A Taxonomy of NFS v4 Performance

Tom Haynes
Darrell Suggs

Overview

- State of v4 for Performance Testing
- Performance Goals of v4
- Challenges of v4 Performance Dimensions
- v3 vs v4 Performance Matrix Outline
- Some early performance returns
- Points of performance leverage
- What we have not measured
- Summary and Q&A

Network Appliance - Tom Haynes / Darrell Suggs

A Guiding Principle

- **Why we care**

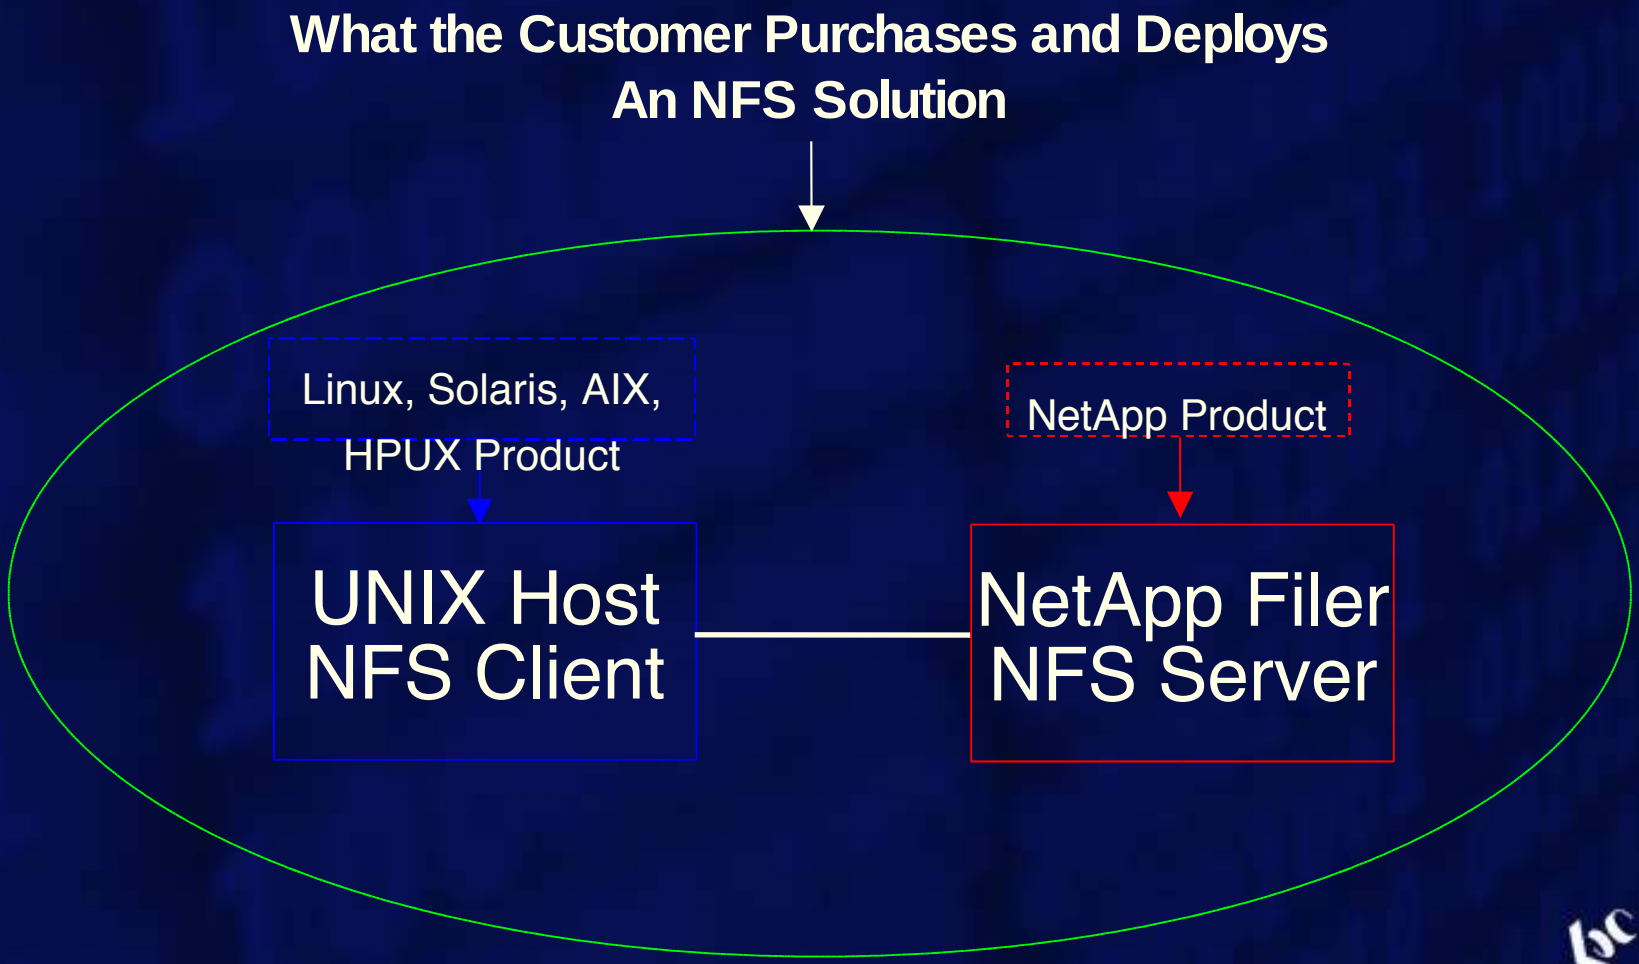

Status of v4 for Performance Testing

- Multiple clients and servers available
- Mostly functional
- Some issues under heavy stress
- Definitely ready for performance comparisons

Performance Goals of NFS v4

- First and most important v4 is about features
  - Considerable new functionality
- But performance is also important
- Primary performance improvement opportunities
  - Compound operations
  - Delegations
  - General code path enhancements
  - Other misc stuff

Network Appliance - Tom Haynes / Darrell Suggs

Performance Goals of NFS v4 (cont)

- General: Performance parity with NFS v3
  - Customers should not see degradation
- Some workloads may see large improvement
- Reality
  - Typical performance challenges with new protocol
  - Some features have performance cost
- Questions
  - Which metrics to measure
  - Which workloads to use

Challenges of Measuring v4 Performance

- What is an NFS v4 OP?
  - In v3, operation is easily identified
