CHAPTER 4

Network Security

In this chapter, I present the various network controls available to you in Azure environments, their function, and how to apply them. I discuss the best practices on implementing these controls for infrastructure-as-a-service and platform-as-a-service workloads.

Azure Virtual Networks

In this section, I introduce you to the key concept of Azure networking – the Azure virtual network.

Microsoft Global Network

To provide cloud services, Microsoft operates a physical network of over 160 **datacenters** globally. The datacenters are grouped into **regions** which operate within interconnected regional networks. These regional networks are connected to[1] through Microsoft global-wide area network (WAN) over private fiber-optic cables.

[1] https://docs.microsoft.com/en-us/azure/expressroute/expressroute-locations-providers#expressroute-locations

K. Ots, *Azure Security Handbook*, https://doi.org/10.1007/978-1-4842-7292-3_4

Customer traffic from the Internet to Microsoft's global WAN enters or exists through **point of presence** locations, with even broader physical footprint than of datacenter regions. Alternatively, enterprise customers can create ExpressRoute connections to connect to the Microsoft global WAN without leaving their dedicated network. ExpressRoute locations are co-location facilities where Microsoft Enterprise edge (MSEE) devices are located.

When connecting through the Internet, network traffic is protected with Azure DDoS Protection. With Azure DDoS Protection, Microsoft applies traffic monitoring and real-time mitigation against network-layer attacks, such as network flushing, before the traffic is routed to customer instances.

IP Addresses in Azure

IP addresses under your control are considered resources in Azure. As such, they are deployed into a resource group in a subscription. IP addresses are subject to the same role-based access control and Azure policies like any other resources. Figure 4-1 illustrates the different resource types related to IP addresses and how the access control and life cycle of IP addresses are independent.

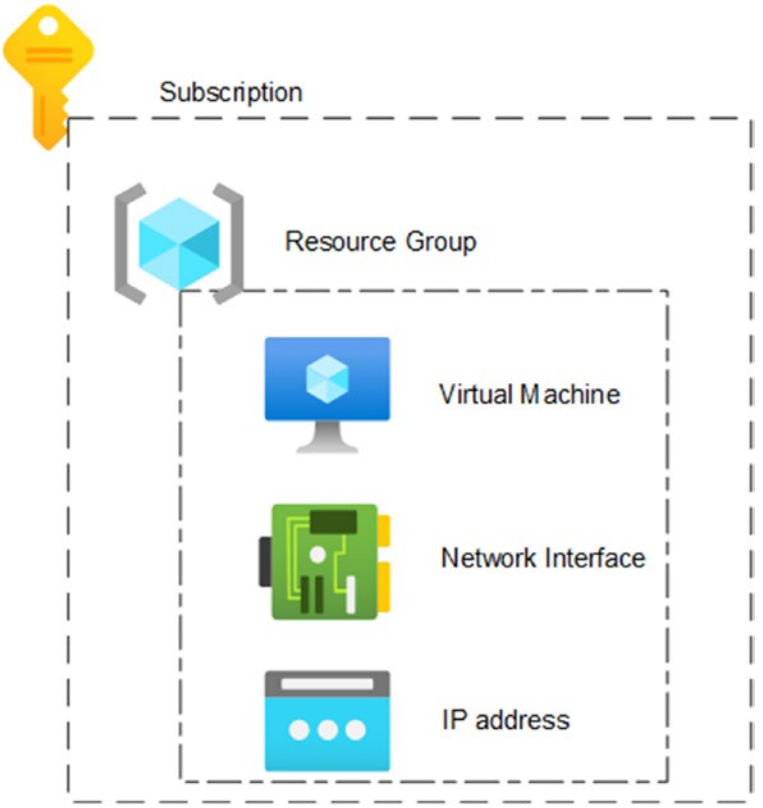

Figure 4-1. *IP address resource*

While managed as independent resources, IP addresses are often associated to other resource types, such as a

- Virtual machine network interface
- Internet-facing load balancer
- VPN gateway
- Application Gateway

Azure IP addresses support either dynamic or static address allocation. You can use both IPv6 and IPv4 addresses in Azure. The actual IP address is tied to the life cycle of the IP address resource. If you want to reallocate an IP address, you need to disconnect its association to another resource, before deleting it. Once that IP address is no longer associated with the resource, you are free to associate it with another resource.

