



Bridging the Gap between Application Performance and Virtual/Physical Infrastructure Management Silos

By Chia-Chee Kuan

Abstract

Application performance is critical to business objectives, and data center managers are tasked with optimizing application performance to the fullest. However, application servers are typically provisioned with virtualized compute, storage and network resources, and the ever-increasing challenge to data center IT operations teams is the ability to grasp virtual infrastructure's impact on application performance. Virtualization management teams simply do not have the real-time application level visibility and analytics required to ensure that application performance is optimized and well supported by data center virtual infrastructure and the underlying physical resources.

In this article, we'll look at the challenges in managing and optimizing virtualized data center infrastructure, and the requirements for a unified full stack application, virtualization, and infrastructure visibility solution.

Application Interdependencies and Visibility

Application visibility including interdependencies among application services themselves as well as infrastructure has long been a cap for IT infrastructure managers. Virtual resources and allocations most directly impact application performance, long before physical resources impact that performance. A physically over provisioned data center can be a virtually under provisioned infrastructure, causing application performance bottlenecks unresolvable by adding unlimited amount of physical resources. For example, if an application is not allocated adequate vCPU (virtual CPU) or memory, its performance will be adversely impacted regardless of the amount of physical CPU resources available.

Unfortunately, data center management occurs in silos (virtualization management, storage management, network management, server management, and application management), and it is difficult for any one discipline to see outside its silo. What's needed is instrumentation that can see across these multiple silos.

The virtualization layer is the center of the data center and is best suited for instrumentation for all-inclusive operational visibility. Above the virtualization layer are the applications and guest OS, and below the virtualization layer is the physical layer (compute, storage, and network resources). The virtualization layer itself contains the virtual machines, virtual storage, and virtual networking resources.

Now, let's look at the requirements for visibility in each of the management silos.

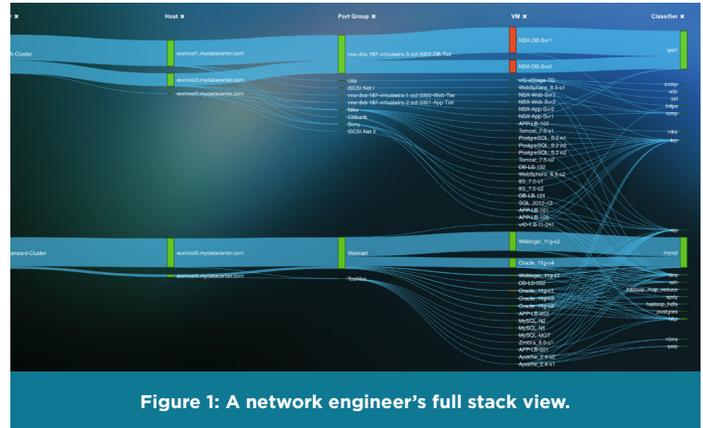
The Network Engineer's View

The network engineer's full stack visibility should include:

- Physical network switch & port connectivity
- Physical host network interface
- Virtual Switch (vSwitch or Distributed vSwitch)
- Virtual Networks (VLAN or VXLAN or SDN)
- Virtual Machine network interface
- Virtualized Application service and topology

Such a network view might look like the one in Figure 1.

A network engineer should be able to view the entire stack from the physical layer to the applications running on VMs. This should include the following as illustrated in Figure 1: Top of Rack switch, physical hosts connected to the Top of Rack switch, Distributed vSwitches and port groups (VLAN, VxLAN, etc), virtual machines deployed on the physical hosts and the kinds of application traffic passing through all these layers. This at-a-glance view of the full stack enables the Network Engineer to quickly identify the VM and application services that have issues (color coded in red) and correlate them to the physical host and the switch that are associated with this VM (application servers).