  - In v4, OP is more ambiguous
- Is an OP simply a compound?
  - Makes sense on a certain level
  - Easy to count
- Or do we need to count the internal ops?
  - More comparable to v3
  - More representative of “work” required on CPUs
- In ONTAP, we simply count them both
  - But which to compare to v3? Which to boast to customers?

NetApp

Challenges of Measuring v4 Performance

- We chose to use component counts
  - Not compound counts
- Why?
  - Perception
    - Compound counts << Component counts
    - Customers might view v4 as slower than v3
  - More comparable
    - Directly comparable to v3 results

v3 vs v4 Performance Dimensions

- Basic dimension of interest
  - v3 vs v4 – simply change the mount option
- Three workloads
  - Random reads
  - Random writes
  - "Metadata Operations"
    - Opens, reads, closes, locks, etc

v3 vs v4 Performance Dimensions (cont)

- Five metrics of interest
  - Host uSec/IO – client CPU needed per OP
  - Filer uSec/IO – filer CPU needed per OP
  - Throughput – Ops/Sec
  - Latency – Average access time (ms)
  - Host Ops / Filer ops
    - Think efficiency

Measurement Techniques

- I/O load generator – SIO (Simulated I/O)
- Concurrency Level
  - Set to 1 thread for Read/Write tests
    - Avoids possible queuing effects
    - More accurate comparisons for metrics
  - Set to 4 threads for the metadata test
    - Need to capture more complex actions

