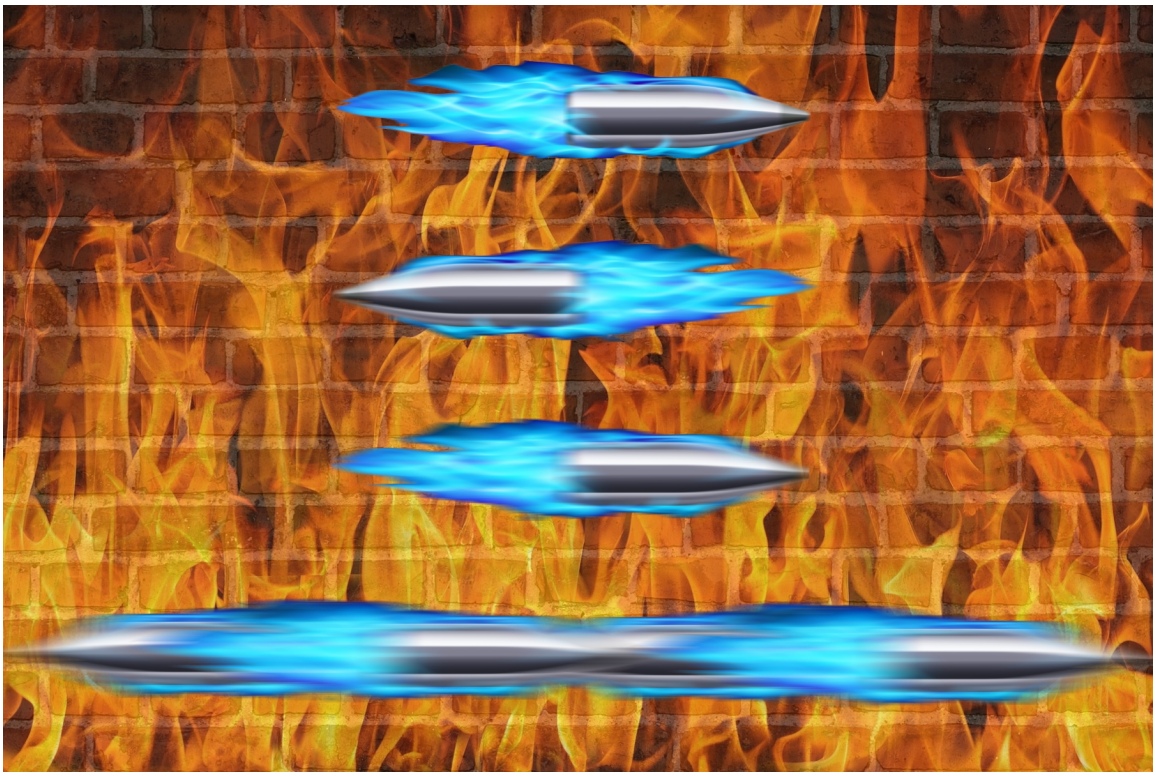

Max Power: Check Point Firewall Performance Optimization Addendum – 6/21/2017

Additional Tips &Tricks

R80.10 Gateway Updates

R80/R80.10 Management Supplement

Timothy C. Hall

Introduction

It has been just over two years since the release of my book *Max Power: Check Point Firewall Performance Optimization*. This third addendum will share with the Check Point community reader-submitted tips as well as other useful techniques and utilities I've discovered in the meantime. The additional content provided in this document is a roll-up of the two previous addendums, new tips and tricks I've run into during my consulting work since then, and updates for the new R80.10 gateway and management. All added or updated content from the previous addendum dated 4/11/2016 will be highlighted like this to clearly show what has changed.

The long-anticipated R80.10 release is for both Security Management Servers and Security Gateways/firewalls, so the first part of this addendum will provide additional tips, tricks & updates as well as note the performance-related differences in R80.10 gateway. In the second section, **R80/R80.10 Management Updates**, SmartConsole GUI differences will be documented so that the optimization steps described in *Max Power* can be performed from the R80/R80.10 SmartConsole or equivalent GUI tool.

Check Point has also updated their Check Point Security Administrator (CCSA) and Check Point Security Engineering (CCSE) classes for the R80.10 release, with the R80.10 version of Check Point Security Master (CCSM) coming soon. *Shameless plug alert*: I teach all official Check Point classes for my company Shadow Peak Inc, which are available both live online and classroom-led in Colorado, USA. **A limited number of seats per class are available for booking via Check Point's CO-OP system, which allows net new customers and Check Point partners to attend courses hosted by ATCs like Shadow Peak free of charge!**

The "R80 book" mentioned in the last addendum's introduction has morphed into "The CPUG Papers", please check out the following link for more information:

http://www.cpugpapers.com

Supplementary Material by Page Number

Page 16: If your site does not have the Monitoring Blade present, be sure to check out the **nmon** tool discussed in the Page 58 entry below. Check Point has explicitly disclaimed support for use of the **nmon** tool, see sk108122: Using the monitoring tool 'nmon' is not supported.

