Microsoft Azure Network Security

Microsoft

Abstract

This document is a guide to enhancing network communications security to better protect virtual infrastructure and data and applications deployed in Microsoft Azure.

The intended audience for this whitepaper includes:

- IT and network administrators interested in deploying applications in Azure
- Developers interested in creating applications that run in Azure
- Technical decision makers (TDM) considering Azure to support new or existing service offerings

NOTE: Certain recommendations contained herein may result in increased data, network, or compute resource usage, and increase your license or subscription costs.

Version 3, Published February 2015

Table of Contents

1 Overview

Microsoft Azure (Azure) networking provides the infrastructure necessary to securely connect Virtual Machines (VMs) to one another, and be the bridge between the cloud and on-premises datacenter.

Azure's network services maximize flexibility, availability, resiliency, security, and integrity by design. This white paper provides details on the networking functions of Azure and information on how customers can use Azure's native security features to help protect their information assets.

2 Guidelines for Securing Azure Virtual Machines

Azure is a multi-tenant platform that uses shared infrastructure to support millions of simultaneous customers across more than 80 global datacenters. Because Azure's shared infrastructure hosts hundreds of millions of active VMs, protecting the security and confidentiality of network traffic is critical.

Azure Virtual Networks use a combination of logical isolation, firewalls, access controls, authentication, and encryption to protect customer data in-transit. Microsoft's Azure datacenter operations implement comprehensive information security policies and processes using standardized industry control frameworks such as ISO 27001, SOC 1, and SOC 2. Third-party auditors regularly certify Microsoft's adherence to these standards for both the physical and virtual aspects of Azure infrastructure.

In the traditional datacenter model, a company's Information Technology (IT) organization controls networked systems, including physical access to networking equipment. Company employees or contractors are responsible for deployment, configuration, and management duties, such as physically altering network topology, changing router settings, deploying firewall devices, and so on.

In the cloud service model, the responsibilities for network protection and management are shared between the cloud provider and the customer. Customers do not have physical access – they cannot walk into a cloud provider datacenter and rewire a server rack – but they implement the logical equivalent within their cloud environment through tools such as Guest operating system (OS) firewalls, Virtual Network Gateway configuration, and Virtual Private Networks. This physical and logical separation enables customers to rely on the fundamental security capabilities delivered by Azure as they build their infrastructure.

2.1 Private Networks

The logical isolation of customer infrastructure on a public cloud is fundamental to maintaining security. Azure accomplishes this primarily through a distributed virtual firewall. Further, a customer may deploy multiple logically isolated private networks. These sub-divided networks generally fall into one of two categories:

- **Deployment network**: Each deployment can be isolated from the other deployments at the network level. Multiple VMs within a deployment can communicate with each other through private IP addresses.

- **Virtual network**: Each virtual network is isolated from the other virtual networks. Multiple deployments inside the same subscription can be placed on the same virtual network, and then communicate with each other through private IP addresses.

Figure 1 illustrates an example of a virtual network topology.

