

Deploying F5 with Microsoft Remote Desktop Gateway Servers

Welcome to the F5 deployment guide for Microsoft® Remote Desktop Services included in Windows® Server 2012 and Windows Server 2008 R2. This document provides guidance on configuring the BIG-IP Local Traffic Manager (LTM) for directing traffic and maintaining persistence to Microsoft Remote Desktop Gateway Services. It also provides guidance on how to configure Access Policy Manager to act as a secure HTTP proxy for RDP connections, as well as how to use the BIG-IP Advanced Firewall Manager (AFM) to provide a sophisticated layer of security for your Remote Desktop Gateway Server deployment.

Remote Desktop Services enables users to remotely access full Windows desktops, or individual Windows-based applications, on Remote Desktop Session Host computers. In an environment using BIG-IP LTM system, a farm of Remote Desktop Session Host servers has incoming connections distributed in a balanced manner across the members of the farm. BIG-IP APM can securely proxy RDP connections if using version 11.6 or later.

To provide feedback on this deployment guide or other F5 solution documents, contact us at solutionsfeedback@f5.com.

Visit the Microsoft page of F5's online developer community, DevCentral, for Microsoft forums, solutions, blogs and more: <http://devcentral.f5.com/Microsoft/>.

Products and versions

Product	Version
BIG-IP LTM, AFM	11.4 - 11.6
BIG-IP APM for acting as a secure HTTP proxy for RDP connections	11.6
Microsoft Windows Server Remote Desktop Services	2012 R2, 2012, 2008 R2
iApp version	f5.microsoft_rds_remote_access.v1.0.0 ¹
Deployment Guide version	1.2 (see <i>Document Revision History</i> on page 39)
Last updated	05-28-2015

¹ This iApp replaces the Remote Desktop Gateway/APM Native Proxy iApp (f5.microsoft_remote_desktop_gateway.v1.0.0rc1) on DevCentral and on downloads.f5.com.

Important: Make sure you are using the most recent version of this deployment guide, available at <http://www.f5.com/pdf/deployment-guides/microsoft-remote-desktop-gateway-dg.pdf>

For previous versions of this and other guides, see the Deployment guide *Archive tab* on f5.com: <https://f5.com/solutions/deployment-guides/archive-608>

Contents

What is F5 iApp?	3
Prerequisites and configuration notes	3
Configuration example	5
BIG-IP LTM only configuration example	5
BIG-IP APM configuration example	5
Using this guide	6
Configuring the iApp template for Remote Desktop Access with Remote Desktop Gateway Servers	7
Downloading and importing the new iApp	7
Starting the iApp for Microsoft Remote Desktop Gateway Remote Access	7
Next steps	20
Creating an NTLM Machine Account	21
Troubleshooting	22
Manual configuration tables	23
Supporting RemoteFX for Remote Desktop Gateway (optional)	24
Configuring the BIG-IP system for secure HTTP Proxy with BIG-IP APM	26
Manually configuring the BIG-IP Advanced Firewall Module to secure your RDG deployment	30
Appendix A: Configuring WMI monitoring of the RDS servers for LTM only	34
Appendix B: Using X-Forwarded-For to log the client IP address in IIS 7.0, 7.5, and 8 (optional)	36
Appendix C: Configuring DNS and NTP settings on the BIG-IP system	38
Document Revision History	39

What is F5 iApp?

New to BIG-IP version 11, F5 iApp is a powerful set of features in the BIG-IP system that provides a new way to architect application delivery in the data center, and it includes a holistic, application-centric view of how applications are managed and delivered inside, outside, and beyond the data center. The iApp template acts as the single-point interface for managing this configuration.

For more information on iApp, see the White Paper *F5 iApp: Moving Application Delivery Beyond the Network*:
<http://www.f5.com/pdf/white-papers/f5-iapp-wp.pdf>.

Skip ahead **Advanced**

If you are already familiar with this configuration, you can skip directly to the relevant section *after* reading the prerequisites:

- If using the iApp template, see *Configuring the iApp template for Remote Desktop Access with Remote Desktop Gateway Servers on page 7*
- If configuring the BIG-IP system manually, see *Manual configuration tables on page 23*

Prerequisites and configuration notes

Use this section for important items you need to know about and plan for before you begin this deployment. Not all items will apply in all implementations, but we strongly recommend you read all of these items carefully.

BIG-IP system and general prerequisites

- The BIG-IP LTM system must be running version 11.4 or later. If you want to use BIG-IP APM to securely proxy Remote Desktop connections, you must be using version 11.6 or later. For more detailed information on the BIG-IP system, see <http://www.f5.com/products/bigip/>.
- Although our examples and diagrams show external users connecting to the system in a routed configuration, the steps described in this guide are equally valid for a one-armed configuration, and both topologies may be used simultaneously.
- The BIG-IP LTM offers the ability to mix IPv4 and IPv6 addressing; for instance, you might want to use IPv6 addressing on your internal networks even though connections from clients on the Internet use IPv4.
- Be sure to see *Appendix A: Configuring WMI monitoring of the RDS servers for LTM only on page 34* and *Appendix B: Using X-Forwarded-For to log the client IP address in IIS 7.0, 7.5, and 8 (optional) on page 36* for optional configuration procedures.
- The third-party Web site information in this guide is provided to help you find the technical information you need. The URLs are subject to change without notice.
- For the BIG-IP LTM only configuration, you must create an RD Gateway Server Farm and add all members of farm (must match those in the BIG-IP LTM pool). Enable HTTPS - HTTP Bridging. For the SSL Certificate any setting will work, the BIG-IP LTM does SSL processing.
- If you are using BIG-IP APM as an RDP proxy and using security groups to determine host access, in Active Directory, you must create the security group and add the appropriate hosts you want to allow.

Remote Desktop Services and Windows Server prerequisites

- You must be using Windows Server 2008 R2 or 2012 or 2012 R2 Remote Desktop Services. If you are using a previous version see the Deployment Guide index at: <http://www.f5.com/solutions/resources/deployment-guides.html>.
- For information on Microsoft Windows Server, including Windows Remote Desktop Services, see one of the following:
 - » Windows Server 2012: technet.microsoft.com/library/hh831447
 - » Windows Server 2008 R2: technet.microsoft.com/en-us/library/dd647502%28WS.10%29.aspx
- For the BIG-IP LTM only configuration, you must install the Remote Desktop Gateway role on at least one server; for load balancing connections, you need at least two servers. See the Deploying Remote Desktop Gateway Step-by-Step Guide at: technet.microsoft.com/en-us/library/dd983941%28WS.10%29.aspx
- You must install the Remote Desktop Session Host role, or enable remote desktop services on hosts supporting Remote Desktop connections.
- If Broker services are required, install the Remote Desktop Connection Broker role on at least one server.

- Each user's Remote Desktop Connection client needs to be configured to use an RD Gateway Server. The configured Server Name must correspond to the fully-qualified DNS name that is associated with the Client SSL profile that you create on the BIG-IP LTM. Additionally, the certificate associated with that name and profile must be trusted by the client computer, and the client computer must be able to resolve the DNS name to the IP address assigned to the BIG-IP virtual server. Instructions for the various methods of client configuration can be found in the following Microsoft TechNet article: <http://technet.microsoft.com/en-us/library/cc772479.aspx>.

In our example, we show a manually configured Remote Desktop Connection client.

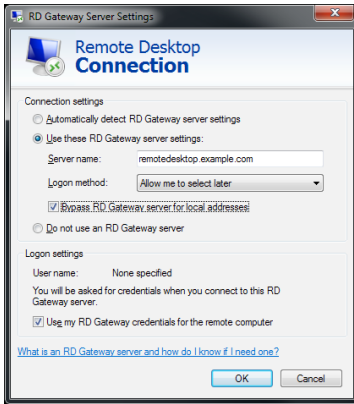


Figure 1: RD Gateway Server settings

For the BIG-IP LTM only configuration, in the following images, we show an example of a RD Gateway server that has been properly configured to participate in a RD Gateway server farm. In Figure 3, you can see that SSL Bridging has been enabled. Figure 4 shows that two members have been added to the farm.

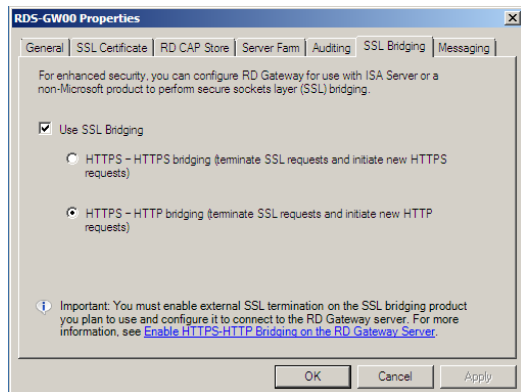


Figure 2: Configuring HTTPS-HTTP bridging on the TS Gateway server

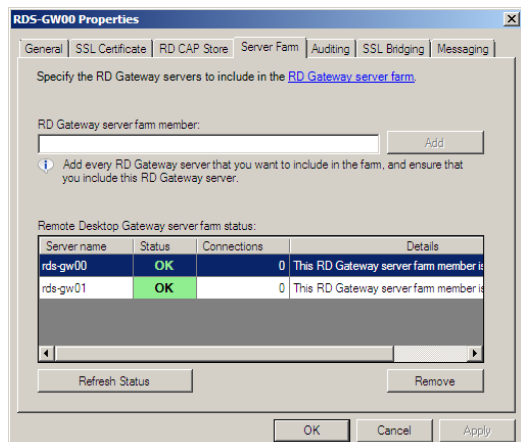


Figure 3: Configuring the Server Farm properties

For more information on configuring the Gateway Server role, see the Microsoft documentation.

Optional Modules

This iApp allows you to use four modules on the BIG-IP system. To take advantage of these modules, they must be licensed and provisioned before starting the iApp template. For information on licensing modules, contact your sales representative.

- **BIG-IP AFM**

BIG-IP Advanced Firewall Manager (AFM) is a high-performance, stateful, full-proxy network firewall designed to guard data centers against incoming threats that enter the network on the most widely deployed protocols—including HTTP/S, SMTP, DNS, and FTP. By aligning firewall policies with the applications they protect, BIG-IP AFM streamlines application deployment, security, and monitoring. For more information on BIG-IP AFM, see <https://f5.com/products/modules/advanced-firewall-manager>.

- **BIG-IP APM**

BIG-IP Access Policy Manager (APM) is a flexible, high-performance access and security solution that provides unified global access to your business-critical applications and networks. By consolidating remote access, web access management, VDI, and other resources in a single policy control point—and providing easy-to-manage access policies—BIG-IP APM helps you free up valuable IT resources and scale cost-effectively. See <http://www.f5.com/products/big-ip/big-ip-access-policy-manager/overview/>.

Configuration example

The iApp template and this deployment guide describe how to configure the BIG-IP system for Remote Desktop Gateway services. You can configure the BIG-IP system to use LTM only, or add BIG-IP APM to act as a secure HTTP proxy for RDP connections.

BIG-IP LTM only configuration example

In this scenario, we create a pool for the farm of RD Gateway Servers on the BIG-IP system. While still using the Remote Desktop Connection client, user RDP sessions are now encapsulated in HTTPS, which is more likely to be allowed through firewalls. When the HTTPS sessions arrive at the BIG-IP system, they are decrypted and passed to the pool of RD Gateway servers using HTTP. The RD Gateway Servers remove the HTTP, and forward the RDP sessions to the destination Remote Desktop server specified by the client.

BIG-IP APM configuration example

In this scenario, we use the BIG-IP Access Policy Manager to securely proxy Remote Desktop connections, so the deployment of Remote Desktop Gateway servers is not required. Continuing to use the Remote Desktop Connection client, the RDP sessions of the users are encapsulated in HTTPS, which is more likely to be allowed through firewalls. When the HTTPS sessions arrive at the BIG-IP system, they are decrypted and inspected to determine the Remote Desktop Session Host. The BIG-IP APM removes the HTTP, and forwards the RDP sessions to the destination Remote Desktop host specified by the client. APM offers several ways to verify host connection requests: you can specify an explicit list, you can use an AD security group where RD Session Hosts added to the group are allowed, you can allow any client requested host, or you can use a combination of an explicit list and group membership.

