

Includes
Six Easy Steps
for Firepower
Network Analysis
every morning!



CCIE/CCNP SECURITY

Exam 300-710: Securing Networks with
Cisco Firepower (SNCF)
Volume I

Todd Lammle
Donald Robb

CCIE/CCNP Security Exam 300-710: Securing Networks with Cisco Firepower (SNCF)

Volume 1

Todd Lammle Donald Robb

Copyright © 2021 by Todd Lammle
All rights reserved. ISBN: 9798635481059

About the Author

is the authority on Cisco certification and **Todd Lammle** internetworking. He is a world-renowned author, speaker, trainer, and consultant and is Cisco certified in most Cisco certification categories.

Todd has over three decades of experience working with LANs, WANs, and large enterprise licensed and unlicensed wireless networks. Lately, he's been implementing advanced security technologies in large data centers and organizations in the USA and abroad with technologies including Cisco ISE, StealthWatch, and Cisco Firepower/FTD.

Todd's many years of real-world experience is evident in his writing. He isn't just another author—he's an accomplished networking engineer, having cultivated extensive practical experience working on some of the world's largest networks at companies like Xerox, Hughes Aircraft, Texaco, AAA, Cisco, and Toshiba, among many others.

Todd has published over 100 books, including the uber-popular *CCNA: Cisco Certified Network Associate Study Guide*, *CCNA Wireless Study Guide*, *CCNA Data Center Study Guide*, *SSFIPS Firepower Study Guide*, *Firepower*

6.x Study Guide, and this volume as well.

Todd runs an international consulting and training company based in Evergreen, Colorado. He spends his free time with family, friends, and his two golden retrievers playing in the mountains, rivers, and lakes that surround his home.

iii

You can reach Todd through his website and find more Firepower/ FTD material, videos, and classes at www.lammle.com/firepower.

Don't forget to download your bonus chapter, "Six Easy Steps to Network Analysis Every Morning" from [Lammle.com/ firepower!](http://Lammle.com/firepower)

Donald Robb , also known as the-packet-thrower online, has been the industry for over 15 years doing everything from help desk to network architect. He has worked in most areas of IT to an expert level including networking, security, collaboration, data center, wireless, and service providers. Currently he is a Principal Consultant for a major Cisco Value Added Reseller (VAR) and focuses on complex projects that usually involve several different technologies. He also acts as the company's Subject Matter Expert (SME) for Automation/ DevOps, SDN, and Security topics.

During his time, he has worked with most of the big vendors and some of the smaller ones too and have earned many advanced certifications and specializations. These include most of the cloud certifications such as the Azure Solutions Architect Expert, most of the Cisco certifications, and some random ones across vendors like Fortinet, Palo Alto, and HPE. He is also recognized as a Cisco FireJumper proving his proficiency with Cisco's security solutions, and worked with Todd Lammle on several books and courses.

You can reach Donald through www.lammle.com or check out his YouTube channel at: <https://www.youtube.com/c/ThePacketThrower>

Table of Contents

About the Authoriii

Acknowledgments.....	vii
Introduction	viii
What Is This Book Really About?	xiii
What Does This Book Cover?	xv

Securing Networks with Cisco Firepower (SNCF 300-710) Exam Objectives	xix
Chapter 1: Chapter 2:	

Chapter 3: Chapter 4: Chapter 5: Chapter 6: Chapter 7: Chapter 8: Chapter 9: Firepower Management Center (FMC)	1
--	---

Cisco Firepower Management Center (FMC) Configuration	20
---	----

Firepower Management Center (FMC) Actions	63
Licensing & Health Policy	79
Chassis Manager	95
Firepower Devices	155
High Availability	217
Objects	251
Access Control Policy	309
Chapter 10: Malware and File Policy	376
Chapter 11: Firepower Network Discovery	418

v

Chapter 12: Intrusion Prevention System (IPS) Policy	443
Chapter 13: DNS Policy	496
Chapter 14: Prefilter	514
Chapter 15: Network Address Translation (NAT)	529
Chapter 16: Identity Policy	547
Chapter 17: User Management	570
Chapter 18: Advanced Network Analysis	596

Acknowledgments

I want to thank Donald Robb for his contribution to chapter 5 on Chassis

Manager and RadWare, and in addition, to helping me edit this book when I was just too tired to even talk about editing anymore! Don has become a great friend!

Also, I wanted to mention my good friend John Gay, who also helped me pen the first ever published Sourcefire/Firepower Study guide: SSFIPS, four years ago. He's been in inspiration during good and bad times. Thank you, John, for your friendship and faith!

I'll also mention Alex Tatistcheff, even though he was out selfishly working on his own Firepower book during the writing of this book and didn't work on this book at all! However, that said, he still is an enormous vast knowledge of Snort information that is 2nd to none. He's the best of the best in Firepower/Snort. Thank you for your advanced technical support over the years, Alex!

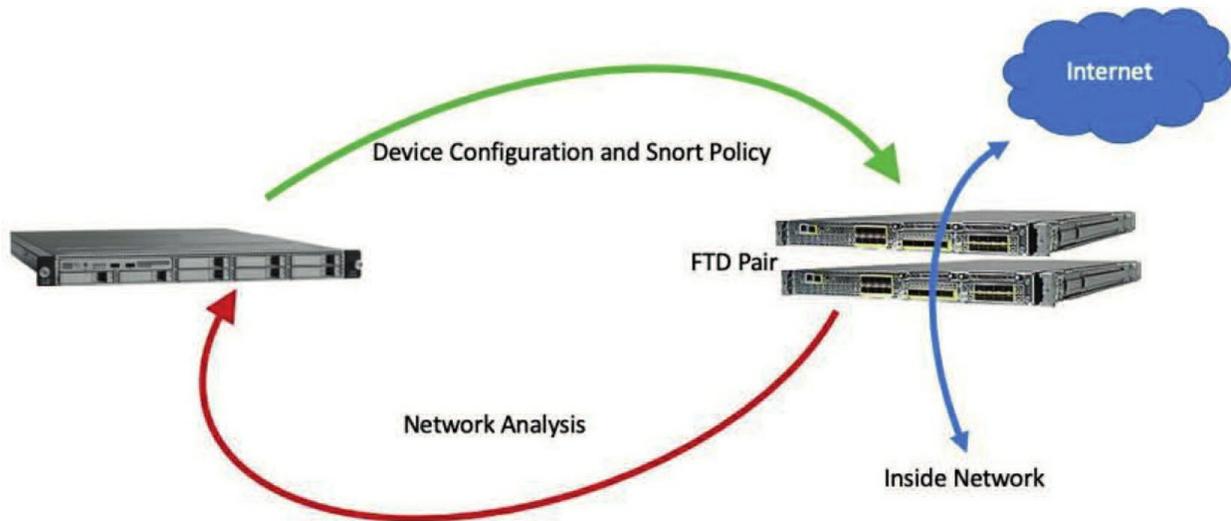
vii

Introduction

Cisco Firepower is an integral part of the suite of Cisco security products. There are Firepower managers and the various Firepower devices that are configured, managed, and monitored from the managers.

The devices are further categorized into Firepower appliances like 7000/8000s, which are all EOL, but there's a legion of SourceFire appliances out there, so Cisco still covers them in this exam for now.

In the figure below, you can see how the manager—an FMC in this example—sends configuration and Snort security policy out to the devices. The devices then make decisions about the packets traveling through them based upon the Snort Security policy and finally sends the Snort verdict back to the manager:



The new Adaptive Security Appliance (ASA), called Firepower Threat Defense (FTD), is definitely all the rage now! Even if you really just want to run plain, powerful ASA code, you’ll want the new devices with their tremendous power and inspection throughput that

viii

can run either ASA or FTD code. So, the ASA code isn’t going away anytime soon, although most people run FTD when they get the new devices.

I’m going to talk about the various Firepower managers in this introduction, but I’m only going to use one of them throughout the book—think exam objectives! The other Firepower devices will still be discussed thoroughly throughout this book though.

History of SourceFire/Firepower

In early 2012, Sourcefire introduced version 5 of the “SourceFire System.” Along with this new version came several new brands and one of them was FirePOWER.

FirePOWER was used to represent the advanced network interface hardware in the latest detection devices. The new Netronome Flow Processor (NFP) interface included far more advanced technology than a typical network interface card. Technically this is still the case—the real power behind the

detection speed of those expensive devices is still FirePOWER. Today, this term has changed a bit because it's now used in conjunction with the Cisco ASA. Now, when you see *FirePOWER*, it's almost always used to describe "FirePOWER Services on ASA." This could mean software services, or the FirePOWER blade installed on the ASA 5585-X.

What you won't see is the term *FirePOWER* being used in accordance with the 7000/8000 appliances or FTD devices. They just refer to everything as Firepower now, so I'll be doing that in this book to reference all products as well.

FireSIGHT is another term introduced with version 5. Historically, (pre-Cisco), the term *FireSIGHT* referred to the passive detection capabilities of the Sourcefire System. In version 4.x, these capabilities were called Realtime Network Awareness (RNA) and Realtime User Awareness (RUA). When version 5 hit the scene, these two names were rebranded as *FireSIGHT* to refer to the new and improved RNA/RUA.

Prior to the Cisco acquisition, FireSIGHT never meant Snort or IPS detection. It was all about network and user awareness. But today, the term *FireSIGHT* has been expanded to Firepower, which encompasses the entire NGIPS/NGFW system. The "Firepower System" is the new Cisco IPS, and as I said, when we're talking about this "system," we just use the term *Firepower* now.

Managing Firepower

Even though there are many ways to manage your Firepower appliances and/or Firepower Threat Defense (FTD) devices, only the Firepower Management Center (FMC) is covered in the exam objectives. So again, that's the only manager I'll use throughout this book. Here are all the options:

- Firepower Device Manager (FDM)
- Firepower Management Center (FMC)
- Cisco Defense Orchestrator (CDO)
- Adaptive Security Device Manager (ASDM)

Let's explore each one:

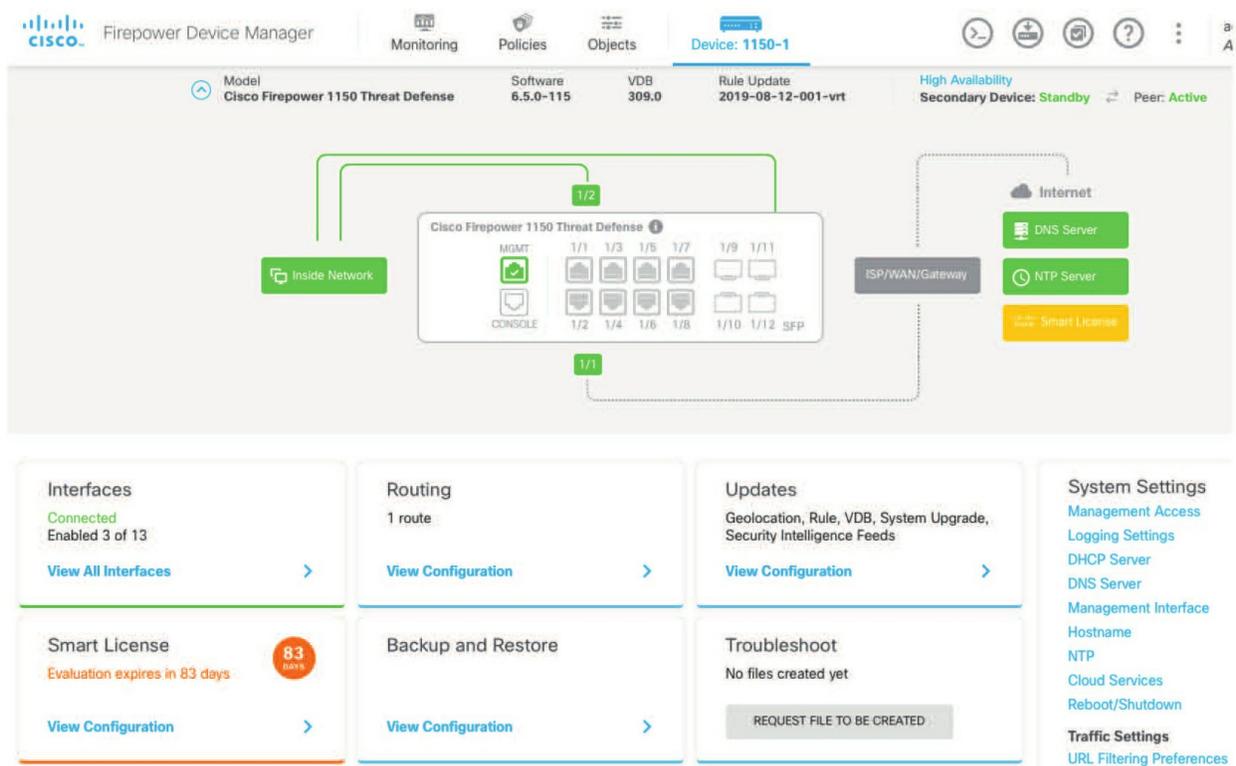
Firepower Device Manager (FDM)

This little power manager is used for SOHO environments or single-device configuration where you have no FMC available. Over the last five years, FDM has become more useful and easier

x

to configure a single device with than in the past. Here's a look:

You can see here in my FDM output for my 1150 device above, that I received a nice GUI output of the interfaces and basic configuration. You can do quite a bit with the FDM, just not as much as you can with the FMC at this point.

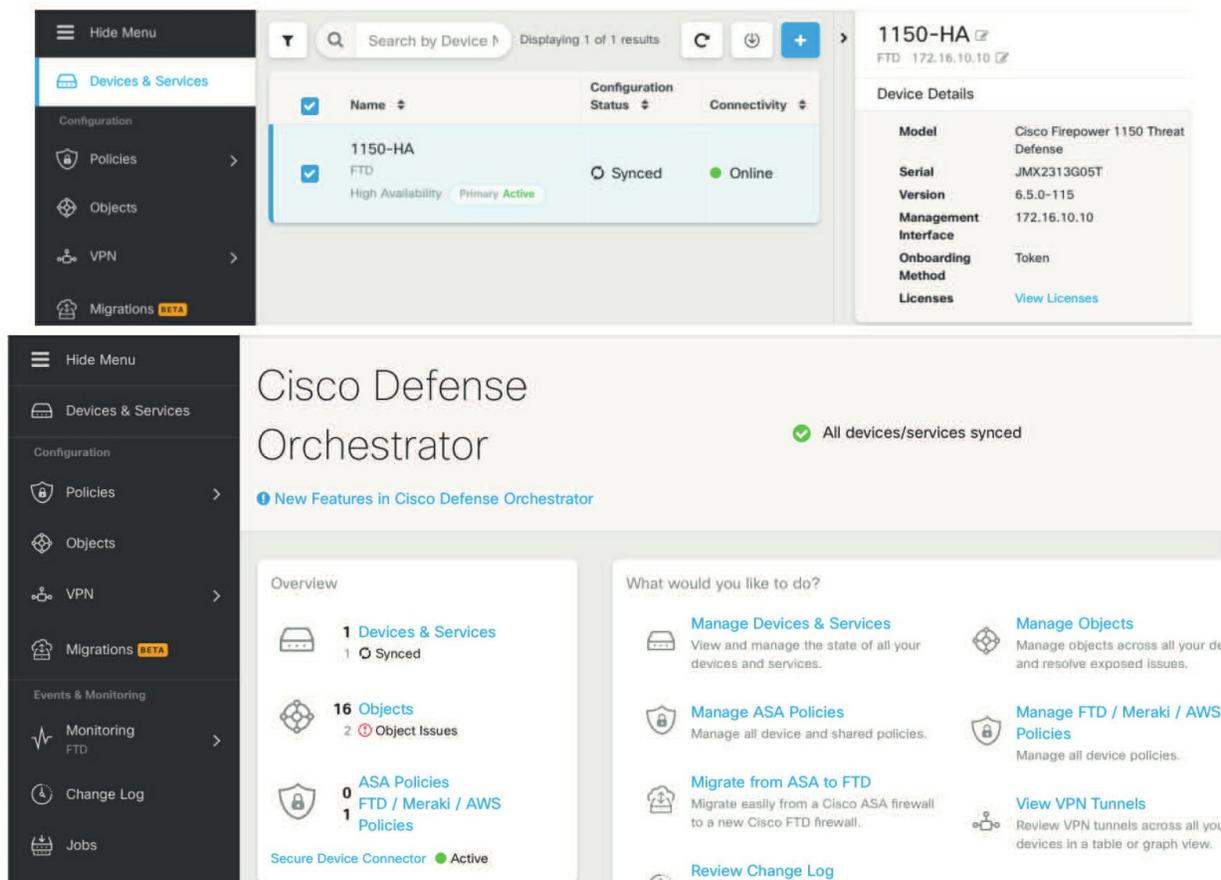


Firepower Management Center (FMC)

This is the most prevalent Firepower manager, and its available in both virtual and hardware versions. The difference really comes down to how many devices you need to configure/manage. No need to show it here because you have 36 chapters of FMC coming up in two books!

Cisco Defense Orchestrator (CDO)

New and upcoming, this cloud-based management system is the future of managing Cisco security products like the ISE, Firepower, StealthWatch, AMP ASAs, and more:



This product will be very useful for managing a suite of Cisco security products for a company and even more helpful for companies that have devices requiring management over a large geographic spread. I didn't cover CDO in this book but look for this in the future as your manager for all things Firepower.

Adaptive Security Device Manager (ASDM)

This powerful little beauty has been around for a long time and has really helped configure, manage, and troubleshoot our small to large ASA deployments for more than a decade. ASDM provides some of the configuration and management capability of FirePOWER but not enough to use on its own. I've also used something called Cisco Security Manager (CSM) for centralized management of very large environments with hundreds

of ASAs.

xii

What Is This Book Really About?

Of course, I definitely penned this book with the CCNP Security SNCF objectives in mind. Even so, I included a few chapters that go beyond the objectives and really dig deep into all my years of real-life Firepower experience!

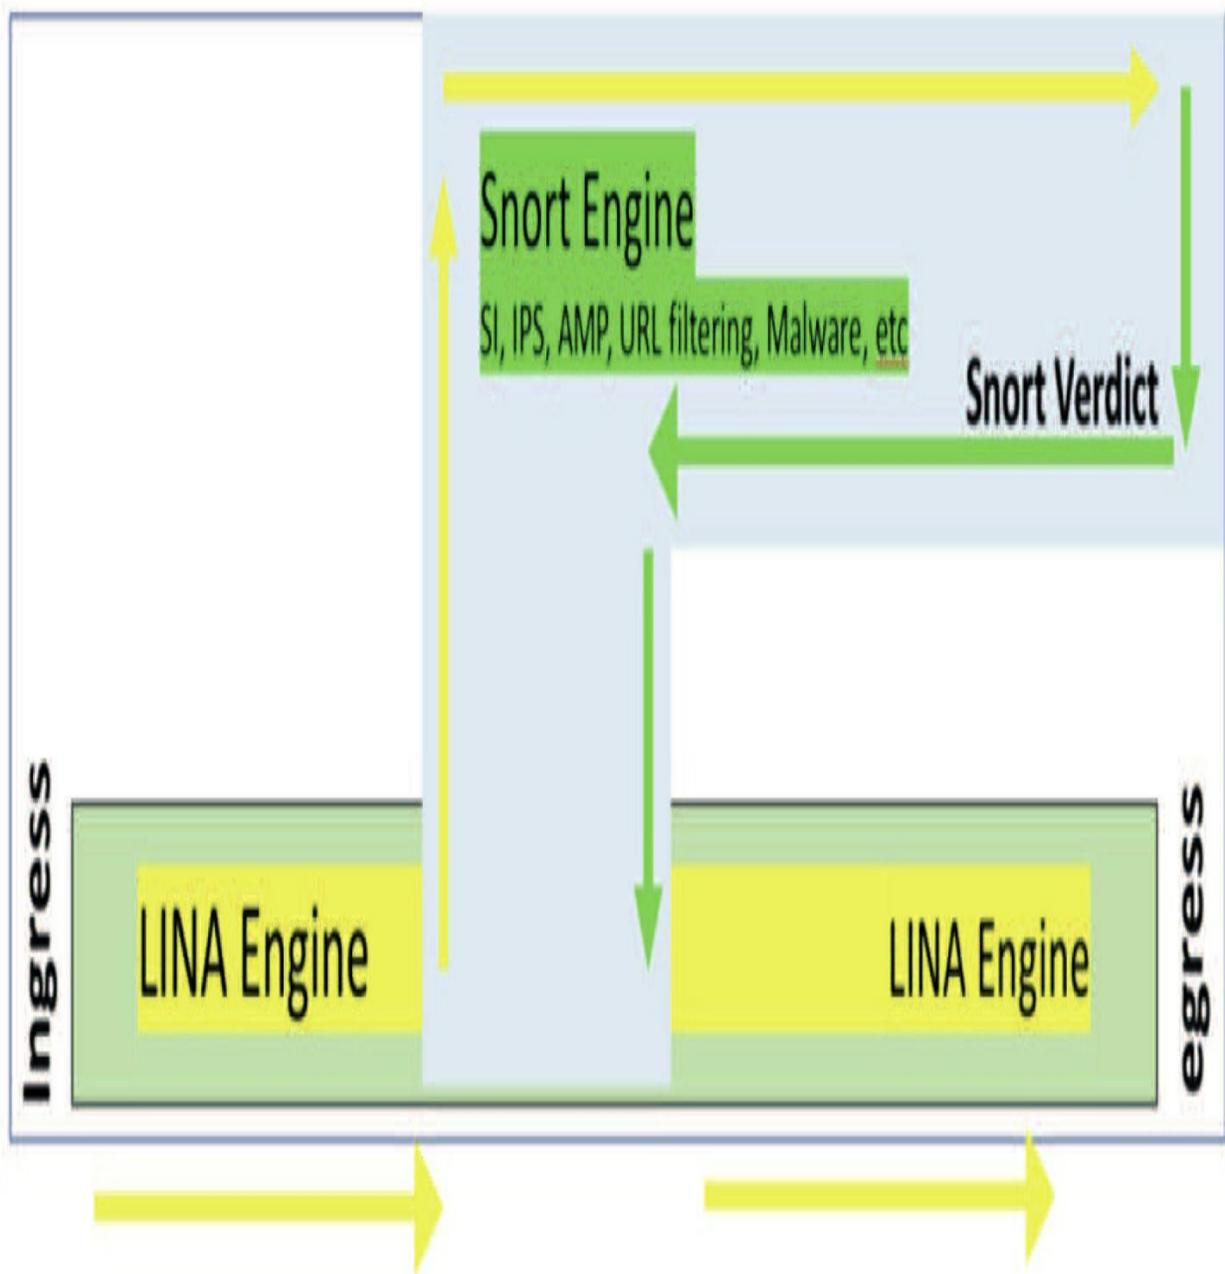
The information in those is very valuable because it's based upon my real-world experience at hundreds of customers after installing well over 10,000 FTD devices in Fortune 50 and 500 companies all over the world.

We'll follow the Cisco FTD packet flow that took me quite a while to draw out because when I began working with this product, they didn't have FTD. Plus, when FTD finally made its appearance, it still took years for documentation worth reading to come out!

Whether you have a Firepower appliance (7000/8000), ASA with a FirePOWER module, or an FTD device, the Snort engine is basically the same for all models and configured mostly the same way through the FMC.

The figure below is based on FTD code, and you can tell because you can see the LINA engine, which is really the integrated base ASA code configured through the FMC. That's what makes the FTD different.

xiii

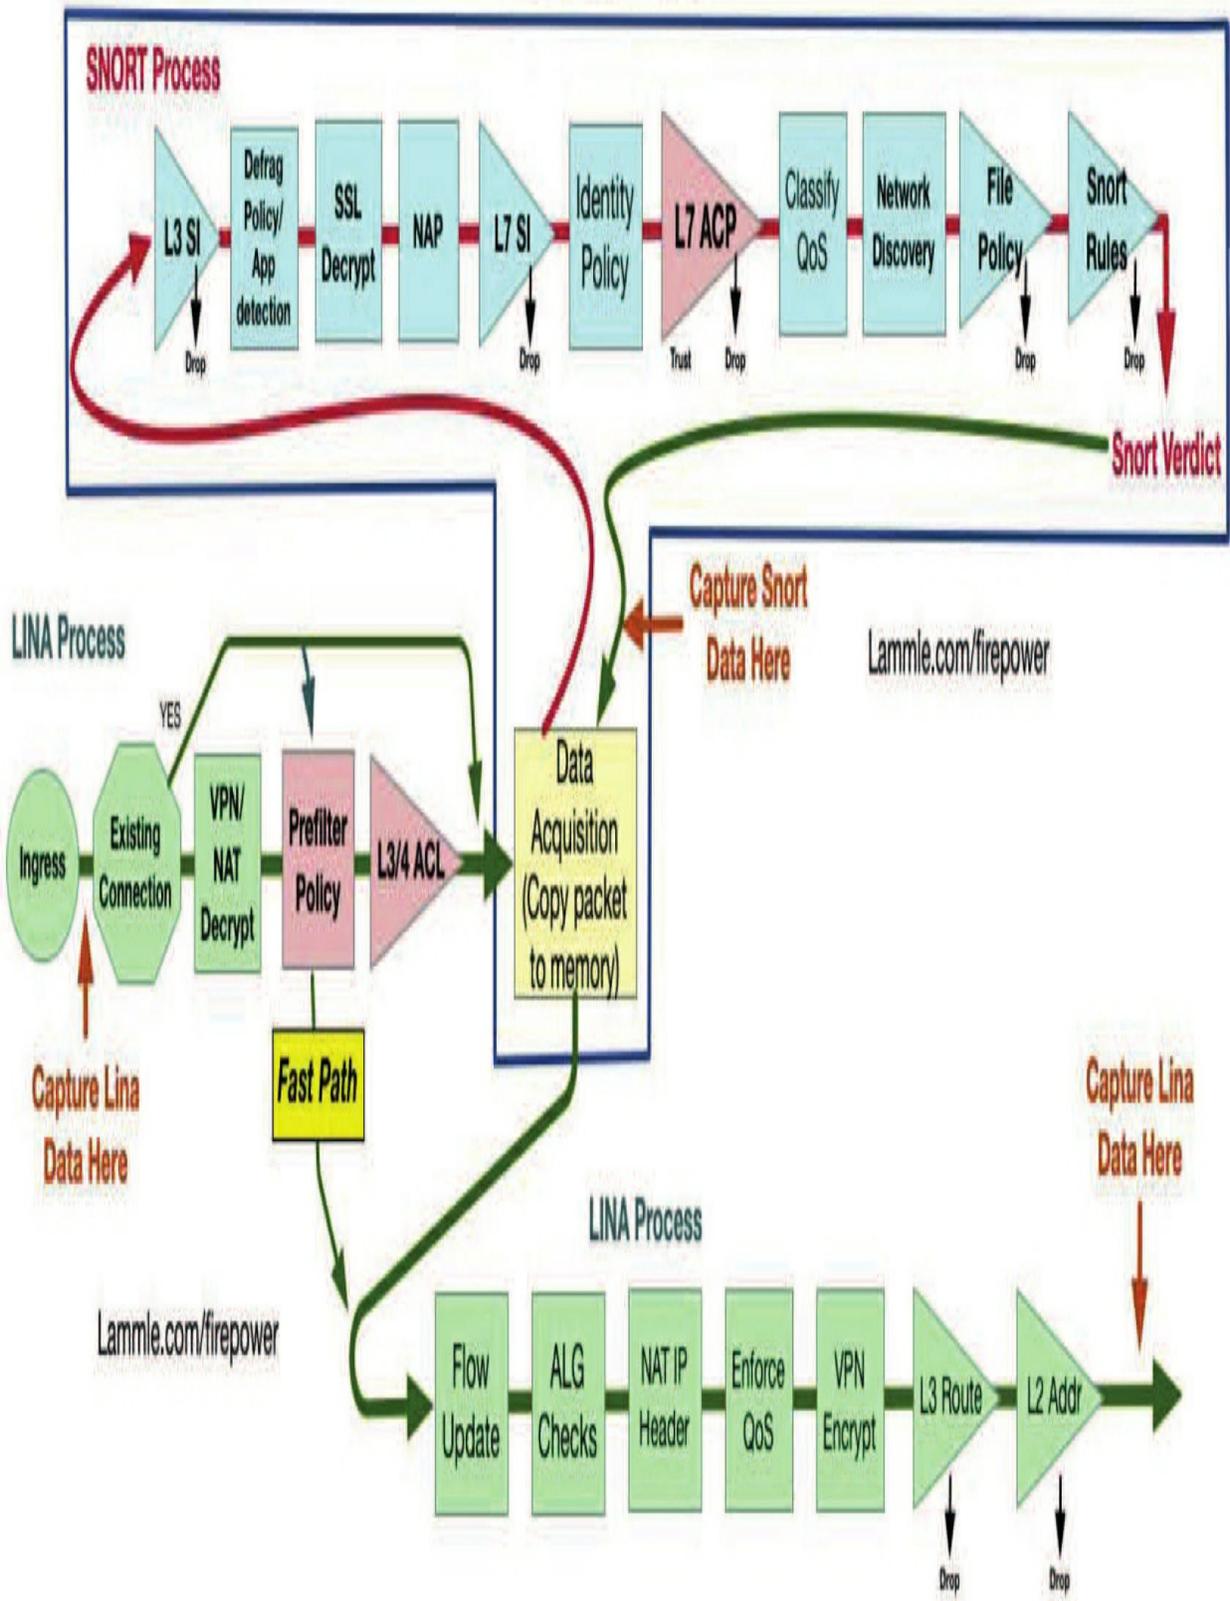


Based on packet traces and lots of network analysis, I was able to draw a really accurate packet flow for FTD that I'll use for reference throughout the book as we go through each policy.

I am well aware that the Cisco objectives also include the old, outdated Firepower appliances and they still follow the same flow, but only in the Snort process.

The appliances don't have a prefilter or LINA process at all, but that is the only difference.

You can see in the next drawing of the Ingress LINA process that the packets are delivered to the Snort process via a memory location. Then, after the packets traverse the Snort process, the Snort verdict is sent to the Lina egress for processing.



We'll be traveling through the drawing above step-by-step in this book.

What Does This Book Cover?

You'll be learning all of the following in this book:

Introduction

What is Firepower? What was FireSIGHT? What was SourceFire? Understand Firepower by building a solid foundation in defining key, industry-wide, plus Cisco-specific terms that we'll be using throughout this book, like FMC and FTD.

Chapter 1: Firepower Management Center (FMC) This chapter will cover the steps on how to install and configure both the virtual FMC and hardware FMC.

Chapter 2: Cisco Firepower Management Center (FMC) Configuration This chapter will cover how to log in and configure your Firepower Management Center (FMC) using System Configuration. It covers how to apply settings on the systems to control user preferences, NTP, time zones, and other key factors.

Chapter 3: Firepower Management Center (FMC) Actions A detailed and important chapter, though actions are something you don't configure much—typically only when you're bringing up your FMC. This allows you to create alerts using email, SNMP, and syslog to help you manage your network and find issues faster.

Chapter 4: Licensing & Health Policy

In this chapter, I'll cover the licensing for Firepower, both the Classic and Smart Licensing. We will also cover the Health settings, which have remained largely unchanged for over a decade, though they're still important to understand and configure.

xv

Chapter 5: Chassis Manager

This is a more advanced chapter and very important if you have 4100/9300

devices. Don't forget there are exam objectives on FXOS and Chassis Manager! We'll go through the configuration, troubleshooting, upgrades, and finally, clustering of devices in Chassis Manager.

Chapter 6: Firepower Devices

This very important chapter will go into detail on how to bring up the various Firepower devices like an appliance and the different FTD devices: 1000/2100/4100/9300. We'll start at the beginning by demonstrating how to reboot and reset the devices, then go step-by-step in configuring the devices and bringing them into their FMC.

Chapter 7: High Availability

This is a really great chapter about providing HA on hardware FMCs and the various Firepower devices. Once we configure the HA pairs, we'll go through verification and troubleshooting and finally finish with Firepower upgrades for the FMC and devices.

Chapter 8: Objects

This chapter will provide you with the understanding of the object types that are used by the Firepower system. Firepower employs reusable configuration components—objects—to provide an easier way to use values across policies, searches, reports, dashboards, and so on.

Chapter 9: Access Control Policy

This chapter covers the heart of the Firepower system. An AC policy acts kind of like a central traffic cop for Firepower because all traffic passing through a device is processed through it. You'll find plenty of great tips in this chapter, along with real-world examples on how to configure and implement this at work!

Chapter 10: Malware and File Policy

A nickname derived from the term *malicious software*, malware comes in a variety of disgusting flavors, from coded weapons fashioned to damage, control or, disable a computer system to reconnaissance, stealing data, and identity theft. Firepower's Advanced Malware Protection (AMP) is designed to tackle one of the worst and arguably most prevalent threat vectors of today—malware!

Chapter 11: Firepower Network Discovery

Firepower Network Discovery (formerly FireSIGHT) is the name given to a technology built into the Cisco Firepower NGIPS to provide us with contextual awareness regarding events, IP addresses, users on the network, and even background about the hosts in the system. Once you've been acquainted with this awesome technology, we'll move on to explore discovery components like the discovery policy, type of data collected, connection events, and host attributes associated with it.

Chapter 12: Intrusion Prevention System (IPS) Policy This chapter provides you with the background necessary for success in the real world with a thorough presentation of IPS policy management. This in-depth chapter covers IPS policies, which precisely describe the suspicious and/or malicious traffic that the system must watch out for. The Intrusion policy also controls how evil traffic is dealt with when it's discovered.

Chapter 13: DNS Policy

This chapter's focus will be on one of the newer features in Firepower—the DNS policy, which gives us visibility beyond typical packet sniffing

xvii

detection by granting additional insight into potentially compromised hosts or encrypted data.

Chapter 14: Prefilter

While it's not actually part of the Access Control policy, the Prefilter policy actually processes traffic first. This policy only applies to Firepower Threat Defense (FTD) devices.

Chapter 15: Network Address Translation (NAT) I'll survey the different NAT types available to managed devices in this chapter and show you how to continue to configure your devices. I'm going to cover some best practices from the real world in it too!

Chapter 16: Identity Policy

This chapter covers identity—the ability for Firepower to take different actions depending on the user associated with a connection. The concept of using Firepower to block or allow traffic based on users or groups is very

popular because it's really helpful to many organizations. It keeps employees more productive and allows blocking access to websites or services based on job function.

Chapter 17: User Management

In this chapter, we're going to cover a variety of administrative functions for user account management. You'll learn all about creating and managing both internal and external users.

Chapter 18: Advanced Network Analysis

This is an awesome chapter chock-full of pro tips and essential, real-world knowledge and skills. We'll dive deep into using the Firepower System to analyze intrusion event data and explore some of the workflows available when analyzing events. I'll show you lots of examples of how to drill into relevant event data.

xviii

Securing Networks with Cisco Firepower (SNCF 300-710) Exam Objectives

(SNCF 300-710) Exam Objectives 710 exam. Still, I'm going to cover a lot more administration and troubleshooting than what's on the exam, so you can definitely use this book as a guide for your real-life network easily!

Here's the listing of objectives, which is the foundation of the chapters in this book.

Exam Description

Securing Networks with Cisco Firepower v1.0 (SNCF 300-710) is a 90-minute exam associated with the CCNP Security and Cisco Certified Specialist - Network Security Firepower certifications.

This exam tests a candidate's knowledge of Cisco Firepower Threat Defense and Firepower 7000/8000 Series virtual appliances, including policy

configurations, integrations, deployments, management and troubleshooting.

This book will help you prepare for this exam in full and includes: ▪ Policy configurations

- Integrations
- Deployments
- Management and troubleshooting

The CCNP Security SNCF exam objectives covered in this study guide include:

1.0 Deployment 30% 1.1 Implement NGFW modes

1.1.a Routed mode

1.1.b Transparent mode

1.2 Implement NGIPS modes

1.2.a Passive

1.2.b Inline

1.3 Implement high availability options

1.3.a Link redundancy

1.3.b Active/standby failover

1.4 Describe IRB configurations

2.0 Configuration 30%

2.1 Configure system settings in Cisco Firepower Management Center

2.2 Configure these policies in Cisco Firepower Management Center

2.2.a Access control

2.2.b Intrusion

2.2.c Malware and file

2.2.d DNS

2.2.e Identity

2.2.g Prefilter

2.3 Configure these features using Cisco Firepower Management Center

xx

2.3.a Network discovery

2.3.d Actions

2.4 Configure objects using Firepower Management Center

2.4.a Object Management

2.4.b Intrusion Rules

2.5 Configure devices using Firepower Management Center

2.5.a Device Management

2.5.b NAT

4.0 Integration 15% 4.1 Configure Cisco AMP for Networks in Firepower Management Center

4.2 Configure Cisco AMP for Endpoints in Firepower Management Center

What Objectives Will the Second Book in This Series Cover? CCNP

Security SNCF series part II, we will explore these important exam objectives and topics:

1.0 Deployment 30% 1.3.c Multi-instance

2.0 Configuration 30% 2.2.f SSL

2.3.b Application detectors (OpenAppID)

2.3.c Correlation

2.5.c VPN

2.5.d QoS

2.5.e Platform Settings

2.5.f Certificates

3.0 Management and Troubleshooting 25%

3.1 Troubleshoot with FMC CLI and GUI

3.2 Configure dashboards and reporting in FMC

3.3 Troubleshoot using packet capture procedures 4.0 Integration 15%

4.3 Implement Threat Intelligence Director for third-party security intelligence feeds

4.4 Describe using Cisco Threat Response for security investigations

4.5 Describe Cisco FMC PxGrid Integration with Cisco Identify Services Engine (ISE)

4.6 Describe Rapid Threat Containment (RTC) functionality within Firepower Management Center

What chapters are covered in Part II in this CCNP Security

SNCF series?

Chapter 19: Platform Settings

Chapter 20: Domains

Chapter 21: Dashboards and Reporting

Chapter 22: FTD Quality of Service

Chapter 23: Network Analysis Policy (NAP)

Chapter 24: Correlation Policy

Chapter 25: SSL

Chapter 26: Cisco Threat Response

xxii

Chapter 27: Multi-Instance

Chapter 28: PxGrid

Chapter 29: Rapid Threat Containment (RTC) Chapter 30: Threat Intelligence Director (TID) Chapter 31; Remote Access VPN S2S

Chapter 32: Virtual Private Networks AnyConnect Chapter 33: Certificates

Chapter 34: Troubleshooting

Chapter 35: Troubleshooting Part II

Chapter 36: Application Detectors

Chapter 37: Firepower SafeSearch

Chapter 38: Daily Network Analysis Steps

xxiii

Chapter 1: Firepower Management Center (FMC)

The following CCIE/CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.5 Configure devices using Firepower Management Center

2.5.a Device Management

We're going to begin exploring this amazing technology by diving into the Cisco FMC installation and configuration. I'll open by telling you how to

install a Firepower network as well as how the FMC works with your devices. Then we'll journey into the various FMCs and how to install a Firepower Management Center (FMC)—both hardware and virtual. I'll wrap the chapter up by showing you how to configure the FMC at startup so you can access a graphical user interface (GUI) to manage your network. Anywhere from 1 to 750 devices can be managed this way!

Our focus through this entire book series will be zeroed in on only the Cisco Firepower Management Center (FMC) for managing your Firepower network. Yes, of course there are a couple other ways, options, and approaches that I talked about in the introduction, but the exam, my own networks, and all my clients use the FMC exclusively. If you happened to skip the introduction, at least go back and skim over the parts that explain the differences between an FMC, FTD, and Firepower appliance.

Because the virtual FMC is actually a low-cost and easy solution to install and manage, I think that it is by far the better solution for all Cisco networks. And what the exam wants you to know is definitely important too, isn't it?

It's also really key for you understand the difference between the virtual and hardware FMC, so I'll be going deep into both of these in this chapter too.

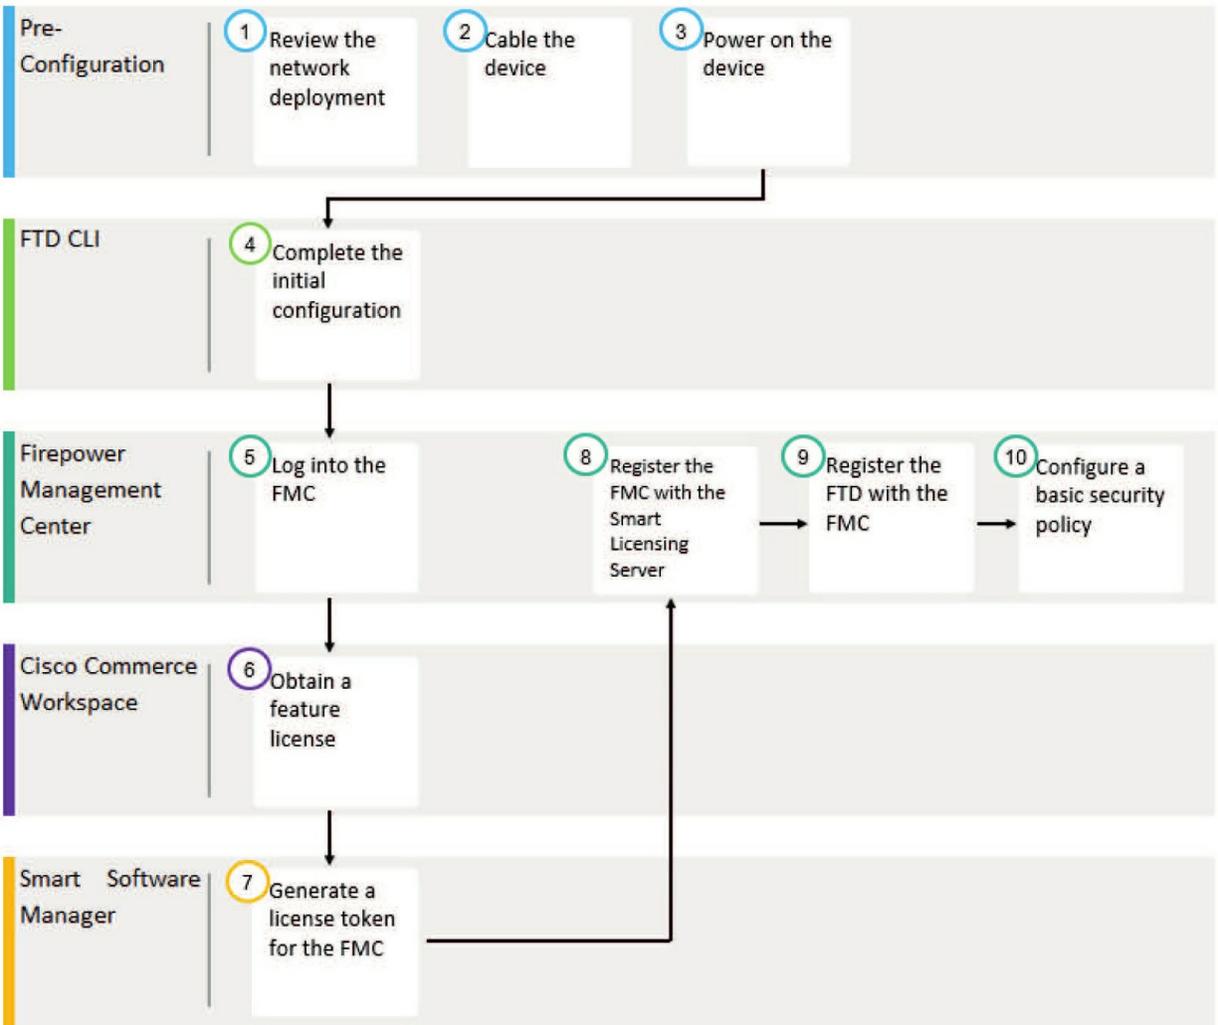
To find exam study material like videos, downloadable supplemental material, and practice questions, please see www.lammle.com/firepower

Deploying a Cisco Firepower Network

Before we actually get into installing a new FMC, I want to cover some steps that Cisco has laid out regarding how to install and deploy a new Firepower system.

As you probably know, it's still kind of tough to find helpful, reliable documentation about Cisco Firepower. Still, there are some hidden gems out there like the layout shown here, which Cisco calls the end-to-end deployment:

Let's walk through these steps now.



Steps 1, 2, and 3: Pre-Configuration

This is where we review our network deployment and create a management VLAN that your Firepower devices will communicate on using TCP 8305. We'll physically cable the devices with the management port of each device placed into a switch port configuration with the management VLAN, then power up the device.

Step 4: FTD CLI

Log in to the command-line interface (CLI) of each device and complete the initial configuration of the Firepower devices from the CLI, which configures and enables the management ports.

Step 5: Log In to FMC

Once the pre-configuration for the FMC is complete (what this chapter is all

about), then you need to log in so you can configure the FMC and prepare the appliance for managing Firepower devices like FTD.

Step 6: Cisco Commerce Workspace – Give us money! Obtain licenses for the Firepower Management Center and buy feature licenses like malware and URL filtering.

Step 7: Smart Software Manager

Obtain licenses for the Firepower Management Center: Generate a license token for the FMC and obtain licenses for the Firepower Management Center.

Steps 8, 9, 10: Firepower Management Center

Log in to the FMC and register the appliance with the Smart Licensing server for the devices you want the FMC to manage. The licenses stay on the FMC and are not transferred to the devices. Chapter 4 is all about licensing, so hang in there regarding licensing details!

Finally, after Step 9, we can start configuring some basic security policies on the FMC like File & Malware, Security Intelligence (SI), and IPS, which can then start dropping bad packets.

I'm going to discuss what the FMC actually does, and then we'll get started on the FMC installation.

What Is a Firepower Management Center (FMC)?

Cisco refers to the FMC as “an administrative nerve center for managing critical Cisco network security solutions.” That’s a mouthful for sure, but the FMC does provide complete and unified management for Firepower devices, referred to as Next Generation Firewalls (NGFWs). These provide application visibility and control (AVC), intrusion prevention (IPS), URL filtering, and advanced malware protection along with lots of other important things that I’ll cover in detail throughout this book.

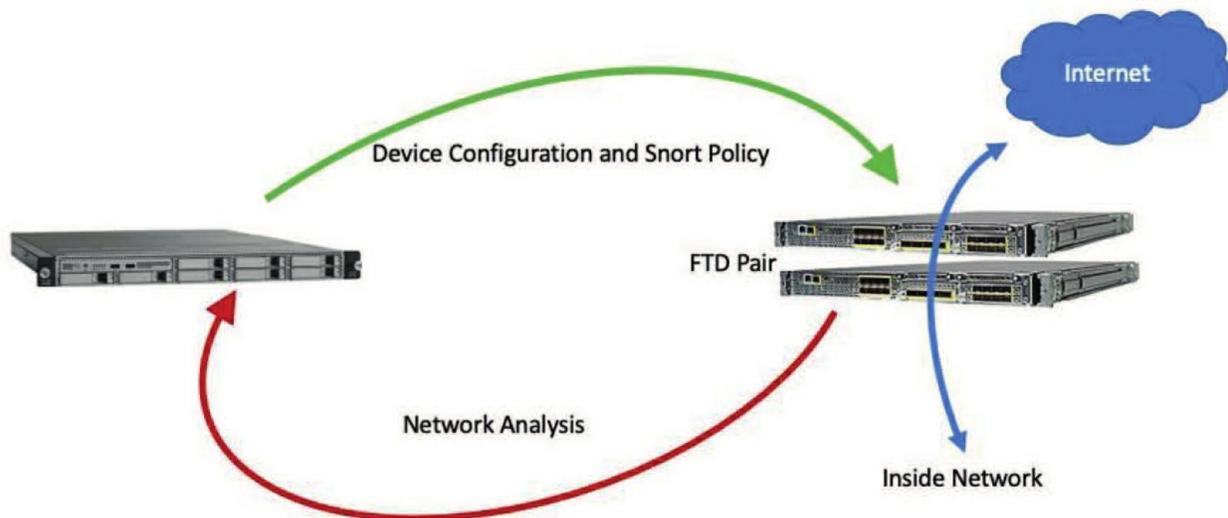
The single management solution provided by the FMC gives us unparalleled visibility and correlation of events that no other product on the market can match! This is really what this book, and the corresponding exam, is all

about.

Your job is to learn how to correctly configure potentially multiple traffic-sensing, managed devices like Firepower Threat Defense (FTD) that monitor traffic for analysis and reports to a manager. And don't forget, we can use the Firepower Device Manager (FDM) as a manager, which provides us with a centralized management console, complete with a nice, graphical user interface (GUI). This GUI is what we'll use to perform administrative, management, analysis, correlation of events, and reporting.

The next figure shows how the FMC configures devices and policies and then deploys these policies to the device(s).

It's important to understand that the FMC makes no decisions on what happens to the packets going through your network. Only the devices make these decisions, based on the Snort process that I discussed in the introduction of this book.



In this simple example, the FMC is sending the configuration and policies to the FTD devices, which make decisions based on Snort policy configuration. If a Snort event occurs, the logging and the packets themselves will be sent to the FMC for analysis.

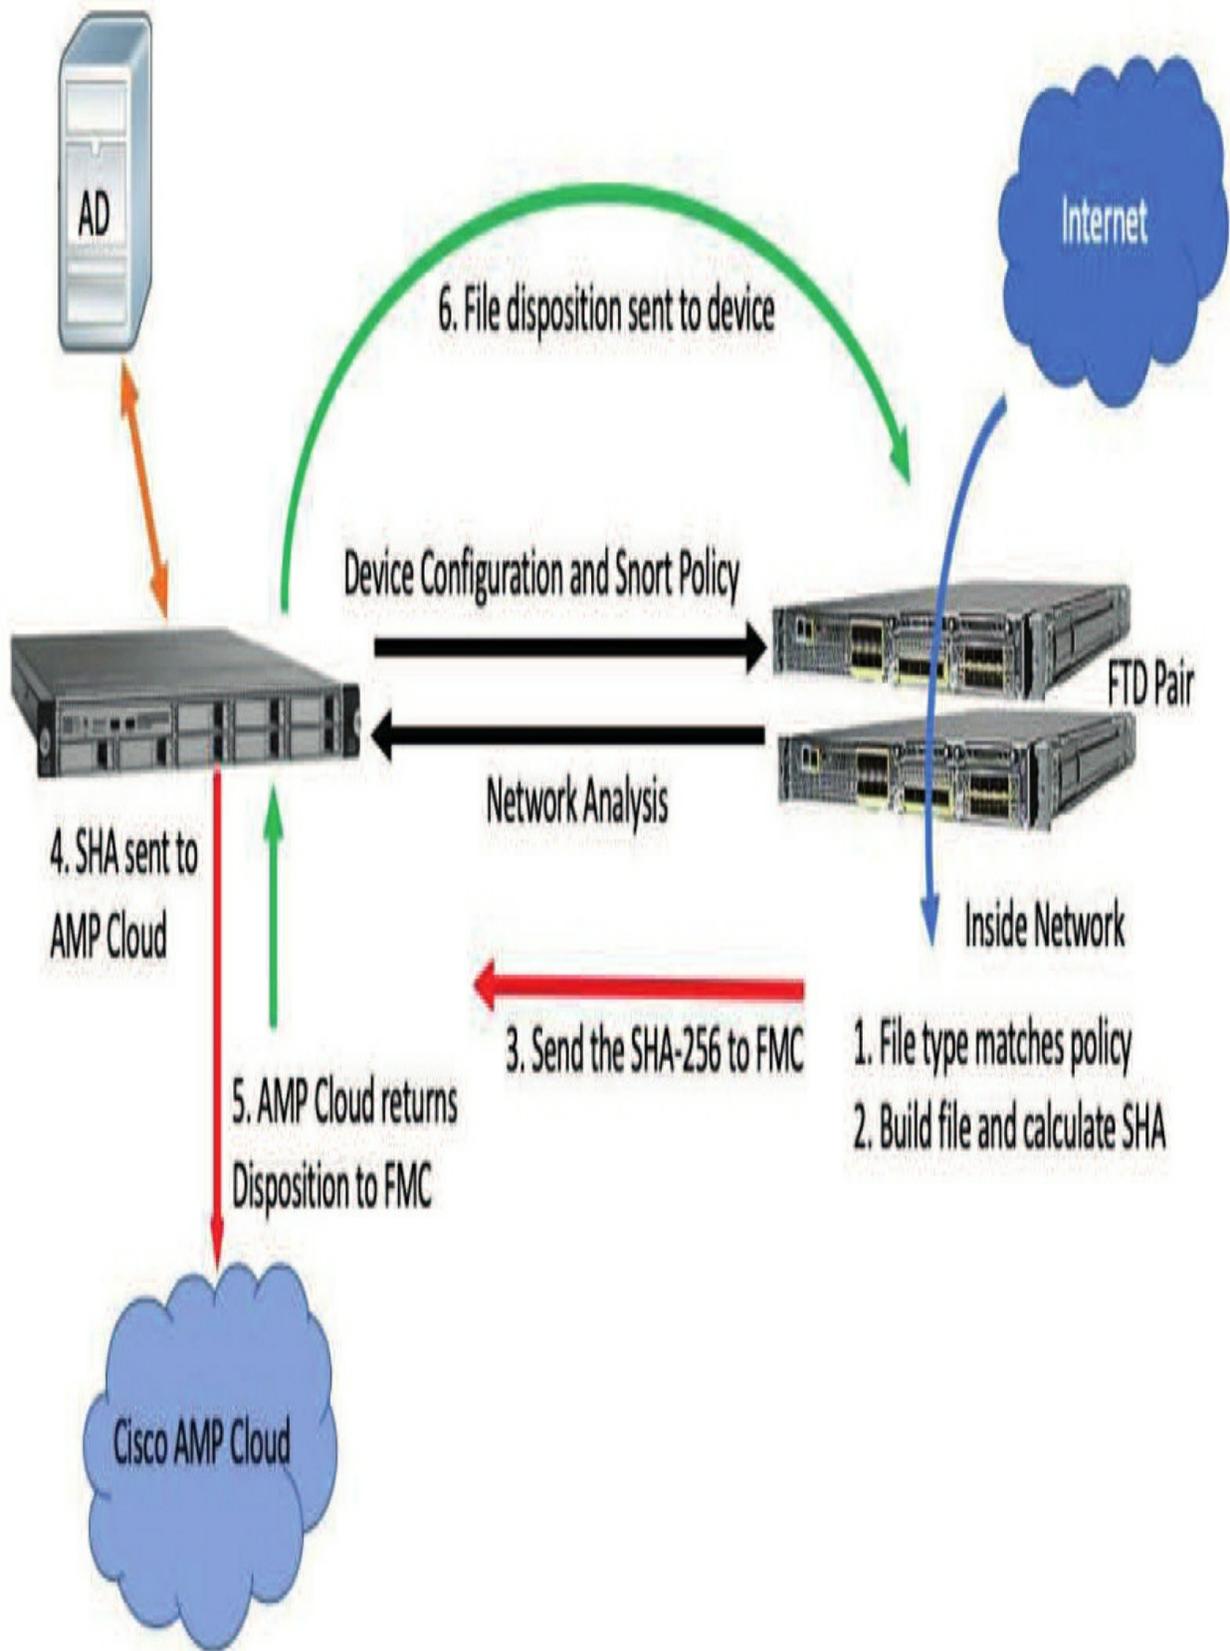
What If Your FMC Goes Down?

Life being less than perfect, what happens if you lose your FMC for some reason? Well, once the configuration and policies have been deployed and working on the devices, the answer is, not much! Of course, you won't be getting network analysis and you won't be able to make any changes to the devices, but the Firepower device will actually keep passing/dropping packets based on the configuration and policies it's already received from the FMC.

There are two things you lose access to when your FMC goes down:

1. File SHA-256 Hash checks
2. Any rule created using AD users and groups will no longer be able to be filtered because the AD integration only connects to the FMC, and the Firepower device needs to query the FMC for the AD information.

To clarify this process shown in the figure above:



1. The Firepower/FTD device reads the data and finds packets that match a file policy rule.
2. The device builds the file out of the network flow and calculates the SHA-256 of the file.
3. The device transmits the hash to the FMC—only the SHA-256 hash is transmitted, not the file.
4. The FMC checks its local cache and if necessary, transmits the hash to the Cisco AMP cloud.
5. A disposition is then returned to the FMC.
6. The disposition is forwarded to the device.

If the disposition returned is malware, the device can block the file and maybe even store the file on the FMC depending on the file policy settings in place. This lookup happens really fast—usually well under 600 milliseconds! Know that if this process takes more than 2 seconds to find the disposition, the file will just be passed.

One last thing here...If you have rules in your access control policy (ACP) that use AD users/groups, the devices must query the FMC to see if the rule matches because only the FMC integrates with the AD using the Identity policy. So, if you lose your FMC, the users would then be permitted to say, social media, which you were blocking via AD groups. (No worries, I'll cover Identity policy in depth later in this book!)

Okay—time to explore the two types of FMCs, the virtual and hardware versions, so you can discover the differences between them.

Virtual FMCs

Since late 2019, Cisco now has two virtual FMCs on the market: the original version, which managed up to 25 devices, and its new virtual FMC that runs in beast mode, managing up to 300 devices!

We'll get into both FMCs here before we move on to the hardware FMCs.

FMC Virtual

The FMC Virtual has been around a pretty long time in technology years, probably because it's extremely cost effective for starters. Virtual FMCs only

use a small amount of server resources, but even though claims are that it can manage up to 25 devices, I'd make sure to only manage 14 to 16 (7 to 8 pairs) at the most using the FMC Virtual. This depends on your traffic, of course.

So, let's install the vFMC on vCenter now. I'll be referring to and using this FMC throughout the book along with my hardware FMCs that we'll install next.

1. Okay—from your vCenter, choose to deploy an OVF template:

2. Choose the following files if you're using a vCenter, after you expand the tar file:

Navigator

Back

vcsa.aps.lab
LammleDC

FMC

IN

IN 0

IN 02

IN 03

FMC

Getting Started Summary

Virtual Machines VM Templ

New Virtual Machine...

Actions - FMC

New Virtual Machine

New vApp

New Folder...

Deploy OVF Template...

Move To...

Rename...

Tags & Custom Attributes

Add Permission...

Alarms

Remove from Inventory

Update Manager

01

02

03

04

05

06

07

09

10

11

17-65

 Cisco_Firepower_Mgmt_Center_Virtual_VMware-6.6.0-35.tar

 Cisco_Firepower_Mgmt_Center_Virtual_VMware-6.6.0-35-disk1.vmdk

 Cisco_Firepower_Mgmt_Center_Virtual_VMware-VI-6.6.0-35.mf

 Cisco_Firepower_Mgmt_Center_Virtual_VMware-VI-6.6.0-35.ovf

3. Name and choose the location for the new vFMC: 4. Now choose the host to install the FMC on.

Deploy OVF Template

- ✓ 1 Select an OVF template
- 2 Select a name and folder**
- 3 Select a compute resource
- 4 Review details
- 5 Select storage
- 6 Ready to complete

Select a name and folder

Specify a unique name and target location

Virtual machine name: FMC 03

Select a location for the virtual machine.

- ▼  10.9.200.1
 - ▼  DC
 - >  Discovered virtual machine
 - >  SD-WAN Servers
 -  StudentVMs

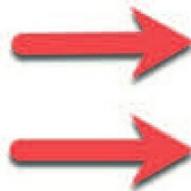
Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- 4 Review details**
- 5 License agreements
- 6 Select storage
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

Review details

Verify the template details.

Publisher	No certificate present
Product	Firepower Management Center Virtu
Version	6.6.0
Vendor	Cisco Systems, Inc.
Description	Cisco Firepower Management Cente Tasman Dr San Jose, CA 95134 USA
Download size	2.2 GB
Size on disk	4.7 GB (thin provisioned) 250.0 GB (thick provisioned)



5. Review details—notice it’s 2.2 GB in size and will take 250 GB of disk space using thin provisioning:

6. Accept the license agreements

7. Choose your disk and provisioning. I typically use thin provisioning, but you need to research and make your own decision here:

The screenshot shows the 'Select storage' step of a VM creation wizard. On the left is a progress list with '6 Select storage' highlighted. The main area has a 'Select storage' section with a dropdown menu for 'Select virtual disk format' showing options: 'Thin Provision' (selected), 'Thick Provision Lazy Zeroed', and 'Thick Provision Eager Zeroed'. Below is a table for 'VM Storage Policy'.

Name	Capacity	Free Space	Used Space
datastore1	92.5 GB	299.7 GB	78.45 GB

8. Select your management VLAN that you configured in advance.

9. Finish customizing your vFMC and then click Next. You don’t have to configure anything here except the password, but you can set the DNS and NTP servers if desired - more on this in chapter 2.

The screenshot shows the 'Password' configuration screen. It has a title bar 'Password' with '1 settings'. Below is a section '01. Password' with a text area containing password requirements: 'admin password Password must meet the following criteria: - At least 8 characters - At least 1 lower case letter - At least 1 upper case letter - At least 1 digit - At least 1 special character such as @#*-_+! - No more than 2 sequentially repeated characters - Not based on a simple character sequence or a string in password cracking dictionary You can provide a password that does not meet these criteria, but on initial login you will be forced to change your password to meet these requirements.' At the bottom are two input fields: 'Password' and 'Confirm Password'.

Notice the default IP address of 192.168.45.45/16

is assigned to the vFMC:
Now after you hit Save
and the virtual machine
installs, it's time to power the
vFMC on.

You should receive a login
screen when you https:// to
the IP address you just set, or
the default of 192.168.45.45.

I'll show you the login in
the next chapter.

Network		13 settings
Hostname		Fully Qualified Domain Name <hr/>
DNS1		Primary DNS Server <hr/>
DNS2		Secondary DNS Server <hr/>
NTP1		Primary NTP Server <hr/>
NTP2		Secondary NTP Server <hr/>
IPv4 Configuration		IPv4 Configuration <input type="button" value="Manual"/>
IP Address		IPv4 Address 192.168.45.45 <hr/>
Netmask		IPv4 Netmask 255.255.0.0 <hr/>
Gateway		IPv4 Gateway 192.168.1.1 <hr/>
IPv6 Configuration		IPv6 Configuration <input type="button" value="Disabled"/>



FMC Virtual 300

Most people that use the virtual FMC do so because of cost. The FMC Virtual is super cheap compared to the hardware models! Of course, the ease of setting up and managing it comes into play too. The biggest caveat is that the virtual FMC can only support up to 25 devices, which as I mentioned, is majorly pushing its limits.

The FMC Virtual 300 is where the true power is. It was created and designed to give serious, commanding glory to the virtual world of FMCs, and with support of up to 300 devices, it does that in spades!

But it all that power sure isn't free—it's way more expensive than the original FMC virtual and is comparably priced to some hardware FMCs at this point. Plus, it requires a whole lot more server resources. So if you go with the v300, you basically need to have a serious server setup as well, and all of this adds up to parting with some serious cash!

But even considering the cost, would it still be better than going with a hardware FMC? Maybe. The UCS chassis used for the FMCs will eventually always go EOL, so at some point, you'll lose support and have to buy a new one. I'll discuss the difference between the old and new hardware FMCs in a bit, but for now just know there really aren't a lot of changes. Interestingly, the UCS chassis just went EOL, so they've come out with new hardware FMCs and new numbers. Support for the old ones will be around for a while, so what's the problem? Well, for one thing, some of the new features in the new codes are only available with the new FMCs, so there's that!

Let's talk about what you need to have available for the FMC v300. The smaller FMC virtual version can run with as little as 8 GB of RAM and 4 cores, with 250 GB of disk space. But the minimum specs for the new virtual 300 are as follows:

- A minimum 4.8 GB of disk space (thin provisioned).
- 64 GB of RAM.
- 32 CPUs.

VM Hardware	
CPU	32 CPU(s), 0 MHz used
Memory	65536 MB, 0 MB memory active
Hard disk 1	2.25 TB

Wow! Better start looking for some new server hardware to run this bad boy.

Hardware FMCs

In the next chapter, I'll cover the process of initially configuring your FMC, which will get you to the point at which you can register your Firepower devices into the FMC and manage them. Once the devices are connected to the FMC, we'll configure the interfaces, routing, and device settings.

The FMC will only have management interfaces. Remember, the FMC doesn't actually do any detection itself. Instead, it pushes policy and then receives events from the devices to enable us to perform correlation of events and network analysis.

Once you've gotten the FMC physical appliance racked and powered, your first task will be to connect and configure the management network interface. This procedure varies depending on the device type and may require more than one person.

Let's take a quick look at the differences between hardware FMCs. The old and new FMCs are pretty similar, with same number of sensors managed and the same amount of RAM. There are some small CPU differences. The new FMCs have more storage as well as hot swappable drives and two 10 Gbps LAN interfaces, but these factors still don't add up to a big enough reason to upgrade at this point in my opinion.

Here are the specs for the various hardware FMC's:

FMC1000/1600

Up to 50 sensors managed

30 million maximum IPS events

90 million connection events

900 GB event storage

32 GB RAM

Events per second: 12,000

Network map up to 50K hosts, 50K users

One Intel E5-2620 V4 CPU / 1 Intel Xeon 4110 processor Two 900 GB SAS drives /Two 1.2 TB 10-K SAS HDDs RAID 1 100 Mbps/1 Gbps Ethernet / 100 Mbps/1 Gbps/10 Gbps Ethernet

FMC2500/2600

Up to 300 sensors managed

60 million maximum IPS events

300 million connection events

1.8 TB event storage

64 GB RAM

Events per second: 12,000

Network map up to 150K hosts, 150K users

Two Intel E5-2620 V4 / Two Intel Xeon 4110 processors Four 600 GB SAS drives /Four 600 GB 10K SAS HDDsRAID 5 100 Mbps/1 Gbps Ethernet / 100 Mbps/1 Gbps/10 Gbps Ethernet

FMC4500/4600

Up to 750 sensors managed 300 million maximum IPS events 1 billion connection events
3.2 TB event storage
128 GB RAM
Events per second: 20,000
Network map up to 600K hosts, 600K users
Two Intel E5-2640 V4 CPUs / Two Intel Xeon 4116 processors Six 800 GB SSDs /Ten 1.2 TB SAS SSDsRAID 6
100 Mbps/1 Gbps Ethernet / 100 Mbps/1 Gbps/10 Gbps Ethernet

The primary management port for all of the hardware FMCs is eth0, and you can use eth1, eth2, and eth3 as secondary management or event ports. Most people don't know that you can use these other ports to separate management and event traffic on your hardware FMCs, so I'll definitely cover this in more detail in the next chapter.

Okay—so let's power a hardware FMC up now.
For this book, I'll be using the two 2500 hardware FMCs that I have available in my rack that I use in my classes.

Starting the Firepower Management Center

The hardware FMCs are considered devices, but an FMC isn't actually a device—it's an appliance. So, all we just need to do is turn it on and set the IP. We'll log in through the GUI and configure the 2500 FMC in the next chapter.

The procedure to configure the management IP is the same for all hardware FMC types and involves logging in from the console and setting the management IP address from the command line. There are three options available to access the console:

- (1) Connect a USB keyboard and VGA monitor.
- (2) Connect to the serial console port.
- (3) Connect via SSH to the default IP address of 192.168.45.45.

For virtual FMCs, it's much easier. You access the console through your vCenter or other virtual application as I demonstrated earlier.

Even though the first option for the hardware FMC is pretty simple, you still need to find a keyboard and monitor and plug it into the appropriate ports on the rear of the appliance.

The challenge here is actually more of a physical thing. These are UCS chassis, which are heavy, long, and just really unwieldy, so it's a good idea to have someone help you. Plus, you need to stay in your data center to start the configurations on these hardware devices. You might be fine on your own if you have a lot of room to maneuver in the location you're installing them.

If you're going with the second option, connect a computer to the serial console port on the appliance using a rollover serial cable. This is definitely the simplest method if you have a laptop handy!

If you don't have a laptop, or a keyboard and monitor for some reason, you can also connect from a computer via SSH. The FMC ships with the default IP address of 192.168.45.45, which is assigned to its management port. You can just use a normal Ethernet cable to connect directly between your computer and the FMC.

Using whichever method works for you, once you connect to the console, you'll be greeted with a login prompt. Log in with the default credentials:

Username: **admin**

Password: **Admin123**

Run the configure network script with the following command:

```
>expert
```

```
>sudo /usr/local/sf/bin/configure-network
```

We trust you have received the usual lecture from the local System Administrator. It usually boils down to these three things:

#1)Respect the privacy of others.

#2)Think before you type.

#3)With great power comes great responsibility.

Password:

Do you wish to configure IPv4 (y or n) y

Management IP address? 172.16.10.20

Management netmask? 255.255.255.0

Management Default gateway? 172.16.10.1

Are these settings correct? (y or n) y

Do you wish to configure IPv6 (y or n) n

Updated network configuration.

Please go <https://172.16.10.20/> to finish installation. And that's it! Now we are ready to move on to the next chapter, where we'll log in to the FMCs, and then start configuring them.

Summary

This chapter opened with a talk about how to install a Firepower network and how the FMC works with your devices.

You then learned about the various FMCs and how to install both the virtual and hardware versions.

I showed how to configure the FMC at startup so that a graphical user interface (GUI) can then be accessed and used to manage your network.

We ended the chapter by covering how to start up the FMC.

Chapter 2: Cisco Firepower Management Center (FMC) Configuration

The following CCIE/CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.1 Configure system settings in Cisco Firepower Management Center

In this chapter, I'll continue where we left off in Chapter 1, the installation of the FMC. I'll be working with the Firepower Management Center, or FMC, and I'll start by covering the system configuration of the FMC, which can be used to harden it, for example. Since we will be using the FMC to perform virtually all the management tasks, it makes sense to spend time up front on its configuration before bringing the 7000/8000 appliances and Firepower Threat Defense (FTD) devices into the FMC.

Also, for this chapter, I'll use both a hardware FMC and a virtual FMC so you can see the small but important differences.

To find exam study material like videos, downloadable supplemental material, and practice questions, please see www.lammle.com/firepower

Initial FMC Login

It really doesn't matter if you have a virtual FMC or a hardware FMC when it comes to basic configuration. These next first steps laid out in the beginning of this chapter are the same, so let's start right where we left off from the last chapter.

After you configure the management IP address as discussed in Chapter 1, your next step will be to connect to the Web UI and complete the initial setup. To do this, you will connect to `https://<yourFMC-IP>`, where you will be presented with the initial login screen.



The default username/password for a hardware FMC is **admin** for the username and **Admin123** for the password. For a virtual FMC, you can no longer install the FMC without setting a password, although it doesn't have to be a hard one, you'll just be asked to change it when you first login.

Once logged in, you will see the initial settings page where you can configure your password and license. This is the only time you will see these pages. Upon subsequent logins, you will be taken to the main FMC landing page (the Dashboard by default).

This next figure shows the initial password page.

Notice it says Step 1 of 3 on the top right. However, as of this writing, there are only two steps...go figure.

In the figure, you can see the criteria that is now needed on the Cisco FMC to set your admin password.

After you set your password, click Next.

* Be advised, this password is no longer

For security and privacy reasons, you must change the default password before configuring this appliance.

New Password

Generate password

Confirm Password

Show password

Password must meet the following criteria:

- × 8 characters
- × One lower case letter
- × One upper case letter
- × One digit
- × One special character such as @#*-_+!
- × No more than 2 sequentially repeated characters
- × Passwords match

Next

recoverable, so do not lose it!

The next figure shows the license page you'll then receive.

Configure Smart Licensing

You can configure licensing now or later. To configure licensing now:

- Register device with Cisco Smart Software Manager
 1. Create or log into your [Cisco Smart Software Manager](#) account.
 2. In your assigned virtual account, click the "General tab", click on "New Token".
 3. Copy the token and paste it here:
- Start 90-day evaluation period without registration.
Please make sure you register with Cisco before the evaluation period ends.
Otherwise you won't be able to make any changes on the device.

Enable Cisco Success Network

Cisco Success Network enablement provides usage information and statistics to Cisco which are essential for Cisco to provide technical support. This information also allows Cisco to improve the product and to make you aware of unused available features so that you can maximize the value of the product in your network. Check out the [sample data](#) that will be sent to Cisco.

Management Center establishes a secure connection to the Cisco Cloud so that it can participate in additional...

The FMC uses Smart Licensing for all devices except for the legacy ASA with Firepower platform. If you have Smart Licensing setup, you can paste in your registration token here.

Cisco does provide a 90 evaluation that you can use if you aren't setup for Smart Licensing or this is a lab, it supports everything, but you can only use DES encryption for your VPNs due to export laws.

Hold on there, partner! If you are using a code previous to 6.5, then when bringing up a FMC, you will get the screen shown here first before you can move on.

Change Password

Use these fields to change the password for the admin account. Cisco recommends that you use a password that has at least eight alphanumeric characters of mixed case and includes at least one numeric character. Avoid using words that appear in a dictionary.

New Password

Confirm

Network Settings

Use these fields to specify network-related information for the management interface on the appliance.

Protocol

IPv4 IPv6 Both

IPv4 Management IP

10.11.10.15

Netmask

255.255.255.0

IPv4 Default Network Gateway

10.11.10.1

Hostname

firepower

Domain

Primary DNS Server

Secondary DNS Server

Tertiary DNS Server

Time Settings

Use these fields to specify how you want to set the time for the Defense Center.

Set My Clock

Via NTP from 0.sourcefire.pool.ntp.org, 1.sourcefire

Manually 2019 / September / 2 , 16

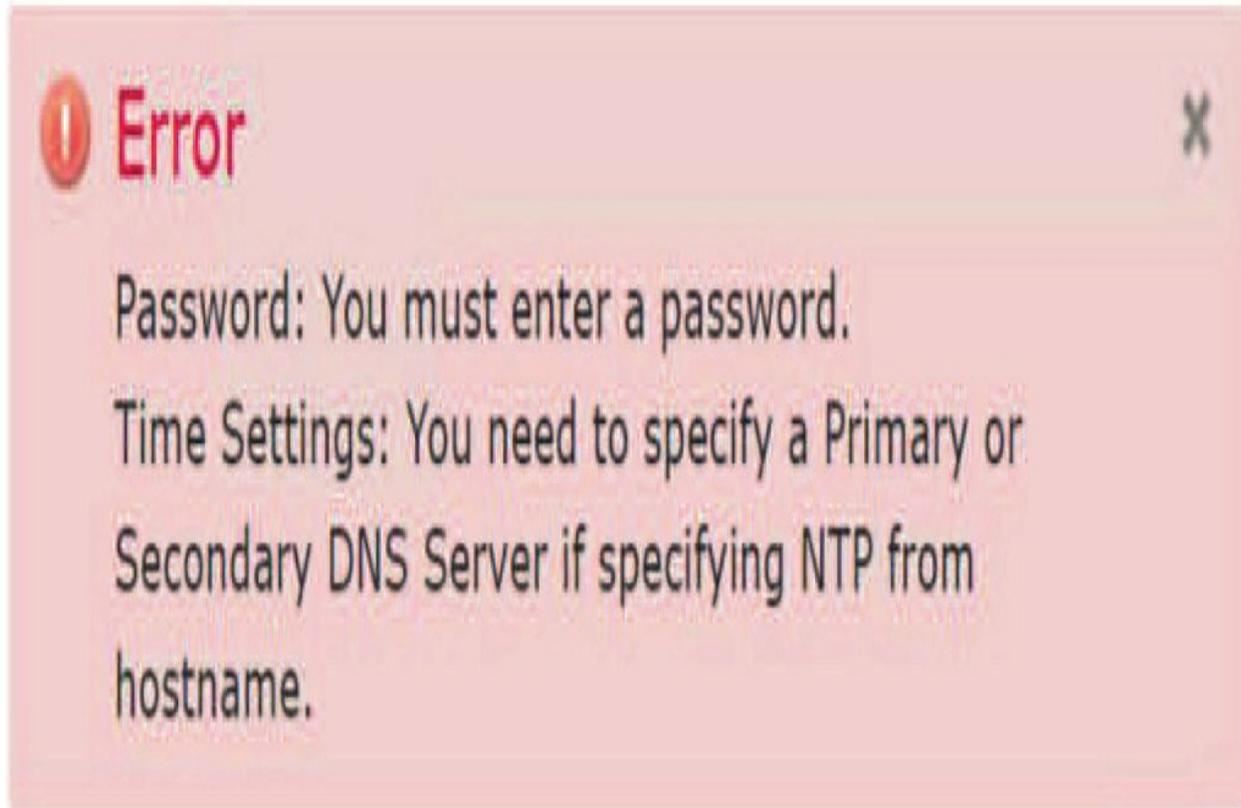
Current Time

2019-09-02 16:49

Set Display Time Zone

[America/New York](#)

To get to this page, you either had to use the default of 192.168.45.45 or change the IPs during initial setup to get into this GUI, and if you didn't assign any other options in the initial setup configuration, you'll need to configure a password, and two DNS servers as well, if you didn't change the preconfigured NTP server to an IP address, or you'll receive the error shown here.

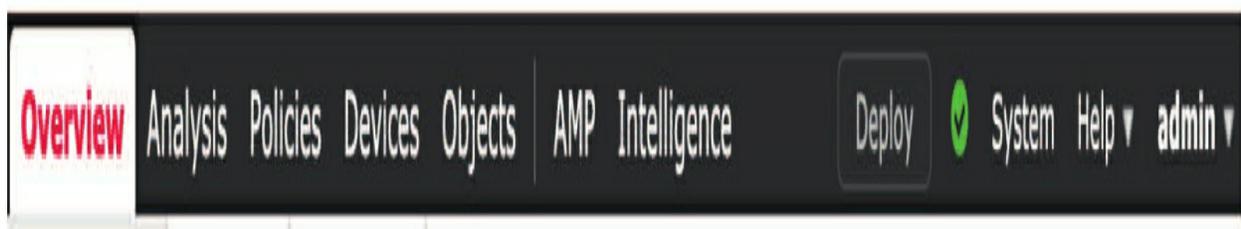


Finally, after you finish your minimum configuration, scroll to the bottom and accept the EULA and click Apply.

Navigation Overview

The Firepower menu system is fairly intuitive—especially once you’ve been using it for several years! There is a top menu bar with sub-menu items. There are also various tabbed windows for access to different types of settings. One thing you will get used to quickly is hovering. Most of the time, you will want to navigate down to a submenu item by hovering over the menus as they appear rather than start out by clicking on the top-level items.

Top-level menu items are divided into two groups on the right and left.



The items on the right—Deploy, System, Help, and <user-name>— are focused on “operational” tasks. Here you’ll find items such as health, licensing, user management, system integration, backups, task monitoring, and so on.

However, starting in code 6.5, there is a new Light theme available; the right-side operational menu is shown here (the only difference is the look).

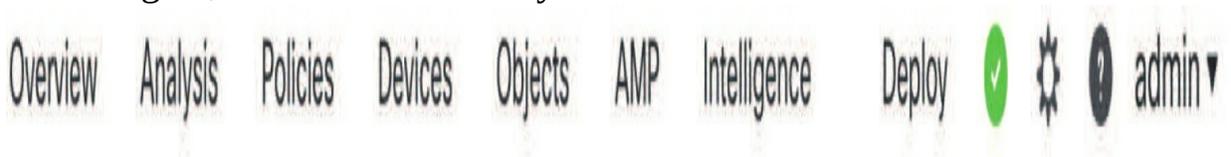


They are the same, except the System menu has been changed to a gear icon. Now, the items on the left—Overview, Analysis, Policies, Devices, Objects, AMP, and Intelligence—are where you will go to configure the actual detection and prevention settings as well as analyze events.

Generally speaking, you will spend more of your time on the left side of the menu bar after the first few days of system setup.



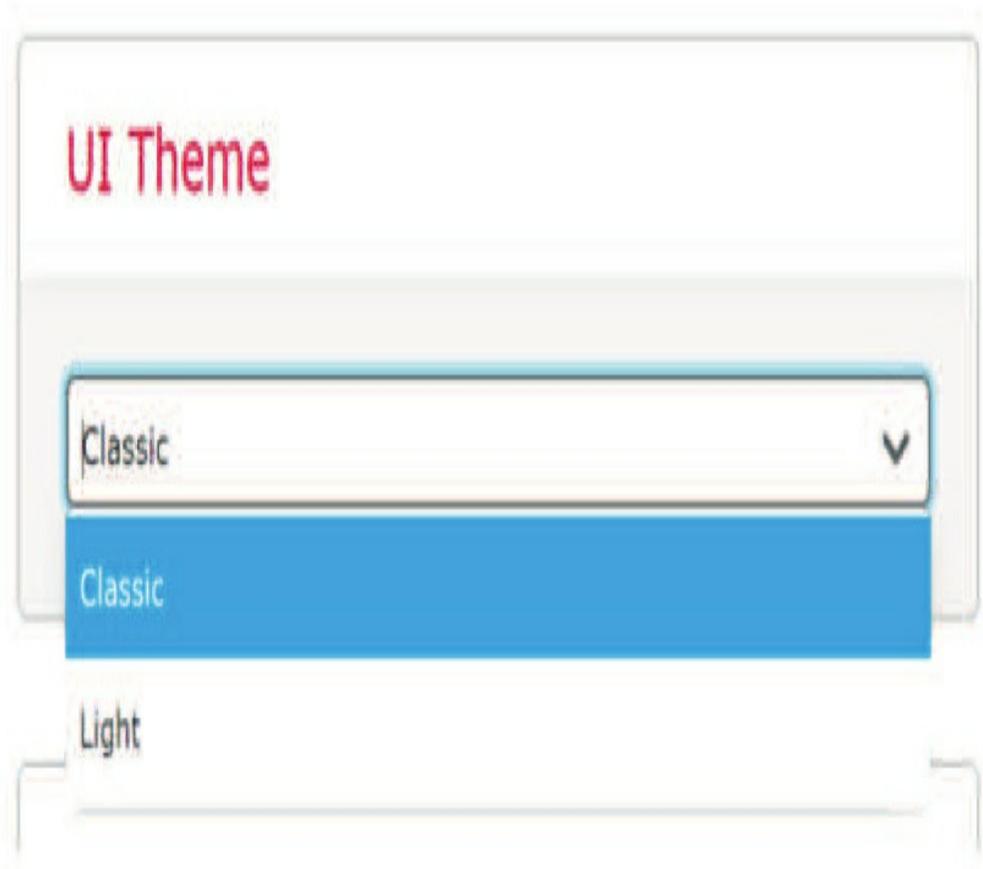
The menus on the left and right are shown in the next figure with the Light theme. Again, no differences really.



To change from one theme to the other, go to **user>User Preferences**, as here.



From User Preferences, you can leave it at the default of Classic or choose the new Light theme.



Firepower System Configuration

Let's start configuring the FMC by looking at the FMC configuration settings. You can find these by navigating to **System>Configuration**. This is where you will find all of the items that apply to the FMC itself.

When you arrive at this screen, you will find a list of the configuration items running down the left side of the screen. Selecting an item will display the screen to modify these settings.

There is one very important note to remember: the settings here go into effect

right away. Instead of several changes being made and then applied, every setting you change is applied to the FMC as soon as you click the Save button. Something to keep in mind is that navigating between links on the left does not automatically save your work.

For some of the settings, you will receive a warning pop-up if you click away without saving, but for some you will find that the information you modified will be silently discarded.

Key tip—use the Save button! Anytime you change a parameter, you must click Save before navigating away from the page.

Information

When you first navigate to the System Configuration page, you will find FMC information as the default page that opens up. This includes the name (firepower is default), model, serial number, etc.

Configuration Users Domains Integration Updates

- Access List
- Access Control Preferences
- Audit Log
- Audit Log Certificate
- Change Reconciliation
- DNS Cache
- Dashboard
- Database
- Email Notification
- External Database Access
- HTTPS Certificate
- Information**
- Intrusion Policy Preferences
- Language
- Login Banner
- Management Interfaces
- Network Analysis Policy Preferences
- Process
- REST API Preferences
- Remote Storage Device
- SNMP
- Shell Timeout
- Time
- Time Synchronization
- UCAPL/CC Compliance
- User Configuration

Name	firepower
Product Model	Cisco Firepower Management Center for VMWare
Serial Number	None
Software Version	6.5.0
Operating System	Cisco Fire Linux OS
Operating System Version	6.5.0
IPv4 Address	10.11.10.175
IPv6 Address	Disabled
Current Policies	Health Policy Initial Health Policy 2019-08-27 20:07:08
Model Number	66

Save Refresh

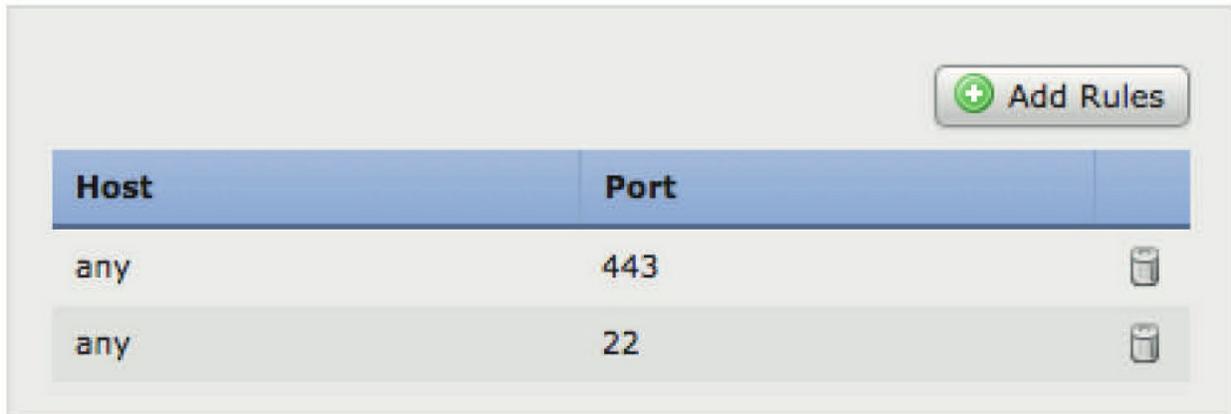
All you can really change here is the FMC name.

You can use whatever name you want, and your name can even have spaces, if desired. If you need quick access to other information, like version or serial number (shown with hardware FMCs only), this is a good page to remember.

In the output of my hardware FMC shown above, notice that the only difference is that the serial number shows and in the virtual FMC there isn't one. You can also get this information from Help.

Access List

This setting configures the local host (iptables) firewall on the FMC. By default, any source IP address is allowed to connect on port 443 or 22 for management.



Host	Port	
any	443	
any	22	

To restrict access, simply click the **Add Rules** button and enter the IP/CIDR address to you want to allow. Then check one, two, or all of the three ports—SSH, HTTPS, or SNMP. You will do this as many times as needed to add each source IP or CIDR block.

In the figure below, I am adding the 10.0.0.0/24 network and allowing SSH and HTTPS access.

IP Address

Port SSH HTTPS SNMP

When you are finished, you will have your allowed IP ranges as well as the “any” rules, as shown in the next figure.

 Add Rules

Host	Port	
any	443	
any	22	
10.0.0.0/24	22	
10.0.0.0/24	443	

Now, we haven't actually restricted access yet since the "any" rules are still there. The last step is to delete these rules, which will then restrict access to just the remaining entries (**note that you can shoot yourself in the proverbial foot here**). If you remove the "any" rules and have not added the IP address you are currently logged in from, you will immediately lose access to the FMC! Once you click the Save button over on the right side of the screen, these iptables entries will go into effect immediately. You have been warned! If you're sure your entries are correct, then click Save and move on to the next setting.

Also, be advised that you *must* put in your NMS station IP address set as SNMP here if you want to send traps from the FMC to the NMS station! However, you must set the SNMP configuration further down before you can configure the NMS station IP address here. Most people miss this setting.

Access Control Preferences

These are simple settings that control rule comments. There are three settings: Disabled, Optional, and Required.



- Disabled (default): When you edit an access control rule, nothing happens upon saving the rule.
- Optional: When you save an access control rule, the comment dialog appears, and you can either add a comment or use the Cancel button to decline.
- Required: When you save an access control rule, the comment dialog appears, and you must enter a comment.

Comments inserted this way are saved with the rule in the Access Control policy.

There is also a check box for writing changes to the audit log, which is checked by default. This increases the level of detail when an Access Control policy is updated. If you don't need this level of detail, you can uncheck this.

Audit Log Certificate and Audit Log

These two groups of settings are related, so we'll talk about them together. The Audit Log settings allow you to stream the FMC audit log to an external system.

The audit log is updated as users log in and use the FMC Web UI to administer the system. User activity of various types is recorded in the audit log, including each page visited, policy changes, and user account activity.

All of these entries are saved in the FMC database automatically. This setting allows you to send these to an external system as well.

Send Audit Log to Syslog

Disabled

Host

Facility

USER

Severity

INFO

Tag (optional)

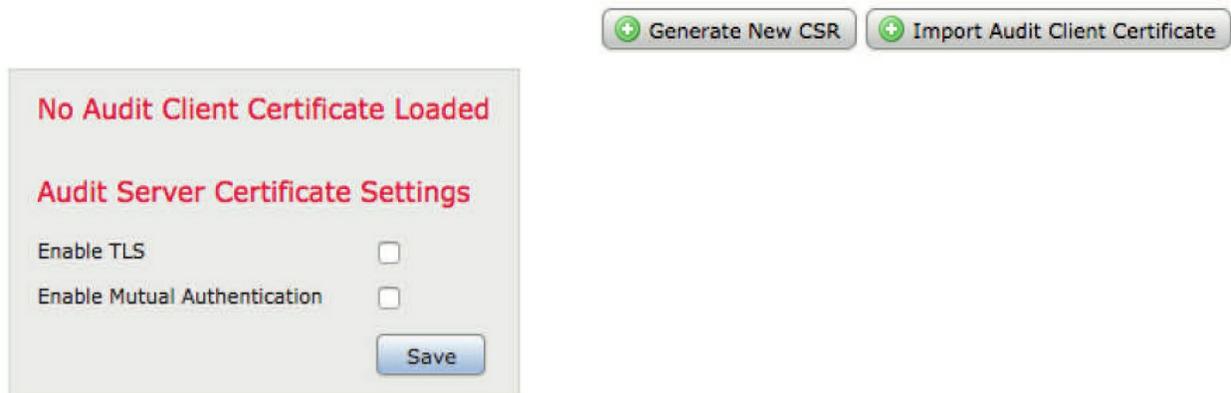
Send Audit Log to HTTP Server

Disabled

URL to Post Audit

You can send your audit entries via syslog (UDP/514) or HTTP.

If using syslog, you can also add an optional tag, which will accompany each event. You can also log via plain HTTP or HTTP with SSL. The Audit Log Certificate settings are where you would generate a Certificate Signing Request (CSR) to be signed by your certificate authority.



Once signed, you would return to this page and import the signed client certificate. You can also optionally specify a certificate revocation list (CRL) URL.

Change Reconciliation

This group of settings allows you to send a daily Change Reconciliation report with any configuration changes to the FMC in the previous 24 hours. Note that to use this, you must first configure a mail relay, which we will see in a moment. Once you do so, you can enter the email address you want the report sent to. Once you have sent at least one report, you will also see a Resend Last Report button on this page that allows you to do pretty much what it says—resend the last report.

Enable	<input type="checkbox"/>
Time to Run	0 : 00
Email to	<u>Not available. You must set up your mail relay host.</u>
Include Policy Configuration	<input type="checkbox"/>
Show Full Change History	<input type="checkbox"/>
<input type="button" value="Save"/>	

I absolutely setup *all* my customers to run this report every night and have it emailed for storage and so should you. Why? Because of the information it provides.

It's extremely important information to know and understand if you are restoring from a backup. The software version, Snort version, and VDB must all be the same on the FMC you want to restore to and the FMC the backup was taken with. In addition, if you are installing a code less than 6.5, the FMC must be the same model as well.

Here is the change reconciliation front page.
Daily Change Reconciliation Report

System Name		FCM20.sfgtc.local
Model		Cisco Firepower Management Center for VMWare
Software Version		6.4.0.3 (build 29)
OS		Cisco Fire Linux OS 6.4.0 (build2)
Snort Version		2.9.14.3 GRE (Build 15301)
Rule Update Version		2019-08-19-001-vrt
Rulepack Version		2286
Module Pack Version		2588
Geolocation Update Version		None
VDB Version		build 309 (2019-02-08 19:48:25)
IPV4		10.11.10.205
IPV6		
System Policy		Local System Configuration
Health Policy		Pod20-FTD_Health
Time Period		2019-08-20 08:00:00 - 2019-08-21 08:00:00

Since most people won't always look and understand their Snort update version and VDB version, this report will store that information for you, and then you can match this info to a backup.

Console Configuration

This setting is only on hardware FMCs and allows you to use a Linux system console for remote access on the FMC supported systems by either the VGA port (which is the default) or the serial port on the physical appliance.

Notice the Console Configuration menu item shown in my 2500 FMC output here.



However, there is still one more setting here.

Console Configuration

Console VGA Physical Serial Port Lights Out Management

Lights-Out Management Settings

IPv4 Settings

Configuration

DHCP 

IP Address

192.168.227.177

Netmask

255.255.255.0

Default Gateway

192.168.227.1

Lights-Out Management Users

Username	Status	Action
admin	Access will be granted on next login	 Edit

Save

Refresh

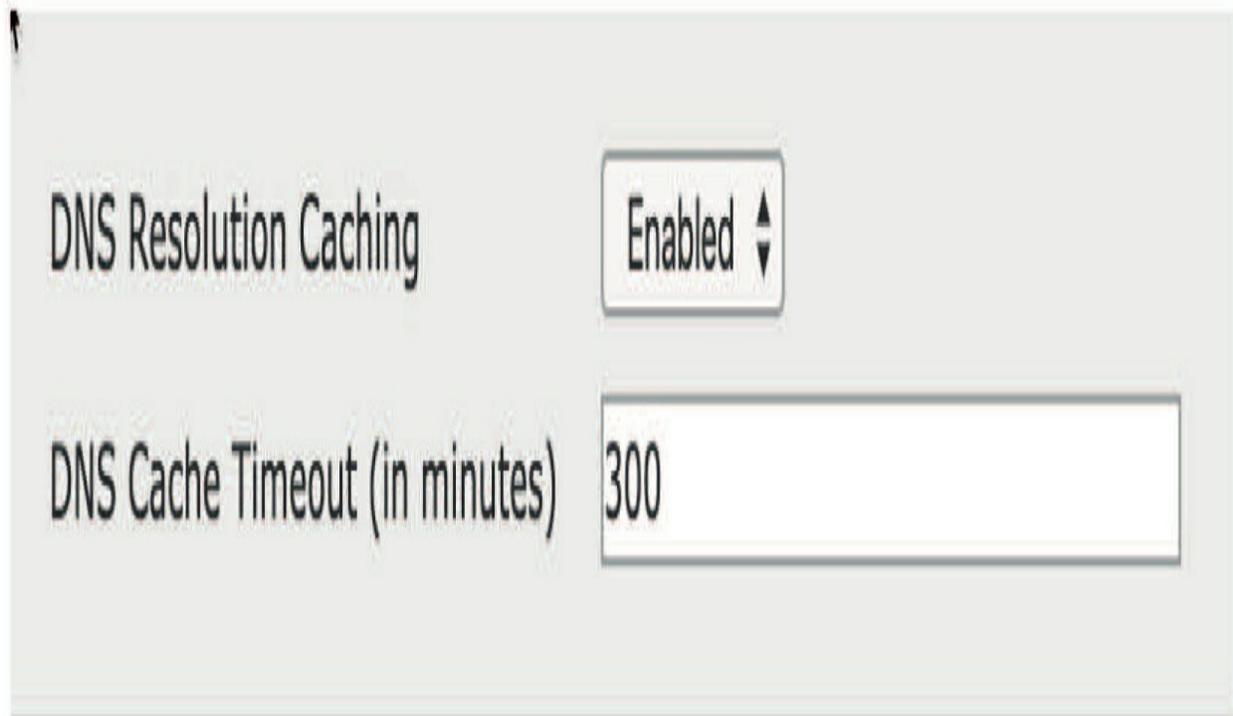
On supported physical-hardware-based Firepower systems, you can use Lights-Out Management (LOM) on the default Ethernet 0 management interface on a Serial Over LAN (SOL) connection to remotely monitor or manage the system without logging into the management interface of the system.

This is a very limited console, but it does provide a CLI as an outof-band connection. You must both enable LOM for the FMC system using the IP and provide the user that you want to manage the system. Admin is the default user.

After you enable the system and the user here, you then still need to use a third-party Intelligent Platform Management Interface (IPMI) utility to access and manage your system.

DNS Cache

The DNS Cache settings control the length of time hostnames will be cached within the UI for the various event types.



The screenshot shows a configuration panel for DNS Cache. It contains two settings:

- DNS Resolution Caching**: A dropdown menu currently set to **Enabled**.
- DNS Cache Timeout (in minutes)**: A text input field containing the value **300**.

This prevents the FMC from constantly performing DNS lookups every time you load an event view page. Shorten this default if you have a more dynamic environment; otherwise, the default of 300 minutes (5 hours) is generally just fine.

Dashboard

This contains one setting, which honestly is a throwback to a time when some low-powered appliances could not handle displaying all the Dashboard widget types. This defaults to being enabled—you should leave it that way.

Database

The Database page is where you will configure the maximum size of the various databases within the FMC. Each of the databases listed is a “circular” database, meaning it keeps the most recent n number of events. New events will be added until the threshold is reached, after which older events will be purged to make room, keeping the database at the maximum record count.

Most databases default to keeping the most recent 1 million events. A quick look down the page shows there are exceptions, such as connection summaries and audit events.

The two databases you are most likely to ever change are those holding intrusion events and those holding connections. This is because a lot of folks want to keep a long history of intrusion events. That is usually possible because, once tuned, your system should not generate a large number of intrusion events.

However, when it comes to connection events (logging of the rules), even though a long history may not be needed, the default number may only represent a few hours or even a few minutes on a busy network. Because of this, the default number is often increased.

Here I show the default database screen. To find your maximum connection events, add a large number and the system will provide you with what is supported. The 50,000 is for the virtual FMC. All other FMCs handle more.

Error ✕

Connection Events Database must be between 0 and 50,000,000, inclusive.

- Access List
- Process
- Audit Log Certificate
- Audit Log
- Login Banner
- Change Reconciliation
- DNS Cache
- Dashboard
- Database**
- External Database Access
- Email Notification
- Access Control Preferences
- HTTPS Certificate
- Information
- Intrusion Policy Preferences
- Console Configuration

Intrusion Event Database

Supported Platforms Firepower Management Center

Maximum Intrusion Events

Discovery Event Database

Supported Platforms Firepower Management Center

Maximum Discovery Events (0 = do not store)

Connection Database

Supported Platforms Firepower Management Center

Maximum Connection Events (0 = do not store)

Maximum Security Intelligence Events

Connection Summary Database



For example, next I show the output on my 2500 FMC, providing the maximum number of connection events.



However, understand that this maximum number is connection events and security intelligence events combined.

We can spend a lot of time discussing the factors influencing these settings, but suffice it to say, a larger number means more history. That is, you can look back further in time for events. In the end, the amount of history is related to the event rate and the database capacity. You will have to experiment with these settings and balance the need for historical event storage with the performance of your FMC. The more events you store, the further back in history you can look but the slower your FMC will perform these searches. There is no “right” answer here; the history available for the various event types will be determined by factors such as these:

- Access control policy connection logging settings
 - Volume of network traffic
 - IPS policy
 - IPS rule tuning
 - FMC model
-
- External logging capability (maybe you can store the events externally rather than on the FMC)

For now, let's look at how the FMC model impacts these settings. Each FMC provides a specified amount of event storage and processing power. A large appliance such as the MC4600 can store and process many times more events than an MC1600.

Understand that just because an MC4500/4600 can store 1 billion (yes, that's a *b*) connection events doesn't mean you should automatically pick that! Although you can store that many events, even a beefy MC4600 will *never* return a query of that many rows. It will timeout before the query finishes. Cisco recommends—as do I—that you use care in increasing the defaults for any of your databases. Don't be greedy; just as with the buffet, take only what you can eat (use). Your FMC will thank you and your admins will not be frustrated with slow query performance.

If you are performing an initial configuration of your FMC, you should probably leave these numbers as is for now. Once you start collecting events, you will have a better idea if you need to come back here and change anything.

Email Notification

This is another group of fairly simple settings. If you want your FMC to send email notifications (and everyone does), then you have to configure an SMTP mail relay. Once you configure the settings here, you can use the Test Mail Server Settings button to send yourself a test email.

If you are just setting up a test FMC and don't have an enterprise email server, you can use gmail; however, it's possible that you may get an email from Google if you don't have a business gmail account.

The settings for testing would be smtp.gmail.com, port: 587 or 465, TLS, and use your username and password for authentication. In the settings, the from address is not important, but your authentication address is. A sample configuration is shown here:

- Access List
- Access Control Preferences
- Audit Log
- Audit Log Certificate
- Change Reconciliation
- DNS Cache
- Dashboard
- Database
- Email Notification**
- External Database Access
- HTTPS Certificate

Mail Relay Host: smtp.gmail.com

Port Number: 587

Encryption Method: TLS ▼

From Address: fmc20@acme.com

Use Authentication:

Username: tlammler@acme.com

Password:

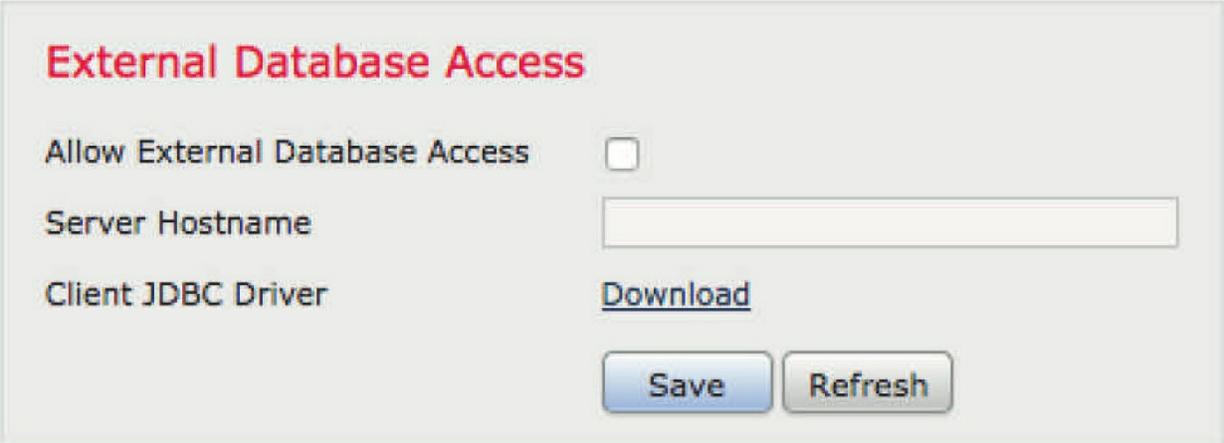
Test Mail Server Settings Message sent



External Database Access

This group of settings controls access to the FMC database for a custom reporting tool. Most people will not need this, as the built-in report generator is quite capable. However, if you want to use an external tool such as Crystal Reports or iReport Designer, then read on.

To enable external access, you will enter the hostname or IP address of your report server (the host that will be connecting to the FMC to query the database). You also must add the IP address to the access list. You will also need to download the client JDBC driver, which is provided as the file `client.zip`. This archive contains a Java query app and client certificate installation tool. You will install these on your reporting host to provide the ability to query the database.



The screenshot shows a configuration panel titled "External Database Access" in red text. It contains three main settings: "Allow External Database Access" with an unchecked checkbox, "Server Hostname" with an empty text input field, and "Client JDBC Driver" with a blue "Download" link. At the bottom of the panel are two buttons: "Save" and "Refresh".

Rather than me spending multiple pages covering this, your best source of information is the *Firepower System Database Access Guide*, which is available on the Cisco.com site. The FMC External DB configuration page is shown above.

HTTPS Certificate

This is where you can replace the self-signed HTTPS certificate with one signed by your own certificate authority.

It includes a button to generate a Certificate Signing Request (CSR) and another one to import the signed certificate.

Note that the Enable Client Certificates check box is used to require clients (like you) to use a known certificate, which is stored on the server. People have lost access to their FMC by unknowingly checking this box.

Don't do this unless you have uploaded your client certificate(s) to the FMC and know what you are doing!

 Generate New CSR

 Import HTTPS Server Certificate

Current HTTPS Server Certificate

Subject	commonName firepower	countryName US	organizationName Cisco Systems, Inc	organizationalUnitName Intrusion Management System
Issuer	commonName firepower	countryName US	organizationName Cisco Systems, Inc	organizationalUnitName Intrusion Management System
Validity	Not Before Jun 30 18:57:55 2018 GMT	Not After Jun 30 18:57:55 2038 GMT		
Version	3			
Serial Number	BFB3075E3FC25197			
Signature Algorithm	sha256WithRSAEncryption			

HTTPS Client Certificate Settings

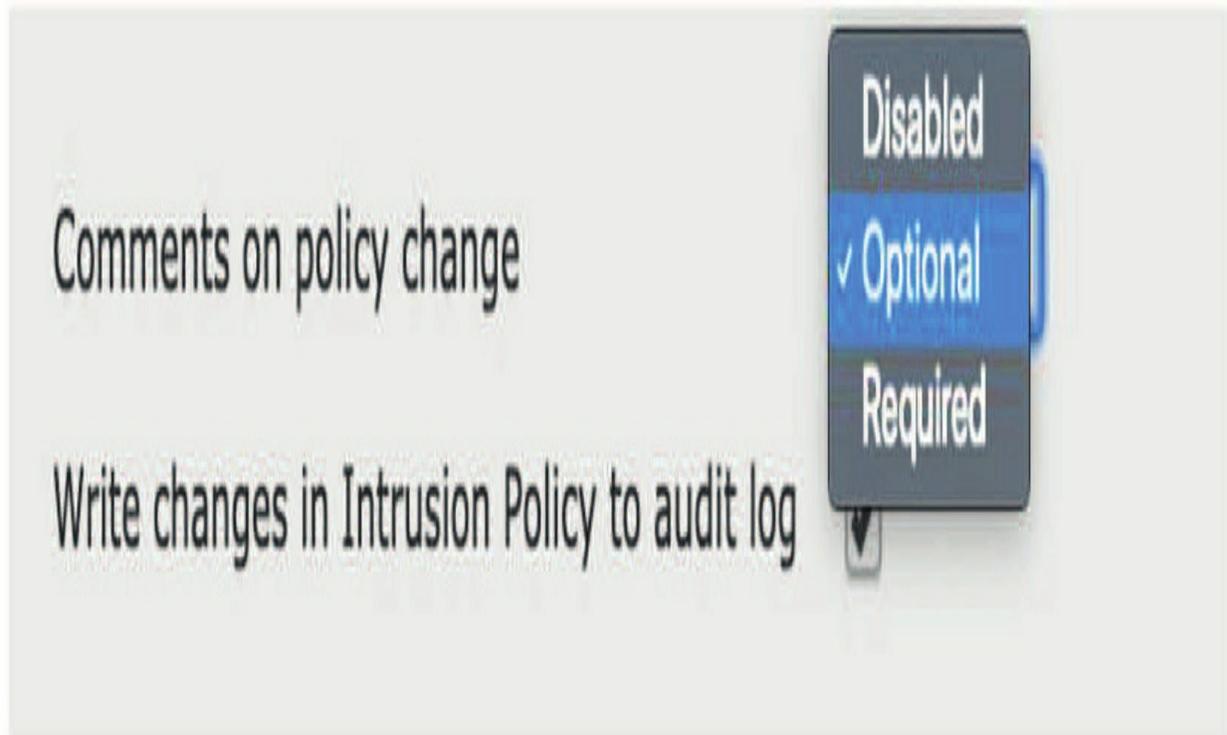
Enable Client Certificates

Save

Intrusion Policy Preferences

This group of settings controls comment and logging behavior for Intrusion policies. The Comments on Policy Change setting has three options: Disabled, Optional, and Required.

When you edit an Intrusion policy, this setting controls what happens when you commit the changes.



- Disabled: Nothing happens; your policy is committed.
- Optional (default): A comment dialog appears where you can enter comments or click the Cancel button.
- Required: A comment dialog appears, and you must enter some text before committing the policy.

Admins sometimes think this is a good idea at first to require comments. Later, they tire of constantly being required to insert comments upon changing the Intrusion policy and revert to Optional or Disabled. Any comment entered here is inserted into the audit log.

There is also a check box for writing changes to the audit log, which is checked by default. This increases the level of detail and can mean hundreds of audit log entries when an Intrusion policy is updated. If you don't need this level of detail, you can uncheck this.

Language

English, Japanese, and three others—not much more to say about that, other than if you can't read what language it is you probably shouldn't use it.



However, since work can be tedious at times, and to have fun at work, just before you leave on vacation, change the language configuration to one of these other languages. What a laugh riot that will be!

Login Banner

The login banner is a simple concept. This is text that will display on the Web UI and the console prior to users logging in. Generally, this is some verbiage provided by your legal department and is standardized across the organization.

Enter your text and then click Save. After you log out and back in, a pop-up window will appear with the banner message asking if you accept this message.



Yes, I am aware that I wrote a ridiculous banner message, but when you're writing a Cisco security book, you have to amuse yourself sometimes. I do this probably more than other authors for sure.

If you now log in to the shell of the FMC, the banner will also show at the top.

However, be advised that before the 6.5 code was released, the banner message would only show in the window where the pretty picture was displayed, and it covered up the picture:

This isn't a problem; I just wanted you to see the differences.

```
This is Todd Lammle's vFMC.
```

```
Only people with big PO's can come into my FMC
```

```
firepower login: admin
```

```
Password:
```

```
Copyright 2004-2019, Cisco and/or its affiliates. All rights reserved.
```

```
Cisco is a registered trademark of Cisco Systems, Inc.
```

```
All other trademarks are property of their respective owners.
```

```
Cisco Fire Linux OS v6.5.0 (build 1036.INTEG)
```

```
Cisco Firepower Management Center for VMWare v6.5.0 (build 39)
```

```
>
```

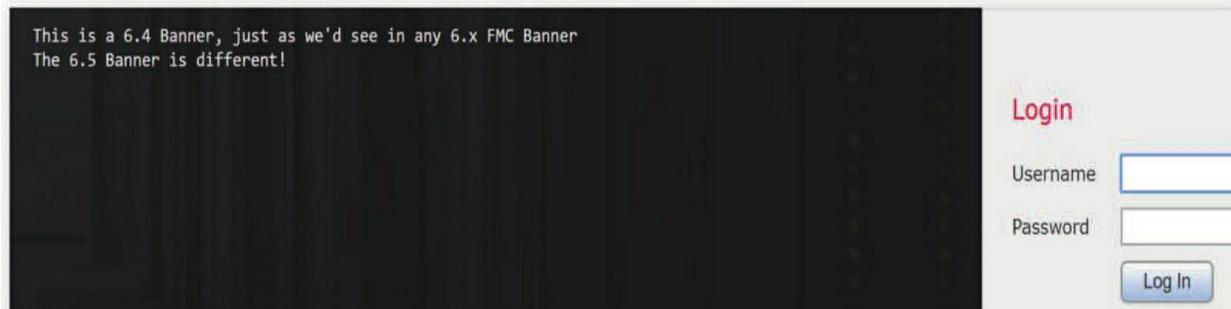
Management Interfaces

This is where you will set the network connection details for the management interface(s), and this also is the longest section in the chapter since it is the

most important and has the most possible configurations, depending on what type of FMC you have.



For technical/system quest
or call us at 1-800-55



By default, the Firepower Management Center manages all devices on a single management interface named eth0, and it uses this eth0 interface for initial setup, HTTP access for administrators and management of devices as well as the smart licensing server and to download updates. Eth0 is also the interface used to log into the FMC as an administrator.

Devices managed by the FMC (7000/8000 appliance or FTD) also have a single management interface for performing the initial setup as well as communicating with the FMC.

Let's start with the items you are most likely to look at first in a virtual FMC. Since you are connected with the Web UI, I know you already have at least some of the network information configured.

Your screen may look similar to the screen shot shown next. Notice the eth0 interface on the virtual FMC output here; I'll come back to that in a minute.

▼ Interfaces

Link	Name	Channels	MAC Address	IP Address	
	eth0	Management Traffic Event Traffic	00:50:56:A7:A1:1A	10.11.10.175	

▼ Routes

IPv4 Routes

Destination	Netmask	Interface	Gateway	
*			10.11.10.1	

IPv6 Routes

Destination	Prefix Length	Interface	Gateway	
-------------	---------------	-----------	---------	--

▼ Shared Settings

Hostname	<input type="text" value="firepower"/>
Domains	<input type="text"/>
Primary DNS Server	<input type="text" value="10.11.11.250"/>
Secondary DNS Server	<input type="text" value="1.1.1.1"/>
Tertiary DNS Server	<input type="text"/>
Remote Management Port	<input type="text" value="8305"/>

Now you can set the hostname, search domain and DNS servers. You may want to enter a tertiary DNS server as well. This is really all you need for a fully functional FMC in many environments.

However, as long as traffic can reach the FMC and the FMC can communicate with the managed devices and Cisco's Internet-based update servers, you can use the FMC in your network.

ICMPv6

After the Shared Settings section, you'll find the ICMPv6 section. This is a newer group of settings found in the latest code for allowing ICMPv6 with the management of the FMC instead of just IPv4, and they're enabled by default.



Proxy

Some environments will require configuration of the Proxy settings to allow the FMC to download the various security feeds and updates. Firepower supports unauthenticated or authenticated proxies

To Split or Not to Split

I started to discuss this in chapter 1: If your FMC has more than one management interface (found on hardware FMC's), you also have the option to use multiple ports and decide what type of traffic to use each one for. Clicking on the pencil icon by one of the interfaces will bring up this dialog:

▼ Interfaces

Link	Name	Channels	MAC Address	IP Address	
	eth0	Management Traffic Event Traffic	A0:93:51:52:C9:E2	172.16.10.20	
	eth1	Management Traffic Event Traffic	A0:93:51:52:C9:E3		
	eth2	Management Traffic Event Traffic	90:E2:BA:FB:BE:C8		
	eth3	Management Traffic Event Traffic	90:E2:BA:FB:BE:C9		

▼ Routes

IPv4 Routes



Destination	Netmask	Interface	Gateway	
*			172.16.10.1	

In the figure, the output from my hardware FMC shows that I have four interfaces available for management.

Understand that regardless of what FMC you have, the eth0 interface is the default and must be used for all FMCs; however, eth1 – 3 can be used on hardware FMC. For example, if you have a large number of devices to manage on your FMC, adding more management interfaces can improve throughput and performance.

Nicely, you can also configure these additional interfaces on the same management network or on different management networks.

Depending on your FMC model, you may have more than one interface available. All FMCs will default to a single 1 GB copper management port. The FMCs 4500/4600 and 2500/2600 also have optional 10 GB fiber ports

The eth0 interface shown above is called management0, and this is the internal name of the Management 1/1 interface.

Note that only the FMC eth0 interface supports DHCP IP addressing. Other management interfaces only support static IP addresses.

You can click on the pencil of an interface and edit the interface. You can set some of the properties of the physical interface here, but you also have check boxes for the channels the interface will use, as shown in below.

By default, both management traffic and event traffic use the same interfaces. However, you can split these if you have more than one interface in your FMC. This allows event and management traffic to each have its own dedicated 1 Gb or 10 Gb port.

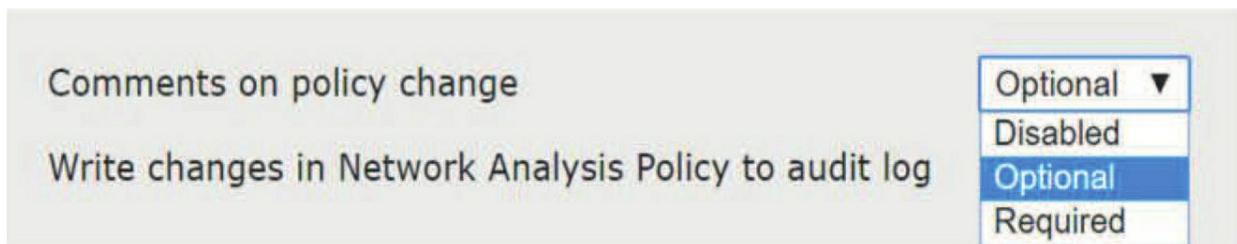
If you are just interested in redundancy in your management interfaces, you can leave both Channels settings checked on your additional interface(s). This will use both interfaces for both traffic types and also provide redundancy in case one of the network paths goes down.

On a practical note, even on a busy MC4600, a single 1 Gb management interface will not normally represent a bottleneck even at the appliance's

maximum event rate (20,000 events/sec). Writing to the disk is still the choke point that limits the rate at which the device can process event data.

Network Analysis Policy Preferences

These settings apply to the Network Analysis policy and are identical to the setting for comments in the Intrusion policy, and after your network is up and running you want to enable this so no one can make changes without it recording name/date/time. The next figure shows the NAP Preferences screen. Writing to the audit log is enabled by default.



Process

The Process section is not actually a configuration page. It is a way to shut down or restart the FMC or the web services. These would only be used for maintenance or if directed by Cisco TAC. Clicking the green arrow will run the command selected. Note that if you select Shutdown, you will need some way to start the FMC back up! You will have to either use the lights-out functionality or have someone actually press the power button to start the appliance again.

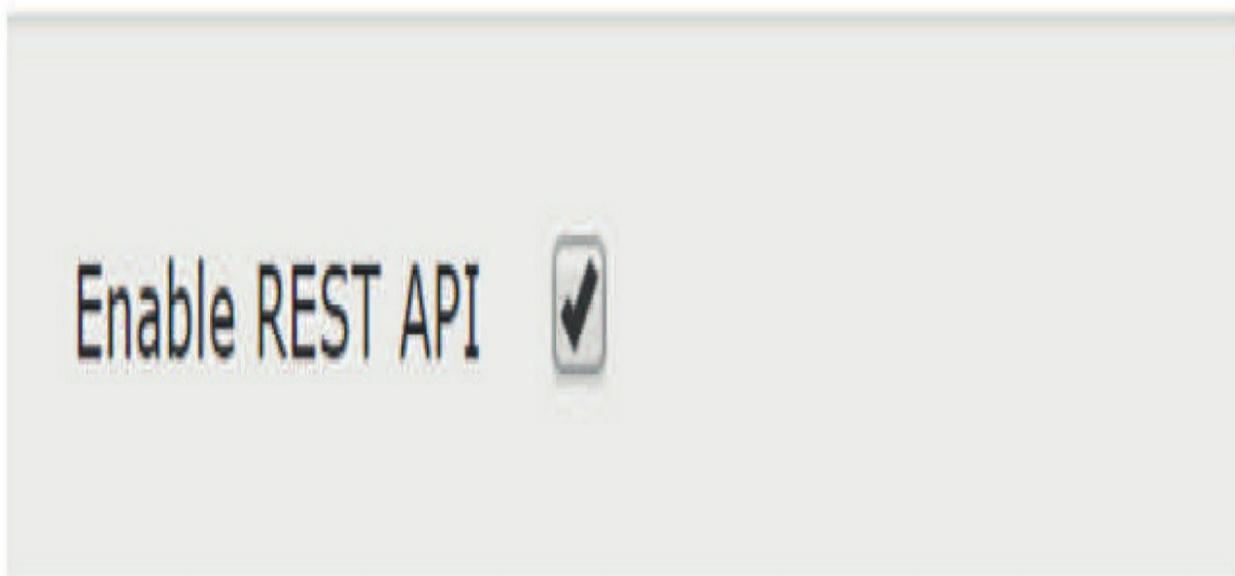
Name	
Shutdown Management Center	➔ Run Command
Reboot Management Center	➔ Run Command
Restart Management Center Console	➔ Run Command

When working with a customer or at your job late at night with TAC on the phone with a ticket for an FMC issue, and they tell you just to restart the FMC, which just restarts the processes, don't do it. Many hours later they finally say, "Hey maybe we should just try and reboot it."

Yes, I do understand that some people don't want to reboot because they want to spend all night finding the "root cause" of the issue you're troubleshooting, but most of the time when you've restarted the processes and all seemed fine, it still has a problem in the morning and you'll have to reboot at that point anyway, which is after you stayed up all night of course!

REST API Preferences

This is a feature introduced in version 6.1 allowing a lightweight interface for third-party applications to view and manage configurations using a REST (representational state transfer) client.



By default, the FMC accepts REST clients; to disable this feature, uncheck the Enable REST API check box.

Remote Storage Device

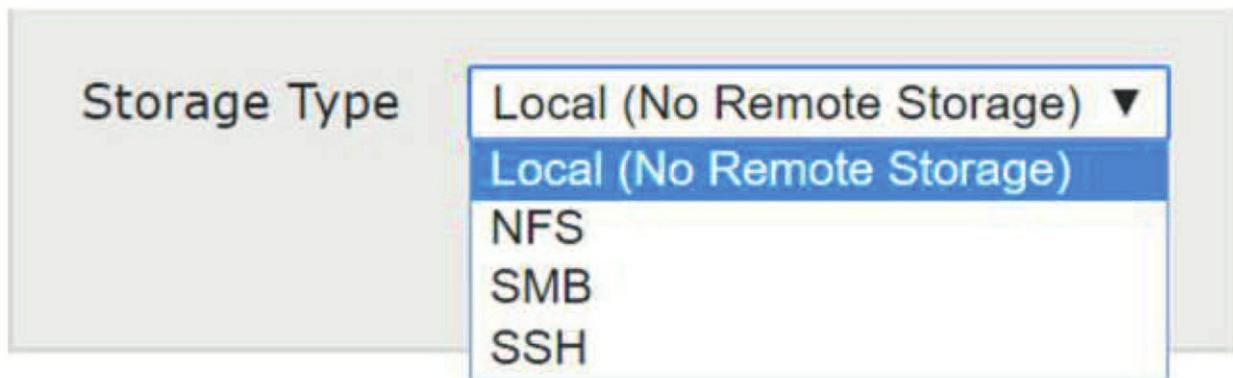
Remote storage can be configured to store backups and reports. This is a fairly common configuration and allows backups to be written off-box as

soon as they occur. This saves the step of copying the backup files off of the FMC and ensures that you are not storing the backups on the appliance you are backing up.

It allows automatically generated reports to be stored on a remote file share where they can be made accessible to a wider audience without worrying about granting access to the FMC.

There are three types of remote storage supported as shown here:

- NFS – Network File System
- SSH – Secure Shell (SCP)
- SMB – Server Message Block (Windows) file systems

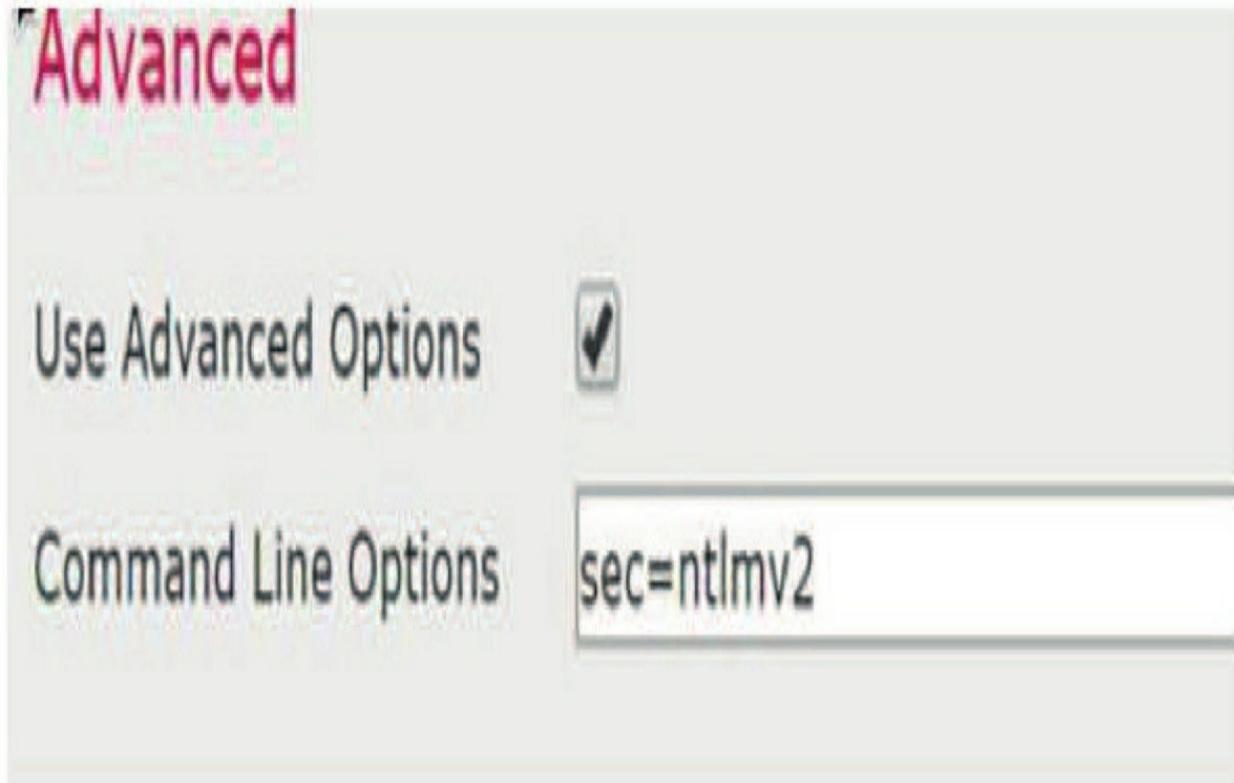


Note that this is *not* for event storage; it is only for backups and reports.

There is (currently) no way to store the FMC event databases anywhere but on the FMC (yes, you can move the logging to another location).

Depending on which one you select, you will see the appropriate dialog to collect the connection information. Understand that if you use the SMB option, it only uses SMBv1 by default.

Use the following settings under the Advanced field and enter the command-line option as shown in the figure to run SMBv2.



After that, the process is fairly simple: you complete the Connection information and use the Test button to test it. When you do, the FMC creates a directory structure on the share where the reports/backups can be stored.

Check the appropriate System Usage check boxes and note the Disk Space Threshold setting. It's important to note that the FMC will not store data to the share if the utilization exceeds the percentage shown. This prevents the FMC from filling up your storage.

When this feature is enabled, reports and backups will now automatically be saved to the remote storage location.

Note: It's possible that this SMBv2 feature has been fixed by the time you are reading this, and the command string shown may not be needed any longer.

SNMP

An SNMP manager can be used to query the FMC for appliance health status. The FMC uses a fairly standard Linux management information base (MIB) to define what types of information can be queried. Remember, this is simply

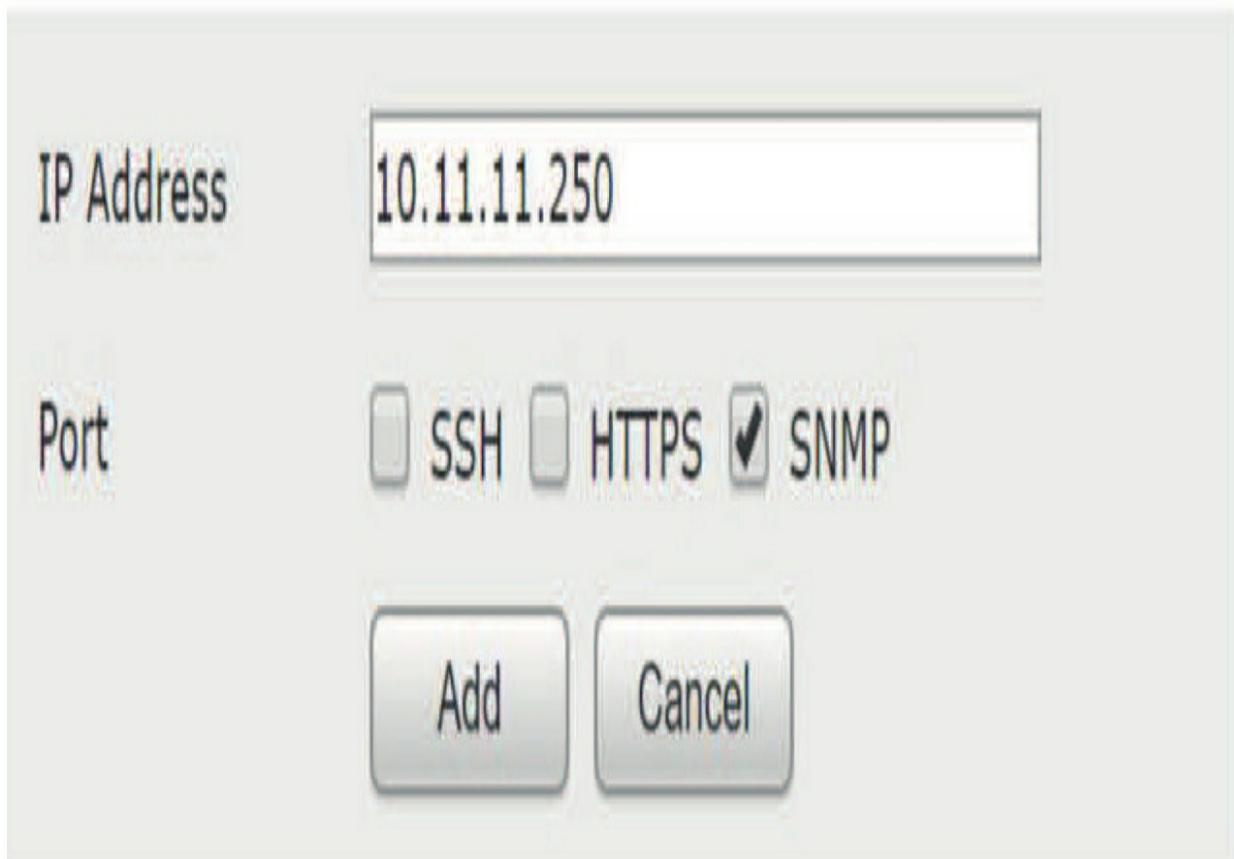
operating system/hardware data; we're not talking about event data here. This next figure shows the SNMP configuration.



The image shows a configuration form with two main sections. The first section is labeled 'SNMP Version' and contains a dropdown menu. The dropdown menu is currently open, showing the following options: '(Disabled)' (selected), '(Disabled)', 'Version 1', 'Version 2', and 'Version 3'. The second section is labeled 'Community String' and contains a text input field that is currently empty.

To enable, first select the SNMP version and then enter the appropriate community string or user information.

But wait, there is more! Don't forget to get back up to the Access List section and add the IP address of your NMS station or the FMC won't send and receive traps and alerts.



The image shows a configuration form with two main sections. The first section is labeled 'IP Address' and contains a text input field with the value '10.11.11.250'. The second section is labeled 'Port' and contains three radio buttons: 'SSH', 'HTTPS', and 'SNMP'. The 'SNMP' radio button is selected. Below the radio buttons are two buttons: 'Add' and 'Cancel'.

Click **Add** and then **Save**.

Shell Timeout

These settings have to do with the browser session and FMC SSH shell timeouts, where the GUI is set to 60 minutes and can be changed all the way up to 1440 minutes (24 hours) The shell is set to never time out by default.

The default settings are shown in the next figure. If you need to modify these based on your organization's policies, you can do so here.

Browser Settings

Browser Session Timeout (Minutes)

Shell Settings

Shell Timeout (Minutes)

Time

The Time settings are not so much configuration items. They control the display of the current time and a history of the last few NTP server connections.

Current Setting Via NTP (based on System Configuration [Time Synchronization](#))

Current Time 2019-09-02 12:31

NTP Server	Status	Authentication	Offset	Last Update
127.127.1.1	Unknown	none	0.000(milliseconds)	6d(seconds)
195.21.137.209	Being Used	none	-0.154(milliseconds)	24(seconds)

Clicking on the Time Synchronization link takes you to our next item, where you can actually configure the time settings.

Time Synchronization

Time is very important in Firepower. It's important that the FMC and managed devices are in sync with each other. Time synchronization is also important for Smart Licensing to function correctly, and to also stop your

Health policy from screaming about a critical alert every 5 minutes.

Because of this, you should always use an accurate NTP time source.

Serve Time via NTP

Enabled ▼

Set My Clock

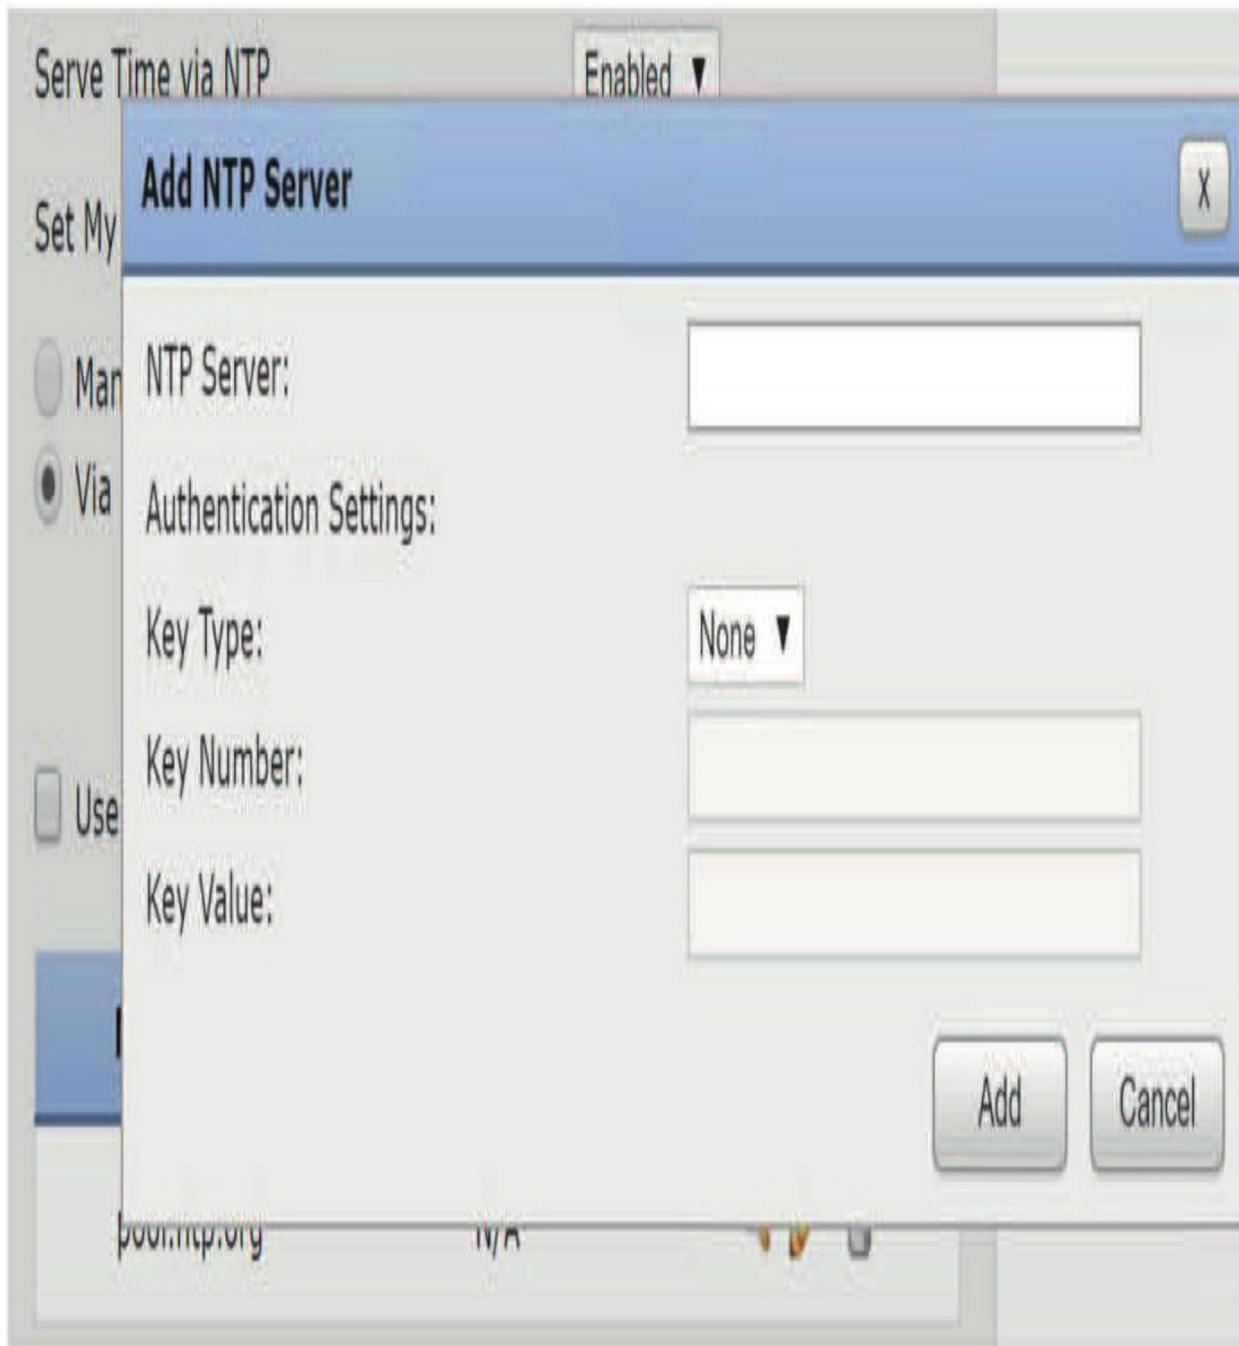
- Manually in Local Configuration
- Via NTP

Use the authenticated NTP server only

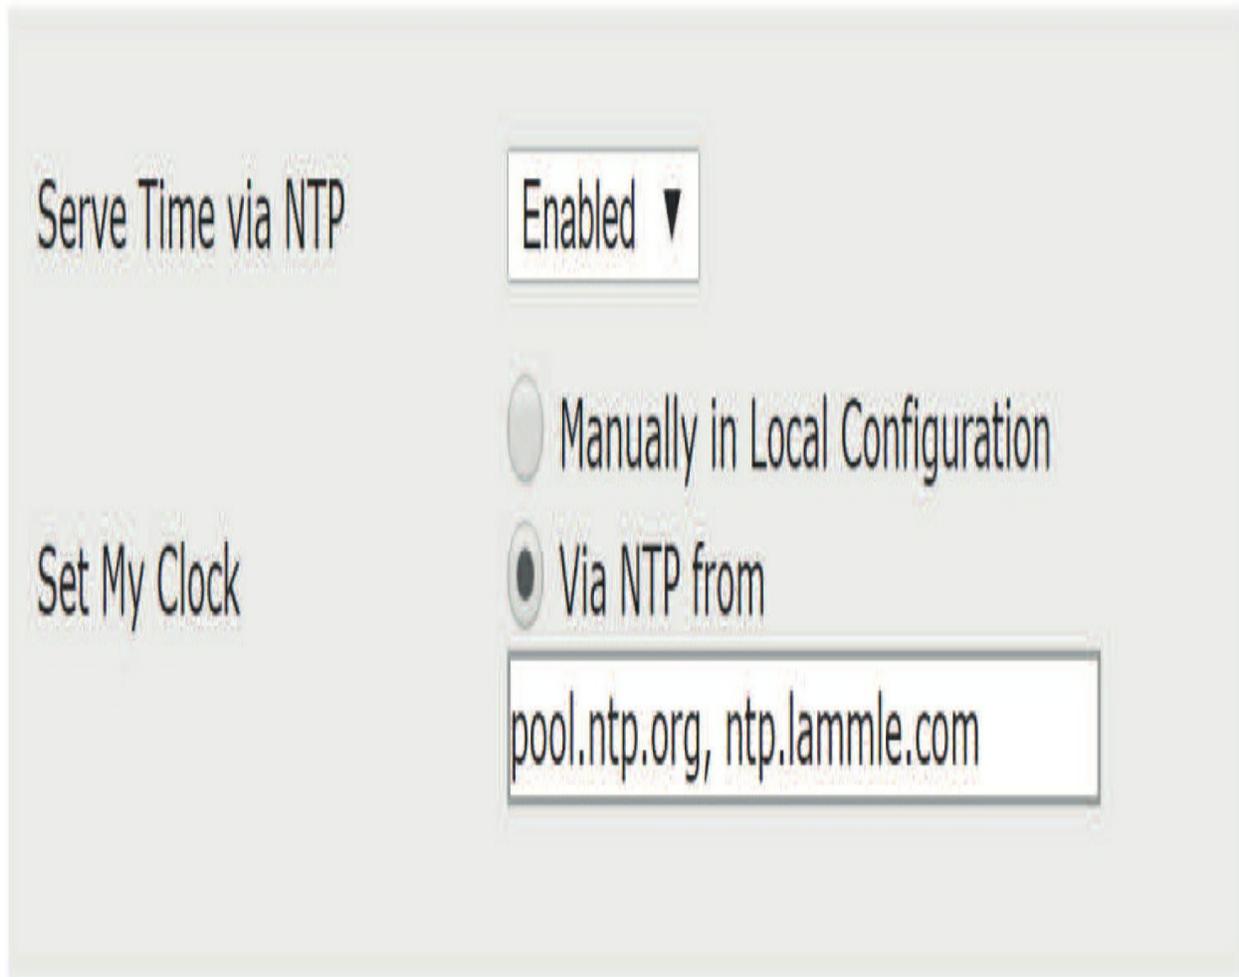
Add

NTP Server	Authentication	Action
0.sourcefire.pool.ntp.org	N/A	
1.sourcefire.pool.ntp.org	N/A	

The options are fairly spartan. Delete the defaults and click Add to enter one or more NTP servers the FMC will use as time sources.



Any “serving” of time by the FMC only occurs over the management connection to the devices and is tunneled over port 8305/TCP. Before 6.5 code, you had a slightly different and classic-looking way of adding multiple NTP servers, as shown here:



The NTP servers would be configured as a comma-separated list.

UCAPL/CC Compliance

These settings apply configurations to the FMC to make it compliant with the US Department of Defense Unified Capabilities Approved Products List (UCACL) or the ISO/IEC Common Criteria (CC).

Changing these will first bring up a warning dialog that they are nonreversible. If you continue, the changes will be made, and the FMC will reboot.

Note that you should *not* make these changes unless you are *sure* you need to. Oftentimes these introduce limitations in functionality or performance in the system, which could ruin your whole day, week, month, or maybe even your

year!

User Configuration

These are new global configuration settings for your FMC admin users that started in code 6.4. Very self-explanatory, as shown here:

Password Reuse Limit

Limit (0 = no limit)

Track Successful Logins

Days (0 = no tracking, 365 = max value)

Max Number of Login Failures

Maximum number of failures before temporary lockout (0 = no lockout, 999 = max value)

Set Time in Minutes to Temporarily Lockout Users

Minutes (0 = no wait time, 1440 = max value, 1 day)

Max Concurrent Sessions Allowed

Maximum sessions for users with Read Only privileges (0 = no limit, 1024 = max value)

Maximum sessions for users with Read/Write privileges/CLI users (0 = no limit, 1024 = max value)

The details for adding users are configured in `System>Users`. Not sure why this configuration isn't found in that same area.

There are still three other options that aren't changed too often:

- VMware Tools
- Vulnerability Mapping
- Web Analytics



Let's take a look at them:

VMware Tools

This is a simple check box to enable/disable VMware Tools. If you install the

virtual FMC on VMware, it will automatically enable this. If you use the virtual FMC, you should probably leave it enabled.

Vulnerability Mapping

This is a setting that always takes a long time to describe, but in the end it's always the same—leave it alone. If you trust me, then move on. If you want the gory details, keep reading.

This has to do with the way the FMC maps vulnerabilities to hosts. This is basically a list of “applications” for which the FMC cannot accurately determine the vendor or version based on the network traffic. I put “applications” in quotes because you'll see that the list contains a lot of things that look like websites. Since websites have a lot of the same characteristics as applications, they are often treated like them within Firepower.

Since vulnerabilities are almost always specific to a given version/ vendor of an application, just because Firepower detects a given application doesn't mean it's the vulnerable one. I always use DNS—the Domain Name System—as my example. Firepower can detect the DNS protocol on a server and deduce that the server is offering DNS services; however, DNS is a pretty lightweight protocol, so there is nothing in the DNS traffic that indicates who the vendor is. Is this a BIND server? Is it Microsoft? In addition, there's no way to tell which version of the vendor's application we're seeing: BINDv7, BINDv9, etc.

You can be sure that the vulnerabilities in Microsoft's DNS services are different than the ones in BIND. Again, between BIND versions these vulnerabilities will be different. So if we're going to map vulnerabilities to a DNS server, we really need to know which vendor and version of DNS it's running. But, since we can't learn that from the network traffic, we have two options:

1. Map all the possible DNS vulnerabilities to any host offering DNS services.
2. Don't map any DNS vulnerabilities to DNS servers.

The default behavior for all applications in this list is #1 above. Mapping all of those extra vulnerabilities will probably just cause more false-positive,

high-impact intrusion events to be generated. Rather than do that, Cisco has opted not to map these vulnerabilities.

To be clear, this has nothing to do with the number of intrusion events generated. If the particular DNS rules are enabled, you will still get intrusion events for a BINDv7 attack on a Microsoft DNS server. However, instead of that being a red (high) impact event, it will have a lower impact assigned. If your brain hurts now, just take my word for it and leave these all disabled. You'll sleep better and you won't be chasing down as many high-impact false-positive events.

Web Analytics

This setting permits your FMC to send non-personally identifiable information to Cisco. The types of data include things like web browsers used, web pages viewed in the FMC along with the length of visit, product versions, device names, and IP addresses. This information is used by Cisco to help improve the products. If you want to opt out of sending this info, then uncheck the box!

Summary

In this chapter, I began working with the Firepower Management Center, or FMC, and I started by covering the system configuration of the FMC, which can be used to harden the FMC, for example.

Since we will be using the FMC to perform virtually all the management tasks, it makes sense to spend some time up front on its configuration before bringing the 7000/8000 appliances and Firepower Threat Defense (FTD) devices into the FMC.

Also, for this chapter, I used both a hardware FMC and a virtual FMC so you could see the small but important differences.

Chapter 3: Firepower Management Center (FMC) Actions

The following CCIE/CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.3 Configure these features using Cisco Firepower Management Center

2.3.d Actions

This chapter will continue to build upon the FMC configuration in Chapters 1 and 2, and our focus will be on how to configure the potential actions responses that the FMC can be configured to use.

Even though the objective for this chapter is all about actions, as a huge bonus along the way you'll be introduced to lots of FMC features that will also really help you to administrate your Firepower network!

Just for a heads-up, we'll cover remediations later in chapter 24 on correlation policy.

To find exam study material like videos, downloadable supplemental material, and practice questions, head over to www.lammle.com/firepower.

Firepower Management Center (FMC) Actions

Both the Responses and Remediations menu items are under the **Policies>Actions** menu, as shown here:

Policies | Devices | Objects | AMP | Intelligence

Network Discovery

Application Detectors

Correlation

Actions ▾

Responses

Alerts

Scanners

Groups

Remediations

Modules

Instances

board

ity on the appliance



s x

Intrusion Events x

Status x

Geolocation

er Time

- x

Top Web Appl

We could configure responses and remediations here, but as I said, we're going to focus on the very effective and popular alerts in this chapter.

Firepower Management Center Alert Responses

It's vital to have external event notification in place via SNMP, syslog, or email for critical-system monitoring. The FMC can use configurable alert responses to interact with your external servers, so alerts happen to be one of the first policies that I configure on my own FMCs.

Events detected by Firepower are called responses because you can use them to send alerts in response to an email, SNMP, or syslog server. This sounds great but you've got to approach the feature thoughtfully. It's just way too easy to overdo things regarding the amount of alerts you configure. Be sure you don't configure multiple alert responses to send different types of alerts to different monitoring servers and/or people. I don't know about you, but I sure get enough emails per day already and definitely don't need a few hundred more!

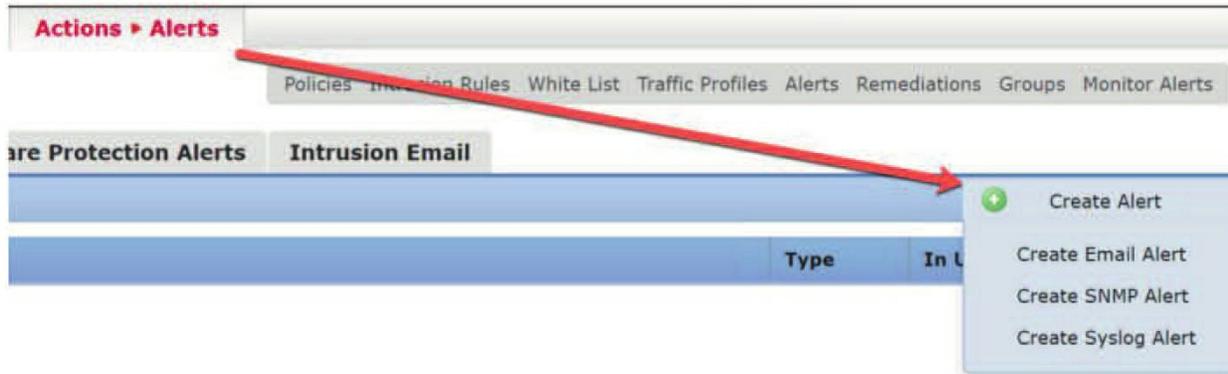
If you have intrusion rules set to send SNMP and syslog alerts, they'll be sent by the actual devices (appliance or FTD) to the servers configured from something called the Platform policy. This really can be the best way to configure your alerts, and I'll cover the platform policy in detail in Chapter 19.

Be advised that once we configure our alerts, I'll use them throughout the rest of the chapters.

Alert Responses

Action>Alerts is used more often than any other Action item found in this chapter. The new figure shows the three alerts that I'll configure:

As shown in the figure below, the three alerts are as follows:



- Email
- SNMP
- Syslog

Email Alerts

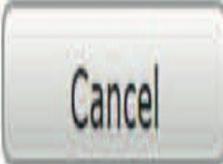
Before you can configure email alerts, you need to make sure that your DNS is working on your FMC and you can reverse-resolve the FMC's own IP address. You just can't configure the email alert if you haven't set up an SMTP relay in the System Configuration settings as I demonstrated in Chapter 2.

Last, there are no license requirements for any of the alerts covered in this chapter, so using a base license will work for these just fine.

Edit Email Alert Configuration



Name	Lammle_Email_Alert
To	fmcgroup@lammle.com
From	2500FMC@lammle.com
Relay Host	smtp.gmail.com 

But you do need to have FMC admin access. The figure shows the configuration I'll use for the email alert on my FMC. Click Create Email

Alert.

Notice that the relay host is already configured, or you wouldn't be able to save your configuration. Add the group email address you want to send the emails to, and then just make up the from address and you're set.

You can now use this alert in your policies. I'll use all the alerts as soon as the next chapter in the section "Health Monitor Alerts."

SNMP Alerts

Before you can configure the SNMP alert, you need to configure SNMP in System Configuration and also configure your NMS station in the Access List section as we did back in Chapter 2. Once you've done that, you can come back here and configure the alert.

You can create SNMP alert responses using SNMPv1, SNMPv2, or SNMPv3, but there are some caveats. SNMPv2 only supports RO communities and SNMPv3 only supports read-only users and only communities and SNMPv3 only supports read-only users and only bit monitoring, you must use SNMPv2 or SNMPv3 if you want to monitor 64-bit values with SNMP.

The next figure shows the configuration I'll use for the SNMP alert on my FMC. Click Create SNMP Alert.

Edit SNMP Alert Configuration



Name	<input type="text" value="Lammie_SNMP_Alert"/>
Trap Server	<input type="text" value="10.11.11.250"/>
Version	<input type="text" value="v2"/>
Community String	<input type="text" value="firepower"/>

Now hit Save to make this alert available to use in your policies and rules.

If your network management system requires the FMC's management

information base (MIB) file, obtain it by connecting to your FMC console. The file is at /etc/sf/DCEALERT.MIB.

Syslog

Syslog can be a real beast to deal with because there are just so many options that you can configure. It's worth it though—I definitely find syslog more useful than SNMP for this very reason.

When configuring a syslog alert response, you can specify the severity and facility associated with the syslog messages to ensure they're processed properly by the syslog server. The facility indicates the subsystem that creates the message and the severity defines the severity of the message. Facilities and severities aren't displayed in the actual message that appears in the syslog but are instead used to tell the system that receives the syslog message how to categorize the messages.

Although you can choose any type of facility when creating a syslog alert response, you should choose one that makes sense based on your syslog server configuration. Understand that not all syslog servers support all facilities, so when in doubt, just add SYSLOG here as shown in the next figure. For UNIX syslog servers, the *syslog.conf* file should indicate which facilities are saved to which log files on the server.

This figure shows the configuration I'll use for the syslog alert on my FMC. Click Create Syslog Alert and you can see that there are more configuration options here than in the SNMP configuration. The facility is typically SYSLOG, but there are many other options. Again, when in doubt, just choose SYSLOG here.

In the next figure, you can see that there are eight severity levels, and unfortunately, they're in alphabetical, not severity level, order. Worse, as you move through the Firepower system, you'll see severities in alphabetical order like this and yet sometimes listed in severity level.

This can be annoying.

If you choose debugging, which is “everything on,” it's fine as long as you're sending the data to a syslog server and not to your console—choose wisely

here.

This table shows the severity levels and the description of each.