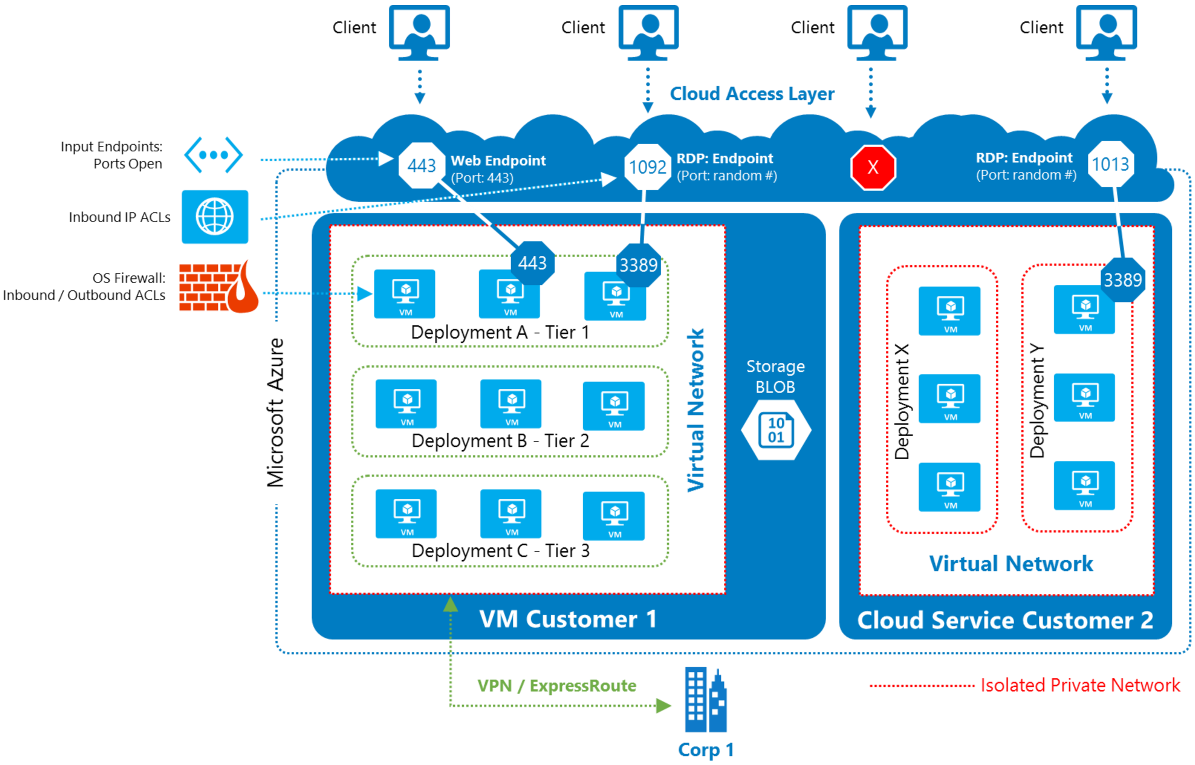

Figure 1. An example of isolated multi-tier IaaS applications hosted within Azure.

Network administrators can manage these isolated private networks in a way similar to the management of on-premises private networks.

The mechanisms for administrators to manage network security on their Azure private networks are in the Azure Cloud Access Layer, which is comparable to the edge of a corporate network that faces the Internet. The Cloud Access Layer includes a firewall, load-balancer, and network address translation (NAT) functionality managed by the customer administrator.

2.1.1 Enabling Internet Communications

By default, VMs inside the private network do not receive inbound traffic from the Internet. Administrators can enable Internet communications in one of three ways:

- The administrator can define an input endpoint that specifies which VM port mapping should receive inbound traffic initiated from outside a deployment's isolated network, including Internet traffic and traffic from other VMs inside Azure.
- The administrator can further enhance security by defining Azure Security Groups that specify which IP addresses should receive inbound traffic initiated from outside a virtual network. Note, that the administrator can either define input endpoint Access Control Lists (ACLs) or Security Groups, but not both.
- The administrator can assign an instance-level public IP address to a Virtual Machine. Then, all ports of the VM are accessible from the Internet.

NOTE: The term "inbound traffic", used in this paper, refers to traffic initiated by a computer on the Internet or outside of a customer's private network in Azure. This is also called unsolicited inbound traffic to distinguish it from inbound traffic that is a response to a request, also known as solicited inbound traffic.

2.1.2 Securing Communications

Securing communications between VMs inside the private network

Virtual Machines inside a deployment network can communicate internally via private IP addresses. Communication security between VMs in multiple deployments of a subscription can be enhanced by using Virtual Networks.

If an application sends or receives sensitive data over an internal private network, such as through a VPN, the data can be encrypted using IPsec, SSL/TLS, or other application-level encryption technologies. Customers with higher confidentiality or privacy concerns (such as for compliance with different industry regulations and standards) should ensure that all private communications between VMs inside a region are encrypted.

For more information on configuring Virtual Networks with encryption, see the Azure Virtual Network documentation on MSDN.

Securing inbound communications from the Internet

By default, Azure blocks inbound traffic from the Internet on a VM created through the Azure Management Portal, except for remote management ports.

Administrators can decide which VM ports and IP addresses can be accessed from the Internet via inbound traffic. Additionally, administrators can implement a number of configuration changes to help secure remote access at the network level from the Internet to ports on VMs or VNETs, including:

- Defining input endpoints at the Cloud Access Layer to open ports only when needed. Administrators can specify access control lists (ACLs) on input endpoints to control the source IPs from which the VM will allow traffic.

- Defining Security Groups to control open inbound traffic to specific VMs inside a virtual network. The administrator can either define input endpoint ACLs or Security Groups, but not both.
- Using a third-party proxy firewall (such as the Barracuda Web Application Firewall Vx or NG Firewall Vx virtual appliances) that runs on a VM to filter traffic to other VMs. Add the VMs to a virtual network, and then define an input endpoint for a port on the proxy firewall.
- Defining open ports in the firewall inside the Guest OS VM.

