Flash back to reality – Myths and Realities
SSD Industry trends perspectives and tips

Presented by Greg Schulz, Founder & Sr. Advisory Analyst
The Server and StorageIO Group (StorageIO)
Author: Cloud and Virtual Data Storage Networking (CRC Press)
Greg@storageio.com | StorageIOblog.com | Facebook.com/StorageIO | @storageio

Introduction

Who is Greg Schulz, contact and other information

Has been IT Customer

Application systems development
Systems programming/management
Performance and Capacity Planning
Data Protection/Backup/BC/DR
Electric Power, Financial, Transportation

Has been Vendor

Storage, Network, SSD, Disk & Tape
Backup/BC/DR, RAID, Replication
NAS, SAN, LAN, MAN and WAN
Hardware, Software & Services
Sales Engineering, Tech Marketing

MUST-READ I.T. BLOG Biztech

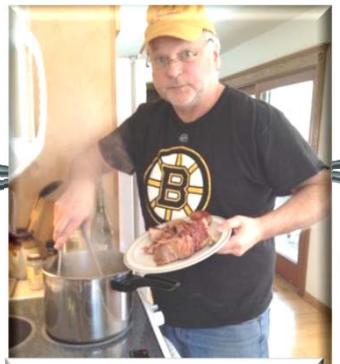

vmware vEXPERT
Five Time

Industry Analyst/Advisor

Cloud, virtualization/VDI, servers, HW, SW, servers, software defined, services, archive, backup/BC/DR, performance/capacity planning
Five time VMware vExpert
StorageIO.com

Author and Consultant

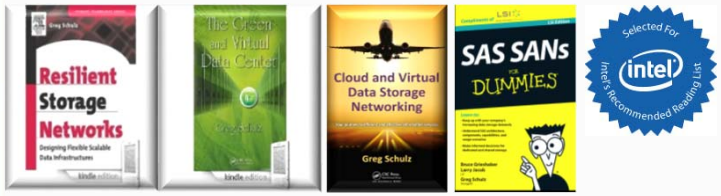

Syndicated columnist & blogger
StorageIOblog.com & StorageIO.TV

www.storageio.com/downloads

Twitter @storageio | Facebook.com/storageio | Storageioblog.com | storageio.com/newsletter

StorageIO

Industry Trends: SSD Walking the Talk

My experiences with SSD, spanning a "*few*" decades ;)

- ✓Launch customer for DEC ESE20 ram based SSD (late 80s, early 90s)
- ✓AS a vendor sold various SSD solutions across various industry's
- ✓As a vendor also partnered with SSD providers to provide connectivity
  - Now I cover them from an analyst/advisor/consult basis
  - This means using them, researching them and other things
  - I have a mix of flash, RAM SSD in various packages
  - Enterprise and consumer class devices
  - From laptop to servers, not to mention phones, USB thumb drives
  - These get used in physical, virtual and I even use cloud SSDs
  - Workloads: database (little data) to Hadoop (big data) to exchange & others
  - Not to mention vdbench, iorate and others even when needed iometer ;)

Performance
Tool Box
Hardware
Software Tools
And
Techniques

Industry Trends: SSD yesterday & today

For some DejaVu, for others revolutionary, or "technolutionary"

Long list of past (DEC,HP, Imperial, Memorex, Quantum, Solid Data and others), current and emerging vendors. Some will survive on their own, some will be acquired, some will end up on the future "Where are they now list"

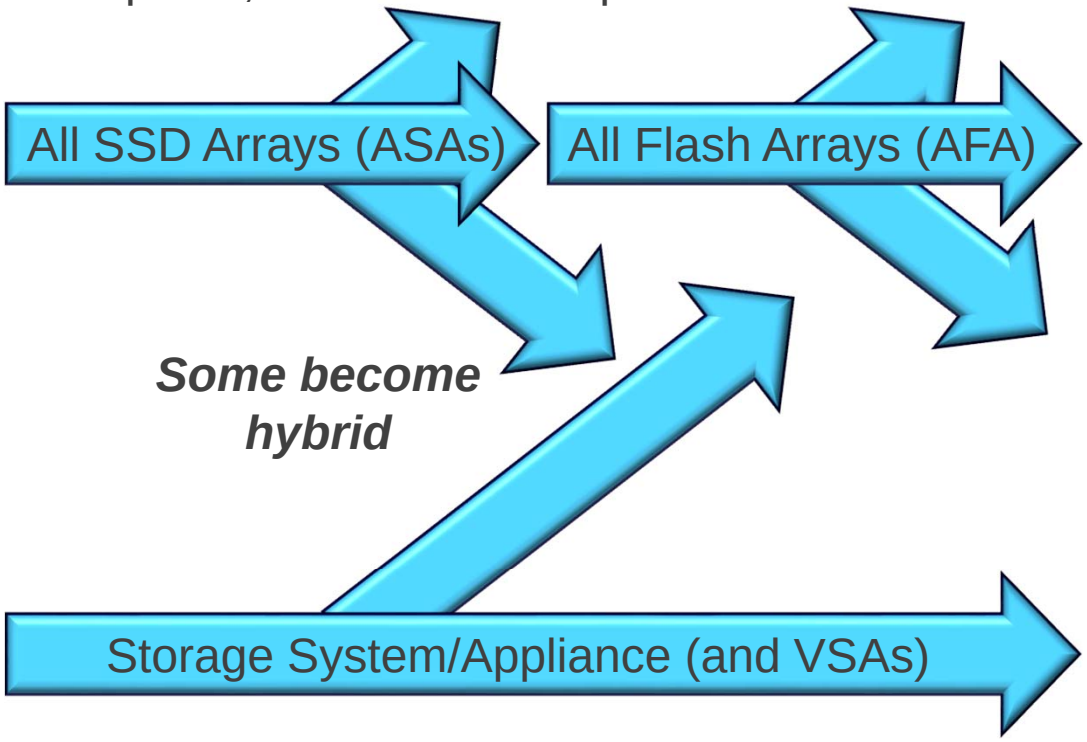

Storage Arrays/Appliances
Multi-Protocol
Multi-types of media
Mix of cache and no cache

Reality (someday in the future)

Future Hybrid may be mix of:

- DRAM, flash and their successors
- Perhaps even some legacy magnetic
- Some using in old ways
- Some using in new ways

Reality (Today) = Hybrid Home Run

Hybrid can be mix of:

- Legacy storage and AFA
- HDD and SSD in system
- HHDD and SSHD
- Local and cloud storage
- Block, file and object
- DRAM, flash and magnetic
- New and old items
- How you use it all…

Industry Trends: Taking a step back

What's your objective? Solve problem, find problem to solve?

Step or pull back so you have a better view of the big picture

What are you trying to accomplish, address, enable or verify? Is what you are doing relevant to your needs or simply what somebody else does, says or recommends?

High capacity, low cost bulk? Small low-latency high-priority?

All or nothing vs. hybrid (e.g. mix of what's best for your needs)

Companion Material

Industry Trends: Context - IOPs per device?

How many IOPs can a device do? What's the context?

Device	Vendor	Make	Model	Form Factor	Capacity	Interface	RPM Speed	Raw Test Result	FileSys Test Result
HDD	HGST	Desktop	HK250-160	2.5	160GB	SATA	5.4K	Here	Soon
HDD	Fujitsu	Desktop	MHWZ160BH	2.5	160GB	SATA	7.2K	Here	Soon
HDD	WD/Dell	Enterprise	WD1003FBYX	3.5	1TB	SATA	7.2K	Here	Soon
HDD	Seagate	Momentus	ST9160823AS	2.5	160GB	SATA	7.2K	Here	Soon
HDD	Seagate	MomentusXT	ST95005620AS	2.5	500GB	SATA	7.2K(1)	Here	Soon
HDD	Seagate	Savio 10K.3	ST9300603SS	2.5	300GB	SAS	10K	Here	Soon
HDD	Seagate	Savio 15K.2	ST9146852SS	2.5	146GB	SAS	15K	Here	Soon
HDD	Seagate	Barracuda	ST3000DM01	3.5	3TB	SATA	7.2K	Here	Soon
HDD	Seagate	Barracuda	ST3500320AS	3.5	500GB	SATA	7.2K	Here	Soon
HDD	Seagate	Cheetah	ST3146855SS	3.5	146GB	SAS	15K	Here	Soon
SSHD	Seagate	Ent. Turbo	ST600MX0004	2.5	600GB	SAS	SSHD	Here	Soon
HDD	Seagate	Ent. 15K	ST600MP0003	2.5	600GB	SAS	15K	Here	Soon
SSD	Seagate	600 SSD	ST120FP0021	2.5	120GB	SATA	SSD	Soon	Soon
SSD	Seagate	1200 SSD	ST400FM0073	2.5	400GB	12GSAS	SSD	Soon	Soon
SSD	Samsung	840 PRo	MZ-7PD256	2.5	256GB	SATA	SSD	Soon	Soon
HDD	Seagate	Desktop	ST4000DM000	3.5	4TB	SATA	HDD	Here	Soon

Performance characteristics 1 worker (thread count) for RAW IO (non-file system)

	Workload Pattern of test	Avg. Resp (R+W) ms	Avg. IOP Sec (R+W)	Bandwidth KB Sec (R+W)
16KB	100% Seq 100% Read	0.1	7,658	122,537
16KB	60% Seq 100% Read	4.7	210	3,370
16KB	30% Seq 100% Read	7.7	130	2,080
16KB	0% Seq 100% Read	10.1	98	1,580
16KB	100% Seq 60% Read	3.5	282	4,522
16KB	60% Seq 60% Read	7.7	130	2,090
16KB	30% Seq 60% Read	9.3	107	1,715
16KB	0% Seq 60% Read	11.1	90	1,443
16KB	100% Seq 30% Read	6.0	165	2,644
16KB	60% Seq 30% Read	9.2	109	1,745
16KB	30% Seq 30% Read	11.0	90	1,450
16KB	0% Seq 30% Read	11.7	85	1,364
16KB	100% Seq 0% Read	8.5	117	1,874
16KB	60% Seq 0% Read	10.9	92	1,472
16KB	30% Seq 0% Read	11.8	84	1,353
16KB	0% Seq 0% Read	12.2	81	1,310

Dell/Western Digital (WD) 1TB 7200 RPM SATA HDD (Raw IO) thread count 1 16K IO size

http://storageioblog.com/part-ii-iops-hdd-hhdd-ssd/

Industry Trends: Where's the Problem?

Lead with the solution (or cure) before ailment is known?

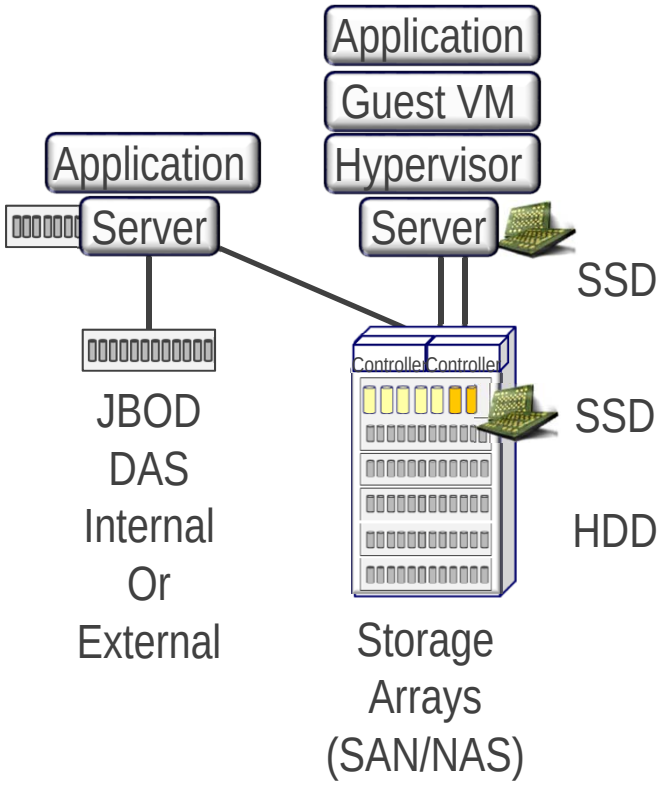

Where are the bottlenecks/problems?

