

Setting Up an F5[®] NVGRE Gateway Environment

BIG-IP Version 11.5, BIG-IQ Version 4.4



Table of Contents

Legal Notices.....	5
Chapter 1: Setting Up an F5 Networks NVGRE Gateway Environment.....	7
Overview: Setting up an F5 Networks NVGRE gateway environment.....	8
About network virtualization using generic routing (NVGRE).....	8
About customer addresses.....	8
About provider addresses.....	10
About virtual subnets.....	10
About routing domains.....	11
About logical networks.....	12
About IP address pools.....	13
About logical switches with port profiles.....	14
About virtual port profiles.....	15
About VM networks.....	15
Before you begin the installation.....	16
Task summary.....	18
Creating the BIG-IQ device resolver group.....	18
Installing the F5 Networks HNV Gateway PowerShell Module.....	19
Configuring the VM gateway BIG-IP system to forward packets.....	21
Configuring the F5 gateway in SCVMM.....	21
Viewing F5 Networks HNV Gateway PowerShell Module logs.....	29
Example of F5 Networks NVGRE gateway environment.....	29

Legal Notices

Publication Date

This document was published on August 26, 2014.

Publication Number

PUB-0301-00

Copyright

Copyright © 2013-2014, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

AAM, Access Policy Manager, Advanced Client Authentication, Advanced Firewall Manager, Advanced Routing, AFM, APM, Application Acceleration Manager, Application Security Manager, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Certified [DESIGN], F5 Networks, F5 SalesXchange [DESIGN], F5 Synthesis, f5 Synthesis, F5 Synthesis [DESIGN], F5 TechXchange [DESIGN], Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, iApps, IBR, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, Local Traffic Manager, LTM, LineRate, LineRate Systems [DESIGN], LROS, LTM, Message Security Manager, MSM, OneConnect, Packet Velocity, PEM, Policy Enforcement Manager, Protocol Security Manager, PSM, Real Traffic Policy Builder, SalesXchange, ScaleN, Signalling Delivery Controller, SDC, SSL Acceleration, software designed applications services, SDAC (except in Japan), StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TechXchange, TMOS, TotALL, Traffic Management Operating System, Traffix Systems, Traffix Systems (DESIGN), Transparent Data Reduction, UNITY, VAULT, vCMP, VE F5 [DESIGN], Versafe, Versafe [DESIGN], VIPRION, Virtual Clustered Multiprocessing, WebSafe, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at:

<http://www.f5.com/about/guidelines-policies/patents>

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Chapter

1

Setting Up an F5 Networks NVGRE Gateway Environment

- *Overview: Setting up an F5 Networks NVGRE gateway environment*
- *About network virtualization using generic routing (NVGRE)*
- *Before you begin the installation*
- *Task summary*
- *Example of F5 Networks NVGRE gateway environment*

Overview: Setting up an F5 Networks NVGRE gateway environment

This document provides instructions for installing the F5 Networks HNV Gateway PowerShell Module in the System Center Virtual Machine Manager (SCVMM) for integration into a Microsoft Hyper-V environment. The plug-in allows you to use a BIG-IP® device or VE to act as a gateway between virtual networks in SCVMM and external networks. That is, with this plug-in, virtual machines can connect to the outside world, using NVGRE tunnels. After you have made this connection, you can configure the BIG-IP systems to provide application services for those virtual machines. The BIG-IQ® system is the management endpoint for the BIG-IP systems. By default, all communication from the F5 Networks HNV Gateway PowerShell Module occurs through the BIG-IQ system.

About network virtualization using generic routing (NVGRE)

Using generic routing encapsulation (GRE) for policy-based, software-controlled network virtualization supports multitenancy in public and private clouds. NVGRE encapsulates Ethernet frames in an NVGRE-formatted GRE packet. You can combine virtual network segments managed by NVGRE with segments managed by VXLAN in either or both multicast and unicast modes.

NVGRE serves most data centers deploying network virtualization. The system encapsulates packets inside another packet, and the header of the new packet has the appropriate source and destination provider address (PA) IP address in addition to the virtual subnet ID (VSID), which is stored in the Key field of the GRE header. The VSID allows hosts to identify the customer's virtual machines for any given packet.

NVGRE is a policy-driven solution, so the provider addresses (PAs) and customer addresses (CAs) on the packets can overlap without problems. Consequently, all virtual machines on the same host can share a single PA.

These concepts are important for deploying NVGRE with Microsoft System Center Virtual Machine Manager (SCVMM):

- Customer address (CA)
- Provider address (PA)
- Virtual subnets
- Routing domains
- Logical networks
- IP pools for each logical network site
- Logical switches with port profiles
- Virtual port profiles
- VM networks

For additional information about network virtualization concepts, you can consult Microsoft documentation, for example: http://blogs.msdn.com/b/microsoft_press/archive/2014/03/24/free-ebook-microsoft-system-center-network-virtualization-and-cloud-computing.aspx

About customer addresses

In NVGRE deployments with System Center Virtual Machine Manager (SCVMM), the *customer address (CA)* is the IP address assigned by the customer or tenant, based on the subnet, IP address range, and network

topology. This IP address is visible only to the virtual machine and, eventually, other virtual machines within the same subnet VM network, if you allow routing.

In this example, ERICVM1 is a virtual machine currently running on Hyper-V host MTCPARIS-2. Its IPv4 address 192.168.0.2 is visible only to this virtual machine, and not to the underlying network fabric.

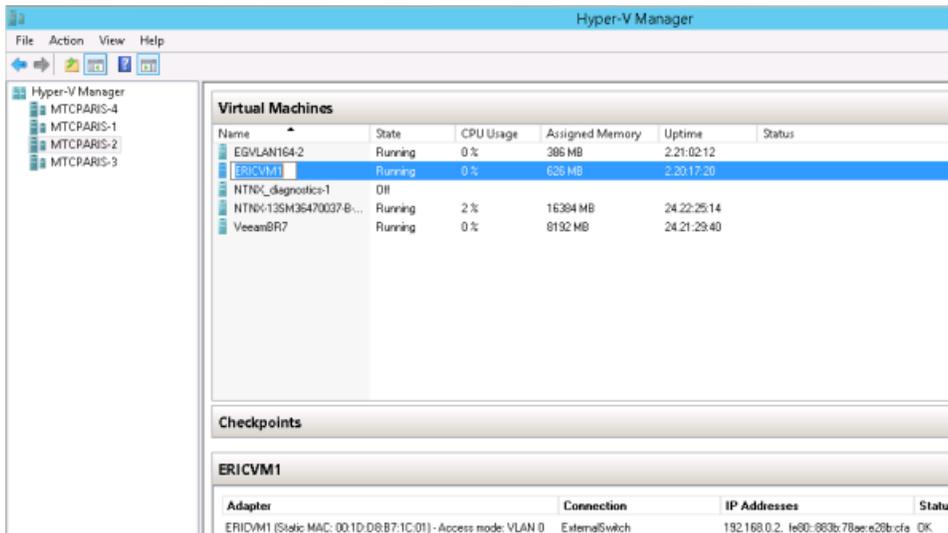


Figure 1: Screen snippet showing visibility of customer address to VM

You can double-check this concept by connecting directly to the virtual machine, as in this example.

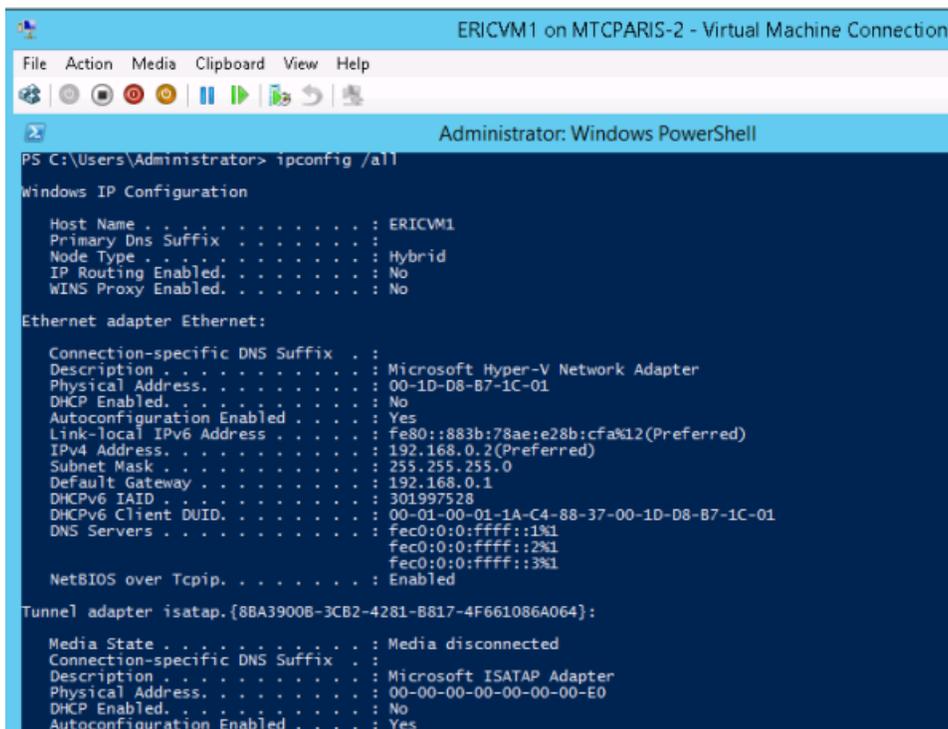


Figure 2: Command line verification of customer address visibility to VM

About provider addresses

In NVGRE deployments with System Center Virtual Machine Manager (SCVMM), the *provider address (PA)* is the IP address assigned by the administrator or by SCVMM, based on the physical network infrastructure. This IP address, visible only on the physical network, is used when Hyper-V hosts (either standalone or clustered) and other devices are exchanging packets, when participating in network virtualization.

