

Microsoft Cloud Networking for Enterprise Architects

What IT architects need to know about networking in Microsoft cloud services and platforms

This topic is 1 of 6 in a series 1 2 3 4 5 6

Evolving your network for cloud connectivity

Cloud migration changes the volume and nature of traffic flows within and outside a corporate network. It also affects approaches to mitigating security risk.



Before the cloud

Most networking infrastructure investments were spent on ensuring available, reliable, and performant connectivity to on-premises datacenters. For many organizations, Internet connectivity was not critical for internal business operations. Network boundaries were primary defenses against security breaches.

Network infrastructure investments begin with connectivity. Additional investments depend on the category of cloud service.



After the cloud

With new and migrated productivity and IT workloads running in the cloud, infrastructure investments shift from on-premises datacenters to Internet connectivity, which is now critical for internal business operations. Federated connectivity shifts security strategy to protecting identities and data as they flow through the network and points of connectivity to Microsoft cloud services.

SaaS Software as a Service

Microsoft SaaS services include Office 365, Microsoft Intune, and Microsoft Dynamics 365. Successful adoption of SaaS services by users depends on highly-available and performant connectivity to the Internet, or directly to Microsoft cloud services.

Network architecture focuses on reliable, redundant connectivity and ample bandwidth. Ongoing investments include performance monitoring and tuning.

Azure PaaS Platform as a Service

In addition to the investments for Microsoft SaaS services, multi-site or geographically distributed PaaS applications might require architecting Azure Application Gateway or Azure Traffic Manager to distribute client traffic. Ongoing investments include performance and traffic distribution monitoring and failover testing.

Azure IaaS Infrastructure as a Service

In addition to the investments for Microsoft SaaS and PaaS services, running IT workloads in IaaS requires the design and configuration of Azure virtual networks that host virtual machines, secure connectivity to applications running on them, routing, IP addressing, DNS, and load balancing. Ongoing investments include performance and security monitoring and troubleshooting.

Microsoft 365 is a combination of Office 365, Enterprise Management + Security (EMS), and Windows 10. Microsoft 365 combines multiple SaaS and Azure services for a complete, intelligent solution that empowers everyone to be creative and work together securely.

Areas of networking investment for success in the cloud

Enterprise organizations benefit from taking a methodical approach to optimizing network throughput across your intranet and to the Internet. You might also benefit from an ExpressRoute connection.

Optimize intranet connectivity to your edge network

Over the years, many organizations have optimized intranet connectivity and performance to applications running in on-premises datacenters. With productivity and IT workloads running in the Microsoft cloud, additional investment must ensure high connectivity availability and that traffic performance between your edge network and your intranet users is optimal.

Optimize throughput at your edge network

As more of your day-to-day productivity traffic travels to the cloud, you should closely examine the set of systems at your edge network to ensure that they are current, provide high availability, and have sufficient capacity to meet peak loads.

For a high SLA to Microsoft cloud services, use ExpressRoute

Although you can utilize your current Internet connection from your edge network, traffic to and from Microsoft cloud services must share the pipe with other intranet traffic going to the Internet. Additionally, your traffic to Microsoft cloud services is subject to Internet traffic congestion.

For a high SLA and the best performance, use ExpressRoute, a dedicated WAN connection between your network and Azure, Office 365, Dynamics 365, or all three.

ExpressRoute can leverage your existing network provider for a dedicated connection. Resources connected by ExpressRoute appear as if they are on your WAN, even for geographically-distributed organizations.

[ExpressRoute for Office 365](#)
[ExpressRoute for Azure](#)

The scope of network investments depend on the category of cloud service. Investing across Microsoft's cloud maximizes the investments of networking teams. For example, investments for SaaS services apply to all categories.

Investment area	SaaS	PaaS	IaaS
Architect reliable, redundant Internet connectivity with ample bandwidth	■	■	■
Monitor and tune Internet throughput for performance	■	■	■
Troubleshoot Internet connectivity and throughput issues	■	■	■
Design Azure Traffic Manager to load balance traffic to different endpoints		■	■
Architect reliable, redundant, and performant connectivity to Azure virtual networks			■
Design secure connectivity to Azure virtual machines			■
Design and implement routing between on-premises locations and virtual networks			■
Architect and implement load balancing for internal and Internet-facing IT workloads			■
Troubleshoot virtual machine connectivity and throughput issues			■

Microsoft Cloud Networking for Enterprise Architects

What IT architects need to know about networking in Microsoft cloud services and platforms

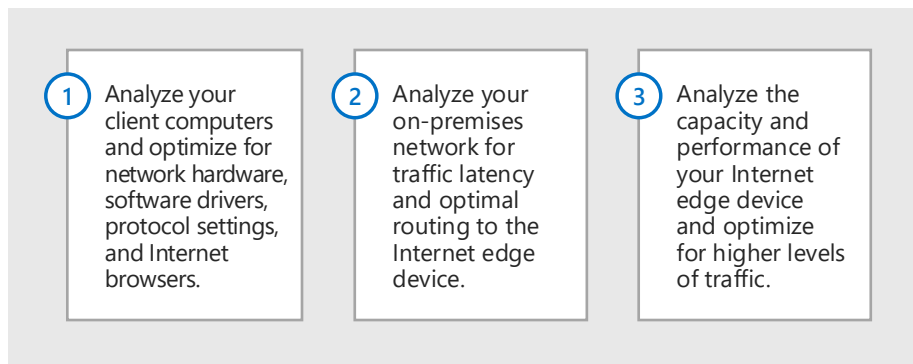
This topic is 2 of 6 in a series 1 2 3 4 5 6

Common elements of Microsoft cloud connectivity

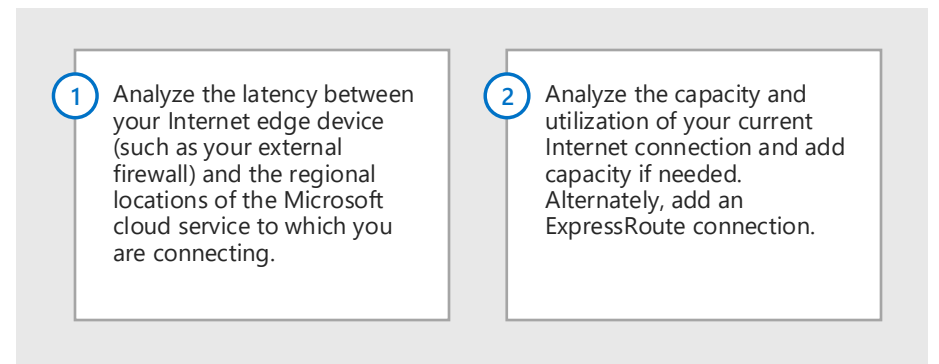
Integrating your networking with the Microsoft cloud provides optimal access to a broad range of services.

Steps to prepare your network for Microsoft cloud services

On-premises network

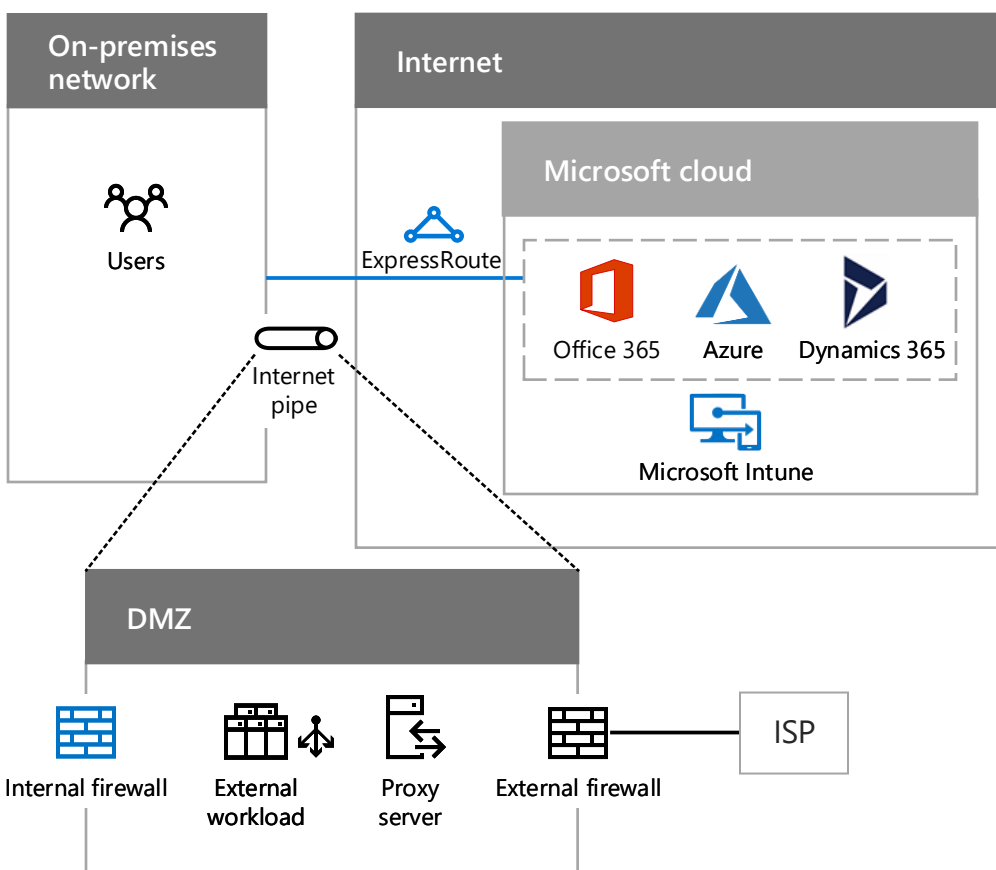


Internet

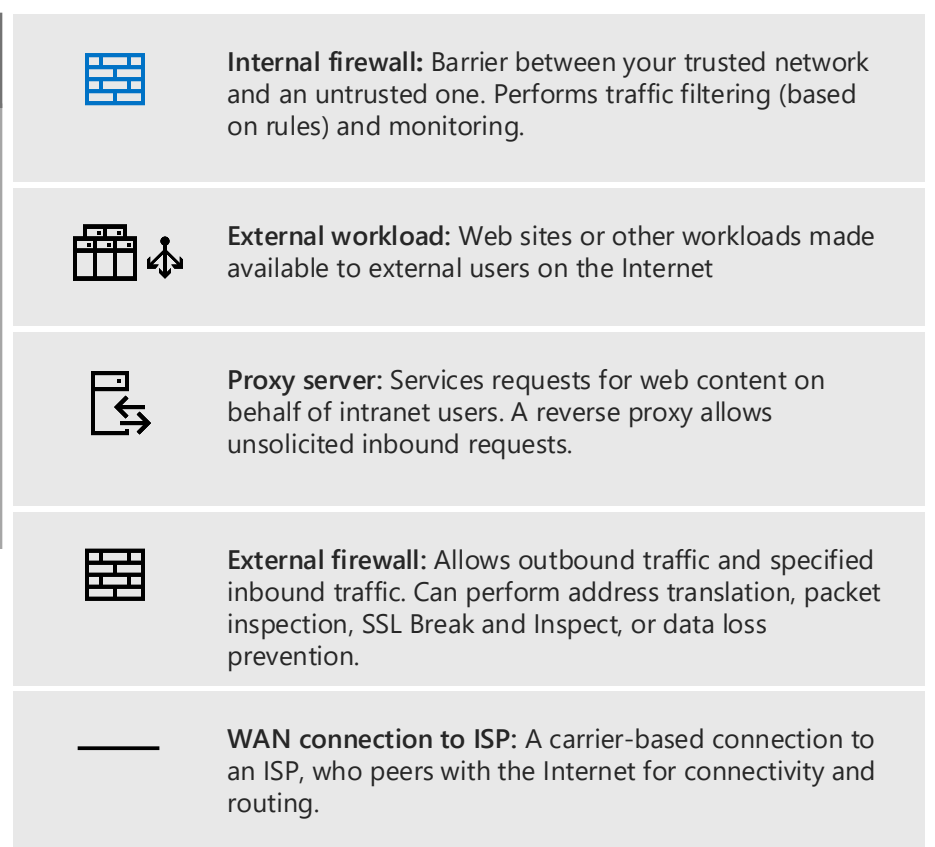


Microsoft cloud connectivity options

Use your existing Internet pipe or an ExpressRoute connection to Office 365, Azure, and Dynamics 365.



