VMware® Avi™ Load Balancer

Multi-Cloud Application Services: Load Balancing, Application Security, Container Ingress and Analytics

Key Benefits

- 97% Faster Service Provisioning
- Rapid Problem Resolution in seconds through app health scores, application analytics, security, and client insights
- 30% Reduction in TCO through on-demand application scaling and support for any bare metal server, VM, or container on-premises or in the cloud
- A single license

What's Included

A single platform that provides

- L4-L7 load balancing
- Web application firewall (WAF)
- Container ingress
- Global server load balancing (GSLB)
- Real-time application analytics
- On-demand application autoscaling in a customer-managed or SaaS way

Application Services Accelerate Business Agility

Modern enterprises need an on-demand, fast-to-deploy, easy-to-use app delivery solution that facilitates multi-cloud consistency across on-premises and cloud environments. Multi-cloud and microservices are driving business IT agility. Legacy infrastructure lacks the elasticity, flexibility, and agility needed to deliver applications securely and reliably. The rise of containers, APIs, and observability needs presents an opportunity for infrastructure to become composable, automated, and intelligent without the limitations of the appliance-based approach.

Platform Capabilities

The VMware Avi Load Balancer is a software-defined architecture that separates the central control plane (Controller) from the distributed data plane (Service Engines). Avi Load Balancer has comprehensive REST APIs, making it fully automatable and seamless with the CI/CD pipeline for application delivery. Avi is a unified platform designed to deliver the business IT needs that are required in today's world of digital transformation. It provides elasticity, security and is operationally easy to manage. Avi Load Balancer scales out applications on demand and detects failures for a fault-tolerant self-healing application infrastructure.

These functions can be automated for a hands-off operational management model through a closed-loop monitoring process. Advanced analytics/ observability optimize the application delivery and protect them along with their data with context-aware application and API security. Security policies are kept current through live threat updates via Avi Load Balancer with Cloud Services which includes the Cloud Console.

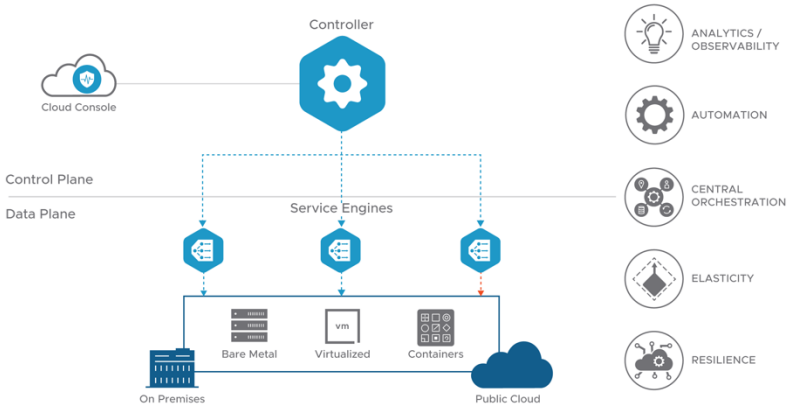

vmware® by Broadcom

Local and Global Load Balancing

VMware Avi Advanced Load Balancer uses the Controller as the “brain” of the entire system and acts as a single point of intelligence, management, and control across a distributed fabric of enterprise-grade load balancing, application security, container ingress and analytics. The Controller provides decision automation based on closed-loop telemetries and presents actionable insights based on based on application monitoring, end-to-end timing, searchable traffic logs, security insights, log insights, client insights, and more. The Cloud Console, also delivers an always-on, as-a-service consumption model for operational capabilities such as central licensing, security feeds, and proactive support. See Figure 1.

App and API Security

For Web Application and API Protection (WAAP), Avi Load Balancer features a web application firewall (WAF), bot management and API protection with a distributed web application security fabric. Customers can enforce security through closed-loop analytics and application learning mode that covers OWASP CRS protection, support compliance regulations such as PCI DSS, HIPAA, and GDPR, and signature-based detection. The WAF provides an optimized security pipeline with a positive security model to maximize the efficiency of resource-intensive operations. Cloud Console provides live feeds of new threat updates including IP reputation, bot detection, signatures, and more, and automatically minimize false positives with advanced security analytics, detection, and enforcement modes. With real-time app security insights and analytics provide actionable insights on performance, end-users and security events in a single dashboard with end-to-end visibility.
See Figure 2.