This example shows the virtual machines running on the Hyper-V host MTCPARIS-2, which includes the ERICVM1 virtual machine. The PA associated with the ERICVM1 virtual machine is 10.10.0.5, which is never visible to the ERICVM1 virtual machine itself.

```
PS C:\Windows\system32> hostname
mtcparis-2
PS C:\Windows\system32> Get-NetVirtualizationLookupRecord

CustomerAddress : 192.168.0.2
VirtualSubnetID : 7829576
MACAddress      : 001dd8b71c01
ProviderAddress : 10.10.0.5
CustomerID      : <E90CB765-2866-47E4-BD2E-3DD36909CAD4>
Context         : SCUMM-MANAGED
Rule            : TranslationMethodEncap
UMName         : ERICVM1
UseUmMACAddress : False
Type           : Static

CustomerAddress : 192.168.0.3
VirtualSubnetID : 7829576
MACAddress      : 001dd8b71c02
ProviderAddress : 10.10.0.6
CustomerID      : <E90CB765-2866-47E4-BD2E-3DD36909CAD4>
Context         : SCUMM-MANAGED
Rule            : TranslationMethodEncap
UMName         : ERICVM2
UseUmMACAddress : False
Type           : Static

CustomerAddress : 192.168.0.1
VirtualSubnetID : 7829576
MACAddress      : 004fdc5fe168
ProviderAddress : 1.1.1.1
CustomerID      : <E90CB765-2866-47E4-BD2E-3DD36909CAD4>
Context         : SCUMM-MANAGED
Rule            : TranslationMethodEncap
UMName         : GW
UseUmMACAddress : False
Type           : Static

CustomerAddress : 10.254.254.2
VirtualSubnetID : 11713130
MACAddress      : 00155d27ec07
ProviderAddress : 10.10.0.254
CustomerID      : <E90CB765-2866-47E4-BD2E-3DD36909CAD4>
Context         : SCUMM-MANAGED
Rule            : TranslationMethodEncap
UMName         : GW-External
UseUmMACAddress : False
Type           : Static
```

Figure 3: Example of provider address for VM participating in network virtualization

About virtual subnets

In NVGRE deployments with System Center Virtual Machine Manager (SCVMM), a unique virtual subnet ID (VSID) identifies an IP subnet at Layer 3 and a broadcast domain boundary at Layer 2, similar to VLAN technology. The VSID must be unique within the data center and within the range of 4096 to $2^{24}-2$. Two customers in a hosted data center cannot both use the same VSID, even if they have different routing domains.

The VSID is a setting of the port of the virtual switch (vSwitch). However, it is presented to you as a property of the virtual network interface (VNI) of a VM.

In this example, the VSID for the ERICVM1 virtual machine is 7829576.

```
PS C:\Windows\system32> Get-UMNetworkAdapter -UMName ERICVM1 | fl
Name           : ERICVM1
Id             : Microsoft:D14226D1-EBE7-4026-8B63-B49AEB8A998D\3B5D3E6B-C104-41DD-9504-C5D1CFC81640
IsLegacy       : False
IsManagementOs : False
ComputerName   : MICPARIS-2
UMName        : ERICVM1
UMId          : d14226d1-ebe7-4026-8b63-b49aeb8a998d
SwitchName     : ExternalSwitch
SwitchId      : 927495c7-dad8-4e29-93f4-4a3dd425eb8a
Connected      : True
PoolName       :
MacAddress     : 001DD8B71C01
DynamicMacAddressEnabled : False
MacAddressSpoofing : Off
AllowTeaming  : Off
RouterGuard    : Off
DhcpGuard      : Off
StormLimit     : 0
PortMirroringMode : None
IeeePriorityTag : Off
VirtualSubnetId : 7829576
DynamicIPAddressLimit : 0
UMQWeight     : 100
UMQUsage      : 0
IOUWeight     : 0
IOUUsage      : 0
IoQueuePairsRequested : 1
IoQueuePairsAssigned : 0
IOUInterruptModeration : Default
IPsecOffloadMaxSA : 512
IPsecOffloadSAUsage : 0
UPDataPathActive : False
MaximumBandwidth : 0bps
MinimumBandwidthAbsolute : 0bps
MinimumBandwidthWeight : 0(weight)
BandwidthPercentage : 0%
MandatoryFeatureId : <>
MandatoryFeatureName : <>
Status        : <Ok>
IPAddresses   : <192.168.0.2, fe80::883b:78ae:e28b:cfa>
```

Figure 4: Example including VSID for a virtual machine

About routing domains

In NVGRE deployments with System Center Virtual Machine Manager (SCVMM), a routing domain defines a relationship between the virtual subnets created by the tenants, and identifies the VM network.

- The routing domain ID (RDID) has a globally unique ID (GUID) within the data center.
- The network virtualization stack enables Layer 3 routing between the subnets with a default gateway (always x.x.x.1), which cannot be disabled or configured.
- Hyper-V network virtualization (HNV) addresses distribute Layer-3 routing between virtualized subnets by including a network virtualization routing extension natively inside Hyper-V virtual switches running on each Hyper-V host.
- This distributed router can make cross-subnet routing decisions locally within the vSwitch to directly forward traffic between VMs on different virtualized subnets within the same virtual network or routing domain.
- To manage and distribute the appropriate routing policies to each Hyper-V host, System Center 2012 R2 VMM performs as the routing policy server, enabling the configuration of distributed routers across many Hyper-V hosts to be easily coordinated from a single, centralized point of administration.

This example shows two different routing domains on the same Hyper-V host.

```

PS C:\Windows\system32> hostname
mtcparis-2
PS C:\Windows\system32> Get-NetVirtualizationCustomerRoute

RoutingDomainID : <E90CB765-2866-47E4-BD2E-3DD36909CAD4>
VirtualSubnetID : 7829576
DestinationPrefix : 192.168.0.0/24
NextHop           : 0.0.0.0
Metric            : 0

RoutingDomainID : <E90CB765-2866-47E4-BD2E-3DD36909CAD4>
VirtualSubnetID : 11713130
DestinationPrefix : 0.0.0.0/0
NextHop           : 10.254.254.2
Metric            : 0
    
```

Figure 5: Example of two routing domains on a single Hyper-V host

About logical networks

In NVGRE deployments with System Center Virtual Machine Manager (SCVMM), a logical network can contain one or more associated network sites. A *network site* is a user-defined named grouping of IP subnets, VLANs, or IP subnet and VLAN pairs, which is used to organize and simplify network assignments. Logical networks are useful in large environments for mapping and streamlining network connectivity and dependencies in the configuration.

Uses for logical networks include but are not limited to these:

- Management: Contains the IP subnet used for management. Typically, both VMM and the Hyper-V servers are connected to this physical network. If you have more than one site and/or several VLANs, you can add all of these to the same logical network.
- Cluster: Contains the IP subnet and VLAN for cluster communication. Live Migration
- Front end: Contains the IP subnet used for public IP addresses.
- PA network: Contains the IP subnet used for provider addresses.

The logical network is dedicated to network virtualization. It is enabled at the logical network level. This network must be isolated. Do not use any of the other networks for this purpose.

The logical network in this example has an associated IP pool, so that SCVMM can manage IP address assignments to the hosts dedicated to network virtualization, the virtualization gateway VMs, and the virtualization hosts running virtual machines connected to VM networks.

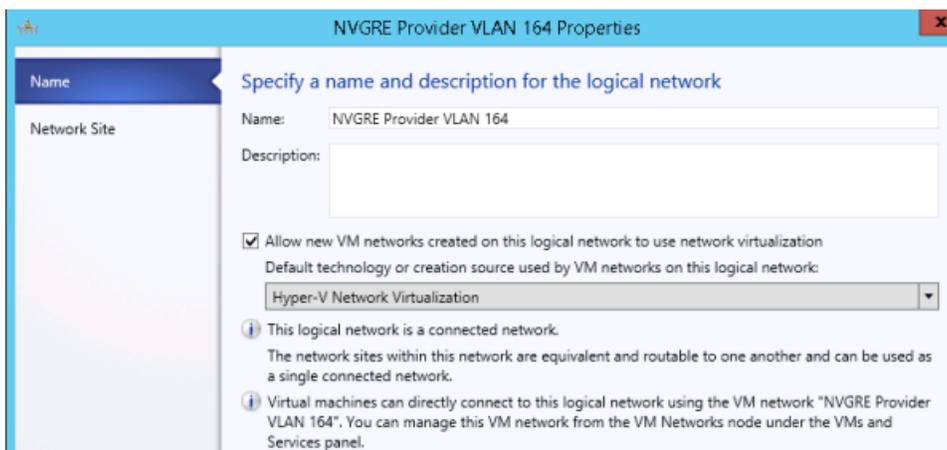


Figure 6: Specifying a logical network

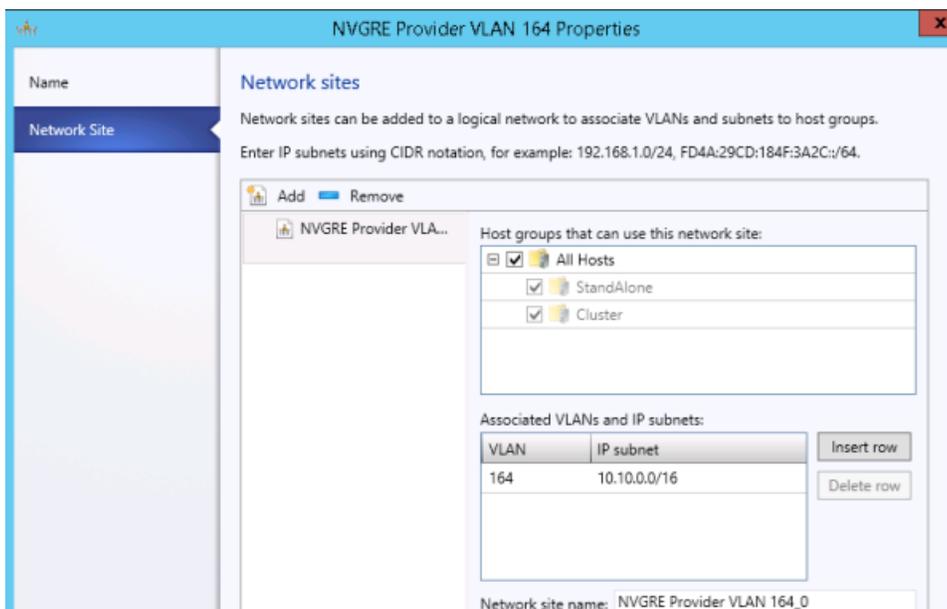


Figure 7: Adding network sites to a logical network

About IP address pools

In NVGRE deployments with System Center Virtual Machine Manager (SCVMM), You must have IP address pools for each logical network site, so that VMM can assign the right IP configuration to its resources within this network.