Components of a typical DMZ



Areas of networking common to all Microsoft cloud services

Intranet performance

Performance to Internet-based resources will suffer if your intranet, including client computers, is not optimized.

Edge devices

Devices at the edge of your network are egress points and can include Network Address Translators (NATs), proxy servers (including reverse proxies), firewalls, intrusion detection devices, or a combination.

Internet connection

Your WAN connection to your ISP and the Internet should have enough capacity to handle peak loads.

You can also use an ExpressRoute connection.

Internet DNS

Use A, AAAA, CNAME, MX, PTR and other records to locate Microsoft cloud or your services hosted in the cloud. For example, you might need a CNAME record for your app hosted in Azure PaaS.

Microsoft Cloud Networking for Enterprise Architects

What IT architects need to know about networking in Microsoft cloud services and platforms

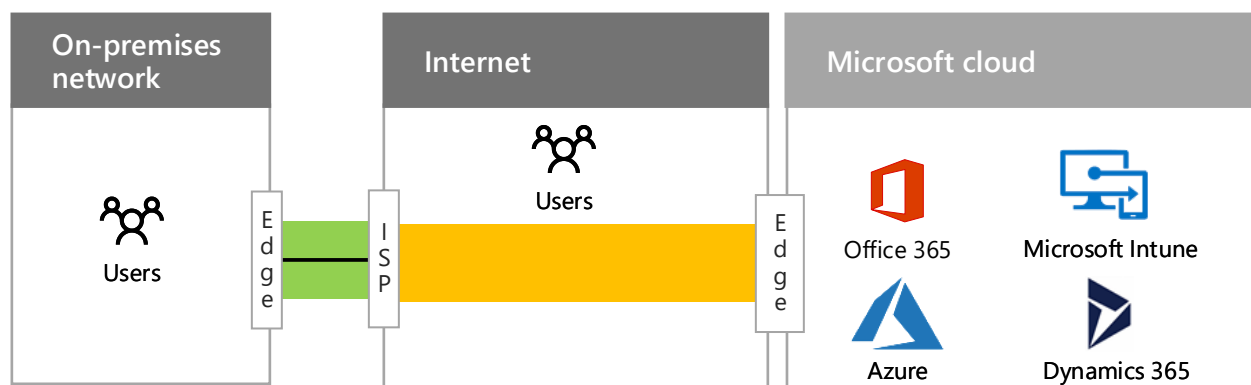
This topic is 3 of 6 in a series 1 2 3 4 5 6

ExpressRoute for Microsoft cloud connectivity

ExpressRoute provides a private, dedicated, high-throughput network connection to Microsoft's cloud.

ExpressRoute to the Microsoft cloud

Without ExpressRoute

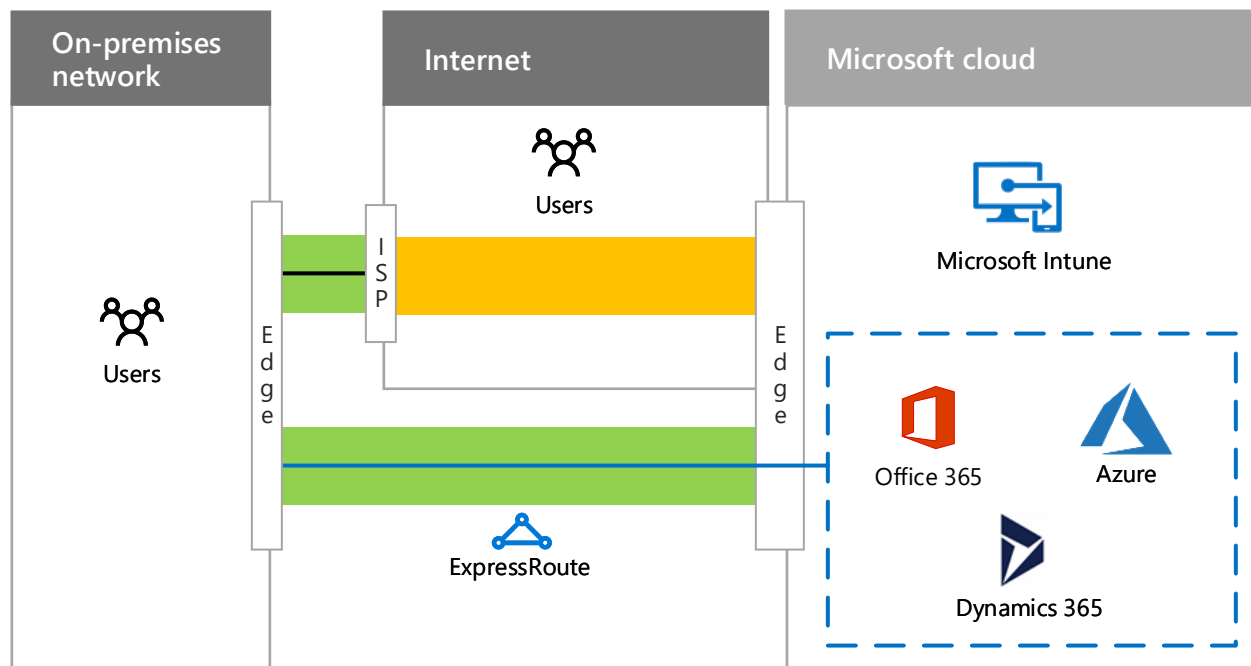


With an Internet connection, the only part of the traffic path to the Microsoft cloud that you can control (and have a relationship with the service provider) is the link between your on-premises network edge and your ISP (shown in green).

The path between your ISP and the Microsoft cloud edge is a best-effort delivery system subject to outages, traffic congestion, and monitoring by malicious users (shown in yellow).

Users on the Internet, such as roaming or remote users, send their traffic to the Microsoft cloud over the Internet.

With ExpressRoute



With an ExpressRoute connection, you now have control, through a relationship with your service provider, over the entire traffic path from your edge to the Microsoft cloud edge. This connection can offer predictable performance and a 99.9% uptime SLA.

You can now count on predictable throughput and latency, based on your service provider's connection, to Office 365, Azure, and Dynamics 365 services. ExpressRoute connections to Microsoft Intune are not supported at this time.

Traffic sent over the ExpressRoute connection is no longer subject to Internet outages, traffic congestion, and monitoring.

Users on the Internet, such as roaming or remote users, still send their traffic to the Microsoft cloud over the Internet. One exception is traffic to an intranet line of business application hosted in Azure IaaS, which is sent over the ExpressRoute connection via a remote access connection to the on-premises network.

Even with an ExpressRoute connection, some traffic is still sent over the Internet, such as DNS queries, certificate revocation list checking, and content delivery network (CDN) requests.

See these additional resources for more information: [ExpressRoute for Office 365](#) | [ExpressRoute for Azure](#)

Advantages of ExpressRoute for Azure

Predictable performance

With a dedicated path to the edge of the Microsoft cloud, your performance is not subject to Internet provider outages and spikes in Internet traffic. You can determine and hold your providers accountable to a throughput and latency SLA to the Microsoft cloud.

Data privacy for your traffic

Traffic sent over your dedicated ExpressRoute connection is not subject to Internet monitoring or packet capture and analysis by malicious users. It is as secure as using Multiprotocol Label Switching (MPLS)-based WAN links.

High throughput connections

With wide support for ExpressRoute connections by exchange providers and network service providers, you can obtain up to a 10 Gbps link to the Microsoft cloud.

Lower cost for some configurations

Although ExpressRoute connections are an additional cost, in some cases a single ExpressRoute connection can cost less than increasing your Internet capacity at multiple locations of your organization to provide adequate throughput to Microsoft cloud services.


An ExpressRoute connection is not a guarantee of higher performance in every configuration. It is possible to have lower performance over a low-bandwidth ExpressRoute connection than a high-bandwidth Internet connection that is only a few hops away from a regional Microsoft datacenter.

For the latest recommendations for using ExpressRoute with Office 365, see [ExpressRoute for Office 365](#).

[Continued on next page](#)


ExpressRoute connectivity models

Co-located at a cloud exchange



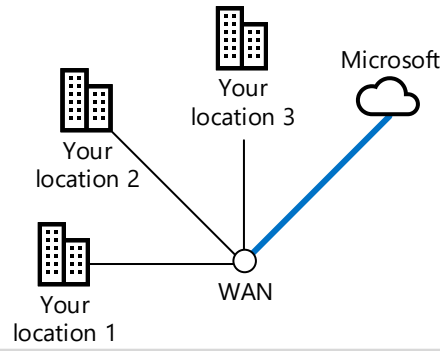
If your datacenter is co-located in a facility with a cloud exchange, you can order a virtual cross-connection to the Microsoft cloud through the co-location provider's Ethernet exchange.

Point-to-point Ethernet



If your datacenter is located on your premises, you can use a point-to-point Ethernet link to connect to the Microsoft cloud.

