



Deploying the F5 BIG-IP LTM with IBM Lotus iNotes

Table of Contents

Deploying the BIG-IP LTM with IBM Lotus iNotes	1
Prerequisites and configuration notes	1
Product versions and revision history	1
Configuration example	2
Configuring the BIG-IP LTM	3
Creating the health monitor	3
Creating the pool	4
Using SSL certificates and keys	5
Creating profiles	6
Creating the virtual server	9
Appendix A: Optional configuration for highly available implementations	12
Creating Data Group Lists	12
Creating the iRule	13
Modifying the virtual server to reference the iRule	15

Deploying the BIG-IP LTM with IBM Lotus iNotes

Welcome to the F5 and IBM Lotus iNotes deployment guide. This guide shows you how to configure the BIG-IP LTM system for a highly available and easily scalable iNotes deployment. The BIG-IP LTM provide users with a seamless failover experience. The user never realizes if the original server with which they were interacting is no longer available; rather, the BIG-IP seamlessly detects any failure and sends the request on to another appropriate server.

IBM® Lotus® iNotes 8.5 software provides a security-rich messaging and collaboration platform for sharing data, connecting your employees and extended communities. It provides a Web browser alternative for accessing IBM Lotus Domino applications, including email calendar, and personal information management (PIM) capabilities, as well as instant messaging and presence awareness.

For more information on iNotes, see:

<http://www-01.ibm.com/software/lotus/products/inotes/>

Prerequisites and configuration notes

The following are prerequisites and notes about this configuration.

- ◆ Working deployment of IBM Lotus Domino 8.5 Email Service, and Lotus Notes 8.5 with the iNotes Web client option installed.
- ◆ F5 BIG-IP LTM must be running version 10.0 or later.
- ◆ For more information on the iNotes configuration, see the IBM Redbook: <http://www.redbooks.ibm.com/redbooks/pdfs/sg246518.pdf>

Product versions and revision history

Product and versions tested for this deployment guide:

Product Tested	Version Tested
BIG-IP LTM	v10.0.1, v10.1
IBM Lotus iNotes	v8.5 (applies to 8.5.1)

Revision history:

Version	Description
1.0	New deployment guide
1.1	Removed support for BIG-IP LTM versions prior to 10.0. For this guide, you must be running LTM version 10.0 or later.

Version	Description
1.2	Corrected the optional iRule on page 14 to add missing spaces in the HTTP Response section.
1.3	Corrected the optional iRule on page 14 to the correct name of the Data Group.

Configuration example

The following is a sample network architecture depicting the BIG-IP managing traffic to the iNotes clients and the iNotes Domino servers. The BIG-IP provides server load balancing, high availability, server health monitoring, and SSL offload services. Additionally, the BIG-IP provides TCP and HTTP protocol optimizations, enabling a superior user experience. The BIG-IP LTMs are deployed as an active-standby pair to provide high availability.




Figure 1 Simple, logical configuration example

Configuring the BIG-IP LTM

In this section, we configure the BIG-IP LTM device for directing traffic to the iNotes devices.

◆ Note

In the following procedures, we assume you are using the BIG-IP LTM to offload SSL from the iNotes devices. If you are not offloading SSL to the BIG-IP LTM, you do not need to import a certificate ([Using SSL certificates and keys](#), on page 5) or create an SSL profile ([Creating a Client SSL profile](#), on page 8). In this case, the BIG-IP LTM virtual server should use port 80 instead of port 443.

Creating the health monitor

The first task is to create a health monitor for the iNotes devices. This procedure is optional, but very strongly recommended. In our example, we create a basic HTTP health monitor.

To create a health monitor

1. On the Main tab, expand **Local Traffic**, and then click **Monitors**.
2. In the upper right portion of the screen, click the **Create** button. The New Monitor screen opens.
3. In the **Name** box, type a name for the Monitor.
In our example, we type **inotes-monitor**.
4. From the **Type** list, select **http**.
5. In the Configuration section, in the **Interval** and **Timeout** boxes, type an Interval and Timeout. We recommend at least a (1:3) +1 ratio between the interval and the timeout (for example, the default setting has an interval of **5** and an timeout of **16**). In our example, we use a **Interval** of **30** and a **Timeout** of **91**.
6. In the **Send String** box, type a send string using the following syntax:

```
GET / HTTP/1.1\r\nHost: <FQDN of your iNotes deployment>\r\nConnection: Close\r\n
```

In our example, we type

```
GET / HTTP/1.1\r\nHost: domino.inotes.example.com\r\nConnection: Close\r\n
```

7. Click the **Finished** button (see Figure 2).
The new monitor is added to the Monitor list.

Local Traffic >> Monitors >> New Monitor...

