Statistical Math for VDI Design

Underlying principles for VDI sizing…

Daniel Beveridge – Sr. View Architect, Office of CTO

Designing to scale

What's the best way to determine the IOPS load a storage subsystem should accommodate?

- Average IOPS
- Peak IOPS
- Avg of peak IOPS
- Someplace in between?
  - How far in between?

- **How can you project a sense of confidence in your sizing estimates?**

Introduction to Standard Deviation

- **A short story about Standard Deviation**
  - First presented at VMworld 2009, DV3266 “Efficient VDI Design, The Math behind VDI”
- **Once you understand and appreciate Standard Deviation, you will see it everywhere**
- **Harnessing the power of Standard Deviation provides you:**
  - Clear direction on what performance data is relative to design
  - A quantifiable rationale for sizing estimates
  - Confidence in your estimates
  - A means of adjusting designs to suit constraints while ensuring SLA targets are met.

Workload basics: Comparing Server to Desktop

Server	Desktop
Diversity on cluster	Homogeneous on cluster
Peaks randomized across apps	Synchronized peaks
App driven I/O	User driven I/O
Maintenance is manual or targeted	Maintenance is often synchronized on many machines
Batch-oriented	Real-time oriented
Performance fluctuations have low to moderate visibility and impact	Performance fluctuations have high visibility and real-time impact
Low to moderate workload volatility (peak < 10x avg)	High workload volatility (peak > 50x avg)

- **Key differences between workloads**
  - Vastly different volatility
  - Users react in real-time to performance fluctuations
  - Severity of reaction to performance fluctuations

