WHITE PAPER: June 2025

Plug-and-Play Load Balancing for VMware Cloud Foundation with VMware® Avi™ Load Balancer

vmware® by Broadcom

Table of Contents

VMware Cloud Foundation

VMware Cloud Foundation (VCF) simplifies self-service private cloud with a consistent operating, governance, and consumption model. It provides quick time to value by enabling self-service consumption and delivery of private cloud resources, providing users and developers with a unified and consistent self-service layer. VCF customers can take advantage of the self-service private cloud with solutions to enable use cases and help drive increased business and IT agility, productivity, and efficiency.

VMware Avi Load Balancer

For those who have realized the value of VCF, why not extend the same principles you have come to believe into application delivery? With its genesis in the software-defined principles and backed by VMware's validated designs, guided by simplified operations, self-service automation, holistic visibility, and overarching security that complements the network security, VMware Avi Load Balancer (Avi) offers a superior plug and play load balancing solution like none other.

Executive Summary

Plug-and-Play Load Balancing with a Cloud Operating Model

Businesses today are under immense pressure to deliver applications rapidly, driving a fundamental shift in how IT infrastructure and operations are designed and managed. Adoption of the cloud operating model—a framework that brings agility, self-service, automation, and elasticity to all environments is the foundation of this transformation.

Infrastructure platforms have rapidly evolved too: from the hardware-defined era to the software-defined cloud era and now accelerating into an AI-defined era. This latest phase is propelled by the rise of AI applications, including agentic workloads, which demand intelligent, scalable, and highly automated infrastructure solutions.

As business needs change, not all workloads can be handled the same way. Compliance requirements, integration complexities, and cost considerations often dictate the choice of the environment. Yet, users and developers expect a seamless, on-demand experience regardless of where their workloads reside. The cloud operating model addresses this by abstracting underlying services and enabling resource delivery through self-service portals or APIs, while IT maintains control, compliance, and optimal use of the infrastructure.

Traditional IT operational models are no longer sufficient to meet these evolving demands. To stay competitive, IT organizations must embrace:

- Self-service access to resources, available anytime and anywhere
- Elastic scalability, including auto-scaling and auto-healing capabilities
- Consumption-based resource management for cost efficiency
- Automated lifecycle management to streamline operations

Successfully adopting the cloud operating model also requires a shift in organizational mindset. IT teams need to become more customer-centric, develop multidisciplinary skills, and collaborate across functions to deliver value quickly and efficiently.

Foundational services such as load balancing must also evolve. Modern load balancing solutions are designed to accelerate resource consumption via automation, provide elasticity to handle dynamic workloads, and utilize application latency analytics for rapid issue detection and resolution—ensuring optimal performance, especially for AI-driven and next-generation applications.

By fully embracing the cloud operating model, IT organizations can deliver resources at the speed of business, enabling rapid application development and deployment while maintaining robust control and operational excellence. This strategic evolution is essential for organizations aiming to succeed in the AI-defined era.

Why Load Balancing is Critical for VCF

Every Application Needs Load Balancing

Enterprises are facing an unprecedented need to provide infrastructure that matches the agility of ever-changing applications. The network team is constantly pressured to support a much higher velocity of changes and updates to meet time-to-market considerations. Delivering great application experience is synonymous with business growth, better customer satisfaction, and workforce productivity. Organizations are realizing the need for a modern, software-defined load-balancing solution to enable a better application experience. Load balancing is a fundamental computing, storage, and network building block. Without load balancing, one cannot realize the true value of applications and digital transformation and deliver applications that business need to succeed.

Avi Enables the Cloud Operating Model

VMware Avi Load Balancer (Avi) is designed to provide flexible deployment and simplified operations through automation and centralized management and support the Cloud Operating Model strategy. Avi decouples the control plane from the data plane with a software-defined approach to application delivery infrastructure. It provides consistent application delivery services across clouds and heterogeneous infrastructure, including bare metal servers, virtual machines (VMs), and containers. Unlike legacy load balancers, which are highly overprovisioned and wasteful, Avi enables optimal capacity management and highly elastic fabric through automatic zero-touch autoscaling features. Yet another disadvantage of the legacy load balancers is the inability to troubleshoot issues in time, often resulting in friction between the application and the network teams. Avi makes it easy to troubleshoot complex problems through rich and contextual analytics and stops the blame game. By completely integrating with the VCF ecosystem, you can enjoy the benefits of unified automation workflows, validated designs, and visibility for complete peace of mind, knowing your applications will be delivered as intended.

