The 2010
Cloud Networking Report

Dr. Jim Metzler, Ashton Metzler & Associates

Management & Call to Action

Management

Background

As pointed out in ***Virtualization: Benefits, Challenges and Solutions***, server virtualization creates a number of management challenges. For example, the need to manually reconfigure the network to support VM migration that was previously mentioned can be regarded as either a LAN challenge or a management challenge. Additional management challenges include:

- **Breakdown of Network Design and Management Tools**
  The workload for the operational staff can spiral out of control due to the constant stream of configuration changes that must be made to the static date center network devices in order to support the dynamic provisioning and movement of VMs.

- **Limited VM-to-VM Traffic Visibility**
  The first generation of vSwitches doesn't have the same traffic monitoring features as does physical access switches. This limits the IT organization's ability to do security filtering, performance monitoring and troubleshooting within virtualized server domains.

- **Poor Management Scalability**
  Many IT organizations have experienced VM proliferation sometimes called VM sprawl. In addition, the normal best practices for virtual server configuration call for creating separate VLANs for the different types of traffic to and from the VMs. The combined proliferation of VMs, and VLANs places a significant strain on the manual processes traditionally used to manage servers and the supporting infrastructure.

- **Multiple Hypervisors**
  It is becoming common to find IT organizations using multiple hypervisors, each of which comes with their own management system and their own management interface. In addition, the management functionality provided by each hypervisor varies as does the degree to which each hypervisor management system is integrated with other management systems.

- **Management on a per-VM Basis**
  IT organizations typically perform management tasks such as discovery, capacity planning and troubleshooting on a per server basis. While that is still required, IT organizations must also perform those tasks on a per-VM basis.

The Webtorials Respondents were asked to indicate how important it is to their organization to get better over the next year at managing some key tasks related to server virtualization. They were given the following five-point scale:

1. Not at all important
2. Slightly important
3. Moderately important
4. Very important
5. Extremely important

Included in Table 6.1 are the tasks and the percentage of The Webtorials Survey Respondents who indicated that the task was either very or extremely important for their organization to get better at over the next year.

Server Virtualization Management Task	Importance: Very or Extremely
Perform traditional management tasks such as troubleshooting and performance management, on a per VM basis	49%
Keep track of VMs as they are moved between physical servers	38%
Dynamically move VMs, and all of the supporting management functionality, between physical servers	37%
Discover VMs	33%
Manage the traffic that goes between virtual machines (VMs) on a single physical server	31%

Table 6.1: Importance of Managing Server Virtualization

As shown in Table 6.1, The Webtorials Survey Respondents indicated that getting better at each of the individual challenges associated with server virtualization is important to their organization. In addition, it is reasonable to look at the five challenges contained in Table 6.1 as being a single challenge - managing server virtualization. When looked at that way, getting better at server virtualization is extremely important to The Webtorials Survey Respondents.

The Evolving Management Environment

One of the primary management challenges associated with cloud networking is the movement to focus on managing services as defined in section two of this report. To put that challenge into perspective, The Webtorials Survey Respondents were asked to indicate the approach their organization takes to management. They were given the following choices and allowed to choose all that applied to their environment.

- We have a focus primarily on individual technology domains such as LAN, WAN and servers

- We have a focus on managing the performance of applications as seen by the end user
- We have a focus on managing the performance of services as seen by the end user, where service refers to multiple, inter-related applications
- Other

Their responses are shown in Figure 6.1.