Level Description

ALERT A condition that should be corrected immediately CRIT A critical condition

DEBUG Messages that contain debugging information EMERG A panic condition broadcast to all users

ERR An error condition

INFO Informational messages

NOTICE Conditions that are not error conditions but require attention

WARNING Warning messages

Once you've chosen the port, facility, and severity, you can finally add a tag that sends vendor information about the device, as shown here:

Edit Syslog Alert Configuration

? X

Name	<input type="text" value="Lammle_Syslog_Alert"/>
Host	<input type="text" value="10.11.11.250"/>
Port	<input type="text" value="514"/>
Facility	<input type="text" value="SYSLOG"/>
Severity	<input type="text" value="INFO"/>
Tag	<input type="text" value="2500FMC"/>

Save

Cancel

Finally, your alerts should look something like what I've got here:

Name	Type	In Use	Enabled	
Lammie_Email_Alert	Email	Not Used		
Lammie_SNMP_Alert	SNMP	Not Used		
Lammie_Syslog_Alert	Syslog	Not Used		

FMC Detailed Alerts

Also found on the **Policies>Action** page are miscellaneous alerts you can configure, as shown here:

Alerts	Impact Flag Alerts	Discovery Event Alerts	Advanced Malware Protection Alerts	Intrusion Email
---------------	--------------------	------------------------	------------------------------------	-----------------

It's key to understand that most administrators typically won't use these unless they're a full-time NGFW administrator because you receive a ton of data with these with alerts! But if your job is to work on Firepower and nothing else, these tabs will be useful to you.

That said, you can use something like intrusion emails for a bit if, say, you're looking to make sure you solved an issue and want to be instantly alerted for a specific problem. But just turning all these on will completely inundate you

with data!

Let me show you what I mean by going through the following tabs:

- Impact Flag Alerts
- Discovery Events Alerts
- Advanced Malware Protection (AMP) Alerts
- Intrusion Email

Impact Flag Alerts

- Supported devices: Any
- License needed: Threat, Protection
- Access: Admin

That you can configure the system to alert you whenever an intrusion event with a specific impact flag occurs is a great advantage.

Impact flags help you evaluate the impact an intrusion has on your network by correlating intrusion data, network discovery data, and vulnerability information. Their values are from 0 to 5, with the lower number being considered a more urgent issue.

The next figure shows the Impact Flag Alerts page. Notice I didn't enable the email!

Impact Flag Alerts

Discovery Event Alerts

Advanced Malware Protection Alerts

Alerts

Select alert responses to use for Impact Flag alerts. You must click Save after you make your selections.

Syslog

Lammie_Syslog_Alert ▼

Email

None ▼

SNMP

Lammie_SNMP_Alert ▼

To help you evaluate the impact an event has on your network, the Firepower Management Center displays an impact level in the table view of intrusion events. For each event, the system adds an impact level icon whose color

indicates the correlation between intrusion data, network discovery data, and vulnerability information.

Impact Vulnerability Color Description

0 Unknown Gray Neither the source nor the destination host is on a network that's monitored by network discovery.

1 Vulnerable Red Either the source or the destination host is in the network map, and a vulnerability is mapped to the host or the source. Can also mean the destination host is potentially compromised by a virus, Trojan, or other piece of malicious software.

2 Potentially

Vulnerable

3 Currently Not Vulnerable

4 Unknown Target

Orange Either the source or the destination host is in the network map and one of the following is true: for port-oriented traffic, the port is running a server application protocol; for non-port-oriented traffic, the host uses the protocol.

Yellow Either the source or the destination host is in the network map and one of the following is true: for port-oriented traffic (example, TCP or UDP), the port isn't open; for non-port-oriented traffic (example, ICMP), the host doesn't use the protocol.

Blue Either the source or destination host is on a monitored network, but there is no entry for the host in the network map.

Discovery Events Alerts

- Supported devices: Any
- License needed: Any
- Access: Admin

Now here's another option that would provide a lot of data you may not need. Even my most security-minded clients don't use this, but if you had a data

center in a bank that no hosts should ever be added to without your knowledge, this would be the alert for you.

You can configure the system to alert you whenever a specific type of discovery event occurs. Here I show you the Discovery Event Alerts page. Remember, I didn't enable the emails.

Impact Flag Alerts

Discovery Event Alerts

Advanced Malware Protection Alerts

Alerts

Select alert responses to use for discovery event alerts. You must click Save after you make your selections.

Syslog

Lammle_Syslog_Alert ▼

Email

None ▼

SNMP

Lammle_SNMP_Alert ▼

Scrolling further down this same page, in the next figure, notice we still have even more detailed configurations available to us.

Events Configuration

Select the event types for which you want to generate alerts.

Event	Syslog Notification	Email Notification	SNMP Notification
	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Add Client	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Add Host	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Add Port	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Add Protocol	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Add Scan Result	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Advanced Malware Protection (AMP) Alerts

You can configure the system to alert you whenever any malware event is generated by AMP for Networks, including a retrospective event, but you can't alert on malware events generated by AMP for Endpoints.

- Supported devices: Any
- License needed: Malware
- Access: Admin

This next figure illustrates the Advanced Malware Protection Alerts page. We have the options to receive just retrospective or all network malware events.

Alerts

Select alert responses to use for advanced malware detection event alerts. You must click Save after you make your selections.

Syslog: ▼

Email: ▼

SNMP: ▼

Event Configuration

Select the event types for which you want to generate alerts.

Event	Syslog	Email	SNMP
Retrospective Events	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
All network-based malware events	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Save

Intrusion Email

You can click on Intrusion Email and configure email alerting on an IPS event if needed. Since no one wants to get more emails, clicking on Coalesce Alerts and Summary Output will let you receive a summary of events in one email instead of an email every time there is an event. So yes, definitely enable those!

Alerts

Impact Flag Alerts

Discovery Event Alerts

Advanced Malware Protection Alerts

Intrusion Email

Email Alerting

State on off

From Address

To Address

Max Alerts

Min Frequency

Note: Frequency is set by number of seconds

Coalesce Alerts on off

Summary Output on off

Time Zone

Rules [Email Alerting per Rule Configuration](#)

Save

Other Connections You Can Log

You really need to pay attention to the FMC logging options and how much an FMC can handle (covered in Chapter 1), then set logging on your rules as needed with respect to the FMC's capacity.

As we travel through the policies in this book, the following factors are vital when logging configuration on a per-rule basis:

- Prefilter Policy: Rules and Default Action
- SSL Policy: Rules and Default Action
- Access Control Policy: Security Intelligence Decisions
- Access Control Policy: Rules and Default Action

Configuring Response Groups

The groups are used when you configure a correlation policy, which is actually a pretty advanced policy. Still, configuring groups here can really make things easier for you. Again, no worries—we'll explore this further in the chapter on correlation policy.

Supported Devices: Any

License needed: Any

Access: Admin/Discovery Admin

You can create a correlation response group of alerts and

remediations, then activate and assign the group to a correlation rule within your correlation policy. The Firepower system launches all the grouped responses when network traffic matches the correlation rule.

When done in an active correlation policy, changes to an active group or any of its grouped responses take effect immediately.

Summary

This chapter covered the alert functions in the Firepower system. You learned that it's vital to configure the FMC system configuration before you can configure the alerts discussed in this chapter. We configured the most

important alerts and created the email, SNMP and Syslog alerts.

Chapter 4: Licensing & Health Policy

The following CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.1 Configure system settings in Cisco Firepower Management Center

In this chapter, we will discuss licensing of the FMC as well as the health policy. To install and maintain your licenses for your devices, you must install a Smart License to use FTD devices (Classic for other devices), and the FMC then must be able to talk to the Smart License server at all times. Understand that licensing for your devices happens only at the FMC. For FTD you must install the license on the FMC before you can bring the device into the FMC to be managed. If you remove a device from an FMC, the license stays on the FMC, not the device. So, if you want to move your device to another FMC, you must install a license on that new FMC first. But what if your FMC is not connected to the Internet? An airgapped FMC for example? In this chapter I'll provide two solutions to help you meet this type of business requirement.

The health of the Firepower system should be monitored to ensure proper operation of the devices involved. Health is monitored by using a series of checks configured through the health policy. The health policy settings have essentially not changed since the old Sourcefire days. It is the one policy/setting that needs to be applied in an “old school” way.

In situations where you do not look at the FMC every day, Health Monitor alerts can be created to automatically alert you to a change in your systems health conditions.

Licensing

With FTD, Cisco introduced the Smart Licensing model. Smart Licensing centrally manages a pool of licenses. Unlike the older method (Classic Licensing) of using a product authorization key (PAK), Smart Licenses are not tied to a specific serial number or license key. Each customer account

receives a token code, which is used to register your FMC with the Cisco central licensing manager.

After this, any licenses you purchase are automatically available on your FMC. Firepower Threat Defense devices use Smart Licensing, while classic devices (Firepower 7000/8000, FirePOWER on ASA, NGIPSv, Meraki with FirePOWER) use Classic Licensing.

One very nice thing about the FMC and Smart Licensing is that Cisco provides a 90-day grace/evaluation period when you first install your FMC. During this period, you can try out any of the features except non-exportable options such as remote access.

This figure shows our FMC currently in this evaluation period.

Welcome to Smart Licenses

Before you use Smart Licenses, obtain a registration token from [Cisco Smart Software Manager](#), then click Register

Register

Smart License Status

[Cisco Smart Software Manager](#)

Usage Authorization:	N/A
Product Registration:	Evaluation Period (Expires in 79 days)
Assigned Virtual Account:	Evaluation Mode
Export-Controlled Features:	Enabled

Smart Licenses

Filter Devices...

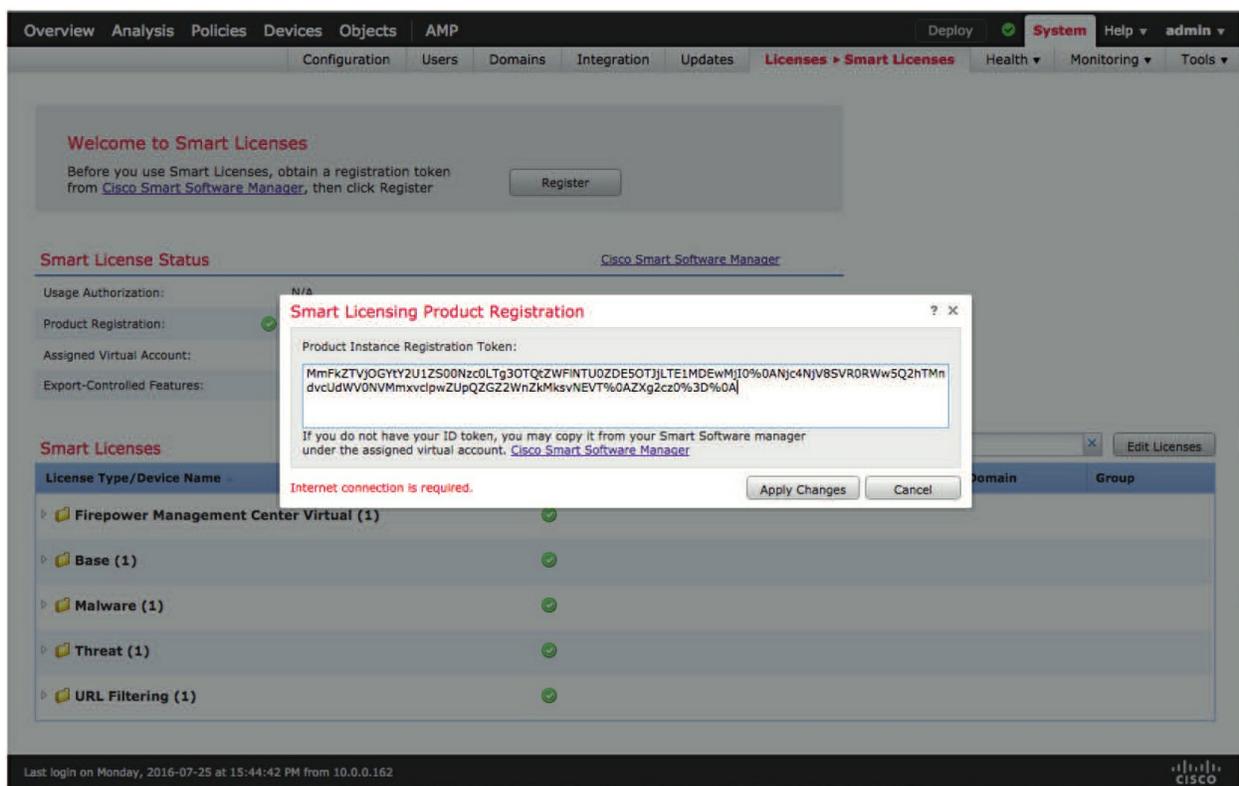
License Type/Device Name	License Status	Device Type	Domain	Group
Firepower Management Center Virtual (1)				
Base (1)				
Malware (1)				
Threat (1)				
URL Filtering (1)				

The number next to each folder shows how many licenses are in use.

The process of registering your FMC is pretty straightforward. It will require a visit to the online Cisco Smart Software Manager. There you will log in and generate a new token.

To find the manager URL, you can navigate to **System>Licenses> Smart Licenses** on your FMC. There you will find a link to the Smart Software Manager.

Then simply click the Register button and paste the token into the field.



After this, any licenses you have will automatically be downloaded and synchronized with your FMC. The FMC must have an active Internet connection to maintain your Smart Licenses.

If this connection is interrupted or you have been connected for too long (over 90 days), your licenses could be deactivated. That said, there are two different ways to install and maintain licenses with no Internet connection:

- Smart Satellite server
- Specific License Reservation
- Let's take a quick look at each of these

Under `System>Integration` there is a tab called Smart Satellite server, as shown next. This allows you to have a server with an interface on the Internet and then a connection to a private network that has no Internet configuration.

I could easily write a whole chapter on this subject, but I'll post documentation on the configuration steps on my website since this is way above the objectives of this book. Understand that the advantage of this method is that it will dynamically update your Smart License on your FMC, but the disadvantage is the difficult and time-consuming installation.

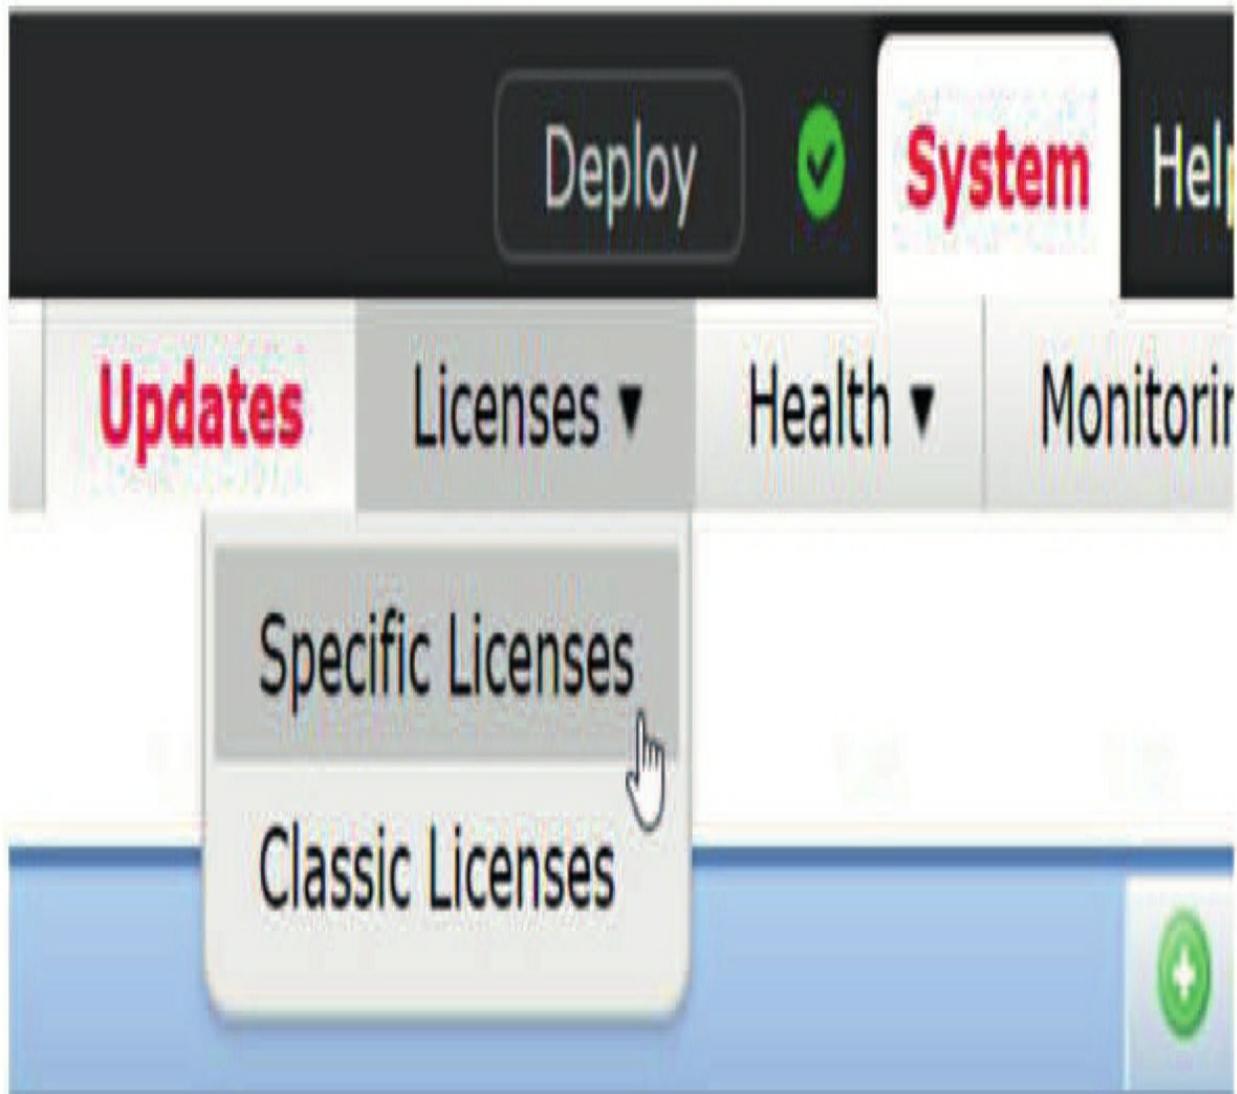
Isn't there a better way then? Yes, but this very easy way doesn't dynamically update like the Satellite server does, and this is found in a blog post on my site as well. There is a new Specific License Reservation available for approved customers. This allows using the Firepower Management Center (FMC) on an air-gapped network.

Previously, the only way to use the smart license feature was to have either Internet access or a satellite license server. Now a special license can be installed that is valid without the requirement to access the Cisco site or the Smart Software Satellite Server.

This is as simple as it comes:

1. Log in to the CLI of the FMC.
2. Enter the `expert` command to access the shell.
3. Execute the following command to access the Specific License Reservation options:
4. Chose option 2 to enable SLR.
5. Chose option 0 to exit.

Now log out and back into your FMC and notice the change in the license's available options, shown here.



Now you just need to get a license token from Cisco at this point, which is put into your CCO in order to finish installing the license.

Health Policy

This part of the chapter will cover the following health-related items from the **System>Health** menu, as shown in this figure.

Health ► Monitor

Monitor 

Policy

Events

Blacklist

Monitor Alerts

ary

I'll discuss all items in the menu in the following sections:

- Health Monitor
- Health policy
- Health events
- Blacklists
- Monitor alerts

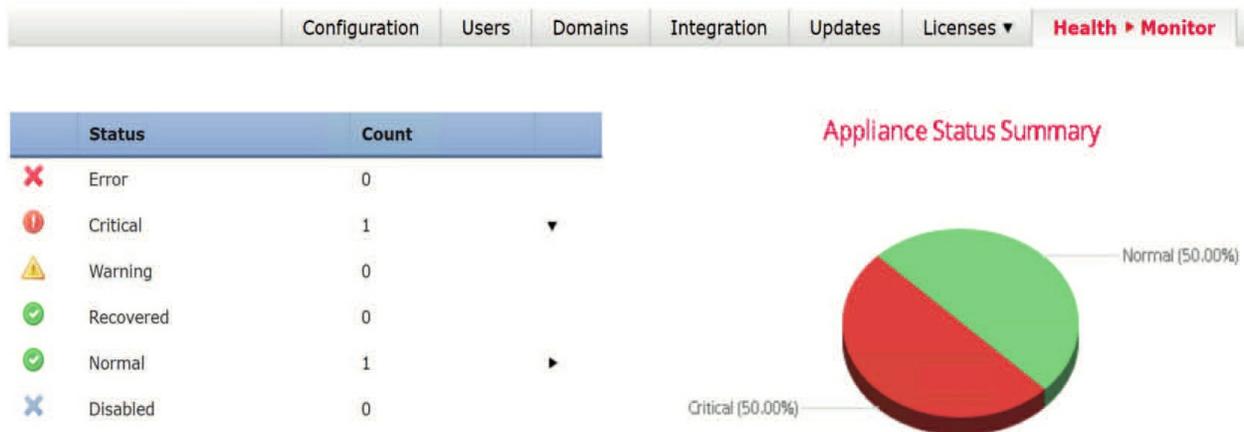
Health Monitor

The first item under **System>Health** is the *Health Monitor*, which reveals details about the health of your system. There are several different status messages that can be displayed, with Normal being optimal.

Critical indicates just what you would think—that there are issues that probably require immediate attention, and Warning indicates something worth keeping an eye on as well.

A status of Recovered indicates that a health check previously reported a problem condition (Critical or Warning) but it's now okay. Disabled indicates that there's no health policy loaded, and Error tells you that there's an issue with the health system itself.

Notice in this figure that there's a small triangle to the right of the number for the Normal status. Clicking it will expose details of the appliances.



The Count column next to a particular status shows the number of appliances in that specific state. However, the question I always get at this point is, “Why are my health alerts always critical!?”

One of the first items to disable in your new health policy would be interface status if you are using FTD or if you are using ASA with FirePOWER and the ASA is in an HA pair.

Once you bring these types of devices into the FMC, you'll get critical alerts

as shown next.



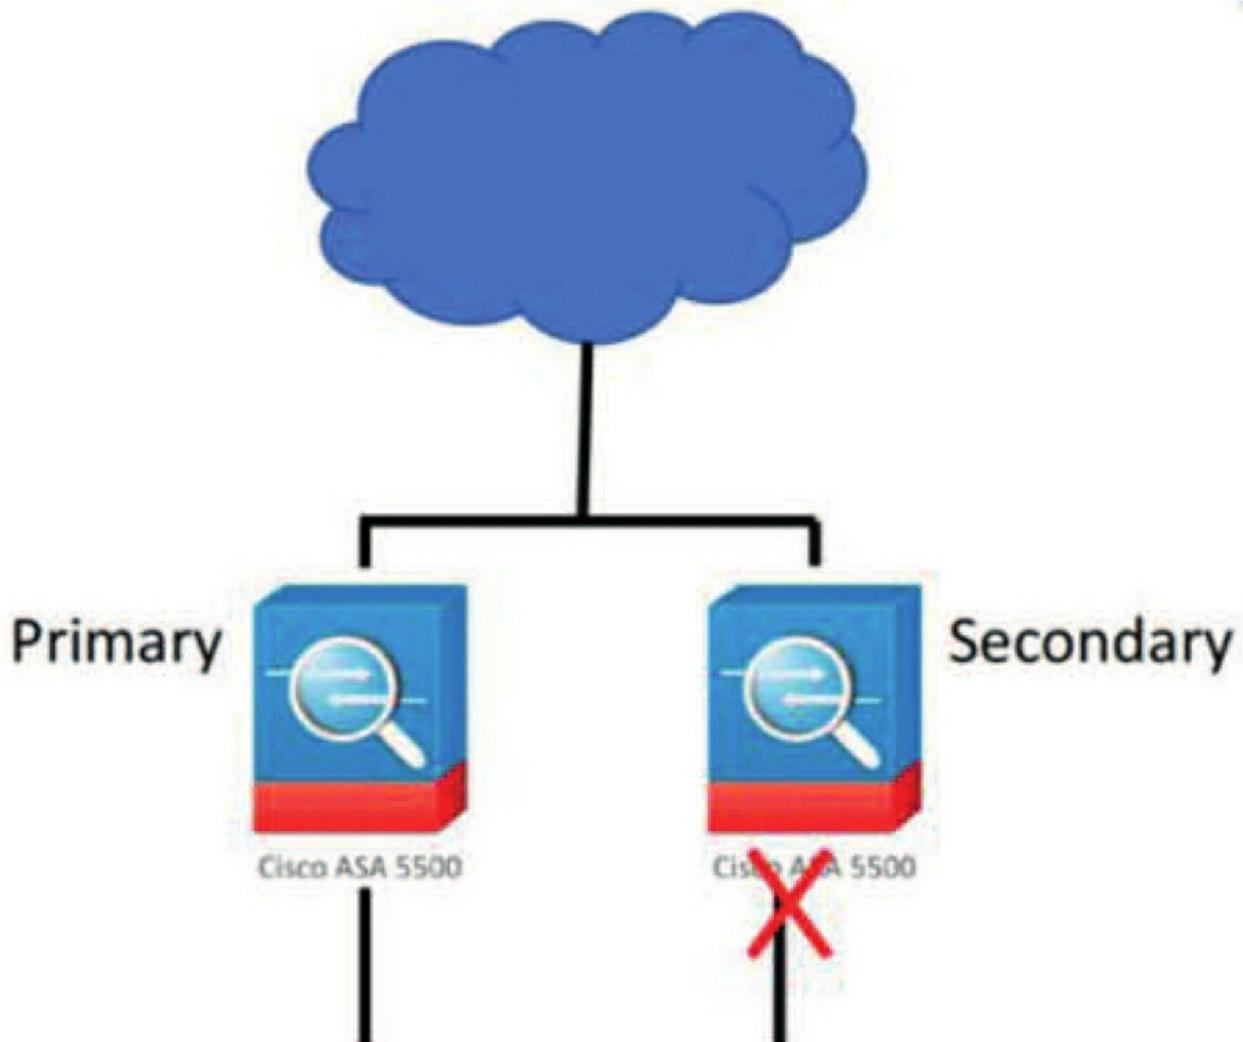
Let's take a look at why.

ASA with FirePOWER

If you have ASAs, you'll have a primary and a secondary group for HA, as shown here:

The secondary in the group does not communicate on the inside interface, so the device will start issuing a critical alert: Interface

'DataplaneInterface0' is not receiving any packets". This is annoying because what it is saying is, "I'm working correctly, and I just wanted you to know that,



so here's a Critical alert message every 5 minutes!".

Firepower Threat Defense (FTD)

The FTD device is providing a slightly different Critical error, but the solution is the same. The interface is receiving packets, unlike the secondary ASA, but it is not registering because you are using a subinterface. This is an

issue that Cisco needs to resolve. In the meantime, just apply the same health policy to your FTD that was applied to the secondary ASA and you'll stop receiving these messages.

Health Policy

Data for the Health Monitor comes from the *health policy*, which is a collection of checks that are executed every 5 minutes by default. You can modify the health policies to meet the needs of each environment. Health checks can also be disabled.

To fix the issues mentioned above, edit or create a new Health policy by navigating to **System>Health>Policy** and click the Create Policy button.

Create Policy

Copy Policy

Initial_Health_Policy 2019-05-31 20:05:28 ▼

Initial_Health_Policy 2019-05-31 20:05:28

New Policy Name

Default Health Policy

New Policy Description

Save

Cancel

When creating a new policy, you get two options to copy an existing health policy. Choose either one; it doesn't matter.

There are many different checks in a list to the left of the policy. Each setting can be enabled or disabled, and some allow you to specify thresholds for warnings or critical alerts.

Most of the default settings work just fine, but if you make any changes, be sure to use the **Save Policy** and **Exit** button on the bottom and then reapply the policy.

There is one other monitor I turn off and this is the Smart License Monitor shown in the figure, if I am using an Eval license. This will stop the critical alerts for this happening every 5 minutes.

▶ Smart License Monitor

This license alert is Cisco just telling you, “Hey, you owe us money!” but they do it every 5 minutes, which drives you batty. However, understand that it’s important you don’t turn off the Smart License Monitor when you apply a real token to your FMC, so you’ll be alerted if you have a license issue.

Notice in the next figure that since I am using FTD devices in this chapter, I’ll disable the interface status and then go to the bottom and click on Save Policy and Exit.

The image shows a configuration interface for the Smart License Monitor. On the left, there is a list of monitors:

- Policy Run Time Interval
- AMP For Endpoints Status
- AMP for Firepower Status
- Appliance Heartbeat
- Automatic Application Bypass Status
- Backlog Status
- CPU Usage
- Card Reset
- Cluster/Failover Status
- Disk Status
- Disk Usage
- FMC HA Status
- Hardware Alarms
- Health Monitor Process
- Host Limit
- Inline Link Mismatch Alarms

Below this list, the **▶ Interface Status** monitor is selected. On the right, the configuration for this monitor is shown:

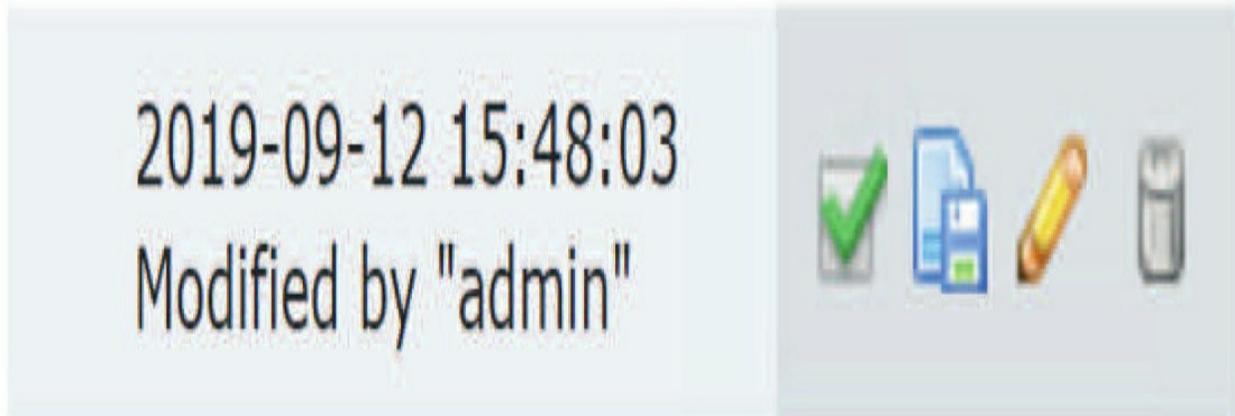
Description	Monitors if the interfaces are receiving traffic
Enabled	<input type="radio"/> On <input checked="" type="radio"/> Off

It is worth noting here that this is the one policy that is not deployed with the Deploy button. As mentioned, this policy is a dinosaur and still follows the old ways!

To deploy the health policy, you click on the



icon in the health policy list.



This will the present a list of all the appliances, including the FMC itself, as shown in here.

▼ Ungrouped (2 total)

FTD10
10.11.10.10 - Cisco Firepower Threat Defense for VMWare v6.4.0.1

fmc1.sfgtc.local
10.11.10.15 - Cisco Firepower Management Center for VMWare v6.4.0



Select the appropriate appliances and then click the Apply button at the bottom of the screen, and in 5 minutes your alerts should go away.

Health Events

To view current and past health events, select **System>Health> Events**. The health events will be displayed from all your devices, as shown here.

No Search Constraints (Edit Search)

Module Name	Test Name	Time	Description	Value	Units	Status
Memory Usage	Memory Usage - Memory Test	2019-09-12 16:29:01	Used 8057.53MB of 21966.09MB (Physical + Swap)	37		✓
VPN Status	VPN Status	2019-09-12 16:29:01	Process is running correctly	0		✓
Smart License Monitor	Smart License Monitor	2019-09-12 16:29:01	Smart Licensing in evaluation mode.	0	Licenses	✓
Appliance Heartbeat	Appliance Heartbeat	2019-09-12 16:29:01	All appliances are sending heartbeats correctly	0		✓
Security Intelligence	Security Intelligence	2019-09-12 16:29:01	Process is running correctly	0		✓
Threat Data Updates on Devices	Threat Data Updates on Devices	2019-09-12 16:29:01	Process is running correctly	0	hours	✓
AMP for Firepower Status	AMP for Firepower Status	2019-09-12 16:29:01	Successfully connected to cloud	0		✓
RRD Server Process	RRD Server Process	2019-09-12 16:29:01	The server is functioning normally.	0	n/a	✓
Interface Status	Interface Status	2019-09-12 16:29:01	All interfaces are working correctly	0	Packets Received	✓
Disk Usage	Disk Usage - Disk Test	2019-09-12 16:29:01	/ using 37%	37	%	✓

If you're looking for a specific health event, use the Search feature and specify your search criteria. You can also adjust the time window to a specific range.

Blacklist

The health blacklist (**System>Health>Blacklist**) allows you to set up exclusions to health checks. You can blacklist an entire device or just individual checks.

This can be helpful when there's a known issue, but you don't need alerts nagging you about it every 5 minutes, and it looks as shown in this figure



Then the device shows, as you can see here:

A good example of when you can use this, is when you've had a power supply fail but you've contacted Cisco for a replacement.

Click the pencil next to the device when in this screen, and you would

just blacklist the power supply health check for the appliance until you get the replacement, as shown here.

Health Monitor Alerts

This section will use the alerts we

Modules

- Intrusion and File Event Rate**
- Link State Propagation**
- Local Malware Analysis**
- Memory Usage**
- Platform Faults**
- Power Supply**
- Process Status**
- Realm**
- Reconfiguring Detection**
- RRD Server Process**
- Security Intelligence**
- Smart License Monitor**
- Threat Data Updates on Devices**
- Time Series Data Monitor**
- Time Synchronization Status**
- URL Filtering Monitor**
- User Agent Status Monitor**
- VPN Status**

Save

Back

created in Chapter 3.

In some environments, the FMC console may not be manned 24/7. In these situations, it's a good idea to have external alerts set up to notify people immediately of any health conditions that crop up. This is where Health Monitor alerts come in—with them, you can send emails, SNMP traps, or syslog messages.

As discussed, and shown in Chapter 3, to configure Health Monitor alerts, you must first set up the responses known as alerts. Navigate to **System >Health >Monitor Alerts** to display the configuration page. No need to go over those again here.

Once the alerts are created, go to the Health Monitor Alerts screen. To create the Health Monitor alert, select a severity based on when you want the alert to occur, the module(s) to monitor, and finally, the alert you want generated. There are six steps to this:

1. Specify the health alert name at the top.
2. Choose the severity level.
3. Highlight and choose all modules.
4. Choose the alert or alerts you want to use.
5. Save the alert(s).
6. Choose and load the alert.

Keep in mind that you can opt to include multiple severities, modules, and alerts.

Summary

In this chapter, I discussed licensing of the FMC as well as the Health policy. To install and maintain your licenses for your devices, you must install a Smart License to use FTD devices (Classic for other devices), and the FMC then must be able to talk to the Smart License server at all times.

Licensing for your devices happens only at the FMC. But what if your FMC is not connected to the Internet? I provided two solutions to help you meet this type of business requirement.

You should monitor the health of the Firepower system to ensure proper operation of the devices involved. You can do so by using a series of checks configured through the health policy.

You may not look at the FMC every day, so you can create Health Monitor alerts to automatically alert you to a change in your system's health conditions.

Chapter 5: Chassis Manager

This is an *advanced chapter* covering the Cisco 4100 and 9300 devices, and you're probably going to feel like you were suddenly thrown into the deep end! But I promise to keep your head above water and get you through this vital configuration. It's just really important to have a handle on it before we get to the next chapter where we'll add the devices to the FMCs we configured in Chapter 1, "Firepower Management Center (FMC)."

It's important to understand that there are no CCIE/CCNP Security SNCF exam objectives covered in this chapter. This chapter instead includes important foundational information because it is still is very likely you will see exam questions from this chapter. So even though you won't be checking off any CCIE/CCNP Security SNCF exam objectives here, we'll be covering really important, foundational configuration that's absolutely key to understanding the rest of this book.

Firepower eXtensible Operating System (FXOS) runs the Firepower Chassis Manager for the 4100/9300's, which provides GUI-based management of the hardware for the device. A component called the supervisor is responsible for the physical interfaces and for managing the ASA or FTD images that'll run on the device.

We'll get started with a hardware overview of the architecture used with FXOS, and then I'll jump into the boxes and reset the 4100s used in this book back to factory default. I'll then configure the hardware using Chassis Manager, and after that, we'll tour the FXOS architecture and different contexts available from the CLI.

Building further, I'll show you how to add images to the Chassis Manager and perform upgrades of FXOS, including an image for something called Radware, which is a really helpful tool when you're dealing with denial of service (DoS) attacks.

Wait, there's more! We're going to log in to Firepower Chassis Manager (FCM), go through the platform settings, and configure the FXOS for ASA, FTD, and Radware. I'll end the chapter by clustering two 4100 chassis together for the grand finale!

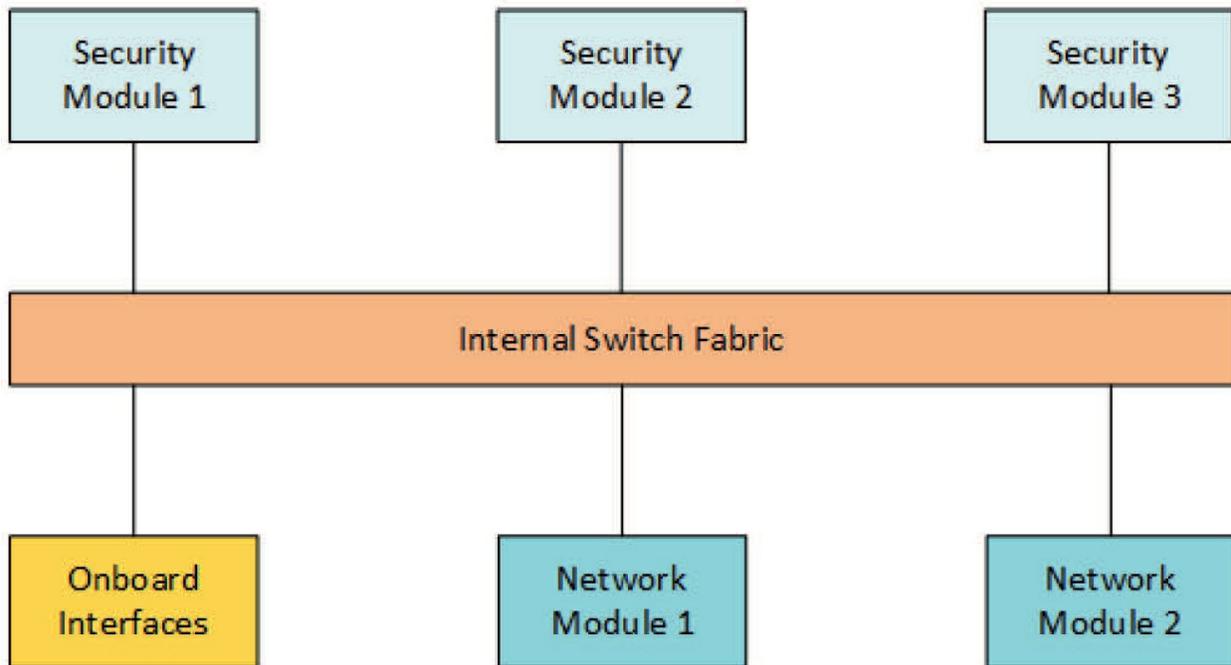
So, get ready because this is going to be quite a ride. I promise to make as painless and fun as I can!

To find exam study material such as hands-on labs access, videos, downloadable supplemental material, and practice questions, head over to www.lammle.com/firepower

Hardware Overview

You might be wondering just what's so special about the 4100/9300 series that it requires a whole chapter just to talk about it! Well, aside from offering very high throughput as well as plenty of hardware bells and whistles, it takes full advantage of Cisco Unified Computing System (UCS) architecture to offer Cisco security solutions in a modular fashion.

The 9300 supports 8x10 gbps onboard interfaces, plus additional network modules if you want to add some 40/100 gbs interfaces to the appliance.



The appliance also supports three security modules, which lets you install up to three images without creating any instances since these are basically hard drives that you can deploy a solution to.

Since the 9300 is a Firepower appliance, you probably guessed that the security module supports FTD. But that's not all; you can also have an ASA deployment. Why would you want that? Mostly, in case you need to support a feature that hasn't been fully ported over to Firepower yet, like a Virtual Tunnel Interface (VTI) or clientless VPNs.

Cisco has partnered with RadWare to provide some pretty fantastic DDoS protection by letting you install its DefensePro solution into a security module. We'll take a look at this solution later in this chapter.

The 4100 is pretty much the same idea as the 9300, only on a smaller scale. For example, it only supports one security module instead of three. Even so, we can add multiple instances of ASA and FTD on both the 4100 and 9300.

Now that you've got a picture of the hardware architecture, let's move on to getting logged in!

Resetting our 4100s

First, and to get ready for exploring the FXOS and Chassis Manager ahead, I'm going to set my 4140 boxes back to factory default. You've got to be able to log in to your devices to do this, so if you don't know or you forgot the password, I'll show you how to reset it. After that, I'll erase the devices, do an initial configuration, and then get them ready for configuration into the FMC that's coming up in the next chapter.

Resetting the Password on 4100/9300

Okay—here's how you reset the password:

1. Reboot the device.
2. Use `BREAK`, `ESC`, or `CTRL+L` to interrupt the boot.
3. Find the boot flash command and make a note of the kickstart image and system image.
4. Load the kickstart image, which will be something like this:
Load the kickstart image, which will be something like this:
`kickstart.5.0.3.N2.3.14.69.SP`
5. This will take you to the `switch(boot)#` prompt.
6. Now you can change the password:

```
switch(boot)# config t
switch(boot)(config)# admin-password erase
Your password and configuration will be erased
Do you want to continue? (y/n) [n] y
```

7. Exit to the `Switch(boot)#` prompt and load the system image saved earlier to complete the procedure:

```
switch(boot)(config) # exit
switch(boot)# load bootflash:/installables/switch/fxos-k9 system.5.0.3.N2.3.14.69.SPA
```

8. Configure the device.

Setting the 4100/9300 Devices to Factory Default

Now that you know the password, or have reset it, we'll go ahead and set the FTD box (4100/9330) back to factory default:

1. Log in to your FXOS.

2. Type `connect local-mgmt`
3. Type `erase configuration`

```
cisco4140-1# connect local-mgmt
cisco4140-1(local-mgmt)# erase configuration
All configurations will be erased and system will reboot. Are you sure? (yes/no): yes
Removing all the configuration. Please wait....
Configurations are cleaned up. Rebooting....
```

[output cut]
Type Ctrl-C at any time for more options
or to abort configuration and reboot system.
You have chosen to setup a new Security Appliance.

Continue? (yes/no):yes
Now we'll wait for this to reboot and then continue with the initial
configuration afterward.

Initial Configuration for 4100/9300

Things are going to get more complicated now than they were on the
previous boxes—at least for the first time you perform this setup. After you
agree to the continue question in the last line of the above output, you'll be
presented with these options:

Type 'reboot' to abort configuration and reboot system, Type 'show_mgmt_port' to display
management port info Type 'show_dhcp_lease' to display DHCP lease
Type 'show_serial_num' to display Chassis Serial Number or press Ctrl+D to continue.

I'll be going with Ctrl+D here to configure the device:

Type Ctrl-C at any time for more options
or to abort configuration and reboot system.
You have chosen to setup a new Security Appliance.

Continue? (yes/no): yes
Enforce strong password? (yes/no) [y]: n
Enter the password for "admin":
Confirm the password for "admin":

Next we're prompted to configure a system name, which requires beginning
it with an alphanumeric character:

Enter the system name: 4140
The system name must start with a alphabetic character and end with a alphanumeric character and may
have alphanumeric characters including hyphen.

Enter a system name : A4140-1
Supervisor Mgmt IP address : 10.11.10.207
Supervisor Mgmt IPv4 netmask : 255.255.255.0
IPv4 address of the default gateway : 10.11.10.1

So I can more easily use SSH and HTTP for the devices, I'm going to keep the addresses broad for the sake of these labs. In the real world, I'd definitely recommend configuring your devices more securely!

The system cannot be accessed via SSH if SSH Mgmt Access is not configured.
SSH Mgmt Access host/network address (IPv4/IPv6): 10.0.0.0
SSH Mgmt Access IPv4 netmask: 255.0.0.0
Firepower Chassis Manager cannot be accessed if HTTPS Mgmt Access is not configured.
Do you want to configure HTTPS Mgmt Access? (yes/no) [y]: HTTPS Mgmt Access host/network address (IPv4/IPv6): 10.0.0.0
HTTPS Mgmt Access IPv4 netmask: 255.0.0.0
Configure the DNS Server IP address? (yes/no) [n]: y DNS IP address : 10.11.11.250
Configure the default domain name? (yes/no) [n]: yes

Default domain name : sfgtc.local
Following configurations will be applied:
Switch Fabric=A
System Name=A4140-1
Enforced Strong Password=no
Supervisor Mgmt IP Address=10.11.10.207
Supervisor Mgmt IP Netmask=255.255.255.0
Default Gateway=10.11.10.1
SSH Mgmt Access Configured=yes
SSH Mgmt Access IP Address=10.0.0.0 SSH Mgmt Access IPv4 Netmask=255.0.0.0
HTTPS Mgmt Access Configured=yes
HTTPS Mgmt Access IP Address=10.0.0.0 HTTPS Mgmt Access IPv4 Netmask=255.0.0.0
DNS Server=10.11.11.250
Domain Name=sfgtc.local

Apply and save the configuration (select 'no' if you want to re-enter)?
(yes/no):yes

Applying configuration. Please wait.. Configuration file - Ok

.....

With the basic configuration complete, I'll go ahead and set the 4140-2 back to factory default now. I'm also going to do the initial setup without showing all the output here. The IP address I'll use for 4140-2 is 10.11.10.209.

Here's the lab layout information we'll be using throughout this book. I highlighted the 4140s. Their vFMC will be 10.11.10.205 right there at the top:

Lab Layout Information

Device IP Address Device IP Address FMC 2500-1 172.16.10.20 vFMC 10.11.10.205

FMC 2500-1 172.16.20.21 4140-1 10.11.10.207

FTD 1010-1 172.16.10.10 4140-2 10.11.10.209

FTD 1010-2 172.16.10.11 Firepower Appliance 10.11.10.210

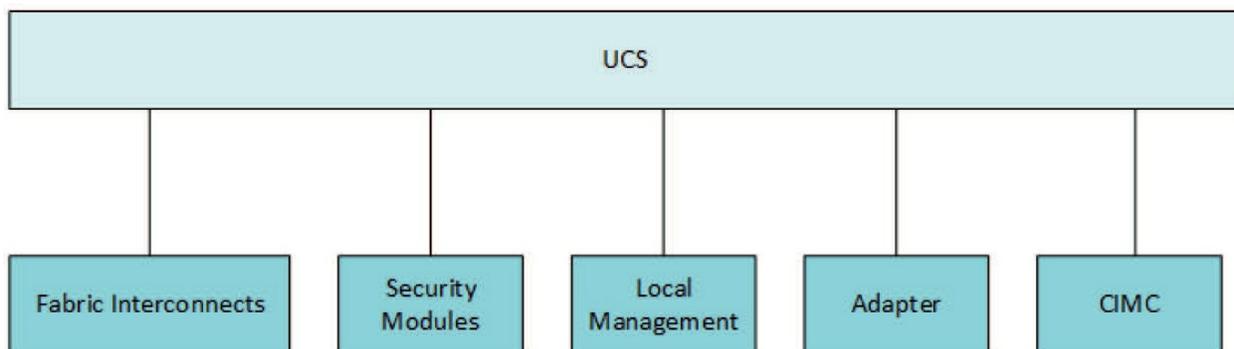
FTD 1150-1 172.16.10.12 vFTD-19 10.11.10.190

FTD 1150-2 172.16.10.13 vFTD-20 10.11.10.200

FXOS Overview

Now that we can access our 4100s, let's explore the Firepower eXtensible Operating System (FXOS) that runs on the box. Unlike the simpler OS that the ASA runs, FXOS is a hybrid OS that's made from the UCS servers because these boxes use the UCS architecture, like fabric interconnects, internally.

FXOS has several contexts that it glues together to provide our Firepower Cluster Manager that we'll use to configure our physical boxes. No worries—we'll have a more detailed look at each of them next.



UCS Context

This is the default context that you'll see if you SSH or console into the box after the wizard is done. If you've ever used UCS Manager's CLI before, then you'll be right at home! For our purposes, right now you really only need to know how to do basic tasks:

A4140-1 # ?
acknowledge backup
clear
commit-buffer

connect
decommission Acknowledge
Backup
Clear managed objects
Commit transaction buffer Connect to Another CLI
Decommission managed objects

decommission-secure Decommission managed objects in secure delete
discard-buffer

end
exit
recommission remove

restore-check scope
set

show
terminal
top

ucspe-copy

Delete managed objects Discard transaction buffer Go to exec mode Exit from command interpreter
Recommission Server Resources Remove Check if in restore mode Changes the current mode Set
property values Show system information Set terminal line parameters Go to the top mode Copy a file
in UCSPE

up Go up one mode

where Show information about the current mode For example, if I wanted to change my
4140's MGMT IP, I'll just do the following to change the fabric interconnect
address:

```
A4140-1 # scope fabric-interconnect a
A4140-1 /fabric-interconnect # set out-of-band ip 10.11.10.207 netmask 255.255.255.0 gw 10.11.10.1
A4140-1 /fabric-interconnect # commit
```

Note: You need to commit changes before they'll take effect in this CLI! We can change CLI contexts by using the **connect** command like this:

```
A4140-1 # connect ?
adapter Mezzanine Adapter
cimc Cisco Integrated Management Controller fxos Connect to FXOS CLI
local-mgmt Connect to Local Management CLI
```

module Security Module Console

FXOS Context

The next context I'm going to show you is **fxos**, which is actually the fabric interconnect that lives in the 4100/9300 and runs the Nexus OS since it's the switch fabric. So first, we covered UCS and now Nexus...surprise! You're suddenly reading a CCNP Data Center book, right?

```
A4140-1# connect fxos
```

We can't adjust any configuration here, but we can check out how the switch is set up if we're troubleshooting something or maybe just curious. For example, I can use **show run** to see what features are enabled on the switch like this:

```
A4140-1(fxos) # show run | include ^feature feature-set virtualization
feature npiv
feature tacacs+
feature private-vlan
feature port-security
feature udd
feature lacp
feature vmfex
feature lldp
feature fex
feature network-segmentation-manager
```

Security Module Context

This context is where the security modules live. We're really not far enough to have anything to look at yet, but for now I'll point out that we can use either **console** or **telnet** to connect to the module.

It doesn't really matter which one I pick here because there's only one module for the 4100. Just know that if you want to connect to a couple modules at once, then telnet would be the way to go.

```
A4140-1 # connect module 1 ?
console Console
telnet Telnet
```

Local Management Context

The Local Management context is pretty similar to the privilege mode in an ASA or a Cisco router. It lets you use some basic troubleshooting tools like **ping** or **traceroute**, check some logs, and you can do some cluster stuff from here as well:

```
A4140-1# connect local-mgmt A4140-1(local-mgmt)# ? cd
clear
cluster connect copy
cp
delete dir
enable Enable
end Go to exec mode
erase Erase
erase-log-config Erase the mgmt logging config file exit
fips
ls
mgmt-port
mkdir
move
mv
ping
ping6
pwd
reboot
restore-check rm
rmdir
run-script
show
shutdown
ssh
tail-mgmt-log telnet
terminal
top
traceroute
traceroute6 verify
Change current directory Clear managed objects Cluster mode
Connect to Another CLI Copy a file
Copy a file
Delete managed objects Show content of dir
```

Exit from command interpreter FIPS compliance

Show content of dir

Management Port

Create a directory

Move a file

Move a file

Test network reachability Test IPv6 network reachability Print current directory

Reboots Fabric Interconnect Check if in restore mode Remove a file

Remove a directory

Run a script
Show system information
Shutdown
SSH to another system
tail mgmt log file
Telnet to another system Set terminal line parameters Go to the top mode
Traceroute to destination Traceroute to IPv6 destination Verify Application Image

To give you a good example, I'll ping a DNS server on the Internet. Keep in mind that this is communicating through the management interface, so things can get a little muddier if the management network takes a different path than your FTD interfaces!

```
A4140-1(local-mgmt)# ping 1.1.1.1 count 5
PING 1.1.1.1 (1.1.1.1) from 10.11.10.207 eth0: 56(84) bytes of data. 64 bytes from 1.1.1.1:
icmp_seq=1 ttl=58 time=9.32 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=58 time=9.00 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=58 time=9.03 ms
64 bytes from 1.1.1.1: icmp_seq=4 ttl=58 time=8.94 ms
64 bytes from 1.1.1.1: icmp_seq=5 ttl=58 time=8.95 ms

--- 1.1.1.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4004ms rtt min/avg/max/mdev =
8.940/9.051/9.329/0.186 ms
```

Adapter Context

Frankly, you'll probably never use this context unless you're on the phone with Cisco TAC and they ask you to come here. The Adapter context basically provides low-level information on the interfaces in the box.

```
A4140-1# connect adapter 1/1/1
adapter 1/1/1 # help Available commands:

connect
exit
help
history
show-fwlist
show-identity show-phyinfo
- Connect to remote debug shell
- Exit from subshell
- List available commands
- Show command history
- Show firmware versions on the adapter
- Show adapter identity
```

- Show adapter phy info

show-systemstatus - Show adapter status

CIMC Context

Hopefully, you won't use the Cisco Integrated Management Controller (CIMC) context either unless you need to do a low-level restore. This context does provide some debugging tools for the CIMC on the box.

```
A4140-1 # connect cimc 1/1
```

```
Trying 127.5.1.1...
```

```
Connected to 127.5.1.1.
```

```
Escape character is '^'.
```

```
CIMC Debug Firmware Utility Shell [ support ] [ help ]# help
```

```
Debug Firmware Utility
```

```
Command List
```

```
_____ alarms
```

```
cores
```

```
dimmb1
```

```
exit
```

```
i2cstats
```

```
images
```

```
mctools
```

```
memory
```

```
messages
```

```
mrcout
```

```
network
```

```
obfl
```

```
post
```

```
power
```

```
programmables
```

```
sensors
```

```
sel
```

```
fru
```

```
tasks
```

```
top
```

```
update
```

```
users
```

```
version
```

```
cert
```

```
sldp
```

help
help [COMMAND]

Notes:

“enter Key” will execute last command “COMMAND ?” will execute help for that command

Image Management

Before you configure instances, you have to download the image files you want to use from Cisco. There are two types of files that you’ll have to work with.

The FXOS image file has a SPA extension that’s used to update the Firepower Chassis Manager. and it’s up to you to stay up to date in order to support the newer instance versions. For example, to deploy a FTD 6.6 VM, you’ll need FXOS 2.8.1 to be running on the FCM.

The instance image file has a CSP extension and contains the actual solution image we want to deploy. As of today, this can be an ASA, FTD, or Radware’s VDP file.

Adding FXOS

To upload a new FXOS image in the Chassis Manager, go to **System>Updates** and then click **Upload Image**.

You’ll then need to select the SPA file you downloaded from Cisco:

 Refresh

Upload Image

Filter..



26 Aug 2019, 08:42 PM



26 Aug 2019, 08:42 PM



Upload Image



Select File:

Choose File

fxos-k9.2.8.1.52.SPA

Upload

Cancel

Upgrading FXOS

To upgrade the FXOS on the 4100/9300, click the left-most icon on the right side of the image name in the update page—it's the square with the arrows:

Available Updates

Image Name	Type	Version	Status	Build Date	Image Integrity	
fxos-k9.2.6.1.157.SPA	platform-bundle	2.6(1.157)	Not-Installed	04/17/2019	✓ Verified - Mon 26 Aug 2019, 08:42...	 
fxos-k9.2.7.1.52.SPA	platform-bundle	2.7(1.52)	Installed	06/18/2019	✓ Verified - Mon 26 Aug 2019, 08:42...	
fxos-k9.2.8.1.52.SPA	platform-bundle	2.8(1.52)	Not-Installed	11/08/2019	✓ Verified - Sun 29 Dec 2019, 02:18...	 
cisco-ftd.6.4.0.102.csp	ftd	6.4.0.102	Not-Installed	04/23/2019		
cisco-ftd.6.3.0.85.csp	ftd	6.3.0.85	Not-Installed	01/19/2019		
cisco-ftd.6.5.0.32.csp	ftd	6.5.0.32	Installed	06/21/2019		

You'll get a prompt letting you know that the Chassis will be down during the upgrade and will reboot at the end.

To start the upgrade, just click Yes:

Not-Installed

04/17/2019

Installed

05/10/2019

Update Bundle Image



Please ensure Application configuration is saved. All existing sessions will be terminated and FCM will not be accessible during the process. It may take several minutes. Chassis will reboot after upgrade, please re-login to FCM after upgrade completes.

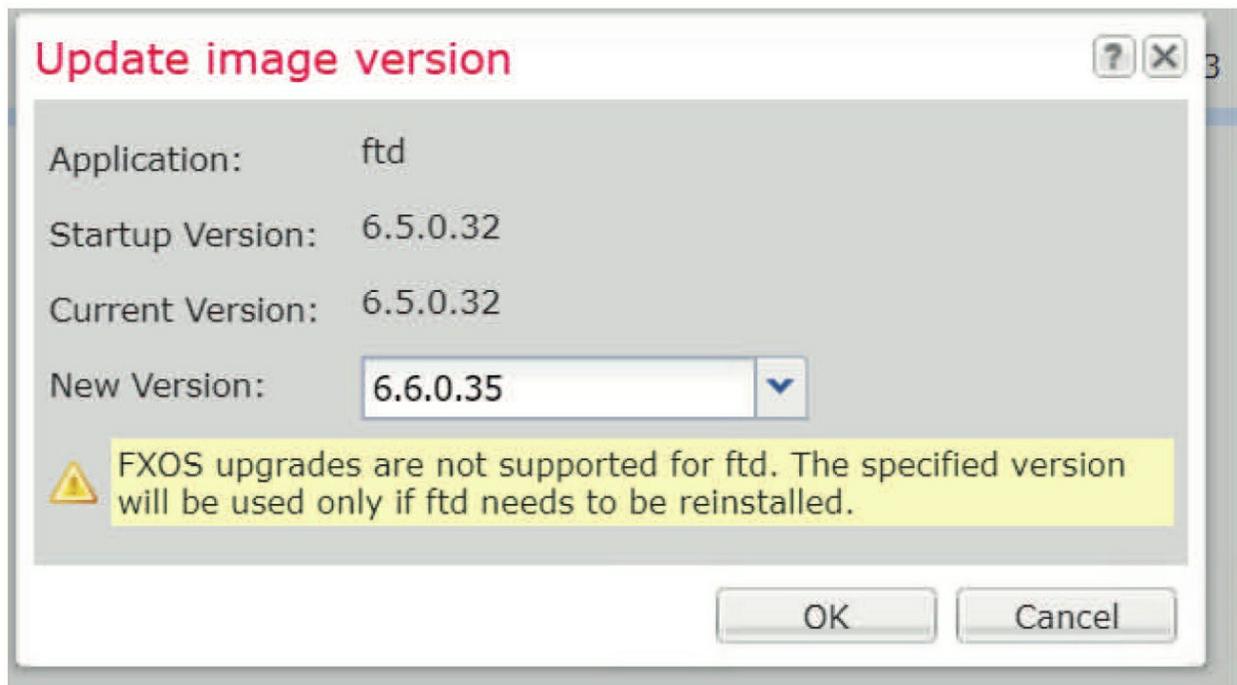
Selected version 2.8(1.52) will be installed. Do you want to proceed?

Yes

No

If you already have a logical device deployed, the prompt might tell you that your instance(s) are out of date after the upgrade and won't be usable until you upgrade the Logical Device to a supported version or reinstall the image.

You can see below that I have 6.5 FTD images loaded, but the FXOS upgrade doesn't support the new 6.6 image. This won't be an issue because I'll upgrade the FTD image after the FXOS is installed:

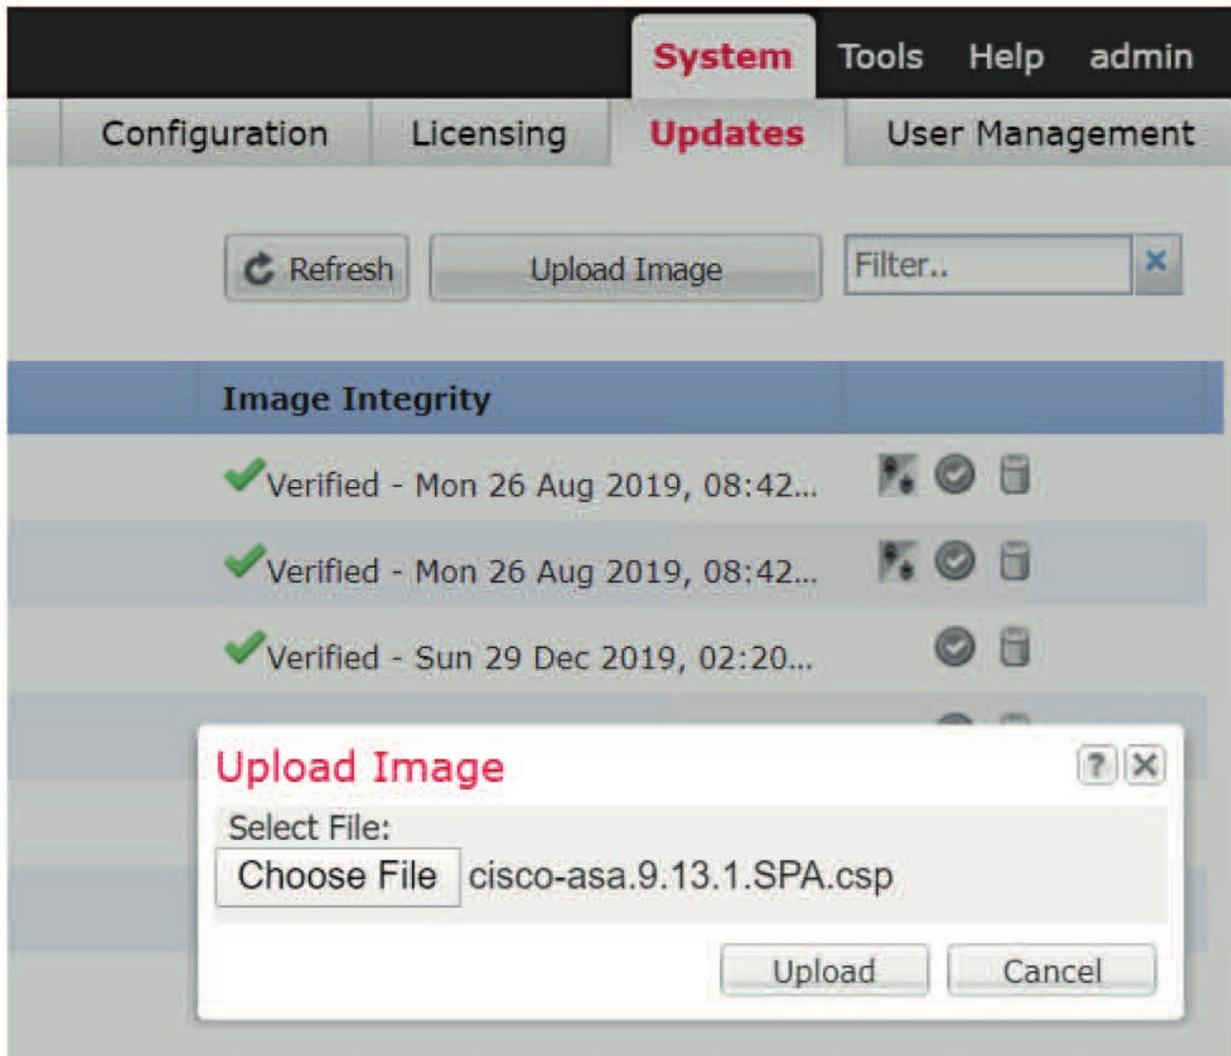


I'm jumping ahead a bit here, but if you've a deployed instance that becomes unsupported, you can go to Logical Devices and then click the Set Version button or select it from the menu.

Adding a ASA Image

The first instance image we're going to add is the ASA just because it's nice and alphabetical.

To upload the image, simply press the Upload Image button again and select the ASA CSP file you've downloaded:



Once the image has been uploaded, the FCM will automatically verify the image to make sure it looks healthy. If you want to run the verification again, click the green circle with a check mark on the right-hand side of the entry.

Adding an FTD Image

To add an FTD image, I'll actually do the same thing and just upload an image. The only difference here from the ASA is that I'll have to accept a license agreement at the end of the upload process.

Once you sign your life away, the image is ready to use!

Successfully Uploaded cisco-ftd.6.6.0.35.SPA.csp

End User License Agreement

Effective: May 22, 2017

This is an agreement between You and Cisco Systems, Inc. or its affiliates ("**Cisco**") and governs your Use of Cisco Software. "**You**" and "**Your**" means the individual or legal entity licensing the Software under this EULA. "**Use**" or "**Using**" means to download, install, activate, access or otherwise use the Software. "**Software**" means the Cisco computer programs and any Upgrades made available to You by an Approved Source and licensed to You by Cisco. "**Documentation**" is the Cisco user or technical manuals, training materials, specifications or other documentation applicable to the Software and made available to You by an Approved Source. "**Approved Source**" means (i) Cisco or (ii) the Cisco authorized reseller, distributor or systems integrator from whom you acquired the Software. "**Entitlement**" means the license detail; including license metric, duration, and quantity provided in a product ID (PID) published on Cisco's price list, claim certificate or right to use notification. "**Upgrades**" means all updates, upgrades, bug fixes, error corrections, enhancements and other modifications to the Software and backup copies thereof.

This agreement, any supplemental license terms and any specific product terms at www.cisco.com/go/softwareterms (collectively, the "**EULA**") govern Your Use of the Software.

1. **Acceptance of Terms.** By Using the Software, You agree to be bound by the terms of the EULA. If you are entering into this EULA on

I understand and accept the agreement

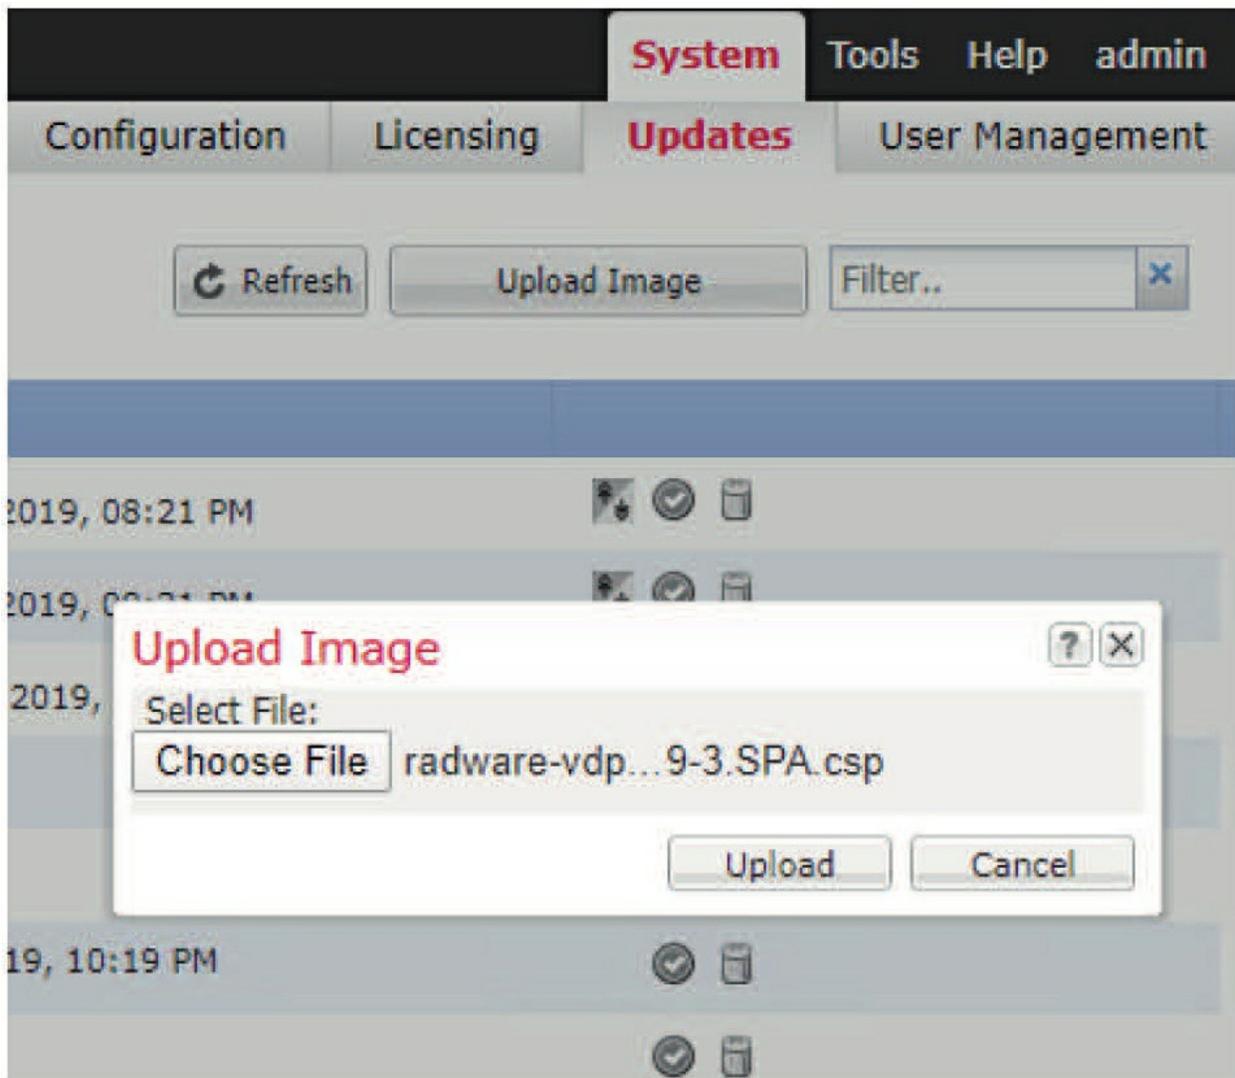
Ok

Cancel

Adding RadWare

Adding a RadWare image is the same as with the ASA—just upload the image. There's not even a license agreement to worry about:

Now that we have the ASA, FTD, and Radware images loaded into FXOS, let's log in to the Chassis Manager and start configuring them.



Logging into the Chassis Manager

Using the IP address configured for the management of our 4100s, (10.11.10.207 and 10.11.10.209), we should be able to **https:** to those IP addresses and log in using the username/passwords we set to get into the Chassis Manager:

https://10.11.10.207



Firepower Chassis Manager

admin

••••••••

Login

Support : tac@cisico.com | 1-800-553-2447

Once we get the splash page and log in, we'll be taken to the Chassis Manager home page.

Here's what the dashboard looks like when you first log in. It gives you a summary of the interfaces, system health, and any alarms you might want to know about:

The screenshot shows the Cisco Firepower Chassis Manager dashboard. At the top, there's a navigation bar with 'Overview' selected and other tabs like 'Interfaces', 'Logical Devices', 'Security Engine', and 'Platform Settings'. Below the navigation bar, the device information is displayed: 'A4140-1' with IP '10.11.10.207', 'Model: Cisco Firepower 4140 Security Appliance', 'Version: 2.7(1.52)', and 'Operational State: Power-problem'. The dashboard is divided into several sections: 'CONSOLE', 'MGMT', and 'USB' ports; 'Power 1 - Unknown' and 'Power 2 - Running' status indicators; 'Network Module 1' with ports 1, 3, 5, 7, 2, 4, 6, 8; 'Network Module 2: Empty' and 'Network Module 3: Empty'. Below these are summary cards for 'FAULTS' (0 Critical, 2 Major), 'INTERFACES' (8 Down, 0 Up), 'DEVICES & NETWORK' (0 Down, 0 Up), 'LICENSE' (Smart Agent Disabled), and 'INVENTORY' (1 Security Engine, 6 Fans, 2 Power Supplies). A table at the bottom lists faults:

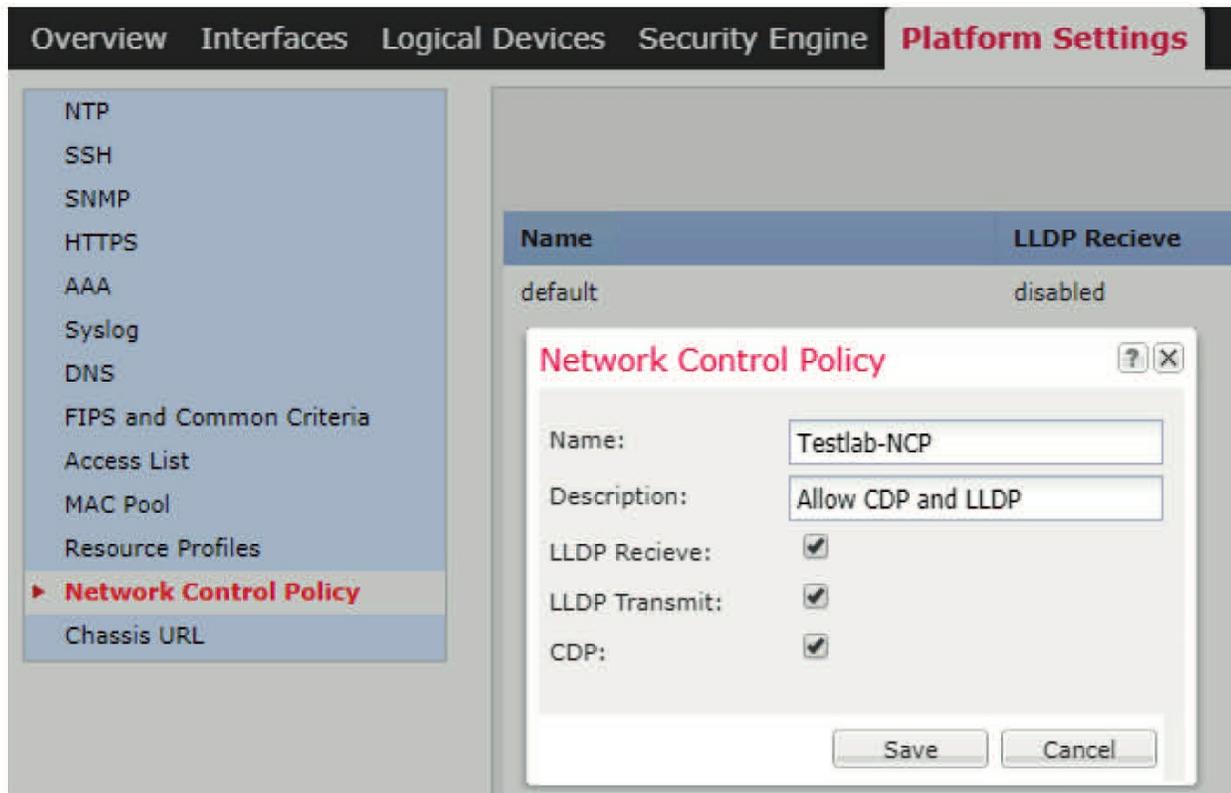
Severity	Description	Cause	Occurrence	Time	Acknowledged
MAJOR	The password encryption key has not been set.	password-encryption-key...	1	2019-11-03T20:21:19.074	no
MAJOR	Power state on chassis 1 is redundancy-failed	power-problem	1	2019-11-03T20:22:42.322	no

Notice the Network Module 1 that I'm using. You can see the gray and white boxes. Before you can start configuring your 4100/9300s, you must have a network design and the network cabled up like I do.

Ports 1 and 2 are my inside and outside network, respectively, and port 3 is my management port—it's mandatory that you have a separate port for this in the 4100/9300! I also have a direct fiber connection on port 8 to my 4140-2 for our Clustering and High Availability (HA) labs.

Platform Settings

A good place to start configuring the Firepower Chassis Manager (FCM) is the **Platform Settings** because this is where all the global settings for the box are configured. Here's an example of the Network Control Policy that configures LLDP and CDP on the device:



Most of the settings are pretty self-explanatory, especially if you've configured other boxes. For example, NTP is where you specify the NTP servers the platform uses to sync its clock.

Here's a brief on some of the less obvious ones:

Mac Pool

This setting comes from the UCS side of the platform. It allows you to pick what MAC address prefixes the platform will use when it's spinning up instances for the security module. You usually don't need to change it, but you just might want the option some day!

Resource Profiles

This setting lets you control how many resources, like RAM and CPU, an instance can consume when you create it. As an example, you might need to give containers more memory than an instance can use.

Network Control Policy

This is actually a first for Cisco firewalls and it allows you to send and

receive CDP or LLDP from the chassis interfaces. Very handy for troubleshooting!

Chassis URL

This setting lets you customize the URL that's sent to the FMC when you register the chassis.

Interfaces

Before we can install and get an instance up and running, we've got to tell the FMC about the interfaces that'll be used in the deployment. From this page we can edit interfaces, create port channels or subinterfaces, and enable/disable interfaces:

As I mentioned, interface 1/1 will be my Inside zone, 1/2 my Outside, and 1/3 will be used for management, with 1/8 connected for High Availability.

Before I create my first device, I'm going to enable the ports. Here's the output when editing interface Ethernet1/1:
Okay—so the first thing I'm going to do is click the Enable box.

Edit Interface - Ethernet1/1



Name:

Ethernet1/1 Enable

Type:

data

Admin Speed:

data

Auto Negotiation:

mgmt

Admin Duplex:

firepower-eventing

Network Control
Policy:

data-sharing

default

OK

Cancel

We don't need to tell the chassis manager about any interface IP addresses, but we'll have to give it the physical information like the link speed.

We'll also need to tell it about which mode the interface will be running, and the available choices are as follows:

Data

This is the default port type for sending traffic and it's the one you'll be using in most cases.

Mgmt

This specifies the interface that'll be used as a management port by the security modules. We'd use this to join the FTD to FMC. You need a management interface defined before you can run an instance.

Firepower-eventing

This one's an optional interface that lets Firepower send out all event data through a dedicated interface.

Data-sharing

This one's the same as the data interface type except it can be used to share a connection between containers if you are running multiple instances...more on that in the multi-instance chapter 28!

One thing I want to point out here is that if you want to use a port channel in your instance, you must create the port channel on this page and not in the FTD configuration!

Port channels also give you access to the Cluster interface type, which predictably is used if you're clustering things. Port-Channel 48 will be created by default for clustering and we'll be using that a bit later on.

Ports Ethernet1/1 and Ethernet1/2 are just data interfaces, so I'll enable and configure both of them as data and 1gbps as shown here:

Remember that Ethernet1/3 is configured as my dedicated management interface. This is a mandatory interface configuration for the 4100/9300.

Overview **Interfaces** Logical Devices Security Engine Platform Settings System Tools

The screenshot displays the 'Interfaces' configuration page. At the top, there's a hardware bypass diagram showing Network Module 1 with ports 1-8, and Network Modules 2 and 3 which are empty. Below this is a table of all interfaces.

Interface	Type	Admin Speed	Operational Speed	Instances	VLAN	Admin Duplex	Auto Negotiation	Operation State	Admin State
MGMT	Management								<input checked="" type="checkbox"/>
Port-channel48	cluster	10gbps	indeterminate			Full Duplex	no	failed	<input type="checkbox"/>
Ethernet1/1	data	1gbps	1gbps			Full Duplex	no	up	<input checked="" type="checkbox"/>
Ethernet1/2	data	1gbps	1gbps			Full Duplex	no	up	<input checked="" type="checkbox"/>
Ethernet1/3	mgmt	1gbps	1gbps			Full Duplex	no	up	<input checked="" type="checkbox"/>
Ethernet1/4	data	10gbps	10gbps			Full Duplex	no	sfp-not-present	<input type="checkbox"/>
Ethernet1/5	data	10gbps	10gbps			Full Duplex	no	sfp-not-present	<input type="checkbox"/>
Ethernet1/6	data	10gbps	10gbps			Full Duplex	no	sfp-not-present	<input type="checkbox"/>
Ethernet1/7	data	10gbps	10gbps			Full Duplex	no	sfp-not-present	<input type="checkbox"/>
Ethernet1/8	data	1gbps	1gbps			Full Duplex	no	up	<input checked="" type="checkbox"/>

In the screen below, we see that all ports are up: E1/1, E1/2, E1/3 management, and E1/8 for HA. These enabled ports are the only ones that'll be available when you go configure your instance:

Logical Devices

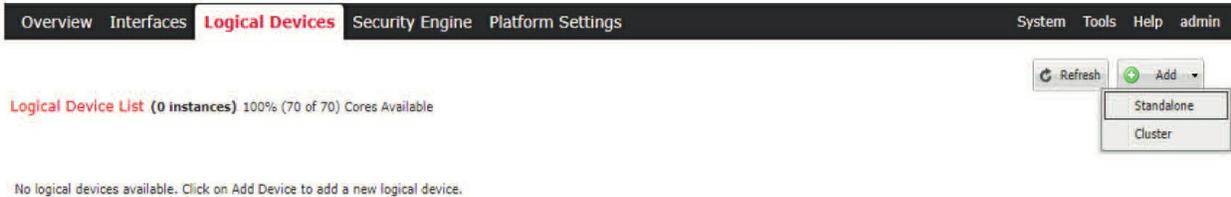
Now that we have our Firepower Chassis Manager set up and we've uploaded some images, it's time to let it run some security solutions for us!

I'm sure you've noticed that throughout this chapter, I've mostly referred to Logical Devices as instances. This is because the security module creates a Virtual Machine instance when you create a Logical Device and you'll see both terms in the Logical Device page.

Adaptive Security Appliance (ASA)

The first Logical Device we'll create is the ASA. We're also going to include Radware as an add-on coming up soon.

One thing to point out here is that the page says we have 70 cores available to us... that's a lot! Higher model numbers will give us even more resources to play with:



To create a Logical Device, click the Add button, which will give us a choice between Standalone and Cluster. I'm going to pick Standalone since we'll look at the Cluster option later in this chapter.

Next, the FCM is going to ask us for the following information:

Device Name

This is the name of the device. It can be anything that makes sense to you.

Template

This specifies which platform is being installed. It can be either ASA or FTD. Radware isn't an option here because it's an add-on.

Image Version

This is the firmware version to install. It will be what you've uploaded, and if you have multiple versions on the box, you can pick the one you want.

Instance Type

The Instance Type controls whether or not the instance will be a container or not. No worries—I'll tell you more about multiinstances in multi-instance in Chapter 28.

Add Standalone ? X

Device Name: ASA01

Template: Cisco: Adaptive Security Appliance ▼

Image Version: 9.13.1 ▼

Instance Type: Native ▼

OK Cancel

Okay, so after we press **OK** here, we'll be taken to a pretty cool page where we'll tell FCM which interfaces the Logical Device is going to use.

I'll go with the same interfaces I did earlier, so E1/1 for the Inside, E1/2 for the Outside, leaving E1/8 to be used for High Availability. We'll get this done by just clicking the interfaces that'll be used and the page will update to show them in the device topology. Next, I'll click the big ASA square where it says, "ASA- 9.13.1 Click to configure," so we can configure the device-specific settings.

Provisioning - ASA01



Standalone | Cisco: Adaptive Security Appliance | 9.13.1



Data Ports

- Ethernet1/1
- Ethernet1/2
- Ethernet1/4
- Ethernet1/5
- Ethernet1/6
- Ethernet1/7
- Ethernet1/8



Decorators



Application	Version	Resource Profile	Management IP	Gateway	Management Port	Status
ASA	9.13.1					
Interface Name		Type				
Ethernet1/1		data				
Ethernet1/2		data				
Ethernet1/8		data				

In the **General Information** section, it will want to know what to use for the management interface, plus the IP information that the management interface will use.

The image displays two overlapping windows from the Cisco Adaptive Security Appliance - Bootstrap Configuration utility. The left window is titled "Cisco: Adaptive Security Appliance - Bootstrap Configuration" and shows the "General Information" section. It includes a "Management Interface" dropdown set to "Ethernet1/3", an "Address Type" dropdown set to "IPv4 only", and input fields for "Management IP" (10.11.10.211), "Network Mask" (255.255.255.0), and "Network Gateway" (10.11.10.1). The right window is also titled "Cisco: Adaptive Security Appliance - Bootstrap Configuration" and shows the "Settings" section. It includes a "Firewall Mode" dropdown set to "Routed", and two password input fields labeled "Password" and "Confirm Password", both containing masked characters (dots).

The **Settings** section asks us what firewall mode to run – either routed or transparent, and also the password to set.

When we're done with that, we'll press **OK** to go back to the main design page.

RadWare

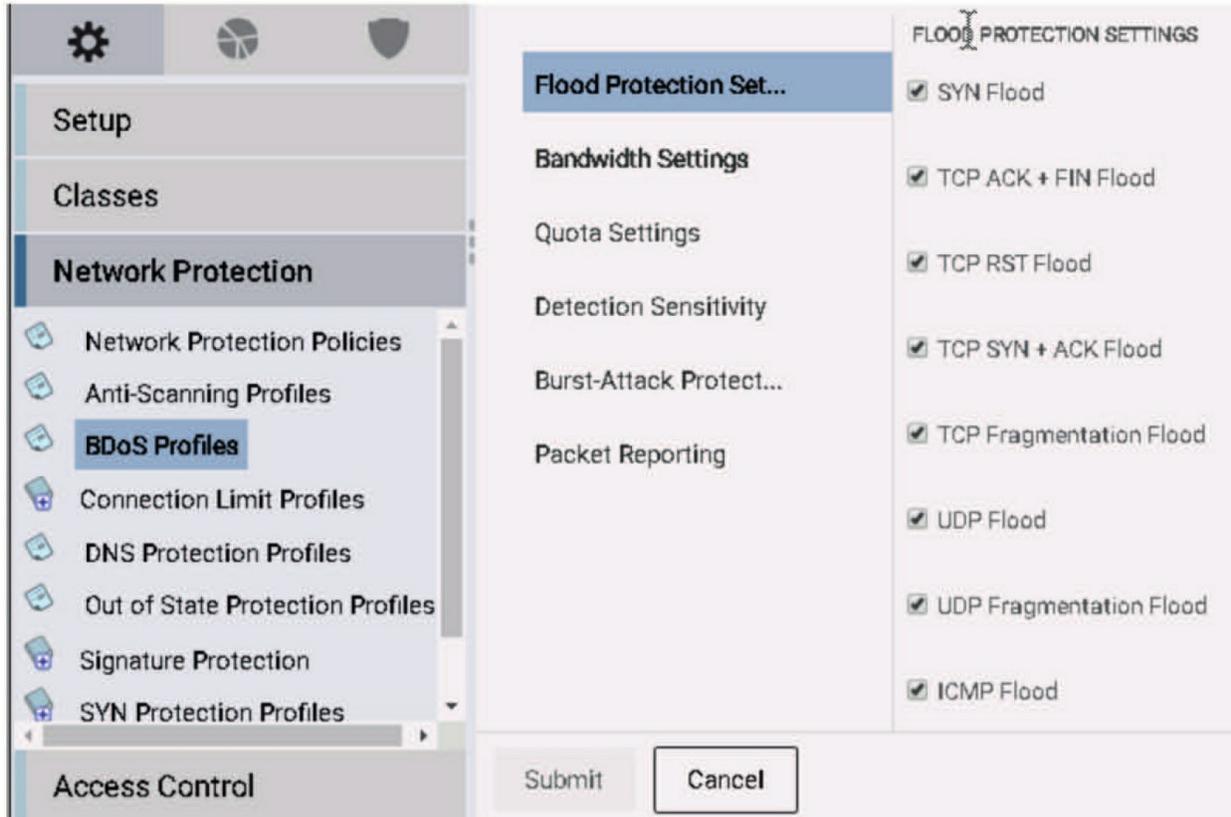
A DDoS attack generally works by overwhelming your firewall and/or your web solution to the point where it can't process legitimate traffic anymore. vDefensePro tries to prevent this from happening by acting like a transparent firewall of sorts. It transparently connects inline between the Internet connection and the ASA/FTD's outside interface so that it can prevent DDoS traffic from ever hitting your firewall or your company's online resources vulnerable to this kind of attack.

Here's an example showing an inline configuration of Radware, which can be seeing by going to the port configuration for Radware. I'll demonstrate more in a minute:



Source Port	Destination Port	Operation	Inbound Port
1	2	Process	Source
2	1	Process	Destination

Once we're inside APSolute Vision, we can configure the policies to suite our environment. A great example here would be some of the Flood Protection settings we can enable:



Of course, no anti-DDoS solution is going to be perfect because if the attack is big enough, it can actually affect your ISP and Radware can't really help you if your ISP's router blows up! Radware does have a cloud service that can try to stop a DDoS before it gets to your on-premises equipment though.

Just so you know, I could actually write a whole book on Radware, alone! Radware even has its own certification tracks that climb all the way to the expert level. But you don't really need to go there for day-to-day stuff. Radware can still add a lot of value to your Firepower appliance deployment.

So let's add Radware now. To do that, I'm going to click the VDP box under **Decorators** on the bottom left of the page to open up a similar configuration box.

Once in this page, I'll select the version to install and choose a resource profile—in this case, the default works for us. Then, just like during the ASA config, we'll be asked for the management interface and IP info.

Radware: Virtual DefensePro - Configuration ? ×

General Information

Version: 8.13.01.09-3

Resource Profile: DEFAULT-RESOURCE ⓘ

Security Module: Security Module-1

Management Interface: Ethernet1/3

DEFAULT

Address Type: IPv4 only

IPv4

Management IP: 10.11.10.210

Network Mask: 255.255.255.0

Network Gateway: 10.11.10.1

Data Ports:

- Ethernet1/1
- Ethernet1/2
- Ethernet1/8

OK Cancel

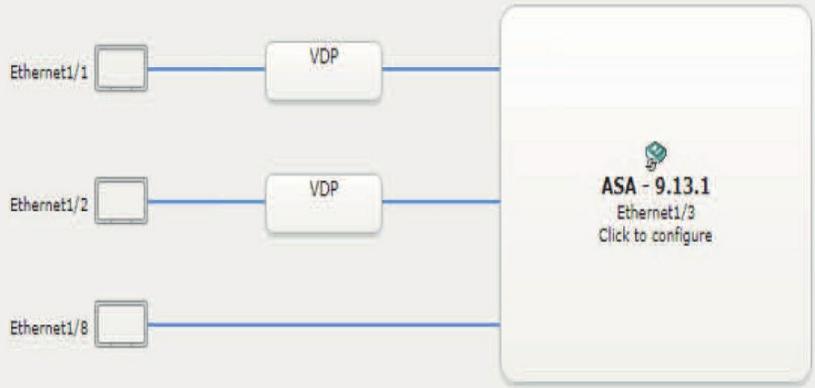
Finally, it'll ask us which interfaces the VDP should connect to. Because it's an anti-DDOS solution, I'll connect it to the outside interface of the ASA so

it can provide inline protection to incoming traffic.

Keep in mind that detailed and proper Radware design is out of scope for these books:

When I click OK, the VDP will be added to the design page. When I'm happy with everything, I'll go ahead and save the topology to deploy it. FCM will then deploy the images and configure them based upon the settings we gave it.

- Ethernet1/1
- Ethernet1/2
- Ethernet1/4
- Ethernet1/5
- Ethernet1/6
- Ethernet1/7
- Ethernet1/8



Decorators

VDP

Application	Version	Resource Profile	Management IP	Gateway	Management ...	Status
ASA	9.13.1	DEFAULT-RESOU...	10.11.10.211	10.11.10.1	Ethernet1/3	
Interface Name		Type				
Ethernet1/1		data				
Ethernet1/2		data				
Ethernet1/8		data				
VDP	8.13.01.09-3	DEFAULT-RESOU...	10.11.10.210	10.11.10.1	Ethernet1/3	
Interface Name		Type				
Ethernet1/1		data				
Ethernet1/2		data				

This takes a few minutes, so this is a good time to grab a coffee or a nap, but I don't recommend trying to do both!

FXOS CLI

Rested and recharged? Great because now that it's been installed, we can access our shiny new ASA device through the CLI. Let's take a look! First, we're going to connect to the security module and pick either console or telnet. It doesn't really matter which one for us today:

```
A4140-1 # connect module 1 console
Telnet escape character is '~'.
Trying 127.5.1.1...
Connected to 127.5.1.1.
Escape character is '~'.

CISCO Serial Over LAN:
Close Network Connection to Exit
Firepower-module1>
```

Then, when we're in the **Firepower-module1** context, we can connect to the ASA instance. Because there's only a single instance installed, we can simply say **connect asa** because it's the only type available. Of course, if we had multiple containers in the instance, we'd have to specify the container name as well:

```
Firepower-module1>connect asa
Connecting to asa(ASA01) console... hit Ctrl + A + D to return to bootCLI
```

From here, I'm going to configure the instance just as I would any other ASA. As an example, I can configure a hostname:

```
asa# conf t
asa(config)# hostname asa01
```

Now I'm going to verify the management IP is properly configured:

```
asa01(config) # show ip add
System IP Addresses:
Interface Name IP address Subnet mask Method Ethernet1/3 management 10.11.10.211 255.255.255.0
manual Current IP Addresses:
Interface Name IP address Subnet mask Method Ethernet1/3 management 10.11.10.211 255.255.255.0
manual
```

And then configure our inside interface:

```
asa01(config) # interface e1/1
```

```
asa01(config-if)# nameif inside
INFO: Security level for "inside" set to 100 by default. asa01(config-if)# ip add 192.168.10.1
255.255.255.0
asa01(config-if)# no shut
```

As well as our outside interface:

```
asa01(config-if) # interface e1/2
asa01(config-if)# nameif outside
INFO: Security level for "outside" set to 0 by default. asa01(config-if)# ip add 20.1.1.1 255.255.255.0
```

If you're keeping track of this inception, we SSH'd to the UCS management, then connected to the security module through reverse telnet. We finally connected to the ASA console...wild! At this point, it wouldn't be weird if you're wondering how you get out of this console.

To exit, press **Ctrl+A** then **D**, which will bring you back to the security module context:

```
Disconnected from asa(ASA01) console!
Firepower-module1>
```

And if you want to get all the way back to the UCS CLI, press **Shift+~** to get to the telnet prompt and then press **q** to quit:

```
telnet> q
Connection closed.
A4140-1#
```

Next, we're going to check out our Radware VM. I'll connect from the security module with **connect vdp**:

```
Firepower-module1> connect vdp
Connected to domain vDP
Escape character is ^A
```

The first thing that's going to happen here is that it'll make you set a new password... pro-tip, it accepts only letters and numbers! The password of user "radware" must be changed. Please enter the new password: *****

```
Confirm password: *****
Password change OK.
Generic Version: 01.00-00.00
Init completed successfully.
```

```
login
```

User: radware
Password: *****

Now we can use the DefensePro CLI as we please:

DefensePro#? classes

device

dp

help

login

logout

manage

net

ping

ping6

reboot

security

services

shutdown

ssh

statistics system Sets telnet

trace-route Measures hops and latency to a given destination. trace-route6 Measures hops and latency to a given ipv6 destination.

Configures traffic attributes used for classification Device Settings

DefencePro Security settings

Displays help for the specified command

Login into the device

Logout of the device

Device management configuration

Network configuration

Pings a remote host.

I don't want to get too carried away here, so I'll just enable the web interface to add it to the management product:

DefensePro#**manage secure-web status set enable**

Pings a remote IPv6 host.

Reboot the device

Device Security

General networking services

Shutdown

Connect via SSH to a remote host. Device statistics configuration. system parameters.

Connects to a remote host via telnet.

Updated successfully

Web Server SSL Status: enable

After the web interface was enabled, I added the VDP to Radware's

APolute Vision, its central management solution that lets you configure the DefensePro instance:



Sites and Devices



Status ▼ Type ▼ Name IP Address 🔍

▼ Default [2/2]

▼ Cisco-FP-vDP [2/2]

Alteon-AppWall (AW)

DefensePro (vDP)

Mixed Device Types



You have selected 2 devices:

1 Alteon Devices

1 DefensePro Devices



Summary



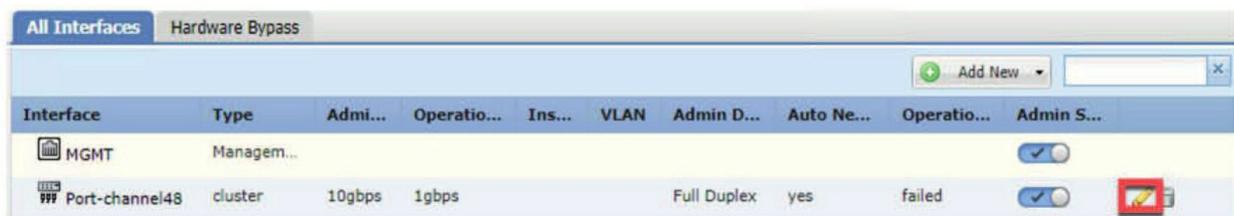
Selected Devices

I'm going to go ahead and delete the Logical Device so that we can try out FTD.

FTD Cluster

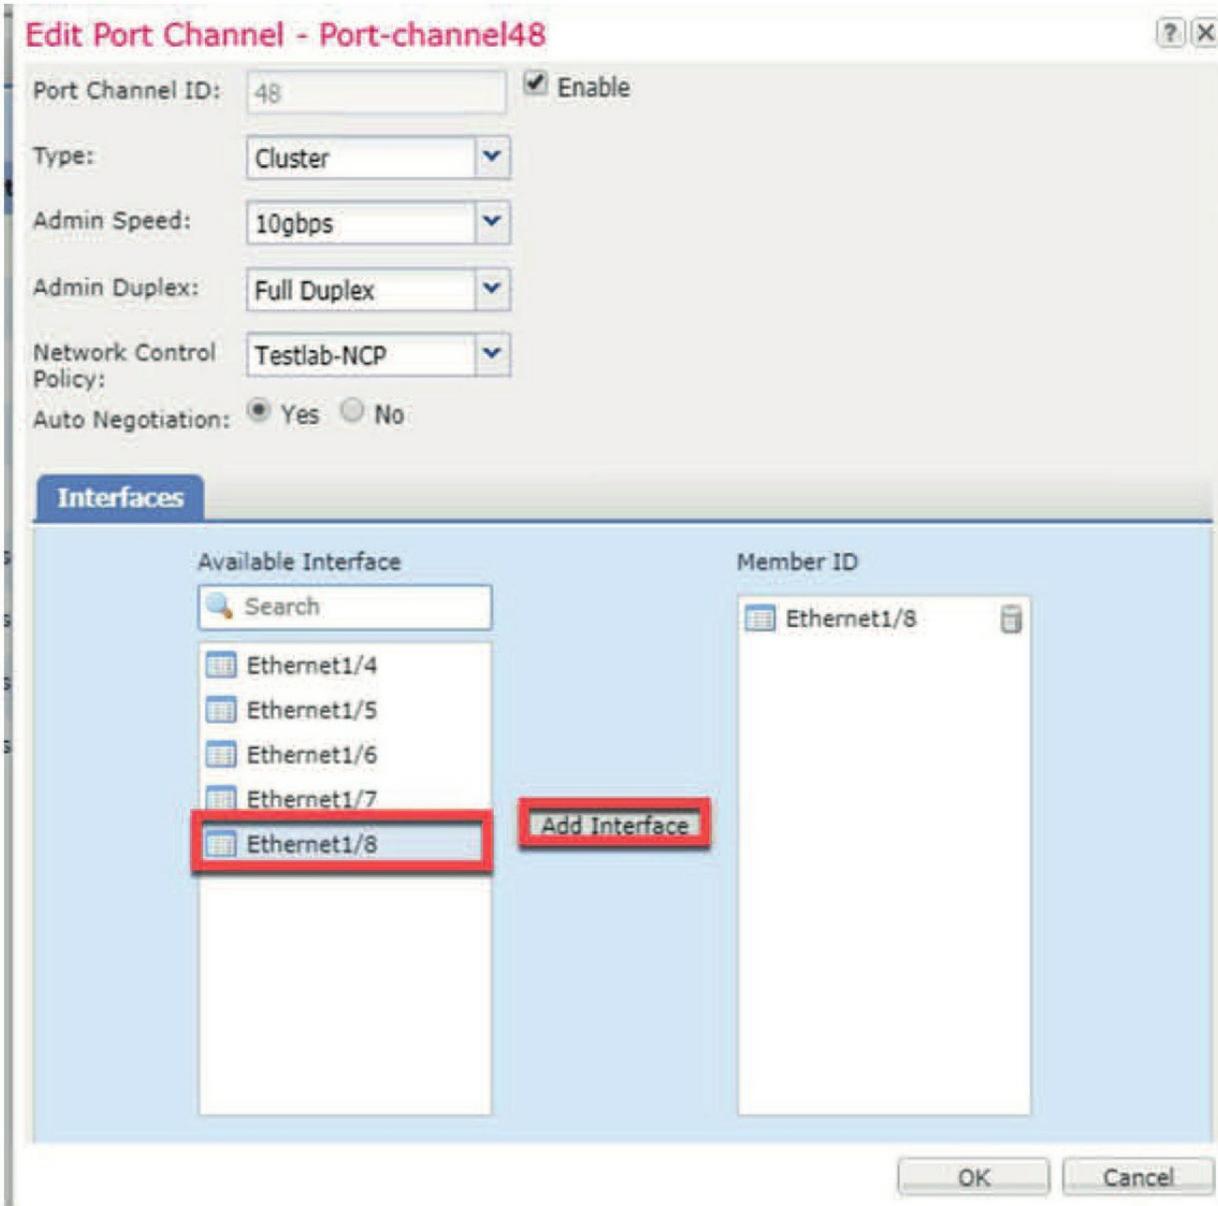
The next option we're going to explore is all about deploying FTD across two 4100/9300s. Called an inter-chassis cluster, this option is the most complicated and expensive one because it requires not one, but two pricey Firepower appliances to get it done!

The first step we'll take is to adjust the interfaces a bit on both of our 4140s. Next up is to add Ethernet1/8 into Port-Channel 48 by pressing the pencil icon next to Port-Channel 48, demonstrated below. This'll work really well for our cluster link because it's directly connected between the units.



Interface	Type	Admi...	Operatio...	Ins...	VLAN	Admin D...	Auto Ne...	Operatio...	Admin S...
MGMT	Managem...							<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Port-channel48	cluster	10gbps	1gbps			Full Duplex	yes	failed	<input checked="" type="checkbox"/> 

Notice the interface type is set to Cluster, which will be used for inter-chassis cluster communication. To add Ethernet1/8 into the port channel, simply select it, press **Add Interface**, then press **OK**.



Because the Inter-Chassis Cluster feature requires that all our data interfaces are in port-channels, we'll have to move Ethernet1/1 into Port-Channel 1 and Ethernet1/2 into Port-Channel 2.

To create a new port-channel, just press Add New and then Port Channel: I'm going to give our inside port-channel an ID of 1, add E1/1, and then I'll click OK:

Add Port Channel



Port Channel ID: Enable

Type:

Admin Speed:

Mode:

Admin Duplex:

Network Control Policy:

Auto Negotiation: Yes No

Interfaces

Available Interface

- Ethernet1/1
- Ethernet1/2
- Ethernet1/4
- Ethernet1/5
- Ethernet1/6
- Ethernet1/7
- Ethernet1/8

Member ID

Ethernet1/1

Add Interface

OK

Cancel

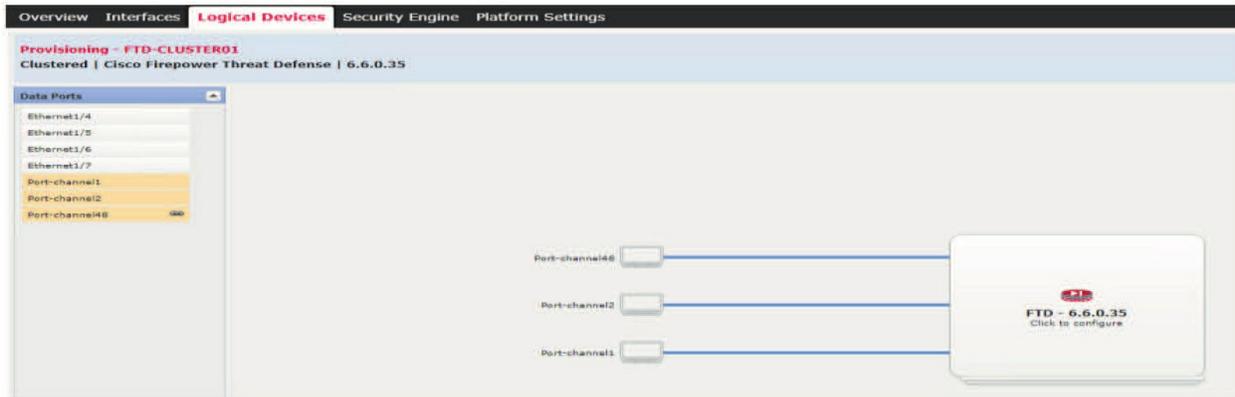
I'm going with an ID of 2 for the outside port-channel. I'll add E1/2 and then hit OK:

When we're done here, we should have a total of three port-channels: one for the inside interface, one for the outside interface, and the default one for the cluster:

Interface	Type	Admin Speed	Operational Speed	Instances	VLAN	Admin Duplex	Auto Negotiation	Operation State	Admin State
 MGMT	Management								
 Port-channel1	data	10gbps	1gbps			Full Duplex	yes	failed	
 Ethernet1/1	data	10gbps				Full Duplex	yes	down	
 Port-channel2	data	10gbps	10gbps			Full Duplex	yes	indeterminate	
 Ethernet1/2	data	10gbps				Full Duplex	yes	down	
 Port-channel48	cluster	10gbps	1gbps			Full Duplex	yes	up	
 Ethernet1/8	cluster	10gbps				Full Duplex	yes	up	
 Ethernet1/3	mgmt	1gbps	1gbps			Full Duplex	yes	up	
 Ethernet1/4	data	10gbps	10gbps			Full Duplex	no	sfp-not-present	
 Ethernet1/5	data	10gbps	10gbps			Full Duplex	no	sfp-not-present	
 Ethernet1/6	data	10gbps	10gbps			Full Duplex	no	sfp-not-present	
 Ethernet1/7	data	10gbps	10gbps			Full Duplex	no	sfp-not-present	

I'll go ahead and make the same changes to the other unit.

So, now I'm going to head over to the Logical Devices page where I'll delete any deployments already there. Then, I'll add a new one and go with Cluster.



The Add Cluster menu offers the same fields we're used to seeing, but it also lets us opt to create a new cluster or join an existing one. We're going to choose **Create New Cluster**. The rest of the steps have to do with FTD.

The process of installing an FTD Logical Device is a lot like it was for the ASA, except this time, we're going to select the FTD template. Also, the instance type is set to native, but we have an option to change it to container. We didn't get that option for the ASA because ASAs don't support the feature.

They use contexts to support multiple instances. Anyway, you can see our newly created

Add Cluster ? X

I want to: Create New Cluster

Device Name: FTD-CLUSTER01

Template: Cisco Firepower Threat Defense

Image Version: 6.6.0.35

Instance Type: Native

OK Cancel

cluster configuration here:

We'll make sure our port channels are selected in the design page as shown below, and then we click the FTD box to configure the primary settings.

Okay—so this time we're going to start out with a Cluster Information section. The Chassis ID value can be whatever number you want but it has to be unique between the pair. The Site ID value is used for inter site clustering, and we can just leave it at 1 because it doesn't really matter in our deployment.

Cisco Firepower Threat Defense - Bootstrap Configuration



Cluster Information Settings Interface Information Agreement

Security Module-1

Interface Information

Chassis ID:

Site ID:

Cluster Key:

Confirm Cluster Key:

Cluster Group Name:

Management Interface:

CCL Subnet IP:

OK

Cancel

Giving the cluster a name and a secret key so the other node can securely join is important, so we'll do that now.

We're also being asked for a Cluster Control Link (CCL) Subnet IP, but we can just enter 0.0.0.0 to have Firepower select the default IP range— (127.2.0.0/16 in case you're curious). As always, we need to specify the management interface.

As you can see, the **Settings** section for the FTD cluster sure has a lot more questions for us than the ASA did! In addition to the standard fields we saw last time, it wants to know what the FMC IP address is, the registration key, the device we will use, and the FQDN for the instance:

The **Interface Information** section is where we enter in the management IP address, netmask, and default gateway for the FTD.

Cisco Firepower Threat Defense - Bootstrap Configuration



Cluster Information **Settings** Interface Information Agreement

Registration Key:

Confirm Registration Key:

Password:

Confirm Password:

Firepower Management Center IP:

Search domains:

Firewall Mode: ▼

DNS Servers:

Firepower Management Center NAT ID:

Fully Qualified Hostname:

Eventing Interface: ▼

OK

Cancel

We don't need to do anything in the Agreement section because we already accepted the license when we imported the FTD CSP file! But if you want to have a look anyway, here it is:

Cisco Firepower Threat Defense - Bootstrap Configuration



End User License Agreement

Effective: May 22, 2017

This is an agreement between You and Cisco Systems, Inc. or its affiliates ("**Cisco**") and governs your Use of Cisco Software. "**You**" and "**Your**" means the individual or legal entity licensing the Software under this EULA. "**Use**" or "**Using**" means to download, install, activate, access or otherwise use the Software. "**Software**" means the Cisco computer programs and any Upgrades made available to You by an Approved Source and licensed to You by Cisco. "**Documentation**" is the Cisco user or technical manuals, training materials, specifications or other documentation applicable to the Software and made available to You by an Approved Source. "**Approved Source**" means (i) Cisco or (ii) the Cisco authorized reseller, distributor or systems integrator from whom you acquired the Software. "**Entitlement**" means the license detail, including license metric, duration, and quantity provided in a product ID (PID) published on Cisco's price list, claim certificate or right to use notification. "**Upgrades**" means all updates, upgrades, bug fixes, error corrections, enhancements and other modifications to the Software and backup copies thereof.

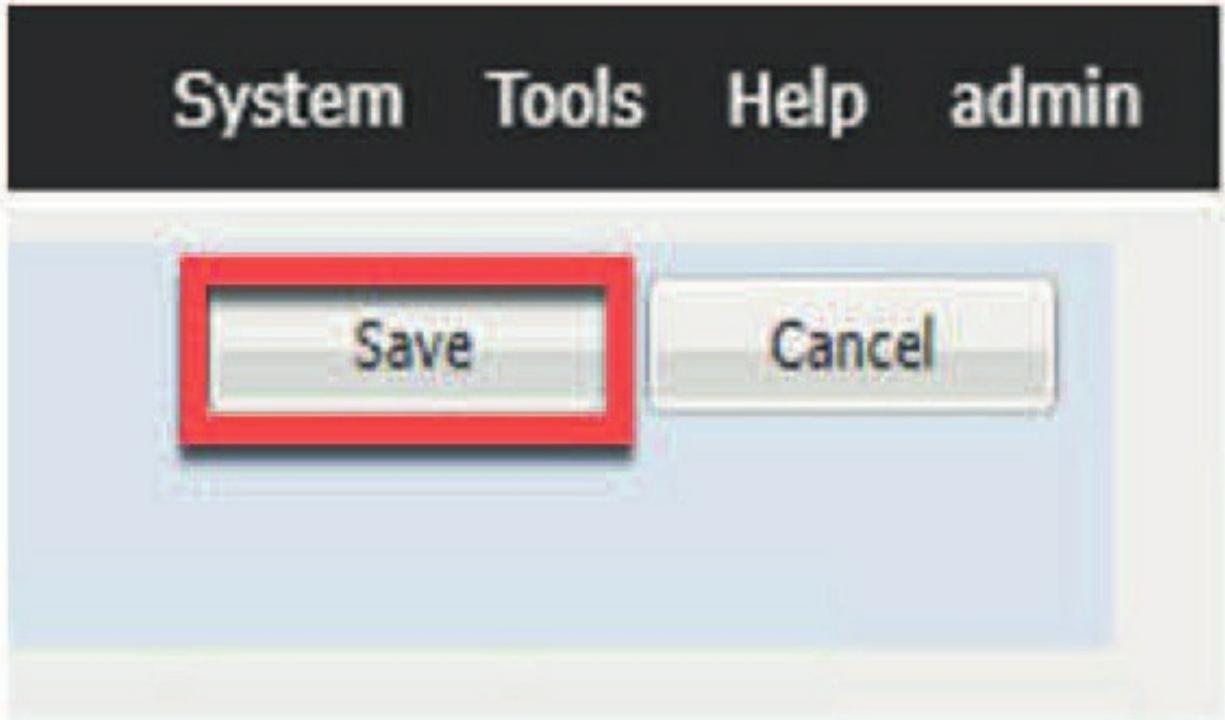
This agreement, any supplemental license terms and any

I understand and accept the agreement

OK

Cancel

To start the install, click OK to accept the config settings, and then click Save on the top right of the screen:



At this point, we just wait until the primary is done installing and showing online:

Overview Interfaces **Logical Devices** Security Engine Platform Settings System Tools Help admin

Refresh Add

Logical Device List (1 instances) 0% (0 of 70) Cores Available

FTD-CLUSTER01 Clustered Status:ok

Application	Version	Resource Profile	Management IP	Gateway	Management Port	Status
FTD	6.6.0.35		10.11.10.211	10.11.10.1	Ethernet1/3	Online

Interface Name	Type	Attributes
Port-channel1	data	Cluster Operational Status: in-cluster
Port-channel2	data	FIREPOWER-MGMT-IP: 10.11.10.211
Port-channel49	cluster	CLUSTER-ROLE: master
		CLUSTER-IP: 127.2.1.1
		MGMT-URL: https://10.11.10.205/
		UUID: 8d6f3d04-2a79-11ea-b35c-095f434ff066

Once that's done, we'll copy the cluster config so we can paste it right into the other node. To do that, choose Cluster Configuration from the menu on the right:

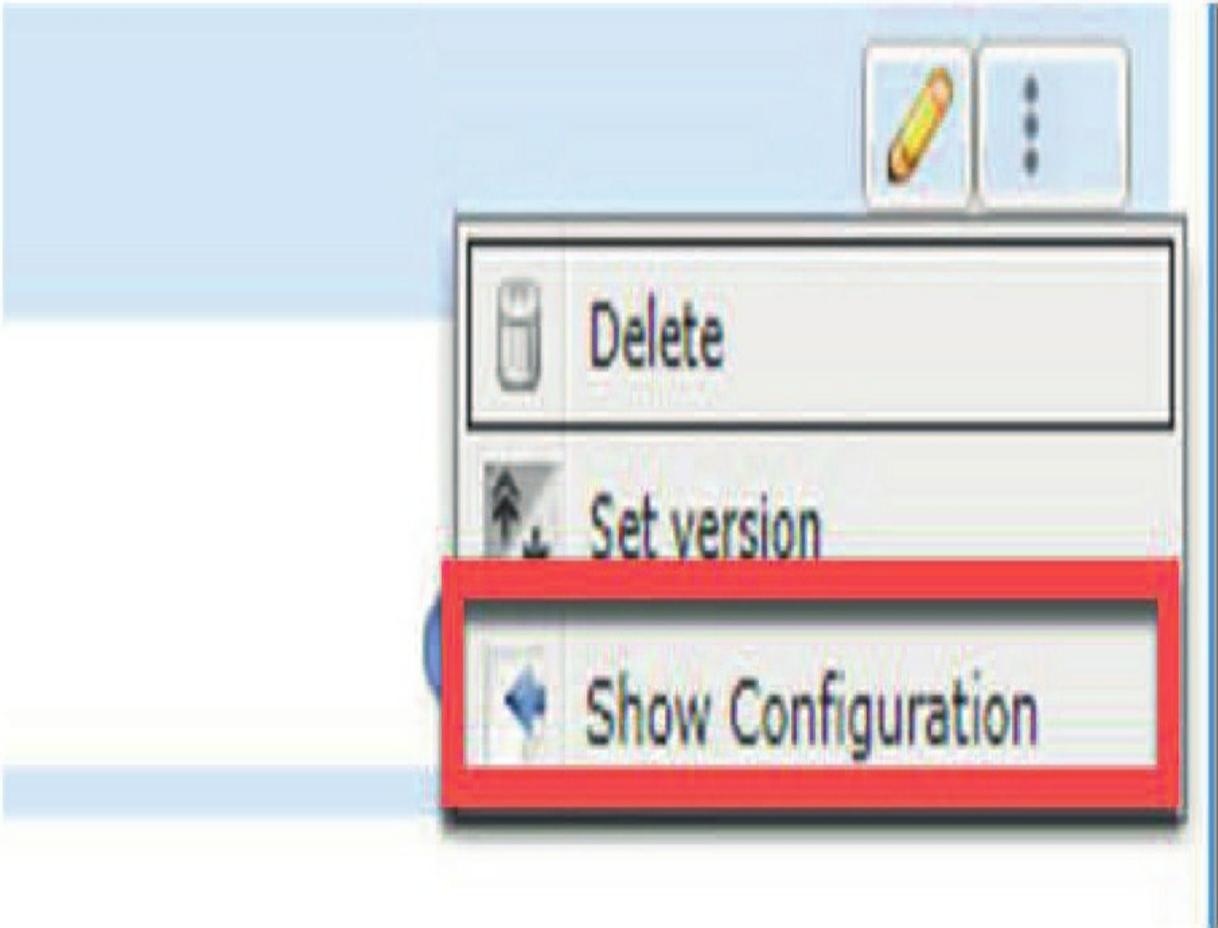


Logical Device List

(1 instances) 0% (0 of 70) Cores Available

FTD-CLUSTER01		Clustered	Status:ok			
Application	Version	Resource Profile	Management IP	Gateway	Management Port	Status
FTD	6.6.0.35		10.11.10.212	10.11.10.1	Ethernet1/3	Online

Then selecting Show Configuration



We're rewarded with a pop-up that has a bunch of JSON code we'll need to copy for the other node:

Cluster Configuration(copy to clipboard)

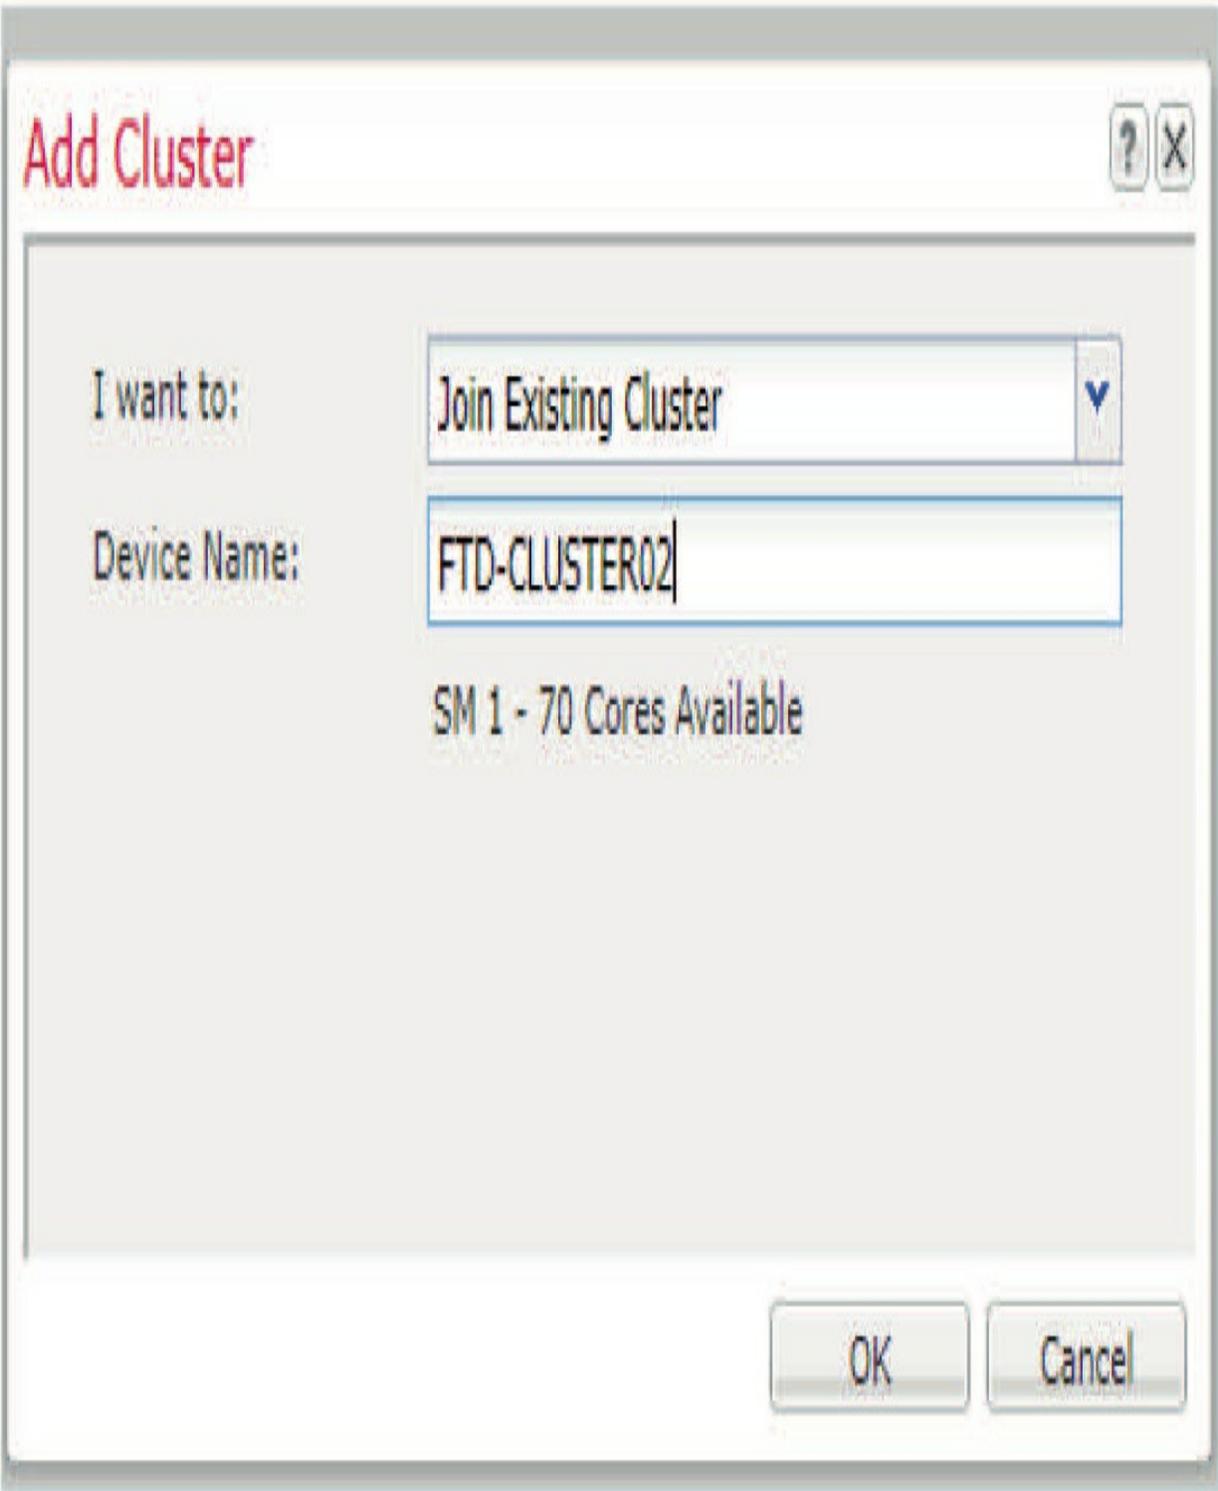


```
{ "smLogicalDevice": { "smExternalPortLink": { "smSystemMac": null, "appName": "ftd", "portDn": "fabric/lan/A/phys-slot-1-port-3", "description": "", "name": "Ethernet13_ftd", "portName": "Ethernet1/3", "linkDecorator": "", "rn": "ext-portlink-Ethernet13_ftd"}, { "smSystemMac": null, "appName": "ftd", "portDn": "fabric/lan/A/pc-48", "description": "", "name": "PC48_ftd", "portName": "Port-channel48", "linkDecorator": "", "rn": "ext-portlink-PC48_ftd"}, { "smSystemMac": { "macAddress": "28:6F:7F:02:B2:AE"}, "appName": "ftd", "portDn": "fabric/lan/A/pc-1", "description": "", "name": "PC1_ftd", "portName": "Port-channel1", "linkDecorator": "", "rn": "ext-portlink-PC1_ftd"}, { "smSystemMac": { "macAddress": "28:6F:7F:02:B2:DD"}, "appName": "ftd", "portDn": "fabric/lan/A/pc-2", "description": "", "name": "PC2_ftd", "portName": "Port-channel2", "linkDecorator": "", "rn": "ext-portlink-PC2_ftd"}, "smMgmtBootstrap": { "smIPv6": null, "smEncryptedKey": { "value": null, "key": "REGISTRATION_KEY", "rn": "encrypted-key-REGISTRATION_KEY"}, { "value": null, "key": "PASSWORD", "rn": "encrypted-key-PASSWORD"}, "smKey": { "value": "10.11.10.205", "key": "FIREPOWER_MANAGER_IP", "rn": "key-FIREPOWER_MANAGER_IP"}, { "value": "testlab.com", "key": "SEARCH_DOMAINS", "rn": "key-SEARCH_DOMAINS"}, { "value": "routed", "key": "FIREWALL_MODE", "rn": "key-FIREWALL_MODE"}, { "value": "208.67.220.220", "key": "DNS_SERVERS", "rn": "key-DNS_SERVERS"}, { "value": null, "key": "FQDN", "rn": "key-FQDN"}, "smIP": { "gateway": "10.11.10.1", "mask": "255.255.255.0", "slotId": "1", "mgmtSubType": "firepower", "ip": null, "rn": "ip-1-firepower"}, "appName": "ftd", "rn": "mgmt-bootstrap-ftd"}, "smClusterBootstrap": { "name": "CLUSTER01", "gatewayv4": "0.0.0.0", "virtualIPv4": "0.0.0.0", "gatewayv6": "", "virtualIPv6": "", "netmaskv4": "0.0.0.0", "poolEndv6": "", "poolStartv4": "0.0.0.0", "prefixLength": "", "poolStartv6": "", "poolEndv4": "0.0.0.0", "key": "", "chassisId": "1", "mode": "spanned-etherchannel", "rn": "cluster-bootstrap", "siteId": "1", "cdNetwork": "0.0.0.0"}, "dn": "Id/FTD-CLUSTER01", "description": "", "name": "FTD-CLUSTER01", "operationalState": "ok", "slotId": "1", "templateName": "ftd", "ldMode": "clustered", "errorMsg": "", "isDecorator": null}}]EOFJSON{"fabricEthLanPc": [{"adminSpeed": "10gbps", "adminState": "enabled", "bandwidth": "1", "clusterName": "", "descr": "", "dn": "ports/pc/48", "dtagVlan": "1048", "fabricEthLanPcEp": [{"adminState": "enabled", "aggrPortId": "0", "chassisId": "N/A", "ifType": "physical", "locale": "external", "membership": "up", "name": "Ethernet1/8", "operState": "up", "portId": "8", "rn": "ep-slot-1-port-8", "slotId": "1", "switchId": "A", "transport": "ether", "type": "lan", "udldOperState": "admin-disabled"}, {"flowCtrlPolicy": "default", "ifType": "aggregation", "lACPPolicyName": "default", "locale": "external", "mtu": null, "name": "Port-channel48", "operLACPPolicyName": "org-root/lACP-default", "operSpeed": "1gbps", "operState": "up", "portId": "48", "spannedCluster": "disabled", "ssaPortType": "cluster", "ssaVlanId": "1048", "stateQual": "", "switchId": "A", "transport": "ether", "type": "lan", "urlink": "https://10.11.10.207/api/ports/pc/48", "vlanStatus": "ok", "adminDuplex": "fullDuplex", "autoNeg": "yes", "fabricSubIf": null}, {"adminSpeed": "10gbps", "adminState": "enabled", "bandwidth": "10", "clusterName": "CLUSTER01", "descr": "", "dn": "ports/pc/2", "dtagVlan": "1002", "fabricEthLanPcEp": [{"adminState": "enabled", "aggrPortId": "0", "chassisId": "N/A", "ifType": "physical", "locale": "external", "membership": "suspended", "name": "Ethernet1/2", "operState": "up", "portId": "2", "rn": "ep-slot-1-port-2", "slotId": "1", "switchId": "A", "transport": "ether", "type": "lan", "udldOperState": "admin-disabled"}, {"flowCtrlPolicy": "default", "ifType": "aggregation", "lACPPolicyName": "default", "locale": "external", "mtu": null, "name": "Port-channel2", "operLACPPolicyName": "org-root/lACP-default", "operSpeed": "1gbps", "operState": "failed", "portId": "2", "spannedCluster": "enabled", "ssaPortType": "data", "ssaVlanId": "1002", "stateQual": "NonOperational"}]
```

OK

Cancel

Now over on the other node, we'll again delete any deployments there and add a cluster. Only this time, we'll pick the Join Cluster option and give the deployment a name, which can be whatever you want. I went with FTD-CLUSTER02 because it just makes sense:



When you press OK, you'll get a pop-up where you can paste all that JSON.

Copy Cluster Details



```
{ "autoNeg": "yes", "fabricSubIf": null } } EOFJSONnativeEOFJSONONE_RUEOFJSON2.8
on": "6.6.0.35", "operationalState": "online", "adminState": "enabled", "clusterOperatio
100, "urlink": "https://10.11.10.207/api/slot/1/app-inst/ftd-FTD-
d": "ftd_001_JAD2052043V5XWU5014", "clearLogData": "available", "currentJobStat

gementURL": null, "errorMsg": "", "resourceProfileName": "", "smAppAttribute":
-attribute-firepower-mgmt-ip", "urlLink": null, "value": "10.11.10.211"}, {" dn": null, "ke
null, "value": "master"}, {" dn": null, "key": "cluster-ip", "rn": "app-attribute-cluster-
", "key": "mgmt-url", "rn": "app-attribute-mgmt-url", "urlLink": null, "value": "https://10.
", "urlLink": null, "value": "a4c076cc-2b45-11ea-9d86-
ier": "FTD-CLUSTER01", "smResource":
allocatedRAM": "225150", "allocatedDataDisk": "195313", "allocatedCoreNR": "70"} } } }
```

OK

Cancel

Click OK again and after a minute you'll be taken to that design page we know and love:

Provisioning - FTD-CLUSTER02

Clustered | ftd | 6.6.0.35

Data Ports

Ethernet1/4

Ethernet1/5

Ethernet1/6

Ethernet1/7

Port-channel1

Port-channel2

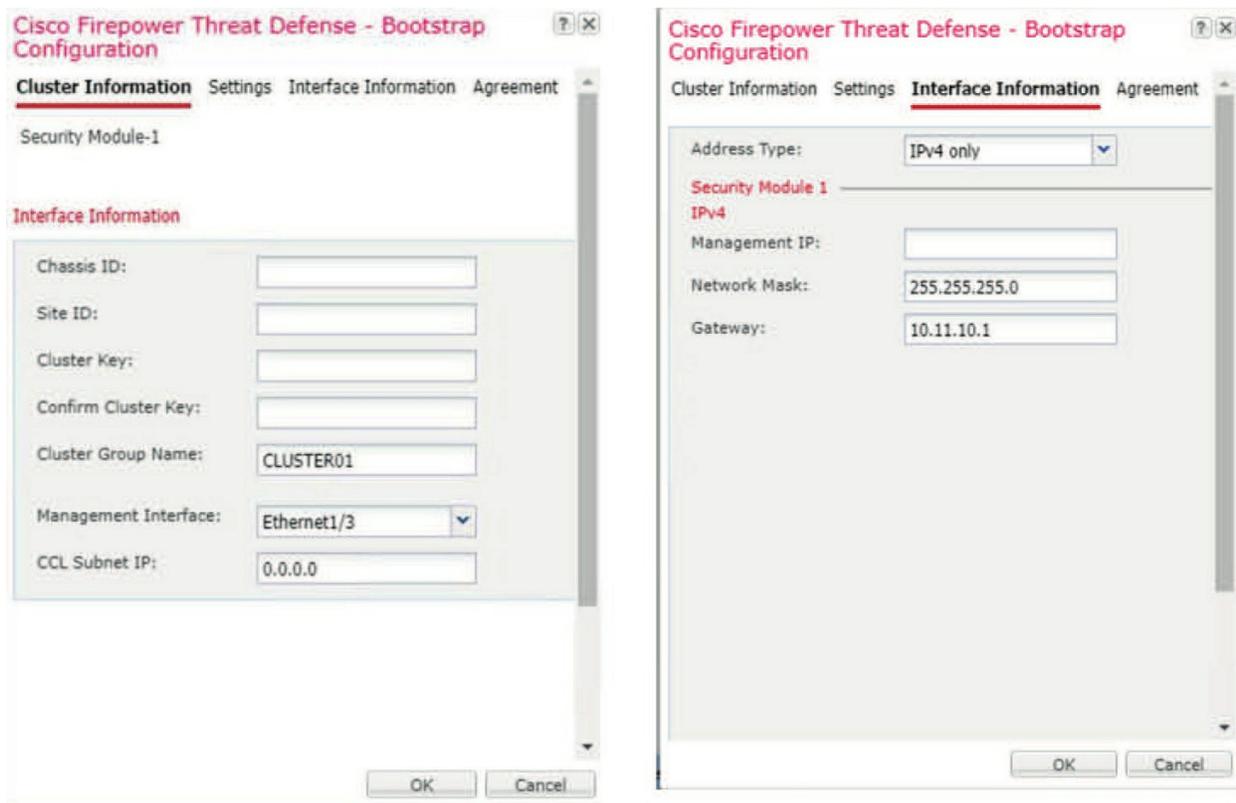
Port-channel48



Once again, we're going to select the port-channels we want to use and press the FTD to configure the settings:

It's really nice that we don't have quite as much to fill out since it prepopulates some information for us!

The main thing here is to change the chassis ID because it needs to be unique—I'll set mine to 2. We'll also have to enter the site ID (1 works fine since we aren't using it) and fill in the cluster key.



We've also got to give the secondary node its own management IP and then we can click OK and Save. After 10 min or longer, we'll have a working cluster that we can add to FMC.

One thing I want to point out here, though, is that if we connect to the FTD instance, we can see that High Availability isn't configured. That's because it's carried out at the chassis level. More on this in the HA configuration coming up in chapter 7.

When you add an inter-chassis cluster to the FMC, it automatically figures

out that it's a cluster and adds the secondary node as well:

```
A4140-1 # connect module 1 console
```

```
Telnet escape character is '~'.
```

```
Trying 127.5.1.1...
```

```
Connected to 127.5.1.1.
```

```
Escape character is '~'.
```

```
CISCO Serial Over LAN:
```

```
Close Network Connection to Exit
```

```
Firepower-module1> connect ftd
```

```
Connecting to ftd(FTD-CLUSTER01) console... enter exit to return to bootCLI
```

```
>
```

```
> show high-availability config
```

```
Failover Off
```

```
Failover unit Secondary
```

```
Failover LAN Interface: not Configured
```

```
Reconnect timeout 0:00:00
```

```
Unit Poll frequency 1 seconds, holdtime 15 seconds
```

```
Interface Poll frequency 5 seconds, holdtime 25 seconds Interface Policy 1
```

```
Monitored Interfaces 1 of 1291 maximum
```

```
MAC Address Move Notification Interval not set
```

Okay—now because the rest of this book won't be covering the cluster deployment, let's just take a minute to clean things up so we can deploy a standard FTD instance.

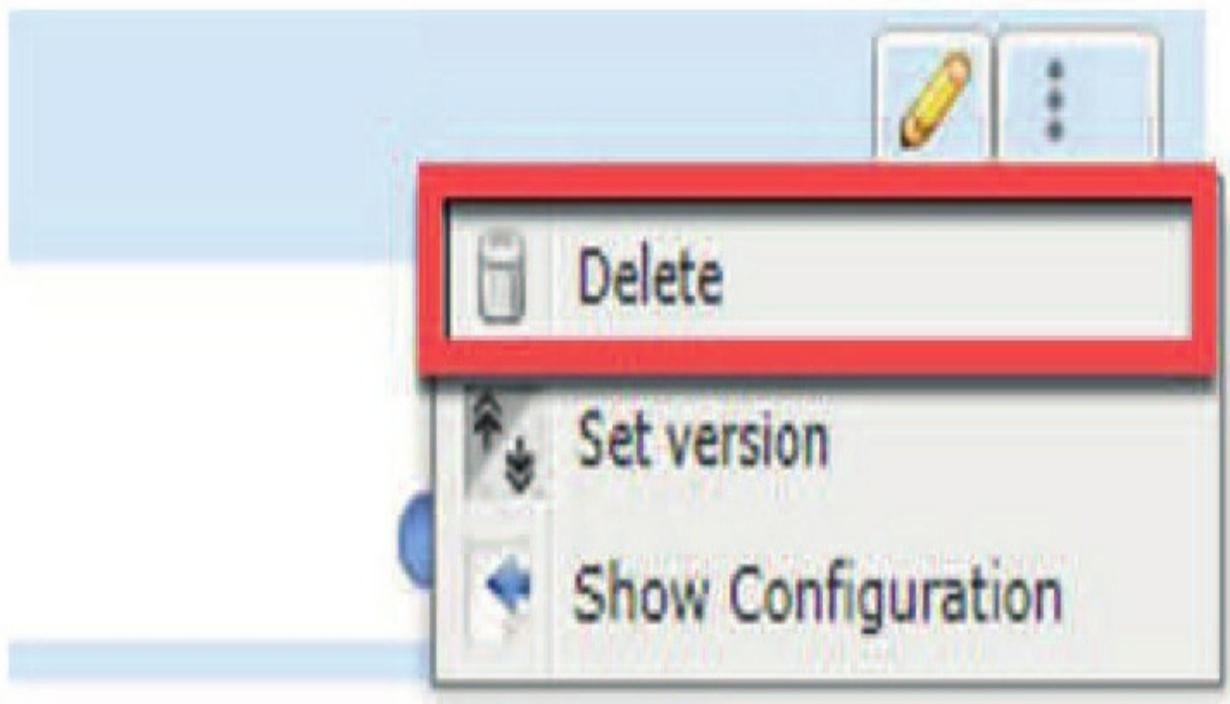


Logical Device List

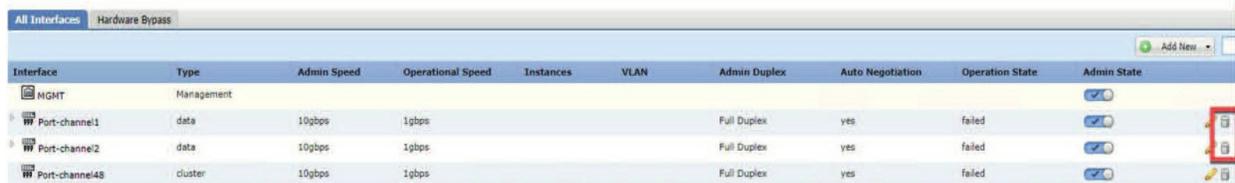
(1 instances) 0% (0 of 70) Cores Available

Application	Version	Resource Profile	Management IP	Gateway	Management Port	Status
FTD	6.6.0.35		10.11.10.212	10.11.10.1	Ethernet1/3	Online

To delete the cluster, press the menu button on the device and then click Delete:



After a minute the Logical Device will be removed. Next, we've got to set our interfaces back to normal because we made them all port channels. So, let's go back to the interface page and get rid of them! To delete Port-Channel 1 and 2, just click the trashcan icon next to them.



Interface	Type	Admin Speed	Operational Speed	Instances	VLAN	Admin Duplex	Auto Negotiation	Operation State	Admin State
MGMT	Management								<input checked="" type="checkbox"/>
Port-channel1	data	10gbps	1gbps			Full Duplex	yes	failed	<input checked="" type="checkbox"/> 
Port-channel2	data	10gbps	1gbps			Full Duplex	yes	failed	<input checked="" type="checkbox"/> 
Port-channel48	cluster	10gbps	1gbps			Full Duplex	yes	failed	<input checked="" type="checkbox"/> 

You can actually delete Port-Channel 48 too, but since it was created by default, it's probably better to just remove Ethernet1/8 from it. Let's do that now by editing the interface and removing E1/8 from PC48. Once that's done, PC48 will go back to its standard failed state.

FTD Standard

If you're thinking that installing a standard FTD instance is pretty similar to the cluster example, you're right! But this time around, we just choose Standalone:

Add Standalone ? X

Device Name:	<input type="text" value="FTD01"/>
Template:	<input type="text" value="Cisco Firepower Threat Defense"/> ▼
Image Version:	<input type="text" value="6.6.0.35"/> ▼
Instance Type:	<input type="text" value="Native"/> ▼

Just as we did back in the ASA design page, we'll select Ethernet1/1, Ethernet1/2, and Ethernet1/8 from the Data Ports menu:

Provisioning - FTD01

Standalone | Cisco Firepower Threat Defense | 6.6.0.35

Data Ports

- Ethernet1/1
- Ethernet1/2
- Ethernet1/4
- Ethernet1/5
- Ethernet1/6
- Ethernet1/7
- Ethernet1/8



Next, click on the FTD box to configure the instance settings. Just like before, the **General Information** section is asking for the management interface and IP info:

Cisco Firepower Threat Defense - Bootstrap Configuration ? X

General Information Settings Agreement

Interface Information

Management Interface:

Management

Address Type:

IPv4

Management IP:

Network Mask:

Network Gateway:

OK

Cancel

The **Settings** section for the FTD has a lot more questions for us than the ASA did!

In addition to the standard fields we saw last time, it wants to know if the management type is going to be Firepower Management Center or the standalone Firepower Device Management.

It's also asking what the FMC IP address and registration key the device will use.

Again, just as in the cluster example, we don't need to do anything with the license

Cisco Firepower Threat Defense - Bootstrap Configuration ? X

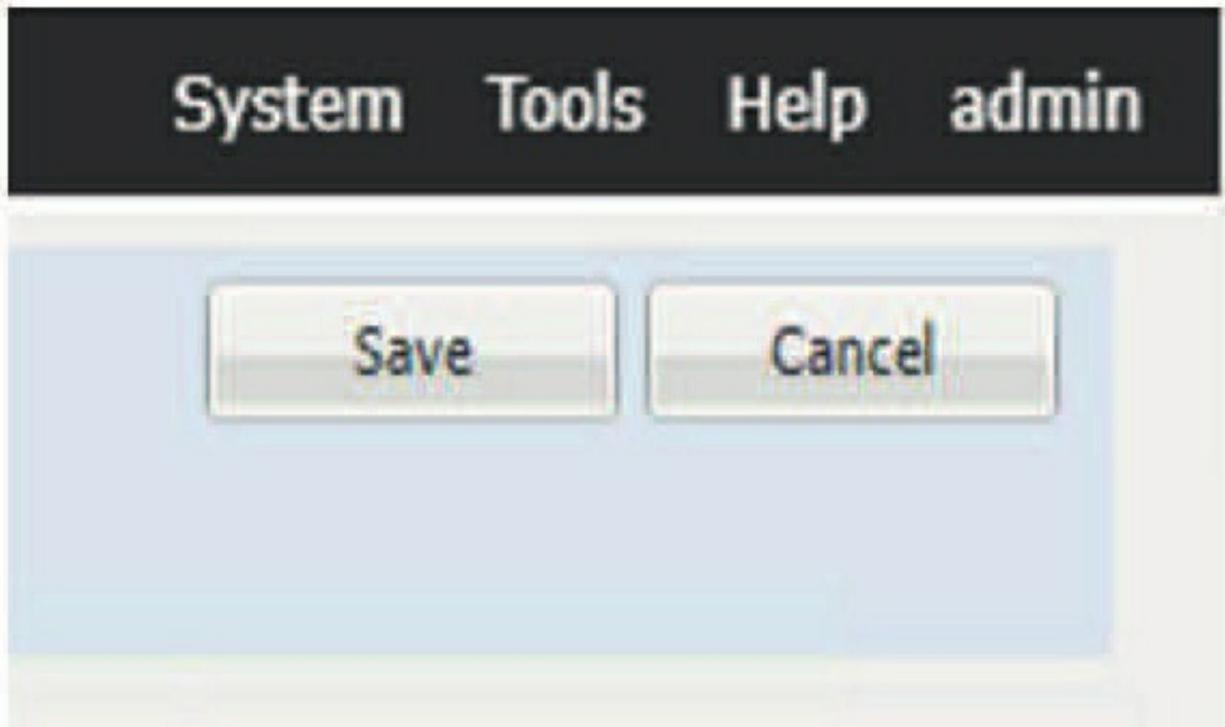
General Information Settings Agreement

Management type of application instance:	<input type="text" value="FMC"/>
Search domains:	<input type="text" value="testlab.com"/>
Firewall Mode:	<input type="text" value="Routed"/>
DNS Servers:	<input type="text" value="208.67.220.220,208.67.220.220"/>
Fully Qualified Hostname:	<input type="text" value="ftd01.testlab.com"/>
Password:	<input type="password" value="....."/>
Confirm Password:	<input type="password" value="....."/>
Registration Key:	<input type="password" value="....."/>
Confirm Registration Key:	<input type="password" value="....."/>
Firepower Management Center IP:	<input type="text" value="10.11.10.205"/>
Firepower Management Center NAT ID:	<input type="text"/>
Eventing Interface:	<input type="text"/>

OK Cancel

agreement since we already agreed to it.

Just hit **OK** and then press the **Save** button on the top right to start the deployment:



Now I'm going to configure the 4140-2 exactly the same way, except I'm going to use 10.11.10.209 as the IP. Still, both 4140s will be ready to be put into their FMC with IP 10.11.10.205. We'll move on to configure and managed them in Chapter 6.

See? That wasn't so bad! You really need to get your hands on some 4100s so you can practice with the FXOS and Chassis Manager. Until you can do that, just follow along as I take these devices into the FMC and finish configuring them in the network.

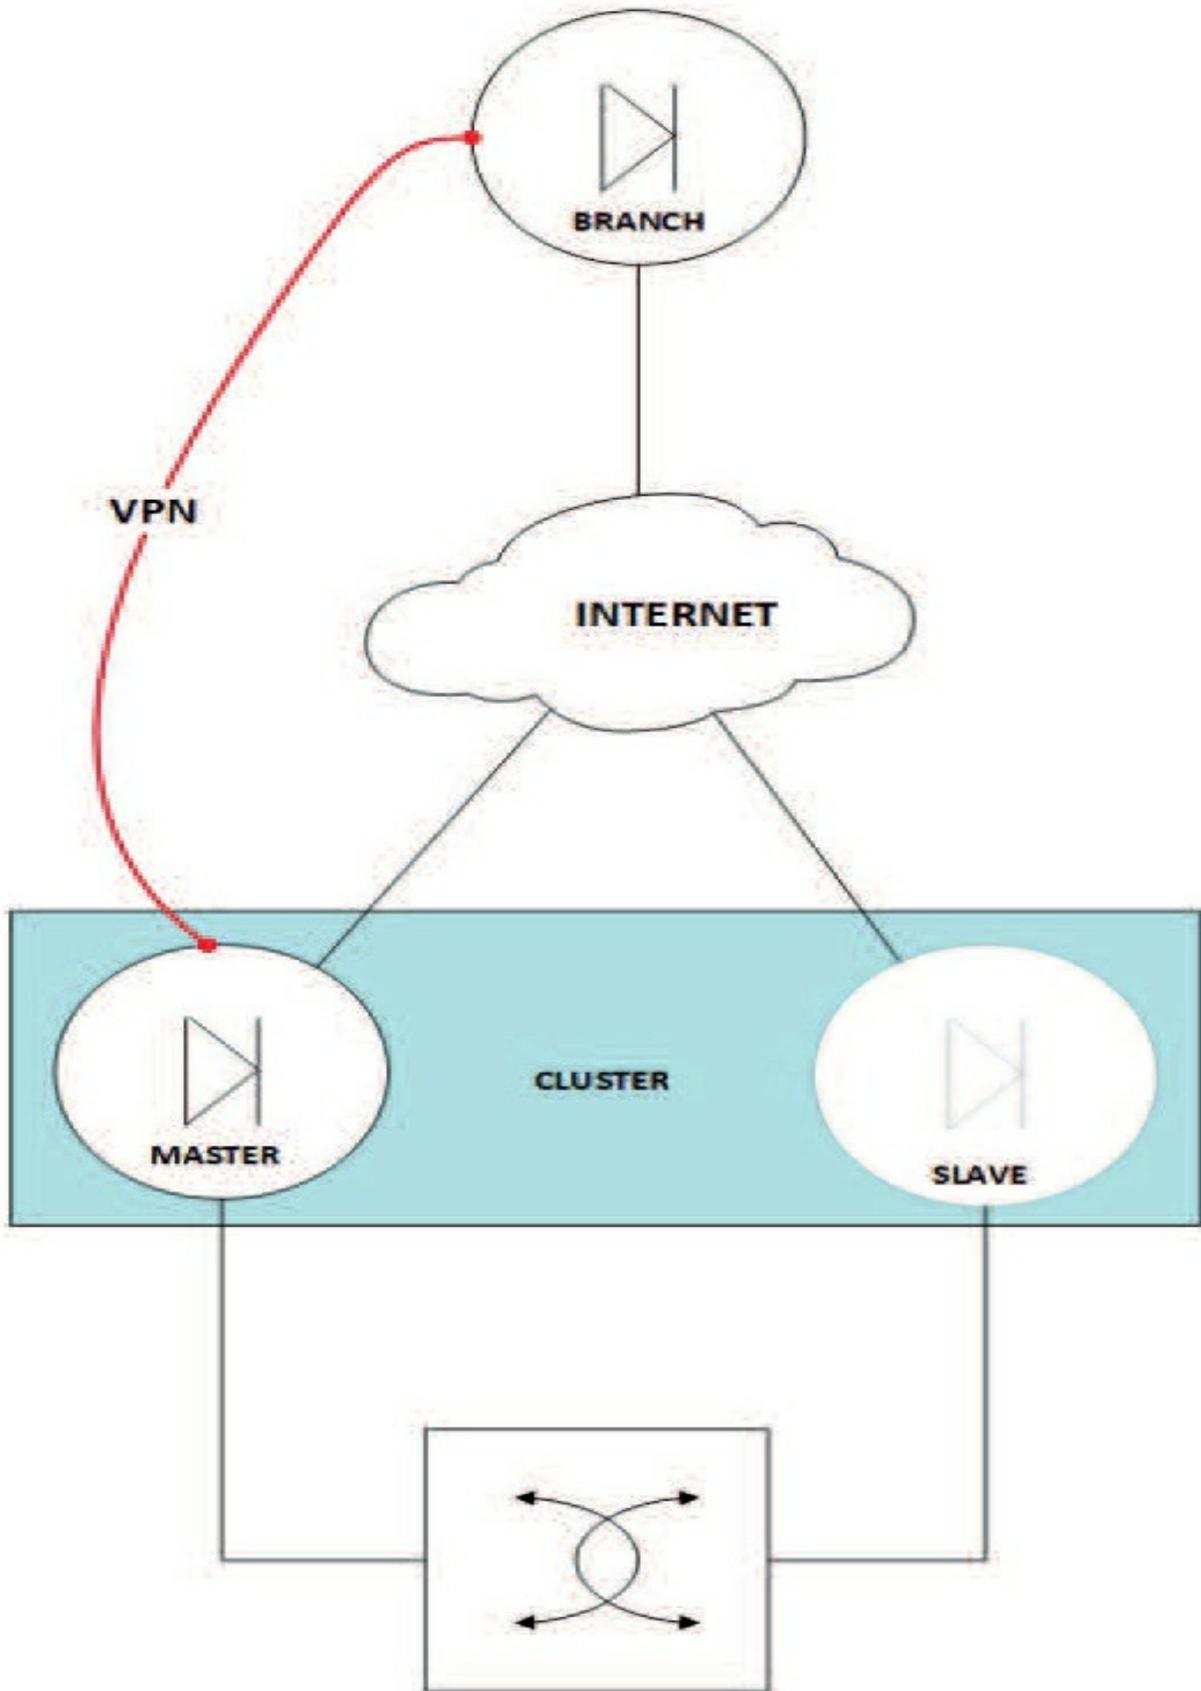
Oh and by the way, I'm just going to skip over the CLI verification here because it's the same as in the cluster example.

Cluster Traffic Flow

While the Firepower Chassis cluster runs in **active/activem**ode for most simple traffic, there are a bunch of rules that control how what node can process the workload.

Before we dive into this, we first must talk about the cluster election process. When you create a Firepower cluster, a master is elected and all other firewalls in the cluster are slaves.

The election is just like the OSPF Designated Router election, when the cluster comes up the devices will listen for messages from an existing master on the Cluster Control Link (CCL). If nothing is heard, then the device with highest priority is chosen, if all else fails then the device with the highest serial number is picked.



Also, just like OSPF, there is no preemption, once a master is chosen it will stay master as long as it is up even if a higher priority device is added to the cluster deployment. Though you can manually change the cluster master if you want but do it in a change window since it will drop all connections until the new master is up!