If an administrator opens input endpoints or IP addresses, the administrator should follow the same security model as if the VMs were running open on the Internet. If the application sends or receives any sensitive data on the input endpoints, then all input endpoints should use server and client authentication, and communication should be encrypted. If an application sends or receives sensitive data over public networks (including over Public IP addresses within a region), then communications should be encrypted using SSL or similar application-level encryption technologies.

Securing communications across subscriptions

A customer may have multiple subscriptions, and there may be a need for communication between VMs in different subscriptions. In these cases, VMs can be configured to communicate via public virtual IP addresses. Additionally, IP ACLs would need to be configured on input endpoints to allow VMs to initiate connections only with each other.

It is important to use Reserved IP Addresses for the public virtual IP addresses so that the IP address ACLs do not change.

For more information on configuring IP address ACLs, see About Network Access Control Lists.

Securing communications to on-premises networks

When workloads require secure communications between the Azure Virtual Network and on-premises systems, it is best to protect those channels using a Virtual Network Gateway. There are two (2) deployment scenarios:

1. **Internal Multi-Tier Application**: A multi-tier application (such as a web-based records processing system) deployed on Azure where the application does not need any inbound connectivity from the Internet, but does need connectivity to servers and applications in the customer's on-premises network, as shown in Figure 2.

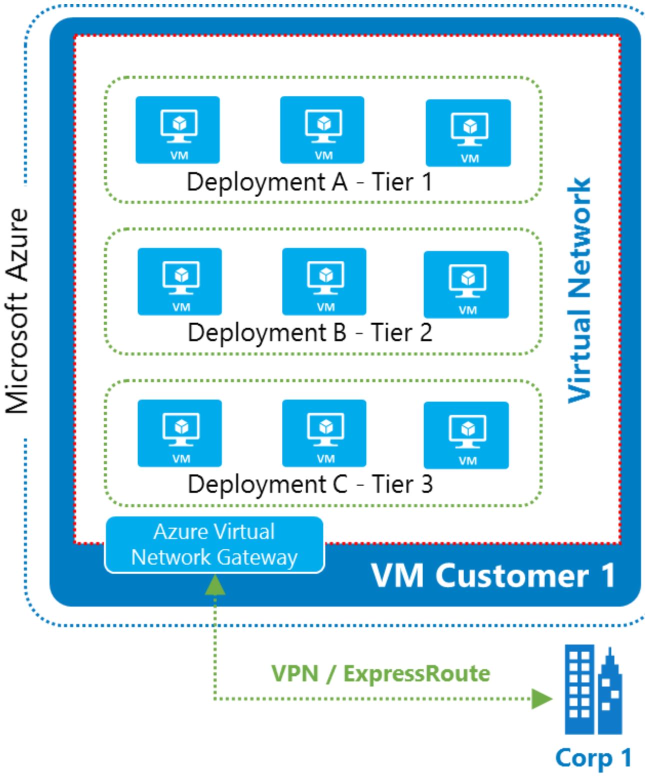

Figure 2. VPN connection between a corporate network and Microsoft Azure.

If the administrator needs to create a VNET to VNET connection, the administrator can create a virtual network and add the VMs of the application tiers to the virtual network, but does not need to define any input endpoints. Additionally, the administrator should:

- Remove the remote management input endpoints or lock them down using the guidance provided below in the section titled "Isolating VMs within a virtual network for defense-in-depth" to secure management endpoints.
- Configure the Virtual Network Gateway or ExpressRoute so that traffic destined for the corporate network flows through the VPN.

2. **Public-Facing Multi-Tier Application**: A multi-tier application deployed in Azure, and the front-end tier requires inbound connectivity from the Internet (over SSL port 443). The back-end tiers do not need inbound connectivity from the Internet, but do need connectivity to the customer's corporate network, as shown in Figure 3.