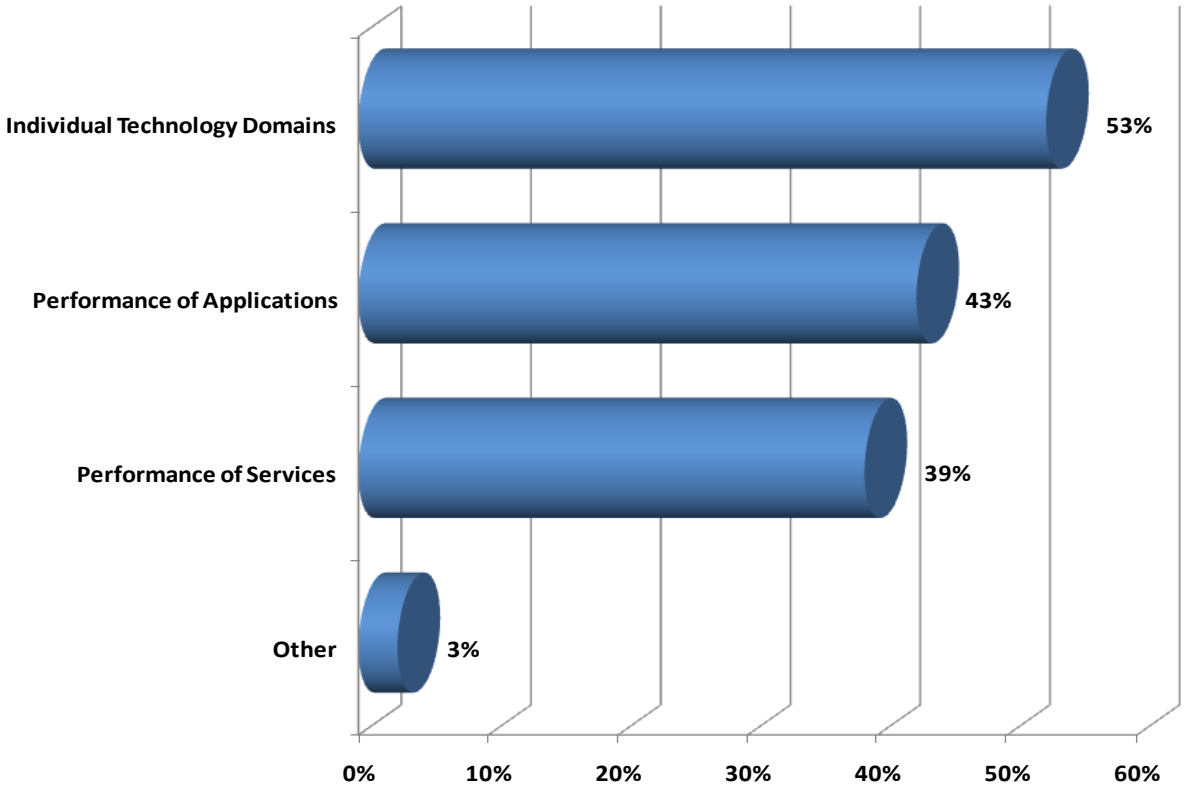

Figure 6.1: Focus of Management

The data in Figure 6.1 indicates that the most frequent approach that IT organizations take to management is to focus on individual technology domains. However:

A significant percentage of IT organizations focus their management activities on the performance of applications and/or services.

The Webtorials Survey Respondents were also asked to indicate how important it is to their organization to get better at twenty different management tasks over the next year. They were given the same five-point scale as was discussed relative to Table 6.1. Included in Table 6.2 are the tasks and the percentage of The Webtorials Survey Respondents who indicated that the task was either very or extremely important for their organization to get better at over the next year.

Management Task	Importance: Very or Extremely
Rapidly identify the root cause of degraded application performance	76%
Identify malicious traffic and eliminate it	71%
Effectively manage QoS	67%
Prevent large scale DDOS attacks	66%
Identify the components of the IT infrastructure that support the company's critical business applications	66%
Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems	64%
Effectively manage services, where services are comprised of multiple, inter-related applications	61%
Effectively manage SLAs for one or more business critical applications	61%
Obtain real-time, or nearly real-time, insight into how specific applications and end user sessions are performing	59%
Track end user experience and relate it to factors such as Web response time	51%

Table 6.2: The Importance of Improving Management Tasks

More detail on the management challenges facing IT organizations can be found in the report ***Application Delivery: A Reality Check***.

The fact that sixty one percent of The Webtorials Respondents indicated that effectively managing services, where services are comprised of multiple, inter-related applications was either very or extremely important to them underscores the growing importance of IT organizations having a focus on services.

Communications Based Applications

The fact that two thirds of The Webtorials Respondents indicated that effectively managing QoS was either very or extremely important to them is consistent with some of the results contained in the previous section of this report. In particular, the data in Table 5.8 indicates that two thirds of The Webtorials Respondents indicated that ensuring acceptable performance for VoIP traffic was either very or extremely important to them. In addition, over half of The Webtorials Respondents indicated that ensuring acceptable performance for video or telepresence traffic was either very or extremely important to them. QoS is an important technique to allow IT organizations to ensure the performance of communications based applications such as VoIP, video and telepresence.

The survey data discussed in the preceding paragraph clearly indicates the importance of managing communications based applications such as VoIP, traditional video conferencing, as well as high definition video conferencing and telepresence. This importance will only grow as the use of those services grows. To illustrate the growth in communications based applications, Table 6.3 shows the percentage of The Webtorials Respondents that currently make at least some use the application and the percentage that expect to increase their use of the application over the next year.