Note When you delete an Azure IP address resource, the IP address is released to be used by other Azure customers!

Azure Virtual Network

A key component in Azure networking is the Azure virtual network. Virtual networks bring a level of control similar to on-premises networks into that multitenant cloud world. Any resource deployed into a virtual network are isolated from other cloud users and even your own virtual networks. You can control inbound and outbound traffic, IP addresses, and routing for Azure virtual networks. Figure 4-2 illustrates the virtual network isolation.

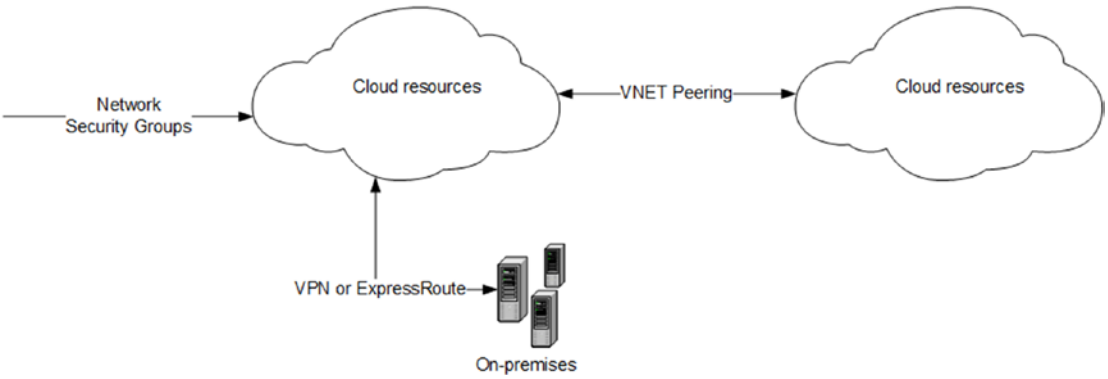

Figure 4-2. *Virtual networks isolate network connectivity within the cloud, as well as inbound or outbound to/from on-premises or public networks*

To control network traffic within Azure, Azure virtual networks can communicate with other Azure virtual networks by using virtual network peering. For example, an application deployed to multiple Azure regions could leverage virtual network peering to extend the logical virtual network across regions.

Azure virtual networks are used to communicate with on-premises resources, too. You can connect your virtual networks to other networks outside of Azure using the Azure virtual private network (VPN) or Azure ExpressRoute.

Like IP addresses, virtual networks exist within the Azure Resource Manager as resources. Therefore, access control and life cycle considerations should be kept in mind when planning network topologies. Virtual network resources have several properties

to be configured, such as IP address spaces, gateways, or user-defined routes. Crucially, subnets are not considered properties of virtual network resources, but rather sub-resources. Role-based access control can be granted for in the sub-resource scope, too. The logical hierarchy is illustrated in Figure 4-3.

