Evaluating Network Attached Storage Units

Benchmarking Strategies for Home Users, SOHOs and SMBs

Ganesh T S

October 23, 2015

- Sr. Editor @ AnandTech
- 5+ years reviewing multimedia systems, storage and networking products
  - Communicating with NAS vendors - receiving product pitches & marketing collateral
  - Communicating with readers - both power users and SMB IT administrators
  - First-hand view of NAS market evolution and features gaining market traction
- 10+ years as an ASIC verification engineer
- Coverage primarily from an engineering perspective
- UX aspects also noted

- Wide variety of NAS units to target different market segments
  - RISC-based (ARM / PowerPC) vs. x86
  - Price vs. Performance vs. Feature Set
  - Underlying file system - XFS, EXT4, btrfs, ZFS
- Success dictated by multiple external factors beyond testing results
  - Channel presence
  - Bundles - hard drives and NAS, networking gear (switches) and NAS etc.
  - Word of mouth - reliability and ease of use

- Subjective metrics
  - Management interface UX
  - Quality of mobile app(s)
  - Value added services
    - Media server
    - Dropbox-like sync, backup and replication support
    - Extensibility - third-party apps, virtualization / containers
    - Private cloud - Google Docs / MS Office 365 replacement
    - Centralized management
- Objective metrics
  - Transfer rates, latency, response times
  - Power efficiency
  - Failure handling

- NAS connected to a power measurement device (Ubiquiti Networks mFi mPower Pro)
- Operated in a diskless state, followed by initializing of a single-disk volume, shares configuration and population with a few media files
- Disks added one by one to test out online RAID migration (from JBOD to RAID-1 to RAID-5)
- Collected metrics
  - Power consumption at the wall every second for each operation
  - Time taken for each operation
  - Extent of disruption in streaming of stored media during each operation

- Simulating disk failure
  - Randomly selected member disk physically pulled out of the NAS during read / write operation
  - Ensure no disruption in data access despite physical failure
  - Fresh disk plugged in as replacement
  - Power consumption / time taken for rebuild process recorded
- Evaluating data recovery options when the NAS hardware fails, but disks are OK
  - Connect disks to a PC using JBOD DAS
  - Use Ubuntu + mdadm or Windows + commercial software like UFS Explorer to access and save data
  - Chance for NAS vendors to differentiate

- Benchmarking with the OS file copy utility for various sets of files - subset of a very limited use-case
- Workload generation - single client vs. multi-client with switch
- Workloads - artificial vs. real-world traces
- Options for benchmarking
  - IOMeter, IOZone, fio, Vdbench etc.
  - Load testing tools - Login VSI, hIOmon Disk I/O Ranger
  - Appliances - LoadDynamiX
  - Intel NAS Performance Toolkit (NASPT)
  - SPEC SFS 2014
- Very important to keep storage media consistent across NAS units for benchmark comparisons

- Most NAS units marketed with transfer rates, but easy to saturate network links with appropriate artificial workload traces
- Increasing prevalance of multiple clients (streaming devices, IP cameras etc.)
- Only high-end SMB NAS units marketed with IOPS numbers - not easy to convey importance to home users
- Many workloads run out of IOPS long before bandwidth saturates - storage media is important

- Focus on evaluation of CIFS, NFS and iSCSI features across different client platforms
- Used only Intel NASPT in the early days
  - Evaluates CIFS and iSCSI single client performance for real-life workloads
  - Supplied traces include video streaming, recording, office productivity, photo album viewing etc.
  - Determining performance penalties for encrypted volumes / shares
- NFS evaluation using IOZone and a CentOS client
- Unfortunately, not great choices for multi-client scenarios and business workloads

Moved to multi-client testing in late 2012..

- Testbed hardware - 2P Xeon-based system with 6x (1Gbps x 4) PCIe network adapters
- Switch configuration depending on solution under test