Application	Percent Currently Using	Percent that Expect to Increase Use
Traditional voice (POTS)	93.0%	18.2%
VoIP	91.6%	76.9%
Traditional videoconferencing	86.2%	36.5%
HD Videoconferencing or telepresence	60.1%	58.3%

Table 6.3: Use of Communications Based Applications

Internal SLAs

As recently as two years ago, few IT organizations offered an SLA to the company's business and functional managers; a.k.a., an internal SLA. However, that situation has changed and now it is common for IT organizations to offer internal SLAs. To understand the prevalence and effectiveness of internal SLAs, The Webtorials Respondents were asked to indicate their agreement or disagreement with three statements. The three statements and the percentage of The Webtorials Respondents that agreed with the statement are shown in Table 6.4.

Statement	Percent that Agree
We provide an SLA internally for every application that we support	30.0%
We provide an SLA internally for at least some applications	69.9%
We do a good job of managing our internal SLAs	55.8%

Table 6.4: Status of Internal SLAs

The data in Table 6.4 highlights the growing interest that IT organizations have in providing internal SLAs for at least some applications. However, as previously noted, the SLAs that are associated with public cloud computing services such as Salesforce.com or Amazon's Simple Storage System are generally weak or non-existent.

The lack of meaningful SLAs for public cloud services is a deterrent to the Global 2000 adopting these services for delay-sensitive, business-critical applications.

Root Cause Analysis

It is not surprising that rapidly identifying the root cause of degraded application performance is so important to IT organizations in part because on an ever increasing basis a company's key business processes rely on a handful of business critical applications. That means that if those applications are not running well, neither are those key business processes.
Even in the traditional IT environment[1] when the performance of an application is degrading the degradation is typically noticed first by the end user and not by the IT organization. In addition, when IT is made aware of the fact that application performance has degraded, the process to identify the source of the degradation can be lengthy.

Unfortunately:

The adoption of cloud computing makes troubleshooting application performance an order of magnitude more difficult.

For example, assume that a hypothetical company called SmartCompany has started down the path of implementing private cloud computing by virtualizing their data center servers. Further assume that one of SmartCompany's most important applications is called BusApp and that the users of the application complain of sporadic poor performance and that BusApp is implemented in a manner such that the web server, the application server and the database server are each running on VMs on separate physical servers which have been virtualized using different hypervisors.

One of the challenges associated with managing a private cloud environment was referred to in section 2 of this report. In particular, due to the combination of the dynamic nature of IP and the meshed nature of enterprise networks, it is often not possible in a traditional IT environment to know what path the traffic took from origin to destination. This lack of knowledge complicates tasks such as troubleshooting. The difficulty of knowing the path from origin to destination is greatly increased in a cloud environment because services can be dynamically moved between servers both within and between data centers.

Another one of the challenges associated with managing a private cloud environment is that in order to manage BusApp in the type of virtualized environment described above, an IT organization needs detailed information on each of the three VMs that support the application and the communications amongst them. For the sake of example, assume

[1] This refers to an IT environment prior to the current wave of virtualization and cloud computing.

that the IT organization has deployed the tools and processes to gather this information and has been able to determine that the reason that BusApp sporadically exhibits poor performance is that the application server occasionally exhibits poor performance. However, just determining that it is the application server that is causing the application to perform badly is not enough. The IT organization also needs to understand why the application server is experiencing sporadic performance problems. The answer to that question might be that other VMs on the same physical server as the application server are sporadically consuming resources needed by the application server and that as a result, the application server occasionally performs poorly. Part of the challenge associated with troubleshooting this scenario is that as previously noted, in most cases once an IT organization has virtualized its servers it loses insight into the inter-VM traffic that occurs within a physical server.

Staying with this example, now assume that SmartCompany has decided to evaluate the viability of deploying BusApp using either a public or hybrid cloud computing solution.

One of the fundamental issues relative to managing either a public or hybrid cloud computing service is that the network topology becomes even more complex and hence understanding the end-to-end path becomes more difficult.

For example, as described below, in some hybrid cloud environments the client request goes to a web server hosted by a cloud computing service provider that then queries the internal database over an MPLS network. Part to the complexity here is a result of the fact that:

- For added reliability, there are likely to be multiple BGP-based Internet peering points
- MPLS is difficult to manage using most traditional management tools
- All of the tiers of the application (i.e., Web, application, database) are running on VMs that are being dynamically moved between servers, both within and between data centers.

Another fundamental issue relative to managing either a public or hybrid cloud computing service is that the service has at least three separate management domains: the enterprise, the WAN service provider(s) and the various cloud computing service providers.

For the sake of this example, consider two alternative approaches that SmartCompany might evaluate. Those approaches are:

1. Public Cloud Computing

SmartCompany acquires BusApp functionality from a SaaS provider. The employees of SmartCompany that work in branch and regional offices use an MPLS service from a network service provider (NSP) to access the application, while home office workers and mobile workers use the Internet.