Avi Makes VCF Deployment Better

Customers continue to extract tremendous value from the VCF solution, keeping in line with the principles of the cloud operating model. However, the last-mile challenges with application delivery remain. These are compounded by geographically disparate data centers with broader attack surfaces and reduced analytics and visibility. This is where Avi adds tremendous value to the existing VCF customers. The solution provides plug and play integration with VCF to deliver load balancing at the speed of applications. Customers can leverage built-in automation to deploy load balancing as code as part of DevOps workflows for rapid infrastructure roll-out. In addition to excellent server load balancing capabilities, the Avi enables Global Server Load Balancing with Intelligent Traffic distribution across Data centers and easily addresses essential use cases such as Disaster Recovery. With the inbuilt WAF, Avi can even provide application protection before traffic reaches the network. Lastly, it proactively analyzes and optimizes the end-to-end application experience while simplifying troubleshooting and automatically building a self-healing infrastructure.

Here are a few ways Avi makes VCF investments better:

Enterprise Scale and Resiliency

As the mandate for application modernization takes effect, it becomes increasingly necessary to load balance applications across geographies and clouds for high availability and address critical issues such as disaster recovery and application upkeep. Avi provides global server load balancing (GSLB) services, which provide load balancing of applications across multiple geographically dispersed locations while providing centralized GSLB configuration, application monitoring, and analytics. That includes centralized provisioning with automated discovery of applications across sites and centralized application monitoring, logs, and analytics. In addition, integration with VCF ensures load balancing as a service is available with VCF Cloud Management and identity services like LDAP, Radius, and DNS servers for better scale and resiliency.

Cloud Experience for DevOps

Avi improves the cloud experience for DevOps in VMware Cloud Foundation (VCF) by providing a software-defined, plug-and-play load balancing solution that streamlines deployment and operations across traditional and modern workloads. Unlike legacy hardware-based load balancers, Avi reduces latency by eliminating hair-pinning of traffic, a common cause of sluggish application performance, and further accelerates response times by offloading resource-intensive tasks such as connection

management and Access Control List (ACL) lookups from routers and switches. Its deep integration with VCF enables automation throughout the entire lifecycle—from initial deployment to configuration and ongoing management—allowing DevOps teams to consume load balancing as code and through self-service catalogs. This empowers developers to rapidly onboard applications, embed load balancing into CI/CD pipelines, and manage resources programmatically, all while IT maintains centralized control and oversight. With built-in analytics and end-to-end visibility, Avi also enables proactive resource optimization and troubleshooting, ensuring consistent, high-performance application delivery in a modern cloud operating model

Comprehensive Web App Security

In recent years, web application security has become increasingly important, especially as web application attacks are the most common reason for breaches. In fact, starting March 2025, WAF is a mandatory requirement for public facing web applications per PCI 6.4 compliance. WAFs is indeed a critical component of web application security, guarding against web application vulnerabilities while allowing customization of the security rules for each application. As WAF is in line with traffic, some functions are conveniently implemented by a load balancer. This is an additional layer of protection at Layer 7 before the traffic reaches your network. The WAF can protect your applications from common threats, including Cross-site Scripting (XSS), SQL injection, cookie poisoning, Layer 7 DoS, and Web scraping. Knowing that time to deploy and operational efficiency are critical to a successful security strategy is essential. Avi continuously provides consistent security policies and holistic visibility while reducing complexity and cost, thanks to its unique integration with VCF Networking by NSX.

IDC Business Value Study of Avi

IDC interviewed organizations using the Avi to deploy application services to understand how they are using the platform to support their business operations. IDC's whitepaper results show that study participants are achieving significant business value. Learn how Avi helps enterprises achieve*:

- 27% higher application developer productivity
- 90% faster to scale capacity
- 54% fewer outages
- 43% lower cost of operations