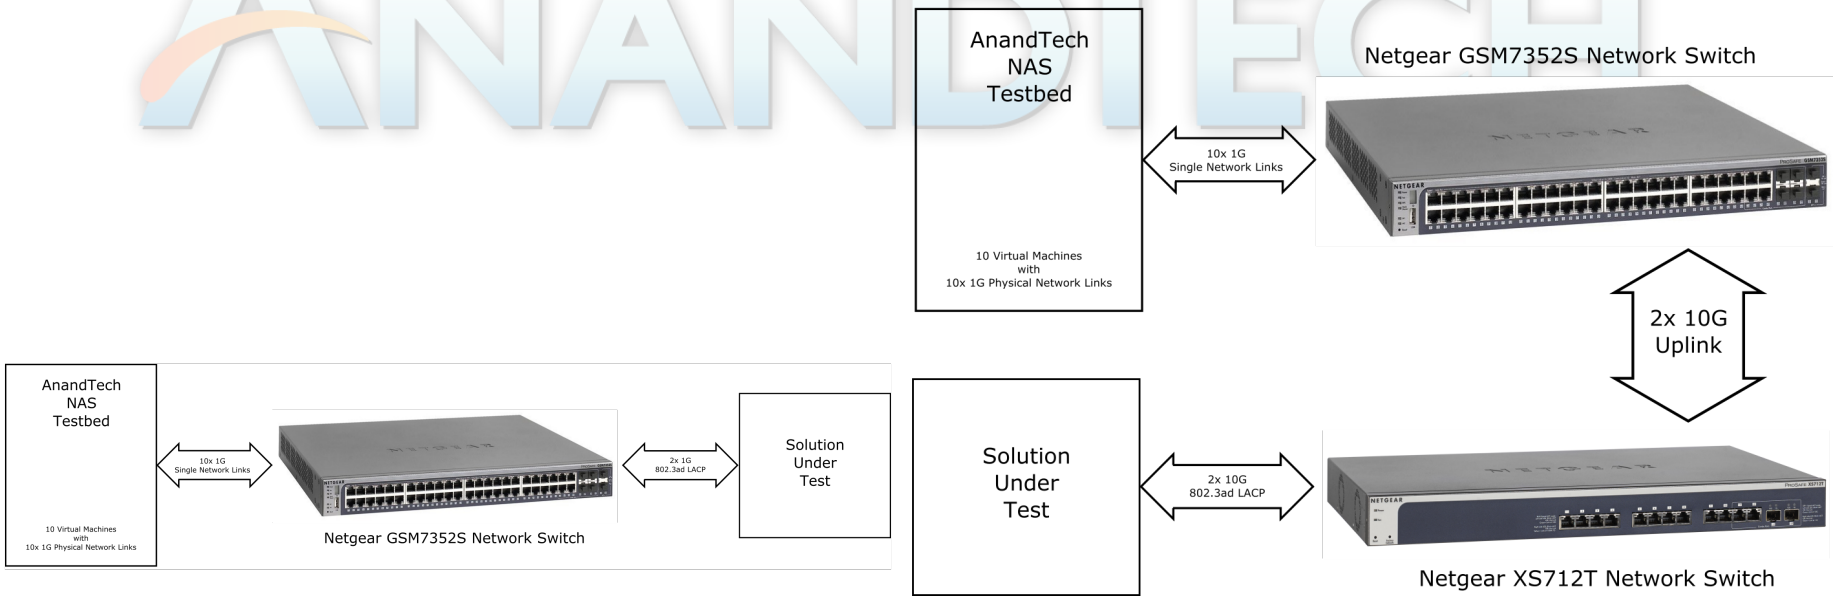

- Transfer rates and latencies from IOMeter with artificial workloads for benchmarking
  - Max. Throughput (Sequential) - 100% Reads
  - Max. Throughput (Sequential) - 50% Reads
  - Random 8K - 70% Reads
  - Real Life - 60% Random, 65% Reads
- Benchmark numbers difficult for end-users to relate to
- Doesn't answer typical end-user questions
  - How many simultaneous videos can the NAS stream out?
  - How many users can work on spreadsheets and documents directly off the NAS simultaneously?
  - How many IP cameras can record to the NAS at the same time reliably?
- New approach based on the concept of business metrics

- NAS Performance Toolkit - released by Intel in 2007
  - Ships with a library of traces representing home usage scenarios
  - Traces replayed on target devices and response times / performance metrics recorded
  - Includes visualizer for dissecting results
- EOL software, but source code available
- Very susceptible to client caching, requires RAM limiting / source code modifications for running on modern systems
- Important to ensure client is same across different evaluations

Test	# files	% seq.	Bytes Rd/Wr	Description
HD Video Play	1	99.5%	2.0GB Rd	256kB reads
HD Video Record	1	99.9%	2.0 GB Wr	256kB writes
Directory Copy From NAS	2833	52.5%	0.20 GB Rd	64kB reads
Directory Copy To NAS	2833	52.5%	70B Rd 0.25GB Wr	Predominantly 64kB writes, wide scattering under 16kB
File Copy From NAS	1	100%	4.3GB Rd	64kB reads
File Copy To NAS	1	100%	4.3GB Wr	64kB writes
Photo Album	169	80%	0.81GB	All reads – wide distribution of sizes
Office Productivity	607	81.3%	1.4GB Rd 1.4GB Wr	Reads & writes; small, 1kB & 4kB reads; Mostly 1kB writes
Content Creation	98	38.6%	12MB Rd 14MB Wr	95% writes; 1k, 4k & little reads; Writes up to 64kB

Additional workloads (extensions of the above) available . . .

