



F5 White Paper

# Creating a Hybrid ADN Architecture with both Virtual and Physical ADCs

The virtualization of network and application network infrastructure is the second wave of the virtualization tsunami to hit the shores of the data center. Unlike server virtualization, because of its unique role in the data center, Application Delivery Controller (ADC) virtualization brings with it architectural implications that make a simple virtual-for-physical replacement strategy unacceptable. But there are appropriate places across the data center and organization where virtualized ADCs can be leveraged as stand-alone solutions, as well as in conjunction with its physical predecessor, to enable a more dynamic data center without compromising reliability, scalability, and performance.

**by Lori MacVittie**  
Technical Marketing Manager



# Contents

<b>Introduction</b>	<b>3</b>
<hr/>	
<b>Who Should Deploy a vADC</b>	<b>4</b>
Enterprise Data Center Deployment	4
Independent Software Vendors	4
Cloud Computing Environments	5
<hr/>	
<b>Architectural Challenges</b>	<b>5</b>
Scalability	6
Adaptability	8
Mobility	8
<hr/>	
<b>Recommended Best Practices</b>	<b>9</b>
Virtual ADC Architectural Considerations	10
Physical ADC Architectural Considerations	11
<hr/>	
<b>Conclusion</b>	<b>12</b>



## Introduction

Virtualization has long since moved from an emerging data center technology into a standard deployment tool across organizations of all sizes, regardless of industry. A virtualization survey from Rackspace in August 2007 indicated a strong trust in virtualization, with 72 percent of respondents noting a willingness to deploy virtualized applications in a production environment.<sup>1</sup> Since then, a recent Shavlik Technologies' survey found that 75 percent of respondents had already virtualized half of their production servers.<sup>2</sup> Other recent surveys, including those from respected analyst firms, confirm such trends, with virtualization adoption rates proving that most organizations have embraced data center virtualization.

In spite of the widespread adoption, virtualization still presents a variety of challenges. Research conducted by CIO.com in early 2008 uncovered numerous challenges in both deploying and managing virtualized environments, with the top challenge (64 percent) being balancing server workloads and maintaining application performance.<sup>3</sup>

As virtualized technologies have become more widely deployed, adoptees have turned to virtualizing other data center components in an effort to increase IT agility and better respond to changing business needs. Network and application network vendors are under the gun to provide virtualized versions of traditionally physical data center devices, leading to an increase in virtual network appliance offerings.

Successful deployment of these virtual network appliances (VNAs), however, can be difficult without the necessary architectural guidance. Server virtualization has minimal impact on the overall data center architecture. The deployment of VNAs, however, can be highly disruptive due to the unique interdependencies of network infrastructure components. In some cases it may not be desirable or technically feasible to replace a traditional physical appliance with a VNA. Just as some workloads are not a good fit for server virtualization, not every network component should be replaced with a VNA. Because of this, organizations will end up with a network architecture comprised of both physical and virtual versions of a solution. Such architectures will be necessary to support the flexibility and scalability required by organizations without sacrificing reliability and availability.

The availability of a virtual Application Delivery Controller (vADC) affords organizations the opportunity to include such a solution in a variety of ways and at all points along the application development lifecycle. In many cases deployment of a vADC in place of its physical counterpart requires little or no changes to the

Organizations with 500 or more employees are significantly more likely to measure success based on increased business agility.

Source: CDW's Server Virtualization Life Cycle Report, January 2010



architecture. In others, however, the insertion of a vADC into the network in place of a physical Application Delivery Controller (pADC) can be problematic.

## Who Should Deploy a vADC

There are three distinct organization types that will benefit from the deployment of vADC:

- Enterprise data centers
- Independent Software Vendors
- Cloud and hosting providers

### Enterprise Data Center Deployment

#### **Network administrators and architects**

Network administrators and architects benefit most from a vADC when it is deployed in testing and QA environments. A vADC, in conjunction with server virtualization technology, enables the organization to replicate production environments without a significant investment in physical components. By employing virtualization across all components of the architecture, testing of new solutions and optimization of existing policies can be accomplished in an isolated environment and then migrated to production.

#### **Developers and application architects**

While developers and application architects are aware of the benefits of integrating application delivery technologies with applications, the cost of providing developers with an accessible physical component has prevented adoption. The use of production-deployed pADCs is discouraged, with good reason, and thus the integration of application delivery technology has been largely left unexplored. The ability to use these technologies by deploying a vADC opens the way for developers to include and leverage acceleration, security, optimization, and customization of the application delivery platform in conjunction with the design and deployment of their specific applications.

