



October 12-14, 2004

File System Suite of Benchmarks

John Corbin

President

EP Network Storage Performance Lab

jcorbin@nsplab.com



October 12-14, 2004

Overview

- File System Benchmark Types
- File System Suite of Benchmarks
- NFS Client Benchmark
- Workload Capture Tools
- Product Plans



October 12-14, 2004

File System Benchmark Types



October 12-14, 2004

What Type of Benchmark?

- System Level
- Component Level
- Single Benchmark
- Suite of Benchmarks



October 12-14, 2004

System Level Benchmark

- Protocol Independent
- OS Independent
 - Work on UNIX and Windows[®]
- Scalable
 - Across Multiple Computers
 - Across Multiple CPUs per Computer
- Compare System Level Performance
 - Local File System vs. NAS
 - NFS vs. CIFS



October 12-14, 2004

Component Level Benchmark

- Protocol Dependent
 - One for NFS, One for CIFS, ...
- Scalable
- Compare Component Performance
 - Ex: NFS Server Implementations



October 12-14, 2004

Benchmark Suite

- **System vs. Component**
 - End Users Want System Benchmarks
 - Developers Want Component Benchmarks
 - Mixture of EU/VAR/Developers In-Between
- **Deliver Suite with System and Component Benchmarks**
 - Single Tool
 - Single Interface