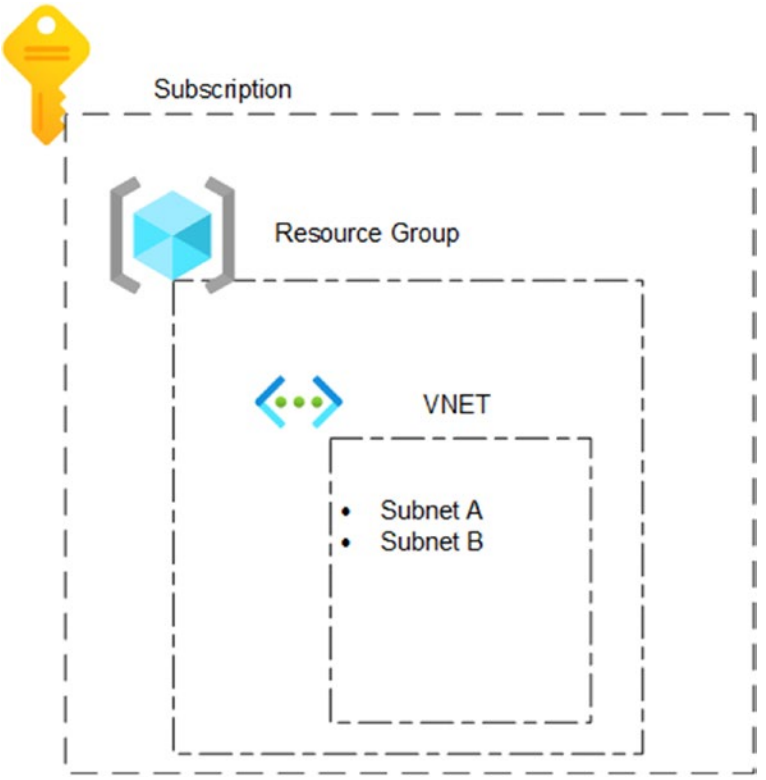

Figure 4-3. *Logical hierarchy of virtual network resource and subnet sub-resource in Azure Resource Manager*

Network Controls for Infrastructure as a Service

Network Security Groups

The native way to control network traffic to and from resources in your Azure virtual network is the network security group. Like an access control list, a network security group consists of a list of security rules that allow or deny traffic. If you are familiar with VLAN segmentation, virtual network subnets and network security groups offer a

comparable solution in Azure. For each network security group security rule, you can specify

- **Traffic source or destination**: IP address range, service tag, or application security group
- **Protocol**: TCP, UDP, ICMP, ESP, AH, or any
- **Port range**
- **Rule priority** between 100 and 4096 (lower numbers are higher priority)
- **Rule action** (either allow or deny)

Service tags are Microsoft-managed lists of IP addresses[2] for multitenant Azure services. You can use service tags to create network security group rules that deny or allow traffic to the public endpoints of Azure PaaS services, while still preventing access to or from the Internet. For example, the service tag **AppService.WestEurope** can be used to allow outbound traffic to Azure App Service. Microsoft updates service tags and IP address ranges periodically. Service tag information is available for download as JSON files, as well as programmatically through the Service Tag Discovery API.

Application security groups are logical groupings of security rules for your own applications. For example, if all your front-end virtual machines need connectivity blocked to the Internet but allowed to your back-end virtual machines, application security groups can be used.

Azure evaluates incoming or outgoing traffic against network security groups. Traffic flows are interrupted when connections are stopped, and no traffic is flowing in either direction.

Note Existing connections are not interrupted when you remove a security rule that enabled the flow. NSG flows are evaluated for new connections.[3]

---

[2] www.microsoft.com/en-us/download/details.aspx?id=56519

[3] https://docs.microsoft.com/en-us/azure/virtual-network/network-security-groups-overview#security-rules

Network security groups can be assigned in the subnet or the network interface scope. For inbound traffic, Azure processes NSG rules configured in the subnet scope first and then the rules in a network security group associated to the network interface. For outbound traffic, the order of processing is reversed. Figure 4-4 illustrates the combination of different network security group destination and source types and cumulative network security group rules.

