

ARX-VE Installation Guide

810-0064-00



Publication Date

This manual was published on February 12, 2013.

Legal Notices

Copyright

Copyright 2011 - 2012, F5 Networks, Inc. All rights reserved.

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

Trademarks

Access Policy Manager, Advanced Client Authentication, Advanced Routing, APM, Application Security Manager, ARX, AskF5, ASM, BIG-IP, BIG-IQ, Cloud Extender, CloudFucious, Cloud Manager, Clustered Multiprocessing, CMP, COHESION, Data Manager, DevCentral, DevCentral [DESIGN], DNS Express, DSC, DSI, Edge Client, Edge Gateway, Edge Portal, ELEVATE, EM, Enterprise Manager, ENGAGE, F5, F5 [DESIGN], F5 Management Pack, F5 Networks, F5 World, Fast Application Proxy, Fast Cache, FirePass, Global Traffic Manager, GTM, GUARDIAN, IBR, Intelligent Browser Referencing, Intelligent Compression, IPv6 Gateway, iApps, iControl, iHealth, iQuery, iRules, iRules OnDemand, iSession, L7 Rate Shaping, LC, Link Controller, Local Traffic Manager, LTM, Message Security Manager, MSM, OneConnect, OpenBloX, OpenBloX [DESIGN], Packet Velocity, Policy Enforcement Manager, PEM, Protocol Security Manager, PSM, Real Traffic Policy Builder, Rosetta Diameter Gateway, ScaleN, Signaling Delivery Controller, SDC, SSL Acceleration, StrongBox, SuperVIP, SYN Check, TCP Express, TDR, TMOS, Traffic Management Operating System, Traffix Diameter Load Balancer, Traffix Systems, Traffix Systems (DESIGN), Transparent Data Reduction, UNITY, VAULT, VIPRION, vCMP, virtual Clustered Multiprocessing, WA, WAN Optimization Manager, WebAccelerator, WOM, and ZoneRunner, are trademarks or service marks of F5 Networks, Inc., in the U.S. and other countries, and may not be used without F5's express written consent.

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

Patents

This product may be protected by U.S. Patents 7,877,511; 7,958,347. This list is believed to be current as of February 12, 2013.

Export Regulation Notice

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

RF Interference Warning

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

FCC Compliance

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

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

Canadian Regulatory Compliance

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

Standards Compliance

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

Acknowledgments

This product includes software from several third-party vendors. Each vendor is listed below with the applicable copyright.

Copyright (c) 1990, 1993, 1994, 1995 The Regents of the University of California. All rights reserved.

Copyright 2000 by the Massachusetts Institute of Technology. All Rights Reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

Copyright 1993 by OpenVision Technologies, Inc.

Copyright (C) 1998 by the FundsXpress, INC.

All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

Copyright (c) 1995-2001 International Business Machines Corporation and others

All rights reserved.

Copyright (c) 1990-2003 Sleepycat Software. All rights reserved.

Copyright (c) 1995, 1996 The President and Fellows of Harvard University. All rights reserved.

Copyright (c) 1998-2004 The OpenSSL Project. All rights reserved.

Unless otherwise noted, the companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted in examples herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

Revision History

January 2011 — Rev A.

February 2011 — Rev B.

June 2011 — Rev C.

September 2011 — Rev D.

October 2012 — Rev E.



Table of Contents

I	
Introduction	
	Audience for this Manual I-3
	Document Conventions I-3
	Related Documents I-3
	Contacting Customer Service I-4
2	
Product Overview	
	ARX-VE Overview 2-3
	Configuring Redundant Pairs 2-3
	Connecting Multiple Switches 2-3
	Managing the Switch 2-4
3	
Installing the ARX-VE	
	Tools and Equipment 3-3
	Hypervisor Requirements 3-3
	Installing the ARX-VE 3-3
	Sample: Installing the ARX-VE onto an ESX 4.0 Hypervisor 3-4
	Sample: Installing the ARX-VE onto an ESX 4.1 Hypervisor 3-8
	Booting the Switch 3-13
	Connecting the Ethernet Management Interface 3-20
	Powering Down the ARX-VE 3-21
	Converting From the Trial Version to the Production Version 3-21
A	
Best Practices For ARX-VE	
	Best Practices For ARX-VE A-3
Index	



Introduction

This manual describes the installation process for the F5 Adaptive Resource Switch Virtual Edition, or ARX-VE. This is a software product, downloaded from a web site and installed as a guest on a hypervisor. This manual describes the requirements for the hypervisor, the installation process, and the process of connecting the ARX-VE to your network.

- [Audience for this Manual](#)
- [Document Conventions](#)
- [Related Documents](#)
- [Contacting Customer Service](#)

Audience for this Manual

This manual is intended for field engineers and network administrators responsible for installing the ARX and adding it to the network.

Document Conventions

This manual uses the following conventions, when applicable:

- `console` text represents system output
- **bold** text represents user input
- *italic* text appears for emphasis, new terms, and book titles

◆ Note

Notes provide additional or helpful information about the subject text.

◆ Important

Important notices show how to avoid possible service outage or data loss.

◆ WARNING

Warnings are instructions for avoiding damage to the equipment.

Danger notices help you to avoid personal injury.