### Independent Software Vendors

The lack of management and orchestration systems capable of managing both

To make virtual networks flexible and manageable, programmability of the network elements is of utmost importance. Only through programmable network elements, it will be possible for the service providers to implement customized protocols and deploy diverse services. Hence, the design decisions: “how much programmability should be allowed,” and “how it should be exposed” must have satisfactory answers.

Source: “A Survey of Network Virtualization”, October 2008, N.M. Mosharaf Kabir Chowdhury, et al.



physical and virtual application and network components complicates the adoption of many virtualization and cloud computing models. A lot of the benefits of cloud computing and virtualization come from the ability to automate and orchestrate IT processes related to on-demand provisioning. Much in the same way a vADC benefits developers and architects in the enterprise, a vADC provides independent software vendors with a cost effective way in which application delivery can be integrated into management solutions and leveraged to create new, innovative uses of application delivery technology.

## Cloud Computing Environments

Cloud and traditional hosting providers, and by extension their customers, will see as much benefit from a vADC—if not more—as enterprise adopters. The highly dynamic nature of such environments requires the flexibility and rapid scalability associated with virtualized solutions. While traditional pADC components provide both flexibility and rapid scalability, they are less able to efficiently address the need to isolate the application delivery policies of thousands of customers.

Cloud and hosting providers can use a vADC to differentiate their offerings from competitors by providing customers with the ability to deploy enterprise and carrier-class application delivery technologies in conjunction with their cloud-based applications. Giving customers the ability to deploy their own vADC alleviates the need for the provider to open up its pADC to customers. This then mitigates concern about the isolation of configuration and components within the cloud computing provider's infrastructure.

The customers of cloud and hosting providers can leverage the capabilities of best-in-class pADCs in a virtual form. With a vADC in the target cloud environment, the process of migrating applications that might be dependent on application delivery technologies for security, availability, or performance-enhancing capabilities is a much less painful prospect. In some cases it could mean the difference between needing to rewrite the application for a cloud environment and simply deploying both the application and the vADC instead. The ability to bundle together application delivery technology with the application gives organizations more choice in cloud providers as they will no longer be dependent on the provider for access to application delivery components.



## Architectural Challenges

It would be premature to assume that virtualized Application Delivery Controllers will replace their existing physical Application Delivery Controller counterpart. In some cases a pure replacement strategy might be possible, but in others the impact would be such that retaining (or obtaining) a pADC and augmenting its capabilities via a vADC would be recommended.

The optimal Application Delivery Network architecture takes advantage of both physical and virtual Application Delivery Controllers to bring the necessary mobility, scalability, and adaptability to enable a dynamic data center. A hybrid approach to architecting a flexible and adaptable, yet highly scalable and well-performing Application Delivery Network garners most success for organizations embracing the virtualization of network and Application Delivery Network components.

### Scalability

Scalability, particularly on-demand or “auto” scalability drives most of the demand for VNAs. This stems primarily from the way in which a pADC is scaled upon reaching capacity—either as a physical replacement or as an additional deployment. The time and consequential time to acquire a physical replacement can be highly disruptive. It is assumed that if one could simply “spin up” additional vADCs to address capacity constraints it would be faster, less expensive, and more seamless transition. Scaling to a vADC, while would faster less costly, would not be less disruptive. The transition might, in fact, be detrimental to the performance and availability of the network and the applications being delivered.

The scalability of VNAs automatically assumes that it can also distribute requests/traffic across multiple instances of the virtual appliance, a role typically assigned to a capable pADC. That role, however, does not automatically migrate from the pADC to the VNA simply because it is now virtualized. The result is that scalability of vADC, and VNAs in general, will continue to require a pADC.

### Scaling Out

Roles do not migrate automatically due to the core way in which network components virtualize other components,—for example, how they distribute traffic across multiple instances of a device. In a completely virtualized architecture it is often thought that scaling out is possible by simply by spinning up a new virtual appliance, resulting in what is known as an active-active (or n-active) configuration



in which all instances are actively accepting and processing traffic/requests. Herein lays the problem: all instances are accepting and processing all traffic/requests because to implement such a configuration from a network perspective, all instances must share a common MAC address to which all traffic/requests are directed. The individual solutions must then determine which instance will process the traffic/request. Essentially, this architecture turns upstream switches into hubs and duplicates network traffic. The more a solution is scaled out, the more bandwidth will be consumed by this broadcasting behavior. This can negatively impact network capacity and performance, not to mention making troubleshooting more difficult, especially in environments where visibility is already limited such as cloud computing.

### **Scaling Up**

Scaling up eliminates network issues by leveraging additional compute resources as a means to increase capacity. This approach, while valid, has limitations and unlike scaling out, is just as disruptive in implementation as scaling up traditional network components. Increasing compute resources often requires new hardware, and the difference between replacing commodity hardware and specialized hardware is nonexistent.

