

Practical Advice for
Improving Security, Performance,
Manageability, and High Availability

2017

PALO ALTO NETWORKS BEST PRACTICES

fuel
PALO ALTO NETWORKS
USER GROUP

2017

Palo Alto Networks

Best Practices

Practical Advice for Improving Security,
Performance, Manageability, and High Availability

By Barry J. Stiefel

with contributions and support from the Fuel
User Group Education Advisory Committee.

© Palo Alto Networks Fuel User Group, Inc.



Table of Contents

Master List of Best Practices and Goals.....	4
How to Use This Book to Improve Security, Performance, Manageability, and High Availability	18
The Need for This Book.....	18
The Intended Audience for This Book	18
How to Use this Book.....	18
How to Send Feedback	19
Acknowledgements	20
About the Author	20
UNIVERSAL CONSIDERATIONS	21
Licensing.....	21
Relationships.....	28
Baseline Your Environment	34
Troubleshooting.....	37
Preparing for Disaster Recovery	40
Topology	44
Other	49
PRODUCT: NGFW	50
Initial Setup	50
Management Interface	63
Administrators	69
Management.....	80
Data or Traffic Interfaces	102
Security Zones.....	107
Zone Protection Profiles	112
Virtual Routers	121
Rulebases (General)	122
Security Rulebase.....	138

Objects	176
Address and Address Group Objects	182
Services Objects	184
Security Profile #1: Antivirus	185
Security Profile #2: Anti-Spyware	188
Security Profile #3: Vulnerability Protection	190
Security Profile #4: URL Filtering	192
Security Profile #5: File Blocking.....	200
Security Profile #6: WildFire Analysis	205
Security Profile #7: Data Filtering	212
Security Profile #8: DoS Protection	213
Evasion Prevention	215
Saving NGFW Changes	225
Dynamic Updates.....	232
App-ID	234
User-ID	238
DNS	242
Dynamic Routing.....	248
IPv6	248
Virtual Private Networks (VPNs).....	250
VM-Series NGFWs.....	271
QoS.....	273
Monitoring and Logging.....	276
High Availability	284
While Troubleshooting	295
Upgrading	297
Deployment Scenario: Public Cloud	302
General.....	302
Amazon Web Services (AWS)	304

Master List of Best Practices and Goals

	Security	Performance	Manageability	High Availability
All Products				
Licensing				
Get Organized on Your Palo Alto Networks Licensing and Subscriptions			X	X
Relationships				
Stay Current With Your Support and Feature Licenses			X	X
Build a Good Relationship With Your Reseller			X	X
Build a Good Relationship With Your Systems Engineer (SE)			X	X
Build a Good Relationship With Your Technical Account Manager (TAM)			X	X
Build a Good Relationship With Your Dedicated Technical Support Engineer (TSE)	X		X	X
Build a Good Relationship With Your Resident Engineer (RE)	X	X	X	X
Ask for Help From Professional Services When You Need It	X	X	X	X
Receive and Review E-Mails Sent to Your "Abuse" Address	X			X
Preparing for Troubleshooting				
Document Your Configuration			X	X
Document What "Normal" Traffic and Behavior Looks Like			X	X

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
While Troubleshooting				
Use the Documentation Resources at the Palo Alto Networks Website			X	
Use the Resources at the Palo Alto Networks Live Community			X	
Use the Resources at the Fuel User Group			X	
Preparing for Disaster Recovery				
Make Regular Backups of Your Configuration and Device State				X
Create and Maintain Your Disaster Recovery "Go Bag"				X
Other				
Build a Lab for Learning, Experimentation, and Testing			X	X
Next-Generation Firewall				
Initial Setup				
1. Configure Your Management Interface			X	
2. Configure Your DNS Servers			X	
3. Configure Your Time Zone			X	
4. Configure Your Time Servers			X	
5. Change the Default Password So You Don't Embarrass Yourself	X			
6. Configure Service Routes If Necessary			X	
7. Retrieve Your License Keys From the License Server			X	
8. Configure Your Dynamic Updates Refresh Schedule	X			

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
Licenses				
Subscribe to the Threat Prevention License	X			
Subscribe to the URL Filtering License	X			
Subscribe to the WildFire License	X			
Topology				
Create a Thorough and Accurate Network Diagram			X	
Learn the First Rule of Firewall Topology	X		X	
Design Using the "Zero Trust" Philosophy	X			
Isolate Internet-Accessible Servers in DMZs	X			
Isolate Critical Assets in Separate Physical Networks	X			
Management Interface				
Physically and Logically Isolate Your Management Network	X			
Restrict Permitted IP Addresses Connecting to the Management Interface	X			
Restrict Permitted Services Hosted by the Management Interface	X			
Configure Failed Attempts and Lockout Time	X			
Configure the Management GUI and CLI Idle Timeout	X		X	
Administrators				
Create a Unique Administrator Account for Each Administrator	X		X	

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
Audit Your Administrator Accounts Regularly	X			
Create and Use Admin Role Profiles	X			
Configure Minimum Password Complexity	X			
Don't Use Password Profiles	X			
Management				
Configure the Login Banner	X			
Force Admins to Acknowledge the Login Banner	X			
Configure Header and Footer Banners			X	
Configure Custom Logos in PAN-OS			X	
Change the Default Master Key	X			
Encrypt the Master Key with a Hardware Security Module (HSM)	X			
Set Certificate Expiration Check	X			
Configure Your Geographic Location			X	
Check Last Login Time and Look for Failed Login Attempts	X			
Choose SNMP V3 Over V2c	X			
Configure the Statistics Service	X			
Interfaces				
Configure an Interface Management Profile for Each Interface	X		X	
Enable IPv6 Support	X			
Use the Interface Comment Field			X	

	Security	Performance	Manageability	High Availability
Security Zones				
Apply a Strong Zone Protection Profile to Untrusted Security Zones	X	X		
Use the Special Security Zone Tag Trick to Assign a Color to a Security Zone	X		X	
Zone Protection Profiles				
Drop Malformed IP packets	X			
Remove TCP Timestamps on SYN packets	X			
Drop Mismatched Overlapping TCP Segment	X			
Drop Packets With a Spoofed Source IP Addresses	X			
Experiment With Enabling Some or All of the Other Protections in a Zone Protection Profile	X			
Virtual Routers				
Enable Support for Multicast Firewalling	X			
Rulebases (All)				
Give Every Rule a Name Starting With an Action Verb			X	
Create a Meaningful Description for Each Rule			X	
Use Rule Tags to Organize Rules into Groups			X	
Minimize the Number of Rules		X	X	
Put More Frequently Matched Rules Higher in the Rulebase		X		
Configure "Temporary" Rules to Expire on a Schedule	X		X	

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
Check for Unused Rules Regularly	X		X	
Audit Your Rulebases Regularly	X	X	X	
Add Your Workflow Ticketing System ID to the Rule's Description			X	
When Removing a Rule, Disable It First to See If It Breaks anything			X	
Security Rulebase				
Use Continuous Improvement to Improve Your Security Rulebase			X	
Create a Cleanup Rule at the Bottom of Your Security Rulebase	X			
Follow the Principle of Least Privilege	X			
Create an IP Source Blacklist Rule at the Top of Your Security Rulebase	X	X		
Create an IP Destination Blacklist Rule at the Top of Your Security Rulebase	X	X		
Use Geographic IP Filters As Appropriate	X			
Create an Application Blacklist at the Top of Your Security Rulebase	X	X		
Create an Inbound Service Blacklist at the Top of Your Security Rulebase	X	X		
Avoid Using the Any Source Zone in Allow Rules	X			
Avoid Using the Any Source Address in Allow Rules	X			

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
Avoid Using the Any Source User in Allow Rules	X			
Avoid Using the Any Destination Zone in Allow Rules	X			
Avoid Using the Any Destination Address in Allow Rules	X			
Avoid Using the Any Application in Allow Rules	X			
Before Dropping Previously Allowed Traffic, Allow and Log to Analyze the Effects			X	
Prefer the <i>application-default</i> Service in Allow Rules	X			
Restrict Outbound NTP Traffic Destinations	X			
Restrict Outbound DNS Traffic Destinations	X			
Restrict Outbound SMTP Traffic Destinations	X			
Always Use Security Profiles in Allow Rules	X			
Don't Use a Security Profile in a Drop Rule		X		
Prefer the Drop Action Over Deny or Reset	X			
Don't Normally Send ICMP Unreachable Messages	X			
Consider DSRI for Internet-Facing Servers		X		
Block Internet Connections To and From Private Non-Routable IP Addresses	X			
Block Internet Connections To and From Bogons and Fullbogons	X			

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
Objects				
Create and Use an Object Naming Convention			X	
Create and Use an Object Color Convention	X		X	
Address and Address Group Objects				
Create an Address Object for Every Host, Subnet, Range and FQDN in Your Policy			X	
Create a Meaningful Description for Every Address and Address Group Object			X	
Services Objects				
Create a Meaningful Description for Every Service Object			X	
Security Profile #1: Antivirus				
Create a Strict Antivirus Security Profile	X			
Attach an Antivirus Security Profile to Every Allow Rule	X			
Security Profile #2: Anti-Spyware				
Attach an Anti-Spyware Security Profile to Every Allow Rule	X			
Security Profile #3: Vulnerability Protection				
Attach a Vulnerability Protection Security Profile to Every Allow Rule	X			
Security Profile #4: URL Filtering				
Attach a URL Filtering Security Profile to Every Web Browsing Allow Rule	X			
Prefer PAN-DB URL Filtering Over BrightCloud	X			

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
Block Access to Malicious URL Categories	X			
Block Access to "Unknown" URLs	X			
Log HTTP Header Information	X			
Security Profile #5: File Blocking				
Attach a File Blocking Security Profile to Every Allow Rule	X			
Block the Transit of Windows PE Files	X			
Block the Transit of Common Dangerous and Malicious File Extensions	X			
Security Profile #6: WildFire Analysis				
Attach a WildFire Analysis Security Profile to Every Allow Rule	X			
Maximize the WildFire File Size Limits	X			
Configure WildFire to Report Benign Files	X		X	
Configure WildFire to Report Grayware Files	X		X	
Allow Forwarding of Decrypted Content	X			
Security Profile #7: Data Filtering				
If You Need a Real DLP Solution, Don't Use the Data Filtering Security Profile		X	X	
Security Profile #8: DoS Protection				
Create a DoS Protection Policy	X	X	X	X
Configure a Strong DoS Protection Security Profile	X	X	X	X

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
Evasion Prevention				
Upgrade to At Least PAN-OS 7.1.1 and Applications and Threats Version 579	X			
Enable Vulnerability Protection Security Profile Evasion Signatures	X			
Enable Anti-Spyware Security Profile Evasion Signatures	X			
Clear the Urgent Data Flag in the TCP Header	X			
Drop Segments Without Flags	X			
Drop Segments With a Null Timestamp	X			
Don't Forward Segments Exceeding the TCP Out-of-Order Queue	X			
Don't Forward Segments Exceeding the TCP App-ID Inspection Queue	X			
Don't Forward Datagrams Exceeding the TCP or UDP Content Inspection Queues	X			
Don't Allow the HTTP Header Range Option	X			
Commits				
Avoid Letting Uncommitted Changes Linger			X	
Preview Changes Before Committing			X	
Check for Warnings After Committing			X	
Resolve Commit Warnings			X	
Dynamic Updates				

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
Consider Setting the Threshold Value for Antivirus and Applications and Threats Downloads				X
App-ID				
Use App-ID Filters in Your Security Policy Rules	X			
Blacklist the <i>unknown-tcp</i> , <i>unknown-udp</i> , and <i>unknown-p2p</i> Applications	X			
Create App-ID Signatures for Custom Applications	X		X	
User-ID				
Enable User-ID in Your Security Rules	X		X	
Enable User Identification in Trusted Security Zones	X		X	
Don't Enable User Identification in Non-Trusted Security Zones Unless You're Using Captive Portal	X		X	
Applications				
Replace Telnet in Your Organization With SSH	X			
Replace FTP in Your Organization With SCP/SFTP	X			
DNS				
Configure the Firewall to Be a DNS Proxy	X			
Configure a DNS Sinkhole	X			
Enable Passive DNS Collection for Improved Threat Intelligence	X			

	Security	Performance	Manageability	High Availability
Dynamic Routing				
Consider Separating Firewalling and Dynamic Routing			X	X
IPv6				
Enable IPv6 Firewalling	X			
Virtual Private Networks (VPNs)				
IKE/IPSec Crypto Profile: Configure Strong Authentication	X			
IKE Crypto Profile: Configure Strong Encryption	X			
IKE Crypto Profile: Configure Strong DH Groups	X			
IKE Crypto Profile: Configure Short Key Lifetimes	X			
IKE Crypto Profile: Configure a Low IKEv2 Authentication Multiple	X			
IPSec Crypto Profile: Configure Strong Encryption	X			
IPSec Crypto Profile: Prefer ESP Over AH	X			
IPSec Crypto Profile: Configure Strong DH/PFS	X			
IPSec Crypto Profile: Enable Lifesize Limiting	X			
IKE Gateway: Prefer IKEv2 Over IKEv1	X			
IKE Gateway: Prefer Certificates Over Pre-Shared Keys	X			
IKE Gateway: Prefer Main Exchange Mode Over Aggressive	X			

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
VM-Series NGFWs				
Use Hypervisor Assigned MAC Addresses			X	
QoS				
Use QoS to Match Your Traffic Flows to Your Priorities		X		
Begin Your QoS Configuration With a Simple Plan		X	X	
Monitoring and Logging				
Prefer "Log at Session End" Over "Log at Session Start"		X	X	
Prefer to Not Log DNS, NTP, DHCP and LDAP Traffic		X	X	
Think Hard About What to Log		X	X	
Enable "Resolve Hostname" While Viewing Logs			X	
Get Skilled With the Log Filter Language			X	
Review Your Logs Regularly	X		X	
Select "Enable Log on High DP Load"		X	X	
High Availability				
Configure HA				X
Prefer Active/Passive HA Over Active/Active			X	X
Configure an HA1 Control Link Backup				X
Configure an HA2 Data Link Backup				X
Configure HA1 Control Link Encryption	X			
Configure Dynamic Updates Schedules to "Sync To Peer"			X	

2017 Palo Alto Networks Best Practices

	Security	Performance	Manageability	High Availability
When Manually Downloading Software or Dynamic Updates, Choose "Sync to HA Peer"			X	
Test Your HA Configuration With Real Failures				X
While Troubleshooting				
Consider Logging at Both Session Start and at Session End			X	
Consider Sending ICMP Unreachable for Both Drops and Resets			X	
Upgrading				
Stay Mostly Current With Minor Software Releases	X		X	
Let Major Software Releases Mature for a While Before Installing			X	
Verify Update Server Identity	X			
Deployment Scenario: Public Cloud				
General				
Think Carefully When Choosing the PAYGO Versus BYOL License Models			X	
Amazon Web Services (AWS)				
Think Carefully When Choosing Your Instance Type			X	
Understand the Special AWS Routing Model			X	
Configure Identity and Access Management (IAM) Roles Carefully	X			
Control Access to the Management Port with Security Groups	X			

How to Use This Book to Improve Security, Performance, Manageability, and High Availability

The Need for This Book

Most successful technology companies eventually get a technical book written about their products, but the only books available about Palo Alto Networks so far are study guides for the PCNSE exam. No general interest books about these products or services exist.

But even if there were already a book about Palo Alto Networks products, most technical books are designed to answer the questions “Tell me about X”, or “How do I do Y?” They almost never answer the question, “**What should I do, and why?**”

And that’s the question this Best Practices book answers. It consists entirely of **specific, immediately useful, practical advice on what to do and why**. It’s not an exhaustive enumeration of features and functions, but instead specific advice to help you be as effective as possible as a Security Administrator.

The Intended Audience for This Book

Firewall and Security Administrators interested in implementing best practices will use this book to improve the security, performance, manageability, and high availability of their Palo Alto Networks products.

Potential customers will understand that simply deploying Palo Alto Networks products is not the end of the process in delivering security, performance, manageability, and high availability in their environment availability.

Palo Alto Networks employees will use the book to quickly ramp up their technical skills, better enabling them to service their customers in designing the best Palo Alto Networks infrastructure for their customers’ environment.

How to Use this Book

This book has only one goal, and that’s to help you become the most effective possible administrator of Palo Alto Networks products.

To meet that goal, there are two workflows into this book, both of which involve starting with the Master List of Best Practices, located right after the Table of Contents.

Workflow #1: Focus on a Product and Functional Area

In this workflow, you know the Product and Functional Area you wish to optimize. Find the section corresponding to that area and then refer to the included Best Practices.

Workflow #2: Focus on a specific area

This book is designed to help you achieve improvement four specific areas

- Security

- Performance
- Manageability
- High Availability

In this workflow, you start with one of these specific areas in mind, then follow that goal's column down in the Master List of Best Practices to see which Best Practices help you to achieve that goal.

How to Send Feedback

Continuous improvement works, and we reserve the right to become smarter. If you have a way to make this book more useful, clear, or comprehensive, please send an e-mail to:

BestPracticesBook@FuelUserGroup.org

Acknowledgements

Technical Contributors:

Thanks to David Leitzel, Consulting Engineer, for the Best Practices for AWS deployment.

Thanks for Brian Adams for editing the manuscript.

Thanks to Reese Warner and Martin Markovich, fellow authors, for guidance.

Technical Reviewers:

Thanks to the Technical Reviewers:

Marcel Hoffmann

Justin Scaggs

Fuel User Group Education Advisory Committee:

Steiner Aandal-Vanger

Martinien Betwa

Jack Crowder

Barry Hofecker

Ian Johnston

Mario Perez

Eugene Purugganan

Jason Rakers

Darin Sutton

Jason Reverri

About the Author

Barry J. Stiefel (“Stee-ful” or “Shtee-ful”)

Barry has a **B.S., MBA**, a year of **Post-Graduate Study in Operations Research** and **Certificate in Internet and E-Commerce Security** from the University of California and has earned the numerous industry certifications including PCNSE6, CISSP, and NSA IAM.

In addition to this, he has written several books on technical subjects and been heavily involved in user groups for over 15 years. Some of his expertise include creating curricula, courseware, lab environments, online communities, software tools, certification exams and accreditation exams to improve the productivity of Sales Engineers, as well as training Sales Engineers and organizing teams of Consulting Engineers for special projects.

UNIVERSAL CONSIDERATIONS

Licensing

□ Get Organized on Your Palo Alto Networks Licensing and Subscriptions

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

Background Information:

The full security functionality of Palo Alto Networks' products is partitioned among a handful of individually licensable products. These products aren't going to work unless you have a valid license and that valid license is properly installed.

Why This Best Practice Is Important:

Getting organized on your Palo Alto Networks licensing and subscriptions is important for three reasons.

- You want to ensure you've licensed all the components you need.
- You want to ensure your licenses are properly installed and recognized.
- You want to ensure you're not hit with a surprise functionality outage because a license has inadvertently expired.

How to Implement It:

Step 1: Figure out which licenses you're entitled to.

The easiest way to view and your licenses is on the Palo Alto Networks support website. You'll be able to see exactly what you've got and when it will expire. The display will highlight already-expired licenses.

You can view and manage your support contract and licenses on the Technical Support web page at <https://www.paloaltonetworks.com/services/support>.

Step 2: Ensure your licenses are installed and recognized.

1. Go to **Device > Licenses**.
2. If your firewall has Internet access—not all do, for security reasons—then click on the *retrieve license keys from license server* command to refresh the data.

You'll see a result similar to this:

PA-VM Date Issued July 11, 2016 Date Expires Never Description Standard VM-300	AutoFocus Device License Date Issued October 19, 2016 Date Expires October 11, 2021 Description AutoFocus Device License
BrightCloud URL Filtering Date Issued July 12, 2016 Date Expires July 12, 2019 Description BrightCloud URL Filtering Active No (Activate)	GlobalProtect Gateway Date Issued July 12, 2016 Date Expires July 12, 2019 Description GlobalProtect Gateway License
PAN-DB URL Filtering Date Issued July 12, 2016 Date Expires July 12, 2019 Description Palo Alto Networks URL Filtering License Active Yes Download Status 2016-12-19 15:47:55 PAN-DB download: Finished successfully. Re-Download	Threat Prevention Date Issued July 12, 2016 Date Expires July 12, 2019 Description Threat Prevention
WildFire License Date Issued July 12, 2016 Date Expires July 12, 2019 Description WildFire signature feed, integrated WildFire logs, WildFire API	License Management Retrieve license keys from license server Activate feature using authorization code Manually upload license key Deactivate VM

Viewing all your licenses.

□ Stay Current with Your Support and Feature Licenses

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

Background Information:

Your support contract and the separately-licensed add-ons to your platform require current licenses with Palo Alto Networks.

Why This Best Practice Is Important:

If you're not current with your support contract, you won't be able to open tickets, call for help when you need it, or upgrade to the current version.

If you're not current with your licenses, critical additional functionality won't work.

How to Implement It:

You can view and manage your support contract and licenses on the Technical Support web page at <https://www.paloaltonetworks.com/services/support>.

Schedule calendar reminders about your renewals 30-60 days before they expire.

While it's true that both Palo Alto Networks and your reseller will probably remind you in plenty of time, it's best to schedule calendar reminders to ensure you renew your licenses before they expire. The last thing you want to be doing is trying to get a Purchase Order approved during a service outage.

□ Subscribe to the Threat Prevention License

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The Threat Prevention License allows you to receive Dynamic Updates for these three key functionalities:

- Antivirus
- Anti-Spyware
- Vulnerability Prevention

Why This Best Practice Is Important:

You really need those three key functionalities.

How to Implement It:

Talk with your Sales Engineer or reseller and make sure your firewalls have the Threat Prevention license.

What It Looks Like After You've Implemented It:

Go to **Device > Licenses > License Management**.

The screenshot shows the 'License Management' page in the Palo Alto Networks management console. It displays a list of installed licenses on the left and a detailed view of the 'Threat Prevention' license on the right. The 'Threat Prevention' license is circled in red.

License Name	Date Issued	Date Expires	Description	Active
PA-VM	July 11, 2016	Never	Standard VM-300	
BrightCloud URL Filtering	July 12, 2016	July 12, 2019	BrightCloud URL Filtering	No (Activate)
PAN-DB URL Filtering	July 12, 2016	July 12, 2019	Palo Alto Networks URL Filtering License	Yes
WildFire License	July 12, 2016	July 12, 2019	WildFire signature feed, integrated WildFire logs, WildFire API	
AutoFocus Device License	October 19, 2016	October 11, 2021	AutoFocus Device License	
GlobalProtect Gateway	July 12, 2016	July 12, 2019	GlobalProtect Gateway License	
Threat Prevention	July 12, 2016	July 12, 2019	Threat Prevention	

License Management Actions:

- Retrieve license keys from license server
- Activate feature using authorization code
- Manually upload license key
- Deactivate VM

A valid Threat Prevention license

❑ Subscribe to the URL Filtering License

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The URL Filtering License allows you to use Security Policy rules to enforce web access based on dynamic URL categories.

Why This Best Practice Is Important:

URL Filtering can reduce two risks to your organization:

- Security Risk: By blocking access to sites categorized as *malware* and *phishing*.
- Legal Risk: By blocking access to websites that can incur Human Resources and other legal liabilities.

How to Implement It:

Talk with your Sales Engineer or reseller and make sure your firewalls have the URL Filtering license.

What It Looks Like After You've Implemented It:

Go to **Device > Licenses > License Management**.

The screenshot displays the 'License Management' section of the Palo Alto Networks management console. It lists several licenses for a device named 'PA-VM':

- AutoFocus Device License:** Issued July 11, 2016, expires never, description 'Standard VM-300'.
- BrightCloud URL Filtering:** Issued July 12, 2016, expires July 12, 2019, description 'BrightCloud URL Filtering', active status 'No (Activate)'.
- PAN-DB URL Filtering:** Issued July 12, 2016, expires July 12, 2019, description 'Palo Alto Networks URL Filtering License', active status 'Yes', download status '2016-12-19 15:47:55 PAN-DB download: Finished successfully. Re-Download'.
- WildFire License:** Issued July 12, 2016, expires July 12, 2019, description 'WildFire signature feed, integrated WildFire logs, WildFire API'.
- GlobalProtect Gateway:** Issued July 12, 2016, expires July 12, 2019, description 'GlobalProtect Gateway License'.
- Threat Prevention:** Issued July 12, 2016, expires July 12, 2019, description 'Threat Prevention'.

The 'PAN-DB URL Filtering' license is circled in red, indicating it is the focus of the best practice. The 'License Management' section on the right provides options to 'Retrieve license keys from license server', 'Activate feature using authorization code', 'Manually upload license key', and 'Deactivate VM'.

A valid PAN-DB URL Filtering license

□ Subscribe to the WildFire License

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Access to WildFire is one of the best parts of being a Palo Alto Networks customer. WildFire is the cloud-based malware analysis environment that executes potential malware in virtual machines. It uses Dynamic Analysis, Static Analysis, Machine Learning and bare metal analysis, observing hundreds of possibly malicious behaviors and watching for detonations. Before it blows up in *your* network, let it blow up in *a* virtual machines.

Every Palo Alto Network firewall and every Traps endpoint user forwards files up to WildFire, giving WildFire a **million distant early warning sensors throughout the Internet**. When a new piece of malware is released, it only has to touch **one** of those sensors anywhere in the world and a signature for it is created and then new signatures are published every five minutes and pushed down to every WildFire subscriber.

WildFire discovers **150,000 pieces of never-before-seen malware per day**, and **every one of them is dead within minutes on every Palo Alto Networks firewall that has a WildFire subscription**.

Why This Best Practice Is Important:

You have to be a WildFire subscriber to get these every-five-minutes downloads. Without it, you have to wait for the once per day Antivirus signature downloads. This is a community that you want to be part of.

How to Implement It:

Talk with your Sales Engineer or reseller and make sure your firewalls have the WildFire subscription.

What It Looks Like After You've Implemented It:

Go to **Device > Licenses > License Management**.

2017 Palo Alto Networks Best Practices

PA-VM Date Issued July 11, 2016 Date Expires Never Description Standard VM-300	AutoFocus Device License Date Issued October 19, 2016 Date Expires October 11, 2021 Description AutoFocus Device License
BrightCloud URL Filtering Date Issued July 12, 2016 Date Expires July 12, 2019 Description BrightCloud URL Filtering Active No (Activate)	GlobalProtect Gateway Date Issued July 12, 2016 Date Expires July 12, 2019 Description GlobalProtect Gateway License
PAN-DB URL Filtering Date Issued July 12, 2016 Date Expires July 12, 2019 Description Palo Alto Networks URL Filtering License Active Yes Download Status 2016-12-19 15:47:55 PAN-DB download: Finished successfully, Re-Download	Threat Prevention Date Issued July 12, 2016 Date Expires July 12, 2019 Description Threat Prevention
WildFire License Date Issued July 12, 2016 Date Expires July 12, 2019 Description WildFire signature feed, integrated WildFire logs, WildFire API	License Management Retrieve license keys from license server Activate feature using authorization code Manually upload license key Deactivate VM

A valid WildFire license

Relationships

☐ Build a Good Relationship With Your Reseller

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

Palo Alto Networks sells through resellers and distributors and these partners are crucial local resources for customers worldwide.

Why This Best Practice Is Important:

Your local reseller maintains a close relationship with Palo Alto Networks, has up-to-date information, and is your partner for upgrades and new services. They may also provide Technical Support services and Professional Services.

Your best strategy is to maintain a close relationship with your reseller, let them know your goals and plans, and let them help you meet your objectives.

□ Build a Good Relationship With Your Systems Engineer (SE)

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

Palo Alto Networks Systems Engineer (SEs) are very skilled and experienced and are regularly receiving additional training. They know the company's products and services well and can help with designing your security architecture and answering technical questions in the sales process.

They also have access to many internal resources including an even more elite team of Consulting Engineers (CEs) who further specialize in specific product areas.

Why This Best Practice Is Important:

Your best strategy is to maintain a close relationship with your SE, let them know your goals and plans, and let them help you meet your objectives.

□ Build a Good Relationship With Your Technical Account Manager (TAM)

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

The Palo Alto Networks Technical Account Manager (TAM) program provides customers who require a more proactive level of support with a dedicated customer advocate. The Technical Account Manager will be assigned to manage and prioritize all customer cases, questions, OS recommendations, and regularly scheduled account calls. They will also be familiar with the customer's specific implementation, acting as a bi-directional communications conduit between Palo Alto Networks and the customer.

Why This Best Practice Is Important:

Here are some of the benefits of having a Technical Account Manager:

- Designated Technical Account Manager to coordinate all aspects of the customer interaction
- Familiarity with customer business objectives and deployment plans
- Coordinated access to resources, including white papers, best practices, product updates, known issues and resolutions, and recommended upgrades
- Liaison between deployment and support teams
- Improved case resolution and accelerated turnaround times on Return Material Authorization (RMA) and Failure Analysis (FA), as well as prioritized call routing
- Faster access to Palo Alto Networks top Engineers and Subject Matter Experts (SMEs)

As your configuration grows larger, more complex, or more mission-critical, the benefits of having a Technical Account Manager become obvious.

How to Implement It:

Contact Palo Alto Networks Technical Management at:

<https://www.paloaltonetworks.com/resources/datasheets/technical-account-management>

□ Build a Good Relationship With Your Dedicated Technical Support Engineer (TSE)

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

The Premium Plus level of support contract assigns a dedicated Technical Support Engineer (TSE) to your account. This Engineer, who supports only a limited number of customers, is going to be fully read in and current on your equipment, configuration and issues.

Why This Best Practice Is Important:

Tech support calls go much faster because you don't have to get anybody up to speed, and you've got an Engineer whose job it is to think about your account, rather than always having to move on the next ticket.

How to Implement It:

Upgrade to Premium Plus support, and start making a good friend in Technical Support.

□ Build a Good Relationship With Your Resident Engineer (RE)

Improve Security	X	Improve Manageability	X
Improve Performance	X	Improve High Availability	X

Background Information:

Even though recruiting and retaining staff is vital to the security of your company, it can sometimes be difficult and time-consuming to find people with the right skills in the geography that you need to cover. The Palo Alto Networks Resident Engineer Program can provide customers with on-site product experts to aid your team.

Why This Best Practice Is Important:

As an on-site member of your security team, a Resident Engineer is uniquely qualified to advise your team how to use Palo Alto Networks products and services inside your organization. By understanding your business needs from the inside out, a Resident Engineer can help match requirements with solutions.

Here are some of the tasks a Resident Engineer can help with:

Engineer Responsibility Tasks

Advise	<ul style="list-style-type: none"> • Serve as the customer's go-to resource for all matters related to the Palo Alto Networks Next-Generation Firewall • Deliver best practices guidance for managing Palo Alto Networks firewalls • Assist in identifying customer-specific requirements and provide solutions • Answer questions about the product
--------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

2017 Palo Alto Networks Best Practices

	capabilities, features and configuration
Plan	<ul style="list-style-type: none"> • Assist the customer team with the design and placement of Palo Alto Networks devices • Evaluate deployment scenarios, both locally and globally • Help with proof of concept and new feature evaluation in the customer's lab • Assist with technology integration questions
Migrate	<ul style="list-style-type: none"> • Plan migration tasks • Analyze existing rules and objects • Optimize and migrate policies and objects from existing environment to Palo Alto Networks firewalls • Test and validate the migration environment • Coordinate and execute cutover to production
Maintain	<ul style="list-style-type: none"> • Work with security teams to provide ongoing firewall management recommendations • Facilitate the development of new application and threat signatures • Provide ongoing support for customized tools and scripts developed by professional services • Help the customer with managing licenses • Work with the customer to request new features • Facilitate discussions between the customer and the product management team to review roadmaps
Troubleshoot	<ul style="list-style-type: none"> • Take the lead on escalation issues • Work with support to troubleshoot product issues
Train	<ul style="list-style-type: none"> • Provide on-site assistance for knowledge transfer with operational personnel • Train the customer on how to access support, knowledge base and other available services

□ Ask for Help From Professional Services When You Need It

Improve Security	X	Improve Manageability	X
Improve Performance	X	Improve High Availability	X

Background Information:

In addition to providing Resident Engineers, Palo Alto Networks can also provide these services:

- Professional Services
- Migration Consulting Services
- Architecture Consulting Services
- Health Check & Configuration Audit Services
- Proof of Concept Testing
- Validation Testing

In addition to in-house Professional Services Engineers, Palo Alto Networks has also developed the Certified Professional Services Provider (CPSP) program to enable and promote qualified partners who have demonstrated world-class professional services capabilities and expertise.

Why This Best Practice Is Important:

Palo Alto Networks Professional Services Engineers do this full-time and are regularly receiving ongoing training. It's often just simpler and easier to bring in specialists than try to develop your own talent.

How to Implement It:

Talk to Palo Alto Networks Professional Services:

<https://www.paloaltonetworks.com/services/consulting>.

See also the Certified Professional Services Provider (CPSP) program at:

<https://www.paloaltonetworks.com/services/cpsp-partners>

Baseline Your Environment

☐ Receive and Review E-Mails Sent to Your "Abuse" Address

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	X

Background Information:

The standards for good etiquette on the Internet include sending an e-mail to the address abuse@domain.com whenever you spot spam, malware, or other bad behavior coming from a particular domain. This implies, of course, that another part of the standards for good etiquette is receiving and reviewing e-mails sent to *your own* abuse reporting address.

Why This Best Practice Is Important:

While, yes, it can be annoying to receive e-mail at your abuse reporting address, it's also a great way to receive feedback about misbehaving hosts, devices, and users. Think of it as an early warning system to help you get on top of issues quickly.

Also, ignoring complaints is a fast way to get your domains on blocklists, which can cause all sorts of problems.

How to Implement It:

Forward e-mails sent to abuse@domain.com for each of your domains to someone who can act on them quickly.

What Else You Need to Know:

Remember that it's not done until it's tested.

In that spirit, I sent a test message to abuse@PaloAltoNetworks.com and a real person e-mailed me back within twenty minutes, offering her assistance. "You're doing it right," I said.

□ Document Your Configuration

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

Why This Best Practice Is Important:

There are two benefits to documenting your configuration:

- **It Will Help In Planning:** A prerequisite for planning for increased capacity or functionality is knowing what you've already got.
- **It Will Help In Disaster Recovery:** When something goes wrong and you've got an outage, time is critical and brainpower is limited. You need to know what you've got before you can start troubleshooting, so it's best to have this ready in advance..

How to Implement It:

80% of life is showing up

--Woody Allen

80% of the benefit of documenting your configuration is just in the doing of it. The form is much less important than its mere existence. These methods work well:

- Microsoft Word documents
- Microsoft Excel spreadsheets (By using cells borders creatively, you can actually make a rather good network diagram and *everybody* has the software to read it and edit it.)
- Visio (on the PC) or OmniGraffle (on the Mac)
- Pen and Paper: This remains surprisingly effective. Just tell everyone you're using an early pre-alpha version of Visio.

What It Looks Like After You've Implemented It:

Your goals should be thoroughness and accuracy. Think hard about what you're going to wish you had in the event of an outage, and then create it.

What Else You Need to Know:

To remain effective, you're going to have to update it regularly. After major changes and once per quarter might be a reasonable schedule.

□ Document What "Normal" Traffic and Behavior Looks Like

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

Why This Best Practice Is Important:

When you're looking down a wire, quite often a properly functioning firewall is indistinguishable from an unplugged cable. When the time comes for troubleshooting and you try to ping something and don't get a response, does that mean there's something terribly wrong, or that the firewall is functioning perfectly?

When you measure traffic flow on a link, is that a normal amount of traffic, or is the link being flooded with something unusual, or is that flow smaller than normal, which means something has failed?

You can learn a lot about a system by learning what its "normal" behavior is and then comparing it to real time observations.

How to Implement It:

Automatic network monitoring tools help a lot. You'll learn which devices are usually up, or down, and which IPs and services are usually up or down.

Keep a spreadsheet or monitoring history that tracks normal traffic flows.

What It Looks Like After You've Implemented It:

You should be able to step into your network and make observations and instantly be able to tell whether things look "normal" or not. This will make it a lot easier to isolate problems.

Troubleshooting

□ Use the Documentation Resources at the Palo Alto Networks Website

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The company's public website home page is at <https://PaloAltoNetworks.com> and the technical documentation is at <https://www.PaloAltoNetworks.com/documentation>. There's a lot of useful material there, it's well organized, and the Search function works well.

Why This Best Practice Is Important:

The company itself is the single source of truth about their products and the public website is fast and free. They work hard to put as much useful technical information out there as possible. The documentation section is a great source for understanding what's going on when you're troubleshooting.

□ Use the Resources at The Palo Alto Networks Live Community

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The Live Community at <https://live.PaloAltoNetworks.com> is the community portal for Technical Support. It has discussion boards, articles and blogs.

Why This Best Practice Is Important:

There's a tremendous amount of information available in the Live Community and it's all free and easily searchable. It's really useful when you're troubleshooting.

□ Use the Resources at the Fuel User Group

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The official Palo Alto Networks user group is the Fuel User Group at <https://www.fuelusergroup.org/>. It contains forums, webinars, User Summits and local chapter events.

Why This Best Practice Is Important:

The Fuel User Group provides lots of useful information by and for your fellow firewall administrators and helps you meet up with fellow security professionals in similar roles and facing similar issues.

Preparing for Disaster Recovery

□ Make Regular Backups of Your Configuration and Device State

Improve Security		Improve Manageability	
Improve Performance		Improve High Availability	X

Background Information:

Your firewall's configuration is stored in XML format on the box itself and is not automatically backed up anywhere else.

Why This Best Practice Is Important:

There are two types of failures that you need to prepare for:

- **Type I:** Your *running configuration* gets corrupted or you make a change that you can't undo.
- **Type II:** The whole box fails and you lose the configuration and device state.

How to Implement It:

How to Prepare for Type I Failures:

Save a named configuration snapshot:

- On a regular basis
- Whenever you're about to begin a major set of configuration changes

To save a named configuration snapshot, go to **Device > Setup > Operations > Configuration Management > Save named configuration snapshot**. This will save the configuration file locally on the box. It will be easy to load this configuration later if you need to.

How to Prepare for Type II Failures:

Export named configuration and snapshot and device state:

- On a regular basis
- Whenever you're performing an update or major configuration changes

To export a named configuration snapshot, go to **Device > Setup > Operations > Configuration Management > Export named configuration snapshot**. This will pull down the XML file through your web browser and save it locally on your machine. It will be easy to import this configuration later if you need to.

To Export the device state, go to **Device > Setup > Operations > Configuration Management > Export device state**.

What Else You Need to Know:

When saving a configuration snapshot, the name field is picky about how to name the file. You can only use letters and numbers; no symbols or spaces.

When exporting a configuration snapshot, you can only choose among the configuration snapshots you've already

saved on the device.

□ Create and Maintain Your Disaster Recovery “Go Bag”

Improve Security		Improve Manageability	
Improve Performance		Improve High Availability	X

Why This Best Practice Is Important:

When things go wrong, you’re going to be under a lot of time and cognitive pressure to fix things as quickly as possible. The last thing you want to be doing is scrambling to find your hardware, software, and information tools.

How to Implement It:

This is a pretty good list of things that should go in your Go Bag:

Hardware:

Productivity:

- High capacity USB hard drive
- Bootable Linux USB thumb drive with tools
- Bootable Linux DVD with tools
- USB Mouse
- USB keyboard
- USB CD/DVD reader/writer
- Blank DVDs
- Label maker

Networking:

- Long Category 6 Ethernet cable
- Short Male-Female Category 6 crossover cable
- Serial cable
- Fiber cables and SFPs appropriate for your data center
- Cable ties

Tools:

- Flashlight and spare batteries
- Needle nose pliers (large, with cutter blades at the base; and small)
- Screwdrivers (flat heads and Philips heads)
- Micro screwdrivers
- Pocket knife or multi-tool

Support the Human:

- Power cable for laptop
- Charging cable for mobile phone
- Headset for mobile phone
- Spare battery for the mobile phone
- Ear plugs or hearing protectors
- Two cans of Red Bull
- Granola bars
- Cash for vending machine
- Credit cards for the computer store

Miscellaneous:

- Power cable
- Anti-static bag for components

Software:

Install packages for:

- PuTTY
- WinSCP
- Nmap
- Wireshark

Information:

- Backups of all your configuration files
- Lists of the physical locations of all your gear (“PA-5260 1 racked in DC1, Row 6, rack 3, position 17-20”)
- Copies of all your network configuration information
- Encrypted list of your user accounts and passwords
- Contract and contact information for all your Technical Support contracts

What It Looks Like After You’ve Implemented It:

Pack all this into a soft-sided tool bag and you’re ready to deploy anywhere to fix things.

Topology

□ Create a Thorough and Accurate Network Diagram

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Modern networks are highly complex and multi-layered. With anything other than the most simple of configurations, there's no way for an administrator to remember it all.

"We are losing the ability to understand anything that's even vaguely complex."

--Chuck Klosterman, Sex, Drugs, and Cocoa Puffs: A Low Culture Manifesto

Why This Best Practice Is Important:

A map says to you, 'Read me carefully, follow me closely, doubt me not.' It says, 'I am the earth in the palm of your hand. Without me, you are alone and lost.'

--Beryl Markham, West with the Night

How to Implement It:

The method by which you create your network diagram is so much less important than just the fact that you're creating it. Visio does a pretty good job, as does OmniGraffle in the Apple world. By clever use of borders, Excel can be made to work fine and it has the added advantage that the diagrams can be edited on any PC. But paper and pencil are also fine. What matters is that you do it, and the trial-and-error and hard thinking that it requires forces you to understand your network.

Once you have it, you'll need to regularly update it with adds, moves, and changes, but when it comes time to architect or troubleshoot, you'll find it invaluable.

A good network diagram usually contains three classes of objects:

- **Network Objects:** A good start would be to include key devices that have Layer 2 or Layer 3 addresses.
- **Links between Network Objects:** Ethernet, Wi-Fi, WAN links, etc.
- **Zones:** Layer 3 subnets, Security Zones, zones delineated by physical architecture.

What It Looks Like After You've Implemented It:

You want enough detail to show what's going on, but not so much it's cluttered. Remember that its two goals are to be an aid in planning and a reference while troubleshooting. You really can't be effective as a Security Administrator without a thorough and accurate network diagram.

□ Learn the First Rule of Firewall Topology

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The First Rule of Firewall Topology is that the firewall can't protect connections that don't transit the firewall.

Why This Best Practice Is Important:

There are two concerns here:

1. It's easy in a highly complex network for the topology to migrate from "tree-like" to "mesh-like", especially if you're using dynamic routing, and before you know it there's an alternative route that some traffic can take that bypasses firewall inspection. This can be particularly dangerous.
2. Another Best Practice discusses the "Zero Trust Model" in which even traffic between internal hosts get inspected. Plenty of Advanced Persistent Threats (APT) these days can spend months moving laterally through your network and unless you're specifically forcing this traffic to transit a firewall, it's not going to get inspected.

How to Implement It:

Follow the Best Practice about maintaining a thorough and accurate network diagram and ensure that as much traffic as possible transits a firewall.

□ Design Using the "Zero Trust" Philosophy

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

In the old days, network security meant following the Turtle Strategy: crunchy on the outside, soft on the inside. The idea was that it was both necessary and sufficient to have good perimeter security; because you were keeping all the bad stuff out you didn't have to care so much about security on the inside.

Times have changed. With universal web and e-mail access, USB drives, Bring Your Own Device (BYOD) policies, etc., there are many more opportunities for malware to infect patient zero on the inside and then move laterally within your organization.

Why This Best Practice Is Important:

Threats aren't just coming from the outside anymore, they're also coming from within your network. You need to protect traffic *between* all your devices, regardless of where those devices are located, physically or topologically.

How to Implement It:

By using Layer 3 subnets, VLANs, or even just cabling topology (when using Virtual Wire interfaces), partition your network into smaller subsets and force all inter-subset traffic to transit a firewall. If one of your devices is infected, you want to provide as small a topological area as possible for free movement before it runs into an inspection.

Example: A large consumer data company's server farm contained multiple layers of web servers, middleware servers, and databases. By creating 240 (mostly VLAN) interfaces on their firewall, they were able to ensure that *every connection between every pair of devices* had to transit the firewall.

What Else You Need to Know:

Converting your topology from the Turtle Strategy to the Zero Trust strategy can increase the traffic transiting your firewalls. This may lead to performance or sizing issues. Your Sales Engineer (SE) can help you figure out the best strategy.

□ Isolate Internet-Accessible Servers in DMZs

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Why This Best Practice Is Important:

Servers that receive connections from the Internet are particularly vulnerable to attack and malware infection. Because of their higher risk, they should be isolated in physically and logically separate network segments.

How to Implement It:

The design goal is that connections from the DMZ into the internal network—if any are permitted at all—should face restrictions and scrutiny at least as rigorous as if they were coming the public Internet.

It's best to separate out the DMZ on to its own physical network with its own dedicated port on the firewall.

□ Isolate Critical Assets in Separate Physical Networks

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Why This Best Practice Is Important:

Devices containing the most critical data, such as customer passwords or financial or health record information, warrant a higher level of protection.

How to Implement It:

Similar to the extra physical and topological isolation given to Internet-accessible servers, servers containing critical information need their own separate networks. Give them a separate physical port on the firewall, or perhaps even a separately-managed firewall.

Other

□ Build a Lab for Learning, Experimentation, and Testing

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

Why This Best Practice Is Important:

Once a production firewall is in place, everyone's afraid to touch it. Nobody wants to be the admin who knocked Engineering off the Internet because he used it for learning, experimentation, or testing.

But you still need to do these things, particularly as your network and security requirements get more complex. It's important that you have a sandbox where you can play without fear and perhaps even stage changes before rolling them out to your main firewalls.

How to Implement It:

If you've got the budget to match your lab equipment to your production hardware, that's a very good start. If not, start smaller. Even if you have to just use a firewall from the PA-200 series, it's worth it to play with all the options, experiment with upgrading, and try out lots of new Best Practices. Remember that PAN-OS behaves essentially the same on every platform.

What It Looks Like After You've Implemented It:

If you're doing it right, you'll have a lab environment where you can learn, test out hypotheses, stage potential configuration changes, and feel free to move fast and break things. Get to work.

PRODUCT: NGFW

Initial Setup

☐ Configure Your Management Interface

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The first of the eight steps you need to complete for minimal functionality is to configure the Management Interface (MGT). The Management Interface is a physically separate Ethernet port from the “traffic” ports.

Why This Best Practice Is Important:

If you can’t connect to the MGT interface via Ethernet, you won’t be able to view or manage your security policy or monitor your NGFW.

How to Implement It:

Go to **Device > Setup > Management > Management Interface Settings**.

The settings that are truly required in this first step are:

- IP Address
- Netmask
- Default Gateway
- Service: HTTPS

What It Looks Like After You've Implemented It:

The screenshot shows the 'Management Interface Settings' window. On the left, under 'IP Type', 'Static' is selected. The 'IP Address' field contains '192.168.44.2', 'Netmask' is '255.255.255.0', and 'Default Gateway' is '192.168.44.254'. The 'Speed' dropdown is set to 'auto-negotiate' and 'MTU' is '1500'. On the right, the 'Services' section has checkboxes for HTTP, HTTP OCSP, HTTPS (checked), Telnet, SSH (checked), Ping (checked), SNMP, User-ID, User-ID Syslog Listener-SSL, and User-ID Syslog Listener-UDP. To the right of the services is a 'Permitted IP Addresses' list with 'Add' and 'Delete' buttons at the bottom. 'OK' and 'Cancel' buttons are at the bottom right of the window.

Configuring just the basics to get started

What Else You Need to Know:

Other Best Practices will cover how to secure your management interface.

□ Configure Your DNS Servers

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The second of the eight steps you need to complete for minimal functionality is to configure the device's DNS servers.

Why This Best Practice Is Important:

The management interface is going to need to resolve Fully Qualified Domain Names (FQDN) into IP addresses for the many types of connections it needs to make to the Internet.

How to Implement It:

Go to **Device > Setup > Services**.

If your organization doesn't have private DNS servers, a reasonable public choice is provided by Google's public DNS servers:

- 8.8.8.8
- 8.8.4.4

What It Looks Like After You've Implemented It:

The screenshot shows the 'Services' configuration window in the Palo Alto Networks management interface. The 'DNS' section is active, with 'Servers' selected. The 'Primary DNS Server' is set to 8.8.8.8, the 'Secondary DNS Server' is set to 8.8.4.4, and the 'Update Server' is set to updates.paloaltonetworks.com. The 'Verify Update Server Identity' checkbox is checked. The 'Proxy Server' section is also visible, with fields for Server, Port (1 - 65535), User, Password, and Confirm Password. The 'OK' and 'Cancel' buttons are at the bottom right.

Now your NGFW can resolve FQDNs into IP addresses.

□ Configure Your Time Zone

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The third of the eight steps you need to complete for minimal functionality is to configure the device's time zone.

Why This Best Practice Is Important:

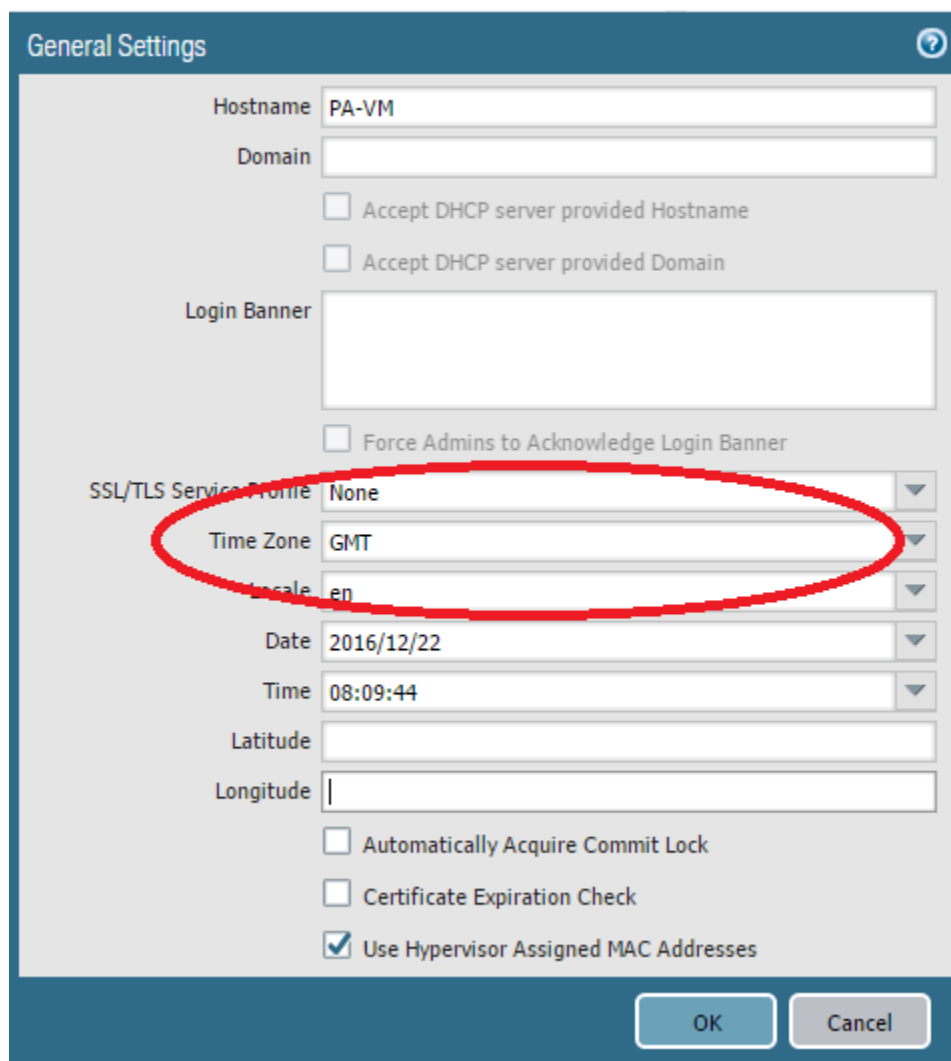
If you configure your time zone properly then a whole bunch of problems won't visit you while you're troubleshooting. Your logs will make sense and you'll be able to correlate events across multiple devices.