Using this guide

This deployment guide is intended to help users deploy the BIG-IP system. This document contains guidance configuring the BIG-IP system using the iApp template, as well as manually configuring the BIG-IP system.

Using this guide to configure the iApp template

We recommend using the iApp template to configure the BIG-IP system for your Microsoft RDS implementation. The majority of this guide describes the iApp template and the different options the template provides for configuring the system for Microsoft RDS Remote Access.

The iApp template configuration portion of this guide walks you through the entire iApp, giving detailed information not found in the iApp or inline help. The questions in the UI for the iApp template itself are all displayed in a table and at the same level. In this guide, we have grouped related questions and answers in a series of lists. Questions are part of an ordered list and are underlined and in italics or bold italics. Options or answers are part of a bulleted list, and in bold. Questions with dependencies on other questions are shown nested under the top level question, as shown in the following example:

1. *Top-level question found in the iApp template*

- *Select an object you already created from the list* (such as a profile or pool; not present on all questions. Shown in bold italic)
- **Choice #1** (in a drop-down list)
- **Choice #2** (in the list)

a. *Second level question dependent on selecting choice #2*

- **Sub choice #1**
- **Sub choice #2**

a. *Third level question dependent on sub choice #2*

- **Sub-sub choice**
- **Sub-sub #2**

a. *Fourth level question (rare)*

Advanced options/questions in the template are marked with the Advanced icon: **Advanced**. These questions only appear if you select the Advanced configuration mode.

Using this guide to manually configure the BIG-IP system

Users already familiar with the BIG-IP system can use the manual configuration tables to configure the BIG-IP system for the Remote Desktop implementation. These configuration tables only show the configuration objects and any non-default settings recommended by F5, and do not contain procedures on specifically how to configure those options in the Configuration utility. See *Manual configuration tables on page 23*.

Configuring the iApp template for Remote Desktop Access with Remote Desktop Gateway Servers

Use the following guidance to download the iApp template and configure the BIG-IP system for Microsoft Remote Desktop Gateway services. If you are configuring the BIG-IP system manually, see the *Manual configuration tables on page 23*.

Downloading and importing the new iApp

The first task is to download and import the iApp template.

To download and import the iApp

1. Open a web browser and go to downloads.f5.com.
2. Click **Find a Download**, and then click **BIG-IP v11.x / Virtual Edition**.
3. If necessary, select a BIG-IP product version from the list, and then click **iApp-Templates**.
4. Accept the EULA, and then download the iapps zip file to a location accessible from your BIG-IP system.
5. Extract (unzip) the **f5.microsoft_rds_remote_access.v<latest version>.tmpl** file.
6. Log on to the BIG-IP system web-based Configuration utility.
7. On the Main tab, expand **iApp**, and then click **Templates**.
8. Click the **Import** button on the right side of the screen.
9. Click a check in the **Overwrite Existing Templates** box.
10. Click the **Browse** button, and then browse to the location you saved the iApp file.
11. Click the **Upload** button. The iApp is now available for use.

Starting the iApp for Microsoft Remote Desktop Gateway Remote Access

To begin the Remote Desktop Gateway iApp Template, use the following procedure.

1. Log on to the BIG-IP system.
2. On the Main tab, expand **iApp**, and then click **Application Services**.
3. Click **Create**. The Template Selection page opens.
4. In the **Name** box, type a name. In our example, we use **rds-remote-access_**.
5. From the **Template** list, select **f5.microsoft_rds_remote_access.v1.0.0**. The iApp template opens.

Advanced options

If you select **Advanced** from the **Template Selection** list at the top of the page, you see Device and Traffic Group options for the application. This feature, new to v11, is a part of the Device Management configuration. This functionality extends the existing High Availability infrastructure and allows for clustering, granular control of configuration synchronization and granular control of failover. To use the Device and Traffic Group features, you must have already configured Device and Traffic Groups before running the iApp. For more information on Device Management, see the product documentation.

1. Device Group

To select a specific Device Group, clear the **Device Group** check box and then select the appropriate Device Group from the list.

2. Traffic Group

To select a specific Traffic Group, clear the **Traffic Group** check box and then select the appropriate Traffic Group from the list.

Template Options

This section contains general questions about the way you configure the iApp template.

1. Do you want to see inline help?

Choose whether you want to see informational and help messages inline throughout the template, or if you would rather hide this inline help. If you are unsure, we recommend having the iApp display the inline help.

Important and critical notes are always shown, no matter which selection you make.

- **Yes, show inline help text**

Select this option to see all available inline help text.

- **No, do not show inline help text**

If you are familiar with this iApp template, or with the BIG-IP system in general, select this option to hide the inline help text.

2. Which configuration mode do you want to use?


Select whether you want to use F5 recommended settings, or have more granular, advanced options presented.

- **Basic - Use F5's recommended settings**

In basic configuration mode, options like load balancing method and parent profiles are all set automatically. The F5 recommended settings come as a result of extensive testing with web applications, so if you are unsure, choose Basic.

- **Advanced - Configure advanced options**

In advanced configuration mode, you have more control over individual settings and objects, such as server-side optimizations and advanced options like Slow Ramp Time and Priority Group Activation. You can also choose to attach iRules you have previously created to the Application Service. The Advanced option provides more flexibility for experienced users.

As mentioned, advanced options in the template are marked with the Advanced icon: . If you are using Basic/F5 recommended settings, you can skip the questions with the Advanced icon.

3. Do you want to deploy BIG-IP APM as an RDP proxy?

This question only appears if you are using BIG-IP version 11.6 or later and have licensed and provisioned BIG-IP APM.

Select whether you want to deploy the BIG-IP APM to act as a Microsoft Remote Desktop Gateway to securely proxy Remote Desktop connections.

- **Yes, deploy BIG-IP APM as an RDP proxy**

Select this option if you want the template to configure the BIG-IP APM system to act as an RDP proxy. If you choose this option, the rest of the questions in this section disappear. Continue with the next section.

- **No, do not deploy BIG-IP APM as an RDP proxy**

Select this option if you do not want the template to configure BIG-IP APM as an RDP proxy.

- a. Which version of Windows Server are you deploying?

Select the version of Microsoft Windows Server you are using: 2008 R2, 2012, or 2012 R2. This determines the health monitor the iApp creates.

- **Windows Server 2008 R2**

Select this option if you are using Windows Server 2008 R2. No further information is needed.

- **Windows Server 2012**

Select this option if you are using Windows Server 2012.

- a. Will clients be connecting via UDP?

Select whether your clients will attempt to connect to the Windows Server 2012 or 2012 R2 servers using UDP. The BIG-IP system uses this information to create an additional virtual server for UDP traffic.

- **Windows Server 2012 R2**

Select this option if you are using Windows Server 2012.

- a. Will clients be connecting via UDP?

Select whether your clients will attempt to connect to the Windows Server 2012 or 2012 R2 servers using UDP. The BIG-IP system uses this information to create an additional virtual server for UDP traffic.

Advanced Firewall Manager (BIG-IP AFM)

This section gathers information about BIG-IP Advanced Firewall Manager if you want to use it to protect this implementation. For more information on configuring BIG-IP AFM, see <http://support.f5.com/kb/en-us/products/big-ip-afm.html>, and then select your version. This section only appears if you have BIG-IP AFM licensed and provisioned on your system.

1. Do you want to use BIG-IP AFM to protect your application?

Choose whether you want to use BIG-IP AFM, F5's network firewall, to secure this Remote Desktop Gateway deployment. If you choose to use BIG-IP AFM, you can restrict access to the Remote Desktop Gateway virtual server to a specific network or IP address. See the BIG-IP AFM documentation for specific details on configuring AFM.

- **No, do not use AFM to secure your application**

Select this option if you do not want to enable BIG-IP AFM at this time. You can always re-enter the template at a later date to enable BIG-IP AFM. Continue with the next section.

- **Select an existing AFM policy from the list**

If you already created a BIG-IP AFM policy for this implementation, select it from the list. Continue with **c**.

- **Yes, use F5's recommended AFM configuration**

Select this option if you want to enable BIG-IP AFM using F5's recommended configuration.

- a. Do you want to restrict access to your application by network or IP address?

Choose whether you want to restrict access to the Remote Desktop Gateway implementation via the BIG-IP virtual server.

- **No, do not restrict source addresses (allow all sources)**

By default, the iApp configures the Advanced Firewall module to accept traffic destined for the Remote Desktop Gateway virtual server from all sources. If you do not have a need to restrict access to the virtual server, leave this option selected and then continue with **b**.

- **Restrict source addresses**

Select this option to restrict access to the Remote Desktop Gateway virtual server by IP address or network address.

- a. What IP or network addresses should be allowed to access your application?

Specify the IP address or network access that should be allowed access to the Remote Desktop Gateway virtual server. You can specify a single IP address, a list of IP addresses separated by spaces (not commas or other punctuation), a range of IP addresses separated by a dash (for example **192.0.2.10-192.0.2.100**), or a single network address, such as **192.0.2.200/24**.

- b. How do you want to control access to your application from sources with a low reputation score?

The BIG-IP AFM uses an IP intelligence database to categorize IP addresses coming into the system. Choose what you want the system to do for sources that are attempting to access the Remote Desktop Gateway virtual server with a low reputation score. For more information, see the BIG-IP AFM documentation.

 **Important** You must have an active IP Intelligence license for this feature to function. See <https://f5.com/products/modules/ip-intelligence-services> for information.

- **Allow all sources regardless of reputation**

Select this option to allow all sources, without taking into consideration the reputation score.

- **Reject access from sources with a low reputation**

Select this option to reject access to the Remote Desktop Gateway virtual server from any source with a low reputation score.

- **Allow but log access from sources with a low reputation**

Select this option to allow access to the Remote Desktop Gateway virtual server from sources with a low reputation score, but add an entry for it in the logs.

- c. Would you like to stage a policy for testing purposes?

Choose whether you want to stage a firewall policy for testing purposes. A staged policy allows you to evaluate the effect a policy has on traffic by analyzing the system logs, without actually modifying traffic based on the firewall rules. You must already have a policy on the system in order to select it.

- **Do not apply a staging policy**

Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

- *Select an existing policy from the list*

If you have already created a firewall policy for this implementation, select it from the list. Only policies that already exist on the system appear in the list. To create a new policy, on the Main tab, click **Security > Network Firewall > Policies**. Specific instructions for creating a firewall policy is outside the scope of this iApp and deployment guide.

d. *Which logging profile would you like to use?*

Choose whether you or not you want to use a logging profile for this AFM implementation. You can configure the BIG-IP system to log detailed information about BIG-IP system Network Firewall events and store those logs on the BIG-IP system or a remote logging server (supports syslog and Splunk). If you want to use a logging profile, we recommend creating one outside this template. The list only contains profiles with Network Firewall enabled.

- **Do not apply a logging profile**

Select this option if you do not want to apply a logging profile at this time. You can always re-enter the template at a later date to add a logging profile. Continue with the next question.

- *Select an existing logging profile from the list*

If you have already created a logging profile for this implementation, select it from the list. You must create a profile before it is available in the list. To create a logging profile, on the Main tab, click **Security > Event Logs > Logging Profiles**. Specific instructions for creating a logging profile is outside the scope of this iApp and deployment guide. See the online help or the *About Local Logging with the Network Firewall* chapter of the **BIG-IP Network Firewall: Policies and Implementations** guide for more information.

Access Policy Manager

This section appears if you selected to deploy the BIG-IP APM as a RDP proxy. If you did not choose to deploy APM as a RDP proxy, continue with the next section. To use APM, it must be fully licensed and provisioned before starting the template.

1. ***Do you want to create a new AAA server, or use an existing AAA server?***

Choose whether you want the system to create a new BIG-IP APM AAA Server object, or if you have already created a custom AAA Server outside the template. The AAA server contains information about your Active Directory implementation. If you are unsure, select **Create a new AAA Server**.


- *Select the AAA Server you created from the list*

If you have previously created an AAA Server for your Remote Desktop implementation, select the existing object you created from the list. Only AAA Server objects with a Type of Active Directory appear in the list.

a. ***Which NTLM Machine Account should be used for Kerberos delegation?***

The NTLM Machine Account creates a computer account for this BIG-IP system in your Active Directory domain. You must have an existing NTLM Machine Account on this BIG-IP system in order to select it from this list. You cannot create a NTLM Machine Account from an iApp template.