Volatility Examples – High and Low

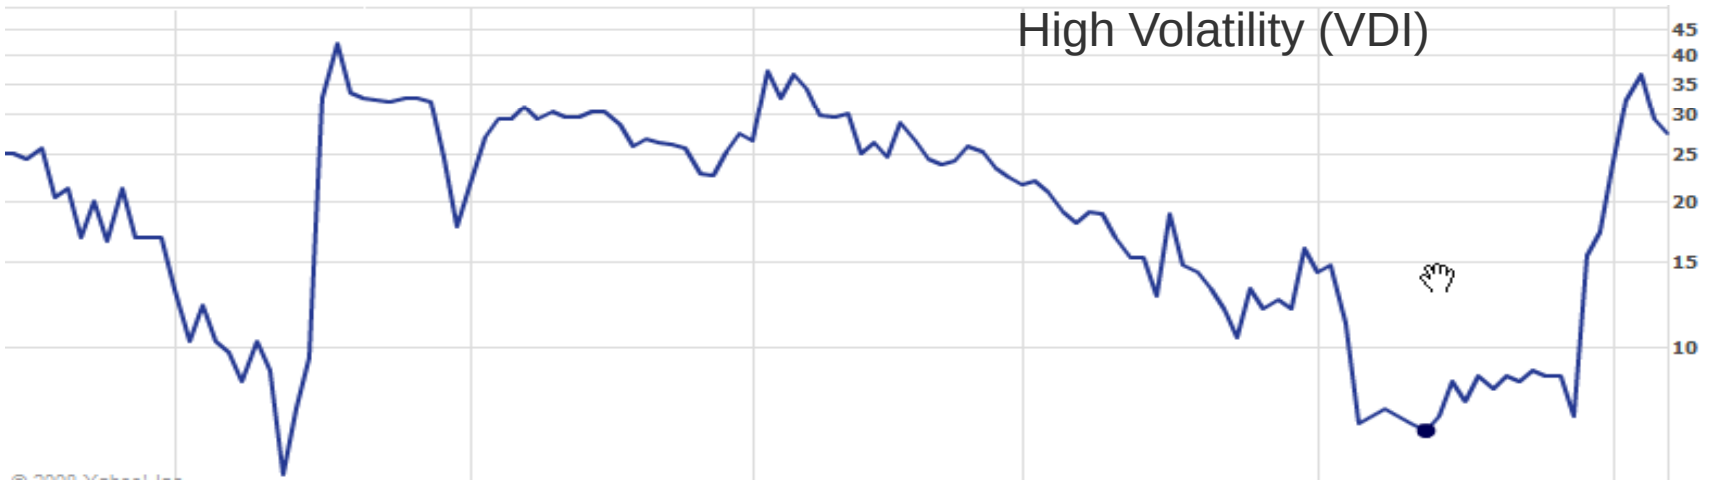

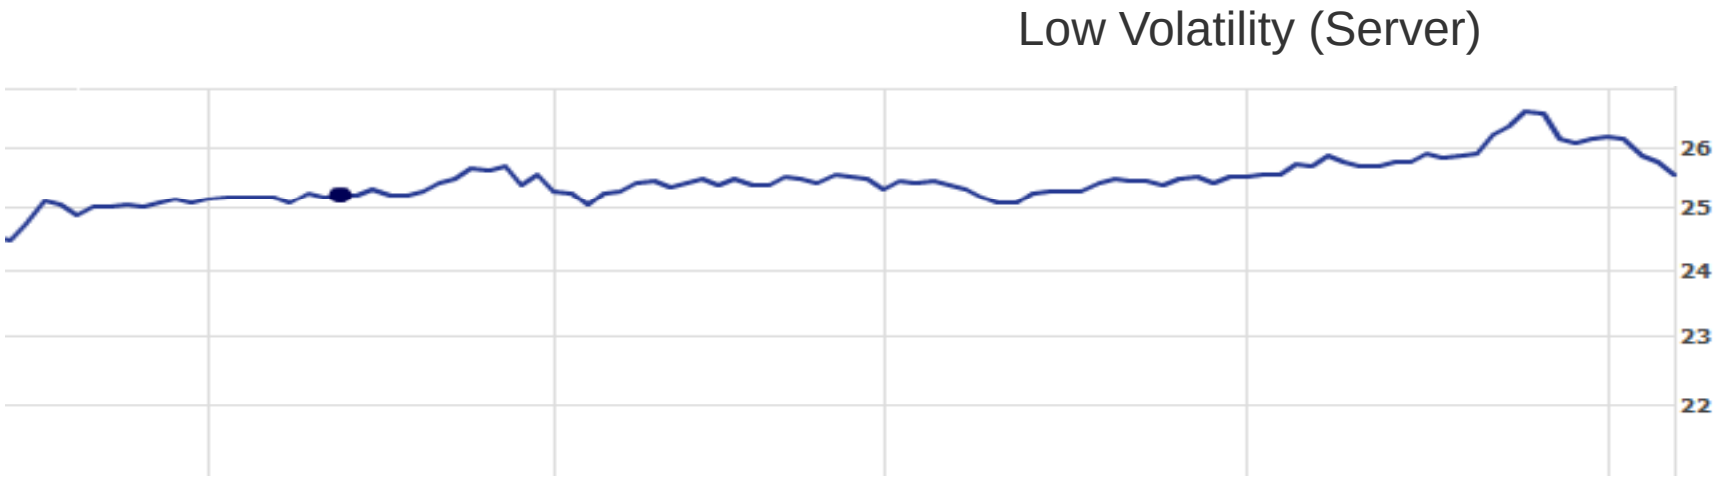

Standard deviation explained

- **Before analysis—obtain the data (which contains peaks and valleys)**
- **Step 1: Calculate the mean (average) of the dataset**

Consider a population consisting of the following eight values:

$$2,\ 4,\ 4,\ 4,\ 5,\ 5,\ 7,\ 9$$

These eight data points have the mean (average) of 5:

$$\frac{2+4+4+4+5+5+7+9}{8} = 5$$

- **Step 2: Compute the difference of each data point from the average, then square the result**

$$\begin{array}{ll} (2-5)^2 = (-3)^2 = 9 & (5-5)^2 = 0^2 = 0 \\ (4-5)^2 = (-1)^2 = 1 & (5-5)^2 = 0^2 = 0 \\ (4-5)^2 = (-1)^2 = 1 & (7-5)^2 = 2^2 = 4 \\ (4-5)^2 = (-1)^2 = 1 & (9-5)^2 = 4^2 = 16 \end{array}$$

- **Step 3: Take the average of those results, and take the square root**

$$\sqrt{\frac{9+1+1+1+0+0+4+16}{8}} = 2$$

Source: http://en.wikipedia.org/wiki/Standard_deviation

Another way to express Std. Dev: Coefficient of Variation (CV)