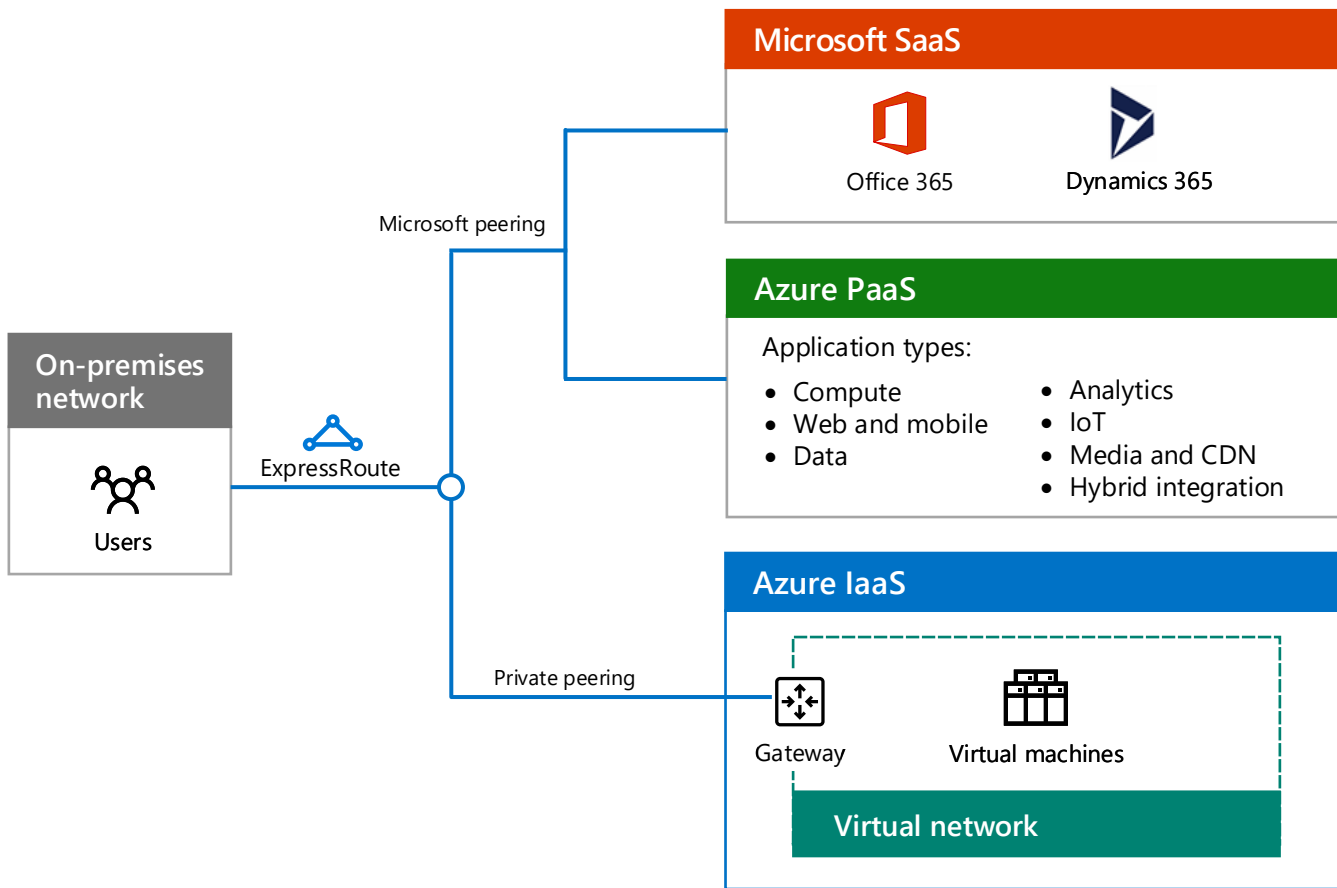
Any-to-any (IP VPN) connection



If you are already using an IP VPN (MPLS) provider to connect the sites of your organization, an ExpressRoute connection to the Microsoft cloud acts like another location on your private WAN.

ExpressRoute peering relationships to Microsoft cloud services

A single ExpressRoute connection supports up to two different Border Gateway Protocol (BGP) peering relationships to different parts of the Microsoft cloud. BGP uses peering relationships to establish trust and exchange routing information.



Microsoft peering

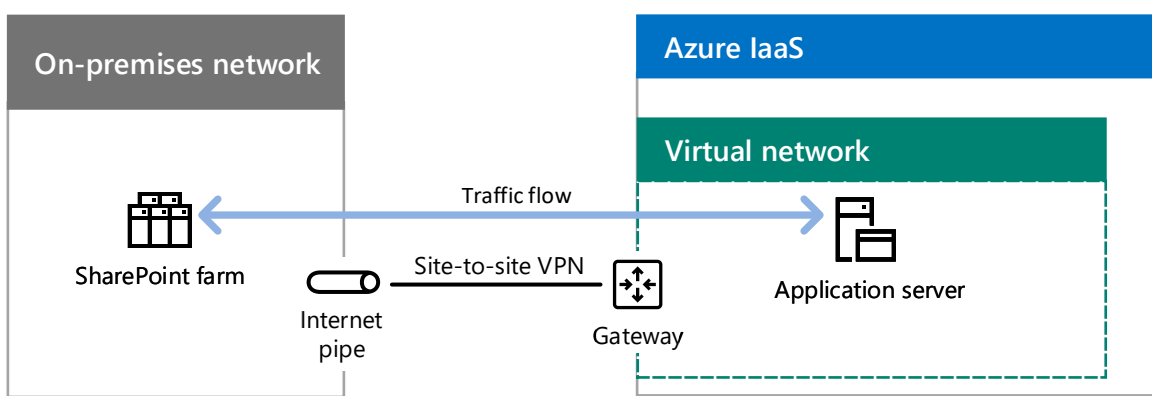
- Is from a router in your DMZ to the public addresses of Office 365, Dynamics 365, and Azure services.
- Supports bidirectional-initiated communication.

Private peering

- Is from a router on the edge of your organization network to the private IP addresses assigned to your Azure VNETs.
- Supports bidirectional-initiated communication.
- Is an extension of your organization network to the Microsoft cloud, complete with internally-consistent addressing and routing.

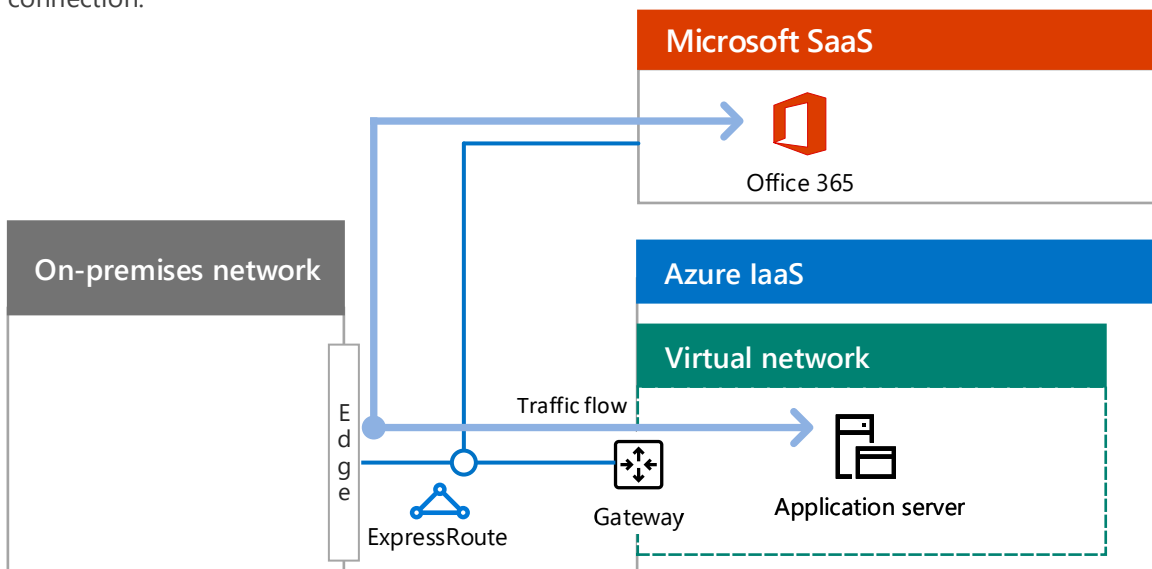
Example of application deployment and traffic flow with ExpressRoute

How traffic travels across ExpressRoute connections and within the Microsoft cloud is a function of the routes at the hops of the path between the source and the destination and application behavior. Here is an example of an application running on an Azure virtual machine that accesses an on-premises SharePoint farm over a site-to-site VPN connection.



The application locates the IP address of the SharePoint farm using the on-premises DNS and all traffic goes over the site-to-site VPN connection.

This organization migrated their on-premises SharePoint farm to SharePoint Online in Office 365 and deployed an ExpressRoute connection.



With the Microsoft and private peering relationships:

- From the Azure gateway, on-premises locations are available across the ExpressRoute connection.
- From the Office 365 subscription, public IP addresses of edge devices, such as proxy servers, are available across the ExpressRoute connection.
- From the on-premises network edge, the private IP addresses of the Azure VNet and the public IP addresses of Office 365 are available across the ExpressRoute connection.

When the application accesses the URLs of SharePoint Online, it forwards its traffic across the ExpressRoute connection to a proxy server in the edge.

When the proxy server locates the IP address of SharePoint Online, it forwards the traffic back over the ExpressRoute connection. Response traffic travels the reverse path. The result is hair pinning, a consequence of the routing and application behavior.

ExpressRoute and Microsoft's cloud network

With ExpressRoute

How traffic travels between your organization network and a Microsoft datacenter is a combination of:

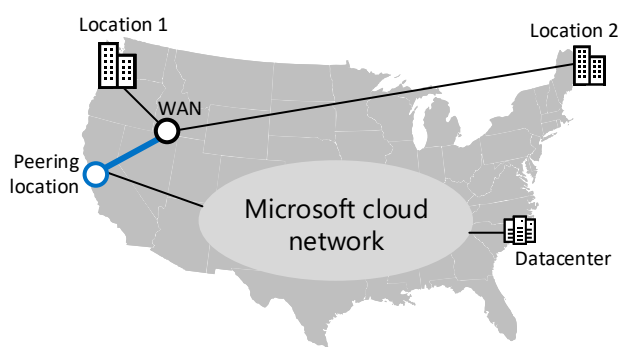
- Your locations.
- Microsoft cloud peering locations (the physical locations to connect to the Microsoft edge).
- Microsoft datacenter locations.

Microsoft datacenter and cloud peering locations are all connected to the Microsoft cloud network.

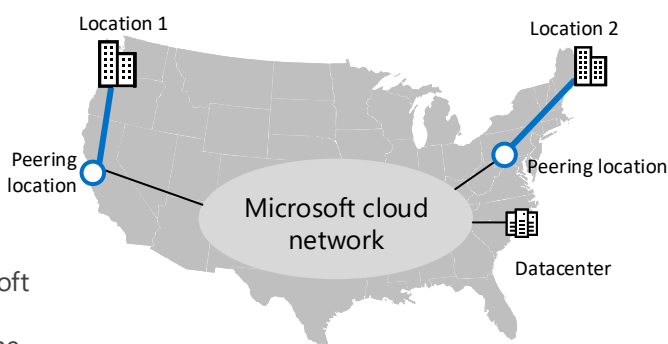
When you create an ExpressRoute connection to a Microsoft cloud peering location, you are connected to the Microsoft cloud network and all the Microsoft datacenter locations in the same continent. The traffic between the cloud peering location and the destination Microsoft datacenter is carried over the Microsoft cloud network.

This can result in non-optimal delivery to local Microsoft datacenters for the any-to-any connectivity model.

In this example, traffic from the east coast branch office has to go across the country to a west coast Microsoft cloud peering location and then back across to the East US Azure datacenter.



For optimal delivery, use multiple ExpressRoute connections to regional Microsoft cloud peering locations.



This can provide:

- Better performance to regionally local Microsoft datacenter locations.
- Higher availability to the Microsoft cloud when a local ExpressRoute connection becomes unavailable.

This works well for organizations in the same continent. However, traffic to Microsoft datacenters outside the organization's continent travels over the Internet.

For intercontinental traffic over the Microsoft cloud network, you must use ExpressRoute Premium connections.

