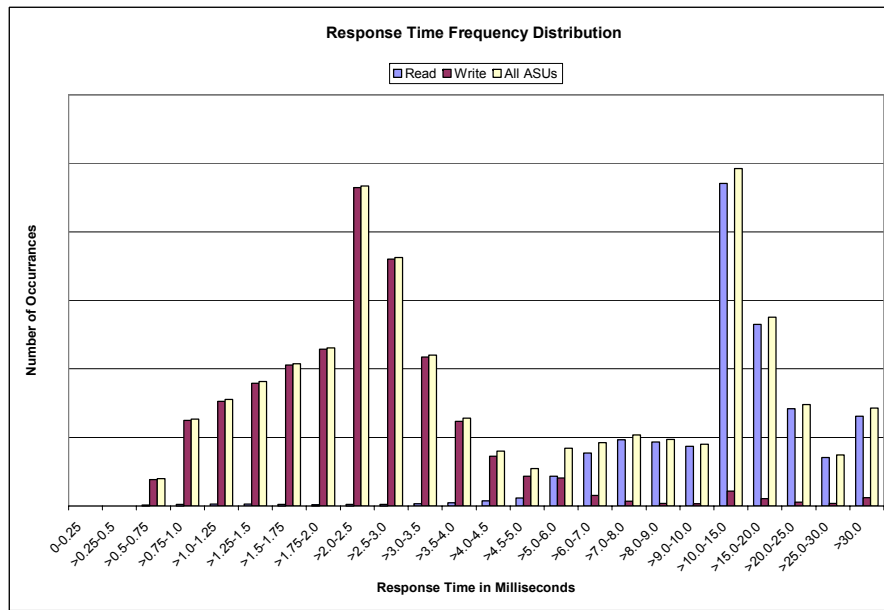
Result 2



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


Comparing SPC-1 Results

SPC-1 IOPS™ Reads and Writes Response Time Distribution



	<i>0-5 ms</i>	<i>>5-10 ms</i>	<i>>10-30 ms</i>	<i>>30 ms</i>
Reads	2.88%	20.44%	68.09%	8.59%
Writes	94.69%	2.86%	1.93%	0.52%

	<i>0-5 ms</i>	<i>>5-10 ms</i>	<i>>10-30 ms</i>	<i>>30 ms</i>
Reads	36.05	7.57	30.91%	25.47%
Writes	39.19%	28.48%	27.76%	4.57%





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.
- Storage capacity utilization:
 - Unused storage ratios – possible “short stroking”
 - “Cost” of usable storage – various overhead capacities
- Pricing
 - By various storage capacities
- ...

SPC-1 Results

Randy Kerns




SPC-1 Results

- As of 11/21/03
- Two categories:
 - Single image storage systems
 - Cache centric
 - Distributed
 - Clustered distributed
 - Open platform
 - Multiple storage systems
- Excluded from the following charts (for brevity)
 - Non mirrored cache results

SPC Posted Benchmark Results 11/21/03

(Single System Image - Mirrored cache results)