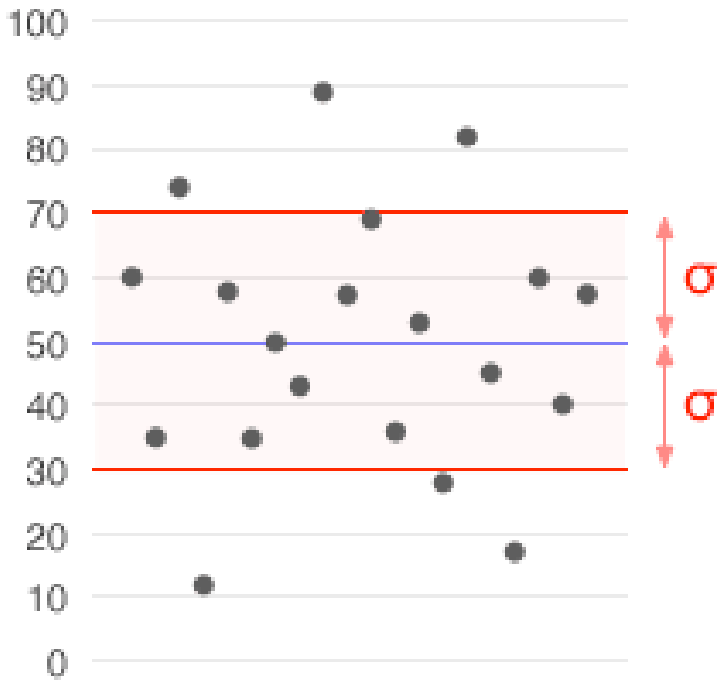

- **Average or Mean (μ) is 50,**
- **Standard deviation (σ)** is measurement of *dispersion* from average
- **Coefficient of Variation (CV)** expresses standard deviation as a percentage of volatility

$$c_v = \frac{\sigma}{|\mu|}.$$

CV = 20 / 50 = .4 or 40%

Example: Bursty vs. Steady IOPS

Time	Bursty Workload	Steady Workload
9AM	330	115
10AM	15	135
11AM	25	110
12PM	330	145
1PM	22	110
2PM	18	135
3PM	27	125
4PM	13	100
5PM	345	150
Total IOPS	1125	1125
Std Dev	149	16
CV	**119%**	**13%**

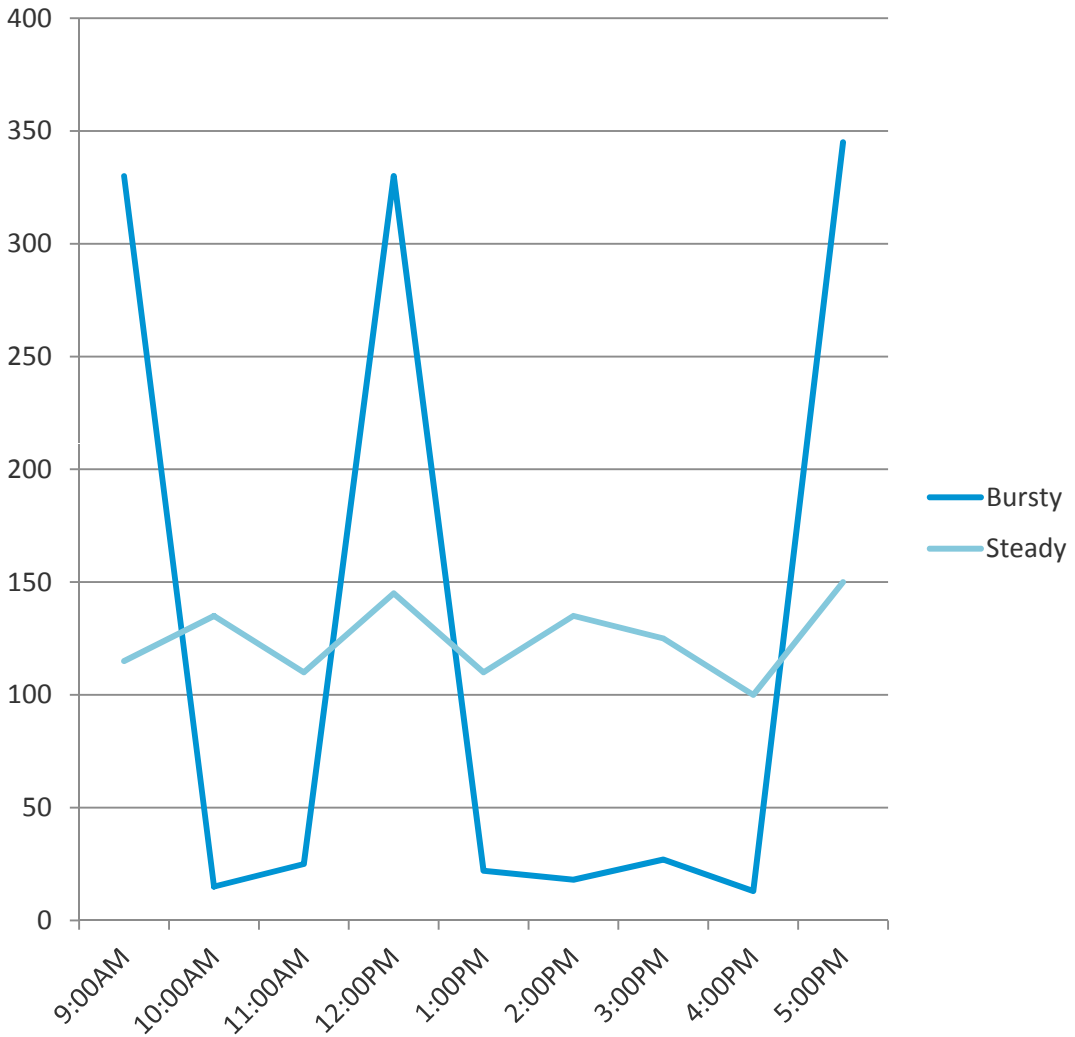

CV – Real World Example

Large VDI Client:

- VDI 1.0 design on underpowered NAS.
- Periodic testing of IOPS availability from VDI desktop showed swings for max IOPS of between 70 and 550 during workday.

CV for VDI: Average Max IOPS was around 240. Std. Dev of 113. CV is **47%**.

Reference CV for Laptop: Average Max IOPS was 305, Std. Dev of 17. CV is 6%

CV & Std. Dev: Smaller is always Better!

Implications of larger CV and Std Dev:

- Larger swings in workload
- Larger physical infrastructure investment
- Increased likelihood of inadequate capacity
- More risk to user experience
- More performance driven user hostility

Action Items

- Find ways to reduce CV will result in less infrastructure needed, less risk to user experience, less cost per user
- Lowering CV results in predictable consumption patterns aligned with virtualization's shared resource assumptions
- Designing VDI with lower cost per unit of CV is also critical.

Chebyshev's Inequality Theorem & SLA Management

The Man: Russian pioneer in probability theory and statistics, number theory.

Definition: A data point is rarely more than a few standard deviations away from the average. No more than $1/k^2$ of the values are more than k standard deviations away from the mean.

So then:

- At least 75% of values within 2 Std. Dev.
- At least 89% of values within 3 Std. Dev.
- At least 94% of values within 4 Std. Dev.
- At least 96% of values within 5 Std. Dev.
- At least 97% of values within 6 Std. Dev.

Pafnuty Chebyshev
Russia: 1821 - 1894

CV Impact: File Transfer Example – 1.25GB file

Summary: By better data flow distribution, a 2Gbit link is sufficient rather than a 10gbit link.

	Pattern 1 Values		Pattern 2 Values
	9991		1475
	1		1050
	1		1200
	1		1025
	1		500
	1		700
	1		820
	1		950
	1		1302
	1		978
Total	**10000**		**10000**
AVG	**1000**		**1000**
Std. Dev.	**2997**		**271**
CV	**300%**		**27%**
2.5 Std: 2.5 * 2997 + 1000	**8493**	2.5 x 271 + 1000	**1678**

