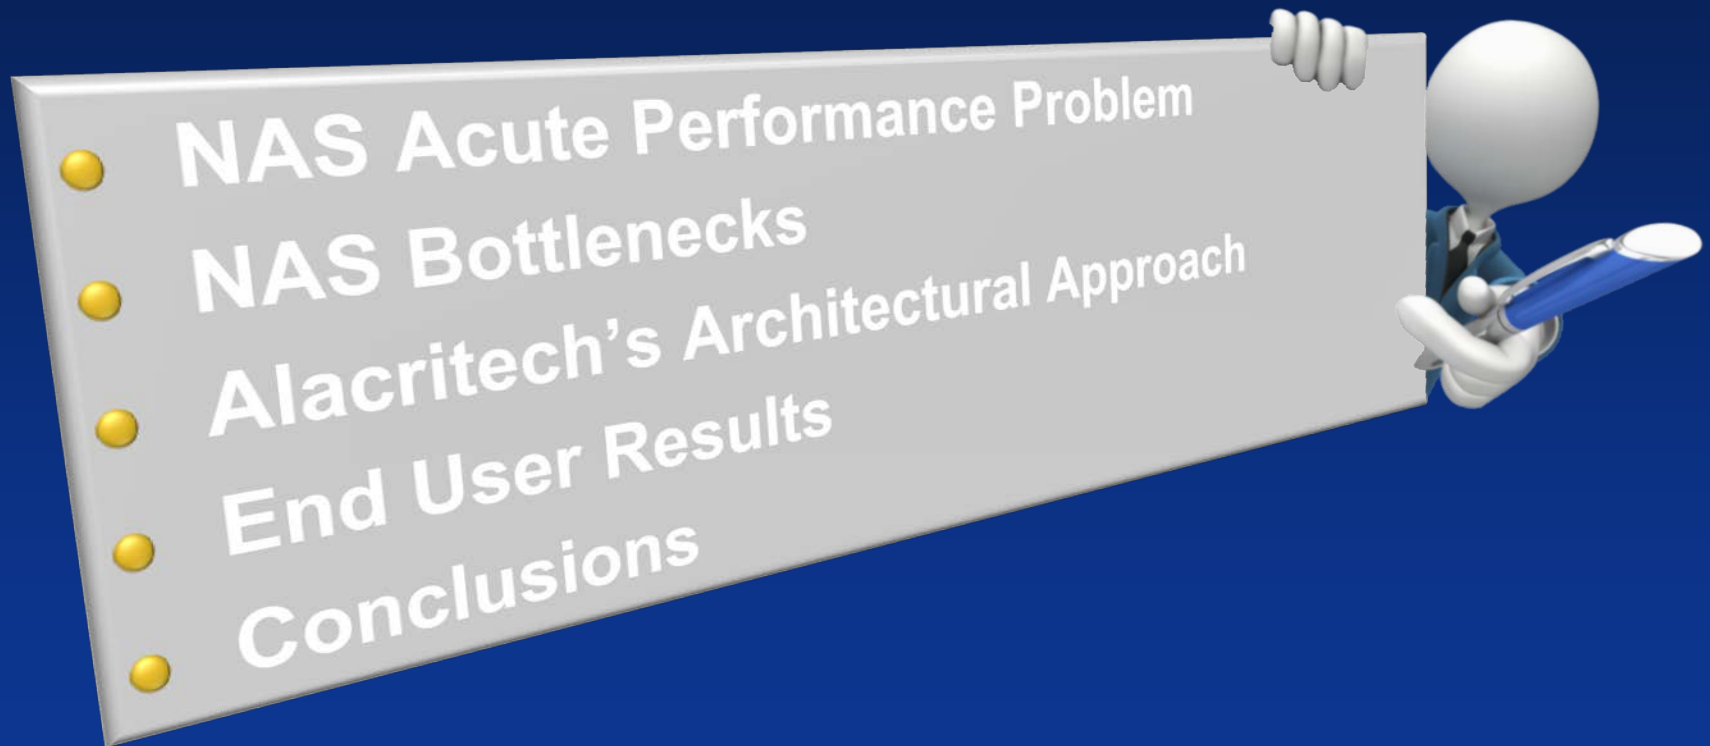# Boost Performance and Extend NAS Life

Doug Rainbolt

Vice President of Marketing

Alacritech, Inc.

