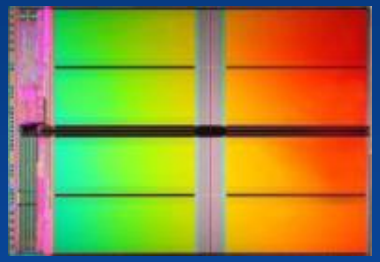

Why Flash Should Be Used For Every Workload

Even Secondary Ones

9 August 2016

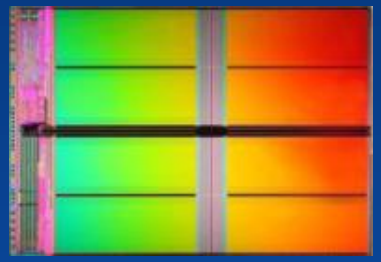

Dragon Slayer Consulting – Marc Staimer CDS

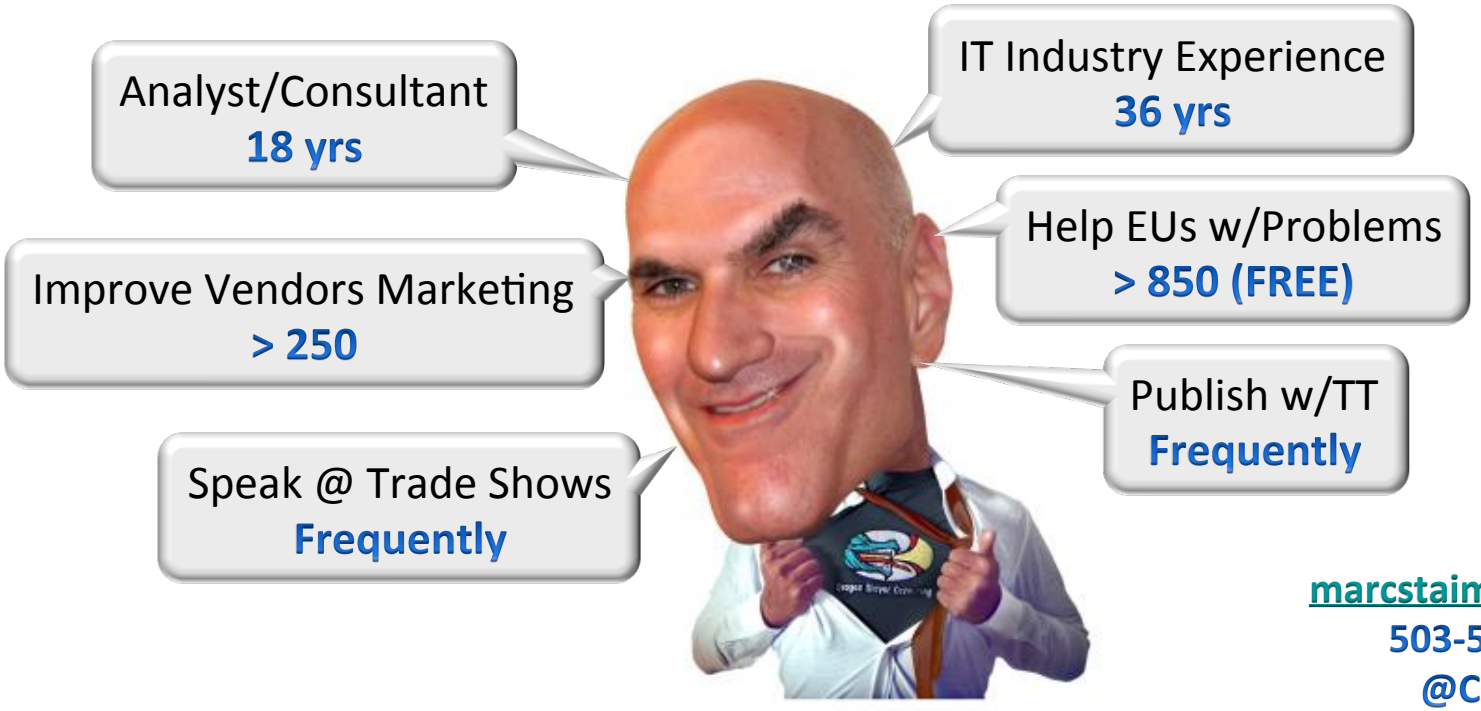

marcstaimer@me.com
503-579-3763
@ChiefDS

Yes, Yes You Will Be Drinking From A Fire Hose

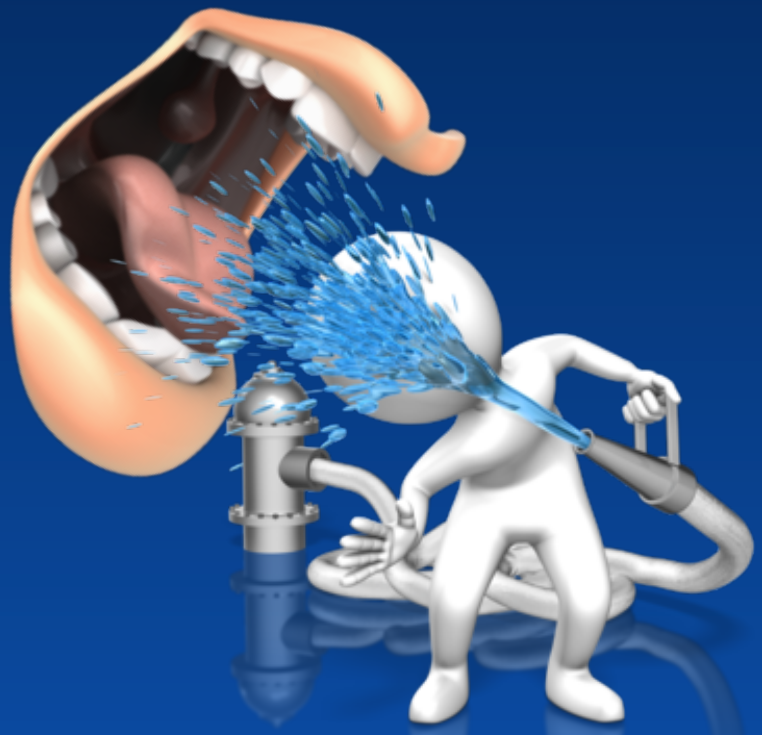

Agenda

- Flash vs. HDD Myths & Facts
- Performance Analysis
- Technical Analysis
- Financial Analysis
- Business Case
- Conclusions
- Q & A

We See But We Do Not Observe

Flash vs. HDD Myths & Facts

Myths	Facts
HDDs are more reliable	False
HDDs last longer	False
All Flash are alike	Completely untrue
HDDs are cheaper	In most cases…not true
Can't cost justify Flash	Untrue
HDD performance good enough	Usually not
Infrastructure costs =	Nope

Performance Comparison

- There's no comparison
- Flash up to 1000x > than fastest HDD