- Focus of NASPT on single-user performance, workload traces are a bit dated
- Trace replay component is still a valuable resource
- Taking advantage of NASPT for multi-client scenarios
  - Reuse existing workload traces and trace replay program
  - Add wrapper to farm out and synchronize across multiple clients
  - Run each workload on multiple clients simultaneously
- Determine number of clients that can provide acceptable performance
  - Failure to scale throughput linearly without saturating link
  - Sudden spike in average response times

Netgear ReadyNAS RN202 - Folder Copy from NAS

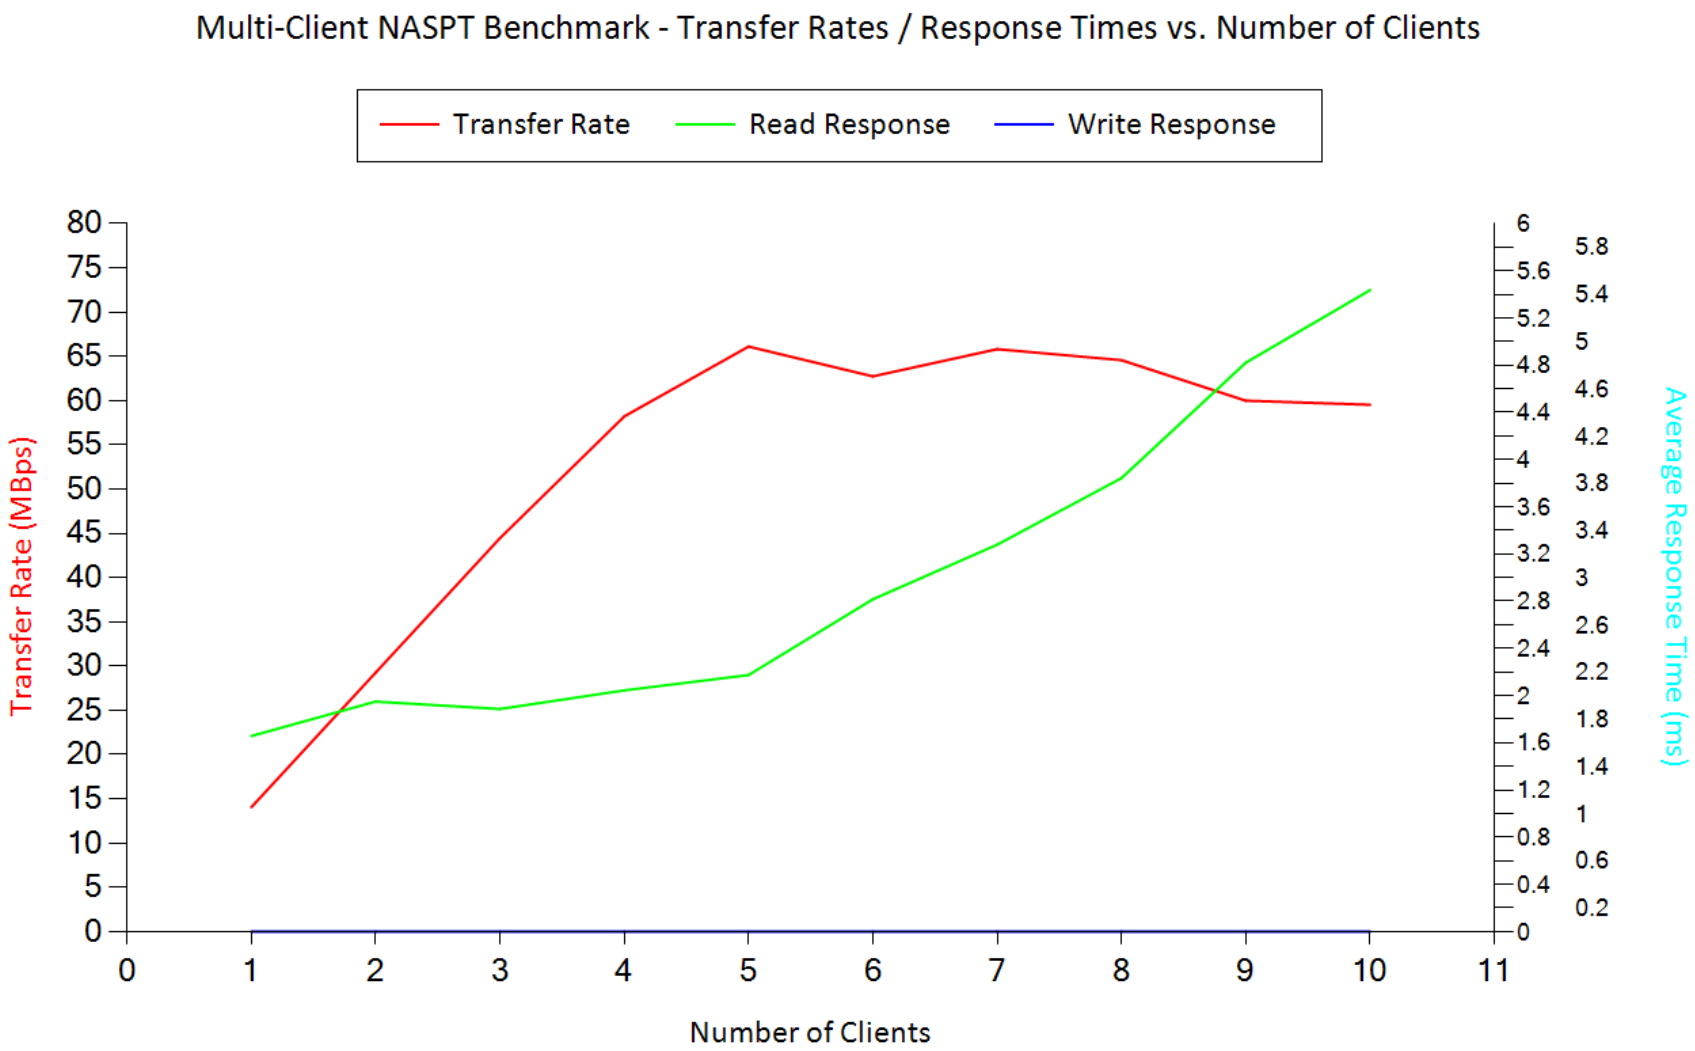

Standard Performance Evaluation Corporation SFS Benchmark

- SPEC benchmarks - long-standing, accepted industry-wide
- SFS (Solution File Server) benchmarks since 1993 - initially NFS-only
- SPEC SFS 2014 - multi-platform / CIFS & NFS evaluation
- Benchmark binaries & source (based on IOZone) licensed from SPEC
- Workloads beased on real-world application traces, measures quality of service
- Simulates multi-client workloads, records op rate, throughput and response times

Real-life Workload Traces

- Video Data Acquisition (IP cameras)
  - VDA1 - high bitrate sequential writes
  - VDA2 - companion applications / user access
- Virtual Desktop Infrastructure (hypervisors)
  - VDI - data-heavy workload, direct I/O, compressible large files
- Software Build (software project compilation)
  - SWBUILD - reads & writes to 573K highly compressible files
