VMware Cloud on AWS: Network Security

VMware General

Table of contents

VMware Cloud on AWS: Network Security

Overview

Understanding the network security model of an SDDC is a critical part of designing and managing a VMware Cloud on AWS solution. While this model isn’t complicated, there are a few concepts that may be unfamiliar to individuals who are only accustomed to managing traditional hardware firewalls and security.

The following sections are designed to provide the fundamental knowledge required to successfully design and deploy a network security policy for an SDDC.

The Network Security Model

Network security within an SDDC is enforced by 2 different types of firewalls: the NSX gateway firewall and the NSX distributed firewall (DFW).

Figure 1 - Gateway and Distributed Firewall

These firewalls are designed to address network security policy enforcement in 2 different ways, with the gateway firewalls enforcing policy at the network border at tier-0 or tier-1 routers for north-south traffic and DFW enforcing policy within the SDDC for east-west traffic protection. Keep the following points in mind with the NSX firewalls:

1. As with most firewalls, NSX firewall rules are evaluated top-to-bottom with the first matching rule being applied to a new connection.
2. The firewall policy is applied based on the condition where new connections are evaluated in the direction specified by the rule. The direction here means that response traffic for that same flow is allowed, It doesn't mean that new sessions initiated in the opposite directions are also allowed.
3. The firewalls are stateful. This means that once a session is established, traffic, in either direction, that matches that flow is only permitted while new sessions initiated in the opposite directions are denied. For instance, with TCP protocol there is a formal session negotiation process. For other traffic types firewall evaluation that do not have a state are based on 5-tuple matching and timers.
4. NSX uses group and service definitions as part of firewall rule creation. This provides the ability to dynamically apply rules to objects, and the NSX FW automatically & dynamically translates those objects into the IP addresses associated with those objects. With this, It is operationally easier to manage than a huge list of IP addresses.

Network security within an SDDC is configured from the Security tab of the NSX UI console. The following sections will provide more details on the concepts used within each firewalls.

The Gateway Firewalls

Gateway firewalls may seem like the traditional centralized firewalls where security is enforced at the perimeter of the network for north-south traffic. Similarly within the SDDC, gateway firewalls are enforced on the NSX edge routers and are responsible for enforcing network security policy at the border of their respective networks. There are two distinct types of networks within the SDDC: the management network and the compute networks.

The management network is part of the SDDC infrastructure and provides network connectivity to the various infrastructure components of the SDDC. Within these, the appliance subnet is used by the vCenter, and NSX appliances in the SDDC. Other VMware appliance-based services added to the SDDC also connect to this subnet. The infrastructure subnet is also utilized by the ESXi hosts in the SDDC. Due to the architecture of the SDDC, the IP address space is fixed for this network and its layout may not be altered and no customer workloads can be deployed into the management gateway segments.

The compute networks are used by the compute workloads VMs of the SDDC. Customers have the ability to add and remove network segments within the compute networks as needed. The compute networks can either be attached to default compute gateway or custom tier-1 gateways.

Figure 2 - Management, Compute and Custom Tier-1 Gateway Firewall

It is important to understand the following points regarding the gateway firewalls:

- The default security policies of the management and compute gateway firewalls are “default deny” while for custom tier-1 gateways are “default allow”.
- There are 3 points of enforcement for the gateway firewalls: on the tier-0 edge router, the management gateway, and custom tier-1 user created gateways.

Since there is a default 'Deny All' for management and compute gateway, the security administrator must create additional rules to explicitly allow the traffic that should be permitted through the firewalls. Keeping in mind that the firewall rule set is applied based on the direction of initiation, the security administrator must define a permitted rule with only required source and destination details. Since the firewalls are stateful, responses for permitted traffic will always be permitted. Regarding the point of enforcement, it is important to understand how and why the gateway firewalls are applied as they are.

The entry point for the entire SDDC is its tier-0 edge router. This is the first point of enforcement for the gateway firewall, and as a result, this edge device has the potential to protect all networks within the SDDC. However, a design decision excludes the management network of the SDDC from the gateway firewall of the tier-0 edge. This means that the gateway firewall of the tier-0 edge only protects the compute networks of the SDDC. This design decision makes sense when the management gateway (which borders the management network) has its own gateway firewall. Since the management gateway is already protecting the management network, the protection of the management network at the tier-0 edge would be redundant.

