The 2013 Application & Service Delivery Handbook

Part 3: Management and Security

By Dr. Jim Metzler, Ashton Metzler & Associates
Distinguished Research Fellow and Co-Founder
Webtorials Analyst Division

Platinum Sponsors:

cisco

ipanema Technologies

NETSCOUT

Gold Sponsors:

A10 Networks

agility made possible
ca technologies

radware

riverbed

Silver Peak

Produced by:

Webtorials

Management and Security

Executive Summary

The ***2013 Application and Service Delivery Handbook*** (The Handbook) will be published both in its entirety and in a serial fashion. This is the third of the serial publications. The first publication focused on describing a set of factors that complicate the task of ensuring acceptable application delivery. The second publication described the technologies and services that are available to improve the performance of applications and services. This is the third publication in the series and the primary goal of this publication is to describe the technologies and services that are available to improve the management and security of applications and services. The fourth and final publication in the series will include an executive summary as well as a copy of the complete document.

The first section of The ***2013 Application and Service Delivery Handbook*** described the surveys that were administered to the subscribers of Webtorials. Throughout this document, the IT professionals that responded to those surveys will be referred to as The Survey Respondents.

Management

Background

As will be discussed in this section of the handbook, in order to respond to the myriad challenges facing them, IT organizations need to adopt an approach to management that focuses on the services that IT provides. In this context, a *service* is comprised of the following four components:

- Either a multi-tier application or multiple applications
- Supporting protocols
- Enabling network services; e.g., DNS, DHCP
- The end-to-end network

Market Research

As was mentioned in the preceding sections of The Handbook, in early 2013 two surveys were given to the subscribers of Webtorials. One of the surveys focused on identifying the optimization and management tasks that are of most interest to IT organizations. With that goal in mind, The Survey Respondents were given a set of twenty optimization tasks and twenty management tasks and asked to indicate how important it was to their IT organization to get better at these tasks over the next year. The Survey Respondents were given the following five-point scale:

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

Some of the responses of The Survey Respondents were included in the preceding section of The Handbook and some others will be highlighted in this section of The Handbook. For completeness, **Table 1** shows how The Survey Respondents answered the question about the management tasks that are of most interest to their IT organization.

Table 1: The Importance of Getting Better at 20 Key Management Tasks

	Not at All	Slightly	Moderately	Very	Extremely
Rapidly identify the root cause of degraded application performance	2.4%	4.2%	14.5%	46.4%	32.5%
Identify the components of the IT infrastructure that support the company's critical business applications	3.1%	10.0%	18.8%	40.6%	27.5%

Table 1: The Importance of Getting Better at 20 Key Management Tasks					
	Not at All	Slightly	Moderately	Very	Extremely
Obtain performance indicator metrics and granular data that can be used to detect and eliminate impending problems	4.3%	6.7%	27.6%	43.6%	17.8%
Monitor the end user's experience and behavior	3.1%	12.3%	22.2%	45.1%	17.3%
Relate the performance of applications to the impact on the business	4.4%	14.6%	17.1%	46.8%	17.1%
Effectively manage SLAs for one or more business critical applications	6.3%	10.1%	25.3%	41.1%	17.1%
Manage the use of VoIP	7.4%	13.5%	23.9%	30.1%	25.2%
Perform traditional management tasks such as troubleshooting and performance management, on a per VM basis	3.2%	15.4%	25.0%	42.3%	14.1%
Monitor and manage the performance of applications delivered to mobile users	10.3%	13.3%	27.9%	32.1%	16.4%
Manage a business service, such as CRM, that is supported by multiple, inter-related applications	11.7%	13.0%	30.5%	31.2%	13.6%
Manage end-to-end in a private cloud computing environment	12.3%	13.5%	31.6%	30.3%	12.3%
Manage the traffic that goes between virtual machines on a single physical server	9.1%	17.5%	35.1%	27.9%	10.4%
Support the movement of VMs between servers in different data centers	12.2%	20.9%	25.7%	31.8%	9.5%

Table 1: The Importance of Getting Better at 20 Key Management Tasks

	Not at All	Slightly	Moderately	Very	Extremely
Effectively monitor and manage an application acquired from a SaaS provider such as Salesforce	17.8%	15.1%	26.3%	29.6%	11.2%
Manage end-to-end in a public cloud computing environment	21.1%	13.8%	28.3%	25.7%	11.2%
Manage end-to-end in a hybrid cloud computing environment	20.0%	16.0%	28.0%	25.3%	10.7%
Manage the use of telepresence	20.4%	21.7%	24.2%	26.8%	7.0%
Effectively monitor and manage computing services acquired from a IaaS provider such as Rackspace	22.3%	18.0%	24.5%	29.5%	5.8%
Manage the use of traditional video services	19.2%	30.8%	22.4%	19.9%	7.7%
Effectively monitor and manage storage services acquired from a IaaS provider such as Rackspace	27.3%	21.7%	21.7%	24.5%	4.9%

Some of the conclusions that can be drawn from the data in **Table 1** include:

Rapidly identifying the root cause of degraded application performance is the most important management task facing IT organizations; followed closely by related tasks such as identifying the components of the IT infrastructure that support the company's critical business applications.

Some traditional management tasks, such as managing the use of VoIP remain very important and some new tasks, such as managing the performance of applications delivered to mobile users have become very important.

Being able to perform traditional management tasks such as troubleshooting and performance management on a per VM basis is almost as important as being able to monitor the user's experience and behavior.

Managing the use of services acquired from an IaaS provider such as Rackspace is relatively unimportant.

Forces Driving Change

Previous sections of this handbook described the traditional and emerging service and application delivery challenges. This subsection will identify how some of those challenges are forcing a change in terms of how IT organizations manage services.

Server Virtualization

Until recently, IT management was based on the assumption the IT organizations performed tasks such as monitoring, baselining and troubleshooting on a server-by-server basis. Now, as highlighted by the data in Table 1, IT organizations understand that they must also perform management tasks on a virtual machine (VM)-by-VM basis. Another assumption that underpinned the traditional approach to IT management was that the data center environment was static. For example, it was commonly assumed that an application resided on a given server, or set of servers, for very long periods of time. However, part of the value proposition that is associated with server virtualization is that it is possible to migrate VMs between physical servers, both within the same data center and between disparate data centers.

IT organizations need to adopt an approach to management that is based on the assumption that the components of a service, and the location of those components, can and will change frequently.

Cloud Balancing

IT management has historically been based on the assumption that users of an application accessed that application in one of the enterprise's data centers and that the location of that data center changed very infrequently over time. The adoption of IaaS solutions in general, and the adoption of cloud balancing in particular demonstrates why IT organizations need to adopt an approach to IT management that is based on gathering management data across myriad data centers, including ones that are owned and operated by a third party. The adoption of cloud balancing is also another example of why IT organizations need to adopt an approach to management that is based on the assumption that the components of a service, and the location of those components, can and will change frequently.

Delay Sensitive Traffic

Voice and video are examples of applications that have high visibility and which are very sensitive to transmission impairments.   As previously mentioned, getting better at managing VoIP is one of the most important management tasks facing IT organizations.

As part of the traditional approach to IT management it is common practice to use network performance measurements such as delay, jitter and packet loss as a surrogate for the performance of applications and services. A more effective approach is to focus on aspects of the communications that are more closely aligned with ensuring acceptable application and service delivery. This includes looking at the application payload and measuring the quality of the voice and video communications. In the case of Unified Communications (UC), it also means monitoring the signaling between the components of the UC solution.