2. Hybrid Cloud Computing

SmartCompany hosts the application and data base servers in one of their data centers and the web servers are provided by a cloud computing service provider. All of the users access the web servers over the Internet and the connectivity between the web server layer and the application server layer is provided by an MPLS service.

In order to monitor and manage either deployment, consistent and extensive management data needs to be gathered from the cloud computing service provider(s), the MPLS provider(s) and the provider(s) of Internet access. In the case of the first option (public cloud computing) similar management data also needs to be gathered on the components of the on-site infrastructure that are used by SmartCompany's employees and supported by the IT organization. In the case of the second option (hybrid cloud computing) similar management data also needs to also be gathered on both the on-site infrastructure as well as the web and application servers that are supported by the IT organization. In either case, effective tools are also necessary in order to process all of this data so that IT organizations can identify when the performance of the application is degrading before end users are impacted and can also identify the root cause of that degradation.

Management Solutions

Route Analytics

In a traditional IT environment it is sometimes difficult to know the end-to-end path that packets take across a network. This management complexity comes in part from the distributed nature of IP. In particular, routers exchange reachability information with each other via a routing protocol such as OSPF (Open Shortest Path First). Based on this information, each router makes its own decision about how to forward a packet. There is, however, no single repository of routing information in the network. As described in the preceding section, the difficulty of understanding the end-to-end path is magnified in a cloud network.

As shown in Figure 6.2, route analytics provides IT organizations and service providers with insight into the routing layer.

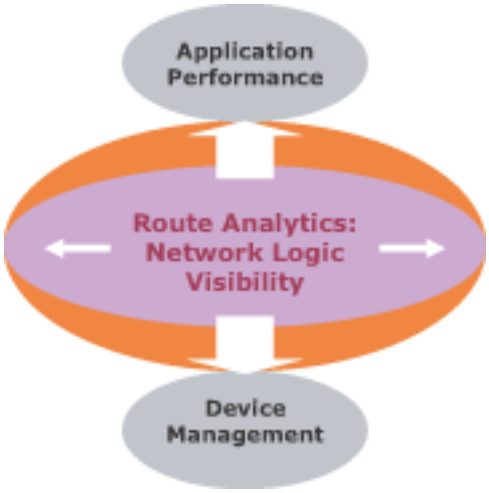

Figure 6.2: The Positioning of Route Analytics

In particular, route analytics provides visibility, analysis, and diagnosis of the issues that occur at the routing layer in complex, meshed networks. A route analytics appliance draws its primary data directly from the network in real time by participating in the IP routing protocol exchanges. This allows the route analytics device to compute a real-time Layer 3 topology of the end-to-end network, detect routing events in real time, and correlate routing events or topology changes with other information, including application performance metrics. As a result, route analytics can help both IT organizations and service providers determine the impact on performance of both planned and actual changes in the Layer 3 network.

Dynamic Infrastructure Management

A dynamic virtualized environment can benefit greatly from a highly scalable and integrated DNS/DHCP/IPAM solution, which is also well integrated with the virtual server management system.

Where DNS/DHCP/IPAM share a common database, the integration obviates the need to coordinate records in different locations and allows these core services to accommodate any different addressing and naming requirements of physical and virtual servers. Potential advantages of this approach include the automated generation of IP addresses for newly created VMs, the automated allocation of subnets for new VLANs, and the population of an IP address database with detailed information about the current location and security profiles of VMs. The integration of infrastructure utilities with the virtual server management system can also facilitate automated changes to the DHCP and DNS databases.

Virtualized Performance and Fault Management

Another example of a management capability in the traditional physical environment that is important to implement in a virtual environment is adaptive performance thresholding. This capability identifies systemic deviations from normal as well as time over threshold violations, and can automatically update thresholds based on changes to historic levels of utilization. That same capability is needed in a virtualized environment so that IT organizations can monitor the performance of individual VMs.

Virtual switches currently being introduced into the market can export traffic flow data to external collectors in order to provide some visibility into the network flows between and among the VMs in the same physical machine. Performance management products are currently beginning to leverage this capability by collecting and analysing intra-VM traffic data. Another approach to monitoring and troubleshooting intra-VM traffic is to deploy a virtual performance management appliance or probe within the virtualized server. This approach has the advantage of potentially extending the fault and performance management solution from the physical network into the virtual network by capturing VM traffic at the packet level, as well as the flow level.

While changes in the virtual topology can be gleaned from flow analysis, a third approach to managing a virtualised server is to access the data in the virtual server management system. Gathering data from this source can also provide access to additional performance information for specific VMs, such as CPU utilization and memory utilization.