General Properties

Name	inotes-monitor
Type	HTTP
Import Settings	http

Configuration: Basic

Interval	30 seconds
Timeout	91 seconds
Send String	GET / HTTP/1.1\r\nHost: domino.inotes.example.com\r\nConnection: Close\r\n
Receive String	
User Name	
Password	
Reverse	<input type="radio"/> Yes <input checked="" type="radio"/> No
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No

Cancel Repeat Finished

Figure 2 New Monitor configuration

Creating the pool

The next task is to define a load balancing pool for the iNotes devices. A BIG-IP pool is a set of devices grouped together to receive traffic according to a load balancing method. This pool uses the monitor you just created.

To create the pool

1. On the Main tab, expand **Local Traffic**, and then click **Pools**.
2. Click the **Create** button. The New Pool screen opens.
3. In the **Name** box, type a name for your pool. We type **inotes-pool**.
4. In the **Health Monitors** section, select the name of the monitor you created in *Creating the health monitor*, and click the Add (<<) button. In our example, we select **inotes-monitor**.
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 (node)**.
6. In this pool, we leave the Priority Group Activation **Disabled**.
7. In the New Members section, make sure the **New Address** option button is selected.

8. In the **Address** box, add the first iNotes device to the pool. In our example, we type **10.132.81.100**.
9. In the **Service Port** box, type **80** or select **HTTP** from the list.
10. Click the **Add** button to add the member to the list.
11. Repeat steps 8-10 for each server you want to add to the pool.
In our example, we repeat these steps four times for the remaining servers, **10.132.81.101 - .104**.
12. Click the **Finished** button.

Figure 3 New pool configuration

Using SSL certificates and keys

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 iNotes connections on the BIG-IP LTM device. For this 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*.

◆ **Important**

If you are not using the BIG-IP to offload SSL, you do not need to complete this procedure.

Importing keys and certificates

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.

Creating profiles

The BIG-IP system use configuration objects called profiles. A *profile* is an object that contains user-configurable settings for controlling the behavior of a particular type of network traffic, such as HTTP connections. Using profiles enhances your control over managing network traffic, and makes traffic-management tasks easier and more efficient.

Although it is possible to use the default profiles, we strongly recommend you create new profiles based on the default parent profiles, even if you do not change any of the settings initially. Creating new profiles allows you to easily modify the profile settings specific to this deployment, and ensures you do not accidentally overwrite the default profile.

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. In the following example, we base our HTTP profile off of the **http-wan-optimized-compression-caching** parent.

To create a new HTTP profile

1. On the Main tab, expand **Local Traffic**, and then click **Profiles**. The HTTP Profiles screen opens.
2. In the upper right portion of the screen, click the **Create** button. The New HTTP Profile screen opens.
3. In the **Name** box, type a name for this profile. In our example, we type **inotes-http-opt**.
4. From the **Parent Profile** list, select **http-wan-optimized-compression-caching**.
5. Modify any of the other settings as applicable for your network. In our example, we leave the settings at their default levels.
6. Click the **Finished** button.

Creating the TCP profiles

The next profiles we create are the TCP profiles. If most of the end 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).

Creating the LAN optimized TCP profile

First we configure the LAN optimized 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. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **inotes-tcp-lan**.
5. From the **Parent Profile** list, select **tcp-lan-optimized**.
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.

Creating the WAN optimized TCP profile

Now we configure the WAN optimized profile. 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.

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. In the upper right portion of the screen, click the **Create** button. The New TCP Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **inotes-tcp-wan**.
5. From the **Parent Profile** list, select **tcp-wan-optimized**.
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.

Creating a Client SSL profile

The next step in this configuration is to create a Client SSL profile. This profile contains the SSL certificate and Key information for decrypting the SSL traffic on behalf of the servers.

◆ Important

If you are not using the BIG-IP to offload SSL, you do not need to complete this procedure.

To create a new Client SSL 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 SSL menu, select **Client**. The Client SSL Profiles screen opens.
3. In the upper right portion of the screen, click the **Create** button. The New Client SSL Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **inotes-clientssl**.
5. In the Configuration section, check the **Certificate** and **Key Custom** boxes.
6. From the **Certificate** list, select the name of the Certificate you imported in the *Importing keys and certificates* section.

-
7. From the **Key** list, select the key you imported in the *Importing keys and certificates* section.
 8. Click the **Finished** button.