Moreover, scaling up has an upper bound capacity limit that the VNA cannot traverse. Limitations on addressable memory constrain scalability as will inherent performance degradations caused by network software that necessarily focuses on maintaining connection-oriented “lists” internally that, when grown too large, take longer and longer to access.

Scaling up a VNA also completely reverses any performance and efficiency improvements gained through integration with specialized hardware. The packet-processing engines, algorithmic acceleration, and protocol optimizations found in specialized hardware cannot achieve the same performance rates when moved to commodity hardware.

### **Hybrid Scalability**

A good rule of thumb in determining whether it is appropriate to deploy a pADC or a vADC (or any VNA versus its physical counterpart) is to determine its core functionality. If the solution aggregates traffic/requests and distributes across devices or networks, for example, routers and load balancing, then it not a good candidate for virtualization because of the challenges associated with scaling VNAs. A second guideline to follow is this: If the solution is already being load balanced by



a pADC, it will almost certainly still require a pADC when it is deployed in a virtual form. Even if the solutions—for example, a network firewall—are not currently load balanced that it might be necessary to implement a load balancing solution to scale the VNA when moving from a physical to a virtual deployment model.

That does not mean that there is no room in a production environment for a vADC. On the contrary, a vADC will complement a pADC and other VNAs quite well. For some functions, however, it is not architecturally advantageous to deploy virtualized versions to support those functions.

## Adaptability

Adaptability, or flexibility, is another often cited reason for implementation of virtual solutions. It is certainly far easier, and faster, to deploy a vADC than a pADC when there is an immediate business or technical need to address a problem. This is particularly true in cases where a temporary “fix” is necessary. The ability to dynamically deploy application delivery-related technologies such as web application security in an on-demand fashion enables IT to adjust to changing business and technical demands with little disruption.

The use of a vADC to augment existing pADC implementations helps further separate of functionality based on need, at an application or department level. The separation helps prevent various needs within the organization from overriding one another. This architecture is similar to a cloud computing or hosting provider architecture in which a multi-tenant solution is implemented through a combination of pADC with specialized or customer-specific vADCs deployed in a separate application delivery tier.

## Mobility

Mobility can be used to describe an attribute of both users and applications. In the case of the former, the challenge of application delivery is to identify and apply the right delivery strategy for each application based on user context— location, device, network conditions, and so on. In an environment where a pADC is employed, this contextual information can be shared across multiple application delivery function modules in order to ensure the proper access is allowed— or denied—and to apply appropriate acceleration and optimization policies based on that context. While a vADC would have the same capabilities, it is debatable whether the underlying physical hardware would provide the resources necessary to deploy the same set of functions in a single instance. While separating functionality through the



deployment of multiple vADCs is certainly an architecturally sound option, this separation necessarily hinders the sharing of context across functionality while simultaneously degrading performance by requiring many more connections or “hops” in the data path. Careful consideration must be made to determine what functions will require access to user context and which ones might not before it would be wise to move any functions to a separate, vADC instance.

In the case of applications, mobility refers to the act of moving an application across physically separate locations—for example, data center to data center to cloud computing provider. The use of a vADC is almost prescriptive in current cloud computing environments due to lack of customer access to underlying pADC capabilities. Applications that have come to rely upon pADC for acceleration, offloading, integrated network-scripting, and security will need to examine the viability of cloning policies and deploying in an application-specific “package” that can be migrated from one environment to another.

## Recommended Best Practices

The recommended best practice for a dynamic infrastructure is to deploy pADCs as key aggregation points to perform server and application offloading functions, to support high throughput application workloads, and for complex deployments using multiple advanced ADC functions such as application security, web acceleration, and access control. vADCs should be targeted for specific application workloads where the application requires more complex and compute intensive processing such as network-side scripting. In this case, dedicated processing for the particular workload is desirable to enable scalability through dedication of compute resources. These vADCs should be deployed in a tier behind the pADCs where offload functions would be performed.

Additionally, vADCs are best-suited for lab and QA environments to support rapid development, integration, and reduce time to market. In these environments, vADCs will further enhance collaboration between network, security, and application teams by making this technology more accessible to a broader audience without incurring the capital and management costs associated with a pADC.