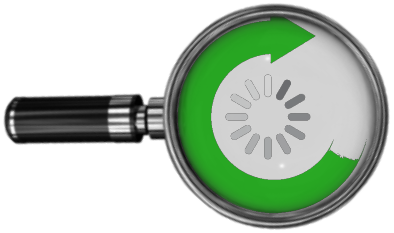

HDD vs SATA SSD Performance Comparison

Performance	HDD	Flash
Read Latency	~ 2 to 8.5 ms	~ 115 µs
Random Read IOPS	~ 200 - 400	~ 99,000
Sequential Reads	~ 160 to 250MBps	~ 540MBps
Random Write IOPS	~ 150 - 300	~ 18,000
Sequential Writes	~ 120 to 200MBps	~ 500MBps

Reliability Comparison

- HDD vs. “wear-life” (PE-Cycles) of Flash
- HDDs also have a wear-life
- Common measure DWPD
  - Complete drive writes/day 3-5 yr. warranty
- SSDs & HDDs both state their DWPDs
  - HDDs on avg. coming in lighter than SSDs

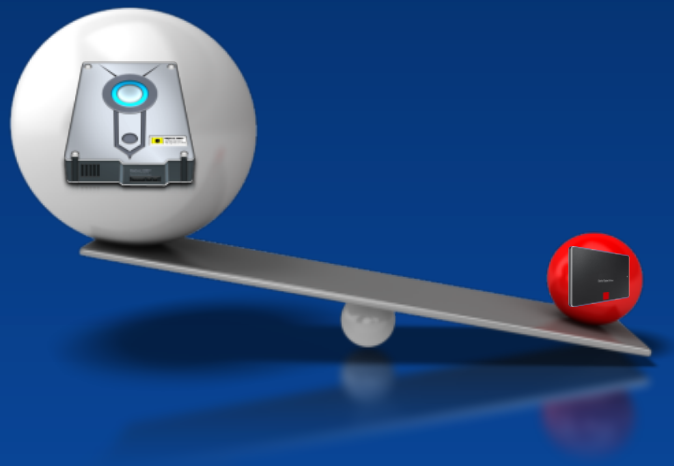

Reliability: Deeper

Reliability	HDD	Flash
MBTF	Rated @ ~ 2M hrs	Rated @ ~ 2.5M hrs.
Real World MTBF	Much < 100K hrs	~ 1M hrs
Failure types	Entire drive	Mostly write blocks
UBER	SAS 10^{16} SATA 10^{15}	SAS 10^{17} to 10^{18} SATA 10^{17}
RAID	√	√
Wide-stripe RAID	√	√
MCM	√	√
Erasure Codes	√	√