Page 21: If you are unlucky enough to be forced to utilize Emulex NICs (driver name be2net) on your open hardware firewall, be aware that a nasty firewall stability issue involving these NICs was fixed in R77.30 and R77.20 jumbo hotfix Take 94 and later. You'll definitely want to install this fix if using Emulex NICs on your firewall.

Page 26: The book recommended always using an even number of physical interfaces in a bonded aggregate Ethernet interface. After some reader questions I dug into it a little further, as this has been an unofficial recommendation floating around for quite some time. While I was not able to learn the exact nature of the issue, I was assured that it was an Intel driver issue and that it was fixed in R77.30. However of the four main Intel drivers shipped with Gaia R77.20 (e1000, e1000e, igb, ixgbe), only the e1000e driver was updated (from version 1.2.20 to 2.1.4) in the R77.30 release. So unless your firewall is using the e1000e driver (igb and ixgbe are by FAR the most common though) this recommendation does not appear to be valid. It is also possible that this recommendation is a bit of a myth, created by the fact that some networking vendors do not support using an odd number of physical interfaces when aggregating them using the older EtherChannel technique.

Based on the CPUG thread below it seems this is just a general recommendation rather than a strict requirement:
https://www.cpug.org/forums/showthread.php/20588-Amalgamating-Joining-Bonds

Page 34: One other potential STP-related issue pointed out by a student of mine is that different variants of the spanning tree algorithms don't mix well. As an example if two switches are connected together and one of them is using the original 802.1D standard STP and the other is using Rapid STP, the various timers will be radically different between the two and possibly cause network stability issues.

A reader did point out that newer versions of STP are supposed to go into a sort of "backward compatibility" mode when they detect an older version of STP present, but this should probably not be relied upon if at all possible.

Page 49: ICMP isn't just all about **ping** and **traceroute**; the various types and codes of ICMP datagrams can sometimes indicate that performance-impacting conditions are occurring within the network. Running a **netstat -s** on the firewall shows counters for how many different types of ICMP messages have been received by the firewall. Particular ones that can impact performance and be helpful to investigate further are:

- Fragmentation required but DF set (Type 1, Code 4)
- Precedence cutoff in effect (Type 1, Code 15)
- Source Quench (Type 4, Code 0) – very rare
- Redirect (Type 5)
- Time Exceeded (Type 11)

If nonzero values are noted for any of these in the **netstat -s** output, it is entirely possible they came from the Internet and you have no control over their generation. However seeing these types of ICMP datagrams arriving on the firewall's internal interfaces via **tcpdump** should be checked out. To display all ICMP traffic on an internal interface that is not associated with ping testing traffic, use this command:

tcpdump -eni (interface name) icmp and not icmp[0]=0 and not icmp[0]=8

Page 58: One additional built-in CPU profiling tool brought to my attention is **nmon:**

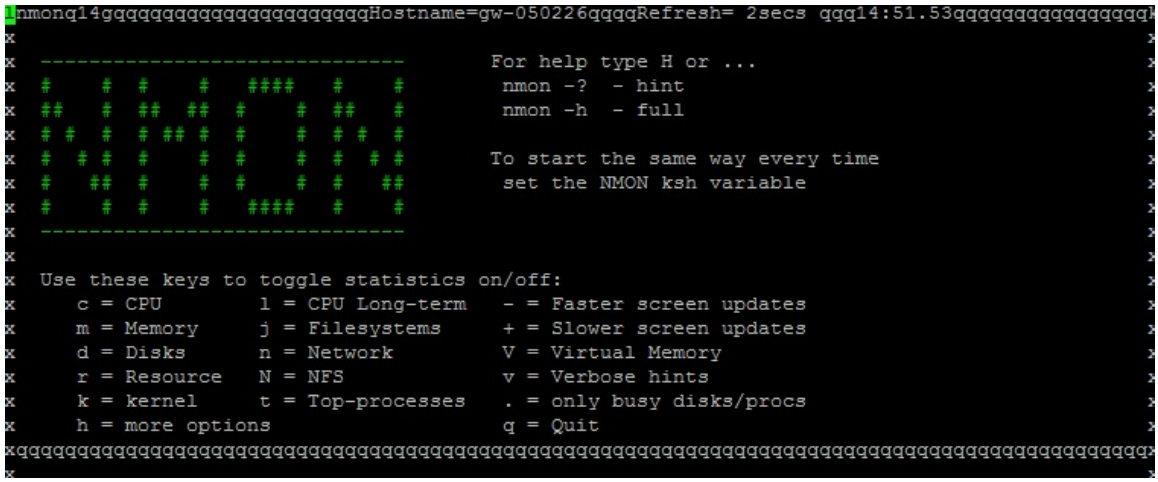

Added to Gaia in R76 it serves many of the same functions as the **top** command for monitoring CPU usage on the firewall but in a more graphical format. As an example, typing lowercase "L" provides a CPU usage graph; hitting "c" will show a graph for a for multi-core systems. While most of these CPU monitoring statistics are also available in **top**, a significant value-add of **nmon** is the ability to monitor and graph disk utilization which is very handy if the **wa** percentage shown by **top** is excessive. **Nmon** can also graph and monitor network interface activity, and serve as a useful stand-in for the **Traffic** and **System Counters** reports that are only available via the SmartView Monitor when a Monitoring Blade license is present. Thanks to Yasushi Kono of Arrow ECS for submitting this tip. Check Point has explicitly disclaimed support for use of the **nmon** tool, see sk108122: Using the monitoring tool 'nmon' is not supported.