In these configuration screen examples, note that there is no direct mapping of the PA network to the hosts. The PA network is available to the hosts only through this configuration, together with Uplink port profiles and logical switches.

Important: Do not configure network virtualization on any other logical networks that you present to the same hosts.

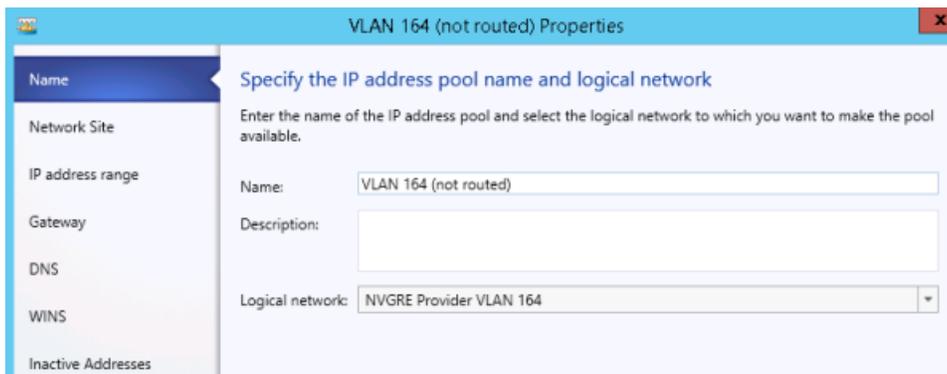


Figure 8: Specifying the IP address pool name and logical network

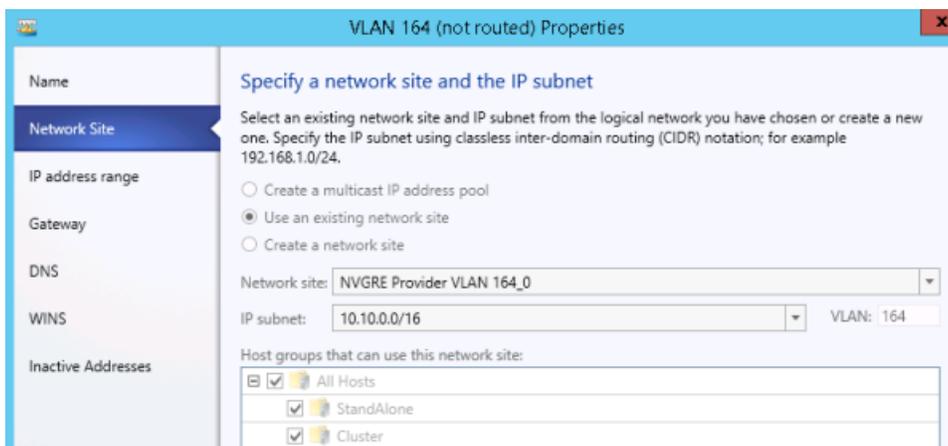


Figure 9: Specifying a network site and the IP subnet

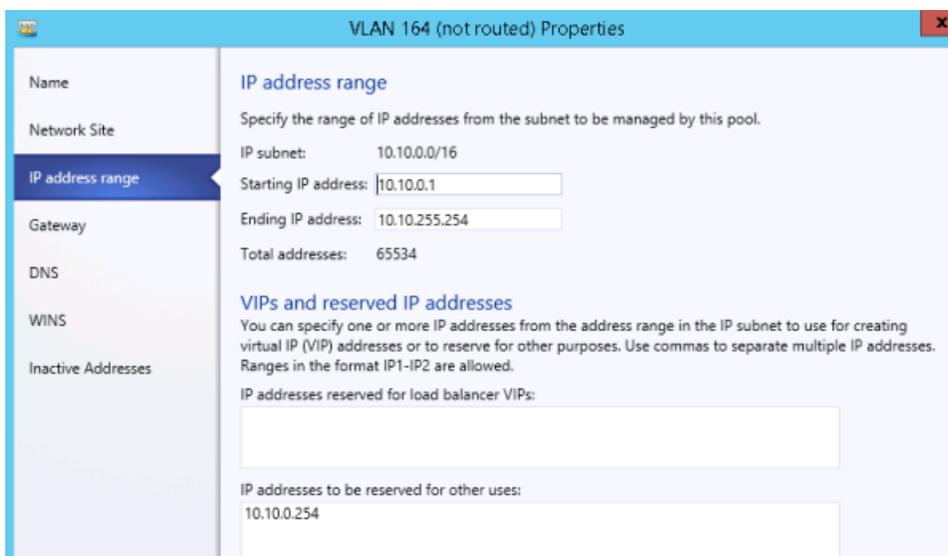


Figure 10: Specifying the range of IP addresses for a pool

About logical switches with port profiles

In NVGRE deployments with System Center Virtual Machine Manager (SCVMM), you can use port profiles and logical switches to create identical capabilities for network adapters across multiple hosts. Port profiles and logical switches act as containers for the properties or capabilities that you want your network adapters to have. Instead of configuring individual properties or capabilities for each network adapter, you can specify the capabilities in port profiles and logical switches, and then apply these capabilities to the appropriate adapters. This can simplify the configuration process and ensure that your hosts are using correct load balancing algorithm and the virtual adapters have the right settings related to capabilities and QoS.

About virtual port profiles

In NVGRE deployments with System Center Virtual Machine Manager (SCVMM), you can take advantage of several port profiles that are shipped with SCVMM and use the existing profiles for host management, cluster, and live migration.

You can see the profiles in SCVMM by navigating to **Port Profiles** on the networking tab in **Fabric**.

For example, on the Security Settings screen, you can enable **Allow guest specified IP addresses**, so that VMM can detect changes made to tenants within the guests, and update the NVGRE policy in the environment.

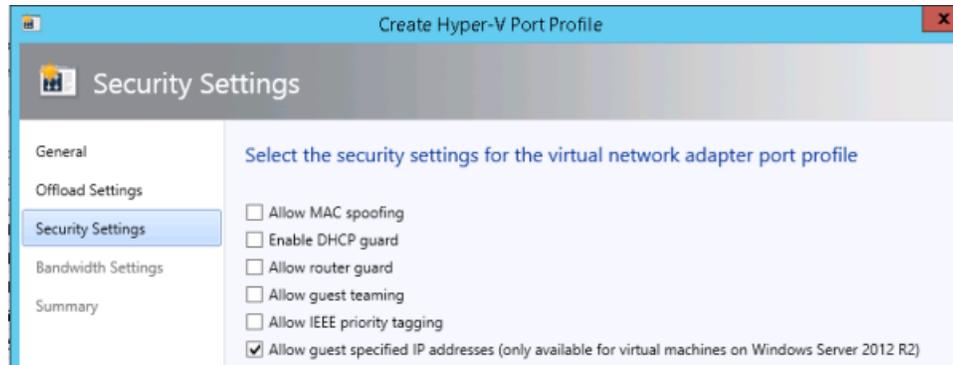


Figure 11: Setting security for a virtual port profile

About VM networks

Note: This step is not necessary if, when you created the logical networks, you selected **Create a VM network with the same name to allow virtual machines to access this logical network directly**.

You need to create VM networks with 1:1 mapping to your logical networks in the fabric.

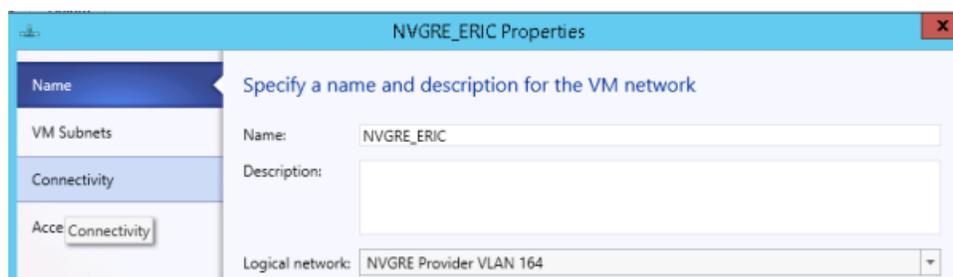


Figure 12: Creating a VM network

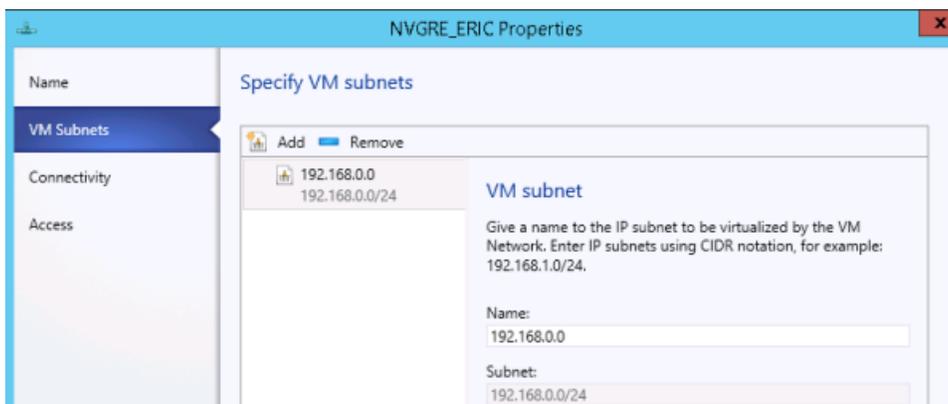


Figure 13: Specifying VM subnets

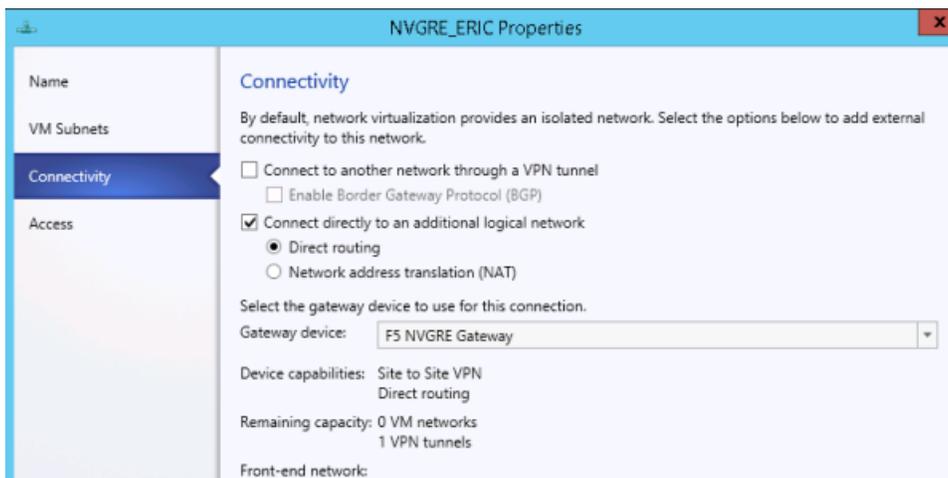


Figure 14: Adding external connectivity to a VM network

Before you begin the installation

Before you install the F5 Networks HNV Gateway PowerShell Module, you need to prepare the BIG-IP® system. F5 Networks strongly recommends using Engineering Hotfix 121.14 for v11.5.1-HF2. It includes bug fixes that ensure that monitors on VIPs work correctly in an HA setup that also uses tunnels.