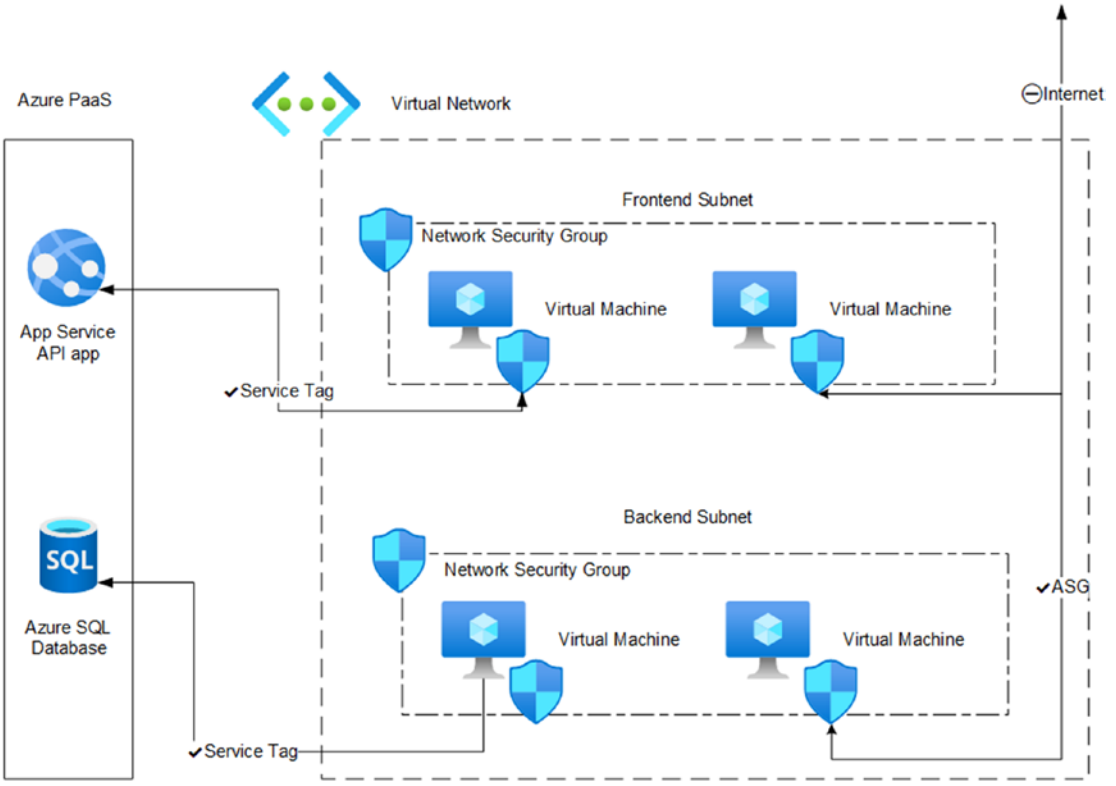

Figure 4-4. *Network security groups, service tags, and application security groups*

Note Network security groups cannot be assigned at a virtual network scope, but only on the subnet scope!

Securing Administrative Access to Virtual Machines

One of the key attack vectors for workloads running in public cloud is misconfigurations. Exposed management ports are often scanned by adversaries as part of the initial discovery tactics.[4]

As the IP address ranges of Azure datacenters are relatively easy to obtain, it is safe to assume that if you leave a management port exposed to the Internet, it will be picked up by adversaries in the matter of minutes. Once discovered, your virtual machines are susceptible to password spray attacks. Azure DDoS Protection basic or other protocol-level network protections will not help you to prevent this. To protect yourself from this and still allow legitimate access, you need to introduce additional network controls.

As part of Azure Defender for Servers, Azure Security Center just-in-time (JIT) access to virtual machines reduces the exposure of your virtual machine management ports by managing the network security group or Azure Firewall rules on demand. Once configured, Azure security center JIT enforces that inbound traffic is denied. To open management port access from their network location, users need to perform a just-in-time activation in Azure Portal. This verifies user role-based access control, audits the access, and opens the management ports for a predetermined time. Specifically, the requesting user needs to have access to the `Microsoft.Security/locations/jitNetworkAccessPolicies/initiate/action` resource provider action in the scope of the virtual machine resource group.

As discussed earlier, network security groups interrupt traffic flows when connections are stopped, so just-in-time access to virtual machines does not prevent your privileged users from keeping their connections open even after the just-in-time access time has ended.

An alternative to Azure Security Center just-in-time virtual machine access would be to create a point to site VPN connection to the virtual network of the virtual machine. Azure point-to-site VPN supports most client operating systems. The point-to-site VPN access can be authenticated using Azure AD (Windows 10), RADIUS, or certificate authentication. It supports the following protocols:

- OpenVPN®
- Secure Socket Tunneling Protocol (SSTP)
- IKEv2 VPN

[4] https://attack.mitre.org/techniques/T1046/

As an alternative to client-based point-to-site VPN, you can use Azure Bastion. Azure Bastion provides HTML5-based web client that is streamed to your administrative users through the Azure Portal. Behind the scenes, Azure Bastion is a Microsoft-managed and hardened service that is deployed into a separate subnet in your virtual network. Azure Bastion then automatically manages network security group rules to allow connectivity between the Azure Bastion subnet and your virtual machines.