Agenda

- NAS Acute Performance Problem
- NAS Bottlenecks
- Alacritech's Architectural Approach
- End User Results
- Conclusions

NAS Acute Performance Problem

- Shows up as increased client response times
  - Reads & writes
  - Root cause…*LATENCY* or delay
    - Leading to frustration & angst
    - Causing much lower productivity
    - Especially when intermittent

From a NAS Admin Perspective

- Real NAS admin quotes
  - Users complain until fixed
  - Managers complain until the users stop complaining
  - Phones light up like a Christmas tree
  - Vendors blame everyone else but themselves

3 Typical Primary NAS Latency Bottlenecks

- NAS controller utilization

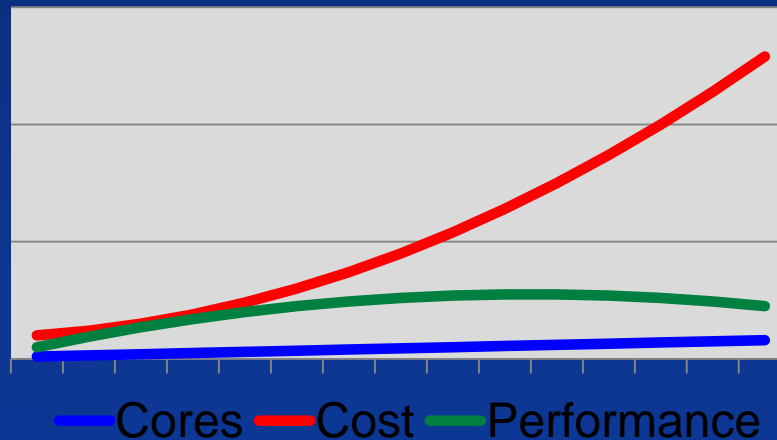- Hard Disk Drives (HDDs)

- NFS metadata handling

NAS Controller Utilization Bottleneck

- NAS required to do a lot more than just file serving
  - Including but not limited to:
    - Snapshot, snap rollback, mirroring, de-duplication, thin-provisioning, replication, virtual infrastructure integrations, backup software integration, continuous data storage tiering, RAID 6, bit scrubbing, & file serving
  - Spiking controller utilization
    - Noticeably increasing latency & response times…intermittently
   - Escalating client frustrations & reduced productivity

Common NAS Controller Utilization Workarounds

- More CPUs, cores, or controller nodes
  - Has serious shortcomings
  - Called "**Law of Diminishing Marginal Returns**"
    - Each add'l CPU, core, or controller node adds less than one before
    - Eventually next addition actually reduces performance
  - Significantly increases cost of IO

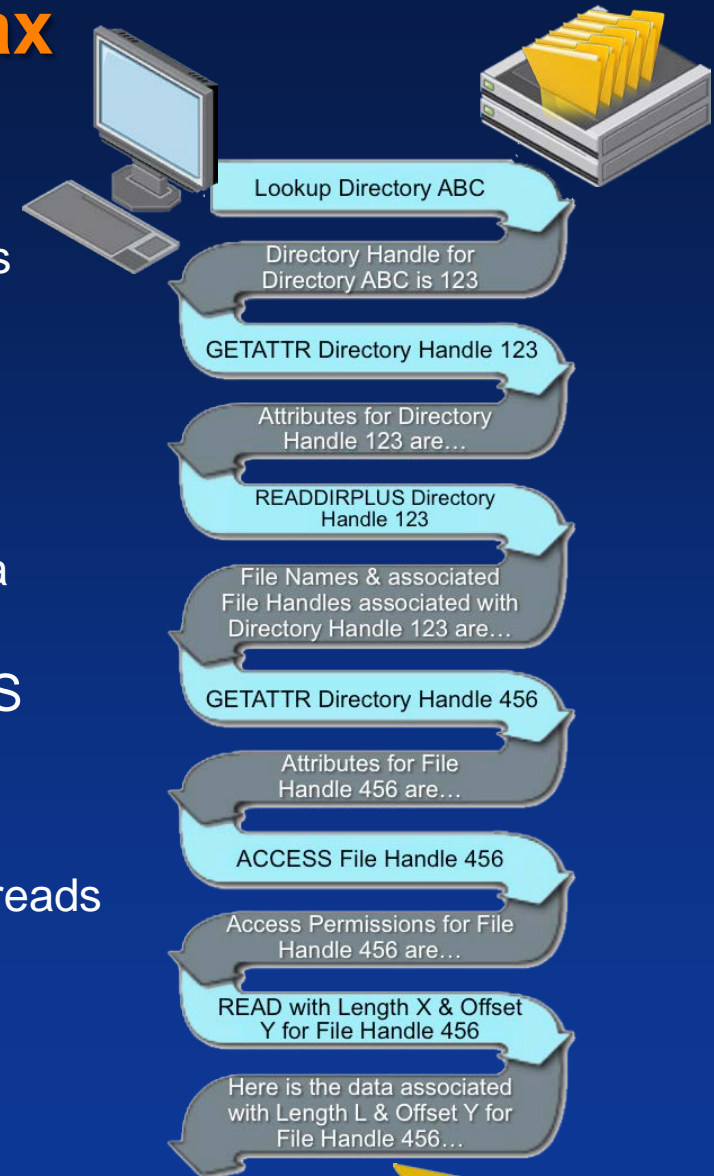Cores — Cost — Performance

Hard Disk Drives (HDDs)