Make sure that the BIG-IP system is configured with these considerations in mind.

- At least one BIG-IP system is configured with at least one IP address, preferably on the management interface.
- Verify that at least one VLAN has connectivity to the provider network.

When you configure the SCVMM gateway using a config file, you specify the following BIG-IQ parameters:

- The IP address of the F5 BIG-IQ system
- The name of a BIG-IQ device resolver group that contains either one standalone BIG-IP system or two BIG-IP systems in a device cluster

For a standalone BIG-IP system

If you are setting up a standalone BIG-IP system, verify that you have not configured a masquerading MAC address.

For a pair of BIG-IP systems in a device group

If you are setting up a pair of BIG-IP devices as a device group, verify the following.

- You are using a Sync-Failover device group.
- Auto-Sync and Network Failover are turned on for the device group.
- You have configured a masquerading MAC address.
- Make a note of the traffic group used for floating objects; you need to provide it in the configuration file.

When you are not using route domains

When you do not use route domains; that is, when `UseRouteDomains` is set to `false` in the F5 Networks HNV Gateway PowerShell Module configuration file, you must create a forwarding virtual server on each of your BIG-IP systems.

Here is an example using the `tmsh` command line utility.

```
create ltm virtual scvmm-vs destination 0.0.0.0:0 mask any ip-forward
source-address-translation { type automap }
```

Additional information

The following files are shipped with the plug-in.

- Plug-in binaries:

```
C:\Windows\System32\WindowsPowerShell\v1.0\Modules\F5GatewayProvider
```

- Sample configuration file for one BIG-IP system:

```
C:\Windows\System32\WindowsPowerShell\v1.0\Modules\F5GatewayProvider\gateway-one-bigip.cfg
```

- Sample configuration file for two BIG-IP systems in a redundant (HA) pair:

```
C:\Windows\System32\WindowsPowerShell\v1.0\Modules\F5GatewayProvider\gateway-bigip-ha-pair.cfg
```

- Script to create a BIG-IP device resolver group:

```
C:\Windows\System32\WindowsPowerShell\v1.0\Modules\F5GatewayProvider\Setup-Device-Group.ps1
```

- Log file:

```
C:\Program Files\Microsoft System Center 2012 R2\Virtual Machine
Manager\bin\F5-SCVMM-Gateway.log
```

Task summary

Before you start this installation, you need to acquire the file `F5GatewayPowerShellSetup.msi`, which you can find on the BIG-IQ® system. You must use an account with administrative privileges to complete the installation. Ensure that the BIG-IP® system includes a local IP address in the provider IP address space, and make a note of this address. Also, ensure that the network interface (NIC) to be used for the provider addresses is named WNVNIC.

Task list

Creating the BIG-IQ device resolver group

Installing the F5 Networks HNV Gateway PowerShell Module

Configuring the VM gateway BIG-IP system to forward packets

Configuring the F5 gateway in SCVMM

Viewing F5 Networks HNV Gateway PowerShell Module logs

Creating the BIG-IQ device resolver group

When you configure The F5 Networks HNV Gateway PowerShell Module, you need to create a BIG-IQ device resolver group for each gateway you configure.

1. Locate the setup script file

`C:\Windows\System32\WindowsPowerShell\v1.0\Modules\F5GatewayProvider\Setup-Device-Group.ps1`

2. Run the file.

The setup script takes the parameters shown in the table. If you do not specify the credentials as command-line arguments, the system prompts you for them.

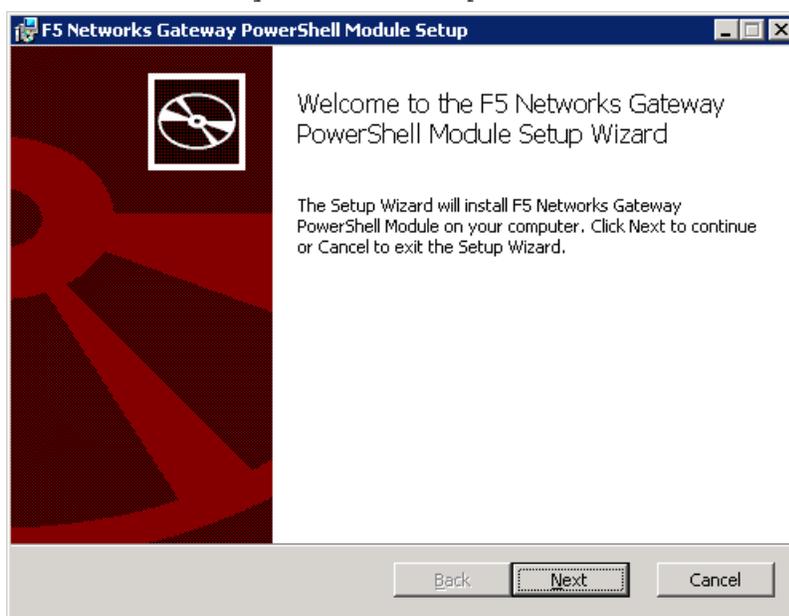
Parameter	Required?	Description
-BigIpAddress	Yes	IP address of the BIG-IQ system. This should match the setting <code>BigIpAddress</code> in the configuration file.
-BigIqCred	No	Credential for the admin account on the BIG-IQ system. (This is a <code>PSCredential</code> , as are the other credentials.)
-Force	No	If the group exists, force it to be recreated.
-ActiveBigIp	Yes	The IP address of the active BIG-IP system
-StandbyBigIp	No	The IP address of the standby BIG-IP system, if used
-GroupName	Yes	The name of the group to create. This should match the setting <code>BigIqDeviceGroup</code> in the configuration file.

Parameter	Required?	Description
-ActiveBigIpAdminCred	No	Credential for the admin account of the active BIG-IP system.
-ActiveBigIpRootCred	No	Credential for the root account of the active BIG-IP system.
-StandbyBigIpAdminCred	No	Credential for the admin account of the standby BIG-IP system.
-StandbyBigIpRootCred	No	Credential for the root account of the standby BIG-IP system.

Installing the F5 Networks HNV Gateway PowerShell Module

The F5 Networks HNV Gateway PowerShell Module provides a setup wizard to install the BIG-IP® system as an NVGRE gateway for System Center Virtual Machine Manager (SCVMM). After the installation is complete, you must restart the SCVMM services, or reboot the SCVMM server.

1. Run the file `F5GatewayPowerShellSetup.msi`.



2. Click **Next**.
3. Accept the EULA, and for the **Installation Type**, select **Complete**.
4. Configure the F5 Networks Gateway.
 - a) Open the sample configuration file that corresponds to your setup, either a standalone BIG-IP system (`gateway-one-bigip.cfg`) or a BIG-IP HA pair (`gateway-bigip-ha-pair.cfg`), at this location.

```
{SYSDIR}\WindowsPowerShell\v1.0\Modules\F5GatewayProvider
```

For default installations, `{SYSDIR}` is `C:\Windows\System 32`.
 - b) Edit the file, as indicated.

```
<GatewaySettings>
  <BigIqAddress>0.0.0.0</BigIqAddress>
  <BigIqDeviceGroup>scvmm</BigIqDeviceGroup>
```

```

<ActiveProviderAddress></ActiveProviderAddress> [BIG-IP HA pair only]
<StandbyProviderAddress></StandbyProviderAddress> [BIG-IP HA pair only]
<ProviderFwEnforcedPolicy></ProviderFwEnforcedPolicy>
<ProviderFwStagedPolicy></ProviderFwStagedPolicy>
<FloatingTrafficGroup>traffic-group-1</FloatingTrafficGroup> [BIG-IP HA pair only]
<UseRouteDomains>>true</UseRouteDomains>
<RouteDomainRange first="1" last="500"/>
<UseInboundTunnelMode>>false</UseInboundTunnelMode>
<CreateForwardingVirtual>>true</CreateForwardingVirtual>
<ForwardingVirtualSNAT>automap</ForwardingVirtualSNAT>
<CustomerSelfIpAllow>all</CustomerSelfIpAllow>
<CustomerSelfFwStagedPolicy></CustomerSelfFwStagedPolicy>
<CustomerSelfFwEnforcedPolicy></CustomerSelfFwEnforcedPolicy>
<TunnelMtu>0</TunnelMtu>
<TunnelProfile>nvgre</TunnelProfile>
<DumpGatewayState>>true</DumpGatewayState>
</GatewaySettings>

```

Considerations for these settings:

- The setting `BigIqAddress` is the management IP address of the BIG-IQ system.
- For a standalone system, `BigIqDeviceGroup` is the name of a BIG-IQ device group containing a single BIG-IP system. This is the name of the BIG-IQ device resolver group that you created previously.
- For a BIG-IP HA pair, `BigIqDeviceGroup` is the name of a BIG-IQ device group that contains two BIG-IP systems: one active, one standby. This is the name of the BIG-IQ device resolver group that you created previously.
- If you create more than one gateway instance on a pair of BIG-IP systems, you must ensure that each of them uses the same non-floating provider address. To do this reliably, you pre-create the provider IP address and specify it in the `ActiveProviderAddress` and `StandbyProviderAddress` settings. Note that you can attach a single gateway to multiple virtual networks without needing this setting. It is applicable only when using a device cluster, not a standalone BIG-IP system.
- If you have not provisioned Advanced Firewall Manager™ AFM™, or you do not want to use AFM on provider self IP addresses, leave empty the values for `ProviderFwEnforcedPolicy` and `ProviderFwStagedPolicy`, which specify the enforced firewall and staged firewall policies for the provider. If you specify a policy name, it is your responsibility to ensure that the policy is created before it is needed.
- For a BIG-IP HA pair, `FloatingTrafficGroup` is the name of the traffic group to use for objects that are synced between BIG-IP devices in the device cluster.
- If you might ever have two different VM networks that contain the same IP address, set `UseRouteDomains` to `true`. If you are not sure, retain the setting `true`. If you are sure that you will never reuse IP addresses in different VM networks, you could set this value to `false` to avoid a small scalability hit in configuring the BIG-IP system.
- The setting `RouteDomainRange` is relevant only if `UseRouteDomains` is set to `true`. It specifies the range of route domains that can be used by this gateway on the BIG-IP systems. If you create multiple gateways on the same set of BIG-IP systems (common in an active-active gateway setup), this setting ensures that each gateway does not use route domains belonging to other gateways that share the same BIG-IP device. Make sure that for each gateway that shares a set of BIG-IP devices, you specify a unique range of route domains. Gateways on different BIG-IP devices can use the same range of route domains. Note that this range is inclusive; all the route domains from `first` to `last` are used.
- If you are using a standard BIG-IP system running software v11.5.1 or earlier, set the `UseInboundTunnelWorkaround` to `false`. If you are using the engineering hotfix that makes the NVGRE tunnel pair look like a single tunnel, set `UseInboundTunnelWorkaround` to `true`. If you are not sure, ask F5 Networks support, or set it to `false`. Although both values should work, there are advantages to setting it to `true`, when available.

- If `CreateForwardingVirtual` is set to `true`, when a gateway is added to a customer virtual subnet, a forwarding virtual server is added. Note that this happens only if `UseRouteDomains` is also set to `true`.
 - The `ForwardingVirtualSNAT` setting specifies how you want to handle SNAT on the forwarding virtual, if you requested one. The possible values are `none` and `automap`. If you are not sure, retain the setting `automap`.
 - The `CustomerSelfIpAllow` setting specifies which traffic is allowed when the plug-in creates a self IP address on a customer's virtual subnet. The possible values are `all`, `none`, and `default`. Custom values are not allowed.
 - If you have not provisioned Advanced Firewall Manager™ AFM™, or you do not want to use AFM on customer self IP addresses, leave empty the values for `CustomerSelfFwStagedPolicy` and `CustomerSelfFwEnforcedPolicy`, which specify the staged firewall and enforced firewall policies for the customer. If you specify a policy name, it is your responsibility to ensure that the policy is created before it is needed.
 - If you need to set the MTU on the NVGRE tunnels, specify a value for the `TunnelMtu` setting. If you retain the value `0`, the system automatically sets the MTU.
 - Typically, you do not change the `TunnelProfile` setting. However, if you need to change it to an alternate NVGRE tunnel profile, you can set it using this parameter. The tunnel must be an NVGRE tunnel that already exists on the BIG-IP system.
 - If the setting `DumpGatewayState` is `true`, the system creates a file in the same directory as the log file, named `F5-Gateway-CONFIGFILENAME-state.txt`. The file contains information about the mapping of SCVMM routing domains to route domains, and the set of VSIDs managed within each routing domain. Although it takes some time to generate this file, it is useful if you are building a BIG-IP configuration on top of the configuration generated by the plug-in. It might be less useful if `UseRouteDomains` is `false`.
 - The BIG-IP partition setting is currently ignored.
- c) Save or copy the file to the same location.
When you create the VM gateway, you need to specify the name of this file as the network service connection string.
5. After the installation has completed, restart the SCVMM.



Configuring the VM gateway BIG-IP system to forward packets

After you add the BIG-IP® system as a VM gateway, you need to configure the system to forward packets.

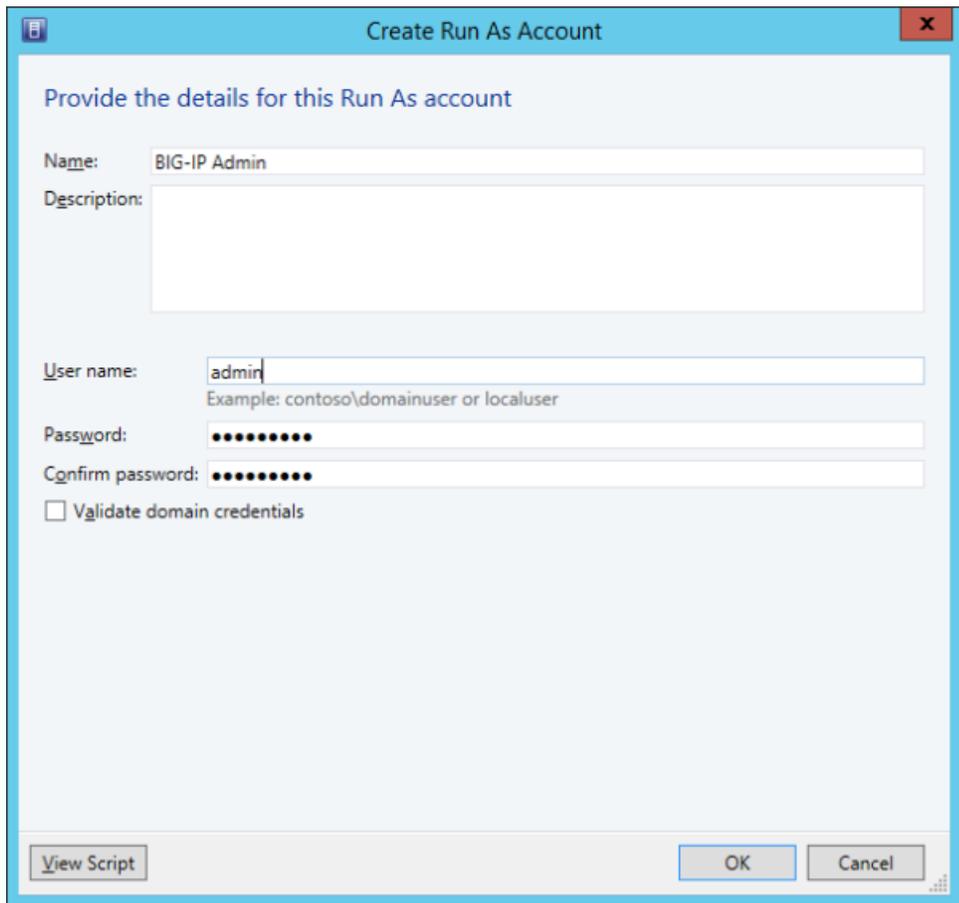
1. Create an external VLAN on the BIG-IP system that has external access.
2. Create a default route on the BIG-IP system that directs traffic outward.

Configuring the F5 gateway in SCVMM

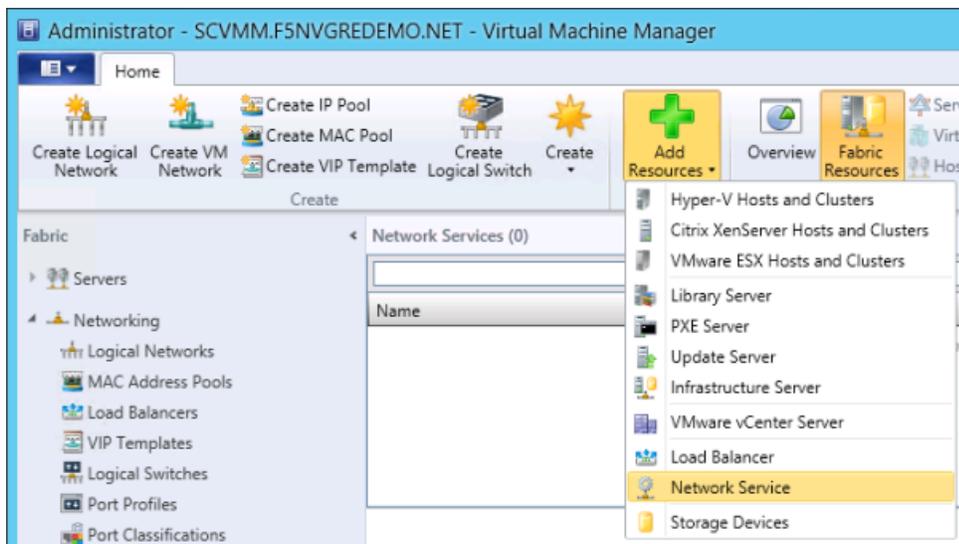
Before starting this task, you must install and load the BIG-IP® F5 Networks HNV Gateway PowerShell Module.

After you install the gateway, you can configure the VM network to use the gateway.

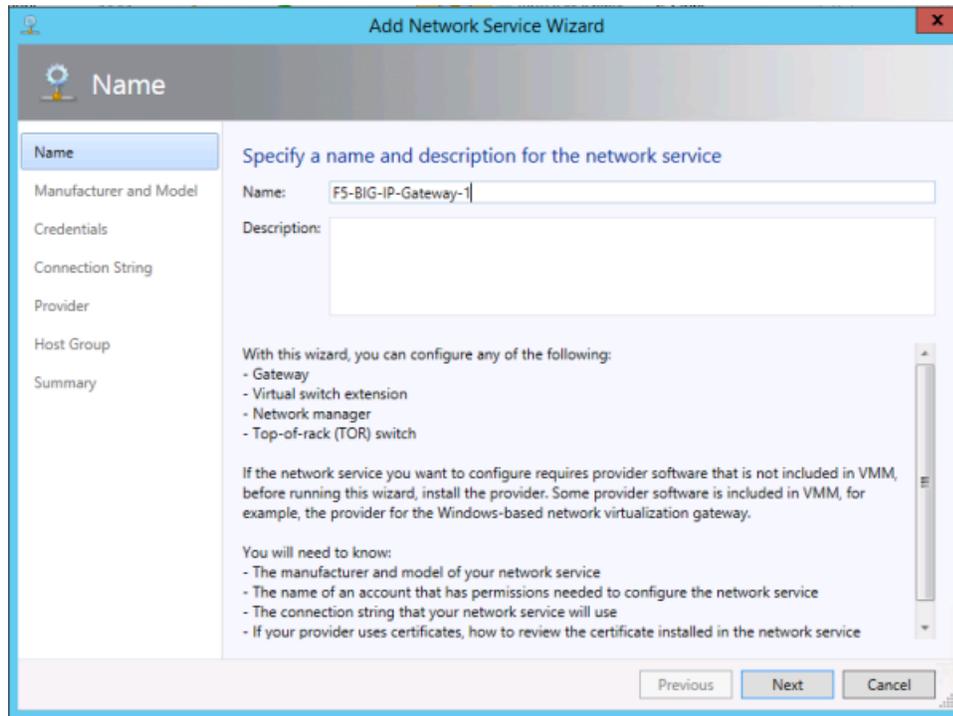
1. In the VMM's Settings area, create a new Run As account. This account includes the user name and password of the BIG-IP system you are using. It does not use domain credentials. You will select this Run As account later in the configuration process. If you already have an appropriate Run As account, you do not need to create another. Multiple gateways can refer to the same Run As account.



