IntelliMagic

SHARE
Technology • Connections • Results

Best Practices for Mainframe I/O SLA and Efficiency Optimization

Lee LaFrese
IntelliMagic

Thursday, March 15, 2012
Session 10461

IntelliMagic

About IntelliMagic

- A world leader in Storage Performance Management software solutions
- Developing SPM solutions since 1991
- Private, no debt
- Headquarters in Leiden, NL   US office in Dallas, TX International Partners

Storage Intelligence
www.IntelliMagic.net

IntelliMagic

About the Speaker

Lee LaFrese

- Recently joined IntelliMagic as a Senior Performance Consultant.
- Worked at IBM for 32+ years and was the technical lead in product development for Enterprise Disk Storage Performance.
- Has written over 20 whitepapers and made numerous technical presentations on a wide variety of performance topics spanning both mainframe and distributed storage.

Contents

IntelliMagic

Contents

1. The Shift in I/O Queuing
2. Best Practices for z/OS Storage Performance Management (SPM)
3. Case Study: Storage Efficiency
4. Case Study: I/O vs. CPU Case Study

IntelliMagic

The Evolution of I/O Queuing on z/OS

Direct
Attached
Storage
Device

The Evolution of I/O Queuing on z/OS

Multiple LPARS connected through switches and directors to disk subsystem(s)

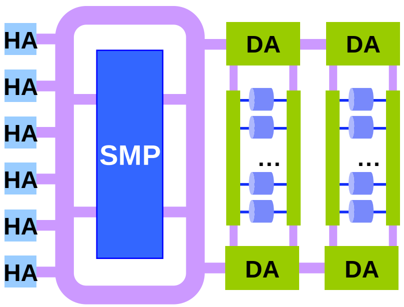

The Evolution of I/O Queuing on z/OS

I/O Queuing is rarely on the host side anymore. Today's bottlenecks are now **inside** of the disk subsystem

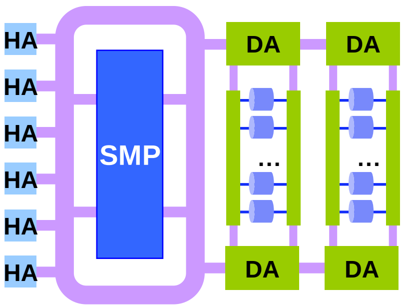

I/O Queuing has shifted due to features like:

- FICON vs. ESCON
- PAV
- Multiple Allegiance

The Evolution of I/O Queuing on z/OS

I/O Queuing is rarely on the host side anymore. Today's bottlenecks are now **inside** of the disk subsystem

"Visibility Gap"

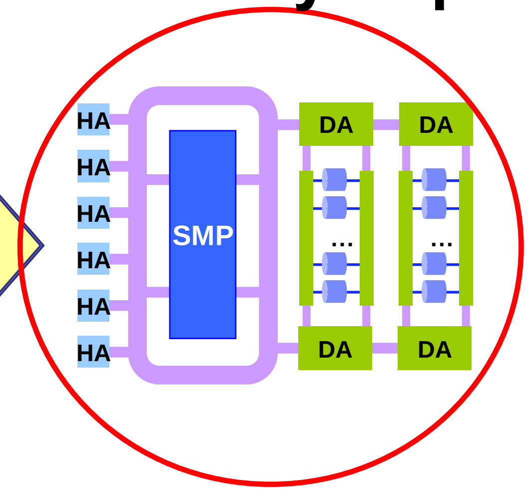

I/O Queuing has shifted due to features like:

- FICON vs. ESCON
- PAV
- Multiple Allegiance

But I/O Reporting is still:

- LPAR (host) centric, not DASD centric
- About symptoms, not root causes

IntelliMagic

Storage Intelligence
www.IntelliMagic.net

Consequences of the Visibility Gap

1. **Performance (SLA)**
   - Higher risk of service level disruptions
   - Problems discovered ***after*** SLA violations
   - Root cause analysis requires hardware vendor

2. **Efficiency**
   - Only a few resources in play for peak workloads
   - Boxes replaced prematurely
   - Over-dependence on CPU vs. faster I/O

IntelliMagic

Performance Consequences of the Visibility Gap

- Current reporting alert about symptoms, not root causes
- IT staff often learns about performance issues from end-users (reactive) rather than by monitoring the health
- Negative surprises occur quicker – the response time “knee of the curve” is much sharper than it used to be

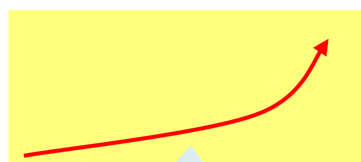

In the days of host-centric bottlenecks, you could more easily see how close to the edge you were as your response times grew in a linear fashion

Nowadays, the curve is more sudden - it spikes up when any component in the storage system reaches a critical level of utilization

IntelliMagic

Most Common Root Cause of Delays:

In today's architectures, from two places *inside* the storage system (hard to see):

The "Front End"

When **H**ost **A**dapters become over utilized during peak periods and can not keep up with the amount of work requested

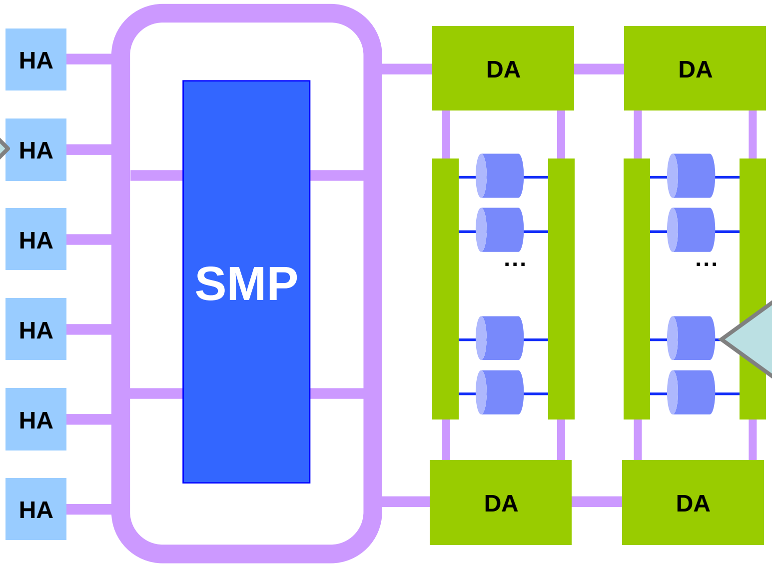

The "Back End"

When the disk devices themselves cannot keep up with requests to stage data into cache or to de-stage data from cache onto the drive

The **Front End**: Host Adapters - Workload

The workloads on this DSS are almost all using only 8 of the 16 ports, and it shows up in high response times in the chart at right

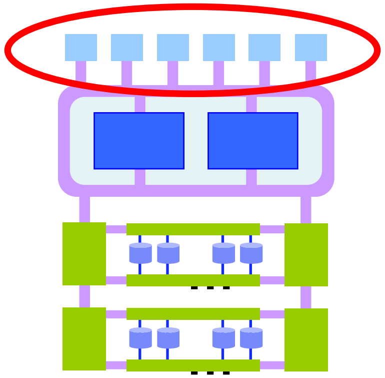

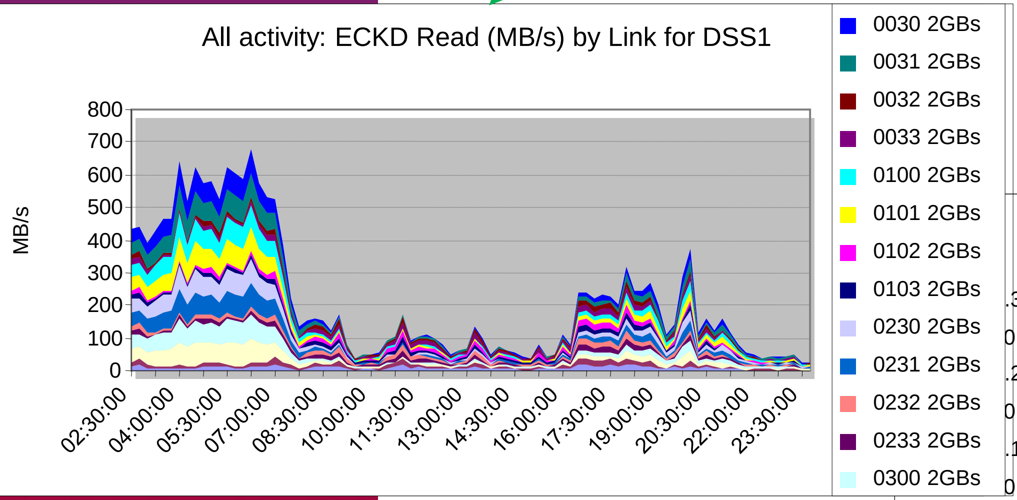

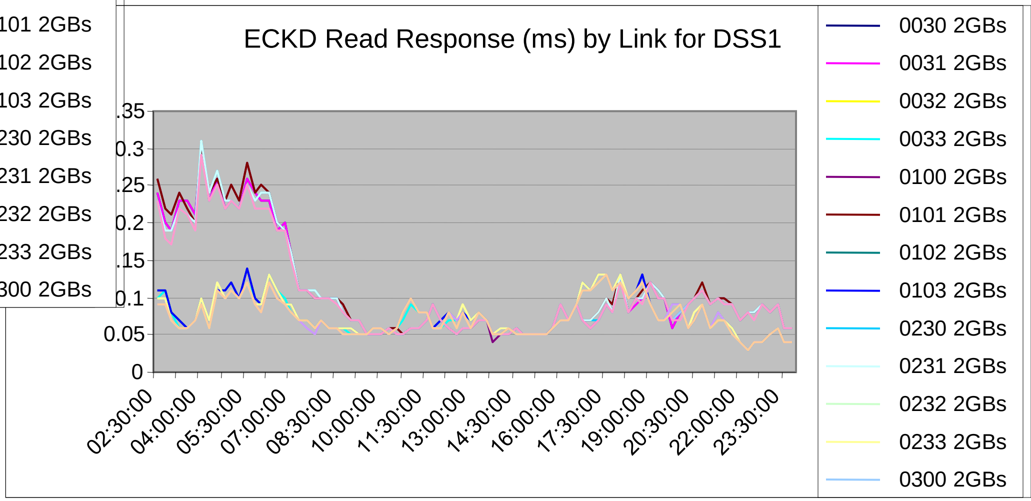

The **Back End**: Disk capabilities

Drive Density					
Space in GB:	73	146	300	450	600
Perf. in RPM:	15k	15k	15k	15k	15k
Safe IO's / drive:	150	150	150	150	150
Max I/O per GB:	2.05	1.03	0.5	0.33	0.25

The "Access Density Gap" where space grows but performance remains constant

Space has been the greatest constraint

Now Performance is the greatest constraint

This is why SSD's are more and more necessary – because back-end access density per drive would be too high for many workloads on large drives