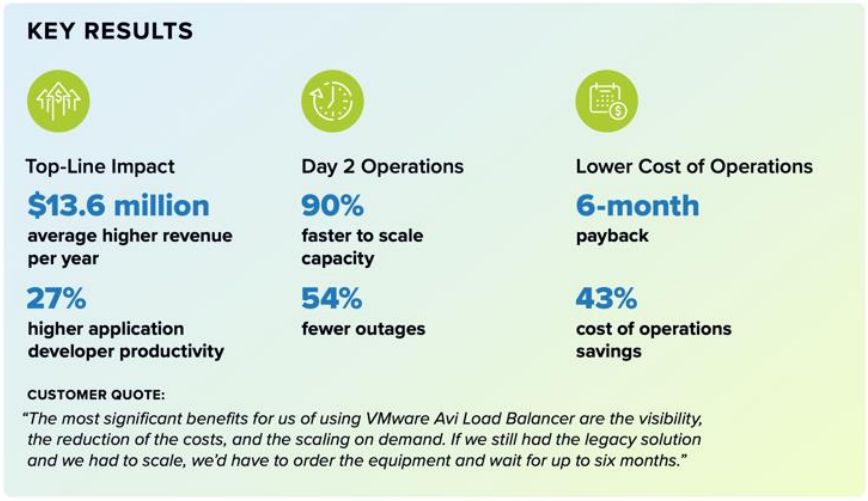

SOURCE: "IDC Business Value White Paper, sponsored by VMware by Broadcom February 2024 | IDC #USUS51843824"

Avi and VCF Use Cases

Avi's software-defined architecture makes it the perfect candidate for an infrastructure-agnostic platform for any application – VMs, bare metal, or containers. The abstraction allows the same APIs, workflows, and automation to offer consistent load-balancing services regardless of the underlying environments.

Integrated Experience with Full Visibility for VCF

Customers demand better application experience round the clock with complete end-to-end analytics and automation from Layer 2 through 7. This requires full-stack integration with VCF and the backing of rigorously tested, validated, and supported deployment designs. Deploying Avi considerably reduces operational and training efforts, and this integrated experience with VCF ensures load balancers are never an afterthought nor deployed without self-service built-in for DevOps. Avi is the only load balancer offering true plug-and-play load balancing for VCF workloads through VCF Operations integration. This enables automated lifecycle management for the Avi load balancing infrastructure including the Avi controller.

Disaster Recovery across Multiple Sites with GSLB

Enterprise workloads are often mission-critical, demanding rigorous business continuity plans and application resilience. However, DR plans can be quite expensive with a standby site that is always idle. This translates to double the licensing cost and complex testing required to ensure this high-availability and failover process works in times of emergency. Avi's software-defined architecture and on-demand auto-scaling of capacity dramatically simplifies the DR solution. Avi Cloud Console monitors and manages a pool of active licenses deployed across multiple sites – on-prem or cloud. With GSLB set up, activating a standby site is as easy as reassigning licenses from one Avi Controller to another, eliminating the need for idle capacity and unused licenses. Maintaining consistent policies in the Avi Controller and ensuring data plane traffic is not impacted during the switchover is much more simplified.

Web Application Security for VCF with WAF

Security is only as good as its weakest point. With VCF Networking by NSX offering complete network security from L2-L3. Avi WAF enhances the protection for your workloads by offering excellent L4-7 web app security built into the platform without separate licenses. Avi provides comprehensive protection for applications hosted in private cloud environments and ensures compliance with GDPR, HIPAA, and PCI-DSS. WAF significantly enhances security by defending against sophisticated web-based threats, such as SQL injection, cross-site scripting (XSS), and other vulnerabilities in the OWASP Top 10. Key advantages of WAF with VCF include:

- Advanced Threat Mitigation: WAF inspects HTTP/S traffic in real-time, blocking malicious requests while ensuring legitimate traffic is unaffected.
- Customizable Security Policies: Administrators can create tailored rules to meet specific application requirements, ensuring maximum protection and flexibility.
- Automated Deployment: WAF deployment can be automated through VCF Operations, enabling rapid provisioning and consistent configuration across environments.
- Centralized Management: Security policies, monitoring, and updates are managed centrally, simplifying operations and reducing administrative overhead.
- Synergies with VMware vDefend: WAF compliments vDefend products, enhancing application security alongside existing lateral security like micro-segmentation and distributed firewalls.

