

Storage Performance Council (SPC)

www.StoragePerformance.org

SPCadmin@StoragePerformance.org





Presentations and Speakers

- ❑ ***Introductions***
Mel Boksenbaum (*Hitachi Data Systems*)
- ❑ ***An Introduction to the SPC***
Leah Schoeb (*Sun Microsystems, Inc.*)
- ❑ ***SPC-1/SPC-2 Update***
Walter E. Baker (*SPC Administrator and Auditor*)
- ❑ ***SPC-1C / SPC-2C (Component-Level Benchmarks)***
Craig Parris (*Seagate Technology LLC and SPC-C Chair*)
- ❑ ***SPC-3 Development Status***
Bruce McNutt (*IBM Corporation and SPC-3 Chair*)
Jack Stephens (*SPC Administrator*)
- ❑ ***2008 SPC Benchmark Activity, Q&A***
Walter E. Baker (*SPC Administrator & Auditor*)



An Introduction to the SPC

Leah Schoeb
Sun Microsystems, Inc.



Council Mission

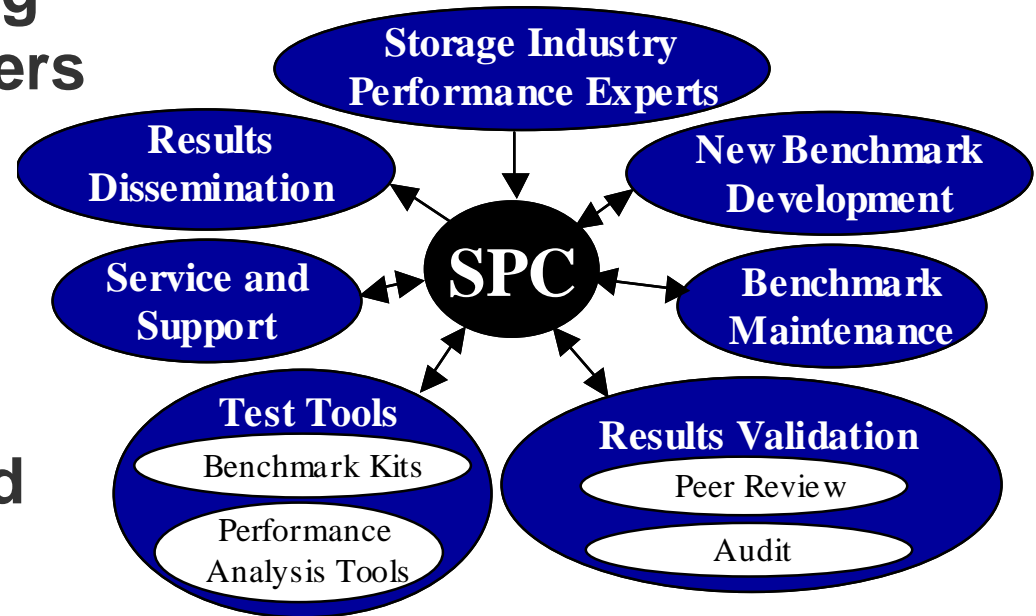
The SPC is a non-profit corporation founded in 1998 to

- ❑ define, standardize, and promote the first industry standard storage benchmarks
- ❑ to disseminate objective, verifiable performance data
- ❑ to developers, product managers, and customers of computer systems



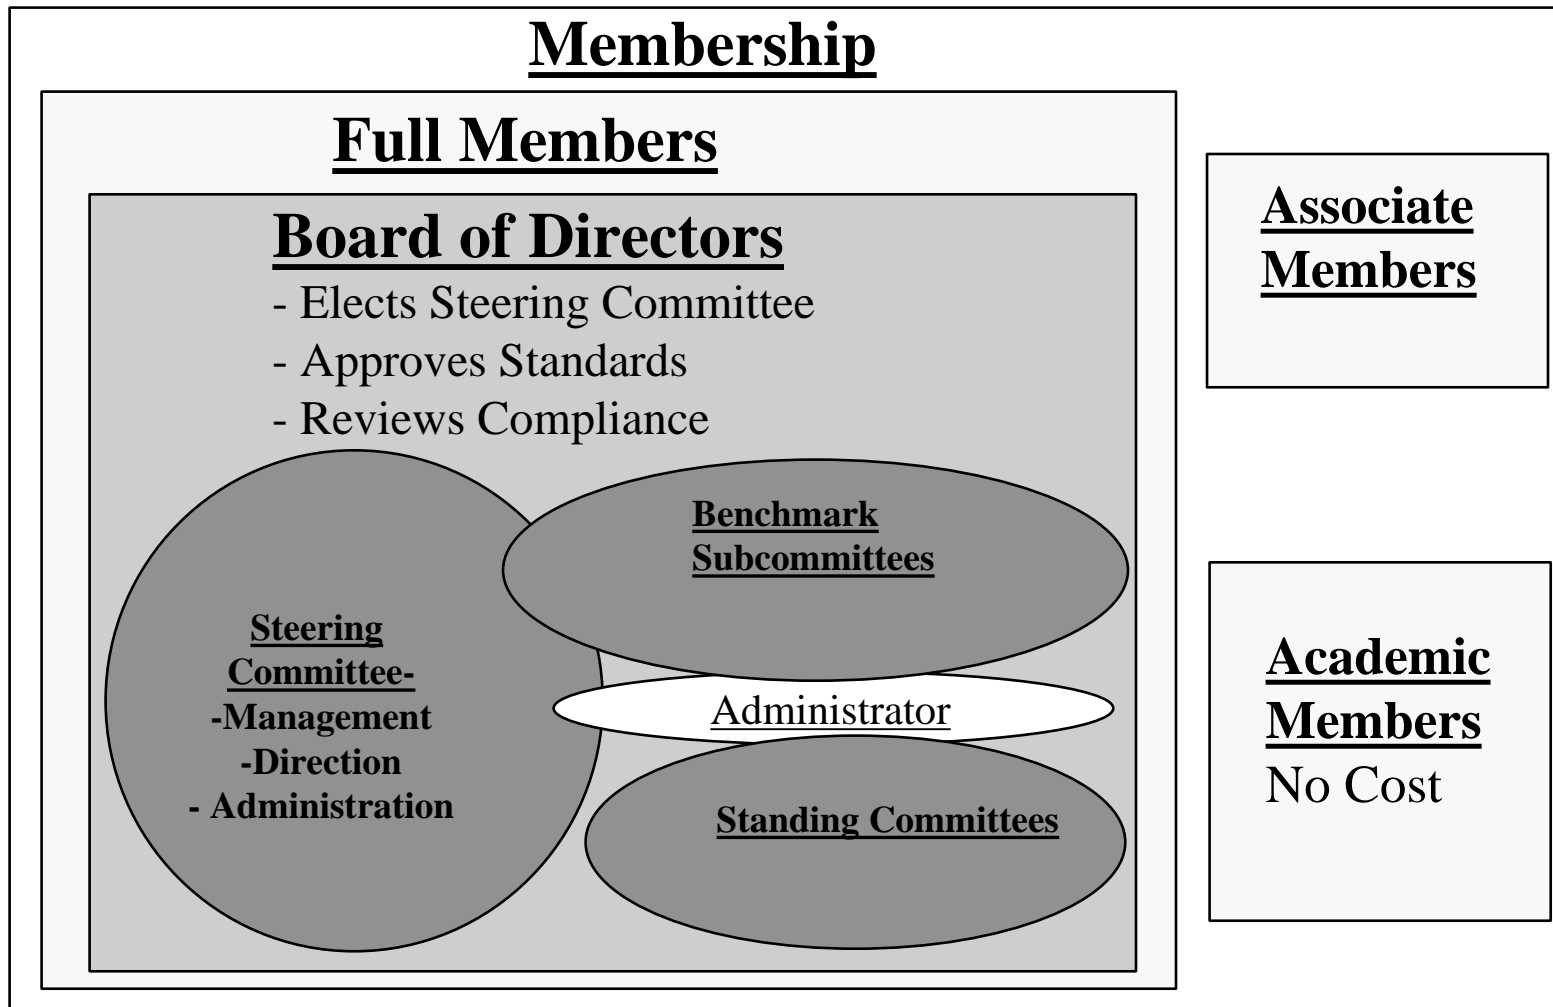
SPC Objectives

- Drive performance improvement in the computer systems industry for storage
- Insure that customers can accurately compare products in a multi-vendor world
- Establish a level “playing field” for all manufacturers
- Publicize benchmark results
- Ensure accuracy and authenticity of published results





SPC Organizational Structure





SPC Membership





SPC-1/SPC-2 Update

Walter E. Baker
SPC Administrator and Auditor




SPC-1 Overview

- A single workload to demonstrate storage subsystem performance typical of business critical application.
- Applications characterized by predominately random I/O operations and require both queries as well as update operations.
- Application examples include OLTP, database operations, and mail server implementations



SPC-1 Results

- Continued submission of SPC-1 Results across a wide range of configurations.





