

White Paper

Application Delivery as an Infrastructure Service

Guaranteed Performance in Shared Environments

APV Series Application Delivery Controllers

White Paper

APV Series I Application Delivery as an Infrastructure Service



Introduction	3
Application Delivery-as-a-Service	3
Figure 1: Typical Integrated Application Delivery Functions	4
Architecture	5
Platform Options	6
Figure 2: vAPV, APV and AVX Platform Options	6
Management	7
Figure 3: Cloud Management Integration Options	8
Return on Investment (ROI)	9
Conclusion	10

Introduction

The move to the cloud is in full swing; businesses are using cloud services to drive productivity, as gaming and social media startups are building businesses on cloud infrastructure and consumers are taking advantage of the cloud for anytime, anywhere access to content and services. For organizations that wish to forgo ownership of IT infrastructure and businesses that want a flexible and cost-effective approach to supplementing existing IT infrastructure, Infrastructure-as-a-Service (IaaS) has become a go-to solution.

In the beginning, IaaS providers offered mainly on-demand storage and servers; more recently, however, they are supplementing their offerings with high-value networking services. With the addition of cloud-based networking services such as load balancing, WAN optimization and secure access, IaaS providers can offer customers a full set of building blocks to ensure fast, available and secure delivery of applications and services.

To provide these networking services, IaaS providers need to make correct decisions when selecting underlying products and solutions. This paper sheds light on the attributes IaaS providers should look for when selecting application delivery solutions to power networking service offerings.

Application Delivery-as-a-Service

For the sake of simplicity, the scope and focus of this white paper will be limited to Load Balancing-as-a-Service (LBaaS). That said, principals described apply equally to other application networking services such as secure access and WAN optimization.

Many factors contribute to customer demand for load balancing services:

- **Best Practices** – Load balancing and application delivery have been a staple of network architectures for over a decade. Wherever there are Web and application servers, it is accepted best practice to front-end them with load balancers.
- **Availability** – By load balancing across a pool of servers, applications remain up and running in the event that one or more servers become unresponsive or are taken offline for maintenance.
- **Performance** – Load balancers distribute requests evenly to keep servers operating in their power band. What's more, they offload compute-intensive SSL processing and provide caching, compression and other acceleration technologies to improve application performance.
- **Security** – As a proxy for applications and services, load balancers provide a wall of separation between users and back-end servers. In addition, integrated Web and application security provide a first line of defense for external-facing Web operations.
- **Traffic Management** – Advanced load balancers have the ability to manipulate Layer-7 content to enable custom functionality for applications and services.




Figure 1: Typical Integrated Application Delivery Functions

Whether deployed in a traditional enterprise data center or in the cloud, the benefits of load balancing and application delivery are largely identical. The same cannot be said for deployment requirements. For cloud providers seeking to offer load balancing as an infrastructure service, the following criteria merit special consideration when selecting the load balancing and application delivery solutions that will power LBaaS offerings.

- **Architecture** – A sizable percentage of businesses moving to the cloud are Web based. Examples include online gaming, software-as-a-service (SaaS), eCommerce and social media. Companies with large external-facing Web operations require solutions with robust Layer-7 capabilities, a core competency in SSL (used to secure almost all traffic) and the ability to rapidly incorporate new features in response to continuous application development.
- **Platform Options** – Load balancing and application delivery solutions come in many shapes and sizes: dedicated hardware appliances, software virtual appliances and virtualized appliances. For IaaS providers, selecting a platform that supports all three form factors is critical to offering services to a broad cross-section of customer requirements including performance, reliability, on-demand agility and cost-efficiency.
- **Management** – Automation is critical to the IaaS provider's competitive advantage. What used to take hosting and managed service providers days or weeks to provision can be provisioned on-demand by an IaaS provider. To facilitate agility in service provisioning, application delivery solutions must be able to seamlessly integrate with any open source, 3rd party or proprietary cloud management system.

- **Return on Investment** – To maintain and maximize profitability, IaaS providers must align their capital and operational expenditures with customer demand. For virtual appliances, this means solutions that support pay-per-use. For dedicated and virtualized appliances, this means cost-effective hardware; it also means selecting solutions that support guaranteed performance for shared environments.