Container Ingress with Integrated Security

Container security with Avi ensures robust protection for modern, containerized applications while delivering high-performance traffic management. Avi acts as a Kubernetes Ingress Controller, managing ingress and securing inter-cluster traffic via GSLB. It includes a WAF to protect against web-based threats like SQL injection and XSS, while API security safeguards microservices. Avi supports TLS/SSL offloading, automates certificate management, and enables secure communication. With centralized visibility, it offers actionable insights into application performance and security metrics. Avi integrates seamlessly with DevOps workflows, ensuring security within CI/CD pipelines. Avi provides a scalable and resilient solution for securing containerized workloads by uniting advanced security features with simplified operations. Avi offers a comprehensive and consolidated set of Layer 4 to Layer 7 services for any Kubernetes environment. Avi also seamlessly integrates with Gateway API to deliver enterprise-grade application services for Kubernetes environments. The Avi controller interacts with the Gateway API to automate traffic routing, handle dynamic and support advanced networking policies, all while leveraging its software-defined distributed architecture for scalability and elasticity.

vSphere Kubernetes Service (VKS)

VKS is enterprise-grade, CNCF certified and conformant Kubernetes runtime, which is available with VCF. VKS was designed with enterprise functionality in mind, including an easier installation process, self-service, and the continued inclusion of components that practitioners of Kubernetes have come to expect. This gives customers unprecedented abilities to run enterprise-grade Kubernetes on-prem in a more consumable and scalable fashion.

Avi combined with VKS on VCF delivers the fastest path to production-ready Kubernetes clusters by consolidating L4-7 container networking services on a single scalable platform. This integration provides local and global traffic management, Kubernetes ingress, WAF, and container ingress, all tightly integrated with VCF for consistent security and performance across VMs and containers. Avi's cloud operating model enables automated lifecycle management, self-service consumption, and multi-tenancy, simplifying operations and accelerating application delivery. Together, Avi and VKS on VCF unify container networking and security, enhancing scalability, agility, and resilience for modern Kubernetes workloads.

Avi and VCF Integration Advantages

Unlike the legacy load balancers that are either hardware-based or derivatives of the hardware-based solutions, Avi's software-defined load balancer brings elastic scale, robust performance, and intelligent analytics to every data center and cloud. Customers get enterprise-grade L4-L7 features with self-service automation and rich analytics while solving the significant operational challenges of appliance-based hardware load balancers. Avi offers the following distinct advantages as compared to legacy load balancers:

One Platform vs. Disjointed Point Products

Avi and VCF integration deliver a unified platform that streamlines operations and reduces complexity compared to managing separate point products. This approach simplifies automation, security, and analytics across workloads, minimizing operational overhead. Organizations benefit from easier scaling, faster deployments, and end-to-end visibility. Overall, the plug and play approach to load balancing enhances agility and reliability while lowering costs versus disjointed solutions.

Controller vs. Instance Manager

Avi's software-defined architecture separates the control, and data planes and delivers load balancing and WAF as an elastic fabric that auto-scales based on real-time traffic.

Decision Automation vs. Task Scripting

Built-in automation coupled with an API-based approach makes intelligent decisions, scales, and simplifies hybrid cloud application deployments. This is supported by Closed-loop analytics that helps automate decisions with over -eight hundred unique application metrics.

Optimal Resource Utilization vs. Over-provisioning

Avi's elastic, fabric-based approach with active-active high availability (HA) ensures optimal resource utilization by dynamically scaling capacity to match demand, minimizing unutilized infrastructure. This contrasts with traditional over-provisioning, where excess resources are reserved as contingency leading to inefficiency and higher costs.

Built-in Analytics vs. Bolt-on Solutions

Compared to a bolt-on implementation by legacy load balancers, analytics-first architecture with built-in visibility ensures complete end-to-end L2-7 visibility with VCF integration. Thanks to the VCF and Avi integration by mapping apps to hosts, networks, and flows, Avi breaks visibility silos with holistic L2 to L7 visibility and analytics. This provides contextual Insights for faster Root Cause Analysis, thus reducing MTTR and exceeding customer SLA and expectations.

Avi and VCF Integration Details: Plug and Play Throughout Application Lifecycle

Avi delivers a true plug-and-play experience for VCF by tightly integrating with VCF Operations automating the deployment and lifecycle management of load-balancing infrastructure. Infrastructure administrators use VCF Operations to deploy and manage Avi Controllers and Service Engines, ensuring consistent, resilient, and scalable load-balancing across VCF workload domains. Once the load-balancing infrastructure is in place, DevOps teams can self-service the creation of virtual services directly from a catalog, accelerating application deployment without needing deep load-balancer expertise. Application owners benefit from multitenancy features, allowing them to create organizations and manage their own isolated environments within the shared infrastructure, supporting secure, scalable, and efficient operations for multiple teams or business units. This division of responsibilities where infrastructure admins manage the infrastructure, DevOps provision services, and app owners handle tenancy for lines of business—streamlines operations, enhances agility, and maintains enterprise-grade security and visibility throughout the application lifecycle.