Then “Why enable the gateway firewall on the management gateway?". Since it is desirable to protect the management network from not only the external world but also from the compute network, enforcement of security policy at the border of the management network is necessary. By enabling the gateway firewall on the management gateway, the security administrator can protect the management network from both external and internal (compute) networks.

From the perspective of the end-user, the NSX UI console abstracts details of where these policies are applied and presents a simplified view that displays rulesets from entities: labeled as Management Gateway, Compute Gateway, and tier-1 Gateways. In actuality, the default compute gateway policies are enforced on the tier-0 router unlike management policies and tier-1 gateways which are enforced on the respective tier-1 gateway uplinks.

Gateway firewall policies can be applied through the NSX UI->Security->Gateway Firewall that is accessible through the “Open NSX manager” option available on VMware Cloud on AWS SDDC console.

Management Gateway Firewall

Management gateway firewall policies can be applied to access management resources. By default all inbound access is denied to management resources. Management gateway firewall rules specify actions to take on network traffic based on the source and destination addresses, and service.

- Either the source or destination must be a system-defined inventory group.
- When sources is a system-defined group, services is “Any”.
- List of available ports and services is limited and managed by VMware and dependent on the type of management resource. For example, to access Vcenter Server appliance from user defined source, services available to access are only

HTTPS, ICMP and SSO.

- When any add-on VMware application is added to the SDDC, the system-defined inventory group also gets created.

Refer Add or Modify Management Gateway Firewall Rules for more information.

Compute Gateway Firewall

Based on the traffic type and security policy requirements, compute gateway firewall policies can be applied to one or multiple available interfaces uplink available through the UI parameter option “Applied To” as displayed in Figure 3.

Figure 3 - Compute Gateway Firewall Uplinks

Interface uplinks types and of what kind of traffic can be controlled through each of them is as follows:

- VPC Interface - Network traffic between VMs in the SDDC and native AWS instances and services with network addresses in the Connected VPC
- Internet Interface - Traffic over the internet and policy-based VPN over public IP. Default SNAT rules for compute workloads and dedicated NAT rules management workloads are in place over this interface. Customers can request public IPs for their SDDC and create custom NAT rules for SNAT and/or DNAT. If the default route goes through a Direct Connect, VPN, or VTGW connection or has been added as a static route to a VPC, NAT rules run for inbound traffic but not for outbound traffic, creating an asymmetric path that leaves the VM unreachable at its public IP address. When the default route is advertised from the on-premises environment, you must configure NAT rules on the on-premises network, using the on-premises Internet connection and public IPs.
- Direct connect Interface - For Direct connect, SDDC grouping with traffic over VMware Transit Connect (vTGW) (with on-premises, SDDCS and native VPC connected over vTGW), and policy-based VPN over private IP.
- VPN tunnel Interface - Traffic specific to tier-0 route-based VPN.

Note that all uplinks selection does not include the VPN tunnel interface and have to be selected explicitly for traffic over route-based VPN.

Custom Tier-1 Gateway Firewall

For user created custom tier-1 gateways, the default firewall rule is to allow all, meaning firewall functionality is effectively disabled by default. To implement tier-1 gateway rules, the security administrator must specifically construct deny rules with either “drop” or “reject” option.

For custom tier-1 traffic leaving the SDDC, it must be allowed by both the custom tier-1 firewall and the default compute gateway firewall.

The Distributed Firewall

The distributed firewall (DFW) is different from the centralized firewalls such that it is not enforced in a centralized location (as with the gateway firewalls) but rather enforced at the vNIC of each VM in the compute networks. DFW is also where the Layer-7 and Distributed IDS/IPS security are also enforced as mentioned in a later section of this article. By enforcing security policy at the absolute edge of the network, it becomes possible to manage network security policy in ways that are difficult to replicate in a traditional data center network.

The distributed firewall is available from the Security tab of the NSX manager UI view as displayed in Figure 4.

Figure 4 - Distributed Firewall

The ruleset for DFW is organized around the concept of sections. As highlighted in figure-4 highlighted section 1, there are 4 pre-created sections for DFW. These are conveniences that have been added to the UI to direct the security administrator into the good practice of organizing the rules of the security policy. The key thing to remember is that sections are an organizational tool only. Rules within NSX firewalls are evaluated top-to-bottom independently of sections. This means that rules in the Emergency Rules section will be evaluated before rules in the sections below it (and so on). Keep this in mind particularly when creating “deny” or “reject” rules.