Related Documents

In addition to this guide, the following F5 Data Solutions documentation is available:

- *ARX-VE Quick Installation*
- *ARX CLI Reference*
- *ARX CLI Network-Management Guide*
- *ARX CLI Storage-Management Guide*
- *ARX CLI Maintenance Guide*

Contacting Customer Service

You can use the following methods to contact F5 Networks Customer Service:

F5 Networks Online Knowledge Base Online repository of answers to frequently-asked questions.	http://support.f5.com
F5 Networks Services Support Online Online customer support request system	https://websupport.f5.com
Telephone	Follow this link for a list of Support numbers: http://www.f5.com/support/support-services/contact/



2

Product Overview

- [ARX-VE Overview](#)
- [Configuring Redundant Pairs](#)
- [Connecting Multiple Switches](#)
- [Managing the Switch](#)

ARX-VE Overview

The Virtual Edition of the F5 Adaptive Resource Switch, called the ARX-VE, enables enterprises to globally access, manage, deliver and optimize information resources. The ARX-VE is a cost-effective Adaptive Resource Switch designed for use in small data centers and branch/remote offices. The ARX-VE demonstrates application processing and control in a portable Virtual Machine (VM). It offers the same software features as the other ARX platforms, differing only in performance and scale.

Specifically, the ARX-VE provides the following features:

- Access. Simplified, flexible, location-independent access to enterprise-wide data.
- Data protection. Failover of large file systems, centralized backup, and scaled backup performance.
- Data migration. Seamless, transparent data migration across heterogeneous NAS.
- Dynamically tiered storage.
- Storage aggregation. Aggregation of multiple shares into a single client volume.
- Capacity management. Inline management of storage capacity to adapt the storage to client demands.
- Seamless storage addition/removal. Ability to add or remove storage without any effect on clients.

The ARX-VE supports Gigabit Ethernet throughput for connectivity to network infrastructure, network-attached storage (NAS) devices, and file servers with direct-attached storage (DAS).

Configuring Redundant Pairs

You can use standard VM failover mechanisms with the ARX-VE. Assign the VM to a hypervisor cluster. If the currently-running hypervisor fails, the ARX-VE resumes processing on one of the hypervisor's peers.

Connecting Multiple Switches

You can connect multiple ARX units to form a Resilient Overlay Network (RON), a series of IP tunnels between the switches. The switches in a RON can be of any platform type.

The RON provides a network for distributing and accessing file storage. An ARX can replicate storage to another ARX in the same RON, updating the replicas periodically as the writable master files change. This process is called *shadow copy*. With shadow copy, clients are granted read-only access to shadow target volumes at multiple geographic locations, independent of where the shadow source volume resides.

For information about configuring RON tunnels, see the *ARX CLI Network-Management Guide* and the *ARX CLI Reference*.

Managing the Switch

For local and remote management, the ARX-VE provides the following management interfaces:

- Serial console port for accessing and managing the switch through the VM client's console terminal. This accesses the CLI.
- The Inband Ethernet interface, which runs on the same virtual NIC (VNIC) used for the client/server networks.

For general information about using the CLI, see the *ARX CLI Network-Management Guide* and *ARX CLI Reference*.



3

Installing the ARX-VE

- Tools and Equipment
- Hypervisor Requirements
- Installing the ARX-VE
- Connecting the Ethernet Management Interface
- Powering Down the ARX-VE
- Converting From the Trial Version to the Production Version

Tools and Equipment

You need a hypervisor client, where you can install the ARX-VE as you would install a standard Virtual Machine (VM), and where you can access the ARX Console at the end of the installation.

Hypervisor Requirements

The hypervisor where you install ARX-VE must support OVF templates for its VM installations. VMware ESX and VMware ESXi support OVF templates, but desktop and end-user versions (such as VMware Player) do not.

Installing the ARX-VE

Obtain a base registration key for your ARX-VE instance from the F5 web site at the following URL:

<https://www.f5.com/trial/>

The ARX-VE is packaged in a standard OVF template. Download the OVF template file from the F5 web site at the following URL:

<https://downloads.f5.com/>

You require a valid username and password to access this URL; you can register with the web site to get these credentials. After you log in, find the OVF template in the above directory.

From your hypervisor client, install the ARX-VE using the above OVF template. The ARX-VE requires the following resources from the hypervisor:

- 2 CPU cores, 64-bit
- 4 GB of memory
- 1 Virtual NIC (VNIC) interface
- 40 GB or more of disk space

These are defined in the OVF template. Please contact F5 technical support prior to making any change to the settings in the OVF template.

Attach the VNIC to a vSwitch within the hypervisor, so that the VNIC is on the same (or otherwise routable) network as the ARX's clients and back-end filers. The VNIC should be untagged (no VLAN tags).

Sample: Installing the ARX-VE onto an ESX 4.0 Hypervisor

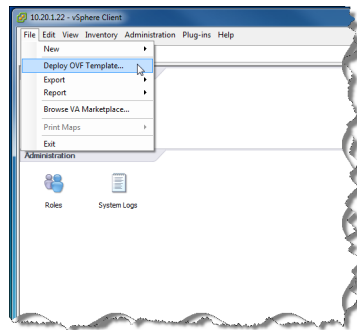
This is a sample installation of the ARX-VE as a guest VM on an ESX 4.0 hypervisor. In this example, we have downloaded the following file from the F5 web site:

arxve-esx-13103.ova

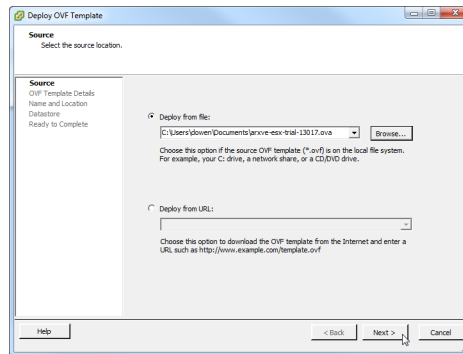
◆ Note

The configuration of your ESX server may vary from the one in this example. Your installation may therefore include variations of the sample screens below.

