VIRTUAL INSTRUMENTS

WHITEPAPER

Select Use Cases for VirtualWisdom Applied Analytics

April 7, 2017

CONTRIBUTORS

Andrew Sebastian

Roberta Sorbello

Mark Samuelson

Ricardo Negrete

Ryan Perkowski

Jim Bahn

Introduction

VirtualWisdom, our App-centric Infrastructure Performance Management (IPM) platform collects data from the instrumented datacenter environment and performs complex, multi-variate analysis, which we call Applied Analytics. With Applied Analytics, we've taken the expertise gained from working with hundreds of enterprise clients and built that wisdom right into the platform. The analysis that used to take experts hours or even days is now executed in minutes or seconds.

Applied Analytics focuses on the critical findings that exist beyond your raw data. The Applied Analytics module is delivered on a subscription basis, specific to the scope and scale of your unique VirtualWisdom implementation. This licensing model provides customers with our standard analytics including Balance Finder, Event Advisor, Trend Matcher, Queue Solver, VM Coordinator, VM Deployment Advisor, Seasonal Trend, and Workload Analysis. It also ensures that, as these analytics evolve, and new ones are added, our customers can immediately benefit. Applied Analytics deliver the authoritative insights and consistency you need to maintain agility and sustain your competitive edge.

In this paper, we're sharing several real-world examples of how some of our enterprise customers have used VirtualWisdom's Applied Analytics to prevent performance and availability problems, shorten the time to find problem causes, and quickly resolve trouble tickets.

#1: Application Slowdown and Impact of 4 Gbps infrastructure on servers and buffer credit starvation

This customer had an alarm set for Buffer to Buffer credit depletions. This alert was generating incident tickets to the NOC for resolution. They were seeing buffer credit depletions at multiple time during the day of over 40%, causing I/O response times to go above 30ms. The team ran Event Advisor on the storage port for the alert and produced the report below. This graph shows % of time spent at zero receive credits.

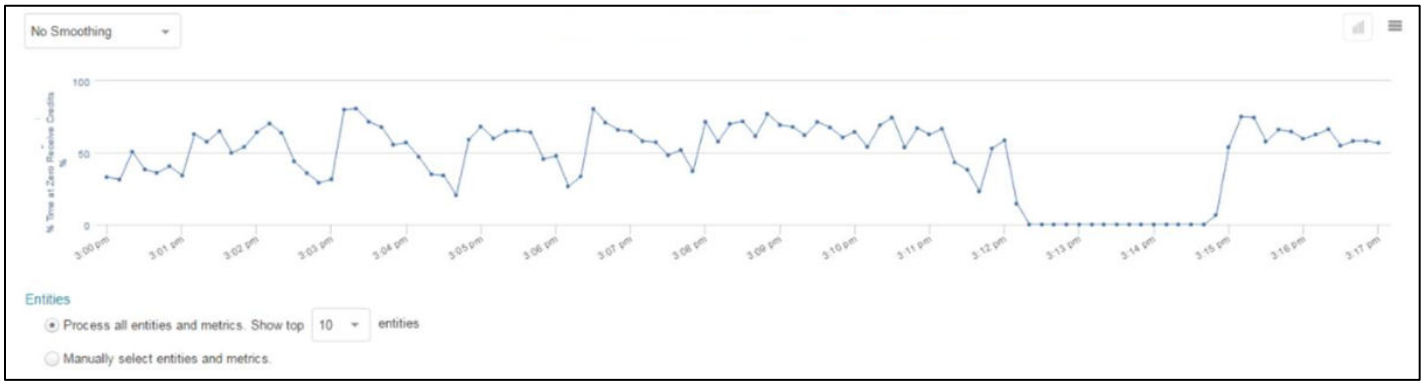

The Trend Matcher then identified the likely HBA ports (below) as the correlated causes, and zeroed in on one in particular.

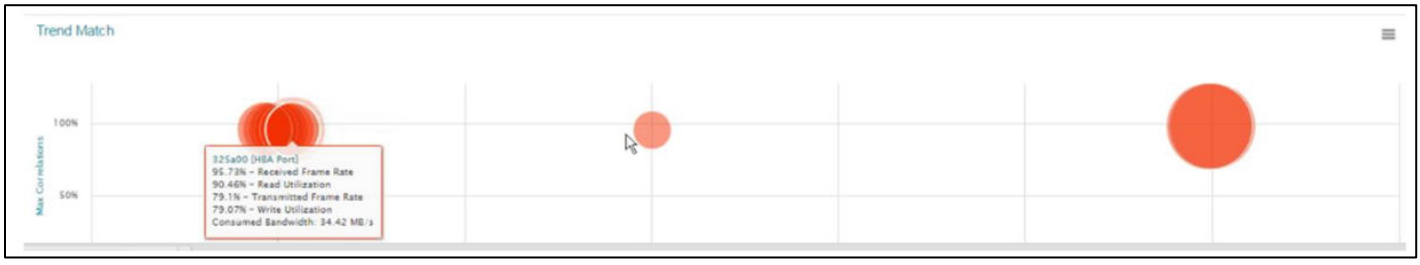

They then looked at the PID for that HBA port, and found that the uplink is a 4 Gbps HBA talking to an 8 Gbps storage array. *Topology View in later versions of VirtualWisdom would immediately identify the mismatch.* As I/O reads hit 350 MBps, the delays in the fabric increased significantly. The eventual solution was the elimination of the 4 GB infrastructure. A shorter resolution option was to reduce the queue depth on these blades to reduce the outstanding I/O. Total time elapsed to identify the problem was approximately 5 minutes.

#2: Database - Sporadic I/O Slowdowns During Production Hours

A user of a homegrown application was complaining about response time. The NOC had been working on problem isolation for at least 2 weeks without resolution and had engaged other members of the engineering team.

First, they reviewed alerts on each port, and noticed an error on one of the array ports. They ran Trend Matcher analytic and the heat map (below) found the SAP servers that were driving a lot of I/Os. It was noted that there was very high read I/O processing from an SAP application server. This caused a buffer credit starvation and impacted the app running on the other database. The solution was to move either the SAP database servers or the homegrown app database to different storage ports.