Azure Security Center just-in-time access is the easiest of these to set up. It reduces the exposure time and source addresses of administrative access. However, it does not remove them altogether. Azure point-to-site VPN adds an additional layer of protection by keeping that management ports exposed and requiring user authentication for the VPN gateway connection. Azure Bastion removes any need for exposed IP addresses. Based on your user experience and performance requirements, however, Azure Bastion might not meet your needs. As with all cloud security, securing management port access comes down to balancing your need for agility and control.

Securing Outbound Access from Virtual Machines

The federated nature of Azure networking adds complexity in monitoring and securing outbound network traffic of your virtual machines. The default NSG rules allow outbound Internet connectivity on any port. For virtual machines that are not connected to your centralized virtual networks, you might not have a straightforward way to determine malicious traffic from legitimate traffic.

Azure Security Center and Azure Firewall allow you to filter out and alert on outbound communication with malicious IP addresses. Azure Security Center's adaptive network hardening also provides recommendations for narrowing down any NSG rules, based on your actual traffic patterns.

For some Azure offer types, outbound connectivity is limited even further. Specifically, outbound connectivity over port 25 (SMTP) is disabled for most offers. When you are using one of the paid subscription types (pay as you go, CSP, and EA), you can request opening of this port through support.

Network Controls for Platform as a Service

In this section, I discuss the network controls available in Azure application and storage PaaS services.

Application PaaS Networking

Most Azure application platform-as-a-service resources are multitenant by nature. From the networking perspective, this means that the services are by default reachable by anyone through the public endpoints. Furthermore, controlling outbound network traffic can also be quite limited. Having the option to deploy your platform-as-a-service resources inside your virtual network often requires committing to a higher pricing tier.

Controlling Inbound Traffic to App Service

Azure App Service is one of the most widely used application platform-as-a-service resource types in Azure. In the multitenant App Service, inbound traffic can be restricted using Access Restrictions,[5] a feature that provides allow and deny rules for inbound access. Access Restriction rules are set separately for management operations and content. Access Restriction rules are evaluated separately from your App Service instance, by the managed App Service front end, so any denied traffic never reaches your application code. Access Restriction rules can be set to limit incoming traffic based on

- IP address ranges
- Virtual network service endpoints
- Service tags
- HTTP headers

Controlling Outbound Traffic from App Service

Outbound traffic from Azure App Service can be secured through virtual network integration or the App Service Hybrid Connections feature.

[5] https://docs.microsoft.com/en-us/azure/app-service/app-service-ip-restrictions

App Service virtual network integration allows **outbound** connectivity from your App Service to your virtual network. Your application can then reach resources in the virtual network. Virtual network integration does not grant inbound private connection from your virtual network. App Service can be configured to route all its outbound traffic to the virtual network.

Figure 4-5 illustrates the Azure App Service networking controls. First, Access Restrictions are used to deny inbound connectivity from the Internet and allow inbound connectivity from the specific IP address range. Second, virtual network integration is configured for controlling outbound traffic from the App Service application. Third, integration subnet is configured to allow traffic from the App Service endpoint, and the back-end subnet's network security group is configured to allow traffic from the integration subnet. And finally, all outbound traffic is enforced to go through the virtual network using the `WEBSITE_VNET_ROUTE_ALL` setting in App Service.

