SDC
STORAGE DEVELOPER CONFERENCE
SNIA SANTA CLARA, 2015

Application-Level Benchmarking with SPEC SFS® 2014

Nick Principe – EMC
Vernon Miller – IBM

Agenda

- Why application-level benchmarking?
- What is application-level benchmarking?
- The SPEC SFS 2014 Workloads
  - Reporting SFS 2014 Results
- Testing a “Storage Solution”
- Ramifications of application-level benchmarking
  - Concepts
  - Real-life examples
- Key takeaways
- Q&A

SPEC
Standard Performance Evaluation Corporation

- **The Standard Performance Evaluation Corporation (SPEC)** is a non-profit corporation formed to establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers. SPEC develops benchmark suites and also reviews and publishes submitted results from member organizations and other benchmark licensees
- www.spec.org
- SPEC and SPEC SFS are registered trademarks of the Standard Performance Evaluation Corporation. Additional product and service names mentioned herein may be the trademarks of their respective owners.

Why application-level benchmarking?

- The focus of the SFS benchmark has changed in SFS 2014
  - Load is now generated at the application level
  - The aim is to measure the storage performance of the environment as a whole
    - We call this the “Storage Solution”
  - Vendors now have the flexibility to configure the benchmark to match their environment
    - Put the bottleneck where you want to show value

Why application-level benchmarking?

- Increased flexibility addresses the market
  - Complexity of storage solutions
  - Diversity of architectures and protocols
  - Fairness to all implementations
- Any other approach is no longer appropriate for industry-standard benchmarking of storage solutions

What is Application-Level Benchmarking?

- Prior to SFS 2014, benchmark generated its own NFS or SMB traffic
  - Bypassed load generator operating system
  - Focus: performance of monolithic NAS server
- SFS 2014 uses native OS calls to generate application-level load
  - Data and metadata ops processed by OS
  - Focus: storage performance of the environment as a whole (Storage Solution)

The SPEC SFS 2014 Workloads

For more details, see:
- SDC 2014 presentation: SPEC SFS 2014: An Under-the-Hood Review
- The SPEC SFS 2014 website http://www.spec.org/sfs2014

- DATABASE
  - Simulates OLTP database consolidation
  - Measured in # of concurrent **DATABASES**
- SWBUILD
  - Simulates large software project compilation
  - Measured in # of concurrent **BUILDS**
- VDA
  - Simulates acquisition of streaming data
  - Measured in # of concurrent **STREAMS**
- VDI
  - Simulates heavy steady-state VDI workload
  - Measured in # of concurrent **DESKTOPS**

Reporting SFS 2014 Results

- Disclosure of SPEC SFS 2014 results must meet the requirements of
  - SPEC SFS 2014 License
  - SPEC SFS 2014 Run and Reporting Rules
  - SPEC Fair Use Rules
  - Submission to SPEC for review encouraged
- Certain information is required to be disclosed
  - Do not use this presentation as a guide for public disclosure of SFS 2014 results
  - Created for education under auspices of SPEC using “generic” environments

Testing a “Storage Solution”

- More attention to benchmark configuration required
  - You must put the bottleneck in the right place
    - Understanding the whole system, from load generator to the disks, is a requirement
  - For publication, more configuration details must be recorded and disclosed
- SFS 2014 is still a storage benchmark
  - There is no attempt to simulate compute load

Ramifications of application-level benchmarking: Concepts

- Your load generators matter
  - Any config detail can affect performance
    - Storage connectivity, OS version, patch level, memory, client count, tuning parameters
- You can test anything that provides a file API to an application
  - Traditional NAS server, block storage with a file system on load generators, hyper-converged solutions, a single server with storage

Ramifications of application-level benchmarking: Concepts

- With great power comes great responsibility
  - Understand where your bottleneck is
    - SFS 2014 allows great flexibility in load placement as the workload scales
    - Getting this right is the key to getting the performance you expect
      - Likely you want to spread load as evenly as possible across ALL resources as the benchmark ramps up load
- The key config parameter in SFS 2014?
  - CLIENT_MOUNTPOINTS

Ramifications of application-level benchmarking: Concepts

- Measuring performance at multiple levels of the solution under test is key to understanding your solution's performance and bottlenecks
  - SFS 2014 reports application-level performance
  - Other statistics that are helpful to collect:
    - Storage array statistics, NAS server statistics, Hypervisor statistics, LG OS statistics
  - You may see different performance at the different levels
    - Each layer of the solution under test may change the workload

Ramifications of application-level benchmarking: Real-life Examples (Env 1)

- Configured an environment for testing
  - Midrange Storage Array
    - FC drives, FC frontend
  - Large Windows Server 2012 R2 NAS Server
    - FC backend, 48 cores, 256GB memory, 10GbE frontend
  - 20 Windows 8.1 VMs; 10 physical servers
    - FC or 10GbE SMB3 connectivity, 2 cores, 2 GB memory
  - Testing was done in two ways
    - Via SMB3 shares from NAS server (10GbE)
    - Via local E: drive, FC LUN via RDM to VM (FC)

Ramifications of application-level benchmarking: Real-life Examples (Env 1)

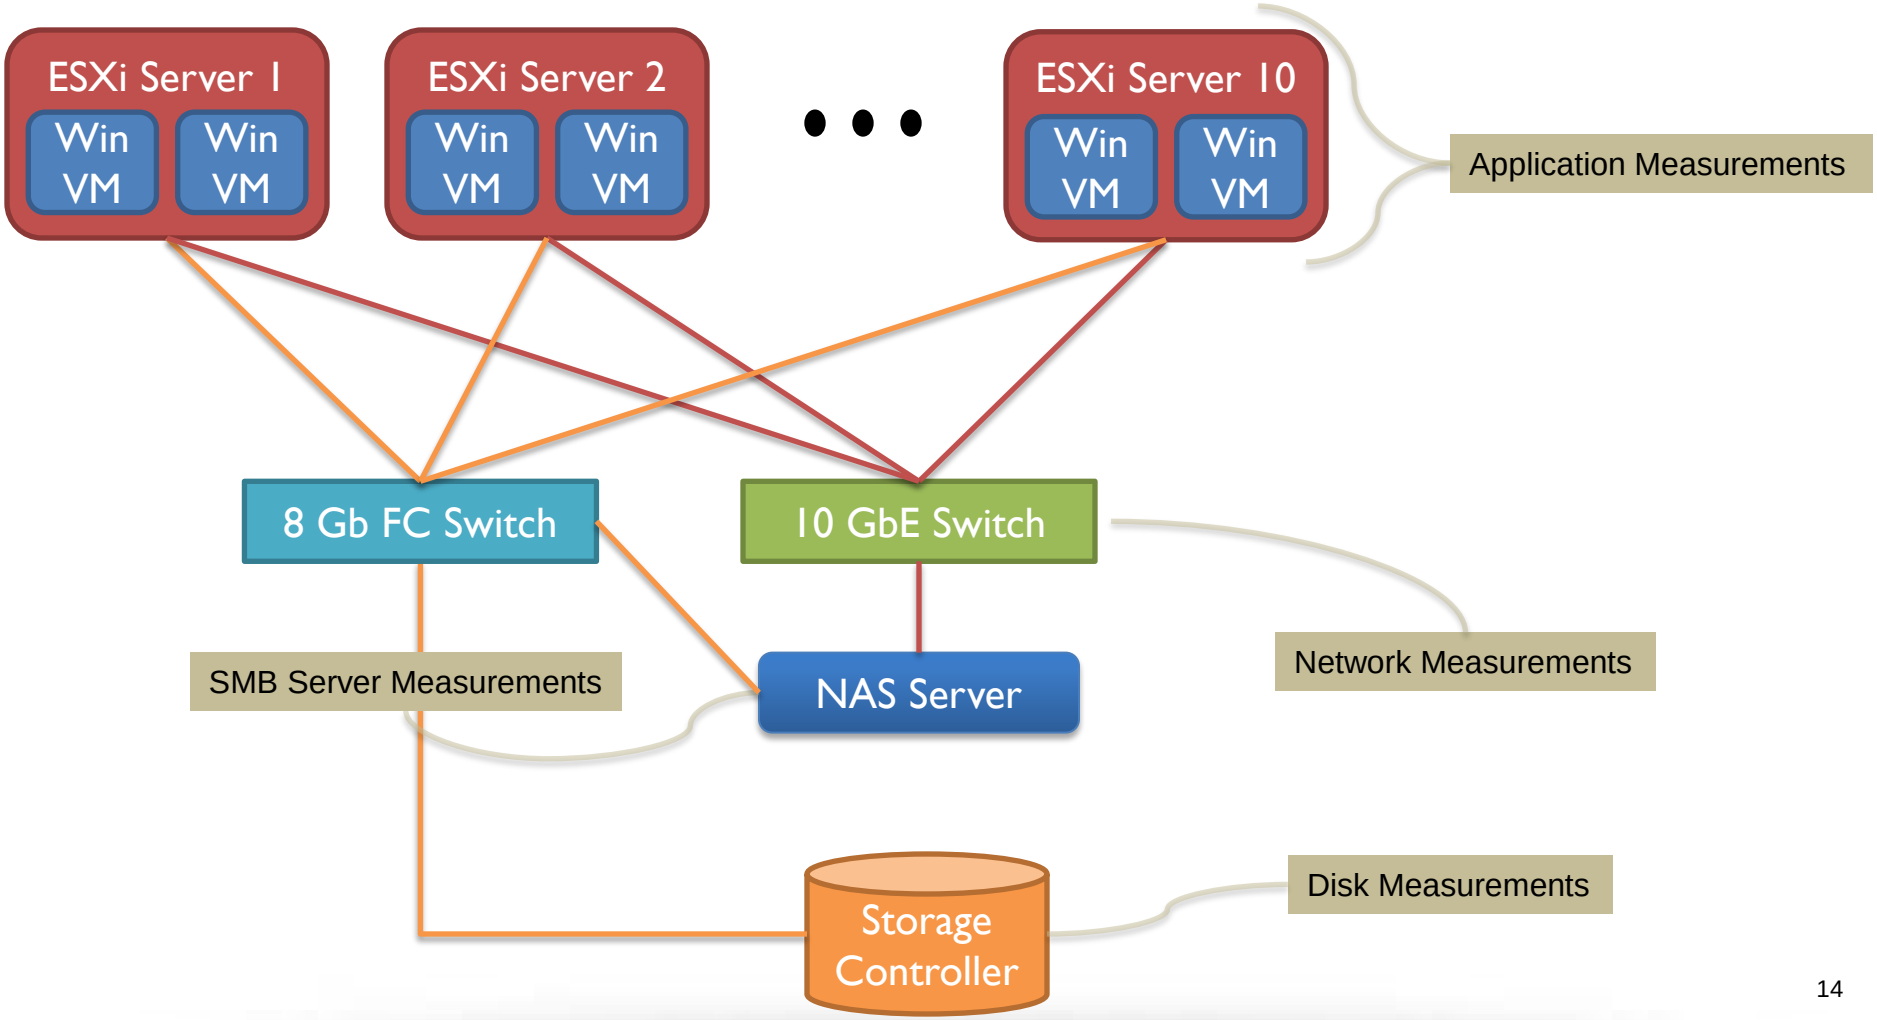

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

- Configured an environment for testing
  - Midrange Storage Array
    - SAS drives, FC frontend
  - 4 Node Distributed Filesystem
    - 2 nodes with FC backend, DDR IB cluster network
    - 2 nodes acting as NFS server, 10 GbE frontend
  - 2 NFS Clients
    - RHEL 6.5, 10GbE connectivity, 4 cores, 32 GB memory
  - Testing was done in two ways
    - Via NFSv3 exports from NAS server (10GbE)
    - Via local filesystem: single namespace on 2 nodes