With ExpressRoute Premium

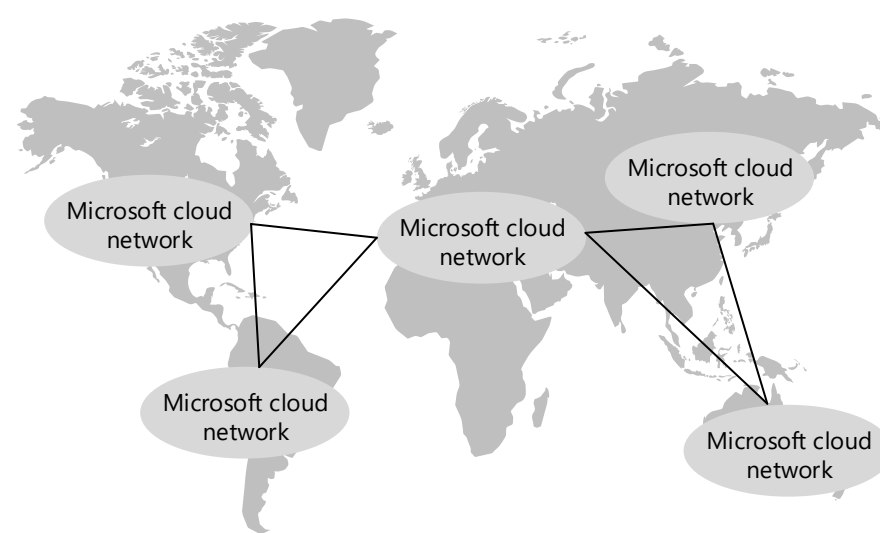
For organizations that are globally distributed across continents, you can use ExpressRoute Premium.

With ExpressRoute Premium, you can reach any Microsoft datacenter on any continent from any Microsoft peering location on any continent. The traffic between continents is carried over the Microsoft cloud network.

With multiple ExpressRoute Premium connections, you can have:

- Better performance to continentally local Microsoft datacenters.
- Higher availability to the global Microsoft cloud when a local ExpressRoute connection becomes unavailable.

Example of ExpressRoute Premium connections for a global enterprise using Office 365



With a portion of the Microsoft cloud network in each continent, a global enterprise creates ExpressRoute Premium connections from its regional hub offices to local Microsoft peering locations.

For a regional office, appropriate Office 365 traffic to:

- Continental Office 365 datacenters travels over the Microsoft cloud network within the continent.
- Office 365 datacenters in another continent travels over the intercontinental Microsoft cloud network.

[Network planning and performance tuning for Office 365](#)

ExpressRoute options

Security at your edge

To provide advanced security for the traffic sent and received over the ExpressRoute connection, such as traffic inspection or intrusion/malware detection, place your security appliances in the traffic path within your DMZ or at the border of your intranet.

Internet traffic for VMs

To prevent Azure VMs from initiating traffic directly with Internet locations, advertise the default route to Microsoft. Traffic to the Internet is routed across the ExpressRoute connection and through your on-premises proxy servers. Traffic from Azure VMs to Azure PaaS services or Office 365 is routed back across the ExpressRoute connection.

WAN optimizers

You can deploy WAN optimizers on both sides of a private peering connection for a cross-premises Azure virtual network (VNet). Inside the Azure VNet, use a WAN optimizer network appliance from the Azure marketplace and user-defined routing to route the traffic through the appliance.

Quality of service

Use Differentiated Services Code Point (DSCP) values in the IPv4 header of your traffic to mark it for voice, video/interactive, or best-effort delivery. This is especially important for the Microsoft peering relationship and Skype for Business Online traffic.

More information

ExpressRoute for Office 365

<http://aka.ms/expressrouteoffice365>

ExpressRoute for Azure

<https://azure.microsoft.com/services/expressroute/>

Microsoft Cloud Networking for Enterprise Architects

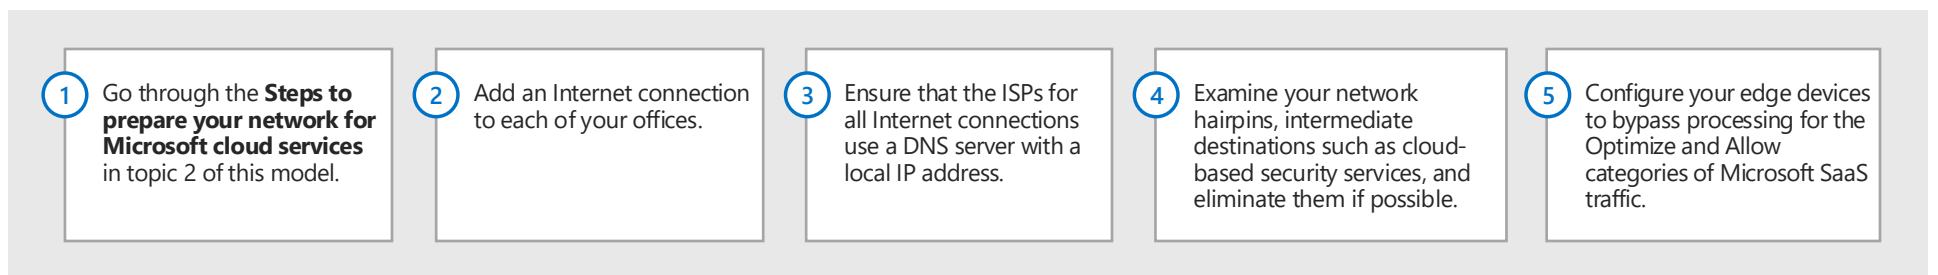
What IT architects need to know about networking in Microsoft cloud services and platforms

This topic is 4 of 6 in a series 1 2 3 4 5 6

Designing networking for Microsoft SaaS (Office 365, Microsoft Intune, and Dynamics 365)

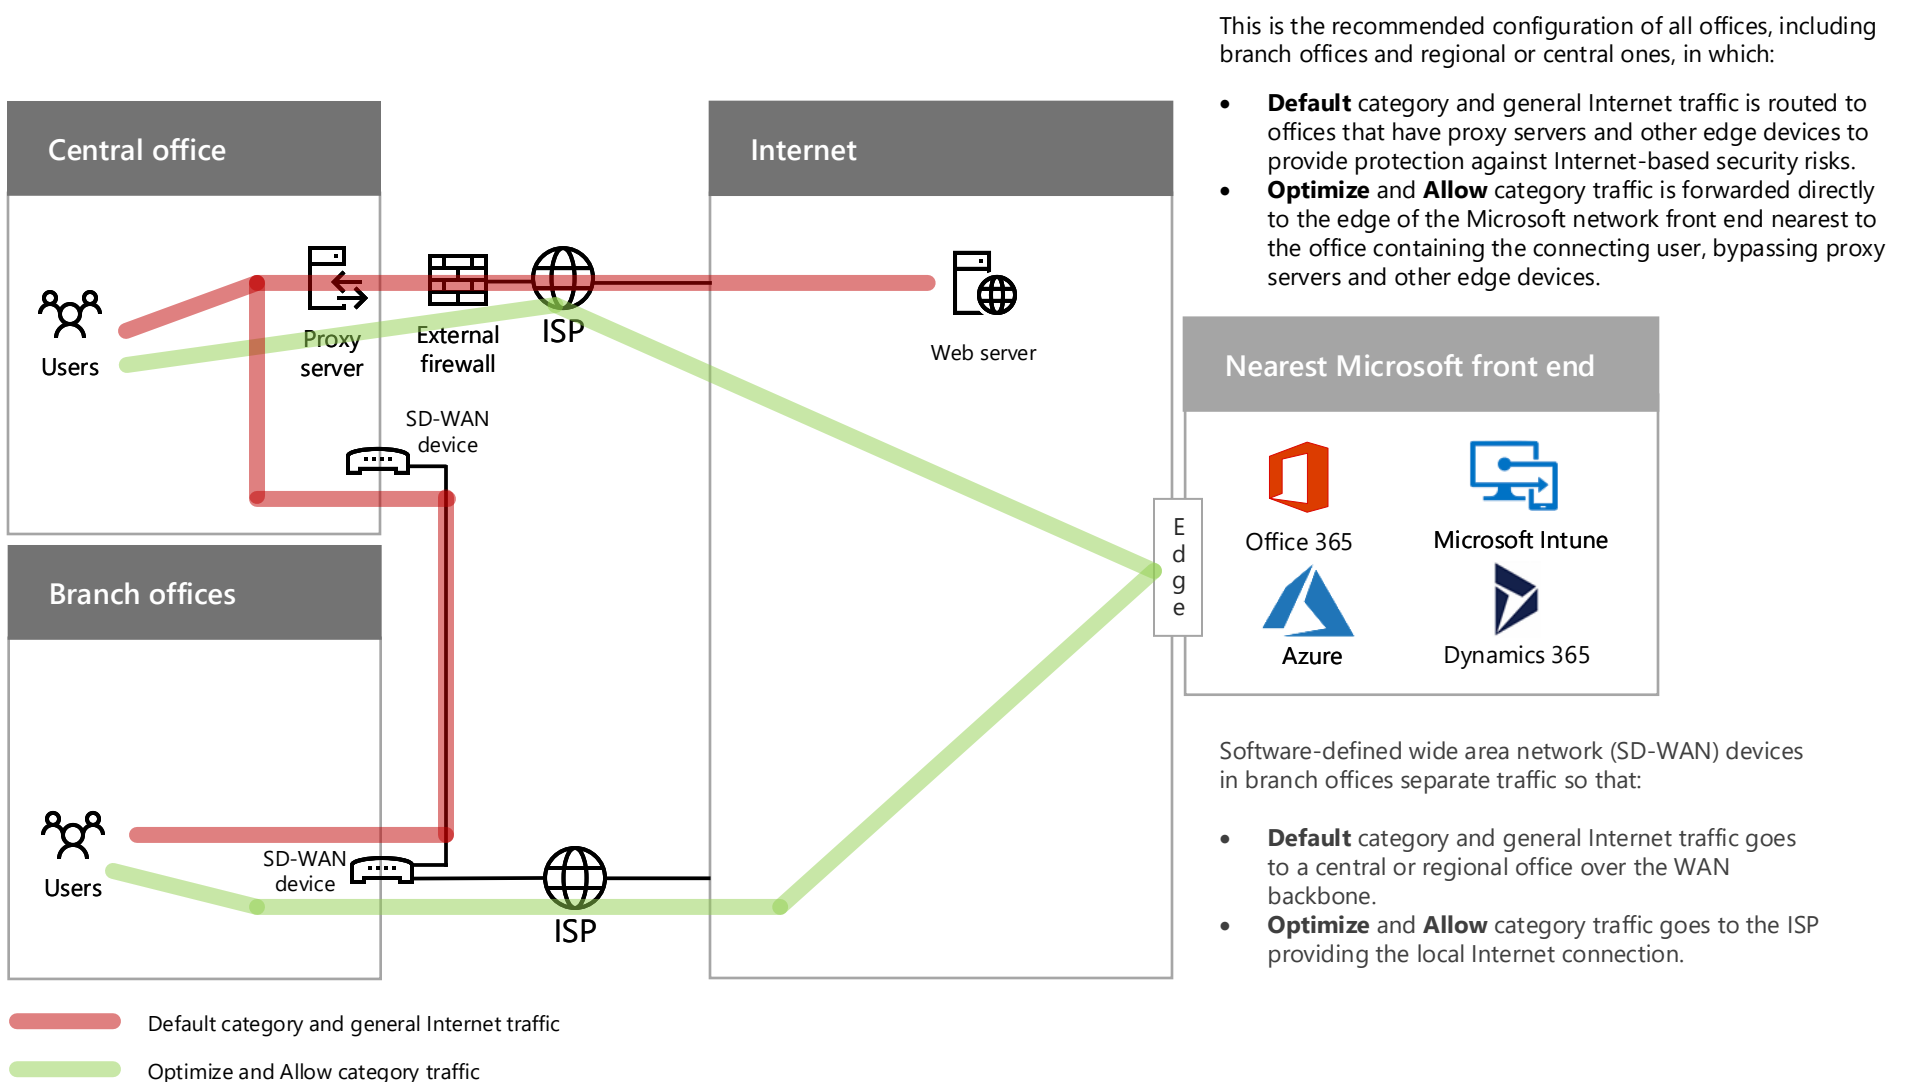
Optimizing your network for Microsoft SaaS services requires the configuration of internal and edge devices to route the different categories of traffic to Microsoft SaaS services.

