WHITE PAPER–AUGUST 2019

Best Practices to Load Balancing on Microsoft Azure

Table of Contents

ABOUT THIS DOCUMENT

Enterprises moving their applications to the cloud want to make sure these applications are available, secure, and reliable. In addition to using Avi Networks for enterprise-grade load balancing that integrates out of the box with Microsoft Azure, there are numerous design and deployment considerations and best practices that can help enterprises get the most out of their cloud-based enterprise applications.

MICROSOFT AZURE CUSTOMERS NEED ENTERPRISE-GRADE LOAD BALANCING

Microsoft Azure is the fastest growing public cloud platform, attracting enterprises with its worldwide reach, support, security, and compliance benefits. While enterprises are adopting Azure to run enterprise applications in the cloud, they need an enterprise-grade load balancing solution to ensure applications are available, secure, and reliable, meeting the metrics defined in their service level agreements (SLAs) (see Table 1).

SLA	METRICS
Application Availability	• Percentage of time when the application is functioning (single VM/across Availability Sets/across Availability Zones)
Load Balancing	• Transactions per second • End-to-end latency • Response times
Security	• Secure Socket Layer (SSL) terminations • Open web application security project core role set (OWASP CRS) • SSL/transport layer security (TLS) certifications
Analytics & Monitoring	• Number of metrics tracked • Statistical algorithms used

Table 1: Common SLA Metrics for Enterprise Applications

Traditional appliance-based load balancers or their virtual editions don’t offer the automation and elasticity necessary to meet SLAs in the cloud. They are difficult to set up and maintain. And, while open source tools appear less expensive, they are not feature-complete, and setup and support are very “do it yourself,” which require a deep bench of engineers.

Avi Networks brings enterprise-grade, multi-cloud application services to Azure with the VMware NSX® Advanced Load Balancer™ (rebranded from the Avi Vantage Platform). NSX Advanced Load Balancer is an elastic, multi-cloud load balancing and web application security solution for Microsoft Azure with built-in application and analytics. The enterprise-grade, software-defined solution includes:

- A software load balancer
- An intelligent web application firewall (iWAF)
- An elastic service mesh for container-based applications.

ENTERPRISE-GRADE LOAD BALANCING USING SOFTWARE-DEFINED PRINCIPLES

NSX Advanced Load Balancer brings enterprise-grade load balancing to any cloud or data center through its multi-cloud, automation, and intelligence features. The software-defined solution uses a declarative model that allows the enterprise to focus on the desired business outcomes by specifying policies, freeing administrators from the repetitive and error-prone manual inputs of the past.

"Many of our customers choose Avi Networks for high performance."

"You deploy your applications for your business and you deploy load balancers to make sure applications are real, secure, and highly performant."

ASHISH SHAH,
SENIOR DIRECTOR OF PRODUCT,
AVI NETWORKS

Avi Networks is a Key Partner in Adobe's Adoption of Microsoft Azure.

One of Microsoft Azure's largest customers, Adobe, has chosen NSX Advanced Load Balancer to deploy both on- premise and in Azure. Reasons why Adobe chose Avi Networks as a key partner in its adoption of Azure include Avi Networks':

- Consistent, universal platform for both on-premise and public cloud.
- Enterprise-grade load balancer/ADC.
- Simplified operations with application visibility and performance insights.
- Elastic scale-out of load balancers and application servers based on real-time traffic data.

The NSX Advanced Load Balancer (see Figure 1) is built from the ground up using a distributed architecture, unlike monolithic legacy appliances. The architecture separates the data and control planes to deliver L4-L7 application services deployed across any environment – cloud or on-premise – with central management. Software-defined principles and 100% REST APIs take automation to the next level, enabling businesses to set the rules and thresholds so NSX Advanced Load Balancer can do the heavy lifting, spinning up additional service engines (SEs) – and shutting them down – on demand.

