



Evolving Application Delivery Solutions for
the Virtualization Imperative - Solution Paper




Table of Contents

Executive summary	3
WAN has become critical to business and IT	3
Datacenter consolidation	4
Business Continuity/Disaster Recovery readiness	4
Workforce mobility	4
Customer/supplier integration over the network and uptake of online services	4
Voice and video over IP	4
Conclusion.....	4
Challenges in designing and managing a WAN optimized for application delivery	4
Availability and Business Continuity management	4
Service level management (performance).....	5
Capacity planning and financial management	6
Optimizing application delivery WANs with multi-WAN load balancers	6
Ensuring availability and business continuity	6
Optimizing service levels and performance	7
Creating a flexible, scalable and cost-effective WAN	8
Gaining control with bandwidth management	8
Case Studies	9
Case study: Calculating the cost of downtime.....	9
Case study: Regency Hospital	9
Case study: Unicom.....	9
The Radware advantage: what to look for in a good multi-WAN load-balancing solution.....	10

Introduction to Virtualization Market Trends

As the virtualization market matures, new technologies and applications are driving fundamental changes in IT infrastructure. This white paper discusses how virtualization impacts data centers and the business-critical applications they support.

According to Frank E. Gillett's Forrester report of February 2009, "Adoption of x86 server virtualization neared 50% ." To remain competitive, data centers must meet the challenges posed by virtualization and even capitalize on nascent opportunities this new market offers.

Virtualization is a growing market with a potential size of 18 billion USD by 2010 (according to IDC 2006) and with an expected 73% CAGR (according to Morgan Stanley CIO survey May 2008).

IT organizations are approaching server virtualization in order to cut server costs. However, organizations that have a mature server virtualization deployment in place are leveraging virtualization for much more. All market indicators suggest that the virtualization market is evolving from consolidating servers to achieving more complex tasks that require automation. According to Gartner, the current phase of virtualization focuses on server consolidation but the upcoming second phase of virtualization is more strategically important, more complex to implement and provides far more value for the customer. This phase concentrates on the flexible allocation of resources to respond to business demand, such as starting up more virtual machines or assigning more resources to a virtual machine with a priority task. In this phase, the focus shifts to delivering new services or improving the quality and speed of service.

Even though virtualization is expected to grow rapidly, the data center is not expected to be fully virtualized in the foreseeable future. According to Gartner, 18% of [server workloads](#) in 2009 run on virtualized servers; that share will grow to 28% in 2010 and reach almost half by 2012. IDC also adds that 16.5% of all new servers shipped in the second quarter of 2009 were virtualized.

As virtualization is a fairly new technology and introduces many new challenges, IT organizations are reluctant to virtualize their entire data center and tend to rely on the dedicated, physical realm when deploying mission-critical applications. As a result, as organizations continue to virtualize their data centers, mission-critical applications will be deployed in a hybrid environment where some of the application components are deployed on a dedicated physical infrastructure while others are deployed in the virtual environment or are dynamically provisioned on-demand only when required. A CIO Research survey (2008) shows that the highest ranked challenge for a successful virtualization project is balancing servers workloads and maintaining application service levels. Furthermore, in the same survey, 43% of respondents stated that they are tackling the technical challenges in virtualization by investing in new maintenance tools to balance workload and maintain application service levels.

Ensuring Virtualized Application Delivery and Quality with Radware ADC

Challenges

High Costs of Downtime in Virtual and Hybrid Environments

A typical SLA of 99.5% uptime permits 43.7 hours annually of application downtime. Considering the fact that one hour of downtime per line of business application, typically costs about \$120K emphasizes the high cost involved in application downtime.

Due to the nature of the virtual environment, multiple applications are deployed on one physical host. Virtualization vendors claim a 1:6 or even 1:10 consolidation ratio (meaning each physical host runs 10 virtual machines simultaneously). In this type of environment, the risk of application failure is higher, since a physical server

failure implies downtime for multiple applications. The high availability (HA) capabilities integrated into server virtualization solutions suffer from several deficiencies:

- Typically, availability is monitored at the OS and Host levels, meaning that action is taken only when a server fails or the OS installed on the virtual machine fails, but not when there is a failure of the OS or the application software running on top of the virtual machine. Failure in the network connection to the server might also not be detected.
- Since all server virtualization solutions are software based, they monitor availability by implementing a software agent on each physical server, thus consuming resources on all servers in the virtual infrastructure for this purpose only.
- Failure recovery time is typically as long as restarting a virtual machine on a different server. This may vary based on the load on that machine and can, in some cases, take up to several minutes. This means that during this period, application service is denied to the users that were working on the failing virtual machine.

In order to fully address business application availability, all the above challenges must be addressed. Failures in all levels (server, OS, network, application and data) must be detected and, when a failure occurs, users need to be serviced by an alternate resource instantly without any delays.