Page 59: An easier way to see if cpwd has restarted any firewall processes since the last cpstart or firewall boot is to run **cpview** then select **Overview.** Down-arrow to the bottom of the page, and you will see the counter "# of monitored daemons crashes since last cpstart":

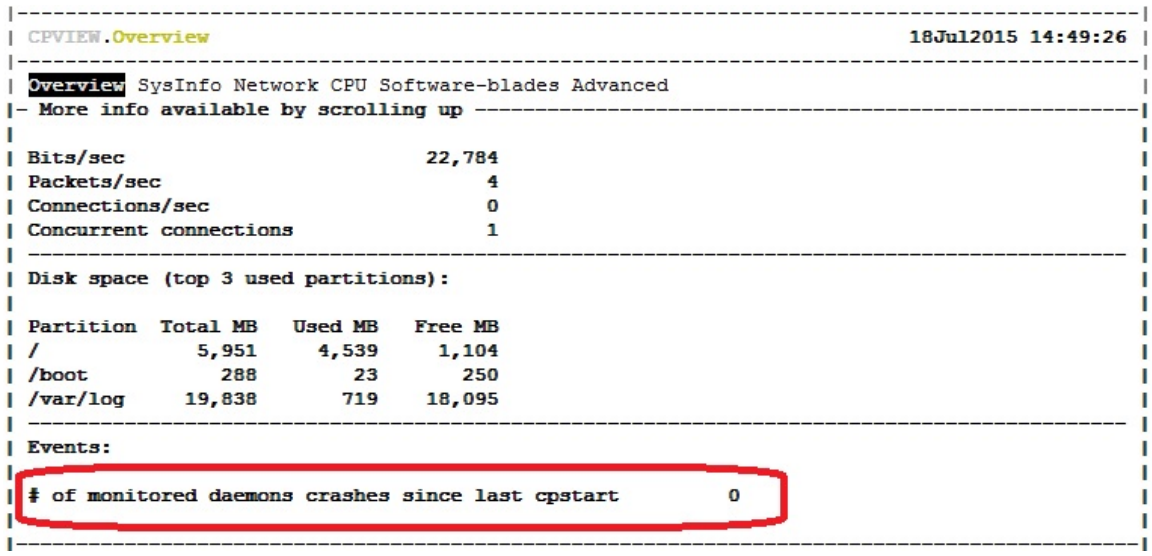

If this value is nonzero run **cpwd_admin list** to determine which daemon(s) are having a problem.

Pages 59-60: If while running **top** you notice a process called **kipmi0** consuming an excessive amount of CPU on an open hardware firewall, this is a known issue and you should consult sk104316: kipmi0 daemon consumes CPU at 100% on Open Servers running Gaia OS.

Page 69: If SecureXL is currently disabled on your firewall, you won't be able to use the "Top Talkers" script mentioned in the book to determine which internal IP address is hogging connection table slots. Use this command instead:

fw tab -u -t connections | awk '{ print $2 }' | sort -n | uniq -c | sort -nr | head -10

This will show the top ten source IPs consuming slots in the connection table in descending order, however you will need to manually convert the IP addresses displayed from hex to decimal like so: 0a1e0b53 = 10.30.11.83

For the top 10 destinations, substitute $4 for $2 in the awk command above.

Page 73-76: While the historical data reported by the **sar** command can be useful, it can definitely be overwhelming unless you have a good idea of what you are looking for. A picture is worth a thousand words, so if you find yourself staring down the barrel of hours and hours of poring through **sar** data trying to find a performance problem, consider the **ksar** tool. **ksar** can be used to graph **sar** data like this:

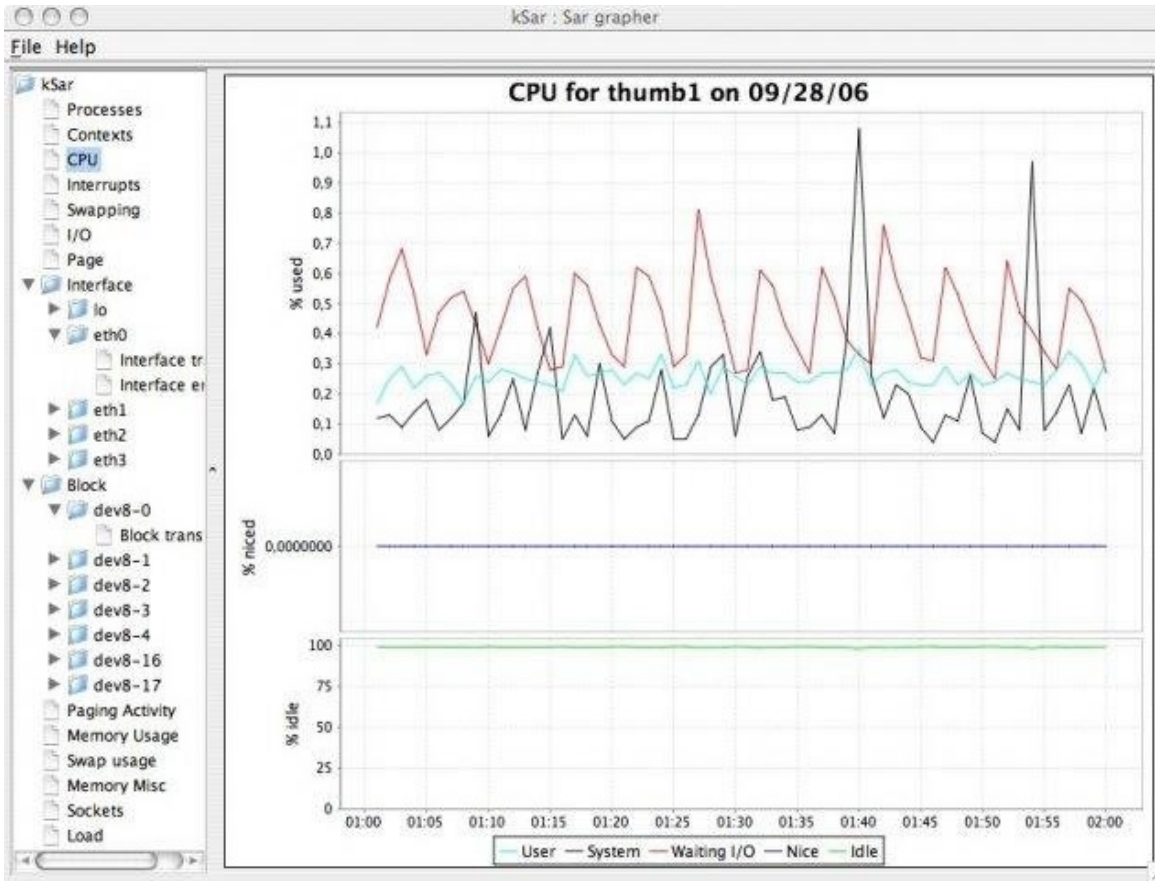

While not included with the Gaia operating system, **ksar** can be downloaded for free from https://sourceforge.net/projects/ksar/. If trying to figure out a past performance issue, keep in mind that the incredibly useful **cpview** tool can be launched in history mode with the **-t** option; 30 days of **cpview** historical data are available by default on the firewall.

Page 76: In addition to hitting "1" while running **top** to see individual core utilizations, the command **cpstat os -f multi_cpu** can also be used to obtain this information. Thanks to Yasushi Kono of Arrow ECS for submitting this tip.

Page 84: At the bottom of this page, a classic recommendation is made to put the rules with the highest hit counts towards the top of the rulebase. This recommendation is based on the so-called "top-down, first fit" rulebase matching algorithm that has remained mostly unchanged since Stateful Inspection was patented by Check Point in the early 90's. However in a major paradigm shift for R80.10 gateway, security policy layer evaluation in the Firewall path (F2F) has been significantly enhanced with something called "Columnar Matching" or "Column-based Matching". This new optimized matching implementation works with all types of policy layers such as ordered, unified, and inline; it is also valid for all Access Control and Threat Prevention layers.

Disclaimer: This is an emerging & undocumented feature of R80.10 gateway, and its implementation details are subject to change at any time. Some of the information presented here is based on assumptions and observations I have made that may not be 100% correct.

How can we best take advantage of this new R80.10 gateway feature to enhance rulebase lookup performance? ***Let me cut to the chase right here: Try to avoid using "Any" in all your policy layer rules for as long as possible, but most especially in the Destination field.***

Suppose your rule base starts like this:

No.	Name	Source	Destination	VPN	Services & Applications
Security Gateways Access (1-2)					
1	Administrator Access to Gateways	Admins	Corporate-GW	Any	Manage Services
2	Stealth rule	Any	Corporate-GW	Any	Any
VPN (3-4)					
3	Remote Access of employees to Data Center servers	Remote Access Users	Data Center LAN	RemoteAccess	Any
4	VPN between Internal LANs and Branch office LAN	Corporate LANs Branch Office LAN	Branch Office LAN Corporate LANs	Site2Site	Any
Access To Internet (5-6)					
5	Access to Internet according to Web control policy	InternalZone	Internet	Any	Any
6	block stupid attacks	attackers	Any	Any	Any

When a R80.10 gateway loads a new security policy package into the Firewall Worker Cores, it creates an indexed table (or perhaps a hash lookup table) to represent all fields used as matching criteria in the policy. For our sample rulebase, suppose the destination IP for a new connection is an internal (Non-Internet) address that does not match the destination of the first five rules. Prior to even looking at the rulebase for this new connection the gateway consults its Columnar Matching table, and determines that since it is impossible for this connection to match rules 1-5 based on destination IP address it can skip these rules and commence top-down, first-fit rulebase matching at rule 6. Why did it have to start at rule 6? Because rule 6 uses "Any" in its destination field!

