

HPE InfoSight: Cross-stack Analytics Cut Through Infrastructure Silos

JUNE 2018



Accurate and action-oriented predictive analytics have long been the Holy Grail of IT management. Predictive analytics that bring together large amounts of real-time data with powerful analytical capabilities have the potential to provide IT managers with real-time, data-driven insights into the health and performance of their overall environment, enabling them to anticipate and remediate looming issues and optimize resource utilization. While these potential benefits have long been understood, it has only been recently that major innovations in cloud, Internet of Things (IoT), data science, and AI/machine learning have paved the way for predictive analytics to become a reality in the data center.

The IoT now enables companies to collect and monitor real-time sensory or operational data at the edge—whether in online financial systems, retail locations, or on the factory floor. This raw data is typically streamed to the cloud, where it can be tabulated and analyzed. Powerful advances in edge-to-cloud networks and global learning capabilities make the cloud an optimal location for the analytics to take place. Informed by data science and increasingly driven by AI and machine learning technologies, these analytics can help IT managers to monitor key system metrics and understand how well specific infrastructure elements—such as servers or storage—are performing.

But analytics that are focused on a single infrastructure element at a time can only go so far. Sure, it is helpful to monitor the health and performance of specific IT resources, such as CPU heartbeat or storage latency, but infrastructure resources do not operate independently or in isolation. Analytics must go beyond one dimension, and take into account how resources such as servers and storage interact with and depend on one another. This is especially critical in virtualized infrastructures, in which the interaction of virtual machines with hosts, networks and storage makes IT management even more challenging. Ideally, using the power of AI, analytics can cross these various layers of the IT stack to reveal the impact of resource interactions and interdependencies among all the layers. This would take analytics to a whole new level, transcending the limits of human intelligence to enable dynamic, multi-dimensional analysis of complex, virtualized IT environments.

Think about the implications of AI-driven, cross-stack analytics for IT management. For example, such a capability has the potential to transform technical support from a reactive, always-playing-catch-up function to a proactive and forward-looking capability. In this scenario, built-in analytics are capable of connecting the dots between infrastructure layers to automatically anticipate, diagnose, and fix technical issues before they become major problems. Cross-layer analytics might also help to improve system performance by predicting looming configuration issues and recommending changes to address them.

One product—HPE InfoSight—is already embracing these possibilities, fast-forwarding to bring AI-driven, cross-layer analytics to virtualized environments today. HPE InfoSight has proven its value in delivering predictive storage analytics to customers for many years, while extending its capabilities across the infrastructure stack. In this piece we'll explore the key characteristics that customers should look for in an analytics solution for virtual infrastructure and then look at the HPE InfoSight architecture and its capabilities, and how they are helping customers transform IT management in virtualized environments today. Specifically, we will demonstrate how one customer uses cross-stack

analytics delivered by HPE InfoSight to save tremendous time and money in their HPE 3PAR Storage environment.

AI-DRIVEN, CROSS-STACK ANALYTICS INCREASE IT EFFICIENCY AND LEAD TO BETTER BUSINESS OUTCOMES

Virtualization can bring many benefits to enterprise server infrastructure, including improved utilization, increased IT efficiency and agility, and reduced operating costs. But the technology can also introduce new management challenges. As many IT admins have learned firsthand, server virtualization can make infrastructure more complex and less predictable, particularly in the interactions with networking and storage. For example, companies may experience intermittent performance or availability issues arising from IO blending or “noisy neighbors,” which can be difficult to diagnose even in environments running a modest number of VMs.

To cut through this complexity, enterprises need predictive analytics that can span the entire infrastructure stack, to identify interdependencies and provide cross-layer visibility and insights. This capability allows IT managers to understand interactions between VMs, networking, and storage, and help pinpoint the root cause of any issues. For instance, when an application suffers from periodic bouts of high latency in a cluster running 100s of virtual servers, the issue might be caused by a few resource-hungry or misbehaving VMs and manifest itself only intermittently in various layers of the stack. Attempting to localize and characterize this type of problem with traditional operations management tools can be like finding a proverbial needle in a haystack, frustrating even the most dedicated IT manager. Some newly emerging analytics solutions are designed to help IT overcome these challenges.