Early Performance Returns

- Next chart contains results of comparing two clients
  - Names are removed for simplicity
- Results
  - 1.00 is parity
  - Depending on metric > 1.00 is good or bad

Early Performance Returns

		4K Client 1	4K	4 thds	4K Client 2	4K	4thds
		RDS	WRS	META	RDS	WRS	META
host ì s/io	v3	26	25	195	177	189	2324
host ì s/io	v4	27	26	270	204	219	3211
< is better	v4/v3	1.04	1.04	1.38	1.15	1.16	1.38
filer ì s/io	v3	44	87	609	41	90	883
filer ì s/io	v4	51	85	918	48	91	1125
< is better	v4/v3	1.16	0.98	1.51	1.19	1.01	1.27
latency	v3	0.28	0.28	7.10	0.35	0.42	8.61
latency	v4	0.28	0.28	18.00	0.41	0.45	17.67
< is better	v4/v3	1.00	1.00	2.54	1.17	1.07	2.05
tput	v3	14076	13552	563	11367	9523	464
tput	v4	13894	14144	222	9876	8979	228
> is better	v4/v3	0.99	1.04	0.39	0.87	0.94	0.49
hops/fops	v3	1.0	1.0	11.1	1.0	1.0	12.0
hops/fops	v4	2.0	2.0	34.0	2.0	2.0	44.0
< is better	v4/v3	2.0	2.0	3.1	2.0	2.0	3.7

Points of Performance Leverage

- First point of leverage is clear
  - Optimizing operations in a compound
  - This will impact the other 4 metrics heavily
    - latency, throughput, host and filer CPU / op
  - Must be efficient
- Next leverage point?
  - For basic functionality, probably in code path length
  - For new and advanced functionality? TBD

What we have NOT measured

- Still early in the performance process
  - Lots of optimization opportunities
- We have not yet measured
  - High concurrency basic workloads
  - Delegations
  - Multiple host access
  - WAN performance
  - Security features

Summary

- Early performance reports are ok
  - Some optimization opportunities
  - Lots of additional results to gather
- NetApp working closely with Vendors
  - To define performance framework
  - Share testing setup
  - Jointly optimize performance

Questions and Answers ?