Centralized Features

In a Firepower cluster there is a concept called a centralized feature, this is a fancy way of saying that only the firewall running in the master role can run the service, if the master goes down then all the service connections will be dropped until a new master is running!

If a slave receives traffic destined for a centralized feature, then it will forward the traffic to the master through the CCL. Clustering does support a rebalancing feature that might send traffic to slaves before they are fully classified, once the slaves realizes the traffic is a centralized feature send it on its way to the master.

The following is the list of centralized features:

- Some Application Inspections
- Dynamic Routing
- Static Route monitoring
- Site to Site VPNs

For the most part, the master-only application inspections are just legacy protocols that you probably won't be sending through your firewall all that much. Don't use RSH anymore please.

The affected application inspections are:

- DCERPC ▪ NetBIOS
- RSH ▪ TFTP
- SUNRPC ▪ XDMCP

Let's go through all the cluster traffic scenarios that you'll need to know about if you are using a cluster, and of course what you might see in the exam!

Dynamic Routing in a Cluster

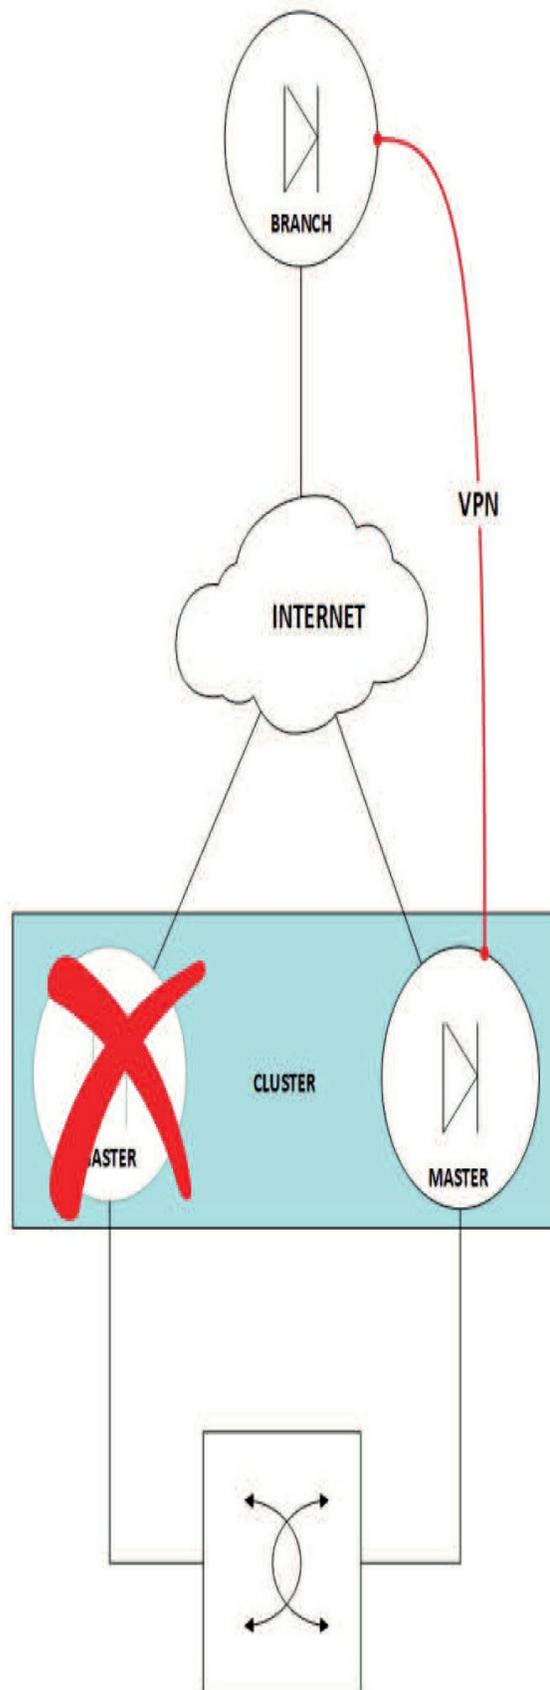
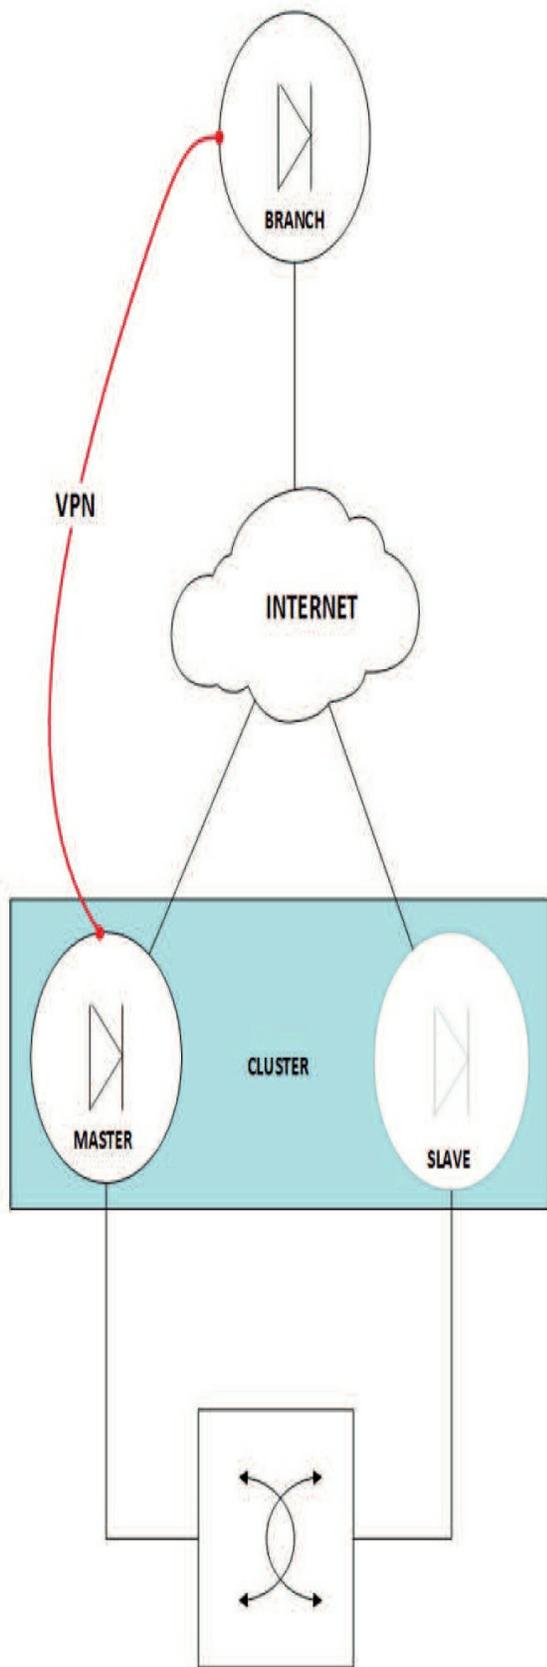
The routing process only runs on the master unit; however, the learned routes are replicated to the slaves so they will have full reachability and they can make their own forwarding decisions... though they should arrive at the same conclusion. As far as the neighboring routers are concerned, the slave firewalls are invisible.

The routing protocol database isn't replicated though, so if you have to switch masters the neighboring routers will detect the restart and kill your neighbor relationship, the master router ID will also change, unless you have statically set the router ID in your configuration. You can configure Non-Stop Forwarding to help prevent any interruptions during a master switchover.

If you choose to use the static route monitoring feature to track route reachability with SLA monitors, then that traffic will be done only on the master.

Site to Site VPNs in a Cluster

Site to Site VPNs can only be setup on the master. In the event the master goes down, the VPN will be moved to the newly elected master as shown in the next two graphics.



NAT in a Cluster

NAT runs in **active/active** mode but the catch is that it can lead to performance issues because the returning NAT traffic has to be sent to the firewall that is the NAT session owner, if it lands on another firewall it has to be sent to the proper device through the CCL.

Because of this, you should try to avoid doing NAT on the cluster and instead do it on an upstream device.

There are a couple other things to keep in mind for NAT in a Cluster:

- If you are using NAT pools, then the cluster evenly distributes the NAT pool IPs between all the cluster devices. If the cluster member runs out of IPs in the NAT pool, then the packet is dropped instead of being sent to another member.
- Dynamic NAT xlates are managed by the master and replicated to the slaves. If the slave needs to make a new dynamic xlate, it requests one from the master and then the slave owns the NAT connection.

Lastly, you can't do Static PAT for the following application inspections:

- FTP ▪ TFTP
- RSH ▪ XDMCP
- SQLNET ▪ SIP

SIP in a Cluster

SIP can run in **active/active** so any member can handle traffic, but all child data must flow through the same device. This isn't a big deal since you don't want to randomly load balance SIP sessions or the user's next call will be to help desk!

Syslog in a Cluster

Each Firepower device generates and sends its own syslog messages. Though

you can choose to configure syslog, so all devices use the same device ID, so your Network Monitoring System thinks they all come from the same device.

SNMP in a Cluster

SNMP is also done on a per device basis. You should monitor each device through its diagnostic interface instead of by the main cluster IP. The reason why is that if a new master is elected, your polling will fail.

There is no way to get consolidated cluster information from SNMP.

FTP in a Cluster

FTP is fairly straightforward; it can run in **active/active** and the FTP data channel will periodically update idle timeout for the FTP control channel owner. However, if the control channel owner reloads and gets moved to another device...then flow relationship will break, and the timeout will no longer be updated.

Trustsec in a Cluster

Trustsec is a very complicated feature, fortunately it is nice and simple in a cluster! The master learns all the Security Group Tag (SGT) information and replicates the info to all the slaves.

Unsupported Features

In addition to the various limitations we just talked about, some features just flat out don't work in a cluster. These can be a showstopper if you are thinking about using a cluster and may drive you to use high availability instead.

- Remote Access VPN (Both SSL and IPSEC)
- Any DHCP (client, server, or proxy) aside from DHCP relay
- High Availability
- Integrated Routing and Bridging.

Summary

This was an advanced chapter covering the Cisco Firepower Threat Defense (FTD) 4100 and 9300 devices.

You were given a tour of the hardware involved and an overview of the architecture used with FXOS before jumping straight into the boxes to reset the 4100s used in this book back to factory default. After that, we configured the hardware using Chassis Manager and then dove into the FXOS architecture and different contexts available from the CLI.

Next, I showed you how to add images to the Chassis Manager and perform upgrades of FXOS, including an image for Radware, a great tool that helps us fend off denial of service (DoS) attacks.

After logging into the Firepower Chassis Manager (FCM), I went through the platform settings and configured the FXOS for ASA, FTD, and Radware. Then we had some fun by clustering two 4100 chassis together! Then, before ending the chapter, we discussed some of the Cluster gotchas that you can encounter.

Our 4100s are now ready to go into an FMC, which is exactly what we'll be doing in Chapter 6!

Chapter 6: Firepower Devices

The following CCIE/CCNP Security SNCF exam objectives are covered in this chapter:

1.0 Deployment

1.1 Implement NGFW modes

1.1.a Routed mode

1.1.b Transparent mode

1.2 Implement NGIPS modes

1.2.a Passive

1.2.b Inline

1.4 Describe IRB configurations

2.0 Configuration

2.5 Configure devices using Firepower Management Center

2.5.a Device Management

We're going to spend a lot of time on Firepower Threat Defense (FTD) in this chapter. We'll also start on the configuration of the LINA process that the Snort process runs on top of. Once we get our routers and devices configured, we'll be all set to begin configuring the Snort policies in upcoming chapters.

I'm still going to guide you on a little tour of the old 7000/8000 appliances and how they can be configured because the exam objectives still cover these Firepower devices. I'll also configure and bring a Firepower appliance into the vFMC.

We're going to be configuring a pair of Firepower 1010s and a pair of 1150s as well as continuing with the 4140s I configured with Chassis Manager in Chapter 5. I'm even going to add two virtual FTDs to top it off!

Once they are configured, we'll add these nine FTD devices into the FMCs we configured back in Chapter 1—the vFMC and the hardware 2500s. Once through this wild ride of a chapter, you'll be all set to begin configuring the Snort process in the next one!

Firepower Threat Defense (FTD) on the 1000/1100/2100 and 4100/9300 Devices

The Firepower 1000/1100/2100/4100/9300 devices are the first of a new breed of network security appliances from Cisco that have replaced the legacy ASA models.

It's good to know is that these new devices can run the legacy ASA operating system or Firepower Threat Defense (FTD). So just because FTD is the most talked-about system installed right now, the ASA code isn't going anywhere anytime soon!

The new devices are different than any previous hardware you've installed before. They are not just new; they run a new operating system called the Firepower eXtensible Operating System, or FXOS. This operating system forms the foundation on which the ASA or FTD code will run.

Even though the 1000/1100/2100 run the FXOS like their big brothers do, they don't support the Firepower Chassis Manager I demonstrated for you in Chapter 5. The FXOS layer is effectively hidden from view and is only accessible from the serial console. No worries though—there should be little or no reason to perform any management tasks in FXOS on these models.

Okay—time to configure our devices!

Configuration for 1000/1100/2100

As I've done before, I'm going to start by erasing the configuration on my devices, and then configuring each one anew. By the way, if you have a brand-new box, just skip the first section here. I have two 1150s and two 1010s for this section that I'm going bring into my hardware FMC.

Resetting to Factory Default from ROMMON

Just in case, we all need to know how to recover a device if we don't know the password on the 1000/1100 and 2100 by setting the devices to factory default, so I'll show you that first.

Back in the day, there was a way to reset just the password in older codes, but that `password_reset` command at rommon is no longer valid in the codes I'm using in this book (6.5 and 6.6). This changed in the ROMMON versions as shown here:

For ROMMON version 1.0.06 or later:

```
rommon 2 > factory-reset
```

For ROMMON version 1.0.04:

```
rommon 2 > password_reset
```

And here are the steps to completely reset your FTD 1000/1100 and 2100 devices:

1. Make sure you have a working console connected to the device.
2. Power-cycle the device.
3. Press ESC when you see the prompt to interrupt boot—you have 10 seconds to do this.

4. At the `rommon 1>` prompt, type `factory_reset`:

```
rommon 1 > factory_reset
```

Warning: All configuration will be permanently lost with this operation and application will be initialized to default configuration.

This operation cannot be undone after booting the application image.

Are you sure you would like to continue ? yes/no [no]: **yes**

Please type 'ERASE' to confirm the operation or any other value to cancel: **ERASE**

Performing factory reset...

Warning: filesystem requires journal recovery

File size is 0x0000003b

Located .boot_string

Image size 59 inode num 15, bks cnt 1 blk size 8*512

Rommon will continue to boot disk0:installables/switch/ fxos-k8-fp2k-lfbff.2.6.1.133.SPA

Are you sure you would like to continue ? yes/no [no]: **yes**

5. The device will now reboot and reinstall the original code that it shipped with, which will take at least 20 minutes or so... ..Okay, maybe 30 minutes or more.

6. Once the device is finally done booting, the username and password will be admin/Admin123:

Setting the Devices to Factory Default from Firepower

If you know your username and password so you can log in, connect directly to the device with a console. I'll set the configuration on my two 1010s and two 1150s back to factory default and then configure them from scratch. The steps are exactly the same on the 1000/1100/2100. This approach won't reinstall the OS as shown in the preceding section, so if you know your username/password, so much the better for you!

1. Log in with your known username and password.

2. Type `connect ?` to view the options, then connect to the local management configuration and erase the configuration, like this:

```
1010 # connect ?
```

```
ftd Connect to FTD Application CLI
```

```
local-mgmt Connect to Local Management CLI
```

```
1010 # connect local-mgmt
```

```
1010(local-mgmt)# erase
```

configuration System configuration

1010(local-mgmt) # **erase configuration**

All configurations will be erased and system will reboot. Are you sure? (yes/no): **yes**

Removing all the configuration. Please wait....

Initial Configuration for 1000/1100/2100

Once the devices have been erased and they're finished booting, or if you're bringing up a brand-new box, you'll arrive at the `firepower#` login prompt. Enter `admin/Admin123` to log in for the first time:

```
Cisco FPR Series Security Appliance
firepower login: admin
Password: (enter Admin123 here)
Successful login attempts for user 'admin' : 1
[output cut]
```

Typing a question mark (?) will show you all the commands you can run here:

```
firepower # ?
acknowledge Acknowledge
backup Backup
clear Clear managed objects
commit-buffer Commit transaction buffer
connect Connect to Another CLI
discard-buffer Discard transaction buffer
end Go to exec mode
exit Exit from command interpreter
scope Changes the current mode
set Set property values
show Show system information
terminal Set terminal line parameters
top Go to the top mode
up Go up one mode
where Show information about the current mode
```

You want to go with the `connect` command here:

```
firepower # connect
ftd Connect to FTD Application CLI
local-mgmt Connect to Local Management CLI
```

You can use the `local-mgmt` like I did to reset the password or go right into the

FTD. I'm going to configure the FTD image now, but first I'll have to accept the EULA and set my new password:

```
firepower # connect ftd
```

```
You must accept the EULA to continue.
```

```
Press <ENTER> to display the EULA:
```

```
[output cut]
```

```
Please enter 'YES' or press <ENTER> to AGREE to the EULA: yes
```

```
System initialization in progress. Please stand by. You must change the password for 'admin' to continue. Enter new password:
```

```
Confirm new password:
```

Next up is to set the IP information. I'll choose the default for configuring IPv4, press Enter to choose `no` on IPv6, and then fill in the rest of the information:

```
You must configure the network to continue.
```

```
You must configure at least one of IPv4 or IPv6.
```

```
Do you want to configure IPv4? (y/n) [y]:y
```

```
Do you want to configure IPv6? (y/n) [n]:
```

```
Configure IPv4 via DHCP or manually? (dhcp/manual) [manual]: Enter an IPv4 address for the management interface [192.168.45.45]:172.16.10.10
```

```
Enter an IPv4 netmask for the management interface [255.255.255.0]: Enter the IPv4 default gateway for the management:[data-interface] 172.16.10.1
```

```
172.16.10.1
```

```
1.lammle.com
```

```
Enter a comma-separated list of DNS servers or 'none' [208.67.222.222,208.67.220.220]:
```

```
Enter a comma-separated list of search domains or 'none' []: If your networking information has changed, you will need to reconnect.
```

```
Setting DNS servers: 208.67.222.222 208.67.220.220
```

```
No domain name specified to configure.
```

```
Setting hostname as 1010-1.lammle.com
```

```
DHCP Server Disabled
```

```
Setting static IPv4: 172.16.10.10 netmask: 255.255.255.0 gateway: 172.16.10.1 on management0
```

```
Updating routing tables, please wait...
```

```
All configurations applied to the system. Took 6 Seconds. Saving a copy of running network configuration to local disk. For HTTP Proxy configuration, run 'configure network http-proxy'
```

Okay, this next part is pretty important to nail down so pay attention here! The first question you'll get is if you want to use Firepower Device Manager (FDM) or not. Since I'm using nothing but the FMC in this book, I'll say `no` to this.

The last question you'll receive is about firewall mode. The devices can run

in layer 2 switch mode, called Transparent Firewall mode, or in layer 3 Routed Firewall mode. Although I need to demonstrate the transparent mode for the objectives, routed mode is what most people go with in production with a few exceptions—like a data center with east<>west traffic.

I'll take the defaults of routed mode here so these devices will run as a layer 3 router:

```
Manage the device locally? (yes/no) [yes]: no
DHCP Server Disabled
Configure firewall mode? (routed/transparent) [routed]: Configuring firewall mode ..
```

I'll verify the configuration with the `show network` command:

```
> show network
===== [ System Information ] =====
Hostname : 1010-1.lammle.com
DNS Servers : 208.67.222.222

208.67.220.220
Management port : 8305
IPv4 Default route

Gateway : 172.16.10.1

===== [ management0 ] ===== State
Channels
Mode
MDI/MDIX
MTU
MAC Address
----- [ IPv4 ] ----- Configuration : Manual
Address : 172.16.10.10
Netmask : 255.255.255.0
Broadcast : 172.16.10.255
----- [ IPv6 ] ----- Configuration : Disabled

: Enabled
: Management & Events : Non-Autonegotiation : Auto/MDIX
: 1500
: 5C:5A:C7:B8:E7:80

===== [ Proxy Information ] ===== State : Disabled
Authentication : Disabled

>
```

Last up, I'm going to add the IP address of my FMC so the device can be

brought into the manager. This is my 2500 FMC with the IP address 172.16.10.20. (We'll get to my other virtual FMC soon.)

```
> configure manager add 172.16.10.20 cisco
```

Manager successfully configured.

Please make note of reg_key as this will be required while adding Device in FMC.

Believe it or Not, There's Yet Another Way to Completely Reset a Firepower Device

The `Secure Erase` command overwrites the device by erasing all data on the SSDs, which makes the data extremely difficult to recover! So you should only execute a secure erase when decommissioning the device.

For the Firepower 2100, the software image is not erased, so you can still boot into the ASA. For the Firepower 1000, the software image is erased, so the device will boot into ROMMON, where you'd need to download a new image if you were wanting to rebuild.

Changing the Admin Password If It's Known

If you know your admin password and just want to change it from FTD on the 4100/9300, you'd go into FXOS of the device. This next command isn't available on 1000/2100 series unless they're running in ASA mode, not FTD. First, type `exit` to return to FXOS, and then use the `scope` command to go into security as shown here:

```
> exit
ftd1010-1# scope security
ftd1010-1 /security # show local-user
User Name First Name Last name
-----
admin none
ftd1010-1 /security # enter local-user admin
ftd1010-1 /security/local-user # set password
Enter a password:
Confirm the password:
ftd1010-1 /security/local-user* # commit-buffer
```

Notice that after you change the password, an asterisk (*) appears before the #, which tells us that there are changes that need saving. So, we'll use the `commit-buffer` command here.

Again, the reason I added this info here is that if your 1000/2100 is running ASA code, you can still use this command.

Shutting down the 1000/2100

You never just want to turn off or unplug a Firepower FTD device. From the Firepower prompt (>), you can just use the `shutdown` command, which gracefully shuts down the device, like this:

```
> shutdown
```

```
This command will shutdown the system. Continue?
```

```
Please enter 'YES' or 'NO': yes
```

```
System is stopped.
```

```
It is safe to power off now.
```

```
Do you want to reboot instead? [y/N] n System is stopped.
```

```
It is safe to power off now.
```

You can do this from the FMC GUI as well.

Firepower Devices used in this book

For all the devices used in my lab, Table 6.1 shows the IP management VLAN information. The FTD devices used are 1010s, 1150s, 4140s, Firepower appliance, and vFTDs. The FMCs I'm using are two 2500s and a vFMC. Be sure to make a note of these devices so you can follow along through this book.

Since I already showed you the initial configuration of the 1010-1 FTD device, and the others are the exact same process, I'll configure them without showing process output in this chapter.

On the left side of the table are the four FTD devices I'm using with the 2500 FMCs, and their management IPs that I configured for each device.

Lab Layout Information

Device	IP Address	Device	IP Address	FMC 2500-1	172.16.10.20	vFMC
10.11.10.205	FMC 2500-2	172.16.20.21	4140-1	10.11.10.207	FTD 1010-1	172.16.10.10
4140-2	10.11.10.209					

FTD 1010-2	172.16.10.11	Firepower	10.11.10.210	appliance
------------	--------------	-----------	--------------	-----------

FTD 1150-1 172.16.10.12 vFTD-19 10.11.10.190
FTD 1150-2 172.16.10.13 vFTD-20 10.11.10.200

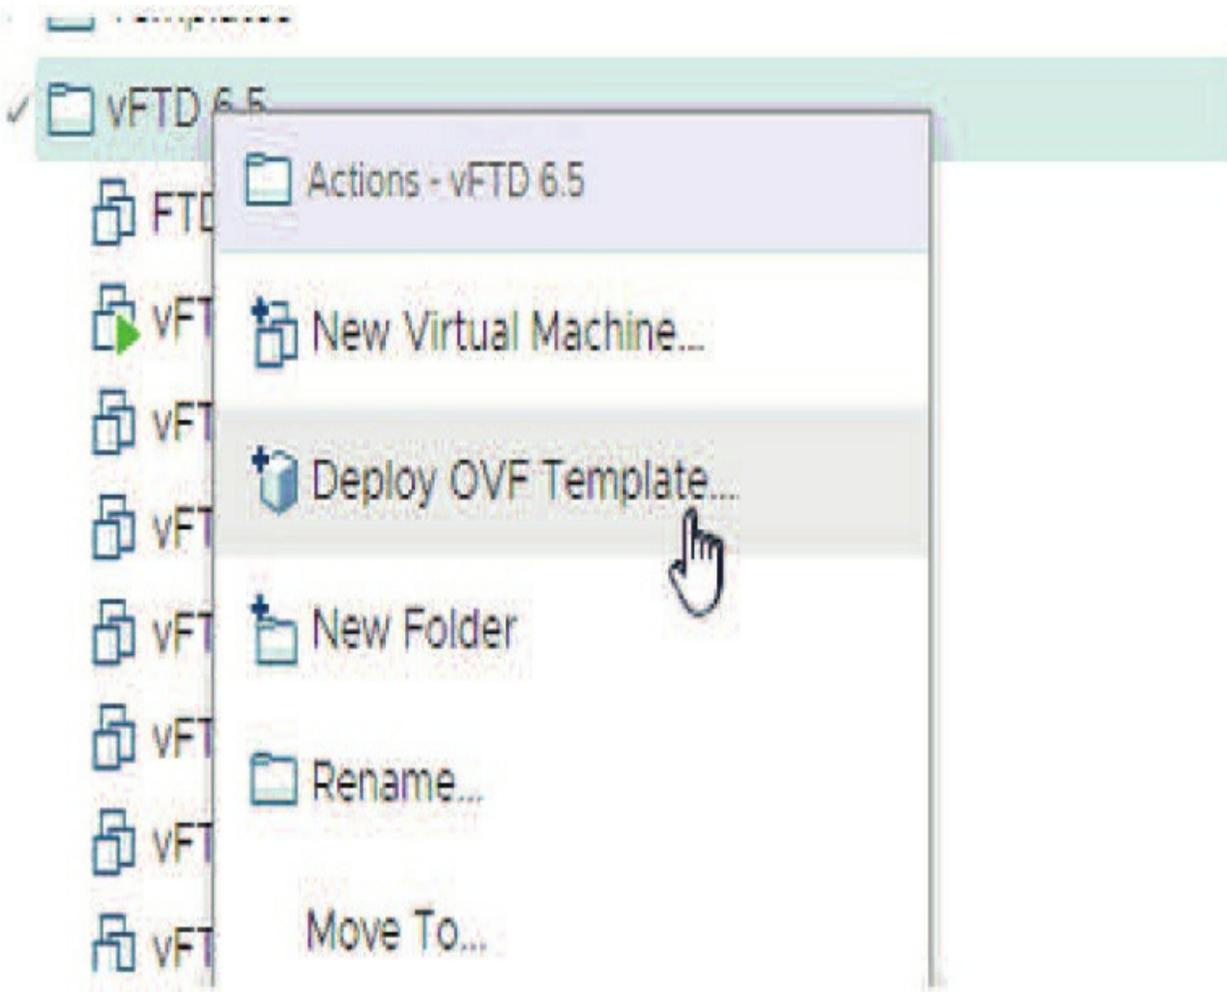
I've configured the second 1010, 1150-1, and 1150-2 with the same manager as 1010-1: the FMC 2500-1, which is 172.16.10.20. Now we'll get to bringing up a Firepower appliance.

7000/8000 Appliances

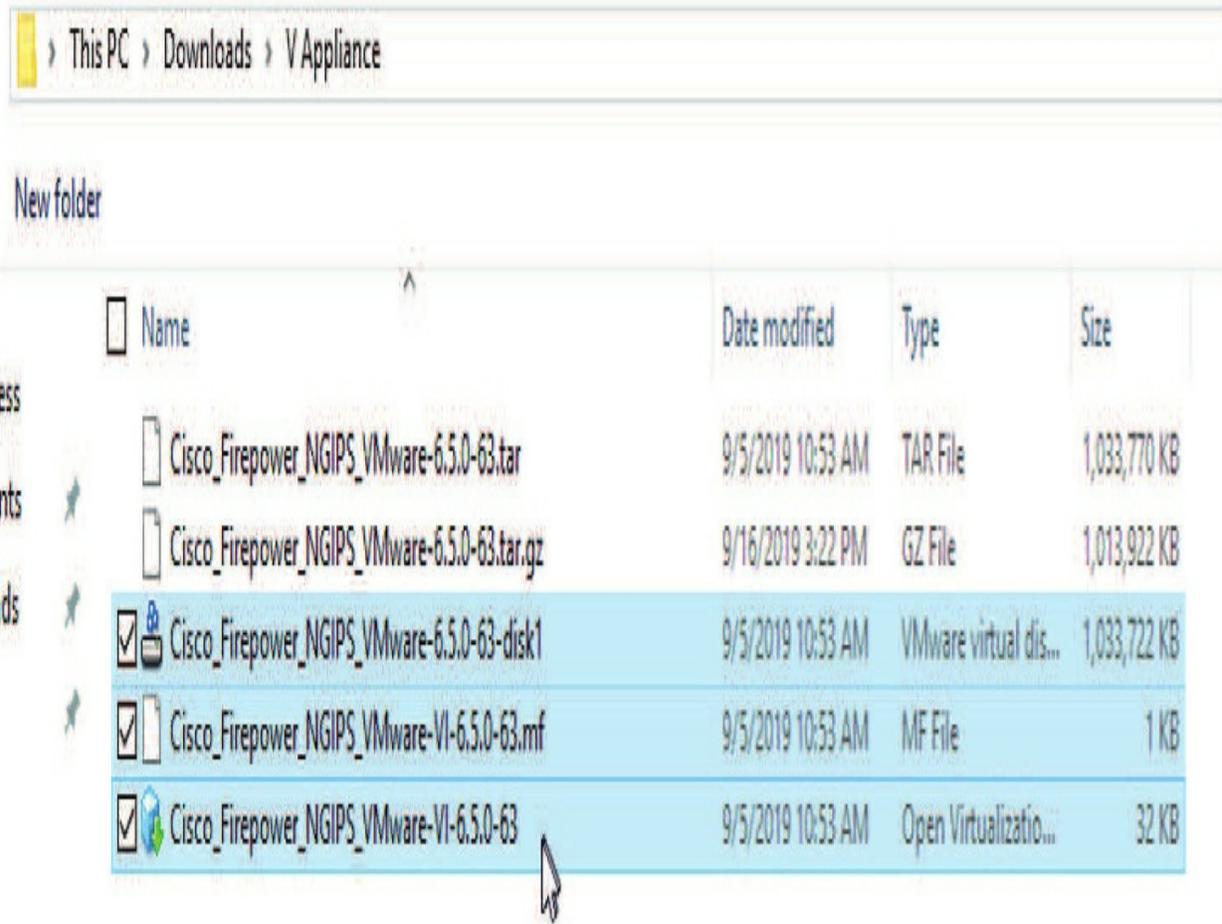
The default IP address on all Firepower FTD devices, appliances, and FMCs is 192.168.45.45, so you can either HTTPS into the device using that IP or change it via the CLI if you have your device or appliance set to the defaults.

However, I'm configuring a virtual appliance from scratch and configuring the IPs that I want during the install. You'll see how easy this is!

1. Go to your vCenter, choose your folder, and then deploy the OVF template:

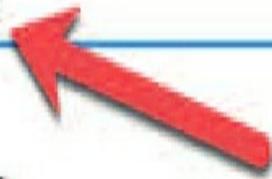


2. Traverse to the folder where your files are stored.
3. You need to expand the tar.gz and then do it again on the .tar file. For vCenter installation, choose the three files that are highlighted:



4. Choose the folder you want to run the Firepower appliance from and name the device:

Virtual machine name: FTD Appliance



Select a location for the virtual machine.



5. Choose the host you want to run the device on. When you hit Next, it will validate the storage and memory:

✓ 1 Select an OVF template

Select a compute resource

✓ 2 Select a name and folder

Select the destination compute resource for this operation

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

✓ LammleDC

> esxi01.ap.s.lab

> esxi02.ap.s.lab

> esxi03.ap.s.lab

> esxi04.ap.s.lab

6. Accept the EULA, click Next, and then choose your type of provisioning.
7. Select your networks that are configured on your vCenter and then select your disk format:

Encrypt this virtual machine (No encryption policies available)

Select virtual disk format: Thin Provision v

VM Storage Policy:

Thick Provision Lazy Zeroed
Thick Provision Eager Zeroed

Name	Capacity	Thin Provision	VM
 sas01	1.64 TB	2.64 TB	1.25 TB

8. Set the password, FQDN, and DNS servers and then scroll down to finish configuring the rest:

9. Here is where I'll configure the search domain, IP address, mask, and gateway for FTD appliance 21. After that, I'll scroll down one more time:

05. Search Domains	DNS Search Domains <u>sfgtc.local</u>
06. IPv4 Configuration	IPv4 Configuration Manual ▾
07. IP Address	IPv4 Address <u>10.11.10.210</u>
08. Netmask	IPv4 Netmask <u>255.255.255.0</u>
09. Gateway	IPv4 Gateway <u>10.11.10.1</u>
10. IPv6 Configuration	IPv6 Configuration [. . .]

10. I'll then choose the detection mode. I'll be using Inline detection mode in this book, which allows us to drop inspected packets.

▼ Detection Mode

1 settings

Detection Mode

Initial Detection Mode

Inline ▼

Passive

Inline

Network Discovery

▼ Registration

Manager

Managing Defense Center

10.11.10.215

Registration Key

Registration Key

cisco

NAT ID

NAT ID

Let's talk about this for a bit before we move on.

Passive NGIPS Modes

Think of passive detection mode as more of an IDS instead of a needed IPS, except that it won't drop any traffic.

In a passive IPS deployment, the Firepower system monitors traffic flowing across a network using a switch SPAN or mirror port. The SPAN or mirror port allows for traffic to be copied from other ports on the switch. This provides the system with visibility within the network without being in the flow of network traffic.

When configured in a passive deployment, the system can't do certain actions such as blocking or shaping traffic. Passive interfaces receive all traffic unconditionally and no traffic received on these interfaces is retransmitted.

Inline NGIPS modes

In an inline IPS deployment, you configure the Firepower system transparently on a network segment by binding two ports together. This allows the system to be installed in any network environment without the configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of an inline set unless explicitly dropped.

Now, back to configuring out virtual appliance:

11. Okay—once I've selected Inline shown in the figure, I'll set the FMC address that will manage this appliance and add the registration key. This key will be used in the configuration when the appliance is brought into the FMC:

12. Now I'm going to click Next and then let the appliance install for 20 minutes or so. Later in this chapter, I'll bring the appliance into the FMC so I can finish configuring the device and detection mode of the appliance.

13. Notice that I set the FMC address that the Firepower appliance will connect into as 10.11.10.215 with a registration key of cisco. Well, that's wrong, but since I've already saved and powered up the devices, I'm just going to redo that quickly from the CLI:

NOTE: After I installed the device, I noticed that I set the manager with the wrong IP, so I just went to the console and typed the following:

```
> configure manager delete  
> configure manager add 10.11.10.205 cisco
```

Okay, now the Firepower appliance is ready to go into its FMC, but before we do that, let's create a couple of virtual FTD devices.

Virtual FTD on vCenter

This is pretty much the same as the appliance installation above with some small but important differences. Let's take a look: 1. Go to your vCenter, choose your folder, and then deploy the OVF template:

2. Travers to the folder where your files are stored. You need to expand the tar.gz and then do it again on the .tar file. For the virtual FMC vCenter installation, choose the three files that are highlighted:

<input checked="" type="checkbox"/>	 Cisco_Firepower_Threat_Defense_Virtual-VI-6.5.0-115	9/25/2019 7:15 PM	Open Virtualizatio...	36 KB
<input checked="" type="checkbox"/>	 Cisco_Firepower_Threat_Defense_Virtual-VI-6.5.0-115.mf	9/25/2019 7:15 PM	MF File	1 KB
<input checked="" type="checkbox"/>	 Cisco_Firepower_Threat_Defense_Virtual-6.5.0-115	9/25/2019 7:15 PM	VMware virtual dis...	1,099,516 KB
	 Cisco_Firepower_Threat_Defense_Virtual-6.5.0-115.tar	9/28/2019 10:34 AM	TAR File	1,099,570 KB
	 Cisco_Firepower_Mgmt_Center_Virtual_VMware-VI-6.5.0-115	9/25/2019 7:51 PM	Open Virtualizatio...	30 KB
	 Cisco_Firepower_Mgmt_Center_Virtual_VMware-VI-6.5.0-115.mf	9/25/2019 7:51 PM	MF File	1 KB
	 Cisco_Firepower_Mgmt_Center_Virtual_VMware-6.5.0-115-disk1	9/25/2019 7:51 PM	VMware virtual dis...	2,295,899 KB
	 Cisco_Firepower_Mgmt_Center_Virtual_VMware-6.5.0-115.tar	9/28/2019 10:34 AM	TAR File	2,295,950 KB

3. Choose the folder you want to run the FTD device from, and name the device:

Virtual machine name: vFTD 20

Select a location for the virtual machine.

- ▼  vcsa.aps.lab
 - ▼  LammleDC
 - >  FMC
 - >  IN
 - >  ISE
 - >  MGMT
 - >  Templates
 - >  vFTD 6.5



4. Choose the host you want to run the device on. When you click Next, the storage and memory will be validated:
5. Accept the EULA and press Next.
6. You now get to choose how many cores and how much RAM you'll use. Choose your option and click Next:

Select a deployment configuration

Encrypt this virtual machine (No encryption policies available)

Select virtual disk format:

Thin Provision

VM Storage Policy:

Thick Provision Lazy Zeroed
Thick Provision Eager Zeroed

Name	Capacity	Thin Provision	Thick Provision Lazy Zeroed	Thick Provision Eager Zeroed	Type
 sas01	1.64 TB	2.64 TB	1.25 TB		Virtual Disk

7. Choose your disk provision type:

8. Now choose your network interfaces. This can be a bit tricky. For example, the ports listed as G0-4 and G0-5 are my interfaces g0/1 and g0/2 on the device.

These will be my inside and outside interfaces.

Each device may be different, so you have to keep trying this until you get it right:

Source Network	Destination Network
GigabitEthernet0-0	108 - Pod 8 IN
Management0-0	510 - MGMT
GigabitEthernet0-1	108 - Pod 8 IN
GigabitEthernet0-2	108 - Pod 8 IN
GigabitEthernet0-3	108 - Pod 8 IN
GigabitEthernet0-4	120 - Pod 20 IN
GigabitEthernet0-5	512 - OUT
Diagnostic	510 - MGMT
GigabitEthernet0-6	108 - Pod 8 IN
GigabitEthernet0-7	108 - Pod 8 IN

10 items

Notice above that I set the management VLAN on two interfaces. You need to do this as well.

9. Here, I'm going to set the password and DNS servers, but

1. Password 1 settings

Password

admin password

Password

Confirm

Password

2. Network 13 settings

01. Hostname

Fully Qualified Domain Name

ftd20.sfgtc.local

02. DNS1

Primary DNS Server

10.11.11.250

03. DNS2

Secondary DNS Server

8.8.8.8

04. DNS3

Tertiary DNS Server

CANCEL

BACK

NEXT

I won't configure anything else because this way, I can show you how to get this done via the CLI instead.

At this point I just clicked Next and then powered on the device, so I haven't finished configuring the device yet.

Before we bring all the devices into the FMCs, let's take a look at the CLI of the FTD devices and finish configuring the vFTD 20.

CLI of the FTD Devices

I also want you to check out the CLI of the various FTD devices that I have configured in this chapter so far. Just remember that the 4100/9300s were covered in Chapter 5 already, so they're all set and ready to go into their FMC.

I'll start by fixing the virtual FTD 20 from the CLI because I've only set the password and DNS server on this device so far. After I log in, the `show network` command provides the defaults I talked about. The greater than sign (`>`) lets you know you're in FTD:

```
> show network
===== [ System Information ] =====
Hostname : ftd.sfgtc.local
DNS Servers : 10.11.11.250

8.8.8.8
Management port : 8305
IPv4 Default route

Gateway : 192.168.45.1

===== [ management0 ] ===== State
Channels
Mode
MDI/MDIX
MTU
MAC Address : Enabled
: Management & Events : Non-Autonegotiation : Auto/MDIX
: 1500
: 00:50:56:A7:8C:34

----- [ IPv4 ] ----- Configuration : Manual
Address : 192.168.45.45
```

Netmask : 255.255.255.0
Broadcast : 192.168.45.255

Here is how you configure the IP information and hostname from the CLI of any type of FTD device:

> **configure network ipv4 manual 10.11.10.200 255.255.255.0 10.11.10.1** Setting IPv4 network configuration.

Network settings changed.

> **configure network hostname ftd20.sfgtc.local**

> **show network**

=====[System Information]===== Hostname : ftd20.sfgtc.local DNS Servers : 208.67.222.222

208.67.220.220

Management port : 8305

IPv4 Default route

Gateway : 10.11.10.1

=====[management0]===== State

Channels

Mode

MDI/MDIX

MTU

MAC Address : Enabled

: Management & Events : Non-Autonegotiation : Auto/MDIX

: 1500

: 5C:5A:C7:B8:E7:80

-----[IPv4]-----

Configuration : Manual

Address : 10.11.10.200

Netmask : 255.255.255.0

Broadcast : 10.11.10.255

-----[IPv6]-----

Now I also need to set the FMC address that the device will be managed from with the `configure manager add` command: And this can be verified with the `show managers` command:

> **show managers**

Host : 10.11.10.205

Registration Key : *****

Registration : pending

RPC Status :

>

Also, the default firewall mode of the FTD devices is Routed, but because I want to make this a transparent device, I'll cover the objectives with you:

```
> show firewall
```

```
Firewall mode: Router
```

To change the firewall mode, use the `configure firewall` command:

```
> configure firewall transparent
```

```
The firewall mode cannot be changed when a manager is configured.
```

I need to delete the configured manager first:

```
> configure manager delete
```

```
Manager successfully deleted.
```

```
> configure firewall transparent
```

```
This will destroy the current interface configurations, are you sure that you want to proceed? [y/N] y
```

Now I just need to add the manager IP address back in and I'll have a layer 2 FTD device in the network:

Basic FTD CLI Commands

This short section will have you practice some basic FTD CLI commands that are important for the exam objectives, such as troubleshooting.

Log in to the console of your FTD and type the following commands. Even though they won't provide much output at this time in your configuration, they'll be useful later:

```
> show ?
```

```
>show running-config
```

```
>show route
```

```
>show logging
```

```
>show interface ip brief
```

```
>show firewall
```

```
>configure ?
```

To get to the LINA CLI, type the following, and then run through some basic commands:

```
> system support diagnostic-cli
```

```
firepower>enable
```

```
Password: (press enter)
```

```
firepower# term pager 10 (this stops the commands from rolling) firepower# show run access-list
```

```
firepower# show run snmp-server
```

```
firepower# show run int
```

```
firepower# show run int g0/1
```

```
firepower# exit
firepower>exit
```

CLI Troubleshooting commands

Cisco exams love to verify you understand how to troubleshoot a device and the Firepower exam is no different. Although we are dedicating a huge chapter 22 to troubleshooting, just get familiar with these commands for now, so be sure and run through them but just understand that I'll dig into more detail on these in chapter 22.

First, the **system support firewall-engine-debug utility has an entry for each packet being evaluated by the ACP** and is used to generate firewall debug messages on a Cisco Firepower sensor. This also shows the rule evaluation process taking place, along with why a rule is matched or not matched. Here is an example:

Back in the FTD prompt type this command:

```
> system support firewall-engine-debug
Please specify an IP protocol: tcp
Please specify a client IP address: 10.17.117.20
Please specify a client port: 23
Please specify a server IP address: 10.11.11.250
Please specify a server port: 21
Monitoring firewall engine debug messages
^C
Caught interrupt signal
Exiting.
```

Also, starting in version 6.2 and above, the **system support trace** tool can be run. It uses the same parameters but includes more details in the output. Be sure to enter 'y' when prompted with "**Enable firewall-engine-debug too?**".

```
> system support trace
Enable firewall-engine-debug too? [n]: y
[output cut for now until chapter 22]
```

Lastly, here is an ASA old faithful and favorite CLI command of mine is the **packet tracer** command.

```
> packet-tracer input inside tcp 10.17.117.20 25 10.11.11.250 80
```

```
Phase: 1
Type: CAPTURE
```

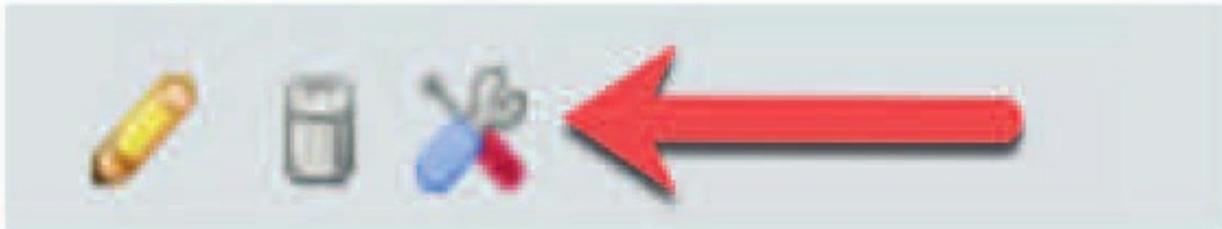
Subtype:
Result: ALLOW
Config:
Additional Information:
MAC Access list
Phase: 2
[output cut for now until chapter 22]

Download the Advanced Troubleshooting File from the GUI

Important! You need to know how to download a troubleshooting file from the CLI and the GUI, so let's look at both ways.

Again, just giving you a taste of what you need to know for the objectives, but I want to quickly show you one more thing, and that is the same commands can be found in the GUI

Go to **Devices>Device** and click on the Troubleshooting tool for the appliance/sensor/device:



You'll end up at Health Monitor, where you can click on advanced Troubleshooting

Health Monitor



Now you can find the tabs for helping you troubleshoot your devices from the GUI.'

Advanced Troubleshooting

FTD 1



Download the Advanced Troubleshooting File from the CLI

TIP: Pay Attention here

Enter the following command on Firepower devices/modules and virtual

managed devices in order to generate a troubleshoot file:

```
> system generate-troubleshoot all  
Starting /usr/local/sf/bin/sf_troubleshoot.pl...  
Please, be patient. This may take several minutes.  
The troubleshoot option code specified is ALL.  
Troubleshooting information successfully created at  
/var/common/xxxxxx.tar.gz
```

Did you notice in the last line above where the file is stored? I know I did!

Okay! Now that we've installed and provided the initial configuration of the various devices, (including the 4100s in chapter 5), and had some initial fun with the CLI, let's add the devices into their prospective managers.

Adding the 1010s and 1150s to the 2500 FMC

We have the FTD 1010s and 1150s configured and ready to go into their manager. The next table is a reminder of the IPs for these devices in the 172.16.10.0 network.

I'm going to log in to the FMC to bring these already-configured devices in now. Again, the following table shows the devices along with the corresponding management IPs for each one:

172.16.10.0 Firepower Devices

Device IP Address

FMC 2500-1	172.16.10.20
FMC 2500-1	172.16.20.21
FTD 1010-1	172.16.10.10
FTD 1010-2	172.16.10.11
FTD 1150-1	172.16.10.12
FTD 1150-2	172.16.10.13

We'll start with the 1150s... shown below is the configuration for 1150-1 required to bring it into the 2500-1 FMC.

Once I choose Add Device on the Devices page, I then apply the IP address

and the name and the registration key of `cisco`. Since this is a new FMC, I had to create a new Access Control policy (ACP).

When you add your first device, you'll always be asked to create a new ACP because you just can't add a device without an Access Control policy (ACP) attached at any time.

Add Device

? X

Host:†	<input type="text" value="172.16.10.12"/>
Display Name:	<input type="text" value="1150-1"/>
Registration Key:*	<input type="text" value="cisco"/>
Group:	<input type="text" value="None"/> ▼
Access Control Policy:*	<input type="text" value="Create new policy"/> ▼
Smart Licensing	
Malware	<input type="checkbox"/>
Threat	<input type="checkbox"/>
URL Filtering	<input type="checkbox"/>
Advanced	
Unique NAT ID:†	<input type="text"/>
Transfer Packets	<input checked="" type="checkbox"/>

 On Firepower Threat Defense devices version 6.2.1 onwards, AnyConnect VPN licenses can be enabled from [smart license page](#)

Access control policy is required.

Register

Cancel

You also need to name the ACP and then choose a default action.

The default action is the rule found at the end of your layer 7 ACP rule set, and you can see this in the figure below. If you choose Block All Traffic, it will add a deny IP any any to the end of your ACP, and Intrusion Prevention will add a Permit IP any any along with an IPS policy of Balanced Security and Connectivity.

We'll cover the IPS policies in detail later on in Chapter 12, "Intrusion Prevention System (IPS) Policy."

Notice I used Intrusion Prevention because, for now, I want a permit IP any any at the end of my list. It really doesn't matter which one you use right now, as long as you understand the difference!

Finally, I'll click to add my licenses and then click Register.

Lastly, you need to understand that the Transfer Packets check box is enabled by default. You'll usually want this enabled because it means that when an IPS event occurs, the packets of the file that triggered the event will be transferred and stored on the FMC. This is an important factor needed for Network Analysis.

Even so, you might want to disable it if your device is located across a long slow WAN link and you don't have enough bandwidth to handle all the packets.

Once the first device shows up on the Devices page and starts the registration phase, the FMC will deploy the basically empty ACP to the device:



1150-1 172.16.10.12	FTD on Firepower 1150	6.5.0	N/A	Base, Threat (2 more...)	None	
------------------------	--------------------------	-------	-----	-----------------------------	------	--

While that's happening, I can go and bring my second 1150 into the FMC. I'll add the IP, name the FTD device, add the registration key and the ACP, choose my licensing, and then click Register.

Once both 1150s are in the FMC, we can highlight over the name to get some

information.

The screenshot displays a network management interface with two rows of device information. The first row shows device 1150-1 with IP 172.16.10.12, running FTD on Firepower 1150, version 6.5.0, N/A, Base, Threat (2 more...), and FMC2500-ACP. The second row shows device 1150-2 with IP 172.16.10.13, running FTD on Firepower 1150, version 6.5.0, N/A, Base, Threat (2 more...), and FMC2500-ACP. A detailed view for device 1150-2 is open, showing 'Cisco Firepower 1150 Threat Defense (Version 6.5.0)'. Under 'Device Information', there are links for 'Context Explorer' and 'Health Dashboard'. Under 'Health Modules', there is a table:

Health Module	Count
Normal	13
Disabled	24

From the screen we can see both devices are Routers and running 6.5.0 code. Also, both have the ACP deployed:

Once I add the 1010s in, the devices' pages will provide all the information about the devices, along with their health, which shows all good here because we have a green check mark on the left. Check it out:

✓ 1010-1 172.16.10.10 - Routed	FTD on Firepower 1010	6.4.0	N/A	Base, Threat (2 more...)	FMC2500-ACP	  
✓ 1010-2 172.16.10.11 - Routed	FTD on Firepower 1010	6.4.0	N/A	Base, Threat (2 more...)	FMC2500-ACP	  
✓ 1150-1 172.16.10.12 - Routed	FTD on Firepower 1150	6.5.0	N/A	Base, Threat (2 more...)	FMC2500-ACP	  
✓ 1150-2 172.16.10.13 - Routed	FTD on Firepower 1150	6.5.0	N/A	Base, Threat (2 more...)	FMC2500-ACP	  

Now that we have our 2500s with our four FTD devices being managed and ready to be configured, let's go over to my virtual FMC to bring in the four FTD devices and one Firepower appliance into that manager.

Adding the Firepower Appliance, 4140s and vFTDs into the Virtual FMC

Now that the 1010s and 1150s are in the 2500 FMC and ready to be configured, let's add the Firepower appliance, the 4140s, and the vFTDs into the virtual FMC. In the next section, we'll configure all of the IP addresses and routing for all devices.

The devices in this section are all in the 10.11.10.0/24 management VLAN. Remember that I already installed the Smart and Classic Licenses, (Classic for the appliance), on the vFMC in Chapter 4. This must be completed before you can add devices!

The table below shows the management IPs for all the devices in the 10.11.10.0 management VLAN.

Device IP Address

vFMC 10.11.10.205

FTD 4140-1 10.11.10.207

FTD 4140-2 10.11.10.209

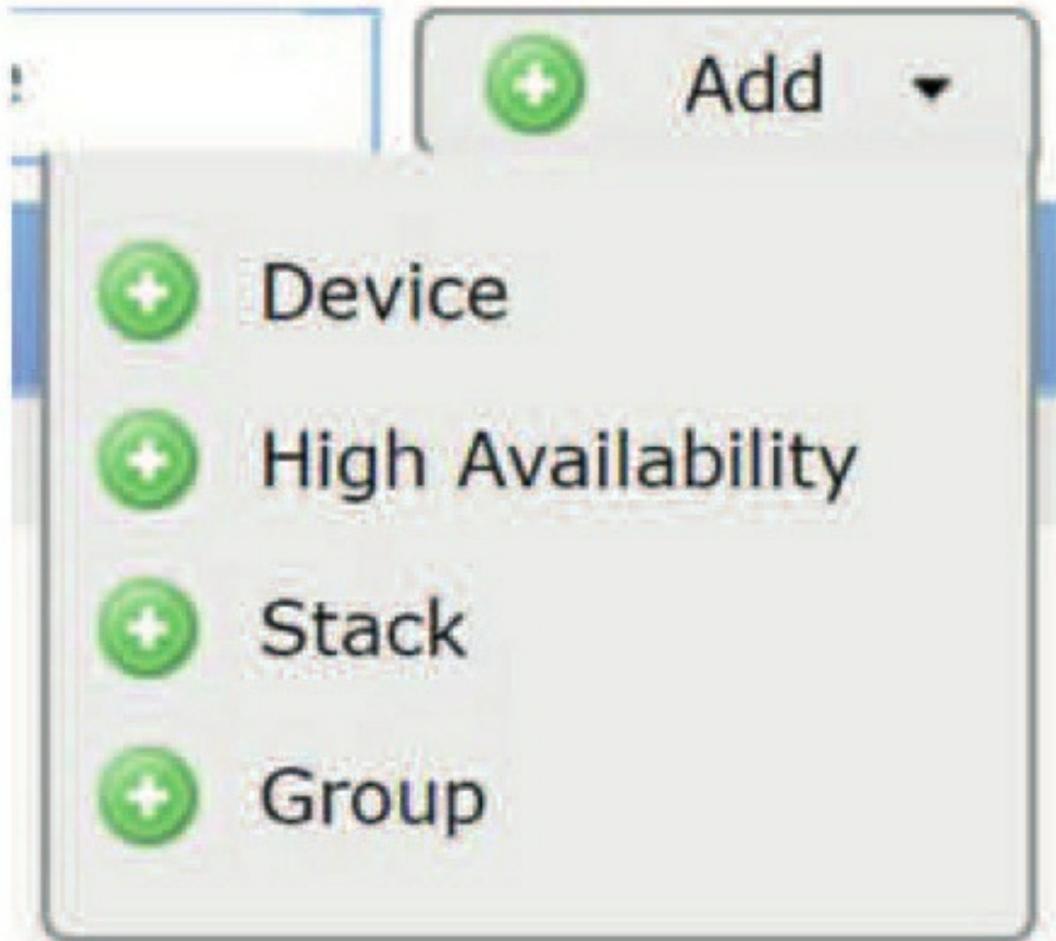
vFTD-19 10.11.10.190

vFTD-20 10.11.10.200

Firepower Appliance 10.11.10.210

Let's start by adding the Firepower appliance with 10.11.10.210 as its management IP into the vFMC of 10.11.10.205.

From the FMC, go to **Devices>Add>Device**:



Now you enter the IP address of the device you want to add into the FMC. In this example, I'm adding the Firepower appliance with IP 10.11.10.210. I named the device and then added the registration key that's the same on all devices: cisco:

Add Device

Host:†	<input type="text" value="10.11.10.210"/>
Display Name:	<input type="text" value="Firepower_Appliance"/>
Registration Key:*	<input type="text" value="cisco"/>

When you add your first device, you will be asked to create a new ACP, just like in the previous section.

Notice I used Intrusion Prevention because again, because I want a **permit IP any any** at the end of my list:

New Policy

Name:

Description:

Select Base Policy:

Default Action: Block all traffic Intrusion Prevention Network Discovery

Access Control Policy: *

Save

My final configuration is below. Notice I added the licensing to the device. The Group field allows you to create administrative folders to put your devices into. If you have dozens of devices in lots of locations, this makes administration much easier!

Add Device



Host:†	<input type="text" value="10.11.10.210"/>
Display Name:	<input type="text" value="Firepower_Appliance"/>
Registration Key:*	<input type="text" value="cisco"/>
Group:	<input type="text" value="None"/> ▼
Access Control Policy:*	<input type="text" value="VFMC20-ACP"/> ▼
Smart Licensing	
Malware	<input checked="" type="checkbox"/>
Threat	<input checked="" type="checkbox"/>
URL Filtering	<input checked="" type="checkbox"/>
Advanced	
Unique NAT ID:†	<input type="text"/>
Transfer Packets	<input checked="" type="checkbox"/>

On Firepower Threat Defense devices version 6.2.1 onwards, AnyConnect VPN licenses can be enabled from [smart license page](#)

Now that the appliance has been added into the FMC, if I go to **System>Licenses>Classic Licenses**, I can see that the Classic License is now in use. In the previous **Add Device** figure, it actually said Smart Licensing, which was not the case as you can see here:

Maximum VirtualDevice64bit Licenses

Protection (Used)	1 (1)
Control (Used)	1 (1)
URL Filtering (Used)	1 (1)
Malware (Used)	1 (1)
VPN (Used)	0 (0)

Maximum Cisco Firepower Management Center for VMWare Licenses

Firepower Host (Used)	50000 (0)
Firepower User (Used)	50000 (0)

Okay—so let's go ahead and add the other devices: two 4140s and two vFTDs. Here's the page for the 4140-1. There's no need to show you all of them being added into the FMC because they're all configured the same way:

Add Device



Host:†	<input type="text" value="10.11.10.207"/>
Display Name:	<input type="text" value="4140-1"/>
Registration Key:*	<input type="text" value="cisco"/>
Group:	<input type="text" value="None"/> ▼
Access Control Policy:*	<input type="text" value="VFMC20-ACP"/> ▼
Smart Licensing	
Malware	<input checked="" type="checkbox"/>
Threat	<input checked="" type="checkbox"/>
URL Filtering	<input checked="" type="checkbox"/>
Advanced	
Unique NAT ID:†	<input type="text"/>
Transfer Packets	<input checked="" type="checkbox"/>

On Firepower Threat Defense devices version 6.2.1 onwards, AnyConnect VPN licenses can be enabled from [smart license page](#)

Finally, I have all five devices into the FMC and showing in the Devices page. If this was a production network, I might add a folder called Groups into the page just to administratively break up the devices by type. I typically use Groups when I have a customer with hundreds of sites, and I'll create a Group for each location for easier management.

Now that all our devices are in their prospective FMCs and ready to be configured, let's add the IPs and routing for each routed device— all of them except the vFTD20 appliance:

4140-1 10.11.10.207 - Routed	FTD on Firepower 4140	6.5.0	A4140-1.sfgtc.local:443 Security Module - 1	Base, Threat (2 more...)	VFMC20-ACP	  
4140-2 10.11.10.209 - Routed	FTD on Firepower 4140	6.5.0	c4140-2.sfgtc.local:443 Security Module - 1	Base, Threat (2 more...)	VFMC20-ACP	  
Firepower_Appliance 10.11.10.210	NGIPSv for VMware	6.5.0	N/A	Protection, Control (2 more...)	VFMC20-ACP	  
vFTD-19 10.11.10.190 - Routed	FTD for VMWare	6.5.0	N/A	Base, Threat (2 more...)	VFMC20-ACP	  
vFTD20 10.11.10.200 - Transparent	FTD for VMWare	6.5.0	N/A	Base, Threat (2 more...)	VFMC20-ACP	  

Configuring the IPs on the 172.16.10.0 Devices

In this section, we'll get started by configuring the 1150s first. Go to **Devices>Device Management** to see all your devices:



From the Devices page, click the pencil on the far right to choose the device you want to configure. You can delete the device from here or go in to troubleshooting:



Clicking the pencil will open these tabs to configure your Firepower device:

1150-1

Cisco Firepower 1150 Threat Defense



The default landing page is the Interfaces page, and there's a lot to configure here! Below is the page shown under the Interface tab of the 1150-1.

1150-1



Cisco Firepower 1150 Threat Defense

Device	Routing	Interfaces	Inline Sets	DHCP	SNMP
<input type="text" value="Search by name"/> <input type="button" value="Sync Device"/> <input type="button" value="Add Interf"/>					
Interface	Logical Name	Type	Security Zones	MAC Address (Active/Standby)	IP Address
Diagnostic1/1	diagnostic	Physical			
Ethernet1/1		Physical			
Ethernet1/10		Physical			
Ethernet1/11		Physical			
Ethernet1/12		Physical			
Ethernet1/2		Physical			
Ethernet1/3		Physical			
Ethernet1/4		Physical			
Ethernet1/5		Physical			
Ethernet1/6		Physical			
Ethernet1/7		Physical			
Ethernet1/8		Physical			
Ethernet1/9		Physical			

There are 12 Ethernet ports and a management port, which is the only port enabled by default.

The 1150s are relatively new devices, and all the ports can run 10/100/1000. However, ports 1/9 and 1/11 can actually run 10 Gig, so those are the ports I'm going to use and configure. I'll click the pencil next to those ports to configure them, but before I do, I want to talk about security zones.

Security Zones

ASAs are configured as security level firewalls and the higherlevel interfaces can automatically go out lower-level interfaces. For example, inside interfaces default to 100 and outside interfaces default to 0. This means that inside hosts can go out (and back in because of stateful inspection), but a host from the lower-level interface can never initiate to the inside by default. This worked well for nearly two decades!

But Firepower is a zone-based firewall, which allows for more security and flexibility in your configuration. Think of zones as property, and to get to another property, you'd need to cross a fence by asking for permission. There are no defaults on Firepower interfaces like the ASA had—everything is considered equal, and all have to be configured to allow packets from one zone to another.

Interface objects segment your network to help you manage and classify traffic flow. An interface object simply groups interfaces. These groups may span multiple devices, and you can also configure multiple interface objects on a single device.

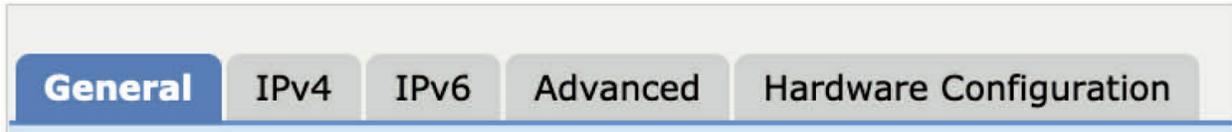
There are two types of interface objects:

- Security zones—An interface can belong to only one security zone.
- Interface groups—An interface can belong to multiple interface groups, and to one security zone.

You only use interface groups in FTD NAT policies, prefilter policies, and QoS policies.

Interfaces Tab

Opening an interface brings up even more tabs that we can use to configure this particular interface. Let's take a look at some of them now:



From the General tab, the first thing you definitely don't want to forget is the Enabled box, which is the `no shutdown` command. A lot of people do, so if you forget to enable this and deploy, you'll have to go back to enable it and completely deploy again:

From here you can name it as I did. I'm only using two interfaces on this device—an inside and outside, so the configuration is pretty simple. Notice that I created and set a zone called Inside here.

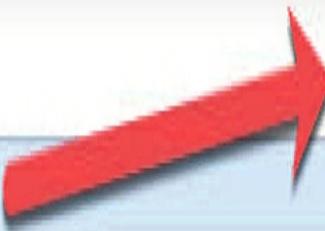
General IPv4 IPv6 Advanced Hardware Configuration

Name: Enabled

Description:

Mode: ▼

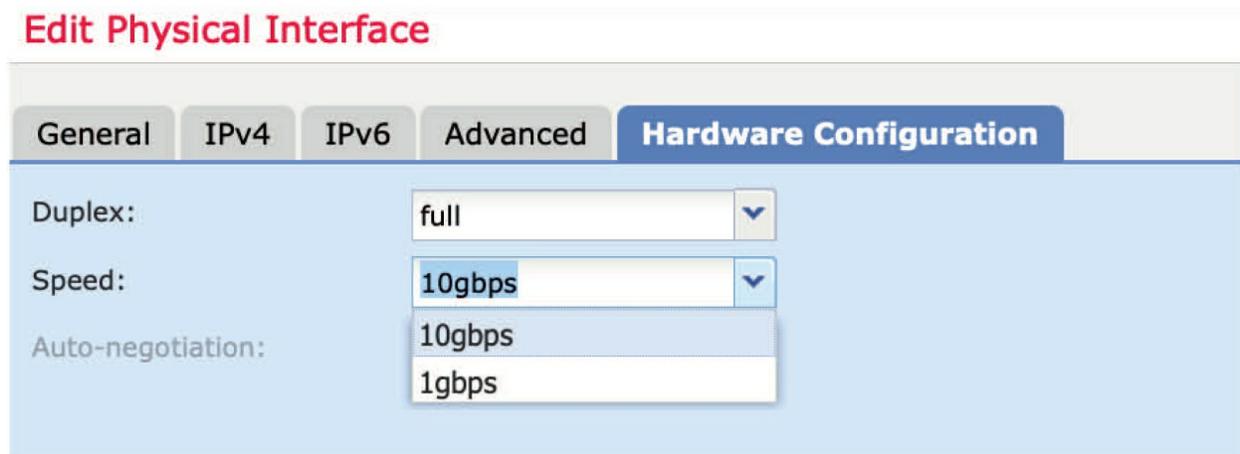
Security Zone: ▼



From the IPv4 tab, enter the IP address of the interface along with the subnet mask:

Next you can configure an IPv6 configuration from the IPv6 tab, but I'm not going to do that here. I'm also going to skip the Advanced tab until we get to the Platform policy in Chapter 19, found in the Volume II book.

The Hardware Configuration tab isn't used so much. The default on the 1/9 and 1/11 ports is 10gbps but understand that these can only be set down to 1Gbps – so good troubleshooting spot to remember! All the other ports can be set to 10/100/1000:



For the objectives, remember that the only thing that is requirement when configuring a routed interface with FTD is Speed and Duplex – that's it.

Integrated Routing and Bridging (IRB)

A lot of my clients want lots of physical interfaces configured to belong to the same VLAN. To make this happen for them, I use a cool feature called IRB because it lets users configure bridge groups when in routed mode.

IRB also allows devices to carry out L2 switching between all interfaces, and all bridge groups must have a Bridge Group Virtual Interface (BVI) that you have to configure an IP address for. FTD will use the BVI address as the source address for packets coming in from the bridge group, and keep in mind that this address has to be located on the same subnet as the connected network is.

The BVI IP address is a must to allow IPv4 traffic, but you also have to at least configure the link-local addresses for IPv6. I typically go with a global

management address because doing that lets me carry out remote management plus other key tasks.

You've also got to name the BVI for routed mode to work because your bridge group will be secluded and behave as though it were in transparent firewall mode if you don't. In addition, you're not allowed to create Access Control policies to the BVI, and you can't hook it up to a security zone either. Instead, I typically place my AC policy on bridge group member interfaces dependent on which zone they're in.

Okay—with that, here's how to simply configure a BVI: 1. Select **Devices>Device Management** and click the pencil edit icon for your FTD device. The Interfaces tab will be the default screen.

2. Choose **Add Interfaces > Bridge Group Interface**.

3. From the Interfaces tab, choose the interfaces you want included in the BVI:

4. If in Routed mode, enter a name up to 48 characters in length in the Name field.

Again, you absolutely must name the BVI if you want to route traffic outside the bridge group members—to the outside interface or to members of other bridge groups! I tried to write that last sentence more dramatically, but that's all I have.... also, the name isn't case-sensitive.

5. In the Bridge Group ID field, enter the bridge group ID between 1 and 250. Notice I just started with 1.

6. In the Description field, enter an optional description for this bridge group. Pretty simple really. Speaking of simple, now let's configure a DHCP server on my 1150 FTD device.

Configuring a DHCP Server

From 1150-1 I'll set a DHCP for my internal network. It's relatively simple. From the device management page, just click on the DHCP tab:

Device	Routing	Interfaces	Inline Sets	DHCP	SNMP
▶ DHCP Server					
DHCP Relay					
DDNS					
Ping Timeout		50			
Lease Length		3600			
Auto-Configuration		<input type="checkbox"/>			
Interface		[Dropdown]			
Override Auto Configured Settings:					
Domain Name		Lammle.com			
Primary DNS Server		[Input] 			
Secondary DNS Server		[Input] 			

Notice that I added the domain name, and I want to add my DNS servers. I can't do this with an IP address—I must create objects and add only the objects into this section.

So, I'm going to create two objects here: one for Google DNS and one for Cloud Flare out of Australia (1.1.1.1). Here's an example of creating the Google DNS object:

New Network Object

Name	Google-DNS
Description	
Network	<input checked="" type="radio"/> Host <input type="radio"/> Range <input type="radio"/> Network
	8.8.8.8

You can see the domain and two DNS objects there at the upper left. From here, scroll down and click Add Server.

Add the interface you want to use to hand out addresses and create your address pool:

Override Auto Configured Settings:

Domain Name	<input type="text" value="lammle.com"/>	Primary WINS Server	<input type="text"/>
Primary DNS Server	<input type="text" value="Aus-DNS"/>  	Secondary WINS Server	<input type="text"/>
Secondary DNS Server	<input type="text" value="Google-DNS"/>  		

Server Advanced

Interface	Address Pool	Enable DHCP Server
-----------	--------------	--------------------

No records to display

Add Server ? X

Interface*	<input type="text" value="Inside"/> 
Address Pool*	<input type="text" value="172.16.10.50-172.16.10.2"/> (2.2.2.10-2.2.2.20)
Enable DHCP Server	<input checked="" type="checkbox"/>

Last, don't forget to enable the DHCP server! Easy to miss. If your DHCP server needs to reach remote networks, go with the DHCP Relay configuration.

Configuring Routing

Although I am going to configure routing for my devices listed here, this isn't a process I'll demonstrate at this point on the 172.16.10.0 devices. I'll detail that process coming up soon in the next configuration for my 10.11.10.0 devices. There's no need to duplicate the process in this book; just pay attention when you get to that section and know that I performed the configuration here as well as there.

Verifying the Configuration from the CLI

Of course, we verify from the GUI, but it's important to understand that you can type in almost everything at the CLI of the FTD that you could from an ASA, so let's work through a few commands. Notice I'm at the FTD prompt (>).

The `show interface` command is useful, but did you catch that I used a zone name with the command? Doing this will succinctly show the `Inside` configured interface information and statistics:

```
> show interface inside
```

```
Interface Ethernet1/9 "Inside", is up, line protocol is up
```

```
Hardware is EtherSVI, BW 10000 Mbps, DLY 10 usec
```

```
MAC address 6c8b.d327.66ac, MTU 1500
```

```
IP address 172.16.10.1, subnet mask 255.255.255.0
```

```
Traffic Statistics for "Inside":
```

```
61716714 packets input, 10845422498 bytes
```

```
128505161 packets output, 172723677237 bytes
```

```
299867 packets dropped
```

```
1 minute input rate 140 pkts/sec, 10557 bytes/sec 1 minute output rate 523 pkts/sec, 718422 bytes/sec 1  
minute drop rate, 1 pkts/sec
```

```
5 minute input rate 134 pkts/sec, 12960 bytes/sec 5 minute output rate 389 pkts/sec, 520221 bytes/sec 5  
minute drop rate, 1 pkts/sec
```

Oh, and because I'm in the FTD, it's important to use detailed commands

since we can't use the terminal pager command as we did with the ASA. This is used to prevent the output from displaying more than one page at a time.

Notice that when I used the show run command, the command must be the full command! I then added the detailed information to check out interface E1/9:

```
> show running-config interface Ethernet 1/9
```

```
!  
interface Ethernet1/9  
  
nameif Inside  
cts manual  
propagate sgt preserve-untag  
policy static sgt disabled trusted  
security-level 0  
ip address 172.16.10.1 255.255.255.0
```

Unlike with the IOS using the show ip interface brief command, the ASA and FTD use **show interfaces ip brief** to provide interface information:

```
> show interface ip brief
```

```
Interface IP-Address OK? Method Status Protocol Internal-Data0/0 unassigned YES unset up up
```

```
Ethernet1/1  
[output cut]
```

```
Ethernet1/9 Ethernet1/10 Ethernet1/11 Ethernet1/12 Internal-Control1/1 unassigned YES unset up up  
Internal-Data1/1 169.254.1.1 YES unset up up Internal-Data1/2 unassigned YES unset up up  
Management1/1 unassigned YES unset up up >
```

And last up, here's the **show route** command output displaying the routing table:

```
> show route
```

```
unassigned YES unset admin down down
```

```
172.16.10.1 YES manual up up unassigned YES unset admin down down  
192.168.227.2 YES manual up up unassigned YES unset admin down down
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, [output cut]
```

```
Gateway of last resort is 192.168.227.1 to network 0.0.0.0 L 172.16.10.1 255.255.255.255 is directly  
connected, Inside C 192.168.227.0 255.255.255.0 is directly connected, Outside L 192.168.227.2  
255.255.255.255 is directly connected, Outside
```

```
S* 0.0.0.0 0.0.0.0 [1/0] via 192.168.227.1, Outside C 172.16.10.0 255.255.255.0 is directly connected,  
Inside
```

```
>
```

Now that we have the IPs, zones, and routing configured on the devices in the 172.16.10.0 network, let's add the 10.11.10.0 devices into their FMC and do the same.

Configuring the IPs on the 10.11.10.0 Managed Devices

Look back to Table 6.3 for a handy reminder of all the devices I have in the vFMC. Counting the devices in the 2500 FMC plus this virtual FMC sums up as an impressive number of devices for this book for sure! You just won't find anything like this in another book.

Next, I'm going to show you the configuration of the Firepower appliance since it's unlike the other devices, and then I'll demonstrate the configuration for the 4140-1. I'm going to go ahead and configure the others without showing the output because they are all the same configuration—with different IPs of course!

When I click on the pencil for the Firepower appliance, I'll be taken to a screen with only three tabs. There really isn't much to configure here because I already configured the box to Inline mode when I installed the appliance:

You can configure your device in either a passive or an inline IPS deployment. In a passive deployment, the system is positioned out of band from the flow of network traffic so it can't stop or drop bad traffic.

Firepower_Appliance
NGIPSv for VMware

Device		Interfaces	Inline Sets
Name	Security Zones	Used By	MAC Address
 eth1	Internal	Default Inline Set	00:50:56:a7:8b:d7
 eth2	External	Default Inline Set	00:50:56:a7:f9:6b

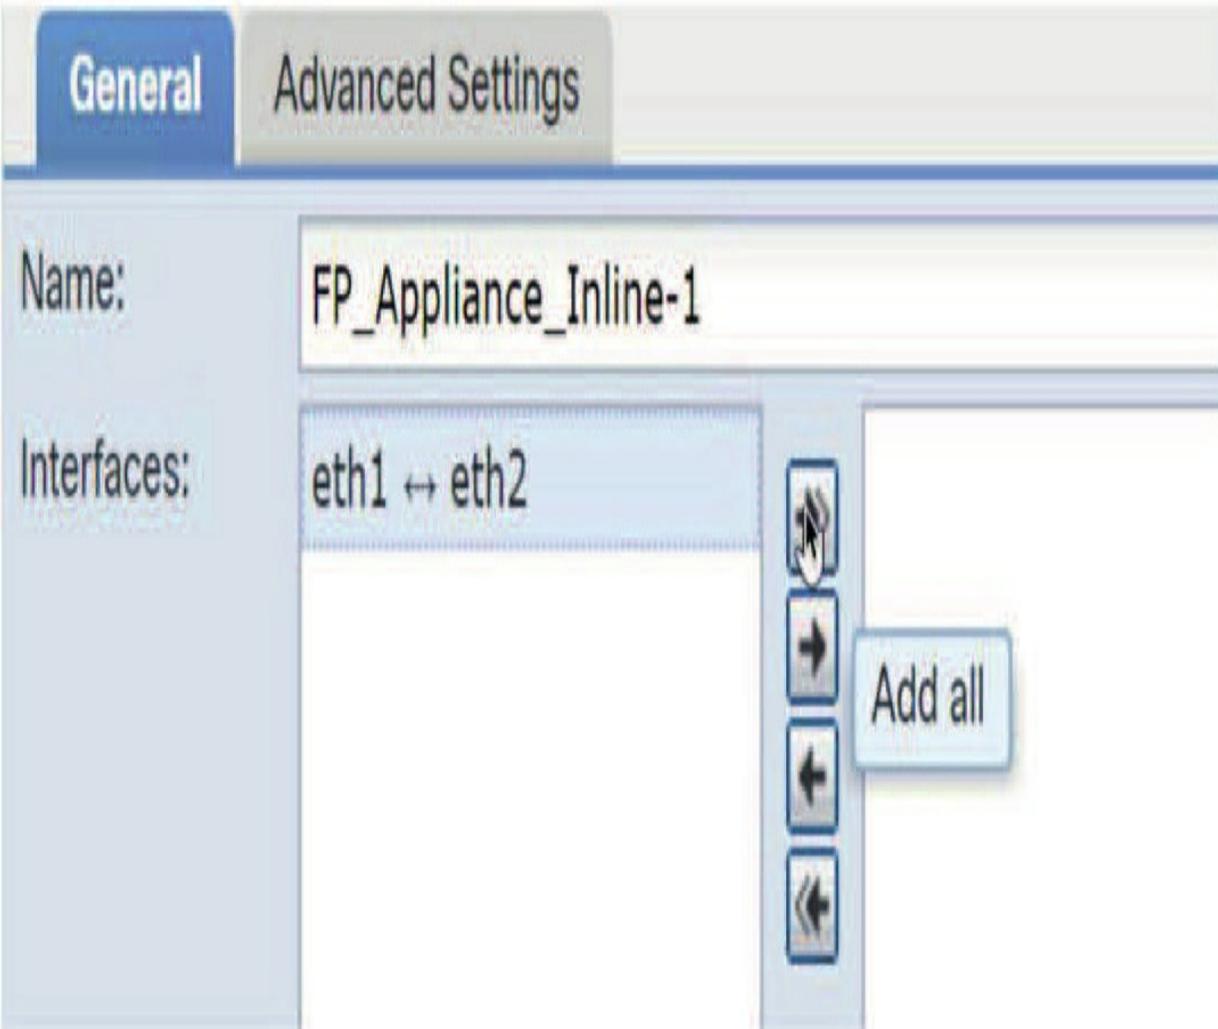
But in an inline deployment like I have here, you configure the system transparently on a network segment by binding two ports together in an inline pair. In the screen below, the device I've installed can only run in Transparent Firewall mode:

Edit Interface



The dialog box features a header with three tabs: 'None', 'Passive', and 'Inline'. The 'Inline' tab is currently selected. Below the tabs, there are three configuration fields: 'Security Zone' is a dropdown menu set to 'Internal'; 'Inline Set' is a dropdown menu set to 'Default Inline Set'; and 'Enabled' is a checked checkbox. At the bottom right, there are 'Save' and 'Cancel' buttons.

If you need to, make any changes here and then press **Save**. Choose Add Inline Set and then add the two ports for your pair. Since I only have two enabled ports, I can only choose the defaults:



Once I've chosen my security zone and inline set, I'll move on to the Advanced tab:

Edit Inline Set

A screenshot of a software dialog box titled "Edit Inline Set". The dialog has two tabs: "General" and "Advanced", with "Advanced" selected. Below the tabs, there are four settings, each with a checkbox:

- Tap Mode:
- Propagate Link State:
- Transparent Inline Mode:
- Strict TCP Enforcement:

At the bottom right of the dialog are two buttons: "OK" and "Cancel".

Here's a brief look at the options on this page:

Tap Mode

This option restricts the inline pair to viewing the data and reporting only. It doesn't enable blocking or denying any packets.

Propagate Link State

A useful option in the event you lose one side of the link switch port. Propagate Link State informs the other side of the inline pair that the data link has lost its connection.

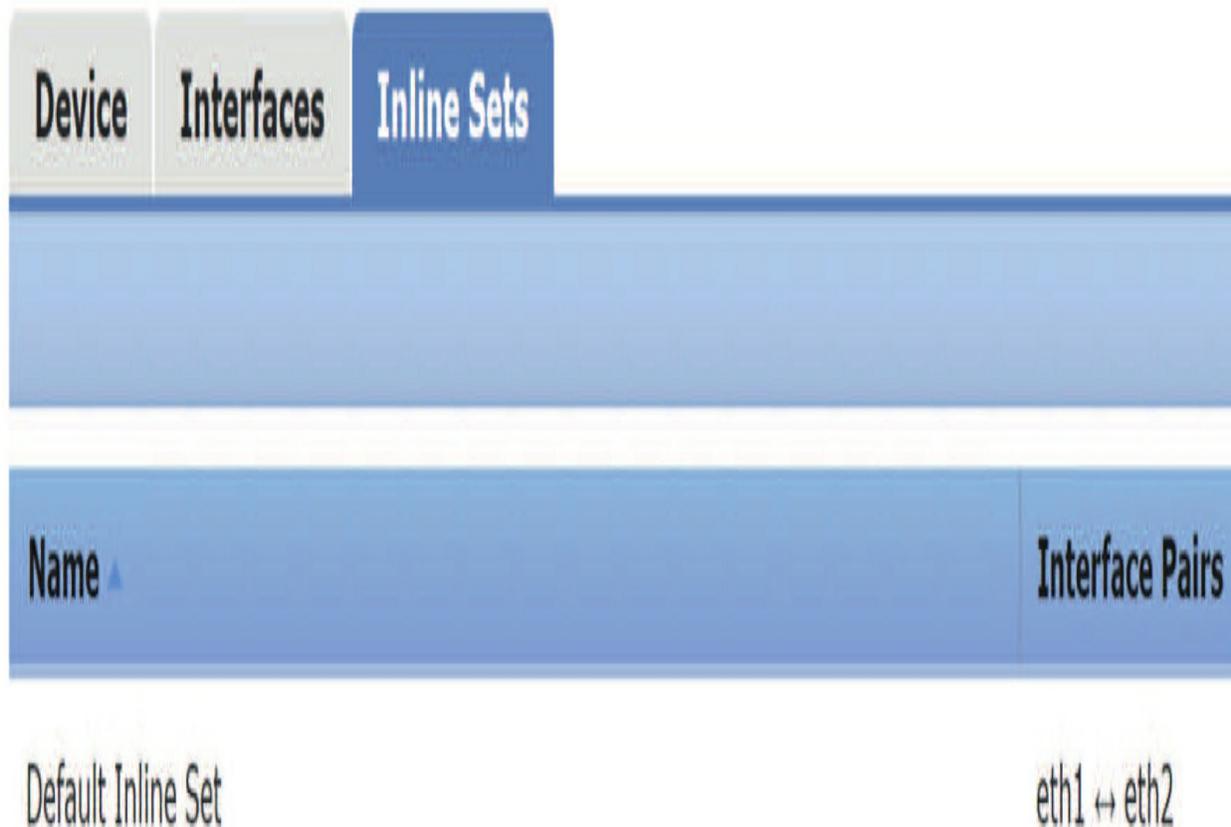
Transparent Inline Mode

The only option on by default, Transparent Inline Mode enables the device to act as a “bump in the wire,” meaning the device will forward all the network traffic it sees, regardless of its source and destination.

Strict TCP Enforcement

This option allows the inline pair to look for and deny any traffic that doesn't start with a TCP 3-way handshake.

From this current page, I'll click OK and then see this.



Device	Interfaces	Inline Sets
Name	Interface Pairs	
Default Inline Set	eth1 ↔ eth2	

Okay- now let's look at the 4140.

I want to point out the lack of interfaces even though the 4140-1 is an eight-

port device. That's because we only see the interfaces that we've actually configured and enabled in the Chassis Manager.

You should also notice that interface 1/3 is already enabled because it was configured in FXOS as my management port.

Here's a short list of the IPs:

Interface Zone IP Address

1/1 Inside 10.17.117.1/24

1/2 Outside 10.11.12.170/24

Let's configure them now by clicking on the pencil to the left of each interface.

4140-1



Cisco Firepower 4140 Threat Defense

Device	Routing	Interfaces	Inline Sets	DHCP	
Search by name Sync Device Add Interfaces ▾					
Interface	Logical Name	Type	Security Zones	MAC Address (Active/Standby)	IP Address
Ethernet1/1		Physical			
Ethernet1/2		Physical			
Ethernet1/3	diagnostic	Physical			
Ethernet1/8		Physical			

When I clicked on the pencil for E1/1, I received this screen for configuring the interface:

Edit Physical Interface

General IPv4 IPv6 Advanced

Name: Enabled

Description:

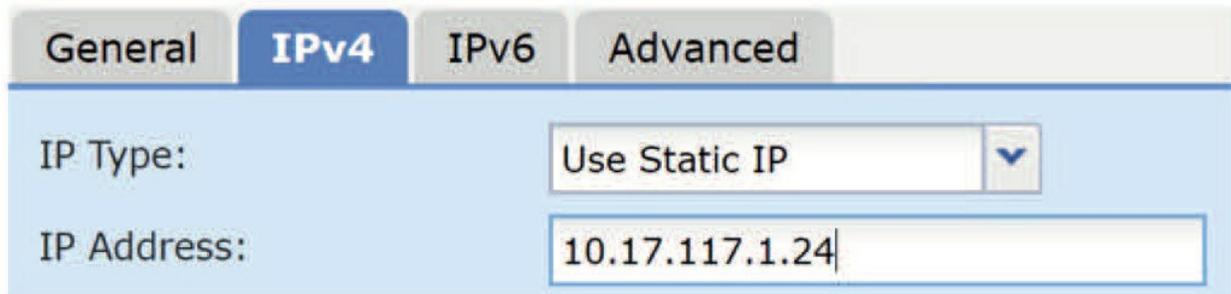
Mode: ▼

Security Zone: ▼

Interface ID:

MTU: (64 - 9184)

Notice the four tabs. From the General tab, I'll choose the name and create a zone, (or add an existing zone), then I proceeded to click on the IPv4 tab and enter the IP:



The image shows a configuration interface with four tabs: General, IPv4, IPv6, and Advanced. The IPv4 tab is currently selected and highlighted in blue. Below the tabs, there are two input fields. The first is labeled 'IP Type:' and has a dropdown menu with 'Use Static IP' selected. The second is labeled 'IP Address:' and contains the text '10.17.117.1.24'.

I'll cover the Advanced Tab when we configure the Platform Settings policy, so for now I'll just configure 1/2, the outside interface. Notice that I didn't forget to enable the interface in the check box on the upper right side. It's too easy to do, so make sure you don't forget either.

General IPv4 IPv6 Advanced

Name: Enabled

Description:

Mode: ▼

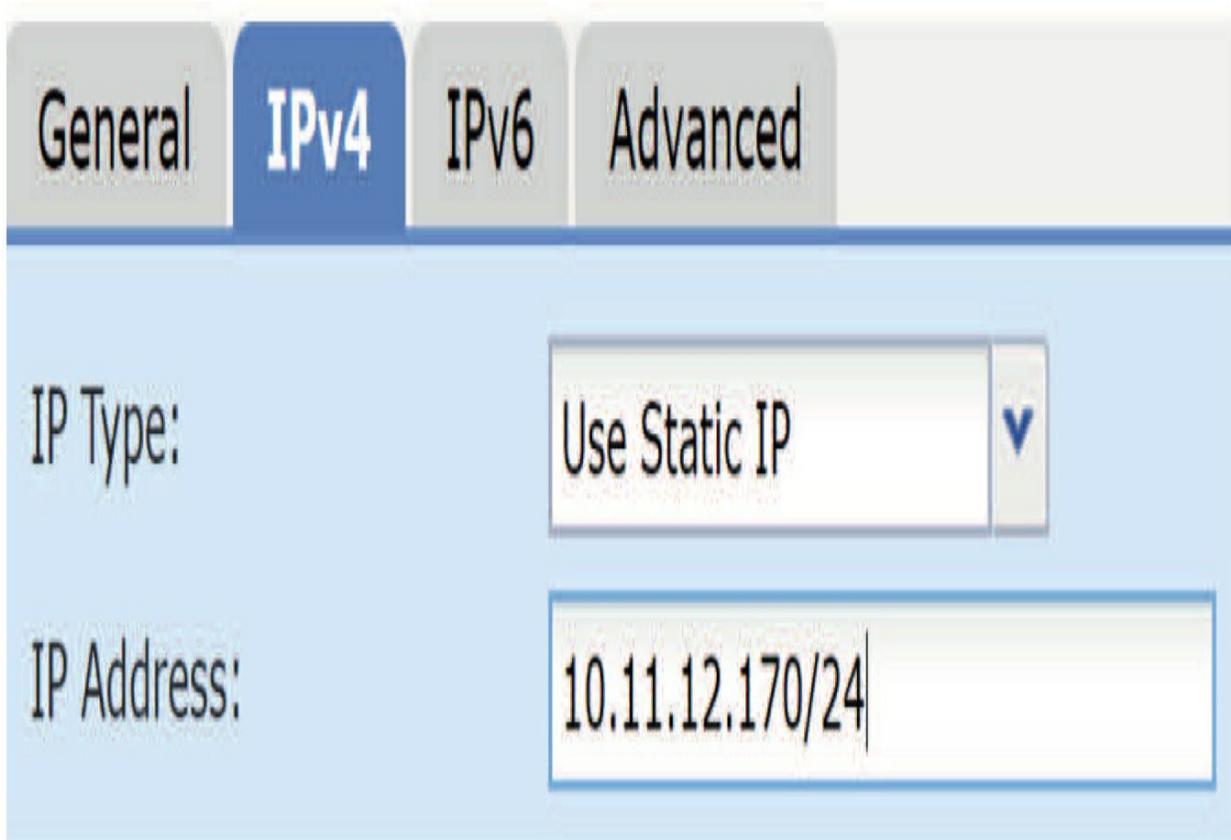
Security Zone: ▼

Interface ID:

MTU: (64 - 9184)

I'll edit the outside interface by adding the IP address and mask on the IPv4

tab:



The image shows a network configuration interface with four tabs: General, IPv4, IPv6, and Advanced. The IPv4 tab is selected and highlighted in blue. Below the tabs, there are two configuration fields. The first field is labeled 'IP Type:' and has a dropdown menu set to 'Use Static IP'. The second field is labeled 'IP Address:' and contains the text '10.11.12.170/24'.

After the IPs are set on 1/1 and 1/2, I also enabled 1/8. I didn't perform any other configuration on that port because it's going to be my HA interface, so you just need to enable it.

Here's another place you need to remember to go to the upper right and click Save.

4140-1



Cisco Firepower 4140 Threat Defense

Device Routing **Interfaces** Inline Sets DHCP

Search by name Sync Device Add Interfaces ▾

Interface	Logical Name	Type	Security Zo...	MAC Address (Active/...	IP Address
Ethernet1/1	Inside	Physical	Inside		10.17.117.1/24(Static)
Ethernet1/2	Outside	Physical	Outside		10.11.12.170/24(Static)
Ethernet1/3	diagnostic	Physical			
Ethernet1/8		Physical			

Displaying 1-4 of 4 interfaces < < Page 1 of 1 > > ↻

Now, you'll want to check out the 4140-2 Interface page as well after configuration and verify the IP's.

4140-2

You have unsaved changes

Please save the configuration to make the changes available for use. ✕

Cisco Firepower 4140 Threat Defense

Device Routing **Interfaces** Inline Sets DHCP

Search by name Sync Device Add Interfaces ▾

Interface	Logical N...	Type	Security ...	MAC Address (Activ...	IP Address
Ethernet1/1	Inside	Physi...	Inside		10.18.118.1/24(Static)
Ethernet1/2	Outside	Physi...	Outside		10.11.12.180/24(Static)

At this point, I'm going to go ahead and configure the virtual FTDs 19 and 20 the same way, using the addressing I've documented.

Configuring Routing

Configuring routing in the Firepower system is really simple and straightforward. You don't even have to understand routing well to make this work, which was actually the idea!

For this book, I'm going to use OSPF and static default routing on all devices. To configure OSPF, click on the Routing tab and the first protocol listed is OSPF—the version required for IPv4.



This screen will only provide two processes numbers: 1 and 2, and you can't use any other ones.

OSPF process IDs are irrelevant numbers anyway, and typically, the defaults work for most networks. Still, detailed configuration for advanced OSPF

networks is available if you need it.

Okay—so after you choose a process ID, the default Internal Router works just fine, and I then clicked on Add down on the right. Doing this allowed me to add the networks of 0.0.0.0 with an object called any-ipv4:

The only change I had to make was the Area ID, which for some reason is set to 1 by default even though most routers use area 0 by default. After that simple change, I just clicked Save.

Then, to also create a static route, I clicked on Add Static Route Configuration on the left-hand side to add my default zone and networks. Again, the value is just 0.0.0.0 since it's a default route:

Add Area

Area

Range

Virtual Link

OSPF Process:

Area ID:*

Area Type:

Summary Stub

Redistribute

Summary NSSA

Default Information orig

Metric Value:

Metric Type:

Available Network



Selected Network

Last, I configured the next hop address in the Gateway field and that's it!
Now we have routing set up and configured on our router.

Add Static Route Configuration

Type:

IPv4 IPv6

Interface*

Outside

Available Network



Search

- any-ipv4
- IPv4-Benchmark-Tests
- IPv4-Link-Local
- IPv4-Multicast
- IPv4-Private-10.0.0.0-8
- IPv4-Private-172.16.0.0-12
- IPv4-Private-192.168.0.0-1
- IPv4-Private-All-RFC1918
- IPv6-to-IPv4-Relay-Anycast

Selected Network

any-ipv4

Add

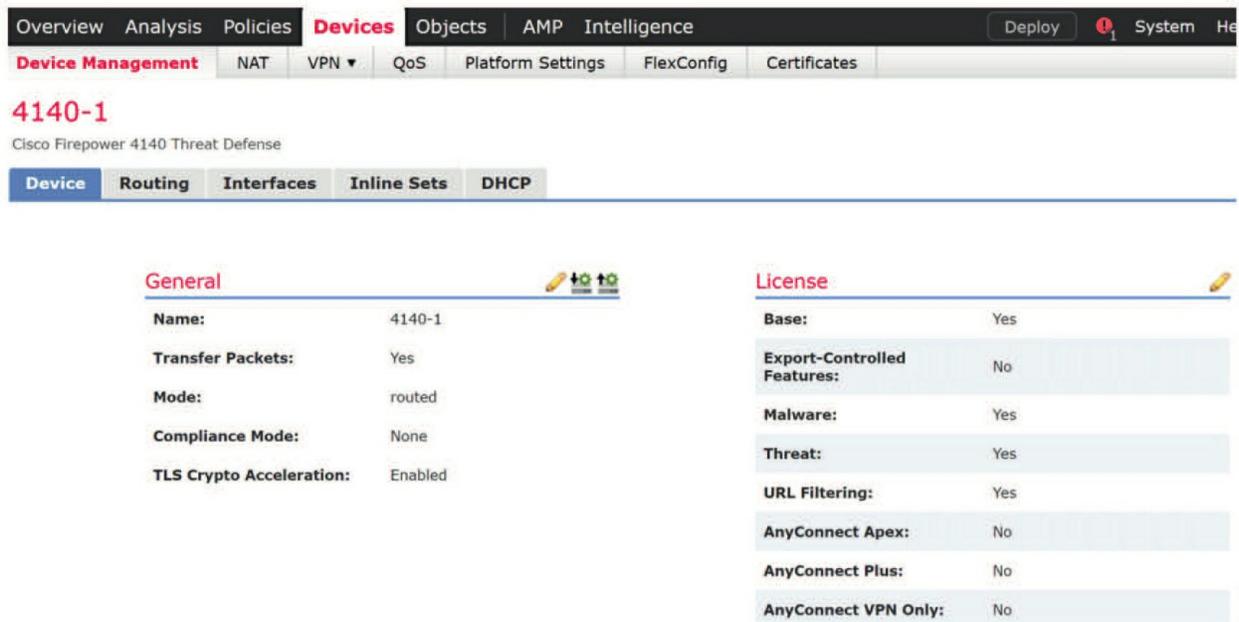
Gateway*

10.11.12.1

Device Tab

There's one more tab that's important and it must be addressed for each device you add—the Device tab.

First, on the right side, you can add or delete unused licenses for your device. You can also manage the licenses for the device from the **System>Licenses** page:



The screenshot shows the Cisco Firepower Management Center (FMC) interface. At the top, there is a navigation bar with tabs for Overview, Analysis, Policies, **Devices**, Objects, AMP, and Intelligence. Below this, there is a sub-navigation bar with tabs for Device Management, NAT, VPN, QoS, Platform Settings, FlexConfig, and Certificates. The main content area is titled "4140-1" and "Cisco Firepower 4140 Threat Defense". Below this, there is a sub-navigation bar with tabs for **Device**, Routing, Interfaces, Inline Sets, and DHCP. The main content area is divided into two sections: "General" and "License".

General	
Name:	4140-1
Transfer Packets:	Yes
Mode:	routed
Compliance Mode:	None
TLS Crypto Acceleration:	Enabled

License	
Base:	Yes
Export-Controlled Features:	No
Malware:	Yes
Threat:	Yes
URL Filtering:	Yes
AnyConnect Apex:	No
AnyConnect Plus:	No
AnyConnect VPN Only:	No

When you click the pencil on the General menu here, you can't do much other than change the name. However, you can disable Transfer Packets if you no longer want to receive IP packets when a Snort event occurs.

Last, there's actually a hidden command that allows you to deploy to a device even when the device is "up to date" and you can't deploy policy with the Deploy command:

General ? X

Name::	<input type="text" value="4140-1"/>
Transfer Packets:	<input checked="" type="checkbox"/>
Mode:	routed
Compliance Mode:	None
Force Deploy:	

Now I have to tell you that I've never found any use for this hidden command, meaning that deploying when the devices are already up to date hasn't ever solved any issues for me.

Okay—so scrolling further down the page provides a few more options for us when managing our device. For most devices, the System menu provides a helpful Shutdown and Restart button for the FTD device, but for the 4100/9300, this can only be done from the Chassis Manager.

Do keep this in mind for all your other devices though. The Health and Management menus provide information as well as a way to stop our FMC from temporarily managing the FTD device.

System



Model: Cisco Firepower 4140 Threat Defense

Serial: FLM2048QJJ1

Time: 2019-11-05 22:41:05

Time Zone: UTC (UTC+0:00)

Version: 6.5.0

Health

Status: 

Policy: [Initial Health Policy 2019-10-27 12:18:37](#)

Blacklist: [None](#)

Management



Host: 10.11.10.207

Status: 

Scrolling further, we can now get inventory details and information about the

policies that are applied to our devices. This is a new and very useful feature included starting in the 6.5 code, so if you're looking at a previous, older code version, you won't see these.

Inventory Details		Applied Policies	
Cpu Type:	CPU Xeon E5 series 2300 MHz	Access Control Policy:	VFMC20-ACP
Cpu Cores:	2 CPUs (72 cores)	Prefilter Policy:	Default Prefilter Policy
Memory:	137864 MB RAM	SSL Policy:	
Storage:	NA	DNS Policy:	Default DNS Policy
Chassis Url:	https://A4140-1.sfgtc.local:443//	Identity Policy:	
Chassis Serial Number:	JAD2052043V	Nat Policy:	
Chassis Module Number:	1	Platform Settings Policy:	
Chassis Module Serial Number:	NA	NGFW QoS Policy:	
		FlexConfig Policy:	

I have absolutely found the Applied Policies section to be really helpful! Even though we don't really have much applied yet, we will later on. Just wait, you'll see!

Scrolling all the way to the bottom, we find one more section—the Advanced menu with the Automatic Application Bypass (AAB) section. It's vital enough to make it worth your time to go to the Device tab on each and every device and appliance to enable this!

This parameter cannot be set globally:

Advanced



Application Bypass:	No
Bypass Threshold:	3000 ms

I want to stop and tell you about the little ? that the arrow is pointing to in the following screen shot. This feature is so awesome! When you click on the ? you'll be taken to a web page that defines exactly what you're working on. Yep, the ? does *not* disappoint and you can use it on pretty much every page

and configuration:

Clicking on the pencil allows us to enable AAB, and no one can tell me why this is disabled by default.

Advanced



Automatic Application Bypass:

Bypass Threshold (ms):

Save Cancel

The Automatic Application Bypass (AAB) feature limits the time permitted to process packets through an interface, and it allows packets to bypass detection if the time is exceeded. The feature functions with any deployment, but it really shines in inline deployments.

You should enable AAB on every device. It should be an exception when it is not enabled.

You balance packet processing delays with your network's tolerance for packet latency. When events occur like a malfunction within Snort or a device misconfiguration causes traffic processing time to exceed a specified threshold, AAB causes Snort to restart within ten minutes of the failure and

generates troubleshooting data that can be analyzed to investigate the cause of the excessive processing time.

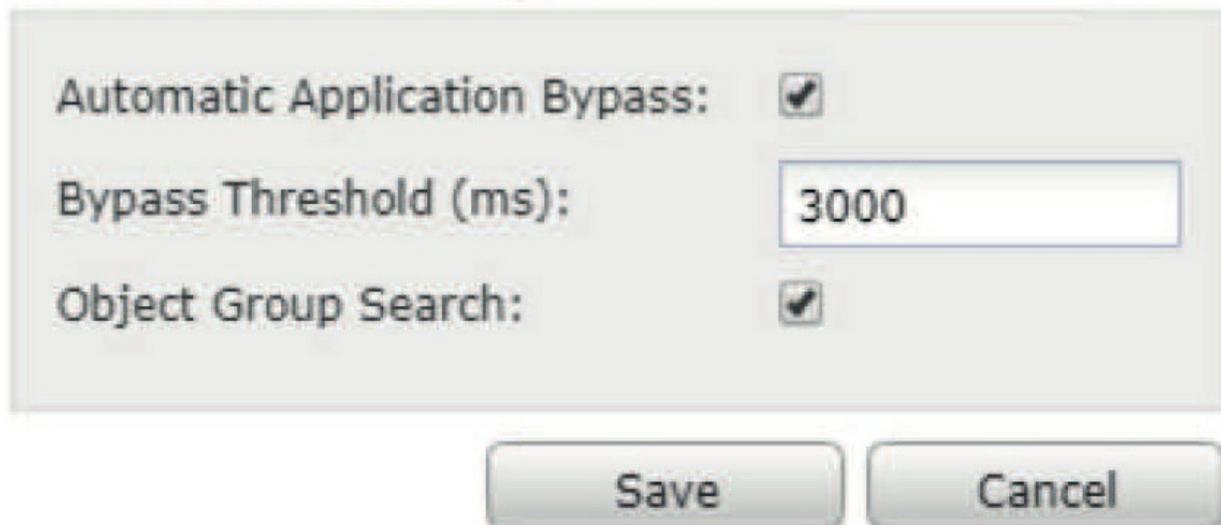
But Wait! There's More!

One last thing I want to mention in this chapter—with the newer codes you'll obtain yet another option in the Advanced Settings section of Devices called Object Group Search.

This is basically an old ASA command that they now support on the Firepower FTD devices for configurations with hundreds of thousands or millions of ACL rules.

This command reduces the memory required to search access control rules by enabling object group search, but it comes at the expense of lookup performance and increased CPU utilization.

Advanced Settings



Automatic Application Bypass:	<input checked="" type="checkbox"/>
Bypass Threshold (ms):	<input type="text" value="3000"/>
Object Group Search:	<input checked="" type="checkbox"/>

Save Cancel

When enabled, Object Group Search doesn't expand network or service objects into RAM, it instead searches access rules for matches based on those group definitions.

Just like AAB, Object Group Search is disabled by default. You can enable it on one device at a time; you can't enable it globally.

Warning! If you enable Object Group Search, and it's been in use for a while, be aware that disabling the feature might lead to some pain. When you disable Object Group Search, your existing access control rules will then be expanded in the device's running configuration. If the expansion requires more memory than is available on the device, your device can be left in an inconsistent state and you'll probably experience impacted performance.

Summary

In this chapter, I spent a lot of time on Firepower Threat Defense (FTD) and the configuration of the LINA process that the Snort process runs on top of. Once we got our routers and devices configured, we began configuring the Snort policies we'll use in the chapters ahead.

We took a quick look at the old 7000/8000 appliances and how they're configured because the exam objectives still cover these Firepower devices for now. So, we configured and brought a Firepower appliance into the vFMC.

We moved on to configure a pair of Firepower 1010s and a pair of 1150s and we also continue with the 4140s I configured with Chassis Manager in Chapter 5. Finally, I added two virtual FTDs to the mix to top it off.

Once they were configured, we added all nine FTD devices into the two FMCs that I configured back in Chapter 1.

Chapter 7: High Availability

The following CCIE/CCNP Security SNCF exam objectives are covered in this chapter:

1.0 Deployment

1.3 Implement high availability options

1.3.a Link redundancy

1.3.b Active/standby failover

In this chapter, we're going to take off from Chapter 1, where I configured the FMCs, and Chapter 6, where the Firepower devices were all configured, and add High Availability into the mix. I've configured way too much to lose

my configurations at this point if there was a failure!

A lot of people ask me how to create an HA-type environment with their virtual FMC. I tell them that they should just do a backup every single night with any FMC they have, then make sure it's all copied off their FMC system because the backups are stored on the FMC by default.

I'm sure the issue here regarding the vFMC is pretty clear: Doing a snapshot on your vFMC is a bad idea because it just grows way too big! I've had clients contact me in a panic because an admin deleted the snapshot and corrupted the vFMC, with no backup....oops! I also happen to have customers that use vMotion and tell me it's worked for them without issue. Just be advised that Cisco does not support this configuration.

One last thing before we get started... I get a lot of questions on when you can deploy the HA on the Firepower devices. Well, for the FMC, you can basically do that anytime, even though you lose connection for a bit while they reboot. This doesn't bring your network down, but you'll still lose malware SHA checks and AD integration for a very short period. It's really not a problem because most people do this in a maintenance window anyway. When it comes to the Firepower and FTD devices, once you have a HA FTD pair, you can perform upgrades most anytime because the upgrade is performed on one device at a time automatically, which is great! You'll see this in the System Upgrades section at the end of this chapter. People still usually do the updates within a maintenance window as well—just in case. With all that in hand, I'm going to start configuring HA on the hardware 2500 FMC, then move over to the 1150s and 1010s. After that, I'm going to upgrade the 2500 FMCs to a new code and then upgrade the FTD devices, and I'm going to do all of this while keeping my production network running! How fun is this chapter going to be?

High Availability

Configuring High Availability, sometimes called Failover, requires two identical FMCs, Firepower Appliances or Firepower Threat Defense devices that are connected to each other through a dedicated link. You can also add an additional state link between devices, but that's not required.

Cisco's Firepower appliances and Firepower Threat Defense support active/standby failover, where one unit is the active one that passes traffic.

The standby unit doesn't actively pass traffic; instead it synchronizes configuration and other state information from the active unit. If a failover occurs, the active unit will then fail over to the standby unit, which will then become the active unit.

So, what are the hardware requirements for HA on any Firepower device? The two units in a High Availability configuration must be exactly the same model with the exact same interfaces. There's one exception to this: the Firepower 9300 devices as High Availability is only supported between same-type modules, but the 9300 chassis can include mixed modules and each chassis has an SM-36 and SM44. So you can create High Availability pairs between the SM-36 modules and between the SM-44 modules. Most of us can't afford these devices anyway, but now you know in case they ask about that on the exam.

Additionally, and picking up from Chapter 5, the 4100/9300 chassis must have all identical interfaces preconfigured in FXOS before you can enable High Availability. If you need to change the interfaces after you enable HA, make the changes in FXOS on the standby unit first, then make the same changes on the active unit.

Licensing

In an HA configuration, each Firepower Threat Defense device must have the same licenses. It doesn't matter which licenses are assigned to the secondary/standby device before HA is established because the FMC will release any unnecessary licenses assigned to the standby device and replace them with the identical licenses assigned to the primary/active device.

Determining the Active Unit

You'll absolutely determine the primary and secondary units during configuration, but what happens on a reboot? These points may sound pretty obvious, but we need to cover the exam objectives!

The active unit is determined by the following:

- If a unit boots and detects a peer already running as active, it becomes the

standby unit.

- If a unit boots and does not detect a peer, it becomes the active unit.
- If both units boot simultaneously, then the primary unit becomes the active unit and the secondary unit becomes the standby unit.

Okay, let's do some HA!

High Availability on the Hardware FMC

This might actually be the easiest section of the book, so let's breeze through it. To create an HA pair with the hardware FMCs, all that's needed is for each part of the pair to connect together directly, or through a switch VLAN with an IP address on each management port. Just remember that they have to be the same exact hardware, same code, and same IPS version. You'll find out pretty quick if you've got something mismatched!

To log in to each FMC and verify that all is well code-wise, go to **System>Integration>High Availability**.

It's important to start the upgrade on the secondary device first and then configure the primary. As I said, it's a relatively painless process, as long as your codes and hardware match.

The screenshot shows a navigation bar with tabs: Cloud Services, Realms, Identity Sources, High Availability (selected), eStreamer, Host Input Client, and Smart Software Satellite. Below the navigation bar, the text reads: "Select a role for this Management Center and specify peer details to setup high availability." Underneath, there is a section titled "Role For This FMC:" with three radio button options: "Standalone (No High Availability)" (which is selected), "Primary", and "Secondary".

So, as you can see here, the default mode is Standalone. I'm going to click on Secondary at this point and type in the Primary information.

Peer Details:

Primary FMC
Host:

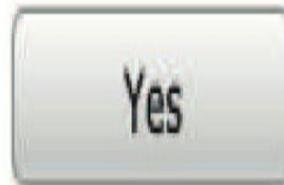
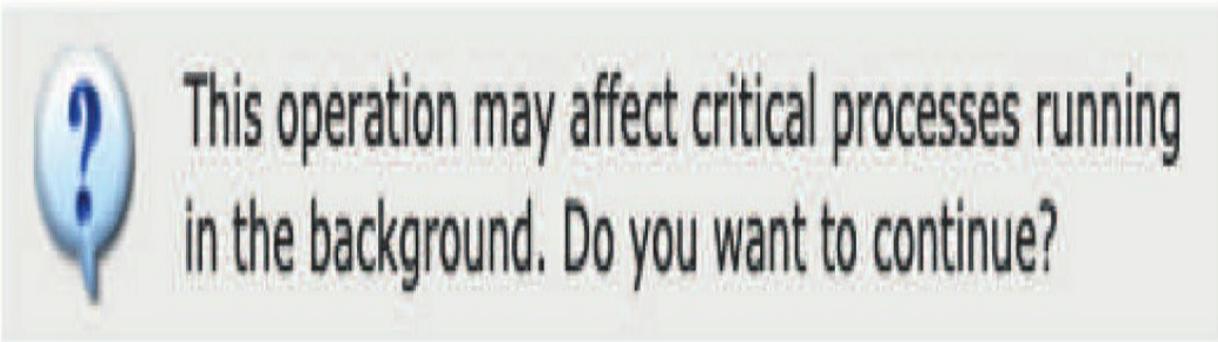
Registration
Key*:

Unique NAT ID:

† Either host or NAT ID is required.

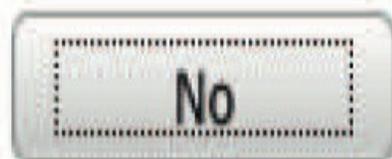
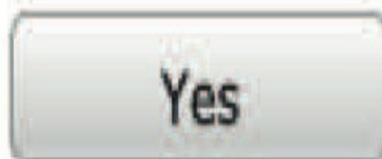
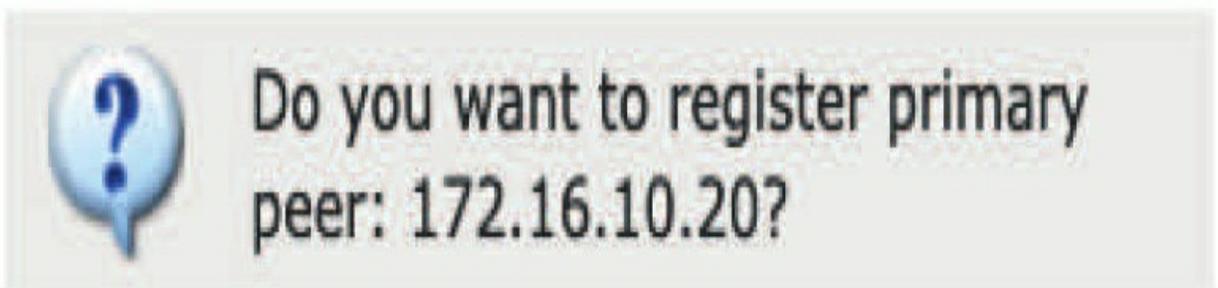
Once you click Register,
you'll get this warning:

Warning

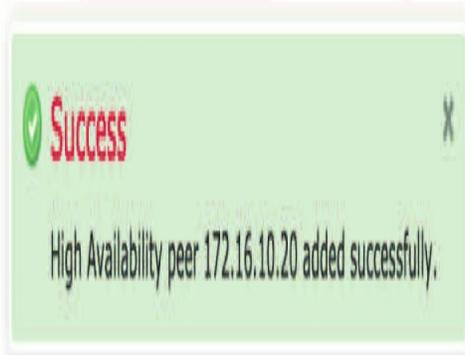


...and then another message;
just say Yes to both.

Warning



Until finally the secondary is ready to connect to the primary.



Host	Last Modified	Status	State
172.16.10.20	2019-11-16 13:24:35	Pending Registration	  

Now, let's do the same configuration on the primary FMC. Just add in the Secondary IP, registration key, and then click Register.

You basically get the same messages, except the second message tells you that the secondary FMC configuration will be cleared out— that's what we want at this point.

Role For This FMC:

Standalone (No High Availability)

Primary

Secondary

Peer Details:

Configure the secondary Management Center with details of the primary, before registration.

Secondary FMC Host:

172.16.10.21

Registration Key*:

cisco

Unique NAT ID:

Register

Warning



Do you want to register secondary peer: 172.16.10.21? Secondary peer configuration and policies will be removed.

Yes

No



Switch Peer Roles



Break HA



Pause Synchronization

High availability operations are in progress. The status messages and alerts on this page are temporary. Please check after high availability operations are complete.

These operations include file copy which may take time to complete.

Database files synchronization: 100% of 126MB transferred

Summary

Status



Temporarily degraded- high availability operations are in progress.

Synchronization



Failed

Active System

172.16.10.20

Standby System

172.16.10.21

System Status

Local

Remote

Active - Primary
(172.16.10.20)

Standby - Secondary
(172.16.10.21)

Operating System

Fire Linux OS 6.5.0

Fire Linux OS 6.5.0

Software Version

6.5.0-39

6.5.0-39

Model

Cisco Firepower
Management Center
2500

I'll be logged out as the FMC is rebooted. After you log back in, we can open the HA page to see the status.

And now's a great time for a nice, long coffee break!

Okay, so after about 20 minutes or more, I can see that the 2500 FMCs are now synchronized. Also notice in this screen shot from the secondary FMC that the secondary is now missing the tabs that would provide configuration for that FMC. This is because the FMC configuration is now only executed from the primary.

The screenshot shows the HA page with the following components:

- Navigation tabs: Cloud Services, High Availability (selected), eStreamer, Host Input Client.
- Action buttons: Switch Peer Roles, Break HA, Pause Synchronization.
- Summary section:
 - Status: Healthy
 - Synchronization: OK
 - Active System: 172.16.10.20 (HA synchronization time : Sat Nov 16 18:46:49 2019)
 - Standby System: 172.16.10.21 (HA synchronization time : Sat Nov 16 18:48:43 2019)
- System Status table:

	Local Standby - Secondary (172.16.10.21)	Remote Active - Primary (172.16.10.20)
Operating System	Fire Linux OS 6.5.0	Fire Linux OS 6.5.0
Software Version	6.5.0-39	6.5.0-39
Model	Cisco Firepower Management Center 2500	Cisco Firepower Management Center 2500

Now, I want you to notice there at the upper right, that we can switch peer roles, break the HA, or pause synchronization. You want to test this by switching roles at times.

We can also tell that our synchronization worked, and the FMC is managed only by the primary by looking at the Devices page. All four FTD devices are there, but the pencil, trashcan, and troubleshooting icons are no longer here.

The screenshot shows the Devices page with the following table:

Ungrouped (4)						
1010-1 172.16.10.10 - Routed	FTD on Firepower 1010	6.4.0	N/A	Base, Threat (2 more...)	FMC2500-ACP	
1010-2 172.16.10.11 - Routed	FTD on Firepower 1010	6.4.0	N/A	Base, Threat (2 more...)	FMC2500-ACP	
1150-1 172.16.10.12 - Routed	FTD on Firepower 1150	6.5.0	N/A	Base, Threat (2 more...)	FMC2500-ACP	
1150-2 172.16.10.13 - Routed	FTD on Firepower 1150	6.5.0	N/A	Base, Threat (2 more...)	FMC2500-ACP	

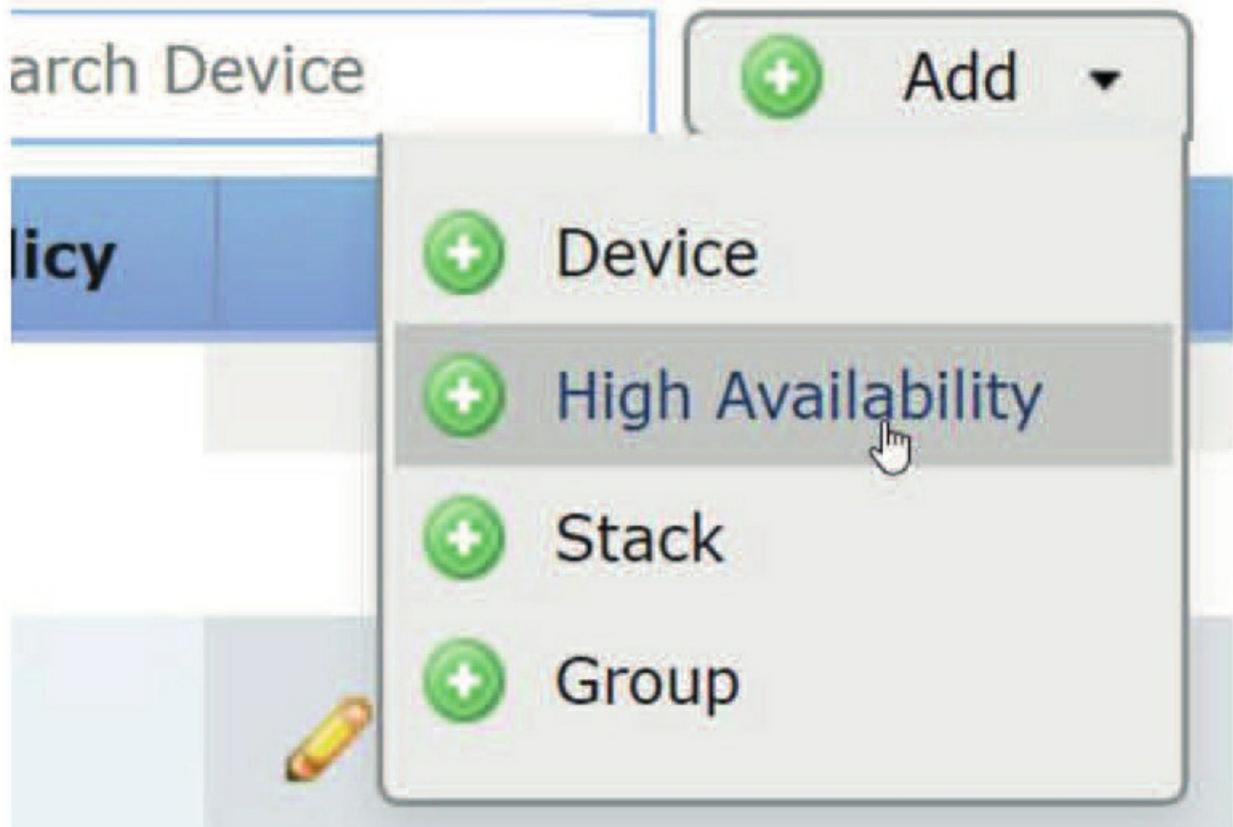
Looking good—now that my 2500s are running in HA mode, I'm going to make my network even better by running HA between all my hardware FTDs!

High Availability on the Firepower Devices

Unlike the FMCs, where it's pretty simple to create an HA pairing, the FTDs require a little more understanding and configuration. That's not all—there's also the risk of bringing down your network while configuring Firepower device HA pairs. I've rarely experienced issues, but just once would definitely be one too many!

So now that I've got your attention, go to your Devices page, and there on the right side, click **Add>High Availability**.

And again, it's vital to make sure the devices you're going to add into a pair are 100% the same hardware and code version. Also key is that you can't include any devices that need an outstanding deploy on them. So before you start, always run a deploy first!



Okay—the first thing you need to do now is to name the HA pair and then choose Firepower appliances or FTD devices depending on which type you’re pairing. For us, that would be FTD.

Add High Availability Pair

Name:*	<input type="text" value="4140-HA"/>
Device Type:	<input type="text"/>
Primary Peer:	<input type="text" value="Firepower Threat Defense"/>
Secondary Peer:	<input type="text" value="Firepower"/>

Next, choose your primary device:

Add High Availability Pair



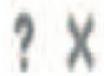
Name:*	4140-HA
Device Type:	Firepower Threat Defense
Primary Peer:	4140-1
Secondary Peer:	4140-1 4140-2 vFTD-19 vFTD20

Threat Defense High Availability configuration. Licenses must be purchased for both devices to support their high availability configuration.

When you click on the secondary, you'll only receive the options that are compatible with your primary. Only the 4140-2 showed up as an option when I clicked on the Secondary Peer option as you can see.

Now press Continue and it'll tell you that it's going to restart Snort.

Add High Availability Pair



Name:*

Device Type:



Primary Peer:

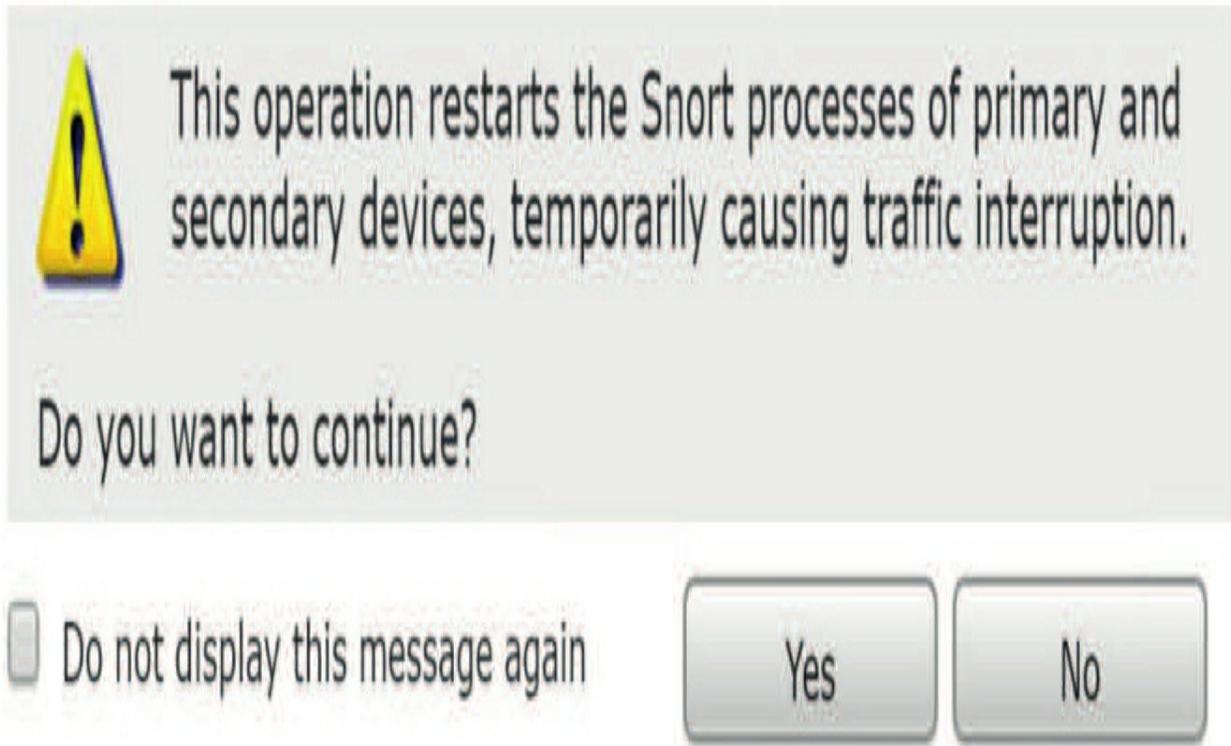


Secondary Peer:



 Threat Defense High Availability pair will have primary configuration. Licenses from primary peer will be converted to their high availability versions and applied on both peers.

Warning



Now this is where you're going to receive an error if everything isn't properly matched up or if you've got an outstanding deployment. Because everything was in order here, I received the next configuration screen.

First, I'll configure the HA interface that I want to use to communicate between the units.

Since these are my 4140s, I only have one more interface available to configure (from Chapter 5) for my HA link.

Add High Availability Pair

High Availability Link

Interface:*

- Select -



Logical Name:*

- Select -

Ethernet1/8

Primary IP:*

Use IPv6 Address

Secondary IP:*

Subnet Mask:*

The E1/8 interfaces on each 4140 are hooked up directly point-to-point with a fiber connection.

It's important to remember that the following information is communicated over the failover link, which is why this link is so crucial:

- The unit state (active or standby)

- Hello messages (keep-alives)
- Network link status
- MAC address exchange
- Configuration replication and synchronization

Once I've chosen the interface, I can select the name and IPs for

the direct link, as shown here:

I used to go with the 1.1.1.1 and 1.1.1.2 IPs for my HA link, but we don't want to do that anymore as you're well aware of by now. So instead, I'll just choose RFC 5735 reserved addresses of 198.18.0.0/15.

High Availability Link

Interface:*	Ethernet1/8
Logical Name:*	4140-HA
Primary IP:*	198.18.0.1
	<input type="checkbox"/> Use IPv6 Address
Secondary IP:*	198.18.0.2
Subnet Mask:*	255.255.255.252

As long as the connection happens to be one that's directly connected to one another as mine are here, you can choose the same IP addresses for all your pairs. You can't do that if the connection is via a switch or VLAN.

In my experience with FMC pairs, clients usually spread the devices out between data centers, and they're usually doing this through an extended VLAN. But since the FTD devices here are usually in the same location, using a direct Ethernet/fiber connection works great.

I need to point out that if you don't place a switch between the units and an interface fails, the link will be brought down on both peers. This makes troubleshooting harder because you can't easily tell which unit has the failed interface that caused the link to come down.

NOTE: You cannot specify an interface for your failover link that's

currently configured with a name.

So, I recommend that failover links never use the same switch as

the data interfaces. Instead, either use a different switch if possible or opt for a direct cable to connect the failover link.

For the secondary configuration here, I'll just choose Same as LAN Failover Link, or E1/8 (it doesn't matter), and my configuration is complete.

Add High Availability Pair



High Availability Link

Interface:*

Ethernet1/8

Logical Name:*

4140-HA

Primary IP:*

198.18.0.1

Use IPv6 Address

Secondary IP:*

198.18.0.2

Subnet Mask:*

255.255.255.252

State Link

Interface:*

- Select -

Logical Name:*

- Select -

Same as LAN Failover Link

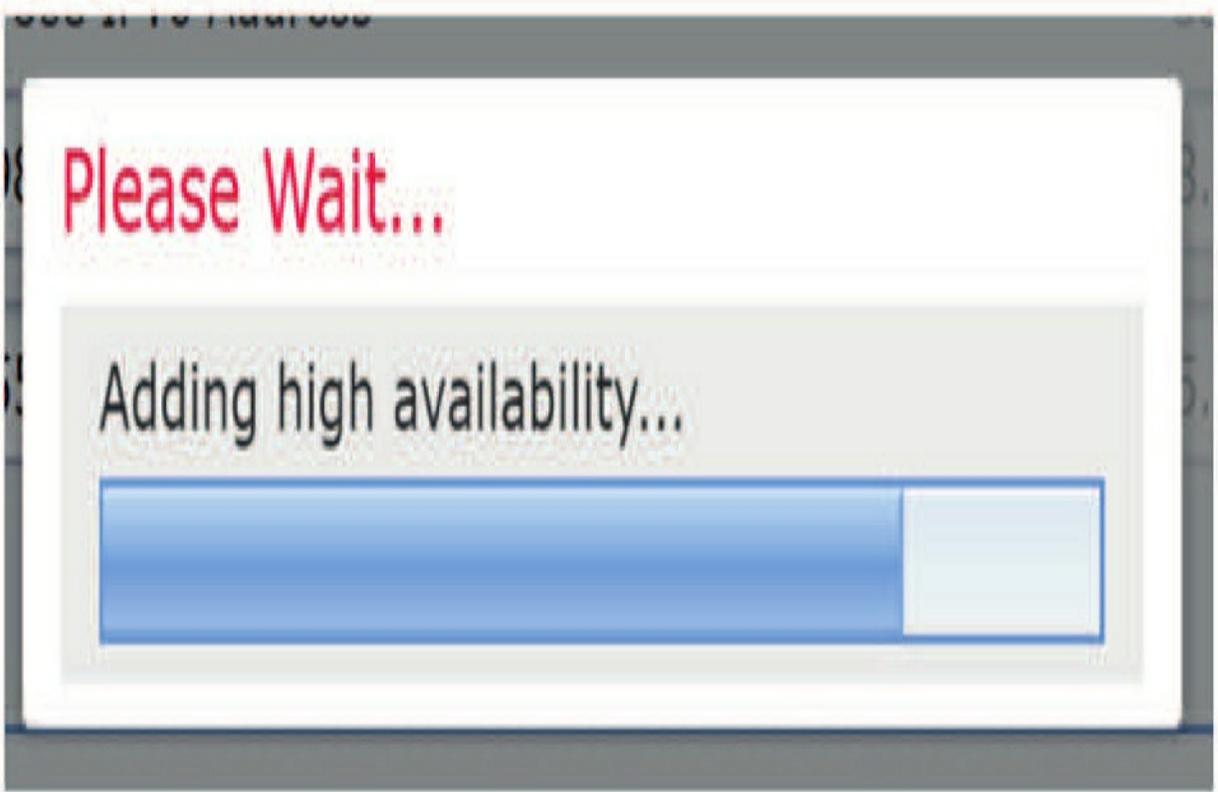
Primary IP:*

Ethernet1/8

Secondary IP:*

Subnet Mask:*

Here, we can see the HA paring beginning. We'll just let it go for about 20 minutes.



Okay, after another extended coffee break, I'm back and we can see that 4140s are now in a HA pair—nice!

4140-HA High Availability

4140-1 (Primary, Active)
10.11.10.207 - Routed

FTD on Firepower
4140

6.5.0

A4140-1.sfgtc.local:443
Security Module - 1

4140-2 (Secondary, Standby)
10.11.10.209 - Routed

FTD on Firepower
4140

6.5.0

c4140-2.sfgtc.local:443
Security Module - 1

Of course, we could just leave the defaults in place and let them run the HA configuration, and we will—for now. After I go through the process of adding the FTD HA pair in my 2500, I'll demonstrate some of the configuration options we can add to fine-tune our HA's configuration.

So here, I also added the two pairs of 1150s and 1010s in my 2500 FMC, and now they look like the screen below. I skipped the configuration because they're all exactly the same as it was with the 4140 HA configurations. This looks good:

In a perfect world, these always come up and go green first time like mine did!

1010-HA
High Availability



1010-1(Primary, Active)
172.16.10.10 - Routed
FTD on
Firepower
1010
6.4.0 N/A
Base, Threat (2
more...)
[FMC2500-ACP](#)



1010-2(Secondary, Standby)
172.16.10.11 - Routed
FTD on
Firepower
1010
6.4.0 N/A
Base, Threat (2
more...)
[FMC2500-ACP](#)



1150-HA
High Availability



1150-1(Primary, Active)
172.16.10.12 - Routed
FTD on
Firepower
1150
6.5.0 N/A
Base, Threat (2
more...)
[FMC2500-ACP](#)



1150-2(Secondary, Standby)
172.16.10.13 - Routed
FTD on
Firepower
1150
6.5.0 N/A
Base, Threat (2
more...)
[FMC2500-ACP](#)



Because it's not a perfect world, troubleshooting is a pretty big deal for HA. Let's get into that now.

Oops now wait—before we dive into monitoring and troubleshooting, I did actually get one message when I was adding my 1150s that would be good to show you.

Peer Configuration Mismatch



Review the configuration mismatch list. High availability can be created only after all configurations between active and standby peers match.

Summary

- ✓ Peers are under the same device group
- ✓ Peers have same type of interfaces
- ✓ Peers have same number of interfaces
- ✓ Peers do not have interfaces with a DHCP or PPPoE configuration
- ✗ There are pending deployment tasks on peers
- ✓ No ongoing deployments are in progress for the selected peers
- ✓ Secondary Peer is not configured in any of the VPN topologies
- ✓ All required certificates are enrolled in Secondary Peer
- ✓ Peers are in same compliance mode
- ✓ Peers are having same resources size

I thought I had deployed everything but looks like I forgot to deploy for this pair. Just look at all those checks the system goes through before you can get HA up and running on your devices!

For a nice map to the exam objectives, here are all the details for the two units in a High Availability configuration that must come up. The devices must:

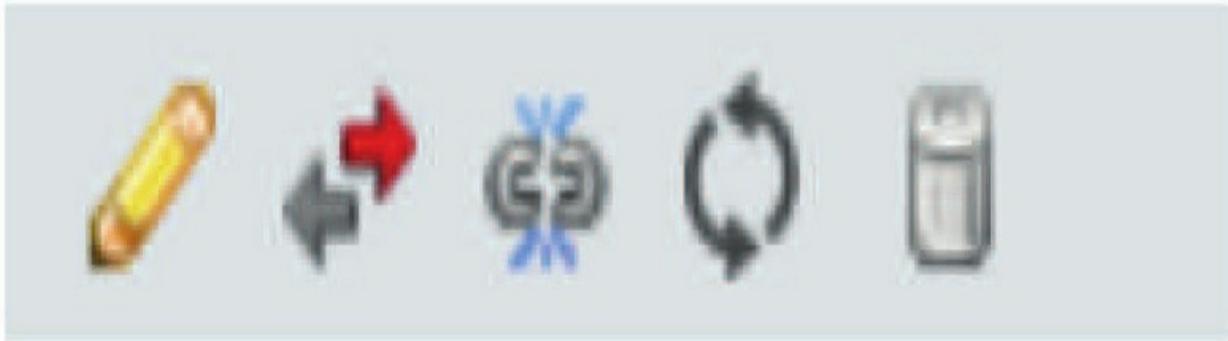
- Be in the same firewall mode (routed or transparent).
- Have the same software version.
- Be in the same domain or group on the Firepower Management Center.
- Have the same NTP configuration

- Be fully deployed on the Firepower Management Center with no uncommitted changes.
- Not have DHCP or PPPoE configured in any of their interfaces. ▪ (Firepower 4100/9300) Have the same flow offload mode, either both enabled or both disabled.

Now it's time to dive into this chapter's most important section!

Monitoring and Troubleshooting

Let's navigate back to the Devices screen... On the right side of the HA pair in the Devices page, you'll see these icons:



The pencil allows you to edit the HA pair, the arrows are used to switch pairs; then you can break the pair, refresh the status, and with that last one, you can delete the pair. Let's find out what we can edit by clicking the pencil.

Note: It's always good to test your HA by switching pairs to make sure you don't lose connection when they fail over. First thing to notice is there's a new tab in the configuration of the

device. You can see the statistics by clicking on the icon on the right, which will provide you with some statistics to establish if you're really sending and receiving packets.

Notice that you can go to each device individually to check its stats.

Stateful Failover Statistics

Select a device: 4140-1(Primary,Active) ▼

Parameters	4140-1(Primary,Active)	4140-2(Secondary,Standby)	Transmission Errors	Received Packets
General	1607	0		1606
System Commands	1606	0		1606
ASA Up Time	0	0		0

I have to tell you that I cut the output here because it's really, really long. No worries, though, I gave you enough to get a good idea of the kind of information you get with the pencil icon.

Monitoring Interfaces

Next up on the page, you can see the monitored interfaces. You need to add all your hardware interfaces in here if they aren't enabled to be monitored by

default. Just remember not to monitor your subinterfaces or other logical interfaces.

You can also edit the monitored interfaces and add a standby IP address, which is recommended but not required. You definitely want to configure this, so I'll cover this topic soon.

By default, monitoring is enabled on all physical interfaces—for some devices, all VLAN interfaces that have logical names configured, too, like the 1010.

It's also a good idea to exclude interfaces attached to less critical networks so they don't affect your failover policy.

NOTE: Firepower 1010 switch ports are not eligible for interface monitoring.

Right there, under the Monitor Interfaces options, is where we set the failover trigger criteria.

Most people leave the defaults as is, but whether or not you leave the defaults in place really depends on your network! Just so you know, I've set this as high as three at a customer, but when in doubt, leave it alone.

Standby Interface IP Addresses

When configuring your HA pair, it's important to set a standby IP address and a virtual MAC address for each of the physical interfaces so switchovers will be more efficient. Nonetheless, these are still optional.

Be aware that they do take some time to configure, which normally would be fine. Thing is, if you ever break your pair and need to bring them back together, all that configuration will be completely gone, so you'll have to do them all over again, every time!

To set a standby address, just edit the interface and add an unused IP in the same subnet. You'll want to do this for every monitored interface.

NOTE: Without a standby IP address, the active unit can't perform network tests to check the standby interface health. It can only track the

link state, so clearly this can be important to do!

Monitored Interfaces

Interface Name	Active IPv4	Standby IP...	A
 Inside	10.17.117.1	10.17.117.2	
 diagnostic			
 Outside	10.11.12.170	10.11.12.172	

Active/Standby IP Addresses and MAC Addresses

For active/standby High Availability, take a look at this list about IP address and MAC address usage during a failover event:

- The active unit always uses the primary unit's IP addresses and MAC addresses.

- When the active unit fails over, the standby unit assumes the IP addresses and MAC addresses of the failed unit and begins passing traffic.
- When the failed unit comes back online, it will now be in a standby state, so it will take over the standby IP addresses and MAC addresses.

Note: The IP address and MAC address for the state link does not change at failover.

Now if the secondary unit boots without detecting the primary unit, the secondary unit becomes the active unit and uses its own MAC addresses because it doesn't know the primary unit MAC addresses. When the primary unit becomes available, the secondary (active) unit changes the MAC addresses to those of the primary unit, which can cause an interruption in your network traffic. Similarly, if you swap out the primary unit with new hardware, a new MAC address will be used.

Virtual MAC addresses guard against this snag because the active MAC addresses are known to the secondary unit at startup and they'll remain the

same in case there's new primary unit hardware. So, if you don't configure virtual MAC addresses, you might need to clear the ARP tables on connected routers to restore traffic flow. The Firepower Threat Defense device doesn't send gratuitous ARPs for static NAT addresses when the MAC address changes, so connected routers don't learn of the MAC address change for these addresses.

NOTE: The FTD device has multiple ways to configure virtual MAC addresses, but I recommend using only one method for all your interfaces, so you don't get unstable results.

Verifying with the CLI

I'll just look at my 4140 pair and go through some commands used to verify HA on the devices, although I usually just use the GUI for troubleshooting my HA pairs, but the exam really loves the CLI for troubleshooting.

Here I'll connect to the FTD image and look at the `show failover` command:

```
A4140-1 # connect module 1 console Telnet escape character is '~'. Trying 127.5.1.1...
Connected to 127.5.1.1.
Escape character is '~'.
```

```
CISCO Serial Over LAN:
Close Network Connection to Exit
```

```
Firepower-module1> connect ftd
Connecting to ftd(4140-1) console... enter exit to return to bootCLI
status(UpSys)
```

```
> show failover
```

```
Failover On
Failover unit Primary
Failover LAN Interface: 4140-HA Ethernet1/8 (up) Reconnect timeout 0:00:00
Unit Poll frequency 1 seconds, holdtime 15 seconds Interface Poll frequency 5 seconds, holdtime 25
seconds Interface Policy 1
Monitored Interfaces 3 of 1291 maximum
MAC Address Move Notification Interval not set failover replication http
Version: Ours 9.13(0)26, Mate 9.13(0)26
Serial Number: Ours FLM2048QJ1, Mate FLM2051RJ2K Last Failover at: 19:20:38 UTC Nov 16
2019
```

```
This host: Primary - Active
Active time: 11546 (sec)
```

slot 0: UCSB-B200-M3-U hw/sw rev (0.0/9.13(0)26)

status(UpSys)

Interface diagnostic (0.0.0.0): Normal (Waiting) Interface Inside (10.17.117.1): Normal (Waiting)

Interface Outside (10.11.12.170): Normal (Waiting)

slot 1: snort rev (1.0) status (up)

slot 2: diskstatus rev (1.0) status (up)

Other host: Secondary - Standby Ready

Active time: 0 (sec)

slot 0: UCSB-B200-M3-U hw/sw rev (0.0/9.13(0)26)

Interface diagnostic (0.0.0.0): Normal (Waiting) Interface Inside (0.0.0.0): Normal (Waiting)

Interface Outside (0.0.0.0): Normal (Waiting)

slot 1: snort rev (1.0) status (up)

slot 2: diskstatus rev (1.0) status (up)

[output cut]

The `show high-availability` command is the same

command as `show failover`. But the `show failover` command has all the options we want, so that's the command we'll go with:

> **show failover**

descriptor Show failover interface descriptors.

exec Show failover command execution information

history Show failover switching history

interface Show failover command interface information state Show failover internal state information

statistics Show failover command interface statistics information | Output modifiers

<cr>

The `show failover state`, `show failover statistics`,

and `show failover interface` commands serve up some good information to help you here:

> **show failover state**

State Last Failure Reason Date/Time

This host - Primary

Active None

Other host - Secondary

Standby Ready Comm Failure 19:20:55 UTC Nov 16 2019

====Configuration State====

Sync Done

====Communication State====

Mac set

> **show failover statistics**

tx:22500

rx:19824

> **show failover interface**

interface 4140-HA Ethernet1/8

System IP Address: 198.18.0.1 255.255.255.252 My IP Address : 198.18.0.1

Other IP Address : 198.18.0.2

Disabling, Suspending and Resuming HA Replication There are a few commands you can use from the CLI of an FTD device to configure HA availability, let's take a look at these:

> **configure high-availability ?**

disable Disable high-availability configuration

resume Resume temporarily suspended high-availability configuration suspend Temporarily suspend high-availability configuration

Once you break a HA FTD pair and remove a device from the FMC, you need to also use the CLI command **configure highavailability disable** before you can add the FTD to another manager, so it's important to remember this command.

If you want to stop the synchronization of a HA pair because you are troubleshooting, you can pause the synchronization with the **configure high-availability suspend** command.

To reenble the HA synchronization, you can use the **configure high-availability resume** command.

Upgrading a High Availability Pair

In this section, I'll show you how to upgrade your code in your HA pairs. Before we go there, first the FMCs need to be upgraded so the Firepower devices will then be upgraded.

Here are my 1010 and 1150 HA pairs. They look good and ready to go! The 1010s are both 6.4 code and the 1150s are 6.5 code and, now I want to upgrade everything to 6.6 code.

1010-HA High Availability							
1010-1(Primary, Active) 172.16.10.10 - Routed	FTD on Firepower 1010	6.4.0	N/A	Base, Threat (2 more...)	FMC2500-ACP		
1010-2(Secondary, Standby) 172.16.10.11 - Routed	FTD on Firepower 1010	6.4.0	N/A	Base, Threat (2 more...)	FMC2500-ACP		
1150-HA High Availability							
1150-1(Primary, Active) 172.16.10.12 - Routed	FTD on Firepower 1150	6.5.0	N/A	Base, Threat (2 more...)	FMC2500-ACP		
1150-2(Secondary, Standby) 172.16.10.13 - Routed	FTD on Firepower 1150	6.5.0	N/A	Base, Threat (2 more...)	FMC2500-ACP		

Before I can do anything, I need to get the codes needed onto the FMC. Go to **System>Updates** and then click on **Upload Updates**.



Choose your file and then click on Upload.

Product Updates

Rule Updates

Geolocation Updates

Currently running software version: **6.5.0**

Updates

updates and patches here.

Browse... No file selected.

Upload

Cancel



Verify your files in the Updates page.

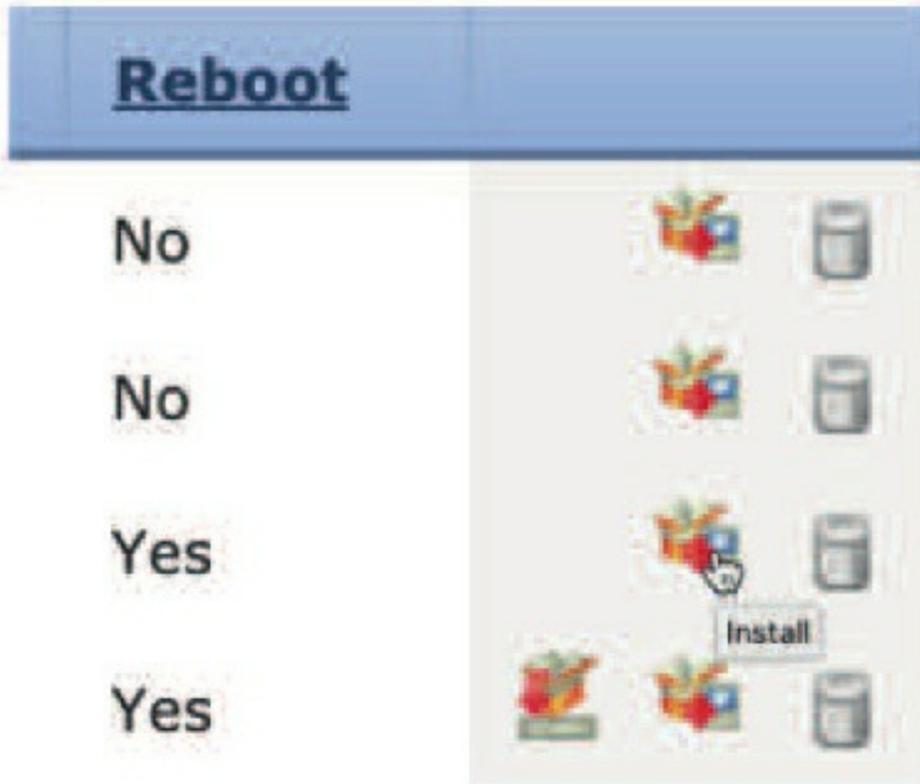
Currently running software version: 6.5.0

Updates

Type	Version	Date	Release Notes	Reboot	
Cisco Vulnerability And Fingerprint Database Updates	329	Tue Nov 12 20:05:20 UTC 2019		No	 
Cisco Vulnerability And Fingerprint Database Updates	327	Tue Aug 27 14:35:22 UTC 2019		No	 
Cisco Firepower Mgmt Center Upgrade(v6.2.1 and above)	6.6.0-35	Wed Dec 4 11:11:33 UTC 2019		Yes	 
Cisco FTD SSP FP1K Upgrade(v6.2.1 and above)	6.6.0-35	Wed Dec 4 11:50:40 UTC 2019		Yes	  

Download updates

After that, click on the icon that looks like an exploding package on the right-hand side. I'm going to click on the FMC upgrade package. You can see when I hover over it says Install in the smallest letters possible.



Now I'll get the first of my many error messages. I need to pause my synchronization with my secondary FMC. You won't need to do this with the devices:

To pause the synchronization, go to **System>Integration>High Availability**.

 Switch Peer Roles
  Break HA
  Pause Synchronization



System Status

Synchronization incomplete
 as fewer devices
 Synchronization failed on the
 Peer Management
 (Registered.)

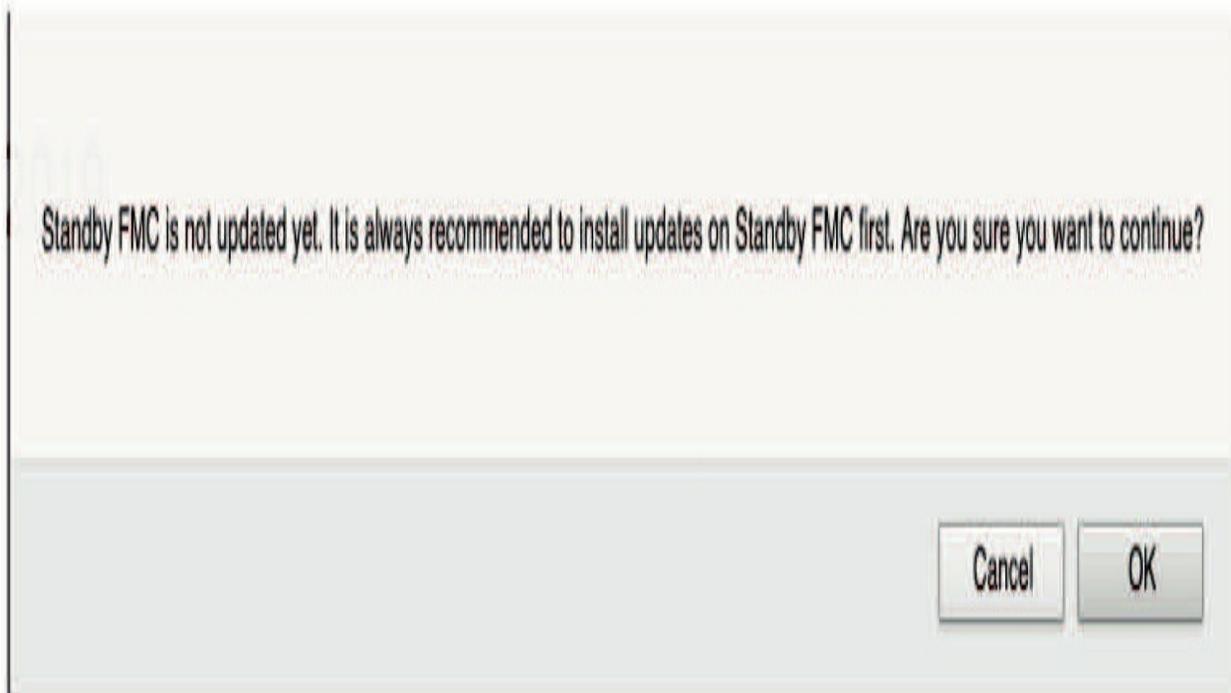
	Local	Remote
	Active - Primary (172.16.10.20)	Standby - Secondary (172.16.10.21)
Operating System	Fire Linux OS 6.5.0	Fire Linux OS 6.5.0
Software Version	6.5.0-39	6.5.0-39
Model	Cisco Firepower Management Center 2500	Cisco Firepower Management Center 2500

Sun Nov 17 15:06:46

Click on Pause Synchronization and then you'll need to click **OK** twice. Now I'll go back and try and install the package again.

Once I click on the update package, I'll choose the FMC that I want to upgrade by clicking on the check box and then choosing to install.

But wait, it's recommended that I go to the secondary box and upgrade that first...



So, I've logged into the secondary FMC and downloaded the files in the Update page. Then I clicked on the install package.