SPC-2 Overview

- Three distinct workloads to demonstrate storage subsystem performance of business critical applications that require large-scale, sequential movement of data.
- Those applications are characterized predominately by large I/Os organized into one or more concurrent sequential patterns.



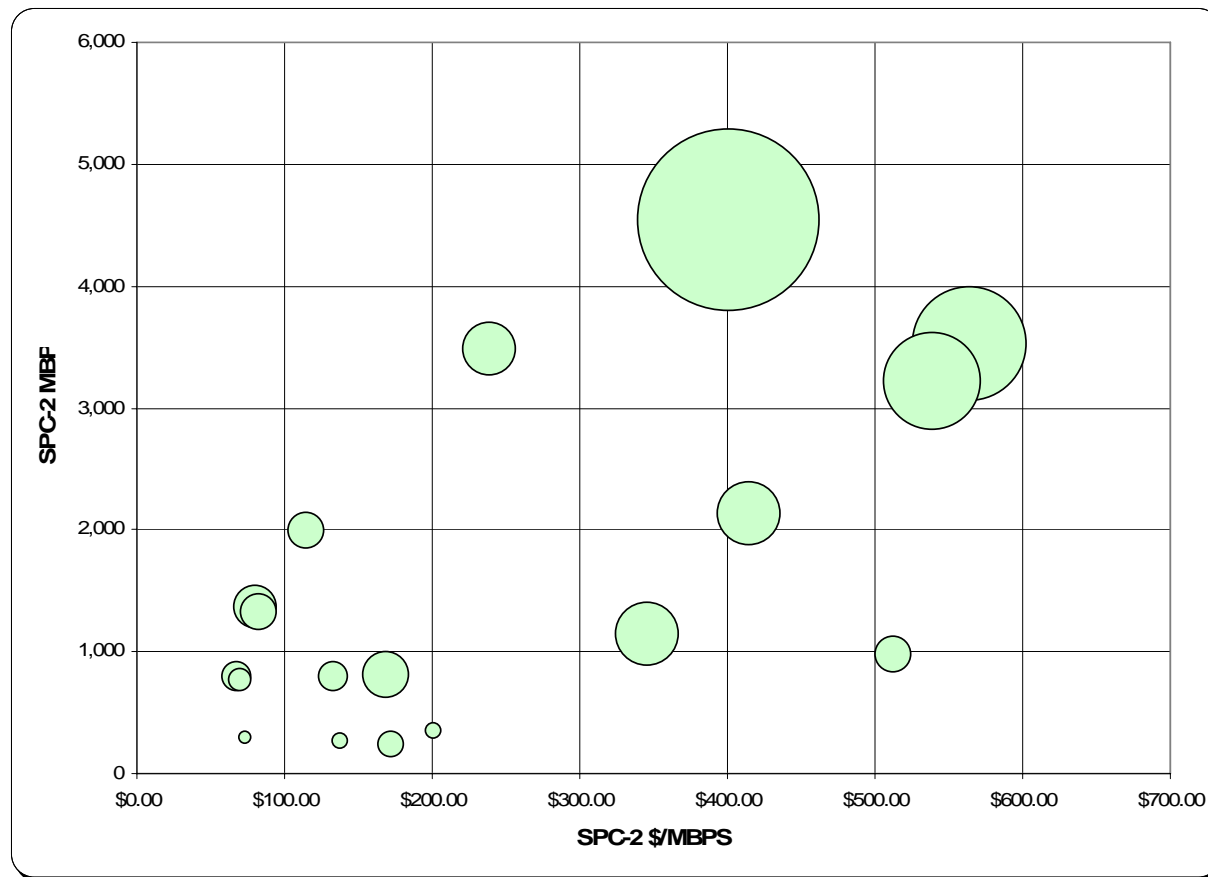
SPC-2 Overview

□ SPC-2 Workloads:

- **Large File Processing:** Simple sequential processing of one or more large files (*scientific computing/large-scale financial processing*).
- **Large Database Query:** Scans or joins of large relational tables (*data mining/business intelligence*).
- **Video on Demand:** Delivery of individualized video entertainment to a community of subscribers by drawing from a digital film library.

SPC-2 Overview

- SPC-2 Results continue to be submitted across a wide range of configurations.





SPC-2 Public Availability

- The SPC-2 Toolkit is currently available only to SPC member companies.

- Public availability of the SPC-2 Toolkit has been delayed due to both licensing and packaging issues for non-member distribution.
 - All licensing agreements have been completed and approved.
 - Packaging for non-member distribution and support is in final testing.



SPC-2 Public Availability

- The SPC-2 Toolkit will be available in two versions:
 - “Locked”: The SPC-2 workload parameters are all fixed at the values required by the SPC-2 specification.
 - “Unlocked”: A number of SPC-2 workload parameters may be varied by an end-user.
 - ‘Customized’ the SPC-2 workloads for internal measurements.
 - Parameters include: transfer sizes, read percentage, etc.



SPC-2 Public Availability

- The SPC-2 Toolkit will be available for non-member purchase in January 2008.
- Details for purchasing the toolkit will be posted on the SPC website prior to the public availability date.



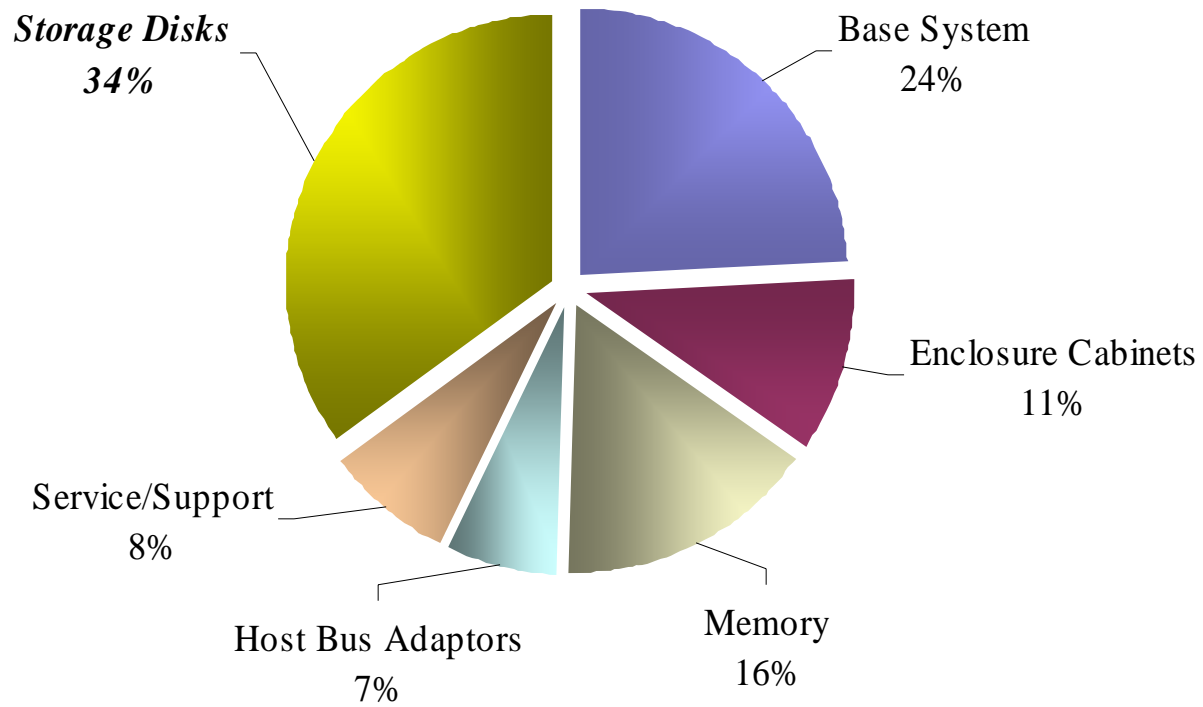
SPC-C: SPC-1C and SPC-2C

(component-level benchmarks)