Ensure SLA and Improve QoE

Service level agreements (SLAs) and Quality of Experience (QoE) are important metrics for businesses whether dealing with internal or external stakeholders. Maintaining SLAs and improving the QoE of end users are primary responsibilities entrusted to the care of the IT organization, which oversees IT service delivery. In a virtualized environment, this responsibility becomes much more complex since hundreds or even thousands of different virtual machines (VMs) may share the same physical resource pool and therefore compete for the same resources. This competition for resources only increases the level of difficulty involved in monitoring, controlling, and enhancing application performance.

Different tools like VMware's Distributed Resource Scheduler (DRS) tries to address these challenges, but DRS suffers from several deficiencies:

- DRS optimizes each virtual machine (VM) placement individually. A common use case is that a virtual application is deployed on several VMs. In order to optimize such a virtual application, the entire set of VMs that support the virtual application must be viewed as one application. Optimizing each one of the application's VM separately may not provide the best results.
- DRS acts on information gathered from inside the virtual environment, such as CPU and memory, and does not take into account other important metrics like network latency and real user-experience related metrics. In many cases internal information like CPU and Memory are not enough to detect poor responsiveness of an application.
- DRS is bound by limitations imposed by the available resources. DRS cannot add or remove resources to support each application load. DRS is less effective in cases where more physical or virtual resources are required to support an application desired QoE.

Scalability Challenges in a Virtualized Dynamic Business Environment

Mission-critical applications serve the corporate business needs. The applications are typically very dynamic, based upon company growth and the number of users, whether they are employees, partners, or customers. There may also be a requirement to decrease application capacity due to changing conditions in the market. In addition there are inherent seasonal peaks in business-application usage such as end of quarter activity, end of financial year, etc. These peaks can also influence capacity requirements.

Each of these changes is part of normal business practice and therefore must be implemented with minimal costs and maximum efficiency. This translates into the challenge of building the application infrastructure so that capacity can be changed, and resources such as servers, storage, or bandwidth can be added or removed transparently, on the fly, without any application service interruption and without any application code changes.

On top of the above challenges, regular maintenance requires IT organizations to seamlessly remove from operation a VM or a server that requires maintenance. Removing the VM or server must be done in a way that will be completely transparent to the end users in order to ensure 24x7 application delivery.

Server virtualization solutions contribute significantly to addressing these challenges by providing a quick and easy way to add new virtual machines or stop their operation, and by providing quick and transparent virtual machine mobility from server to server without service interruption. However, typically, the integrated clustering capabilities in a server virtualization solution have the following limitation:

- The integrated clustering is based on load balancing functions that are done by software running on all servers. Requests are broadcast to all virtual machines in a cluster and the selected virtual machine is chosen in coordination.
- The clustering software on the respective virtual machine forwards the request to the application. This solution wastes resources and requires changes in the configuration of the servers.
- All servers must handle all packets, even if they do not process them in the application, thus traffic peaks impact all the servers in the cluster unnecessarily. Scalability is typically limited to 32 virtual servers in a cluster.
- When stopping or suspending a virtual machine, all users' sessions that were started with that VM are cleared and end users must recreate the session with another VM. This typically translates into the users need to re-login and start over their session losing all their previous state.

Over Provisioning of Virtual Machines

The process of physical to virtual migration (P2V) involves a capacity planning process. The capacity planning process tries to estimate the number of VM instances needed per each virtualized application. Due to the complexity of virtualized environments, estimating the right amount of VMs is not an easy task and usually involves some educated guesswork. As a result, the number of VMs deployed in the virtual environment may exceed the actual number of VMs required to satisfy QoE and comply with SLAs.

Unlike physical machines, VMs are quick and easy to deploy. The process of moving or copying a VM can be done in minutes, and often without appropriate IT oversight. This ease of deployment may lead to VM sprawl - the proliferation of virtual machines without adequate IT control. Both the above arguments translate into direct costs. Virtual sprawl can consume the entire software license budget, more and more administrator time, and eventually require the purchase of more physical servers and disks than needed, negating a major benefit of virtualization and one of the foundations of its return on investment.

Overspending to support future business growth

As virtualization is a fairly new technology, companies tend to start with a small virtualization project and plan to expand it over time. Due to its nature, a small virtualization project requires a smaller outlay on the supported hardware (hosting servers as well as network devices), than a larger virtual deployment that serves thousands or even millions of concurrent users. In such a challenging deployment, IT organizations must prevent overspending on hardware up-front but, at the same time, build an environment that will be able to support future business demand.

High labor costs

As more applications migrate to the virtual environment, the virtualization environment becomes more complex to manage manually. More applications sharing the same physical resource pool, competition for resources, monitoring for QoE, and preserving good QoE for virtualized applications makes tasks more complex to achieve manually. IT personnel require new sets of skills to manage virtualization environments; training a large IT staff to manually support the complexity of virtual environments implies high costs of labor.

High Fees Involved in Consuming Resources from a Cloud Service Provider

Forrester defines cloud computing as “a pool of abstracted, highly scalable, and managed computer infrastructures capable of hosting end customer applications and billed by consumption”. The cloud service providers business model is based on consumption; the more resources consumed by the organization, the more the organization must pay for the cloud-based service. Organizations that use cloud-based services must minimize usage of the cloud service to save on costs.