The following behaviors are possible for the DFW security policy:

1. default permit (default option) – This means that traffic is permitted unless specifically blocked by a “deny” rule.
2. default deny – This means that all traffic is denied unless specifically allowed by a “permit” rule.

Key Takeaways:

- DFW defaults to permit behavior, meaning it is effectively disabled by default.
- DFW is applied only to the compute networks.
- DFW may filter network traffic both north-south and east-west.To prevent confusion with the dual layers of gateway and distributed firewalls, the default security policy of DFW is set such that it is effectively disabled. To implement DFW, the security administrator must specifically construct deny rules with either “drop” or “reject” option.

The purpose of DFW is to enable the security administrator to construct security policies that may be applied within the compute network itself. While source, destination and services are important in defining DFW rules, for DFW rules implementation, parameters like “Applied To '' and “DFW exclusion list as discussed below are necessary to understand worth considering as this plays a critical role in DFW rule optimization with rules placement only at the required VMs since DFW are enforced at the vNIC of each VM in the compute networks.

Applied To - This effectively enables a default policy specific to a group/network, e.g. creating a deny all rule with an applied to of a group that has a network segment or segments as members. This is a flexible option that can't be done by physical firewalls and allows customers to implement a full DFW/micro-segmentation policy in stages. Applied to can only be used for objects in the environment (VMs, network segments, etc.) and not IP addresses, as it determines which NICs receive the policy rules.

DFW exclusion list - The DFW exclusion list allows specify inventory groups to exclude from distributed firewall coverage. Adding VMs to the exclusion list further optimize the system by reducing overhead and preventing it from having to scrutinize unnecessary traffic.

As an example, a typical use of DFW is to provide security between the tiers of a multi-tiered application or to provide security between separate tenants within the SDDC. DFW is unique in that it is applied at the absolute edge of the network. This feature effectively decouples network security from network architecture. To illustrate this point, imagine a traditional data center network. Typically, if security was needed between tenants or between tiers of an application, then network architecture would be designed to reflect this separation (i.e. a VLAN per tenant or application tier). In this model, if security requirements changed then the network infrastructure would need to be altered and workloads migrated and re-IPed. With DFW, security policy is agnostic of network architecture and may be enforced regardless of VM placement within the SDDC. Due to this decoupling, it is possible to provide security between tenants or application tiers even when the workloads reside within the same subnet as shown with micro-segmentation in Figure 5. If security policy changes, then workload migrations and IP changes are not necessarily required

Figure 5 - Micro-segmentation using DFW

For more information on DFW rules implementation, refer to Add or Modify Distributed Firewall Rules.

The Connected VPC

The Connected VPC provides access to native AWS services and resources directly from VMs running in a VMware Cloud on AWS SDDC through a high-bandwidth, low-latency connection that provides no-cost same-AZ data transfer. These native AWS services and resources are hosted in a customer managed AWS account..

Since network security between the SDDC and the connected customer VPC is managed in multiple places, it is worth specifically calling it out as a stand-alone topic. The security administrator must consider all of the points where security policy may be enforced for traffic between the SDDC and connected VPC:

- DFW - The security policy defined by DFW would be enforced at the vNIC level of all VMs within the compute networks.
- Gateway Firewall - Management gateway policies would affect connectivity to/from the management network, and compute gateway policies would affect connectivity to/from the compute networks.
- AWS Security Groups - The security groups of the connected VPC itself and security groups applied to any service/instances will impact connectivity to/from the VPC.
- AWS network Access Control List (ACL) - ACLs defined at VPC subnet levels will impact connectivity to/from the VPC.

The policies of the gateway firewalls and DFW have already been discussed, this section will focus on Security Groups within the connected VPC itself.

Cross-Account Elastic Network Interfaces (ENIs) are created at the time of SDDC deployment. There are several ENIs created when the SDDC is deployed but not all are active. Even though it is possible, you should avoid modifying or deleting these ENIs since doing so may impact the cross-link to the SDDC.

Connected VPC cross-linking is accomplished using the Cross-Account ENI feature of AWS and creates a connection between every host within the base Cluster of the SDDC to a subnet within the connected VPC. As part of this setup, these ENIs are configured to utilize the default Security Group of the VPC. An ENI is attached to each ESXi host in the SDDC's management cluster and has a status of "In-Use".One of the in-use ENIs will have a secondary IP address assigned and connected to the host running the active default edge. In case of an Edge HA failover or vMotion of active Edge, this mapping is changed to ENI of the host hosting the Edge.