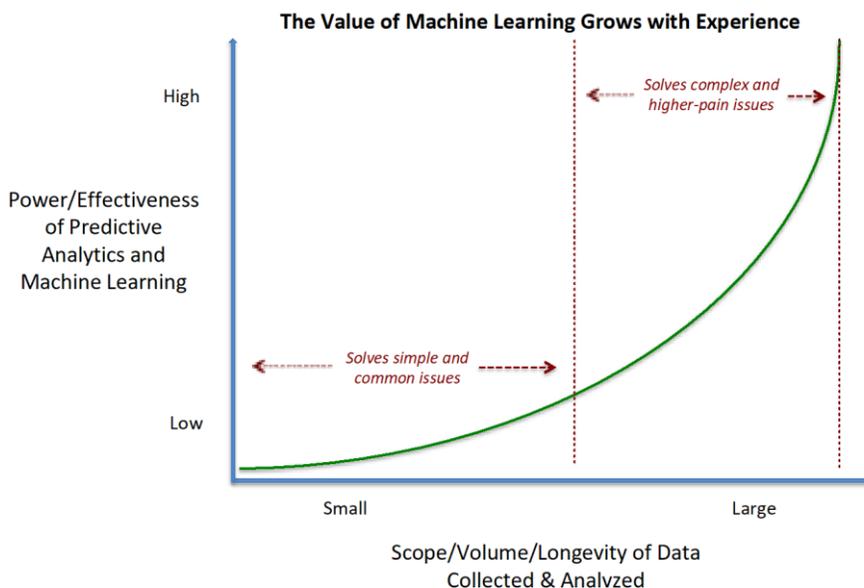
But as we’ll see, not just any analytics tool will do. To gain the highest-value management visibility and intelligent insights in a virtualized environment, customers should focus on solutions with the following characteristics:

- **Cloud-based.** This allows the analytics solution to learn from infrastructure telemetry and application behavior across a global installed base.
- **Cross-stack capable.** To key in on the true cause of complex performance and availability issues in a virtualized environment, analytics must be able to monitor interactions and identify dependencies across all infrastructure layers simultaneously, as opposed to focusing ad hoc on just a single layer or resource. Focus on solutions whose scope of data collection extends wide and deep, covering a broad range of installations and traversing every layer of physical and virtual infrastructure.
- **Embedded within infrastructure.** To enable comprehensive and continuous visibility, collection of telemetry data should be instrumented into the infrastructure itself rather than handled separately through independent sensors or tools.
- **Benefiting from a long history of data collection.** History and experience is really the difference when it comes to one analytical solution vs. another: the greater the history/experience, the smarter the solution will be, and the better and more robust its predictions. Look for solutions with data sets that have been accumulated over an extended period of time.
- **Powered by AI.** Advanced machine learning drives the ability to develop “cause-and-effect” insights and more accurately predict future behavior and potential issues. In addition, as analytics become more AI-driven, they enable greater automation of administrative and support tasks by progressively reducing the need for manual human intervention.

As shown in Figure 1, the power and benefit of machine learning increases with the size, scope and longevity of collected data. For this reason, we recommend that customers focus on solutions that

have been proven over many years across hundreds of deployments, to take advantage of the accumulated learning from analyzing massive historical data sets.

Figure 1: Value of Machine Learning Based on Size and Scope of Collected Data



SOURCE: TANEJA GROUP

Cross-infrastructure analytics solutions should also be capable of addressing several key use cases in virtualized environments. One of the most powerful applications is in technical support, where such solutions help improve the customer experience by anticipating and remediating problems before they turn into support incidents. This, in turn, improves virtual infrastructure reliability while reducing support costs. A second use case is performance optimization, where the analytics solution spots current or developing configuration problems in one or more infrastructure layers and recommends or makes changes to improve performance. To work most effectively in each of these use case examples, an analytics solution should provide single-pane resource visibility and root cause diagnostics along with the ability to help customers optimize resource spend.

Most analytics tools on the market today are relatively limited, delivering on no more than a couple of the key characteristics outlined above. In particular, while the best IT Operations Management (ITOM) solutions are now cloud-accessible and embedded in infrastructure, they are not cross-stack capable or AI-driven, and so are unable to make data-driven predictions and recommendations for action. In fact, only HPE InfoSight delivers on all of these requirements.

Let's now take a closer look at HPE InfoSight and examine in detail what sets it apart from other cross-stack analytics solutions.

