Microsoft Azure Network Security

Microsoft

Abstract

This document provides guidance on securing network communications for applications deployed in Microsoft Azure, and will help customers understand how best to protect their virtual infrastructure and data.

The intended audience for this whitepaper includes:

- IT and network administrators interested in deploying applications in Microsoft Azure
- Developers interested in creating applications that run in Microsoft Azure
- Technical decision makers (TDM) considering Microsoft 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 2, Published October 2014

Table of Contents

1 Overview

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

Azure's network services have been designed for flexibility, availability, resiliency, security, and integrity. This white paper uncovers some of Azure's inner-workings, and provides insights on how customers can take advantage of the platform's native security features to best protect their information assets.

- Section 2 of this white paper provides details on securing Virtual Networks (VNET) for Azure VMs, (also known as Infrastructure as a Service, or IaaS).
- Section 3 of this paper provides information on the differences between network security for VMs and network security for Azure Cloud Services (also known as Platform as a Service, or PaaS).

2 Guidelines for Securing Azure Virtual Machines

Azure is a unified, multi-tenant platform that uses shared infrastructure to support millions of simultaneous customers across more than 80 datacenters around the world. With hundreds of millions of active Virtual Machines, network traffic security and confidentiality are critical.

Azure Virtual Networks use a combination of logical isolation, firewalls, access controls, authentication, and encryption to protect customer data in-transit. Comprehensive information security policies and processes are part of Microsoft's datacenter operations, critical measures of which are regularly audited by third-parties (through industry control frameworks such as ISO 27001, SOC 1, SOC 2, and others). This includes both the physical and virtual aspects of Azure infrastructure.

In the traditional datacenter model with which most customers are familiar, a company's Information Technology (IT) organization retains ultimate control over networked systems, including physical access to networking equipment. Deployment, configuration, and management are handled by individuals with a direct relationship to the company; network administrators can physically alter the network topology, change router settings, deploy 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. An Azure subscriber cannot walk into a cloud provider datacenter and rewire a server rack, but they can do the equivalent within their cloud environment through a number of different virtual mechanisms, including Guest OS firewalls, VNET Gateway configuration, and Virtual Private Network (VPN). This approach enables customers to rely on the fundamental logical security capabilities delivered by Azure infrastructure as they build out a secure and compliant cloud solution.

2.1 Private Networks

Fundamental to any cloud architecture is the isolation provided to customers. The distributed virtual firewall in Azure helps customer's private network traffic remain separated from other customers' data.

In addition, a customer subscription can contain multiple logically isolated private networks:

- **Deployment network**: Each deployment can be isolated from others 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 other virtual networks. Multiple deployments (inside the same subscription) can be placed on the same virtual network, and then allowed to communicate with each other through private IP addresses.

An example of such a topology is shown in Figure 1 below.