The **Back End**: Disk capabilities

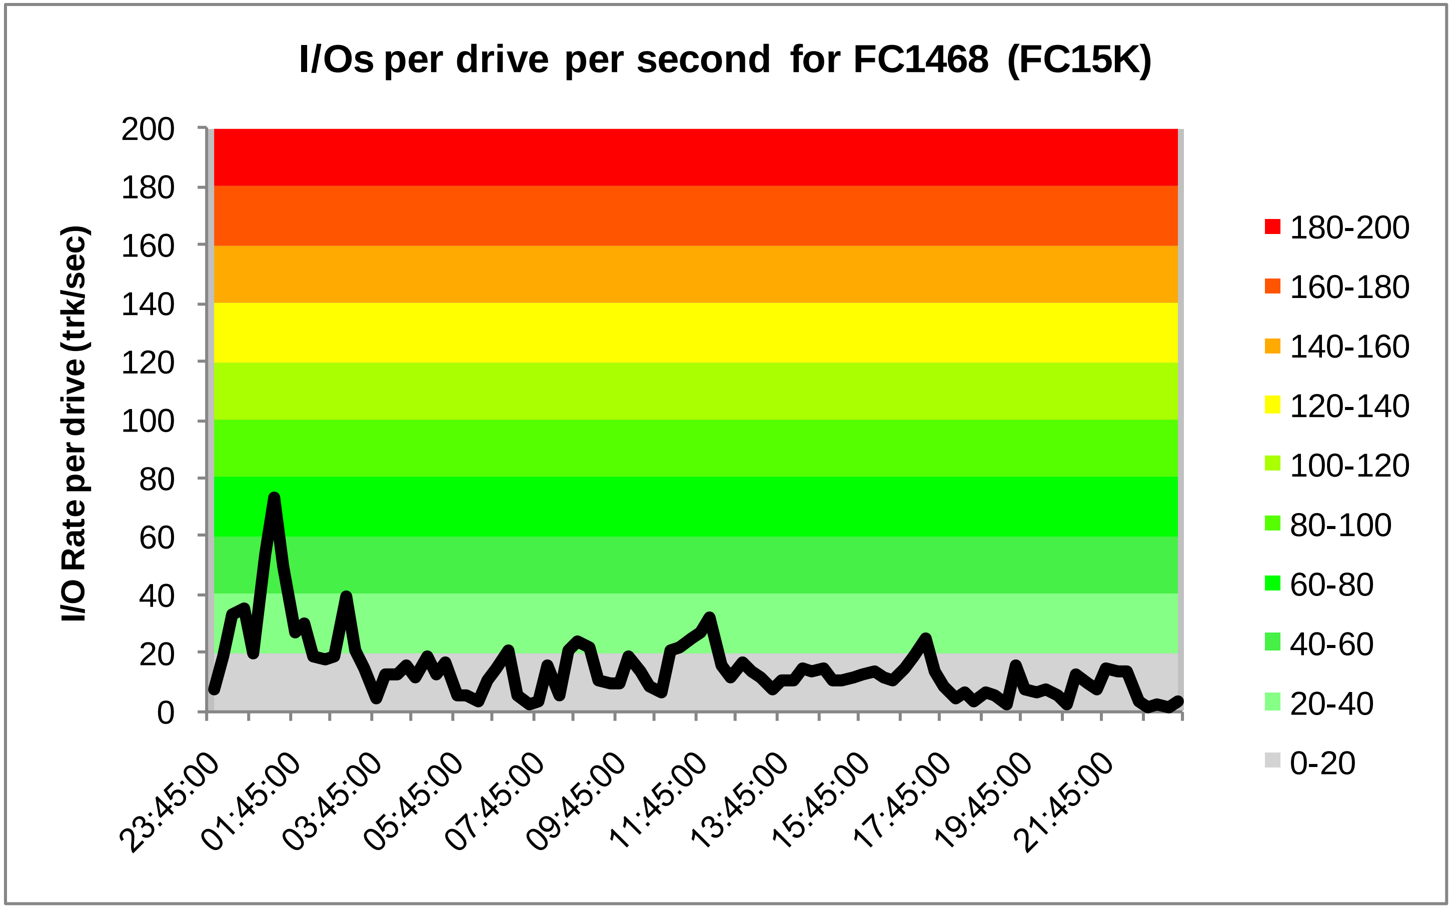

Storage Intelligence
www.IntelliMagic.net

Contents

1. The Shift in I/O Queuing
2. Best Practices for z/OS Storage Performance Management (SPM)
3. Case Study: Storage Efficiency
4. Case Study: I/O vs. CPU Case Study

IntelliMagic

SPM Best Practice #1
The Right Metrics

Use the right storage centric metrics for today's architectures.

"If you can't measure it, you can't manage it!"

Familiar Territory

- Pending time
- Disconnect time
- Connect time

These metrics have been available "forever"

But these metrics :

1. Have taken on a new meaning
2. Are no longer enough to assess disk subsystem health

New Metrics

- Host Adapter busy
- Back-end Array Group busy
- Connect time elongation
- FICON effective data rate

RMF does not report any of these, but they are critical to understanding DSS health and must be computed!

SPM Best Practice #1
The Right Metrics

Sample Host Adapter Utilization chart:

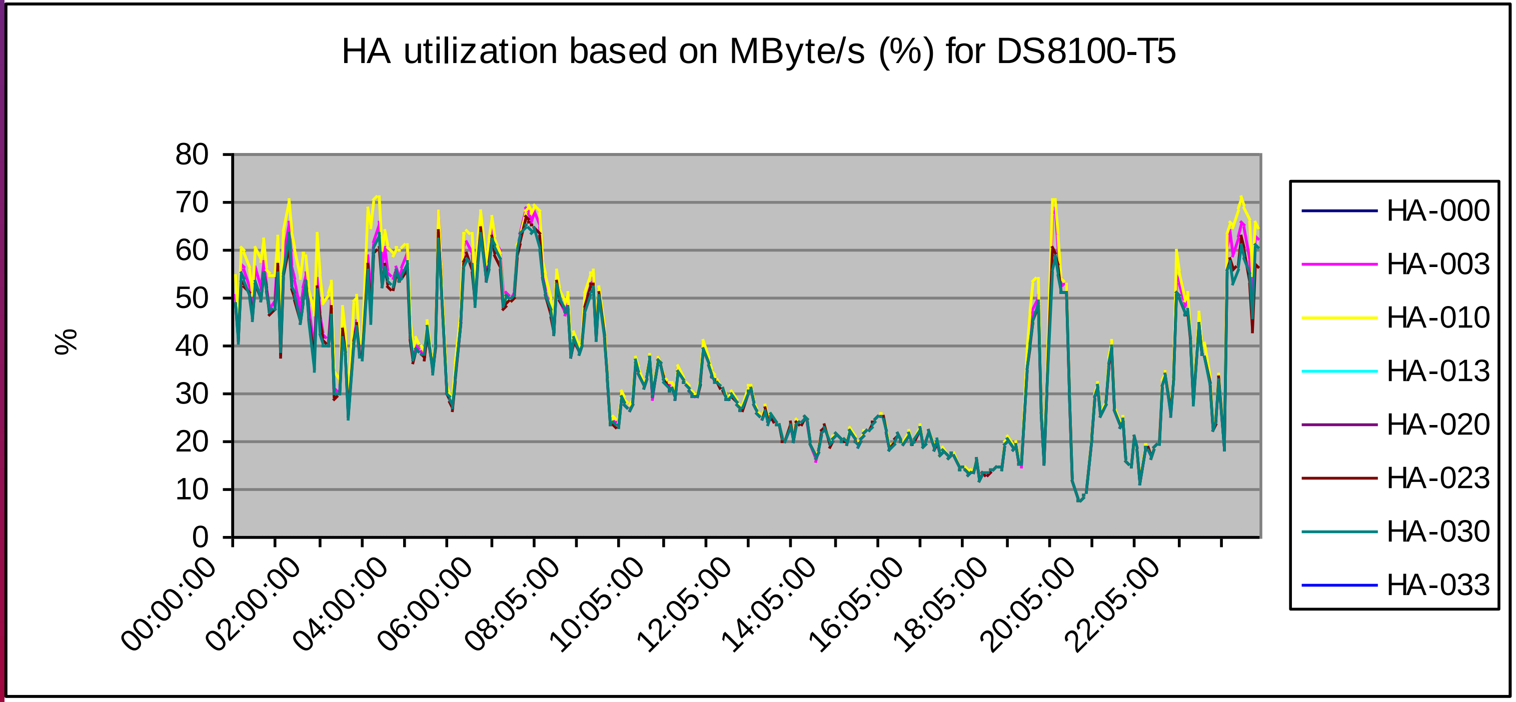

In general, utilization levels over 60% will start to introduce delays, over 80% will be unacceptable

SPM Best Practice #2
Contextual Interpretation

Interpret the right metrics in the context in which they occur:

- based on the host environment
- based on the capabilities of the storage hardware

For example, is disconnect time of 2ms good or bad?

- good for DB2 workload with remote mirroring active
- bad for z/OS system pack
- May be good for older generation hdw, but bad for new

SPM Best Practice #2
Contextual Interpretation

Example of viewing metrics with context sensitive thresholds:

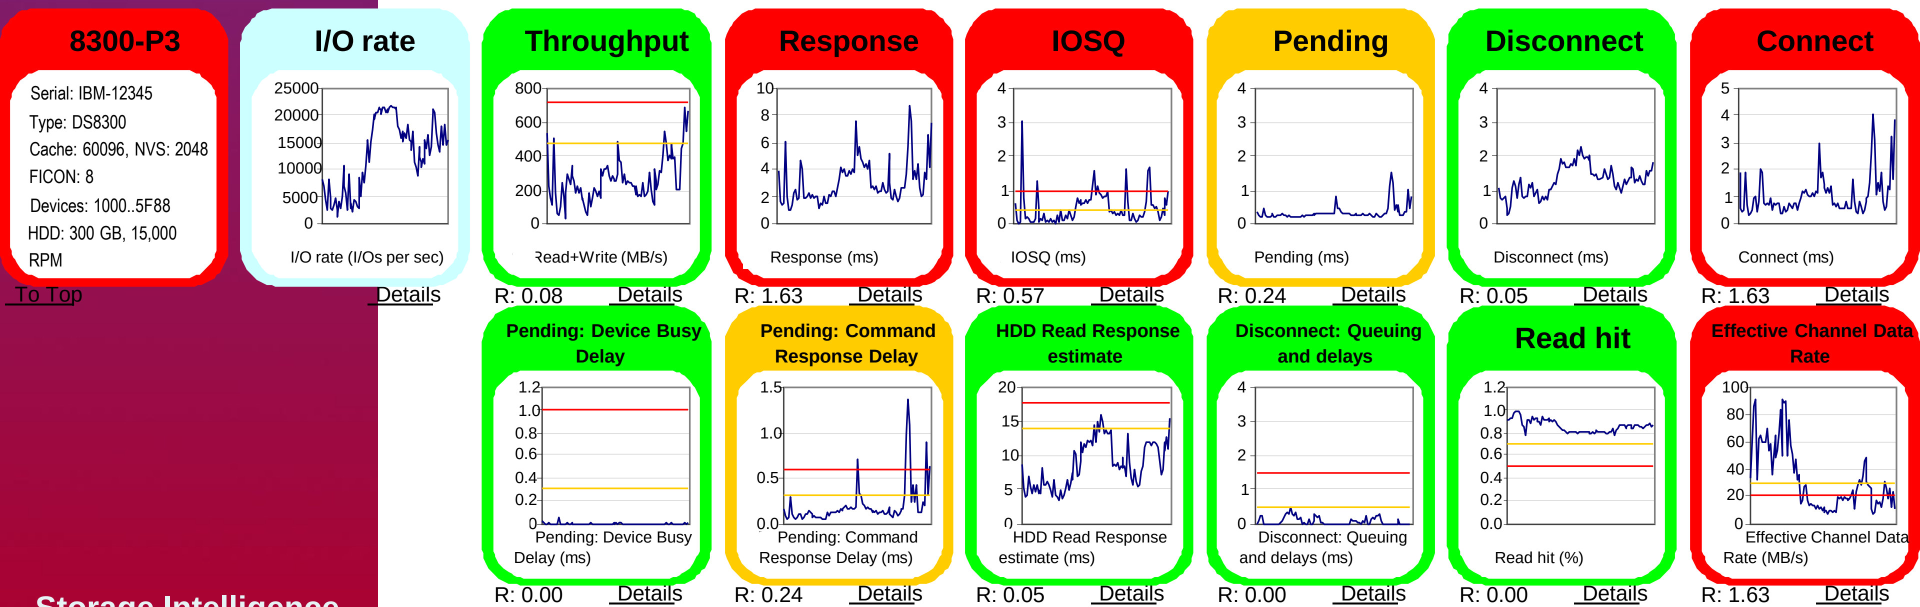

SPM Best Practice #3
Proactive Evasion

See and address I/O performance problems ***before*** they disrupt production users.

Status quo today is to react after pain is reported.

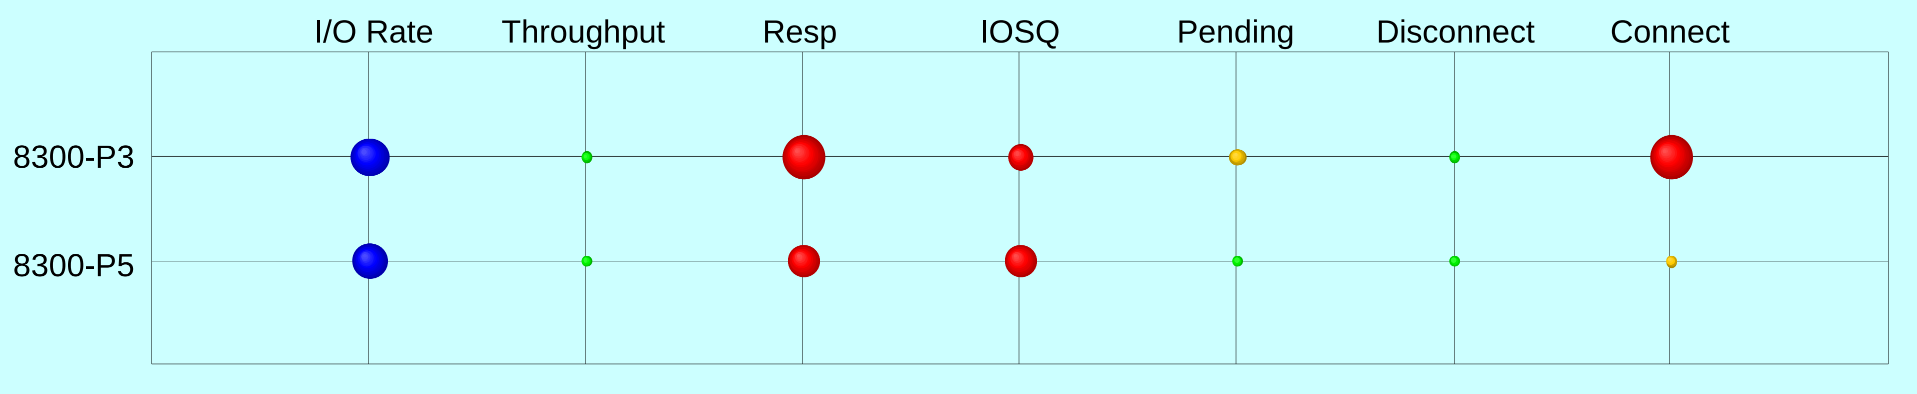

IntelliMagic

SPM Best Practice #4
Simplified Diagnosis

Simplify investigations into root causes of performance degradation.