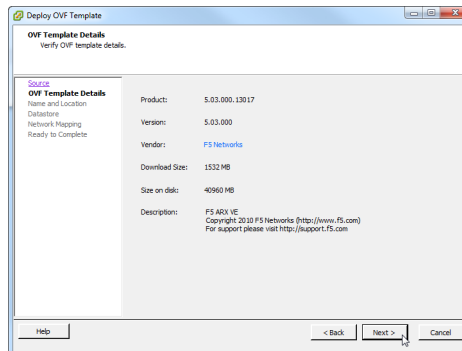
From a VSphere client, select File -> Deploy OVF Template...:



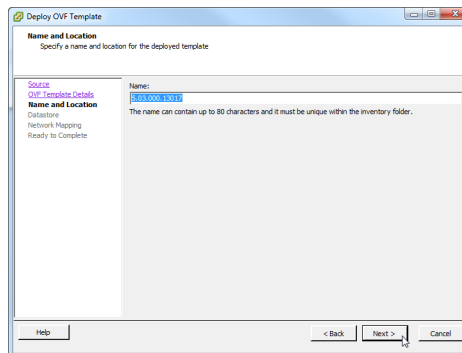
Type the path to the downloaded .ova file, or click **Browse...** to search for it. Then click **Next**:




View the details of the OVF template, and click **Next** to proceed:




Type a meaningful name for the OVF template:



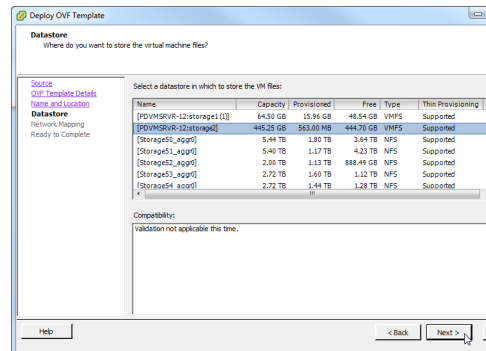
Select a deployment configuration from the Configuration dropdown menu; choose **Trial** or **Production** as appropriate for your installation, then click **Next**.



Select a host or cluster on which to run the deployed template:



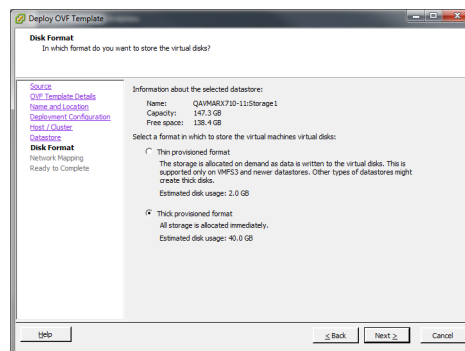
Select a datastore on which to store the VM files:



Note

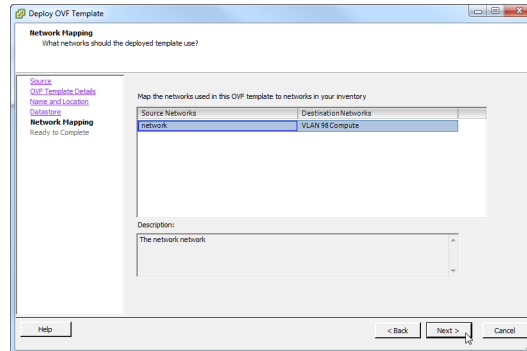
Network latency between the ARX-VE and its datastore should be kept to 250 microseconds or less to ensure optimal performance. After installation, use the CLI command `probe metalog latency` to ensure that latency is within the desired limits. If the latency is consistently greater than 400 microseconds, you should consider re-installing the ARX-VE and using a different datastore; 250 microseconds latency or less is recommended for best performance, if that is feasible.

Select a storage format for the virtual disks:

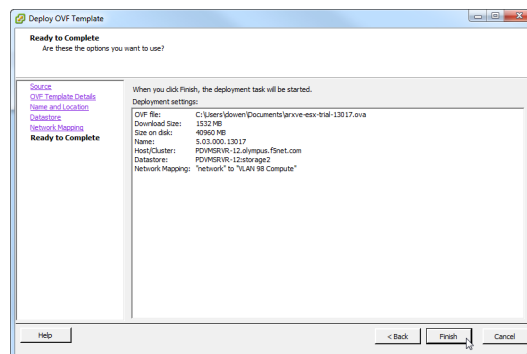


Associate the network used by the template with a network in your inventory.