Wait There's More: HDD Silent Data Corruption

- Torn/lost/misplaced writes HDD controller misses
- Doesn't inform RAID controller
- Parity pollution locks in no correction possible
- 10x > on SATA vs SAS
  - SAS costs more

Flash SSD UBER Advantage

HDD Unrecoverable Bit Errors

- Treated as drive failure
- Drives taken offline, rebuilt

SSD Unrecoverable Bit Errors

- Treated as P/E block failure
- Capacity over-provisioned
  - Replace failed P/E block
  - From internal overage

HDD Rebuilds Take A Very Long Time

Rebuilds used to be measured in hours

- Today it's days, weeks, or months
  - Depending on priority
- During rebuilds
  - Controller performance declines
  - Data loss risk increases
- Flash SSDs rarely need to be rebuilt

Why Erasure Codes Are Popular For HDDs

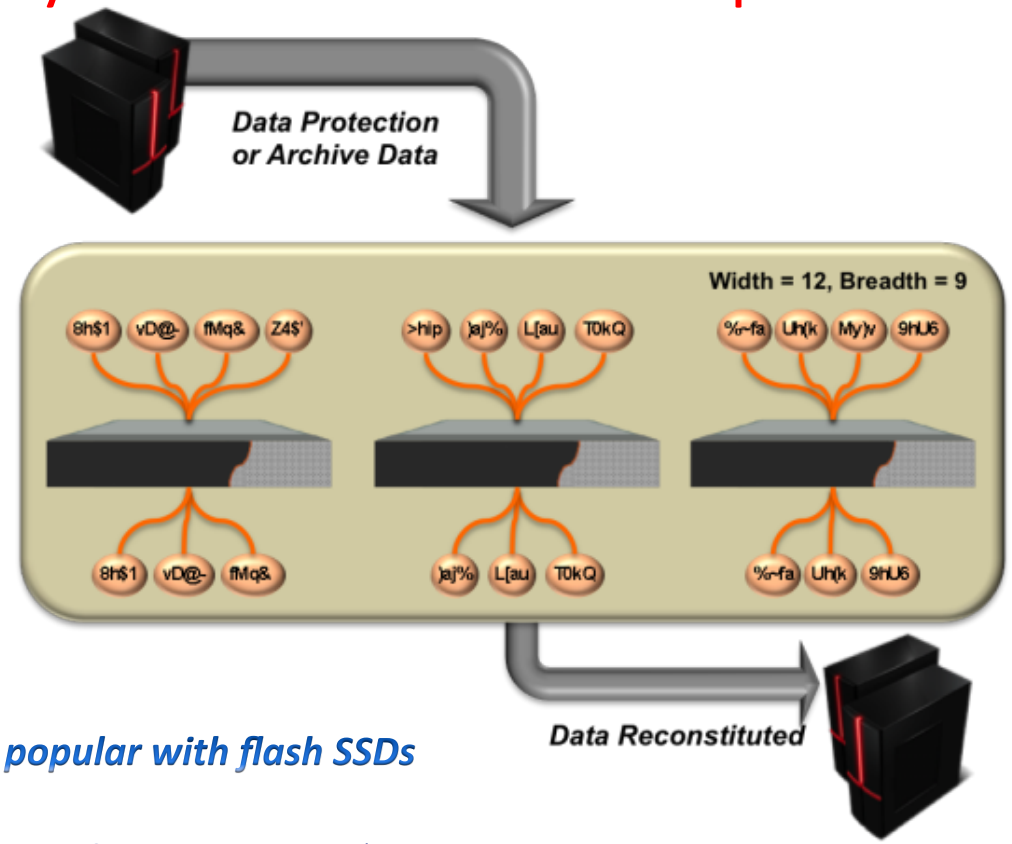

Will be just as popular with flash SSDs

Even Expensive Multi-Copy Mirroring (MCM)

