Why SSDs Should Be Utilized For All Data Center Workloads

Even Secondary Ones

Dragon Slayer Consulting – Marc Staimer CDS

Analyst
17 yrs

Consultant
> 250 clients

Speak @ Trade Shows
Frequently

IT Industry Experience
35 yrs

Help EUs w/Problems
> 1000 (FREE)

Publish w/TT
Frequently

marcstaimer@me.com
503-579-3763

Agenda

- Conventional HDD & SSD Wisdom
- Fact Checking Conventional Wisdom
- The Value Of Time
- Comparing HDDs & SSDs
- Conclusions
- Q&A

Pay Attention, You Might Miss Something

HDD Conventional Wisdom

HDD Perceptions

- Reliable
- High capacity/density
- Balanced costs
- Good enough performance

HDD Conventional Wisdom

HDD Perceptions

- Avg failure rates
- High weight
- Slow performance/> latency
- Cost declinations decreasing

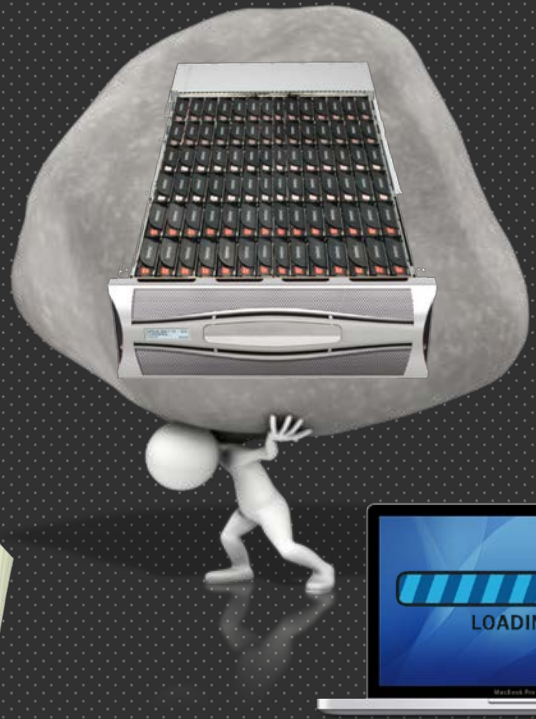

HDD Reality

HDD Reliability

Reliability	Perception	Reality
MTBF	High @ 1.5M hrs	< 100,000 hrs
HDD Failures	Individually	In Bunches
BER	SAS 10^{16}	> BER/GB
	SATA 10^{15}	
RAID	Known, proven, & good enough	Rebuilds too long
		> risk data loss
Wide-stripe RAID	Faster rebuilds	+ > % of data loss
MCM	Dial up/down DP	Consumes > storage
Erasure Codes	Dial up/down DP	Adds latency

HDD Unique Corruption

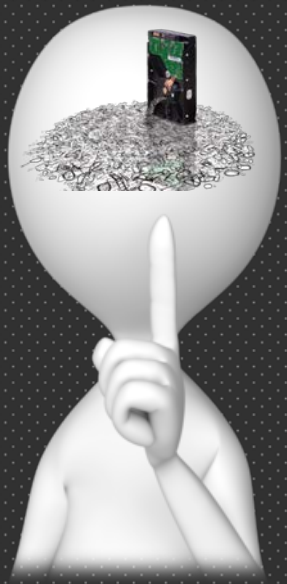

Silent Data Corruption

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

HDD Rebuilds Take A Very Long Time

Rebuilds used to be measured in hours

- Today rebuilds are measured in days, weeks, even months
  - Depends on priority
- During rebuilds
  - Controller performance declines
  - Data loss risk increases

HDD Failure Rates Increase Over Time

Traditional parity RAID

- Not keeping pace w/HDD density
- Part of that is unrecoverable BER
- BER unchanged
  - 1, 2, 3, 4, 6, 8, 10 TB the same
  - SATA – 10^{15}
  - SAS – 10^{16}
- > capacity = > BER/GB