Craig Parris
Seagate Technology LLC and SPC-C Chair



System Hardware Costs



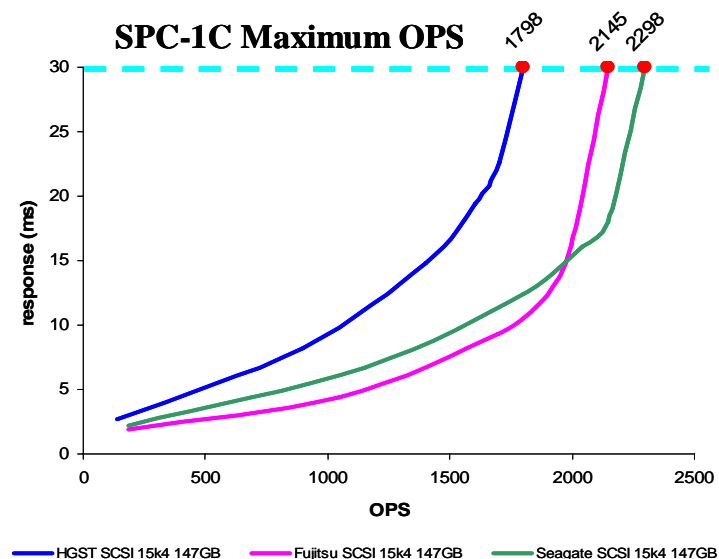

SPC-1C Re-cap

SPC-1C is comprised of a set of I/O operations designed to demonstrate the performance of a small storage subsystem while performing the typical functions of a business critical application.

SPC-1C represents a segment of applications characterized by predominately random I/O operations and requiring both queries as well as update operations (for example: OLTP systems, Database systems, or Mail Server applications).

