VMWARE HORIZON WITH VIEW AND F5 BIG-IP VS. CITRIX XENDESKTOP WITH NETSCALER

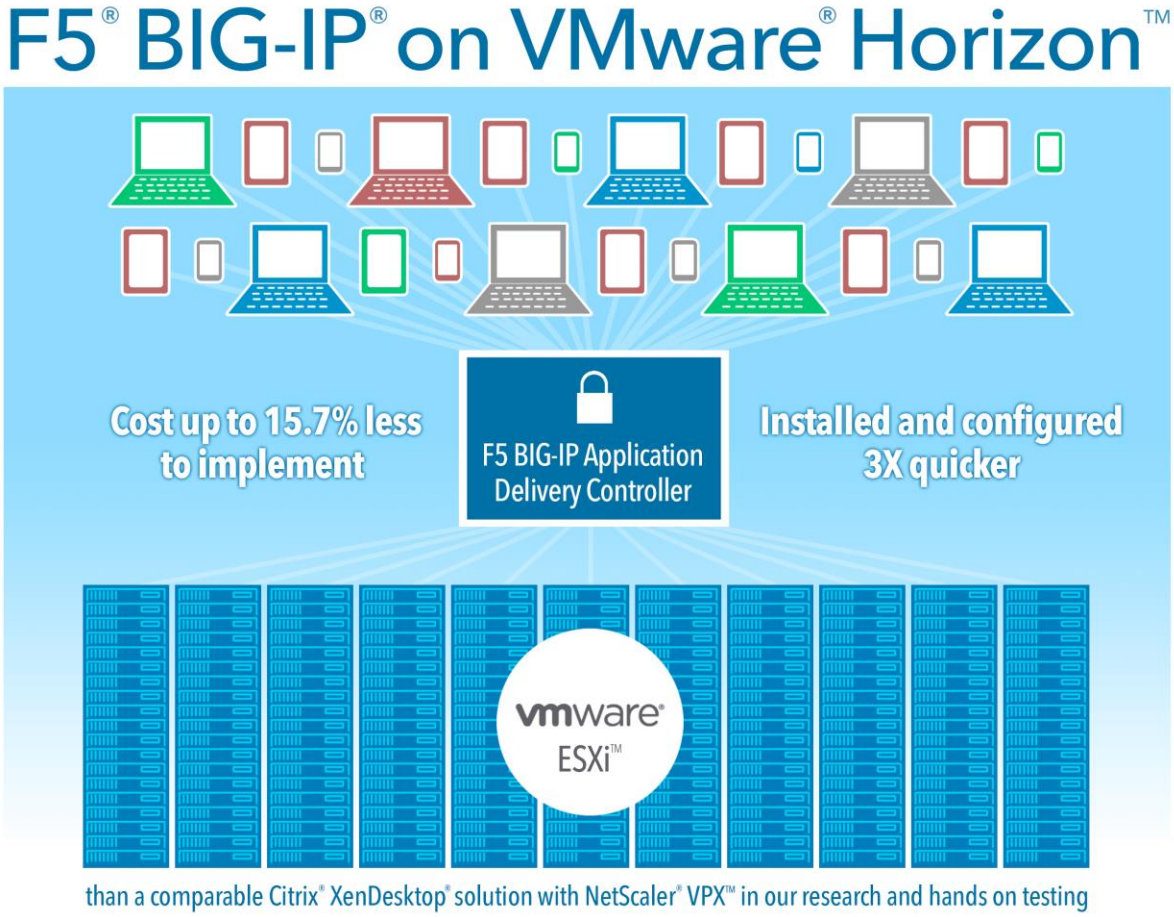

Enterprises increasingly rely on virtual desktop infrastructure (VDI) to meet the desktop computing needs of their workers. These VDI deployments benefit from having more than one connection server to provide high availability. Some enterprises choose to use Application Delivery Controllers (ADCs) to help balance the load across the multiple connection servers in their environment. ADCs, which also provide a number of other capabilities such as access control, SSO, and protocol proxy, vary significantly in cost and in how easy they are to implement and use.

In the Principled Technologies labs, we tested two VDI solutions, each with its own ADC, to see how easy each was to install and configure: VMware Horizon with View 5.3 with F5 BIG-IP 11.5 and Citrix XenDesktop 7.1 with Citrix NetScaler 10.1. We also performed a cost analysis.

We found that BIG-IP was considerably easier and three times faster to integrate with an existing Horizon VDI environment than Citrix NetScaler with its respective VDI environment, Citrix XenDesktop. In addition, the VMware-F5 solution was up to 15.7 percent less expensive than the Citrix solution. Together, these findings make VMware Horizon and F5 BIG-IP an attractive choice compared to alternatives on the market.

APRIL 2014

A PRINCIPLED TECHNOLOGIES TEST REPORT

Prepared for VMware and F5

VMWARE HORIZON AND F5 BIG-IP—A GREAT COMBINATION

VMware Horizon provides organizations with the ability to virtualize user desktops by delivering them to individual clients over the network from a central location. Those desktops are stored and run in the data center, rather than on individual desktop and laptop computers running localized operating systems, and are much easier for IT to manage. As long as performance is strong, users do not notice this seamless virtualization.

In a VDI environment, performance must compare favorably to a conventional desktop while availability and security must be even greater. F5 offers a variety of solutions to meet requirements for secure access, a single namespace, load balancing, server health monitoring, and more.

VMware Horizon Optimized Secure Access & Traffic Management by F5 provides valuable secure remote access, load balancing, and health monitoring that can lead to increased system availability and greater scalability.

F5 also provides detailed and easy to follow deployment guides to assist during the installation and configuration of BIG-IP, eliminating the guesswork often associated with the deployment and integration of ADCs.

What ADCs can do

An Application Delivery Controller typically goes between the firewall and one or more application servers in a data center. While earlier ADCs primarily worked to accelerate application performance and balance load between servers, the latest generation of ADCs, such as the F5 BIG-IP product family, does much more, including SSO, WAN optimization, rate shaping, and SSL offloading, as well as serving as a Web application and data center network firewall.

The BIG-IP system is an Application Delivery Controller and full proxy between users and application servers, creating a layer of abstraction to secure, optimize, and load balance VMware Horizon traffic. BIG-IP can make in-depth application decisions without creating bottlenecks for clients and ensure that only allowed services pass through to the View servers.

iApps templates make BIG-IP easy to implement

F5 iApps® templates make configuration straightforward, simplifying setup by providing the recommended settings and helping to prevent human error.

iApps greatly simplifies deploying BIG-IP and is extremely configurable to accommodate any level of complexity required by a particular installation. iApps could also help to automate different types of tasks by implementing a simple to use scripting language and a question-driven interface which allows users to interact with the application.

VMWARE HORIZON WITH F5 BIG-IP—EASIER TO SET UP AND LESS EXPENSIVE THAN CITRIX XENDESKTOP WITH CITRIX NETSCALER

In the Principled Technologies labs, we set out to compare the installation and setup experience of VMware Horizon and F5 BIG-IP with that of a competing ADC-enhanced VDI solution, Citrix XenDesktop with Citrix NetScaler. To do so, we started with a fully configured VDI environment in our labs and then deployed and configured each ADC solution, measuring the time necessary for each. As Figure 1 shows, setting up and installing the BIG-IP solution on top of View took less than half an hour, while doing the same with the Citrix solution took 1 hour and 25 minutes. That is a time savings of 56 minutes.

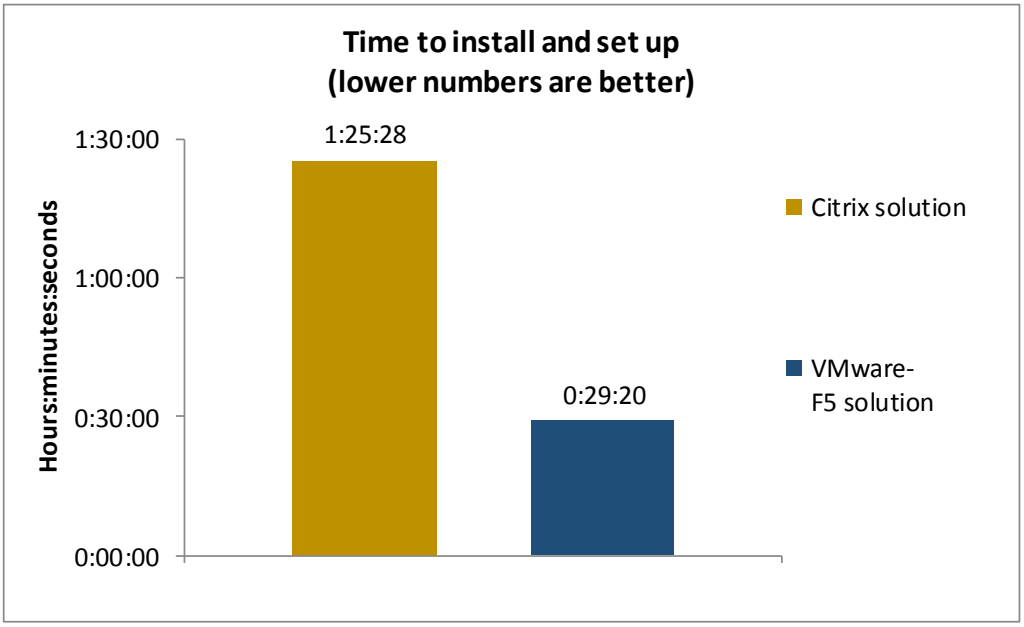

Figure 1: Time to install and set up the two solutions. Lower numbers are better.

In addition to the time saved during installations and setup, the VMware-F5 solution has the potential to save organizations money as well. As Figure 2 shows, the cost per concurrent user with the VMware-F5 solution is from 14.5 percent to 15.7 percent less than with the Citrix XenDesktop-NetScaler solution.