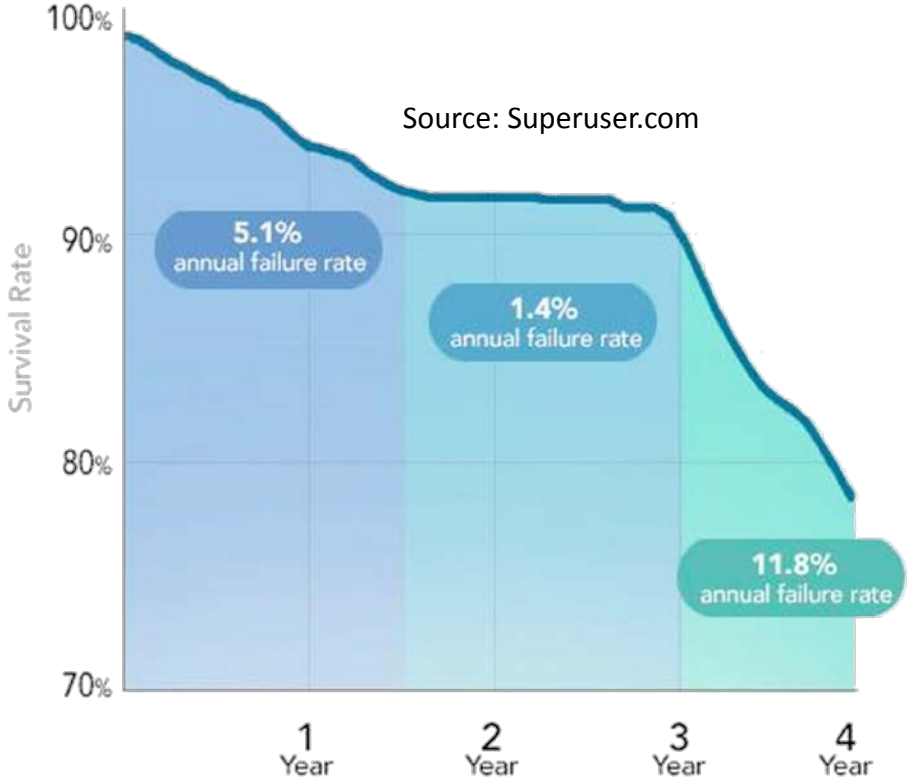

What Exactly Are Erasure Codes?