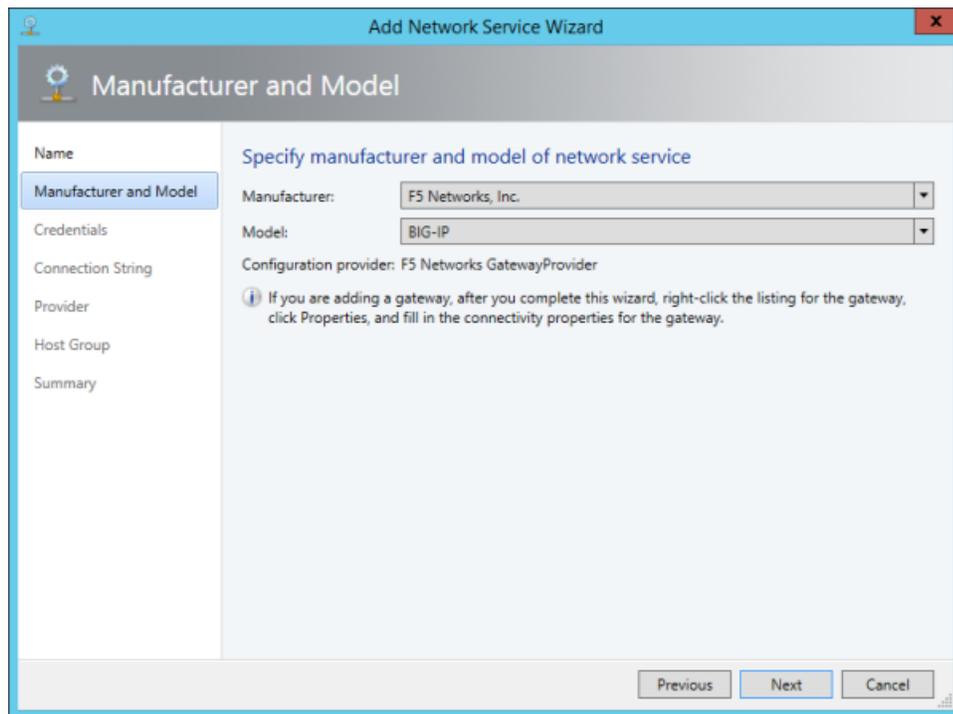
2. In the Fabric portion of the user interface, click **Add Resources**, and select **Network Service**, as shown.



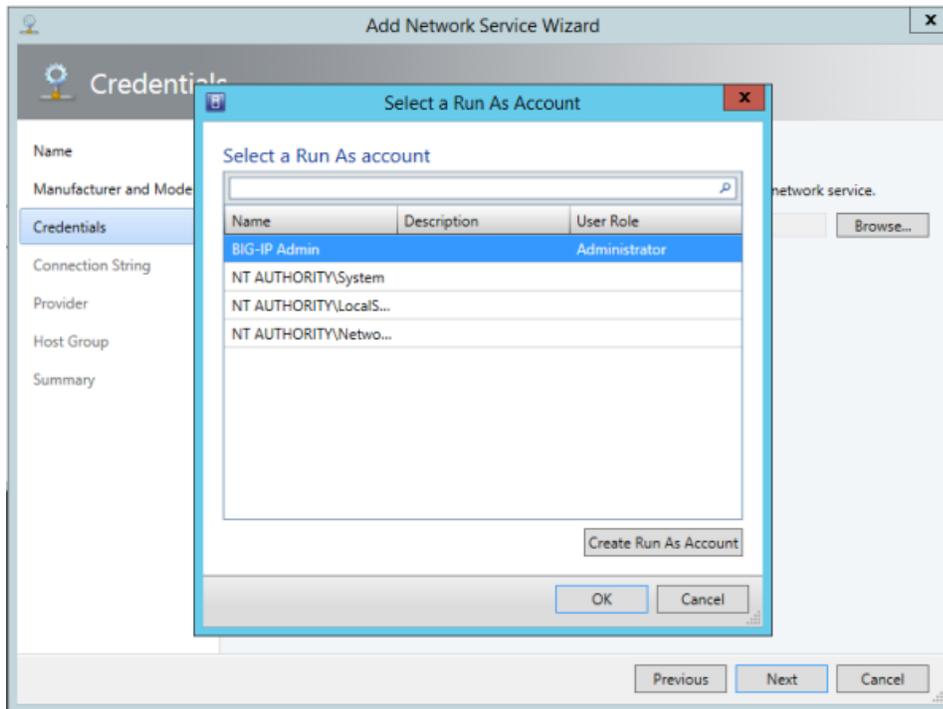
3. Type a name for the gateway, as shown, and then click **Next**.



4. For the **Manufacturer**, select **F5 Networks, Inc.**, and for the **Model**, select **BIG-IP**, as shown, and then click **Next**.

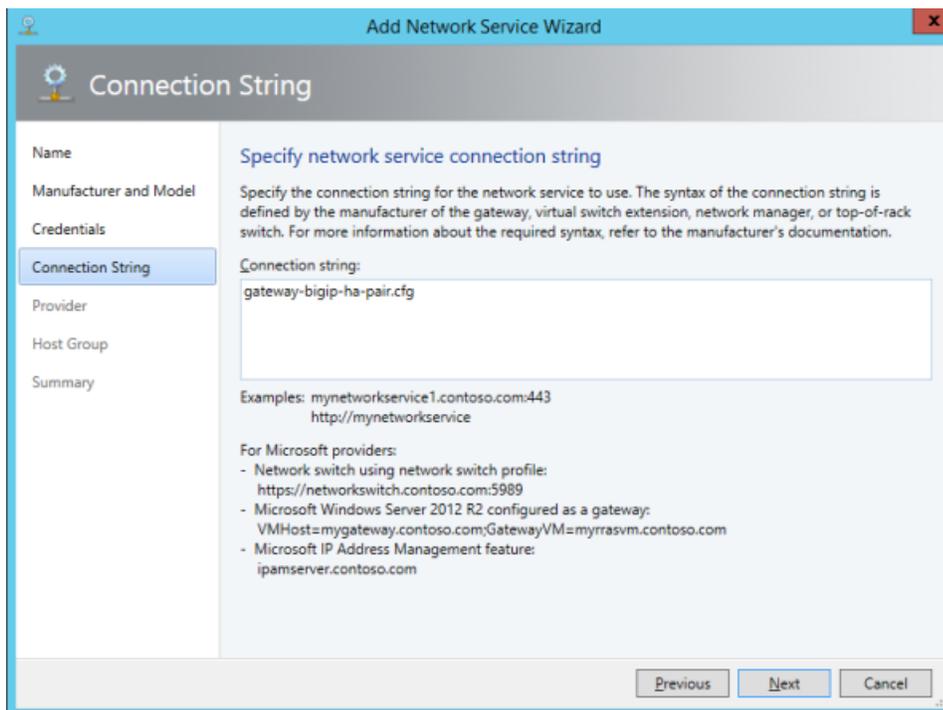


5. Select the Run As account you created previously, and then click **OK**.



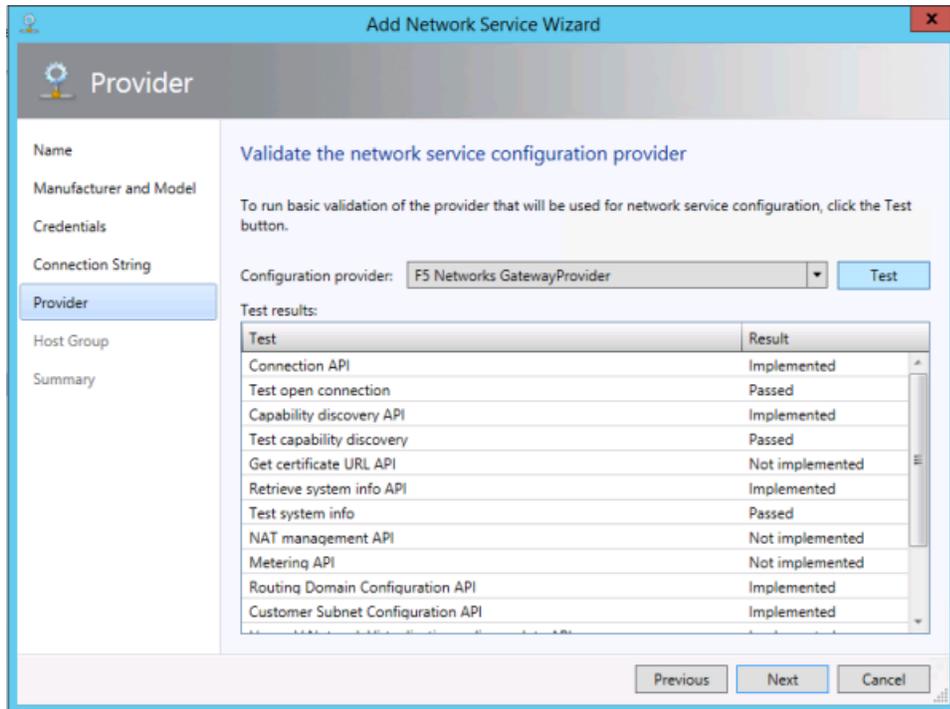
6. Specify the connection string for the gateway.

This is the name of the configuration file you edited and saved.