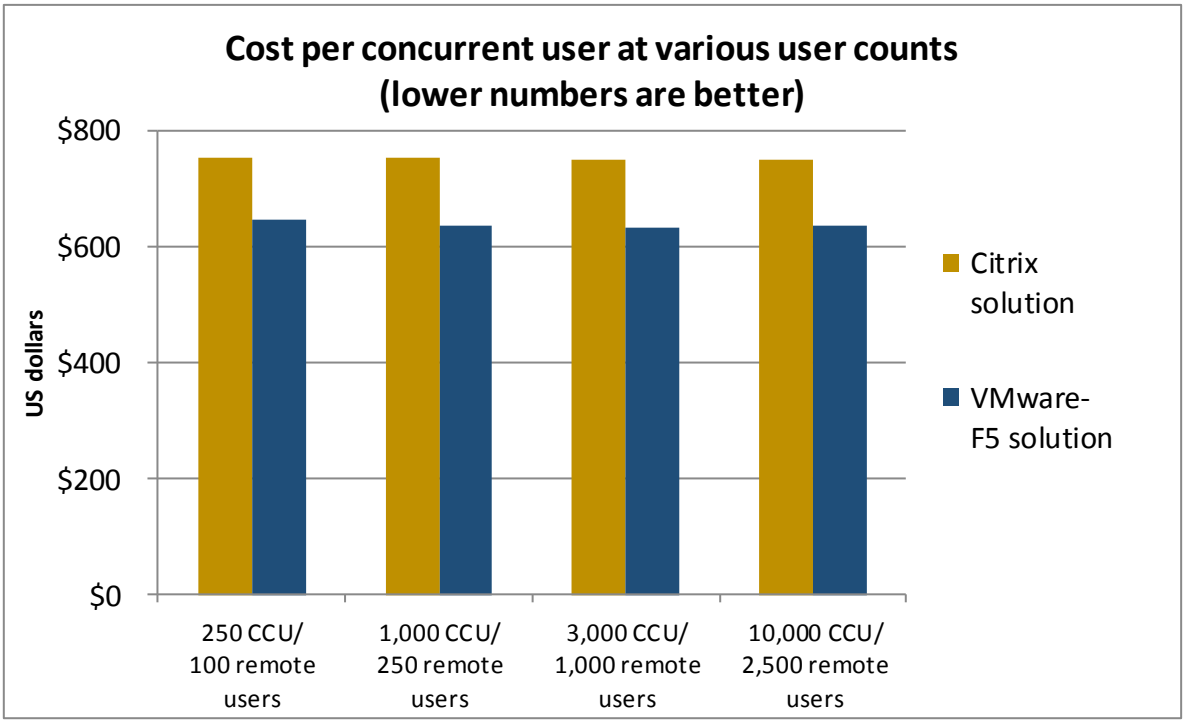

Figure 2: Cost per concurrent user with the two solutions. Lower numbers are better.

Below, we present more detailed results and discussion of the setup and installation advantages and the pricing benefits of the VMware-F5 solution.

See Appendix A for more about our cost analysis, and Appendix B for the specifics of our testing.

DETAILED TEST RESULTS

VMware Horizon and F5 BIG-IP—Quicker and easier to deploy

Less time

Figure 3 shows the steps involved in setting up and installing the two solutions and the time (in hours:minutes:seconds) that each step took our technicians to perform. As it shows, the time to carry out the steps on the VMware-F5 solution was considerably less than on the XenDesktop with NetScaler solution. Setting up and installing the VMware-F5 solution took less than half an hour, while doing the same with the Citrix solution took 1 hour and 25 minutes. That is 3 times faster.

VMware Horizon with F5 BIG-IP		Citrix XenDesktop and Citrix NetScaler	
Obtain and install SSL certificates	0:04:03	Obtain and install ISS SSL certificates / configure HTTPS	0:08:27
Configure View standard and replica server	0:01:15	Configure / join StoreFronts	0:25:25
Configure locked.properties file	0:00:56	Deploy NetScaler VPX 1 & 2	0:02:04
Deploy BIG-IP systems 1 and 2	0:04:51	Set management IP on NetScaler 1	0:01:29
Set management IP on BIG-IP systems 1 and 2	0:01:26	Install license on Netscaler 1	0:01:58
Install license on BIG-IP system 1 and set APM	0:00:37	Set management IP on NetScaler 2	0:01:37
Add device SSL certificate to BIG-IP system 1 / configure host name	0:01:16	Install license on Netscaler 2	0:01:36
Configure network on BIG-IP system 1 / configure failover	0:03:39	Request and add SSL certificates	0:08:11
Install license on BIG-IP system 2 and set APM	0:00:39	Add Load Balance Virtual server	0:13:47
Add device SSL certificate to BIG-IP system 2 / configure host name	0:01:22	Add Gateway Virtual server	0:04:39
Configure network on BIG-IP system 2/configure HA	0:03:11	Configure LDAP	0:04:28
Install View SSL certificate and import iApps	0:01:37	Configure Gateway	0:08:37
Create iApps application service and Configure iApps template	0:04:07	Add node and synchronize NetScalers	0:03:00
Synchronize BIG-IP systems	0:00:21	Enable Fail-Safe / HA and force synchronize	0:00:10
Total	**0:29:20**	**Total**	**1:25:28**

Figure 3: Steps and time required to install and set up the two solutions. Lower numbers are better.

VMware Horizon and F5 BIG-IP—Less expensive to implement

As we have shown, the VMware-F5 solution is easier to deploy and integrate than Citrix XenDesktop with NetScaler, delivering savings in IT staff time and IT operational costs for VDI deployments that need to support remote user connections. These operating expense savings are not the only advantages. You also save on software costs because F5 BIG-IP and VMware Horizon are also much less expensive than the Citrix products.

We compared costs for the two solutions using four scenarios based on VDI concurrent user (CCU) counts between 250 and 10,000. We included VDI software costs for those users. We sized and calculated costs for Citrix NetScaler and F5 BIG-IP access servers for the minority of those users connecting remotely.

Across all four scenarios, VMware Horizon Enterprise with Production Support and F5 BIG-IP offered significant cost savings compared to Citrix XenDesktop Platinum with Premier Support and Citrix NetScaler Gateway VPX. As Figure 4 shows, these savings range from $109 to $118 per CCU (see Figure 4).

Concurrent VDI users*	Remote users **	XenDesktop Platinum with Premier Support and 2 x Citrix NetScaler Gateway VPX	VMware Horizon Enterprise with Production Support and F5 BIG-IP solutions	Savings	Percentage savings	Savings per CCU
250	100	$188,565	$161,210	$27,355	14.5%	$109
1,000	250	$752,130	$634,920	$117,210	15.6%	$117
3,000	1,000	$2,254,260	$1,899,800	$354,460	15.7%	$118
10,000	2,500	$7,510,650	$6,354,148	$1,156,502	15.4%	$116

Figure 4: Costs for the two solutions and savings for the VMware and F5 BIG-IP solution. Prices are rounded to the nearest dollar.
*** Concurrent VDI users is the total number of concurrently connected users, both those accessing the solution remotely and those accessing the solution using an internal LAN.**
**** For both ADC solutions, only the remote users count against the total number of sessions. Internal LAN-based users do not count against the session license limit.**

See Appendix A for cost analysis details.

CONCLUSION

Many enterprises use Application Delivery Controllers to balance the load on their VDI servers and to provide other capabilities. In our tests, we found that the solution with VMware Horizon and F5 BIG-IP was roughly three times quicker to set up and install than a competing solution, Citrix XenDesktop with NetScaler. The VMware-F5 solution was also more cost effective, saving up to 15.7 percent per concurrent user. These findings make VMware Horizon and F5 BIG-IP a very compelling choice.

APPENDIX A – COST ANALYSIS DETAILS: SOFTWARE COST SAVINGS WITH F5 BIG-IP AND VMWARE HORIZON

Scenarios

In typical VDI deployments, the majority of the users directly connect to the VDI systems over the company's network via the LAN or WAN while a minority of users connect to the VDI systems remotely via the Internet. Only the remote users go through the Citrix NetScaler or F5 BIG-IP access servers to reach the VDI systems. For each of the four scenarios, we assumed between 25 percent and 40 percent of the CCUs would be remote connections.

We calculated VDI software costs and access server costs. We include costs for licenses and one-year support for each product.

- **VDI software.** We use concurrent-user licensing for each Citrix XenDesktop and VMware Horizon and calculate those costs based on the CCU counts for each scenario.
- **Access servers.** We selected the F5 BIG-IP solutions to match our estimated remote connection user counts. We compare that to costs for the minimum number of NetScaler Gateway VPX instances required support these remote users, assuming one NetScaler Gateway VPX license per 500 remote connections.

Assumptions

- We used list prices before any discounts.
- We selected support packages that provide 24-hour/7-day/year-round support.
- We compared costs for four concurrent-user seat counts between 250 and 10,000 seats. We used the seat counts to calculate the number of VMware Horizon and Citrix XenDesktop licenses needed.
- For each of those counts, we estimated between 25 percent and 40 percent of the users are remote users and used different percentages within that range across the four scenarios.
- F5 recommended the F5 BIG-IP solutions for each remote user count and calculated the quantities needed to provide F5 BIG-IP instances for those users.
- We included costs for one Citrix NetScaler Gateway VPX instance per 500 remote connections. Citrix XenDesktop Platinum includes Citrix NetScaler Gateway universal license, so we do not include additional costs for it.

Software cost data

We collected software cost data from vendor and a tier 1 reseller.

Citrix costs

Citrix XenDesktop Platinum with Premier Support

For each concurrent-user connection, we included costs for Citrix XenDesktop Platinum at $700 for the license and one year Citrix Subscription Advantage and $50 for one-year Premier Support. Citrix Subscription Advantage provides access to product updates during the membership term. Citrix Premier Support provides 24 x 7 x 365 unlimited worldwide support and unlimited support incidents on software products. Total costs for Citrix XenDesktop are $750 per CCU.

Citrix NetScaler Gateway VPX

Citrix NetScaler Access Gateway is available as a virtual platform (VPX) and as a hardware platform. We included prices for the VPX platform only. We included the costs of licenses for virtual machine instances of Citrix NetScaler Gateway VPX Enterprise Edition at $995 per license, which includes one year of Subscription Advantage. We added $70 for one year Premier Support. Total costs for Citrix NetScaler Gateway VPX are $1,065 for each instance. We assumed each instance supports up to 500 remote connections. Figure 5 shows the costs we calculated for the four scenarios.

Scenarios	Number Citrix NetScaler Gateway VPX instances	License cost
250 CCU/100 remote connections	1	$1,065.00
1,000 CCU/250 remote connections	1	$2,130.00
3,000 CCU/1,000 remote connections	2	$4,260.00
10,000 CCU/2,500 remote connections	5	$10,650.00

Figure 5: Citrix NetScaler Gateway VPX costs.

VMware Horizon and F5 BIG-IP costs

VMware Horizon Enterprise with Production Support

VMware Horizon Enterprise with 1-year Production support is available for $6,250 for 10 CCU, an average of $625 per user. Production support offers 24 x 7 x 365 worldwide support.