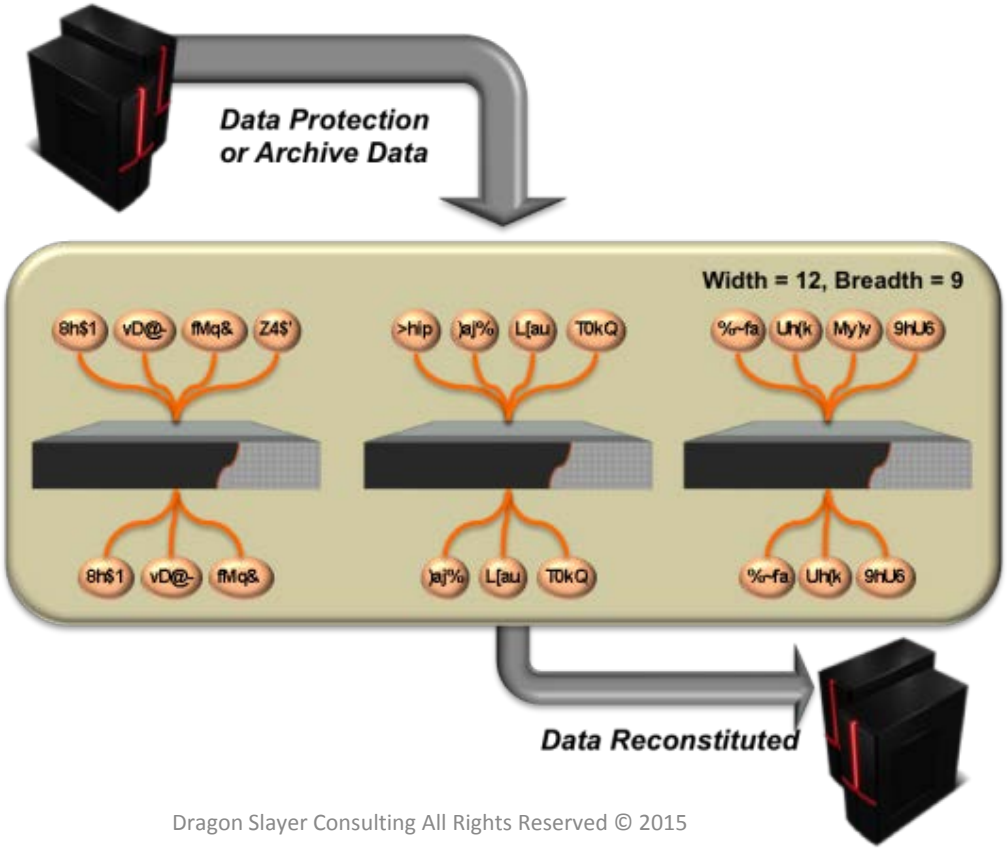

HDD Density

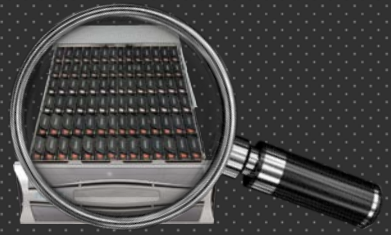

Density	Perception	Reality
Capacity	Highest	LFF up to 10TB, SFF up to 1.2TB.
Density Footprint	Highest	LFF raw yes up to 840TB/4U: no for primary data reduced.
Weight	Manageable	SFF true, LFF not. High density racks can weigh > 250lbs.
Data Reduction	Good	Secondary data true (~4:1 or betterr); Primary data not true (~2:1)
Serviceability	Easy	SFF true, LFF HD racks, not true, ladders & server lifts.

HDD Performance

Performance	Perception	Reality
Latency	Fair	Poor
IOPS	OK w/high RPM SFF & short-stroking or DRAM caching	Not good w/random IOPS. Short-stroking kills 2/3 capacity
Throughput	Good	Sequential reads & writes yes, random IOPS no

HDD Cost

Cost: $/GB	Perception	Reality
Acqusition	Lowest	LFF true, SFF sometimes
Power & cooling	Highest	True for all form factors
$/IOPS/GB	Highest	Getting higher
Productivity	Adequate	Large productivity losses
TCO	Lowest	LFF mostly true, often not, & not true for SFF

HDD Revenue Impact

Rev Impact	Perception	Reality
Time-to-market	On time	Late
Competitive	Yes	No

Flash SSD Conventional Wisdom

SSD Perceptions

- Very fast
- Very low latency
- Low power/cooling
- Light weight

Flash SSD Conventional Wisdom

SSD Perceptions

- Small capacity/density
- Low reliability
- High cost

SSD Reality

SSD Reliability

Reliability	Perception	Reality
MTBF	Low to medium	DWPD ranges from .5 to 35
SSD Failures	Equivalent to HDDs	Much < HDDs. Whole drives rarely fail
BER	Equivalent to HDDs	Much > HDDs
RAID	Required	Unnecessary. Parity reduces SSD wear life
Wide-stripe RAID	Faster rebuilds	Rebuilds not an SSD issue
MCM	Works as well w/SSDs	Consumes > storage
Erasure Codes	Works as well w/SSDs	Adds latency

