



Deploying the BIG-IP LTM System with Citrix XenDesktop

Important: This guide has been archived. While the content in this guide is still valid for the products and versions listed in the document, it is no longer being updated and may refer to F5 or third party products or versions that have reached end-of-life or end-of-support. For a list of current guides, see <https://f5.com/solutions/deployment-guides>.

Table of Contents

Deploying the BIG-IP LTM with Citrix XenDesktop

Prerequisites and configuration notes	1
Product versions and revision history	2
Configuration example	3
Configuring the BIG-IP LTM system	4
Creating the health monitors	4
Creating the Citrix pools	7
Creating Profiles	8
Creating the Citrix Web Interface virtual server	13
Creating the Desktop Delivery Controller virtual server	15
Modifying the Citrix Web Interface configuration	16
Appendix A: Creating a Server SSL profile	17
Modifying the virtual server to use the Server SSL profile	17

Deploying the BIG-IP LTM with Citrix XenDesktop

Welcome to the F5 BIG-IP deployment guide for Citrix® XenDesktop®. This guide contains step-by-step procedures for configuring the BIG-IP Local Traffic Manager (LTM) for directing traffic, ensuring application availability, improving performance and providing a flexible layer of security for Citrix XenDesktop version 5.0.

Citrix XenDesktop lets you create virtualized desktops quickly and easily, then make them available to users on demand through any device.

The BIG-IP LTM provides mission critical availability, enhanced security, simple scalability and high operational resiliency to the Citrix XenDesktop deployment.

In a Citrix XenDesktop environment, the BIG-IP LTM provides intelligent traffic management and high-availability by monitoring and managing connections to the Citrix Web Interface. In addition, the built-in performance optimization capabilities of the LTM provide faster operations to facilitate a better end-user experience. The LTM also keeps persistence records for certain connections to always be directed to the same server for a specified period of time, to ensure that the workflow in the XenDesktop environment is fully preserved.

For more information on the F5 BIG-IP LTM, see www.f5.com/products/big-ip/product-modules/local-traffic-manager.html

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

Prerequisites and configuration notes

All of the procedures in this Deployment Guide are performed on the BIG-IP system. The following are prerequisites for this solution:

- ◆ For this deployment guide, the Citrix XenDesktop installation must be running version 5.0.
- ◆ For this deployment guide, the BIG-IP LTM system should be running version 10.2 or later. If you are using a previous version of the BIG-IP LTM system see the [Deployment Guide](#) index.

Important: If you are using version 10.2.1, you **must be running version 10.2.1 Hotfix 1** or later for the configuration in this guide.

- ◆ If you are using the BIG-IP system to offload SSL, we assume you have already obtained an SSL certificate and key, but it is not yet installed on the BIG-IP LTM system. For more information, see *Offloading SSL from the XenDesktop servers*, on page 11.
- ◆ Citrix Session configuration must be set to Direct mode (see Figure 1, on page 2). For specific information on configuring the Citrix Session mode, see the Citrix documentation.




Figure 1 Citrix Session configuration

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP LTM	10.2, 10.2.1 HF1, 10.2.2
Citrix XenDesktop	5.0

Revision history:

Document Version	Description
1.0	New deployment guide
1.1	<ul style="list-style-type: none"> - Removed support for v10.2.1, added support for 10.2.1 HF-1 and 10.2.2. - Added note that the Citrix Session configuration must be set to Direct mode. - Added additional information on tuning the TCP WAN optimized profiles for users with low bandwidth or high latency connections.

Configuration example

This configuration example describes the typical configuration of the BIG-IP LTM system to monitor and manage the critical component of a Citrix XenDesktop environment: the Web Interface servers (WI) and Desktop Delivery Controllers (DDC).

In this implementation, traffic to the Citrix WI and DDC servers are managed by the BIG-IP LTM system. When necessary, the BIG-IP LTM ensures that each client connects to the same member of the farm across multiple sessions using persistence. The BIG-IP LTM system is also setup to monitor the Citrix WI and DDC servers to ensure availability, authentication and to automatically mark down servers that are not operating properly.