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

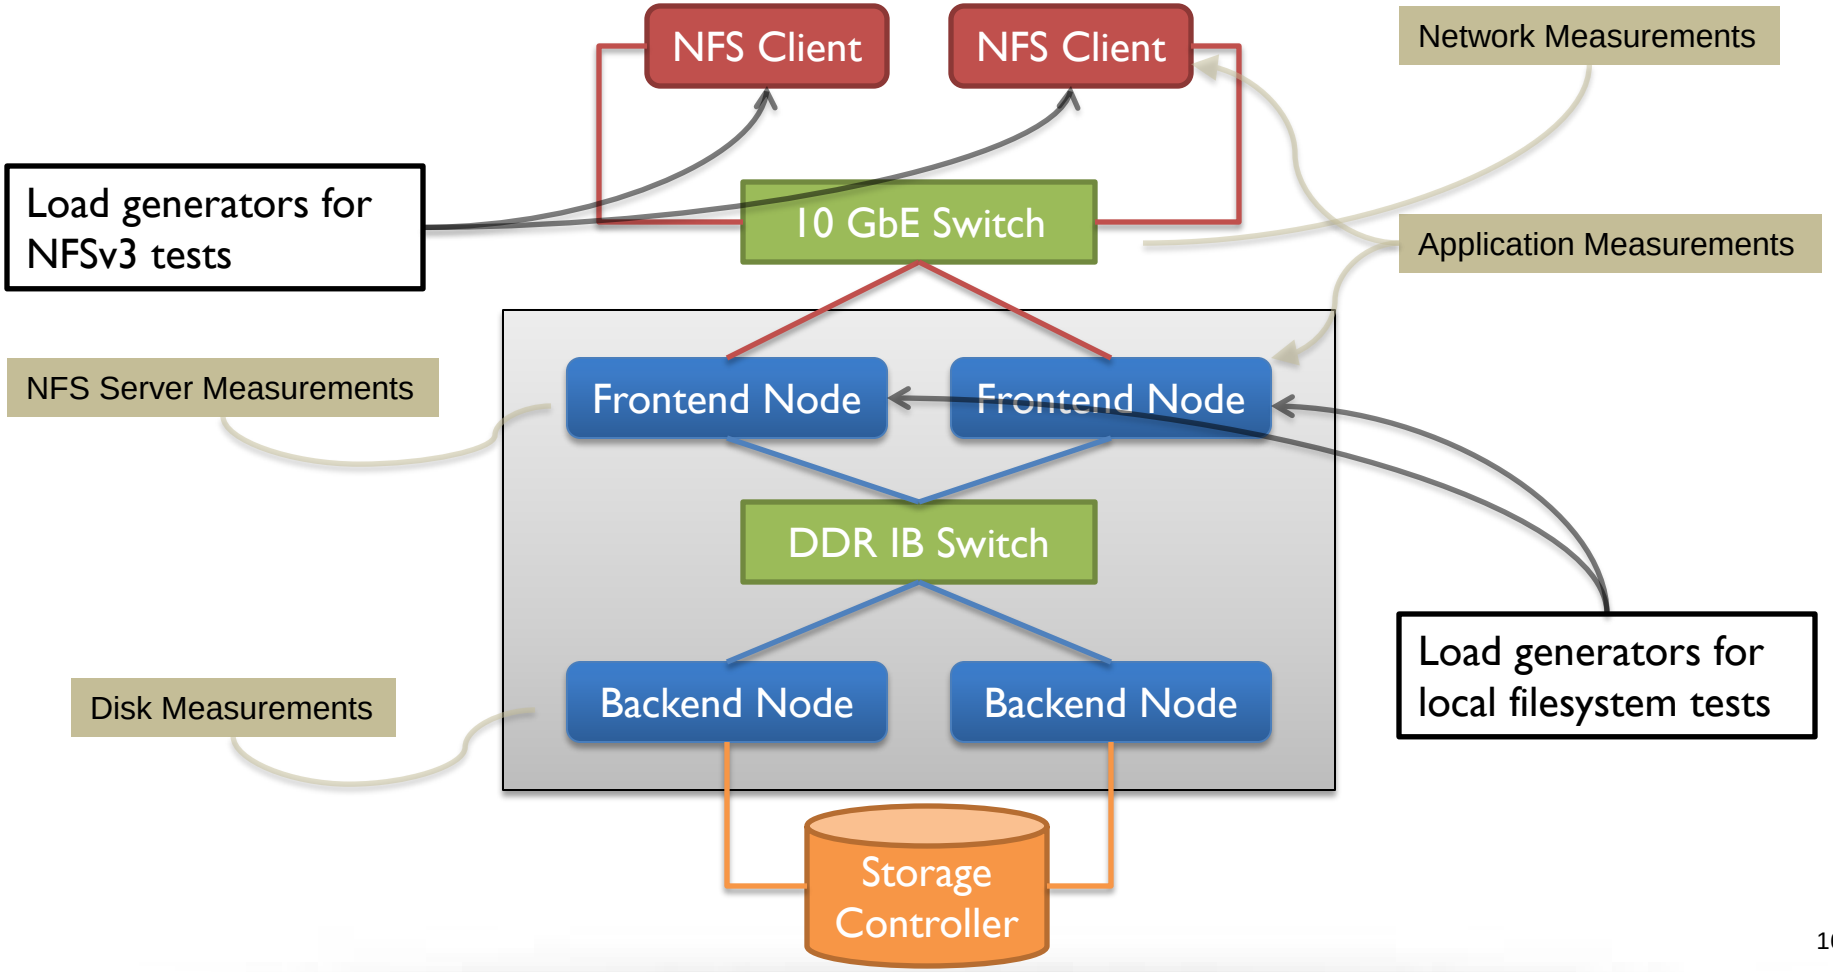

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

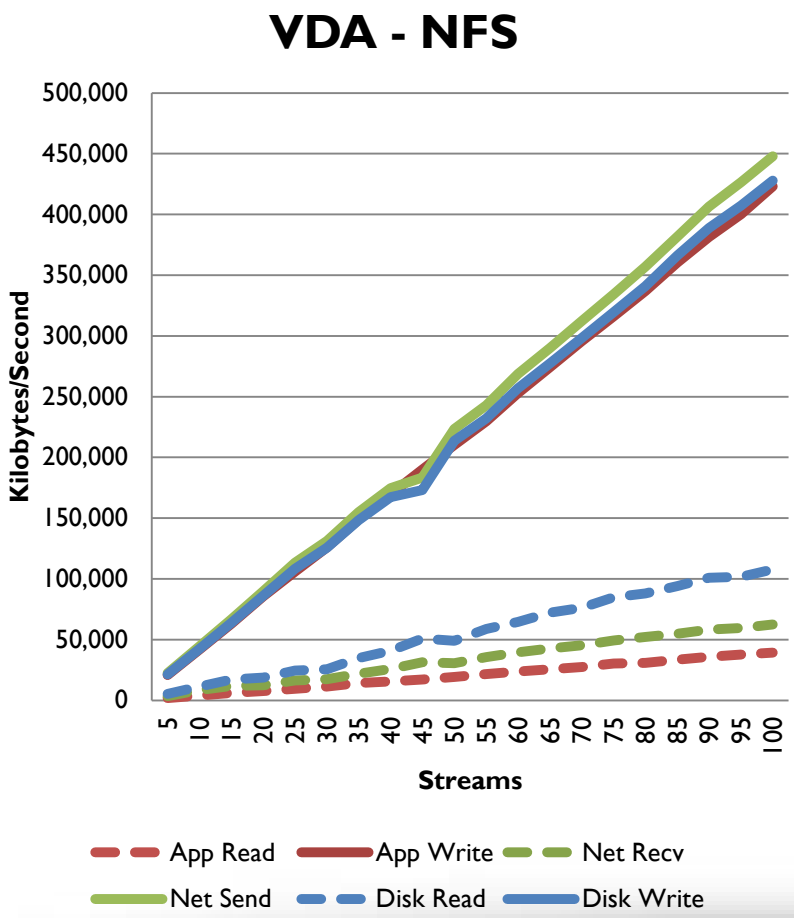

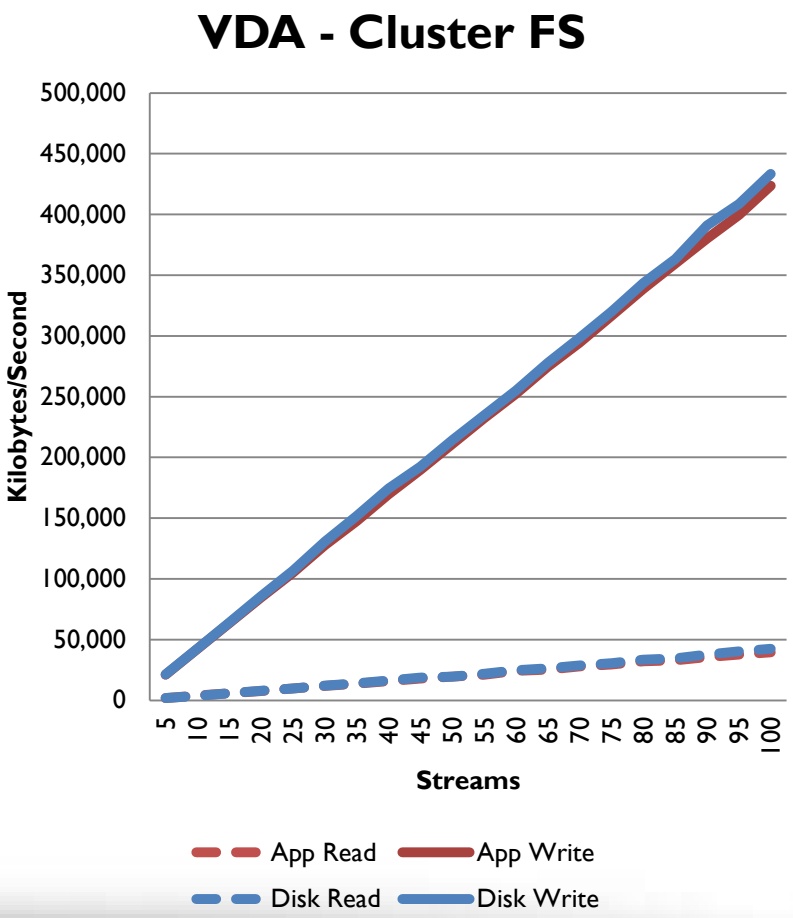

Ramifications of application-level benchmarking: Real-life Examples (Env 1)