October 12-14, 2004

File System Suite of Benchmarks

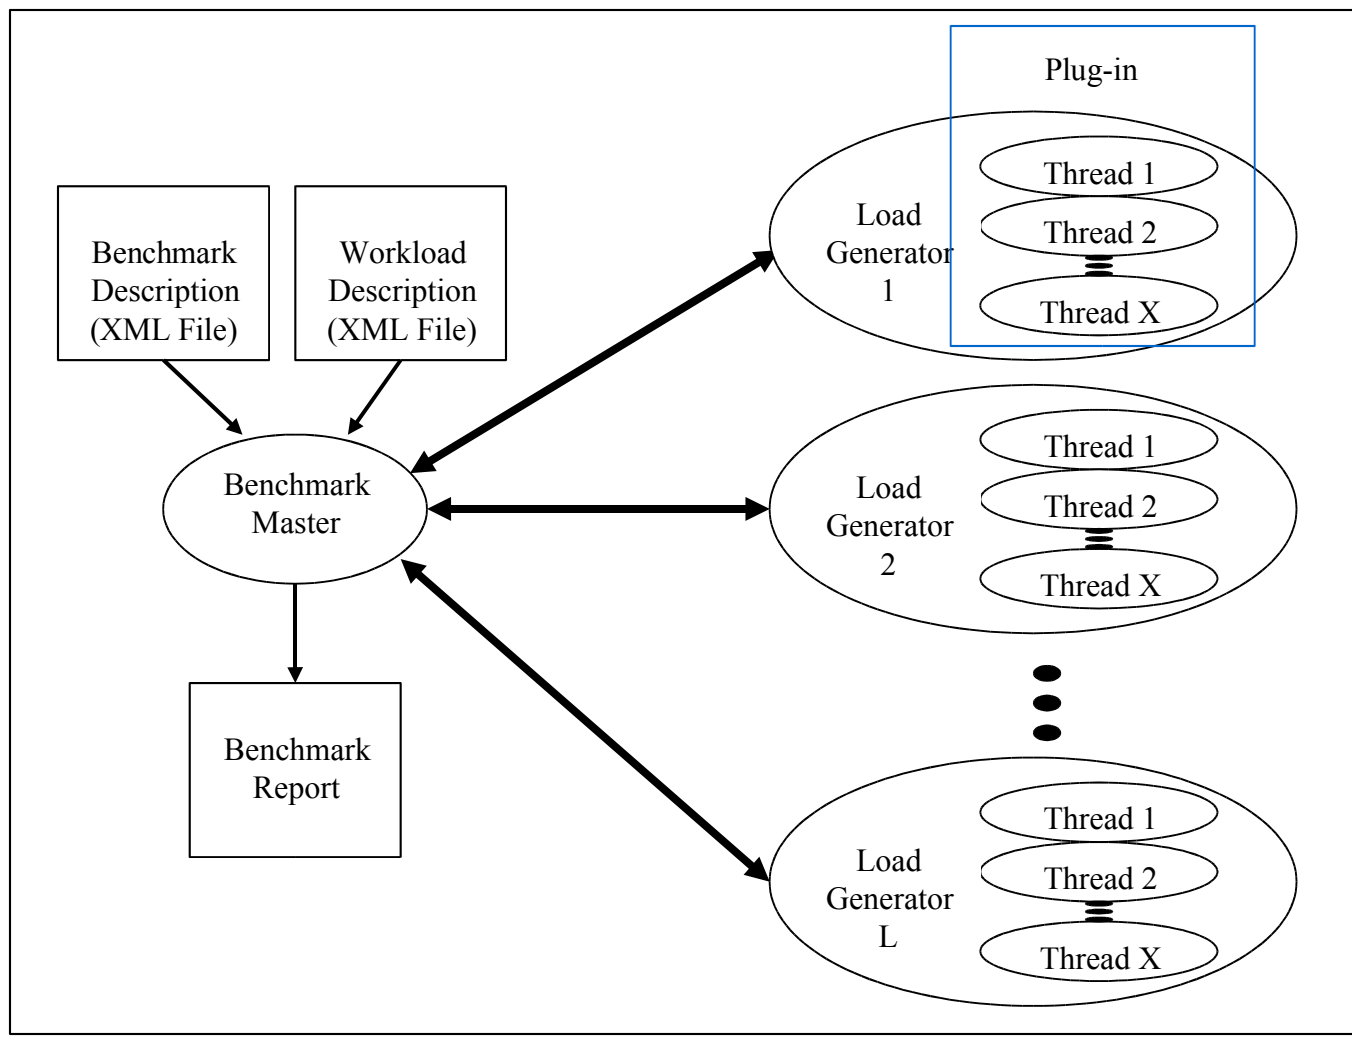
Technical Report TR-2004-001

www.nsplab.com



October 12-14, 2004

Architecture





October 12-14, 2004

Benchmark Master

- Process Benchmark and Workload Description Files
- Spawn the Load Generators (LGs)
- Reliably Steps LGs Through Benchmark Phases
- Report Results