While all fields of a policy layer that are considered matching criteria (Source, Destination & Service for example) can be potentially "skipped over" in this fashion, ***the Destination column is the first one examined by the firewall to see if any "skipping" can occur. Waiting as long as possible to use "Any" in the Destination field in your policy layers will maximize the number of rules that can be skipped over, and can significantly reduce policy lookup overhead (and therefore CPU utilization) on the Firewall Worker Cores.*** Rules that use object "Internet" and/or negations do not affect the Columnar Matching feature and can still potentially be "skipped over". Note that this feature is completely separate from SecureXL connection templating, however SecureXL certainly received its share of enhancements for R80.10 gateway (see Pages 239-240 notes for more details).

Page 89: As stated in the book, setting **fw_rst_expired_conn** to 1 should always be tried first to gracefully terminate application-based connections that aren't closing properly and impacting perceived application performance. In some cases however this will not fully remediate the situation, and you be forced to go one step further with this: **fw ctl set int fw_reject_non_syn 1.** A classic example of an application that requires this firewall setting is SAP HANA traffic. This setting also handles client port reuse out of state errors when RST packets from the server to the clients get lost (e.g. due to policy install or packet loss).

Bear in mind however that this setting is quite likely to make your friendly auditor/penetration tester upset with you, since the firewall will now issue a TCP RST for ***all*** received packets that are out of state and have the ACK flag set. An auditor running a TCP ACK nmap scan will have it light up like a Christmas tree with tens of thousands of ports showing up as filtered instead of closed. For this reason, using this setting is generally not recommended on an Internet perimeter firewall but may be acceptable on internal firewalls. Thanks to Andrew Craick of Dimension Data for submitting this tip.

Page 90: I found out the hard way with a customer that the TCP State Logging function was introduced in R77, and is not available on firewalls running older versions of code. An alternative to this feature on pre-R77 firewalls is using the **Account** option in the **Track** column of the rule matching the problematic traffic. When this option is set for a rule, an Accept entry is created at the start of the connection just as it is when the **Track** is set to **Log**. However once the connection finishes (FIN, RST, idle time out, etc.) or every 10 minutes if the connection exists that long, the existing log entry is converted from a **Log** type to an **Account** type.

While the Accounting option can be used to infer connection and session behavior as described in the book, new to R80.10 management is a feature called "Session Logging". When enabled this feature will correlate multiple individual connections into a "session" that has additional logging information. See the following for more information: Infinity R80.10 "Cool Feature of the Day" - *Session logging*

Note that in the R80/R80.10 SmartConsole the Account option is now enabled by selecting the "Accounting" checkbox in a rule's Track column. On security gateway/cluster objects with a version of R80.10 or later, configuring TCP State Logging can now be performed directly from the SmartConsole as shown here:

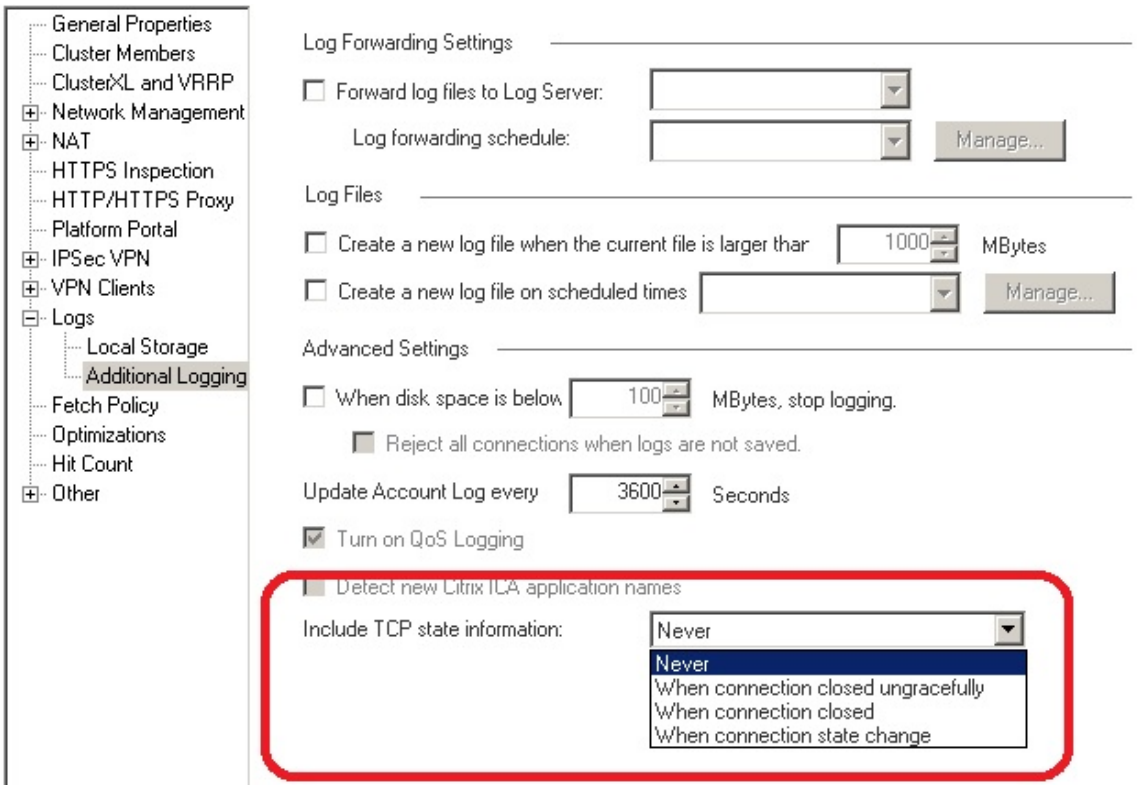

Page 97: R77.30 has added the ability to set the "Magic MAC" value via the Gaia web interface instead of by hand-editing the fwkern.conf file. During the firewall's post-installation dialog in the Gaia web interface if "Unit is part of a cluster" is checked, the new field "Cluster Global ID" will become editable:

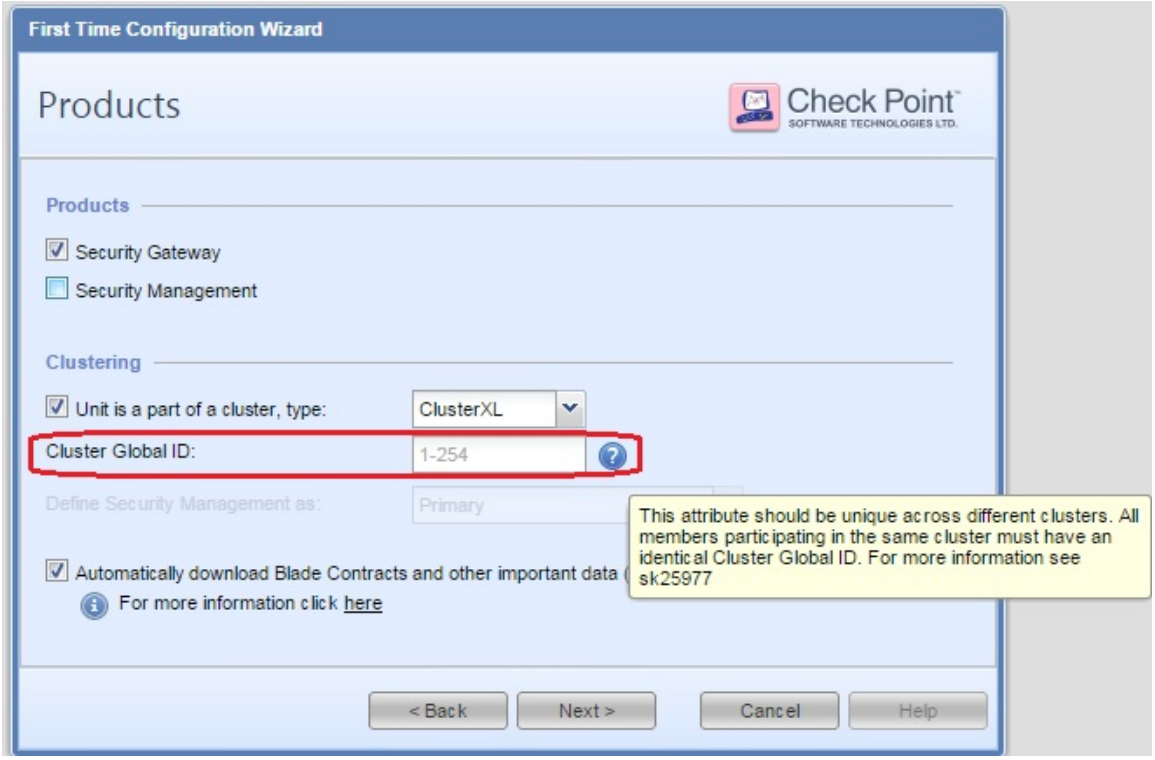

The Cluster Global ID should be set identically on all members of the same cluster, but be a different value for different clusters. R80.10 gateway clusters use a new feature called "Automatic MAC Magic" by default to automatically derive a unique Cluster Global ID, and prevent conflicts with other gateway clusters on the same network. The status of this new feature can be checked with the **cphaprob mmagic** command. This feature can also be monitored from a new ClusterXL-based screen of the **cpview** tool on a R80.10 gateway under **Advanced...ClusterXL**, and is backward compatible with gateways that had their Cluster IDs configured manually in earlier versions.