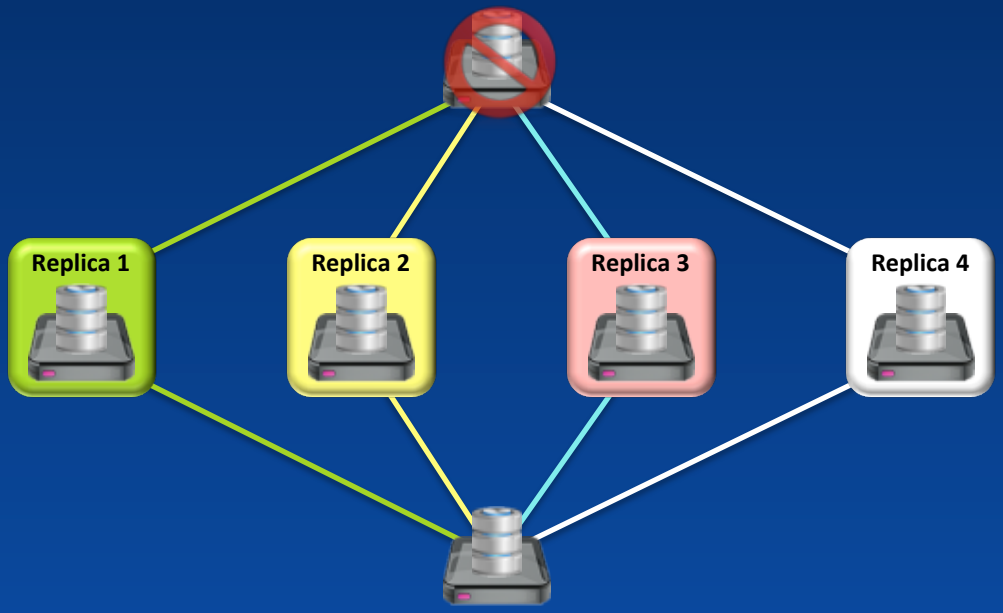

Infrastructure Costs Are Quite Different

Flash storage sold differently from HDDs

- HDDs sold by # drives required for performance =
  - Over capacity, & > shelves, rack space, floor space
  - + > infrastructure (switches, cables, conduit, power/cooling)
  - Which consumes yet more DC real estate
- Flash SSDs sold on performance & capacity requirements
  - Require much < drives & < capacity

Density For Secondary Storage A Factor

All technology advances – some faster than others

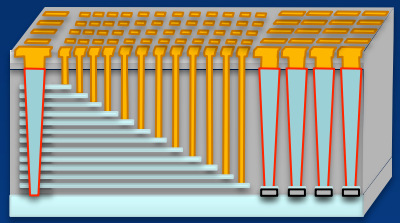

- HDDs
  - PMR – perpendicular magnetic recording
  - Helium sealed
  - SMR – shingled magnetic recording
  - HAMR – heat assisted magnetic recording
- Flash SSDs
  - 15nm 2D planar MLC
  - CFMs – custom form modules
  - 32/48/64/128 layer TLC 3D NAND
  - 3D QLC 3D NAND coming