Product Updates

Currently running software version: **6.5.0**

Selected Update

Type Cisco Firepower Mgmt Center Upgrade(v6.2.1 and above)
Version 6.6.0-35
Date Wed Dec 4 11:11:33 UTC 2019
Release Notes
Reboot Yes

By Group ▾

▾ Ungrouped (1 total)

2500-2.Lammle.com
172.16.10.21 - Cisco Firepower Management Center 2500 v6.5.0

Health Policy

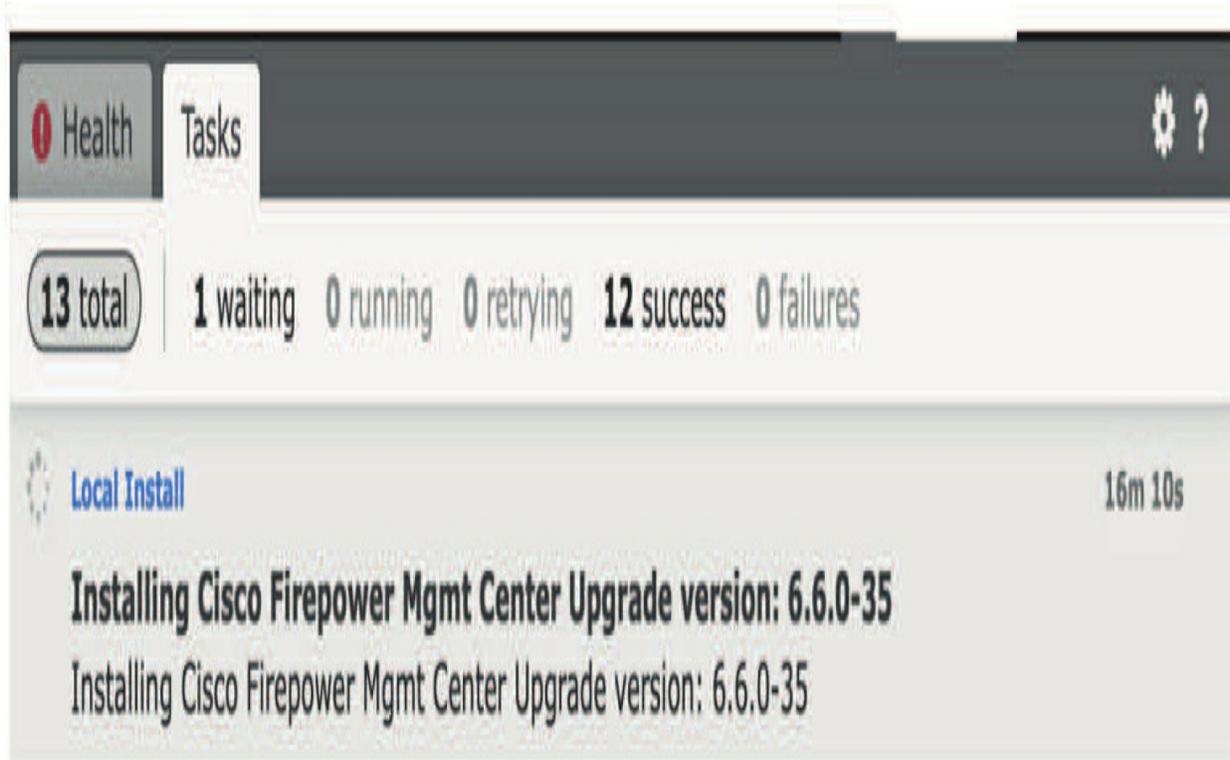
[Initial Health Policy 2019-07-29](#)
20:52:07 (Remotely authored by
[172.16.10.20](#))

Launch Readiness Check

Install

Cancel

Before it started upgrading, I received a couple warnings. I'm just going to click Close and then OK and wait an hour...



I want to point out that when this FMC is upgraded, the primary FMC still needs to be upgraded so the HA can start syncing again. Once my FMCs are upgraded, I can upgrade my FTDs.

Now that my upgrades to my FMC is finished, by going to **Overview>Dashboard**, we can get a summary of the software versions in the Status tab.

Just go to **System>Integration>High Availability** and make your primary the active FMC, which will then copy the configuration to the standby unit.

High Availability

eStreamer

Host Input Client

Smart Software Satellite



Make Me Active

This high availability pair is in split brain. Make one Management Center active by clicking 'Make Me Active'.

Warning



This operation may affect critical processes running in the background. The local peer will be active and the other peer will become a standby. The active peer will overwrite configuration and policies present on the standby peer. Do you want to continue?

Yes

No

Management Center is active
Management Center has fewer devices
Management Centers are configured to run in
Management Centers configured for high availability
Management Centers has fewer devices registered.
()

Local

Pair - Primary

(16.10.20)

Linux OS 6.6.0

6.6.0-35

Use the **Overview>Dashboard>Status** tab's Appliance Information section to get the FMC information.

Appliance Information



Name	2500-1
IPv4 Address	172.16.10.20
IPv6 Address	Disabled
Model	Cisco Firepower Management Center 2500

Versions

Software	6.6.0
Rule Update	2019-08-19-001-vrt
Geolocation Update	2019-09-03-001
VDB	329

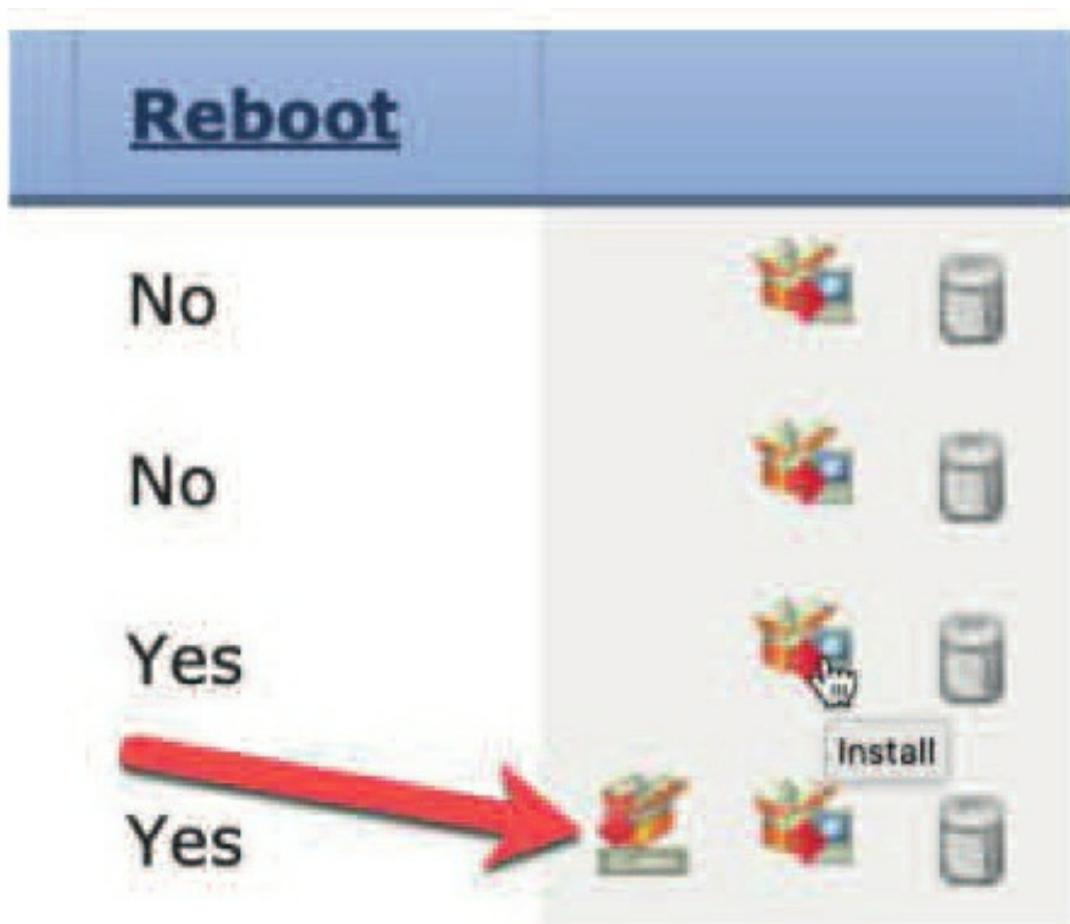
High Availability

Peer	172.16.10.21
Model	Cisco Firepower Management Center 2500
Software	6.6.0
FMC Role	Active - Primary
Status	Healthy
Detail Status	---
Last Contact	37 seconds

Upgrading a Firepower Device HA Pair

Unlike with the FMC, you don't want to break the HA pair or stop the synchronization of your devices. So nice that the Firepower install packages upgrade one device at a time for you!

From **System>Updates**, you want to push the software update to the devices before you start the install. This can even be automated in Task Management. As shown in the next screen shot, click the icon that the red arrow points to.



It's great that I have all 1000 series FTD devices because they'll all use the same software update file – sweet! So, I'll just click to update all devices.

Currently running software version: 6.6.0

Selected Update

Type Cisco FTD SSP FP1K Upgrade
Version 6.6.0-35
Date Wed Dec 4 11:50:40 UTC 2019
Release Notes
Reboot Yes

By Group ▾

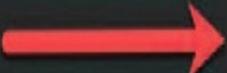
▼ Ungrouped (2 total)

▼ 1010-HA Cisco Firepower 1010 Threat Defense Cluster		
1010-1 (active) 172.16.10.10 - Cisco Firepower 1010 Threat Defense v6.4.0	Health Policy Initial_Health_Policy 2019-07-29 20:52:07	✓ ✓
1010-2 172.16.10.11 - Cisco Firepower 1010 Threat Defense v6.4.0	Health Policy Initial_Health_Policy 2019-07-29 20:52:07	✓ ✓
▼ 1150-HA Cisco Firepower 1150 Threat Defense Cluster		
1150-1 (active) 172.16.10.12 - Cisco Firepower 1150 Threat Defense v6.5.0	Health Policy Initial_Health_Policy 2019-07-29 20:52:07	✓ ✓
1150-2 172.16.10.13 - Cisco Firepower 1150 Threat Defense v6.5.0	Health Policy Initial_Health_Policy 2019-07-29 20:52:07	✓ ✓

→ Push Cancel

Remember, this will update the secondary for each pair first, then make the primary a secondary and update that device too. Then it will finally make the device primary again, all without losing connection!

Click Push and then wait for just a few minutes. Click on the green check mark to get status of the push.

Deploy 

Deployments Health Tasks

17 total | 0 waiting 1 running 0 retrying 16 success 0 failures

 Update Push
Push to 1150-2
Pushing upgrade...

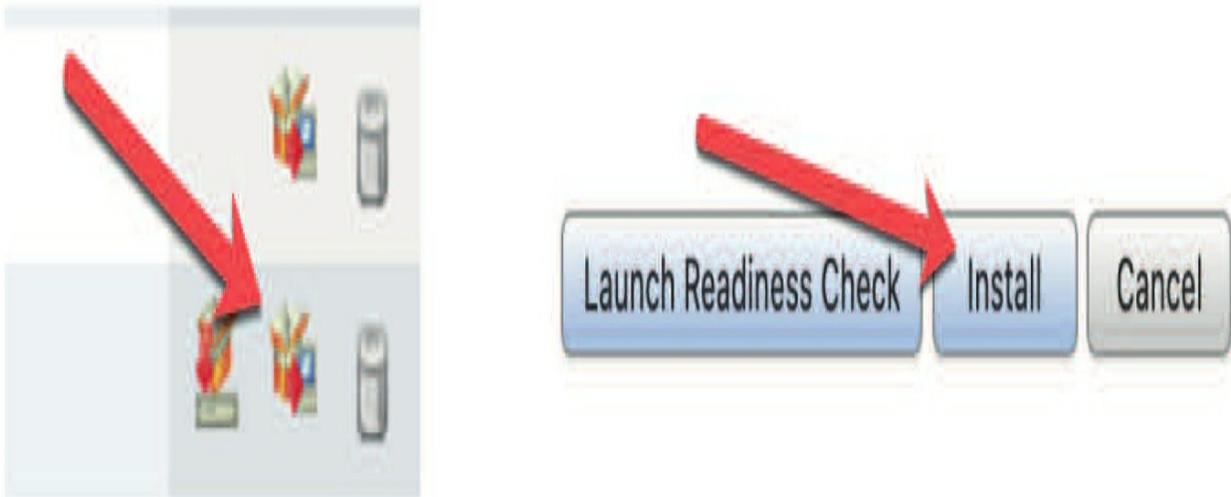
 Update Push
Push to 1150-1
Complete

 Update Push
Push to 1010-2
Complete

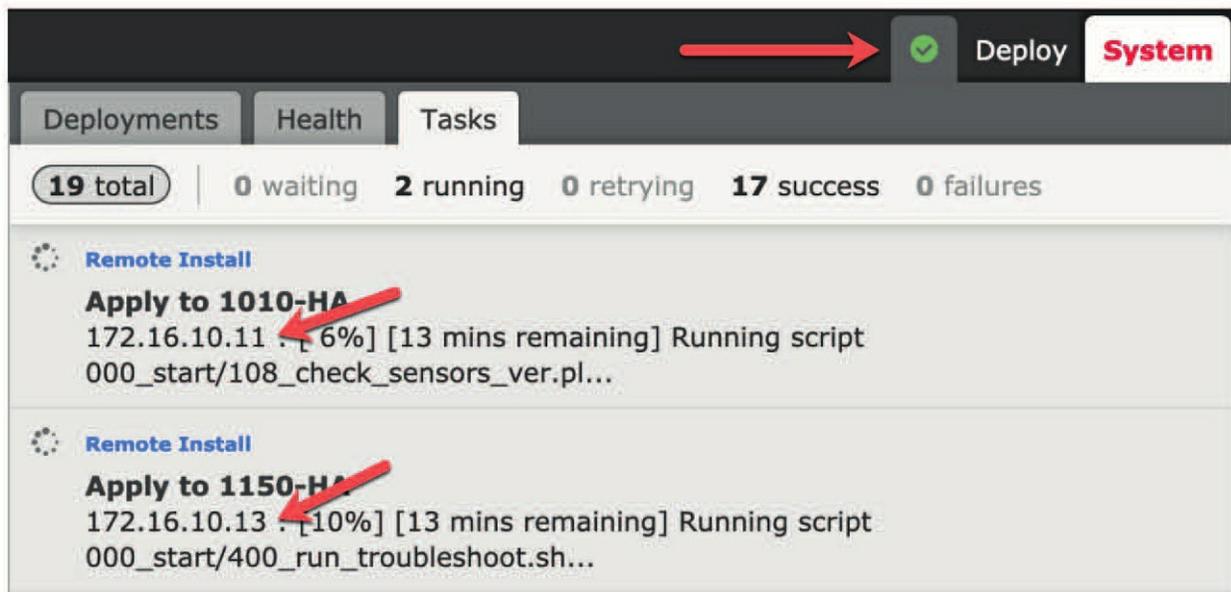
 Update Push
Push to 1010-1
Complete

When that's done, you can start the install again. I'll do this for all devices here. Again, be assured that I absolutely will *not* bring down the network when I do this upgrade.

Now I'll click on the exploding package, choose all devices, and then click on Install.



Click on the green check mark again to get status...



See the IP addresses of the devices? These are the secondary units' IPs. Again, this process will update the secondary, then switch to be primary, and

then the original primary will be updated to finally become primary again.

Now I'm going to verify that no interruptions take place. Time for me to play guitar and watch some Internet TV to verify that everything keeps working.... this process will take a bit...

Nice—okay, no interruptions! It took about 30 minutes to upgrade the secondaries to this point. So, after another show and another 30 or so minutes, I checked my devices and found that they are all now the latest code! And I never lost the connection on my TV



1010-HA High Availability			
✓ 1010-1(Primary, Active) 172.16.10.10 - Routed	FTD on Firepower 1010	6.6.0	N/A
✓ 1010-2(Secondary, Standby) 172.16.10.11 - Routed	FTD on Firepower 1010	6.6.0	N/A
1150-HA High Availability			
✓ 1150-1(Primary, Active) 172.16.10.12 - Routed	FTD on Firepower 1150	6.6.0	N/A
✓ 1150-2(Secondary, Standby) 172.16.10.13 - Routed	FTD on Firepower 1150	6.6.0	N/A

Finally, right there in the **Dashboard>Status** tab, you can see that the FMCs and all the devices are running the latest code. Sweet, and not a single outage!

Type	Current	Latest
Geolocation Update		
Local Geolocation Update	2019-09-03-001	2019-11-12-002.sh.REL
Rule Update		
Local Rule Update	2019-08-19-001-vrt	Unknown
Software		
1 Management Center	6.6.0	6.6.0
4 Devices	6.6.0	6.6.0
VDB		
1 Management Center	329	329



Summary

We picked up from Chapter 1, where I configured the FMCs, and also Chapter 6, where the Firepower devices were all configured, and added High Availability into the mix. If you're like me, you don't want to lose your configurations at this point in the event of a failure!

You learned when you can perform the HA on the Firepower devices, which for the FMC, is basically anytime, even though you'll lose connection for a bit while they reboot. You discovered that this doesn't bring your network down, but you'll still lose malware SHA checks and AD integration for a very short period of time. Most people do this in a maintenance window.

As for the Firepower and FTD devices, once you have an HA FTD pair, you can perform upgrades pretty much anytime because the upgrade is performed on one device at a time automatically. Most people do the updates in a maintenance window—just in case.

I began with the hardware 2500 FMC HA and then configured HA on the 1150s and 1010s. After that, I demonstrated upgrading the 2500 FMCs to a new code. We wrapped the chapter by upgrading the FTD devices, all while keeping my production network running—and my TV!

Chapter 8: Objects

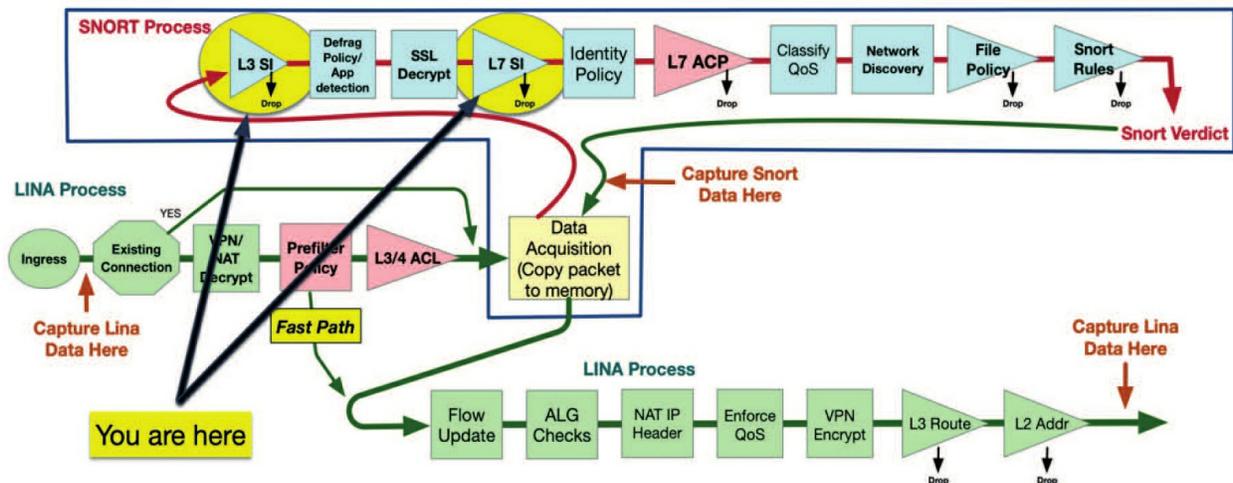
The following CCIE/CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.4 Configure objects using Firepower Management Center

2.4.a Object Management

2.4.b Intrusion Rules



Welcome to objects! We are finally getting into the meat of the Firepower system, and I really mean that because this is a super long chapter!

Firepower uses reusable configuration components—objects—to provide an easier way to use values across policies, searches, reports, dashboards, etc. You may think a chapter describing these components would be boring. However, understanding the objects themselves will require an understanding of how they are used within the rest of the system. Because of this, we will wind up digging into several interesting areas in our discussion of this subject.

We're going to go down the list of object types as shown in the user interface. Some descriptions, for well-known object types, will be rather brief, but for other object types a more detailed discussion is in order.

Toward the end of this chapter we will hit several object types that apply only

to Firepower Threat Defense and that relate to some of the legacy ASA routing and VPN features.

Finally, I'll discuss the Intrusion Rules tab found in the Object Manager.

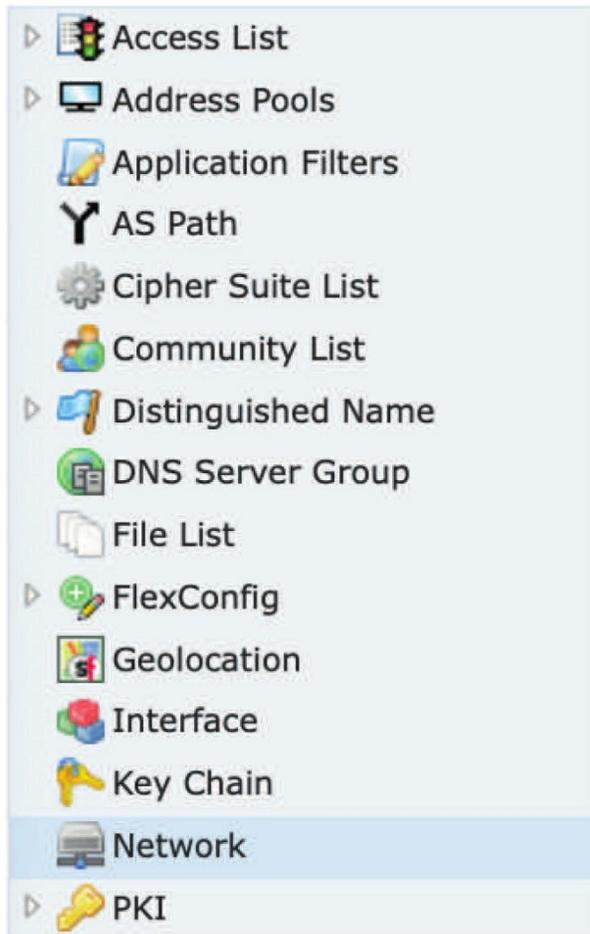
Objects

When you click on the Objects tab at the top of the FMC GUI, this will take you to Object Management, where you'll find all the available types of objects you can manage.

This first figure shows the Objects tab and the Object Management screen you're taken to. Take note of the Intrusion Rules tab as we'll discuss that at the end of this chapter.



On the next figures, you'll find the objects are displayed in alphabetical order. Firepower codes before 6.4 has them in usage order, and that is the order I'll follow in this chapter, not alphabetical.



The Object Management page is somewhat long, so I am showing the available objects side by side as well as compressed, because they are all expanded by default, which makes this screen even longer and harder to read.

I always try to compress the Objects listings by clicking on the little black arrow



on the left; the list will then compress, and the arrow will look

as so:



This makes the objects easier to manage. However, if you use the new Light Theme, they are now automatically compressed. Nice.

Network

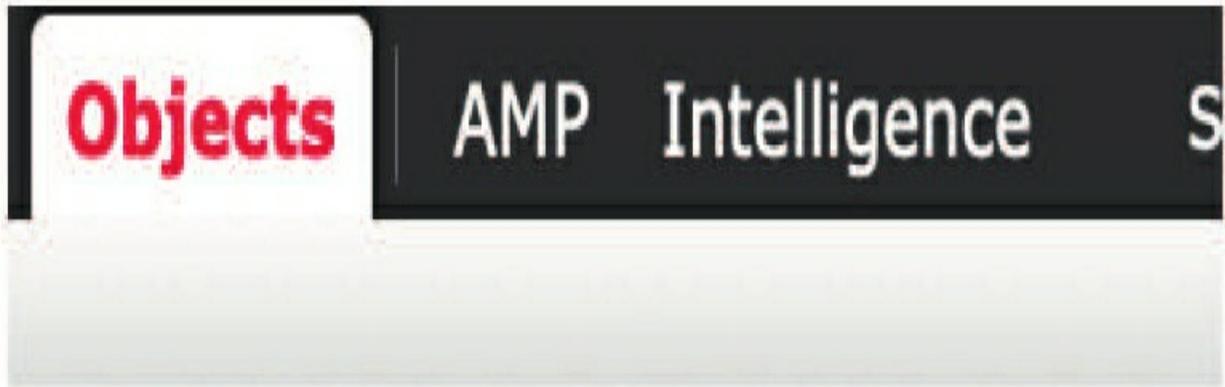
Although the objects are now in alphabetical order starting in 6.5 code, the screen showing the Network objects is still the default screen, and as you can see in the previous figure, Network is highlighted by default.

Network objects are simply IP addresses. They can be single IPs, CIDR blocks, and even ranges. There are quite a number of default Network objects already present.

These will be useful in many deployments to identify ranges such as RFC-1918 private addresses or multicast addresses.

To create your own Network object, click the Add Network button. You can then select Add Object or Add Group. A group is simply a number of objects bundled together.

The Add Network Object dialog is shown here:



Once you click Add Object, the next splash screen shown here opens up.

In the Network field, you can enter an IP address in several formats:

New Network Object



Name:

Description:

Network: Host Range Network FQDN

Allow Overrides:

Save Cancel

- Single IP host address 10.0.1.2
- An arbitrary range such as 10.0.0.8-10.0.0.12
- Network address such as 192.168.20.0/24
- FQDN

Let's discuss the last option listed here and then Allow Overrides, as the rest are self-explanatory. **FQDNs** must begin and end with a digit or letter, and only letters, digits, and hyphens are allowed as internal characters, no spaces.

The FQDN objects can only be used in access control rules and prefilter rules. The rules match the IP address obtained for the FQDN by doing a DNS lookup. This means you must have configured the DNS server settings in DNS Server Group Objects and the DNS platform settings in Configure DNS.

The **Allow Overrides** box also deserves some explanation. Checking this allows overriding this object's value at the device level if desired. When you check the Allow Overrides box, another section of the screen appears. You can then select Devices or Domains, add one or more, and select a different value for this object when used in the selected context.

As an example, the object "bob" has the value of 10.0.0.1 except when it's used on the device named "FTDv 6.1 Routed," where the value is 10.0.0.2.

New Network Objects ? X

Name:

Description:

Network:
Format: ipaddr or ipaddr/len or range (ipaddr-ipaddr)

Allow Overrides:

Override (1)			
Override On	Content	Type	
FTDv 6.1 Routed	10.0.0.2	Host	

Save Cancel

It unlikely you want to provide a secondary address to an object since you need to then add the secondary variable to a physical FTD device, although obviously you can, but just be prepared for a complex troubleshooting routine when used.

Last, as mentioned, a Network group is a collection of one or more objects. Use the Add Group button to create a group and add Network objects to it.

You can nest object groups within object groups—up to 10 levels!

Port

Like Network objects, Port objects are a pretty basic concept. These are most commonly the TCP/UDP port numbers that live at the Transport layer of the OSI model. Port objects can also be ICMP or IPv6-ICMP Type/Code combinations or any other IP protocol number.

Clicking on the Port link in the upper left of the screen brings you to the list of default Port objects. All the common TCP/UDP ports are here. These default objects cannot be edited or deleted. To create a new Port object, click the Add Port button and select Add Object.

This figure shows the New Port Objects dialog.

The image shows a dialog box titled "New Port Objects" with a red title bar. The dialog contains the following fields and controls:

- Name:** A text input field with a red border.
- Protocol:** A group of radio buttons for "TCP", "UDP", "ICMP", "IPv6-ICMP", and "Other". The "TCP" radio button is selected.
- Other:** A dropdown menu currently displaying "All".
- Port:** A text input field.
- Allow Overrides:** An unchecked checkbox.
- Buttons:** "Save" and "Cancel" buttons at the bottom right.

Again, rather than discuss the mundane steps to create a Port object, I will

just note one bit of flexibility offered in Firepower.

In the Port field, you can enter a single port in the range 1 to 68838, or you can enter an arbitrary range such as 28–38. You can also create Port Group objects, which are a combination of one or more Port objects.

There are a couple of caveats with Port objects:

1. You cannot add any protocol other than TCP or UDP for source port conditions in access control rules.
2. You cannot mix transport protocols when setting both source and destination port conditions in a rule.
3. If you create a Port object containing both TCP and UDP ports, then add it as a source port condition in a rule, you cannot add a destination port, and vice versa.

The caveats sound more complicated than they are, but I'll explain and demonstrate these rules in the chapter on Access Control policy chapter, which is the next chapter.

Interface

There are two types of Interface objects—security zones and interface groups. Each interface must be assigned to a security zone and/or interface group. You then apply your security policy based on zones or groups.

A *security zone* is a descriptive name for an interface and is different than security levels. Security levels on your ASAs were set high to low, where the high could go out the lower level, but the lower level couldn't go out the higher level by default. Zones are more like a fence between properties, where every interface is equal and then you create rules to say who can go over the fence and what they can do when they get to that other side.

There is a slight difference between FTD and a Firepower appliance device here. In FTD your interface has a logical name, which is assigned through the **Device Management>Interfaces** page. These interfaces can also have a security zone.

However, with a Firepower device you do not have this capability, so the

security zone becomes your only friendly name for the interface. The main benefit of using zones in your appliance is that you get a friendly or descriptive name for the interface. So, if you are connecting a firewall and a switch through an inline interface pair, you can name the zone on the one side “Firewall” and the zone on the other side “Switch.” This friendly name shows up in event table views such as those for connection or intrusion events adding valuable context for the analyst.

A security zone can also contain more than one interface. This way, similar interfaces could have the same security zone assigned. This could be handy if, for example, your Access Control policy contains a rule with the source security zone set to “Firewall.”

By using this same security zone across devices, this rule will match traffic coming from the Firewall anywhere the policy is applied. You also can be specific if desired and use names like NY-Firewall, MSP-Firewall, etc. This way, when reviewing events, the Ingress and Egress zones will provide the analyst with more source/destination context surrounding the event.

By using the new Interface Group feature we can even go a step further in providing flexibility and ease of management. While a security zone can contain multiple interfaces, an interface can belong to just one security zone.

However, an interface can belong to multiple interface groups. In the case of multiple sites with similar network designs, you can have an interface group that contains all your firewall interfaces across the organization. You can also create an interface group with all your interfaces in a site.

You can use interface groups in FTD NAT policies, Prefilter policies, and QoS policies.

In our example above, the MSP-Firewall interface could then be a member of the worldwide “Firewall” group as well as the “MSPInterfaces” group.

Let’s walk through an example. First, we will navigate to **Objects>Interface**, click the Add button, and select Security Zone. This is shown here:



You will then need to select the Security Zone type as shown here:

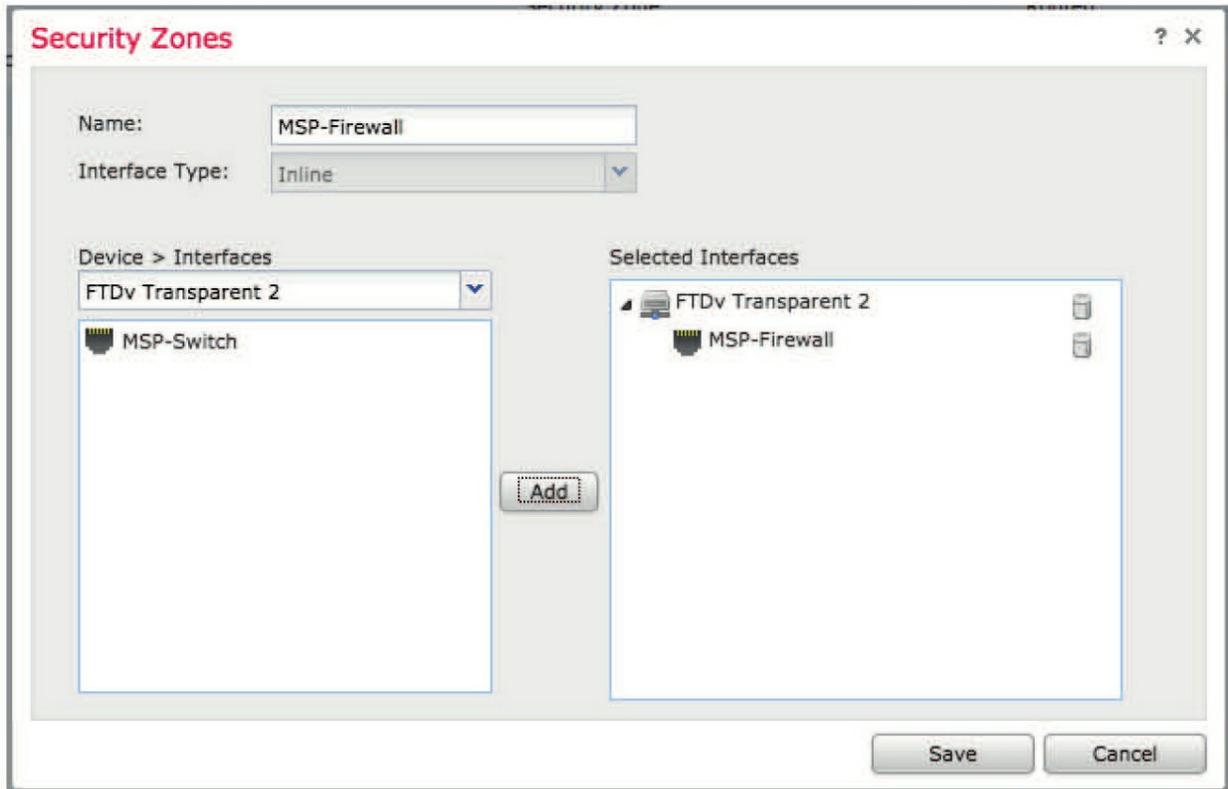
Security Zones

The screenshot shows a web-based configuration interface for Security Zones. It includes a 'Name' text input field, an 'Interface Type' dropdown menu with a red border and a blue arrow, and a list of available interface types: Passive, Inline (highlighted in blue), Switched, Routed, and ASA. Below the dropdown is a 'Select Device' button and a table for 'Available Interfaces'. To the right is a table for 'Selected Interfaces'. An 'Add' button is positioned at the bottom right of the interface.

When adding a security zone, all interfaces in this zone must be of the same type. Your choices are Passive, Inline, Switched, Routed, and ASA.

Once you have the type selected, choose your device from the **Device > Interfaces** drop-down. At this point, any interfaces not already members of another zone will appear. In the example below, I am creating an inline zone I'll name Inside.

I've selected the FTDv Transparent 2 device and there are two interfaces I can choose from: MSP-Firewall and MSP-Switch. I have added the MSP-Firewall interface to this zone as shown here.

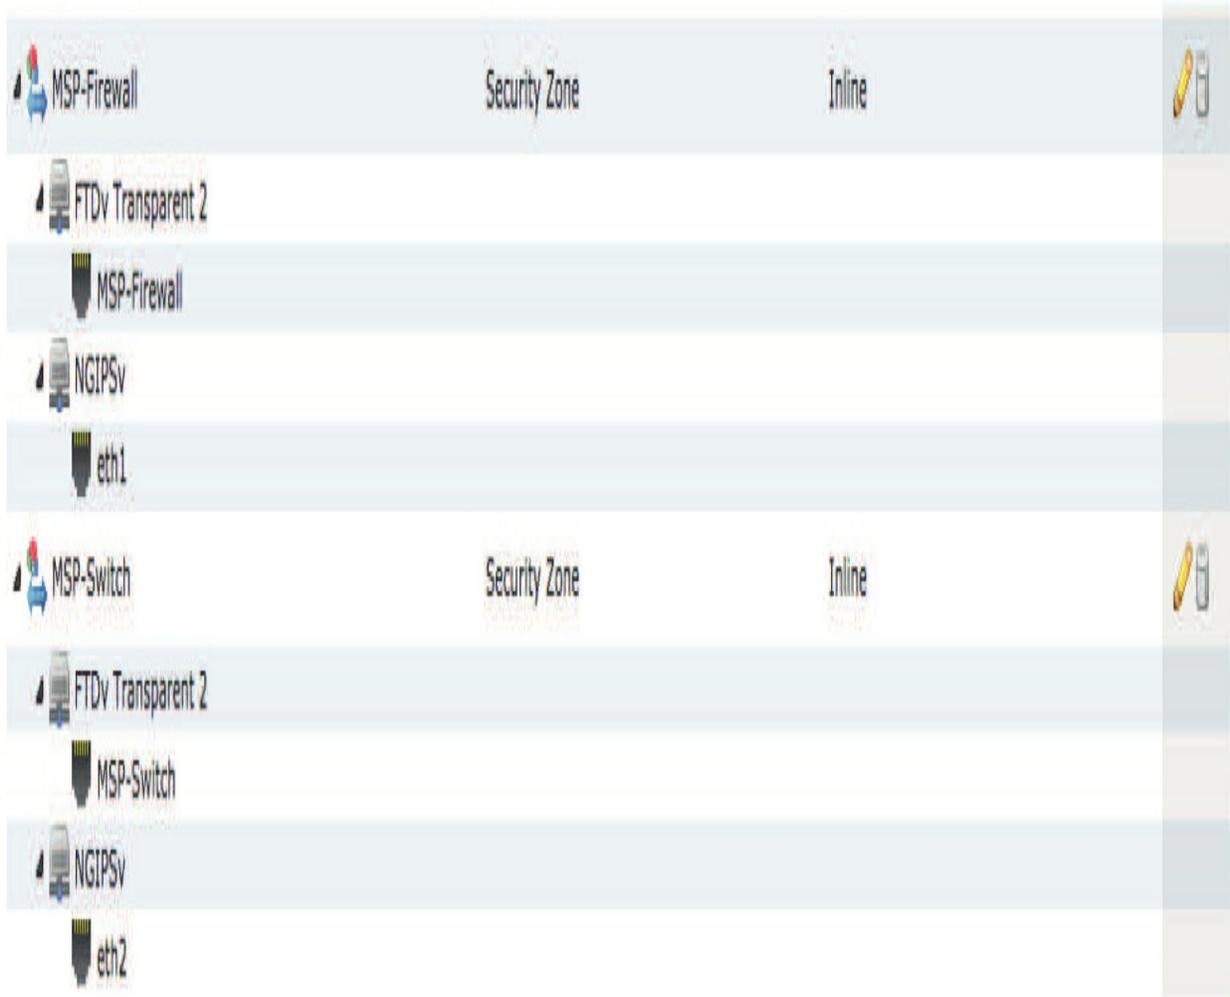


Now, if you have other inline firewall-facing interfaces in the MSP location, you can add them to this same zone. They can even be on different devices.

In the example below, I have created the MSP-Firewall and MSPSwitch zones and added inline interfaces from an FTD and a NGIPSv device. Notice the logical names for the FTD device and the systemassigned name for the Firepower NGIPSv.

Remember, this is because you can name the actual interface on

FTD but on a Firepower device you will use the slot/port notation such as s1p1, s1p2, etc.



This next figure shows an output of FMC interface zones for both Firepower appliances using Inline mode and FTD devices configured as routers.

Name	Type	Interface Type
External	Security Zone	Inline
Inside	Security Zone	Routed
Internal	Security Zone	Inline
Outside	Security Zone	Routed
FTD19		
FTD20		

Tunnel Zone

The Tunnel Zone object is used in conjunction with the Prefilter policy.

Tunnel zones can be assigned to tunneled traffic in Prefilter and then referred to in access control rules. This will be discussed in more detail in the chapter on the Prefilter policy.

Tunnel Zone objects are simply arbitrary names you create to assign to tunneled traffic.

Application Filters

One of the benefits of a next-generation firewall is application awareness. At the time of this writing, Firepower can identify over 3,800 applications. These are further categorized according to risk, business relevance, and type. In addition, applications are assigned tags such as evasive, blog, webmail, etc. The purpose of an Application Filter object is to allow you to define your own criteria to allow or block applications in your environment. This can be done in an access control rule without using a pre-created Application Filter object, but if you need to reuse these filters across multiple policies, then creating an object is the way to go.

When you click the Add Application Filter button, you are presented with the dialog shown here:

Application Filter

? X

Name:

Application Filters

Search by name

Risks (Any Selected)

<input type="checkbox"/>	Very Low	1225
<input type="checkbox"/>	Low	877
<input type="checkbox"/>	Medium	1005
<input type="checkbox"/>	High	245
<input type="checkbox"/>	Very High	170

Business Relevance (Any Selected)

<input type="checkbox"/>	Very Low	960
<input type="checkbox"/>	Low	629
<input type="checkbox"/>	Medium	1303
<input type="checkbox"/>	High	346
<input type="checkbox"/>	Very High	284

Types (Any Selected)

Available Applications (3522)

Search by name

- 050plus
- 1&1 Internet
- 1-800-Flowers
- 1000mercis
- 100Bao
- 100ye.com
- 12306.cn
- 123Movies
- 126.com
- 17173.com
- 1fichier

Add to Rule

Viewing 1-100 of 3522

Selected Applications and Filters (0)

any

Save

Cancel

By checking the appropriate boxes under Risks, Business Relevance, Types, etc., you narrow down the list of available applications in the center column. You can also use the Search by Name field in either the left or center columns to zero in on the apps you want. Once you have what you're looking for, click the Add to Rule button to add them.

The mechanics of the user interface are easy enough to figure out. However, there are some important points you should keep in mind as you're setting up these filters.

First, the blue information icon (



) is your friend.

Next, I'll click on the icon next to LinkedIn Job Search. Here you can see how this application is classified according to risk, business relevance, type, etc. Notice this is also tagged as "not work related."

Available Applications (6)

Selected Applications and

 linked 

-  All apps matching the filter

- All apps matching the filter 
- LinkedIn Contacts 
-  LinkedIn Inbox 
- LinkedIn Job Search 
-  LinkedIn Profile 
- LinkedIn Upload 

any

LinkedIn Job Search 

The job search facility on LinkedIn.

Risks: Low
Business Relevance: Very Low
Types: Web Application
Categories: business, social networking
Tags: not work related

 [Wikipedia](#),  [Google](#),  [Yahoo!](#),  [Bing](#)

Also, notice the lock icon (



). This indicates applications that the

system can only detect if the traffic is decrypted. Keep in mind that this is not all-inclusive. You will find applications that do not have this icon but may still be encrypted. With more and more applications and websites defaulting to SSL, decryption is becoming an even more important tool for network traffic inspection.

Finally, if you have been living on a deserted island for the past 10 years (please tell me how you did that!), you can use one of the search engines links to find out what this LinkedIn application is all about. Actually, there may be some applications in the list that are not as well-known, so these search links do come in handy from time to time.

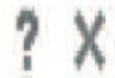
When you are finished, you will have application filter(s) that you can then use in your access control rules to allow or restrict access for your users.

VLAN Tag

The VLAN Tag object is simply a VLAN number or range. You can set these up ahead of time for use in your access control rules. You can create VLAN Tag objects as well as groups.

Here is an example of the New VLAN Tag Objects dialog.

New VLAN Tag Objects



Name:

Description:

VLAN Tag:

Allow Overrides:

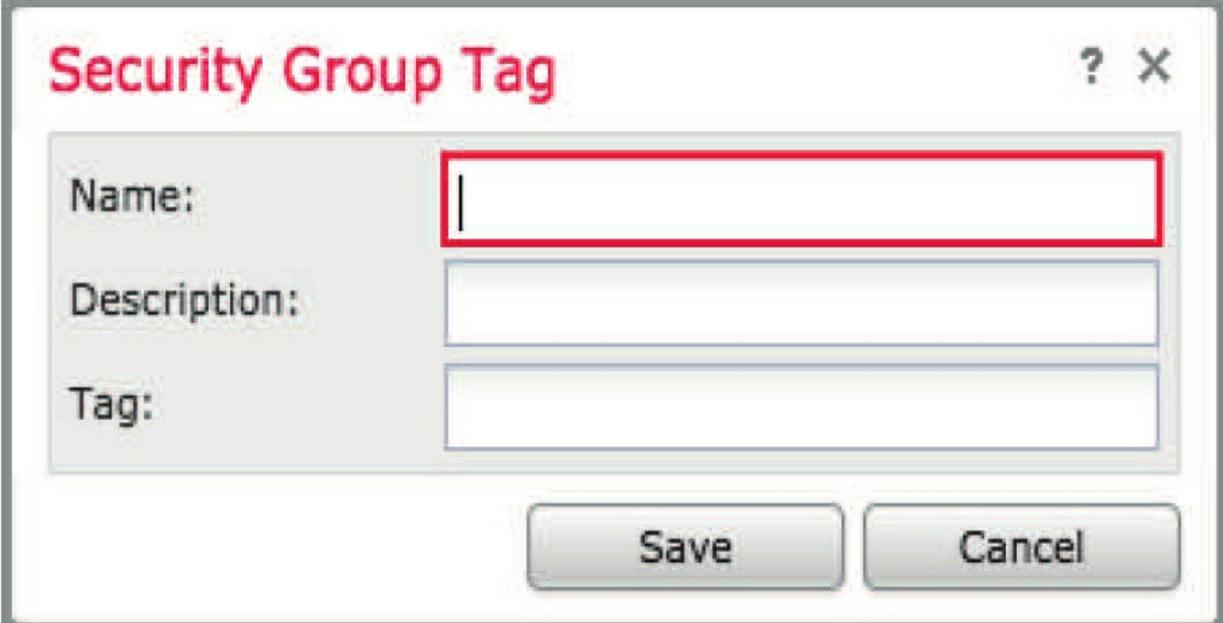
Save

Cancel

Security Group Tag

The Security Group Tag (SGT) object is used in conjunction with Cisco's Identity Services Engine (ISE) or TrustSec. The SGT is a field added to network packets that is used to identify the source machine or user. This object is used if you have not connected your FMC to an ISE or pxGrid server. It allows you to specify a tag based on your knowledge of what SGTs will be passing through your devices. Because of this, it is only available if you have *not* configured ISE/ pxGrid as an identity source.

If you have configured one of these identity sources, then the available SGTs are pulled from the ISE/pxGrid server rather than using custom SGT objects.



The image shows a dialog box titled "Security Group Tag" with a red title bar. It contains three input fields: "Name:" (highlighted with a red border), "Description:", and "Tag:". At the bottom, there are "Save" and "Cancel" buttons.

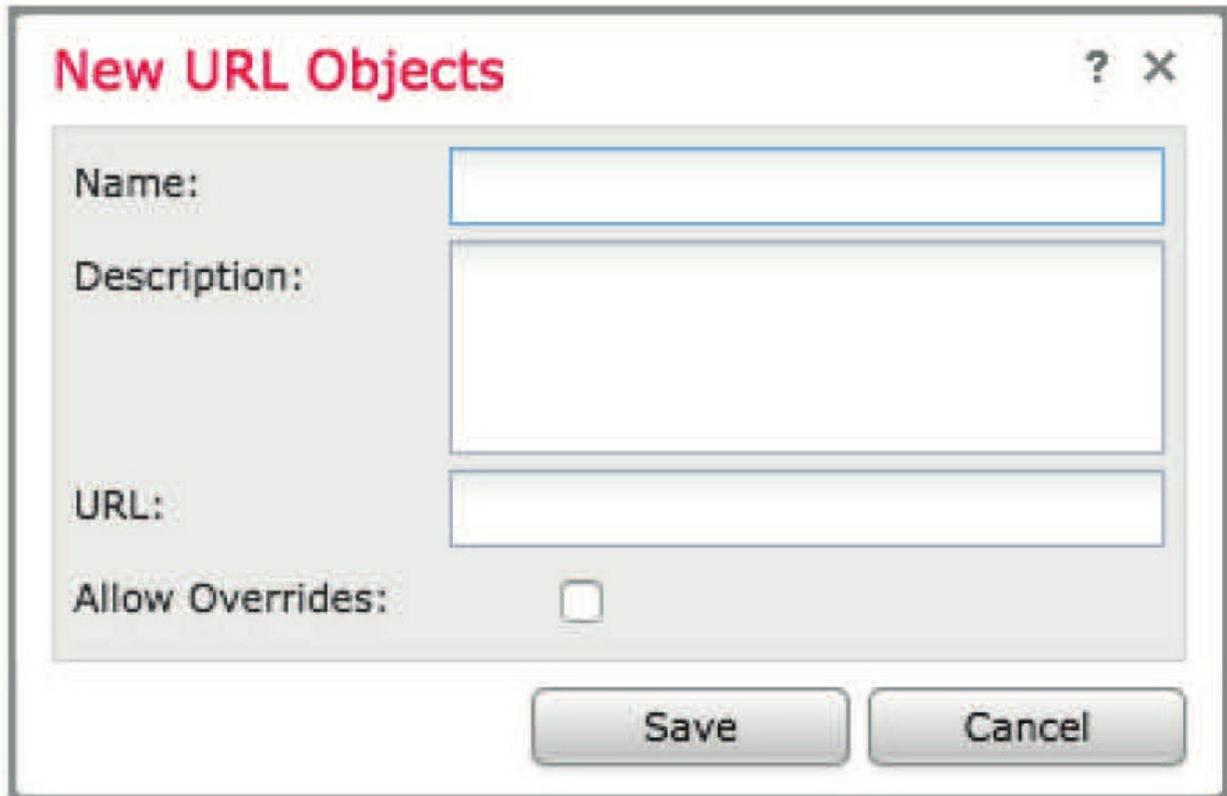
Yes, the identity subject is quite involved, and for now, I will limit the discussion to just the use of the Security Group Tag object. Creating an object is very simple; just give it a name, an optional description, and the tag.

URL

The URL object is another pretty simple concept. You give a URL a friendly name you can then reuse in various access control rules. The process for creating these objects is straightforward.

There are a couple of things you should remember when creating URL

objects. First, if you plan to use the object to match HTTPS traffic, then create the object based on the common name in the public key certificate used by the server. Also, the system disregards subdomains within the common name. For example, use **badsite.com** rather than **www.badsite.com**. Second, the system performs a substring match on the URL.



The image shows a dialog box titled "New URL Objects" with a standard window title bar (question mark and close button). The dialog contains the following fields and controls:

- Name:** A single-line text input field.
- Description:** A multi-line text input field.
- URL:** A single-line text input field.
- Allow Overrides:** A checkbox, currently unchecked.
- Buttons:** "Save" and "Cancel" buttons at the bottom right.

For example, if you create an object to match **ign.com**, it will match *any* URL that contains this text. This means you will also match **verisign.com**. This can yield unexpected results if you are not careful.

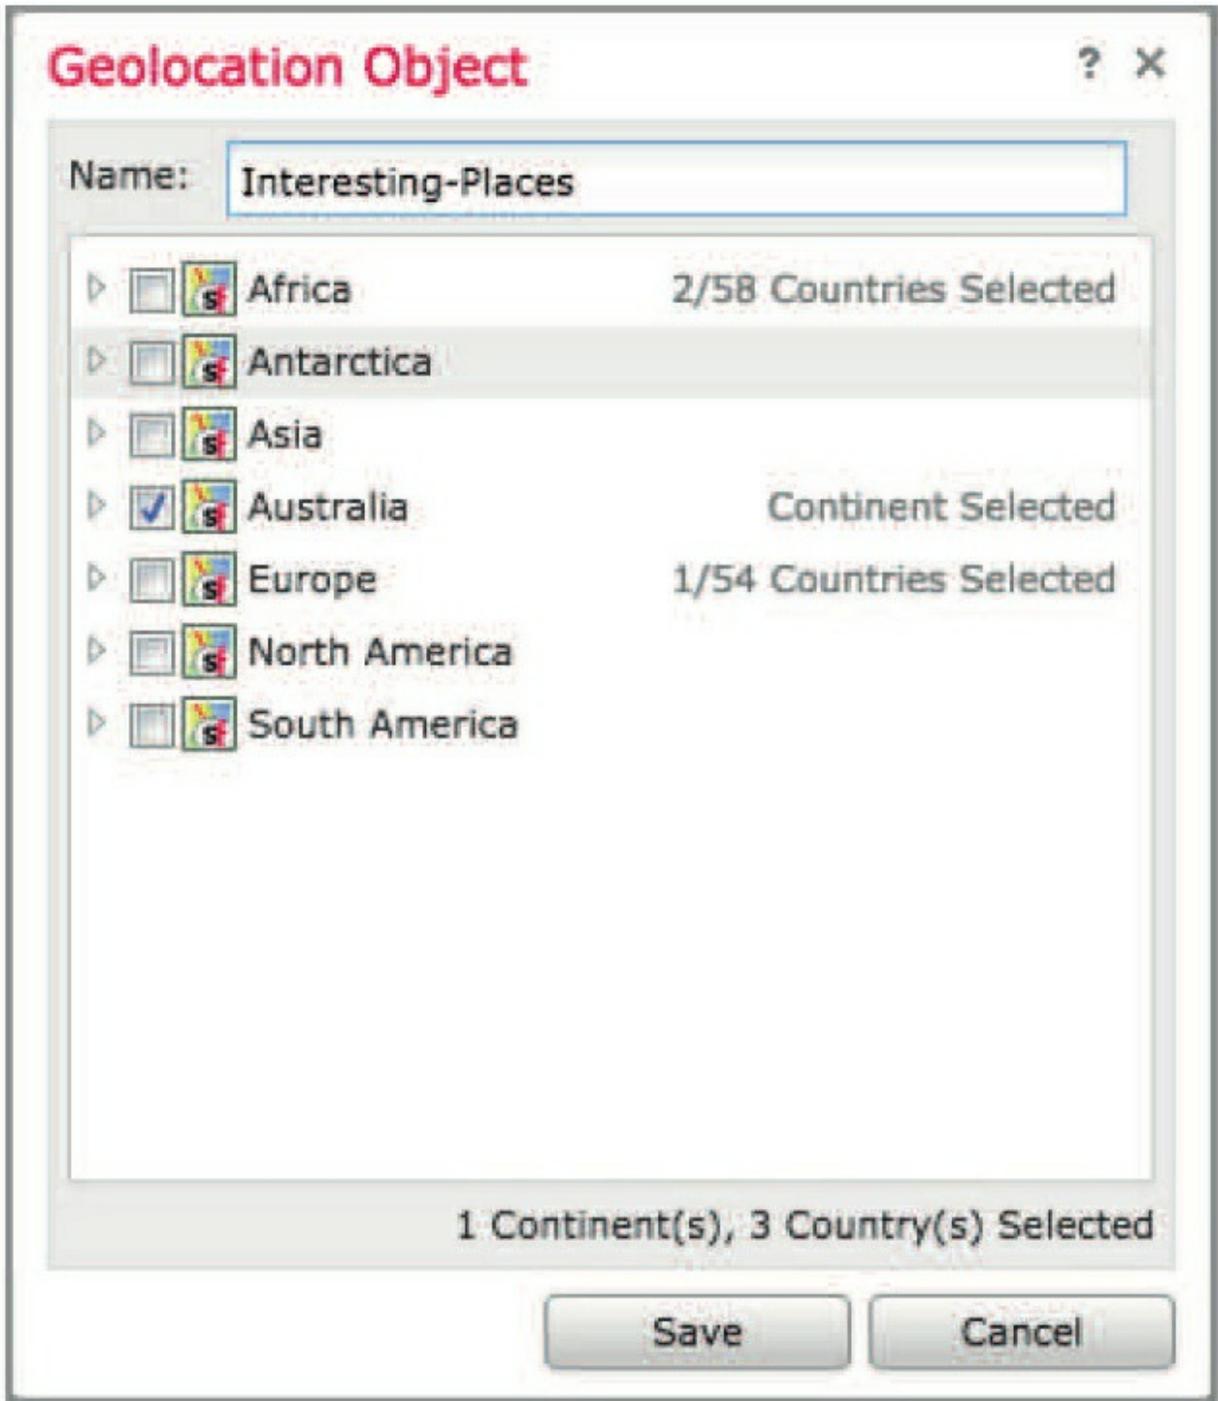
The fact that the system ignores subdomains for HTTPS coupled with the substring matching behavior can provide some challenges in filtering exactly the URLs you are looking for. My advice is to be thoughtful and test extensively!

Geolocation

The Geolocation object allows for defining country and continent constraints

for use in your access control rules. This is based on an IP address-to-geolocation mapping database that is updated periodically on your FMC using **System>Updates**.

Creating a new object is a simple process of clicking on the Add Geolocation button, giving your object a name, then checking



the boxes by the continents or countries you want to include. This object will then be available for you to use as you wish in access control rules.

Unlike Network and Port objects, you cannot create a Geolocation object

inside the Access Control policy (ACP).

Variable Set

The Variable Set object relates directly to the use of Snort rules in the Intrusion policy. Each Snort rule contains two parts—the header and the body. The rule header determines things like the IP protocol, source/destination ports, and source/destination IP addresses to which the rule applies.

The body is where the detection keywords are located. These do the magic of inspecting the packet contents looking for... well, whatever the rule writer wants to find. We will look at rules in more detail in the chapter on Intrusion policy. For now, let's examine a simple rule header to get a basic understanding of what these variables are all about.

Consider the following Snort rule header:

```
alert tcp $EXTERNAL_NET any -> $HOME_NET
$HTTP_PORTS
```

Let's break down the various keywords used in this rule header.

alert

This is the rule action; in this case the rule will generate an event (alert) but will not drop the packet.

tcp

The protocol, can be ip, tcp, or udp.

\$EXTERNAL_NET

The source IP address(es).

any

The source port.

->

The directional operator; this means the packet must be traveling from the source (on the left) to the destination (on the right).

\$HOME_NET

The destination IP address(es).

\$HTTP_PORTS

The destination port(s).

As you probably guessed, some of the items above are variables. See them? They are the ones that start with the dollar sign (\$). These are defined by the Variable Set object.

If you only remember one thing about this section, remember this: Many Snort rules contain the variables `$HOME_NET` and `$EXTERNAL_NET`. This is because you are typically looking at traffic either to or from your “protected network.” Sometimes Snort rules are designed to detect an attack in a packet traveling to your protected network. Sometimes they are designed to detect the evidence of an attack, data leakage, or maybe signs of a malware infected host in a packet traveling from your protected network.

As you probably guessed by now, the variable `$HOME_NET` contains the IP address range(s) of your protected network. Oftentimes we include all the RFC-1918 private address space plus any public IP space owned by your organization.

Conversely, the `$EXTERNAL_NET` variable points to the unprotected network or maybe the “suspect” network if you would rather. This is commonly the Internet, but it may also include your `$HOME_NET` as well. You do not have to exclude the `$HOME_NET` range from the `$EXTERNAL_NET` variable. In fact, in many cases, I recommend you leave `$EXTERNAL_NET` at the default value of “any.” By doing this, you can protect yourself from any internal hosts that may misbehave and start spreading mischief.

One last thing to keep in mind. The default value of both `$HOME_NET` and `$EXTERNAL_NET` is “any.” In addition, nearly all of the other IP variables default to the value of `$HOME_NET`. Some of these are variables like the following:

- `DNS_SERVERS`
- `HTTP_SERVERS`
- `SIP_SERVERS`
- `SMTP_SERVERS`
- `SQL_SERVERS`

All of these default to the value of `$HOME_NET`. They are provided to allow for further narrowing down an IP range for a specific rule type if desired. For example, you will find that rules written to protect web servers contain the variable `$HTTP_SERVERS` as their destination IP in the rule header. I find that in

the vast majority of installations, there is no need to refine these IP variables. My advice—leave them alone. Define `$HOME_NET` with your internal IP address ranges, leave `$EXTERNAL_NET` at “any,” and go from there. You will be getting the best detection and also adding efficiency by inspecting traffic in the direction each Snort rule was designed to operate.

Another option is to leave both `$HOME_NET` and `$EXTERNAL_NET` to “any.” One problem with this is now you are forcing Snort to inspect traffic without regard to its source or destination IP. Rules designed to inspect inbound traffic will be forced to inspect outbound traffic as well. This will reduce the efficiency of the system and increase false positive intrusion events. However, in some cases when inspecting East-West (internal) traffic, you may have little choice but to leave these defaults unchanged. Overall, it’s usually better to cover more than you need than to run the risk of leaving some assets unprotected.

Now that you have an understanding of variables and their importance in Snort, let’s look at the Variable Set object. First, you will notice that there is a default set (named Default-Set) already created. This is the system-provided variable set. If your deployment is small or you do not need to specify different variables for different devices in your network, then you can simply make any changes to this Default-Set. It will then be used automatically by any access control rules where you implement an Intrusion policy.

Your other option is to create your own Variable Set objects and customize them for use in various access control rules. These new Variable Set objects will start off using the values in the Default-Set but you can override any or all of them to create your custom object.

Let’s take a look at the Default-Set provided with the system.

Edit Variable Set Default-Set



Name:

Description:

 Add

Variable Name	Type	Value	
---------------	------	-------	--

Customized Variables

This category is empty

Default Variables

AIM_SERVERS	Network	[64.12.31.136/32, 205.188.210.203/32, 6...]	  
DNS_SERVERS	Network	HOME_NET	  
EXTERNAL_NET	Network	any	  
FILE_DATA_PORTS	Port	[HTTP_PORTS, 143, 110]	  
FTP_PORTS	Port	[21, 2100, 3535]	  
GTP_PORTS	Port	[3386, 2123, 2152]	  
HOME_NET	Network	any	  

Save

Cancel

Notice our two rock stars `$HOME_NET` and `$EXTERNAL_NET`? They default to “any” as mentioned previously. You can modify these by clicking the pencil icon and adding your own IP ranges. You can use an existing Network object (which is what I recommend) or type in your own IP address information on the fly.

Once you change the values in the Default-Set, any new custom Variable Set objects will start out with these values.

Edit Variable HOME_NET

? X

Name: HOME_NET

Type: Network

Available Networks

Search by name or value

- AIM_SERVERS
- DNS_SERVERS
- EXTERNAL_NET
- HTTP_SERVERS
- SIP_SERVERS
- SMTTP_SERVERS
- SNMP_SERVERS
- SQL_SERVERS
- SSH_SERVERS
- TELNET_SERVERS
- any
- Enterprise1
- IPv4-Private-All-RFC1918
- SQL_Administrators
- 10.0.0.64 38 NAT Depl

Include

Exclude

Included Networks (0)

any

Excluded Networks (0)

none

Network Enter an IP address

Add

Network Enter an IP address

Add

Save

Cancel

Also, keep in mind that once you override the default value for any of these variables, any subsequent updates from Cisco will not take effect. You will find that some of the variables, particularly the port variables, are updated by Cisco from time to time.

Time Range

If you want a policy to apply only during a specified time range, create a Time Range object, then specify that object in the VPN group policy. You can specify time range objects only in VPN Group Policy objects.

All times are in UTC and entered in the 24-hour clock format. For example, you would enter 3:30 p.m. as 18:30.

First, click Add to create a new time range and choose when the new Time Range object will be effective with a start date and time. The New Time Range dialog is shown here:

New Time Range

Name:

Mon-Fri_Time_Range

Effective From
(UTC):

2020-2-24



To:

Never End



February 2020

S	M	T	W	T	F	S
26	27	28	29	30	31	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
1	2	3	4	5	6	7

Today

17:00

Now create a further constraint to recurring intervals by clicking on Add Recurring Interval:



Choose either Daily Interval or Range. Here you can see that the object runs daily from 8 a.m. to 5 p.m.

Edit Recurring Interval



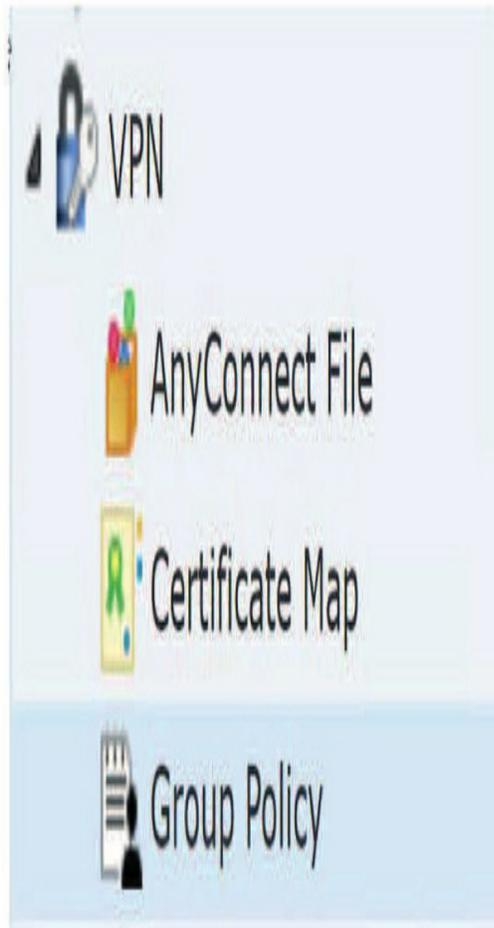
The screenshot shows a dialog box titled "Edit Recurring Interval". It has three main sections:

- Recurring Type:** Two radio buttons are present. The first is labeled "Daily Interval" and is selected (indicated by a black dot). The second is labeled "Range" and is unselected.
- Days of Week:** A row of seven green buttons labeled "Mon", "Tue", "Wed", "Thu", "Fri", "Sat", and "Sun". To the right of these buttons is a checkbox labeled "Daily", which is checked.
- Effective Time (UTC):** Two text input fields are shown. The first contains "08:00" and the second contains "17:00". Between them is the word "To". To the right of the second field is an unchecked checkbox labeled "All Day".

At the bottom right of the dialog box are two buttons: "Add" and "Cancel".

Click Add and then you will see the name, the effective dates, and the time the object is configured for, as shown in the next figure.

To implement this Time Range object, you can only place it in one policy, and as mentioned, that is the VPN group policy. Go the VPN Group Policy object and click on the Add Group Policy button.



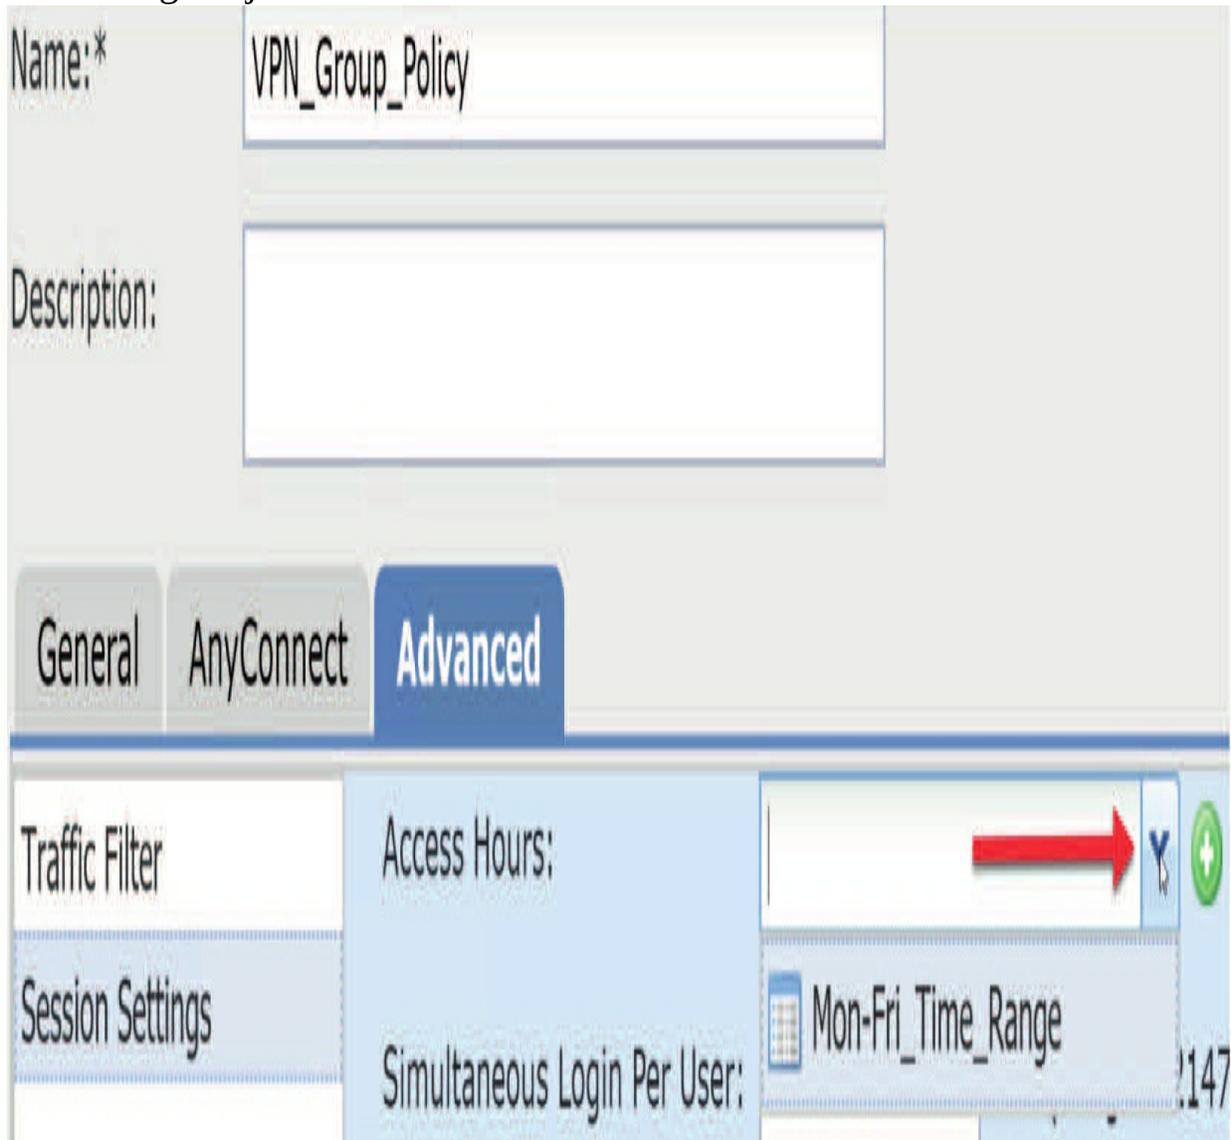
Name	Effective Dates	Time
------	-----------------	------

Mon-Fri_Time_Range

2020-02-24, 17:00 to Never End

Daily Interval: Daily, 08:00 to 17:00

Last, click on the Advanced tab and you can use the pull-down to add your Time Range object.



The screenshot shows a configuration window for a VPN Group Policy. The 'Name' field is filled with 'VPN_Group_Policy'. Below it is an empty 'Description' field. At the bottom, there are three tabs: 'General', 'AnyConnect', and 'Advanced'. The 'Advanced' tab is selected. Under the 'Advanced' tab, there are two sections: 'Traffic Filter' and 'Session Settings'. The 'Access Hours' field is highlighted in blue and has a dropdown menu open. A red arrow points to the dropdown arrow. The dropdown menu shows a calendar icon and the text 'Mon-Fri_Time_Range'. To the right of the dropdown is a green plus icon. The 'Simultaneous Login Per User' field is also visible, with the value '147'.

Security Intelligence Overview

Security Intelligence is an object category that contains three different categories of lists and feeds:

1. Network
2. DNS
3. URL

Network Security Intelligence

Network security intelligence objects are simply IP addresses. These can be used to blacklist or whitelist traffic based on the source/destination IP address. The intent is to block malicious or compromised hosts and also provide a safety net to prevent inadvertently blacklisting critical assets.

DNS Security Intelligence

These are domain names with a bad reputation. These are designed to prevent hosts from resolving and subsequently connecting to evil or compromised servers. They can also be used to identify internal hosts that may already be compromised so they can be remediated.

URL security intelligence

These consists of URLs with a poor reputation. Two of the three objects above—Network and URL—are implemented through the Access Control policy. DNS security intelligence is implemented through the DNS policy.

Each of these is capable of receiving a Security Intelligence feed from Cisco Talos. This feed is a dynamic collection of IP addresses, URLs, or domain names that the FMC downloads at an interval you configure. Changes to this feed are immediately applied to all devices without requiring a manual policy deployment.

You can also create your own custom feeds for any of these categories. Custom feeds are hosted on an HTTP or HTTPS server and—like the Cisco Talos feed—they can be automatically downloaded at a customizable interval. This method also applies the updated feed data to all devices immediately.

You can also create custom lists that you then upload as a Security Intelligence object. However, this is a rather inflexible method, as any change requires re-uploading the entire list. In addition, if a custom list is added or changed, you must deploy policies for it to take effect. This method may be more appropriate for static information such as a whitelist containing your critical routers or gateways.

Network Security Intelligence (SI)

Network Lists and Feeds have been included with the Firepower product since version 8.x. They are simply a list of IP addresses. They are implemented on the Security Intelligence tab in the Access Control policy where they can be used to blacklist or whitelist traffic.

The system comes with four Network Security Intelligence objects by default:

1. The Cisco-Intelligence-Feed
2. Cisco-TID-Feed
3. The Global-Blacklist
4. The Global-Whitelist

The Cisco-Intelligence-Feed

This is actually several lists of IP addresses in different categories. You will see this later when we look at the Access Control policy in chapter 9, where the SI is implemented. This feed is updated by Cisco constantly, and you can download it as often as once every 8 minutes. Because of the fast-paced nature of some malware campaigns, I recommend taking advantage of this update frequency and setting it to the minimum interval. If you upgraded from a previous version of Firepower, you may find that the interval is still set to the old minimum window of 2 hours.

The Cisco Threat Intelligence Director (TID) feed Starting with Cisco Firepower 6.2.2, the TID feature was added to Firepower and is intended to supplement other Firepower functionality, offering even more defenses against threats. TID can bring in data from threat intelligence *sources* and publishes the data to all configured managed devices (*elements*).

TID enhances the Firepower system's ability to block connections based on SI from third-party sources as follows:

TID adds supports for additional SI filtering with SHA-286.

Cisco's Security Intelligence allows you to filter traffic based on IP address, URL, and domain name. TID also adds support for filtering on SHA-286 hash values.

TID supports manual file uploads or STIX and TAXII.

You can manually upload flat files or configure the system to retrieve flat files from a third-party host in order to import threat intelligence into the Firepower system. In addition, TID can download intelligence provided in Structured Threat Information eXpression (STIX), which can use TAXII feeds as sources. STIX and TAXII are standards developed in an effort to improve prevention and mitigation of cyberattacks. STIX states the *what* of threat intelligence, while TAXII defines *how* information is relayed.

TID provides granular filtering actions.

The SI allows you to only blacklist or whitelist by object, not by individual components of an object. With TID, you can configure filtering actions for individual simple indicators or individual observables.

TID does not require redeployment.

After you modify any settings in the Access Control policy, you must always redeploy the changes to the managed devices. TID allows you to configure sources, indicators, and observables without redeploying, and the system automatically publishes new TID data to the elements.

FMC in a High Availability configuration

Last, if you have your hardware FMC in a High Availability configuration, the system does not synchronize TID configurations and TID data to the standby Firepower Management Center. You need to make sure that you back up your FMC every night on your active FMC so that you can restore the TID data after failover. When you create your backup profile, TID backups is a check box selection you need to choose to enable TID backups.

The Global-Blacklist and Global-Whitelist

These objects are empty by default. They are designed to be populated with your own custom IP addresses during the process of event analysis. By right-clicking on an IP address in any of the various event views or in the Context Explorer, you can add the address to either of these lists on the fly.

The scenario goes something like this:

“Holy malware Batman, this intrusion event is really bad. We should block all communications to/from that host”

<right-clicks on IP address – selects Blacklist IP Now>

“Whew, that was close!”

<u>Source IP</u>	<u>Message</u>
191.235.95.108	MALWARE-CNC Win.Trojan.Pmabot outbound connection (1:37215:2)
66.111.12.100	CNC Win.Trojan.Pmabot outbound connection (1:37215:2)
66.111.12.100	CNC Win.Trojan.Pmabot outbound connection (1:37214:2)
10.0.0.1	NT SMB INVALID SHARE (133:26:2)
220.222.11.1	Bash CGI environment variable injection attempt (1:31978:1)
222.78.1.1	EBAPP Joomla JDatabaseDriverMysqli unserialize code execut
172.17.0.1	UNESCAPED SPACE IN URI (119:33:2)
172.17.0.1	BARE BYTE (119:4:1)
91.188.1.1	SESSION HIJACKED SERVER (129:10:1)
91.188.1.1	SESSION HIJACKED CLIENT (129:9:1)
91.188.1.1	SMALL SEGMENT (129:12:2)
91.188.1.1	SESSION HIJACKED CLIENT (129:9:1)

- Open in New Window
- Whois
- View Host Profile
- Blacklist IP Now
- Whitelist IP Now
- Open in Context Explorer
- Generate Events
- Drop and Generate Events
- Disable Rule
- Threshold
- Suppression
- Rule documentation
- Edit rule
- Exclude



The above figure shows what the right-click menu looks like for an intrusion event, notice how you can edit and even disable an IPS rule from this menu.

...and here is a connection event menu.

Open in New Window

Exclude

Open in Context Explorer

Whois

View Host Profile

Blacklist IP Now

Whitelist IP Now

AlienVault IP

IBM X-Force Exchange IP

Looking Glass IP

Recorded Future IP

Talos IP

Threat Grid IP

Threat Response IP

Umbrella IP

Virus Total IP

In the figure, the menu has vastly improved where you can not only blacklist and whois, but now use third-party and Cisco products to verify and test an IP.

Whois and Blacklist IP Now are my go-to items that I use.

When you select this menu item you receive an “Are you sure?” confirmation dialog. Selecting Blacklist IP Now adds the IP to the Global-Blacklist and updates this list on all your devices within seconds.

You can also use the Whitelist IP Now option to add the host to the Global-Whitelist and allow the IP to bypass any Network Security Intelligence blacklists. It’s important to understand that you only bypass SI and not the whole Snort process when you whitelist a host.

When you click on Network Lists and Feeds under the Security Intelligence category, you will see the default lists and the CiscoIntelligence-Feed. This is shown here:



Name	Type
Cisco-Intelligence-Feed <i>Last Updated: 2019-09-18 20:14:53</i>	Feed
Cisco-TID-Feed <i>Last Updated: 2019-09-18 20:53:34</i>	Feed
Global-Blacklist	List
Global-Whitelist	List

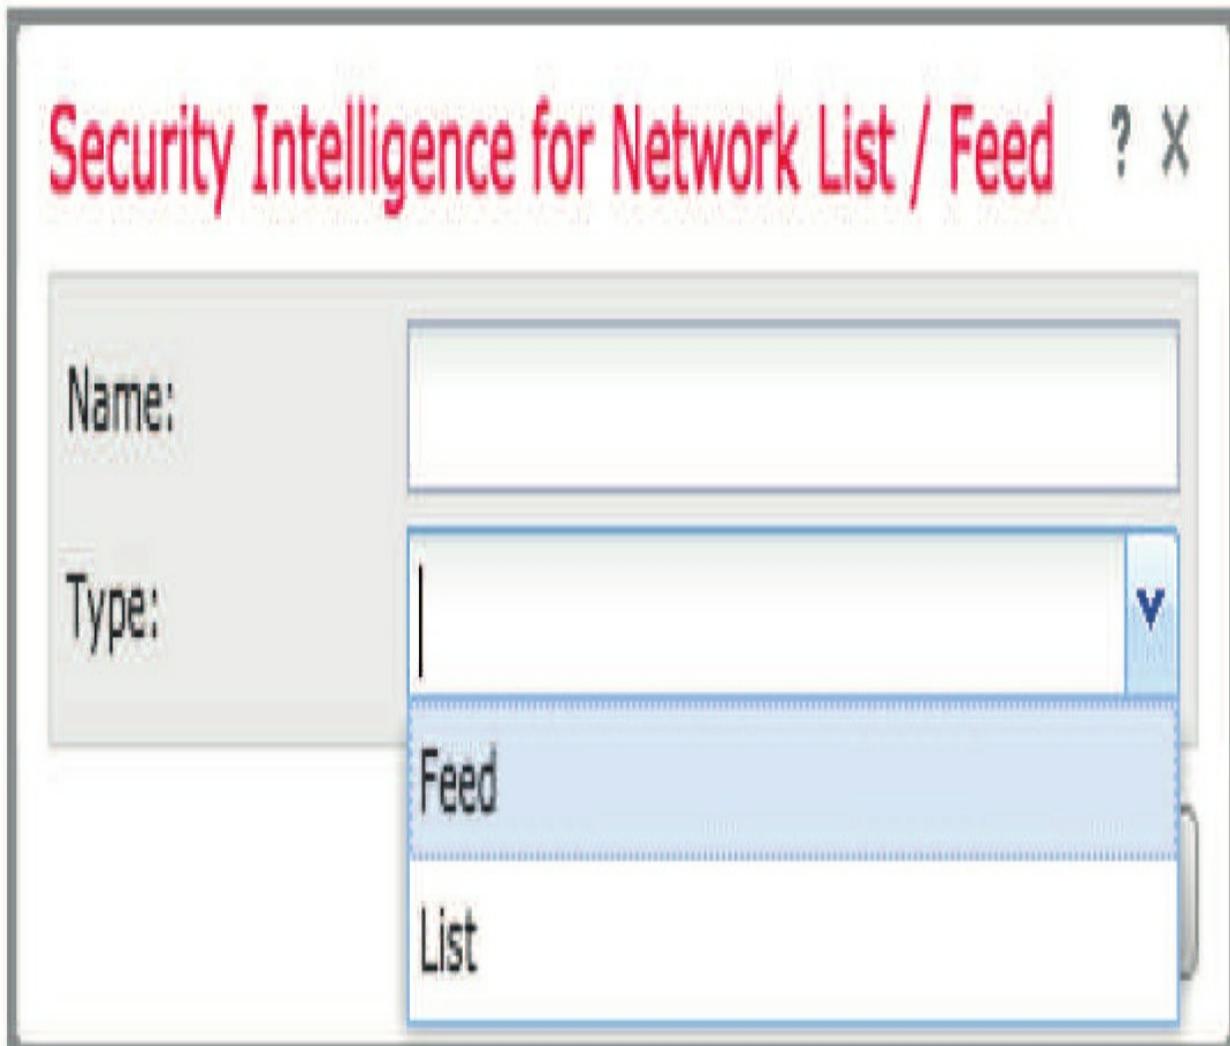
While it appears that you can edit the Global-Blacklist and Global-Whitelist here, you are actually limited to removing IP address entries only. They can only be added through the right-click method mentioned previously.

Note: According to the Firepower online help, adding an entry to either global list causes the devices to update immediately. However, removing an entry from either list requires a policy deployment. This is mostly true; however, what really happens is when you add an entry to a global list, the entire list is pushed out immediately. So if you have also just removed an item from a list, that change will also be pushed out with the new list. What I

just wrote here matches Cisco’s documentation.

However, there is a “but” to this as well. Starting in 6.3 code, if you either add or remove an IP address from the list, it is automatically and silently pushed out immediately. This was not in Cisco’s release documentation for 6.3 and I stumbled on this information by accident while at a customer site, and then I wrote and posted a blog on the new changes in the 6.3 and 6.4 code that were never in the release notes.

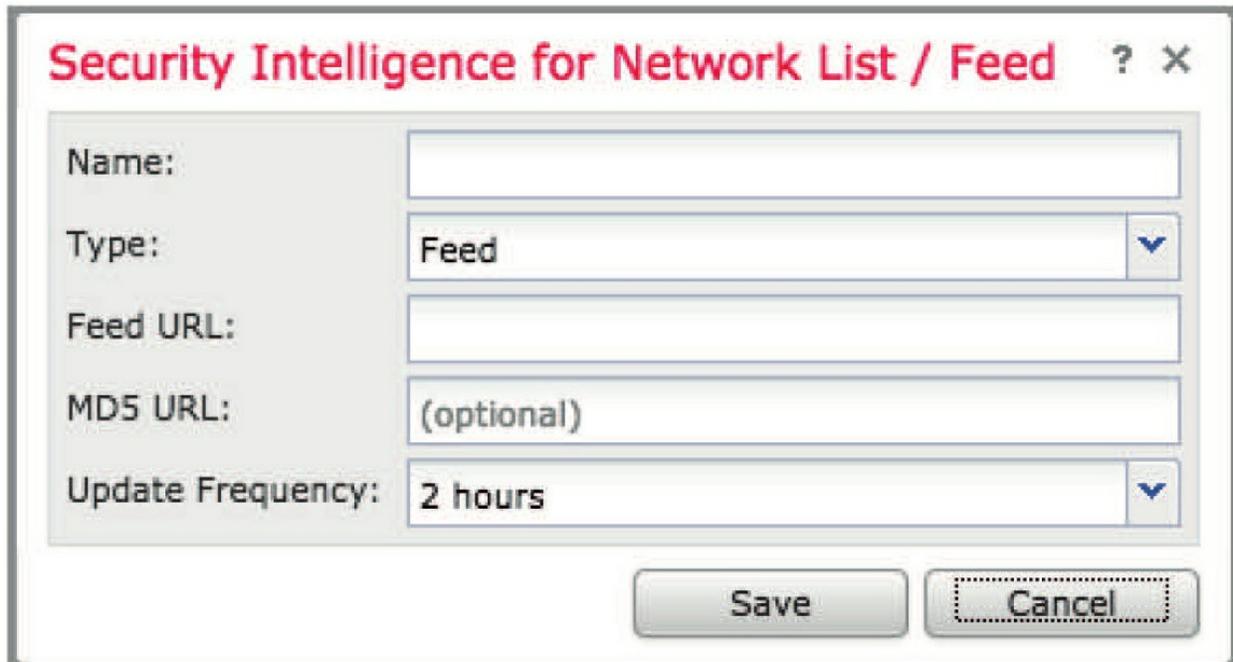
To add a custom list or feed, click the Add Network Lists and Feeds button. Next, enter the name and pick the type of object—either Feed or List.



The Security Intelligence for Network List / Feed dialog is shown here:

If you selected List, you will be presented with a dialog to browse to the list, which should be a text file with one IP/CIDR entry per line. The file can be a maximum of 800 MB in size.

If you select Feed, you will see the following dialog:



The screenshot shows a dialog box titled "Security Intelligence for Network List / Feed". It contains the following fields and controls:

- Name:** An empty text input field.
- Type:** A dropdown menu currently set to "Feed".
- Feed URL:** An empty text input field.
- MD5 URL:** A text input field containing the text "(optional)".
- Update Frequency:** A dropdown menu currently set to "2 hours".

At the bottom of the dialog are two buttons: "Save" and "Cancel".

Here you enter the URL where the feed file is located and an optional MD8 URL that can be used to determine if the file has been updated. If the MD8 hash has not changed since the previous update, the feed file will not be downloaded. The Update Frequency field will default to 2 hours, but you can set it as low as 30 minutes to update your feed more quickly.

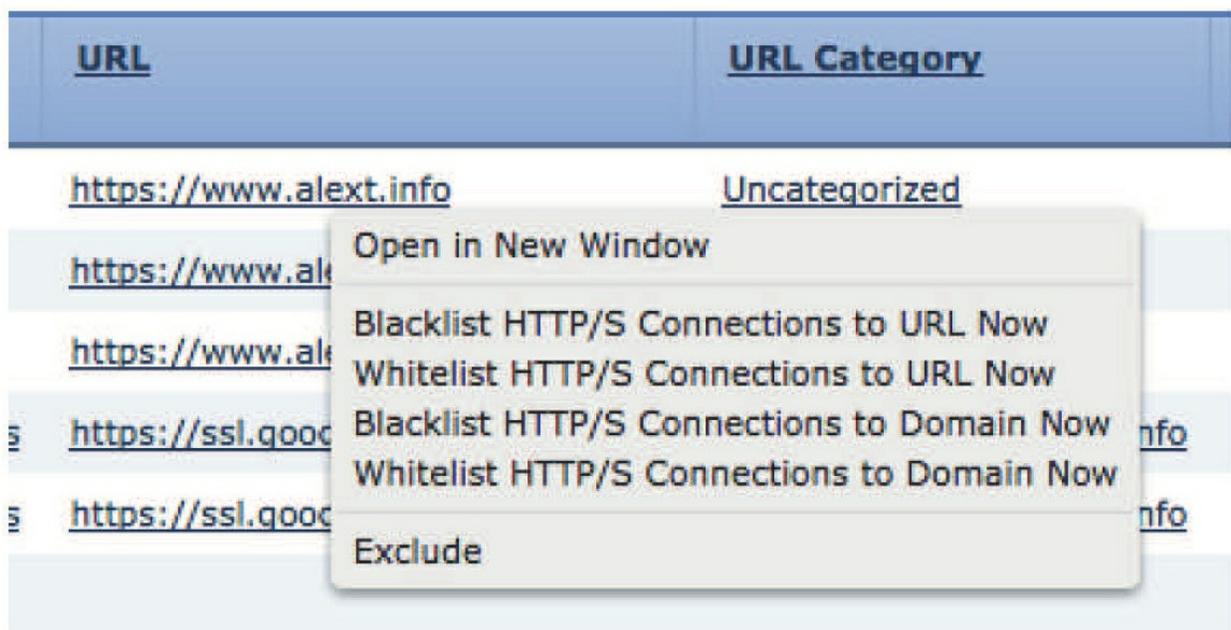
DNS Security Intelligence

The DNS Security Intelligence feature was added in Firepower version 6. It works like its Network sibling by receiving a feed from Cisco. However, these are domain names with a poor reputation rather than IP addresses. Also similar to Network Security Intelligence, there are a Global-Blacklist and Global-Whitelist. You can add your own lists and feeds, which operate pretty much the same way, by either uploading a text file full of domains or using a custom feed that can be updated at an interval you select. DNS Security Intelligence is implemented via the DNS policy.

We will cover how this works in detail later in the chapter on DNS, chapter 13.

The dialogs for adding/removing DNS lists and feeds are in nearly every way identical to those for Network lists and feeds. The only difference is how you go about adding entries to the Global Whitelist and Blacklist. Instead of right-clicking on an IP address, you do so on a URL.

This figure shows an example of the menu you will see if you right-click on a URL in a connection event.



By selecting **Blacklist HTTP/S Connections to URL Now**, you will add this to the URL Global Blacklist. By selecting **Blacklist HTTP/S Connections to Domain Now**, you will add this entry to the DNS Global Blacklist.

URL Security Intelligence

The last intelligence category is URL Security Intelligence. As you may have guessed, this is for URLs with a poor reputation. When you select URL Lists and Feeds from the Security Intelligence objects list, you will see a page similar to the ones for Network and DNS. However, one thing is missing—there is no entry for the Cisco feed. This is because the DNS and URL entries are combined into a single feed, which is managed under DNS Lists and Feeds. If you return to the DNS Lists and Feeds page, you will notice that the

built-in feed is called **Cisco-DNS-and-URL-Intelligence-Feed**. So, while there are three security intelligence categories (Network, DNS, and URL), there are only two Cisco feeds.

As with URL objects, the URL Security Intelligence Lists and Feeds entries will match any URL that contains the entry. Thus a feed entry of `www.lammle.info` will match any page on the site.

Security Intelligence Under the Hood

Since this is an advanced book, let's talk about some of the underthe-hood mechanics when it comes to security intelligence. A common question is, "Where can I find a list of the IP, URL, DNS entries in the Cisco feed?" There is no way within the graphical user interface to find this information.

The best you can do is to see how many entries are in a given feed category. To find out how many entries are in a given category, navigate to **Policies>Access Control>Access Control**.

Then edit one of your policies and click the Advanced tab. You can then use your mouse to hover over one of the Network, DNS, or URL categories. A pop-up will indicate how many entries are currently in this category.



But what about the actual entries? To find these you must SSH to either a device or the FMC. You will find the three types of security intelligence entries in the following three locations:

Network – `/var/sf/iprep_download`

DNS – /var/sf/sidns_download

URL – /var/sf/siurl_download

Here you will find separate text files for each security intelligence category. You will also find text files for any of your custom feeds as well.

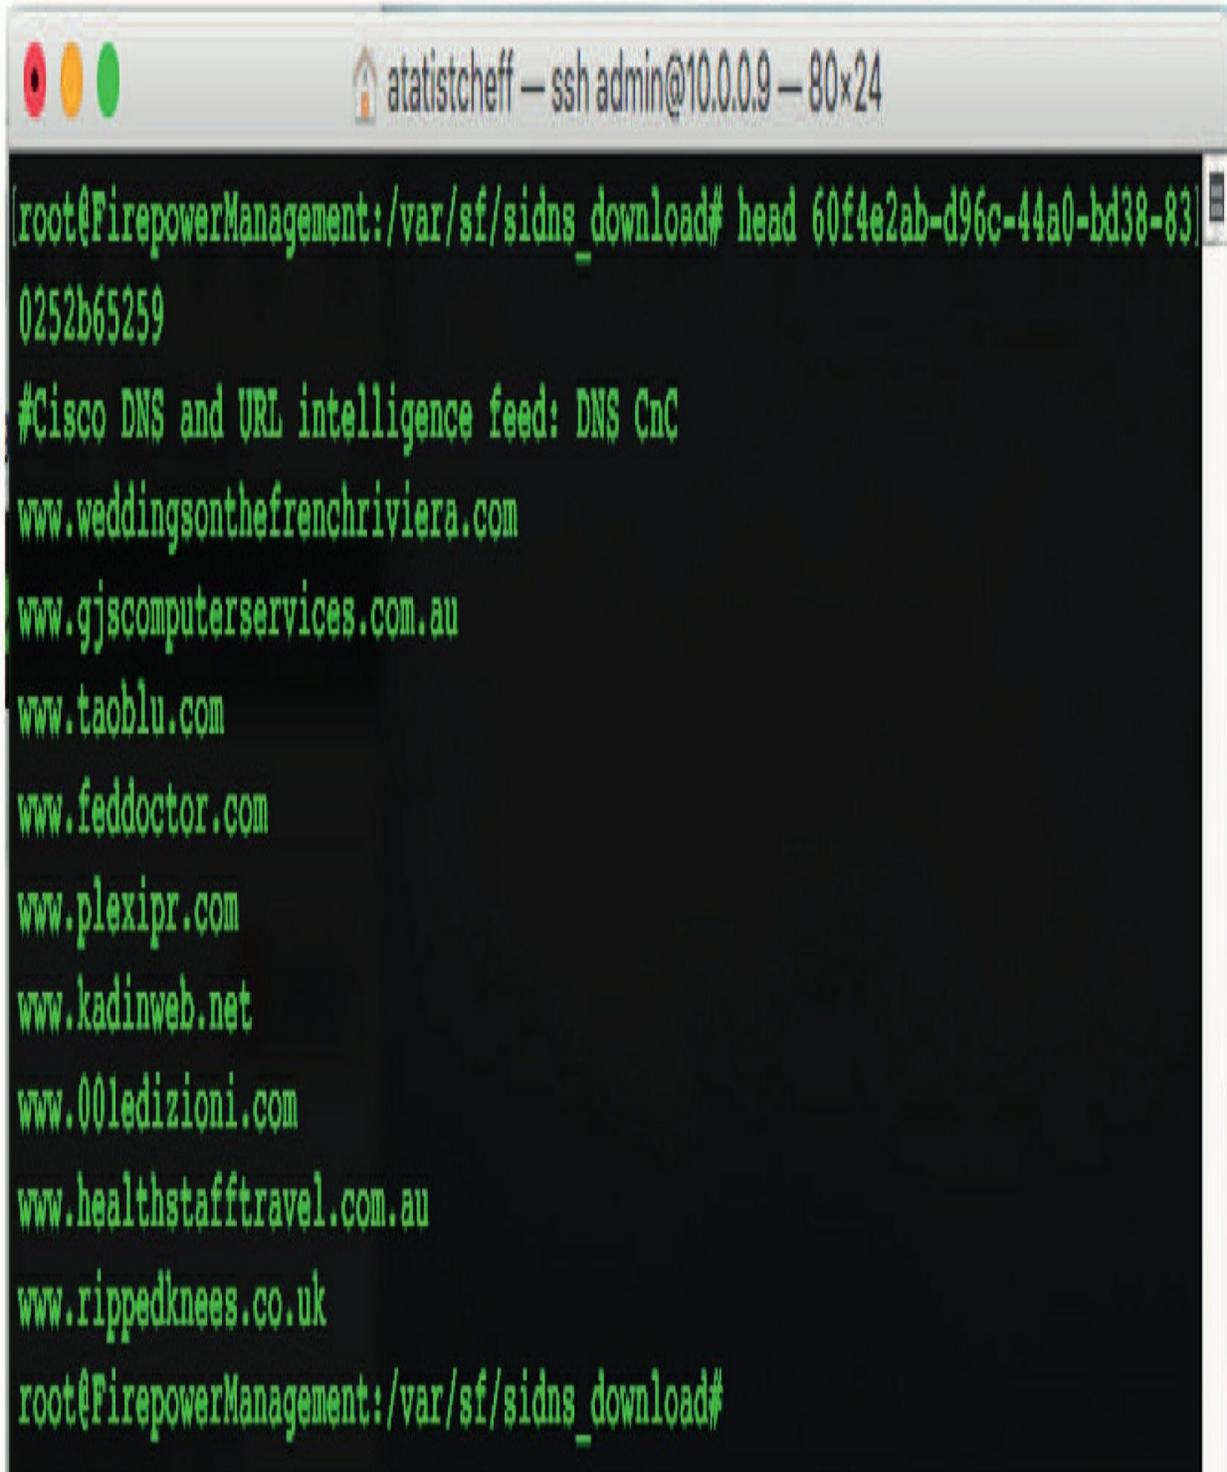
The output here shows the contents of the `sidns_download` folder.

```
ataticsheff — ssh admin@10.0.0.9 — 80x24
[root@FirepowerManagement:/var/sf/sidns_download# ls
032ba433-c295-11e4-a919-d4ae5275b77b 6ba968f4-7a25-4793-a2c8-7cc77f1f1074
1b117672-7453-478c-be31-b72e89ca2dde A27C6AAE-8E52-4174-A81A-47C59FECd3a5
23f2a124-8278-4c03-8c9d-d28fe08b71ab Cisco_DNS_Intelligence_Feed
2CCDA18E-DDFF-4F5C-AF9A-F00985219707 IPRVersion.dat
2b15cb6f-a3fc-4e0e-a342-ccc5e5804576 bldf3aa8-2841-4c88-8e64-bfaacec7111f
30f9e69c-d64c-479c-821d-0e4edab8348d d7d996a6-6b92-4a56-8f10-e8506e432fb8
3e2af68e-5fc8-4b1c-b5bc-b4e7cab5abcd health
5a0b6d6b-e2c3-436f-b4a1-48248b331d39 peers
5f8148f1-e5e4-427a-aa3b-e01c2745d663 rep_dd.yaml
60f4e2ab-d96c-44a0-bd38-830252b65259 tmp
root@FirepowerManagement:/var/sf/sidns_download#
```

The files have unrecognizable UUID (Universally Unique Identifier) names, but if you use `cat`, `head`, or `tail` to look at their contents, you will see they are simply text files. Each one contains the name of the list as a comment in the

first line.

The next output shows the contents of one of these. Turns out it is the file for the Command and Control (CnC) DNS list.



```
atalistcheff — ssh admin@10.0.0.9 — 80x24
[root@FirepowerManagement:/var/sf/sids_download# head 60f4e2ab-d96c-44a0-bd38-83]
0252b65259
#Cisco DNS and URL intelligence feed: DNS CnC
www.weddingsonthefrenchriviera.com
www.gjscomputerservices.com.au
www.taoblu.com
www.feddoctor.com
www.plexipr.com
www.kadinweb.net
www.00ledizioni.com
www.healthstafftravel.com.au
www.rippedknees.co.uk
root@FirepowerManagement:/var/sf/sids_download#
```

Using this technique, you can find out the contents of any of the security intelligence download files for each of the three categories. One huge caveat, however: these files are updated frequently.

Depending on the update frequency you have selected, an entry that was here 30 minutes ago may be gone now. If you're trying to troubleshoot an issue or predict whether a given IP, domain, or URL will be blocked, this may not be a viable technique. It's nice to know, however, and now you can impress your friends with your in-depth knowledge of Firepower!

Sinkhole

The Sinkhole object is another one that is new in version 6. I will go into detail on how this works in the chapter 13 on DNS. For now, I'll whet your appetite with a brief description of the Sinkhole object and a bit of information on how to create one.

A sinkhole is a DNS server that is designed to return nonroutable addresses in response to DNS queries. More accurately, these IP addresses—routable or not—do not resolve to an actual server. Since DNS resolution is the first step in virtually any TCP/IP connection, this is designed to prevent a user/host from successfully establishing a connection. By making it impossible to resolve a name to an IP address, the connection is stopped dead in its tracks.

You might ask, “Why not just block the DNS request? Why use a sinkhole?” Those are good questions. One reason has to do with the placement of the IPS in relation to your internal DNS server. Oftentimes, the IPS sees a DNS request for an evil domain coming from your DNS server, who forwarded it on behalf of an infected client (infected by malware or just infected by a user).

The IPS can easily block this request using a Snort rule. However, what you really want to know is, “Which of my hosts made that DNS query?” Since the host is trying to resolve an “evil” hostname, it would be very helpful to know which host this is and maybe send in the incident response team.

By using a sinkhole, we can return a bogus IP address to the DNS server who relays that to the (infected) host. Now when the host tries to actually connect

to this sinkhole IP—BAMO—we've got him! The IPS can see this request and generate an alert that somebody is trying to connect to our sinkhole IP. The only way a host would get this IP address in the first place is if we gave it to him in response to an evil DNS request.

I said this would be a short description, so let's stop here while we're ahead. I'll explain this further with fancy pictures and diagrams in the DNS chapter. For now, let's talk about how you would go about setting up a Sinkhole object.

There are no Sinkhole objects created by default. To create one, click the Add Sinkhole button.

This will display the dialog shown in here:

Sinkhole



Name:

IPv4 Address:

IPv6 Address:

Log Connections to Sinkhole:

Block and Log Connections to Sinkhole:

Type:

Save

Cancel

The options for your sinkhole are as follows:

- **Name** – Friendly name; you can use spaces if desired.
- **IPv4 Address** – The IPv4 address. Pick an address that resolves to something North of your IPS (outside your network) but that no host should ever try to connect to. One suggestion is to use the IPv4 space reserved for “documentation” according to RFC 8737. There are three ranges: TEST-NET-1

192.0.2.0/24, TEST-NET-2

198.81.100.0/24, TEST-NET-3

203.0.113.0/24. While technically public space, some

people also used **1.1.1.1** because it’s so easy to pick out when scanning an event view page; however, even though we used this IP address for years, you no longer want to do this as it is now in use. Keep in mind this should be a single IP address entry, not a range.

- **IPv6 Address**– IPv6 does not have a specifically reserved range for documentation, but given the size of the IPv6 address space, you should be able to come up with something unique!

- **Log Connections to Sinkhole or Block and Log Connections to Sinkhole** – Fairly straightforward. Do you want to allow this packet to continue or block it at the IPS? If you are using an actual server listening on port 83/UDP for your sinkhole IP address, you may want to allow the connection so you can log the request there. Either way, Firepower will log it as a Security Intelligence event.

- **Type** – The DNS Security Intelligence category breaks down into three types (and None of course):

- Command and Control (CnC)

- Malware

- Phishing

The purpose of this selection is to allow you to have different

DNS sinkhole IPs for each type of DNS request. You may want to prioritize CnC events and remediate these more quickly. This way you can set up special alerting for any hosts triggering your CnC sinkhole IP.

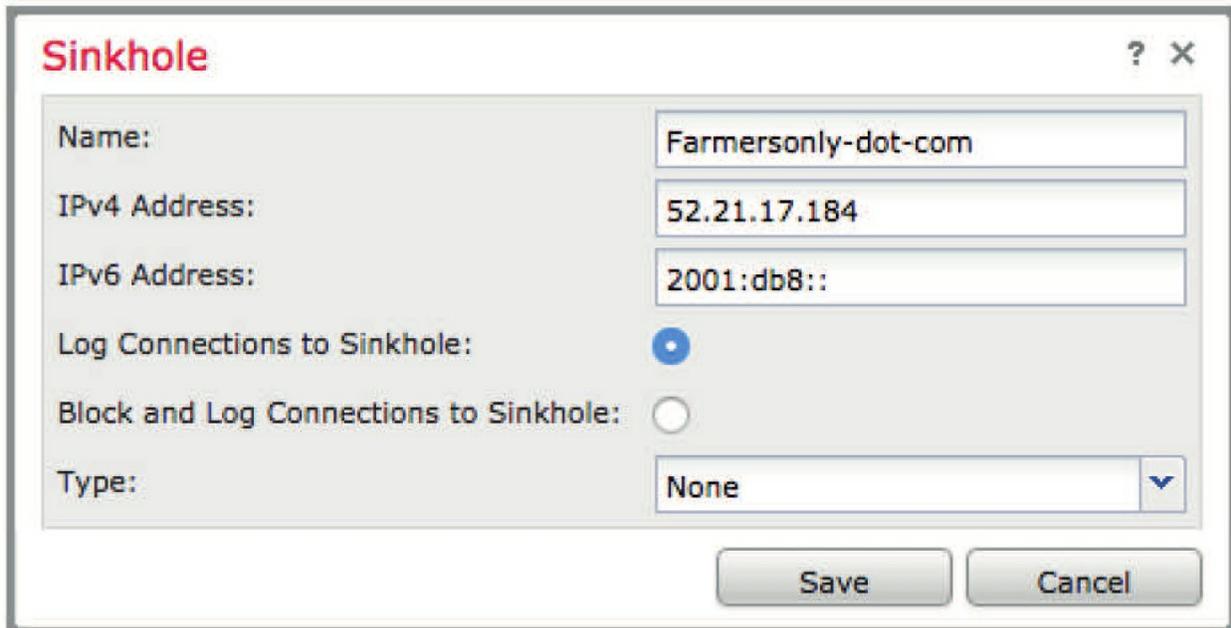
Once your sinkhole objects are added, you can use them in your DNS policy to catch much evil on your network!

Sinkhole Reloaded

Now that you know how a sinkhole works, let's think of ways we can (ab)use this feature. Since you can now make the DNS server return whatever IP you want, you control *everything!* Say you have a list of sites that you'd rather nobody go to. Maybe you want everyone to use only your favorite dating site—farmersonly.com. To implement your evil plan using the DNS sinkhole feature, here are the steps to take.

1. Create a text file with all the domains of all those other sites like eharmony.com, match.com, okcupid.com, etc. 2. Upload this file as a custom DNS Security Intelligence object.
3. Create a Sinkhole object and use the actual farmersonly.com site for the IP address. Make sure you don't block the connections to this sinkhole.
4. In your DNS policy, add a rule with an action of Sinkhole, and for the Sinkhole, select your Farmersonly-dot-com Sinkhole object. On the DNS tab, add only your custom DNS Security Intelligence object.
5. Deploy policies

What you did is tell the device to sinkhole any request for one of those other dating sites and instead return the farmersonly.com IP address. The result will be that anyone trying to go to match.com, eharmony.com, etc. will find themselves looking at the farmersonly.com site. Then get ready for the helpdesk calls!



The image shows a configuration dialog box titled "Sinkhole". It contains the following fields and controls:

- Name:** Farmersonly-dot-com
- IPv4 Address:** 52.21.17.184
- IPv6 Address:** 2001:db8::
- Log Connections to Sinkhole:**
- Block and Log Connections to Sinkhole:**
- Type:** None (dropdown menu)

At the bottom right, there are two buttons: "Save" and "Cancel".

File List

The File List object has not changed since version 5 of Firepower. It simply provides a way to customize file detection much like the blacklist or whitelist in security intelligence. There are two objects here, the Clean-List and the Custom-Detection-List. Both of these contain SHA-256 file hashes. They are designed to override the system's default behavior for a given SHA-256 hash.

The Clean-List contains SHA-256 hashes that should be considered clean regardless of their disposition in the Cisco cloud. Conversely, the Custom-Detection-List contains hashes considered to be malicious. Both of these lists are checked by the FMC prior to querying the Advanced Malware Protection (AMP) cloud when performing malicious file checks. The option to use this feature is controlled from the Malware & File policy.

If you click the pencil icon to the right of either of the two lists, you get a dialog like the one here:

File List



Note: For file lists to take effect, a file policy containing a rule with either a Malware Cloud Lookup or Block Malware action must be deployed to your devices.

Name:

Add by:

List of SHAs

Description

Calculate SHA

Enter SHA Value

No data to display



Save

Cancel

There are three ways to add SHA-256 values to this object:

1. List of SHAs – This is a text file containing SHA-256 hash values, one per line. You can give this entry a description when you upload the file. You can upload multiple SHA lists in this manner.
2. Calculate SHA – If you have the file but do not know the SHA256 hash, you can upload the file and let the FMC calculate the hash value. The file is not saved on the FMC; it's only used to calculate the hash for this entry.
3. Enter SHA Value – If you know the SHA-256 value, you can use this option and paste it into the SHA-256 field.

Cipher Suite List

The Cipher Suite List object can be used if you are performing SSL decryption on your device. The object comprises one or more cipher suites, each representing a method used to negotiate SSL or TLS encrypted sessions. You can use this list in your SSL rules to control which cipher suites the rule applies to.

This figure shows an example of a Cipher Suite List object.

Cipher Suite List



Name: Interesting-Cipher-Suites

Available Cipher Suites

Search

- SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA
- SSL_RSA_FIPS_WITH_DES_CBC_SHA
- TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA256
- TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA
- TLS_DHE_RSA_WITH_CAMELLIA_128_CBC_SHA256
- TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA
- TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA256
- TLS_DHE_RSA_WITH_DES_CBC_SHA
- TLS_DH_Annon_WITH_AES_128_GCM_SHA256
- TLS_DH_Annon_WITH_AES_256_GCM_SHA384
- TLS_DH_Annon_WITH_CAMELLIA_128_CBC_SHA
- TLS_DH_anon_WITH_CAMELLIA_128_CBC_SHA256

Add

Selected Cipher Suites

- TLS_RSA_WITH_NULL_MD5
- TLS_RSA_WITH_NULL_SHA
- TLS_ECDHE_RSA_WITH_NULL_SHA
- TLS_ECDHE_ECDSA_WITH_NULL_SHA

Save

Cancel

Distinguished Name

Distinguished Name is another object related to SSL/TLS decryption. Each object represents the distinguished name for a public key certificate's subject or issuer. The idea is that you can use these objects in your SSL policy to control decryption based on server certificates with a certain subject or issuer. Oftentimes these are used to determine which SSL sites *not* to perform decryption for. This typically applies to outbound connections.

This is illustrated by the names of the default object that come with the system.

Here shows the list of default Distinguished Name individual objects.

Add Distinguished Name

Filter

- Network
- Port
- Interface
- Tunnel Zone
- Application Filters
- VLAN Tag
- Security Group Tag
- URL
- Geolocation
- Variable Set
- Security Intelligence
 - Network Lists and Feeds
 - DNS Lists and Feeds
 - URL Lists and Feeds
- Sinkhole
- File List
- Cipher Suite List
- Distinguished Name
 - Individual Objects
 - Object Groups
- PKI
- SLA Monitor
- Prefix List

Name	Value	
CN=*.citrixonline.com	CN=*.citrixonline.com	
CN=*.data.toolbar.yahoo.com	CN=*.data.toolbar.yahoo.com	
CN=*.fedoraproject.org	CN=*.fedoraproject.org	
CN=*.itunes.apple.com	CN=*.itunes.apple.com	
CN=*.logmein.com	CN=*.logmein.com	
CN=*.mozilla.org	CN=*.mozilla.org	
CN=*.rhn.redhat.com	CN=*.rhn.redhat.com	
CN=*.sls.microsoft.com	CN=*.sls.microsoft.com	
CN=*.update.microsoft.com	CN=*.update.microsoft.com	
CN=*.update.microsoft.com	CN=*.update.microsoft.com	
CN=account.live.com	CN=account.live.com	

Displaying 1 - 11 of 11 rows Page 1 of 1

All of the default objects are sites where the use of SSL decryption will break the underlying application. This usually stems from the fact that some sites add extra or nonstandard functionality to their SSL communications. Because these deviate from “normal” SSL, any attempt to decrypt these communications will cause the underlying application or process to fail.

Other sites that may be good candidates to bypass decryption include online banking or healthcare sites—in other words, any site that you trust as legitimate and you do not want to decrypt for privacy reasons. Remember, we’re talking mostly about outbound SSL, so one of your users initiated the connection. Cracking open someone’s online banking transaction—even in the name of cyber security—is pretty much always a bad idea.

You can add your own objects and object groups. As you can see from the examples, wildcards are also supported.

PKI

Public Key Infrastructure (PKI) objects are used when you are performing SSL decryption on your device(s). Rather than going into detail on how PKI works, I’ll give a short description of each of the objects:

- **Internal CAs** – If you are performing outbound decryption, it means you will have to re-sign SSL communications. In this re-signing process, you will use a certificate that was either (a) generated by your FMC or (b) imported into your FMC from an external certificate authority (CA). The Internal CA object allows you to perform either operation.
- **Trusted CAs** – These objects represent the CA public certificate that belongs to a trusted CA. You can import your own trusted CA objects. The system comes with over 200 preconfigured Trusted CAs.
- **External Certs** – These objects represent a server public key certificate that does not belong to your organization. They consist of the object name and the certificate. You can use these in SSL rules to control whether or not you decrypt traffic using the server certificate. An example would be a self-signed server certificate that you trust but cannot verify because it’s not signed by a

trusted CA.

- **Internal Certs** – These objects represent server public key certificates that belong to your organization. These consist of the object name, the public key certificate, and the paired private key. These are used for the following purposes:

- Decrypting incoming traffic to one of your organization's servers using the known private key
- Identity Services Engine (ISE) integration
- Captive portal configuration to authenticate the identity of your captive portal device when users connect via web browser

FTD Only Settings

The following objects are only available when using FTD devices. They are more commonly associated with some of the advanced routing or VPN features courtesy of the system's ASA lineage. This included the following objects:

- SLA Monitor
- Prefix Lists
- Route Map
- Access Lists
- AS Path
- Community List
- Policy List
- VPN
- Address Pools
- FlexConfig
- Radius Server Group

We will touch on each one in the following sections. For more information, check out the online help or the user guide.

SLA Monitor

The SLA (Service Level Agreement) Monitor object is used to monitor connectivity to a monitored address through a given interface. It tracks the availability of the route to the address. This is accomplished by periodically

sending ICMP echo request (ping) packets and waiting for a response. If the request times out, the route is removed from the routing table and replaced with a backup route.

These objects are used in the Route Tracking field of an IPv4 Static Route policy.

You'll find this under the routing tab on your FTD device.

New SLA Monitor Object ? X

Name:

Description:

Frequency:

SLA Monitor ID*:

Threshold:

Timeout:

Data Size:

ToS:

Number of Packets:

Monitor Address*:

Available Zones ↻

- Gig0.0-ESXi
- Gig0.1-Switch

Selected Zones/Interfaces

Prefix Lists

The IPv4 and IPv6 Prefix List objects are used when configuring route maps, policy maps, OSPF filtering, and BGP Neighbor filtering.

Route Map

The Route Map object is used when redistributing routes into any routing process. They are also used when generating a default route into a routing process. A route map defines which of the routes from the specified routing protocol are allowed to be redistributed into the target routing process.

Access Lists

Access control lists (ACLs) come in two types – standard and extended. You can use them when configuring particular features, such as route maps.

Traffic identified as allowed by the ACL is provided to the service. Blocked traffic is excluded from the service. Excluding traffic from a service does not mean that it is blocked altogether.

Standard Access List

A standard access list only uses a source IPv4 address to define an entry, and the action can only be Allow or Block.

Extended Access List

An extended access list allows matching traffic based on source and destination IP, protocol, and port and also supports IPv6 addresses.

AS Path

An AS path is mandatory when setting up BGP routing. It consists of a sequence of AS numbers through which a network can be accessed.

Community List

A community is an optional attribute for configuring BGP. It is a group of destinations sharing a common attribute and is used for route tagging. Once again, check out the help for more information on this feature. It is only available on FTD devices.

Policy List

The Policy List object is used when configuring route maps. When a policy list is referenced within a route map, all of the matching statements in the list are evaluated and processed.

VPN

The objects under the VPN category are used in configuring the key exchange and encryption protocols used for VPN communications. There are several pre-created policies and proposals. You can also create your own.

Address Pools

These are the IPv4 and IPv6 address pools that can be used with the diagnostic interface with device clusters or for VPN remote access profiles.

FlexConfig

FlexConfig is a feature that allows the use of certain ASA configuration settings that are not yet available through the FMC user interface. It was added to allow configuration of these less-used features and because there is no “config” command available for the ASA side of FTD. The FlexConfig objects are pre-created templates you can copy and use to implement additional features on your FTD device.

You don't need to use the templates if you are an ASA guru and know just what commands you want to execute on your device. In the next figure, I created a new FlexConfig object to enable Netflow reporting to a StealthWatch Flow Collector.

You can hard-code the values as shown here or use the Insert button to insert existing variables or even ASA system variables.

Add FlexConfig Object

? X

Name:

Netflow-to-Stealthwatch

Description:

 Insert ▾



Deployment:

Once ▾

Type:

Append ▾

```
flow-export destination $interface 10.0.0.33 4444
flow-export active refresh-interval 1
flow-export template timeout-rate 1
flow-export delay flow-create 15
policy-map global_policy
class class-default
  flow-export event-type all destination 10.0.0.33
```

Variables

Name	Dimension	Default Value	Property (Type:...	Override	Description
------	-----------	---------------	--------------------	----------	-------------

No records to display

Save

Cancel

Radius Server Group

The RADIUS Server Group object is used to configure RADIUS authentication for remote access VPN users. It's fairly easy to configure. Just give your group a name and add up to 16 RADIUS servers. You can optionally configure an LDAP realm if desired.

You would then select this realm in your identity policies to access the associated RADIUS server group when determining the VPN authentication identity source for a traffic flow.

Intrusion Rules

This seems like an odd place to discuss intrusion prevention system (IPS) rules since we haven't covered the IPS policy at all yet, but this part is actually found in the CCNP Security SNCF Object objectives for the exam, because inside Object Management is actually found a tab for intrusion rules, which can help you manage, edit, and configure your IPS rules.

This tab is called the Intrusion Rule Editor since you can edit rules here, and I am only going to cover the minimum for the exam at this point, and then I'll provide more details on editing intrusion rules in the IPS chapter, but you must know what this tab does.

This Intrusion Rules tab is shown here:



So, when do you use this? Well, not too much for a normal Firepower administrator, but if you tweak and edit your IPS policy rules, then this is

another place to search, find, change, and save a rule. For example, if I want to make a rule an advanced Pass rule, this is the place I would come to do that configuration.

You can search for a rule, change the rule, and you can even create a new rule by clicking on the Create Rule button on the right.

You'll receive then receive the screen shown here:

Create New Rule

Message

Classification A Client was Using an Unusual Port ▼
[Edit Classifications](#)

Action alert ▼

Protocol icmp ▼

Direction Directional ▼

Source IPs Source Port

Destination IPs Destination Port

Detection Options

ack ▼ Add Option Save As New

From here you can write the message you want to use, set a classification for the rules, and define

Action

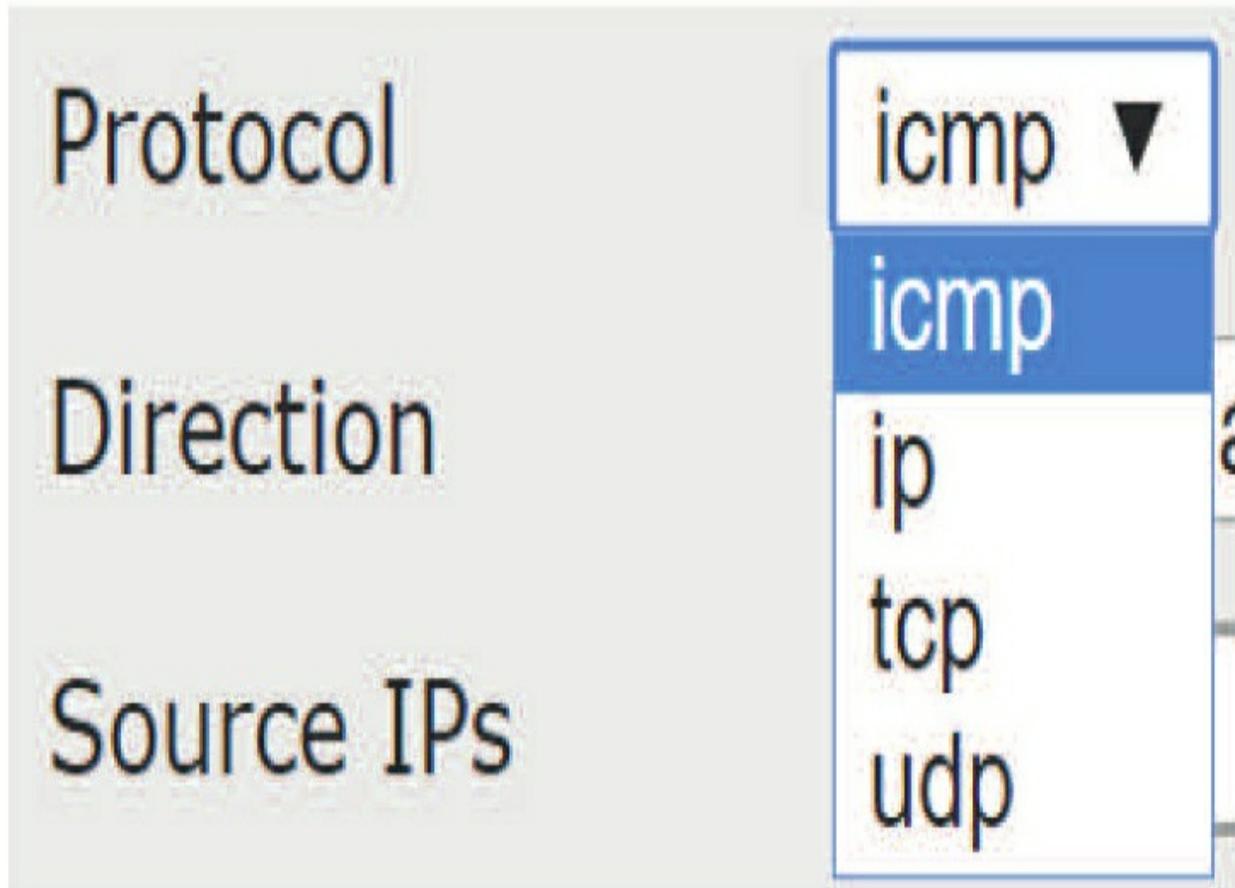
alert ▼

alert

pass

an action for the rule. The options are alert and pass.

The protocols you can use are IP, ICMP, TCP, and UDP.



Going back to the Classification setting again, this is an important part of the rules, and it has many, many options; shows just part of them.

Classification	A Client was Using an Unusual Port
	A Client was Using an Unusual Port
Action	A Network Trojan was Detected
	A Suspicious Filename was Detected
Protocol	A Suspicious String was Detected
	A System Call was Detected
Direction	A TCP Connection was Detected
	Access to a Potentially Vulnerable Web Application
Source IPs	An Attempted Login Using a Suspicious Username was Detected
	Attempt to Login By a Default Username and Password
Destination IPs	Attempted Administrator Privilege Gain
	Attempted Denial of Service
Detection Options	Attempted Information Leak
	Attempted User Privilege Gain
	Decode of an RPC Query
ack	Denial of Service
	Detection of a Denial of Service Attack
	Detection of a Network Scan
	Detection of a Non-Standard Protocol or Event
	Executable Code was Detected
	Generic ICMP Event

PAY ATTENTION! You MUST know what a pass rule is! To elaborate, an intrusion rule is a set of keywords and arguments that the system uses to detect attempts to exploit vulnerabilities on your network. As the system analyzes network traffic, it compares packets against the conditions specified in each rule. If the packet data matches all the conditions specified in a rule, the rule triggers.

- If a rule is an **alert** rule, it generates an intrusion event.
- If it is a **pass** rule, it ignores defined source or destination traffic. You can create pass rules in order to prevent packets defined in the pass rule from triggering the alert rule in specific situations, this can be better than disabling the alert rule. By default, pass rules override alert rules.
- For a **drop** rule in an inline deployment, the system drops the packet and generates an event.

So, to double check that you were paying attention, answer this question:

What do you think you can put in place in your IPS rules to exclude traffic from one host or group of hosts that is being blocked by an intrusion rule, but you want all other traffic to continue on with the default action of the rule? I'll go over the answer to this in the IPS policy in chapter 12.

When you create a custom intrusion rule, the system assigns it a unique rule number, which has the format GID: SID: Rev. The elements of this number are:

GID

Generator ID. For all standard text rules, this value is 1 (Global domain or legacy GID) or 1000 - 2000 (descendant domains). For all shared object rules you save as new, this value is 1.

SID

Snort ID. Indicates whether the rule is a local rule or a system rule. When you create a new rule, the system assigns the next available SID for a local rule. SID numbers for local rules start at 1000000, and the SID for each new local rule is incremented by one.

Rev

The revision number. For a new rule, the revision number is one. Each time you modify a custom rule the revision number increments by one.

Before we finish the chapter, I want you to understand that the Firepower system can create two types of intrusion rules—shared object rules and standard text rules—but you can only create one.

Shared object rules can only be created by the Cisco Intelligence Group (Talos). These are used to detect attacks against vulnerabilities in ways that traditional standard text rules cannot. You can delete instances of a shared object rule that you create, but you cannot delete shared object rules created by Talos.

When the rest of us write our own intrusion rules, we create a standard text rule. When writing a rule, understand that the most successful rules target traffic that may attempt to exploit known vulnerabilities rather than specific known exploits.

By writing rules and specifying the rules' event messages, you can more easily identify traffic that indicates attacks and policy evasions.

Summary

In this chapter we finally got into the meat of the Firepower system, objects. Firepower uses objects, which are reusable configuration components, to provide an easier way to use values across policies, searches, reports, dashboards, etc. You saw how understanding the objects themselves requires an understanding of how they are used within the rest of the system.

We went through the list of object types as shown in the user interface. Some descriptions for well-known object types were rather brief, but there was a more detailed discussion for other object types.

Toward the end of this chapter, I described several object types that apply only to Firepower Threat Defense and that relate to some of the legacy ASA routing and VPN features.

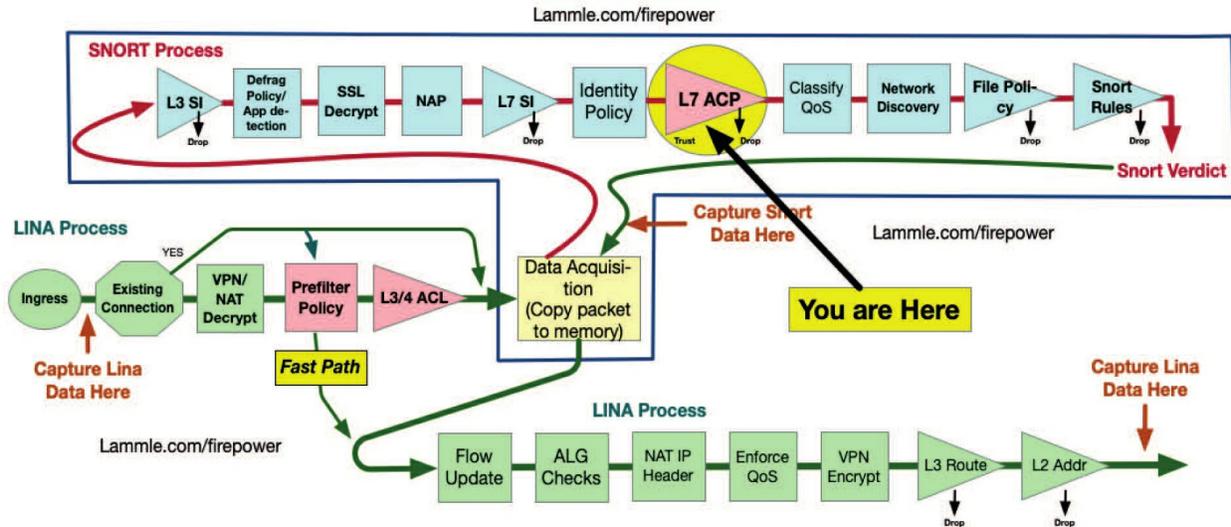
Chapter 9: Access Control Policy

The following CCIE/CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.2 Configure these policies in Cisco Firepower Management Center

2.2.a Access control



If there's one policy that's arguably at the heart of the Firepower system, it's the Access Control policy. Virtually all traffic entering a Firepower or FTD device will be processed through one or more rules in this policy. Firepower's central traffic cop!

The Access Control policy determines which traffic will be logged, allowed, or blocked by rules you create—the main source of logging that's found in **Analysis>Connection** Events. The ACP is also used to implement Security Intelligence lists, IPS rules and File policies as well as SSL, Identity, DNS, and your Network Analysis policy (NAP).

All traffic passing through a device is processed through the AC policy. If you've had experience with packet filtering firewalls, this technology will feel familiar to traditional firewall access control lists (ACLs), except that in addition to allowing/blocking connections based upon IP address, port, and protocol, the AC policy adds new ways to inspect traffic, greater flexibility, and unmatched visibility via Firepower's GUI.

Speaking to that flexibility factor, you can deploy your Access Control policies in a number of ways by opting to create a single policy for deployment to all your devices or going with a separate policy for each device or group of them. Applying a policy replaces any existing policy with the new one, and you can only apply one policy to a device at a time.

After we spend some time going deep into all things AC policy, we'll walk

through a real-life sample of AC policy rules at the end of this chapter.

To find exam study material such as videos, downloadable supplemental material, and practice questions, head over to www.lammle.com/firepower.

Overview

While it's true that the Access Control (AC) policy is the central hub, it's definitely dependent on several other policies. The figure below reveals almost all of the policies that are configured in the ACP.

Check out the legacy GUI on the left compared with the same screen with the new Light Theme on the right—pretty snappy upgrade, which to me is easier to navigate *most of the time!*

Overview Analysis Policies

Access Control  Network Disco

- Access Control
- Intrusion
- Malware & File
- DNS
- Identity
- SSL
- Prefilter

Use Smart License Status

Policies Devices

- Access Control
- Access Control
- Intrusion
- Malware & File
- DNS
- Identity
- SSL
- Prefilter

Let me quickly brief you on each of the policies shown in the preceding figure here:

- **Access Control policy:** The main and mandatory policy. All policies listed below it are configured into the AC policy.
- **Intrusion policy:** Certain rules in the AC policy can send traffic to the Snort engine for inspection via an Intrusion policy.

- **Malware & File policy:** Traffic can also be routed for file and malware inspection through one or more AC rules.
- **DNS policy:** Each AC policy has a DNS policy associated with it that inspects traffic prior to processing by AC rules and can be configured to allow, block, or sinkhole DNS traffic.
- **Identity policy:** AC rules can be created to match traffic based on user identity. For these rules to function, an Identity policy is associated with an AC policy.
- **SSL policy:** To carry out SSL decryption, this has to happen before processing by the AC policy. It's key, especially when it comes to identifying applications within encrypted sessions.
- **Prefilter policy:** Prefilter rules may block or trust traffic prior to the AC policy. Plus, we can assign tunnel zone names to assorted types of tunneled traffic and then create AC rules to deal specifically with them.
- **Network Analysis policy:** You won't find this one in the figure because it's kind of a hidden policy. The AC policy specifies one or more Network Analysis policies with settings that impact Snort preprocessors and subsequent Intrusion policy processing.

Through the Access Control policy, Firepower can invoke one or more of the above policies to process network traffic. Here's an example: We could decrypt an SSL session, identify the application protocol, then inspect the application traffic via Snort and Advanced Malware Protection (AMP) before allowing it to proceed to its destination.

Think of it like this... Once all the dependent/interdependent policies are configured, they're then basically glued together with the Access Control policy. Each FTD device can only have a single AC policy deployed to it at a time, but a single policy can be targeted for one or more devices. So, you can choose to deploy one policy per device or assign the same policy to multiple devices. If your deployment is bare bones to basic and you're simply performing intrusion inspection on all traffic, then a single AC policy across

all your devices is probably all you need. It all comes back to that key adage “Complexity is the enemy of security.” Plus, using fewer policies will make configuration

and maintenance a whole lot easier down the line!

Policy hierarchy can really help us to keep settings consistent between policies. By using a master/child policy design, you can enforce certain settings in your child policies by defining them in the master. You can even allow child policies to override various settings from the master.

Policy Creation

So, let’s dive right into Access Control policy now! To get started, navigate to **Policies>Access Control>Access Control**. There’s no default AC policy.

To create one, click the **New Policy** button and you’ll be presented with the dialog shown here

Here’s the place to insert your policy name with an optional description and select a base policy if you want one. If this is your first policy, leave the Select Base Policy option at None. Next up are three choices for your policy’s default action:

New Policy

Name:

Description:

Select Base Policy:

Default Action: Block all traffic Intrusion Prevention Network Discovery

Targeted Devices

Select devices to which you want to apply this policy.

Available Devices

- 4140-HA
- Firepower_Appliance
- vFTD-19
- vFTD20

Selected Devices

- **Block All Traffic:** If traffic doesn't match any AC rules, it won't be allowed to pass through the device. It'll insert a "*deny ip any any*" at the end of your rule set.

- **Intrusion Prevention:** Traffic that doesn't match any rules will be processed through Snort rules in an Intrusion policy. This inserts a "*permit ip any any*" at the end of your rule set. Here's an example from an FTD CLI output—notice the permit IP any any:

```
access-list CSM_FW_ACL_line 10 advanced permit ip any any rule-id 268434432 (hitcnt=0)
0xa1d3780e
>
```

- **Network Discovery:** Traffic that doesn't match any rules will be allowed through the device and may be subject to analysis via the Network Discovery policy. I've never used this option and I really don't recommend it because there are better ways to perform an IDS configuration.

I need to point out that there are a couple more options that aren't available on this screen. The first one is "which Intrusion policy will be used." and the

second is the “trust action.” Both of these can be configured from the main policy configuration screen.

Okay—so after configuring the settings above, you have the option of targeting the policy for one or more devices. Select one or more devices on the left and click the **Add To Policy** button.

If you target a device that already has an AC policy defined, when you click Save, you’ll see a dialog similar to this:

Error

 Following devices already have assignments listed below. These devices will be reassigned to current policy

device:ASA-5515 FTD - policy:ASA-5515 Inline

 Do you want to continue with above changes?

Yes

No

Clicking Yes will move the device assignment from the existing policy to your new policy. No worries here; if you actually meant to remove an existing device from another AC policy and add it to your new one, this isn't really an error.

This behavior stems from the fact that Firepower won't allow you to configure a device with no Access Control policy. And as a result, the system won't let you just remove a device from an AC policy. The only way to do so is to add the device to another AC policy, which will remove it from the old policy assignment. Call it a "positive handoff" if you like. It results in a message similar to the one shown in the preceding figure.

After you have assigned your devices, click the Save button to begin editing your new policy.

Policy Editing

When you start editing a policy, the first screen you'll be presented with shows the empty rule set along with tabs for the various policy settings. This is displayed in figure below.

Before we dive straight into access control rules, I want to show you some of the other policy settings available. Again, this follows our strategy of ensuring that all the dependent settings and policies are in place before we get to the heart of the policy—the rules! It's really tempting to go straight into the AC rules, but we're going to make sure all our ducks are in a row first. Check out the following figure:

See Default Action on the very bottom? This was set when we were bringing our devices in. I set them to Intrusion Prevention, and if you look at your CLI output you'll see this is a "*permit ip any any*" at the end of your layer 7 access list, called the ACP rules.

After my network is up and running, I usually make this a "*deny IP any any*" rule, but whether or not I do that, depends on my customer.

General Settings

Let's take a minute to look at some of the settings across the middle of the Access Control policy page starting with Prefilter policy on the left. These settings aren't specific to the tab that's selected, and all the first three listings here are shortcuts to their location in the Advanced tab on the same screen.

Prefilter Policy

The quick link for Default Prefilter Policy on the top left, allows you to select a Prefilter policy, which is actually found in the Advanced Settings tab. This policy will be applied to the devices targeted by this Access Control policy. Even though this setting only applies to Firepower Threat Defense (FTD) devices, you still have to select a policy here. Every policy starts with Default Prefilter Policy selected and you can actually just leave this as is if you're not using FTD with the default Prefilter policy selected.

SSL Policy

This quick link on the top-middle of the page, allows you to select your SSL policy, which again is actually configured in the Advanced Settings tab. Of course, if you're not performing SSL decryption, you can leave this at the default setting of None. SSL decryption is performed after layer 3 SI inspection, but before traffic is processed from the L7 SI, and then through access control rules.

Identity Policy

Just as with Prefilter and SSL, this quick link allows you to select an Identity policy to deploy to your devices. The default is None.

Inheritance Settings

Clicking Inheritance Settings displays the dialog shown in the figure below. Inheritance allows us to nest multiple Access Control policies. You select a base policy and optionally select Child Policy Inheritance Settings. Leaving these boxes unchecked means child policies can accept or override the various settings in the base policy. By checking a box, you are forcing the child policy to use the settings for that particular area:

When policies are nested, rules in the base policy are inherited by the child

policies. The Mandatory and Base rule categories are placed at the beginning and end of the child policy rule set. More on this when we get to rules later in this chapter!

For now, know that policies can be nested up to 10 levels deep, and this can wind up as an administrative nightmare because the various rule layers are nested within each other. One of the most effective uses of policy nesting is to maintain consistency in the Advanced and Security Intelligence settings across your policies.

I also want you to be aware that these settings can tend to drift over time, so maintaining them in the single base layer is a smart way to keep all of your policies synchronized. Also, if you decide to modify these settings, you only have to do it in the base policy—nice!

Policy Assignments

The Policy Assignments link is what we use to modify the device assignments of the policy, and it's the same setting used when we first created the policy.

The next figure shows the Policy Assignments dialog:

Policy Assignments

Targeted Devices

Select devices to which you want to apply this policy.

Available Devices

- 4140-HA
- Firepower_Appliance
- vFTD-19
- vFTD20

Selected Devices

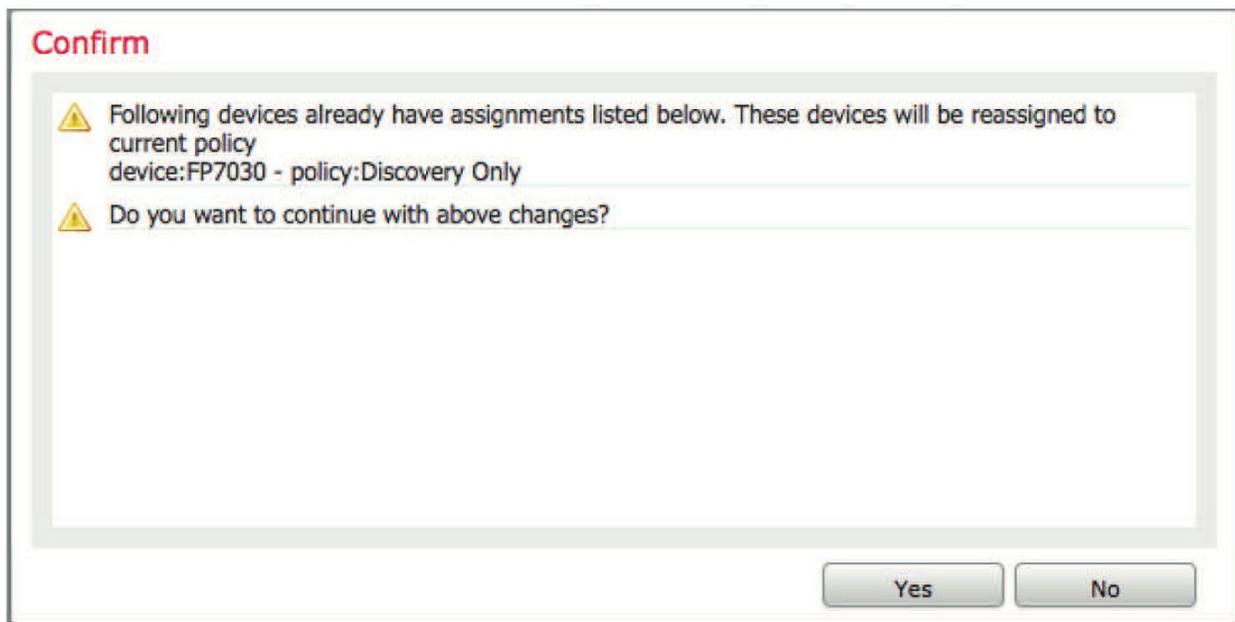
- 4140-HA
- Firepower_Appliance
- vFTD-19
- vFTD20

Add to Policy

Impacted Devices

Selected Devices reveals the devices where the policy is currently assigned. If this policy is used as a base for others, then the Impacted Devices section lists the policies that would be impacted by changing this one.

Because each device must always have an Access Control policy assigned to it, the system won't allow you to remove a device from a policy and then save that policy. To assign a different policy to a device, you have to edit your target policy and *pull* the device from the previous policy into the new one. When you do this, you'll get a confirmation dialog like the one shown below, but it's nothing to worry about. This is just confirming that you really meant to remove this device from the previous policy assignment:



Security Intelligence

We're going to take a second to review some of the other configuration tabs before we get into the rules. A really vital one to always consider is the Security Intelligence tab, covered in detail in Chapter 8, "Objects," but the AC policy is where the SI is actually implemented. This is the tab you'll use to configure the behavior of the various Security Intelligence lists.

The primary purpose of a Security Intelligence list is to blacklist or block something. This capability has been expanded over subsequent Firepower

versions and now includes DNS, URL, and IP address lists.

Blacklists and Whitelists

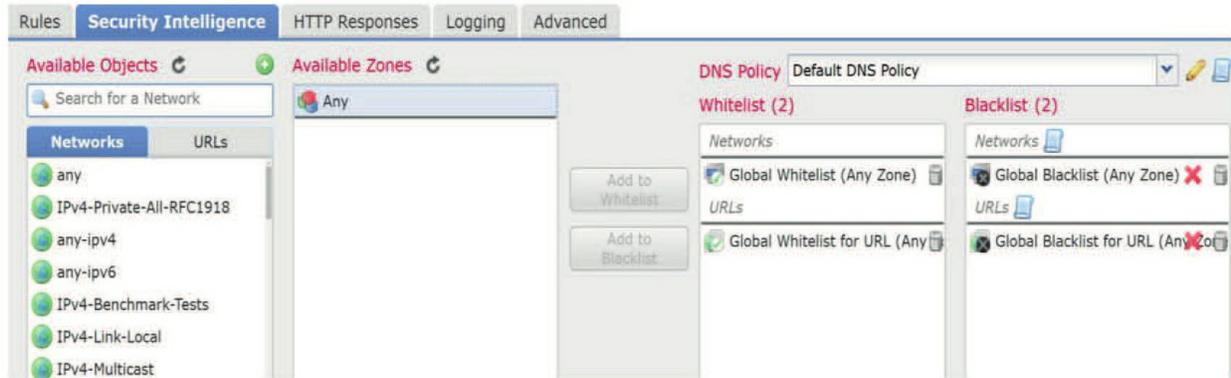
Using Security Intelligence, you have the option of adding a given DNS request, URL, or IP address to either a blacklist or whitelist. The blacklist is pretty self-explanatory—adding an entry to a blacklist means the system will generate a Security Intelligence event and optionally block the traffic. But what does a whitelist do?

The easiest way to think about a whitelist is that a whitelist entry really only exists to override a blacklist entry. The purpose of a whitelist is to prevent a blacklist entry from messing with critical traffic and ruining your day. Take an IP list for example... Just thinking about what would happen if the IP address of a core router on my network gets mistakenly blacklisted makes me kind of lose it. If my IPS started blocking all traffic to and from that router, it could easily make the entire network grind to a halt—game over! To prevent disasters like this, definitely add the IP addresses of key systems to a whitelist. The reason this works so well is that whitelist entries always override blacklist entries, so if there's ever a conflict, the whitelist will prevail.

Keep in mind that whitelist entries only override blacklist entries, so putting an entry into a whitelist doesn't mean Firepower will give that traffic a free pass. It just means Firepower won't blacklist the traffic, but it'll still be processed by the appropriate access control rules and could still undergo intrusion and/or file inspection. It could also be blocked by an AC rule.

The Security Intelligence Tab

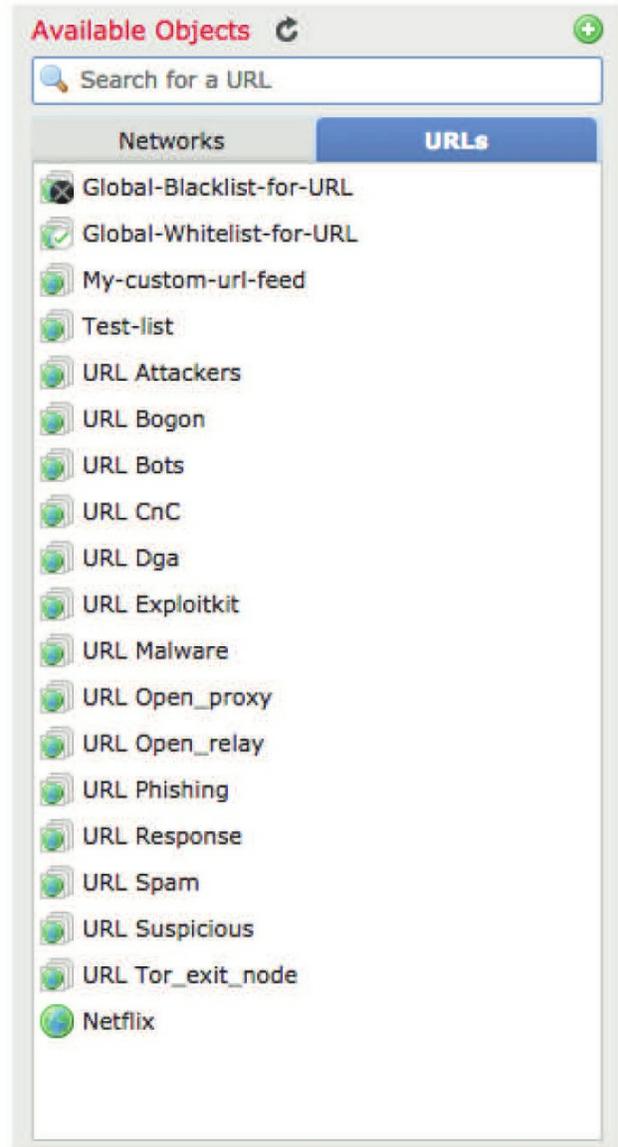
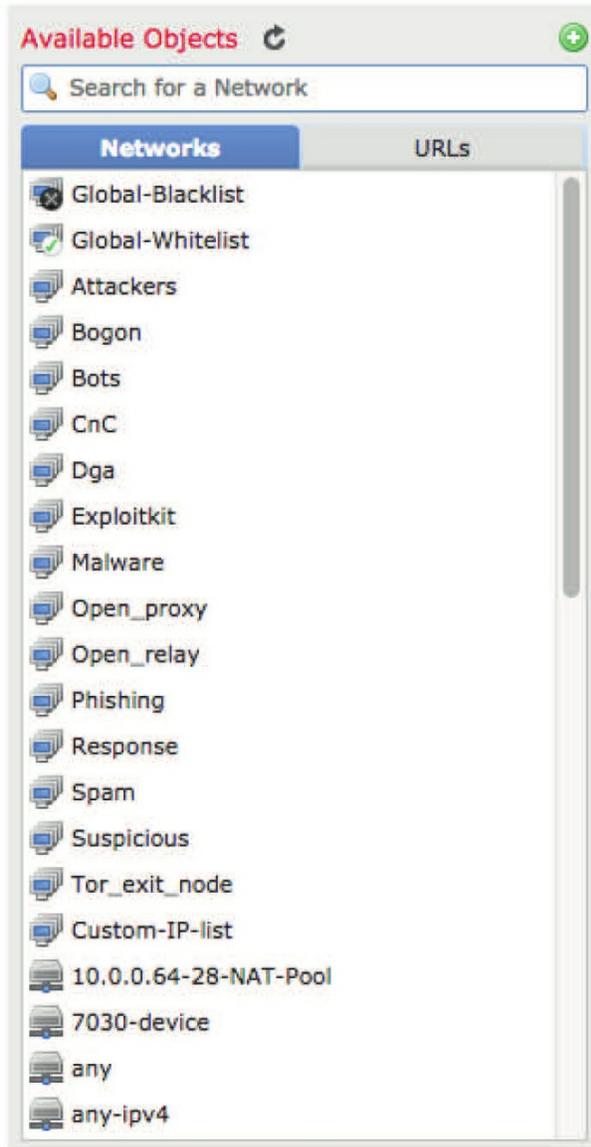
The Security Intelligence tab definitely looks a bit busy when you first see it:



On the left, you can see the Available Objects and Available Zones columns where you'll select your network or URL objects. You can also optionally add a zone qualification to your object, which will make sure the traffic matches the object (IP address or URL) and is also traveling to or from the selected zone. You're not forced to select a zone, and if you don't, the default zone of Any will be used. On the right, we have the Whitelist and Blacklist columns that are also divided horizontally into Networks and URLs sections. Adding a Network or URL object will cause it to appear in the appropriate section within each list.

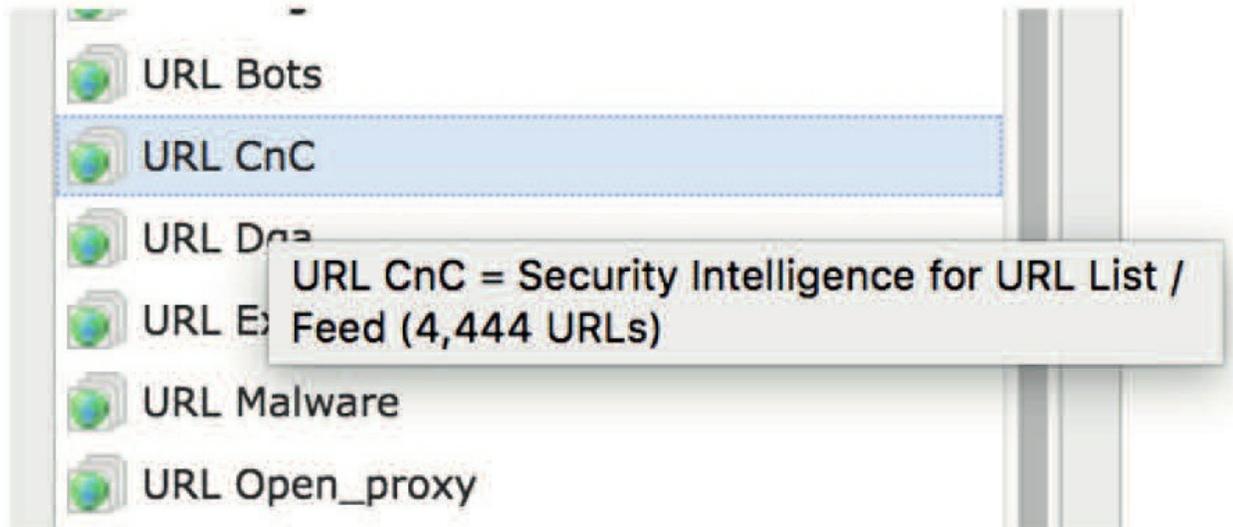
Available Objects

Let's take a closer look in the next two figures called the Available Objects column, which is further divided into two subtabs, one for networks and one for URLs:



Each of these lists contains various Security Intelligence categories populated as a result of the Cisco Security Intelligence feed your FMC downloads periodically to populate these categories.

Even though you can't check into the individual entries from here, you can see how many entries are in a given category. Hovering your mouse over a category will cause a pop-up to appear listing the number of entries in the category:



Keep in mind that this feed is updated constantly, so this count is likely to change within minutes or hours, depending on the update frequency delimited in the intelligence feed object. If this is a bit blurry to you, head back to Chapter 8 to clear things up.

The Available Objects column will also show all of your custom Network or URL objects as well as custom feeds. Any of these can be added to either the Whitelist or Blacklist column by selecting them and clicking either Add to Whitelist or Add to Blacklist. You'll also notice the green plus icon above the search field. This allows you to add a Security Intelligence feed or list on the fly without having to navigate back to the Objects menu. If you do add an object in this

manner, use the refresh icon (



) to refresh the list and show your new entry.

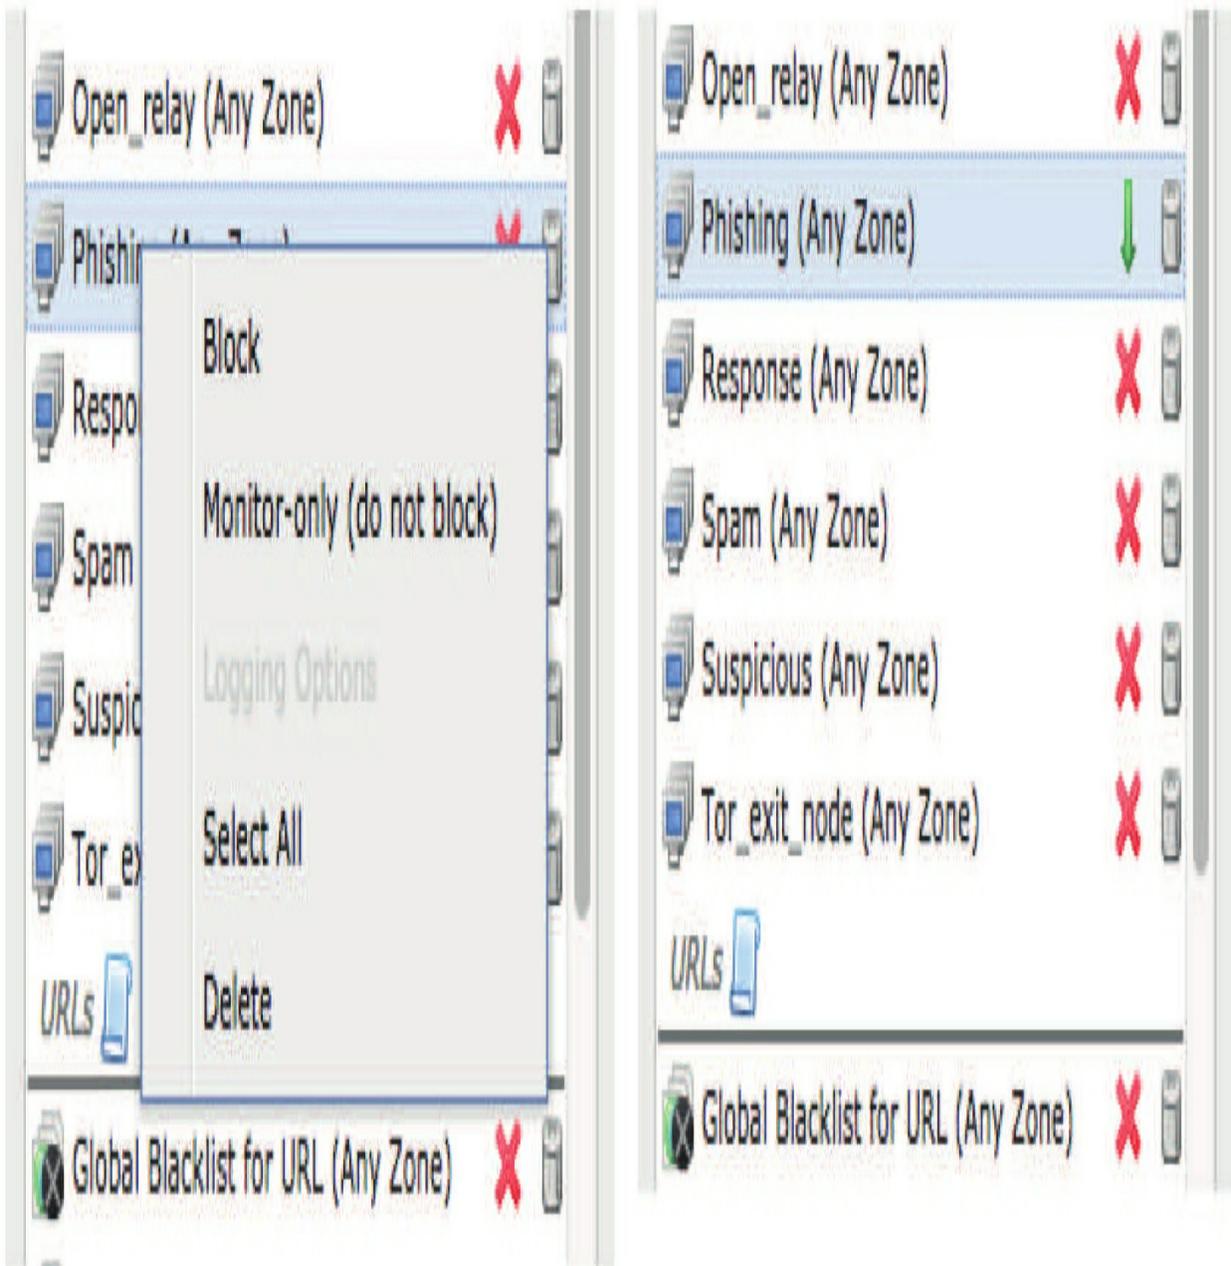
Blacklist Actions

When you first add entries to your Blacklist column, they'll default to a drop action indicated by the red X next to the list. For some lists you might not want the connection blocked and just want to be notified about it instead.

