IntelliMagic

Availability Intelligence

for z/OS & Open Systems

Agenda

- **Intros and Overview**
  - Curtis Ryan, Vice President
- **Deep Dive on z/OS and Open Systems**
  - Lee LaFrese, Senior Performance Consultant

We are inspired by creating intelligence that illuminates the risks hiding inside your IT infrastructure.

-Availability Intelligence

Availability Intelligence for IT Infrastructure

- **What?**
  Foreknowledge about hidden threats to availability
- **Why?**
  To better protect continuous availability at primary site by
  1. Avoiding Incidents (Make More Predictable)
  2. Accelerating the Resolution (Reduce MTTR)
- **How?**
  *Use built-in **expert domain knowledge** in automatic analysis of the performance and configuration data*

Reporting Vs. Analytics Vs. Intelligence

Data + Analytics ≠ Intelligence

	Reporting	Analytics	Availability Intelligence
Primary Method	Data Visualization	Statistical Analysis	Infrastructure Knowledge
Shows	Tables, Graphs, etc.	Anomalies, Correlations	Interpretation of Good and Bad
What Answers From the Data	What is it?	What Patterns?	What Does it Mean?

Availability Intelligence for z/OS Disk & Systems

Avoid Issues & Accelerate Fixes for all Major z/OS Disk & Systems Areas

Disk Storage
- Front End
- Back End
- Channels & zHPF
- etc

Coupling Facility
- CF / XCF Health
- CF / XCF Analysis
- Trending
- etc

Disk Replication
- Replication Status
- Rating Over Time
- Verify Balance
- etc

Jobs and Datasets
- Data Sets
- Address Spaces
- Trending
- etc

Systems, Paging, WLM
- Processor Reporting
- Paging Reports
- Workload Manager
- etc

FICON Directors
- Director Health
- Service Statistics
- Channel Health
- etc

Availability Intelligence for z/OS Tape and Virtual Tape

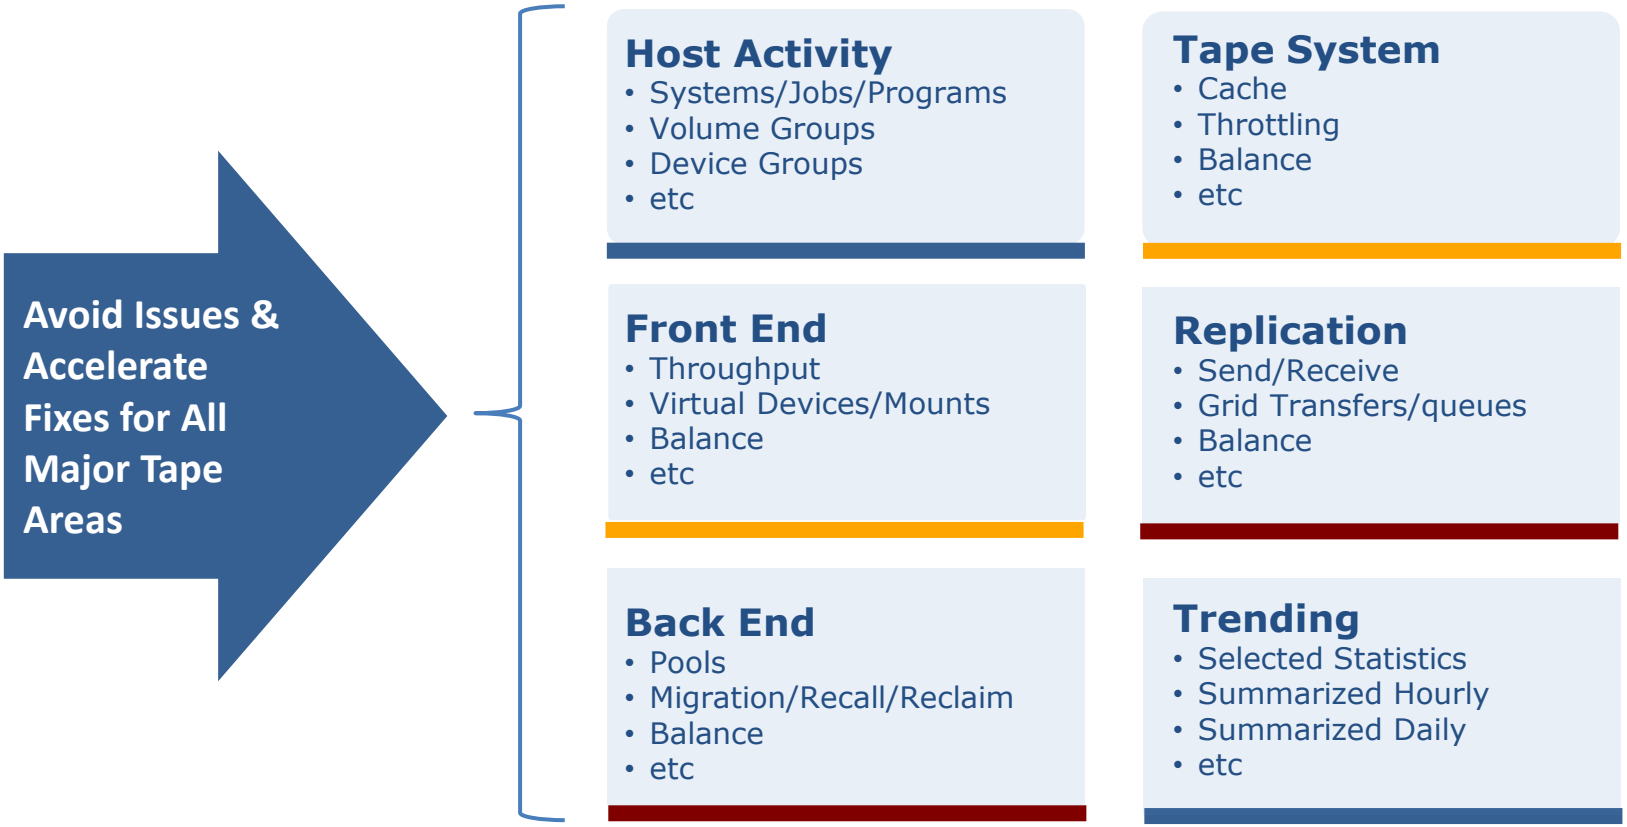

Availability Intelligence for Open Systems Storage & Fabric

Avoid Issues & Accelerate Fixes for all Major SAN Storage & Fabric Areas

SAN Storage
- DSS Health
- Replication Health
- etc

Groups & Clusters
- Volumes Groups
- Cluster Health
- etc

Tier Dashboards
- DSS Tier
- RAID Group Tier
- Disk Drive Tier
- etc

SVC Dashboards
- Front end
- Back end
- Nodes
- Ports
- etc

Brocade/Cisco Fabric
- DSS Port Errors
- Switch Port Error
- etc

Trending
- Selected Statistics
- Summarized Hourly
- Summarized Daily
- etc

Typical Customer Priorities

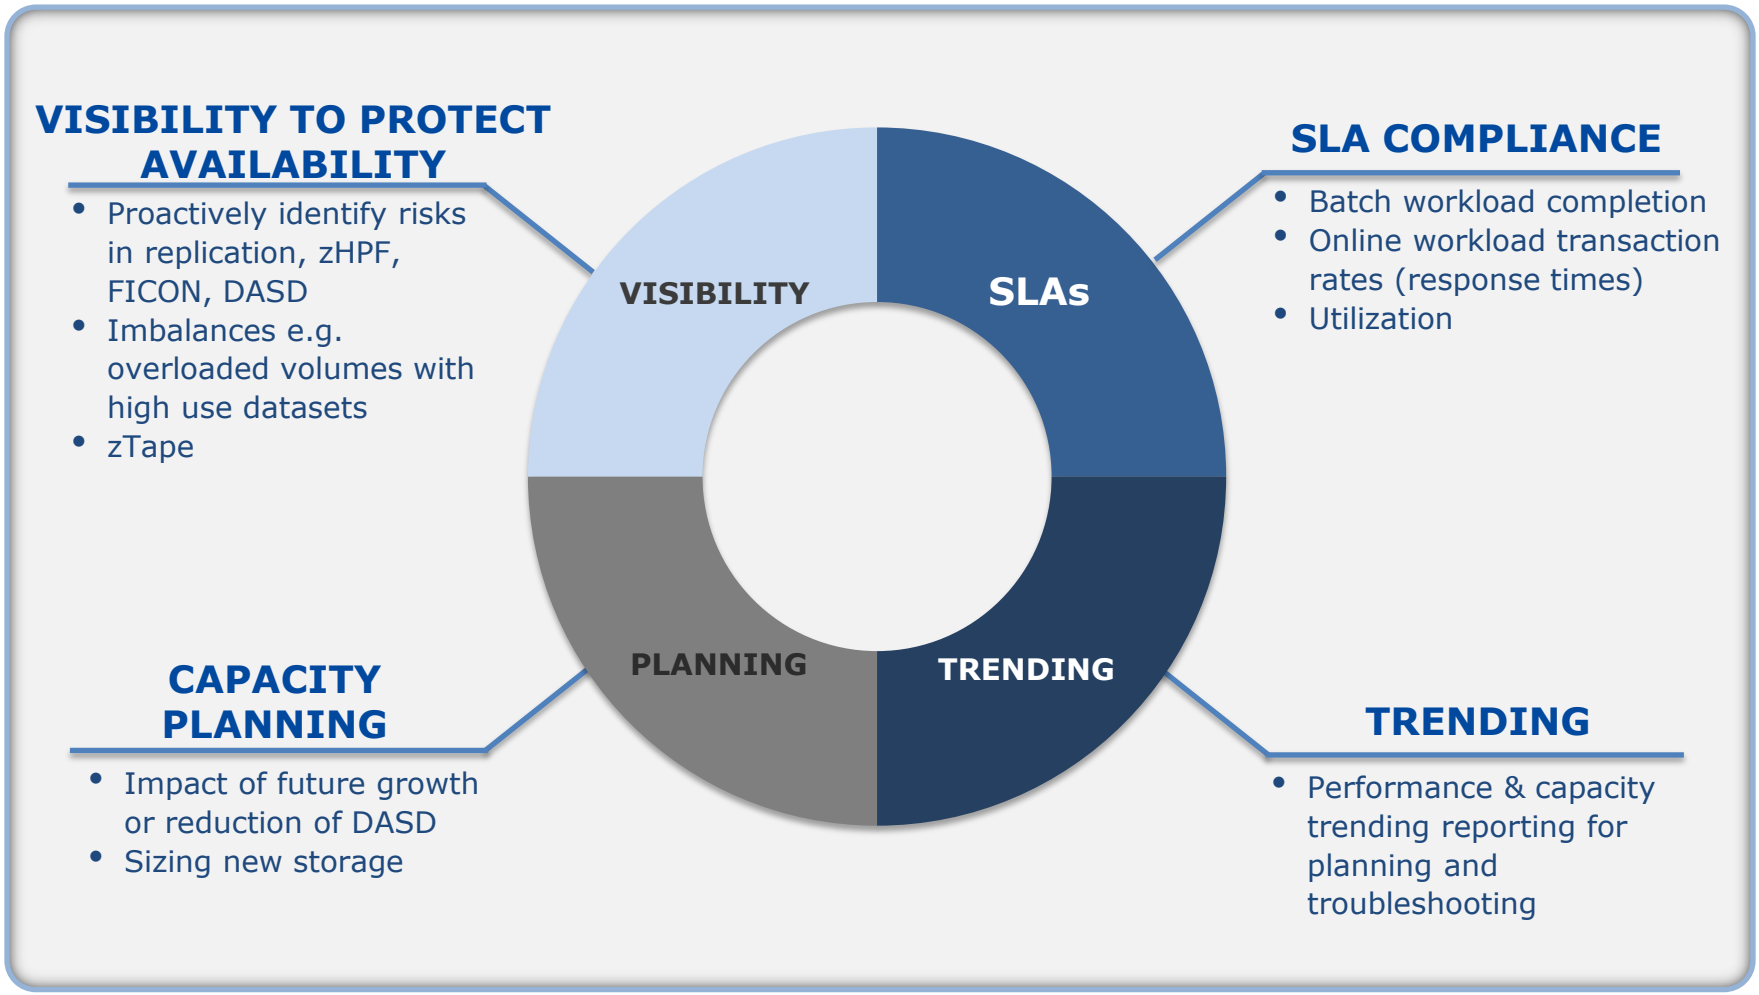

Key Capabilities to be Discussed

- **Greater Visibility**
  - Into storage and fabric performance, usage, availability
  - Ingest a ton of data from large storage and fabric infrastructure, analyze it and present it in a single-pane-of glass that can be easily understood and drilled down for insight

- **Early Warning to Protect Availability**
  - Proactively identify hot spots and risks in the infrastructure

- **Easy Troubleshooting**
  - Accelerate problem resolution with **quick and intuitive** drill down of issues

IntelliMagic

IntelliMagic Vision for z/OS and SAN

Lee LaFrese
Senior Consultant, Americas
lee.lafrese@intellimagic.com
+1-214-432-7920 x718

Availability Intelligence

IntelliMagic Solutions

- z/OS Storage and Systems
  - IntelliMagic Vision for z/OS
    - Covers entire z/OS infrastructure
  - IntelliMagic Direction Service for z/OS
    - What-if performance modeling of z/OS storage

- Open Systems Storage and Fabric
  - IntelliMagic Vision for SAN
    - Covers distributed storage and switches
    - Applicable to **z/VM, zLinux**, Windows, AIX, VMWare, etc.
  - IntelliMagic Direction Services for SAN
    - What-if performance modeling of SAN storage

IntelliMagic Vision for z/OS Architecture and Design Overview

Covers Entire z/OS Infrastructure

- Storage Systems (74.1, 74.5, 74.8, 78.3)
- Replication: GDPS GM, EMC SRDF/A, XRC (105, 206, 42.11)
- Host channels (73), FICON Director (74.7)
- Processors
  - CEC, LPARs (70, 72, 113)
  - WLM goals (72)
  - Paging (75)
- Coupling Facility (74.4)
- XCF (74.2)
- Tape and Virtual Tape (SMF 14,15,21,30, IBM BVIR, Oracle StorageTek)
- Job records (SMF 30) and Dataset records (SMF 42).
- Supports all z/OS mainframe disk storage (EMC, IBM, HDS, HP)

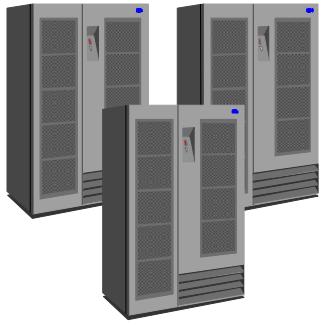

Note: Not all features listed are included in every license

Solution Architecture Overview – z/OS

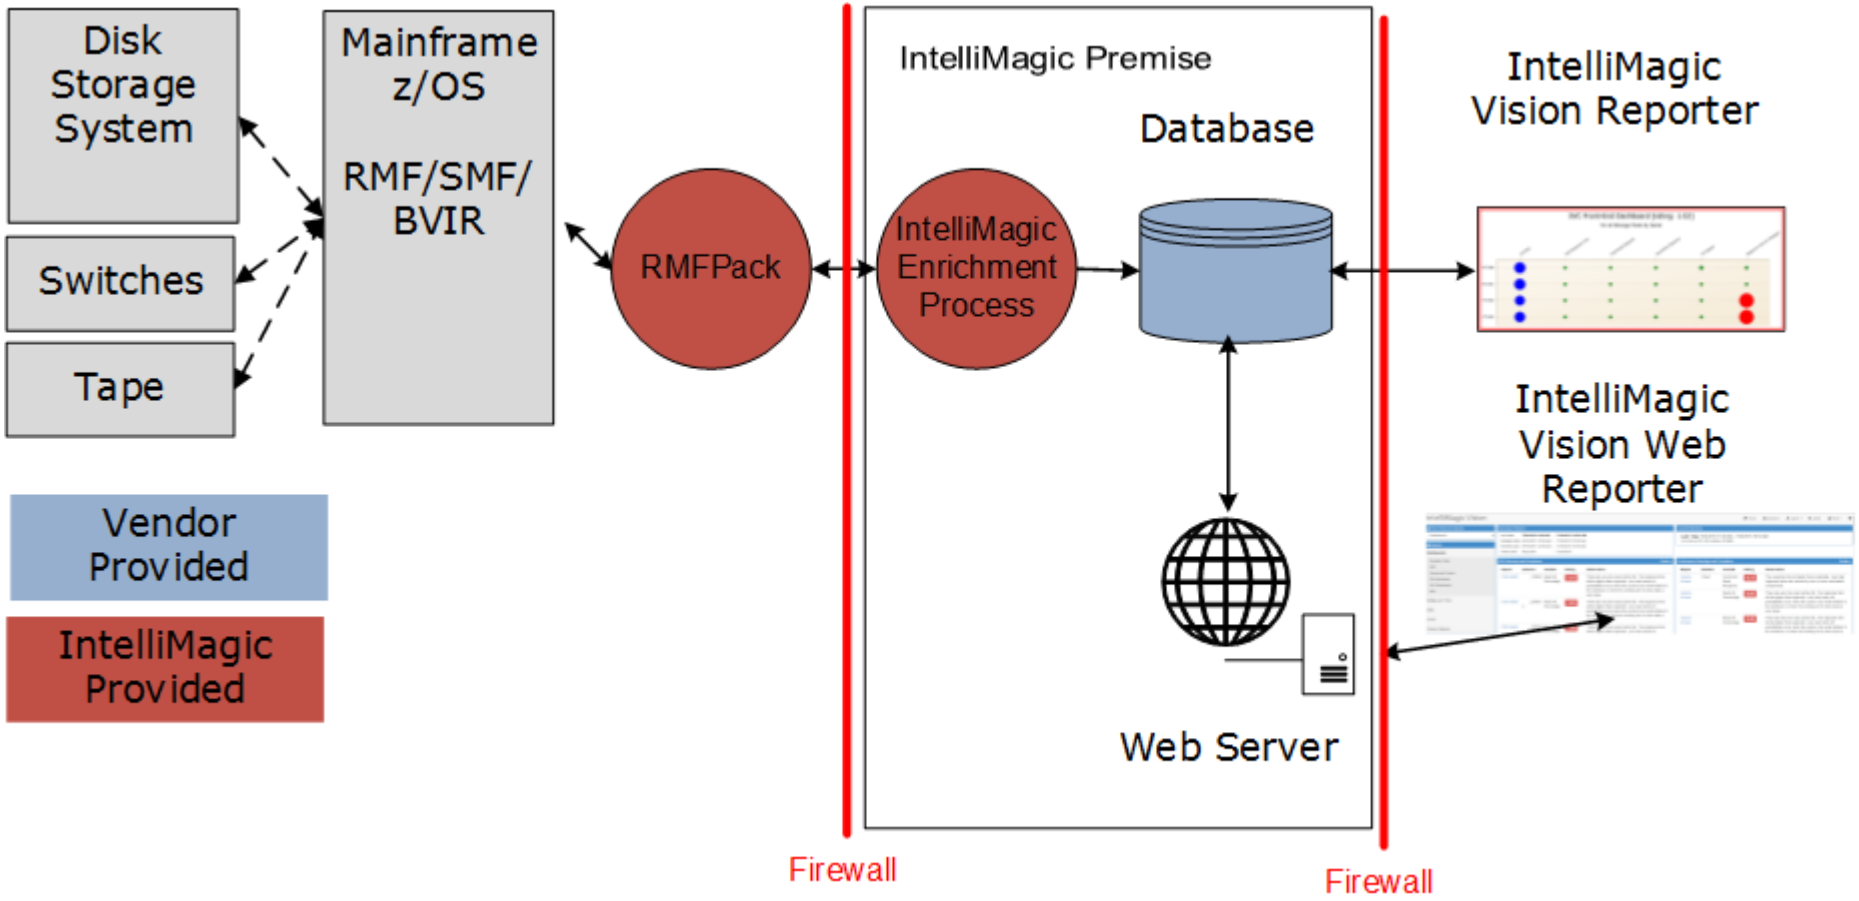

IntelliMagic Vision as a Service

Combine RMF Data

IntelliMagic Vision filters and enriches RMF data for data mining:

- **Align** native z/OS and external data
  - E.g. 74.1 device and 74.5 cache counters
- **Eliminate** redundant data across z/OS images
  - For each system there is unique as well as duplicate data
  - No easy way to 'sum' data from multiple systems
- **Supplement** with external data
  - Sometimes needed to get the complete picture
  - e.g. EMC SQ Mirror, DCOLLECT