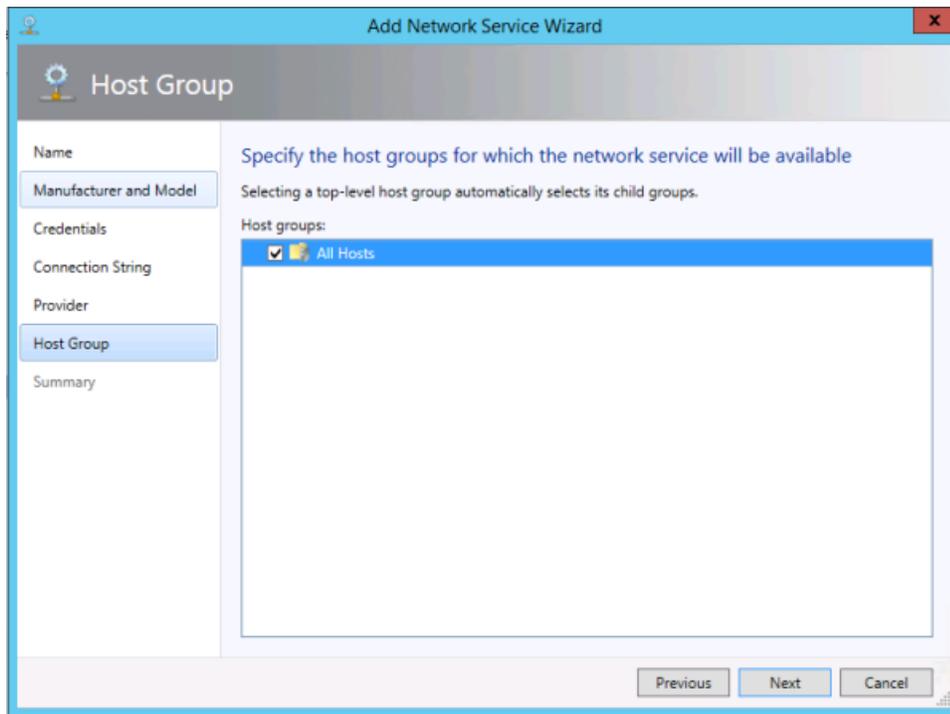


7. Click **Test** to test the gateway.

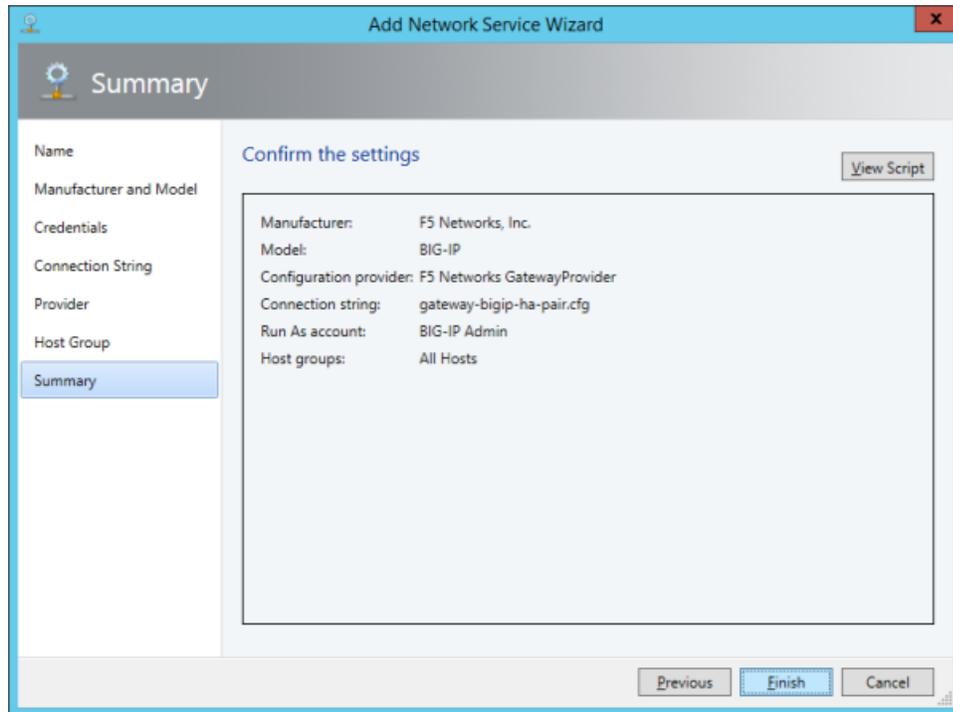
SCVMM exercises the gateway functions, as shown in this example. Only some of the network service functionality will be implemented.



8. Select a host group for the plug-in, as shown.

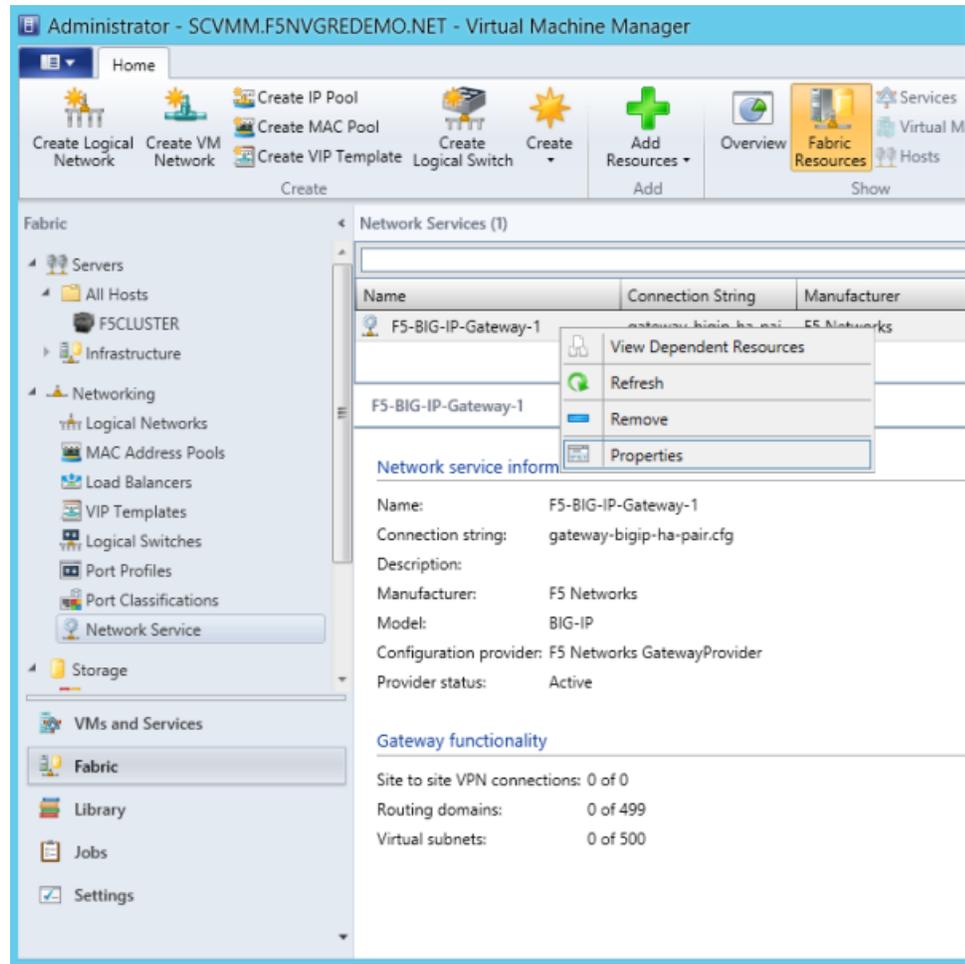


9. Confirm the settings you configured, as shown in this example, and then click **Finish**.



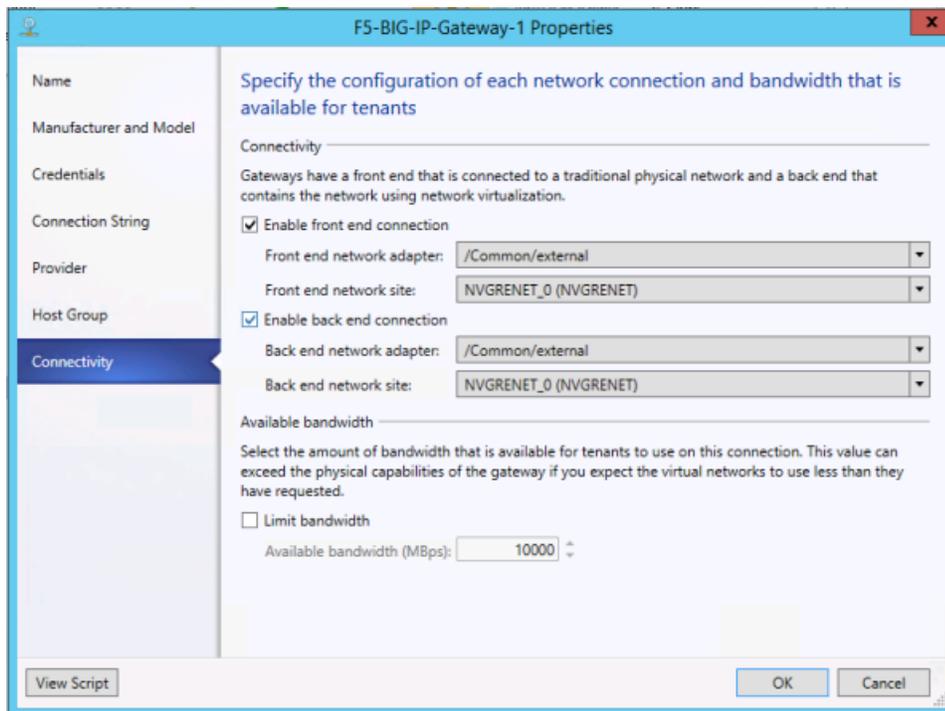
The SCVMM performs an initial configuration of the BIG-IP system. This might take a few seconds.