F5 BIG-IP Application Delivery Controller

F5 suggested the following BIG-IP solutions and quantities and provided the prices. Costs include one-year support. Product choices are based on bandwidth estimates, which are 25Mb for 100 remote users, 200Mb for 250 remote users, 1Gb for 1,000 remote users, and 3Gb for 2,500 remote users. The BIG-IP product names identify the bandwidth license limits. Figure 6 summarizes the F5 costs we use in this analysis.

Scenario	F5 BIG-IP product	License cost per unit	Number units	Total license cost	1st year cost (24% of license cost)	Total license plus one-year support cost	Average cost per CCU
250 CCU/100 remote connections	F5-BIG-APMVE-VW-25M	$2,000.00	2	$4,000.00	$960.00	$4,960.00	$19.84
1,000 CCU/250 remote connections	F5-BIG-APMVE-VW-200M	$4,000.00	2	$8,000.00	$1,920.00	$9,920.00	$9.92
3,000 CCU/1,000 remote connections	F5-BIG-APMVE-VW-1G	$10,000.00	2	$20,000.00	$4,800.00	$24,800.00	$8.27
10,000 CCU/2,500 remote connections	F5-BIG-VE-BT-3G-V12	$41,995.00	2	$83,990.00	$20,157.60	$104,147.60	$10.41

Figure 6: F5 costs.

F5 BIG-IP provides remote connections at a low cost

You get the advantage of remote connections at a low cost with F5. Figure 6 shows the average cost to provide remote connections with F5 BIG-IP. We calculated this cost by comparing the cost of the standard VMware Horizon Enterprise & Security Server with the cost of the same package with F5 BIG-IP included. As Figure 7 indicates, the cost of including F5 BIG-IP was less than $20 per CCU for all four scenarios.

Concurrent users	Remote users	VMware Horizon Enterprise & Security Server per CCU	VMware Horizon Enterprise & Security Server with F5 BIG-IP average per CCU	Additional cost necessary for F5 BIG-IP per CCU
250	100	$625.00	$644.84	$19.84
1,000	250	$625.00	$634.92	$9.92
3,000	1,000	$625.00	$633.27	$8.27
10,000	2,500	$625.00	$635.41	$10.41

Figure 7: Average cost per CCU for VMware Horizon with and without F5 BIG-IP.

APPENDIX B – DETAILED TEST METHODOLOGY

Application Delivery Controllers tested

ADC	Model
F5 BIG-IP VE	11.5.0.0.0.221
Citrix NetScaler VPX	NS10.1 Build 119.7

Figure 8: ADCs we tested.

Equipment

We set up an iSCSI storage array to be used by the infrastructure server and the virtual desktop host server. The storage was connected via 10G redundant connections to a Dell™ PowerConnect™ 6248 that provided 1Gb access to the infrastructure server and virtual desktops host server. Storage management connections and server connections were handled by 1Gb connections to a Dell PowerConnect 5448. We installed VMware vSphere 5.5 on the Infrastructure and Virtual Desktop host servers. Figure 9 shows detailed information for the servers.

Server name	Model	Processors	Memory (GB)	Functional role
Infrastructure	Dell PowerEdge R620	(2) Intel Xeon E5-2660	64	Infrastructure host
Virtual Desktop Host	Intel Black Box Server	(2) Intel Xeon E5-2640	256	VDI host

Figure 9: Detailed server specifications.

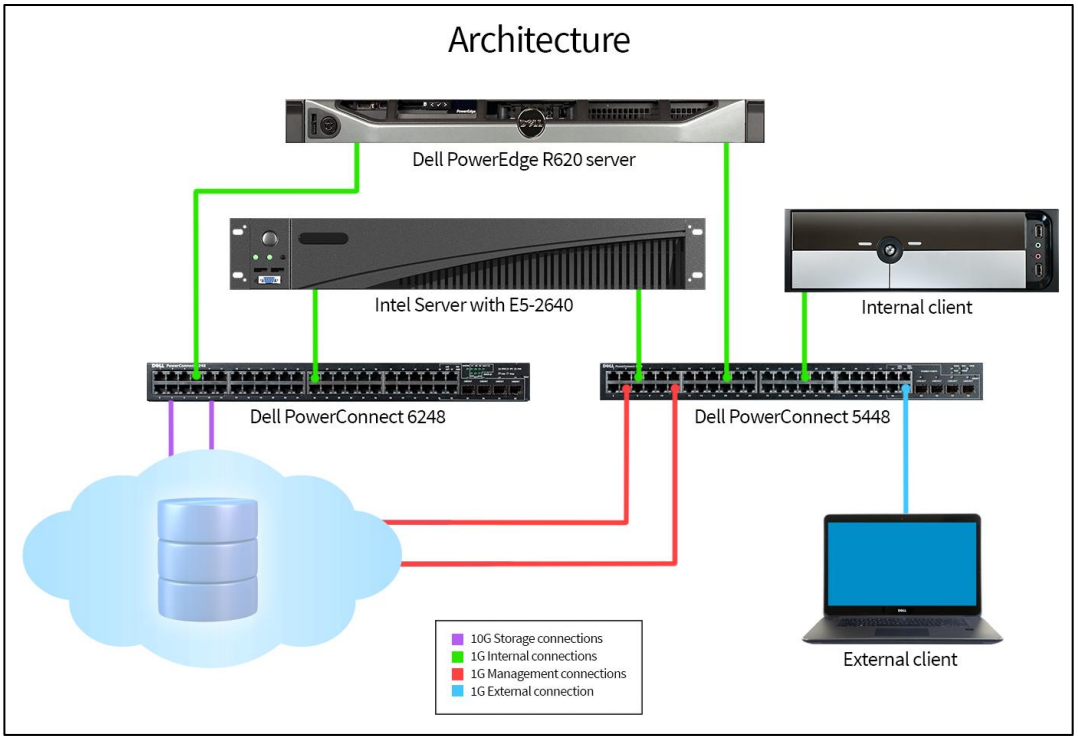

Figure 10: Test architecture.

VM and storage configuration

VM name	Qty.	OS	Role	Host(s)	LUN	vCPUs	Mem (GB)	vDisk (GB)
AD	1	Windows 2008 R2	Active Directory, DNS, DHCP, CA.	Infra	Infra	2	4	40
SQL	1	Windows 2008 R2	Windows SQL 2008 R2	Infra	Infra	2	2	40
vCenter	1	Windows 2008 R2	VMware vCenter	Infra	Infra	2	8	100
Composer	1	Windows 2008 R2	View Composer Server	Infra	Infra	2	4	40
Connection	2	Windows 2008 R2	View Connection Server	Infra	Infra	2	4	40
F5-BIG IP	2		F5 BIG-IP VE	Infra	Infra	2	4	100
View001 - 100	100	Windows 7 (x86)	VMware Virtual Desktops	Virtual Desktop Host	VDT	1	1	24
License	1	Windows 2008 R2	Citrix Licensing Server	Infra	Infra	2	4	40
Delivery Controller	2	Windows 2008 R2	XenDesktop Delivery Controller	Infra	Infra	2	4	40
StoreFront	2	Windows 2008 R2	XenDesktop StoreFront	Infra	Infra	2	4	40
NetScaler	2		Citrix NetScaler	Infra	Infra	2	2	20
XD001 - 100	100	Windows 7 (x86)	XenDesktop Virtual Desktops	Virtual Desktop Host	VDT	1	1	24

Figure 11: Detailed VM configuration.

LUN name	Size (GB)
Infra	1,024
VDT	2,048

Figure 12: Detailed storage configuration.

We configured a Dell EqualLogic PS6110 array and provision two separate storage pools. The first pool hosted infrastructure virtual machines and the second one hosted virtual desktops exclusively.

Prerequisites and configuration notes

This guide assumes you have a Virtual infrastructure (View Composer, View Connection Servers, Citrix License Server, Delivery Controllers, StoreFront), Active Directory domain, Certificate Authority, SQL and vCenter already in place. Installing either is outside of the scope of this deployment guide.

Installing and integrating F5 BIG-IP ADC with Horizon

- Virtual servers and virtual desktops
  - 2 View Connection Servers
  - 1 View Composer

    - 1 MSSQL server
    - 100 Horizon virtual machines
- Two BIG-IP Virtual Editions
    - HA configuration requires two or more BIG-IP systems.
    - This guide sets up an active/passive HA configuration.