- *Applications, Database, File systems*
- *OS, Hypervisor, Drivers, Configuration*
- *CPU, Memory, PCIe, Adapters*
- *Connectivity, Enclosures, Controllers*

Where to find, then fix problem…

- *Insight, awareness, baseline*

Vs.

Move and mask the problem
Or have a solution looking for problem ;)

Industry Trends: Server and Storage I/O

Some fundamentals and common sense around flash & SSD

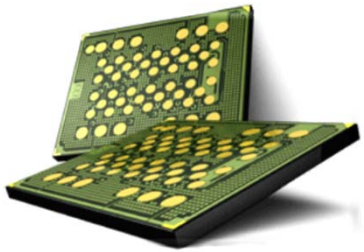

- ✓This is not an all or nothing, one size fits all value proposition
- ✓Memory is storage and storage is persistent memory
  - The best I/O is the one you don't have to do
  - The second best I/O is the one with least overhead
  - The importance of locality of reference
- ✓Cache and SSD is like real-estate, location matters and impacts cost
  - A little bit of cache (RAM, flash, etc.) in the right place goes a long way
  - A lot of cache (RAM, flash, etc.) should have a benefit yet cost cash
- ✓Just because something is new doesn't mean its better or faster
- ✓Fast applications need fast servers (and software), drivers, adapters, I/O paths, storage systems and devices

Can we get a side of context with them IOPS and other storage metrics?
http://storageioblog.com/side-context-iops/

Companion Material

Industry Trends: Locality of Reference

Memory is storage, storage is persistent memory

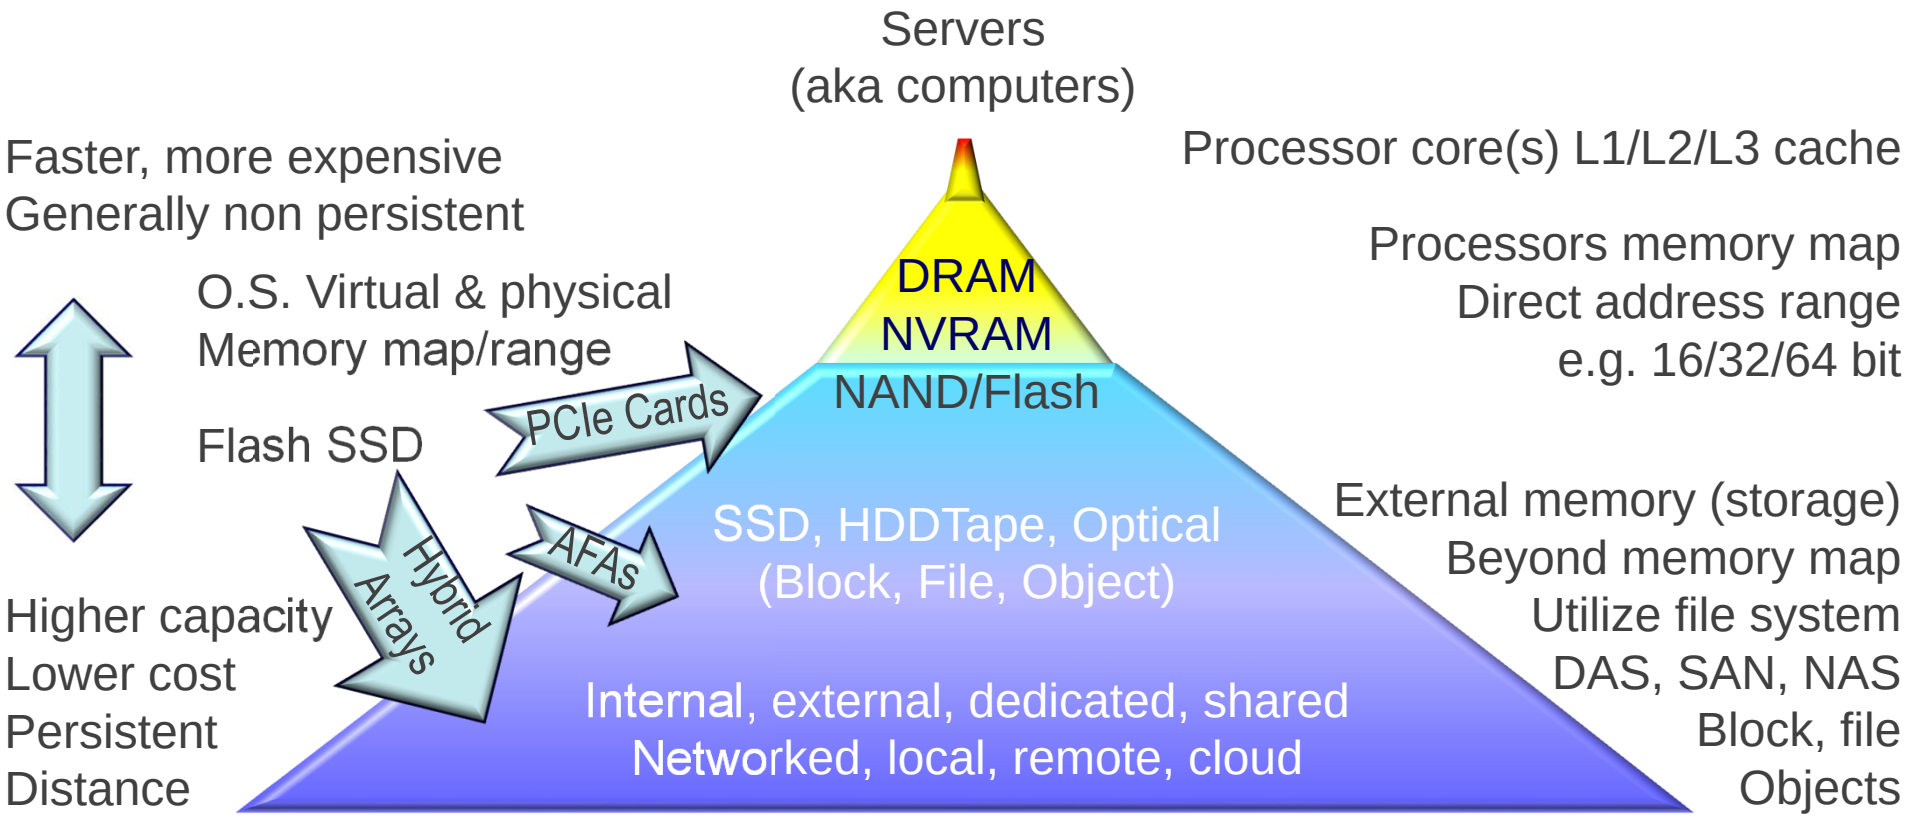

Can we get a side of context with them IOPS and other storage metrics
http://storageioblog.com/side-context-iops

Companion Material

Industry Trends: SSD and flash myths

Some common SSD / flash related myths and realities

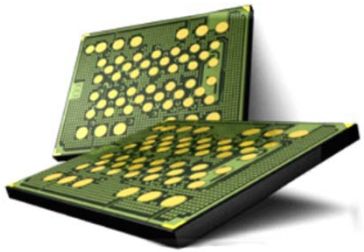

- Myth: SSD and flash are new
  - DRAM SSD has been around for decades
  - Flash is now over 20 years old
- Myth: Only All Flash Arrays (AFAs) can deliver performance
  - Not necessarily true, particular if not a good implementation
- Myth: Only "new" AFAs from startups can deliver performance
  - Care to guess who likes to keep this myth going? ;)
  - Some existing legacy systems will not benefit from flash or SSD
  - Some existing legacy systems greatly benefit from flash or SSD
  - Some systems have been back-end device starved (e.g. not controller)
- Myth: SSD too is expensive
  - This is true if compared on cost per space capacity
  - This is false if you compare on cost per work done (IOP, transaction, etc)

www.storageio.com/ssd

Companion Material

Industry Trends: SSD and flash myths

Some common SSD / flash related myths and realities

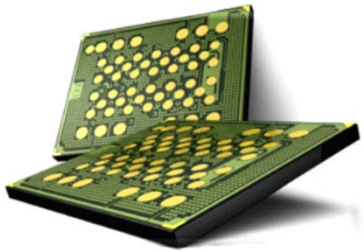

- Myth: Flash SSD does not consume power or generate heat
  - This is true when they are powered off
  - However they can run cooler than some HDDs
- Myth: Flash SSD does not break or wear out like HDD or tape
  - The flash cells do wear our from program/erase (p/e) cycles
  - Hence look at wear or durability, new metric such as TBytes/Written (TBW)
  - Also look for solutions that manage the wear, write gathering, optimization
  - Btw, also if benchmarking, “condition” them before use to avoid surprises
- Myth: Flash SSD wears out and will cause data loss
  - True however like other media, you can manage and plan accordingly
  - Don’t be scared of them or of some of the fud such as power loss (data loss)
- Myth: Flash SSD needs to replace HDDs
  - True for some applications and some vendors

www.storageio.com/ssd

Companion Material

Industry Trends: SSD and flash myths

Some common SSD / flash related myths and realities

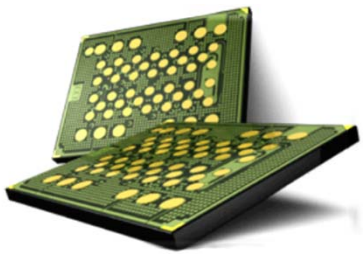

- Myth: Flash SSD will magically fix all your data center issues
  - With the right amount of budget, perhaps…
  - However watch out for moving problems or bottlenecks elsewhere
- Myth: Flash SSD is all about more IOPs
  - We need more context around IOPs, IOPs without context are BS
  - However more than IOPs, also bandwidth, latency, reads, writes etc.
  - Also, what's the system/server CPU and other impact when doing work
- Myth: Flash SSD eliminate need for performance and capacity planning
  - Hope you have a large bank account ;)
- Myth: Flash SSD cannot be used for backup/data protection
  - Why not, it's a great way to speed up snaps, copies, catalog and other things
  - However vendors and their pundits don't know how to tell the story
  - Its not in the play-book yet ;)

Can we get a side of context with them IOPS and other storage metrics
http://storageioblog.com/side-context-iops

Companion Material

Industry Trends: SSD and flash myths

Some common SSD / flash related myths and realities

- Myth: SSD in the cloud is magical or wont work
  - First if your application is in the cloud that SSD can be good
  - However, if your application is remote, keep response time in mind
  - Also, there are different types of cloud SSD, understand the type of service
  - For example there are AWS EBS SSD backed volumes and IOP limits
  - Then there are AWS EC2 high-io instances (servers with SSD)
  - What do you need, what are you trying to do, understand measurements
  - For example, what is the cost per capacity of the service
  - What are the IOP limits or constraints, also what size IOPs are used for billing
  - Will the IOPs be deterministic or variable, yes they can vary by service
  - This gets back to basic performance engineering and capacity planning

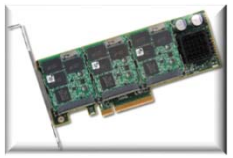

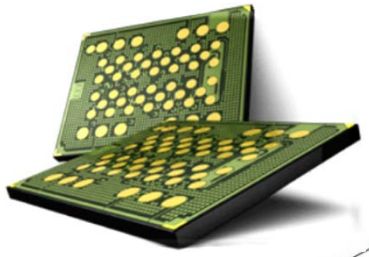

www.storageio.com/ssd

Industry Trends: Where's the Problem?

Lead with the solution (or cure) before ailment is known?