Right-clicking on an entry will bring up a context menu that'll allow you to

select Monitor-Only (Do Not Block). Doing this will change the action of the list to logging an event without blocking the traffic. I really don't recommend doing this because, well, these are 100% guaranteed evil, so why would you let them in just to watch them break your stuff and steal things?

But if you do decide to do this for reasons unknown to sane people, the icon will change to a green arrow as shown below:



Logging Options

There are three blue scroll icons on the page that control the logging for each type of Security Intelligence list, and they're pictured in the figure below. The default setting for each type of Security Intelligence is to log the connection. This means any matching traffic will be logged as Security Intelligence events.

These can be viewed by navigating to **Analysis > Connections > Security Intelligence Events**.

Check out the next figure:

DNS Policy Default DNS Policy

Whitelist (4)

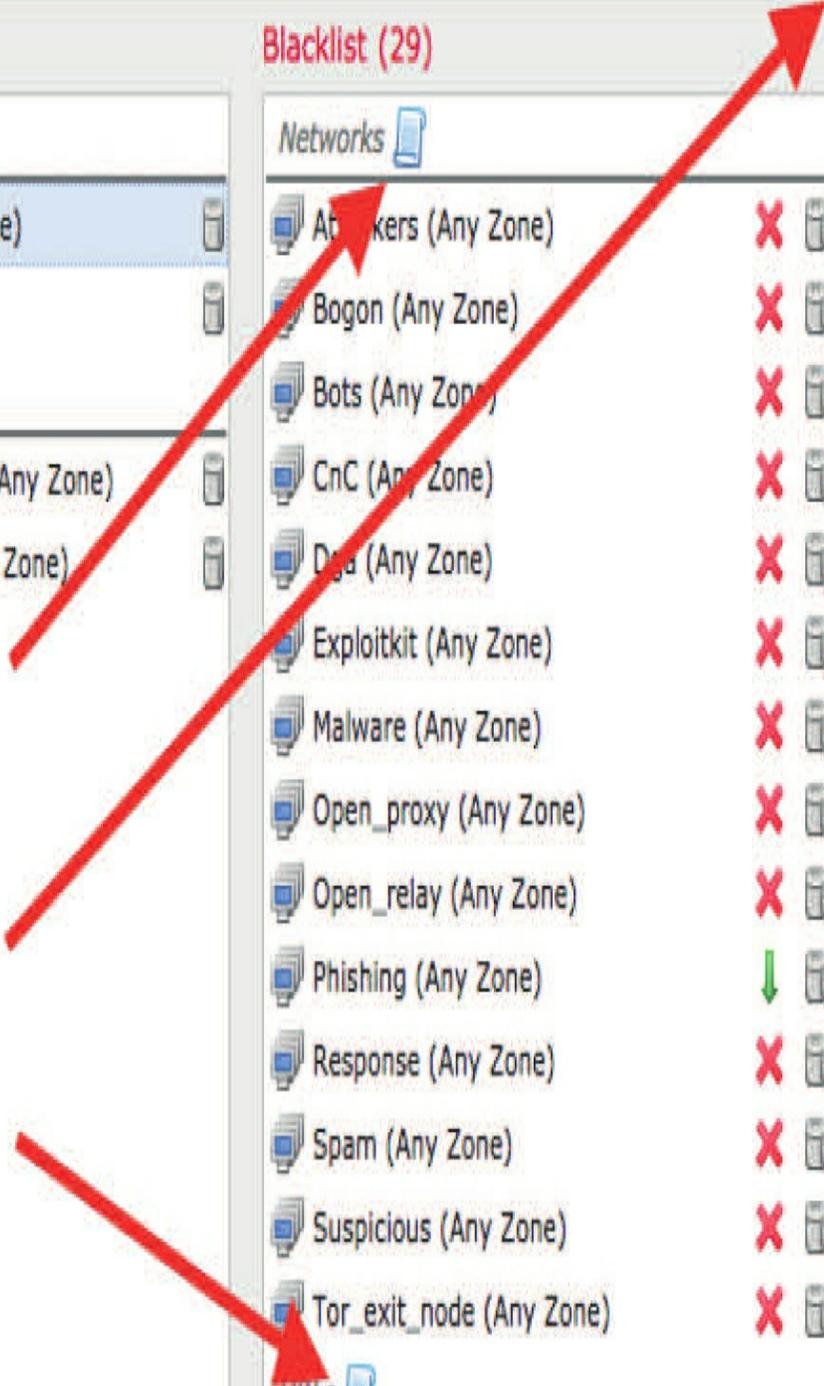
Networks	
Global Whitelist (Any Zone)	
Critical-IPs (Any Zone)	

URLs	
Global Whitelist for URL (Any Zone)	
My-custom-url-feed (Any Zone)	

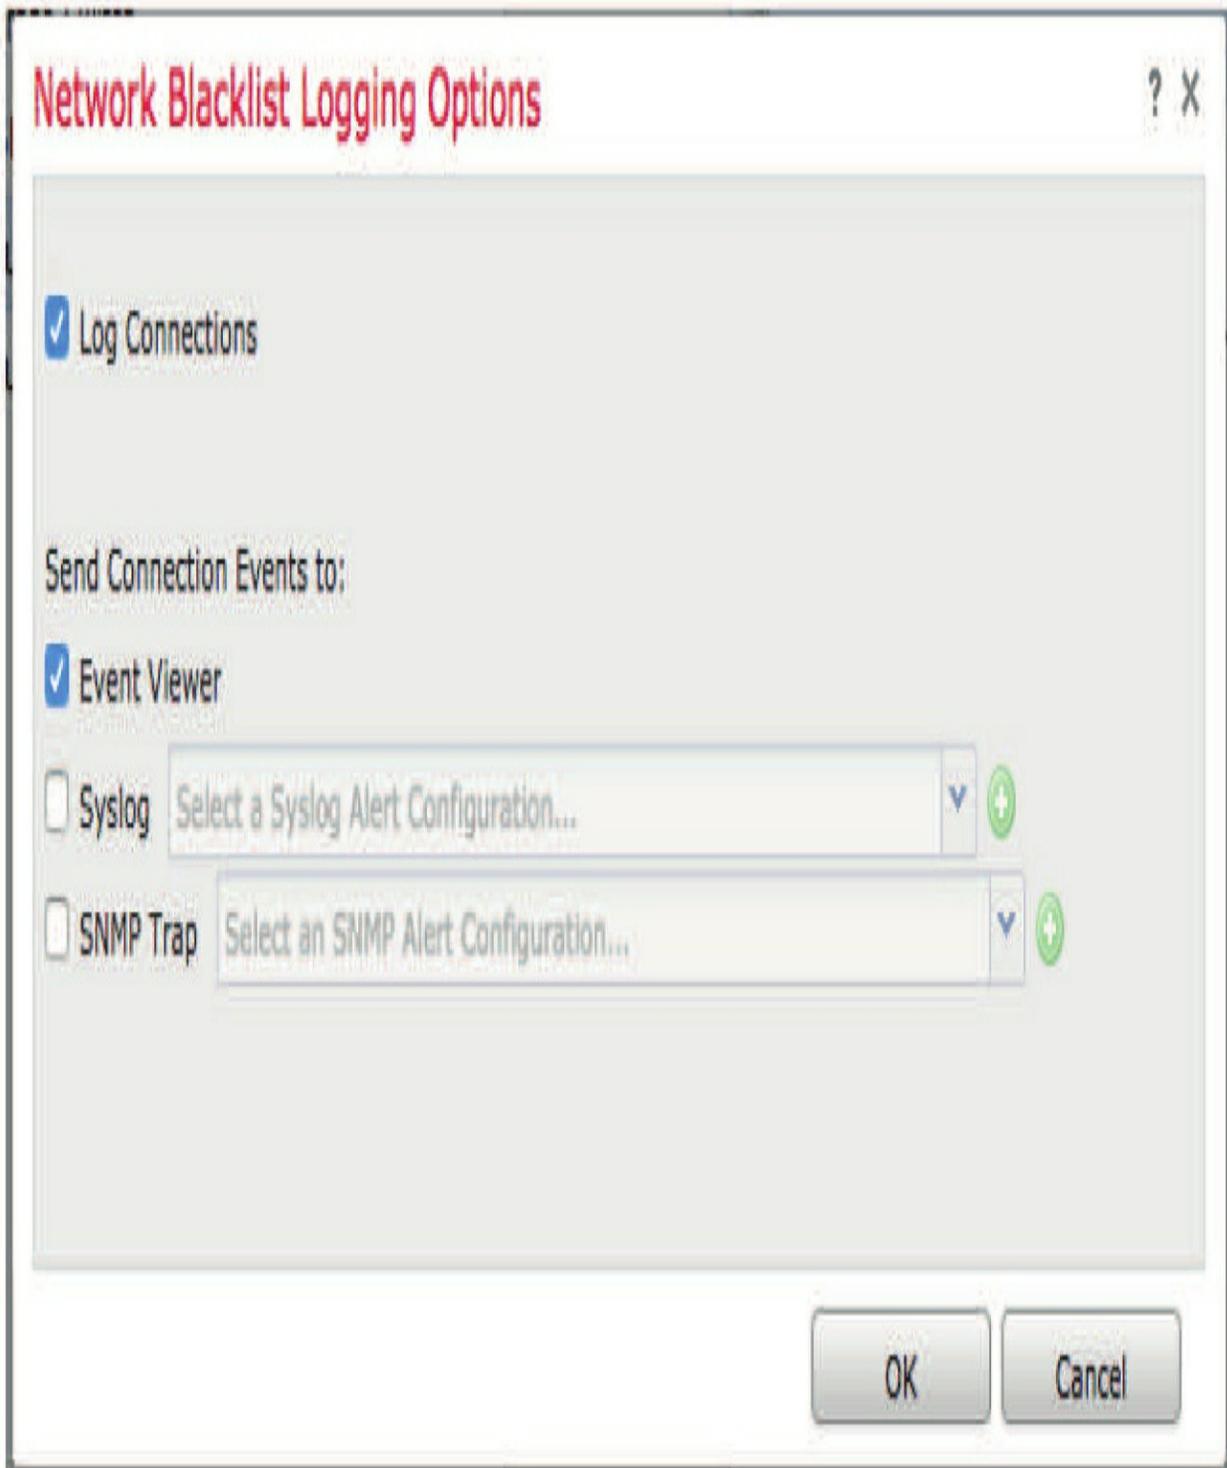
Blacklist (29)

Networks	
Attackers (Any Zone)	
Bogon (Any Zone)	
Bots (Any Zone)	
CnC (Any Zone)	
Dga (Any Zone)	
Exploitkit (Any Zone)	
Malware (Any Zone)	
Open_proxy (Any Zone)	
Open_relay (Any Zone)	
Phishing (Any Zone)	
Response (Any Zone)	
Spam (Any Zone)	
Suspicious (Any Zone)	
Tor_exit_node (Any Zone)	

URLs	
Global Blacklist for URL (Any Zone)	
Global-Blacklist (Any Zone)	



Clicking on any of the icons displays the logging options dialog, shown next:



This allows you to disable logging if you want, or direct the logging to a syslog or SNMP destination.

DNS Policy

DNS Security Intelligence is controlled by the DNS policy, which we'll explore thoroughly in Chapter 13. For now, though, I do want to quickly preview a couple of things about it.

Selecting a DNS policy is as easy as clicking the DNS Policy drop-down and selecting your previously created policy as you can see here:

Prefilter Policy: Default Prefilter Policy

SSL Policy: None

Identity Policy: None

 Inheritance Settings |  Policy Assignments (8)

Rules Security Intelligence HTTP Responses Logging Advanced Settings

 Filter by Device  Show Rule Conflicts  Add Category  Add Rule

Name	Sou...	Des...	Sou...	Des...	VLA...	Users	App...	Sou...	Des...	URLs	Sou...	Des...	Act..	
#..														      
▼ Mandatory - VFMC20-ACP (-)														
There are no rules in this section. Add Rule or Add Category														
▼ Default - VFMC20-ACP (-)														
There are no rules in this section. Add Rule or Add Category														
Default Action										Intrusion Prevention: Balanced Security and Connectivity				

You can also edit the DNS policy directly from here by clicking the pencil icon, which will open your DNS policy in a new browser tab.

HTTP Responses

The HTTP Responses tab has not changed since version 5 of the Sourcefire System was first introduced. This tab is used to configure the web page returned when a user attempts to visit a blocked website. When using the URL Filtering feature, it's generally a good idea to present the user with a web page explaining why a particular site is blocked. The alternative—simply dropping their HTTP request—will usually leave the user wondering if there is something wrong with their Internet connection. Attempting to visit a page and experiencing a time-out or an unfriendly browser error message will just increase help desk calls. So instead, Firepower can return a web page notifying the user that their attempt was blocked.

There are two available options for configuring the HTTP response on this tab: Block Response Page and Interactive Block Response Page:

Rules

Security Intelligence

HTTP Responses

Logging

Advanced

Block Response Page

This page will be displayed when HTTP traffic is blocked.

None



Interactive Block Response Page

This page will be displayed when HTTP traffic is blocked, but the user may choose to continue.

System-provided



None

System-provided

Custom...

Block Response Page

The Block Response page will be returned anytime a user's HTTP request is blocked by a rule using the Block or Block with Reset action. For this page, you've got three options: None, System-Provided, or Custom. The default is None.

This means that if you use an access control rule to block a user's HTTP request, they won't get any information about why the page won't display. If the rule does not include the reset action, the user will probably see their browser spin for a while before returning a browser-specific message telling them the site couldn't be loaded.

So better, you can send the user a generic notification that the site was blocked or return a customized response page instead. The default system-provided response page is shown below:



Access Denied - Mozilla Firefox



File Edit View History Bookmarks Tools Help



http://www.cnn.com/



Google



Access Denied

You are attempting to access a forbidden site.

Consult your system administrator for details.

Done

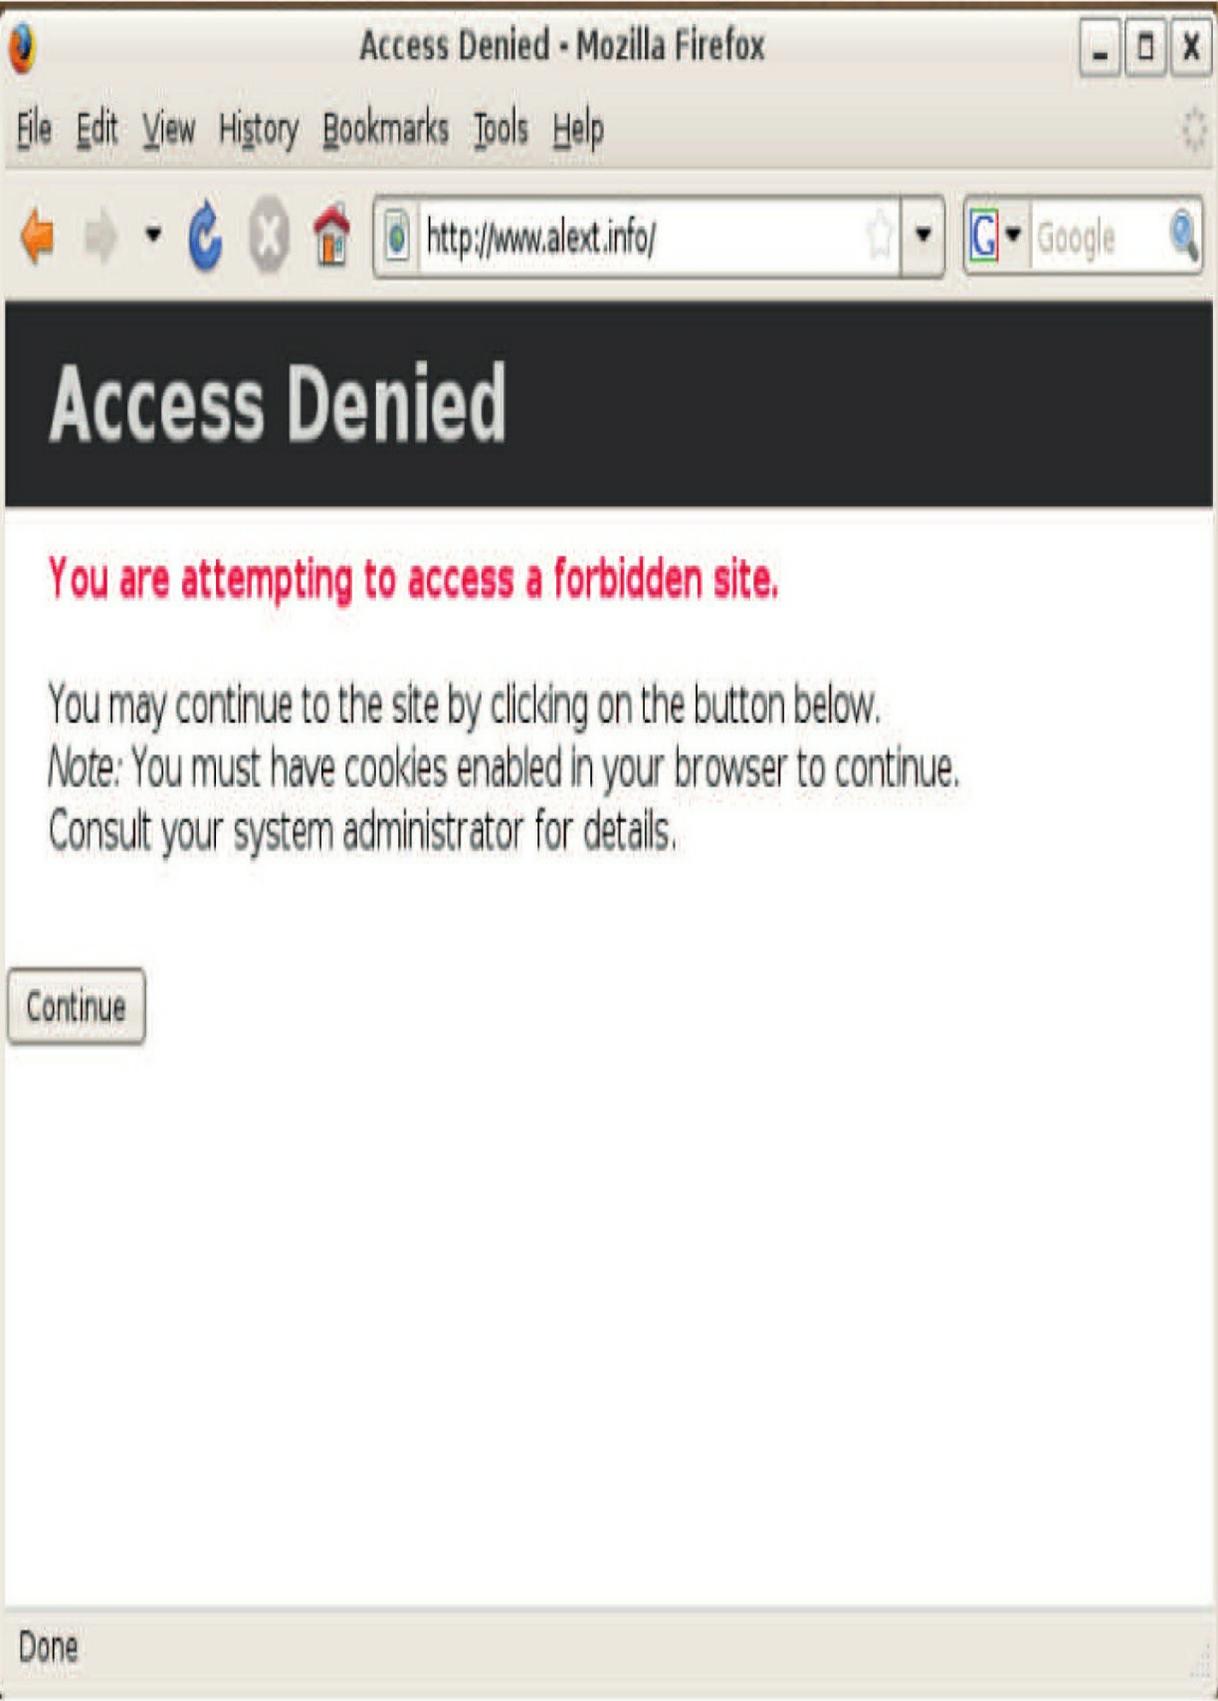
To select the standard response page, choose System-Provided from the pull-down list. To view the page source, click the magnifying glass icon next to the field.

To customize the HTTP response, select Custom. This will load the existing page—system-provided by default—and allow editing or pasting the HTML code you choose.

Interactive Block Response Page

The Interactive Block Response Page will be returned anytime a user's HTTP request is intercepted by a rule using the Interactive Block or Interactive Block with Reset action. This page functions a lot like the Block Response Page above except that the interactive rules give the user the option of proceeding to the site anyway. The Interactive Block Response page includes a JavaScript-driven button that if clicked allows the user to proceed to the blocked website.

The default System-Provided Interactive Block Response page is shown here:



Access Denied

You are attempting to access a forbidden site.

You may continue to the site by clicking on the button below.
Note: You must have cookies enabled in your browser to continue.
Consult your system administrator for details.

Continue

Done

Once the user clicks the Continue button, they can load the blocked website. By default, the duration of this access is 600 seconds or 10 minutes. After this period, the interactive block page will reappear.

Note: Without the use of SSL decryption, neither of the HTML response pages will work if the site blocked is using HTTPS, unless you've configured your SSL policy, and SSL decryption is occurring.

Advanced

The Advanced tab has boatloads of settings that can be modified as you can see in the figure below.

Rules	Security Intelligence	HTTP Responses	Logging	Advanced
General Settings		Intelligent Application Bypass Settings		
Maximum URL characters to store in connection events	1024	Intelligent Application Bypass Settings	Off	
Allow an Interactive Block to bypass blocking for (seconds)	600	Total Apps and Filters Configured: All applications including unidentified applications		
Retry URL cache miss lookup	Yes	Transport/Network Layer Preprocessor Settings		
Enable Threat Intelligence Director	Yes	Ignore the VLAN header when tracking connections		
Inspect traffic during policy apply	Yes	No		
Identity Policy Settings				

Settings are grouped into various sections on the page, and most of them can be left at their defaults. Several, like the Identity, SSL, and Prefilter policies, can be set at any time by clicking the links near the top of the page.

General Settings

The General Settings section is a catchall of settings that really don't fit into the other categories. Check it out:

General Settings

Maximum URL characters to store in connection events

Allow an Interactive Block to bypass blocking for (seconds)

Retry URL cache miss lookup



Enable Threat Intelligence Director



Inspect traffic during policy apply can be set to one of the following: Enabled [inspect]: A few packets might drop (depending on device load and capacity). Disabled [do not inspect]: Traffic is dropped or not (depending on how the device is configured).



Revert to Defaults

OK

Maximum URL characters to store in connection events Default: 1024
Maximum: 4096

This defines the maximum length of the field used to store URLs in connection events for HTTP. Setting this value to zero would clearly disable storing of URLs in connection events, but you might just want to go with reducing or disabling this setting to improve performance if your system logs tons of HTTP connections.

Allow an Interactive Block to bypass blocking for (seconds)
Default: 600 (10 minutes)

Maximum: 31536000 (365 days)

This is the default time during which a user can access a URL that triggered an Interactive Block rule, and the clock starts ticking when the user clicks the Continue button. Setting this value to zero requires the user to bypass the block using the Continue button for each request.

Retry URL cache miss lookup
Default: Checked

When this is disabled, the system will immediately allow traffic to a URL when the URL category isn't in the cache. A cloud lookup to determine the URL category then takes place in the background. The URL is treated as uncategorized until the cloud lookup completes with a different category. In short, unchecking this will result in a faster user experience at the cost of some accuracy and timeliness in URL categorization. Keep in mind that this only applies if the URL Filtering license is in use.

Enable Threat Intelligence Director
Default: Checked

When enabled, it allows data from Cisco's Threat Intelligence Director to be passed to your managed devices. Threat Intelligence is configured under the Intelligence>Settings tab.

Inspect traffic during policy apply
Default: Checked

When policies are deployed, the Snort process may either be reloaded or restarted, depending on the type of changes in the policy. When Snort is reloaded, traffic will continue to pass and will continue to be inspected. There may be some increased packet latency as the system works to reload Snort and inspect traffic at the same time, but no packets will pass uninspected.

Disclaimer here: When Snort is restarted, traffic will pass through the device uninspected between the time Snort is stopped and the time the Snort process is up and running again. This is normally just a matter of seconds, but there will be *some* traffic that passes through the device uninspected.

Anyway, this setting allows you some control over how this restart/ reload takes place. When you leave this setting at its default state of checked, the Snort process will be reloaded if possible, but certain policy setting changes will still cause Snort to restart.

If you uncheck this setting, Snort will always restart when policies are deployed. The effect is that there will always be a few packets that'll squeeze through uninspected, but at least you won't experience the increased latency that'll be the result of a Snort reload.

What Are All the Possible Things That Restart Snort?

The list of policy changes that cause a Snort restart is a long one, and there's no reliable rule of thumb that'll help determine if a given policy change will cause a reload or a restart.

But I'm guessing you still want to see the list, right? Okay, here goes:

1. Add a custom Network Analysis policy to AC policy
2. Add URLs the first time to an AC rule
3. Add/delete a File policy to/from an AC rule
4. Add/delete Intrusion policy to/from an AC rule
5. Deploy a new AC policy
6. Enable/disable adaptive profiles in AC policy
7. Enable/disable Identity policy in AC policy
8. Enable/disable SSL policy in AC policy

9. Make changes to default values under File and Malware Settings
10. Change IMAP preprocessor depth in Network Analysis policy
11. Change POP preprocessor depth in Network Analysis policy
12. Change SMTP preprocessor depth in Network Analysis policy
13. Enable/disable Inspect Archives in File policy
14. Enable/disable Store Files in File policy
15. Add custom DNS policy to AC policy
16. Modify the MTU size on device interfaces
17. Activate or deactivate an existing application detector
18. Create a new application detector
19. Install a vulnerability database (VDB) update
20. Install a Snort rule update (SRU)
21. Deploy new shared object rules in the Intrusion policy
22. Create an identity rule with an action of passive authentication
23. Add/remove a URL category/reputation condition in an AC rule
24. Revert to default values under File and Malware Settings on the Advanced tab of AC policy
25. Restore a single default value under File and Malware Settings on Advanced tab of AC policy

Now remember, this list is definitely not exhaustive because there are likely several other modifications, especially in the Network Analysis policy, that'll precipitate a Snort restart.

The setting you decide on here is largely a trade-off of detection versus connectivity. To get the best detection at the expense of a little latency during policy deployment, keep the default setting. To ensure minimum latency at the cost of some detection, disable the setting.

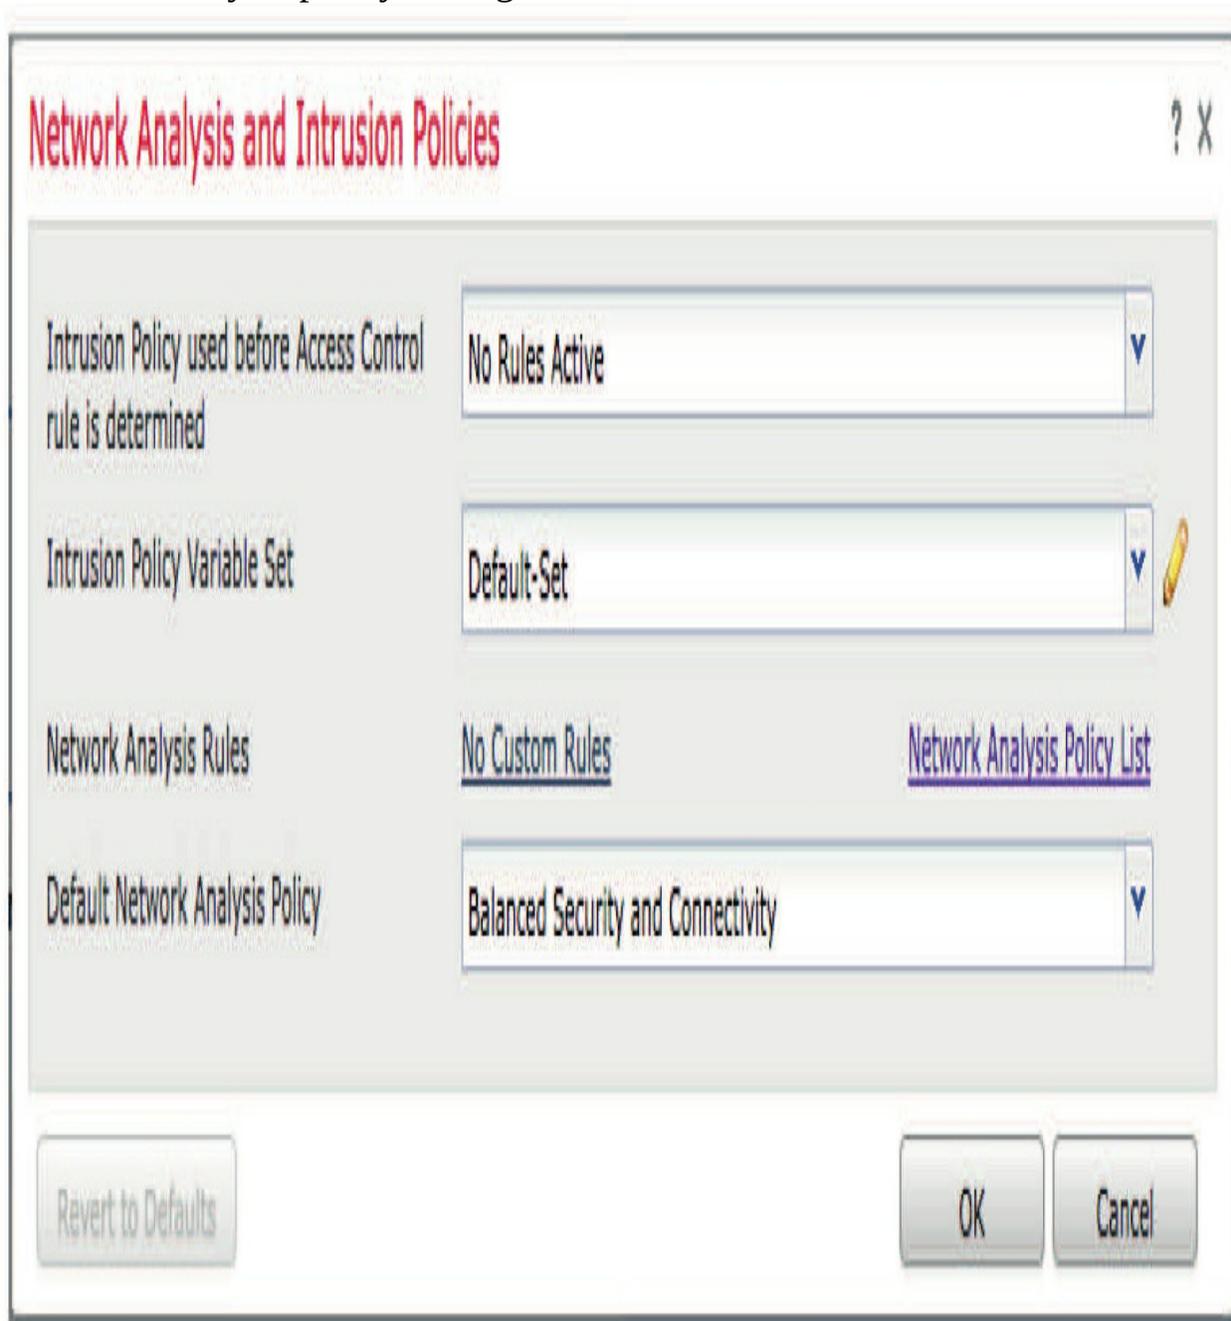
Identity, SSL, and Prefilter

These three are pretty straightforward and each one offers a simple drop-down menu to select the associated Identity, SSL, or Prefilter policies.

Network Analysis and Intrusion Policies

These settings affect how intrusion rules are processed as well as which

Network Analysis policy settings are in effect:



Intrusion Policy used before Access Control rule is determined Default:
(Depends on default action)

This setting determines what'll happen to traffic that matches an

AC rule containing an application or URL condition. The challenge in application identification is that it can't be accomplished by inspecting just a

single packet. The number of packets required depends on the application detector, but the system must allow at least some of the traffic to pass until it can decide whether a given application is present. Once the application has been identified, the AC rule or rules designed to take an action will then be enforced.

This setting determines what happens to the packets allowed to pass prior to application detection. The default setting just allows them to pass using the No Rules Active Intrusion policy, which is essentially an empty Intrusion policy.

I'm going to suggest that you go ahead and perform intrusion inspection on this traffic, and to do that, just use the drop-down to select an Intrusion policy with your preferred rule set.

There's another factor here when you're deciding on which rule set to use. Say you've created a custom Intrusion policy that you plan to use in your AC rules—we'll call it My IPS Policy. When you deploy policies to your devices, your custom Intrusion policy will be applied, but if you choose a different policy, like Balanced Security and Connectivity, in this setting, it'll result in two intrusion policies being applied to your devices and each of them will use system resources (CPU and memory). To increase the efficiency of your policy and reduce some overhead, only apply as many different intrusion policies as you really need. This is why I recommend using an Intrusion policy you're already using in your AC rules for this setting. In the example here, you'd go with My IPS Policy from the drop-down.

The default setting for this item depends on the policy default action you selected when you first created this AC policy. If you chose Intrusion Prevention, this'll be set to Balanced Security and Connectivity, but if you went with block or network discovery, it'll be set to No Rules Active.

Intrusion Policy Variable Set

Default: Default-Set

This is the variable set object that'll be used in conjunction with the Intrusion policy immediately above it. Recall from Chapter 6 that this is what

determines the values for the variables used by Snort in intrusion rules.

Network Analysis Rules and Default Network Analysis Policy Default: Balanced Security and Connectivity

These two settings are closely related, so I'm going to cover them together. One way to look at the Network Analysis policy is that it contains advanced Snort settings. Most of the settings in this policy shouldn't be changed from their defaults because doing this can negatively impact detection, performance, or both! But there are a few settings that Cisco does recommend adjusting, which is pretty much based upon whether the device will be used in passive or inline mode. And when I say inline here, I'm also including transparent, switched, or routed mode.

We spent a lot of time in Chapter 23 going over how the Network Analysis policy works and warning you about making modifications, so here's my basic guidance for the Network Analysis policy:

- Start with the Talos base policy that corresponds to your Intrusion base policy: Connectivity, Balanced, Security.
- Enable inline mode and the inline normalization preprocessor if your device is inline, including switched or routed.
- Restrain yourself from making other changes unless you know *exactly* how they will impact detection and performance!

What I mean here is if your device uses only inline interfaces, you probably only need one Network Analysis policy for that device. The same thing applies if your device uses all passive interfaces. If this is the case, *select your customized Network Analysis policy here and move on*. You don't need any network analysis rules! The vast majority of deployments fall into this category.

So the real purpose of network analysis rules is to allow us to specify a different Network Analysis policy for traffic in different zones, networks, or VLANs. This means that if you have both inline and passive interfaces on the same device, you probably want to specify two different Network Analysis policies. In this case, using zones is probably the best idea. Zones translate to physical interfaces, so these are natural traffic delineators between passive

and inline interfaces. To specify these different policies, you need custom rules.

To create custom rules, click the No Custom Rules link shown in the preceding figure. This expands the dialog as shown here:

Network Analysis and Intrusion Policies

? X

Intrusion Policy used before Access Control rule is determined

No Rules Active

Intrusion Policy Variable Set

Default-Set

Network Analysis Rules

[No Custom Rules](#)

[Network Analysis Policy List](#)

 Add Rule

#	Source Zo...	Dest Zones	Source Networ...	Dest Networks	VLAN T...	Network Analysis ...
---	--------------	------------	------------------	---------------	-----------	----------------------

No Rules

Default Network Analysis Policy

Balanced Security and Connectivity

Revert to Defaults

OK

Cancel

This allows you to create rules based on zone, source/destination network, and/or VLAN tag. Each of these rules has an associated Network Analysis policy, so you could specify your inline Network Analysis policy as the default and then add a rule that includes all of your passive zones. In that rule you'd specify your passive Network Analysis policy—the one without the inline preprocessor enabled.

Even so, if you had both inline and passive interfaces, it probably wouldn't make much difference if you just used a Network Analysis policy with inline mode enabled for all the traffic. In security, we really have to pick our battles, and you need to decide how much time you have to tweak your policies and settings and whether that time could be used more effectively somewhere else. Also, consider that adding these rules increases the complexity of your policy and that's to be avoided when you can. But, if you really want to dial in your detection and take advantage of all the flexibility that Firepower has to offer, custom analysis policy rules are one way to get that done.

Files and Malware Settings

The Files and Malware Settings section allows you to tweak the behavior of your Advanced Malware Protection (AMP) and file detection. Most of the settings here are designed to balance performance and detection, but be very careful, especially when it comes to increasing these values. They were set by folks who know an awful lot about how malware behaves and how to squeeze the most performance out of the system while still providing the maximum detection possible!

The Files and Malware Settings dialog is pictured here:

Files and Malware Settings



Limit the number of bytes inspected when doing file type detection

Do not calculate SHA256 hash values for files larger than (in bytes)

Allow file if cloud lookup for Block Malware takes longer than (seconds)

Minimum file size to store (bytes)

Maximum file size to store (bytes)

Minimum file size for dynamic analysis testing (bytes)

Maximum file size for dynamic analysis testing (bytes)

Revert to Defaults

OK

Cancel

Limit the number of bytes inspected when doing file type detection

Default: 1460 bytes

Range: 0–4294969295 (4 GB)

When performing file type detection, Firepower looks for a signature near the beginning of the file. For example, a GIF file will always begin with the string “GIF8.” Other file types, such as EXE or PDF, also have similar telltale strings or byte patterns. This setting limits how many bytes will be inspected to make this determination. The default size of 1460 is equivalent to the data segment of a typical Ethernet TCP packet and a value of zero removes the restriction altogether.

Do not calculate SHA256 hash values for files larger than (in bytes)

Default: 10485760 (10 MB)

Range: 0–4294969295 (4 GB)

This setting will prevent the system from storing, performing malware cloud lookups for, or blocking files larger than this size. This value has to be greater than or equal to the settings for “Maximum file size to store (bytes)” and “Maximum file size for dynamic analysis testing (bytes).” Again, setting this value to zero removes the size restriction.

At first this can really look like a good number to increase. After all, it seems like we’re allowing malware larger than 10 MB to traverse our network. Isn’t this opening up a hole in our detection? Well, the answer is yes and no. Yes, this does mean we do not do malware lookup on files larger than 10 MB, but it turns out that 99.9% of all malware is under 10 MB. This helps preserve the detection resources of the devices by not requiring them to collect and calculate hashes for large files. Statistically, these large files are extremely unlikely to contain malware, so including them would probably just get you some ugly performance with virtually no benefit!

Allow file if cloud lookup for Block Malware takes longer than (seconds)

Default: 2 seconds

Range: 0–30 seconds

This setting determines how long a device will hold on to the last piece of a file while waiting for the FMC to perform a malware cloud lookup on the SHA-256. This setting prevents a large delay in a file transfer in case there's latency in performing the cloud lookup. If this time elapses with no response from the FMC, the device will allow the file transfer to complete and register a disposition of Unavailable for the file's status. As with most settings, Cisco recommends leaving this at the default.

Minimum file size to store (bytes)

Default: 6144 (6 KB)

Range: 0–10485960 (10 MB)

The minimum file size that can be stored using a file rule, and again a setting of zero disables file storage. Storing gobs of really tiny files is a potential drain on system resources with dubious benefits. This field must be less than or equal to “Maximum file size to store (bytes)” and “Do not calculate SHA256 hash values for files larger than (bytes).”

Maximum file size to store (bytes)

Default: 1048596 (1 MB)

Range: 0–10485766 (10 MB)

The maximum file size that can be stored using a file rule. A setting of zero disables file storage. Must be greater than or equal to “Minimum file size to store (bytes)” and less than or equal to “Do not calculate SHA256 hash values for files larger than (bytes).”

Minimum file size for dynamic analysis testing (bytes) Default: 6144 (6 KB)

Range: 6144 (6 KB) – 2099152 (2 MB)

The minimum file size the system will submit to the cloud for

dynamic analysis. This field must be less than or equal to “Maximum file size for dynamic analysis testing (bytes)” and “Do not calculate SHA256 hash values for files larger than (in bytes).”

Maximum file size for dynamic analysis testing (bytes) Default: 1048576 (1 MB)

Range: 6144 (6 KB) – 2099152 (2 MB)

The maximum file size the system will submit to the cloud for

dynamic analysis. This field must be greater than or equal to “Minimum file size for dynamic analysis testing (bytes)” and less than or equal to “Do not calculate SHA256 hash values for files larger than (in bytes).”

Intelligent Application Bypass Settings

Intelligent Application Bypass (IAB) is a new-ish and cool feature in Firepower. It was first introduced in version 6.0 and it allows the system to automatically bypass or trust traffic flows based on criteria you set. The idea is that large flows like nightly backup scans cause excessive utilization and increase packet latency through the device, which not only affects the backup, but all other traffic too.

Intelligent Application Bypass Settings

? X

State	<input type="text" value="Off"/>
Performance Sample Interval (seconds)	<input type="text" value="5"/>
Bypassable Applications and Filters	<input type="radio"/> <u>0 Applications/Filters</u> <input checked="" type="radio"/> All applications including unidentified applications
Inspection Performance Thresholds	Hide
Drop Percentage	<input type="text" value="5"/>
Processor Utilization Percentage	<input type="text" value="95"/>
Packet Latency (microseconds)	<input type="text" value="1000"/>
Flow Rate (flows/second)	<input type="text" value="0"/>
Flow Bypass Thresholds	Hide
Bytes per Flow (kbytes)	<input type="text" value="500000"/>
Packets per Flow	<input type="text" value="0"/>
Flow Duration (seconds)	<input type="text" value="0"/>
Flow Velocity (kbytes/second)	<input type="text" value="250000"/>

Revert to Defaults

OK

Cancel

It would definitely be better to *not* inspect these backups since there is no security value in doing so.

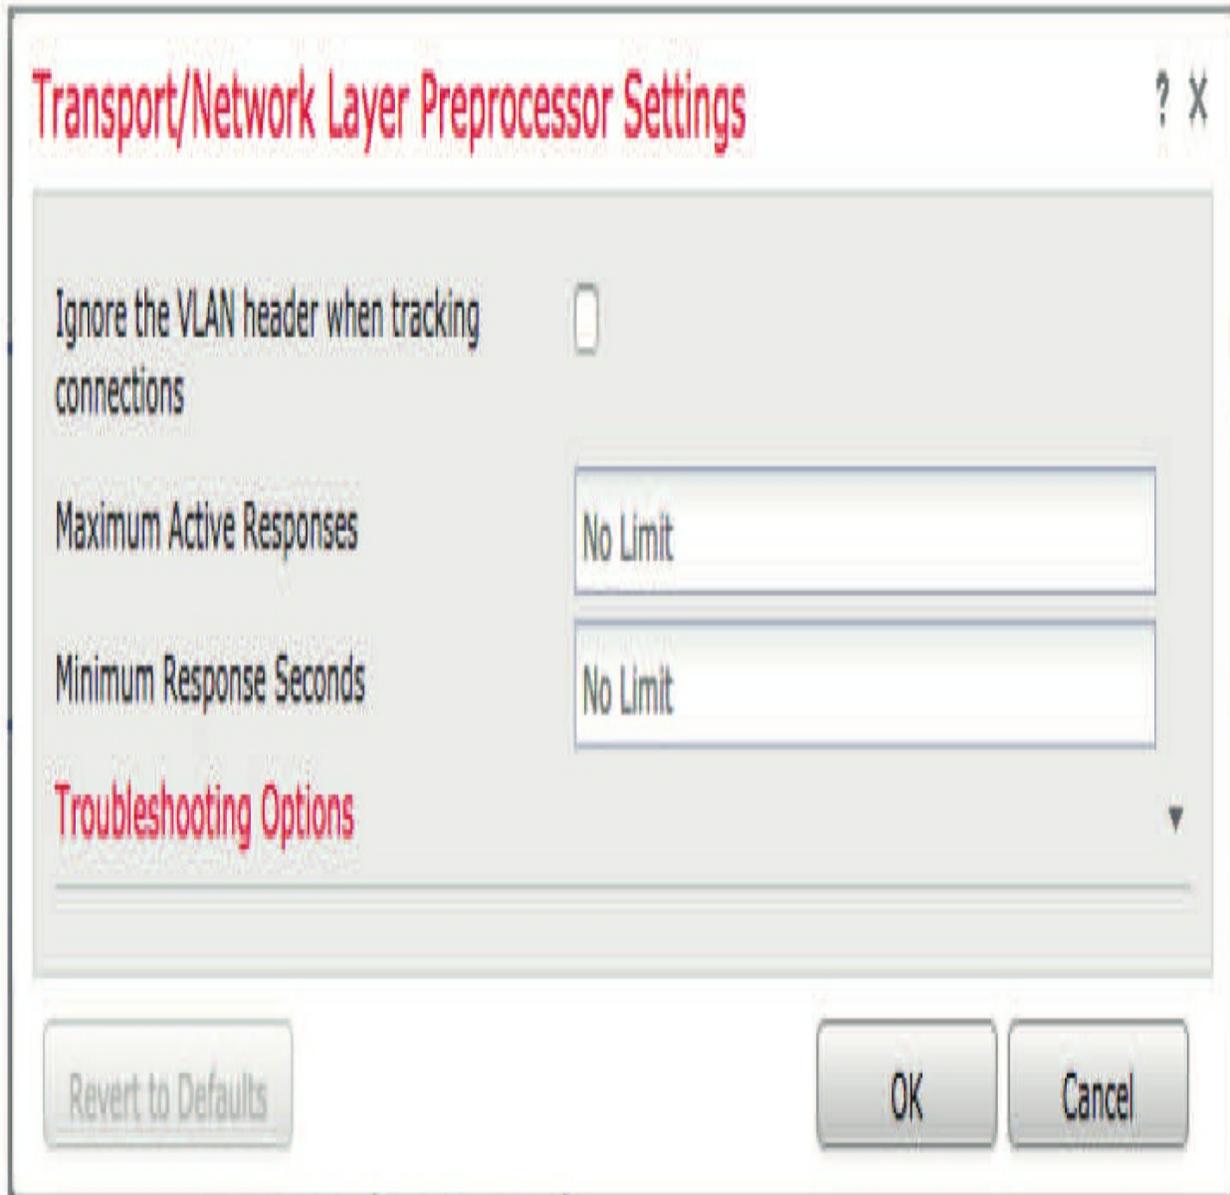
You can trust these flows with explicit trust rules or you can do it automatically using IAB. Even if you do add manual trust rules, IAB can help by automatically trusting other flows you may not have anticipated. The next figure shows the available IAB settings:

To use IAB, first set the *State* option to Test or On, then select a non-zero *Performance Sample Interval (seconds)* value. After this, click the link next to *Bypassable Applications and Filters* and add the applications you want to be eligible for bypass. Once that's done, adjust the *Inspection Performance Thresholds* and *Flow Bypass Thresholds* settings. You've got to set at least one condition in each of these two categories, and both conditions must be exceeded for bypass to kick in on a given flow.

It's a very good idea to start in Test mode because in this mode, no traffic will be bypassed but connection events will still indicate any flows that would've been bypassed using the current settings. You should confirm that these flows *should* have been bypassed prior to enabling this feature!

Transport/Network Layer Preprocessor Settings

These settings control some seldom-used behaviors of Snort's transport/network layer preprocessor (stream5). The figure below shows the Transport/Network Layer Preprocessor settings. The first check box is *Ignore the VLAN header when tracking connections*. You want to check this if your device might process packets that are part of the same TCP connection but have different VLAN tags.



True, this is unusual, but if you don't check this, these connections won't be correctly reassembled.

The next two items have to do with Snort's active response capability. The *Maximum Active Responses* setting limits the number of responses Snort will generate for a specific connection. *Minimum Response Seconds* determines the minimum number of seconds to wait before initiating subsequent active responses for a connection.

There are two situations that are being addressed here.

The first one is drop rules. When deployed inline, (including switched or routed), Snort will actively terminate a connection in the case where a drop rule matches a packet. It does this by silently dropping the offending packet and then continuing to drop subsequent packets in that same connection by sending a RST.

The second one has to do with Snort rules written with the *resp* or *react* keywords. These keywords can be used in rules to tell Snort to take an active role in terminating the connection. It does this by sending a TCP RST or ICMP port unreachable (type 3) packet to the source host. Rules with either of these keywords will reset the connection regardless of the settings selected here. Still, these settings determine whether Snort continues to initiate additional active responses as it does for drop rules.

If you want to disable Snort's active responses altogether, just enter a zero in the *Maximum Active Responses* field. This will cause drop rules to silently drop only the offending packets that they match. It will also disable active responses for the *resp* and *react* keywords.

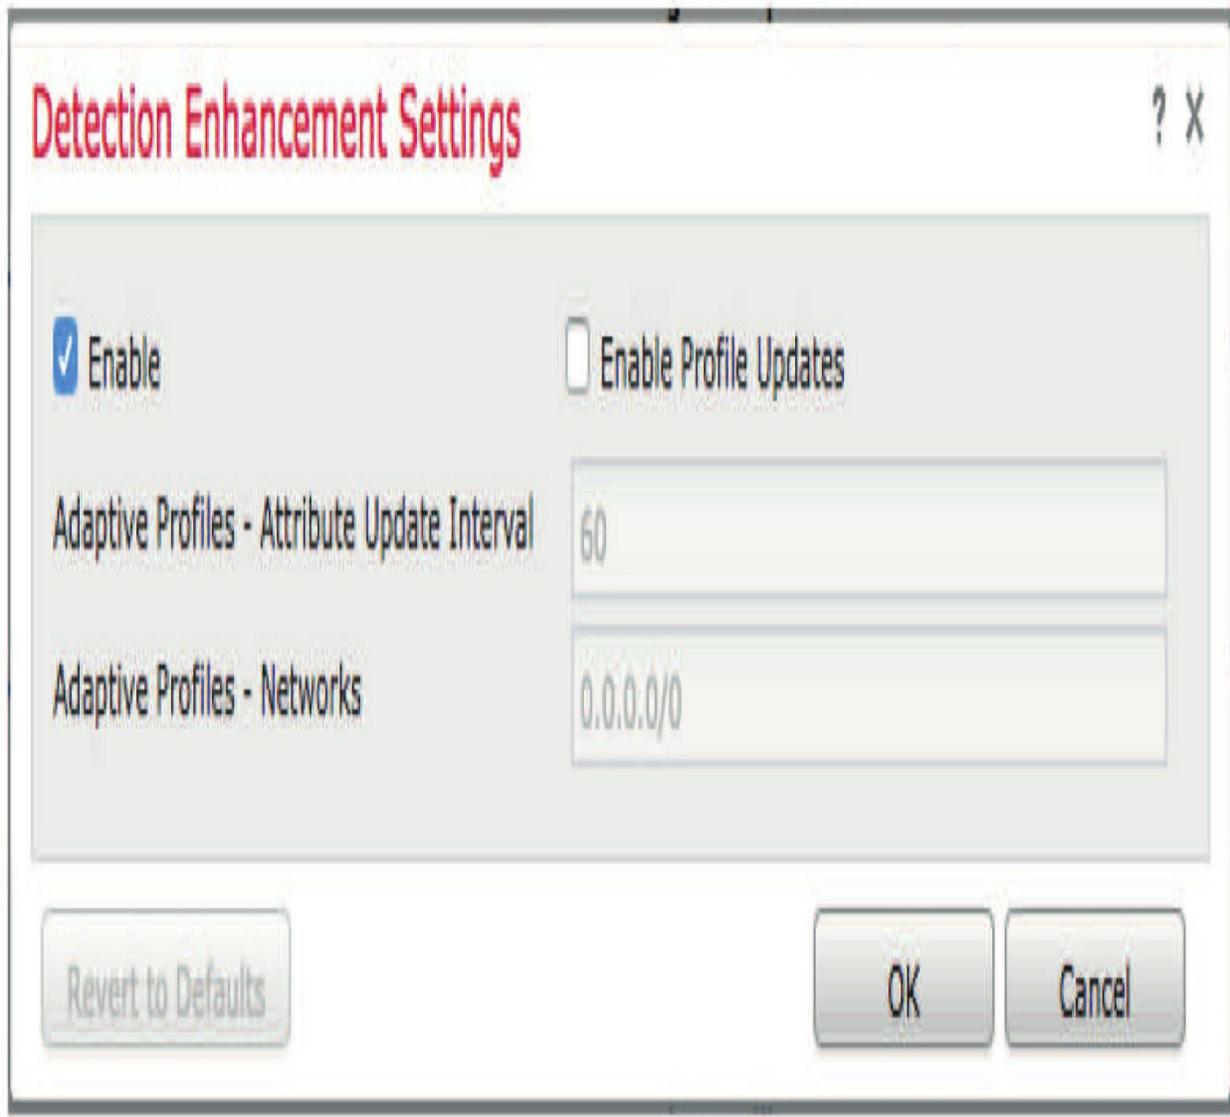
Keep in mind that the active response settings don't apply to ASA with FirePOWER Services modules.

Depending on your FMC model, you may have more than one interface available. All FMCs will default to a single 1 GB copper management port. The FMCs 4500/4600 and 2500/2600 also have optional 10 GB fiber ports

Detection Enhancement Settings

Detection Enhancement Settings has been updated a bit for version 6.x. Previously, this setting was only used to enable adaptive profiles and set the update interval. The default was to disable adaptive profiles. Starting with 6.1, we now have an Enable and an Enable Profile Updates check box. The former is enabled by default and the latter is disabled. But what does it all mean?

Check out this figure for an example of the Detection Enhancements Settings dialog:



First, let me talk about the classic definition of adaptive profiles in Firepower. The purpose is to allow the device to “adapt” its traffic processing based upon the various packet reassembly behaviors of different operating systems. Different operating systems perform both TCP and IP fragment reassembly differently, and this is especially true when there are overlapping data offsets in the packets. Never mind that this should never happen—the fact that it can happen opens up the possibility that an attacker could evade detection by fooling the IPS into reassembling packets incorrectly!

NOTE: If you’re curious and want to know more about this, check out the 1998 paper by Thomas Ptacek and Timothy Newsham called “Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection.” To

find out how this is implemented in Snort, read “Target-Based Fragmentation Reassembly,” written in 2005 by Judy Novak. Both of these papers can be found online by searching for their titles.

The key factor with this is that if we know the operating system performing the reassembly, we can predict what the resulting data will be. The Adaptive Profiles feature allows Firepower to do just that. Because Firepower passively builds a host database, it presumably knows the operating systems of all the hosts in the protected network. This information is saved on the FMC. So when you enable Adaptive Profiles and periodic updates are sent to the managed devices, Snort can leverage this intelligence when performing packet and stream reassembly.

Now, in previous versions of Firepower, we only enabled Adaptive Profiles if we had passive devices. When the devices are inline, the inline normalization preprocessor handles this without being updated with information from the FMC’s host database.

But starting in version 6.1, we have a new *Enable* check box, which is checked by default. This box must remain checked for access control rules to perform application and file control, including AMP, and for Snort rules to use service metadata. So yep, this box is a huge deal because unchecking it turns off most of the valuable features you bought Firepower for in the first place! Would you ever want to disable these beauties and uncheck that box? Well sure—if your system is held captive in a painfully cramped environment with no application control, no AMP, and you don’t use any of the most popular protocols like HTTP, FTP, and SMTP, then yes—knock yourself out and shut it all down!

But you still have the *Enable Profile Updates* option if your policy is deployed to passive devices, and you should also set *Adaptive Profiles – Attribute Update Interval*. This is the interval in minutes when the FMC is triggered to update the managed device with new profile data from the host database. If you want to be even more efficient, use the *Adaptive Profiles – Networks* setting to define a subnet range or variable where the protected hosts are for this device. Otherwise, information on all the hosts in the database will be pushed to the target device or devices.

Performance Settings

I'm going to cover the settings shown below on an "FYI only" basis because performance settings are something that you just don't touch. If you can't stand it and are now dying to know what these things do, PLEASE consult the online help and proceed with caution! Be well warned that changing parameters in areas like regular expression or pattern matching limits can create detection nightmare scenarios that are so extremely hard to troubleshoot it will make you cry.

Here's a snapshot of these little piranhas—hands off!

Performance Settings



Pattern Matching Limits - Max Pattern Match States to Analyze Per Packet	5
Performance Statistics - Sample Time (seconds)	300
Regular Expression - Limit	Default
Regular Expression - Recursion Limit	Default
Intrusion Event Logging Limits - Max Events Stored Per Packet	8

Latency-Based Performance Settings

Latency-Based Performance Settings control Snort's behavior in cases where packet processing is slowing down. These settings determine how Snort behaves when there's something causing packets to move through the

inspection process too slowly. There are two tabs available: Packet Handling and Rule Handling, both shown below.

Latency-Based Packet Handling

The Packet Handling tab controls the per-packet latency threshold. Each packet is timed because it's processed through Snort's preprocessors and intrusion rules. The timer starts at zero and is checked at various points throughout the inspection process. If at any time the latency number exceeds the microsecond setting here, the packet is released to exit the device.

The purpose of this setting is to reduce or prevent negative impact on applications. The default setting is 256 microseconds and, according to the Firepower help, equates to a traffic flow of about 100 Mbps.

The following table shows the recommended minimum packet latency settings:

Data Rate Microsecond Threshold

1 Gbps	100
100 Mbps	250
5 Mbps	1000

By default, packets that exceed this threshold pass through and out of the device without generating any alerts. Yes—hearing for the first time that packets might be slipping through your devices can be a bit unnerving, but in this case, there's actually something you can do to know for sure!

To view alerts for packets triggering the packet latency threshold, *enable Snort rule 134:3*. The easiest way to find this beauty is to go to your Intrusion policy and search for “GID:134” in your rule filter bar.

This is shown in the following figure. Set the rule state for 134:3 to Generate Events and stand back!



I say that this because if you've got a lot of traffic and you're hitting this threshold repeatedly, you're likely to see a legion of events generated from this rule. So just be ready to disable it if it sends your event analysis group into a frenzy!

In my experience, the default tends to favor application performance, meaning when utilization spikes and latency numbers climb, there'll likely be some packets that aren't fully inspected as they pass through the device. Because of this, you may find yourself scrambling to raise the default number, which will potentially increase packet latency but also reduce the number of packets that exit the device prematurely.

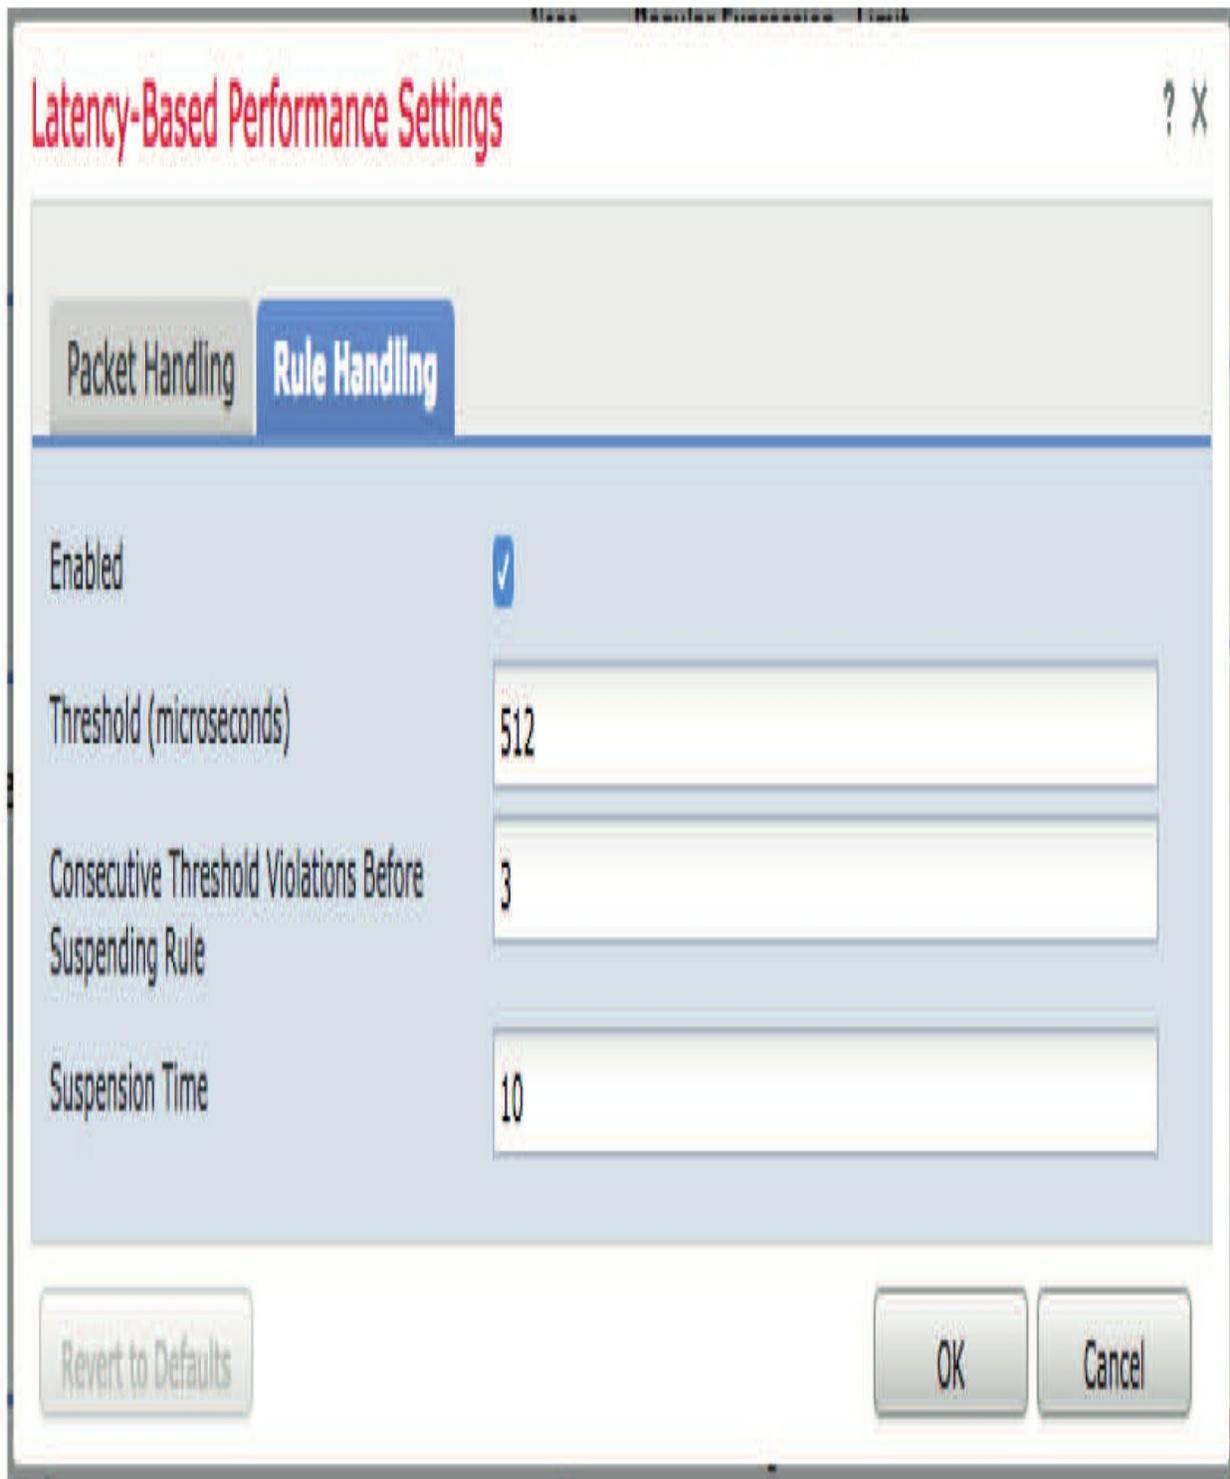
The idea here is to adjust this number through trial and error—yep, trial and error—until you arrive at a happy place where you obtain maximum traffic visibility with little or no impact on applications.



Unless your devices are woefully underpowered, you should be able to perform 100% inspection with no application impact—with practice.

Latency-Based Rule Handling

The Rule Handling tab controls Snort's behavior when there are specific rule groups that are introducing excessive latency.



Like Packet Handling, these settings are designed to prevent Snort from causing application problems due to poor performance. In this case, we're zeroing in on poorly performing rules. There are four parameters that can be

configured here:

Enabled

Default: On

Enables or disables the Rule Latency setting. If this is disabled, Snort will never disable poorly performing rules due to excessive latency.

Threshold (microseconds)

Default: 59

The default microsecond threshold allowed for a rule group to process a single packet. Exceeding this threshold will increment the violation count for the group.

Consecutive Threshold Violations Before Suspending Rule Default: 3

The number of consecutive threshold violations before the rule group is automatically disabled.

Suspension Time

Default: 10 seconds

The number of seconds a rule group will be suspended. The documentation for this setting speaks of “rule groups.” So,

what are these groups? When Snort first starts up, it builds decision trees from the rule set. The purpose of these trees is to leverage the fact that many rules look for similar packet conditions. Think about it—if a bunch of rules are all looking for the same content, it would be really inefficient to repeatedly check for that as Snort evaluates a horde of rules individually against a given packet, right?

So by grouping similar rules together, Snort can inspect a lot more efficiently by qualifying or eliminating a number of rules with a single content check. To this done, Snort builds rule trees or groups with similar rules grouped together for inspection; the process is completely transparent and occurs in the Snort start-up process. This typically just lasts a few seconds.

Genius-level knowledge of this rule grouping behavior is beyond the scope of our current talk, but understanding it comes in very handy for rule writers setting out to create more efficient Snort rules.

During packet inspection, if a rule group exceeds the rule handling latency threshold, the violation counter for that group is incremented. If the violation count meets the consecutive violations number, the rule group is temporarily disabled and packets won't be inspected by it. At the end of the time specified in the Suspension Time setting in seconds, the rule group will be re-enabled.

Know that for a rule group to be disabled, the violations must be consecutive. Any packets inspected by the rule group that don't exceed the threshold will reset the violation counter to zero. This means a stray packet or two that exceed the threshold won't trigger any change in Snort's inspection.

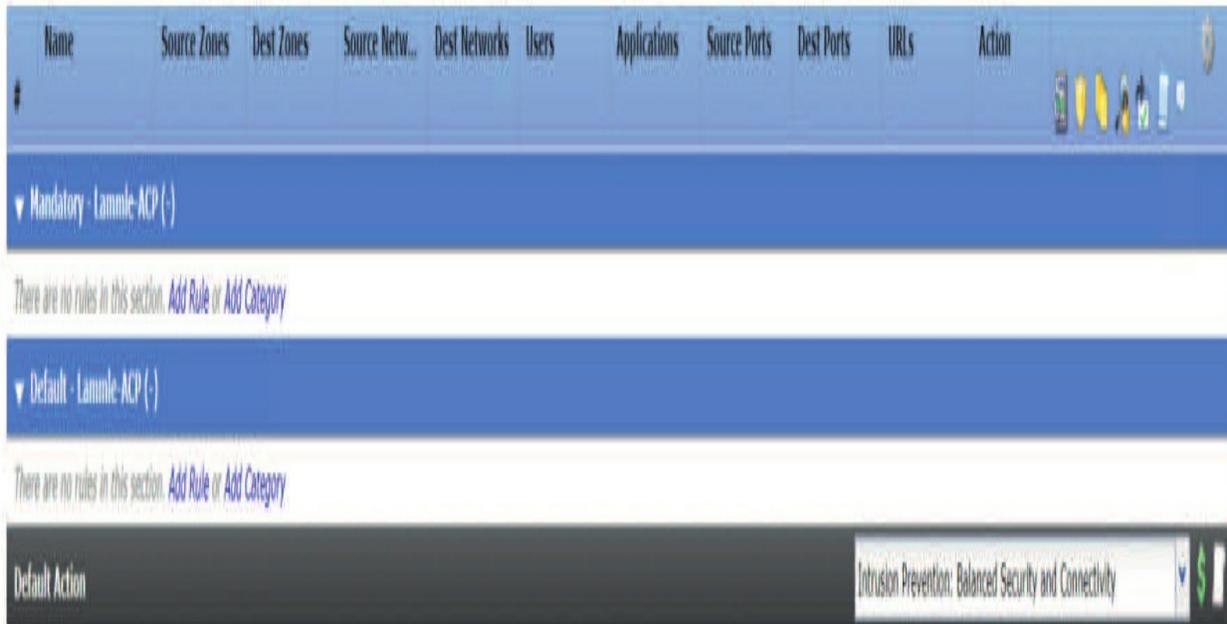
As with packet latency, the default setting is silence if a rule group is disabled. To enable event logging for this setting, enable events 134:1 and 134:2 using the technique mentioned previously. The 134:1 intrusion event will trigger when a rule group is disabled, and 134:2 will trigger when the rule group is re-enabled—10 seconds later by default.

Getting these alerts indicates that your device may be oversubscribed, or you have poorly written rules. I've seen plenty of cases where sloppy rules are repeatedly disabled even though the device is humming along fine otherwise. Because this setting is enabled by default with no alerting, it can dupe you into thinking your fancy custom rule is working great when it's actually being disabled nearly 100% of the time due to high latency!

The Rules Tab

If Firepower has one location that could be considered the heart of the system, that location is the Access Control Rules tab pictured below. This is the central rule set determining what happens to traffic because it's processed through the devices.

Check it out:



Rule Categories

Access control rules are evaluated from top to bottom. By default, there are two rule categories, Mandatory and Default, we use to group similar types of rules together, and additional categories are added by clicking the Add Category button. A good example is when you have a Trust category that you've placed near the beginning of your rule set. The Add Category dialog is shown below.

Adding additional categories can really help organize rules. Important to keep in mind is that the category has nothing to do with the order rules are processed. Traffic is always processed from top to bottom regardless of whether you use multiple categories or place all your rules into a single category. When adding a custom category, you have the option to select the existing category where it will be inserted. Choose carefully here—while you can reorder individual rules, you can't reorder categories!



The image shows a dialog box titled "Add Category" with a close button (X) in the top right corner. The dialog contains two main input fields: "Name:" followed by a text input box, and "Insert:" followed by a dropdown menu currently showing "into Mandatory". At the bottom right of the dialog are two buttons: "OK" and "Cancel".

Shown in the figure of the Rules tab above is a child policy that's inheriting settings from a base policy aptly named Base Policy. Any rules in the base policy are added to Mandatory and Base categories at the beginning and end of the child policy. The effect is to give the base policy the first and last word when it comes to rules. Clicking the white triangle icon next to the category name will expand or collapse the rule listing in that category.

Default Action

The Default Action is the action that the device takes if none of the access control rules match a given traffic flow. Clicking the dropdown field displays the available actions:

- **Inherit from base policy:** This option only appears for a child policy so the default action becomes whatever's configured in the base or master policy.
- **Access Control: Block All Traffic:** Simply stated, if traffic makes it this far, block it. It's a typical choice for a device deployed as a firewall.
- **Access Control: Trust All Traffic:** If traffic makes it this far, let it pass through the device with no further inspection performed.
- **Network Discovery Only:** Traffic is inspected via the Network Discovery policy only. This is not recommended.
- **Intrusion Prevention:** Allow the traffic, but first inspect it with Intrusion and Network Discovery policies. All of the Cisco-provided and user-created policies are listed here. Depending on the Intrusion policy, some traffic could still be dropped if it matches a Snort rule with the "Drop and Generate" action.

It's worth noting that you can't perform file and malware inspection with the default action. Those things must be performed in an access control Allow rule.

Rules!

To add a new rule, click the Add Rule button.

Add Rule

? X

Name

Enabled

Insert into Mandatory

Action



- Zones**
- Networks
- VLAN Tags
- Users
- Applications
- Ports
- URLs
- SGT/ISE Attributes
- Inspection
- Logging
- Comments

Available Zones

- Arris-Router
- Garage-Switch
- MSP-Firewall
- Routed-ESX
- Routed-Switch
- GRE-tunnel

Add to Source

Add to Destination

Source Zones (0)

any

Destination Zones (0)

any

Add Cancel

Name

Your rule deserves a name but keep it short because you only get 30 characters here! The name field turns an angry red if you enter too many characters.

Enabled

New rules are enabled by default. Unchecking this box will leave the rule in its current location but it won't process traffic. **Insert**

Choose the category or position of the rule here. If this is the first rule, you'll only have category choices, but if there are other rules in your policy already, you can insert your rule above or below an existing rule number.

Rule Actions

Allow

Use this action for AMP or Snort inspection. Traffic matching Allow rules is allowed to pass through the device after being inspected by the appropriate policies. You can inspect this traffic using an Intrusion policy, a File policy or both. Depending on the rules in these policies, traffic could be blocked. Technically, you can use an Allow rule with no inspection, but doing that doesn't really make sense because that's what a Trust rule is for.

Trust

The Trust action allows traffic to pass through the rule set unmolested. It's typically used for network communications you don't want to inspect because of an impact to either the device or the application. Trust rules are usually inserted early in the rule set.

Monitor

This is the only action that doesn't affect traffic flow because the only purpose of a Monitor rule is to log a connection event. Traffic matching a Monitor rule continues to be processed by subsequent rules. It's a good idea to insert a Monitor rule early in the rule set to ensure connection events are generated. This also prevents the need to remember to check the logging

option for some of your other rules.

Block and Block with Reset

These two actions deny traffic without further inspection. The Block action simply stops the traffic from passing, while Block with Reset stops it and also resets the connection. Blocked traffic is not inspected by Intrusion, File, or Discovery policies, and even if you have a Monitor rule upstream, you'll probably still want to enable logging for this rule type.

Interactive Block and Interactive Block with Reset

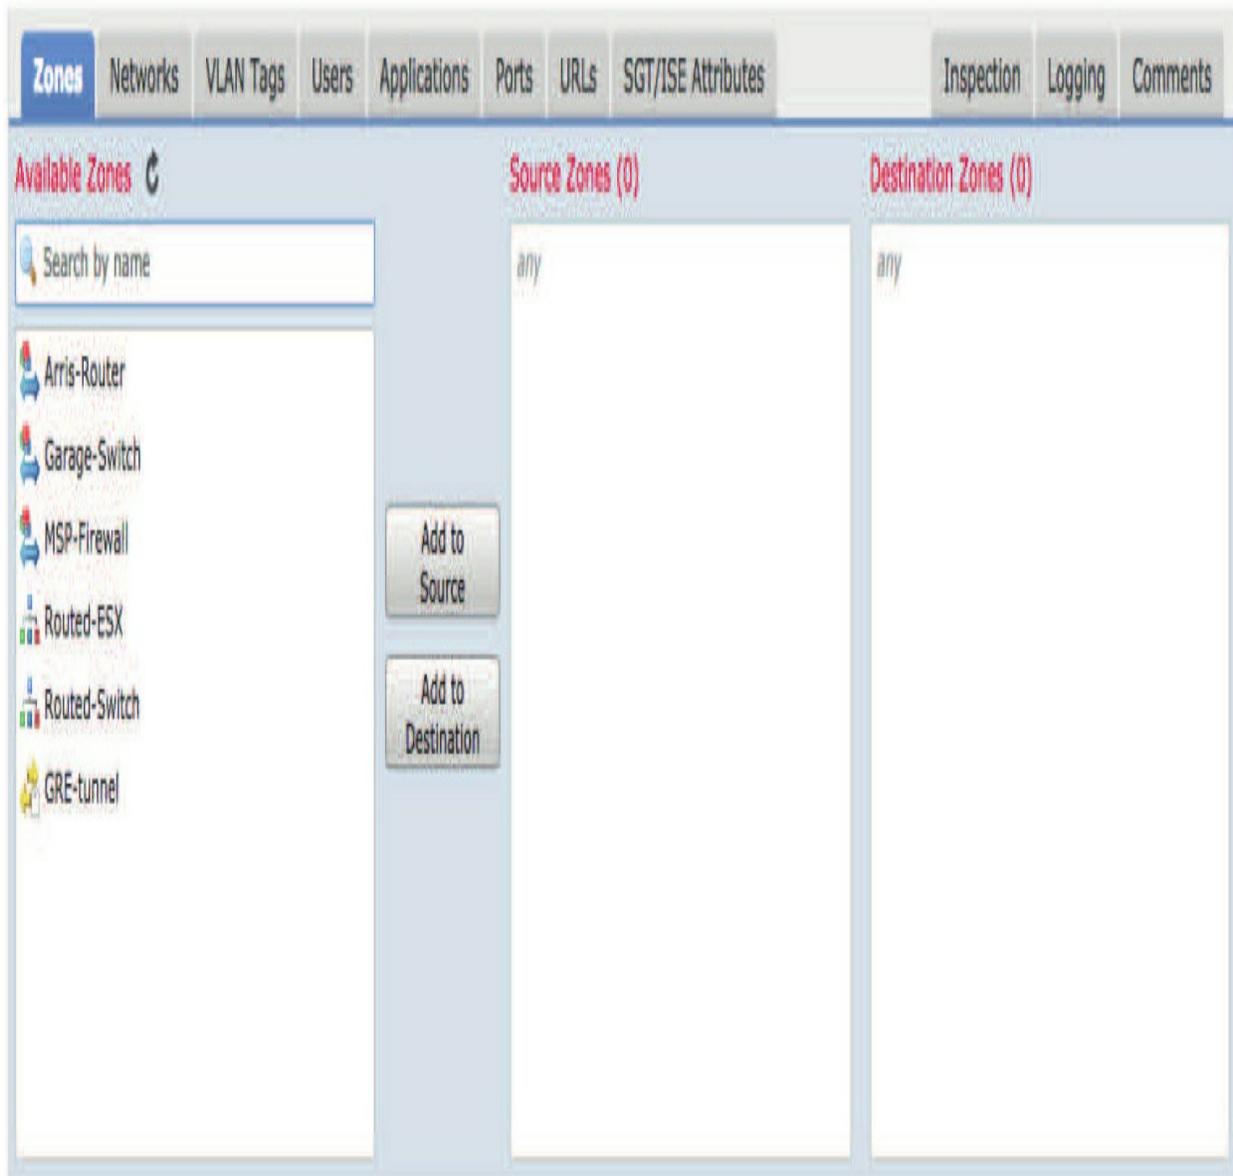
In the case of HTTP traffic, these actions allow users to bypass a website block by clicking through a warning page. If the user chooses not to bypass the block, the rule functions like a Block rule in that traffic is denied without further inspection. But if somebody goes ahead and bypasses the block, the rule functions like an Allow rule and can perform intrusion, file, and discovery inspection. These rules can be associated with Intrusion and File policies in the same way that Allow rules can be. For non-HTTP traffic, these rules work just like a Block or Block with Reset rule.

Rule Criteria

The group of eight tabs to the left of the dialog is used to determine the criteria for matching traffic. By default, all of the tabs are empty, meaning your rule will match all traffic.

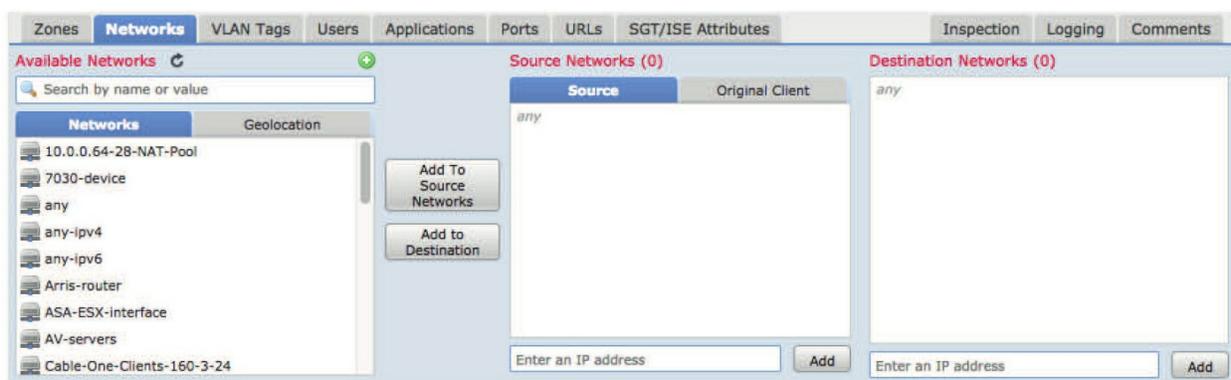
Zones

The Zones tab, shown below, allows you to select the source or destination zone or tunnel tag. This is helpful when your rule is designed for traffic to or from a certain interface or tunneled traffic type. Select the applicable zone or zones and use the Add to Source or Add to Destination button as appropriate:



Networks

The Networks tab, pictured below, give us several options. First, you can select an existing network object from the left and add it to the source or destination networks column.



You can also enter an IP address or netmask manually below the source or destination column. If you want to add a new network object from here, use the green plus icon () to add an object on the fly. After you do so, the object will appear in the Networks column.

The Geolocation sub-tab provides access to all available locations plus any Geolocation objects you've created. This can be used instead of or in addition to the other network criteria.

There's also a very cool sub-tab in the Source Networks column. The default setting of Source simply matches the IP address of the packet's source IP field in the header, but the Original Client tab allows matching on the original client IP address. This applies to HTTP traffic only and equates to the X-Forwarded-For, True-Client-IP, and other HTTP header fields as configured in the Network Analysis policy HTTP Configuration settings. Why this is cool is because if your device is positioned outside your outbound web proxy, you can still take action on traffic based on the IP address of your internal hosts.

VLAN Tags

VLAN Tags allows you to specify a number from 0 to 4094 to identify a network by VLAN. You can use objects you have created, add objects on the fly using the green plus icon, or enter VLAN numbers directly below the Selected VLAN Tags column as shown here:

Zones Networks **VLAN Tags** Users Applications Ports URLs SGT/ISE Attributes Inspection Logging Comments

Available VLAN Tags   Selected VLAN Tags (0)

Search by name or value

VLAN-100

any

Zones Networks VLAN Tags **Users** Applications Ports URLs SGT/ISE Attributes Inspection Logging Comments

Available Realms  Available Users  Selected Users (0)

Search by name or value

Search by name or value

any

Special Identities

Server 2012

Server 2012/*

- Account Operators
- Administrators
- Backup Operators
- Cert Publishers
- Distributed COM Users
- DnsAdmins
- DnsUpdateProxy
- Domain Admins
- Domain Computers

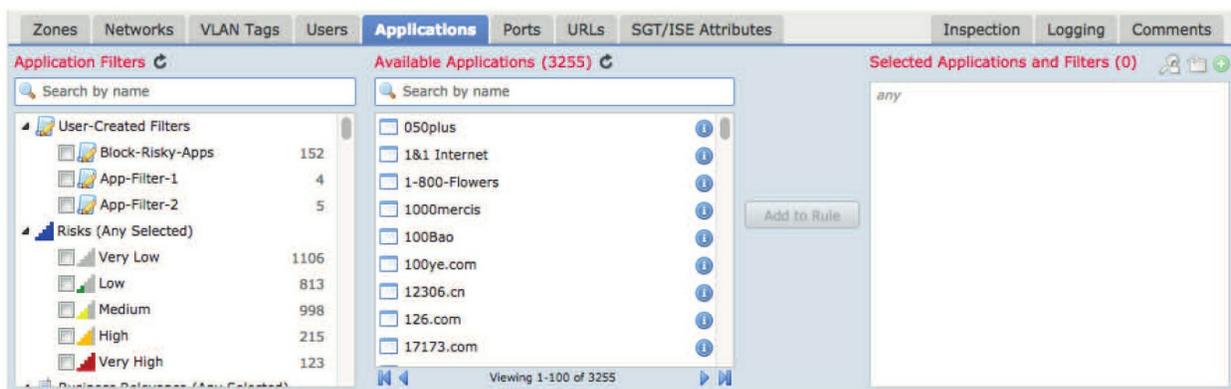
Add to Rule

Users

This tab, shown below, allows you to specify which users your rule should apply to. These users and groups are retrieved from a Microsoft Active Directory server. Before you can use this condition, you must configure an Identity policy and it must be selected using the Identity Policy link near the upper right of the main Access Control policy screen or in the Advanced tab:

We'll cover Identity policy thoroughly in Chapter 16. **Applications**

The Applications tab seen in the following figure allows for filtering your rule based upon built-in applications, user-defined applications, and the application filter objects you've created. You'll find the same application criteria (Risks, Business Relevance, Types, etc.) available here as you saw under Application Filters on the Object Management page back in Chapter 8. You'll also see any user-created filters previously added here:



In the upper right corner, there are two icons that control the Safe Search and YouTube EDU settings. These icons are only clickable for rules with the Allow action and they're visible but disabled for other rule actions.

Safe Search

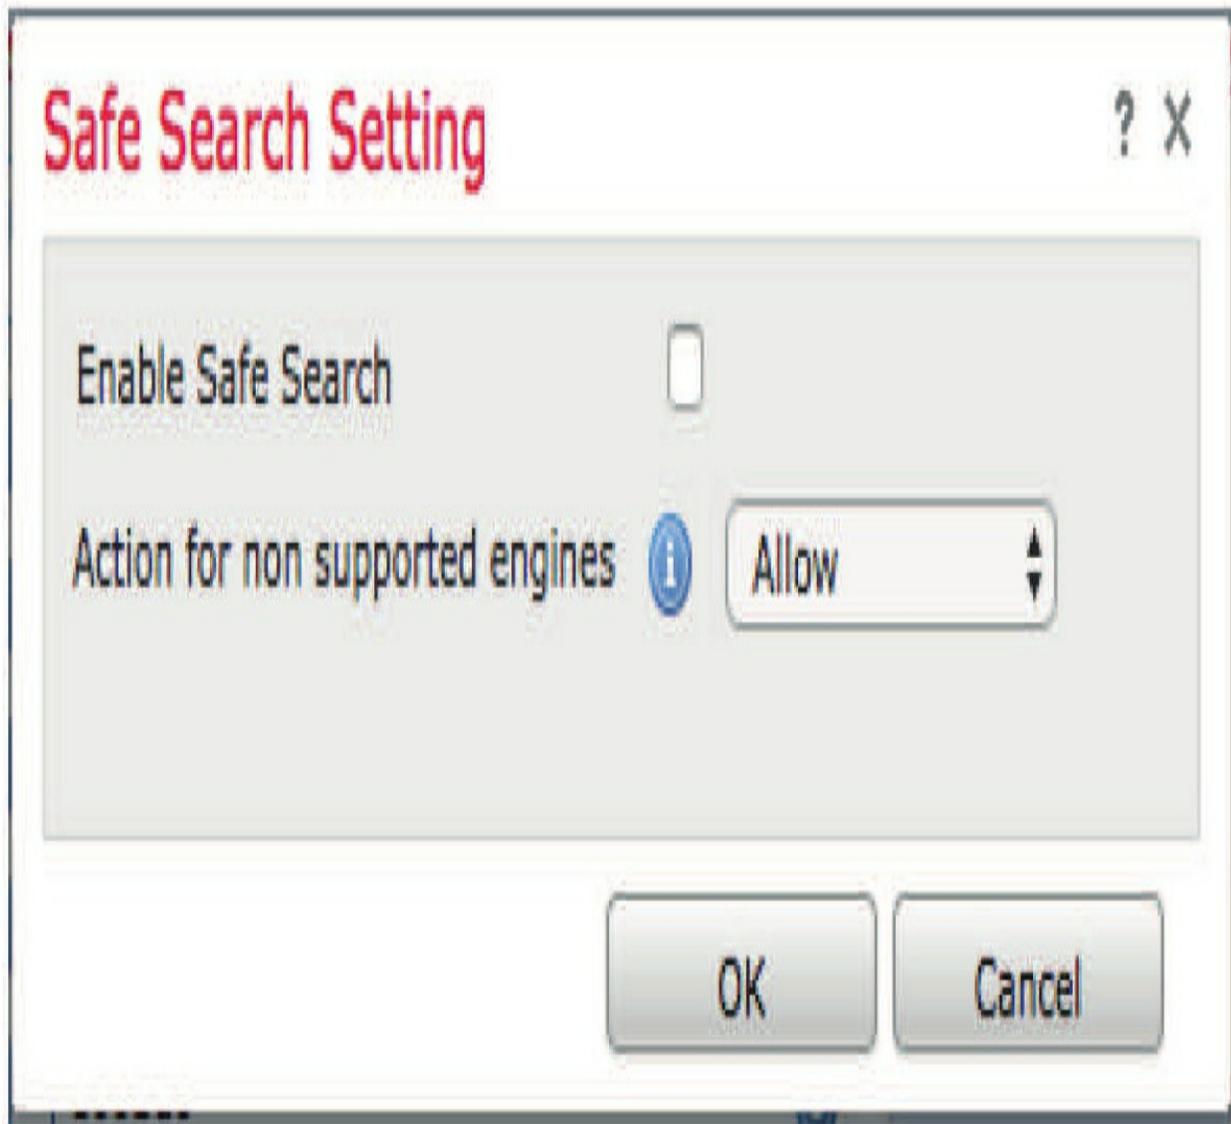
Safe Search is supported by major search engines and involves filtering out explicit and adult-oriented content that businesses and education environments, for example, would find objectionable. The way this works is by communicating the restricted status to the search engine via the request URI, cookie, or a custom HTTP header element. When you configure this in

an AC Allow rule, Firepower makes the appropriate modifications to traffic matching the rule.

Clicking Safe Search icon (



) displays the dialog shown here:



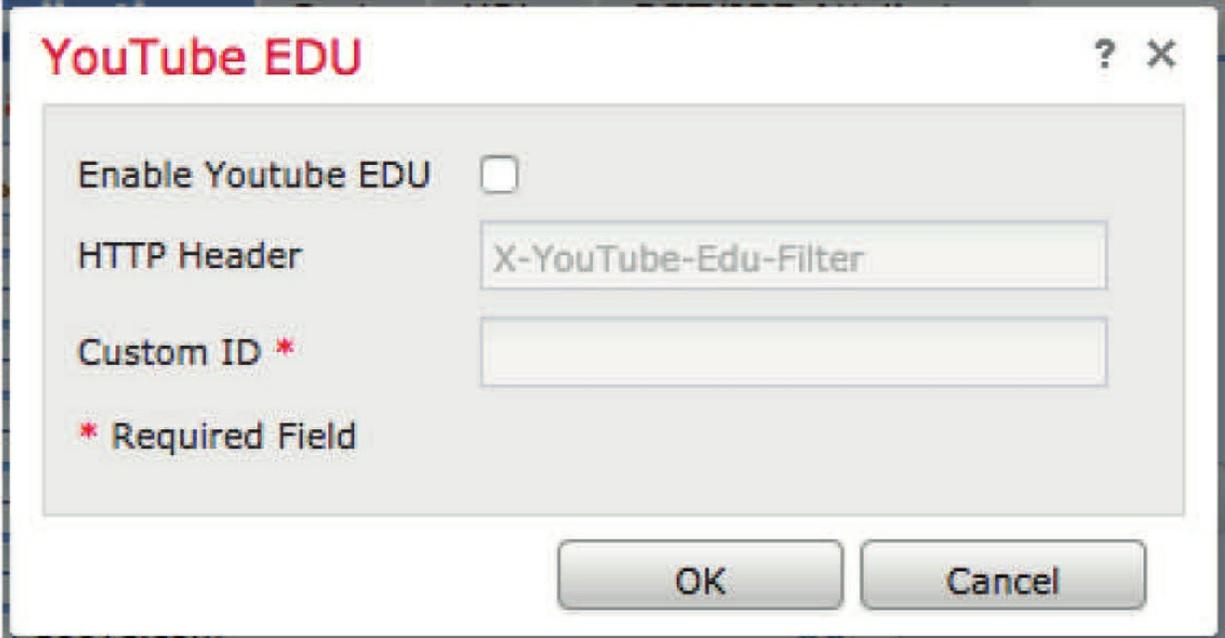
Once you enable Safe Search, you must choose an action for traffic to non-supported search engines. If you choose Block or Block with Reset, you've

also got to configure the HTTP response page that the system displays when blocking restricted content.

There are specific application tags for search engines that do or don't support Safe Search. To find these, look under the Application Filters column, expand the Tags category and look for "safesearch supported" or "safesearch unsupported." These settings can be used in conjunction with the Safe Search option to tailor the system to your requirements.

YouTube EDU

The YouTube EDU service filters YouTube content for an educational environment. This is different than YouTube Restricted Mode, which is a subset of Google's Safe Search feature. When using YouTube EDU, users access the YouTube EDU home page rather than the standard YouTube home page.



The image shows a configuration dialog box titled "YouTube EDU" in red text. The dialog has a light gray background and a white border. At the top right, there are icons for help (?) and close (X). The main content area contains the following elements:

- "Enable Youtube EDU" with an unchecked checkbox.
- "HTTP Header" with a text input field containing "X-YouTube-Edu-Filter".
- "Custom ID *" with an empty text input field.
- A red asterisk followed by the text "* Required Field" below the Custom ID field.

At the bottom of the dialog, there are two buttons: "OK" and "Cancel".

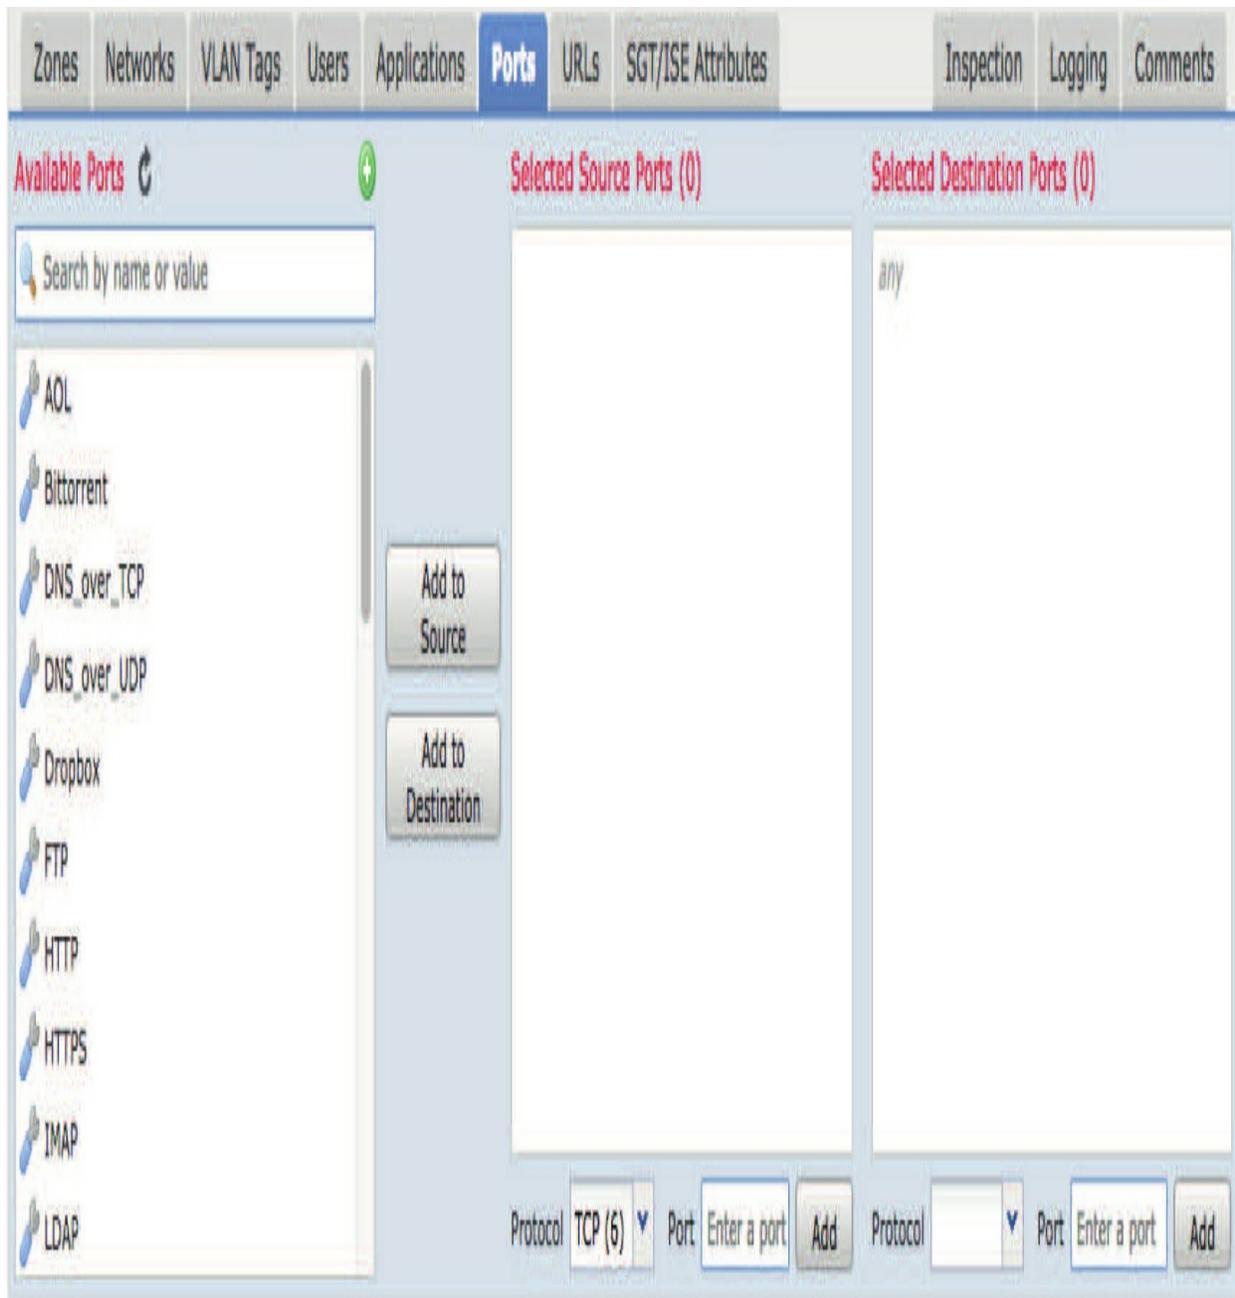
To use this feature, you must enter a custom ID that uniquely identifies a school or district network. This ID is provided by YouTube when the school registers for a YouTube EDU account. If you configure both Safe Search and YouTube EDU, you should place the YouTube EDU rules higher in your policy to avoid preemption.

Clicking the YouTube EDU icon (



) displays the dialog shown in the next figure. Once you check the Enable check box, you must then enter the custom ID for the school or district: **Ports**

The Ports tab is shown below. It is fairly straightforward and allows selecting source or destination ports for your rule:



Port selection works just as you would expect. You'll notice the source ports are limited to UDP and TCP, while destination ports include all of the IP protocol ID numbers. This isn't really a limitation; it's just the way IP protocols work.

These other protocols don't have port numbers, so if you select one of them, the destination port field is grayed out. Selecting ICMP (Internet Control Messaging Protocol) will display an additional dialog allowing you to select

the Type and Code fields in the ICMP packets.

URLs

The URLs tab allows filtering based on Cisco-provided categories and reputations as well as user-created URL objects. There are two sub-tabs, one for Category and the other for URLs. The Category tab will only be populated if you've purchased and installed the URL filtering license. Reputations are only available when selecting URL categories and this column is grayed out when the URLs sub-tab is selected.

If you are configuring URL filtering on Cisco FTD, you need to make sure that Outbound port TCP/80 and Outbound port TCP/443 are able to communicate with a Cloud Service Provider (CSP).

The URLs sub-tab allows us to use a previously created URL object. Alternatively, you can enter a URL in the field at the bottom of the Selected URLs column. Remember that URLs you enter are treated as a substring match, meaning that if the text you enter is found *anywhere* in the URL, the rule will match!

The URLs tab is shown here:

Add Rule



Name Enabled

Insert into Mandatory

Action Allow

- Zones
- Networks
- VLAN Tags
- Users
- Applications
- Ports
- URLs**
- SGT/ISE Attributes
- Inspection
- Logging
- Comments

Categories and URLs

- | Category | URLs |
|--------------------------|------|
| Dating | |
| Dead Sites (db Ops only) | |
| Dynamic Comment | |
| Educational Institutions | |
| Entertainment and Arts | |
| Fashion and Beauty | |
| Financial Services | |
| Food and Dining | |
| Gambling | |

Reputations

- Any
- 5 - Well Known
- 4 - Benign sites
- 3 - Benign sites with security risks
- 2 - Suspicious sites
- 1 - High Risk

Selected URLs (0)

any

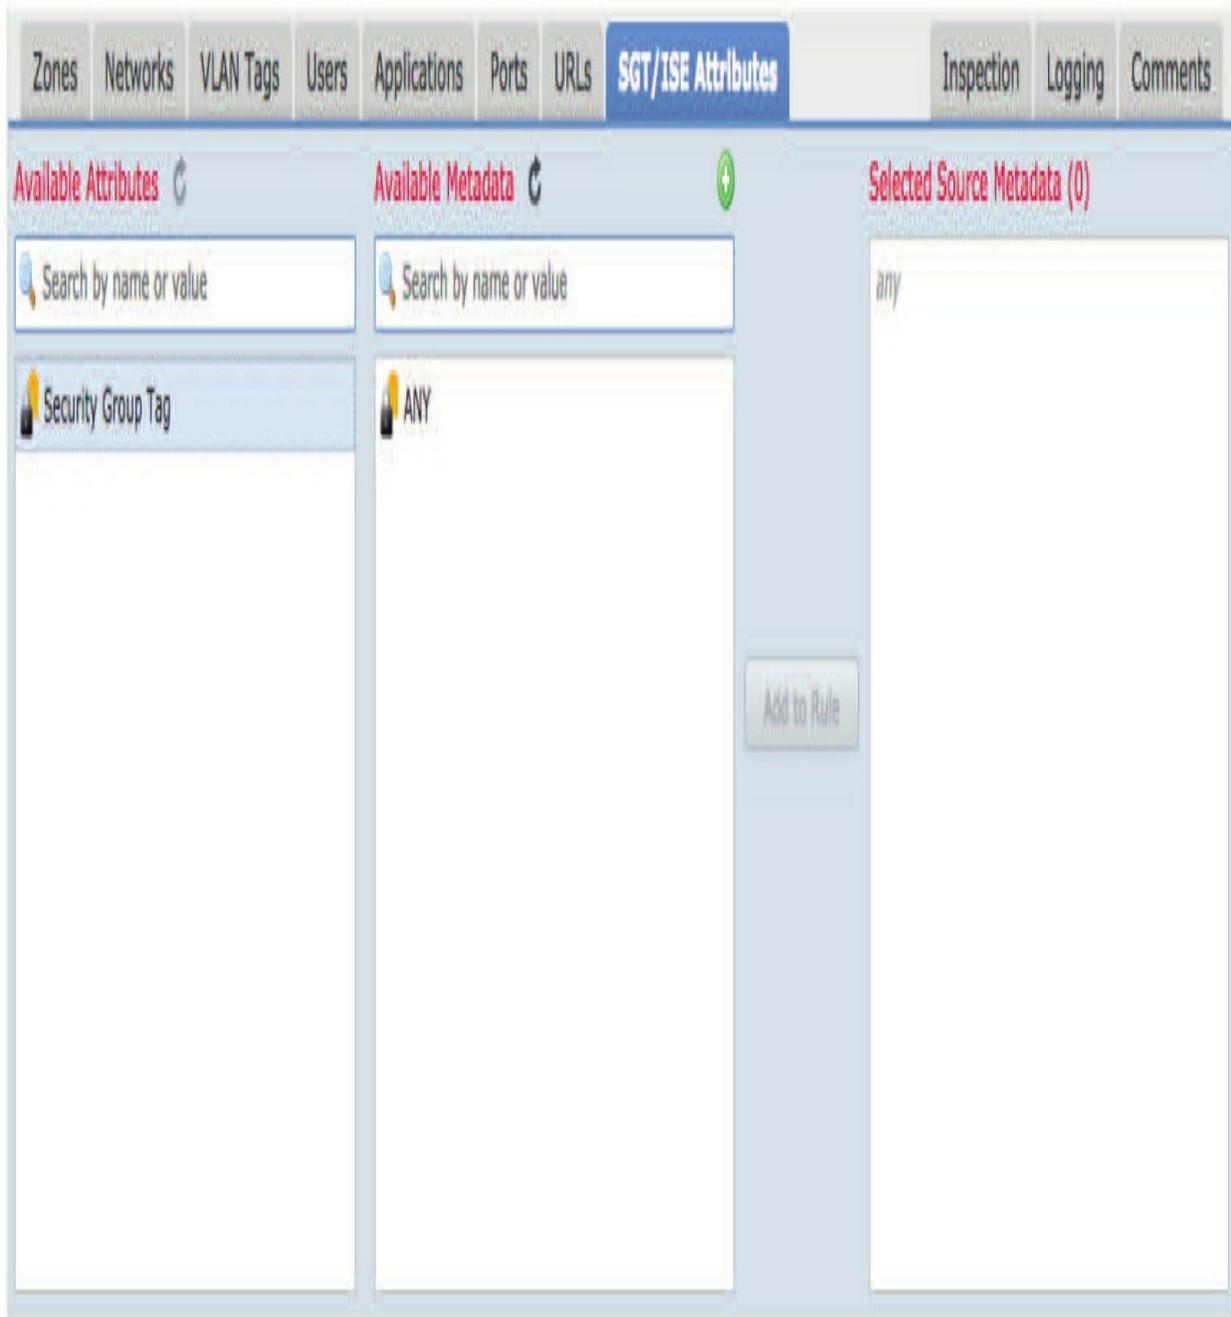
Add to Rule

SGT/ISE Attributes

The SGT/ISE Attributes tab provides a method to match traffic based on Security Group Tags (SGTs) or through the Cisco Identity Services Engine (ISE). If you have an ISE server configured, the Available Metadata column is populated by querying ISE for available tags.

If you don't have an ISE server configured, you can create custom SGT objects and use them as conditions on this tab.

The SGT/ISE tab is shown in the following figure:

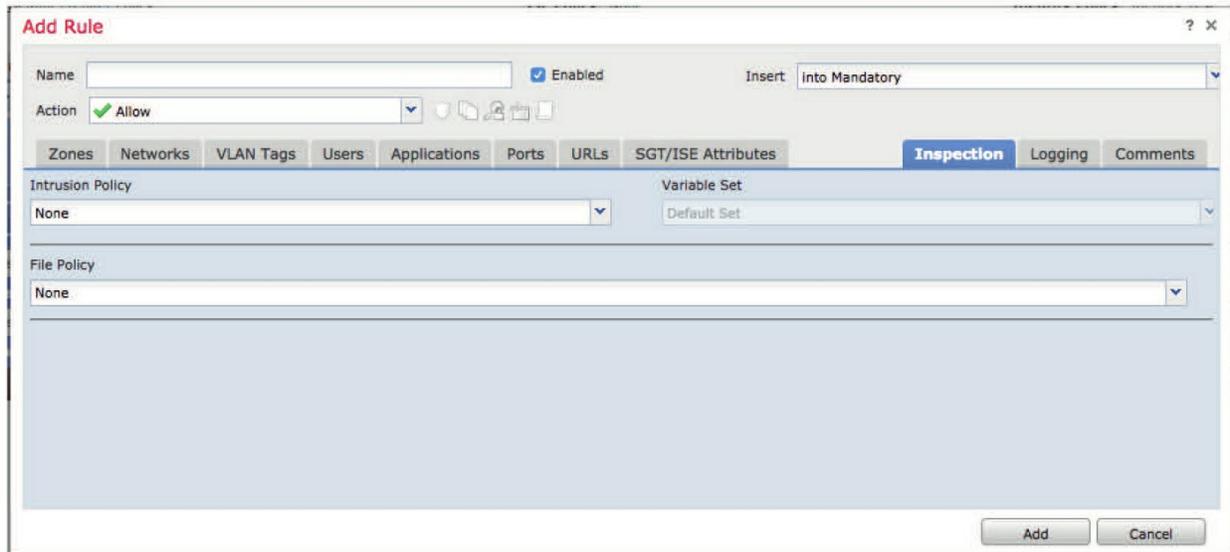


Inspection, Logging, Comments

The three tabs to the right side of the Add Rule dialog don't impact which traffic a rule will match so you really could call them additional rule settings.

Inspection

The Inspection tab controls intrusion and file inspection for rules that provide this capability.



that all of the options on this tab are unavailable and will be grayed out for the rule actions of Trust, Monitor, Block, and Block with Reset, but our Allow rule shown below, is not greyed out.

This is because these rules don't allow for Intrusion or File policy inspection of traffic, only Allow does. The following figure pictures the Inspection tab dialog:

Use the drop-down fields to specify the Intrusion Policy, Variable Set and File Policy settings used to inspect this traffic. Selecting a File policy automatically enables and selects the File Events check box on the Logging tab. It's important to keep in mind that if you want to perform intrusion and file inspection on a given traffic flow, you've got to enable both of these options in a single rule. You can't perform intrusion inspection in one rule and file inspection in another because the second rule will be preempted by the first one!

Logging

The Logging tab shown below controls the logging of connection and file events for the rule. This tab doesn't impact the logging of intrusion or malware events because these events are generated based upon settings in the Intrusion and File policies, respectively. To clarify, a connection event records information about a network connection and a file event records

information about a file transfer. Connection events are generated by traffic matching AC rules with the log option enabled, including Monitor rules. File events will be generated by file transfers matching File policy rules that have the Detect Files action configured.

There are two options for connection logging: you can log at either the beginning or end of the connection. You could also check both boxes and receive events at the beginning and end of the connection, but we use this option mainly for troubleshooting connection logging issues. Logging at the end of the connection is the usual choice.

These end-of-connection events include the beginning and ending timestamps as well as connection data like the number of bytes, packets, protocol, URL, and a bunch of other attributes. If this is a Block or Block with Reset rule, only the Log at Beginning of Connection option is available.

Add Rule

? X

Name Enabled Insert into Mandatory

Action Allow    

Zones

Networks

VLAN Tags

Users

Applications

Ports

URLs

SGT/ISE Attributes

Inspection

Logging

Comments

Log at Beginning of Connection

Log at End of Connection

File Events:

Log Files

Send Connection Events to:

Event Viewer

Syslog 

SNMP Trap 

Add

Cancel

I'll cover inspection in more detail in Chapters 10 and 12 coming up.
Comments

The Comments tab displays existing comments and allow us to add new comments to the rule. Clicking the New Comment button displays a dialog where free-form text comments can be entered. The User and Date fields are automatically populated. After a comment is entered, you can delete or modify it only until the rule is saved. After that, comments can only be viewed but they can't be modified or deleted.

ACP Rules Example

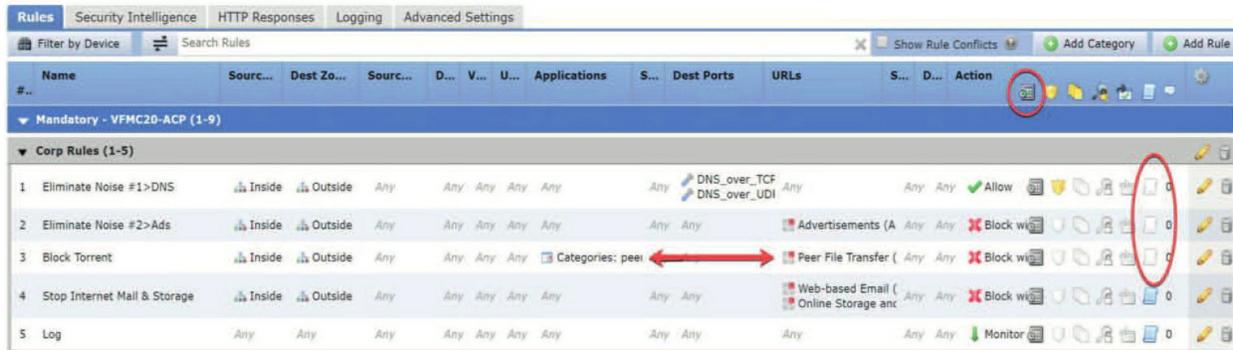
Okay—with all that, I'm going to provide a real-life example of an ACP that I'd build for a typical customer to get started with. Keep in mind that you'd need to customize this for your network and add to it accordingly.

First, always configure your Security Intelligence (SI), HTTP responses, and optional global logging and *then* come back to rules. For SI, add all the Cisco objects for Network IPs, URL, and DNS feeds.

Let's start at the top as shown in the figure of the Rules tab below. Notice that I'm using categories. Remember that these don't make some rules more important than others—rules are always read top down. Categories are only for administrative purposes.

My first Corp rules allow me to make global rules that affect everyone, and if you had parent>child ACPs, these would always be at the top of every ACP.

I want to point out that rules 1 and 2 are used to eliminate noise in the **Analysis>Connection>Events** table by not logging DNS requests or ads on web pages. Trust me, these will fill up your events table even with just one host!



On the right-hand side, you can see that I didn't enable logging on the first three rules. Also, up there on the top, I circled a new feature out in the newest code called Time Range. This allows me to create a time-based object and then apply it to a rule to make a rule active only from 8 a.m. to 5 p.m., for example. I only mention this new feature here because this feature actually just came out while I was writing this chapter! Maybe you'll find a business case for this new feature.

For rule 3, pay special attention to the Application and URLs configuration. You *cannot* have both and expect it to work! I'm just showing you that you can perform blocking with peer-to-peer two different ways in a rule with the same results. But again, don't do both! If you choose the application configuration, it takes about eight packets for Firepower to determine the application flow, so users will get to the website but they won't be able to download torrents. The URL works better but, well, it costs a lot more too, so if you don't have the URL license, you can just use the application.

Remember not to use the Application and URL filter in the same rule because of latency issues!

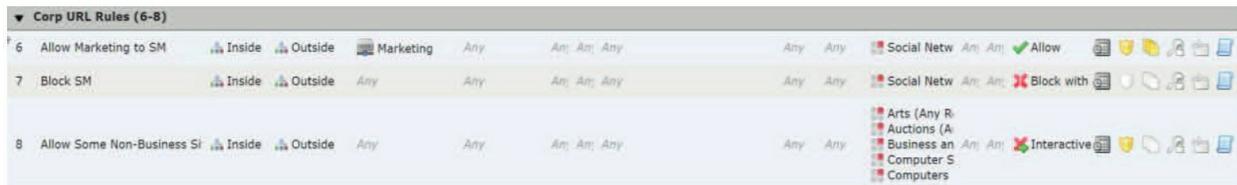
I've used rule 4 with some customers, but this may or may not be useful at your company. Just notice that it blocks any Internet mail (Yahoo or Gmail for example) and cloud storage like Dropbox, iCloud, etc.

Rule 5 will just log every packet. This means for any allow rule under rule 5, you'll no longer need to log those rules because you'd be getting double logging on those packets.

This next section (rules 6–8) is just a broad example of what you can do with

the URL filtering. This assumes you have a URL filter license of course.

Notice that I'm only allowing social media for the Marketing VLAN and the next rule is to block everyone else from these platforms. You can see in rule 6 that I also have inspection on. The basic rule of thumb is that you never let traffic out/in uninspected! Check it out:



Two key things here... First, I'm using the VLAN object that I created for the Marketing VLAN in rule 7. Using an AD group would be a much better solution, but that isn't covered until Chapter 16 in the Identity policy discussion, so I'll come back to this rule in Chapter 16!

Last up is another broad example of allowing non-work traffic with an Interactive Block rule using the Cisco URL categories. You can put up to 50 categories in a rule, so you can create more than one rule if you are considering something like rule 8. I don't recommend putting all Cisco categories in an Interactive Block rule because that just doesn't work well at all. But if I just put a solid block here instead, well, that would definitely work. Your phone would start ringing bigtime, so you'll need to be available for tuning! So be careful to only put in the categories that you really want to block.

Here's another way to stop non-business sites, if that's what you need. Check out this rule, an optional block for non-work sites:

Name

Block Non-Business Sites

Enabled

[Move](#)

Action

 Interactive Block

Time Range

Zones Networks VLAN Tags Users **Applications** Ports URLs SGT/ISE Attributes Inspection Logging Comm

Application Filters C

[Clear All Filters X](#)

Available Applications (360) C

Selected Applications and Filters (1)

Q Search by name

Q Search by name

▼ Risks (2 Selected)

- Very Low 443
- Low 457
- Medium 315
- High 226
- Very High 134

▼ Business Relevance (2 Selected)

- Very Low 253

All apps matching the filter

- 100Bao 
- AcFun 
- Addictive Mobility 
- Adult Friend Finder 
- Adult World 
- AJP 
- Amazon Cloud Player 

[Add to Rule](#)

Filters

Risks: High, Very High; Business Relevance

poor man's firewall, just as we've been using for well over 20 years!



I enabled logging on rule 12 and remember that it's very important to not forget to log the default action as well so I used a red arrow so you'll see it and hopefully always remember to log the default action!

Summary

You learned all about the Access Control policy, which determines which traffic will be logged, allowed, or blocked by rules you create. It's the main source of logging, and it's found in Analysis>Connection Events.

You now know that the AC policy is also used to implement Security Intelligence lists, IPS rules, and File policies as well as SSL, Identity, DNS, and your Network Analysis policy (NAP).

All traffic passing through a device is processed through the AC policy, which bears some resemblance to a traditional firewall access control list (ACL).

The ACP is the main policy and the only policy you have to have, so it's great that you now understand how this works and how it deploys to your devices!

The last thing you were shown in this chapter was a solid, realworld example of an ACP I'd start with at a typical customer to bring home the practical application of this key aspect to Firepower security.

Chapter 10: Malware and File Policy

The following CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

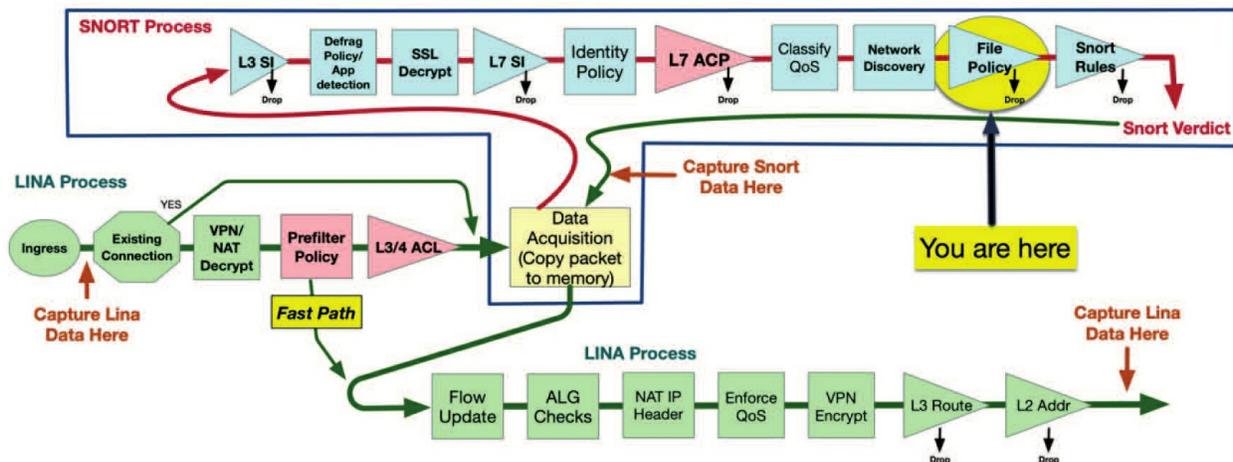
2.2 Configure these policies in Cisco Firepower Management Center

2.2.c Malware and file

4.0 Integration

4.1 Configure Cisco AMP for Networks in Firepower Management Center

4.2 Configure Cisco AMP for Endpoints in Firepower Management Center



I've got to say, one of my favorite things about Firepower is that it actually lets us carve out files from network flows, reassemble them, and then perform an action. There are two very cool tools that combine to give us this impressive ability, Advanced Malware Protection (AMP) and File Control, and they're both controlled by a single policy type.

AMP is designed to address the most prevalent threat vector in use today—malware. This malicious software is intended to damage or disable a computer system, steal data, or carry out some other ugly deed on an unsuspecting victim. It can exploit a software bug or other weakness, but by far, the vulnerability malware targets the most is the user!

Cisco has integrated AMP capabilities into several products, including Firepower and AMP for Endpoints, plus the email (ESA), and Web Security Appliance (WSA). All of these are a lot alike when it comes down to their detection methods, and they all use the centralized Cisco Collective Security Intelligence (CSI) cloud as their main source of information.

The second tool, File Control, gives Firepower the ability to perform actions on files based on the file type, like blocking executable file types transferred

via HTTP. Unlike malware detection, this action doesn't require reassembling the entire file, just the first part that contains the header information.

I'm going to briefly cover AMP for Endpoints before moving into some hands-on lab examples. After that I'll configure my pair of 2500 FMCs with a File/Malware policy and push it to my 1050s' FTDs, then verify the policy is actually blocking malware. This is going to be a fun chapter!

To find exam study material like videos, downloadable supplemental material, and practice questions, just head over to www.lammle.com/firepower.

Advanced Malware Protection (AMP) Basics

Modern networks definitely make it super easy to communicate, share information, and collaborate with other users and businesses. The barriers and hoops we had to jump through to download and run software before today's effortless connectivity have been reduced to a click of the mouse. This is all really great, except that way too many users are totally unaware that there's just no way to tell exactly what a binary executable program is going to do before executing it. Most people just trust it to do what it's supposed to do and click away! Worse, the program can usually do anything within the context of the user's permissions. So if that user has root or admin privileges, then the sky is pretty much the limit on what that malware can do on the system, resulting in potentially devastating damage!

A key component of AMP is cloud intelligence. The Collective Security Intelligence cloud continuously processes samples of files it gets from lots of sources. These files are run through a series of checks comparing them to known malware or executing them in a safely contained environment called a sandbox. The files are then assigned a disposition. Sometimes, third-party intelligence feeds provide hash values and dispositions for malicious software without the actual files themselves. For these trusted sources, the AMP cloud may convict the file without actually testing and analyzing it in a sandbox environment.

File Analysis

AMP uses this cloud intelligence to block and/or alert on known malicious software before it reaches the endpoint target using these techniques:

- SHA-256 hash
- Static file fingerprint (Spero)
- Local malware analysis
- Dynamic analysis

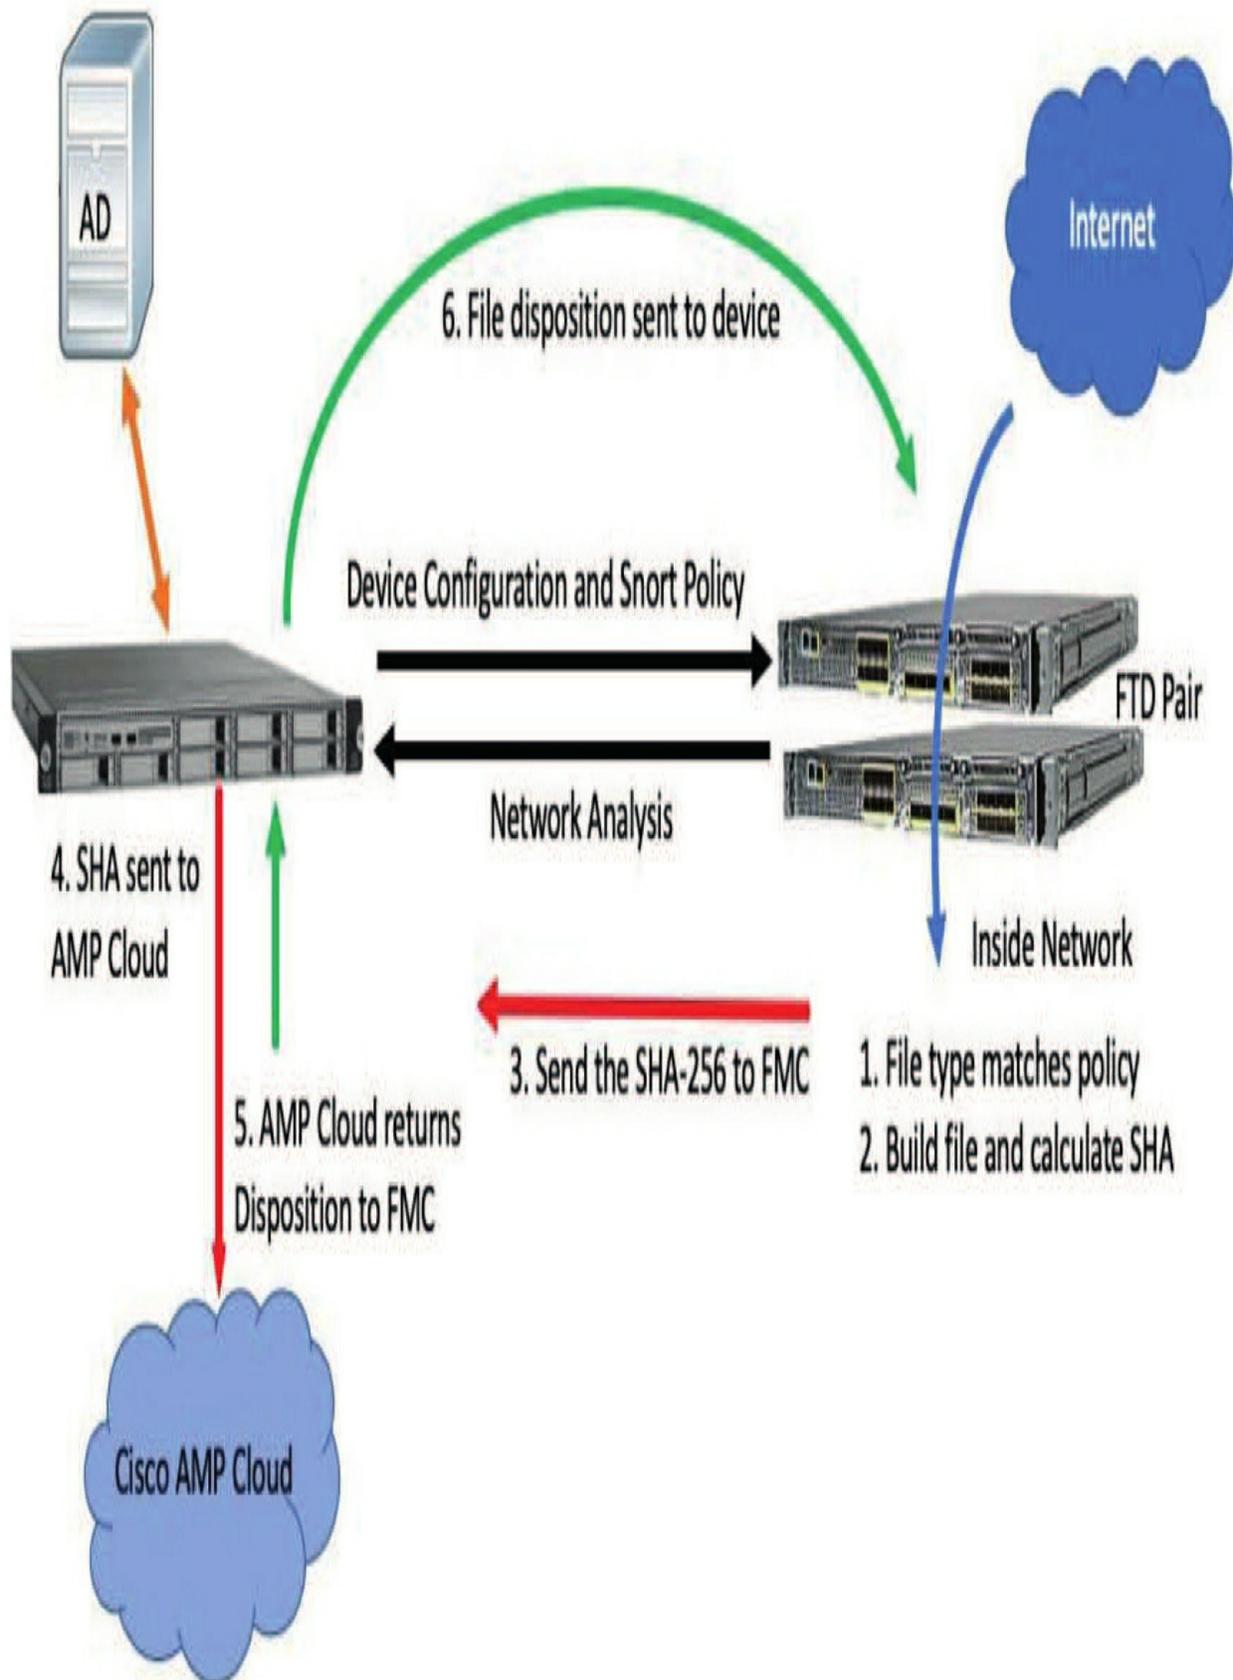
SHA-256

The first technique used is to calculate the SHA-256 hash for the file in question. Each file has a unique hash value. The Firepower/FTD device carves the file out of the network flow, calculates the SHA-256, and transmits it to the FMC. In turn, the FMC checks its local cache and, if necessary, transmits the hash to the Cisco AMP cloud.

A disposition is then returned to the FMC, which it then forwards to the device. If the disposition is malware, the device can block the file and maybe even store it depending on file policy settings. This lookup happens really fast—usually under 600 milliseconds- more like about 200ms!

I want to point out that the file itself never leaves the device in this case—only the SHA-256 hash is transmitted.

Here is the figure, and the 6 steps, that I used in chapter 1 to describe the SHA lookup when packets match a File policy rule on a Firepower device.



Spero Analysis

The second technique is a static file analysis called Spero, which collects a number of file attributes and generates a signature of these values. They include file header information, dynamic-link libraries (DLLs), and other static metadata information. The idea here is that even if the file isn't exactly the one we're looking for, there are probably signs within it that'll peg it as malware. The Spero signature is analyzed and the Cisco cloud arrives at a malware score.

This score is returned to the FMC, which forwards it to the device. If this score is high enough, the file could be convicted as malware. As with the SHA-256 method, Spero analysis doesn't require the file itself to be transmitted to the cloud.

Local Malware Analysis

This method involves the device inspecting a file using a local malware inspection engine known as ClamAV, an open source antivirus and anti-malware toolkit maintained by Cisco. Even though this isn't actually a full implementation of ClamAV, it runs the files through a number of high-confidence checks. Local malware analysis can generate a malware alert and block the file if it's suspected to be malware. In addition, a file composition analysis is performed detailing a file's properties, embedded objects, and possible malware characteristics. You can find this report via the FMC analysis user interface.

Dynamic Analysis

The last technique differs from the previous three in a few keyways. First, it takes time—7 to 10 minutes or more to return a disposition. Plus, dynamic analysis requires the file to be uploaded to the Cisco cloud and involves executing the file in a sandbox virtual machine environment. As the file executes, its actions are analyzed, including behaviors like these:

- Host IPS/firewall/operating system protection evasion
- Persistence and installation behavior
- Anti-debugging

- Boot survival
- Data obfuscation
- Remote access functionality
- Virtual machine detection
- Network connections

Here again, if the analysis score is high enough, the file will be convicted as malware. By this time the file has already passed through the device, so clearly it can't be blocked, but because the hash is now classified as malware, any future detections of this file will immediately return a malware disposition. And if the score is really high, the conviction will be universal within the Cisco cloud. This means that if the file is ever seen by any AMP-enabled product again, it'll get a malware disposition returned on it, which demonstrates the real power of the cloud. Protection will be provided across the entire Cisco customer base, even to those who haven't actually encountered a specific piece of malware themselves!

Retrospective Events

Another interesting AMP feature is called a retrospective event, which occurs when the malware disposition for a previously detected file changes. When this happens, it pretty much always refers to a file that was previously assigned the disposition "unknown" has been since discovered to be malicious.

A retrospective event can occur minutes or even days after a file was first detected and the change will be reported to the FMC from the Cisco Security Intelligence cloud. The FMC will then update all previous detection events for this SHA-256 value with the new disposition.

So retrospective detection is a powerful AMP feature for sure. Even if a new piece of malware escapes detection at first, once it's true nature has been revealed as malicious, you'll be alerted so that you can go back and find out exactly where it traversed your network and block it on all your Firepower devices!

ESA, WSA, and Endpoint AMP all allow you to block or quarantine

malicious files, and we can only do that because of the AMP cloud database.

This database actually knows every network and every endpoint that has encountered a particular file and shouts out that vital information to all affected customers—nice!

File Dispositions

All of these techniques exist to determine and return one of five possible file dispositions:

Clean

The file is benign, indicating that it's a known, good file. The clean disposition is either returned by the AMP cloud or it has been manually added to the Clean-List.

Unknown

A definitive disposition could not be determined, meaning the file hasn't been seen anywhere yet or a sandbox score wasn't high enough to convict it.

Malware

The file has been categorized as malware either by the cloud or by local malware analysis, or it's exceeded the malware score threshold in the file policy.

Unavailable

This disposition isn't actually returned from Cisco, but it means something prevented the cloud lookup. Cisco tells us a small percentage of lookups return this disposition.

Custom detection

The user has added the file to the Custom-Detection-List.

Archive Files

These receive the disposition of the lowest rated file in the archive. For example, an archive containing all clean files will be marked as clean, but if there's even one unknown file there, it'll be marked as unknown. An archive with only a single malware file in it will be marked as malware no matter

what other files it contains.

File Disposition Caching

To minimize the number of cloud lookups, file dispositions are cached on the FMC. Once a file disposition is returned from the cloud, it's cached so that subsequent lookups of this same hash value won't trigger repeated cloud communications. Here's a list of the cache time-to-live values:

- Clean: 4 hours
 - Unknown: 1 hour
 - Malware: 1 hour
 - Unavailable: Not cached
-
- Custom Detection: This disposition never results in a cloud communication since the FMC maintains the CustomDetection-List locally.

Cloud Communications

The communications architecture for this process is fairly simple and nearly all cloud communications are initiated by the FMC. By default, this communication takes place over port 443.

You can change this port on the FMC under **System>Integration>Cloud Services**.

The next figure pictures the options available for AMP network communications.

AMP for Networks

Last Local Malware Detection Update: Jan 5, 2020 6:13 PM

Enable Automatic Local Malware Detection Updates

Share URI from Malware Events with Cisco

Use Legacy Port 32137 for AMP for Networks

Save

You can change the cloud communications port to 32137 if you want, and even though this is an unusual port that might require a special firewall rule, it can be a good option. Cloud communications are encrypted end to end by the FMC, and when using port 443, this encrypted traffic is encrypted a second time with Secure Sockets Layer (SSL). Of course, encrypting traffic that's already encrypted means more overhead, so if you want efficient cloud communications, using port 32137 would seem like a nice way to go. Just keep in mind that you can't proxy communications on port 32137.

So clearly, if your FMC communicates to the Internet via a proxy, you're stuck with port 443. Plus, Cisco cloud infrastructure happens to be moving away from the use of port 32137 anyway, so efficiency aside, I've got to recommend using port 443 for new installations so you don't have to change it later.

Now I know I said *nearly* all communications are initiated by the FMC, but there's one exception: If you enable Dynamic Analysis in your file policy, files will be uploaded directly from the appliance or FTD device to the cloud. This means you'll have to open outgoing ports on the firewall for your Firepower/FTD devices. Devices can connect directly or through a proxy.

Malware & File Policy

The policy controlling all this behavior is called the Malware & File policy, and I'm just going to call it file policy for short. It controls which application protocols will undergo file inspection, the direction of file transfer, the type of files to inspect, and the action. File policy isn't applied directly to managed devices, it's applied to traffic via a rule in the Access Control policy instead. Usually you add file inspection to the same rule or rules where you specify an intrusion policy, and I'll go much deeper into this in Chapter 14, "Access Control Policy." For now, just know there are no policies created by default.

To create a new file policy, go to **Policies>Access Control> Malware & File** and click the **New File Policy** button to start a new policy.

Fill in the policy name and optional description in the dialog as shown here:

New File Policy



Name

Lammle_File_Policy

Description

Cancel

Save

Advanced Settings

There are two tabs in the file policy: Rules and Advanced Settings. Let's start with the Advanced Settings tab.

Clicking it displays the options shown here.

Lammle_File_Policy

Enter Description

Rules

Advanced Settings

General

- First Time File Analysis
- Enable Custom Detection List
- Enable Clean List

If AMP Cloud disposition is Unknown, override disposition based upon threat score

Disabled

Archive File Inspection

- Inspect Archives
- Block Encrypted Archives
- Block Uninspectable Archives

Max Archive Depth

Enter a value between 1 and 3

2

Okay—the settings shown are the defaults, so you'll probably want to modify these based on your requirements. Let's take a look at your options now.

First Time File Analysis

Submit for file analysis a file that's been detected on the system for the first time, which has an unknown disposition. For this to work, the file must match a rule configured to perform malware cloud lookup as well as Spero, local malware, or dynamic analysis. If you disable this option, files detected for the first time will be given an unknown disposition.

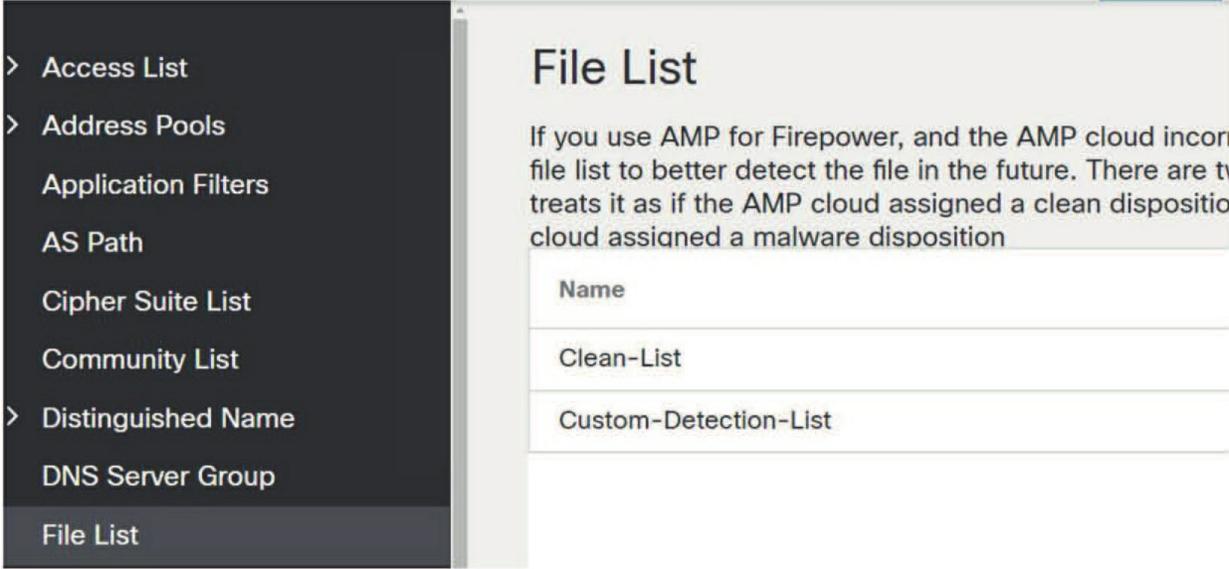
Enable Custom Detection List

You'd check this box if you going to use your Custom-DetectionList object found in **Objects>File List**.

The Custom Detection List contains a list of SHA-256 values that'll always be considered malware and that disposition is stored on the FMC.

Enable Clean List

Choosing this box enables your Clean-List object. These SHA256 values will always be considered clean regardless of their cloud disposition.



The screenshot shows the Cisco FMC Object Management interface. On the left is a dark sidebar with a list of object types: Access List, Address Pools, Application Filters, AS Path, Cipher Suite List, Community List, Distinguished Name, DNS Server Group, and File List (which is highlighted). The main content area is titled 'File List' and contains the following text: 'If you use AMP for Firepower, and the AMP cloud incor file list to better detect the file in the future. There are t treats it as if the AMP cloud assigned a clean dispositio cloud assigned a malware disposition'. Below this text is a table with the following content:

Name
Clean-List
Custom-Detection-List

Here is a reminder of where these objects are found in the main Object management screen. Unlike the SI objects, you can add and removes SHAs from these objects here:

Override AMP Cloud Disposition Based upon Threat Score This should stay disabled now because it's a legacy setting that's no longer used, but you can change it if you want to override the default threat score from Talos.

This option has four settings:

- Disabled (Files will not be considered malware based on Dynamic Analysis.)
- Medium
- High
- Very High

Medium, High, and Very High are the Dynamic Analysis scores that will result in marking a file as malware and will apply to any files analyzed in the Cisco Intelligence Cloud.

Inspect Archives

I definitely recommended enabled, and I'm actually not sure why it's disabled by default. I've just never found a good reason not to look into

archived and compressed files!

Block Encrypted Archives

After enabling Inspect Archives, you can choose to look into the archived and compressed files and block them if they have encrypted contents.

Block Uninspectable Archives

This option blocks archive files that can't be inspected for reasons other than encryption or if they're password protected. It usually applies to corrupt files or files that exceed your maximum archive depth—even to unknown or custom file types. Seemed like a good idea at the time, but I enabled this at a customer site and most of their files were dropped because they used custom applications and Cisco didn't know the file type... Live and learn! Now I usually avoid enabling this option because of that lovely experience.

Max Archive Depth

Block nested archives exceeding the specified depth. The toplevel archive is not counted, so an archive with one nested archive file would be a depth of one. I've set this to 3 at most customers' sites and have not encountered an issue.

File Rules

File policy rules are configured on the Rules tab, and to create a rule, just click the Add Rule button. Doing that brings up the View Rule dialog. Let's go through the options on this page shown here:

Application Protocol: Any

Action: Detect Files Store files

Direction of Transfer: Any

File Type Categories

<input type="checkbox"/> Office Documents	20
<input type="checkbox"/> Archive	20
<input type="checkbox"/> Multimedia	30
<input type="checkbox"/> Executables	15
<input type="checkbox"/> PDF files	2
<input type="checkbox"/> Encoded	2
<input type="checkbox"/> Graphics	6

File Types

Search name and description

- 7Z (7-Zip compressed file)
- 9XHIVE (Windows 9x registr...)
- ACCDB (Microsoft Access 2...)
- ALZ (Archive file for Micros...)
- AMF (Advanced Module For...)
- AMR (Adaptive Multi-Rate ...)

Add

Selected File Categories and Types

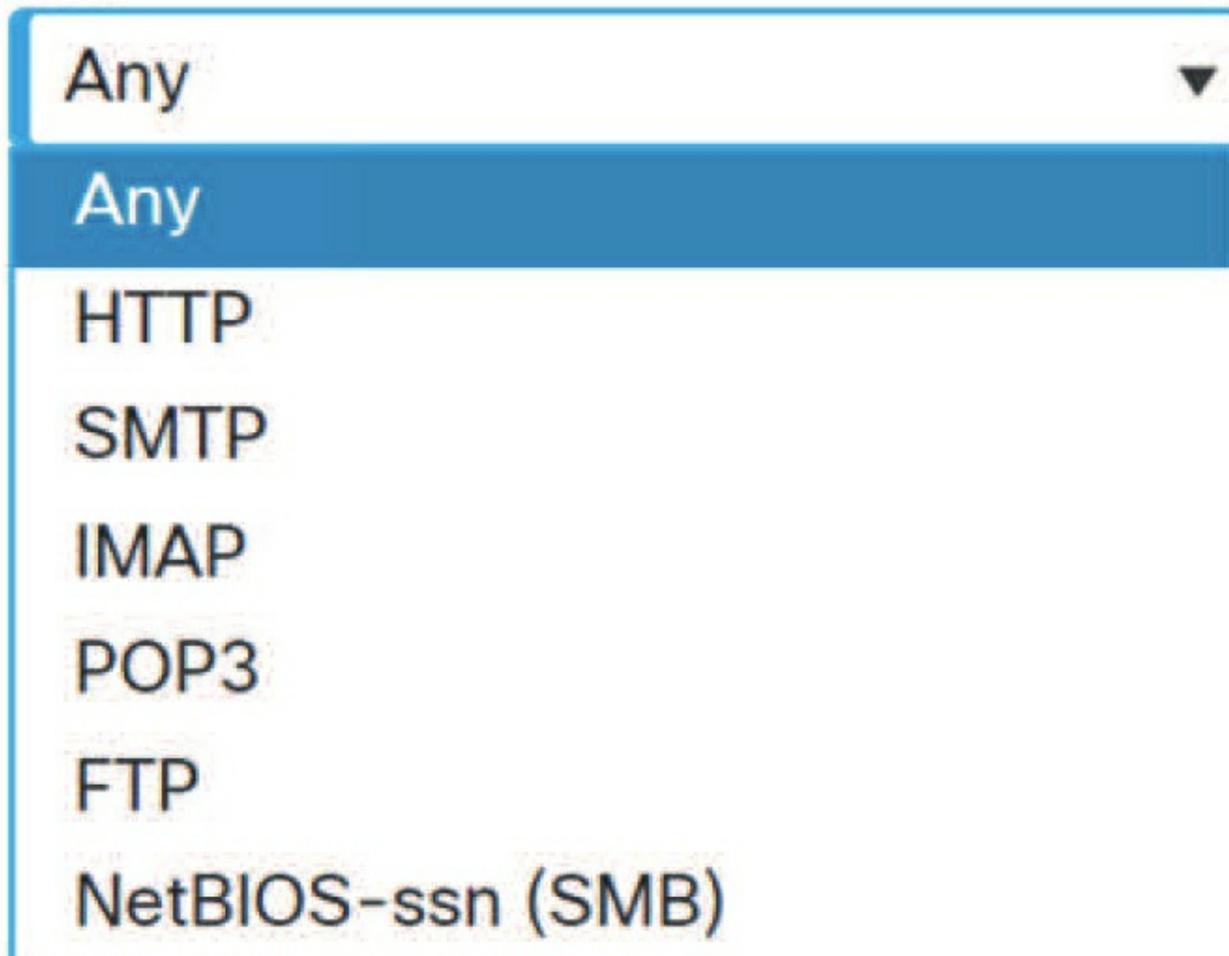
Cancel Save

Application Protocol

The first item on this screen is Application Protocol, and clicking the drop-down reveals the protocols available for file inspection.

Okay so these are pretty self explanatory and they encompass the most common protocols associated with file

Application Protocol



A screenshot of a web-based configuration interface. At the top, the text "Application Protocol" is displayed. Below it is a drop-down menu. The menu is currently open, showing a list of options. The first option, "Any", is highlighted with a blue background and white text, indicating it is the selected option. Below "Any", the following options are listed in black text: "HTTP", "SMTP", "IMAP", "POP3", "FTP", and "NetBIOS-ssn (SMB)". A small black triangle icon is visible at the top right of the menu box, indicating it can be closed.

transfers. Notice they're not encrypted protocols, though, so without the help of a separate decryption solution, Firepower can't peer into connections using protocols like SCP, SFTP, SSL, and TLS.

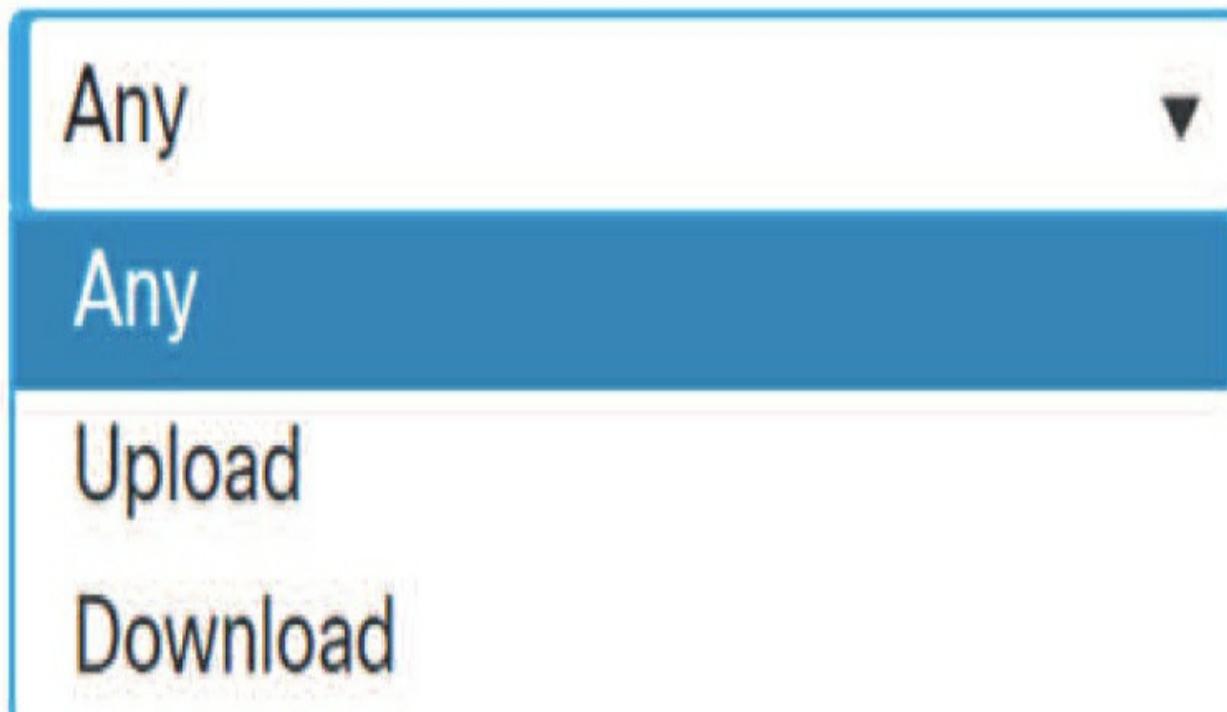
Application protocol selection is what you need to improve performance because it allows for only inspecting a certain protocol with your rule. Choosing Any will detect files over multiple protocols regardless of the file transfer direction.

Direction of Transfer

There, below the Application Protocol option, you'll see the Direction of Transfer option. The values available in this drop-down depend on what's

been selected in the Application Protocol option above it.

Direction of Transfer



If Any is selected for a protocol, then the options here are as follows:

- Any
- Upload
- Download

Know that not all application protocols are treated equally, and some of them are restricted in the direction of file transfer supported. You'll find that the protocols dealing with email (SMTP, POP, IMAP) are restricted to upload or download.