Key Features

- Point-and-click simplicity for security policies with central control
- Elastic scale with high performing, load based automatic scale-out architecture
- Granular security insights on traffic flows and rule matches for precise policies
- Automated threat updates through Cloud Services
- Real-time app security insights and analytics
- Protects applications from DDoS attacks and OWASP Top 10 threats

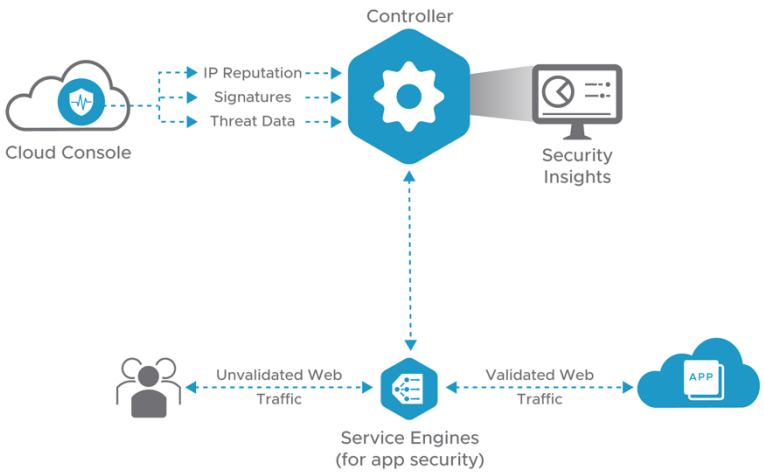

Kubernetes Ingress Services

Modern application architectures based on microservices have made appliance-based load balancing solutions obsolete. Containerized applications deployed in Kubernetes clusters need a scalable and enterprise-class solution for load balancing, global and local traffic management, service discovery, monitoring/analytics, and security. However, this should not be done in a disparate way with siloed DIY products to be stitched together all by the platform teams. Enterprises adopting Kubernetes need a cloud-native approach for traffic management and application networking services. For modern container-based applications, Avi Load Balancer offers a consolidated set of container services including cloud-native, scalable, enterprise-class container ingress traffic management, dynamic service discovery, and security. See Figure 3

Key Features

Traffic Management & Service Discovery

- Local and global load balancing
- DNS / IPAM / Circuit Breaking
- Health Monitoring
- TLS termination, Cert management / automation
- CI/CD and Blue-Green / Canary deployments

Security & Observability

- WAF
- Authentication
- Allowlist / Denylist
- Rate Limiting
- DOS detection / mitigation
- Application and infra performance metrics
- Transaction tracing & fine-grained logging

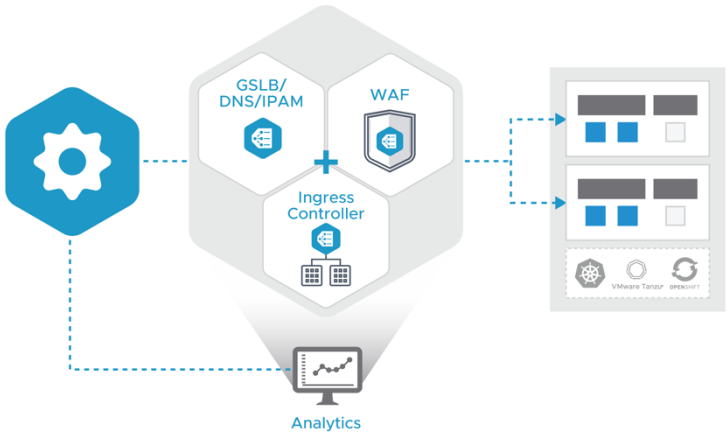

FIGURE 3: Kubernetes Ingress Services

Avi Use Cases and Ecosystems

Universal Solution

Avi Load Balancer is designed to work for different scenarios. It can provide a platform for multi-cloud environments to offer a single management console to deal with the various environments. Web application security is delivered as a critical component of the solution to protect the applications and data. The technology is designed to work for traditional application and container microservices alike.

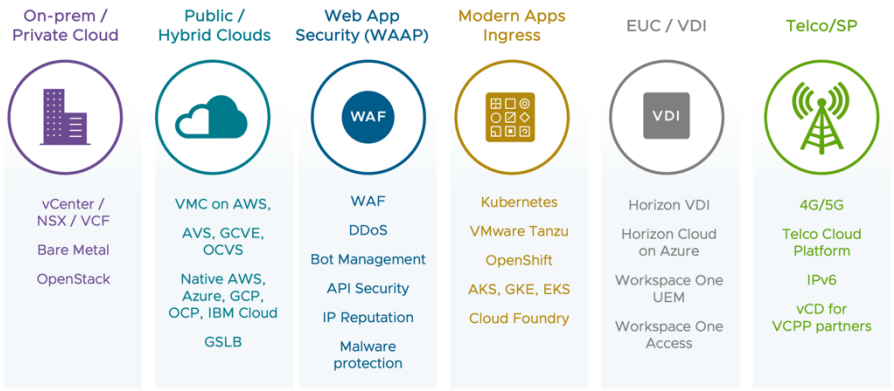

The solution integrates with VMware solutions including VMware Cloud Foundation, vCenter, Aria, Horizon VDI, Workspace One, NSX and vSphere, to name a few. For telco environments, The Avi Load Balancer is designed to integrate with the VMware Telco Cloud Platform (TCP) with support for IPv6 and telco-specific requirements.