This guide also addresses SSL offload - the ability of the BIG-IP system to terminate SSL sessions in order to offload this CPU-intensive processing from the XenDesktop WI servers. We strongly recommend SSL offload for XenDesktop deployments, which is available with a simple addition of the Client SSL profile to the WI virtual server, referred to in this guide. For organizations that would prefer not to offload, we describe how to perform re-encryption in *Appendix A: Creating a Server SSL profile*, on page 17.

F5 Application Delivery Control for XenDesktop provides high availability in conjunction with advanced monitoring that looks at XenDesktop farm availability on DCC servers and authentication through WI servers provides the ultimate flexibility to deliver a resilient and available environment.




Figure 2 Logical configuration example

Configuring the BIG-IP LTM system

Use the following procedures to configure the BIG-IP LTM system for Citrix XenDesktop.

Creating the health monitors

To ensure traffic is directed only to those servers that are responding to requests, it is important to configure health monitors on the BIG-IP LTM to verify the availability of the servers being load balanced.

For Citrix XenDesktop, we create two advanced monitors. The first monitor is for the Web Interface servers and attempts to login to the servers by using the user name and account of a test user. We recommend you create a test user that reflects users in your environment for this purpose. If a particular server fails authentication, traffic is diverted from those servers until those devices are fixed. If all authentication is down, users will not be able to connect. We recommend setting up a Fallback Host for these situations. Please see F5 product documentation on setting up Fallback Hosts in your pools

The second monitor is for the Desktop Delivery Controller servers. This monitor determines the availability of the Desktop Farm to which users connect. If the farm is not available on the controller, it is taken out of service.

◆ Note

The first monitor uses a user account (user name and password) that can retrieve applications from the XenDesktop server. Use an existing account for which you know the password, or create an account specifically for use with this monitor.

For the second monitor, you need to know the name of your farm. This information can be found in your Citrix XenDesktop Management Console.

Both health monitors are created using a script, available on DevCentral <http://devcentral.f5.com/wiki/default.aspx/tmsh/CitrixXenDesktopMonitor.html>.

Download the script to a location accessible by the BIG-IP device.

Optionally, you can cut and paste the script directly into the TMSH editor on the BIG-IP device. However, cutting and pasting is error-prone and therefore we provide instructions here on how to copy the file to the BIG-IP device using secure-copy (SCP).

To create the Web Interface Monitor and the Desktop Delivery Controller Monitor using the script, you must first copy the script into the BIG-IP device. The following procedures show you how to copy the file both on a Windows platform using WinSCP, and on Linux, UNIX or MacOS system using SCP.

To import the script on a Windows platform using WinSCP

1. Download the script found on the following link to a computer that has access to the BIG-IP device:

<http://devcentral.f5.com/wiki/default.aspx/tmsh/CitrixXenDesktopMonitor.html>

1. Open a Windows compatible SCP client. We recommend WinSCP. It is available as a free download from <http://winscp.net/>. The login box opens.
2. In the **Host name** box, type the host name or IP address of your BIG-IP system.
3. In the **User name** and **Password** boxes, type the appropriate administrator log on information.
4. Click **Login**. The WinSCP client opens.
5. In the left pane, navigate to the location where you saved the script in step 1.
6. In the right pane, navigate to **/shared/tmp/** (from the right pane drop-down list, select **root**, and then double-click **shared**, and then double-click **tmp**) (see Figure 3).
7. In the left pane, select the script and drag it to the right pane.
8. You can now safely close WinSCP.