IntelliMagic Vision for SAN Architecture and Design Overview

IntelliMagic Vision for SAN Logical Architecture

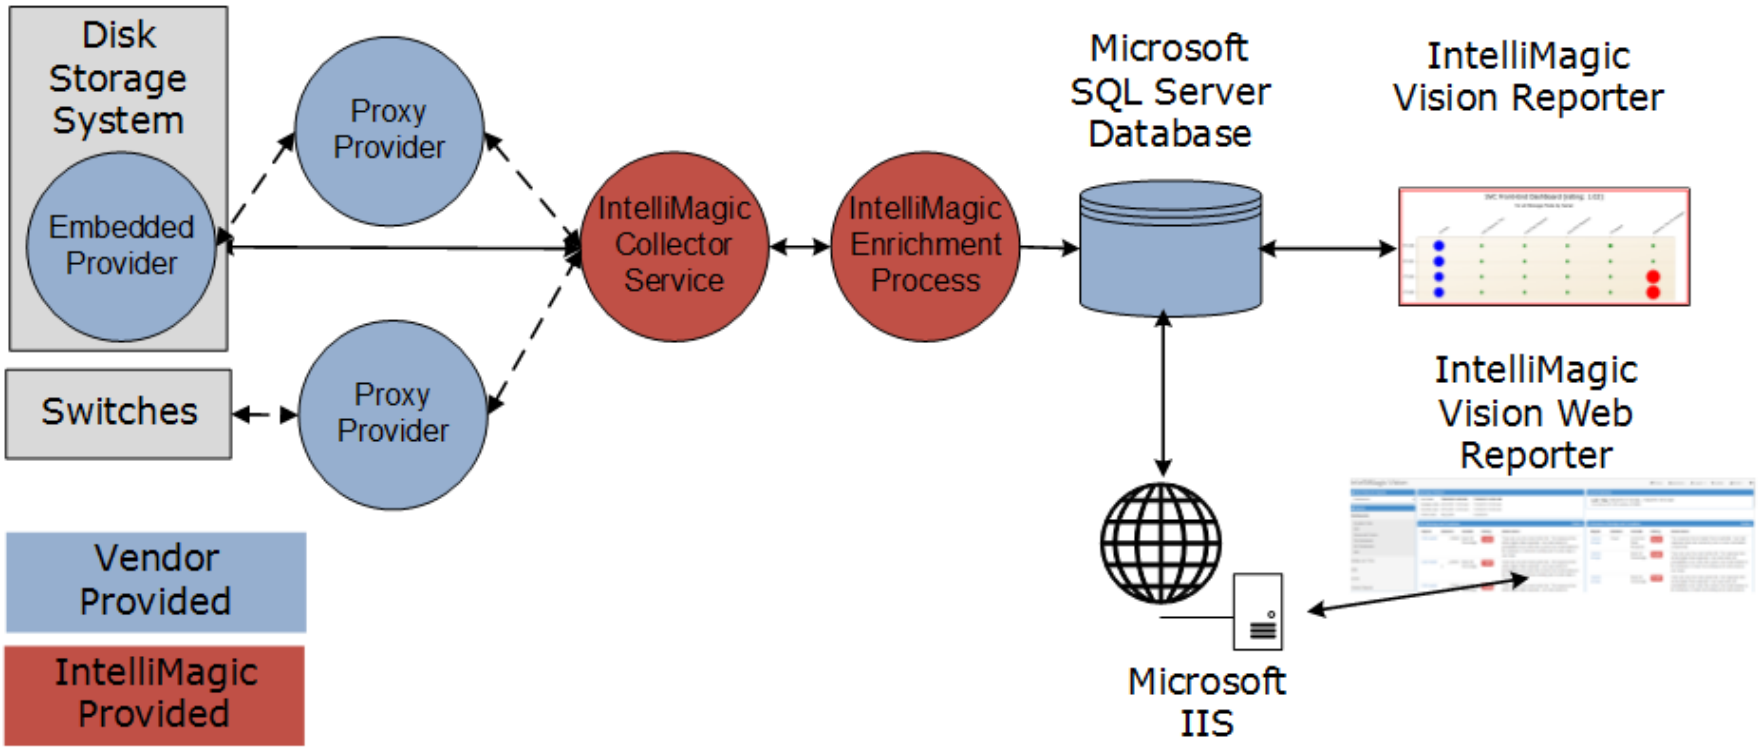

IntelliMagic Vision as a Service Architecture

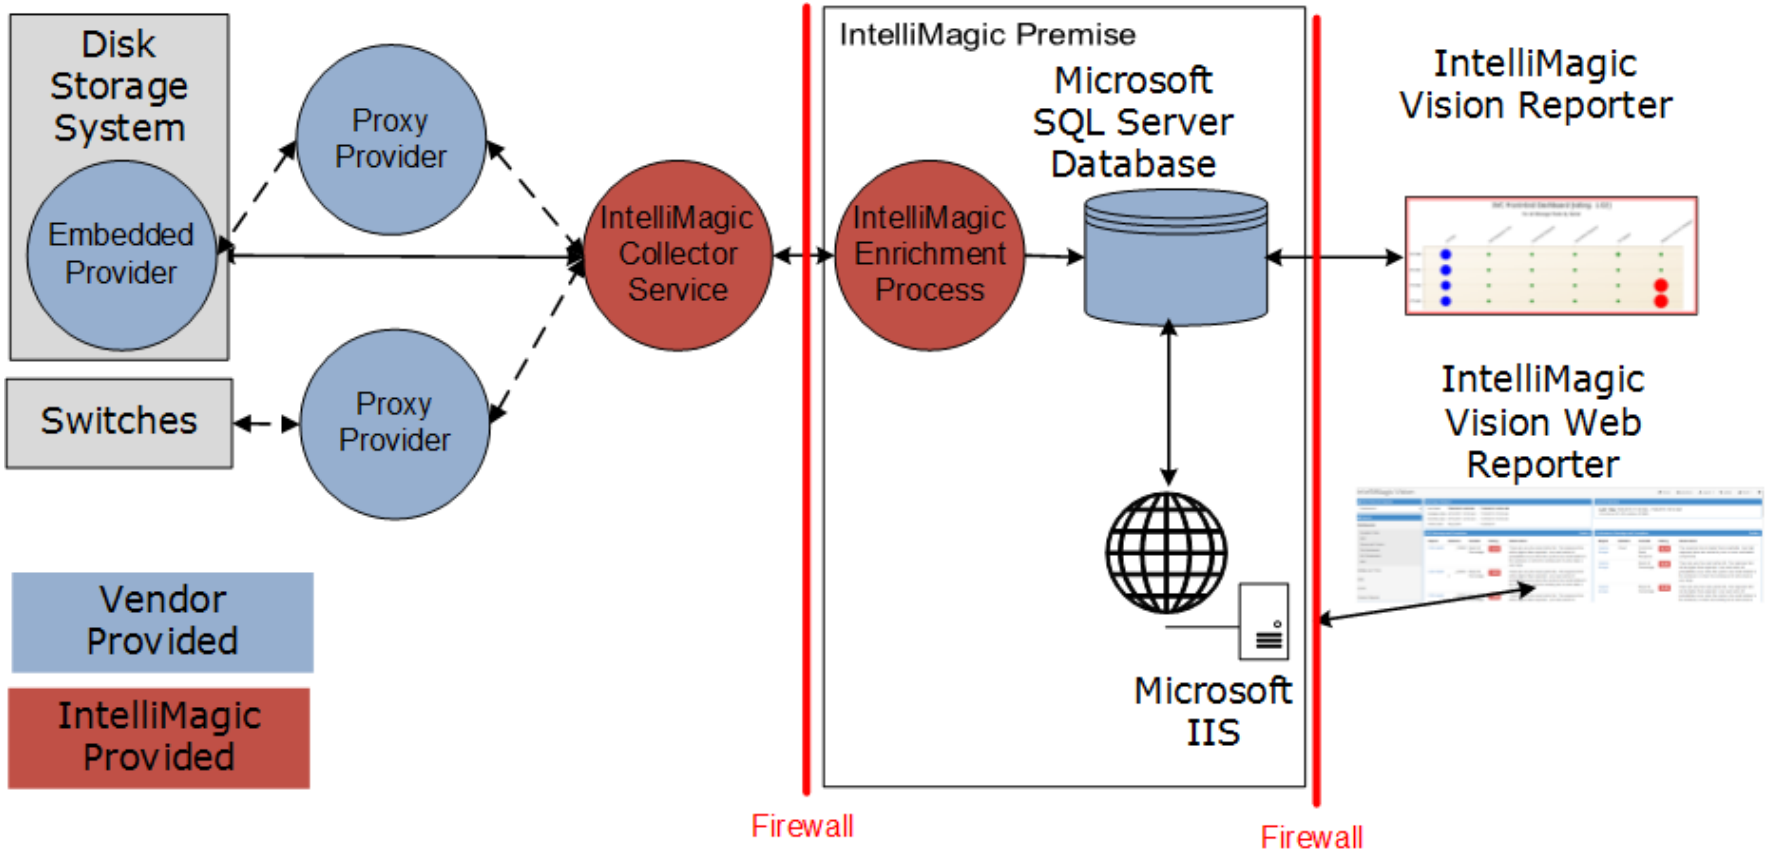

How IntelliMagic Vision Provides Availability Intelligence

Detect Risks

- Health rules are applied to enriched RMF data (z/OS) or data collected directly from storage and switches (SAN)

- This results in **Ratings** that show risk levels

Green = Healthy
Yellow = Early Warning
Resource getting busier, availability risk to application
Red = Exception
Busy or slow resource, impact to application

- This process is applied to 100s of metrics, with many thousands of data points across every collection interval!

Dashboards - Visualizing Key Risk Indicators

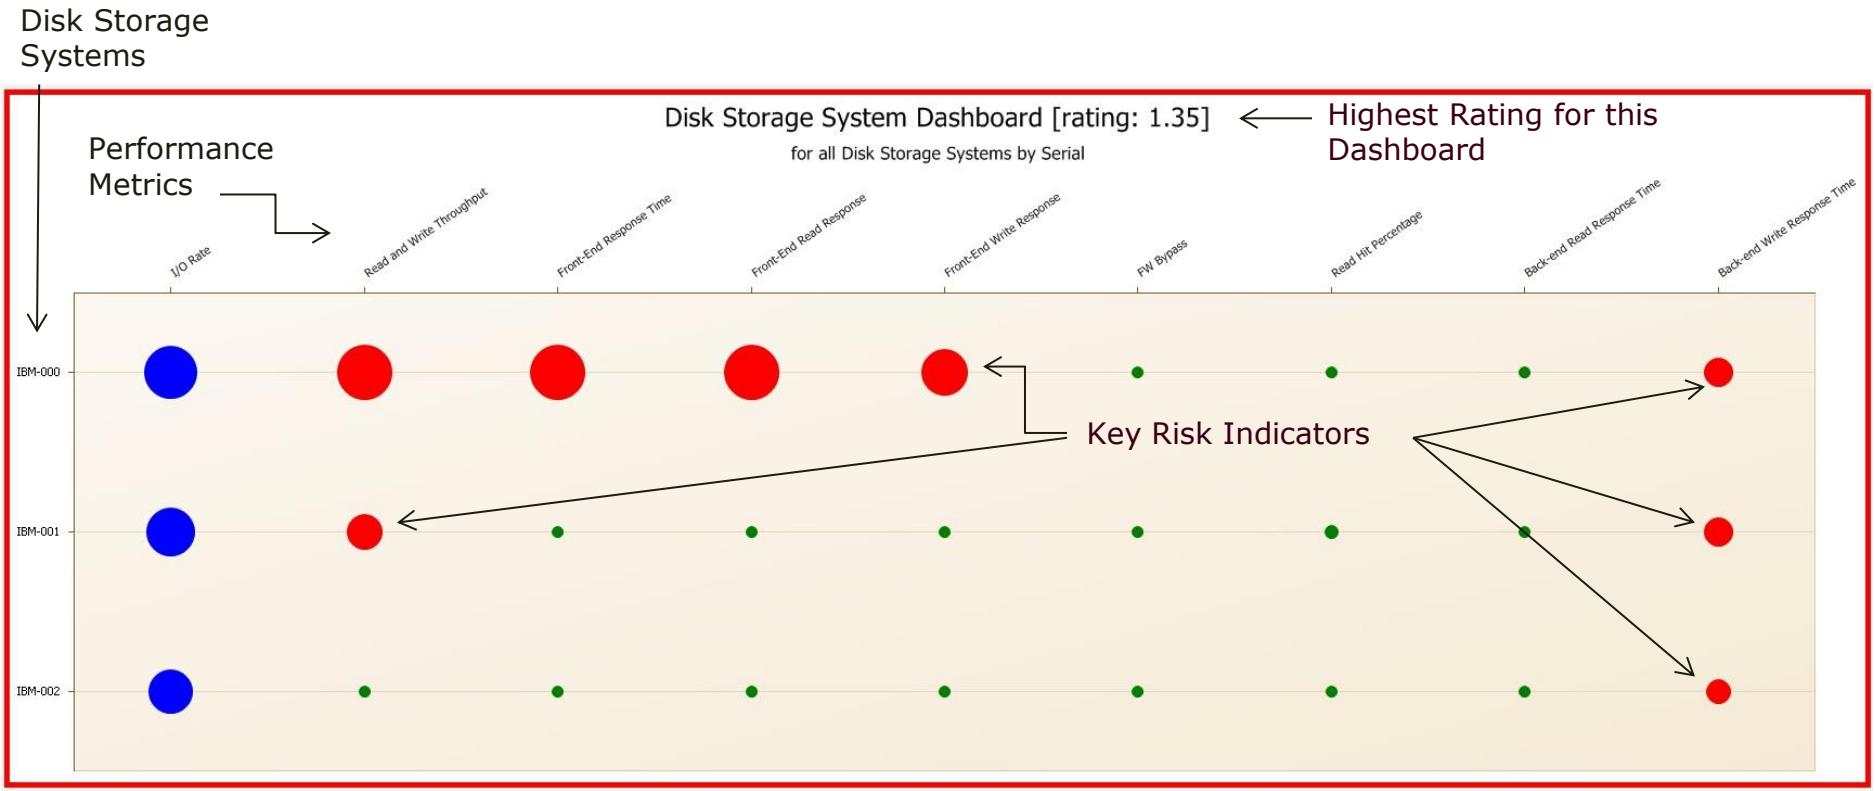

Knowledge-Rated Performance Reports

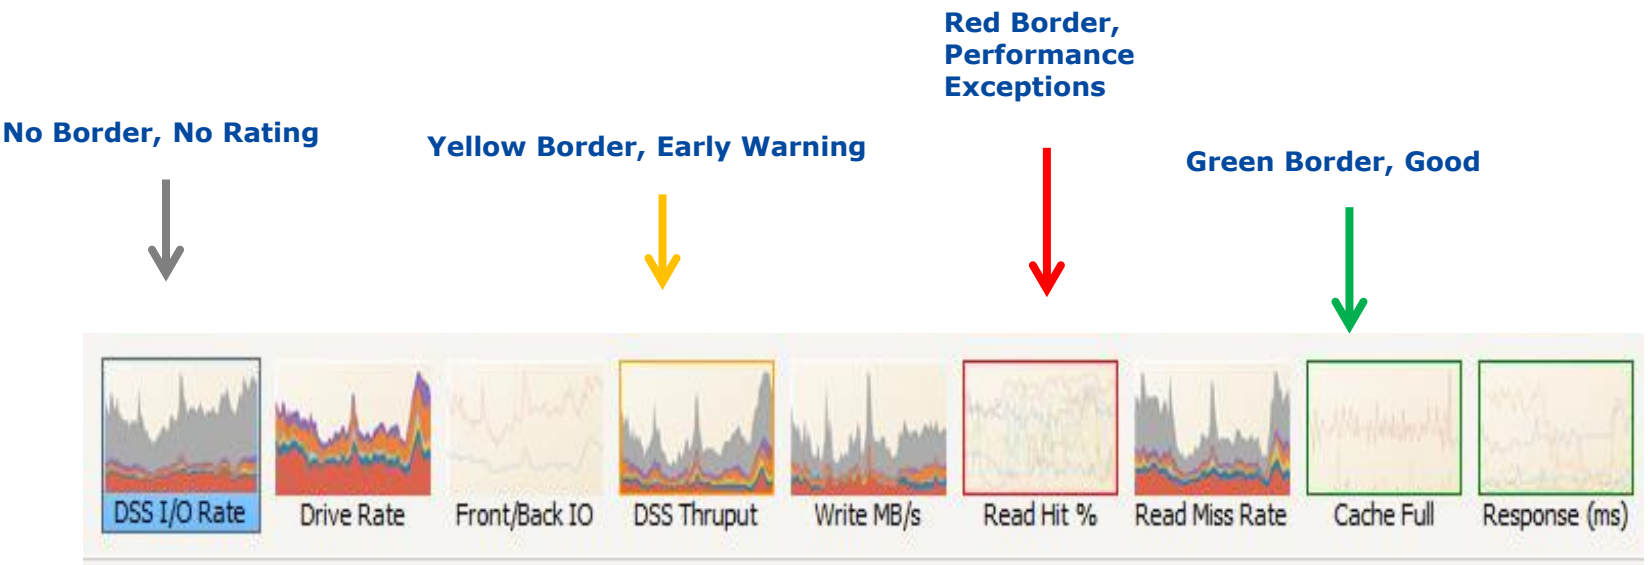

Reports for key metrics are rated according to adaptive thresholds defined per platform, providing pro-active warning of potential performance issues.

Multi-level Thresholds

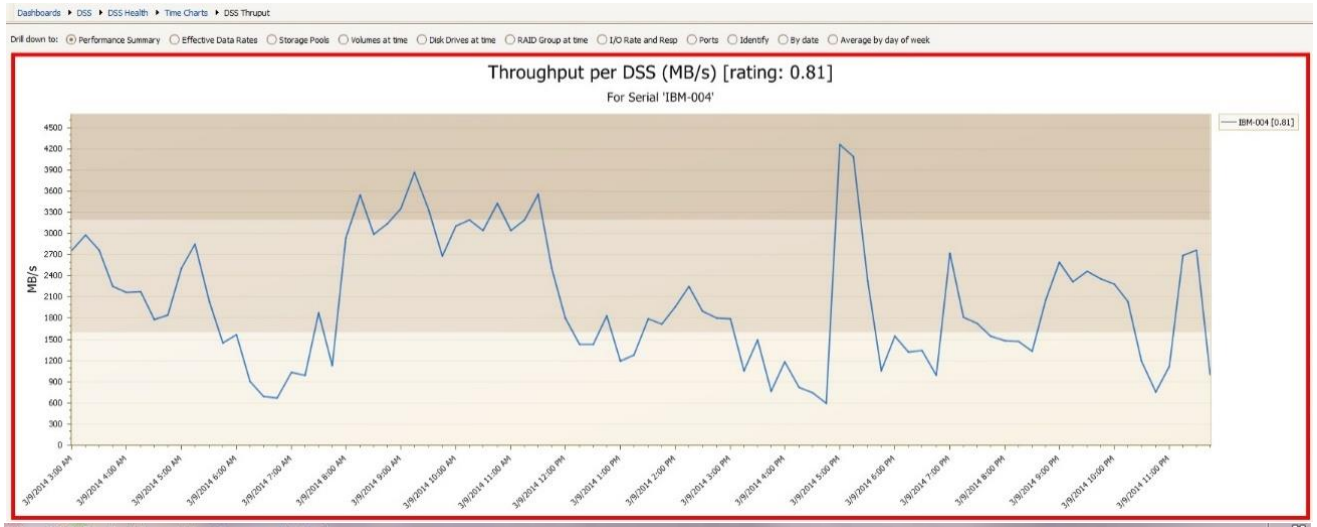

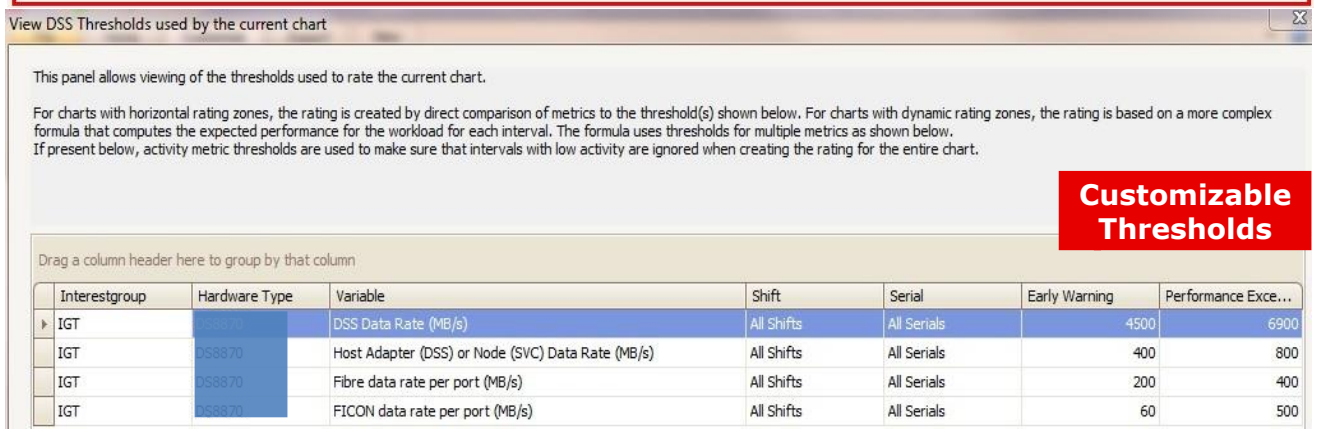

Interestgroup	Hardware Type	Variable	Shift	Serial	Early Warning	Performance Exce...
IGT	DS8870	DSS Data Rate (MB/s)	All Shifts	All Serials	4500	6900
IGT	DS8870	Host Adapter (DSS) or Node (SVC) Data Rate (MB/s)	All Shifts	All Serials	400	800
IGT	DS8870	Fibre data rate per port (MB/s)	All Shifts	All Serials	200	400
IGT	DS8870	FICON data rate per port (MB/s)	All Shifts	All Serials	60	500

Dynamic Workload-Based Thresholds

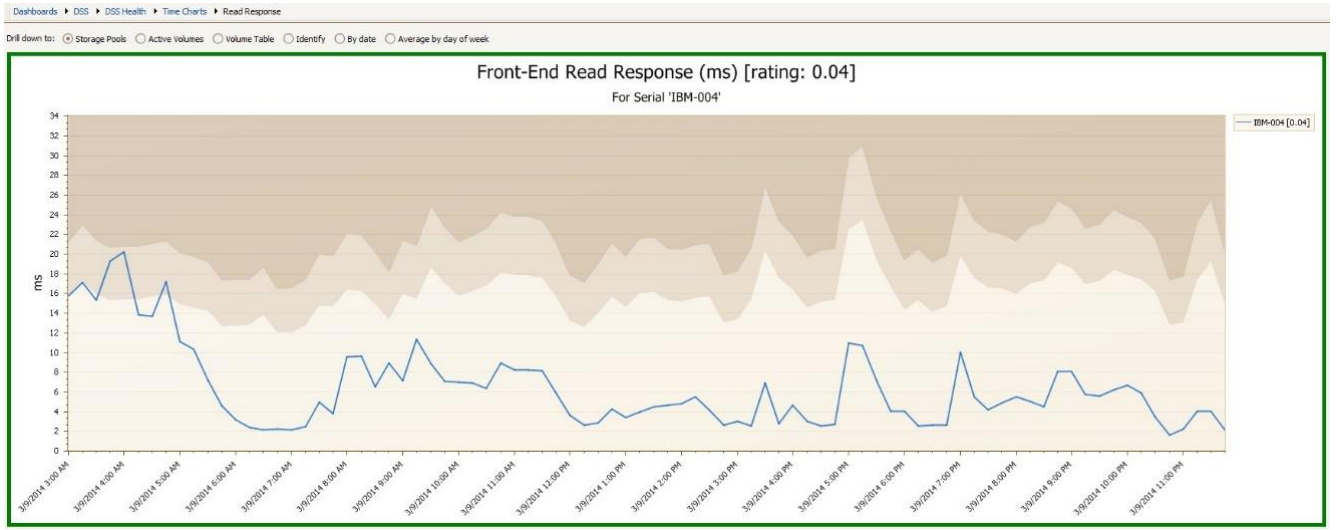

Drag a column header here to group by that column

Interestgroup	Hardware Type	Variable	Shift	Serial	Early Warning	Performance Exce...
IGT	DS8[redacted]	Connected data rate (MB/s)	All Shifts	All Serials	50	35
IGT	DS8[redacted]	Utilization of Drives (%)	All Shifts	All Serials	50	60
IGT	DS8[redacted]	Group I/O activity for interval to be included in the rating (I/Os ...	All Shifts	All Serials	10	50
IGT	DS8[redacted]	Overall cap for acceptable response time (ms)	All Shifts	All Serials	998	999
IGT	DS8[redacted]	Storage System overhead for read or write I/Os (ms)	All Shifts	All Serials	1.5	3

Other thresholds, like those for Front-end Read Response time are based upon the capabilities of the controller, activity and workload.

