

Insight and Management in the Virtual Infrastructure – what you can't see can hurt you.

February 2009

Looking around us, we've seen a change happen in the enterprise that has overtaken many administrators with silent footsteps. While the IT practitioner's every day is a swim through waves of invisible bits, there has long been some comfort to be found in the "physicality" and accessibility of key devices. When problems arise, we have always been able to identify a switch port for examination, a server at the end of a wire that might be causing problems, an HBA for inspection, or any number of other physical things we can turn to for further examination. But in today's data center, that comfort has long ago vanished.

In part, this is due to the virtualization of technologies surrounding us, and while this trend is spearheaded by server virtualization, the variations are almost uncountable, including the likes of application virtualization, network device virtualization, IO virtualization, storage virtualization, and more.

At the same time, today's IT infrastructure has scaled out of control, with more interdependencies between systems, while continuously operating closer to the performance limits of the infrastructure than ever before. For a given application, dependencies may cross multiple applications, servers, SAN fabrics, IO adapters, network hops, and even data centers. When the performance limit of a single component in one of these systems is breached, common experience indicates these IO laden systems will not see performance gradually degrade, but rather will see latency and performance rapidly spiral out of control under a bombardment of increasingly delayed, dropped and repeated IO attempts that can no longer be queued, cached, backed-off by congestion controls, or otherwise gracefully handled. Moreover, traditional management tools have quickly become obsolete, as today's systems operating at immense scale have outpaced their capabilities. Faced with complexity and potentially catastrophic impacts from any change, the administrator literally faces the unknown.

The problem is a lack of instrumentation

Regardless of whether systems are real or virtual, it has been hard to find any way to evaluate the big picture and then drill down to perform detailed inspection upon the infrastructure. We simply lack the instrumentation to give us the data for planning, performance management, routine monitoring, compliance, or troubleshooting.

How is the administrator dealing with these challenges today? From our conversations with end users, we estimate over 85% of businesses today rely upon their initial testing of known

I N D U S T R Y A R T I C L E

good configurations or arbitrary rules of thumb rather than real data when they manage and make decisions about their virtual infrastructures. This makes planning an exercise in waste that over-provisions resources, provides no guarantees that SLAs can be met, obfuscates troubleshooting, and restricts flexibility. Moreover, as configurations change over time, the administrator has no guarantees about what their current infrastructure capabilities are. We find that over 80% of VMware customers have still not deployed VMotion, and these users are most often deterred by their inability to determine performance impact when making changes to their virtual infrastructure.

The challenge facing administrators is really one of getting visible, meaningful data about their infrastructure on an on-going basis, and this leaves the IT administrator in the position of driving through unfamiliar terrain with fog on the windshield. But times are changing.

Bring on the visibility – Virtual Infrastructure Optimization

Today, many increasingly sophisticated technologies hope to clear this fog from the IT administrator's windshield. Such solutions fall under the Virtual Infrastructure Optimization banner. Virtual Infrastructure Optimization holistically assesses the entire virtual infrastructure, and provides the administrator with the data necessary to make intelligent decisions about capacity, utilization, and performance for every layer of the infrastructure - network, server, storage, and applications.

Some of the solutions in this area peer into single dimensions of the infrastructure, such as capacity, and spot clean the windshield. Tried and true solutions as well as newer vendors fall into this category, and our long list includes BMC's VSM solutions, Novell's PlateSpin, HP's VSE, VMware's Capacity Planner (from their acquisition of BeeHive), Virtugo's virtualSuite, Hyperic, TwinStrata's Clarity AP, and Computer Associates' ASM. While these solutions have a place in routine planning, they do not address an equal or more important aspect of today's virtual infrastructure management needs: *data about what is happening in the moment*.

We're more excited about a number of solutions that enable administrators to peer into multiple dimensions of the infrastructure in real or near real-time. These solutions deliver the integrated monitoring and analytics required to optimize or troubleshoot virtual infrastructure performance holistically, in real time, and across every involved system – from the application to the spindle. Moreover, such solutions provide the granular data necessary for good decision-making. Without intelligent assessment of performance, capacity and utilization can only be planned on top of assumptions that may or may not apply to any single system or that may change at any given point in time. In our view, real-time or near real-time performance-based analytics *is the required foundation for building and managing a virtual infrastructure*.