The following table shows the directions supported for each protocol:

Application Protocol	Direction of Transfer
Any	Any, Upload, Download
HTTP	Any, Upload, Download
SMTP	Upload
IMAP	Download
POP3	Download

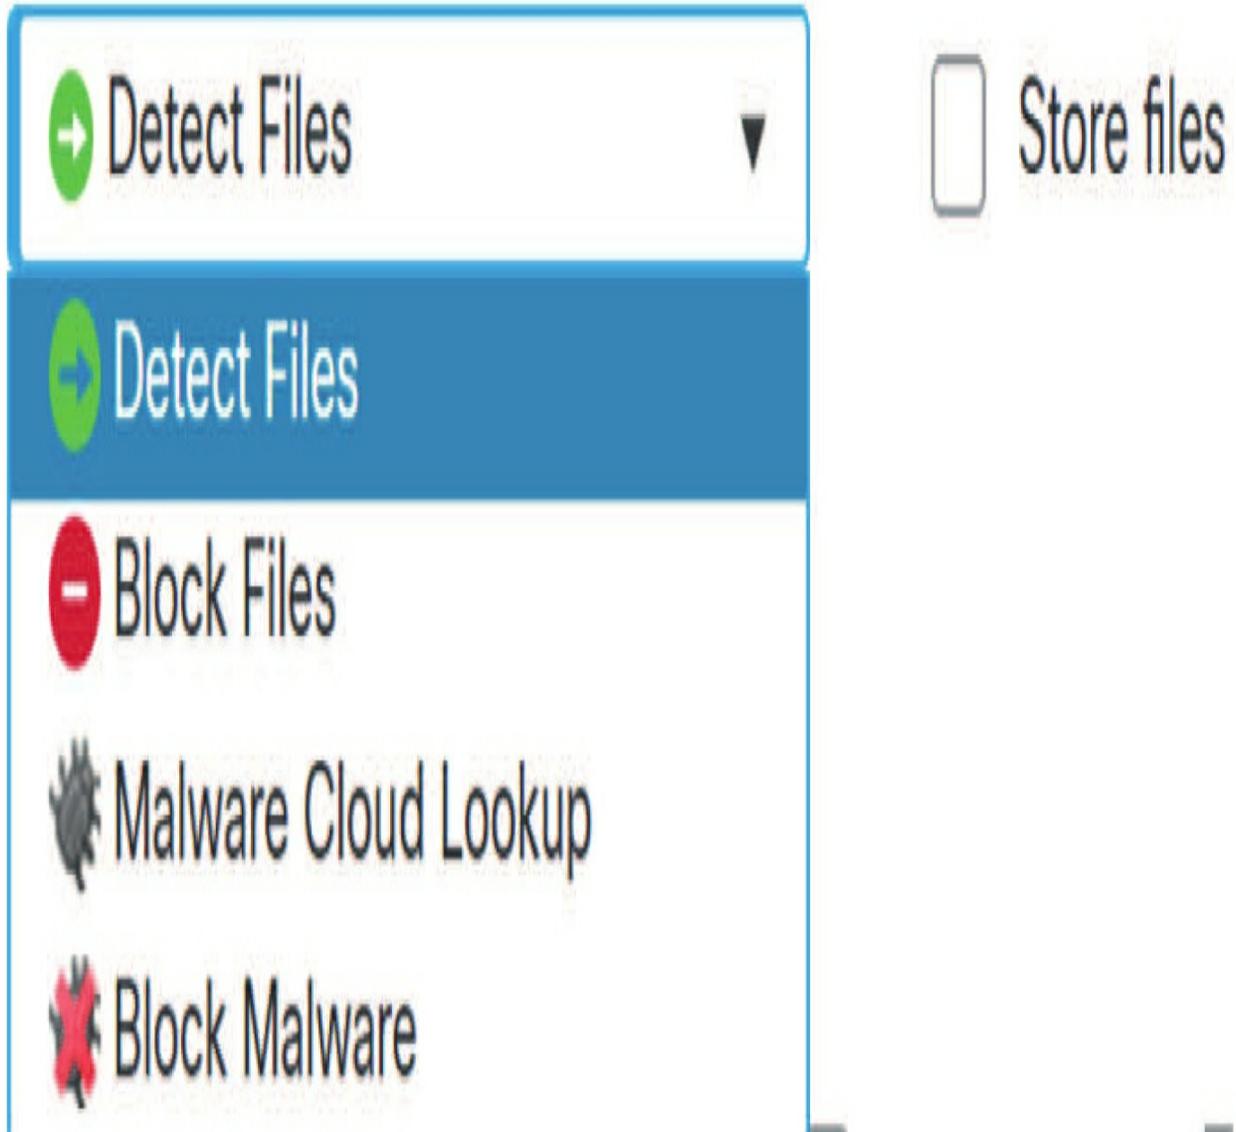
FTP Any, Upload, Download

NetBIOS-ssn (SMB) Any, Upload, Download

File Policy Actions and Licensing

It's really important to understand what you can do with the file policy using actions and the licensing needed for each of these features. When you choose a file action, the licenses are then enforced. The actions, including the storage of files are shown here.

Action



So now, let's look deeper into each of the actions and the licenses needed for them.

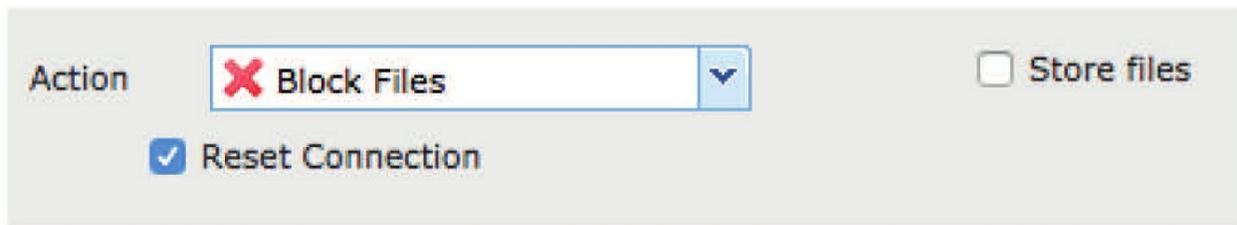
Detect Files

Detect Files rules let you log the detection of specific file types to the database while still allowing their transmission. These rules only log the passage of the file type through the device, and no malware or other analysis of the file will be performed.

This action also doesn't collect the entire file as it passes through the device—only the first 1460 bytes are collected and examined for file header information. This is the information used to determine the file type. If the Detect Files option is checked, files matching the rule will be stored on the device. The corresponding license is Base.

Block Files

Block Files rules allow you to block specific file types. You can configure options to reset the connection when a file transfer is blocked as well as store captured files to the managed device.



This option allows for the detection of and reporting on file types. It also allows you to block or allow all files of a certain type, like all PDFs or all EXE files.

The following licenses are needed to use the Detect, Allow, Block, and Block with Reset functions:

- Threat (for FTD devices)
- Protection (for Classic devices)

Malware Cloud Lookup

The Malware Cloud Lookup action adds a number of options that are pictured here, but it doesn't actually block anything, so why?

Action	 Malware Cloud Lookup	
<input type="checkbox"/>	Spero Analysis for MSEX	
<input type="checkbox"/>	Dynamic Analysis	
<input type="checkbox"/>	Capacity Handling 	
<input type="checkbox"/>	Local Malware Analysis	
		Store Files
		 Malware <input type="checkbox"/>
		 Unknown <input type="checkbox"/>
		 Clean <input type="checkbox"/>
		 Custom <input type="checkbox"/>

This action will log files with a malware disposition while still allowing their transfer through your network. Malware Cloud Lookup requires these three licenses:

- Threat (for FTD devices)
- Protection (for Classic devices)
- Malware

The options for this action are as follows:

- Spero Analysis for MSEX: Checking this box causes the device to send the Spero signature to the FMC if the initial SHA-256 comes back as unknown. It only applies to MSEX (Microsoft Windows executable) file types.
- Dynamic Analysis: Checking this box tells the device to upload the file to the cloud for dynamic analysis. This is only for certain file types and applies to files that have an Unknown file disposition.

Did you know that you're limited to sending 100 files to the dynamic cloud a day? True story. If you need to send more than that, you've got to call Cisco, tell them you want more, then buy an Enterprise Level Agreement (ELA) that allows for the inspection of 3000 files a day, per device! Just don't forget to mention that "money is no object," and you better mean it!

- Capacity Handling: This option is grayed out unless the Dynamic Analysis option has been checked. Hovering over the information icon will get you this note:

Capacity Handling will attempt to store files that cannot be submitted for dynamic analysis to disk for later submission. Selecting this option will use disk space that is available for file storage.

Yes, it's a little weird, but it has to do with the action taken when a file can't be submitted for Dynamic Analysis.

The reasons why can be due to a cloud communication issue or it could be because you've exceeded the file submission limit. If this happens and this option isn't checked, the file won't be saved and it won't be submitted for analysis.

Checking this box will cause the device to save the file so it can be submitted later (I do not recommend this!) The device will resubmit files to the cloud when either of the following conditions is met:

- The device couldn't communicate with the cloud but later reestablishes cloud communications.
- The device reached the maximum number of submissions, but a sufficient time has passed.
- **Local Malware Analysis:** Runs a small version of ClamAV local malware checks on the file. It also provides information used for the file composition analysis that becomes a report available in the file analysis screen on the FMC.
- **Store Files:** This option permits the storage of files based on their

disposition. I don't recommend storing all file types because it can have a ugly effect on device performance and storage ability.

Block Malware

The Block Malware action is a lot like Malware Cloud Lookup except this rule will also block transmission of files marked as malicious. It also adds the option to reset the connection, as shown in the figure.

The Block Files action blocks files based on the file type. And just as with the Detection action, it doesn't collect the whole file. It also permits storing the file. These options are here:

Application Protocol	Action	Store Files
<input type="text" value="Any"/>	<input checked="" type="checkbox"/> Block Malware	<input type="checkbox"/> Malware
Direction of Transfer	<input type="checkbox"/> Spero Analysis for MSEXE	<input type="checkbox"/> Unknown
<input type="text" value="Any"/>	<input type="checkbox"/> Dynamic Analysis	<input type="checkbox"/> Clean
	<input type="checkbox"/> Capacity Handling ⓘ	<input type="checkbox"/> Custom
	<input type="checkbox"/> Local Malware Analysis	
	<input checked="" type="checkbox"/> Reset Connection	

I totally recommend choosing to keep the default of Reset Connection box to prevent blocked sessions from remaining open until the TCP connection resets itself. If not, a low-quality user experience will be the result.

You'll need the same licenses as you did for Malware Cloud Lookup:

- Threat (for FTD devices)
- Protection (for Classic devices)
- Malware

Store Files

With Store Files selected on your FMC, if you want to be able to store files with the any of the file rule actions, or to a Malware storage pack (7,000/10,000 appliances or 2,100 devices), you'll need these three licenses:

- Threat (for FTD devices)
- Protection (for Classic devices)
- Malware

File Blocking Behavior

Now is a great time to talk about exactly how Firepower blocks files classified as malware because you really need to understand this when troubleshooting or observing packet flows on the network.

Blocking a file type is pretty straightforward. The file type is identified by the first 1460 bytes, and Firepower knows right away if it should block the file transfer. If it's blocked, all subsequent packets in the file transfer will fail and an optional but highly recommended TCP reset is sent to both the source and destination hosts.

But blocking malware isn't so simple because to make that happen, the entire file has to be collected in order to analyze it, right? Clearly, this means Firepower has to allow the transfer to continue until the file is complete, which begs the very good question: If we allow the transfer to complete, how the heck do we block the file? The answer to this dilemma comes through allowing the file to *almost* complete. How does that work? Well, Firepower watches the file transfer until the end-of-file marker is seen and then the packet containing it is held. The Firepower/FTD device then completes the file reassembly, calculates the SHA-256, and transmits it to FMC that forwards it to the Cisco CSI cloud. The file disposition is returned and the device can then make a decision on what to do with the last packet. If the file is malware, the packet will be dropped. The file transfer won't complete successfully and the endpoint will never receive the complete malicious file. And for most protocols, any remaining fragments will be removed from the destination host.

The process I just described happens pretty fast. By default, the device will wait no longer than two seconds for a file disposition. If it doesn't receive anything in this time frame, the last packet will be released and the file transfer will be allowed to complete.

File Blocking Notes and Limitations

I'm going to quote these file blocking notes from Firepower online help here because they really illustrate the important nuances of the complex process to blocking behavior for different protocols:

- If an end-of-file marker is not detected for a file, regardless of transfer protocol, the file will not be blocked by a Block Malware rule or the Custom-Detection-List. The system waits to block the file until the entire file has been received, as indicated by the end-of-file marker, and blocks the file after the marker is detected.
- If the end-of-file marker for an FTP file transfer is transmitted separately from the final data segment, the marker will be blocked and the FTP client will indicate that the file transfer failed. The file will actually completely transfer to disk.
- File rules with Block Files and Block Malware actions block automatic resumption of file download via HTTP by blocking new sessions with the same file, URL, server, and client application detected for 24 hours after the initial file transfer attempt occurs.
- In rare cases, if traffic from an HTTP upload session is out of order, the system cannot reassemble the traffic correctly and therefore will not block it or generate a file event.
- If you transfer a file over NetBIOS-ssn (such as an SMB file transfer) that is blocked with a Block Files rule, you may see a file on the destination host. However, the file is unusable because it is blocked after the download starts, resulting in an incomplete file transfer.

File Types and Categories

The final criteria for a file rule is the file type. There are three columns in the rule related to file types, and they're pictured here: The File Types list is in the center column and it shows us all the supported file types that can be inspected with a file rule.

File Type Categories

<input type="checkbox"/>	Office Documents	16
<input type="checkbox"/>	Archive	19
<input type="checkbox"/>	Multimedia	4
<input type="checkbox"/>	Executables	12
<input type="checkbox"/>	PDF files	1
<input type="checkbox"/>	Encoded	0
<input type="checkbox"/>	Graphics	1

File Types

Q Search name and description	
7Z (7-Zip compressed file)	
ACCDB (Microsoft Access ...)	
ALZ (Archive file for Micros...)	
ARJ (Compressed archive fi...)	
BINARY_DATA (Universal Bi...)	
BINHEX (Macintosh BinHex ...)	

Add

Selected File Categories and Types

--

If you know the specific files you want to inspect, you can select one or more from this column and click the Add button to add them to the Selected File Categories and Types column. The File Type Categories column, on the left, exists to make this selection a lot easier because by checking the boxes on the left, you can select all the file types matching a given category. The number to the right of each category indicates how many file types are currently in this category. There can sometimes be some overlap with File Type Categories; for instance, the MSEXE file type is present in both the Dynamic

Analysis Capable and Executables categories.

Let's zoom in on the very cool Dynamic Analysis Capable file category because it contains file types that are eligible for upload to the Cisco cloud for Threat Grid sandbox analysis. The eligible file types are as follows:

- MSEXE (Windows/DOS executable file)
- MSOLE2 (Microsoft Office applications OLE document)
- NEW_OFFICE (Microsoft Office Open XML document format, DOCX, PPTX, XLSX)
- PDF (Adobe Portable Document Format)

Before we can upload a file for dynamic analysis, it has to be preclassified as malware through the local malware analysis or because of its Spero score. Of course, this means that if you enable dynamic analysis and include all the file types above, Firepower will *not* upload all of the Office documents or PDF files it sees. But, the system will still *automatically* upload some files containing information that's confidential or sensitive. Before you enable dynamic analysis for all of these file types, just make sure your organization understands any privacy ramifications surrounding automatically uploading these documents!

Enabling the MSEXE file type for dynamic analysis is usually a lot less risky because executable files don't usually contain any type of sensitive or confidential data. Plus there's added gain to going this route because unknown executables represent a serious risk, so they're prime candidates for dynamic analysis!

Now, once you've selected the criteria for the rule, you're ready to add it to your policy. The rule shown below is configured to inspect executables and system files within all supported protocols in all directions for malware. It'll perform Spero and local malware analysis and upload the file for dynamic analysis if required. Files that exceed the threat score or return a malware disposition from the cloud will be blocked and the connection reset. Just for fun, we're also storing any file classified as malware so we can analyze it ourselves later.

Depending on the Access Control policy rule, a file/malware event will also be logged to the FMC.

Here is a good example of a Block Malware rule. Notice the application and direction are set to “any”, and I’m only storing files that get a Malware disposition, but most importantly I have all options except capacity handling enabled....that options would use ALL your storage up quickly so don’t do this! However, this is a good rule example that works well for most companies.

The screenshot shows a configuration interface for a Block Malware rule. It is organized into several sections:

- Application Protocol:** A dropdown menu set to "Any".
- Direction of Transfer:** A dropdown menu set to "Any".
- Action:** A dropdown menu set to "Block Malware".
- Store Files:** A list of checkboxes: "Malware" (checked), "Unknown" (unchecked), "Clean" (unchecked), and "Custom" (unchecked).
- File Type Categories:** A list of categories with counts: Office Documents (16), Archive (19), Multimedia (4), Executables (12), PDF files (1), Encoded (0), and Graphics (1). Each has an unchecked checkbox.
- File Types:** A search box "Search name and description" and a list of file types: 7Z (7-Zip compressed file), ACCDB (Microsoft Access ...), ALZ (Archive file for Micros...), ARJ (Compressed archive fi...), BINARY_DATA (Universal Bi...), and BINHEX (Macintosh BinHex ...). Each has a trash icon and a checked checkbox.
- Selected File Categories and Types:** A list of selected categories: Category: System files, Category: Graphics, Category: Encoded, Category: PDF files, Category: Executables, Category: Multimedia, and Category: Archive. Each has a trash icon.

Remember, after you send up the first 100 files (by default), you may get some health warnings that no more files are being sent. No big deal, and if you get that a lot you can turn down the health alerts for that issue.

Rule Precedence

File rules are in and of themselves unordered, meaning they’ll be placed into the policy in the order they’re created. No worries here, though, because their order doesn’t affect how they operate.

It’s also possible to have conflicting file rules. For example, if a file policy includes a rule to inspect a certain file type but another rule blocks the same

file type, these rules would clearly be in conflict.

The example next shows a picture of such a file policy that you'd want to fix.

See those warning triangles to the left of the rules? They're pointing out the problem with this policy; rule 1 will block all categories so rule 2 won't perform malware lookup on any files.

Remember that file type analysis only requires the first 1460 bytes of the file? Firepower can determine the type of file that's being transferred by just checking out the file header. It's enough information to make the decision to block a file based on file type. Once the transfer is blocked, the rest of the file won't pass, so we can't calculate the SHA-256 for malware analysis. You just can't block a file and also analyze it for malware!

Keep in mind that you can still apply and use a file policy with warnings as long as you really understand how the policy will behave! However, my example would just block all files. If your feeling bored at work, this policy example would get you some needed attention.

Sample Policy

With everything in hand you've learned this far, let's create a basic File & Malware policy now. Here's a list of the business requirements our policy needs to map to:

- Maximize malware detection and blocking capabilities by using all available detection methods.
- Disallow transfer of executable or system files via HTTP.
- Block encrypted archive files.
- Store malicious files for internal investigation.
- Company policy doesn't allow sending potentially sensitive files to offsite storage locations.

Do you see any problems with our business requirements? We definitely have a bit of a conflict there! The first one says we want maximum detection but the last one says we can't send potentially sensitive files off site. We'll

address this by assuming that the privacy requirement trumps the detection requirement because that's usually the case.

Of course, this means we can't use dynamic analysis on all the possible file types.

So, this policy is one way to address our business requirements. I always start with the Advanced tab so I don't forget to configure this.

Our file policy is set up like this:

- On the Advanced tab, the default settings already include not blocking encrypted archives, which is the default. We also checked Inspect Archives and left Block Uninspectable Archives off. We changed the default Max Archive Depth value to 2.
- Rule 1 blocks the System files and Executables categories for the HTTP protocol.
- Rule 2 performs malware analysis for the MSOLE2, NEW_OFFICE, and PDF file types. This rule does not perform Spero (only for MSEXEX) or dynamic analysis. This addresses the privacy requirement by disabling dynamic analysis for these file types.
- Rule 3 performs malware analysis including Spero, local malware, and dynamic analysis for all file types except the three document file types in rule 2.

File Types	Application Protocol	Direction	Action
Warning 1 Category: System files Category: Executables	HTTP	Any	Block Files with Reset
MSOLE2 NEW_OFFICE PDF	Any	Any	Block Malware with Reset Local Malware Analysis Store files of disposition: Malware
Warning 2 Category: Local Malware Analysis Capable Category: System files Category: Graphics Category: Encoded (19 more...)	Any	Any	Block Malware with Reset Spero Analysis Dynamic Analysis Capacity Handling Local Malware Analysis Store files of disposition: Malware

But there are those warning triangles, alerting us to the fact that we've told

the policy to block some file types in the first rule and inspect them in the third rule! Is this a bad thing? Not really, because in the first rule we're only blocking the System files and Executables categories transferred via HTTP. Of course, the protocol Any in the third rule also includes HTTP, meaning we're telling this rule to inspect the files that were blocked in the first rule, but the fact that the third rule will never inspect these file types is something we're willing to accept so we have a simpler policy.

The limitation that's really causing this conflict here is the fact that a rule can only contain one entry for the application protocol. It would be so great if we could create an inspection rule and specify every protocol except HTTP! But it's basically an all-or-one proposition for now. Each rule can contain Any for the application protocol or it can contain a single protocol.

Just so you know, we could build a file policy with no warnings, but it would require about six more rules. Because we can only include a single protocol (or all of them) in a rule, we would need to add separate rules for each application protocol. Each one would have to include all file types except the three office documents in rule 2 and each rule would be for a specific application protocol. They would all be the same except for the HTTP rule, which would exclude the System files and Executables categories. This is how we could remove the conflict and the yellow triangles, but we'd also have a much more complex policy resulting in increased odds for errors and confusion down the road.

Our sample policy demonstrates a pretty solid and safe way to go about addressing the business requirements. Another way we could go is to use the Access Control policy to apply one file policy to HTTP traffic and another one to all other traffic, but this would require two file policies.

Of course, we could also add the additional six rules to our policy if we really needed to eliminate the warnings. The idea here is to show you that there are often several ways to address a situation with Firepower, which is great! Always use the one that makes the most sense in your specific situation.

Cisco AMP for Endpoints in Firepower Management Center

Cisco's AMP for Endpoints is a separate malware-protection product that can supplement malware protection provided by the Firepower system and be

integrated with your Firepower deployment.

Integrating Firepower with AMP for Endpoints

So far, we've covered AMP for Networks by adding the malware license and then the malware rules in the file policy. We could optionally add AMP for Endpoints on our inside hosts and, also optionally, integrate this product with our FMC without adding any new licenses.

If we did go with adding AMP for Endpoints within our Firepower network, we would gain two key benefits:

- A central black- and whitelist configured in our AMP for Endpoints can determine verdicts for file SHAs sent from Firepower to the AMP cloud for disposition.
- AMP for Endpoints can detect and send the malware events detected into our FMC, making it easy to then manage these events along with malware events generated by the Firepower system.

The thing is, when I've implemented this for clients, they usually send back the disposition pretty quickly because all this data includes scans, malware detections, quarantines, blocked executions, and cloud recalls as well as indications of compromise (IOCs).

That's a lot of data! So if you go with this, plan on your FMC being overloaded with information that you can't keep up with. Implement carefully!

AMP for Endpoints and AMP Private Cloud

If you configure a Cisco AMP private cloud to collect the AMP endpoint data from your network, all AMP for Endpoints connectors (all clients) send data to the private cloud, which forwards that data to your FMC. Understand that this is a private connection, so no data will be sent to any public cloud.

Even if you have configured and implemented an AMP for Networks Endpoints, you need to configure a separate connection for AMP for

Endpoints in the FMC as shown here.

From your FMC, go to **AMP>AMP Management>Add AMP Cloud Connection**.



Name your connection using the closest geographical location. Options are APJC (Asia), Europe, U.S. and a private cloud.

Add AMP Cloud Connection



Cloud Name: APJC Cloud

Use for AMP for Firepower:

APJC Cloud

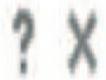
EU Cloud

US Cloud

Private Cloud

If you want to use this cloud for both AMP for Networks and AMP for Endpoints, choose the **Use for AMP for Firepower** check box under the name.

Add AMP Cloud Connection



Cloud Name:

Use for AMP for Firepower:

After doing that, you'll be redirected to log in to the AMP for Endpoints Management console.

Creating AMP Cloud Connection...



Do you want to allow redirection to another page to complete cloud registration?

Yes

No

Log in using your management console credentials, and then click Allow to authorize the AMP to FMC connection.

Understand that if you use an AMP private cloud, you can't use the SI black and whitelists configured in AMP for Endpoints. You'll also have no visibility in AMP for Endpoints on malware events generated from Firepower.

Okay, so at this point, you should see that your connection is now both AMP for Endpoints and AMP for Networks.



Here's the Dashboard on the management console for AMP for Networks.

Dashboard

Dashboard [Inbox](#) [Overview](#) [Events](#) [iOS Clarity](#)

Refresh All Auto-Refresh ▾

Teste-MARCU...(Default)

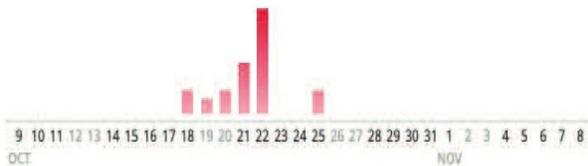
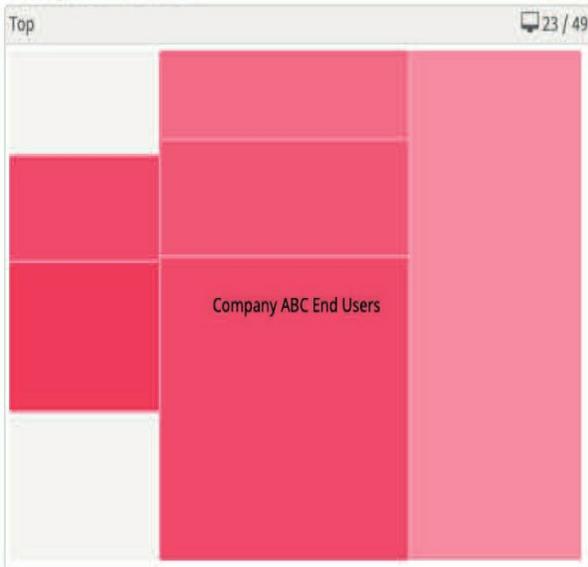
46.9% compromised

Inbox Status

21 Require Attention 2 In Progress 1 Resolved

Compromises

[Inbox](#)



Quarantined Detections

[Quarantine Events](#)



Significant Compromise Artifacts

38 artifacts muted

FILE	1eb15091...22d59900	3	
FILE	b7b28e85...9d0f1d96	ceoflm.exe	
FILE	7c9d5724...f4a6e27a	4543543.exe	
FILE	dd6d4fed...d0de3f91	MspthrdHash.exe	
FILE	92a6e18d...d8655a7a	a.exe	

« < 1 2 3 4 5 6 7 ... > »

Compromise Event Types

3 event types muted

High	Executed malware	13	
High	Cloud Recall Detection	5	
High	Cloud Recall Quarantine Successful	5	
High	DFC Threat Detected	3	
Medium	Quarantine Failure	3	

« < 1 2 3 4 5 > »

Putting It All Together

Here's where I'm going to take everything, I've shown you in this chapter so far with my pod full of great Firepower equipment and build an amazing File policy, implement it, then test it. It would be best if you can log in to a system and follow along!

From your FMC, traverse to **Policies>Access Control>Malware & File**

Overview Analysis Policies

Access Control ▾

Network Discov

Access Control

Intrusion

Malware & File

DNS

Identity

SSL

Prefilter

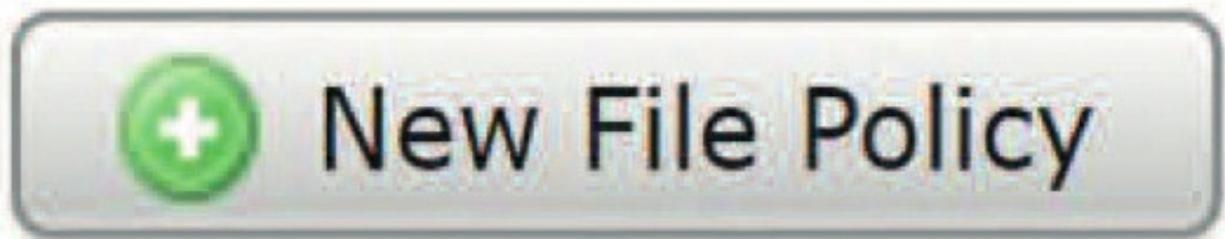
es

Realms

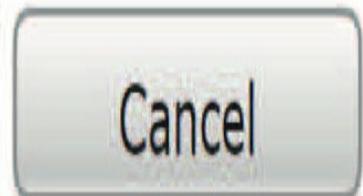
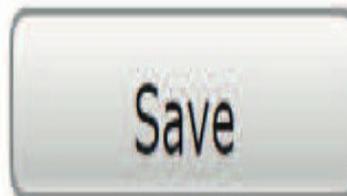
Filtering

URL Filtering Update: S

Click on New File Policy and then name your policy as shown in next:



Name	<input type="text" value="TL_File_Policy"/>
Description	<input type="text"/>



Now click on Add Rule on the right side. From the new screen, choose Block Malware from the Action option, and then select Spero Analysis, Dynamic Analysis, Local Malware Analysis, and Reset Connection. Choose and add all file types, and also choose to store malware. Your rule should look like

this:

Application Protocol: Any

Direction of Transfer: Any

Action: Block Malware

Spero Analysis for MSEXE
 Dynamic Analysis
 Capacity Handling 
 Local Malware Analysis
 Reset Connection

Store Files

Malware
 Unknown
 Clean
 Custom

File Type Categories

<input checked="" type="checkbox"/> Office Documents	15
<input checked="" type="checkbox"/> Archive	17
<input checked="" type="checkbox"/> Multimedia	2
<input checked="" type="checkbox"/> Executables	9
<input checked="" type="checkbox"/> PDF files	1
<input checked="" type="checkbox"/> Encoded	0
<input checked="" type="checkbox"/> Graphics	0
<input checked="" type="checkbox"/> System files	2
<input type="checkbox"/> Dynamic Analysis Capable	5

File Types

Search name and description

All types in selected Categories

- 7Z (7-Zip compressed file)
- ACCDB (Microsoft Access 2007 file)
- ARJ (Compressed archive file)
- BINARY_DATA (Universal Binary/Java Bytecode)
- BINHEX (Macintosh BinHex 4 Compressed archive)
- BZ (bzip2 compressed archive)

Add

Selected File Categories and Types

- Category: System files
- Category: Graphics
- Category: Encoded
- Category: PDF files
- Category: Executables
- Category: Multimedia
- Category: Archive
- Category: Office Documents

Now create a rule that will passively detect all file types after they pass through the Access Control policy (ACP). Again, click **Save**. Your two rules should look like the next figure. Notice the no Access Control Policies use this Malware & File policy. Let's fix that!



Traverse to your Access Control policy, and then click on Add Rule. Create a permit IP any any rule Allow rule (not shown), and inside the Inspection tab (yes, shown), choose your file policy in the File Policy drop-down. After that, click on the Logging tab on the right of Inspection

Name: Inspect all Files for Malware Enabled Insert: into Mandatory

Action: Allow

Zones Networks VLAN Tags **Users** Applications Ports URLs SGT/ISE Attributes **Inspection** Logging

Intrusion Policy: None Variable Set: Default Set

File Policy: None

TL_File_Policy

Finally, choose the Logging tab and then click on log at the beginning and at the end (again, not shown, but you should know this!) and then click **Add**.

Your rule should be similar to the one shown here.

Notice the Allow rule, the File policy enabled, and the fact that

Logging is enabled.

The last thing is to deploy your configuration to the device(s). Next, I'm going to verify the policy and see if this is all working well.

Verifying a File Policy

After the deploy has finished, from an inside host on your network, traverse to <http://2016.eicar.org/105-0-Download.html>. From here, I'll download sample malware files.

Now I'm going to click on one of the standard protocols http download files shown above, and the screen should return something like what's shown next. If yours doesn't come back like this, check your configuration!



This site can't be reached

The connection was reset.

Try:

- Checking the connection
- Checking the proxy and the firewall
- Running Windows Network Diagnostics

ERR_CONNECTION_RESET

Next, I'll traverse to **Analysis>Files>Malware Events** on my FMC, as shown.



Finally, from here, I'm going to verify that I've received a malware event. Here's an example:

Malware Summary [\(switch workflow\)](#)

[Malware Summary](#) > [Table View of Malware Events](#)

No Search Constraints [\(Edit Search\)](#)

Jump to... ▼

	Detection Name	File Name	File SHA256
↓	EICAR	eicar.com	 275a021b...f651fd0f

This is going perfectly... Notice the file SHA256. That red icon tells me that I've detected malware! If there was a period in the middle of that red icon, that tells me the file was downloaded to the FMC, and in this case it didn't download it because this was a known test file.

Summary

In this chapter, I covered the Cisco Firepower File & Malware policy in

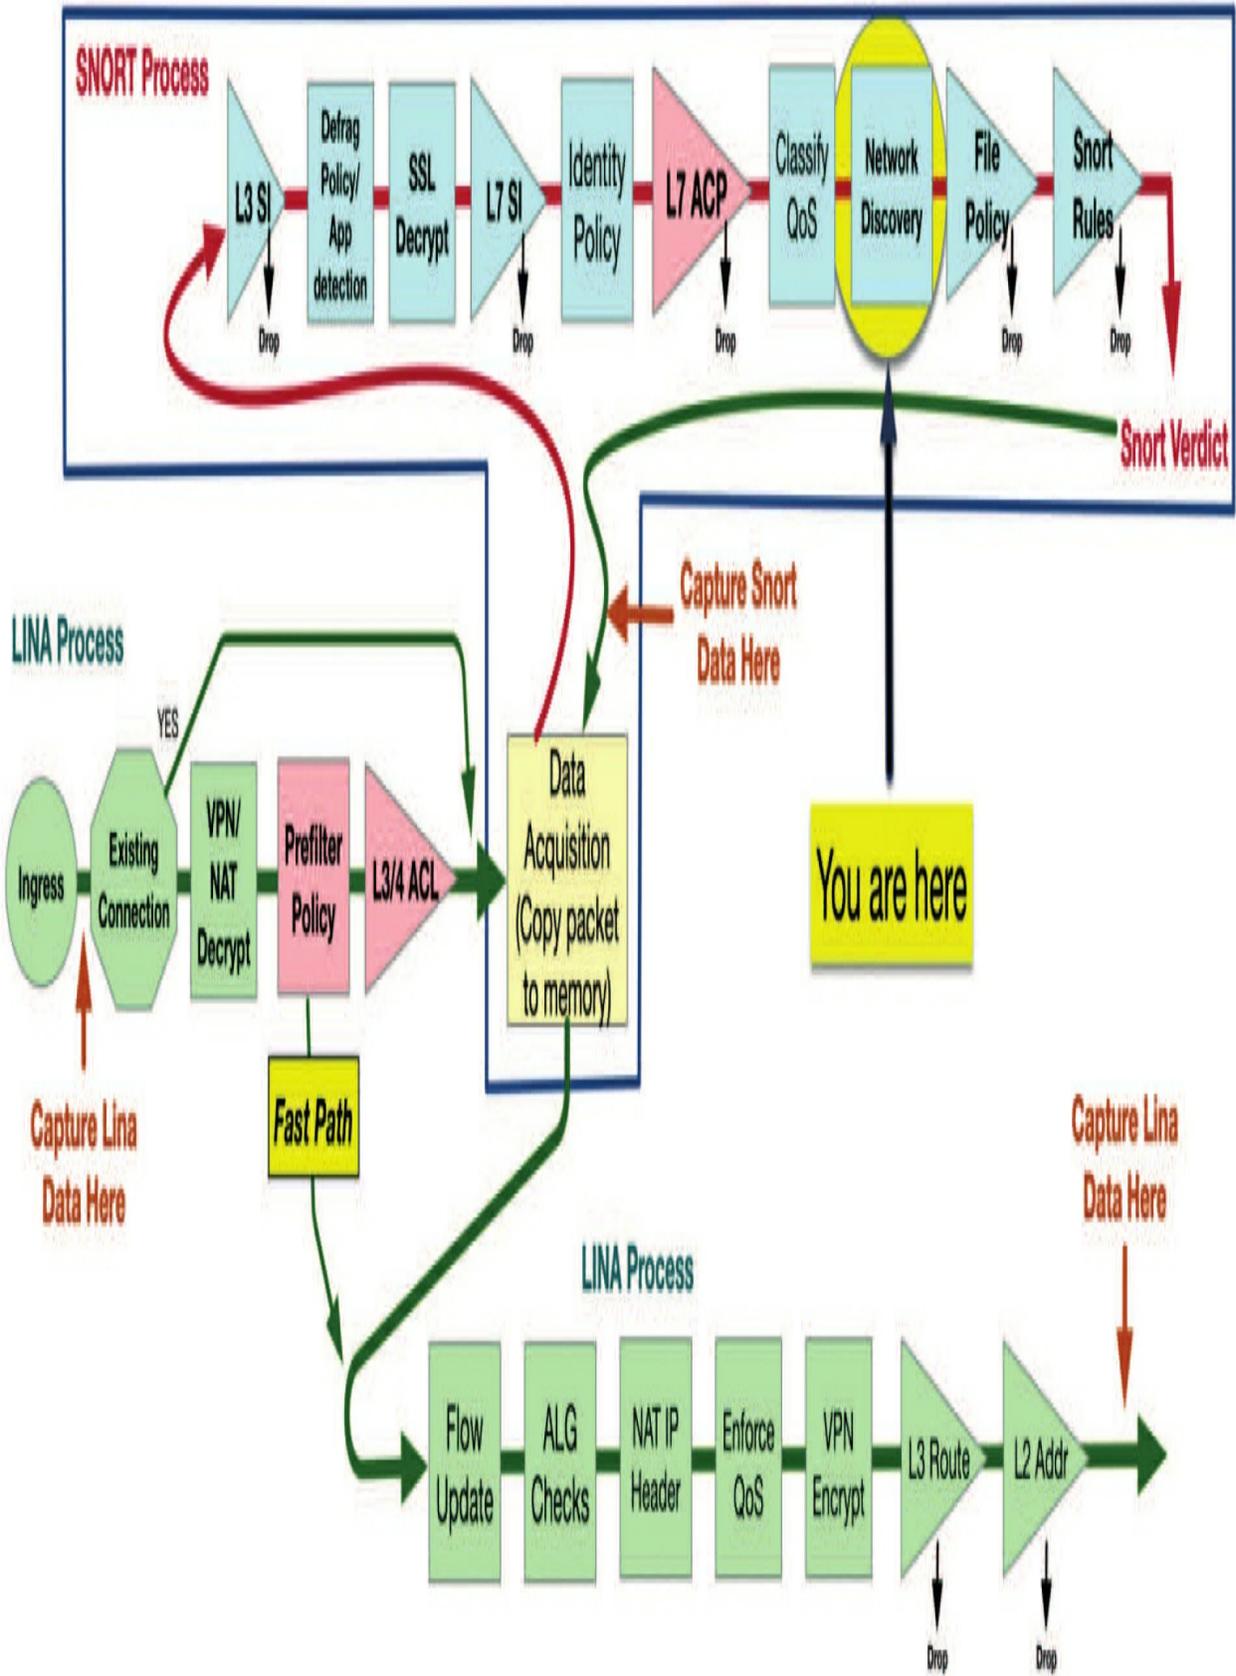
depth. I wrapped it up by configuring my pair of 2500 FMCs with a File/Malware policy, pushed it to my 1050s' FTDs and, verified that the policy is blocking malware. This chapter was fun!

Chapter 11: Firepower Network Discovery

The following CCNP Security SNCF exam objectives are covered in this chapter:

2.3 Configure these features using Cisco Firepower Management Center

2.3.a Network discovery



We'll be going deep into Firepower Network Discovery, formerly known as FireSIGHT, in this chapter. I really think FireSIGHT is actually a better name for this technology. Back in the day, it was originally dubbed Realtime Network Awareness (RNA) and Realtime User Awareness (RUA). This powerful technology generously serves up the key information we need to analyze our data and fine-tune our policies. Cisco just calls it Firepower these days.

After I give you a proper introduction to this awesome technology, we'll move on to explore discovery components like the policy, type of data collected, connection events, and the host attributes associated with it.



Networks	Zones	Source Port Exclusions	Destination Port Exclusions	Action
0.0.0.0/0 ::/0	any	none	none	Discover: Applications

By the end of this chapter, you'll have gained valuable insight into exactly how Firepower Network Discovery is used to powerfully enhance event analysis!

To find exam study material such as hands-on lab access, videos, downloadable supplemental material, and practice questions, please go to www.lammle.com/firepower.

Firepower Technologies

FireSIGHT was the name given to a technology built into the Cisco Firepower NGIPS. It gives us contextual awareness about events, IP addresses, and users on the network and even background about the hosts in the system.

This powerful technology collects information about each IP address and builds a host profile that includes the operating system, services, applications users, and network connections. Firepower will even include assumed vulnerabilities in profiles based upon those factors and additional data it's collected. All of this great information is then used to automatically present

an impact flag in the IPS analysis views that tells us whether or not the systems involved in the IPS events are susceptible to threats.

All of this vital intelligence is gained as a result of analyzing the packets on the wire and leveraging patented passive fingerprinting technology. For this to happen, either packets must traverse the managed device or the device must actually see the traffic itself during passive deployments.

A key benefit to this automated collection process is that it requires no additional software and doesn't probe the network. Plus, the more traffic that's seen moving to and from hosts, the more accurate the information entered into the database will be. And if that's not enough, you can even supplement the Firepower information with active techniques!

On the other hand, when there's only a limited amount of traffic to collect, we can leverage the host input API—a strategy that allows us to import data from a separated values file, Nmap, or even third-party vulnerability scanners—very cool.

Network Discovery Policy

The Firepower *network discovery policy* is configured on the FMC and controls the Firepower technology. There's one discovery policy per FMC, which should be specific to the environment it's being deployed in.

To configure the policy, go to **Policies>Network Discovery**, as shown here. Notice that only Applications are discovered by default.

Did you notice there's no button to create a new policy? That's because you can only have one per FMC, as I already mentioned.

There are three main tabs across the top of the policy: **Networks**: Allows you to define the IP addresses that you want to perform discovery on.

Users : Includes a list of protocols to discover users. **Advanced**: Contains a variety of settings used to further tweak the discovery settings.

Networks

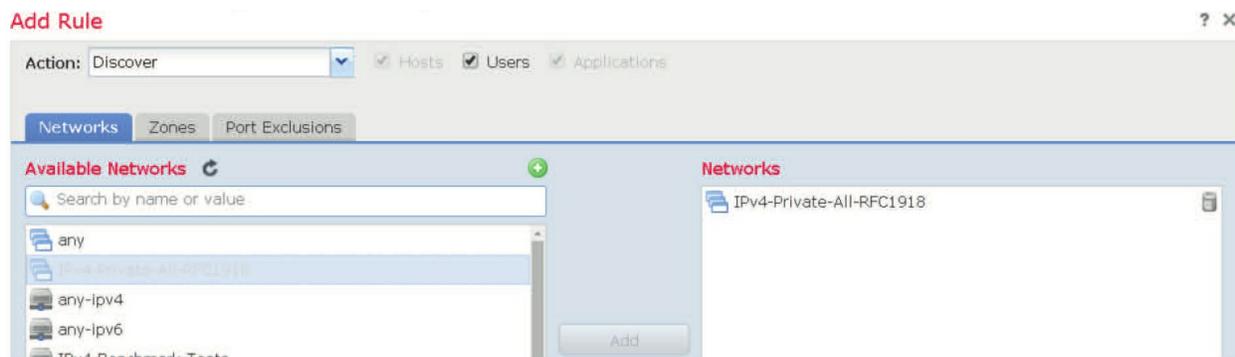
The default discovery takes place on all IPv4 and IPv6 networks as noted, with specified networks 0.0.0.0/0 and ::/0. So, if we see a packet originating to and from anywhere, a host profile will start being generated for the IP addresses involved. The more paranoid among us will be tempted to track absolutely everything but be aware that the FMC is licensed based on model type to only handle a certain number of hosts.

For instance, the Firepower vFMC is set to only 50,000 hosts and 50,000 users. Considering that the clients on your internal network could easily be communicating to hundreds of web IP addresses each, your license could max out its limit pretty fast—choose wisely!

The FMC 1100/1600 is also set to a maximum of 50,000 hosts/ users, the FMC's 2500/2600 can handle 150,000 hosts/users, and the 4500/4600 has a maximum of 600,000 hosts and users. That only seems like a lot until you misconfigure your policy!

The good news here is that you can modify or delete the built-in rule and even add rules to create a specialized policy ideally suited to your particular environment's needs.

Looking at the next figure, notice it's showing some of the options available for the rules. So, if licensing worries are keeping you up at night, just create rules limiting discovery to only the high-value networks in your organization. By adding rules like this, you're basically entering guidance to either discover or exclude individual IP addresses, networks, or network objects. You can also configure discovery based on zones.



By looking at the rule, you can see that I only configured the RFC 1918

addresses to be discovered. I enabled Hosts, Users, and Applications to be added to the database on the top. These check boxes are really important to include!

Another way to free up some resources is to exclude specific ports from discovery. You can also exclude protocols that probably aren't all that important to keep track of because they're unlikely to contain threat data.

While you can specify whether user or host information is collected, application inspection is automatically set to *on* because the Firepower system is designed to be application aware. So just because there's check box that makes it seem like you can deselect it, you really can't. You also don't get to deselect host discovery unless user discovery is also deselected.

User Discovery

This figure shows the Users tab. This tab allows you to focus on the specific protocols that you want to detect user logins on, enabling or disabling them at will.

Typically, you'd just leave these at default, but let's discuss it anyway.

Networks

Users

Advanced

Traffic-Based Detection



aim	Yes
imap	Yes
ldap	Yes
oracle	Yes
pop3	Yes
sip	Yes
ftp	Yes
http	Yes
mdns	Yes
Capture Failed Login Attempts	Yes

And here again, the number of users is restricted just as it was with IP addresses based upon the Defense Center model. Restricting the protocols up for detection here is really helpful for managing your license count.

Here's a list of the supported protocols:

- AIM • SIP
- IMAP • FTP
- LDAP • http
- Oracle • mdns
- POP3

You can also pick up users from Active Directory, but that's handled through an agent and configured in **System>Integration** which we'll cover in the Identity Policy (chapter 16) and the PxGrid policy (chapter 27).

Advanced Discovery Settings

The Advanced tab in the next figure contains settings key to tweak the actual inspection that's being performed on the traffic.

Networks Users **Advanced**

General Settings

Capture Banners	No
Update Interval (seconds)	3600

Identity Conflict Settings

Generate Identity Conflict	No
Automatically Resolve Conflicts	(Disabled)

Vulnerabilities to use for Impact Assessment

Use Network Discovery Vulnerability Mappings	Yes
Use Third-Party Vulnerability Mappings	Yes

Indications of Compromise Settings

Enabled	Yes
Rules	40 / 40

NetFlow Devices

NetFlow Device	
----------------	--

Network Discovery Data Storage Settings

When Host Limit Reached	Drop hosts
Host Timeout (minutes)	10080
Server Timeout (minutes)	10080
Client Timeout (minutes)	10080

Event Logging Settings

All events enabled.

OS and Server Identity Sources

Name	Type	Timeout
Nmap	Scanner	0 hours

Here's another list of the sections these settings are sorted into:

- General Settings
- Identity Conflict Settings
- Vulnerabilities to use for Impact Assessment
- Indications of Compromise Settings
- NetFlow Devices
- Network Discovery Data Storage Settings
- Event Logging Settings
- OS and Server Identity Sources

General Settings

In the General Settings section, you'll find the Capture Banners and Update Interval options. Capture Banners is off by default, but you can enable it to collect the protocol banners for an array of services. Update Interval is set to 3600 seconds by default and specifies how often to refresh data in the database. This data includes things like the last time an application was seen, the last time an IP address was seen, how many times a certain protocol was used, and more. Setting Update Interval to a lower value will display more

recent info on the Firepower manager, but just know that will likely result in more overhead.

Identity Conflict Settings

Identity conflicts can crop up if you're leveraging third-party data like Nmap, host input, and so on for information about the host OSs in your environment. When the data gathered with Firepower conflicts with data from these third-party sources, you can generate an alert indicating that a conflict has occurred. Conflicts can be manually resolved in the host profiles or resolved automatically based upon the options you pick. For example, you can choose Keep The Passive Information From Firepower or Use The Active Information From Other Sources.

Vulnerabilities to Use for Impact Assessment

One of Firepower's greatest features is its ability to automatically correlate vulnerability information with intrusion data. By leveraging this information, you can quickly eliminate false positives from analysis. You can choose to discontinue this feature with Firepower or third-party vulnerability mappings, but just don't ever do that!

Indications of Compromise Settings

Indications of compromise (IOC) offer another way to make hosts stand out in analysis. This next figure shows some of the indications of compromise that are enabled by default. This set includes 31 different kinds of rules that perform important correlations by analyzing data about IPS, vulnerability, file activity, security intelligence, and malware events that point to a "compromised host."

Edit Indications of Compromise Settings

Note: To detect Indications of Compromise, you must enable each IOC rule here and also enable the features, such as Security Intelligence logging and intrusion and malware protection, that the rules below depend on.

Enable IOC 40 out of 40 Rules Enabled

Category	Source	Event Type	Description	Enabl...
Adobe Reader Compromise	Malware Events	PDF Compromise Detected by AMP for Endpoints	Generic Adobe Reader Compromise	<input checked="" type="checkbox"/>
Adobe Reader Compromise	Malware Events	Adobe Reader launched shell	A shell was launched on the host by Adobe Reader	<input checked="" type="checkbox"/>
CnC Connected	Security Intelligence Events	Security Intelligence Event - CnC	The host may be under remote control	<input checked="" type="checkbox"/>
CnC Connected	Intrusion Events	Intrusion Event - malware-cnc	The host may be under remote control	<input checked="" type="checkbox"/>
CnC Connected	Intrusion Events	Intrusion Event - malware-backdoor	The host may be under remote control	<input checked="" type="checkbox"/>
CnC Connected	Malware Events	Suspected Botnet Detected by AMP for Endpoints	The host may be under remote control	<input checked="" type="checkbox"/>
CnC Connected	Security Intelligence Events	Security Intelligence Event - DNS CnC	The host may be under remote control	<input checked="" type="checkbox"/>
CnC Connected	Security Intelligence Events	Security Intelligence Event - URL CnC	The host may be under remote control	<input checked="" type="checkbox"/>

Compromised hosts will show up in any analysis view with a red icon instead of a normal, blue one, making them easy to spot. You can disable these rules individually if you want to.

NetFlow Devices

If you have NetFlow devices in your environment, you can export information from them to your Firepower device to supplement connection information. Keep in mind that the IP addresses involved will deduct from your license count, so definitely choose wisely here again!

Network Discovery Data Storage Settings

These two settings deal primarily with the retention of data in these two ways:

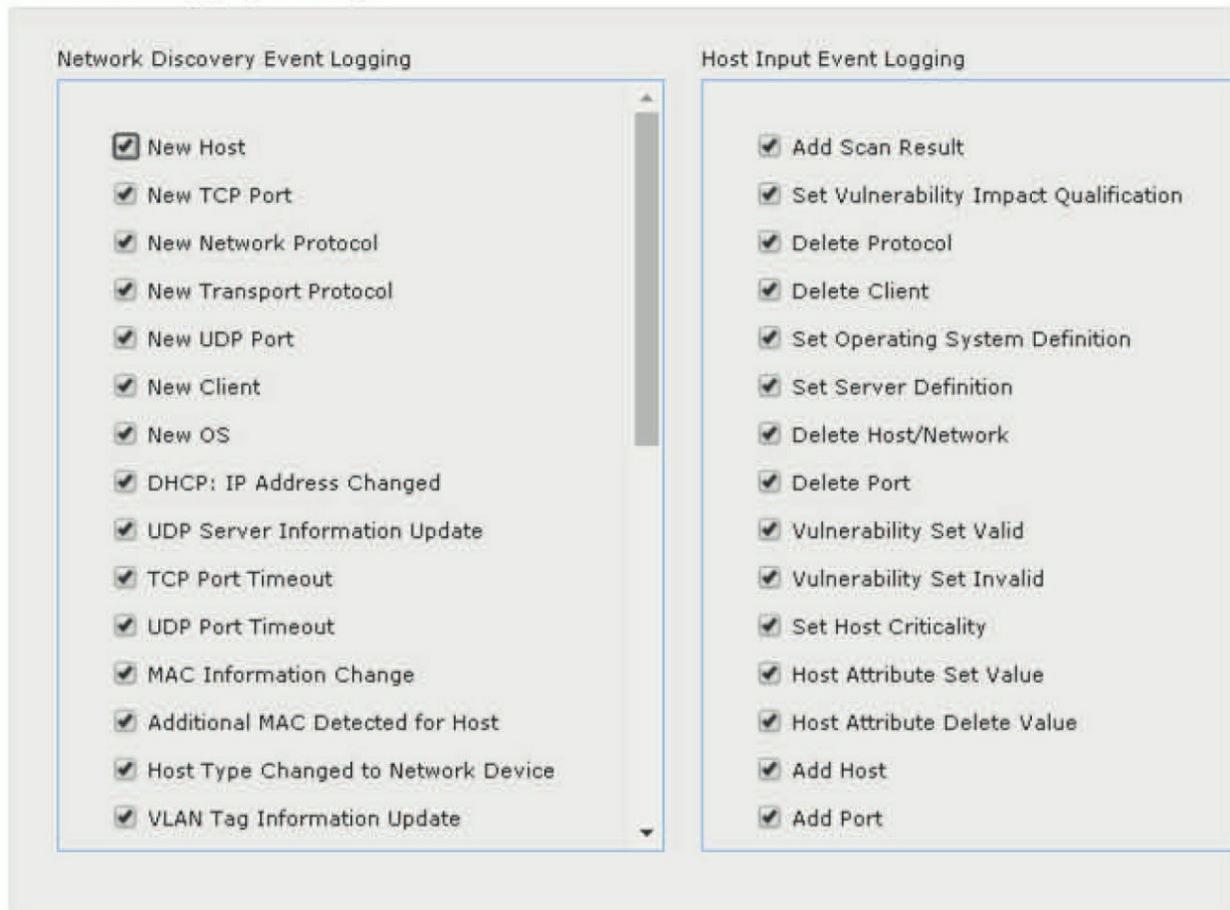
1. **When Host Limit Reached:** When you reach your host license limit, you can choose to either drop newly detected hosts or overwrite older ones. The default is to drop new hosts.

2. **Host, Server, and Client Timeout**—This time-out occurs in minutes and indicates when you'd like to remove information from the database. If one of the items hasn't been seen for the duration of the specified time-out value, the information will be removed. The default time is 110111 minutes—seven days.

Event Logging Settings

There are 33 different types of data that can be logged with Firepower and 20 different settings that can be leveraged through the host input API. All of them are enabled by default, but these data types can be turned off or on as needed. Here are the settings.

Edit Event Logging Settings



The screenshot displays the 'Edit Event Logging Settings' interface, which is divided into two main sections: 'Network Discovery Event Logging' and 'Host Input Event Logging'. Each section contains a list of settings, all of which are checked (indicated by a small square icon with a checkmark).

Network Discovery Event Logging

- New Host
- New TCP Port
- New Network Protocol
- New Transport Protocol
- New UDP Port
- New Client
- New OS
- DHCP: IP Address Changed
- UDP Server Information Update
- TCP Port Timeout
- UDP Port Timeout
- MAC Information Change
- Additional MAC Detected for Host
- Host Type Changed to Network Device
- VLAN Tag Information Update

Host Input Event Logging

- Add Scan Result
- Set Vulnerability Impact Qualification
- Delete Protocol
- Delete Client
- Set Operating System Definition
- Set Server Definition
- Delete Host/Network
- Delete Port
- Vulnerability Set Valid
- Vulnerability Set Invalid
- Set Host Criticality
- Host Attribute Set Value
- Host Attribute Delete Value
- Add Host
- Add Port

OS and Server Identity Sources

This is where you can add additional sources of host identities like Nmap or other third-party applications. You can also specify a time period after which the data becomes stale and the other identity sources take priority over it.

Keep in mind that once you've made changes to the policy, you need to click the Apply button in the upper right portion of the interface to make them stick.

Firepower Discovery Information

Once the Firepower network discovery policy has been created and applied, the managed devices will begin sending information to the Firepower Management Center (FMC). This information can be viewed in a bunch of ways, but to get started, we're going to take a look at **Analysis>Hosts** now.

These are the different views available:

Network Map

A tree view of all IP hosts discovered on the network that's broken out by subnet. You can also type in an IPv4 or IPv6 address/subnet at the top of the list to check out information about specific hosts/networks.

Hosts

A list of hosts organized by operating system.

Indications of Compromise

A list of the IOCs that have been

Hosts ▼

Users ▼

Corr

Network Map

Hosts

Indications of Compromise

Applications

Application Details

Servers

Host Attributes

Discovery Events

Vulnerabilities

Third-Party Vulnerabilities

triggered by category.

Applications

A list of applications along with the number of hosts the applications have been detected on.

Application Details

An inventory of detected application client software and web applications.

Servers

An inventory of application server types along with the application vendors.

Host Attributes

An index of hosts by attribute, which are user-created, definable fields. We'll get into more about these fields in a bit.

Discovery Events

A list of items that were either seen for the first time, (discovery events) or changed in the database, (change events).

Discovery Events

Discovery events are a great way to find out exactly what's popped up on your network.

You can create searches based on subnets and event types that'll present you with a list of hosts based upon their first appearance on the network or network segment, as seen in the next figure.

Discovery Events

[Table View of Events](#) > [Hosts](#)

No Search Constraints ([Edit Search](#))

Jump to... ▼

<input type="checkbox"/>	Time ×	Event ×	IP Address ×	User ×	MAC Address ×	MAC Vendor ×
<input type="checkbox"/>	2017-03-08 16:46:54	New Transport Protocol	 10.131.160.13		34:88:5D:6D:14:46	Logitech Far East
<input type="checkbox"/>	2017-03-08 16:46:54	New Network Protocol	 10.131.160.13		34:88:5D:6D:14:46	Logitech Far East
<input type="checkbox"/>	2017-03-08 16:46:54	New Host	 10.131.160.13		34:88:5D:6D:14:46	Logitech Far East
<input type="checkbox"/>	2017-03-08 16:46:54	Additional MAC Detected for Host	 10.110.125.4			
<input type="checkbox"/>	2017-03-08 16:46:54	Additional MAC Detected for Host	 10.131.101.7			
<input type="checkbox"/>	2017-03-08 16:46:54	Additional MAC Detected for Host	 192.168.36.10			
<input type="checkbox"/>	2017-03-08 16:46:43	New Transport Protocol	 10.120.106.9		04:18:94:65:9D:58	Host Mobility AB
<input type="checkbox"/>	2017-03-08 16:46:43	New Network Protocol	 10.120.106.9		04:18:94:65:9D:58	Host Mobility AB
<input type="checkbox"/>	2017-03-08 16:46:43	New Host	 10.120.106.9		04:18:94:65:9D:58	Host Mobility AB

When you click on any of the items in this list, you're actually drilling down in a workflow. This is also restricting the view to the items you've selected. Ultimately, you'll arrive at a host profile.

Host Profile

The *host profile* contains the most detail about a system, and you can get to it from any of the analysis views.

The host profile will serve up detailed information about the IP address, hostname, indications of compromise, applications, services, attributes, and potential vulnerabilities that exist on the hosts.

Host Profile

Scan Host

Generate White List Profile

Domain Global \ Cisco_Backend \ Cisco_SOC
IP Addresses 10.120.10.254
NetBIOS Name
Device (Hops) vNGIPS.dcloud.cisco.com (0)
MAC Addresses (TTL) 00:15:F7:7D:89:F8 (Wintecronics Ltd.) (64)
00:55:44:33:22:11 (64)
28:6A:BA:17:9B:5C (Apple, Inc.) (64)
... [\(show all\)](#)
Host Type Host
Last Seen 2017-03-08 16:09:23
Current User
View [Context Explorer](#) | [Connection Events](#) | [Intrusion Events](#) | [File Events](#) | [Malware Events](#)

Indications of Compromise (1) ▾

Edit Rule States

Mark All Resolved

Category	Event Type	Description	First Seen	Last Seen	
Malware Detected	Threat Detected in File Transfer	The host has encountered malware	2017-03-07 20:36:51	2017-03-08 12:38:57	

Operating Systems (3) ▾

Edit Operating System

View Operating Systems

Vendor	Product	Version	Source
Linux	Linux	2.6	Firepower
CentOS	Linux	5.5	Firepower
IBM	AIX	5.x, 5L 5.x	Firepower

Applications (1) ▾

Application Protocol	Client	Version	Web Application	
HTTP	Firefox	2.0.0.17	Web Browsing	

Connection Events

Devices can also collect connection data, and you can check out the events by going to **Analysis>Connection Events**. When you do that, you'll get information about protocols, applications, bytes transferred, URLs, and more, as shown below.

Connection Events (Switch workflow)

[Connections with Application Details](#) > [Table View of Connection Events](#)

2017-03-08 10:07:58 - 2017-03-08 16:07:58

Expanding

No Search Constraints [\(Edit Search\)](#)

Jump to... ▼

<input type="checkbox"/>	First Packet	Last Packet	Action	Reason	Initiator IP	Initiator Country	Responder IP	Respo Country
	2017-03-08 16:07:57	2017-03-08 16:07:58	Interactive Block with Reset		10.131.1.72		149.81.1.9	USA
	2017-03-08 16:07:56	2017-03-08 16:07:56	Interactive Block with Reset		10.0.10.1		10.120.10.120	
	2017-03-08 16:07:54	2017-03-08 16:07:55	Block	Intrusion Block	10.0.10.197		10.120.10.12	
	2017-03-08 16:07:54	2017-03-08 16:07:54	Interactive Block		84.8.10.21	GBR	172.16.10.124	
	2017-03-08 16:07:54	2017-03-08 16:07:54	Interactive Block with Reset		10.131.1.168		149.81.1.54	USA
	2017-03-08 16:07:53	2017-03-08 16:07:55	Block	Intrusion Block	10.0.10.197		10.120.10.11	
	2017-03-08 16:07:53	2017-03-08 16:07:54	Block	Intrusion Block	10.0.10.197		10.120.10.11	

Be warned that the data will only be found here if the appropriate logging is enabled in the Access Control policy!

One of the very cool things you can do with this information is look at it in the context of different workflows. Choosing `Switch` Workflows next to Connection Events will get you the options shown here.

Here is a picture of a great example of the Traffic Over Time workflow. The graphs are interactive, and you can drill down on individual elements by clicking them.

Overview

Analysis

Policies

Devic

Context Explorer

Connections ▶ **Events**

Connection Events ×

Global

▶ [Table V](#)

Connection Events

Connections by Application

Connections by Initiator

Connections by Port

Connections by Responder

Connections over Time

Traffic by Application

Traffic by Initiator

Traffic by Port

Traffic by Responder

Traffic over Time

Unique Initiators by Responder

Unique Responders by Initiator

Fast Packet

[2017-03-08 :](#)

[2017-03-08 :](#)

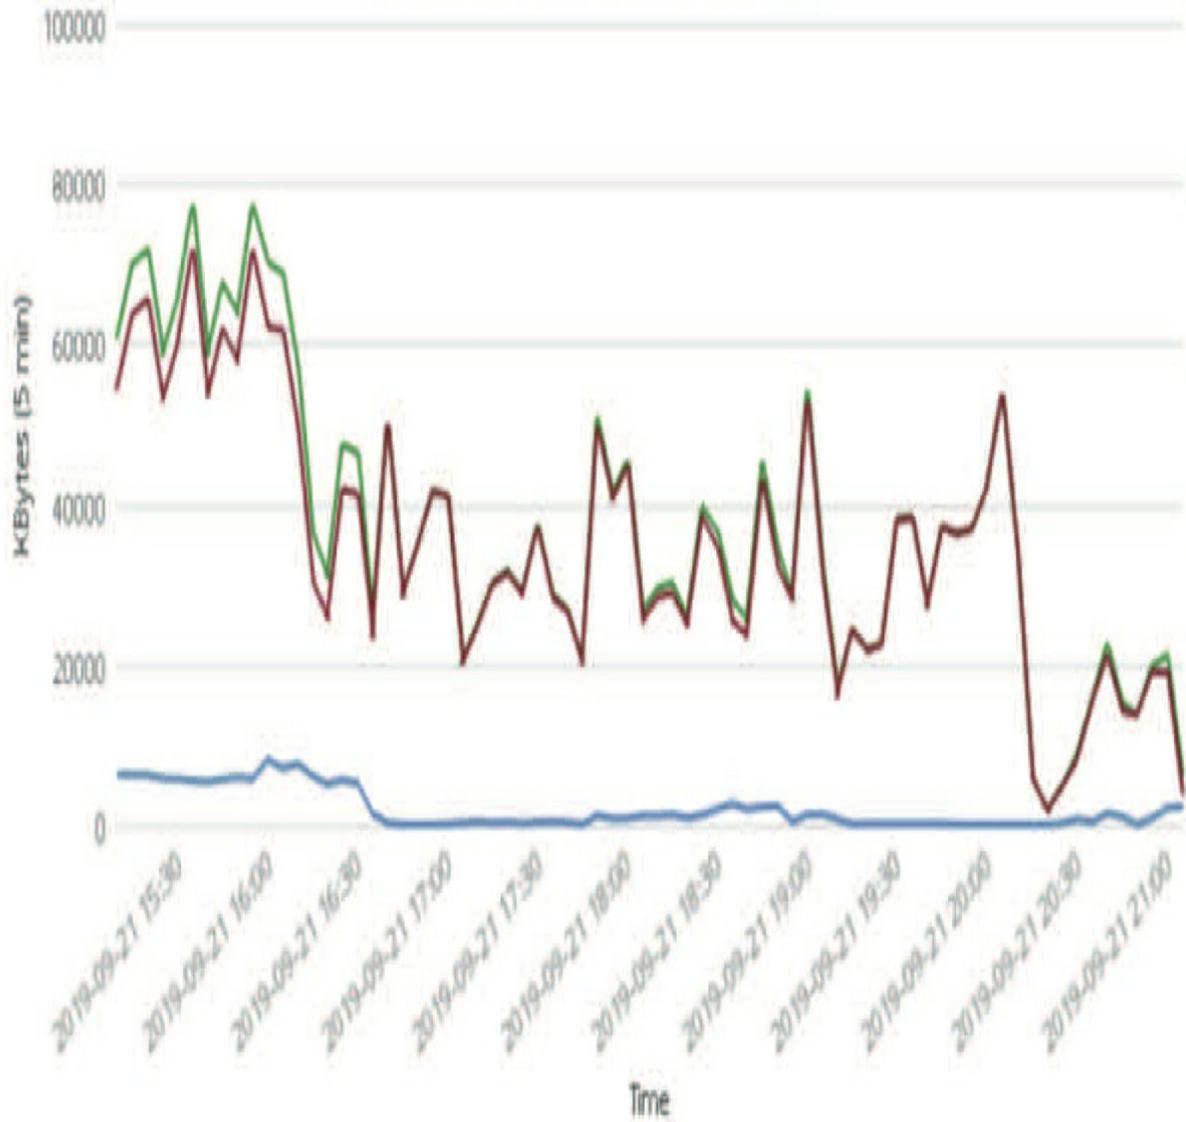
[2017-03-08 :](#)

[2017-03-08 :](#)

[2017-03-08 :](#)

[2017-03-08 :](#)

KBytes over Time (2019-09-21 15:05:00 - 2019-09-21 21:05:00)



Connection Summaries

Another way to view connection information is via a Dashboard. Choose **Overview>Dashboard>Connection Summary** to get the following tabs in the summary.

Connections

Displays information on the number of connections by initiator, application, port, and responder and over time.

Traffic

Refers to traffic bytes by initiator, application, port, and responder and over time.

Geolocation

Displays information about the source and destination countries and continents.

SSL

Provides tables and charts for the SSL activity on your monitored network segment

URL

Provides tables and charts for the URL activity on your monitored network segment

Each of these graphs is interactive, and you can use any or all of them to get even more detailed information. Nice, huh?



Overview Analysis Policies Devices Objects



Dashboards ▼ Reporting Summary ▼

Connection Summary

Provides tables and charts of the activity on your monitored network segment organized by different criteria



Connections x Traffic x Geolocation x SSL x URL x +

User Information

User information is collected either passively, using the protocols in the discovery policy, or via an Active Directory user agent.

The Active Directory agent communicates with specified AD servers to collect login information based upon the audit logs. An AD agent can be

installed on a Windows-based system inside your environment, but you've got to provide the name of an account with the capability to read the logs on the AD servers plus a password. And for some companies, I've found getting this access pretty challenging!

In addition to reading the logs, you can connect with an AD server via LDAP to read other attributes on accounts. Information including the first and last name, email, department, and phone number would be included provided the data is populated in the AD server. You can also use this LDAP connection to pull user/group account information for use in AC policies.

It's also good to know that connections can be made to non-AD LDAP servers and user data can be pulled for the other discovered user types as long as there's corresponding information in LDAP.

Firepower User Agent

To configure the FMC to communicate with the agent, go to **System>Integration>Identity Sources**. This User Agent is only available through 6.5 code. You need to use the Identity Service Engine (ISE) with pxGrid, as I'll show you in chapter 27.

Cloud Services Realms **Identity Sources** eStreamer Host Input Client Smart Software Satellite

Identity Sources

Service Type

Support for Cisco Firepower User Agent is deprecated and will be removed in a future release

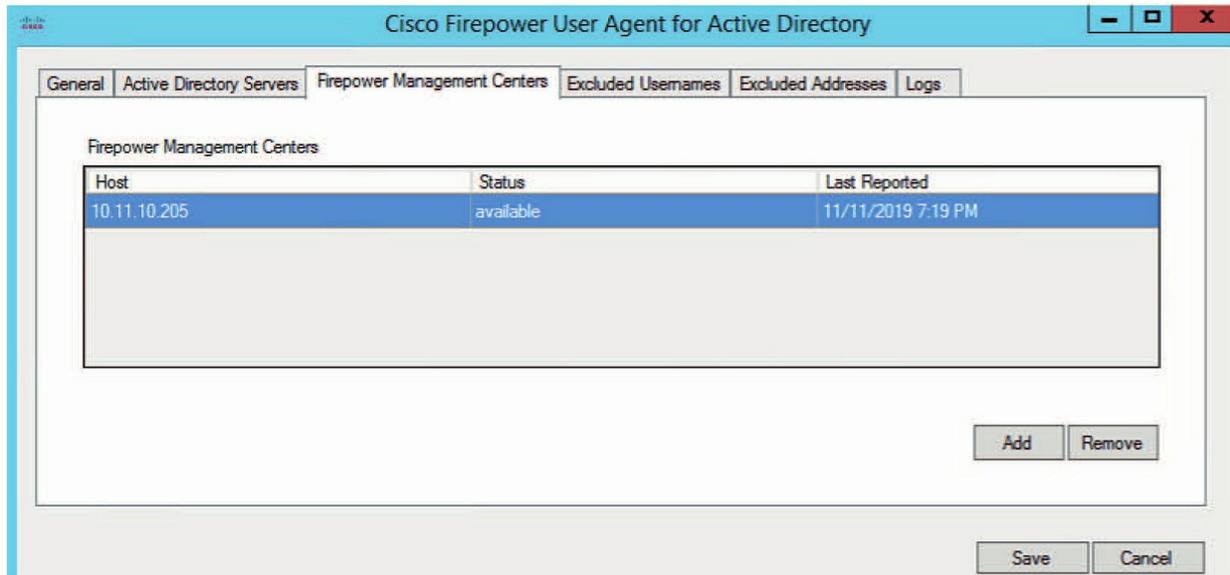
Host Name/IP Address

Click "Add" to add a new User Agent

Once here, just click **User Agent** and then **New Agent**. Enter a name for the agent plus the hostname or IP address of the FMC that you want the agent to perform monitoring on.

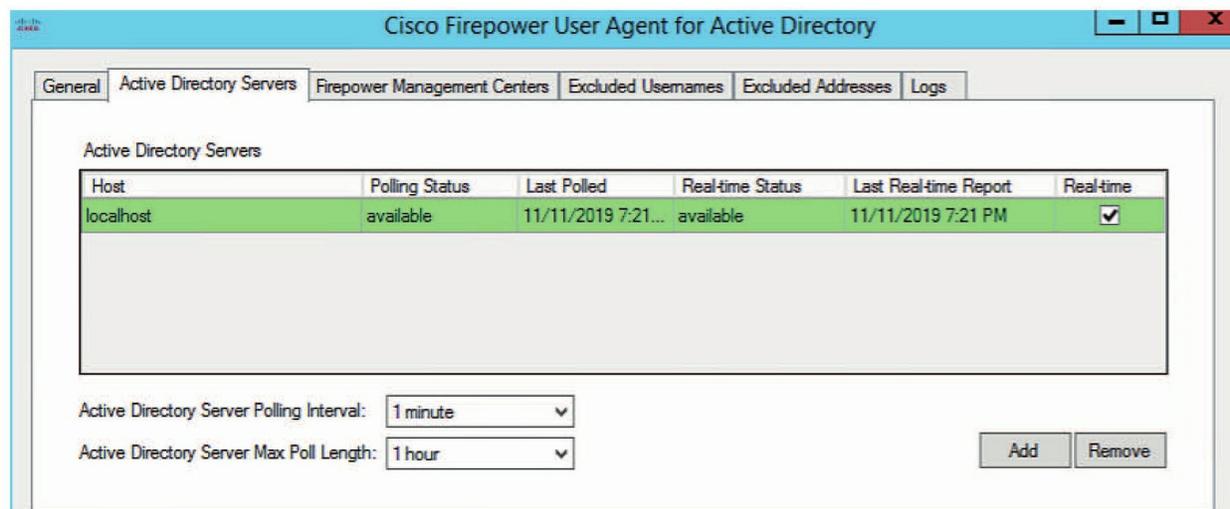
If you are using the User Agent, make sure all the key prerequisites are installed on the host on which you're installing the AD agent, including the

.NET Framework and SQL CE from Microsoft.



The install of the agent itself is actually really straightforward, but once it's installed, you've still got to configure it, which is also simple.

You can see the various settings by clicking through the tabs across the top of the dialog. You'll need to specify the agent name, AD servers, and FMC IP address. You can add up to five FMCs here.



Here you can see the dialog where you add the AD server. You can also exclude names and addresses from the discovery. If you run the supplicant on

your AD server and use the name localhost, the IP address will *not* work!

But you can actually run this on any host and then just point it at your AD server(s). One AD agent can communicate with up to five AD servers.

LDAP Configuration

You configure the LDAP connection on the same page that you added the User Agent, with the Realms tab.

The top half of the page contains the AD settings.

Just click Add New Realm, fill out the information, and then click Add.

Add New Realm

Name *	Lammle-AD
Description	
Type *	AD ▼
AD Primary Domain *	sfgtc.local
AD Join Username	
AD Join Password	
Directory Username *	administrator@sfgtc.local
Directory Password *	●●●●●●●●
Base DN *	dc=sfgtc,dc=local
Group DN *	dc=sfgtc,dc=local
Group Attribute	Member ▼

All of this is again pretty straightforward.

Just so you know, if you fill out the AD Join Username and AD Join

Password fields, a really useless test button shows up that doesn't tell you anything worthwhile, unless you're using Kerberos, which is what that test is for. Mostly just leave it blank.

User Analysis

Okay, so once user discovery has been enabled, the users table and user activity in the database will populate. You can see their contents by navigating to **Analysis>Users>Users**. This'll display all users that have been identified by the system.

[Table View of Users](#) > Users

No Search Constraints ([Edit Search](#))

Jump to...

User	Last Seen	Realm	Username	First Name	Last Name	E-Mail
aaron.con (D:\CLOUD-SOC\acon, LDAP)	2019-11-08 14:58:47	D\CLOUD-SOC	acon	aaron	con	acon@dcloud.cisco.com
aaron.kempf.f (D\CLOUD-SOC\lkemp, LDAP)	2019-11-08 15:17:00	D\CLOUD-SOC	lkemp	aaron	kempf	aaron.l.kempf@dcloud.cisco.com
aaron.smith (D\CLOUD-SOC\zsmit, LDAP)	2019-11-08 13:56:00	D\CLOUD-SOC	zsmit	aaron	smith	aaron.z.smith@dcloud.cisco.com
abbey.endres (D\CLOUD-SOC\kendr, LDAP)	2019-11-08 14:04:14	D\CLOUD-SOC	kendr	abbey	endres	abbey.k.endres@dcloud.cisco.com
abbey.esquivel (D\CLOUD-SOC\lesqu, LDAP)	2019-11-08 14:29:57	D\CLOUD-SOC	lesqu	abbey	esquivel	abbey.l.esquivel@dcloud.cisco.com
abbie.simms (D\CLOUD-SOC\gsimm, LDAP)	2019-11-08 15:27:49	D\CLOUD-SOC	gsimm	abbie	simms	abbie.g.simms@dcloud.cisco.com
abbie.wald (D\CLOUD-SOC\gwald, LDAP)	2019-11-08 14:15:42	D\CLOUD-SOC	gwald	abbie	wald	abbie.g.wald@dcloud.cisco.com
abby.kendrick (D\CLOUD-SOC\skend, LDAP)	2019-11-08 14:13:33	D\CLOUD-SOC	skend	abby	kendrick	abby.s.kendrick@dcloud.cisco.com
abby.ussery (D\CLOUD-SOC\kusse, LDAP)	2019-11-08 15:30:39	D\CLOUD-SOC	kusse	abby	ussery	abby.k.ussery@dcloud.cisco.com
abdul.cundiff (D\CLOUD-SOC\pcund, LDAP)	2019-11-08 14:24:22	D\CLOUD-SOC	pcund	abdul	cundiff	abdul.p.cundiff@dcloud.cisco.com

Next you'll see the User Activity view that's opened by going to **Analysis>Users>User Activity**. This view will reveal precisely when all those users were seen performing activities such as newly discovered users, and users logging in or logging out.

I want to point out that you can only see logouts if you are using the User Agent.

User Activity

[Table View of Events](#) > [Users](#)

No Search Constraints ([Edit Search](#))

	Time	Event	Username	Realm	Authentication Protocol	Authentication Type	IP Address
	2017-03-08 18:52:12	User Login	 marre	DCLOUD-AD 1	 LDAP	Passive Authentication	 10.112.196.113
	2017-03-08 18:52:06	User Login	 choar	DCLOUD-AD 1	 LDAP	Passive Authentication	 10.131.117.48
	2017-03-08 18:52:05	User Login	 ddevo	DCLOUD-AD 1	 LDAP	Passive Authentication	 172.16.10.208
	2017-03-08 18:52:04	User Login	 jsarm	DCLOUD-AD 1	 LDAP	Passive Authentication	 10.110.84.46
	2017-03-08 18:52:04	User Login	 khend	DCLOUD-AD 1	 LDAP	Passive Authentication	 10.0.91.111
	2017-03-08 18:52:01	User Login	 hmund	DCLOUD-AD 1	 LDAP	Passive Authentication	 10.0.83.146
	2017-03-08 18:52:00	User Login	 yloom	DCLOUD-AD 1	 LDAP	Passive Authentication	 10.131.10.104
	2017-03-08 18:52:00	User Login	 qdesm	DCLOUD-AD 1	 LDAP	Passive Authentication	 10.0.30.52

Host Attributes

There's one more data type available in the Firepower Manager: *host attributes* are elements that you can create on your own and either automatically or manually assign to systems.

Create Host Attribute

Name:

Type: List

List Values + Add Value

Name
<input type="text" value="Los Angeles"/>
<input type="text" value="Phoenix"/>
<input type="text" value="Chicago"/>
<input type="text" value="New York"/>

Auto-Assign Networks + Add Network

Value	IP Address	Netmask
<input type="text" value="Los Angeles"/>	<input type="text" value="192.168.1.0"/>	<input type="text" value="24"/>
<input type="text" value="Phoenix"/>	<input type="text" value="192.168.2.0"/>	<input type="text" value="24"/>
<input type="text" value="Chicago"/>	<input type="text" value="192.168.3.0"/>	<input type="text" value="24"/>
<input type="text" value="New York"/>	<input type="text" value="192.168.4.0"/>	<input type="text" value="24"/>

These attributes can then be used as sorting and search criteria to make it easier to locate systems or tag machines you want to flag as “special.”

Here's a list of the four types of attributes:

- **Text:** Creates a text box for a user to enter information
- **URL:** Creates a

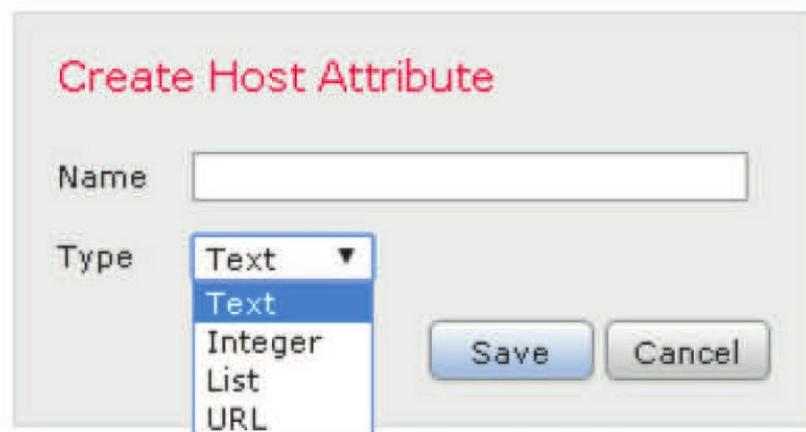
field where a URL can be entered

- **List:** Allows you to access a drop-down list with choices you create
- **Integer:** Creates a field for a numeric value to be assigned

To create attributes, go to **Analysis>Host Attributes** and choose **Host Attribute Management** in the upper-right corner. You'll see a default attribute called White List there,. On the upper right, there will be a button labeled Create Attribute.

Clicking the **Create Attribute** button will open a screen offering a menu of attributes you can create.

Creating the URL, text, and integer attributes will cause those characteristics to appear in the host profiles of each host in your network. The list attribute also allows for the option to automatically assign the attribute based upon a network/subnet combination.

A dialog box titled 'Create Host Attribute' with a red title. It contains a 'Name' text input field, a 'Type' dropdown menu, and 'Save' and 'Cancel' buttons. The dropdown menu is open, showing options: Text (selected), Integer, List, and URL.

Create Host Attribute

Name

Type

- Text
- Integer
- List
- URL

Save Cancel

Okay—so first, create the list attributes, and then specify which network the attributes belong to. If anything falls outside of those parameters, it would still show up in a category called Unassigned.

Once you've created your attributes, you can check up on them in several ways. You can use the host profile, but you can also check out the Network Map Host Attribute view.

There's even a link on the Create Attribute page. And if that's not enough, you can access them by going to Analysis>Hosts>Network Map and clicking the Host Attributes tab.

This gets you to a drop-down list where you can pick out any of the host attributes you or someone else created. The three attributes associated with whitelists are compliant, non-compliant, and not assigned—options that can't be assigned by an administrator. These characteristics can only be assigned based on the actual whitelist values, which we'll cover in detail in the chapter on correlation policies.

Summary

We covered Firepower Network Discovery, which was called FireSIGHT at one point, extensively in this chapter. It was originally called Realtime Network Awareness (RNA) and Realtime User Awareness (RUA). You learned this powerful technology equips us with the in-depth information needed to analyze data and fine-tune policies.

After you were briefed on this awesome, multifaceted tool, we explored discovery components like the policy, type of data collected, connection events, and the host attributes associated with it.

We wrapped up the chapter by demonstrating exactly how Firepower Network Discovery is used to greatly enhance event analysis.

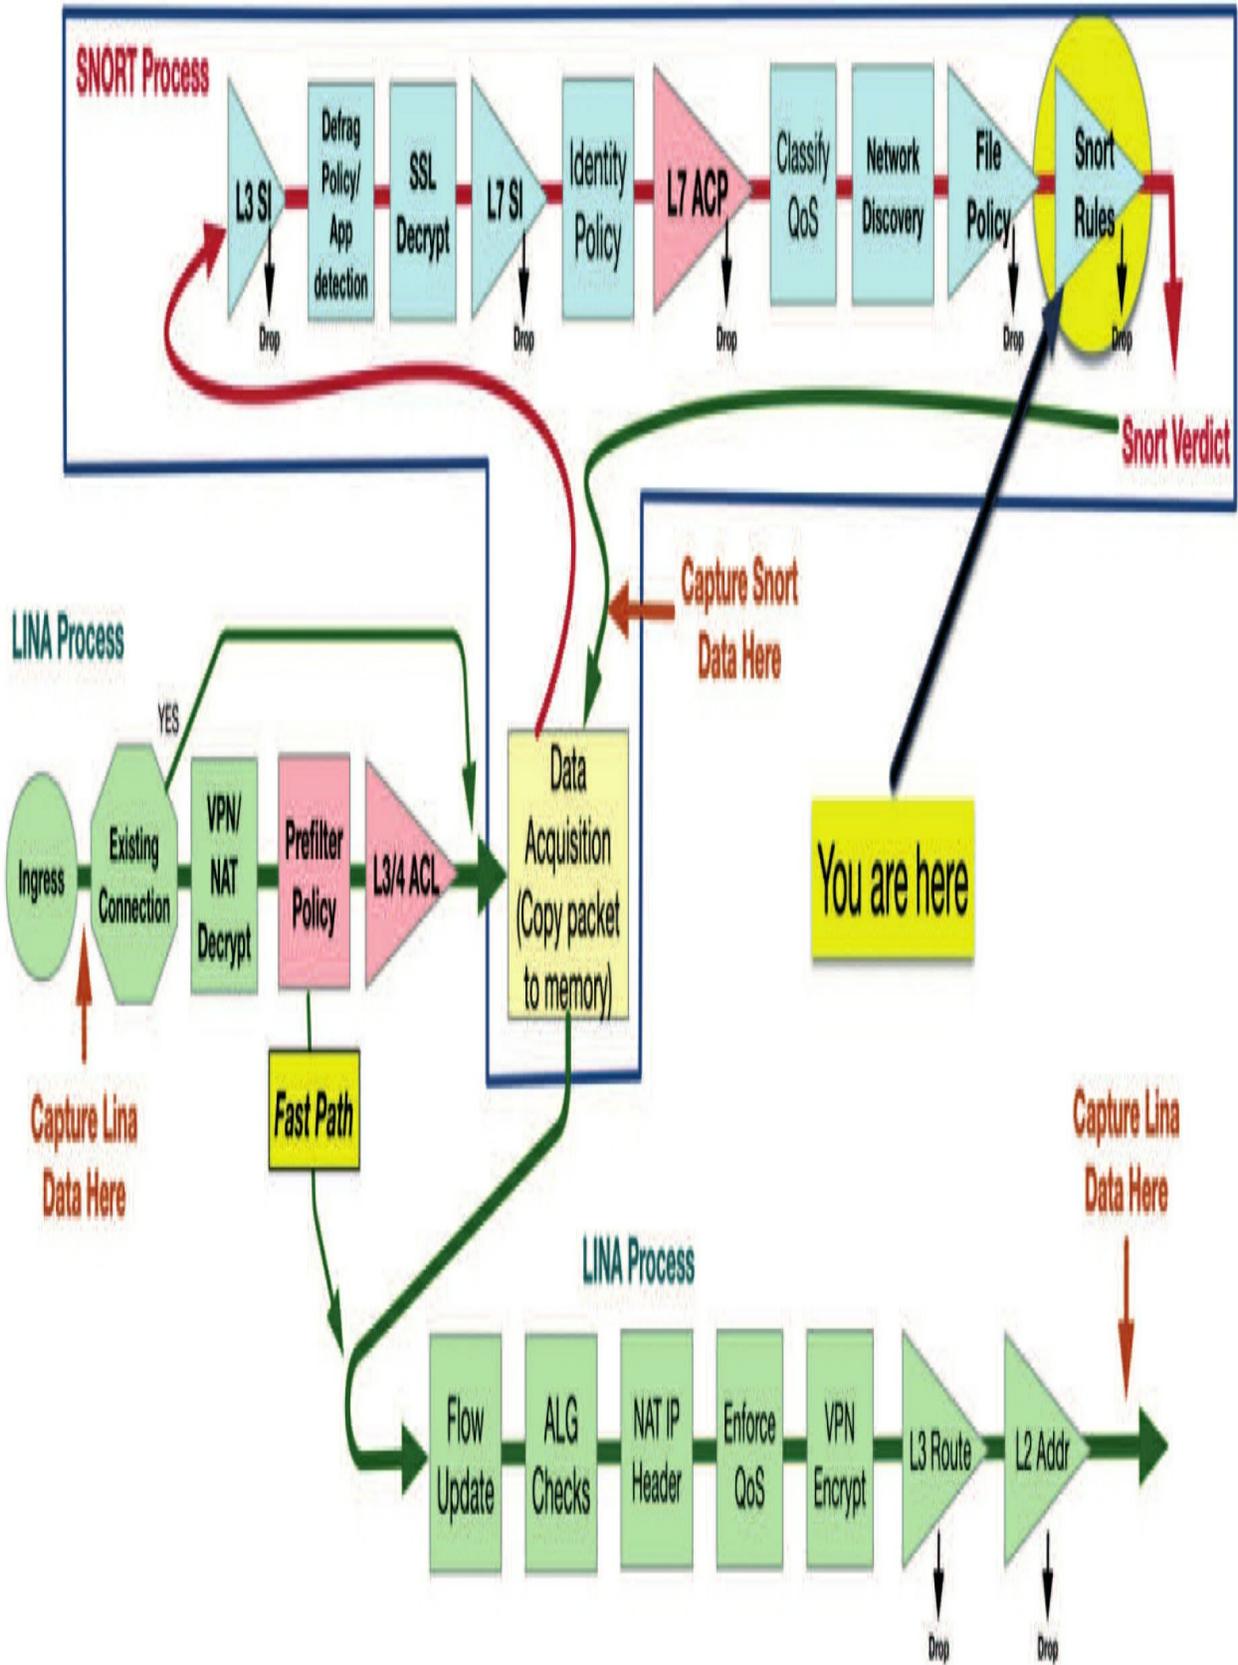
Chapter 12: Intrusion Prevention System (IPS) Policy

The following CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.2 Configure these policies in Cisco Firepower Management Center

2.2.b Intrusion



In this chapter, I'll turn the spotlight upon the Intrusion policy that's initially created on the FMC. The IPS policy is then configured into an ACP allow rule to be deployed to your devices, equipping them to come to a Snort verdict on the packets traversing your Firepower appliances and FTD devices.

You can definitely think of this policy as your Snort rule configuration because that's actually its primary purpose. Yes, it's true there's a sprinkle of advanced settings included, but for the most part, it's all about Snort rules plus how to implement and verify your policy.

We'll cover all the key concepts and I'll share some sound advice to guide you through creating and managing your policies. And of course, our journey wouldn't be complete without a nice demonstration on how to implement and verify a new IPS policy.

Just to let you know now, because we cover this in depth in this chapter; CVSS stands for Common Vulnerabilities Scoring System and does just that, vulnerability scoring. There's also the CVE, which stands for Common Vulnerabilities and Exposures and is kind of like a dictionary for the vulnerabilities. According to Cisco, CVE's stated goal is to "standardize the names for all publicly known vulnerabilities and security exposures".

Okay, let's get this awesome show on the road!

Policy Basics

The Intrusion policy's mission is to determine the set of Snort rules that'll be enabled to execute packet inspection. Organizations can create anywhere from one to hordes of Intrusion policies, allowing them to customize Snort rulesets for specific devices and/or portions of the network.

So maybe you're thinking... "Why not just enable all the rules? Wouldn't that be more secure?" Of course it would, but things are rarely that simple. Snort rules run the gamut from detecting an iTunes login attempt to triggering on evidence of a CryptoLocker infected workstation. Can you even imagine the sheer deluge of false positive event alerts coming at you if you

enabled all the rules? This alert flood would quickly overwhelm your intrusion analysts, making it nearly impossible to ferret out the alerts with actual security implications.

No worries if I got ahead of you at “false positive event alert.” We’ll get into the difference between a false positive and a true positive event in detail later, in Chapter 18, “Advanced Network Analysis,” but I think you get the idea.

At this writing there are close to 40,000 rules shipping with Firepower, and over 50,000 available Snort rules, and enabling all of them would epically overwhelm the device and completely crush your performance! If the device is inline, you’re definitely going to experience major latency (or will you? - we’ll discuss this!), and it’ll probably start dropping packets all over the place too, as it desperately tries to evaluate such a huge ruleset.

To address this tiny issue, Cisco’s Talos group created three basic base policies you can use as a starting point to tailor for your environment. We like to call them small, medium and large, but there’s also a fourth one called maximum detection. Here’s a description of those base policies now:

Connectivity over Security:

This is the small ruleset policy and only enables the low overhead rules. In this policy, connectivity is king and only a small subset of the total rules are enabled. The actual number fluctuates as rulesets are updated, but you can expect somewhere around 500 rules will be set to generate alerts or drop and generate alerts.

Balanced Security and Connectivity:

This is the medium ruleset and enable low and medium overhead rules. As you might guess, it balances the needs of connectivity and security. It’s hands down the most popular policy with Firepower customers but it won’t turn on High or Very High overhead rules by default, which can be a problem. The ruleset is designed to alert or block on the most important vulnerabilities or attacks while at the same time providing good performance. Good to know is that the throughput rating for Firepower devices is based upon the use of this ruleset. Expect roughly 11,000 rules in this set.

Security over Connectivity (Todd Recommended): Another popular

ruleset is the “shoot first, ask questions later” security ruleset and enables all the low, medium and high overhead rules. This will enable about five thousand more rules than the balanced ruleset, and these additional rules will cover threats that may be slightly older, but still needed. It also adds a few additional rule categories to the mix. You can expect to see around 16,000 enabled or so rules in this set. Bear in mind that this ruleset isn’t typically recommended by Cisco for inline devices due to the reduced throughput, but it’s still definitely worth considering in higher security environments. Cisco will yip that your Firepower device is *notguaranteed* to perform at its rated speed if you use this ruleset, but that’s just a bunch of...baloney. If you use Balanced, then you’ll probably call them less, period, which could be their real motive, but I’ll teach you how to make this SoC policy sing in your network while stopping more attacks than their Balanced endorsed policy does!

Maximum Detection:

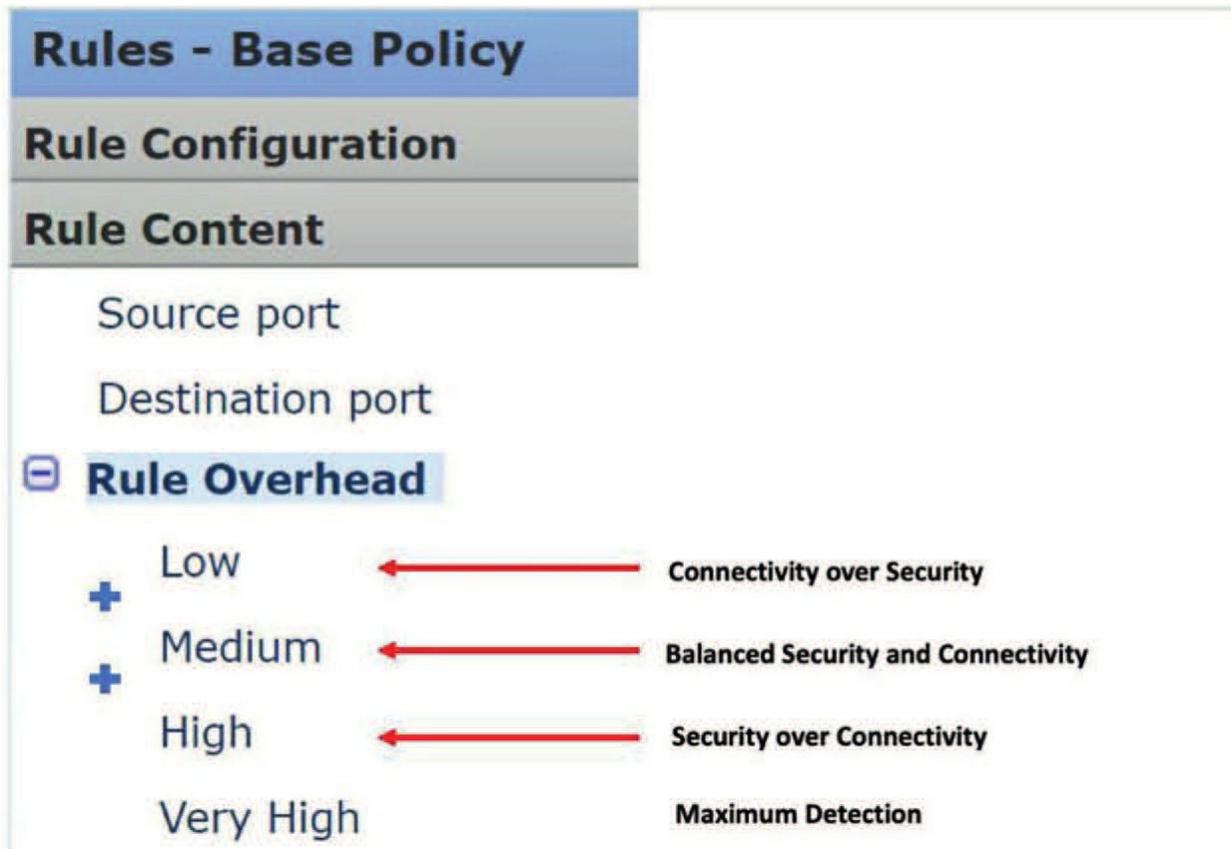
You can expect to see around 30,000 rules enabled or so in this set, and this can be a lot of fun as long as you’re not attempting this on a production network with Drop when Inline on until this can be thoroughly tuned. So what overhead rule set does this base policy use anyway? None. Max Detect includes enabled all rules above Sid:10000, which I think is awesome... Interestingly enough though, when you enable this base policy, this also disables the latency settings, meaning it inspects all traffic fully regardless of the latency it might introduce, This is why Max Detect is called “maximum”. Instead you’ll see pretty massive latency introduced by the Firepower device if it’s under any kind of load. By the way, what I *just* wrote here - you won’t find anywhere...I found this out when testing.

Talos has some objective criteria for selecting the enabled rules in each of the Cisco policies I just described, which I’ll go through in a minute, but first I want to talk about how to see which rules are enabled in each Cisco authored IPS policy by default.

What Rules Are Enabled by Default?

This is a good question that you probably don’t know the answer to yet, but I’m going to show you an easy way to find out.

If you go into your policy and open your rules (yes, I know you are saying “what are you talking about?!” ...please just hang in here), then get over to the Rule Accordion and finally click on Rule Content, you’ll see that Cisco breaks the rules into overhead categories as shown in the following figure:



So, although what I just showed you above is all-so-very true, Cisco does *not* like it when I discuss the enabled Snort rules in this manner, because it can be confusing. However, understand this: If you use the default of balanced base policy, the High and Very High overhead rules would *never* be enabled by default, and this is not good because, yes, you need at least some of them!

Okay, I know what you’re thinking: “I have the Cisco default policy enabled, and I have configured automatic rule updates, so I’m good and sleep well at night. Besides, I trust Cisco explicitly, and they have my best interest at heart, so they’d turn on the rules that are very important, even with high and very high overhead” ...NO! THEY WON’T! This is a common misconception. I have multiple blogs on Lamble.com that describe this issue in detail.

Now hold on, understand that the Rule Overhead shown above with values of Low, Medium, High, and Very High were only intended to provide customers with a way of limiting the rulesets enabled when using what is called Firepower Recommended Rules (RR) so RR wouldn't enable rules in more security-conscious base Intrusion policies. We'll do RR rules later in this chapter, and this is why we just went through the Firepower Network Discovery chapter before this one – we need that data!

The rulesets resulting from these base policy filters are determined programmatically by the FMC at the time the Snort Rule Update (SRU) is imported. They aren't pre-built by Talos.

Okay - so the idea of this Rule Overhead setting was initially designed to help you see what was enabled using Firepower RR, however, before I get 100 angry emails from Cisco not that I showed you this, the overhead is not the only determination of how a rule gets enabled in a base policy. However, this does not change the fact that they will *not* turn on important rules that are listed as high and very high overhead with Balanced,.

Now it appears this “feature,” which was originally intended to be applicable for Recommended Rules, was also included in the Rule Selection Editor in the Intrusion policy. This is a bit of a blunder because constraining the rulesets based upon Talos-provided base Intrusion policies is already provided via the Base Policy selector when the Intrusion policy is first created.

Moreover, “Rule Overhead” is just a bad way of describing this in the first place because it gives the impression Talos has some universal metric used to quantify the performance impact of a given rule. It doesn't.

The base Intrusion policies rules that are enabled by default (meaning their overhead) is determined by a combination of the vulnerability's CVSS score that the rule uses to protect hosts from being exploited, and also based upon how old or how prevalent threats are.

Connectivity over Security Base Policy

Quoting Cisco again, here are the criteria for selecting the rules in the

Connectivity over Security base policy:

“This policy is specifically designed to favor device performance over the security controls in the policy. It should allow a customer to deploy one of our devices with minimal false positives and the full rated performance of the box in most network deployments. In addition, this policy should detect the most common and prevalent threats our customers will experience.”

Here’s Todd’s quote about this base policy: “if you *never* want to be bothered with any IPS events - ever, then use this”. Here is his next awesome quote: “if you are thinking of using this base policy, just don’t bother creating a policy at all and go to happy hour instead”. Got it?

Criteria :

CVSS Score = 10

CVE year is current – 2 (So, for example, 2020, 2019, 2018)

Balanced Security and Connectivity Base Policy

Loosely quoting Cisco—Here are the criteria for the Balanced Security and Connectivity base policy.

This policy is the default policy that’s recommended for initial deployments, which attempts to balance security needs and performance characteristics. Customers should be able to start with this policy and get a good block rate with public evaluation tools plus achieve a relatively high-performance rate with evaluation and testing tools.

NOTE: This policy is for people who want to just turn Firepower on and leave it, which is really bad policy for any NGFW product. It’s way better to tune your NGFW no matter who the manufacturer is!

This is the default shipping state of the Snort Subscriber Ruleset for Open-Source Snort.

Criteria:

- CVSS Score ≥ 9
- CVE year is current – 2 (for example, 2020, 2019, 2018)
- MALWARE-CnC rules

- EXPLOIT-KIT rules
- SQL Injection rules
- Blacklist rules
- Also includes the rules in the **Connectivity over Security** policy.

Security over Connectivity Base Policy

Again, loosely quoting Cisco, here are the criteria for the Security over Connectivity policy.

This policy is designed for the Cisco customer base that's exceptionally concerned about organizational security (translation: people who care enough to look into network analysis on their IPS and actually tune something, instead of ignoring it like a bad toothache).

Customers deploy this policy in protected networks with possibly lower bandwidth requirements but much higher security requirements. (Cisco says this, but I don't agree with this statement at all.)

Additionally, customers care less about false positives and noisy signatures because they will tune them out. Application control and locked-down network usage are also concerning to customers deploying this policy. It should provide maximum protection and application control, but it shouldn't bring the network down.

Anyway, this is a great policy and I used it at all my customers. Just understand you just can't set it and forget it! It needs to be tuned, and once you do that well, your network will be awesome.

Criteria:

- CVSS Score ≥ 8
- CVE year is current -3 (for example, 2020, 2019, 2018, 2017) ▪

MALWARE-CNC rules

- EXPLOIT-KIT rules
- SQL Injection rules
- Blacklist rules
- App-detect rules

- Also includes the rules in the **Connectivity over Security** and **Balanced** policies.

Maximum Detection Base Policy

Loosely, per Cisco—This ruleset is meant to be used in testing environments, so it's not optimized for performance. False positives for many of the rules in this policy are tolerated and/or expected, and Firepower investigations will normally not be undertaken.

Criteria:

The coverage is required for in field testing and includes rules from the Security, Balanced, and Connectivity rulesets. Okay, but which ones? There are approximately 40,000 on every system!

The answer is that the Max Detect includes all active rules above Sid:10000, unless otherwise specified, and this includes high and very high overhead rules. Love this! Why? Because I love events that allow me to tune an IPS policy, and events happen in spades with this policy!

Final Thoughts on Base Policies

So the various criteria for rules come with some caveats.

First, Talos reserves the right to not follow these criteria for any rule. Threats and vulnerabilities are individual, and a long-running threat may mean the rule will be enabled for a longer or shorter time frame than specified in the policy.

Plus, not all the rules in the categories listed will automatically be enabled. The idea is that the particular category will be considered, and rules enabled if deemed necessary. These policies are constantly tuned, ensuring the best protection is provided while staying within the performance constraints of the Firepower devices.

Of course, there's still another option. You can start with *no rules active* and go through the entire rule set yourself to decide which rules to turn on based

on OS, Application, protocols, web browsers, etc. If you have eons of time and the expertise equivalent to the entire Cisco Talos group or get paid by the hour, then this might be an option for you. To be real, nobody has that much time and nobody is as intimately familiar their own Snort rules as Talos is.

So no... Starting from scratch really isn't a viable option. But just in case you get paid by the hour and have nothing but time to tune rules, there's the Talos provided policy called *No Rules Active* to allow you to do just that. The rest of us might want to tweak one of the provided rules sets I just talked about instead.

Rule States

Before we move onto Policy Layers, let's talk about the options available for enabling rules in a policy (this is different than creating or editing a rule).

There are three states a Snort rule can have in Firepower:

- **Disabled:** The rule will not be deployed to devices.
- **Generate Events:** The rule will generate an alert if it encounters a packet matching the rule criteria.
- **Drop and Generate Events:** The rule will generate an alert and drop a packet matching the rule criteria. By default, this is a "silent" drop, meaning no TCP reset or UDP port unreachable response is sent to the source or destination.

You'll probably hear these three states described as disabled, alert, and drop, which comes from the fact that the actual keyword in the rule itself for a generate events state is *alert*. Likewise, the keyword for a Drop and Generate rule state is *drop*

Layers

To understand how you can use a ruleset provided by Cisco and still be able to modify it yourself, you've got to understand the concept of layers. In Firepower, an Intrusion policy consists of at least two layers, and you can add more layers if you want. Keep these principles in this in mind as I explain this concept:

- Settings in higher layers always override settings in lower layers.
- Layers can be private or shared across multiple policies.
- Layers can be merged together.
- The base layer in a policy is read-only and can't be edited.
- The base layer in a policy is actually another policy.

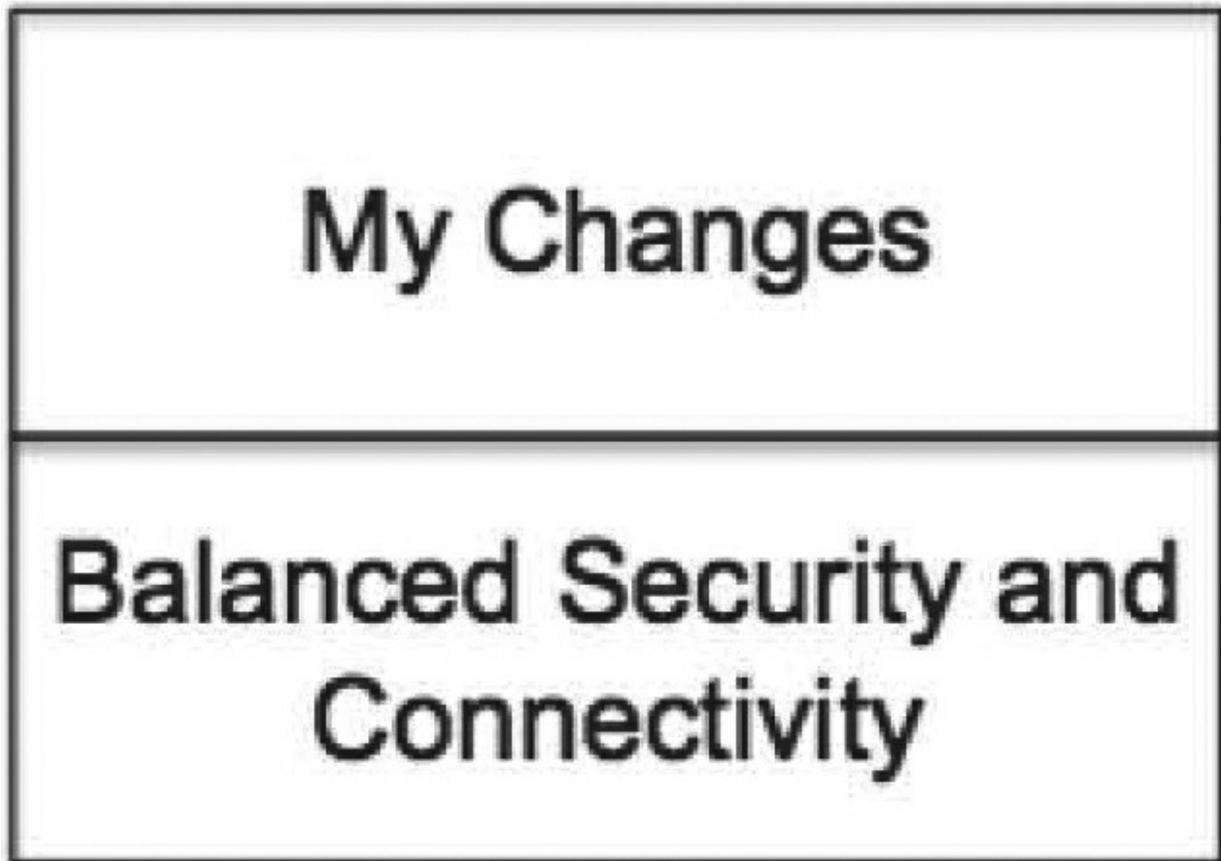
Okay—with that, let's make some sense of this now. First, when you create a new policy, it will *always* be based on some other policy. This is your base layer. That other policy can be one you've previously created, or it can be one of the Talos base policies.

If you use a custom policy as your base, that custom policy also has a base layer. The bottom line is if you trace any Intrusion policy back to its roots, it originally started with a Talos policy somewhere.

In addition to the base layer, each policy comes with a single private layer called My Changes. The My Changes layer sits atop the base policy or base layer as illustrated here:

In this example, I chose the Cisco Balanced Security and Connectivity policy as our base layer, which means we're starting with the balanced rule set. Initially, the My Changes layer makes no modifications to the Cisco policy. In any policy, the resulting ruleset will be the combination of settings

Intrusion Policy



in the layers. Since My Changes has no settings yet, the resultant ruleset will have only the Cisco balanced rules enabled.

Keeping in mind Principle #1 from the list above, if we make any change to a rule state in the My Changes layer, it'll override the rule state from the layer, or layers, below. This is how you start with a Cisco-provided policy and modify rule states to customize it. Basically, once you make your own decision on a rule state, you've overridden Cisco's default state for that rule.

From that point forward, no rule state changes for that rule will be inherited from the base layer but understand this will not prevent rule structure

changes.

So, if Talos decides to modify the rule to improve detection or maybe add documentation, those updates would still take effect. *Only the rule state or alerting capability is affected in the My Changes layer.*

Private and Shared Layers

Buckle up because things get a little bumpy from here! Using a single private layer as I did in the preceding example is pretty straightforward, but we can do more. In addition to using private layers, you can add shared layers, which can be added to a policy and then shared across other policies.

By using shared layers, you can make a single change to a shared layer and change rule states across multiple policies. This eliminates the need to visit multiple policies to make the same change over and over.

The following figure offers an example of three policies sharing a layer:

Intrusion Policy A

Intrusion Policy B

Intrusion Policy C

My Changes	My Changes	My Changes
Shared Layer	Shared Layer	Shared Layer
Balanced Security and Connectivity	Balanced Security and Connectivity	Balanced Security and Connectivity

Here you can see that when a change is made to the shared layer, it affects all

three policies. At the same time, if you need to make a change to just one of the policies, you still have that option. To do that, just make the change to the specific policy's My Changes layer.

I used a pretty generic name for our shared layer in this example, but you can add multiple shared layers with descriptive names across a number of policies as part of your design.

Before we go much further, I want to go over some of the considerations for using shared layers. First, even though the layer is shared across multiple policies, it can only be changed in the policy where it was first created. Back to the example, if the layer was created in policy A, it's *read-only* in the other policies. This means you have to remember the policy where you actually created the shared layer if you want to modify it. One way around that is to create an Intrusion policy that does nothing but house your shared layers. Name it something like "Shared Layers" and now you know where to go to edit them!

Another caveat is that the FMC doesn't have a good way of displaying the fact that a policy has a shared layer, or which other policies are sharing layers. The policy list view is decidedly flat and contains only the policy name and description.

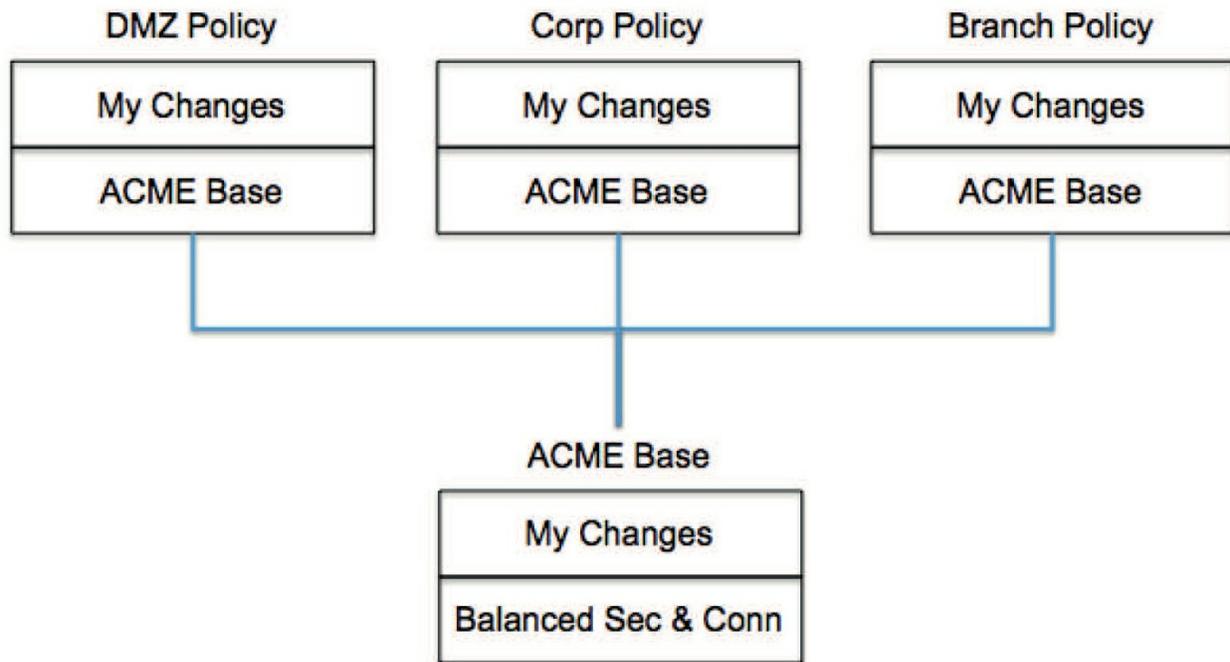
Because of this, your external documentation is key to keeping you and the other Firepower administrators sane by helping everyone remember where the shared layers are, and which policies will be affected if you change one! If you edit a policy, you can see the layers more clearly, but continuously editing policies to understand what layers are present eats up a lot of time.

A Shared Layer Alternative

So, let's now explore a really cool way to get similar results in a more user-friendly way. Instead of using shared layers to house common rule settings, you can create your own base policies, and doing so gains a couple of advantages over using shared layers.

First and foremost, it's easier to see within the Firepower user interface. Even though layers aren't visible from the policy list, a separate base policy does

show up as a distinct entry. Along with this, it's easy to tell which policy you need to edit to affect the other policies. You also don't need to worry about making the change in the right layer because each policy has only one editable layer by default—the private layer named My Changes. This technique is illustrated in the following figure:



As you can see, there are four Intrusion policies here. The ACME Base is used as the base layer in the top three policies. By doing this, we can start with the Cisco Balanced Security and Connectivity ruleset. Then you just make any changes you want to take effect throughout your organization to the ACME Base policy and these settings will then be inherited by the policies above. If you want to make a change to only the DMZ policy, you edit that policy directly. Any change made to the top three policies will override the setting inherited from below.

I want to point out that in this design, the ACME Base policy will never be deployed to a device—it exists only as a base for the other policies. I'd mention this in the policy description just so everyone is clear. Something like "This policy is never deployed by itself!" would be good.

By using this policy hierarchy, you get the best of all worlds:

- You leverage the expertise of the Talos team in determining which rules are

most effective.

- You have the option to override any of the Talos default rule states across your organization by editing the ACME Base policy just once.
- You have the option to customize your specific policies if needed.

Of course, this can also be expanded with additional base policies if you want. If you have a whole bunch of Intrusion policies deployed, creating an overall base and another set of intermediate base policies probably won't work for you. And keep in mind, the saying "Complexity is the enemy of security" is in full force here. Making your policy structure too complex can lead to unintended rule states, horrible confusion among your Firepower administrators, cats and dogs living together—mass hysteria!

The Intrusion Policy Interface

Okay—we've covered some basics of the Intrusion policy and some options for helping you manage rules across multiple policies. So next, we'll dive into the mechanics of what this looks like in the user interface.

Intrusion Policy Editing

To get things rolling, find the Intrusion policy under **Policies>Access Control>Intrusion**.

Create Intrusion Policy



Policy Information

Name *

vFMC_Lammle_IPS

Description

Drop when Inline



Base Policy

Balanced Security and Connectivity ▼

--System-Provided Policies--

Balanced Security and Connectivity

Connectivity Over Security

Maximum Detection

No Rules Active

Security Over Connectivity

* Required

and Edit Policy

Cancel

To create a new policy, click the Create Policy button.

In the dialog you'll see the following options:

- Name: The display name for your policy. Make it descriptive.
- Description: Optional description. This will appear in the policy list.
- Drop when Inline: Checking this box allows "Drop and Generate" rules to actually drop packets. Unchecking it means these rules will only generate alerts. This comes in handy for initial deployments before you've tuned your rules to ensure you don't block legitimate traffic.
- Base Policy: The base policy or base layer for this policy. This must be another policy, and it can be a system-provided (Talos) policy or a user-created policy.

The Create Policy button simply creates the policy and returns you to the policy list, which is an easy way to create and implement a new policy.

Using the Create and Edit Policy button creates the policy and takes you into the policy management interface:

Policy Information

Rules

Firepower Recommendations

 Advanced Settings

 Policy Layers

 My Changes

 Balanced Security and Connectivity

Policy Information

Name

vFMC_Lammie_IPS

Description

Drop when Inline

 **Base Policy** Balanced Security and Connectivity ▼

 The base policy is up to date (Rule Update 2020-01-13-001-vrt)

 **This policy has 11302 enabled rules**

 98 rules generate events

 11204 rules drop and generate events

[No recommendations have been generated. Click here to set up Firepower recommendations.](#)

Commit Changes

Discard Changes

On the left is the navigation pane that includes the various policy sections:

- **Policy Information:** Displays a summary screen that lists the basic settings of the policy. This one is important because if you make any changes, you've got to return here to either commit the changes (save the policy) or discard them. Otherwise you or other admins will be unable to edit other policies. Understand that choosing Commit Changes here only saves the changes in the FMC database—it doesn't actually push the settings to the managed device!
- **Rules:** Displays all the rules available to the policy.
- **Firepower Recommendations:** Allows Firepower to recommend which rules to enable or disable based on host data collected from your network.
- **Advanced Settings:** Contains logging, global rule threshold, and sensitive data settings. Really, not much here.
- **Policy Layers:** Allows you to expand and view the policy layers, including the base layer and any user-defined layers.

Rule Management

Clicking the Rules link brings you to the rule management interface. This is where you will spend most of your time when editing an Intrusion policy:

Edit Policy: Demo Policy

Policy Information ⚠

Rules

Firepower Recommendations

⊕ Advanced Settings

⊕ Policy Layers

Rules < Back

Rule Configuration

Filter:

Rule Content

Category

- app-detect
- blacklist
- browser-chrome
- browser-firefox
- browser-ie
- browser-other
- browser-plugins
- browser-webkit
- content-replace
- decoder
- exploit-kit
- file-executable

Classifications

Microsoft Vulnerabilities

Microsoft Worms

Platform Specific

Preprocessors

Priority

Rule Update

➔ ⏏ ⚠ ! 💬
Policy ▼

Rule State Event Filtering Dynamic State Alerting Comments

	GID	SID	Message		
<input type="checkbox"/>	1	37062	APP-DETECT 12P DNS request attempt	➔	
<input type="checkbox"/>	1	28071	APP-DETECT 360.cn SafeGuard local HTTP management console access attempt	➔	
<input type="checkbox"/>	1	28068	APP-DETECT 360.cn Safeguard runtime outbound communication	➔	
<input type="checkbox"/>	1	32845	APP-DETECT Absolute Software Computrace outbound connection - 209.53.113.223	➔	
<input type="checkbox"/>	1	32846	APP-DETECT Absolute Software Computrace outbound connection - absolute.com	➔	
<input type="checkbox"/>	1	32847	APP-DETECT Absolute Software Computrace outbound connection - bh.namequery.com	➔	
<input type="checkbox"/>	1	32848	APP-DETECT Absolute Software Computrace outbound connection - namequery.nettrace.co.za	➔	
<input type="checkbox"/>	1	26286	APP-DETECT Absolute Software Computrace outbound connection - search.dnssearch.org	➔	
<input type="checkbox"/>	1	26287	APP-DETECT Absolute Software Computrace outbound connection - search.namequery.com	➔	
<input type="checkbox"/>	1	32849	APP-DETECT Absolute Software Computrace outbound connection - search.us.namequerv.com	➔	

< of 611 >

Let's take a look at each section:

- **Filter panel (also called the Rule Accordion):** This is the vertical column that lets you view rules associated with specific categories, classifications, priorities, or rule updates. Selecting an option from the filter panel will populate the filter bar at the top of the rules list. Clicking the sections will expand or contract other sections.
- **Filter bar:** Up there at the top of the interface, use this to manually input search criteria for rules. You also get the option to populate the filter bar by selecting items from the rules filter panel.
- **Rules list:** This makes up the main part of the display. It shows you all the rules based upon the criteria specified in the filter bar.

Policy drop-down: Located on the right, below the filter bar and above the rules list. This allows you to examine a specific layer of your Intrusion policy. The default view, Policy, shows the cumulative values of all layers.

- **Show Details:** If you have a specific rule selected in the rules list, you'll see a Show Details button appear at the bottom of the interface. Clicking it will bring up the details of that rule, including the text of the rule itself, documentation, overhead, and reference websites.
- **Button bar:** Directly under the filter bar you'll find a set of buttons for settings that can be applied to rules. Settings can be applied to single or multiple rules by selecting their associated check boxes in the rules list.

The Rule Button Bar

The button bar is used to make changes to rule settings within the policy. It's important to remember that whatever you select in this bar only applies to rules with their check boxes selected. If you attempt to make a selection from the button bar without first checking at least one rule, you'll get an error. Here is a close-up of the rule button bar:



Let's take a look at what each of these does:

Rule State

The first menu drop-down is Rule State, which lets you specify one of three conditions: Generate Events (alert only), Drop and Generate Events (alert and drop matching packets), or Disable (turn the rule off). Remember that if your device is inline, you gain the ability to drop traffic if the rules are set to Drop and Generate and the Drop when Inline check box is selected in the policy.

Event Filtering

The Event Filtering settings allow you to specify whether you want to decrease the frequency of the alerts generated by the rule. When you click this button, you have the option to adjust Suppression or Threshold settings.

Suppression

Suppression settings are straightforward and used when you don't want alerts generated for a specific rule. This can be based on the rule, source, or destination IP address. The Suppression dialog is shown below.

“warning”

WARNING: Suppression in a rule only suppresses the alert generated by the rule; it has no effect on the ability of the rule to drop packets. That bears repeating, so let me put it another way: If you add a suppression to a Drop and Generate Events rule, the rule will *continue to drop traffic and you won't be notified*. Do Not. Ever. Suppress. A. Rule! I don't think this should even be an option. If you just said to yourself “why don't you

use a Pass rule instead?”...you’re close to being a Cisco CCNP Security.

Because of this, you shouldn’t use the suppression option for a rule that is set to drop because you’ll end up with a rule that stops generating alerts for the suppressed source or destination IP address but still continues to drop packets—nightmare.

See this:

Just say NO!



One more thing about suppression... You have the option to suppress the rule without adding a source or destination IP address. Think long and hard about why you would want to do this. Normally, if you don't want a rule to generate an alert, you would just disable the rule. Suppression only suppresses the rule's output, but the rule continues to process traffic. Why leave the rule in the ruleset, using valuable computing resources, if it's never going to generate an alert?

Thresholding

There are three types of thresholds to choose from: Limit, Threshold, and Both. They have a common dialog box, shown below, and each includes Track By, Count, and Seconds parameters. Each one is tracked based on a source or destination IP address, but not both.

Limit:

Allows you to limit the number of alerts generated within a selected number of seconds. Once that limit is reached, no new alerts will occur until the time period has expired. For example, by setting Count to 10 and Seconds to 300, you'll receive no more than 10 alerts every 5 minutes for the selected Source/Destination IP address.



The image shows a dialog box titled "Set Threshold for 1 rule". It contains the following fields and controls:

- Type:** A dropdown menu with "Limit" selected.
- Track By:** A dropdown menu with "Source" selected.
- Count:** An empty text input field.
- Seconds:** An empty text input field.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

Threshold:

This will set the number of times the rule must match before generating an alert within a given time window. For example, if the count was set to 10 and the seconds set to 600, the result would be an alert for every 10 occurrences of that event within 600 seconds.

Both:

This is a combination of the Limit and Threshold parameters. The count specifies the number of times the rule must match before an alert is generated. Once that count is reached, the rule will generate one alert; another alert won't be generated within the Seconds window. So a Both threshold with Count set to 10 and Seconds set to 300 means trigger an alert after seeing 10 matches within 300 seconds, and after triggering the alert, don't trigger again for 300 seconds. If the rule continues to match and exceed the threshold, you'll get a maximum of one event every 300 seconds.

Dynamic State

Dynamic State allows the device to dynamically change the state of the rule from its current one to any of the other available states: Drop and Generate Events, Generate Events, or Disabled. The state changes based upon the source or destination address as defined in the Track By drop-down. Another option that can appear in the Track By drop-down is Rule. If you specify Source or Destination, then the Network field must be populated. If Rule is selected in the Track By field, the Network field won't appear. The Rate field lets you specify that the rule must fire a certain count within a given number of seconds. There's also a Timeout field that resets the rule to the previous state.

The Dynamic State dialog is shown below:



The screenshot shows a dialog box titled "Add Rate-Based Rule State for 1 rule" with a question mark and close button in the top right corner. The dialog contains the following fields and controls:

- Track By:** A dropdown menu currently set to "Source".
- Network:** An empty text input field.
- Rate:** A text input field followed by the text "Count /" and another empty text input field, followed by the text "Seconds".
- New State:** A dropdown menu currently set to "Drop and Generate Events".
- Timeout:** An empty text input field.

At the bottom right of the dialog, there are two buttons: "OK" and "Cancel".

Alerting

To enable or disable SNMP alerts for a specific rule, use the Alerting option.

You configure the SNMPdestination in Advanced settings. SNMP alerts are generated by the device.

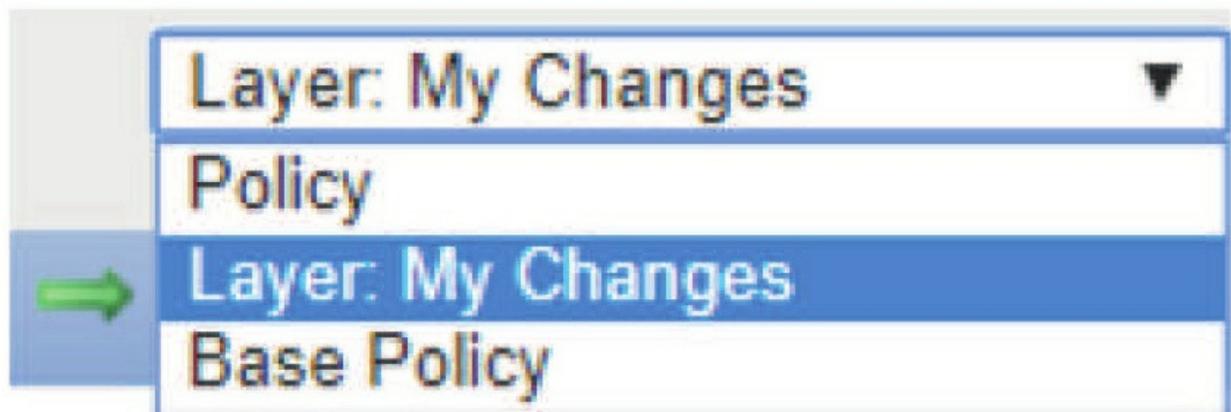
Comments

You can add comments to the individual rules that can be viewed by the FMC's users and analysts who inspect the traffic. Adding a comment automatically stamps the entry with your username and the date/time.

The Policy Drop-Down

The Policy drop-down on the far right of the button bar gives you a view of the rules based upon states set within the individual policy layers.

The default view is just Policy, but you can also check out any policy layers like My Changes, Firepower Recommendations (if used), and the base policy. Colors are important indicators here:



- All rules will appear white by default.
- When you look at individual layers, pink indicates that the rule's state has been modified in a higher layer.
- A rule highlighted in yellow means its state was adjusted in a lower layer.
- Rules highlighted in orange are ones you've clicked on.

Creating, Importing, Deleting, & Editing a Snort Rule

Remember that when creating or editing a rule, you'll receive other options as discussed in chapter 8: Objects.

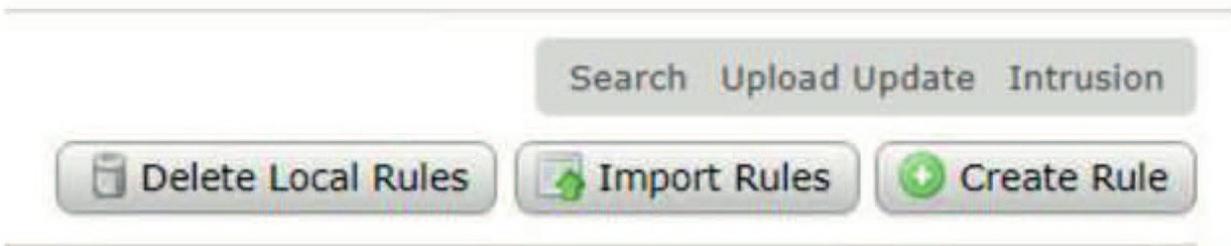
- If a rule is an **alert** rule, it generates an intrusion event.
- If it is a **pass** rule, it ignores defined source or destination traffic. You can create pass rules in order to prevent packets defined in the pass rule from triggering the alert rule in specific situations, this can be better than disabling the alert rule. By default, pass rules override alert rules.
- For a **drop** rule in an inline deployment, the system drops the packet and generates an event.

I am bringing up the options here because it is important to understand the difference between *alert* and *pass*, so let's go through this again. In chapter 8: Objects, I started discussing the alert and pass rules at the end of the chapter. It's imperative you understand these topics and how to create, import, delete and edit a rule.

Let's start by going to **Objects>Intrusion Rules**.



To import new snort rules, delete or create a new rule, the righthand menu is where you need to start. The very top right, however, is quick links to other pages.



If you want to edit an existing snort rule, you'd start on the lefthand side of the page.

For you can search for rules by going to search bar on the left and just typing in a SID number, which is what I usually do, or search by group as shown, or on the right of that you can search by categories, lots of categories!



Group Rules By

Category



Category

Local Rules

Microsoft Vulnerabilities

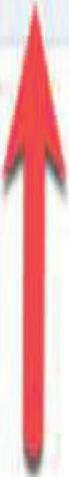
Microsoft Worms

Platform Specific

Priority

SANS Top 20 (version 5.0)

SANS Top 20 (version 6.01)



Category (52811)

▶ app-detect (165)

▶ browser-chrome (77)

▶ browser-firefox (296)

▶ browser-ie (2639)

▶ browser-other (97)

▶ browser-plugins (2552)

▶ browser-webkit (130)

▶ content-replace (23)

▶ decoder (153)

▶ deleted (12242)

▶ exploit-kit (758)



Notice on the top of “Category” shows how many rules are in my Firepower system- 52,811! You can even search for deleted rules here and reenable them.

I am going to just type in 53093 and open that existing enabled rule. So, this is an active rule that is looking for MULTIMEDIA TRUFFLEHUNTER attack attempts, and this is a rather new rule so we take it serious, however, R&D has been doing their own testing with these attack attempts and we’re getting false positive form them.

Business requirement: We want to stop getting any alerts from R&D but leave the rule alone to drop offending packets from any other zone. Here is the current rule that is:

```
alert tcp $EXTERNAL_NET $FILE_DATA_PORTS -> $HOME_NET any
```

Rule state: **Alert**

Protocol: **TCP**

Source Net: **External_Net zone**

Source Port: **\$File_Data_Ports**

Dest Net: **\$Home_Net**

Dest Port: **Any**

Here is the Snort rule in full, which may be harder to read at first:

```
alert tcp $EXTERNAL_NET $FILE_DATA_PORTS -> $HOME_NET any (msg:"FILE-  
MULTIMEDIA TRUFFLEHUNTER TALOS-2020-1012 attack attempt"; sid:53093; gid:3; rev:1;  
classtype:attempteduser; flowbits:isset,file.quicktime; reference:url,www.  
talosintelligence.com/reports/TALOS-2020-1012/; metadata:engine shared, soid 3|53093, service ftp-  
data, service http, service imap, service pop3, policy max-detect-ips drop, policy securityips drop; )
```

Here is the rule shown in the editor output. Notice how the same rule is formatted here:

Edit Rule 3:53093:1

[\(View Documentation\)](#), [Rule Comment](#)

Message	<input type="text" value="FILE-MULTIMEDIA TRUFFLEHUNTER TALOS-2020-1012 attack attempt"/>		
Classification	<input type="text" value="Attempted User Privilege Gain"/> ▼		
	Edit Classifications		
Action	<input type="text" value="alert"/> ▼		
Protocol	<input type="text" value="tcp"/> ▼		
Direction	<input type="text" value="Directional"/> ▼		
Source IPs	<input type="text" value="\$EXTERNAL_NET"/> ⓘ	Source Port	<input type="text" value="\$FILE_DATA_PORT"/>
Destination IPs	<input type="text" value="\$HOME_NET"/>	Destination Port	<input type="text" value="any"/>

Detection Options

flowbits

Operator	<input type="text" value="isset"/> ▼
State	<input type="text" value="file.quicktime"/>
Group	<input type="text"/>

reference

metadata

▼

Going into my IPS policy and from the GUI we can see it's a drop rule, with High overhead, meaning it's enabled with the SoC policy and the Max policy.

In the metadata of the rule it shows this as well, including that this is a drop rule by default

```
metadata:engine shared, soid 3|53093, service ftp-data, service http, service imap, service pop3, policy max-detect-ips drop, policy security-ips drop; )
```

We cannot change the Cisco shared rule, only edit them by renaming them, so I am going to change this rule and then save it.

Let's see what happens, but remember before I do that, I do not want this rule triggered when packets from R&D (172.16.10.0/24) to any destination or port. Here is my rule before save.

Edit Rule 3:53093:1

[\(View Documentation, Rule Comment\)](#)

Message

Classification

[Edit Classifications](#)

Action

Protocol

Direction

Source IPs

Source Port

Destination IPs

Destination Port

Detection Options

Nice, it saved successfully as shown below.

So, I didn't put an arrow towards anything to point anything out, so what is the one difference in this rule before save and after save? Do you see it?

Only the SID rule number changed at this point. User created rules start at 1,000,000, and even though we were only editing it can only be saved as a new rule.

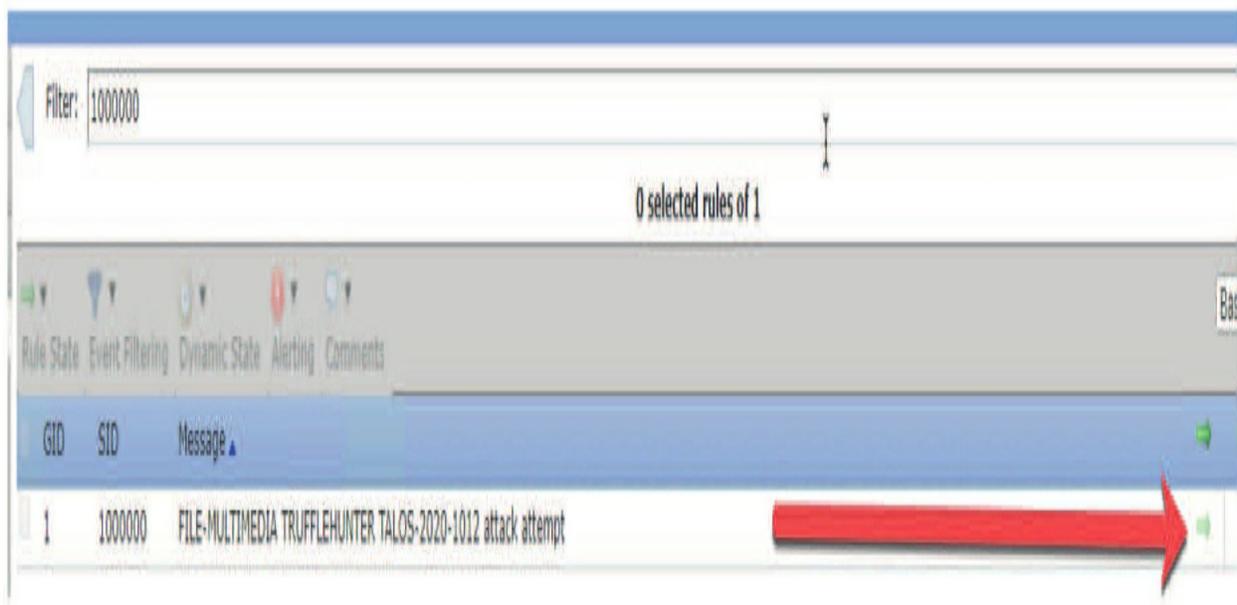
So, I took an existing rule, and created a new rule using the same snort rule, but changed the source and destination network, however, the huge change is that it is now a Pass rule.

This will still work in conjunction with the old rule, except the source network is 172.16.10.0/24, and that traffic will take precedence over the existing rule, and finally not alert on this traffic at all, but still alert on all other traffic matching that rule.

So. Very. Cool.

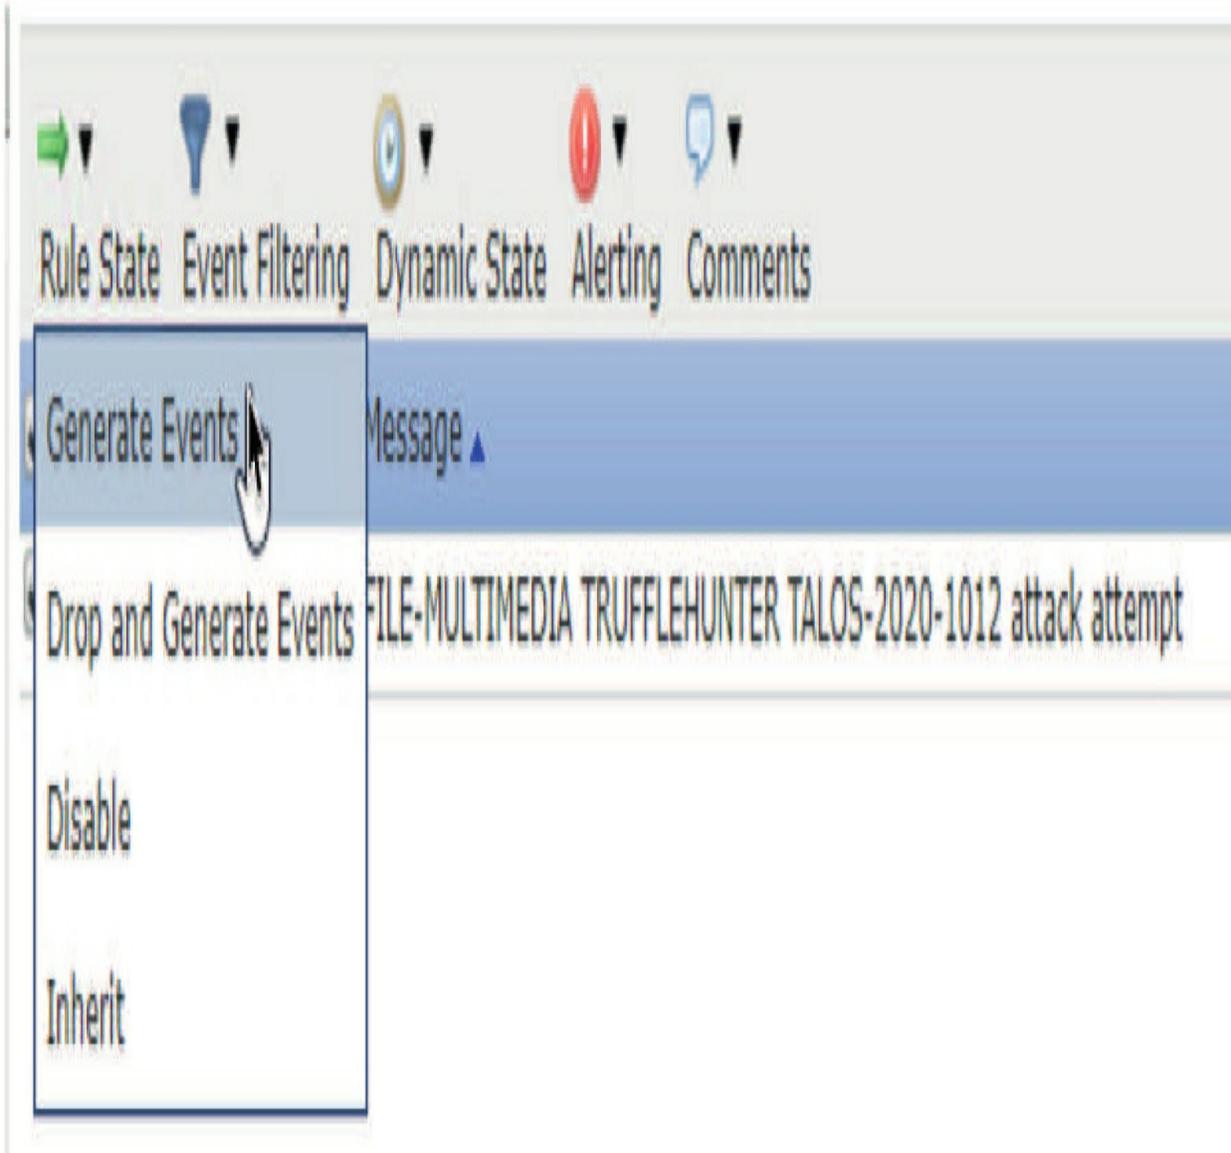
What did you learn? That Pass rules should be used instead of Suppression of a rule.

But wait, this rule is not enabled by default, so we still need to do that still. Here is the disabled rule.



I'll choose to make this an Alert rule only, as it's a Pass rule and we don't

want to drop anything for this rule.



Firepower Recommendations

The recommendations feature has been included with the Firepower product almost from its inception. It has gone by several names over the years—RNA (Realtime Network Awareness) Recommendations, FireSIGHT Recommendations, and now Firepower Recommendations, but the basic feature set and goal of the setting has remained the same. The idea is to take advantage of the host intelligence gained from the passive detection capability to provide tuning recommendations for Snort rules.

For the networks configured with host discovery in the Network Discovery policy, Firepower passively collects host information like the following:

- IP address
- Operating system
- Applications
- Protocols
- Servers (also known as services; these are listening ports)
- Vulnerabilities based on the above information

For the Firepower Recommendations feature, the key item in the list above is vulnerabilities. By comparing the vulnerabilities present on your hosts with the vulnerabilities covered by Snort rules, Firepower can arrive at a recommended ruleset tailored to your environment. And Firepower can recommend which rules should be enabled. It can also recommend which rules should be disabled because they aren't needed in a particular network.

When the Firepower Recommendations process is executed, recommendations are generated, and depending on the policy settings, they can be previewed or implemented.

If you decide to use the recommendations, a policy layer is inserted immediately above the base layer. This layer then enables or disables rules based on the recommendations generated.

To configure recommendations, click the Firepower Recommendations link on the left. If you want, you can expand the Advanced Settings item to go a bit deeper into the internals of the process. This screen is shown here:

The screenshot displays the Cisco Firepower Management Center (FMC) interface for editing a policy. The top navigation bar includes 'Overview', 'Analysis', 'Policies', 'Devices', 'Objects', and 'AMP'. The 'Policies' tab is active, and the 'Access Control > Intrusion' sub-tab is selected. The main content area is titled 'Edit Policy: Demo Policy' and is divided into two sections: 'Policy Information' on the left and 'Firepower Recommended Rules Configuration' on the right. The 'Policy Information' section includes 'Rules', 'Firepower Recommendations', 'Advanced Settings', and 'Policy Layers'. The 'Firepower Recommended Rules Configuration' section shows a message: 'No recommendations have been generated.' Below this, there is a checkbox for 'Include all differences between recommendations and rule states in policy reports' and a sub-section for 'Advanced Settings'. The 'Advanced Settings' section includes a 'Networks to Examine' field with a text input box and a label '(Single IP address, CIDR block, or comma-separated list)'. Below this is a 'Firepower Recommended Rules Configuration' section with a 'Recommendation Threshold (By Rule Overhead)' slider set to 'Medium' and a checkbox for 'Accept Recommendations to Disable Rules' which is checked. At the bottom of this section are two buttons: 'Generate and Use Recommendations' and 'Generate Recommendations'.

Clicking the Generate Recommendations button will run the process to generate rule state recommendations, but it won't actually insert the recommendations layer into the policy.

Firepower Recommended Rules Configuration < Back

Firepower recommends 17732 rule state settings for 47 hosts

- Set 84 rules to generate events
- ✘ Set 4849 rules to drop and generate events
- Set 12799 rules to disabled

View Recommended Changes

View

View

View

Policy is not using the recommendations. Click to change recommendations

Last generated: 2016 Sep 11 11:41:48

Include all differences between recommendations and rule states in policy reports

Advanced Settings

After you click **Generate Recommendations** and allow the process to complete, the screen will update with **View** links pictured here:

Clicking on one of the magnifying glass icons will take you to the **Rules** section of the policy and will insert the appropriate search into the filter bar depending on which recommendation you select: (**Generate**, **Drop and Generate**, or **Disabled**). This allows you to see what the rule states will be if you accept the recommendations.

After you generate recommendations once, the buttons on the page change to **Use Recommendations** and **Update Recommendations**. If you click **Use Recommendations**, the **Firepower Recommendations** layer is inserted into the policy just above the base layer. After this, the buttons change again and you

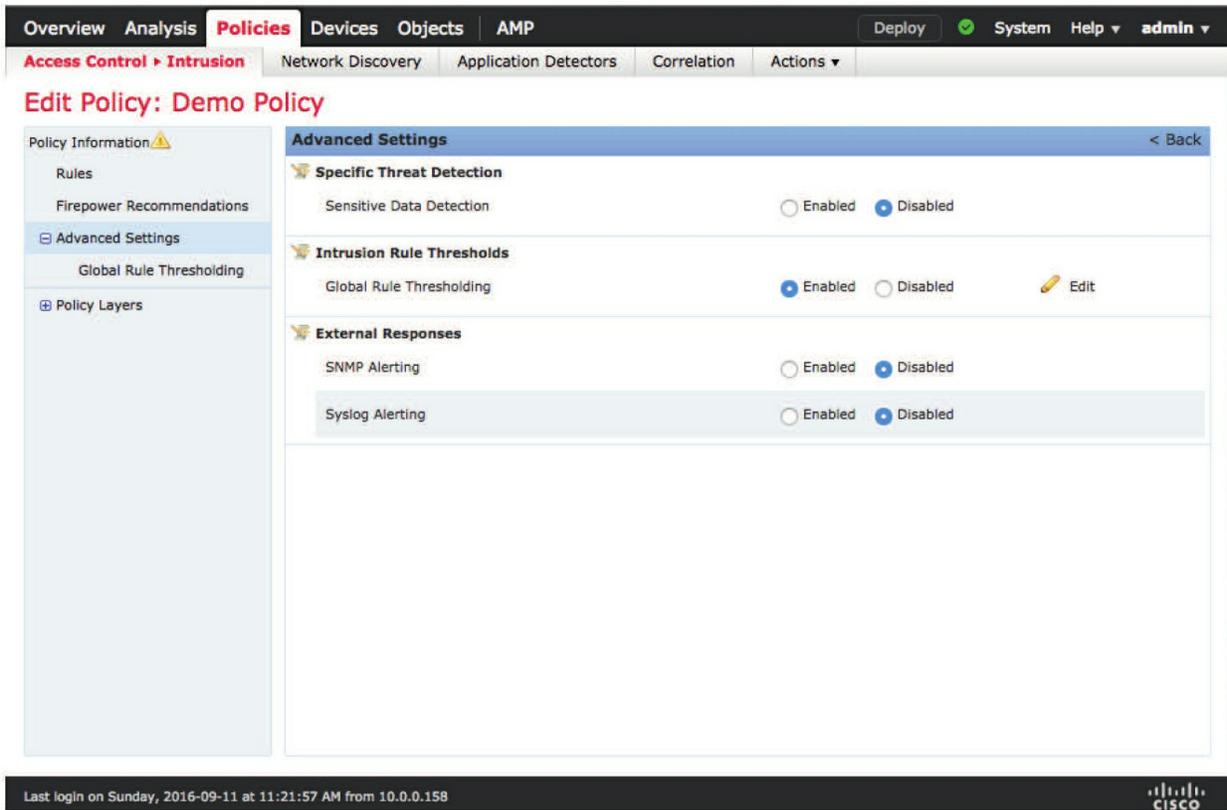
now have Do Not Use Recommendations and Update Recommendations. Clicking the former will remove the Firepower Recommendations layer from the policy, reverting back to the pre-recommendation state.

Firepower Recommendations is a useful tool, but there are a few caveats to be aware of. First of all, it depends on quality Firepower discovery data for your hosts. The system's passive detection capability works best when devices are placed close to your hosts, which means installing Firepower/FTD devices throughout your network. But in the real world, devices are often deployed only at the network's edge. As a result, the only packets the devices process are traveling into or out of the network, which makes it hard to collect reliable information on things like operating systems, services, and applications. The takeaway? If your deployment only includes edge devices, your host data may not be all that useful for Firepower Recommendations.

What's more, Firepower recommendations aren't a one-shot setting. If you use them, they should be updated regularly to make sure the most up-to-date host information is being used. To help you with this, you can schedule a regular task to automatically update recommendations under **System>Tools>Scheduling**.

Advanced Settings

There are only a handful of advanced settings in the Intrusion policy. In fact, if you expand the Advanced Settings item by clicking the plus icon to the left, you'll only see one setting: Global Rule Thresholding.



This highlights the somewhat strange behavior of this Advanced Settings “flattened” view. Only enabled settings will appear in the list. To see *all* the advanced settings, you have to click the Advanced Settings item itself. After this, you can see all the settings—whether they are enabled or disabled—to the right. The next figure shows all the available advanced settings:

Specific Threat Detection

This category of settings contains one item: Sensitive Data Detection and it’s disabled by default. It allows searching for sensitive data within specific protocols. It has been called a “poor man’s DLP” (data loss prevention) solution. The idea is you can detect when a certain volume of specific sensitive data starts flowing through the device, which could indicate a breach in progress as this data is siphoned off to some unknown evil destination.

If you enable this setting, the edit pencil icon will appear. Clicking on it will reveal the configuration options shown here:

Sensitive Data Detection < Back

Configure Rules for Sensitive Data Detection

Global Settings

Mask (Obscure all but the last four credit card or Social Security numbers.)

Networks

Global Threshold (Total occurrences across all data types combined before generating a global threshold event.)

Targets

Data Types
Credit Card Numbers
Email Addresses
U.S. Phone Numbers
U.S. Social Security Numbers (w/out dashes)
U.S. Social Security Numbers (with dashes)

Configuration

Data Type

Threshold (Total occurrences of this data type before generating an event.)

Destination Ports

Application Protocols ✎ Edit

This configuration is contained in the layer: My Changes

There's a bunch of pre-built rules to detect various types of information, like credit card numbers, email addresses, US phone numbers, and Social Security numbers. For each, you select the threshold, destination ports, and application protocols. You can also create your own custom detections using a subset of regular expression pattern-matching syntax.

Each one of the data types, including any custom detections, corresponds to a Snort rule. These rules must be enabled to activate this feature.

Near the top of the page you'll see this link:



Clicking this takes you to the rule management page and filters on rules that

use this preprocessor. Once there, you can enable these rules by setting them to Generate Events.

In practice, the Sensitive Data Detection feature has a couple of disadvantages. First, it has a tendency to produce false positive alerts. When you're looking for a suspicious pattern of numbers across millions or billions of packets, the odds of finding this pattern in benign data are pretty darn high.

To combat this, the threshold for each rule can be adjusted to reduce false positive alerts as much as possible. Of course, this also reduces the sensitivity for detecting an actual breach.

Second, this feature tends to be a bad performer. Yes, this is much better than trying to use regular Snort rules with regular expression matches, but there's still a pretty decent amount of overhead caused by turning on this feature.