In addition to having a single set of tools and more of a focus on application payload, IT organizations need to implement management processes that understand the impact that each application is having on the other applications and that can:

- Analyze voice, video, UC and data applications in consort with the network
- Support multi-vendor environments
- Support multiple locations

Converged Infrastructure

One of the characteristics that is frequently associated with cloud computing is the integration of networking, servers and computing in the data center. While a converged data center infrastructure offers a number of benefits, it does create a number of management challenges. In particular, the converged infrastructure requires a management system and management processes that have the same level of integration and cross-domain convergence that the infrastructure has. For example, in order to support the requirement for the dynamic provisioning and re-allocation of resources to support a given IT service, the traditional manual processes for synchronizing the required server, network and storage resources will have to be replaced with integrated, automated processes. In order to enable this change, the provisioning and change management processes will need to be integrated and will need to feature the automatic configuration of network and storage resources when additional infrastructure services are deployed or when additional physical or virtual servers are brought on line or are moved. In a similar fashion, operations management needs to be consolidated and automated to keep service quality in line with user expectations.

While not a requirement, the cross-domain integrated management of a converged infrastructure will bring the greatest benefit in those instances in which a single administrator has the authority to initiate and complete cross-domain management tasks, such as provisioning and modifying infrastructure services. For example, the use of a single administrator can eliminate the considerable delays that typically occur in a traditional management environment where the originating administrator must request other administrators to synchronize the configuration of elements within their domains of responsibility. However, in many cases the evolution from the current approach of having separate administrators for each technology domain to an approach in which there is a single administrator will involve organizational challenges. As a result, many IT organizations will evolve to this new approach slowly over time.

Application Performance Management

Background

This section of The Handbook will outline an approach that IT organizations can utilize to better manage application and service delivery, where the term *service* was previously defined. However, in an effort to not add any more confusion to an already complex topic, instead of using a somewhat new phrase *application and service delivery management*, this section will use the more commonly used phrase *application performance management*.

Since any component of a complex service can cause service degradation or a service outage, in order to effectively perform application performance management IT organizations need a single unified view of all of the components that support a service. This includes the highly visible service components such as servers, storage, switches and routers, in both their traditional stand-alone format as well as in their emerging converged format; i.e., Cisco's UCS and VCE's Vblock platforms. It also includes the somewhat less visible network services such as DNS and DHCP, which are significant contributors to application degradation. Multiple organizational units within an IT organization have traditionally provided all of these service components. On an increasing basis, however, one or more network service providers and one or more cloud computing service providers will provide some or all of these service components. As a result, in order to achieve effective service delivery management, management data must be gathered from the enterprise, one or more Network Service Providers (NSPs) and one or more Cloud Computing Service Providers (CCSPs). In addition, in order to help relate the IT function with the business functions, IT organizations need to be able to understand the key performance indicators (KPIs) for critical business processes such as supply chain management and relate these business level KPIs to the performance of the IT services that support the business processes.

IT organizations must also be able to provide a common and consistent view of both the network and the applications that ride on the network to get to a service-oriented perspective. The level of granularity provided needs to vary based on the requirements of the person viewing the performance of the service or the network. For example, a business unit manager typically wants a view of a service than is different than the view wanted by the director of operations, and that view is often different than the view wanted by a network engineer.

As shown in **Table 1**, being able to monitor the end user's experience and behavior is a very important management task. One of the reasons for that importance is that in spite of all of the effort and resources that have gone into implementing IT management to date:

It is the end user, and not the IT organization who typically is the first to notice when the performance of an application begins to degrade.

Monitoring actual user transactions in production environments provides valuable insight into the end-user experience and provides the basis for an IT organization to be able to quickly identify, prioritize, triage and resolve problems that can affect business processes.

An effective approach to application performance management must address the following aspects of management:

- The adoption of a system of service level agreements (SLAs) at levels that ensure effective business processes and user satisfaction for at least a handful of key applications.

- Automatic discovery of all the elements in the IT infrastructure that support each service. This functionality provides the basis for an IT organization to being able to create two-way mappings between the services and the supporting infrastructure components. These mappings, combined with event correlation and visualization, can facilitate root cause analysis, significantly reducing mean-time-to-repair.

As was previously discussed, getting better at identifying the components of the IT infrastructure that support the company's critical business applications and services is one of the most important management tasks facing IT organizations.

If IT organizations can effectively identify which components of the infrastructure support a particular application or service, monitoring can much more easily identify when services are about to begin to degrade due to problems in the infrastructure. As part of this monitoring, predictive techniques such as heuristic-based trending of software issues and infrastructure key performance indicators can be employed to identify and alert management of problems before they impact end users – which the data in **Table 1** indicates is one of the most important management tasks. In addition, outages and other incidents that generate alerts can be prioritized based on their potential business impact. Prioritization can be based on a number of factors including the affected business process and its value to the enterprise, the identity and number of users affected and the severity of the issue.

As was also previously discussed, getting better at rapidly identifying the causes of application degradation is the most important management task facing IT organizations. Once the components of the infrastructure that support a given application or service has been identified, triage and root cause analysis can be applied at both the application and the infrastructure levels. When applied directly to applications, triage and root cause analysis can identify application issues such as the depletion of threads and pooled resources, memory leaks or internal failures within a Java server or .NET server. At the infrastructure level, root cause analysis can determine the subsystem within the component that is causing the problem.

As part of an effective approach to application performance management, the automated generation of performance dashboards and historical reports allow both IT and business managers to gain insight into SLA compliance and performance trends. The insight that can be gleaned from these dashboards and reports can be used to enhance the way that IT supports key business processes; help the IT organization to perform better capacity and budget planning; and identify where the adoption of new technologies can further improve the optimization, control and management of application and service performance. Ideally, the dashboard is a single pane of glass that can be customized to suit different management roles; e.g., the individual contributors in the Network Operations Center, senior IT management as well as senior business management.

Application Performance Management in the Private Enterprise Network[1]

Enterprise IT organizations can choose among several types of tools for monitoring and managing application performance over a private enterprise network. These include: application agents, monitoring of real and synthetic transactions, network flow and packet capture, analytics, and dashboard portals for the visualization of results.

At a high level, there are two basic classes of tools. The first class of tool monitors global parameters such as user response time or transaction completion time and provides alerts when thresholds are exceeded. These tools include agents on end user systems and monitoring appliances in the data center. The second class of tool supports triage by monitoring one or more of the components that make up the end-to-end path of the application. These tools include devices that capture application traffic at the flow and packet levels, agents on database, application, and web servers, as well as agents on various network elements.

The ultimate goal of application performance management is have a single screen that integrates the information from all of the tools in both categories. The idea being that a dashboard on the screen would indicate when user response time or transaction completion time begins to degrade. Then, within a few clicks, the administrator could determine which component of the infrastructure was causing the degradation and could also determine why that component of the infrastructure was causing degradation; e.g., high CPU utilization on a router.

Each type of individual tool has its strengths and weaknesses. For example, agents can supply the granular visibility that is required for complex troubleshooting but they represent an additional maintenance burden while also adding to the load on the servers and on the network. Monitoring appliances have more limited visibility, but they don't require modification of server configurations and don't add traffic to the network. Taking into consideration these trade-offs, IT organizations need to make tool decisions based on their goals for application performance management, their application and network environment as well as their existing infrastructure and network management vendors.

Independent of the approach that IT organizations take towards application performance management, a critical component of application performance management is end-to-end visibility. One of the challenges with discussing end-to-end visibility is that the IT industry uses the phrase end-to-end visibility in various ways. Given that one of this handbook's major themes is that IT organizations need to implement an application-delivery function that focuses directly on applications and not on the individual components of the IT infrastructure, this handbook will use the following definition of end-to-end visibility:

End-to-end visibility refers to the ability of the IT organization to examine every component of IT that impacts communications once users hit ENTER or click the mouse button until they receive a response back from the application.