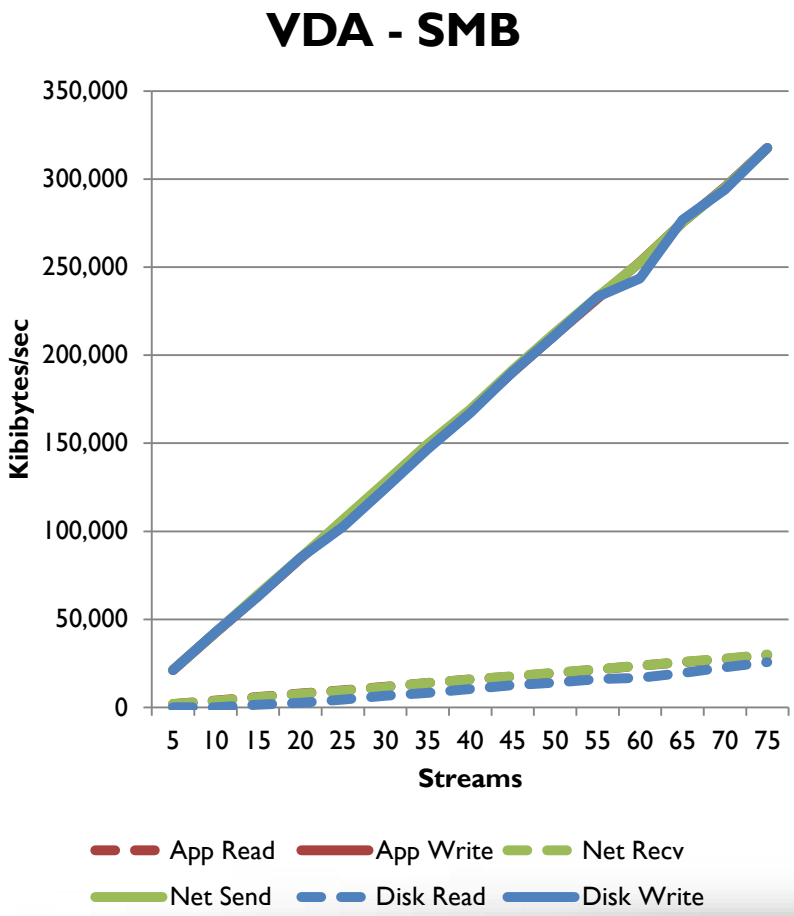

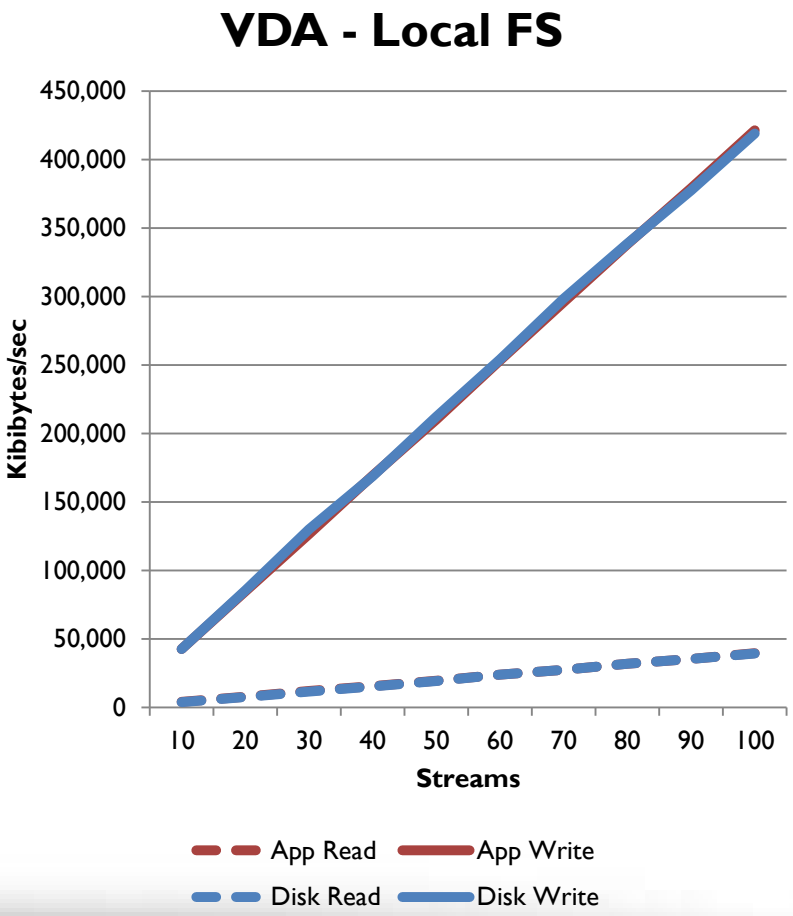

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

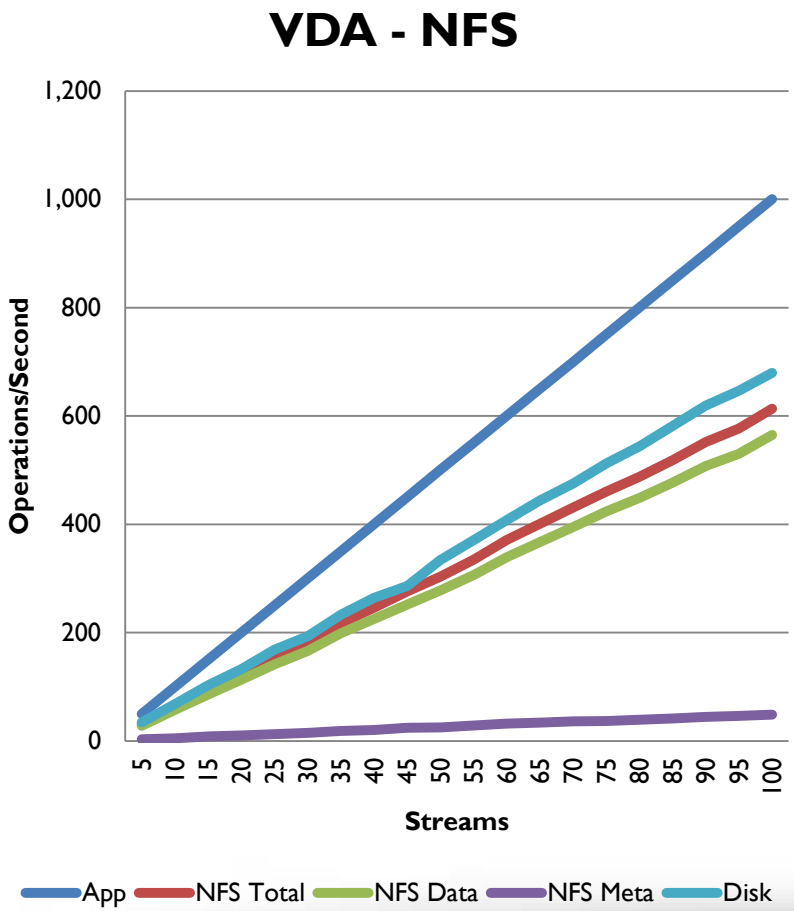

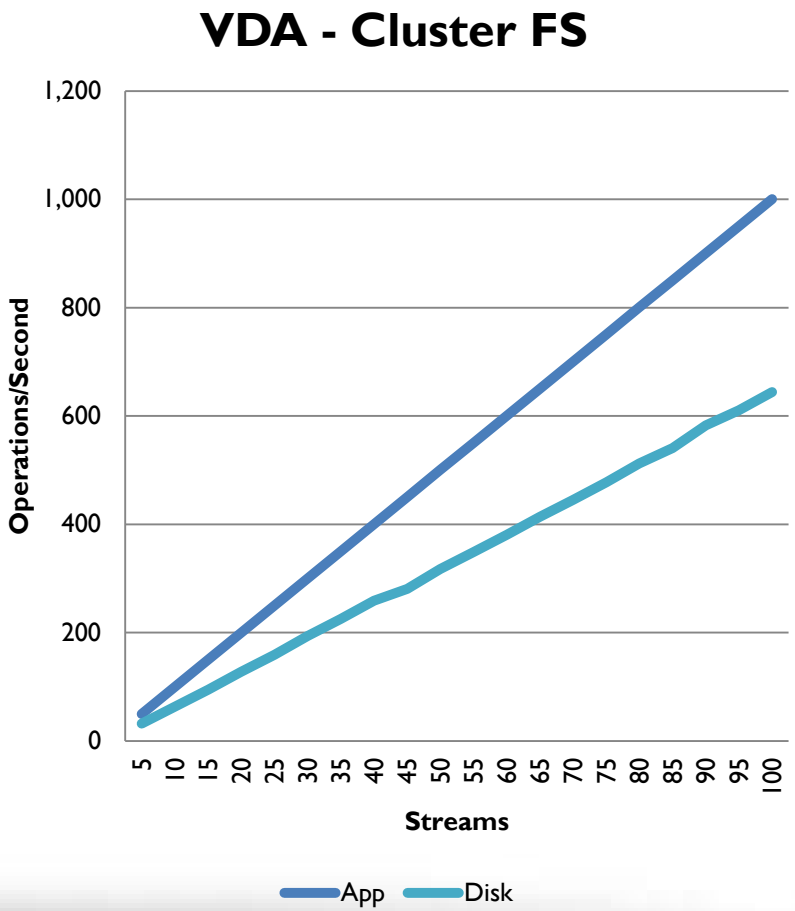

Ramifications of application-level benchmarking: Real-life Examples (Env 1)