HPE INFOSIGHT: AI FOR THE DATA CENTER

Initiated in 2010, HPE InfoSight was created by Nimble Storage with the goal of transforming the support experience. They believed that multi-dimensional machine-learning techniques informed by globally collected data could be applied to managing and analyzing the complex web of data center infrastructure components. HPE InfoSight is a cloud-based platform, and is initially focused on solving infrastructure issues that impact storage performance and availability. Interestingly, because it correlates all layers across the infrastructure stack, 54% of the support issues detected by HPE InfoSight have nothing to do with the storage device itself. This fact demonstrates that issues that may

at first seem to be storage performance or availability related are becoming too complex for most humans to fix without some level of advanced AI.

HPE InfoSight enables a truly transformative support experience, in which 86% of would-be support issues in Nimble Storage environments are automatically predicted and resolved without any action on the customer’s part. The remaining 14% of issues that result in support incidents are addressed directly to a Level 3 support engineer, eliminating the need for iterative cycles of problem solving and escalation. This approach has resulted in highly satisfied customers, earning HPE Nimble Storage one of the highest Net Promoter Scores (NPS) in the industry.

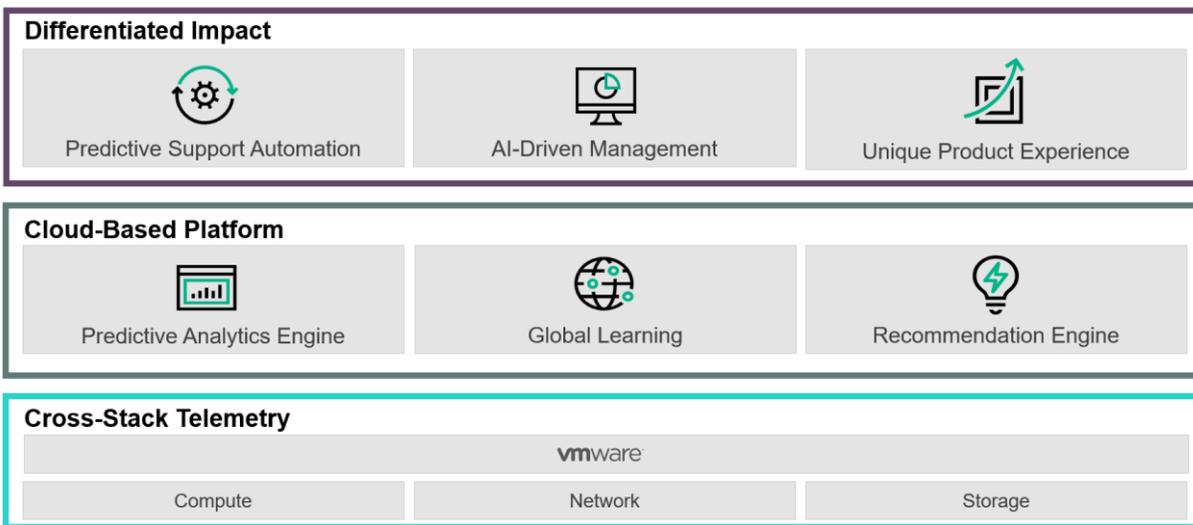
When HPE acquired Nimble Storage in early 2017, they seized the opportunity to increase the intelligence of HPE InfoSight by integrating with over a decade of existing rich telemetry collected from the 3PAR installed base. This formed the industry’s most experienced AI and enabled HPE to rapidly extend its predictive analytics and cross-stack algorithms to 3PAR.

HPE InfoSight now supports HPE 3PAR Storage and the customer benefits seen thus far are significant. They have demonstrated 85% prediction and auto-resolution for more than 1000 complex cases, driving increased availability across its installed base. The following section outlines the key elements of the InfoSight architecture that enable a transformative product experience across both HPE Nimble Storage and 3PAR Storage.

The HPE InfoSight Design Approach

The key design characteristics of HPE InfoSight are shown in Figure 2 below, followed by a brief explanation of how the various layers work together to drive better business outcomes for customers.