SPC-1C focuses on Small storage solutions e.g. storage media devices, adapters

The Average Response Time, for the IOPS cannot exceed 30ms



SPC-1C Locality




Chart Courtesy of Microsoft Xperf

Example not to scale



SPC-1C/2C Configuration Groups

Single Storage Device



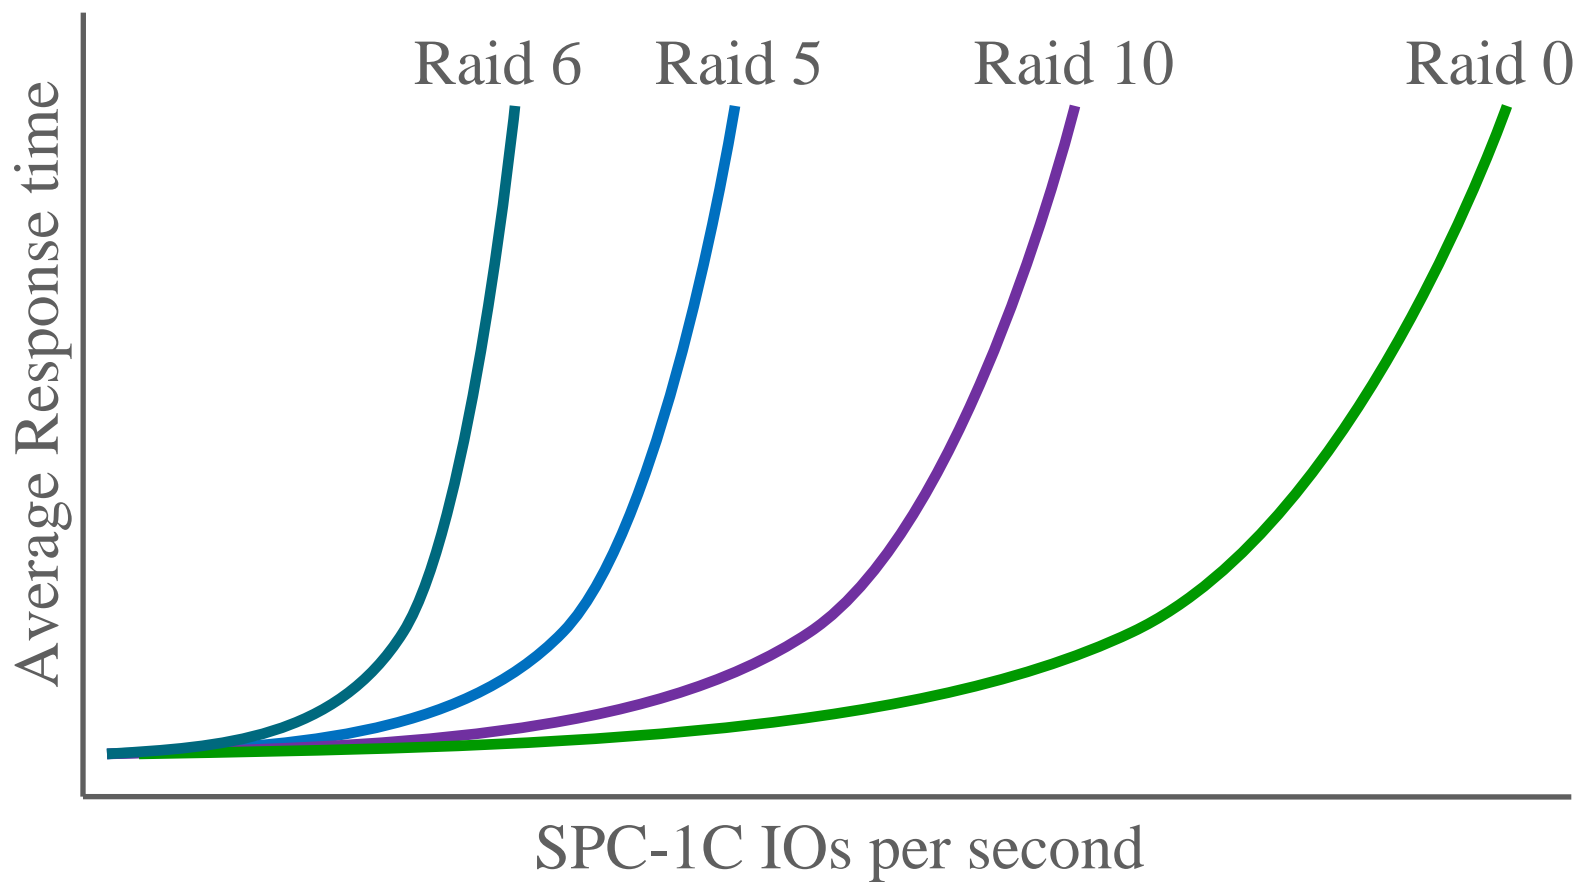
Four Storage Devices



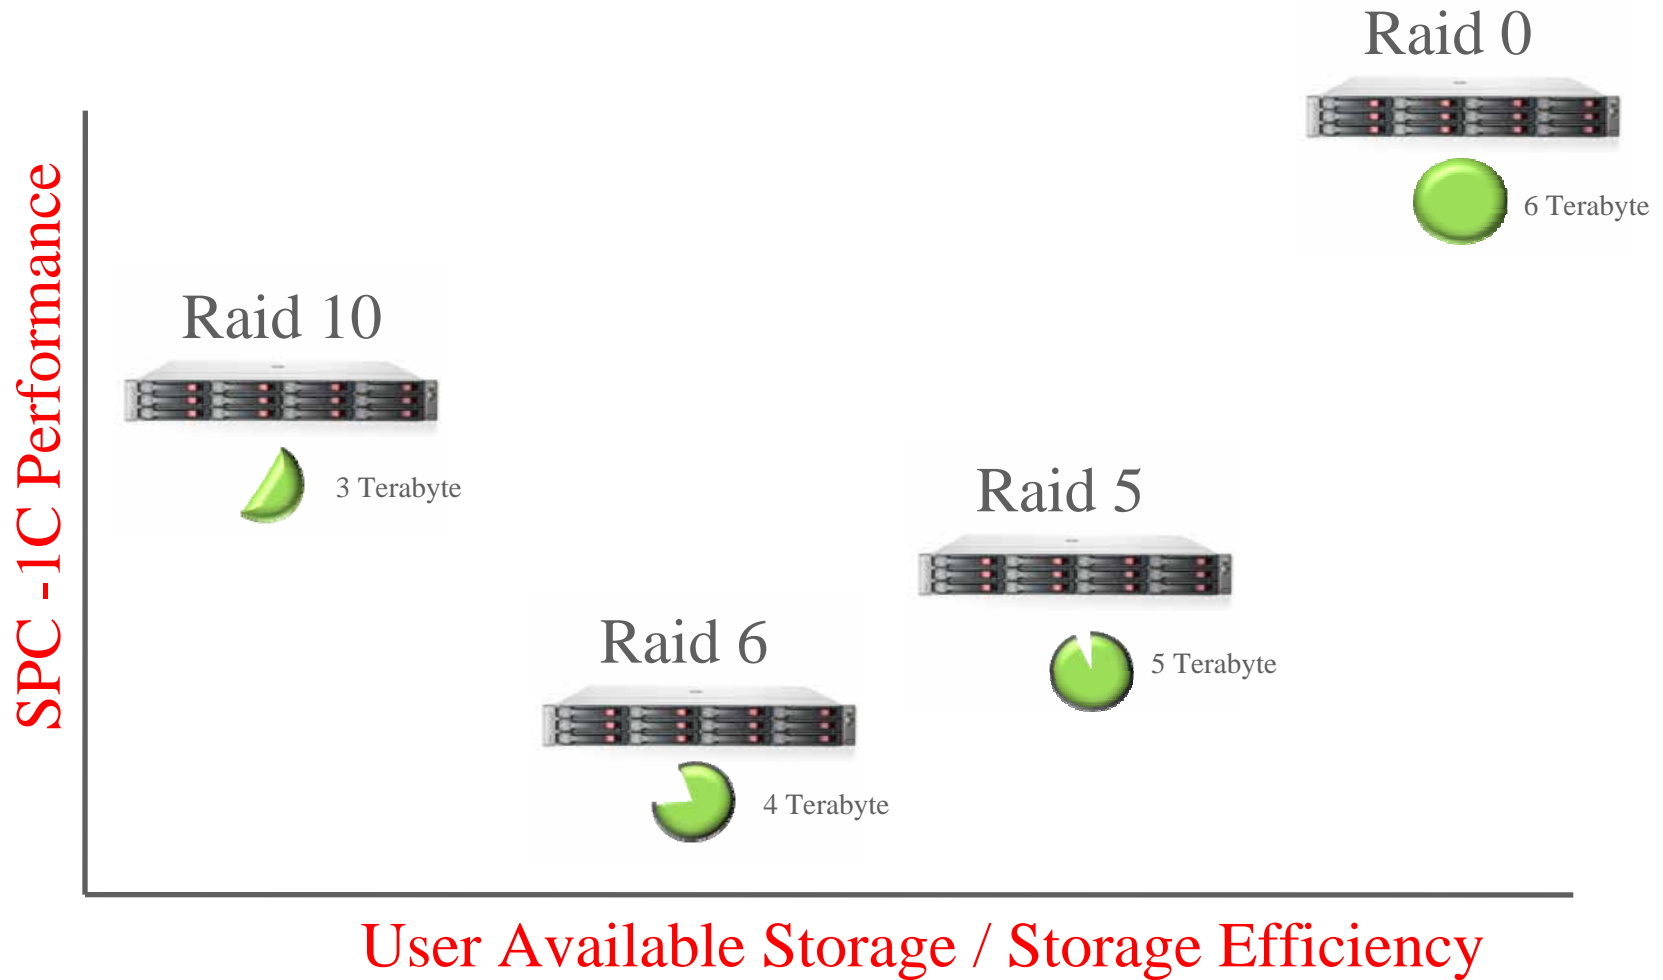
Small Storage Subsystem



Using SPC-1C



Using SPC-1C

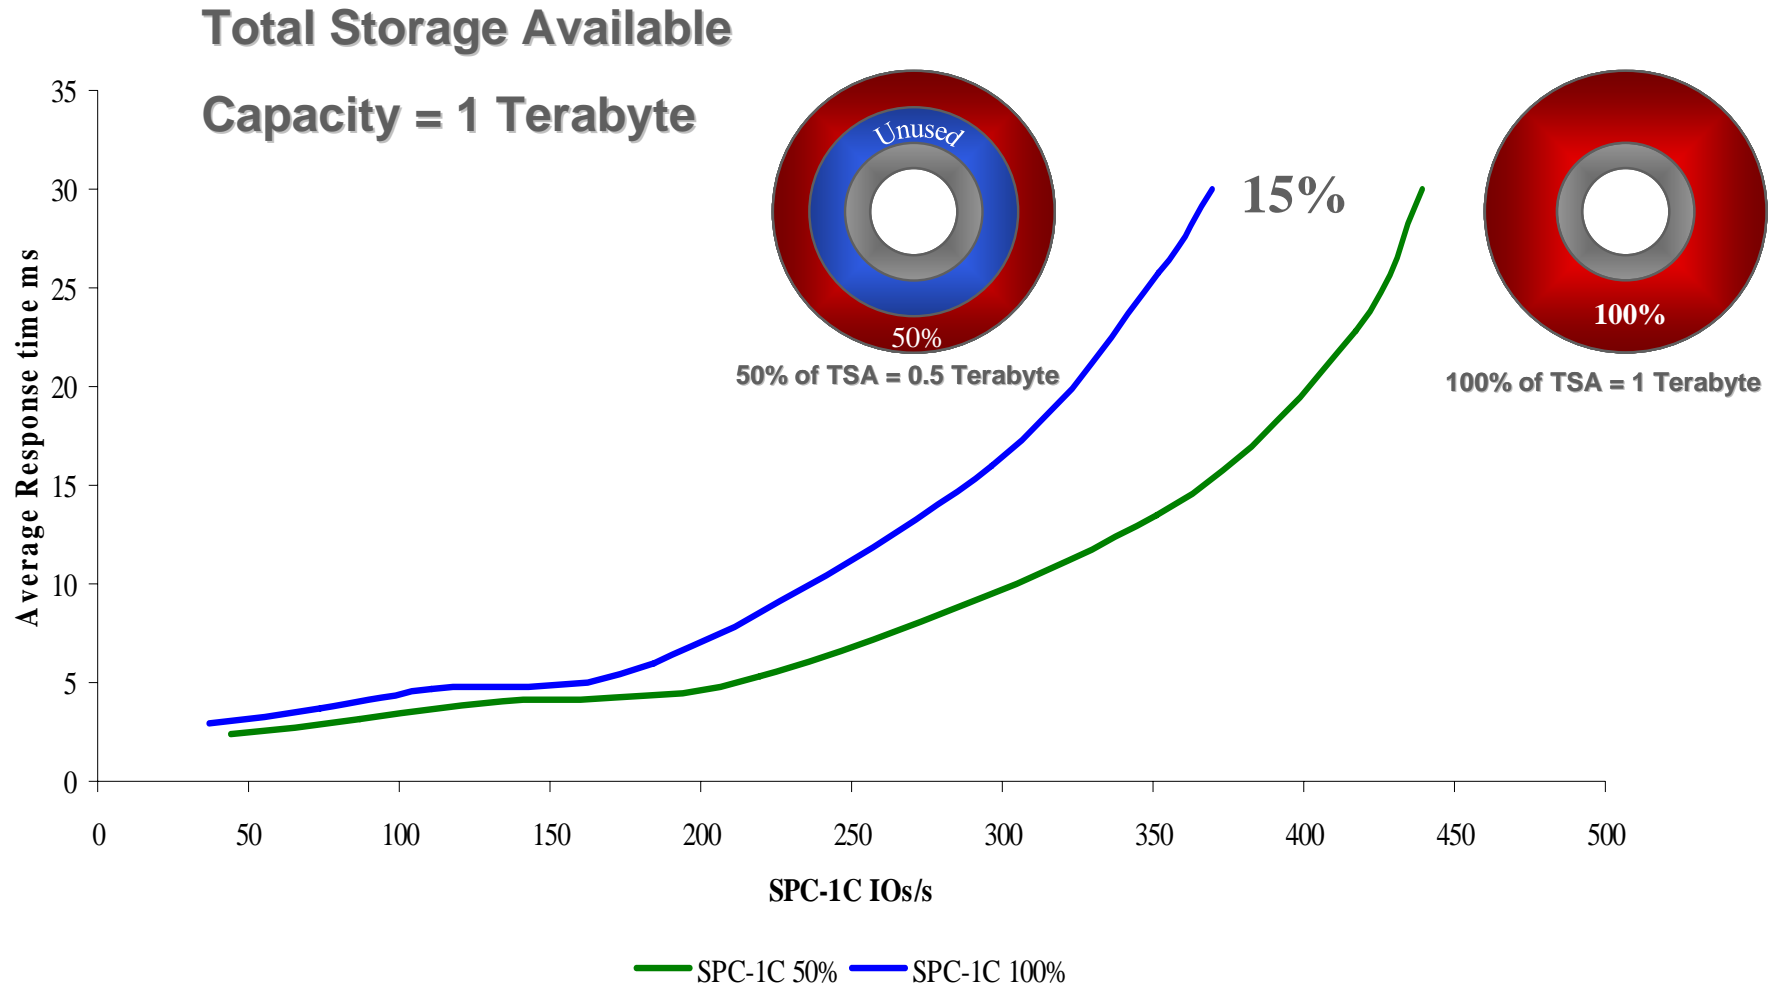




SPC-1C Differentiate Products



SPC-1C Capacity Example







SPC-2C Re-examined

SPC-2 is comprised of a set of I/O operations designed to demonstrate the performance of a storage subsystem running business critical applications that require large-scale, sequential movement of data.