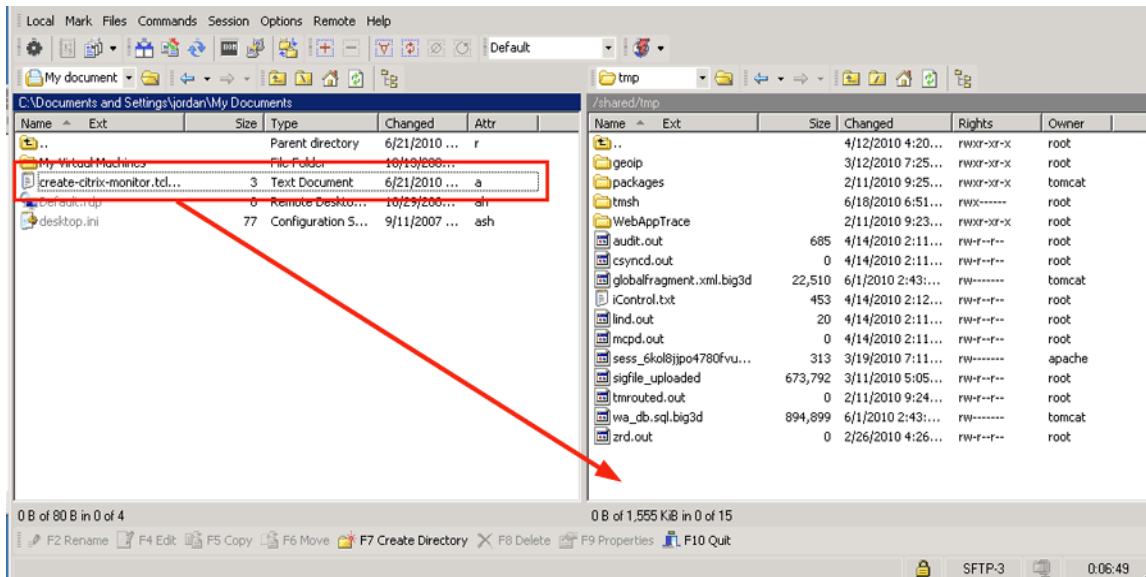


Figure 3 WinSCP client showing the monitor

To import the script using Linux/Unix/MacOS systems

1. Download the script:
<http://devcentral.f5.com/wiki/default.aspx/tmsh/CitrixXenDesktopMonitor.html>.
2. Open a terminal session.
3. Use your built in secure copy program from the command line to copy the file. Use the following syntax:

```
scp <source file> <username>@<hostname>:<Destination Directory and filename>
```

In our example, the command is:

```
scp create-citrix-monitor.tcl root@bigip.f5.com:/shared/tmp/create-citrix-monitor.sh
```

The next task is to import the script you just copied to create the monitor. The following tasks are performed in the BIG-IP Advanced Shell (see the BIG-IP manual on how to configure users for Advanced shell access).

To import and run the monitor creation script

1. On the BIG-IP system, start a console session.
2. Type a user name and password, and then press Enter.
3. Change to the directory containing the creation script. In our example, we type:

```
cd /shared/tmp/
```

If you copied the script to a different destination, Use the appropriate directory.

4. Change the permissions on the script to allow for execute permission using the following command:
5. Run the script, which imports the monitor creation script to your local configuration, using the following command:

```
chmod 755 create-citrix-monitor.sh
```

```
./create-citrix-monitor.sh
```

With a successful run you will see the following output:

```
Reading configuration from -.
Loading the configuration ...
TMSH script loaded successfully.
```

You have now successfully imported the script. The next step is to run the script and provide the parameters to create the Citrix XenDesktop monitor for your environment.

To run the monitor script

1. At the system prompt, type **tmsh** and then press Enter. This opens the Traffic Management shell.
2. Enter CLI Script mode by typing **cli script**. The prompt changes to `root@bigip-hostname (Active) (tmsh.cli.script) #`
3. From the command prompt, type the following command:

```
run create-citrix-xendesk-monitor.tcl
```

◆ Important

The monitor script name above (`create-citrix-xendesk-monitor.tcl`) differs from the file name you initially uploaded to BIG-IP. Follow the instructions in step 3 exactly.

-
- The script starts, you are prompted for four arguments. You are automatically switched to interactive mode
4. At the **What is the User Name** prompt, type the user name of the XenDesktop user.
 5. At the **What is the Password** prompt, type the associated password.
 6. At the **What is the Farm name** prompt, type the name of the farm of your XenDesktop farm you would like to check is available. In our example, we use **HOME**.
 7. At **What is the domain name** prompt, type the Windows domain used for authentication of users. In our example, we use **corpdomain**. Do not use the fully-qualified-domain-name from DNS here; this is referring to Windows Domain only.

The script creates the monitor. You can view the newly created monitor from the web-based Configuration utility from the Main Tab, by expanding **Local Traffic** and then clicking **Monitors**. The name of the monitors starts with the farm name you configured in step 6.