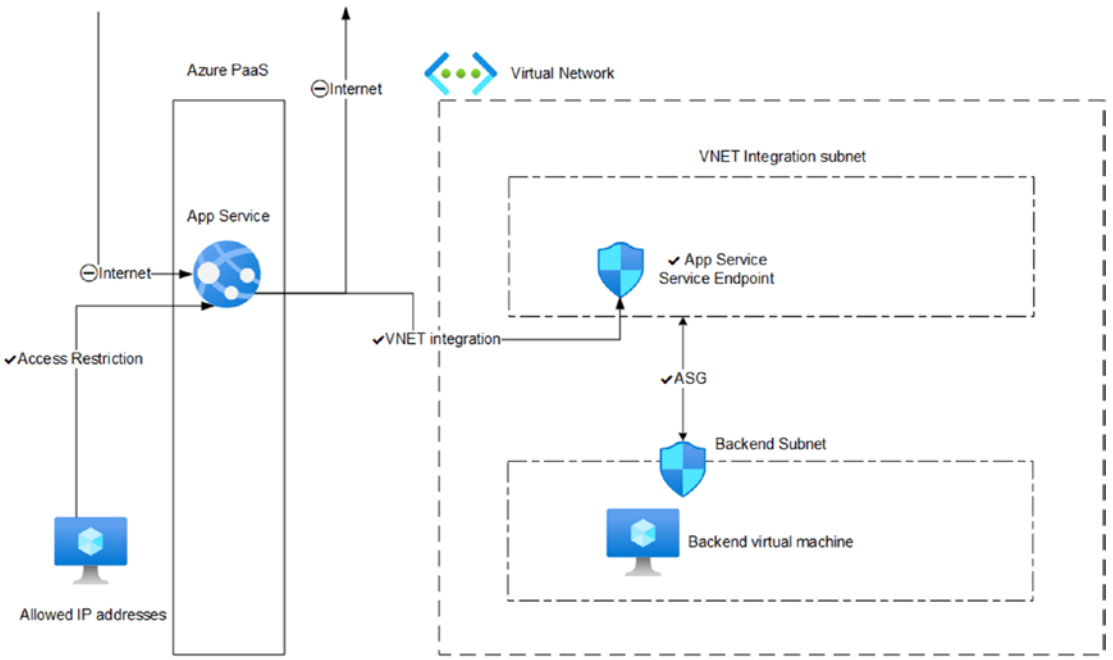

Figure 4-5. *App Service networking controls*

Cross-Network Connectivity

App Service Hybrid Connections enables access to your on-premises resources, without changes to on-premises inbound firewall rules. Instead, Hybrid Connections depends on a relay agent to be installed on your on-premises network and outbound TCP connectivity to Azure over port 443. Once configured, any DNS requests in your App Service application that match your Hybrid Connection endpoint will be redirected through the Hybrid Connection relay.

Finally, the single tenant tier of Azure App Service, App Service Environment, is in fact placed in a virtual network that you control yourself. This means that you can use the same virtual networking connectivity options that you can for any other virtual network resources, such as virtual machines.

Network Controls for Data Platform as a Service

You can store your data on several platform-as-a-service data stores in Azure. Just like application platform as a service, data platform-as-a-service resources are often hosted in a multitenant environment. Gaining access to full virtual network injection capabilities is either not possible at all or available only in the highest pricing tiers.

By default, most multitenant Azure data services are available through the public endpoints in an unrestricted manner. It is our responsibility as the cloud consumer to configure traffic restrictions. When combined with limitations on the other parts of our workloads, such as legitimate end users or continuous deployment pipelines, this can become quite complex.

Storage Account Firewall

Most Azure data services offer a variation of a **firewall** functionality. For example, the firewall of Azure storage account lets you control inbound traffic by creating access rules that target IP address ranges or virtual networks. This is done by changing the **default network access rule** to deny all incoming traffic by default. Once access is restricted, you can configure storage account to allow access only from certain public IP address ranges. Additionally, private addresses can be granted by allowing virtual network access, provided that the target virtual network allows outbound traffic from the virtual network to storage account service endpoint. Figure 4-6 illustrates a storage account that denies default public access and allows access from a specific subnet and a single public IP address.