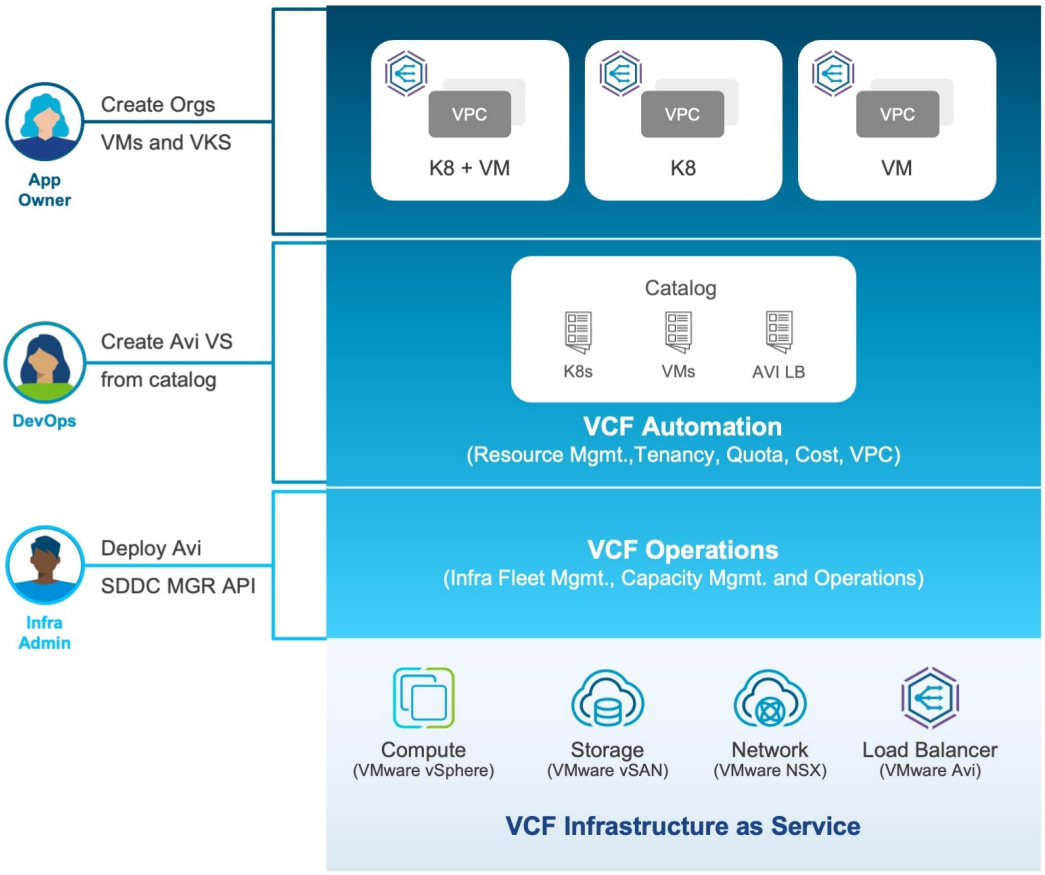

In this section, we will examine capabilities that VCF customers can leverage out of the box to accelerate the adoption of the Cloud Operating model throughout the application lifecycle: build | operate | consume | protect.

Build and Operate through VCF Operations Integration

Organizations adopting private cloud solutions aim to cut costs and simplify operations by using a single declarative API across virtual infrastructure and containers alike to increase agility. They need an integrated platform with a unified operating model, which VCF provides. Avi is the only load balancer offering true plug-and-play load balancing for VCF workloads through VCF Operations integration. This enables automated lifecycle management or the Avi load balancing infrastructure including the Avi controller. Here are a few out of the box capabilities that VCF customers can expect as they build and operate their private cloud.

Infrastructure readiness through native integrations: VCF admins can opt to import vSphere clusters to VCF domains or start with new VCF workloads, then deploy Avi via VCF Operations into the VCF infrastructure. Through the integrations with vCenter and NSX, Avi auto-configures and auto-discovers infrastructure objects and network inventory including virtual IP allocation, DNS registration and routing table updates. Because of this integration, Avi understands the environment and deploys Service Engines that drives us application availability and enhances resiliency and scalability.

