



Storage Performance Council (SPC)

www.StoragePerformance.org

SPCadmin@StoragePerformance.org





Presentations and Speakers

- ❑ ***Introductions, Moderator***
Mel Boksenbaum (*Hitachi Data Systems*)
- ❑ ***An Introduction to the SPC***
Leah Schoeb (*Sun Microsystems, Inc.*)
- ❑ ***SPC-3 Development Status***
Chuck Paridon (*Hewlett-Packard Company*)
- ❑ ***SPC-2 Update***
Bruce McNutt (*IBM Corporation & SPC-2 Chair*)
- ❑ ***SPC-1C / SPC-2C Introduction***
Craig Parris (*Seagate Technology LLC*)
- ❑ ***2007 SPC Benchmark Development, Q&A***
Walter E. Baker (*SPC Administrator & Auditor*)



An Introduction to the SPC

Leah Schoeb

Sun Microsystems, Inc.



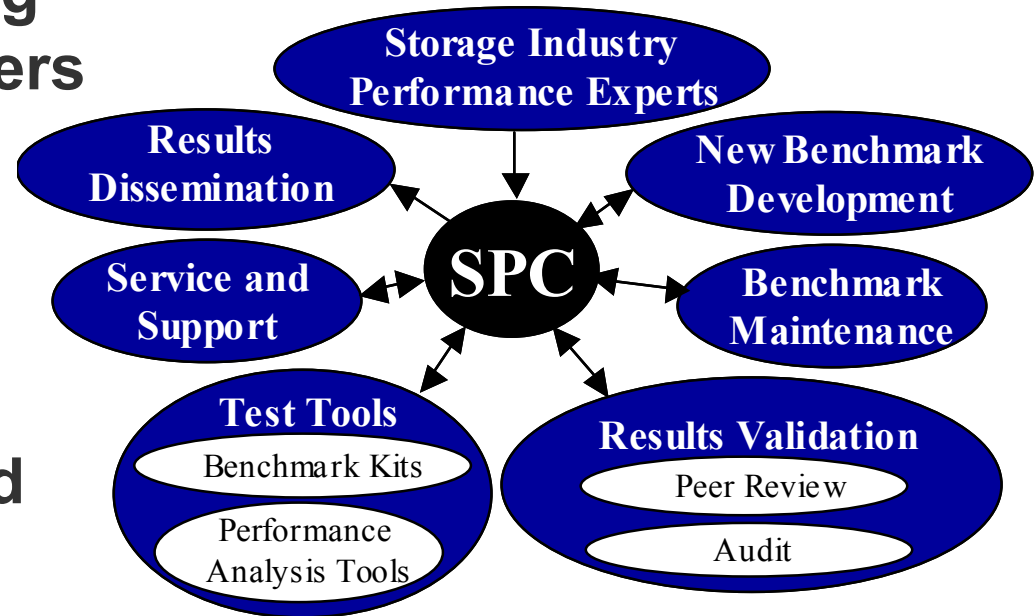
Council Mission

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

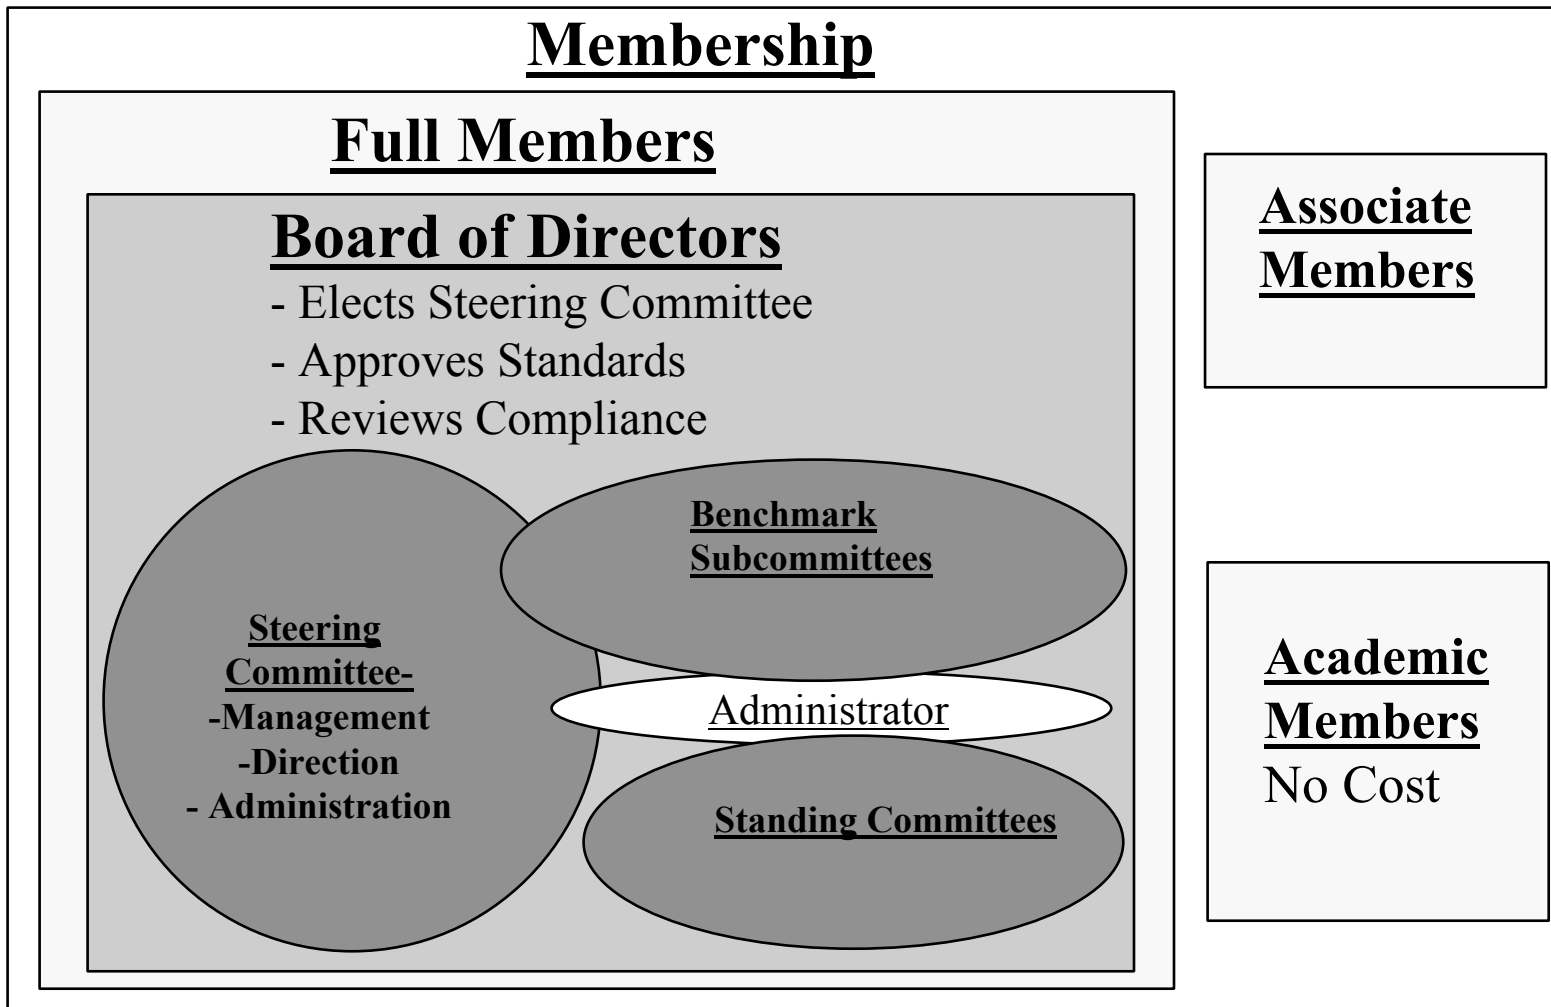
SPC Objectives

- 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-3 Development Status

Chuck Paridon
Hewlett-Packard Company



Description of effort

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 virtualized storage environment.
 - Storage from the top down



Environments and workloads

File Services

- ❑ **Models Corporate Document Repository**
 - Realistic File Set
 - Varied file types and sizes
 - “Aged” meta data
- ❑ **Two Workload Phases**
 - File Transfer (e.g., “Get” and “Put”)
 - File System Traversal (e.g., backup)
- ❑ **Two User Classes**
 - Interactive Session
 - Random walk of file system
 - Batch/Background
 - Models streaming data (e.g., on-line training)



Critical Questions

Is a measurement at the application level 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



SPC-2 Update

Bruce McNutt

IBM Corporation

Chair, SPC-2 Technical Subcommittee



SPC-2 Public Availability

- The SPC-2 Toolkit is currently available only to SPC member companies.
- Beginning January 8, 2007, the toolkit will be publicly available for purchase.
 - Details for purchasing the toolkit will be posted on the SPC website prior to the public availability date.