It's an interesting question, though, as to which time zone to choose. If your entire operation is within a single time zone, that's probably the best choice, but you can still get confused if your region makes changes for Daylight Saving Time.

If your organization spans multiple time zones or you want to be more professional and avoid DST issues, then choose GMT.

How to Implement It:

Go to **Device > Setup > Management > General Settings**.



The screenshot shows the 'General Settings' dialog box in the Palo Alto Networks management interface. The 'Time Zone' dropdown menu is highlighted with a red oval and is set to 'GMT'. Other settings visible include Hostname (PA-VM), Domain, Login Banner, SSL/TLS Service Profile (None), Date (2016/12/22), Time (08:09:44), and various checkboxes for system configuration.

Setting	Value
Hostname	PA-VM
Domain	
Accept DHCP server provided Hostname	<input type="checkbox"/>
Accept DHCP server provided Domain	<input type="checkbox"/>
Login Banner	
Force Admins to Acknowledge Login Banner	<input type="checkbox"/>
SSL/TLS Service Profile	None
Time Zone	GMT
Locale	en
Date	2016/12/22
Time	08:09:44
Latitude	
Longitude	
Automatically Acquire Commit Lock	<input type="checkbox"/>
Certificate Expiration Check	<input type="checkbox"/>
Use Hypervisor Assigned MAC Addresses	<input checked="" type="checkbox"/>

By setting the Time Zone to GMT you can easily correlate events throughout your organization.

□ Configure Your Time Servers

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The fourth of the eight steps you need to complete for minimal functionality is to configure the device's time servers.

Why This Best Practice Is Important:

If the clocks on your devices are accurate then a whole bunch of problems won't visit you while you're troubleshooting. Your logs will make sense and you'll be able to correlate events across multiple devices.

Use your organization's internal time servers if you have them, or use public time servers. Here are two suggestions:

- utcnist.colorado.edu > 128.138.140.44
- utcnist2.colorado.edu > 128.138.141.172

These are at the U.S. National Institute of Standards (NIST). See <http://tf.nist.gov/tf-cgi/servers.cgi>.

Another good source is the NTP Pool Project at <http://www.pool.ntp.org/en/>.

How to Implement It:

Go to **Device > Setup > Services**.

The screenshot shows the 'Services' configuration page with the 'NTP' tab selected. It contains two main sections: 'Primary NTP Server' and 'Secondary NTP Server'. Each section has a text field for 'NTP Server Address' and a dropdown menu for 'Authentication Type'. The Primary NTP Server address is set to '128.138.140.44' and the Secondary NTP Server address is set to '128.138.141.172'. Both 'Authentication Type' dropdowns are set to 'None'. At the bottom right, there are 'OK' and 'Cancel' buttons.

Now we're synchronized to atomic clocks.

❑ Change the Default Password So You Don't Embarrass Yourself

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The fifth of the eight steps you need to complete for minimal functionality is to change the default password. The default password is **admin**, which is awful, not only because it's short and guessable, but *because it's the same on every new box*.

Why This Best Practice Is Important:

Anyone who stumbles across a Palo Alto Networks firewall login screen anywhere will be able to search on the Internet for "Palo Alto Networks default login" and get a password within seconds that they can immediately try.

If your box gets pwned because you never changed the default password, you'll never live it down.

How to Implement It:

Click on "admin" in the bottom left corner of your GUI to get this dialog box:

If you haven't seen this dialog box yet, start looking for it now.

□ Retrieve Your License Keys from the License Server

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The sixth of the eight steps you need to complete for minimal functionality is to retrieve your license keys from the license server.

Why This Best Practice Is Important:

Downloading any of the Dynamic Updates or any versions of the Software requires that you not only have a valid license but that your firewall has a downloaded copy of your licenses.

How to Implement It:

Go to **Device > Licenses > License Management => Retrieve licenses keys from license server.**

What It Looks Like After You've Implemented It:

PA-VM Date Issued July 11, 2016 Date Expires Never Description Standard VM-300	AutoFocus Device License Date Issued October 19, 2016 Date Expires October 11, 2021 Description AutoFocus Device License
BrightCloud URL Filtering Date Issued July 12, 2016 Date Expires July 12, 2019 Description BrightCloud URL Filtering Active No (Activate)	GlobalProtect Gateway Date Issued July 12, 2016 Date Expires July 12, 2019 Description GlobalProtect Gateway License
PAN-DB URL Filtering Date Issued July 12, 2016 Date Expires July 12, 2019 Description Palo Alto Networks URL Filtering License Active Yes Download Status 2016-12-19 15:47:55 PAN-DB download: Finished successfully. Re-Download	Threat Prevention Date Issued July 12, 2016 Date Expires July 12, 2019 Description Threat Prevention
WildFire License Date Issued July 12, 2016 Date Expires July 12, 2019 Description WildFire signature feed, integrated WildFire logs, WildFire API	License Management Retrieve license keys from license server Activate feature using authorization code Manually upload license key Deactivate VM

These licenses are installed and current

❑ Configure Your Dynamic Updates Refresh Schedule

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The seventh of the eight steps you need to complete for minimal functionality is to configure the dynamic updates refresh schedule.

Your Palo Alto Networks firewall is just one member of a constellation of sensors and analytics that acts in concert to rapidly disseminate updates and malware intelligence.

Why This Best Practice Is Important:

For maximum effectiveness, your firewall needs to stay synchronized with this constellation and therefore be configured to regularly download and install the freshest dynamic updates.

How to Implement It:

Go to **Device > Dynamic Updates**.

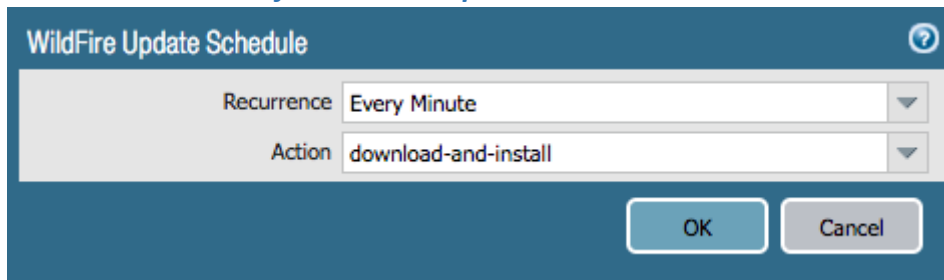
Version	File Name	Features	Type	Size	Release Date
Antivirus Last checked: 2016/12/22 08:37:03 PST Schedule: Every hour at 37 minutes past the hour (Download and Install, sync-to-peer)					
2100-2585	panup-inc-antivirus-2100-2585		Incremental	15 MB	2016/12/21 13:57:12 PST
2099-2584	panup-inc-antivirus-2099-2584		Incremental	15 MB	2016/12/20 13:55:33 PST
Applications and Threats Last checked: 2016/12/22 00:19:02 PST Schedule: Every day at 00:19 (Download and Install)					
647-3753	panupv2-all-contents-647-3753	Apps, Threats	Full	35 MB	2016/12/20 03:42:18 PST
648-3755	panupv2-all-contents-648-3755	Apps, Threats	Full	35 MB	2016/12/21 07:07:52 PST
GlobalProtect Data File Schedule: Every hour at 7 minutes past the hour (Download and Install)					
1482421206					2016/12/22 15:40:06
WildFire Last checked: 2016/12/22 08:46:22 PST Schedule: Every minute (Download and Install)					
99345-100399	panupv2-all-wildfire-99345-100399	PAN-OS 7.1 and later	Full	6 MB	2016/12/22 08:42:00 PST
99344-100398	panupv2-all-wildfire-99344-100398	PAN-OS 7.1 and later	Full	6 MB	2016/12/22 08:37:04 PST

The Dynamic Updates page showing the update schedules.

Here are some guidelines for configuring the schedules:

- Choose the shortest (most frequent) Recurrence so you get updates as quickly as possible.
- Consider the bandwidth requirements for frequent updates. The size of a WildFire update is documented on the Dynamic Updates page.
- Choose the *download-and-install* Action so the updates take effect immediately. The only exception is with Applications and Threats in which you may wish to enable *Disable new apps in content update* to enable you to manually review App-ID updates before they're installed.

What It Looks Like After You've Implemented It:



The WildFire Update Schedule Dialog Box, one of several you need to configure

□ Configure Service Routes If Necessary

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The eighth of the eight steps you need to complete for minimal functionality is to configure the Service Routes if necessary.


The firewall itself needs to make approximately 23 different types of connections to the Internet to support its own internal processes and functionality. These include the usual suspects like DNS and NTP, as well as many others. By default, these connections exit the firewall by the management port.

Why This Best Practice Is Important:

Depending upon your topology, the management interface might not have a route to the Internet. In this case, you need to provide a route from one of your “traffic” ports. By changing the service route configuration, you tell the firewall which physical port to use. Service routes allow the NGFW to use other interfaces beside the default Management interface to send specific Service traffic. For instance, perhaps your Management interface is not allowed to access the Internet from your edge Palo Alto Networks firewall, a Service Route using the external interface of the firewall might be needed in order to ensure your NTP traffic reaches their Internet destinations.

How to Implement It:

Go to **Device > Setup > Services > Services Features**.



The Service Route Configuration Dialog Box with its default setting.

What It Looks Like After You've Implemented It:

The screenshot shows the 'Service Route Configuration' window in Palo Alto Networks firewalls. The 'IPv4' tab is active, and the 'Customize' radio button is selected. Below the tabs is a table of service routes. All services listed are checked, and all are configured to use 'ethernet1/1' as the source interface and '172.16.31.253' as the source address. At the bottom of the table is a button labeled 'Set Selected Service Routes'. The window has 'OK' and 'Cancel' buttons at the bottom right.

<input checked="" type="checkbox"/>	Service	Source Interface	Source Address
<input checked="" type="checkbox"/>	AutoFocus	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	CRL Status	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	DNS	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	Email	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	HSM	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	Kerberos	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	LDAP	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	MDM	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	Netflow	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	NTP	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	Palo Alto Updates	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	Panorama	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	Proxy	ethernet1/1	172.16.31.253
<input checked="" type="checkbox"/>	RADIUS	ethernet1/1	172.16.31.253

The service routes are now configured to a traffic interface and not the management interface.

Management Interface

□ Physically and Logically Isolate Your Management Network

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Your Palo Alto Networks firewall can be managed only through the physical Management port (MGT), which is a separate physical Ethernet port on the appliance. In virtual firewalls it's also a separate interface.

There are three reasons why the Management port is a physically separate port:

- The Management Plane and the Data Plane are separate, so it makes sense to have them use separate interfaces.
- You don't want heavy traffic on the Data Plane interfaces to degrade performance on the Management Plane connections.
- A separate physical port for the Management Plane is the first step in physically and logically isolating this security sensitive network.

Why This Best Practice Is Important:

Anyone who can log in to the firewall's GUI interface has the keys to the kingdom, so it's reasonable to protect this interface with additional measures. By making it physically and logically separate from not only the Internet but also the rest of your network, you're providing enhanced protection.

How to Implement It:

Physical Isolation:

Connect this port to an isolated physical network. Provide a management workstation with a web browser, but ensure there are no physical connections to any other networks. If you need remote access, use a secure remote access gateway that requires two factor authentication.

Logical Isolation:

Use a unique non-routable IP subnet to ensure that even if someone accidentally connects an Ethernet cable where they shouldn't, no IP protocol routing path exists to connect to the management interface from anywhere else.

What It Looks Like After You've Implemented It:

In order to connect to the Management interface, you should be forced to be in the right physical location, connected to the right physical network, and provide the right credentials, and all that's before you get a chance to present credentials to the firewall GUI interface.

What Else You Need to Know:

Don't forget that except for some extreme high security configurations, the firewall needs to be able to connect to the Palo Alto Networks cloud for various updates, etc. It's best to configure Service Routes so these connections don't have to originate from the Management Interface.

❑ Restrict Permitted IP Addresses Connecting to the Management Interface

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The firewall allows you to configure the Management Interface to accept connections only from a fixed set of source IP addresses.

Why This Best Practice Is Important:

By restricting Management Interface access to a fixed set of source IP addresses, you further lock down access to the Management GUI.

How to Implement It:

Go to **Device > Setup > Management > Management Interface Settings**

The screenshot shows the 'Management Interface Settings' window. On the left, there are fields for IP Type (Static selected), IP Address (192.168.44.2), Netmask (255.255.255.0), Default Gateway (192.168.44.254), IPv6 Address/Prefix Length, Default IPv6 Gateway, Speed (auto-negotiate), and MTU (1500). In the center, there is a 'Services' list with checkboxes for HTTP, HTTP OCSP, HTTPS (checked), Telnet, SSH (checked), Ping (checked), SNMP, User-ID, User-ID Syslog Listener-SSL, and User-ID Syslog Listener-UDP. On the right, there is a 'Permitted IP Addresses' list with a single entry, 192.168.44.1, which is circled in red. At the bottom right of the list are 'Add' and 'Delete' buttons. At the bottom of the window are 'OK' and 'Cancel' buttons.

Connections from only a single IP address are permitted to this Management Interface

❑ Restrict Permitted Services Hosted by the Management Interface

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The firewall allows you to configure which services will respond on the Management Interface.

Why This Best Practice Is Important:

Because we want to provide maximum security to the Management Interface, it's best to lock it down so it responds to as few services as necessary.

How to Implement It:

Go to **Device > Setup > Management > Management Interface Settings**

Some guidance:

- Always prefer HTTPS over HTTP.
- Enable SSH because it's secure and it allows you to use the CLI.
- Enable Ping when you want to test connectivity.
- Disable everything else.
- If you enable Telnet, you should consider other equally rewarding careers.

What It Looks Like After You've Implemented It:

The screenshot shows the 'Management Interface Settings' window. The 'IP Type' is set to 'Static' with IP Address '192.168.44.2' and Netmask '255.255.255.0'. The 'Default Gateway' is '192.168.44.254'. The 'Speed' is 'auto-negotiate' and 'MTU' is '1500'. In the 'Services' section, the following services are checked: HTTPS, SSH, and Ping. All other services are unchecked. On the right, the 'Permitted IP Addresses' list contains '192.168.44.1'.

This is a good, minimal set of Services to allow on the Management Interface

❑ Configure Failed Attempts and Lockout Time

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The firewall allows you to configure the Management Interface to temporarily lock out further authentication attempts after a configured number of failed attempts.

Why This Best Practice Is Important:

In a Brute Force Attack, an attacker simply tries all possible passwords in an authentication system's password space. Given that the password space is finite, we know that our attackers can always defeat us if they have enough time and computers, so our job is to bury them in delay and extra work.

The first defense against a Brute Force Attack is to use a long, complicated password. This Best Practice helps with the second defense, which is to introduce a time delay into the attack.

How to Implement It:

Go to **Device > Setup > Management > Authentication Settings**

The screenshot shows the 'Authentication Settings' dialog box. The 'Failed Attempts' field is set to 10 and the 'Lockout Time (min)' field is set to 30. These two fields are circled in red. Other fields include 'Authentication Profile' (None), 'Certificate Profile' (None), and 'Idle Timeout (min)' (240). The 'OK' and 'Cancel' buttons are at the bottom right.

These are reasonable settings

Even the most generous settings, of 10 attempts and 60 minutes, are still exponentially better than no settings at all. We don't really mind if they try 10 passwords per hour, we just won't let them try *10,000 passwords per second*.

What Else You Need to Know:

The range of values for Failed Attempts is [0-10], and a value of 0 means "unlimited attempts permitted and there won't be any lockouts."

The range of values for Lockout Time is [0-60] minutes, and a value of 0 means "permanently locked out until another administrator successfully logs in."

❑ Configure the Management GUI and CLI Idle Timeout

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The firewall lets you configure the idle timeout for the Management GUI and CLI.

Why This Best Practice Is Important:

Assuming you always log out when you get up from your chair, you should never need this, so the idle timeout is designed to automatically log you out only after so many minutes of inactivity when you forget to do it yourself.

Here's the complicated algorithm for deciding on the correct value:

One the one hand...

You want enhanced security, so if you've forgotten to log off you'd like the idle timeout to log you out rather quickly. This is an argument for making this value small.

But on the other hand...

Sometimes when you're deep into troubleshooting, you have many tools and tasks going at once and you might spend a bit of time away from one while you're working on the others, and it's annoying to keep getting logged out. This is an argument for making the value large.

A good compromise might be somewhere between 30 and 120 minutes.

How to Implement It:

Go to **Device > Setup > Management > Authentication Settings**.

The screenshot shows the 'Authentication Settings' dialog box. The 'Idle Timeout (min)' field is circled in red, indicating the default value of 60 minutes. The other fields are: 'Authentication Profile' (None), 'Certificate Profile' (None), 'Failed Attempts' (0), and 'Lockout Time (min)' (0). The 'OK' and 'Cancel' buttons are at the bottom right.

The default value is 60 minutes

What Else You Need to Know:

Valid settings range from 1 to 1440 minutes (24 hours).

The default value is 60 minutes.

A setting of 0 means "never timeout".

Administrators

□ Create a Unique Administrator Account for Each Administrator

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Administrator accounts control access to NGFWs and Panorama. To gain access to the management GUI, the administrator must first:

- *Identify* themselves with a username, and
- *Authenticate* themselves with a password or other method

PAN-OS allows the creation and use of multiple Administrator accounts.

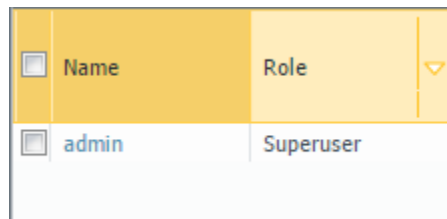
Why This Best Practice Is Important:

Giving each firewall administrator a unique Administrator account gives you these benefits:

- **Access Control:** You can give different permissions and access to different administrators, and can modify or delete these permissions without affecting other administrators.
- **Auditing:** You can view and audit the actions of specific administrators.

How to Implement It:

Go to **Device > Administrators**.



Name	Role
admin	Superuser

By default, the firewall has only a single Superuser administrator. This is what it looks like after a fresh install.

What It Looks Like After You've Implemented It:

<input type="checkbox"/>	Name	Role
<input type="checkbox"/>	admin	Superuser
<input type="checkbox"/>	bstiefel	Device administrator
<input type="checkbox"/>	rhall	Device administrator (read-only)
<input type="checkbox"/>	jsieber	Superuser (read-only)
<input type="checkbox"/>	nasingh	Superuser

Here are multiple administrator accounts, with different roles

□ Audit Your Administrator Accounts Regularly

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Firewall administrators hold the keys to the kingdom. When someone leaves the organization or steps away from firewall administration duties, it's easy to forget to delete their administrator account.

Why This Best Practice Is Important:

Orphaned administrator accounts are a security risk. By eliminating them, you're reducing your attack surface.

How to Implement It:

Set up a scheduled event that reoccurs every 90-180 days. When it's time, go to **Device > Administrators** and review the administrator accounts. Delete any that aren't currently needed. It's OK to be aggressive; you can always add one back later.

□ Create and Use Admin Role Profiles

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The security model built in to PAN-OS provides for fine-grained permission setting for administrators. Whether it's for access using the Web UI, XML API, or the Command Line, you can specify in great detail the permissions granted for each individual administrator.

Why This Best Practice Is Important:

Following the *Principle of Least Privilege*, it's important to limit each administrator's permissions to those required to perform their duties.

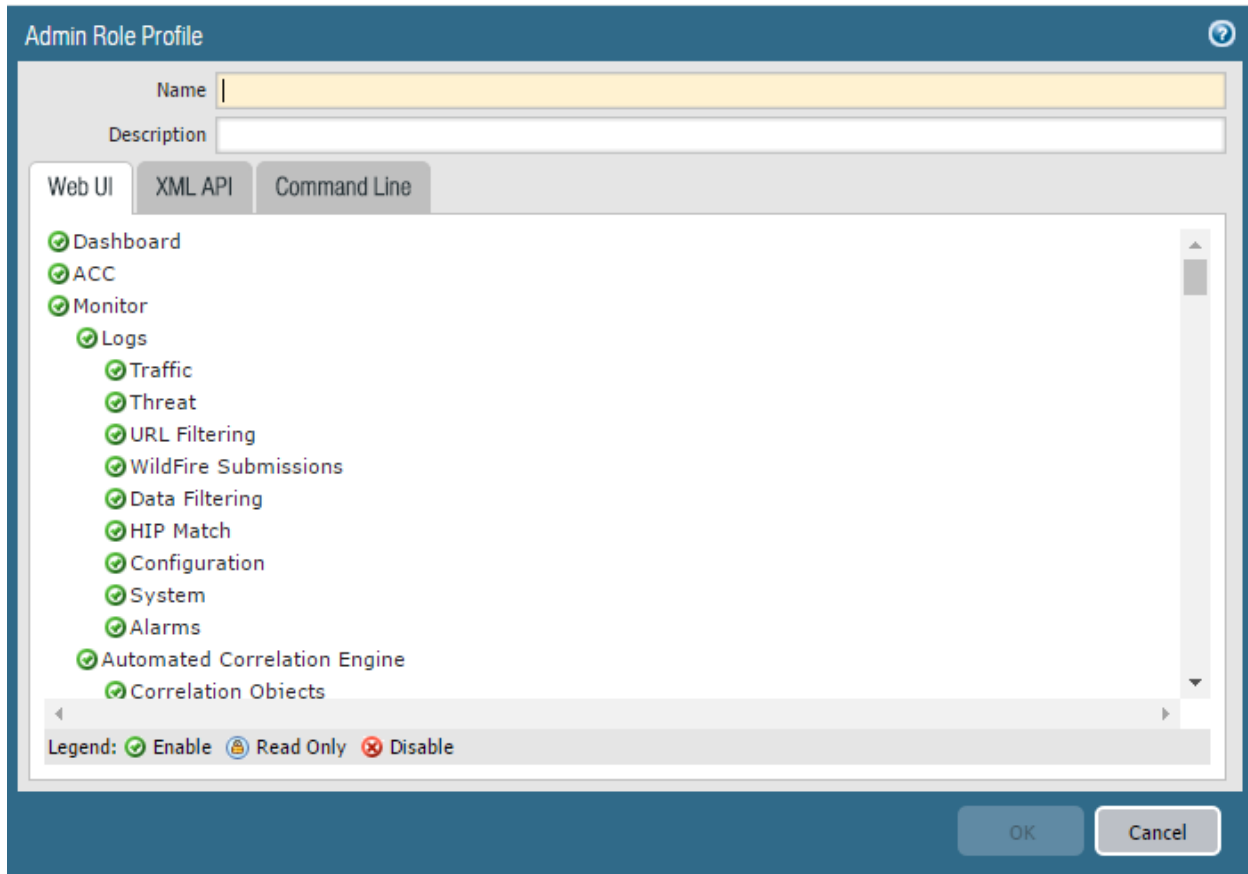
How to Implement It:

You can view and manage the Admin Role Profiles by going to **Device > Admin Roles**. On this page you'll see the three default read-only roles:

<input type="checkbox"/>	Name	Description	Role	CLI Role
<input type="checkbox"/>	auditadmin	Audit Administrator for Common Criteria	device	
<input type="checkbox"/>	cryptoadmin	Crypto Administrator for Common Criteria	device	
<input type="checkbox"/>	securityadmin	Security Admin for Common Criteria	device	

The three default read-only Admin Roles

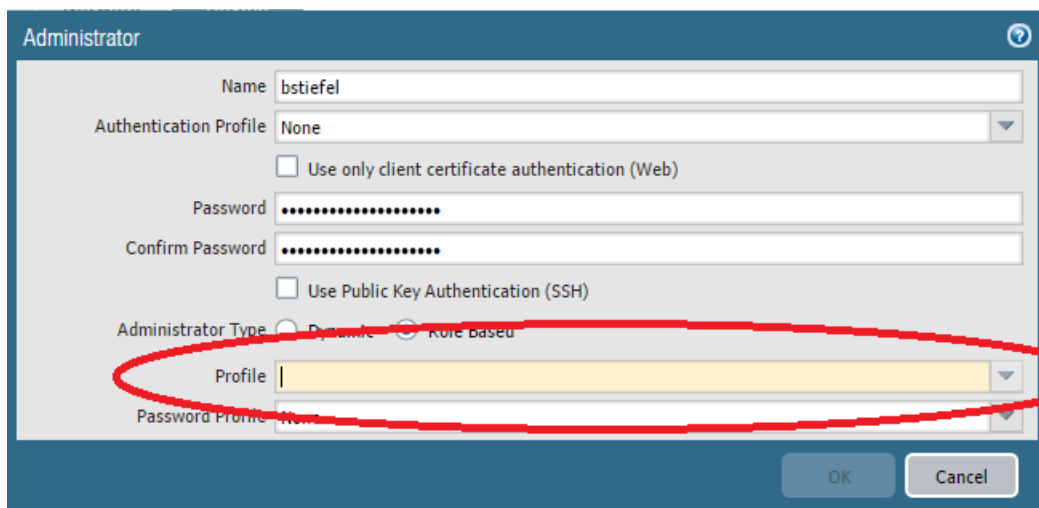
The *Admin Role Profile* dialog box allows you to view and manage the permissions within a Profile:



The **Admin Role Profile** dialog box is shown. It has a title bar with a question mark icon. Below the title bar are fields for **Name** and **Description**. There are three tabs: **Web UI** (selected), **XML API**, and **Command Line**. The **Web UI** tab contains a list of features with checkboxes: **Dashboard**, **ACC**, **Monitor**, **Logs**, **Traffic**, **Threat**, **URL Filtering**, **WildFire Submissions**, **Data Filtering**, **HIP Match**, **Configuration**, **System**, **Alarms**, **Automated Correlation Engine**, and **Correlation Objects**. All checkboxes are checked. At the bottom of the list is a legend: **Legend: ☒ Enable ☐ Read Only ☒ Disable**. At the bottom right are **OK** and **Cancel** buttons.

The Admin Role Profile Dialog Box

Once you've created and configured the Admin Role Profile, you need to assign it to an administrator account. Go to **Device > Administrators** and edit the account:



The **Administrator** dialog box is shown. It has a title bar with a question mark icon. Below the title bar are fields for **Name** (bstiefel), **Authentication Profile** (None), **Use only client certificate authentication (Web)** (unchecked), **Password** (masked), **Confirm Password** (masked), **Use Public Key Authentication (SSH)** (unchecked), **Administrator Type** (Dynamic), **Profile** (empty), and **Password Profile** (None). A red oval highlights the **Profile** field. At the bottom right are **OK** and **Cancel** buttons.

Configuring the Admin Role Profile for an Administrator account

□ Configure Minimum Password Complexity

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

A Brute Force Attack is a trial-and-error method used to guess a user password by employing automated software to generate a large number of consecutive guesses.

Why This Best Practice Is Important:

A cursory look at the mathematics shows us that our attackers can always defeat us if they have enough time and computers, so our best defense is to bury them in extra work. The PAN-OS Minimum Password Complexity settings allow you to force administrators to use long, complex, hard-to-guess passwords to slow the effectiveness of these attacks.

How to Implement It:

Go to **Device > Setup > Management > Minimum Password Complexity**. This is what the settings look like by default:

Minimum Password Complexity ⓘ

☐ Enabled

Password Format Requirements

Minimum Length

Minimum Uppercase Letters

Minimum Lowercase Letters

Minimum Numeric Letters

Minimum Special Characters

Block Repeated Characters

☐ Block Username Inclusion (including reversed)

Functionality Requirements

New Password Differs By Characters

☐ Require Password Change on First Login

Prevent Password Reuse Limit

Block Password Change Period (days)

Required Password Change Period (days)

Expiration Warning Period (days)

Post Expiration Admin Login Count

Post Expiration Grace Period (days)

Functionality requirements can be overridden by password profiles

OK Cancel

The Minimum Password Complexity Dialog Box, after a fresh install

Both the help page associated with this dialog box and the online documentation provide details of each setting.

What It Looks Like After You've Implemented It:

Here are reasonably strong settings for Minimum Password Complexity:

Minimum Password Complexity

☒ Enabled

Password Format Requirements

Minimum Length: 15

Minimum Uppercase Letters: 1

Minimum Lowercase Letters: 1

Minimum Numeric Letters: 1

Minimum Special Characters: 1

Block Repeated Characters: 3

☒ Block Username Inclusion (including reversed)

Functionality Requirements

New Password Differs By Characters: 8

☒ Require Password Change on First Login

Prevent Password Reuse Limit: 50

Block Password Change Period (days): 1

Required Password Change Period (days): 90

Expiration Warning Period (days): 15

Post Expiration Admin Login Count: 3

Post Expiration Grace Period (days): 30

Functionality requirements can be overridden by password profiles

OK Cancel

These reasonably strong settings are much more secure than the defaults

❑ Avoid Password Profiles

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Password Profiles allow you to override some of the global Minimum Password Complexity settings for a specific administrator.

To configure Password Profiles, go to **Device > Password Profiles** or **Panorama > Password Profiles**. When you create or edit a Password Profile, you'll see the Password Profiles dialog box:

Configuring a Password Profile, which you don't want to do

To assign a Password Profile to an administrator, go to **Device > Administrators** and edit the administrator account:

This is where you assign a Password Profile to an administrator account, which you don't want to do

Why This Best Practice Is Important:

An important clue comes in small print at the bottom of the Minimum Password Complexity dialog box, found at

Device > Setup > Management > Minimum Password Complexity:

Minimum Password Complexity

☒ Enabled

Password Format Requirements

Minimum Length: 15

Minimum Uppercase Letters: 1

Minimum Lowercase Letters: 1

Minimum Numeric Letters: 1

Minimum Special Characters: 1

Block Repeated Characters: 3

☒ Block Username Inclusion (including reversed)

Functionality Requirements

New Password Differs By Characters: 8

☒ Require Password Change on First Login

Prevent Password Reuse Limit: 50

Block Password Change Period (days): 1

Required Password Change Period (days): 90

Expiration Warning Period (days): 15

Post Expiration Admin Login Count: 3

Post Expiration Grace Period (days): 30

Functionality requirements can be overridden by password profiles

OK Cancel

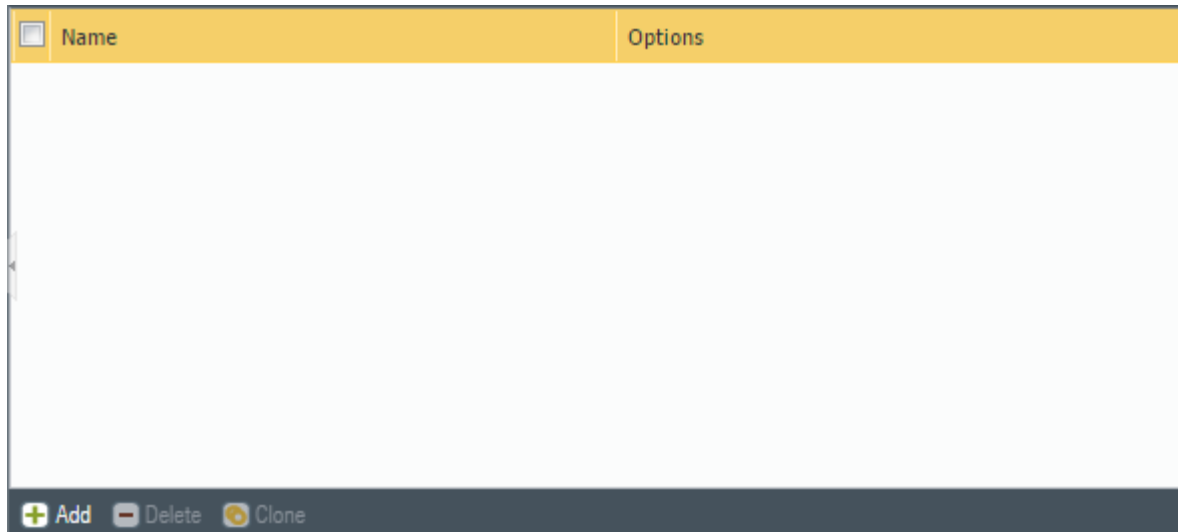
Read the fine print

Any Password Profile assigned to an administrator account will overrule the global settings for Minimum Password Complexity. This is dangerous as it could quietly reduce the security compliance requirements for passwords. By ensuring you have no Password Profiles configured, you maintain compliance with the Minimum Password Complexity settings.

How to Implement It:

Go to **Device > Password Profiles** or **Panorama > Password Profiles** and ensure you have no Password Profiles configured.

What It Looks Like After You've Implemented It:



When the list is empty like this, that means you're doing it right

Management

□ Configure the Login Banner

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The login procedure is a chokepoint where you can force the user to read Terms of Use before logging in.

Why This Best Practice Is Important:

Here is your chance to explicitly state what sort of users and activities are permitted on your firewall. By making things clear here you can set expectations, frighten off evildoers and protect your legal rights.

How to Implement It:

The Login Banner is your chance to make explicit statements such as:

- Only authorized use is permitted
- All use may be monitored
- The user should have no expectation of privacy
- While not using words like "Welcome"
- While not identifying the owner, name, location, or use of the equipment

A reasonable Login Banner might look something like this:

WARNING: The use of this system is restricted to authorized users only.

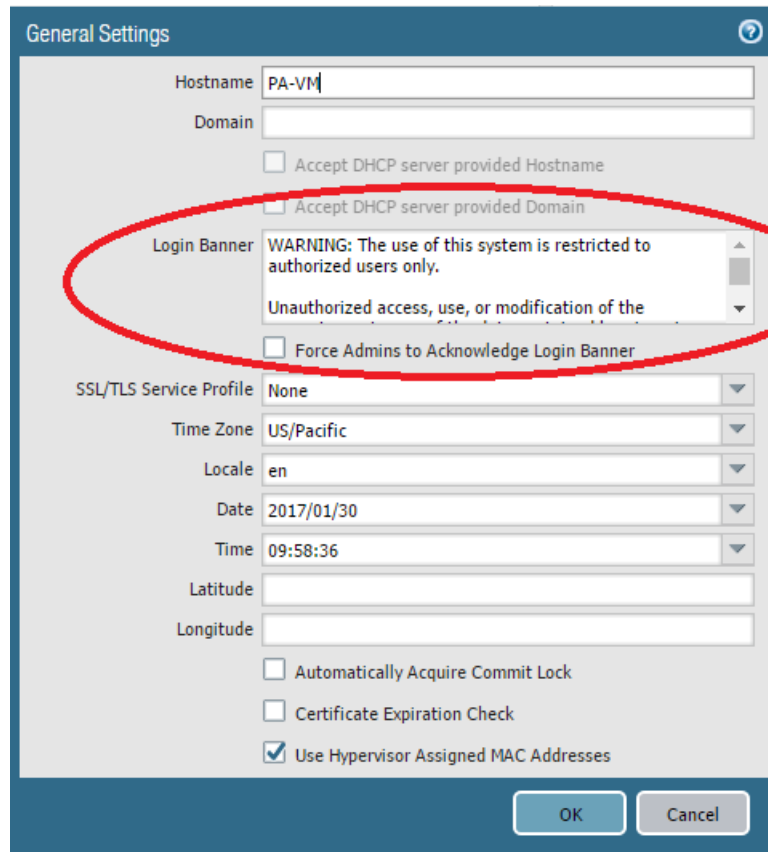
Unauthorized access, use, or modification of the computer system or of the data contained herein or in transit to/from this system may subject you to criminal prosecution.

These systems and equipment are subject to monitoring to ensure proper performance of applicable security features or procedures.

Such monitoring may result in the acquisition, recording and analysis of all data being communicated, transmitted, processed or stored in the system by a user.

If monitoring reveals possible evidence of criminal activity, such evidence may be provided to law enforcement personnel.

Go to **Device > Setup > Management > General Settings**



The screenshot shows the 'General Settings' dialog box in the Palo Alto Networks management interface. The 'Login Banner' section is highlighted with a red oval. It contains a text area with the following text: 'WARNING: The use of this system is restricted to authorized users only.' and 'Unauthorized access, use, or modification of the'. Below this text area is a checkbox labeled 'Force Admins to Acknowledge Login Banner'. Other settings visible include 'Hostname' (PA-VM), 'Domain', 'Accept DHCP server provided Hostname', 'Accept DHCP server provided Domain', 'SSL/TLS Service Profile' (None), 'Time Zone' (US/Pacific), 'Locale' (en), 'Date' (2017/01/30), 'Time' (09:58:36), 'Latitude', 'Longitude', 'Automatically Acquire Commit Lock', 'Certificate Expiration Check', and 'Use Hypervisor Assigned MAC Addresses' (checked). The 'OK' and 'Cancel' buttons are at the bottom right.

Now we've made ourselves clear

The image shows the Palo Alto Networks login page. At the top is the Palo Alto Networks logo, which consists of a blue square with white diagonal lines and the text "paloalto NETWORKS" in blue and green. Below the logo are two input fields: "Name" and "Password". The "Name" field is a white box with a blue border, and the "Password" field is a white box with a grey border. Below these fields is a grey "Login" button. At the bottom of the page is a warning banner with a blue border and a scroll bar. The banner contains the following text: "WARNING: The use of this system is restricted to authorized users only. Unauthorized access, use, or modification of the computer system or of the data contained here. These systems and equipment are subject to monitoring to ensure proper performance of applicab. Such monitoring may result in the acquisition, recording and analysis of all data being commun. If monitoring reveals possible evidence of criminal activity, such evidence may be provided to".

paloalto
NETWORKS

Name

Password

Login

WARNING: The use of this system is restricted to authorized users only.
Unauthorized access, use, or modification of the computer system or of the data contained here.
These systems and equipment are subject to monitoring to ensure proper performance of applicab.
Such monitoring may result in the acquisition, recording and analysis of all data being commun.
If monitoring reveals possible evidence of criminal activity, such evidence may be provided to

And here it is in action

What Else You Need to Know:

Check with your Legal and Human Resources departments for their guidance on what should be included in the Login Banner.

❑ Force Admins to Acknowledge the Login Banner

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Why This Best Practice Is Important:

In some legal jurisdictions it might not be enough to merely post the Terms of Use on the login page and claim the user has thus “agreed” to follow them. This setting allows you to force the user to click acknowledgement of the terms, providing additional binding to a “click-through” agreement.

How to Implement It:

Go to **Device > Setup > Management > General Settings**

The screenshot shows the 'General Settings' window in the Palo Alto Networks management interface. A red circle highlights the 'Login Banner' section. Within this section, the checkbox 'Force Admins to Acknowledge Login Banner' is checked. The login banner text is visible, showing a warning and a line about unauthorized access. Other settings like Hostname, Domain, Time Zone, and various checkboxes are also visible.

The second step of configuring the Login Banner, forcing the admins to acknowledge it

What It Looks Like After You've Implemented It:

A screenshot of the Palo Alto Networks login page. At the top center is the Palo Alto Networks logo. Below it are two input fields: 'Name' and 'Password'. A 'Login' button is positioned below the password field. At the bottom of the page, there is a rectangular box containing a warning statement. To the left of this box is a checkbox with the text 'I Accept and Acknowledge the Statement Below'.

Now a “click-through” agreement is required to proceed

What Else You Need to Know:

Here's a secret “feature”: If the “Force Admins to Acknowledge Login Banner” box is unchecked, then the box containing the Login Banner on the Login page will not expand to show all the text. If the box is checked, the box will expand properly and show all the text.

□ Configure Header and Footer Banners

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

PAN-OS allows you to configure optional Header and Footer Banners that display on every page in the GUI.

Why This Best Practice Is Important:

It doesn't take much hopping between firewall user interfaces before they all start to look alike and before you know it you're making configuration changes on the wrong firewall because you forgot where you are.

By configuring a Header or Footer Banner you can have a visual reminder on every page of the GUI showing you exactly which machine you're configuring. You can also provide other information or reminders to firewall administrators.

How to Implement It:

Go to **Device > Setup > Management > Banners and Messages**.

Banners and Messages

☒ **Message of the Day**

Message of the Day

☐ Allow Do Not Display Again

Title

Background Color

Icon

Banners

Header Banner

Header Color

Header Text Color

☐ Same banner for header and footer


Footer Banner

Footer Color


Footer Text Color

OK Cancel

Some reasonable settings for Header and Footer Banners



Header, as displayed



Footer, as displayed

❑ Configure Custom Logos in PAN-OS

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

PAN-OS allows you to configure optional custom logos for the following:

- Login Screen
- Main UI
- PDF Report Title Page
- PDF Report Footer

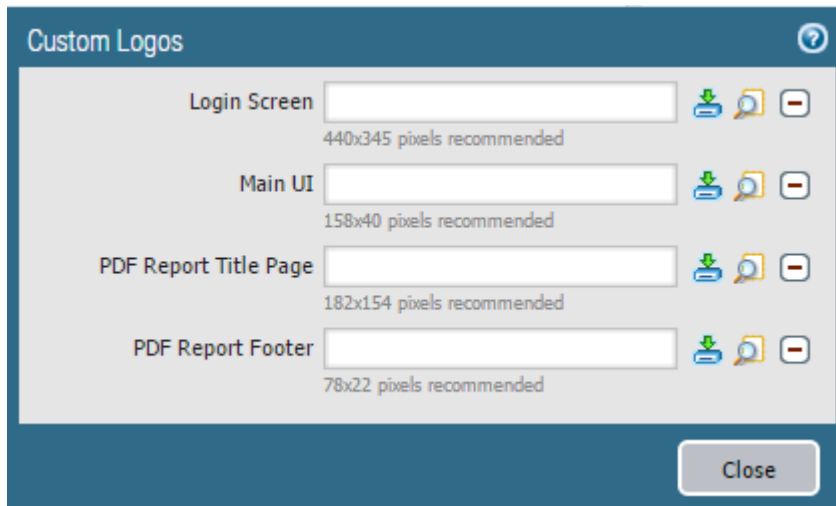
Why This Best Practice Is Important:

For the Login Screen, you might want to remove the Palo Alto Networks logo to reduce the information that leaks back to an unauthorized person who attempts to login.

For the other logos, you might want to replace them with your own organization's logo to present more consistent branding.

How to Implement It:

Go to **Device > Setup > Operations > Miscellaneous > Custom Logos**.



Configure them as you wish

❑ Change the Default Master Key

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Every PAN-OS and Panorama device has a Master Key that encrypts all the passwords and private keys in the configuration. Upon initial installation, the key is set to its default value.

Why This Best Practice Is Important:

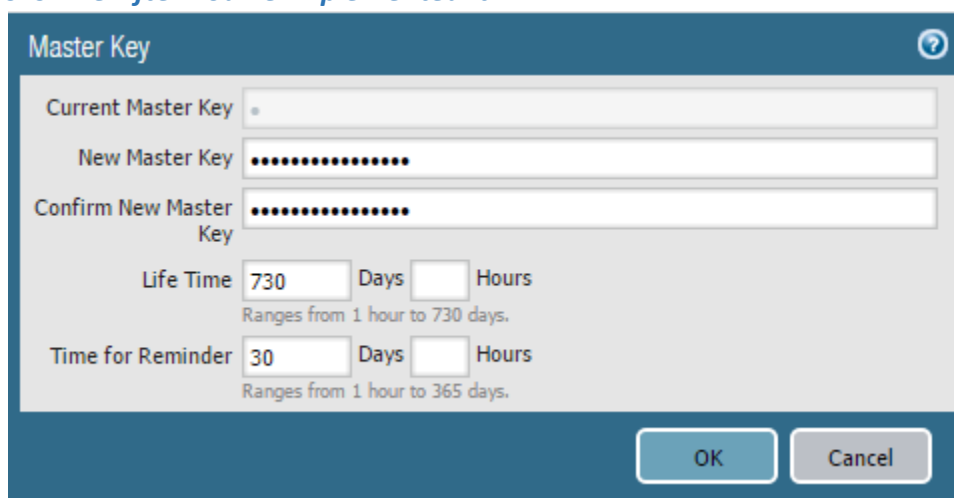
The default Master Key is insecure, because it is set to a known default value. You need to change it.

How to Implement It:

Go to **Device > Master Key and Diagnostics**

Configure the Master Key and its settings here

The Master Key length must be exactly 16 characters. Acceptable characters include letters, numbers, and special characters except '\$' and '&'.

What It Looks Like After You've Implemented It:The image shows a 'Master Key' configuration window. It has a title bar with a question mark icon. Inside, there are three text input fields: 'Current Master Key' (empty), 'New Master Key' (filled with dots), and 'Confirm New Master Key' (filled with dots). Below these are two time settings. The first is 'Life Time' with a value of '730' in a box, followed by 'Days' and 'Hours' boxes. A note below says 'Ranges from 1 hour to 730 days.' The second is 'Time for Reminder' with a value of '30' in a box, followed by 'Days' and 'Hours' boxes. A note below says 'Ranges from 1 hour to 365 days.' At the bottom right are 'OK' and 'Cancel' buttons.

A new Master Key and timeout and reminder settings

What Else You Need to Know:

For HA to synchronize properly, both PAN-OS and Panorama devices must have the same Master Key.

Store the Master Key in a safe location as it cannot be recovered. The only way to restore the Default Master Key is to do a configuration reset to factory default settings.

If you are using Panorama to manage your firewalls, you must use the same Master Key on Panorama and all the firewalls to enable proper communication.

❑ Encrypt the Master Key with a Hardware Security Module (HSM)

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Every PAN-OS and Panorama device has a Master Key that encrypts all the passwords and private keys in the configuration.

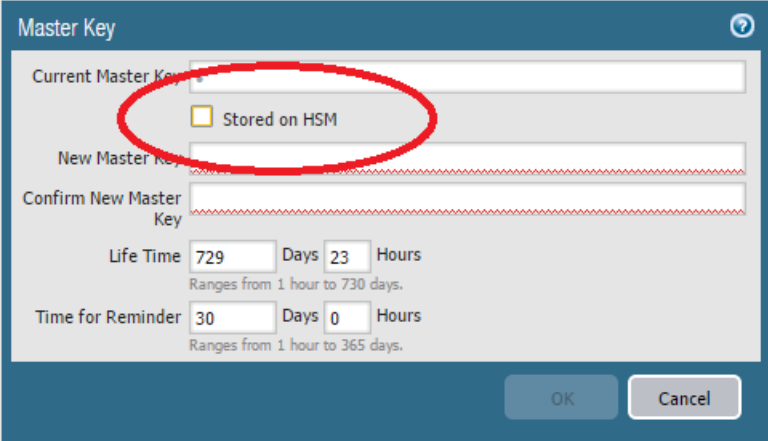
If your security requirements include storing your private keys in a secure location, you can encrypt the Master Key using an encryption key – known as a wrapping key – that is stored on an external Hardware Security Module (HSM). When the firewall or Panorama needs access to the Master Key, it requests the HSM to decrypt it. The HSM is a physically separate secure device for storing the Master Key.

Why This Best Practice Is Important:

Some organizations require enhanced protection of private keys. A particularly effective way to do this is to use an HSM.

How to Implement It:

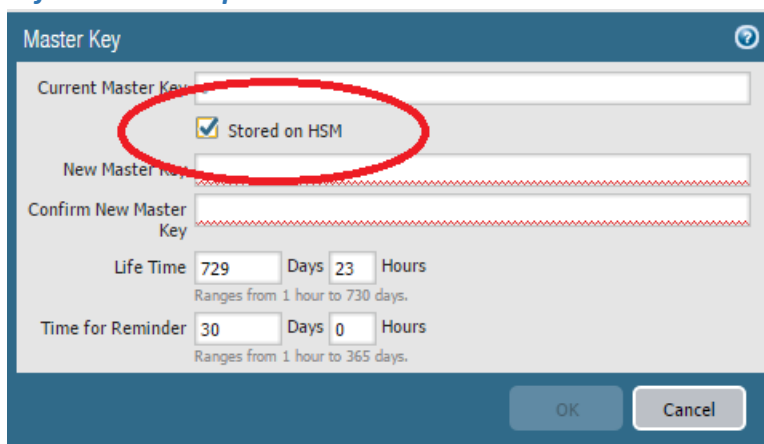
Go to **Device > Master Key and Diagnostics**



The screenshot shows the 'Master Key' configuration window. It includes fields for 'Current Master Key', 'New Master Key', and 'Confirm New Master Key'. A checkbox labeled 'Stored on HSM' is highlighted with a red circle. Below these fields are 'Life Time' and 'Time for Reminder' settings, each with a numeric input and a dropdown for 'Days' and 'Hours'. The 'Life Time' dropdown is set to '729' and the 'Time for Reminder' dropdown is set to '30'. At the bottom right are 'OK' and 'Cancel' buttons.

Once the HSM is installed, this checkbox will appear

What It Looks Like After You've Implemented It:



The screenshot shows the 'Master Key' configuration window. It includes fields for 'Current Master Key', 'New Master Key', and 'Confirm New Master Key'. A red circle highlights the 'Stored on HSM' checkbox, which is checked. Below these fields are 'Life Time' and 'Time for Reminder' settings, each with a numeric input, 'Days', and 'Hours' sub-inputs. The 'Life Time' is set to 729 days and 23 hours, with a range of 1 hour to 730 days. The 'Time for Reminder' is set to 30 days and 0 hours, with a range of 1 hour to 365 days. 'OK' and 'Cancel' buttons are at the bottom right.

Now you're good to go

What Else You Need to Know:

Firewalls configured in FIPS/CC mode do not support master key encryption using an HSM.

□ Set Certificate Expiration Check

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Palo Alto Networks' firewalls keep several certificates on-box. Each of these certificates has an expiration date. When they expire, they will, at a minimum, generate a warning upon use, if they're accepted at all.

Why This Best Practice Is Important:

You don't want to have a certificate expire because it will cause either a warning or a service outage.

Configuring Certificate Expiration Check ensures the firewall creates warning messages when on-box certificates are near their expiration dates.

How to Implement It:

Go to **Device > Setup > Management > General Settings**.