Connection servers

- con1.view.mycompany.com
- con2.view.mycompany.com

MSSQL server

- sql1.view.mycompany.com

View Composer

- com1.view.mycompany.com

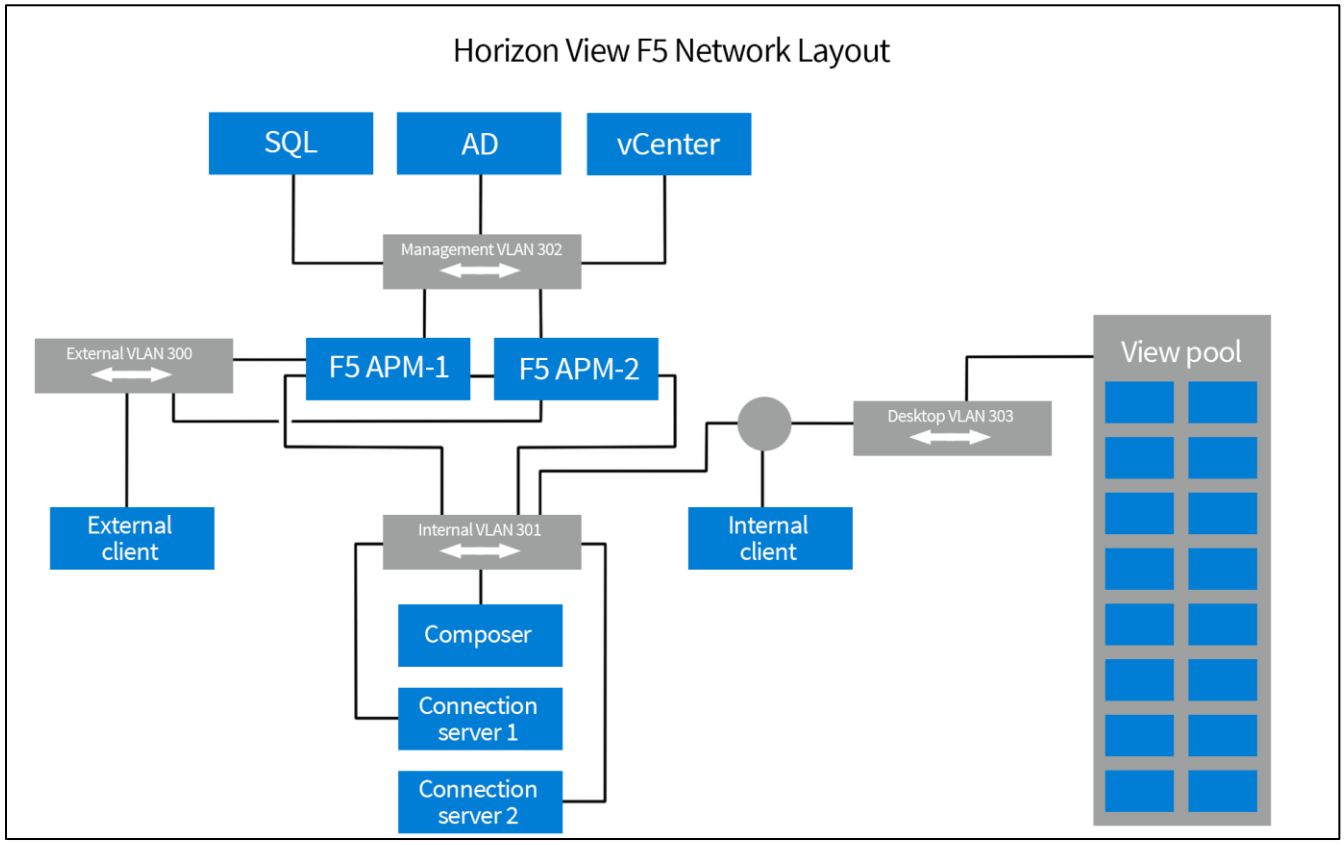

Figure 13: VMware Horizon and F5 BIG-IP network layout.

Obtaining and installing the SSL certificate for your View environment

Although there are many ways to obtain SSL certificates, the following section explains how to create a Certificate Signing Request (CSR) file, which can then be used to obtain an SSL certificate from your internal or third-party certificate authority (CA).

The following example CSR uses **certreq** (available on 2008 R2 servers).

1. Build the request.inf file and save it onto con1.view.mycompany.com using the example below. Make sure to modify the example as described in the following. Required changes are in red text in the example.
    a. Replace CN=remote.view.mycompany.com with the FQDN used by clients to access your Horizon environment. Note that in our reference architecture, we suggest using split DNS, which allows you to use the same FQDN for both remote (untrusted) and local (trusted) client connections.
    b. Fill in the appropriate information for OU=OU, O=Org, L=City, S=State, C=Country.
    c. Modify the Key Length value to the appropriate setting. We suggest using at least a length of 2048.

```
;----------------- request.inf -----------------
[Version]
Signature="$Windows NT$
[NewRequest]
Subject = "CN=remote.view.mycompany.com, OU=OU, O=Org, L=City, S=State, C=Country"
KeySpec = 1
KeyLength = 2048
; KeyLength is usually chosen from 2048, 3072, or 4096. A KeyLength
; of 1024 is also supported, but it is not recommended.
Exportable = TRUE
MachineKeySet = TRUE
SMIME = False
PrivateKeyArchive = FALSE
UserProtected = FALSE
UseExistingKeySet = FALSE
ProviderName = "Microsoft RSA SChannel Cryptographic Provider"
ProviderType = 12
RequestType = PKCS10
KeyUsage = 0xa0
[EnhancedKeyUsageExtension]
OID=1.3.6.1.5.5.7.3.1 ; this is for Server Authentication
;-------------------------------------------------
```

2. Create a CSR using Certreq:
    a. On con1.view.mycompany.com, a Windows 2008 R2 server, open a command prompt as an administrator.
    b. Navigate to the directory where you saved your request.inf file; for example, `cd c:\cert`
    c. Type `certreq –new request.inf viewcert.txt`
    d. Open viewcert.txt with Notepad and copy the contents.
3. Using the guidelines for your organization, submit the CSR to an internal or third party CA by pasting the copied contents when prompted for the CSR.
4. Copy the certificate that is returned onto con1.view.mycompany.com and save it as view.cer, using the same folder used to create the CSR. In some cases, the certificate is returned in base 64 digital format rather than a file. Copy the returned ASCII characters into a file you create with the name view.cer.
5. Create and install the certificate and key onto CON1 using Certreq.
    a. Make sure both the CSR file and returned certificate are located in C:\cert.

b. Use the following command: `Certreq –accept view.cer`

6. Open the Microsoft Management Console with the certificates snap-in using the following guidance.
    a. Click Start, Run and then type `mmc`. The Console opens.
        i. Click File, and then Add/Remove Snap-in.
        ii. From the list of Snap-ins, select Certificates and then click Add.
        iii. When asked what type of certificates to manage, click Computer account.
        iv. When asked to select the computer you want the snap-in to manage, select Local Computer.
        v. Click Finish.
        vi. Click OK.
    b. View the locally installed certificates using the following guidance.
        i. Under Console Root, expand Certificates and Personal, and then click Certificates
        You should now see a certificate with the FQDN you entered into your request.inf file. In our example we see a certificate named remote.view.mycompany.com.
        ii. Verify you see a key symbol in the upper left hand corner of the certificate. This indicates the Cert includes the private key.
    c. Add a friendly name to certificate.
        i. Right-click the certificate and then click Properties.
        ii. In the Friendly Name field, type `vdm`
        iii. Vdm is used by View to indicate which certificate should be used in the View environment. Make sure you have only one certificate installed with a friendly name of vdm.
        iv. Click OK.
    d. Export the Certificate and Key.
        i. Right-click the certificate, and then from the All tasks menu, click Export. The Export Wizard opens.
        ii. On the Export Private Key page, select Yes, export the private key.
        iii. On the Export File Format page, check the Include all certificates in the certificate path if possible box.
        iv. On the Password page, type and then confirm a password.
        v. On the File to Export page, specify the file name, and specify a secure location to save the exported certificate.
        vi. Click Finish.
    e. Add intermediate and root certificates.
        i. Make sure the CA that issued your certificate is located as a root authority.
        ii. Note: Root and/or intermediate certificates can be obtained from the Certificate Authority.
        iii. Under Console Root, expand Certificates and Trusted Root Certification Authorities, and then click Certificates.
        iv. Import the Root server if not present.
            - In the left panel, right-click Certificates and from the All tasks menu, click Import. The Import Wizard opens.
            - On the File to Import page, select the file you to import.
            - Walk through the rest of the wizard, and then click Finish.
        v. Import Intermediate Certificates if required.
        vi. In some cases, an intermediate CA is used rather than the root server to protect the identity of the CA root servers. If so, import intermediate CA certificate using import wizard.
            - Under Console Root, expand Certificates and Intermediate Certification Authorities, and then click Certificates.
            - Right-click Certificates and from the All tasks menu, click Import. The Import Wizard opens.
            - On the File to Import page, select the file you to import.
            - Walk through the rest of the wizard, and then click Finish.

7. Apply the exported certificate to all View servers. In our environment, we import remote.view.mycompany.com to servers con2 and com1using the import wizard in the MMC certificates snap-in. You also need to install root and intermediate server certificates if they are not present.
    a. Go to the server (com2 in our example) and open MMC (see step 6a above on how to open MMC if necessary).
        i. Click File, and then Add/Remove Snap-in.
        ii. From the list of Snap-ins, select Certificates and then click Add.
        iii. When asked what type of certificates to manage, click Computer account.
        iv. When asked to select the computer you want the snap-in to manage, select Local Computer.
        v. Click Finish.
        vi. Click OK.
    b. Import the certificates using the following guidance.
        i. Under Console Root, expand Certificates and Personal, and then click Certificates
        ii. In the left panel, right-click Certificates and then from the All tasks menu, click Import. The Import Wizard opens.
        iii. On the File to Import page, select the certificate you previously exported.
        iv. Walk through the rest of the wizard, and then click Finish.
    c. Import root and intermediate certificates if necessary.
    d. Repeat on the other server (com1.view.mycompany.com in our example).

Configure View Connection Servers

1. Modify the Connection Server to use the remote FQDN supplied in the SSL certificate:
    a. From the View Configuration tab, select Servers.
    b. Click Connection Servers.
    c. Highlight CON1 and then click Edit.
    d. Modify the HTTP External URL and BLAST External URL to match the URL of your SSL certificates. In our example, we use https://remote.view.mycompany.com:443.
        i. Clear the check from Use Secure Tunnel connection to desktop and Use Blast Secure Gateway for HTML access to desktop after modifying the External URLs.
    e. Click OK.
    f. Highlight CON2 and then click Edit.
    g. Modify the HTTP External URL and BLAST External URL to match the URL of your SSL certificates. In our example, we use https://remote.view.mycompany.com:443.
        i. Clear the check from Use Secure Tunnel connection to desktop and Use Blast Secure Gateway for HTML access to desktop after modifying the External URLs.
2. Verify the settings from the Dashboard tab.
    a. Click the Dashboard menu item.
    b. Review the Connection Server status to confirm the status has changed from red to green. If necessary, click refresh as status changes can take a couple of minutes to complete.

Allowing HTTP connections to intermediate servers

When SSL is offloaded to an intermediate server, you can configure View Connection Server instances to allow HTTP connections from the client-facing BIG-IP system. The BIG-IP system must accept HTTPS for View Client connections.

To allow HTTP connections between View servers and BIG-IP system, you must configure the locked.properties file on each View Connection Server instance on which HTTP connections are allowed.

Even when HTTP connections between View servers and intermediate devices are allowed, you cannot disable SSL in View. View servers continue to accept HTTPS connections as well as HTTP connections.

To configure the locked.properties file

1. Create or edit the locked.properties file in the SSL gateway configuration folder on the View Connection Server host. For example: install_directory\VMware\VMware View\Server\sslgateway\conf\locked.properties
2. To configure the View server's protocol, add the serverProtocol property and set it to http. The value http must be typed in lower case.
3. Optional: Add properties to configure a non-default HTTP listening port and a network interface on the View server.
    i. To change the HTTP listening port from 80, set serverPortNonSSL to another port number to which the intermediate device is configured to connect.
    ii. If the View server has more than one network interface, and you intend the server to listen for HTTP connections on only one interface, set serverHost to the IP address of that network interface.
4. Save the locked.properties file.
5. Restart the View Connection Server service to make your changes take effect.