If the ARX-VE's assigned VNIC is not already on the network, attach the VNIC to a vSwitch within the hypervisor. The VNIC should be on the same (or otherwise routable) network as the ARX's clients and back-end filers:




The next screen summarizes the configuration options you have chosen above. Click **Finish** if the information is correct:



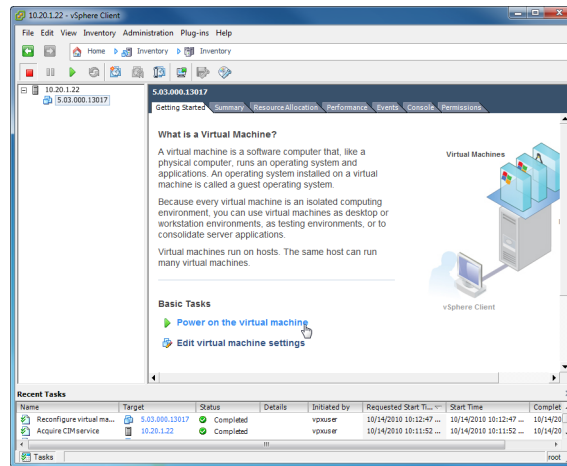
This invokes the deployment of the ARX-VE on your hypervisor. A small status window provides a progress meter:



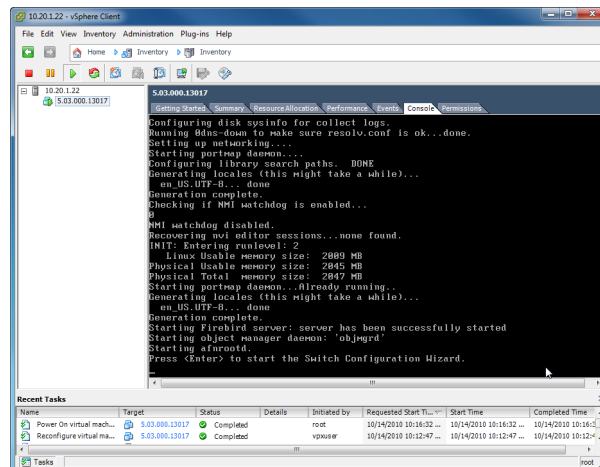
This final window indicates successful deployment:



From the same VSphere client, power up the new ARX-VE:



Access the Console window. Boot-up messages appear on the screen, followed by a prompt to start the initial-boot script:



The section after the next sample explains how to answer the questions in the initial-boot script.

Sample: Installing the ARX-VE onto an ESX 4.1 Hypervisor

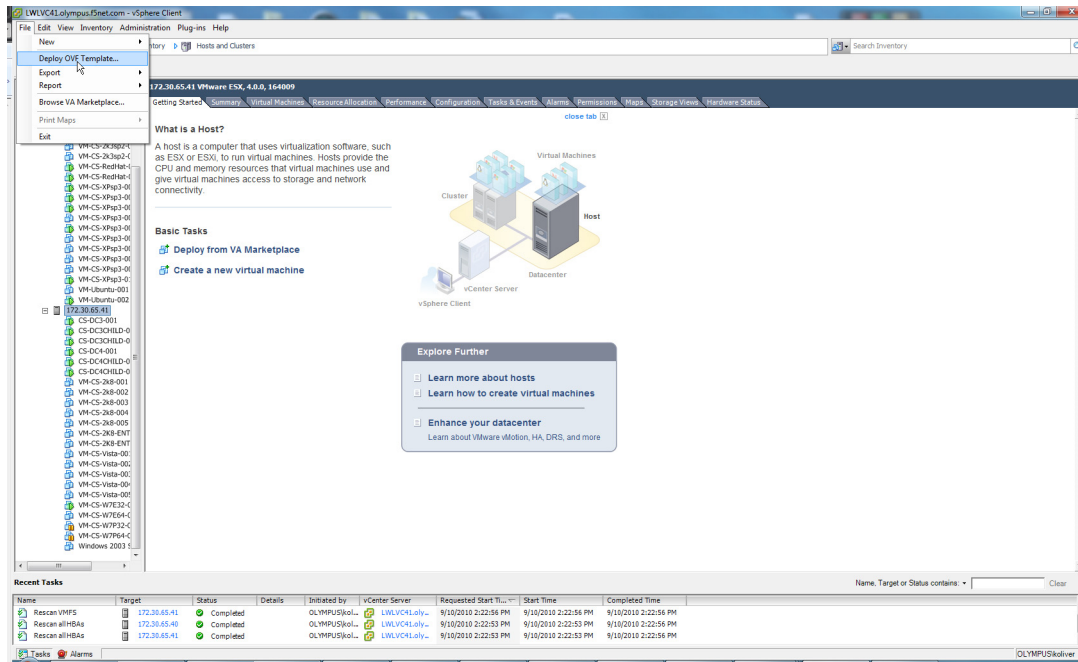
This is a sample installation of the ARX-VE as a guest VM on an ESX 4.1 hypervisor. In this example, we have downloaded the following file from the F5 web site:

arxve-esx-13103.ova

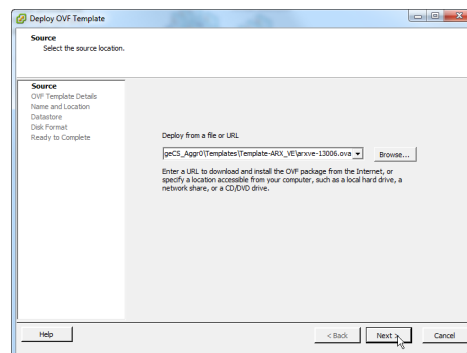
◆ Note

The configuration of your ESX server may vary from the one in this example. Your installation may therefore include variations of the sample screens below.