You cannot apply a custom Security Group at the time of creation but can change it to a custom one based on your requirements. If you ever need additional ENIs (e.g. if the management cluster has to grow > 16 hosts due to uncontrolled storage growth), additional ENIs will be created by VMware for every host added, and those ENIs will be assigned the default Security group. In this case again the additional ENIs should be added to the custom Security Group.

Routing between the SDDC and the VPC is facilitated through static routes, which are created on-demand as networks are added to the SDDC. By default, these static routes are incorporated into the main routing table of the customer-owned connected VPC, utilizing the active Cross-Account ENIs as the next hop for the routes. It is crucial to note that the next-hop ENI designated for the static routes will consistently be that of the ESXi host hosting the active edge appliance within the SDDC. Consequently, should the edge appliance migrate to a different host (as can occur during failover events or when the SDDC undergoes upgrades), the next-hop for the static routes will be updated to reflect this change.

It is important to keep the following points in mind with Security Groups within the connected VPC:

Direction - As mentioned, the cross-account ENIs used by the SDDC have the default Security Group applied to them, the directions are relative to the SDDC. It is important to visualize this setup and remember that the “inbound” rules of the Security Group apply to traffic from the VPC toward the SDDC, and that “outbound” rules apply to traffic from the SDDC toward the VPC.

Default Rules - By default, Security Groups allow all traffic to go outbound, and no traffic inbound. Therefore if you have connections being initiated from the AWS VPC towards a VM in the SDDC, you may modify the default Security Group to allow that traffic. The destination is the subnets of VMs in the SDDC, which are not members of the default Security Group by our implementation. Note Security Groups are stateful, so traffic must be defined in the direction it originates, but allow bi-directional traffic flow between those endpoints once the flow is established.

Non-Default Security Groups - Oftentimes, other services within the VPC are utilizing custom Security Groups. In these cases, the security administrator may need to modify both the custom Security Group as well as the default Security Group to ensure that the required connectivity is permitted.

Refer Connected VPC and Connecting EC2 With SDDC Workload for further information.

NAT

Network Address Translation (NAT) services are provided to the SDDC by the tier-0 and tier-1 custom gateways.

To provide a secure environment, Internet access is blocked at edge firewalls for compute workloads for both inbound and outbound access, but firewall policy rules can be created to allow the required access. For management resources, inbound access is blocked and requires firewall policy rules while outbound access is allowed by default.

For outbound requests, by default, the workloads of the compute network will utilize a dedicated NAT IP that exists on the internet uplink of the tier-0 edge. Source NAT (SNAT) is automatically applied to all workloads in the SDDC to enable Internet access. This default SNAT IP is visible from the Home-> Overview section in the NSX UI console. Management resources like vCenter and HCX manager use different public IPs other than the default SNAT rule.

You can also request a public IP for compute workloads and create custom NAT policies for them. For default compute gateway attached workloads, you can create 1:1 or 1:many DNAT to access VM services directly from the internet. In cases where you require the workload VM to use the same IP for inbound and outbound traffic, you may create 1:1 reflexive NAT and use a service of “Any'' which will override the default SNAT rule. For custom tier-1 gateways attached workloads you have multiple options available like SNAT, DNAT, reflexive, “No SNAT '' and “No DNAT”.

In addition to NAT rules, the firewall rules must also allow the traffic. For custom tier-1 traffic leaving the SDDC, it must be allowed by both the custom tier-1 firewall and the default compute gateway firewall. It is important to note that Tier-0 NAT rules are only applicable when the default route is over the Internet uplink.

For more details on NAT refer to this link.

SDDC Traffic Inspection through AWS VPC or an On-premises Security Gateway

In some cases, it may be desired to perform traffic monitoring and inspection outside SDDC to meet security and compliance requirements as per the Organisation policy. SDDC supports this via advertising routes either through static routing over VPC or dynamically learning the default route pointed to on-premises gateway over direct connect or route based VPN.

Security VPC

Security VPC is a connectivity/security model where all the SDDC traffic is forwarded to this VPC for inspection and logging before being routed to the Internet or on-premises. This is achievable through manually configuring the static route on the VMware Transit Connect attachment to VPC where it requires a 3rd party instance to inspect and route the traffic.

The third party firewall or security appliance can be provisioned in one of the subnets of the security VPC and it can monitor the traffic between SDDC and Internet, SDDC and AWS native VPC,etc. and SDDC routed through the security VPC. Such an instance could be your existing firewall appliance configured with multiple zones(Trusted, Untrusted, DMZ etc.) hosted in native AWS VPC that you may already be using for security native AWS VPC infrastructure. You can extend the security of SDDC traffic with this security appliance as one of the trusted zones. Also these appliances with multiple vendors are available in the AWS marketplace that you may choose as based on your requirement.