Re-iterate – What's the focus and problem so we can solve or address

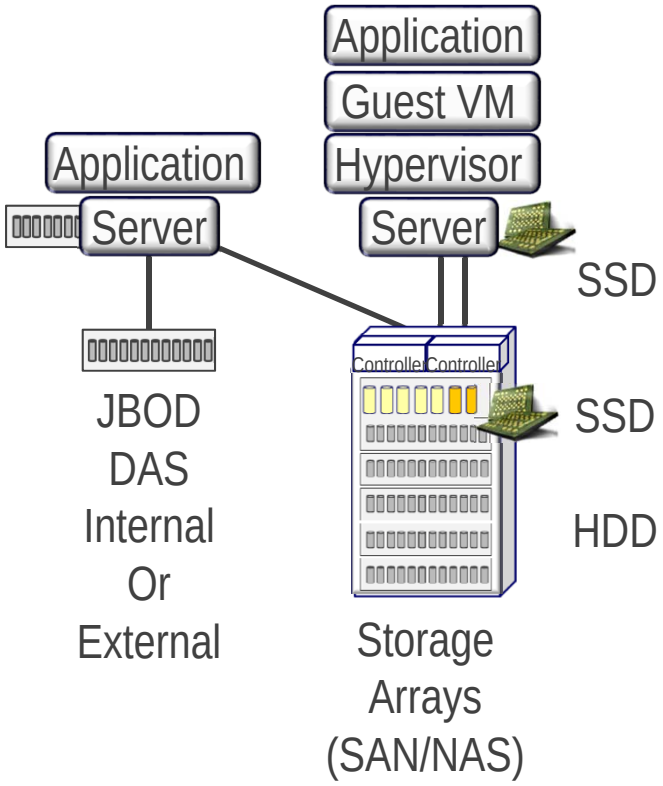

Where are the bottlenecks/problems?

- *Applications, Database, File systems*
- *OS, Hypervisor, Drivers, Configuration*
- *CPU, Memory, PCIe, Adapters*
- *Connectivity, Enclosures, Controllers*

Where to find, then fix problem…

- *Insight, awareness, baseline*

Vs.

Move and mask the problem
Or have a solution looking for problem ;)

Companion Material

Where to use SSD?

SSD as cache, target or hybrid, which is right or best for you?

Answer: It depends on your needs, budget and preferences!

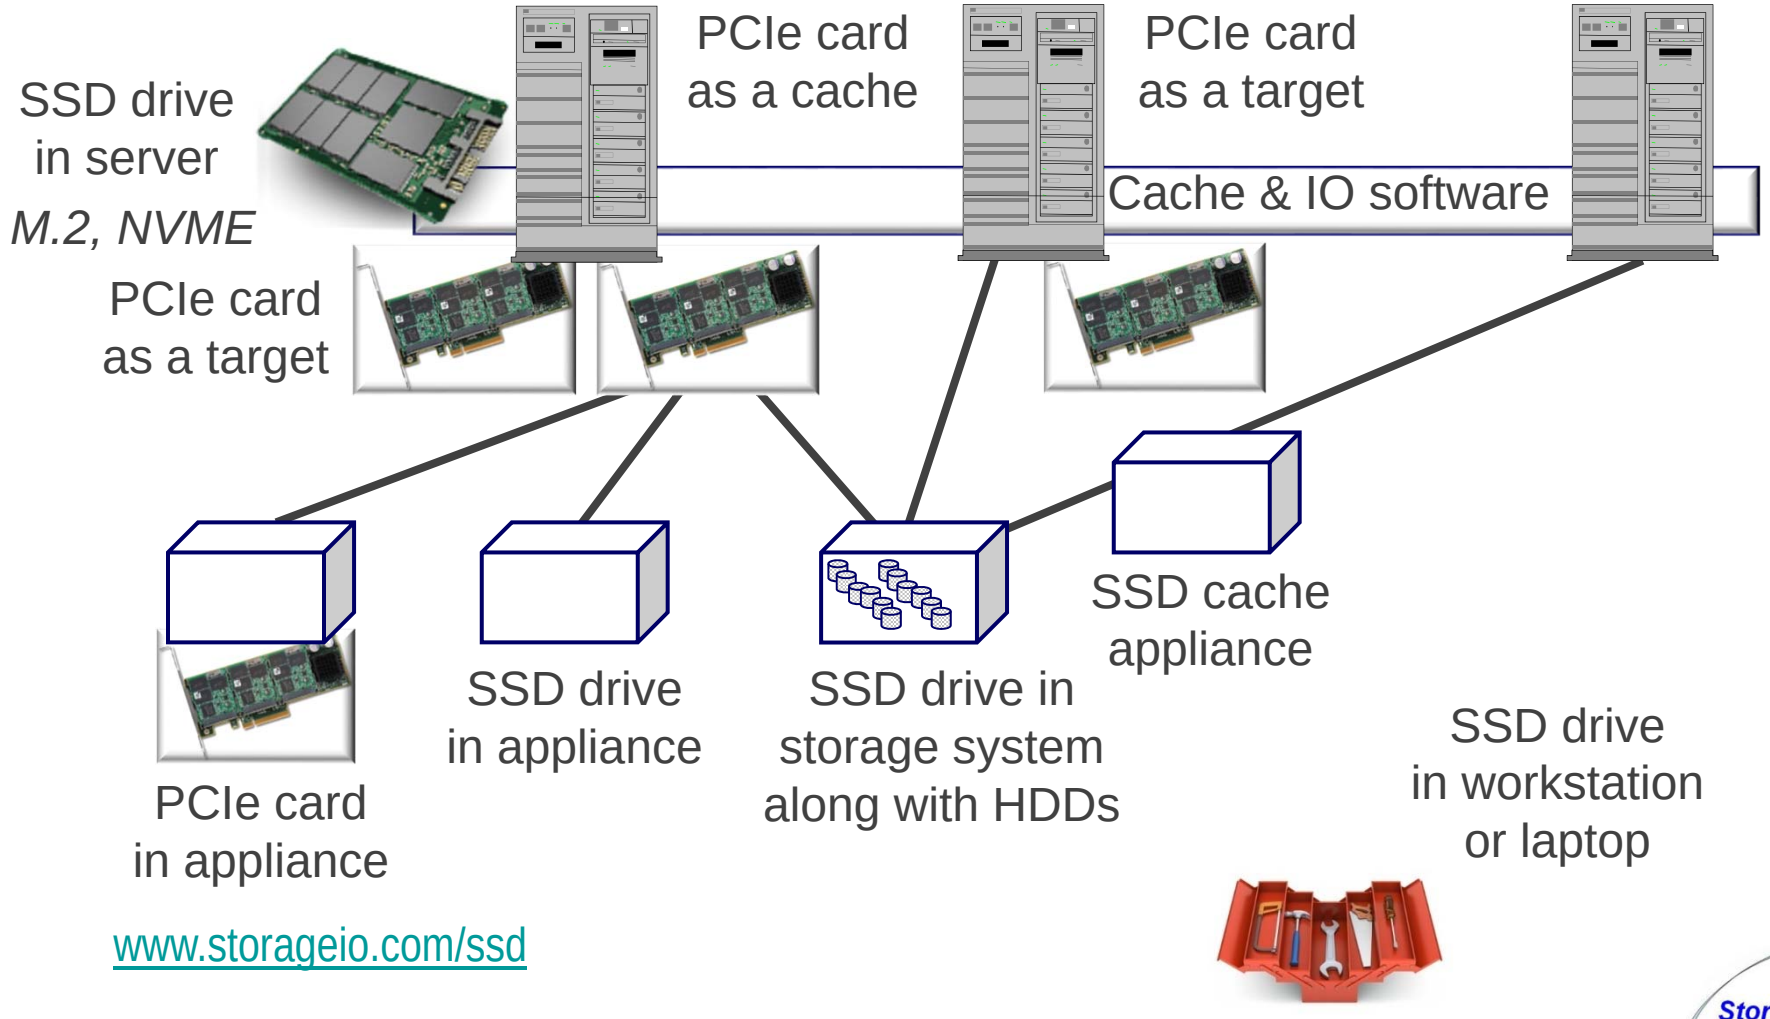

www.storageio.com/ssd

Industry Trends: Storage I/O tools

What's in your toolbox, what to use for different things?

- Insight and awareness
  - OS tools (e.g. Perfmon etc), 3rd party (Spotlight on Windows/*nix), etc.
  - VM tools (ESXTOP, Visual ESXTOP) and storage system based
  - Others including HiMon (e.g. from HyperIO)
  - Database tools (e.g. Microsoft SQL Server Studio and others)
- Workload generators (benchmarks etc.)
  - The best = Your application under real or applicable workloads!
  - Second best = Variation of your application under realistic workloads
  - Synthetic that span application software, server, storage and I/O paths
  - Traces or sub-system or component specific tools
  - Some examples include among others

    Benchmark Factory, Dedisbench, DFSIO (Hadoop), Fio, Hammer, Iometer, Iorate, Iozone, Jet, Login VSI, PCMark, SNIA Emerald (uses vdbench), SPEC, SQLIO, Tera (Teragen, Terasort, Teravalidate for Hadoop), TPC, Vdbench, VMmark (based on DVD Store) and many others.

http://storageioblog.com/iops-hdd-hhdd-ssd-vmware/

Companion Material

Focus: Guests OS

Focus area: Performance, availability, resource usage?

Tools Not In A Specific Order Below!

- Plugins
- I/O Cache Tools Accelerators
- Path Managers
- Drivers

- Business Applications
- Database / Key Value Repositories
- File system Volume Manager **Operating System**
- Hypervisors
- Hardware

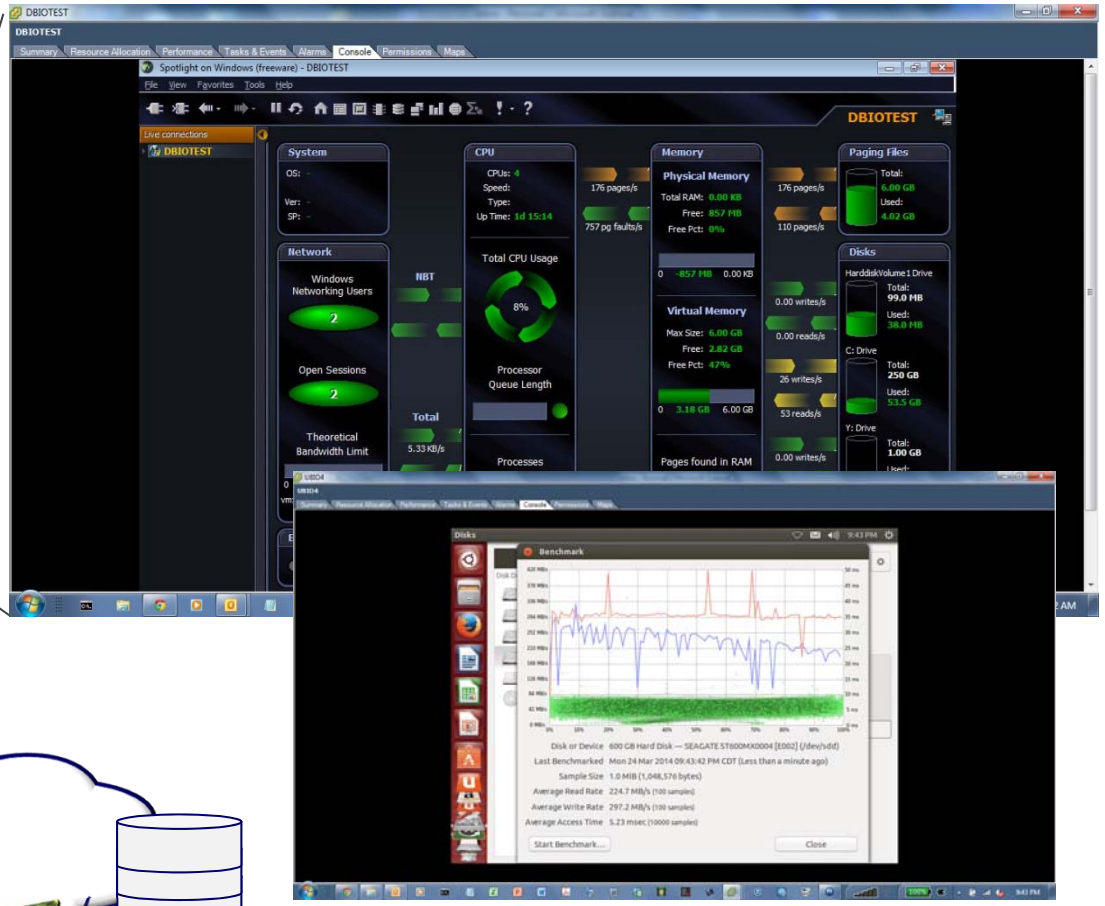

Companion Material

Focus: Hypervisors

Focus area: Applications, CPU, Memory, Storage I/O network?