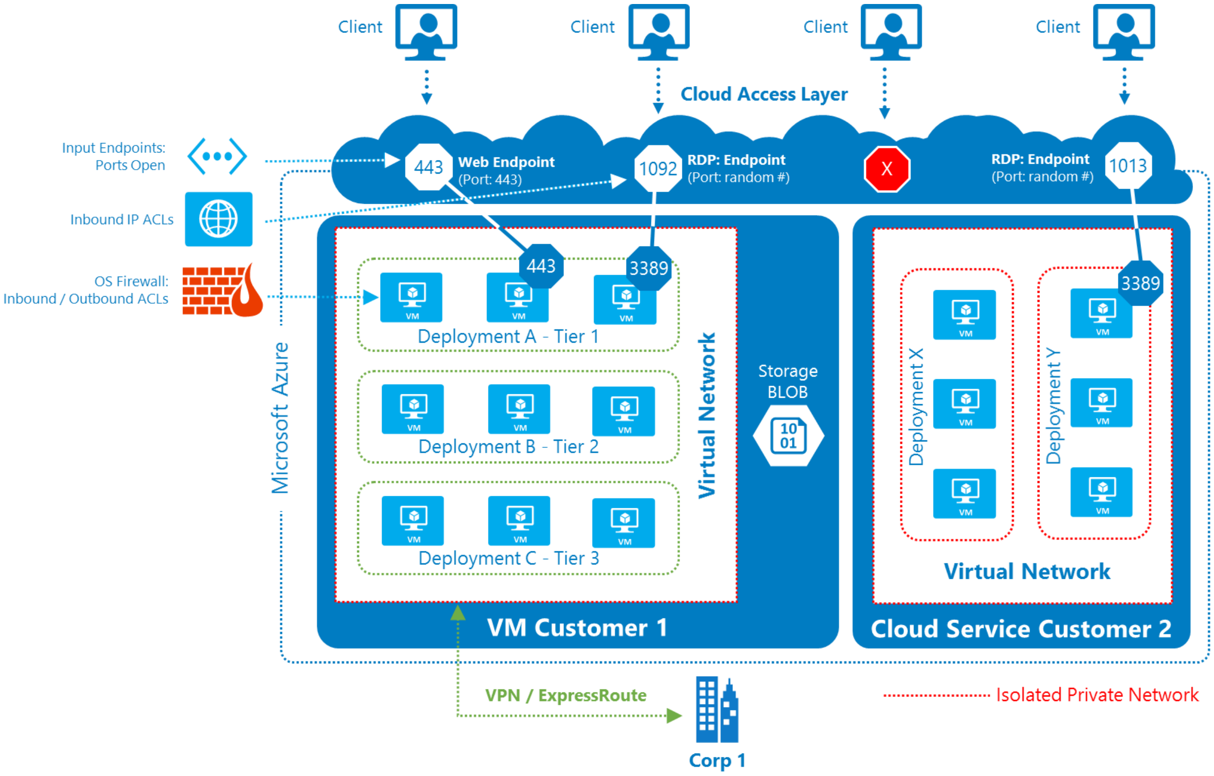

Figure 1. An example of isolated multi-tier IaaS applications hosted within Microsoft Azure.

This short video provides details on how Microsoft Azure provides security for virtual machines.

These isolated private networks can be managed by administrators similarly to how private networks are managed in an on-premises network.

The Cloud Access Layer in Azure is similar to the edge of a corporate network that faces the Internet. This layer includes a firewall, load-balancing and network address translation (NAT) functionality that is managed by the customer administrator.

By default, Virtual Machines inside the private network do not receive inbound traffic from the Internet.

- 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—traffic that could come from the Internet, as well as from other VMs inside Microsoft Azure.
- Alternatively, 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 sometimes referred to as unsolicited inbound traffic to distinguish it from inbound traffic that is a response to a request, also known as solicited inbound traffic.

Following are the principal network security considerations for administrators when deploying or migrating Virtual Machines to Azure.

Securing communications between VMs inside the private network

Virtual Machines inside a deployment network are allowed to communicate with each other via private IP addresses. Communication between VMs in multiple deployments of a subscription can be secured by using Virtual Networks.

If an application sends or receives sensitive data over an internal private network (e.g., a VPN), then it 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, a VM created through the Azure Management Portal has inbound traffic flow blocked from the Internet, except for remote management ports.

By configuring input endpoints, administrators can decide which VM ports can be accessed from the Internet via inbound traffic. Below are some configuration changes to help secure remote access at the network level from the Internet to ports on VMs or VNETs.

IT administrators can restrict access by:

- Defining input endpoints to only open ports that you need at the Cloud Access Layer. In addition, you can specify access control lists (ACLs) on input endpoints to control the source IPs from which the VM will allow traffic.
- Using a third-party proxy firewall (such as the Web Application Firewall Vx or NG Firewall Vx virtual appliances available from Barracuda Networks) that runs on a Virtual Machine to filter traffic to your 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 you open input endpoints, you should follow the same security model as if it 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 VMs may need to communicate between multiple subscriptions. VMs can be configured to communicate via public virtual IP addresses. IP ACLs would need to be configured on input endpoints to allow VMs to initiate connections only with each other.

IP address ACLs must be updated 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 de-allocated when a Virtual Machine is shut down, or after a deployment is deleted.

However, using in-place upgrade enables administrators to deploy new versions of their service without the public IP addresses of the VMs changing, and the administrator doesn't need to know what those addresses are.

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

Securing communications to on-premises networks

When your workloads require secure communications between the Azure Virtual Network and your on-premises systems, it is best to protect those channels using a Virtual Network Gateway (VNG). Two scenarios for deploying VNG are:

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

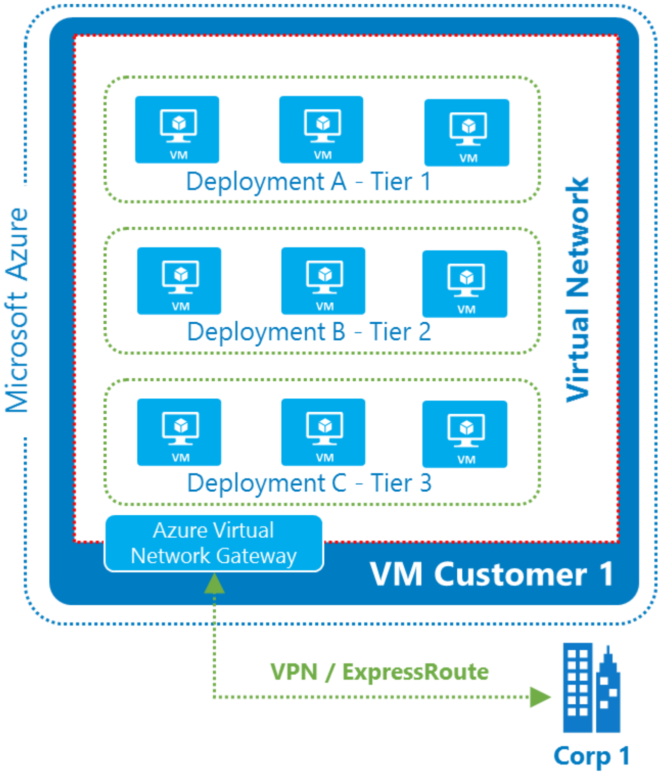

Figure 2. VPN connection between a corporate network and Microsoft Azure.

In this case for a VNET to VNET connection, you can create a virtual network and add the VMs of the application tiers to the virtual network, but you do not need to define any input endpoints:

- Remove the remote management input endpoints or lock them down using the guidance provided below to secure management endpoints.
- Configure the Virtual Network Gateway so that traffic destined for the corporate network flows through the VPN connection to the target servers / network devices on the corporate network.

2. **Public-facing Multi-tier Application**: A multi-tier application is 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.

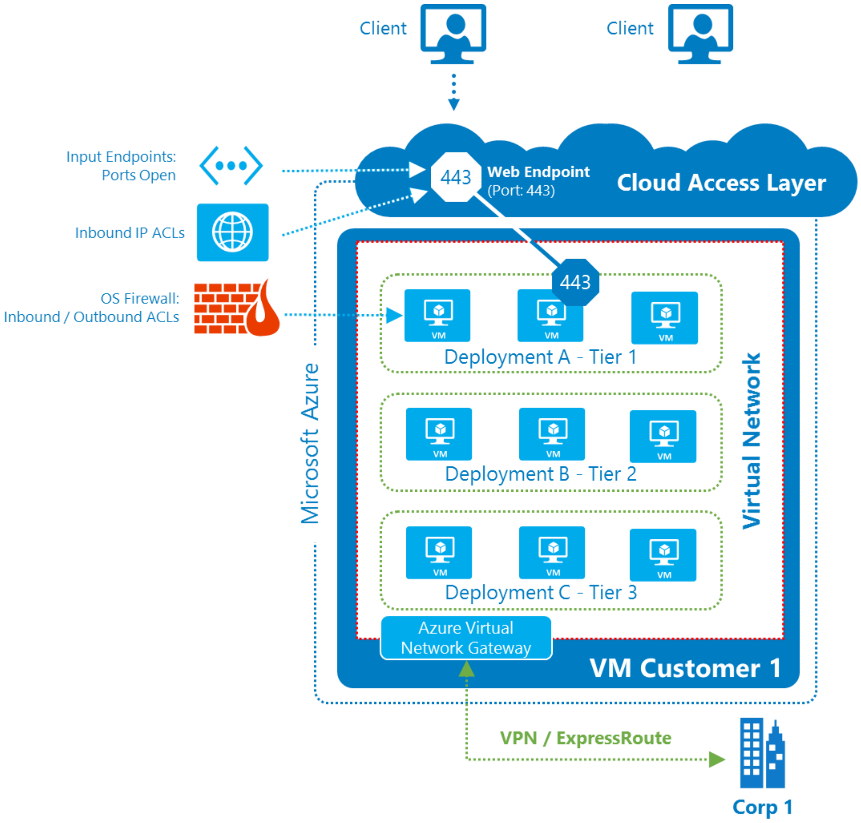

Figure 3. Addition of Internet-facing input endpoints to allow Internet access to the front-end tier.

In this case:

a. Create a virtual network with the appropriate VMs from each of the application tiers
b. Define input endpoints for the inbound Internet traffic for the VMs in the front-end tier
c. Remove the remote management input endpoints for all VMs or lock them down
d. 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 between the virtual network and the customer's VPN device (which can be either a hardware VPN device or a software VPN such as Windows Server 2012 Routing and Remote Access Services), and routes traffic appropriately.