The Solution: Radware ADC Solution

Radware APSolute is a product suite designed to make networks business smart and to ensure application delivery to users anytime, anywhere. Radware APSolute integrates seamlessly into a virtualized server environment providing applications with full availability, enhanced performance, and further cost savings for a virtualized data center as well as for hybrid data centers that consist of virtual and physical resources.

Radware’s Application Delivery devices include:

- AppDirector: This intelligent application delivery controller (ADC) for the data center enables high availability, optimized application performance, application security, and infrastructure scalability. AppDirector provides application acceleration features, such as web compression, secure socket layer (SSL) offloading, TCP optimization, and caching for the fastest application and transaction response times.
- VirtualDirector is an application delivery optimization solution for the virtualized data center providing real time dynamic allocation of data center resources, based on QoE metrics and business events, ensuring positive user experience and improving response time. The solution aligns virtual data center operations with business policies while optimizing the use of virtual as well as physical resources to further generate CAPEX and OPEX savings.

Radware ADC integration into server virtualization environment:



Radware ADC integration into a Virtual Server Environment

Radware ADC, a member of the Radware APSolute product suite, seamlessly integrates into a virtual server environment.

Normally in virtual server solutions each virtual machine has its own IP and the applications running on top of each virtual machine have standard or proprietary ports. Radware ADC is designed to work with this type of application environment, and also supports multiple instances of the same application running on the same virtual machine.

Radware AppDirector can maintain session persistency so that a user transaction starting on one virtual machine will continue to be served by the same virtual machine as long as this virtual machine and the application on top of it is up and running.

In addition, Radware ADC transforms a virtual server into a virtual application environment by the abstraction of the entire infrastructure underlying a certain enterprise application service including servers, virtual machines, operating systems and application software instances. For the end user it appears as a single application service.

AppDirector is deployed in front of the hybrid or fully virtual applications ensuring 24x7 availability and high performance for those applications.

VirtualDirector is deployed passively in the network (connected to a TAP or a copy port on a switch) – outside the virtualization environment, unaffected by virtualization constraints, and providing full visibility to the virtual and hybrid application transaction flow without introducing any latency or risks to the production environment.

VirtualDirector is connected to AppDirector as well as to the virtualization management application. This connectivity provides VirtualDirector with the ability to take actions on AppDirector as well as on the virtualization layer. Furthermore, utilizing this connection, VirtualDirector is capable of retrieving KPI information from the virtualization layer.

Solution Benefits and Value

Full Virtual Application Availability

Using Radware ADC, application availability can be obtained by using the following:

- Advanced health monitoring - Radware ADC has an advanced module for health monitoring, which includes an extensive library of pre-defined health checks. Using these checks AppDirector can be configured to identify any type of failure, whether it is a server hardware failure, a virtual machine failure, network cable disconnect, an operating system problem, a specific application failure, or a back-end database failure. Each component that is required to complete an ERP or CRM transaction can be monitored to ensure that users are never directed through a transaction path which is not fully operational. Furthermore the back-end database health can be monitored to ensure a correct result is returned. All the checks are easily set up without any need for scripting or programming. This ensures a straight-forward and error-free implementation and reduces the management costs involved. All health monitoring checks are performed by Radware ADC, without any software agents running on the server, avoiding any server overhead or software compatibility issues.
- Transparent failure bypassing to ensure 24x7 availability - Once any type of failure (server, network, virtual machine, OS, application, or database) is identified, Radware ADC automatically and instantly redirect users to operational resources, without waiting for a virtual machine restart process. This means that if one of the components required for completing a transaction is down, AppDirector directs users through an alternative path thus ensuring end-to-end transaction completion. Radware's health monitoring and failure bypassing ensures complete availability of end-to-end application operations. Users are directed to the least loaded virtual machine in the pool, without wasting stand-byspare resources. During all times all the virtual machine and servers are active and providing service to the users.

- Full disaster recovery and business continuity with Radware's global solution – Ensuring maximal availability of application resources in a single data center may not be enough in the face of natural or man made disaster. In these cases all the resources (servers, applications and WAN connections) are lost at the same time. In order to ensure critical application availability even in extreme conditions, application resources and infrastructure must be built across more than a single data center - in multiple locations. In a multiple data center implementation, the same ADC can be used for global and local traffic redirection, ensuring that users will be directed to operational resources globally even if there is a complete data center failure. Radware's global solution complements virtualization vendors' disaster recovery solutions due to its ability to control the traffic and redirect it to the operational site. The combination of all the above-mentioned capabilities provides an unmatched solution to address all enterprise virtual application availability challenges. All these technologies are completely transparent to the virtualization solution and to the applications; therefore they do not require any special development effort to achieve maximal application availability.