Architecture

Online gaming, SaaS, eCommerce, social media and other external-facing Web operations represent a large and growing segment of the IaaS user community. Consequently, it makes sense for IaaS providers to offer load balancing services built on the strongest possible architecture for Web application delivery. Three essential attributes to consider include:

- **Layer-7 Agility** – Unlike well-known enterprise applications such as Microsoft Exchange, there generally are no preexisting deployment guides for the gaming, SaaS, social media, eCommerce and other external-facing applications being deployed in the cloud. IaaS providers therefore need load balancing and application delivery solutions that offer robust Layer-7 functionality via a simplified approach that does not require complex, costly and time-consuming scripting.

Array's SpeedCore® architecture is built for Web application delivery. Leveraging an extensive built-in Layer-7 command library, customers can combine and nest functionality using CLI or WebUI commands to achieve custom Layer-7 functionality without scripting. In addition, for providers automating the service creation process, CLI and WebUI commands are far easier to integrate with cloud management systems as compared to custom scripting.

- **SSL Performance** – From SaaS applications that need to protect business data, to social media platforms that need to protect personal information, to eCommerce operations that need to secure online transactions, 2048-bit SSL is the current standard for Web-based encryption. For the IaaS provider and its customers, selecting load balancing and application delivery solutions with the highest level of SSL performance, scalability and security is essential.

Unlike load balancing and application delivery solutions that rely on OpenSSL, Array developed its own proprietary SSL stack. The result is an SSL implementation that achieves both superior performance and superior security. Leveraging the advantages of the SpeedCore architecture, Array SSL simultaneously achieves industry-leading scalability and full immunity from OpenSSL-related vulnerabilities such as Heartbleed and Man-in-the-Middle (MITM).

- **Software Architecture** – Developers working on external-facing Web applications sometimes operate in a vacuum, having no idea whether existing infrastructure is capable of supporting their latest features. Because Web-based applications undergo continuous development, new application delivery features may be required to support new features and new architectures at any time.

Array's SpeedCore is a pure software architecture that does not rely on ASICs or custom silicon. In the event that an IaaS provider needs to support a new application delivery feature to meet the needs of its customers, Array can build the new functionality into its next software release and enable the feature with full performance and scalability.

Platform Options

For many years, load balancers and application delivery controllers were available only as dedicated hardware appliances. Today, ADCs are available in physical, virtual and virtualized form factors. For the IaaS provider, these platform options translate to meeting the needs of a larger and more diverse base of customers. Therefore, it is in the best interest of IaaS providers to seek out load balancing and application delivery solutions that demonstrate proficiency for each potential use case.




Figure 2: vAPV, APV and AVX Platform Options

- **Dedicated physical ADCs** – The strengths of dedicated ADCs are performance and reliability. Where heavy lifting is required to handle high-volume traffic or where there is a need to support HTTPS (SSL) traffic, dedicated physical ADCs remain the most reliable, scalable and cost-efficient approach. Drawbacks of physical appliances include upfront costs and human factors involved in physical deployment, as well as space and power considerations associated with modern data center environments.

For IaaS customers with larger workloads and longer term application delivery requirements, Array's dedicated APV Series appliances deployed in the IaaS cloud deliver all the benefits of a traditional stand-alone ADC without the time and expense of enterprise owned and operated infrastructure.

- **Virtual ADCs** – The strengths of virtual ADCs are on-demand flexibility and maximum resource utilization. For workloads that are smaller or time-bound, virtual ADCs provide the ability to spin up resources almost instantly. For IaaS customers, resources are available at a moment's notice

and may be purchased in exact proportion to current requirements; for IaaS providers, virtual ADCs cost next to nothing to purchase and deploy. Drawbacks of virtual appliances include lack of scalability and diminished performance for HTTPS (SSL) traffic¹. In addition, as with many on-demand services, prolonged usage may reduce cost-efficiency.