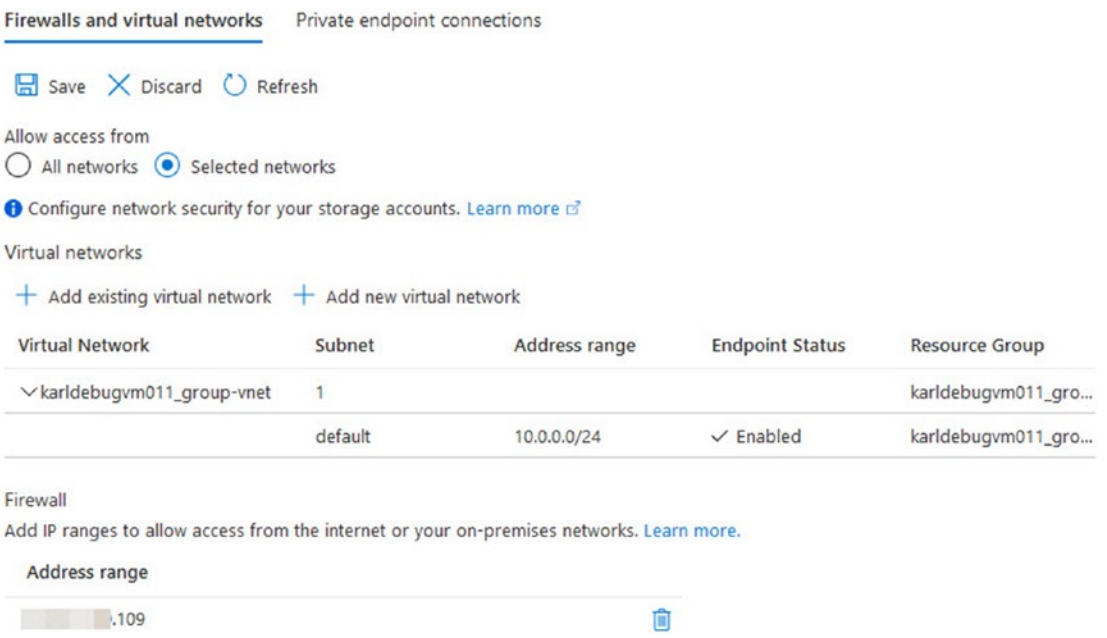

Figure 4-6. *Storage account firewall configuration*

Azure Monitor Network Isolation

The firewall functionality of Azure Monitor service, including Log Analytics Workspace and Application Insights, allows to configure separate network rules for data ingress and egress. Figure 4-7 illustrates this functionality. Like for storage account, you need to allow egress traffic from your network to Azure monitor. To do this, you can use the `AzureMonitor` service tag.

Figure 4-7. *Azure monitor network isolation controls*

Private Endpoints

As an alternative to IP address or virtual network access rules, traffic between Azure platform-as-a-service resources and your virtual network cannot be secured using private endpoints.

Private endpoints expose your instance of a multitenant Azure service as private IP addresses within your virtual networks. Using private endpoints protects from data exfiltration and gives you more control of both inbound and outbound traffic. Instead of service tags or service endpoints, your access rules can be set in a more granular level.

Private endpoint configuration is more complex than just enabling access rules or creating service endpoints. As all traffic between your private endpoint resource and your network is within your control, you can no longer rely on Microsoft to resolve your resource names. Creating a private endpoint connection requires you to also configure private DNS zones. For some services, such as Azure storage, you will need to configure the DNS zones for each storage endpoint.

While adding an additional layer of control, private endpoints also add an additional price component. As the traffic is being transferred across your dedicated connection, there is a data transfer fee for both inbound and outbound data.

Azure Firewalls

One of your key network control decisions is about whether to use Azure-native controls or third-party controls. Azure-native controls provide tighter integration and fast time to market. If you want to build more complex networks or leverage existing licenses or skills come up, you can bring in network virtual appliances from existing vendors through the Azure marketplace. The best solution depends on your cloud strategy. If you plan to take advantage of modern platform-as-a-service components, the benefits you gain from third-party network virtual appliances are more limited. However, if your cloud strategy involves a large footprint on infrastructure as a service, bringing in advanced network virtual appliances can be beneficial.

Azure Firewall is a stateful, Layer 3–7 firewall that controls virtual network access. It is a separate service from the stateless network security group rule sets and provides additional capabilities, such as

- Application FQDN filtering rules
- Network traffic filtering rules

- Threat intelligence
- Outbound SNAT support
- Inbound DNAT support
- Forced tunneling
- TLS inspection
- Intrusion detection
- Intrusion prevention

Azure Web Application Firewall

Azure Web Application Firewall is a managed, OSI Layer 7 firewall. Azure WAF is based on OWASP Core Rule Set,[6] updated and managed by Microsoft, protecting your web applications against

- SQL-injection protection
- Cross-site scripting protection
- Common Apache and IIS misconfigurations

Azure Web Application Firewall functionality can be deployed to Azure Application Gateway or Azure Front Door. The firewall rules can be managed centrally using Web Application Firewall Policies. Azure Front Door is deployed on Azure network edge locations, inspecting incoming requests before they even reach Microsoft global network.