For example, the following locked.properties file allows non-SSL HTTP connections to a View server. The server uses the default port 80 to listen for HTTP connections. The value http must be lower case.

serverProtocol=http
serverHost=<server's client-facing IP>

Configuring the BIG-IP system

Configuring the BIG-IP system is broken into three distinct sections:

- Deploying the BIG-IP OVF on this page
- Performing the initial BIG-IP system configuration
- Installing and configuring the iApps template

Deploying the BIG-IP OVF on ESXi 5.5

Use this section to deploy the BIG-IP software onto a virtual machine.

1. Launch VMware vCenter and log in.
2. From the File menu, click Deploy OVF Template. The Deploy OVF Template wizard opens. Complete the following.
    a. Click Browse and go to the location you saved the BIG-IP system OVF. Click Open, and then back on the Source Page, click Next.
    b. On the OVF Template Details page, review the details and then click Next.
    c. Read the End User Software License. When finished, click Accept, and then click Next.
    d. From the Name and Location page, in the Name box, type a unique name for this VM, and then select an inventory location. Click Next.
    e. On the Deployment Configuration page, select the appropriate number of CPUs and amount of RAM for this deployment and then click Next.
        i. Use the appropriate settings based on your license. The VMware virtual machine guest environment for the BIG-IP Virtual Edition (VE), at minimum, must include:
            - 2 x virtual CPUs
            - 4 GB RAM
            - 1 x VMXNET3 virtual network adapter or Flexible virtual network adapter (for management)

- 1 x virtual VMXNET3 virtual network adapter (three are configured in the default deployment for dataplane network access)
- 1 x 100 GB SCSI disk, by default
- 1 x 50 GB SCSI optional secondary disk, which might be required as a datastore for specific BIG-IP modules. For information about datastore requirements, refer to the BIG-IP module's documentation.
  Important: Not supplying at least the minimum virtual configuration limits will produce unexpected results.

f. On the Host/Cluster page, select appropriate cluster and then click Next.
g. On the Resource Pool page, select the appropriate pool and then click Next.
h. On the Storage page, select the appropriate destination storage for the virtual machine and then click Next.
i. On the Disk Format page, click Next.
j. On the Network Mapping page, select the appropriate Destination Networks for each of the Source Networks. By default, there are four source networks: Management, Internal, External, and HA.
   i. Management is used specifically for managing this BIG-IP instance and is not used to pass production traffic.
   ii. Internal, for this View reference architecture, is used to reach your backend View Connection servers and Virtual Desktop networks. It is also used by internal trusted View clients.
   iii. External, for this View reference architecture, is used for remote untrusted View Client connections.
   iv. HA is used for communication to and from the secondary BIG-IP system.
   v. When you have mapped all of the networks, click Next.
k. On the Ready to Complete page, review the deployment settings. Use the Back button to make any changes. If the settings are correct, check Power on after deployment and then click Finish.
l. After vCenter deploys and powers on the BIG-IP system, click the Summary tab of the new virtual machine and then click Open Console.
m. At the localhost login prompt, type `root`. At the password prompt, type `default`
n. At the prompt, type `config` and do the following:
   i. Press Enter to start the Configuration Utility.
   ii. On the Configure IP Address page, use the Tab key to select No and then press Enter.
   iii. On the IP Address page, type the IP address for management and then select OK.
   iv. On the Subnet mask page, type the appropriate subnet mask for the management network and the select OK.
   v. On the Management Route page, if you do not need to supply a gateway (default route), select No and then continue with the next step.
   vi. If you need to supply a gateway, select Yes, and then enter the gateway address for the management network on the next page. Select OK to continue.
   vii. On the Confirm Configuration page, review the settings and select Yes to continue.
o. At the prompt, type exit and then press Enter. You return to the login page. You may now close the console session.

3. Return to step 1 and repeat this entire process for the second BIG-IP system.

Performing the initial BIG-IP system configuration

In this section, we walk through using the Setup Utility wizard to deploy an internal, external, and high availability network. In our example, we name the first BIG-IP system, bigip1.mycompany.com, and the second BIG-IP system, bigip2.mycompany.com.

Configuring the first BIG-IP system

1. Open https://bigip1.mycompany.com using your preferred web browser. If you have not yet configured your DNS settings, use the management IP address for the BIG-IP system. Ignore any warnings about untrusted connections at this point.
2. For both the Username and Password fields, type `admin`. The BIG-IP Configuration utility opens to the Setup Utility page.
3. License the system using the following guidance:
    a. On the Welcome page, click Next. The License page opens.
    b. On the License page, click the Activate button to Activate your F5 license.
    c. In the Base Registration Key field, type (or copy/paste) the base registration key you received from F5.
    d. In the Activation Method row, make sure Automatic is selected.
    e. In the Outbound Interface row, make sure mgmt is selected.
    f. Click Next.
    g. Read the End User Software License and then click Accept.
    h. Wait while the system verifies the license and then click Log in.
4. Provision the F5 BIG-IP Access Policy Manager® (BIG-IP APM) and deprovision the F5 BIG-IP Local Traffic Manager™ (BIG-IP LTM) using the following guidance:
    a. From the Module table, find the Local Traffic row, and clear the box to deprovision BIG-IP LTM.
    b. In the Access Policy row, check the box to provision BIG-IP APM. Ensure Nominal is selected.
    c. Click the Next button.
    d. When the reprovisioning warning displays asking if you want to proceed, click OK. The system loads and verifies the new configuration.
5. Configure the Device certificates using the following guidance.
   Note that the Certificate subject needs to match the FQDN host name of the BIG-IP systems. Modify the request.inf file and request new certificates using **certreq**.
    a. On the Device Certificates page, click the Import button.
    b. From the Import type list, select PKCS 12 (IIS).
    c. From the Certificate Source row, click Choose File and then select the appropriate file.
    d. In the Password field, type the associated password.
    e. Click Import.
    f. Select the correct certificate subject and then click Next.
6. Configure the Platform options using the following guidance.
    a. In the Host Name field, type a host name. In our example, we use bigip1.mydomain.com.
    b. From the Time Zone list, select the correct time zone.
    c. In the Root Account row, type and confirm a password for the root account in the associated fields.
    d. In the Admin Account row, type and confirm a password for the admin account in the associated fields.
    e. Optional: If you want to restrict SSH access that are allowed to access this system to a specific range of IP addresses, from the SSH IP Allow list, select Specify Range and then type a range of IP addresses.
    f. Click Next.
    g. When the updated password warning displays, click OK. You are logged out. Log in again with your new credentials.
7. Configure the Network options using the following guidance:
    a. Under Standard Network Configuration, click the Next button.

b. In the Config Sync row, make sure the Display configuration synchronization options box is checked.
c. In the High Availability row, make sure the Display failover and mirroring options box is checked and the Failover Method is set to Network.
d. Click Next.
e. Complete the Internal Network Configuration using the following guidance:
    i. In the Netmask box, type the associated mask.
    ii. In the Floating IP row, in the Address box, type an IP address that is part of the trusted network set you during the BIG-IP OVF installation. This is the address that both units in a redundant system share, and must be different than the Self IP address.
f. Complete the Internal VLAN Configuration using the following guidance:
    i. In the VLAN Tag ID field, type an ID if the VLAN associated with the internal network is configured to use tagging.
    ii. In the VLAN interfaces row, from the Available list, select interface 1.1, and click the Add button to move it to either Untagged or Tagged.
g. Click Next.
h. Complete the External Network Configuration using the following guidance:
    i. In the Self IP row, in the Address box, type an IP address that is part of the untrusted network you configured during the BIG-IP OVF installation.
    ii. In the Netmask box, type the associated mask.
    iii. In the Default Gateway field, type the default gateway for the system.
    iv. In the Floating IP row, in the Address box, type an IP address that is part of the untrusted network you set during the BIG-IP OVF installation. This is the address that both units in a redundant system share, and must be different that the Self IP address.
i. Complete the External VLAN Configuration using the following guidance:
    i. In the VLAN Tag ID field, type an ID if the VLAN associated with the external network is configured to use tagging.
    ii. In the VLAN interfaces row, from the Available list, select interface 1.2, and click the Add button to move it to either Untagged or Tagged.
    iii. Click Next.
j. Complete the High Availability Network Configuration using the following guidance:
    i. In the Self IP row, in the Address box, type an IP address that is part of the trusted HA network you configured during the BIG-IP OVF installation.
    ii. In the Netmask box, type the associated mask.
k. Complete the High Availability VLAN Configuration using the following guidance:
    i. In the VLAN Tag ID field, type an ID if the VLAN associated with the HA network is configured to use tagging.
    ii. In the VLAN interfaces row, from the Available list, select interface 1.3, and click the Add button to move it to either Untagged or Tagged.
    iii. Click Next.
l. On the Config Sync Configuration page, from the Local Address list, select the IP address on the HA VLAN to use for configuration synchronization and then click Next.
m. Complete the Failover configuration using the following guidance:
    i. In the Failover Unicast Configuration section, click the Add button.
    ii. From the Address list, select the address on the internal VLAN, and then click Repeat.
    iii. From the Address list, select the address on the external VLAN, and then click Finish.
    iv. In the Failover Unicast Configuration table, check the box for the addresses on the Management Address and HA VLANs, and then click Delete. Click OK to confirm.

v. Note: The internal and external local addresses will be monitored for availability by the second BIG-IP system. If the second BIG-IP device is unable to reach either address, system automatically fails over.
vi. Click Next.

n. Complete the Mirroring configuration using the following guidance. Mirroring (or Connection mirroring) on the BIG-IP system is the mechanism by which connections on one system are essentially replicated on another system. Configuring mirroring helps ensure that in-process connections are not dropped when failover occurs. You enable mirroring on each relevant device. To set up mirroring, you specify, on the local device, the primary and secondary IP addresses that you want the system to use for mirroring. These are typically self IP addresses.
   i. From the Primary Local Mirror Address list, select the HA IP address.
   ii. Optional: From the Secondary Local Mirror Address list, select another local IP address to mirror connections.
   iii. Click Next.

o. Complete the Active/Standby Pair configuration using the following guidance:
   i. Under Standard Pair Configuration, click the Next button.
   ii. Under Configure Peer Device, click the Finished button.

This completes the configuration for the first BIG-IP system.