For IaaS customers with small or time-bound workloads, Array's vAPV virtual appliances deliver all the features and functionality of a dedicated application delivery controller, but in a purely software form-factor. For IaaS providers, vAPV virtual appliances eliminate risk and up-front hardware costs and provide the flexibility to extract maximum utilization of virtualized data centers.

- **Multi-tenant ADCs** – Multi-tenancy aims to deliver the best attributes of both physical and virtual appliances. In essence, multi-tenant ADCs support multiple customers or application workloads on a single physical appliance. When implemented correctly, a multi-tenant ADC can deliver performance levels higher than that of virtual ADCs while at the same time reducing the overhead associated with racking and stacking individual appliances.

For IaaS customers with smaller workloads, a dedicated Array APV appliance with multi-virtual IP (VIP) support is the most cost-effective and efficient means to enable multi-tenancy. As compared to supporting multiple virtual ADCs, the multi-VIP approach does not incur hypervisor management overhead and requires only a single software image. Using multi-VIP support, a dedicated Array appliance may be divided into multiple virtual IPs to support up to 500 individual customers on a single platform, each with a guaranteed number of connections per second.

For IaaS customers with larger workloads, an Array AVX Series virtualized appliance may be divided into up to 32 fully independent vAPV instances, each with its own dedicated CPU, SSL cores, memory and I/O resources for ensuring guaranteed performance in a shared environment. As compared to the multi-VIP approach, multiple ADC software images are used, giving each customer the ability to run versions and configurations best suited to their individual requirements. Although fewer customers may be supported on each appliance, a far greater level of performance may be enabled for each customer. In addition, the AVX Series offers the flexibility to support vxAG secure access gateway instances, or 3rd party instances such as NGFW and IDS/IPS, in conjunction with or in place of vAPV instances to support secure remote/mobile/desktop access and BYOD.

Management

Automation and orchestration are two of the most important principles of cloud architecture. To fully utilize available resources and to be able to provision resources and services on demand, it is essential to implement an overarching cloud management system (CMS).

¹ Note, however, that Array also supports a hybrid virtual/dedicated SSL offloading option that can help assure performance while maintaining the flexibility of virtual appliances. [See our white paper](#) for more information.

There is no one-size-fits-all management system. Today, there are a number of different approaches to providing cloud management. Some providers create their own proprietary management system, others build their management system based on open-source platforms such as OpenStack, still others utilize systems such as VMware vRealize Orchestrator or Microsoft System Center. Ultimately, the choice of CMS will depend upon the goals of the IaaS cloud provider and the needs of their customers.

Regardless of the cloud management system that is in place, it is crucial that IaaS providers make sure that the load balancing and application delivery vendor they select has the ability to integrate with the provider's chosen CMS. To ensure that APV Series application delivery controllers are able to integrate with the widest possible range of cloud management systems, Array has developed a robust set of APIs and integrations with leading CMS providers. Our eCloud™ API enables management integration in each of the following potential management use cases.




Figure 3: Cloud Management Integration Options

- **Proprietary CMS** – For IaaS providers that have developed a proprietary management system, Array enables integration via either XML-RPC or the eCloud RESTful API. Using XML-RPC, cloud providers gain the most comprehensive control over Array application delivery controllers. Using the eCloud API, cloud providers gain a faster path to integration for a smaller subset of ADC functionality.
- **OpenStack** – For IaaS providers that have chosen to develop their management system based on OpenStack standards, Array has created a plug-in for the eCloud API that integrates with the OpenStack load balancing-as-a-service (LBaaS) API. Although limited to controlling functionality

as defined by the OpenStack community, utilizing the LBaaS interface is perhaps the fastest way to bring load balancing into a cloud management framework.

- **VMware & Microsoft** – For IaaS providers that have chosen to manage their cloud environment using VMware vRealize or Microsoft System Center, Array has developed plug-ins for bringing Array into orchestration systems such as vRealize Orchestrator (vRO). Utilizing vRO and similar systems, IaaS providers are able to create workflows that enable automated provisioning and configuration of Array application delivery controllers.