SPC-2: Large Block, Sequential I/O

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:


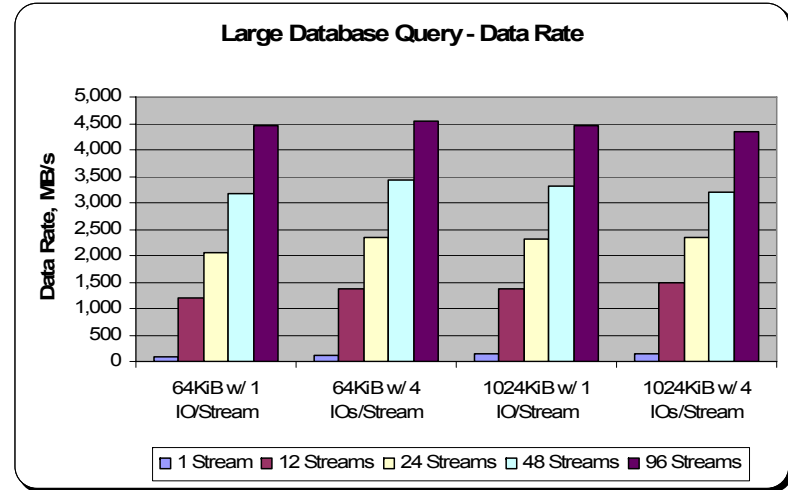

scans or joins of large relational tables.

Video on demand:


individualized playback from a digital film library.



SPC-2 measurements (Example)



Published SPC-2 Results





SPC-2 Overview

- ❑ **Test of storage system performance for environments that require large scale data movement.**
- ❑ **Developed using a series of application-oriented sequential workloads.**
 - Large File Processing (*LFP*)
 - Large Database Query (*LDQ*)
 - Video on Demand (*VOD*)



SPC-1C / SPC-2C Introduction

Craig Parris
Seagate Technology LLC



Storage Component Performance

- There is a distinct need to provide objective and verifiable performance measurement and comparison of storage components such as disk drives, HBAs/controllers, small storage subsystems (*single enclosure*), etc.

- SPC-1 and SPC-2 provide that type of objective and verifiable performance measurement and comparison, but are typically used for large, complex storage configurations.



SPC-1 and SPC-2

- An SPC-1 or SPC-2 result could be produced with a 'small' configuration to isolate storage component performance, but:
 - The SPC-1 and SPC-2 reporting and pricing requirements would be difficult to meet in many cases.
 - A 'small' SPC-1 result would be "buried" among the larger results.
 - It would typically be cost-prohibitive to produce such a 'small' configuration result.



SPC-1C and SPC-2C

- SPC-1C and SPC-2C will be the first industry-standard benchmarks to provide objective and verifiable performance measurement and comparisons for individual storage components such as disks, HBAs/controllers, small storage subsystems (*single enclosure*), storage software, etc.



SPC-1C and SPC-2C

- ❑ SPC-1C will be based on the SPC-1 specification and will utilize the SPC-1 workload.
- ❑ SPC-2C will be based on the SPC-2 specification and will utilize the three SPC-2 workloads.
- ❑ Reporting requirements will be similar to SPC-1 and SPC-2.
- ❑ Pricing requirements will be simplified to reflect the focus on storage components.



SPC-1C/SPC-2C Producer Audience

- Storage component vendors comprise the primary producer audience for SPC-1C and SPC-2C.
- SPC-1C and SPC-2C will provide performance evaluation and comparison for storage components such as:
 - Disk drives
 - HBAs/controllers
 - Small Storage Subsystems (*single enclosure*)
 - Processors used in the above components
 - Storage software such as Logical Volume Managers



SPC-1C/2C Consumer Audiences

- Storage solution providers (*vendors, resellers, etc.*), as well as end users, are the primary consumer audiences for SPC-1C and SPC-2C.
- SPC-1C and SPC-2C will allow storage solution providers to evaluate the storage components that comprise the larger storage solutions offered by the provider.
 - Disks used in an array
 - HBAs/Controllers
 - Small Storage Subsystems
 - Logical Volume Managers



SPC-1C/SPC-2C Configurations

- ❑ SPC-1 storage configuration requirements do not actively encourage “head-to-head” storage product comparisons.