Configuring the second BIG-IP system

1. Open https://bigip2.mycompany.com using your preferred Web browser. If you have not yet configured your DNS settings, use the management IP address for the BIG-IP system. Ignore any warnings about untrusted connections at this point.
2. For both the Username and Password fields, type `admin`. The BIG-IP Configuration Utility opens to the Setup Utility page.
3. License the system using the following guidance:
   a. On the Welcome page, click Next. The License page opens.
   b. On the License page, click the Activate button to Activate your F5 license.
   c. In the Base Registration Key field, type (or copy/paste) the base registration key you received from F5.
   d. In the Activation Method row, make sure Automatic is selected.
   e. In the Outbound Interface row, make sure mgmt is selected.
   f. Click Next.
   g. Read the End User Software License and then click Accept.
   h. Wait while the system verifies the license and then click Log in.
4. Provision the Access Policy Manager (BIG-IP APM) and deprovision the Local Traffic Manager (BIG-IP LTM) using the following guidance:
   a. From the Module table, find the Local Traffic row, and clear the box to deprovision BIG-IP LTM.
   b. In the Access Policy row, check the box to provision BIG-IP APM. Ensure Nominal is selected.
   c. Click the Next button.
   d. When the reprovisioning warning displays asking if you want to proceed, click OK. The system loads and verifies the new configuration.
5. Configure the Device certificates using the following guidance.
   Note that the Certificate subject needs to match the FQDN host name of the BIG-IP system.
   a. On the Device Certificates page, click the Import button.
   b. From the Import type list, select PKCS 12 (IIS).
   c. From the Certificate Source row, click Choose File and then select the appropriate file.
   d. In the Password field, type the associated password.
   e. Click Import.
   f. Select the correct certificate subject and then click Next.

6. Configure the Platform options using the following guidance:
    a. In the Host Name field, type a host name. In our example, we use bigip2.mydomain.com.
    b. From the Time Zone list, select the correct time zone.
    c. In the Root Account row, type and confirm a password for the root account in the associated fields.
    d. In the Admin Account row, type and confirm a password for the admin account in the associated fields.
    e. Optional: If you want to restrict SSH access that are allowed to access this system to a specific range of IP addresses, from the SSH IP Allow list, select Specify Range and then type a range of IP addresses.
    f. Click Next.
    g. When the updated password warning displays, click OK. You are logged out. Log in again with your new credentials.
7. Configure the Network options using the following guidance:
    a. Under Standard Network Configuration, click the Next button.
    b. In the Config Sync row, make sure the Display configuration synchronization options box is checked.
    c. In the High Availability row, make sure the Display failover and mirroring options box is checked and the Failover Method is set to Network.
    d. Click Next.
    e. Complete the Internal Network Configuration using the following guidance:
        i. In the Self IP row, in the Address box, type an IP address that is part of the trusted network you configured during the BIG-IP OVF installation.
        ii. In the Netmask box, type the associated mask.
        iii. In the Floating IP row, in the Address box, type the same IP address you used on the first BIG-IP system.
    f. Complete the Internal VLAN Configuration using the following guidance:
        i. In the VLAN Tag ID field, type an ID if the VLAN associated with the internal network is configured to use tagging.
        ii. In the VLAN interfaces row, from the Available list, select interface 1.1, and click the Add button to move it to either Untagged or Tagged.
    g. Click Next.
    h. Complete the External Network Configuration using the following guidance:
        i. In the Self IP row, in the Address box, type an IP address that is part of the untrusted network you configured during the BIG-IP OVF installation.
        ii. In the Netmask box, type the associated mask.
        iii. In the Default Gateway field, type the default gateway for the system.
        iv. In the Floating IP row, in the Address box, type the same IP address you used on the first BIG-IP system.
    i. Complete the External VLAN Configuration using the following guidance:
        i. In the VLAN Tag ID field, type an ID if the VLAN associated with the external network is configured to use tagging.
        ii. In the VLAN interfaces row, from the Available list, select interface 1.2, and click the Add button to move it to either Untagged or Tagged.
        iii. Click Next.
    j. Complete the High Availability Network Configuration using the following guidance:
        i. In the Self IP row, in the Address box, type an IP address that is part of the trusted HA network you configured during the BIG-IP OVF installation.
        ii. In the Netmask box, type the associated mask.
    k. Complete the High Availability VLAN Configuration using the following guidance:
        i. In the VLAN Tag ID field, type an ID if the VLAN associated with the HA network is configured to use tagging.

ii. In the VLAN interfaces row, from the Available list, select interface 1.3, and click the Add button to move it to either Untagged or Tagged.
iii. Click Next.

l. On the Config Sync Configuration page, from the Local Address list, select the IP address on the HA VLAN to use for configuration synchronization and then click Next.

m. Complete the Failover configuration using the following guidance:
i. In the Failover Unicast Configuration section, click the Add button.
ii. From the Address list, select the address on the internal VLAN, and then click Repeat.
iii. From the Address list, select the address on the external VLAN, and then click Finish.
iv. In the Failover Unicast Configuration table, check the box for the addresses on the Management Address and HA VLANs, and then click Delete. Click OK to confirm.
v. Note: The internal and external local addresses will be monitored for availability by the first BIG-IP system. If the first BIG-IP device is unable to reach either address, system automatically fails over.
vi. Click Next.

n. Complete the Mirroring configuration using the following guidance.
Mirroring (or Connection mirroring) on the BIG-IP system is the mechanism by which connections on one system are essentially replicated on another system. Configuring mirroring helps ensure that in-process connections are not dropped when failover occurs. You enable mirroring on each relevant device. To set up mirroring, you specify, on the local device, the primary and secondary IP addresses that you want the system to use for mirroring. These are typically self IP addresses.
i. From the Primary Local Mirror Address list, select the HA IP address.
ii. Optional: From the Secondary Local Mirror Address list, select another local IP address to mirror connections.
iii. Click Next.

o. Complete the Active/Standby Pair configuration using the following guidance:
i. Under Standard Pair Configuration, click the Next button.
ii. Under Discover Configured Peer Device, click the Next button.
iii. Under Remote Device Credentials, specify the Management IP address, Administrator Username, and Administrator Password for the first BIG-IP system.
iv. Click Retrieve Device Information.
v. Click Finished after the peer device information has been retrieved.
vi. Click Awaiting Initial Sync.
vii. Select bigip2.mycompany.com (Self).
viii. Verify Sync Device to Group is selected.
ix. Select Sync. Sync status should now read In Sync

p. If a reboot is required:
i. If you have a message indicating the system needs to reboot for some changes to take effect do the following:
ii. Right-click the BIG-IP virtual machine, select Power, and then click Restart Guest. Alternatively, you can reboot the system using the BIG-IP GUI by going to System → Configuration → Device → and then clicking Reboot.

Installing and configuring the iApps template

Use the following procedure to import the SSL certificate and key, as well as the iApps template, onto the BIG-IP system, and then configure the iApps template for Horizon.

1. Log in to the active BIG-IP system and import the Horizon server SSL certificate and key you created in Obtaining and installing the SSL certificate for your View environment. The certificate and key must be in a location accessible by the BIG-IP system.

a. On the Main tab, click System > File Management → SSL Certificate List → Import.
b. From the Import Type list, select PKCS 12 (IIS).
c. In the Certificate Name field, type a name. In our example, we use remote.
d. In the Certificate Source row, click Choose File and browse to the PCKS file you created.
e. In the Password field, type the password you assigned.
f. Click Import. The system imports the certificate and key.

2. Import the View optimized iApps template onto the BIG-IP system.
   a. On the Main tab, click iApps → Templates.
   b. Click the Import button on the right.
   c. Click the Choose File button and then browse to the location of the iApps template.
   d. Click Upload.
3. Create a new iApps Application Service using the following guidance
   a. On the Main tab, click iApps → Application Services.
   b. Click the Create button.
   c. In the Name field, type a name for this application service. In our example, we use `view`
   d. From the Template list, select the f5.vmware_view_optimized_solution.vx.x.x template, where the x.x.x corresponds to the version of the iApps template, such as v1.1.0.
   e. Read the welcome information and prerequisites.
4. Configuring the iApps template for Horizon.
   Use the following guidance to help configure the iApps template. Each sub-step corresponds to a section of iApps.
   a. Configure the Template Options section:
      i. Inline Help: We recommend leaving inline help set to Show inline help text unless you are already very familiar with this iApps template.
      ii. DNS Servers: Specify the IP address of at least one DNS server unless you have already configured DNS services on the BIG-IP system outside of the iApps template. Some services, such as authentication, require server-to-IP resolution or IP-to-host name lookups to work properly. The BIG-IP system uses the list of servers you enter to resolve hosts to IP addresses or reverse lookup IP to hosts when required.
      iii. NTP Servers: Specify the IP address of at least one NTP server unless you have already configured NTP on the BIG-IP system outside of the iApps template. You can specify the FQDN of a NTP server farm, or individual NTP server IP address(es).
   b. Configure the BIG-IP Access Policy Manager:
      i. NAT address: This is optional. If your remote clients use a network translated address to connect to the View environment, specify this IP address. Only enter a value if the remote IP address is translated on another device prior to communicating with the BIG-IP system.
      ii. NetBIOS domain: Enter the NetBIOS domain(s) used for this Horizon installation. In our domain example (view.mycompany.com), we use view.
      iii. AAA Server object: The AAA Server performs user name look ups against Active Directory. If you already created an AAA server object outside of this iApps template, you select it from the list, otherwise, complete the following:
         - Active Directory Servers: Specify each of your Active Directory domain controllers, both FQDN and associated IP address, used for this View environment. Click the Add button for additional rows.
         - Active Directory domain name: Specify the fully qualified domain name (FQDN) used for this View environment.
         - Anonymous binding: Select whether anonymous binding is allowed in your Active Directory environment. If it is not allowed, specify an Active Directory user with administrative permissions in the fields that appear.