OPEX Savings by Optimizing QoE and Preventing SLA breaches

Radware's ADC ensures, in spite of server virtualization overhead and WAN communications delays, the fastest response time to remote end users, optimal application performance, and maximal network and server utilization. These benefits translate into quality of experience (QoE) improvement for the end user, improved employee productivity, and further cost savings by reducing the number of physical servers required to support the virtualized applications while still meeting the SLA requirements. In order to provide the best QoE for applications Radware ADC integrates a set of acceleration and optimization mechanisms into the virtual application infrastructure, including:

- By learning the real user quality of experience of applications and automatically taking actions accordingly, Radware's ADC solution preserves the users' QoE, and prevents SLA breaches for all the virtual applications or hybrid applications that it manages.
- By learning each application's real business impact, Radware's solution aligns the physical resources as well as the virtual resources with the real business needs in real-time, and prioritizes virtual and physical resources used by each application according to the application's current business needs.
- Load balancing and persistency - AppDirector optimizes server performance by redirecting users to the least loaded and best performing virtual machine. • Global service response time optimization - virtualized applications can be implemented over multiple data centers, and further service response time optimization can be done by the AppDirector global solution. In this scenario, users performing a transaction will always be redirected to the best performing virtual machine within the data center, which will give that specific user the best response time.
- End-to-end bandwidth management and prioritization - Radware ADC includes an integrated Bandwidth Management module. The APSolute bandwidth management module allows allocation of bandwidth according to business priorities and application importance. In a virtualized environment, where bandwidth is shared between different applications,
- Radware's bandwidth management solution ensures that mission-critical applications are never starved by non-critical applications and that they get the allotted bandwidth according to the business policies and priorities. By utilizing VirtualDirector, different bandwidth policies can be automatically applied in realtime, aligning the bandwidth available for each application with its real, ever changing, business priority.

- Offloading CPU and memory-intensive tasks from the virtual machines– Radware’s ADC solutions offers the following:
 - SSL offloading - SSL processing is very CPU intensive, thus it takes away resources from the virtual application processing itself and delays transaction response time. Furthermore, in a virtual environment where physical resources are shared between different virtual machines, offloading non-core functionality to the network, such as SSL processing, frees shared resources to be used by other virtual applications.
 - Web compression - AppDirector’s Web compression can accelerate response time dramatically. Compression also reduces the size of the replies, hence reducing the amount of bandwidth consumed.
 - Caching served information - AppDirector improves virtual application performance by caching requested information.
 - HTTP and TCP connection optimization - AppDirector optimizes HTTP and TCP connections, thus improving response time to end users as connections are managed faster by AppDirector, and offloads the virtual applications from unnecessary overhead that is not application related.
- As it is a network solution that is deployed outside the virtual environment, the Radware solution can monitor and optimize applications that are fully virtual and applications that are deployed on both virtual and physical resources.

Dynamic Business Oriented Scalability

Companies may increase the number of users, whether they are employees, partners or customers, or may need to decrease application capacity due to changing conditions in the market or seasonal demands. This needs to be done transparently to the application’s users. Using AppDirector’s load balancing and traffic redirection capabilities, virtual machines and servers can be added or removed on the fly, completely transparently to end users and without any interruption to application services. This enables enterprises to increase or decrease application capacity according to dynamic business needs, while maximizing their virtualization solution capabilities of virtual machine mobility and further reducing ownership costs.

Seamless and transparent maintenance operations

Often in any application environment there are standard maintenance operations such as virtualization solution upgrades, application software upgrades or operating system service packs installations that are required to ensure the ongoing operation of software applications as well as server operating system and hardware. In such cases, resources must be removed from operation seamlessly and transparently to users in order to ensure 24x7 application delivery. AppDirector has a built-in mechanism to enable such maintenance operations, allowing virtual machines to be taken out of service transparently by redirecting all users to alternative resources and notifying when all active users have finished their transactions on the ‘planned to be removed’ server. This operation can be done manually or automatically through an API.

Increase CAPEX and OPEX Savings by Reducing the VM Count

By taking action automatically, VirtualDirector can optimize the VM count, and align each virtual application-supporting VM count with the application’s real business needs. By adding, removing, stopping, and starting VMs according to QoE-based metrics, Radware solution can actively adjust the amount of VMs required by each application in order to comply with the business requirements - over provisioning of VMs can be detected, and unneeded VMs can be decommissioned, saving on licenses, management costs, and physical hardware required to support those unused VMs.

Reduce CAPEX and Prevent Over-spending

AppDirector's unique, pay-as-you-grow licensing model reduces the risk involved in physical-to-virtual (P2V) migration. By enabling customers to pay only for the capacity they really need, IT organizations can save on costs when creating the initial virtual environment. As the virtualization deployment grows, more virtual machines are added to support the increased load, and more physical servers are added to support the increased number of VMs. As a result of the growth in VMs and physical servers, the network throughput increases, and more capacity is required from the ADC solution. In this case, the IT organization can easily purchase a new license key and perform a software-based upgrade without any downtime, thus increasing the ADC capacity to support the new load. Radware's zero-downtime, pay-as-you-grow licensing approach, in combination with virtualization technologies, provides a zero-downtime upgrade for the entire virtual and physical infrastructure, thus ensuring that organizations can purchase capacity only when it is really required.