Predictive Modeling with IntelliMagic Direction Services

Why is Modeling Important?

- Support Storage Decisions
  - Add, Replace or Repurpose Disk Systems
  - Add Drives or Adapters
  - Increase Cache
  - Implement Automated or Manual Tiering
  - Deploy Replication
- Storage Capacity Planning
  - Evaluate Workload Growth
  - Examine Workload Changes
  - Anticipate Bottlenecks
- Vendor Performance Comparisons

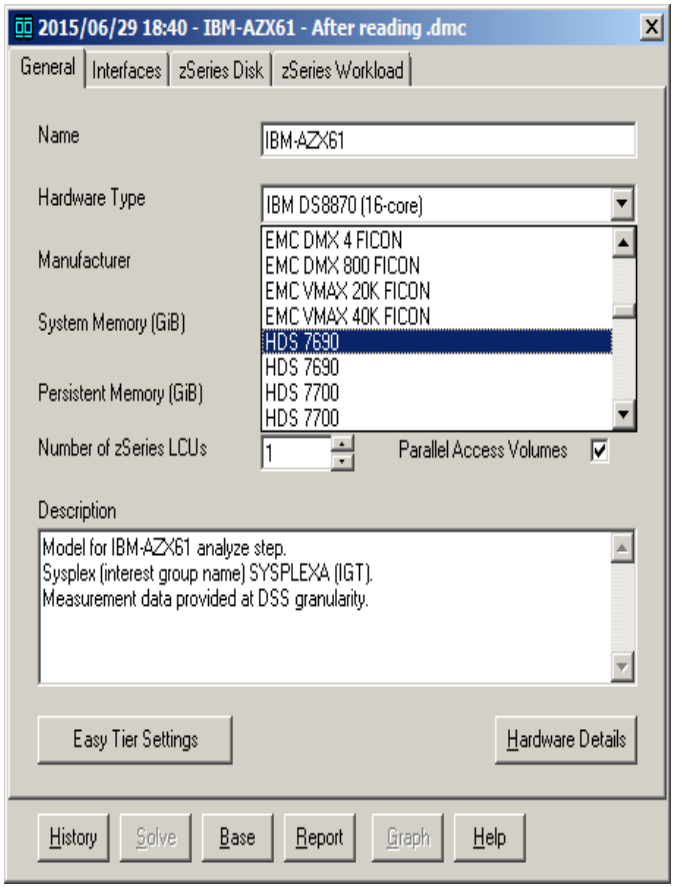

IntelliMagic Direction: Process

1. **Collect Data:** Input performance data from a representative interval as well as the current storage configuration to IntelliMagic *Direction*. IntelliMagic *Vision* can help you automate this process.
2. **Baseline:** IntelliMagic *Direction* will estimate unknown parameters to be consistent with the data provided
3. **Model Configuration Changes:** Solve model for various "what if" scenarios as desired to see the effect of configuration options on performance expectations
4. **Model Growth:** Produce graphs showing performance as a function of I/O growth for targeted configuration(s)
5. **Review and Repeat:** Sanity check results and repeat as needed for other systems, intervals or configuration options

IntelliMagic Direction Modeling Example

- Example Workload
  - Peak Throughput Interval - 1:00AM July 21st
- Example Scenario: What would the impact to performance be if a DS8870 configuration was changed to a smaller/cheaper option
  - From 16 core to 8 core (lower cost)
  - From 300GB drives to 600GB drives (600GB = less floor space, less power)
- Impact on current performance
- Impact of future I/O growth

Impact on Current Performance of Changing Configuration

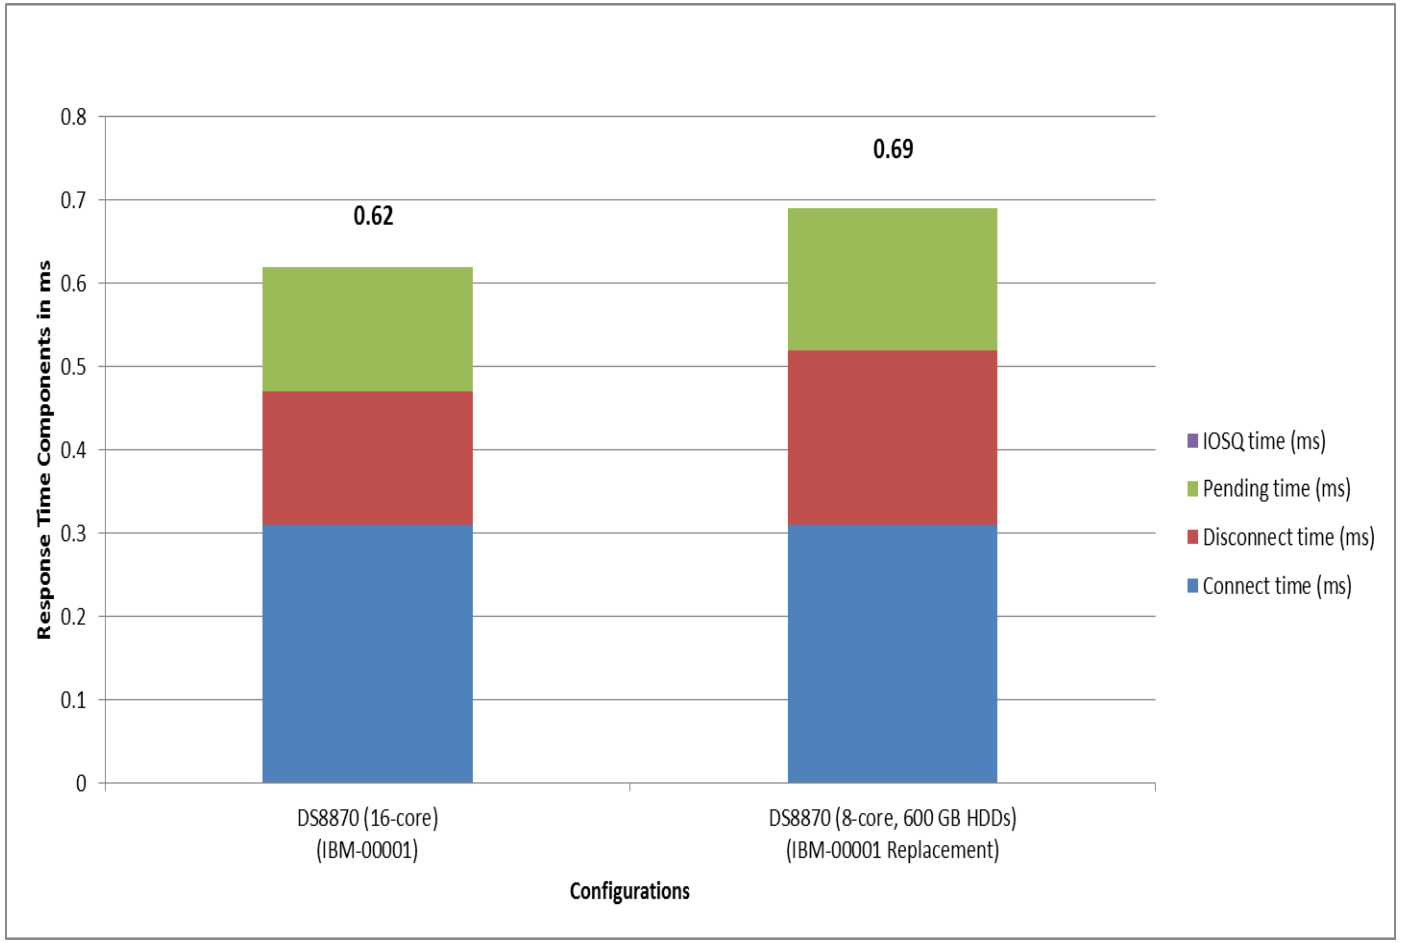

Projected Performance with I/O Growth

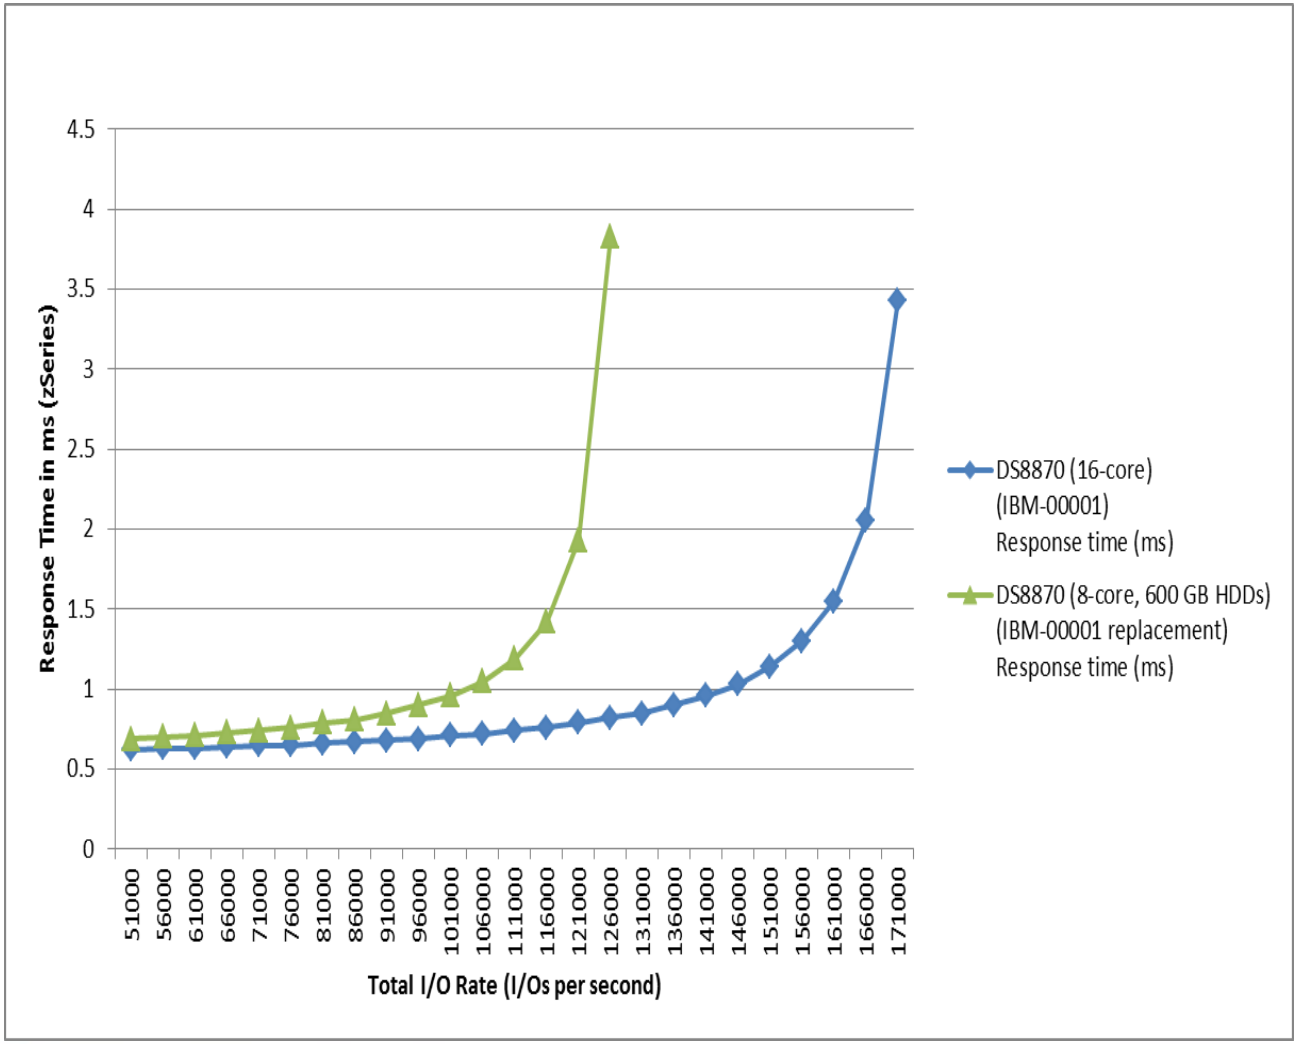

Resource Utilization Projections

Utilization Overview (zSeries)																		
16-way (IBM-00001)			**Total I/O Rate (I/Os per second)**															
Utilizations	Amber Threshold	Red Threshold	51000	81000	106000	111000	116000	121000	126000	131000	136000	141000	146000	151000	156000	161000	166000	171000	176000
Average SMP	60%	80%	28.8%	45.7%	59.8%	62.7%	65.5%	68.3%	71.1%	73.9%	76.7%	79.6%	82.4%	85.2%	88.0%	90.8%	93.7%	96.5%	99.3%
Average Bus	70%	90%	27.2%	43.1%	56.4%	59.1%	61.7%	64.4%	67.1%	69.7%	72.4%	75.0%	77.7%	80.3%	83.0%	85.7%	88.3%	91.0%	93.6%
Average Logical Device	n/a	n/a	0.9%	1.4%	1.8%	1.9%	1.9%	2.0%	2.1%	2.2%	2.3%	2.4%	2.5%	2.5%	2.6%	2.7%	2.8%	2.9%	3.0%
Highest DA	60%	80%	14.9%	23.6%	30.9%	32.3%	33.8%	35.2%	36.7%	38.1%	39.6%	41.0%	42.5%	44.0%	45.4%	46.9%	48.3%	49.8%	51.2%
Highest HDD (FC/SAS)	60%	80%	3.9%	6.1%	8.0%	8.4%	8.8%	9.1%	9.5%	9.9%	10.3%	10.6%	11.0%	11.4%	11.8%	12.2%	12.5%	12.9%	13.3%
Average FICON HA	50%	70%	18.4%	29.2%	38.3%	40.1%	41.9%	43.7%	45.5%	47.3%	49.1%	50.9%	52.7%	54.5%	56.3%	58.1%	59.9%	61.7%	63.5%
Highest FICON Port	50%	70%	3.6%	5.7%	7.4%	7.8%	8.1%	8.5%	8.9%	9.2%	9.6%	9.9%	10.3%	10.6%	11.0%	11.3%	11.7%	12.0%	12.4%
Average Fibre HA	60%	80%	22.7%	36.0%	47.0%	49.3%	51.5%	53.7%	55.9%	58.1%	60.3%	62.5%	64.8%	67.0%	69.2%	71.4%	73.6%	75.8%	78.1%
Average PPRC Link	60%	80%	16.5%	26.2%	34.2%	35.9%	37.5%	39.1%	40.7%	42.3%	43.9%	45.5%	47.2%	48.8%	50.4%	52.0%	53.6%	55.2%	56.8%

Utilization Overview (zSeries)									
8-way, 600 GB HDDs (IBM-00001 replacement)			**Total I/O Rate (I/Os per second)**						
Utilizations	Amber Threshold	Red Threshold	51000	81000	106000	111000	116000	121000	126000
Average SMP	60%	80%	39.8%	62.4%	81.6%	85.4%	89.3%	93.1%	96.9%
Average Bus	70%	90%	27.2%	43.1%	56.4%	59.1%	61.7%	64.4%	67.1%
Average Logical Device	n/a	n/a	0.9%	1.4%	1.9%	2.0%	2.1%	2.2%	2.2%
Highest DA	60%	80%	14.9%	23.6%	30.9%	32.3%	33.8%	35.3%	36.7%
Highest HDD (FC/SAS)	60%	80%	9.7%	15.5%	20.2%	21.2%	22.1%	23.1%	24.0%
Average FICON HA	50%	70%	18.4%	29.2%	38.3%	40.1%	41.9%	43.7%	45.5%
Highest FICON Port	50%	70%	3.6%	5.7%	7.4%	7.8%	8.1%	8.5%	8.9%
Average Fibre HA	60%	80%	22.7%	36.0%	47.0%	49.3%	51.5%	53.7%	55.9%
Average PPRC Link	60%	80%	16.5%	26.2%	34.2%	35.9%	37.5%	39.1%	40.7%

Sample IntelliMagic Vision Use Cases

High Connect Time (z/OS)

Disk Storage System Dashboard [rating: 0.01]

for all Disk Storage Systems by Serial

Two days before the issue (January 30), all dashboard indicators were showing acceptable performance

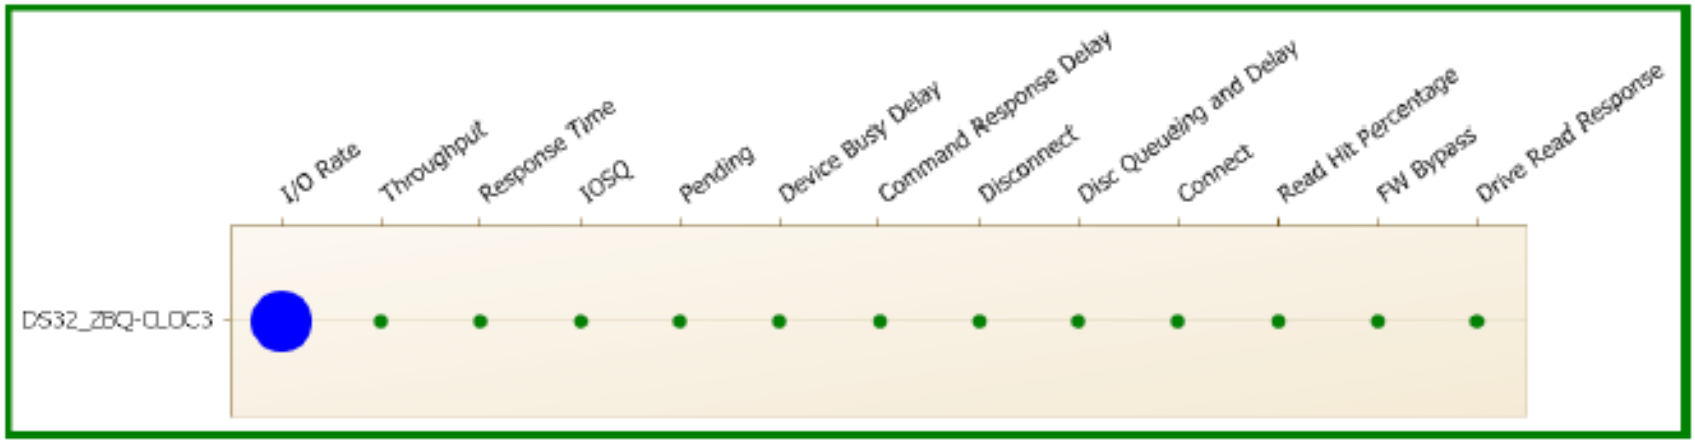

Disk Storage System Dashboard [rating: 0.18]

for all Disk Storage Systems by Serial

On January 31st, the Connect Time indicator changed to “yellow” which is an IntelliMagic Vision early warning.