End-to-end visibility is one of the cornerstones of assuring acceptable application performance. This functionality is important because it:

[1] This refers to managing the performance of applications that are delivered over WAN services such as Frame Relay, ATM and MPLS.

- Provides the information that allows IT organizations to notice application performance degradation before the end user does.

- Identifies the symptoms of the degradation and as a result enables the IT organization to reduce the amount of time it takes to identify and remove the causes of the degraded application performance.

- Facilitates making intelligent decisions and getting buy-in from other impacted groups. For example, end-to-end visibility provides the hard data that enables an IT organization to know that it needs to add bandwidth or redesign some of the components of the infrastructure because the volume of traffic associated with the company's sales order tracking application has increased dramatically. It also positions the IT organization to manage the recreational use of the network.

- Allows the IT organization to measure the performance of a critical application before, during and after a change is made. These changes could be infrastructure upgrades, configuration changes or the adoption of a cloud computing delivery model. As a result, the IT organization is in a position both to determine if the change has had a negative impact and to isolate the source of the problem so it can fix the problem quickly.

The value of providing end-to-end visibility is maximized if two criteria are met. One criterion is that all members of the IT organization use the same tool or set of tools. The second criterion is that the tool(s) are detailed and accurate enough to identify the sources of application degradation. One factor that complicates achieving this goal is that so many tools from so many types of vendors all claim to provide the necessary visibility. A second factor that complicates achieving this goal is the complexity and heterogeneity of the typical enterprise network. The typical enterprise network, for example, is comprised of switches and routers, access points, firewalls, ADCs, WOCs, intrusion detection and intrusion prevention appliances from a wide range of vendors. An end-to-end monitoring solution must profile traffic in a manner that reflects not only the physical network but also the logical flows of applications, and must be able to do this regardless of the vendors who supply the components or the physical topology of the network.

Application Performance Management in Public and Hybrid Clouds

There are a number of possible ways that an IT organization can adjust their application performance management strategies in order to accommodate accessing services hosted by a CCSP. These include:

- Extend the enterprise monitoring solutions into the public cloud using agents on virtual servers and by using virtual appliances. This option assumes that the CCSP offers the ability to install multiple virtual appliances (e.g., WOCs, and ADCs) and to configure the virtual switches to accommodate these devices.

- Focus on CCSPs that offer either cloud resource monitoring or application performance management as a service. Basic cloud monitoring can provide visibility into resource utilization, operational performance, and overall demand patterns. This includes providing metrics such as CPU utilization, disk reads and writes and network traffic. The value of cloud monitoring is increased where it is tied to other capabilities such as automated provisioning of instances to maintain high availability and the elastic scaling of capacity to satisfy demand spikes. A possible issue with this option is integrating the cloud monitoring and enterprise monitoring and application performance management solutions.

- Increase the focus on service delivery and transaction performance by supplementing existing application performance management solutions with capabilities that provide an outside-in service delivery view from the perspective of a client accessing enterprise applications or cloud applications over the Internet or mobile networks. Synthetic transactions against application resources located in public clouds are very useful when other forms of instrumentation cannot be deployed. One option for synthetic transaction monitoring of web applications is a third party performance monitoring service with end user agents distributed among numerous global ISPs and mobile networks.

Security

How IT Organizations are Implementing Security

The security landscape has changed dramatically in the last few years. In the very recent past, the typical security hacker worked alone, relied on un-sophisticated techniques such as dumpster diving, and was typically motivated by the desire to read about their hack in the trade press. In the current environment, sophisticated cyber criminals have access to malware networks and R&D labs and can use these resources to launch attacks whose goal is often to make money for the attacker. In addition, national governments and politically active hackers (hacktivists) are engaging in cyber warfare for a variety of politically motivated reasons.

The sophistication of computer attacks has increased dramatically in the last few years.

Security is both a first and a second-generation application and service delivery challenge and it will remain a significant challenge for the foreseeable future. Rapid changes in IT, such as those created by the adoption of cloud computing, social networking and the new generation of mobile devices, combined with the ongoing evolution of regulations pose a spate of new challenges for IT security systems and policies in much the same manner that they present challenges to the IT infrastructure.

IT security systems and policies have evolved and developed around the traditional application delivery architecture in which branch offices are connected to application servers in a central corporate data centers. In this architecture, the central corporate data center is a natural location to implement IT security systems and policies that provide layered defenses as well a single, cost efficient location for a variety of IT security functions. With the adoption of public cloud computing, applications and services are moving out of the central corporate data center and there is no longer a convenient single location for security policies and systems.

In addition, IT security systems and policies have traditionally distinguished between people who were using IT services for work versus those who were using it for personal use. The use of an employer provided laptop was subject to the employer's IT security policies and systems. In this environment, the use that employees made of personal laptops was generally outside of the corporate IT security policy. With the arrival of smartphones and tablet computers, the ownership, operating systems and security capabilities of the end user devices have changed radically. IT security policies and standards that were developed for PCs are no longer effective nor optimal with these devices. Most corporations have embraced the BYOD movement and end users are less willing to accept strict corporate security policies on devices they own. Additionally, strict separation of work and personal usage for security on an employee owned device is impractical.

The current and emerging environment creates a set of demanding security challenges.

The demands of governments, industry and customers have historically shaped IT security systems and policies. The wide diversity of organizations that create regulations and standards can lead to conflicts. For example, law enforcement requires access to network communications (Communications Assistance for Law Enforcement Act – CALEA) which may in turn force the creation of locations in the network that do not comply with the encryption requirements of other standards (e.g. Health Insurance Portability Accountability Act – HIPPA). In order to determine how IT organizations are responding to the traditional and emerging security challenges, The Survey Respondents were asked a series of questions. For example, to get a high level view of how IT organizations are providing security, The Survey Respondents were asked to indicate which of a number of network security systems their organization supports. The Survey Respondents were asked to check all of the alternatives that applied in their environment. Their responses are shown in **Table 2**.

Table 2: The Network Security Systems in Use

Network Security Systems	Percentage
Remote Access VPN	86.30%
Network Access Control	73.50%
Intrusion Detection/Protection Systems (IDS/IPS)	65.70%
Next Generation Firewalls (Firewall+IPS+Application Control)	56.90%
Secure Web Gateways	46.10%
Web Application and/or XML Firewalls	36.30%
Mobile Device Security/Protection	36.30%
Security Information Event Management	31.40%
Data Loss Prevention	24.50%
Password Vault Systems (either local or portal based)	12.70%
SAML or WS-Federation Federated Access Control	8.80%

One obvious conclusion that can be drawn from **Table 2** is that IT organizations use a wide variety of network security systems. A slightly less obvious conclusion is that on average, IT organizations use 4.8 of the network security systems listed in the preceding table.

The Survey Respondents were asked to indicate the approach that best describes how their company uses data classification to create a comprehensive IT security environment. Their responses are shown in **Table 3**.

Table 3: Approach to Comprehensive IT Security

Approach	Percentage
We have a data classification policy and it is used to determine application access/authentication, network and end user device security requirements.	42.90%
We do not have a data classification policy.	33.00%
We have a data classification policy and it is used to determine application security requirements.	13.20%
We have a data classification policy, but it is not used nor enforced.	11.00%

The data in **Table 3** represents a classic good news/bad news situation. The good news is that the majority of IT organizations have a data classification policy that they use to determine requirements. The bad news is that 44% of IT organizations either don't have a data classification policy or they have one that isn't used or enforced.