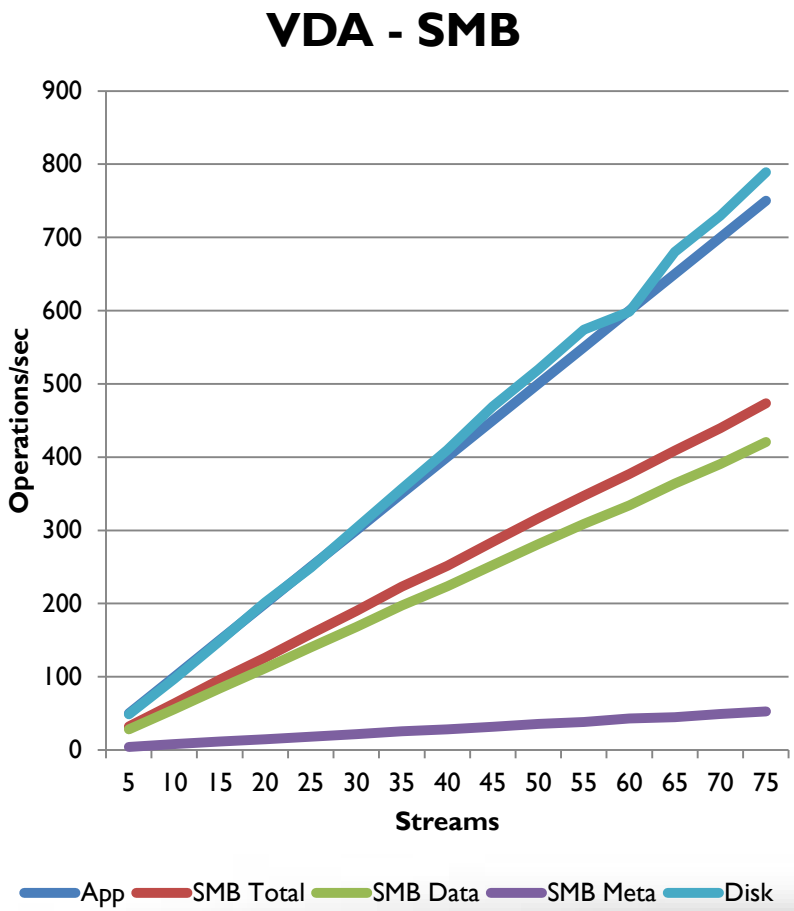

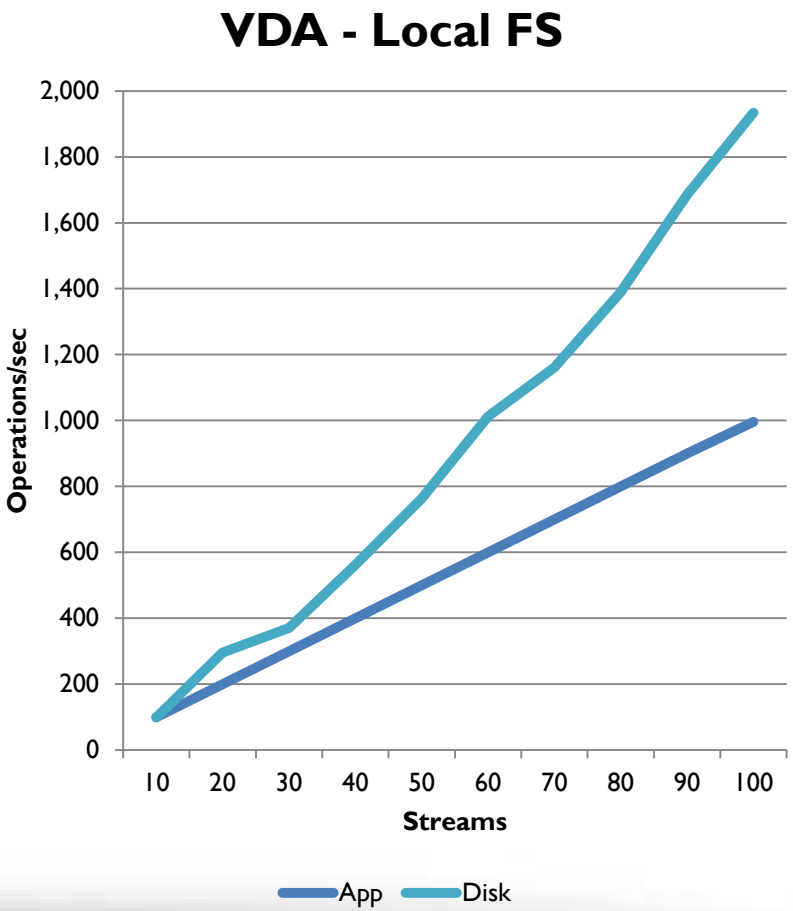

Ramifications of application-level benchmarking: Real-life Examples

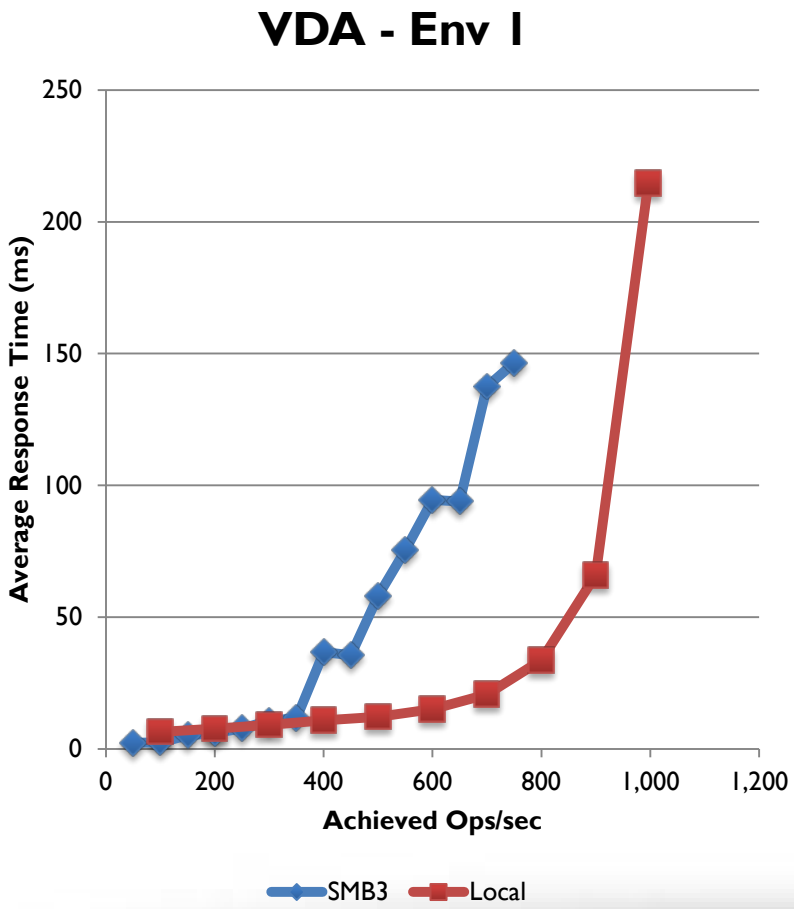

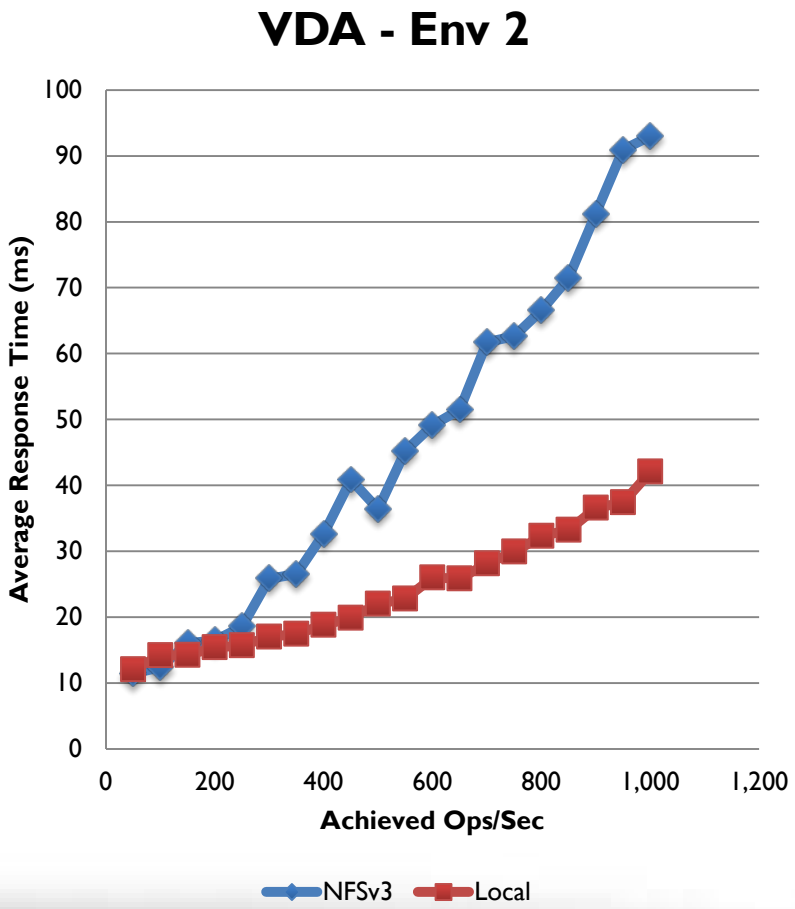

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

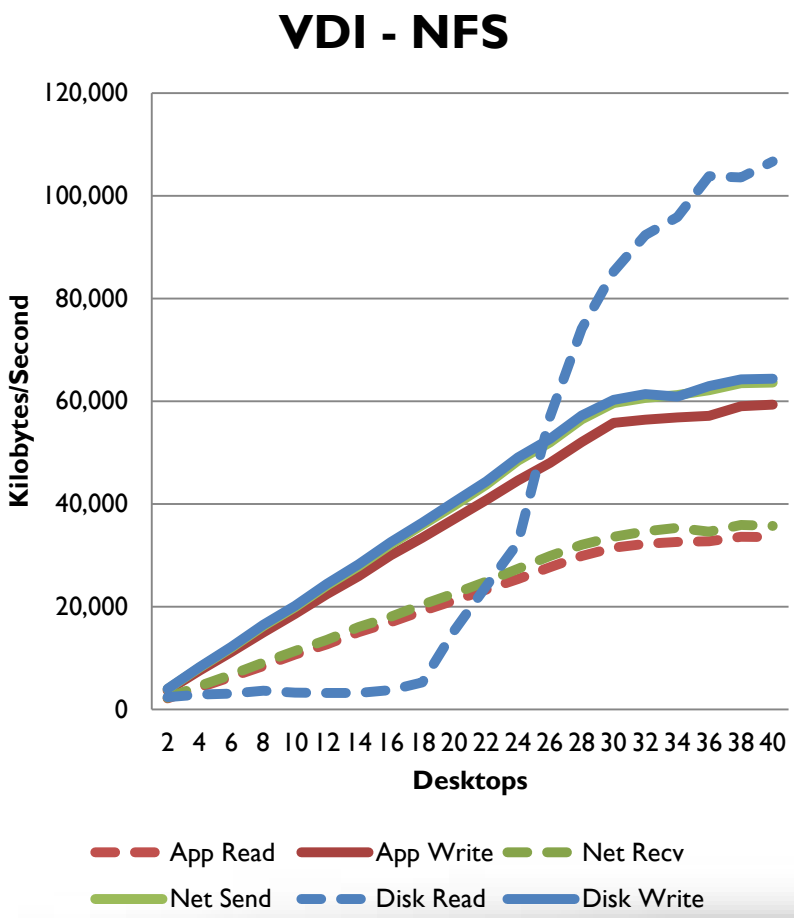

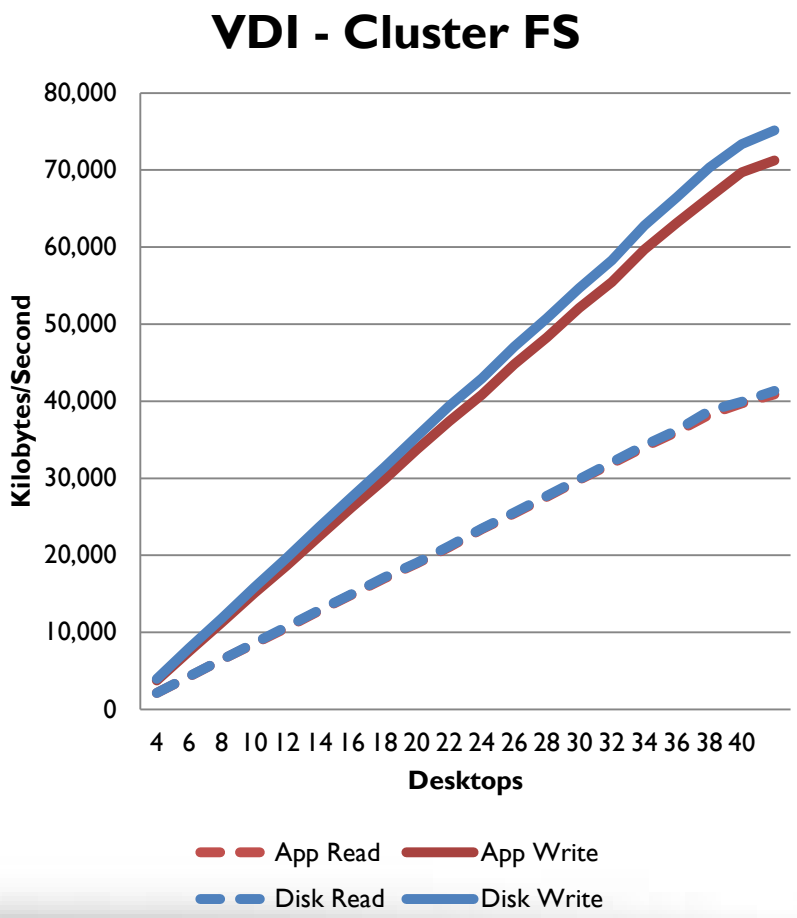

Ramifications of application-level benchmarking: Real-life Examples (Env 1)