- Huge gap between HDD & processing performance
  - Processing performance doubles every 18 mos to 2 yrs.
    - Moore's Law
    - While HDD performance stays flat
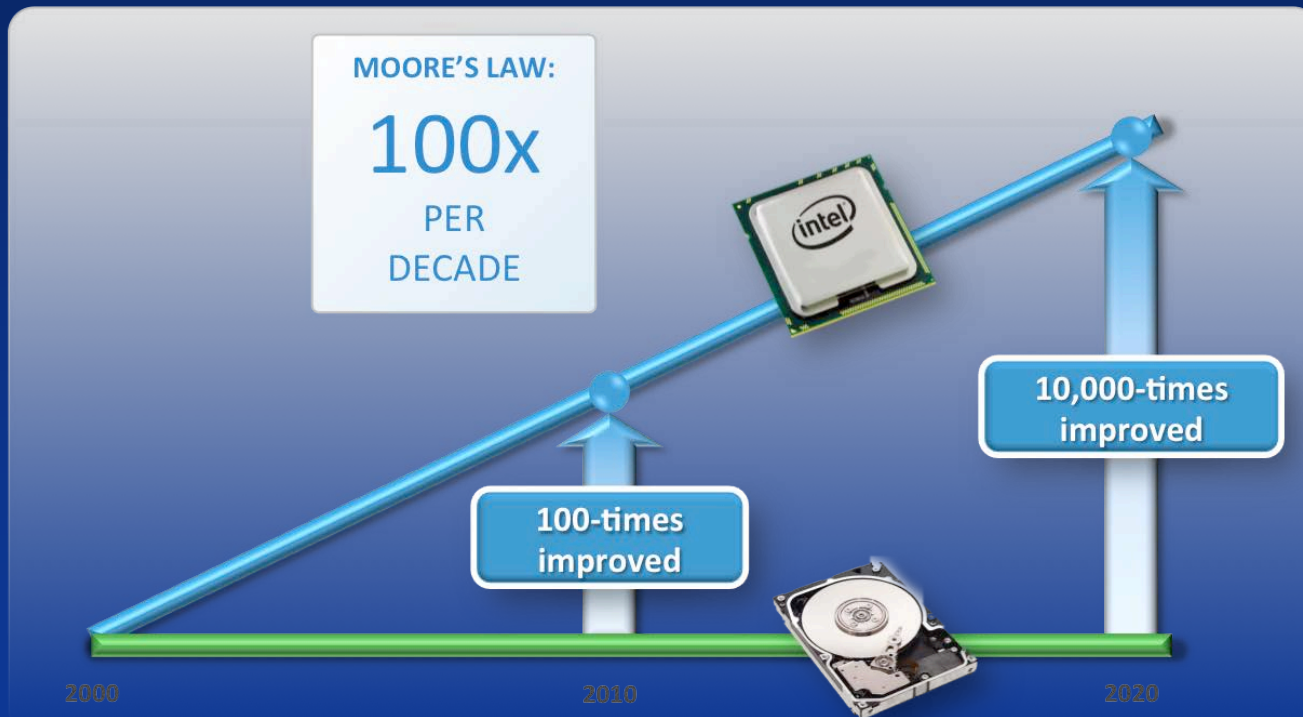
NFS Metadata NAS Performance Tax

- NFS chatty stateless protocol
  - Lots of client/server info exchange for file status
    - When a file was last modified
    - Permission accesses
    - Pathnames stored in symbolic links
- Each NFS conversation adds multiple RTs
  - Equals 100s of 1000s to millions NFS metadata ops/s
- NFS metadata ops easily => 80 to 90% of NFS total
- Clients must know workflow shared file status
  - Ensuring coherency especially onerous status reads
    - 100s to 1000s of clients
    - Huge file #s & gridlocked NAS traffic
    - Severely impacts performance
    - Increased response times

Lookup Directory ABC

Directory Handle for Directory ABC is 123

GETATTR Directory Handle 123

Attributes for Directory Handle 123 are…

READDIRPLUS Directory Handle 123

File Names & associated File Handles associated with Directory Handle 123 are…

GETATTR Directory Handle 456

Attributes for File Handle 456 are…

ACCESS File Handle 456

Access Permissions for File Handle 456 are…

READ with Length X & Offset Y for File Handle 456

Here is the data associated with Length L & Offset Y for File Handle 456…

Alacritech Vision

- Enabling maximum application performance gains with minimal infrastructure or ecosystem changes
  - i.e. Empowering optimal application performance via the smallest possible least-intrusive "**linchpin**"

Why This Vision

- Current chronic market problem
  - Improving application performance requires major changes
  - Typically requires "**Ripping Out and Replacing**"
    - Servers, networks, storage, and more
   - Causing massive disruptions with excessive scheduled downtime