Immediate drill down to root causes at the deepest levels:

- without more data collection & manipulation
- Without having to depend on the vendor or other specialists

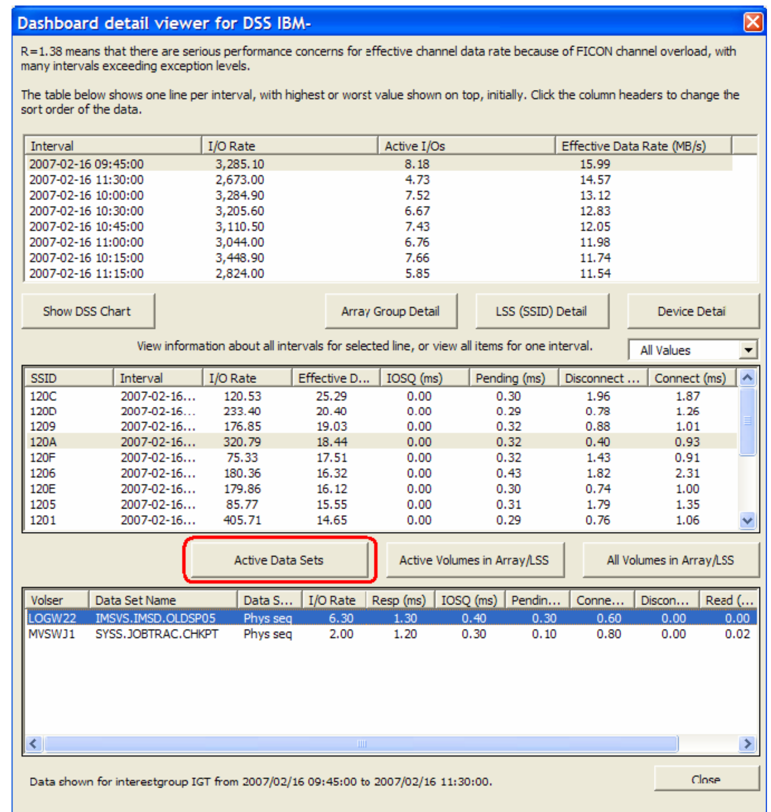

SPM Best Practice #5
Imbalance Identification

Address imbalance rather than adding hardware.

In this example, it is easy to see imbalance on back-end drives

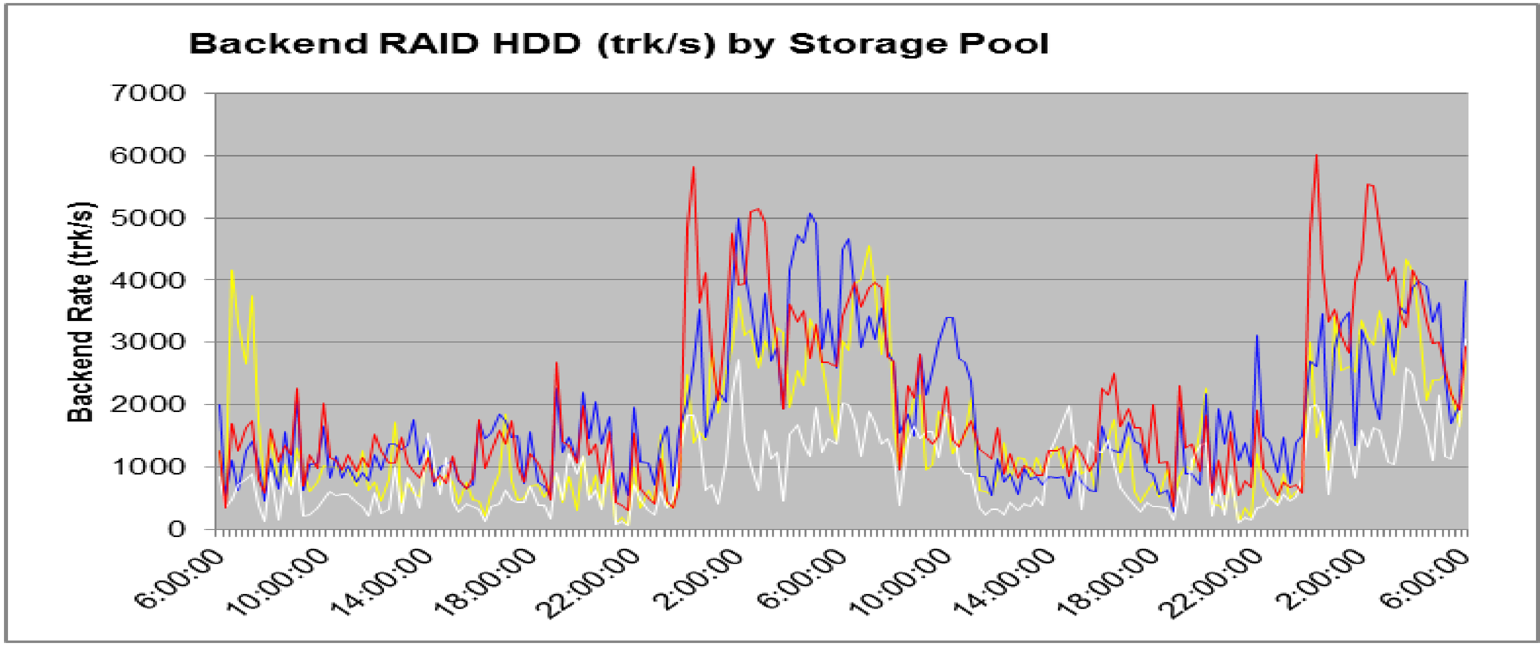

IntelliMagic

SPM Best Practice #6
Smart Placement

Calculate optimal volume placement on hardware resources in data migrations.

- Creates better hardware utilization, saving money
- Enables pain free storage tiers
- Avoid risk of creating new hot spots in data moves

IntelliMagic

SPM Best Practice #7
Analyze Trends

Understand workload performance characteristics over time in multiple dimensions with a historical database for the metrics.

- Allow ad-hoc reporting on historical data
- See how application I/O has been changing over time
- Quickly compare present and historical data to diagnose where differences are occurring

IntelliMagic

SPM Best Practice #8
Predict Performance

Model the performance impact of changes.

- What spikes in workloads can your current hardware handle?
- What hardware configuration options provide the best

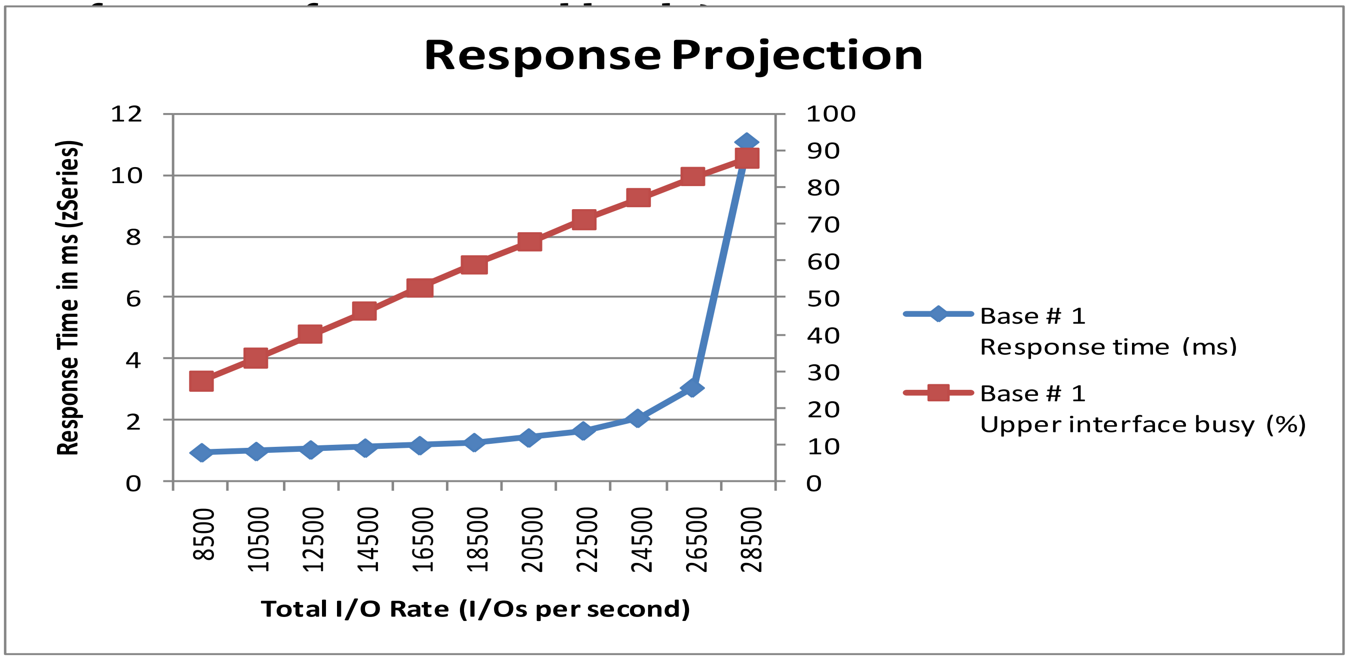

Storage Intelligence
www.IntelliMagic.net

Contents

1. The Shift in I/O Queuing
2. Best Practices for z/OS Storage Performance Management (SPM)
3. Case Study: Storage Efficiency
4. Case Study: I/O vs. CPU Case Study

Storage Efficiency Objective:
Maximize *Safe* Growth Potential

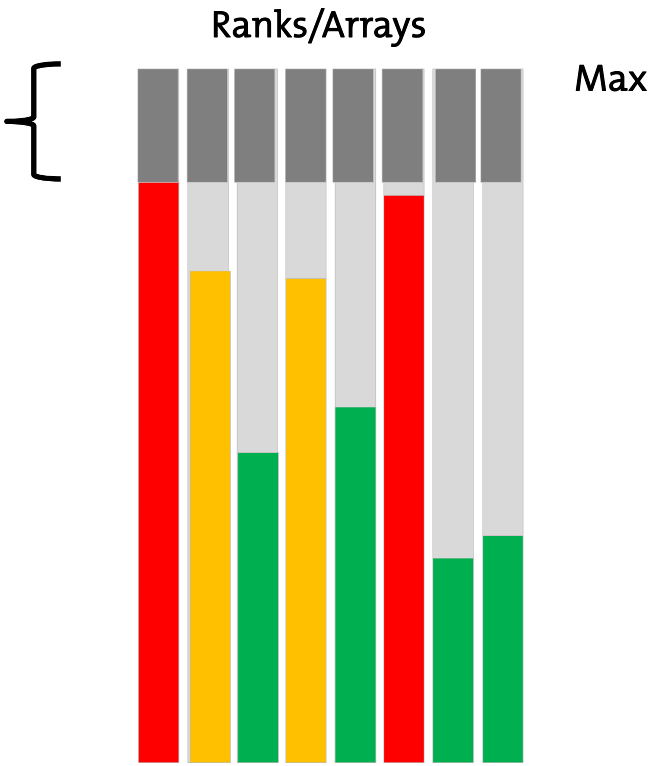

On which arrays would you like your volumes to be?