Timeline – Capacity Comparison

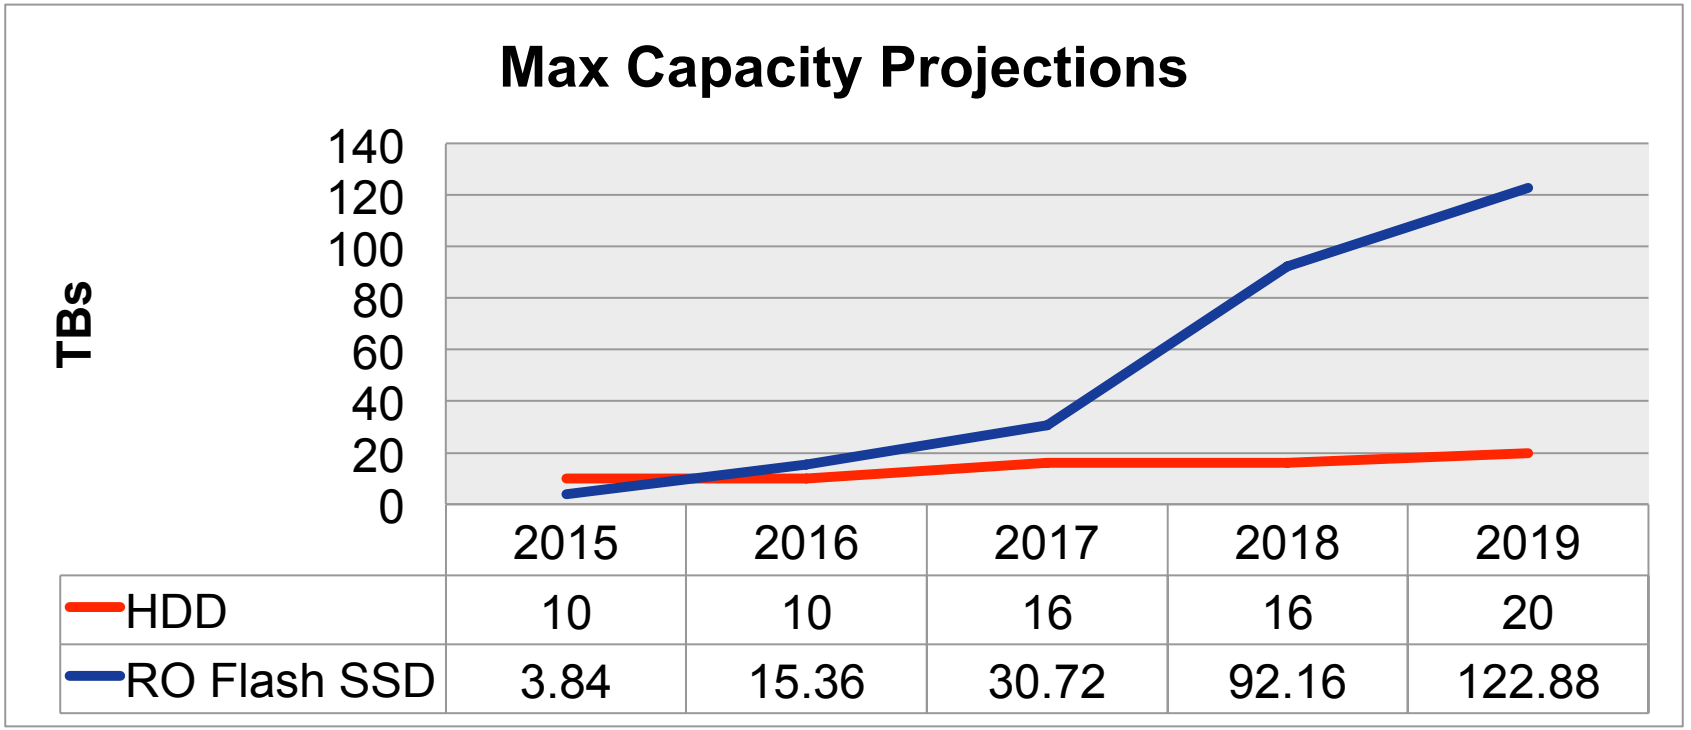

	2015	2016	2017	2018	2019
HDD	10	10	16	16	20
RO Flash SSD	3.84	15.36	30.72	92.16	122.88

Timeline – Density Comparison

Max Density/2U

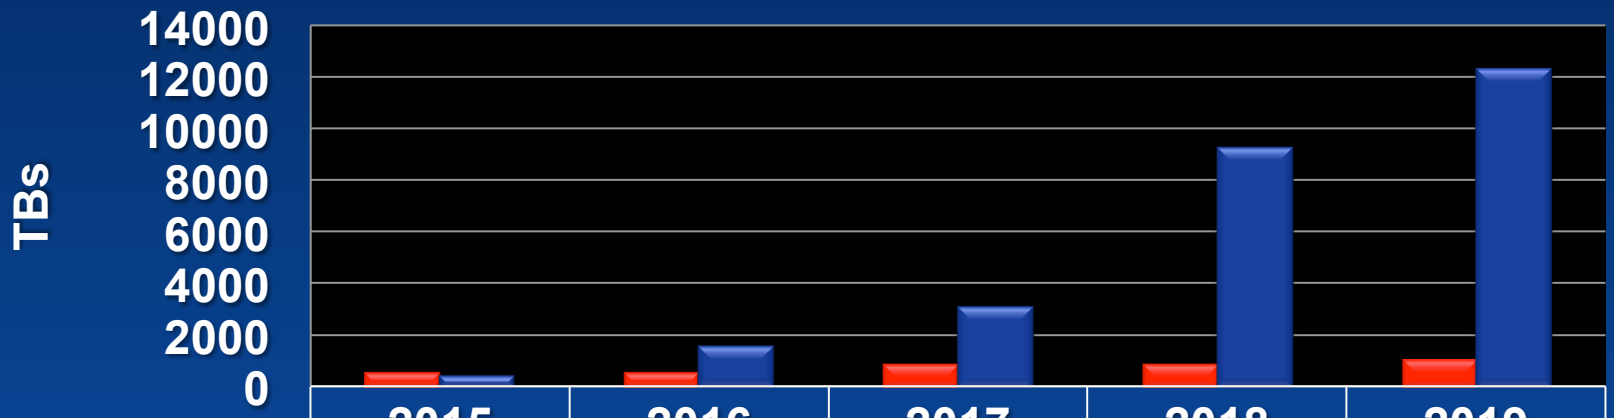

	2015	2016	2017	2018	2019
HDD	490	490	784	784	980
Flash SSD	384	1536	3072	9216	12288