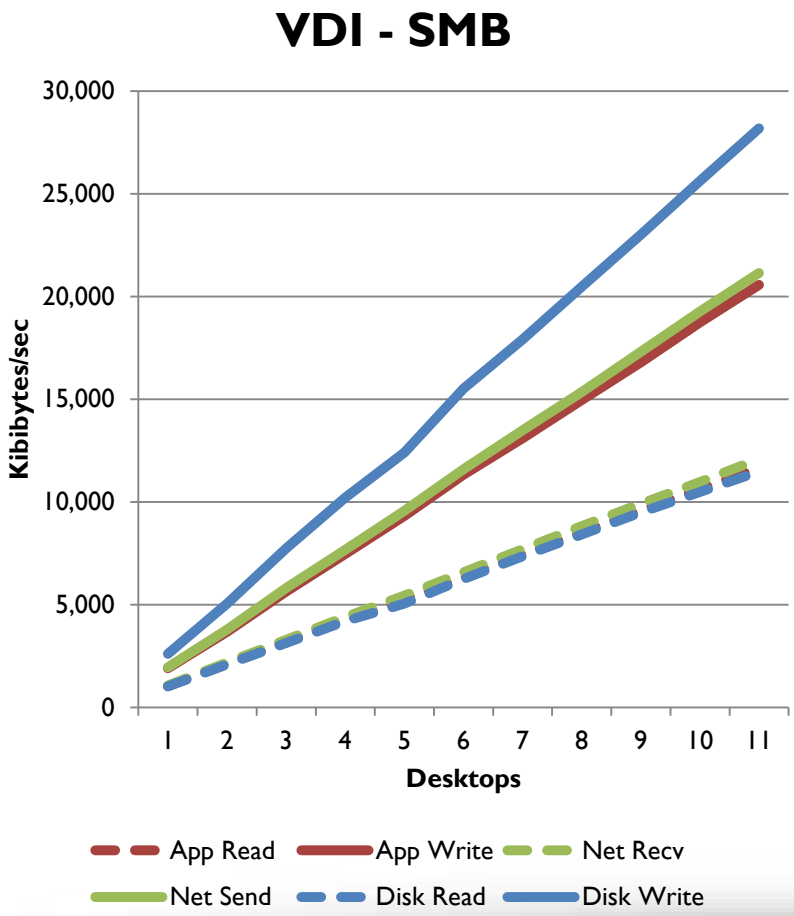

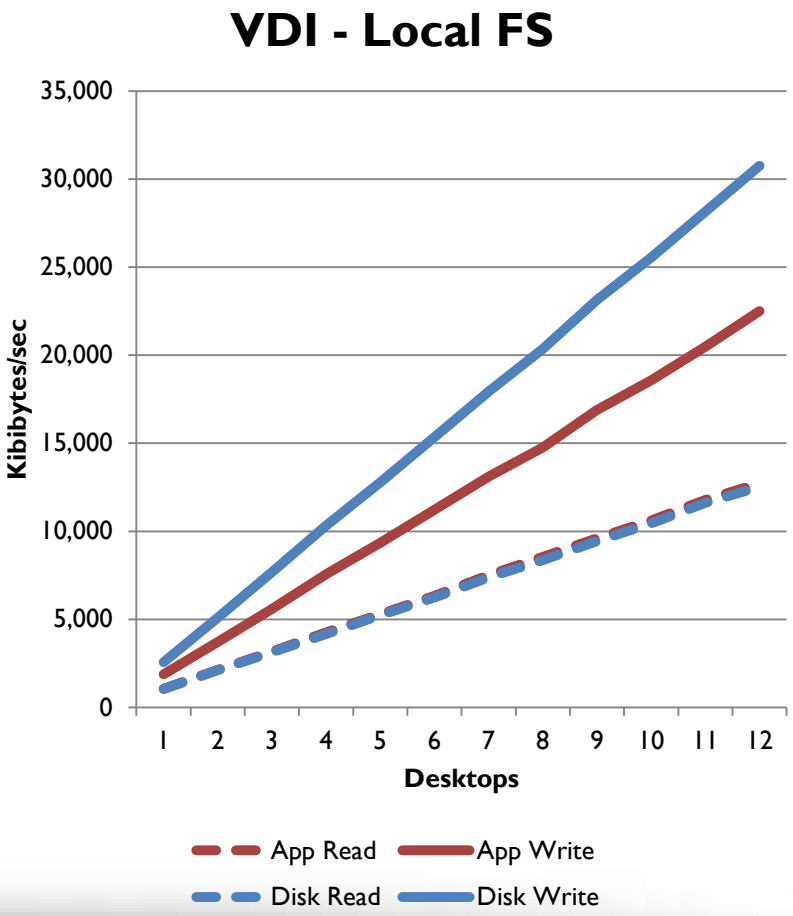

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

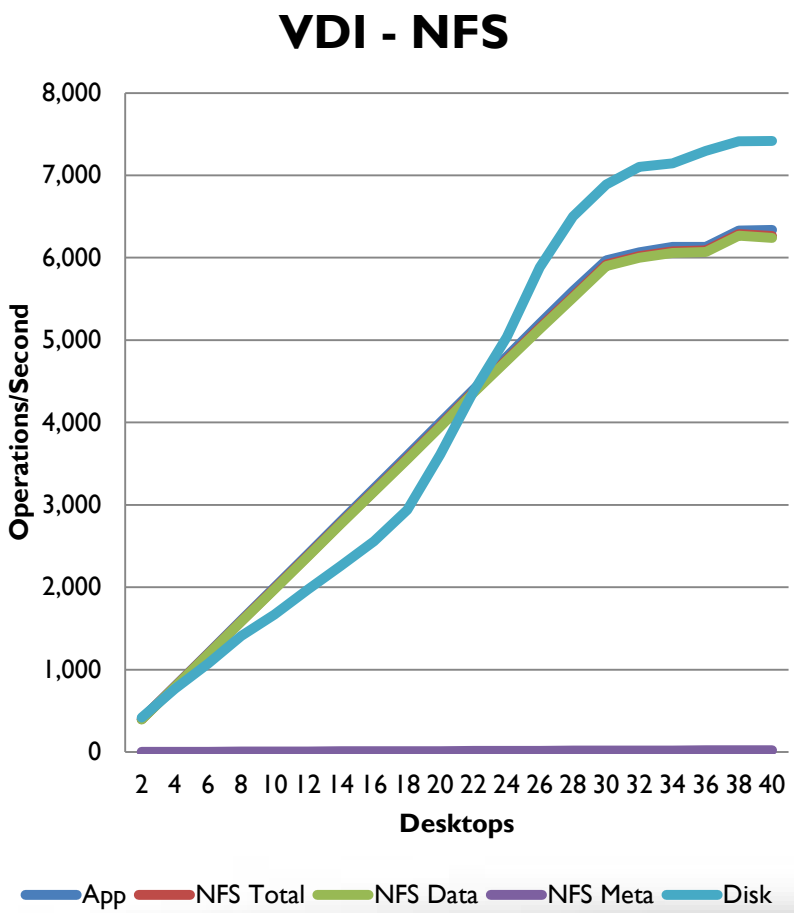

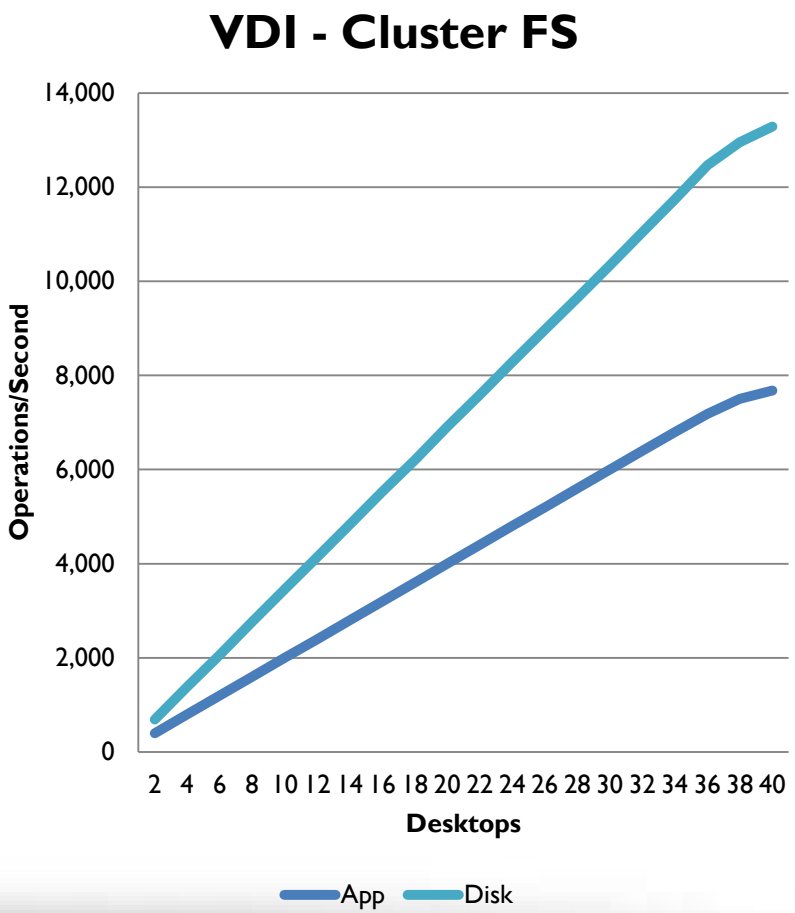

Ramifications of application-level benchmarking: Real-life Examples (Env 1)

