

# VMware Cloud™ on AWS - Reference Architecture

Leveraging Amazon Application Load Balancer (ALB) with VMware Cloud on AWS

This reference architecture details how Amazon Application Load Balancers can be used to load balance secured web traffic destined for web servers hosted on VMware Cloud on AWS infrastructure. The Three-tier architecture shown here leverages AWS native services such as Amazon Route 53, Amazon CloudFront, Web Application Firewall, AWS Shield, AWS Certificate Manager, Amazon S3, Amazon RDS, AWS CloudTrail and Amazon CloudWatch.

- 1 Deploy and configure web and application servers in VMware Cloud on AWS. Deploy and configure **Amazon RDS MSSQL Database server** in Highly Available Multi-AZ setup that serves the DB requests to web and app servers residing on VMware Cloud on AWS.
- 2 Setup **VPN connections and/or Direct Connect connections** with private VIF so that on-prem users and VMware Cloud admins can access resources on the SDDC as WAN private networks. Ensure routing is configured appropriately.
- 3 Setup **Compute Gateway** with appropriate firewall rules to route the web/app traffic "to and from" web and application servers and resources from Customer VPC. Setup **Management Gateway** with appropriate firewall rules to route all the administrative traffic to and from Management Appliances/VMs. Configure appropriate security group rules for ALB, RDS and other resources in Customer VPC.
- 4 Configure **Application Load Balancer (ALB)** for web servers to load balance and serve web traffic by using Target Group that are configured to use IP addresses of web servers hosted on VMC on AWS. Setup ALB and web servers to use appropriate SSL certificates so all the data in transit are encrypted and secured.
- 5 Configure a second **Application Load Balancer (ALB)** for application servers to load balance and serve application traffic coming from web servers, by using Target Group that are configured to use IP addresses of application servers hosted on VMC on AWS.
- 6 Configure **Amazon S3** – as an origin for static contents like video/audio media files, manuals etc so it can serve the static contents to Amazon CloudFront. It can also store the log files generated by CloudFront and web and application servers hosted on VMC on AWS.
- 7 Configure **VPC endpoint for S3** in customer VPC so the web and application servers on VMC on AWS can leverage low latency and high bandwidth connections provided by Cross VPC ENI when accessing S3.
- 8 Configure **Amazon CloudFront** with appropriate origin servers and behaviours. Setup CloudFront to fetch static contents from S3 and dynamic contents from Application Load Balancer.
- 9 Integrate AWS edge services, so **Route 53** record set points to CloudFront distribution, with Shield as DDoS detection & mitigation, WAF for L7 traffic firewalling/whitelist/blacklist. Configure **AWS CloudTrail** for collection of relevant logs about user activities on AWS resources and **Amazon CloudWatch** for monitoring native AWS resources.
- 10 **VMware Cloud Admin** manages/administers VMC on AWS resources over the internet connection or from On-Premises networks using VPN or Direct Connect connections.
- 11 **End users** access highly available and optimized web sites hosted on VMC on AWS, that is well integrated with multiple AWS Native services and VMC on AWS resources.