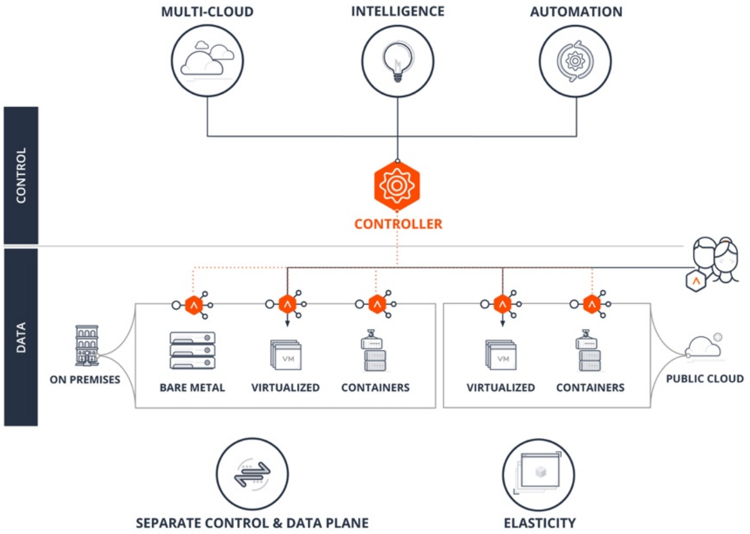

Figure 1: NSX Advanced Load Balancer Architecture for Application Delivery

AVI INTEGRATES OUT-OF-THE-BOX WITH MICROSOFT AZURE

NSX Advanced Load Balancer offers out-of-the-box integration with Microsoft Azure (see Figure 2), simplifying deployment of the elastic, enterprise-grade load balancing and web application security solution. NSX Advanced Load Balancer automatically integrates with objects within Azure.

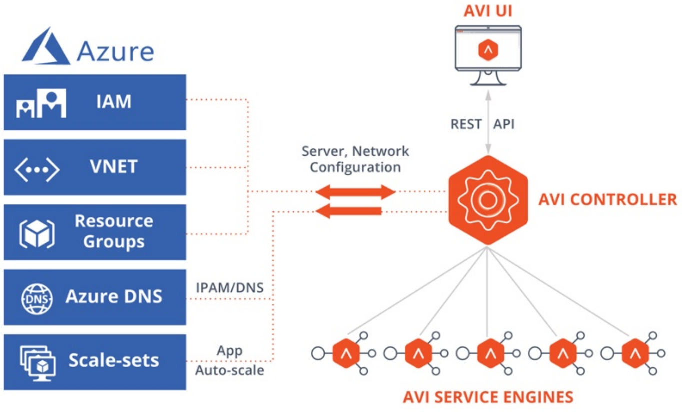

Figure 2: NSX Advanced Load Balancer and Microsoft Azure Integration

Not only is the integration easy, but it is less expensive than appliance-based solutions. NSX Advanced Load Balancer achieves high performance without pricey custom hardware, scaling to millions of SSL transactions per second.

NSX Advanced Load Balancer supports any application-level deployment decisions made by the enterprise, including deployment type, architecture, capacity sizing and elasticity, and high availability (HA).

DEEP DIVE INTO AVI AND AZURE INTEGRATION

Avi Controller communicates with the Azure components via standard Azure APIs to achieve the following (see Figure 3 for deployment details):

1. Subscription and Resource Group

The Avi Controller, when deployed on Azure, is within a subscription and resource group. In addition, the Azure cloud(s) configured in the Avi Controller are associated with a subscription and a particular resource group within the subscription. This resource group will be used to create Avi SEs and the load balancing infrastructure.

2. Service Engines

The Avi SE is the load balancing entity on which the virtual IP (VIP) is programmed. The SEs are automatically spawned by the Avi Controller when an application is configured. The entire lifecycle of the SE is automated.

The SE can be of varied instance types (bare metal, VMs, and containers), depending on your needs. Depending on your traffic and load requirements, the appropriate instance type can be configured. In addition, multiple SEs can be spawned (scale-out) on-demand to handle more incoming traffic.

3. Networking: VNets, NICs, and IP Addresses

Avi utilizes Azure Load Balancer internally to route incoming traffic to appropriate SEs. The management of the Azure Load Balancer is completely automated and taken care of by the Avi Controller. The user/operator is not required to configure or tweak any of these configurations.

The Avi Controller interacts with the following Azure networking constructs:

- VNets: An Azure cloud configured on the Avi Controller is associated with a particular VNet. This VNet is the one where the load balancing service is expected to be provided.
- In addition, the SEs communicate with the Avi Controller using a specified Management VNet (and subnet). This Management VNet could be different from the one containing the workloads.
- IP address for the SE’s Interface.
- IP address for the VS, called VIP. A public IP can be provisioned for the VS, in addition to the mandatory private IP.
- NICs: Avi automatically creates NICs and assigns IP addresses from the appropriate subnet, and attaches these to the requisite SE.