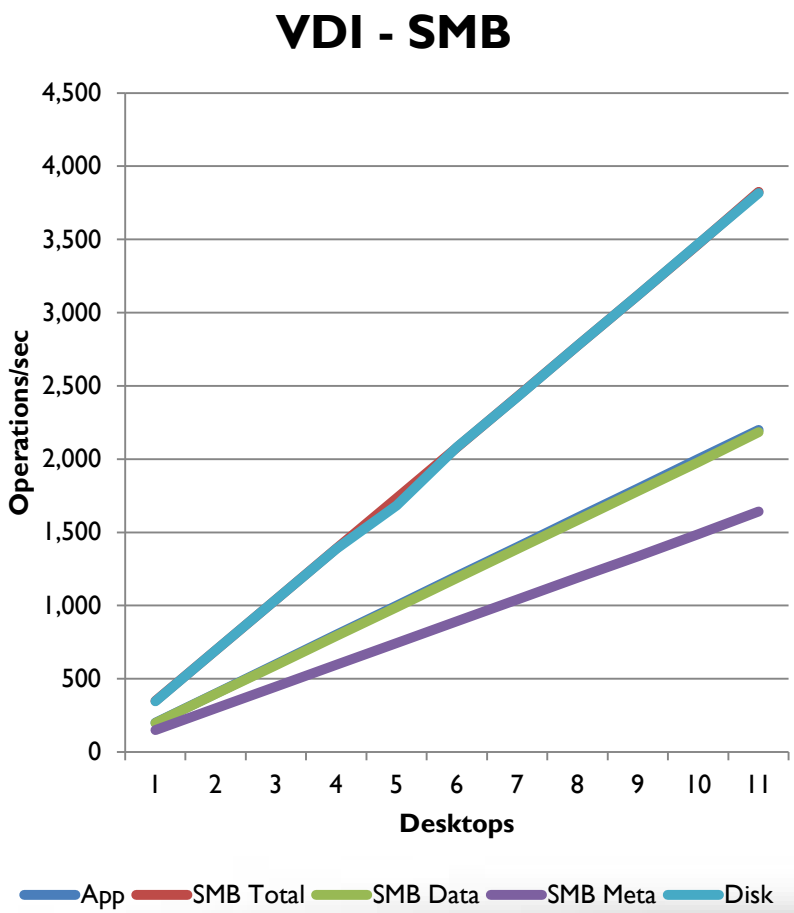

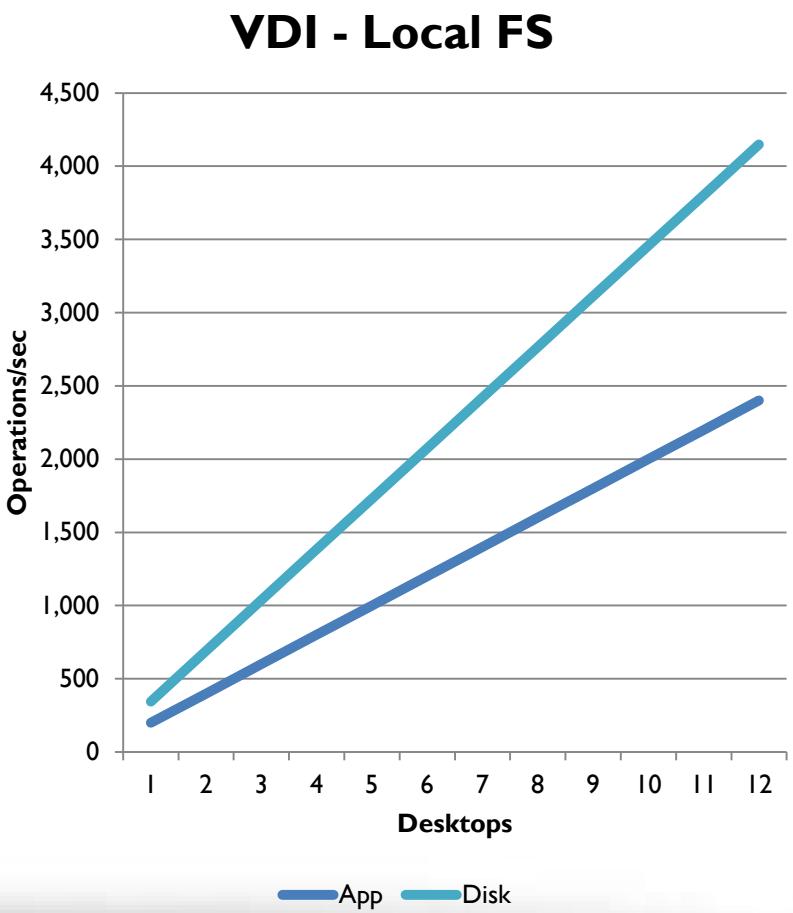

Ramifications of application-level benchmarking: Real-life Examples

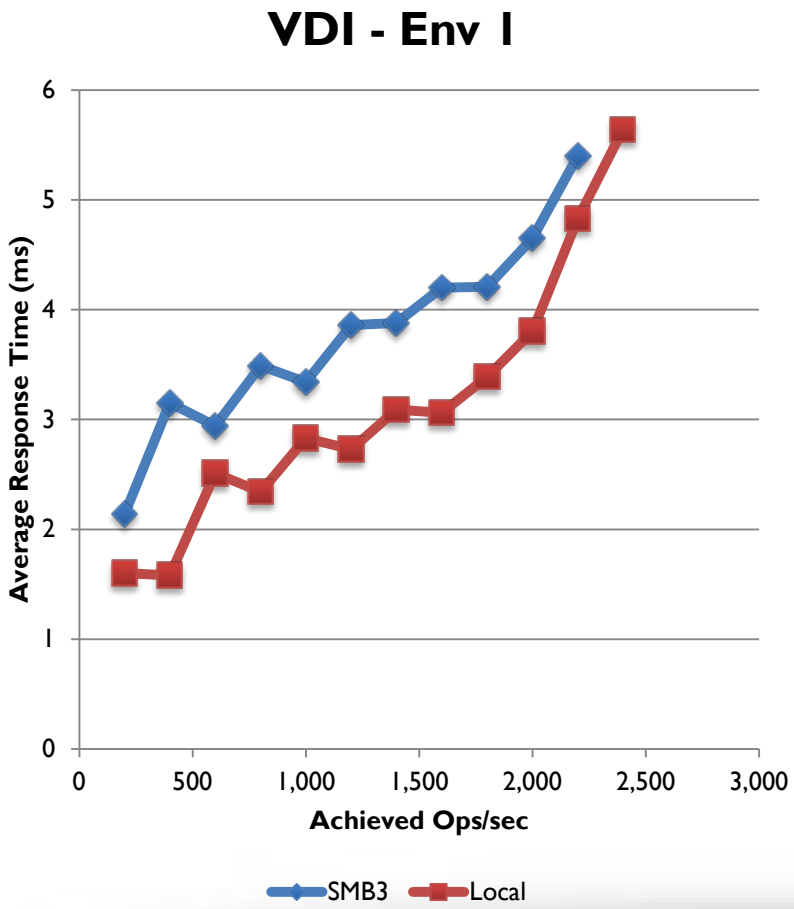

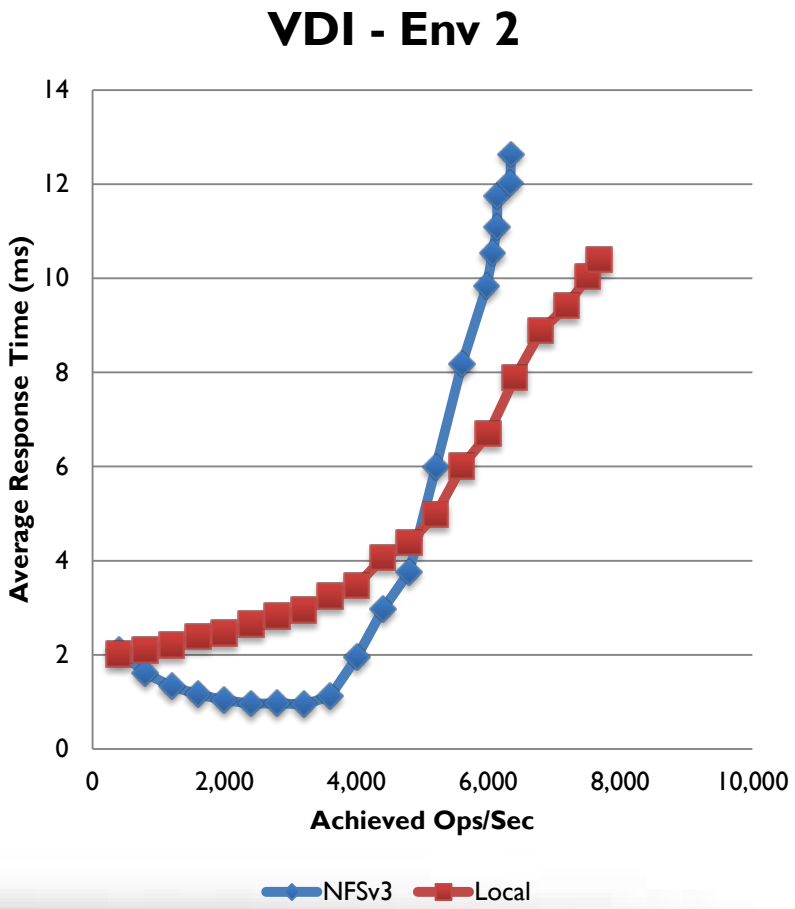

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

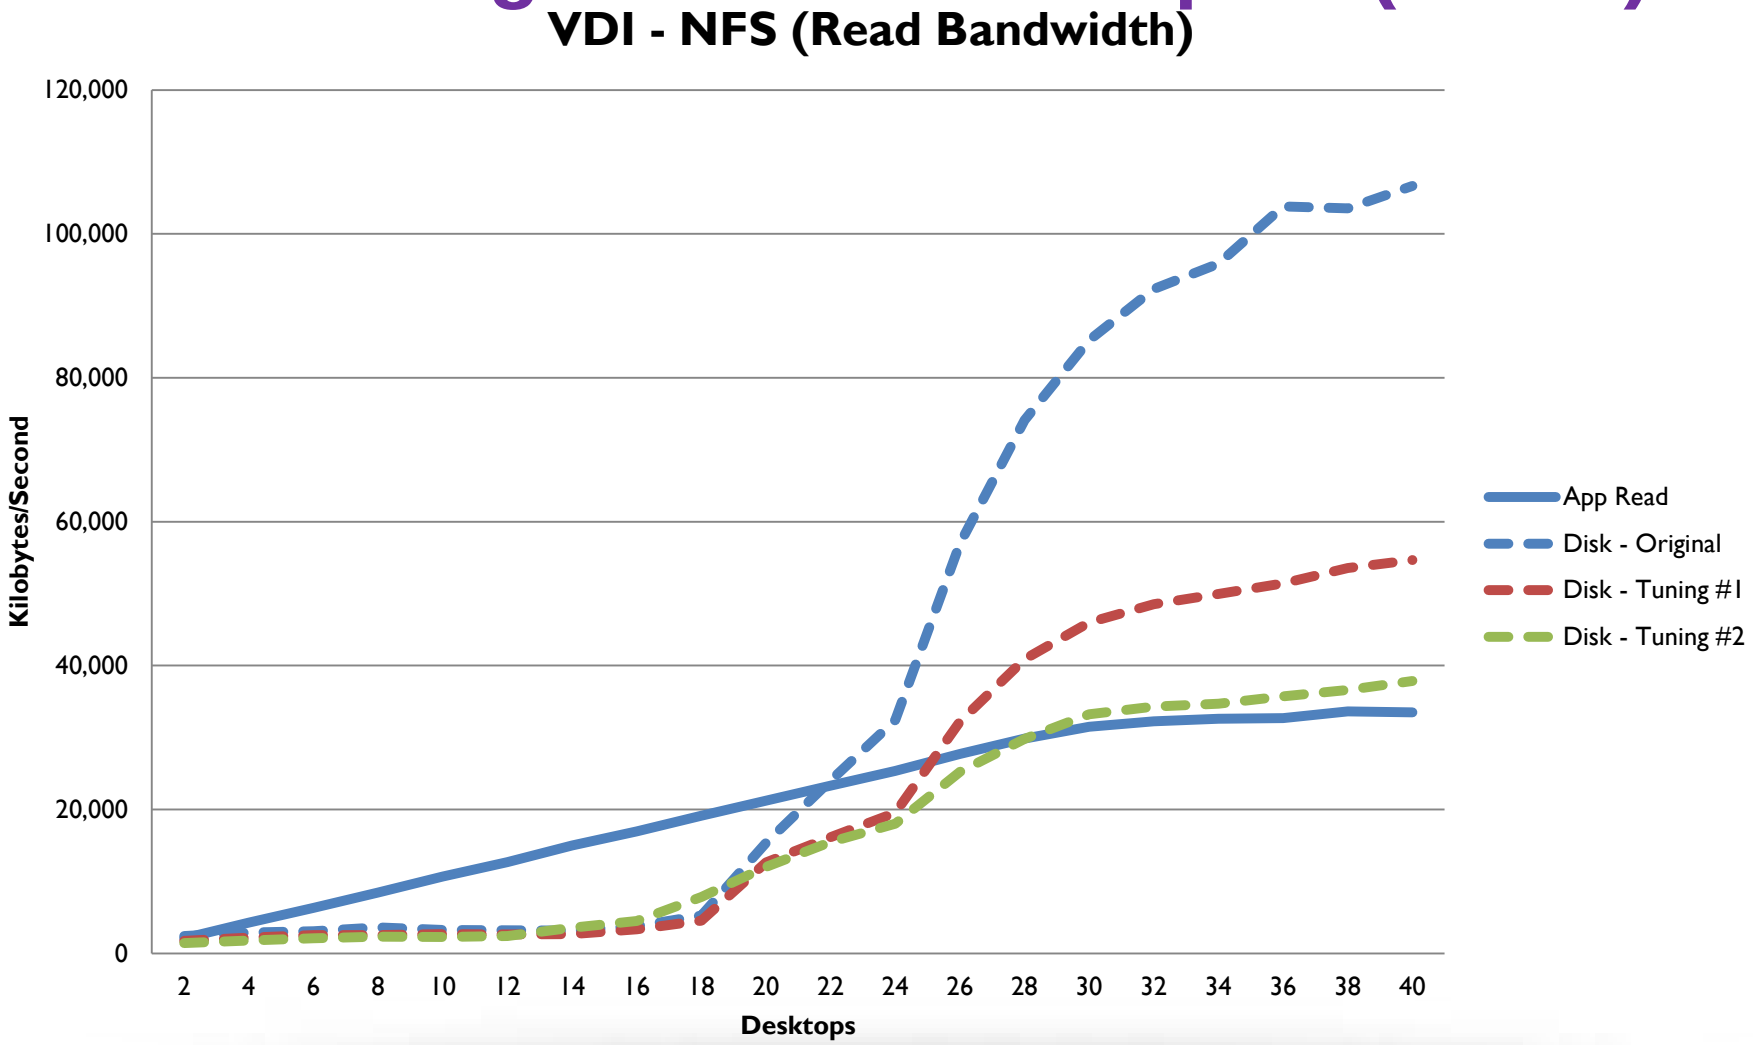

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

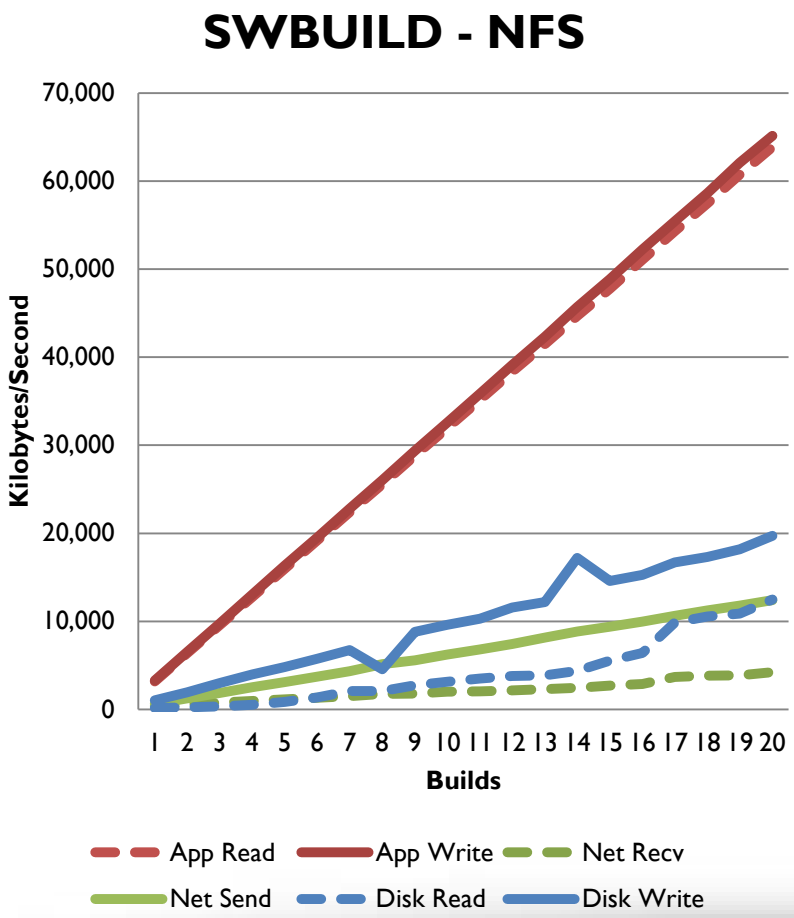

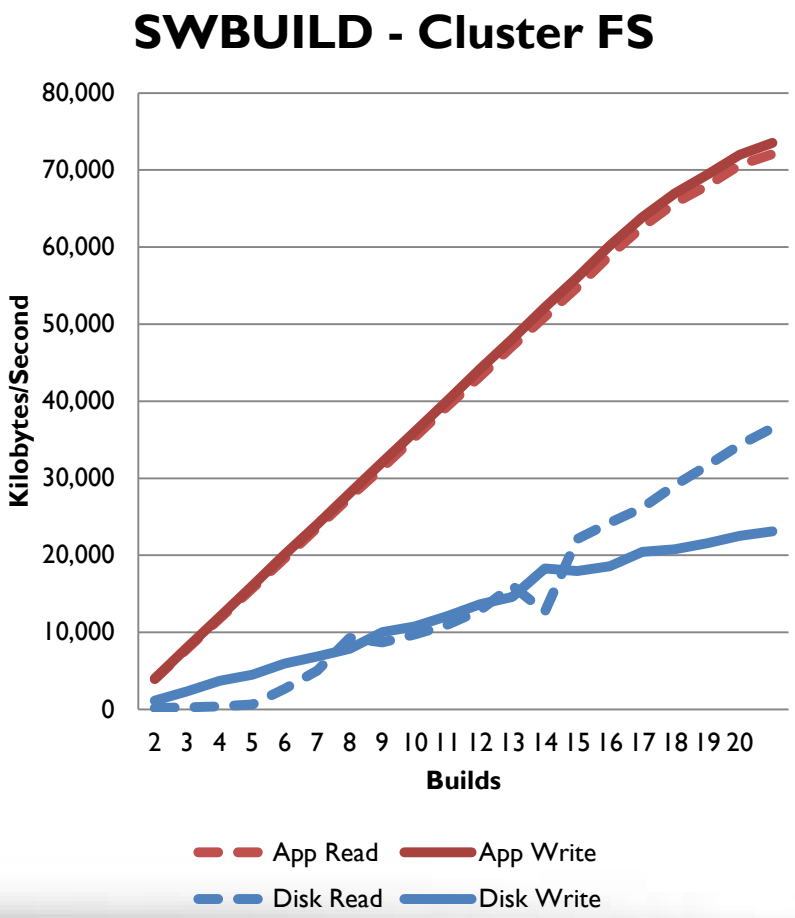

Ramifications of application-level benchmarking: Real-life Examples (Env 1)