Select the NTLM Machine Account you created from the list.

 **Important** *You must have an existing NTLM Machine Account object on your BIG-IP APM. If you do not see any items in this list, you must exit this template and create the NTLM Machine Account. See [Creating an NTLM Machine Account on page 21](#).*

- **Create a new AAA Server**

Select this option if you want the system to create a new AAA Server object for this implementation.

a. ***Which Active Directory servers in your domain can this BIG-IP system contact?***

Specify both the FQDN and IP address of each Active Directory server you want the BIG-IP APM to use for servicing authentication requests. Click **Add** to include additional servers.

b. ***What is the FQDN of your Active Directory implementation for your Remote Desktop users?***

Specify the FQDN of the Active Directory deployment for your Remote Desktop users. This is the FQDN for your entire domain, such as **example.com**, rather than the FQDN for any specific host.

c. ***Does your Active Directory domain allow anonymous binding?***

Select whether anonymous binding is allowed in your Active Directory environment.

- **Yes, anonymous binding is allowed**

Select this option if anonymous binding is allowed. No further information is required.

- **No, credentials are required for binding**

If credentials are required for binding, you must specify an Active Directory user name and password for use in the AAA Server.

- a. Which Active Directory user with administrative permissions do you want to use?

Type a user name with administrative permissions.

- b. What is the password associated with that account?

Type the associated password.

- d. Create a new health monitor for the Active Directory servers?

Specify whether you want the template to create a new LDAP monitor or a new ICMP monitor, or if you select an existing monitor. For more accurate monitoring, we recommend using an LDAP monitor.

- **Select an existing monitor for the Active Directory pool**

Select this option if you have already created a health monitor (only monitors with a Type of LDAP or External can be used) for the Active Directory pool that will be created by the template. If you want to create a health monitor, but have not already done so, you must exit the template and create the object before it becomes available from the list. The iApp allows you to select monitors that are a part of another iApp Application Service. If you select a monitor that is a part of another Application Service, be aware that any changes you make to the monitor in the other Application Service will apply to this Application Service as well.

- a. Which monitor do you want to use?

From the list, select the LDAP or External monitor you created to perform health checks for the Active Directory pool created by the template. Only monitors that have a Type value of LDAP or External appear in this list.

- **Yes, create a simple ICMP monitor**

Select this option if you only want a simple ICMP monitor for the Active Directory pool. This monitor sends a ping to the servers and marks the server UP if the ping is successful. Continue with the "What text should appear in the user access login prompt" question on this page.

- **Yes, create a new LDAP monitor for the Active Directory servers**

Select this option if you want the template to create a new LDAP monitor for the Active Directory pool. You must answer the following questions:

- a. Which Active Directory user name should the monitor use?

Specify an Active Directory user name for the monitor to use when attempting to log on as a part of the health check. This should be a user account created specifically for this health monitor, and must be set to never expire.

- b. What is the associated password?

Specify the password associated with the Active Directory user name.

- c. What is the LDAP tree for this user account?

Specify the LDAP tree for the user account. As noted in the inline help, ADSI editor, a tool for Active Directory LDAP administration, is useful for determining the correct LDAP tree value. For example, if the user name is 'user1' which is in the organizational unit 'F5 Users' and is in the domain 'f5.example.com', the LDAP tree would be: ou=F5 Users, dc=f5, dc=example, dc=com.

- d. Does your Active Directory domain require a secure protocol for communication?

Specify whether your Active Directory implementation requires SSL or TLS for communication, or does not require a secure protocol. This determines the port the health monitor uses.

- **No, a secure protocol is not required**

Select this option if your Active Directory domain does not require a secure protocol.

- **Yes, SSL communication is required**

Select this option if your Active Directory domain requires SSL communication. The health check uses port 636 as the Alias Service Port.


- **Yes, TLS communication is required**

Select this option if your Active Directory domain requires TLS communication. The health check uses port 389 as the Alias Service Port.

- e. Which NTLM machine account should be used for Kerberos delegation?

The NTLM Machine Account creates a computer account for this BIG-IP system in your Active Directory domain. You must have an existing NTLM Machine Account on this BIG-IP system in order to select it from this list. You cannot create a NTLM Machine Account from an iApp template.

Select the NTLM Machine Account you created from the list.

 **Important** You must have an existing NTLM Machine Account object on your BIG-IP APM. If you do not see any items in this list, you must exit this template and create the NTLM Machine Account. See [Creating an NTLM Machine Account on page 21](#).

Network

This section contains questions about your networking configuration.

1. What type of network connects clients to the BIG-IP system? **Advanced**

Choose the type of network that connects your clients to the BIG-IP system. If you choose WAN or LAN, the BIG-IP system uses this information to determine the appropriate TCP optimizations. If you choose WAN through another BIG-IP system, the system uses a secure and optimized tunnel (called an iSession tunnel) for traffic between BIG-IP systems on opposite sides of the WAN. Only choose this option if you have another BIG-IP system across the WAN that will be a part of this implementation.

- **Local area network (LAN)**

Select this option if most clients are connecting to the BIG-IP system on a LAN. This field is used to determine the appropriate TCP profile which is optimized for LAN clients. In this case, the iApp creates a TCP profile using the *tcp-lan-optimized* parent with no additional modifications.

- **Wide area network (WAN)**

Select this option if most clients are connecting to the BIG-IP system over a WAN. This field is used to determine the appropriate TCP profile which is optimized for WAN clients. In this case, the iApp creates a TCP profile using the *tcp-wan-optimized* parent with no additional modifications.

2. Which VLANs transport client traffic? **Advanced**

The BIG-IP system allows you to restrict client traffic to specific VLANs that are present on the system. This can provide an additional layer of security, as you can allow traffic only from the VLANs you choose. The VLAN objects must already be configured on this BIG-IP system before you can select them.

By default, all VLANs configured on the system are enabled and appear in the Selected list. Use the Move buttons (<<) and (>>) to adjust list membership. Only VLANs in the Selected list are allowed.

3. Where will the virtual servers be in relation to the Remote Desktop Gateway servers?

Select whether your BIG-IP virtual servers are on the same subnet as your Remote Desktop Gateway servers, or on different subnets. This setting is used to determine the SNAT (secure NAT) and routing configuration.

- **BIG-IP virtual server IP and Remote Desktop Gateway servers are on the same subnet**

If the BIG-IP virtual servers and Remote Desktop Gateway servers are on the same subnet, SNAT is configured on the BIG-IP virtual server and you must specify the number of concurrent connections.

- a. How many connections to you expect to each Remote Desktop Gateway server?

Select whether you expect more or fewer than 64,000 concurrent connections to each server. This answer is used to determine what type of SNAT that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device. The system configures SNAT Auto Map (fewer than 64,000) or a SNAT pool (more than 64,000).

- **Fewer than 64,000 concurrent connections**

Select this option if you expect fewer than 64,000 concurrent connections per Remote Desktop Gateway server. With this option, the system applies SNAT Auto Map, doesn't require any additional IP addresses, as an existing self IP address is used for translation. Continue with the next section.

- **More than 64,000 concurrent connections**

Select this option if you have a very large deployment and expect more than 64,000 connections at one time. The iApp creates a SNAT Pool, for which you need one IP address for each 64,000 connections you expect.

- a. Create a new SNAT pool or use an existing one?

If you have already created a SNAT pool on the BIG-IP system, you can select it from the list. Otherwise, the system creates a new SNAT pool with the addresses you specify.

- **Create a new SNAT pool**


Select this option (the default) to enable the system to create a new SNAT pool. You must specify the appropriate number of IP addresses in the next question.

a. What are the IP addresses you want to use for the SNAT pool?

Specify one otherwise unused IP address for every 64,000 concurrent connections you expect, or fraction thereof. Click **Add** for additional rows. Do not use any self IP addresses on the BIG-IP system.

- **Select a SNAT pool**

Select the SNAT pool you created for this deployment from the list.

 **Important** *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per Remote Desktop Gateway server is reached, new requests fail.*

- **BIG-IP virtual servers and Remote Desktop Gateway servers are on different subnets**

If the BIG-IP virtual servers and servers are on different subnets, the following question appears.

a. How have you configured routing on your Remote Desktop Gateway servers?

If you chose different subnets, this question appears asking whether the Remote Desktop Gateway servers use this BIG-IP system's self IP address as their default gateway. Select the appropriate answer.

- **Servers have a route to clients through the BIG-IP system**

Choose this option if the servers use the BIG-IP system as their default gateway. In this case, no configuration is needed to support your environment to ensure correct server response handling. Continue with the next section.

- **Servers do not have a route to clients through the BIG-IP system**

If the Remote Desktop Gateway servers do not use the BIG-IP system as their default gateway, SNAT is configured on the BIG-IP virtual server and you must select the expected number of concurrent connections in the next question.

a. How many connections to you expect to each Remote Desktop Gateway server?

Select whether you expect more or fewer than 64,000 concurrent connections to each server. This answer is used to determine what type of SNAT that system uses. A SNAT is an object that maps the source client IP address in a request to a translation address defined on the BIG-IP device. The system configures SNAT Auto Map (fewer than 64,000) or a SNAT pool (more than 64,000).

- **Fewer than 64,000 concurrent connections**

Select this option if you expect fewer than 64,000 concurrent connections per Remote Desktop Gateway server. With this option, the system applies SNAT Auto Map, doesn't require any additional IP addresses, as an existing self IP address is used for translation. Continue with the *SSL Encryption* section.

- **More than 64,000 concurrent connections**

Select this option if you have a very large deployment and expect more than 64,000 connections at one time. The iApp creates a SNAT Pool, for which you need one IP address for each 64,000 connections you expect.

a. Create a new SNAT pool or use an existing one?

If you have already created a SNAT pool on the BIG-IP system, you can select it from the list. Otherwise, the system creates a new SNAT pool with the addresses you specify.

- **Create a new SNAT pool**


Select this option (the default) to enable the system to create a new SNAT pool. You must specify the appropriate number of IP addresses in the next question.

a. Which IP addresses do you want to use for the SNAT pool?

Specify one otherwise unused IP address for every 64,000 concurrent connections you expect, or fraction thereof. Click **Add** for additional rows. Do not use any self IP addresses on the BIG-IP system.

- **Select a SNAT pool**

Select the SNAT pool you created for this deployment from the list.

 **Important** *If you choose more than 64,000 connections, but do not specify enough SNAT pool address(es), after the maximum connection limit of 64,000 concurrent connections per Remote Desktop Gateway server is reached, new requests fail.*

SSL Encryption

Before running the template you should have already imported a certificate and key onto the BIG-IP system. While the BIG-IP system does include a self-signed SSL certificate that can be used internally or for testing, we strongly recommend importing a certificate and key issued from a trusted Certificate Authority for processing client-side SSL.

For information on SSL certificates on the BIG-IP system, see the online help or the *Managing SSL Certificates for Local Traffic* chapter in the **Configuration Guide for BIG-IP Local Traffic Manager** available at <http://support.f5.com/kb/en-us.html>.

1. **How should the BIG-IP system handle SSL traffic?**

This question only appears if you are not using BIG-IP APM as an RDP proxy. If you are using APM as an RDP proxy, start with step a under SSL offload.

There are four options for configuring the BIG-IP system for SSL traffic (only two are available if you deployed BIG-IP APM in the previous section. Select the appropriate mode for your configuration.

- **Encrypt to clients, plaintext to servers (SSL Offload)**

Choose this method if you want the BIG-IP system to offload SSL processing from the servers. You need a valid SSL certificate and key for this method.

- a. **Which Client SSL profile do you want to use?**

Select whether you want the iApp to create a new Client SSL profile, or if you have already created a Client SSL profile which contains the appropriate SSL certificate and key.

Unless you have requirements for configuring specific Client SSL settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : SSL : Client** to create a Client SSL profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Select an existing Client SSL profile**

If you created a Client SSL profile for this implementation, select it from the list. If you select an existing Client SSL profile, the rest of the questions in this section disappear. Continue with the next section.