Orchestration and Provisioning

> ***Service orchestration is an operational technique that helps IT organizations automate many of the manual tasks that are involved in provisioning and controlling the capacity of dynamic virtualized services.***

By automatically coordinating provisioning and resource reuse across servers, storage, and networks, service orchestration can help IT organizations streamline operational workloads and overcome technology and organizational silos and boundaries. Orchestration engines use business policies to define a virtual service and to translate that service into the required physical and virtual resources that are needed for deployment. The orchestration engine then disseminates the needed configuration commands to the appropriate devices across the network in order to initiate the requested service. The orchestration engine can automatically initiate the creation of the required virtual machines while simultaneously deploying the network access and security models across all of the required infrastructure components. This includes routers, switches, security devices, and core infrastructure services. The entire process can allow setup and deployment of network routes, VPNs, VLANs, ACLs, security certificates, firewall rules and DNS entries without any time consuming manual entries via device-specific management systems or CLIs.

Orchestration engines are generally limited in the range of devices with which they can interface due to differences in device and/or vendor management interfaces. Therefore, orchestration solutions mirror to some extent the constraints of virtual data center solutions that result from vendor partnerships among manufacturers of virtual server software, networks, and networked storage. The initial focus of such partnerships has been on promulgating validated network designs and architectures rather than on fully integrated or automated management. The next logical step for such partnerships is to include orchestration capabilities.

Orchestration solutions would benefit greatly from the emergence of an open standard for the exchange of information among the full range of devices that may be used to construct a dynamic virtual data center. In the Cloud Computing arena there are a number of standards under development, including the Open Cloud Computing Interface (OCCI) from the Open Grid Forum. These standards activities may also provide value within the enterprise virtual data center, since the stated scope of the specification is to encompass "all high level functionality required for the life-cycle management of virtual machines (or workloads) running on virtualization technologies (or containers) supporting service elasticity".

IF-MAP is another emerging standard proposed by the Trusted Computing Group and implemented by a number of companies in the security and network industries. It is a publish/subscribe protocol that allows hosts to lookup meta-data and to subscribe to service or host-specific event notifications. IF-MAP can enable auto-discovery and self-assembly (or re-assembly) of the network architecture. As such, IF-MAP has the potential to support automation and dynamic orchestration of not only security systems but also other elements of the virtual data center. For example, IF-MAP could facilitate automation of the processes associated with virtual machine provisioning and deployment by publishing all of the necessary policy and state information to an IF-MAP database that is accessible by all other elements of the extended data center.

Summary & Call to Action

For the foreseeable future IT organizations will increasingly adopt cloud computing. Cloud networking is the LAN, WAN and management functionality that enables IT organizations to support cloud computing. The key characteristics of a cloud network are that a cloud network:

- Has the same goal as cloud computing
- Supports the characteristics of a cloud computing solution
- Does no harm to cloud computing solutions
- Provides solutions that are good enough

Data Center LANs

The majority of IT organizations have the goal of evolving their data center infrastructure to be one that can dynamically provide each application and network service with the required resources. Because of the complexity and risk associated with achieving that goal, many IT organizations will choose to implement a Greenfield data center and then gracefully cut over to a new data center.

IT organizations that are going to implement a Greenfield data center should:

- Give primary consideration to two-tier designs that avoid the spanning tree protocol and which are based on switch virtualization and MC LAG.

- Consider new servers with multi-core processors that are capable of supporting a large number of VMs, and which incorporate dual 10 GbE LAN connections. For example, with four-processor, 48 core servers, a realistic goal for the number of VMs per server would fall in the range of ten to fifty, depending on the characteristics of the applications.

- Base the data center LAN design on both high-density modular access and core switches that can provide non-blocking support for 40 and 100 GbE when available. If TOR switches are used due to cost or cabling considerations, IT organizations should make sure that the switches can provide 10 GbE server connections and can support the desired over-subscription ratios.

- Acquire switches from a vendor whose product roadmap includes TRILL/SPB as this allows for possible modification of the switch topology.

- Acquire switches from a vendor whose product roadmap includes some form of early EVB/VEPA support during 2011 or early 2012.

In addition, if storage access (i.e., NAS, SAN) over Ethernet is an important consideration, IT organizations should make sure that all the switches have a solid

roadmap for supporting DCB. DCB will also benefit other applications including real-time video and voice.

IT organizations need to adopt a storage networking strategy. Three possibilities are:

- Continue to deploy Fibre Channel.
- Cap the use of Fibre Channel and implement FCoE. This allows IT organizations to both preserve the investment in Fibre Channel and to migrate to a unified fabric.
- Cap the use of Fibre Channel and deploy 10/40/100 GbE iSCSI in large part because this is a simpler approach to fabric unification.