Creating persistence profile

The next profile we create is a Persistence profile. 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**.
2. On the Menu bar, click **Persistence**.
3. Click the **Create** button. The New Persistence Profile screen opens.
4. In the **Name** box, type a name for this profile. In our example, we type **inotes-cookie**.
5. From the **Persistence Type** list, select **Cookie**.
The configuration options for cookie persistence appear.
6. We recommend leaving the **Cookie Method** list set to **HTTP Cookie Insert**.
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 virtual server

Next, we configure a virtual server that references the profiles and pool you created in the preceding procedures.

To create the 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 **inotes-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 **192.168.10.125**.
6. In the **Service Port** box, type **443**, or select **HTTPS** from the list. If you are not using the BIG-IP LTM to offload SSL, you should type **80** or select **HTTP** from the list.
7. From the Configuration list, select **Advanced**.

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 **inotes-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 **inotes-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 **inotes-http-opt**.
12. From the SSL Profile (Client) list, select the profile you created in *Creating a Client SSL profile*. In our example, we select **inotes-clientssl**.

The screenshot shows the configuration page for a new virtual server. The breadcrumb trail is 'Local Traffic >> Virtual Servers : Virtual Server List >> New Virtual Server...'. The 'General Properties' section includes: Name (inotes-vs), Destination (Type: Host, Address: 192.168.10.125), Service Port (443, HTTPS), and State (Enabled). The 'Configuration' section is set to 'Advanced' and includes: Type (Standard), Protocol (TCP), Protocol Profile (Client) (inotes-tcp-wan), Protocol Profile (Server) (inotes-tcp-lan), OneConnect Profile (None), NTLM Conn Pool (None), HTTP Profile (inotes-http-opt), FTP Profile (None), Stream Profile (None), SSL Profile (Client) (inotes-clientssl), and SSL Profile (Server) (None). At the bottom right, there are 'Enabled' and 'Available' checkboxes.

Figure 4 Virtual Server configuration (truncated)

13. From the **SNAT Pool** list, select **Automap**.

-
14. In the Resources section, from the **Default Pool** list, select the pool you created in the *Creating the pool* section. In our example, we select **inotes-pool**.
 15. From the **Default Persistence Profile** list, select the persistence profile you created in the *Creating persistence profile* section. In our example, we select **inotes-cookie**.
 16. Click the **Finished** button.

	Up	Down
Default Pool	+	inotes-pool
Default Persistence Profile		inotes-cookie
Fallback Persistence Profile		None

Cancel Repeat Finished

Figure 5 *End of the Virtual Server configuration*

This completes the BIG-IP LTM configuration. If you are using a highly available BIG-IP configuration, continue with the following section.

Appendix A: Optional configuration for highly available implementations

Lotus Domino Notes servers can be deployed in several architectures. When deploying Notes in a High Availability architecture, one of these configurations is referred to as a Non-Mirrored Cluster. When configured in this manner, a user's mailbox data exists on more than one member of the cluster, but not all of the members in the cluster, as the mailbox is not replicated to all members of the cluster.

IBM and F5 have created a joint solution to support this advanced architecture. There are 2 requirements for this:

- ◆ The creation of the “Load Balancer Assistance Service”. This is an additional web form, running on each server in the cluster, that provides information to the BIG-IP about the exact URL location of a user's mailbox. It inserts a custom HTTP Header containing a list of members in the cluster that have a copy of a user's mailbox.
- ◆ The creation of the BIG-IP iRule. This is high performance runtime software that will query the cluster members, and using the information provided in the custom HTTP header, correctly route each user's request to the appropriate server.

◆ Important

*You must read and understand the details of this architecture and solution before attempting to configure it in your environment. For more information on how this is configured, see the IBM Developer Works article **Achieving high availability with IBM Lotus iNotes:***

<https://www.ibm.com/developerworks/lotus/library/inotes-avail/>

Creating Data Group Lists

Before we create the iRule, we create two Data Group Lists that the iRule uses. Changing the number of members or modifying IP addresses is much easier in the Data Group List than editing the iRule directly.

It is important to name the Data Group carefully as it is referenced by the iRule we create in the next procedure. If you modify the Data Group name, you must also modify it in the iRule.

To create an address data group

1. On the Main tab, expand **Local Traffic**, and then click **iRules**.
2. On the menu bar, click **Data Group List**.
3. In the upper right corner of the screen, click **Create**.
4. In the **Name** box, type **NSREPLICALOOKUP**.
5. From the **Type** list, select **Address**.

6. In the Records section, select **Host**.
7. In the **Address** box, type the first IP address for the data group.
8. In the **Value** box, type the server's FQDN. In our example, we type **server1.inotes.example.com**.
9. Click **Add**. The entry appears in the Address Records box.
10. Repeat steps 7 - 9 until you have entered all IP addresses. In our example, we add our 4 servers.
11. Click **Finished**.