- Health monitor: Unless you have already created an Active Directory health monitor on this BIG-IP system, we recommend you allow the iApps template to create a new monitor. You can choose between no monitor, a simple ICMP ping monitor, or a more sophisticated Active Directory monitor. The Active Directory monitor requires a valid user account which it uses to log in to each server as part of the health check. If you choose the Active Directory monitor, we recommend creating a user account specifically for this health monitor that is set to never expire. Specify the user name and password in the appropriate fields, as well as the LDAP tree for the account. In the final monitor question, specify whether your Active Directory domain requires a secure protocol.

c. Configure SSL Encryption:
   i. SSL Offload: With SSL offload, traffic is encrypted to and from the client to the BIG-IP system. The BIG-IP system decrypts the traffic for processing, and then communicates with the View servers unencrypted. This offloads the task of processing SSL from the View servers, saving resources.
   ii. SSL key: Select the associated key you imported.

d. Configure the Virtual Servers and Pools:
   i. Public IP address: Specify the IP address your remote, untrusted clients will use to access the View servers.
   ii. Private IP address: Specify the IP address the local, trusted clients will use to access the View servers.
   iii. Port: Specify the associated service port, if different than 443.
   iv. FQDN: Specify the FQDN clients use to resolve to the remote and internal addresses. This solution uses a split DNS architecture where internal clients use local DNS, which points the FQDN to the local trusted virtual server address, and remote clients use external DNS, which points the same FQDN to the remote untrusted virtual server address.
   v. Server IP addresses: Specify the IP addresses of the View servers. In our example, we type the IP addresses for con1.view.mycompany.com and con2.view.mycompany.com.

e. Configure Client optimization:
   i. HTTP compression profile: Select whether you want the system to compress HTTP responses. If you want the system to use compression, you can specify a compression profile you have already created, or use the F5 recommended compression profile.

f. Configure Application Health:
   i. Health Monitor: Unless you have created a health monitor specifically for this implementation, leave the default: Create a new health monitor.
   ii. Interval: Specify the interval at which the BIG-IP system should monitor the View servers. We recommend the default of 30 seconds.

5. Final steps:
   a. Synchronize the configuration of the BIG-IP systems.
      i. In the upper left corner of the BIG-IP system, click Changes Pending.
      ii. Select the BIG-IP system with a status of Changes Pending.
      iii. Click Sync.
      iv. Verify the iApps configuration exists on the other BIG-IP system. Log in to the system, and on the Main tab, click iApps → Application Services. You should see the application service you created.

Installing and integrating Citrix NetScaler VPX with XenDesktop 7.1

- Virtual servers and virtual desktops
  - 2 Citrix Delivery Controller
  - 2 Citrix StoreFront Servers
  - 1 Citrix License Server
  - 1 MSSQL server
  - 100 Citrix XenDesktop virtual machines

- 2 NetScaler VPX
    - HA configuration requires two NetScaler systems.
    h. This guide sets up an active/passive HA configuration.

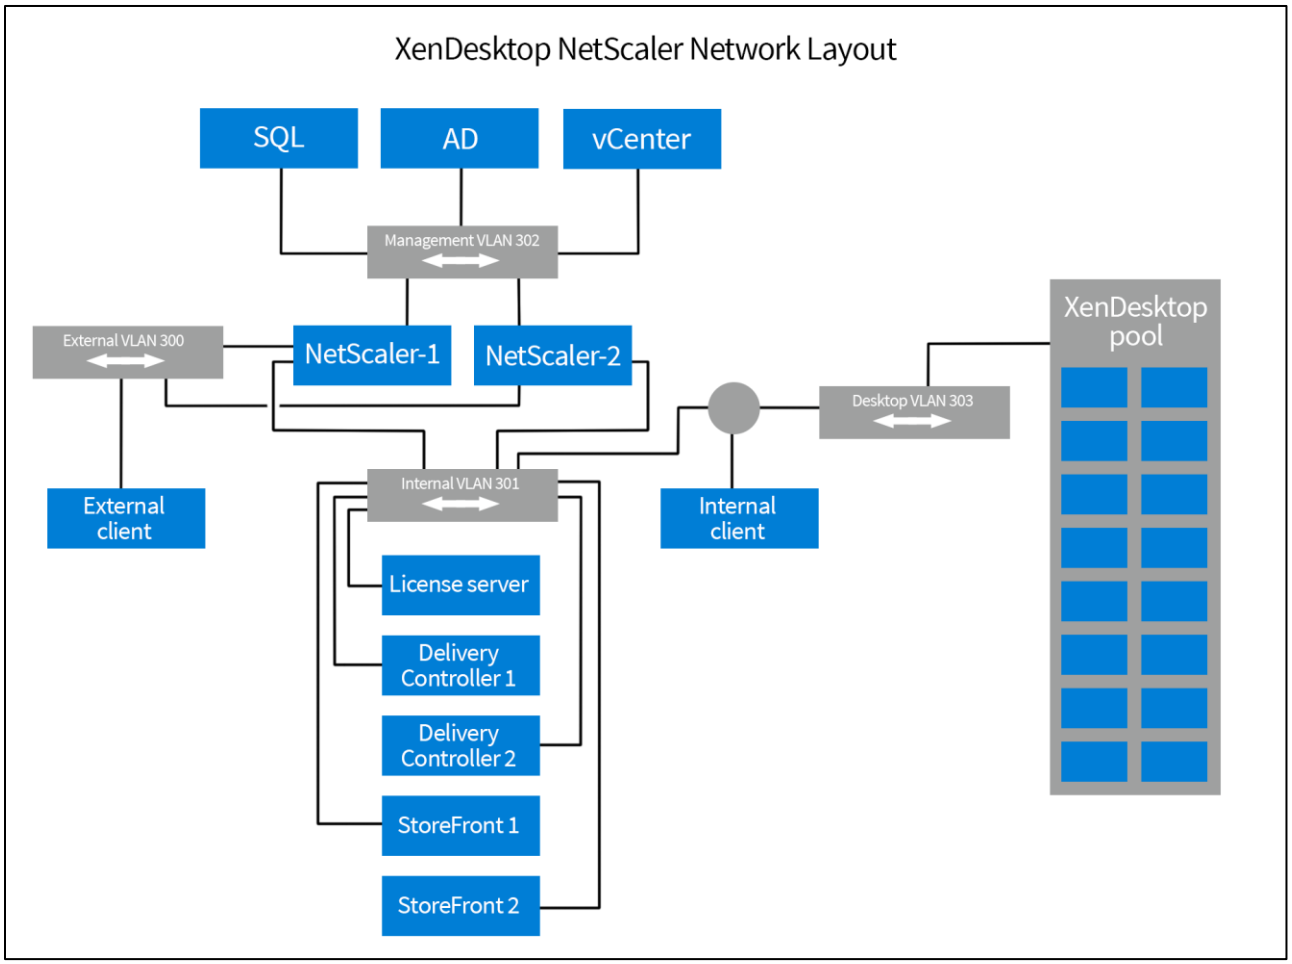

Figure 14: Citrix XenDesktop NetScaler network layout.

Delivery Controllers

- ddc1.xd.mycompany.com
- ddc2.xd.mycompany.com

StoreFront

- sf1.xd.mycompany.com
- sf2.xd.mycompany.com

License server

- lc1.xd.mycompany.com

MSSQL server

- sql1.xd.mycompany.com

Obtaining and installing ISS SSL certificate and configuring HTTPS access on StoreFront

1. Log in to the sf1.xd.mycompany.com.
2. Click Start → Administrative Tools → Internet Information Manager.
3. Double-click Server Certificate under IIS.

4. Under the Actions, select Create Certificate Request.
5. Provide the required information for the certificate:
    a. Common name
    b. Organization
    c. Organizational Unit
    d. City/locality
    e. State/providence
    f. Country/region
6. Click next.
7. Under Cryptographic service provider select:
    a. Microsoft RSA SChannel Cryptographic Provider
8. Under Bit length enter:
    a. 2048
9. Click Next.
10. Specify a file name and for the certificate request, and click Next.
11. Select a location, and click Next.
12. Open certificate request with Notepad and copy the contents.
13. Using the guidelines for your organization, submit the CSR to an internal or third party CA by pasting the copied contents when prompted for the CSR.
14. Copy the Base 64 encoded certificate that is returned onto sf1.xd.mycompany.com.
15. Go back to the ISS Server Certificate screen and click Complete Certificate Request.
    a. Browse to the file certification file.
    b. Enter a friendly name.
16. Click OK.
17. Verify that the new certificate has been imported.
18. Click Sites → Default Web Sites.
19. Click Bindings.
20. Click Add…
21. At the Add Site Binding change the type to HTTPS and select the appropriate SSL certificate.
22. Click OK.
23. Repeat steps 1 – 22 for on second storefront.

Configuring Citrix StoreFront

1. Log in to the sf1.xd.mycompany.com.
2. Click select Start → Citrix StoreFront.
3. Click Create a new deployment.
    a. Confirm the Base URL, and click Next.
    b. Enter a Store name, and click Next.
    c. Click add and enter StoreFront load balance vServer information:
        i. Enter a Display name.
        ii. Click Add… and enter the load balance VIP.
        iii. Click OK.
    d. Click Next.
    e. On the Remote Access screen, select no VPN tunnel.
    f. Under NetScaler Gateway appliances, click Add…
        i. Enter a Display name.
        ii. Enter a NetScaler Gateway URL.
        iii. Chose Version 10.0 (Build 69.4) or later.
        iv. Enter the Subnet IP address.

v. Select the appropriate logon type.
vi. Enter Callback URL.

g. Click Next.
h. Under Secure Ticket Authority URLs screen click Add…
i. Enter the two DDCs as Secure Ticket Authorities (STAs).
i. Click OK.
j. Click Create twice.
k. Click Finish.

Configuring StoreFront Authentication Methods

1. Open StoreFront and Select Authentication.
2. Under Actions select Authentication → Add/Remove Methods.
   a. Choose the authentication methods applicable, and click OK.

Adding StoreFront to Server Group

1. Log in to sf1.xd.mycompany.com.
2. Click Start → Citrix StoreFront.
3. Select Server Group.
4. Under Actions, select Add Server.
5. Copy the Authorizing server and Authorization code.
6. Log in to sf2.xd.mycompany.com.
7. Click Start → Citrix StoreFront.
8. Select Join existing server Group.
9. At the join Server group screen enter:
   a. Authorizing sever
   b. Authorization code
10. Wait for server to synchronize and confirm availability.
11. Select Authentication.
12. Under User Name and password, select Configure Trusted Domains.
    a. Select Trusted domain only.
    b. Click Add… and enter domain information.
    c. Select the Default domain, and click OK.
13. Under User Name and passwords, select Manage Password options.
    a. Select the appropriate password option, and click OK.

Configuring ISS default site redirection

1. Log in into sf1.xd.mycompany.com.
2. Click Start → Administrative tools → Internet information Services (ISS) Manager.
3. Under connections, select Default Web Site.
4. Under ISS, double-click HTTP Redirect.
   a. Select Redirect request to this destination and enter a valid URL.
   b. Under Redirect Behavior select:
      i. Redirect all request to exact destination (instead of relative destination)
      ii. Only redirect requests to content in this directory (not subdirectories)
      iii. Status code Found (302)
5. Click Apply.