Save on Labor Costs

Virtualization introduces new management challenges. A new set of technical skills is required to support the virtual environment which tends to be more complex than a physical environment. To manage this complex environment, more trained IT workers are assigned to monitor and manage the virtual environment, especially if manual or semi-automated management is required. Radware's VirtualDirector enables IT organizations to automate the monitoring and QoE optimization tasks involved in managing virtual or hybrid environments. By offloading the optimization tasks from IT personnel, VirtualDirector frees those highly trained personnel from doing manual tasks, and enables them to manage more strategic tasks for the organization.

Minimize the Fees Paid for Cloud Service Providers

One of the uses of cloud-based services is to provide more capacity for applications or services that are hosted in the organization's own local data center. Radware's solution ensures that local resources are fully utilized before redirecting users to the cloud. Radware's ADC solution ensures that the organization uses the cloud-based services only when really needed, and only for applications with lower business impact. By redirecting traffic to the cloud only when local resources cannot provide the desired QoE, organizations optimize their use of the cloud service, which reduces the fees the organization has to pay for the cloud service provider.

Virtual Desktop Availability and Performance with Radware Solutions

Challenges:

High Costs Due to Link Connectivity Failure

In a virtual desktop environment the employees' desktops run in the data center and not on the employee's own PC. A thin presentation client running on each employee machine relies on an open remote desktop protocol (RDP) connection in order to send commands and present the results back to the user. This means that virtual desktop deployments rely on stable and reliable connectivity between the client devices and the data center. As the number of users that use virtual desktops increases, connectivity links become more congested. Congested links significantly degrade the quality of experience of virtual desktop users and may even prevent those users from accessing their virtual desktops. A broken connection or a heavily congested link translates immediately into the employees not being able to use their desktops and therefore:

- The employees cannot use mission-critical or revenue-generating applications; this translates into revenue losses.
- The employees cannot perform their tasks - high costs related to employee idle time and productivity loss

Even though different vendors offer link HA solutions to prevent broken connections, these solutions introduce seconds or even minutes of downtime during which open virtual desktop sessions are terminated and no new sessions can be opened.

Maintaining a reliable, 24X7 connectivity between all employees and their respective desktop instances running in the data center is a major challenge for all virtual desktop deployments.

High Costs Due to Virtual Desktop Broker Application Failure

Most of the virtual desktop deployments require a broker or management application to provide a single access point for all virtual desktop users. Virtual desktop users use this management application to log in to their virtual desktop and start their virtual desktop session. Furthermore, in most cases all the virtual desktop traffic passes through this management application, which serves as a gateway between the virtual desktop instance running in the data center and the end user. As a core component in every virtual desktop deployment, any downtime or slow responsiveness of this management application has a direct effect on the productivity of employees. If the management application is down, employees cannot start their desktop session and virtual desktop traffic cannot reach its destination.

Any downtime or poor performance of this management application translates immediately into employees not being able to use their desktops, and therefore not able to use revenue generating applications.

As a single point of failure, IT organizations that deploy virtual desktops are required to ensure the availability and performance of this management application.

High Costs of Hardware Required to Support Virtual Management (VDM) Deployment for Large Number of External Users

As the main access point for any virtual desktop session, the virtual desktop management application (VDM) should be able to handle the large number of users who require its services. Since usually all employees start new desktop sessions in the morning when the working day begins, and terminate their sessions in the evening before going home, it is crucial for the management application to be able to handle those concentrated peaks. Furthermore, one of the benefits of VDM is its ability to secure the connection between the end client and the data center using the SSL protocol. Parsing SSL transactions is a CPU-intensive task, VDM deployments that handle SSL termination must be robust, and usually cost more to implement since more hardware is required to support users at peak times.

High Cost of Investment in the Network Infrastructure to Support the Growth in Virtual Desktop Users

As virtual desktop technologies mature and organizations move towards larger deployments, the number of end users that use virtual desktops will increase and, as a result, the traffic that those users generate will also increase. Furthermore, as more organizations utilize VDI technologies in remote branches, the demand for bandwidth in those branches will increase. In this ever changing environment, IT organizations must continuously upgrade their link capacity and their network infrastructure to support the higher amount of traffic. Since any downtime in the virtual desktop environment translates into employee idle time, IT organizations should aim at preventing any downtime while still being able to continually perform the required upgrades.