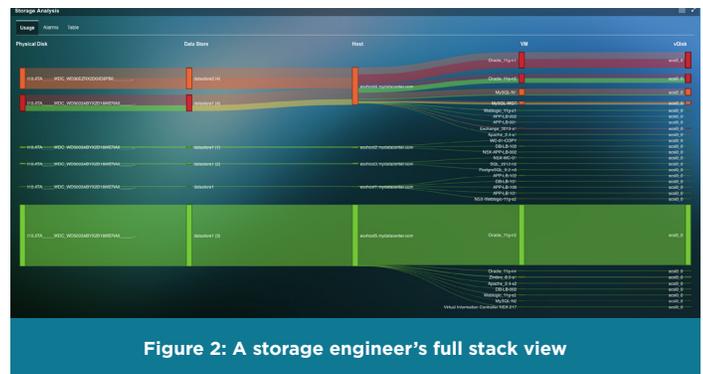
The Storage Engineer's View

For the storage engineer, full stack visibility should include:

- Physical and centralized storage devices (fiber channel, iSCSI, etc.)
- Physical host storage
- Virtual storage (logical datastore)
- Virtual Machine storage (vDisk)
- Virtualized Application services

Such a view might look like the one in Figure 2.

A storage engineer should be able to view the entire stack all the way from the data store to the performance relating to storage of applications running on VMs. This should include the following as illustrated in Figure 2: Logical data store (which groups together a number of physical hosts), physical hosts associated with these data stores, virtual machines deployed on the physical hosts and the virtual disk accessed by these VMs. This at-a-glance view of the full stack enables the storage engineer to quickly identify the vDisk and VM that have issues (color coded in red) and correlate it to the logical data stores and physical hosts.



A Server Administrator's View

For the server administrator, the full stack view should include:

- Physical host clustering
- Physical host

Such a view might look like the one in Figure 3.

A server administrator should be able to view all the clusters in a virtualized data center. This should include the following as illustrated in Figure 3: the physical hosts, associated VMs operating on these hosts as well as the applications that are running on these VMs. With this view, any application performance issue can easily be visualized and tracked to the virtual and physical infrastructure hosting the application. Key metrics such as response time, transactions per minute, and traffic for each application enables the Server Administrator to proactively identify performance issues.

- Virtual Machine
- Virtualization application services

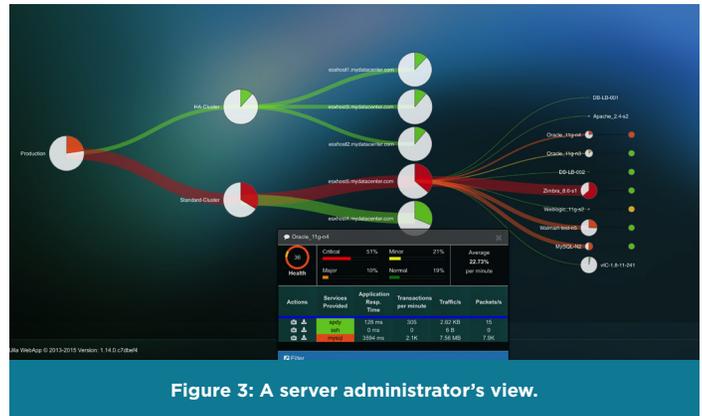


Figure 3: A server administrator's view.

A Virtualization Engineer's View

For the virtualization engineer, the view should include:

- Automatically identified applications and services
- Application and virtualized resource usage and performance heatmap
- Application service topology

Such a view might look like the one in Figure 4.

A virtualization engineer's full stack view should consist of key metrics about application performance as well as the virtual and physical compute, storage and network layers that support these applications. This converged data center full stack visibility that cuts across compute, network and storage domains with contextual metrics across these layers enables the virtualization engineer to spot root causes and resolve application performance issues in the underlying infrastructure very quickly, and to manage the data center with confidence.

As mention in the beginning of this article, application interdependencies are key to application and data center performance management. Such interdependencies and analytics are best attained with application topology visualization and dependency maps such as the one illustrated in Figure 5.