Creating a virtual private network within Azure provides the ability to more securely extend your on-premises network into Azure. This connection can be either site-to-site VPN 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, then this traffic uses MPLS technology, which is considered safer compared to sending the traffic over the Internet. Customers with additional security concerns should ensure that this communication is encrypted using IPsec, SSL/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

By default, if a VM is created through the Azure Management Portal, Remote Desktop Protocol (RDP) and remote Windows PowerShell ports are opened. However, when a VM is created through Windows PowerShell, RDP and remote Windows PowerShell ports must be *explicitly* opened.

- The Management Portal assigns RDP and remote Windows PowerShell port numbers using a random number to reduce the chances of a password dictionary attack.
- You can choose to keep the RDP and remote Windows PowerShell ports open to the Internet, but at a minimum, secure the accounts allowed to create RDP and remote Windows PowerShell connections with strong passwords.
- Consider using the general options to secure inbound communication from the Internet mentioned above.

Protecting against DDoS

Threat mitigation and protection of customer environments is similar to that used in many on-premises datacenters.

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

The following are examples of several different kinds of DDoS attacks that the system focuses on:

1. Network-layer high volume attacks choke network pipes and packet processing capabilities. 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 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, impacts can be minimized by:
    - 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 Virtual Machine, providing protection against a broad range of DoS and other attacks (e.g. low-rate, HTTP, and application-layer threats). Some virtualized solutions available are also capable of both intrusion detection and prevention (such as Barracuda). Virtual appliances should work on Azure as long as they are certified by the vendor.
    - Web Server add-ons that protect against certain DoS attacks
    - Network ACLs which can prevent packets from certain IP addresses from reaching your deployment.

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

Securing internal VM names with internal DNS

To allow VMs within a cloud service to be addressed 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 repair. Due to this, 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 can be assumed to be 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

Application isolation helps prevent exploits on neighboring tenants or deployments from compromising your cloud data and operations. This includes front-end and back-end application-tier scenarios. For example, Virtual Machines in a back-end network or sub-network may allow only certain client computers to connect to a particular endpoint based on an IP address whitelist.

Following are options to isolate the tiers in an application at the network layer, and they are typically used in combination for defense-in-depth:

1) ACLs inside the guest OS firewall
    - Deploy the tiers in a single virtual network, and restrict access to machines using a guest OS firewall (e.g., Windows Firewall with Advanced Security).
2) Network ACLs on public IP addresses
    - Deploy the tiers in separate virtual networks, and restrict access between the tiers using network ACLs at public IP addresses.
3) Network ACLs at the corporate firewall
    - Deploy the tiers in separate virtual networks, and connect these virtual networks via VPN to your corporate network. Then set ACLs at the corporate firewall to restrict access between the tiers.
4) IPsec inside the guest OS
    - With IPsec integration, Windows Firewall with Advanced Security provides a simple way to enforce authenticated, end-to-end network communications.

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, firewall rules must be defined in Azure SQL Database allowing the public IP address of the VM in Microsoft Azure to communicate with the data source.

IP address ACLs must be updated 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 de-allocated when a Virtual Machine is shut down, or after a deployment is deleted.

However, using 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 the Internet by default, as well as remote management ports. When a role is enabled for remote desktop, the RDP port is opened. RDP port numbers are assigned using a random number (as shown in Figure 4) to reduce the chances of a broad scanning/password dictionary attack.