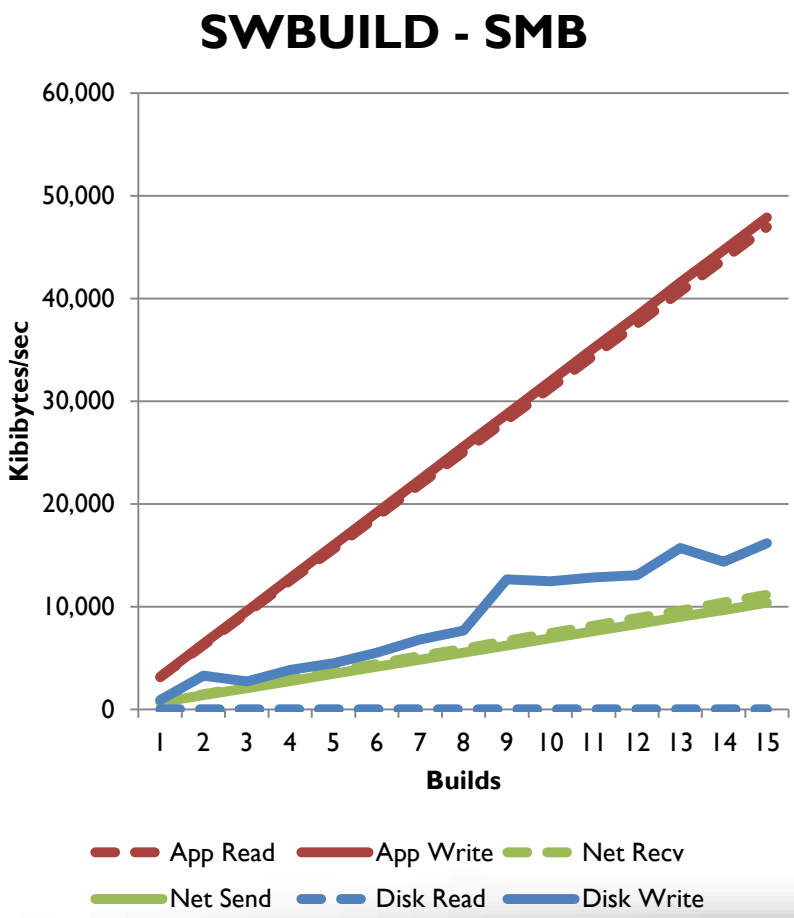

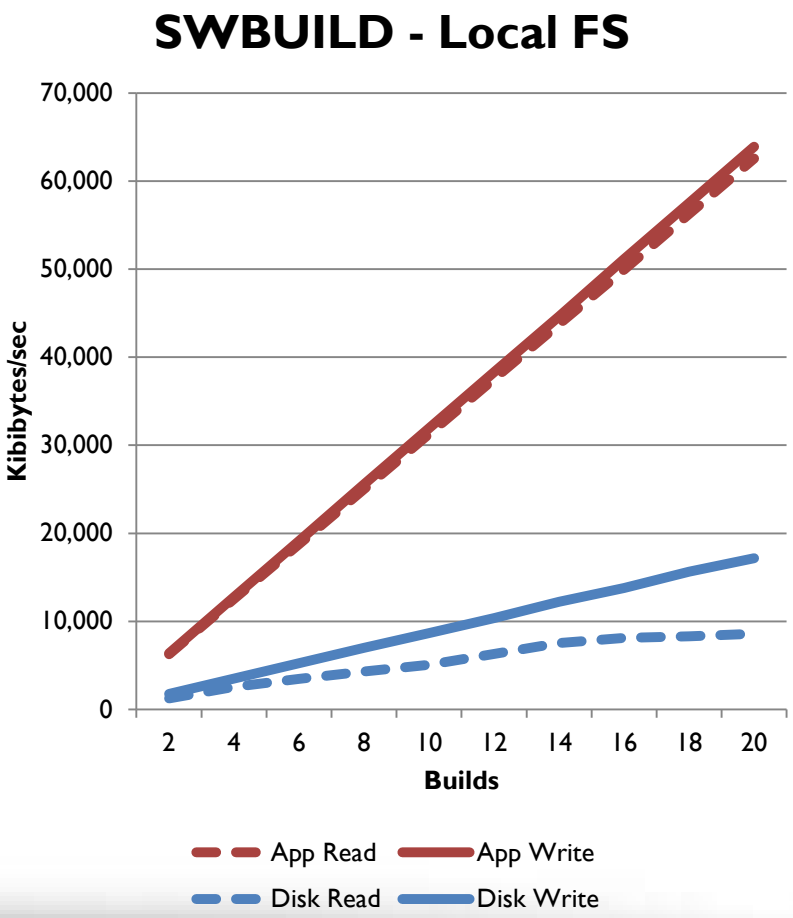

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

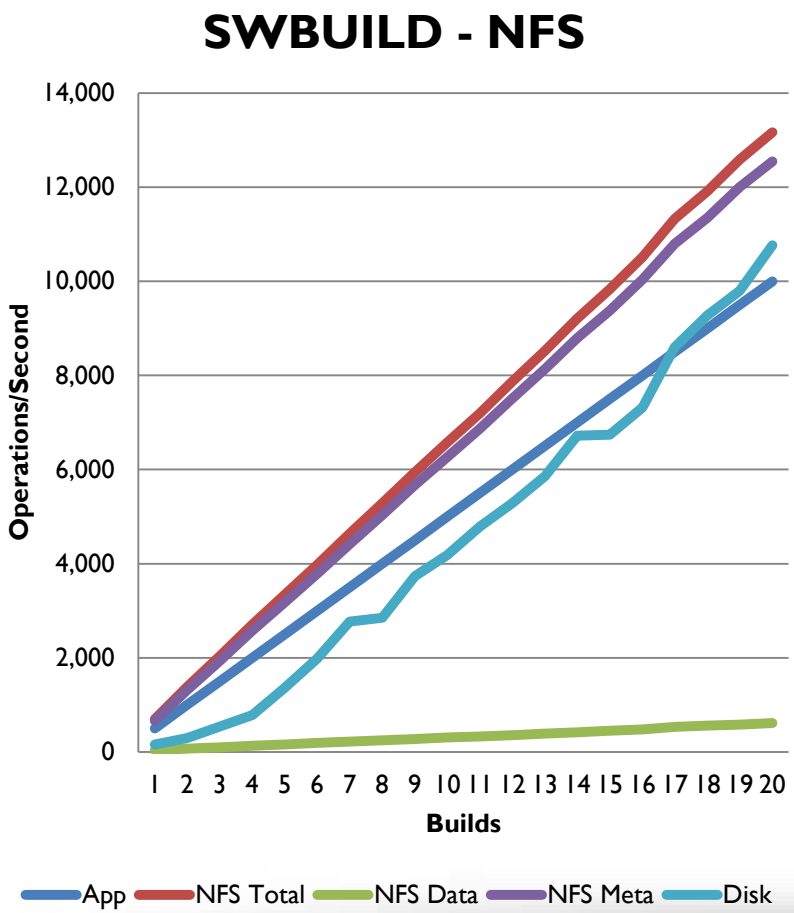

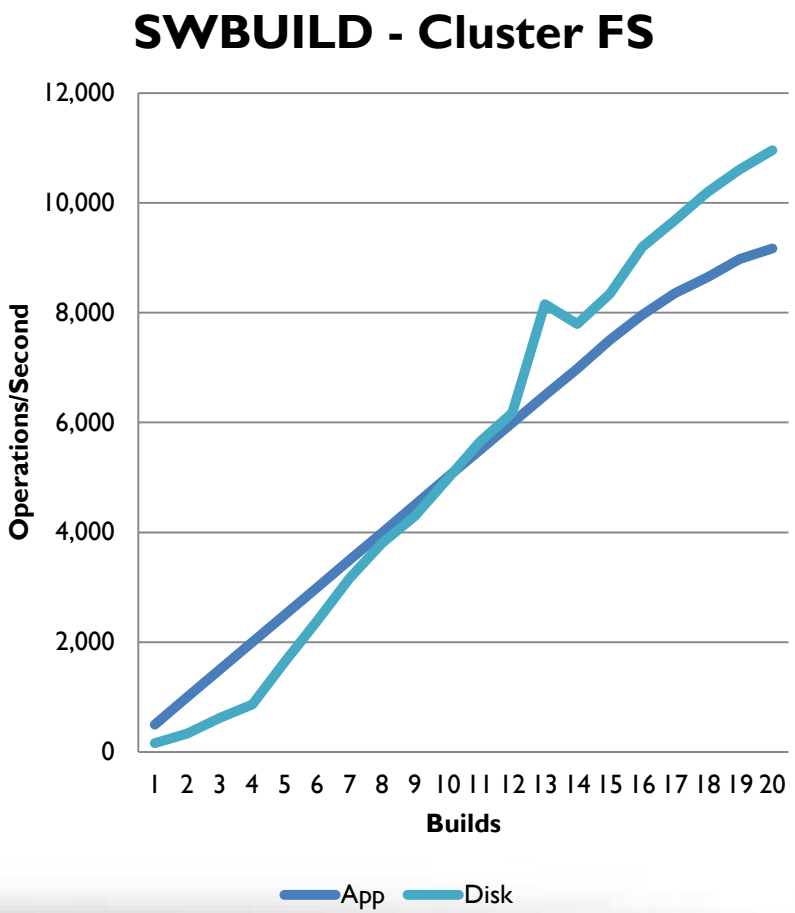

Ramifications of application-level benchmarking: Real-life Examples (Env 1)