Figure 6 - Security VPC for SDDC Traffic Inspection

Security with Traffic routed to On-premises over Direct Connect, Route based VPN or HCX

Direct Connect/Route based VPN

Routing all the SDDC traffic to on-premises over Direct Connect or route-based VPN is another option to forward the traffic for inspection and monitoring. In this case, the remote end has to advertise the default route, and the same is learned and installed at the SDDC.

Figure 7 -SDDC Connectivity to On-Premises over VPN and Direct Connect

HCX Network Extension

HCX Network Extension offers another way to control security, allowing traffic to pass through layer-2 extension to the on-premises security appliances for monitoring.

Note that egress charges may apply to VM traffic on extended networks communicating from VMware Cloud to on-premises. These charges will vary depending on whether your HCX service mesh is running over the internet or a Direct Connect. Also traffic performance is limited by the Internet/DX connectivity in place along with the per HCX Network Extension appliance throughput while additional appliances can be deployed to scale throughput. Using this method may increase overall complexity to the infrastructure, increasing additional traffic overhead on the connected path while providing limited routing options for the local access.

For more information refer HCX Network Extension.

Key takeaways:

- Protect traffic using existing on-premises hardware/virtual security appliances without any additional capital expenditure.
- Operations and maintenance of security appliances is not an additional overhead.
- Utilize existing configurations in place with minimal changes for security monitoring.

Traffic over VMware Transit Connect with SDDC Group

SDDCs in an Organisation can be grouped to connect with each other using VMware Transit Connect(vTGW). VMware Transit Connect uses the AWS Transit Gateway (TGW) construct. It provides high bandwidth and low latency connectivity between SDDCs in the SDDC Group with inter region and intra region support.

vTGW is a VMware managed connectivity solution that enables VMware Cloud on AWS SDDCs to connect between SDDCs in the same Organisation and also connect to external networks. It can also be used with multiple SDDCS or just one SDDC added to the SDDC group to connect to Native AWS Transit Gateways (TGWs), VPCs, and Direct Connect Gateway (DXGW) as a path to on-Premises networks.

Figure 8 - VMware Transit Connect with SDDC Group

For security considerations, traffic between connected SDDCs is considerably secure as it is using underlying AWS infrastructure.

SDDC Groups can also be connected together. When connecting multiple SDDC groups, only networks within the SDDCs themselves are shared across the connection.

NSX Advanced Firewall

NSX Advanced Firewall is also included as part of Distributed Firewall with VMware Cloud on AWS subscription. NSX Advanced Firewall functionalities provide the capability to define and enforce security policies at Layer-7 and enable deep packet inspection across all virtual networking endpoints within the SDDC. The decryption of network traffic is not supported using this functionality however, the parameters of the encryption used by network connections can be enforced to meet enterprise security requirements. For example, it is possible to restrict SSL-based applications according to the version of TLS protocol used.

NSX Advanced Firewall includes:

- Layer-7 Application ID Firewall
- Distributed Firewall with Fully Qualified Domain Name (FQDN) Filtering
- Distributed Firewall with Active Directory based User ID Identity Firewall

For more information on NSX advanced firewalls refer here.

NSX Distributed IDS/IPS

NSX Distributed IDS/IPS is an add-on capability that is separately billed when enabled on the SDDC. NSX Distributed IDS/IPS provides security operators with a software-based IDS/IPS solution that enables them to achieve regulatory compliance, create virtual zones, and detect lateral movement of threats on east-west traffic. Few important points to consider to use IDS/IPS functionality:

- Advanced IDS/IPS add-on enablement gives users access to provision the IDS/IPS functionality for enhanced security.
- IDS/IPS can be set with “detect only” for IDS mode or “detect and prevent” for IDPS mode. Even in detect-only mode, there are limits of how much traffic can be processed, and in case of oversubscription above the threshold limits results in traffic drop.
- Signatures available with severity classification from critical, medium, high, low and suspicious.
- Signatures can be customized at granular level based on the applications deployed.
- Provision to exclude signatures for inspection globally within the SDDC for custom use cases.
- Use the “Applied-to” field, directionality, and IP Protocol (IPv4/IPv6) where appropriate to limit unnecessary traffic to the IDPS Engine.
- VMs added to the Exclusion List in the Distributed Firewall are also excluded from any Distributed IDS/IPS Policy.
- Provision to auto-update the signatures automatically in sync the latest available signature from NSX Threat Intelligence Cloud (NTIC) service
- Support for offline signatures upload to the IDS/IPS, in such cases it has to be uploaded and updated manually from the NSX UI.
- IDS/IPS if enabled on all the clusters, there is no impact of vMotion/migration to applications running on workload VMs protected by IDS/IPS.