Figure 2: HPE InfoSight: AI for the Data Center



SOURCE: HPE

BASE LAYER: CROSS-STACK TELEMETRY AND INSTRUMENTATION IS FOUNDATIONAL

The goal of cross-stack telemetry is to have the right instrumentation at the appropriate level of detail customized for the device under management. This means that, to solve more complex issues, the specific device under management, such as a storage system, must have deep instrumentation both to automate the detection of mundane failures (e.g. simple hardware faults) and to monitor and measure critical SLA attributes (such as IOPS, cache efficiency, storage memory, and CPU utilization). The instrumentation list is long for HPE 3PAR and Nimble storage, in which thousands unique hardware sensors and system attributes are continuously monitored and the resulting telemetry data sent home to the cloud for deep analysis. HPE InfoSight also uniquely adds the appropriate set of instrumentation to enable data from non-storage (e.g. server and networking) infrastructure to be

captured, correlated and sent home for analysis. Examples of key cross-stack instrumentation related to storage include virtualization performance metrics, network and switch metrics, IO drivers, etc. Specifically, for a VMware virtualized environment, InfoSight monitors, collects, and sends home hundreds of key data elements across all the virtualized layers that may impact what may be perceived as storage-related issues. As HPE expands InfoSight to other hardware elements, the solution will provide an increasingly more comprehensive view of data center infrastructure availability and performance, thereby enabling an even better support experience for customers.

AI AND DEEP-LEARNING LAYER: A CLOUD-BASED PLATFORM ENABLES CONTINUOUS GLOBAL LEARNING

Initially, storage device administration was handled completely as an on-premises function. Now, most devices phone home, where many typical hardware failures and configuration issues can be remotely taken care of. However, to truly transform the support and management experience, one must collect data globally and analyze it based on a mix of machine learning and expert human input. HPE implements this by combining a team of data scientists with support engineering to build algorithms and signatures that benefit the customer experience.

To automate and enhance the customer's support and management experience, HPE has developed a **predictive analytics engine** to identify potential issues before they arise, as well as a **recommendation engine** that give customers clear-eyed, automated guidance on a variety of topics that improve business outcomes. HPE InfoSight constantly analyzes data and continues to build predictive analytics and recommendations that provide counsel based on InfoSight analysis of these data streams, such as performance resolution options that may or may not involve storage itself.

DIFFERENTIATED IMPACT LAYER: ENABLING BETTER BUSINESS OUTCOMES

Great technology is for naught if it does not deliver results that customers desire and appreciate. HPE delivers customer business value in the following key areas:

Predictive Support Automation: HPE InfoSight goes beyond mundane fault detection and resolution to automatically predict and remediate complex issues before they impact the customer. For example, the solution might determine that a particular combination of firmware and hypervisor versions has the potential to degrade or even crash a customer's system. Once such an issue has been identified, InfoSight draws on its machine learning and intelligent insights to determine the subset of HPE customers that are running this environment and then automatically triggers notifications and remedies to be dispatched to susceptible customers, preventing the issue from occurring.

AI-driven Management: The rich set of cross-stack telemetry data along with cloud-based machine-learning allows HPE to deliver unique AI-driven features for day-to-day management of infrastructure. These AI-driven features take provisioning and monitoring to levels not seen before. Cross-stack analytics for virtualized environments is one such major capability highlighted in the rest of this paper.

Unique Product Experience: With these capabilities, HPE InfoSight delivers customers a unique product experience that they cannot find with other infrastructure providers. We describe below our in-depth interview with one HPE InfoSight customer who validated the unique experience that cross-stack analytics has brought to the HPE 3PAR Storage environment.

HPE InfoSight Cross-stack Analytics for VMware

Virtualized environments, while extremely powerful, flexible, and efficient, can also be quite complex, as the physical and virtual elements of the solution are layered on top of each other. Issues that seemingly relate to one element of the solution may, in fact, be the product of another unrelated element. It is through this lens that the cross-stack analytics capability of InfoSight was developed. While cross-stack analytics is just one capability of HPE InfoSight, it is one of the most valued by customers and demonstrates well the power of the HPE InfoSight platform. Normally in a paper like this, Taneja Group would evaluate and demonstrate the power of cross-stack analytics by

systematically describing a hypothetical or lab-generated situation where the power of cross-stack analytics solves a particular problem. For this profile, we had the added benefit of gaining access to an HPE 3PAR customer using HPE InfoSight and cross-stack analytics to solve a complex performance issue. The following section provides an in-depth view of how cross-stack analytics works through the eyes of a real customer solving a real problem.

CROSS-STACK ANALYTICS THROUGH THE EYES OF A CUSTOMER

To analyze the impacts of HPE InfoSight in the real world, we interviewed HPE 3PAR customer Gary Stewart, a Storage Architect for a large chemical company. Gary granted us a real-time insight into a very complex virtualization performance issue he was resolving, one with the potential to save his company the purchase of a new storage array. We found that HPE InfoSight was instrumental in helping Gary understand the cross-stack interactions of his environment. As Gary so enthusiastically stated, “This is a great product, I want to get the story out.” What follows is a detailed report on how Gary is using HPE InfoSight cross-stack analytics and the impact it is having on his business.