Unmatched Application visibility: Avi speeds up troubleshooting of applications deployed in VCF with industry-leading application visibility. It provides rich telemetric insights to the application owners, DevOps teams and the VCF Cloud admins, including how the applications are performing and how to prevent web app attacks with anomaly detection.

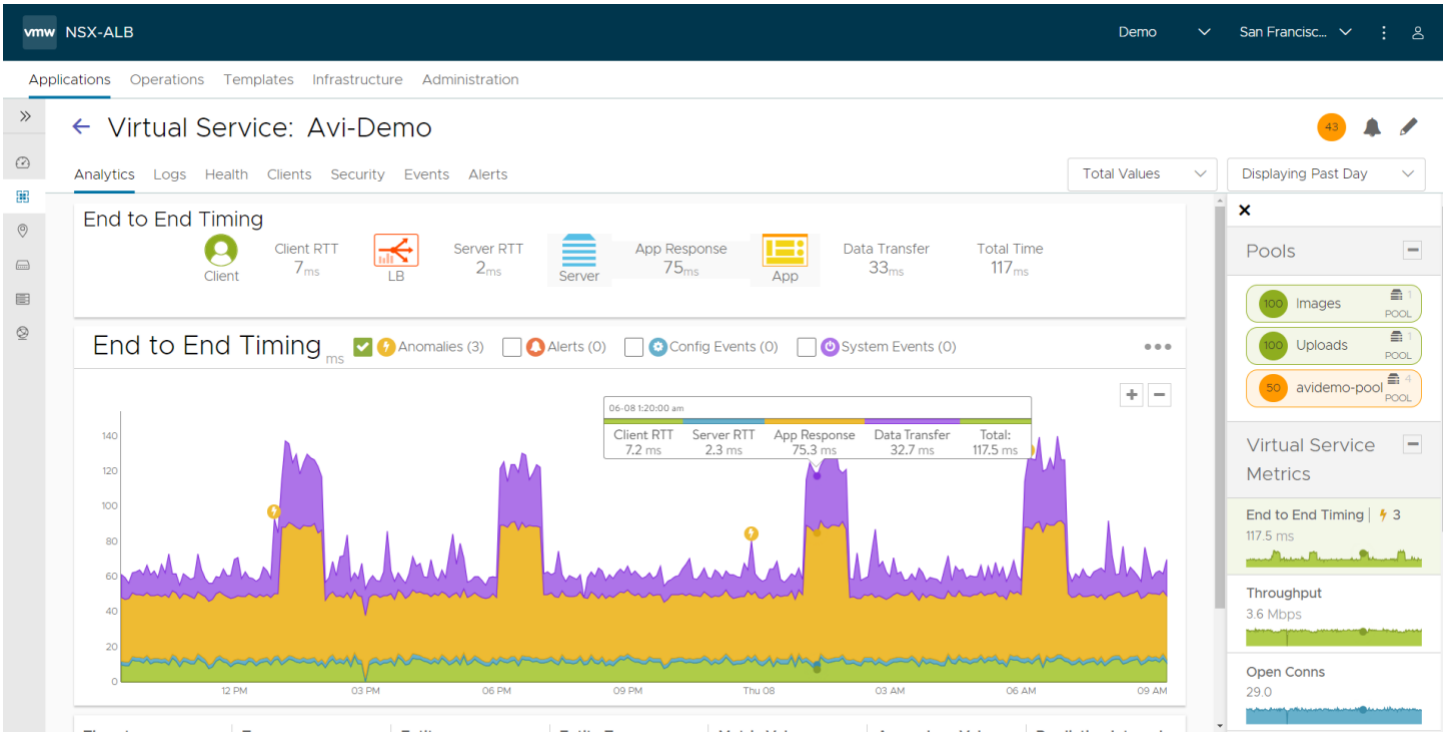

Avi streamlines troubleshooting by providing comprehensive end-to-end visibility through key metrics like Client RTT, Server RTT, and application response time in a unified dashboard. Easy to follow intuitive color-coded Health Scores offer a holistic view of application and network health by integrating metrics such as CPU and memory utilization and license usage, enabling proactive infrastructure management. Powered by AI and machine learning, Avi prioritizes critical logs and categorizes common issues like HTTP errors and high latency, helping IT teams focus on key problems and resolve them faster. Additionally, Avi correlates performance metrics with configuration changes in real time, allowing rapid root cause identification, minimizing downtime, and improving operational efficiency. This integrated approach reduces guesswork and accelerates problem resolution for IT teams.