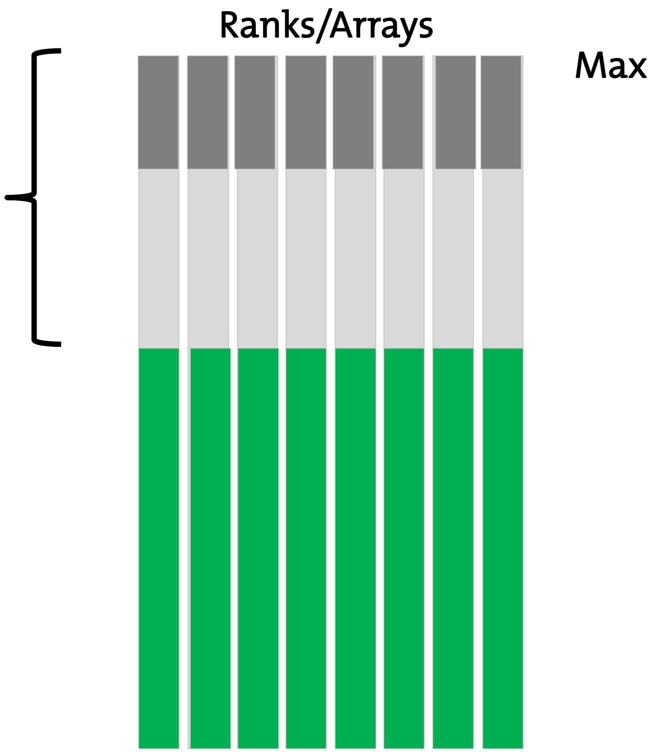

Anywhere will do!

Storage Efficiency

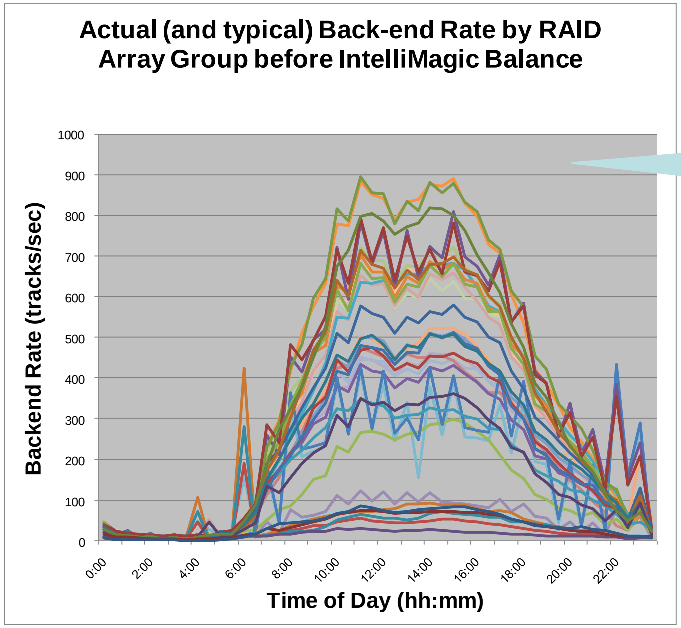

Time to replace this box?

Maximum rate:
900 tracks / second / array group (8 disks)

Storage Efficiency

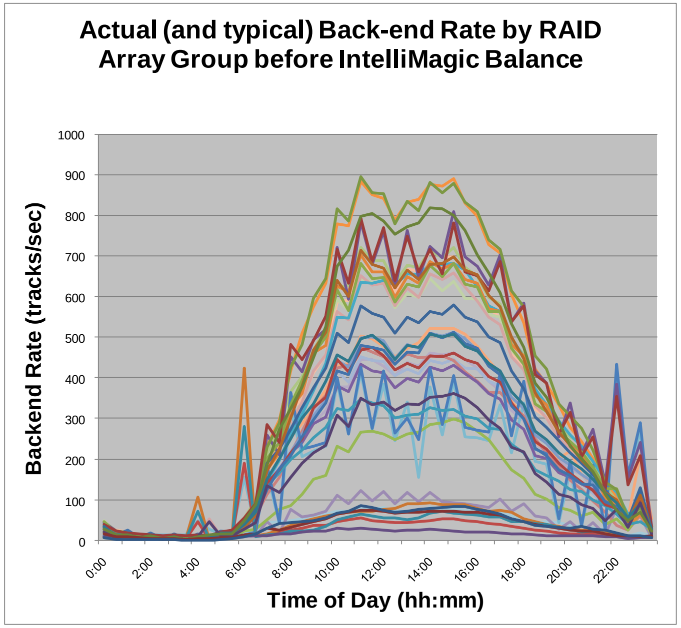

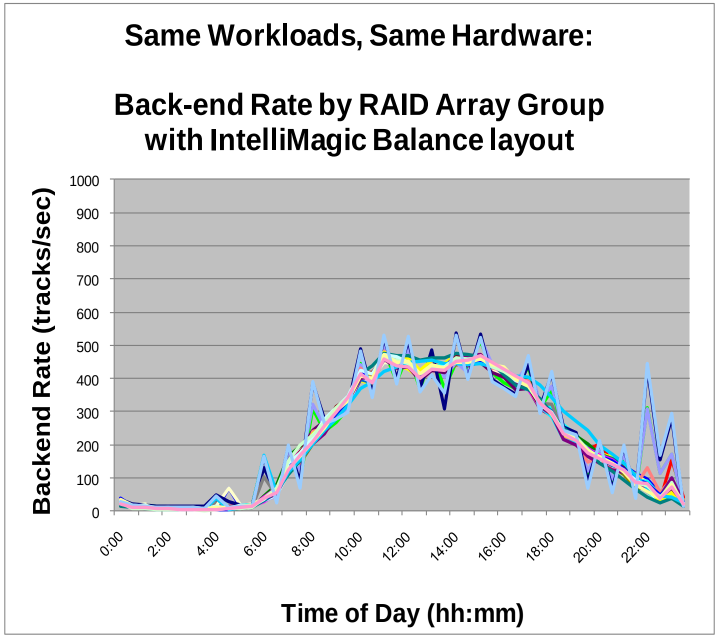

Old Max Rate:
900 tracks per second
per array group

New Max rate:
550 tracks per second
per array group

About a 40% reduction in peaks!

IntelliMagic

Storage Intelligence
www.IntelliMagic.net

Contents

1. The Shift in I/O Queuing
2. Best Practices for z/OS Storage Performance Management (SPM)
3. Case Study: Storage Efficiency
4. Case Study: I/O vs. CPU Case Study

Evaluating I/O vs. CPU Optimization
The Objective:

1. Can I defer CPU upgrade plans if disk I/O performance improves?
2. Can I reduce batch run times significantly storage tuning or upgrades?
3. Will our planned storage hardware upgrade allow me to defer a CPU upgrade?

Evaluating I/O vs. CPU Optimization

The Process:

1. Identify key workloads that need additional resources to meet SLA windows
2. Examine their relative disk I/O and CPU contributions
3. If I/O contribution is significant, quantify the gap between current I/O service times and best-in-class I/O
4. Apply the I/O gap differential in order to estimate its impact on the execution of high-value jobs
5. Contrast the cost of closing the I/O gap with obtaining similar performance benefit with CPU

IntelliMagic

Evaluating I/O vs. CPU Optimization
Example Case Study:

Critical batch workload needs a performance gain:

1. Identify a key job to assess potential impact
2. Examine the current DASD response times for the storage hardware this job is running on.
3. Model performance improvement I/O
4. Contrast with a CPU approach

IntelliMagic

IntelliMagic

The Critical Batch service class has a significant "I/O to CPU bound ratio" with 58.9% of time on I/O. Let's use it to illustrate the I/O optimization potential