Alacritech ANX

## The Cure for the NAS Performance Blues

Alacritech's ANX Solution

- Significantly improves NAS performance by quantifiably < latency
  - Frees up NAS system controller cycles
  - Increases NAS usable capacity
  - Solves the NFS metadata handling problem
 - None of the drawbacks of the workarounds
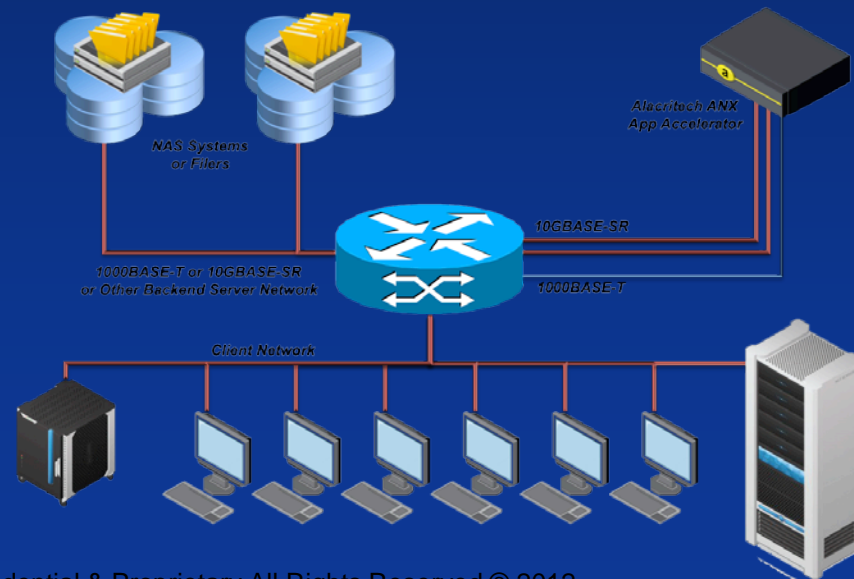
How ANX Solves
The NAS Performance Blues

- Leverages proven unique Alacritech acceleration technology
  - 2 on-board custom ASIC Alacritech 10GbE accelerators
    - Highly efficient data transfer
  - Customized NFS processing & TCP offload
    - High throughputs & low latencies
    - Minimal impact on local processors
- Large Flash caching capacity
  - 2 or 4TB SSD capacity + 48GB of DRAM
  - Serves most working sets in the cache
- NAS transparent write-through appliance
  - Captures hot data on reads
 - Data ownership remains with NAS
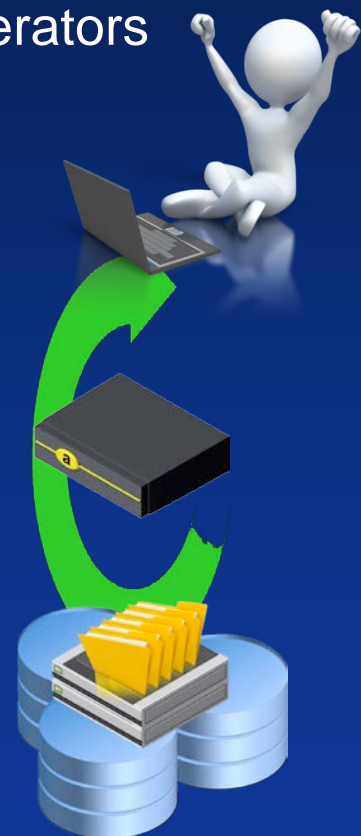
ANX 1500 Architecture

- Five components of ANX 1500

1. Forty eight GB DRAM

2. Solid state disks
   - 2TB – 10 X 200GB
   - 4TB – 20 X 200GB

3. Dual 10GbE TCP offload NICs

4. Conventional Hardware
   - Dual Intel Quad Core Xeon, 5500s
   - Dual auto switch power supplies
   - Quad 10/100/1000 BASE-T (management connectivity)

5. Software
   - Alacritech NFS Bridge Patented Technology
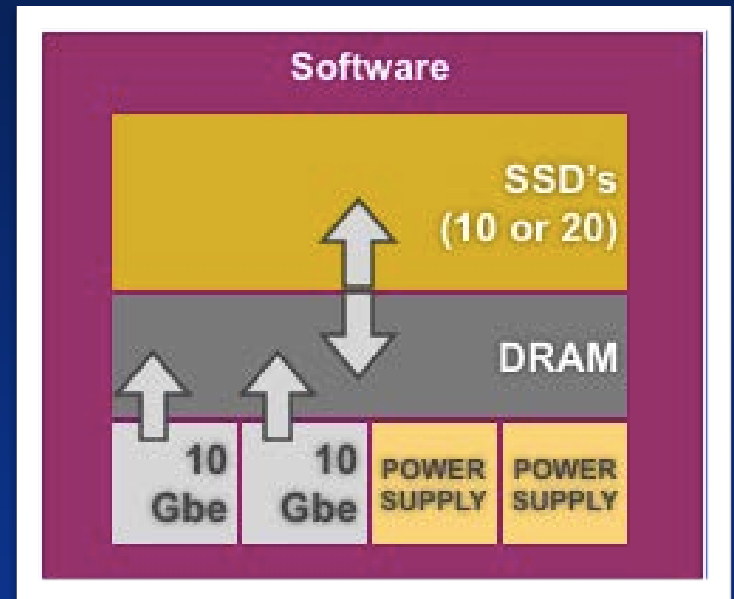
ANX NFS Metadata Latency Tax Cut