In addition to providing the means to integrate load balancing solutions with the cloud provider's CMS, it is also important to support the capability for remote management of infrastructure – both for customer and provider IT staff. To meet this requirement, Array's suite of application delivery solutions includes SSL VPN capability with the ability to partition customer and provider connections to enable access to infrastructure while maintaining separation and security.

Return on Investment (ROI)

For the cloud provider, any strategy that can reduce risk and up-front costs is invaluable. This is because significant investment must be undertaken before providers are able to offer services and realize a return on investment. For this reason it is critical that IaaS providers make the correct decision when it comes to the underlying infrastructure that will power their services. Hardware and software solutions must meet all technical requirements, but they must do so at the lowest possible cost. Where possible, they must support strategies that do not require the IaaS provider to pay until such time as their customers pay for services. In addition, infrastructure must support the ability to deliver guaranteed performance in a shared environment in order for SLAs to be offered to customers to enable billable services.

- **Lowest-Cost Enterprise & Service Provider Class ADCs** – Where hardware is required to provide services with superior performance and reliability, Array's APV Series application delivery controllers provide the highest levels of scalability and feature functionality at the lowest possible price point. Array's ADCs routinely cost 40% less as compared to other top-level ADCs and support SSL that costs up to 75% less per transactions per second (TPS).
- **Pay-Per-Use & Consumption-Based Billing** – Where virtual load balancing services are being deployed to meet customer demand for flexible, on-demand high availability, Array's vAPV virtual application delivery controllers allow IaaS providers to roll out services with minimal risk or up-front costs. Array's vAPV software is provided to IaaS providers free of charge, requiring only integration between the vAPV and the cloud provider's CMS. IaaS providers may bill customers based on concurrent connections or similar metrics, and pay Array after customers have been billed and in direct proportion to customer demand.
- **Guaranteed Performance in Shared Environments** – For some customers, "best effort" is good enough. However, for many customers, especially customers with mission-critical applications, SLAs will be requested. To achieve ROI, many IaaS providers must be able to service this latter

type of customer. Array's approach to multi-tenancy and shared environments places an emphasis on the ability to guarantee performance.

Where a single APV appliance and VIPs are used to support multiple customers, connections may be limited on a per-VIP basis to ensure each customer receives a guaranteed portion of overall system resources. Where a virtualized AVX Series appliance is used to support up to 32 fully independent vAPV instances on a single hardware appliance, each virtual appliance is assigned its own CPU, SSL, memory and I/O resources to ensure that each supports guaranteed performance metrics.

Conclusion

As an increasing number of businesses turn to the cloud to gain on-demand access to infrastructure and application services, IaaS providers must support a full range of server, storage and networking services. Within an IaaS service portfolio, there is increased demand for load balancing and application delivery offerings driven by the need to provide high availability for cloud-based applications and services.

For IaaS providers, selecting the correct platform to drive their load balancing and application delivery services is critical to competitive advantage and profitability. The chosen solution must support feature functionality aligned with customer requirements and it must be available in physical, virtual and virtualized form factors to strike the best balance of performance, reliability, flexibility and cost-efficiency. Cloud solutions must also provide the means to efficiently integrate with IaaS provider management systems for the purpose of automation and orchestration. Finally, solutions deployed to power service offerings must reduce risk and up-front costs and support guaranteed services to ensure ROI for the IaaS provider.

For IaaS providers seeking to offer load balancing services, Array leads the way. By leveraging distinct advantages of Array APV Series application delivery controllers, including modern features, platform options, management integration and price-performance, IaaS providers can offer services that meet wide-ranging customer requirements while maximizing competitive advantage, profitability and ROI.

White Paper

APV Series I Application Delivery as an Infrastructure Service

About Array Networks

Array Networks is a global leader in application delivery networking with over 5000 worldwide customer deployments. Powered by award-winning SpeedCore® software, Array application delivery, WAN optimization and secure access solutions are recognized by leading enterprise, service provider and public sector organizations for unmatched performance and total value of ownership. Array is poised to capitalize on explosive growth in the areas of mobile and cloud computing, analysts and thought leaders including Deloitte, IDC and Frost & Sullivan have recognized Array Networks for its technical innovation, operational excellence and market opportunity.