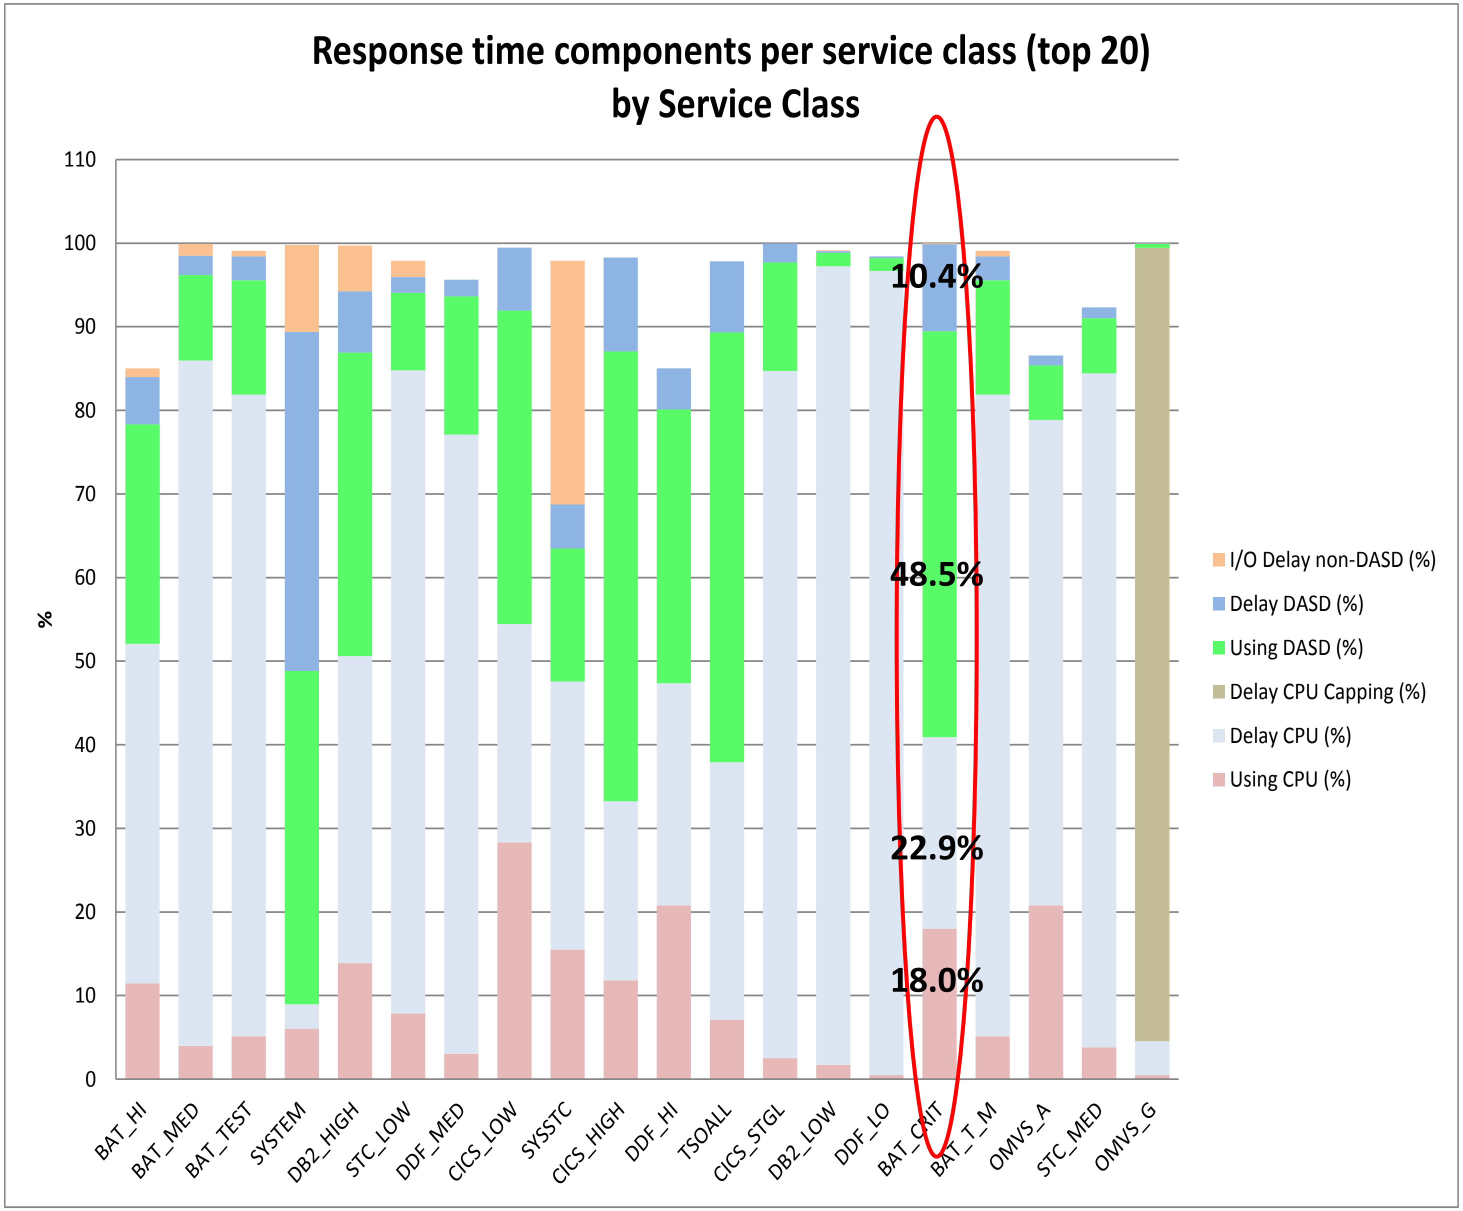

Storage Intelligence
www.IntelliMagic.net

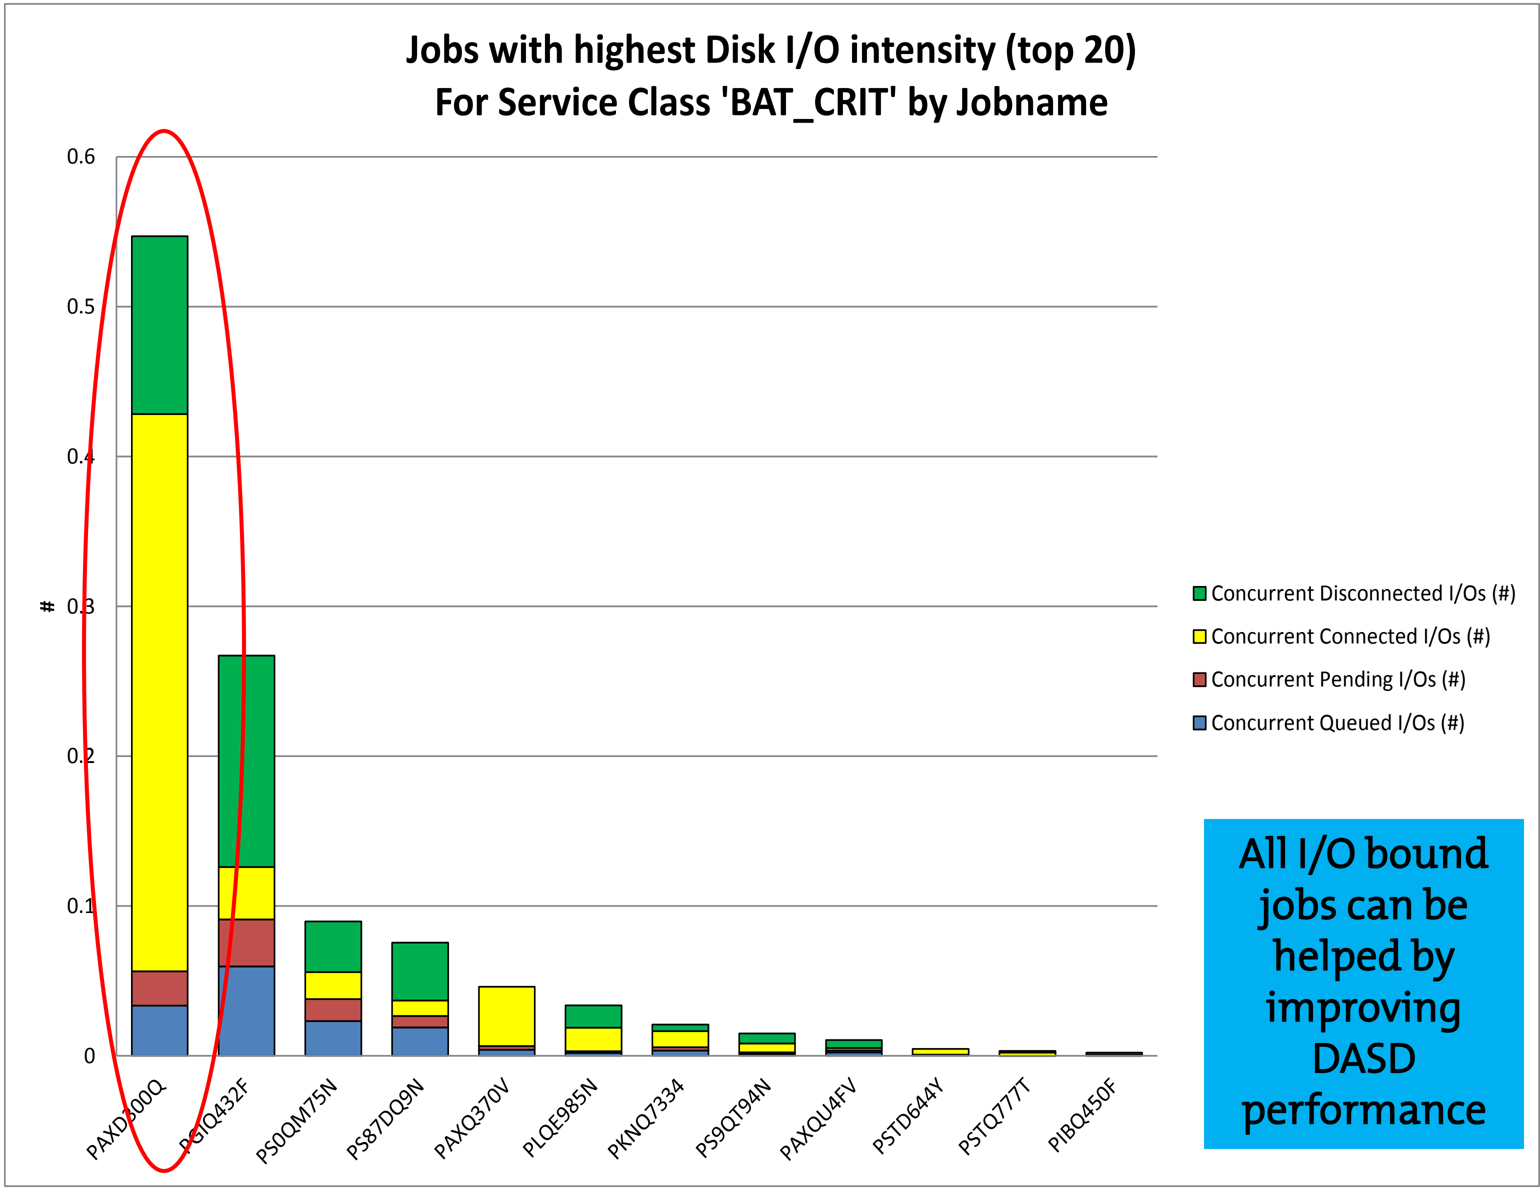

There is one job in this service class driving the overall I/O times in the service class far more than the other jobs – lets illustrate the impact on that one.

All I/O bound jobs can be helped by improving DASD performance

IntelliMagic

Storage Intelligence
www.IntelliMagic.net

IntelliMagic

The job currently runs for almost two hours.

What I/O service times is it currently getting?

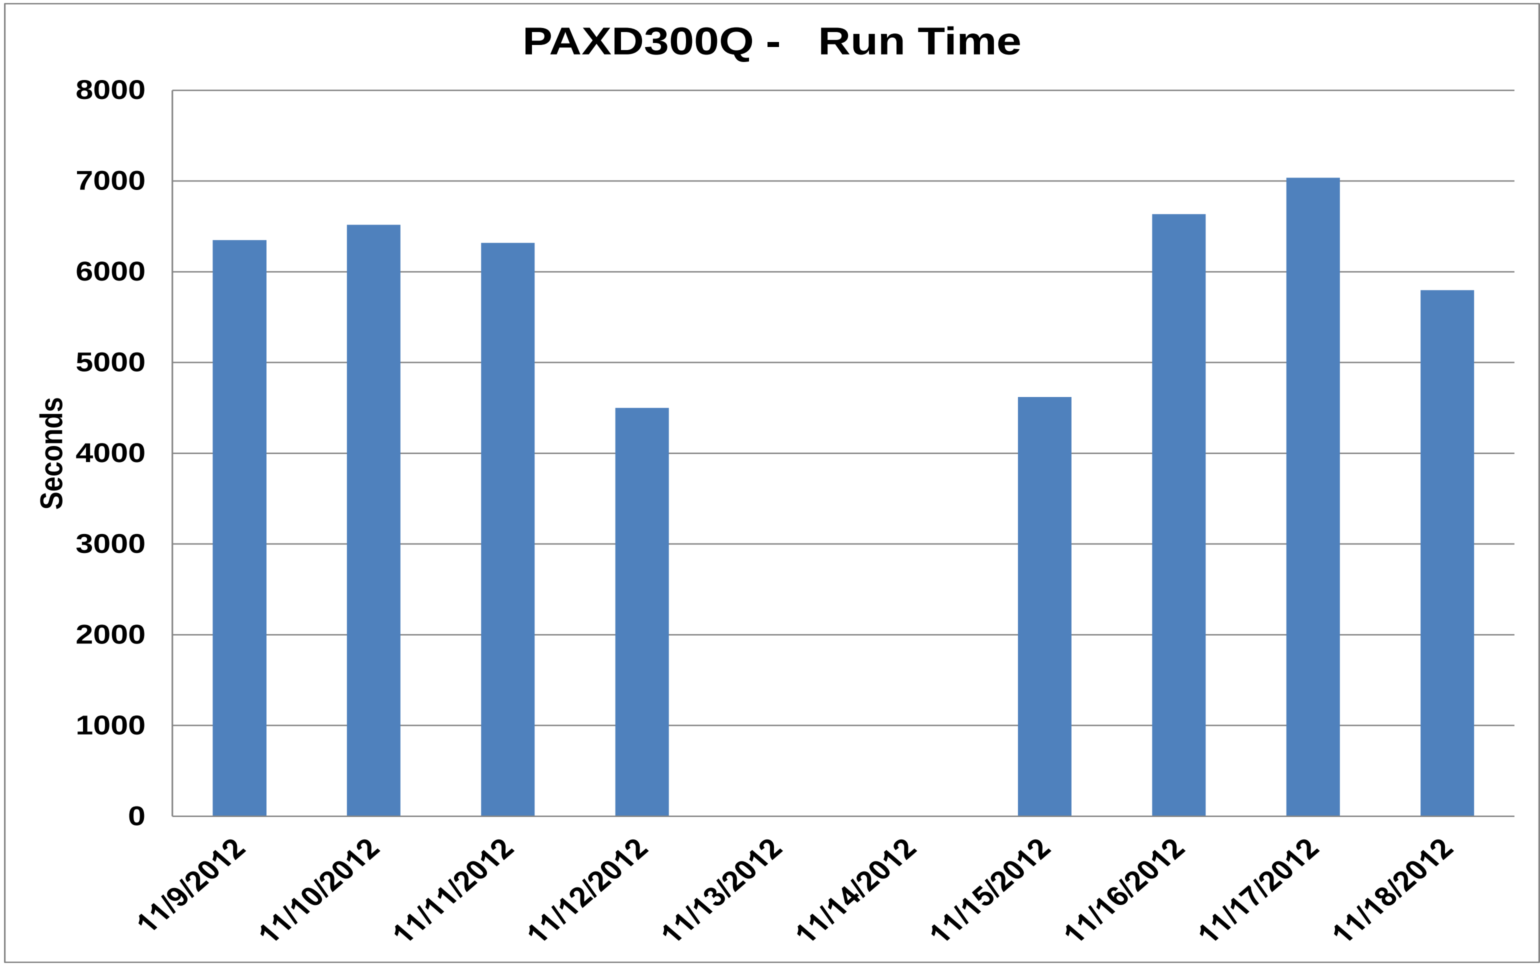

Response time components

For Jobname 'PAXD300Q', for Service Class 'BAT_CRIT'