Steps to prepare your network for Microsoft SaaS services



Optimizing traffic to Microsoft's SaaS services

Three categories of Microsoft SaaS traffic:	Optimize	Allow	Default
	Required for connectivity to every Microsoft SaaS service and represent over 75% of Microsoft SaaS bandwidth, connections, and volume of data.	Required for connectivity to specific Microsoft SaaS services and features but are not as sensitive to network performance and latency as those in the Optimize category.	Represent Microsoft SaaS services and dependencies that do not require any optimization. You can treat Default category traffic like normal Internet traffic.



More information

Network connectivity principles
<http://aka.ms/pnc>

Network planning and performance tuning for Office 365
<http://aka.ms/tune>

Microsoft Cloud Networking for Enterprise Architects

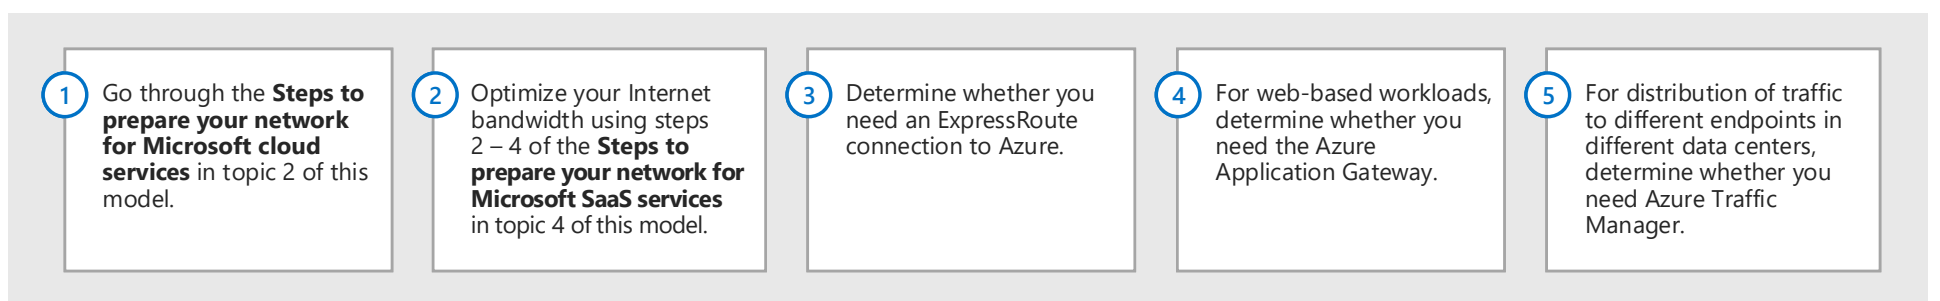
What IT architects need to know about networking in Microsoft cloud services and platforms

This topic is 5 of 6 in a series 1 2 3 4 5 6

Designing networking for Azure PaaS

Optimizing networking for Azure PaaS apps requires adequate Internet bandwidth and can require the distribution of network traffic across multiple sites or apps.

Planning steps for hosting organization PaaS applications in Azure

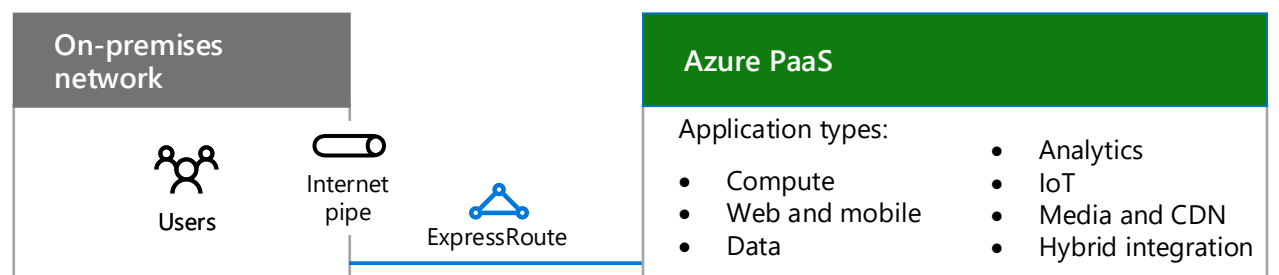


Internet bandwidth for organization PaaS applications

Organization applications hosted in Azure PaaS require Internet bandwidth for intranet users.

Option 1 Use your existing pipe, optimized for Internet traffic with the capacity to handle peak loads. See page 4 of this model for Internet edge, client usage, and IT operations considerations.

Option 2 For high-bandwidth or low latency needs, use an ExpressRoute connection to Azure.

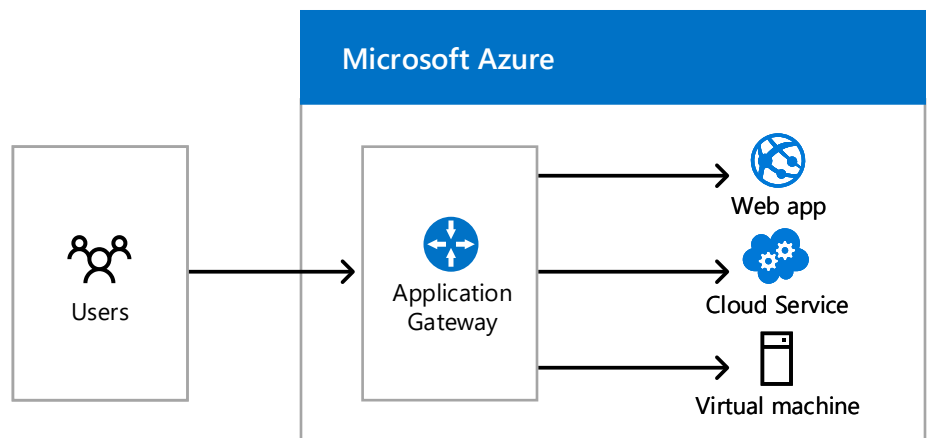


Azure Application Gateway

Application-level routing and load balancing services that let you build a scalable and highly-available web front end in Azure for web apps, cloud services, and virtual machines. Application Gateway currently supports layer 7 application delivery for the following:

- HTTP load balancing
- Cookie based session affinity
- SSL offload

[Application Gateway](#)

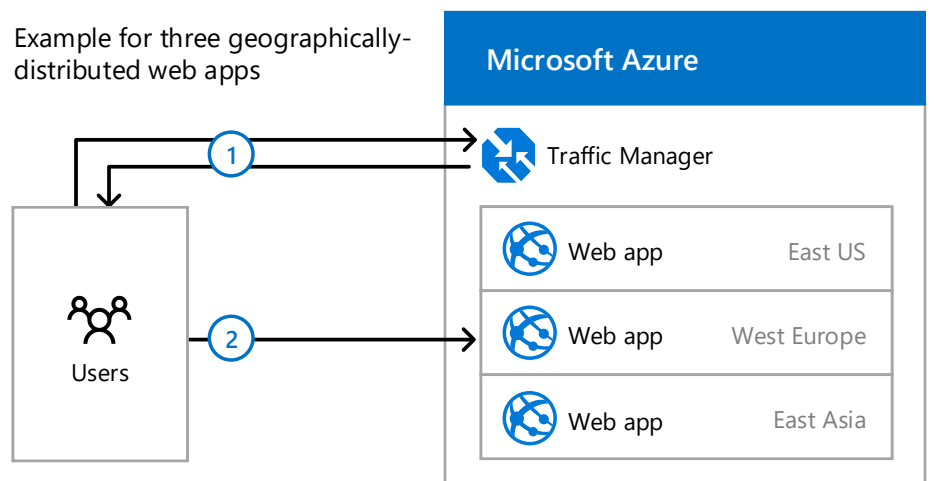


Azure Traffic Manager

Distribution of traffic to different endpoints, which can include cloud services or Azure web apps located in different data centers or external endpoints.

Traffic Manager routing methods
Failover The endpoints are in the same or different Azure datacenters and you want to use a primary endpoint for all traffic, but provide backups in case the primary or the backup endpoints are unavailable.
Round robin You want to distribute load across a set of endpoints in the same datacenter or across different datacenters.
Performance You have endpoints in different geographic locations and you want requesting clients to use the "closest" endpoint in terms of the lowest latency.

Example for three geographically-distributed web apps



1. A user DNS query for a web site URL gets directed to Azure Traffic Manager, which returns the name of a regional web app, based on the performance routing method.
2. User initiates traffic with the regional web app. [Traffic Manager](#)

Microsoft Cloud Networking for Enterprise Architects

What IT architects need to know about networking in Microsoft cloud services and platforms