Large file processing : represents large file read and write activity such as that encountered when processing large CAD files. Workloads reflect the behavior that occurs during file copy operations.



Large database queries: scans or joins of large relational tables, such as those performed for data mining or business intelligence.



Video on demand: individualized video entertainment provided to a community of subscribers, by drawing from a digital film library.



SPC-2C Locality

Large File Processing example 5 Streams


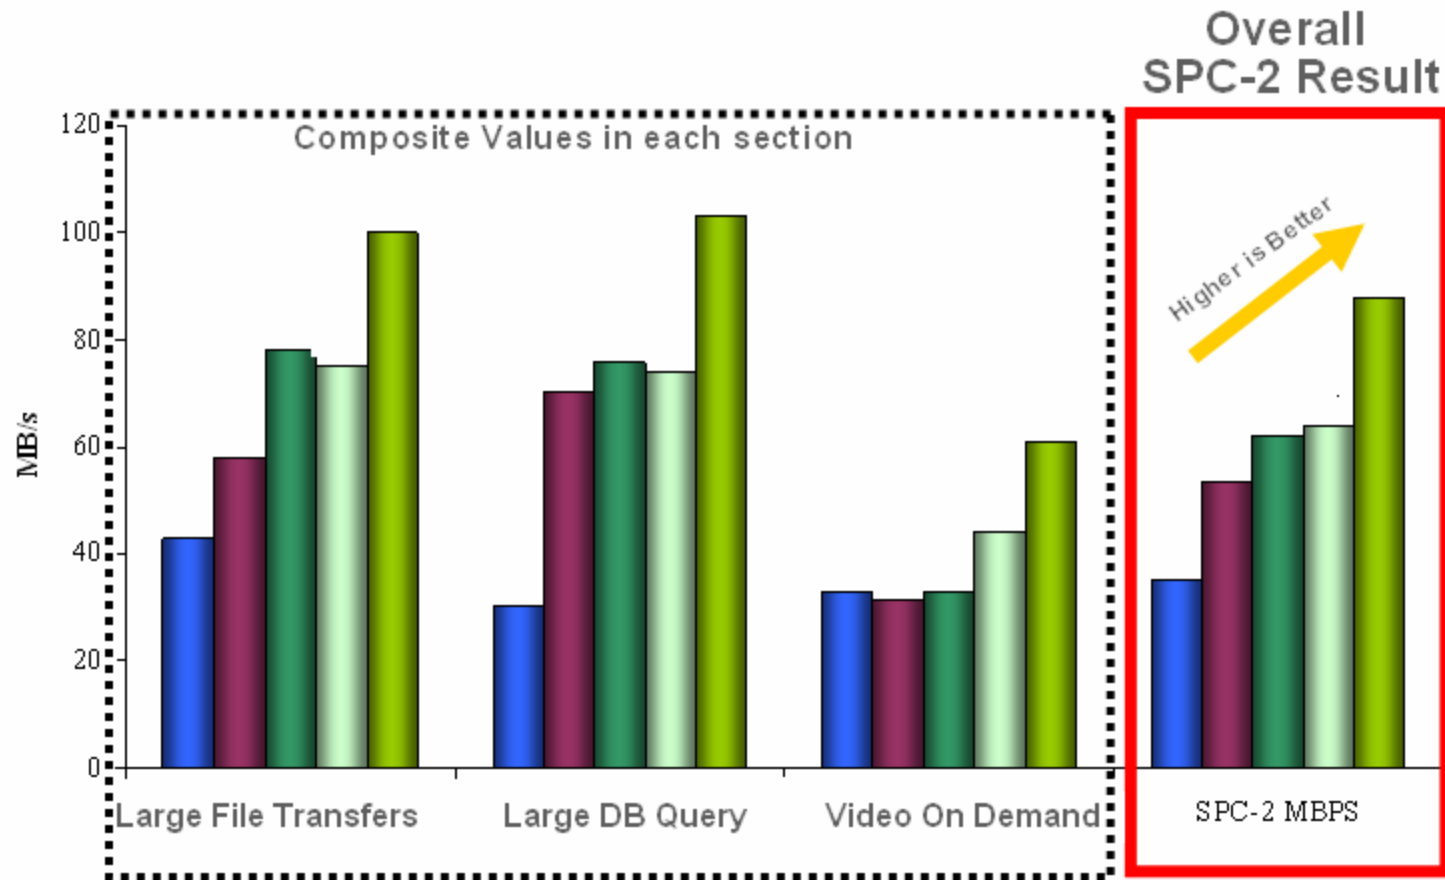




Chart Courtesy of Microsoft Xperf



SPC-2C Product Comparison



Large File Processing Product Comparison Example





SPC Green Power



Green Power Proposal

- Workload Generator

- Storage Performance Component benchmarks

- Performance metrics

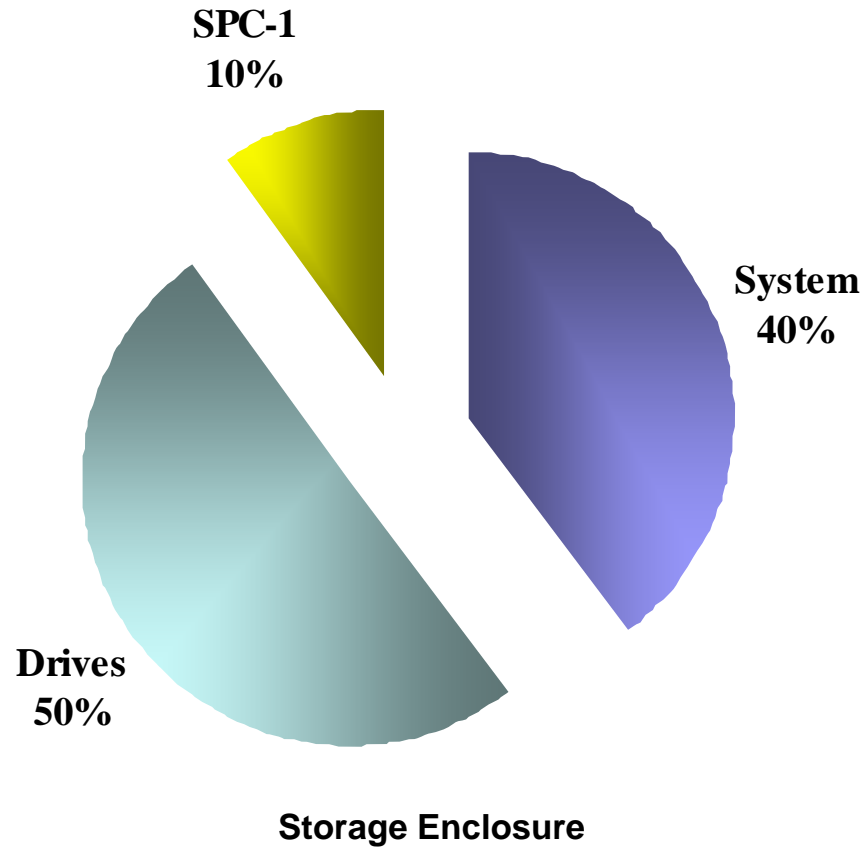
- SPC-1C IOPs per Watt

- Measure AC rms Watts

- Idle/SteadyState (no commands, drives @ full RPM ready to accept commands without delay)
- Maximum SPC-1C IOs/second (complex transactional workload)
- Maximum MB/s Large File Processing (throughput)
- Watt accuracy 2%
- Measured at Ambient room temperature 23°C +-5°C