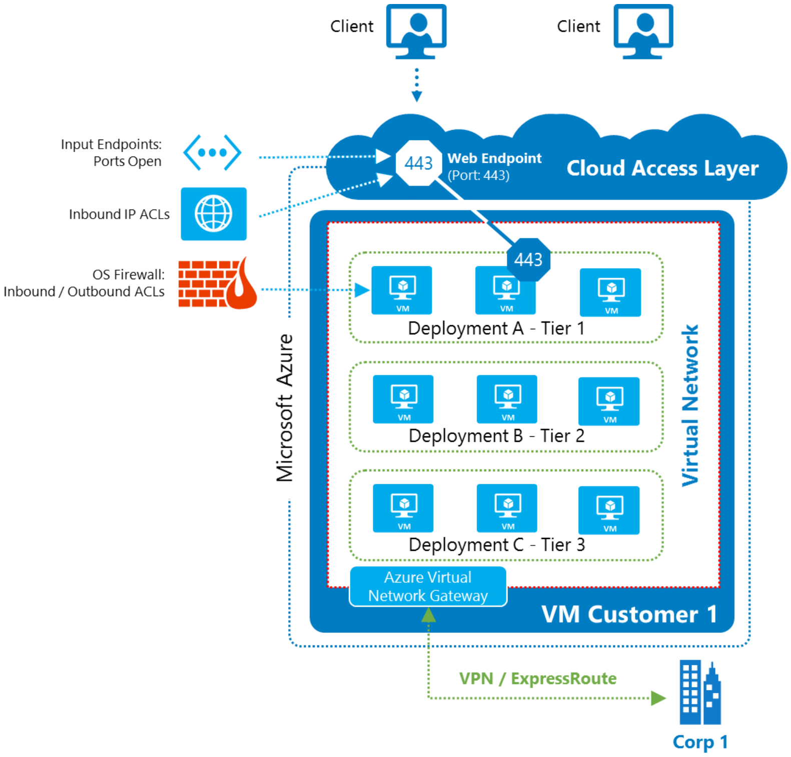

Figure 3. Addition of Internet-facing input endpoints to allow Internet access to the front-end tier.

In this case, the administrator should:

- Create a virtual network with the appropriate VMs from each of the application tiers.
- Define input endpoints for the inbound Internet traffic for the VMs in the front-end tier.
- Remove the remote management input endpoints for all VMs or lock them down.
- Configure the Virtual Network Gateway so that traffic destined to the corporate network flows through the VPN connection and to the corporate network.

The Virtual Network Gateway establishes an IPsec tunnel to route traffic between the virtual network and the customer's VPN device. This can be a hardware VPN device or a software VPN such as Windows Server 2012 Routing and Remote Access Services.

Creating a virtual private network within Azure provides the ability to more securely extend on-premises networks into Azure. This connection can be a site-to-site or point-to-site VPN.

If the VPN within a region is connected to a corporate network over the Internet through a virtual network gateway, then those communications are encrypted by default with a standard such as AES-256, although the configuration is dependent on the site-to-site VPN gateway on the corporate network.

If the virtual private network within a region uses direct connectivity technology such as Azure ExpressRoute to connect to corporate network, this traffic is generally considered more secure because it traverses the ISP over an MPLS network. Customers with additional security concerns should encrypt communications using IPsec, TLS, or other application-level encryption technologies, such as BitLocker when moving Virtual Hard Disk (VHD) files.

For more information on virtual network gateway configuration, see Configure a Virtual Network Gateway in the Management Portal.

2.2 Security Management and Threat Defense

Securing remote management of VMs

Administrators can create a VM using either the Azure Management Portal or Windows PowerShell.

When an administrator uses the Azure Management Portal to create a VM, Remote Desktop Protocol (RDP) and remote Windows PowerShell ports are opened by default. The Azure Management Portal then assigns RDP and remote Windows PowerShell random port numbers to reduce the chances of a password dictionary attack.

When an administrator uses Windows PowerShell to create a VM, RDP and remote Windows PowerShell ports must be *explicitly* opened.

- The administrator can choose to keep the RDP and remote Windows PowerShell ports open to the Internet, but at a minimum, the administrator should secure the accounts allowed to create RDP and remote Windows PowerShell connections with strong passwords.

- The administrator should also consider using the general options to secure inbound communication from the Internet mentioned above.