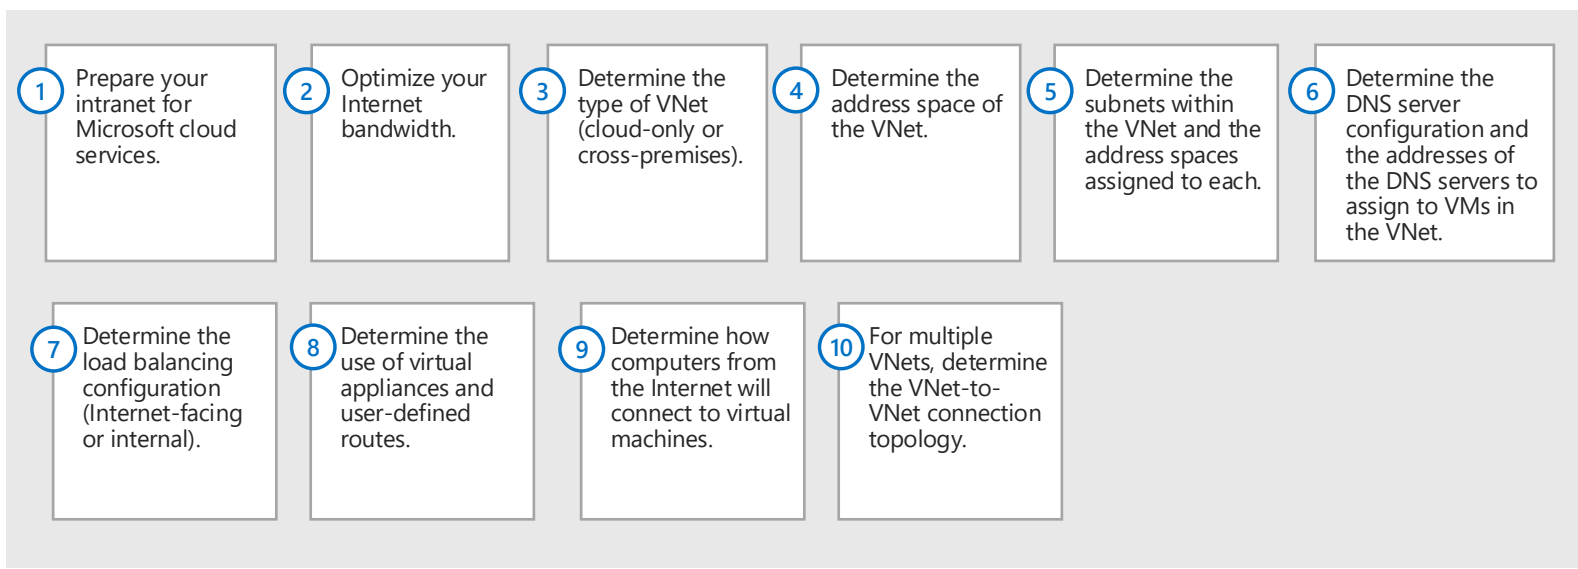
This topic is 6 of 6 in a series [1](#) [2](#) [3](#) [4](#) [5](#) [6](#)

Designing networking for Azure IaaS

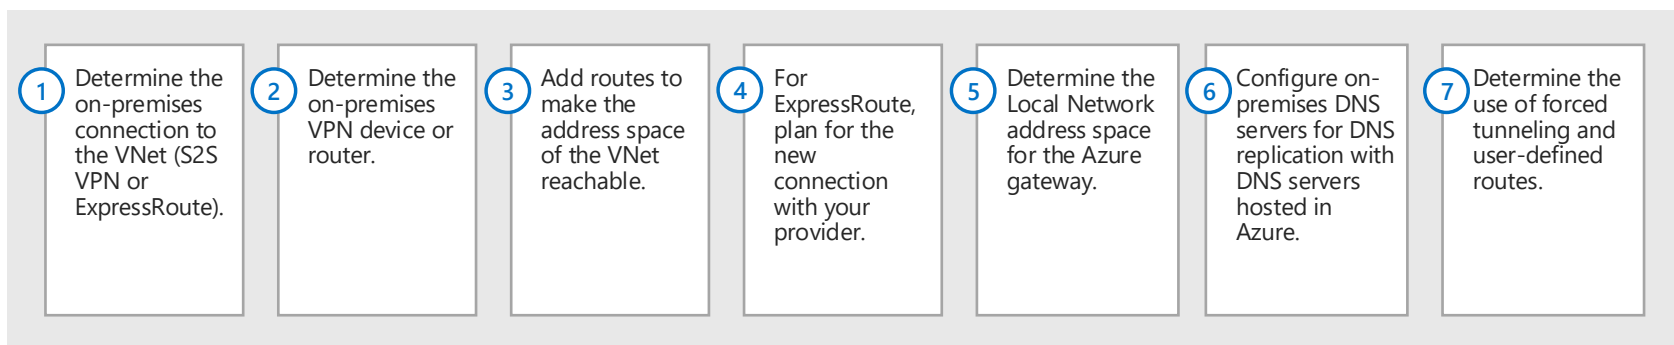
Optimizing networking for IT workloads hosted in Azure IaaS requires an understanding of Azure virtual networks (VNets), address spaces, routing, DNS, and load balancing.

Planning steps for hosting an IT workload in an Azure VNet

Planning for any VNet



Planning for cross-premises VNets



Planning steps for any Azure VNet

Step 1: Prepare your intranet for Microsoft cloud services.

Go through the **Steps to prepare your network for Microsoft cloud services** in topic 2 of this model.


Step 2: Optimize your Internet bandwidth.

Go through steps 2 – 4 of the **Steps to prepare your network for Microsoft SaaS services** in topic 4 of this model.

Step 3: Determine the type of VNet (cloud-only or cross-premises).

Cloud only

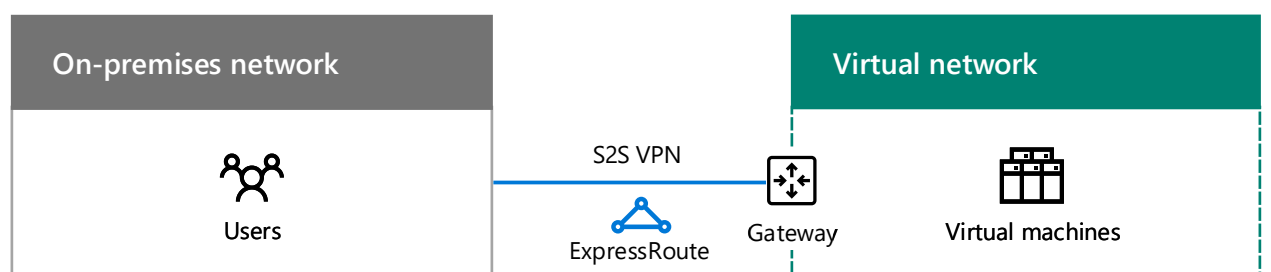
A VNet with no connection to an on-premises network.



[Continued on next page](#)

Cross-premises

A VNet with a Site-to-Site (S2S) VPN or ExpressRoute connection to an on-premises network through an Azure gateway.



See the additional **Planning steps for a cross-premises Azure VNet** in this topic.

Step 4: Determine the address space of the VNet.

Addressing for virtual networks

Type of VNet	Virtual network address space
Cloud only	Arbitrary private address space
Interconnected cloud-only	Arbitrary private, but not overlapping with other connected VNets
Cross-premises	Private, but not overlapping with on-premises
Interconnected cross-premises	Private, but not overlapping with on-premises and other connected VNets

Addressing for virtual machines

Virtual machines are assigned an address configuration from the address space of the subnet by DHCP:

- Address/subnet mask
- Default gateway
- DNS server IP addresses

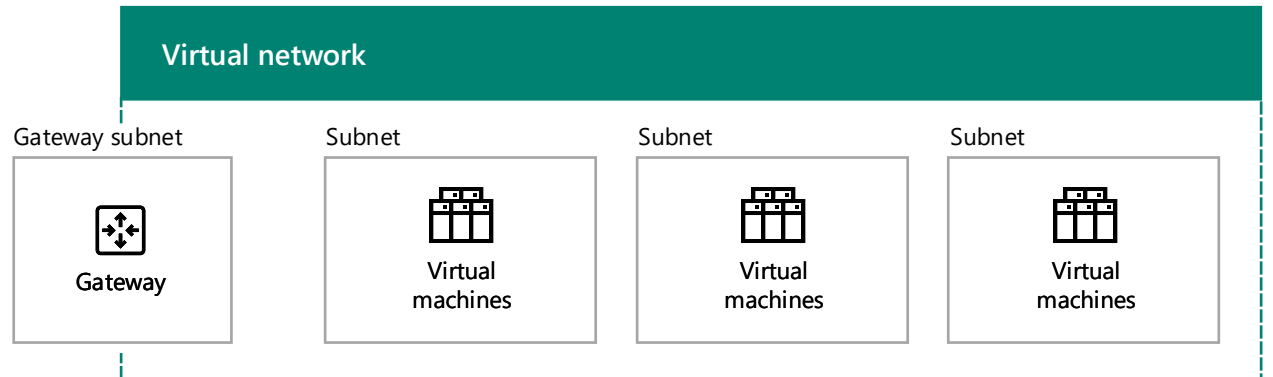
You can also reserve a static IP address.

Virtual machines can also be assigned a public IP address, either individually or from the containing cloud service (for classic deployment machines only).

Step 5: Determine the subnets within the VNet and the address spaces assigned to each.

Azure gateway subnet

Needed by Azure to host the two virtual machines of your Azure gateway. Specify an address space with at least a 29-bit prefix length (example: 192.168.15.248/29). A 27-bit prefix length is recommended, especially if you are planning to use ExpressRoute.



Best practice for determining the address space of the Azure gateway subnet:

1. Decide on the size of the gateway subnet.
2. In the variable bits in the address space of the VNet, set the bits used for the gateway subnet to 0 and set the remaining bits to 1.
3. Convert to decimal and express as an address space with the prefix length set to the size of the gateway subnet.

With this method, the address space for the gateway subnet is always at the farthest end of the VNet address space.

[Address space calculator for Azure gateway subnets](#)

Example of defining the address prefix for the gateway subnet

The address space of the VNet is 10.119.0.0/16. The organization will initially use a site-to-site VPN connection, but will eventually get ExpressRoute.

Step	Results
1. Decide on the size of the gateway subnet.	/28
2. Set the bits in the variable portion of the VNet address space: 0 for the gateway subnet bits (G), otherwise 1 (V).	10.119. bbbbbbbb . bbbbbbbb 10.119. vvvvvvvv . vvvvGGGG 10.119. 11111111 . 11110000
3. Convert result from step 2 to decimal and express as an address space.	10.119.255.240/28

Virtual machine-hosting subnets

Place Azure virtual machines in subnets according to typical on-premises guidelines, such as a common role or tier of an application or for subnet isolation.

Azure uses the first 3 addresses on each subnet. Therefore, the number of possible addresses on an Azure subnet is $2^n - 5$, where n is the number of host bits.

[Networking limits](#)

Virtual machines	Host bits	Subnet size
1-3	3	/29
4-11	4	/28
12-27	5	/27
28-59	6	/26
60-123	7	/25

Step 6: Determine the DNS server configuration and the addresses of the DNS servers to assign to VMs in the VNet.

Azure assigns virtual machines the addresses of DNS servers by DHCP. DNS servers can be:

- Supplied by Azure: Provides local name registration and local and Internet name resolution
- Provided by you: Provides local or intranet name registration and either intranet or Internet name resolution

[Name Resolution for resources in Azure virtual networks](#)

Type of VNet	DNS server
Cloud only	Azure-supplied for local and Internet name resolution
	Azure virtual machine for local and Internet name resolution (DNS forwarding)
Cross-premises	On-premises for local and intranet name resolution
	Azure virtual machine for local and intranet name resolution (DNS replication and forwarding)

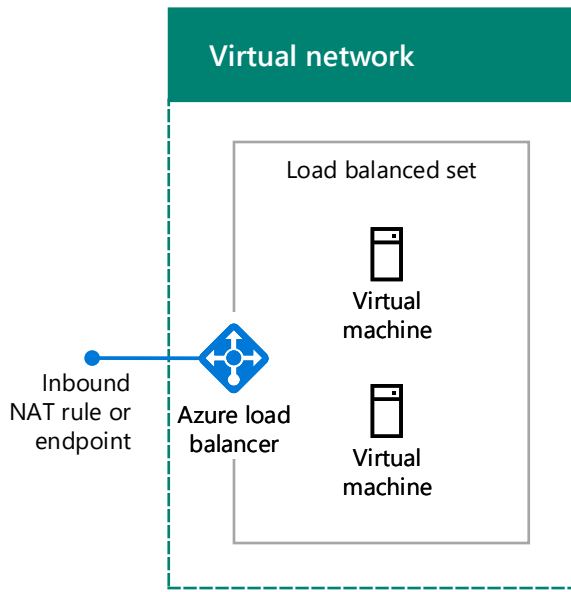
[Continued on next page](#)

Step 7: Determine the load balancing configuration (Internet-facing or internal).