Rack Density Last Year 2015 Comparison

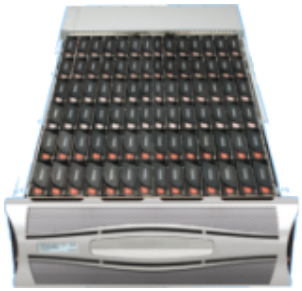

HDD

- Max Density
  - 980TB 4U
  - > 300 lb drawer
  - Requires server lift + ladder

Flash SSD or CFM

- Max Density
  - 512TB (raw) 3U – CFM
  - 384TB (raw) 2U – 2.5” SSD
  - Requires ladder

Rack Density 2016 Comparison

HDD

- Max Density
  - 980TB 4U
  - > 300 lb drawer
  - Requires server lift + ladder

Flash SSD or CFM

- Max Density
  - 1.024PB (raw) 3U – CFM
  - 1.536PB (raw) 2U – 2.5” SSD
  - Requires ladder

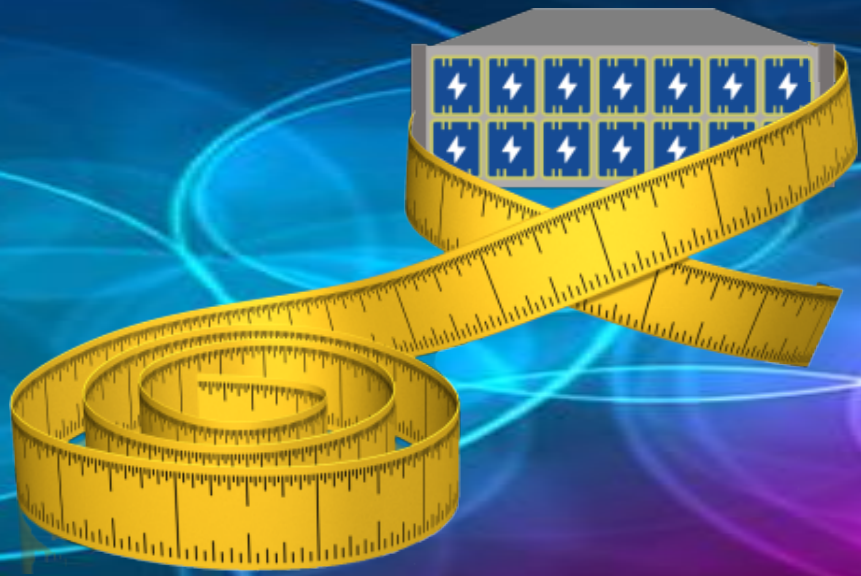

What About Measuring Cost?

Cost Focus Frequently Misplaced

Tends to be $/TB or comparative $/TB

- Acquisition costs
- Mostly $/raw TB

Should Be On Total Cost of Ownership (TCO)

TCO has multiple moving parts

Understanding Effective Capacity

- Raw capacity
  - Untouched out of box
- Usable capacity – after
  - Formatting
  - File system
  - RAID
- Effective capacity – after
  - Compression
  - Dedupe
  - Encryption

Crucial Storage Cost Metrics

Levels cost playing field

- Acquisition $ +
- 3-5 yr OpEx -
- Infrastructure reduction savings ÷
- IOPS (or MBps) ÷
- Effective usable TBs =
- $/IOPS/TB or
- $/MBps/TB

Comparative Metrics

Levels playing field

- BUT…
  - Only looks at system costs
  - Doesn't look at > productivity costs
  - Doesn't look at much > increased revenue