The screenshot shows the 'General Settings' window in the Palo Alto Networks management interface. The 'Certificate Expiration Check' checkbox is checked and highlighted with a red circle. Other visible settings include: Hostname (PA-VM), Domain (empty), Login Banner (WARNING: The use of this system is restricted to authorized users only.), Force Admins to Acknowledge Login Banner (checked), SSL/TLS Service Profile (None), Time Zone (US/Pacific), Locale (en), Date (2017/01/30), Time (16:01:31), Latitude (empty), Longitude (empty), Automatically Acquire Commit Lock (unchecked), and Use Hypervisor Assigned MAC Addresses (checked).

Now you'll be warned when on-box certificates near their expiration

□ Configure Your Geographic Location

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

PAN-OS can provide global maps to show you the origin and destination of traffic and threats. You can find these maps here:

Monitor > App Scope > Threat Map

Monitor > App Scope > Traffic Map

For reasons that remain a mystery, PAN-OS by default assumes you're in the southern Indian Ocean:



Your firewall is on the Endurance with Ernest Shackleton

Why This Best Practice Is Important:

It's a minor point, but you might as well have the map match reality by information PAN-OS of the approximate latitude and longitude of your firewall. After too many late nights troubleshooting networking problems you might forget which continent you're on.

How to Implement It:

Go to **Device > Setup > Management > General Settings**.

It's best to fill this in

An easy way to determine your latitude and longitude is to go to Google Maps and search for your street address. In the URL, you'll see the latitude and longitude ready for copying.

For example, a search for the headquarters of Palo Alto Networks – 4401 Great America Pkwy, Santa Clara, CA 95054 – generates this URL:

<https://www.google.com/maps/place/4401+Great+America+Pkwy,+Santa+Clara,+CA+95054/@37.3953852,-121.97878...>

From that you can extract the following:

Latitude: 37.3953852

Longitude: -121.97878

What It Looks Like After You've Implemented It:

The screenshot shows the 'General Settings' window in the Palo Alto Networks management interface. A red oval highlights the following fields:

- Date: 2017/01/31
- Time: 07:13:30
- Latitude: 37.3935531
- Longitude: -121.9788409

Other visible settings include:

- Hostname: PA-VM
- Domain: (empty)
- Accept DHCP server provided Hostname: ☐
- Accept DHCP server provided Domain: ☐
- Login Banner: WARNING: The use of this system is restricted to authorized users only. Unauthorized access, use, or modification of the system is prohibited.
- Force Admins to Acknowledge Login Banner: ☒
- SSL/TLS Service Profile: None
- Time Zone: US/Pacific
- Locale: en
- Automatically Acquire Commit Lock: ☐
- Certificate Expiration Check: ☐
- Use Hypervisor Assigned MAC Addresses: ☒

Now we know where we are



Land ho!

❑ Check Last Login Time and Look for Failed Login Attempts

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

PAN-OS provides two methods for helping to detect when an unauthorized person has attempted to log into the firewall:


- The “Last Login Time” field and warning symbol at the bottom left of every page in the GUI
- The “auth-fail” and “auth-success” Events in the System Log

Why This Best Practice Is Important:

Anyone to logs into your firewall GUI has the keys to the kingdom. It’s important to monitor who has been trying to log in to the system, regardless of whether they were successful.


How to Implement It:

Last Login indicator is enabled by default. This is what the Last Login Time indicator looks like when the previous login attempt was successful:



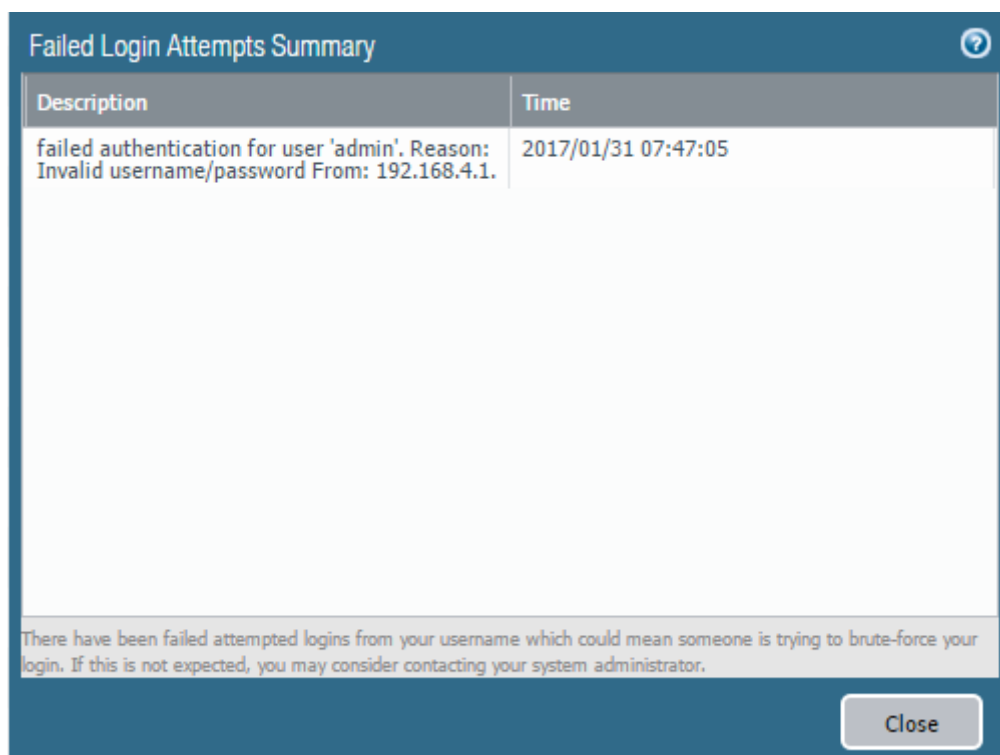
There is no warning symbol

This is what the last Login Time indicator looks like when the previous login attempt failed:



Notice the warning symbol

If you click on the triangular warning symbol, you’ll get this dialog box:



The Failed Login Attempts Summary dialog box

To view the authentication events in the System Log, go to **Monitor > Logs > System** and enter (**eventid eq auth-fail**) or (**eventid eq auth-success**) in the search field.

(eventid eq auth-fail) or (eventid eq auth-success)					
Receive Time	Type	Severity	Event	Object	Description
01/31 07:47:21	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
01/31 07:47:05	general	medium	auth-fail		failed authentication for user 'admin'. Reason: Invalid username/password From: 192.168.4.1.
01/31 07:45:26	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
01/31 06:53:36	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
01/30 15:55:14	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
01/30 14:46:02	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
01/30 10:36:41	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
01/30 10:26:51	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
01/30 09:14:00	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
01/30 09:12:49	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
01/30 09:09:33	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
12/22 08:09:28	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
12/21 14:16:43	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
12/20 13:29:29	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
12/20 13:29:17	general	medium	auth-fail		failed authentication for user 'admin'. Reason: Invalid username/password From: 192.168.4.1.
12/20 07:38:07	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
12/19 15:02:38	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.
12/19 15:02:27	general	medium	auth-fail		failed authentication for user 'admin'. Reason: Invalid username/password From: 192.168.4.1.
12/19 13:50:13	general	informational	auth-success		authenticated for user 'admin'. From: 192.168.4.1.

The System Log showing recent authentication events

By monitoring these logging events you can track attempted firewall logins.

❑ Choose SNMP V3 Over V2c

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Simple Network Management Protocol (SNMP) is a networking protocol for collecting and modifying data from managed network devices. It runs on UDP/161 and is documented in multiple Requests for Comments (RFCs).

Why This Best Practice Is Important:

SNMP Version V3 is newer than V2C and doesn't add any changes to the protocol other than encryption and authentication, but these two new features are critical for a secure configuration.

If you're going to configure your firewall to use SNMP, use V3 instead of V2C.

How to Implement It:

Go to **Device > Setup > Operations . Miscellaneous > SNMP Setup**

How to choose SNMP version V3 over V2c

❑ Configure the Statistics Service

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The Application and Threat Statistics Collection Service allows your firewall to send anonymous application and threat information to the Palo Alto Networks research team.

Why This Best Practice Is Important:

The research team is always trying to improve the company's products and services. By sending this anonymous real-world information on applications and threats, they're able to increase the speed of their continuous improvement, benefitting all customers.

How to Implement It:

Go to **Device > Setup > Operations > Miscellaneous > Statistics Service Setup**.

Statistics Service

The Application and Threat Statistics Collection Service allows your firewall to send anonymous application and threat information to the Palo Alto Networks research team. The collection of real-world information on applications and threats is an invaluable resource of research and will help Palo Alto Networks in our continuous efforts to extend our leadership in these areas.

The Statistics Collection Service is disabled by default and once enabled, information will be uploaded every 4 hours. The Statistics Collection Service can be disabled at any time. To see the data that will be submitted, view the Report Sample.

Settings | **Report Sample**

Application And Threat Reports

- ☐ Application usage
- ☐ Threats by destination ports
- ☐ Threats by attacking countries

URL Reports

- ☐ Unknown categories by URLs
- ☐ Malware categories by URLs
- ☐ Dataplane Cache URLs

Unknown Application Reports

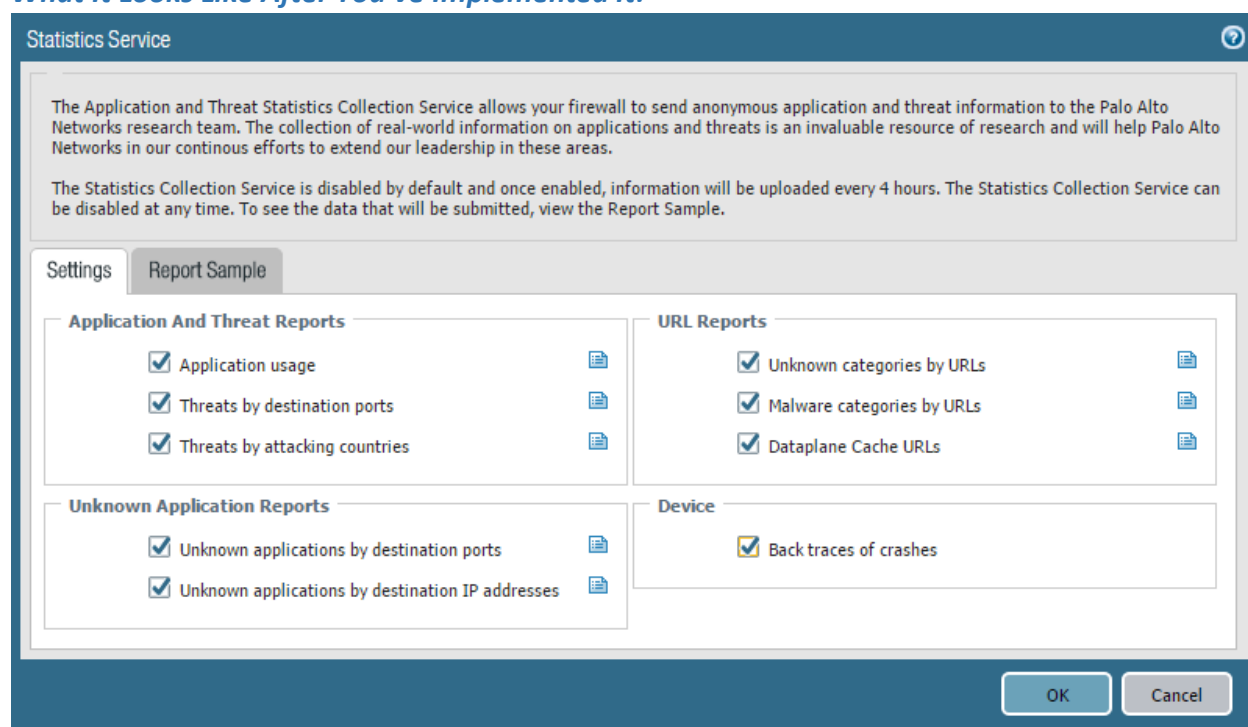
- ☐ Unknown applications by destination ports
- ☐ Unknown applications by destination IP addresses

Device

- ☐ Back traces of crashes

OK Cancel

You're ready to step up and help the community

What It Looks Like After You've Implemented It:*Continuous improvement just got faster***What Else You Need to Know:**

- This service is disabled by default.
- Once enabled, it will send information every four hours.

You can disable it at any time.

•

Data or Traffic Interfaces

□ Configure an Interface Management Profile for Each Interface

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Sometimes a firewall interface needs to do more than just pass packets (or not) to other interfaces. Sometimes it has to act as an IP host and participate as the server side in a client-server connection.

Why This Best Practice Is Important:

Not only do you want to specify precisely how the firewall will behave when a packet tries to *transit* an interface, you also want to specify precisely how the firewall behaves when a packet tries to *connect* to an interface.

How to Implement It:

Step 1:

Go to **Network > Network Profiles > Interface Mgmt** and create or modify an Interface Management Profile:

Interface Management Profile

Name: Allow-Only-Ping

Permitted Services

- ☒ Ping
- ☐ Telnet
- ☐ SSH
- ☐ HTTP
- ☐ HTTP OCSP
- ☐ HTTPS
- ☐ SNMP
- ☐ Response Pages
- ☐ User-ID
- ☐ User-ID Syslog Listener-SSL
- ☐ User-ID Syslog Listener-UDP

Permitted IP Addresses

+ Add - Delete

Ex. IPv4 192.168.1.1 or 192.168.1.0/24 or IPv6 2001:db8:123:1::1 or 2001:db8:123:1::/64

OK Cancel

This is a fairly tight Interface Management Profile that will accept only a ping connection

Step 2:

Go to **Network > Interfaces** and edit each interface you wish to attach this Interface Management Profile to.

In the dialog box, go to **Advanced > Other Info**:

The screenshot shows the 'Ethernet Interface' configuration window. The 'Interface Name' is 'ethernet1/1' and the 'Comment' is 'Outside NIC'. The 'Interface Type' is 'Layer3' and the 'Netflow Profile' is 'None'. The 'Advanced' tab is selected, and within it, the 'Other Info' sub-tab is active. The 'Management Profile' dropdown is set to 'Allow-Only-Ping' and is circled in red. Below it, the 'MTU' is set to '[576 - 1500]'. There are also fields for 'Adjust TCP MSS' with 'IPv4 MSS Adjustment' at 40 and 'IPv6 MSS Adjustment' at 60. At the bottom, there is an 'Untagged Subinterface' checkbox. The 'OK' and 'Cancel' buttons are at the bottom right.

Attaching an Interface Management Profile to an Interface

What Else You Need to Know:

Remember, an Interface Management Profile is for a “traffic” port, not the MGT port.

□ Enable IPv6 Support

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Why This Best Practice Is Important:

Enabling this feature will give you three benefits:

1. It will allow access to IPv6 hosts.
2. It will filter IPv6 packets that are encapsulated in IPv4 packets.
3. It will prevent IPv6 over IPv4 multicast addresses from being leveraged for network reconnaissance.

How to Implement It:

1. Go to **Network > Interfaces > Ethernet > IPv6** and edit an interface
2. Go to **IPv6** and check the box next to “Enable IPv6 on the interface”

The screenshot shows the 'Ethernet Interface' configuration window. The 'Interface Name' is 'ethernet1/1' and the 'Comment' is 'Outside NIC'. The 'Interface Type' is 'Layer3' and the 'Netflow Profile' is 'None'. The 'IPv6' tab is selected, and the checkbox 'Enable IPv6 on the interface' is checked and circled in red. Below this, there is a table for IPv6 addresses with columns for 'Address', 'Enabled', 'Interface ID as host', 'Anycast', and 'Send RA'. The 'Interface ID' is set to 'EUI-64 (default 64-bit Extended Unique Identifi)'. Below the table are buttons for '+ Add', '- Delete', 'Move Up', and 'Move Down'. The 'Address Resolution' section has 'DAD Attempts' set to 1, 'Reachable Time (sec)' set to 30, and 'NS Interval (sec)' set to 1. The 'Enable Duplication Address Detection' checkbox is unchecked. The 'Enable Router Advertisement' section has several settings: 'Min Interval (sec)' 200, 'Max Interval (sec)' 600, 'Hop Limit' 64, 'Link MTU' unspecified, 'Reachable Time (ms)' unspecified, 'Retrans Time (ms)' unspecified, 'Router Lifetime (sec)' 1800, and 'Router Preference' Medium. There are also checkboxes for 'Managed Configuration', 'Other Configuration', and 'Consistency Check', all of which are unchecked. At the bottom right are 'OK' and 'Cancel' buttons.

Enabling IPv6 support on the Interface

What Else You Need to Know:

If you'd like to learn more about IPv6, I recommend "**IPv6 Essentials: Integrating IPv6 into Your IPv4 Network, 3rd Edition**" by my friend Silvia Hagen of Zurich.

□ Use the Interface Comment Field

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

“There is nothing to writing. All you do is sit down at a typewriter and bleed.”

--Ernest Hemingway

Why This Best Practice Is Important:

A written note is a message to your future self, and your future self will thank you when, months from now, you read the informative descriptions you’ve left in the Comment fields.

How to Implement It:

Go to **Network > Interfaces**.

Edit each of your Interfaces and fill out the Comment field. A dozen words are plenty. Here are some hints:

- What purpose does this interface serve?
- What does this interface connect to?
- What’s the one special thing you have to remember about this interface?

The screenshot shows the 'Ethernet Interface' configuration window. The 'Interface Name' field contains 'ethernet1/1'. The 'Comment' field contains 'Outside NIC to untrusted Internet'. The 'Interface Type' dropdown is set to 'Layer3'. The 'Netflow Profile' dropdown is set to 'None'. Below these fields are tabs for 'Config', 'IPv4', 'IPv6', and 'Advanced', with 'Config' being the active tab. Under the 'Assign Interface To' section, the 'Virtual Router' dropdown is set to 'VR-01' and the 'Security Zone' dropdown is set to 'SZ_Untrust'. At the bottom right, there are 'OK' and 'Cancel' buttons.

It's not Shakespeare or Goethe, but future readers will appreciate it anyway

Security Zones

□ Apply a Strong Zone Protection Profile to Untrusted Security Zones

Improve Security	X	Improve Manageability	
Improve Performance	X	Improve High Availability	

Background Information:

There are many types of attacks at the lower levels of the network stack that don't target specific applications. These include:

- Flood attacks
- Reconnaissance attacks
- Other packet-based attacks

Why This Best Practice Is Important:

Zone Protection Profiles offer protection against these attacks and are easy to configure.

How to Implement It:

Step 1: Create a Zone Protection Profile:

Go to **Network > Network Profiles > Zone Protection** and create a new Zone Protection Profile:

Creating a new Zone Protection Profile

Step 2: Configuring your Security Zone to have a Zone Protection Profile:

After you've created your Zone Protection Profile, go to **Network > Zones** and edit each untrusted Security Zone.

Apply the correct Zone Protection Profile:

The screenshot shows the 'Zone' configuration window. The 'Name' field is 'SZ_Untrust' and the 'Type' is 'Layer3'. Under 'Interfaces', 'ethernet1/1' is listed. The 'Zone Protection Profile' dropdown is highlighted with a red circle and set to 'BJS-Zone-Protection-Profile'. The 'Log Setting' is 'None'. The 'Enable User Identification' checkbox is unchecked. The 'User Identification ACL' section on the right has empty 'Include List' and 'Exclude List' fields.

Configuring a Security Zone to have a Zone Protection Profile

What It Looks Like After You've Implemented It:

The screenshot shows the 'Zone Protection Profile' configuration window for the profile 'BJS-Zone-Protection-Profile'. The 'Flood Protection' tab is selected. Under 'Flood Protection', the 'SYN' checkbox is checked, and the 'Action' is set to 'Random Early Drop'. The 'Alert (packets/sec)' is 10000, 'Activate (packets/sec)' is 10000, and 'Maximum (packets/sec)' is 40000. Under 'Reconnaissance Protection', the 'ICMP' checkbox is checked, with 'Alert (packets/sec)' at 10000, 'Activate (packets/sec)' at 10000, and 'Maximum (packets/sec)' at 40000. The 'ICMPv6' checkbox is also checked with the same values. Under 'Packet Based Attack Protection', the 'Other IP' and 'UDP' checkboxes are checked, both with 'Alert (packets/sec)' at 10000, 'Activate (packets/sec)' at 10000, and 'Maximum (packets/sec)' at 40000.

Every protection is enabled on the Flood Protection tab

Zone Protection Profile

Name: BJS-Zone-Protection-Profile

Description:

Flood Protection Reconnaissance Protection Packet Based Attack Protection

Scan	Enable	Action	Interval (sec)	Threshold (events)
TCP Port Scan	<input checked="" type="checkbox"/>	alert	2	100
Host Sweep	<input checked="" type="checkbox"/>	alert	10	100
UDP Port Scan	<input checked="" type="checkbox"/>	alert	2	100

OK Cancel

Every protection is enabled on the Reconnaissance Protection tab

Zone Protection Profile

Name: BJS-Zone-Protection-Profile

Description:

Flood Protection Reconnaissance Protection Packet Based Attack Protection

IP Drop TCP Drop ICMP Drop IPv6 Drop ICMPv6 Drop

☒ Spoofed IP address

☒ Strict IP Address Check

☒ Fragmented traffic

IP Option Drop

<input checked="" type="checkbox"/> Strict Source Routing	<input checked="" type="checkbox"/> Security
<input checked="" type="checkbox"/> Loose Source Routing	<input checked="" type="checkbox"/> Stream ID
<input checked="" type="checkbox"/> Timestamp	<input checked="" type="checkbox"/> Unknown
<input checked="" type="checkbox"/> Record Route	<input checked="" type="checkbox"/> Malformed

OK Cancel

Every protection is enabled on the Packed Based Attack Protection tab

What Else You Need to Know:

The zone protections are enforced only on packets that don't match already-approved and already-existing sessions. Packets that match an existing session will bypass these protections.

There have been some known instances of some useful but probably poorly written applications that can get broken by some of these settings, so it's probably best to proceed slowly and enable these settings one at a time and look for problems. The settings that caused these problems included Spoofed IP Address and Fragmented Traffic.

□ Use the Special Security Zone Tag Trick to Assign a Color to a Security Zone

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

PAN-OS allows you to assign a color to a tag, and then assign a tag to many types of objects, and then that object takes on the color of the tag.

However, Security Zones don't let you assign a tag, so you normally can't color a Security Zone to mark it by risk level, so we need to get one level more clever to make it work.

Why This Best Practice Is Important:

When you're creating any sort of security perimeter, from medieval castle to Facebook sharing permissions to a modern firewall, it's always important to keep clear whether you're dealing with the inside or the outside. Color is a particularly good way to distinguish descriptors, so it's especially useful to use color to mark each Security Zone. There's a special trick in PAN-OS that allows you assign a color to each Security Zone.

How to Implement It:

If you create a colored tag *with the exact same name as a Security Zone*, then something magic happens: in your rulebases, every instance of the Security Zone will have that color. This allows you to assign colors to Security Zones.

Go to **Objects > Tags** and create a new tag:

The name matches that of a Security Zone

What It Looks Like After You've Implemented It:

	Name	Tags	Type	Source			
				Zone	Address	User	HIP Profile
5	Allow access from Host 1	RB_Host	universal	SZ_Untrust	GRP-LW-Monitoring GRP-LW-Offices	any	any

The Zone is now displayed in the same color as the tag with the same name

Zone Protection Profiles

These profiles are not enabled by default for the default untrust zone. Remember to try this in a test environment before you implement any of these in a production environment.

☐ Drop Malformed IP packets

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The standards for assembling IP packets are decades old. Everybody knows how to do it. There are no legitimate reasons for malformed IP packets to be transiting the Internet any longer, only illegitimate reasons.

Why This Best Practice Is Important:

Maybe, a long time ago, it was best to be flexible in what you'd accept when it came to IP packets, but there's no justification for that now, especially because malformed packets are often specially crafted to evade a firewall's protections and trick the Security Policy into coming to the wrong conclusions.

How to Implement It:

1. Go to **Network > Network Profiles > Zone Protection** and create or edit a Zone Protection Profile.
2. Within the Zone Protection Profile, go to **Packet Based Attack Protection > IP Drop > IP Option Drop**.
3. Check the box next to "Malformed".

Zone Protection Profile

Name: BJS-Zone-Protection-Profile

Description:

Flood Protection **Reconnaissance Protection** **Packet Based Attack Protection**

IP Drop **TCP Drop** **ICMP Drop** **IPv6 Drop** **ICMPv6 Drop**

☐ Spoofed IP address

☐ Strict IP Address Check

☐ Fragmented traffic

IP Option Drop

☐ Strict Source Routing

☐ Loose Source Routing

☐ Timestamp

☐ Record Route

☐ Security

☐ Stream ID

☒ Unknown

☒ Malformed

OK Cancel

Sorry, malformed IP packets: the 1980's called and they want you back

What Else You Need to Know:

Ensure you attach a properly configured Zone Protection Profile, at a minimum, to all your untrusted Security Zones.

❑ Remove TCP Timestamps on SYN packets

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

One way to confuse a less careful TCP host is to deliver multiple packets with different timestamps but the same sequence number. This is abnormal and unexpected.

However if the SYN packet (the first packet in a TCP connection) has the TCP timestamp option disabled, then the TCP stacks at both ends of the connection will not support TCP timestamps.

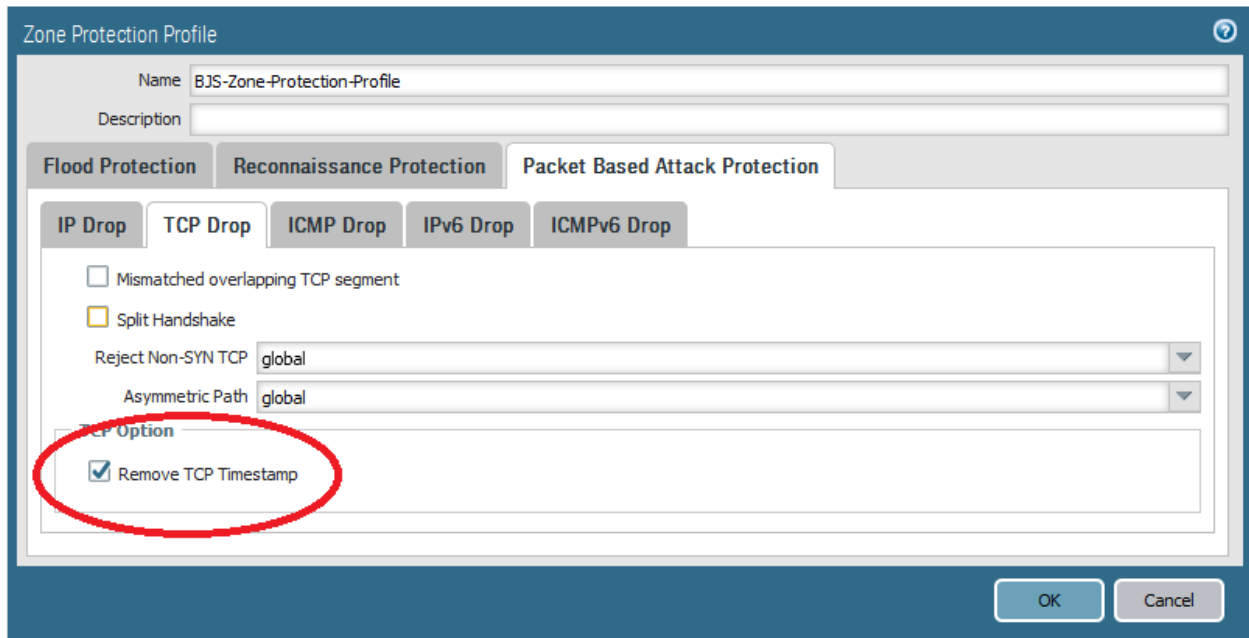
With this setting enabled, PAN-OS will disable the TCP timestamp option in all TCP SYN packets.

Why This Best Practice Is Important:

Disabling the TCP timestamp option helps protect less careful TCP hosts from specially crafted intentionally confusing TCP streams that might lead to an evasion.

How to Implement It:

1. Go to **Network > Network Profiles > Zone Protection** and create or edit a Zone Protection Profile.
2. Within the Zone Protection Profile, go to **Packet Based Attack Protection > TCP Drop > TCP Option**.
3. Check the box next to "Remove TCP Timestamp".



This protection is timeless

What Else You Need to Know:

Ensure you attach a properly configured Zone Protection Profile, at a minimum, to all your untrusted Security Zones.

❑ Drop Mismatched Overlapping TCP Segment

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Selecting this option causes PAN-OS to discard TCP frames with mismatched and overlapping data. The scenarios where the received segment will be discarded are:

- The segment received is contained within another segment
- The segment received overlaps with part of another segment
- The segment completely contains another segment

Why This Best Practice Is Important:

By deliberately constructing TCP connections containing overlapping but different data, an attacker can attempt to cause misinterpretation of the intent of the connection. This can be used to deliberately induce false positives or false negatives in a Security Policy. An attacker can use IP spoofing and sequence number prediction to intercept a user's connection and inject their own data into the connection.

How to Implement It:

1. Go to **Network > Network Profiles > Zone Protection** and create or edit a Zone Protection Profile.
2. Within the Zone Protection Profile, go to **Packet Based Attack Protection > TCP Drop > TCP Option**.
3. Check the box next to "Mismatched overlapping TCP segment".

Zone Protection Profile

Name: BJS-Zone-Protection-Profile

Description:

Flood Protection **Reconnaissance Protection** **Packet Based Attack Protection**

IP Drop **TCP Drop** **ICMP Drop** **IPv6 Drop** **ICMPv6 Drop**

☒ Mismatched overlapping TCP segment

☐ Split Handshake

Reject Non-SYN TCP: global

Asymmetric Path: global

TCP Option

☒ Remove TCP Timestamp

OK Cancel

Like mismatched socks, there are some things we just won't allow

What Else You Need to Know:

Ensure you attach a properly configured Zone Protection Profile, at a minimum, to all your untrusted Security Zones.

❑ Drop Packets With a Spoofed Source IP Addresses

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Because packets are routed across networks and the Internet based upon their destination IP address, the originator of a packet could fake, or “spoof”, the source IP address, and it would still get forwarded as appropriate. This gives an attacker some room to maneuver. Palo Alto NGFW will use the configured routing table to determine spoof address matches. If a match is made for an IP address on a route coming from the inside zone, but it is received on the outside zone it will be dropped.

Why This Best Practice Is Important:

When evaluating a proposed new connection, the firewall takes the first packet, and then starting at the top of the rulebase, starts comparing the packet against the rules, taken in order from the top down, to see which is the first rule that “matches” the packet.

The “match” is based upon comparing the packet to the values in the various elements of the rule. One of these elements is the Source Address.

The problem occurs in the firewall, when one of the key pieces of information used to determine which rule matches a packet is something that can be faked by the packet’s originator. In short, an attacker can spoof a source IP address to trick the firewall into matching the wrong rule.

A packet with a spoofed source IP address is very dangerous. This will only occur in one of two circumstances:

- You’ve got some sort of routing problem, like a loop, in which case you need to fix it.
- You’re being fed specially crafted packets with spoofed source IP addresses designed to fool your firewall.

Because it might be the second case, you need to always drop packets with spoofed source IP addresses.

How to Implement It:

1. Go to **Network > Network Profiles > Zone Protection**.
2. Open or create a Zone Protection Profile.

Zone Protection Profile

Name: BJS-Zone-Protection-Profile

Description:

Flood Protection Reconnaissance Protection Packet Based Attack Protection

IP Drop TCP Drop ICMP Drop IPv6 Drop ICMPv6 Drop

☒ Spoofed IP address

☐ Strict IP Address Check

☐ Fragmented traffic

IP Option Drop

☐ Strict Source Routing ☐ Security

☐ Loose Source Routing ☐ Stream ID

☐ Timestamp ☐ Unknown

☐ Record Route ☐ Malformed

OK Cancel

At a minimum, enable this protection

❑ Experiment With Enabling Some or All of the Other Protections in a Zone Protection Profile

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

There are a lot of attacks that aim for your Interfaces but come in below the level of an application, and the Zone Protection Profile has many types of protections against them, including:

- Flood Protection
- Reconnaissance Protection
- Packet Based Attack Protection

None of these protections are enabled by default. In addition to the protections recommended in other Best Practices, you may wish to consider enabling some or all of these other protections.

Why This Best Practice Is Important:

All of these protections block harmful attacks.

How to Implement It:

There isn't yet a clear answer to the question of what's best to enable, because of several complexities.

On the one hand...

- All other things being equal, more protection is better than less, which is an argument for preemptively enabling *everything*.
- If you're being hit with one of these attacks, other than noticing a degradation of throughput or high CPU usage on the firewall, it might not be easy to figure out exactly what's going on, so waiting to enable a protection until you need it might not work so well.

But on the other hand...

- There might be performance degradation from enabling some of these protections, which is an argument for being cautious when enabling them.
- Not every Internet application follows the rules exactly and some might get broken by dropping all fragmented traffic, for example.

Therefore, you'll have to do some experimenting to see if enabling these protections causes an unacceptable resource burden or blocks some otherwise legitimate traffic.

Here's some reasonable, middle-of-the-road advice:

- Keep notes
- Do lots of testing

- Enable a few protections at a time
- Watch for performance issues

Instructions:

1. Go to **Network > Network Profiles > Zone Protection**.
2. Open or create a Zone Protection Profile

One of the tabs in a Zone Protection Profile, with everything enabled

Virtual Routers

❑ Enable Support for Multicast Firewalling

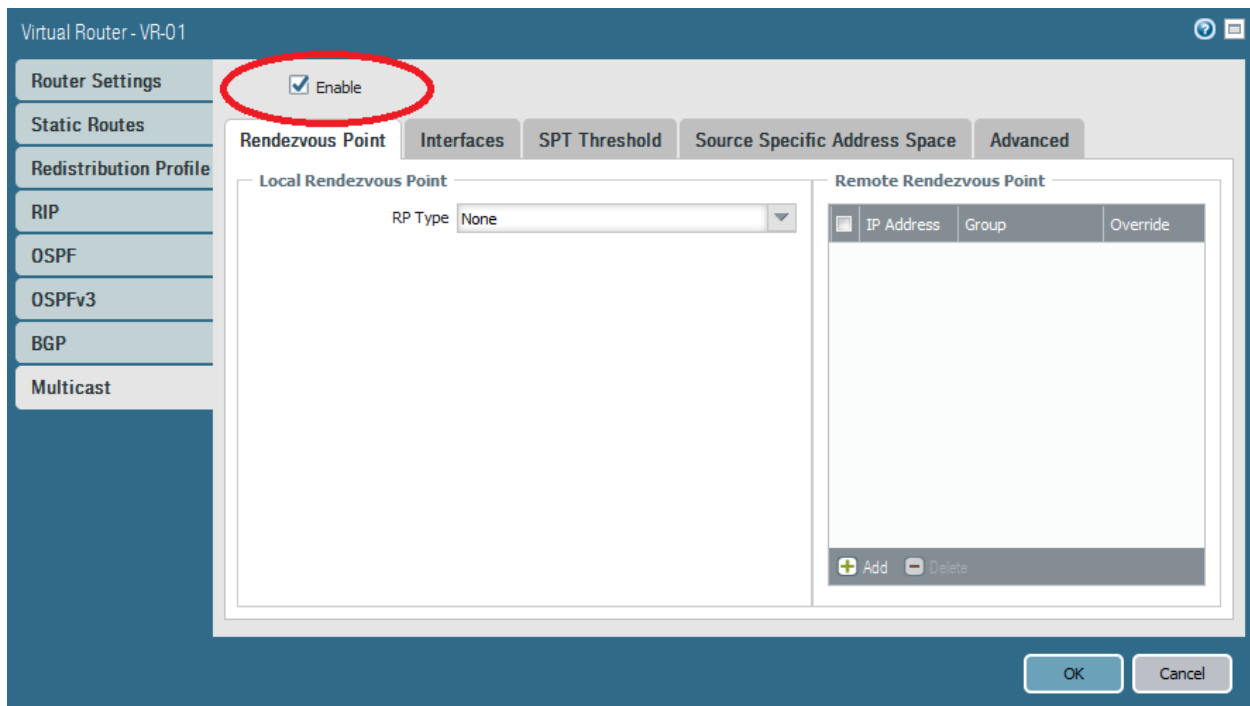
Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Why This Best Practice Is Important:

Enabling support for multicast traffic allows the firewall to enforce policy on multicast traffic.

How to Implement It:

1. Go to **Network > Virtual Routers** and edit or create a Virtual Router.
2. Got to **Multicast** and check the box next to “Enable”.



Enabling multicast routing in a Virtual Router

Rulebases (General)

❑ Give Every Rule a Meaningful Name

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Every rule has a *Name*, and the *Name* is the first element in the rulebase. The *Name* is your chance to create a prominent human-readable label for each rule.

Why This Best Practice Is Important:

As clearly as possible, the *Name* should immediately tell you exactly what the rule does, so when you're scanning the rulebase you don't have to examine the other elements in each rule to try to understand what's going on. It's a message to your future self, and your future self will thank you.

How to Implement It:

Edit the *Name* field in the *Security Policy Rule General* tab. It's best to always start the name with an action verb so you'll know exactly what it does. Good action verbs include:

- Deny
- Allow
- Drop
- Reset
- Block
- Alert
- Log

The screenshot shows the 'Security Policy Rule' configuration window with the 'General' tab selected. The 'Name' field contains 'ping'. The 'Rule Type' is set to 'universal (default)'. The 'Description' field is empty. The 'Tags' field contains 'SRB_Initial_Drop' with a red 'X' icon next to it. At the bottom right, there are 'OK' and 'Cancel' buttons.

This is a terrible Name for a rule. We don't know what this rule does.

What It Looks Like After You've Implemented It:

The screenshot shows the 'Security Policy Rule' configuration window. The 'General' tab is selected. The 'Name' field contains 'Allow ping to Default Gateway'. The 'Rule Type' is set to 'universal (default)'. The 'Description' field is empty. The 'Tags' field contains 'SRB_Initial_Drop' with a red 'X' icon next to it. The 'OK' and 'Cancel' buttons are at the bottom right.

Now this is an informative Rule Name.

What Else You Need to Know:

Changing the *Name* in a rule won't change any of the firewall's behavior, so a *Commit* is not directly required. However, the change is now sitting in the volatile *candidate config* so you'll need to either do a *Commit*, which will save it to the *running configuration*, or explicitly save it as the *candidate configuration*. If you make a name, change, make sure to double-check that it does not impact the detection of hits on a rule or your ACC reporting.

□ Create a Meaningful Description for Each Rule

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The name field in a rule is limited to 31 characters, which often isn't enough to describe what was intended with a rule or what some of the background information or thought processes were leading up to the rule's creation.

Why This Best Practice is Important:

You want future security policy administrators—including yourself—to be able to quickly understand why a rule exists and what it's trying to accomplish.

How to Implement It:

The screenshot shows the 'Security Policy Rule' configuration window with the 'General' tab selected. The 'Name' field is 'Drop these source IPs always', 'Rule Type' is 'universal (default)', and 'Tags' include 'SRB_Initial_Drop'. The 'Description' field is empty and circled in red.

The Security Policy Rule General tab Description field. It looks pretty empty here.

What It Looks Like After You've Implemented It:

The screenshot shows the same 'Security Policy Rule' configuration window, but now the 'Description' field contains the text: 'This rule drops all attempted connections from a list of IP addresses known to be associated with malware.'

Now it's much more useful.

❑ Add Your Workflow Ticketing System ID to the Rule's Description

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Most I.T. departments of any size use a support ticketing system that assigns a unique serial number to each support request. If it's managed properly, it's an easily searchable permanent database of all requested changes.

Why This Best Practice Is Important:

If you put the ticketing system ticket number in the Description field then you have a logical pointer back to the request that motivated the creation or changing of that rule. It's hard enough when you have a large number of rules to remember the details of each; with a ticket number you can quickly look up what's going on.

Also, it's useful when something goes wrong and you're sitting around blamestorming.

What It Looks Like After You've Implemented It:

Security Policy Rule

General Source User Destination Application Service/URL Category Actions

Name: Drop these Source IPs always

Rule Type: universal (default)

Description: Drop all attempted connections from a list of IP addresses known to be associated with malware. See Service Now ticket #SREQ020014.

Tags: RB_Initial_Drops X

OK Cancel

Now we know where to look for more information.

□ Use Rule Tags to Organize Rules into Groups

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

PAN-OS rulebases don't have a native rule group grouping functionality, so the rules are organized and displayed as a single one-dimensional ordered list.

Why This Best Practice is Important:

When you have a lot of rules, it's nice to have a way to "group" them in a way that's more useful than a simple list. By assigning tags to rule, you can get most of the benefits of rule grouping.

How to Implement It:

Tags are your friend because each has both a name and a color, and these are useful in organizing things into groups. One of the elements of each rule is tags, and each rule can have up to 64 Tags. By creating tags for each group of rules, and then assigning these tags to each rule, you create visual elements that help organize rules into groups.



For these tags it's best to use a naming convention. A good suggestion is to start with the prefix "RB_", which stands for rulebase. Tags with this prefix are used exclusively for rules, and not other objects.

It's also best to use a color convention keyed to the Action setting of the rule:

Action Setting: Tag Color:

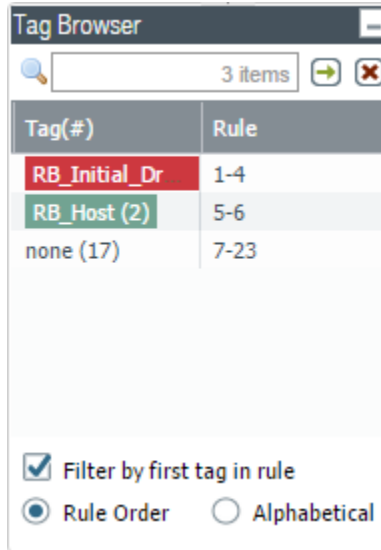
Deny, Drop, Reset client, Reset server, and Reset both client and server	Red
Allow	Green

What It Looks Like After You've Implemented It:

	Name	Tags	Type	Zone
1	Drop these source IPs always	RB_Initial_Drop	universal	any
2	Drop these desti IPs always	RB_Initial_Drop	universal	any
3	Drop these applications always	RB_Initial_Drop	universal	any
4	Drop these services always	RB_Initial_Drop	universal	any
5	Allow access from Host 1	RB_Host	universal	 SZ_Untrust
6	Allow access from Host 2	RB_Host	universal	 SZ_Untrust

Red tags for drop rules, green tags for allow, and they all have the proper prefix.

Once you've configured rule tags, then you can make use of the Tag Browser, at the lower left corner of the rulebase, which lets you filter the rule view by tag:



The tag browser allows you to see rule groupings as defined by tags.

What Else You Need to Know:

Changing the tag in a rule won't change any of the firewall's behavior, so a *Commit* is not directly required. However, the change is now sitting in the volatile *candidate config* so you'll need to either do a *Commit*, which will save it to the *running configuration*, or explicitly save it as the *candidate configuration*.

□ Minimize the Number of Rules

Improve Security		Improve Manageability	X
Improve Performance	X	Improve High Availability	

Background Information:

The rules in a rulebase are evaluated in order, from the top down. Each rule, at a fundamental level, is a set of matching conditions and an associated action. When a rule is evaluated, the matching conditions are compared to the connection, and if *all* of them match, then the firewall executes the specified action and stops evaluating rules further down in the rulebase.

Each rule is also a specific instruction from the firewall administrator, and has to be considered in context with all the other rules.

Why This Best Practice Is Important:

Even though the software code that evaluates these matching conditions is optimized and extremely fast, in high traffic environments the cumulative CPU load of thousands or millions of rule evaluations can be significant.

Also, given that each rule has potential interactions with every other rule, the total number of interactions the administrator has to keep track of grows with the square of the number of rules. It doesn't take a very large rulebase before you can't keep it all in your head, and if gets even larger, it might take days or weeks to get a new administrator or auditor on board with understanding what's going on.

Therefore, for both performance improvement and for reducing complexity, it's best to minimize the number of rules in your rulebases.

How to Implement It:

The better you understand your network and your goals, the better you'll be able to configure a tight, appropriate rulebase.

Here are some specific strategies for reducing the number of rules:

- Delete unused rules.
- Delete disabled rules.
- Combine rules with similar elements.
- Delete rules that get shadowed by other rules.

What It Looks Like After You've Implemented It:

"Everything should be made as simple as possible, but not more so"

--Albert Einstein

What Else You Need to Know:

Obviously, there may be some tradeoffs on the margins between extracting every last bit of specificity and security control by maintaining multiple different rules to cover a complex situation and the performance improvement

and reduced complexity of combining them. You'll learn with experience how to weigh these competing goals.

□ Put More Frequently Matched Rules Higher in the Rulebase

Improve Security		Improve Manageability	
Improve Performance	X	Improve High Availability	

Background Information:

The rules in a rulebase are evaluated in order, from the top down. Each rule, at a fundamental level, is a set of matching conditions and an associated action. When a rule is evaluated, the matching conditions are compared to the connection, and if *all* of them match, then the firewall executes the specified action and stops evaluating rules further down in the rulebase.

Why This Best Practice Is Important:

Even though the software code that evaluates these matching conditions is optimized and extremely fast, in high traffic environments the cumulative CPU load of thousands or millions of rule evaluations can be significant.

By reordering the rules in a rulebase, the *average number of rules evaluated before matching* can be reduced, thus reducing CPU load and improving performance.

Also, don't forget about the potential for shadowing. You don't want your frequently hit generic rule to shadow out your specific app rule.


How to Implement It:

Step 1: The first step is to understand your traffic mix.

Go to **Monitor > Reports > Traffic Reports > Security Rules** to see which Security Rules are matching the most sessions.

Step 2: The second step is to change the order of the rules.

There are two ways to do this. The first is to use the *Move* action menu at the bottom of the rulebase.



The Move action menu lets you change the order of rules in your rulebase.

The second method is to drag-and-drop individual rules. To select a rule to be dragged-and-dropped, simply click and hold on its rule number.

	Name	Tags	Type	Source				
				Zone	Address	User	HIP Prof	Zone
1	Drop these source IPs always	SRB_Initial_Drop	universal	any	GRP-Internet-Scanners	any	any	any
2	Drop these desti IPs always	SRB_Initial_Drop	universal	any	any	any	any	any
3	Drop these applications always	SRB_Initial_Drop	universal	any	any	any	any	any

Click and drag on a rule number to move a rule to a different location in your rulebase.

What Else You Need to Know:

After changing the order of rules in a rulebase, you'll need to do a Commit for the changes to take effect.

❑ Configure "Temporary" Rules to Expire on a Schedule

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Being a firewall administrator means dealing with an endless stream of requests for access justified by often temporary circumstances.

Why This Best Practice Is Important:

Leaving an *Allow* rule in your rulebase after it's no longer needed is a security risk. It's best whenever configuring a rule for a "temporary" project or circumstance to explicitly configure an end date in the rulebase so the rule will self-disable at the appropriate time. A later audit will catch it and allow you to delete it.

How to Implement It:

The screenshot shows the 'Security Policy Rule' configuration window with the 'Actions' tab selected. The 'Action Setting' section shows 'Action' set to 'Drop' and 'Send ICMP Unreachable' unchecked. The 'Log Setting' section shows 'Log at Session Start' and 'Log at Session End' unchecked, and 'Log Forwarding' set to 'None'. The 'Other Settings' section is circled in red, showing 'Schedule' set to 'None', 'QoS Marking' set to 'None', and 'Disable Server Response Inspection' unchecked. The 'Profile Setting' section shows 'Profile Type' set to 'None'. The 'OK' and 'Cancel' buttons are at the bottom right.

Configure a Schedule for a rule

Start Date	Start Time	End Date	End Time
2017/07/17	00:00	2017/08/18	23:45

Here's an appropriate schedule for a month-long audit.

Go to **Objects > Schedules** to view and manage your *Schedule* objects.

What It Looks Like After You've Implemented It:

This security rule is valid only during this non-recurring scheduled event.

What Else You Need to Know:

Ensure your firewall is properly configured for the correct time zone and with NTP to ensure your scheduled events occur when you're expecting them to. The ability to not both enforce daily time and global start and end dates is a significant shortcoming which can negatively expose an environment, when one assumes the schedule is doing an activity it simply is not enforcing.

□ Check for Unused Rules Regularly

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:


A security policy is the living expression of a security administrator's goals, and in most organizations these goals undergo constant change. There are adds and moves; and changes in personnel, equipment and projects. While an add or change that requires new access may easily motivate a new rule, deletions or changes that make a rule obsolete can easily slip by unnoticed. The result can be unneeded rules that never get matched.

Why This Best Practice Is Important:

Of particular concern is an *Allow* rule that never matches. This is an unneeded hole in the firewall that might someday provide access for malware. This violates the *Rule of Least Privilege* and should be closed.

How to Implement It:

At the bottom of every rulebase is the *Highlight Unused Rules* checkbox. Check it and unused rules will be shaded for easy recognition. With the unused rules now visible, you can decide if a rule is truly unneeded and should be disabled or deleted.



The Highlight Unused Rules checkbox

The time scope for this feature is the period since the firewall was last restarted, not since the last Commit.

What Else You Need to Know:

Some non-matching rules can remain in your rulebase for good reason. For example, if you need to provide access to auditors every quarter, or every year, or access for some other rare need, you might have a rule that remains unmatched for extended periods but still shouldn't be deleted. If it's long enough between uses of this rule, it's probably best to leave it disabled until needed, or configure a *Schedule* object in the *Action* element of the rule.

Security Policy Rule

General Source User Destination Application Service/URL Category Actions

Action Setting

Action: Drop
☐ Send ICMP Unreachable

Log Setting

☐ Log at Session Start
☐ Log at Session End
Log Forwarding: None

Profile Setting

Profile Type: None

Other Settings

Schedule: None
QoS Marking: None
☐ Disable Server Response Inspection

OK Cancel

How to configure a Schedule in a Security Policy Rule

□ Audit Your Rulebases Regularly

Improve Security	X	Improve Manageability	X
Improve Performance	X	Improve High Availability	

Background Information:

There are two naturally-occurring processes that contribute to your rulebase acquiring a dense underbrush of unneeded rules over time.

The first is that it's in the nature of organizations that errors of commission are punished more severely than errors of omission, so no-longer-needed rules tend to linger in the rulebase. It's far better to be the firewall admin who was too busy to delete unneeded rules than to be the firewall admin who accidentally knocked the Engineering team off the Internet.

The second is that once a rulebase becomes larger than one or two pages of scrolling, and the interactions between rules increase with the square of the number of rules, the amount of concentration and research to truly understand what's going on quickly becomes greater than that available in the typical firewall administrator's interrupt-driven existence. This means it's going to be especially hard for someone to feel confident enough in their analysis to want to delete a rule.

The result is that rulebases tend to grow by neglect until they're simply unwieldy.