Internet-facing load balancing

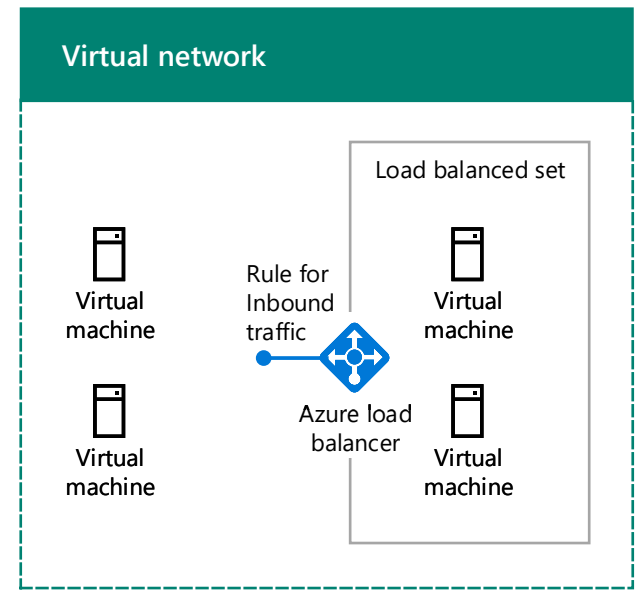
Randomly distribute unsolicited incoming traffic from the Internet to the members of a load-balanced set.

Azure Load Balancer



Internal load balancing

Randomly distribute unsolicited incoming traffic from other Azure VMs or from intranet computers (not shown) to the members of a load-balanced set.




Step 8: Determine the use of virtual appliances and user-defined routes.

User-defined routing

You may need to add one or more user-defined routes to a subnet to forward traffic to virtual appliances in your Azure virtual network.

Virtual network traffic routing



Step 9: Determine how computers from the Internet will connect to virtual machines.

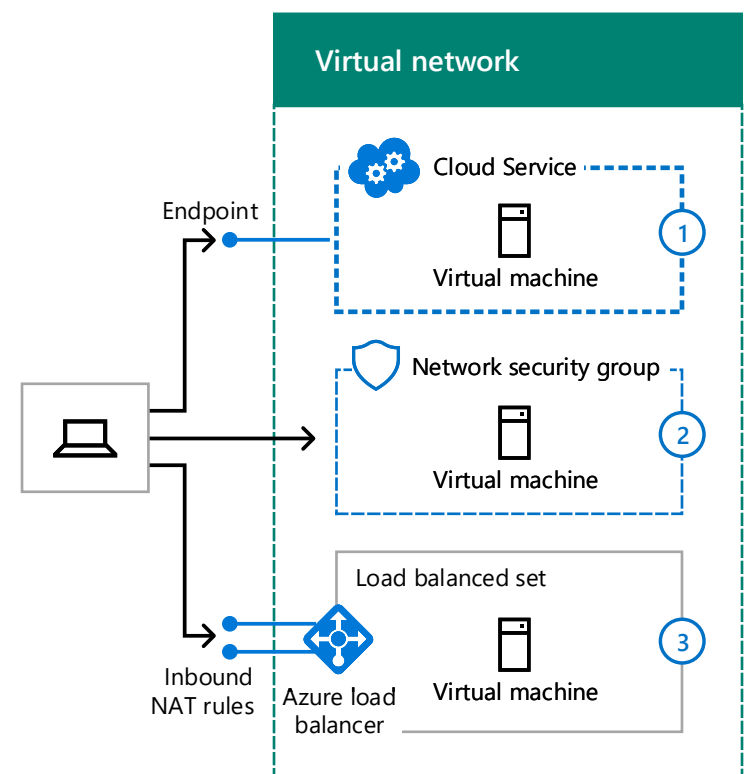
Includes access from your organization network through your proxy server or other edge device.

Methods for filtering or inspecting unsolicited incoming traffic

Method	Deployment model
1. Endpoints and ACLs configured on cloud services	Classic
2. Network security groups	Resource Manager and classic
3. Internet-facing load balancer with inbound NAT rules	Resource Manager
4. Network security appliances in the Azure Marketplace (not shown)	Resource Manager and classic

Additional security:

- Remote Desktop and SSH connections are authenticated and encrypted
- Remote PowerShell sessions are authenticated and encrypted
- You can use IPsec transport mode for end-to-end encryption
- Azure DDOS protection helps prevent external and internal attacks



Step 10: For multiple VNETs, determine the VNet-to-VNet connection topology.

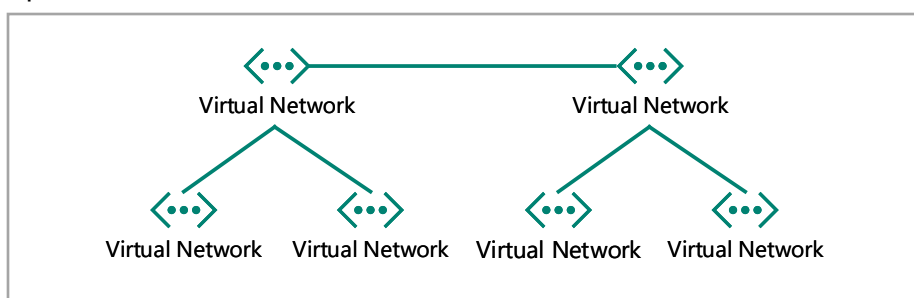
Azure VNETs can be connected to each other using topologies similar to those used for connecting the sites of an organization using VNet peering or VNet-to-VNet (V2V) connections.

Virtual network peering


Daisy chain



Spoke and hub



Full mesh



Continued on next page

Planning steps for a cross-premises Azure VNet

Step 1: Determine the cross-premises connection to the VNet (S2S VPN or ExpressRoute).

	Site-to-Site (S2S) VPN	Connect 1–10 sites (including other VNets) to a single Azure VNet.
	ExpressRoute	A private, secure link to Azure via an Internet Exchange Provider (IXP) or a Network Service Provider (NSP).

Other types of connections:

	Point-to-Site (P2S) VPN	Connects a single computer to an Azure VNet.
	VNet peering or VNet-to-VNet (V2V) VPN	Connects an Azure VNet to another Azure VNet.

Networking limits

VPN devices and IPsec/IKE parameters for site-to-site VPN gateway connections

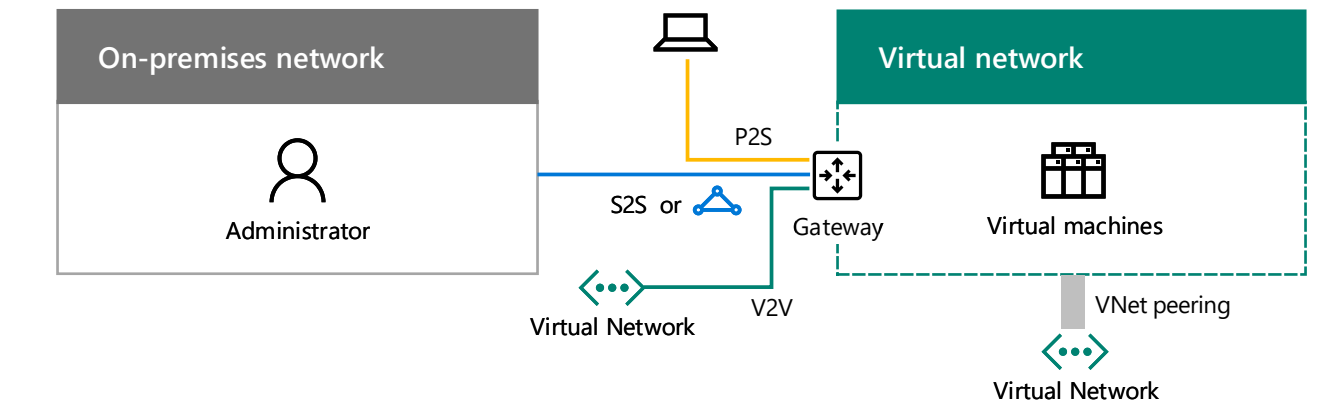
Virtual network peering

Connecting to VMs in the VNet:

- Administration of VNet VMs from your on-premises network or the Internet
- IT workload access from your on-premises network
- Extension of your network through additional Azure VNets

Security for connections:

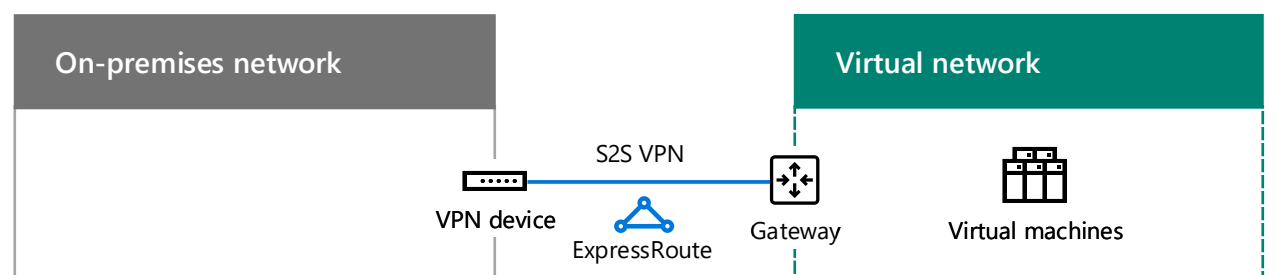
- P2S uses the Secure Socket Tunneling Protocol (SSTP)
- S2S and V2V VPN connections use IPsec tunnel mode with AES256
- ExpressRoute is a private WAN connection



Step 2: Determine the on-premises VPN device or router.

Your on-premises VPN device or router:

- Acts as an IPsec peer, terminating the S2S VPN connection from the Azure gateway.
- Acts as the BPG peer and termination point for the private peering ExpressRoute connection.

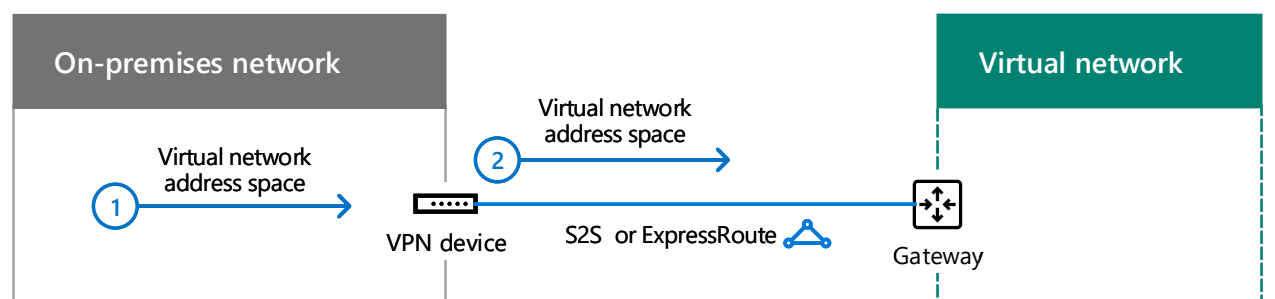


VPN gateways

Step 3: Add routes to your intranet to make the address space of the VNet reachable.