In our example, the two monitors that are created are:
Home-CitrixDDCFarm and *Home-CitrixWICredentials*.

Creating the Citrix pools

In this section, we create the Citrix Web Interface and Desktop Delivery Controller pools.

◆ Note

If your Desktop Controllers are on the same machine as your Web Interface servers, you can apply both monitors to one pool. In this case, there is no need to repeat the procedure to create a pool for the Desktop Controllers. Simply select both monitors in Step 4.

To create the Citrix pools

1. On the Main tab, expand **Local Traffic**, and then click **Pools**. The Pool screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New Pool screen opens.
3. In the **Name** box, enter a name for your pool. In our example, we use **web_pool**.
4. In the Health Monitors section, select the name of the monitor created by the script for the Web Interface servers in *Creating the health monitors*, on page 4, and then click the Add (<<) button. In our example, we select **Home-CitrixWICredentials** and click the Add (<<) button.

5. From the Load Balancing Method list, choose your preferred load balancing method (different load balancing methods may yield optimal results for a particular network). In our example, we select **Least Connections (member)**.
6. In the New Members section, make sure the **New Address** option button is selected.
7. In the **Address** box, type the address of the first server. In our example, we type **192.168.10.1**.
8. In the **Service Port** box, type the service number you want to use for this device, or specify a service by choosing a service name from the list. In our example, we type **80**.
9. Click the **Add** button to add the member to the list.
10. Repeat steps 8-10 for each server you want to add to the pool. In our example, we repeat these steps for the remaining server in this pool, **192.168.10.2**.
11. Click the **Repeat** button.
12. Repeat this entire procedure to create a pool for the Desktop Delivery Controllers, with the following changes:
 - a) In Step 3, give this pool a unique name, such as **desktopcontroller_pool**.
 - b) In Step 4, select the name of the monitor created by the script for the Desktop Controllers in *Creating the health monitors*, on page 4, and then click the Add (<<) button. In our example, we select **Home-CitrixDDCFarm** and click the Add (<<) button.
 - c) In Step 7, type the appropriate IP address for the Desktop Controllers.
 - d) In Step 11, click the **Finished** button.

Creating Profiles

The BIG-IP system uses profiles for greater control over managing network traffic while making network traffic management easy and efficient. A profile is an object that contains user-configurable settings, with default values, for controlling the behavior of a particular type of network traffic, such as HTTP connections.

Although it is possible to use the default profiles, we strongly recommend that you create new profiles based on the default parent profiles. Creating new profiles allows you to easily modify the profile settings specific to your deployment, and ensures that you do not accidentally overwrite the default profile. We also recommend however, that you use the default settings in the cookie persistence profile for this configuration (Insert method, Session based).

Creating an HTTP profile

The first new profile we create is an HTTP profile. The HTTP profile contains numerous configuration options for how the BIG-IP LTM system handles HTTP traffic. Our testing with Citrix XenDesktop shows that while we see a great deal of benefit from compression, caching only produces a minimal improvement. Therefore, we recommend using the **http-wan-optimized-compression** parent profile. This profile uses specific compression (among other) settings to optimize traffic over the WAN.

Citrix XenDesktop must have access to the IP address of the connecting clients in order to be fully functional. Some of the BIG-IP LTM features used in this Deployment Guide obscure this information. To overcome this, we use the following HTTP profile to insert an **X-Forwarded-For** header into the HTTP header. This supplies the IP address of the client so it is available to Citrix XenDesktop.

To create a new HTTP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. Click the **Create** button.
3. In the **Name** box, type a name for this profile. In our example, we type **citrix-http-opt**.
4. From the **Parent Profile** list, select **http-wan-optimized-compression**.
5. From the **Redirect Rewrite** row, click the **Custom** box, and then select **All** from the list.
6. From the **Insert XForward For** row, click the **Custom** box, and then select **Enabled** from the list.
7. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
8. Click the **Finished** button.