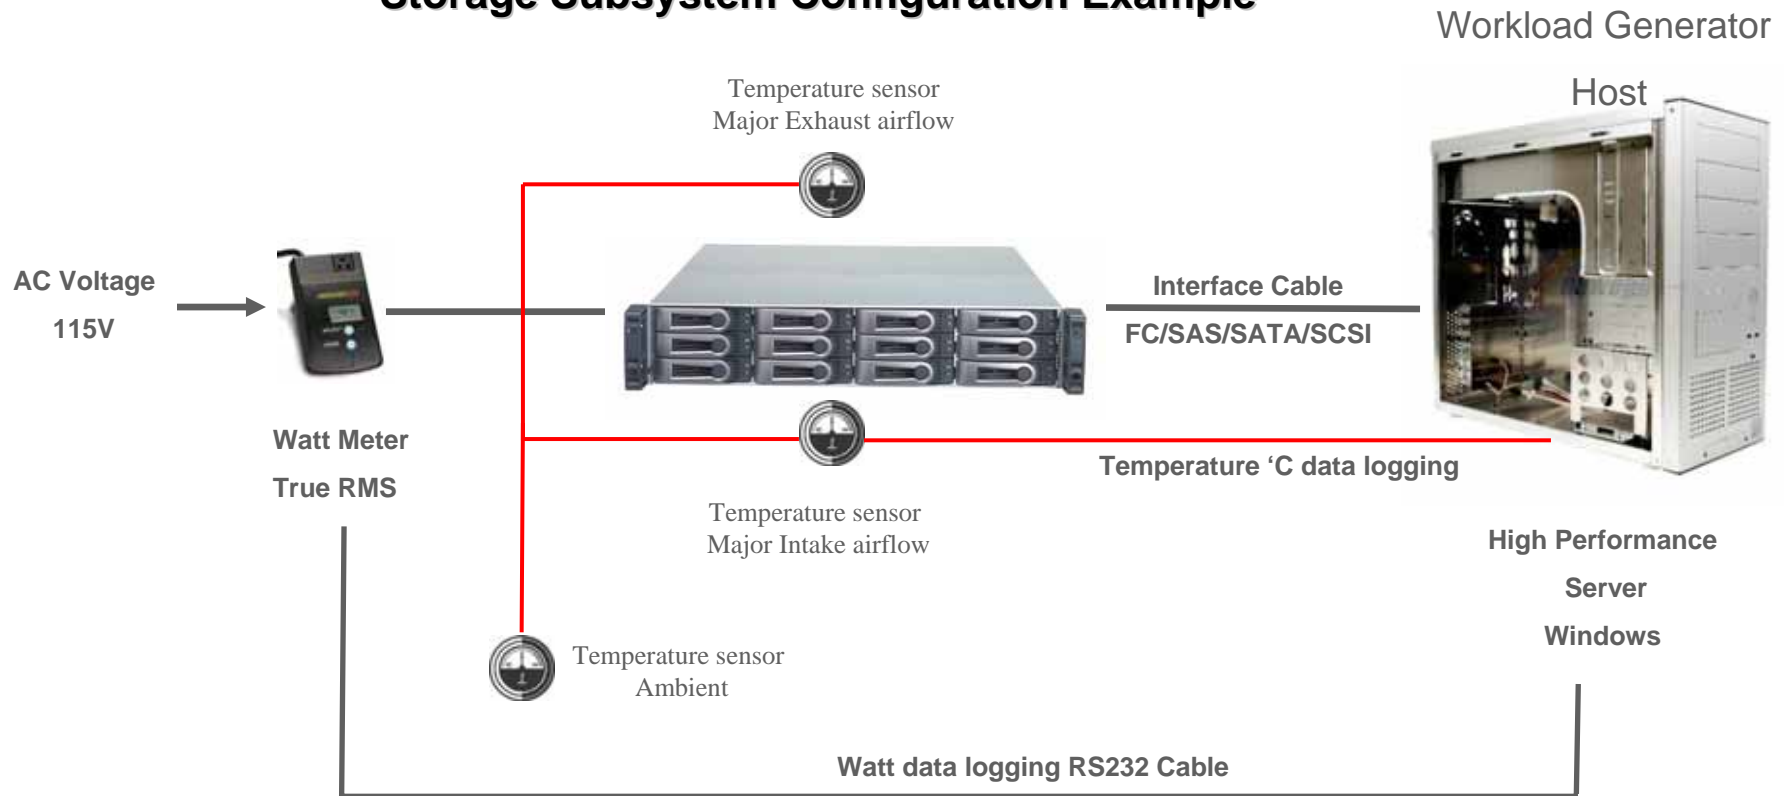
Green Power Proposal



System= Power Consumed without Drives
Drives = System Power with Drives – System Power without Drives
SPC-1 = (System + Drives + SPC1 100%) – (System + Drives)

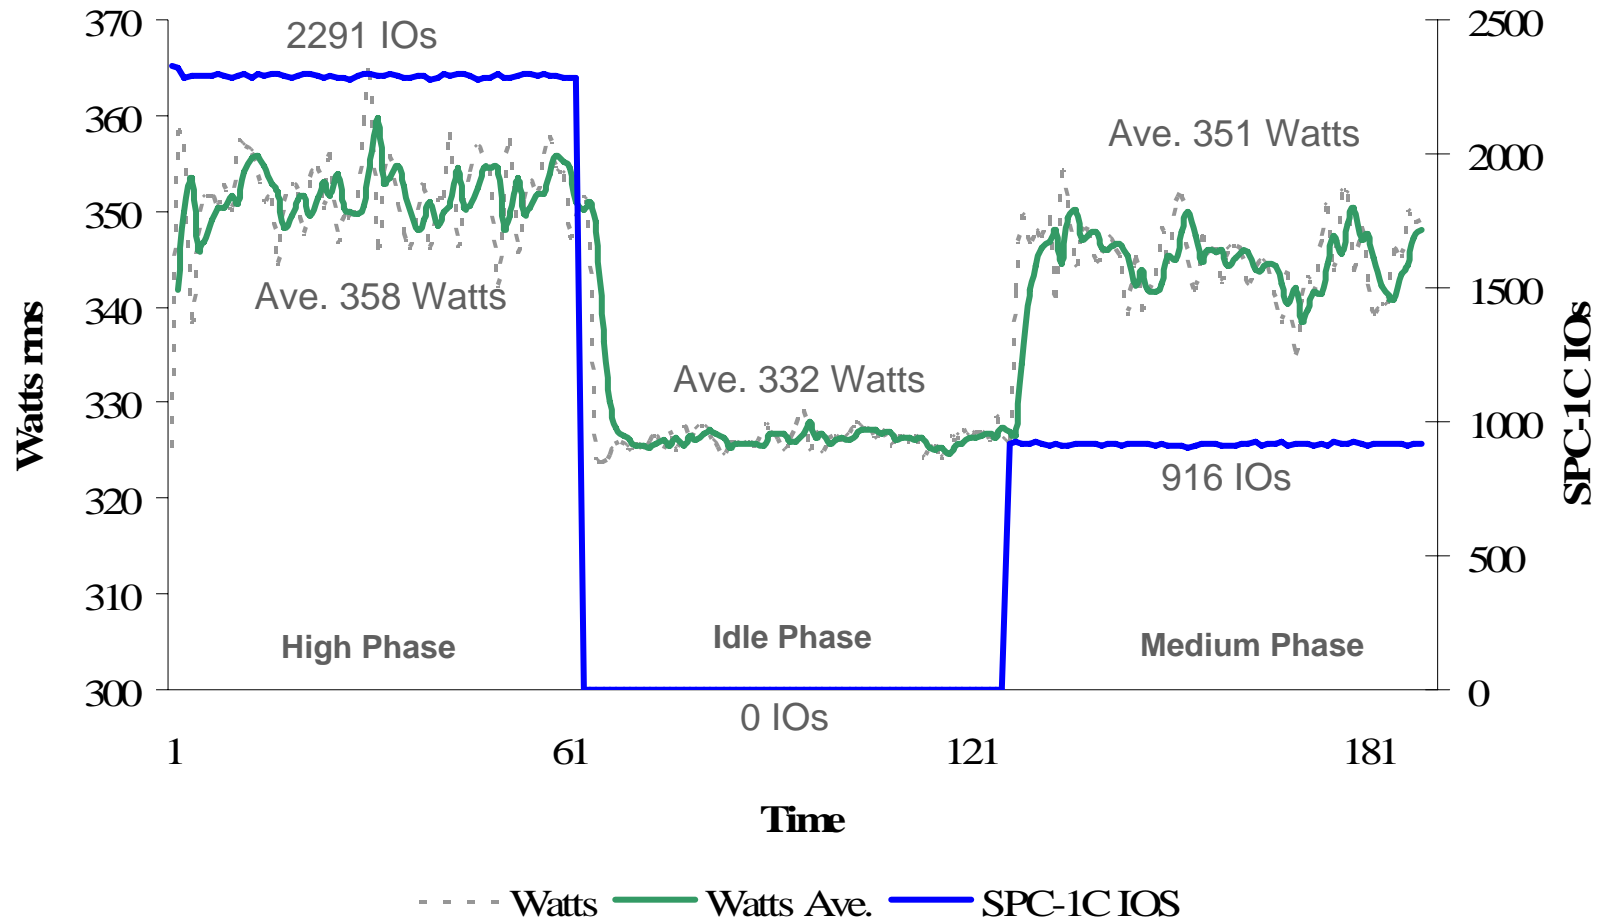
Green Power Proposal

Storage Subsystem Configuration Example



Green Power Proposal

Sub-System RAID





SPC-3

(A Family of File System Benchmarks)

Bruce McNutt
IBM Corporation and SPC-3 Chair

Jack Stephens
SPC Administrator



SPC-3

SPC-3 is a content model and related suite of benchmarks focused on the performance aspects of content management, including storage management, information lifecycle management and hierarchical storage management.