4. IPAM, Azure DNS, and Third-Party Integrations

IPAM for native Azure workloads is utilized by default. As the Avi Controller and SEs are powered up, they interact with the Azure Networking infrastructure to receive appropriate IP addresses.

Integration with various DNS systems (including Azure DNS) allows updating the fully qualified domain name (FQDN) of the

VS automatically, so that the private or public users can access the application via a known URL. When the VS is deleted, all resources including the DNS record are removed.

5. Azure Scale Sets

An Avi pool configuration can reference an Azure Scale Set (see Figure 3). When this is configured, Avi polls the Scale Set for membership changes, and synchronizes the Avi pool configuration accordingly.

In addition, the Avi Controller can take over the autoscale of the VM Scale Set, utilizing any of the hundreds of metrics and events to make a scale-out or scale-in decision regarding the application. See Table 2 for common Azure application deployment considerations.

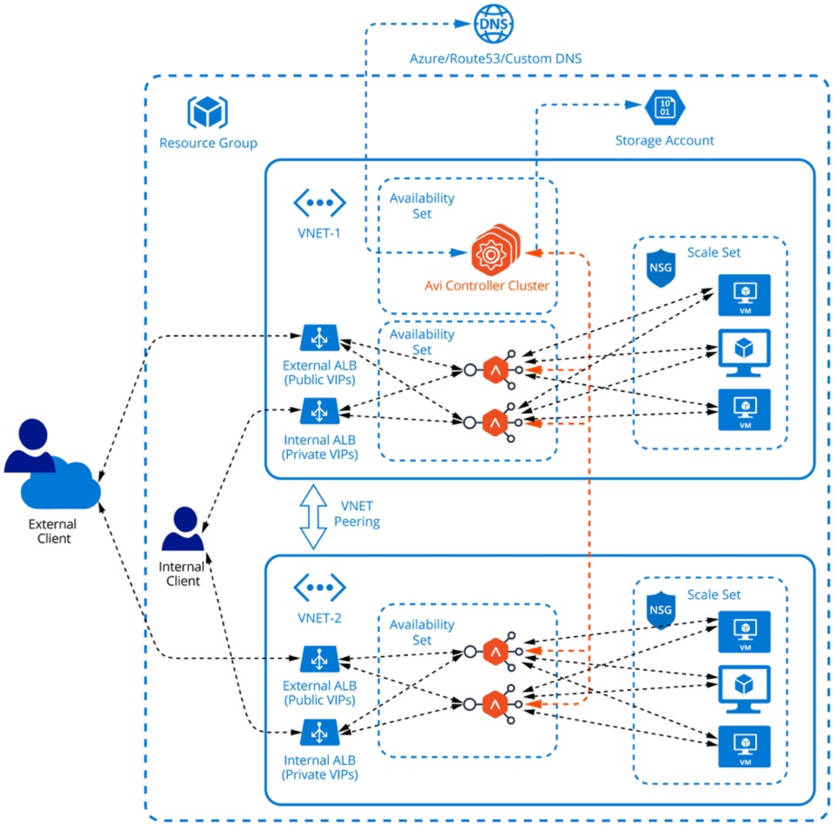

Figure 3: Deep Dive into Avi/Azure Integration

CONSIDERATION	QUESTIONS TO ASK
Cloud Only vs. Hybrid Cloud	• What are the connectivity requirements: Site-to-Site Virtual Private Network (VPN) or ExpressRoute? • What are the failover scenarios?
Application Architecture	• Is the architecture multi-tier? Is it a service mesh with micro-services? • Will security include SSL, WAF, and/or distributed denial of service (DDoS) capabilities? • How will the application be accessed? Will users be external, internal, or both? Will traffic come through the internet?
Capacity Sizing and Elasticity	• What are the traffic patterns and where does traffic come from? • Will virtual machine scale sets be used? • Is an application delivery controller (ADC) scale-out required?
High Availability	• Are availability sets being used for infrastructure? • Are availability zones required? • Is multi-region resilience necessary?

Table 2: Common Azure Application Deployment Considerations

AVI TAKES LOAD BALANCING AUTOMATION TO THE NEXT LEVEL

Deploying a legacy appliance load balancer is a complicated, time-consuming process. It begins with filing a support ticket for the load balancer and is followed by multi-step manual provisioning and operations to get the load balancer up and running. Avi's elastic platform takes load balancing automation and web application security to the next level, allowing for self-service load balancing with a real-time automated provisioning and operations (see Figure 4).