Tools Not In A Specific Order Below!

- Plugins
- I/O Cache Tools Accelerators
- Path Managers
- Drivers

Business Applications

Database / Key Value Repositories

File system Volume Manager Operating System

Hypervisors

Hardware

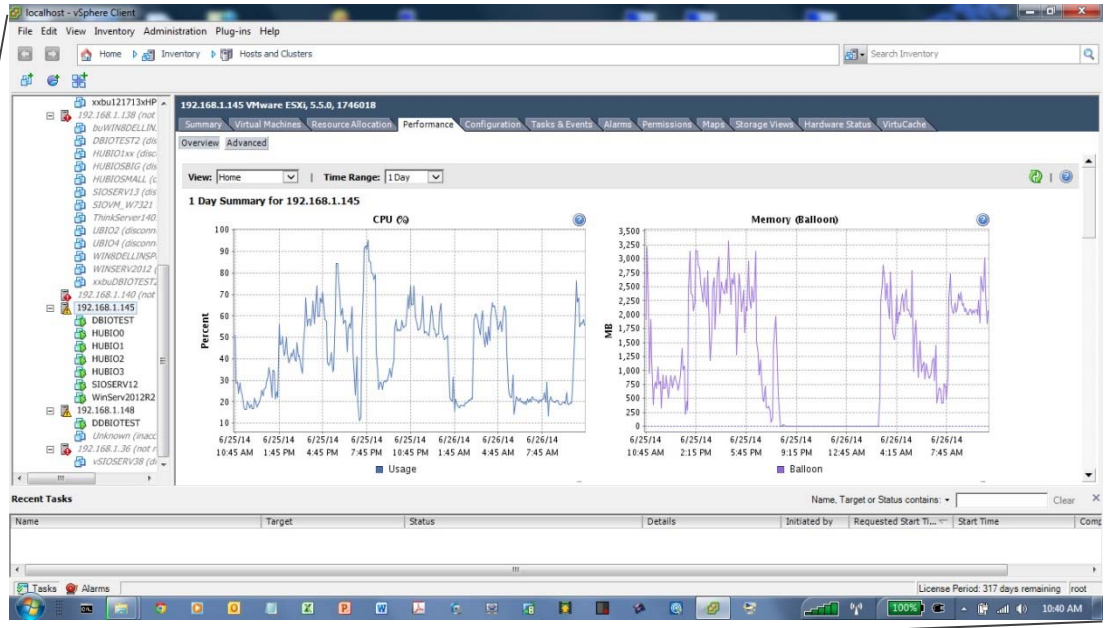

Some tools include Citrix, Dell, JAM, Login VSI, Microsoft, RV Tools, Solarwinds, Spacesniffer, Visual ESX top and VMware *among many others*

Companion Material

Focus: Hypervisors

Focus area: Applications, CPU, Memory, Storage I/O network?

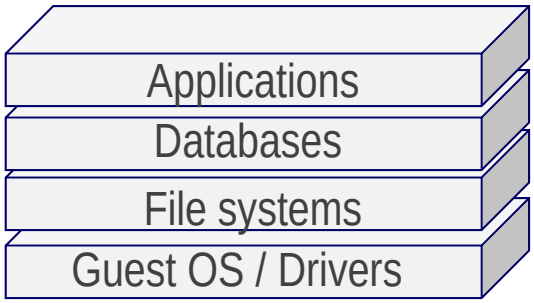

SAP, Exchange, other App centric & 3rd party

Dell (Toad), Oracle, Microsoft and other tools

OS, 3rd party, Treesize, Spacesniffer, etc

Microsoft, *nix and 3rd parties, SOW, SO*nix, etc

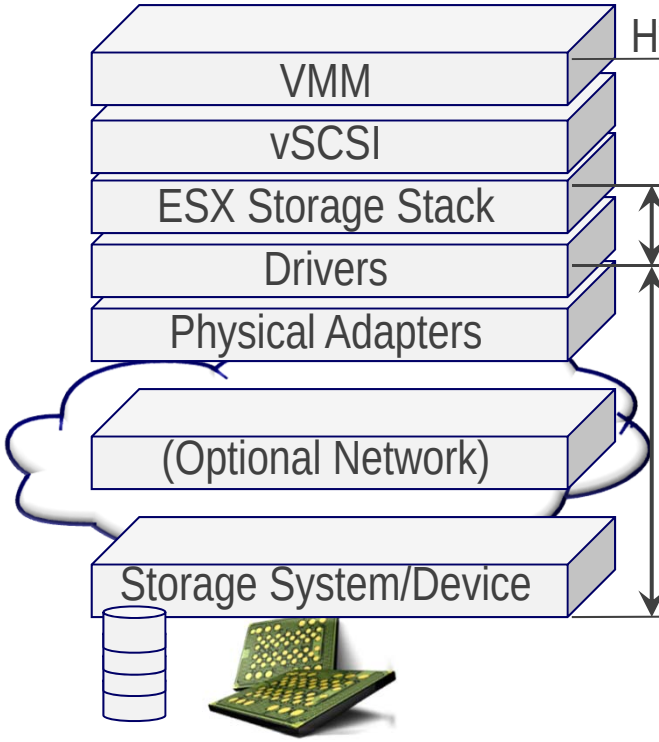

Visual ESXtop and others

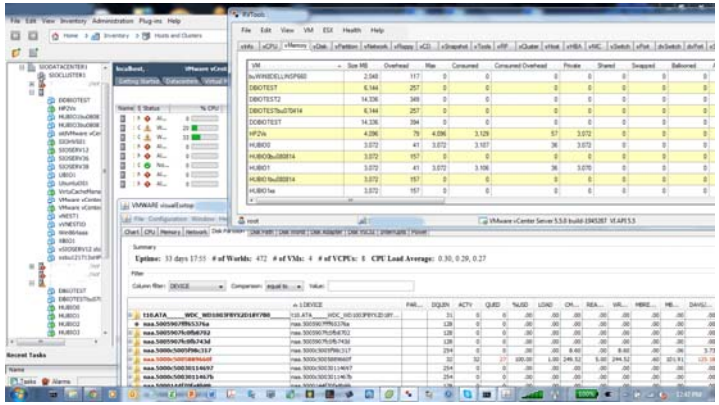

GAVG = Guest avg. resp. time

KAVG = Time in vmkernel

QAVG = Time in vStorage stack

DAVG = Device time (w/o OS or vmw)

Companion Material

Focus: Hypervisors and Cache

Focus area: Cache Tools and Technologies

Tools Not In A Specific Order Below!

- Business Applications
- Database / Key Value Repositories
- File system, Volume Manager, Operating System
- Hypervisors
- Hardware

- Plugins
- I/O Cache Tools Accelerators
- Path Managers
- Drivers

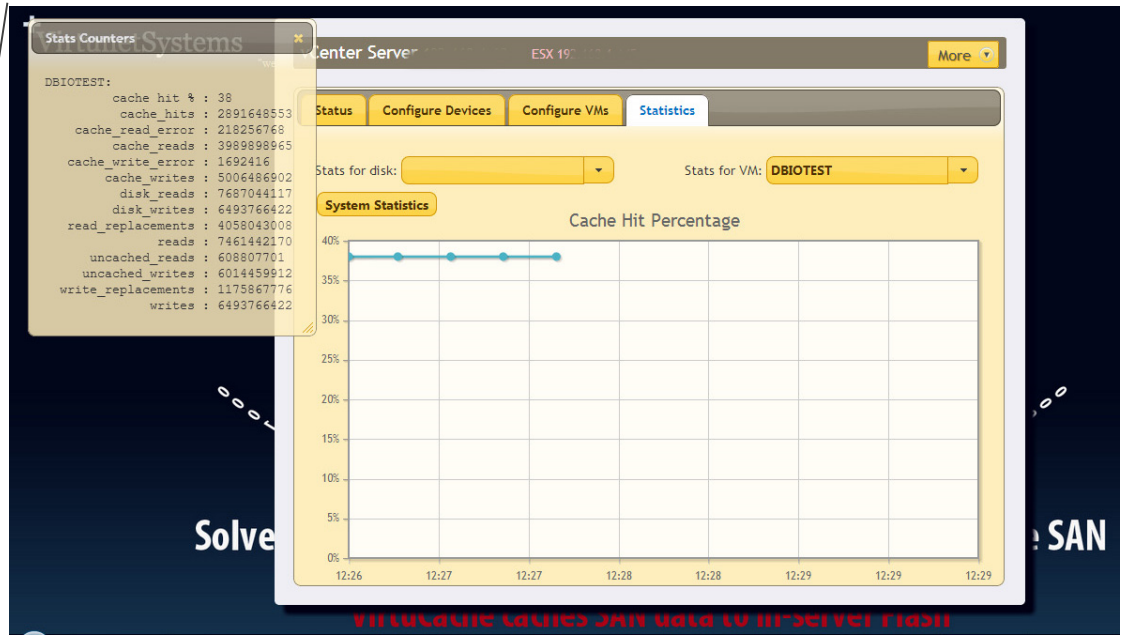

VirtunetSystems VirtuCache

Focus: Storage Performance

Focus area: Server and Storage IO performance

Tools Not In A Specific Order Below!

- Plugins
- I/O Cache Tools Accelerators
- Path Managers
- Drivers

- Business Applications
- Database / Key Value Repositories
- File system Volume Manager Operating System
- Hypervisors
- Hardware

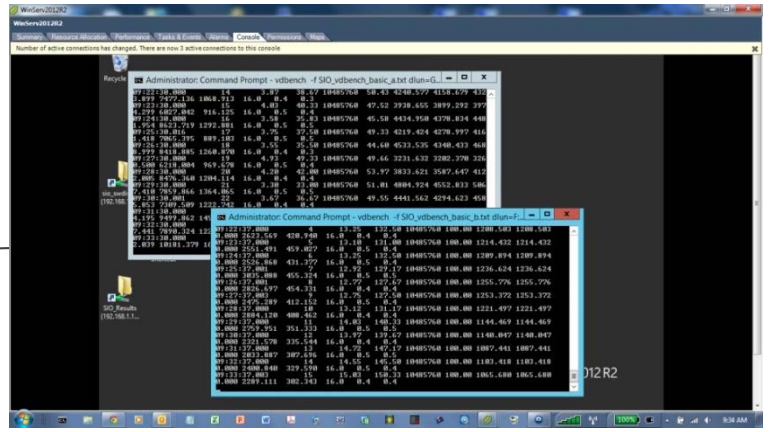

Part II: How many IOPS can a HDD, HHDD or SSD do with VMware?