Routing to VNets from on-premises

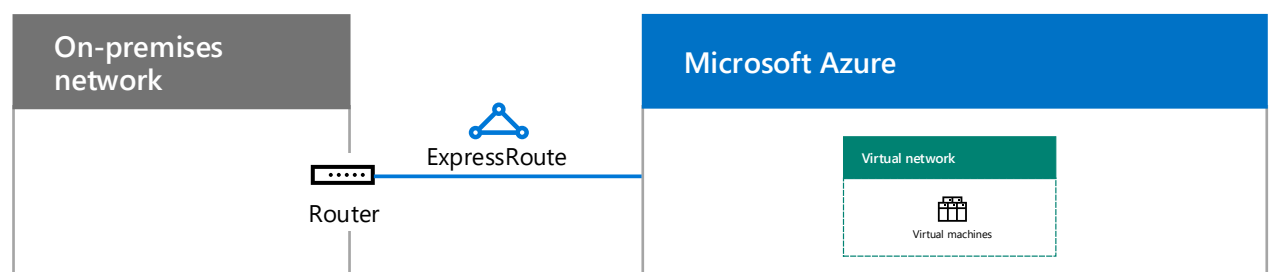
1. Route for the virtual network address space that points toward your VPN device
2. Route for the virtual network address space on your VPN device



Step 4: For ExpressRoute, plan for the new connection with your provider.

You can create an ExpressRoute connection with private peering between your on-premises network and the Microsoft cloud in three different ways:

- Co-located at a cloud exchange
- Point-to-point Ethernet connections
- Any-to-any (IP VPN) networks

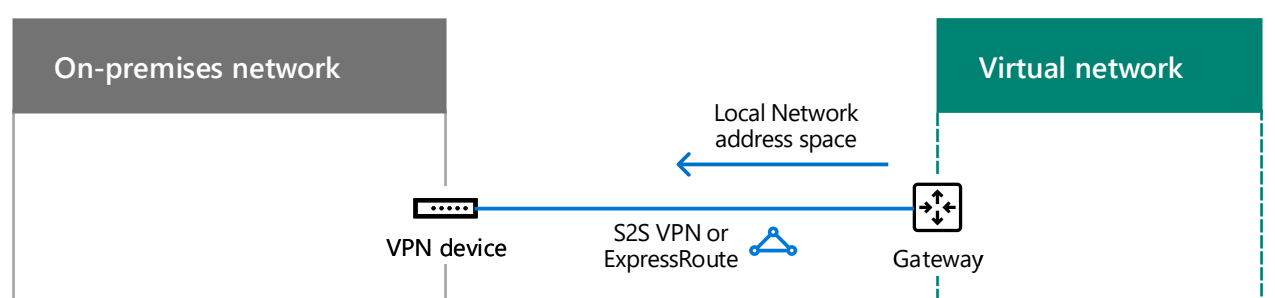


See topic 3, **ExpressRoute**.

Step 5: Determine the Local Network address space for the Azure gateway.

Routing to on-premises or other VNets from VNets

Azure forwards traffic across an Azure gateway that matches the Local Network address space assigned to the gateway.




Continued on next page

Defining the Local Network address space:

Option 1: The list of prefixes for the address space currently needed or in use (updates might be needed when you add new subnets).

Option 2: Your entire on-premises address space (updates only needed when you add new address space).

Because the Azure gateway does not allow summarized routes for S2S VPN connections, you must define the Local Network address space for option 2 so that it does not include the virtual network address space.



■ The virtual network address space
■ The root space

Example of defining the prefixes for the Local Network around the address space "hole" created by the virtual network for S2S VPN connections

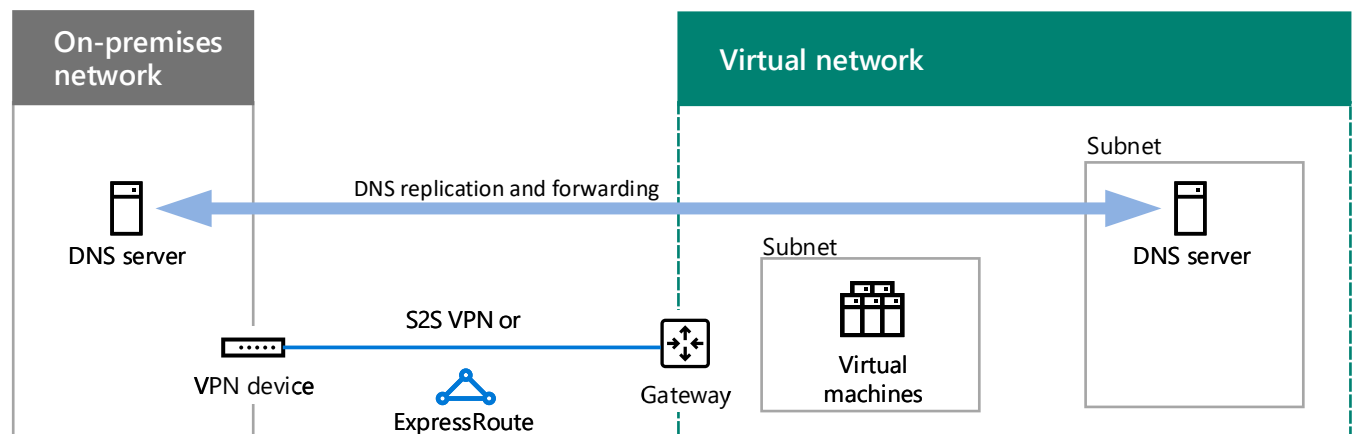
An organization uses portions of the private address space (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16) across their on-premises network. They chose option 2 and 10.100.100.0/24 as their virtual network address space.

Step	Prefixes
1. List the prefixes that are not the root space for the virtual network address space.	172.16.0.0/12 and 192.168.0.0/16
2. List the non-overlapping prefixes for variable octets up to but not including the last used octet in the virtual network address space.	10.0.0.0/16, 10.1.0.0/16... 10.99.0.0/16, 10.101.0.0/16... 10.254.0.0/16, 10.255.0.0/16 (255 prefixes, skipping 10.100.0.0/16)
3. List the non-overlapping prefixes within the last used octet of the virtual network address space.	10.100.0.0/24, 10.100.1.0/24... 10.100.99.0/24, 10.100.101.0/24... 10.100.254.0/24, 10.100.0.255.0/24 (255 prefixes, skipping 10.100.100.0/24)

Step 6: Configure on-premises DNS servers for replication with DNS servers hosted in Azure.

To ensure that on-premises computers can resolve the names of Azure-based servers and Azure-based servers can resolve the names of on-premises computers, configure:

- The DNS servers in your virtual network to forward to on-premises DNS servers.
- DNS replication of the appropriate zones between DNS servers on-premises and in the Azure VNet.




Step 7: Determine the use of forced tunneling.

The default system route for Azure subnets points to the Internet. To ensure that all traffic from virtual machines travels across the cross-premises connection, create a routing table with the default route that uses the Azure gateway as its next-hop address. You then associate the route table with the subnet.

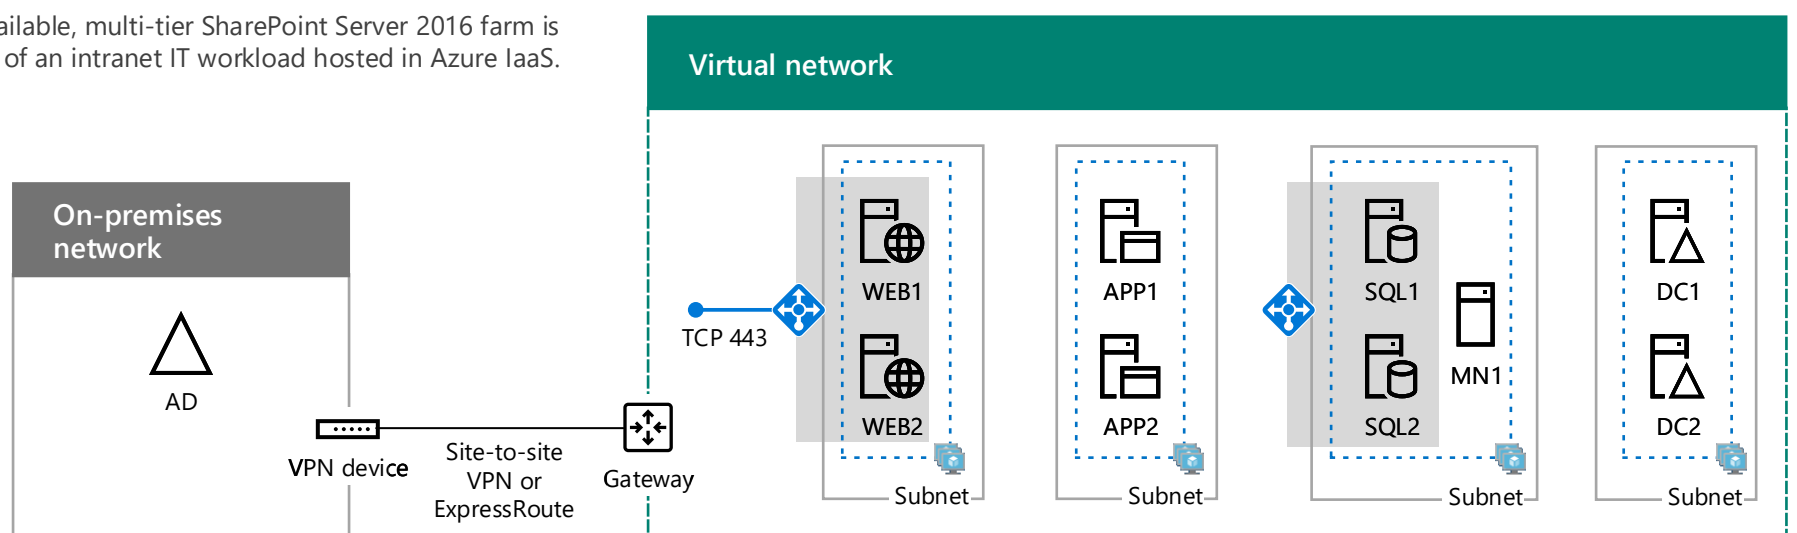
This is known as forced tunneling.

[Configure forced tunneling](#)



SharePoint Server 2016 farm in Azure

A highly-available, multi-tier SharePoint Server 2016 farm is an example of an intranet IT workload hosted in Azure IaaS.



[SharePoint Server 2016 in Microsoft Azure](#)

[Intranet SharePoint Server 2016 in Azure dev/test environment](#)

More Microsoft cloud IT resources

Security
aka.ms/cloudarchsecurity

Hybrid
aka.ms/cloudarchhybrid