Automated Lifecycle Management: VCF provides automated lifecycle management for Avi which minimizes downtime and reduces the risk of errors. The integration of Avi with the VCF operations simplifies mundane yet complex operations such as password or certificate updates and enables version compatibility. It takes the operational complexities out of the hands of VCF admins who may not have load balancing expertise, yet responsible for delivering a cloud experience for load balancing services. This integration also offers existing Avi customers an easy path to migrate from vSphere to VCF and marks a significant shift from conventional virtualization to a true private cloud operating model offered by VCF. The advantages of automation, consistent operations, and integrated lifecycle management make it an attractive progression. The ability to preserve existing networking and storage design offers a practical and seamless migration path that minimizes disruption to ongoing business activities.

Consume Avi Load Balancing with VCF Automation

The Avi integration with VCF provides Cloud admins and developers one consistent experience to consume load balancing as code and as self-service. This enables a cloud-like consumption model for load balancing, empowering developers and application teams to deploy and manage load balancers rapidly and securely without deep load balancer expertise, while IT retains complete control and oversight. Here are a few highlights on how cloud administrators can consume Load Balancing as a plug and play solution in VCF:

Self-Service Catalog: VCF Automation provides a service catalog where DevOps or app admins curate load balancing service offerings in Blueprint tailored to different teams or users based on roles, org, projects, and regions, with clearly defined policies. Avi is available as a catalog tile with easy self-service consumption and faster application onboarding for both virtualized and container workloads.