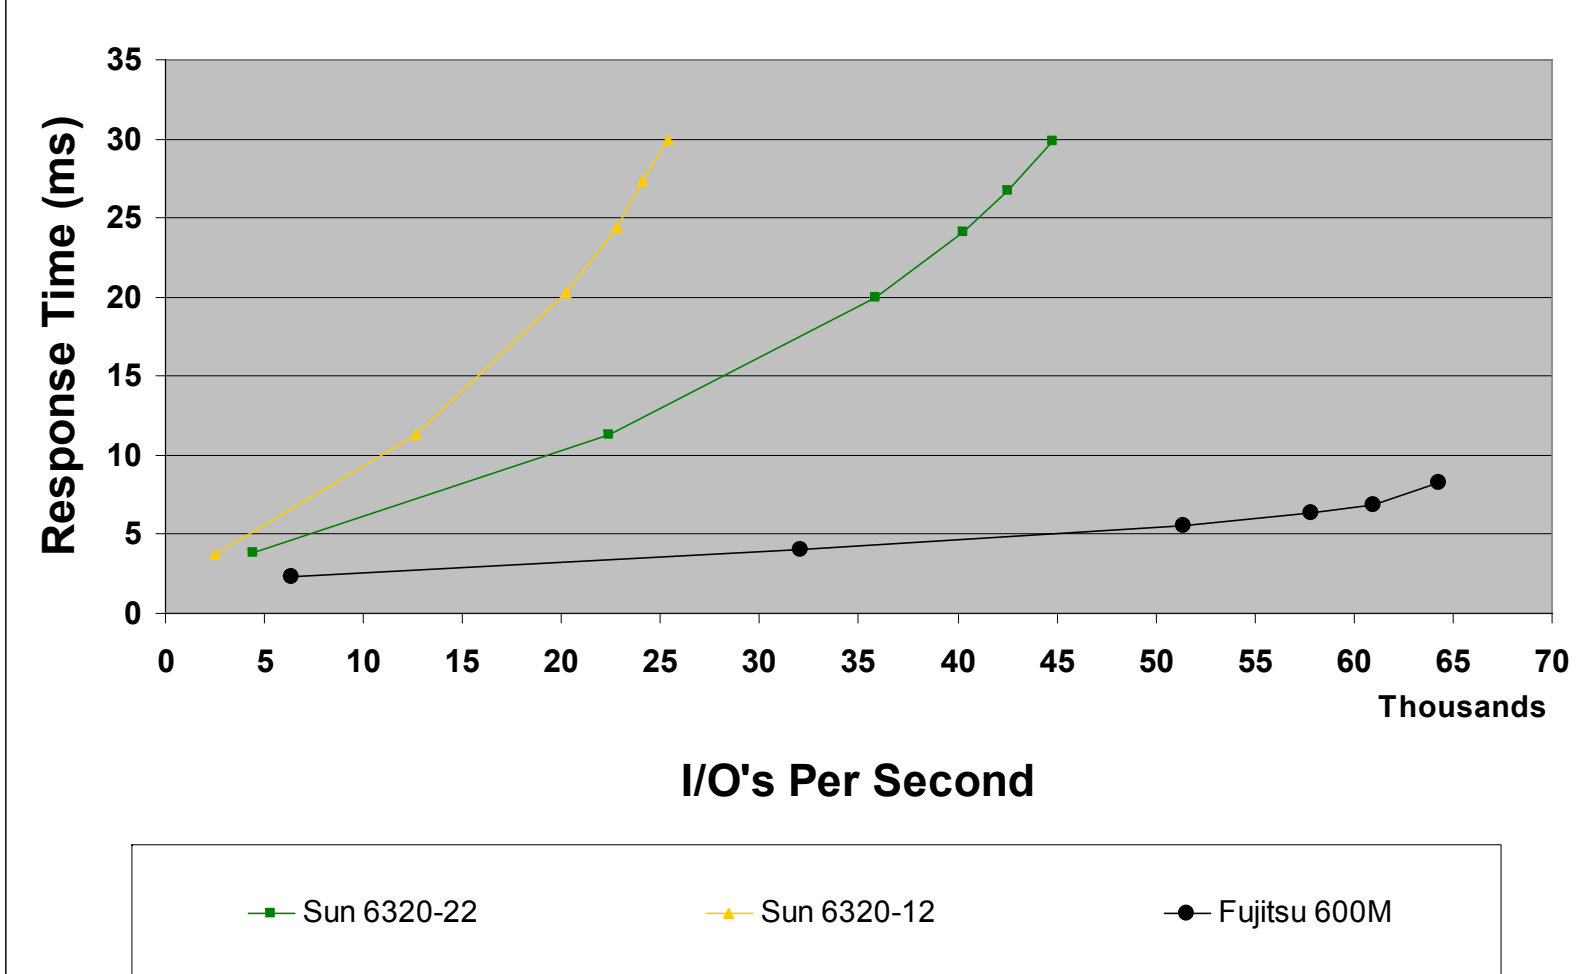


I/O's Per Second



SPC-1 Posted Benchmark Results 11/21/03

(Multiple Systems - mirrored cache results)





SPC-1 Results – Single Image

Vendor / Product	Protection	No. Disk Spindles	Capacity GBs	SPC-1 IOPS	IOPS @ <10ms	Price-Perf (\$/IOP)
Dell Perc3/QC	Mirror	56 18GB 15K rpm	515.2	7,650	~7,000	4.48
LSI 4600	Mirror	112 18GB 15K rpm	482.8	15,708	~14,000	16.01
HP EVA (mirrored cache)	Mirror	168 36GB 15K rpm	2,596.31	20,097	~16,000	23.88
Fujitsu ETERNUS3000 Model 400	Mirror	120 73GB 15K rpm	4,027.64	17,545	~16,500	37.92
Fujitsu ETERNUS3000 Model 600	Mirror	240 73GB 15K rpm	7,849.41	33,497	~33,497	34.61
3Par S800 8-Node	Mirror	640 18GB 10K rpm	4,444.44	47,001	~43,000	24.90
3Par S800 2-Node	Mirror	120 18GB 10K rpm	896	12,905	~12,000	19.71