It is best to consider the specific needs of the environment and weigh the pros and cons associated with deploying vADC or pADC and the advantages of employing a hybrid architecture that leverages both deployment form factors. To assist in making the determination when to use vADC or pADC, the following chart of pros/cons for each is provided:



## Virtual ADC Architectural Considerations

Pros	Description
Rapid deployment	As a software solution, a vADC can be provisioned and ready for inclusion in the development process much quicker than a physical appliance.
Financial efficiency for specific workloads	Because the cost of a physical appliance can be high relative to certain application types, use, and deployment scenarios, organizations sometimes have to choose between doing nothing and running application infrastructure sub-optimally. With a v ADC, cost can be charged more easily to a specific application workload and the vADC can be dedicated to that workload.
Failure isolation	In the event that the failure of a specific application configuration causes the failure of a physical device front-ending many applications, it will failover to the redundant unit. However, all applications could then be affected. By dedicating a vADC to specific application workloads better fault isolation is created.
Management	Being part of the hypervisor vendor's overall management framework can simplify the movement and management of the vADC. Coupling a vADC to specific applications makes it a more integral part of the overall ecosystem.
Cons	Description
High availability	The same degree of high availability achieved with a purpose-built pADC cannot be realized by commodity server hardware.
Security	Instead of a completely hardened system, a shared environment is used in which virtual appliance security is dependent upon the hypervisor vendor and the commodity server vendor.
Scalability	Certain high performance offload services do not have direct access to hardware. Commodity servers also lack purpose-built ASICs for offload. Both impact the scale and throughput of a vADC.



## Physical ADC Architectural Considerations

Pros	Description
High availability	pADC hardware designs are carrier-hardened for rapid failover and reliability. Redundant components (power, fans, RAID, and hardware watchdogs) and serial-based failover make for extremely high up-times and MBTF numbers. Commodity hardware of this type is costly and will not be integrated with the ADC software.
Security	Most pADC appliances and systems are security hardened and proprietary to the vendor. pADCs are not dependent on other vendors' security implementation or lack thereof. With hypervisors, there are known and potentially unknown vulnerabilities. To a certain extent, virtual appliance security is thus dependent upon the hypervisor vendor.
Scalability	Some pADCs have unique high-speed bridge and offload ASICs for such capabilities as high performance L4 processing, SSL, and compression, which enables them to be a cost-effective aggregation point for many applications or high-performance/throughput applications where latency matters a great deal.
Management	A pADC has special lights-out management capabilities so regardless of a physical device issue it can still be accessed, diagnosed, and fixed. Management can be less complex because the application delivery functions are centralized in a single device instead of distributed across the data center.
Cons	Description
Rapid deployment	Shipping a physical product, racking, stacking, and cabling takes time and adds cost to a deployment. It is also not well suited for agile development environments and QA labs.
Failure isolation	In the event that the failure of a specific application configuration causes a physical device front-ending many applications to fail, it will failover to the redundant unit. However, all applications can then be affected. Thus a combination of both physical and virtual ADC can simultaneously provide both failure isolation and scale.

## White Paper

Creating a Hybrid ADN Architecture with both Virtual and Physical ADCs

# Conclusion

The availability of virtual network appliances is certainly a step in the right direction, and vADCs are no exception. It is important to remember, however, that there are different technical challenges associated with the virtualization of network components than those encountered with virtualization of servers, and that these challenges are almost all architectural in nature.

There are many environments and uses within the organization for which vADCs will provide immediate benefits: testing, development, integration, QA, and staging of application delivery policies. Others, such as production networks upon which business and customers rely, may or may not be a good fit for vADC and are certainly not good candidates before any potential architectural issues are identified and addressed. Though many challenges will be shared across organizations and industries, many others will be unique to individual data centers based on specific needs, applications, and architectures currently in place.

<sup>1</sup> <http://www.rackspace.com/downloads/surveys/VirtualizationSurvey.pdf>

<sup>2</sup> <http://www.channelinsider.com/c/a/Virtualization/Virtualization-a-Driver-for-2010-Server-Refresh-360526/>

<sup>3</sup> [http://www.cio.com/article/168401/Virtualization\\_in\\_the\\_Enterprise\\_Survey\\_Your\\_Virtualized\\_State\\_in\\_2008?page=2&taxonomyId=3112](http://www.cio.com/article/168401/Virtualization_in_the_Enterprise_Survey_Your_Virtualized_State_in_2008?page=2&taxonomyId=3112)

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 [www.f5.com](http://www.f5.com)

F5 Networks, Inc.  
Corporate Headquarters  
[info@f5.com](mailto:info@f5.com)

F5 Networks  
Asia-Pacific  
[info.asia@f5.com](mailto:info.asia@f5.com)

F5 Networks Ltd.  
Europe/Middle-East/Africa  
[emeainfo@f5.com](mailto:emeainfo@f5.com)

F5 Networks  
Japan K.K.  
[f5j-info@f5.com](mailto:f5j-info@f5.com)