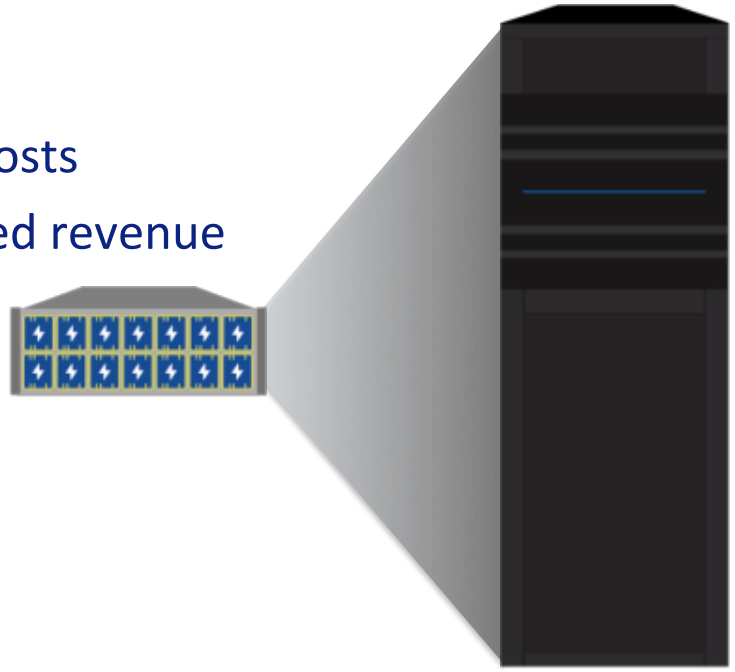

Measuring Productivity

IBM's "The Economic Value of Rapid Response Time"

1st published 1982 repeated since

- Faster app response times = much > productivity
- In addition
  - Application costs plummeted
  - Much > user work satisfaction
  - Much > morale
  - Much > work quality
- Findings surprise
  - as little as .5 seconds = huge impact
- Each study repeat = similar results

Measuring Productivity Gains

Application Response Time (Sec)	Aggregate Productivity Gains	Net Productivity Gains
3	0%	0%
2	16%	16%
1	37%	21%
0.6	47%	11%
0.3	80%	33%

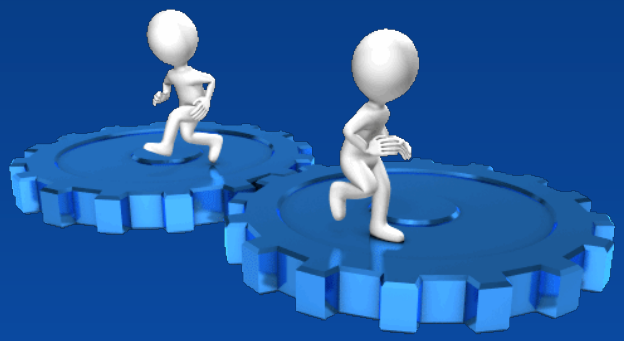

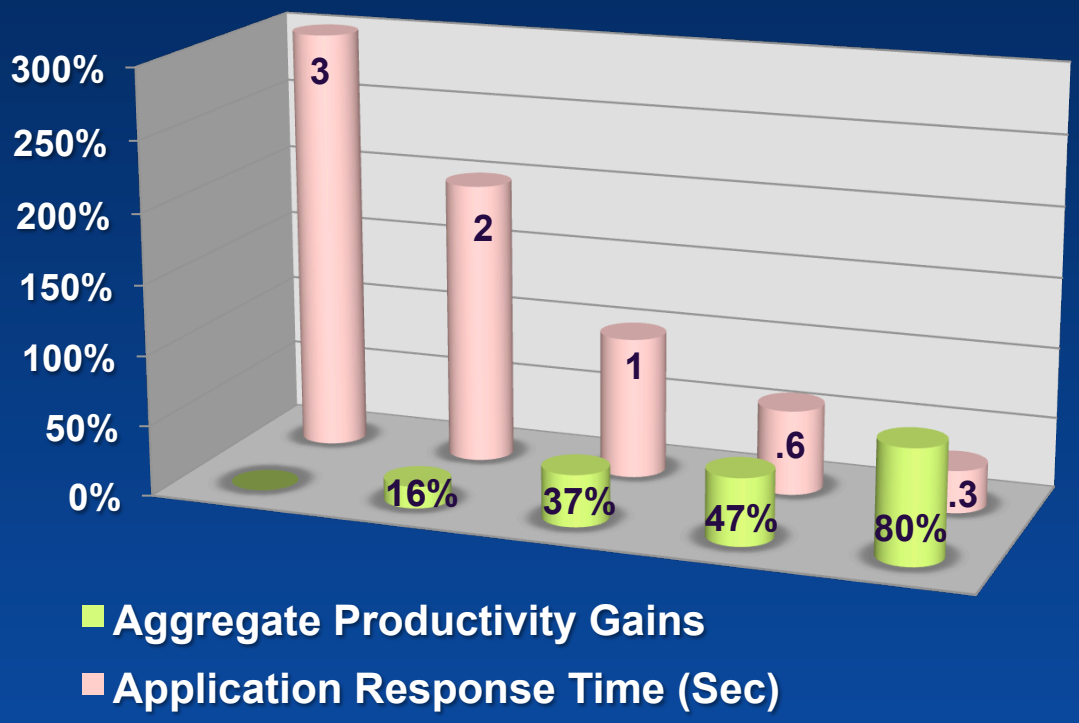

Converting Worker Productivity To Hard $

Before Flash Storage

- Avg app response time = 1 sec
- Requirement = .3 sec
- 44% productivity gained
  - When achieved