In order to understand how IT organizations are responding to the BYOD movement, The Survey Respondents were asked, "If your organization does allow employee owned devices to connect to your network, please indicate which of the following alternatives are used to register employee owned devices and load authentication (e.g. certificate/private key) data onto those devices before they are allowed to connect to your company's network." The Survey Respondents were asked to check all of the alternatives that applied in their environment. Their responses are shown in **Table 4**.

Table 4: Alternatives to Support Employee Owned Devices

Alternative	Percentage
Employees must install a VPN client on their devices for network access	53.90%
IT Administrator and/or Service Desk must register employee owned device for network access	47.40%
Employees can self-register their devices for network access	28.90%
Employees must generate and/or load X.509 certificates & private keys network access	13.20%
Employees must install a token authentication app on their devices for network access	10.50%

The data in **Table 4** indicates that while using a VPN is the most common technique that a wide range of techniques are used. VPN's popularity comes in part from the fact that remote access VPN solutions implemented on the new generation of mobile devices have various capabilities to enforce security policies when connecting to the corporate network. Popular security checks include ensuring that a screen password is present, that anti-virus software is present and is up to date, that there is not rogue software on the device and that the operating system has not been modified.

Two different approaches have emerged to protect against lost devices. For the traditional PC, full disk encryption is typically used to protect data if the PC is lost or stolen. However, on the new generation of mobile devices, remote erase solutions are typically used to protect data. In order to understand how IT organizations have implemented full disk encryption, The Survey Respondents were asked to indicate which alternatives their organization implements relative to using full disk encryption on laptops and desktop PCs. Their responses are shown in **Table 5**.

Table 5: Techniques for Implementing Full Disk Encryption

Alternative	Percentage
We do not use full disk encryption on PCs.	52.5%
We use software based disk encryption on PCs.	49.5%
We use hardware based self-encrypting rotating drives on PCs.	6.1%
We use hardware based self-encrypting Solid State Drives on PCs.	6.1%

The data in **Table 5** indicates that just over half of all IT organizations don't use full disk encryption on PCs. The data also indicates that those IT organizations that do use full disk encryption do so by using a software solution and that a small percentage of IT organizations use multiple techniques.

The Survey Respondents were asked to indicate the approach that best describes their company's approach to Identity and Access Management (IAM). Their responses are shown in **Table 6**.

Table 6: How IAM is Implemented

Approach	Percentage
We do not have a formal IAM program.	36.6%
We have an IAM program, but it only partially manages identities, entitlements and policies/rules for internal users.	25.8%
We have an IAM program and it manages identities, entitlements and policies/rules for all internal users.	20.4%
We have an IAM program and it manages identities, entitlements and policies/rules for end users for internal, supplier, business partner and customers.	17.2%

The data in **Table 6** indicates that only a minority of IT organizations has a IAM program that has broad applicability.

The Survey Respondents were asked to indicate how their company approaches the governance of network and application security. Their responses are shown in **Table 7**.

Table 7: Governance Models in Use

Approach	Percentage
Network Security and Application Security are funded, architected, designed and operated together.	46.9%
Network Security and Application Security are funded, architected, designed and operated separately.	30.2%
Network Security and Application Security are funded jointly, but architected, designed and operated separately.	22.9%

The data in **Table 7** indicates that in the majority of instances, network security and application security are architected, designed and operated separately.

Cloud-Based Security

The Survey Respondents were asked how likely it was over the next year that their company would acquire a traditional IT service from an IaaS provider. Their responses are shown in **Table 8**.

Table 8: Interest in Obtaining IT Services as a Cloud-Based Service

	Will Not Happen	Might Happen	50/50 Chance	Will Likely Happen	Will Happen
VoIP	32.6%	18.6%	15.3%	13.5%	20.0%
Unified Communications	30.2%	22.8%	20.5%	14.9%	11.6%
Security	42.6%	17.1%	14.4%	11.6%	14.4%
Network and Application Optimization	32.1%	28.8%	16.0%	14.6%	8.5%
Network Management	41.4%	22.3%	13.5%	13.5%	9.3%
Application Performance Management	37.9%	26.5%	15.6%	11.4%	8.5%
Virtual Desktops	38.8%	28.0%	15.9%	12.1%	5.1%

As shown in **Table 8**, the interest shown by The Survey Respondents in obtaining security as a Cloud-based service is bimodal. When looking just at the percentage of The Survey Respondents that indicated that it either will happen or will likely happen, security is one of the most likely services that IT organizations will acquire from a CCSP. However, a higher percentage (42.6%) of The Survey Respondents indicated that they will not acquire security from a CCSP than made that indication for any other form of IT service listed in the survey.

One way that a Cloud-based Security Service (CBSS) could provide value is if it provides protection against the growing number of malware attacks.   To effectively protect against malware attacks, a CBSS should be able to identify suspicious content or sites that are either suspicious or are known to distribute malware. In order to be effective, a CBSS that provides Web content filtering or malware protection needs a source of intellectual capital that identifies known and suspected vulnerabilities. This source needs to be both dynamic and as extensive as possible.

One part of the value proposition of a CBSS that provides security functionality is the same as the value proposition of any cloud based service. For example, a security focused CBSS reduces the investment in security that an organization would have to make. In addition, a security focused CBSS reduces the amount of time it takes to deploy new functionality. The speed at which changes can be made to a CBSS adds value in a variety of situations, including providing better protection against zero-day attacks[2]. Another part of the value proposition of a security focused CBSS is that unlike a traditional security solution that relies on the implementation of a hardware based proxy, a CBSS can also protect mobile workers. The

[2] http://en.wikipedia.org/wiki/Zero-day_attack

CBSS does this by leveraging functionality that it provides at its POPs as well as functionality in a software agent that is deployed on each mobile device.

In many instances, the best security solution is a hybrid solution that combines traditional on-premise functionality with one or more Cloud-based solutions. For example, in many cases IT organizations already have functionality such as web filtering or malware protection deployed in CPE at some of their sites. In this case, the IT organization may choose to implement a CBSS just to protect the sites that don't have security functionality already implemented and/or to protect the organization's mobile workers. Alternatively, an organization may choose to implement security functionality in CPE at all of their sites and to also utilize a CBSS as part of a defense in depth strategy.

Other situations in which a CBSS can serve to either be the only source of security functionality, or to compliment CPE based implementations include cloud-based firewall and cloud-based IPS services.   Such a service should support equipment from the leading vendors. Given the previously mentioned importance of hybrid solutions, the service should allow for flexibility in terms of whether the security functionality is provided in the cloud or from CPE as well as for flexibility in terms of who manages the functionality – a CCSP or the enterprise IT organization.

In addition to the specific security functionality provided by the CBSS, the CBSS should also:

- Provide predictive analytics whereby the CBSS can diagnose the vast majority of potential enterprise network and security issues before they can impact network health.

- Incorporate expertise, tools, and processes to ensure that the service that is provided can meet auditing standards such as SAS-70 as well as industry standards such as ITIL.

- Integrate audit and compliance tools that provide the necessary event-correlation capabilities and reporting to ensure that the service meets compliance requirements such as Sarbanes-Oxley, HIPAA, GLB and PCI.

- Provide the real-time notification of security events.

Web Application Firewall Services

The section of this report entitled *Network and Application Optimization*, discussed how a Cloud-based service, such as the one shown in **Figure 1**, can be used to optimize the performance of the Internet. As will be discussed in this sub-section of the handbook, that same type of service can also provide security functionality.