- **Create a new Client SSL profile**

Select this option for the iApp to create a new Client SSL profile

- a. **Which SSL certificate do you want to use?**

Select the SSL certificate you imported for this implementation.

- b. **Which SSL private key do you want to use?**

Select the associated SSL private key.

- c. **Which intermediate certificate do you want to use?** **Advanced**

If your deployment requires an intermediate or chain certificate, select the appropriate certificate from the list. Intermediate certificates are intended to create a chain of trust between the CA that signed the certificate and the CA that is already trusted by the recipient of the certificate. This allows the recipient to verify the validity of the certificates presented, even when the signing CA is unknown.

If you are using BIG-IP APM as an RDP proxy, this completes this section. Continue with Virtual Servers and Pools on page 15.

- **Terminate SSL from clients, re-encrypt to servers (SSL Bridging)**

Choose this method if you want the BIG-IP system to terminate SSL to process it, and then re-encrypt the traffic to the servers (SSL Bridging). You need a valid SSL certificate and key for the client-side, and optionally for the server-side (see #b).

- a. **Which Client SSL profile do you want to use?** **Advanced**

Select whether you want the iApp to create a new Client SSL profile, or if you have already created a Client SSL profile which contains the appropriate SSL certificate and key.

Unless you have requirements for configuring specific Client SSL settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : SSL : Client** to create a Client SSL profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Select an existing Client SSL profile**

If you created a Client SSL profile for this implementation select it from the list. If you select an existing Client SSL profile, the rest of the questions in this section disappear. Continue with the next section.

- **Create a new Client SSL profile**

Select this option for the iApp to create a new Client SSL profile

- a. **Which SSL certificate do you want to use?**

Select the SSL certificate you imported for this implementation.

b. Which SSL private key do you want to use?

Select the associated SSL private key.

c. Which intermediate certificate do you want to use? **Advanced**

If your deployment requires an intermediate or chain certificate, select the appropriate certificate from the list.

Intermediate certificates are intended to create a chain of trust between the CA that signed the certificate and the CA that is already trusted by the recipient of the certificate. This allows the recipient to verify the validity of the certificates presented, even when the signing CA is unknown.

b. Which Server SSL profile do you want to use?

Select whether you want the iApp to create the F5 recommended Server SSL profile, or if you want to choose a Server SSL profile you already created. In this scenario, the BIG-IP system is acting as an SSL client and by default, we assume the servers do not expect the BIG-IP system to present its client certificate on behalf of clients traversing the virtual server. If your servers expect the BIG-IP system to present a client certificate, you must create a custom Server SSL profile with the appropriate certificate and key.

The default, F5 recommended Server SSL profile uses the serverssl parent profile. For information about the ciphers used in the Server SSL profile, see <http://support.f5.com/kb/en-us/solutions/public/8000/800/sol8802.html>.

- **Encrypted traffic is forwarded without decryption (SSL pass-through)**

Choose this method if you do not want the BIG-IP system to do anything with encrypted traffic and simply send it to the Remote Desktop Gateway servers. This is similar to SSL bridging, although in this case the system does not decrypt then re-encrypt the traffic, it only sends it on to the servers without modification.

Virtual Servers and Pools

This section gathers information about your Remote Desktop Gateway deployment that is used in the BIG-IP virtual server and load balancing pool if you did not deploy the BIG-IP APM.

Note: This section appears if you chose not to deploy the BIG-IP APM as a RDP proxy.

If you chose to deploy APM, skip this section and continue with *Virtual Server on page 17*.

1. What IP address do you want to use for the virtual server(s)?

Type the IP address you want to use for the BIG-IP virtual server. This is the address clients use (or a DNS entry resolves to this address) to access the Remote Desktop Gateway deployment via the BIG-IP system. If you are using Windows Server 2012 or 2012 R2, and specified clients will use UDP, the system creates an additional virtual server using this address for UDP traffic.

2. What is the associated service port?

Type the port number to use for the BIG-IP virtual server. For Remote Desktop Gateway deployments, this is typically 80 (HTTP) or 443 (HTTPS). The default port in the box is based on your answer to the How should the system handle SSL traffic question.

3. Which FQDNs will clients use to access Remote Desktop Gateway?

Type each fully qualified domain name clients will use to access the Remote Desktop Gateway deployment. Click the **Add** button to insert additional rows. If you only have one FQDN, do not click Add.

4. Which HTTP profile do you want to use? **Advanced**

The HTTP profile contains settings for instructing the BIG-IP system how to handle HTTP traffic. Choose whether you want the iApp to create a new HTTP profile or if you have previously created an HTTP profile for this deployment.

Unless you have requirements for configuring specific HTTP settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : Services : HTTP** to create a HTTP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Select an existing HTTP profile from the list**

If you already created an HTTP profile for this implementation, select it from the list.

- **Create a new HTTP profile (recommended)**

Select this option if you want the iApp to create a new HTTP profile.

a. Should the BIG-IP system insert the X-Forwarded-For header? **Advanced**

This question only appears if you chose not to deploy APM as a RDP proxy.

Select whether you want the BIG-IP system to insert the X-Forwarded-For header in the HTTP header for logging purposes.

- **Insert the X-Forwarded-For header**

Select this option if you want the system to include the X-Forwarded-For header. You may have to perform additional configuration on your Remote Desktop Gateway servers to log the value of this header. For more information on configuring logging see *Appendix B: Using X-Forwarded-For to log the client IP address in IIS 7.0, 7.5, and 8 (optional) on page 36*.

- **Do not insert the X-Forwarded-For header**

Select this option if you do not want the system to include X-Forwarded-For in the HTTP header.

5. Do you want to create a new pool or use an existing one?

A load balancing pool is a logical set of servers, grouped together to receive and process traffic. When clients attempt to access the application via the BIG-IP virtual server, the BIG-IP system distributes requests to any of the servers that are members of that pool.

- **Select an existing pool**

If you have already created a pool for your Remote Desktop Gateway servers, you can select it from the list.

If you do select an existing pool, all of the rest of the questions in this section disappear.

- **Create a new pool**

Leave this default option to create a new load balancing pool and configure specific options.

- a. Which load balancing method do you want to use? **Advanced**

Specify the load balancing method you want to use for this Remote Desktop Gateway server pool. We recommend the default, **Least Connections (member)**.

- b. Should the BIG-IP system queue TCP requests?

Select whether you want the BIG-IP system to queue TCP requests. TCP request queuing provides the ability to queue connection requests that exceed the capacity of connections for a pool, as determined by the connection limit. Consequently, instead of dropping connection requests that exceed the capacity of a pool, TCP request queuing enables those connection requests to reside within a queue according to defined conditions until capacity becomes available. For more information on TCP Request Queuing, see the *Preventing TCP Connection Requests From Being Dropped* chapter in the **BIG-IP Local Traffic Manager: Implementations** guide, available on support.f5.com.

i Important *TCP Request Queuing is an advanced feature and should be used only if you understand how it will affect your deployment, including application behavior and BIG-IP performance. If you enable TCP Request Queuing, you must have a Connection Limit set on at least one of the nodes when configuring the Address/Port for the nodes.*

- **No, do not enable TCP request queuing (recommended)**

Select this option if you do not want the BIG-IP system to queue TCP requests.

- **Yes, enable TCP request queuing**

Select this option if you want to enable TCP request queuing on the BIG-IP system.

- a. What is the maximum number of TCP requests for the queue?

Type the maximum number of requests you want to queue. We do **not** recommend using 0, which means unlimited and is only constrained by available memory.

- b. How many milliseconds should requests remain in the queue?

Type a number of milliseconds for the TCP request timeout value.

- c. Use a Slow Ramp time for newly added servers? **Advanced**

With Slow Ramp, the BIG-IP system gradually adds connections to a newly-enabled or newly-added Remote Desktop Gateway server over a time period you specify, rather than sending a full proportion of the traffic immediately. Slow Ramp is essential when using the Least Connections load balancing method (our recommended method for Remote Desktop Gateway servers), as the BIG-IP system would otherwise send all new connections to a new server immediately, potentially overwhelming that server.

Select whether you want to use a Slow Ramp time.

- **Use Slow Ramp**

Select this option for the system to implement Slow Ramp time for this pool.

a. How many seconds should Slow Ramp time last?

Specify a duration in seconds, for Slow Ramp. The time period you select for Slow Ramp is highly dependent on the speed of your server hardware and the behavior of your web services. The default setting of 300 seconds (5 minutes) is very conservative in most cases.

- **Do not use Slow Ramp**

Select this option if you do not want to use Slow Ramp. If you select this option, we recommend you do not use the Least Connections load balancing method.

d. Do you want to give priority to specific groups of servers? **Advanced**

Select whether you want to use Priority Group Activation. Priority Group Activation allows you to segment your servers into priority groups. With Priority Group Activation, the BIG-IP system load balances traffic according to the priority number you assign to the pool members. A higher number indicates higher priority. Traffic is only sent to the servers with the highest priority, unless the number of available servers in that priority group falls below the value you specify as the minimum. The BIG-IP then sends traffic to the group of servers with the next highest priority, and so on.

- **Do not use Priority Group Activation (recommended)**

Select this option if you do not want to enable Priority Group Activation.

- **Use Priority Group Activation**

Select this option if you want to enable Priority Group Activation.

You must add a priority to each server in the Priority box described in #c.

a. What is the minimum number of active members for each priority group?

Specify the minimum number of servers that must be active to continue sending traffic to the priority group. If the number of active servers falls below this minimum, traffic will be sent to the group of servers with the next highest priority group number.

e. Which servers are a part of this pool?

Specify the IP address(es) of your Remote Desktop Gateway servers. If you have existing nodes on this BIG-IP system, you can select them from the list, otherwise type the addresses. You can optionally add a Connection Limit. If you enabled Priority Group Activation, you must also specify a Priority for each device. Click **Add** to include additional servers.

Virtual Server

This section gathers information about your Remote Desktop Gateway deployment that is used in the BIG-IP virtual server.

Note: This section appears if you chose to deploy the BIG-IP APM as a RDP proxy. If you did not, continue with the next section.

1. What IP address do you want to use for the virtual server(s)?

Type the IP address you want to use for the BIG-IP virtual server. This is the address clients use (or a DNS entry resolves to this address) to access the Remote Desktop Gateway deployment via the BIG-IP system. If you are using Windows Server 2012 or 2012 R2, and specified clients will use UDP, the system creates an additional virtual server using this address for UDP traffic.

2. What is the associated service port?

Type the port number you want to use for the BIG-IP virtual server. For Remote Desktop Gateway deployments, this is typically 80 (HTTP) or 443 (HTTPS). The default port is based on your answer to the How should the system handle SSL traffic question.

3. Which HTTP profile do you want to use? **Advanced**

The HTTP profile contains settings for instructing the BIG-IP system how to handle HTTP traffic. Choose whether you want the iApp to create a new HTTP profile or if you have previously created an HTTP profile for this deployment.

Unless you have requirements for configuring specific HTTP settings, we recommend allowing the iApp to create a new profile. To select a profile from the list, it must already be present on the BIG-IP system. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : Services : HTTP** to create a HTTP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **Select an existing HTTP profile from the list**

If you already created an HTTP profile for this implementation, select it from the list.

- **Create a new HTTP profile (recommended)**

Select this option if you want the iApp to create a new HTTP profile.

4. How would you like to secure your hosts?

Select how the system should secure your RDS hosts. If you do not want the system to secure your hosts, select Allow all hosts.

- **Allow all hosts**

Select this option if you do not want the system to secure the hosts at all, and want to allow all hosts. Continue with the following section.

- **Secure hosts using an explicit list and group membership**

Select this option if you want to secure your hosts using both an explicit list and Active Directory group membership.

- a. Are users required to use FQDNs?

Select whether your users are required to use FQDNs, or if the host name can be specified in simple (host name only) format as well.

- b. Which hosts should be allowed?

Specify the host names that users are allowed to access. You can use IP address, host name, or FQDN format in this field.

- c. What AD security group contains your allowed hosts?

Specify the AD group you created in your Active Directory domain that contains the list of hosts you want to allow.

- **Secure hosts using an explicit list**

Select this option if you want to secure your hosts using an explicit list that you specify.

- a. Which hosts should be allowed?