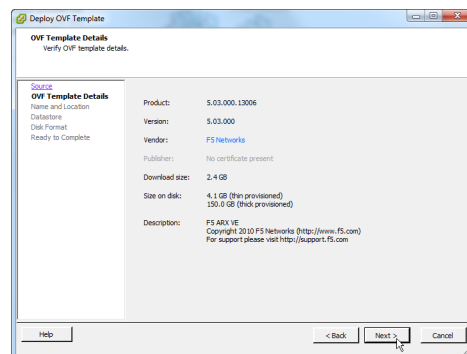
From a VSphere client, select File -> Deploy OVF Template...:




Type the path to the downloaded .ova file, or click Browse... to search for it. Then, click Next:



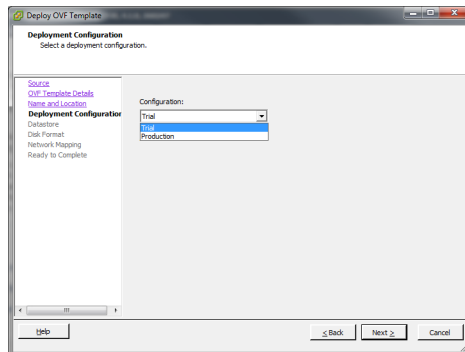
View the details of the OVF template, and click Next to proceed:




Type a meaningful name for the OVF template:




Select a deployment configuration from the Configuration dropdown menu; choose Trial or Production as appropriate for your installation, then click Next.



Select a host or cluster on which to run the deployed template:




Select a datastore on which to store the VM files:



Note

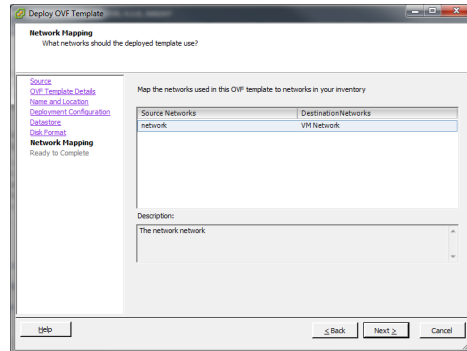
Network latency between the ARX-VE and its datastore should be kept to 250 microseconds or less to ensure optimal performance. After installation, use the CLI command `probe metalog latency` to ensure that latency is within the desired limits. If the latency is consistently greater than 400 microseconds, you should consider re-installing the ARX-VE and using a different datastore; 250 microseconds latency or less is recommended for best performance, if that is feasible.

Choose a storage format (“Thin” or “Thick”) for the virtual disks:

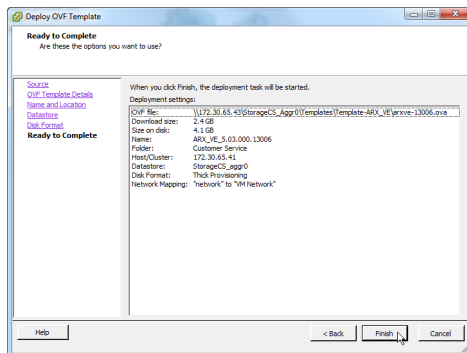


Associate the network used by the template with a network in your inventory:


If the ARX-VE's assigned VNIC is not already on the network, attach the VNIC to a vSwitch within the hypervisor. The VNIC should be on the same (or otherwise routable) network as the ARX's clients and back-end filers. The VNIC should also be untagged (no VLAN tags).




The next screen summarizes the configuration options you have chosen above. Click **Finish** if the information is correct:



This invokes the deployment of the ARX-VE on your hypervisor. A small status window provides a progress meter:



This final window indicates successful deployment:



From the same VSphere client, power up the new ARX-VE and access its Console window. Boot-up messages appear on the screen, followed by the initial-boot script. The next section explains how to answer the questions in this script.

Booting the Switch

The initial-boot script runs automatically at switch start-up, at the end of the ARX-VE installation. It prompts for basic configuration and security information required to access the switch and manage it remotely. It starts with the following prompt:

Press <Enter> to start the Switch Configuration Wizard.

Press the **Enter** key as prompted.

Several questions appear, prompting you for basic network information (such as management-IP address, mask, and gateway). These questions comprise the initial-boot script. Answer these questions as they come up. Examples and instructions appear in the following sections.

Sample: Booting a Switch After its Initial Installation

The following example shows the simplest initial-boot scenario — booting a new (non-replacement) switch.

The answers in the example are *not* appropriate to the following scenarios:

- Re-installing the ARX-VE to replace a defunct one
- Re-booting an in-service ARX-VE after F5 personnel performed a *Manufacturing Installation* on it (which returns the switch to its factory defaults)

Later sections discuss these contingencies and how to handle each of them. The answers below apply to the simplest case only — booting a new (non-replacement) switch. For many of these questions, the default is sufficient.

Sample answers are shown in bold text.

F5 ARX Startup

This F5 ARX does not currently have critical system information programmed. The following wizard prompts you for this information. You can connect to the switch through the management interface when you finish.

To restart the configuration program, enter 'r' at any prompt.

The switch's management port requires an IP address and mask.

1. Enter the management port IP address
in the format nnn.nnn.nnn.nnn or 'none'. # **192.168.66.62**
2. Enter the management port subnet mask
in the format nnn.nnn.nnn.nnn.(default=255.0.0.0) # **255.255.0.0**

The switch's management port requires a gateway IP address.

3. Enter the gateway IP address for the management interface
in the format nnn.nnn.nnn.nnn or 'none'. # **192.168.66.1**

A name server address must be assigned so that the software license can be activated.

4. Enter the DNS name server IP address to access the license server
in the format nnn.nnn.nnn.nnn. # **192.168.66.23**

The switch's internal subnet requires an IP address and mask.