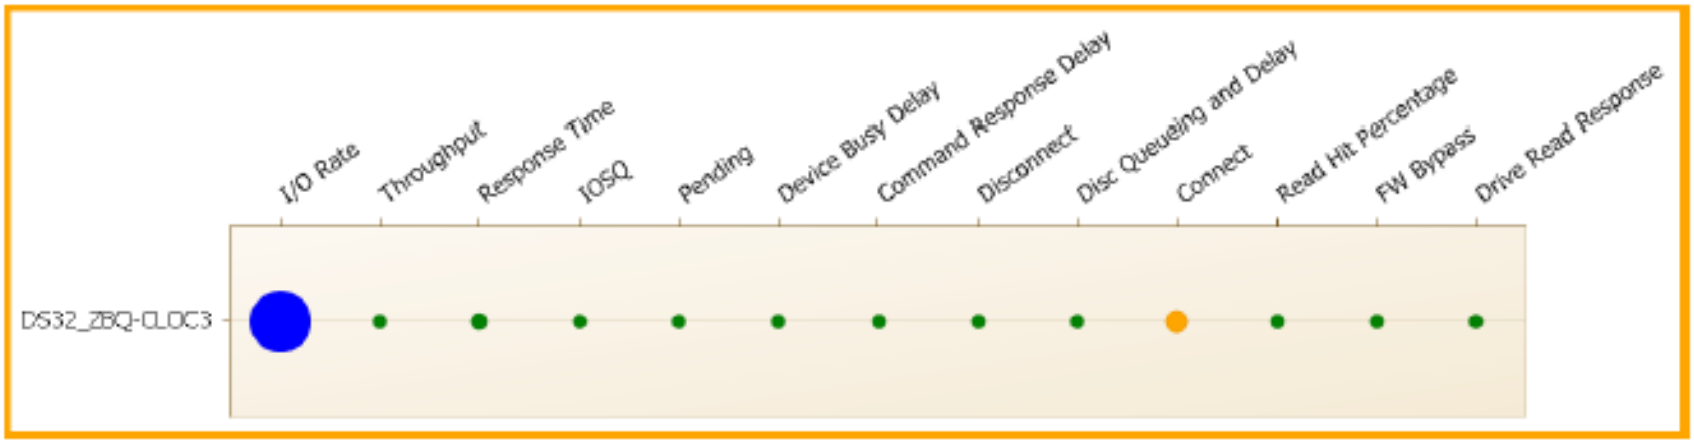

DSS Performance Summary

Charts behind the dashboard indicators

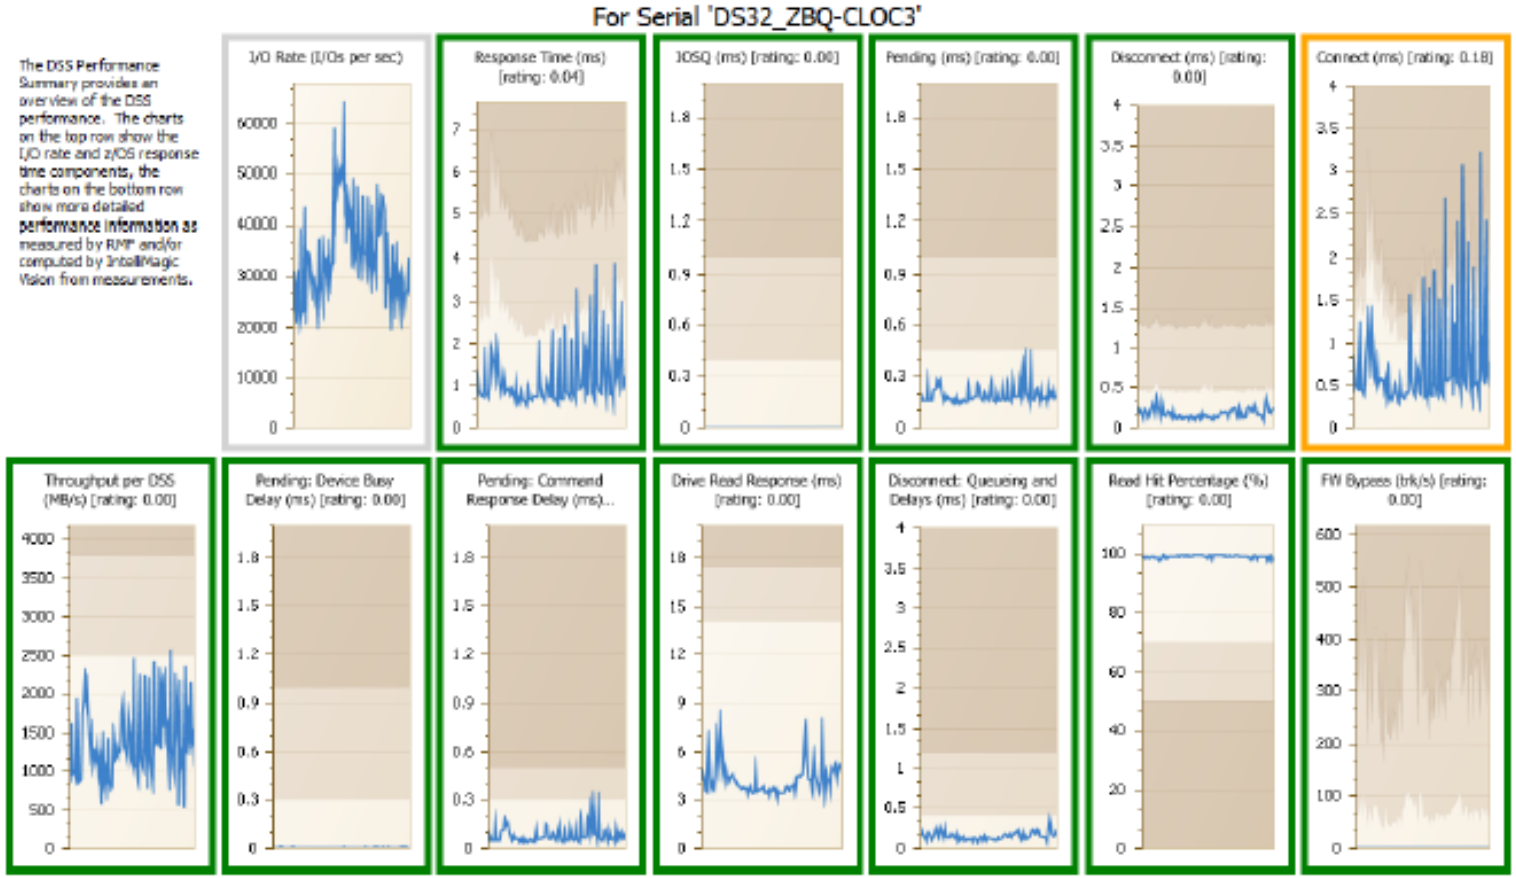

Disk Storage System Dashboard [rating: 0.41]

for all Disk Storage Systems by Serial

The next day, February 1st, showed that the Connect Time was now indicating a performance exception.

DSS Performance Summary

Charts behind the dashboard indicators

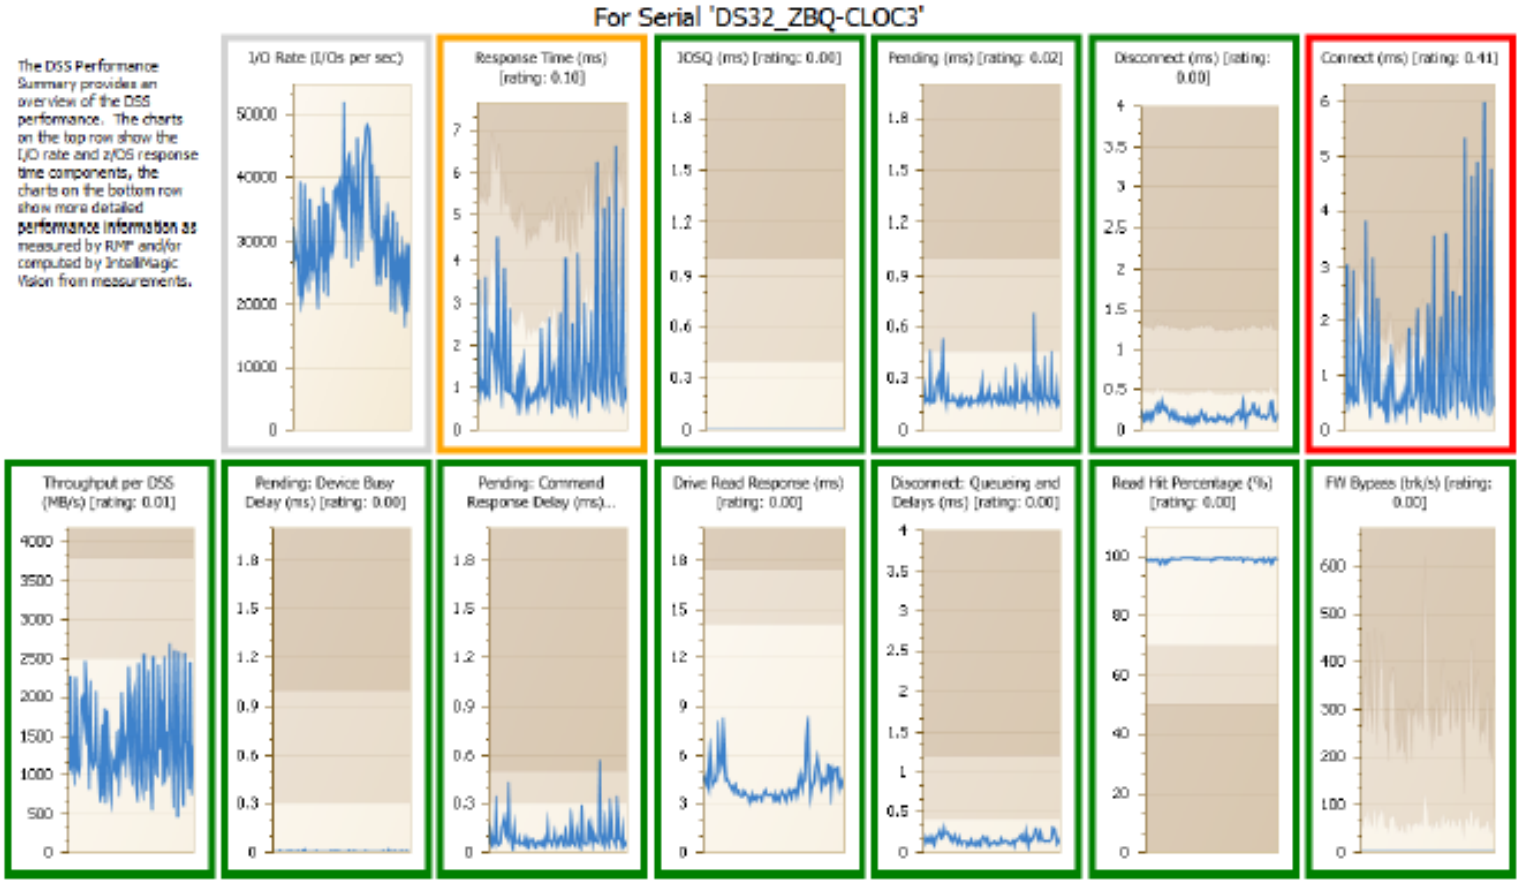

Response Time (ms) [rating: 0.10]

For Serial 'DS32_ZBQ-CLOC3'

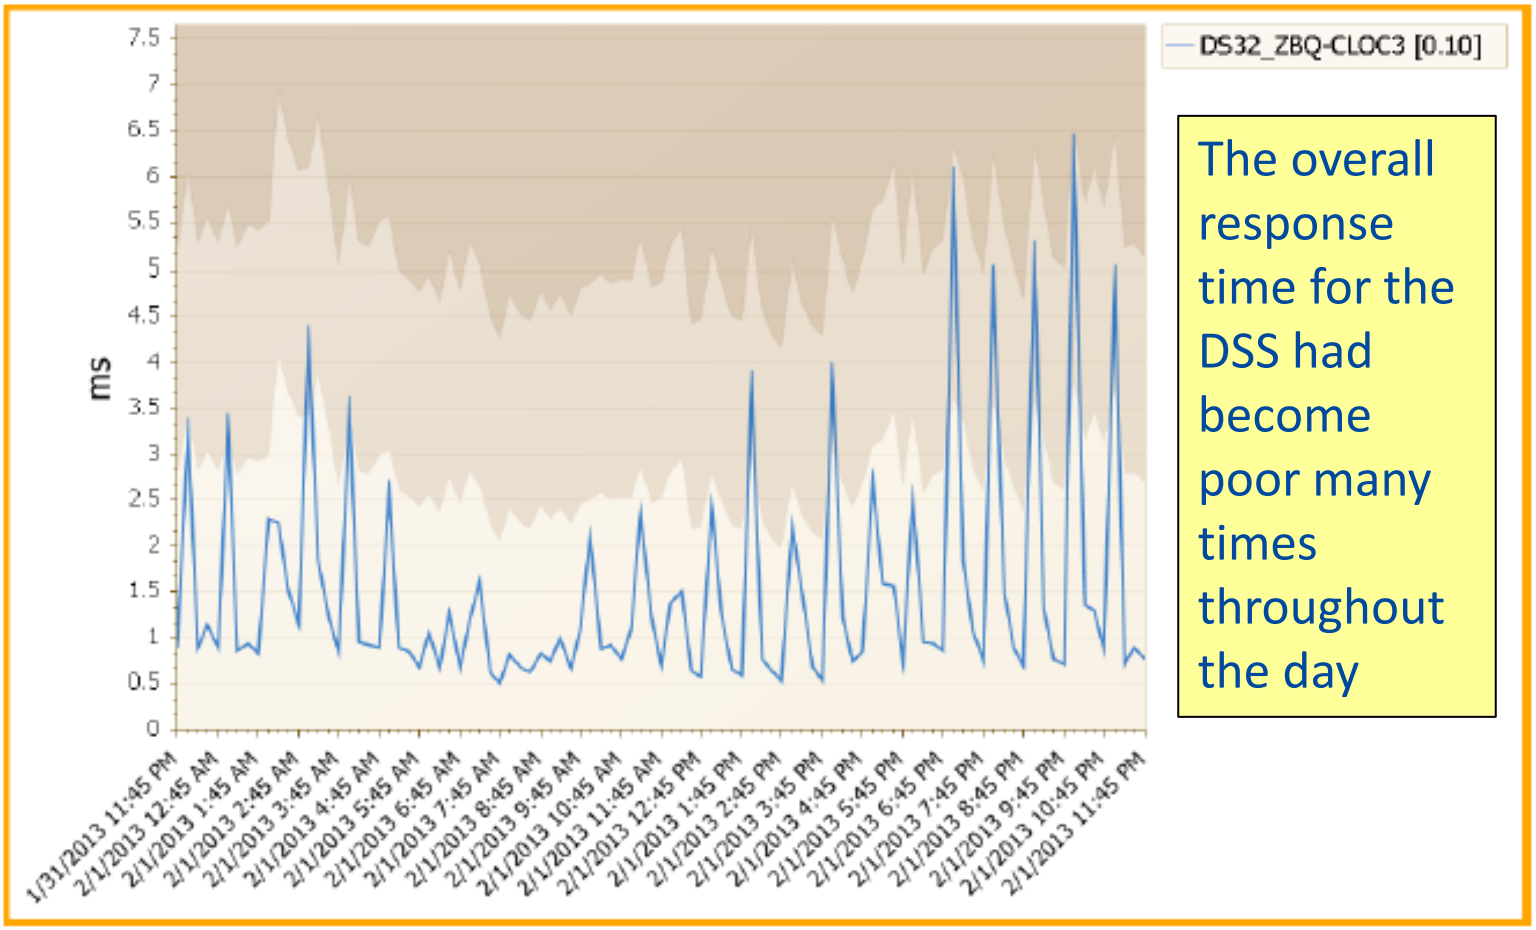

The overall response time for the DSS had become poor many times throughout the day

Connect (ms) [rating: 0.41]

For Serial 'DS32_ZBQ-CLOC3'

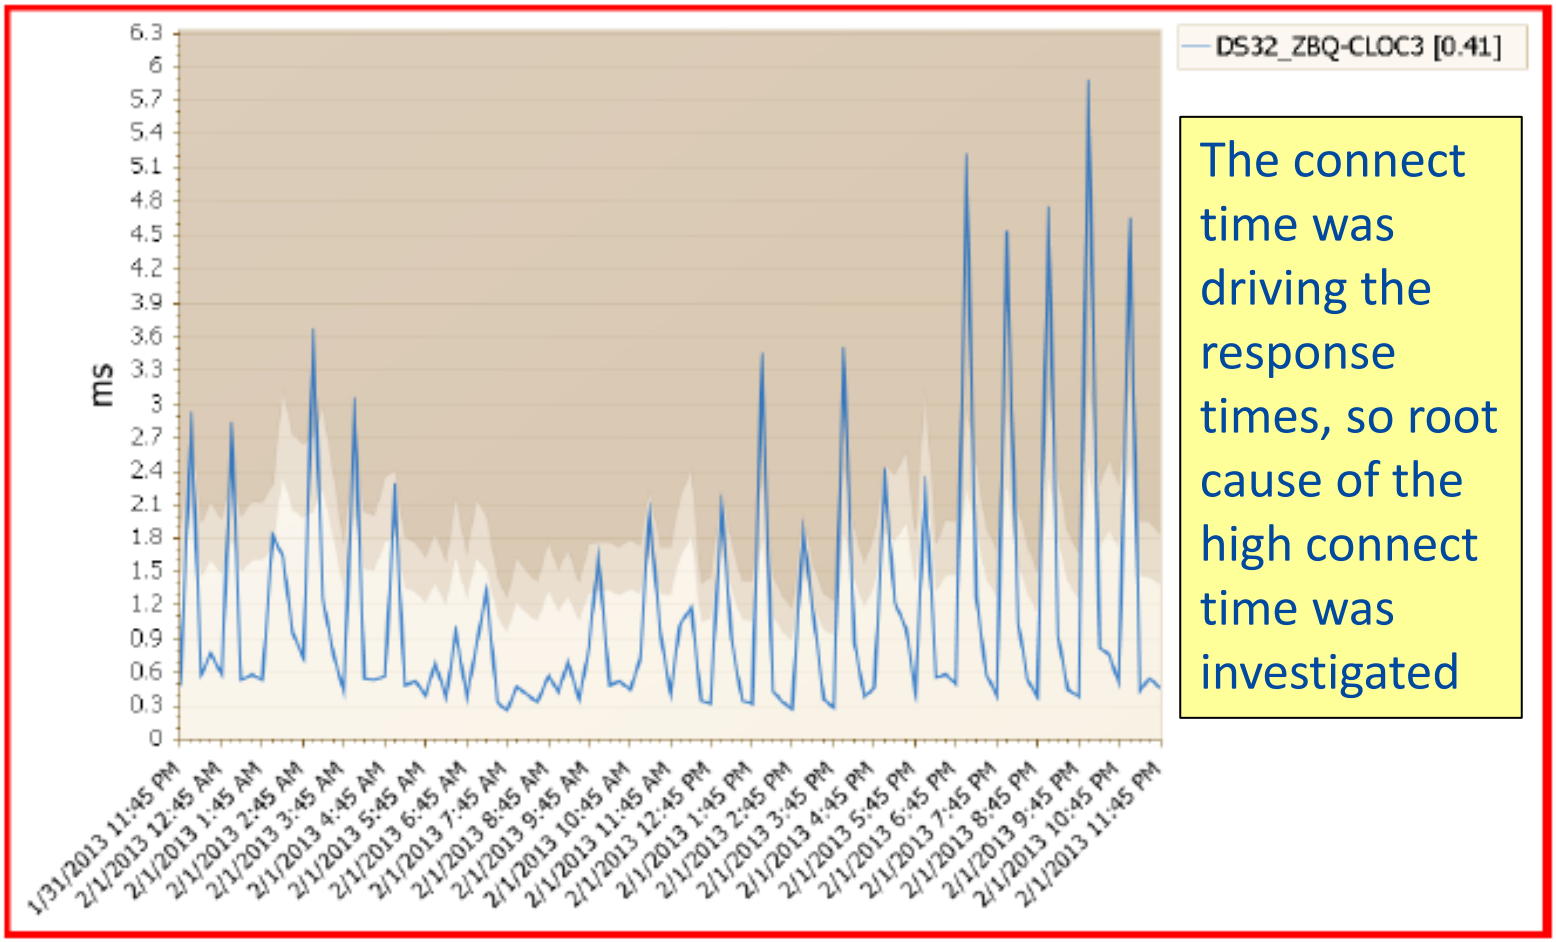

The connect time was driving the response times, so root cause of the high connect time was investigated

Connect (ms) (top 30) [rating: 0.88]