If IT organizations choose to implement FCoE, they need to determine whether they want to connect their Fibre Channel SANs to the access tier or to the core tier of the network. The trade-off between these approaches is primarily between the benefits of sharing storage across the data center vs. the complexity and the impact of loading the uplinks and the core switches with a mixture of storage and traffic.

Wide Area Networking

Unlike the way things were during the entire twenty-year period that began in the mid to late 1980s, today there is not a fundamentally new generation of WAN technology in development. As such, IT organizations need to maximize their use of the existing WAN technologies and leverage whatever new products and services are developed. With this in mind, IT organizations should:

- Consider the use of VPLS. As is typically the case with WAN services, the viability of using this service vs. traditional services will hinge largely on the relative cost of the services. This will vary by service provider and by geography.
- Implement a dynamic virtual WAN/network virtualization. One of the many positive aspects of this approach to wide area networking is that it can be based on a number of different WAN services; e.g., MPLS as well as cable, DSL, T1/E1 and 4G access to the Internet. In addition, this approach can be deployed at a limited number of sites and expanded incrementally if desired.
- Explore cloud bridging solutions if hybrid cloud solutions are of interest.
- Consider local access to the Internet. The trade-off here is that centralized access to the Internet reduces the complexity of providing security, but this approach increases the amount of traffic on the WAN, and hence the cost of the WAN. This approach also adds to overall network delay.
- Implement WOC functionality both to save on the cost of WAN services and to improve application performance. Alternatives include both hardware and software-

based WOCs, whether they are provided by the IT organization itself or by a third party as part of a managed service. Another alternative is to acquire WOC functionality from a SaaS provider.

- Implement ADC functionality both to improve the performance of the data center servers as well as to improve overall application performance. Alternatives include both hardware and software-based ADCs, whether they are provided by the IT organization itself or by a third party as part of a managed service.

Management

Almost every aspect of cloud computing (e.g., server virtualization, public and private cloud computing solutions) create significant management challenges. To respond to these challenges, IT organizations should:

- Analyze solutions for data center management automation and integration initiatives from the perspective of the organizational domains (e.g., servers, storage, and network) as well as the required expertise and staff development that are required to fully exploit vendor-supported APIs and the associated scripting languages.

- Evaluate the viability of implementing a route analytics solution to obtain visibility, analysis, and diagnosis capabilities of the issues that occur at the routing layer in complex, meshed networks, such as those found in public and hybrid cloud computing solutions.

- Increase their focus on managing services vs. focusing on managing individual technology domains.

- Work with the team responsible for data center LANs to determine the best way to get visibility into the traffic that goes between VMs on a given server.

- Implement the ability to perform standard management functions (e.g., troubleshooting, baselining) on a per-VM basis.

- Analyze the offerings of cloud computing service providers to determine if they provide APIs that can be leveraged to better manage public and hybrid cloud computing solutions.

- Implement a dynamic, highly scalable and integrated DNS/DHCP/IPAM solution, which is also well integrated with the virtual server management system.

- Evaluate service orchestration solutions relative to their ability to automate many of the manual tasks that are involved in provisioning and controlling the capacity of dynamic virtualized services.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

Published by Webtorials Editorial/Analyst Division
www.Webtorias.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
taylor@webtorials.com

Professional Opinions Disclaimer

All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

Packet Design Solutions:

Packet Design's IP routing and traffic analysis solutions empower network management best practices in the world's largest and most critical enterprise, Service Provider and Government OSPF, IS-IS, BGP, EIGRP and RFC2547bis MPLS VPN networks, enabling network managers to maximize network assets, streamline network operations, and increase application and service up-time.

Route Explorer: Industry-Leading Route Analytics Solution

Optimize IP Networks with Route Explorer

- Gain visibility into the root cause of a signification percentage of application performance problems.
- Prevent costly misconfigurations
- Ensure network resiliency
- Increase IT's accuracy, confidence and responsiveness
- Speed troubleshooting of the hardest IP problems
- Empower routing operations best practices
- Complement change control processes with real-time validation of routing behavior
- Regain network visibility across outsourced MPLS VPN WANs

Deployed in the world's largest IP networks

400+ of the world's largest enterprises, service providers, government and military agencies and educational institutions use Packet Design's route analytics technology to optimize their IP networks.