Pages 98-99: Other good preexisting SKs for troubleshooting unexpected ClusterXL failovers are: sk62570: How to troubleshoot failovers in ClusterXL - Advanced Guide and sk56202: How to troubleshoot failovers in ClusterXL.

When trying to troubleshoot intermittent ClusterXL failovers, establishing exactly when failovers happened and whether they occurred in any kind of recognizable time interval can be quite time-consuming. However an undocumented clish command option introduced in R77.30 was recently brought to my attention by Gary Lipets (gary.lipets@gmail.com) that provides a an easy way to see when failovers occurred: **show routed cluster-state detailed**. This command is valid even if you are not utilizing any kind of dynamic routing protocol on the firewall itself and only static routes are in use. Routed is tracked as a pnote in all ClusterXL clusters, and it actively logs all cluster state changes:

```
gw> show routed cluster-state detailed

Cluster: Clustered
Master/Slave: Master
Sync IP: 192.168.212.73
Cluster Sync: Synchronized
Last Sent: INIT_STATE_FIN
Last Received: MASTER_INIT_STATE_REQ

Cluster VIPs
(long list of VIPs deleted)

Cluster Members: 2

Member ID Member Address
2 192.168.252.71

Cluster State Change History
Timestamp State Change Type
Jun 5 14:44:05 Slave to Master
Jun 5 14:44:00 Master to Slave
Jun 5 13:23:11 Slave to Master
Jun 5 13:23:03 Master to Slave
Jun 5 12:56:57 Slave to Master

Cluster Routed Pnote Change History
Timestamp Routed Pnote Event Description
Jun 5 14:44:05 OK Slave State Fin [OK]
Jun 5 14:44:00 PROBLEM DR Enabled; Master To Slave [Problem]
Jun 5 13:23:09 OK Slave State Fin [OK]
Jun 5 13:23:03 PROBLEM DR Enabled; Master To Slave [Problem]
Jun 5 12:56:55 OK Slave State Fin [OK]
Jun 5 12:56:50 PROBLEM DR Enabled; Master To Slave [Problem]
```

Page 136: The maximum number of total cores supported by CoreXL has been raised from 30 cores to 40 cores for R80.10 gateway. The limit of 10 Firewall Worker instances for VSX R77.30 and earlier has been lifted in R80.10, see the following for details: sk117375: How to add more than 10 CoreXL FW instances to a Virtual System R80.10.

Page 139: Some additional commands to check CoreXL licensing status are:

```
[Expert]# fw ctl get int fwlic_num_of_allowed_cores
fwlic_num_of_allowed_cores = 8
[Expert]# fw ctl get int fwlic_num_of_allowed_cpus
fwlic_num_of_allowed_cpus = 8
```

Thanks to Yasushi Kono of Arrow ECS for submitting this tip.

Pages 141-148: Shortly after Max Power was published, an all-new Advanced Technical Reference Guide (ATRG) for VPNs was created that includes some performance-related information: sk104760: ATRG: VPN Core. A bit lengthy, but highly recommended reading if you work with and/or troubleshoot Check Point VPNs on a regular basis!

Pages 141-148: A new SK for VPN Performance Best Practices has been created by Check Point: sk105119: Best Practices - VPN Performance. Very similar to what was presented in the Max Power book with a few extra tidbits you may want to check out.

Pages 141 & 146: On these pages it was mentioned that SecureXL can accelerate some IPSec VPN encryption/decryption operations. If SecureXL is enabled on your firewall and you'd like to check if this is occurring run **fwaccel stats**. Nonzero or rapidly incrementing values in the **Accelerated VPN Path** section of the output indicate that SecureXL acceleration of IPSec traffic is occurring. The SHA-384 hash algorithm has not yet been implemented in the R77.30 and R80.10 gateway SecureXL Accelerated Path

code. Any VPN traffic verified using this algorithm will be ineligible for encryption/decryption in the Accelerated Path, and be forced into the Firewall Path (F2F) on the lead (lowest-numbered) Firewall Worker core for processing on pre-R80.10 firewalls.

In addition to the SHA-384 hashing algorithm, two other algorithms not implemented in the Accelerated Path for R77.30 and R80.10 gateway are AES-GCM-128 and AES-GCM-256 encryption. While the Galois/Counter Mode (GGM) versions of these AES algorithms combine integrity and encryption into a single operation that requires less CPU (and is even able to be accelerated directly in hardware with AES-NI), all traffic subject to these two GCM-based algorithms will be ineligible for processing in the Accelerated Path and be forced into the Firewall Path (F2F) on the lead (lowest-numbered) Firewall Worker core for processing on R77.30 gateway. To be clear, the generic versions of the AES-128 and AES-256 algorithms **are** eligible for processing in the Accelerated Path.