Figure 1: Internet Based Security Functionality

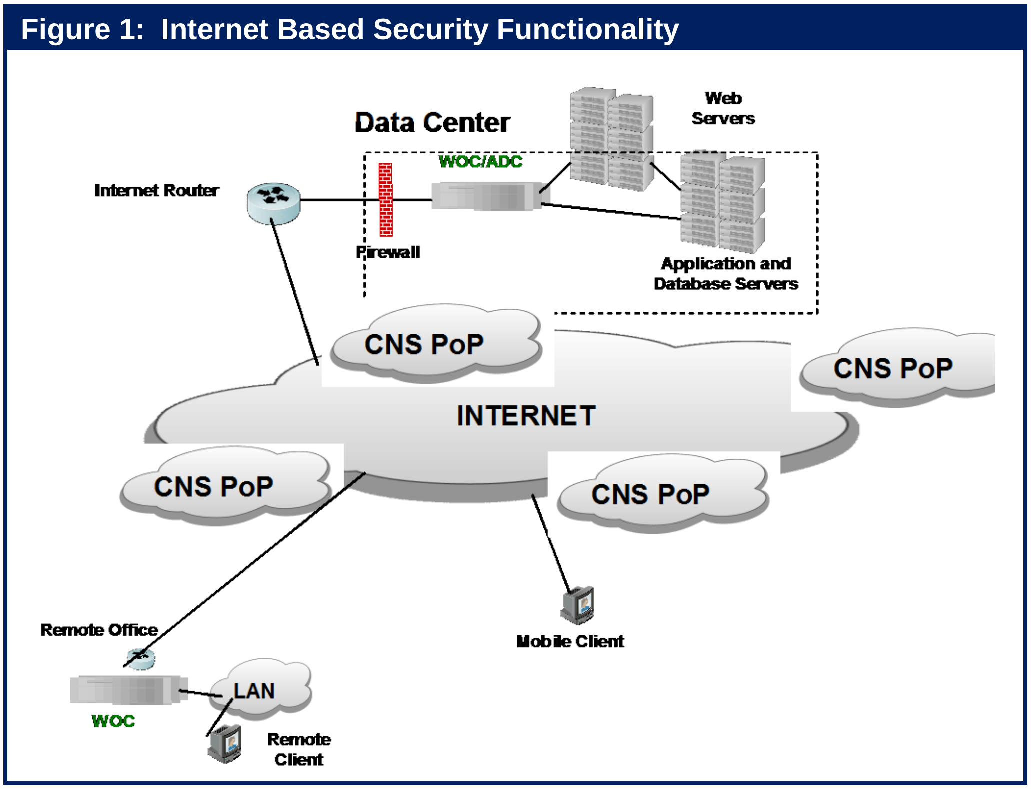

Role of a Traditional Firewall: Protect the Perimeter

Roughly twenty years ago IT organizations began to implement the first generation of network firewalls, which were referred to as packet filters. These devices were placed at the perimeter of the organization with the hope that they would prevent malicious activities from causing harm to the organization.

Today most network firewalls are based on stateful inspection. A stateful firewall holds in memory attributes of each connection. These attributes include such details as the IP addresses and ports involved in the connection and the sequence numbers of the packets traversing the connection. One of the weaknesses associated with network firewalls is that they are typically configured to open up ports 80 and 443 in order to allow passage of all HTTP and SSL traffic. Given that ports 80 and 443 are generally configured to be open, this form of perimeter defense is porous at best.

Whereas network firewalls are focused on parameters such as IP address and port numbers, a more recent class of firewall, referred to as a Web application firewall, analyzes messages at layer 7 of the OSI model. Web application firewalls are typically deployed as a hardware appliance and they sit behind the network firewall and in front of the Web servers. They look for violations in the organization's established security policy. For example, the firewall may look for abnormal behavior, or signs of a known attack. It may also be configured to block specified content, such as certain websites or attempts to exploit known security vulnerabilities. Because of their ability to perform deep packet inspection at layer 7 of the OSI model, a Web application firewall provides a level of security that cannot be provided by a network firewall.

Defense in Depth: The Role of a Web Application Firewall Service

As is well known, there are fundamental flaws with an approach to security that focuses only on the perimeter of the organization. To overcome these flaws, most IT organizations have moved to an approach to security that is typically referred to as *defense in depth*. The concept of defense in depth is not new. What is new in the current environment is the use of a CBSS to provide Web application firewall functionality that is distributed throughout the Internet. This means that Web application functionality is close to the source of security attacks and hence can prevent many security attacks from reaching the organization.

In the current environment, high-end DDoS attacks can generate 300 Gbps of traffic or more[3]. Attacks of this magnitude cannot be prevented by onsite solutions. They can, however, be prevented by utilizing a CBSS that includes security functionality analogous to what is provided by a Web application firewall and that can identify and mitigate the DDoS-related traffic close to attack traffic origin.

There is a wide range of ways that a DDoS attack can cause harm to an organization in a number of ways, including the:

- Consumption of computational resources, such as bandwidth, disk space, or processor time.
- Disruption of configuration information, such as routing information.
- Disruption of state information, such as the unsolicited resetting of TCP sessions.
- Disruption of physical network components.
- Obstructing the communication media between the intended users and the victim so that they can no longer communicate adequately.

Because there are a variety of possible DDoS attacks, IT organizations need to implement a variety of defense in depth techniques. This includes:

- **Minimizing the points of vulnerability**
  If an organization has most or all of its important assets in a small number of locations, this makes the organization more vulnerable to successfully being attacked as the attacker has fewer sites on which to concentrate their attack.
- **Protecting DNS**
  Many IT organizations implement just two or three DNS servers. As such, DNS is an example of what was discussed in the preceding bullet – how IT organization are vulnerable because their key assets are located in a small number of locations.
- **Implementing robust, multi-tiered failover**
  Many IT organizations have implemented disaster recovery plans that call for there to be a stand-by data center that can support at least some of the organization's key

[3] http://nakedsecurity.sophos.com/2013/03/28/massive-ddos-attack-against-anti-spam-provider-impacts-millions-of-internet-users/

applications if the primary data center fails. Distributing this functionality around a global network increases overall availability in general, and dramatically reduces the chance of an outage due to a DDoS attack in particular.

In order to be effective, a CBSS that provides Web application firewall functionality needs to be deployed as broadly as possible, preferably in tens of thousands of locations. When responding to an attack, the service must also be able to:

- Block or redirect requests based on characteristics such as the originating geographic location and whether or not the originating IP addresses are on either a whitelist or a blacklist.

- Direct traffic away from specific servers or regions under attack.

- Issue slow responses to the machines conducting the attack. The goal of this technique, known as tarpits[4], is to shut down the attacking machines while minimizing the impact on legitimate users.

- Direct the attack traffic back to the requesting machine at the DNS or HTTP level.

A CBSS that provides Web application firewall functionality is complimentary to a premise-based Web application firewall. That follows because while the Cloud-based Web application firewall service can perform many security functions that cannot be performed by an on premise Web application firewall, there are some security functions that are best performed by an on premise Web application firewall. An example of that is protecting an organization against information leakage by having an onsite Web application firewall perform deep packet inspection to detect if sensitive data such as a social security number or a credit card number is leaving the site. If sensitive data is leaving the site, the onsite Web application firewall, in conjunction with other security devices, can determine if that is authorized and if it is not, it can prevent the data from leaving the site.

[4] Wikipedia Tarpit(networking)

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.

Jim Metzler has a broad background in the IT industry. This includes being a software engineer, an engineering manager for high-speed data services for a major network service provider, a product manager for network hardware, a network manager at two Fortune 500 companies, and the principal of a consulting organization. In addition, he has created software tools for designing customer networks for a major network service provider and directed and performed market research at a major industry analyst firm. Jim's current interests include cloud networking and application delivery.