http://storageioblog.com/part-ii-iops-hdd-hhdd-ssd/

Industry Trends: vdbench example

Server and Storage I/O metrics and context, more than just IOPs

```
15:23:06.971 Starting RD=write000nseqSIOV; I/O rate: Uncontrolled MAX; elapsed=120; For loops: xfersize=10m

Apr 30, 2014  interval        i/o   MB/sec   bytes   read     resp     read    write     resp     resp queue  cpu%  cpu%
                            rate  1024**2     i/o    pct     time     resp     resp      max   stddev depth sys+u   sys
15:23:36.078        1       37.17   371.67 10485760 100.00  404.306  404.306    0.000 2729.035  210.964  15.1  12.9   9.5
15:24:06.047        2       41.17   411.67 10485760 100.00  387.927  387.927    0.000  455.572   11.864  16.0   1.4   1.2
15:24:36.047        3       40.60   406.00 10485760 100.00  394.411  394.411    0.000  486.506   12.511  16.0   2.5   1.6
15:25:06.046        4       40.23   402.33 10485760 100.00  396.580  396.580    0.000  936.045   66.346  16.0  17.8  10.6
15:25:06.062   avg_2-4      40.67   406.67 10485760 100.00  392.938  392.938    0.000  936.045   39.547  16.0   7.2   4.5

15:25:08.010 Starting RD=write000seqSIOV; I/O rate: Uncontrolled MAX; elapsed=120; For loops: xfersize=4k

Apr 30, 2014  interval        i/o   MB/sec   bytes   read     resp     read    write     resp     resp queue  cpu%  cpu%
                            rate  1024**2     i/o    pct     time     resp     resp      max   stddev depth sys+u   sys
15:25:38.109        1   103335.53   403.65    4096 100.00    0.145    0.145    0.000   94.824    0.556  14.9  55.7  33.9
15:26:08.031        2   117086.83   457.37    4096 100.00    0.133    0.133    0.000   26.966    0.151  15.5  50.2  33.8
15:26:38.031        3   117405.70   458.62    4096 100.00    0.132    0.132    0.000   25.839    0.140  15.5  52.8  34.4
15:27:08.031        4   115595.13   451.54    4096 100.00    0.134    0.134    0.000  103.710    0.356  15.5  53.5  35.2
15:27:08.031   avg_2-4  116695.89   455.84    4096 100.00    0.133    0.133    0.000  103.710    0.237  15.5  52.2  34.5
15:27:08.031 *
15:27:08.031 * Warning: total amount of i/o per second per slave (116696) greater than 50000.
15:27:08.031 * You may need to adjust your total slave count (see 'jvms=' in documentation
15:27:08.031 * This run actively used 1 slaves.
15:27:08.031 *
```

Summary report

Vdbench "manifest" Of output Summary Detail, logs Histograms

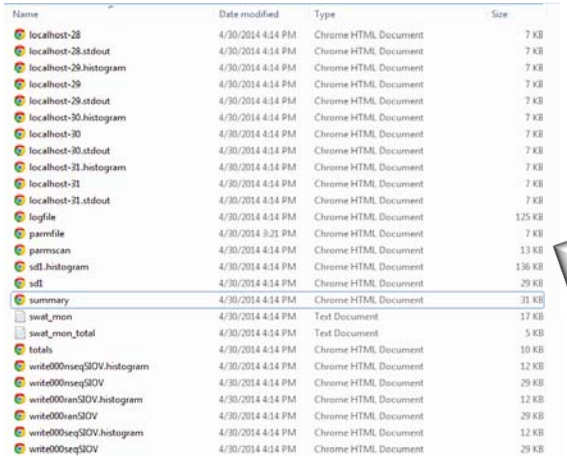

Companion Material

Industry Trends: vdbench jumpstart

Some quick start, hints and tips

Download vdbench from Oracle
http://www.oracle.com/technetwork/server-storage/vdbench-downloads-1901681.html

Create a folder of where to copy the download into

Code and examples for different OS and other items
Install Java 7u55 (or later) - http://java.com/en/download/chrome.jsp?locale=en

Add Java environment settings (if applicable, e.g. Windows)

JAVA_HOME
C:\Program Files (x86)\Java\jre7\bin;

Path
;C:\Program Files (x86)\Java\jre7\bin;

For Windows servers you may need reset performance counters
Via command (with admin rights) Lodctr /R

Read the vdbench documentation…
http://storageioblog.com/part-ii-iops-hdd-hhdd-ssd/

StorageIO

Companion Material

Industry Trends: vdbench example

Example script for exercising workload

Some examples:
SNIA Emerial on-line active disk “hot-band” workload

http://www.snia.org/emerald/download
http://snia.org/sites/default/files/SNIA_Emerald_Script_Version_2014_05_14.txt

Or another script example

* Sample command line
*
* vdbench -f SIO_vdbench_basic.txt devnum=1 devsize=5g devname=X1TMPFILE worktbd=4k,10m etime=120 itime=30 -
061814_BasicX1
*
* sd=sd1,lun=\\.\PHYSICALDRIVE!devnum,size=!devsize,offset=512,thread=16

sd=sd1,\\.\c:\temp\vdbenchJun182014.tmp,size=!devsize,offset=512,thread=16

*
* sd=sd2,lun=\\.\PHYSICALDRIVE2,size=370g,offset=512
*

http://storageioblog.com/part-ii-iops-hdd-hhdd-ssd/

Companion Material

Industry Trends: vdbench example

Example script for exercising workload

```
* Define the workloads:
*
wd=write100seqSIOV,sd=(sd1),seekpct=seq,rdpct=0
wd=write075seqSIOV,sd=(sd1),seekpct=seq,rdpct=25
wd=write050seqSIOV,sd=(sd1),seekpct=seq,rdpct=50
wd=write025seqSIOV,sd=(sd1),seekpct=seq,rdpct=75
wd=write010seqSIOV,sd=(sd1),seekpct=seq,rdpct=90
wd=write000seqSIOV,sd=(sd1),seekpct=seq,rdpct=100
*
wd=write100ranSIOV,sd=(sd1),seekpct=100,rdpct=0
wd=write075ranSIOV,sd=(sd1),seekpct=100,rdpct=25
wd=write050ranSIOV,sd=(sd1),seekpct=100,rdpct=50
wd=write025ranSIOV,sd=(sd1),seekpct=100,rdpct=75
wd=write010ranSIOV,sd=(sd1),seekpct=100,rdpct=90
wd=write000ranSIOV,sd=(sd1),seekpct=100,rdpct=100
*
wd=write050mixSIOV,sd=(sd1),seekpct=50,rdpct=50
*
```

http://storageioblog.com/part-ii-iops-hdd-hhdd-ssd/

StorageIO

Companion Material

Industry Trends: vdbench example

Example script for exercising workload

```
*
* Define the test steps
*
rd=write050mixSIOV,wd=write050mixSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
*
rd=write000seqSIOV,wd=write000seqSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write010seqSIOV,wd=write010seqSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write025seqSIOV,wd=write025seqSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write050seqSIOV,wd=write050seqSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write075seqSIOV,wd=write075seqSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write100seqSIOV,wd=write100seqSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
*
rd=write000ranSIOV,wd=write000ranSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write010ranSIOV,wd=write010ranSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write025ranSIOV,wd=write025ranSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write050ranSIOV,wd=write050ranSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write075ranSIOV,wd=write075ranSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
rd=write100ranSIOV,wd=write100ranSIOV,elapsed=!etime,interval=!itime,iorate=max,forxfersize=(!worktbd),openflags=directio
```