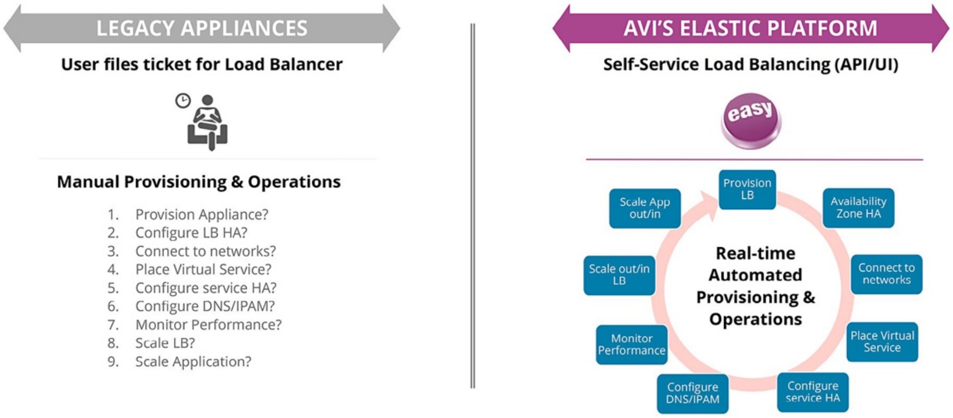

Figure 4: Avi's One-Click Elastic Load Balancing and Web Application Firewall for Azure

Avi's full-featured load balancer automatically integrates with Azure scale-sets and scales based on the policies set. With built-in analytics, NSX Advanced Load Balancer can add and remove service engines as traffic hits certain learned thresholds. The policies can also be set to limit access or perform rate limiting to specific applications and services based on the IP address, adding security in top of the load balancing.

BUILT-IN APPLICATION ANALYTICS BOOST TROUBLESHOOTING AND PERFORMANCE MONITORING

NSX Advanced Load Balancer is the only enterprise-grade load balancing solution that delivers built-in analytics. Analytics insights are used by the platform to make real-time load balancing and security decisions, as well as by administrators for troubleshooting and performance monitoring. Metrics from NSX Advanced Load Balancer include total round-trip time for each transaction, application health scores, errors, user statistics, and security insights (see Figures 5 and 6).