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.Webtorials.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.

For editorial and sponsorship information, contact Jim Metzler or Steven Taylor. The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.

Customer Driven Innovation

Cost Effective Cloud Computing

Reduce Your
Application Delivery Costs

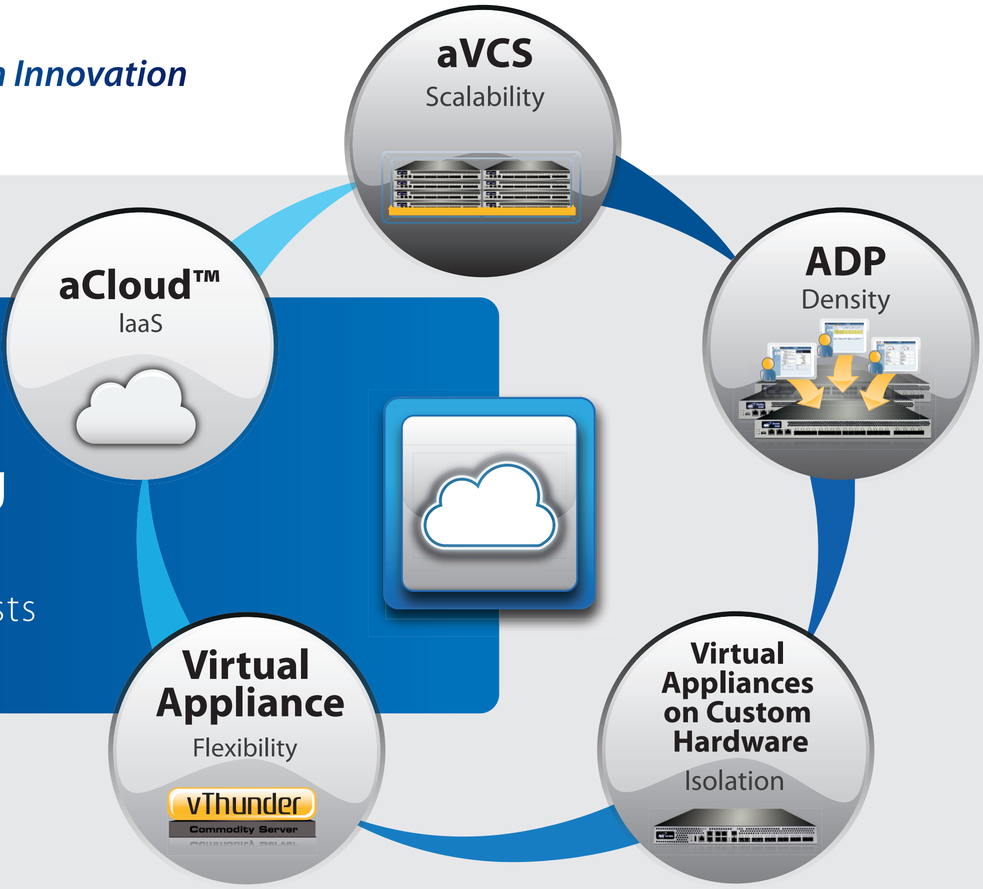

A10 Thunder™ and AX Series Products & Solutions

Based on A10's award-winning Application Delivery Controllers (ADCs) and Advanced Core Operating System (ACOS™) architecture, enterprises and service providers will have the flexibility to choose the following scale-as-you-grow virtualization options.

Virtual Appliances

- Virtual machine (VM) on commodity hardware
- Rapidly scale with commodity hardware
- Reduce hardware costs and upload to compatible cloud providers

Virtual Chassis System

- Cluster multiple ADCs to operate as a unified single device
- Scale while maintaining single IP management
- Reduce costs and simplify management while adding devices as you grow

Virtual Appliances on Custom Hardware

- Hardware appliances with embedded hypervisor running virtual appliances
- Flexibility with hardware performance and reliability

Application Delivery Partitions

- Partition the ADC platform resources for individual applications
- Enable quality multi-tenancy with granular resource allocation
- Reduce the number of appliances to host multiple applications

Mission-critical app? You need CA Application Performance Management

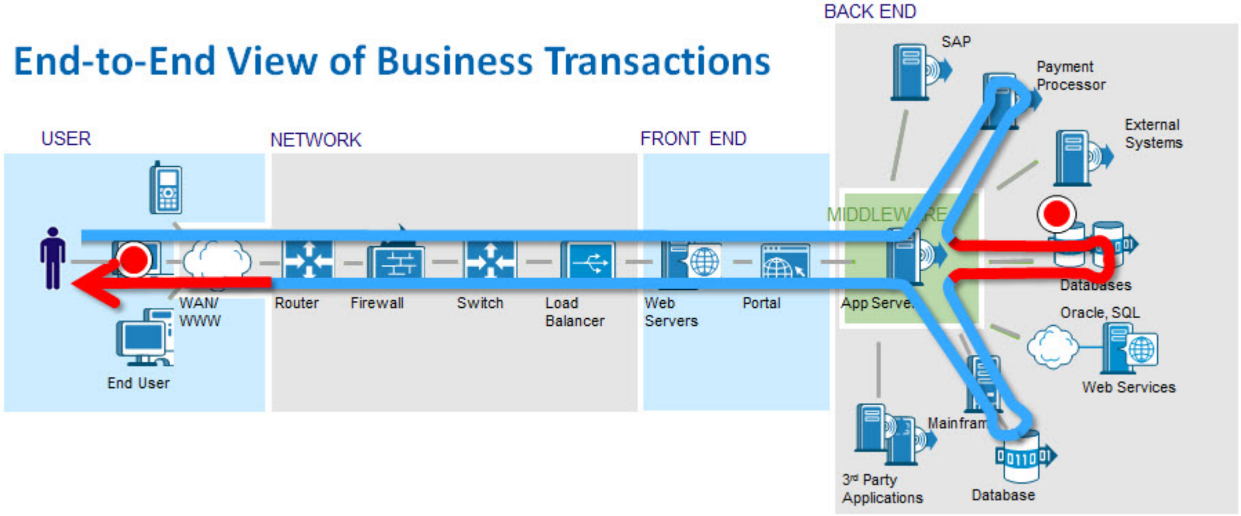

What does an hour of an application outage or downtime cost your business? For internal systems, it's lost productivity. For external systems it's lost customers and lost revenue. Let's face it: neither is good for the bottom line. Applications are the lifeblood of virtually all organizations today, meaning they require a robust application performance management solution that helps ensure end users get the experience they expect and business services are reliably delivered. Today's complex, business-critical applications require CA Application Performance Management (CA APM) to ensure business success.

CA APM delivers 360-degree visibility into and analysis of all user transactions across the hybrid-cloud infrastructure – physical, virtual, cloud and mainframe – to understand the health, availability, business impact and end-user experience of critical enterprise, mobile and cloud applications. Advanced Application Behavior Analytics add deeper visibility into the wealth of metric performance collected by CA APM, giving IT operators another set of eyes to look for potential trouble spots. With CA APM, organizations can proactively identify, diagnose and resolve problems throughout the application lifecycle to put organizations firmly in control of the end-user experience and optimize the performance of critical, revenue-generating services.

