

Storage Performance Challenges and How to Stop Flying Blind in NAS Storage Cloud

Julia Palmer



Table of Contents

- Who are we?
- NAS Storage Performance overview
- Your data workload. How slow CAN you go?
- Performance analytics by using Dtrace .
...I know what you did last Friday!

Where are we coming from?

Company Facts

- There are than 47 million domain names under management.
- We also have more than 9 million customers worldwide!
- Go Daddy is the #1 Web hosting provider in the world. GoDaddy.com has more than 5 million active hosting accounts.
- Go Daddy is now the world's #1 provider of net-new SSL certificates.
- We boast 3,000+ employees with more than 250 in-house developers.
- Go Daddy maintains operations in Arizona, Colorado, Iowa, Washington, D.C., Toronto, Singapore and The Netherlands.



Go Daddy Storage Challenges

- **We grow.**
 - Exponential data growth for unstructured data (Go Daddy – 75%, Industry 30%)
- **We do more with less.**
 - High Utilization of System resources (Go Daddy 70%, Industry 40%)
 - Operational Ratio (Go Daddy 4.5PB/Admin, Industry 500TB/Admin)
- **We go fast.**
 - Go Daddy has a customer driven environment which introduces new Storage requirements: we must react quickly in order to address changing data patterns.
- **We love our customers.**
 - Our highest priority is to address any complaints concerning the customers' experience.



Data Storage Objectives

1. **Continue** to drive down cost for Production Technology Storage Solutions.
2. **Introduce** more scalable and flexible technologies to improve Storage Agility.
3. **Deploy** and Improve Operational Tools for ease of management, monitoring, and continuing to maintain our Operational excellence and 99.999% Uptime.
4. **Improve** Storage Efficiency by introducing new processes and technologies



Storage Performance

“Any sufficiently advanced technology is indistinguishable from magic.”

-Arthur C. Clarke

Sir Arthur Charles Clarke was a British science fiction author, inventor, and futurist, most famous for the novel “2001: A Space Odyssey”

Know your data:
How SLOW can you go?

Match your data &
application requirements
to your storage design

Know your data.

- Random or sequential?
- What % of metadata IOPS?
- Large or small? IO size?
- Growth rate and data change rate?
- Deduplication or compression potential?


Key Performance Characteristics

- Throughput: IOPs, Mb/sec, cache hits
- Disk utilization
- Disk Latency/NAS latency
- Application latency

How do we do it on a budget?

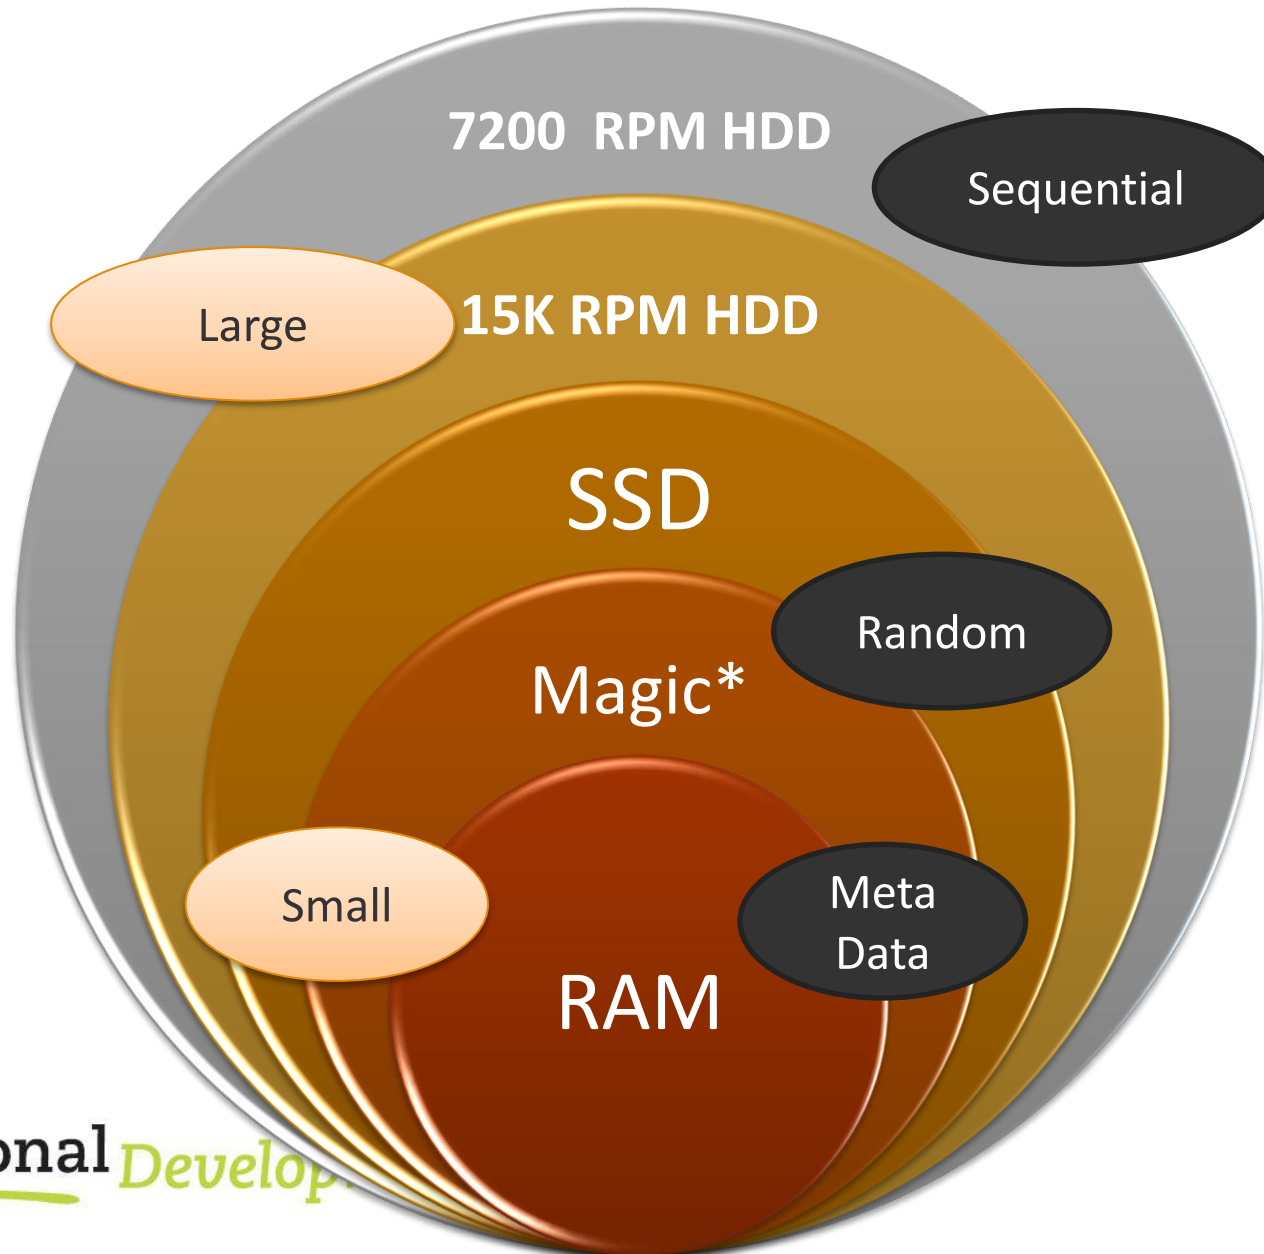
Cost containment:

- Introducing new technologies to enable the deployment of High Performance Data Sets on low cost Data Disks. Replace FC 15K with SATA 7200 rpm by introducing new technologies (SSD/Performance Acceleration Hardware/Magic*)



*By Magic we are referring to various performance acceleration solutions such as caching mechanisms, tiered file systems, data virtualization, read and write optimization hardware and methods.

Where does your data fit?



Performance Troubleshooting Triggers

- Customer calls
- Triggered by monitoring system
- Day to day analysis of system performance data

Performance Troubleshooting

- How many times do you get “Your storage is slow” phone calls?
- How many times have you placed “What are you doing to my storage?” calls?
- What is the usual success rate for finding root cause for performance spikes or service outages?



Diagnostic Tools

What's needed for Storage Engineers:

- Ability to see inside storage set to analyze data workload
- Capability to identify issues both inside the NAS and outside the of the Storage tier in the applications space

Oracle Solaris Dtrace performance analytics


- What is Dtrace?
 - Your answer to diagnosing performance and abuse problems, understanding your data load and resource utilization as well as getting to the bottom of your customer application issues.
- Who needs it?
 - Anyone who wants an insight into storage performance data.
- Where is it?
 - This analytics utility is provided as part of Oracle SUN ZFS Storage Appliance.

Example of analytics charts


Worksheets ▶ Key Performance statistics New Save Clone Close

➕ Add statistic...


Cache: ARC accesses per second broken down by hit/miss



Cache: L2ARC accesses per second broken down by hit/miss



Network: device bytes per second broken down by device




Analytics graphs

➕ Add statistic...

⊕ Protocol: NFSv3 operations per second




At 13:49:40:
786 ops per second



⊕ Protocol: NFSv3 operations per second broken down by client



At 13:49:40:
5869 m1ploftweb00
3452 m1ploftweb00
2724 m1ploftweb00
2536 m1ploftweb00
1142 m1ploftweb00
966 m1ploftweb00
608 m1ploftweb00
567 m1ploftweb00
17906 ops per second




⊕ Protocol: NFSv3 operations per second broken down by latency



At 13:49:40:
10 80.0 ms
20 70.0 ms
26 60.0 ms
66 50.0 ms
135 40.0 ms
300 30.0 ms
789 20.0 ms
1308 10.0 ms
15240 0 us
17912 ops per second

Drilldown on operations exceeding 20.0 ms:

- By type of operation
- By client
- By file name
- By share
- By project
- By size
- By offset



Scenario #1 for Dtrace use:



Scenario #1 continued


Protocol: NFSv3 operations per second broken down by type of operation



At 07:58:17:

17765	getattr
162	access
82	lookup
52	read
18	create
16	remove
14	write

18126 ops per second




Protocol: NFSv3 operations per second of type getattr broken down by project



At 07:58:17:

53285	off_data01
59	off_data02
57	off_data03

53401 ops per second



Protocol: NFSv3 operations per second of type getattr for project 'off_data01' broken down by file name





At 07:58:17:

29668	/export/data01/b
	/PA/Conferences
	/2009/Already in
23428	/export/data01/b
	/PA/Product/Pres
	client conferenc
70	<unknown>
9	/export/data01/b
9	/export/data01/b
5	/export/data01/b
	/PA/Product/Pres
5	/export/data01/b
	/PA/Conferences
	/PeopleAdmin/20
5	/export/data01/b
	/PA/Conferences
5	/export/data01/b
	/PA/Conferences

Show hierarchy

53284 ops per second



Own your data

Q&A

For more information or questions:
Julia@godaddy.com