QoE Degradation due to High Bandwidth Utilization

As virtual desktop technologies mature, emerging technologies will enable virtual desktops to deliver a more rich and personalized virtual desktop to any device (whether thick or thin). This implies that virtual desktops will support all types of multimedia applications as well as other bandwidth-demanding applications. The virtual desktop evolution toward a richer experience for the end user has many benefits but it also introduces a new challenge for the IT organization. In a virtual desktop deployment, branch employees' desktops run in the data

center and not on the employee’s own PC, which is physically located in the branch. This fact implies that branch users who open a connection to a location outside the data center are consuming bandwidth from the data center itself (since the actual desktop is running in the data center). In this environment, high-bandwidth demanding applications like multimedia applications are competing for the same bandwidth resources as mission-critical and revenue-generating applications that are deployed in the data center. Since multimedia or other content-rich applications tend to consume a high amount of bandwidth, mission-critical applications may not have enough bandwidth to operate officially, and therefore users of those revenue-generating applications will suffer from degradation in QoE which in turn may translate into revenue loss.

Ensuring good QoE for revenue-generating applications in such a high-bandwidth demanding environment is a challenge that IT organizations face when deploying virtual desktops.

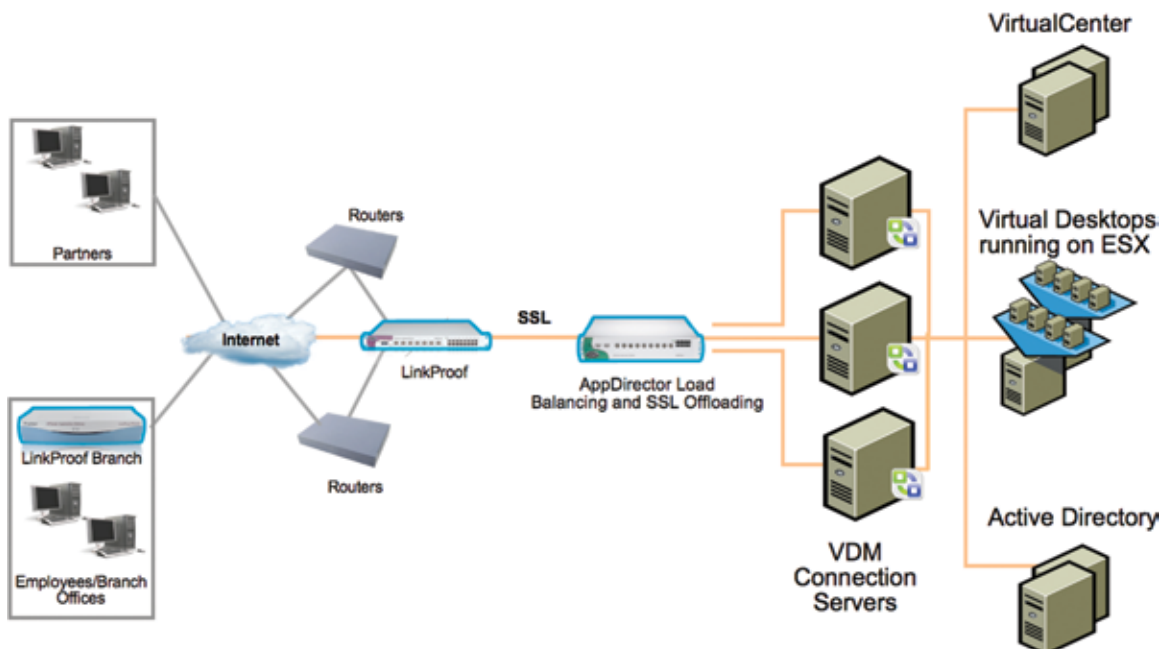
Solutions

Radware’s Application Delivery devices, part of the APSolute™ product suite, provide virtual desktop applications with full availability and enhanced performance.

Radware’s Application Delivery devices include:

- AppDirector: This intelligent application delivery controller (ADC) for the virtual desktop environment enables high availability for virtual desktop deployments, optimized performance and infrastructure scalability. AppDirector provides application acceleration features such as web compression, secure socket layer (SSL) offloading and TCP optimization for the fastest application and transaction response times.
- LinkProof™ is the industry leading multi-WAN load balancing solution, ensuring data center connectivity to the Internet and intranets through multiple ISPs and private lines, overcoming any connectivity loss. LinkProof™ enhances user experience and response time with its sophisticated proximity-based routing, smart link load balancing, end-to-end transaction path health monitoring, application-based routing, and application bandwidth management. With its “Add-link-as-you-grow” approach, LinkProof allows the addition of new Internet links on-demand as the organization grows its business and number of users.

Solution diagram



Radware Integration into VDI environment

- Deploying Radware's LinkProof Branch in each remote branch provides those branches with uninterrupted access to the desktop instances running in the data center. Furthermore, deploying LinkProof Branch in each branch enhances the virtual desktop experience and response time for employees at the branch offices.
- Deploying LinkProof in the perimeter of the data center ensures 24x7 connectivity for the data center. Deploying AppDirector in front of the VDM application servers ensures VDM 24x7 availability and high performance for the VDM application.
- Deploying AppDirector in front of the VDM security servers where SSL processing is required, offloads all SSL processing to AppDirector.
- Deploying LinkProofs in branches and in the data center provides transparent link scalability with no downtime, to support more virtual desktop users as the virtual desktop deployment grows.
- Deploying AppDirector and LinkProof in the data center provides the ability to control bandwidth allocation across different applications.