Why This Best Practice Is Important:

You have to tame your ever-growing rulebases or else they'll just get out of hand, harming security, performance, and manageability. At some number of rules—500?, 1,000?, 5,000?—you just lose control of it.

How to Implement It:

Every rule should be documented somewhere else other than just in the rulebase. Then, at an interval appropriate for your organization—90 days, a year?—you can verify that the rule is still needed. You'll want to apply an especially tough necessity filter to *Allow* rules. Highlighting unused rules is also your friend, as you'll be able to see rules that aren't getting matched.

Lastly, rules that have a *Schedule* and a *Non-recurring Recurrence* and have already passed their *End Date* are prime candidates for deletion.

What It Looks Like After You've Implemented It:

"Everything should be made as simple as possible, but not more so"

--Albert Einstein

A nicely manicured rulebase is a thing of beauty; it does what you want, but not more, and you understand it.

What Else You Need to Know:

If you're convinced a rule is ready to be deleted, consider disabling it first to see if it breaks something you weren't aware of.

❑ When Removing a Rule, Disable It First To See If It Breaks Anything

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

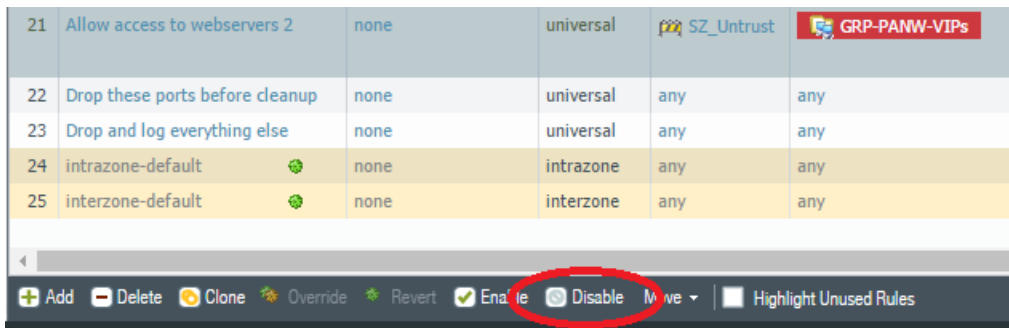
Why This Best Practice Is Important:

Every firewall administrator knows that the fastest way to make your phone ring is to delete a rule on the firewall that you're sure no one is using. Also, once you've deleted a rule, it's very difficult to remember the elements in the rule and the rule's position in the rulebase.

For these reasons it's far better to *disable* the rule first to see if it breaks anything. Be sure to leave helpful comments in the rule's Description, such as the ticket system ID number, or your name and contact information. Then you can let it sit for 30 or 60 days and see if anyone complains.

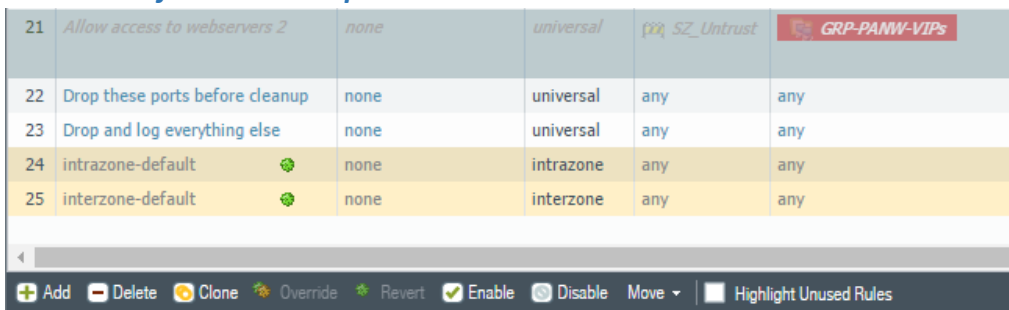
How to Implement It:

At the bottom of every rulebase is the *Disable* action button, which is available when you've selected a rule.



How to disable a selected rule. Rule 21 is selected and now can be disabled.

What It Looks Like After You've Implemented It:



It's a bit hard to see, but Rule 21 is now disabled and its text is now grayed out.

What Else You Need to Know:

After you disable a rule, you need to do a *Commit* for the change to take effect.

Security Rulebase

□ Use Continuous Improvement to Improve Your Security Rulebase

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

A good security policy and firewall configuration is a living document in that you're probably making regular adds, moves, and changes while figuring out better ways to do things.

More importantly, a good security policy is the result of a *process* more than an *event*.

Why This Best Practice Is Important:

Continuous improvement works, and Palo Alto Networks products are complex, sophisticated, and improved upon with every software release, so you're going to have to work hard to stay at the peak of your skills.

Also, it can take a lot of experimentation to truly understand what's going on in your network and to strike the right balance between improving security, improving manageability, and improving performance.

How to Implement It:

Roll up your sleeves and jump in. Do your homework, get your facts straight, and do experiments. You're never going to have the "perfect" security policy, but you can almost always create a *better* security policy.

If you're doing it right, you'll be spending a lot of time on the Monitor tab, seeing what your firewall is telling you, creating IP Address and Address Group objects to better label the hosts you're dealing with, and making more finely delineated Security Rules to permit precisely what you want, but not more.

❑ Create a Cleanup Rule at the Bottom of Your Security Rulebase

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

By default, the process of creating your Security rulebase should start with dropping all traffic and then making *a few, well-considered, narrowly defined exceptions*, which are your Allow rules.

Why This Best Practice Is Important:

The first part, *dropping all traffic*, is especially important. While your organization may want every employee to default to yes on every request, your firewall should definitely default to no.

How to Implement It:

Create a rule at the bottom of your Security Rulebase with the following settings:

Field:	Setting:
Name:	Drop and log everything else
Type:	universal
Zone:	any
Address:	any
User:	any
HIP Profile:	any
Zone:	any
Address:	any
Application:	any
Service:	any
Action:	Drop
Profile:	none
Options:	Log at Session End

Notice how this rule matches and drops *everything*. Therefore, if a packet doesn't match of your *few, well-considered, narrowly defined* Allow rules, it's going to get dropped. Creating this rule is the first step in creating a good security policy.

□ Follow the Principle of Least Privilege

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

This is the Principle of Least Privilege:

Every program and every privileged user of the system should operate using the least amount of privilege necessary to complete the job.

--Jerome Saltzer, Communications of the ACM

Your security policy should always default to drop, and then open up the smallest possible window that permits transit of the traffic that you want.

Why This Best Practice Is Important:

With enough time and computing power, attackers can try literally every possible opening in your network. If you grant a more expansive privilege than is required, you're in effect leaving a door open.

How to Implement It:

Do your homework when an Allow rule is requested. Ask awkward, probing questions to figure out what's the most restrictive possible rule that you can create that will still allow the desired transit.

You're not there to be nice; you're there to create the tightest possible rule that still permits the task at hand. Be like Bartleby the Scrivener in Herman Melville's short story and just keep saying, "I would prefer not to..."

□ Create an IP Source Blacklist Rule at the Top of Your Security Rulebase

Improve Security	X	Improve Manageability	
Improve Performance	X	Improve High Availability	

Background Information:

There are some IP addresses you just already know you don't want to get any connections from. These may include:

- Addresses and subnets on any of many public blacklists
- Countries where you're not legally allowed to do business (죄송합니다, 김정은)
- Addresses on the list of Bogons or Superbogons
- Your competitors

Why This Best Practice Is Important:

For both security and performance reasons, it's best to just drop these proposed new connections right at the top of your firewall policy and not consume any more resources processing them.

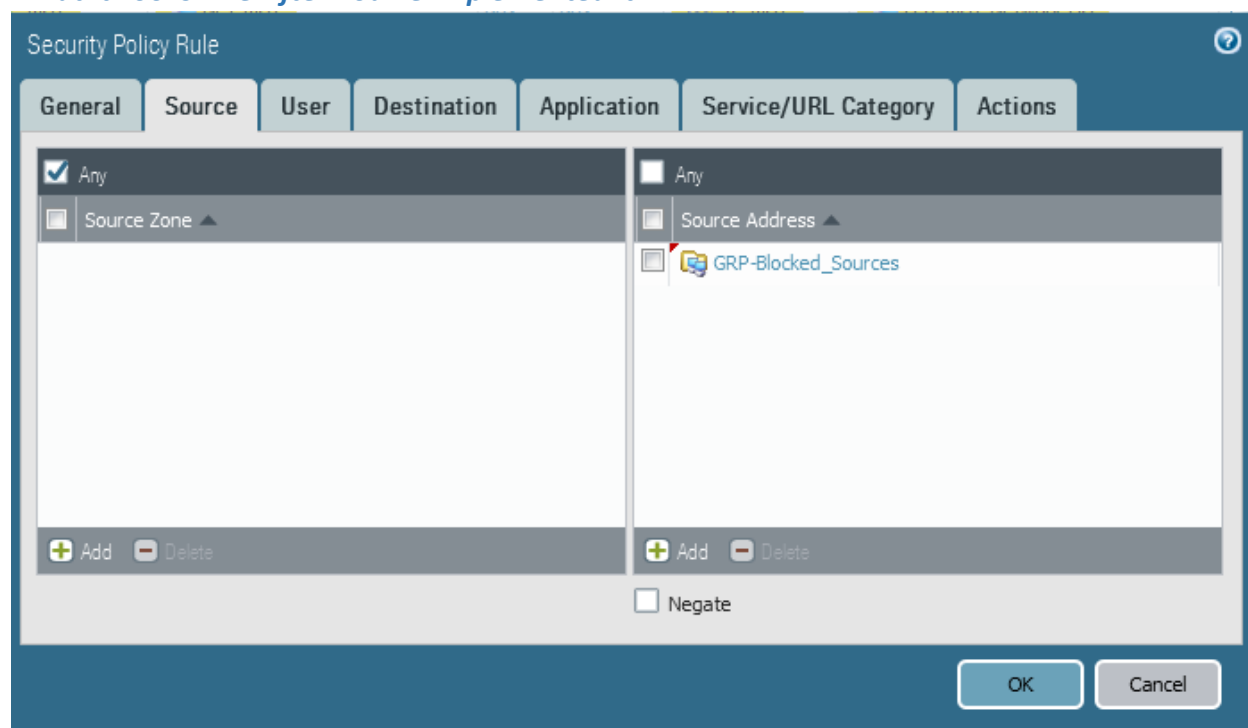
How to Implement It:

Step 1:

1. Go to **Objects > Address Groups**.
2. Create a new Address Group called "GRP-Blocked_Sources".
3. Put everything in there you never want to hear from again.

Step 2:

Create a rule at the top of your Security Policy rulebase that drops everything when the Source Address is a member of GRP-Blocked_Sources.

What It Looks Like After You've Implemented It:

A drop rule that blocks everything from this Address Group

What Else You Need to Know:

You can decide whether you wish to log these attempted connections or not.

□ Create an IP Destination Blacklist Rule at the Top of Your Security Rulebase

Improve Security	X	Improve Manageability	
Improve Performance	X	Improve High Availability	

Background Information:

There are some IP addresses you just already know you don't want permit any connections to. These may include:

- Addresses and subnets on any of many public blacklists
- Countries where you're not legally allowed to do business (ما نمی توانیم به اینجا بروید و در عین حال، هم)
- Addresses on the list of Bogons or Superbogons

Why This Best Practice Is Important:

For both security and performance reasons, it's best to just drop these proposed new connections right at the top of your firewall policy and not consume any more resources processing them.

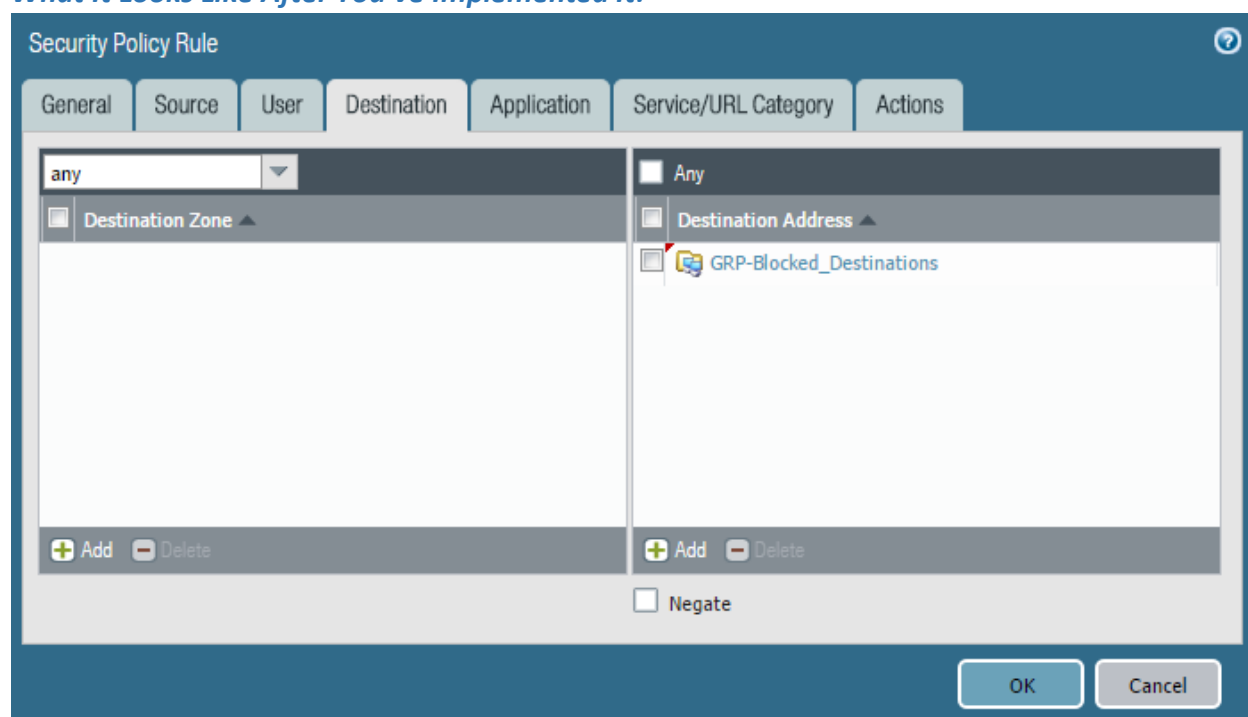
How to Implement It:

Step 1:

1. Go to **Objects > Address Groups**.
2. Create a new Address Group called "GRP-Blocked_Destinations".
3. Put everything in there you never want to connect to again.

Step 2:

Create a rule at the top of your Security Policy rulebase that drops everything when the Destination Address is a member of GRP-Blocked_Destinations.

What It Looks Like After You've Implemented It:

A drop rule that blocks everything to this Address Group

What Else You Need to Know:

You can decide whether you wish to log these attempted connections or not.

❑ Use Geographic IP Filters As Appropriate

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Palo Alto Networks subscribes to a commercial country-to-IP-address mapping service that provides updated lists of networks associated with each country. This allows the firewall administrator to place objects representing individual countries into fields that take IP addresses.

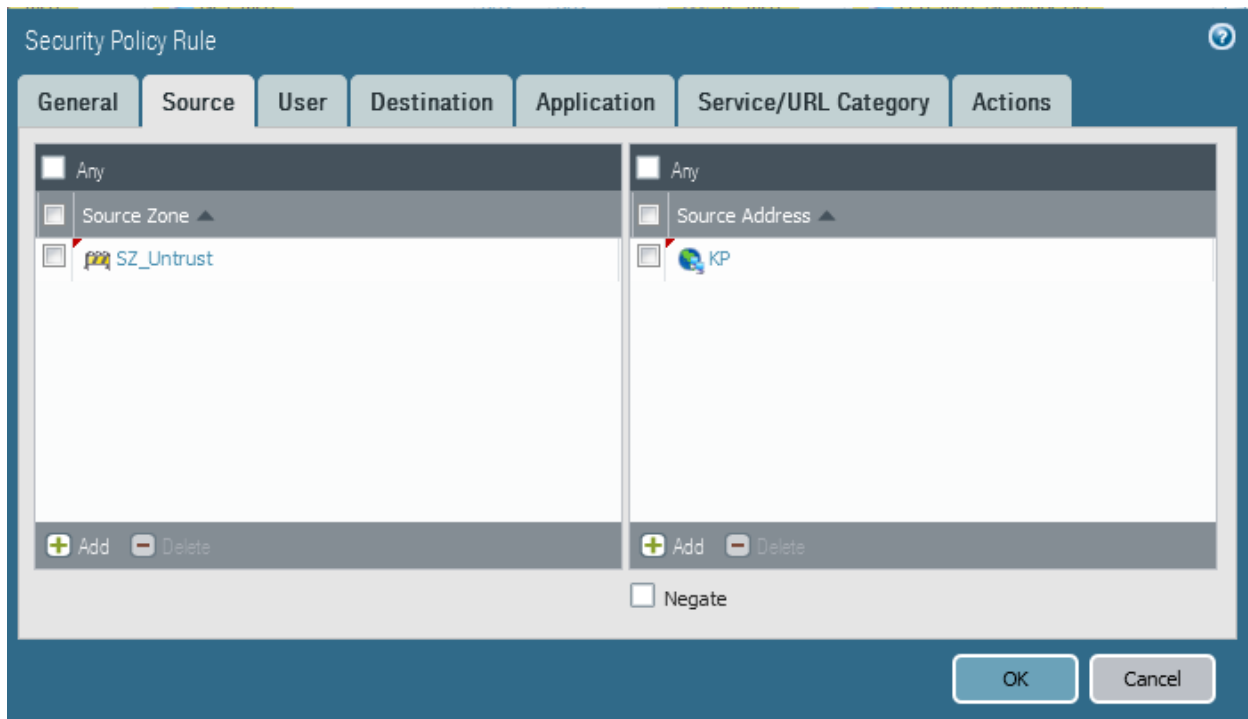
Why This Best Practice Is Important:

There are several reasons why you might want to block traffic to or from specific geographic regions, including:

- Legal requirements such as embargoes
- Knowledge about where an organization does or does not do business
- Known patterns of Internet attack origination

How to Implement It:

Go to **Policies > Security** and edit a security rule:



I just saw "The Interview", so here's a rule blocking connections from North Korea

□ Create an Application Blacklist at the Top of Your Security Rulebase

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

There are some applications that you just don't want transiting your firewall under any circumstances. These applications typically meet one or more of these criteria:

- They're insecure, as in sending passwords or data in the clear
- They've been replaced by better applications
- They aren't recognized by the App-ID engine, which, if they aren't custom in-house applications, may mean they're intentionally deceptive or evasive
- They're excessive chatty but don't provide any value
- They shouldn't be going through a router, which is what a firewall is in Layer 3 mode

Here is a list of the usual suspects:

- telnet (always replace this with SSH)
- rlogin (always replace this with SSH)
- ftp (always replace this with SCP/SFTP)
- rsync (Unix file syncing on TCP/873 and UDP/873)
- netbios-dg (NetBIOS Datagram distribution service on UDP138)
- netbios-ns (NetBIOS Name Service on TCP/137 and UDP/137)
- netbios-ss (NetBIOS Session mode on TCP/139)
- unknown-p2p
- unknown-tcp
- unknown-udp

Why This Best Practice Is Important:

It's best to block these applications right at the top of the firewall policy so they don't accidentally get permitted by a more permissive rule further down.

How to Implement It:

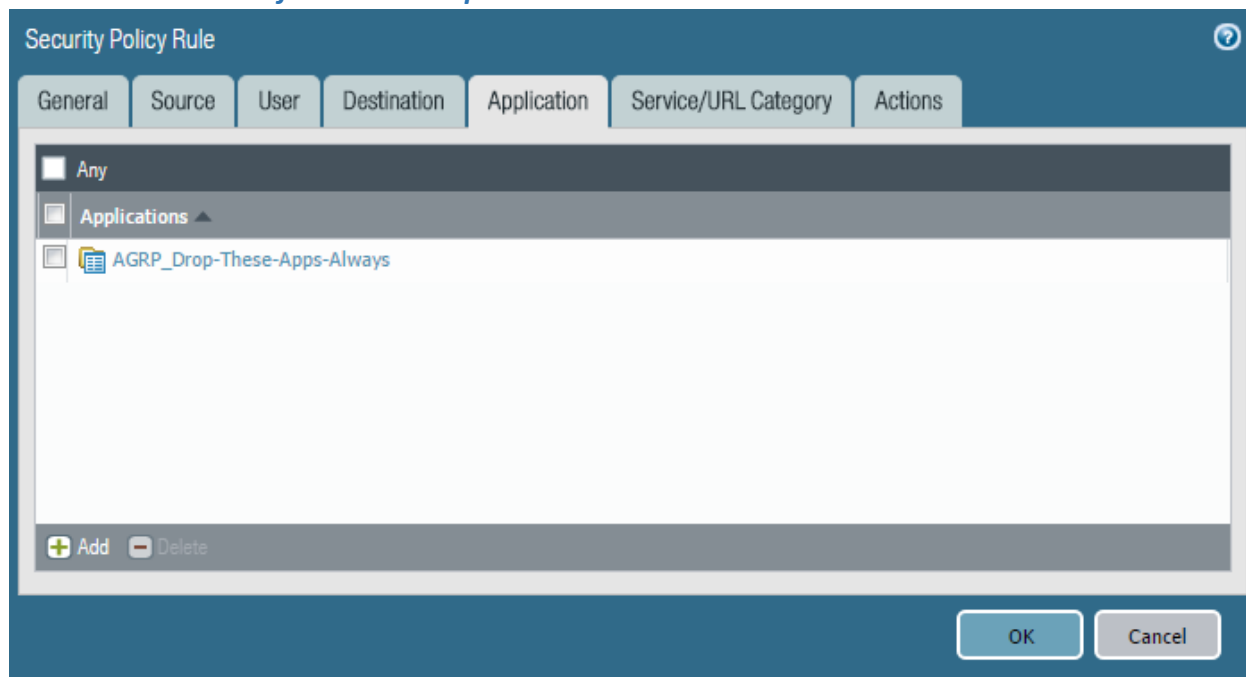
Step 1:

1. Go to **Objects > Application Groups**.
2. Create an Application Group called "GRP-Drop-These-Apps-Always."

Step 2:

1. Go to **Policies > Security**.
2. Add a rule at the top of your rulebase that always blocks these applications.

What It Looks Like After You've Implemented It:



A Security Policy rule that blocks these applications always

What Else You Need to Know:

You can decide whether or not you wish to log these blocked connections.

❑ Create an Inbound Service Blacklist at the Top of Your Security Rulebase

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

In the Palo Alto Networks world, the word “Service” roughly means “Port Number”, and while App-ID is a very powerful engine for identifying applications on the wire, regardless of which port an application is using, sometimes you still want to look at the port number.

Why This Best Practice Is Important:

Any public IPv4 sitting naked on the Internet is being scanned many times per minute, 24/7/365. Most of this is trying to connect to a few specific ports and is really just robots checking to see if you were careless enough to leave something useful naked and exposed.

To improve the signal-to-noise ratio in your log, and to just drop this noise before getting down to the real traffic analysis, it's best to just silently drop all these probes at the top of your rulebase. By dropping them based on their port number, the firewall doesn't even have to let enough packets flow to identify the application.

How to Implement It:

Step 1:

1. Go to **Objects > Services**.
2. Create new Service objects for these ports:

Name	Description	Protocol	Destination Port
TCP-23	Telnet	TCP	23
TCP-137	netbios-ns	TCP	137
TCP-139	netbios-ss	TCP	139
TCP-445	NetBIOS Over TCP	TCP	445
TCP-1433	SQL Server	TCP	1433
TCP-3306	MySQL	TCP	3306
TCP-3389	Windows RDP	TCP	3389
TCP-5060	SIP	TCP	5060
TCP-5061	SIP	TCP	5061
TCP-8080	Alternate Port for HTTP	TCP	8080
UDP-137	netbios-ns	UDP	137
UDP-138	netbios-dg	UDP	138

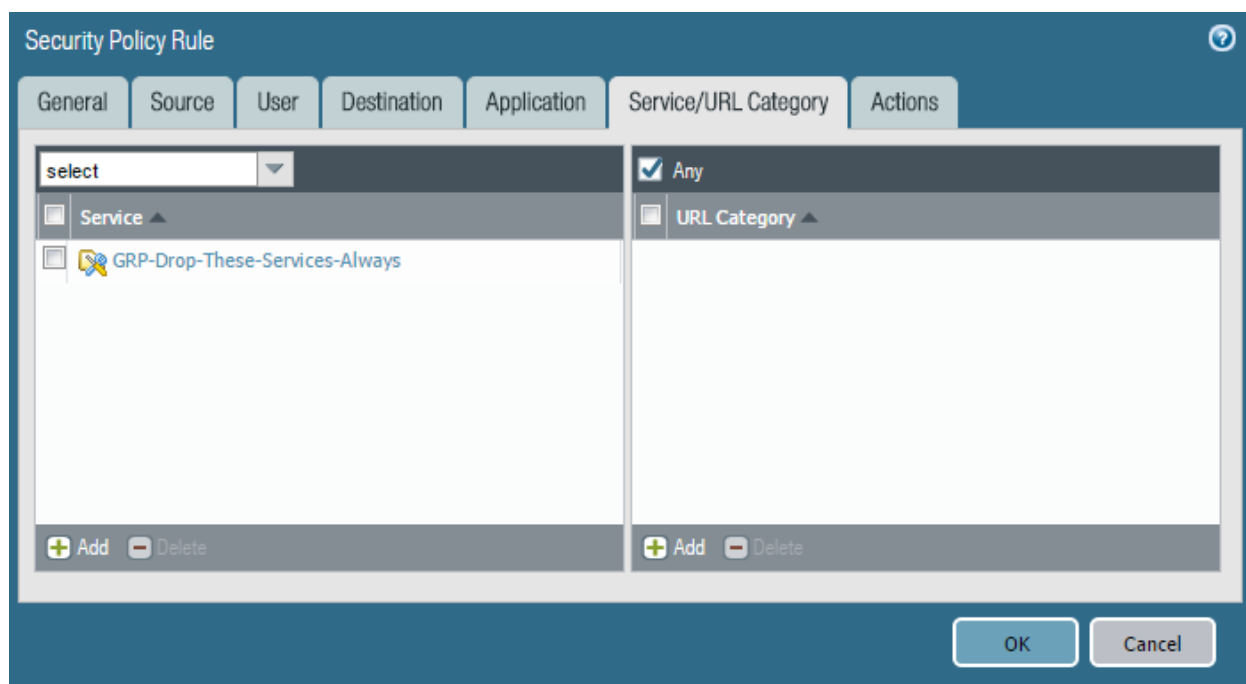
UDP-3306	MySQL	UDP	3306
UDP-3389	Windows RDP	UDP	3389
UDP-5060	SIP	UDP	5060
UDP-5061	SIP	UDP	5061
UDP-53413	Netis Systems Routers Vulnerability	UCP	53413

Step 2:

1. Go to **Objects > Service Groups**.
2. Create a new Service Group called "GRP-Drop-These-Services-Always".
3. Add all these new Services into this group.

Step 3:

Go to **Policies > Security** and add a new rule at the top that drops all connections inbound from the Internet using these Services:



Every proposed new connection from the Internet matching one of these Services will be dropped

What Else You Need to Know:

You can decide whether or not you wish to log these blocked connections.

❑ Avoid Using the Any Source Address in Allow Rules

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

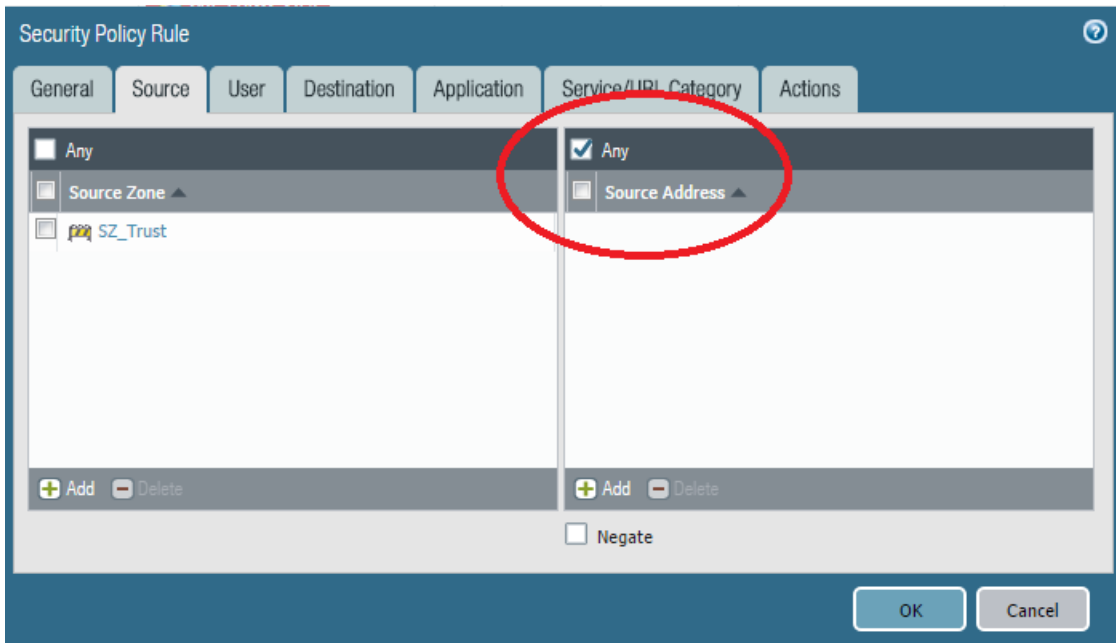
PAN-OS uses a multipart test to determine whether a proposed new connection matches a specific rule in the Security rulebase. One of those tests is to examine the source address.

Why This Best Practice Is Important:

By default, the process of creating your Security rulebase should start with dropping all traffic with a Cleanup Rule at the bottom and then making *a few, well-considered, narrowly defined exceptions*, which are your Allow rules.

“Narrowly defined” means you should avoid using the *any* source address in Allow rules.

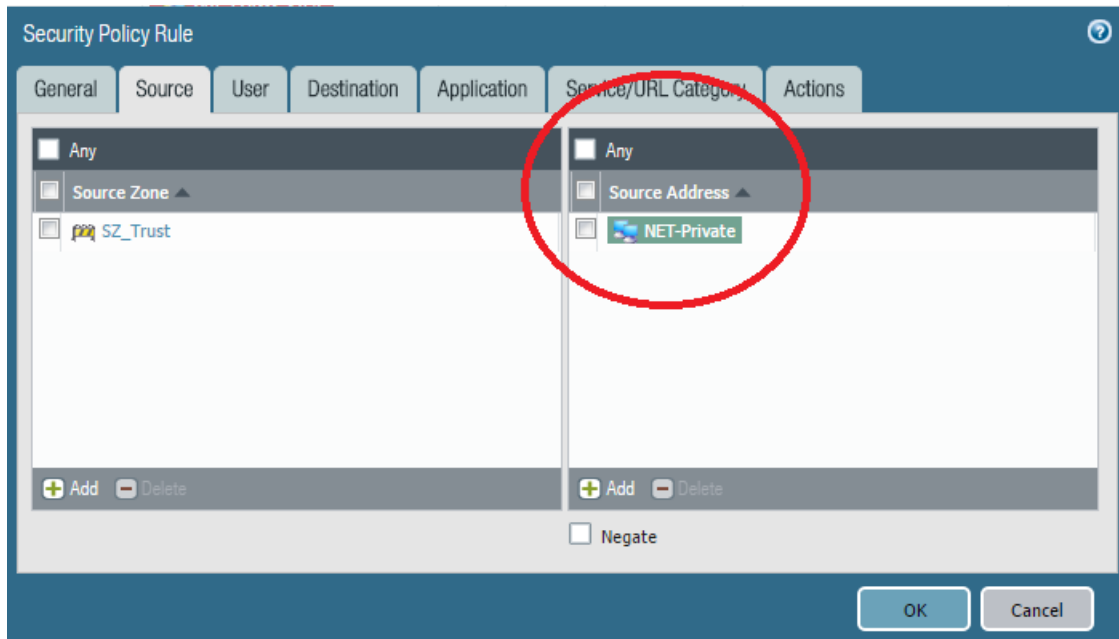
This is what it should **not** look like:



This Allow rule is matching all source addresses

How to Implement It:

This is what it **should** look like:



This Allow rule now must match a specific address or set of addresses

❑ Avoid Using the Any Source User in Allow Rules

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

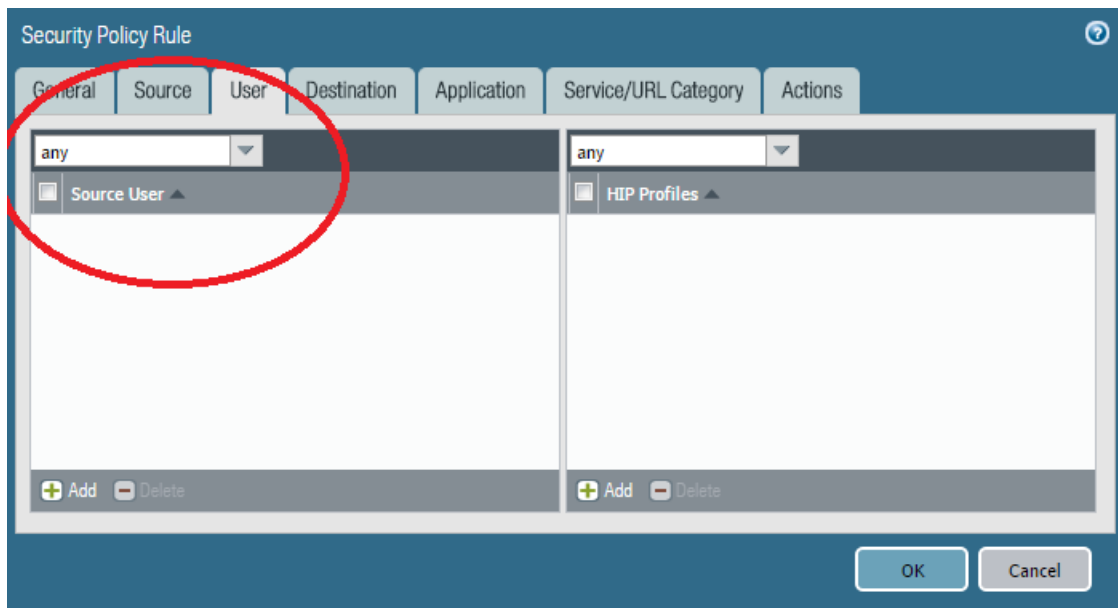
PAN-OS uses a multipart test to determine whether a proposed new connection matches a specific rule in the Security rulebase. One of those tests is to examine the source user.

Why This Best Practice Is Important:

By default, the process of creating your Security rulebase should start with dropping all traffic with a Cleanup Rule at the bottom and then making *a few, well-considered, narrowly defined exceptions*, which are your Allow rules.

“Narrowly defined” means you should avoid using the *any* source user in Allow rules.

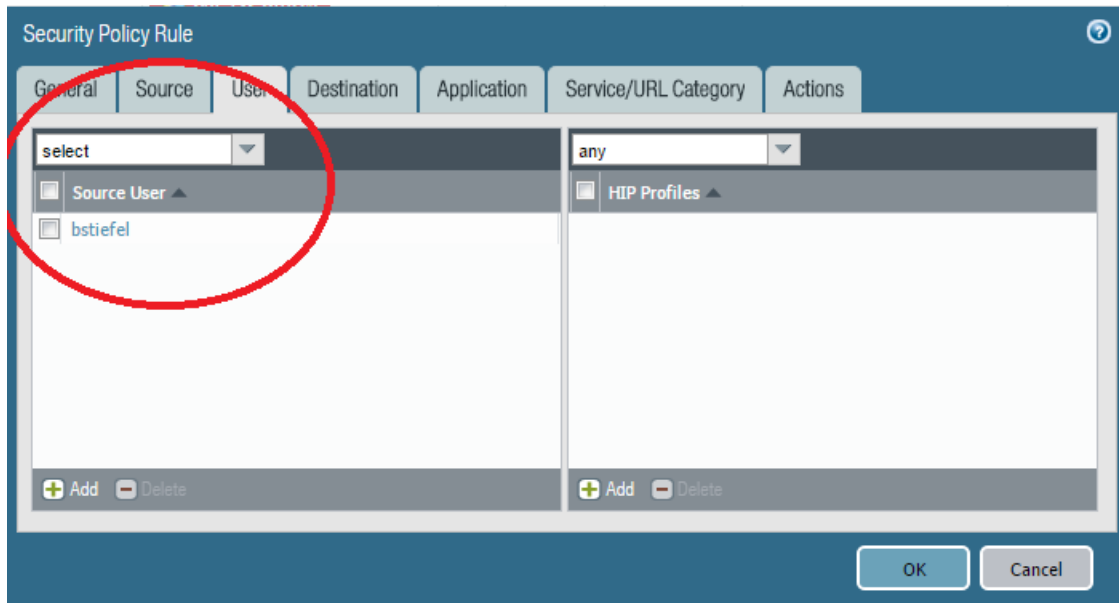
This is what it should **not** look like:



This Allow rule is matching all source users

How to Implement It:

This is what it should look like:



This Allow rule now must match a specific source user

❑ Avoid Using the Any Destination Zone in Allow Rules

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

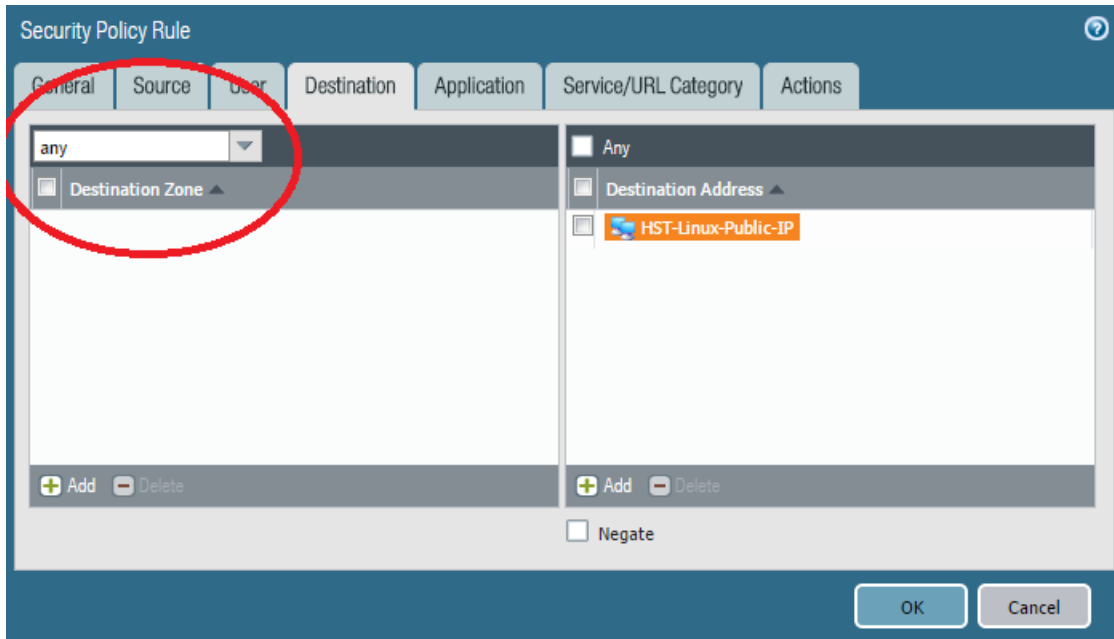
PAN-OS uses a multipart test to determine whether a proposed new connection matches a specific rule in the Security rulebase. One of those tests is to examine the destination interface and determine the destination zone.

Why This Best Practice Is Important:

By default, the process of creating your Security rulebase should start with dropping all traffic with a Cleanup Rule at the bottom and then making *a few, well-considered, narrowly defined exceptions*, which are your Allow rules.

“Narrowly defined” means you should avoid using the *any* destination zone in Allow rules.

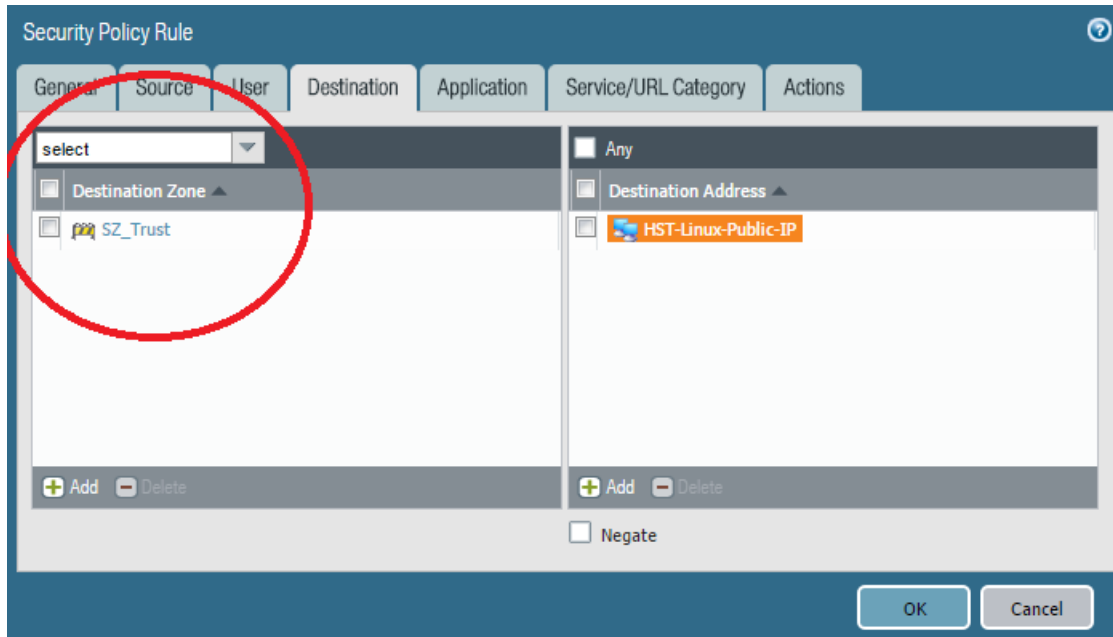
This is what it should **not** look like:



This Allow rule is matching all zones

How to Implement It:

This is what it should look like:



This Allow rule now must match a specific zone

❑ Avoid Using the Any Destination Address in Allow Rules

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

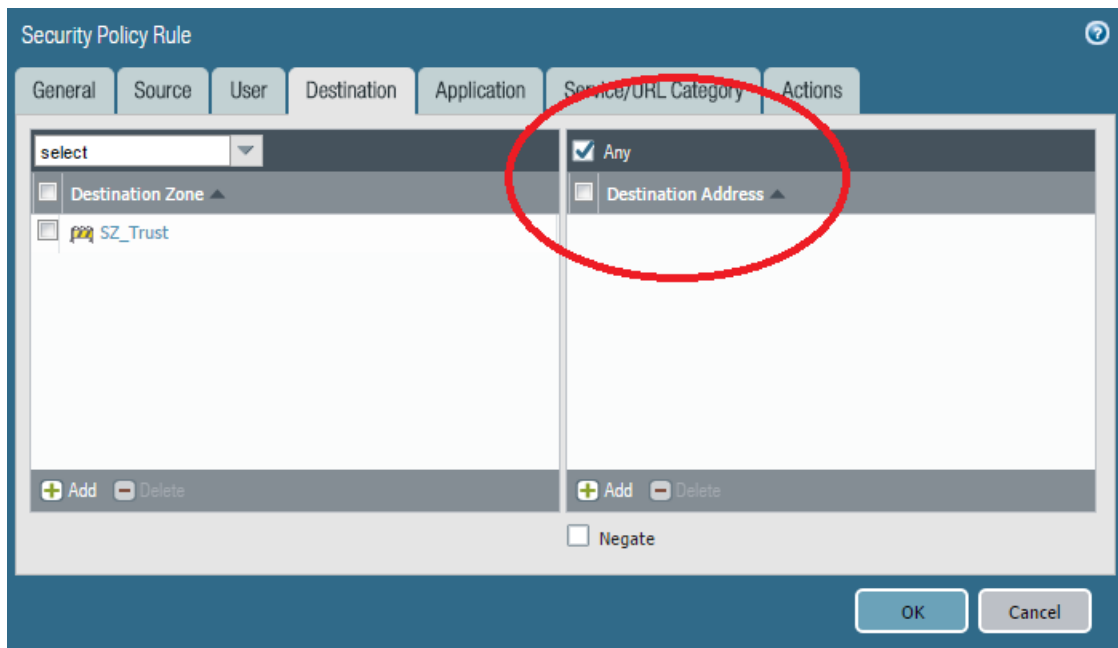
PAN-OS uses a multipart test to determine whether a proposed new connection matches a specific rule in the Security rulebase. One of those tests is to examine the destination address.

Why This Best Practice Is Important:

By default, the process of creating your Security rulebase should start with dropping all traffic with a Cleanup Rule at the bottom and then making *a few, well-considered, narrowly defined exceptions*, which are your Allow rules.

“Narrowly defined” means you should avoid using the *any* destination address in Allow rules.

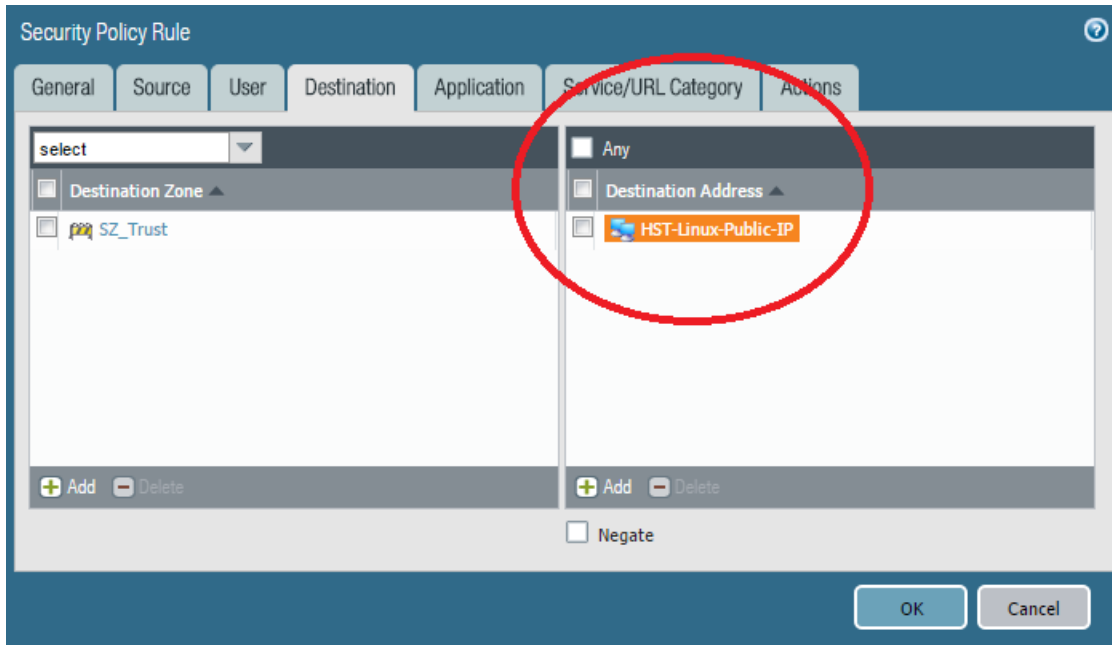
This is what it should **not** look like:



This Allow rule is matching all source addresses

How to Implement It:

This is what it should look like:



This Allow rule now must match a specific address or set of addresses

❑ Avoid Using the Any Application in Allow Rules

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

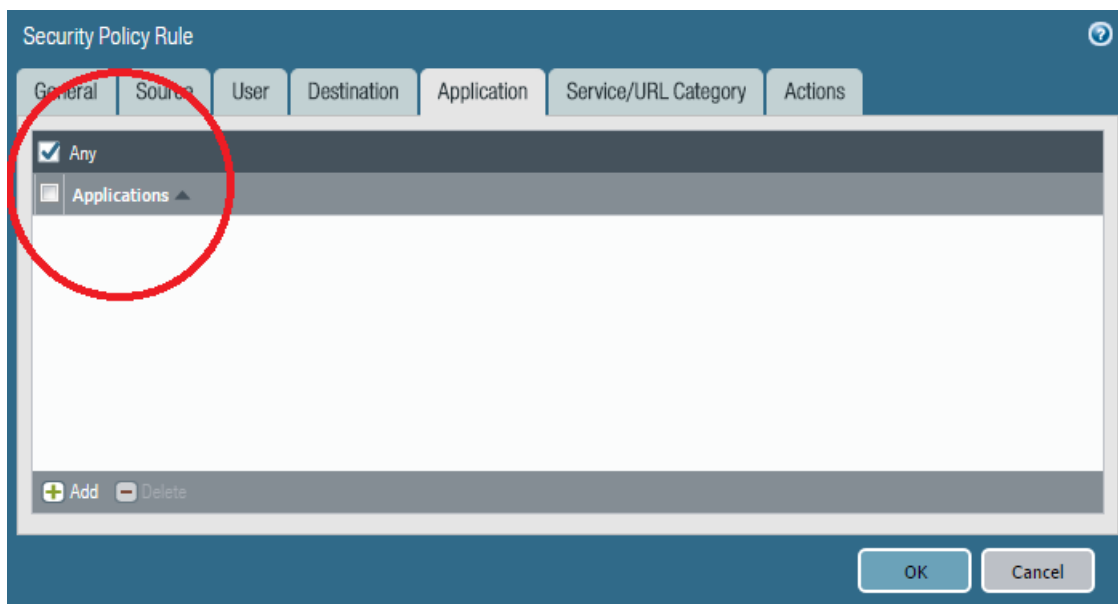
PAN-OS uses a multipart test to determine whether a proposed new connection matches a specific rule in the Security rulebase. One of those tests is to examine the application with App-ID. App-ID is always enabled, so it's just a matter of looking at the result of the examination.

Why This Best Practice Is Important:

By default, the process of creating your Security rulebase should start with dropping all traffic with a Cleanup Rule at the bottom and then making *a few, well-considered, narrowly defined exceptions*, which are your Allow rules.

"Narrowly defined" means you should avoid using the *any* application in Allow rules.