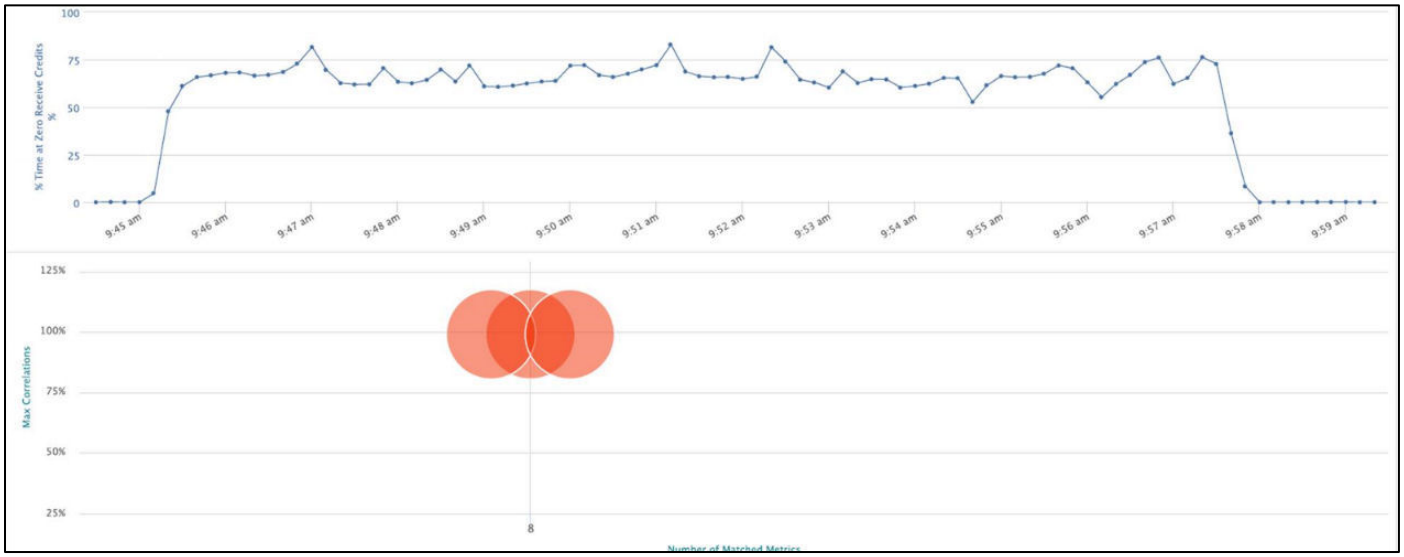

The report graph below confirms that the two SAP servers indeed had much higher read throughput on the HBA.

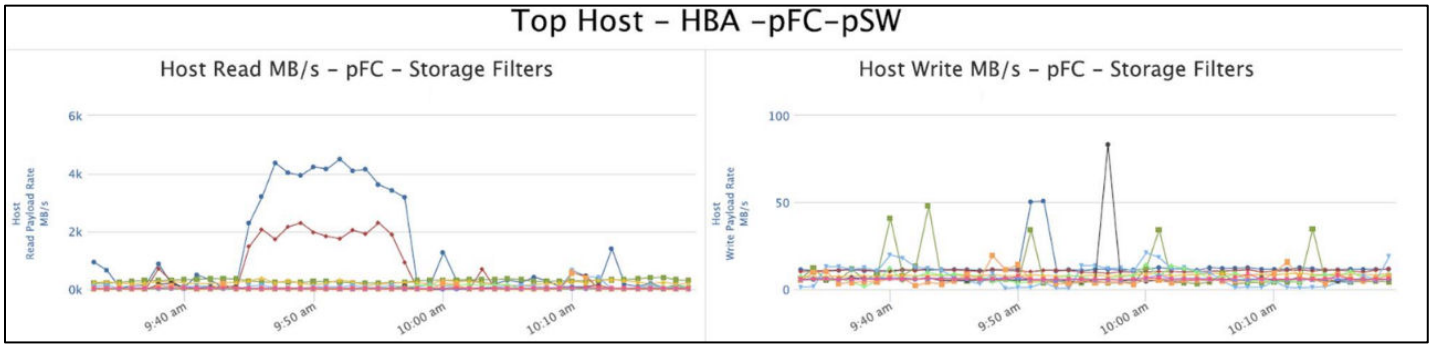

#3: Application Growth Filling up a Virtualized Array

In this example, the operations team from the NOC looked at some cases which had taken what appeared to be a long time to resolve, and they examined the use of VW Topology, Alerts, and Analytics to shorten MTTR. Specifically, there were some all-flash-arrays (AFA) which were not being monitored by the VW hardware probes, and were virtualized behind another vendor's storage controller. The problem was this … the customer was seeing a tremendous uptick in writes and it appeared that the AFA was in danger of running out of space if this continued. Before provisioning new storage, they wanted to discover who was writing all the data and whether it was a legitimate application/user or some runaway process. The first task was in identifying the host associated with the virtualized AFA. With virtualized arrays and internal pools, isolating the server driving this I/O was difficult.

The team ran the Event Advisor analytic, selected Utilization and Write Bandwidth and noted that Consumed Bandwidth exceeded 90% for a two-hour timeslot, then selected that period (highlight) for more analysis.

The above Event Advisor Analytic and all graphs in this paper are scrubbed of customer-identifiers to protect anonymity, so headings, labels, and legends may not contain all the data the customers saw.

Next, the team used the Trend Matcher analytic heat map to identify the possible "offending" host. It looked for all trends in the fabric that matched. In the screenshot below, with the bubble chart, you can see the high degree of correlation with events and the "write consumed bandwidth" metric.