Specify the host names that users are allowed to access. You can use IP address, host name, or FQDN format in this field

- **Secure hosts using group membership**

Select this option if you want to secure your hosts using Active Directory group membership.

- a. Are users required to use FQDNs?

Select whether your users are required to use FQDNs, or if the host name can be specified in simple (host name only) format as well. What AD security group contains your allowed hosts?

- b. What AD security group contains your allowed hosts?

Specify the AD group you created in your Active Directory domain that contains the list of hosts you want to allow.

This completes the iApp template if you are using the BIG-IP APM as an RDP proxy. Continue with *Finished on page 19*.

Application Health

In this section, you answer questions about how you want to implement application health monitoring on the BIG-IP system.

1. Create a new health monitor or use an existing one?

Application health monitors are used to verify the content that is returned by a request. The system uses these monitors to ensure traffic is only sent to available Remote Desktop Gateway servers.

Unless you have requirements for configuring other options not in the following list of questions, we recommend allowing the iApp to create a new monitor. Creating a custom health monitor is not a part of this template; see **Local Traffic >> Monitors**. To select any new monitors you create, you need to restart or reconfigure this template.

- **Select the monitor you created from the list**

If you manually created the health monitor, select it from the list.
Continue with the next section.

- **Create a new health monitor**

If you want the iApp to create a new monitor, continue with the following.

- a. How many seconds should pass between health checks?

Specify how long the system should wait between each health check. This is used as the Interval setting for the monitor.
We recommend the default of 30 seconds.

- b. What user name should the monitor use?

Type the domain and user name for the account you created for the health monitor. You must include the domain in front of the user, such as EXAMPLE\USER.

- c. What is the associated password?

Type the password for the account.

Client Optimization

In this section, you answer a question on how you want to optimize client-side connections. This determines the type of TCP profile the iApp assigns to the virtual server. This entire section only appears if you selected the Advanced configuration mode.

1. How do you want to optimize client-side connections? **Advanced**

The client-side TCP profile optimizes the communication between the BIG-IP system and the server by controlling the behavior of the traffic which results in higher transfer rates, improved connection reliability and increased bandwidth efficiency.

Unless you have requirements for configuring specific TCP optimization settings, we recommend allowing the iApp to create a new profile. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : Protocol : TCP** to create a TCP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **New profile based on tcp-wan-optimized (recommended)**
Select this option to have the system create the recommended TCP profile optimized for WAN connections.
- **Select the TCP profile you created from the list**
If you created a custom TCP profile for the client-side connections, select it from the list.

Server Optimization

In this section, you answer a question on how you want to optimize server-side connections. This determines the type of TCP profile the iApp assigns to the virtual server. This entire section only appears if you selected the Advanced configuration mode.

1. How do you want to optimize server-side connections? **Advanced**

The client-side TCP profile optimizes the communication between the BIG-IP system and the server by controlling the behavior of the traffic which results in higher transfer rates, improved connection reliability and increased bandwidth efficiency.

Unless you have requirements for configuring specific TCP optimization settings, we recommend allowing the iApp to create a new profile. Creating a custom profile is not a part of this template; see **Local Traffic >> Profiles : Protocol : TCP** to create a TCP profile. To select any new profiles you create, you need to restart or reconfigure this template.

- **New profile based on tcp-lan-optimized (recommended)**
Select this option to have the system create the recommended TCP profile optimized for LAN connections.
- **Select the TCP profile you created from the list**
If you created a custom TCP profile for server-side connections, select it from the list.


iRules

In this section, you can add custom iRules to the Remote Desktop Gateway deployment. This section is available only if you selected Advanced mode.

iRules are a scripting language that allows an administrator to instruct the system to intercept, inspect, transform, direct and track inbound or outbound application traffic. An iRule contains the set of instructions the system uses to process data flowing through it, either in the header or payload of a packet.

1. Do you want to add any custom iRules to the APM virtual server? **Advanced**

Select if have preexisting iRules you want to add to your Remote Desktop Gateway implementation.

 **Warning** *Improper use or misconfiguration of an iRule can result in unwanted application behavior and poor performance of your BIG-IP system. We recommended you verify the impact of an iRule prior to deployment in a production environment.*

If you do not want to add any iRules to the configuration, continue with the following section.

If you have iRules you want to attach to the APM virtual server the iApp creates for your Remote Desktop Gateway servers, from the **Options** box, click the name of the applicable iRule(s) and then click the Add (<<) button to move them to the **Selected** box.

Finished

Review the answers to your questions. When you are satisfied, click the **Finished** button. The BIG-IP system creates the relevant objects for the Remote Desktop Gateway application.

Next steps

After completing the iApp Template, the BIG-IP Application Services page opens for the Microsoft Remote Desktop service you just created. To see the list of all the configuration objects created to support the application, on the Menu bar, click **Components**. The complete list of all related objects opens. You can click individual objects to see the settings.

Once the objects have been created, you are ready to use the new deployment.

Modifying DNS settings to use the BIG-IP virtual server address

Before sending traffic to the BIG-IP system, your DNS administrator may need to modify any DNS entries for the RDS implementation to point to the BIG-IP system's virtual server address.

Modifying the iApp configuration

The iApp Application Service you just created can be quickly and easily modified if you find it necessary to make changes to the configuration. The Strict Updates feature of the iApp prevents users from manually modifying the iApp configuration (Strict Updates can be turned off, but use extreme caution). iApp allows you to re-enter the template, make changes, and then update the template. The modifications are automatically made to any of the associated objects.

To modify the configuration

1. On the Main tab, expand **iApp** and then click **Application Services**.
2. Click the name of your Remote Desktop Application Service from the list.
3. On the Menu bar, click **Reconfigure**.
4. Make the necessary modifications to the template.
5. Click the **Finished** button.

Viewing statistics

You can easily view a number of different statistics on the BIG-IP system related to the configuration objects created by the iApp template.

Object-level statistics

If you want to view object-level statistics, use the following procedure.

To view object-level statics

1. On the Main tab, expand **Overview**, and then click **Statistics**.
2. From the **Statistics Type** menu, you can select **Virtual Servers** to see statistics related to the virtual servers.
3. You can also choose **Pools** or **Nodes** to get a closer look at the traffic.
4. To see Networking statistics in a graphical format, click **Dashboard**.

For more information on viewing statistics on the BIG-IP system, see the online help or product documentation.

Creating an NTLM Machine Account

If you chose to use deploy the BIG-IP APM as an RDP proxy, you must have an NTLM Machine Account object configured before you can successfully complete the template. Use the following procedure to create the NTLM Machine Account.

To create the NTLM Machine Account

1. On the Main tab, expand **Access Policy**, and then click **Access Profiles**.
2. On the Menu bar, from the **NTLM** menu, click **Machine Account List**.
3. Click the **Create** button.
4. In the **Name** box, type a name for the BIG-IP Machine Account object. Currently, the NTLM machine account should contain alphanumeric characters and underscores only. Spaces and special characters are not allowed.
5. In the **Machine Account Name** box, type the name of the computer account that will be created in the domain after clicking Join.
6. In the **Domain FQDN** box, type the fully qualified domain name of the domain that you want the machine account to join.
7. In the **Domain Controller FQDN** box, if the machine account should have access to one domain only, type the FQDN for the domain controller for that domain.
8. In the **Admin User** box, type the name of a user with administrative privileges.
9. In the **Password** box, type the associated password.
10. Click the **Join** button.

Troubleshooting

Use this section for troubleshooting advice for common problems.

➤ **You may experience deployment errors when the NTLM Machine Account name contains spaces or special characters**

If you are using BIG-IP APM as an RDP proxy, you must specify an NTLM Machine Account in the iApp template that you created manually. If the name of the NTLM Machine Account object contains spaces or special characters, you may experience errors when trying to deploy the template.

If your NTLM Machine Account object name contains a special character or space, the workaround for this issue is to create an NTLM Machine Account name that only contains alphanumeric characters and underscores, with no spaces. Return to *Creating an NTLM Machine Account on page 21*, and create a new machine account using a name with no spaces or special characters.

➤ **You may experience a Failed to Connect error after attempting to connect to the BIG-IP virtual server using the default Remote Desktop client in Windows 7**

After attempting to connect to the BIG-IP virtual server address using the Windows 7 default Remote Desktop client, you may receive the following error message after about 30 seconds of waiting: "failed to connect".

This message is most likely due to using an unsupported Remote Desktop client. Verify you are using a current client by using the latest BIG-IP APM support matrix, found in the following guide

https://support.f5.com/kb/en-us/products/big-ip_apm/manuals/product/apm-clientcompatmatrix-11-6-0.html.

Download and install the latest client if you are not running a supported version:

<https://support.microsoft.com/en-us/kb/2592687>

Manual configuration tables

The following table contains a list of BIG-IP LTM configuration objects along with any non-default settings you should configure as a part of this deployment scenario. Unless otherwise specified, settings not mentioned in the table can be configured as applicable for your configuration. For specific instructions on configuring individual objects, see the online help or product manuals.

 **Note:** The health monitor requires a user account with permission to access the Remote Desktop Gateway. This is defined in the Client Access Policy on the RDG server. We recommend creating a user specifically for the health monitors.

Health Monitors (Main tab > Local Traffic > Monitors)		
Name	Type a unique name	
Type	HTTP (use HTTPS if configuring SSL Bridging)	
Interval	30 (recommended)	
Timeout	91 (recommended)	
Send String	RPC_IN_DATA /rpc/en-us/rpcproxy.dll HTTP/1.1\r\nHost: rdg.example.com (replace red text with your host name)	
Receive String	200 Success	
User Name	Type the user name of an account with RDG access.	
Password	Type the associated password	
Pools (Main tab > Local Traffic > Pools)		
Name	Type a unique name	
Health Monitor	Select the monitor you created above	
Slow Ramp Time ¹	300	
Load Balancing Method	Choose a load balancing method. We recommend Least Connections (Member)	
Address	Type the IP Address of the <u>RD Desktop Gateway</u> nodes. This can be an IPv4 or IPv6 address.	
Service Port	80 (use 443 if configuring SSL Bridging) Click Add , and repeat Address and Port for all nodes	
Profiles (Main tab > Local Traffic > Profiles)		
TCP (Profiles-->Protocol)	Name	Type a unique name
	Parent Profile	tcp-wan-optimized or tcp-lan-optimized depending on where the clients are located
HTTP (Profiles-->Services)	Name	Type a unique name
	Parent Profile	HTTP
Client SSL (Profiles-->SSL)	Name	Type a unique name
	Parent Profile	clientssl
	Certificate	Select the certificate you imported
	Key	Select the associated key
Server SSL ² (Profiles-->SSL)	Name	Type a unique name
	Parent Profile	serverssl
Virtual Servers (Main tab > Local Traffic > Virtual Servers)		
Name	Type a unique name.	
Address	Type the IP Address for the virtual server	
Service Port	443	
Protocol Profile (client) ¹	Select the TCP profile you created above	
Protocol Profile (server) ¹	Select the TCP profile you created above	
HTTP Profile	Select the HTTP profile you created above	
SSL Profile (Client)	Select the Client SSL profile you created above	
SSL Profile (Server)	<i>If configuring SSL Bridging Only:</i> Select the Server SSL profile you created above	
SNAT Pool ³	Auto Map (optional; see footnote ³)	
Default Pool	Select the pool you created above	

¹ You must select **Advanced** from the **Configuration** list for these options to appear

² A Server SSL profile is only required if you are configuring SSL Bridging

³ If want to use SNAT, and you have a large deployment expecting more than 64,000 simultaneous connections, you must configure a SNAT Pool with an IP address for each 64,000 simultaneous connections you expect. See the BIG-IP documentation on configuring SNAT Pools.

Supporting RemoteFX for Remote Desktop Gateway (optional)

If you are using BIG-IP LTM and Microsoft RemoteFX for Remote Desktop Services, use the following table to configure additional BIG-IP LTM objects for the Remote Desktop Gateway servers.