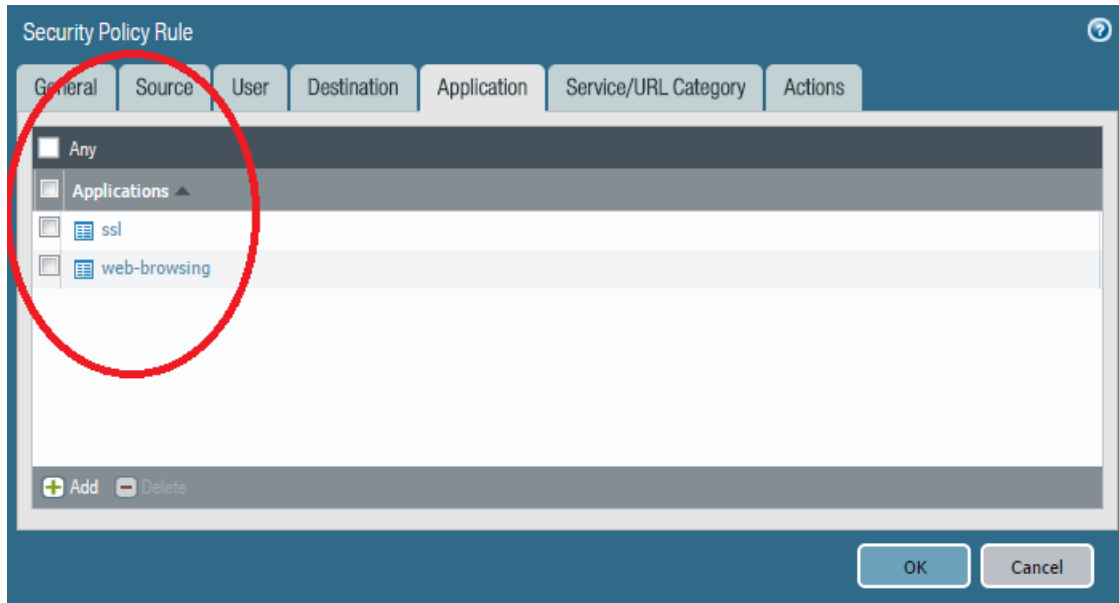
This is what it should **not** look like:



This Allow rule is matching all applications

How to Implement It:

This is what it should look like:



This Allow rule now must match a specific set of applications

❑ Before Dropping Previously Allowed Traffic, Allow and Log to Analyze the Effects

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

It's hard to know all the internal processes in your organization, and which applications they use, and when. Some applications only run daily, weekly, monthly, quarterly, or yearly.

Why This Best Practice Is Important:

Depending upon the size of your organization and its maturity, it might be politically disruptive to simply start blocking connections that might have been working reliably for users for a long time. Something no longer works like it should, and your users have no clue why, leading to inefficient troubleshooting before someone finally starts asking, "Did you do something with the firewall?"

How to Implement It:

If your network has been stable and running for a while, and you've decided to add a Drop rule to drop previously Allowed traffic, it might be best to instead create a separate rule that matches just the traffic you now wish to start dropping, and then continue to Allow it, but log it. In this way you can start seeing exactly what you'll be blocking and choose from among the following actions:

- Work with the user to find a workaround or otherwise become compliant with the security policy again
- Notify the user and explain why these connections will be dropped going forward

Decide to just go ahead and block the traffic anyway (Does Willie the mail boy really need *bittorrent*?)

❑ **Prefer the *application-default* Service in Allow Rules**

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

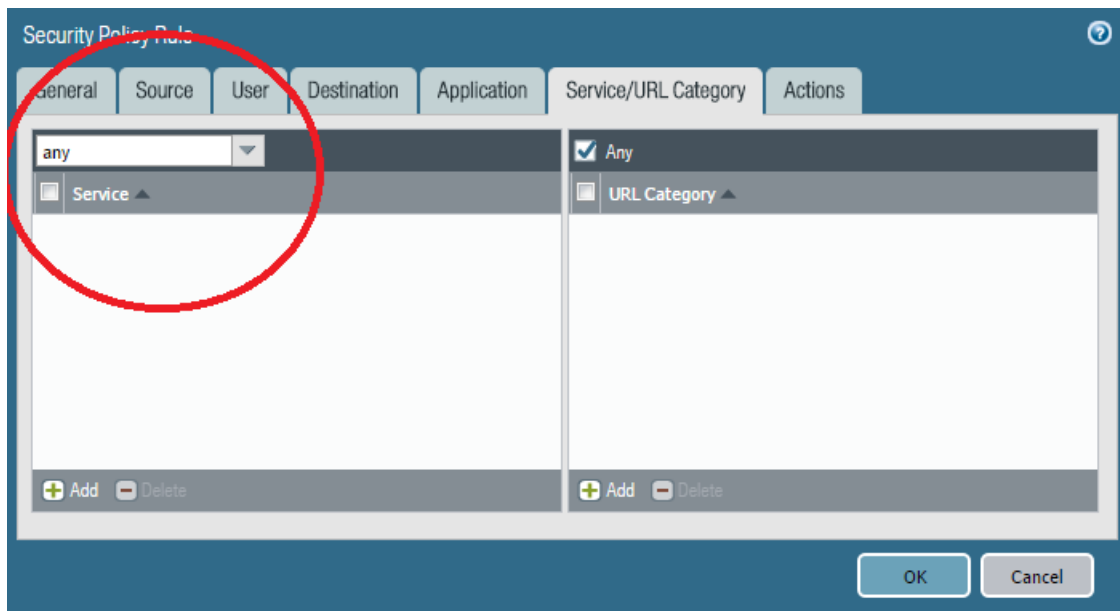
PAN-OS uses a multipart test to determine whether a proposed new connection matches a specific rule in the Security rulebase. One of those tests is to examine the Service.

Why This Best Practice Is Important:

By default, the process of creating your Security rulebase should start with dropping all traffic with a Cleanup Rule at the bottom and then making *a few, well-considered, narrowly defined exceptions*, which are your Allow rules.

“Narrowly defined” means you should avoid using the *any* Service in Allow rules and should choose *application-default* instead. To see what Services are included in an application’s *application-default*, go to <https://applipedia.paloaltonetworks.com/>.

This is what it should **not** look like:



This Allow rule is matching all Services

How to Implement It:

This is what it should look like:

The screenshot shows the 'Security Policy Rule' configuration window. The 'Service/URL Category' tab is active. The 'Service' dropdown menu is set to 'application-default' and is circled in red. The 'URL Category' dropdown menu is set to 'Any'. The 'Add' and 'Delete' buttons are visible at the bottom of each list.

This Allow rule now must match the Services specified in the Applipedia

❑ Restrict Outbound NTP Traffic Destinations

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Network Time Protocol (NTP) is an important tool for network administrators. It allows your network-connected devices to automatically stay synchronized with authoritative clocks in your network or on the Internet.

Why This Best Practice Is Important:

There are two good reasons to limit the destinations your internal hosts can connect to when using the NTP protocol:

- You don't want them connecting to a rogue NTP server, which could return false information.
- You don't want them to be able to launch an amplification DoS attack over the Internet through someone else's NTP server.

How to Implement It:

Step 1: Determine which NTP server you want your internal hosts to connect to:

If you're a small organization, it's fine to use any number of known good public servers.

If you're a larger organization, you may wish to configure your own internal NTP server to serve your internal hosts, which then synchronizes itself with an authoritative NTP server on the Internet.

If you're a really large organization, or have needs for extremely precise time measurements, connect your NTP server to an outdoor GPS receiver. Because each GPS satellite flies with four atomic clocks on board, GPS signals are accurate to within 40 ns, which is much smaller than the time it would take to transmit the time even to an adjacent device, so it's literally more accurate than you could possibly need.

Step 2:

Go to **Policies > Security** and create a rule that allows outbound NTP traffic to connect only to specific NTP servers.

❑ Restrict Outbound DNS Traffic Destinations

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The Domain Name System (DNS) is a global, distributed, hierarchical database that resolves Fully Qualified Domain Names (FQDN) into IP addresses.

Why This Best Practice Is Important:

There are two good reasons to limit the destinations your internal hosts can connect to when requesting DNS lookups:

- You don't want them connecting to a rogue DNS server, which could return false information.
- You don't want them to be able to launch an amplification DoS attack through someone else's DNS server.

How to Implement It:

Step 1: Determine which DNS servers you want your internal hosts to connect to:

If you're a small organization, it's fine to use the DNS servers provided by your Internet Service Provider (ISP).

If you're a larger organization, you may wish to configure your own internal DNS servers to serve your internal hosts.

Step 2:

Go to **Policies > Security** and create a rule that allow outbound DNS traffic to connect only to specific DNS servers.

❑ Restrict Outbound SMTP Traffic Destinations

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Simple Mail Transport Protocol (SMTP) is the standard protocol for sending e-mail.

Why This Best Practice Is Important:

One of the ways malware generates revenue for its creators is to send out spam, quite often in massive quantities. If you were to allow your internal hosts to connect to any open SMTP relay on the Internet, an infected device could send literally tens of thousands of e-mails per hour.

It's important to strictly limit which SMTP servers your internal hosts can connect to reduce the freedom of action of infected internal hosts.

How to Implement It:

Step 1:

Determine which SMTP servers you want your internal hosts to connect to.

Step 2:

Go to **Policies > Security** and create a rule that allow outbound SMTP traffic to connect only to specific SMTP servers.

□ Leverage NGFW Capabilities in Rules

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Content-ID is the Palo Alto Networks big-picture term for all the ways you can do deep packet inspection using Security Profiles. The available Security Profiles include:

1. Antivirus
2. Anti-Spyware
3. Vulnerability Protection
4. URL Filtering
5. File Blocking
6. Wildfire Analysis
7. Data Filtering
8. DoS Protection

Why This Best Practice Is Important:

Old school port-and-protocol firewalls are limited to inspecting traffic at Layer 3 and Layer 4 only, which leaves them blind to most types of attack. Next-Generation Firewalls (NGFWs) allow much more sophisticated deep packet inspection by also examining Layers 5-7 and keeping much more sophisticated state information about sessions.

At a minimum, attach Antivirus, Anti-Spyware, Vulnerability Protection and URL Filtering Security Profiles to your Allow rules.

How to Implement It:

Edit an Allow Security Policy Rule:

Security Policy Rule

General Source User Destination Application Service/URL Category Actions

Action Setting

Action: Allow

☐ Send ICMP Unreachable

Profile Setting

Profile Type: Profiles

Antivirus: default

Vulnerability Protection: default

Anti-Spyware: default

URL Filtering: default

File Blocking: None

Data Filtering: None

WildFire Analysis: None

Log Setting

☐ Log at Session Start

☒ Log at Session End

Log Forwarding: None

Other Settings

Schedule: None

QoS Marking: None

☐ Disable Server Response Inspection

OK Cancel

*Configuring individual Security Profiles****What Else You Need to Know:***

Individual Security Profiles can be bundled into Security Profile Groups for easier attachment to Security Policy rules.

❑ Don't Use a Security Profile in a Drop Rule

Improve Security		Improve Manageability	
Improve Performance	X	Improve High Availability	

Background Information:

Content-ID is the Palo Alto Networks big-picture term for all the ways you can do deep packet inspection using Security Profiles. The available Security Profiles include:

1. Antivirus
2. Anti-Spyware
3. Vulnerability Protection
4. URL Filtering
5. File Blocking
6. Wildfire Analysis
7. Data Filtering
8. DoS Protection

Why This Best Practice Is Important:

The deep packet inspection of Content-ID takes CPU cycles and most of the time once you've decided to drop a packet there's no more benefit to be gained from further inspection.

How to Implement It:

Edit a Drop Security Policy Rule:

The screenshot shows the 'Security Policy Rule' configuration window. The 'Action' dropdown is set to 'Drop' and the 'Profile Type' dropdown is set to 'None'. Both are circled in red. Other settings include 'Log at Session Start', 'Log at Session End', 'Log Forwarding', 'Schedule', 'QoS Marking', and 'Disable Server Response Inspection'.

Let the packets hit the floor

❑ Prefer the Drop Action Over Deny or Reset

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Once the firewall determines that a proposed new connection matches a specific Security policy rule, it takes the action specified in the Actions tab. There are several choices:

- **Allow:** Allows the traffic
- **Deny:** Blocks the connection, and enforces the default Deny action defined for the application being blocked.
- **Drop:** Silently drops the connection, without sending a TCP reset.
- **Reset (Client, Server, or Both):** Sends a TCP reset.

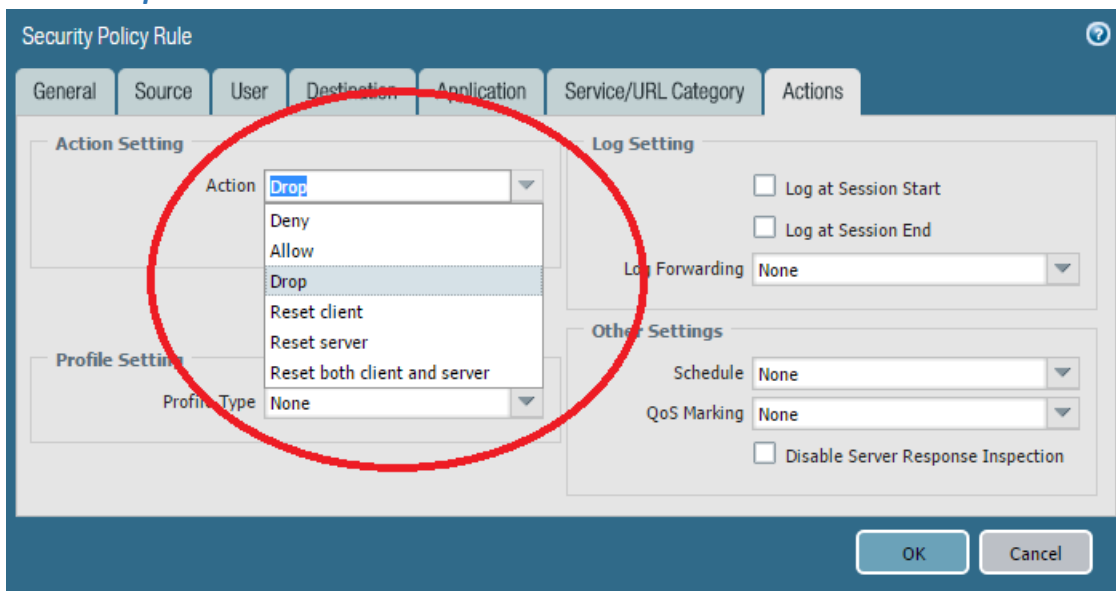
Why This Best Practice Is Important:

It's best to be quiet in the jungle.

When you send a TCP reset, you announce your presence as a live, thinking TCP/IP device on the wire. When you just silently drop a proposed new connection, you are indistinguishable from an unplugged cable and your attacker can't even confirm your presence.

Unless you're trying to be helpful while troubleshooting (see another Best Practice in this book), use Drop instead of Deny or Reset on unwanted connections.

How to Implement It:



The Actions tab and the Action selection drop down choices

❑ Don't Normally Send ICMP Unreachable Messages

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The ICMP Type 3 Destination Unreachable message is generated by a router to inform the source host that the destination unicast address is unreachable. The goal is to inform the source so that it may gracefully close or clear the session and prevents some applications from breaking.

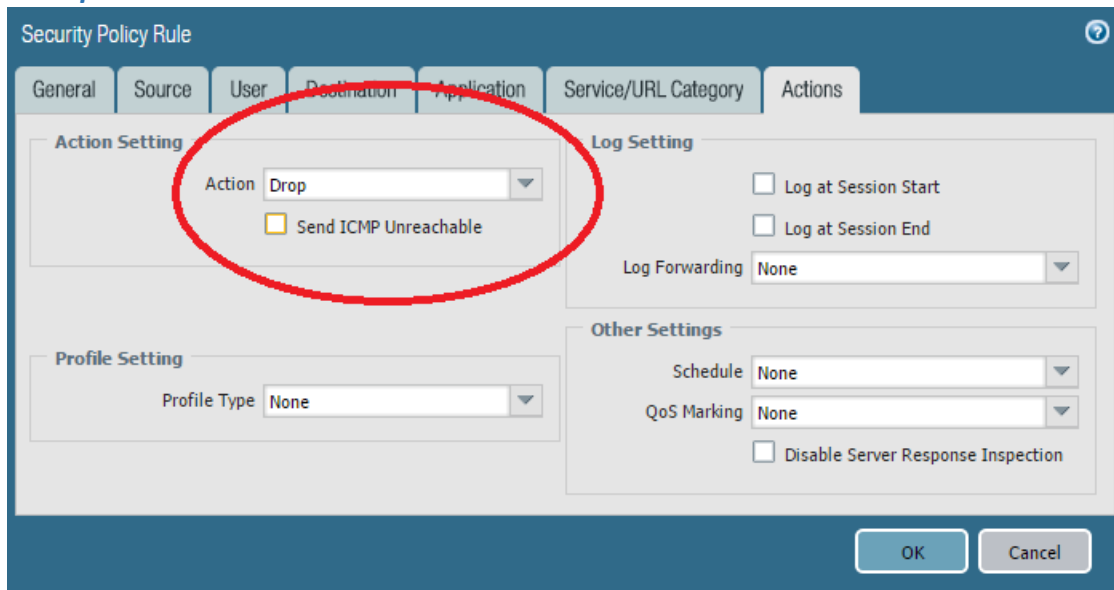
Why This Best Practice Is Important:

It's best to be quiet in the jungle.

When you send an ICMP Unreachable Message, you announce your presence as a live, thinking TCP/IP device on the wire. When you just silently drop a proposed new connection, you are indistinguishable from an unplugged cable and your attacker can't even confirm your presence.

Unless you're trying to be helpful while troubleshooting (see another Best Practice in this book), don't enable Send ICMP Unreachable Messages on unwanted connections.

How to Implement It:



The screenshot shows the 'Security Policy Rule' configuration window. The 'Action' dropdown is set to 'Drop', and the 'Send ICMP Unreachable' checkbox is checked. A red circle highlights the 'Drop' action and the 'Send ICMP Unreachable' checkbox. Other settings include 'Log at Session Start' and 'Log at Session End' (both unchecked), 'Log Forwarding' set to 'None', 'Profile Type' set to 'None', 'Schedule' set to 'None', 'QoS Marking' set to 'None', and 'Disable Server Response Inspection' (unchecked). The 'OK' and 'Cancel' buttons are at the bottom right.

The Send ICMP Unreachable setting

What Else You Need to Know:

The Send ICMP Unreachable option is available only on Layer3 interfaces, and only with the Drop or Reset actions.

❑ Consider DSRI for Internet-Facing Servers

Improve Security		Improve Manageability	
Improve Performance	X	Improve High Availability	

Background Information:

Your Palo Alto Networks firewall does seven layer deep-packet inspection. Even with specialized hardware assistance, this takes CPU cycles.

DSRI (Disable Server Response Inspection, or “*Don’t Sniff Returning Information*”, if you can’t remember the real definition), tells the firewall to not inspect packets originating from the server side of a session. It’s a Boolean condition applied on a per-security rule basis.

Why This Best Practice Is Important:

You may wish to enable DSRI if all of the following three conditions are true:

- Your firewall is nearing its throughput capacity and performance constraints are becoming an issue.
- You’ve tried other reasonable steps to reduce the load on your firewall.
- You have Internet-facing servers that handle a lot of traffic.

How to Implement It:

Edit an Allow Security Rule that permits traffic to an Internet-facing server:

Enabling DSRI in a Security Policy rule

What Else You Need to Know:

- Because there’s always a small chance that something interesting might be caught with a thorough inspection of this return traffic, don’t enable this unless you really need to.
- If you’re really running out of capacity, talk with your Systems Engineer and get a properly-sized firewall in

there. We've got really big boxes waiting for you with your name on them.
DSRI is always disabled by default; you must manually enable it to get it to work.

•

❑ Block Internet Connections To and From Private Non-Routable IP Addresses

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

These are the three famous RFC-1918 non-routable IPv4 address blocks:

IPv4 Subnet: Range:

10.0.0.0 /8	10.0.0.1 – 10.255.255.255
172.16.0.0 /12	172.16.0.0 – 172.31.255.255
192.168.0.0 /16	192.168.255.255

Every network administrator can remember the first and the third subnets, but nobody can remember the middle one.

Why This Best Practice Is Important:

One of the features of the “non-routable” IP address blocks are that they’re never assigned anywhere on the public Internet, and a properly configured Internet BGP router will drop packets heading for these addresses.

Therefore, any packet that is transiting your Internet gateway and heading to or from a non-routable IP address on the Internet side is deeply suspicious and should be dropped.

How to Implement It:

You should already have a pair of rules at the top of your Security Rulebase that block all sessions to and from an IP address/subnet blacklist. Add these three subnets to that blacklist.

What Else You Need to Know:

Obviously, this Best Practice only makes sense on firewalls connected with external public IP addresses. Know where your NAT boxes are.

❑ Block Internet Connections To and From Bogons and Fullbogons

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

A bogon prefix is a route that should never appear in the Internet routing table. A packet routed over the public Internet (not including over VPNs or other tunnels) should never have a source address in a bogon range. These are commonly found as the source addresses of DDoS attacks.

Bogons come in two flavors, Bogons and Fullbogons:

Bogons:

Bogons are defined as Martians (private and reserved addresses defined by RFC 1918, RFC 5735, and RFC 6598) and netblocks that have not been allocated to a regional internet registry (RIR) by the Internet Assigned Numbers Authority.

As of the time of this writing, this is the list of IPv4 bogons:

- 0.0.0.0/8
- 10.0.0.0/8
- 100.64.0.0/10
- 127.0.0.0/8
- 169.254.0.0/16
- 172.16.0.0/12
- 192.0.0.0/24
- 192.0.2.0/24
- 192.168.0.0/16
- 198.18.0.0/15
- 198.51.100.0/24
- 203.0.113.0/24
- 224.0.0.0/4
- 240.0.0.0/4

Notice this list includes the RFC 1918 non-routable addresses.

Fullbogons:

Fullbogons are a larger set which also includes IP space that has been allocated to an RIR, but not assigned by that RIR to an actual ISP or other end-user. IANA maintains a convenient IPv4 summary page listing allocated and

reserved netblocks, and each RIR maintains a list of all prefixes that they have assigned to end-users.

The list of fullbogons is much longer than the list of bogons.

Why This Best Practice Is Important:

Because nothing legitimate ever transits the Internet with a bogon source address, and because bad stuff often does, it's best to drop all packets with a bogon or fullbogon source or destination address.

How to Implement It:

Team Cymru maintains a list of bogons and fullbogons and is the source of the technical information for this Best Practice:

<http://www.team-cymru.org/bogon-reference.html>

You should already have a pair of rules at the top of your Security Rulebase that block all sessions to and from an IP address/subnet blacklist. Add the bogons and fullbogons to that blacklist.

What Else You Need to Know:

It is important to realize that the bogon and fullbogon lists are NOT static lists. You need to either manually update the list on a regular basis, or use an automated method. There is at least one Internet tale circulating of a difficult-to-troubleshoot network problem getting resolved when someone realized that a block of IP addresses had come off the fullbogon list but a firewall had not yet been updated.

The fullbogons list has multiple changes every day.

Objects

❑ Create and Use an Object Naming Convention

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

There are many different types of objects that you can create and name in PAN-OS and use as part of your security policy.

Why This Best Practice Is Important:

It's important to stay clear on which type of object you're dealing with in each part of the GUI.

Because each time you name an object you're sending a message to your future self, it's important to ensure names are clear and concise. An easy way to help disambiguate objects is by naming each with a prefix that identifies its class of object. In this way, whenever you see an object, you'll instantly know not only its "given" name, but also its object type.

This convention is similar to the convention used in the international travel industry of showing the passenger's family name first, in all capitals, followed by the given name, using initial capitalization, as in "**STIEFEL, Barry**". Any alphabetized list of passengers will group family members together and it's always clear which name is the given name.

How to Implement It:

Here is a table of the most commonly created objects in PAN-OS. For each, there's a recommended prefix to go at the beginning of each name.

Object Type:	Object Subtype:	Prefix:
Addresses	An IP host or device	HST_
Addresses	An IP subnet	NET_
Addresses	An IP range	RNGE_
Address Groups		GRP_
Application Groups		AGRP_
Application Filters		AFLT_
Interface Management Profile		IMP_
Security Zones		SZ_
Services		SVC_
Service Groups		SGRP_
Virtual Routers		VR_

Virtual Wires		VW_
Zone Protection Profile		ZPP_

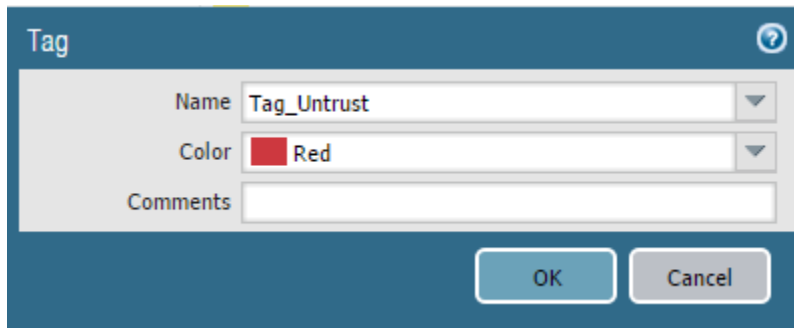
□ Create and Use an Object Color Convention

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Tags Can Acquire a Color:

PAN-OS allows you to attach a color to a tag, like this:



This tag is colored Red

Tagged Objects Acquire the Color of The First Tag in Their Tag List:

PAN-OS allows you to attach one or more tags to the following Objects:

- Any Address or Address Group
- Any Application (built-in, but not custom), and not any Application Group
- Any Service or Service Group
- Any of the rules in the seven Security Policies

You control the order of the tags in the list of tags.

Any object with one or more attached tags will take on the color of the *first* tag in its tag list.

This Address Object has a Red colored tag, so this object will now show as Red wherever it appears

The Special Tag Color Trick for Security Zones:

If you wish to attach a tag to a Security Zone because you'd like the zone to take on the color of the tag, you can't do it directly. Instead, you need to create a tag with the *same name* as the Security Zone (when you finish typing the name into the Name field, the name of the zone will appear and be clickable; click on it), and *when you save the tag the Security Zone will magically take on the color of the tag, even though you can't attach a tag to a Security Zone*.

The Available Tag Colors:

There are 17 possible colors to choose from, including two of them (#2 and #17) that are for reasons unknown both labeled "Green".

Why This Best Practice Is Important:

Being trichromats (look it up), humans are especially good at seeing colors and classifying objects into groups by color. Therefore, if we can use color clues to help classify objects in the firewall GUI, it can increase comprehension and reduce errors.

Of particular benefit is the ability to spot an errant member of a group. For example, if the permitted source addresses in a rule include 12 objects with a color indicating "trusted insiders", and a 13th member has a color saying "untrusted outsider", this should flag your attention for review.

How to Implement It:

The first question is what information do we want to convey with color? Because the purpose of a firewall is to control traffic between zones of differing security posture, it's best to have a color convention that marks objects by security risk.

The second question is to decide which colors to use. A smart choice would be to leverage an existing globally recognized color scheme, and the most well-known must be the 1968 Vienna Convention on Road Signs and Signals, specifying the use of Red, Amber, and Green for "Proceed", "Proceed with caution", and "Stop".

Another existing globally recognized color scheme is that of the spectral colors in the visible spectrum:

Color	Wavelength	Frequency	Photon energy
violet	380–450 nm	668–789 THz	2.75–3.26 eV
blue	450–495 nm	606–668 THz	2.50–2.75 eV
green	495–570 nm	526–606 THz	2.17–2.50 eV
yellow	570–590 nm	508–526 THz	2.10–2.17 eV
orange	590–620 nm	484–508 THz	2.00–2.10 eV
red	620–750 nm	400–484 THz	1.65–2.00 eV

Because the traffic signal convention can be argued to be a subset of the spectral color table, the spectral colors now seem like a good starting point for our analysis.

We next need to find a way to map these six spectral colors to the colors that we have available in PAN-OS. Therefore, we must pare down the list of 17 PAN-OS colors into a more manageable subset and assign meanings to the ones we choose.

Visibility and differentiability are important considerations in choosing a color convention, so the next step is to examine and score the PAN-OS colors:

Color:	Visibility:	Differentiability:
Red	High	High
Green #1	Medium	Medium
Blue	Medium	Too close to Green and Blue Gray, and inferior to Cyan
Yellow	High	High
Copper	Medium	Think of it as “Dark Orange”
Orange	High	High
Purple	High	High
Gray	Medium	Medium
Light Green	Low	Too close to some of the GUI background colors
Cyan	High	A better Blue than Blue
Light Gray	Low	Too close to some of the background colors
Blue Gray	Medium	Too close to Green

Lime	Medium	Medium, but the white text doesn't show up well against this background
Black	High	High
Gold	High	Too close to Orange
Brown	Low	Low
Green #2	Medium	Medium, but an apparent duplicate of the other Green

By looking at how these colors score, and trying to match to the spectral colors, we can now come up with a color scheme that gives the signals that we want using the colors available in PAN-OS, rank ordered from top to bottom from less trusted down to more trusted, plus a special color for firewalls:

PAN-OS Color	Meaning	Tag Name
Red	Untrusted, or the Internet	Untrusted
Orange	DMZ (Higher Risk)	DMZ-Higher Risk
Yellow	DMZ (Lower Risk)	DMZ-Lower Risk
Green	Internal devices (Higher Risk)	Internal-Higher Risk
Cyan (for Blue)	Internal devices (Lower Risk)	Internal-Lower Risk
Purple (for Violet)	Highly secure internal networks	Trusted
Black	Firewalls	Firewalls

The final steps are to create these tags and then assign them to your Addresses, Address Groups, Applications, Services, Service Groups, and Policy Rules.

Address and Address Group Objects

□ Create an Address Object for Every Host, Subnet, Range and FQDN in Your Policy

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

PAN-OS lets you create Address objects of three different types:

1. **IP Netmask:** Consisting of an IP address (identifying a host), or followed by an optional “/24”-style subnet mask (identifying a subnet).
2. **IP Range:** Consisting of the two included boundary IP addresses, in this format: 192.168.1.100-192.168.1.109
3. **FQDN:** Consisting of a Fully Qualified Domain Name that the firewall occasionally updates with a reverse DNS query

Why This Best Practice Is Important:

Humans are terrible at remembering and processing IP addresses, so here’s your chance to assign useful names to your network objects. Your tired brain will thank you later.

How to Implement It:

Go to **Objects > Addresses**.

Create Address objects for every host, subnet, range and FQDN that you expect to add to your Security Rulebase.

□ Create a Meaningful Description for Every Address and Address Group Object

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

PAN-OS provides a Description field for every Address and Address Group Object.

Why This Best Practice Is Important:

Do your future self a favor and always enter a useful description in this field.

How to Implement It:

Go to **Objects > Address** and **Objects > Address Groups**.

Edit each of your Address and Address Group objects and fill in the Description with something clear and informative.

Services Objects

❑ Create a Meaningful Description for Every Service Object

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

PAN-OS provides a Description field for every Service Object.

Why This Best Practice Is Important:

Your future self with thank you for leaving an informative Description when you create this object.

How to Implement It:

1. Go to **Objects > Services**.
2. Edit each of your Services Objects.
3. Create an informative Description.

Service

Name: UDP-53413

Description: Netis Systems routers vulnerability

Protocol: ☐ TCP ☒ UDP

Destination Port: 53413

Source Port: [] >= 0

Port can be a single port #, range (1-65535), or comma separated (80, 8080, 443)

Tags:

OK Cancel

Now here's a clear and informative Service Description

Security Profile #1: Antivirus

❑ Create a Strict Antivirus Security Profile

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The Antivirus and WildFire updates you receive as Dynamic Updates are your defense against *known threats*. The Antivirus Security Profile is how you apply these signatures to your traffic and block transit of malware.

Why This Best Practice Is Important:

You want to drop malware regardless of which protocol it's using for transit. The built-in *default* Antivirus Security Profile doesn't drop (or reset) connections over the *smtp* protocol and only alerts on WildFire signatures, so you need to create a new Antivirus Security Profile called *strict*. IMAP and POP3 connections should not be set to 'reset-both'. Both protocols will retry to fetch the requested email which could result in a denial of service condition. For the SMTP protocol a 541 response will be delivered to the sending SMTP server to prevent it from resending a blocked message.

How to Implement It:

1. Go to **Objects > Security Profiles > Antivirus**.
2. Create an Antivirus Security Profile that looks like this:

The screenshot shows the 'Antivirus Profile' configuration window. The 'Name' field is set to 'strict'. The 'Description' field is empty. The 'Antivirus' tab is selected, and the 'Packet Capture' checkbox is unchecked. The 'Decoders' section contains a table with the following data:

Decoder	Action	WildFire Action
http	default (reset-both)	default (reset-both)
smtp	reset-both	reset-both
imap	default (alert)	default (alert)
pop3	default (alert)	default (alert)
ftp	default (reset-both)	default (reset-both)
smb	default (reset-both)	default (reset-both)

The 'Application Exception' section is empty, showing 0 items. At the bottom right are 'OK' and 'Cancel' buttons.

This strict Antivirus profile resets the connection upon every detection of malware through most of the decoders

What Else You Need to Know:

It's OK to use *reset-both* as the Action here, instead of *drop*, because the firewall cannot hide its presence in these circumstances and the *reset-both* Action might be helpful to both ends of the connection.

□ Attach an Antivirus Security Profile to Every Allow Rule

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The Antivirus and WildFire updates you receive as Dynamic Updates are your defense against *known threats*. The Antivirus Security Profile is how you apply these signatures to your traffic and block transit of malware.

Why This Best Practice Is Important:

If you're not scanning for malware in your Allow rules, then you're not protected. You need to attach an Antivirus Security Profile to every Security Policy rule where the Action is Allow.

How to Implement It:

1. Go to **Policies > Security**.
2. Edit your Allow rules
3. Attach an Antivirus Security Profile to every Allow rule

The screenshot shows the 'Security Policy Rule' configuration window. The 'General' tab is selected. In the 'Action Setting' section, the 'Action' is set to 'Allow'. In the 'Profile Setting' section, the 'Profile Type' is set to 'Antivirus' and the 'Profile' is set to 'default'. The 'Log Setting' section shows 'Log at Session Start' and 'Log at Session End' as unchecked, and 'Log Forwarding' is set to 'None'. The 'Other Settings' section shows 'Schedule' and 'QoS Marking' as 'None', and 'Disable Server Response Inspection' as unchecked. The 'OK' and 'Cancel' buttons are at the bottom right.

This Allow Security Rule now has an Antivirus Security Profile attached to it

Security Profile #2: Anti-Spyware

□ Attach an Anti-Spyware Security Profile to Every Allow Rule

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

An important requirement for most spyware is that it can “phone home” to its Command-and-Control servers. An Anti-Spyware Security Profile can detect and block these connections.

There are two built-in Anti-Spyware Security Profiles, *default* and *strict*.

Why This Best Practice Is Important:

If you’re not blocking “phone home” connections from spyware running in infected hosts, then you’re not protected. You need to attach an Anti-Spyware Security Profile to every Security Policy rule where the Action is Allow.

How to Implement It:

1. Go to **Policies > Security**.
2. Edit your Allow rules
3. Attach an Anti-Malware Security Profile to every Allow rule
4. Use the strict policy for both Anti-Spyware and Vulnerability Protection profiles.

Security Policy Rule

General Source User Destination Application Service/URL Category **Actions**

Action Setting

Action: **Allow**

☐ Send ICMP Unreachable

Profile Setting

Profile Type: Profiles

Antivirus: None

Vulnerability Protection: None

Anti-Spyware: strict

URL Filtering: None

File Blocking: None

Data Filtering: None

WildFire Analysis: None

Log Setting

☐ Log at Session Start

☐ Log at Session End

Log Forwarding: None

Other Settings

Schedule: None

QoS Marking: None

☐ Disable Server Response Inspection

OK Cancel

This Allow Security Rule now has an Anti-Malware Security Profile attached to it

Security Profile #3: Vulnerability Protection

□ Attach a Vulnerability Protection Security Profile to Every Allow Rule

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

A Vulnerability Protection Security Profile stops attempts to exploit system flaws or gain unauthorized access to systems.

Vulnerability Protection profiles help protect against:

- Buffer overflows
- Illegal code execution
- Other attempts to exploit system vulnerabilities

Why This Best Practice Is Important:

The Vulnerability Protection Security Profile has signatures on 9,300+ vulnerabilities that might be lurking in your network and can detect and block all of them. Because your internal devices may not have had patches released for all of the vulnerabilities, and because you almost certainly haven't updated all your devices with their latest patches, you need a Vulnerability Protection Security Profile to block these attacks at your network perimeter.

How to Implement It:

1. Go to **Policies > Security**.
2. Edit your Allow rules
3. Attach a Vulnerability Protection Security Profile to every Allow rule

Security Policy Rule

General Source User Destination Application Service/URL Category Actions

Action Setting

Action: Allow

☐ Send ICMP Unreachable

Profile Setting

Profile Type: Profiles

Antivirus: None

Vulnerability Protection: strict

Anti-Spyware: None

URL Filtering: None

File Blocking: None

Data Filtering: None

WildFire Analysis: None

Log Setting

☐ Log at Session Start

☐ Log at Session End

Log Forwarding: None

Other Settings

Schedule: None

QoS Marking: None

☐ Disable Server Response Inspection

OK Cancel

This Allow Security Rule now has a Vulnerability Protection Security Profile attached to it

Security Profile #4: URL Filtering

□ Attach a URL Filtering Protection Security Profile to Every Web Browsing Allow Rule

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The URL Filtering License allows you to use Security Policy rules to enforce web access based on dynamic URL categories.

Why This Best Practice Is Important:

URL Filtering can reduce two risks to your organization:

- Security Risk: By blocking access to sites categorized as *malware* and *phishing*.
- Legal Risk: By blocking access to websites that can incur Human Resources and other legal liabilities.

How to Implement It:

1. Go to **Policies > Security**.
2. Edit your Allow rules
3. Attach a URL Filtering Security Profile to every Web Browsing Allow rule.

What Else You Need to Know:

Caveat: While this is the best practice, it may not fit your organization's needs. Use your best judgment, and if this is too restrictive, consider just set these URLs to Alert so they can be logged and monitored instead..

Security Policy Rule

General Source User Destination Application Service/URL Category Actions

Action Setting

Action: Allow

☐ Send ICMP Unreachable

Profile Setting

Profile Type: Profiles

Antivirus: None

Vulnerability Protection: None

Anti-Spyware: None

URL Filtering: default

File Blocking: None

Data Filtering: None

WildFire Analysis: None

Log Setting

☐ Log at Session Start

☒ Log at Session End

Log Forwarding: None

Other Settings

Schedule: None

QoS Marking: None

☐ Disable Server Response Inspection

OK Cancel

This Allow Security Rule now has a URL Filtering Security Profile attached to it

❑ Prefer PAN-DB URL Filtering Over BrightCloud

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

PAN-OS supports two URL filtering vendors:

PAN-DB:

This is the Palo Alto Networks developed URL filtering database. It's tightly integrated into PAN-OS and the Palo Alto Networks threat intelligence cloud. It provides high-performance local caching for maximum inline performance on URL lookups, and offers coverage against malicious URLs and IP addresses.

As WildFire, which is a part of the Palo Alto Networks threat intelligence cloud, identifies unknown malware, zero-day exploits, and advanced persistent threats (APTs), the PAN-DB database is updated with information on malicious URLs so that you can block malware downloads, and disable Command and Control (C2) communications to protect your network from cyber threats.

To see how a particular website is categorized, go to:

<https://urlfiltering.paloaltonetworks.com/>

BrightCloud:

This is a third-party URL database that is that is integrated into PAN-OS firewalls. For information on the BrightCloud URL database, visit <http://brightcloud.com>.

Why This Best Practice Is Important:

PAN-DB is just better and we've been recommending for several years now that customers choose PAN-DB over BrightCloud. The BrightCloud option exists now mostly for historic reasons.

How to Implement It:

1. Go to **Device > Licenses**
2. Ensure your PAN-DB license is current.
3. Activate PAN-DB if it's not already activated.
4. Download the seed file if you haven't done it already.

What It Looks Like After You've Implemented It:

PA-VM Date Issued July 11, 2016 Date Expires Never Description Standard VM-300	AutoFocus Device License Date Issued October 19, 2016 Date Expires October 11, 2021 Description AutoFocus Device License
BrightCloud URL Filtering Date Issued July 12, 2016 Date Expires July 12, 2019 Description BrightCloud URL Filtering Active No (Activate)	GlobalProtect Gateway Date Issued July 12, 2016 Date Expires July 12, 2019 Description GlobalProtect Gateway License
PAN-DB URL Filtering Date Issued July 12, 2016 Date Expires July 12, 2019 Description Palo Alto Networks URL Filtering License Active Yes Download Status 2016-12-19 15:47:55 PAN-DB download: Finished successfully. Re-Download	Threat Prevention Date Issued July 12, 2016 Date Expires July 12, 2019 Description Threat Prevention
WildFire License Date Issued July 12, 2016 Date Expires July 12, 2019 Description WildFire signature feed, integrated WildFire logs, WildFire API	License Management Retrieve license keys from license server Activate feature using authorization code Manually upload license key Deactivate VM

Check to ensure there's a valid license expiration date and that "Active" is set to "Yes".

What Else You Need to Know:

Even if you have valid licenses for both URL Filtering databases, enabling one will automatically disable the other.

□ Block Access to Malicious URL Categories

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

PAN-DB partitions websites that it knows about into approximately 65 categories. While there can be no end to debates on what's the proper list of permitted or blocked categories for a particular organization, there do seem to be two categories that everyone should block:

- malware
- phishing

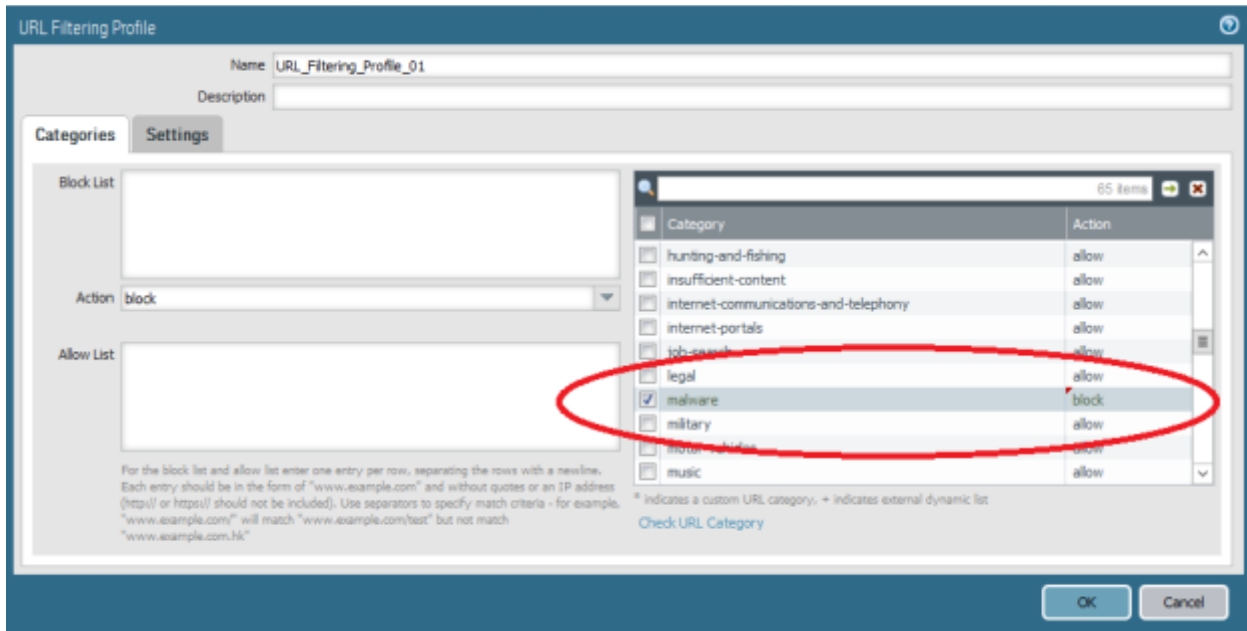
Why This Best Practice Is Important:

Regardless of your political leanings or even your opinion of the proper role of the Human Resources Department, you should be blocking these two categories of URLs.

How to Implement It:

1. Go to **Objects > Security Profiles > URL Filtering**
2. Edit a URL Filtering Security Profile
3. Select at least the *malware* and *phishing* categories
4. Attach this URL Filtering Security Profile to all Security Policy Rules that permit outbound web traffic.

What It Looks Like After You've Implemented It:



Start with the basics: Block malware and phishing sites

□ Block Access to “Unknown” URLs

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

PAN-DB partitions websites that it knows about into approximately 65 categories. One of these categories is the special “unknown” category. URLs in this category are newly created domains that have not yet been categorized by PAN-DB. By itself, that might not be suspicious, but lots of malware uses a constant stream of newly-created domains to keep evading domain blocking.

Why This Best Practice Is Important:

It’s best to block “unknown” URLs because of their higher-than-normal association with malware Command-and-Control domains.

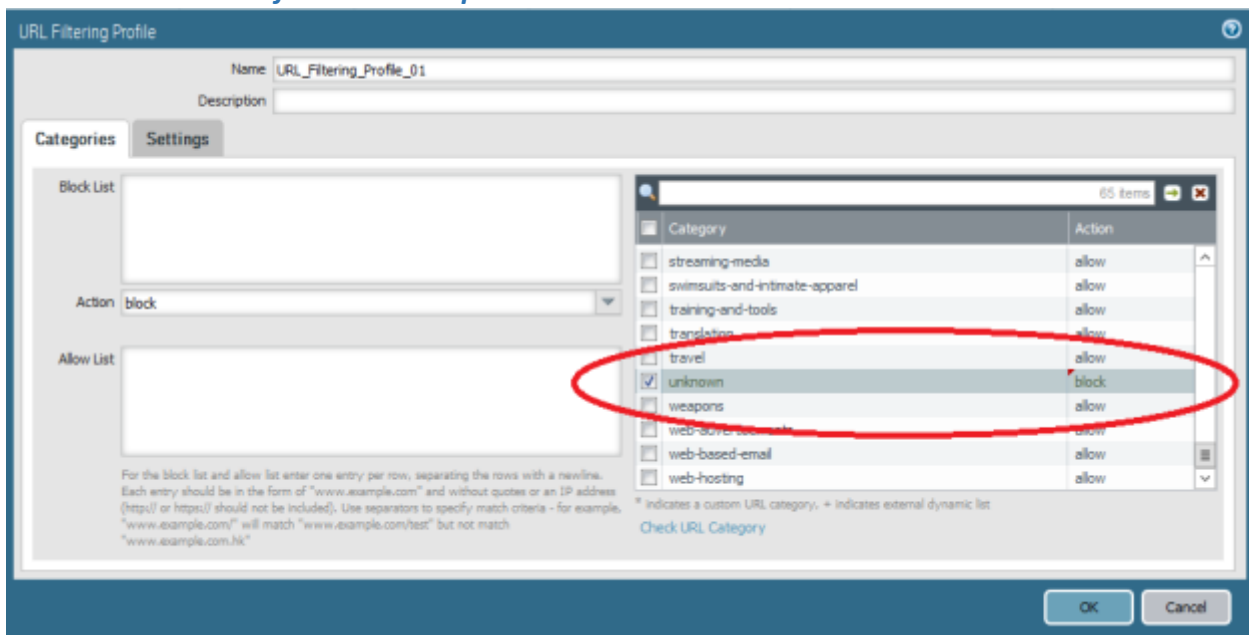
PAN-DB is usually pretty good at programmatically categorizing names quickly after they’re created. If you’d like to accelerate the process, or request a recategorization of a URL, go to:

<https://urlfiltering.paloaltonetworks.com/>

How to Implement It:

1. Go to **Objects > Security Profiles > URL Filtering**
2. Edit a URL Filtering Security Profile
3. Select the *unknown* category
4. Attach this URL Filtering Security Profile to all Security Policy Rules that permit outbound web traffic.

What It Looks Like After You’ve Implemented It:



Blocking the unknown category

□ Log HTTP Header Information

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

A web browser's request for a web page contains more metadata in the header than just the destination URL and IP address. Here are three additional fields that might be useful to a firewall administrator:

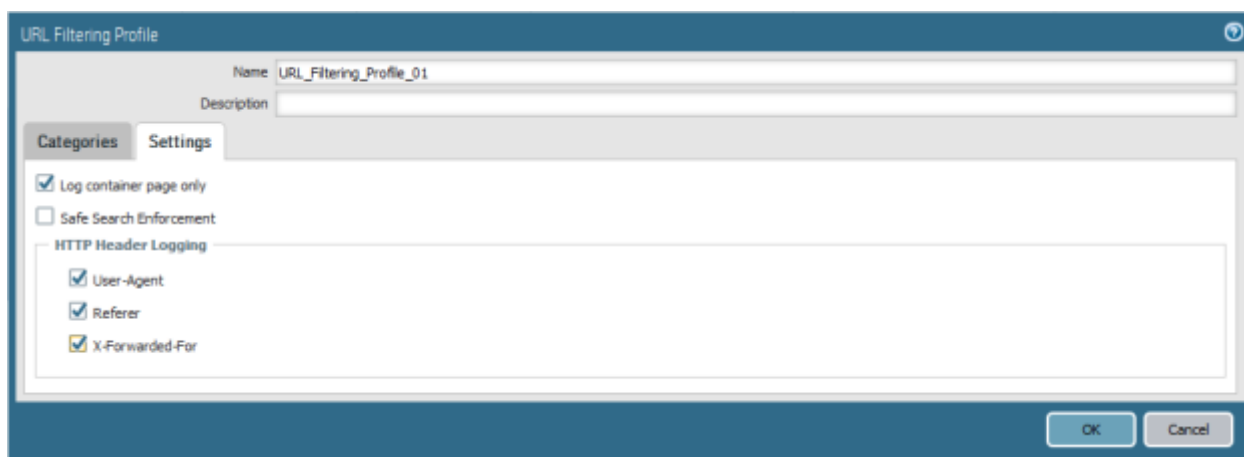
Attribute	Description
User-Agent	The web browser that the user used to access the URL, for example, Internet Explorer. This information is sent in the HTTP request to the server.
Referrer	The URL of the web page that linked the user to another web page; it is the source that redirected (referred) the user to the web page that is being requested.
X-Forwarded-For (XFF)	The option in the HTTP request header field that preserves the IP address of the user who requested the web page. If you have a proxy server on your network, the XFF allows you to identify the IP address of the user who requested the content, instead of only recording the proxy server's IP address as source IP address that requested the web page.

Why This Best Practice Is Important:

Logging these fields gives you improved visibility into web traffic.

How to Implement It:

1. Go to **Objects > Security Profiles > URL Filtering**
2. Edit a URL Filtering Security Profile
3. Go to the Settings tab
4. Check the three HTTP Header Logging options



Logging HTTP header information

Once this URL Filtering Security Profile has been committed, go to **Monitor > Logs > URL Filtering** to see the logged fields.

What Else You Need to Know:

Ensure these settings are enabled in the URL Filtering Security Profile attached to every rule that Allows web traffic.

Security Profile #5: File Blocking

□ Attach a File Blocking Security Profile to Every Allow Rule

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Through the use of a File Blocking Security Profile, PAN-OS allows you to block the transit of specified file types.

Why This Best Practice Is Important:

Some file types are both strongly associated with the transit of malware and unlikely to appear in normal traffic, so blocking them in transit helps prevent infection.