For Serial 'DS32_ZBQ-CLOC3' by Volume Label

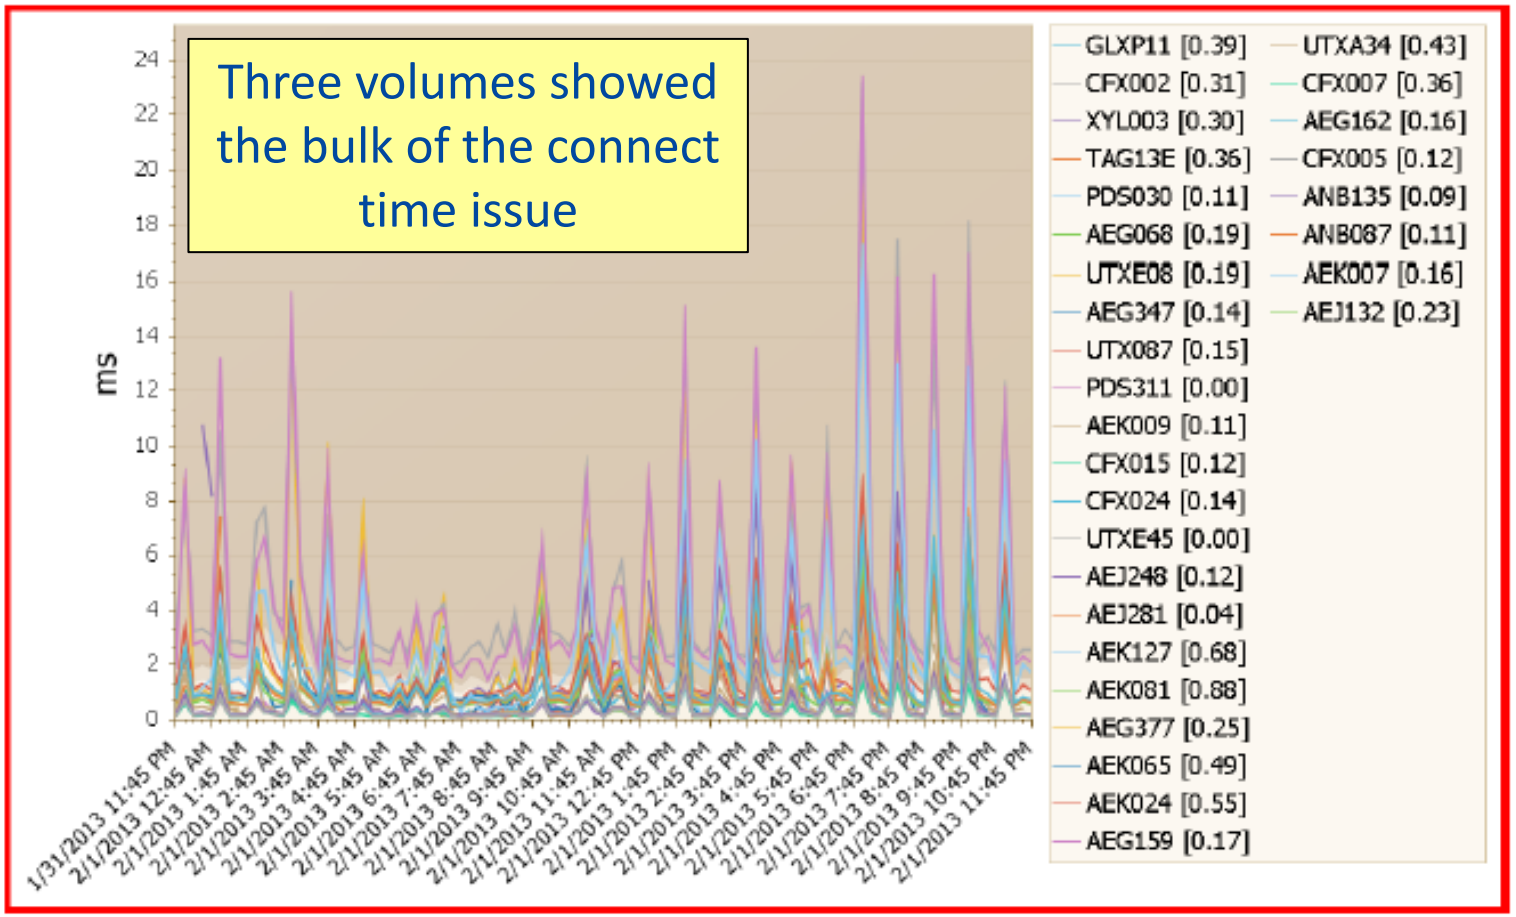

Connect (ms) (top 30) [rating: 0.78]

For Serial 'DS32_ZBQ-CLOC3' by Volume Label

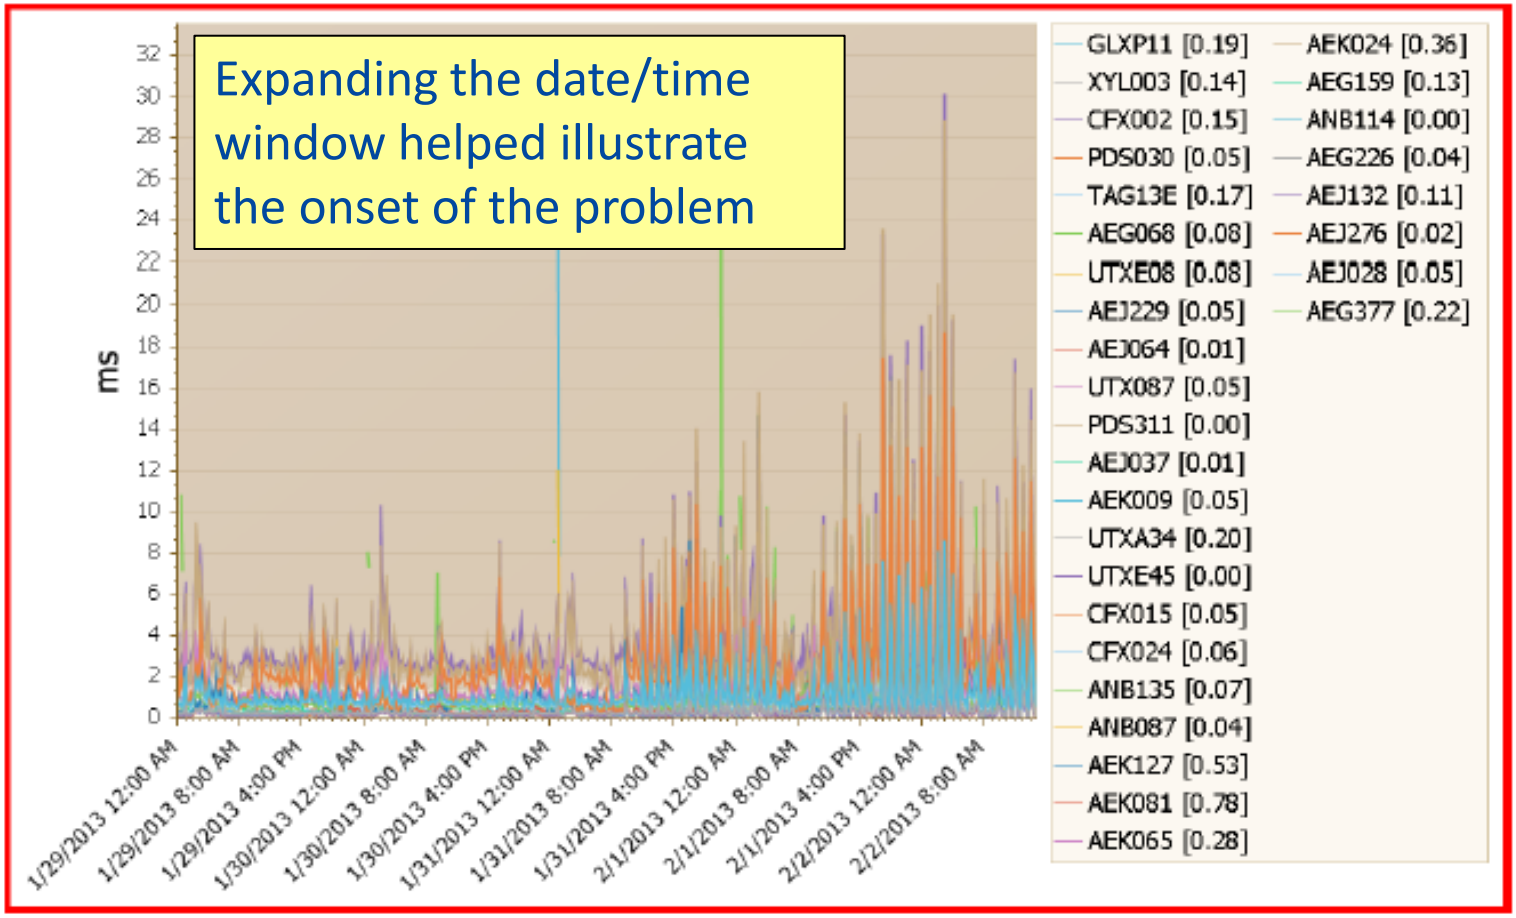

Response Time (ms) [rating: 0.66]

For Serial 'DS32_ZBQ-CLOC3'

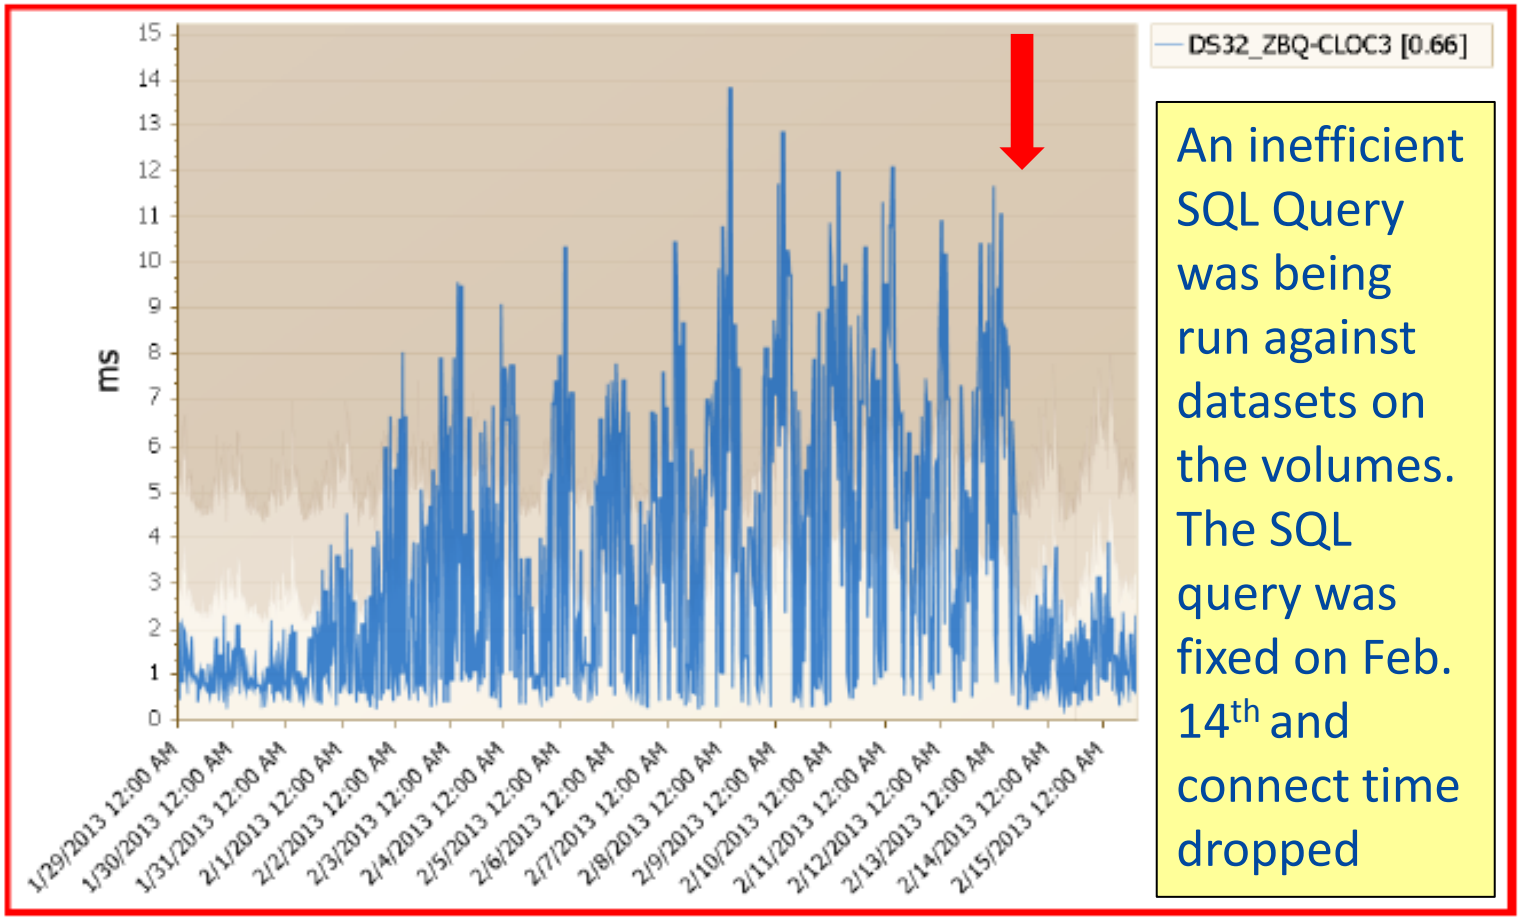

An inefficient SQL Query was being run against datasets on the volumes. The SQL query was fixed on Feb. 14th and connect time dropped

Disk Storage System Dashboard [rating: 0.08]

for all Disk Storage Systems by Serial

After the SQL query was fixed, the dashboard indicators returned to “green”

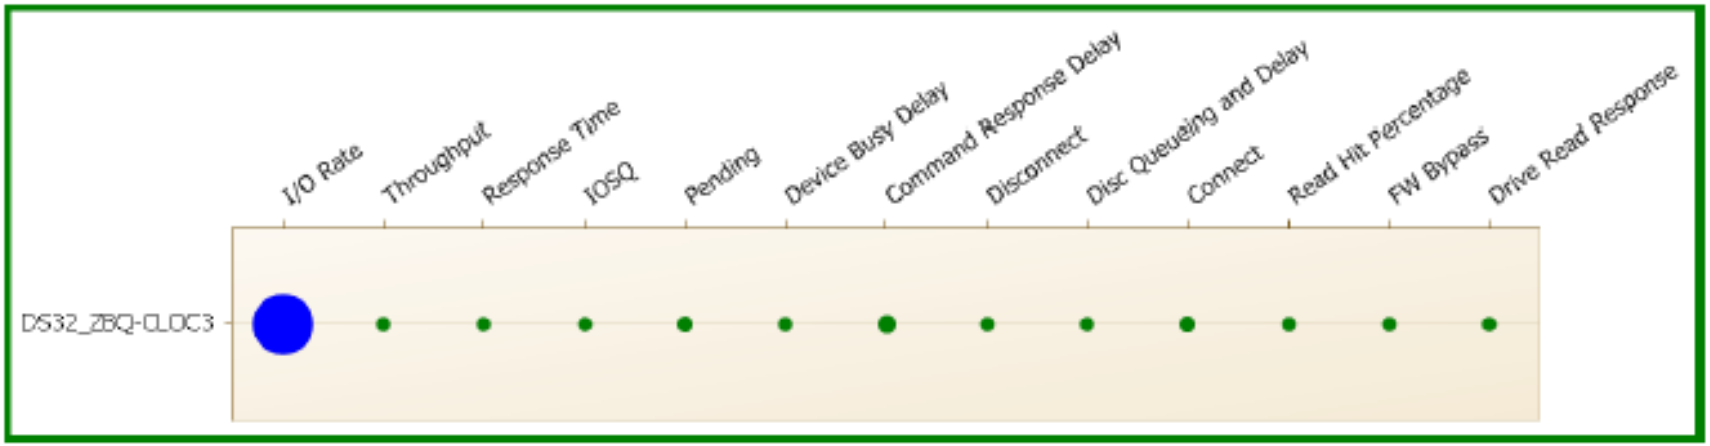

SRDF/A Performance (z/OS)

SRDF/A Dashboard [rating: 0.11] for all SRDF/A Sessions by Session

Early warning for cycle time

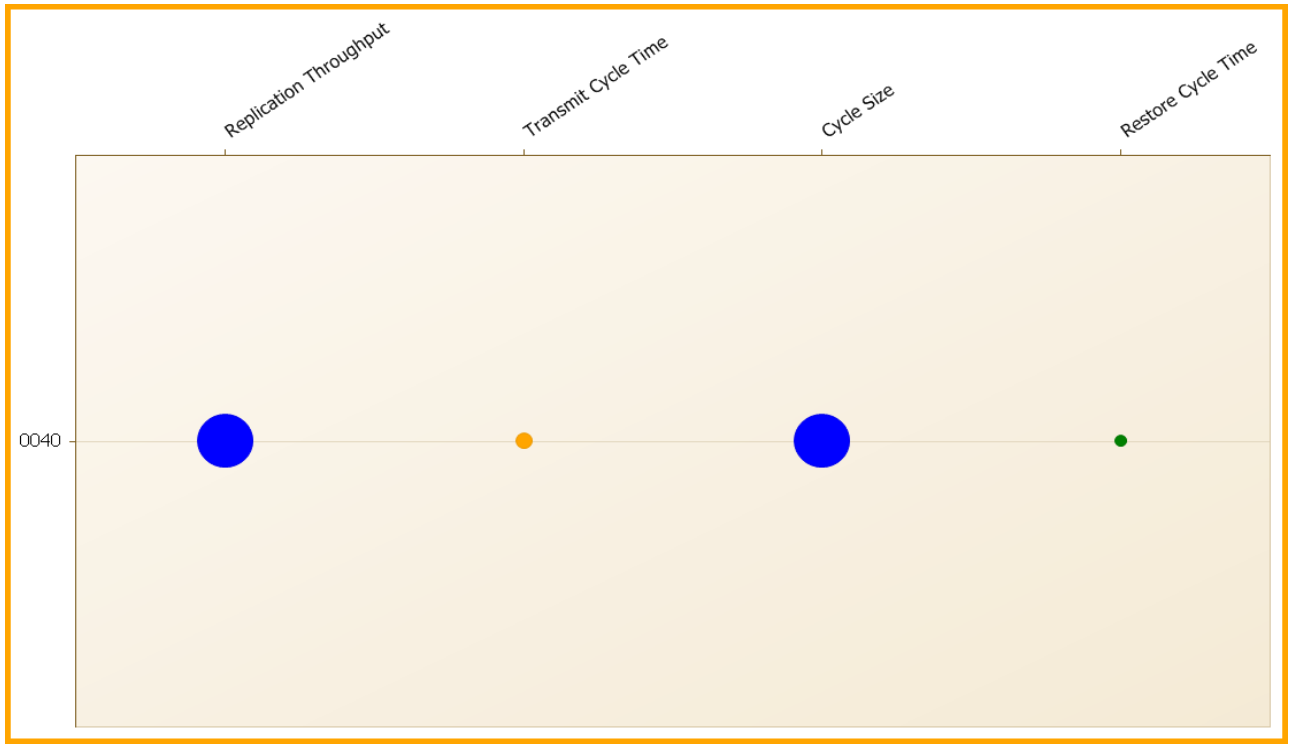

Max Cycle Time (sec) [rating: 0.12] For Session '0040'

Asynchronous replication was struggling during some intervals on April 29th

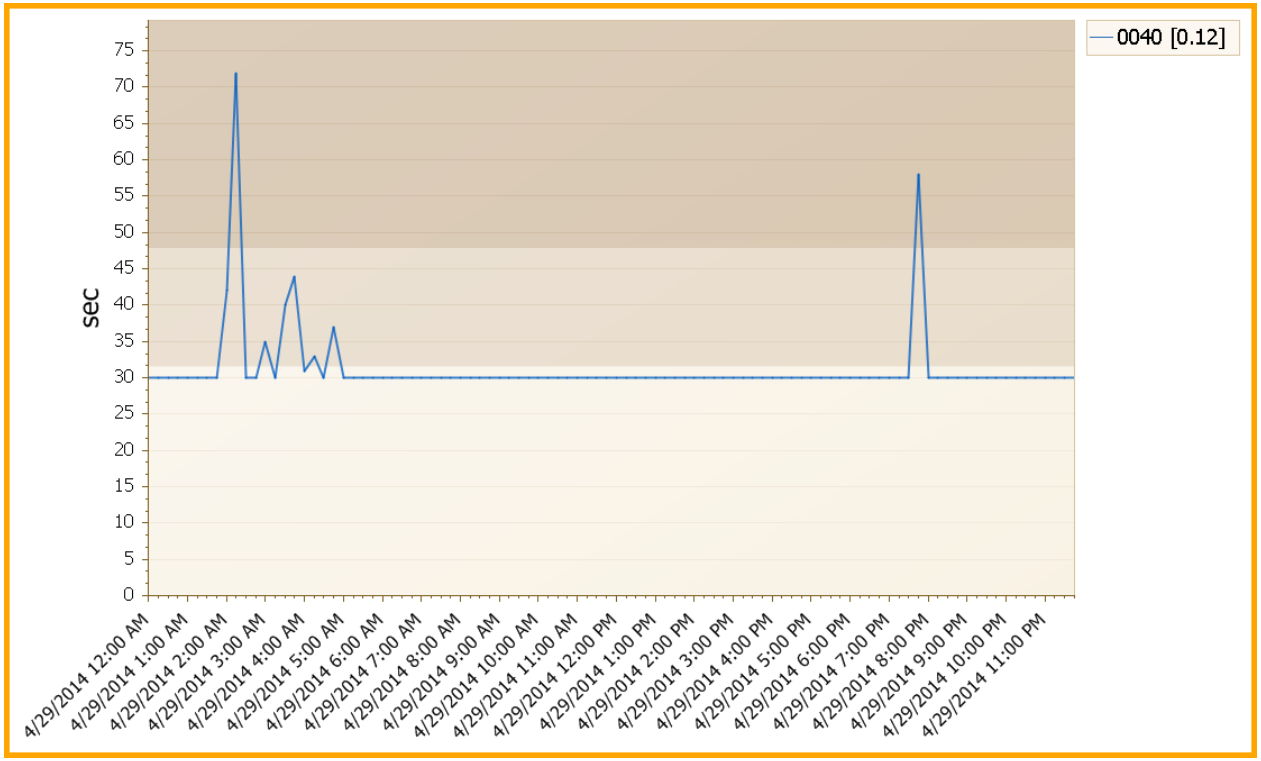

Max Cycle Time (sec) For Session '0040'