SSD Density

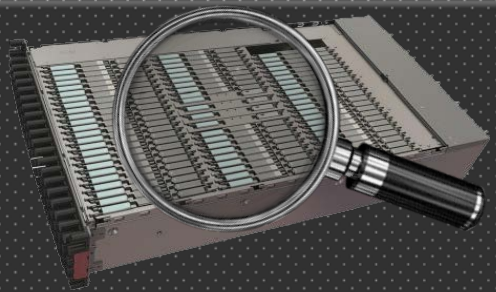

Density	Perception	Reality
Capacity	Lower than HDDs	SFF higher at up to ~ 4TB[1]
Density Footprint	Lower than HDDs	High density racks w/up to 512TB in 3U
Weight	Low	< 10% of equvalent HDD capacities
Data Reduction	Excellent	2 to 3x better than HDDs for primary workloads
Serviceability	Easy	SFF yes, ladder required for HD racks to access SSD

[1]NOTE: 8TB even 12TB SFF SSDs are coming by end of 2016

SSD Performance

Performance	Perception	Reality
Latency	Exceptional	Orders-of-magnitude >HDD
IOPS	Exceptional	Up to 1000x > equivalent capacity HDDs especially random IOPS
Throughput	Exceptional	TRUE

SSD Cost

Cost: $/GB	Perception	Reality
Acqusition	Highest	SFF SSDs < SFF HDDs/GB
Power & cooling	Lowest	True for all form factors
$/IOPS/GB	Lowest	90% < HDDs per GB
Productivity	Best-in-class	Huge Δ w/HDDs
TCO	Highest	< SFF HDDs, > LFF HDDs

Wikibon's Prognostication

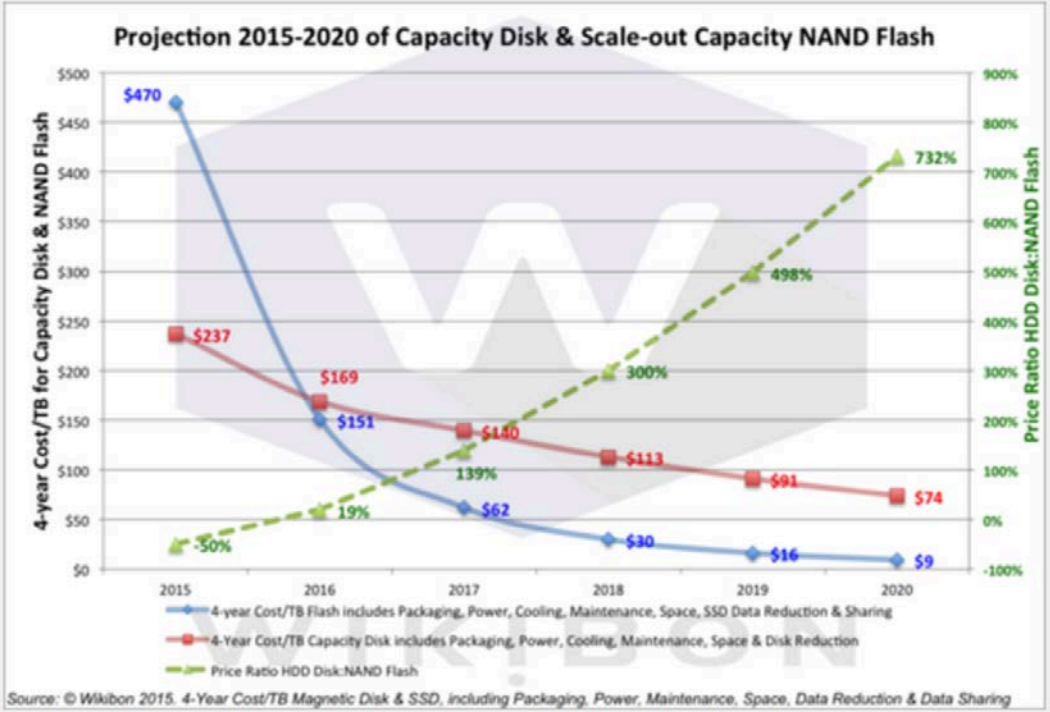

Figure1 – Projection of Technology Costs of Capacity Magnetic Disk Drives and Flash Technologies 2015 - 2020
Source: © Wikibon 2015. See Figure 2 below for more details and assumptions