How to Implement It:

1. Go to **Policies > Security**.
2. Edit your Allow rules
3. Attach a File Blocking Security Profile to every Web Browsing Allow rule

What Else You Need to Know:

While this is the best practice, it may not fit your organization's needs. Use your best judgment, and test in a controlled/lab environment first. Otherwise, you may want to set all files that aren't blocked to be Alert, so that they are logged for visibility and monitoring.

Security Policy Rule

General Source User Destination Application Service/URL Category Actions

Action Setting

Action: ☐ Send ICMP Unreachable

Profile Setting

Profile Type:
 Antivirus:
 Vulnerability Protection:
 Anti-Spyware:
 URL Filtering:
 File Blocking:
 Data Filtering:
 WildFire Analysis:

Log Setting

☐ Log at Session Start
☐ Log at Session End
 Log Forwarding:

Other Settings

Schedule:
 QoS Marking:
☐ Disable Server Response Inspection

OK Cancel

Attaching a File Blocking Security Profile

❑ Block the Transit of Windows PE Files

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

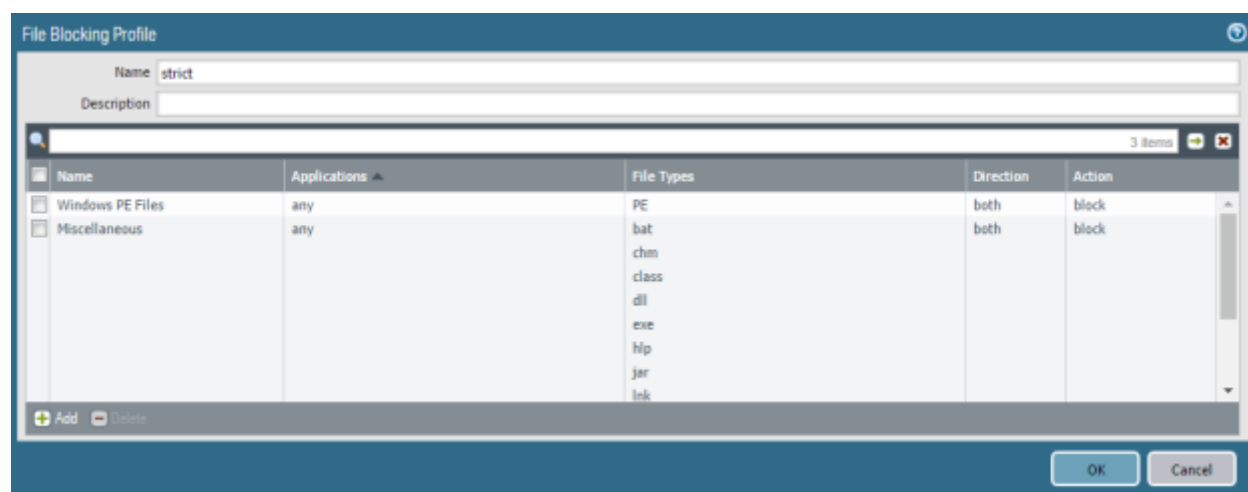
The Portable Executable (PE) format is a file format for executables, object code, DLLs, FON Font files, and others used in Windows operating systems.

Why This Best Practice Is Important:

The transit of Windows PE files is positively associated with the spread of malware and is rarely seen in normal traffic, so blocking them in transit can help prevent infection.

How to Implement It:

Create a rule in your File Blocking Security Profile to block Windows PE files for *any* Application and in *both* Directions.



Blocking Windows PE files

❑ Block the Transit of Common Dangerous and Malicious File Extensions

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Why This Best Practice Is Important:

Some file types are both strongly associated with the transit of malware and unlikely to appear in normal traffic, so blocking their transit helps prevent infection.

How to Implement It:

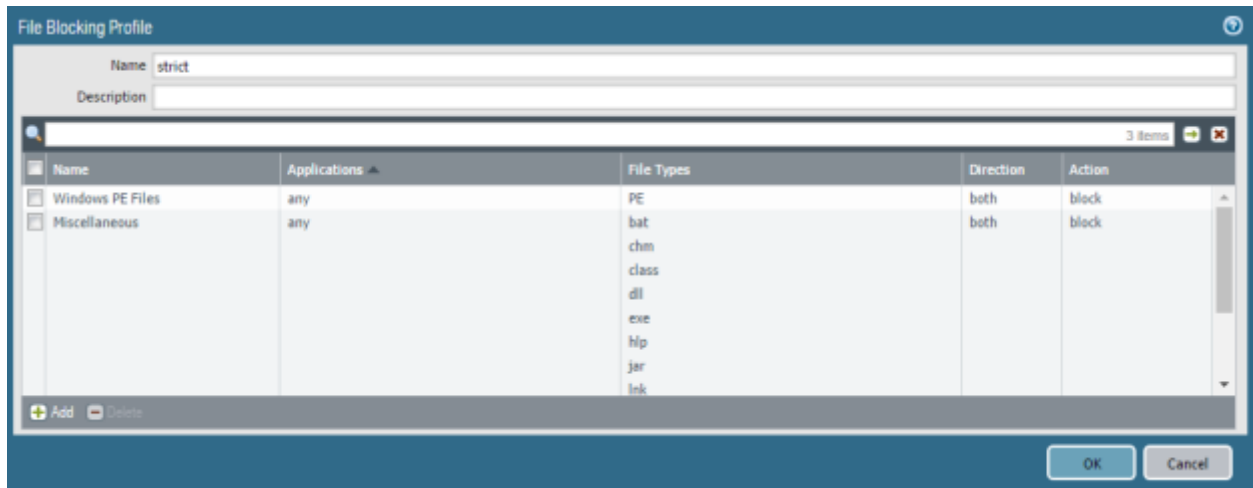
Step 1: Create a list of file types you'd like to block. One reasonable source is at:

<https://www.file-extensions.org/common/dangerous>

Step 2: Take your list of file types you'd like to block and find the intersection of this list with the file types that can be blocked in a File Blocking Security Profile:

- bat
- chm
- class
- dll
- exe
- hlp
- jar
- lnk
- ocx
- scr
- vbe
- wmf

Step 3: Create a rule in your File Blocking Security Profile to block these file types for *any* Application and in *both* Directions.



Blocking common dangerous and malicious file extensions

Security Profile #6: WildFire Analysis

□ Attach a WildFire Analysis Security Profile to Every Accept Rule

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Palo Alto Networks firewalls and Traps endpoints together serve as a million distant early warning WildFire sensors throughout the Internet. By configuring the WildFire Analysis Security Profile, your firewall will forward up to the secure WildFire cloud only those files that have never been seen before and that WildFire determines should be analyzed further.

Why This Best Practice Is Important:

It benefits you and everyone else because if a file which has transited your firewall is malware, WildFire will create a signature and distribute it to everyone with a WildFire subscription within minutes, and everyone else with an Antivirus subscription within a day. With this new signature this piece of malware will never be able to transit a Palo Alto Networks firewall again. We kill malware, *fast*.

How to Implement It:

Step 1:

1. Go to **Objects > Security Profiles > WildFire Analysis**.
2. Create a new WildFire Security Profile or simply review the built-in *default* policy.

Step 2:

1. Go to **Policies > Security Policy**.
2. Edit each of your Allow Security Rules.
3. Attach a WildFire Analysis Security Profile in the Actions tab of the rule.

What Else You Need to Know:

While this is the best practice, it may not fit every organization's needs. Some organizations may choose to attach a profile only to those Policies that show files being transferred.

Security Policy Rule

General Source User Destination Application Service/URL Category Actions

Action Setting

Action:

☐ Send ICMP Unreachable

Profile Setting

Profile Type:

Antivirus:

Vulnerability Protection:

Anti-Spyware:

URL Filtering:

File Blocking:

Data Filtering:

WildFire Analysis:

Log Setting

☐ Log at Session Start

☐ Log at Session End

Log Forwarding:

Other Settings

Schedule:

QoS Marking:

☐ Disable Server Response Inspection

OK Cancel

Ensure you've enabled a WildFire Analysis Security Profile

❑ Maximize the WildFire File Size Limits

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Every Palo Alto Networks firewall has the ability, by using a WildFire Analysis Security Profile, to send suspicious files up to WildFire.

As of PAN-OS 7.1.8, here are the file types and the maximum upload sizes:

File Type	Default Size Limit	Allowable Size Limit Range
PE (Portable Executable)	2 MB	1 Mb – 10 MB
Android APK	10 MB	1 MB – 50 MB
PDF	200 KB	100 KB – 1,000 KB
MS-Office	500 KB	200 Kb – 10,000 KB
Jar	1 MB	1 MB – 10 MB
Flash	5 MB	1 MB – 10 MB
MacOSX	1 MB	1 Mb – 50 MB

Why This Best Practice Is Important:

The more files we send to WildFire, the smarter it gets, and the more malware it catches and provides a signature for in its every-five-minutes updates. All other things being equal, sending more files is better than fewer.

While it's true that the distribution of malware file sizes is skewed toward the smaller sizes, particularly those below the default sizes, malware creators are aware of the file size limits and have an incentive to go just beyond the default sizes.

How to Implement It:

Go to **Device > Setup > WildFire > General Settings** and increase the Size Limits to their maximum values.

What It Looks Like After You've Implemented It:

General Settings

WildFire Public Cloud: wildfire.paloaltonetworks.com

WildFire Private Cloud:

☐ Use Proxy Settings for Private Cloud

File Size Limits

File Type	Size Limit
pe (MB)	10
apk (MB)	50
pdf (KB)	1000
ms-office (KB)	10000
jar (MB)	10
flash (MB)	10
MacOSX (MB)	50

☒ Report Benign Files

☒ Report Grayware Files

OK Cancel

Even larger files now get inspected

❑ Configure WildFire to Report Benign Files

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

One of the possible verdicts returned by WildFire is “Benign”, a rare English word that means “harmless”. By default, files found to be benign are not reported in the **Monitor > WildFire Submissions** log.

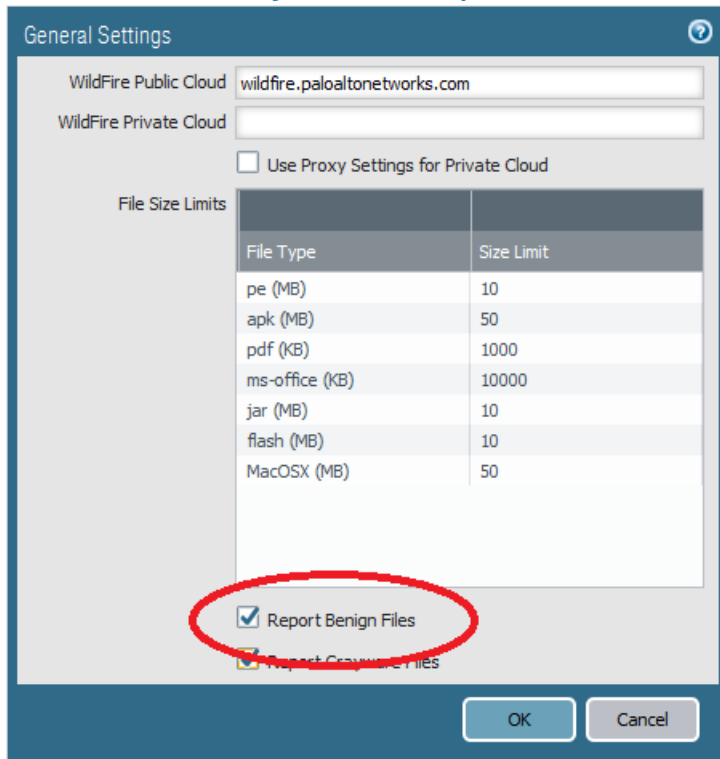
Why This Best Practice Is Important:

The Best Practice is among the less important, but it still might be useful, for the sake of learning or completeness, to see what files were submitted to WildFire and subsequently found to be benign. By enabling this feature, these files will be logged in the **Monitor > WildFire Submissions** log.

How to Implement It:

Go to **Device > Setup > WildFire > General Settings** and enable “Report Benign Files”.

What It Looks Like After You’ve Implemented It:



Benign files will now appear in the **Monitor > WildFire Submissions** log

❑ Configure WildFire to Report Grayware Files

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

One of the possible verdicts returned by WildFire is “Grayware”. Files categorized as grayware do not pose a direct security threat, but might display otherwise obtrusive behavior. Grayware can include, adware, spyware, and Browser Helper Objects (BHOs).

By default, files found to be grayware are not reported in the **Monitor > WildFire Submissions** log.

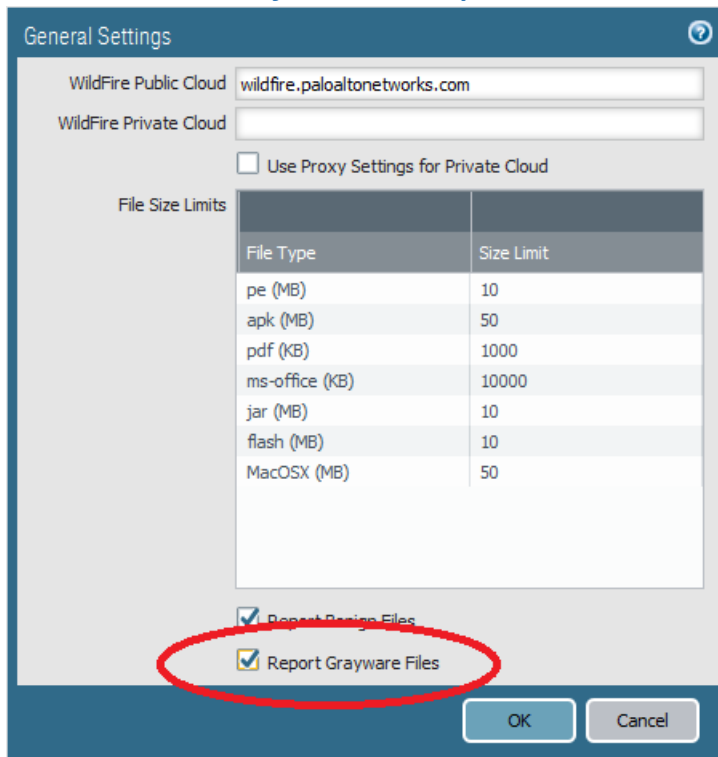
Why This Best Practice Is Important:

The Best Practice is among the less important, but it still might be useful, for the sake of learning or completeness, to see what files were submitted to WildFire and subsequently found to be grayware. By enabling this feature, these files will be logged in the **Monitor > WildFire Submissions** log.

How to Implement It:

Go to **Device > Setup > WildFire > General Settings** and enable “Report Grayware Files”.

What It Looks Like After You’ve Implemented It:



Grayware files will now appear in the **Monitor > WildFire Submissions** log

□ Allow Forwarding of Encrypted Content

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

PAN-OS can decrypt some encrypted traffic. While most traffic transiting the firewall can be forwarded off to an external service, such as WildFire, for security reasons the default behavior for decrypted traffic is that it can't be forwarded off the firewall.

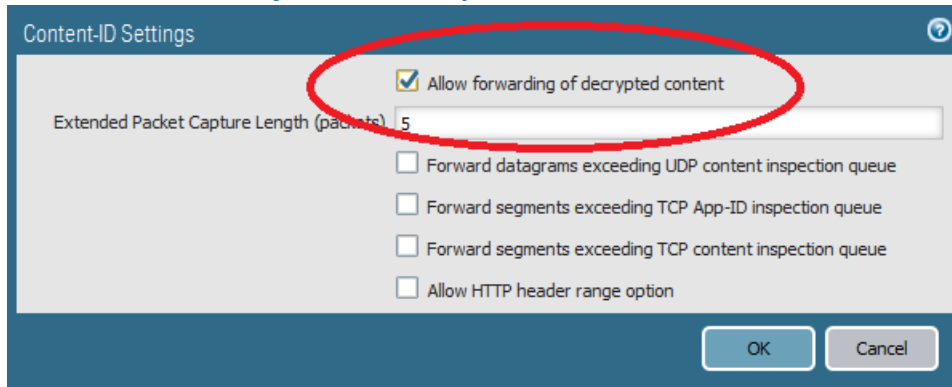
Why This Best Practice Is Important:

Using encryption is a well-known way for malware to tunnel through a firewall and avoid inspection. Once this traffic is encrypted you'll want to, at a minimum, enable it to be sent up to WildFire.

How to Implement It:

Go to **Device > Setup > Content-ID > Content-ID Settings** and check "Allow forwarding of decrypted content".

What It Looks Like After You've Implemented It:



Now WildFire can see what's come through on an encrypted channel

Security Profile #7: Data Filtering

❑ If You Need a Real DLP Solution, Don't Use the Data Filtering Security Profile

Improve Security		Improve Manageability	X
Improve Performance	X	Improve High Availability	

Background Information:

DLP (Data Loss Prevention) software detects potential data breaches/data ex-filtration transmissions and prevents them by monitoring, detecting and blocking sensitive data while in motion. In data leakage incidents, sensitive data is disclosed to unauthorized parties by either malicious intent or an inadvertent mistake. Sensitive data includes private or company information, intellectual property (IP), financial or patient information, credit-card data and other information.

DLP eats CPUs. It's not just the headers of every packet that need to be inspected, but the payloads, also. It's truly deep packet inspection. The process is so CPU intensive, the normal solution is to buy a separate hardware appliance for the job.

Why This Best Practice Is Important:

The PAN-OS Data Filtering Security Profile was not designed to be a full DLP solution. While it may be useful for simple tasks, like filtering for U.S. Social Security Numbers or credit card numbers, it doesn't have the full feature set or full performance capacity of a dedicated DLP solution.

How to Implement It:

If you need a full DLP solution for compliance or other reasons, the Data Filtering Security Profile will not meet your needs.

Security Profile #8: DoS Protection

□ Configure a Strong DoS Protection Security Profile

Improve Security	X	Improve Manageability	X
Improve Performance	X	Improve High Availability	X

Background Information:

The DoS (Denial of Service) Protection Security Profile allows you to configure the protections used in DoS Protection policies.

There are two DoS protection mechanisms that the Palo Alto Networks firewalls support:

- **Flood Protection:** Detects and prevents attacks where the network is flooded with packets resulting in too many half-open sessions or services being unable to respond to each request.
- **Resource Protection:** Detects and prevents session exhaustion attacks. In this type of attack, a large number of bots are used to establish as many fully established sessions as possible to consume all of a system's resources.

You can enable both types of protection mechanisms in a single DoS protection profile.

The DoS profile is used to specify the type of action to take and details on matching criteria for the DoS policy. The DoS profile defines settings for SYN, UDP, and ICMP floods.

Why This Best Practice Is Important:

You need a strong DoS Protection Profile to be able to mitigate DoS attacks.

How to Implement It:

1. Go to **Objects > Security Profiles > DoS Protection**.
2. Create a strong DoS Protection Security Profile

DoS Protection Profile

Name: DoS-Protection-Profile

Description:

Type: ☒ Aggregate ☐ Classified

Flood Protection | Resources Protection

SYN Flood | UDP Flood | ICMP Flood | ICMPv6 Flood | Other IP Flood

☒ **SYN Flood**

Action: Random Early Drop

Alarm Rate (packets/s): 10000

Activate Rate (packets/s): 10000

Max Rate (packets/s): 40000

Block Duration (s): 300

OK Cancel

A DoS Protection Profile showing SYN Flood protections enabled

Evasion Prevention

□ Upgrade to At Least PAN-OS 7.1.1 and Applications and Threats Version 579

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Continuous improvement works, and Palo Alto Networks improves its hardware, software, and services with every release. PAN-OS version 7.1.1 and the Applications and Threats content release Version 579 contain internal changes that improve protections against evasion. Reference the TAC recommended document to determine what Maintenance release should be enabled.

Why This Best Practice Is Important:

Both the attacks and the defenses just keep getting smarter, so if you want to stay defended, you're going to have to keep up with current releases. Getting up to at least these versions is a particularly important milestone.

How to Implement It:

Signatures should be updated and a Saved Named Configuration Snapshot should be taken before upgrading PAN-OS.

To upgrade PAN-OS, go to **Device > Software** and click on "Check Now" to show you what versions are available for downloading and installing.

To upgrade the Applications and Threats content release, go to **Device > Dynamic Updates** and click on "Check Now" to show you what versions are available for downloading and installing.

What Else You Need to Know:

Obviously, there are other implications of upgrading PAN-OS. Some of them are discussed in other Best Practices.

□ Enable Vulnerability Protection Security Profile Evasion Signatures

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Evasions are packets or sessions designed to trick the firewall into matching the wrong rule or coming to the wrong conclusion when evaluating traffic and thus allowing malware to slip through an otherwise properly configured firewall.

Why This Best Practice Is Important:

Evasions are a particularly insidious type of attack. You want to stop them.

How to Implement It:

1. Go to **Objects > Security Profiles > Vulnerability Protection**
2. Create or edit a Vulnerability Protection Security Profile
3. Go to **Exceptions**
4. Search for “evasions” and check the box next to “Show all signatures”. This will show you 56 threats.
5. Click on the Action column header, then click the white triangle drop down activator, then hover over “Update Action” and select “drop”

What Else You Need to Know:

- The built-in *strict* and *default* Vulnerability Protection Security Profiles are read-only, so you’ll have to either clone one of them or create a new Security Profile to be able to make changes.
- Ensure this Vulnerability Protection Security Profile is attached to all Allow rules.

Ensure you’ve configured a DNS Proxy or else you might get some weird false positives in some HA configurations.

□ Enable Anti-Spyware Security Profile Evasion Signatures

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Evasions are packets or sessions designed to trick the firewall into matching the wrong rule or coming to the wrong conclusion when evaluating traffic and thus allowing malware to slip through an otherwise properly configured firewall.

Why This Best Practice Is Important:

Evasions are a particularly insidious type of attack. You want to stop them.

How to Implement It:

1. Go to **Objects > Security Profiles > Anti-Spyware**
2. Create or edit an Anti-Spyware Security Profile
3. Go to **Exceptions**
4. Search for “evasions” and check the box next to “Show all signatures”. This will show you 2 threats.
5. Click on the Action column header, then click the white triangle drop down activator, then hover over “Update Action” and select “drop”

What Else You Need to Know:

- The built-in *strict* and *default* Anti-Spyware Security Profiles are read-only, so you’ll have to either clone one of them or create a new Security Profile to be able to make changes.
- Ensure this Anti-Spyware Security Profile is attached to all Allow rules.

Ensure you’ve configured a DNS Proxy or else you might get some weird false positives in some HA configurations.

❑ Clear the Urgent Data Flag in the TCP Header

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

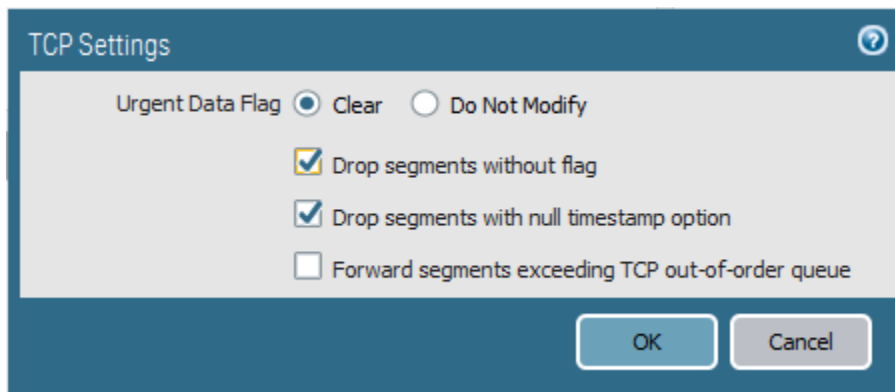
Many hosts use the urgent data flag in the TCP header to promote a packet for immediate processing, removing it from the processing queue and expediting it through the TCP/IP stack. This process is called out-of-band processing. However, the implementation of the urgent data flag varies from host to host.

Why This Best Practice Is Important:

Configuring the firewall to clear this flag eliminates ambiguity in how the packet is processed on the firewall and the host, ensuring the firewall sees the same stream in the protocol stack as the host for which the packet is destined. When the firewall clears this flag, it prevents the packet from being processed urgently.

How to Implement It:

1. Go to **Device > Setup > Session > TCP Settings**.
2. Check the “Clear” radio button next to “Urgent Data Flag”



Clearing the Urgent Data Flag

❑ Drop Segments Without Flags

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

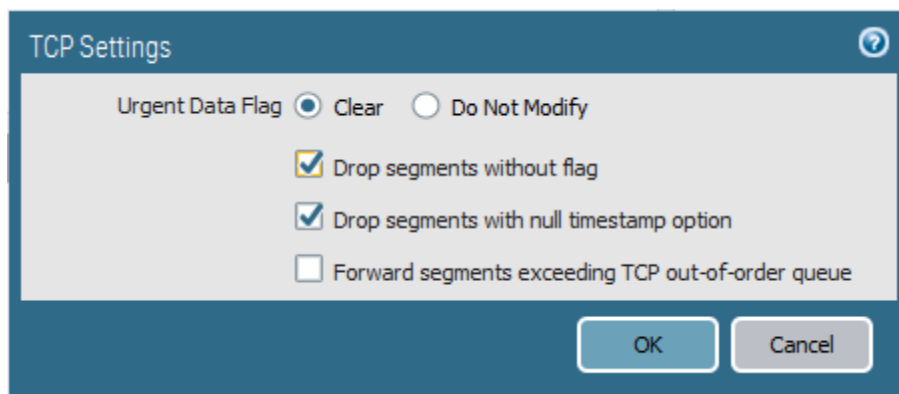
Illegal TCP segments without any flags set can be used to evade content inspection.

Why This Best Practice Is Important:

When you enable this option, the firewall will drop packets that have no flags set in the TCP header.

How to Implement It:

1. Go to **Device > Setup > Session > TCP Settings**.
2. Check the box next to “Drop segments without flag”



Dropping segments without flags

❑ Drop Segments With a Null Timestamp

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The TCP timestamp records when the segment was sent and allows the firewall to verify that the timestamp is valid for that session, preventing TCP sequence number wrapping.

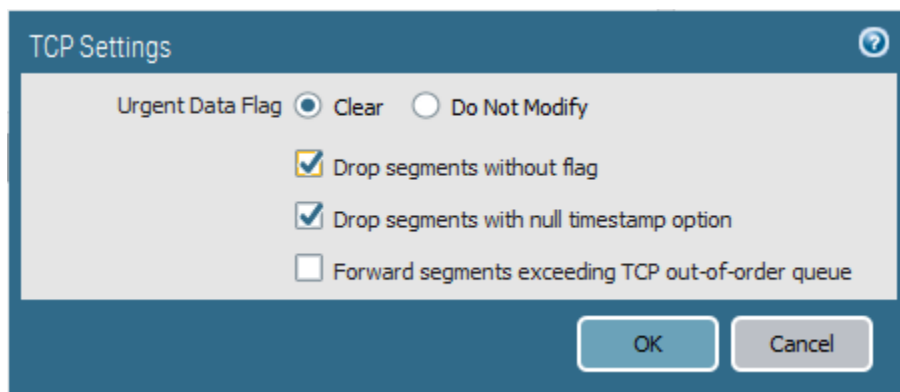
The TCP timestamp is also used to calculate round trip time.

Why This Best Practice Is Important:

When a TCP Timestamp is set to 0 (null) it could confuse either end of the connection, resulting in an evasion. With this setting enabled, the firewall drops packets with null timestamps.

How to Implement It:

1. Go to **Device > Setup > Session > TCP Settings**.
2. Check the box next to “Drop segments with null timestamp option”



Dropping segments with a null timestamp

❑ Don't Forward Segments Exceeding the TCP Out-of-Order Queue

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

By default, the firewall forwards segments that exceed the TCP out-of-order queue limit of 64 per session.

Situations in which a large number of out-of-order TCP packets might arrive include:

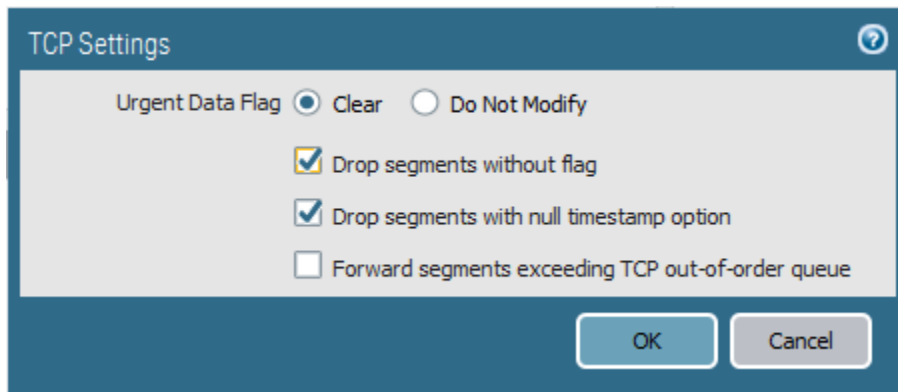
- Misconfigured load-balancers
- Dynamic routing
- A denial-of-service attack
- An evasion attack

Why This Best Practice Is Important:

By disabling this option, the firewall instead drops segments that exceed the out-of-order queue limit.

How to Implement It:

1. Go to **Device > Setup > Session > TCP Settings**.
2. Uncheck the box next to "Forward segments exceeding TCP out-of-order queue".



Not forwarding segments exceeding TCP out-of-order queue

❑ Don't Forward Segments Exceeding the TCP App-ID Inspection Queue

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

By default, when the App-ID inspection queue is full the firewall skips App-ID inspection—classifying the application as unknown-tcp—and forwards the segments.

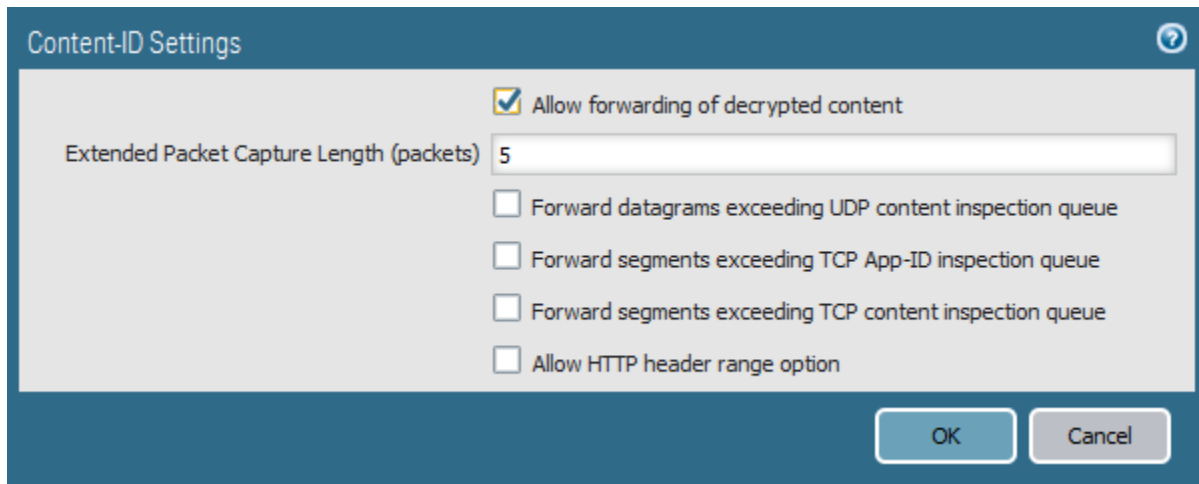
Why This Best Practice Is Important:

By disabling this option, the firewall instead drops segments when the App-ID inspection queue is full.

You don't want to give an attacker who has the ability to flood you with traffic the additional ability to degrade the specificity of your App-ID engine.

How to Implement It:

1. Go to **Device > Setup > Content-ID . Content-ID Settings**.
2. Uncheck the box next to "Forward segments exceeding TCP app-ID inspection queue".



Content-ID Settings

☒ Allow forwarding of decrypted content

Extended Packet Capture Length (packets)

☐ Forward datagrams exceeding UDP content inspection queue

☐ Forward segments exceeding TCP App-ID inspection queue

☐ Forward segments exceeding TCP content inspection queue

☐ Allow HTTP header range option

OK Cancel

Not forwarding segments exceeding the TCP App-ID inspection queue

❑ Don't Forward Datagrams Exceeding the TCP or UDP Content Inspection Queues

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

By default, when the TCP or UDP content inspection queue is full the firewall skips Content-ID inspection for TCP segments or UDP datagrams that exceed the queue limit of 64.

Why This Best Practice Is Important:

By disabling these options, the firewall instead drops TCP segments and UDP datagrams when the corresponding TCP or UDP content inspection queue is full.

You don't want to give an attacker who has the ability to flood you with traffic the additional ability to shut down your Content-ID inspection engine.

How to Implement It:

1. Go to **Device > Setup > Content-ID . Content-ID Settings**.
2. Uncheck the boxes next to "Forward segments exceeding TCP content inspection queue" and "Forward segments exceeding UDP content inspection queue".

Content-ID Settings

☒ Allow forwarding of decrypted content

Extended Packet Capture Length (packets)

☐ Forward datagrams exceeding UDP content inspection queue

☐ Forward segments exceeding TCP App-ID inspection queue

☐ Forward segments exceeding TCP content inspection queue

☐ Allow HTTP header range option

OK Cancel

Not forwarding segments exceeding the TCP content inspection queue or datagrams exceeding the UCP content inspection queue

❑ Don't Allow the HTTP Header Range Option

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The HTTP Range option allows a client to fetch part of a file only.

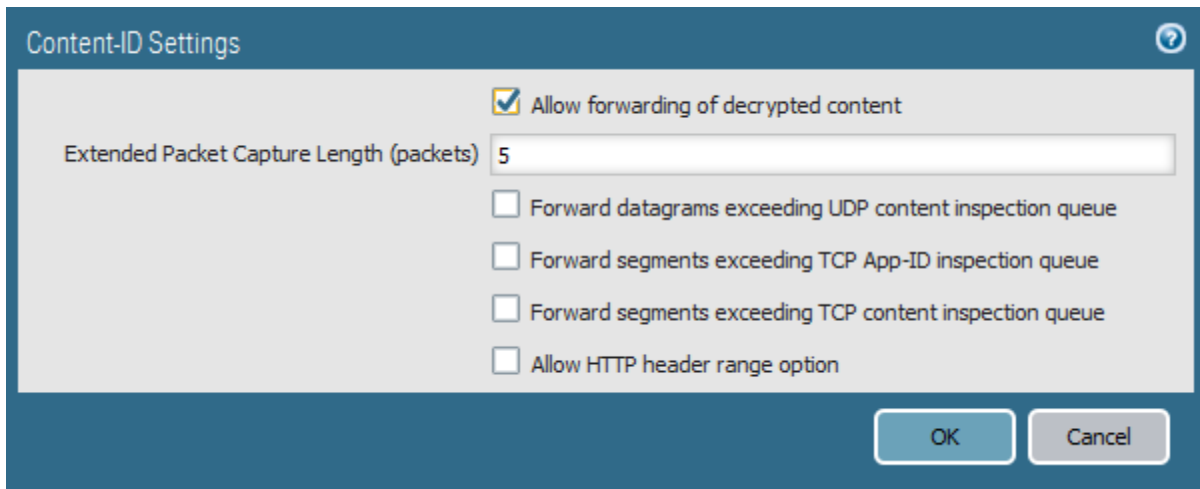
When a next-generation firewall in the path of a transfer identifies and drops a malicious file, it terminates the TCP session with a RST packet. If the web browser implements the HTTP Range option, it can start a new session to fetch only the remaining part of the file. This prevents the firewall from triggering the same signature again due to the lack of context into the initial session, while at the same time allowing the web browser to reassemble the file and deliver the malicious content.

Why This Best Practice Is Important:

Disabling this option prevents clients from fetching part of a file only and possibly avoiding malware detection in this way.

How to Implement It:

1. Go to **Device > Setup > Content-ID . Content-ID Settings**.
2. Uncheck the box next to “Allow HTTP header range option”.



Not allowing the HTTP header range option

What Else You Need to Know:

Keep in mind that disabling this option should not impact device performance; however, HTTP file transfer interruption recovery may be impaired.

In addition, disabling this option could also impact streaming media services, such as Netflix, Windows Server Updates Services (WSUS), and Palo Alto Networks content updates.

Saving NGFW Changes

❑ Avoid Letting Uncommitted Changes Linger

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

When you make a change to the configuration, the change is stored in the volatile *candidate config*. It isn't yet implemented or stored, so you'll need to either do a *Commit*, which will implement it and save it to the *running configuration*, or explicitly save it as the *candidate configuration*. If the firewall is restarted before taking one of these two steps, the changes will be lost.

Why This Best Practice Is Important:

It's easy to forget that you've made changes to the *candidate config*. Because they're not saved, it's easy to lose them, and because they not yet implemented, they can sit there like a land mine, waiting for the next administrator to do a *Commit*.

How to Implement It:

Partition your firewall changes into small groups. Make changes for a specific purpose, then do a *Commit*, then test your changes. Don't leave uncommitted changes hanging.

□ Preview Changes Before Committing

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

All of the customer-changeable configuration information in a Palo Alto Networks firewall is stored in an underlying XML file.

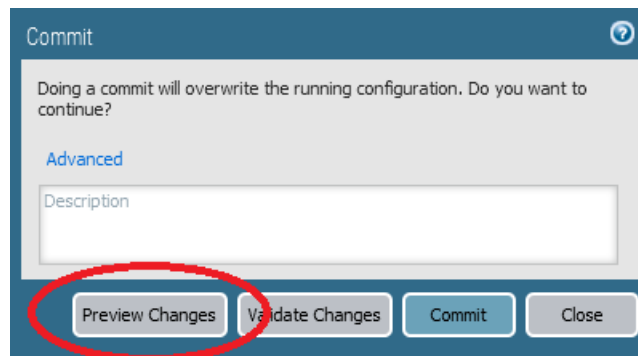
The configuration that is currently running on the firewall is known as the *running configuration*. As you begin to make changes, the *running-configuration* is copied to the *candidate config* and that's what gets edited. When you do a commit, the *candidate-config* is installed and then copied back to the *running-configuration*.

As part of the commit process, you're given the opportunity to preview the changes you've made before committing.

Why This Best Practice Is Important:

An experienced carpenter will tell you: Measure twice, cut once.

How to Implement It:



Here's your chance to look for errors before committing

Device Config Audit (PA-VM) - Mozilla Firefox

https://192.168.4.2/php/device/show-config-diff.php?isGecko=1&width=850&height=500&type=device&filepath=8914655714150957/curly-diff.out

Device Config Audit (PA-VM)

Wed Mar 8 8:49:31 PST 2017

Legend: Added Modified Deleted

Local Device Changes			
Running Configuration		Candidate Configuration	
864	hip-profiles any;	864	hip-profiles any;
865	action drop;	865	action drop;
866	log-end no;	866	log-end no;
867	tag RB_Initial_Drops;	867	tag RB_Initial_Drops;
868	}	868	}
869	"Drop these services always" {	869	"Drop these services alwaysasdsdsw" {
870	to any;	870	to any;
871	from any;	871	from any;
872	source any;	872	source any;
873	destination any;	873	destination any;
874	source-user any;	874	source-user any;

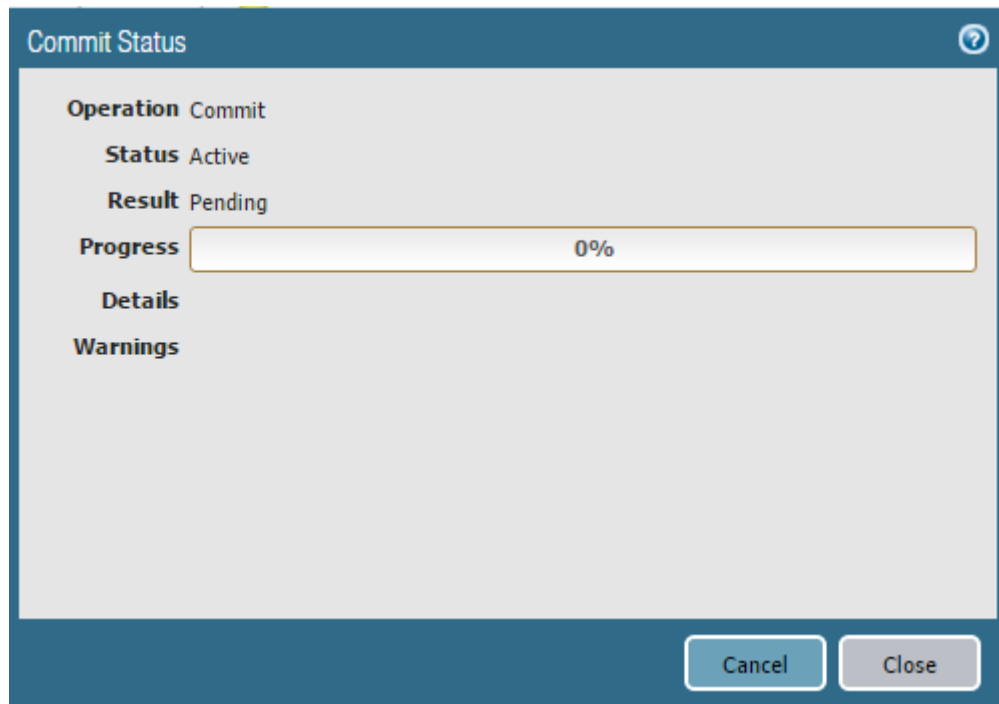
Whoops, it looks like I fat-fingered a security rule name; best to fix it before committing

❑ Check for Warnings After Committing

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Depending upon your hardware and the complexity of your configuration, a *Commit* might take several minutes. You can spend a lot of time starting at this:



Sometimes it takes a while

It's tempting to press the *Close* button and immediately get back to work.

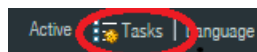
Why This Best Practice Is Important:

If you close out the Commit Status dialog box before the commit is complete, you take on a small risk that the commit won't complete properly and that you haven't been informed. It's important to follow through and verify the commit completed successfully.

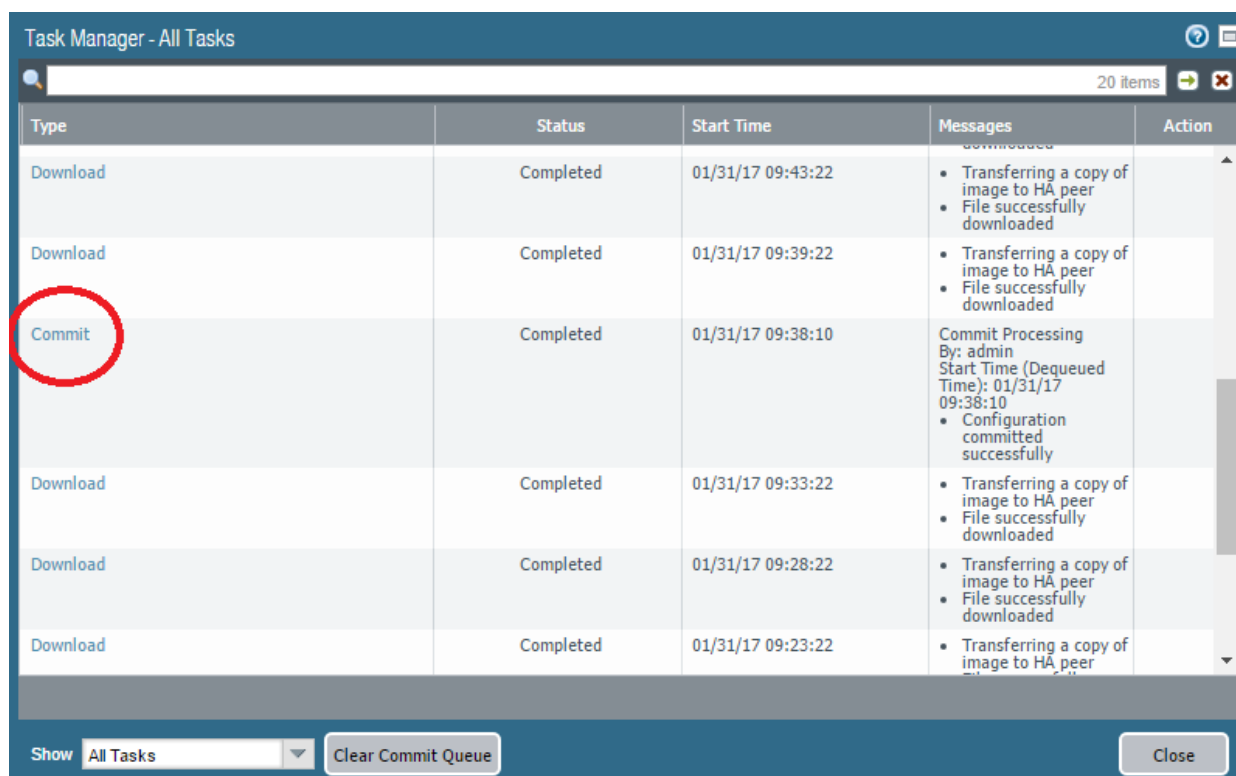
How to Implement It:

If you've closed out the Commit Status dialog box but later want to see if the commit completed successfully, you can always reopen it by through the Task Manager.

The Task Manager is located at the bottom right corner of every page in the GUI:




When you click on it, you get the Task Manager dialog box:



The Task Manager dialog box, showing the recent Commit completed successfully

If you want, you can click on the word Commit in the Task Manager dialog box and restore the original Job Status dialog box:



The Job Status dialog box showing a successful Commit completion

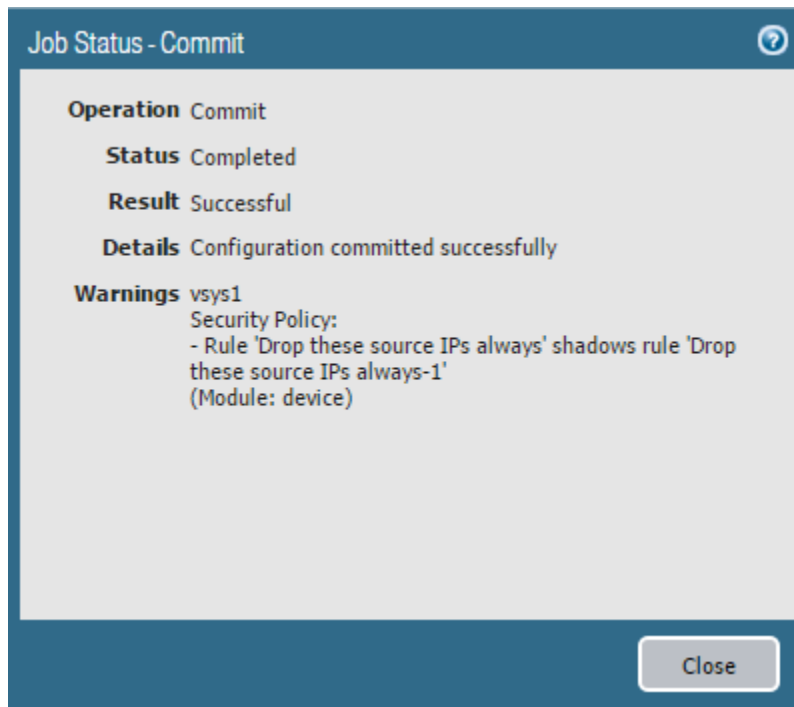
□ Resolve Commit Warnings

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

As part of the Commit process the firewall performs limited additional analysis of your configuration to look for things that are technically correct but indicate some confusion or contradiction in the underlying logic. These errors aren't severe enough to block successful completion of the Commit, but the firewall issues warnings to help keep you informed.

Here is an example of a Commit warning that didn't prevent successful completion of the Commit:



A Commit warning created by absent-mindedly cloning Rule #1

Why This Best Practice Is Important:

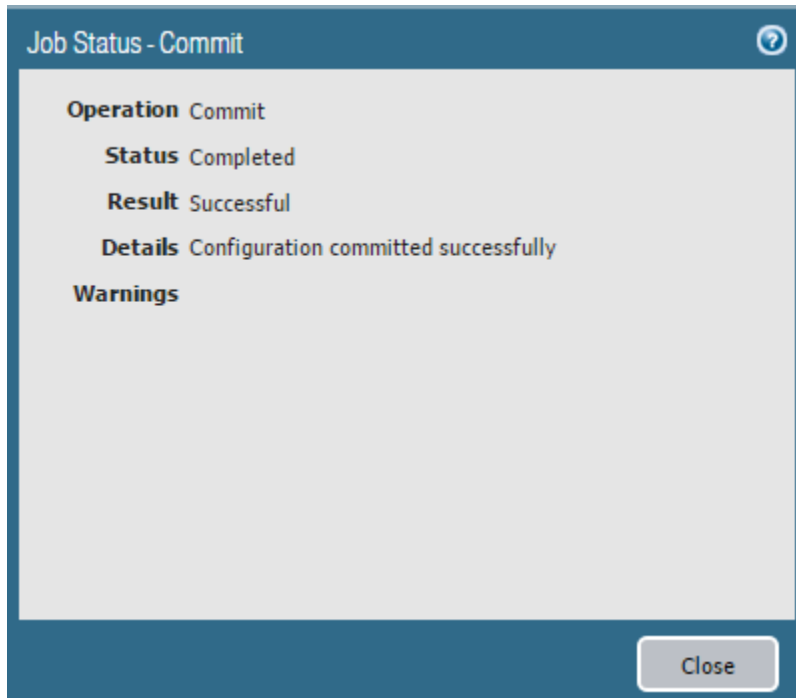
If you understand correctly how the firewall works, and your objects, rules and configuration are logically consistent, you'll never get one of these warnings. While it's true the Commit completed successfully, it's obvious you're misunderstanding something. This is probably an indicator of some other underlying problem. You need to track this down and fix it.

Also, ignored warnings tend to multiply, and once you get more than just a few you won't be able to keep track of them or even notice when new ones are added to your growing list.

How to Implement It:

Check every Commit to ensure it completed without warnings, and resolve every warning as soon as you see it.

What It Looks Like After You've Implemented It:



Everything's looking good again

Dynamic Updates

□ Consider Setting the Threshold Value for Antivirus and Applications and Threats Downloads

Improve Security		Improve Manageability	
Improve Performance		Improve High Availability	X

Background Information:

Palo Alto Networks is constantly updating and publishing new Dynamic Updates and your firewall should be configured to automatically download and install them as they become available.

However, on rare occasions, a Dynamic Update is discovered, typically within a few hours, to have an unexpected issue and a replacement is quickly provided.

For the Antivirus and Applications and Threats downloads, you have the option to set a “Threshold”, or minimum number of hours that the update must age before you’re willing to download and install it.

Why This Best Practice Is Important:

On the one hand...