I N D U S T R Y A R T I C L E

Into this latter category we place products such as Virtual Instruments' real-time VirtualWisdom, Akorri's BalancePoint, Onaro's SANscreen VMInsight, and BlueStripe's FactFinder.

Why we need VIO in the infrastructure

It is obvious from this long list of vendors that VIO is a critical dimension of operating an infrastructure, but what is the set of challenges that these emerging VIO solutions can help the administrator tackle better than other solutions? In our view, they build a more comprehensive data set that is the right foundation for virtual infrastructure management. To shed light on this, we've identified a set of 5 key strategic questions that every end user should ask themselves about their virtual infrastructure management tools. These questions will help end users identify the benefits of VIO technologies, and assess how well their VIO or management platform will address the universal pain points common to increasingly virtual infrastructures.

Can I efficiently plan, design, and make decisions about my infrastructure that I know with certainty will make the best use of expensive systems, and have the desired result 100% of the time?

VIO will decrease operational cost and complexity around planning infrastructure and making operational changes. Today, needless hours are spent identifying initial configurations, and assessing the potential performance and utilization impacts each time a change is required. VIO will provide a data set that gives the administrator visibility into infrastructure capabilities and impacts at the touch of a button, and leave no doubt about how to optimize system configurations.

Can I immediately drill into the root cause of performance issues in my environment, and discover what happened or changed?

While a VIO solution may arm an organization with the right data to avoid misconfigurations in the first place, VIO solutions can also provide real or near real-time visibility into what is happening in an environment, enabling administrators to immediately identify performance anomalies and root causes. The right VIO solutions can capture history, providing an audit trail that identifies when problems started, and what happened.

Can I tell with certainty that I am getting the optimal use of expensive servers, storage, hypervisors, and other infrastructure?

The right VIO solution will identify operational peaks and troughs, and save capital dollars that are today wasted on over-provisioning systems. The right near real-time or real-time solutions can help dynamically but intelligently balance an infrastructure, making sure that even the slimmest infrastructure never runs into performance problems.

I N D U S T R Y A R T I C L E

Can I tell with certainty that virtual technologies don't adversely impact my infrastructure, or concretely identify issues without vendor finger pointing?

VIO solutions can give you the data you need to hold vendors accountable, and determine which systems are the root causes of problems or are having detrimental impacts upon your infrastructure.

Can I safely implement and make use of the full capabilities of the virtualization technology that has been purchased to increase my operational efficiency and improve IT capabilities?

Concern about the inability to see what happens when a technology is used is at the heart of why technologies like VMotion are so infrequently used today. Without the right supporting data, well-meaning automation frameworks and policies can just as often be dangerous landmines when they unintelligently take action on an infrastructure. It is no wonder that policies and tools for VMotion automation, storage changes, guest reconfigurations, IO management, and other solutions are not widely used today. But the right VIO solution can provide the right intelligence to responsibly take action while avoiding potentially catastrophic impacts. VIO provides predictive visibility into what will happen when changes are made, and provides assurances that what was expected did indeed happen.

VIO for better control

If you answer any of these questions with a no, and/or these issues sound familiar, then you may indeed be driving your infrastructure over unfamiliar terrain and through a foggy windshield. The cost of doing so in wasted time and effort or wasted resource utilization can be enormous. The broad number of VIO solutions available today take aim at permanently clearing away this fog. But in our book, some excel better than others for this particularly sticky set of problems that face the virtual infrastructure. These solutions build VIO on top of deep, on-going, performance-centric analysis and correlation of configurations, changes, and events across all layers of the datacenter – from the application layer to the storage spindles, irrespective of application, hypervisor, OS, network, storage, or SAN vendor. Using collected, instrumented data about cross-domain performance and events that are often captured directly over the wire, these VIO solutions can act as master consoles for monitoring operations, viewing configurations, and triggering changes across the entire datacenter. Much more than lifecycle management, configuration management, or other point solutions that each may have unintended repercussions by unintelligent automation, VIO is about total datacenter orchestration. Let's take a look at the core capabilities of this type of solution, and how we see current emerging solutions differentiating themselves around each of these capabilities.