Protecting against DDoS

To protect Azure platform services, Microsoft provides a distributed denial-of-service (DDoS) defense system that is part of Azure's continuous monitoring process, and is continually improved through penetration-testing. Azure's DDoS defense system is designed to not only withstand attacks from the outside, but also from other Azure tenants:

1. Network-layer high volume attacks. These attacks choke network pipes and packet processing capabilities by flooding the network with packets. The Azure DDoS defense technology provides detection and mitigation techniques such as SYN cookies, rate limiting, and connection limits to help ensure that such attacks do not impact customer environments.
2. Application-layer attacks. These attacks can be launched against a customer VM. Azure does not provide mitigation or actively block network traffic affecting individual customer deployments, because the infrastructure does not interpret the expected behavior of customer applications. In this case, similar to on-premises deployments, mitigations include:
    - Running multiple VM instances behind a load-balanced Public IP address.
    - Using firewall proxy devices such as Web Application Firewalls (WAFs) that terminate and forward traffic to endpoints running in a VM. This provides some protection against a broad range of DoS and other attacks, such as low-rate, HTTP, and other application-layer threats. Some virtualized solutions, such as Barracuda Networks, are available that perform both intrusion detection and prevention.
    - Web Server add-ons that protect against certain DoS attacks.
    - Network ACLs, which can prevent packets from certain IP addresses from reaching VMs.

If a customer determines that their application is under attack, they should contact Azure Customer Support immediately to receive assistance. Azure Customer Support personnel prioritizes these types of requests.

Securing internal VM names with internal DNS

To address VMs within a cloud service by name, Azure provides an internal DNS service. VM names are resolved to private IP addresses within a cloud service while maintaining privacy across cloud services, even within the same subscription.

The private IP addresses assigned to both Cloud Service roles and Virtual Machines can change during a repair of cloud infrastructure. Because of this possibility, communications between roles within an Azure hosted service must be resolved via DNS name, and not by IP address. The one exception to this rule is when virtual networks are being used for custom IP address spaces. In those cases, IP addresses are static. Also, as private IP addresses can change, the DNS Time-to-Live (TTL) values of the DNS responses should be honored in the client.

Isolating VMs within a virtual network for defense-in-depth

The administrator can segment intranet traffic at the network layer within a Virtual Network by using Network Security Groups. Network Security Groups can be applied to a subnet in a Virtual Network.

Figure 4 depicts a sample multi-tier application deployment.

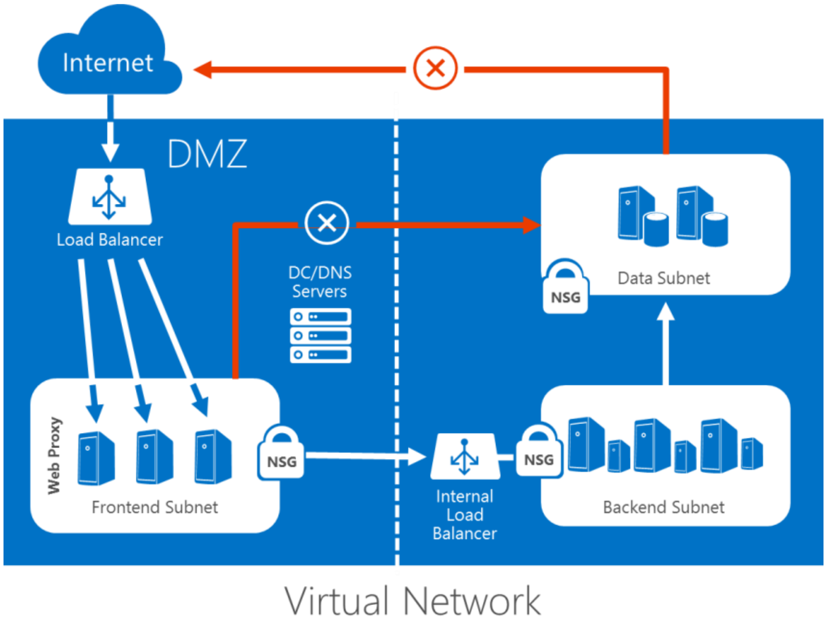

Figure 4: Using Network Security Groups in a multi-tier application VNET.

In addition to segmenting intranet traffic, Network Security Groups (NSGs) can also control traffic going to and coming from the Internet.

Network Security Groups can be applied to a VM or a subnet, and in some cases to both. Some important aspects of the NSGs include:

- NSG rules contain a 5-tuple (Source IP, Source port, Destination IP, Destination port, protocol), as shown in Figure 5.
- NSG rules are stateful. Thus, if there is an inbound rule that allows traffic on a port, then a matching rule on the outbound side is not required for the packets to flow on the same port.
- Every NSG contains default rules that allow connectivity within the Virtual Network and outbound access to the Internet. The customer administrator can override the default rules.
- NSG processes rules based on priority. Rules with small (meaning higher priority) values are processed before rules with larger (meaning lower priority) values.