Intrusion Rule Thresholds

As with Specific Threat Detection, the Intrusion Rule Thresholds section contains just one item: Global Rule Thresholding, and it's the only advanced setting that's enabled by default. The configuration for this item is shown in the following figure:

The screenshot shows the Cisco Firepower Management Center (FMC) interface. The top navigation bar includes 'Overview', 'Analysis', 'Policies', 'Devices', 'Objects', and 'AMP'. Below this, there are tabs for 'Access Control > Intrusion', 'Network Discovery', 'Application Detectors', 'Correlation', and 'Actions'. The main heading is 'Edit Policy: Demo Policy'. On the left, a sidebar menu shows 'Policy Information' (with a warning icon), 'Rules', 'Firepower Recommendations', 'Advanced Settings', 'Global Rule Thresholding' (selected), 'Sensitive Data Detection', and 'Policy Layers'. The main content area is titled 'Global Rule Thresholding' and has a '< Back' link. Under 'Settings', there are three radio buttons for 'Type': 'Limit' (selected), 'Threshold', and 'Both'. Below that, 'Track By' has two radio buttons: 'Source' and 'Destination' (selected). There are two input fields: 'Count' with the value '1' and 'Seconds' with the value '60'. A 'Revert to Defaults' button is located below the input fields. At the bottom of the configuration area, a note states: 'This configuration is contained in the base policy: Balanced Security and Connectivity'. The footer of the interface shows 'Last login on Sunday, 2016-09-11 at 11:21:57 AM from 10.0.0.158' and the Cisco logo.

Yes, this looks a lot like the same type of thresholding available in intrusion rules and that's because it is—it's a global threshold on all intrusion events. A common misconception is that when you deploy Firepower, you'll get an alert every time a Snort rule matches a packet. This is not the case! Because this Global Rule Thresholding setting is enabled by default, what you'll actually get is one event per rule, per destination, every 60 seconds.

This means if there's a situation where multiple packets traveling to a single host will match a certain Snort rule, the system will only alert on the first one over a 60-second window. If this were a longrunning attack over several minutes, you would see one event every 60 seconds as the attack continued. Of course, if there were multiple destination hosts involved or multiple Snort rules matched the traffic, you'd get more alerts.

The purpose for this setting is to prevent a huge influx of events from exceeding the event processing capability of the device or the FMC. The downside is that you lose some visibility on the extent or volume of actual attack traffic. I highly recommend leaving this enabled until you have tuned your intrusion events to a reasonable level. After that, you can disable this

setting, which will prevent you from being overwhelmed. It'll also give you in-depth visibility into your network traffic.

External Responses

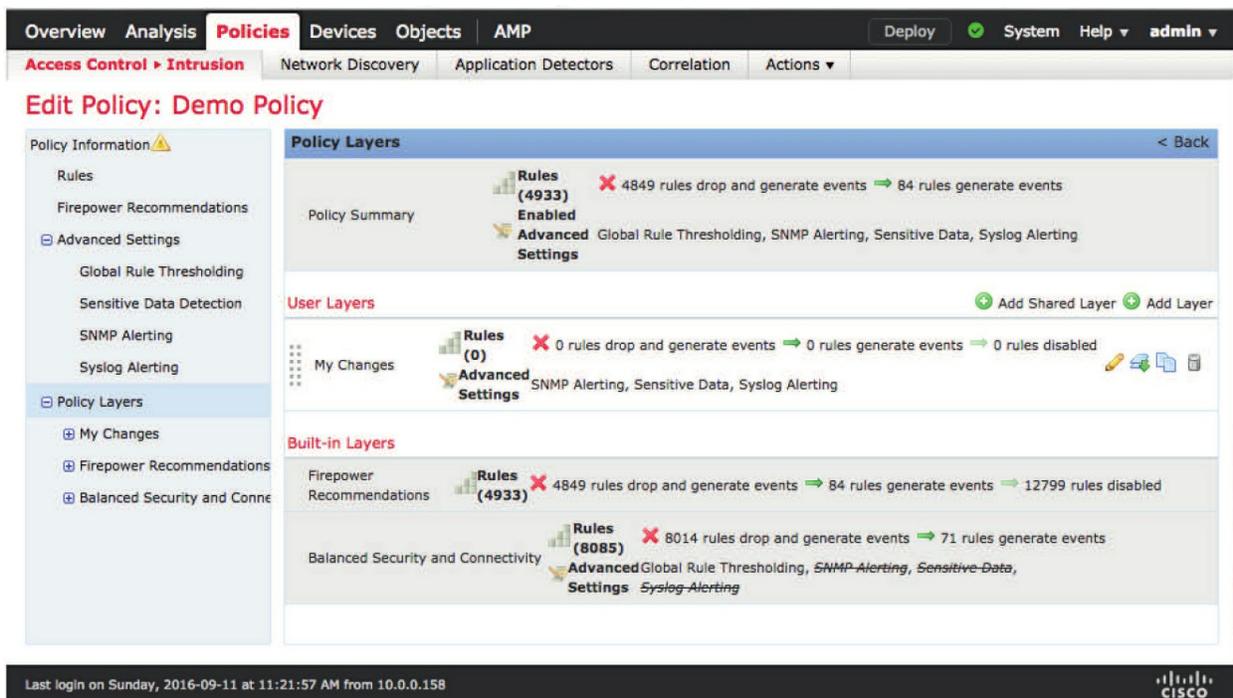
The External Responses section allows the creation of SNMP or syslog alerts in addition to the alerts already displayed on the FMC. The SNMP Alerting section simply defines an SNMP v2 or v3 destination trap server. This destination is used for any rules that have SNMP Alerting enabled.

If you enable Syslog Alerting, you can enter an IP address or comma-separated list of addresses to send UDP 514 syslog messages. Unlike with SNMP, which is on a per-rule basis, enabling this feature will cause any intrusion event to be sent directly to the syslog server from a device in addition to being logged on the FMC.

Policy Layers

The Policy Layers section is a good way to get a visual of the layers in your policy. You can also perform other layer operations like the following:

- Create/delete layers
- Copy layers
- Move layers up or down
- Merge layers together
- Add shared layers from another policy
- Designate a layer as shared



Clicking the Policy Layers link displays the current layers as well as the policy summary as shown below. Here you can see that the base layer is Balanced Security and Connectivity.

Immediately above that is a Firepower Recommendations layer. Then comes the My Changes layer, and at the top, the policy summary lists the enabled rules and advanced settings. Check it out:

Did you notice that you can also tell which advanced settings are enabled or disabled in each layer? If the setting is listed in normal text, it's enabled, if the text is strikethrough, the feature is disabled at that layer.

Use the Add Layer link to add private layers to the policy. You can designate a layer as shared after it's been created if you want to add it to other policies. The Add Shared Layer link is used to add a shared layer from another policy to this one. Clicking it loads a selection list of all of the other shared layers from Intrusion policies on the FMC.

In the next figure, you can see that I created a new layer called Demo Shared Layer 1. Initially, this is a private layer, but I want it to be shared by other policies:

Policy Layers < Back

Policy Summary

Rules (4933) ✖ 4849 rules drop and generate events ➔ 84 rules generate events

Enabled Advanced Settings Global Rule Thresholding, SNMP Alerting, Sensitive Data, Syslog Alerting

User Layers ➕ Add Shared Layer ➕ Add Layer

Demo Shared Layer 1

Rules (0) ✖ 0 rules drop and generate events ➔ 0 rules generate events ➔ 0 rules disabled

Advanced Settings ✎ ➕ 📄 🗑️

My Changes

Rules (0) ✖ 0 rules drop and generate events ➔ 0 rules generate events ➔ 0 rules disabled

Advanced Settings SNMP Alerting, Sensitive Data, Syslog Alerting ✎ ➕ 📄 🗑️

Built-in Layers

Firepower Recommendations

Rules (4933) ✖ 4849 rules drop and generate events ➔ 84 rules generate events ➔ 12799 rules disabled

Balanced Security and Connectivity

Rules (8085) ✖ 8014 rules drop and generate events ➔ 71 rules generate events

Advanced Settings Global Rule Thresholding, *SNMP-Alerting*, *Sensitive-Data*, *Syslog-Alerting*

To fix this, click the pencil icon in the new layer, which brings us to the screen here.

Here we can check the box marked “Allow this layer to be used by other policies,” which turns the private layer into a shared layer:

Layer: Demo Shared Layer 1

< Back

Name

Description

Sharing Allow this layer to be used by other policies

 Rules (0)

→ 0 rules generate events

✗ 0 rules drop and generate events

→ 0 rules disabled

 Manage Rules

 View

 View

 View

 Specific Threat Detection

Sensitive Data Detection Enabled Disabled Inherit

 Intrusion Rule Thresholds

Global Rule Thresholding Enabled Disabled Inherit

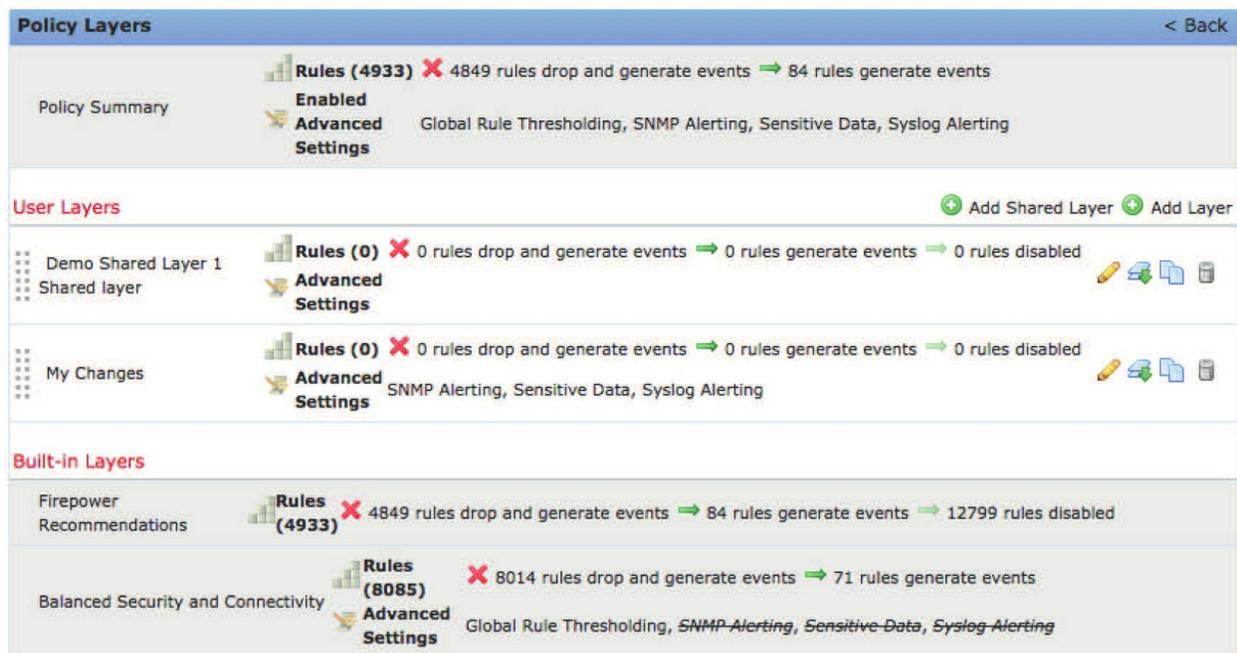
 External Responses

SNMP Alerting Enabled Disabled Inherit

Syslog Alerting Enabled Disabled Inherit

Color legend Above Below

Going back to the Policy Layers list, the layer now contains the text “Shared layer,” letting us know it’s no longer private. To add this layer to another policy, first commit changes to this policy, edit the other policy, and then add the shared layer, Demo Shared Layer 1. The next figure shows our new shared layer:



Clicking the dual rows of vertical dots to the left of a layer and dragging up or down allows layers to be reordered. Only user layers can be reordered this way because built-in layers will always remain at the bottom. The icons to the right also allow editing, merging, copying, or deleting the layer. Keep in mind that only private layers can be merged together.

Committing Changes

When you’re finished editing an Intrusion policy, it’s time to commit the changes, and this can only be done from the Policy Information screen. To navigate there, click on Policy Information in the upper left. You’ll then see the Commit Changes and Discard Changes buttons at the bottom. To save the changes, just click Commit Changes and then add any comments if prompted.

It’s important to understand that as soon as you edit an Intrusion policy, it’s

placed into a “changed” state. It doesn’t matter if you’ve actually changed anything or not, and it doesn’t matter if the yellow

triangle (



) is present next to Policy Information. Each user can only have one policy in this state at a time.

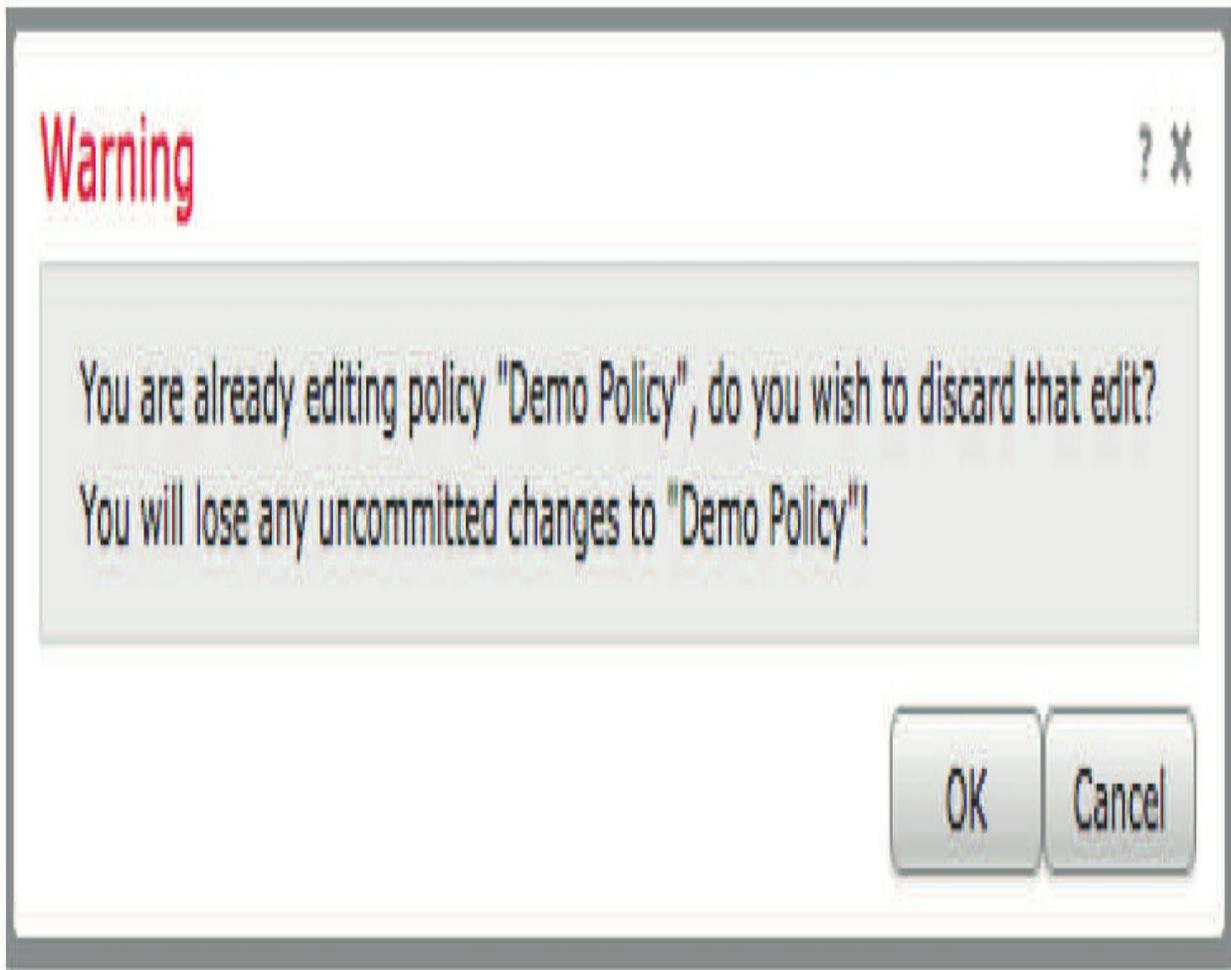
Here’s why that’s important... What if you see an Intrusion policy and want to know more about some of the settings? You click on the policy just to look at which rules are enabled. You didn’t change anything, so you don’t bother to commit or discard changes. You then navigate away or maybe even log off of the FMC. Later, you or another administrator visits the Intrusion policy interface and sees an asterisk by the policy with the text “*You have unsaved changes for this policy.*”

Check out this figure for an example:

Intrusion Policy	Drop when Inline	Status	Last Modified	
Balanced IPS Policy based on the Talos Balanced rule set	Yes	<u>Used by 2 access control policies</u> Policy up-to-date on all 1 devices	2016-09-10 15:26:16 Modified by "alex"	
Demo Policy*	Yes	<i>You have unsaved changes for this policy</i> <u>No access control policies use this policy</u> Policy not applied on any devices	2016-09-09 21:17:53 Modified by "alex"	
Secure IPS	Yes	<u>Used by 3 access control policies</u> Policy up-to-date on all 2 devices	2016-09-10 16:33:10 Modified by "alex"	

So now the questions begin flying. “Who changed this policy? What did they change?” Well, the truth is you didn’t change anything; you just looked at the policy but doing that actually placed the policy in a “changed” state.

Moreover, if you try to create or edit another policy, you’ll see this dialog:



To avoid this unpleasant experience, here's some advice: Always either commit or discard changes before navigating away from an Intrusion policy. And even if you only looked, click Discard Changes as you leave.

Yes, I know you didn't make any changes, but it's just so much better this way!

Configuring and Verifying an Intrusion Policy

In this part of the chapter, I'm going to use the same lab setup and configuration that we used in the last few chapters. Let's create a new policy using Maximum Detection. Why? Because it's fun! Just know that if you're going to use this in production, you really need to disable the Drop when Inline check box.

Okay—let's go to **Policies>Access Control>Intrusion** and then click Create Policy:

Overview

Analysis

Policies

Access Control ▾

Network Disco

Access Control

Intrusion



Malware & File

DNS

Identity

SSL

Prefilter

Management

s currently registered



(1)



After you choose Create Policy, you'll get this screen where you'll enter your information. It's always best to disable Drop when Inline when you choose a base policy that's anything but Balanced. Remember, you have to *tune* the system!

Choose your name and base policy as shown:

Create Intrusion Policy



Policy Information

Name *

Description

Drop when Inline 

Base Policy 

- Maximum Detection ▼
- System-Provided Policies--
- Balanced Security and Connectivity
- Connectivity Over Security
- Maximum Detection**
- No Rules Active
- Security Over Connectivity

* Required

Now click Create and Edit Policy and go get a snack or something because this will take a while to load. After all, it's enabling over 30,000 rules!

You can now see the number of enabled rules as shown below. Again, when using Maximum Detection, it'll enable *all* rules above SID 10,000.

Policy Information 

Rules

Firepower Recommendations

 Advanced Settings

 Policy Layers

Policy Information

Name

Description

Drop when Inline



 **Base Policy**

 The base policy is up to date (Rule Update 2020-01-13-001-vrt)

 **This policy has 31322 enabled rules**

 183 rules generate events

 31139 rules drop and generate events

[No recommendations have been generated. Click here to set up Firepower recommendations.](#)

Commit Changes

Discard Changes

You won't have anything to tune at this point because you need to get packets through the Snort process first.
To do this, add your new IPS policy to an Allow rule Inspection tab in your ACP:

Editing Rule - Inspect all Files for Malware

Name Enabled

Action     

Zones Networks VLAN Tags  Users Applications Ports URLs

Intrusion Policy



None

--System-Provided Policies--

Maximum Detection

Connectivity Over Security

Balanced Security and Connectivity

Security Over Connectivity

--User Created Policies--

TL_IPS_Policy

Add your configured *Variable Set* object (from Chapter 8, “Objects”) to the inspection as well as shown here:



After you add the Variable Set object, click **Add**, then **Save**, and then **Deploy**. Take another break.

At this point you won't have to do much to generate intrusion events, but remember, we're not dropping any traffic since we took off the Drop when Inline check box on our IPS policy.

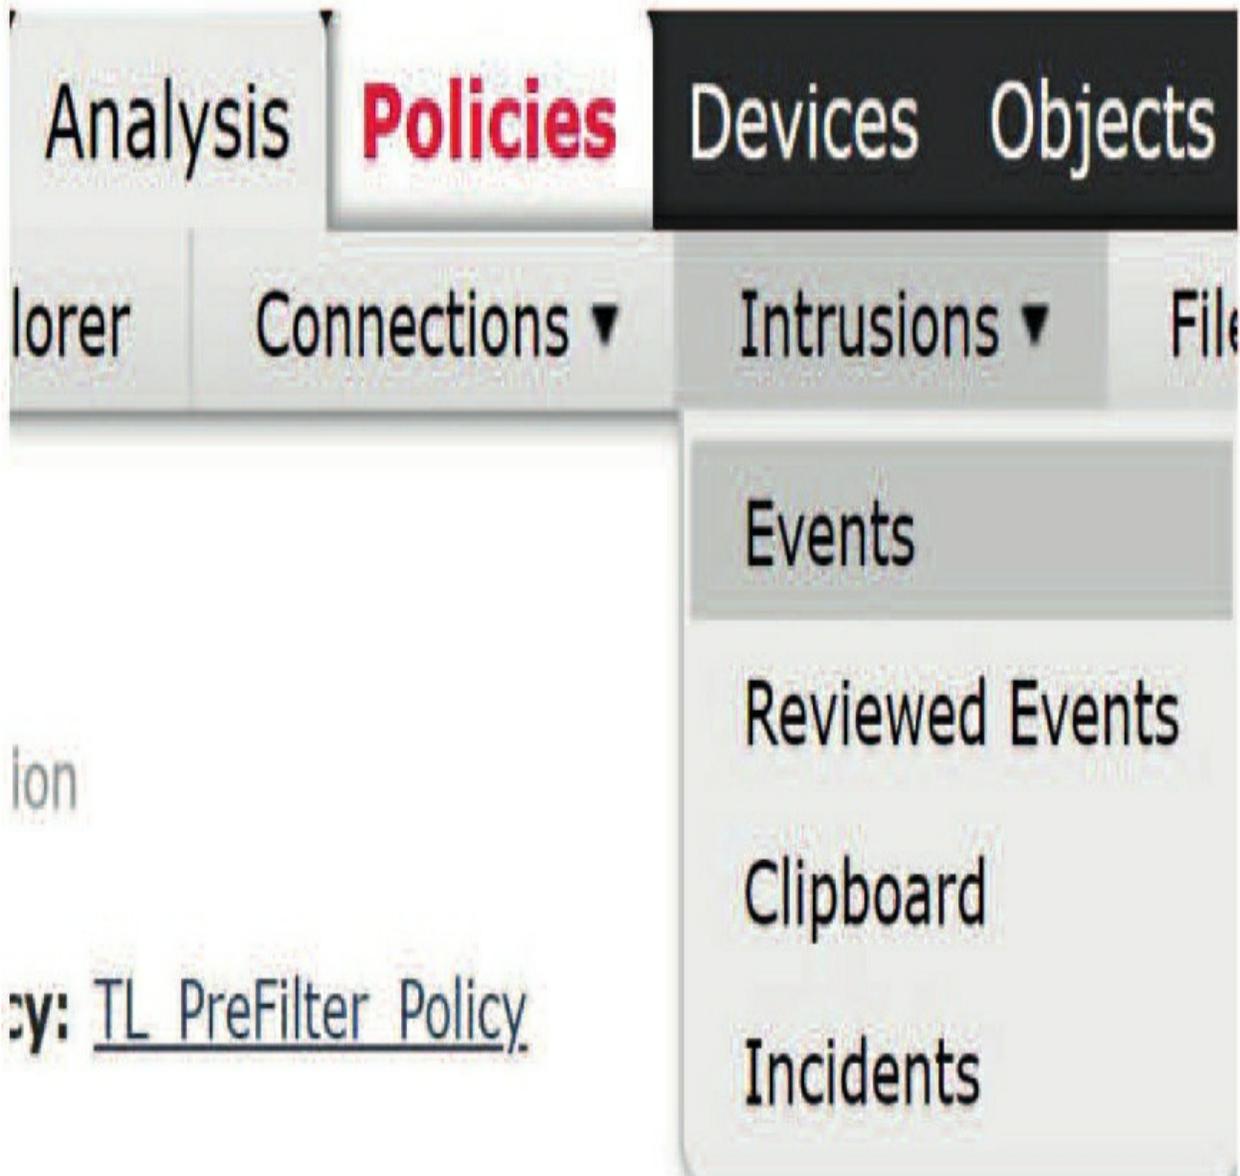
After you deploy, log in to any inside host and just generate some traffic going to various websites. You can also try these sites to generate traffic: <https://www.wicar.org/test-malware.html> and <http://2016.eicar.org/85-0-Download.html>. These two sites *should* work well for you to generate some traffic plus IPS events.

Know that I perform this process every day, so you might not find it as easy as I just made it seem. But please just keep playing with this and you'll get some events, especially with max detection on!

Verifying Your IPS Policy

Once you have finished deploying to your Firepower device(s) and send some packets through the Snort process, you need to verify and tune your policy. Understand that this is just an introduction, and this analysis is fully covered in its own chapter, "Advanced Event Analysis" (Chapter 18).

First traverse to Analysis>Intrusion>Events as shown here:



Now you might have to continue to try to create events by going to more sites, but again, you're using max detection, so you should get a couple right away. Notice that I have two:

Message	Priority	Classification	Count
<input type="checkbox"/> PROTOCOL-DNS-TMG-Firewall-Client-long-host-entry-exploit-attempt.(3:19187:7)	high	Attempted User Privilege Gain	8
<input type="checkbox"/> SERVER-WEBAPP-Lets-Encrypt-SSL-certificate-issuer-detected.(1:43496:2)	medium	Misc Attack	2

There on the left side, I'll click on the blue down arrow to get the Snort verdict for the packet and the rule information that triggered. The next figure displays the result, which is "Would have dropped." This means if the check

box Drop when Inline was checked, the packet would have been dropped.

I understand how hard it is to read that little box that says “would have dropped”, but configuring policy to not drop packets when you first implement the policy allows you to tune out false positives, and not disrupt traffic in the meantime.

Events By Priority and Classification (switch workflow)

[Drilldown of Event, Priority, and Classification](#) > **[Table View of Events](#)** > [Packets](#) 

▶ Search Constraints ([Edit Search](#))

Jump to... ▼

	<u>Time</u> ✕	<u>Priority</u> ✕	<u>Impact</u> ✕	<u>Inline</u> ✕ <u>Result</u>	<u>Source IP</u> ✕
↓	 2019-09-22 11:35:04	medium	0	↓  199.34.228.69 <small>would have dropped</small>	
↓	 2019-09-22 11:19:55	medium	0	↓  199.34.228.69	

This figure shows how you can click further into the packet in order to get more information about the event:

Event Information ▼

Event	SERVER-WEBAPP Lets Encrypt SSL certificate issuer detected (1:43496:2)
Timestamp	2019-09-22 11:35:04
Classification	Misc Attack
Priority	medium
Ingress Security Zone	Outside
Egress Security Zone	Inside
Device	FTD11
Ingress Interface	Outside
Egress Interface	Inside
Source IP	<u>199.34.228.69</u>
Source Port / ICMP Type	443 (https) / tcp
Source Country	 USA
Destination IP	<u>10.11.111.20</u>
Destination Port / ICMP Code	51468 / tcp
Intrusion Policy	TL_IPS_Policy
Access Control Policy	11ACP
Access Control Rule	Inspect all for Malware & IPS
Rule	alert tcp \$EXTERNAL_NET 443 -> \$HOME_NET any (msg:"SERVER-WEBAPP Lets Encrypt SSL certificate 0B 30 09 06 03 55 04 06 13 02 55 53 31 16 30 14 06 03 55 04 0A 13 0D 4C 65 74 27 73 20 45 6E 63 ' 20 41 75 74 68 6F 72 69 74 79 20 58 33 "; fast_pattern:only; metadata:policy max-detect-ips drop, se

Now once you finish tuning your IPS policy, which can take from one day to months, you can set the policy back to Security over Connectivity (SoC) if you're in a production network. In my office, I leave it at max detection because I can tune it every day, and that's the key: Having the time to tune!

Summary

In this chapter you learned all about the Intrusion policy that you create on your FMC. You found that it's configured into an ACP allow rule and finally deployed to your devices where the devices can then make a Snort verdict on the packets moving through your Firepower appliances or FTD devices.

This policy can be thought of as your Snort rule configuration because that's its primary purpose. There are a few advanced settings included, but for the most part, it's all about Snort rules and how to implement and verify your policy.

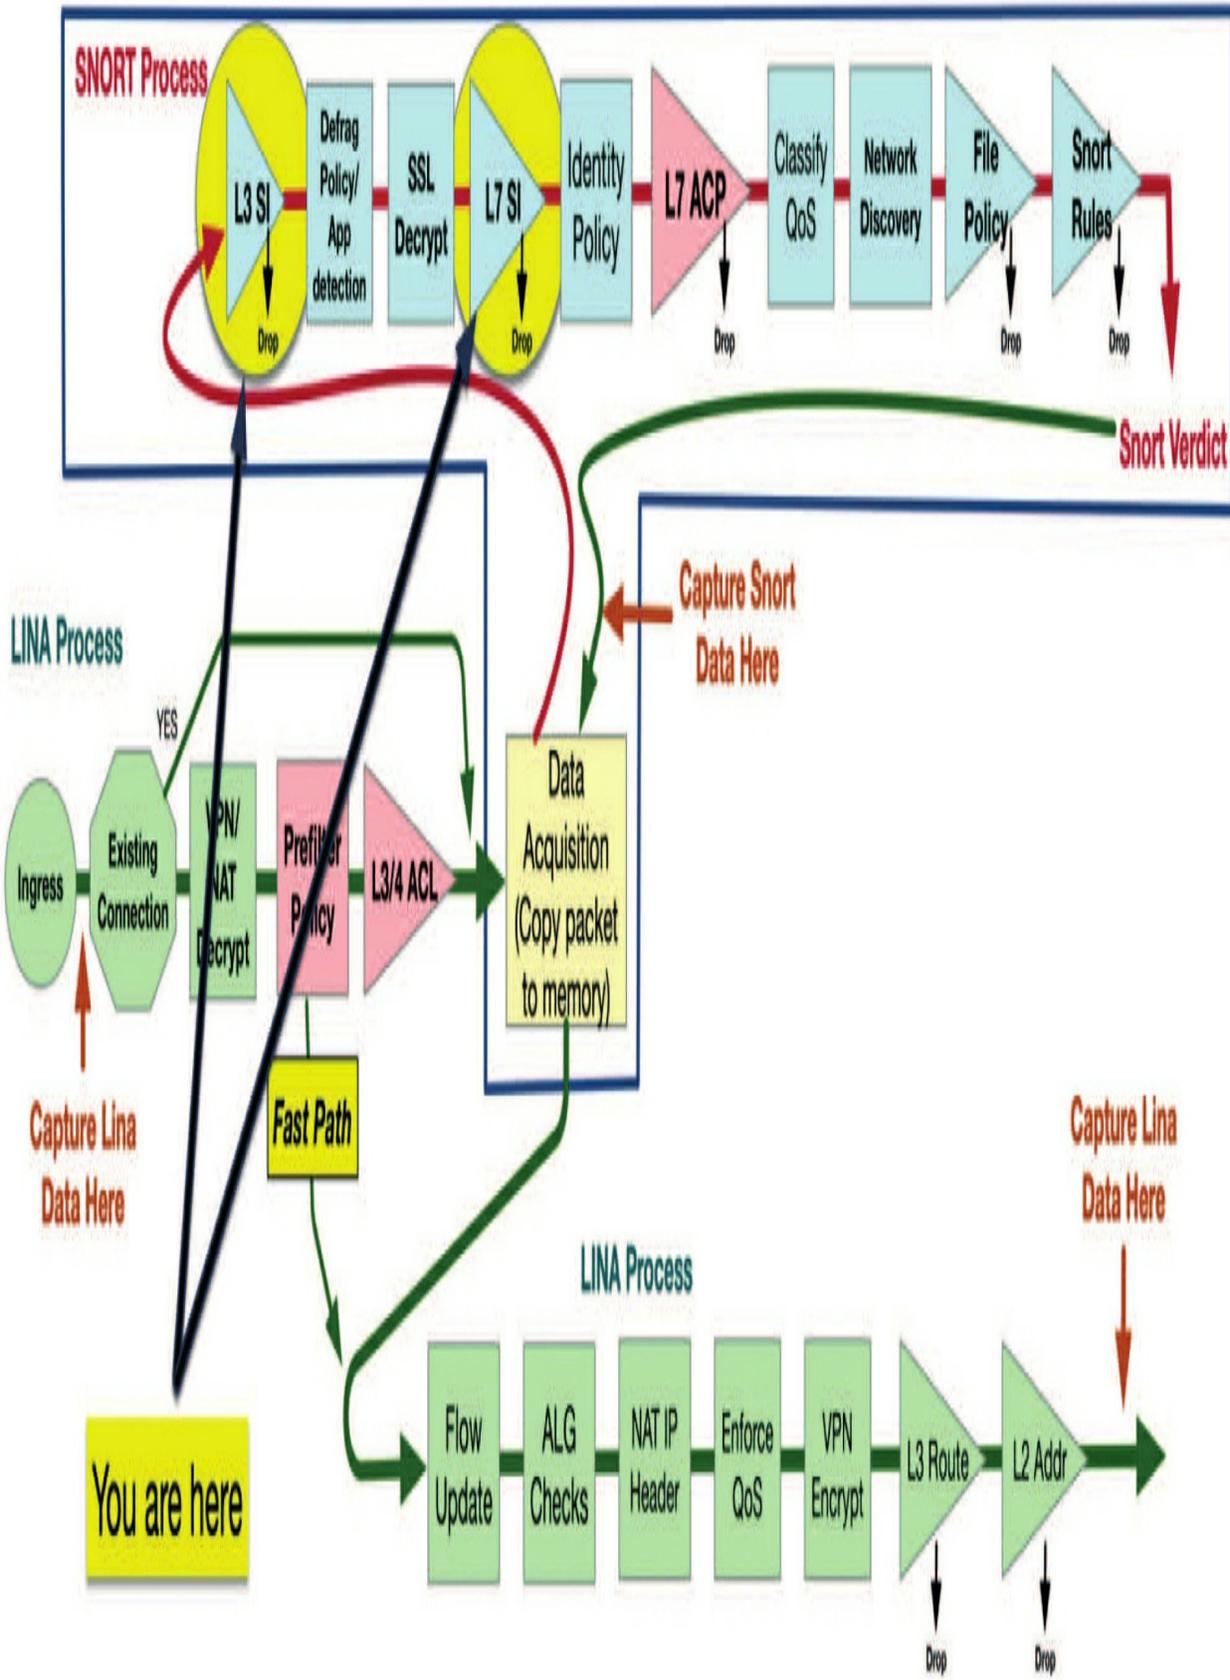
Chapter 13: DNS Policy

The following CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.2 Configure these policies in Cisco Firepower Management Center

2.2.d DNS



The DNS policy is configured and then implemented in the Access Control policy, within the Security Intelligence (SI) tab.

There are two layers of SI in the Firepower system, as you can see in the FTD packet flow figure above. The very first security in the Snort process is the layer 3 SI. Next, the layer 7 SI is processed, which includes URL and DNS objects, if configured.

Both of these enable visibility beyond typical packet sniffing detection by providing additional insight into potentially compromised hosts or encrypted data.

Domain Name System (DNS)

It's often very much taken for granted, but the Domain Name System (DNS) is pretty much the glue that holds the Internet together. It provides the link between the "friendly" name of a host or website and its numeric IP address. Of course, this holds true for both legitimate and malicious sites.

Malware authors and botnet herders usually operate and deploy their malicious software riding on the same infrastructure protocols as everybody else. Since these protocols are well-known, we can scan them for evidence of abuse and even use them against attackers. Now, I'm not talking about actual retaliation against the bad guys—it's more along the lines of leveraging the existing features in the DNS protocol to identify infected hosts or malicious/compromised sites. We can also block malicious DNS activity and help save a few hosts from falling victim in the first place.

The Kill Chain

To fully understand why the DNS policy is so cool, I'll take you on a quick tour showing how hosts get infected with malware. The term *kill chain* comes from the military description of the progression of an attack. It goes something like this:

1. Find: Locate the target.
2. Fix: Make it hard for them to move.

3. Track: Monitor movement.
4. Target: Select the appropriate weapon system to achieve the desired effect.
5. Engage: Shoot them.
6. Assess: Evaluate battle damage and intelligence.

This conceptualization plan lends itself nicely to cyberwarfare as well as the conventional battlefield.

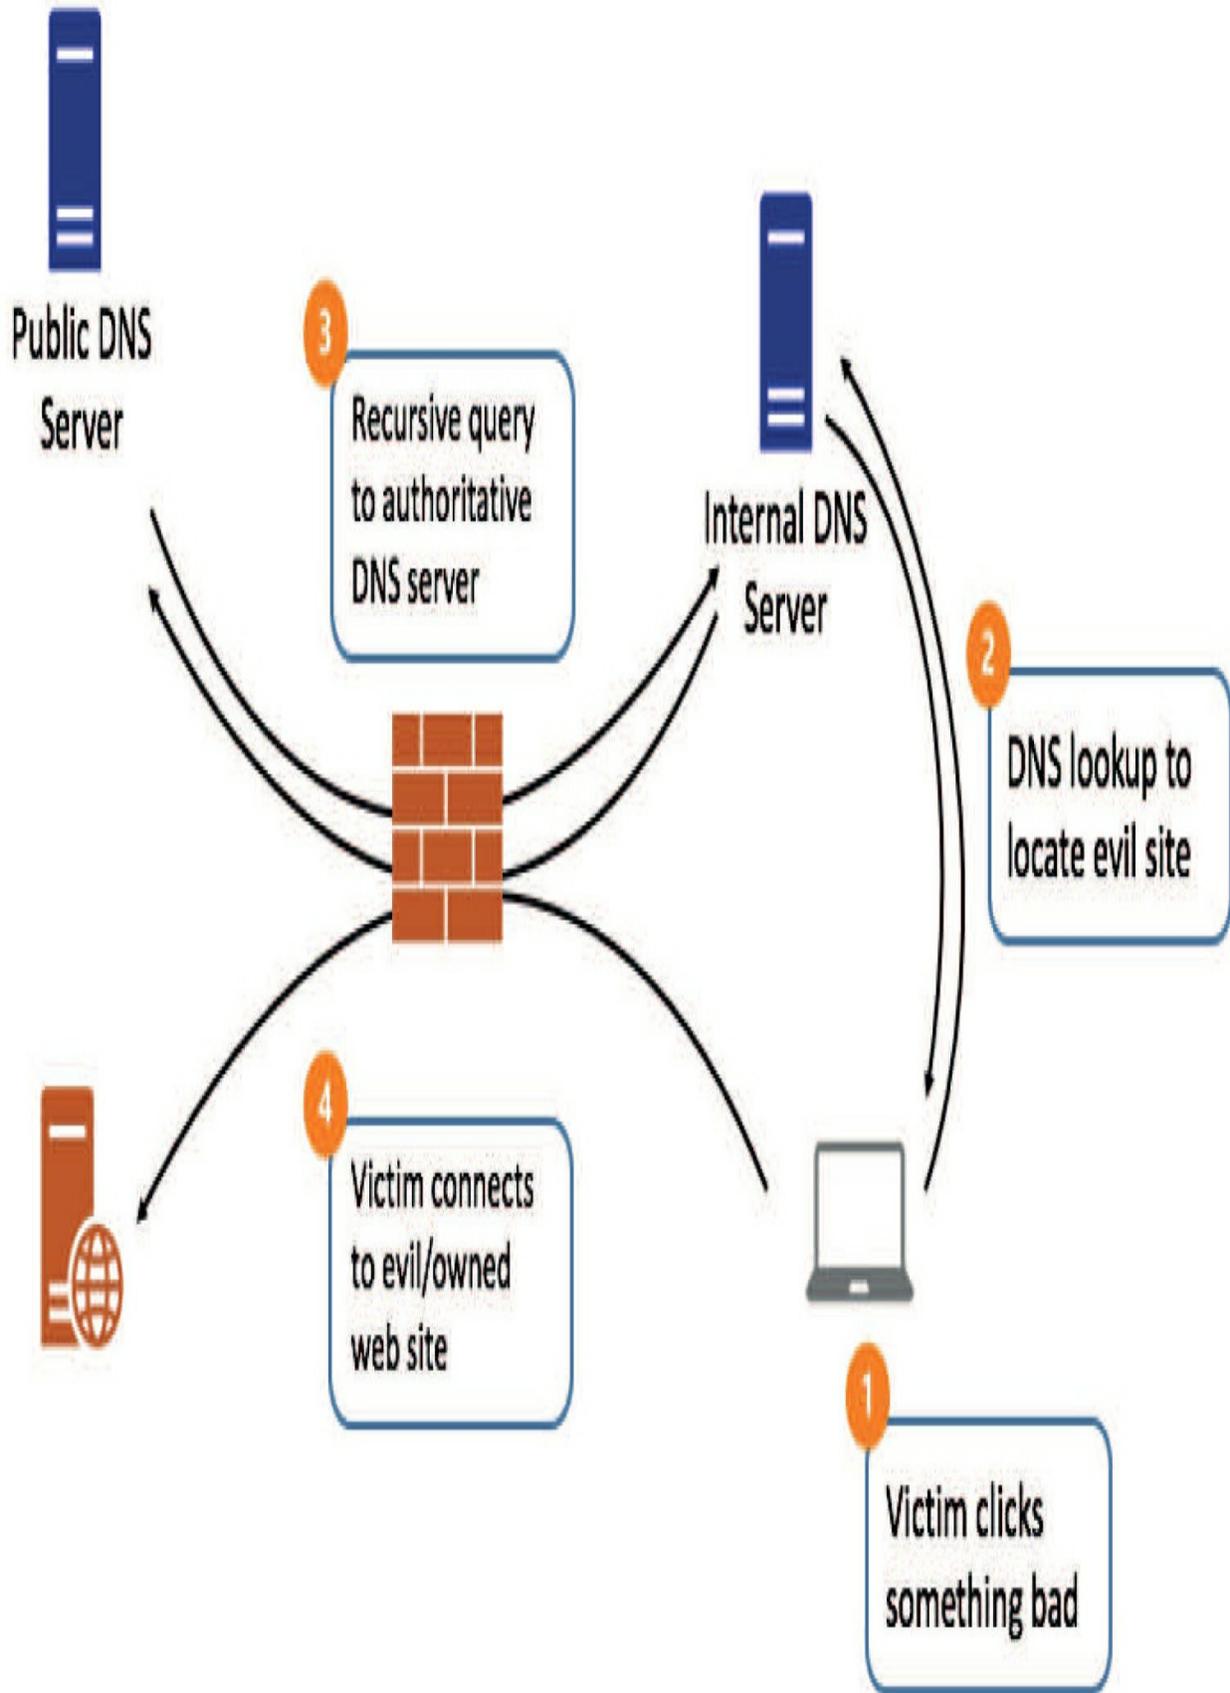
Here's one model of the cyber kill chain:

1. Reconnaissance: Select the target, research and identify vulnerabilities.
 2. Weaponization: Create the weapon—virus, worm, malware, etc.
 3. Delivery: Transmit the weapon to the target.
 4. Exploitation: Execute the malware code: taking action.
 5. Installation: Stay persistently on the target.
 6. Command and Control: Enable attacker to remotely control the weapon.
 7. Actions upon Objective: Take desired action upon the target with data exfiltration, destruction encryption (ransomware), etc.
- Isn't it a relief to know that we can track and potentially even block some of these steps using Firepower's DNS Security Intelligence? I'll admit that our DNS capabilities regarding the first two are limited since much of the work there is research and development and don't require network communications with the victim. But once the actual onslaught is underway, there are a number of ways we can identify and/or block the attack!

A Typical Malware Infection

Let's start off with a simple malware infection. A user on a corporate network clicks on a banner ad or an email link, which directs them to a website that has either been created exclusively for this evil purpose or is a legitimate site that's been compromised to serve malware.

The following figure depicts one scenario on how this can work:



Notice the first thing that happens after the user clicks the baited content is a DNS lookup (2), which is typically routed to the organization's internal name server. Assuming the name server doesn't have the IP address cached, this request will generate a recursive lookup to the name server that's authoritative for the domain in question (3).

This can actually require several lookups even though I've only pictured a single one here. The authoritative name server returns the IP address for the domain, and the IP address is then forwarded down to the victim. The victim then connects directly to the evil/compromised web server and the malware is delivered.

This is analogous to the Delivery step in the cyber kill chain I listed previously. Keep in mind that this doesn't have to be a website. The DNS lookup steps would be the same if the payload were to be delivered via some other method like via File Transfer Protocol (FTP). Beware—delivery can even be achieved through a secure protocol like HTTPS!

So now that you've got a basic picture of how a typical malware attack begins, let's look at some of the options for stopping it.

IP Blacklists

One method of dealing with malware is to use IP blacklists—lists of IP addresses containing compromised or known malicious hosts. These are updated very frequently as new hosts are constantly being infected and infected hosts are taken offline for cleanup. Firepower has just such a feature that we talked about in Chapter 8, "Objects."

IP blacklists give us an effective method of blocking known malicious IP addresses. With them, we can dramatically reduce new malware infections and impede communications for malware that's already running on a victim host. This is all great stuff, but IP blacklists do have some weaknesses too. Here are a few things that make IP blacklisting less effective:

- **Fast-flux DNS:** This refers to the technique used by some command and control systems to hide sites behind an everchanging network of IP addresses.

These typically have a very short time to live (TTL), creating a constantly changing list of IP addresses for a hostname.

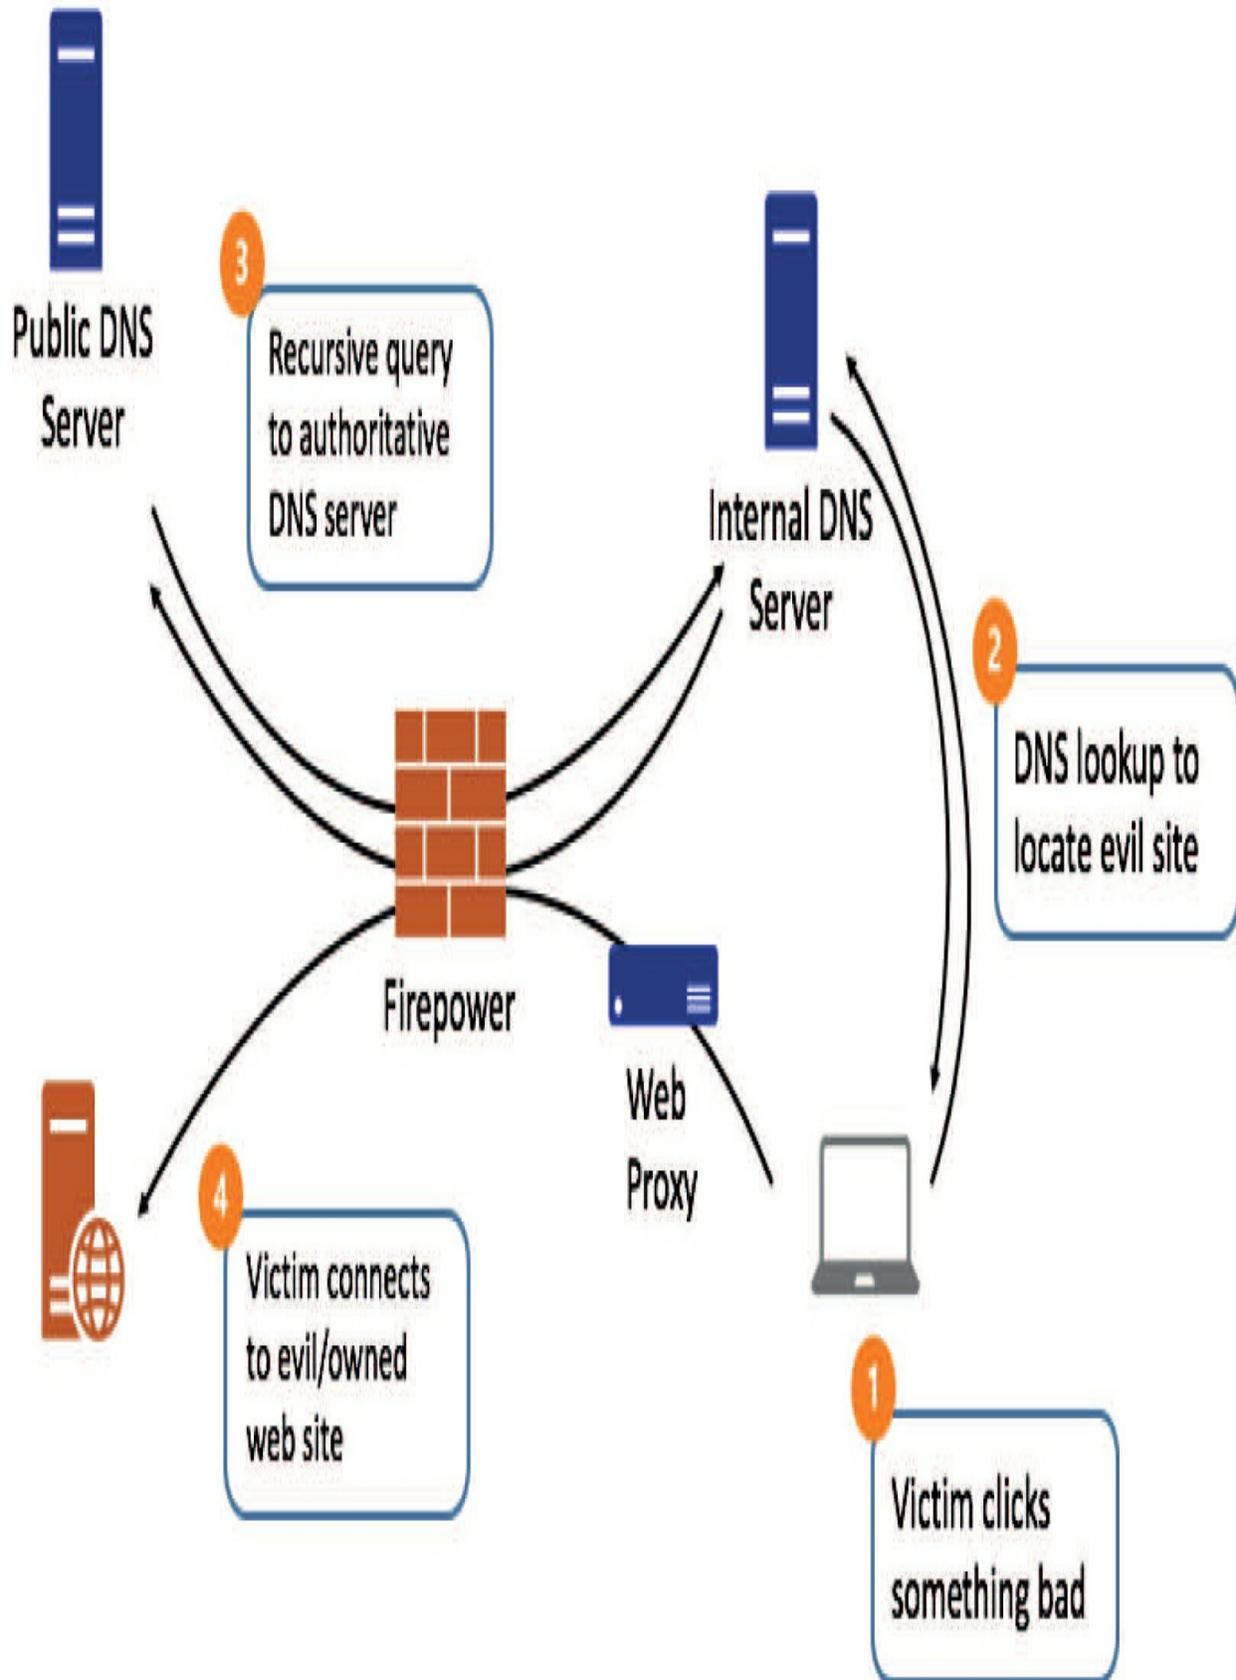
- Update speed: IP addresses can only be blacklisted after they are known to be malicious. Once they're added to the blacklist, the list must be propagated to the various Firepower systems worldwide. The delay in updating your lists depends on the update frequency selected for the Cisco-Intelligence-Feed object, which we covered in Chapter 8. It's during this delay that victim hosts are vulnerable.
- Blacklists aren't secret because attackers have access to the various blacklist feeds from security vendors or services too. If they're monitoring these things, they'll know when one of their compromised systems has been identified and blacklisted. Chances are once they see that, they'll decommission the host and bring in another pre-compromised system with a clean IP reputation.

The Proxy Problem

Another key weakness to IP blacklisting is seen when networks use outbound web proxies. Proxying your users' web traffic is a great way to enforce organizational policies regarding non-business-related web activities.

Client requests go to the web proxy where they're vetted, and if the site is allowed, the proxy initiates the session to the original server on behalf of the client. From that point all traffic to and from the original web server flows through the proxy server.

The following figure pictures the original malware infection, but this time the network has a proxy server implemented:



Now on the surface, this seems a lot like the original malware attack. The only changes I've made are to add the web proxy and to point out that the Internet firewall is actually running Firepower.

So far, we've really only seen the beginning of a malware attack. There's also plenty of activity that occurs post-attack, when the host becomes part of a botnet or the malware checks into its command and control (CnC) servers. So it's nice to know that much of this activity can be detected by Firepower too!

In this case, the proxy's presence doesn't impact Firepower's ability to detect or stop the attack. However, the proxy can impact Firepower's ability to identify the victim. When Firepower detects activity that indicates an active malware compromise, the question isn't, "Can we stop the malware from working?" Because we probably can't. The real question is, "Which host is infected?" The response to a malware-infected host varies a bit, but it's typically along the lines of wipe and reimage. So here, alerting and stopping at least some of the communications is all well and good, but the real value is locating the infected host.

The example shown in the figure above demonstrates that the problem is if we detect the connection to the evil website (4), the source IP address for this traffic will be the proxy server. And we need to know the address of the victim if we're going to deal with its malware infection, right?

What if the IP blacklist contains the public DNS server's address? In this case, we would detect and trace the DNS lookup to the blacklisted IP (3), but the source IP is actually the internal DNS server. Here again, we have to know which client made the original DNS request, but we're being confounded by intermediate devices in the network.

The Proxy Solution

Fortunately, Firepower has a couple of solutions to these proxy and DNS issues:

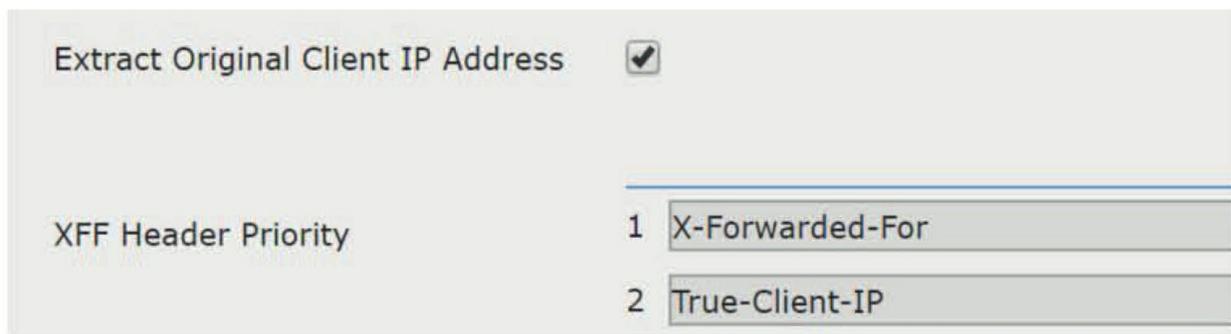
- XFF
- Sinkhole

- So we'll explore these two solutions now.

X-Forwarded-For (XFF) and True-Client-IP

The first solution relates to the web traffic as it exits the network through the proxy. As you now know, while we have intrusion rules and/or IP blacklist entries to identify the traffic, what we really need is to get to ground zero and identify the original, infected host. Most web proxies can add an optional field to the HTTP header for proxied traffic containing the original client IP address of the client.

I'm going to cover this setting during the HTTP configuration in Chapter 23 "Network Analysis Policy (NAP)," but to jump ahead really quick, here's what this configuration looks like in the HTTP Preprocessor:



By default, Snort will pull this original client IP information from the X-Forwarded-For or the True-Client-IP field if either one exists. This information is then populated in the Original Client IP field of the intrusion event.

In prior versions of Firepower, the use of these headers was limited only to intrusion events. A really sweet change in Firepower 6.1 is the extension of this Original Client IP field to Connection and Security Intelligence events. It provides an even greater use case for enabling this feature on your proxy because Firepower can now attribute most event types to the original client with this tool. This means that even Security Intelligence IP blacklist events can now reveal who the original client was!

DNS Sinkhole

Another excellent new Firepower feature is the ability to respond to malicious DNS requests. With the addition of the DNS policy and the Cisco-DNS-and-URL-Intelligence-Feed, we can now quickly respond to attacks in a way that allows identification of infected endpoints—even through forwarded DNS lookups!

The feature that provides this capability is called DNS sinkhole. Through the DNS policy, Firepower can respond to a DNS request in a number of ways:

- Whitelist: Allow the request.
- Monitor the request: Allow and optionally log a Security Intelligence event.
- Return a “domain not found” message: Spoof a response to appear as if it came from the DNS server saying the domain name isn’t found.
- Drop: Silently drop the DNS query.
- Sinkhole: Return the IP address of a sinkhole object in response to the DNS query.

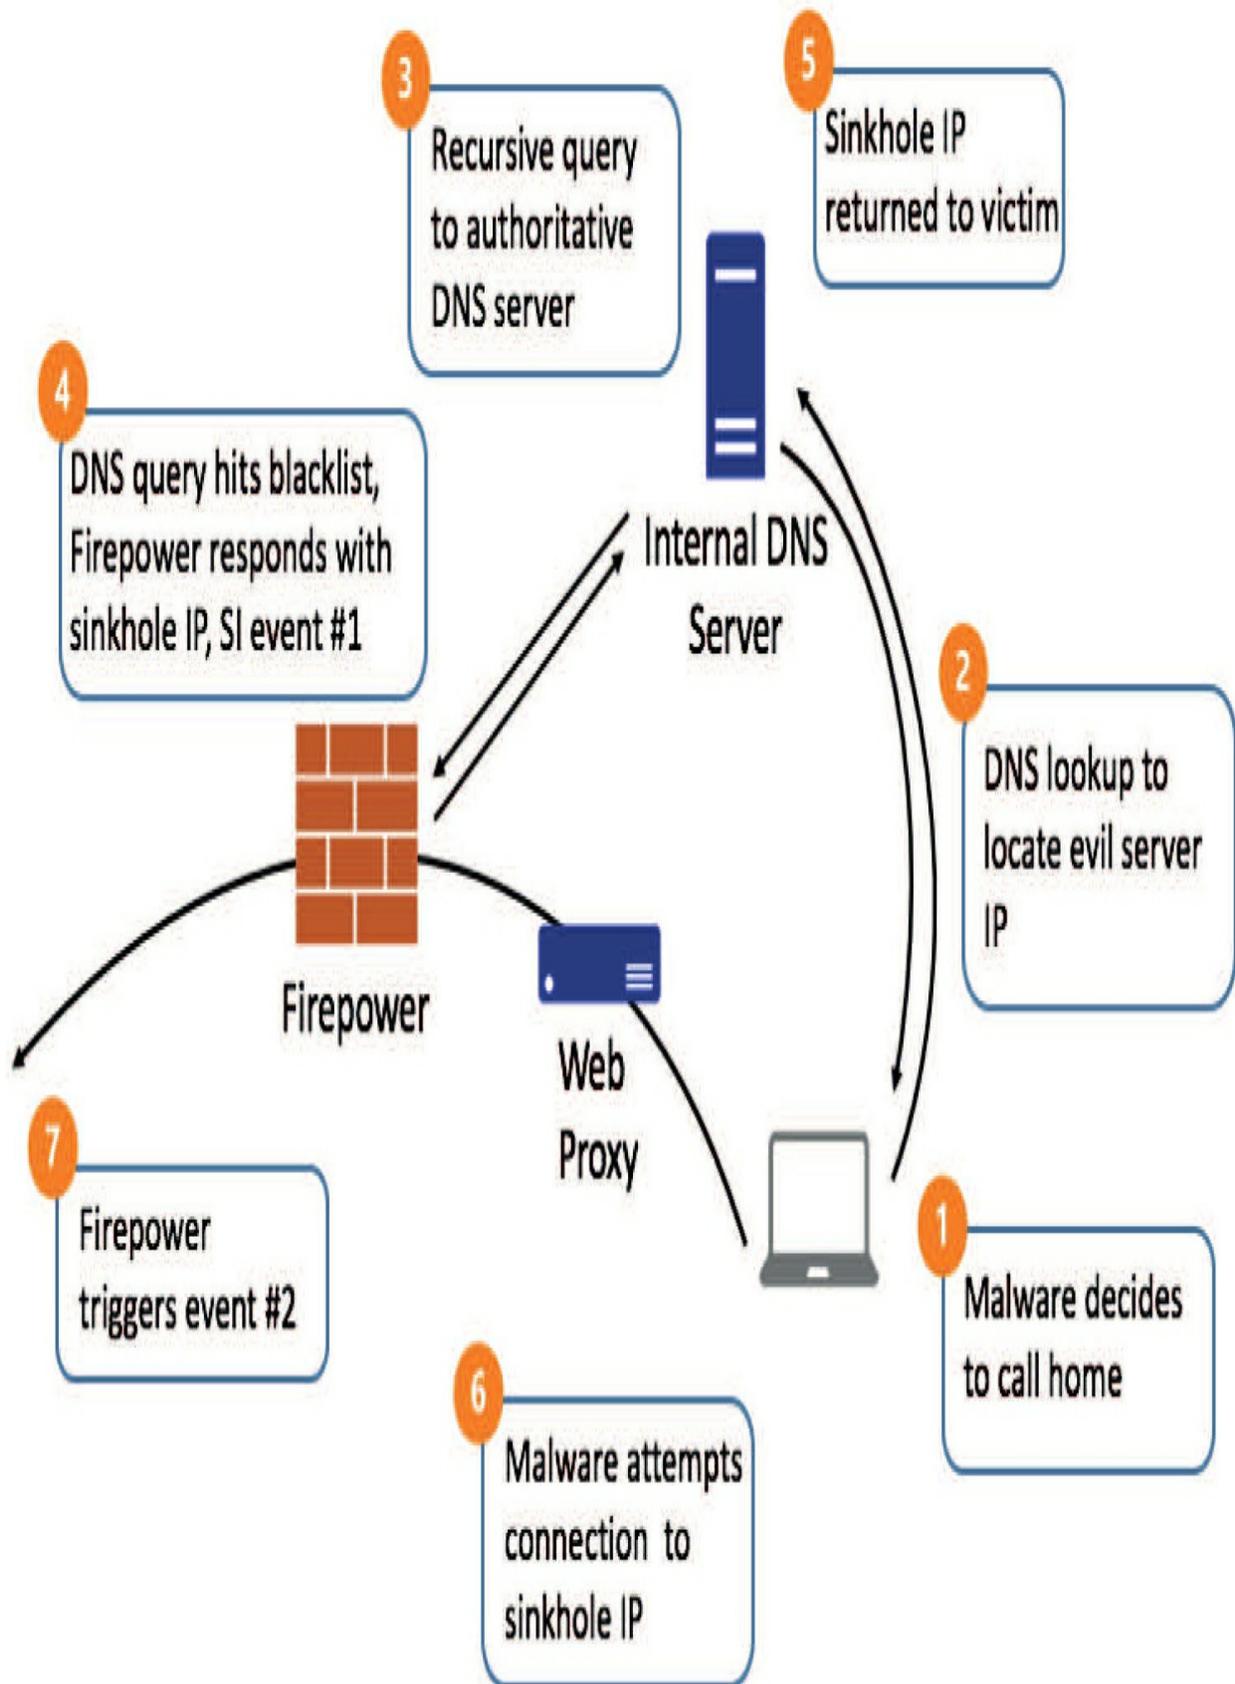
That last option—DNS sinkhole—is definitely the most interesting. It comes to the rescue in the situation where you have an internal DNS server forwarding external name resolution requests on behalf of your internal hosts.

As I said, even if you could just see or maybe block the DNS query, your event intelligence scope would be limited. All you’d get would be the destination IP (a public DNS server), the domain lookup (www.evilsite.com), and the source IP address (your internal DNS server), and none of these offer a clue about which one of your hosts tried to go to evilsite.com in the first place!

And even if you blocked the request...What if the malware just retried the lookup using an alternate DNS name like www.backupevilsite.com? Unless you have DNS blacklist rules for every possible domain it might try, odds are that at least one request is going to slip through. Blocking these requests is good but not foolproof. The only real solution is to get a fix on the host where the requests originated and clean it up!

This is exactly the kind of situation where DNS sinkhole really shines

because it can help you locate the internal host.
Check out the next figure to see how this might work:



Let's walk through the communications shown in the figure in the steps below:

1. A malicious process running on the victim host decides to communicate with an external evil host: `www.evilsite.com`. Reasons could be to check in with its CNC server, get instructions, download additional malware, etc.
2. The operating system performs a DNS lookup to the local DNS server to find the IP address of this host.
3. The local server forwards the query to the authoritative name server for `evilsite.com`.
4. Sure enough, `evilsite.com` is on a DNS blacklist. Firepower intercepts the request and returns a spoofed response containing the IP address of a DNS sinkhole object (1.1.1.1). A Security Intelligence event is generated noting the DNS blacklist that was matched as well as the response (sinkhole).
5. The local DNS server receives this response, adds it to its local cache, and forwards the IP (1.1.1.1) to the victim.
6. The malware attempts a connection via HTTP, FTP, HTTPS, etc. to the sinkhole IP of 1.1.1.1. This may be through the proxy or directly to the Internet, depending on the protocol and network architecture.
7. Firepower identifies the traffic going to the sinkhole, generating event #2—a connection event. Now we can identify the infected host using the event Source IP or Original Client IP field.

Of course, there's room for some flexibility in the configuration and sinkhole behavior. In the diagram, the malware connection to 1.1.1.1 seems to go out to the Internet.

You can choose to block the connection to the sinkhole address to make sure none of the malware traffic actually leaves your network.

You can also have the sinkhole resolve to an actual host if you want. If you set up your own sinkhole system, you can log connections on it as well and gain additional insight into what might be going on with your victim.

Ever so clearly, your sinkhole system shouldn't host any services for general users because the idea is that there should never be a legitimate reason to initiate connections to the sinkhole address.

DNS Policy Configuration

The DNS policy is how we implement the protections I talked about previously. Back in Chapter 8, I demonstrated how to set up and configure various Security Intelligence objects, including DNS Sinkhole objects. In this section, I'm going to show you how to implement the Cisco objects in the DNS policy.

You'll find the DNS policy under **Policies>Access Control>DNS:**

Overview

Analysis

Policies

Access Control ▶ **Access Control**

Access Control

Intrusion

Malware & File

DNS



Policy

Identity

SSL

Prefilter

Here, you'll see the default DNS policy, which you can edit, or you can create a new one. I find it's a good idea to create a new policy because you can always go back to the default if any issues pop up.

So, to create a new policy, click the Add DNS Policy button in the upper right. You'll be prompted to enter your policy's name and an optional description.

Once you load the policy, you'll be greeted by the following screen:

Lammle-DNS



Enter Description

Rules



#	Name	Source Z...	Source Netwo...	VLAN ...	DNS Lists	Action	
Whitelist							
1	Global Whitelist for DNS	any	any	any	Global-Whitelist-for-DNS	Whitelist	
Blacklist							
2	Global Blacklist for DNS	any	any	any	Global-Blacklist-for-DNS	Domain Not Fou...	

The policy comes with two default rules for the global DNS whitelist and blacklist. The Global Blacklist for DNS returns a “domain not found”

response for domains, and the Global Whitelist for DNS permits the traffic, but it'll be subject to further access control inspection.

Notice that the Whitelist rule category is positioned above the Blacklist. This arrangement reflects the default behavior—a whitelist overrides a blacklist. Keep in mind that these two rules can be disabled, but they can't be modified in any other way.

To add a new rule, click the **Add DNS Rule** button on the right, which will display the Add Rule dialog:

Add Rule

Name: Enabled

Action: ✖ Domain Not Found ▼

- ✔ Whitelist
- ↓ Monitor
- ✖ Domain Not Found
- ✖ Drop
- ✖ Sinkhole

Zones

Available Available Zones (0)

Search

- External
- Inside
- Internal
- Outside

Add to Source

Here is where you give your rule a name. Then you can select from the

following actions listed in order of appearance:

- **Whitelist:** Allow the traffic, subject to further access control action/inspection.
- **Monitor:** Log a Security Intelligence event. Allow the traffic, subject to further access control action/inspection.

- **Domain Not Found:** Return a “domain not found” DNS response to the query and log beginning of connection Security Intelligence and connection events. The query will end with no further inspection.

- **Drop:** Drop the DNS query with no further inspection. Log beginning of connection security intelligence and connection events.

- **Sinkhole:** When you select this option, an additional dropdown appears to allow selecting a sinkhole object. Firepower returns the configured sinkhole IP address in response to the query. Logging depends on the configuration of the sinkhole object.
 - If your sinkhole object is configured to log sinkhole connections, the system logs an end of connection event for the follow-on connection.
 - If your sinkhole object is configured to log and block sinkhole connections, the system logs a beginning of connection event for the follow-on connection and blocks that connection.

Okay, so keep in mind that when using a Sinkhole object, your object configuration should be different if your sinkhole is a real host versus a nonexistent one. If your host actually exists and you want to gain additional intelligence from the follow-on connection, your object should be configured to log but not block.

The following figure shows how to add the Cisco objects. This is your minimum DNS policy configuration and should be the minimum default configuration on every Firepower install:

Add Rule

Name: Enabled

Action:  Domain Not Found

Zones Networks VLAN Tags **DNS**

DNS Lists and Feeds  

-  DNS Attackers
-  DNS Bogon
-  DNS Bots
-  DNS CnC
-  DNS Cryptomining
-  DNS Dga
-  DNS Exploitkit
-  DNS Malware
-  DNS Open_proxy
-  DNS Open_relay
-  DNS Phishing
-  DNS Response
-  DNS Spam
-  DNS Suspicious
-  DNS Tor_exit_node
-  Global-Blacklist-for-DNS

Selected Items (15)

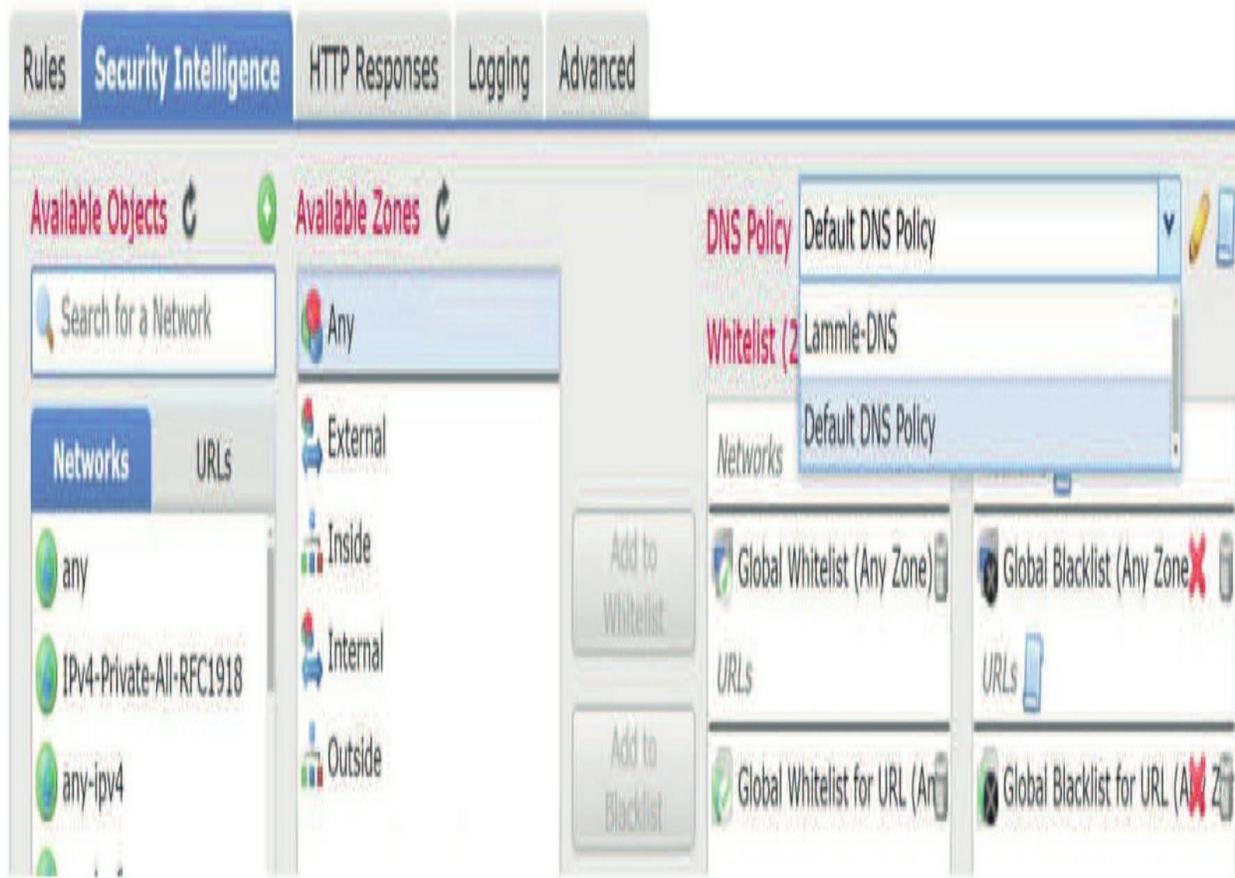
-  DNS Attackers
-  DNS Bogon
-  DNS Bots
-  DNS CnC
-  DNS Cryptomining
-  DNS Dga
-  DNS Exploitkit
-  DNS Malware
-  DNS Open_proxy
-  DNS Open_relay
-  DNS Phishing
-  DNS Response
-  DNS Spam
-  DNS Suspicious
-  DNS Tor_exit_node

Click Add at the bottom and you'll be taken to this screen:

The screenshot shows the 'Rules' configuration page. At the top right, there is a green button labeled 'Add DNS Rule'. Below this is a table with the following columns: '#', 'Name', 'Source Z...', 'Source Netwo...', 'VLAN ...', 'DNS Lists', and 'Action'. The table is divided into two sections: 'Whitelist' and 'Blacklist'.

#	Name	Source Z...	Source Netwo...	VLAN ...	DNS Lists	Action
Whitelist						
1	Global Whitelist for DNS	any	any	any	Global-Whitelist-for-DNS	Whitelist
Blacklist						
2	Global Blacklist for DNS	any	any	any	Global-Blacklist-for-DNS	Domain Not Found
3	Add Cisco Objects	any	any	any	DNS Attackers DNS Bogon DNS Bots DNS CnC (11 more...)	Domain Not Found

Click Save there at the upper right, then go over to your ACP and add the DNS policy into your Security Intelligence tab:



Again, click **Save** and then **Deploy** to make your DNS policy active. This will deploy both layer 3 and 7 SI objects.

After deployment, generate some SI traffic, and then head over to **Analysis>Connections>Security Intelligence Events**, to check out the DNS blocks, IPs, DNS object category, and zones of dropped packets:

Action	Reason	Initiator IP	Initiator Country	Responder IP	Responder Country	Security Intelligence Category	Ingress Security
Sinkhole	DNS Block	10.112.1.240		172.16.77.12		DNS_Intelligence-Malware	External
Sinkhole	DNS Block	10.0.10.57		10.0.1.108		DNS_Intelligence-Malware	DMZ
Sinkhole	DNS Block	10.0.10.57		10.0.1.108		DNS_Intelligence-Malware	DMZ
Sinkhole	DNS Block	10.0.41.26		10.0.96.83		DNS_Intelligence-Malware	DMZ
Sinkhole	DNS Block	172.16.41.140		172.16.83.128		DNS_Intelligence-Malware	External
Sinkhole	DNS Block	10.112.1.240		172.16.77.12		DNS_Intelligence-Malware	External
Sinkhole	DNS Block	10.120.234.122		10.120.19.163		DNS_Intelligence-Malware	DMZ

Summary

In this chapter, I covered the DNS policy and just how well how it provides protection on your Firepower network!

The DNS policy is configured and then implemented in the Access Control policy, within the Security Intelligence (SI) tab.

I discussed the two layers of SI in the Firepower system, and how the very first security in the Snort process is the layer 3 SI. Next, the layer 7 SI is processed, which includes URL and DNS objects, if configured.

Both of these enable visibility beyond typical packet sniffing detection by providing additional insight into potentially compromised hosts or encrypted data.

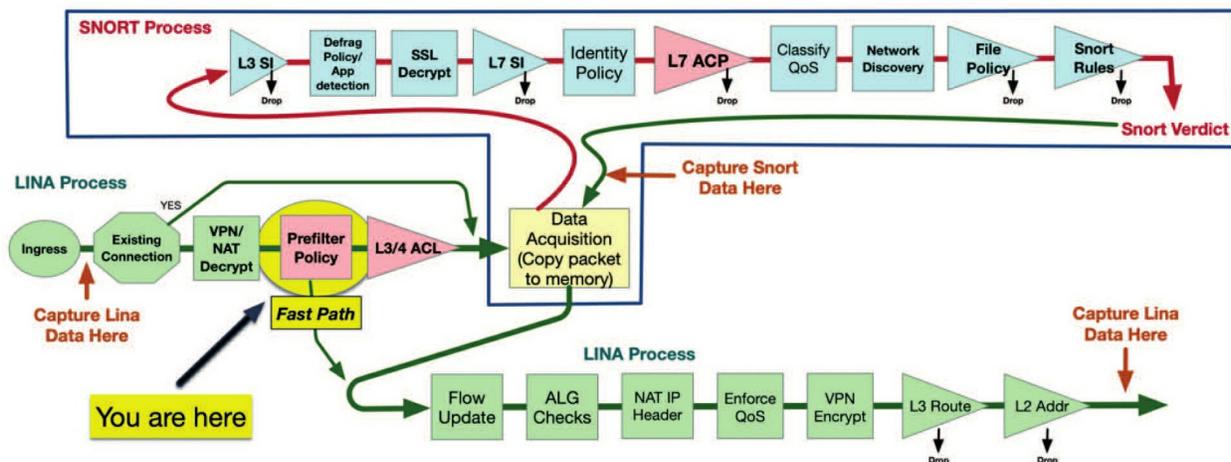
Chapter 14: Prefilter

The following CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.2 Configure these policies in Cisco Firepower Management Center

2.2.g Prefilter



Even though it's not actually part of the Access Control policy, the Prefilter policy is implemented through a setting in the Access Control policy Advanced tab. The Prefilter policy processes traffic first using only layers 3 and 4 to make filtering decisions in the ingress LINA process. You can only have one Prefilter policy per ACP.

This policy is extremely useful if you want to bypass the Snort process with something called FastPath for troubleshooting or other tasks, which we'll cover in detail here.

Overview

Each Access Control policy can have a Prefilter policy associated with it. There's a default Prefilter policy included, which is the one assigned to any new Access Control policy. You can create as many Prefilter policies as you'd like, but again, you can only assign one Prefilter to an ACP.

Prefilter Uses

The Prefilter policy can be broken down into three functions:

- Block traffic
- FastPath traffic
- Rezone tunneled traffic

The first two are pretty straightforward. If you want to block traffic based on criteria like an interface, source/destination IP address, port, or VLAN tag, this is the place to do it. This also applies to traffic you don't want to inspect at all. Passing traffic without inspection using the Prefilter policy is called FastPath.

While blocking traffic is one of the options, the FastPath option is probably the more common use of Prefilter rules. The reason people go with this is usually because there's certain traffic passing through our device that we don't want to inspect or impact in any way.

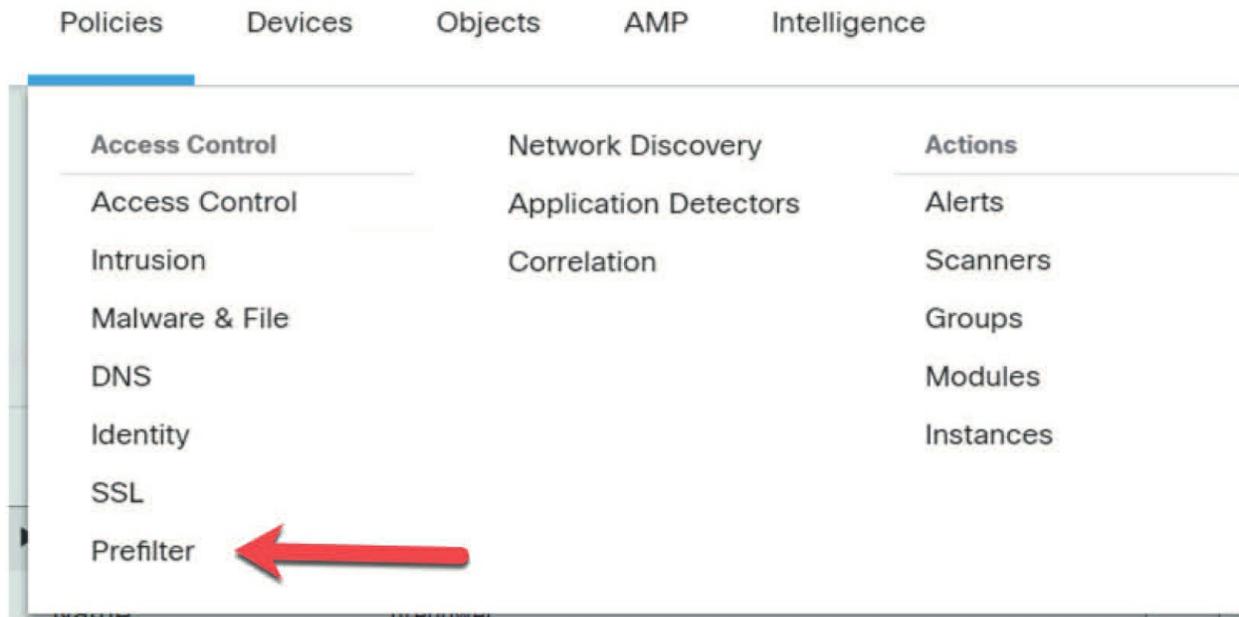
Here's a list of some common use cases for FastPathing traffic:

- Backup applications: They typically transfer high volumes of traffic over long-running connections. Inspecting these flows provides no security benefit and can impact backup speeds as well as device utilization.
- Replication traffic: Database or other types of replication traffic is often encrypted or uses encoding that Snort doesn't understand. These are also long-running high-volume flows that don't need to be inspected.

- Antivirus signature updates: Even though these aren't really long-running flows, they contain data that doesn't lend itself to inspection. Plus, and predictably, we usually see a lot of false positive alerts from Firepower's Advanced Malware Protection file inspection on these flows.
- Voice over IP (VOIP) traffic: Not usually a significant burden to detection, but sometimes latency is an issue with this type of traffic.
- Elephant flows: An elephant flow would be any type of large continuous TCP data flow in terms of bytes. Some examples are the previously mentioned backup or data synchronization traffic and any other traffic that might be found within a data center. With elephant flows, the traffic could be encrypted or of a proprietary protocol, which are considered at low risk of exploitation but can consume large amounts of inspection resources.
- When you want to test an application issue and skip the Snort process. Remember, Prefilter is only layer 3 and 4, but FastPath is super helpful in testing.
- Any other traffic you don't want to inspect.

Policy Creation

Creating rules to FastPath traffic of the types listed above is usually pretty straightforward because if we know which hosts are sending/receiving this traffic, plus their TCP ports, it's easy to create rules to match.



Before we move on to discuss tunneled traffic, let's take a look at the policy itself. You'll find the Prefilter policy under **Policies>Access Control>Prefilter**:

If this is your first visit, you'll notice there's already a default Prefilter policy here. As I said earlier, this policy is used by default for all new Access Control Policies, so it's really best for you to just create your own.

Clicking the edit pencil icon by this policy will bring the screen pictured here:

Default Prefilter Policy

Save Cancel

Default Prefilter Policy with default action to allow all tunnels

Rules

#	Name	Rule Ty...	Source Interface...	Destinati... Interface...	Source Networks	Destinati... Networks	Source Port	Destinati... Port	VLAN Tag	Action	Tunnel Z...

You cannot add rules to the default Prefilter policy. You can change only default action options.

Non-tunneled traffic is allowed

Default Action: Tunnel Traffic

Analyze all tunnel traffic

This policy has a Rules tab as well as a default action. Notice there's no Add Rules button. This is because you can't add rules to the default Prefilter policy, and you can't delete the default policy either. All you can do is change the Default Action and logging options.

Your options for the default action are as follows:

- Analyze all tunnel traffic
- Block all tunnel traffic

Just in case you're wondering what tunneled traffic is, its IP traffic encapsulated within a tunneling protocol. Here's a short list of the protocols the Prefilter policy supports:

- **GRE:** Generic Routing Encapsulation, IP protocol 47. This is a Cisco protocol used to encapsulate a wide variety of Network layer protocols inside virtual point-to-point links over an IP network.
- **IP-in-IP:** Encapsulating an IP packet in another IP packet and covered by RFC 2003.
- **IPv6-in-IP:** Allows tunneling of IPv6 traffic over an IPv4 network.
- **Teredo:** Encapsulating IPv6 packets within IPv4 User Datagram Protocol (UDP) packets.

In the default Prefilter policy, you have the option to analyze traffic within these tunnels or block it. When I say analyze, I mean pass the traffic on to the Access Control policy. Of course, block means... block.

But why do we need a Prefilter policy for this? The reason is that the Access Control policy has no concept of tunnels. As traffic is processed, the tunneling protocol is stripped so the access control, intrusion, and file/malware rules can access the underlying network and application traffic. The Prefilter policy gives you access to these tunneling protocols so you can make decisions on whether to block tunneled traffic or perform some type of Access Control action based on the tunnel characteristics.

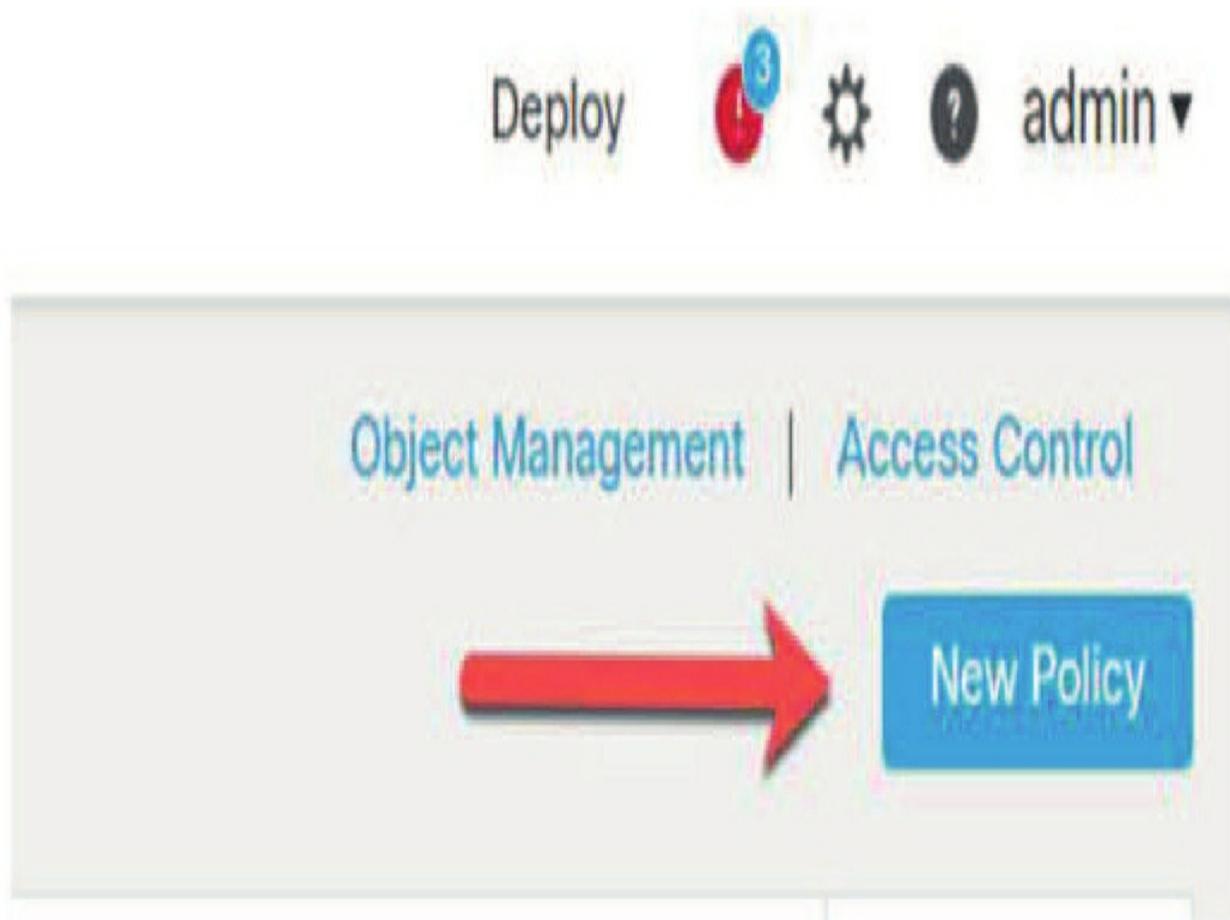
You might still be a little fuzzy on the whole Prefilter policy thing, but just stay with me. One thing that'll help clear things up is to create a new Prefilter

policy and explore the available rules.

Prefilter Rules

First, navigate back to Policies>Access Control>Prefilter. Only this time, click the New Policy button to create your own custom Prefilter policy. Give the policy a name and optional description in the New Policy dialog, shown below.

Then click Save.



The Default Action settings are the same as the settings for the default Prefilter policy, but now we can add rules, as shown next:

Policies Devices Objects AMP Intelligence

Access Control

Network Discovery

Actions

Access Control

Application Detectors

Alerts

Intrusion

Correlation

Scanners

Malware & File

Groups

DNS

Modules

Identity

Instances

SSL

Prefilter



Here we can add tunnel rules and prefilter rules.

Instead of redundantly explaining what's already in the Firepower help, I just included the help page's explanations, which also displays the differences between tunnel and prefilter rules:

Characteristic	Tunnel Rules	Prefilter Rules
Primary function	Quickly fastpath, block, or rezone plaintext, passthrough tunnels.	Quickly fastpath or block any other connection that benefits from early handling.
Encapsulation and port/protocol criteria	Encapsulation conditions match only plaintext tunnels over selected protocols, listed in Encapsulation Conditions .	Port conditions can use a wider range of port and protocol constraints than tunnel rules; see Port and ICMP Code Conditions .
Network criteria	Tunnel endpoint conditions constrain the endpoints of the tunnels you want to handle; see Tunnel Endpoint Conditions .	Network conditions constrain the source and destination hosts in each connection; see Network Conditions .
Direction	Bidirectional or unidirectional (configurable). Tunnel rules are bidirectional by default, so they can handle all traffic between tunnel endpoints.	Unidirectional only (nonconfigurable). Prefilter rules match source-to-destination traffic only.
Rezone sessions for further analysis	Supported, using tunnel zones; see Tunnel Zones and Prefiltering .	Not supported.

We just covered blocking and FastPath, but what's this "rezone" that's mentioned in the Firepower help page above? Well, you remember when I said the Access Control policy has no concept of tunnel protocols? But what if you want to perform some type of inspection only on traffic within certain tunnel links? You'd need a way to define that the traffic is within a tunnel for the Access Control policy, and that's exactly what tunnel rules are for! The process goes something like this:

1. Define a tunnel zone object, which is a specific object type that contains your tunnel name and optional description.
2. Create a tunnel rule to match traffic for this tunnel. Define tunnel encapsulation and optionally interfaces, endpoints, and VLAN tags.
3. Assign your tunnel zone object to the rule.
4. After saving your Prefilter policy, associate it with your Access Control policy and create rules using your tunnel zone object.

You're now all set to specify Access Control actions based upon your tunnel object. The access control rule will have access to the underlying

network/application protocols. It still doesn't know what a tunneling protocol is, but you've identified your tunnel traffic, allowing you to customize your Access Control actions. Let's take a look at some examples.

Our goal is to create a rule to FastPath traffic to our antivirus update servers. When our clients connect to these servers to download updates, we're suddenly getting a bunch of false positive malware events.

Clicking the **Add Prefilter Rule** button displays the dialog in the next figure: Within this dialog, we need to give our rule a name and select an action.

i Prefilter rules perform early handling of traffic based on simple network characteristics. Fastpathed traffic bypasses access control and QoS.

Name

Enabled

Insert

below rule ▼

Action

Analyze ▼

Time Range

▼ +

Interface Objects

Networks

VLAN Tags

Ports

Comment

Available Interface Objects **C**

Q Search by name

External

Firepower_Appliance

Inside

Internal

Outside

Add to Source

Add to Destination

Source Interface Objects (0)

any

Destination Interface Objects (0)

any

The actions available are as follows:

- Analyze: Send the traffic on to the Access Control policy; inner headers will be used for analysis.
- Block: Block the traffic.
- FastPath: Allow the traffic without sending it through the other policies (Access Control, Intrusion, File, etc.).

In our case, we'll select FastPath to allow our antivirus updates to pass without inspection. Using the Interface Objects, Networks, VLAN Tags, and Ports tabs, we'll specify the criteria for the traffic. Be sure to be as specific as possible here because this FastPath rule is actually creating a hole in the detection armor of your Firepower system. So you want to be really sure only the specific traffic you intend to pass gets by with this rule!

As you can see here, I've clicked the Network tab and added our AV-servers network object to the Destination Networks column:

Finally, I'm going to add the port used by the antivirus application—8014/tcp. I used a port object created for this goal and added the destination port to our rule:

The last thing to decide is if we want to log connection events for this traffic, and because there's really nothing interesting to see here, I'll leave this set to the default of no logging.

Clicking the Save button saves our rule.

Next, let's create a sample rule to tag all GRE traffic with a custom tunnel zone. You can create your tunnel zones ahead of time by navigating to **Objects>Object Management**, as shown below.

Object Management Intrusion Rules

Add Tunnel Rule

Tunnel rules perform early ha

Name

Action Analyze

Match tunnels only from source

Match tunnels from source and

Interface Objects Tun

Available Interface Objects

Search by name

- Arris-Router
- Garage-Switch
- MSP-Firewall
- Routed-ESX
- Routed-Switch

- Network
- Port
- Interface
- Tunnel Zone 
- Application Filters
- VLAN Tag

Name
10.0.0.64-28-NAT-Pool
7030-device
any

and QoS.

1

Comment Logging

Interface Objects (0)

Add to Source

Add to Destination

Add Cancel

From there you can click on Tunnel Zone from the list on the left. Give your object a name and optional description.

You can also add this object on the fly without the need to leave your Prefilter policy and visit the Objects menu.

To add a tunnel rule, click the Add Tunnel Rule button. This displays the dialog shown next:

There are several options available for a tunnel rule:

- Name: Descriptive name for your rule.

- Action: Same as a Prefilter rule; actions are Analyze, Block, FastPath.

- Match tunnels only from source: For the rule to match, traffic must originate from one of the source interface objects or tunnel endpoints and leave through one of the destination interface objects or tunnel endpoints. This refers to the session initiation (who sent the SYN). All Prefilter rules are bidirectional, so matching return traffic for established connections is assumed.

- Match tunnels from source and destination: Match traffic from source to destination and destination to source. This means the rule will match sessions initiated in either direction.

- Assign Tunnel Zone: Select a tunnel zone object previously created or create a new one on the fly by clicking the green plus icon to the right.

- Interface Objects: Optional source and destination interface objects.
- Tunnel Endpoints: Optional source/destination IP addresses.
- VLAN Tags: Optional VLAN tags.
- Encapsulation & Ports: This option is required. Check one or more tunnel types to match traffic.
 - GRE
 - IP-in-IP
 - IPv6-in-IP
 - Teredo Port (3544)

- Comment: Optional rule comment.

- **Logging:** Only available for the Block or FastPath actions. Will log a connection event if traffic is FastPathed or blocked.

To add a simple rule to tag all GRE tunnel traffic, we'll leave the Interface Objects, Tunnel Endpoints, and VLAN Tags tabs at their default settings. On the Encapsulation & Ports tab, we'll check GRE and assign our GRE-tunnel tunnel zone object. This rule is shown below:

Add Tunnel Rule

? X

 Tunnel rules perform early handling of non-encrypted encapsulated traffic, using outer IP headers. Fastpathed traffic bypasses access control and QoS.

Name Enabled

Insert

Action

Assign Tunnel Zone  

Match tunnels only from source ()

Match tunnels from source and destination ()

Interface Objects

Tunnel Endpoints

VLAN Tags

Encapsulation & Ports

Comment

Logging

Encapsulation Protocols:

GRE

IP-in-IP

IPv6-in-IP

Teredo Port (3544)

Add

Cancel

When finished, click the Add button to add the rule to our custom Prefilter policy.

We now have two rules in our sample policy as you can see in the following figure. The Prefilter policy processes rules in order from top to bottom. Rules evaluate traffic until there's a match. Once a rule matches traffic, that traffic won't be processed by subsequent rules. In our case, the rules will never conflict, so the order isn't important. But if you do have multiple rules that could match the same traffic, you'll want to place the more specific rules higher in your policy.

This works just like a firewall rule set or router access control list (ACL)—familiar concepts in networking and security.

Demo Prefilter

You have unsaved changes Save Cancel

Enter Description

Rules

Add Tunnel Rule Add Prefilter Rule Search Rules

#	Name	Rule Ty...	Source Interfac...	Destinat.. Interfac...	Source Networks	Destinat.. Networks	Source Port	Destinat.. Port	VLAN Tag	Action	Tunnel Z..	
1	Pass AV updates	Prefilter	any	any	any	AV-servers	any	Symantec	any	Fastpa...	na	0
2	All GRE Traffic	Tunnel	any	any	any	any	any	GRE (47)	any	Analyze	GRE-tunni	0

Non-tunneled traffic is allowed

Default Action: Tunnel Traffic

Analyze all tunnel traffic

1 Row Selected

Displaying 1 - 2 of 2 rows Page 1 of 1

Rule order can be changed by clicking and dragging a rule up or down in the policy. It can also be changed by editing the rule and modifying the rule number in the upper-right corner.

Summary

The Prefilter policy actually processes traffic first using only layers 3 and 4 to make filtering decisions in the ingress LINA process.

You learned how the Prefilter policy is implemented through a setting in the Access Control policy Advanced tab, and that you can only have one Prefilter per ACP.

The Prefilter policy is extremely useful if you want to bypass the Snort process with something called FastPath, either for troubleshooting or other purposes.

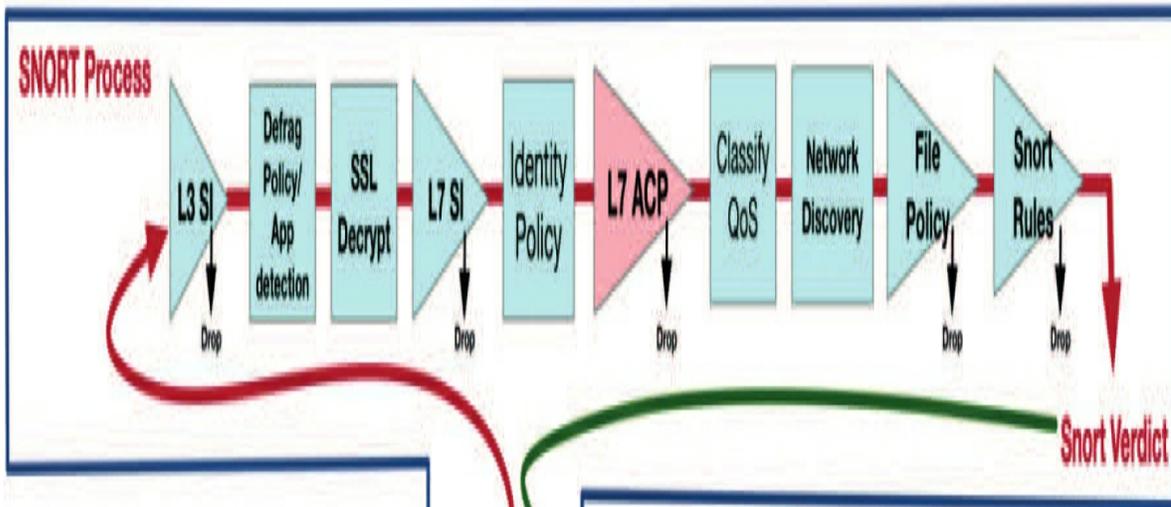
Chapter 15: Network Address Translation (NAT)

The following CCNP Security SNCF exam objectives are covered in this chapter:

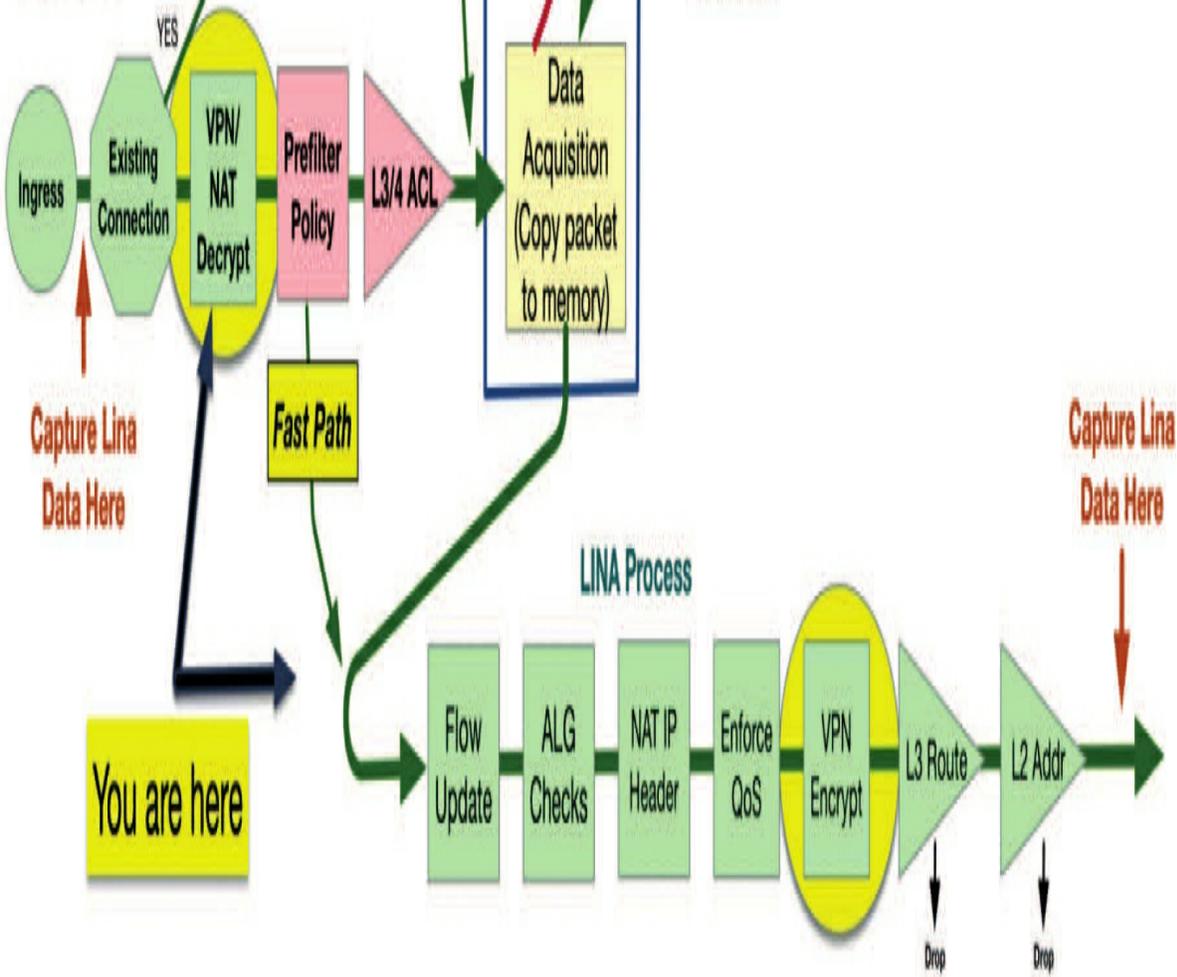
2.0 Configuration

2.5 Configure devices using Firepower Management Center

2.5.b NAT



LINA Process



Now is a perfect time to introduce you to Network Address Translation (NAT), Dynamic NAT, and Port Address Translation (PAT), with the latter also known as NAT Overload.

As always, I'll shore up everything we talk about with examples and figures and by demonstrating all the NAT commands in the exam objectives on Firepower/FTD.

The first thing you'll want to understand is Cisco's straightforward objective for this chapter: You have hosts on your inside corporate network using RFC 1918 addresses and you need to give them Internet access by configuring NAT translations. And because the objectives involve Firepower appliances (7000/8000) as well as FTD, I'm going to cover both in this chapter.

We'll configure NAT on a Firepower system, implement it, and follow up with verification. It's key to remember that these actions don't take place within the Snort process, they happen in the LINA, so we'll need to use the CLI for verification.

We'll get started by going over basic NAT functionality and common terminology.

What Is Network Address Translation (NAT)?

Network Address Translation (NAT) has quite a bit in common with Classless Inter-Domain Routing (CIDR) because NAT was originally created to slow the depletion of available IP address space by enabling multiple, private IP addresses to be represented by a much smaller number of public IP addresses.

Since then, we've discovered that NAT's also really useful for network migrations and mergers, server load sharing, and creating virtual servers. Because NAT decreases the overwhelming amount of public IP addresses required in a networking environment, it comes in really handy when two companies that have duplicate internal addressing schemes merge. NAT is also a great tool for when an organization changes its ISP, but its SysAdmin wants to avoid changing the internal address scheme.

Types of Network Address Translation

There are three kinds of NAT:

Static NAT (one-to-one)

This type of NAT is designed to allow one-to-one mapping between local and global addresses. Keep in mind that the static version requires you to have one real IP address for every host on your network.

Dynamic NAT (many-to-many)

This version gives you the ability to map an unregistered IP address to a registered IP address from a pool of registered IP addresses. You don't have to statically configure your router to map each inside address to an individual outside address as you would using static NAT, but you do have to have enough bona fide IP addresses for everyone that's going to be sending and receiving packets from the Internet at the same time.

Overloading (one-to-many)

This is the most popular type of NAT configuration. Overloading is really a form of dynamic NAT that maps multiple unregistered IP addresses to a single registered IP address (many-to-one) by using different source ports. Why is this so special? It's also known as Port Address Translation (PAT), which is also commonly referred to as NAT Overload. Using PAT enables you to allow thousands of users to connect to the Internet using only one real global IP address! Seriously, NAT Overload is the real reason we haven't run out of valid IP addresses on the Internet.

NAT Names

The names used to describe the addresses used with NAT are easy to remember. Addresses used after NAT translations are called *global addresses*. These are usually the public addresses used on the Internet, which you clearly don't need if you aren't going on the Internet.

Local addresses are the ones used before NAT translation. The inside local address is actually the private address of the sending host that's attempting to get to the Internet. The outside local address would typically be the router

interface connected to your ISP and is also usually a public address used as the packet begins its journey.

After translation, the inside local address is then referred to as the inside global address. The outside global address then becomes the address of the destination host.

The following table lists all this terminology and offers a clear picture of the various names used with NAT. Keep in mind that these terms and their definitions vary a bit based on implementation and this table shows how they're used according to the Cisco exam objectives:

NAT terms

Names Meaning

Inside Source host inside address before translation—typically an RFC local 1918 address.

Outside Address of an outside host as it appears to the inside network. local This is usually the address of the router interface connected to

the ISP—the actual Internet address.

Inside Source host address used after translation to get onto the global Internet. This is also the actual Internet address.

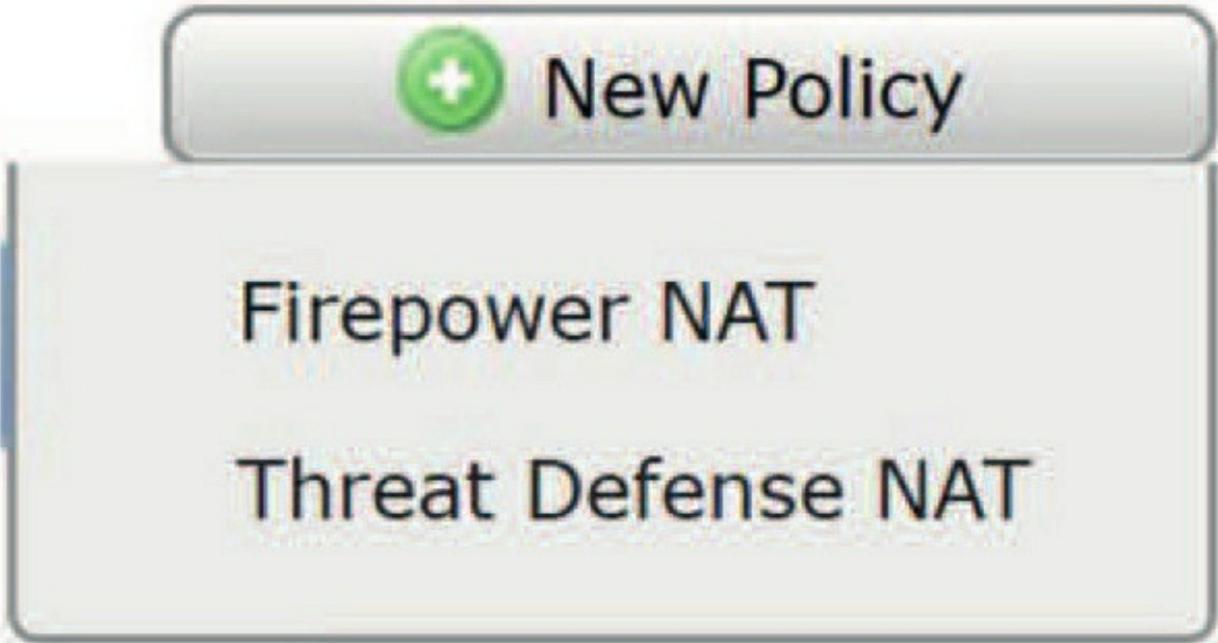
Names Meaning

Outside Address of outside destination host and, again, the real Internet global address.

Firepower NAT

So yes, Network Address Translation (NAT) is a great example of a feature that exists in both Firepower and FTD devices, but there are some key differences between platforms.

To configure NAT, you start with one of two kinds of NAT policy: Firepower and Threat Defense. First, navigate to Devices>NAT, click the New Policy button



in the upper right, and pick your policy type as shown in the next figure. We'll begin with Firepower NAT and then we'll take care of FTD.

When you create a Firepower NAT policy, you'll notice that you can only target Firepower devices with it, which makes sense, right? FTD devices don't show as options in the Available Devices area, only the appliance, as shown below.

New Policy

Name:

Description:

Targeted Devices

Select devices to which you want to apply this policy.

Available Devices

 Firepower_Appliance

Selected Devices

 Firepower_Appliance

But you don't absolutely have to target a device at this point; you can create

the policy and select the device(s) later if you want. Your new policy has no rules yet and the screen looks like so:



You've also got two types of NAT translations listed: Static and Dynamic. Clicking the Add Rule button takes you to the Add Rule dialog:

Add Rule

The screenshot shows the 'Add Rule' dialog box. At the top, there is a 'Name' input field with the placeholder text 'Enter a name...', an 'Enabled' checkbox, and a 'Type' dropdown menu. The 'Type' dropdown is currently set to 'Dynamic IP + Port' and is open, showing options: 'Static', 'Dynamic IP Only', and 'Dynamic IP + Port'. Below the 'Type' dropdown are four tabs: 'Zones', 'Source Network', 'Destination Network', and 'Destination Port'. The 'Zones' tab is selected and highlighted in blue. Under the 'Zones' tab, there is a section titled 'Available Zones' with a refresh icon and a search bar labeled 'Search by name'. To the right of the search bar, the word 'Source' is visible, and below it, the text 'any' is displayed.

The Add Rule dialog offers these options:

- **Name:** Your rule name; maximum of 30 characters.
- **Type:** The choices are Static, Dynamic IP Only, and Dynamic IP + Port.
- **Zones:** Optional source/destination security zones for this NAT rule.
- **Source Network:** This tab is blank if you're configuring a static NAT rule—otherwise, enter the original and translated source network or pick one of the available network objects.
- **Destination Network:** For a static NAT, you can enter both the original destination and the translated destination. For a dynamic NAT, the Translated Destination Network option is grayed out.
- **Destination Port:** For a static NAT, you can enter both the original and translated port. For a dynamic NAT, the Translated Port option is grayed out.

You can create different types of NAT translations using the various rule types and options, and you can either enter IP addresses/ CIDR blocks manually or use network objects.

Firepower NAT Examples

Let's go over a few NAT rule examples now. These aren't exhaustive by any means, but they should still give you a good idea of how to create different types of NAT rules.

Static NAT for a Single Host

To create a static rule, specify the original destination IP and the translated IP on the Destination Network tab. This will translate all traffic for the original host to the translated IP address.

Single NAT for a Single Host on Port 80

You have an internal host with a private IP address, and you want to run a web server on port 80. You don't want to expose the entire host to the big bad Internet, just the HTTP port. How would you get this done?

First, you'd create a static rule as described above. On the Destination Port tab, specify HTTP for the original and translated port. This will cause only traffic destined for port 80 on the original destination IP to be sent to your internal host.

Dynamic NAT Pool

What if you want to hide/change the original IP address for a block of IP address space? Sounds like a job for a dynamic NAT pool. To create this rule, you would pick the rule type Dynamic IP Only. Then on the Source Network tab, you'd specify your original source network and translated source network. These would both be CIDR blocks.

In the next figure, we're translating the 10.0.0.0/24 address space to 172.17.0.0/25:

Editing Rule - Dynamic Pool

? X

Name Enabled Type [Move](#)

Zones **Source Network** Destination Network Destination Port

Available Networks

- IPv4-Private-All-RFC1918
- any
- IPv4-Benchmark-Tests
- IPv4-Link-Local
- IPv4-Multicast
- IPv4-Private-10.0.0.0-8
- IPv4-Private-172.16.0.0-12
- IPv4-Private-192.168.0.0-16
- IPv6-to-IPv4-Relay-Anycast
- any-ipv4

Add to Original

Add to Translated

Original Source Network (1)

10.0.0/24

Translated Source Network (1)

172.17.0.0/25

Add

Add

OK

Cancel

This rule will work fine as long as no more than 126 hosts from the 10.0.0.0 network need to be translated. But the netmask on the 10.0.0.0 network is /24, which means there could be up to 254 hosts on that network!

Firepower will actually allow you to go ahead and create this rule, but if you click the Show Warnings button right there next to the Add Rule button, you'll see that the rule has a warning triangle next to it. Hovering over that triangle will display the following message:

Warning: There are not enough IP addresses specified for source network translation. Consider using "Dynamic IP + Port"

Firepower is just giving you a warning you that you could run out of IP addresses to translate to in your 172.17.0.0/25 network. **Dynamic IP + Port**

This rule type is generally used to allow internal hosts with private IP addresses to access the Internet. Because their addresses aren't routable outside the private network, they must be translated to a public address. Now, nobody but a few large corporations (you know who you are!) has enough public IP addresses to provide one for every computer, thus we have dynamic IP and port translation.

This approach is also sometimes called "Hide NAT" because it "hides" the internal IP addresses behind the external NAT address. We're going to translate a number of internal IP addresses to a single external IP. In addition to translating the source IP, we'll also dynamically translate the source port. This way, even if clients are using the same source port to connect externally, the device can map the connection to the right host based on the dynamic source port assigned on the external network.

This will be a Dynamic IP + Port rule. The Original Source Network field will be your internal address space—say 10.0.0.0/24. The Translated Source Network field will be your single public IP address, and this could also be a CIDR block or even multiple external IP addresses. You could also use zones to specify the source and destination of your traffic. Your source zone would be where your internal hosts reside, and the destination zone would be Internet facing.

Once you're done, your rule will look something like this one:

Rules

Show Warnings Add Rule

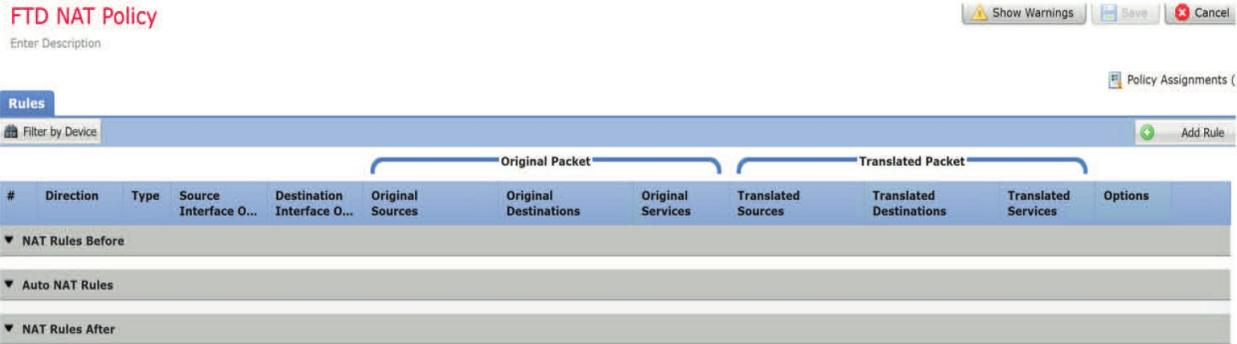
#	Name	Type	Source Zones	Dest Zones	Original Packet			Translated Packet			
					Source Networks	Dest Networks	Dest Ports	Source Networks	Dest Networks	Dest Ports	
Static Translations											
This category is empty											
Dynamic Translations											
1	Hide NAT	Dynamic IP + Port	any	any	10.0.0.0/24	any-ipv4	any	127.23.45.67	original	original	

Firepower Threat Defense (FTD) NAT

FTD includes the basic NAT features available in Firepower devices but it adds more options. You can use NAT on a device in either routed or transparent firewall mode, but I really don't recommend NAT with transparent mode. Also, you should know that you can't configure NAT for interfaces operating in inline, inline tap, or passive mode.

To create a new policy, navigate to Devices>NAT, click the New Policy button, and select Threat Defense NAT for the policy type. Give your policy a name and optionally select one or more FTD device targets.

Your new policy will look like the next figure:



You can see that this policy is similar but not quite the same as the Firepower NAT policy. One difference is in the rule order. In both Firepower and FTD, static NAT rules are evaluated first, but in Firepower, these go into their own rule section in the policy.

In the Firepower policy, static NAT rules are listed above dynamic rules. In FTD, however, manual static and dynamic rules are mixed together. This really doesn't matter because static NAT rules will still be evaluated before dynamic rules.

Looking at the next figure, the order of evaluation seems like it would be 1, 2, 3. But it's not because rules 1 and 3 are static, so the order they would be evaluated in would be 1, 3, 2.



Another interesting difference is that NAT rules don't have names. You can add a rule description but not a name, and the description is only visible when you're editing the rule.

Auto NAT or Manual NAT?

Another difference between Firepower and FTD is Auto NAT and manual NAT rules, also known as object-based NAT or twice NAT after the way each one is configured.

Auto NAT has some limitations, which is fine because it's just really intended to be an easy way to create a NAT rule. When creating an Auto NAT rule, your original source must be a network object, and you can't use network object groups. The object can be a single IP or CIDR block. You also can't configure a translated destination because it only works by translating the source address. Under the hood, the Auto NAT is actually a parameter for the network object, a lot like the Network Object NAT on an ASA. The difference is that you can't actually see the NAT in the object itself, but it's there!

You can use an Auto NAT rule to translate a single source to another source either statically or dynamically or to perform Port Address Translation (PAT). Keep in mind that if all you need to do is translate source IP addresses, then an Auto NAT is easier to use. But, if you need to use multiple CIDR blocks in your rule or specify destination addresses, you'll need to do a manual NAT rule.

FTD NAT Examples

Of course, you can perform the same NAT translations in FTD as you can on Firepower devices, so I'm not going to repeat all the different NAT types here again.

The only one that's significantly different is that Dynamic IP and Port, or "Hide NAT," rule, which isn't NAT; it's actually PAT (Port Address Translation). FTD gets this right by calling the rule what it really is. To create a Hide NAT rule, follow these steps:

Select Dynamic for the rule type.

1. On the Translation tab, enter your internal address space as the original source. Leave the Translated Source field blank.
2. On the PAT Pool tab, check the Enable PAT Pool box.
3. In the Address field, enter your external (public) IP address. You have more options here for

the port selection, like Round Robin, Extended PAT Table, and Flat Port Range. So, here's a simple example of using an inside local address to an inside global and overloading the translated packets. I can do this either with dynamic or static, but I'll just use a simple static mapping since static mappings are always evaluated before dynamic NAT rules. In the figure below, understand that the type is static, and that I used two objects: 19-Inside and 19-Outside. Objects are mandatory in Firepower NAT.

19-Inside	10.19.119.20
19-Outside	10.11.12.191

I'm just referencing a host with the inside address, but you can easily make the variable used in a rule a subnet, which would usually be the case. The outside address is just an unused address on the outside subnet that'll be automatically overloaded.

I'll then add these two objects into the rule used for this inside network to allow them to traverse the router out to the Internet. Other than setting my zones in the Interface Objects tab, there are no other configs once you've set your objects correctly in the rule.

NAT Rule: Auto NAT Rule

Type: Static Enable

Interface Objects Translation PAT Pool Advanced

Original Packet

Original Source:* 19-Inside

Original Port: TCP

Translated Packet

Translated Source: Address

19-Outside

Translated Port:

Finally, notice that Auto NAT rules are in-between the twice-NAT before and after rules and would always be evaluated first:

#	Direction	Type	Source Interface O...	Destination Interface O...	Original Sources	Original Destinations	Original Services	Translated Sources	Translated Destinations	Translated Services	Options
▼ NAT Rules Before											
▼ Auto NAT Rules											
#	↔	Static	Inside	Outside	19-Inside			19-Outside			Ons:false
▼ NAT Rules After											

Testing and Troubleshooting NAT

Cisco's Firepower NAT gives you some serious power even without a whole lot of effort because the configs are really pretty simple. Even so, nothing's perfect, so when something goes wrong, you can ferret out some of the more common culprits by running through this list of potential causes:

- Check the dynamic pools. Are they composed of the right scope of addresses?
- Check to see if any dynamic pools overlap.
- Check to see if the addresses used for static mapping and those in the dynamic pools overlap.
- Make sure there aren't any addresses left out that need to be there and ensure that none are included that shouldn't be.
- Check to make sure you've got both the inside and outside zones delimited properly.

One of the most common problems with a new NAT configuration is often

not specific to NAT at all—it usually involves a routing blooper. So, because you’re changing a source or destination address in a packet, make sure your router still knows what to do with the new address after the translation!

With that, we can perform the verification once our NAT is configured on our Firepower appliance and/or FTD devices. Remember this needs to be done at the LINA CLI level.

Obviously, the best way to test NAT is from an inside host to verify connectivity, but what if that particular host doesn’t work? Let’s take a look at our options.

NAT Verification

Sadly, there isn’t nearly the amount of CLI commands that we’d find on a router/switch IOS CLI output for verifying NAT, but we still have some of our old ASA commands available.

First, the `show nat translation` doesn’t parse here, but the `show xlate` does, and it’s the same output as the ASA. I only have the one inside host, but you can easily see the mapping here:

```
> show xlate
1 in use, 1 most used
Flags: D - DNS, e - extended, I - identity, i - dynamic,
r - portmap, s - static, T - twice, N - net-to-net NAT from Inside:10.19.119.20 to Outside:10.11.12.191
flags s idle 0:00:00 timeout 0:00:00
```

There are some good options with the `show xlate` command as you can see here:

```
> show xlate ?
count Show translation count
detail Enter this keyword to display xlate details global Enter this keyword to specify global ip range
gport Enter this keyword to specify global port(s) id Enter this keyword to specify xlate id
interface Enter this keyword to specify an interface local Enter this keyword to specify local ip range
lport Enter this keyword to specify local port(s) type Enter this keyword to specify xlate type
| Output modifiers
<cr>
```

`Show nat` isn’t a command on the Firepower devices, but you can use this command with its many options:

> **show nat ?**

A.B.C.D Match original IP address
detail Expand object/ object group
divert-table Display Nat Divert table
interface Specify an original interface
object Specify an original network/service object object-group Specify an original network object-
group pool Display NAT/ PAT pool usage
proxy-arp Display proxy ARP table
translated Specify the translated parameters
| Output modifiers
<cr>

Here's an example of the `show nat` command and the translation on a particular interface:

> **show nat translated interface outside**

Auto NAT Policies (Section 2)
1 (Inside) to (Outside) source static 19-Inside 19-Outside translate_hits = 121033, untranslate_hits = 4961

Notice that I can see the same information with the `show nat interface` *option* command:

> **show nat interface Inside**

Auto NAT Policies (Section 2)
1 (Inside) to (Outside) source static 19-Inside 19-Outside

translate_hits = 120565, untranslate_hits = 4959

Another command that I really like that's better than the last one ^{is} `show nat detail`:

> **show nat detail**

Auto NAT Policies (Section 2)
1 (Inside) to (Outside) source static 19-Inside 19-Outside

translate_hits = 120565, untranslate_hits = 4959
Source - Origin: 10.19.119.20/32, Translated: 10.11.12.191/32

Last, I like using the `show running-config` command with some options to verify my configuration too. The *Objects* and *NAT* options provide just what I need:

> **show running-config object**

object network 19-Inside
host 10.19.119.20
object network 19-Outside
host 10.11.12.191

```
> show running-config nat
!  
object network 19-Inside  
  
nat (Inside,Outside) static 19-Outside
```

Supposedly the sky's the limit regarding the number of mappings the NAT table can hold. But this being reality, things like memory and CPU, or even the boundaries set in place by the scope of available addresses or ports, can limit the actual number of entries. Oh... And keep in mind that each NAT mapping devours about 160 bytes of memory.

Summary

In this chapter you learned about static NAT, dynamic NAT, and Port Address Translation (PAT) also known as NAT Overload, and we covered some key terms as well. You also learned about the differences about NAT on the Firepower Appliance and the FTD sensor.

We moved into demonstrating configurations and wrapped up the chapter with verification in the LINA, using the CLI.

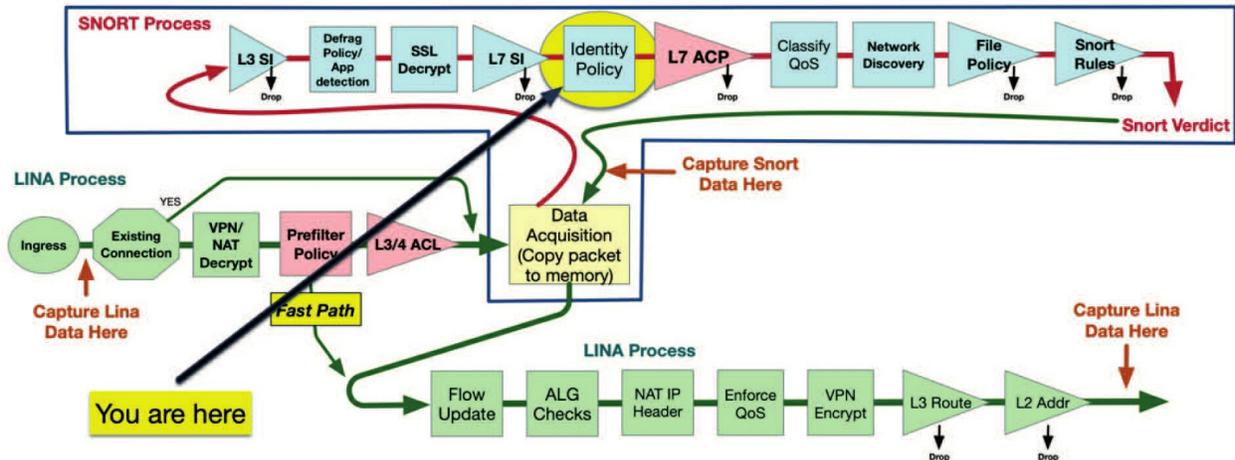
Chapter 16: Identity Policy

The following CCNP Security SNCF exam objectives are covered in this chapter:

2.0 Configuration

2.2 Configure these policies in Cisco Firepower Management Center

2.2.e Identity



This chapter covers identity, which translates to the ability of Firepower to take different actions depending on the user associated with a connection.

To many an organization's management personnel, the idea of using Firepower to block and allow traffic based upon users or groups is absolutely riveting. Why wouldn't it be? It sounds like an excellent way to keep employees productive by allowing or blocking access to websites and services based on job function, and it is. But the actual implementation of user-based control is way more complex than that, as you'll soon see.

I'll start off the chapter by defining identity sources. After that, we'll configure a realm, and then build an actual Identity policy.

To find exam study material such as hands-on labs access, videos, downloadable supplemental material, and practice questions, head over to www.lammle.com/firepower.

Identity with Firepower

The basic goal behind identity with Firepower is for it to enable us to tailor the actions taken on the network regarding the user. It allows you as an administrator to control which users or groups are allowed or denied access to specific network services or applications.

To make this happen, some kind of mapping has to occur between the traffic passing through the device and the user who's sending it, which in this case,

manifests in associating an IP address with a user. There are several ways to get this done, and the method you go with will depend on several factors: the type of directory services you use, the location of your devices, and interface types.

Identity Sources

The first decision to make is which identity sources you'll use. There are two, broad categories, non-authoritative and authoritative. Because there's only one non-authoritative-based form of detection in Firepower, I'll cover that here. I'm only going to briefly touch upon various authoritative sources since these we'll dig deep into these in the next book.

Non-Authoritative (Traffic-Based Detection)

So the single, non-authoritative identity source available in Firepower is traffic-based detection. It's called non-authoritative because the user information is gained via passive detection from a bunch of supported applications and protocols. These are all configured in the Network Discovery policy as I described in Chapter 11: "Firepower Network Discovery." Remember this list of supported protocols:

- LDAP • FTP
- AIM • POP3
- IMAP • HTTP
- Oracle • MDNS
- SIP (VoIP) • SMTP

Both successful and failed login attempts can be recorded too, and in fact, traffic-based detection provides user information based on logins using the above protocols. This information is then saved in the FMC database. As events occur, Firepower associates this user information, providing us with another data point when investigating a possible incident.

The username is pulled directly from logins in the network traffic. Additional information on the username can also be stored by querying a Lightweight Directory Access Protocol (LDAP) directory.

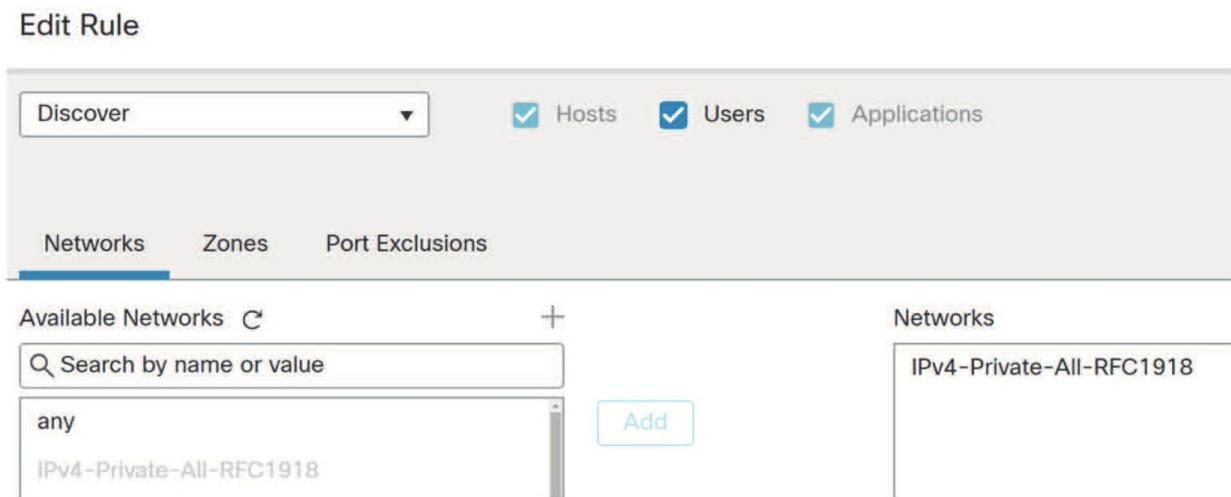
When a device detects a login using one of the supported services or applications, it sends the following information to the FMC:

- The username identified
- Time of the login
- IP address involved
- For POP3, IMAP, and SMTP, the user's email address
- The name of the device detecting the login

If the user was previously detected, the user's login history will be updated with the new event; if it's a previously undetected user, the FMC adds the user to the users database.

To configure traffic-based detection, you've got to enable user discovery in your discovery policy rule or rules. You can find this on the FMC under **Policies>Network Discovery**.

Check your network discovery rules and ensure that both the Hosts and Users boxes are selected as I showed you in chapter 11, Firepower Network Discovery, and as shown in the following figure.



After editing your discovery rules, go back to the discovery policy and click the Users tab. Configure the services and applications where you will be performing this detection—by default these are all enabled. Also, decide if you want to capture failed logging attempts.

Click the pencil icon to enable or disable the various settings. Here's the screen showing the Users tab:

Networks

Users

Advanced

Traffic-Based Detection



aim	Yes
imap	Yes
ldap	Yes
oracle	Yes
pop3	Yes
sip	Yes
ftp	Yes
http	Yes
mdns	Yes
Capture Failed Login Attempts	Yes

That's all there is to it! Most deployments leverage this traffic-based user detection even if they don't use the more complex authoritative sources. Even so, you need to remember that you can't customize your access control rules based upon user data gathered solely with this method. It's really most often employed as an additional data point when analyzing or investigating events.

Authoritative Sources

Authoritative sources can be used to positively associate a username with an IP address. You must use an authoritative source if you want to use user-based access control rules. This method involves periodically updating the FMC with a list of known users and groups, which can then be used when creating access control rules. You then use one of the supported authoritative sources to associate an IP address with a username. Access control rules can be selectively enforced based on this user-to-IP mapping approach.

Two of the sources we're going to cover are the User Agent and Identity Services Engine (ISE). As we're talking about these, one really important thing to remember is that you *can't* use both the User Agent and ISE sources simultaneously. You must choose one of them for your deployment.

The Firepower User Agent

Formerly called the Sourcefire User Agent (SFUA), the Firepower User Agent is one type of authoritative source. This agent reads the login events from a Microsoft Active Directory server and forwards them to the FMC. The agent is actually a Windows service that can be installed locally on an Active Directory server or on a separate Windows computer. When installed on a separate Windows machine, the agent can query up to five Active Directory servers and update up to five Firepower FMCs. You'll definitely want to cover all your Active Directory servers with either local or remote agents. Doing this ensures that a user login will be reported regardless of the server where they're authenticated.

To set up a User Agent connection, go to **System>Integration>Identity Sources** on the FMC. You'll see three choices for identity sources on this page:

Identity Sources

Service Type

Support for Cisco Firepower User Agent is deprecated and will be removed in a future release

Host Name/IP Address
10.11.11.250

None, Identity Services Engine, and User Agent:

You can see that the User Agent is being deprecated because ISE with pxGrid is better. But for now, I'm still using it in this example, and it's still a valid agent to use.

To add a new User Agent, I'll click the **New Agent** button and then be prompted for the hostname or IP address of the agent, as shown below. This is the Windows computer or directory server where I'll install the User Agent software.

User Agent



Host Name/IP Address

Cancel

Add

After adding all of my User Agents here, I'll click the Save button in the upper right, which tells the FMC that I'll be using this method of authoritative user identification.

The next step is to download the User Agent installer from the Cisco support site and install it on my Windows computers. During the installation process, I'll point the agent to my FMC(s).

Later, when my identity configuration is complete, these User Agents will constantly update the FMC with information about which users are logged into which computers. The FMC will, in turn, update the devices where user-based access control rules can be implemented.

This is some powerful information to use in your network analysis! **Identity Services Engine**

The Identity Services Engine (ISE) is another authoritative source of user login information. When an ISE identity source is configured, user information can be updated on the FMC in a manner similar to how it's

updated to the User Agent. ISE server provides the same user-to-IP mapping as the User Agent.

What's more, ISE can use Security Group Tags (SGTs) to tag packets as they traverse the network. These SGTs can be used as conditions in access control rules. So, with ISE we have two methods of controlling access: through the user/group and/or through the SGT in the packet.

To configure an ISE server connection, go to **System > Integration > Identity Sources** on the FMC. Then click the Identity Services Engine button, which displays the ISE settings:

Identity Sources

Service Type

Primary Host Name/IP Address

Secondary Host Name/IP Address

pxGrid Server CA +

MNT Server CA +

FMC Server Certificate +

ISE Network Filter (Only applicable for Session Directory Topic)

Subscribe To:

Session Directory Topic

SXP Topic

Required Field

I'll need the necessary certificates to enable communication with the ISE server. Again, we go deep in the next book—it's all ISE with pxGrid!

The Terminal Services Agent

The Terminal Services Agent assigns port ranges to each user for various applications, thereby allowing the FMC to identify users by IP and source port. Contact your Cisco account management representative if you'd like access to this feature. It'll be included in a future Firepower version.

So even though the Terminal Services Agent is mentioned in the FMC documentation, you've got to have at least Firepower 6.2 at a minimum, or better (6.2.1), to run this. This agent is required if you want to perform user identification in a terminal services environment because a simple IP-to-user mapping won't work when multiple users log in to a single host.

Captive Portal Authentication (Active Authentication)

A Captive portal is another authoritative identity source that can be leveraged in Firepower. It requires users to authenticate to the network through a managed device, but optionally, it can also allow guest access. Active authentication is performed only on HTTP and HTTPS traffic, and to perform active authentication on HTTPS traffic, you must use an SSL policy to decrypt the traffic from the users you want to authenticate. If this sounds like a lot of work, you're right, it is

Your Identity policy can be configured to first look for passive authentication via the User Agent or ISE. It can then redirect the user to a captive portal authentication page if there's no user associated with the traffic.

To perform captive portal authentication, your device must serve up the portal page to your users. Because of this, you can't use a passive or transparent device since they do not have IP addresses assigned to their sensing interfaces.

Realms

If you're going to do identity right, you need a realm. A realm consists of one or more LDAP or Microsoft Active Directory servers. You must configure a realm if you want to leverage user-based control or configure an authoritative identity source.

In fact, you cannot configure an Identity policy until you have configured one or more realms!

Think of a realm as the connection to your user directory. There's very little that can be done without this connection.

All you can really do is glean a few usernames via Firepower's passive discovery of login traffic.

Add Realm

So, to create a realm, navigate to **System>Integration>Realms**. When you click the New Realm button, you'll be greeted by the screen shown below.

Here you'll enter the realm-specific information like type (AD or LDAP), domain, username, password, and distinguished names (DNs).

Identity Sources

Service Type

Primary Host Name/IP Address

Secondary Host Name/IP Address

pxGrid Server CA +

MNT Server CA +

FMC Server Certificate +

ISE Network Filter (Only applicable for Session Directory Topic)

Subscribe To:

Session Directory Topic

SXP Topic

Required Field

Did you notice there's no information on the hostname or IP address of the directory server? That's because you can have multiple directory servers in a realm as long as they host the same directory.

The AD Join Username and Password fields are only needed if you'll be configuring Kerberos captive portal active authentication and they don't work otherwise. The Test button is grayed out unless you fill in these two fields.

Add Directory

After adding your realm, you'll be taken to the Directory tab. This is where you point to

your actual directory servers.
Click the Add Directory button
on the right to add a new entry,
which brings up the page in the
following figure:

Fill out the hostname/IP
and click the Test button. If
you can't reach the server, or

Add directory ?

Hostname / IP Address

Port

Encryption

STARTTLS

LDAPS

None

SSL Certificate
 +

the server isn't listening on 389 (or, most likely, your configuration is wrong), then this will fail.

You can repeat by adding as many directory servers as your heart desires. If there are multiple directory server entries, Firepower will use the top entry. If the server is down, it'll proceed down the list until it finds a responsive server or runs out of entries. You can drag servers up and down this list to change their priority.

User Download

After adding and activating the realm, you can stop right there if you want, but doing that won't accomplish much. Firepower will begin looking up usernames discovered through passive traffic analysis in your directory, and if available, additional information like first name, last name, email, department, and phone number will be populated in the FMC user database. But again, you really don't want to just stop with the realm, so to configure user-based control in your Access Control policy, you'll have to proceed with user download.

Now after adding the directory as I mentioned earlier, click the User Download tab. As soon as you do this, Firepower will attempt to connect to the directory and enumerate the groups available.

If there's a problem with your realm or directory configuration, you'll find out right away because your user download will fail! But if everything was entered correctly, you'll see a screen similar to this one:

Directory Realm Configuration **User Download**

Download users and groups

Begin automatic download at America/New York Repeat Every

[Download Now](#)

Available Groups

- Remote Management Users**
- Print Operators
- Server Operators
- Performance Monitor Users
- DnsUpdateProxy
- RAS and IAS Servers
- Replicator
- Enterprise Admins

|< < Viewing 1-50 of 50 > >|

On this page, the available groups are shown in the left column. By default, no groups are added to the **Groups to Include** column because you probably don't want to pull down all the users/groups in your entire directory to the FMC.

NOTE: You can add up to 3000 groups, with each group having up to 5000 members!

Okay—now to populate users on your FMC, select the groups on the left and click the Add to Include button. If you have larger-sized groups and you want to exclude some sub-groups, you can add these to the Groups to Exclude column.

By default, the users will be updated on your FMC every 24 hours, but you can modify the start time and interval for this update with the settings at the top of the page.

Of course, when you're finished, click the Save button in the upper right,

which will send you to your Realms list. The final and critical step is to enable the realm by clicking the slider in the State column. Once your realm is enabled, Firepower will use it for user identity and optionally download users at the interval specified.

To have the AD groups show up in your ACP rules, you must pull down the groups manually at least twice for some reason.

Identity Policy

Okay finally, after setting up your realm, directory, and authoritative user sources, you all set to proceed with the Identity policy itself! The Identity policy brings together all the various components and associates' traffic on your network with an authoritative identity source and a realm. Remember that you absolutely must have already configured at least one realm before the system will allow you to create an Identity policy.

You'll find the Identity policy on the FMC under **Policies>Access Control>Identity**.

Clicking New Policy brings up the dialog in the next figure, where you simply give your policy a name and optional description.

Your new policy contains no rules by default. To add an authentication rule, click the Add Rule button, which will bring up the page shown here:

New Identity policy



Name

Lammle-Identity

Description

Used in my CCNP book!

Cancel

Save

So first, you need to give your rule a name. After that, you can use the four tabs on the left to configure the traffic that'll match this rule. By default, the rule will match all traffic.

The Realm & Settings tab on the right determines how to authenticate users for traffic matching the rule. If you click this tab, you can select the realm you want to use to authenticate traffic matching the rule. Check out the following figure:

Name: Passive Rule Enabled Insert: into Category Standard Rules

Passive Authentication **Realm:** AD (AD) Authentication Protocol: HTTP Basic Exclude HTTP User-Agents: None

Zones Networks VLAN Tags Ports **Realm & Settings**

Realm * AD (AD)

Use active authentication if passive or VPN identity cannot be established

To enable passive authentication using a realm, simply select the realm from the drop-down list. The configured identity source and realm will then be used—either the User Agent or ISE.

Add Rule

Name: Passive Rule Enabled Insert: Into Category Standard Rules

Passive Authentication **Realm:** No realm **Authentication Protocol:** HTTP Basic **Exclude HTTP User-Agents:** None

Zones Networks VLAN Tags Ports

Available Zones

- External
- Firepower_Appliance
- Inside
- Internal
- Outside
- VRF1-Inside
- VRF1-Outside
- VRF2-Inside

Add to Source Add to Destination

Source Zones (0) any

Destination any

Now if passive authentication fails, meaning there's no user associated with the IP address, then clearly the traffic will not be associated with a user and any user-based access control rules would then fail to match the traffic.

It's just really important to understand the preceding statement when configuring user-based control. In order to allow or block connections based on user identity, the traffic must be associated with a user account before it's processed through the Access Control (AC) policy. If an AC rule contains a user or group criteria but the traffic is not from a known user, the rule will not match the traffic!

But there actually is a situation where this is acceptable... If you only want to allow a specific group of users access to a given application or website, they must be authenticated to gain that access. Any other unauthenticated traffic would be blocked by another AC rule. Still, if you want to authenticate *all* traffic passing through your device, then you may need to actively authenticate unknown connections.

To augment your passive authentication with active authentication, check the

“Use active authentication if passive or VPN identity cannot be established”
check box. When you check this box, you’ll see a bunch of additional
options:

Add Rule

Name: Enabled

Insert into Category:

Passive Authentication Realm: *No realm* Authentication Protocol: *HTTP Basic* Exclude HTTP User-Agents: *None*

Zones Networks VLAN Tags Ports

Available Zones

-
- External
 - Firepower_Appliance
 - Inside
 - Internal
 - Outside
 - VRF1-Inside
 - VRF1-Outside
 - VRF2-Inside

Add to Source

Add to Destination

Source Zones (0)

- any

Destination

- any

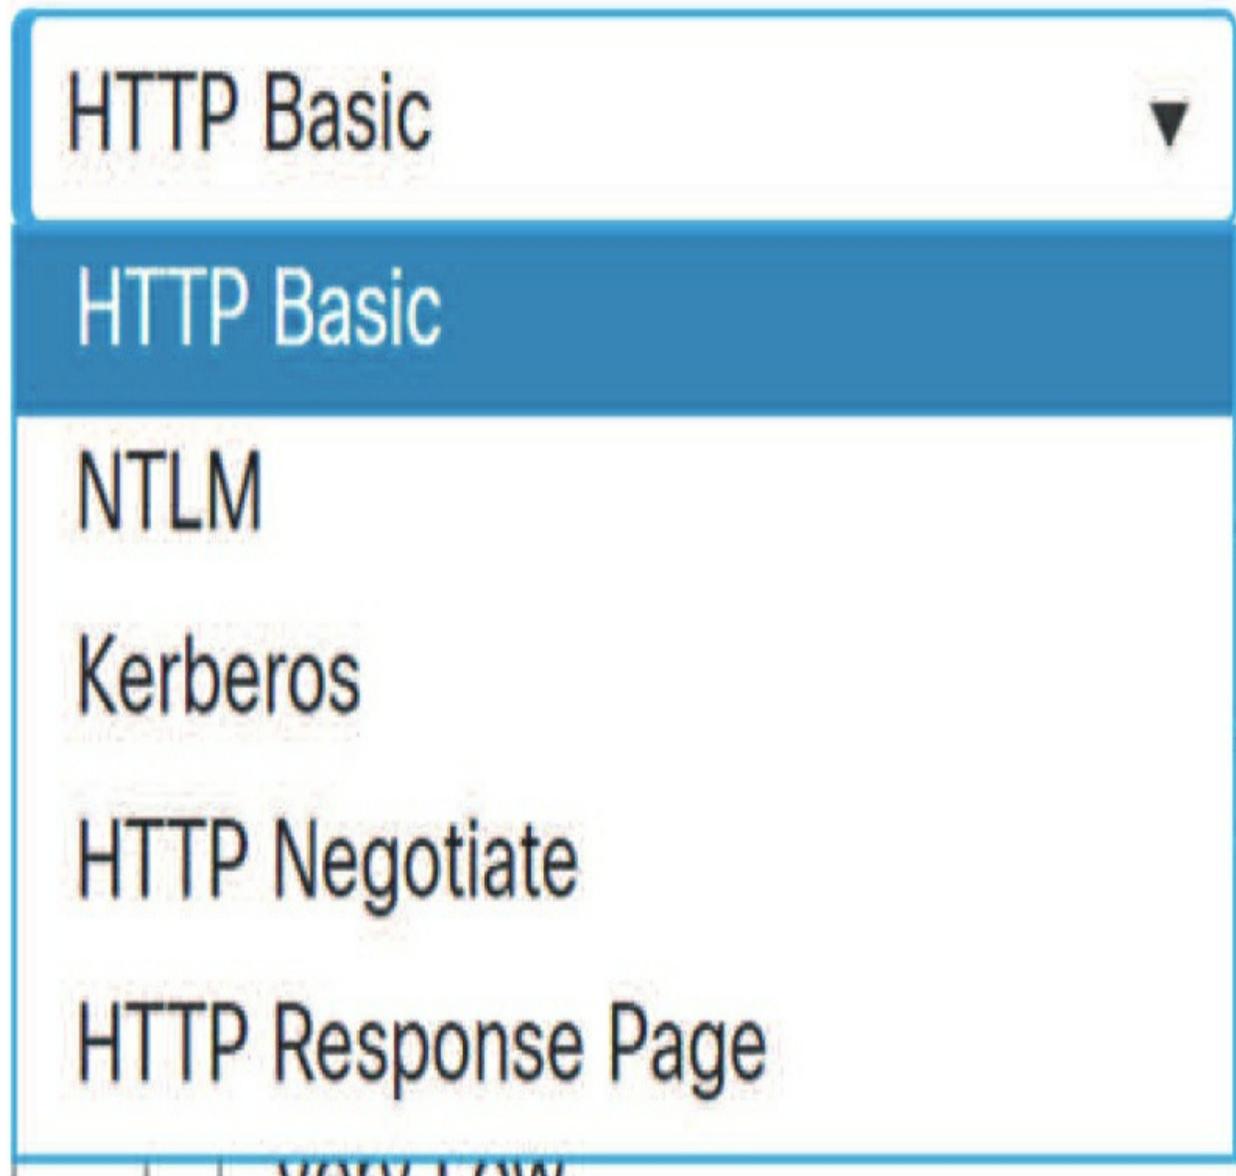
Let's explore some of the options on this page now.

There's also a check box to automatically identify the connection as guest if authentication doesn't successfully identify the user. This is fairly self-explanatory and there are several special identities you can choose in your AC rule matching criteria:

- Failed Authentication
- Guest
- No Authentication Required
- Unknown

There's an Authentication Type drop-down that contains several options, which I'll explain in the list below the figure:

Authentication Protocol



- HTTP Basic: uses standard fields in the HTTP header of each HTTP request.
- NTLM: Requires a browser that supports this type of authentication.
- Kerberos: The same as NTLM—it must be enabled in the browser.
- HTTP Negotiate: Yet another method that can be used on supported browsers. It allows the browser to choose between HTTP Basic, Kerberos,

and NTLM.

- HTTP Response Page: Unlike the preceding options, this is actually a web page served by the Firepower device. Instead of using the credentials from the browser/operating system, this page prompts for a username and password.

The remaining section of the rule is to allow for excluding traffic from active authentication, which is carried out based upon the User Agent that's being used by the client.

Under the Available Applications column, you'll find a number of applications identified by their User Agent that allows these applications to function instead of redirecting them to an authentication page that'll probably break them.

One example is the Advanced Packaging Tool that's used by some Linux distributions, including Ubuntu's `apt-get` command. Beware that if you use this feature, it's actually pretty easy for some mischievous user to spoof a User Agent and get by your AC rules!

If you enable active authentication and save your rule, you'll run straight into this warning:

Error

Please provide server certificate, redirect port, and maximum login attempts in Active Authentication tab when using Active Authentication

The reason why is that we haven't configured the Active Authentication tab yet. Clicking that tab will bring up the page in the following figure:

Lammle-Identity
Used in my CCNP book!

Rules **Active Authentication**

Server Certificate * +

Port * (885 or 1025 - 65535)

Maximum login attempts * (0 or greater. Use 0 to indicate unlimited login attempts)

Active Authentication Response Page

This page will be displayed if a user triggers an identity rule with HTTP Response Page as the Authentication Type.

🔍

* Required when using Active Authentication

As you can see, there's one red asterisk item that hasn't been configured yet—Server Certificate, and we've got to have a server certificate because the authentication mechanism used is going to be SSL encrypted.

And just so you know, there aren't any certificates included with Firepower by default. You can create your own certificate authority (CA) and server certificate, but it'll be a self-signed certificate, which will alert users that there's a problem. Implementing authentication this way is a just bad idea because you'd basically be training users to ignore these SSL errors. So what you really need is a server certificate that has been generated or signed by your own CA and trusted by your users' web browsers. This way, your devices can serve up this secure web page without forcing users to deal with certificate errors—nice!

To add the certificate used here, go to Objects>PKI> Internal Certs and add the certificate. Or—even easier—just click the green



plus icon () to the right of the Server Certificate drop-down to get sent directly to the Add Known Internal Certificate dialog.