Chapter 3

Installing the ARX-VE

5. Enter the switch's private IP address
in the format nnn.nnn.nnn.nnn.(default=169.254.13.0) #
6. Enter the switch's private subnet mask
in the format nnn.nnn.nnn.nnn.(default=255.255.255.0) #
7. Enter the switch's UUID
in the format
xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxx.(default=8fa98111-55ec-d1c8-9380-8dtu78fab47d) #
8fa98111-55ec-d1c8-9380-8dtu78fab47d

The base registration key is used to activate the software license for this system.

8. Enter the switch's base registration key
in the format xxxxxxx-xxxxxx-xxxxx-xxxx-xxxxxxx. # **CRJGV-QPDYW-SATNK-RGBYY-DMTMOBL**

The crypto-officer is the most privileged user in the system.

9. Enter the crypto-officer username
in the format text (1-28 characters). # **admin**
10. Enter the crypto-officer password
in the format text (6-28 characters). # *********
Confirm the crypto-officer password # *********

A system password is required for access to the master key.

11. Enter a system password
in the format text (12-28 characters). # *********
Confirm the system password # *********

The master key is used to encrypt critical security parameters.

12. Enter the master key
in the format base64-encoded key or keyword 'generate'.(default=generate) # **generate**

The system displays a configuration summary. See the following example.

```
Configuration Summary
Management IP Address    192.168.66.62
Management IP Mask      255.255.0.0
Management Gateway       192.168.66.1
DNS IP Address           192.168.66.23

Chassis GUID             8fa98111-55ec-d1c8-9380-8dtu78fab47d
Chassis Base Reg Key     CRJGV-QPDYW-SATNK-RGBYY-DMTMOBL
Switch Password          #####
Switch Master Key        generate
Crypto-officer Username  admin
Crypto-officer Password  #####
```

Enter 'yes' to load the configuration or 'r' to redo the interview #yes

You have completed the switch startup configuration.
The switch will now initialize the local database.
When the login prompt appears, log into the switch using
the crypto-officer's username and password.

Closing configuration file.
Processing configuration file. (boot-config)

Completed initializing the system.
...

The boot-up prompts continues to the **Username** prompt. Confirm that an administrator can log in by using the Crypto-Officer username and password that you entered in the initial-boot script, as in the following example.

...

User Access Authentication

```
Username: admin
Password: mypassword
SWITCH>
```

The switch is now ready for configuration through the CLI or GUI. For configuration instructions, see the *ARX GUI Quick Start: Network Setup* or the *ARX® CLI Network-Management Guide*.

Preparing for Switch Replacement (Re-Installation)

The process of replacing a defunct switch is more complicated than the initial-boot process for a new (non-replacement) switch. It starts the same way, with an installation of the (possibly, same) OVF template onto the hypervisor. The differences are in the initial-boot wizard that runs at the end of the guest-VM installation.

There are a few things you must have done prior to the switch failing. This includes saving your running and global configs, UUID, and master key and associated passwords as described in the *ARX Site Planning Guide, Best Practice: Regularly Saving the Configuration*, on page 1-68. If the global config of the failed switch contained any managed volumes, they must all re-import when the global-config is replayed on the replacement switch.

Matching the Private Subnet

After the questions about the management address and gateway, the next set of questions ask for the switch's *private subnet*. If the failed switch was in a redundant pair and/or Resilient-Overlay Network (RON), the private subnets of the replacement switch should match those of the switch that failed. Each ARX uses its private subnet for communication with other ARXes in the same RON and/or the switch's redundant peer. This private subnet must be reserved for ARX traffic only.

The private-subnet information appears at the top of the output of the **show running-config** command for a failed switch. The following example shows the top of a running-config file from a failed switch. The private-subnet information is in **bold** text.

```
; ARX-VE
; Version 5.03.000.13014 (Sep 24 2010 18:15:38) [nbuids]
; Database version: 503000.24
; Generated running-config Sun Sep 26 08:09:31 2010
; System UUID 8fa98111-55ec-d1c8-9380-8dtu78fab47d
; ip private subnet 169.254.170.0 255.255.255.0
;
terminal character-set unicode-utf-8

...
```

Entering the Private Subnet

Enter the private subnet and mask of the failed switch. See the following example.

```
...
The switch's internal subnet requires an IP address and mask.
4. Enter the switch's private IP address
   in the format nnn.nnn.nnn.nnn.(default=169.254.203.0) # 169.254.170.0
5. Enter the switch's private subnet mask
   in the format nnn.nnn.nnn.nnn.(default=255.255.255.0) # <Enter>
...
```

Finding the UUID of the Failed Switch

When a switch imports storage from file servers, it marks each share with its Universally-Unique ID (UUID). A replacement switch must use the same UUID or it rejects all of the shares imported by its predecessor. Also, you must set the UUID if the switch is brought back to its factory defaults; a *Manufacturing Installation* by F5 personnel resets the switch and its UUID.