- Database (OLTP database consolidation scenario)
  - DB_TABLE - Random reads & writes to same dataset from multiple threads
  - DB_LOG - Mostly sequential writes

- Metrics collected in each load run
  - Average latency
  - Per-process oprate
  - Read and write throughputs
- Final results measured in business metrics
  - VDA: Number of concurrent STREAMS
  - VDI: Number of concurrent DESKTOPS
  - SWBUILD: Number of concurrent BUILDS
  - DB: Number of concurrent DATABASES
- Success criteria at each load point
  - Per-process oprate
  - Overall oprate
  - Component workload variance
- Publishable results require success with at least 10 load points
- Most small-scale NAS units with 7200 RPM SATA drives fail oprate criteria

- Software investigation
  - Reduce caching impacts / force direct I/O in NASPT
  - Generate new NASPT traces for contemporary workloads
  - Formalise *metrics* determination for NASPT workloads
  - Investigate / minimize hypervisor effects on performance
  - Move to Windows 8+ VMs for SMB 3.x testing
- Hardware upgrades
  - Residential lab setting - always looking to drive down power and noise
  - Increase number of virtual machines - looking at Xeon-D based systems with a 10G switch
- Monitor feedback from NAS vendors and readers

Further Reading

- Intel NASPT - User Guides, Whitepaper, IDF Presentation
- SPEC SFS 2014 Benchmark Home Page
- AnandTech NAS Coverage - News & Reviews

Q & A

Thank You!