Creating the TCP profiles

The next profiles we create are the TCP profiles. If most of the Citrix XenDesktop users are accessing the devices via a Local Area Network, we recommend using the **tcp-lan-optimized** (for server-side TCP connections) parent profile. If the majority of the users are accessing the system from remote or home offices, we recommend using an additional TCP profile, called **tcp-wan-optimized** (for client side TCP connections). The use of these optimized profiles is optional, you can alternatively use the base TCP parent profile if appropriate for your configuration.

With regard to the LTM TCP profiles and XenApp, Citrix maintains keepalives using its own clients. This keepalive is configurable on a per client basis (see Citrix documentation instructions on adjusting this

timeout). As an alternate approach, if premature session termination is a concern, we recommend setting the **Idle Timeout** value to a longer time period to prevent idle desktop sessions from being terminated prematurely.

◆ **Important**

*Setting TCP timeout to **Indefinite** may lead to session exhaustion and should be used with care.*

Optional: Certain WAN conditions such as users connecting over low bandwidth or high latency can be optimized further by using different options for the TCP WAN profile. We recommend that you review the following solutions for environments where users are connecting from more challenging WAN conditions. Significant improvements are possible. Specifically, we recommend setting **Nagle's Algorithm** to **Disabled** and setting **Congestion Control** to **Scalable**.

<http://support.f5.com/kb/en-us/solutions/public/7000/400/sol7402.html>

<http://support.f5.com/kb/en-us/solutions/public/7000/400/sol7405.html>

Creating the LAN optimized TCP profile

First we configure the LAN optimized profile. If you do not want to use this optimized profile, you can choose the default TCP parent profile.

To create a new TCP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **TCP**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this profile. In our example, we type **citrix-tcp-lan**.
5. From the **Parent Profile** list, select **tcp-lan-optimized**.
6. *Optional:* In the **Idle Timeout** row, click the **Custom** box, set a time that matches that your desired timeout maximum. We recommend a setting of **600** to **900** seconds (use Indefinite with care, see note above).
7. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
8. Click the **Finished** button.

Creating the WAN optimized TCP profile

Now we configure the WAN optimized profile. Remember, if most of the users are accessing the system over the LAN or other low latency links, you do not need to create this profile. Again, we set the Idle Timeout value to Indefinite to prevent idle desktop sessions from being terminated prematurely.

To create a new TCP WAN optimized profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, from the **Protocol** menu, click **tcp**.
3. Click the **Create** button.
4. In the **Name** box, type a name for this profile. In our example, we type **citrix-tcp-wan**.
5. From the **Parent Profile** list, select **tcp-wan-optimized**.
6. *Optional:* In the **Idle Timeout** row, click the **Custom** box, set a time that matches that your desired timeout maximum. We recommend a setting of 600 to 900 seconds (use Indefinite with care, see note above).
7. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
8. Click the **Finished** button.

Creating the persistence profile

The next profile we create is a Persistence profile. We recommend using persistence for Citrix devices, although the type of persistence depends on your configuration. In our example, use cookie persistence (HTTP cookie insert).

To create a new cookie persistence profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. On the Menu bar, click **Persistence**. The Persistence Profiles screen opens.
3. Click the **Create** button.
4. In the **Name** box, type a name for this profile. In our example, we type **citrix-cookie**.
5. From the **Persistence Type** list, select **Cookie**. The configuration options for cookie persistence appear.
6. Modify any of the settings as applicable for your network. In our example, we leave the settings at their default levels.
7. Click the **Finished** button.

Offloading SSL from the XenDesktop servers

We strongly recommend terminating SSL connections on the BIG-IP LTM and offloading SSL from Citrix XenDesktop servers, this requires a Client SSL profile.

If for some reason you have requirements that traffic is encrypted all the way to the XenDesktop servers, then in order to preserve persistence and benefits from all F5 functionality, we recommend you terminate SSL on the BIG-IP and then re-encrypt the traffic to the Citrix server. This requires an both a Client SSL profile and a Server SSL profile. This is not a typical configuration.

In the following procedure, we describe creating a client SSL profile and offloading SSL from the XenDesktop servers. If your configuration scenario requires you to re-encrypt the traffic, see *Appendix A: Creating a Server SSL profile*, on page 17 for instructions.