Health Monitors (Main tab > Local Traffic > Monitors)		
UDP Monitor		
Name	Type a unique name	
Type	UDP	
Interval	30 (recommended)	
Timeout	91 (recommended)	
ActiveSync monitor		
Name	Type a unique name	
Type	Gateway ICMP	
Interval	30 (recommended)	
Timeout	91 (recommended)	
Pools (Main tab > Local Traffic > Pools)		
Name	Type a unique name	
Health Monitor	Select both monitors you created above (ensure Availability Requirement is set to All (the default)	
Slow Ramp Time¹	300	
Load Balancing Method	Choose a load balancing method. We recommend Least Connections (Member)	
Address	Type the IP Address of a Remote Desktop Gateway device	
Service Port	3391 Click Add, and repeat Address and Port for all Remote Desktop Gateway devices	
Profiles (Main tab > Local Traffic > Profiles)		
Persistence (Profiles-->Persistence)	Name	Type a unique name
	Persistence Type	Source Address Affinity
	Match Across Services	Enabled
Virtual Servers (Main tab > Local Traffic > Virtual Servers)		
Name	Type a unique name.	
Address	Type the same IP Address you used for the Remote Desktop Gateway virtual server on the previous page.	
Service Port	3391	
SNAT Pool ²	Auto Map (optional; see footnote ²)	
Default Pool	Select the Remote Desktop Gateway pool you created above	
Default Persistence Profile	Select the Persistence profile you created	

¹ You must select **Advanced** from the **Configuration** list for these options to appear

² If want to use SNAT, and you have a large deployment expecting more than 64,000 simultaneous connections, you must configure a SNAT Pool with an IP address for each 64,000 simultaneous connections you expect. See the BIG-IP documentation on configuring SNAT Pools.

Modifying the RD Gateway virtual server to use the Persistence profile you created

The final task is to modify the Remote Desktop Gateway virtual server you configured (using the guidance on the previous page) to use the persistence profile you just created for RemoteFX as a fallback method.

To modify the virtual server

1. Expand **Local Traffic** and then click **Virtual Servers**.
2. Click the name of the TCP RD Gateway virtual server you created using the guidance from the table on page 5.
3. On the Menu bar, click **Resources**.
4. From the **Fallback Persistence Profile** list, select the name of the Source Address Affinity persistence profile you just created.
5. Click **Update**. This completes the RemoteFX configuration.

Deploying a reverse proxy virtual server for Remote Desktop Gateway

This section describes how to publish Remote Desktop Gateway services in a perimeter or DMZ network. This virtual server forwards traffic to the internal virtual server you just created.

Configuration table for the reverse proxy virtual server

Health Monitors (Main tab > Local Traffic > Monitors)		
Name	Type a unique name	
Type	HTTP (use HTTPS if configuring SSL Bridging)	
Interval	30 (recommended)	
Timeout	91 (recommended)	
Send String	RPC_IN_DATA /rpc/en-us/rpcproxy.dll HTTP/1.1\r\nHost: rdg.example.com (replace red text with your host name)	
Receive String	200 Success	
User Name	Type the user name of an account with RDG access.	
Password	Type the associated password	
Pools (Main tab > Local Traffic > Pools)		
Name	Type a unique name	
Health Monitor	Select the monitor you created above	
Address	Only add the IP address for the internal Remote Desktop Gateway virtual server you created	
Service Port	443	
Profiles (Main tab > Local Traffic > Profiles)		
TCP (Profiles-->Protocol)	Name	Type a unique name
	Parent Profile	Use tcp-wan-optimized or tcp-lan-optimized depending on where your clients are located.
Client SSL (Profiles-->SSL)	Name	Type a unique name
	Parent Profile	clientssl
	Certificate	Select the certificate you imported
Server SSL ² (Profiles-->SSL)	Key	Select the associated key
	Name	Type a unique name
	Parent Profile	serverssl
Virtual Servers (Main tab > Local Traffic > Virtual Servers)		
Name	Type a unique name.	
Address	Type the External IP address for Remote Desktop Gateway connections as the IP address for this virtual server	
Service Port	443	
Protocol Profile (client) ¹	Select the TCP profile you created above	
Protocol Profile (server) ¹	Select the TCP profile you created above	
SSL Profile (Client)	Select the Client SSL profile you created above	
SSL Profile (Server)	<i>If configuring SSL Bridging Only:</i> Select the Server SSL profile you created above	
SNAT Pool ²	Auto Map (optional; see footnote ²)	
Default Pool	Select the pool you created above	

¹ You must select **Advanced** from the **Configuration** list for these options to appear

² If want to use SNAT, and you have a large deployment expecting more than 64,000 simultaneous connections, you must configure a SNAT Pool with an IP address for each 64,000 simultaneous connections you expect. See the BIG-IP documentation on configuring SNAT Pools.

This completes the configuration for this scenario.

Configuring the BIG-IP system for secure HTTP Proxy with BIG-IP APM

Use the following table to create the BIG-IP APM configuration. Create the objects in the order they appear in the table. For specific instructions on configuring individual objects, see the product documentation.

DNS and NTP	
See <i>Appendix C: Configuring DNS and NTP settings on the BIG-IP system on page 38</i> for instructions.	
AAA Server Health Monitor (Main tab > Local Traffic > Monitors)	
Name	Type a unique name
Type	LDAP
Interval	30 (recommended)
Timeout	91 (recommended)
User Name	Type the user name of a valid Active Directory user account. This should be in Canonical Name format. For example, CN=user1,CN=Users,DC=my,DC=domain,DC=com
Password	Type the associated password
Base	Specify your LDAP base tree. For example, CN=Users,DC=my,DC=domain,DC=com
Filter	Specify the filter. We type cn=user1, using the example above: user1 in OU group "Users" and domain "my.domain.com"
Security	Select a security option appropriate for your environment (either None, SSL, or TLS).
Alias Address	*All Addresses
Alias Service Port	389 (for None or TLS) or 686 (for SSL)
AAA Server (Main tab-->Access Policy-->AAA Servers)	
Name	Type a unique name
Type	Active Directory
Server Connection	Use Pool
Domain Controller Pool Name	Default is based on the name you entered above. You can optionally change it.
Domain Controllers	IP Address: Type the Ip address of a Domain Controller Hostname: Type the host name for the Domain Controller Click Add and repeat for each domain controller.
Server Pool Monitor	Select the monitor you created
Admin Name	If required for authentication, type the admin name
Admin Password	If required, type the associated password
NTLM Machine Account (Access Policy-->Access Profiles-->NTLM)	
Name	Type a unique name. Do not use spaces or special characters in this name.
Machine Account Name	The name of the account which will be joined to the Active Directory domain. This must be different than the account name specified in Kerberos SSO Configuration (such as bigip_machine_acct).
Domain FQDN	Type the FQDN for Active Directory (such as mydomain.com) you want to join
Admin User	Type the user name of a user with permissions to join a computer account to the AD domain.
Admin Password¹	Type the associated password. Click Join
NTLM Auth Configuration (Access Policy-->Access Profiles-->NTLM)	
Name	Type a unique name
Machine Account Name	Select the NTLM Machine Account you created above
Domain Controller FQDN List	Type the fully qualified name of your Active Directory domain controller and then click Add .
VDI Profiles (Access Policy-->Application Access > Remote Desktop > VDI Profiles)	
Name	Type a unique name
Parent Profile	/Common/vdi
MSRDP NTLM Configuration	Select the NTLM Auth Configuration you created
Connectivity Profile (Access Policy -->Secure Connectivity)	
Name	Type a unique name
Parent Profile	/Common/vdi

Access Profile (Access Policy->Access Profiles)	
Name	Type a unique name, such as rdg-apm-access .
Profile Type	All
Languages	Move the appropriate language(s) to the Accepted box.
Repeat this step to create a second Access policy (named rdg-remote-access-policy). The Profile Type <u>must be</u> RDG-RAP .	
Edit the Access Policy	
Edit the Access Profile you just created using the Visual Policy Editor. Continue now with Editing the Access Policy.	

Editing the Access Policy

After creating the objects in the table above, use the following procedure to edit the Access Policy on the BIG-IP APM using the Visual Policy Editor (VPE). This configuration requires two Access Policies, the first which changes based on how you want to secure the hosts, and second that is the same for all scenarios and references the first one.

Configuring the APM Access Policy for the Remote Access Policy

Use this procedure if you want to allow all hosts. This policy only restricts the target port to port 3389.

- On the Main tab, expand **Access Policy**, and click **Access Profiles**.
- Locate the Remote Access Policy Access Profile you created using the table above, and then, in the Access Policy column, click **Edit**. The VPE opens in a new window.
- Click the **+** symbol between **Start** and **Deny**. A box opens with options for different actions.
- Click the General Purpose tab, click the **Empty** option button, and then click **Add item**.
 - In the **Name** field, type a name like **Restrict Target Port**.
 - Click the Branch Rules tab, and then click the **Add Branch Rule** button.
 - In the **Name** field, type **Successful**.
 - In the **Expression** area, click the **change** link.
 - Click the Advanced tab, and then copy and paste the following code: `expr {[mcget {session.rdg.target.port}] == 3389}`
 - Click **Finished** and then click **Save**.
 - If you want to allow all hosts, continue with Step 7. If you want to restrict hosts based on an explicit list, continue with Step 5.
- Optional:* Use this step if you want to secure your hosts based on an explicit list, or based on explicit list and AD security group.
 - Click the **+** symbol between **Restrict Target Port** and **Deny**. A box opens with options for different actions.
 - Click the General Purpose tab, click the **Empty** option button, and then click **Add item**.
 - In the **Name** field, type a name like **Explicit Host Allowed**.
 - Click the Branch Rules tab, and then click the **Add Branch Rule** button.
 - In the **Name** field, type **Successful**.
 - In the **Expression** area, click the **change** link.
 - Click the Advanced tab, and use the following syntax to create the list of host names:
`expr {[mcget {session.rdg.target.host}] contains "hostname-1" || [mcget {session.rdg.target.host}] contains "hostname2"}`

For each additional host name, add this code snippet before the closing bracket:
`|| [mcget {session.rdg.target.host}] contains "additional-hostname"`
 - Click **Finished** and then click **Save**.
- Optional:* Use this step if you want to secure hosts using an Active Directory security group or using an Active Directory group and the explicit list you configured in step 5.
 - Click the **+** symbol between **Start** and **Deny**. A box opens with options for different actions.

- b. Click the **AD Query** option button, and then click **Add item**.
 - In the **Name** field, you can optionally type a unique name
 - From the **Server** list, select the AAA Server object you created using the table.
 - In the **SearchFilter** field, if your users can connect using FQDN only, type **(DNShostName={session.rdg.target.host})**
If your users can connect using FQDN or simple names, type:
(!(name={session.rdg.target.host})(DNShostName={session.rdg.target.host}))
 - Click the Branch Rules tab, delete any existing Branch rules, and then click the **Add Branch Rule** button.
 - In the **Name** field, type **Successful**.
 - In the **Expression** area, click the **change** link.
 - Click **Add Expression**.
 - From the **Condition** list, select **User is a Member of**.
 - In the **User is a Member of** field, type the AD security group that contains the allowed hosts, such as **CN=Example Group**.
 - Click **Add Expression**, click **Finished**, and then click **Save**.
7. Click the **Deny** link in the box to the right of the Successful path(s).
8. Click **Allow** and then click **Save**. Your VPE will look like one of the examples after Step 10.
9. Click the yellow **Apply Access Policy** link on the upper left. You have to apply an access policy before it takes effect.
10. Click the **Close** button on the upper right to close the VPE. Continue with

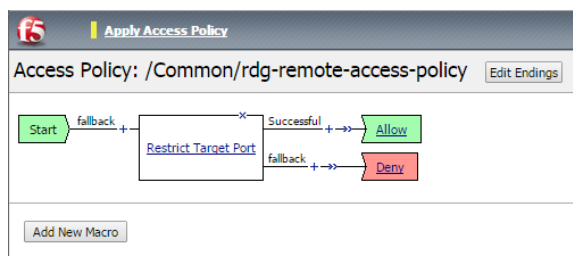


Figure 4: Access Policy if you allowed all hosts

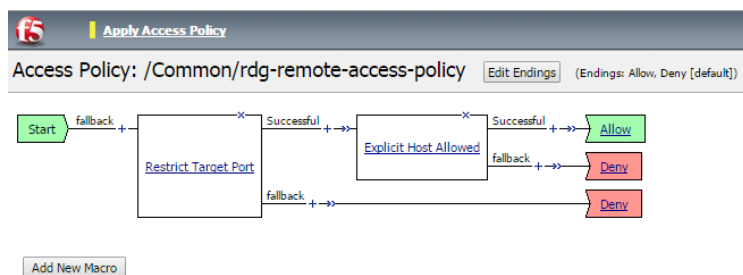


Figure 5: Access Policy if you chose to secure hosts using an explicit list only




Figure 6: Access Policy if you chose to secure hosts by AD security group only




Figure 7: Access Policy if you chose to secure hosts by explicit list and AD security group

Configuring the APM Access Policy for APM access

Next, use this procedure to edit the APM policy for APM access you created. You **must** edit this policy for the configuration to work.