Wikibon's Prognostication

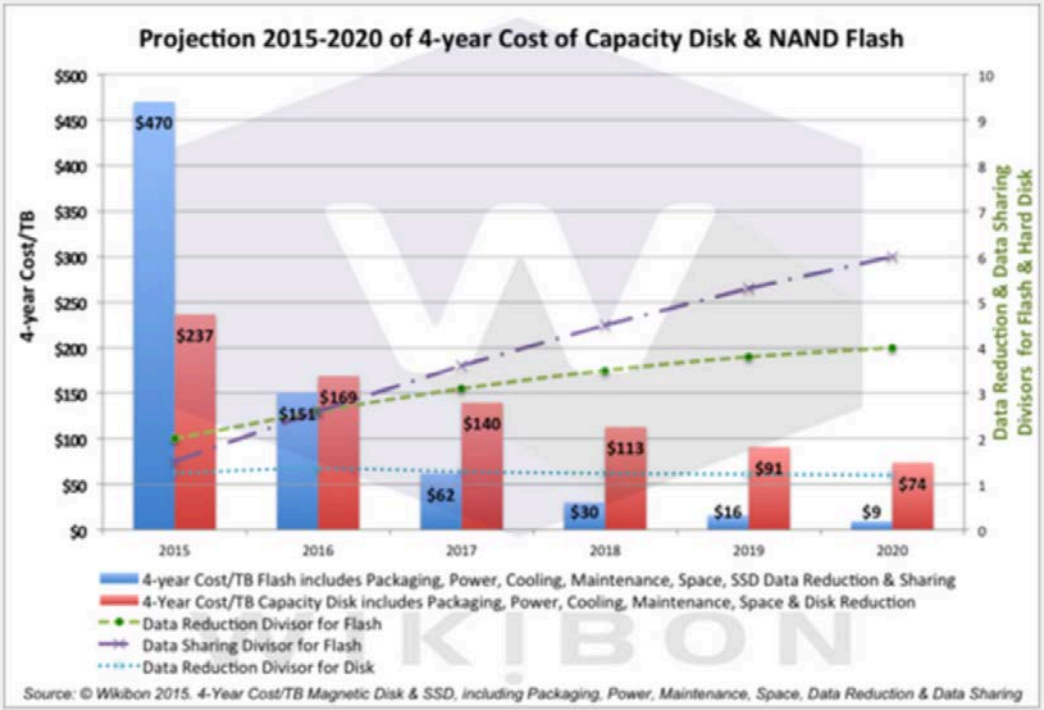

Figure 2: Projection of Technology Costs of Capacity Magnetic Disk Drives and Flash Technologies 2015 - 2020

Sources: © Wikibon 2014. See Figure Footnotes-1 and Table Footnotes-1 in the Footnotes below.

SSD Revenue Impact

Rev Affect	Perception	Reality
Time-to-market	Faster	Much faster
Competitive	Yes	Yes