CA APM is uniquely designed for today's large, complex and heterogeneous production and pre-production environments. CA APM deploys rapidly with little overhead, scales to manage billions of transactions and is easy for both business and IT executives to use, accelerating time to value while increasing return on investment. CA APM is trusted by more than 2,500 large enterprises and service providers to manage the performance and availability of their critical and revenue-generating services. *Learn more at ca.com/apm*.

agility
made possible™

ca
technologies

Simplify and Accelerate Virtual Application Deployments with Cisco Cloud Network Services

Cisco and a Multi-vendor Ecosystem Provide Cloud-Ready Application Services

ROLE OF THE NETWORK FOR THE CLOUD

Access to Critical Data, Services, Resources and People

- Core fabric connects resources within the data center and across data centers to each other.
- Pervasive connectivity links users and devices to resources and each other.
- The network provides identity- and context-based access to data, services, resources, and people.

Granular Control of Risk, Performance, and Cost

- Manage and enforce policies to help ensure security, control, reliability, and compliance.
- Manage and enforce service-level agreements (SLAs) and consistent quality of Service (QoS) within and between clouds, enabling hybrid models and workload portability.
- Meter resources and use to provide transparency for cost and performance.

Robustness and Resilience

- Supports self-healing, automatic redirection of workload and transparent rollover.
- Provide scalability, enabling on-demand, elastic computing power through dynamic configuration.

Innovation in Cloud-Specific Services

- Context-aware services understand the identity, location, proximity, presence, and device.
- Resource-aware services discover, allocate, and pre-position services and resources.
- Comprehensive insight accesses and reports on all data that flows in the cloud.

Overview

Cisco has announced the evolution of its network services strategy for virtual and cloud networks, Cisco® Cloud Network Services, a complete portfolio of application networking and security services built on top of the Nexus® 1000V virtual networking portfolio and part of the Cisco Unified Data Center architecture. Cloud Network Services simplifies and accelerates cloud network deployments without compromising the critical security and application delivery services that critical data center applications require.

Introducing Cisco Cloud Network Services

Advances in cloud computing, data center consolidation, mobility, and big data are imposing new demands on the network, along with demands for greater network simplification and automation.

As virtual networking and programmable overlay networks evolve to meet these challenges, a similar evolution needs to take place in Layer 4 through 7 application networking services to support widespread virtualization, application mobility, cloud architectures and network automation.

Cisco's solution to this challenge is Cisco Cloud Network Services, a portfolio of integrated, application-aware network services offerings designed for virtual and cloud environments. The Cloud Network Services framework eliminates the obstacles of physical service appliances to accommodate the requirements of virtual applications and cloud deployments, such as:

- Limited scalability of physical services in fixed locations

CISCO VIRTUAL NETWORK PORTFOLIO

Routing and Switching

- Cisco Nexus 1000V virtual switch
- Cisco Cloud Services Router (CSR) 1000V

Security and VPN

- Cisco Virtual Security Gateway for Nexus 1000V (included in Nexus 1000V Advanced Edition)
- Cisco Adaptive Security Appliance (ASA) 1000V Cloud Firewall
- Imperva SecureSphere Web Application Firewall

WAN Optimization

- Cisco Virtual Wide Area Application Services (vWAAS)

Network Analysis and Monitoring

- Cisco Prime Virtual Network Analysis Module (NAM)

Application Delivery Controllers

- Citrix NetScaler VPX virtual application delivery controller

Virtual Services Deployment Platform

- Cisco Nexus 1100 Series Cloud Services Platform

Cloud Orchestration and Management

- Cisco Intelligent Automation for Cloud (IAC)
- Cisco Prime Network Controller
- OpenStack

To learn more about Cisco's complete virtual networking portfolio, see http://cisco.com/go/1000v

- Inconsistent application performance based on workload location relative to services
- Difficulty in inserting security and network services into virtual networks
- Lack of control over services and policies for applications deployed at cloud service providers

The Cisco Cloud Network Services portfolio includes the Cisco Adaptive Security Appliance (ASA) 1000V Cloud Firewall, Cisco Virtual Security Gateway (VSG) virtual firewall, Cisco Virtual Wide Area Application Services (vWAAS) WAN optimization solution, and Cisco Prime™ Virtual Network Analysis Module (vNAM). This new architecture provides a complete services portfolio while delivering scale-out architecture, elastic instantiation, and multi-tenant operation, all with a common approach to service provisioning and management.

Cloud Network Services also includes best-in-class third-party virtual service offerings that integrate transparently into the framework. It now includes the Citrix NetScaler VPX virtual application delivery controller (ADC), and the Imperva SecureSphere Web Application Firewall (WAF).

Cisco Cloud Network Services form the virtual network services strategy of the larger Cisco Unified Data Center framework, which brings together a seamless architecture of virtualization and cloud-ready compute servers (Unified Compute Servers), network fabric (Unified Fabric) and automation platform (Unified Management).

Cloud Network Services, based on the Nexus 1000V virtual switch, are designed to run across major hypervisors, including VMware vSphere, Microsoft Hyper-V, and Linux Kernel-based Virtual Machine (KVM). It is also designed to support multiple cloud orchestration and virtualization management platforms, such as VMware vCenter and Microsoft Systems Center Virtual Machine Manager (SCVMM), giving customers greater flexibility.

Cisco Cloud Services Platform

With the evolution to Cisco Cloud Network Services as the Layer 4 through 7 framework for virtual and cloud networks, organizations are increasingly looking for a flexible platform on which to deploy virtual service nodes rather than use existing application servers. The Nexus Cloud Services Platform was created to address this need.

The Cisco Nexus 1100 Cloud Services Platform is a group of Cisco Unified Computing System™ (Cisco UCS) appliances dedicated to running Cloud Network Service nodes. In addition to the virtual services listed earlier, the Nexus 1100 series runs the management platforms for the virtual network, the virtual security module (VSM), and the Cisco Data Center Network Manager (DCNM) application. The Cloud Services Platform can be dynamically configured to allocate its virtual CPUs to each service as needed based on current application and performance requirements. Current models of the Nexus 1100 Cloud Services Platform include the Nexus 1110-S and 1110-X.

vPath: Enabling Services in Virtual and Cloud Networks

vPath is a component of the Cisco Nexus 1000V virtual switch that directs traffic to appropriate virtual service nodes, such as firewalls and ADCs, in the correct order for each application, independent of the topology of the network or the location of the network services. This feature allows greater application mobility and more reliable service delivery (Figure 1).

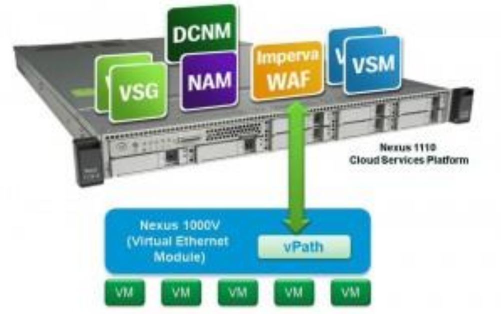
DCNM
VSG
NAM
Imperva WAF
VSM
Nexus 1110 Cloud Services Platform
Nexus 1000V (Virtual Ethernet Module)
vPath
VM
VM
VM
VM
VM

Figure 1 – vPath Connects Virtual Applications to Services Running on the Cisco Cloud Services Platform

Nexus 1000V InterCloud: Enable Hybrid Cloud Connectivity and Cloud-Based Virtual Services

As virtual networks extend from the data center to cloud service providers, organizations are concerned about the consistency of security and application delivery policies and about how these policies are enforced in the cloud. Applications that migrate from the data center to cloud providers can expect different behavior, and organizations may struggle to address compliance issues.

Cisco Nexus 1000V InterCloud complements Cisco Cloud Network Services, allowing seamless hybrid cloud connectivity between data centers and cloud providers and creating one extended network for application and Cloud Network Service deployments.