Overview of Route Explorer

Route Explorer works by passively monitoring the routing protocol exchanges (e.g. OSPF, EIGRP, IS-IS, BGP, RFC2547bis MPLS VPNs) between routers on the network, then computing a real-time, network wide topology that can be visualized, analyzed and serve as the basis for actionable alerts and reports. This approach provides the most accurate, real-time view of how the network is directing traffic, even across MPLS VPNs. Unstable routes and other anomalies – undetectable by SNMP-based management tools because they are not device-specific problems – are immediately visible. As the network-wide topology is monitored and updated, Route Explorer records every routing event in a local data store. An animated historical playback feature lets the operator diagnose inconsistent and hard-to-detect problems by "rewinding" the network to a previous point in time. Histograms displaying past routing activity allow the network engineer to quickly go back to the time when a specific problem occurred, while letting them step through individual routing events to discover the root cause of the problem. Engineers can model failure scenarios and routing metric changes on the as-running network topology. Traps and alerts allow integration with existing network management solutions. Route Explorer appears to the network simply as another router, though it forwards no traffic and is neither a bottleneck or failure point. Since it works by monitoring the routing control plane, it does not poll any devices and adds no overhead to the network. A single appliance can support any size IP network, no matter how large or highly subdivided into separate areas.

Traffic Explorer: Network-Wide, Integrated Traffic and Route Analysis and Modeling Solution

Optimize IP Networks with Traffic Explorer

- Monitor critical traffic dynamics across all IP network links
- Operational planning and modeling based on real-time, network-wide routing and traffic intelligence
- IGP and BGP-aware peering and transit analysis
- MPLS VPN service network traffic analysis
- Network-wide and site to site traffic analysis for enterprise networks utilizing MPLS VPN WANs
- Visualize impact of routing failures/changes on traffic
- Departmental traffic usage and accounting
- Network-wide capacity planning
- Enhance change control processes with real-time validation of routing and traffic behavior

Traffic Explorer Architecture:

Traffic Explorer consists of three components:

- **Flow Recorders:** Collect Netflow information gathered from key traffic source points and summarize traffic flows based on routable network addresses received from Route Explorer
- **Flow Analyzer:** Aggregates summarized flow information from Flow Recorders, and calculates traffic distribution and link utilization across all routes and links on the network. Stores replayable traffic history
- **Modeling Engine:** Provides a full suite of monitoring, alerting, analysis, and modeling capabilities

Traffic Explorer Applications

Forensic Troubleshooting: Traffic Explorer improves application delivery by speeding troubleshooting with a complete routing and traffic forensic history.

Strengthened Change Management: Traffic Explorer greatly increases the accuracy of change management Processes by allowing engineers to model planned changes and see how the entire network's behavior will change, such as if there will be any congestion arising at any Class of Service.

Network-Wide Capacity Planning: Using its recorded, highly accurate history of actual routing and traffic changes over time, Traffic Explorer allows engineers to easily perform utilization trending on a variety of bases, such as per link, CoS, or VPN customer. Traffic Explorer ensures application performance and optimizes capital spending by increasing the accuracy of network planning.

Disaster Recovery Planning: Traffic Explorer can simulate link failure scenarios and analyze continuity of secondary routes and utilization of secondary and network-wide links.

Overview of Traffic Explorer

Traffic Explorer is the first solution to combine real-time, integrated routing and traffic monitoring and analysis, with "what-if" modeling capabilities. Unlike previous traffic analysis tools that only provide localized, link by link traffic visibility, Traffic Explorer's knowledge of IP routing enables visibility into network-wide routing and traffic behavior. Powerful "what-if" modeling capabilities empower network managers with new options for optimizing network service delivery. Traffic Explorer delivers the industry's only integrated analysis of network-wide routing and traffic dynamics. Standard reports and threshold-based alerts help engineers track significant routing and utilization changes in the network. An interactive topology map and deep, drill-down tabular views allow engineers to quickly perform root cause analysis of important network changes, including the routed path for any flow, network-wide traffic impact of any routing changes or failures, and the number of flows and hops affected. This information helps operators prioritize their response to those situations with the greatest impact on services. Traffic Explorer provides extensive "what-if" planning features to enhance ongoing network operations best practices. Traffic Explorer lets engineers model changes on the "as running" network, using the actual routed topology and traffic loads. Engineers can simulate a broad range of changes, such as adding or failing routers, interfaces and peerings; moving or changing prefixes; and adjusting IGP metrics, BGP policy configurations, link capacities or traffic loads. Simulating the affect of these changes on the actual network results in faster, more accurate network operations and optimal use of existing assets, leading to reduced capital and operational costs and enhance service delivery.

For more information, contact Packet Design at:

Web: http://www.packetdesign.com
Email: info@packetdesign.com
Phone: +1 408-490-1000

Managing End-user Experience, Application Performance Across The Cloud Infrastructure

Overview

The fundamental challenge of any IT organization today is aligning its technology with the business goals. In order to achieve alignment, IT organizations need to have visibility into how the performance and changes in the infrastructure impact application and business service delivery. With the emergence of Cloud Computing, this actionable insight is even more important to bridging the gap between business goals, customer experience and IT technology.