- Acts as a NAS proxy
  - Offloads NFS metadata requests for active data
  - Caching & accelerating metadata responses
  - Can process ~ 2.75 million NFS metadata ops/s
  - Passes WRITE requests directly to NAS systems
  - NAS freed up to process non-cacheable requests faster
  - Quieting the NFS metadata chatter = much faster NAS response times

90%　　　　**10%**

NFS Metadata Requests

NFS Metadata Fulfillment

File updates
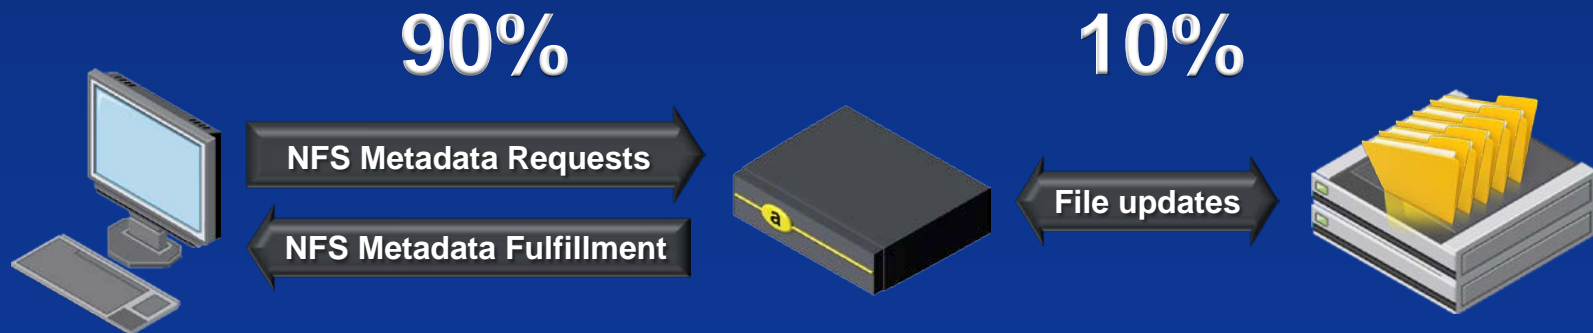
ANX's New NAS Caching Paradigm

- Removes Flash cache – NAS or Server tight coupling
  - Sits between clients & NAS
  - Dedicated NAS performance layer separate from its capacity
  - Multiple ANXes can sit between clients & NAS
  - Each ANX supports multiple clients & NAS systems
    - Same or different vendors
  - Eliminates majority of cache misses
  - Reduces NAS controller resource utilization
    - Increasing NAS usable capacity
  - No "Rip-Out-And-Replace" HW or SW replacements
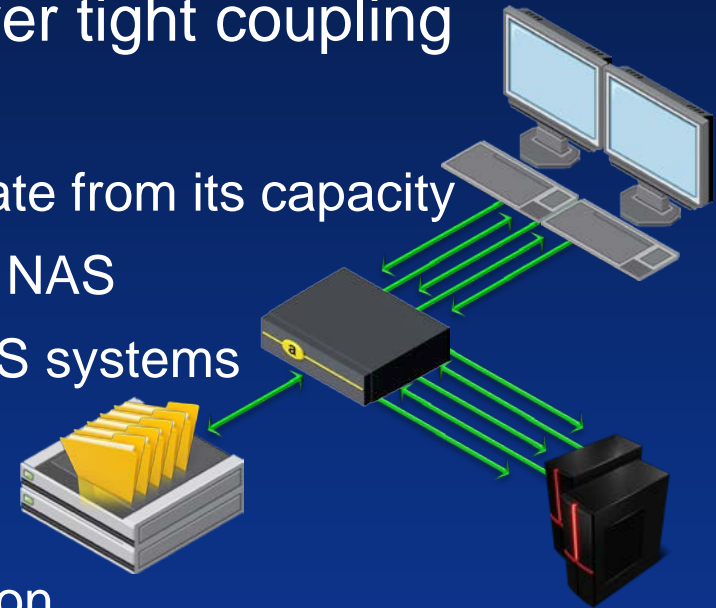
ANX Has an Incredible Impact