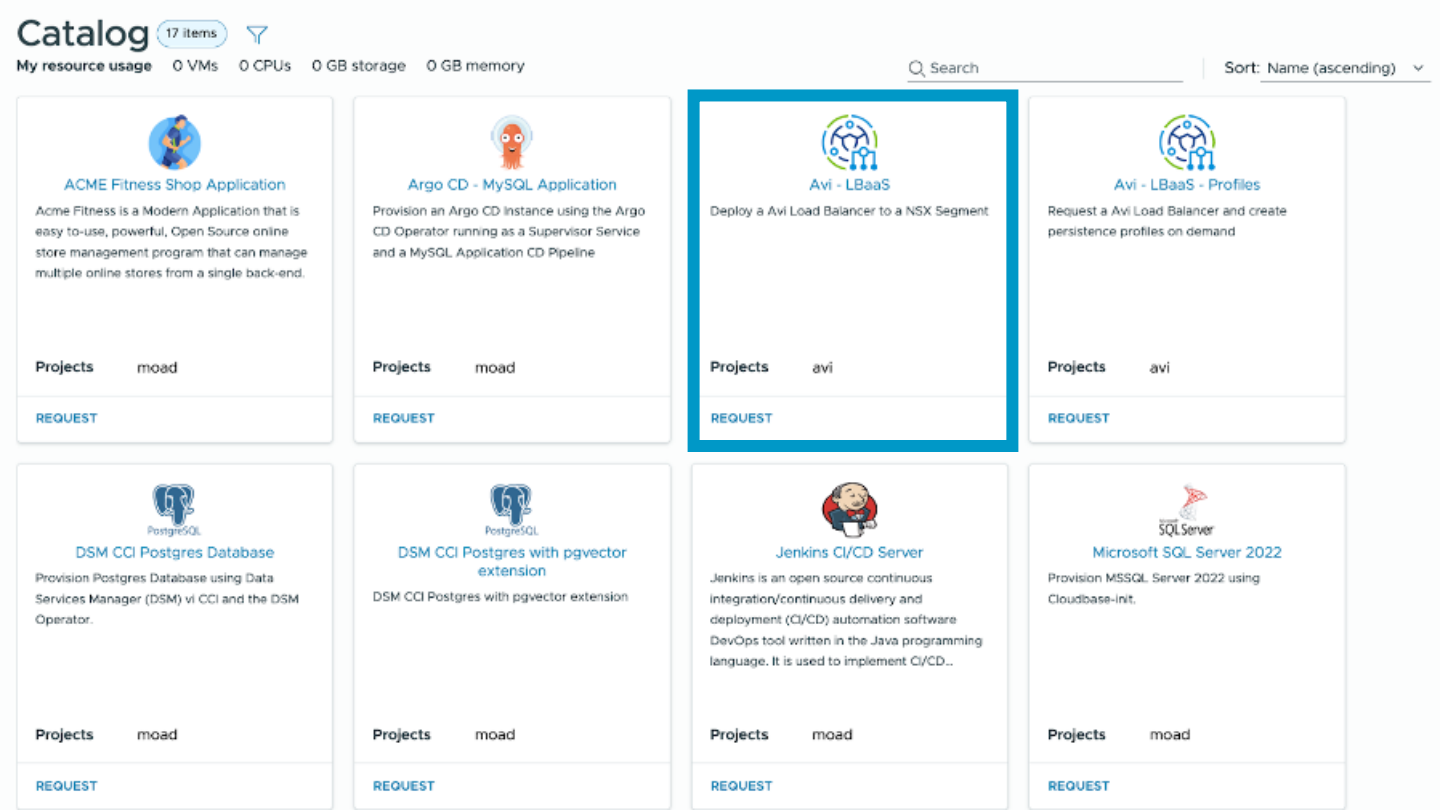

Load Balancing as Code: VCF Automation supports YAML-based Infrastructure as Code (IaC) templates that define load balancing infrastructure and configurations, enabling repeatable and version-controlled deployments. Alternatively, Avi provides plugins for automation tools like Ansible and Terraform, allowing customers to consume load balancing as code programmatically via APIs and automation workflows. This enables DevOps teams to embed load balancing provisioning into their CI/CD pipelines and infrastructure automation scripts.

Multi-Tenancy and VPC: Avi provides robust multitenancy and Virtual Private Cloud (VPC) support within VMware Cloud Foundation (VCF) through deep automation and native integration. Avi's architecture enables infrastructure teams to provision and manage load balancing on a per-tenant basis, allowing each tenant to have isolated environments with dedicated resources, policies, and quotas. The role-based access control (RBAC) enables tenants to self-manage their segmented infrastructure, including creating and managing their own virtual services, while infrastructure owners retain oversight and can enforce utilization policies. VCF automation further streamlines this process, enabling one-click deployment of Avi load balancing as a service, seamless integration with VCF workload domains, and support for native VPC constructs—delivering resource isolation, scalability, and self-service capabilities for DevOps and application teams within a unified cloud operating mode.

Protect VCF Workloads Against Web Attacks with Avi Web Application and API Security

Avi is fully integrated with VCF, providing one secure plug and play application resilient platform for integrated local and global load balancing, container ingress, web application and API security. The integration provides a centralized and scalable platform for managing load balancing and WAF policies across multiple VCF domains, simplifying infrastructure management and improving application security. Let us look at few key highlights in terms of granular security for virtual, containers and bare metal workloads alike:

Granular Web App Security with Avi WAF: Starting March 2025, WAF is a mandatory requirement for public facing web applications per PCI 6.4 compliance. Avi provides a centralized and scalable platform for managing web application security policies across multiple VCF workload domains. Avi offers WAF to protect against OWASP Top 10 attacks but also a comprehensive set of security capabilities including, DDoS protection, rate limiting and bot detection.

DDoS & BOT Detection for defense against advanced threats: Avi WAF offers defense against DDoS and bot attacks by analyzing traffic at the application layer, beyond traditional firewalls. Cloud Console provides live feeds of new threat updates including IP reputation, bot detection, signatures, and more, and automatically minimizes 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.

Kubernetes Ingress Security: Avi offers both ingress controller and Gateway API, a next-gen Kubernetes ingress with standardized and native deployment to K8s. For modern container-based applications, Avi offers a consolidated set of container services including cloud-native, scalable, enterprise-class container ingress traffic management and security. The bottom line is you don't need a separate solution for container workloads and Avi provides one platform approach to all applications.

Summary

Avi is a key component of the VCF ecosystem, enabling the benefits businesses want to achieve from cloud architectures. Operational support of clouds is complicated, and Avi with VCF simplifies the operations and expertise required to manage the cloud and its applications. Application delivery in the cloud requires a change in how people, processes, and technology are aligned, and the tight integrations of Avi and VCF improve and simplify this alignment.

Through enhanced automation, improved security, and application visibility, Avi provides plug-and-play load balancing for VCF, helping adopt the mindset required to embrace the Cloud Operating model businesses are trying to achieve.

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