While Cloud Computing is a significant technology turn that impact how future business services will be delivered to the enterprise, there are a number of challenges to overcome, before broad adoption in the enterprise. One of the characteristics of Cloud Computing is the ability to provision services on demand, and the flexibility to increase capacity on demand. This means a very dynamic environment with changes taking place. While change management is not a new concept within IT, dealing with the nature and volume of change are new.

As stated in the Cloud Networking Report, the key challenges created for the network to support Cloud Computing include:

- Manual network re-configuration to support Virtual Machine (VM) migration
- Maintaining the performance and controlling the cost of the wide area network (WAN)
- Services supported by a virtual and dynamic infrastructure

Each of these challenges creates potential application performance issues impacting the end users and the business. While IT organizations are faced with managing these issues, they still have to deliver consistent application and business services.

Best practices for managing change includes:

- Establishing what is normal so one can easily tell whether there is a change
- Measuring and identifying the impact of infrastructure changes on application performance and end user experience

If there are indeed performance issues, the support teams needs to:

- Identify performance degradation incidents
- Determine the root cause of degradation
- Resolve the problem

Actionable Intelligence is Key

As the report had identified, the task of moving a VM is simple in a virtual server management system. The challenge is in making sure the VM's network configuration state is also transferred. A best practice in managing change is to ensure end users are minimally impacted includes establishing a baseline of expected normal performance of the applications. When changes to the VM or network configurations are made in error, the support team can be alerted to the deviation in performance before the phone rings at the helpdesk.

When performance degradation occurs, the support team needs to be able to assess and identify the user, locations and maybe even impact to the business. Cloud Services, which could be a single application or group of applications, are supported by a virtual and dynamic infrastructure that can expand and move to accommodate changing business capacity. It is even more important to establish performance baseline so that the impact of the changes to the underlying infrastructure can be properly managed to maintain service levels.

Since one of the characteristics of Cloud Computing is the centralization of resources, it will drive more traffic over the WAN and have an impact on performance and cost. Having visibility of the WAN traffic and application performance, in relationship to the remotes sites supported provides the intelligence to make informed decisions about whether additional bandwidth is required to maintain the prior levels of end user experience.

A Unified Performance Management System

Application and network performance management systems are not new to IT support teams. The current breed of products is designed for silo used. This means the network team focus on network performance, and the application support team focuses on application performance. Managing application performance and end user experience in a Cloud infrastructure requires the information to be combined and presented in the context of the Enterprise and not in terms of the technology.

A unified performance management system is a solution that brings together data sources that are useful in tackling performance problems. The architecture of the unified solution includes a common data model and the ability to correlate data from multiple sources. These two aspects are crucial to providing a view of the various domains (application, database, server and network) in context with each other during the time when the performance problem occurred.

Increasing the business value of IT

For many CIOs, one key concern is not only increasing the business value of IT, but also quantifying the positive impact. While many organizations view IT as a cost center with a focus on reducing the expenses, many leading companies view IT as a strategic asset. These organizations focus on how IT can improve overall business value to the organization.

A comprehensive strategy for managing the impact of Cloud infrastructure on application and business services delivery allows organizations to focus on the strategic asset and align technology with business goals. With a complete understanding of the impact of infrastructure performance and changes on the business and users, an enterprise can reduce the risk of downtime and degradation, reduce the cost of operations and troubleshooting, and optimize IT support staff.

What is Enterprise Service Intelligence (ESI)? Why is it Important?

The term Enterprise Service Intelligence is the Visual Networks Systems vision to help IT professionals and business stakeholders understand the true impact of the IT infrastructure on mission critical applications and business services. The implementation of ESI demonstrates the value of IT in the business context by delivering insight into individual user experience, application and network performance. Find out more, and how you can begin to put ESI into action in your environment, at

www.visualnetworksystems.com/ESI.

IT organizations have begun implementing business service dashboards and automating service desk workflow. A business service dashboard provides the line of business owners a clear view of the availability and performance of critical application and services that impact the bottom line. For IT, this dashboard increases the visibility of impact of infrastructure performance and changes on the business and users. This is especially important as the adoption of Cloud Computing expands.

The ideal underlying unified performance management system, supplying the intelligence to the business service dashboard and service desk, needs to be built on an application-aware architecture with the ability to correlate data gathered from a range of instrumentation options covering end user experience monitoring, application and network performance across the enterprise, including the Cloud infrastructure. Identifying and leveraging the right solution will minimize the challenges and limitations presented with the traditional, siloed approach to IT and ultimately help to align them to overall business goals.