The HPE 3PAR Environment

Gary works for a global chemical company with ~\$8B in revenue and over 7,000 employees worldwide. Gary oversees the storage environment in which there are over 3PB of raw capacity across 35 HPE 3PAR Storage systems. He uses the 3PAR management console (called SSMC) that can aggregate many systems together. Gary chose HPE 3PAR since it allowed him to standardize on one architecture that could support a wide variety of workloads, including servicing the most mission-critical SAP environment, while also extending the architecture to remote offices and factories using cost-effective, smaller configurations. When he cited the most important reasons for choosing 3PAR, the following attributes were at the top of his list:

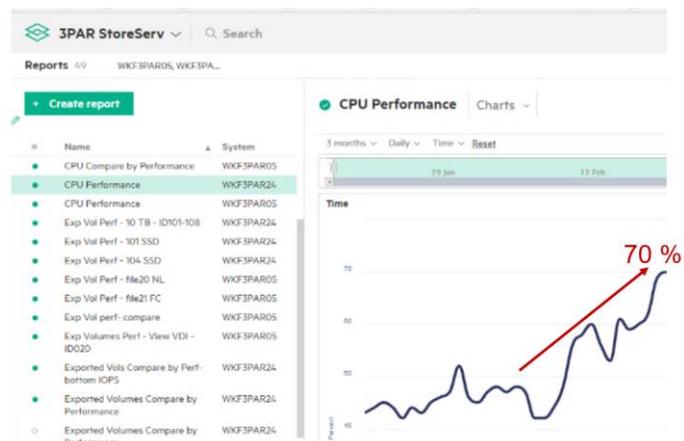
- 1) **Rock solid reliability.** Gary has not had a storage outage in 2 years.
- 2) **HPE support engineering** with remote access so he can take care of keeping the environment up-to-date and focus on strategy and planning.
- 3) **Global data mining of 3PAR systems worldwide.** Gary is very appreciative that HPE is monitoring many thousands of 3PAR devices worldwide and can proactively come back to him with known issues and preventive measures, even in cases where he has not yet seen a problem. HPE demonstrated this when they contacted Gary to replace a set of problematic hard drives that might have eventually impacted a subset of his products.

When Gary found out in late 2017 that HPE InfoSight was going to support HPE 3PAR, he jumped at the chance to be a beta user and is now using InfoSight in full production. When he first adopted a beta version of InfoSight, Gary was in the midst of resolving a 3PAR system issue where CPU utilization had spiked to over 70% (see figure 3). He had a choice to either do the typical thing and throw more hardware at the problem or try to understand what drove the spike in load. What follows is a detailed account of Gary’s use of HPE InfoSight to solve this problem.

The IO Contention Problem

Gary’s utilization problem involved a 1600-seat VMware Horizon View VDI environment. The HPE 3PAR 8200 that was servicing this workload