The UUID appears at the top of the output of a **show running-config** command. The following example shows the top of a running-config file from a switch named **stkbrA**. The UUID is in **bold** text.

```
; ARX-VE
; Version 5.03.000.13014 (Sep 24 2010 18:15:38) [nbuids]
; Database version: 503000.24
; Generated running-config Sun Sep 26 08:09:31 2010
; System UUID 8fa98111-55ec-d1c8-9380-8dtu78fab47d
; ip private subnet 169.254.170.0 255.255.255.0
;
terminal character-set unicode-utf-8
...
```

Applying the UUID

As shown in the preceding example, the initial-boot script has a prompt for the UUID. Enter the UUID of the replaced switch at this prompt. See the following example.

```
...
6. Enter the switch's UUID
   in the format
   xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx.(default=6ef37ece-416f-11df-b23c-11dbddf051df) #
8fa98111-55ec-d1c8-9380-8dtu78fab47d
...
```

Important

*No two running chassis should ever share the same UUID. Enter the UUID **only** in a switch-replacement scenario.*

Applying the Failed Switch's Master Key

A *master key* is an encryption key for all critical-security parameters (CSPs), such as administrative passwords. The global-config file is a list of CLI commands, including commands with encrypted CSPs; the replacement

switch requires the failed switch's master key to decode those CSPs and use them. As mentioned above, the failed switch's master key should have been saved along with its configuration files. Someone also should have saved the *system password* and the *wrapping password* for the master key. Both passwords are required to access and decrypt the master key.

As shown in an example earlier, there is a prompt for the master key in the initial-boot script. You can answer this prompt with the encrypted master key; the script then prompts for the wrapping password. For example,

```
...
A system password is required for access to the master key.
9. Enter a system password
   in the format text (12-28 characters). # d0uble$ecRET
   Confirm the system password # d0uble$ecRET

The master key is used to encrypt critical security parameters.
10. Enter the master key
    in the format base64-encoded key or keyword 'generate'.(default=generate) #
2oftVCwAAAAgAAAApwazSRFd2ww/H1pi7R7JMDZ9SoIg4WGA/XsZP+HcXjsIAAAADDRbMCxE/bc=

The wrapping password is used to encrypt and decrypt the master key.
11. Enter the wrapping password
    in the format text (6-28 characters). # an0ther$cretpw
    Confirm the wrapping password # an0ther$cretpw
...
```

Sample: Rebuilding a Defunct Switch

In the following sample script, a failed switch's private subnet, UUID, and master key are used to replace the failed switch. The initial questions have the same answers as shown in the non-replacement sample:

F5 ARX Startup

This F5 ARX does not currently have critical system information programmed. The following wizard prompts you for this information. You can connect to the switch through the management interface when you finish.

To restart the configuration program, enter 'r' at any prompt.

```
The switch's management port requires an IP address and mask.
1. Enter the management port IP address
   in the format nnn.nnn.nnn.nnn or 'none'. # 192.168.66.62
2. Enter the management port subnet mask
   in the format nnn.nnn.nnn.nnn.(default=255.0.0.0) # 255.255.0.0

The switch's management port requires a gateway IP address.
3. Enter the gateway IP address for the management interface
   in the format nnn.nnn.nnn.nnn or 'none'.(default=192.168.66.1) # 192.168.66.1
```

The next questions are relevant to switch replacement, and use the information saved from the failed switch:

```
The switch's internal subnet requires an IP address and mask.
4. Enter the switch's private IP address
   in the format nnn.nnn.nnn.nnn.(default=169.254.203.0) # 169.254.170.0
```

Chapter 3

Installing the ARX-VE

5. Enter the switch's private subnet mask
in the format nnn.nnn.nnn.nnn.(default=255.255.255.0) # <Enter>
6. Enter the switch's UUID
in the format
xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxx.(default=6ef37ece-416f-11df-b23c-11dbddf051df) #
8fa98111-55ec-d1c8-9380-8dtu78fab47d

The crypto-officer is the most privileged user in the system.

7. Enter the crypto-officer username
in the format text (1-28 characters). # **admin**
8. Enter the crypto-officer password
in the format text (6-28 characters). # **mypassword**
Confirm the crypto-officer password # **mypassword**

A system password is required for access to the master key.

9. Enter a system password
in the format text (12-28 characters). # **d0uble\$ecRET**
Confirm the system password # **d0uble\$ecRET**

The master key is used to encrypt critical security parameters.

10. Enter the master key
in the format base64-encoded key or keyword 'generate'.(default=generate) #
2oftVCwAAAAgAAAApwazSRFd2ww/H1pi7R7JMDZ9SoIg4WGA/XsZP+HcXjsIAAAADDRbMCxE/bc=

The wrapping password in use to encrypt and decrypt the master key.

11. Enter the wrapping password
in the format text (6-28 characters). # **an0ther\$cretpw**
Confirm the wrapping password # **an0ther\$cretpw**

The system displays a configuration summary. See the following example.

```
Configuration Summary
  Management IP Address    192.168.66.62
  Management IP Mask      255.255.0.0
  Management Gateway       192.168.66.1

Chassis GUID              8fa98111-55ec-d1c8-9380-8dtu78fab47d
Switch Password           #####
Switch Master Key
2oftVCwAAAAgAAAApwazSRFd2ww/H1pi7R7JMDZ9SoIg4WGA/XsZP+HcXjsIAAAADDRbMCxE/bc=
Crypto-officer Username   admin
Crypto-officer Password   #####
```

Enter 'yes' to load configuration or 'r' to redo the interview #**yes**

You have completed the switch startup configuration.
The switch will now initialize the local database.
When the login prompt appears, log into the switch using
the crypto-officer's username and password.

Closing configuration file.
Processing configuration file. (boot-config)

...

User Access Authentication

Username: **admin**

```

Password: mypassword
SWITCH>

```

At this point, the switch is ready for configuration through the GUI or CLI.