- Media & Entertainment
  - Rendering
- Oil & Gas
  - Seismic Travel Time Tables
- Life Sciences
  - Protein decoding
  - Genome decoding
- Pharmaceuticals
  - Molecule testing/matching
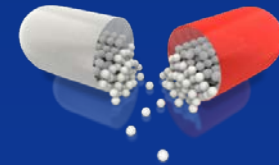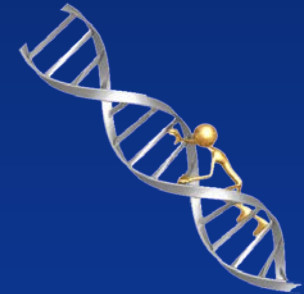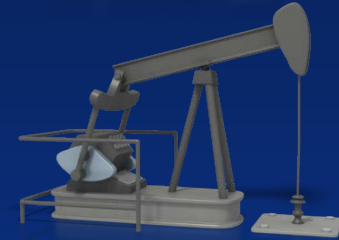
ANX Offloads Backend Filers
Proof Statement

Per your request, here are the latest stats after switching back the production mount point last Friday morning. The total of 11.3 billion OPS and hit rate of 91% during that time is pretty impressive.

ANX Statistics

```
atkstat  -c 2>&1
                         == CACHE STATISTICS ==
                TOTAL            HIT        MISS      HITRATE        STALE      SSD-HIT
------------------------------------------------------------------------------------
NULL:           256767        256767           0         100%            0            0
GETATTR:    6502156869    6294713363   207443506          96%            0            0
SETATTR:      88353023             -           -            -            -            -
LOOKUP:     2327593808    2204666666    53451863          97%     48400121       471888
ACCESS:      977259796     968735523     8524273          99%            0            0
READLINK:      2321293       2320157        1136          99%       632762            0
READ:       1020882641     930926141    89956500          91%      1441256     49825456
WRITE:       171749490             -           -            -            -            -
CREATE:       40590771             -           -            -            -            -
MKDIR:         5984575             -           -            -            -            -
SYMLINK:       4695301             -           -            -            -            -
MKNOD:               0             -           -            -            -            -
REMOVE:       39691731             -
RMDIR:         5455856             -
RENAME:        7264230             -
LINK:             2084             -
READDIR:      10631173             -
READDIRPLUS: 137702515             -
FSSTAT:        3288069             -            -            -            -            -
FSINFO:         248497             -            -            -            -            -
PATHCONF:         1598             -            -            -            -            -
COMMIT:         638743             -            -            -            -            -
------------------------------------------------------------------------------------
TOTAL:     11346768830   10401618617   359377278          91%     50474139     50297344
```
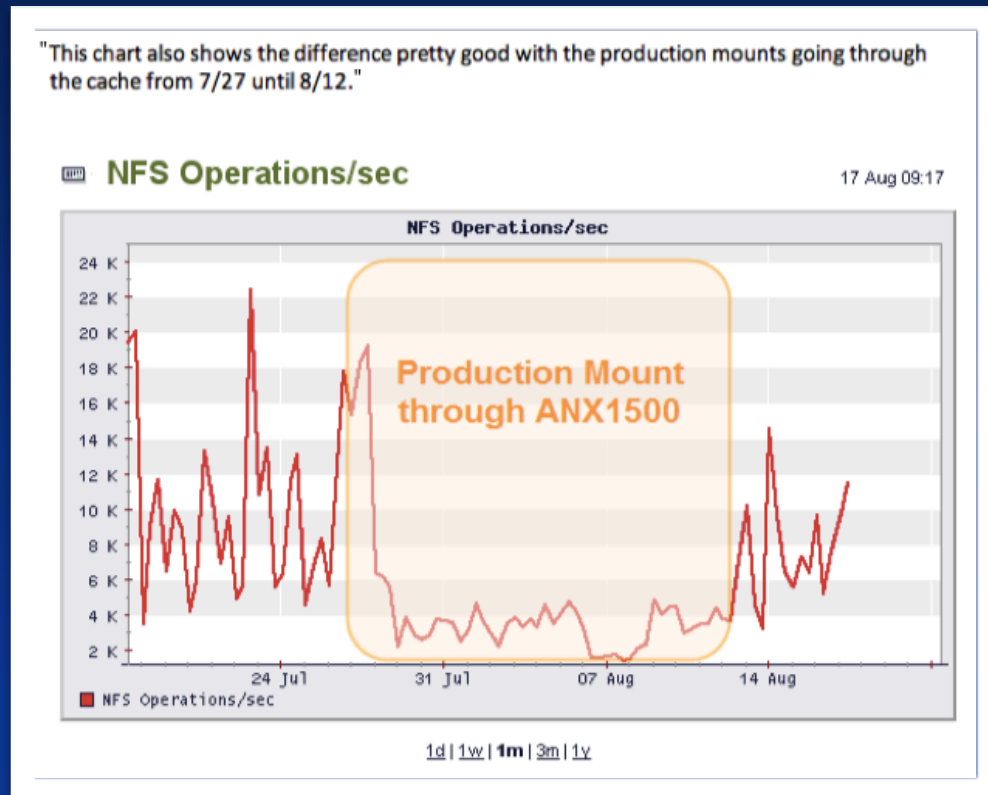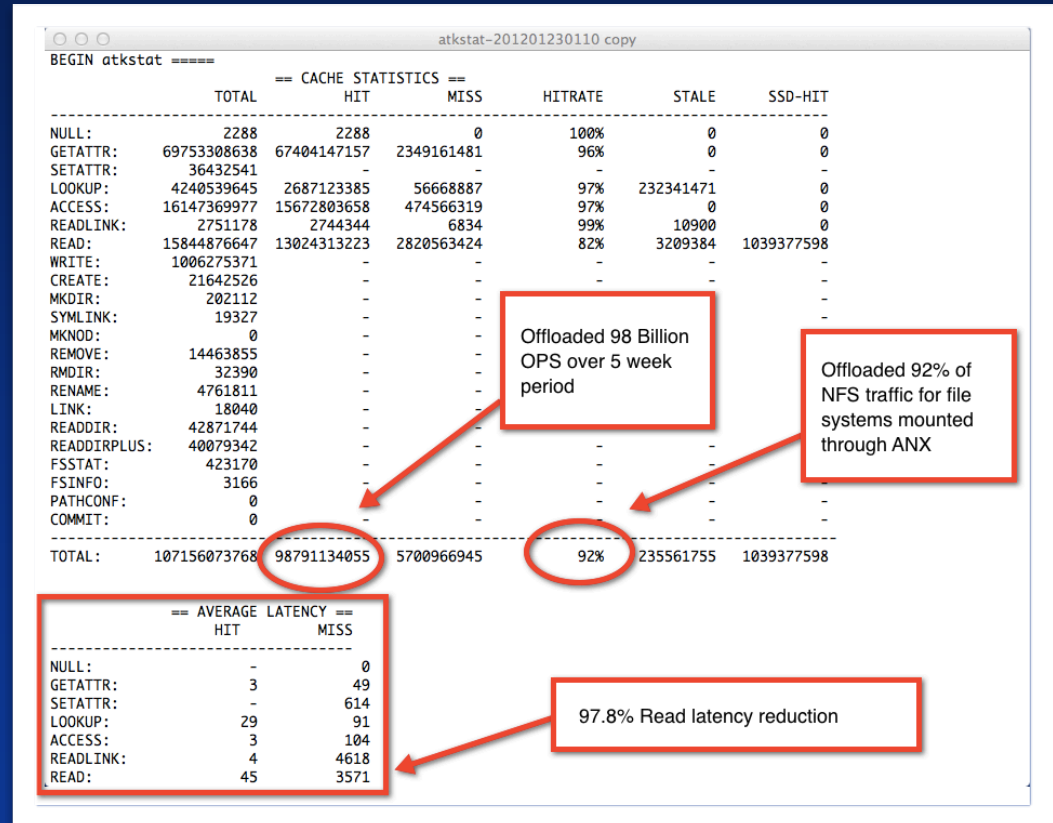
10.4B OPS Offloaded from Backend Filer