Worker Productivity Gains

- Avg worker burdened* cost
  - ~ $74,193 / yr (per USBLS)
- Productivity savings
  - ~ $32,645 / yr / worker
  - 100 workers = ~ $3,264,500 / yr

**Burdened costs include benefits, payroll taxes, and allocated fixed overhead*

Converting IT Productivity To Hard $

Before Flash Storage

- Troubleshooting performance
- Tuning performance
- ~ same 44% productivity savings

IT Admin Productivity Savings

- Avg burdened* cost
  - ~ $100,000 / yr
- Productivity savings
  - ~ $44,000 / yr / IT admin
  - 5 IT admins = ~ $220,000 / yr

**Burdened costs include benefits and allocated fixed overhead*

But This Is Sunk Costs! Isn't It?

No it is not because work not completed in time =

- Overtime = > $
- Or > workers = > $
- Or > outside contractors = > $
- < Morale
- > Turnover
- > Hiring $
- > Training $

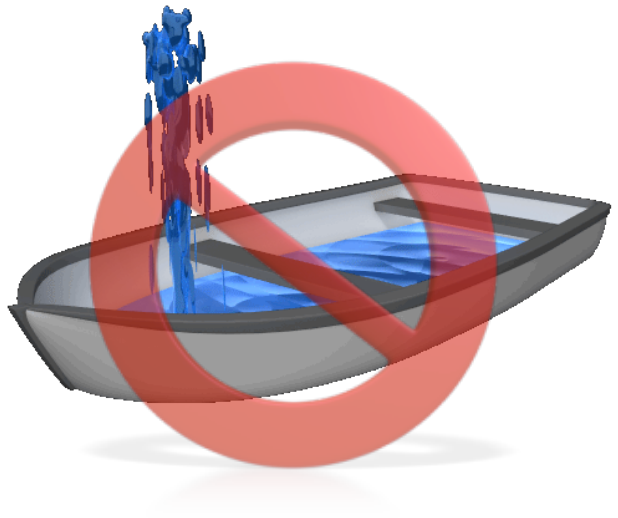

This Is Real Money Being Flushed Away

Measuring Revenue Gains/Losses

Response Time Revenue Implications

Response time variance (≤ .5s) has huge impact

- Google found
  - Std search results are 10/page/search
  - Users wanted 30/page/search
  - Added .5s/response
  - Traffic & revenue dropped 20%

Response Time Revenue Implications

Response time variance (≤ .5s) has huge impact

- Amazon found similar results
  - Experimented w/100ms increments
  - Even 100ms = significant lost revenue
  - Worse, inconsistent response time
  - A 1s slowdown = ~ **$1.6B** lost sales/yr

Response Time Revenue Implications

Response time variance (≤ .5s) has huge impact

- AOL discovered
  - Page views decline as load times >
  - As much as a 33% decline
  - Between best & worst load times

Response Time Revenue Implications

On the flip slide reducing response time has powerful impact

- Shopzilla found response time reduction
  - From ~ 7s to ~2s had positive results
  - 25% > page views
  - 7-12% > revenue
  - 50% < hardware

Every Industry Has Similar Revenue Cases

Unique $$$$ tied to faster time-to-market

- Financial services
- Insurance
- Online gaming
- Web retail
- Healthcare
- Energy
- Life Sciences
- Big Pharma

Calculating That Material Revenue Impact

Projected Revenue Gains Formula

- T = # of quarters gained in time to market
- R = projected revenue/quarter gained
- EQGR = expected quarterly growth rate
- TQ = total # of quarters

Total Net Revenue Gained =

$$[(T*R)*(1+EOQR)^{(TQ)}] - [(T*R)*(1+EOQR)^{(TQ-T)}]$$

Yeah But, What About Secondary Applications?

They're secondary or batch applications for a reason…BUT

- What's your cost for missing the time window?
  - DP (backup) job
  - DBMS copy & replicate
  - Recovery RTOs
  - BI analytics
  - Archive ingest

Huuuuge Costs For Missing Time Windows

- Business
- Revenue
- Reputation
- Clients
- Customers
- Possibly the business

Putting It All Together

Complete Flash Storage Financial Analysis Formula

For accelerated application response time

- Productivity savings +
- Net revenue gains +
- Hardware & infrastructure savings –
- Flash Storage TCO =
- ***Expected Value***

What's Required Frequently ≠ What's Delivered

Thank You!

marcstaimer@me.com
503-579-3763
@ChiefDS