10. Right-click the gateway you created, and select **Properties**, as shown.



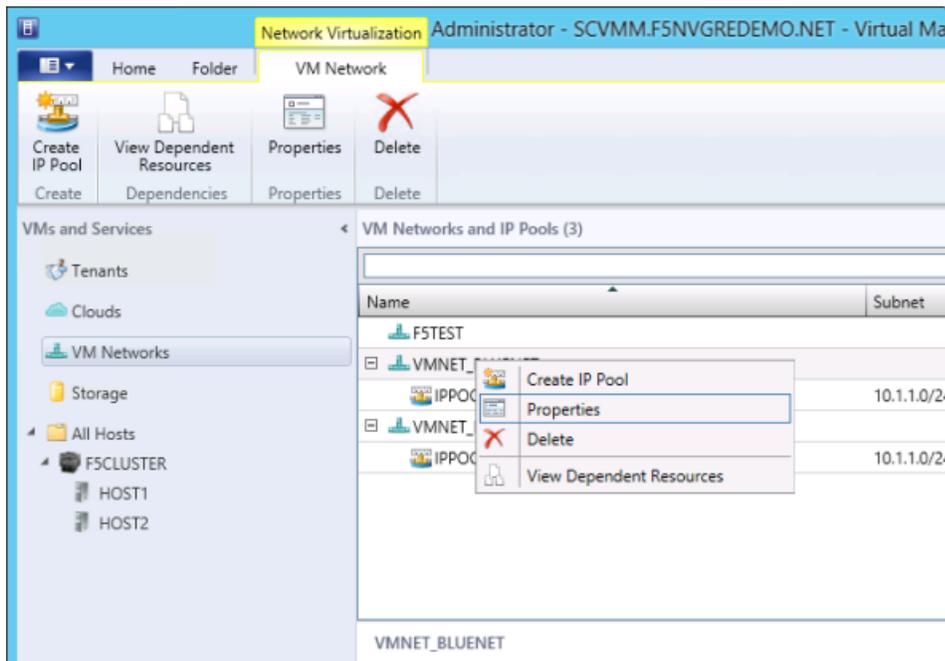
11. In the left navigation pane, select Connectivity, and specify the front end and back end SCVMM interfaces on the BIG-IP system, as shown, and then click **OK**.

Although the front end connection does not matter to the gateway plug-in, you must select one. For the back end connection, select the BIG-IP VLAN that has connectivity on the SCVMM provider network. Note that the BIG-IP system will be configured with a self IP address on this network, using a dynamically allocated address.

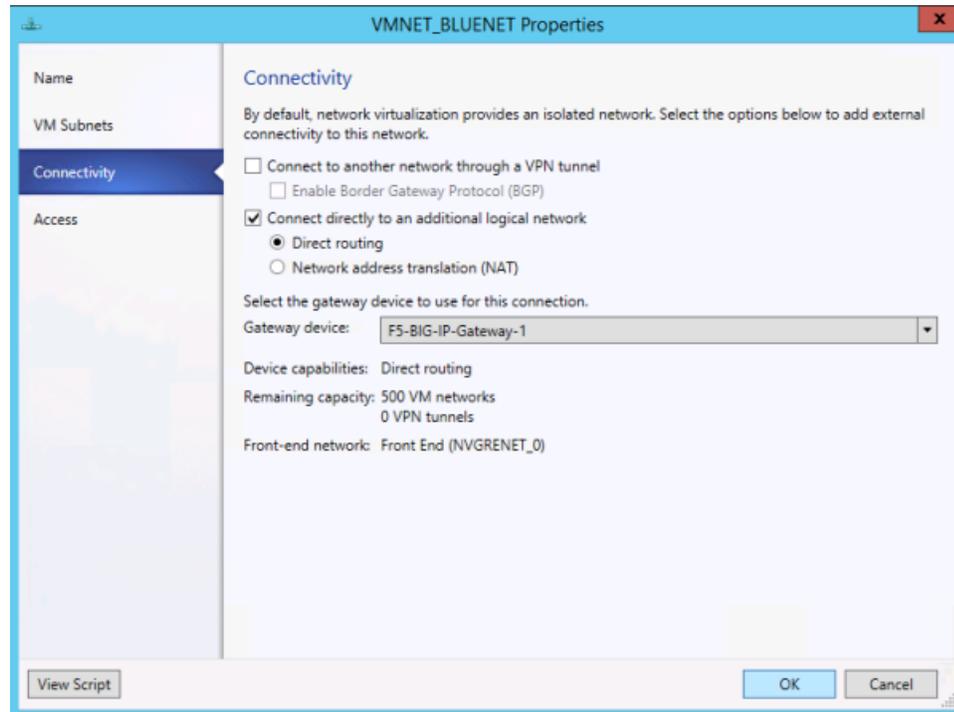


The gateway you created is now ready to use.

- In the VM Network tab, right-click the VM network on which you want the BIG-IP system to provide gateway services, and select **Properties**, as shown.



- In the left navigation pane, select Connectivity, select the check box **Connect directly to an additional logical network**, and then select the gateway you created, as shown.



14. When you have finished, click **OK**.

It might take a few seconds for completion of the BIG-IP configuration.

Viewing F5 Networks HNV Gateway PowerShell Module logs

To view logs for the F5 Networks HNV Gateway PowerShell Module, navigate to

```
C:\Program Files\Microsoft System Center 2012 R2\Virtual Machine
Manager\bin\F5_SCVMM_Gateway.Log
```

Example of F5 Networks NVGRE gateway environment

This illustration is an example of a configured F5 Networks NVGRE gateway environment.

Setting Up an F5 Networks NVGRE Gateway Environment

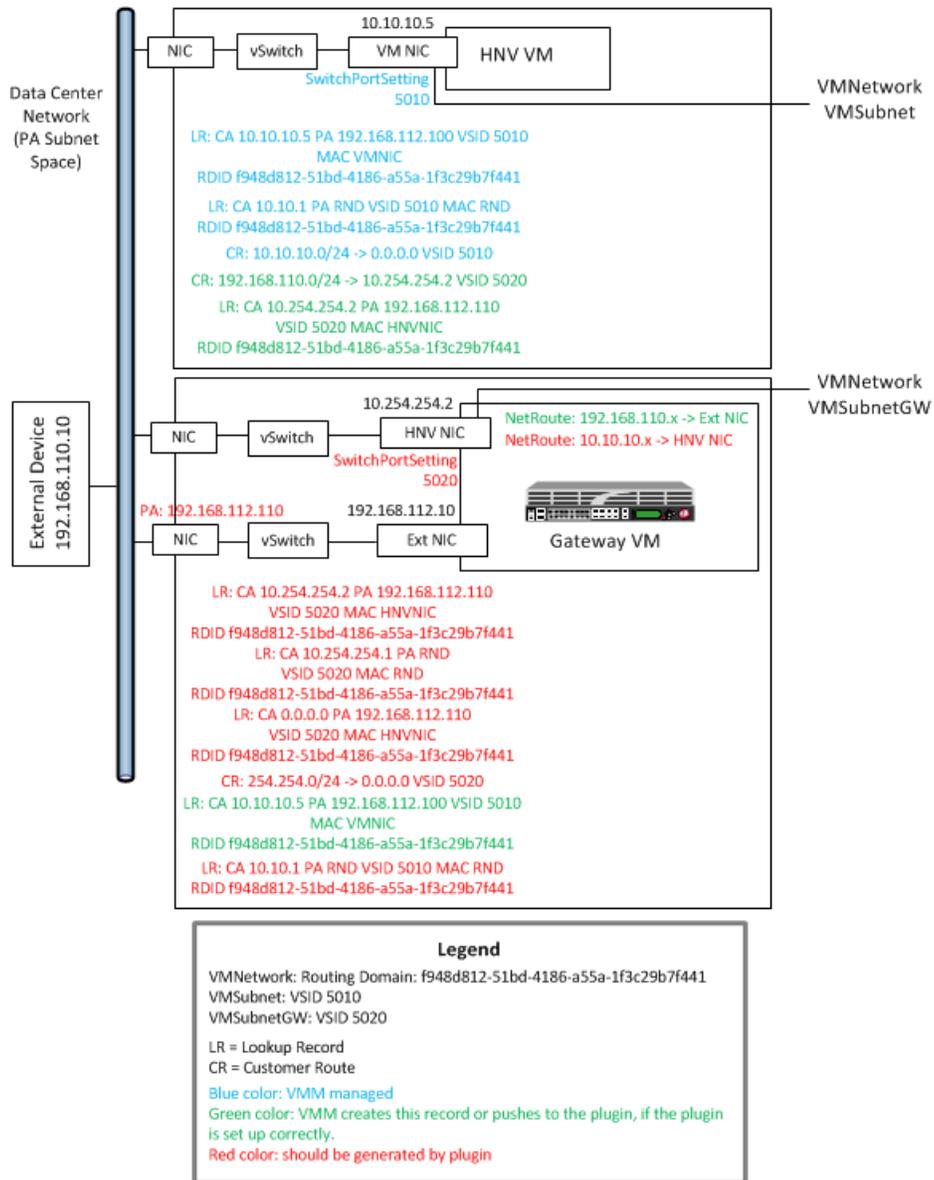


Figure 15: Example of F5 Networks NVGRE Gateway environment

Index

C

customer address (CA)
in NVGRE deployment, defined 8

D

device resolver group
creating for NVGRE gateway plug-in 18

H

HNV Gateway
installing 18
overview 8

I

IP address pools
in NVGRE deployment, described 13

L

logical networks
in NVGRE deployment, described 12
logical switches
in NVGRE deployment, described 14

N

network virtualization
and NVGRE 8
NVGRE
about 8
and customer addresses 8
and IP address pools 13
and logical networks 12
and logical switches 14
and network virtualization 8
and provider addresses 10
and routing domains 11
and virtual port profiles 15
and virtual subnets 10

NVGRE (*continued*)
and VM networks 15
NVGRE gateway plug-in
creating BIG-IQ device resolver group for 18
example of setup 29
overview 8
preparing the BIG-IP system for 16
NVGRE plug-in
viewing logs for 29

O

overlay networks
using NVGRE 8

P

provider address (PA)
in NVGRE deployment, defined 10

R

routing domains
in NVGRE deployment, described 11

S

SCVMM
configuring F5 NVGRE gateway for 21
installing F5 NVGRE gateway for 18
installing F5 NVGRE gateway plug-in for 19

V

virtual port profiles
in NVGRE deployment, described 15
virtual subnets
in NVGRE deployment, described 10
VM gateway
configuring to forward packets 21
VM networks
in NVGRE deployment, described 15