- Azure provides default tags such as INTERNET and VIRTUAL_NETWORK that refer to the Public IP address space outside the Virtual Network and the customer's entire network address space, respectively. The tags can be used as part of an access control rule.

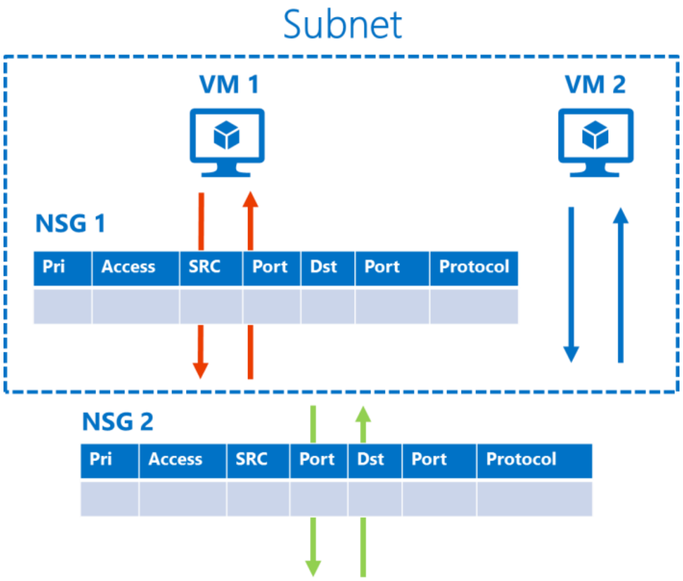

Figure 5: 5-tuple rules processing in a Virtual Network.

Securing communications from VMs to Microsoft Azure SQL Database

Microsoft Azure SQL Database also provides a built-in firewall to filter incoming traffic. Initially, all communications with the SQL Database are blocked. To enable communications with the database, the administrator must define firewall rules in Azure SQL Database allowing the public IP address of the VM in Azure to communicate with the data source.

Additionally, the administrator must update IP address ACLs any time the public virtual addresses change. This can result in service failures, and puts additional burden on the administrator. In addition, public virtual IP addresses can change after compute resources are deallocated when a VM is shut down, or after a deployment is deleted.

However, using an in-place upgrade enables administrators to deploy new versions of their service without the public IP addresses of the VMs changing.

For more information on how to configure IP ACLs, see About Network Access Control Lists. To learn about configuring the SQL Database firewall to specify rules at both the server-level and database-level, refer to the following articles:

- Microsoft Azure SQL Database Firewall
- sp_set_firewall_rule (Microsoft Azure SQL Database)

3 Guidelines for Securing Azure Cloud Services

The above guidelines for Azure VMs and VNETs also apply to Azure Cloud Service Web roles and Worker roles.

As with VMs, every Cloud Service role created through the Azure Portal has inbound traffic flow blocked from both the Internet and remote management ports by default. When an administrator enables a role for Remote Desktop Server, the RDP port is opened. RDP port numbers are assigned using a random number (as shown in Figure 6) to reduce the chances of a broad scanning/password dictionary attack.