Network Monitoring

Azure networks provide logs for both our security information and event management (SIEM) and forensic investigation needs. As with other services, the logging needs to be consciously turned on.

[6] https://owasp.org/www-project-modsecurity-core-rule-set/

Security information and event management should consume logs from NSG diagnostic logs and Azure Firewalls. NSG diagnostic logs contain element information about which security rules are applied to which virtual machines, based on Mac addresses. In addition, Azure activity logs contain logs about configuration changes to public IP instances or NSGs.

Azure Firewall provides log event for each new connection (either accepted or denied). These logs contain evaluated application and network rules. Each log event stores information about the connection source, protocol, ports, and the result of the rule evaluation. If threat intelligence mode is configured, Azure Firewall also logs that read intelligence alerts.

Azure Web Application Firewall (WAF) Access logs provide further application-level information. WAF Access logs contain

- httpMethod
- requestUri
- RequestQuery
- sslProtocol and sslCipher
- UserAgent
- originalHost

Furthermore, Azure Web Application Firewall's Firewall logs contain information about evaluated firewall rule sets.

Logs Supporting Forensic Investigation

Azure network watcher is a service that provides network security group flow log information and point-in-time packet capture functionality.

All virtual network traffic is true network security group. Whenever Azure virtual network traffic passes through a network security group rule set, flow logs can be created. NSG rules are evaluated at OSI Layer 4. The NSG flow logs include the following information:

- Source IP, port, and protocol
- Destination IP, port, and protocol
- Traffic direction
- Traffic decision, indicating either denied or allowed traffic

You should enable NSG flow log collection and export the logs to your intrusion detection systems.

For investigation purposes, Azure network watcher also supports packet capture. This is relatively intrusive agent-based approach, though. Network watcher packet capture is done within a virtual machine extension `AzureNetworkWatcherExtension`. The packet capture can be started manually or based on virtual machine alerts.

Alternative Network Monitoring

As part of Azure security center, Threat Protection for Azure network layer can alert on anomalous activities. These network alerts are raised based on traffic analytics from packet headers and Microsoft Threat Intelligence databases of malicious traffic addresses. Threat Protection for Azure network-layer alerts includes

- Network communication with a malicious machine detected
- Possible incoming SQL brute-force attempts detected
- Possible outgoing port scanning activity detected
- Suspicious incoming RDP/SSH network activity from multiple sources

For raw traffic analytics, you need to use third-party network virtual appliance (NVA) solutions hosted in virtual machines within your virtual networks. Azure marketplace offers third-party solutions for both network packet brokering as well as security analytics and network performance management.[7]

Cloud Adoption Framework

The Cloud Adoption Framework is a series of Microsoft best practice articles and artifacts. As part of the Cloud Adoption Framework, several network architectures are discussed, including traditional hub and spoke and modern networking. The Cloud Adoption Framework's network topology and connectivity section provides guidance for

[7] https://docs.microsoft.com/en-us/azure/virtual-network/virtual-network-tap-overview#network-packet-brokers

- Defining Azure network topology using Azure-native or traditional network solutions
- Connectivity to/from Azure
- Connectivity with Azure PaaS services
- Connectivity with other cloud providers
- Landing zone planning
- Network segmentation
- Traffic inspection
- Private endpoint and DNS configuration at scale

For many key scenarios, such as hub and spoke network topology, the Cloud Adoption Framework also provides infrastructure as code samples and reference architectures.

EXERCISE

You are tasked to secure an application hosted in Azure App Service. One of your requirements is the ability to provide logs of each failed attempts of access. Propose a network architecture that meets this requirement.

Summary

Azure networking at the enterprise scale is immensely complex. You need to solve cross-subscription, cross-region, cross-premises, and cross-cloud connectivity, all while balancing between user flexibility, control, and signal to noise ratio.

While it can be tempting to try to replicate on-premises networking topologies and operations, it is simply not technically feasible. You need to select alternatives that still satisfy your risk appetite and redesign your networking approach for cloud era. If you are compelled to replicate on-premises networking in the cloud world, consider that only a partial solution. If "traditional" workloads need to be segmented in the hub and spoke model, could we find new workloads that suit the cloud networking model better?