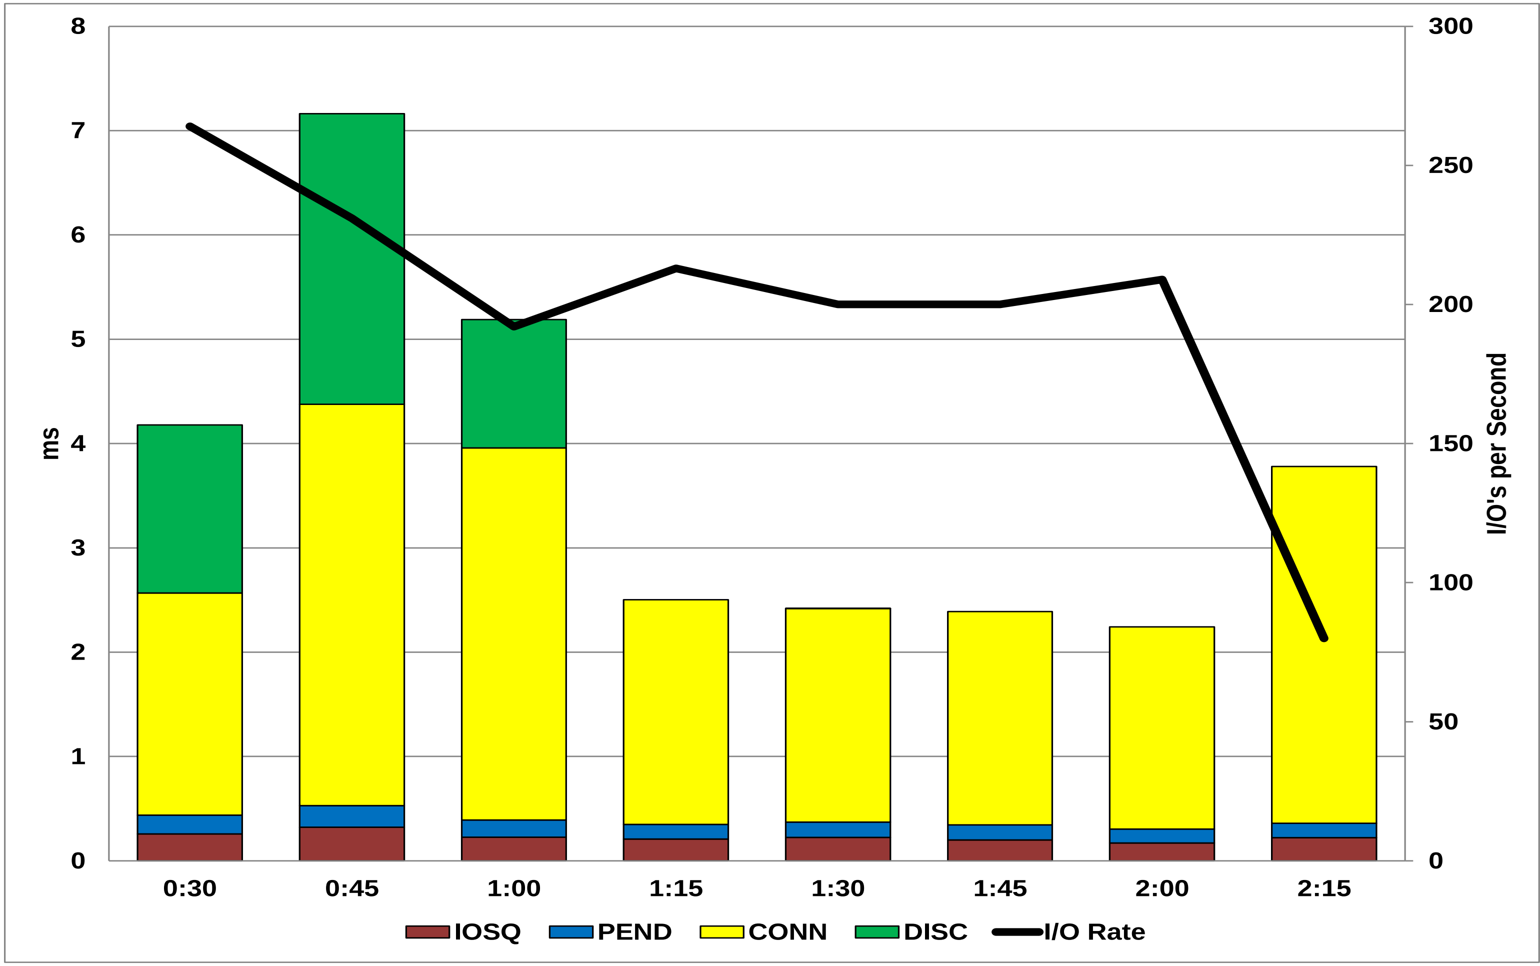

These times are clearly not best-in-class.

We can accurately model what I/O is possible for these specific workloads.

Storage Intelligence
www.IntelliMagic.net

IntelliMagic

Job Component Summary

CPU Using	CPU Delay	DASD Using and Delay	Total Time
1196 seconds	1523 seconds	3915 seconds	6634 seconds

59% of job execution time related to DASD!

Which components can be reduced?

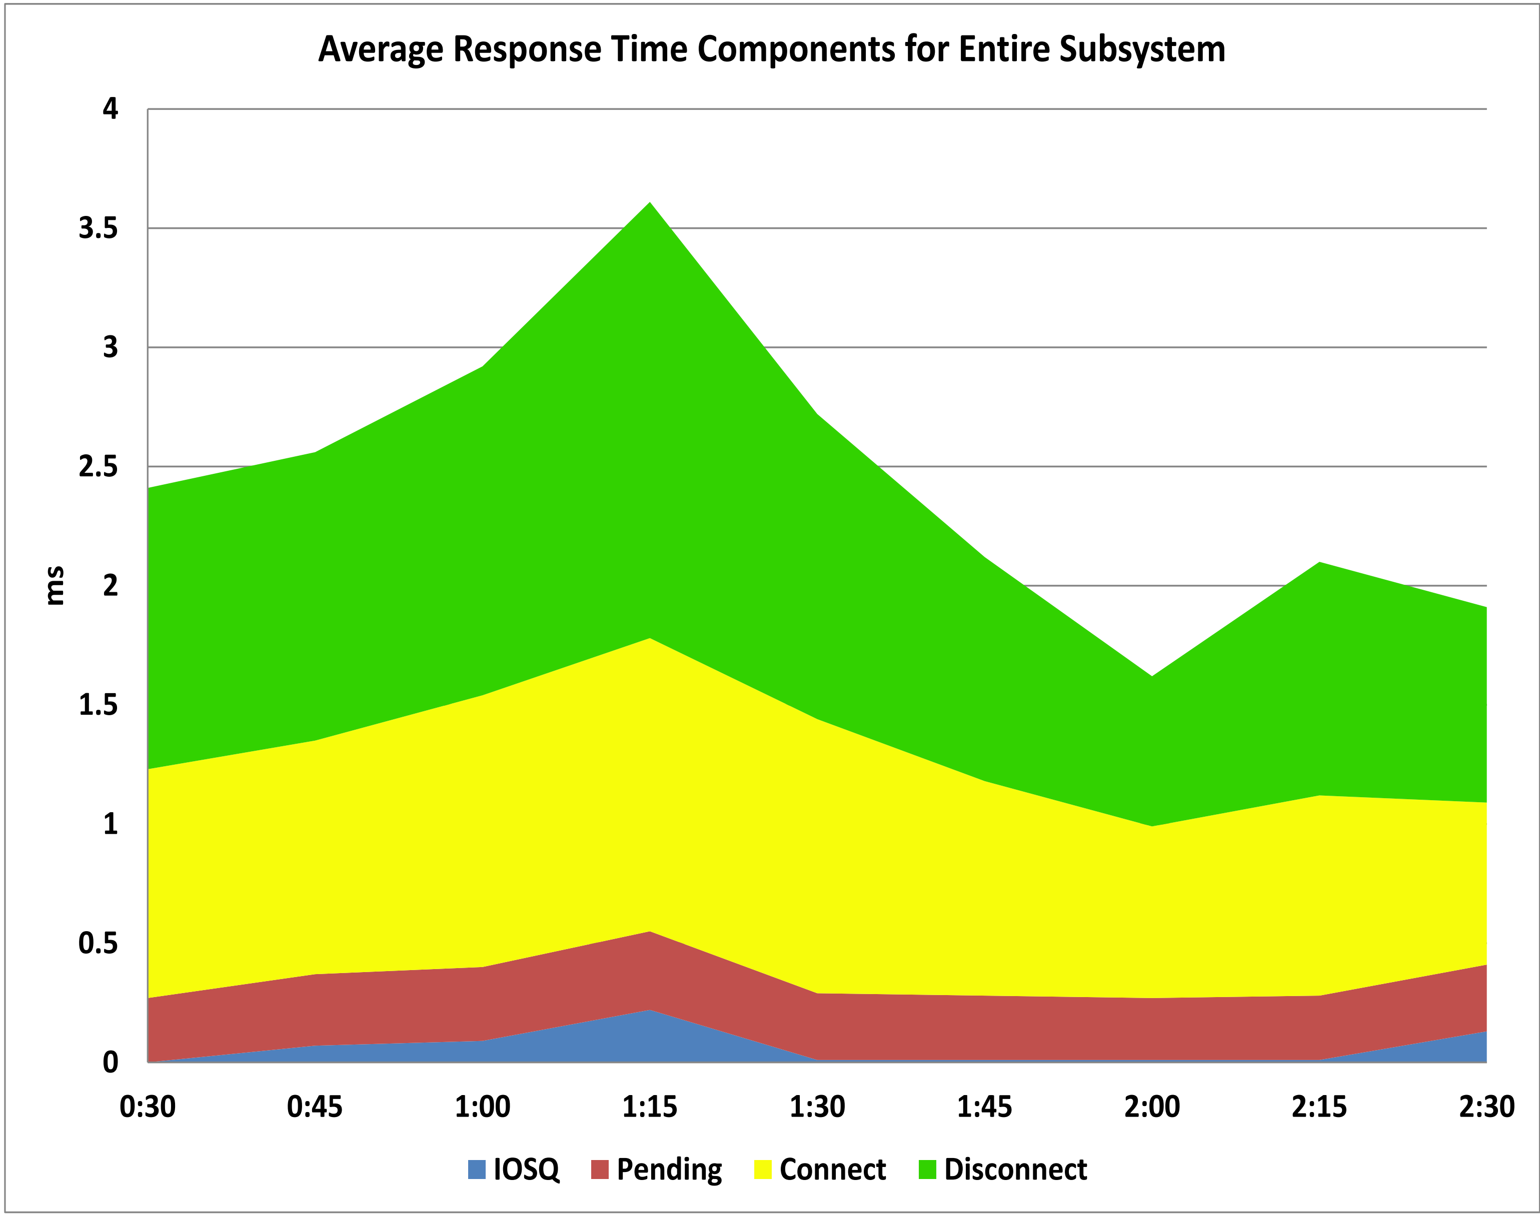

IntelliMagic

"Back-End" Performance

Back-End I/O and Disconnect time

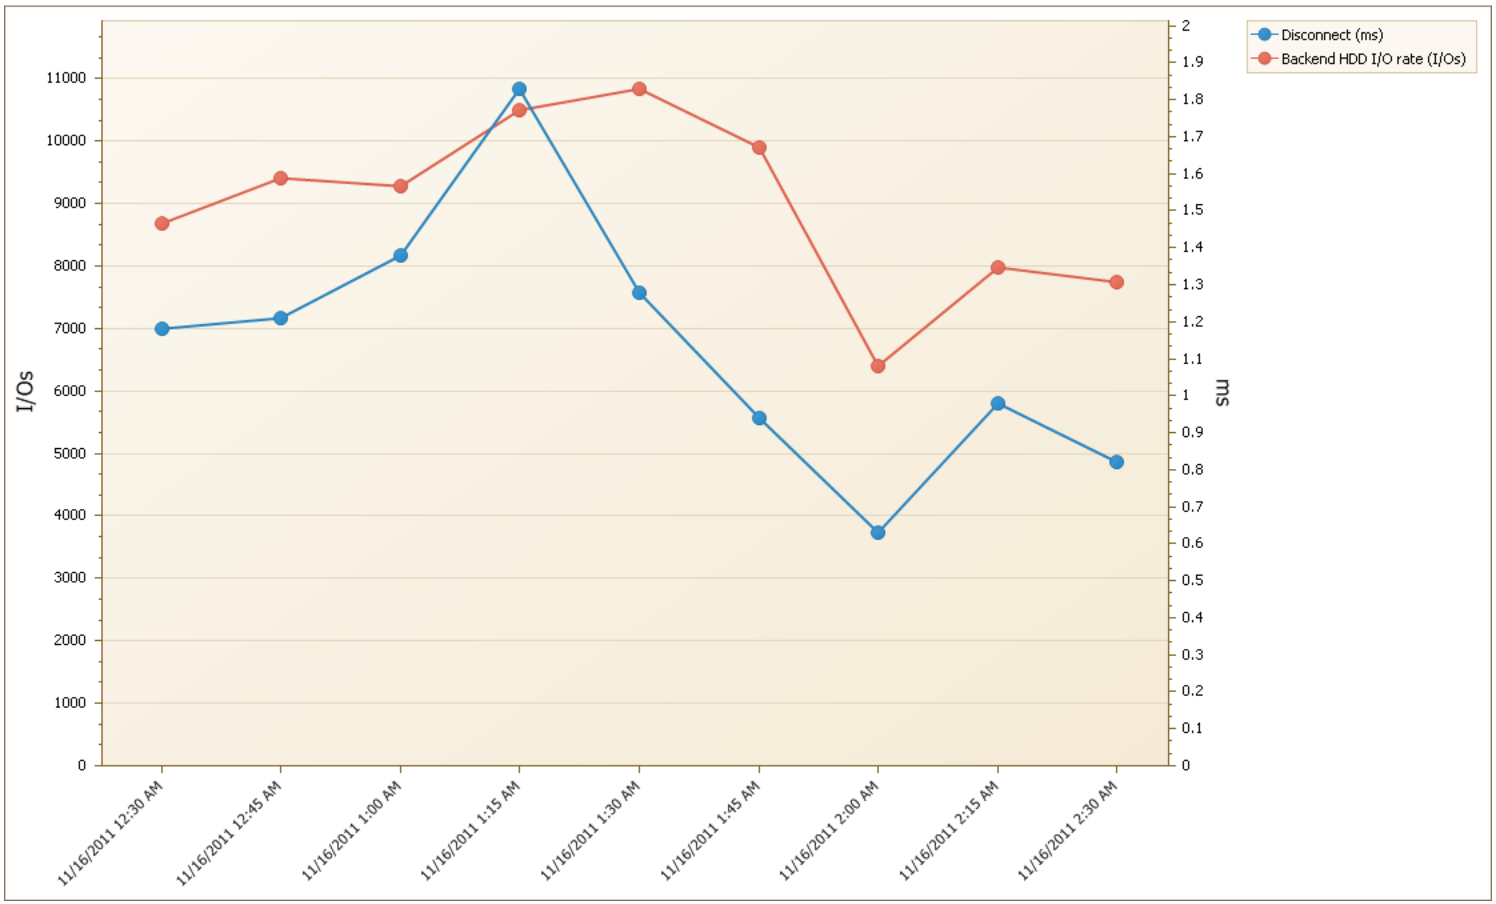

Investigate and correct or eliminate disconnect time. Back-end activity can drive disconnect time