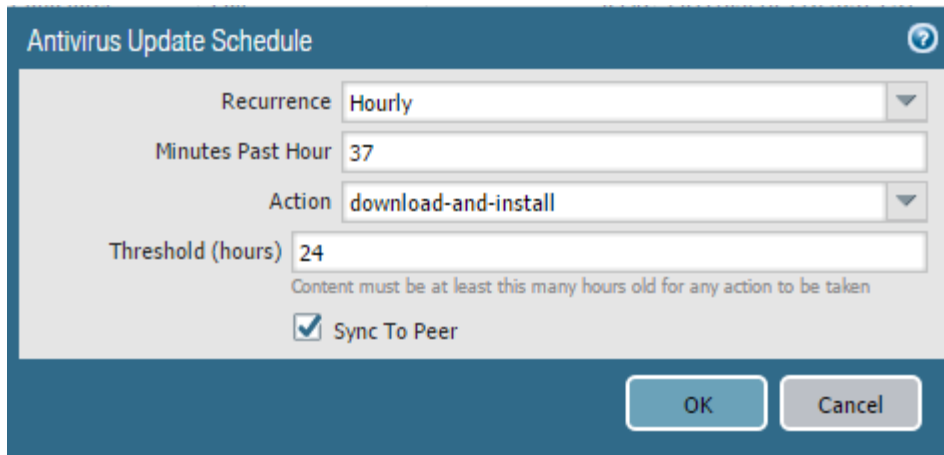
Palo Alto Networks moves *fast* in getting the latest possible updates out to its customers as quickly as possible, so there’s an argument for downloading and installing updates as soon as they’re available.

But on the other hand...

Some particularly cautious customers may wish to let Antivirus and Applications and Threats updates age a few hours before they’re willing to download and install them. The goal is that if there’s a problem, they’d like some other customer to stumble across it first and get it fixed, so they won’t have to.

How to Implement It:

1. Go to **Device > Dynamic Updates**.
2. For the Antivirus and Applications and Threats downloads, click to edit the schedule.
3. Enter the number of hours you wish updates to age before you’re willing to download and install them.



The image shows a dialog box titled "Antivirus Update Schedule". It contains the following fields and controls:

- Recurrence:** A dropdown menu set to "Hourly".
- Minutes Past Hour:** A text input field containing the value "37".
- Action:** A dropdown menu set to "download-and-install".
- Threshold (hours):** A text input field containing the value "24". Below this field is a small note: "Content must be at least this many hours old for any action to be taken".
- Sync To Peer:** A checkbox that is checked.
- Buttons:** "OK" and "Cancel" buttons at the bottom right.

Just a little more time in the oven, please

App-ID

□ Use App-ID Filters in Your Security Policy Rules

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

App-ID is a patented traffic classification system available only in Palo Alto Networks firewalls. It can identify an application regardless of port, protocol or encryption (SSL or SSH), even if the application is evasive. App-ID uses application signatures, application protocol decoding and heuristics to see through all of those possible distractions and accurately identify the application.

App-ID is turned on and classifying traffic all the time, even if you're not using App-ID filters in your Security Policy Rules.

Why This Best Practice Is Important:

Older pre-NGFW firewalls relied on ports and protocols to identify applications and decide which security rule best matched a given proposed session. In response to this, more applications are now being written to use different ports or to be evasive in an attempt to trick firewalls into matching the wrong rule and incorrectly allowing transit. Also, with the rise of SSL encryption, it's harder than ever to see into a session.

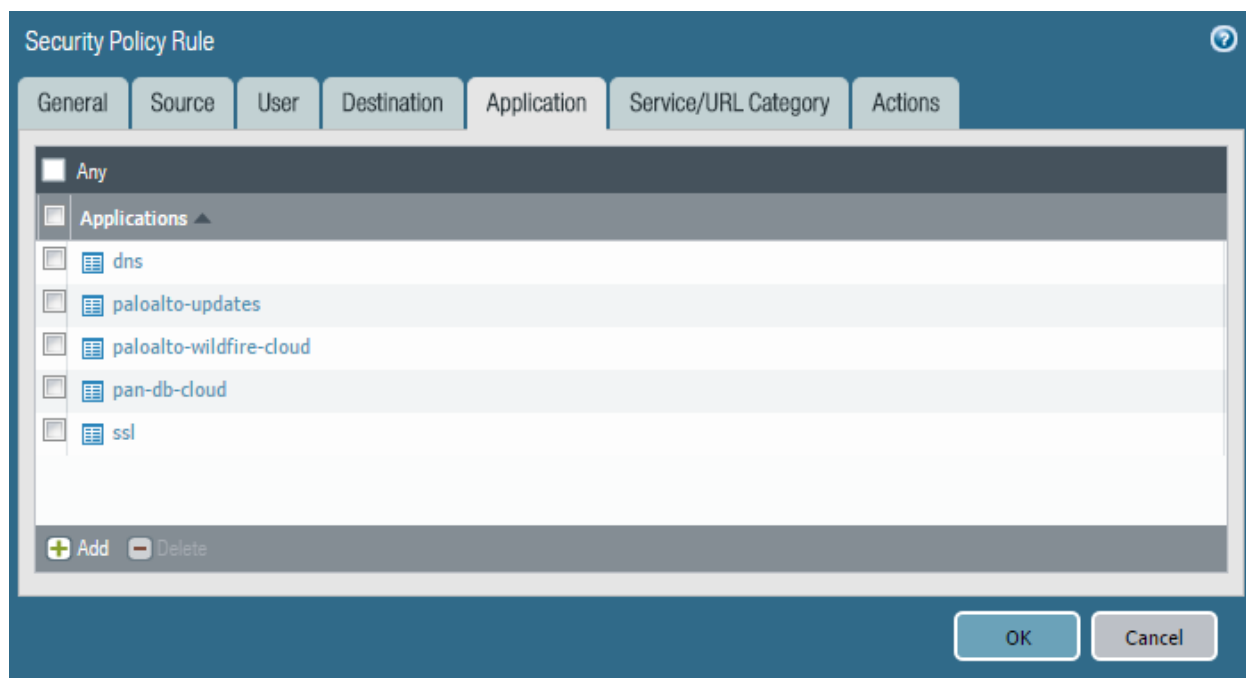
App-ID solves these problems by both using sophisticated mechanisms to identify applications and using decryption to decode SSL and SSH connections.

The benefit is that firewall administrators can now create Security Policy Rules based on Application, not just on port and protocol. This gives you much more specificity in your rules and lets you get a handle on evading applications.

Using App-ID filters in your Security Policy Rules is one of the most important advantages of using a Palo Alto Networks firewall.

How to Implement It:

Because App-ID is turned on and classifying your applications all the time, you don't need to enable it; you just need to start using Application filters in your rules:



Don't click the 'Any' box; explicitly list permitted applications instead

What Else You Need to Know:

The complete list of applications recognized by App-ID is at <https://applipedia.paloaltonetworks.com/>. Included with each is a lot of additional information about each application, including a description, dependencies and characteristics.

❑ Blacklist the *unknown-tcp*, *unknown-udp*, and *unknown-p2p* Applications

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

App-ID understands and can identify some 2,300 applications. Applications it can't identify fall into one of three categories:

1. Rare, highly specialized applications, used in, say the health care or oil & gas industries, that App-ID doesn't know about yet
2. In-house custom applications
3. Custom applications used by attackers to evade and tunnel through firewalls and distribute malware

Why This Best Practice Is Important:

For that first group, it's best to tell Palo Alto Networks about these applications so they can develop signatures to get them recognized. The first thing they'll ask you to provide is a pcap.

If you are seeing unknown traffic for a commercial application that does not yet have an App-ID, you can submit a request for a new App-ID here: <http://researchcenter.paloaltonetworks.com/submit-an-application/>.

For the second group, it's best to develop a custom signature (another Best Practice).

For the third group, it's best to drop these proposed new sessions entirely.

For these reasons, it's best to drop and log the *unknown-tcp*, *unknown-udp*, and *unknown-p2p* applications, both because you'll be dropping the third group and because you'll be forcing yourself to deal properly with the first two groups. These three "applications" are the catch-all designators for everything App-ID can't identify.

How to Implement It:

Either create an explicit rule that drops and logs these three applications, or create a rule that explicitly drops and logs a group of blacklisted applications, and include these three in that group.

What Else You Need to Know:

This may take some tuning to get right. Having Palo Alto Networks create signatures for the first group of applications may take some time, as might creating custom signatures for the second group. Therefore, you might have to start off just logging these applications and working through them one by one.

Also, this level of restriction doesn't work for every organization. Remember to test this approach in a lab environment first, and if it's not a good fit for you, look at setting your policies to log traffic, but have "Allow" as your starting point.

□ Create App-ID Signatures for Custom Applications

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

App-ID understands some 2,300 applications and your firewall automatically gets signature updates weekly, but it obviously can't identify and track custom in-house applications that it's never seen before.

Why This Best Practice Is Important:

Using App-ID to specify applications in your Security Policy Rulebase is incredibly powerful, but you need to be able to disambiguate between unknown applications that you want and those that you don't. For applications that App-ID doesn't have a built-in signature for, you'll do best by creating a custom application signature. This will allow you to continue to drop all *unknown-tcp*, *unknown-udp*, and *unknown-p2p* applications (see another Best Practice).

How to Implement It:

The process is slightly complicated, so it's best to follow these instructions:

<https://www.paloaltonetworks.com/documentation/71/pan-os/pan-os/app-id/create-a-custom-application>

If you need help, Palo Alto Networks has Professional Services Engineers who are experienced at doing this.

User-ID

□ Enable User-ID in Your Security Rules

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

When User-ID buys awesomesauce, it has to order it in the jumbo economy size.

When organizations want to get visibility and control into their network traffic, they want to assign permissions not to IP addresses, but to individual users and groups of users. With Network Address Translation (NAT), wireless networks, desktop virtualization, and mobile devices, there's no longer a 1:1 correspondence between a user and an IP address.

User-ID enables you to identify users on your network using a variety of methods to ensure that you can identify users in all locations using a variety of network connections and operating systems, including Microsoft Windows, Apple iOS, Mac OS, Android, and Linux/UNIX.

PAN-OS provides fourteen methods for establishing a user-to-IP mapping, and the Security Policy Rulebase allows you to specify which users and groups of users will "match" a specific rule.

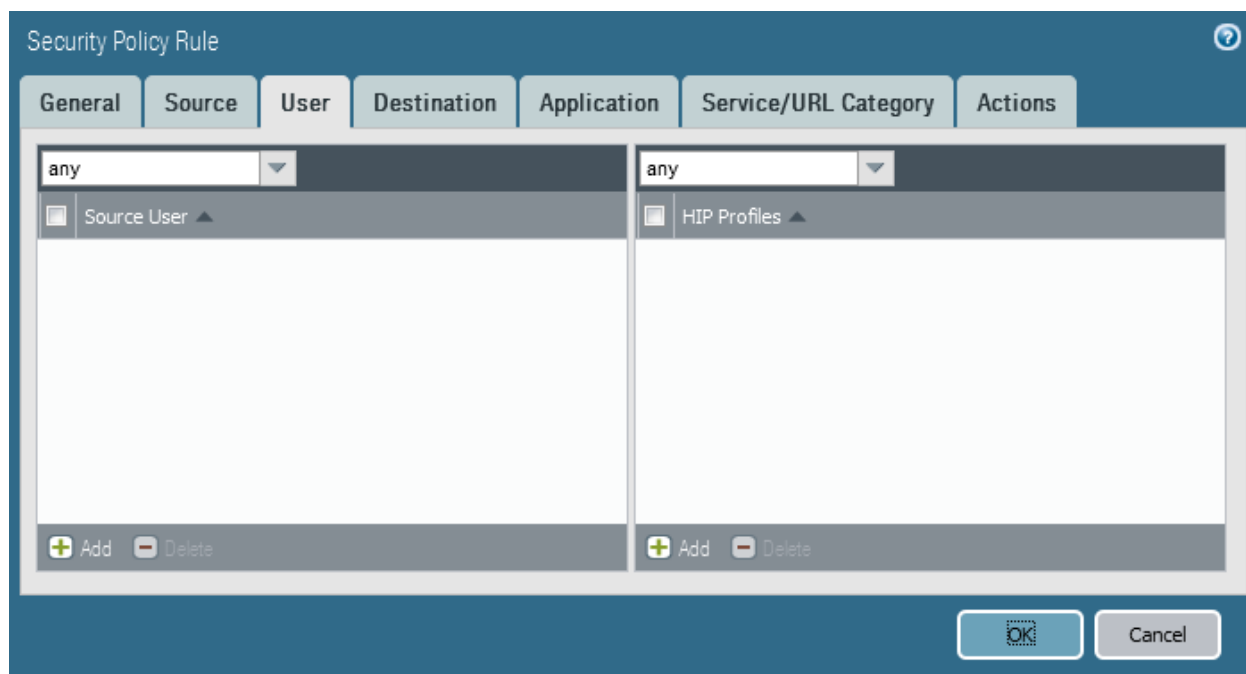
Why This Best Practice Is Important:

User-ID provides two large benefits:

- Organizations no longer have to rely on IP addresses to disambiguate users on their network. With User-ID, specific policies follow users wherever they go, on any device and using any OS. This is a large improvement over a traditional port-and-protocol firewall.
- Each logging event is already tagged with the specific user, making tracking and remediation significantly easier.

How to Implement It:

Edit a Security Policy Rule and click on the User tab:



The User tab in the Security Policy Rule dialog box

Here is where you can specify exactly which users or group of users will match a rule.

What Else You Need to Know:

Implementing the methods that create the User-to-IP mapping is a more complicated process and may require help from Professional Services, but AD integration and GP (Global Protect) work especially well.

❑ Enable User Identification in Trusted Security Zones

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

By default, the methods that PAN-OS uses to create User-to-IP mapping are disabled in every Security Zone. You have to specifically enable it in each zone.

Why This Best Practice Is Important:

In order to get visibility into users in a given Security Zone, you need to enable User Identification in that zone.

How to Implement It:

Go to **Network > Zones** and edit an internal Security Zone:

The screenshot shows the 'Zone' configuration window in the Palo Alto Networks management interface. The 'Name' field is 'SZ_Trust' and the 'Type' is 'Layer3'. Under the 'Interfaces' section, 'ethernet1/2' is listed. The 'Zone Protection Profile' is set to 'BJS-Zone-Protection-Profile'. The 'Log Setting' is set to 'None'. The 'Enable User Identification' checkbox is checked and circled in red. The 'User Identification ACL' section is visible on the right, showing 'Include List' and 'Exclude List' options.

Enabling User Identification in an internal Security Zone

❑ Don't Enable User Identification in Non-Trusted Security Zones Unless You're Using a Captive Portal

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

By default, the methods that PAN-OS uses to create User-to-IP mapping are disabled in every Security Zone. You have to specifically enable it in each zone, but sometimes new firewall administrators accidentally enable User Identification in untrusted Security Zones. It's a common rookie mistake with bad consequences.

Why This Best Practice Is Important:

Enabling User Identification in untrusted Security Zones without using a Captive Portal is a bad idea for two reasons:

- There are no methods available for determining user identification on random Internet visitors.
- Some of the methods for determining user identification (such as WMI Probing) involve probing connections sent out from the firewall that could leak information or potentially open a return communications path.

How to Implement It:

Be careful when enabling User Identification. Only enable it for trusted internal Security Zones, and never for untrusted Security Zones unless you're using a Captive Portal.

DNS

❑ Configure the Firewall to Be a DNS Proxy

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

PAN-OS allows you to create an easy-to-configure DNS Proxy that can provide additional features, performance and security for your DNS lookups. Of particular importance is that a DNS Proxy enables evasion signatures to detect specially crafted HTTP or TLS requests in which a client connects to a domain other than the domain specified in the original DNS query.

Evasions are packets or sessions designed to trick the firewall into matching the wrong rule or coming to the wrong conclusion when evaluating traffic and thus allowing malware to slip through an otherwise properly configured firewall.

Why This Best Practice Is Important:

Evasions are a particularly insidious type of attack. You want to stop them.

By configuring a DNS Proxy, additional anti-evasion signatures become effective.

How to Implement It:

Go to **Network > DNS Proxy > Add**

The screenshot shows the 'DNS Proxy' configuration window in the Palo Alto Networks management interface. The window is titled 'DNS Proxy' and has a search icon in the top right corner. The configuration fields are as follows:

- Enable:** Checked
- Name:** DNS_Proxy_01
- Inheritance Source:** None
- Check inheritance source status:** A link to check the status.
- Primary:** 8.8.8.8
- Secondary:** None

On the right side, there is a list of interfaces with 'ethernet1/1' selected. Below the configuration fields are three tabs: 'DNS Proxy Rules', 'Static Entries', and 'Advanced'. The 'DNS Proxy Rules' tab is active, showing a table with the following columns: Name, Cacheable, Domain Name, Primary, and Secondary. The table is currently empty. At the bottom right of the window are 'OK' and 'Cancel' buttons.

Configuring a DNS Proxy

Once you've configured the DNS Proxy and committed your security policy, you need to reconfigure internal hosts to use the IP address of the firewall's interface as their DNS server.

The Anti-Spyware Profile will now be able to enforce more evasion signatures. This will be discussed in a different Best Practice.

What Else You Need to Know:

There are several other possible benefits to using a DNS Proxy. It may improve performance because it's a local caching resolver, you can add static entries, and you can make rules to direct different queries to different DNS servers. The main purpose of this Best Practice is to get the security benefits of letting the firewall more closely inspect DNS traffic.

□ Configure a DNS Sinkhole

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Malware on infected hosts needs to communicate with Command-and-Control servers out on the Internet, but PAN-DB contains a list of known malware domain names. When a DNS Sinkhole is configured in an Anti-Spyware Profile, DNS lookups from internal hosts, *including from internal DNS servers*, for these known malware domain names will be sent a forged reply by the firewall. The forged reply will be a Sinkhole IP address. If the requestor was an internal DNS server, then this server will now resolve the domain name to the forged Sinkhole IP address.

In this way, the infected hosts, *even if their DNS lookup session to an internal DNS server did not transit the firewall*, will all go to the same forged Sinkhole IP address.

All that's left for the firewall Administrator to do is look in the logs to see which hosts are trying to reach the forged Sinkhole IP address and it will be clear which hosts are infected.

As a courtesy, and perhaps in anticipation of some possible future product enhancement, Palo Alto Networks provides an otherwise unused public IP address you can point to: **71.19.152.112**.

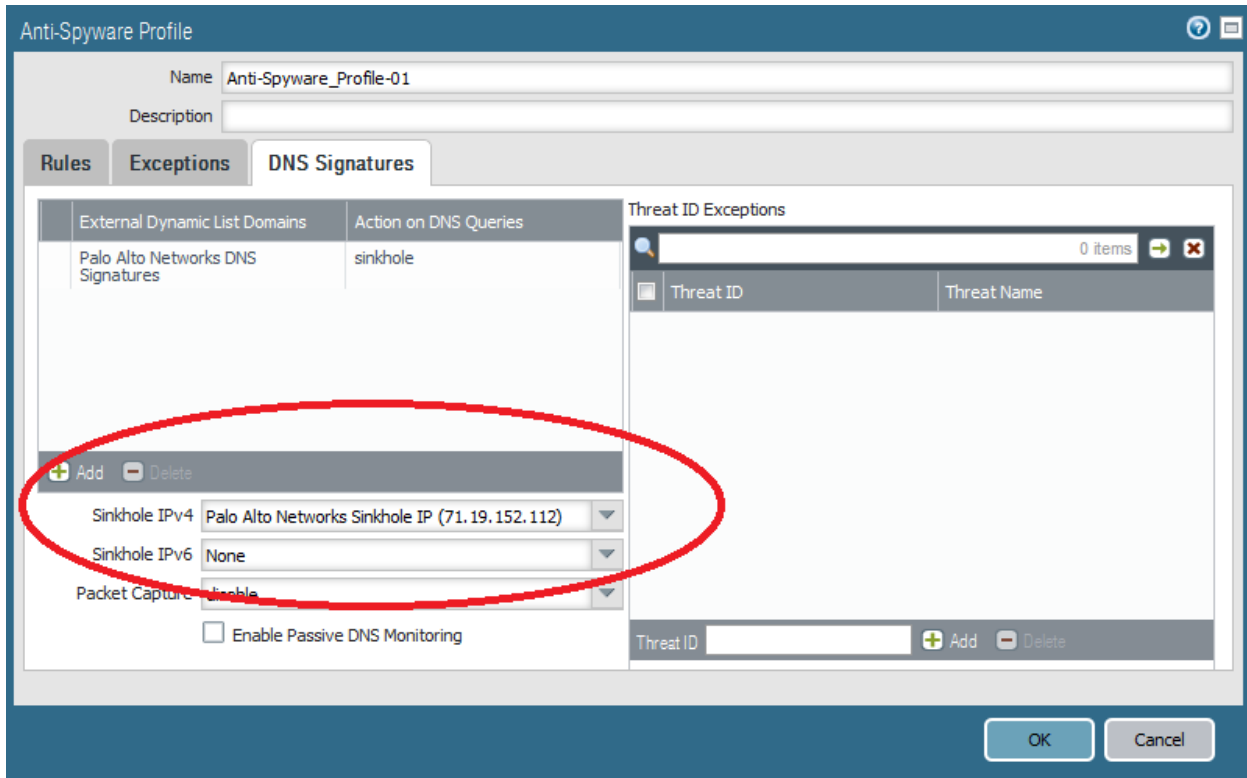
Why This Best Practice Is Important:

Malware on infected hosts needs to communicate with Command-and-Control servers out on the Internet, and a DNS Sinkhole not only disrupts that communication, but also uses the event of the DNS lookup itself to create a beacon that notifies the administrator of the infected host.

How to Implement It:

Go to **Objects > Security Profiles > Anti-Spyware**.

Edit your Anti-Spyware Security Profile:



This Anti-Spyware Profile is configured to use the special Palo Alto Networks provided IP address, but you can use any IP address

Then go to the Monitor tab and start seeing which internal hosts are trying to connect to that IP address.

What Else You Need to Know:

- The two default built-in Anti-Spyware Security Profiles, *default* and *strict*, have DNS Sinkholing disabled and because they're Read-Only, it cannot be enabled. You need to create a new Anti-Spyware Security Profile to enable this feature.

Be sure to attach this Anti-Spyware Security Profile to all rules that permit DNS lookups.

□ Enable Passive DNS Collection for Improved Threat Intelligence

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Passive DNS is an opt-in feature in PAN-OS that enables the firewall to act as a passive DNS sensor and send select DNS information to Palo Alto Networks for analysis in order to improve threat intelligence and threat prevention capabilities.

Data submitted via the Passive DNS Monitoring feature consists solely of mappings of domain names to IP addresses. Palo Alto Networks retains no record of the source of this data and does not have the ability to associate it with the submitter at a future date.

The Palo Alto Networks threat research team uses this information to gain insight into malware propagation and evasion techniques that abuse the DNS system. Information gathered through this data collection is used to improve accuracy and malware detection abilities within PAN-DB URL filtering, DNS-based command-and-control signatures, and WildFire.

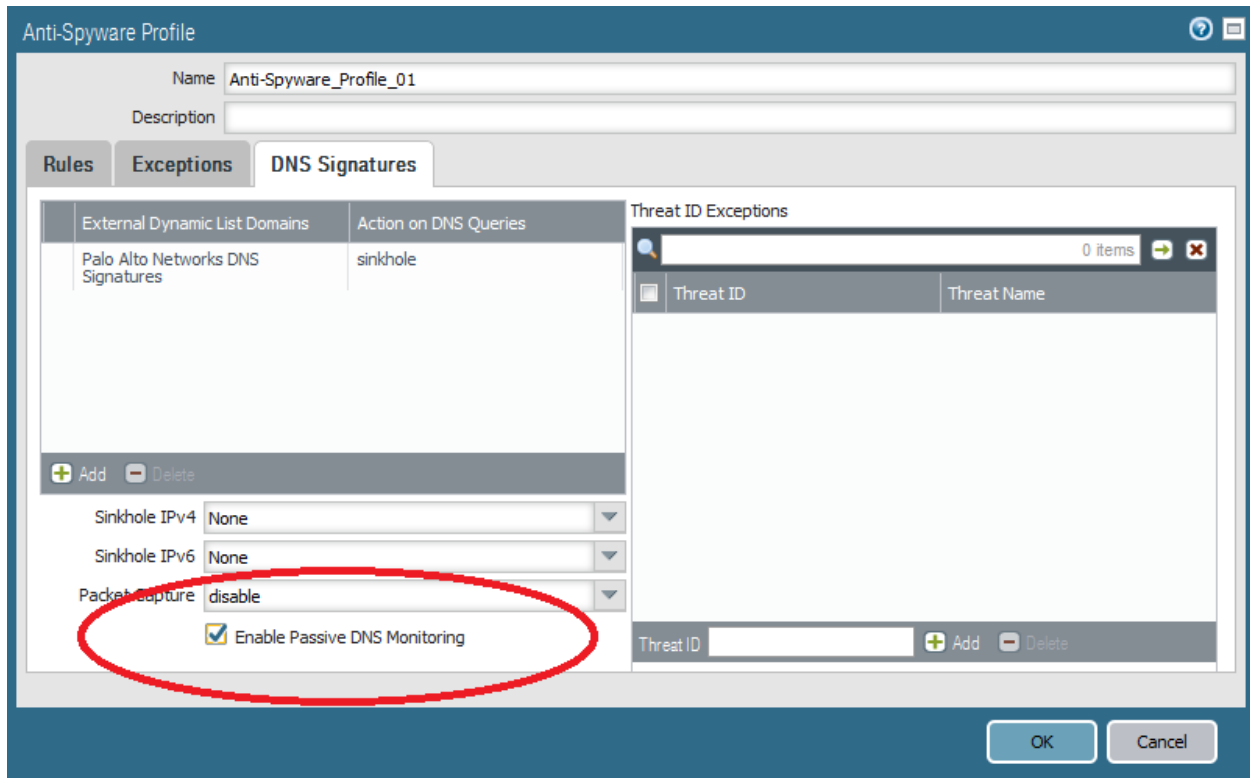
Why This Best Practice Is Important:

It's just good netiquette. Enabling Passive DNS Collection won't immediately improve *your* security, but with the patient accumulation of the information it gathers from thousands of firewalls, it helps improve the security of *every* Palo Alto Networks customer. Join us; you want to be on this team.

How to Implement It:

Go to **Objects > Security Profiles > Anti-Spyware**.

Edit your Anti-Spyware Security Profile:



Now you're helping every Palo Alto Networks customer

What Else You Need to Know:

- The two default built-in Anti-Spyware Security Profiles, *default* and *strict*, have Passive DNS Monitoring disabled and because they're Read-Only, it cannot be enabled. You need to create a new Anti-Spyware Security Profile to enable this feature.

Be sure to attach this Anti-Spyware Security Profile to all rules that permit DNS lookups.

Dynamic Routing

☐ Make Use of Static Routes

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

Background Information:

Dynamic Routing, using protocols such as RIP, BGP, OSPF, IGRP, EIGRP, etc., allows routers to automatically change routing instructions based on link and node availability and link capacity. It can provide high availability to complex networks and reduce transient outages when making network changes.

Why This Best Practice Is Important:

Despite its many advantages, dynamic routing can introduce routing uncertainty into network flows, and this can greatly complicate firewall troubleshooting. Firewall administrator graybeards know in their bones that an awful lot of firewall problems are really routing problems in disguise. Tracking a packet through a firewall is complicated enough without introducing uncertainty about whether the packets are even transiting the firewall at all.

The first step in troubleshooting is to isolate the problem. If you can't even be sure which path the packets are taking, it's much harder to know where to begin.

How to Implement It:

To reduce complexity and uncertainty, consider static routing for traffic passing through your firewall.

What Else You Need to Know:

If you want to try to implement this, the people in your networking team might be happy about it.

IPv6

☐ Enable IPv6 Firewalling

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

IPv6 is the newer, better version of the IP protocol, providing lots of new functionality in addition to exponentially expanding the IP address space. It's catching on, although more quickly outside of the U.S.

Processing packets in IPv6 requires a completely different network stack from IPv4, and this setting in PAN-OS is the master switch that enables all IPv6 functionality.

Why This Best Practice Is Important:

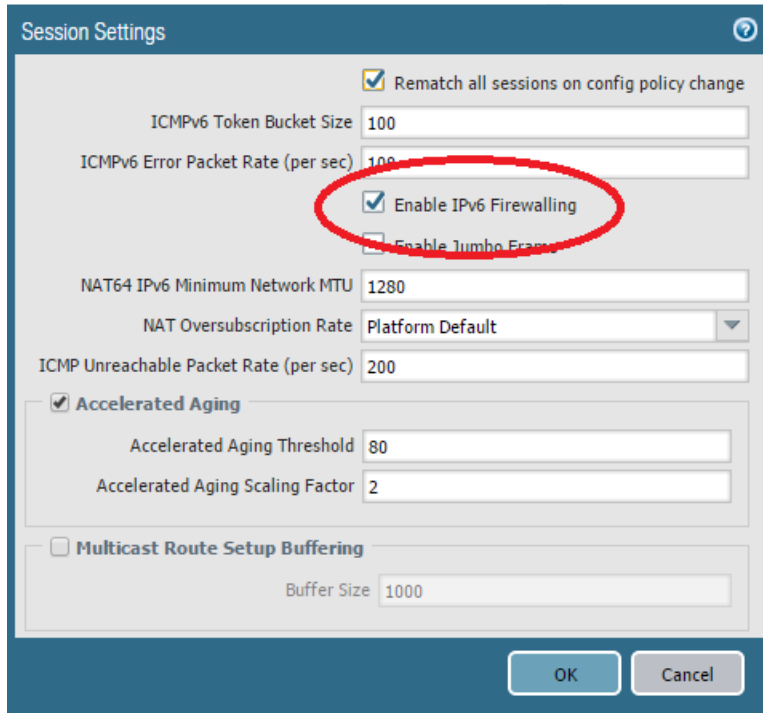
You need to be able to see and control IPv6 traffic in your firewall, so you need to enable IPv6 firewalling. Note: this is only applicable if you are already running IPv6 in your network.

Even if IPv6 is enabled in an Interface (another Best Practice), this option must be enabled for the IPv6 functionality to work.

How to Implement It:

Go to **Device > Setup > Session > Session Settings**.

Check the box next to “Enable IPv6 Firewalling”.



The screenshot shows the 'Session Settings' configuration window. At the top, there is a checkbox labeled 'Rematch all sessions on config policy change' which is checked. Below this are input fields for 'ICMPv6 Token Bucket Size' (100) and 'ICMPv6 Error Packet Rate (per sec)' (100). The 'Enable IPv6 Firewalling' checkbox is checked and highlighted with a red circle. Below it is a checkbox for 'Enable Jumbo Frames' which is unchecked. Further down are input fields for 'NAT64 IPv6 Minimum Network MTU' (1280) and a dropdown for 'NAT Oversubscription Rate' (Platform Default). Below that is an input field for 'ICMP Unreachable Packet Rate (per sec)' (200). There are two expandable sections: 'Accelerated Aging' (checked) with sub-fields for 'Accelerated Aging Threshold' (80) and 'Accelerated Aging Scaling Factor' (2); and 'Multicast Route Setup Buffering' (unchecked) with a sub-field for 'Buffer Size' (1000). At the bottom right are 'OK' and 'Cancel' buttons.

Now you’ve got visibility and control of IPv6 traffic

Virtual Private Networks (VPNs)

❑ IKE/IPSec Crypto Profile: Configure Strong Authentication

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

In the IKE and IPSec Crypto Profiles, here are your choices for the hash algorithm used for authentication:

Hash Algorithm	NIST Rating
md5	Below 80-bit Strength
sha1	128-bit Strength
sha256	256-bit Strength
sha384	Over 256-bit Strength
sha512	Over 256-bit Strength

Why This Best Practice Is Important:

The hash algorithm is used to detect forgery or corruption in the data stream. All other things being equal – and they are except for perhaps a very small performance penalty for using a “longer” hash algorithm – using a stronger algorithm provides better security.

How to Implement It:

You don’t want to use *md5* or *sha1* because they both have known weaknesses and their bit lengths are too short. It’s best to limit yourself to the three strongest algorithms.

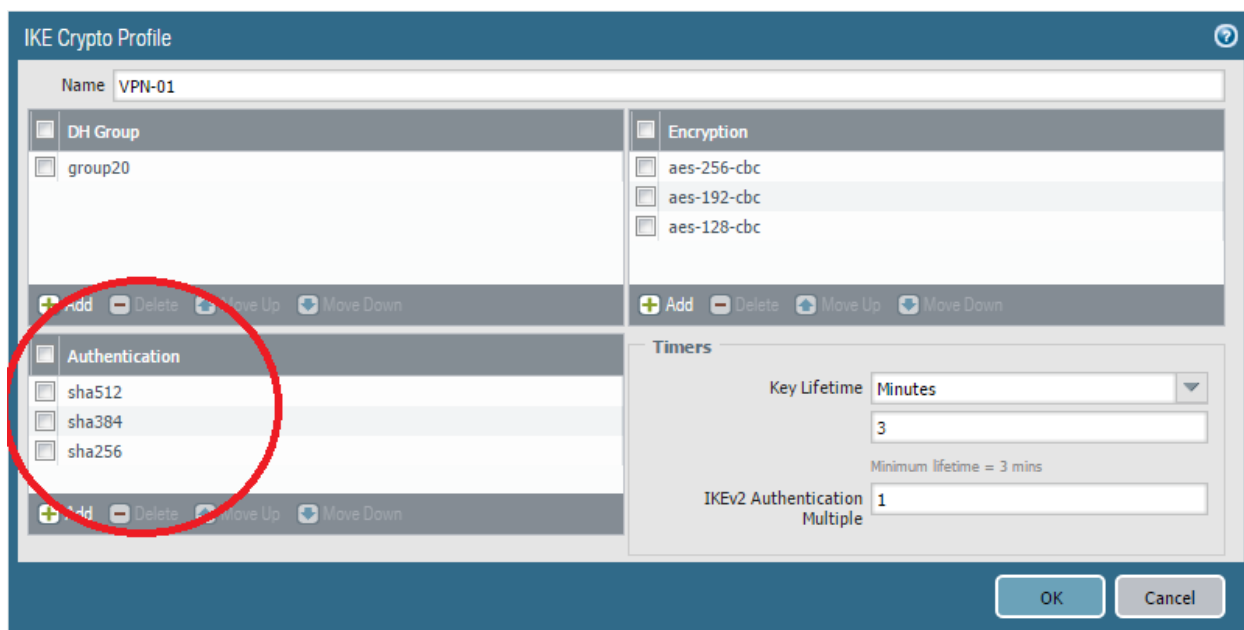
It’s also best to sort your choices, in order from strongest at the top to weakest at the bottom so you’ll always be trying to get the strongest possible algorithm during the negotiation with the remote peer.

Therefore, your choices should be these three, in this order:

- sha512
- sha384
- sha256

To configure your choices:

1. Go to **Network > Network Profiles > IKE Crypto**, and **Network > Network Profiles > IPSec Crypto**.
2. Edit or create an IKE or IPSec Crypto Profile.
3. Configure the Authentication settings.



Choosing the three strongest algorithms, in order from strongest at the top to least strong at the bottom

❑ IKE Crypto Profile: Configure Strong Encryption

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

In the IKE Crypto Profiles, here are your choices for the encryption algorithm:

Encryption Algorithm	Strength
des	Data Encryption Standard (DES) with a security strength of 56 bits
3des	Triple Data Encryption Standard (3DES) with a security strength of 112 bits
aes-128-cbc	Advanced Encryption Standard (AES) using cipher block chaining (CBC) with a security strength of 128 bits
aes-192-cbc	AES using CBC with a security strength of 192 bits
aes-256-cbc	AES using CBC with a security strength of 192 bits

Why This Best Practice Is Important:

This symmetric key encryption algorithm is what keeps your data secret. All other things being equal, a longer key length provides better security. Remember that your enemies can always crack your encryption with a brute force attack if they have enough time and computers, so your job is to bury them in extra work, and each additional bit in the key length doubles the expected value of work required to brute force a password.

How to Implement It:

You don't want to use *des* or *3des* because the key lengths are too short. It's best to limit yourself to the three algorithms with the longest key lengths.

There's also a very good argument to be made, using Information Theory, that, properly configured, *128-bits is always enough* (the summary of the argument is that, even if computers and electricity were free, and even at 1 Kelvin where the thermodynamic operations are most efficient, the energy required to simply iterate through 2^{128} passwords, even without testing them, would require more energy than needed to boil the oceans), but there still may be undetected weaknesses in the current algorithms, so going for the longest key lengths is still best.

It's also best to sort your choices, in order from strongest at the top to weakest at the bottom so you'll always be trying to get the strongest possible algorithm during the negotiation with the remote peer.

Therefore, your choices should be these three, in this order:

- aes-256-cbc
- aes-192-cbc

- aes-128-cbc

To configure your choices:

1. Go to **Network > Network Profiles > IKE Crypto**.
2. Edit or create an IKE Crypto Profile.
3. Configure the Encryption settings.

The screenshot shows the 'IKE Crypto Profile' configuration window for 'VPN-01'. The 'Encryption' section is highlighted with a red circle, showing three selected algorithms: 'aes-256-cbc', 'aes-192-cbc', and 'aes-128-cbc'. The 'Authentication' section shows three selected algorithms: 'sha512', 'sha384', and 'sha256'. The 'Timers' section shows 'Key Lifetime' set to 'Minutes' with a value of '3' and 'Ikev2 Authentication Multiple' set to '1'. The 'OK' and 'Cancel' buttons are at the bottom right.

Choosing the three strongest algorithms, in order from strongest at the top to least strong at the bottom

❑ IKE Crypto Profile: Configure Strong DH Groups

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The Diffie-Hellman Key Exchange (or the “Diffie-Hellman-Merkle Key Exchange”, to be more fair, look it up) allows two hosts to generate a shared secret without sharing a secret. It’s an important step in building a VPN tunnel across an untrusted link.

The strength of the algorithm relies in both the choice of the underlying algorithm (Modular Exponential or Elliptic Curve) and the key length. Elliptic Curve is stronger than Modular Exponential and longer key bit lengths are stronger than shorter key bit lengths.

These are the DH groups available in PAN-OS IKE Crypto Profiles:

DH Group	Key Length
group1	768-bit Modular Exponential (MODP) algorithm
group2	1024-bit Modular Exponential (MODP) algorithm
group5	1536-bit Modular Exponential (MODP) algorithm
group14	2048-bit Modular Exponential (MODP) algorithm
group19	256-bit elliptic curve algorithm
group20	384-bit elliptic curve algorithm

Why This Best Practice Is Important:

Because your IKE Phase I is built on the DH Key Exchange, and all IPSec Phase II session keys are exchanged after being encrypted with these keys, it’s important to use a strong group.

How to Implement It:

Groups 1, 2, and 5 have keys that are too short, so you should limit the choices to Groups 14, 19, and 20. Given that the negotiation with the remote peer always starts with the top choice, you should order your configuration as follows, from strongest at the top, to weakest at the bottom:

1. *group20*
2. *group19*
3. *group14*

To configure your IKE Crypto Profile:

1. Go to **Network > Network Profiles > IKE Crypto**.
2. Edit or create an IKE Crypto Profile.
3. Edit the DH Group settings.

IKE Crypto Profile

Name: VPN 01

DH Group

- ☒ group20
- ☒ group19
- ☒ group14

Encryption

- ☒ aes-256-cbc
- ☒ aes-192-cbc
- ☒ aes-128-cbc

Authentication

- ☒ sha512
- ☒ sha384
- ☒ sha256

Timers

Key Lifetime: Minutes
3
Minimum lifetime = 3 mins

Ikev2 Authentication Multiple: 1

OK Cancel

Configuring for the three strongest DH groups, in order from strongest at the top to weakest at the bottom

❑ IKE Crypto Profile: Configure Short Key Lifetimes

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Remember that our enemies can always guess (brute force) our passwords if they have enough time or computers, so our job is to bury them in extra work.

Why This Best Practice Is Important:

One way to bury your enemies in extra work is to keep changing your keys so the value of a compromised key is limited in scope.

How to Implement It:

The default setting is eight hours. The minimum setting is three minutes, so use three minutes. There are nation-states out there with Application Specific Integrated Circuit (ASIC) fabrication plants and multi-billion dollar budgets, with data centers that can hold a million custom-designed ASICs that do nothing all day but brute force attack passwords. Let's run up their electric bills.

1. Go to **Network > Network Profiles > IKE Crypto**.
2. Edit or create an IKE Crypto Profile.
3. Edit the Key Lifetime setting in the Timers section.

What It Looks Like After You've Implemented It:

The screenshot shows the 'IKE Crypto Profile' configuration window. The 'Name' field is 'VPN-01'. The 'DH Group' section has 'group20' selected. The 'Authentication' section has 'sha512', 'sha384', and 'sha256' selected. The 'Encryption' section has 'aes-256-cbc', 'aes-192-cbc', and 'aes-128-cbc' selected. The 'Timers' section is highlighted with a red circle, showing 'Key Lifetime' set to 'Minutes' with a value of '3'. Below this, it says 'Minimum lifetime = 3 mins'. The 'IKv2 Authentication' section is partially visible with a value of '1'. The 'OK' and 'Cancel' buttons are at the bottom right.

We're throwing out our IKE keys every three minutes, the smallest time interval allowed

□ IKE Crypto Profile: Configure a Low IKEv2 Authentication Multiple

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

When a VPN tunnel is first created, not only are the peers establishing an IKE SA key, but they're also authenticating to each other. As time goes on, when the key lifetime expires, the peers replace the key, but they don't automatically reauthenticate.

The IKEv2 Authentication Multiple setting allows you to specify how many times the IKE SA keys can be replaced before the peers have to also reauthenticate. The permitted values are 0-50. A setting of 0 means "Don't reauthenticate". The default is 0.

In IKEv2, the Initiator and Responder gateways have their own key lifetime value, and the gateway with the shorter key lifetime is the one that will request that the SA be re-keyed.

Why This Best Practice Is Important:

If something goes wrong with one of the peers, or with the connection, and it would be caught by a reauthentication attempt, you're going to want to find out about this as early as possible. Therefore, set this value to 1, which means that every time you throw out the IKE SA keys, while you're at it, force the peers to authenticate again. If somebody finds a way to successfully pull off a Man-In-The-Middle attack or impersonate a peer, you want to find out about it early and fail the tunnel closed quickly.

How to Implement It:

1. Go to **Network > Network Profiles > IKE Crypto**.
2. Edit or create an IKE Crypto Profile.
3. In the Timers section, enter the number 1 in the IKEv2 Authentication Multiple field.

What It Looks Like After You've Implemented It:

The screenshot shows the 'IKE Crypto Profile' configuration window for 'VPN-01'. The window is divided into several sections:

- Name:** VPN-01
- DH Group:** A list containing 'group20', 'group19', and 'group14'. Below the list are buttons for '+ Add', '- Delete', '+ Move Up', and '+ Move Down'.
- Encryption:** A list containing 'aes-256-cbc', 'aes-192-cbc', and 'aes-128-cbc'. Below the list are buttons for '+ Add', '- Delete', '+ Move Up', and '+ Move Down'.
- Authentication:** A list containing 'sha512', 'sha384', and 'sha256'. Below the list are buttons for '+ Add', '- Delete', '+ Move Up', and '+ Move Down'.
- Timers:** A section containing:
 - Key Lifetime:** A dropdown menu set to 'Minutes' and a text input field containing '3'. Below this is the text 'Minimum lifetime = 3 mins'.
 - Ikev2 Authentication Multiple:** A text input field containing '1'. This field is circled in red.

At the bottom right of the window are 'OK' and 'Cancel' buttons.

Papers, please! Now every time the IKE SA keys are replaced, the peers will also have to reauthenticate.

❑ IPsec Crypto Profile: Configure Strong Encryption

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

In the IPsec Crypto Profiles, here are your choices for the encryption algorithm:

Encryption Algorithm	Strength
des	Data Encryption Standard (DES) with a security strength of 56 bits
3des	Triple Data Encryption Standard (3DES) with a security strength of 112 bits
aes-128-cbc	Advanced Encryption Standard (AES) using cipher block chaining (CBC) with a security strength of 128 bits
aes-192-cbc	AES using CBC with a security strength of 192 bits
aes-256-cbc	AES using CBC with a security strength of 256 bits
aes-128-gcm	AES using Galois/Counter Mode (GCM) with a security strength of 128 bits
aes-256-gcm	AES using GCM with a security strength of 256 bits
aes-128-ccm	AES using Counter with CBC-MAC (CCM) with a security strength of 128 bits (Not available on VM-Series firewalls)

Why This Best Practice Is Important:

This symmetric key encryption algorithm is what keeps your data secret. All other things being equal, a longer key length provides better security. Remember that your enemies can always crack your encryption with a brute force attack if they have enough time and computers, so your job is to bury them in extra work, and each additional bit in the key length doubles the expected value of work required to brute force a password.

How to Implement It:

You don't want to use *des* or *3des* because the key lengths are too short. It's best to limit yourself to the six algorithms with the longest key lengths.

There's also a very good argument to be made, using Information Theory, that, properly configured, *128-bits is always enough*, but there still may be undetected weaknesses in the current algorithms, so going for the longest key lengths is still best.

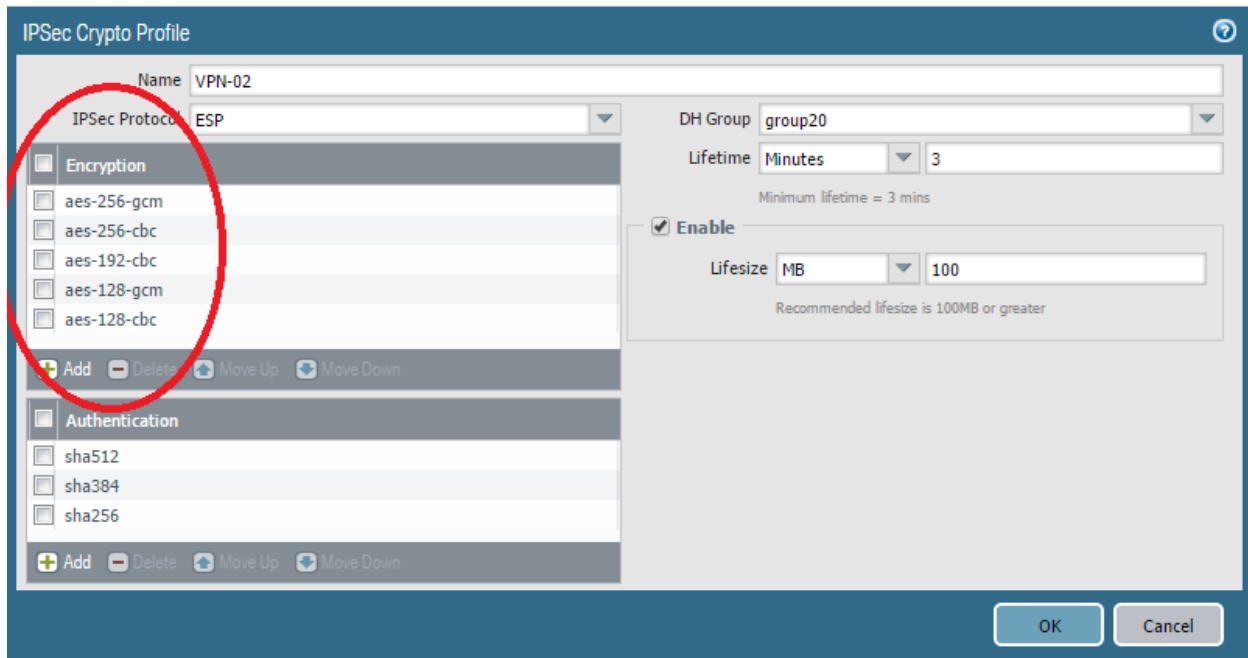
It's also best to sort your choices, in order from strongest at the top to weakest at the bottom so you'll always be trying to get the strongest possible algorithm during the negotiation with the remote peer.

Therefore, your choices should be these six, in this order:

- aes-256-gcm
- aes-256-cbc
- aes-192-cbc
- aes-128-gcm
- aes-128-ccm (Not available on VM-Series firewalls)
- aes-128-cbc

To configure your choices:

1. Go to **Network > Network Profiles > IPSec Crypto**.
2. Edit or create an IPSec Crypto Profile.
3. Configure the Encryption settings.



Choosing the six strongest algorithms (except for aes-128-ccm, which isn't available on this VM-Series firewall), in order from strongest at the top to least strong at the bottom

□ IPSec Crypto Profile: Prefer ESP Over AH

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

IPSec can be configured to use one of two protocols: *ESP* and *AH*. Here are the differences between them:

	ESP	AH
Encrypts the data	X	
Authenticates the source	X	X
Verifies data integrity	X	X

Obviously, the difference is that AH doesn't encrypt the data, which would seem to be a requirement for a VPN.

The situations where AH would be preferred are rare. They include:

- When you're not permitted to use encryption, such as in amateur radio data links on the licensed spectrum, or in some political jurisdictions
- When you want the connection to be inspected, by, say, a firewall on the route
- When you want to be able to compress the packets, without some later encryption

Why This Best Practice Is Important:

Unless you have some highly unusual situation, you'll want to encrypt the data transiting your VPN.

How to Implement It:

1. Go to **Network > Network Profiles > IPSec Crypto**.
2. Edit or create an IPSec Crypto Profile.
3. Configure the IPSec Protocol dropdown box to "ESP".

What It Looks Like After You've Implemented It:

IPSec Crypto Profile

Name: VPN-02

IPSec Protocol: ESP

DH Group: group2

Lifetime: Hours 1

Minimum lifetime = 3 mins

☒ Enable

Lifesize: MB [1 - 65535]

Recommended lifesize is 100MB or greater

+ Add - Delete Move Up Move Down

+ Add - Delete Move Up Move Down

OK Cancel

This VPN will now also encrypt the traffic

❑ IPSec Crypto Profile: Configure Strong DH/PFS

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

VPNs are typically configured so that the DH keys have longer lifetimes than the IPSec session keys that are encrypted and exchanged with those DH keys.

However, every time the IPSec keys are renewed, *a little bit of information* about your underlying DH keys leaks out to your enemies.

Why This Best Practice Is Important:

To improve security, we want to make sure that every time this information leaks – that is, every time the IPSec session keys are changed – we also change out the underlying DH keys. Doing this puts up into a desirable state known as Perfect Forward Secrecy (PFS).

How to Implement It:

1. Go to **Network > Network Profiles > IPSec Crypto**.
2. Edit or create an IPSec Crypto Profile.
3. For DH Group, select *group20*, the most secure group.
4. For Lifetime, select 3 minutes, the shortest lifetime possible.

What It Looks Like After You've Implemented It:

The screenshot shows the 'IPSec Crypto Profile' configuration window. The 'Name' field is 'VPN-02' and the 'IPSec Protocol' is 'ESP'. The 'DH Group' is set to 'group20' and the 'Lifetime' is set to '3' minutes. A red circle highlights the 'DH Group' and 'Lifetime' fields. The 'Enable' checkbox is unchecked. The 'Lifsize' field is set to 'MB' with a range of '[1 - 65535]'. The interface includes 'Add', 'Delete', 'Move Up', and 'Move Down' buttons for the 'Encryption' and 'Authentication' sections.