Solution Benefits and Value

Increase Employee Productivity and Prevent Mission-Critical Application Downtime by Ensuring End-To-End Connectivity

By using Radware's LinkProof solution, organizations are assured that link failures are prevented and as a result, virtual desktop users are provided with 24x7 connectivity to their desktop instances running in the data center.

LinkProof can be deployed in the main data center as well as in every branch that uses virtual desktops, providing the organization with reliable and continuous Internet access and full link protection. In addition LinkProof maximizes the performance of Internet links and as a result improves the virtual desktop usage experience.

Connectivity failure translates directly into employees not being able to use their desktops and therefore renders useless any mission-critical application that the virtual desktop user uses. LinkProof connectivity assurance enables organizations to preserve high productivity levels from all employees, which in turn prevents downtime of mission-critical applications, ensuring that users can access the mission-critical applications at any time.

By deploying LinkProof in remote branches, organizations that deploy virtual desktops for their remote branches eliminate the risk of a link failure between the remote branch and the data center. Since a link failure in a branch prevents all the branch employees from using their desktops, the cost of a link failure in a branch is equal to the revenues lost due to the branch downtime. One of LinkProof's unique features is its ability to provide 24x7 connectivity for VDI users by providing micro-second switching between links in case of a link failure. This capability guarantees that when a link fails, LinkProof will switch to the backup link without losing a single virtual desktop connection. In addition, LinkProof is capable of detecting a highly congested link and re-route virtual desktop traffic to a less congested link in order to increase performance of virtual desktops.

Increase Employee Productivity by Ensuring Zero-Downtime for the Broker Application

To ensure maximal availability for virtual desktop users, the VDM application must perform well and be highly available. By deploying AppDirector in front of a VDM server cluster, organizations eliminate the risk that the VDM application will be a single point of failure. AppDirector provides 24x7 availability and accelerated performance for the VDM application, ensuring that employees, at any time, will be able to start a new virtual desktop session and virtual desktop traffic will be able to pass through the VDM application. AppDirector advanced health monitoring is specifically designed to monitor the health of the VDM application. Once any type of failure is identified, Radware ADC automatically and instantly redirects users to operational resources. Furthermore, at any time, virtual desktop users are directed to the least loaded VDM instance, improving the performance of the VDM application.

By providing 24x7 availability and better performance for the VDM application, Radware's ADC solution maximizes employee productivity and ensures that employees can access their virtual desktop and use mission-critical applications.

Reduce CAPEX Involved with VDM Deployment

By utilizing AppDirector's SSL off-loading capability, IT organizations can eliminate the need for VDM security servers to perform the resource-demanding operation of SSL processing, and off-load the SSL tasks to the AppDirector. Since in most virtual desktop deployments, VDM security servers are required to do the SSL processing, by off-loading SSL tasks to AppDirector, less physical servers are required to support the VDM deployment hence the CAPEX involved in VDM deployment is reduced.

Save on Costs with Transparent Scalability

Radware's on-demand scalability provides IT organizations with a simple method to transparently upgrade connectivity links or network devices that are required to support more traffic. As more capacity is required from Radware's devices to support the ever growing number of virtual desktop users, a transparent upgrade can be done by just providing a new licensing key. This simple and easy software-based upgrade prevents the downtime that is required by any traditional forklift upgrade. Furthermore, by utilizing LinkProof, as more bandwidth is required in the branch or in the data center, the IT organization can transparently, without any downtime add more bandwidth by adding more links to support the new bandwidth requirements. In addition to the above benefits, LinkProof prevents organization from been locked in with a specific Internet service provider enabling a transparent, zero downtime transition between different ISPs. This functionality provides IT organizations with the ability to save on costs by transparently switching, at any time, to the most cost-effective Internet service provider.

Prevent QoE Degradation of Revenue-Generating Applications Caused by Virtual Desktop Deployment

When deployed in front of the VDM application servers, all virtual desktop traffic can be routed through AppDirector. Utilizing AppDirector's advanced bandwidth management module, policies can be defined to provide high bandwidth priority for virtual desktop users who access revenue-generating applications. Bandwidth assigned to content-rich applications, like multimedia applications, can be limited, ensuring that mission-critical applications have enough bandwidth to serve customers thus preventing QoE degradation for those applications. When deployed in the perimeter of the data center, the LinkProof bandwidth management module can limit the entire bandwidth that is consumed by content-rich applications, freeing bandwidth for revenue-generating applications as well as for virtual desktop protocol communication. Furthermore, the LinkProof patented proximity algorithm guarantees that the path selection between the remote branch office and the data center is optimized for best performance ensuring the best possible performance for applications used by virtual desktops users.

On Demand Application Capacity with Radware ADC

Challenges

The Conflict Between Reducing Hardware Costs and Complying with Strict SLAs

Complying with strict SLAs and reducing hardware costs are two business requirements that are usually not aligned with each other. Organizations tend to perform capacity planning based on each application's individual peak behavior; this is done mainly since each application manager needs to ensure that the application complies with a strict SLA as defined by the business. Designing application by peak usually translates into high costs involved in buying the hardware equipment as well as high costs involved in managing it. Reducing the amount of hardware that is required to support an application during its peak time will most likely influence the application ability to provide good quality of experience during the peak times.