Goals for SPC-3

Publish objective test results for backup/restore performance.

Measure performance in the context of cost reductions such as

- Tiered storage
- Content management
- Hierarchical storage management.



SPC-3 Benchmarks

The first two planned “contexts” of SPC-3 are:

- ❑ SPC-3BR is a test of storage management (backup/restore) performance.
- ❑ SPC-3ILM is a test of content access performance in the presence of single-tier and multi-tier storage solutions.



Why Measure ILM Performance?

For the vendor

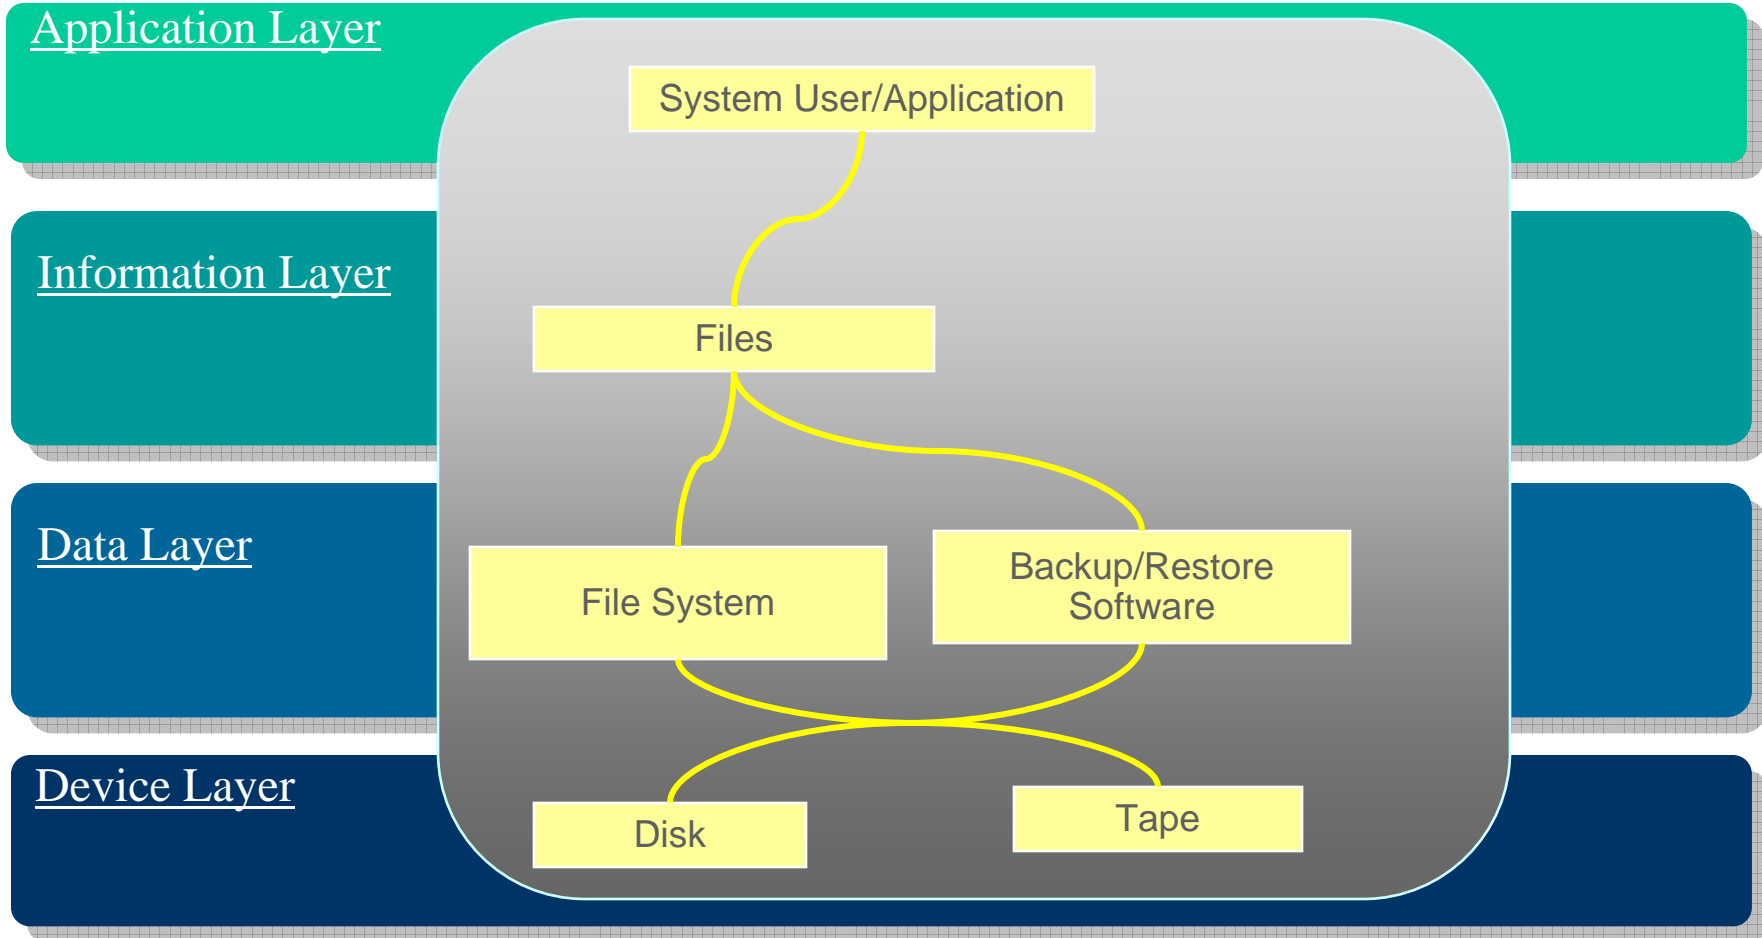
- Demonstrate that a solution provides savings *and* performance.

For the customer

- Find the right balance of these two goals.