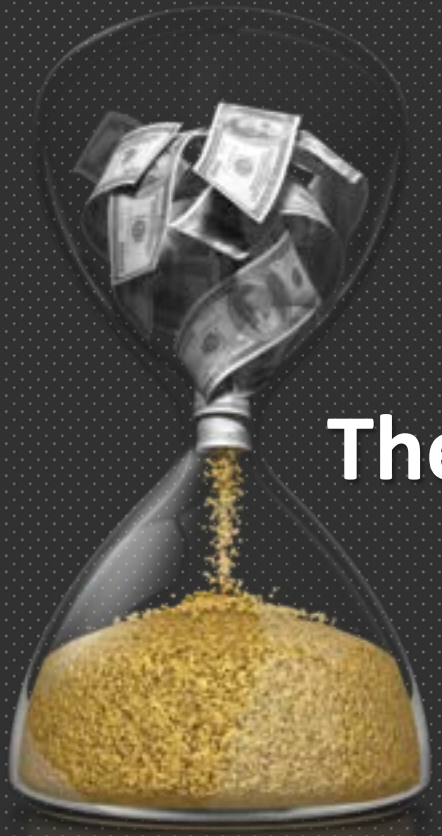

The Value of Time

Quotes About Time Value

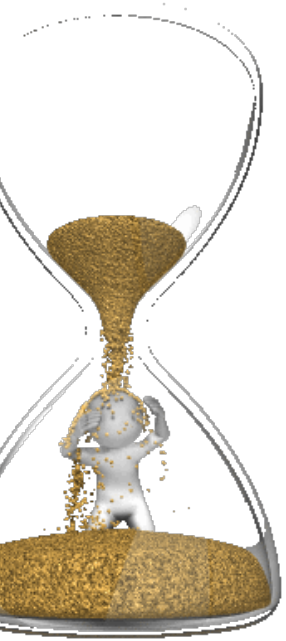

"Time is what we want most, but what we use worst." *William Penn*

"Time is relative." *Albert Einstein*

"Time elapses at different speeds depending on what side of the bathroom door you find yourself." *Albert Einstein*