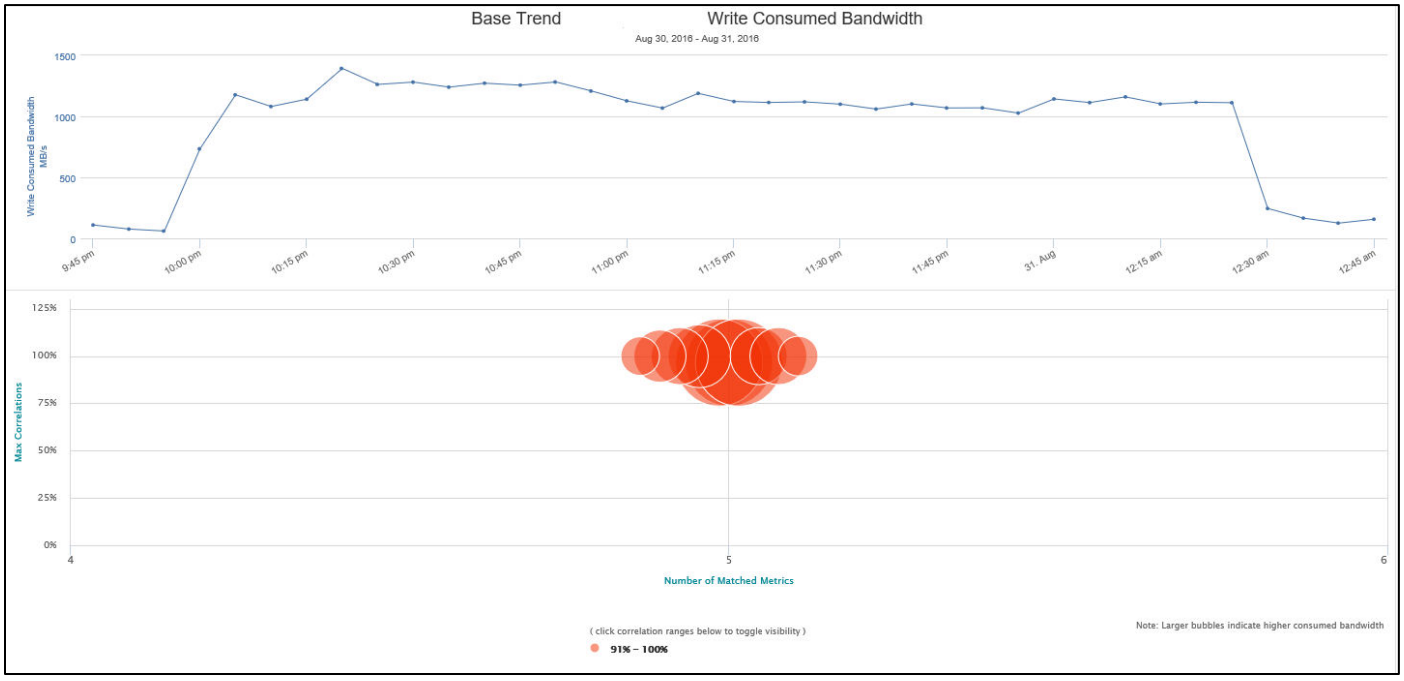

Following this, the Trend Matcher Analytic confirmed the identity of the host below, showing the 96% correlation in detail. Total time to run the analytics? Less than 4 minutes.

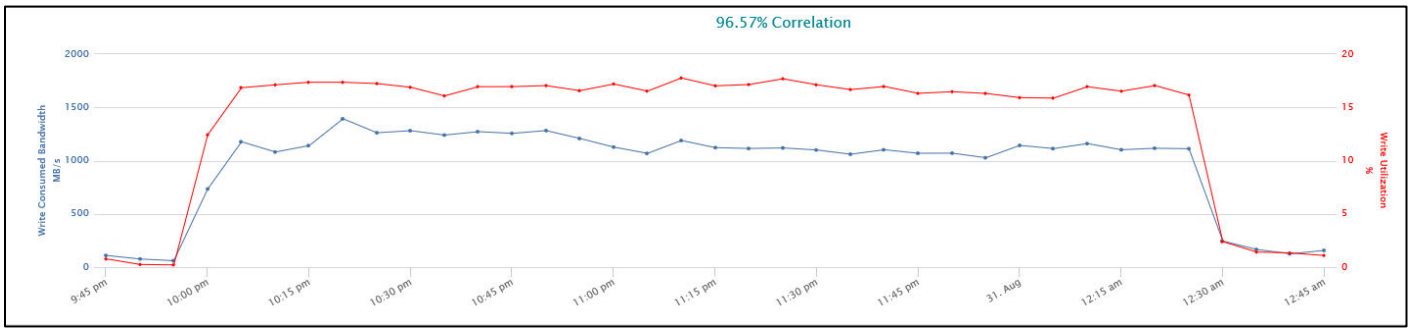

#4: Storage Port Overutilization Host Identification

This customer set VirtualWisdom to generate alerts for excessive port utilization when a port bandwidth is 90% utilized for at least 15 minutes. Often, monitoring tools may be used to identify which host is causing the high utilization by creating a host port performance report. Depending on how this report is formatted and the number of hosts, this approach can take a significant amount of time. With the Event Advisor analytic, it was simple to create a report for consumed read bandwidth (creation window below):

Using the Event Advisor report (*not shown*), the customer saw that the port had 17 events for a total of 260 minutes with the longest event being 46 minutes. By running Trend Matcher, and hovering over the largest spheres, the team quickly identified the problem host. The line graph below shows the correlation for just one of the events, but the entire report shows that this host was responsible for all 17 events. Time to identify this host was less than 3 minutes. Previously, the NOC was not been sure on how to identify and resolve these events and just closed the tickets.