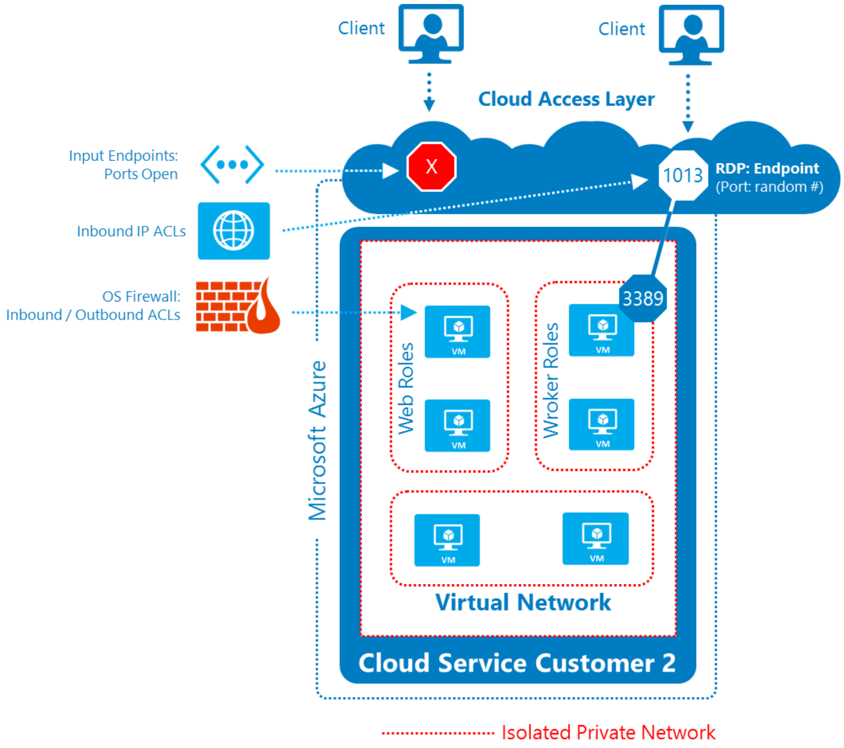

Figure 4. Topology of Azure Cloud Services virtual networks.

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://azure.microsoft.com/en-us/documentation/articles/cloud-services-configure-ssl-certificate/
	IPsec	Configure a VPN for secure cross-premises connectivity	http://msdn.microsoft.com/en-us/library/azure/dn133798.aspx
HOST FIREWALL	IP 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/
GUEST FIREWALL	Windows Firewall	Configure firewalls in VMs to allow only necessary endpoints	http://azure.microsoft.com/en-us/documentation/articles/virtual-machines-set-up-endpoints/
	Application Firewall	Deploy a third-party web application firewall for additional IDS/IPS and DDoS protection	http://azure.microsoft.com/en-us/gallery/virtual-machines/Barracuda/barracudawebapplicationfirewallwaf78/

5 References and Further Reading

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

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

6 Appendix: Microsoft 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 5 shows the different layers of network protection for Azure.

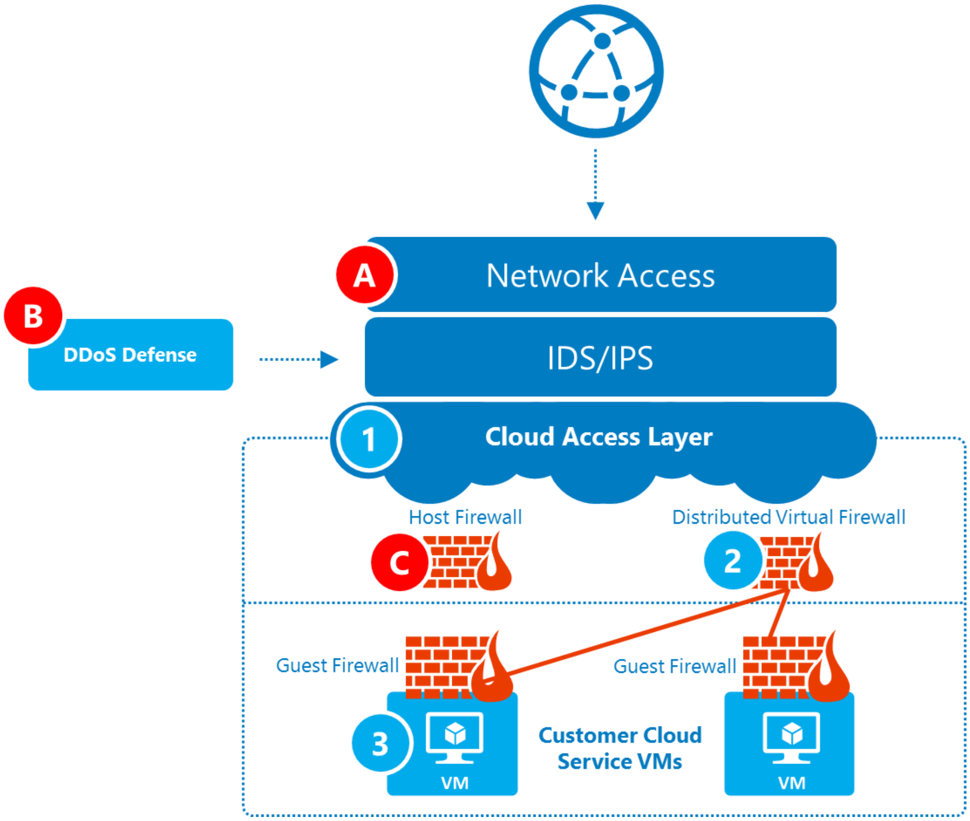

Figure 5. Layers of defense for protecting customers and Azure infrastructure.

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. The cloud access layer 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 Virtual Machine 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, you can connect the Azure virtual network to your on-premises network (such as via VPN) and send the traffic through the devices inside your organization.

6.2 Isolation

Microsoft 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.