"Time is money." *Benjamin Franklin*

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

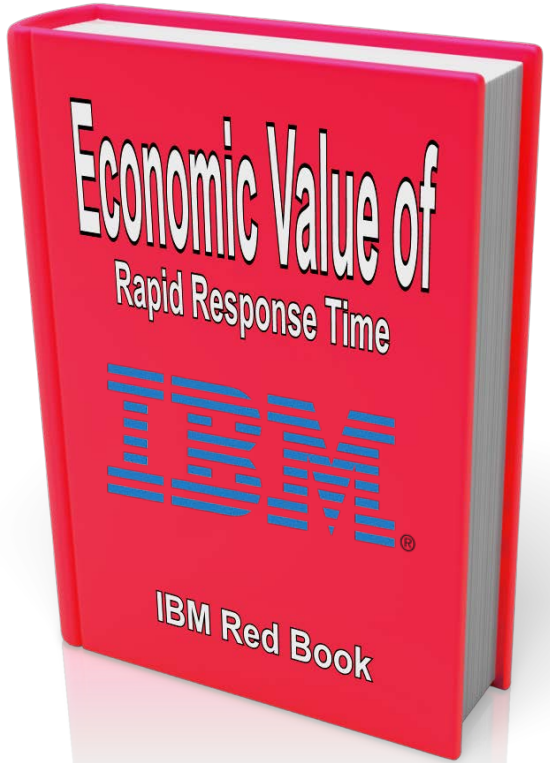

Measuring Productivity Gains

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%

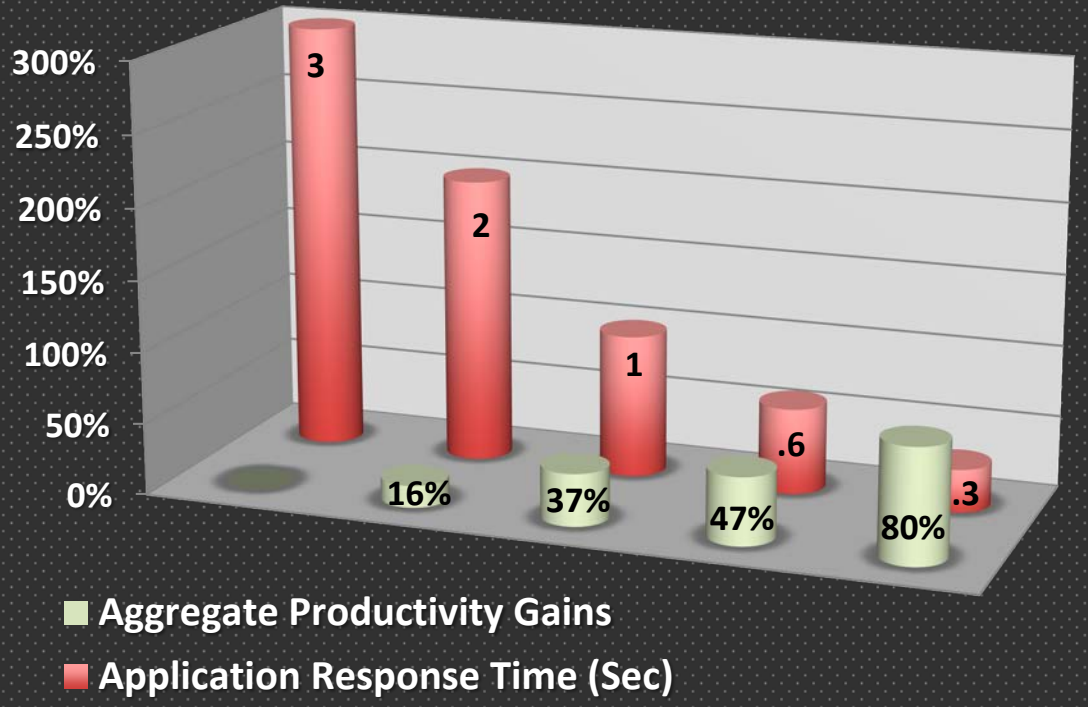

Measuring Productivity Losses

Application Response Time (Sec)	Aggregate Productivity Losses	Net Productivity Losses
0.3	0%	0%
0.6	25%	25%
1	34%	10%
2	52%	17%
3	65%	13%