- ❑ SPC-1C and SPC-2C configurations will be constrained to provide more direct comparison between competing storage component products.



SPC-1C/SPC-2C Configurations

- SPC-1C and SPC-2C configurations will consist of the following two components:
 - One HBA/Controller (single/dual/quad...).
 - One of the following disk drive configurations:
 - One, two, or four disks in a “standalone” configuration.
 - » An external enclosure may be used but only to provide power and/or connectivity for the disks.
 - A small storage subsystem (maximum 3U form factor) with no limit on the number of disks.



HBA/Controller Requirements

- SPC-1C HBA/Controller component
 - This component will consist of a single external adapter that plugs into a Host System or an HBA/Controller embedded in the motherboard of a Host System.
 - Controller functionality may reside in an external enclosure and if that is the case, the enclosure may contain more than one controller.



SPC-1C/2C One Disk Configuration

- ❑ A test of maximum contention.
- ❑ Allows direct comparison between disk products with minimal interpretation of results.
- ❑ One of the basic “building blocks” for a storage solution.



SPC-1C/2C Two Disk Configuration

- ❑ An expansion of the one disk configuration.
- ❑ Allows performance comparisons between software and hardware implementations of storage functionality (mirroring, striping, etc.) in the smallest possible configuration.
- ❑ Allows direct comparison between disk products in a small, easily understood configuration with minimal interpretation of results.



SPC-1C/2C Four Disk Configuration

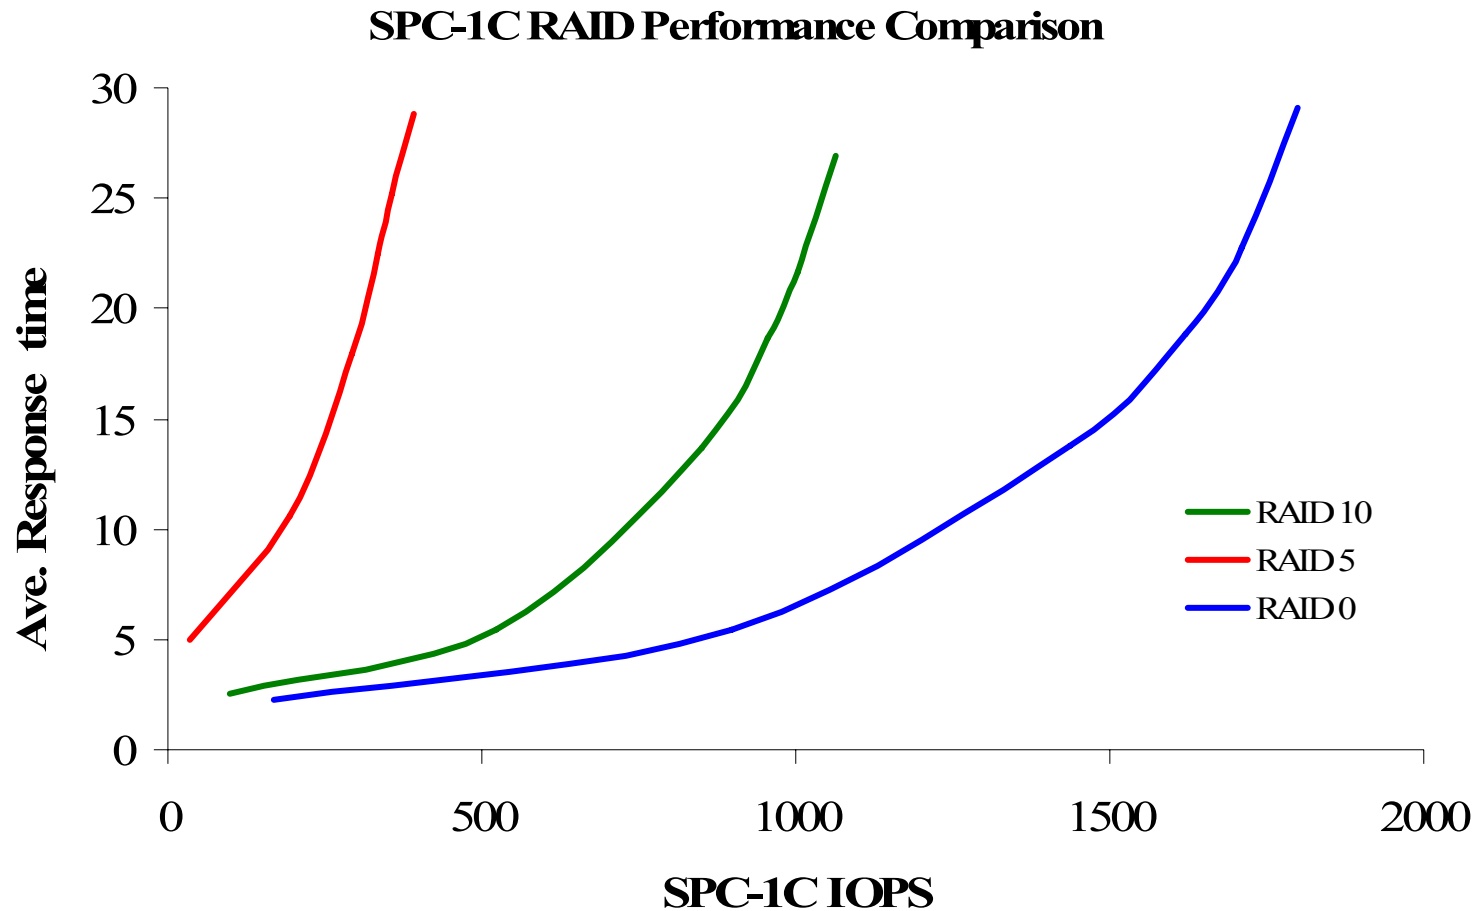
- ❑ Focused on HBA/Controller and disk performance to allow direct comparison of those products.
- ❑ Will allow performance comparisons between software and hardware implementations of storage functionality (mirroring, RAID, etc.).
- ❑ Will allow scaling comparisons with two disk configurations.