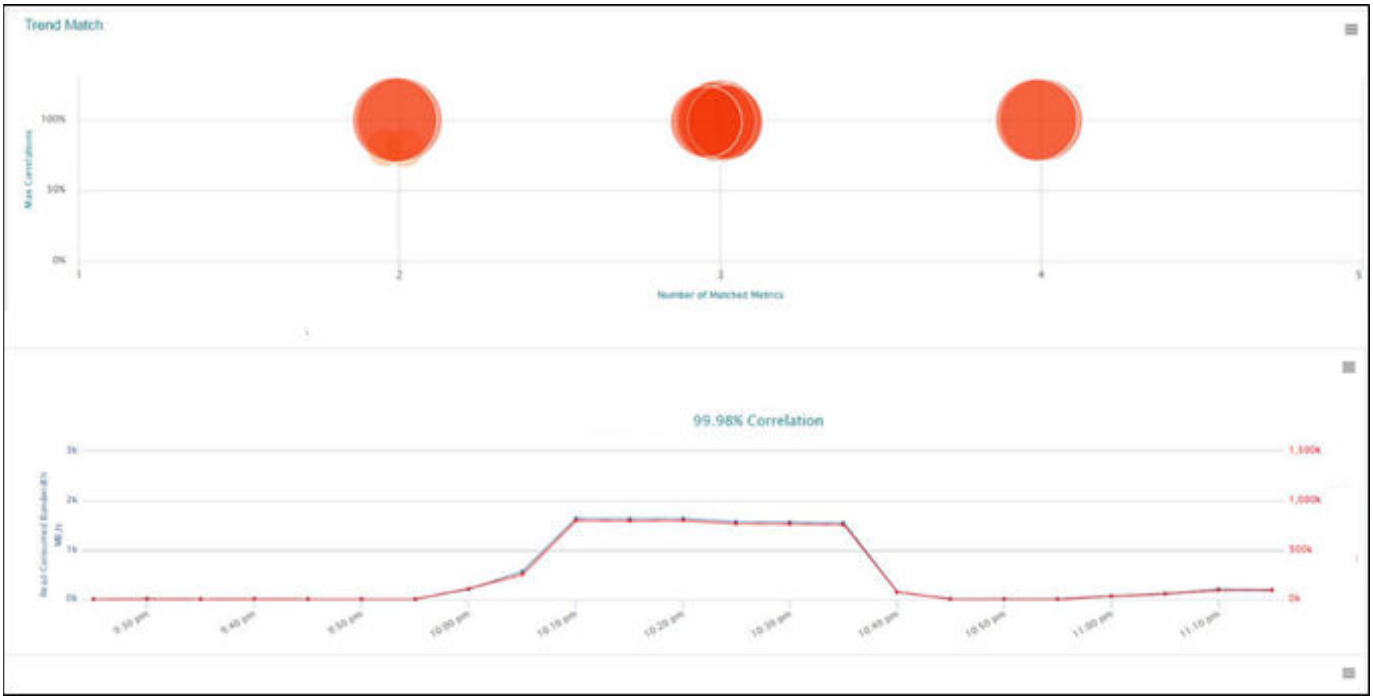

#5: VM CPU utilization

This customer used the VM Coordinator Analytic to avoid a potentially messy performance problem. This use case starts with a ServiceNow ticket. The IT team could see from the tasks that they had a CPU utilization issue, where utilization is over 90%.

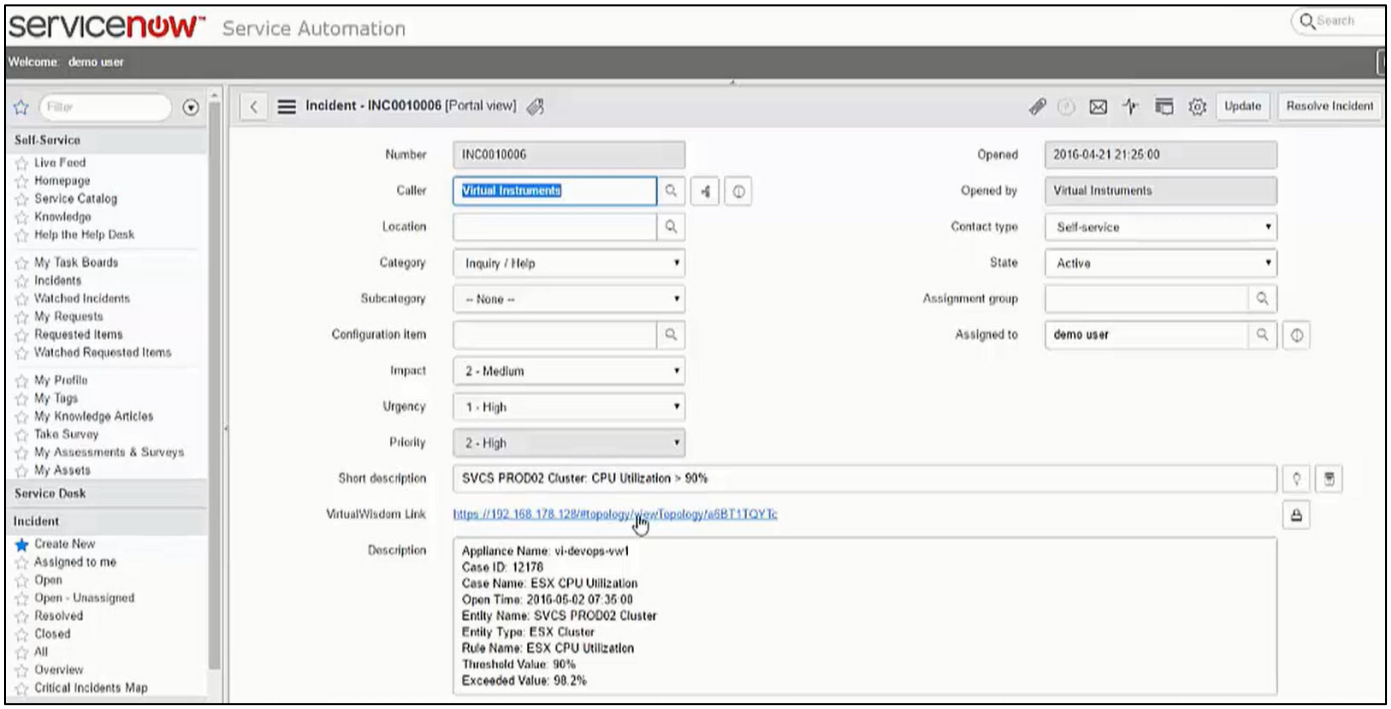

Clicking on the link, the customer viewed the topology *(not shown)* under VirtualWisdom, and found a CPU utilization alarm on a VM cluster. 3 VMs had % CPU Ready alarms and 1 had a utilization alarm. This told them that they had 1 VM that was very busy and 3 which were starving for CPU resources. VirtualWisdom told them to use the Event Advisor analytic to investigate, and the VM Coordinator analytic to remediate, below.