- On the Main tab, expand **Access Policy**, and click **Access Profiles**.
- Locate the Access Profile for APM Access you created using the table above, and then, in the Access Policy column, click **Edit**.
- Click the **+** symbol between **Start** and **Deny**. A box opens with options for different actions.
- Click the End Point Security (Server-Side) tab, click the **Client Type** option button, and then click **Add item**.
 - In the **Name** field, you can type a new name.
 - Delete all of the default branches except **Microsoft RDP Client**, by clicking the **x** button on the right side of each row.
 - Click **Save**.
- Click the **+** symbol between **Client Type** and **Deny**. A box opens with options for different actions.
- Click the Authentication tab, click the **NTLM Auth Result** option button, and then click **Add Item**.
 - Configure the Properties as applicable for your configuration. In our example, we leave the settings at the defaults.
 - Click the **Save** button.
- Click the **+** symbol between **NTLM Auth Result** and **Deny**. A box opens with options for different actions.
- Click the Assignment tab, click the **RDG Policy Assign** option button, and then click **Add Item**.
 - In the **Name** field, you can type a new name.
 - Click the **Add/Delete** link next to **RDG Policy**, and then click the Remote Access Policy you created in the first procedure.
 - Click the **Save** button.
- Click the **Deny** link in the box to the right of **RDG Policy Assign**.
- Click **Allow** and then click **Save**. Your Access policy should look like the example below.
- Click the yellow **Apply Access Policy** link on the upper left. You have to apply an access policy before it takes effect.
- Click the **Close** button on the upper right to close the VPE.

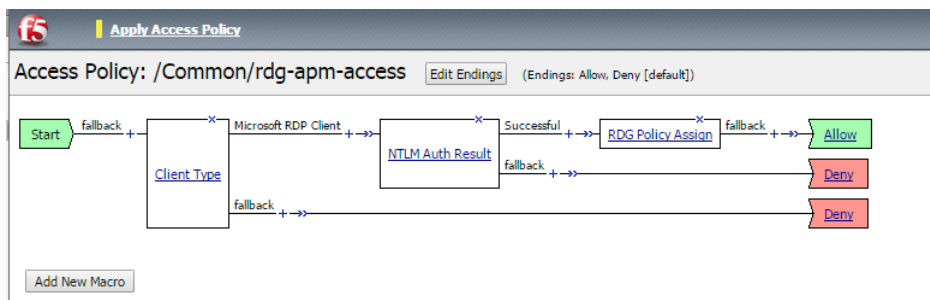


Figure 8: APM Access VPE example

Manually configuring the BIG-IP Advanced Firewall Module to secure your RDG deployment

This section describes how to manually configure BIG-IP AFM, F5's Network Firewall module, to secure your Remote Desktop Gateway deployment. BIG-IP AFM is particularly useful if you want to only allow access from specific clients or networks. Because this configuration can be complex, we recommend using the iApp template in version 11.6 and later to configure BIG-IP AFM.

Network Firewall settings

When configuring the BIG-IP Advanced Firewall Manager, you may want to configure your BIG-IP system to drop all traffic that you have not specifically allowed with firewall rules. This is known as **firewall mode**. By default, your BIG-IP system is set to default-accept, or **ADC mode**. Instructions for configuring your BIG-IP system, and the implications to consider, can be found on AskF5. For example, for BIG-IP v11.5: <http://support.f5.com/kb/en-us/products/big-ip-afm/manuals/product/network-firewall-policies-11-5-0/1.html>

If you have licensed IP Intelligence on the BIG-IP system, you can prohibit connections from sources with low reputation scores.

The following instructions cover a basic firewall configuration that is effective for the most common scenario of wanting to allow connections from a single trusted network. If you have complex requirements, such as the need to schedule different policies for different times of the day, or you want to create complicated rule or address lists, consult the BIG-IP AFM documentation. The basic steps for Policy and Rule creation apply to all scenarios.

To configure the BIG-IP AFM to allow connections from a single trusted network

1. Create a Network Firewall Policy:
 - a. From the Configuration utility, click **Security > Network Firewall > Policies**, and then click **Create**.
 - b. In the **Name** field, type a unique name for the policy, such as **RDG-Policy**.
 - c. Click **Finished**.
2. Create a rule to allow authorized hosts or networks to connect:
 - a. Click **Security > Network Firewall > Policies**.
 - b. Click the name of the policy you just created.
 - c. In the Rule section (below the General Properties section), click the **Add** button.
 - d. Leave the **Type** list set to Rule.
 - e. From the **Order** list, select **First**. The Order list only appears in version 11.5 and later. In 11.4.x, you must reorder the rules from the Policy General Properties page.
 - f. In the **Name** field, type a unique name, for instance **RDG-traffic-Allowed**.
 - g. Ensure the **State** list is set to **Enabled**.
 - h. From the **Protocol** list, select **TCP**. Leave the box to the right of TCP set to **6**.
 - i. In the **Source** section, from the **Address/Region** list, select **Specify**.
You are now able to list the trusted source addresses for your connection.
In the following example, we will configure a single subnet as trusted.
 - Select **Address**.
 - In the box, type the network address you want to allow, including netmask if more than a single host. Specify a network using CIDR notation, such as **10.0.0.0/24**.
 - Do not configure a source port.
 - Optional: If you want to limit inbound connections to a specific VLAN or Tunnel, from the **VLAN / Tunnel** list, select **Specify**, and then move the VLANs or tunnels that are allowed access to the Selected box.
 - Click **Add**.
 - Repeat these steps for additional hosts or networks. Use **Address List** or **Address Range** when appropriate.
 - j. In the **Destination** section, leave the **Address/Region** and **Port** set to **Any**. Because you will be applying your policy to a virtual server that listens only on a single desired address and port, do not specify that information here.
 - k. If necessary, from the **Action** list, select **Accept**.

- l. *Optional:* If you have configured a logging profile and want to log connections, from the **Logging** list, select **Enabled**. Typically, allowed connections do not need to be logged.
- m. Click **Finished**.

3. Creating a firewall rule to block all other traffic

The next task is to create a firewall rule to block all other traffic that you have not allowed. Although this is not a required step if your BIG-IP system is set to default deny (**Firewall mode**), it is required in default-accept (**ADC mode**), and is a good practice to always configure such a rule.

- a. Click **Security > Network Firewall > Policies**.
- b. Click the name of the policy you created in step 1.
- c. In the Rule section (below the General Properties section), click the **Add** button.
- d. Leave the **Type** list set to **Rule**.
- e. Leave the **Order** list, select **Last**.
- f. In the **Name** field, type a unique name, for example **RDG-traffic-Prohibited**.
- g. Ensure the **State** list is set to **Enabled**.
- h. From the **Protocol** list, select **TCP**. Leave the box to the right of TCP set to **6**.
- i. In the **Source** section, leave all the lists set to **Any**.
- j. From the **Action** list, select either **Drop** (to silently discard incoming connections) or **Reject** (to send a Destination Unreachable message to the sender).
- k. If you configured a logging profile as described in *Optional: Configuring the BIG-IP system to log network firewall events on page 32*, from the **Logging** list, select **Enabled**.
- l. Click **Finished**. You return to the Policy Properties page.
- m. On the Policy Properties page, in the Rules section, ensure the rule with the Action of Accept comes before the Drop or Reject rule you just created. If it does not, use the **Reorder** button and drag the rules into the correct order.

4. Apply Your Firewall Policy to your Virtual Server

- a. Click **Security > Network Firewall > Active Rules**.
- b. In the Rule section (below the General Properties section), click the **Add** button.
- c. From the **Context** list, select **Virtual Server**, and then select the virtual server you created for your RDG traffic.
- d. From the **Type** list, select **Policy**, and then select the firewall policy you created.
- e. From the **Policy Type** list, select **Enforced**.
- f. Click **Finished**.

Optional: Assigning an IP Intelligence Policy to your RDG virtual server

If you want to restrict access to your RDG virtual server based on the reputation of the remote sender, you can enable and assign an IP Intelligence policy. This requires an IP intelligence license; contact your F5 Sales representative for more information.

It is outside the scope of this document to provide instructions on configuring an IP Intelligence Policy. Full documentation on enabling and configuring the IP Intelligence feature can be found on AskF5. For example, the manual for BIG-IP AFM v11.5 is: <https://support.f5.com/kb/en-us/products/big-ip-afm/manuals/product/network-firewall-policies-implementations-11-5-0/5.html>

After you have enabled and configured an IP Intelligence policy, use the following steps to assign the policy to your RDG virtual server:

To assign the IP intelligence policy to the RDG virtual server

1. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.
2. Click the name of your RDG virtual server.

- From the **Security** menu, choose **Policies**.
- Next to **IP Intelligence**, select **Enabled**, then select the IP intelligence policy to apply to traffic on the virtual server.
- Click **Update**. The list screen and the updated item display. The IP Intelligence policy is applied to traffic on the virtual server.

Optional: Configuring the BIG-IP system to log network firewall events

If you are using BIG-IP AFM, you have the option of logging network firewall events to one or more remote syslog servers (recommended) or to log events locally. You can either use an iApp template to create the logging profile, or create the logging profile manually.

For specific information on logging on the BIG-IP system, see the appropriate guide for your version. For example, for 11.5.0:

- Remote High-Speed Logging:
https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-implementations-11-5-0/22.html
- Local logging:
https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/tmos-concepts-11-5-0/11.html

Creating the logging profile using the iApp template

Use this section to create the logging profile using the logging profile iApp template. If you have not already downloaded the iApp template, see <https://devcentral.f5.com/wiki/iApp.F5-Remote-Logging-iApp.ashx>.

To configure the logging profile iApp

- Log on to the BIG-IP system.
- On the Main tab, click **iApp > Application Services**.
- Click **Create**. The Template Selection page opens.
- In the **Name** box, type a name. In our example, we use **logging-iapp_**.
- From the **Template** list, select **f5.remote_logging.v<latest-version>**. The template opens.
- Use the following table for guidance on configuring the iApp template. Questions not mentioned in the table can be configured as applicable for your implementation.

Question	Your selection
Do you want to create a new pool of remote logging servers, or use an existing one?	Unless you have already created a pool on the BIG-IP system for your remote logging servers, select Create a new pool .
Which servers should be included in this pool?	Specify the IP addresses of your logging servers. Click Add to include more servers.
What port do the pool members use?	Specify the port used by your logging servers, typically 514 .
Do the pool members expect UDP or TCP connections?	TCP
Do you want to create a new monitor for this pool, or use an existing one?	Unless you have already created a health monitor for your pool of logging servers, select Use a simple ICMP (ping) monitor .
Do your log pool members require a specific log format?	If your logging servers require a specific format, select the appropriate format from the list.

- Click **Finished**.
- On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.
- Click the name of your RDG virtual server.
- From the **Security** menu, choose **Policies**.
- Next to **Log Profile**, select **Enabled**, then select the Logging profile you created.
- Click **Update**. The list screen and the updated item are displayed.

Note: The iApp template creates a log publisher and attaches it to the logging profile. If the publisher does not appear in the BIG-IP Configuration utility (GUI), you can verify the configuration by running the following command from the Traffic Management shell (tmsh): **list security log profile <your profile name>**.

Creating logging profile manually

If you do not want to use the iApp template to create a logging profile, use this section for guidance on configuring the logging profile manually. You must have access to the tmsh command line to use this method.