http://storageioblog.com/part-ii-iops-hdd-hhdd-ssd/

StorageIO

SSD Cache Testing Configuration

Virtual and physical server config

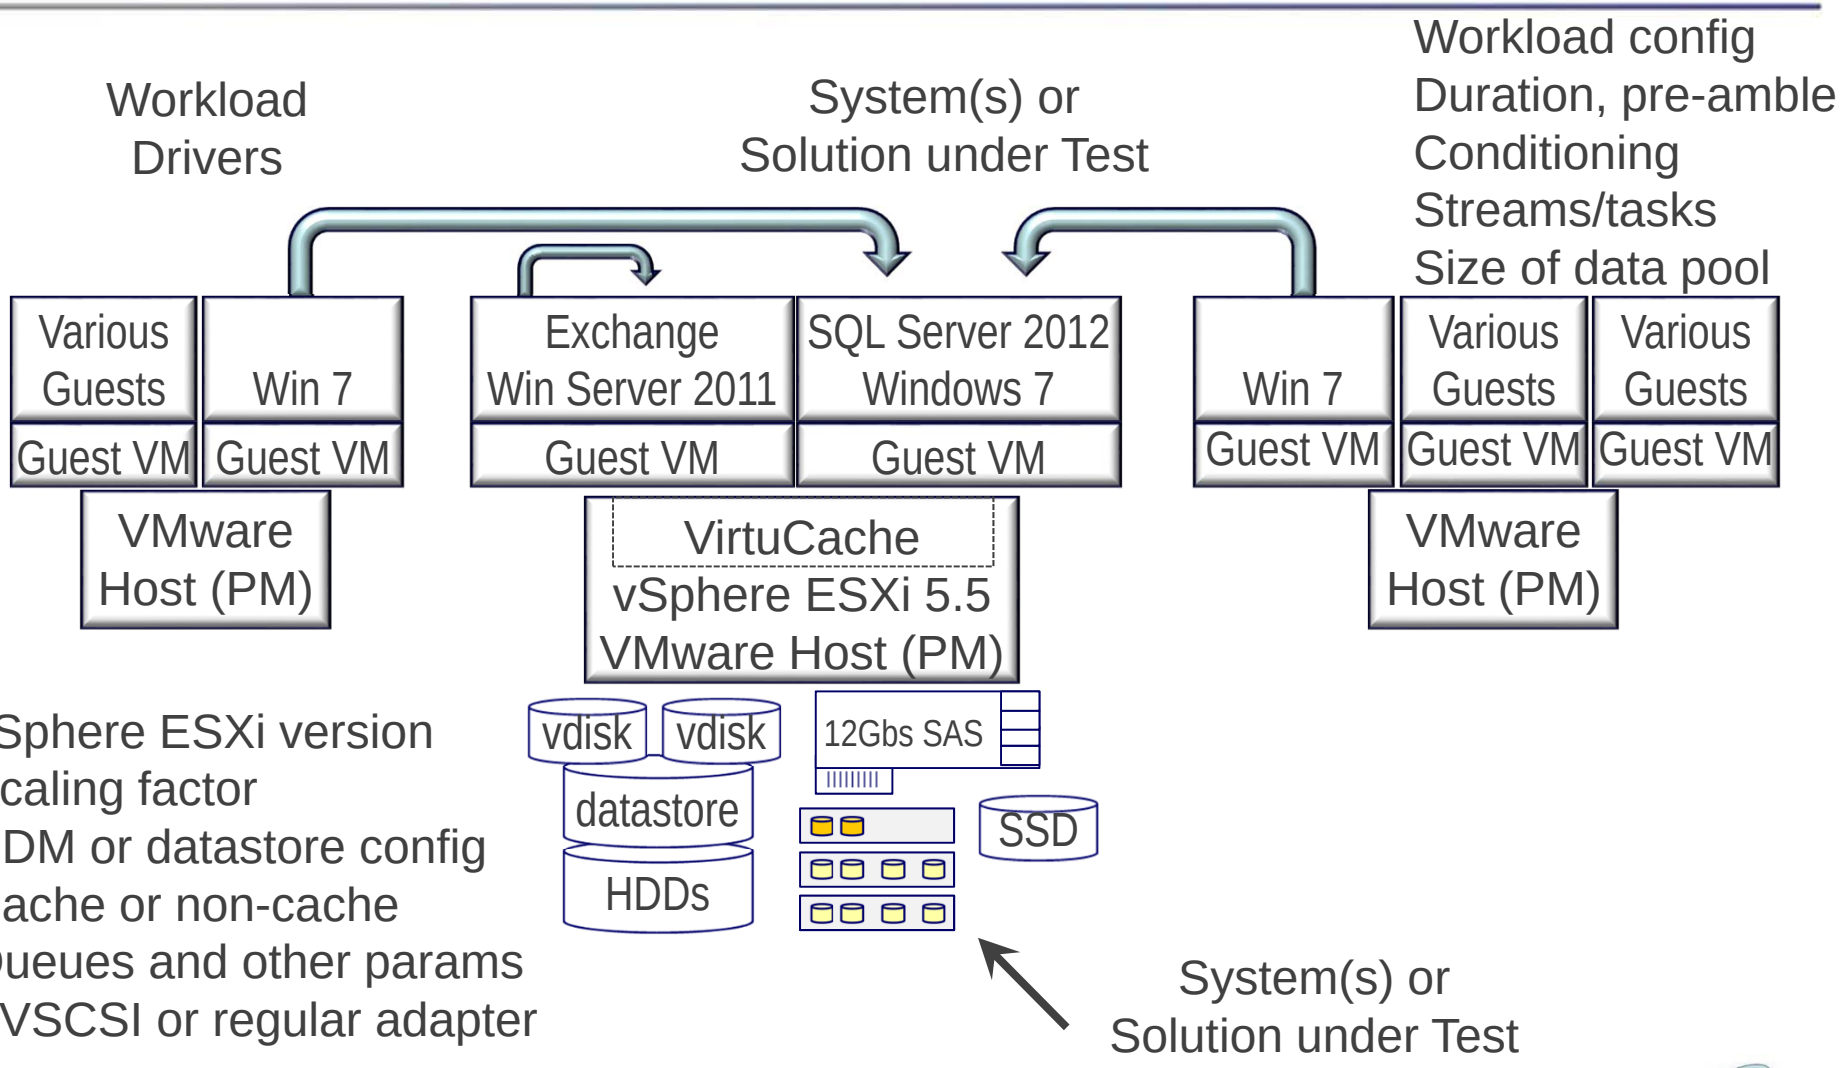

See more info and result examples at http://storageio.com/whitepaper.html

Companion Material

SSD Cache Testing Configuration

Virtual and physical server config

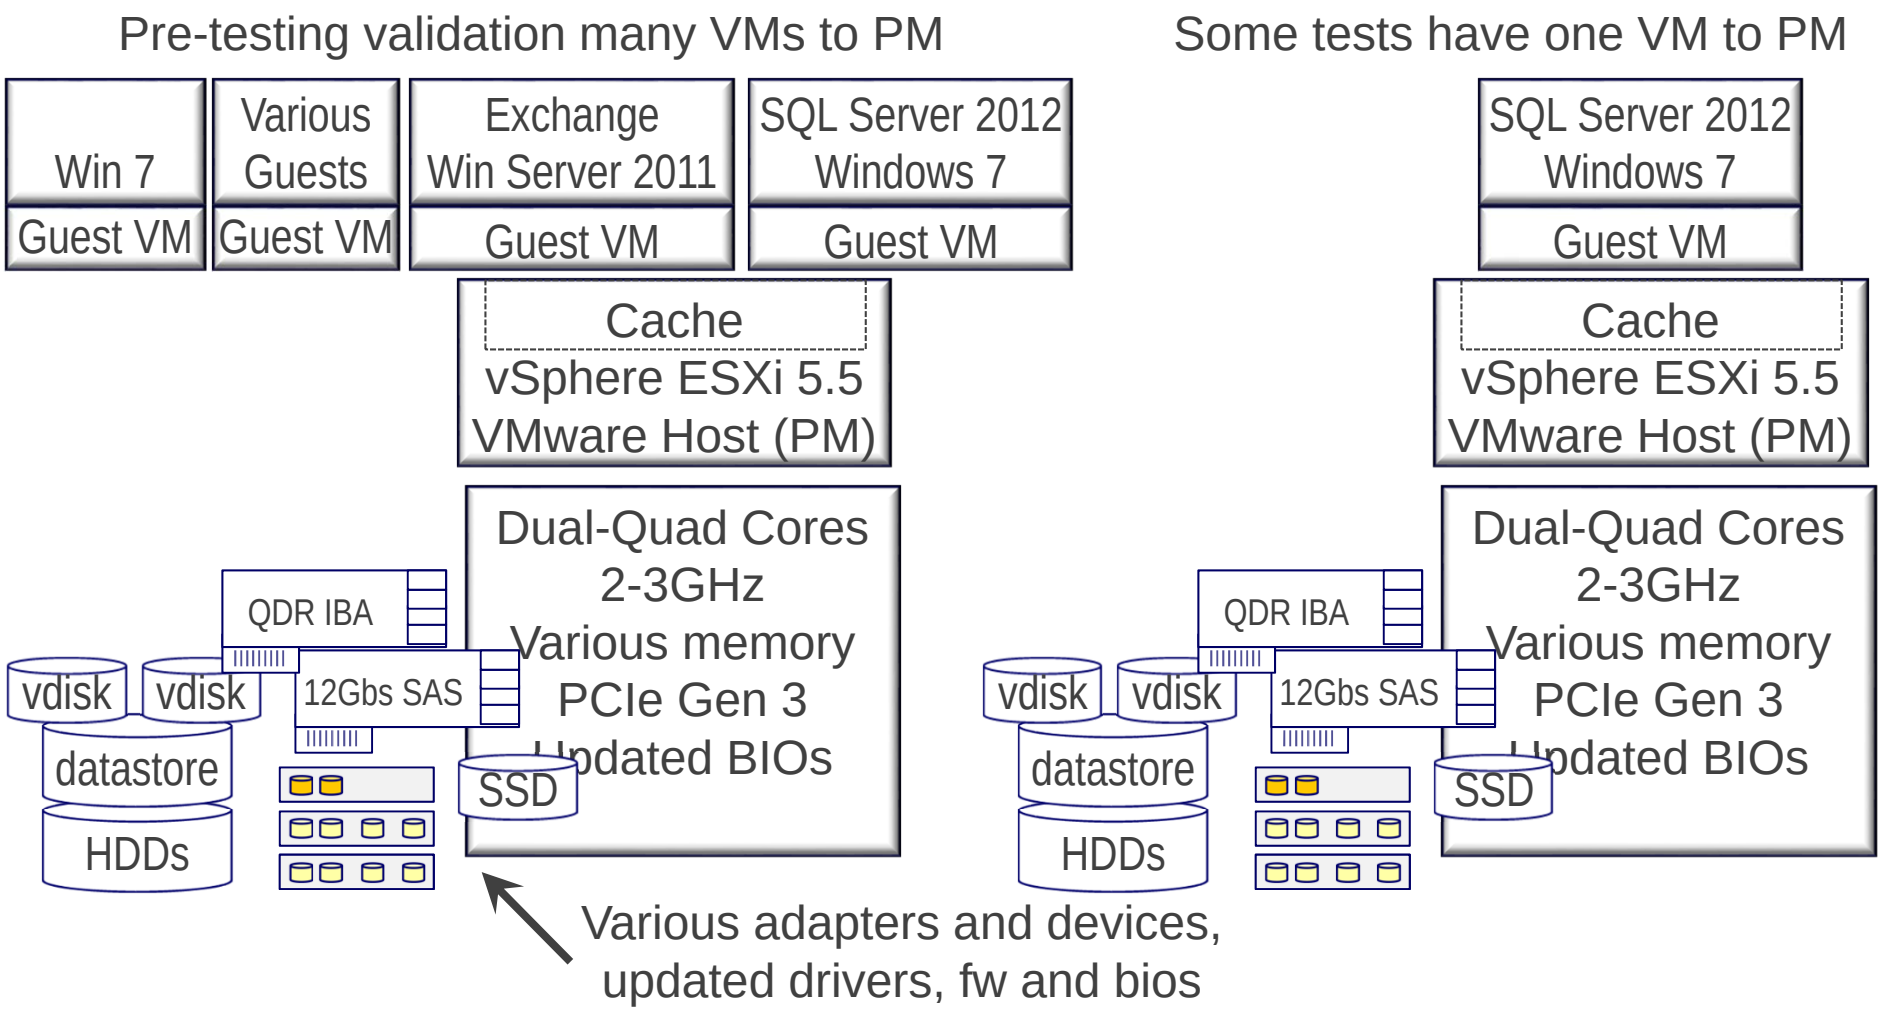

See more info and result examples at http://storageio.com/whitepaper.html

Industry Trends: Some proof points

Examples of what has been discussed here today

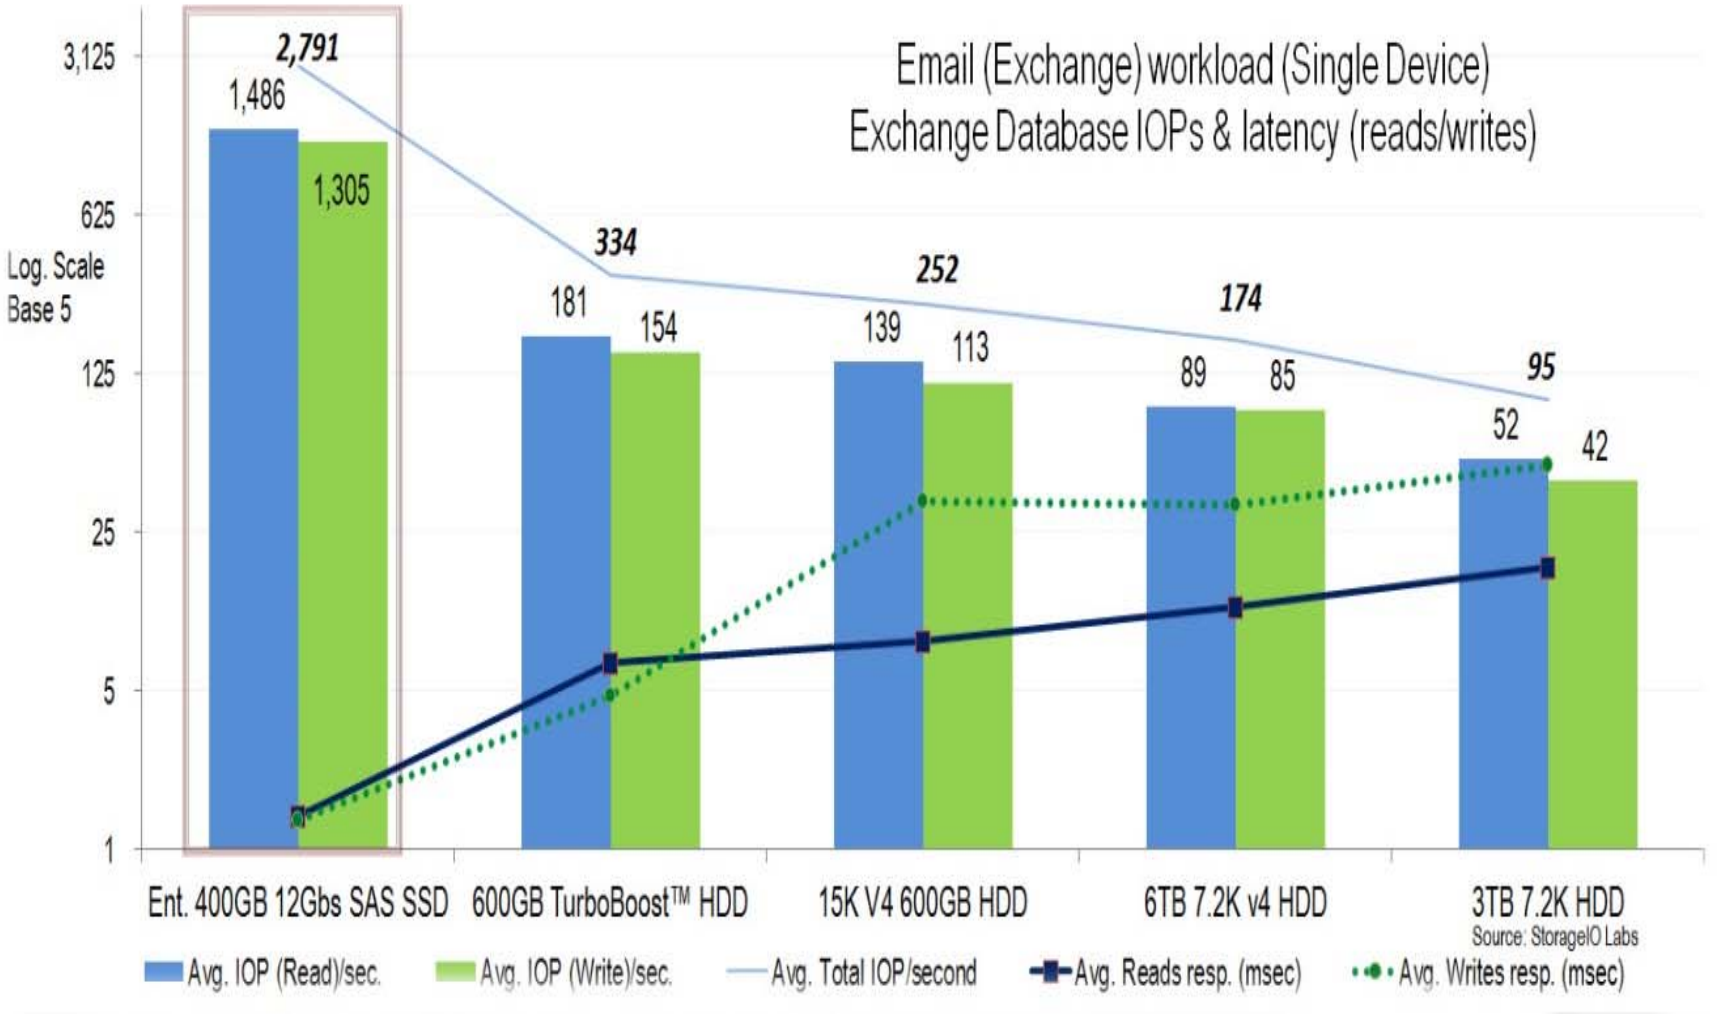

See more info and result examples at http://storageio.com/whitepaper.html

Industry Trends: Some proof points

Examples of what has been discussed here today

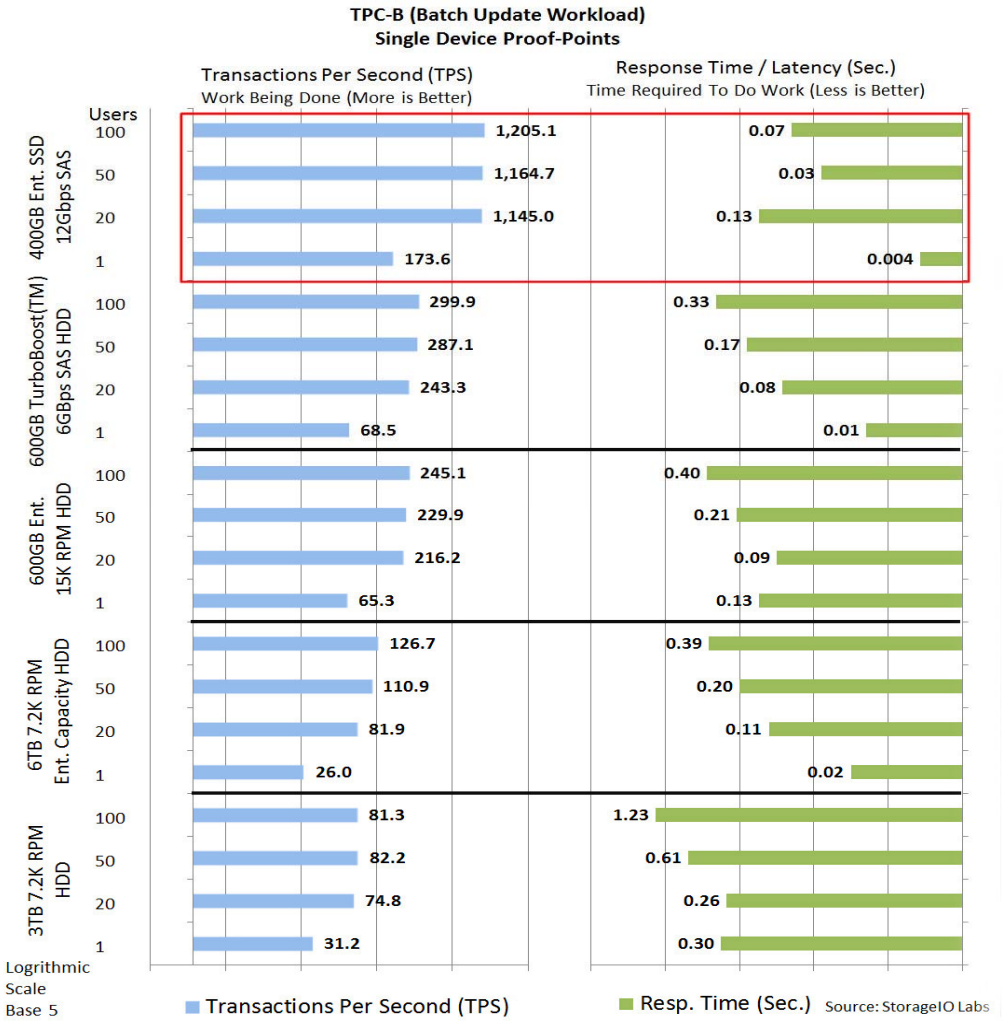

See more info and result examples at http://storageio.com/whitepaper.html

Industry Trends: Some proof points

Examples of what has been discussed here today

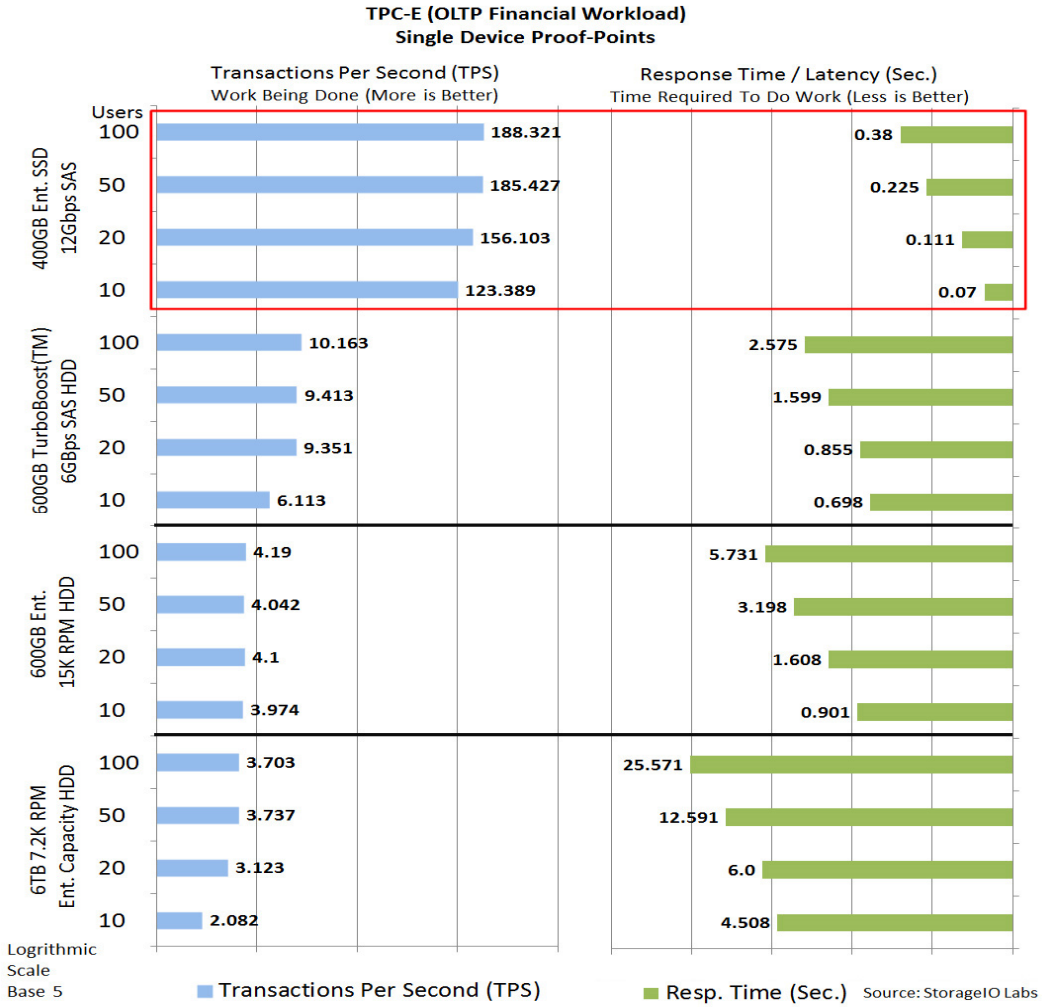

See more info and result examples at http://storageio.com/whitepaper.html

Companion Material

SSD Cache Testing Configuration

Configuration

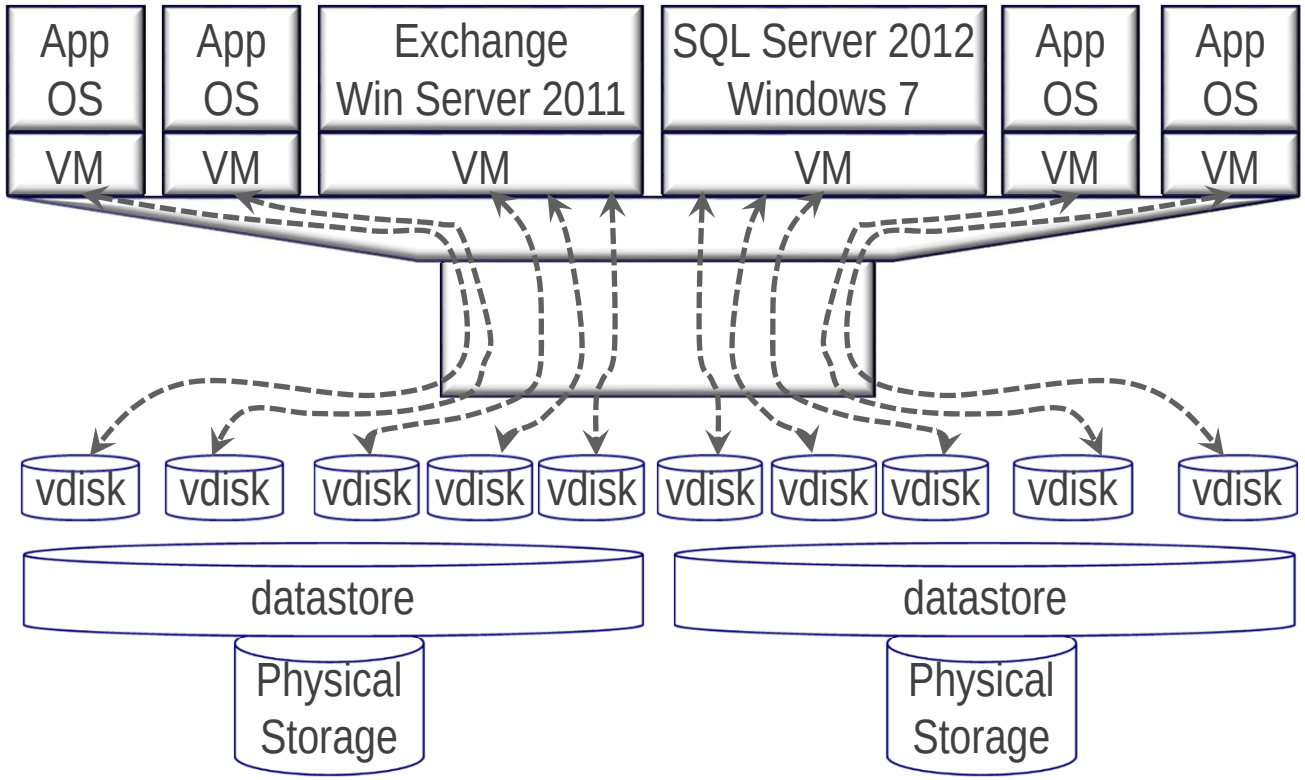

Applications
Operating System
Virtual Machines

Hypervisor
I/O Bottlenecks
Physical Machine

Virtual disks

Virtual datastores

Physical storage

See more info and result examples at http://storageio.com/whitepaper.html

SSD Cache Testing Configuration

Configuration

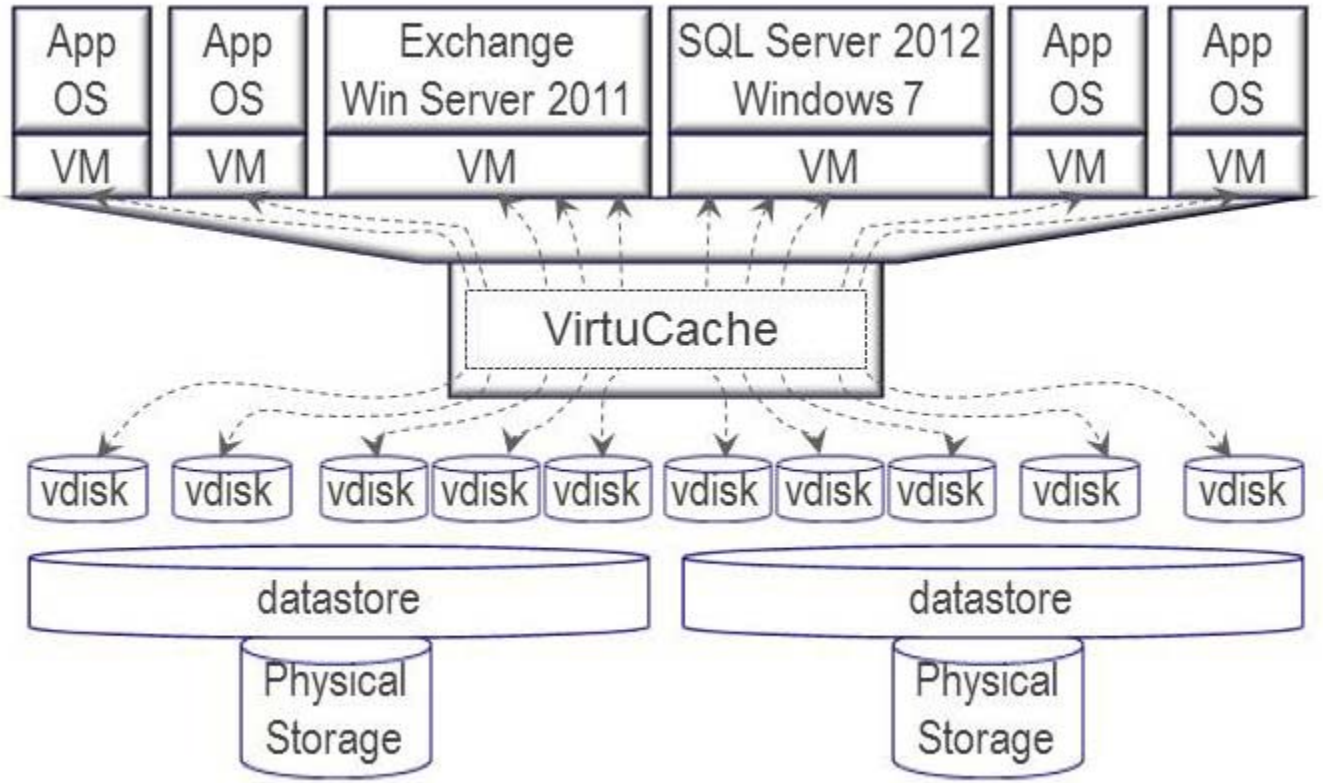

Applications
Operating System
Virtual Machines

Hypervisor
I/O Cache
(SSD & Software)
Physical Machine

Virtual disks

Virtual datastores

Physical storage