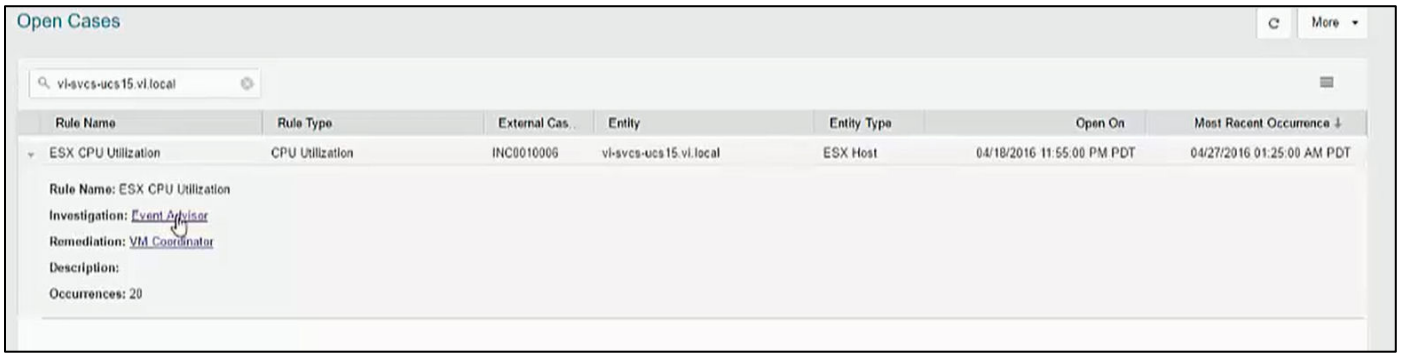

The team changed the timeframe to go back a few days because the alarm starting firing then. That's one of the great things about VirtualWisdom, you can go back at any point in time and replay what happened. They ran the Event Advisor and clicked on the top host, and could see from the screen below that something was running every night on this system.

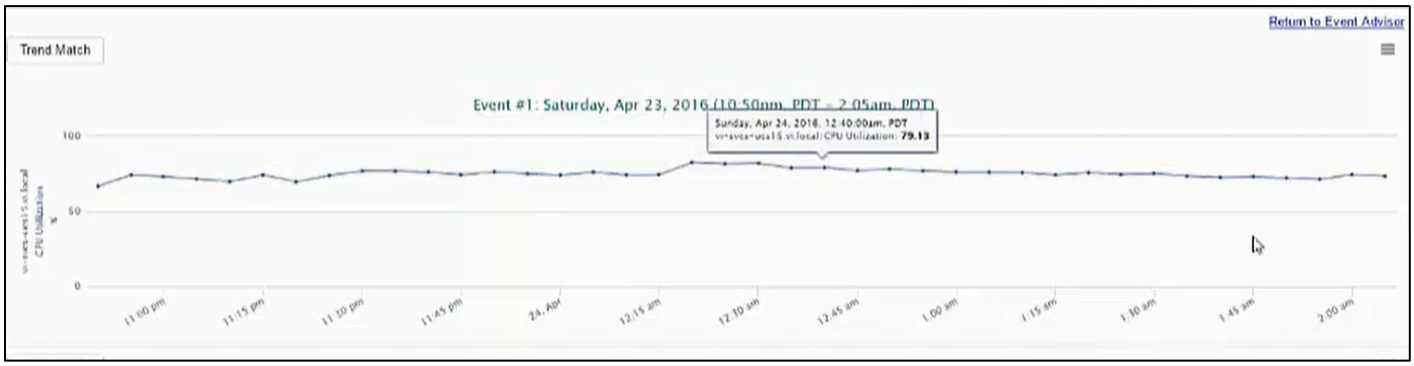

Next, they invoked the Trend Matcher analytic to find all the entities that matched this workload. It gave them a bubble plot showing the correlations, below. There was a strong correlation with the VM they saw on the utilization alarm, so they investigated further.

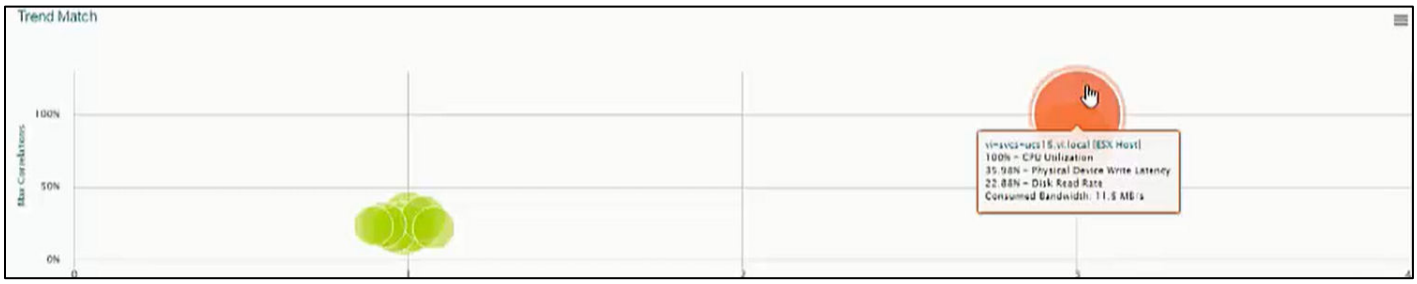

Below, they could see that when the server got busy, the VM got busy at the same time. There was an 81% correlation, which was close. So, they knew the VM that was causing the server problem.

Next, they ran the remediation link in the VM Coordinator analytic. It looked at all the VMs across all the hosts and balanced the VMs, comparing the Before and After configurations. VM Coordinator made several suggestions. In this screen below, you can see that when the recommendations were employed, VirtualWisdom predicted that CPU utilization would go from a range of 18 – 78% to a range of 26 – 47%. So, using the VM Coordinator analytic, headroom was significantly improved, and potential performance problems were avoided.