To manually configure a logging profile

1. Use the following guidance for configuring a health monitor and load balancing pool for the logging servers.

BIG-IP LTM Object	Non-default settings/Notes	
Health Monitor (Local Traffic -->Monitors)	Name	Type a unique name
	Type	ICMP
	Interval	30 (recommended)
	Timeout	91 (recommended)
Pool (Local Traffic -->Pools)	Name	Type a unique name
	Health Monitor	Select the appropriate monitor you created
	Slow Ramp Time	300
	Load Balancing Method	Choose a load balancing method. We recommend Least Connections (Member)
	Address	Type the IP Address of a server.
	Service Port	Type the appropriate port, such as UDP port 514 , the port on which logging typically occurs. Click Add , and then repeat Address and Port for all nodes

2. Log into the BIG-IP system using the command line. Enter the tmsh shell, by typing **tmsh** from the prompt.
3. Create a Remote High Speed Log (HSL) destination:
(tmsh)# create / sys log-config destination remote-high-speed-log [name] pool-name [specified pool] protocol [udp or tcp]

4. If you have a specific log format requirement, create a format-specific log destination, and forward that to the previously-created HSL destination:

(tmsh)# create / sys log-config destination [splunk|arcsight|remote-high-speed-log] [name] forward-to [HSL name]

5. Create a log publisher:

(tmsh)# create / sys log-config publisher [name] destinations add { [logdestination name] }

6. Create the logging profile to tie everything together.

If you chose to log allowed connections, include the green text (as in step 2 substep 1 in *To configure the BIG-IP AFM to allow connections from a single trusted network on page 30*).

If you set the rule to drop incoming connections, include the text in blue.

If you chose to log IP intelligence events, include the text in red to add the parameter that sets the log publisher.

(tmsh)# create / security log profile [name] network add { [name] { filter { log-acl-match-accept enabled log-acl-match-drop enabled log-acl-match-reject enabled } format { field-list { date time action drop reason protocol src ip src port dest ip dest port } type field-list } publisher [logpublisher name] } } ip-intelligence { log-publisher [logpublisher name] }

Assigning the logging profile to the virtual server

The final task is to assign the logging profile to the virtual server.

To assign the logging profile to the RDG virtual server

1. On the Main tab, click **Local Traffic > Virtual Servers**. The Virtual Server List screen opens.
2. Click the name of your RDG virtual server.
3. From the **Security** menu, choose **Policies**.
4. Next to **Log Profile**, select **Enabled**, then select the Logging profile you created.
5. Click **Update**. The list screen and the updated item are displayed.

Appendix A: Configuring WMI monitoring of the RDS servers for LTM only

If you find your RDS servers are under high performance load, you can dynamically load balance between them using F5's WMI monitor. This monitor checks the CPU, memory, and disk usage of the nodes and, in conjunction with Dynamic Ratio load balancing mode, sends the connection to the server most capable of processing it.

For an overview of the WMI performance monitor, see <http://support.f5.com/kb/en-us/solutions/public/6000/900/sol6914.html>.

Installing the F5 WMI handler

The first task is to copy the F5 WMI handler to the RDS server and configure IIS to use the F5 Data Gathering Agent. For instruction on installing the Data Gathering Agent, see:

http://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/itm_configuration_guide_10_0_0/itm_appendixb_monitor_considerations.html#1185026

Be sure to follow the procedures for the version of IIS you are using.

If you want to use the WMI monitor for the Session Host or Connection Broker servers, you must have IIS installed on those devices in order to install the handler.

Creating the WMI Monitor on the BIG-IP LTM

The next task is to create the WMI monitor on the applicable BIG-IP LTM systems. Use the following table:

BIG-IP LTM Object	Non-default settings/Notes	
Health Monitors (Main tab-->Local Traffic -->Monitors)	Name	Type a unique name
	Type	WMI
	Interval	30 (recommended)
	Timeout	91 (recommended)
	User Name	Type the appropriate user name
	Password	Type the associated password
	URL:	/scripts/F5Isapi.dll (for IIS 6, 7, and 7.5)

Create this monitor on all applicable BIG-IP LTM systems.

Apply the monitor on the BIG-IP LTM devices

Next, we apply the monitor to the applicable RDS nodes on the BIG-IP LTM system. This can be any or all of the BIG-IP LTM devices that are sending traffic to the RDS servers.

To apply the monitor to the nodes

1. On the Main tab, expand **Local Traffic** and then click **Nodes**.
2. From the list of nodes, click a node for the external IP address of your RDS server.
3. In the Configuration section, from the **Health Monitor** list, select **Node Specific**.
4. From the Available list, select the WMI monitor you created, and then click Add (<<).
5. Click **Update**.
6. Repeat for all appropriate nodes.
7. Repeat this procedure for all applicable BIG-IP LTM systems.

Modifying the pool(s) to use the Dynamic Ratio load balancing method

The next task is to modify the BIG-IP LTM pools to use the Dynamic Ratio load balancing method. Make this change for each pool that contains the RDS nodes to which you added the WMI monitor.

To modify the load balancing method on the pool

1. On the Main tab, expand **Local Traffic** and then click **Pools**.
2. Click the name of the appropriate Pool. The Pool Properties page opens.
3. On the Menu bar, click **Members**.
4. From the **Load Balancing Method** list, select **Dynamic Ratio (Node)**.
5. Click the **Update** button.
6. Repeat this procedure for all applicable pools on this BIG-IP LTM.
7. Repeat this procedure on all applicable BIG-IP LTM systems.

Appendix B: Using X-Forwarded-For to log the client IP address in IIS 7.0, 7.5, and 8 (optional)

When you configure BIG-IP LTM to use SNAT, the BIG-IP system replaces the source IP address of an incoming connection with its local self IP address (in the case of SNAT Auto Map), or an address you have configured in a SNAT pool. As a result, Microsoft IIS logs each connection with its assigned SNAT address, rather than the address of the client.

Beginning with IIS 7, Microsoft provides an optional Advanced Logging Feature for IIS that allows you to define custom log definitions that can capture additional information such as the client IP address included in the X-Forwarded-For header.

This section is only applicable if you are deploying Remote Desktop Gateway or Remote Desktop Web Access.

Modifying the HTTP profile to enable X-Forwarded-For

The first task is to modify the HTTP profile created by the template to enable the X-Forwarded-For header.

To modify the HTTP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**.
2. From the HTTP profile list, select the HTTP profile you created.
3. In the Settings section, on the **Insert X-Forwarded-For** row, click the **Custom** box. From the list, select **Enabled**.
4. Click the **Update** button.

Deploying the Custom Logging role service

The first task is to deploy the Custom Logging role service. If you do not deploy this role service, you may receive a “Feature not supported” error when trying to edit the log definition in the next section. If you receive this error, ensure that you are editing the log definition at the server level in IIS Manager.

The configuration is slightly different depending on which version of IIS you are running. Use the procedure applicable to your version of IIS.

To deploy the Custom Logging role service for IIS 7.0 and 7.5 (Windows Server 2008)

1. From your Windows Server 2008 or Windows Server 2008 R2 device, open Server Manager.
2. In the Navigation pane, expand **Roles**.
3. Right-click **Web Server**, and then click **Add Role Services**.
4. Under Health and Diagnostics, check the box for **Custom Logging**, and then click **Next**.
5. On the Confirmation page, click **Install**.
6. After the service has successfully installed, click the **Close** button.

To deploy the Custom Logging role service for IIS 8.0 (Windows Server 2012)

1. From your Windows Server 2012 device, open Server Manager.
2. Click **Manage** and then **Add Roles and Features**.
3. Select Role-based or feature-based installation.
4. On the Roles screen, expand **Web Server (IIS)** and **Health and Diagnostics** and then check the box for **Custom Logging**.
5. Click **Next** and then on the Features screen, click **Next** again.
6. Click **Install**.
7. After the service has successfully installed, click the **Close** button.

Adding the X-Forwarded-For log field to IIS

Before beginning the following procedure, you must have installed IIS Advanced Logging. For installation instructions, see http://www.iis.net/community/files/media/advancedlogging_readme.htm

If you are using IIS version 6, F5 has a downloadable ISAPI filter that performs a similar function to the Advanced Logging Feature discussed here. For information on that solution, see the DevCentral post at http://devcentral.f5.com/weblogs/Joel/archive/2009/08/19/x_forwarded_for_log_filter_for_windows_servers.aspx

The following procedure is the same for IIS versions 7.0, 7.5, and 8.0.

To add the X-Forwarded-For log field to IIS

1. From your Windows Server device, open the Internet Information Services (IIS) Manager.
2. From the Connections navigation pane, click the appropriate server on which you are configuring Advanced Logging. The Home page appears in the main panel.
3. From the Home page, under IIS, double-click **Advanced Logging**.
4. From the Actions pane on the right, click **Edit Logging Fields**.
5. From the Edit Logging Fields dialog box, click the **Add Field** button, and then complete the following:
 - a. In the **Field ID** box, type **X-Forwarded-For**.
 - b. From the **Category** list, select **Default**.
 - c. From the **Source Type** list, select **Request Header**.
 - d. In the **Source Name** box, type **X-Forwarded-For**.
 - e. Click the **OK** button.
6. Click a Log Definition to select it. By default, there is only one: %COMPUTERNAME%-Server. The log definition you select must have a status of Enabled.
7. From the Actions pane on the right, click **Edit Log Definition**.
8. Click **Select Fields**, and then check the box for the X-Forwarded-For logging field.
9. Click the **OK** button.
10. From the Actions pane, click **Apply**.
11. Click **Return To Advanced Logging**.
12. In the Actions pane, click **Enable Advanced Logging**.

Now, when you look at the Advanced Logging logs, the client IP address is included.

Appendix C: Configuring DNS and NTP settings on the BIG-IP system


If you are using BIG-IP APM, before beginning the iApp, you must configure DNS and NTP settings on the BIG-IP system.


Configuring DNS and NTP settings

If you are configuring the iApp to use BIG-IP APM, you must configure DNS and NTP settings on the BIG-IP system before beginning the iApp.

Configuring the DNS settings

In this section, you configure the DNS settings on the BIG-IP system to point to a DNS server that can resolve your Active Directory server or servers. In many cases, this IP address will be that of your Active Directory servers themselves.

 **Note:** *DNS lookups go out over one of the interfaces configured on the BIG-IP system, not the management interface. The management interface has its own, separate DNS settings.*

 **Important** *The BIG-IP system must have a self IP address in the same local subnet and VLAN as the DNS server, or a route to the DNS server if located on a different subnet. The route configuration is found on the Main tab by expanding **Network** and then clicking **Routes**. For specific instructions on configuring a route on the BIG-IP system, see the online help or the product documentation.*

To configure DNS settings

1. On the Main tab, expand **System**, and then click **Configuration**.
2. On the Menu bar, from the **Device** menu, click **DNS**.
3. In the **DNS Lookup Server List** row, complete the following:
 - a. In the **Address** box, type the IP address of a DNS server that can resolve the Active Directory server.
 - b. Click the **Add** button.
4. Click **Update**.

Configuring the NTP settings

The next task is to configure the NTP settings on the BIG-IP system for authentication to work properly.

To configure NTP settings

1. On the Main tab, expand **System**, and then click **Configuration**.
2. On the Menu bar, from the **Device** menu, click **NTP**.
3. In the **Address** box, type the fully-qualified domain name (or the IP address) of the time server that you want to add to the Address List.
4. Click the **Add** button.
5. Click **Update**.

To verify the NTP setting configuration, you can use the **ntpq** utility. From the command line, run **ntpq -np**.

See <http://support.f5.com/kb/en-us/solutions/public/10000/200/sol10240.html> for more information on this command.

Document Revision History

Version	Description	Date
1.0	New deployment guide for the fully supported f5.microsoft_rds_remote_access.v1.0.0 iApp.	04-09-2015
1.1	<ul style="list-style-type: none"> - Added a new step-by-step section for creating an NTLM Machine Account object: <i>Creating an NTLM Machine Account on page 21.</i> - Added a new troubleshooting section with the entry: <i>You may experience deployment errors when the NTLM Machine Account name contains spaces or special characters on page 22.</i> Added the same guidance to the manual configuration. 	05-01-2015
1.2	Added a new entry to section <i>Troubleshooting on page 22</i> dealing with an error message that may display if connecting through the BIG-IP system using the default Windows 7 Remote Desktop client.	05-28-2015

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