Figure 3: HPE 3PAR SSMC CPU Performance Report

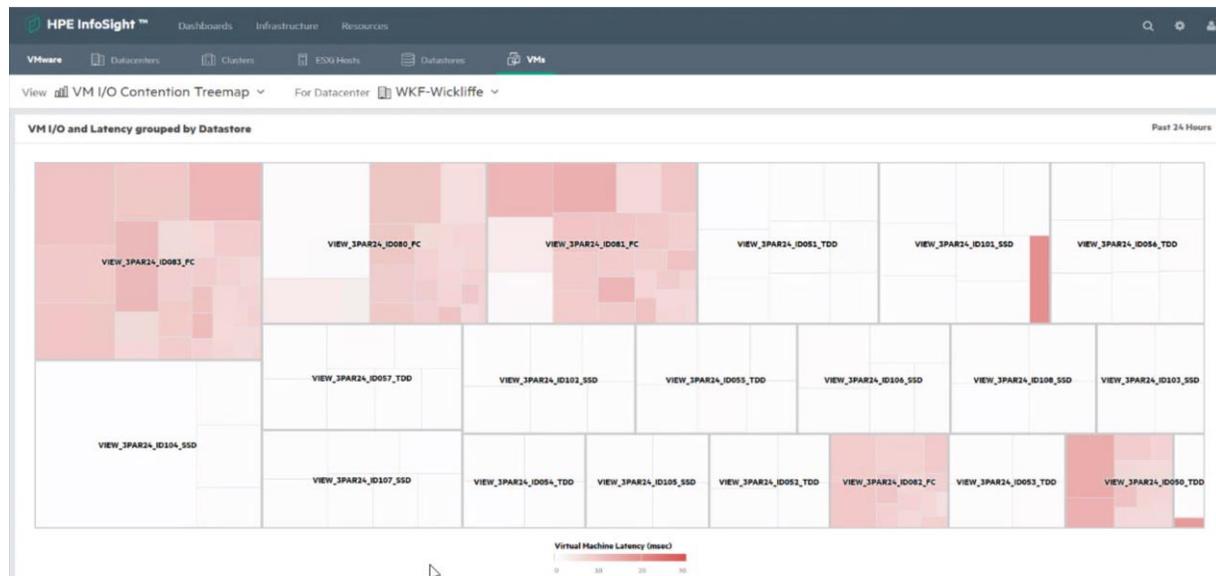


experienced a CPU spike to 70%. The HPE service team proactively reached out to Gary once the CPU utilization exceeded 50% for a sustained period. HPE InfoSight with cross-stack analytics capability helped worked through this issue and avoided the purchase of unnecessary hardware. Gary registered his system with the no-cost, SaaS-based HPE InfoSight portal that is available to use with every HPE Nimble and 3PAR Storage array to monitor performance. Although HPE InfoSight has numerous features, for this performance issue Gary decided to take advantage of the cross-stack analytics feature for more insight into the IO contention on his 3PAR 8200 Storage system.

HPE InfoSight Cross-Stack Views

HPE cross-stack analytics come with several views relative to a complex virtualized software stack. This complexity can be like the parable of the blind only feeling an elephant up close and therefore concluding it is something else. If you narrowly view an issue from one direction, you might not see the bigger picture. The first view Gary looked at was the I/O contention treemap based on I/O and latency issues as seen from a VM datastore. Figure 4 below gives a view of this for Gary's system.

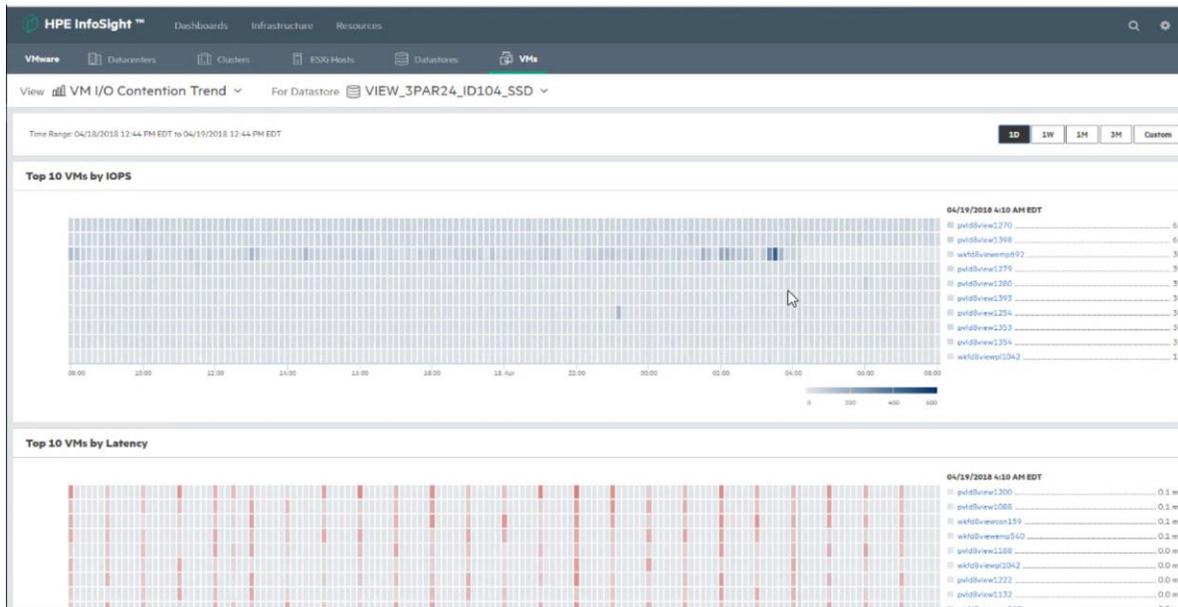
Figure 4: VM I/O and Latency Contention Treemap by Datastore