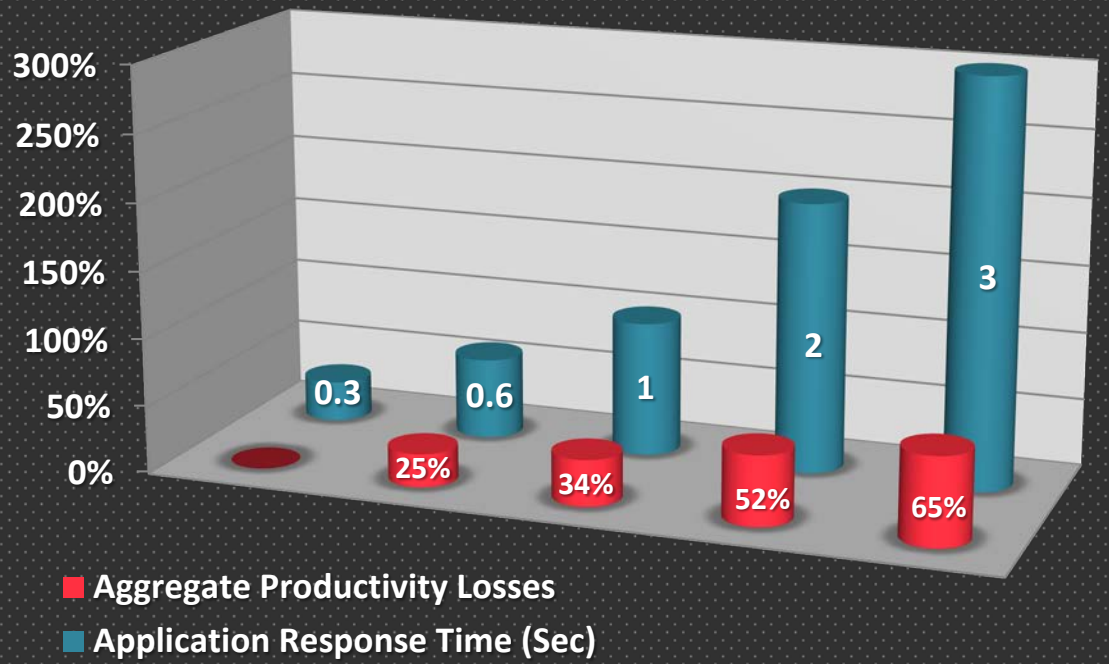

Converting Into Hard $$$$

- 1982 IBM used $35/hr burdened
  - = $87.43/hr burdened[1] in 2015
  - We'll use just $35/hr
- App requirement
  - 1s response time
  - Actual avg response time = 3s
  - 30% productivity loss
- Worker productivity costs
  - $35/70% - $35 = $15/hr
  - $15 * 2080 hrs = $31,200/yr
  - 100 workers = $3,120,000/yr

[1]Based on US Bureau of Statistics calculations

Converting Into Hard $$$$

- For IT workers
  - Burdened costs are higher
  - ~ $48/hr
  - $48/70% - $48 = ~ $20/hr
  - $20 * 2080 hrs = $41,600/yr
  - 100 workers = $4,160,000/yr

And It's Not Sunk Cost

Worker Productivity Hard $

- Work not completed in time
  - Overtime
  - Or more workers
  - Equals > costs
  - 1 women = 1 baby 9 mos
  - 9 women ≠ 1 baby 1 mos
  - Rev permanently lost
  - = < morale > turnover
  - > turnover = > hiring/training $

IT Worker Hard $

- < performance = > trouble shooting
  - Tune storage performance
  - Fix when drives fail
  - Places heavy burden on IT
  - = < morale > turnover
  - > turnover => hiring/training $

Converting Productivity Into Revenue Loss/Gain

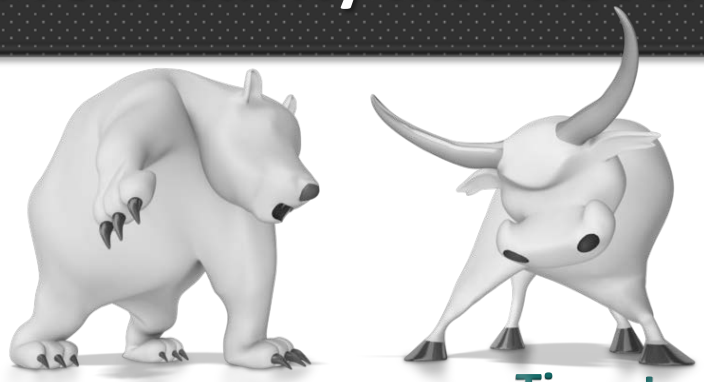

Simple formula

- T = projected time
  - Late or early to market
- R = projected revenue
  - For that “T”
- G = projected CQGR
  - Compounded quarterly growth %

Time-to-market revenue loss/gain

- Calculate T & R
- Compound CQGR of R
- Loss R never made up
  - Losses run into $Millions

What About Secondary Applications?

Time Can Be A Major Cost Factor There As Well

- Example: Job can't complete in time window allowed
  - DP (backup) job
  - Hadoop analytics
  - DBMS copy & replicate
  - Archive ingest

How Do HDDs & SSDs Compare

	2.5" SFF SAS/SATA SSD	2.5" SFF 15/10K RPM HDD	3.5" LFF 7.2K RPM HDD
Reliability			
Density			
Performance			
Cost			
Revenue Impact			

Conclusions

What Have We Learned?

- Flash SSDs are very cost effective for:
  - All primary workloads
  - Some secondary workloads
  - Time value taken into account
  - Should be used for ***ALL*** workloads

Not Magic, Ideology, or Hype, Just Storage Facts

Q & A

Marc Staimer, President & CDS
Dragon Slayer Consulting
marcstaimer@me.com
503-579-3763