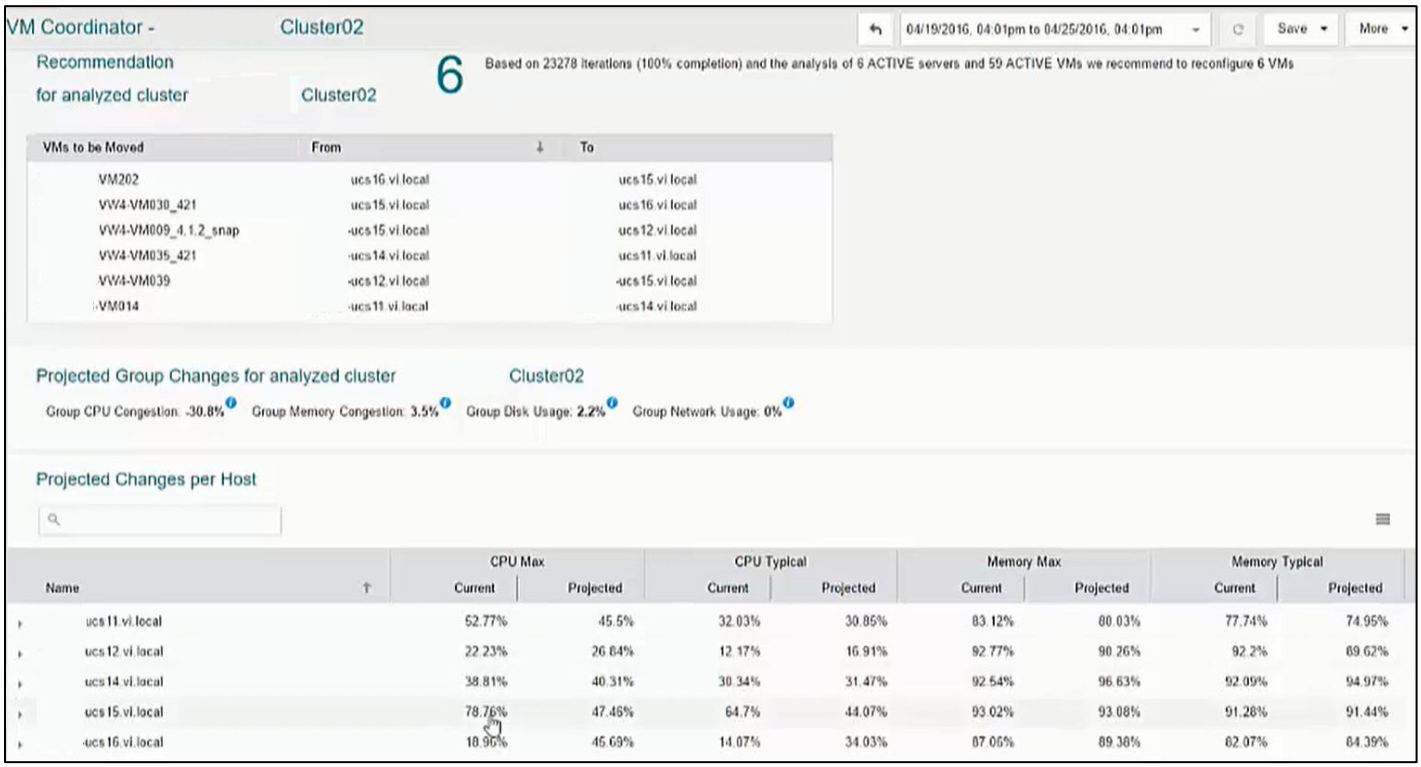

#6: HBA utilization trends driving ISL utilization

This use case begins with an alert for an Inter-Switch Link (ISL) which had exceeded a defined threshold. The engineer logged into the VI-Utilization-SAN report and confirmed that the alert was accurate. But what was driving the utilization? He loaded the ISL which generated the alert into Trend Matcher and generated this bubble chart and line graph, below.

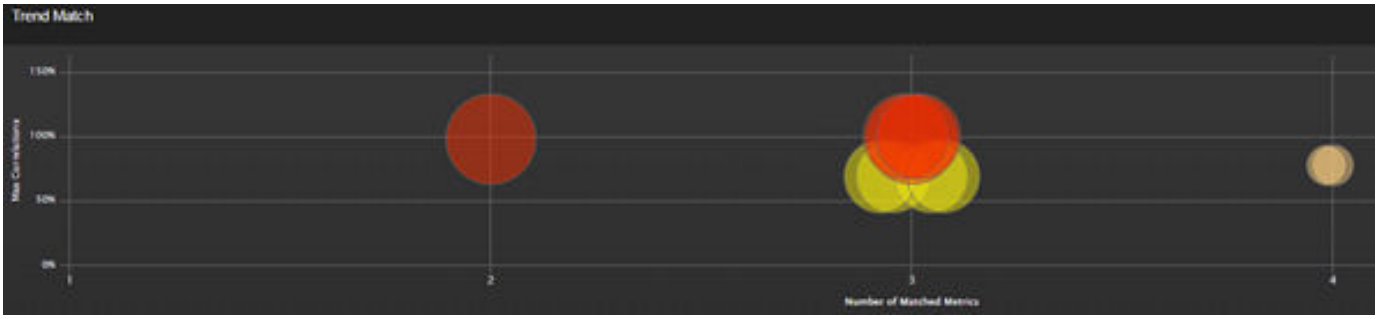

When the HBA was identified, the engineer confirmed that it traversed the ISL in question, via the topology report below. Host is on the left in blue; storage on the right in gold. *Reminder: identifying strings have been scrubbed*.

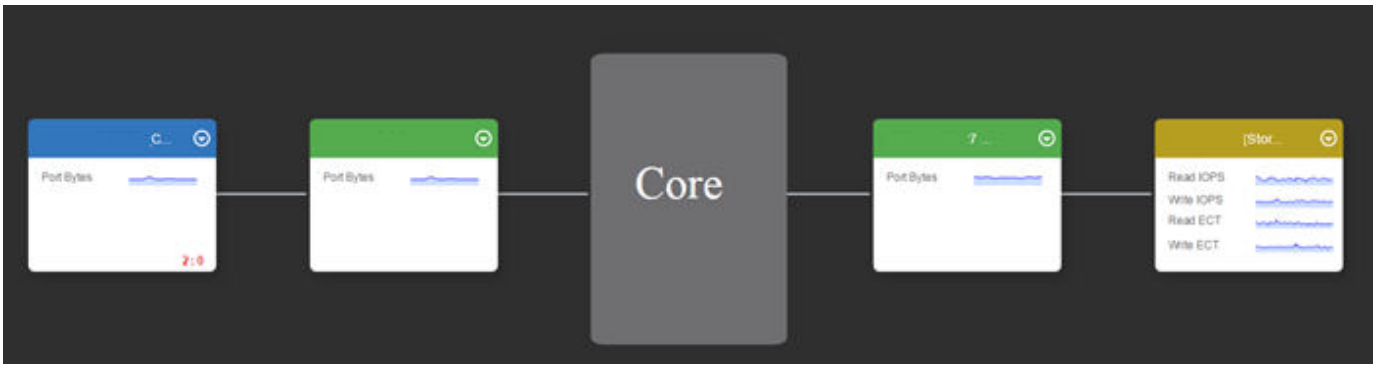

The above Topology report shows the host-to-switch-to-ISL-to-switch-to-storage relationship. Total time to find the source of the overloaded ISL was less than 3 minutes. Steps could then be taken to balance the HBA traffic over the SAN.