SPC-1 Results – Single Image

Vendor / Product	Protection	No. Disk Spindles	Capacity GBs	SPC-1 IOPS	IOPS @ <10 ms	Price-Perf (\$/IOP)
Sun 9910	Mirror	48 73GB 10K rpm	1,600	8,404	8,404	74.29
Sun 3510	Mirror	36 36GB 15K rpm	644	11,050	~8,000	7.95
IBM ESS 800 Turbo	RAID 5	256 18GB 15K rpm	3,328	22,999	~22,000	34.88
IBM FASt900 (mirrored cache)	Mirror	108 36GB 15K rpm	1,196	18,448	~15,500	16.78
IBM FASt600t (mirrored cache)	Mirror	50 36GB 15K rpm	478	9,100	~6,800	11.86
DataCore SANsymphony	Mirror	220 18GB 15K rpm	1,407	50,004	50,004	6.11



SPC-1 Results – Multiple Systems

Vendor / Product	Protection	No. Disk Spindles	Capacity GBs	SPC-1 IOPS	IOPS @ <10ms	Price-Perf (\$/IOP)
Fujitsu ETERNUS3000 Model 600M	Mirror	480 73GB 15K rpm	15,609	64,250	~64,250	32.72
Sun SE6320 12 Tray	Mirror	168 36GB 15K rpm	1,960.13	25,340	~10,000	15.40
Sun SE6320 22 Tray	Mirror	308 36GB 15K rpm	3,275.98	44,805	~20,000	15.56

SPC-2

Bruce McNutt



Storage Performance Council

- Wide storage industry representation.
 - ✓ Over 30 full, associate and academic members.
- An open, objective, and verifiable test process.
 - ✓ All submissions follow a defined SPC-1 standard.
- Open availability of benchmark results.
 - ✓ Nearly 20 published submissions at last count.
- Ultimately, a higher bar for storage performance.
 - ✓ Have you noticed some improvements lately?



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.



Important examples of sequential processing

- ***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.



Typical Large File workload

- Single-threaded for a given file.
- Both read and write processing.
- Uses file system.
- File system may provide readahead/writebehind.
- Wide range of transfer sizes.



Typical Large Query workload

- Multi-threaded for a given file.
- Almost read-only.
- Uses database rather than file system I/O.
- Database provides tunable readahead.
- Wide range of transfer sizes.



Typical VoD workload

- Read-only.
- Special-purpose I/O software.
- Each “stream” must run at a fixed rate.
- Optimized transfer size (e.g. 256K).



Advantages of a broad-based test

- Multiple, different tests probe for the “ceiling” better than any single test.
- Multiple tests show the consistency, flexibility, and special problems (if any) of a specific storage product.
- Multiple tests create the possibility of estimating performance for workloads that are “in between”.

SPC Future

Leah Schoeb



SPC-3 Workgroup

- Investigating a protocol-independent, file system benchmark.
- Target is a broad range of storage configurations and topologies.
 - NAS/SAN/DAS



SPC-3 Workgroup

- Areas of investigation/challenges/issues:
 - How to ensure a “storage” focused benchmark rather than a system benchmark given the potential for multi-tier benchmark configurations.
 - Benchmark Configuration: how many tiers (client/server/storage) will be physically present?



SPC-3 Workgroup

- Areas of investigation/challenges/issues:
 - What are the required file system attributes?
 - Should the file system configuration scale with storage capacity?
 - If so, using what attributes?
 - file size/directory path length/directory size/file type/others?



SPC-3 Workgroup

- Areas of investigation/challenges/issues:
 - How to determine the set of file operations that will comprise the workload.
 - Oriented towards a specific application type or types? If so, what application type(s)
 - Or a reasonable set of file operations that span applications, but not associated with a specific application type.



SPC-3 Workgroup

- The preceding are some of the major topics of investigation by the workgroup.
- The Workgroup's investigation will result in one of two outcomes:
 - A recommendation to begin benchmark development based on a specific proposal from the Workgroup.
 - A recommendation to not pursue the current direction.
- Your feedback and comments are welcome in the course of the SPC-3 Workgroup's investigation.



SPC-1 for File Systems

- SPC-1 does not allow the use of file system configurations.
 - SPC-1 is focused on raw, unblocked I/O.
 - File system functionality is not allowed.
- A new benchmark based on the SPC-1 workload, but focused on file system configurations.



SPC-1 for File Systems

- The new benchmark would generate a separate set of SPC results.
 - The results would not be directly comparable to the current SPC-1 benchmark results.
- Primary advantages:
 - Starting with a predefined workload that would result in a short development cycle.
 - Should result in a less complex benchmark configuration and be storage-focused.
 - No network clients, only server(s) and storage.



SPC-1 for File Systems

- Primary advantages:
 - Will characterize file system storage performance for application types characterized by predominately random I/O and requiring queries as well as updates.
 - OLTP, database, mail server applications are typical of those application types.
- Would this type of benchmark be of value to end users and vendors?
 - We would like your feedback.