By deploying Cloud Network Services in all cloud locations, public and private, organizations help ensure consistent policy enforcement, quality of service (QoS) and compliance independent of the location of the virtual applications. Because Cloud Network Services are virtual machines themselves, they are easily deployed within public cloud providers regardless of the infrastructure they provide, and they provide the service consistency required for mission-critical applications.

For More Information

Learn more about Cisco virtual networking portfolio: http://cisco.com/go/1000v

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

About Ipanema Technologies

Application Performance Guarantee

Go Beyond WAN Optimization

About Ipanema

- Selected by worldwide enterprises across all industry sectors.
- One of the largest customer bases (over 150,000 managed sites).
- Visionary in Gartner's WOC Magic Quadrant 2013.
- Leader for Application-Aware Network services (BT, Colt, Vodafone, KDDI, KPN, OBS, Swisscom, Telecom Italia, Telefónica, Easynet…).

79%*

of organizations suffers application performance problems while increasing their IT budget.

*Ipanema Killer Apps survey 2012

Losing 5 minutes

per day for poor application performance means 1% of productivity drop which can turn down profitability by 10%.

"Thanks to Ipanema, our network is totally aligned with our business requirements. With the flexible application-based managed service delivered by e-Qual, we can guarantee the performance of our business critical applications including our ERP and MS Lync, anytime anywhere while reducing our IT costs".

Philippe Faure, Chief Information Officer, Gemalto

Go Beyond WAN Optimization to guarantee your applications performance

Ipanema provides enterprises with a direct connection between application performance and their business requirements.

With Ipanema Technologies, enterprises automatically understand which applications use the network and deliver guaranteed performance to each user. Enterprises can support their strategic IT transformations (like cloud computing and Unified Communications) and control Internet growth while reducing their IT expenses.

Using Ipanema, **enterprises**:

- Guarantee their business application performance;
- Protect unified communications;
- Enable hybrid cloud applications;
- Deploy hybrid networks;
- Control Internet, social media and video traffic;
- Save on IT costs.

ERP, CRM, Financial Reporting
UNIFIED COMMUNICATIONS
CLOUD APPLICATIONS, BIG DATA
INTERNET & HYBRID NETWORKS
INTERNET
MPLS
ETHERNET
VIDEO TRAFFIC & SOCIAL MEDIA
Bring Your Own Device
TRANSPARENT OPERATIONS & IT GOVERNANCE
IT Spending & Efficiency

Ipanema: the only solution that integrates all the features to guarantee application performance

Ipanema's self-learning and self-optimizing Autonomic Networking System™ (ANS) tightly integrates all the features to guarantee the best application performance:

- Application Visibility,
- Application Control,
- WAN Optimization,
- Dynamic WAN Selection
- And Network Rightsizing.

Do You Have Best-in-Class Application Delivery?

Today's data center challenges extend beyond the traditional needs for application availability, performance and security – challenges well-served previously with classic load balancers / application delivery controllers (ADCs). Nowadays, the adoption of data center virtualization, synchronization with dynamic data center changes, true "awareness" of deployed business applications, and the need for end-to-end visibility– all require a new class of advanced (yet cost-effective) ADCs.

Radware **Alteon® 5224** is an advanced ADC specifically targeted to address all of these challenges. Offering the very latest in next generation application delivery technology with ease of operations, it's simply the best-in-class application delivery choice. Here are four reasons why, we know you'll appreciate:

Reason 1: ADC Virtualization and Consolidation

ADC-VX™, part of Radware's Virtual Application Delivery Infrastructure (VADI)™ strategy, is the industry's first ADC virtualization hypervisor, allowing for the most cost-effective ADC consolidation capabilities. ADC-VX is built on a unique architecture that virtualizes the resources of Radware's ADC including CPU, memory, network and acceleration resources. This specialized hypervisor runs virtual ADC instances (vADC) where each delivers full ADC functionality. Each virtual ADC instance contains a complete and separated environment of resources, OS, configurations and management.

In turn, this allows allocation of a separate, fully-isolated vADC instance for each application. Companies can then maximize application availability and meet application SLA requirement with a resource reservation mechanism. Moreover, this deployment model simplifies operations, reduces the ADC infrastructure footprint, and increases business agility with faster roll out of new vADCs and applications. With vADC per application, application lifecycle management is streamlined and its associated cost is significantly reduced compared to traditional ADC deployment models.

Reason 2: Result-Driven Application Acceleration

Radware's FastView™ result-driven acceleration technology adds Web Performance Optimization (WPO) capabilities on top of standard ADC application acceleration features (e.g., caching, compression, SSL acceleration, etc.), to deliver the fastest Web application response time and ensure best application SLA while offloading server processing. This results in increased revenues, higher conversion rates, higher customer loyalty as well as improved employee productivity when using enterprise web applications. It applies to all browsers, all end-user device types and all users, located anywhere. Radware's leading WPO capabilities include:

- Reducing the # of server requests per page
- Accelerate entire web transaction
- Custom optimization templates for each browser
- Static and dynamic, browser-side caching
- Dedicated, mobile caching based on HTML 5 local storage
- Content minification

Reason 3: End-to-End Application QoE & Performance Visibility

Ensuring applications deliver the best quality of experience requires IT administrators to gain maximum visibility on all application delivery chain components, throughout the life cycle of the application. Radware's multilayer approach for monitoring the application delivery infrastructure, coupled with its integrated application

performance monitoring (APM) module, provide a powerful tool to guarantee continuous high application SLA throughout the entire application life cycle, by displaying actual user transactions and errors. The only APM-integrated ADC on the market, the solution enables easy detection and resolution of SLA degradations, while eliminating the need to manually script synthetic transactions. Cross-ADC infrastructure historical reports on resource utilization provide a holistic view enabling better capacity planning when rolling out new applications. In addition, drilldown-able real-time dashboards, that span multiple ADCs, enable instant visibility for spotting problems and a powerful tool for fast and accurate troubleshooting.

Reason 4: Application Awareness with AppShape™ & AppShape™++

Radware's AppShape technology transforms the ADC into a 'smart' device to accelerate, ease and optimize application deployment on the ADC. With Radware's AppShape, each ADC service is tailored to and *aware* of a specific business application (such as SAP, Microsoft, Oracle, IBM and more). In this way, the ADC can be managed from an application-oriented perspective via application specific configuration templates and wizards – resulting in fast application roll-out and simplified application management. Plus, AppShape offers logs and reports for: compliance, per application trends analysis and resources utilization.

Radware's also provides ADC policy scripting capabilities with its AppShape++ technology to further enable the customization of the ADC service per specific application flows and scenarios. By leveraging scripts, examples in Radware's library and dev-community, customers can easily use AppShape++ to refine various layer 4-7 policies including HTTP, HTTPS, TCP, UDP, SSL and more – with no application modifications to further reduce cost and risk.

Simply the Best-in-Class ADC Choice

The combination of these advantages – along with an industry unique 5-year longevity guarantee, "pay-as-you-grow" approach in throughput, # of vADCs and services, plus performance leadership in all layer 4-7 metrics – makes Alteon 5224 simply your best application delivery choice. Want to see for yourself? We invite you to download our Radware ADC Solution white paper here or contact us at: info@radware.com.

There can be no doubt that increasing user demands, computing trends and new technology are placing incredible expectations on IT professionals. Thanks to Riverbed Technology, you can now "right-place" your IT infrastructure in a way that meets growing application and data needs — without sacrificing performance or blowing your entire budget.

Get started now at **www.datacenter.riverbed.com**

riverbed®
Think fast.®