Storage Intelligence
www.IntelliMagic.net

Physical Drive Performance

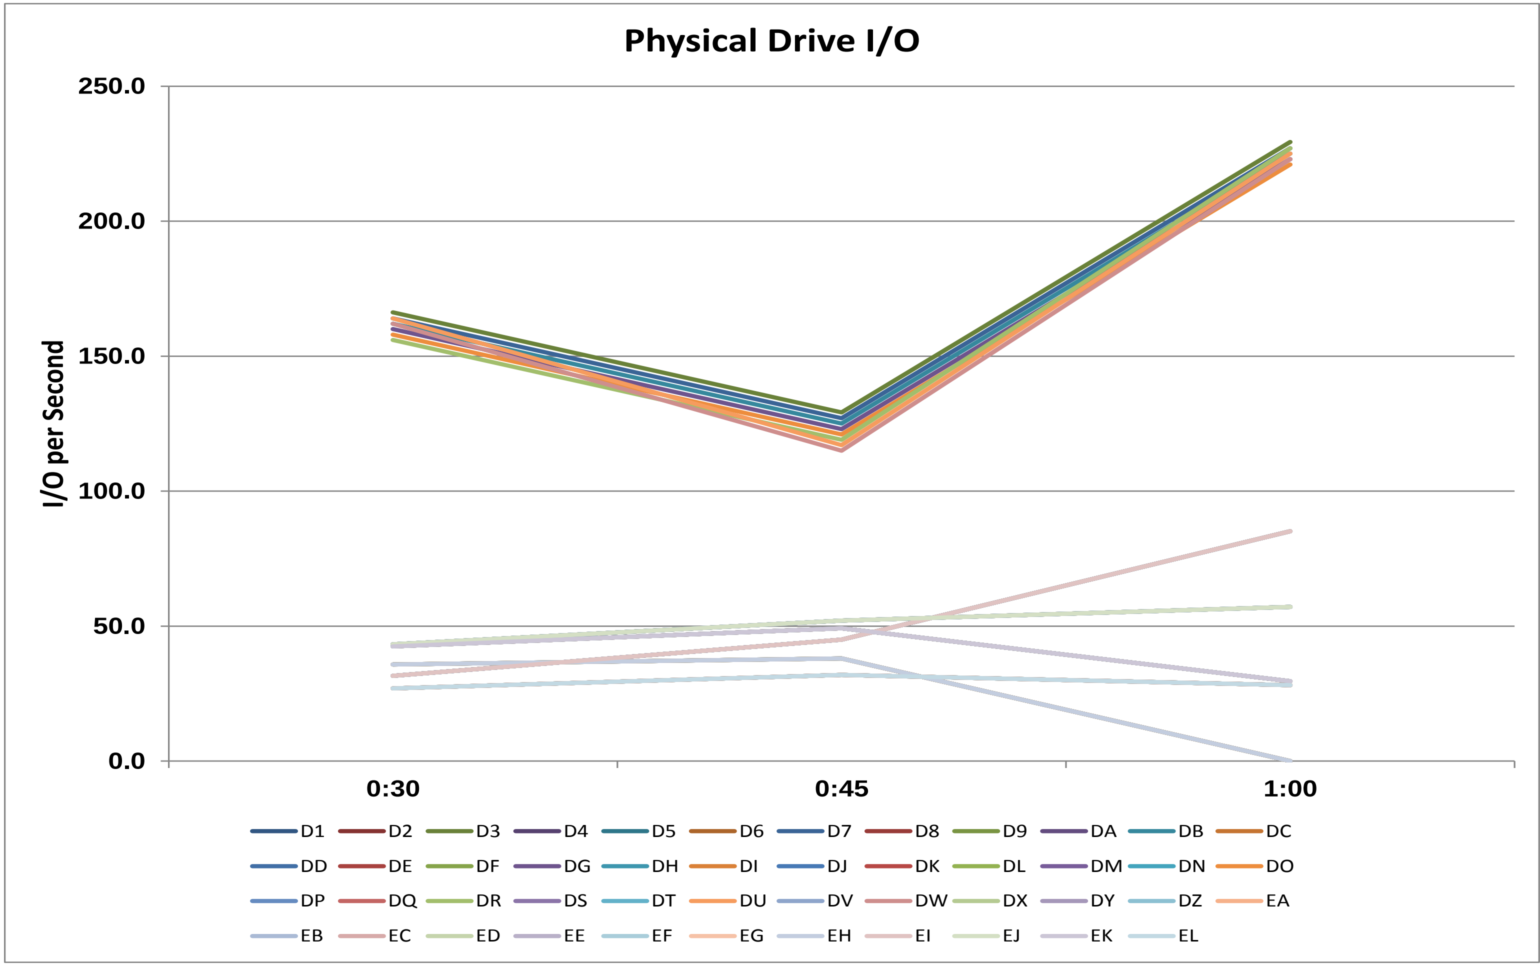

For this job, disconnect time was poor for the first three intervals of the run. The physical drives were exceeding their capabilities

IntelliMagic

Storage Intelligence
www.IntelliMagic.net

Connect Time

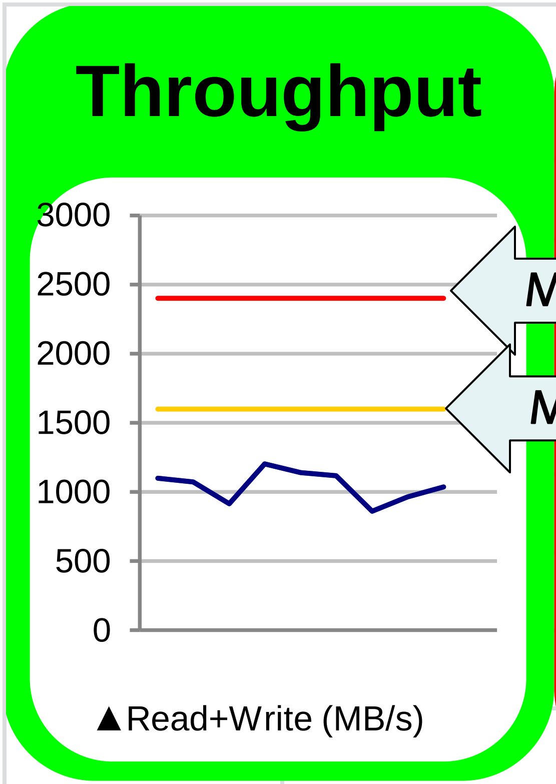

Machine Capabilities Exceeded

Machine Capabilities Stressed

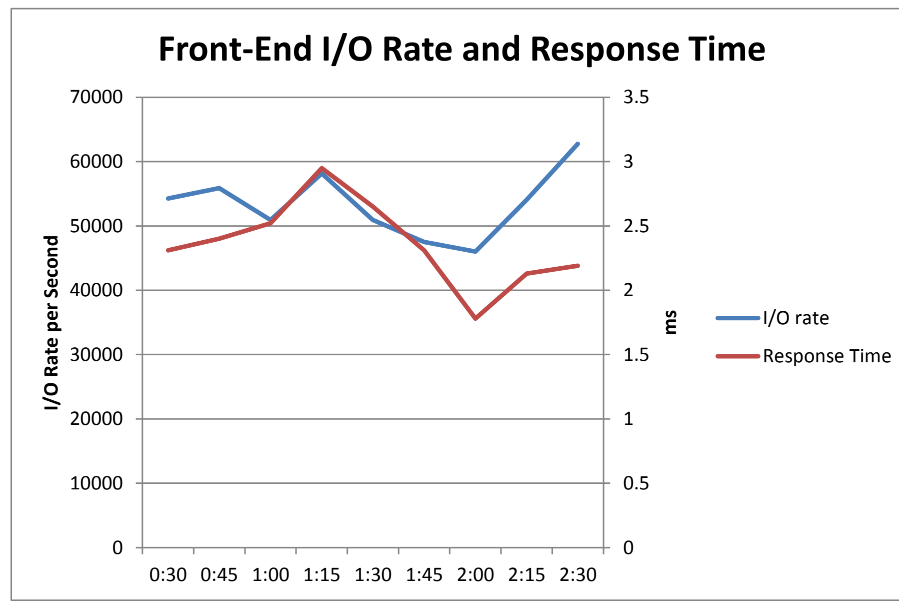

Can connect time be attributed to exceeding the machine capabilities?