#7: Hosts causing buffer-to-buffer starvation on ISLs

An alert on an ISL was investigated by running the Event Advisor analytic below, which showed the percent of time at zero buffer credits (credit starvation) on the B channel.

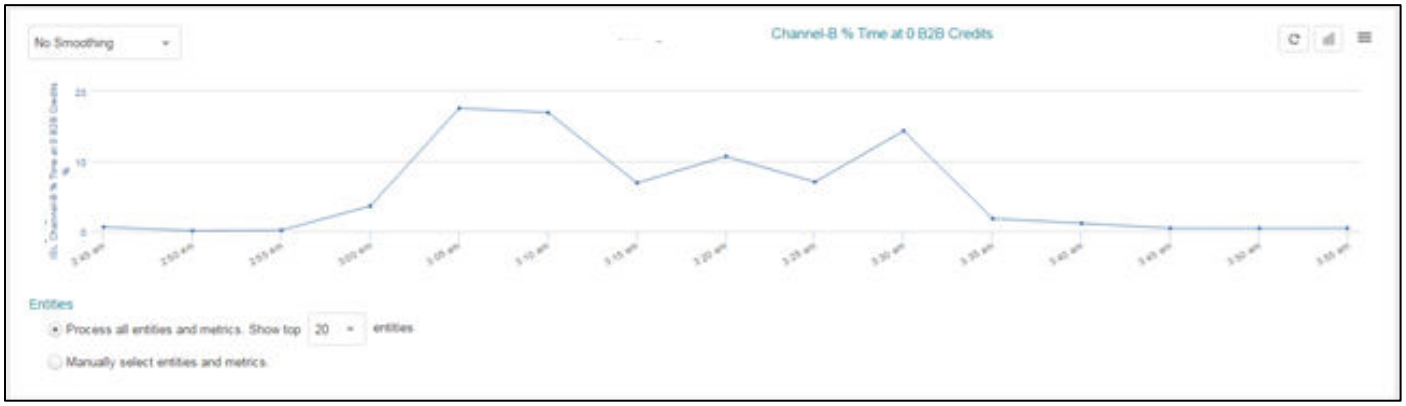

Executing the Trend Matcher analytic yielded the charts below, showing a 99% correlation and matches for host and storage ports.

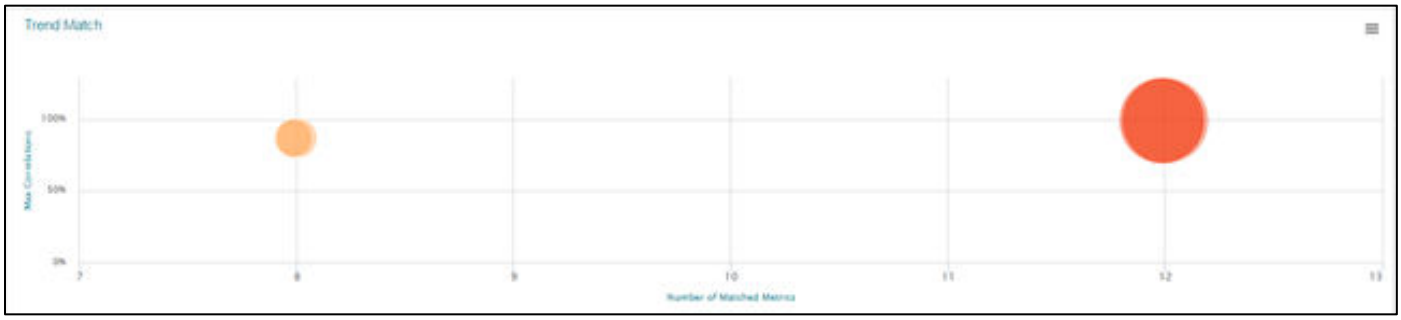

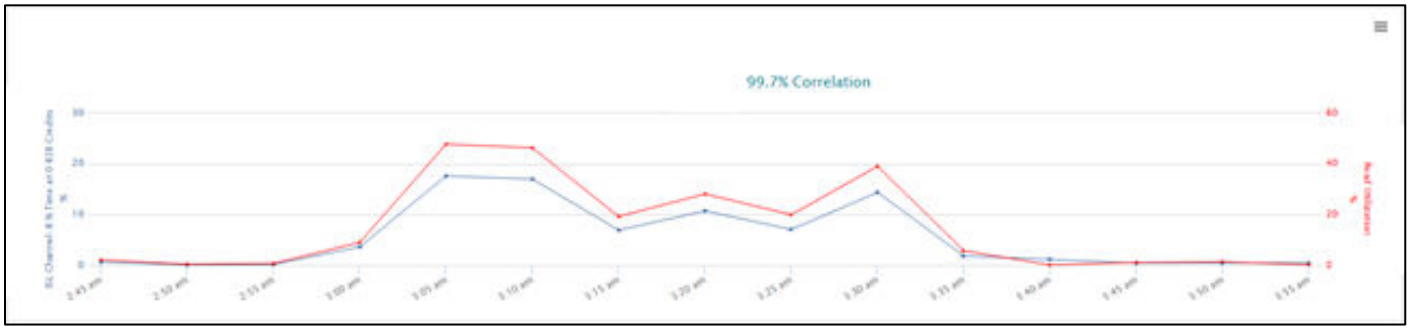

The team then used the Topology view (below) to confirm that the ISL in question was servicing the host and storage port. *Reminder: identifying strings have been scrubbed*