- Application to infrastructure dependency map
- Correlated metrics from application response time to compute, storage & network



Figure 4: A virtualization engineer's view

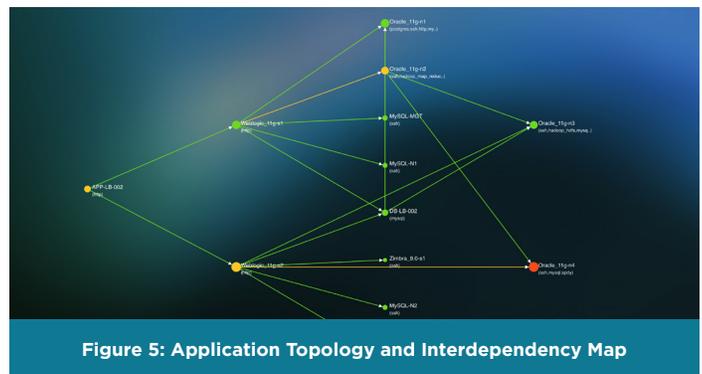
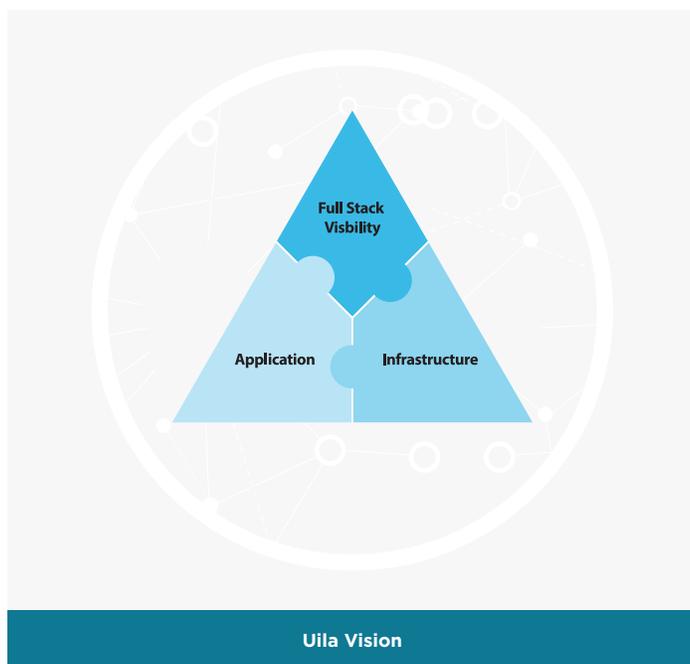


Figure 5: Application Topology and Interdependency Map

Eliminating Data Center Management Silos

Data centers rely on interdependent layers (physical, virtual, and application), and these must all be optimized to support optimum application performance. A proper monitoring and management solution will cover the silos vertically, from the application through virtualization to the physical infrastructure, and horizontally (covering compute, storage, and networking at the virtual and physical layers). The right solution will provide infrastructure views for the network engineer, storage manager, server administrator, and virtualization manager, establishing a common language for these separate teams and making it possible to better optimize application performance.

Chia-Chee Kuan is CEO of Uila, Inc. He has a distinguished 20-plus year career developing unique intellectual property and building game-changing companies focused on computer networking technologies and industries. From early day pre-Internet TCP/IP stack and standards development, to IP multicast video streaming applications, to Wi-Fi security, Mr. Kuan is a pioneering technologist who has been granted 14 U.S. patents. After co-founding AirMagnet in 2001 and merging with Fluke Networks in 2009, Mr. Kuan was named General Manager for the Wireless Business Unit, doubling revenue during his four-year tenure at Fluke Networks. Prior to AirMagnet, he was a founding member at Precept Software, which was later acquired by Cisco Systems, where Mr. Kuan continued to lead the Cisco Video Internet Service Unit product development team. Mr. Kuan holds a Master's degree in Computer Science from Stanford University and a Bachelor's in Information Engineering from National Taiwan University.



Uila, Inc.

2905 Stender Way, Suite 76E
Santa Clara, California 95054

www.uila.com
(408) 819-0777
sales@uila.com