VDI Volatility – Logon IOPS

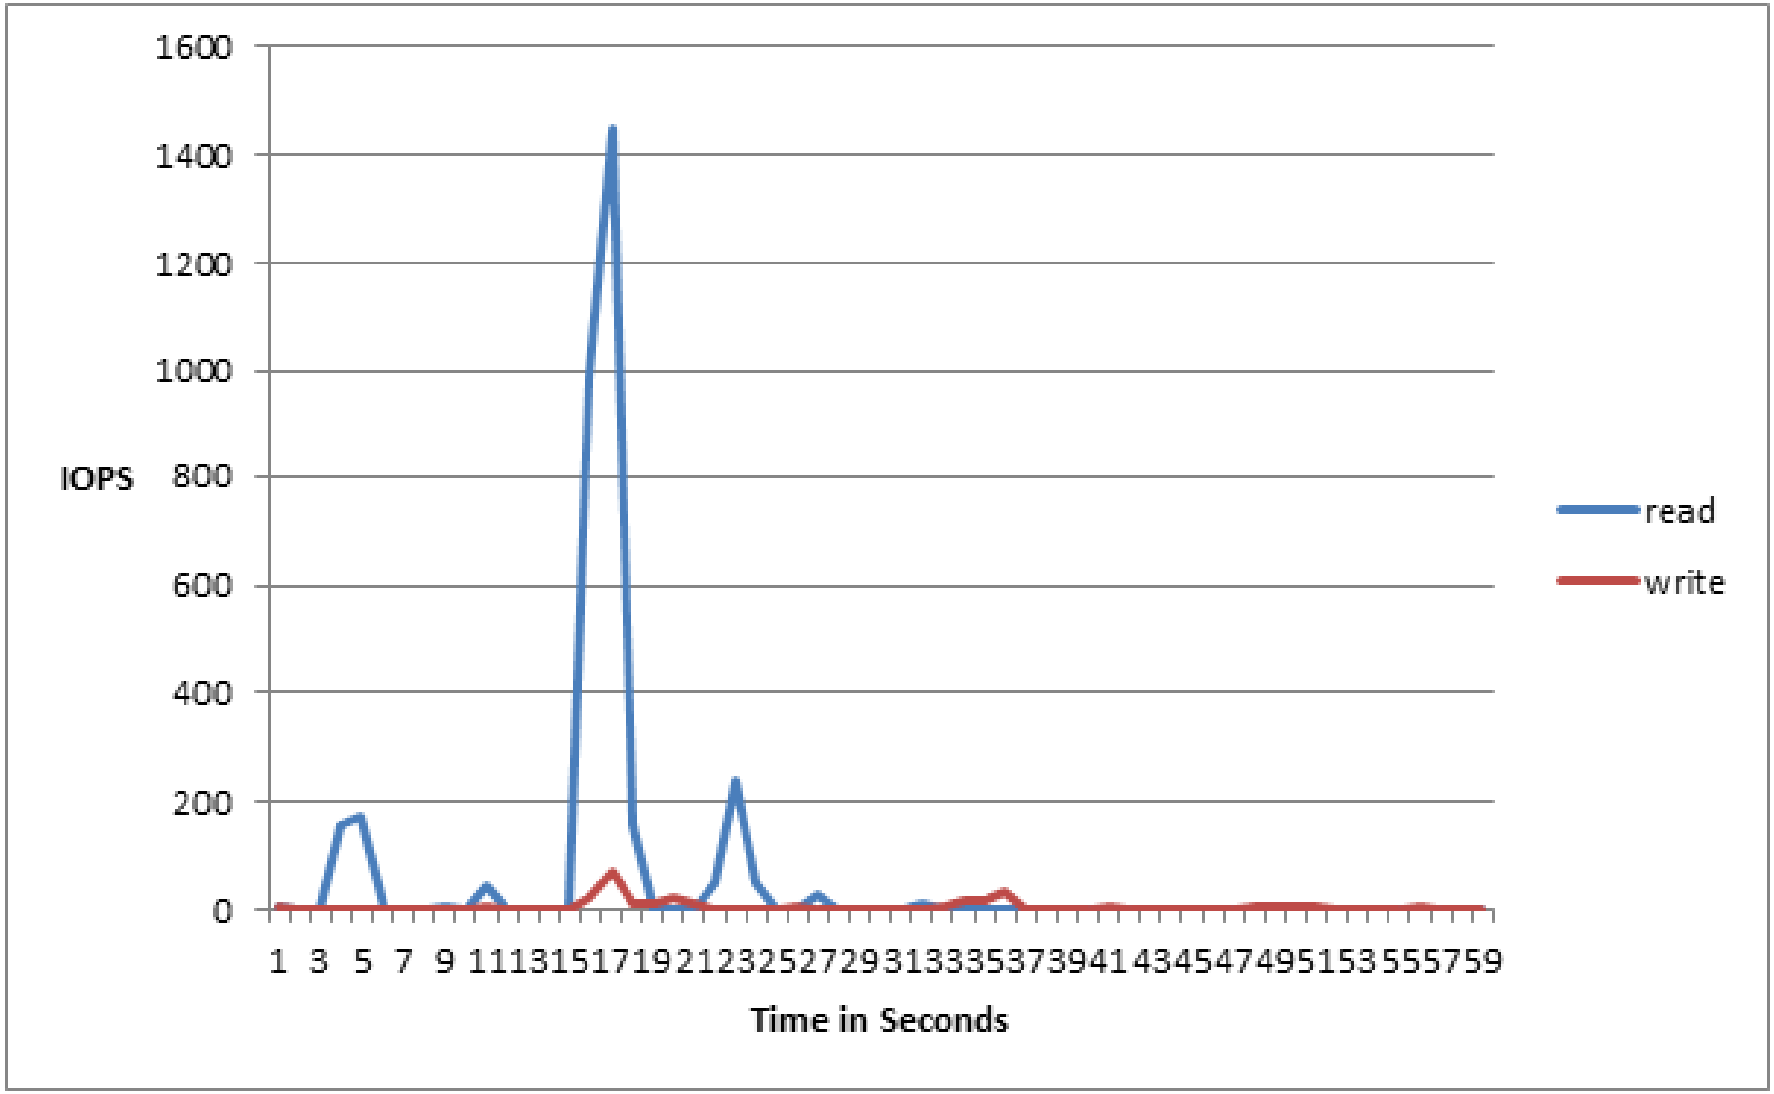

VDI Volatility – Logon IOPS with View 5's Virtual Profiles

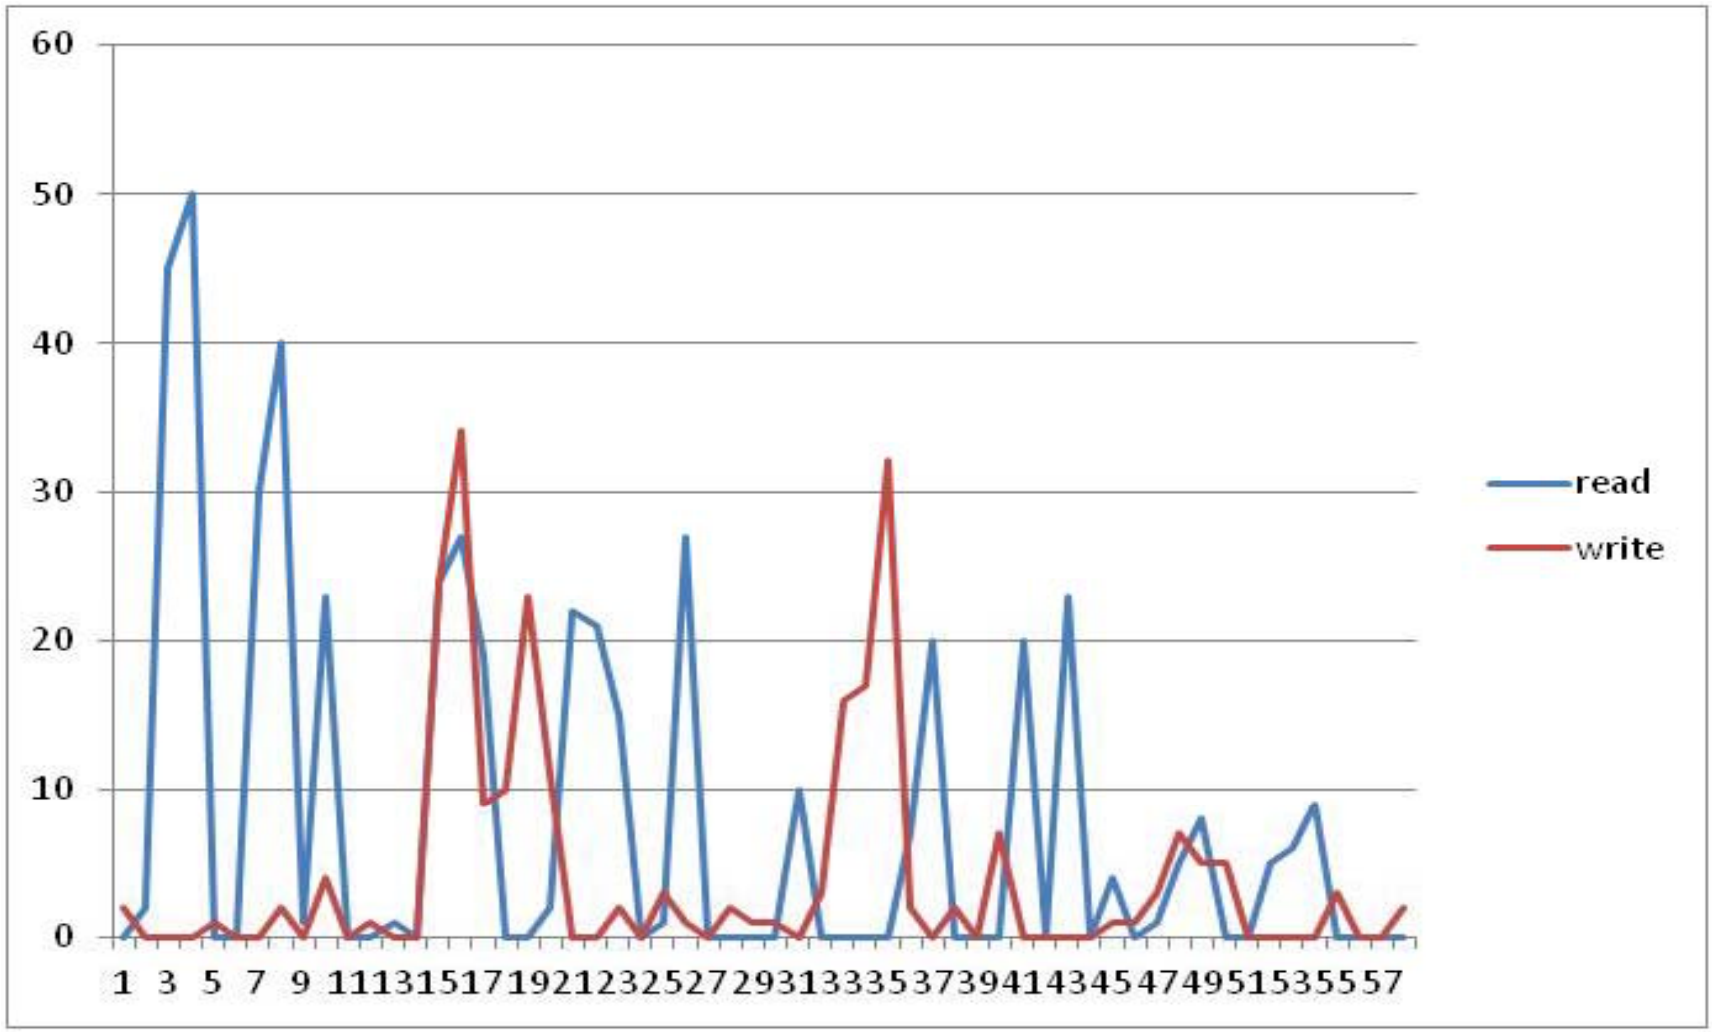

Latency Impact of CV Compression

- **Compressing CV adds Latency:**
  - Failing to accommodate the natural CV of a dataset will add latency. The same workload gets 'stretched' through time, elongating the event duration.
  - CV compression is generally necessary since the OS can often drive infrastructure harder than can be accommodated.
  - For key events, understand the degree of CV compression and latency produced by the design.

How can we estimate impact of forced CV reduction?

- **Solution: Measure surface area above SLA target range, subtract area below target and convert remainder to rectangle. X axis portion of rectangle will determine minimum event extension.**
- **This is an *estimate* and a best case – latency could be worse.**