Virtualization of Mission Critical Applications Introduce a High Business Risk

Due to the apparent cost savings that virtualization introduces, organizations aim to virtualize as many applications as possible. But organizations are reluctant to migrate mission-critical applications to the virtual environment due to the risks that migration introduces. In a virtual environment, different applications share the same physical resources; therefore virtual applications compete for the same physical resources. When mission-critical applications are virtualized, they share the same resources with other virtualized applications. Resource sharing, on one hand, provides cost savings as less physical servers are required to support the same number of applications; on the other hand, IT organizations are worried about the risk it introduces. The traditional “one workload, one box” approach to server provisioning is considered by many IT professionals as a simple way to control the application workload and predict the application behavior. The new complexity challenges that resource sharing introduces, as well as the uncertainty it creates, cause many IT organizations to consider virtualization as a risky technology for mission-critical applications. Some organizations predict that the risk involved in deploying mission-critical applications in a virtualized environment surpasses the benefits that it promises. For that reason IT organizations are reluctant to migrate mission-critical applications to the virtual environment.

Solutions

Radware’s Application Delivery solution, part of the APSolute™ product suite, provides on-demand scalability for mission-critical applications as well as ensuring full availability, enhanced performance, and maximum security for those mission-critical applications.


The solution enables organizations to plan capacity based on average application load instead of the traditional capacity planning by peak. By planning capacity by average, less dedicated servers are required to support each application since the average load on an application is usually less than the load on the same application during the application’s peak times.

When an application requires more physical resources to provide good QoE for its users, Radware’s ADC solution automatically provisions a new dedicated server for the application from a pool of servers.

The Radware ADC solution utilizes the advanced provisioning capabilities that VMware VirtualCenter provides to provision new dedicated servers, but the solution ensures that at any given time each server in the pool is dedicated to only ONE application, eliminating any resource sharing between applications. Radware’s Application Delivery devices include:

- AppDirector: This intelligent application delivery controller (ADC) for the data center enables high availability, optimized application performance, application security, and infrastructure scalability. AppDirector provides application acceleration features, such as web compression, secure socket layer (SSL) off-loading, TCP optimization, and caching, for the fastest application and transaction response times.
- VirtualDirector is an application delivery optimization solution for the data center providing real time dynamic allocation of data center resources, based on QoE metrics and business events, ensuring positive user experience and improving response time. The solution aligns data center operations with business policies while optimizing the use of virtual as well as physical resources to further generate CAPEX and OPEX savings.

Solution diagram



VirtualDirector and AppDirector Integration into Mission Critical Hybrid Environments

- AppDirector is deployed in front of mission-critical applications ensuring 24x7 availability and high performance for those mission-critical applications.
- VirtualDirector is deployed passively in the network (connected to a TAP or a copy port on a switch) providing full visibility to the mission-critical applications’ transaction flow without introducing any latency or risk issues to the production environment. By learning the real QoE of each mission-critical application, VirtualDirector can automatically (on-demand) add more dedicated resources for the mission-critical application when more resources are required to provide good QoE.

Solution Benefits and Value

CAPEX Saving While Complying With Strict SLAs

In a typical data center, different applications peak at different times. Moreover, on average, the dedicated servers supporting those applications are underutilized and operate at about 5%-15% of their total load capacity. As a result, during most of its operational hours, an application requires less dedicated servers than was initially assigned to it by the traditional design by peak. By utilizing Radware’s on-demand capacity solution, IT organizations can reduce the number of servers required to support each application hence increasing CAPEX savings. Radware’s on-demand capacity solution enables organizations to initially deploy fewer dedicated servers per application. Radware’s ADC solution automatically provisions dedicated servers per application (from a shared pool of servers) only when the application requires the additional workload.

By planning the capacity according to average application load and not according to peak, IT organizations can prevent overspending on dedicated servers that are required during the application peak times only. Radware’s on demand capacity solution monitors the QoE of applications and ensures that the QoE of applications is positive even during peak times. The solution aligns, in real-time, the physical resources that are available for each application with the application requirements to support its SLA.

Minimize the Risk Involved in Consolidating Mission-Critical Applications

In an on-demand environment, applications are always deployed on dedicated resources to ensure that the QoE of those applications is good, thereby reducing the risk of SLA breach. The main difference between an on-demand environment and a traditional design-by-peak environment is the number of dedicated resources per application. An on-demand environment usually consists of less pre-provisioned dedicated servers per application. When more resources are required by the application to support its workload, Radware's on-demand solution automatically adds dedicated resources from a shared pool that is large enough to satisfy peak demand. At any given time, each application is deployed on dedicated resources; some resources are especially pre-provisioned for the application and other resources are provisioned automatically by Radware's on-demand solution.