See more info and result examples at http://storageio.com/whitepaper.html

Industry Trends: Some proof points

Examples of what has been discussed here today

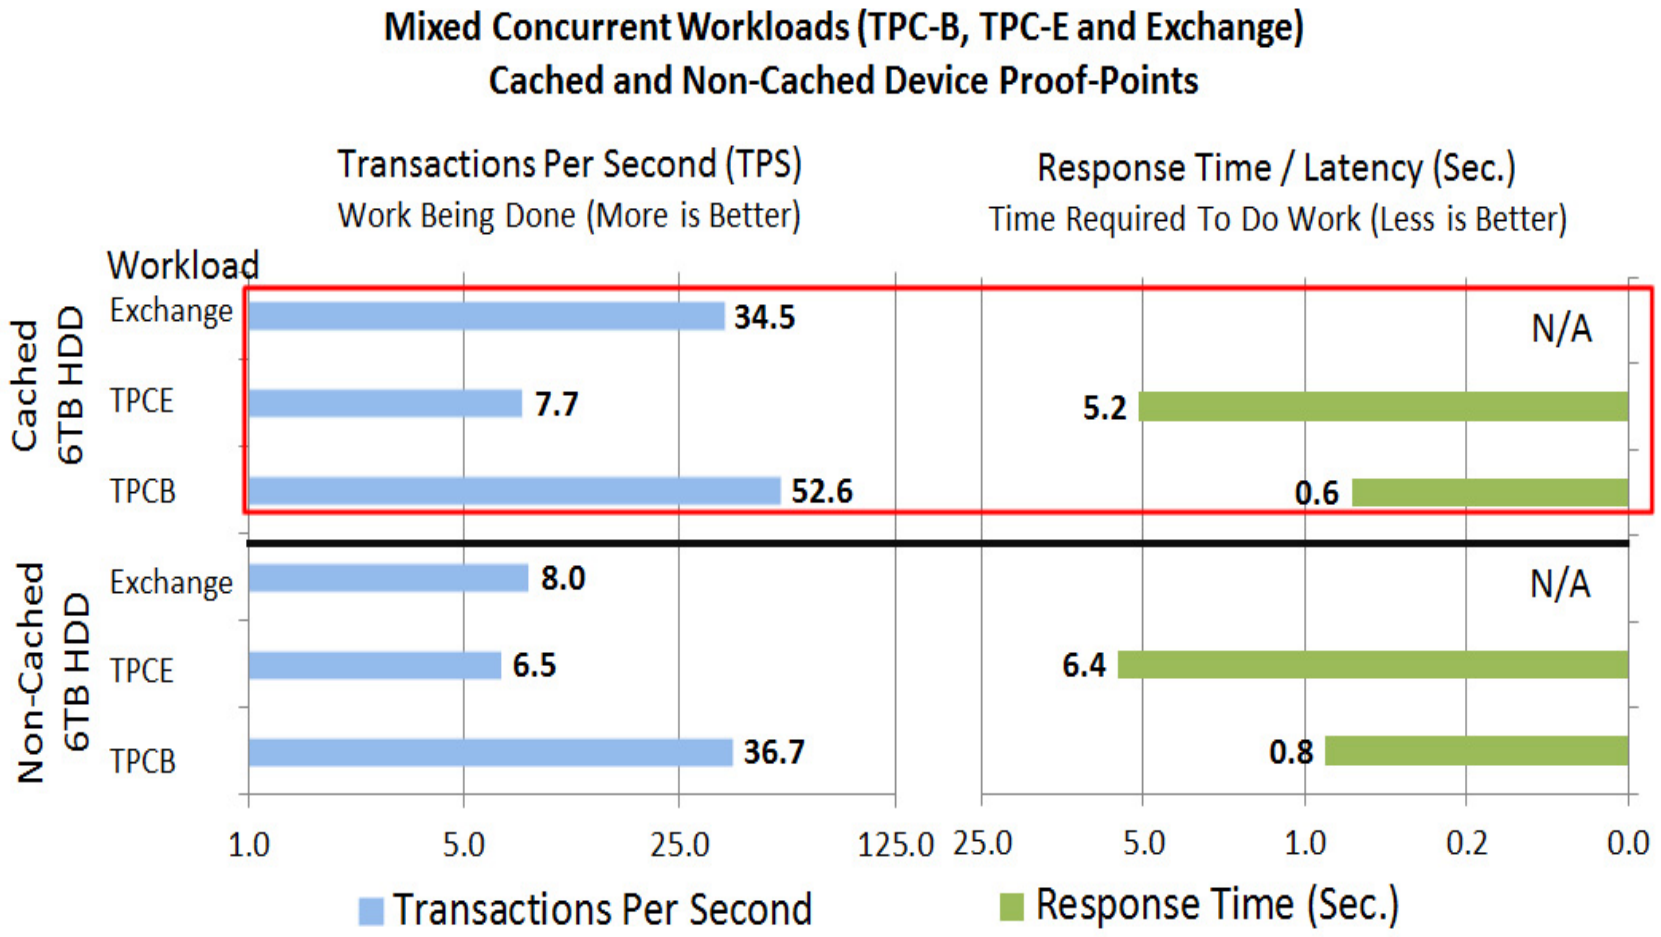

Logrithmic Scale Base 5 TPCE and TPCB = 50 Users

Source: StorageIO Labs

See more info and result examples at http://storageio.com/whitepaper.html

Industry Trends: What about the cloud?

Good question, actually there are several things to consider…

- ✓ How will you be using SSD in the cloud, from within the cloud?
- ✓ Service providers are offering various SSD "backed" e.g. based services
- ✓ AWS, Microsoft/Azure, Rackspace, etc…
  - AWS has many different services including EBS, RDS, and EC2
  - For example most think cloud storage and EBS etc
  - However there are EC2 "high-io" instances with dedicated SSDs
  - With all services understand what IOP or bandwidth limits are in place
  - Also understand if the performance will be deterministic or variable (don't assume)
  - Know how the IOPs are invoiced, for example a 32K IOP might count as two IOPs
  - Understand if there are any space capacity to IOP ratio or requirements
  - Most of the tools mentioned will work in various cloud environments ;)

http://storageioblog.com/cloud-conversations-aws-ebs-optimized-instances/

Industry Trends: Some tips

These may be obvious, however lets state the obvious

- ✓Know your applications, workloads and their characteristics
- ✓Find and identify bottlenecks, avoid simply moving them, fix them
- ✓A little cache (SSD, flash, DRAM) in the right place goes a long way
- ✓Test with or use metrics that have context to your applications
- ✓When testing, use your applications under load