Two days later, SRDF/A dropped (May 1)

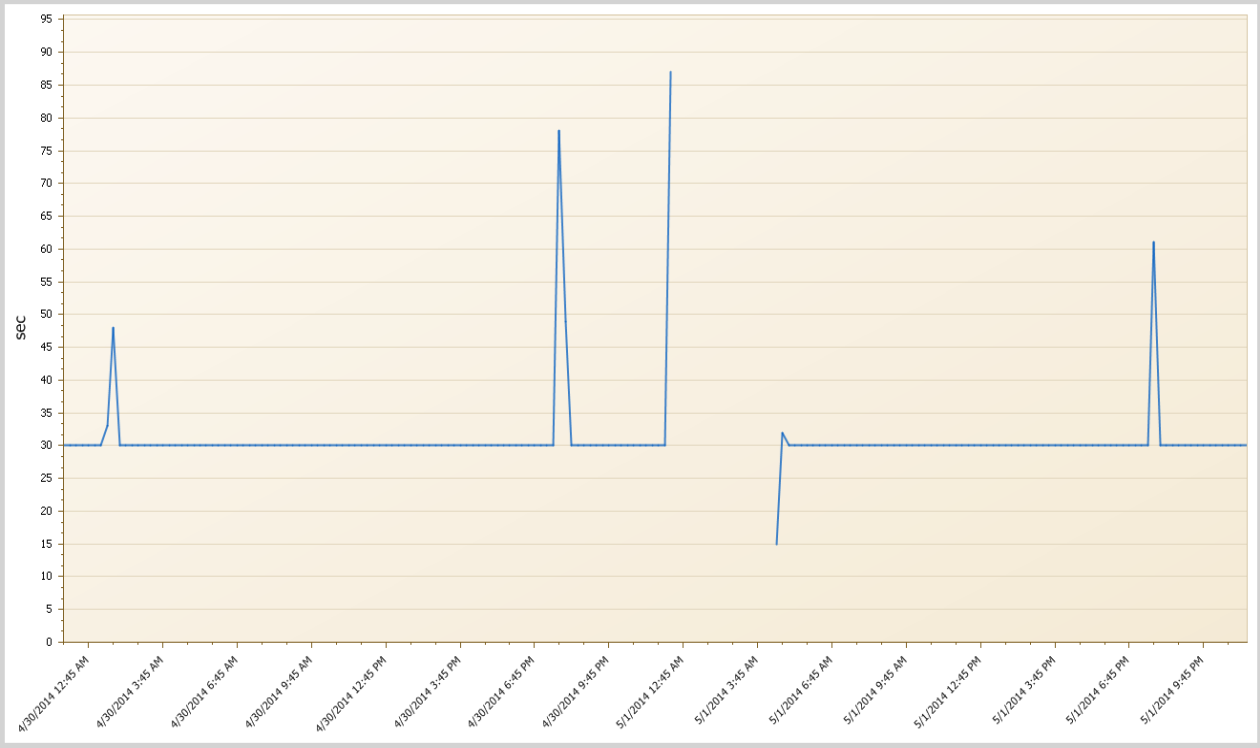

Max Cycle Size (slots)

for all SRDF/A Sessions by Session

The trend for this metric was that the spikes kept worsening

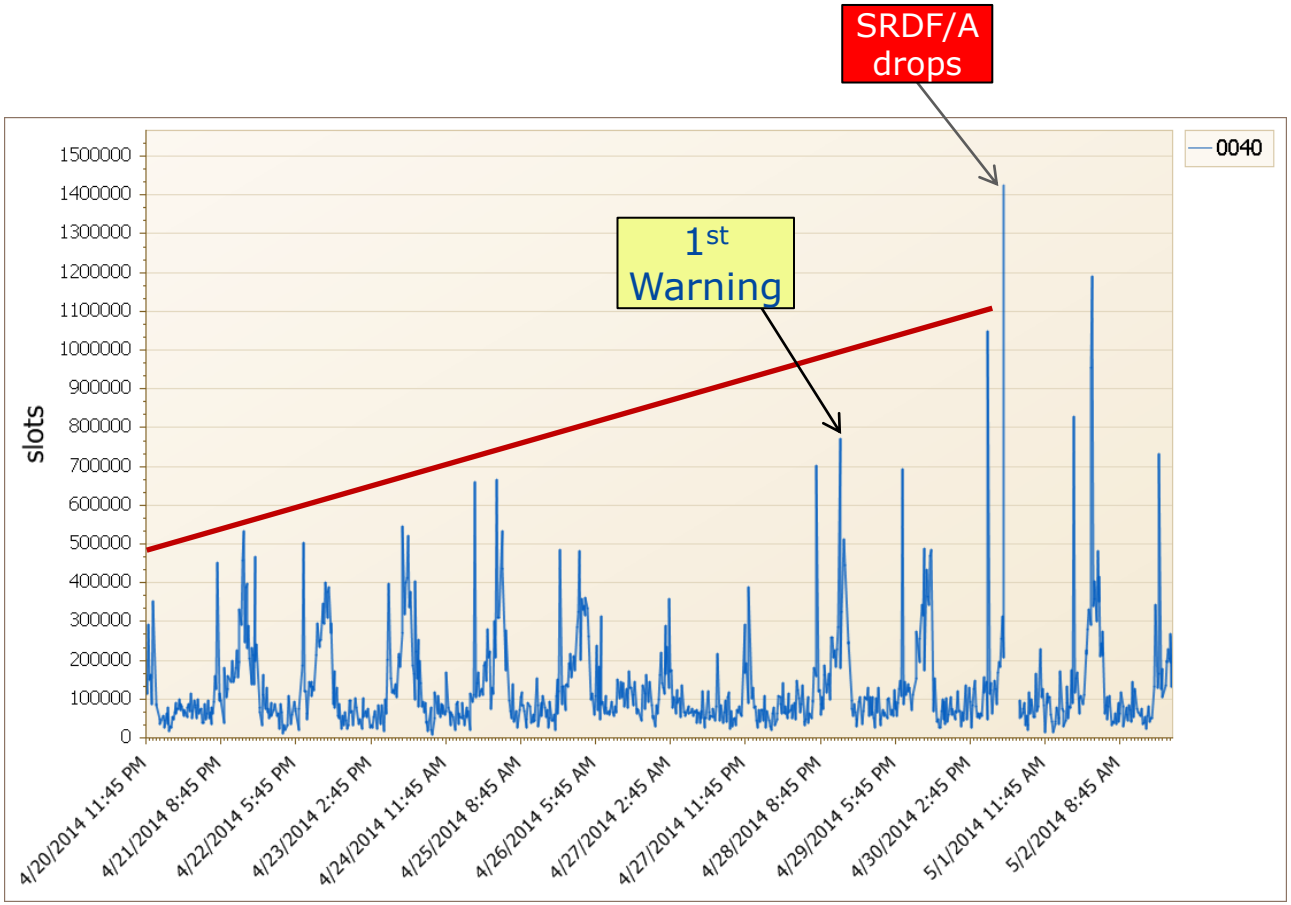

HA Writes (tracks) for all SRDF/A Sessions by Serial

Right before the SRDF/A drop, a burst of write activity came from one VMAX

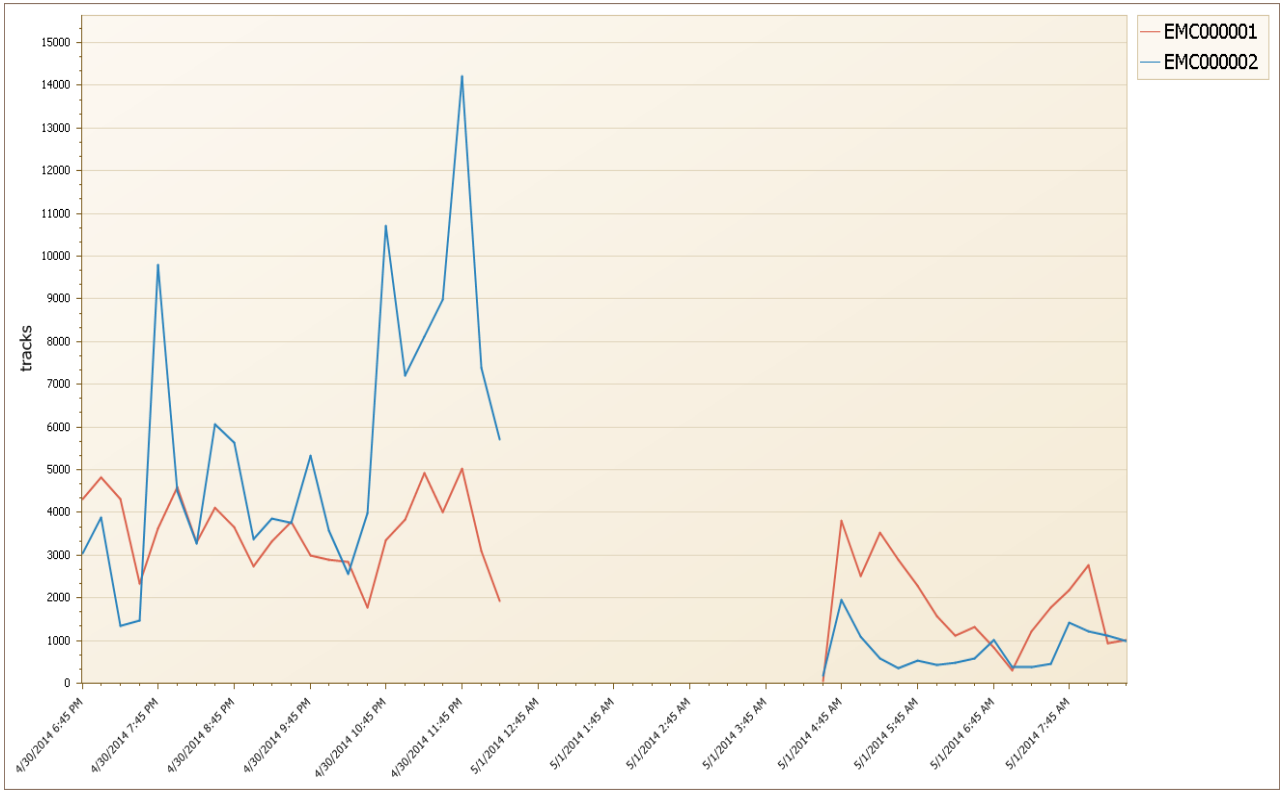

Write Pending High Water Mark (slots)

for all SRDF/A Sessions by Serial

The high water mark for SRDF/A write pending slots spiked.

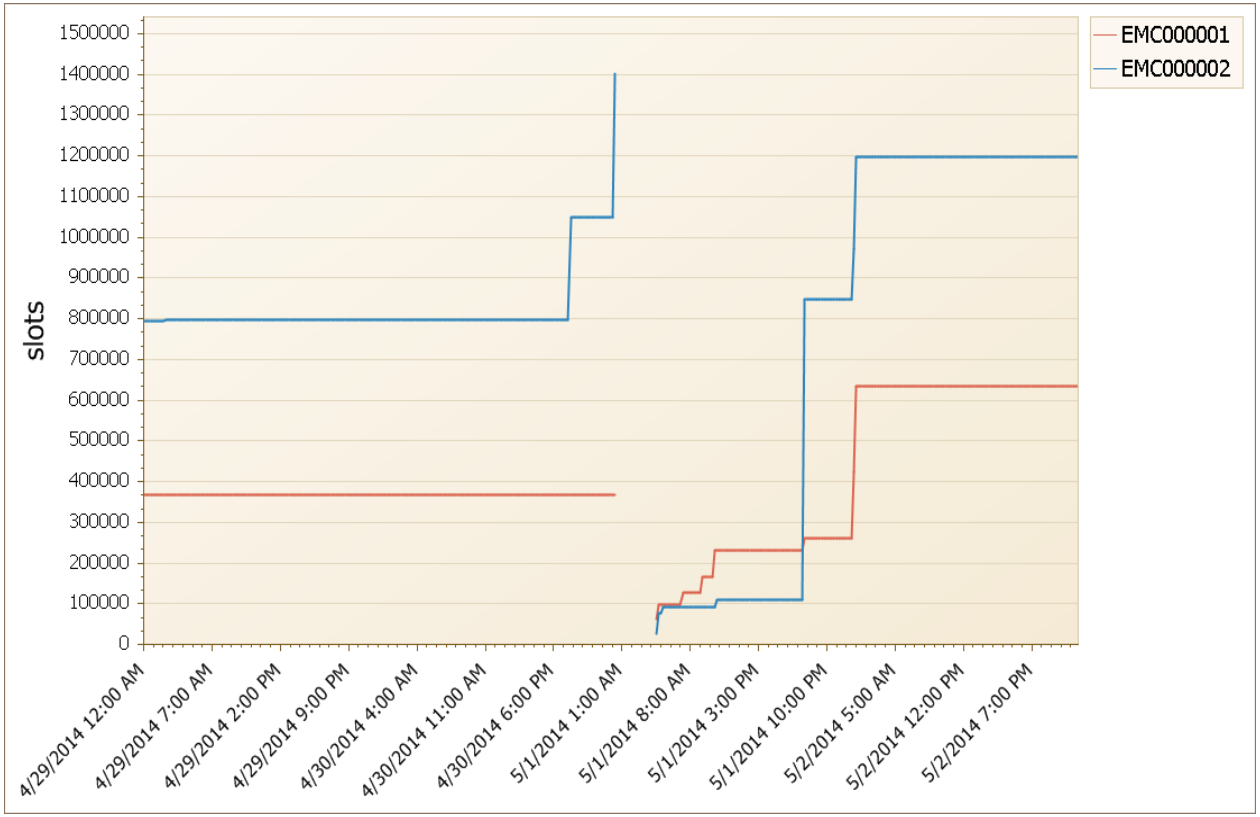

FICON Write (MB/s)

for all Ports by Serial

The writes were very imbalanced between the two VMAX's

Rebalancing could fix the problem.

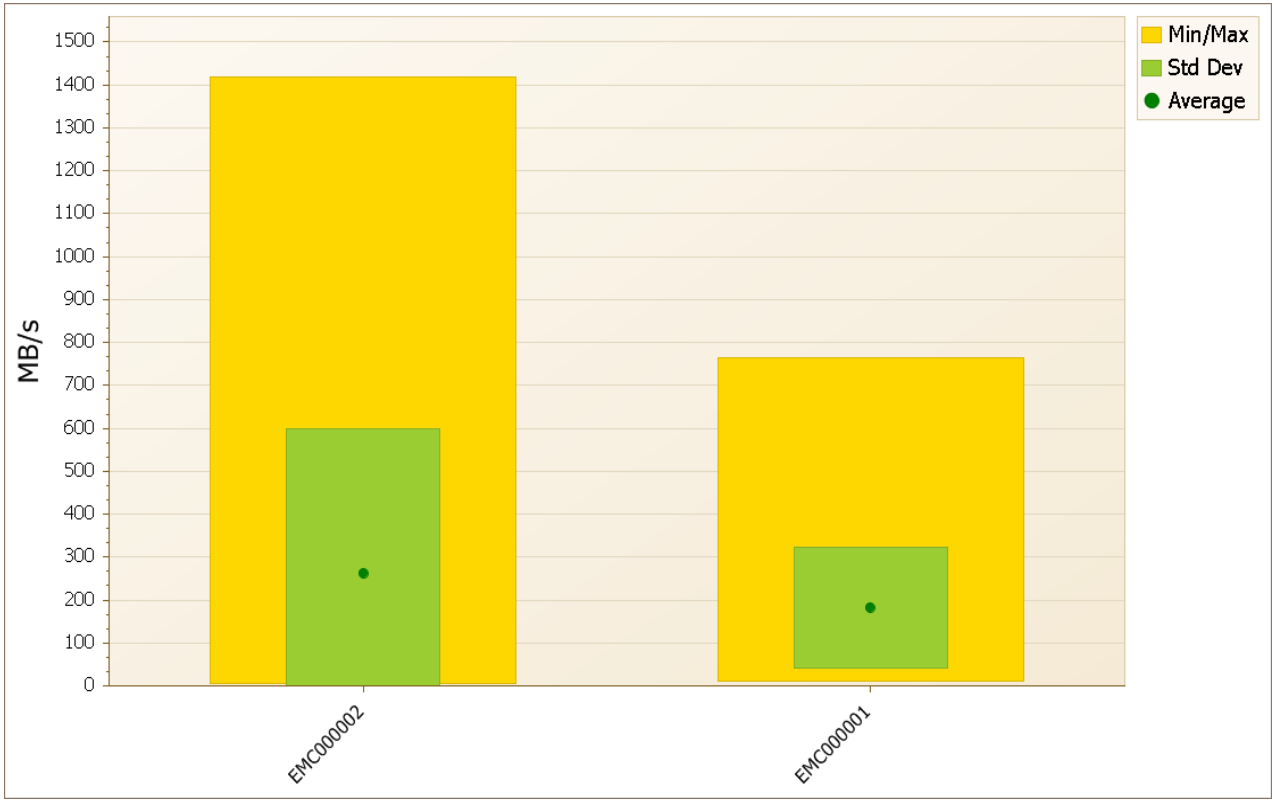

Write (LR/s) (top 20)

for all Storage Groups by Storage Group

The storage groups responsible for most of the writes are easily identified in IntelliMagic Vision

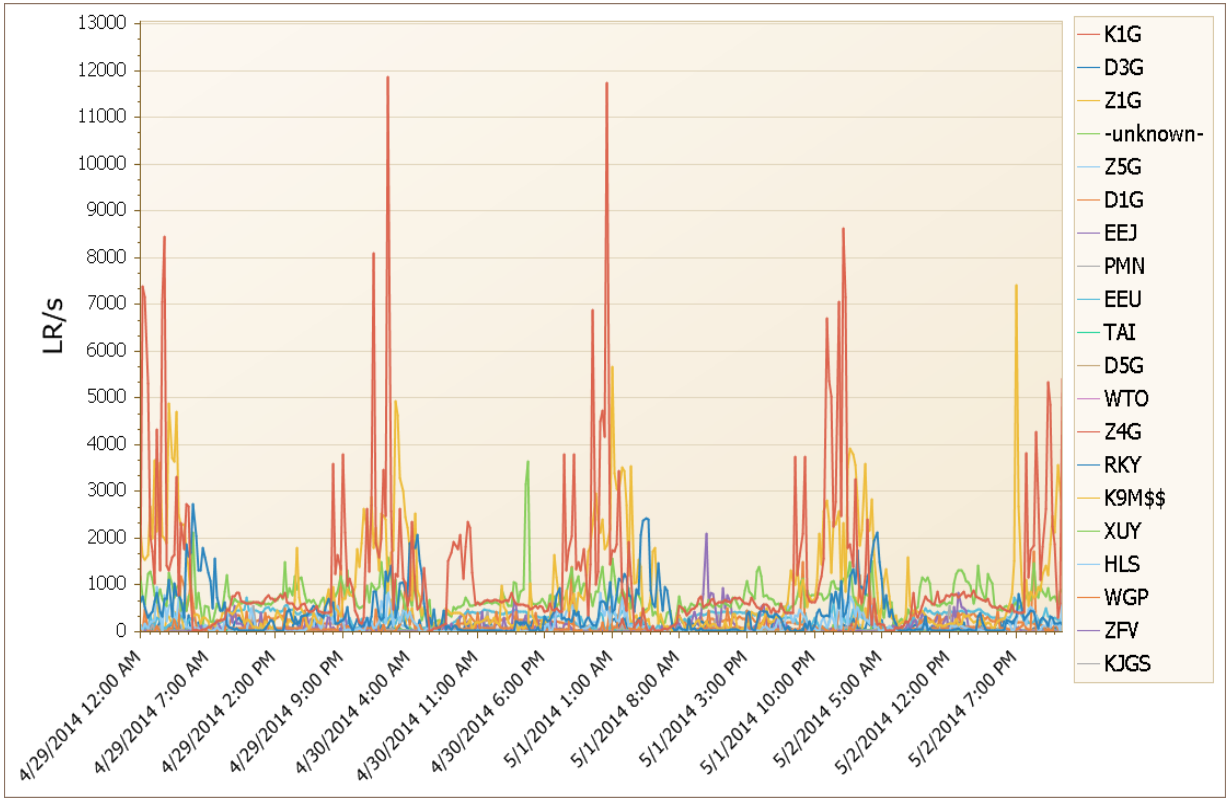

Drives Overloaded (SAN)

Front-End Dashboard [rating: 2.04] for all Storage Pools by Serial

Rating based on DSS Storage Pool data using DSS Thresholds

Response times are generally bad!