Key Capability 1: Detailed visibility through actual data

In our view, VIO solutions that capture real-time data provide the best foundation for holistic infrastructure management, as only this level of detail can capture the changes in systems that occur over time and provide the detail necessary to detect and correct issues caused by changes. VIO solutions available today will vary in several dimensions when it comes to capturing detailed data:

First administrators should assess how much detail a solution captures. Periodic sampling in conjunction with sophisticated algorithms may appear to provide a good foundation for planning, but not provide real-time visibility for troubleshooting. Today, we're seeing an increasing number of solutions on the market, such as Virtual Instruments' VirtualWisdom, that can take action on real-time data. Second, administrators should assess how a solution captures data. Some VIO solutions rely on agent approaches, some passively poll data through available APIs, and some capture data over the wire. For each approach, the complexity of deployment, on-going management, how much detail is captured, and how potentially enormous sets of data are transferred across an infrastructure and reported upon should be examined.

Key Capability 2: Holistic view of the entire virtual infrastructure

VIO solutions must encompass every layer of the enterprise. Taneja Group has often written about the increasing importance of cross-domain correlation technologies. With VIO, cross-domain correlation is absolutely fundamental to taking intelligent action without guesswork about consequences. Only with this comprehensive set of data can an administrator take action to correct one problem, such as excessive LUN traffic, without being worried that they might create another, such as excessive switch port traffic. But VIO solutions vary in their ability to capture comprehensive data, sometimes being limited to select operating systems or storage devices. Such solutions are challenged by their inability to see the full environment, but may be able to compensate through examination of behavior patterns with sophisticated algorithms. Users should evaluate the sophistication of varying approaches, and whether a solution provides visibility into the right aspects of their infrastructure.

Key Capability 3: Actionable Correlated Data

While holistic visibility is one thing, making data actionable is another. Solutions vary in how much data they deliver, which systems they capture data from, and what type of action they can enable or trigger. Solutions that capture over the wire may have an edge here, as they are rest assured that they are seeing everything in an environment, and not just the traffic or performance a single node is exposed to. When a solution is built upon real-time data, it can provide an important perspective into on-going performance, simplify troubleshooting by an order of magnitude, and enable auditing for SLAs or for compliance purposes.

Key Capability 4: Extensibility and Interoperability

Finally, solutions will vary in how well and deeply they integrate with enterprise systems. VIO solutions will vary in their ability to trigger actions themselves. Moreover, with the right capabilities, real-time VIO solutions can provide the right data and APIs to enable better use of other tools such as HP Openview, IBM Tivoli, CA Unicenter, customized scripts, or VMware vCenter Server. The end user should assess their need to integrate a solution with other technologies such as storage management tools, the hypervisor, virtual switches, and various other technologies. Look for extensibility in a solution that matches your needs.

Finding clarity, or living with fog

Some solutions under the umbrella of VIO try to manage the fog on your windshield by providing a framework for automatic navigation, without shedding much light on what is going on around you. You still can't see, but these vendors can provide tremendous management frameworks, that can help you automate routine actions and maintain adherence to policies and best practices. Unfortunately, the detailed visibility into what is happening within the data center is still lacking. The performance management oriented VIO solutions in this space actually clear away the fog, and let you see the road in the form of meaningful, correlated data across multiple elements in the data center. Differentiation today is found in the form of how much detail can be seen, and how that vendor turns that detail into an actionable set of data, either by providing you with meaningful analysis and summarization, or integrating with other systems to automatically take action, or both.

In our view, Virtual Infrastructure Optimization will be one of the most pivotal technologies at play in defining the capabilities of the next generation data center. While today this technology is about creating real-time data center intelligence that can be a foundation for responsible infrastructure actions and reactions, we believe this technology will eventually instigate the evolution of more capable and deeply integrated orchestration across all systems in the data center and help create increasingly autonomic, elastic computing that will respond to dynamically changing business demands. While the future potential of these technologies seems tremendous, they are here and providing key visibility today. From our perspective, these solutions are simply a cornerstone of any virtualization strategy today, and should be a required component for every data center management tool set.

***NOTICE:** The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the 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. The 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 which may appear in this document.*