- ✓Or configure workloads to be realistic of your environments
- ✓Configure system under test (SUT) to remember your environment
  - E.g. RAID, snapshots, replication, thin-provision, dedupe, compress, etc
  - Likewise are your servers and adapters faster enough
  - Pre-condition your storage before use to get realistic results

Can we get a side of context with them IOPS and other storage metrics?
http://storageioblog.com/side-context-iops/

Closing comments

Where to learn more, next steps…

Prepare and plan for your journey

- Have a vision, strategy and plan (e.g. an itinerary and road map)
- Fast servers and applications need fast storage and I/O networks
- Start using new (and old) technology in new ways removing complexity
- Reduce your data footprint impact (pack smartly for your journey)
- A little bit of flash or SSD or cache in the right place goes a long way
- A lot of flash or SSD will help, however it also will cost lots of cash

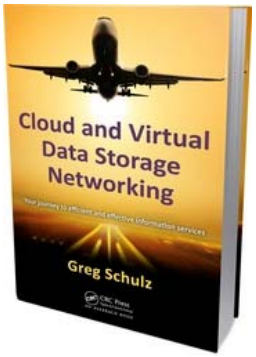

ISBN-13: 9781439851739

Where to learn more

- www.storageio.com (articles, videos & webcasts)
- www.storageioblog.com and twitter @storageio
- Check our other recent and upcoming events www.storageio.com/events
- Facebook.com/storageio and www.StorageIO.com/newsletter
- Feel free to call, IM, tweet, or email greg@storageio.com

StorageIO

Thank You

Flash back to reality – Myths and Realities
SSD Industry trends perspectives and tips

Presented by Greg Schulz, Founder & Sr. Advisory Analyst
The Server and StorageIO Group (StorageIO)
Author: Cloud and Virtual Data Storage Networking (CRC Press)
Greg@storageio.com | StorageIOblog.com | Facebook.com/StorageIO | @storageio