October 12-14, 2004

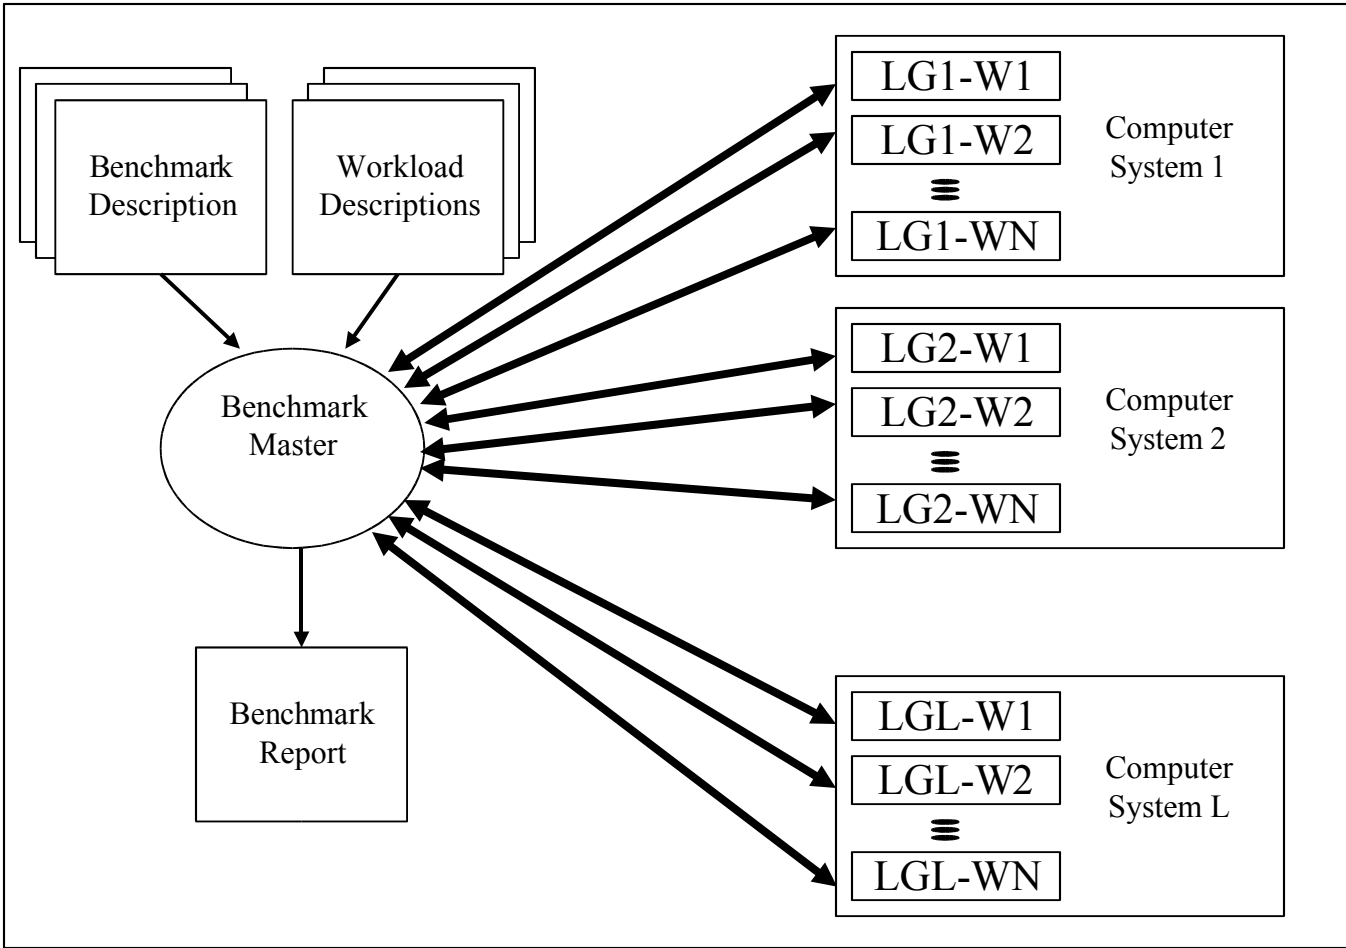
Load Generator

- Thread-based
- Execute the Benchmark Phases
- Generate Operations
- Plugin Operation Specific Modules
 - Examples:
 - NFS Server Plugin
 - POSIX File System Plugin



October 12-14, 2004

Support Multiple Concurrent Workloads





October 12-14, 2004

Operation Clustering

- Chain Operations Together
 - Subsequent ops can use parameters returned from previous ops in the chain.
- Ops Execute Sequentially
 - Good for open, stat, malloc, read, close
 - Not good for mkdir A, create 50 files in parallel in directory A



October 12-14, 2004

Op Scheduling

- Ops Scheduled on Op Queue
- LG Threads Pull Ops Off Queue
 - Op Executed Now or in the Future
- Requires Sequence Points
 - Mkdir A, Create File in Dir A
- Allows Easy Creation of Op Streams
 - Read Ops From Disk
 - Read Ops From Network