Figure 9 - NSX Distributed IDS/IPS

NSX IDS/IPS embraces an all-software distributed approach, moving traffic inspection out to every workload and eliminating the need to hairpin traffic to discrete appliances (i.e., moving traffic to and back from a centralized IDS/IPS capability). The operational simplicity of deploying and managing IDS/IPS functionality at each workload ensures comprehensive coverage without any blind spots.

Figure 10 - IDS/IPS Journey

The IDS/IPS lifecycle starts from enablement and configuration and once you have achieved the “Detect and Prevent” state, the lifecycle management of IDS/IPS progresses to monitoring and day 2 operations.

Classification Constructs

Services

A service definition may be thought of as collections of 1 or more protocols (IP, ICMP, UDP, TCP, etc…) along with their associated ports. Although many of the standard service definitions have been pre-created within the SDDC, it is sometimes necessary or convenient to create custom definitions.

Groups

Groups represent 2 classes of resources:

- VMs within a given network of the SDDC
- IP addresses that are external to a given network within the SDDC

Unlike services, which are defined globally within the SDDC, groups are scoped per the management and compute networks.

Context profiles

Context profiles enable setting attributes and sub-attributes.

Applied To

“Applied to” defines the scope to which the DFW rule will be applied and published. “Applied to” must be applied to groups of objects (e.g. VMs, network segments), and cannot be applied to IP addresses - any IP addresses in a group used for “Applied to” will be ignored.

Recommendations

The following are some basic recommendations for working with network security within an SDDC.

Service Definitions

There are a great many pre-created service definitions within the SDDC. However, it sometimes makes sense to create custom definitions for custom applications. Consider creating a single service definition that encapsulates a given function of a custom application. For example, if an application utilizes a pair of TCP ports then define both ports as part of the service definition.

Using Groups

There are a few different options available for group definitions (with additional new ones on the roadmap) within the SDDC. When creating groups, keep the following in mind:

- Anything external to a given network within the SDDC (management or compute) may only be referenced by IP. Utilize summary addresses as much as possible when defining IP-based groups.
- Anything native to a given network within the SDDC may be referenced by higher-level constructs such as VM name or security tag. Utilize these constructs as much as possible. Doing so will make your security policies more resilient to network changes within the SDDC.

Security tags provide an excellent tool for defining security policy. Put some serious thought into standardizing your tagging scheme. A common approach is to assume a “Lego brick” model for tags: small, atomic tags which may be combined to effectively classify a workload. However, keep the maximum tags-per-VM limit in mind when designing your scheme.

Managing Sections

DFW utilizes sections as a means of organizing firewall rules. Consider organizing sections around your specific business requirements. For example, create a section per application or business unit. Always remember that sections are only a means of organizing rules; rules are evaluated top-to-bottom independently of their parent section.

Porting Existing Rulesets

If your organization is like most organizations, then you have accumulated several years’ worth of defining and using objects in older ways in your existing security policies. This tends to happen with IP-based rulesets. The rules quickly become complex, and security administrators are afraid to remove anything for fear of breaking something.

Avoid the temptation to port existing rulesets to the SDDC. You have a unique opportunity to rework your policies based on higher level grouping constructs which will very likely simplify your security policy. Don’t miss this opportunity.

Managing Context profiles

Recommended to use a specific application or service context-aware to allow traffic is matching to application protocol fingerprint only for enhanced security.

Using Applied To

If Applied to is not optimally utilized, the rule becomes generic and becomes available at every vNIC of VM and every host. Consider using the Applied to the selected groups.

Logging

Gateway and distributed firewall, and Distributed IDS/IPS logging is disabled by default and can only be enabled on a per rule basis. Enable firewall or advanced firewall logs for visibility and troubleshooting in case of any issues. There may also be a need to enable and archive firewall logs for security compliance requirements. Any of these firewall and IDS/IPS logs, once enabled and configured properly can be accessible as part of Aria Operations for Logs.