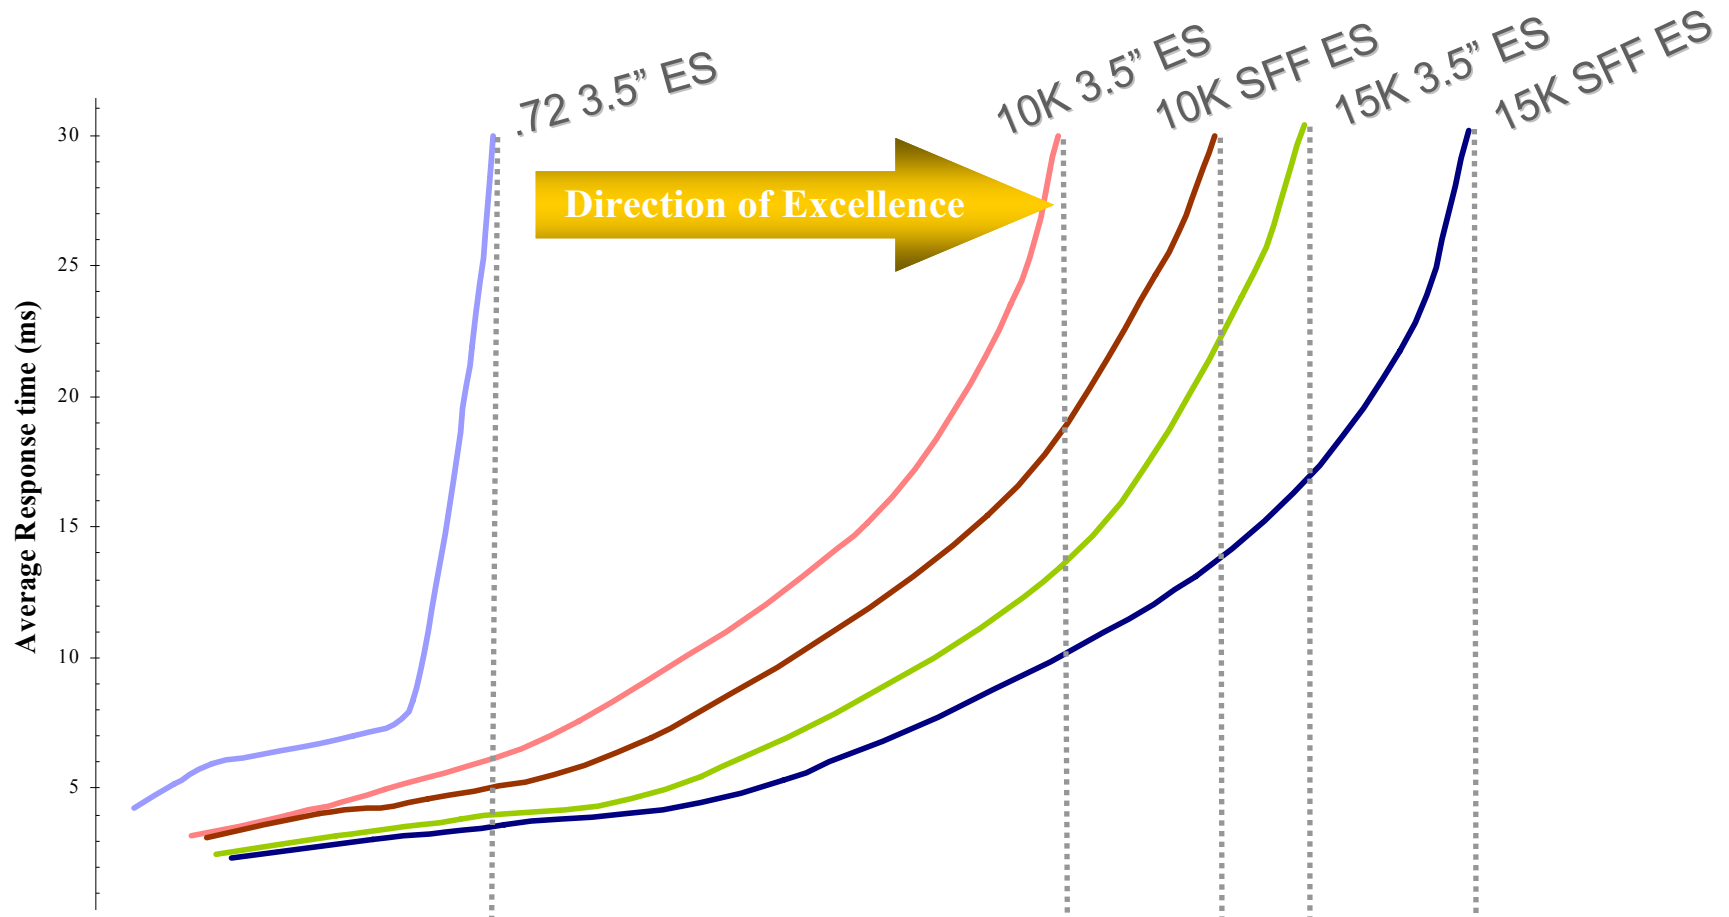
SPC-1C/2C Small Storage Subsystem

- A single enclosure (maximum 3U form factor).
 - No limit on the number of disks in the enclosure.
 - Focused on HBA/Controller and small storage subsystem performance to allow direct comparison of those products.
 - Will allow performance comparisons between software and hardware implementations of storage functionality (mirroring, RAID, etc.) in a 'small' array configuration.
 - Will allow comparison of enclosure functionality and performance.

SPC-1C/2C Prototype Data



SPC-1C/2C Prototype Data





2007 SPC Benchmark Development, Q&A

Walter E. Baker
Gradient Systems, Inc.
SPC Administrator and Auditor



2007 SPC Benchmark Development

- **Workload + Context + Infrastructure = Benchmark**
 - SPC-1 and SPC-2 workloads cover a significant amount of the storage performance “space” (I/O behavior and market).
 - SPC-3, when released, will further extend that coverage.
 - Additional workloads will add minimal coverage.
 - Context provides meaning and perspective to a benchmark for both a Test Sponsor and end-user.
 - Context is relatively easy to define.
 - ”chassis” (workloads) and ”models” (contexts) analogy



2007 SPC Benchmark Development

□ **Workload + Context + Infrastructure = Benchmark**

- There are multiple contexts currently identified, which are appropriate for SPC-1 and SPC-2 to create timely, new benchmarks.
 - SPC-1C, SPC-2C: Storage component performance
 - SPC-1F: File system performance
 - SPC-Multi-Address Space: Storage consolidation
 - Data Services: Replication, Snapshot, etc.



2007 SPC Benchmark Development

□ **Workload + Context + Infrastructure = Benchmark**

- The existing infrastructure can be (re)used for new benchmarks.
 - Pricing model
 - Audit
 - Full Disclosure Reporting
- Again, leveraging prior SPC work.



SPC Q&A

- ❑ Why would non-members purchase the SPC-2 Toolkit?
- ❑ How would VARs and resellers use the information provided by SPC Results?