19

- Operation Manager graph shows:
  - ANX significantly offloading backend NetApp Filer
  - Frees up resources on filer for things like writes and CIFS
- Storage manager commented that CIFS OPS dramatically increased on backend filer after introduction of ANX 1500

"This chart also shows the difference pretty good with the production mounts going through the cache from 7/27 until 8/12."

NFS Operations/sec 17 Aug 09:17

Production Mount through ANX1500

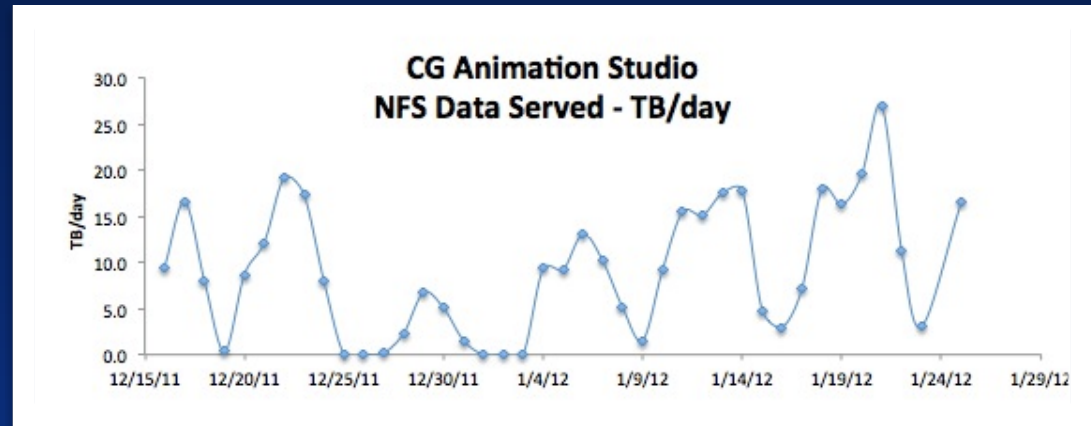NFS Operations/sec

1d | 1w | **1m** | 3m | 1y

Rainmaker Entertainment

- 400 render nodes - in process of doubling to 800

- Was planning to buy third filer head to handle load as existing two were CPU bound....