“THIS VIEW I’VE NEVER HAD BEFORE. THIS IS HUGE – A GAME-CHANGER FOR ME.” - GARY

The larger the rectangle, the more I/Os the datastore is responding to; and the darker red the color in the smaller rectangles, the more latency each VM client is seeing. This I/O contention view includes all the datastores for an instance of vCenter. For Gary, the upper left-hand corner revealed the datastore that was having the issues and belonged to the 3PAR 8200 servicing the VDI solution.. Once he identified the datastore that needed further analysis, he was able to look at that datastore from a variety of perspectives. Gary looks daily at the health of his environment using this I/O contention treemap view. Coincidentally, during the time of the 70% utilization issue, Gary was also getting complaints from his end-user base about VDI performance issues. The next step for Gary was to look at a few other views to get more clues as to what exactly was going on. Figure 5 showed Gary the set of top 10 VMs that, over a selectable period, had consumed the most IOPS or had the most latency issues. This particular view was during a time when things were running relatively stable. During stable times all clients operated under 100 IOPS; however, there was a periodic latency increase across all the clients every few hours.

Figure 5: VM I/O Contention Trend



592

During the times of stress, Gary found that several clients would increase IOPs significantly beyond normal levels. This led Gary to dig further into those clients to see what was happening. Gary identified those systems by using the InfoSight VM list feature that was sorted by performance (see figure 6). What he found with the help of the client team is that they had added some test Windows 10 VDI clients. Gary helped debug the issue by drilling down in these clients and found they had a version of an anti-virus tool that was taking an inordinate amount of system resources.. The security team got involved and helped to get new versions of the AV tool to reduce the issue. The security team is now evaluating new antivirus technology that will reduce the load across all systems. The client team is looking at a way to minimize the total amount of virtual desktops the company needs to deploy. Gary continues to use InfoSight to find those clients that have gone rogue, driving up the load on the 3PAR system. At times he finds some clients that have too little memory, resulting in memory swapping that can cause runaway IOPs. Figure 7 below shows how InfoSight gives detailed performance of a single VM all the way down to a datastore level.

Figure 6: VM List sorted by performance

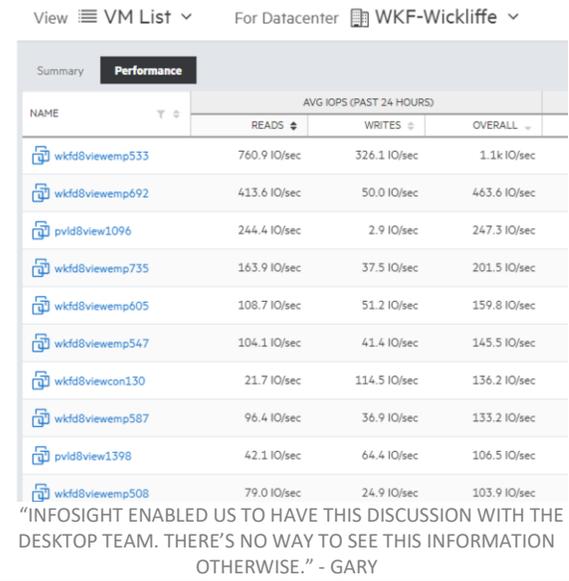
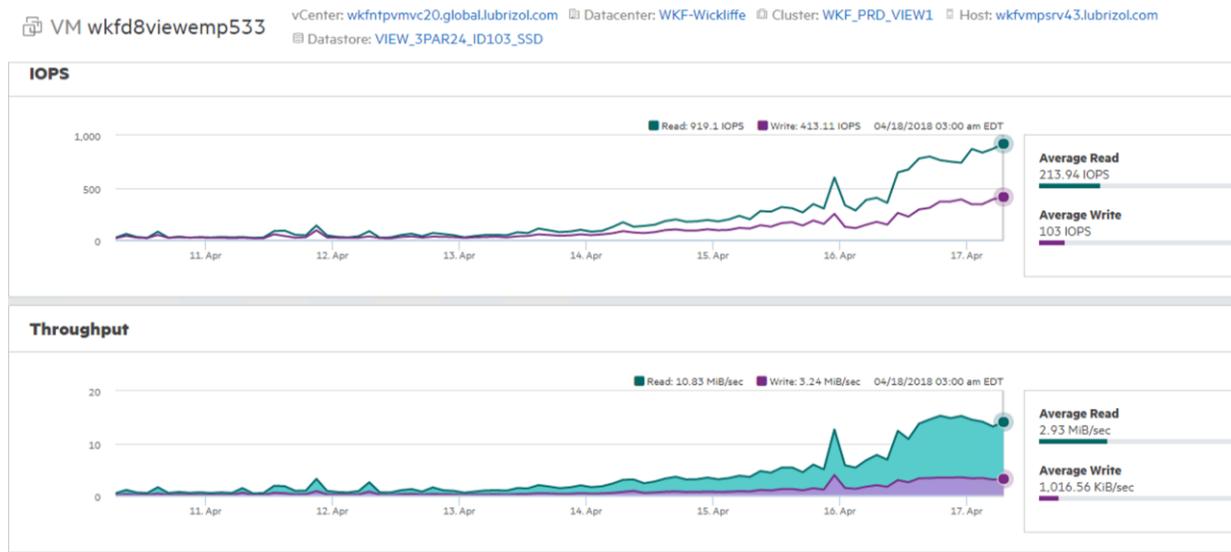