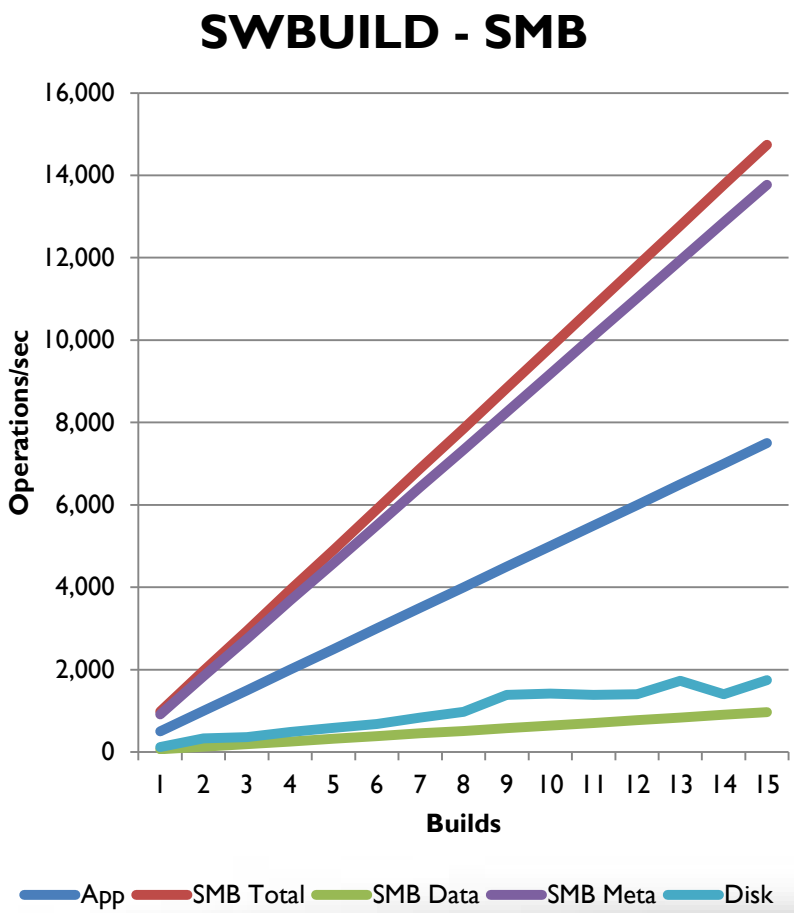

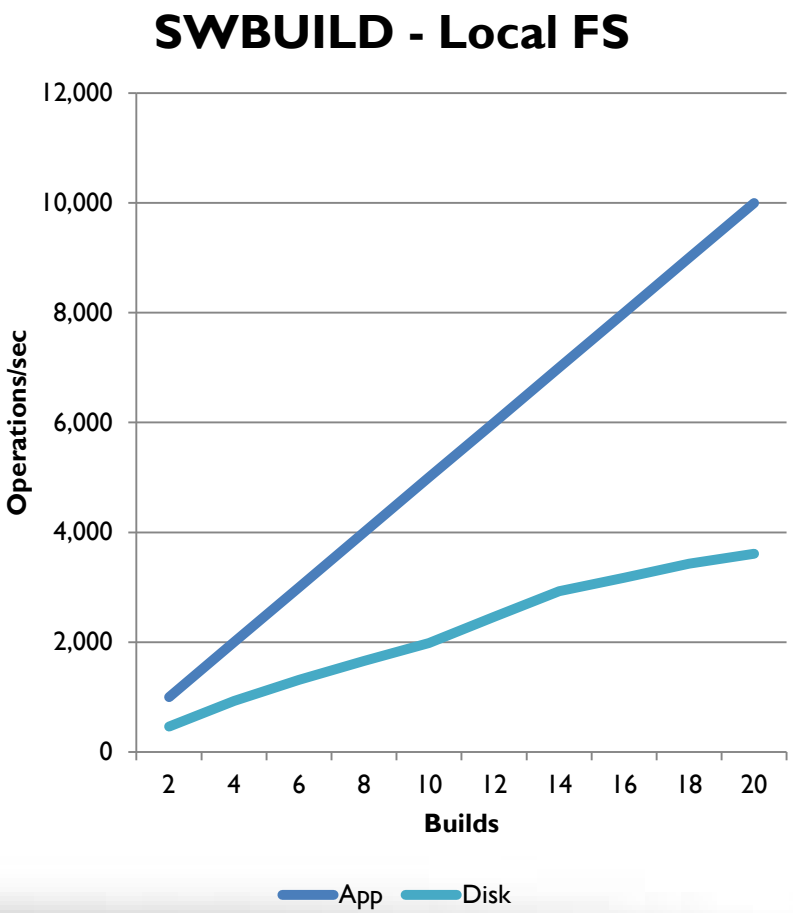

Ramifications of application-level benchmarking: Real-life Examples

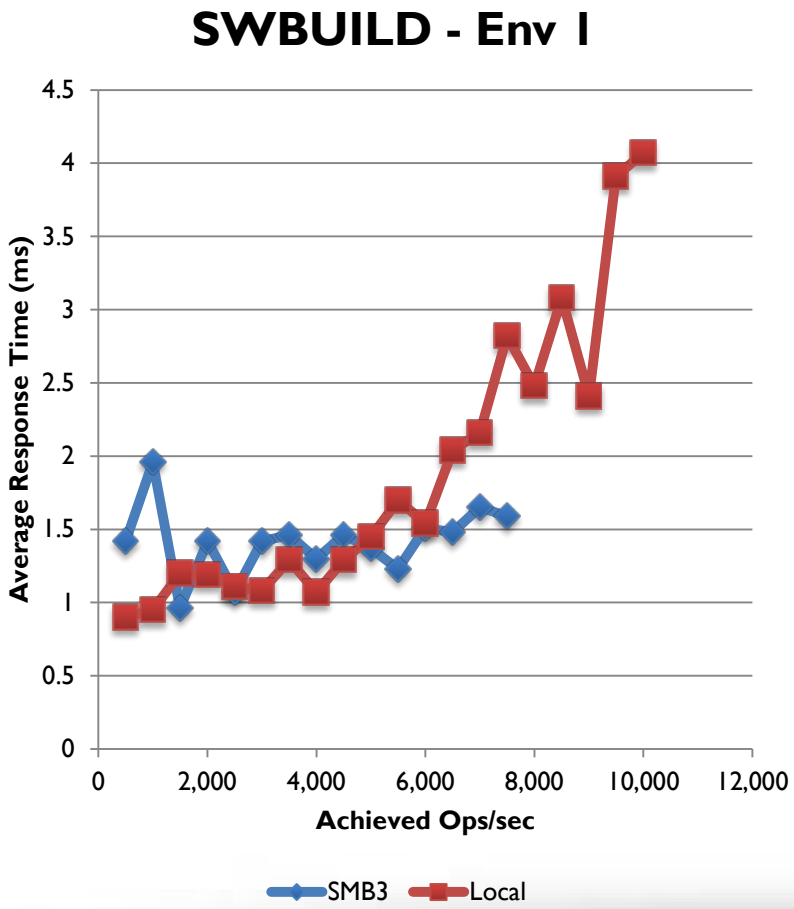

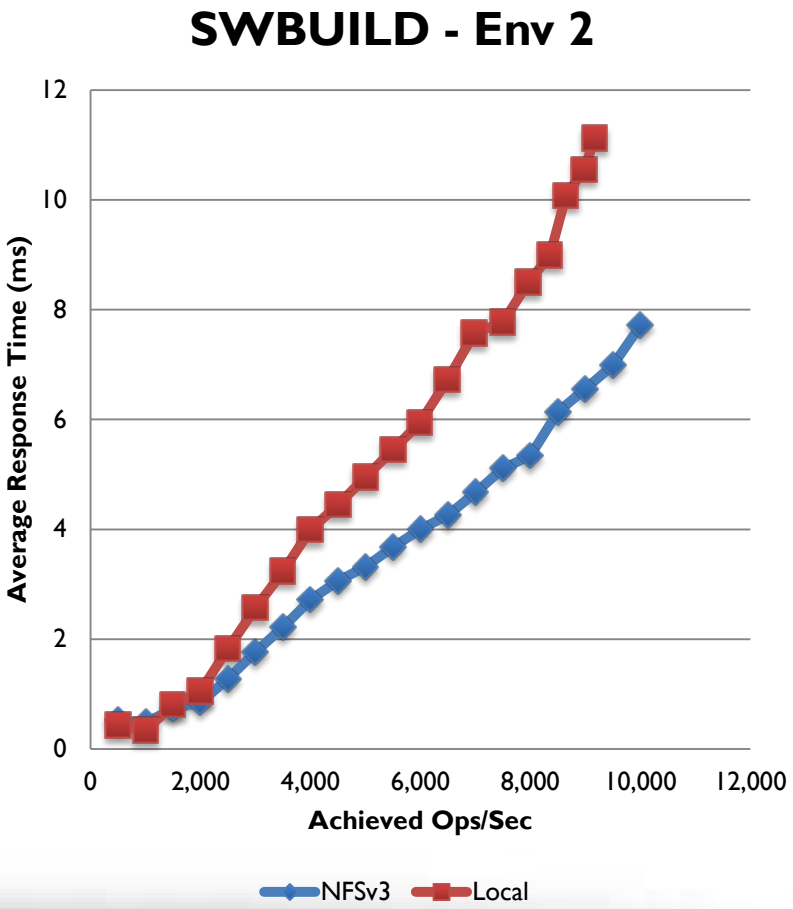

Ramifications of application-level benchmarking: Real-life Examples (Env 2)

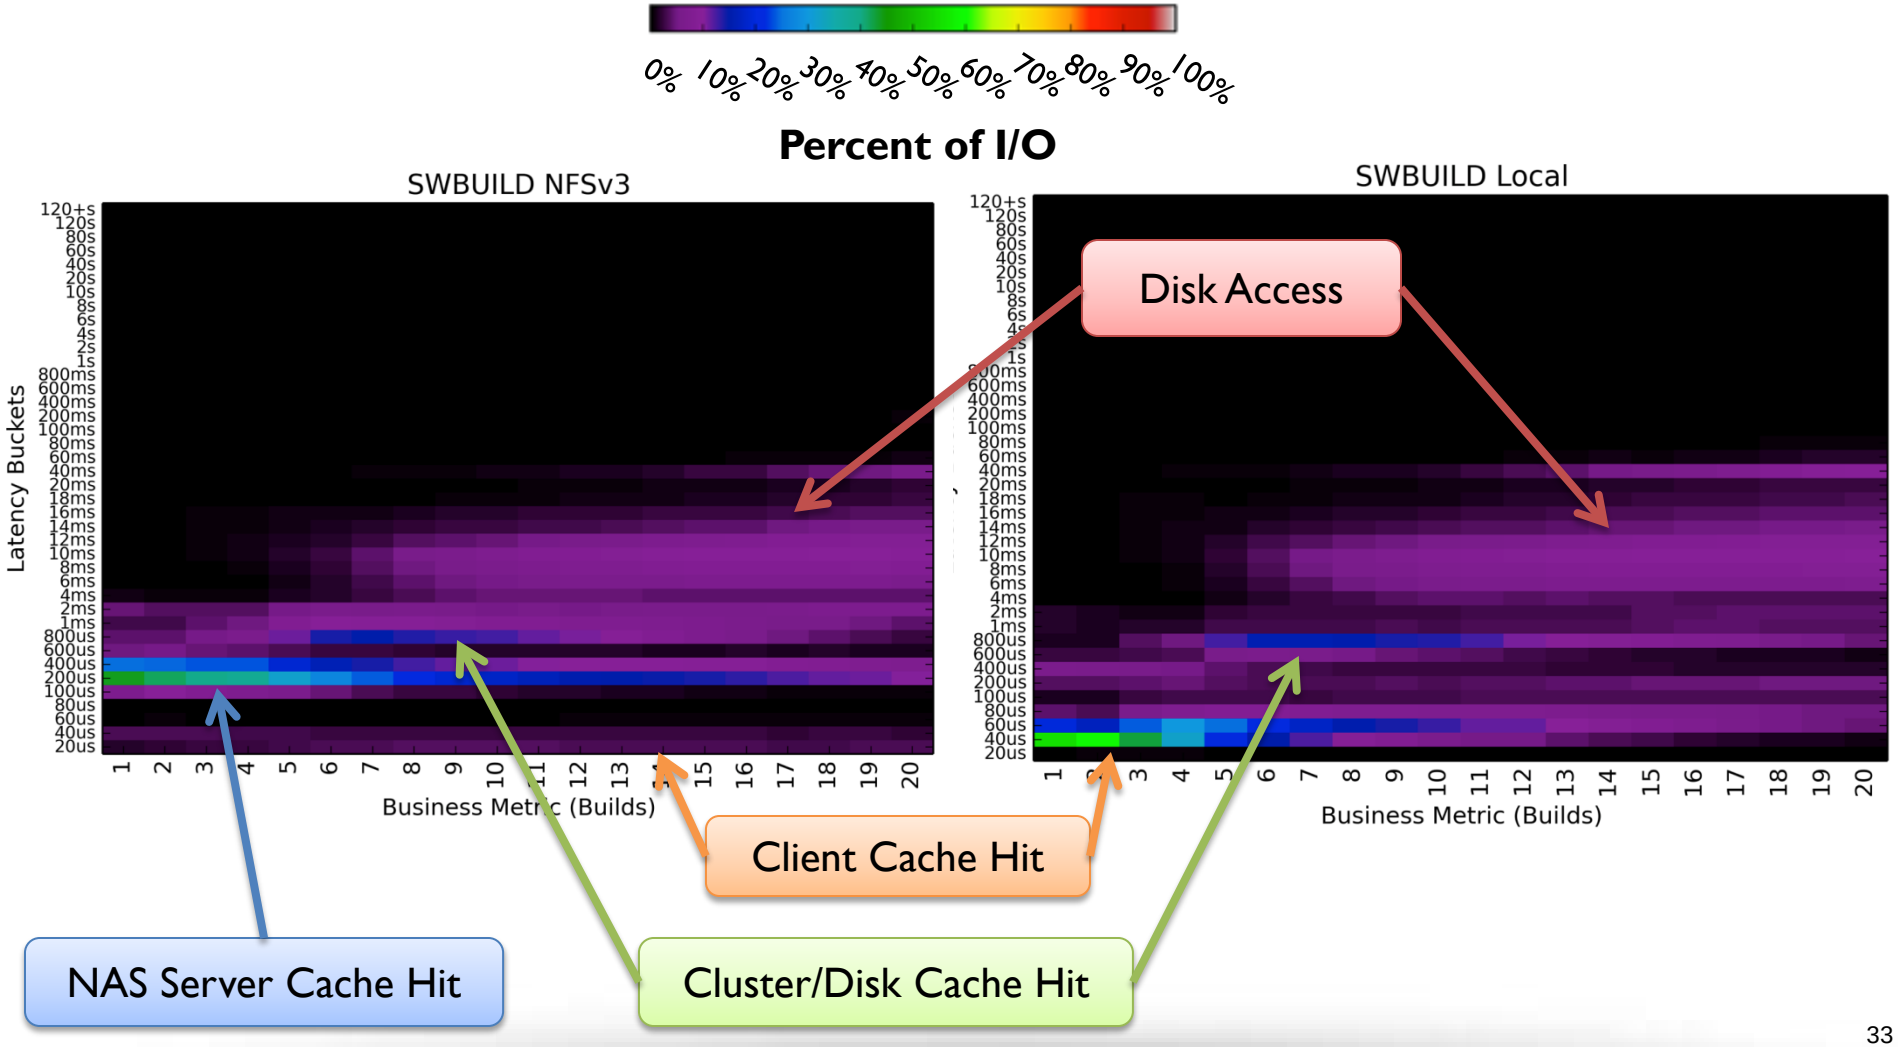

Key takeaways

- The SPEC SFS 2014 is an application-level benchmark that tests the storage performance of an entire storage solution
- Understanding the storage solution under test and bottleneck placement are keys to getting “what you expect” from your storage solution
- The application-level benchmarking provided by SFS 2014 allows testing of a much wider array of products and storage solutions

Q & A

- Any questions?

- Thank you for attending!
  - Please remember to submit feedback!