This is a good place to point out that there are two more rule authentication types when you're adding a new rule. The three rule types available are revealed in



the following figure:

- **Passive Authentication:** This is the one we covered in detail earlier.
- **Active Authentication:** Provides for active authentication instead of passive. The active methods are the same as we discussed earlier.
- **No Authentication:** This rule simply disables the Realm & Settings tab. Any traffic matching these rules is exempt from authentication.

Keep in mind that even though identity rules allow you to select UDP ports for active authentication rules, the system can't enforce active authentication on non-TCP traffic. For that, you have to go with Passive Authentication or No Authentication rules.

Implement your Identity Policy

Just as it is with most other policies, once you've got your Identity policy configured, it's time to implement it via the Access Control policy.

Each Access Control policy can only be associated with a single Identity policy, but you can use the same Identity policy with multiple Access Control policies. Doing this reveals users and groups in the Users tab when creating a rule—nice! So let's do that now.

Go to your Access Control policy and add the Identity policy you've created to the Advanced tab in the ACP following these steps (the figure is shown on the next page):

1. First, go to **Policies>Access Control>Access Control**. 2. Next, go into your ACP to the Advanced tab and add your Identity policy. No need to deploy!
3. Now, create an ACP rule, open the Users tab, and verify that your users and groups are in there. If so, great! Your passive identity is working.
4. If you're using the ACP we created at the end of Chapter 9, go to rule 6, remove the Marketing VLAN object, and add the AD Marketing group.

Rules Security Intelligence HTTP Responses Logging **Advanced Settings**

Identity Policy Settings

Identity Policy



Lammie-Identity

SSL Policy Settings

SSL Policy to use for inspecting encrypted connections

SSL

Prefilter Policy Settings

Prefilter Policy used before access control

Lammie Prefilter Policy

Network Analysis and Intrusion Policies

Intrusion Policy used before Access Control rule is determined

Balanced Security and Connectivity

Intrusion Policy Variable Set

Default-Set

Default Network Analysis Policy

Balanced Security and Connectivity

The AD groups available in the ACP rule are displayed in the next figure.

Name

Allow Marketing to SM

Enabled

[Move](#)

Action

 Allow



Time Range



Zones

Networks

VLAN Tags

Users

Applications

Ports

URLs

SGT/ISE Attr

Available Realms 

Special Identities

AD

Available Users 

AD/*

Access Control Assistance Operators

Account Operators

Administrators

Allowed RODC Password Replication Group

Backup Operators

Cert Publishers

This figure reveals an edit of rule 6 in the ACP created at the end of Chapter 9. I added a VLAN object at that time, but now I've replaced that with an AD group.

The AD groups showed up in my ACP rule because I downloaded the users/groups twice when I was in the realm configuration. Do you remember what I am talking about here?

And yep, the next figure confirms that my rule has been updated with the AD group.



Corp URL Rules (6-8)												
6	Allow Marketing to SI	Inside	Outside	Any	Any	Any	AD/Marketing	Any	Any	Any	Social Networkin	Any
7	Block SM	Inside	Outside	Any	Any	Any	Any	Any	Any	Any	Social Networkin	Any

I want to point out that it's actually not relevant that we built the ACP example in Chapter 9... I just want to make sure you understand what I did here. And don't forget that AD groups are always better than a VLAN object.

Also, do you remember way back in Chapter 1 where I talked about this exact type of rule when the FMC goes down? If you can't remember what will happen to this rule if you lose your FMC, go back to Chapter 1 and review what happens, because it's really important that you understand this!

Summary

The Identity policy is an important policy to have in your quiver of tools because you need it to build a solid Access Control policy.

In this chapter, you learned that with the Identity policy, you can add both passive and active authentication as well as a realm, which allows you to use AD users and groups in an ACP.

Chapter 17: User Management

The following CCNP Security SNCF exam objectives are covered in this chapter:

None. This is a foundational bonus chapter! In this chapter, we're going to cover a variety of administrative functions for user account management. As we progress, you'll learn all about internal and external user account management, how to describe various user roles, and how to create custom user roles.

I'll demonstrate how to configure both internal and external user accounts and show you how to permit a user to escalate their account privileges.

You'll discover that user authentication can be achieved locally through the internal database or via an external authentication server like LDAP or RADIUS.

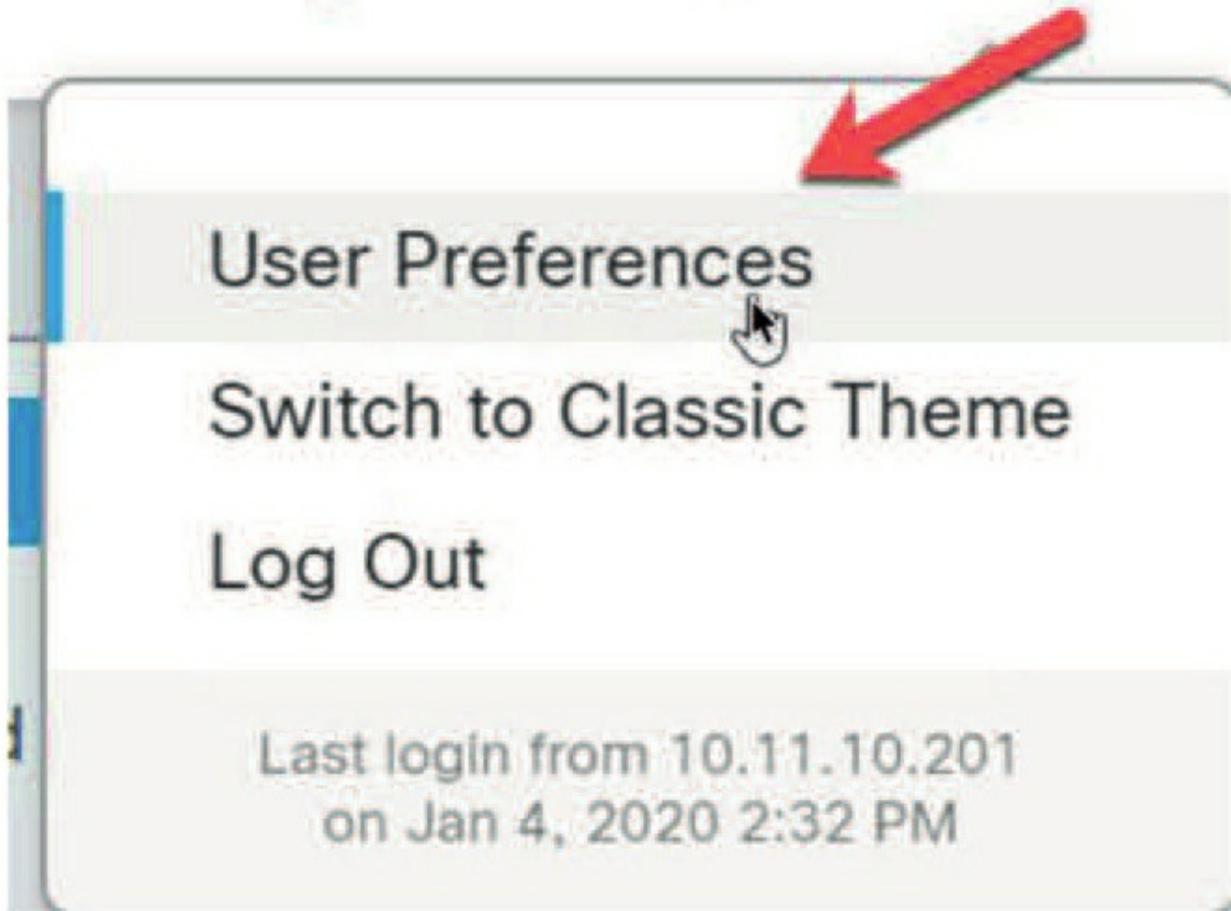
To find exam study material such as videos, downloadable supplemental material, and practice questions, please go to www.lammle.com/firepower.

User Preferences

We're going to get started with a quick overview of user preferences for the admin users on the Firepower system. This is done on a peruser basis; you can't set these commands globally for all users.

From the username on the very right side of the FMC, you can choose User Preferences, Switch Themes, or Log Out as pictured here. I'm going with User Preferences here.

ploy    admin ▼



From the User Preferences General screen, you can change the theme, set your time zone so the network analysis time stamps match your location, and also change your password.

UI Theme

Light

Please provide any Light theme feedback to fmc-light-theme@cisco.com

Time Zone

America/New York

Current time: Jan 4, 2020 4:51 PM

Change Password

Change Password

Although there are some other obvious settings in each tab, I want to tell you about a couple of very useful ones in the Event View Settings tab, shown below.

General Home Page **Event View Settings** Dashboard Settings How-To Settings

Event Preferences

Confirm 'All' Actions Confirm actions that affect all events

Resolve IP Addresses Resolve all IP addresses if possible

Expand Packet View

Rows Per Page Max 1000

Refresh Interval minutes (Set)

Statistics Refresh Interval minutes (Set)

Deactivate Rules

File Preferences

Confirm 'Download File' Actions

Zip File Password Leave blank

Show Zip File Password

It's pretty common to choose to resolve all IP addresses if possible because it comes in really handy in your Connection Events output. I typically add at least 100 in the Rows Per Page box, but you really don't want to set that value too high or your pages can time out!

I'm also going to set the Refresh Interval option to 2 or 3 minutes, which is helpful to ensure you are viewing updated information. After that, I'm free to set a password on any malware I want to download from the FMC. Remember back when we talked about the File & Malware policy in Chapter 10, where I demonstrated how to download malware when the devices detect a file with malware disposition? The default password is "infected."

It's not shown here, but further down that same page are the Default Time Windows options that you can configure to help you get to the pages that help quickly and efficiently.

Note: Remember that all of these settings are set per user—you can't set

these globally.

There's only one place to set some global configurations for users, and that's found at **System>Configuration>User Configuration**. Okay great—let's configure some users now!

Password Reuse Limit

Limit (0 = no limit)



Track Successful Logins

Days (0 = no tracking, 365 = max value)



Max Number of Login Failures

Maximum number of failures before temporary lockout (0 = no lockout, 999 = max value)



Set Time in Minutes to Temporarily Lockout Users

Minutes (0 = no wait time, 1440 = max value, 1 day)



Max Concurrent Sessions Allowed

Maximum sessions for users with Read Only privileges (0 = no limit, 1024 = max value)



Maximum sessions for users with Read/Write privileges/CLI users (0 = no limit, 1024 = max value)



Account Management

When you log in to the FMC's GUI interface with a username and password, it looks for the user ID in the local database first. Users in the local database will predictably be authenticated locally on the appliance, but even if the user account that's asking for access requires external authentication, the system will still try the local database first.

Now if the Firepower system can't find a local user that corresponds, it'll move on and look into an external server for a match. These external servers can either be a Lightweight Directory Access Protocol (LDAP) directory server or a Remote Authentication Dial In User Service (RADIUS) authentication server. Both types can provide a database list of users.

Internal vs. External User Authentication

By default, the appliance will go with the internal database for authentication to check for users' credentials when they log in. Although internal authentication is by far the simplest way to manage users, it's horribly inefficient to the point of not even being a viable option for really large networks. It's this fact that makes the ability to specify whether a given user you're creating will be internally or externally authenticated such a cool feature!

Still, understand that when you're using internal authentication, all user credentials are managed in the internal database and local authentication won't take place until the username and password are verified against the internal FireSIGHT System database. Since you manually create each user in the local database to be authenticated there, you set the user role and access setting for that user at that time as well. This means you don't need to deal with any default settings for users attempting to authenticate.

External authentication occurs when the Firepower Management Center (FMC) or a managed device tries to authenticate users from either an external LDAP or an external RADIUS server database such as Cisco ISE. But you can't use both methods. And if you choose to go with external authentication,

you must create an authentication object for each server to specify the exact location you want to authenticate users from. The authentication object contains your settings for connecting to and checking the user database from the corresponding server.

The appliance will then check each configured authentication server in the order in which they're listed in the system policy when attempting to find the user in the database.

An internal user will be automatically converted to external authentication if the same username and password exists on an external server. Once the user has been converted to external authentication, it can't revert to internal authentication!

User Privileges

You can create users using predefined roles or by creating custom roles to assign to individuals or even a group of users. Let's say you want to create a group of users that only gets access to data to analyze the security events for the monitored network, but you don't want this group to ever gain access to the administrative functions of the Firepower system itself. This is where the ability to create custom roles really shines—you can easily use predefined roles like Discovery Admin and Security Analyst, which allow the users to view network events without being able to change any configuration settings. Creating custom roles provides even more detail on what a particular user or a given group of users can and cannot do—nice!

We'll explore this more later in the section called "Configuring External Authentication." For now, it's key to remember that once the user logs in externally for the first time and receives the default access role, you'll be able to find this user on the User Management page and add or remove access rights for them. You can also opt to not modify the rights, which will grant the user the rights via the default access role. And remember, when you create internal users, you assign the role manually as you create them.

The default access role for an externally authenticated user can be overridden by configuring management for access rights. You can get this done through

either the LDAP or RADIUS groups or objects, where the permissions for users originate from the default access rights assigned to their specific group—assuming they belong to one. If they don't belong to an LDAP group or RADIUS object, they'll assume the default role.

Predefined User Roles

Here's a list of the Firepower System predefined user roles, depending on the features you've licensed:

Access Admins

Can view and modify access control and file policies, but they can't apply their policy changes.

Administrators

Can set up the appliance's network configuration, manage user accounts and Collective Security Intelligence cloud connections, and configure system policies settings. Users with the Administrator role have all the rights and privileges of all other roles, except for lesser, restricted versions of those privileges.

Discovery Admins

Can review, modify, and delete network discovery policies, but they can't apply their policy changes.

External Database

These users can query the Firepower System database using an external application that supports JDBC SSL connections. On the web interface, they can access the online help and user preferences.

Intrusion Admins

Can review, modify, and delete intrusion policies and intrusion rules, but they can't apply their policy changes.

Maintenance Users

Can access monitoring functions like health monitoring, host statistics, performance data, system logs, and maintenance functions, including task scheduling and backing up the system. Maintenance Users don't have access

to the functions in the Policies menu, and they can only access the dashboard from the Analysis menu.

Network Admins

Can review, modify, and apply device configurations as well as review and modify access control policies, but they can't apply their policy changes.

Security Approvers

Can view and apply policy changes, but they can't create configuration and policy changes.

Security Analysts

These users can review, analyze, and delete intrusion, discovery, user activity, connection, correlation, and network change events. They can review, analyze, and when applicable, delete hosts, host attributes, services, vulnerabilities, and client applications. Security Analysts can also generate reports and view health events, but they can't delete or modify these.

Security Analysts (Read Only)

These users enjoy all the same rights as Security Analysts, except that they can't delete events.

Custom user roles

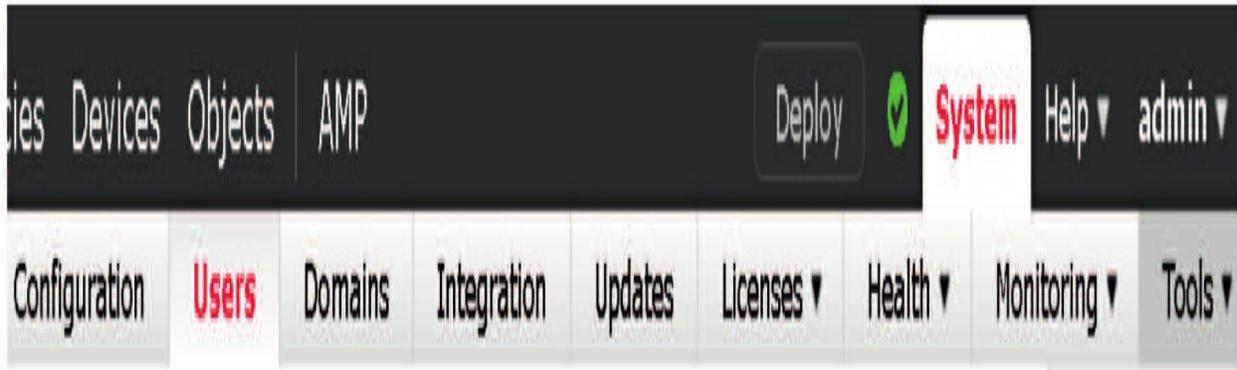
Allow you to customize exactly what users with a designated role can access. Custom user roles also allow you to place precise restrictions on exactly what information users can view.

Keep in mind that any role can be the default access role for externally authenticated users.

Creating New User Accounts

Now's the perfect time for me to demonstrate exactly how to create internal user accounts.

From the menu options on the right side of the menu bar, choose **System>Users**.



Or using the new Light Theme interface, you click the gear now, which is nice.

Deploy



admin ▾

Configuration

Health

Monitoring

Users



Monitor

Audit

Domains

Policy

Syslog

Integration

Events

Statistics

Updates

Blacklist

Monitor Alerts

Tools

Licenses

Backup/Restore

Smart Licenses

Scheduling

Classic Licenses

Import/Export

Data Purge

You'll get the User Management page, which lists the information on existing user accounts.

Username	Real Name	Roles	Authentication Method	Password Lifetime	
admin		Administrator	Internal	Unlimited	/

Here, I want to point out two things you need to focus on. First, there's a default admin user, which uses an internal authentication method, and second, there are those three tabs at the top: Users, User Roles, and External Authentication. I'll be referring to these tabs throughout this section.

Okay—so let's create a new user and assign it a role. From the User Management screen, on the right side, click Create User, which will bring up the screen shown below called the User Configuration Options page:

If you create a user here, you're creating a local user account.

If you click Use External

Authentication Method instead,
most of the options on this page will
disappear. You'll only be able to pick
the user role from the screen pictured.

User Configuration

User Name	<input type="text"/>	
Real Name	<input type="text"/>	
Authentication	<input type="checkbox"/>	Use External Authentication Method
Password	<input type="password"/>	
Confirm Password	<input type="password"/>	
Maximum Number of Failed Logins	<input type="text" value="5"/>	(0 = Unlimited)
Minimum Password Length	<input type="text" value="8"/>	
Days Until Password Expiration	<input type="text" value="0"/>	(0 = Unlimited)
Days Before Password Expiration Warning	<input type="text" value="0"/>	
Options	<input type="checkbox"/>	Force Password Reset on Login
	<input type="checkbox"/>	Check Password Strength
	<input type="checkbox"/>	Exempt from Browser Session Timeout

User Role Configuration

Default User Roles	<input type="checkbox"/>	Administrator
	<input type="checkbox"/>	External Database User (Read Only)
	<input type="checkbox"/>	Security Analyst
	<input type="checkbox"/>	Security Analyst (Read Only)
	<input type="checkbox"/>	Security Approver
	<input type="checkbox"/>	Intrusion Admin
	<input type="checkbox"/>	Access Admin
	<input type="checkbox"/>	Network Admin
	<input type="checkbox"/>	Maintenance User
	<input type="checkbox"/>	Discovery Admin
	<input type="checkbox"/>	Threat Intelligence Director (TID) User

Cancel

Save

User Configuration

User Name

Real Name

Authentication Use **External** Authentication Method

Options Exempt from Browser Session Timeout

User Role Configuration

- Default User Roles
- Administrator
 - External Database User (Read Only)
 - Security Analyst
 - Security Analyst (Read Only)
 - Security Approver
 - Intrusion Admin
 - Access Admin
 - Network Admin
 - Maintenance User
 - Discovery Admin
 - Threat Intelligence Director (TID) User

Cancel

Save

For now, I'm just going to create a simple internal user with a Security Analyst (Read Only) role.

The figure shows how to create an internal user.

You can see that I created a user named **todd** and forced the password to be reset on first login.

User Configuration

User Name	<input type="text" value="todd"/>	
Real Name	<input type="text"/>	
Authentication	<input type="checkbox"/>	Use External Authentication Method
Password	<input type="password" value="....."/>	
Confirm Password	<input type="password" value="....."/>	
Maximum Number of Failed Logins	<input type="text" value="5"/>	(0 = Unlimited)
Minimum Password Length	<input type="text" value="4"/>	
Days Until Password Expiration	<input type="text" value="0"/>	(0 = Unlimited)
Days Before Password Expiration Warning	<input type="text" value="0"/>	
Options	<input type="checkbox"/>	Force Password Reset on Login
	<input type="checkbox"/>	Check Password Strength
	<input type="checkbox"/>	Exempt from Browser Session Timeout

User Role Configuration

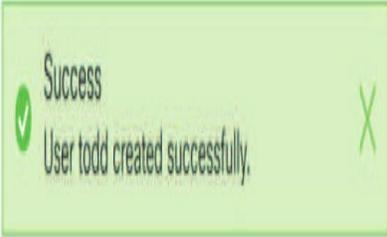
- | | | |
|--------------------|-------------------------------------|------------------------------------|
| Default User Roles | <input type="checkbox"/> | Administrator |
| | <input type="checkbox"/> | External Database User (Read Only) |
| | <input type="checkbox"/> | Security Analyst |
| | <input checked="" type="checkbox"/> | Security Analyst (Read Only) |
| | <input type="checkbox"/> | Security Approver |
| | <input type="checkbox"/> | Intrusion Admin |

I also left the default for the maximum number of failed logins at 5 and changed the minimum password length to 4 characters. I made sure to not set the password to expire.

Another thing you can do when creating a user is to opt for checking password strength. Doing this will require strong passwords, meaning they'll have to be at least eight alphanumeric characters— uppercase and lowercase, including at least one numeric character. The password can't appear in a dictionary or include consecutive, repeating characters.

I chose to exempt from browser session timeout since the user role is a read-only analyst. You probably wouldn't want a user to timeout if they were logged into a screen in a large network operations center (NOC), where the browser is displayed in all its glory on the NOC's huge screen, right?

I clicked Save... you can see that the user was created and that it's active.



Username	Real Name	Roles	Authentication Method	Password Lifetime	
admin		Administrator	Internal	Unlimited	/
todd		Security Analyst (Read Only)	Internal	Unlimited	 / 

So, now that we've created a user using a predefined role and verified that the user is active, let's create a custom user role and assign **Todd** to it.

First, take a look at the figure below and notice all the available menu items on the screen.



Now check out the right side and notice I can choose to log out as admin. Once I log in as **todd**, we'll get a look at the options available as System Analyst (Read Only).

The available menu items that I get after logging in as **todd** are now different! You can see that no Policies, Devices, and so on are available to **todd** right now.



Here, I only have the options for Overview and Analysis. And there on the right side, you can see that I no longer have the System Configuration or Health policy menu options. I'm going to log back in as admin now and then create a custom user role.

This next screen shows the predefined user roles available. Notice that while they can be disabled, they're all enabled by default.

Configure Permission Escalation

Create User Role

User Role	Enabled	Actions
Access Admin System-Provided	<input checked="" type="checkbox"/>	/ + - -
Administrator System-Provided	<input checked="" type="checkbox"/>	/ + - -
Discovery Admin System-Provided	<input checked="" type="checkbox"/>	/ + - -
External Database User (Read Only) System-Provided	<input checked="" type="checkbox"/>	/ + - -
Intrusion Admin System-Provided	<input checked="" type="checkbox"/>	/ + - -
Maintenance User System-Provided	<input checked="" type="checkbox"/>	/ + - -
Network Admin System-Provided	<input checked="" type="checkbox"/>	/ + - -
Security Analyst System-Provided	<input checked="" type="checkbox"/>	/ + - -
Security Analyst (Read Only) System-Provided	<input checked="" type="checkbox"/>	/ + - -
Security Approver System-Provided	<input checked="" type="checkbox"/>	/ + - -
Threat Intelligence Director (TID) User System-Provided	<input checked="" type="checkbox"/>	/ + - -

These predefined roles can be edited, but when you try to save the changes, you'll be asked to name the new custom role. So basically, you can edit an existing role, but you can't save it as a predefined role—only as a new role.

Creating a Custom User Role

Okay, so on the top right side of the main menu, I'm clicking Create User Role, and as shown in the figure, I'm going to create a user role named Helpdesk and assign some individual custom options for this role.

Now once the custom user role is created, we're allowed to assign a user to it.

User Role Configuration

- Administrator
- External Database User (Read Only)
- Security Analyst
- Security Analyst (Read Only)
- Security Approver
- Default User Roles Intrusion Admin
- Access Admin
- Network Admin
- Maintenance User
- Discovery Admin
- Threat Intelligence Director (TID) User
- Custom User Roles Helpdesk

Cancel

Save

Check out the figure and notice that when user **todd** is edited, the custom user role of Helpdesk is now available at the bottom. Once I click Save, user **todd** will only be using the newly created, custom role of Helpdesk.

Managing User Role Escalation

Let's say that user **todd** was working alone on a holiday when something really serious goes wrong. The user **todd** calls you in a panic saying, "What do you want me to do? All I can do is verify information!" If you can't VPN in and take care of the issue yourself, this would be a perfect example of when to rely on permission escalation.

If you take a look back at the predefined user roles, you'll see the blue Configure Permission Escalation button up there at the top?

Clicking that button will get you to this screen.

Here, you can choose to escalate to any predefined or custom user role that you've created. I'm going to choose Administrator. Now I'll go back to User Roles and edit the role I want escalated.

I'm going to edit that Helpdesk custom role I just created.

Configure Permission Escalation

Escalation Target:

Administrator ▼

Access Admin

Administrator

Discovery Admin

External Database User (Read Only)

Intrusion Admin

Maintenance User

Network Admin

Security Analyst

Security Analyst (Read Only)

Security Approver

Threat Intelligence Director (TID) User

Helpdesk

Name

Helpdesk

Description

Created by Todd

Menu-Based Permissions

- ▶ Overview
- ▶ Analysis
- ▶ Policies
- ▶ Devices
- ▶ Object Manager
 - Cisco AMP
 - Intelligence
 - Deploy Configuration to Devices
- ▶ System

System Permissions

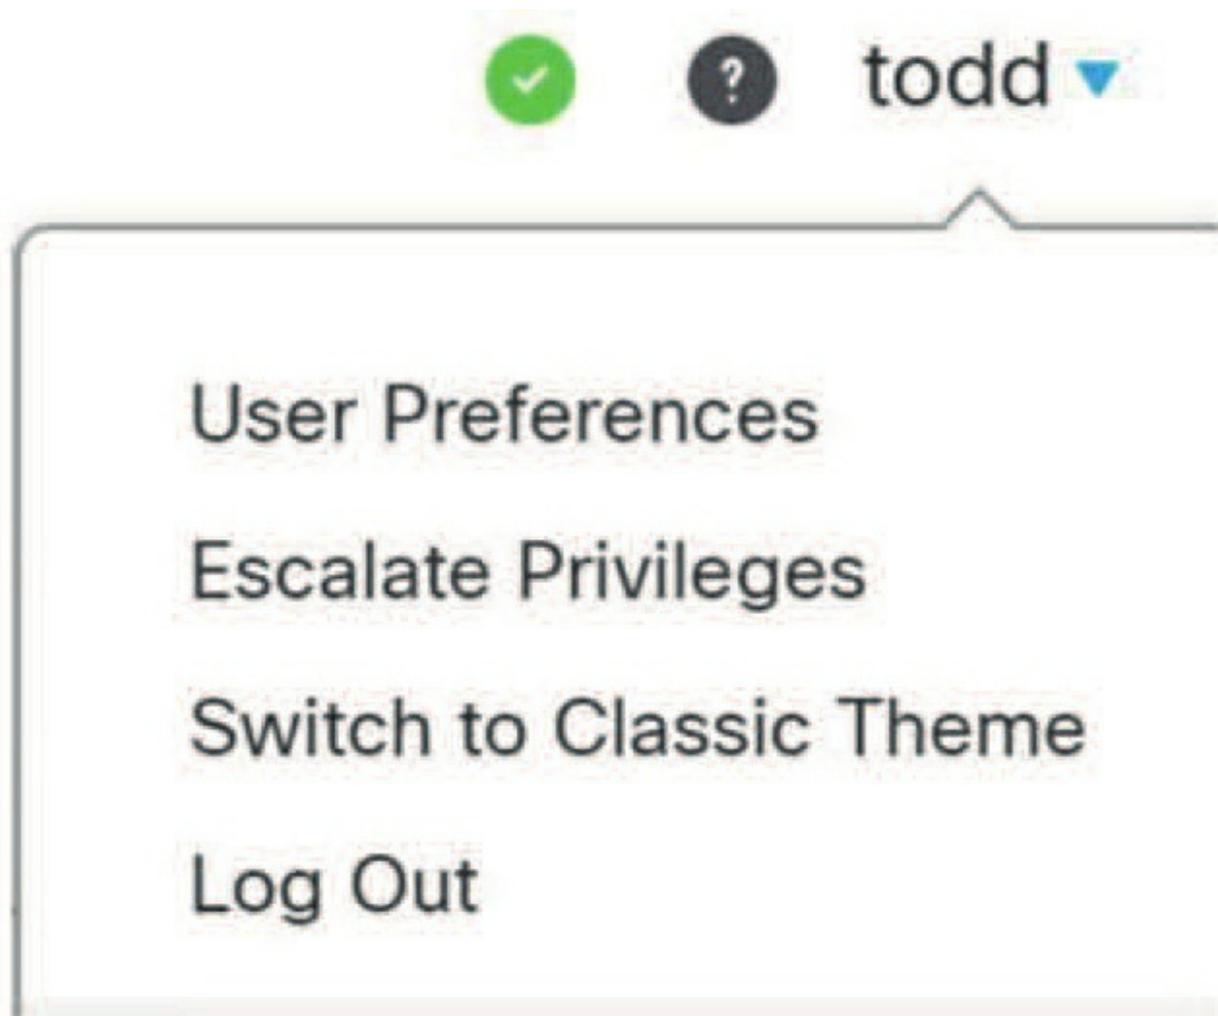
- External Database Access (Read Only)
- Set this role to escalate to: Administrator
 - Authenticate with the assigned user's password
 - Authenticate with the specified user's password

You can see that I choose the Helpdesk role and set the role to escalate to Administrator using the assigned user's password.

This cool feature really comes in handy when you need to escalate a user's privileges on a temporary basis, but I still haven't had any clients use it yet. It's still good to know it's there just in case though!

Okay—so now we're going to log out as Admin and then log back in as **todd** as shown. Notice that we get a new option under the user **todd**. That's right—Escalate Privileges!

Now if I choose the Escalate Privileges option and entered the administrator password when prompted, user **todd** would be escalated to Administrator.



To allow user **todd** to escalate using his own password, you would want to choose the radio button **Authenticate With The Specified User's Password** instead of using the **Authenticate With The Assigned User's Password** radio button.

Configuring External Authentication

Now, why would you go with an external authentication method over an internal one for users logging into the FMC? Because if you do that, you get to set up a directory on your network to organize different objects like user credentials in a convenient, centralized location. This way, if you ever need to change a user profile, you just make the changes in one place and it will affect the user's rights across the network—way better than ploddingly going to each network device, right?

Authentication objects are server profiles for external authentication servers, which contain the settings for the connection as well as authentication filter settings for your external servers. So, when you create an authentication object, you're defining settings that let you connect to an authentication server.

These objects are created, managed and deleted on the FMC, not on the managed devices themselves. Using an external authentication method allows a user to log in to any FMC or managed device with a single login by applying a platform policy on the managed device, because when you apply the policy, the object is copied to the device— sweet!

Creating Authentication Objects

Always remember that when creating an authentication object, you've got to make sure you have a solid IP connection through the network from your FMC to the authentication server(s) you'll be using.

From the menu options on the right side of the menu bar, choose **System> Users** and then choose the Login Authentication tab found in the main panel of the screen.

In the upper-right corner, click the **Add Authentication Object** button to get the screen shown below:

External Authentication

Save Cancel Save and Apply

Shell Authentication Disabled + Add External Authentication Object

Method	Enabled
--------	---------

The first option to choose is what type of authentication object you want to create: LDAP or RADIUS? The figure shows the authentication method, name, and server type that I'm going to use.

In the Server Type field, you can choose OpenLDAP, MS Active Directory, Oracle Directory, or Other. Choosing MS Active Directory, Oracle, or OpenLDAP populates the page with some default values used with those servers. Going with the Other option populates the page with no default values at all. I highly recommend choosing the AD server type because if you don't, you won't be able to populate users in your access control rules! Check out this figure

External Authentication Object

Authentication Method

CAC Use for CAC authentication and authorization

Name *

Description

Server Type

[Set Defaults](#)

After you choose your object type, you've got to set your server IP. You can have a backup server, but that's optional in the configuration. Here shows an LDAP server that I've got running on the local network.

Primary Server

Host Name/IP Address *

Port *

Backup Server (Optional)

Host Name/IP Address

Port

From here, you want to go down the page and configure LDAPspecific parameters like the ones shown.

LDAP-Specific Parameters

Base DN *	<input type="text" value="DC=sfgtc,DC=local"/>	<input type="button" value="Fetch DNs"/>
Base Filter	<input type="text" value="DC=sfgtc,DC=local"/>	
User Name *	<input type="text" value="CN=Configuration,DC=sfgtc,DC=local"/>	
Password *	<input type="text" value="CN=Schema,CN=Configuration,DC=sfgtc,DC=local"/>	
Confirm Password *	<input type="password" value="....."/>	

Really—the hardest part of the LDAP configuration is getting the right

information to configure the LDAP-specific parameters in the first place!

Here's a list to fill you in on what the values above actually represent:

Base DN

Sets a starting point for searching the LDAP directory tree. If you filled out the User Name and Password/Confirm Password fields for the LDAP admin account, you can hit the Fetch DN's button and it'll use the configured credentials to log in to the LDAP server and fetch the DN value.

Base Filter

Sets a filter that retrieves only the objects in the Base DN that have the characters you configured in the filter. You need to use parentheses when defining the characters!

User Name

Type the name of a user that's authorized to access the objects in the LDAP directory using the canonical name. For example, the username we're using in this field is *cn=administrator,cn=users,dc=sfgtc,dc=local*. You can also type this in as *administrator@sfgtc.local*, either way works.

Password / Confirm Password

Type the password for the username you entered.

After you've configured your LDAP-specific parameters, you can click the Show Advanced Options icon. From here, you can configure an encryption method if your LDAP server is configured to support it. Plus, it lets you set a path to the location of the TLS or SSL authentication certificate if you're using one.

The User Name Template option allows you to specify how usernames entered in the user login field should be formatted by mapping the string conversion character (%) to the value of the Pluggable Authentication Module (PAM) login attribute for the user:

▼ Show Advanced Options

Encryption SSL TLS None

SSL Certificate Upload Path No file chosen

User Name Template

Timeout (Seconds)

This username template is the format for the distinguished name used for authentication. When a user enters a username into the login page, the name is substituted for the string conversion character and the resulting distinguished name is then used to search for the user credentials.

Last up, the Timeout field is used to let you to set the amount of time the appliance takes to fail over to the backup LDAP server if one is configured.

After you've finished with the Show Advanced Options portion of the page, then configure the Attribute Mapping options as here.

Attribute Mapping

UI Access Attribute *

Shell Access Attribute *

▼ Group Controlled Access Admin Administrator

Fetch Attrs

- replUpToDateVector
- revision
- sAMAccountName
- sAMAccountType
- samDomainUpdates
- serverReference
- serverReferenceBL
- serverState

Different types of LDAP servers use different attributes to store user data.

For an LDAP server that uses the PAM like mine does, use the login attribute of `uid`. Of course, if the PAM login attribute for the target server is something other than `uid`, set it here. For a Microsoft Active Directory server, use a UI Access attribute

of `sAMAccountName` or `userPrincipalName` to enable user retrieval.

The Fetch Attrs button, which provides the output shown above, will allow you to access the LDAP server via the impersonation account login credentials to obtain UI access and/or shell access attributes—if they exist.

If you prefer to base access permissions on a specific user's membership within an LDAP group, you can specify distinguished names for existing groups on your LDAP server for each of the access roles used by your Firepower System. Do this from the Group Controlled Access Roles configuration.

▼ Group Controlled Access Roles (Optional)

Access Admin

Administrator

Discovery Admin

External Database User

Helpdesk

This is the best way to set an AD group to log in as a predetermined role without having to manually set each user!

You can also configure a default access role for those users detected by LDAP that don't belong to any specified groups. When a user logs in, the Firepower System dynamically checks the LDAP server and assigns access rights per the user's current group membership.

It's vital to remember that you can also add roles to users in the User Management interface. But you can't assign privileges lower than what's already granted to the user by the Group Controlled Access Roles settings!

By configuring the Shell Access Filter settings, you can allow LDAP-authenticated users shell access, but to prevent anyone from having shell access, you need to check the Same as Base Filter box.

This pictures the Shell Access Filter and Additional Test Parameters settings.

Last, as shown, you can test your server settings with the Test button. You can also test specific user credentials on the server from there by entering the username and password.

Shell Access Filter

Shell Access Filter ⓘ Same as Base Filter

(Mandatory for FTD devices)

Additional Test Parameters

User Name

Password ⓘ

After you set up your server and test your connectivity and authentication, click the Save button at the bottom of the screen to save your configuration.



Here are a couple things to remember:

1. If you don't set the default user role, your external users won't be able to log in, so don't skip this! Your other and better option is to set a Group Controlled Access role because then you wouldn't have to set the default user role.
2. The last thing is to enable the LDAP by clicking on the slide bar and click Save and Apply and you're set!

Summary

We covered a bunch of administrative functions for user account management in this chapter. Going through it, you learned all about internal and external user account management and how to describe various user roles, plus how to create custom user roles.

I demonstrated how to configure both internal and external user accounts and showed you how to permit a user to escalate their account privileges.

You discovered that user authentication can be achieved locally through the internal database or via an external authentication server like LDAP or RADIUS.

Chapter 18: Advanced Network Analysis

Just because this chapter doesn't dovetail to specific exam objectives doesn't mean it isn't important! It's full of essential insights into the art and science of Firepower event analysis, which can be really hard to come by. Unlike most of the previous chapters, it's not focused on configuration, but instead zeroes in on actually interpreting the system's output. Once Firepower is up and running smoothly, analyzing the events it generates correctly is an important skill!

On a basic level, event analysis comes down to keenly observing the events the system is generating and responding to them by taking action in a strategic and appropriate way. These actions range all the way from incident response (circle the wagons!) to tuning out false positive events.

So get ready, because I'm going to guide you through critical elements like intrusion, security intelligence, and dealing with file/ malware.

To find exam study material such as hands-on labs access, videos, downloadable supplemental material, and practice questions, head to www.lammle.com/firepower.

Event Analysis Principles

Because intrusion events provide the most complex analysis challenge, I'm going to convey some principles for analyzing this type of event. Firepower has several features to help analysts locate the most critical events and tune out the ones that just amount to noise. Over the years, updates to Snort, as well as improvements in rule keywords and preprocessor configurations, have helped to reduce some of this noise. But still, depending on the rules enabled, each newly deployed intrusion prevention system (IPS) will probably require some rule tuning.

False Positives

When I refer to “tuning,” what I really mean is noise reduction. In an IPS, “noise” is how we refer to events that trigger but are actually of no practical value. You’ll hear them often being referred to as false positives. A simple case of a false positive would be when an IPS rule alerts on traffic that it wasn’t intended to detect. It’s not that the rule is necessarily faulty or there is some problem with Snort itself, because false positives are often the result of the sheer volume of packets that must be inspected. They can also be triggered because the rule has to account for any packets where an attack *might* be present.

Here’s a good example: Say we’ve got a rule that’s looking for the string `cmd.exe` in an HTTP request. It’s probably looking for some type of abuse on a Microsoft Windows host because ancient versions of Microsoft Internet Information Services (IIS) had vulnerabilities where an attacker could actually execute commands through the web server by finding the path to the command processor. To guard against this, you could use a Snort rule that’s simply looking for `cmd.exe` in an HTTP request. Typically, this rule wouldn’t trigger on normal HTTP traffic, but it would detect someone trying to find the Windows command processor.

Sounds good, but what if a website just happened to host a graphic named `cmd.exe.gif`? This would definitely cause the rule to alert and possibly block an HTTP request for the image. Clearly, in this case that alert would be classified as a false positive. The rule did exactly what it was designed to do but here, but `cmd.exe` wasn’t an actual threat—it just happened to be part of an unfortunate name that someone dubbed to an innocuous image file!

False Negatives

The opposite of a false positive is a false negative. It’s defined as the failure of a rule to detect the attack it was designed to catch. While false positives are to be expected, false negatives are clearly something we want to avoid. The failure of a rule to detect an actual attack represents a failure of the IPS as a whole, so it’s definitely something anyone writing a rule is trying hard to

steer clear of! Understanding this philosophy really clarifies why false positives happen.

Looking back at the previous example, the false positive for `cmd.exe.gif` could be avoided by adding the path to `cmd.exe`. So instead, we could rewrite the rule to look for `c:\winnt\system32\cmd.exe` instead and eliminate the false positive triggered by that `cmd.exe.gif` file. Again, sounds great, but what if the Windows host used `c:\Windows` for its system root? This means the path to `cmd.exe`

would be `c:\Windows\system32\cmd.exe`, and if our rule

is too specific, all of a sudden, we're faced with the potential for a horrible false negative. And even though rules are written to avoid as many false positives as possible, they're ideally created in a way so that we get *zero* false negatives! So we'll probably just have to settle with writing the rule so it only looks for `cmd.exe` and live with an occasional false positive because we definitely can't live with false negatives. Trade-offs!

While it's not technically a false negative, you can also miss attacks because a certain rule isn't enabled. Referring back again, rules to detect `cmd.exe` are in the Cisco Snort rule set, but if they're not enabled in the deployed Intrusion policy, our IPS will totally fail to spot this attack. In this case, the reason these rules might not be enabled by default is because modern Windows systems aren't vulnerable to this type of directory traversal attack. But if your network is still living in 2001, you probably need to turn on some of these old rules, right? So knowing when to enable the correct rules isn't really a function of event analysis, it's more the art of understanding which rules you should apply to a specific system so it'll be properly protected relevant to that particular system's actual needs.

Possible Outcomes

Of course, there are two other possibilities—a true positive and a true negative. The true positive is when a rule correctly triggers on the malicious activity it was designed for and a true negative is when rules correctly do not alert for benign traffic. The four possible results of traffic inspection are

shown in the table below:

Traffic is malicious	Traffic is benign	Rule matches traffic	True positive
False positive	Rule does not match traffic	False negative	True negative

Of these four results, three of them are pretty much considered normal and acceptable in an IPS. The unacceptable one is the false negative—the failure to detect an actual attack! As security professionals, we can deal with a little noise in our system tuning, but we *always* try to err on the side of detection versus non-detection.

Now, I have to point out here that false negatives aren't always the fault of the IPS rule set. The device absolutely has to have visibility to all the traffic it's designed to inspect, and it also must have the necessary processing resources to evaluate the traffic fully. Inadequate visibility or detection resources can also be culprits behind portions of traffic getting by inspection, which can also cause a false negative.

The Goal of Analysis

Maybe you're thinking, "Why analyze? Can't we just put our devices inline, load the appropriate rule sets, and trust the system to block all the nasty stuff?" Good question! On the surface it makes some sense because, well, if the main job of the IPS is to stop attacks, can't it just do that regardless of who's checking out the analysis screens?

Well of course, strategically placing inline IPS devices within your network will definitely provide some protection against attacks. It will probably even hamstring malware a bit, but it's definitely not a silver bullet that'll vaporize all attacks and massacre all your network security horrors!

First, recall that some devices can be deployed in passive mode, meaning all they can do is alert. They can't just stop attacks cold on their own. So it comes down to us actually following up on the events generated by these devices to address potential issues. Analysis is key in these situations because we can prioritize our responses, so we don't waste a bunch of time and energy on alerts that are actually false positives. Plus, even though inline devices have the power to stop attacks, they also have the potential to impact

legitimate traffic if the alert is a false positive. So again, it's up to us to analyze the output of these devices to quickly identify and mediate false positives to ensure that there's only minimal impact on the organization's operations and activities. It's our actual analysis of events plus tuning our rules accordingly that are the critical factors for making it all come together.

Another key thing to remember is that the defender has to deal with a barrage of attacks occurring constantly, 24/7/365, and the attacker only has to succeed once to compromise your network! Anyone who thinks that there's *any* security tool that can stop all attacks cold isn't living in a little place we like to call reality. You *will* be compromised, it will happen more than once, and it's probably happened already! And it happened even though you may be have the finest IPS money can buy while tuning it perfectly too.

I say all this because we're talking about analysis as a key facet of identifying which systems on your network have been compromised. The good news is that Firepower has some really excellent features for alerting you about hosts that are already likely to have been compromised. And again, the purpose of these rules is *not* to stop the malware that's already infected your systems, it's to alert you so you can take the appropriate action and clean them up.

I'm so not exaggerating here... I've seriously seen too many Firepower installations veritably screaming about hosts that are regularly connecting to botnets or known malware sites, with no one doing anything about locating and cleaning up those systems. Epic fail! Nobody needs that kind of shame, right? Don't let this be you!

Intrusion Events

Okay, with my little rant over, we're now going to scrutinize the mechanics of intrusion event analysis using workflows.

Workflows

Workflows are an awesome Firepower feature that comes in super handy for customizing your event views. You'll even find that most event types come with built-in workflows specially designed to help us sift through logs to zero

in on particular events. A workflow determines how events are displayed, which columns are displayed, and how they're sorted. Make a mental note that even though the term *event* is the technically correct one, you'll also hear them referred to as alerts. Basically, all everyone is referring to is "the message that's logged when a rule triggers."

To navigate to the default intrusion event workflow on the FMC, go to Analysis → Intrusions → Events. The default workflow is called Events By Priority and Classification.

Intrusion event workflows are all table based—they're really just columns of (mostly) text data. In contrast, you'll find that most of the built-in Connection Event workflows are bar, line, or pie graphs.

The different workflows help analysts (us) highlight different types of events. For example, most of the time you're going to be searching for the intrusion events that represent the most critical threats to your organization. While this is true, sometimes you just might be looking for the events most likely to be false positives. Or maybe you want to focus on events based upon their source IP address. Why? Because these are really useful when you're hunting down malware infections. Events by destination IP are helpful for zeroing in on external-facing server attacks. By using the built-in workflows and/or creating your own custom ones, you can quickly locate the specific types of events you're presently interested in.

To switch workflows, click the Switch Workflow link just to the right of the current workflow name. You'll find the link just to the right of Events By Priority and Classification as pictured in in the following figure:



Events By Priority and Classification [\(switch workflow\)](#)
Drilldown of Event, Priority, and Classification > [Table View of Events](#) > [Packets](#)

Clicking this link displays the workflows available. The next figure shows

the default intrusion event workflows.

If you add your own custom workflows, they'll also be listed here:

Events By Priority and Classification ✕

Destination Port

Event-Specific

Events By Priority and Classification

Events to Destinations

Impact and Priority •

Impact and Source

Impact to Destination

IP-Specific

Source and Destination

Source Port

Built-in Workflows

Here are some examples of built-in workflows:

Destination Port:

Lists the count of events for each destination port. The list is sorted with the port having the highest count of events at the top. Check it out:

The screenshot shows the Cisco AMP interface with the 'Destination Port' workflow selected. The table below displays the results, sorted by count in descending order.

Destination Port / ICMP Code	Count
80 (http) / tcp	10
39022 / tcp	2
35054 / tcp	2
445 (microsoft-ds) / tcp	2
51216 / tcp	2
51260 / tcp	2
55745 / tcp	2
45978 / tcp	1
49434 / tcp	1
63610 / tcp	1
63613 / tcp	1
587 (submission) / tcp	1
52677 / tcp	1
49206 / tcp	1

Displaying rows 1-14 of 14 rows | Page 1 of 1

View Copy Delete Review Download Packets
View All Copy All Delete All Review All Download All Packets

Last login on Friday, 2016-11-04 at 13:40:20 PM from 10.0.0.159

Event-Specific:

Lists the events by count from highest to lowest. This workflow is really sweet when you're hunting for false positives. By the way, it's common to find that the events triggering the most are the ones also the most likely to be false positives. Figure on next page.

Impact and Priority:

This is good—it displays five columns with events sorted by impact and also shows the inline result, priority, and count. Figure on next page.

Events by Priority and Classification:

Now this is the default workflow. It lists events with their priority and classification, and it's sorted by priority from high to low. Even though this is the default workflow, it's, well, honestly not the best. We all like workflows with the Inline Result and Impact columns much better! Here's the Events by Priority and Classification workflow found two pages down:

The screenshot shows the Cisco AMP interface with the 'Events By Priority and Classification' workflow selected. The table below displays the event data:

Message	Priority	Classification	Count
SERVER-IIS cmd.exe access (1:1002:18)	high	Web Application Attack	4
POLICY-OTHER Adobe ColdFusion admin interface access attempt (1:25975:2)	high	Potential Corporate Policy Violation	2
SMTP_COMMAND_OVERFLOW (124:1:2)	high	Attempted Administrator Privilege Gain	1
MALWARE-BACKDOOR x-door runtime detection (1:10185:6)	high	A Network Trojan was Detected	1
MALWARE-CNC Win.Trojan.Agent variant outbound connection (1:25532:1)	high	A Network Trojan was Detected	1
BROWSER-IE Microsoft Internet Explorer 11 VBScript redim preserve denial-of-service attempt (1:32564:4)	medium	Attempted Denial of Service	3
www.alexinfo Test Rule (1:1000010:1)	low	Not Suspicious Traffic	11
*** Testmyids.com test rule *** (1:1000000:5)	low	Misc Activity	3
*** Testmyids.com test rule *** no metadata (1:1000006:1)	low	Misc Activity	2
BLACKLIST_URI - known scanner tool muleblackcat (1:21257:3)	low	Detection of a Network Scan	1

Impact and Source:

This view is sorted by impact and it has six columns, which is the maximum for a built-in workflow. It displays the impact, inline result, source IP, message, priority, and count as you can see:

Impact and Source (switch workflow)

Drill Down of Impact and Source [Drill Down of Source and Destination IPs](#) [Table View of Events](#) [Packets](#)

2015-11-05 19:36:00 - 2016-11-06 19:05:50

Expanding

No Search Constraints [\(Edit Search\)](#)

Jump to...

	Impact	Inline Result	Source IP	Message	Priority	Count
			208.94.116.21 (vhost.phx1.nearlyfreespeech.net)	BROWSER-IE Microsoft Internet Explorer 11 VBScript redim preserve denial-of-service attempt (1:32564:4)	medium	2
			10.0.0.12	MALWARE-CNC Win.Trojan.Agent variant outbound connection (1:25532:1)	high	1
			172.16.0.2	SERVER-IIS cmd.exe access (1:1002:18)	high	4
			10.0.0.11	www.alex.info Test Rule (1:1000010:1)	low	3
			10.0.0.50	www.alex.info Test Rule (1:1000010:1)	low	3
			193.34.8.10	POLICY-OTHER Adobe ColdFusion admin interface access attempt (1:25975:2)	high	2
			10.0.0.12	MALWARE-BACKDOOR x-door runtime detection (1:10185:6)	high	1
			89.248.160.154 (blackhole.d.org)	BLACKLIST URI - known scanner tool muiblackcat (1:21257:3)	low	1
			96.19.3.245 (96-19-3-245.cpe.cableone.net)	www.alex.info Test Rule (1:1000010:1)	low	5
			82.165.177.154 (s193738556.websitehome.co.uk)	*** Testmvids.com test rule *** (1:1000000:5)	low	3
			82.165.177.154 (s193738556.websitehome.co.uk)	*** Testmvids.com test rule *** no metadata (1:1000006:1)	low	2
			10.0.0.13	SMTP_COMMAND_OVERFLOW (124:1:2)	high	1
			208.94.116.21 (vhost.phx1.nearlyfreespeech.net)	BROWSER-IE Microsoft Internet Explorer 11 VBScript redim preserve denial-of-service attempt (1:32564:4)	medium	1

Page 1 of 1 Displaying rows 1-13 of 13 rows

View	Copy	Delete	Review	Download Packets
View All	Copy All	Delete All	Review All	Download All Packets

Event-Specific (switch workflow)

Drill Down of Events > [Drill Down of Source IPs, or Destination IPs](#) > [Table View of Events](#) > [Packets](#)

2015-11-05 19:36:00 - 2016-11-06 18:54:23 ☺

Expanding

No Search Constraints ([Edit Search](#))

Jump to... ▾

<input type="checkbox"/>	Message	Count
<input type="checkbox"/>	www.alexinfo Test Rule (1:1000010:1)	11
<input type="checkbox"/>	SERVER-IIS cmd.exe access (1:1002:18)	4
<input type="checkbox"/>	BROWSER-IE Microsoft Internet Explorer 11 VBScript redim preserve denial-of-service attempt (1:32564:4)	3
<input type="checkbox"/>	*** Testmyids.com test rule *** (1:1000000:5)	3
<input type="checkbox"/>	POLICY-OTHER Adobe ColdFusion admin interface access attempt (1:25975:2)	2
<input type="checkbox"/>	*** Testmyids.com test rule *** no metadata (1:1000006:1)	2
<input type="checkbox"/>	SMTP_COMMAND_OVERFLOW (124:1:2)	1
<input type="checkbox"/>	MALWARE-CNC Win.Trojan.Agent variant outbound connection (1:25532:1)	1
<input type="checkbox"/>	MALWARE-BACKDOOR x-door runtime detection (1:10185:6)	1
<input type="checkbox"/>	BLACKLIST URI - known scanner tool muleblackcat (1:21257:3)	1

Displaying rows 1-10 of 10 rows << Page 1 of 1 >>

View	Copy	Delete	Review	Download Packets
View All	Copy All	Delete All	Review All	Download All Packets

Impact to Destination:

This view is the same as Impact and Source except the destination IP is also displayed.

Impact to Destination (switch workflow)

[Drill Down of Impact and Destination](#) > [Drill Down of Source and Destination IPs](#) > [Table View of Events](#) > [Packets](#)

2015-11-05 19:36:00 - 2016-11-06 19:09:13

Expanding

No Search Constraints [\(Edit Search\)](#)

Jump to... ▾

<input type="checkbox"/>	Impact	Inline Result	Destination IP	Message	Priority	Count
↓	1	↓	10.0.0.27	BROWSER-IE Microsoft Internet Explorer 11 VBScript redim preserve denial-of-service attempt (1:32564:4)	medium	2
↓	1	↓	10.0.0.201	MALWARE-CNC Win.Trojan.Agent variant outbound connection (1:25532:1)	high	1
↓	2	↓	10.0.0.11	SERVER-IIS cmd.exe access (1:1002:18)	high	4
↓	2		10.0.0.11	www.alex.info Test Rule (1:1000010:1)	low	3
↓	2		10.0.0.50	www.alex.info Test Rule (1:1000010:1)	low	3
↓	2	↓	10.0.0.11	POLICY-OTHER Adobe ColdFusion admin interface access attempt (1:25975:2)	high	2
↓	2		10.0.0.201	MALWARE-BACKDOOR x-door runtime detection (1:10185:6)	high	1
↓	2	↓	10.0.0.11	BLACKLIST URI - known scanner tool muleblackcat (1:21257:3)	low	1
↓	3		172.16.0.2	*** Testmyids.com test rule *** no metadata (1:1000006:1)	low	2
↓	3		172.16.0.2	*** Testmyids.com test rule *** (1:1000000:5)	low	2
↓	3	↓	64.147.108.71 (pb-smb2.pobox.com)	SMTP COMMAND OVERFLOW (124:1:2)	high	1
↓	3	↓	10.0.0.162	BROWSER-IE Microsoft Internet Explorer 11 VBScript redim preserve denial-of-service attempt (1:32564:4)	medium	1
↓	3		10.0.0.171	www.alex.info Test Rule (1:1000010:1)	low	1
↓	3		10.0.0.158	*** Testmyids.com test rule *** (1:1000000:5)	low	1
↓	3		10.0.0.13	www.alex.info Test Rule (1:1000010:1)	low	1
↓	3		10.0.0.167	www.alex.info Test Rule (1:1000010:1)	low	1
↓	3		10.0.0.152	www.alex.info Test Rule (1:1000010:1)	low	1
↓	3		172.16.0.2	www.alex.info Test Rule (1:1000010:1)	low	1

<< Page 1 of 1 >> Displaying rows 1-18 of 18 rows

View	Copy	Delete	Review	Download Packets
View All	Copy All	Delete All	Review All	Download All Packets

Impact and Priority (switch workflow)

Impact Based Event Summary > [Drill Down of Source and Destination IPs](#) > [Table View of Events](#) > [Packets](#)

2015-11-05 19:36:00 - 2016-11-06 19:02:54 🕒

Expanding

No Search Constraints [\(Edit Search\)](#)

Jump to... ▼

<input type="checkbox"/>	▼ Impact	Inline Result	Message	Priority	Count
↓	<input type="checkbox"/>	1	↓	BROWSER-IE Microsoft Internet Explorer 11 VBScript redim preserve denial-of-service attempt (1:32564:4)	medium 2
↓	<input type="checkbox"/>	1	↓	MALWARE-CNC Win.Trojan.Agent variant outbound connection (1:25532:1)	high 1
↓	<input type="checkbox"/>	2	↓	www.alexinfo Test Rule (1:1000010:1)	low 6
↓	<input type="checkbox"/>	2	↓	SERVER-IIS cmd.exe access (1:1002:16)	high 4
↓	<input type="checkbox"/>	2	↓	POLICY-OTHER Adobe ColdFusion admin interface access attempt (1:25975:2)	high 2
↓	<input type="checkbox"/>	2	↓	BLACKLIST URI - known scanner tool muleblackcat (1:21257:3)	low 1
↓	<input type="checkbox"/>	2	↓	MALWARE-BACKDOOR x-door runtime detection (1:10185:6)	high 1
↓	<input type="checkbox"/>	3	↓	www.alexinfo Test Rule (1:1000010:1)	low 5
↓	<input type="checkbox"/>	3	↓	*** Testmyids.com test rule *** (1:1000000:5)	low 3
↓	<input type="checkbox"/>	3	↓	*** Testmyids.com test rule *** no metadata (1:1000006:1)	low 2
↓	<input type="checkbox"/>	3	↓	BROWSER-IE Microsoft Internet Explorer 11 VBScript redim preserve denial-of-service attempt (1:32564:4)	medium 1
↓	<input type="checkbox"/>	3	↓	SMTP_COMMAND_OVERFLOW (124:1:2)	high 1

<< Page 1 of 1 >> Displaying rows 1-12 of 12 rows

View	Copy	Delete	Review	Download Packets
View All	Copy All	Delete All	Review All	Download All Packets

So as you can see, there's a candy store of default workflows to choose from. You'll have to determine which one you go with based upon the type of events you are looking for at the moment. Looking for evidence of malware or spyware on hosts? The Impact and Source workflow is what you want because it'll show you the events with the highest impact, whether or not the packet was dropped, and which source IP generated the event.

If it's DMZ or server attacks you are looking for, the Impact to Destination workflow is your daisy because it's very similar to the Impact and Source workflow but you'll also get to see the destination IP.

Looking for false positives? Try the Event-Specific workflow. Even though it won't give you any information on the source, destination, or severity of the event, it'll tell you which rules are triggering the most. Rules that trigger thousands of times in a short period of time tend to be false positives.

I want you to know that all of the example workflow figures were actually generated from the same set of intrusion events! This fact makes it really clear that the workflow you choose to go with has a big impact on which events float to the top in your analysis view. There's just no right or wrong answer to which workflow is best because it depends on the given situation and which one will work best for you at the time.

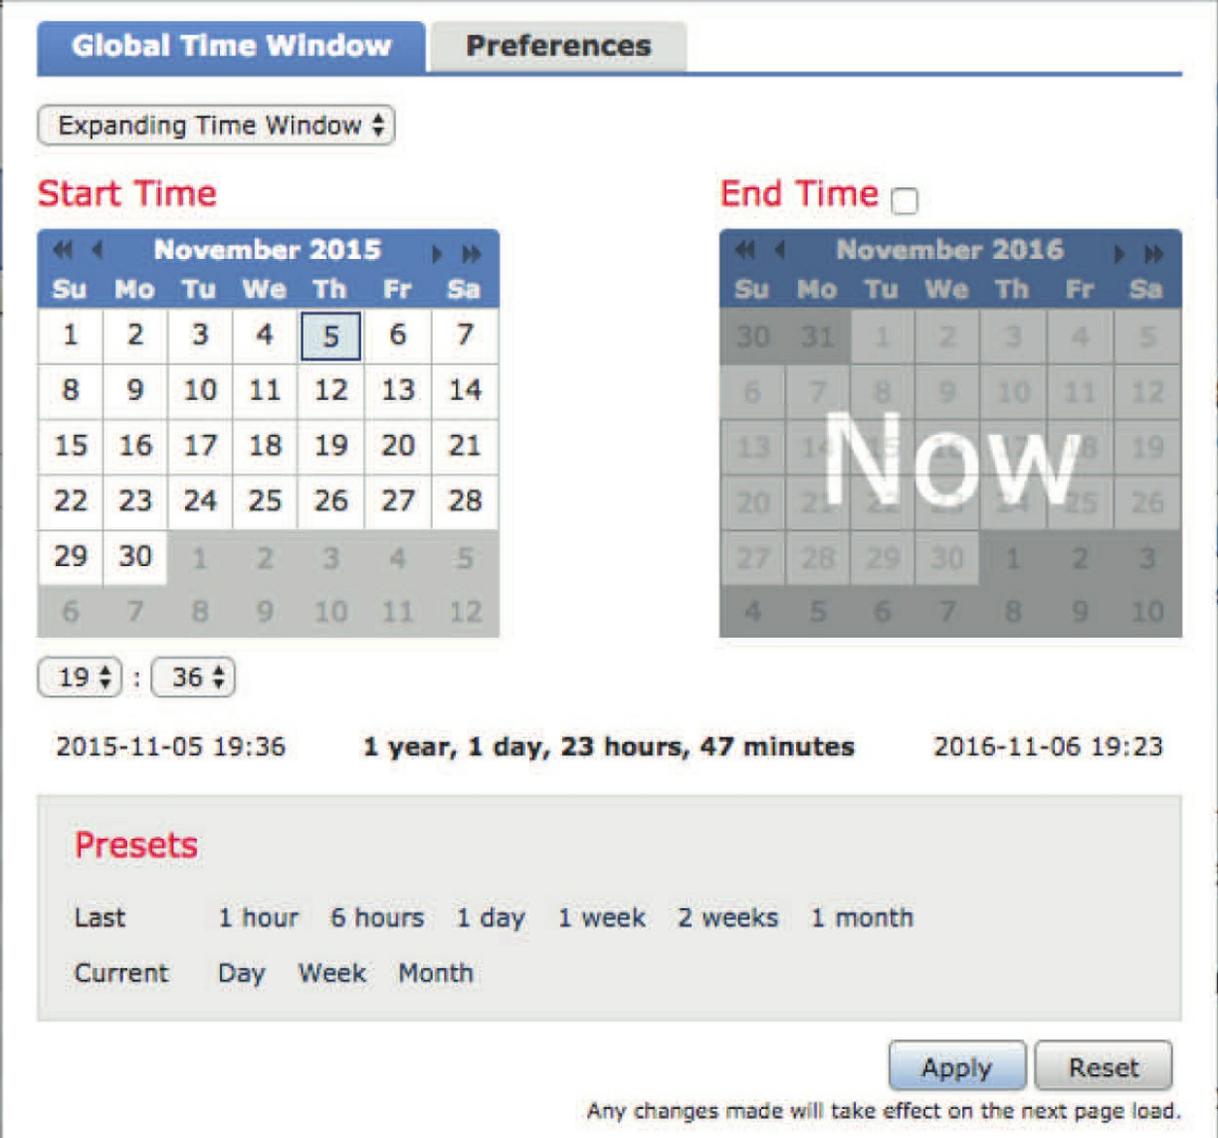
The Time Window

Most event data within Firepower is time-based, meaning the data consists of time-stamped events. When viewing these events, you choose a time window that corresponds to a start and stop time plus a date. Once you've done that, all the events that appear in an analysis workflow will have occurred within the time window you specified. If your time window is really huge—say weeks, even months—you'll clearly get shown more events than it would display if you delimited a small one instead.

In the figure below, you can see that the current time window is displayed as a start/stop date with the time in the upper-right corner of the screen.

To change the time window, click the link that looks like this:

You can click anywhere on the start/stop date or even on the clock icon. Clicking the pause symbol will pause the time windows if you've selected auto-refresh, and clicking the link opens a new window in your browser where you can set the event view time window as shown in the following figure:



Global Time Window Preferences

Expanding Time Window ▾

Start Time

November 2015						
Su	Mo	Tu	We	Th	Fr	Sa
1	2	3	4	5	6	7
8	9	10	11	12	13	14
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	1	2	3	4	5
6	7	8	9	10	11	12

19 : 36

2015-11-05 19:36 **1 year, 1 day, 23 hours, 47 minutes** 2016-11-06 19:23

End Time

November 2016						
Su	Mo	Tu	We	Th	Fr	Sa
30	31	1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28	29	30	1	2	3
4	5	6	7	8	9	10

Presets

Last 1 hour 6 hours 1 day 1 week 2 weeks 1 month

Current Day Week Month

Apply Reset

Any changes made will take effect on the next page load.

There are three types of time windows available: static, expanding, and sliding. You choose between them from the drop-down list in the upper-left

corner.

- **Static:** This time window has fixed start and stop times, so it's useful for locating events that occurred within a specific period. To use the static time window, select it from the dropdown at the top of the window, then delimit the start and stop time/date using the calendar and clock controls.
- **Expanding:** This time window has a fixed start time just like a static window, but instead of a fixed stop time, this window uses right now as the end time. Since the start time is fixed and the end time is constantly moving, this time window is always expanding, hence the name. If you set this time window to 60 minutes and begin analyzing events, after an hour you will have a 120-minute window. This is the most common time window used for interactive analysis because it always displays the most current events. When delimiting an expanding time window, you select the start time using the calendar and clock controls.
- **Sliding:** When you select a sliding time window, you'll find the calendar and clock controls have disappeared from the pop-up window. This is because a sliding window is a fixed duration with the front edge of it always being right now. You can fix the width of the time window and once you do, it'll slide as you continue to view events. To clarify, if you select a time window of 60 minutes, you'll continue to see only the last 60 minutes of events even if you remain logged in for hours. Sliding time windows are most useful for recurring reports. Say you have a weekly report comprising high-impact events. You would configure this report with a seven-day sliding time window so that each time it runs, it reports events for the previous seven days—cool huh?

Notice that there are presets available in the date/time pop-up. These are handy when you just want to look at some recent events because it's a lot faster than trying to use the calendar and clock controls to set a start/stop time.

The Firepower default time window is the last hour expanding, but as I said, once you delimit a new time window, it will become the time window that will remain in effect for the duration of your session. Once you log out and log back in, the time window resets to the default.

You can change this default using the Preferences tab. When you click this

tab, it will display the Time Window Preferences page:

Time Window Preferences

Refresh Interval (minutes)

Set to 0 to disable

Number of Time Windows

Single Multiple

Default Time Window

Show the Last - Static/Expanding ▾

▾

Use End Time

Save Preferences

Any changes made will take effect on the next page load.

The Preferences tab allows you to change your default to any of the three types and select the duration. There's also a radio button to set the number of time windows. By default, there are three: one for health events, one for audit events, and another one for everything else. The idea here is flexibility—you can select a time window of 30 days for your audit event view without affecting the time windows for your health and IPS event views. If you select a single time window, the Events Time Window tab name changes to Global Time Window.

Once your time window is set, you can concentrate on locating the events for analysis.

So clearly your time window has a huge effect on the number of events displayed in the workflow. A really common mistake people make when looking for events is forgetting to check the current time window. If you're using a static window when looking for an event that just happened, your event may be there all right, but your time window ended an hour ago—oops! To prevent this fate, keep the following advice in mind: When you don't see the events you're expecting to see, immediately check the time window!

Navigating the Analysis Interface

Your general goal when it comes to analyzing intrusion events is basically to determine if the event represents a false positive or an actual security issue. Sometimes you can get this done easily by just reading an event message (the meaningful text that describes the event). For instance, if the message clearly indicates that the rule's designed to detect a vulnerability that you know isn't there, it's a screaming false positive. Of course most of the time, it's not so simple and requires drilling into the event detail—even the packet data—to determine what you're really dealing with.

The workflows we just talked about consist of views or pages designed to help you single out the particular events you're interested in. Near the top of the screen you'll see a "breadcrumb trail" showing each page in the workflow. The page you're currently viewing will be a red link.

The next figure shows the Impact and Source workflow.

Impact and Source [\(switch workflow\)](#)

[Drill Down of Impact and Source](#) > [Drill Down of Source and Destination IPs](#) > [Table View of Events](#) > [Packets](#)

From this we can see there are four pages in this workflow:

1. Drill Down of Impact and Source
2. Drill Down of Source and Destination IPs
3. Table View of Events
4. Packets

The page names signify what we would find as we drill down into events. All the default workflows include a Table View of Events option, which contains all of the fields in an event even though some are actually hidden by default. All intrusion event workflows, whether built-in or custom, also include a Packets view.

Priority and Impact

There are several ways to drill down to the next page, and the one you'll want to go with depends on the events you want to focus on. The figure below shows a sample of some intrusion events we just might want to look into!

Impact	Inline Result	Source IP	Message	Priority	Count
1	↓	10.0.0.6	BLACKLIST DNS request for known malware domain proxim.lrcgalaxy.pl - virut (1:16304:6)	high	618
1	↓	187.95.73.80	MALWARE-OTHER sasser attempt (1:9419:9)	high	110
1	↓	10.0.0.27	MALWARE-CNC Win.Trojan.Meac malware component download request (1:29987:1)	high	109
1	↓	10.0.0.48	MALWARE-CNC Win.Trojan.Dampit variant outbound connection (1:29563:2)	high	33
1	↓	182.184.69.240	MALWARE-OTHER self-signed SSL certificate with default Internet Widqlts Pty Ltd organization name (1:19551:9)	high	24
1	↓	10.0.0.26	MALWARE-CNC Win.Trojan.Taldoor variant outbound connection (1:20204:5)	high	16
1	↓	10.0.0.26	MALWARE-CNC Win.Trojan.Rubinurd variant outbound connection (1:33305:3)	high	13
1	↓	10.0.0.6	MALWARE-CNC Win.Trojan.Saltty variant outbound connection (1:19964:9)	high	8
1	↓	14.139.155.135	MALWARE-OTHER sasser attempt (1:9419:9)	high	6
1	↓	60.167.135.207	MALWARE-OTHER sasser attempt (1:9419:9)	high	6
1	↓	10.0.0.126	MALWARE-CNC Win.Trojan.Sosork variant outbound connection (1:26506:3)	high	4
1	↓	10.0.0.48	MALWARE-CNC Phoenix exploit kit post-compromise behavior (1:21850:4)	high	3
1	↓	78.189.214.112	MALWARE-OTHER sasser attempt (1:9419:9)	high	2
1	↓	10.0.0.27	BLACKLIST DNS request for known malware domain godson355.vicp.cc - Win.Trojan.Meac (1:29986:1)	high	2
1	↓	10.0.0.48	MALWARE-CNC Win.Trojan.Zeus encrypted POST Data exfiltration (1:27919:3)	high	2
1	↓	10.0.0.23	MALWARE-CNC Win.Trojan.Dropper.Daws variant outbound connection (1:25099:2)	high	2

This page gives us some really great information. First of all, these are all Impact 1 events. We can also see that the priority is high. So, what's the difference between impact and priority?

- **Priority:** Priority is determined by the rule writer and it's set using the **priority** keyword or determined by the rule classification. Possible values are low, medium, and high. The important thing to remember is that what this is really telling you is how important the rule writer felt the alert would be. For instance, a rule written for a vulnerability that could result in a host being completely compromised will probably be a high priority one, whereas a rule designed to alert on a simple protocol anomaly would be low.
- **Impact:** The level of impact is determined by Firepower based upon what it knows about the rule, the packet, and the victim host. There are several levels of impact and each one is also set by Firepower.

This next figure was taken from the online help and describes how each impact level is determined:

Impact Level	Vulnerability	Color	Description
0	Unknown	gray	Neither the source nor the destination host is on a network that is monitored by network discovery.
1	Vulnerable	red	Either: <ul style="list-style-type: none"> the source or the destination host is in the network map, and a vulnerability is mapped to the host the source or destination host is potentially compromised by a virus, trojan, or other piece of malicious software
2	Potentially Vulnerable	orange	Either the source or the destination host is in the network map and one of the following is true: <ul style="list-style-type: none"> for port-oriented traffic, the port is running a server application protocol for non-port-oriented traffic, the host uses the protocol
3	Currently Not Vulnerable	yellow	Either the source or the destination host is in the network map and one of the following is true: <ul style="list-style-type: none"> for port-oriented traffic (for example, TCP or UDP), the port is not open for non-port-oriented traffic (for example, ICMP), the host does not use the protocol
4	Unknown Target	blue	Either the source or destination host is on a monitored network, but there is no entry for the host in the network map.

Here are some example impact levels for various events:

- A rule triggers on a packet targeted for port 80 on a host, but the target IP address isn't present in a network range defined in the Network Discovery policy. Firepower assigns an impact of 0 (Unknown) to the event.
- A rule triggers on a packet targeted for port 80 on a host and the IP address of this host is present in a discover rule in the Network Discovery policy. So Firepower looks in the host database and finds that there's no entry there. Firepower assigns an impact of 4 (Unknown Target) to this event.
- A rule triggers on a packet targeted for port 80 on a host and the IP address of the host is present in a discover rule in the Network Discovery policy. Firepower looks in the host database and finds a host entry, but the host doesn't show a listening server (service) on port 80. Firepower assigns an impact of 3 (Currently Not Vulnerable) to the event.
- A rule triggers on a packet targeted for port 80 on a host and the IP address of the host is present in a discover rule in the Network Discovery policy. Firepower looks in the host database and finds a host entry, and the database shows a listening server on port 80. Firepower assigns an impact of 2 (Potentially Vulnerable) to the event.
- A rule triggers on a packet targeted for port 80 on a host and the IP address of the host is present in a discover rule in the Network Discovery policy. Firepower looks in the host database and finds a host entry, plus a listening server on port 80. Firepower cross-references the listed vulnerabilities on the host with the vulnerability the rule was written for based upon the vulnerability ID. There's a match—the rule was written for a vulnerability present on the host. Firepower assigns an impact of 1 (Vulnerable) to the event.

So as you can see, Firepower is actively using the data it knows about the host to set an impact level for the event. The purpose of this is to help the analyst pick out the events most likely to be important to them, which is a key ingredient in Firepower's "secret sauce": The ability to dynamically calculate the impact for each event relative to the hosts on *your* unique network—Sweet!

Hard-Coding Impact

There's another thing you should know about impact... It can also be "hard-coded" into a rule. The fact that Firepower calculates the impact as I just described helps determine if a responder (server) is vulnerable to an attack. But what if a host—most likely a workstation— encounters malware? Well, the "vulnerability" in this case is probably sitting right there in front of the keyboard, so it's a given that this vulnerability always exists! Because of this, rules written to look for evidence of malware will generally have impact level 1 written into them. In practice, I find very few false positive events from rules written to detect evidence of malware in outbound traffic. If you see events with names starting with MALWARE, you can be fairly confident the source host is truly running the malware identified by the rule. Because of this, hard-coding impact level 1 into the rule is probably a very good idea.

So I know that was a pretty lengthy talk on priority and impact, but hopefully it helps you see why workflows displaying impact are so useful. The most important events are likely those with both impact level 1 and high priority. By no means are these the only important events, but it's a great place to start!

Drilling into Events

Impact and priority help us narrow down our events to the most critical, but we still need to dig into these and find out just how bad they really are, right?

There are several methods for drilling down to the next page in this workflow and this table describes what will happen for each one of them:

Method

Click one of the links in the table.

We can click any of the items in any column except Count.

Result

Drill to the next workflow page

searching on the item. For example, clicking on a source IP address will load

the next page revealing all the events from that source IP.
Click the blue down-arrow icon on the far left.

Select one or more check boxes, then click the View button at the bottom of the page.

Click the View All button at the bottom of the page.

Click one of the links in the workflow pages list at the top of the page. Drill to the next workflow page showing only that event row.

Drill to the next workflow page revealing the selected event rows.
Drill to the next workflow page revealing all of the events.

Same as above. Drill to the next workflow page revealing all of the events.
Pro tip: You can skip workflow pages using this method.

The following examples refer back to the figure that includes our sample list of intrusion events.

If we wanted to drill into the next page with all the Impact 1 events, click any



of the Impact 1 icons on the left. They look like this: . If you're interested in all the events from source IP 10.0.0.26, click on that address in the Source IP column.

If you're interested in a particular event regardless of the source IP, click on the event message.

And finally, if you're interested in only a specific event row, click the blue down arrow (



) on the far left for the respective row. But let's say we're interested in all the events that have been triggering for a particular source IP.

Here is the view if we decided to click on the source IP of 10.0.0.26. First thing to notice here is that the workflow links at the top of the page have changed. The red bold link is now Drill Down of Source and Destination IPs.

This indicates we're viewing this page.

Also, the columns have changed. This workflow removes some of the other columns and adds Destination IP. Now we can see that all of the events from this source IP involve two destination IP addresses. Like a detective, the analyst gathers all of the data points or context around the event. Knowing these events all involve two external hosts is useful. What do we know about 10.0.0.26? The blue computer icon

(



) tells us that there's an entry in the Firepower host database with more information on this host. This is because the 10.0.0.0/8 range is included in the Network Discovery policy.

To build more context, we can learn about this source or destination if Firepower has been collecting discovery data. Clicking on the computer icon next to an IP address will get us a pop-up with the host entry for that address.

Check out the following very long figure, which shows the result of clicking this icon:

From this screen we can see that Firepower believes this is a Windows XP host. The red outlined box under Indications of Compromise indicates that the host has triggered an intrusion even in the **malware-cnc** category. As a result, Firepower warns that it may be under remote control. The first seen and last seen times/dates indicate when we first saw this type of event on the host and when it was last seen.

We can see several listening servers on the host, including a web server identified as Microsoft-IIS 5.0. Four applications are listed as well.

Now that we have some more information, let's return to the Impact and Source page of our workflow to take a closer look at one of the events for our 10.0.0.26 host. The first two events are in the MALWARE-CNC category, so they were likely the cause of the Indications of Compromise events we saw in

the host entry. The third event is in the BLACKLIST DNS category and tells us the host performed a DNS lookup for a known malicious hostname. This one appears to be associated with the Win.Trojan.Meac malware.

At this point it's looking like our host at 10.0.0.26 has encountered at least one and possibly several malware variants. Even though our IPS probably dropped this traffic, this is still a pretty good indication that the host is compromised and requires some type of remediation. In most cases, the only way to know for certain that an infection has been eradicated is to reimage the host. If your job is limited to intrusion event analysis, you'd probably forward this information to the appropriate incident response team for follow-up.

Okay—so now let's drill even deeper into these events to confirm that they're not false positives. The next workflow page is called Table View of Events. If we select the top event for further analysis and click the blue arrow on the left, it'll bring us to the table view with just the 16 events on that row.

Time	Priority	Impact	Inline Result	Source IP	Source Country	Destination IP	Destination Country	Source Port / ICMP Type	Destination Port / ICMP Code	SSL Status
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1261 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1259 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1256 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1252 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1250 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1249 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1242 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1240 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1239 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1238 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1236 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1235 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1233 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1232 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1230 / tcp	443 (https) / tcp	Unknown
2016-11-08 21:31:05	high	High	↓	10.0.0.26	CHN	223.243.55.48	CHN	1229 / tcp	443 (https) / tcp	Unknown

This view is much wider than almost any screen resolution, so you'll have to

scroll through the page horizontally. Not all the fields in an event entry are populated; for instance, by looking at the figure above, you can see the source country isn't listed. This is because geolocation only works on public routable addresses and the events in the figure are sourced from my own network, which uses non-routable private IP address ranges.

Table view pages contain every field available in the event. You'll notice fields such as Source and Destination Port, VLAN ID, Classification, Application Protocol, Ingress/Egress Interface, Device, etc. That's a lot, but it still isn't all the information available. By default, intrusion event table views still hide some information. So to reveal the hidden fields, click the small black triangle to the left of Search Constraints as shown in the following figure:

Here, you can see we've received a nice, long list of links under the Disabled Columns heading. These are intrusion event fields that are less commonly used and disabled by default. To enable any of these columns for the duration of the session, just click on the link and it will be placed back into the table view. And if you find a column that's useful, you can create a custom workflow to include it.

The screenshot shows a network security dashboard with the following components:

- Navigation:** Overview, Analysis (selected), Policies, Devices, Objects. Sub-navigation includes Context Explorer, Connections, Intrusions > Events, Files, Hosts, Users, Vulnerabilities, Correlation, Custom, Lookup, and Search.
- Event Details:**
 - Impact and Source:** (switch workflow)
 - Drill Down of Impact and Source > Drill Down of Source and Destination IPs > **Table View of Events** > Packets
 - Search Constraints (Edit Search)
 - Table:**

Destination IP	Message	Source IP
223.243.55.48	MALWARE-CNC Win.Trojan.Taldoor variant outbound connection (1:20204:5)	10.0.0.26
- Disabled Columns:**
 - Application Protocol Category
 - Application Protocol Tag
 - Client Category
 - Client Tag
 - Destination User
 - Email Attachments
 - Email Recipient
 - Email Sender
 - HTTP Hostname
 - HTTP Response Code
 - HTTP URI
 - MPLS Label
 - Original Client IP
 - Web Application Category
 - Web Application Tag
- Event Summary Table:**

Time	Priority	Impact	Inline Result	Source IP	Source Country	Destination IP	Destination Country	Source Port / ICMP Type	Destination Port / ICMP Code	SSL Status
2016-11-08 21:31:05	high	1	↓	10.0.0.26		223.243.55.48	CHN	1261 / tcp	443 (https) / tcp	Unknown

In the real world, a table view probably won't give us a whole lot of insight into the event. After all, it's the packet and rule that will really reveal why this event triggered and what our response should be.

So from here, let's take a look at the packet view. From the table view we've got a lot of options for drilling into the individual packets, and one of them is clicking on the blue drill-down arrow. This'll take us to the packet view for that single, specific packet. Another way is to click on the Packets link near the top of the page, which will take us to the packet view and allow us to step through all the packets.

Impact and Source (switch workflow)

Drill Down of Impact and Source > Drill Down of Source and Destination IPs > Table View of Events > **Packets** 2016-11-08 17:19:00 - 2016-11-09 21:22:43 Expanding

Search Constraints [\(Edit Search\)](#)

Event Information

Event	MALWARE-CNC Win.Trojan.Taidoor variant outbound connection (1:20204:5)
Timestamp	2016-11-08 21:31:05
Classification	A Network Trojan was Detected
Priority	high
Device	FTDv 6.1 Transparent
Ingress Interface	Inline-IPS
Egress Interface	Main-IPS
Source IP	10.0.0.26
Source Port / ICMP Type	1261 / tcp
Destination IP	223.243.55.48
Destination Port / ICMP Code	443 (https) / tcp
Destination Country	CHN
Intrusion Policy	Secure IPS
Access Control Policy	vFTD 6.1 Transparent
Access Control Rule	IPS+File Inspection
Rule	alert tcp \$HOME_NET any -> \$EXTERNAL_NET \$HTTP_PORTS (msg:"MALWARE-CNC Win.Trojan.Taidoor variant outbound connection"; flow:to_server,established; content:".php?id=0"; nocase; content:"111D30"; fast_pattern; nocase; http_uri; pcre:"/^[a-z]{5}\.php?id=0\d{5}111D30[a-zA-Z0-9]{6}\$\/"; metadata:impact_flag red, policy balanced-ips drop, policy security-ips drop, service http; reference:cve,2011-0611; reference:url,contagiodump.blogspot.com/2011/06/jun-22-cve-2011-0611-pdf-swf-fruits-of.html; reference:url,www.virustotal.com/file-scan/report.html?id=145d64f38564eafa4fb5da0722c0e7348168024d32ada5cfb37a49f5811cb6b8-1315612892; classtype:trojan-activity; sid:20204; rev:5;)
Summary	This event is generated when an attempt is made to exploit a known vulnerability in acrobat.

Actions

Packet Information

FRAME 1 (Expand All)

- ▶ Frame 1: 248 bytes on wire (248 bytes captured (1984 bits))
- ▶ Ethernet II (Src: 00:0C:29:8F:D5:6D, Dst: 00:0C:29:C0:99:85)
- ▶ Internet Protocol Version 4 (Src: [10.0.0.26](#), Dst: [223.243.55.48](#))
- ▶ Transmission Control Protocol (Src Port: 1261 (1261), Dst Port: 443 (443), Seq: 1, Ack: 1, Len: 194)
- ▶ Packet Text
- ▶ Packet Bytes

Displaying row 1 of 1 rows << Page 1 of 1 >>

Copy	Delete	Review	Download Packet
Copy All	Delete All	Review All	Download All Packets

Clicking either one will take you to a packet view, as seen above. This page contains such a wealth of information about the event that I'm going to take you through it one section at a time.

Event Information

Most of the information here is already available in one of the previous workflow pages. We already know what the source and destination IP are. But it just might reveal some fields that are hidden in the default table view, like HTTP URI, and we can use this information for building context around the event and learn more about the type of traffic we're looking at. In this case, the fact that the source port is above 1024 and the destination port is 443 gives us a nice clue about what kind of packet this is. Traffic destined for port 443 is almost always from an HTTP client to a (secure) web server. Keep that destination port in mind because it also tells us something else about this packet that I'll reveal a bit later.

Rule

Even though it's part of the Event Information section, the Snort rule warrants its own discussion here. Because this is the actual rule that matched one or more packets, understanding exactly what this rule was looking for will really help us figure out how to respond to this event. I zoomed into the rule portion. Check it out:

```
alert tcp $HOME_NET any -> $EXTERNAL_NET $HTTP_PORTS (msg:"MALWARE-CNC Win.Trojan.Taidoor variant outbound connection";
flow:to_server,established; content:".php?id=0"; nocase; content:"111D30"; fast_pattern; nocase; http_uri; pcre:"/^[a-z]{5}\.php\?
id=0\d{5}111D30[a-zA-Z0-9]{6}$/Ui"; metadata:impact_flag red, policy balanced-ips drop, policy security-ips drop, service http;
reference:cve,2011-0611; reference:url,contagiodump.blogspot.com/2011/06/jun-22-cve-2011-0611-pdf-swf-fruits-of.html;
reference:url,www.virustotal.com/file-scan/report.html?id=145d64f38564eafa4fb5da0722c0e7348168024d32ada5cfb37a49f5811cb6b8-
1315612892; classtype:trojan-activity; sid:20204; rev:5; )
```

Let's walk through the rule at a high level to understand more about why it might match a packet.

- The rule header indicates that it's designed to inspect TCP traffic coming from the HOME_NET on any port going to the EXTERNAL_NET on HTTP_PORTS. So basically, it's designed to inspect outbound web browsing from our network.
- The message says this is a MALWARE-CNC category rule designed to detect an outbound connection from malware called Win.Trojan.Taidoor.
- The `flow` keyword restricts the rule to evaluating only traffic that's going from a client to a server. It also only evaluates packets that are post three-way handshake—this rule ignores the SYN, SYN/ACK, and ACK. Now this keyword really just means that the traffic qualifies for further inspection. Since you'll never see a user agent in a SYN, SYN/ACK, or ACK packet and we don't care about server responses (for this rule), there's no reason to waste resources inspecting them.
- Next we have a content check. This is looking for the string `.php?id=0`. The `nocase` keyword indicates the content check isn't case sensitive. Another content check follows looking for `111D30`. This one is also case insensitive and is looking only in the HTTP URI portion of the packet.

After the content checks comes the `pcre` keyword, which performs an additional regular expression match on the content. We won't walk through the entire expression here, but this is actually performing additional checks on the content items that have already been located. Using a regular expression is a more powerful method of searching for content that might deviate slightly but still follows certain predictable patterns. The reason this `pcre` keyword is placed after the content checks is so they can “prequalify” the packet. The `pcre` keyword won't be evaluated unless both of the content checks return true. Regular expression keywords are CPU intensive, and placing `pcre` after the `content` keywords ensures that Snort only performs this check on packets that already appear to be qualified candidates.

There are no other detection keywords in the rule after the `pcre` keyword. The rest are documentation and metadata, which don't look for specific content within a packet.

Okay—now going back to the destination port for this packet, remember it

was port 443? Is there anything strange about that? We know that port 443 is typically used for Secure Sockets Layer (SSL) web communications and that is encrypted traffic. So how was Snort able to inspect the content in a packet that was destined for this port? Clearly, there was some SSL decryption occurring to enable Snort to see this traffic. Either Firepower or some other SSL decryption solution apparently decrypted this session so we could inspect it and find this malware. This really demonstrates the value of SSL decryption when it comes to IPS usage because without it, we wouldn't have a Snort event in this case!

The next section is Actions, but even though it comes next on the screen, it's not always the next place you look. In this case, we're going to take a look at the packet itself first, and that's below Actions, in the Packet Information section.

Packet Information

If you've used a packet analysis tool such as Wireshark, this section will look familiar. Each section of the packet is shown from the Ethernet frame down to the data bytes. Most of the information in the packet header has already been pulled out for us in the workflows. What we really want to examine here is the data portion, or payload, of the packet because that's where the evil is!

The next figure shows the Packet Information portion of the screen.

You can see that the various sections of the packet are here. Clicking the black triangle icon to the left will expand a section's contents:

Packet Information



We already have the information like the source and destination IP address, ports, protocol, etc. What we really want to see is the packet payload and the

next figure reveals the packet with both the Packet Text and Packet Bytes sections expanded:

Packet Information

The screenshot displays a network packet analysis tool interface. At the top, it shows 'FRAME 1 (Expand All)'. Below this, several protocol layers are listed with their respective details: Frame 1 (248 bytes on wire), Ethernet II (Source: 00:0C:29:8F:D5:6D, Destination: 00:0C:29:C0:99:85), Internet Protocol Version 4 (Source: 10.0.0.26, Destination: 223.243.55.48), and Transmission Control Protocol (Source Port: 1229, Destination Port: 443, Sequence: 1, Acknowledgment: 1, Length: 194). The 'Packet Text' section is expanded, showing the raw data of the packet: '..).....).m..E.....E.....', followed by the HTTP request: '.....70.....F..h..P..r>..GET /gmzlk.php?id=031870111D309GE67E HTTP/1.1', and the headers: 'User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1)', 'Host: 211.234.117.141:443', 'Connection: Keep-Alive', and 'Cache-Control: no-cache'. The 'Packet Bytes' section is also expanded, showing a hex dump of the packet data. The hex dump consists of three columns: the offset in hexadecimal (0000 to 00f0), the hexadecimal byte values, and the corresponding ASCII characters. The ASCII characters match the raw data shown in the Packet Text section.

Of course, the Packet Text section is more user friendly, but if you really want to see what's in every byte of the payload, the Packet Bytes section shows you each of the byte values. Looking at this payload, it's clear that this rule probably hit pay dirt with this packet because it's clearly an HTTP GET containing the data the rule was designed to detect. This is hot evidence that the source IP has encountered the Taidoor Windows Trojan!

Actions

Moving up the page just a bit gets us to the Actions section, and if we expand this, we'll see a number of links available:

Actions ▾

Rule Actions

[View Documentation](#)

[Rule Comment](#)

[Edit](#)

[Disable this rule in the current policy \(Secure IPS\)](#)

[Set this rule to generate events in the current policy \(Secure IPS\)](#)

[Disable this rule in all locally created policies](#)

[Set this rule to generate events in all locally created policies](#)

[Set this rule to drop the triggering packet and generate an event in all locally created inline policies](#)

Set Thresholding Options ▶ in the current policy (Secure IPS)

▶ in all locally created policies

Set Suppression Options ▶ in the current policy (Secure IPS)

▶ in all locally created policies

Rule Documentation

To gain more insight into this rule and the alert, the Actions section contains a View Documentation link, and clicking this loads the rule documentation in a pop-up window.

Depending on the rule, there may be a lot of documentation or none at all.

In this case, the following figure indicates there's quite a bit of documentation for the Adobe Acrobat bug leveraged by the Taidoor Trojan.

Rule Documentation (1:20204:5)

This event is generated when an attempt is made to exploit a known vulnerability in acrobat.

Rule	alert tcp \$HOME_NET any -> \$EXTERNAL_NET \$HTTP_PORTS (msg:"MALWARE-CNC Win.Trojan.Taidoor variant outbound connection"; flow:to_server,established; content:".php?id=0"; nocase; content:"111D30"; fast_pattern; nocase; http_uri; pcre:"/^\V[a-z]{5}\.php?id=0\d{5}111D30[a-zA-Z0-9]{6}\$/UI"; metadata:impact_flag red, policy balanced-ips drop, policy security-ips drop, service http; reference:cve,2011-0611; reference:url,contagiodump.blogspot.com/2011/06/jun-22-cve-2011-0611-pdf-swf-fruits-of.html; reference:url,www.virustotal.com/file-scan/report.html?id=145d64f38564eafa4fb5da0722c0e7348168024d32ada5cfb37a49f5811cb6b8-1315612892; classtype:trojan-activity; sid:20204; rev:5;)
Impact	Denial of Service. Information disclosure. Loss of integrity.
Detailed Information	Adobe Flash Player before 10.2.154.27 on Windows, Mac OS X, Linux, and Solaris and 10.2.156.12 and earlier on Android; Adobe AIR before 2.6.19140; and Authplay.dll (aka AuthPlayLib.bundle) in Adobe Reader 9.x before 9.4.4 and 10.x through 10.0.1 on Windows, Adobe Reader 9.x before 9.4.4 and 10.x before 10.0.3 on Mac OS X, and Adobe Acrobat 9.x before 9.4.4 and 10.x before 10.0.3 on Windows and Mac OS X allow remote attackers to execute arbitrary code or cause a denial of service (application crash) via crafted Flash content; as demonstrated by a Microsoft Office document with an embedded .swf file that has a size inconsistency in a "group of included constants," object type confusion, ActionScript that adds custom functions to prototypes, and Date objects; and as exploited in the wild in April 2011.
Affected Systems	adobe acrobat 10.0 adobe acrobat 10.0.1 adobe acrobat 10.0.2 adobe acrobat 9.0 adobe acrobat 9.1 adobe acrobat 9.1.1 adobe acrobat 9.1.2 adobe acrobat 9.1.3 adobe acrobat 9.2 adobe acrobat 9.3 adobe acrobat 9.3.1 adobe acrobat 9.3.2 adobe acrobat 9.3.3 adobe acrobat 9.3.4 adobe acrobat 9.4 adobe acrobat 9.4.1 adobe acrobat 9.4.2 adobe acrobat 9.4.3

Scrolling down to the next large screen show, we find out even more information on attack scenarios and corrective action plus a bunch of Website Reference links for still more information:

Attack Scenarios	<p>Many types of buffer overflow exist, this is a generic term that may apply to many circumstances that result in an overflow of some kind. A parameter overflow for example, means that the attacker is able to supply data as a parameter to the execution of a program. When the program expands the supplied data, if the size of the parameter is not correctly checked, it may exceed a set limit allowing the attacker to overflow the buffer and write data into memory.</p> <p>In a stack overflow, the attacker has the opportunity to overwrite a return memory address which allows them to point the return address to a memory location containing code they wish to execute. This allows the attacker to run code with the full privileges of the program in use. The attacker may also supply the address for a known important call, for example the system() call, with the arguments to the call on the stack. The stack also contains the stack pointer and the frame pointer, overwriting these values may lead to a write-what-where condition.</p> <p>In a heap overflow, it is possible to overwrite function pointers that may be in memory. This may allow the attacker to execute code in memory by changing the function pointer to move to code of their choosing. This can occur even in programs that do not necessarily use function pointers since they may be left in memory at run time. The heap also contains user data which also becomes visible to the attacker.</p>
Ease of Attack	Simple. Exploits exist.
False Positives	None known.
False Negatives	None known.
Corrective Action	Upgrade to the latest non-affected version of the software.
	Apply the appropriate vendor supplied patches.
Contributors	<p>Sourcefire Vulnerability Research Team</p> <p>This document was generated from data supplied by the National Vulnerability Database. A product of the National Institute of Standards and Technology.</p> <p>For more information see http://nvd.nist.gov/</p>
References	<p>Common Vulnerabilities and Exposures Page</p> <p>Website Reference</p> <p>Website Reference</p> <p>Website Reference</p> <p>Website Reference</p>
View	Context Explorer

Close window

You'll find that rules written to detect exploits against known software vulnerabilities will usually have fairly good documentation. This is because the threat vectors and methods of exploiting a particular vulnerability are usually well known. Basically, the vulnerability is a stationary target, and we know what it takes to exploit it. As a result, we can document the conditions the rule was written to detect and provide links to vendor and/or Bugtraq research on the subject. What's more, vulnerabilities in commercial or open-source software will have Common Vulnerabilities and Exposures (CVE) references.

On the other hand, rules written to detect malware are generally not documented. The reason is the lack of documentation on the malware itself. Often we know the effect of the malware—such as changing the User Agent—but there may not be a whole lot available about documenting what the malware actually does or how it operates. In cases like this, search engines

are your friends. Our example event was a rare breed—a malware event that actually leverages a known bug in Adobe Acrobat. Many types of malware don't actually need a software vulnerability, just a cooperative user!

For the purpose of this example, I'm not going to dig into this particular rule and event much further. The bottom line is you would either treat it like a real malware event and follow up accordingly or decide it was a false positive.

Let's talk about how you would proceed if you determined this event was a false positive. You have four options:

- Disable the rule.
- Suppress the alert.
- Threshold the rule.
- Create a Pass rule.

Disable the Rule

If you decide the rule is too prone to false positives or that it's just not applicable to your environment, the best action to take is to disable it. This removes the rule from the applicable intrusion policy or policies. Referring back to the figure in the Actions section earlier, there are two links that will do this – “Disable this rule in the current policy (<policy name>)” and “Disable this rule in all locally created policies.”

Which of these you choose simply depends on which policies will be changed? If you only want the rule removed from the policy that generated this alert, pick the first option. If you want the rule disabled in all your policies, choose the second option. Clicking either one will result in a dialog at the top of the screen indicating what action was taken.

The next figure shows the result when I clicked on the first option to disable the rule only in the policy triggering the alert:



Now that the Intrusion policy has been updated, it will show up as “out-of-date” in the Intrusion Policy list view (Policies>Access Control >Intrusion). The rule will continue to trigger on any matching packets until policy changes are deployed to your devices.

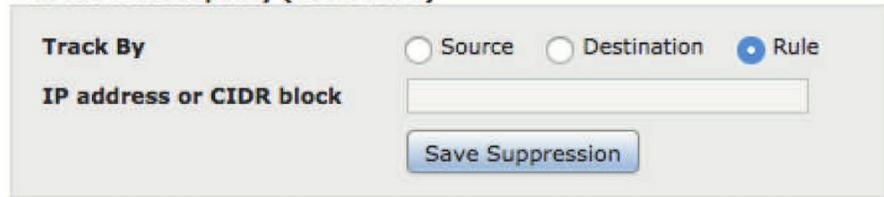
Suppress the Alert

Another option is to suppress the alert, which leaves the rule active but suppresses its output. This is normally done if you want to suppress the rule for a specific source or destination IP address. In our example, what if we found out that the host 10.0.0.26 just happens to use a business application that formats its HTTP GET requests the same way as our malware, but during our research, we decided that this is also an accurate rule to detect the Taidoor Trojan. We’re confident that 10.0.0.26 is the only host running the application, so we want to keep the rule active but stop it from alerting for 10.0.0.26. This is a perfect scenario for going with a suppression.

To suppress this rule for the single source IP address, we have two options: suppress for the current policy or for all policies. This works just like disabling the rule that we talked about earlier. For this example, we’re going to suppress in the current policy only. Clicking the upper black triangle under Suppression Options displays our choices:

Set Suppression Options

▼ in the current policy (Secure IPS)



Track By Source Destination Rule

IP address or CIDR block

Save Suppression

► in all locally created policies

To suppress a source IP address, select the Source radio button and enter the IP address or CIDR block into the field. You can suppress by source IP address, by destination IP address, or by rule. Clicking Save Suppression updates the applicable policy. Don't forget that the Intrusion or Access Control policy must be reapplied for this suppression to take effect.

“warning”

Do not Suppress a rule! Suppress the source/destination IP

Notes on Suppression

So even though suppression is a really cool tool, there are some important things you want to keep in mind:

- Suppressing a rule doesn't remove the rule from the rule set; it simply suppresses the alert output. If the rule action is set to drop, the rule will *silently drop the packet, providing no alert whatsoever*. Yes, I italicized that for a reason—it can create a troubleshooting nightmare! So only do this if you're absolutely sure the rule will *never* trigger a false positive alert. Suppression is primarily useful for rules that only generate events or in passive installations. For inline suppression, use a pass rule, which I'll cover later.

- If you select the Rule option under Track By, you'll cause the rule to continue evaluating traffic but never alert. In passive installations, this is just useless. For inline installations, it also runs the risk of dropping silently if it's a drop rule. So if you really want to suppress the entire rule, you probably should just disable it.

- You can suppress a source IP or a destination IP address but not both. If you need to suppress alerts for traffic from a specific IP address to another specific IP address, you need a pass rule.

Threshold the Rule

Thresholding is another option for reducing the naggy alerting for noisy rules. When you add a threshold, you cause the rule to not trigger until it matches a certain number of packets over a certain period. This is useful for rules that are designed to detect attacks such as brute force password guessing. Often, rules written to detect this activity will look for login failures. One or two failures from a single IP address may just be someone who mistyped their password and we don't need to be alerted to this. But dozens of failures from the same host probably indicates something wicked. To set a threshold on a rule, expand the appropriate section for current or all policies. In the next figure, I've expanded the current policy option:

Set Thresholding Options ▼ in the current policy (Secure IPS)

Type	<input type="radio"/> limit	<input type="radio"/> threshold	<input type="radio"/> both
Track By	<input type="radio"/> Source	<input type="radio"/> Destination	
Count	<input type="text"/>		
Seconds	<input type="text"/>		
	<input type="checkbox"/> Override any existing settings for this rule		
	<input type="button" value="Save Thresholding"/>		

▶ in all locally created policies

To configure thresholding, you've got following options:

- Type: Limit, Threshold, Both
 - Limit will limit notifications to the number of events specified by Count per the time period entered in Seconds below. For instance, alert no more than 5 times in 300 seconds regardless of how many packets actually match the rule.

- Threshold will provide one alert for each number of matching packets specified by Count during the time period. For example, provide one alert for each 5 packets within 300 seconds.
- Both provides a notification once per time period after matching the specified count of packets. For instance, alert once per 300 seconds if you see at least 5 matching packets.
- Track By: Source/Destination track matches to this rule by the source or destination IP address.
- Count: The count in packets matching the rule.
- Seconds: The number in seconds during which the count will be evaluated.
- Override Any Existing Settings for This Rule: Check this box to replace any existing thresholds with the new one. **Create a Pass Rule**

I've already talked about this in chapter 8 Objects and chapter 12 IPS policy. The last option for rule tuning is to create a pass rule, which comes in really handy when you need to suppress a rule based on both the source *and* destination IP addresses. It's also the only real option for suppressing a drop rule in inline mode—recall italics, don't make me do that again! Anyway, the idea of a pass rule is to copy an existing rule, then change the action from alert to pass. You also modify the source/destination IP and/or port to match your specific suppression.

Without going too far into Snort rules, I want you to understand that a rule can have two actions in Firepower: alert and pass. Alert is fairly self-explanatory—do something when traffic is matched. But what does it mean to pass? A rule action of pass tells Snort to ignore this packet. After matching a pass rule, Snort will “pass” the packet, skipping all the remaining rules. Another important point is that regardless of where a pass rule appears in the rule set, Snort will always process traffic through pass rules first. This means that if you have two rules looking for the same thing and one of them is a pass rule, the pass rule will always win.

We can take advantage of this behavior to suppress certain rules. If a rule is generating a false positive for specific traffic, we can take the following

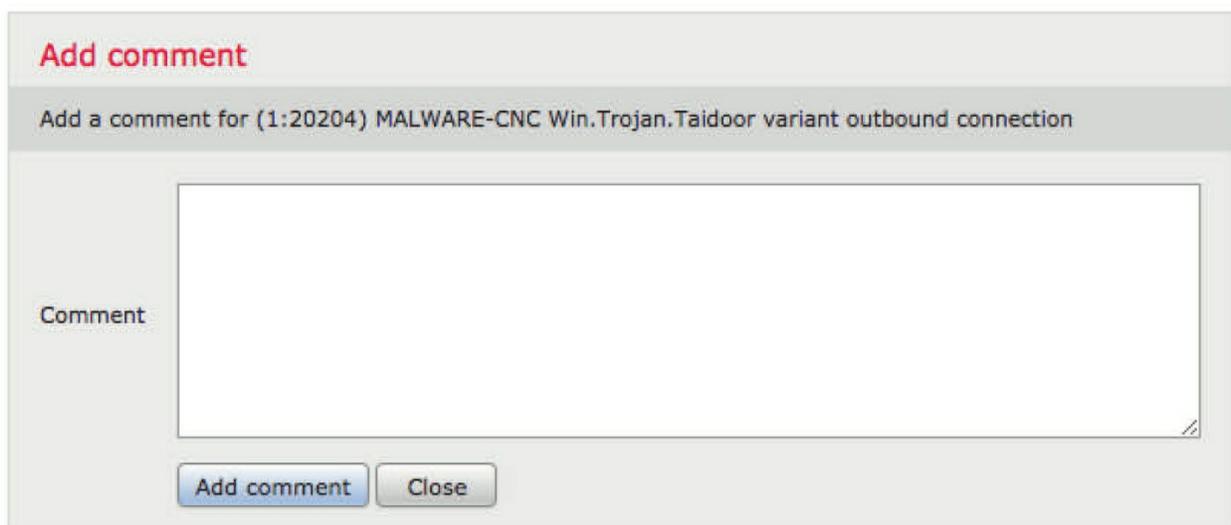
steps:

1. Make a copy of the rule.
2. Modify the rule header to match just the specific false positive hosts(s).
3. Change the action to pass.

Now this pass rule will only evaluate traffic between our specific hosts, and if packets match the rule, (which has the same detection options as the original rule), this traffic will skip all the remaining rules, including the original one.

Rule Comment

The last item I want to cover under Rule Actions is Rule Comment. Clicking this link will bring up the Add Comment page in a pop-up window:



The screenshot shows a pop-up window titled "Add comment" in red text. Below the title bar, there is a header text: "Add a comment for (1:20204) MALWARE-CNC Win.Trojan.Taidoor variant outbound connection". The main area of the window contains a large, empty text input field with the label "Comment" to its left. At the bottom of the window, there are two buttons: "Add comment" and "Close".

This lets us enter free-form text comments regarding a rule and can be used to document your analysis actions. When a comment is added, the system also stamps the comment with the user and time/ date. Keep in mind that rule comments are rule specific, not policy specific, and that comments added here will be visible regardless of where the rule is viewed on the FMC.

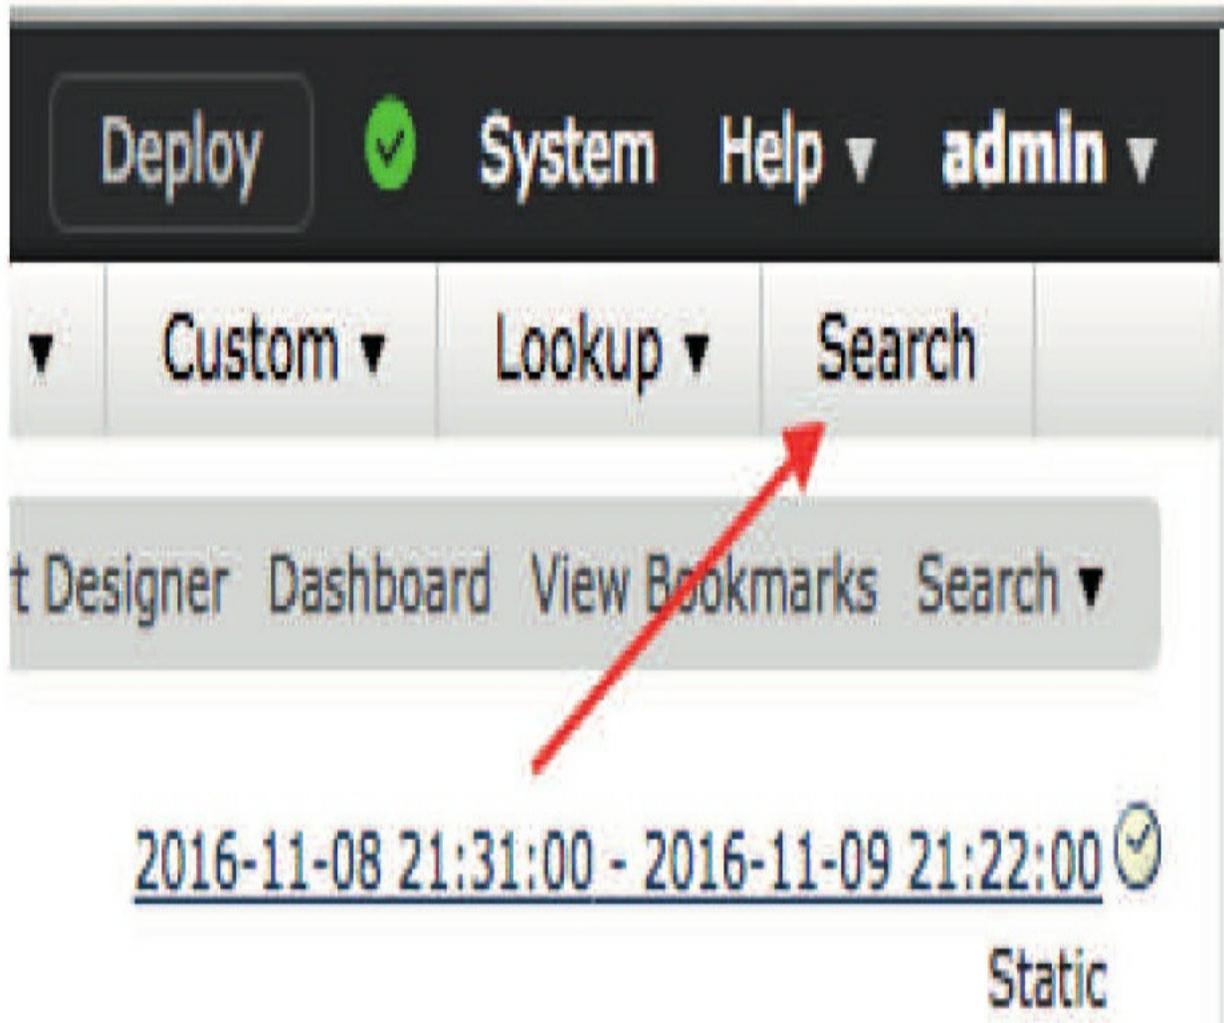
Miscellaneous Intrusion Event Analysis Features

Searching

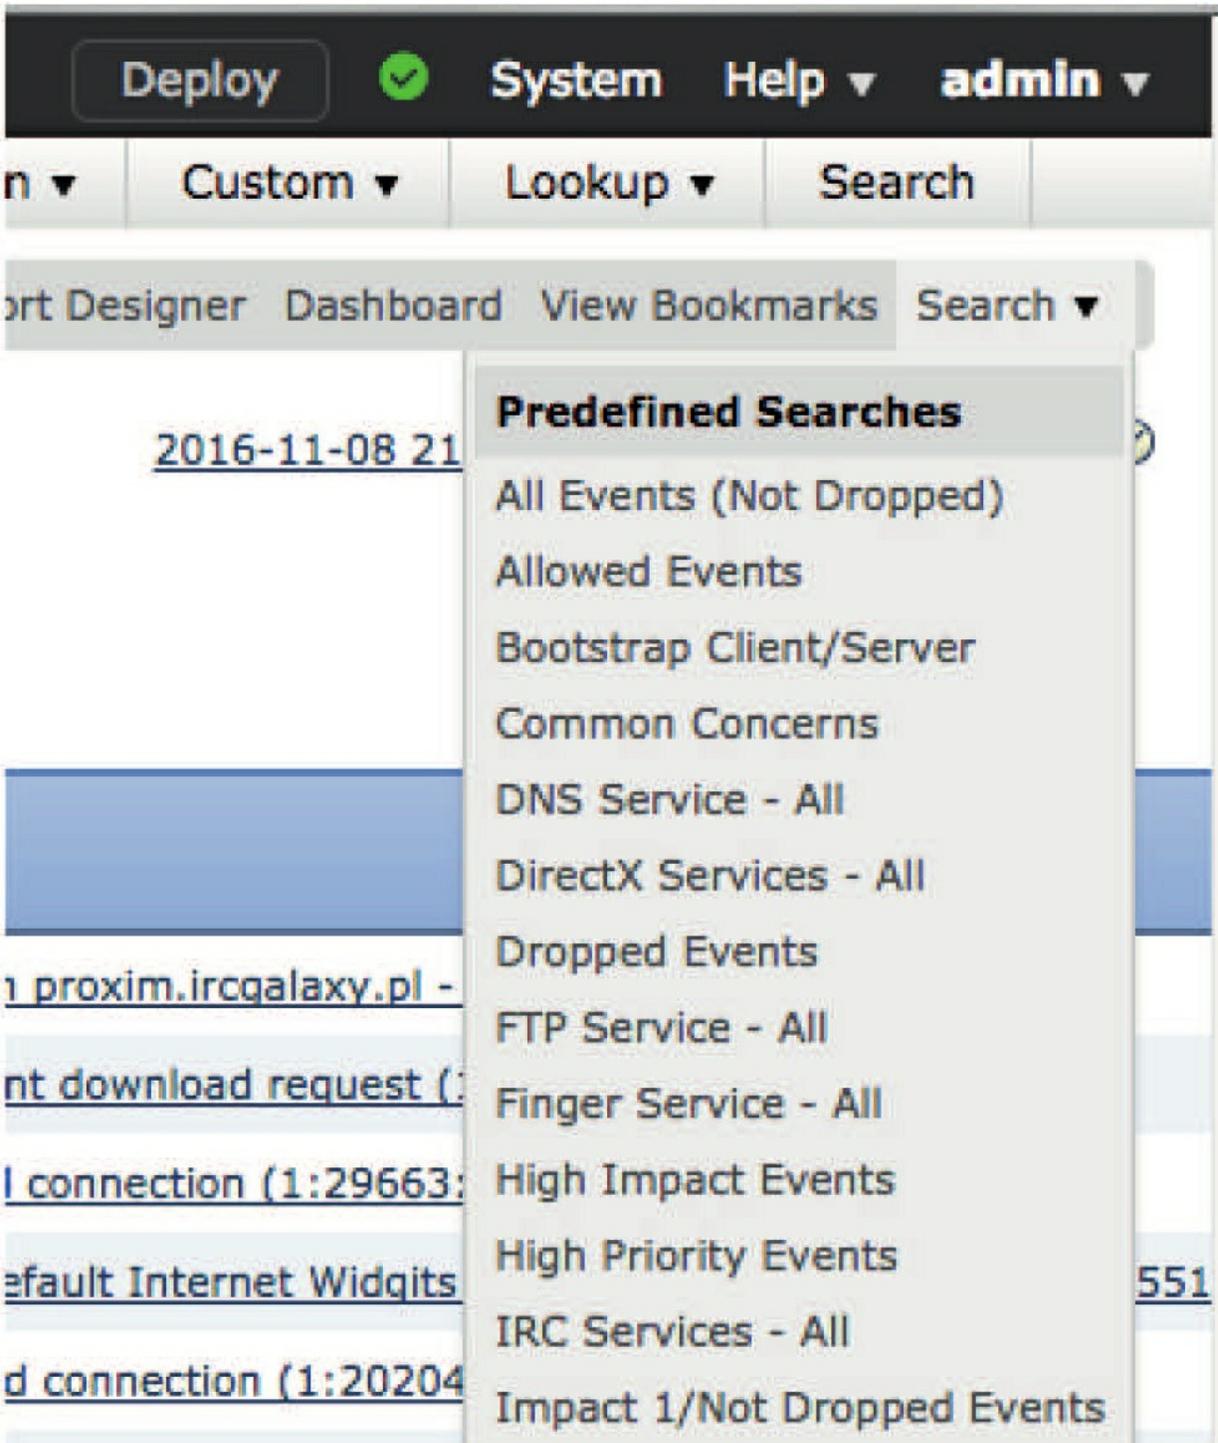
The ability to search through intrusion events is a really powerful feature!

The system allows you to search any field in an event to locate the ones that are relevant to you. From the intrusion event view, there are three ways to navigate to the search page.

The first one is to click Search in the menu bar in the upper right:



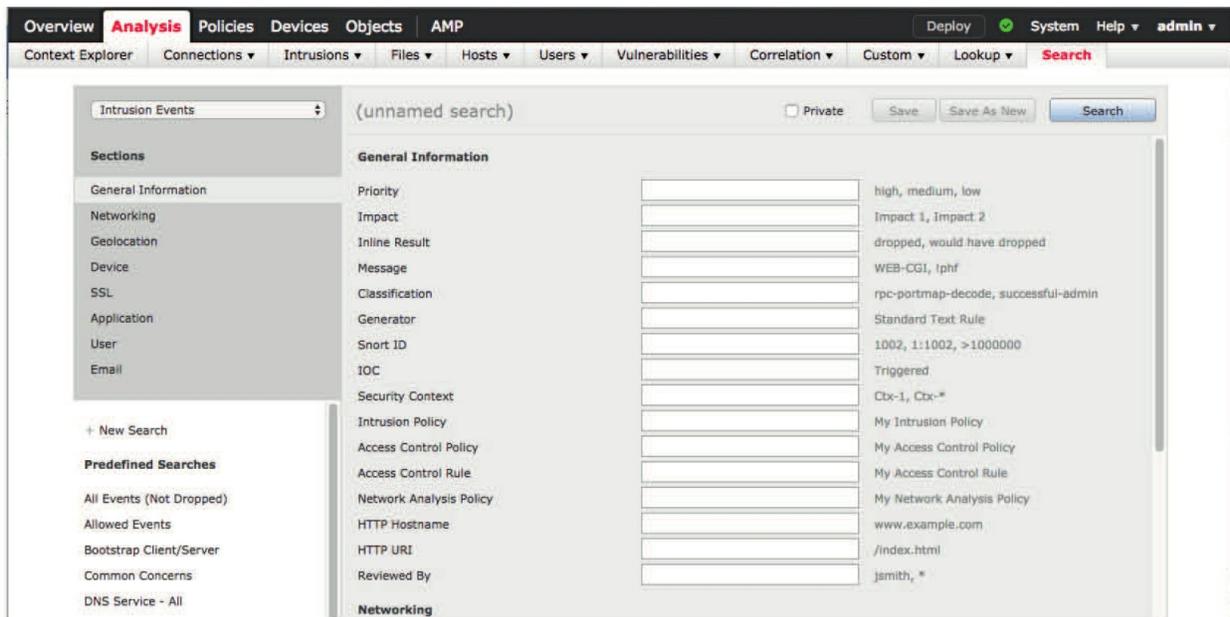
The second way is to use the quick link (also labeled Search) located just below and to the right of the menu bar as you can see in the figure above. Hovering over this link brings up your saved searches along with the builtin searches, which I've shown in the next figure. From here, you can click one of these to execute that search or click Search itself to load the search page:



The third method for entering the search page is to click the Edit Search link on the upper left side of the screen. This also has the effect of editing an existing search query, while the two previously mentioned methods won't preserve any current search conditions. Check it out:



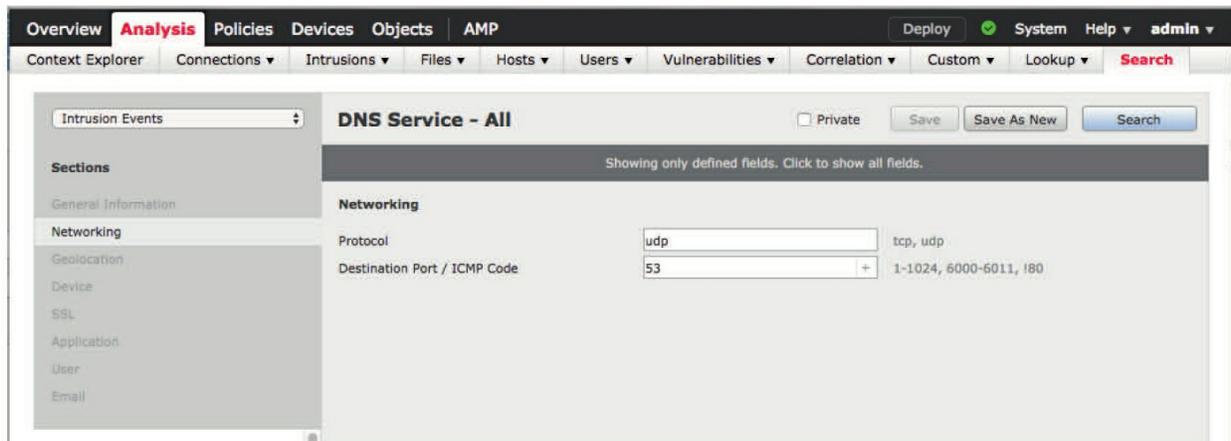
The search page itself, pictured below, is a long list of all the fields you can search on in an intrusion event. The system comes with a number of predefined searches. You can't edit or erase these default searches, but you can create your own. On this page you'll see predefined searches on the left, which is also where you'll also find any custom searches you've created:



The Sections list on the left is a quick way to scroll to the various search fields. This is a handy way to find what you're looking for without scrolling down through the list of search fields, which is quite long.

To load an existing search query, select it from the list on the left. You can then modify the search criteria if desired. You will notice that only the fields

containing search values appear when you are editing an existing search. In the next figure, I am editing a search that looks for UDP packets on port 53.



If you want to add additional search criteria, click the dark bar near the top that says “Showing only defined fields. Click to show all fields.” Doing so will re-enable all the search fields so you can add your additional criteria as needed.

Saving searches is as easy as giving your search a name and clicking the Save or Save As New button. If the Private check box is selected, no other non-admin users will be able to use your searches. To name your search, click unnamed search near the top of the search page shown on the preceding page.

The various search pages in the system also have sample search information to the right of each field. This helps users quickly understand the type of data that’s appropriate for each field. Some fields also have a gray + to the right. Clicking this will load any objects you’ve created. The next figure shows the network objects that resulted from clicking the icon to the right of the Source IP field:

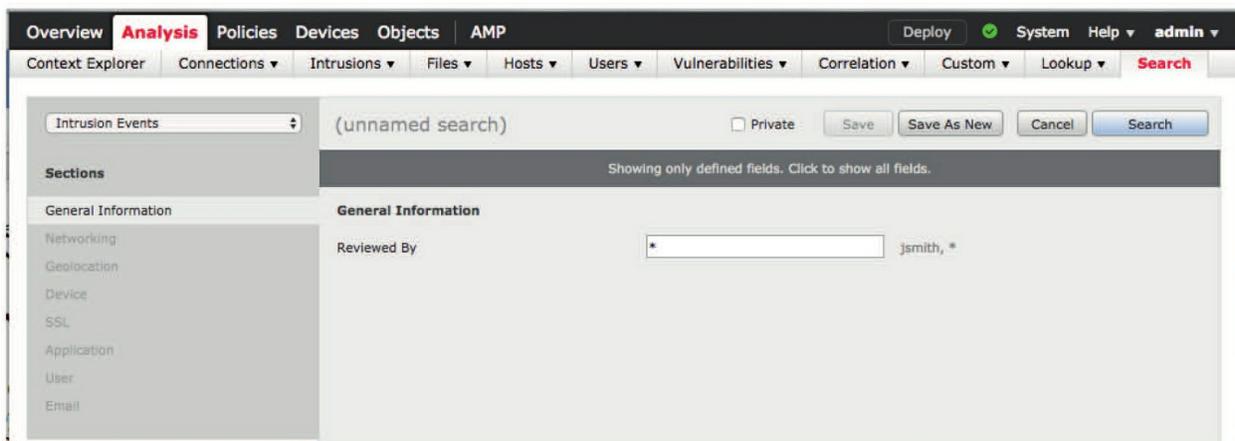
Networking

Source IP		+	192.168.1.0/24, 192.168.1.3, 2001:db8:8...
Destination IP	#{IPv6-IPv4-Mapped}		192.168.1.0/24, 192.168.1.3, 2001:db8:8...
Source / Destination IP	#{IPv4-Private-All}		192.168.1.0/24, 192.168.1.3, 2001:db8:8...
Original Client IP	#{FMC-54_1}		192.168.1.0/24, 192.168.1.3, 2001:db8:8...
Protocol	#{IPv4-Private-192.168.0.0-16}		tcp, udp
Source Port / ICMP Type	#{WDEX2-NAS_1}		1-1024, 6000-6011, 180
Destination Port / ICMP Code	#{IPv6-Private-Unique-Local-Addr}		1-1024, 6000-6011, 180
VLAN ID	#{Ubuntu-ESX-Desktop-2}		10
MPLS Label	#{any-ipv4}		1044495
Ingress Security Zone	#{FMC-61_1}		My Security Zone
Egress Security Zone		+	My Security Zone
Ingress / Egress Security Zone		+	My Security Zone
HTTP Response Code			200

When you're ready to execute your search, just hit Enter or click the Search button in the upper right.

Reviewed Events

Another very cool Firepower feature is the ability to place events into Reviewed status. This is super helpful if you've reviewed an event, made a determination regarding its viability, and want to continue to analyze additional events. To keep track of events you've already evaluated, mark them as Reviewed, which effectively removes the event from your analysis workflow without deleting them from the database.



To review events, use the Review buttons at the bottom of the analysis view. Clicking a Review button will affect the selected events—the rows with the check box selected. The Review All button will place all the events currently listed in Reviewed status. Keep in mind that the All buttons apply to *all* the events selected by the current search query and time window, not just the ones visible on the page! Placing an event into reviewed status removes it from the intrusion event workflows for all users.

Once events have been placed in reviewed status, you can also search by name for events reviewed by a particular user or search for all reviewed events by using a * in the Reviewed By search field.

Once events have been placed into Reviewed status, you can view them two ways. One is to use the Reviewed By search I just mentioned, and another is to select **Analysis>Intrusions>Reviewed Events**, which is actually just a shortcut to loading the Reviewed By * search. You can test this by navigating

to the Reviewed Events view and then clicking the Edit Search link.

Doing so will reveal that the Reviewed Events view is actually a custom search of all reviewed events as shown in the following figure:

Packet Downloads

Now by default, when an intrusion event is triggered, Firepower saves the packet. During event analysis, these packets can be downloaded to the analyst's workstation for archival or further scrutiny with their packet analysis tool of choice. Each workflow page contains a Download button in the row of buttons at the bottom of the page. Clicking this will download the selected event packets. As with the other buttons in this row, the Download All Packets button works for all the packets returned by the time window and search query.

When multiple packets are selected for download, you get a zip archive containing multiple PCAP format files. But when you're drilled down to the packet view, or if you have only one event selected, clicking the Download Packet button downloads a single file in PCAP format. In this case, you can often just select the Open With action in your browser and select an application such as Wireshark to check out the packet's contents. If you're downloading multiple packets, you'll have to unzip the archive first.

If you want to keep a long-term record of packet data, the download feature is a good one to have in your quiver. Because of the circular nature of event storage, intrusion events will eventually be purged from the FMC and along with them that valuable packet data!

Security Intelligence Events

So, like intrusion events, Security Intelligence events are designed to trigger when there's evil about, but the process of identifying and tuning false positives for these events is a lot simpler. There are three types of Security Intelligence events:

- Network Security Intelligence: These are just packets traveling

to or from blacklisted hosts, which are identified based upon their source or

destination IP address.

- **DNS Security Intelligence:** These are DNS lookups for blacklisted domain names.
- **URL Security Intelligence:** These are blacklisted URLs.

Unlike it does with intrusion events, Firepower doesn't capture any packets when generating Security Intelligence events. Also, the concept of a false positive is a bit different than with Snort rules. A false positive Security Intelligence event represents a failure or inaccuracy in the intelligence used to generate the various blacklists. The remediation for a false positive is to either remove the offending entry from the blacklist or add it to a whitelist. If the blacklist is provided externally as it is with the built-in Cisco security feeds, whitelisting is your only option. Of course, you could always disable an entire intelligence category if you're seeing it generate legions of false positives!

Security Intelligence Workflows

There are three built-in workflows for checking out Security Intelligence events, and as it was with other event types, you can also create your own custom workflows to display the data in other meaningful ways.

- **Security Intelligence Events:** This workflow is the default and it reveals a few select fields that are specific enough that you'll only see one event per row. The workflow consists of one page named Security Intelligence with application details and a table view.
- **Security Intelligence Summary:** This is helpful for getting summary information about how many events are triggering in each Security Intelligence category. You can drill into application details and a table view.
- **Security Intelligence with DNS Details:** This is a workflow specific to DNS events. It contains summary information on the number of events in a category and a specific DNS query and also has a table view.

Security Intelligence Tuning

As I mentioned, there really aren't a lot of options when tuning Security

Intelligence events. There's also not a lot information on *why* a given event triggered. If you see an IP block event and the Security Intelligence category is Tor exit node, you can be sure of one thing: At the time the event triggered, the IP address was on the Tor exit node list. And because the Security Intelligence feed from Cisco is fairly dynamic, there's no guarantee that this IP address is still on the list. If you think this might be wrong, you could cross-reference with other sources of TOR exit node hosts, but beyond that, there's not a whole lot of investigation to do. If the blocked connection represents legitimate traffic, your best recourse is to whitelist it. The figure below shows a number of Tor exit node events:

The screenshot shows the Cisco AMP Security Intelligence Events interface. The table displays the following data:

First Packet	Last Packet	Action	Reason	Initiator IP	Initiator Country	Responder IP	Responder Country	Security Intelligence Category
2016-11-18 09:59:42		Sinkhole	DNS Block	10.0.0.158		208.67.222.222 (resolver1.opendns.com)	USA	DNS Malware
2016-11-18 09:48:29		Block	IP Block	10.0.0.163		95.211.205.151 (sh3llbox.hu)	NLD	Tor_exit_node
2016-11-17 17:29:21		Block	IP Block	95.130.11.147 (tor-exit-readme.analyzer.org)	FRA	10.0.0.11		Tor_exit_node
2016-11-17 17:29:10		Block	IP Block	85.248.227.164 (telliana.enn.lu)	SVK	10.0.0.11		Tor_exit_node
2016-11-17 17:29:01		Block	IP Block	163.172.137.174 (174-137-172-163.rev.cloud.scaleway.com)	FRA	10.0.0.11		Tor_exit_node
2016-11-17 16:52:38		Block	IP Block	10.0.0.163		185.100.86.100 (server.saveyourprivacy.is)	FIN	Tor_exit_node
2016-11-17 16:52:19		Block	IP Block	10.0.0.163		185.100.86.100 (server.saveyourprivacy.is)	FIN	Tor_exit_node
2016-11-17 16:50:44		Block	IP Block	10.0.0.163		95.211.205.151 (sh3llbox.hu)	NLD	Tor_exit_node
2016-11-17 16:50:20		Block	IP Block	10.0.0.163		95.211.205.151 (sh3llbox.hu)	NLD	Tor_exit_node
2016-11-17 16:24:45		Block	IP Block	10.0.0.163		178.217.187.39 (srv1044.htdedicated.pl)	POL	Malware
2016-11-17 16:24:31		Block	IP Block	10.0.0.163		178.217.187.39 (srv1044.htdedicated.pl)	POL	Malware

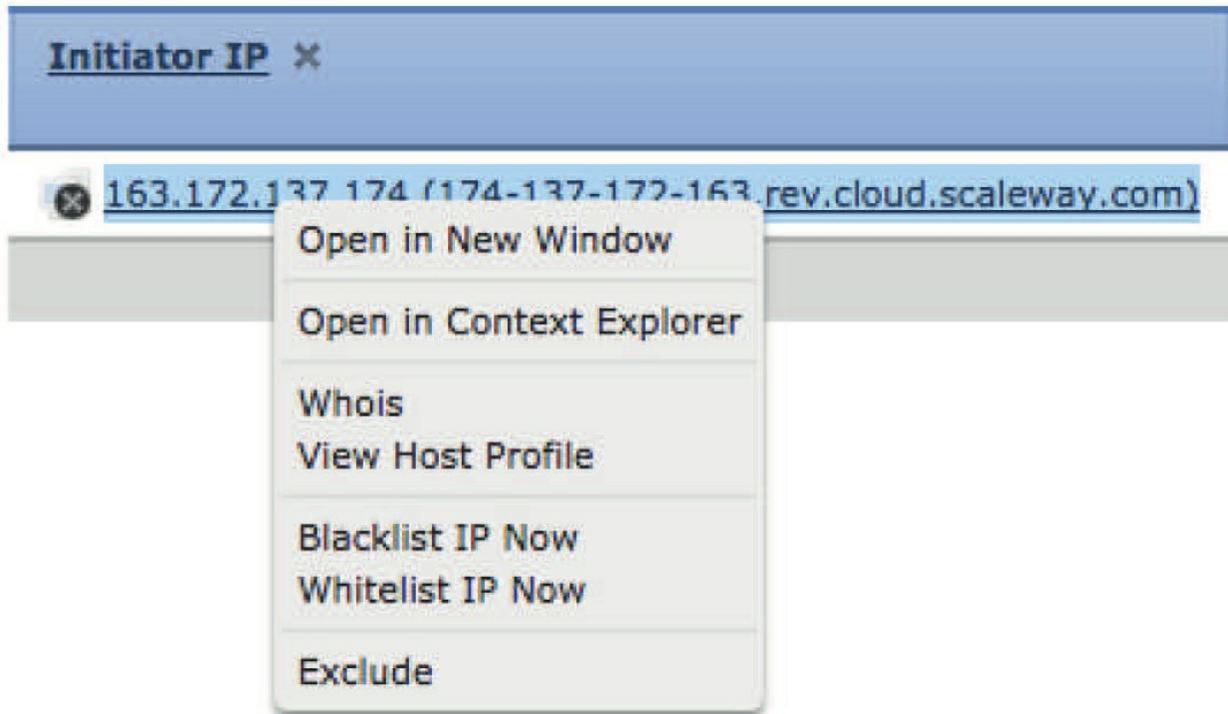
On the other hand, you may see an event involving a given IP address and decide you want to permanently block it. By adding the IP address to the Global-Blacklist, you ensure that even if it is eventually removed from the Cisco feed category, it'll still be blocked.

The options for whitelisting or blacklisting are a bit different depending on the type of Security Intelligence event.

Network Security Intelligence

To whitelist/blacklist an IP address, right-click on the IP address in *any* event. This applies not only to Security Intelligence but also to connection

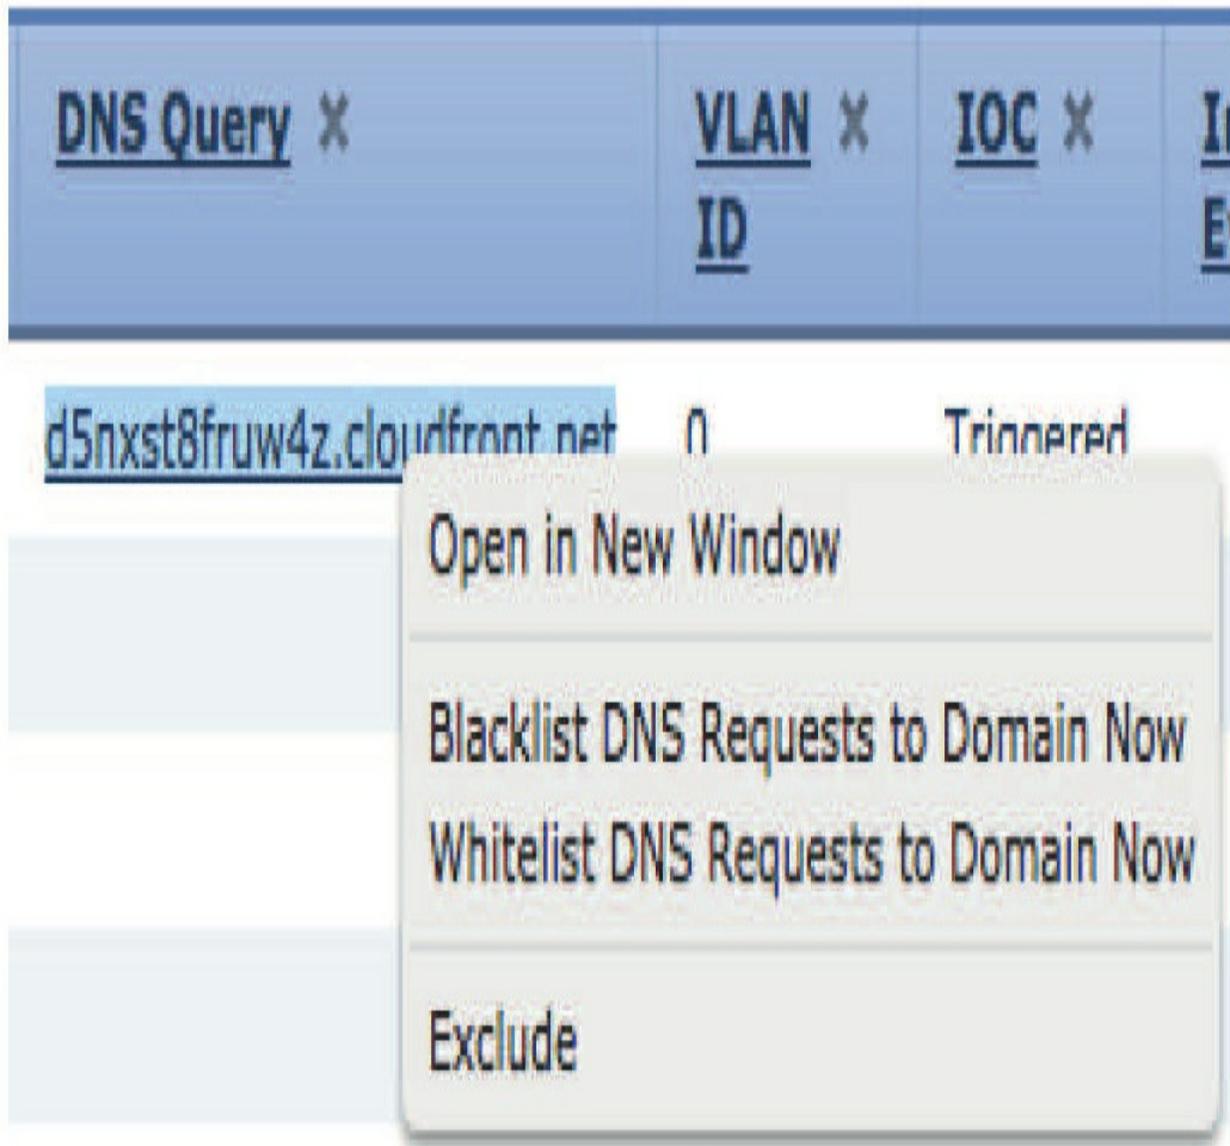
events, intrusion events, correlation events, etc. A rightclick will bring up the menu shown in the next figure:



The options Blacklist IP Now and Whitelist IP Now will add the IP address to the corresponding global list. After confirmation, this list is updated on *all* your devices immediately. No policy deployment is needed.

DNS Security Intelligence

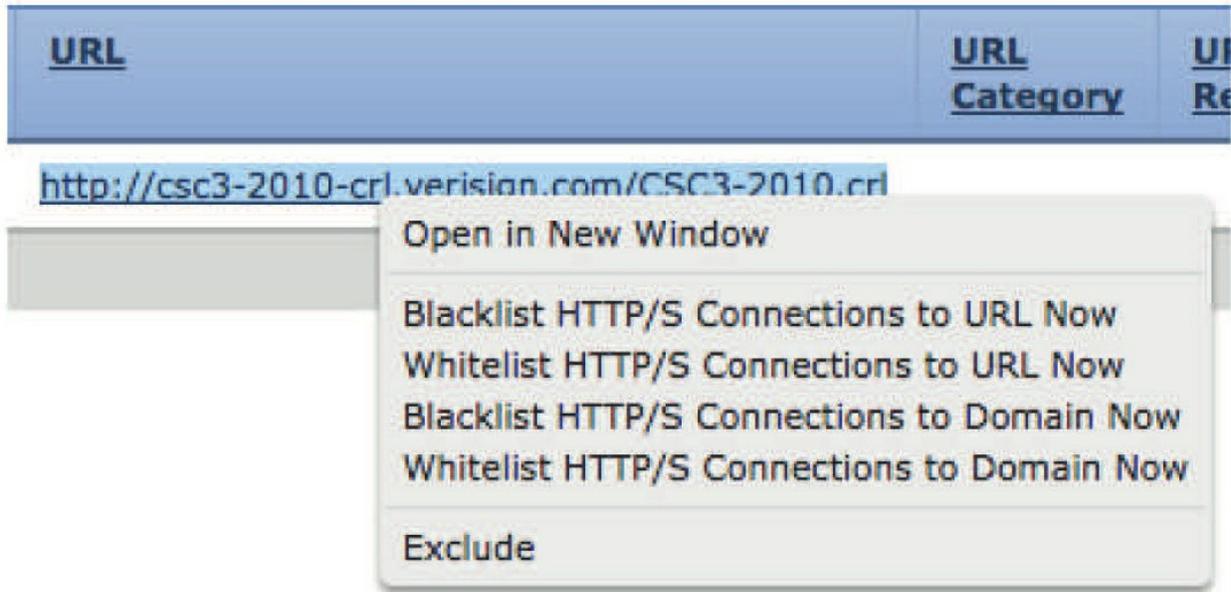
If you want to perform a blacklist/whitelist action for a DNS entry, just right-click the entry in the DNS Query column:



As it was with IP blacklisting, this entry will be added to the Global DNS Blacklist/Whitelist and updated on all your devices immediately.

URL Security Intelligence

In the same way, you can blacklist/whitelist URLs by right-clicking the entry in the URL column. The next figure demonstrates that you can choose to blacklist/whitelist just the URL or the entire domain:



One last thing: There are a whole bunch of categories in the Cisco Security Intelligence feed, ranging from darkly evil categories like Exploitkit, Malware, and CnC (Command and Control) to less sinister lists like Spam and Tor_exit_node.

If your organization has a public-facing website, know that if you block one of the less dangerous categories, you'll probably also impact legitimate traffic to your site. Just because someone uses a TOR-enabled browser doesn't mean they're demonic, although for some high-security sites, blocking this category could still be a good idea. In any case, you don't want to get into hot water for blocking legitimate connections because you implemented overzealous, paranoid Security Intelligence!

File and Malware Events

File/malware events is another category wherein you could come across some tuning challenges. Even though you probably won't have nearly as many false positives, you'll likely find there are still some wonderful opportunities to improve Firepower's file analysis behavior. Remember from Chapter 10 that Firepower uses several methods to identify malware files, with the primary approach is calculating a SHA-256 hash of the file. This hash is unique to a given file, meaning there is no chance of identifying the wrong file with this method, but that doesn't mean it's immune to false positives!

If a file is wrongly convicted by the Cisco cloud, it could be blocked even though it's not actually evil. In the real world, this is pretty rare, and it's usually identified and fixed by Cisco quickly, but it's not out of the question. If you find that a legitimate file is being wrongly convicted, you can correct this by adding the hash to a custom file list, a process covered back in Chapter 8, "Objects." You add the hash of the file to the Clean-List and then deploy your file policy. The CleanList is consulted by your devices before the cached disposition or cloud query, so any SHA-256 found there won't be blocked.

Even though false positive detections are rare, there's another reason you'll want to tune your file detection—performance! You'll find there are certain files that are being checked repeatedly and yet they always come back with an Unknown disposition.

The following figure shows a list of file events. Notice at the top that there are over 3,500 events for the BZ file type! Drilling into these events using the blue arrow on the left shows they are all from the domain canonical.com and were received by a single host—10.0.0.19.

Part of the table view for these events is shown in the next figure:

Bookmark This Page Report Designer Dashboard View Bookmarks Search

File Summary (switch workflow)

File Summary > **Table View of File Events**

2015-11-18 09:30:00 - 2016-11-19 09:59:20

Expanding

▶ Search Constraints [\(Edit Search\)](#)

Disabled Columns

Jump to...

<input type="checkbox"/>	Time <input checked="" type="checkbox"/>	Action <input checked="" type="checkbox"/>	Sending IP <input checked="" type="checkbox"/>	Sending Country <input checked="" type="checkbox"/>	Receiving IP <input checked="" type="checkbox"/>	Receiving Country <input checked="" type="checkbox"/>	Sending Port <input checked="" type="checkbox"/>	Receiving Port <input checked="" type="checkbox"/>
<input checked="" type="checkbox"/>	2016-11-19 06:39:22	Malware Cloud Lookup	91.189.91.26 (hanger.canonical.com)	USA	10.0.0.19		80	53787
<input checked="" type="checkbox"/>	2016-11-19 06:39:22	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:21	Malware Cloud Lookup	91.189.91.26 (hanger.canonical.com)	USA	10.0.0.19		80	53787
<input checked="" type="checkbox"/>	2016-11-19 06:39:21	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:21	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:20	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:20	Malware Cloud Lookup	91.189.91.26 (hanger.canonical.com)	USA	10.0.0.19		80	53787
<input checked="" type="checkbox"/>	2016-11-19 06:39:20	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:19	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:19	Malware Cloud Lookup	91.189.91.26 (hanger.canonical.com)	USA	10.0.0.19		80	53787
<input checked="" type="checkbox"/>	2016-11-19 06:39:19	Malware Cloud Lookup	91.189.91.26 (hanger.canonical.com)	USA	10.0.0.19		80	53787
<input checked="" type="checkbox"/>	2016-11-19 06:39:19	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:18	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:18	Malware Cloud Lookup	91.189.91.26 (hanger.canonical.com)	USA	10.0.0.19		80	53787
<input checked="" type="checkbox"/>	2016-11-19 06:39:18	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:18	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:17	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681
<input checked="" type="checkbox"/>	2016-11-19 06:39:17	Malware Cloud Lookup	91.189.91.26 (hanger.canonical.com)	USA	10.0.0.19		80	53787
<input checked="" type="checkbox"/>	2016-11-19 06:39:17	Malware Cloud Lookup	91.189.88.161 (keeton.canonical.com)	GBR	10.0.0.19		80	52681

Further analysis of the file type and file names reveals that these are update packages and the host 10.0.0.19 is an Ubuntu host.

Another common occurrence is detection of antivirus (AV) updates. For some antivirus products, these updates are distributed within a network via Windows DLL files. This can result in a large number of cloud lookups because the AV servers distribute these updates frequently. The SHA-256 values for these are always unknown, and they are never malicious. You'll likely see a large number of file lookup events and realize that the source of these files is your AV servers.

It's great that there are a number of ways to address these ubiquitous file events. One is to modify your File policy to ignore the file types in question, which will effectively stop the lookups for the AV updates but will also blind your system to any malicious files in the same category—not good at all!

Another approach is to find a way to disable file detection for the specific traffic. In the case of the AV update, sometimes they all source from specific AV servers or retrieve files from a predictable URL, plus the AV software uses a specific source port when sending the files. Armed with this information, you can modify Firepower's behavior to ignore just these specific file transfers—nice!

So to get this done, edit your Access Control policy.

We're going to add a rule specific to this AV update traffic. The policy in the following figure on the next page is a simple one that implements file and IPS inspection on all traffic:

In the policy, there's a single rule that enables both IPS and file inspection for all traffic. To modify the policy and address our AV file events, we're going to insert a rule above the existing rule that'll be designed to match only the AV update traffic. We'll also configure the rule to perform IPS inspection but not file inspection on this traffic.

Okay—first, we'll select the Allow action for the rule, then select the Networks tab and specify the source IP addresses of our AV servers. In the following figure, we've created a custom network object called AV-servers

that contains these IP addresses.

Add Rule

? X

Name AV updates IPS only

Enabled

Insert below rule

1

Action Allow



Zones

Networks

VLAN Tags

Users

Applications

Ports

URLs

SGT/ISE Attributes

Inspection

Logging

Comments

Available Networks

Search by name or value

Networks

Geolocation

10.0.0.64-28-NAT-Pool

7030-device

any

any-ipv4

any-ipv6

Arris-router

ASA-ESX-interface

AV-servers

Cable-One-Clients-160-3-24

Add To
Source
Networks

Add to
Destination

Source Networks (1)

Source

Original Client

AV-servers

Destination Networks (0)

any

Enter an IP address

Add

Enter an IP address

Add

Add

Cancel

Objects are preferred for this type of rule so they can be easily modified later if your list of antivirus servers changes.

Add Rule

? X

Name AV updates IPS only Enabled

Insert below rule

Action Allow

- Zones
- Networks
- VLAN Tags
- Users
- Applications
- Ports**
- URLs
- SGT/ISE Attributes
- Inspection
- Logging
- Comments

Available Ports

- Search by name or value
- SNMP
 - SSH
 - Symantec-AV**
 - SYSLOG
 - TCP_high_ports
 - TELNET
 - TFTP
 - Yahoo_Messenger_Messages
 - YahooMessenger_Voice_Chat_TCP
 - YahooMessenger_Voice_Chat_UDP

Add to Source
Add to Destination

Selected Source Ports (1)

- Symantec-AV

Selected Destination Ports (0)

- any

Protocol TCP (6) Port Enter a port Add

Protocol TCP (6) Port Enter a port Add

Add Cancel

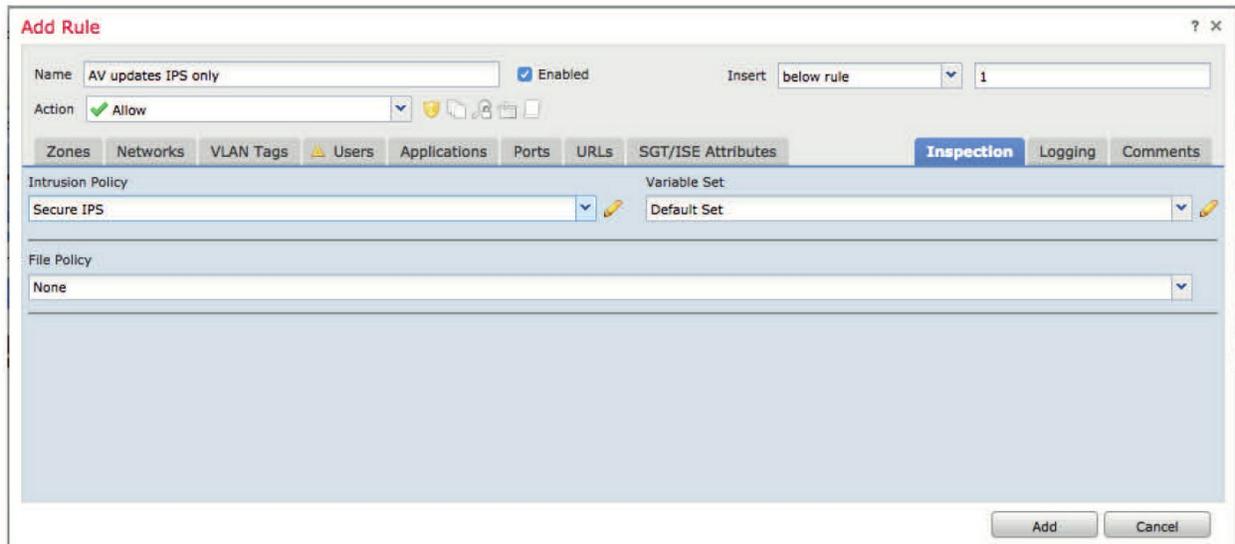
Next, we'll modify the Ports tab and add the source port—8014/ tcp. Again, we're using a custom port object (Symantec-AV) created just for this purpose. This port probably won't change, but it does help in making the rule self-documenting so someone looking at it later can easily understand what we did. The figure below shows the Ports tab:

Finally, on the Inspection tab shown below, we'll specify an Intrusion policy but not a File policy. Doing this allows us to perform IPS inspection but not file inspection on the traffic:

You'll need to decide if you want connection events logged for this rule and configure the Logging tab appropriately.

Wait—one more thing about the rule we just created: The example shows how to address the file inspection behavior only while leaving the IPS inspection intact. If we're being practical, it's really not necessary to even perform IPS inspection on this AV update traffic. If you were to look through all your Snort rules, you'd be hard pressed to find any designed to find attacks in AV update traffic! Because of this, an even more efficient option is to create a Trust rule for AV update traffic. By using the action of Trust instead of Allow, you'll allow this AV update traffic to pass through your devices without any type of inspection at all.

Always remember that complexity is the enemy of security. There are other ways to address situations like the ones in our example. These might include creating custom File policies that don't contain the file type used by our AV updates, but it's really not a good tactic because of the management complexity introduced by multiple File policies. Firepower is a very complex system to start with, and anything you can do to make it easier to understand and streamline it is going to make your life a lot easier!



Summary

In this chapter, I explained the ins and outs of Event Analysis including what False Positives and False Negatives are. Then you learned more about what intrusion events are and how to use the analysis interface to find what you are looking for.

After that, we took a look at Security Intelligence Events, which is a very cool feature that greatly enhances your overall security thanks to Talos for watching out for you!

Finally, we had an in-depth look at File and Malware events, and with that, we'll see you in the next book.