Figure 7: HPE InfoSight individual VM performance trends



“(INFOSIGHT) ENABLES ME TO DRILL DOWN FROM EACH CLIENT ALL THE WAY TO STORAGE.” -- GARY

Gary is an example of just one customer that uses HPE InfoSight daily to achieve goals such as rooting out performance issues that allow him to hold off on buying new infrastructure. In addition, as new clients are added, Gary can quickly provision them to an appropriate 3PAR system. HPE InfoSight is saving Gary much time and providing multi-dimensional insights not seen before, which has helped him save his company the cost of a new storage array. Gary looks forward to exploring how best to use HPE InfoSight to address other infrastructure issues as well.

“I’VE ONLY TOUCHED ON A REALLY SMALL SUBSET OF WHAT’S IN HERE – THERE IS SO MUCH MORE IN (HPE) INFOSIGHT WE COULD TAKE ADVANTAGE OF” -- GARY

TANEJA GROUP OPINION

Virtualized environments introduce complexity into an infrastructure stack, presenting a unique set of management and support challenges. In response, enterprise IT managers are increasingly looking for help from a new class of predictive analytics technologies. But as we’ve seen, only analytics solutions with certain characteristics are capable of addressing these issues in a complex, multi-layer virtual infrastructure. In particular, analytics must be cross-stack capable, enabling visibility into interactions and dependencies across multiple physical and virtual infrastructure layers over time.

HPE InfoSight stands out as a predictive analytics solution that is uniquely qualified to address the significant management and support demands of a virtual infrastructure. From the breadth and depth of its data collection to its advanced, cross-stack machine learning capabilities, HPE InfoSight is designed to cut through infrastructure complexity, enabling IT managers to reduce costs, increase availability, and optimize performance in their virtualized environments. As we’ve seen, InfoSight not only works in the background to anticipate and remediate looming support issues, but also provides IT with a hands-on management tool to investigate subtle or transient cross-layer performance issues and zero in on their root cause. This allows customers to increase the efficiency and extend the life of existing infrastructure, future-proofing their IT investments. The scope and longevity of collected data, built on a multi-year historical data set, further sharpens machine learning capabilities, allowing InfoSight to provide insights that other analytics tools cannot. In our surveys of the market, we have not seen any other analytics solutions that come close to matching the power, experience, and value of InfoSight in enterprise IT environments.

With a track record of proven success across thousands of HPE Nimble systems and a growing presence in HPE 3PAR Storage installations, HPE is well positioned to extend the capabilities of InfoSight to its server and networking product lines. We believe that as a next step, HPE will take advantage of its role as a strategic system-level supplier to instrument in depth all elements of a customer's infrastructure—including servers, networking and storage—and partner with key software providers to enable AI for the “self-driving” data center. This will expand InfoSight's competitive advantages beyond storage, while enabling HPE to play a more prominent role in managing and supporting the data center. More importantly, it will enable IT managers to expend less effort to better administer and support their infrastructure, freeing them up to spend time on higher-value initiatives.

NOTICE: The information and product recommendations made by Taneja Group are based upon public information and sources and may also include personal opinions both of Taneja Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Taneja Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.