Model Effect of "best-in-class" I/O

The case study compared to modeled times of best in class configurations

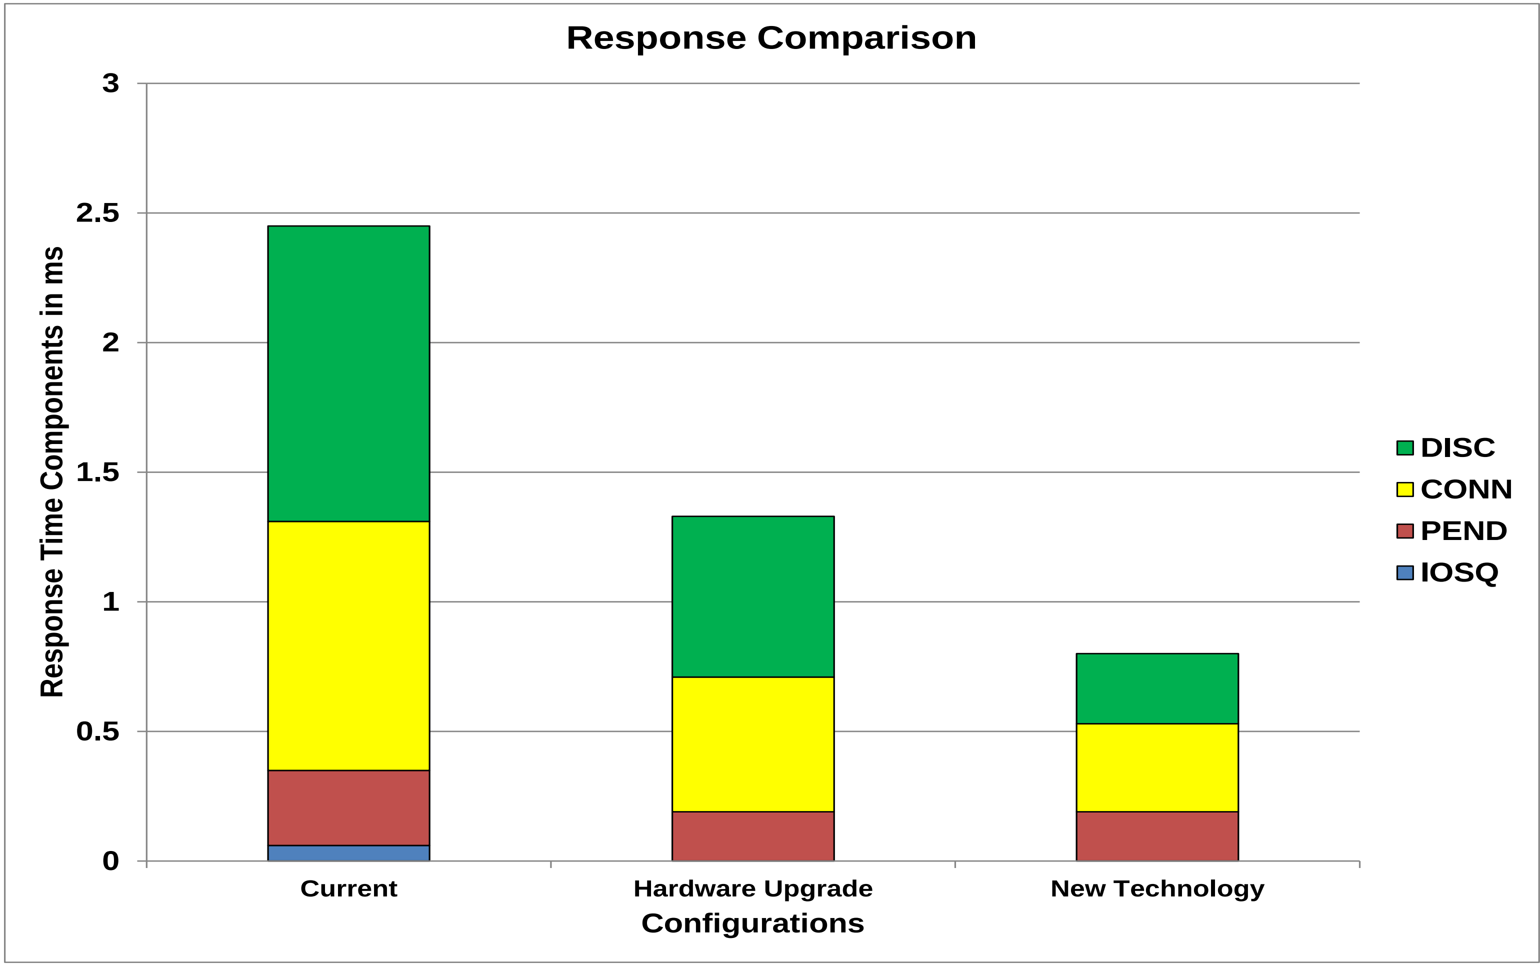

Storage Intelligence
www.IntelliMagic.net

Job Component Summary After DASD Improvement

Original

CPU Using	CPU Delay	DASD Using and Delay	Total Time
1196 seconds	1523 seconds	3915 seconds	6634 seconds

Projected

CPU Using	CPU Delay	DASD Using and Delay	Total Time
1196 seconds	1523 seconds	1027 seconds	3746 seconds

Contrast that performance gain with a similar gain through the CPU side

CPU Upgrade Scenarios

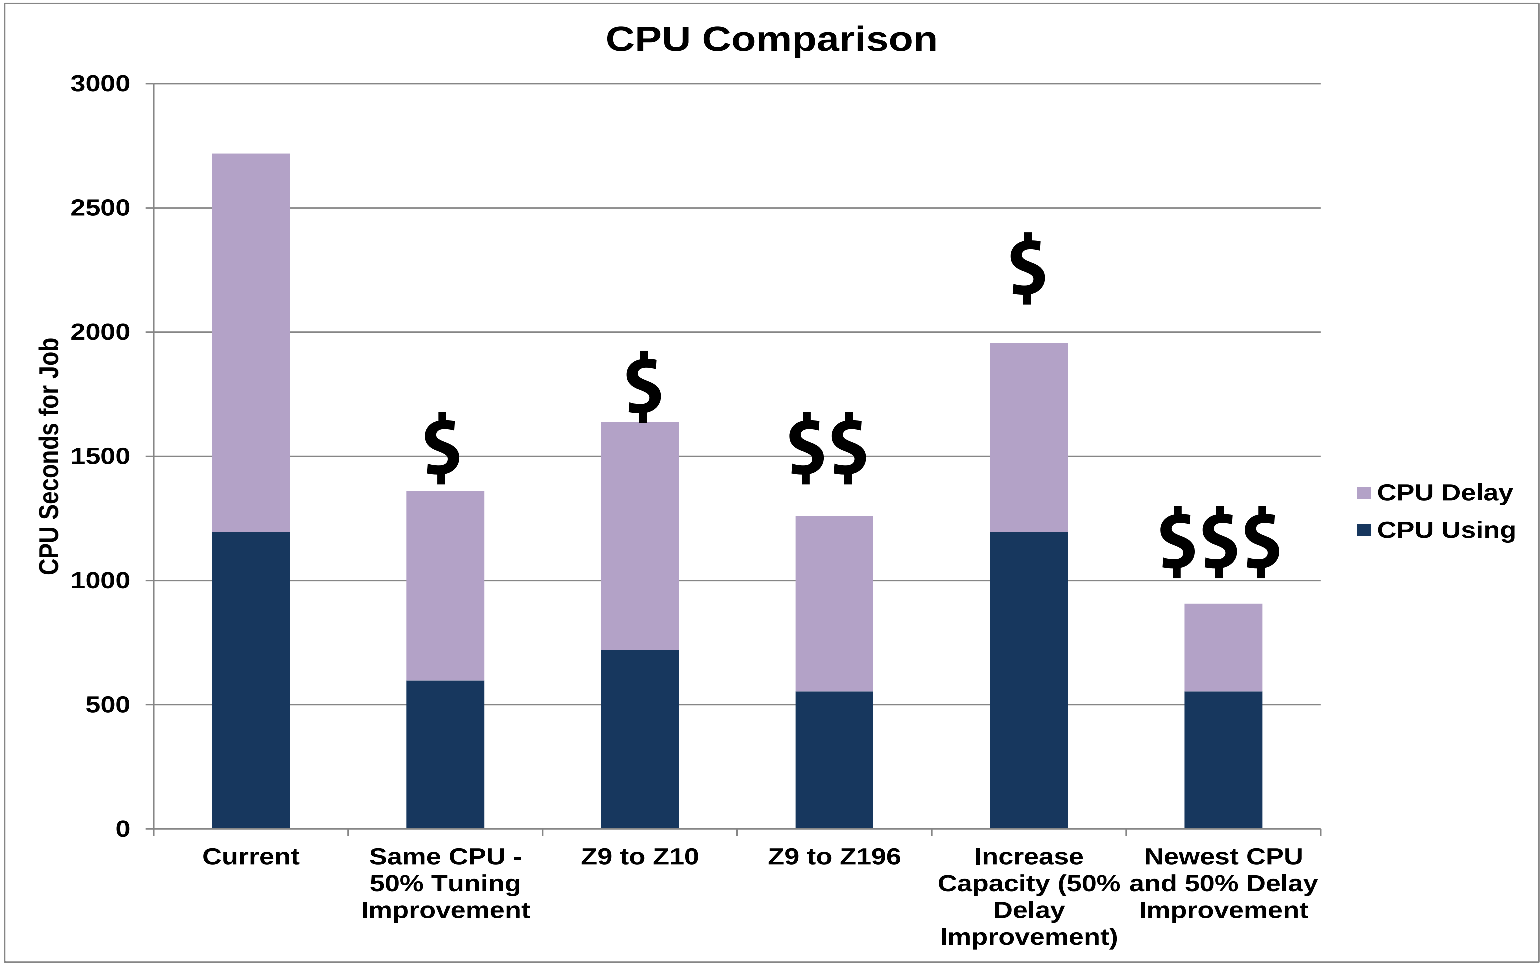

CPU Use/Delay can be solved by Tuning or by upgrades

IntelliMagic

Storage Intelligence
www.IntelliMagic.net

Job Component Summary after CPU Improvement

Original

CPU Using	CPU Delay	DASD Using and Delay	Total Time
1196 seconds	1523 seconds	3915 seconds	6634 seconds

CPU Using	CPU Delay	DASD Using and Delay	Total Time
554 seconds	353 seconds	3915 seconds	4596 seconds

Projected

The new CPU numbers are based on the best improvement possible

Job Run Time Summary

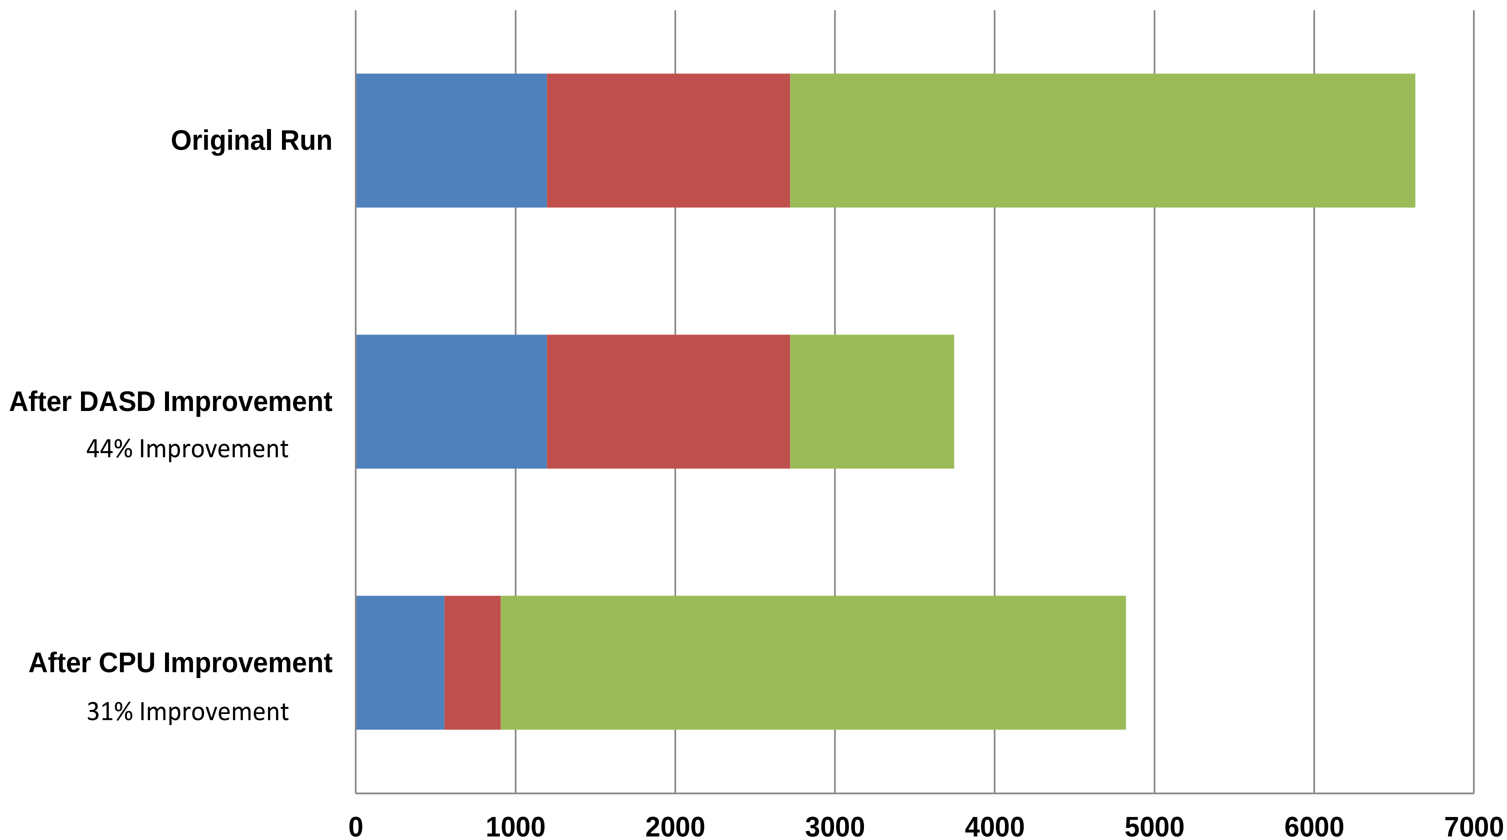

	After CPU Improvement	After DASD Improvement	Original Run
CPU Using	554	1196	1196
CPU Delay	353	1523	1523
DASD Using and Delay	3915	1027	3915

Seconds

Evaluating I/O vs. CPU Optimization
Conclusion:

- The original job ran for **110 minutes**
- Best-in-class I/O can reduce the time to **62 minutes**.
- Upgrading CPU could reduce the time to **80 minutes**.
- Solving it with I/O is likely far cheaper, especially if a disk replacement is scheduled anyway.
- Rule of thumb is that every dollar spent on a CPU upgrade requires at least **10** dollars in software charges

Thank You

Questions?

Lee.LaFrese@intellimagic.net