SVC Storage Pool Front-End Dashboard [rating: 2.04]
For Serial 'V7000-1' by Storage Pool

Rating based on DSS Storage Pool data using DSS Thresholds

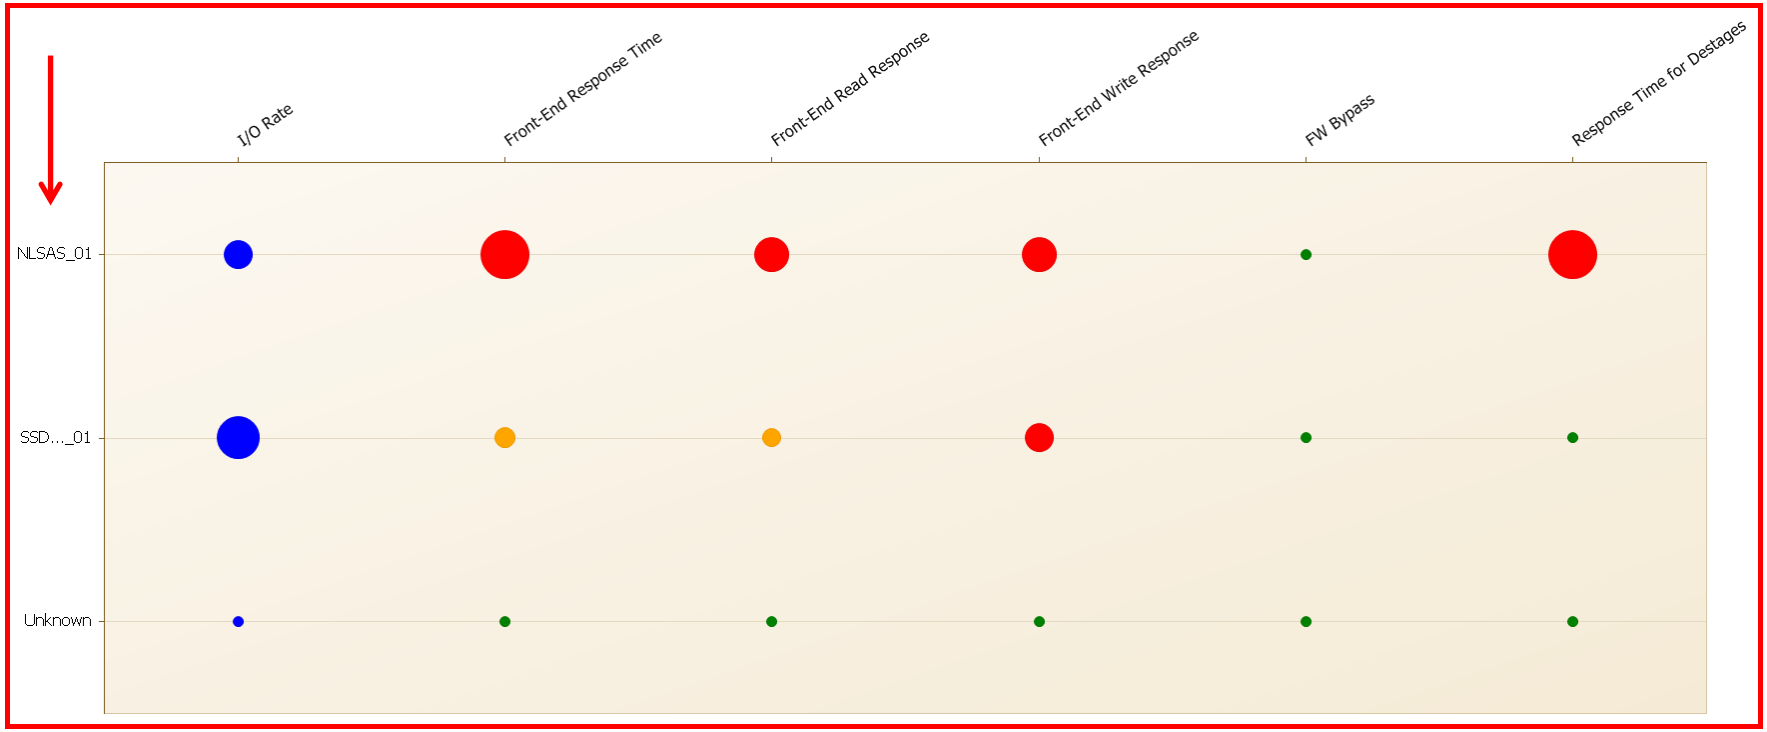

Response Times are Primarily Affecting
NLSAS_01 Storage Pool

SVC Storage Pool minicharts

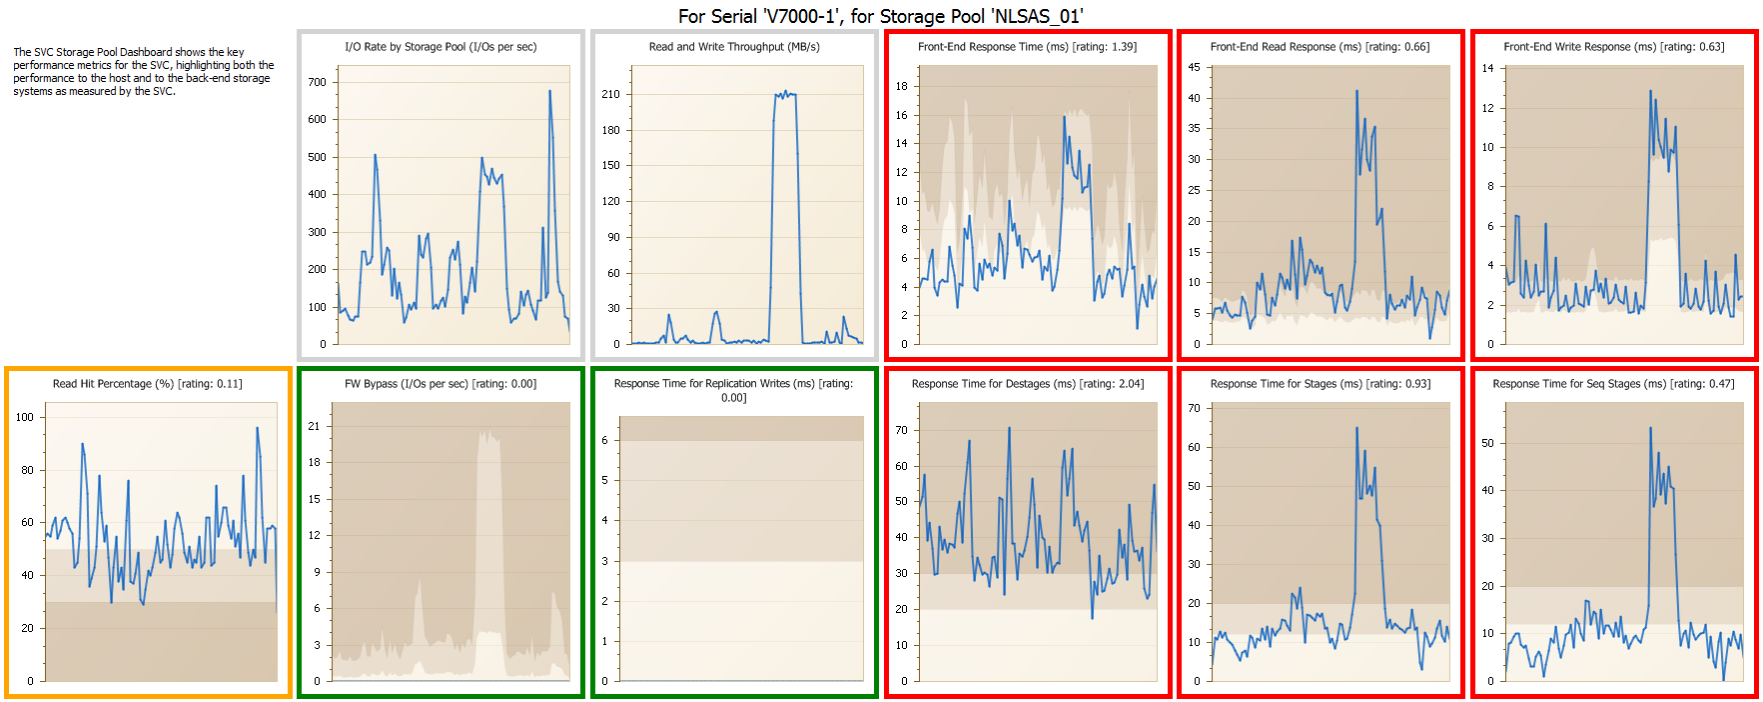

Response times are bad on front-end and back-end (stages/destages)

Drive Read Response time (ms) [rating: 2.33]
For Serial 'V7000-1' by Drive Tier

Rating based on DSS Drives data using DSS Thresholds

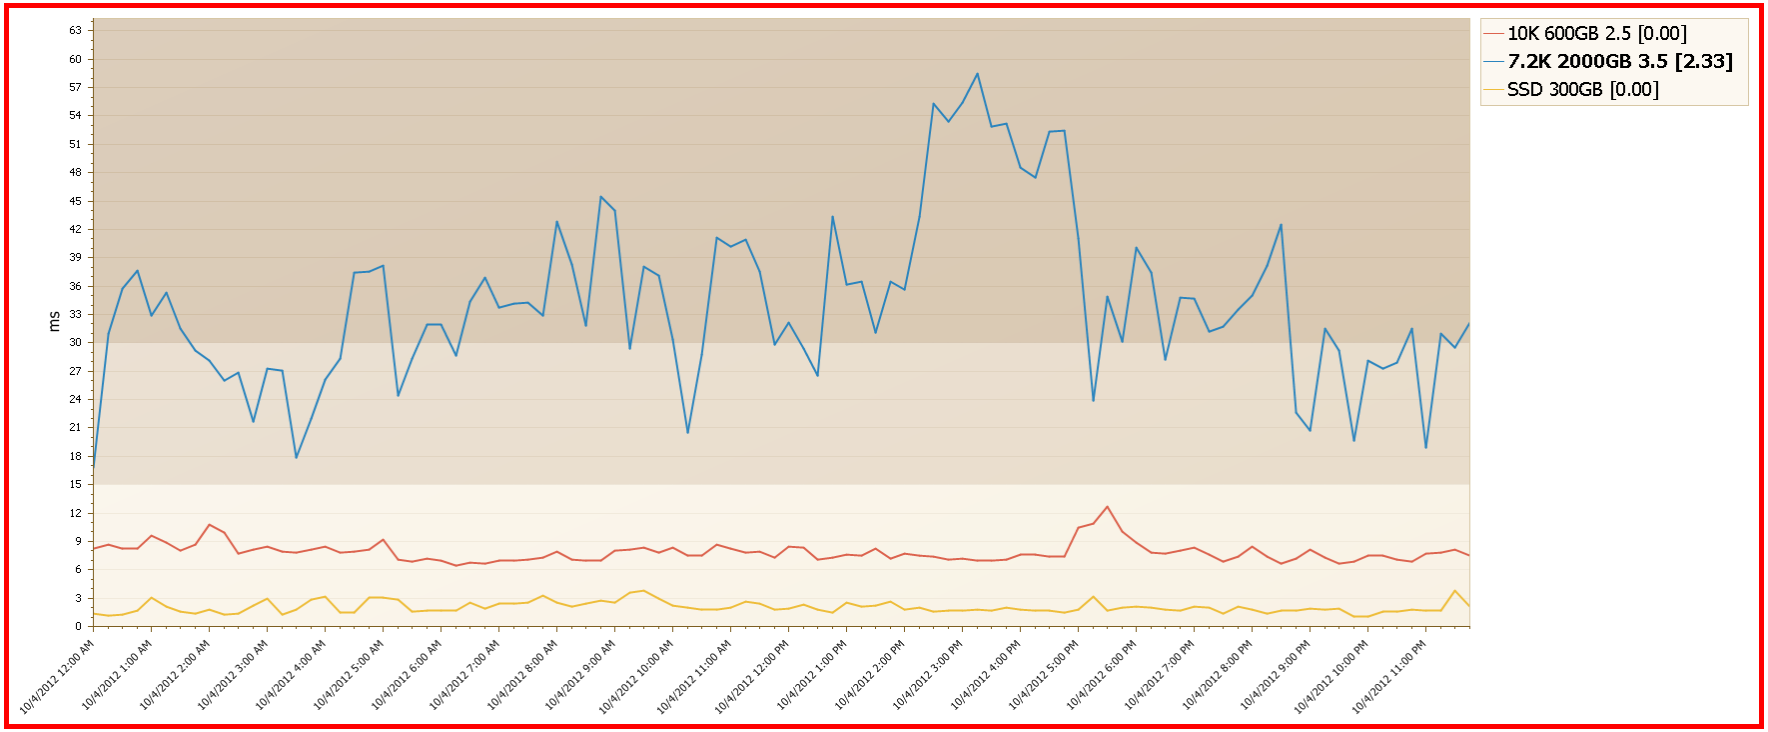

Disk response times are only bad for 7.2K 2TB Drives

Drive Read Response time rated by Drive Tier (ms) [rating: 1.98] For Serial 'V7000-1', for Drive Tier '7.2K 2000GB 3.5', for RAID Group ID 'TBIMO0' by Drive Name

Rating based on DSS Drives data using Drive Tier based Thresholds

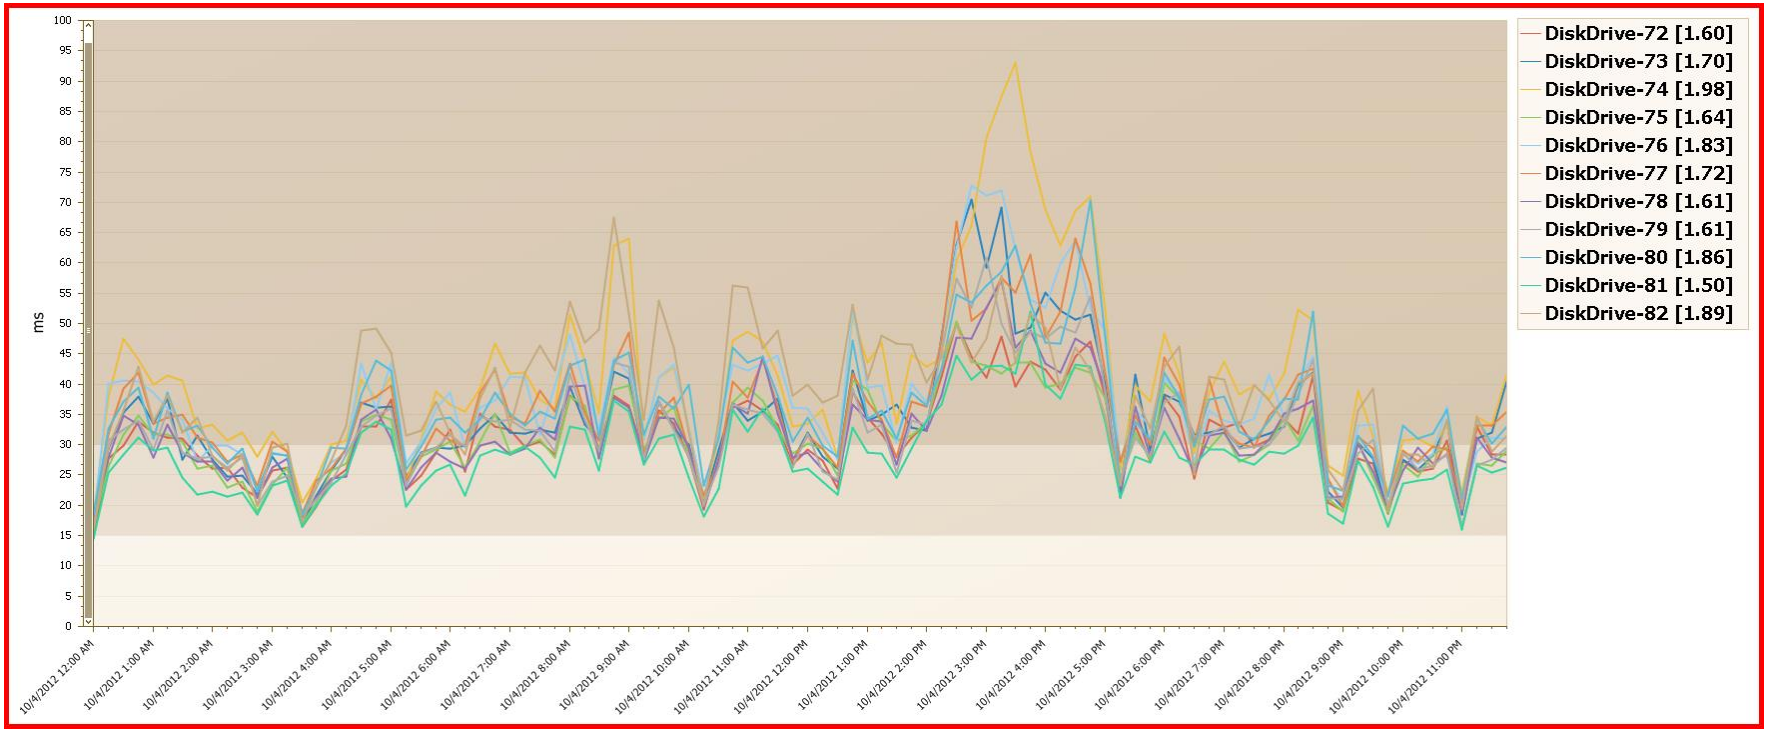

Focusing on a one RAID Group we see all drives have elevated response time.

V7000 Case Study Summary

Finding	Observations
1. Poor front-end response times from 2:00 PM to 5:00 PM due to overloaded 2 TB drives.	1. Spread the load over additional RAID groups.

IBM DS8000 Overloaded Host Adaptors (SAN)

Disk Storage System Dashboard [rating: 0.61]

for all Disk Storage Systems by Serial

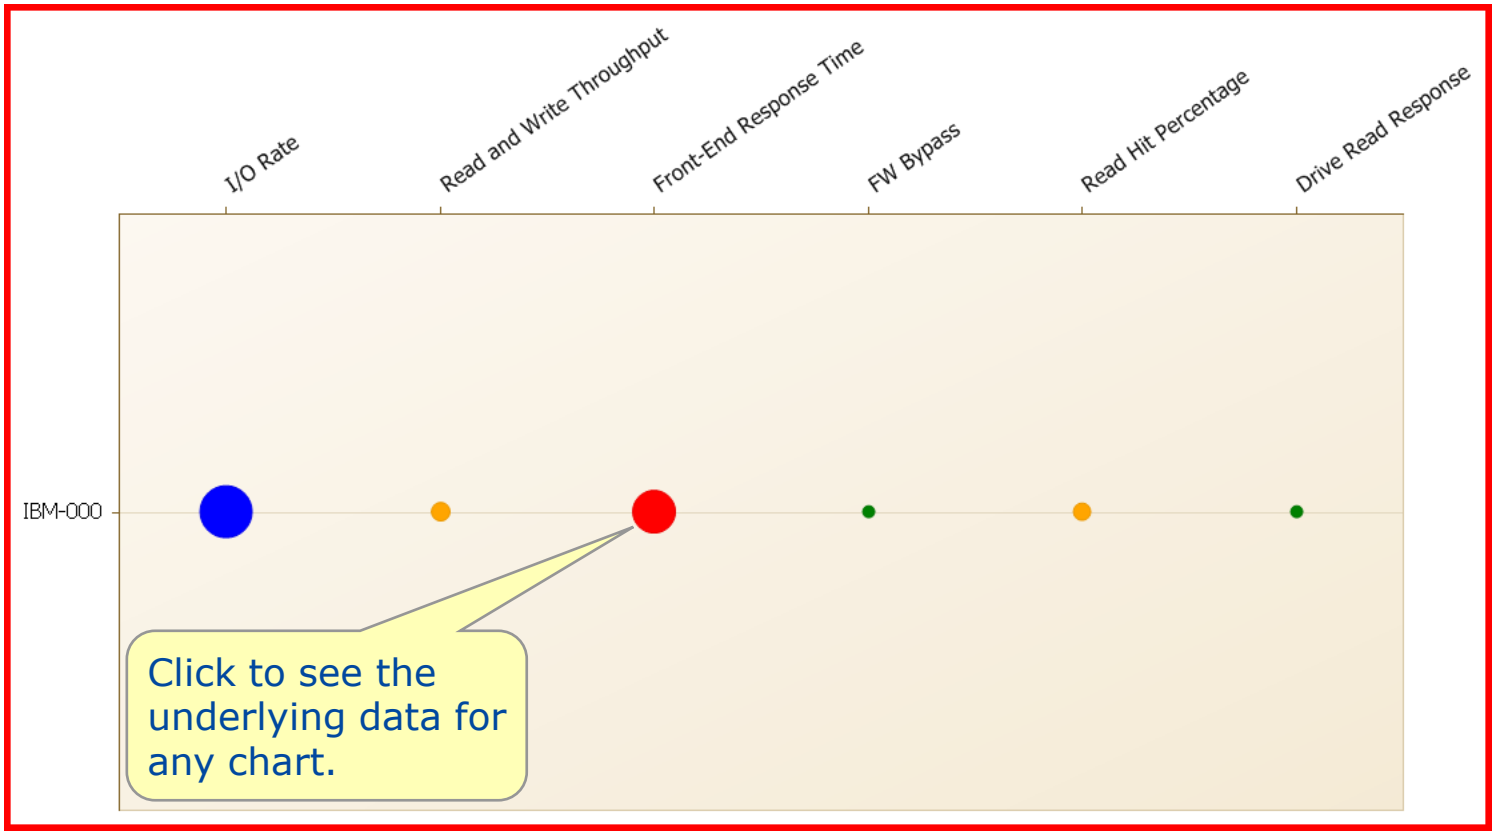

Rated dashboard indicates key risk indicators for performance attributes exceeding warning and exception thresholds.