Before you can enable the BIG-IP LTM system to act as an SSL proxy, you must install a SSL certificate on the virtual server that you wish to use for XenDesktop Web Interface connections on the BIG-IP LTM device. For this Deployment Guide, we assume that you already have obtained an SSL certificate, but it is not yet installed on the BIG-IP LTM system. For information on generating certificates, or using the BIG-IP LTM to generate a request for a new certificate and key from a certificate authority, see the **Managing SSL Traffic** chapter in the *Configuration Guide for Local Traffic Management*.

Once you have obtained a certificate, you can import this certificate into the BIG-IP LTM system using the Configuration utility. By importing a certificate or archive into the Configuration utility, you ease the task of managing that certificate or archive. You can use the Import SSL Certificates and Keys screen only when the certificate you are importing is in Privacy Enhanced Mail (PEM) format.

To import a key or certificate

1. On the Main tab, expand **Local Traffic**.
2. Click **SSL Certificates**. The list of existing certificates displays.
3. In the upper right corner of the screen, click **Import**.
4. From the **Import Type** list, select the type of import (Certificate or Key).
5. In the **Certificate** (or **Key**) **Name** box, type a unique name for the certificate or key.
6. In the **Certificate** (or **Key**) **Source** box, choose to either upload the file or paste the text.
7. Click **Import**.
8. If you imported the certificate, repeat this procedure for the key.

To create a new Client SSL profile

1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then, on the Menu bar, from the **SSL** menu, click **Client**.
2. Click the **Create** button. The New Client SSL Profile screen opens.
3. In the **Name** box, type a name. We type **citrix-clientssl**.

-
4. In the Configuration section, click a check in the **Certificate** and **Key Custom** boxes.
 5. From the **Certificate** list, select the name of the Certificate you imported.
 6. From the **Key** list, select the key you imported.
 7. Click the **Finished** button.


Creating the Citrix Web Interface virtual server

A virtual server with its virtual IP address is the visible, routable entity through which the servers in a load balancing pool are made available to the client (the IP address to give clients or add to DNS).

The next step in the configuration is to configure virtual servers that reference the pools and profiles created in the preceding sections.

To create the Citrix Web Interface virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**. The Virtual Servers screen opens.
2. Click the **Create** button. The New Virtual Server screen opens.
3. In the **Name** box, type a name for this virtual server. In our example, we type **citrix_ps.example.com**.
4. In the **Destination** section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **172.27.92.118**.
6. In the **Service Port** box, type **443**, or select **HTTPS** from the list.



General Properties	
Name	citrix_ps.example.com
Destination	Type: <input checked="" type="radio"/> Host <input type="radio"/> Network
	Address: 172.27.92.118
Service Port	443 HTTPS
State	Enabled

Figure 4 Creating the Citrix virtual server

7. From the Configuration list, select **Advanced**. The Advanced configuration options appear.
8. Leave the **Type** list at the default setting: **Standard**.

9. From the **Protocol Profile (Client)** list select the name of the profile you created in the *Creating the WAN optimized TCP profile* section. If you did not create a WAN optimized profile, select the LAN optimized profile as in the following Step. In our example, we select **citrix-tcp-wan**.
10. From the **Protocol Profile (Server)** list, select the name of the profile you created in the *Creating the LAN optimized TCP profile* section. In our example, we select **citrix-tcp-lan**.
11. From the HTTP Profile list, select the name of the profile you created in the *Creating an HTTP profile* section. In our example, we select **citrix-http-opt**.
12. From the **SSL Profile (Client)** list, select the profile you created in *Offloading SSL from the XenDesktop servers*, on page 11. In our example, we select **citrix-clientssl**.
13. From the **SNAT Pool** list, select **Automap**.