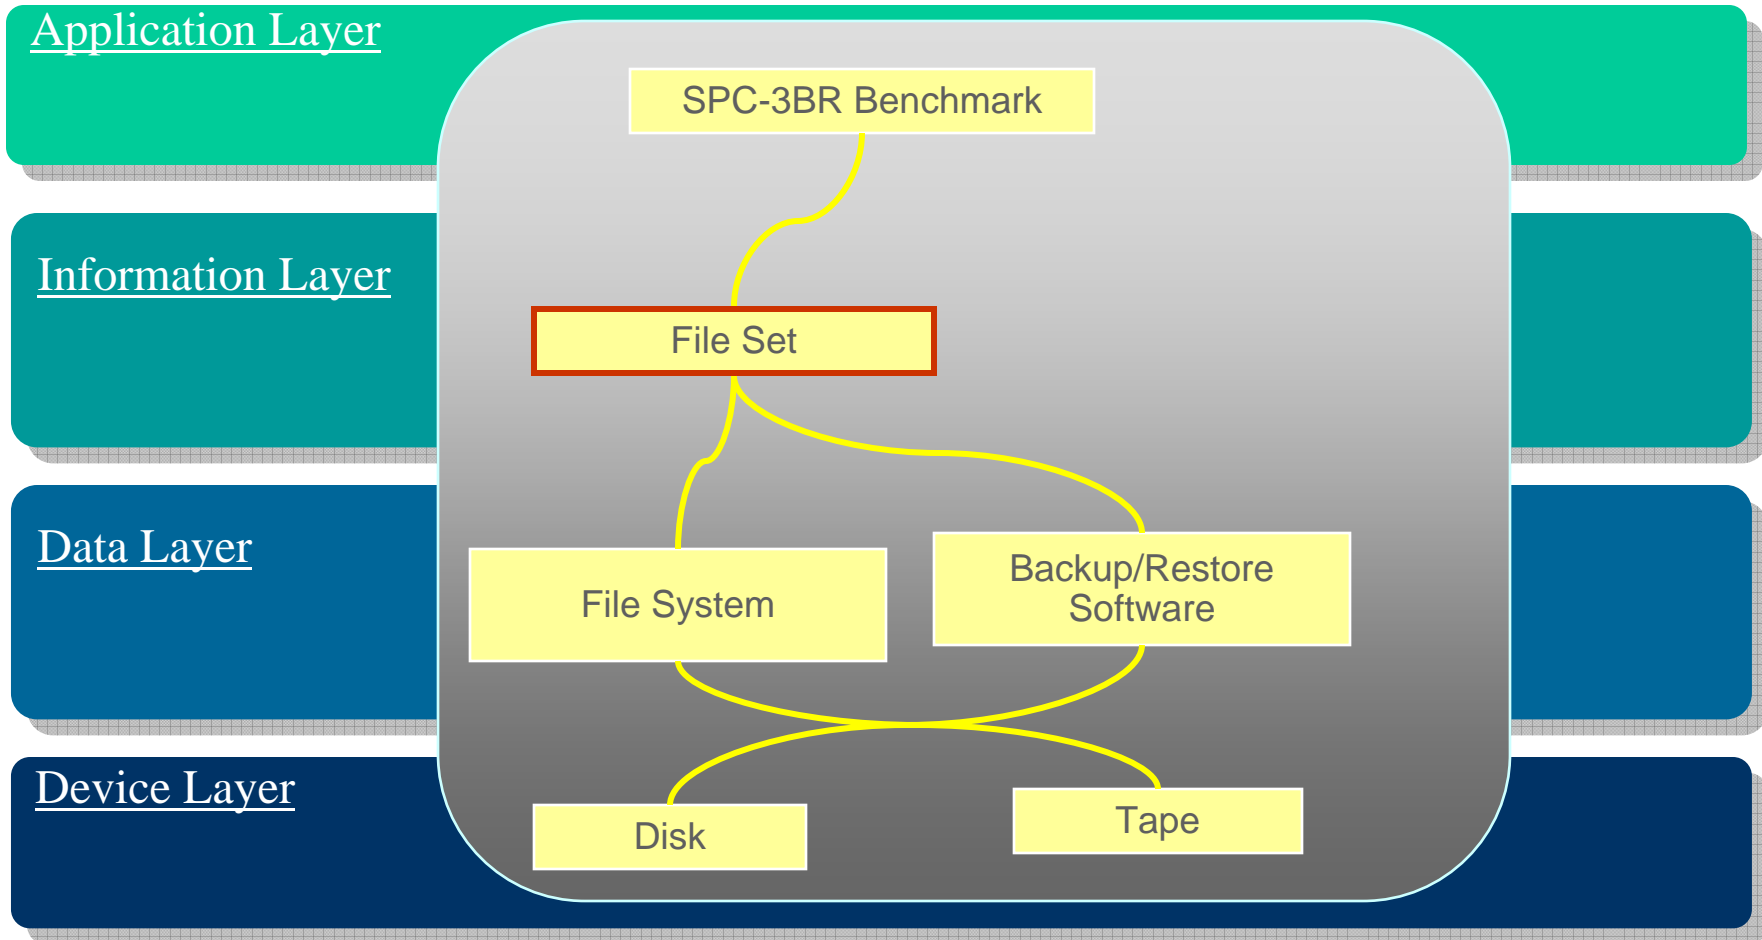


Simplest Picture (Storage Management)





SPC-3 Backup/Restore Test (Example Config)





The FileSet

Metrics are the message,

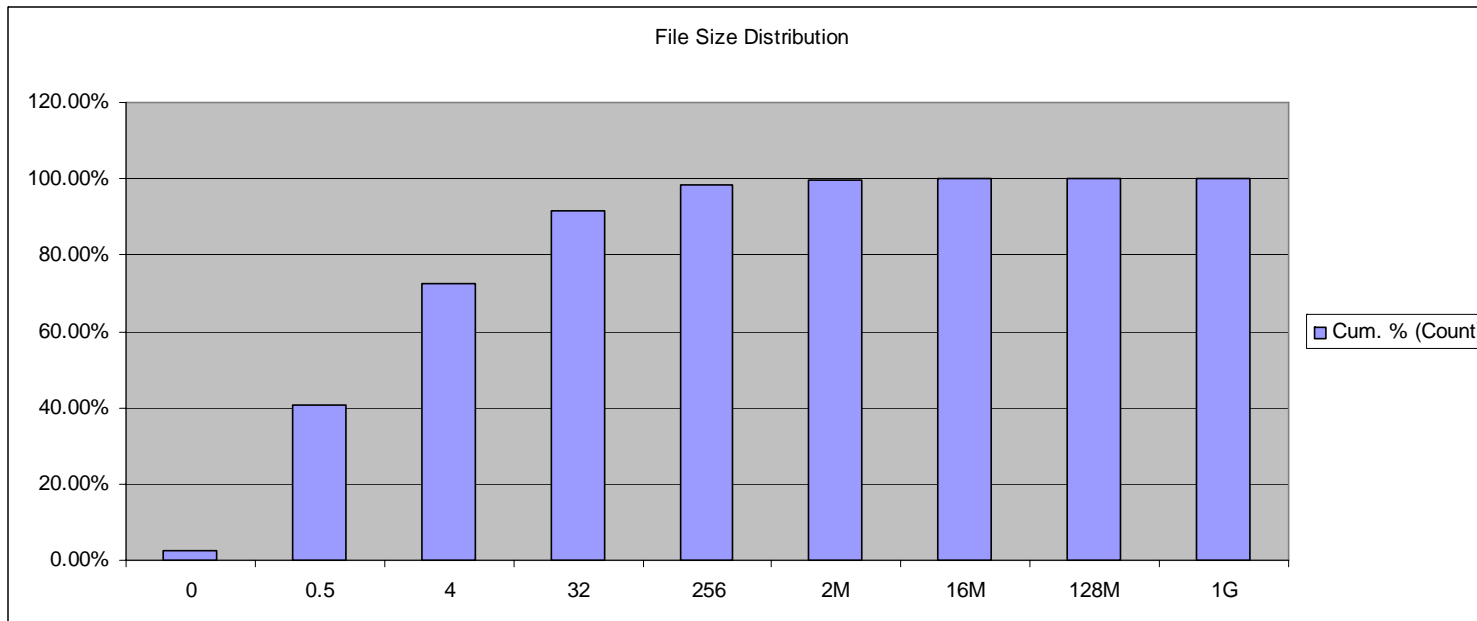
but content is the context.

- ❑ **SPC-3 is the first *content-aware* SPC Benchmark**
- ❑ **Data Set Scale Factors for Head to Head Comparison**
- ❑ **100GB to 100TB**



Realistically Synthetic

- ❑ Seven Important Attributes
- ❑ Realistic Clustering
- ❑ Expandable, Comparable, Reproducible





What's New is Old

Wine gets better with age...

...file systems don't!

- We age it, then we measure it.
- Platform independent approach
 - Fill to 85%
 - Randomly Expand to 95%
 - Contract to 80%
 - Recover to 85%



If We Build It...

Data Generator nearing completion:


- Running on 6 platforms
- Multi-threaded, multi-host (shortly)
 - big data sets *fast*
- Distribution-driven
 - easily tailored for experiments

The Benchmark

Room to maneuver...

... but nowhere to hide!

- ❑ Tests: Create and Restore a Backup
- ❑ Metrics: Throughput and Price/Performance for Backup and Recovery
- ❑ Integrity:
 - Full Disclosure
 - Peer Review
 - Audit





SPC-3BR Public Review

- Public Review Specification Available:
<http://www.storageperformance.org/specs>
- Comments
 - EMAIL: spc3-comments@storageperformance.org



SPC-3 Futures

**SPC-3BR draft spec released for review at
CMG 2007 -- **DONE****

**Data generator for the benchmark nearing
completion**

**Plans for roll-out will be developed based
upon review comments and results of early
testing.**



2008 SPC Benchmark Activity

Walter E. Baker
SPC Administrator and Auditor



2008 SPC Benchmark Activity

- Release of the SPC-2 Toolkit for non-member purchase
- Submission of SPC-1C/SPC-2C Results
- Completion and approval of the SPC-3BR specification
- Development of SPC Green Power metrics
- Review of the SPC-1 workload for possible updates/revisions