DSS Dashboard Mini-Charts

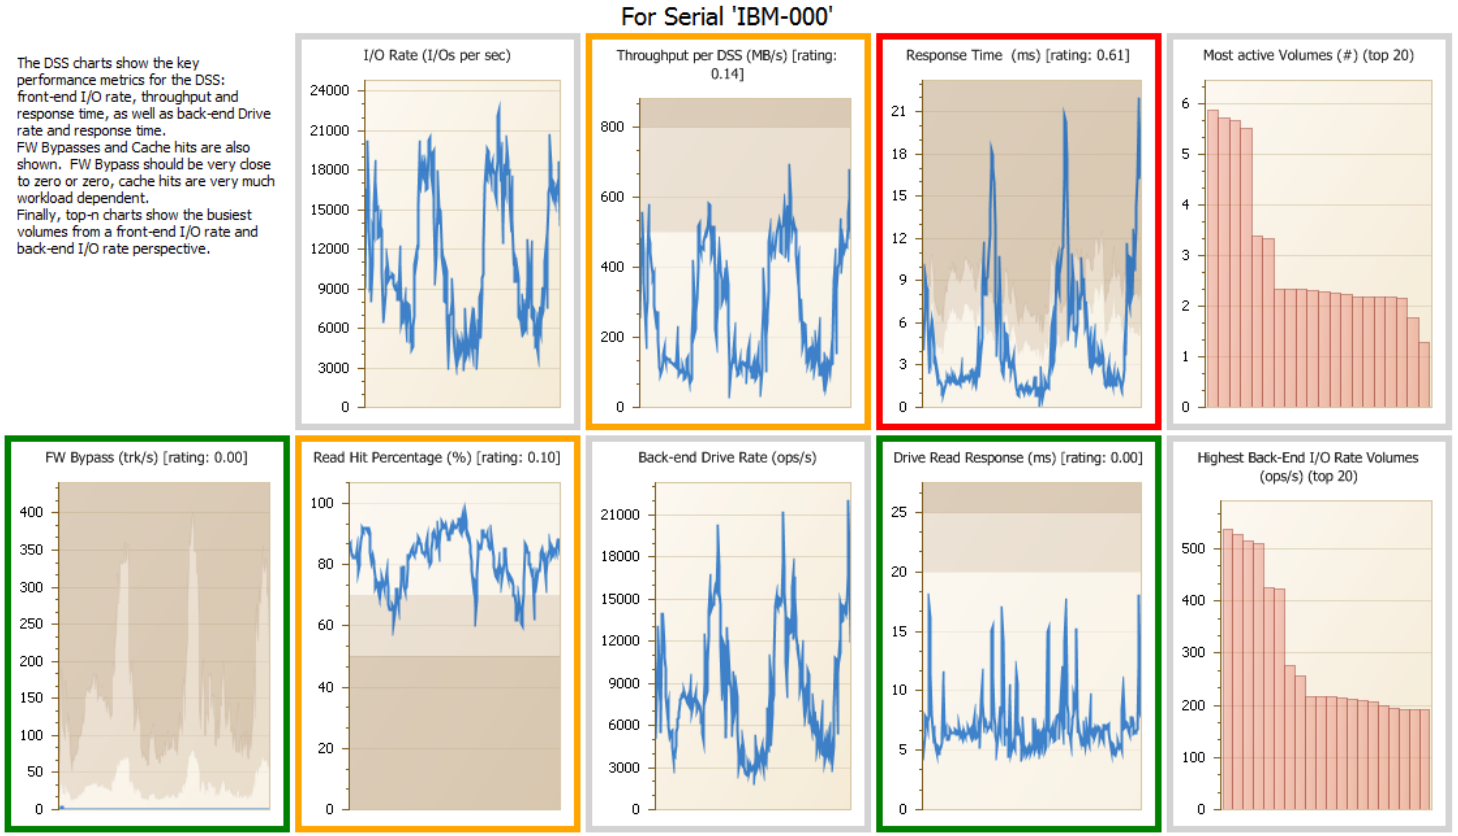

Drilldown shows rated mini-line charts that help you visualize relationships. Note how throughput, response time, and Read Hit % are correlated.

Throughput and Response Time

for all data

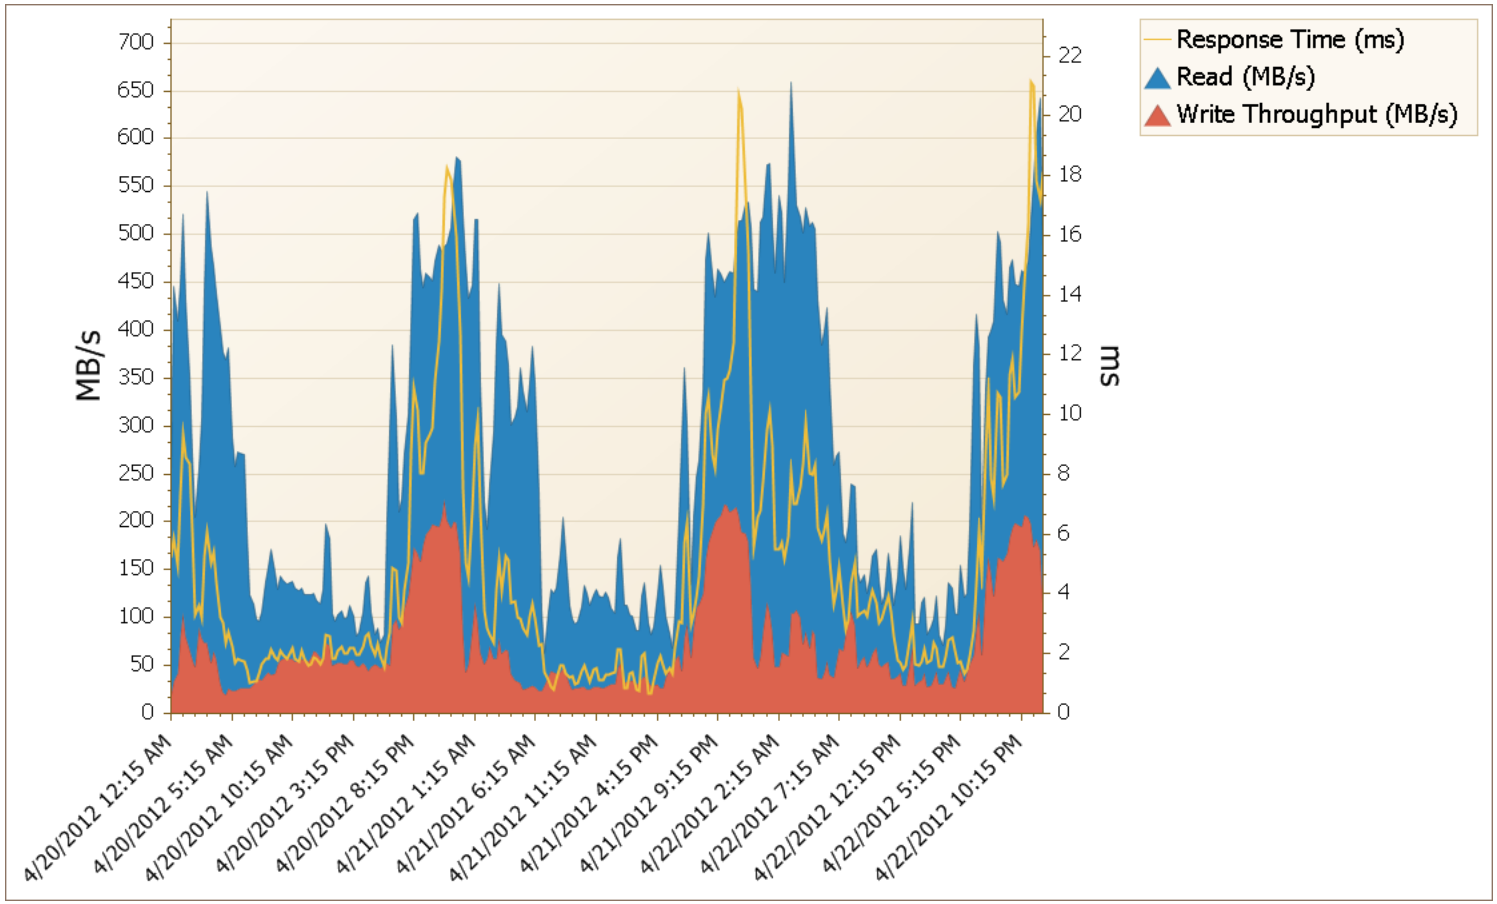

Strong correlation between throughput and response time.

Response Time (ms) [rating: 0.61]

For Serial 'IBM-000'

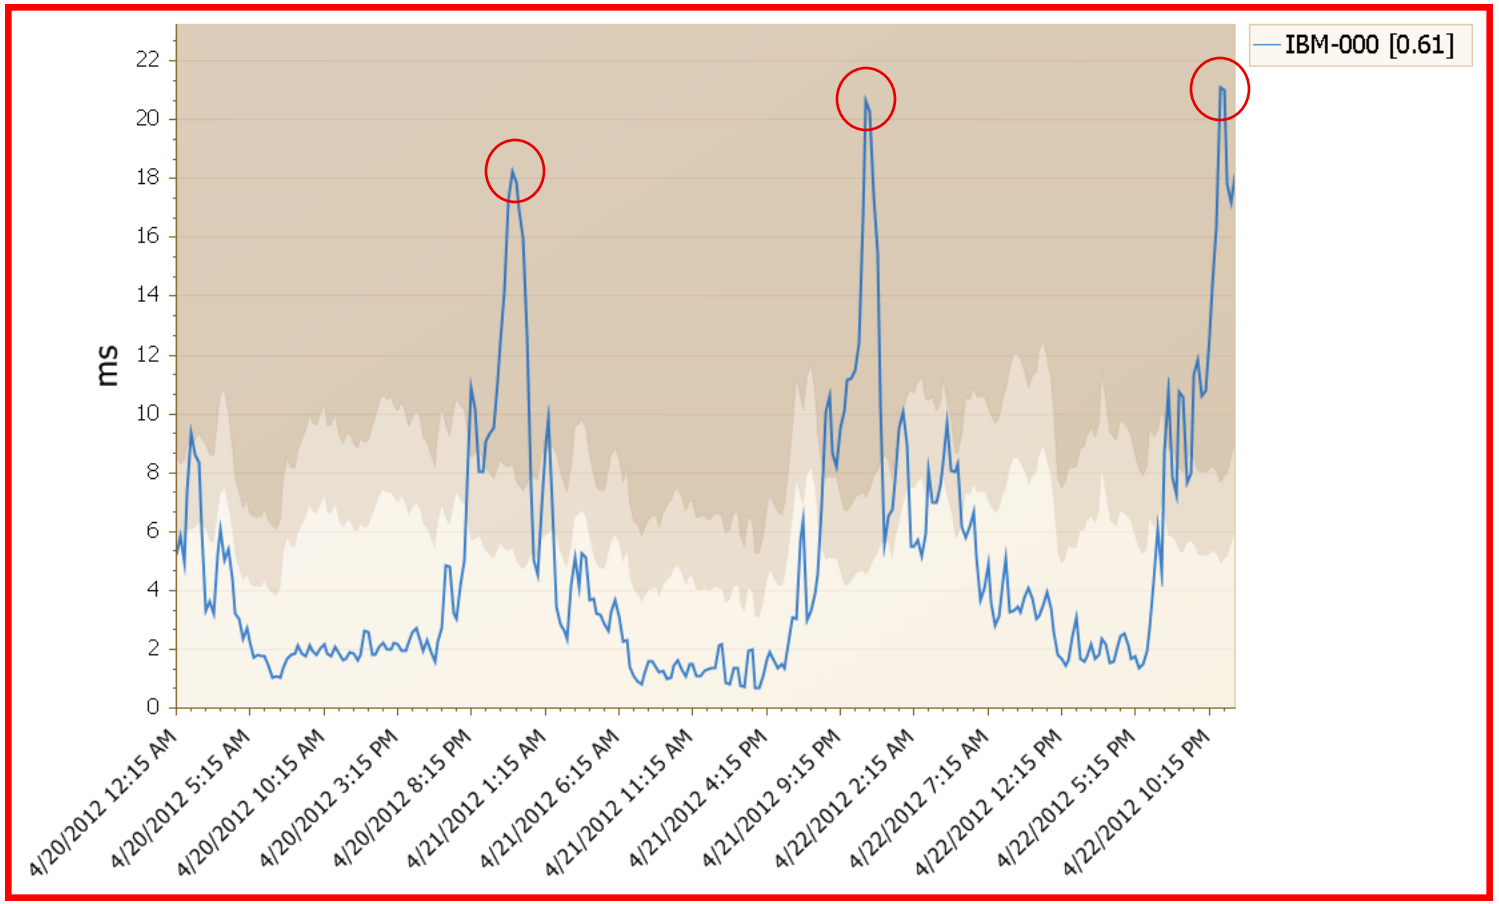

Response time peaks at 11:00 PM each day.

Fibre Front-End Read Response (ms) [rating: 1.62]

for all Ports by Serial

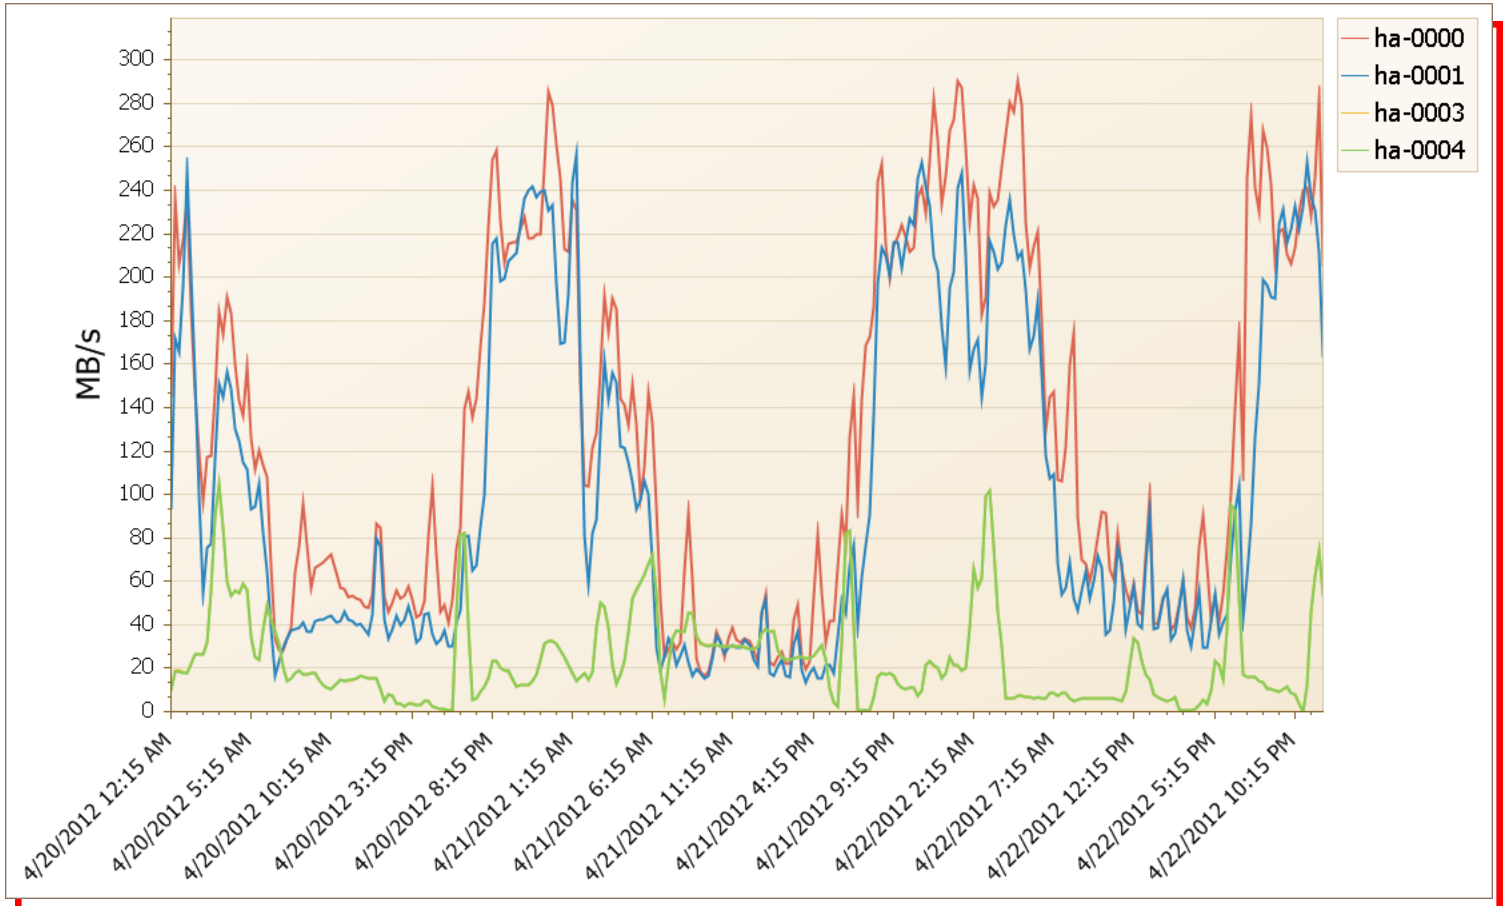

Host adapter throughput is imbalanced and peaks are approaching maximum capability for this type of adapter.

Drive Read Response (ms) [rating: 0.00]

For Serial 'IBM-000'

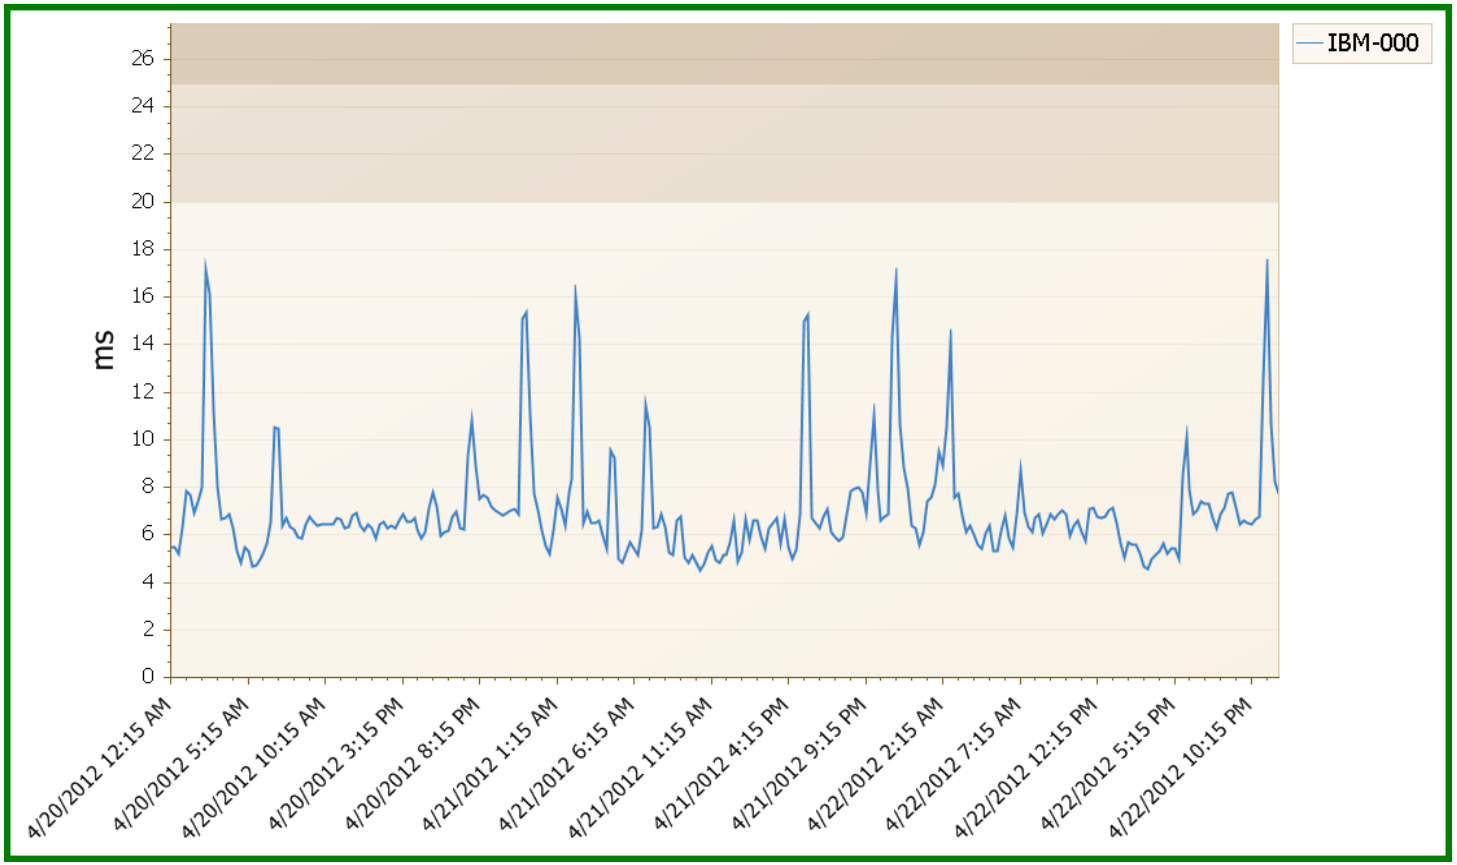

Average back-end response time has lots of peaks but not correlated with front-end response time. Overloaded drives not causing issue.

DS8000 Case Study Conclusion

Finding	Recommendations
HA0000 & HA0001 are saturated during peaks periods. Back-end drive write response times peak during same area but there are no FW bypasses so front-end write response time should not be affected by increases in back-end response time.	The solution is simply to redistribute some of the load from HA0000 & HA0001 to HA0003 & HA0004.

Questions?