Configuration: Advanced	
Type	Standard
Protocol	TCP
Protocol Profile (Client)	citrix-tcp-wan
Protocol Profile (Server)	citrix-tcp-lan
OneConnect Profile	None
NTLM Conn Pool	None
HTTP Profile	citrix-http-opt
FTP Profile	None
SSL Profile (Client)	citrix-clientssl
SSL Profile (Server)	None
Authentication Profiles	<div style="display: flex; justify-content: space-between;"> <div>Enabled</div> <div>Available</div> </div> <div style="display: flex; justify-content: space-between; margin-top: 5px;"> <div><<</div> <div>ssl_cc_idap</div> </div>

Figure 5 Selecting the Citrix profiles for the virtual server

14. In the Resources section, from the **Default Pool** list, select the pool you created in *Creating the Citrix pools*, on page 7. In our example, we select **web_pool**.
15. From the **Default Persistence Profile** list, select the persistence profile you created in the *Creating the persistence profile* section. In our example, we select **citrix-cookie**.
16. Click the **Finished** button (see Figure 6).




Figure 6 Adding the Pool and Persistence profile to the virtual server

Creating the Desktop Delivery Controller virtual server

In this procedure, we create the virtual server for the Desktop Delivery Controllers.

To create the Desktop Delivery Controller virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. Click the **Create** button.
3. In the **Name** box, type a name for this virtual server. In our example, we type **citrix_DDC_vs**.
4. In the **Destination** section, select the **Host** option button.
5. In the **Address** box, type the IP address of this virtual server. In our example, we use **172.27.92.125**.
6. In the **Service Port** box, type **80**, or select **HTTP** from the list.
7. From the **SNAT Pool** list, select **Automap**.
8. In the Resources section, from the **Default Pool** list, select the pool you created for the Desktop Delivery Controllers in *Creating the Citrix pools*, on page 7. In our example, we select **desktopcontroller_pool**.
9. Click the **Finished** button.

This completes the BIG-IP LTM configuration.

Modifying the Citrix Web Interface configuration

You must modify the Web Interface server configuration so the Web Interface devices send traffic to the BIG-IP Desktop Delivery Controller virtual server and not directly to the Desktop Delivery Controllers. You must also make sure *Use the server list for load balancing* is unchecked, as shown below.

To modify the Web Interface servers to point at the Desktop Delivery Controller virtual server

1. From a Web Interface server, open the Access Management Console.
2. In the Navigation pane, expand **Citrix Resources, Configuration Tools, Web Interface** and then your site name.
3. From the middle column, select **Manage server farms**.
4. From the list, select the appropriate farm, and then click **Edit**.
5. In the **Server** box, select each entry and then click the **Remove** button.
6. Click the **Add** button.
7. Type the IP address of the Desktop Delivery Controller virtual server (from step 5 of *Creating the Desktop Delivery Controller virtual server*, on page 15). In our example, we type **172.27.92.125**.
8. Clear the check from the **Use the server list for load balancing** box.
9. Click the **OK** button. Repeat this procedure for any/all additional Web Interface servers.

This concludes the deployment guide configuration. If you are performing SSL re-encryption, continue with Appendix A.

Appendix A: Creating a Server SSL profile

In the recommended configuration with the BIG-IP LTM and XenDesktop, the LTM offloads encryption from the XenDesktop deployment. If for some reason you have requirements that traffic is encrypted all the way to the XenDesktop servers, then in order to preserve persistence and benefits from all F5 functionality, we recommend you terminate SSL on the BIG-IP and then re-encrypt the traffic to the Citrix server. This requires an both a Client SSL profile and a Server SSL profile. This is not a typical configuration.

Re-encrypting the traffic requires a Server SSL profile, and a modification to the virtual server to use this new profile.

To create a new Server SSL profile

1. On the Main tab, expand **Local Traffic**, click **Profiles**, and then, on the Menu bar, from the **SSL** menu, click **Server**.
2. Click the **Create** button. The New Client SSL Profile screen opens.
3. In the **Name** box, type a name. We type **citrix-server-ssl**.
4. Leave the rest of the settings at the default levels.
5. Click the **Finished** button.

Modifying the virtual server to use the Server SSL profile

The final task in this section is to modify the Web Interface virtual server you created in *Creating the Citrix Web Interface virtual server*, on page 13.

To modify the Web Interface virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. Click the Web Interface virtual server you created in *Creating the Citrix Web Interface virtual server*, on page 13. In our example, we click **citrix_ps.example.com**.
3. From the SSL Profile (Server) list, select the name of the profile you created in the preceding procedure. In our example, we click **citrix-server-ssl**.
4. Click the **Update** button.

This completes the configuration.