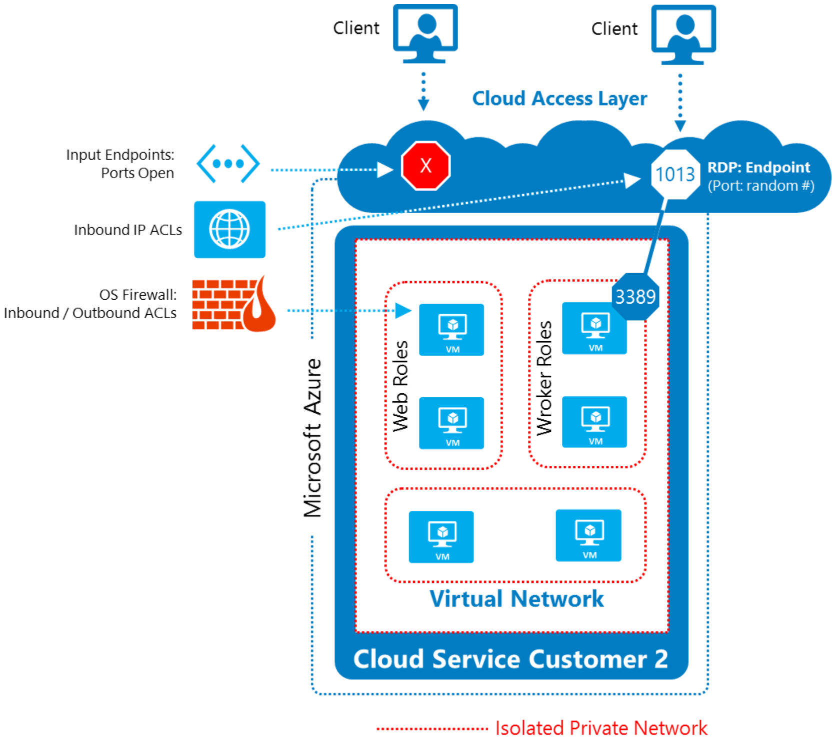

Figure 6. Topology of Azure Cloud Services Virtual Networks.

Customers can choose to keep the RDP ports open to the Internet, but, at a minimum, should secure the role with accounts that use strong passwords. When using RDP, enable the RDP port through the Azure Portal, but then disable the port after use.

Similarly, only open other ports by defining them in the Endpoints element of the WebRole or WorkerRole schema in the service definition file (.csdef). For more information, see the WebRole Schema and WorkerRole Schema guidance on MSDN.

4 Summary

The following table provides pointers to additional information on how to configure Azure Virtual Networks for increased security.

CAPABILITY	TECHNOLOGY	RECOMMENDATION	MORE INFORMATION
ENCRYPTION	SSL / TLS	Secure inbound Internet communications to VMs	http://aka.ms/azure-ssl
	IPsec	Configure a VPN for secure cross-premises connectivity	http://aka.ms/azure-ipsec-vpn
HOST FIREWALL	IP Address ACLs	Create input endpoints to control traffic flow to VMs	http://msdn.microsoft.com/library/azure/dn376541.aspx
ISOLATION	ExpressRoute	Protect remote network traffic with a dedicated fiber link	http://azure.microsoft.com/en-us/services/expressroute/
	Network Security Groups		http://azure.microsoft.com/blog/2014/11/04/network-security-groups/
	Instance level Public IP Addresses		http://azure.microsoft.com/blog/2014/10/22/instance-level-public-ip-address/
GUEST FIREWALL	Windows Firewall	Configure firewalls in VMs to allow only necessary endpoints	http://aka.ms/azure-vm-endpoints
	Application Firewall	Deploy a third-party Web application firewall for additional IDS/IPS and DDoS protection	http://aka.ms/azure-barracuda
MULTIPLE NICS	Multiple NICs and Network Virtual Appliances		http://aka.ms/azure-multi-nic

5 References and Further Reading

The following resources provide general information about Azure and related Microsoft services, as well as specific items referenced in the main text:

- Azure Home – general information and links about Azure
    - http://azure.microsoft.com
- Azure Documentation Center – developer guidance and information
    - http://azure.microsoft.com/en-us/documentation/
- Azure Trust Center
    - http://azure.microsoft.com/en-us/support/trust-center/
- Microsoft Security Response Center [where Microsoft security vulnerabilities, including issues with Azure, can be reported]
    - http://www.microsoft.com/security/msrc/default.aspx
    - Or via email to secure@microsoft.com
- Azure Network Services
    - http://msdn.microsoft.com/library/windowsazure/gg433091
- Companion video for the Windows Azure Network Security white paper
    - http://channel9.msdn.com/Blogs/Windows-Azure/Companion-video-for-the-Windows-Azure-Network-Security-white-paper?format=html5

6 Appendix: Azure Network Security Internals

This section provides additional technical depth for the internals of Azure network security, as well as some guidelines for securing services built on Azure.

6.1 Layers of Protection

Figure 7 shows the different layers of network protection for Azure.