October 12-14, 2004

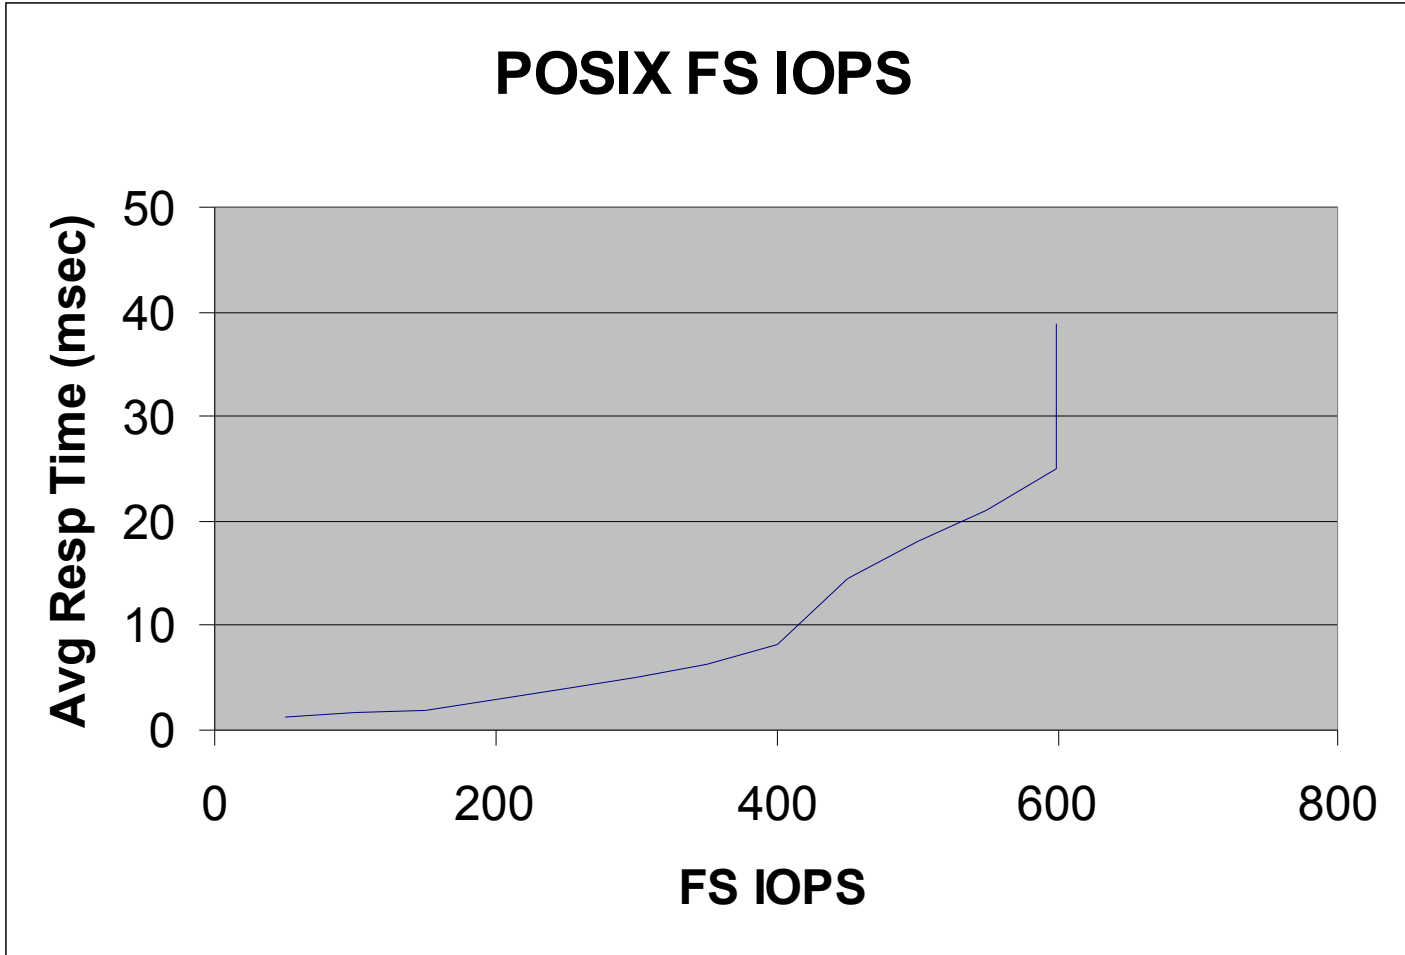
Benchmark Validation

- Validation Features
 - Op Tracing
 - No-Op Execution Flag
- Validate With Workload Capture
- Validate on Different Platforms



October 12-14, 2004

Sample Run





October 12-14, 2004

NFS Client Benchmark



October 12-14, 2004

Requirements

- OS Independent
- Minimal Code Running on Client
- Free of Server-side Effects
 - Client Needs to Send Requests to Server
 - Server on Other Side of Wire
 - Results Independent of Server
- Minimal HW Resources
- Measure NFSOP Efficiency



October 12-14, 2004

Implementation

- Single Client / Multiple Servers
- POSIX File System Benchmark
- Dummy NFS Server
 - TmpFS Requires Too Much Memory
 - Memory-based Metadata Only File System
 - Kernel-Level Implementation
 - Run at Interrupt Level



October 12-14, 2004

Workload Capture Tools



October 12-14, 2004

Workloads

- Most Important Piece of the Puzzle
- Workload Capture
 - Syscall/OS Level Traces
 - I/O Traces
 - Network Traces
- Post Processing
 - Anonymize
 - Statistical Analysis
 - Generate Workload Files for FSSB
- **Must Be Easy to Use!**



October 12-14, 2004

Issues

- What about mmaped files?
 - Is tracking page-in good enough?
 - Will frequency of data access become an issue?
- Simulating Application Load
 - Not All Memory is Available to the OS!



October 12-14, 2004

Product Plans



October 12-14, 2004

Potential Customers

- **Developers**
 - Performance Analysis
 - Sustaining/Regression Testing
 - Replay Captured Workload From Failing System
- **OEM/VARs**
 - Performance Analysis
 - Help Determine Which Components to Use in a Solution
- **End Users**
 - Validate Vendors Proposed Solutions



October 12-14, 2004

Benchmark Products

- Run on Windows[®] and UNIX
- Benchmark Framework
 - Purchase Once
- Benchmark Plugins Shipping 2005
 - POSIX File System
 - NFS Client Benchmark (Pseudo NFS Server - UNIX)
 - NFS V2-4 Server Benchmark
 - CIFS Server Benchmark
- Workload Capture Tool
- Signing Up Alpha Sites for Q105



October 12-14, 2004

EP Network Storage Performance Lab

“We don’t make NAS and SAN devices, we make them better.”

www.nsplab.com