- Purchased ANX HA Pair instead as ANX offloaded 90% of NFS traffic mounted through appliance

- Weighted average latency reduction of 80% across all 22 NFS operations

- 7TBs of NFS data accelerated by cache on average each day

- 27TBs of NFS data accelerated by cache on peak day

- Offloaded 98.7 billion NFS OPS during 5 week period

CG Animation Studio
NFS Data Served - TB/day

- CPU utilization on filers dropped from 100% to 60%

- Network traffic to NAS decreased 35%

- Disk OPS to NAS dropped from 160K to 70K
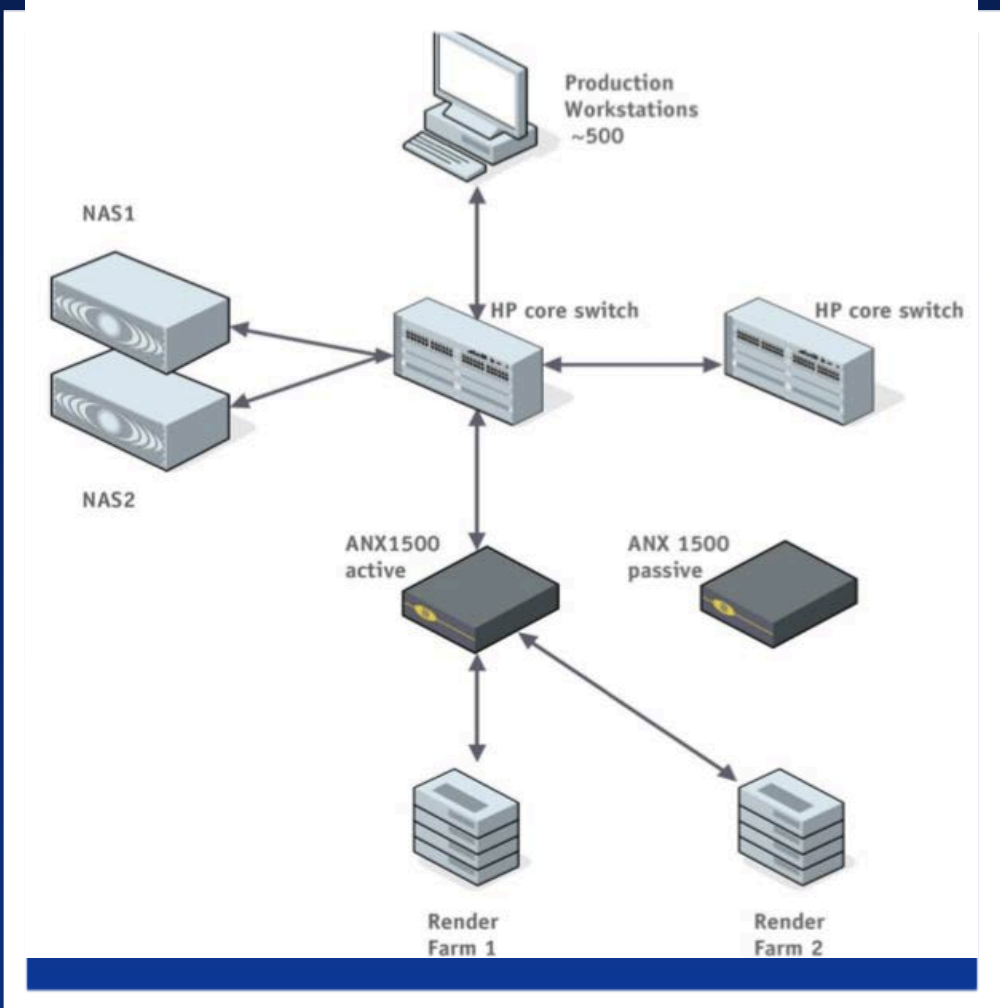
Conclusions

- NAS performance is a problem
- Root cause is latency & metadata handling
- Market offerings don't completely solve problem
  - And do nothing for NFS Metadata handling
- Alacritech ANX solves NAS performance problem
  - By directly attacking the root causes
  - Solving both the latency & NFS metadata problems
- Alacritech ANX cures the NAS performance blues

Questions?

For More Information

- Email: info@alacritech.com
- Website: www.alacritech.com