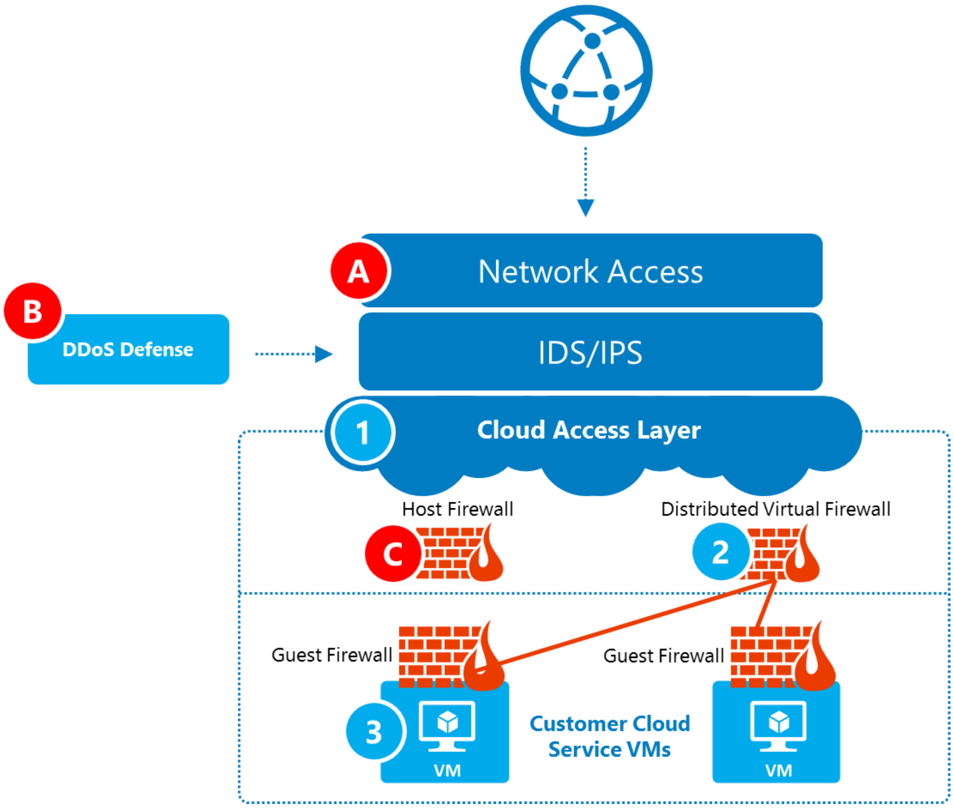

Figure 7. Layers of defense for protecting customers and Azure infrastructure.

There are two (2) separate areas of protection: infrastructure protection and customer protection.

1. Azure platform services infrastructure protection:
    a. Layer A: The Network Access Layer isolates Azure's private network from the Internet.
    b. Layer B: Azure's DDoS/DOS/IDS Layer uses different methods and technologies than on-premises deployments to achieve similar security goals.
    c. Layer C: Host firewalls protect all the hosts, and the VLANs provide additional protection for key assets.
    d. Layer D: Conformance with security and privacy requirements includes two-factor authentication for operators.

2. Customer protection:
    a. Layers 1-2: The distributed firewall isolates one customer's deployment from other deployments at the network level. Multiple deployments can be put inside a virtual network, and each virtual network is isolated from other virtual networks. The Cloud Access Layer acts as the gateway from the Internet into this isolated network, and provides load balancing, NAT, and firewall capabilities that can be configured by the customer.
    b. Layer 3: The virtual network can be managed similar to an on-premises private network.
        i. Inside the VM: Firewalls, IDS, and DoS solutions can be deployed on the guest OS of the VM.
        ii. Virtual network appliances: Proxy-based devices (such as WAFs) that terminate and then forward traffic to endpoints and can run in a VM can provide protection against an even broader range of DoS and other attacks (e.g. low-rate, HTTP, and application-layer threats). If there is a need for bridge-mode security appliances, the administrator can connect the Azure virtual network to on-premises network (such as via VPN) and send the traffic through the devices inside the organization.

6.2 Isolation

Azure provides network isolation for each deployment. Using input endpoints, customers decide which ports can be accessed from the Internet.

- Traffic between VMs always traverses through trusted packet filters.
    a. Protocols such as Address Resolution Protocol (ARP), Dynamic Host Configuration Protocol (DHCP), and other OSI Layer-2 traffic from a VM are controlled using rate-limiting and anti-spoofing protection.
    b. VMs cannot capture any traffic on the network that is not destined to it.
- Customer VMs cannot send traffic to Azure's private interfaces or other customers' VMs, or Azure infrastructure services themselves. Customer VMs can only communicate with other VMs owned or controlled by the same customer and with Azure infrastructure service endpoints meant for public communications.
- When customers put VMs on a virtual private network, those VMs get their own address spaces that are completely invisible, and hence, not reachable from VMs outside of a deployment or virtual network (unless configured to be visible via public IP addresses). Customer environments are open only through the ports they specify for public access; if the VM is defined to have a public IP address, then all ports are open for public access.