We have achieved Perfect Forward Secrecy nirvana

❑ IPSec Crypto Profile: Enable Lifesize Limiting

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Remember that our enemies can always guess (brute force) our passwords if they have enough time or computers, so our job is to bury them in extra work.

One way to bury your enemies in extra work is to keep changing your keys so the value of a compromised key is limited in scope.

Also, when doing cryptographic analysis, as the size of the data set that was encrypted with the same key increases, your analysis gains small statistical advantages, so that's another reason to cycle out the keys regularly.

Why This Best Practice Is Important:

For both of those reasons, you can increase your security by limiting how much data gets encrypted with a single IPSec session key.

How to Implement It:

1. Go to **Network > Network Profiles > IPSec Crypto**.
2. Edit or create an IPSec Crypto Profile.
3. Check the Enable box and set the Lifesize values to 100 MB.

What It Looks Like After You've Implemented It:

The screenshot shows the 'IPSec Crypto Profile' configuration window. The 'Name' field is 'VPN-02' and 'IPSec Protocol' is 'ESP'. Under 'Encryption', several algorithms are listed. Under 'Authentication', 'sha512', 'sha384', and 'sha256' are listed. On the right, 'DH Group' is 'group20' and 'Lifesize' is set to '100' MB. The 'Enable' checkbox is checked. A red circle highlights the 'Enable' checkbox and the 'Lifesize' settings. The 'OK' and 'Cancel' buttons are at the bottom right.

After every 100 MB, we throw away the keys and create new ones

□ IKE Gateway: Prefer IKEv2 Over IKEv1

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

The Internet Engineering Task Force (IETF) originally defined IKE in November 1998 in a series of publications (Request for Comments) known as RFC 2407, RFC 2408 and RFC 2409:

- RFC 2407 defined The Internet IP Security Domain of Interpretation for ISAKMP
- RFC 2408 Internet Security Association and Key Management Protocol (ISAKMP)
- RFC 2409 defined The Internet Key Exchange (IKE)

RFC 4306 updated IKE to version two (IKEv2) in December 2005.

RFC 4718 clarified some open details in October 2006.

RFC 5996 combined these two documents plus additional clarifications into the updated IKEv2, published in September 2010. A later update upgraded the document from Proposed Standard to Internet Standard, published as RFC 7296 in October 2014.

IKEv2 provides the following benefits over IKEv1:

- Tunnel endpoints exchange fewer messages to establish a tunnel. IKEv2 uses four messages; IKEv1 uses either nine messages (in main mode) or six messages (in aggressive mode).
- Built-in NAT-T (NAT Traversal) functionality improves compatibility between vendors.
- Built-in health check automatically re-establishes a tunnel if it goes down. The liveness check replaces the Dead Peer Detection used in IKEv1.
- Supports traffic selectors (one per exchange). The traffic selectors are used in IKE negotiations to control what traffic can access the tunnel.
- Supports Hash and URL certificate exchange to reduce fragmentation.
- Resiliency against DoS attacks with improved peer validation. An excessive number of half-open SAs can trigger cookie validation.

Why This Best Practice Is Important:

IKEv2 is just better. You should prefer it, but be willing to fall back to IKEv1 if necessary.

How to Implement It:

1. Go to **Network > Network Profiles > IKE Gateway**.
2. Edit or create an IKE Gateway Profile.
3. In the “Version” dropdown box, select “IKEv2 preferred mode”.

What It Looks Like After You've Implemented It:

The screenshot shows the 'IKE Gateway' configuration window with the 'General' tab selected. A red oval highlights the 'Name' field (IKE Gateway-03) and the 'Version' dropdown menu (IKEv2 preferred mode). Other fields include Address Type (IPv4), Interface, Local IP Address (None), Peer IP Type (Static), Peer IP Address, Authentication (Pre-Shared Key), Pre-shared Key, Confirm Pre-shared Key, Local Identification, and Peer Identification.

Prefer IKEv2, but be willing to fall back to IKEv1

□ IKE Gateway: Prefer Certificates Over Pre-Shared Keys

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

One reasonably good definition of a VPN is that it's an *authenticated, encrypted, private communications channel over an untrusted network*. This Best Practice is about the authentication part.

It's not enough to just have good encryption in your VPN; both peers also need to be able to prove to themselves that they're communicating with whom they think they're communicating; otherwise you're vulnerable to a man-in-the-middle attack.

There are two ways to authenticate the peers in a VPN:

- Pre-Shared Key
- Certificate

There are two advantages and one disadvantage to using certificates:

- Passwords are usually created by people, and people are lazy (Rationally ignorant? Efficiently optimizing?) when it comes to creating long, secure passwords. By contrast, certificates are randomly distributed among their entire password space so they're always strong.
- Certificates are revocable, which allows an extra level of control, but passwords aren't.
- Certificates are more complex to implement.

Why This Best Practice Is Important:

The security benefits of authenticating your VPNs with certificates outweigh the additional complexities of the initial setup, so you should use them.

How to Implement It:

1. Go to **Network > Network Profiles > IKE Gateway**.
2. Edit or create an IKE Gateway Profile.
3. In the "Authentication" dropdown box, select "Certificate".
4. There are several other steps you'll have to complete, including generating and sharing the certificates. See the Palo Alto Networks technical documentation for the details.

What It Looks Like After You've Implemented It:

IKE Gateway

General **Advanced Options**

Name

Version **IKEv2 preferred mode**

Address Type ☒ IPv4 ☐ IPv6

Interface

Local IP Address **None**

Peer IP Type ☒ Static ☐ Dynamic

Peer IP Address

Authentication ☐ Pre-Shared Key ☒ **Certificate**

Local Certificate

☐ **HTTP Certificate Exchange**

Certificate URL

Local Identification **Distinguished Name (Subject)**

Peer Identification **Distinguished Name (Subject)**

Peer ID Check ☒ Exact ☐ Wildcard

☐ Permit peer identification and certificate payload identification mismatch

Certificate Profile

☐ Enable strict validation of peer's extended key use

OK **Cancel**

This is the first step in configuring your VPN peers to authenticate with certificates instead of pre-shared keys

□ IKE Gateway: Prefer Main Exchange Mode Over Aggressive

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

An IKE session begins when the Initiator sends a proposal to the Responder. The proposal contains information on the various parameters required for the peers to communicate properly, including details of authentication, encryption, key lifetimes, etc.

The Responder chooses which proposal to accept (if any) and responds. In a second exchange, the peers exchange Diffie-Hellman keys and other data. In a third exchange, they authenticate the ISAKMP session.

In IKEv1, these exchanges can occur in one of two modes:

- **Main Mode:** The Phase 1 parameters are exchanged in multiple rounds with encrypted authentication information. Main Mode uses six packets to accomplish all the steps.
- **Aggressive Mode:** The parameters are exchanged in a single message with unencrypted authentication information. Aggressive Mode uses three packets for the entire process.

While Aggressive Mode is faster, it has a security weakness. With Aggressive Mode and pre-shared keys it's possible for an attacker to gather all necessary information in order to mount an off-line dictionary (brute force) attack on the pre-shared keys.

Why This Best Practice Is Important:

Main Mode is more secure, so you should prefer it over Aggressive Mode. Also, because of the security risks, some security audits will return a fail verdict in the presence of Aggressive Mode.

How to Implement It:

1. Go to **Network > Network Profiles > IKE Gateway**.
2. Edit or create an IKE Gateway Profile.
3. Go to the "Advanced Options" tab, and then the "IKEv1" tab.
4. In the "Exchange Mode" dropdown box, select "main" if you're confident all the peers can also use Main Mode. Select "auto" to prefer Main Mode but accept Aggressive Mode if required.

What It Looks Like After You've Implemented It:

The screenshot shows the 'IKE Gateway' configuration window with the 'Advanced Options' tab selected. The 'Common Options' section includes checkboxes for 'Enable Passive Mode' and 'Enable NAT Traversal'. The 'IKEv1' and 'IKEv2' tabs are visible, with 'IKEv2' currently active. The 'Exchange Mode' dropdown is set to 'main' and is circled in red. Below it, the 'IKE Crypt. Profile' is set to 'default'. There is an unchecked checkbox for 'Enable Fragmentation'. The 'Dead Peer Detection' section is checked, with 'Interval' and 'Retry' both set to '5'. At the bottom right are 'OK' and 'Cancel' buttons.

Accepting only Main Mode for the IKE negotiation

VM-Series NGFWs

☐ Use Hypervisor Assigned MAC Addresses

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Virtual machines live inside containers, and one of the places where the VM and the container touch is at network interfaces. Because the VM's network interfaces are virtual, not physical, they don't automatically have a globally unique hardware MAC address that encodes the manufacturer's registered identification number, known as a Burned-In Address (BIA). What's needed is a virtual MAC address.

There are two methods for assigning a virtual MAC address to the VM:

The PAN-OS Custom Schema:

This is an older method and doesn't always work.

Hypervisor Assigned MAC Address:

This is the newer method, is recommended, and became the default setting (at the author's recommendation) partway through the PAN-OS 7.1.x release series.

Why This Best Practice Is Important:

The author once spent three full days trying to understand why packets weren't transiting a VM-300. After many whiteboard network diagrams, lots of detailed troubleshooting to isolate the problem, and numerous conversations with very skilled Palo Alto Networks colleagues, a light bulb went on somewhere and the problem was solved instantly by enabling this setting.

Learn from those who went before you, and enable this setting if it's not already on by default.

How to Implement It:

1. Go to **Device > Setup > Management > General Settings**.
2. Check the box next to "User Hypervisor Assigned MAC Addresses".

General Settings

Hostname: PA-VM

Domain:

☐ Accept DHCP server provided Hostname

☐ Accept DHCP server provided Domain

Login Banner: WARNING: The use of this system is restricted to authorized users only.

Unauthorized access, use, or modification of the computer

☐ Force Admins to Acknowledge Login Banner

SSL/TLS Service Profile: None

Time Zone: US/Pacific

Locale: en

Date: 2017/03/08

Time: 07:57:56

Latitude: 37.3935531

Longitude: -121.9788409

☐ Automatically Acquire Commit Lock

☒ Certificate Expiration Check

☒ Use Hypervisor Assigned MAC Addresses

OK Cancel

Seriously, you want this

What Else You Need to Know:

- If you enable this option and use an IPv6 address for the interface, the interface ID must not use the EUI-64 format, which derives the IPv6 address from the interface MAC address.

In a high availability (HA) active/passive configuration, a commit error occurs if the EUI-64 format is used.

•

QoS

□ Use QoS to Match Your Traffic Flows to Your Priorities

Improve Security		Improve Manageability	
Improve Performance	X	Improve High Availability	

Background Information:

PAN-OS provides basic QoS functionality, controlling traffic leaving the firewall according to network or subnet, and extends the power of QoS to also classify and shape traffic according to application and user.

QoS integrates the output of the App-ID and User-ID engines so you can easily specify applications and users for which you want to manage or guarantee bandwidth.

Why This Best Practice Is Important:

Service quality measurements subject to a QoS implementation are:

- Bandwidth (maximum rate of transfer)
- Throughput (actual rate of transfer)
- Latency (delay)
- Jitter (variance in latency)

The capability to shape and control these service quality measurements makes QoS of particular importance to high-bandwidth, real-time traffic such as voice over IP (VoIP), video conferencing, and video-on-demand that has a high sensitivity to latency and jitter.

Additionally, you can use QoS to achieve outcomes such as the following:

- Prioritize network and application traffic, guaranteeing high priority to important traffic or limiting non-essential traffic.
- Achieve equal bandwidth sharing among different subnets, classes, or users in a network.
- Allocate bandwidth externally or internally or both, applying QoS to both upload and download traffic or to only upload or download traffic.
- Ensure low latency for customer and revenue-generating traffic in an enterprise environment.
- Perform traffic profiling of applications to ensure bandwidth usage.

How to Implement It:

1. Study the details in the technical documentation at <https://PaloAltoNetworks.com>. It's a bit more complicated than most configuration tasks in PAN-OS.
2. Configure a QoS Egress Interface
3. Configure a QoS Profile

4. Configure a QoS Policy

What Else You Need to Know:

Configuring your QoS Policy will benefit from Continuous Improvement. Observe what your traffic flows are, listen to the feedback from your users, and make adjustments until you get it right.

□ Begin Your QoS Configuration With a Simple Plan

Improve Security		Improve Manageability	
Improve Performance	X	Improve High Availability	

Why This Best Practice Is Important:

Configuring QoS can be a large project and can seem daunting, but even a small, simple policy can make big improvements in your traffic management.

How to Implement It:

When you're starting out, partition your traffic into just a few basic categories, ordered from highest to lowest priority. Here's a pretty good start:

Category #1:

Video and audio streaming where there's a live human on the connection who's going to get annoyed if there's latency, jitter, or dropped packets. This category is the most sensitive to QoS.

Category #2:

Web browsing and similar applications where there's a live human on the connection who wants a responsive interface

Category #3:

E-mail, where delays of a few seconds are tolerable

Category #4:

File transfers, where delays of a few minutes are tolerable

Category #5:

Backups, where delays of a few hours are tolerable.

Category #6:

The Penalty Box, where you throttle applications that you want to discourage but you know if you blocked them entirely you'd just encourage people to look for a way to tunnel through or around your QoS engine. Make the pipe small enough that it doesn't adversely affect other applications, but large enough to ensure no one thinks you're "blocking" the application and starts looking for ways to get creative.

This is a moderately sophisticated system of categorization. If you were to implement these priorities as a QoS policy you'd already be assigning your bandwidth in ways that do a pretty good job of supporting your organization's goals.

Monitoring and Logging

❑ Prefer “Log at Session End” Over “Log at Session Start”

Improve Security		Improve Manageability	X
Improve Performance	X	Improve High Availability	

Background Information:

PAN-OS can log at the start of a session (this is disabled by default), or at the end of a session (this is enabled by default), or both, or neither.

Why This Best Practice Is Important:

Logging at the end of the session is better because there’s more information available to the firewall at the end of the session, including information that App-ID and User-ID might need to complete their analysis. Also, because logging consumes resources, logging twice per session has performance implications.

Enabling Log at Session Start is only useful when you’re troubleshooting and you need to see the initial setup for a session that’s not completing properly.

How to Implement It:

Go to **Policies > Security** and edit a Security Policy Rule. On the Actions tab, configure the checkboxes in the Log Setting section.

The screenshot shows the 'Security Policy Rule' configuration window with the 'Actions' tab selected. The 'Log Setting' section is circled in red, indicating the recommended configuration: 'Log at Session Start' is unchecked, and 'Log at Session End' is checked. The 'Log Forwarding' dropdown is set to 'None'. The 'Other Settings' section shows 'Schedule' and 'QoS Marking' both set to 'None', and 'Disable Server Response Inspection' is unchecked. The 'Action Setting' section shows 'Action' set to 'Allow' and 'Send ICMP Unreachable' is unchecked. The 'Profile Setting' section shows 'Profile Type' set to 'Group' and 'Group Profile' set to 'BJS-Security-Profile-Group'.

Configuring the Log Setting options

❑ Prefer to Not Log DNS, NTP, DHCP and LDAP Traffic

Improve Security		Improve Manageability	X
Improve Performance	X	Improve High Availability	

Why This Best Practice Is Important:

The DNS, NTP, DHCP and LDAP protocols tend to generate lots of sessions and the log records they generate usually aren't very interesting to firewall administrators. Because logging these protocols consumes resources and reduces the signal-to-noise ratio in the logs, it's best not to log these sessions unless you're actively troubleshooting an issue.

How to Implement It:

Split the Allow rules that permit these protocols, carving these protocols off into separate rules for which logging is not enabled.

What Else You Need to Know:

Your security compliance team may force you to log all these protocols anyway.

□ Think Hard About What to Log

Improve Security		Improve Manageability	X
Improve Performance	X	Improve High Availability	

Background Information:

On the one hand...

Logging provides lots of useful information for later analysis and troubleshooting.

But on the other hand...

Logging consumes CPU, I/O and storage capacity, all of which might be scarce resources.

Why This Best Practice Is Important:

Unless your security compliance team is requiring you to log absolutely everything, you're going to have to weigh the trade-offs in deciding what to log and what not to.

How to Implement It:

Make a written plan. Run experiments and keep notes on what you've observed. Try things. Use Continuous improvement and keep adjusting it until you've reached the right balance between visibility and resource consumption.

What It Looks Like After You've Implemented It:

You're logging everything you need, and your resources aren't depleted.

□ Enable "Resolve Hostname" While Viewing Logs

Improve Security		Improve Manageability	X
Improve Performance	X	Improve High Availability	

Background Information:

Every packet is routed across your network or across the Internet by its IPv4 or IPv6 address, which looks like **205.219.84.5** or **2001:db8:85a3:8d3:1319:8a2e:370:7348**. While these addresses are precise and compact and easy for a router to understand, it turns out that humans, with our slow, carbon-based processors and limited digital registers, are terrible at remembering or processing long strings of numbers like this.

This deficiency is the main reason that hosts began to use names, then the *hosts* file, then distributed *hosts* files, and then DNS.

Why This Best Practice Is Important:

While reviewing logs, you'll notice the Source and Destination address columns show only IP addresses. Your comprehension will increase significantly if you have the GUI resolve these addresses into hostnames for you.

When resolving hostnames, the GUI uses a three-step algorithm:


Step 1: If the IP address resolves to an Address object, the name of the Address object is shown.

Step 2: If it didn't get resolved in Step 1, then a reverse DNS call is made to see if DNS can provide a hostname. If so, *this* hostname is shown.

Step 3: If Steps 1 and 2 didn't work, then the original IP address is shown.

How to Implement It:

While viewing logs, click the box next to "Resolve hostname" in the lower left corner of the table of log events.



Giving your brain a break

What Else You Need to Know:

Unfortunately, this setting isn't "sticky"; every time you click on a log page, this setting will default to off, and you'll have to manually set it again.

□ Get Skilled With the Log Filter Language

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

Firewall administrator graybeards know that most configuration errors or other mysteries are resolved by spending time in the logs, isolating the problem and figuring out what's going on.

If you go to **Monitor > Logs** you'll see there are nine separate logs you can view, plus a unified union of all of them:

1. Traffic
2. Threat
3. URL Filtering
4. WildFire Submissions
5. Data Filtering
6. HIP Match
7. Configuration
8. System
9. Alarms
10. Unified

When troubleshooting, you'll probably spend most of your time in the Traffic and Threat logs.

Why This Best Practice Is Important:

Without filters, each log view is a full dump of every entry, sorted by Receive Time. This gets cumbersome quickly. To become an ace firewall administrator you must master the jujitsu of the log filter language so you can significantly improve the signal-to-noise ratio of the information in front of you.

How to Implement It:

The Really Easy Method:

For any field in which the contents are a hyperlink, just click on the link and an inclusion filter for that Attribute and Value will be added to the existing filter, joined with the "and" operator. Click on everything you want to match, in sequence, and you can build the filter you want.

One disadvantage of this method is that, without editing the filter manually, you can only add Attribute/Value pairs if the Value is in sight and thus clickable.

Also, negation requires manual editing here, along with any more complex Boolean logic.

The Easy Method:

Just to the right of the log filter text area, there's a '+' icon with "Add Filter" hover text. Click on it to open the Add

Log Filter dialog box:

Connector	Attribute	Operator	Value
and	Action	is present	ping
or	Action Source	equal	
	Address	not equal	
	Application		
	Bytes		
	Bytes Received		
	Bytes Sent		
	CVE		
	Category		
	Content Type		
	Destination Address		
	Destination Country		

☐ Negate

Add Close

This filter will match when the Application is equal to ping

Click Add and this filter will automatically be added and applied.

In this example, the filter created will be **(app eq ping)**.

After you click Add, the dialog box will stay open and you can enter another filter and link them with an “and” or “or” operator.

The Advanced Method:

The Advanced Method is to roll up your sleeves and just write the filter you want in the text box. You may need to use this method when you want to create complex Boolean logic using parentheses and the “and” and “or” operators and negation.

It’s not RegEx, but it works fairly well.

What It Looks Like After You’ve Implemented It:

You’ll know you’ve got it right when the filter logic looks correct and you’re getting exactly what you want in the log, but not more. Now you can focus on problem solving.

□ Review Your Logs Regularly

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Why This Best Practice Is Important:

By reviewing your logs regularly, you'll get three benefits:

- You'll start to understand what "Normal" looks like, so you'll be better able to detect anomalies.
- You'll be able to use continuous improvement to keep improving your Security Rulebase and the filters you use to view your logs.
- You'll start recognizing common hosts that are important enough for you to deserve their own named Object.

How to Implement It:

Start reviewing your logs. Look for things that *Don't Look Right* (DLR). Create objects for Addresses that you see frequently. Examine blocked traffic to see what attacks are being prevented. Examine allowed traffic to see if you can tighten up your rulebase.

What It Looks Like After You've Implemented It:

If you're reviewing your logs regularly and using continuous improvement to improve your policies, you'll feel like you've got things more or less mastered and under control and you'll be much better able to spot something new and out of place.

❑ Select "Enable Log on High DP Load"

Improve Security		Improve Manageability	X
Improve Performance	X	Improve High Availability	

Background Information:

Palo Alto Networks firewalls have separate Management and Data Planes. If the traffic load gets too high, the Data Plane CPU load can reach 100% which could cause performance degradation or dropped packets.

Why This Best Practice Is Important:

If your traffic load is increasing, you want to hear about it early so you can make plans to either tune your firewall for higher performance or migrate to a box with a larger capacity.

Enabling this setting will create a logging event, no more frequently than once per minute, when the Data Plane reaches 100% CPU usage.

How to Implement It:

1. Go to **Device > Setup > Logging and Reporting Settings > Log Export and Reporting**.
2. Check the box next to "Enable Log on High DP Load".

Logging and Reporting Settings

Log Storage | **Log Export and Reporting** | Pre-Defined Reports

Number of Versions for Config Audit: 100

Max Rows in CSV Export: 65535

Max Rows in User Activity Report: 5000

Average Browse Time (sec): 60

Page Load Threshold (sec): 20

Syslog HOSTNAME Format: FQDN

Report Expiration Period (days): [1 - 2000]

Warning: Deletion of logs based on time period may take a long time and during this time the max sustainable log rate will be degraded

☐ Stop Traffic when Log Db Full

☒ **Enable Log on High DP Load**

OK Cancel

A high Data Plane load will create a log event

High Availability

□ Configure HA

Improve Security		Improve Manageability	
Improve Performance		Improve High Availability	X

Background Information:

High availability (HA) is a special deployment of a pair of firewalls that acts as a group, synchronizing their configuration to prevent a single point of failure. A heartbeat connection between the firewall peers ensures seamless failover in the event that one of them goes down.

Why This Best Practice Is Important:

An HA pair provides redundancy and reduces the small risk of a single firewall failure. With HA configured, a single firewall failure doesn't bring down your connections and gives you time to replace the dead unit on a non-emergency basis. It can turn what might be a firefighting exercise into a maintenance task.

How to Implement It:

Configuring two firewalls into an HA pair is not very difficult. The Administrator's Guide is sufficient and Palo Alto Networks Technical Support and Professional Services Engineers are always available to help.

What It Looks Like After You've Implemented It:

The HA Setup dialog box

What Else You Need to Know:

As of Version 7.1, there are a few caveats to remember:

- The PA-200 firewall supports HA Lite only.
- The VM-Series firewall in AWS supports active/passive HA only; if it is deployed with Amazon Elastic Load Balancing (ELB), it does not support HA (in this case ELB provides the failover capabilities).

The VM-Series firewall in Microsoft Azure does not support HA.

•

□ Prefer Active/Passive HA Over Active/Active

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	X

Background Information:

You configure HA on a pair of firewalls in one of two modes: **Active/Passive** or **Active/Active**. Here's how to tell them apart:

Active/Passive:

- One firewall actively handles traffic while the other remains synchronized and ready to take over in the event the Active member fails.
- Active/Passive mode is supported in the virtual wire, Layer 2 and Layer 3 deployments.
- Active/Passive mode is simpler.

Active/Active:

- Both firewalls in the pair are Active and they work together to handle traffic and session setup and session ownership.
- Active/Passive mode is supported in the virtual wire and Layer 3 deployments, but not Layer 2.
- Active/Active doesn't actually load-balance traffic.
- Configuring Active/Active might require additional configuration complexities such as activating networking protocols on both firewalls, replicating NAT pools, and deploying floating IP addresses to provide proper failover. Because both firewalls are actively processing traffic, the firewalls use additional concepts of session owner and session setup to perform Layer 7 content inspection.

Why This Best Practice Is Important:

Active/Passive is simpler and easier to troubleshoot, partially because you can always be certain which firewall is handling the traffic. Given that the first step in troubleshooting is always to isolate the problem, this simpler and disambiguated path and session state leads to faster debugging.

Because in any HA configuration you need to size the firewalls to be able to handle the full traffic load individually in the event one of them fails, you don't gain any throughput capacity by choosing Active/Active over Active/Passive. Many organizations unintentionally go over 50% capacity on each member of an Active/Active configuration, so when one member fails, the remaining member is now at 100% capacity and slowing or dropping traffic.

It could be argued that the main reason Active/Active mode exists at all is for the psychological benefit some less-informed customers might gain from thinking, "We paid for two firewalls, so let's make sure we're putting both to use all the time!". Active/Active is simpler, and when properly sized, doesn't reduce total throughput.

How to Implement It:

Go to **Device > High Availability > General > Setup**.

What It Looks Like After You've Implemented It:

The screenshot shows a 'Setup' dialog box for configuring High Availability (HA) on a Palo Alto Networks device. The dialog has a blue header with the title 'Setup' and a help icon. The main content area is light gray and contains several configuration options:

- ☒ Enable HA
- Group ID:
- Description:
- Mode: ☒ Active Passive ☐ Active Active
- ☒ Enable Config Sync
- Peer HA1 IP Address:
- Backup Peer HA1 IP Address:

The 'Active Passive' radio button is circled in red. At the bottom of the dialog are two buttons: 'OK' and 'Cancel'.

Choosing Active/Passive over Active/Active. You'll thank yourself the next time you're troubleshooting.

□ Configure an HA1 Control Link Backup

Improve Security		Improve Manageability	
Improve Performance		Improve High Availability	X

Background Information:

The purpose of configuring for High Availability is to eliminate single points of failure to reduce the small risk of a service outage.

One possible single point of failure in a firewall High Availability configuration is the HA1 Control Link. This Layer 3 link is used to exchange:

- Hellos
- Heartbeats
- HA state information
- Management plane synchronization for routing
- User-ID information
- Configuration changes

Why This Best Practice Is Important:

If the HA1 Control Link fails, the two firewalls can't communicate and you've lost the protection of High Availability. Fortunately, it's easy to configure a second, redundant link.

How to Implement It:

Go to **Device > High Availability > General > Control Link (HA1 Backup)**.

Specifying which port and IP address to use for “this end” of the backup HA1 Control Link

Go to **Device > High Availability > General > Setup**

Setup

☒ Enable HA

Group ID: 1

Description:

Mode: ☒ Active Passive ☐ Active Active

☒ Enable Config Sync

Peer HA1 IP Address: 192.168.1.2

Backup Peer HA1 IP Address: 192.168.2.2

OK Cancel

Specifying the peer's IP address for "the other end" of the backup HA1 Control Link

Special considerations for a PA-7000 Series firewall:

HA connectivity on the PA-7000 Series mandates the use of specific ports on the Switch Management Card (SMC) for certain functions:

- Use the special HA1-A port for the HA1 Control Link, and the special HA1-B port for the backup HA1 Control Link.
- HA1 cannot be configured on NPC data ports or the MGT port.

What Else You Need to Know:

- The backup HA1 Control Link must use a different interface (an in-band or "traffic" port) and be on a different subnet from the primary HA1 Control Link.
- HA1-backup and HA2-backup ports must be configured on separate physical ports.
- The HA1-backup link uses ports 28770 and 28260.

You should enable heartbeat backup (uses port 28771 on the MGT interface) if you use an in-band port for the HA1 or the HA1 backup links.

•

□ Configure an HA2 Data Link Backup

Improve Security		Improve Manageability	
Improve Performance		Improve High Availability	X

Background Information:

The purpose of configuring for High Availability is to eliminate single points of failure to reduce the small risk of a service outage.

One possible single point of failure in a firewall High Availability configuration is the HA2 Data Link. This Layer 2 link is used to synchronize:

- Sessions
- Forwarding tables
- IPSec Security Associations
- ARP tables

Why This Best Practice Is Important:

If the HA2 Control Link fails, the two firewalls can't communicate and you've lost the protection of High Availability. Fortunately, it's easy to configure a second, redundant link.

How to Implement It:

Go to **Device > High Availability > General > Data Link (HA2 Backup)**.

Specifying which port and IP address to use for “this end” of the backup HA2 Data Link

Special considerations for a PA-7000 Series firewall:

HA connectivity on the PA-7000 Series mandates the use of specific ports on the Switch Management Card (SMC) for certain functions.

Data Link:

- The High Speed Chassis Interconnect (HSCI) ports are Layer 1 Quad Port SFP+ (QSFP+) interfaces used to connect two PA-7000 Series firewalls in an HA configuration. Each port is comprised of four 10 gigabit channels multiplexed for a combined speed of 40 gigabits.

- The traffic carried on the HSCI ports is raw Layer 1, which is not routable or switchable; therefore the HSCI ports must be connected directly to each other. The HSCI-A on the first chassis connects directly to HSCI-A on the second chassis and HSCI-B on the first chassis connects to HSCI-B on the second chassis. This provides full 80 gigabit transfer rates. In software, both ports (HSCI-A and HSCI-B) are treated as one HA interface.
- It's best to use the dedicated HSCI ports for the HA2 link. The HA3 link, required for packet forwarding in an active/active deployment, must use the HSCI port; the HA3 traffic cannot be configured on data ports.

What Else You Need to Know:

- The backup HA2 Data Link must use a different interface (and in-band or "traffic" port) and be on a different subnet from the primary HA2 Data Link.

HA1-backup and HA2-backup ports must be configured on separate physical ports.

❑ Configure HA1 Control Link Encryption

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

In a High Availability configuration, the HA1 Control Link is used to exchange:

- Hellos
- Heartbeats
- HA state information
- Management plane synchronization for routing
- User-ID information
- Configuration changes

Obviously, this information is security sensitive.

Why This Best Practice Is Important:

By default, information sent over the HA1 Control Link is not encrypted. If this data has to travel outside of a datacenter, or even just through a router, it's probably best to encrypt it, but if your two firewalls are physically adjacent and this link is a single Ethernet cable, then encryption might not be necessary.

How to Implement It:

Go to **Device > High Availability > Control Link (HA1)**

The screenshot shows the 'Control Link (HA1)' configuration window. The 'Port' is set to 'ethernet1/6', 'IPv4/IPv6 Address' is '192.168.1.1', 'Netmask' is '255.255.255.252', and 'Gateway' is empty. The 'Encryption Enabled' checkbox is checked and highlighted with a red circle. The 'Monitor Hold Time (ms)' is set to '3000'. 'OK' and 'Cancel' buttons are at the bottom right.

Enabling encryption on the HA1 Control Link

What Else You Need to Know:

When you enable HA1 Control Link encryption, you will be prompted to complete a manual exchange of keys between the two peers.

❑ Configure Dynamic Updates Schedules to "Sync To Peer"

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

In a High Availability configuration, both members of a pair need the same Dynamic Updates installed.

Why This Best Practice Is Important:

By enabling this option you ensure the downloads are synchronized.

How to Implement It:

Go to **Device > Updates**

Configure the schedule for each Dynamic Update and check the box next to "Sync To Peer":

Now they'll always have the same versions

❑ When Manually Downloading Software or Dynamic Updates, Choose “Sync to HA Peer”

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	


Background Information:

In a High Availability configuration, both members of a pair need Software and Dynamic Updates stored locally on the box.


Why This Best Practice Is Important:

Rather than having to log in to both peers and download Software images and Dynamic Updates manually on both, select the “Sync to HA Peer” check box when you’re presented with the downloading dialog box.

How to Implement It:



Easy to select, and will save you some time



It works the same with Dynamic Updates

□ Test Your HA Configuration With Real Failures

Improve Security		Improve Manageability	
Improve Performance		Improve High Availability	X

Background Information:

The whole purpose of configuring High Availability is to eliminate single points of failure and reduce the small risk of a service outage.

Why This Best Practice Is Important:

I.T. Graybeards like to follow a well-known rule: *It's not done until it's tested.*

How to Implement It:

At time of implementation, or in a lab, or during a scheduled maintenance window, unplug a power cable, or the HA1 Control Link, or the HA2 Data Link. Break things. Observe what happens. If the unexpected happens, do a Root Cause Analysis and get to the bottom of it and fix it and prevent it from happening next time. Get completely clear on what types of failures your High Availability configuration can withstand and then prove it can withstand them.

If you're not confident enough in your High Availability configuration to intentionally cause a covered failure, then you're not done implementing High Availability.

While Troubleshooting

❑ Consider Logging at Both Session Start and at Session End

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

PAN-OS allows you to log the matching of a Security Policy Rule at the beginning of a Session, at the end, or both, or neither.

Why This Best Practice Is Important:

If you're in the middle of troubleshooting a difficult problem, you might need additional visibility into how a packet or session is being processed by the firewall. By configuring when to log the session, you can gain additional, and different, information in the log.

How to Implement It:

Configure the Log Setting in the Actions tab of the Security Policy Rule:

The screenshot shows the 'Security Policy Rule' configuration window with the 'Actions' tab selected. A red circle highlights the 'Log Setting' section. In this section, the 'Log at Session End' checkbox is checked, while 'Log at Session Start' is unchecked. The 'Log Forwarding' dropdown is set to 'None'. Below this, the 'Other Settings' section shows 'Schedule' and 'QoS Marking' both set to 'None', and the 'Disable Server Response Inspection' checkbox is unchecked. The 'Action Setting' section shows the 'Action' set to 'Allow' and 'Send ICMP Unreachable' is unchecked. The 'Profile Setting' section shows 'Profile Type' set to 'Group' and 'Group Profile' set to 'BJS-Security-Profile-Group'. At the bottom right are 'OK' and 'Cancel' buttons.

Chose, none, one, or both

❑ Consider Sending ICMP Unreachable for Both Drops and Resets

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

The ICMP Type 3 Destination Unreachable message is generated by a router to inform the source host that the destination unicast address is unreachable. The goal is to inform the source so that it may gracefully close or clear the session and prevents some applications from breaking.

Why This Best Practice Is Important:

When you're troubleshooting, you might want to send this message as a way of being helpful. It indicates there's a live, thinking host on the wire, and helps disambiguate between a properly functioning firewall and an unplugged cable.

How to Implement It:

The screenshot shows the 'Security Policy Rule' configuration window. The 'Action' dropdown is set to 'Drop' and the 'Send ICMP Unreachable' checkbox is checked. These two elements are circled in red. Other settings like 'Log at Session Start', 'Log at Session End', 'Log Forwarding', 'Schedule', 'QoS Marking', and 'Disable Server Response Inspection' are also visible.

The Send ICMP Unreachable setting

What Else You Need to Know:

The Send ICMP Unreachable option is available only on Layer3 interfaces, and only with the Drop or Reset actions.

Upgrading

☐ Stay Mostly Current With Minor Software Releases

Improve Security	X	Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

From time to time, Palo Alto Networks releases minor updates to PAN-OS. These usually contain bug fixes and may occasionally contain small functionality improvements. In the past, the time between each release was something like 6-8 weeks.

Minor software releases are identified by a change in the third portion of the version release number, such as going from 7.1.7 to 7.1.8.

Why This Best Practice Is Important:

On the one hand...

You want to stay current with minor software releases because you want to install bug fixes as quickly as possible, and preferably before they can affect your firewall or your customers. This is an argument for installing new minor releases quickly after they're released, especially if you're having an issue that will get resolved in a specific release.

But on the other hand...

It's rare, but occasionally there's a problem with a release, and some customers, upon later reflection, wish they had waited a bit before installing the latest version, to allow other customers to discover and report an issue and let the company fix it in an interim release.

How to Implement It:

You'll have to find the resolution between those two forces that best suits your needs, but waiting, say, five days after a release might be a reasonable compromise.

Go to **Device > Software** and click on "Check Now" in the bottom left corner. This will update the list of available software releases. You can then download and install the version you want.

2017 Palo Alto Networks Best Practices

Version ▲	Size	Release Date
7.1.0	549 MB	2016/03/29 19:43:48
7.1.1	221 MB	2016/04/18 08:01:40
7.1.2	237 MB	2016/05/12 02:55:17
7.1.3	239 MB	2016/06/29 00:45:47
7.1.4	246 MB	2016/08/12 00:22:55
7.1.4-h2	246 MB	2016/08/19 15:49:07
7.1.5	248 MB	2016/10/03 00:24:29
7.1.6	278 MB	2016/11/17 15:11:23
7.1.7	278 MB	2016/12/30 15:44:14
8.0.0	573 MB	2017/01/29 14:38:41

This is where you look for the most current release

❑ Let Major Software Releases Mature for a While Before Installing

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

From time to time Palo Alto Networks releases a major software update to PAN-OS. This may correspond to the release of new hardware or virtual firewalls, and includes new functionality and performance enhancements. In the past, the time between these releases has been something like 8-12 months.

Major software releases are identified by a change in the first or second portion of the version release number, such as going from 7.0.x to 7.1.x, or from 7.1.x to 8.0.x.

The major releases with a '0' in the second number tend to include more new functionality than major releases with a '1' in the second number, which tend to be more of a maintenance release.

Why This Best Practice Is Important:

On the one hand...

You want to stay current with major software releases because you want to take advantage of the new features and performance enhancements as quickly as possible. This is an argument for installing new major releases quickly after they're released.

But on the other hand...

It's rare, but occasionally there's a problem with a release, and some customers, upon later reflection, wish they had waited a bit before installing the latest version, to allow other customers to discover and report an issue and let the company fix it in an interim release.

How to Implement It:

You'll have to find the resolution between those two forces that best suits your needs, but waiting, say, until the third minor release after a major release might be a reasonable compromise.

Go to **Device > Software** and click on "Check Now" in the bottom left corner. This will update the list of available software releases. You can then download and install the version you want.

2017 Palo Alto Networks Best Practices

Version ▲	Size	Release Date
7.1.0	549 MB	2016/03/29 19:43:48
7.1.1	221 MB	2016/04/18 08:01:40
7.1.2	237 MB	2016/05/12 02:55:17
7.1.3	239 MB	2016/06/29 00:45:47
7.1.4	246 MB	2016/08/12 00:22:55
7.1.4-h2	246 MB	2016/08/19 15:49:07
7.1.5	248 MB	2016/10/03 00:24:29
7.1.6	278 MB	2016/11/17 15:11:23
7.1.7	278 MB	2016/12/30 15:44:14
8.0.0	573 MB	2017/01/29 14:38:41

This is where you look for the most current release

□ Verify Update Server Identity

Improve Security	X	Improve Manageability	
Improve Performance		Improve High Availability	

Background Information:

Your firewall will download and install PAN-OS updates from the Palo Alto Networks Update Server, which is configured by default to be found at **updates.PaloAltoNetworks.com**.

PAN-OS offers you the opportunity to verify that the server from which the software or content package is downloaded has an SSL certificate signed by a trusted authority.

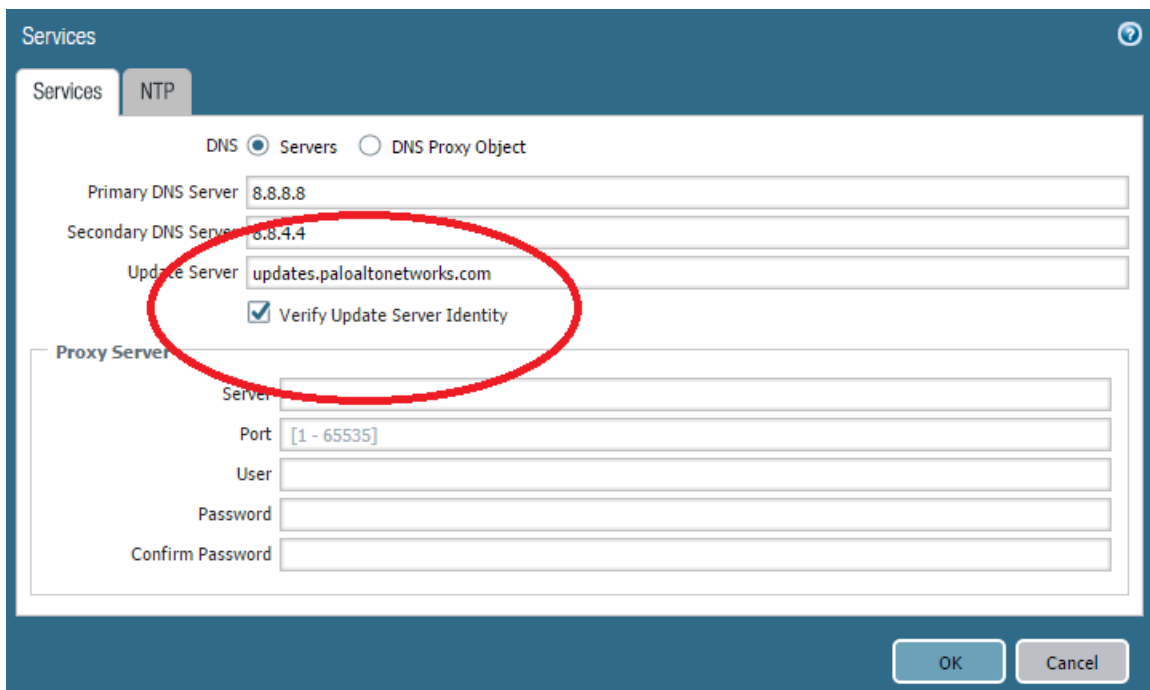
Why This Best Practice Is Important:

If someone could hijack your secure connection to this server and trick you into downloading and installing a cooked version of PAN-OS, this is a major breach and would be hard to detect.

It's best to add an additional level of security to these downloads by verifying the SSL certificate at the update site every time.

How to Implement It:

Go to **Device > Setup > Services > Services**.



The screenshot shows the 'Services' configuration page in the Palo Alto Networks management interface. The 'DNS' section is expanded, and the 'Servers' radio button is selected. The 'Update Server' field is populated with 'updates.paloaltonetworks.com', and the 'Verify Update Server Identity' checkbox is checked. A red circle highlights the 'Update Server' field and the checkbox. The 'Proxy Server' section is also visible but not selected.

Enable this verification to provide additional security

Deployment Scenario: Public Cloud

General

□ Think Carefully When Choosing the BYOL Versus PAYGO License Models

Improve Security		Improve Manageability	X
Improve Performance		Improve High Availability	

Background Information:

You can license the VM-Series firewall in AWS and Azure in two ways:

Bring Your Own License (BYOL): This license is purchased from a partner, reseller, or directly from Palo Alto Networks. BYOL supports capacity licenses, support licenses, and subscription licenses. With this option, you must apply the license after you deploy the VM-Series firewall.

Usage-Based License, also called the Pay-As-You-Go (PAYGO) license: This type of license can be purchased (rented?) from the AWS or Azure Marketplace.

AWS supports hourly and annual PAYGO options; Azure supports the hourly PAYGO option only.

With the usage-based licenses, the firewall is pre-licensed and ready for use as soon as you deploy it; you do not receive an auth code. When the firewall is stopped or terminated on the AWS or Azure console, the usage-based licenses are suspended or terminated.

Usage-based licenses are available in two pricing bundles:

Feature	Bundle #1	Bundle #2
Capacity License	VM-300	VM-300
IPS	X	X
AV	X	X
Malware Prevention	X	X
GlobalProtect		X
WildFire		X
PAN-DB URL Filtering		X
Premium Support	X	X




BYOL licenses can be configured for any of the VM-Series models.

You cannot convert one type of license to the other, but if you have both types of licenses, you can switch them in your VM-Series firewalls.

Why This Best Practice Is Important:

Because it's difficult to change course once you've selected your licensing model, it's best to understand the implications and choose wisely from the beginning.

How to Implement It:

 <p>Free Trial</p>	<p>VM-Series Next-Generation Firewall Bundle 2</p> <p>★★★★★ (3) PAN-OS 8.0 Previous versions Sold by Palo Alto Networks</p> <p>\$1.38/hr or \$4,800/yr (60% savings) for software + AWS usage fees</p> <p>Linux/Unix, Other PAN-OS 8.0 64-bit Amazon Machine Image (AMI) Updated: 2/26/17</p> <p>The VM-Series complements AWS Security Groups and Network ACLs, by uniquely classifying and controlling your AWS traffic based on the application identity, and applying Threat ...</p> <p>More info</p>	<p>Select</p>
 <p>Free Trial</p>	<p>VM-Series Next-Generation Firewall Bundle 1</p> <p>★★★★★ (1) PAN-OS 8.0 Previous versions Sold by Palo Alto Networks</p> <p>\$0.86/hr or \$3,000/yr (60% savings) for software + AWS usage fees</p> <p>Linux/Unix, Other PAN-OS 8.0 64-bit Amazon Machine Image (AMI) Updated: 2/26/17</p> <p>The VM-Series complements AWS Security Groups and Network ACLs, by uniquely classifying and controlling your AWS traffic based on the application identity, and applying Threat ...</p> <p>More info</p>	<p>Select</p>
	<p>VM-Series Next-Generation Firewall (BYOL)</p> <p>★★★★★ (1) PAN-OS 8.0 Previous versions Sold by Palo Alto Networks</p>	<p>Select</p>

Selecting a licensing model within AWS

Amazon Web Services (AWS)

□ Think Carefully When Choosing Your Instance Type

Improve Security	Improve Manageability	X
Improve Performance	Improve High Availability	

Background Information:

In addition to providing VM-Series firewalls, Amazon Web Services (AWS) provides the underlying virtual machines to run them on.

Why This Best Practice Is Important:

You need to properly size and specify the underlying virtual machine to ensure proper performance of your VM-Series NGFW.

How to Implement It:

1. Understand the resource requirements of your VM-Series NGFW.
2. Choose the appropriate Instance Type.

EC2 Instance Type	vCPU	Mem (GiB)	SSD Storage (GB)	Dedicated EBS Bandwidth (Mbps)	EC2 /hr Price
c4.xlarge	4	7.5	EBS-Only	750	\$0.20
c3.xlarge	4	7.5	2 x 40	NA	\$0.21
m4.xlarge	4	16	EBS-only	750	\$0.22
m3.xlarge	4	15	2 x 40	NA	\$0.27
c4.2xlarge	8	15	EBS-Only	1,000	\$0.40
c3.2xlarge	8	15	2 x 80		\$0.42
m4.2xlarge	8	32	EBS-only	1,000	\$0.43
m3.2xlarge	8	30	2 x 80	NA	\$0.53
c4.4xlarge	16	30	EBS-Only	2,000	\$0.80
c3.4xlarge	16	30	2 x 160	NA	\$0.84
m4.4xlarge	16	64	EBS-only	2,000	\$0.86
c4.8xlarge	36	60	EBS-Only	4,000	\$1.59
c3.8xlarge	32	60	2 x 320	NA	\$1.68

AWS instances and pricing

What Else You Need to Know:

In PAN-OS versions before 8.0 the throughput performance of the firewall was only dictated by the number of virtual CPUs and memory on the instance. The firewall software license dictated the number of firewall software objects like concurrent sessions, address objects, etc. Starting with the PAN-OS version 8.0 release the AWS instance type and the VM-Series license type are combined to offer a wider selection of performance options.

□ Understand the Special AWS Routing Model

Improve Security	Improve Manageability	X
Improve Performance	Improve High Availability	

Background Information:

In AWS the routing is simplified for general compute work, but doesn't always easily accommodate easy insertion of network security devices such as VM-Series NGFWs.

There are some special considerations to keep in mind:

- When an AWS VPC is created, the CIDR range of the private IP addresses to be used inside of the VPC is created.
- Only one CIDR range per VPC is allowed.
- All route tables have an automatic route entry defining the VPC CIDR range as locally reachable, meaning any subnet inside the VPC has an IP route to any other subnet inside the VPC.
- You cannot write a more specific route than the VPC CIDR range. For example, if the VPC CIDR range is 10.0.0.0/16 you cannot add a route for 10.0.1.0/24.

Why This Best Practice Is Important:

These routing rules require that some additional design be done to facilitate the required security posture inside the VPC. These designs may include the need to do NAT functions or routing at hosts.

How to Implement It:

Before your implementation, ensure you really understand the special AWS routing requirements and plan your topology out in detail.

❑ Configure Identity and Access Management (IAM) Roles Carefully

Improve Security	X	Improve Manageability
Improve Performance		Improve High Availability

Background Information:

The AWS mechanism for controlling access and authority regarding the infrastructure is controlled by Identity and Access Management (IAM).

Why This Best Practice Is Important:

Best Practices for general IAM use are well defined in the AWS documentation, but you should take special care should when working with security devices such as an NGFW.

How to Implement It:

Think carefully about what permissions should be granted and follow the Principle of Least Privilege.

What It Looks Like After You've Implemented It:

This is an example of an HA policy defined in a role for high availability. As you can see the example is generic in that it defines what attributes about a resource can be acted upon, but does not restrict the resources that can be acted upon.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1443112577000",
      "Effect": "Allow",
      "Action": [
        "ec2:AttachNetworkInterface",
        "ec2:DescribeInstances",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DetachNetworkInterface"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

Specifying permitted Actions

□ Control Access to the Management Port with Security Groups

Improve Security	X	Improve Manageability
Improve Performance		Improve High Availability

Background Information:

The AWS cloud's public IP subnets are well known and receive lots of malicious scanning from the Internet.

Why This Best Practice Is Important:

If your management interface is Internet-facing then it's subject to intensive scanning. By creating a security group just for the firewall's management port, you can lock down exactly which IP addresses are allowed to connect, and over which ports.

How to Implement It:

Create a security group just for the management interface:

The screenshot shows the 'Create Security Group' dialog in the AWS Management Console. The 'Security group name' is 'Firewall Management Security', the 'Description' is 'Used to limit access to the firewall management interface', and the 'VPC' is 'vpc-3adaa15d | test'. Under 'Security group rules', the 'Inbound' tab is selected. Two rules are listed: SSH (Type: SSH, Protocol: TCP, Port Range: 22, Source: Anywhere, 0.0.0.0/0, :::/0) and HTTPS (Type: HTTPS, Protocol: TCP, Port Range: 443, Source: Anywhere, 0.0.0.0/0, :::/0). Both rules have a delete icon (X) to their right. An 'Add Rule' button is at the bottom left. At the bottom right of the dialog are 'Cancel' and 'Create' buttons.

Only SSH and HTTPS can connect to the management port

ISBN: 978-0-692-81971-5

©2017 Palo Alto Networks Fuel User Group, Inc.

Palo Alto Networks, PAN-OS, App-ID, Content-ID, User-ID, Aperture, AutoFocus, GlobalProtect, Panorama, Traps and WildFire are trademarks of Palo Alto Networks.

Version 1.0