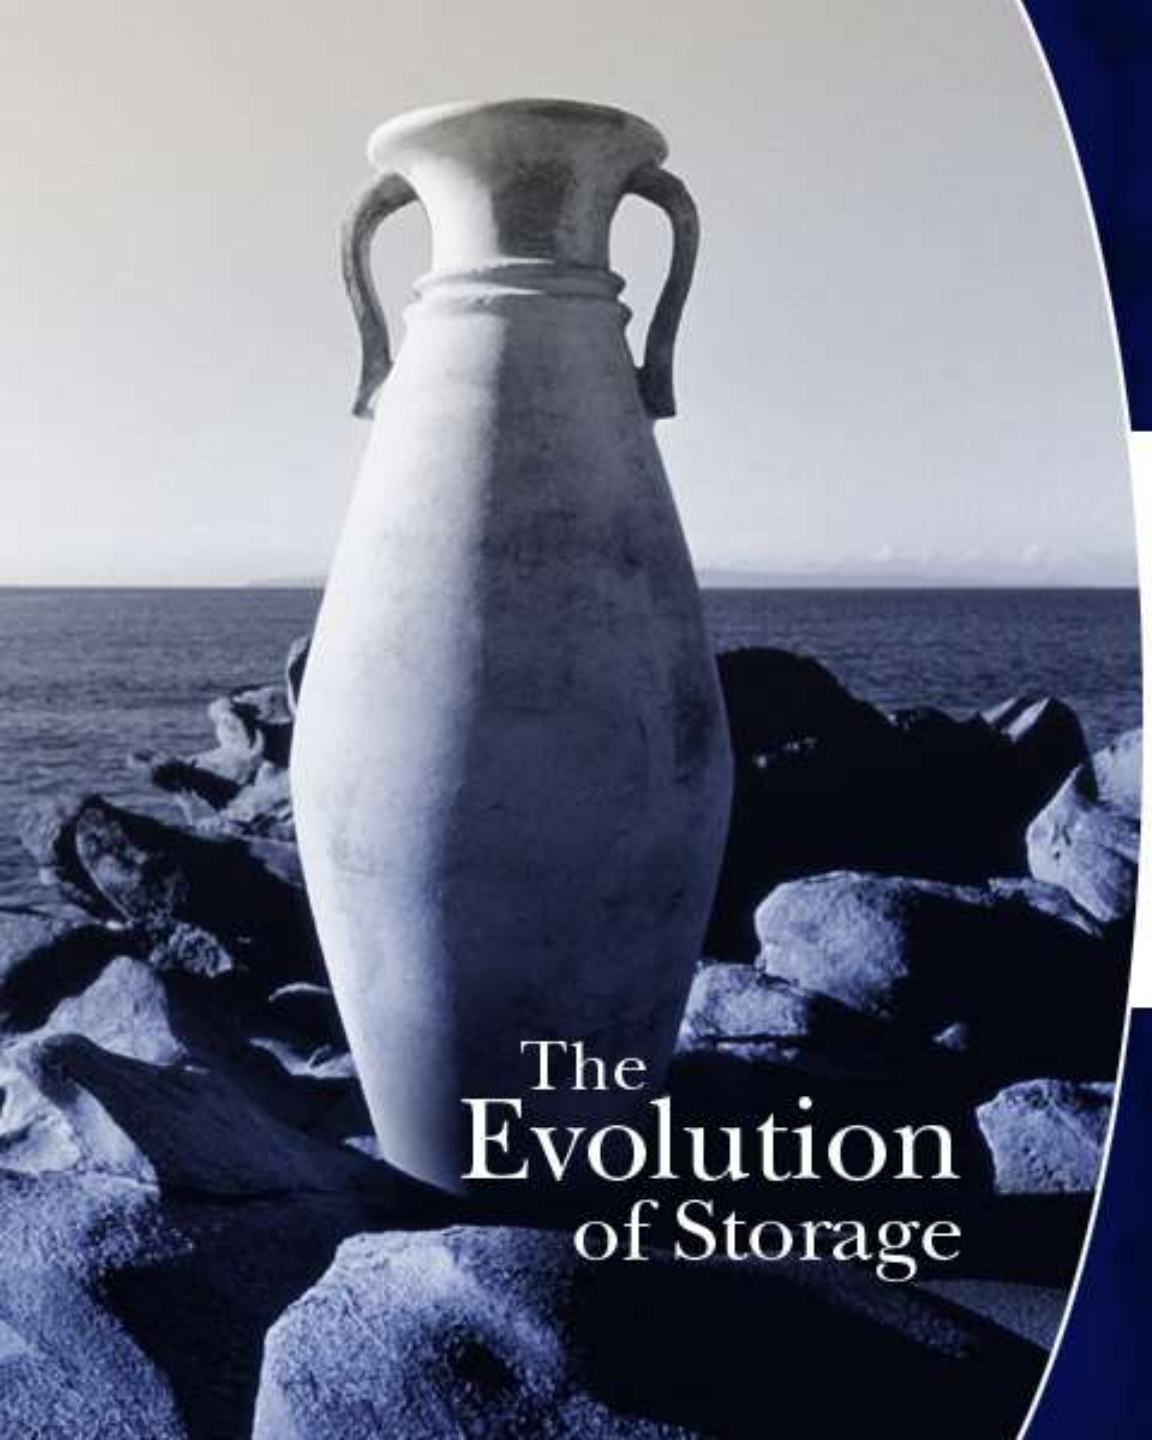

The Evolution of Storage

NetApp®