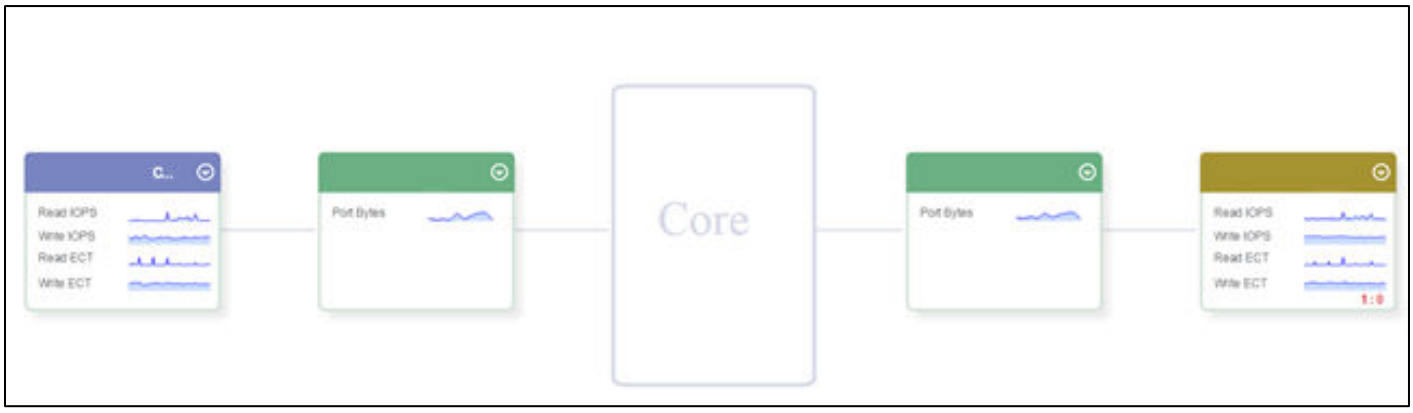

Total time to identify the host was less than 5 minutes.

#8: Heavy port utilization causing latency for systems on other ports

The IT team received a VirtualWisdom severity 1 latency alert on a storage array. They logged into the VirtualWisdom Portal and validated the alert by navigating to Alarms -> Open Cases. The alert offered the 'Event Advisor' button, which they opened to prepopulate the necessary entity information. They defined the appropriate timeframe shown in the green line below, and ran it.

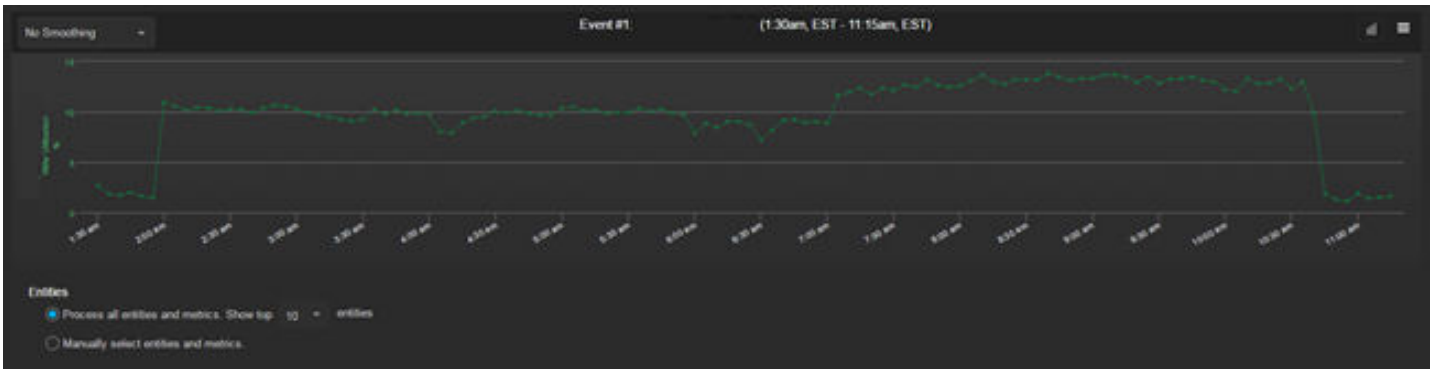

They then invoked the Trend Matcher, and the bubble chart below showed the highest correlations grouped in the upper right.

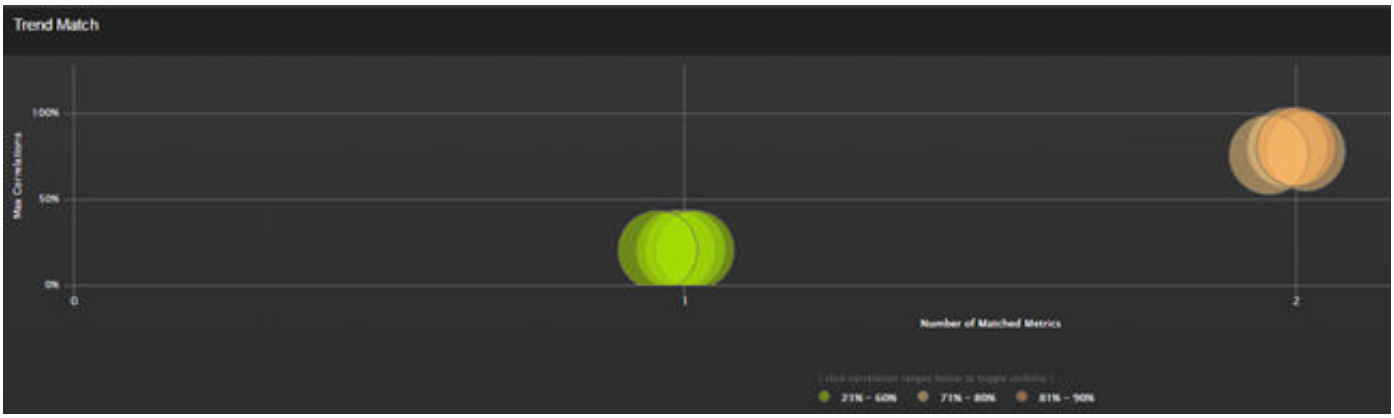

And with an 80% correlation, they identified offending host system.

Through the topology map *(not shown)* they verified the system accessing storage from the storage array (which generated the alert). Once validated, they determined the other host systems accessing storage from the alerting storage array as they were the most likely to experience future performance impact. Armed with this data, they could rebalance workloads before application users were severely impacted. Total analysis time was less than 10 minutes.

Conclusion

With Applied Analytics, we've taken the expertise gained from working with hundreds of enterprise customers and built that wisdom right into the platform. The analysis that used to take experts hours or even days is now executed in minutes – turning data into actionable answers.

In this paper, we demonstrated customer's use of Trend Matcher, Event Finder, Topology View, and VM Coordinator. In future versions of this document, we'll show how customers used Balance Finder to load balance across hosts, and optimized system-wide performance and reduced wait times with Queue Solver.

VIRTUAL INSTRUMENTS

Sales
Sales@virtualinstruments.com
1.888-522.2557

Website
virtualinstruments.com