Deploying the NetScaler OVF on ESXi 5.5

1. Launch VMware vCenter and log in.
2. From the File menu, click Deploy OVF Template. The Deploy OVF Template wizard opens. Complete the following.

3. Click Browse and go to the location of the OVF. Click Open, and then back on the Source Page, click Next.
4. On the OVF Template Details page, review the details and then click Next.
5. From the Name and Location page, in the Name box, type a unique name for this VM, and then select an inventory location. Click Next.
6. Select the storage location, and click Next.
7. Select Thick Provision Lazy Zeroed, and click Next.
8. On the Network Mapping page, select the appropriate Destination Networks for each of the Source.
9. On the Ready to Complete page, review the deployment settings. Use the Back button to make any changes. If the settings are correct, check Power on after deployment and then click Finish.

Repeat all steps 1 – 2 to deploy a second NetScaler.

Configuring management IP on NetScaler

1. Click the summary tab of the new virtual machine, and then click Open Console.
   a. Enter NetScaler's IPv4 address.
   b. Enter Netmask.
   c. Enter Gateway IPv4 address.
2. Enter 4 to Save and quit.
3. After rebooting, open a browser and navigate to the management interface IP.
4. Log in using nsroot and the username and password.
5. Repeat steps 1 – 4 on the second NetScaler.

Licensing NetScaler

1. At the welcome screen enter:
   a. Subnet IP Address
   b. Netmask
   c. Hostname
   d. DNS (IP address)
   e. Time Zone
2. Click Continue.
3. Under Update Licenses, click Browse.
   a. Select the license file, and click Open.
4. Click Continue.
5. Click Done.
6. Click Confirm to reboot the server and save the changes made.
7. Repeat steps 1 – 7 on second NetScaler.

Creating a RSA key

1. Log back in to the NetScaler.
2. Click Traffic Management → SSL.
3. Select Create RSA Key and provide the following:
   a. Provide a Key Filename.
   b. Enter 2048 as the bit Key Size.
   c. F4 as the Public Exponent Value.
   d. DES3 as the PEM encoding Algorithm.
   e. Enter a PEM Passphrase.
   f. Confirm the PEM Passphrase.
4. Click OK.

Creating a Certificate Signing Request

1. Click Traffic Management → SSL.
2. Select Create CRS (Certificate Signing Request) and enter the following:

a. Request File Name
b. Under Key Filename, click Browse and select the key file previously created.
c. Use PEM as the Key Format.
d. Enter the passphrase for the encrypted key previously created.
e. Country
f. State or Providence
g. Organization name
h. City
i. Email Address
j. Organization Unit
k. Common Name
l. Challenge password

3. Click OK.

Downloading the Certificate Signing Request and uploading a Certificate

1. Click Traffic Management → SSL.
2. Click Manage Certificates / Keys / CSRs.
   a. Select the CSR and click Download.
   b. Select a location a save the file, and click Download.
   c. Click Close twice.
   d. Open the request with Notepad and copy the contents.
3. Using the guidelines for your organization, submit the CSR to an internal or third party CA by pasting the copied contents when prompted for the CSR.
4. Download the certificate.
5. Click Manage Certificates / Keys / CSRs.
   a. Select the CSR and click upload.
   b. Upload the new certificate, and click Select.
   c. Click Close.
6. Click Traffic Management → SSL → Certificates.
7. Click Install…
8. At the Install Certificate provide the following:
   a. Certificate-Key Pair Name.
   b. Click Browse on the Certificate File Name and select the file previously uploaded.
   c. Click Open.
   d. Click Browse on the Key File Name and select the previously created Key File.
   e. Click Open.
   f. Select PEM as the Certificate Format.
   g. Enter the Password previously selected.
9. Click Create.
10. Click Close.
11. Click Certificates and select the new certificate.
12. Under the Action dropdown select Link.
13. Link the certificate with the appropriate CA Certificate Name.

Load balancing StoreFront

1. Click Traffic Management → Load Balancing → Load Balancing wizard.
2. At the introduction screen, click Next.
3. At the Create Services screen provide the following:
   a. Service Name

b. Click New…
    i. At the Create Server screen enter a Server Name.
    ii. Enter details for sf1.xd.mycompany.com.
    iii. Click Create.
4. Select SSL as the protocol.
5. Click Add.
6. At the Create Services screen provide the following:
    a. Service Name
    b. Click New…
        i. At the Create Server screen, enter a Server Name.
        ii. Enter details for sf2.xd.mycompany.com.
        iii. Click Create.
7. Select SSL as the protocol.
8. Click Add.
9. Click Next.
10. At the Create Virtual Server page enter the following:
    a. Name of the Storefront Virtual Server
    b. IP of the Virtual server
    c. SSL as the protocol
11. Under Available Services select both services, and click Add to move them to Configured Services.
12. Click Next.
13. Click Finish.
14. Click Exit.
15. Click Traffic Management → Load Balancing → Virtual Servers.
16. Double-click the recently created virtual server.
    a. Click the Method and Persistence tab.
        i. Under LB Method, select Last Connection.
        ii. Under Persistence, select COOKIEINSERT.
    b. Click OK.
17. Click the Save the running configuration bottom on the top right corner.
18. Select Yes to confirm.
19. Click Traffic Management → Load Balancing → Virtual Servers.
20. Double-click the recently created virtual server.
21. Double click the first Service Name.
    a. Select SSL Setting.
    b. Select the appropriate certificate, and click Add.
    c. Click OK.
22. Double click the second Service Name.
    a. Select SSL Setting.
    b. Select the appropriate certificate, and click Add.
    c. Click OK.

Load balancing Delivery Controllers

1. Click Traffic Management → Load Balancing → Load Balancing wizard.
2. At the Introduction screen, click Next.
3. At the Create Services screen provide the following:
    a. Service Name.
    b. Click New…
        i. At the Create Server screen enter a Server Name.

ii. Enter details for ddc1.xd.mycompany.com.
iii. Click Create.
4. Select TCP as the protocol and select a port.
5. Click Add.
6. At the Create Services screen provide the following:
   a. Service Name
   b. Click New…
      i. At the Create Server screen enter a Server Name.
      ii. Enter details for ddc2.xd.mycompany.com.
      iii. Click Create.
7. Select TCP as the protocol and select a port.
8. Click Add.
9. Click Next.
10. At the Create Virtual Server page enter the following:
    a. Name of the DDC XML Virtual Server
    b. IP of the Virtual server
    c. TCP as the protocol and select a port
11. Under Available Services select both services, and click Add to move them to Configured Services.
12. Click Next.
13. Click Finish.
14. Click Exit.
15. Click the refresh bottom on the top right corner.
16. Click Traffic Management → Load Balancing → Virtual Servers.
17. Double-click the recently created virtual server.
18. Double click the first Service Name.
    a. Click the Monitors tab and select tcp.
    b. Click Add.
    c. Click OK.
19. Double click the second Service Name.
    a. Click the Monitors tab and select tcp.
    b. Click Add.
    c. Click OK.

Creating the Access Gateway Virtual Server

1. Click Configuration→NetScaler Gateway→NetScaler Gateway wizard.
2. Click Next at the Introduction screen.
   a. At the Create or choose a virtual server screen, select New and enter the appropriate IP address, port and name.
3. Click Next.
   a. At the Specify a server certificate screen, select Used an installed certificate and private key pair.
   b. Under Server Certificate, select the previously created key pair.
4. Click Next.
   a. At the Configure Name Service screen, enter the DNS server IP and choose DNS as the Name Lookup Priority.
5. Click Next.
   a. At the Configure authentication screen, select LDAP.
   b. Under server, enter the DNS server IP.
   c. Under Connection Settings enter the following:
      i. Base DN
      ii. Administrator Bind DN

iii. Administrator Password
iv. Confirm the Administrator password

d. Click Retrieve Attributes.
e. Click OK when the LDAP attributes have been retrieved successfully.

6. Click Next.
   a. At the Configure additional settings screen, select Allow.
   b. Select Redirect to secure Web address and type the appropriate URL.
7. Click Next.
8. At the Configure Clientless Access screen, click Next. Click Finish.
9. Go to Configuration → Load Balancing → Virtual Servers and double-click the http_redirect server.
   a. Click the advanced tab and change the Redirect URL to the HTTPS external accessible address.
10. Click OK.
11. Go to Configuration → NetScaler Gateway → Authentication → LDAP.
12. Under Policies click Add…
    a. Enter a policy name.
    b. Click New and complete the following:
       i. Server name
       ii. Server IP
       iii. Base DN
       iv. Administrator Bind DN
       v. Administrator Password
       vi. Confirm the Administrator password
    c. Click Create.
    d. Under Named Expressions select General and True Value.
    e. Click Add Expression.
    f. Click Create.
    g. Click Close.
13. Go to Configuration → NetScaler Gateway → Virtual Servers.
14. Double-click the Gateway virtual server.
15. Select the Authentication tab.
16. Select Insert Policy.
17. Select the previously created policy.
18. Click OK.
19. Go to Configuration → NetScaler Gateway → Virtual Servers.
20. Double-click the virtual server.
21. Click the Published Applications tab.
22. Under Secure Ticket Authority click Add.
    a. Enter `https://` followed by the IP of DDC1.
    b. Click Create.
23. Under Secure Ticket Authority click Add.
    a. Enter `https://` followed by the IP of DDC2.
    b. Click Create.
24. Click the Policies tab and select Insert Policy.
25. Select New Policy.
    a. Click New.
    b. Enter name.
    c. Click New.
    d. Enter a name.

e. Click the Published Applications tab.
f. Override the following settings and enter appropriate values:
    i. Override ICA Proxy
    ii. Override Web Interface
    iii. Override Single Sign-on Domain
g. Click the Security Tab.
    i. Override the Default Authorization Action.
h. Click Create.
i. Select the True value Expression and click Create.

Adding a node and Synchronizing NetScaler

1. Go to Configuration → High Availability.
2. Click Add…
    a. Enter the Remote Node IP Address.
    b. Select Configure remote system to participate in High Availability setup.
    c. Select Turn off HA Monitor on interfaces/channels that are down.
    d. Enter System credentials.
3. Click OK.
4. Double click the configure node.
    a. Select Secondary node will fetch the configuration from Primary.
    b. Select Primary will propagate configuration to Secondary.
    c. Select Maintain one primary node even when both nodes are unhealthy.
    d. Set the communication intervals if necessary.
5. Click OK.
6. Go to Configuration → High Availability → Nodes tab.
7. Click Force Synchronization.

ABOUT PRINCIPLED TECHNOLOGIES

Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.