To re-establish network configuration, use the standard practice of copying and applying the failed switch's running config as shown in the following example:

```

SWITCH> enable
SWITCH# copy ftp://juser:jpasswd@ftp.wmed.com/arxVeConfig scripts running
SWITCH# show scripts

```

```

scripts
  running                Oct 12 17:45  2.1k

```

```

SWITCH# run scripts running

```

The running-config script set up all local parameters, such as the hostname and the network settings:

```

SWITCH#; ARX-VE
SWITCH#; Version 5.03.000.13014 (Sep 24 2010 18:15:38) [nbuilds]
SWITCH#; Database version: 503000.24
SWITCH#; Generated running-config Sun Sep 26 08:09:31 2010
SWITCH#; System UUID 8fa98111-55ec-d1c8-9380-8dtu78fab47d
SWITCH#; ip private subnet 169.254.170.0 255.255.255.0
SWITCH#;
SWITCH#terminal character-set unicode-utf-8
SWITCH#;===== vlan =====
SWITCH#config
SWITCH#  vlan 74
SWITCH#    description "personnel dept."
SWITCH#    members 1/4 to 1/4
SWITCH#...
SWITCH#;===== system =====
SWITCH#config
SWITCH(cfg)#  clock timezone America New_York
SWITCH(cfg)#  hostname stkbrgA
stkbrgA(cfg)#  ip domain-list wmed.com
stkbrgA(cfg)#  ...
stkbrgA(cfg)#  exit
stkbrgA#

```

See the *ARX® CLI Network-Management Guide* for detailed configuration instructions.

To re-establish your storage configuration and services, copy and apply the failed switch's global config as shown in the following example:

```

stkbrgA# copy ftp://juser:jpasswd@ftp.wmed.com/arxVeGblConfig scripts gblConfig
stkbrgA# show scripts

```

```

scripts
  running                Oct 12 17:45  2.1k
  gblConfig              Oct 12 17:52  3.4k

```

```
stkbrgA# run scripts gblConfig
stkbrgA# terminal character-set unicode-utf-8
stkbrgA# global
stkbrgA(gbl)# kerberos auto-realm-traversal
stkbrgA(gbl)# kerberos health-check threshold 3500
stkbrgA(gbl)# nfs tcp timeout 30
stkbrgA(gbl)# ...
stkbrgA(gbl)# exit
stkbrgA#
```

See the *ARX® CLI Storage-Management Guide* for detailed configuration instructions.

Connecting the Ethernet Management Interface

After you boot the switch, you can access either the ARX GUI or CLI through its management interface. To access the GUI, launch a web browser to the interface over HTTPS, for example:

```
https://192.168.66.62/
```

Use the crypto-officer username and password, entered above, to log in (for example, **admin** and **mypassword**).

For the CLI, use SSH with the interface and the crypto-officer username. For example:

```
ssh admin@192.168.66.62
```

The [ARX GUI Quick Start: Network Setup](#) guide contains instructions for getting started with the GUI, and the *ARX® CLI Network-Management Guide* contains instructions and best practices for using the CLI.

Powering Down the ARX-VE

You can power down the ARX-VE with a CLI command (`shutdown`). Unlike hardware platforms with a battery-backed NVRAM, you can leave the ARX-VE shut down indefinitely; when you power it back on, no re-imports (of filer storage) are required.

Converting From the Trial Version to the Production Version

Converting an ARX-VE running from a trial version to a production version requires you to purchase an ARX-VE production license from F5 Networks. Clear the evaluation license and enable the production license. You will be warned to increase the memory allocated to the ARX-VE to 4GB, and to provision a second CPU core. You will then need to reload the ARX-VE instance.



A

Best Practices For ARX-VE

- [Best Practices For ARX-VE](#)

Best Practices For ARX-VE

Observe the following best practices when using the ARX-VE:

Issue	Recommendation
VMware snapshots on ARX-VE	Avoid using VMware snapshots on ARX-VE. VMware snapshots are not likely to work correctly, given the ARX's dependencies on external filers and a separate metadata filer.
Live migration of ARX-VE virtual machines	Perform live migration of ARX-VE virtual machines (using VMware VMotion) only when the ARX-VE is idle or lightly loaded (e.g., during off-hours). Live migration of ARX-VE while the virtual machine is processing traffic could produce unexpected results.
VMware DRS environments	In DRS environments, perform live migration of ARX-VE virtual machines (using VMware VMotion) only when the ARX-VE is idle or lightly loaded (e.g., during off-hours). Live migration of ARX-VE while the virtual machine is processing traffic could produce unexpected results. Disable automatic migrations by adjusting the VMware VMotion DRS Automation Level to Partially Automated, Manual or Disabled on a per-ARX-VE basis.
Time synchronization	Configure all ARX-VE systems to use an external time synchronization source. You can do this either by configuring NTP within ARX-VE or by checking the Synchronize guest time with host box within vSphere Client and configuring all VMware hosts to share a single NTP time server or set of related NTP time servers. Note that units within a redundant system configuration must share a common time synchronization source, to prevent inconsistent system behavior.
Default route for management port	Define a default route for the virtual management port.
Network latency	Network latency between the ARX-VE and its metalog datastore should be kept to 250 microseconds or less to ensure optimal performance.



Index

A

audience for this manual 1-3

B

booting the switch 3-13

C

configuring the switch 3-13

R

Redundancy

 between two ARX-VE's 2-3

running the boot wizard 3-13

S

switch boot-up

 for replacement switch 3-15

switch installation

 tools required 3-3

switch installation, unpacking and installing 3-1

T

tools for installation 3-3