Estimating Latency - Logon Example

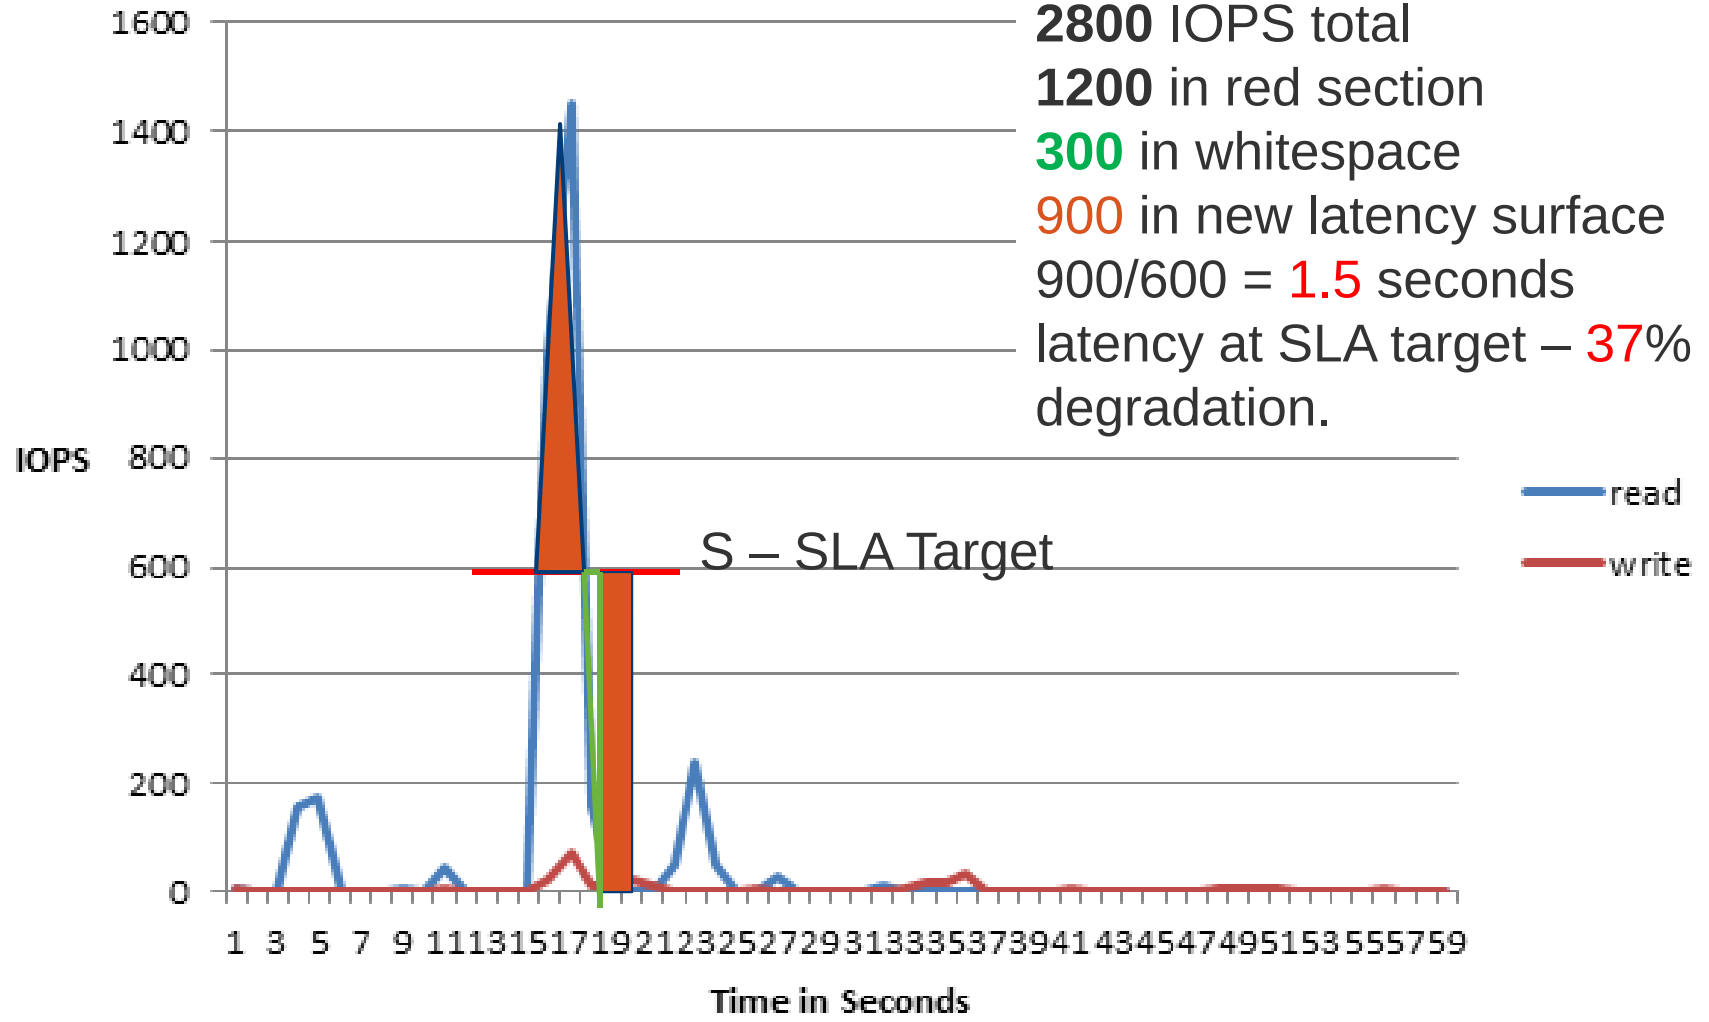

2800 IOPS total
1200 in red section
300 in whitespace
900 in new latency surface
900/600 = 1.5 seconds latency at SLA target – 37% degradation.

Jim Moyle's Logon Chart

How does an Architect use this Math?

- **Identify 'event risks' that raise the standard dev of the infrastructure consumption pattern. Determine SLA target needed.**
  - Boot, login, or logoff storms
  - AV scans or updates
  - Recompose, refresh, software updates
- **Reduce Average & standard deviation—lower the spikes**
  - OS image optimization
  - AV offload or randomization, disable defrag etc.
  - Use streaming profiles to reduce logon burst activity
- **Shift peak load to specialty mechanisms, reducing risk and cost.**
  - VSA (virtual storage appliance) storage designs – accommodate burst I/O at a low cost by leveraging local ESX RAM and the vSwitch as the IO bus.
  - Tiered storage – use the best tier of storage for the I/O pattern – made easy by VMware View's tiered storage feature.
  - Find cheaper ways of accommodating peak loads .

ILIO-VSA w/VMware View Benefits

VMware
View's Tiered Storage:
Tiered storage moves burst workloads components to a separate data path.

User Profiles: Use VMware's Virtual Profile solution, now part of View 5. Streaming profiles reduce synchronized congestion by avoiding burst IO. Can increase user density up to 40%.

Virtualization aware AV: Smarter consolidate of AV activity – removes repetitious scanning at host level.

Atlantis ILIO
•ILIO VSA technology can service I/O inside the ESX Host – limiting interactions between ESX hosts. Host level burst isolation helps with stable performance.
•Limits average I/O load leaving ESX.
•Offers predictably low-latency IOPS – often better than SSD.
•Can accommodate high burst IOPS per user.

Questions?

Thank You.

Q & A ?

* Thanks to Jim Moyle for sharing graphs & John Dodge for insightful feedback.