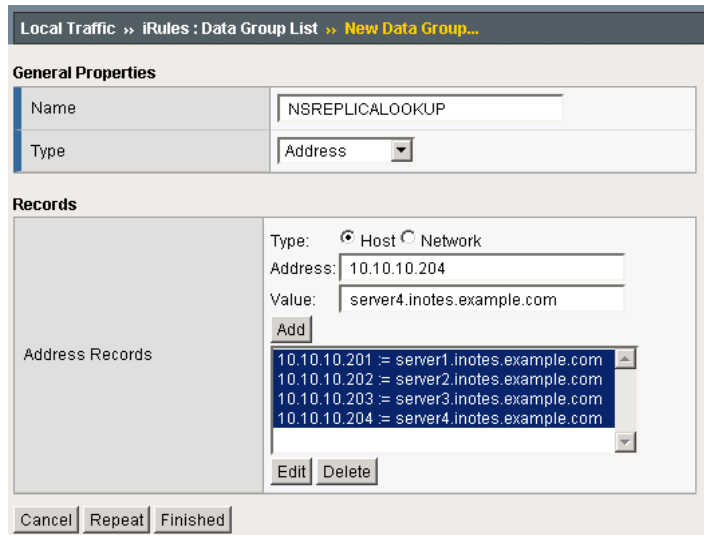


Figure 6 New Data Group

Creating the iRule

The iRule that follows is an example of what is needed to implement this solution. In our example, we have the Log messages commented out. To enable logging, simply remove the comment symbol (#).

To create the iRule

1. On the Main tab, expand **Local Traffic**, and then click **iRules**.
2. Click the **Create** button.
3. In the **Name** box, type a name. We type **inotes-irule**.
4. In the **Definition** box, type the following iRule, omitting the line numbers.

```

1  when CLIENT_ACCEPTED {
2      #set the status - 'needs server' 1 or 0.
3      #log local0. "got initial connect - needs a lookup."
4      set needs_server 0
5  }
6
7  when HTTP_REQUEST {
8      #capture original request - destined for a real server.
9      if { ([HTTP::uri]ends_with ".nsf") and not ([HTTP::uri] contains "names.nsf"){
10         set original_request [HTTP::request]
11         set needs_server 1
12         set nsf "[substr [HTTP::uri] 1 ".nsf"].nsf"
13         HTTP::uri/iwaredir.nsf/ServersLookup?OpenForm&nsfpath=$nsf
14     }
15     } else {
16         set needs_server 0
17     }
18
19     #check to see if we need a server. Else, send to our dest. pool
20     if { $needs_server == 1 } {
21         #dummyServer is our mapping server to query against. It returns the header and its values.
22         pool <inotes-pool>
23     } else {
24         pool <inotes-pool>
25     }
26 }
27 when HTTP_RESPONSE {
28     if { $needs_server == 1 } {
29         set server_list [split [HTTP::header X-Domino-ClusterServers], , ]
30         HTTP::collect[HTTP::header Content-Length]
31     }
32 }
33
34 when HTTP_RESPONSE_DATA {
35     foreach {svr} $server_list {
36         if { "" ne $svr }{
37             set dest [findclass [string trim $svr] :NSREPLICALOOKUP " "]
38             #log local0. "Servername is [string trim $svr]"
39             #log local0. "$dest"
40             #log local0. "server is: $node_addr on $node_port...issuing HTTP::collect"
41             if { [LB::status pool <inotes-pool> member $dest 80 ] eq "up" } {
42                 #log local0. "Selecting $node_addr:$node_port"
43                 pool <inotes-pool> member $dest
44                 HTTP::retry$original_request
45                 break
46             }
47         }
48     }
49 }
50 set needs_server 0
51 }

```

5. Click the **Save** button.

Modifying the virtual server to reference the iRule

The next task is to modify the virtual server you created in *Creating the virtual server*, on page 9 to use the iRule you just created.

To modify the existing virtual server

1. On the Main tab, expand **Local Traffic**, and then click **Virtual Servers**.
2. From the **Virtual Server** list, click the iNotes virtual server you created in the *Creating the virtual server*, on page 9 section. In our example, we click **inotes-vs**.
3. On the menu bar, click **Resources**.
The Resources page for the virtual server opens.
4. In the iRules section, click the **Manage** button.
The Resource Management screen opens.
5. From the **Available** list, select the iRule you created in *Creating the iRule*, on page 13, and click the Add (<<) button. In our example, we select **inotes-irule**.
6. Click the **Finished** button.

This completes the configuration.