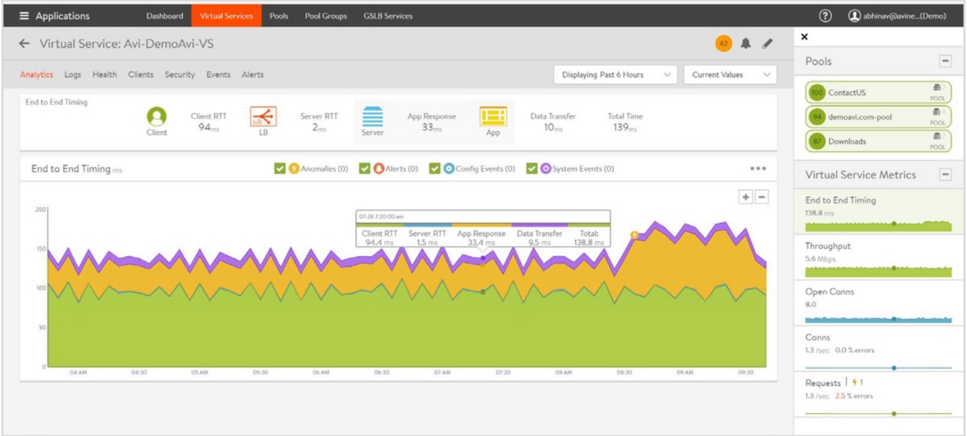

Figure 5: Latency Analytics Offers an Understanding of How Traffic is Served

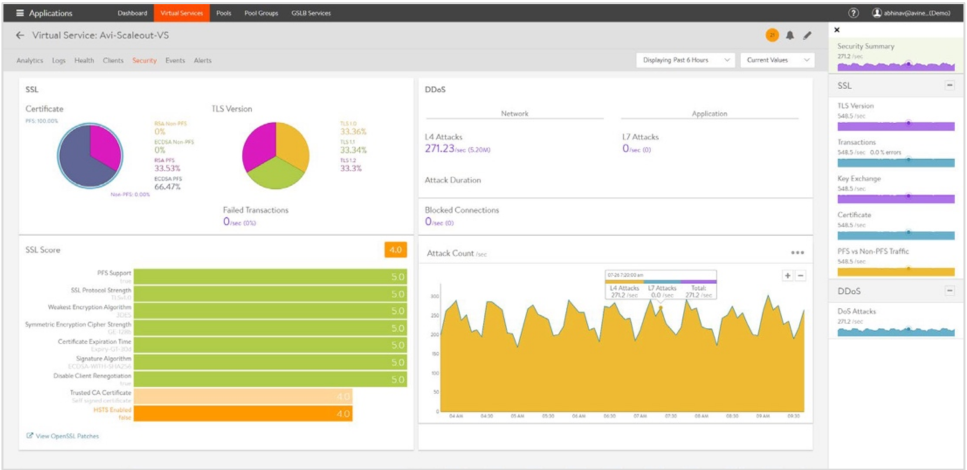

Figure 6: Security Analytics Offers a View into the Security Profile

BEST PRACTICES FOR LOAD BALANCING IN AZURE

Table 3 shows some considerations for designing and deploying an enterprise-grade load balancing service for applications hosted in Azure:

	DESCRIPTION	BENEFIT
Reference Architecture and Design Patterns	Start with a vetted *reference architecture* based on your use case and connectivity requirements.	The reference architectures provide a standardized design reference for your cloud environment. The load balancing solution should plug into this architecture.
Application Security	Enable client access via HTTPS only, for internal (intra-VNet/inter-VNet/VPN connected) as well as external (internet) clients.	HTTP is inherently insecure and prone to man-in-the-middle attacks and data theft. Using HTTPS and SSL for all client-virtual service (VS) communication ensures an optimal privacy and security posture.
	Allow Load balancer to terminate SSL within the private network/VNet. As a VNet provides a secure and isolated network partition, communication between the load balancer and the application servers can be via HTTP, unless there are regulatory or critical security requirements.	SSL processing on the application server increases load on them. Enabling SSL termination is a good balance of security and performance.
Resource Management	Enable auto-scale policies for the Load balancer and Application Servers, so that more resources are made available on-demand. As these resources incur cost on an hourly basis, ensure that the services are able to scale down when the load requirements subside.	Elasticity is a key tenet of cloud architectures, and it is imperative that this extend to application delivery/load balancing services as well. Benefits of an ADC supporting this include: • On-demand capacity management without requiring operator intervention • Cost savings and optimum capacity planning, without need to incur costs on under-utilized resources
On-Premise Application Migration to Azure	Existing on-premise applications have been built with various custom requirements. These include: • Modification of HTTP headers, addition of cookies, etc. during request/response processing by the ADC • Redirection of traffic to certain servers based on Client IP, URL, domain name, etc. • Ensuring SSL termination with approved ciphers Ensuring that the load balancer supports the required capabilities while providing an elastic architecture.	This ensures that migration to cloud provides the best of both worlds: • No compromise on enterprise-grade ADC feature set • Cloud-native deployment and consumption model along with elasticity
	Similar to the above, when hosting an application in multiple environments (e.g. Azure as well as on-premise), ensure that the ADC provides seamless operational support and integration in all environments.	This ensures that a single architecture can provide consistent application delivery across all your environments.

Table 3: Best Practices for Load Balancing in Azure

LEARN MORE

- Find out more about *Microsoft Azure Support in NSX Advanced Load Balancer*.
- Take a 2-hour *Test Drive of NSX Advanced Load Balancer*.
- Read more about the NSX Advanced Load Balancer solution for Microsoft Azure, best practices, deployment, and more information in the *Avi Networks Knowledge Base*.

CONCLUSION

NSX Advanced Load Balancer brings enterprise-grade, multi-cloud load balancing to Microsoft Azure, enabling enterprises to take advantage of Microsoft's cloud compute capabilities while meeting SLAs for applications, regardless of where they are deployed.

NSX Advanced Load Balancer offers:

- Out-of-the-Box Integration with Microsoft Azure that supports the different deployment decisions enterprises can make while providing high performance.
- Automation that simplifies deployment far beyond what legacy appliances offer and that takes load balancing to the next level with software-defined principles.
- Analytics that inform the platform in making real-time load balancing and security decisions, as well as that help enterprises track performance and address problem areas.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com.