AGILITY FROM AUTOMATED PROVISIONING AND SELF-SERVICE

Automated virtual service provisioning with per-app load balancing services

- Application provisioning in seconds
- Full automation with REST APIs to support faster application rollout in Blue/Green and Canary deployments and enable DevOps teams with self-service portals
- Simplified operations with centralized policies

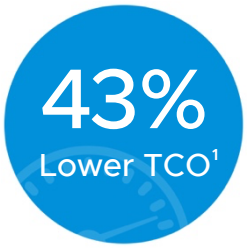

LOWER COST WITH SIMPLIFIED OPERATIONS

Elastic load balancing and just-right-size capacity without overprovisioning

- Flexible, subscription- based licensing model that eliminates static capacity
- Reduced OpEx through simplified operations of central management
- Consistent application services across multi-cloud environments without reconfiguration

RAPID RESOLUTION IN SECONDS

Near real-time visibility into network transactions to troubleshoot quickly

- Application health score for a quick snapshot of network posture
- End-to-end round-trip times with latencies between each hop
- Real-time logging, recording and replaying traffic and security events

[1] IDC Business Value Study of VMware NSX Advanced Load Balancer: A Study of Enterprises Using Next-Generation Application Delivery

VMWARE INTEGRATIONS		
vCenter	Google Cloud VMware Engine	Aria Automation (vRealize Automation)
VMware NSX	Oracle Cloud VMware Service	Aria Automation Orchestrator (vRealize Orchestrator)
VMware Cloud on AWS	VMware Tanzu	Aria Operations (vRealize Operations)
VMware Cloud Foundation (VCF)	VMware Horizon	Aria Operations for Networks (vRealize Network Insight)
Azure VMware Solution	vCloud Director (vCD)	Aria Operations for Logs (vRealize Log Insights)

3rd PARTY INTEGRATIONS	NOTES
OpenStack	Queens, Rocky, Stein, RH OSP, Keystone v3
Bare Metal	RHEL, CentOS, Ubuntu, Oracle Enterprise Linux
Public Cloud	Microsoft Azure, Amazon Web Services (AWS), Google Cloud Platform (GCP), IBM Cloud, Oracle Cloud
Container	Kubernetes, Tanzu, Rancher, OpenShift, Amazon EKS, AKS, GKE

3rd PARTY SUPPORTED PLATFORMS	NOTES
Automation	Ansible, Terraform, Python/Java/Go SDKs, vRO plugin
Analytics / Monitoring	Splunk, Cisco Tetration, Cisco AppDynamics, Graphite, Datadog, Logstash, Elasticsearch, InfluxDB, Syslog, Prometheus, Zabbix
IPAM / DNS	DNS, Azure DNS, Azure DNS Private Zones, AWS Route 53, Infoblox, Custom DNS integration, Custom IPAM Integration

PERFORMANCE - OBSERVED ON A SINGLE SERVICE ENGINE		
	Baremetal Server (24 Core) with xl710 (40 Gbps)	Service Engine running as vCenter VM (6C / 6GB)
Max SSL (EC) Connections	50K per second	12K per second
Max SSL (RSA 2K) Connections	18K per second	4000 per second
Max HTTP Requests	700K per second	185K per second
Max L4 TCP Connections	400K per second	130K per second
Max concurrent connections		
Max SSL throughput	38 Gbps	10 Gbps
Max tenants (shared data plane)	Unlimited	Unlimited
Max tenants (isolated data plane)	200	200
Max Service Engines per cluster	200	200

CATEGORY	FEATURE
Enterprise-class load balancing	TLS 1.3 support, SSL termination, default gateway, GSLB, DNS, wildcard VIP and other L4-L7 services
Multi-cloud load balancing	Intelligent traffic routing across multiple sites and across private or public clouds, global server load balancing supported with Canary upgrades of leader and follower sites
Application performance monitoring	Monitor performance and record and replay network events with granular logging
Predictive autoscaling	Application and load balancer scaling based on real-time traffic patterns
Cloud connectors	VMware, SDN controllers, OpenStack, AWS, GCP, Azure, Linux Server Cloud, VMware Cloud on AWS, Google Cloud VMware Engine, Azure VMware Solution (customer-managed)
Distributed application security fabric	Granular app insights from distributed service proxies to secure web apps in real time
Application security	Bot management (tech preview), Positive security model and learning mode for WAF
SSO / client authentication	SAML 2.0 authentication and authorization for back-end HTTP applications
Automation and programmability	REST API based solution for accelerated application delivery, extending automation from networking to developers with self-service portal enabled
Application analytics	Real-time telemetry from a distributed load balancing fabric that delivers millions of data points in real time
Centralized management and upgrade	Policy-based management and ability to selectively upgrade data plane with Flexible Upgrade
Networking protocols support	BGP, RHI and ECMP, BFD, IPv6, VLAN & trunking, VRF awareness, Radius and SIP
Consolidated container services	Kubernetes Services including ingress, WAF, GSLB, DNS/IPAM on a scalable platform with support for multi-cluster, multi-site and multi-AZ container clusters
Central License and Visibility Platform	Controls all Cloud Services Licensing as well as providing comprehensive Global and Controller dashboarding from a centralized cloud service

vmware® by Broadcom

Item No: vmw-bc-ds-temp-uslet-word-2024 1/24