Page 144: A more graceful way to check the status of AES-NI on your firewall is by running the undocumented command: **sim enable_aesni**

Page 144: Another new SK extolling the virtues of AES instead of the 3DES encryption algorithm has been created, and provides tangible AES-NI performance improvement numbers for the various firewall appliances: sk98950: Slow traffic speed (high latency) when transferring files over VPN tunnel with 3DES encryption

Pages 144-146: Multi-core IPSec VPN processing is finally supported on R80.10 gateway. Previously only available as a special hotfix for R77.20 (but not R77.30), the multi-core IPSec VPN feature is enabled by default in R80.10 and alleviates the lead Firewall Worker Core IPSec VPN processing bottleneck described on these pages. If this value has been manually set to 1 on your firewall as described in the book, the kernel variable **fwmultik_dispatch_skip_global** should probably be set back to 0 after upgrading your gateway to R80.10.

Page 148: On R80.10 gateway, CoreXL is now fully supported with route-based VPN (i.e. using VPN Tunnel Interfaces (VTIs)) instead of only with domain-based VPNs.

Pages 149-151: I'm pleased to report that R77.30 has added the option to substantially improve Firewall Worker Core load distribution via the new Dynamic Dispatcher Feature (sk105261: CoreXL Dynamic Dispatcher in R77.30). This new Firewall Worker Core load-balancing feature is disabled by default in R77.30 but enabled by default on R80.10 gateway; as a general rule of thumb you should consider enabling this feature when the following conditions are present (but note the warnings for R77.30 below):

- Firewall has 6 or more total cores
- Firewall Worker CPU loads consistently vary from each other by >10% **
- Firewall is NOT using a SAM card (i.e. 21000 series)

** Keep in mind that all IPSec VPN and VoIP traffic can only be processed on the lead (lowest-numbered) Firewall Worker Core as specified on page 141 (this is a R77.30 limitation that is no longer present in R80.10 gateway). If there is substantial IPSec and/or VoIP traffic traversing a R77.30 firewall, exclude the lead Firewall Worker Core from consideration when applying the 10% rule of thumb above.

On R80.10 gateway, the Dynamic Dispatcher and Priority Queuing are enabled by default after both a fresh installation or upgrade of the gateway. Use the command **fw ctl multik dynamic_dispatching get_mode** to check the status of these features on R80.10.

If planning to enable the Dynamic Dispatcher on R77.30, ensure that you have loaded the latest Generally Available (GA) jumbo hotfix before enabling the Dynamic Dispatcher on R77.30. The initial Dynamic Dispatcher R77.30 code suffered from various issues such as the ones below that were rectified in the subsequent R77.30 jumbo hotfixes:

- sk108432: Issues with traffic passing through Security Gateway with CoreXL Dynamic Dispatcher enabled
- sk108856: R77.30 cluster member might go Down after disabling CoreXL **Dynamic Dispatcher** only on one member
- sk108894: Difficulties in connecting to untrusted sites when both HTTPS Inspection and CoreXL **Dynamic Dispatcher** are enabled
- sk106665: VoIP traffic, or traffic that uses reserved VoIP ports is dropped after enabling CoreXL **Dynamic Dispatcher**

Pages 159-161: These pages discuss the case for re-enabling SecureXL if it has been disabled on your firewall due to application compatibility issues. Not sure how the heck I missed such a useful tip when preparing the original Max Power book, but there is a technique that **allows SecureXL to be selectively disabled** for certain IP addresses in R77 and later: sk104468: How to disable SecureXL for specific IP addresses

It involves a table.def change with an *f2f_addresses* directive that can be made active with a simple firewall policy push. All traffic matching the IP addresses specified in this directive will always be sent to the Medium/Firewall Paths for processing. There is even a hotfix available for pre-R77 firewalls to implement this functionality as well! Unbelievably useful in environments where SecureXL and all its benefits has to be disabled just to accommodate that one pesky system or application!

Page 162: When attempting to re-enable SecureXL with IPSec VPNs present, watch for this issue: sk102742: When SecureXL is enabled, traffic through the VPN trusted interface is sent encrypted instead of clear. A separate hotfix must be obtained (this fix does not appear to be included in the current R77.20 jumbo hotfix) or upgrade to R77.30.

Page 163: While disabling SecureXL with the **fwaccel off** command and re-enabling it with the **fwaccel on** command might only cause some minor firewall performance degradation, in rare cases it can cause noticeable impacts to production traffic.

The most common scenario for temporarily disabling SecureXL is to perform a packet capture using the tool **fw monitor**. This tool can only capture traffic traversing the Firewall Path and cannot always "see" traffic passing through the Medium or Accelerated Paths. Disabling SecureXL forces all traffic into the Firewall Path and ensures a complete **fw monitor** capture.

But there is a better way using a standard Linux tool: **tcpdump**. This tool is immune to the state of SecureXL and can "see" traffic in all three paths (SXL/PXL/F2F) with one exception: if SecureXL acceleration is being performed with dedicated hardware such as a SAM or Nokia ADP card, tcpdump cannot "see" the traffic and SecureXL must indeed be disabled in that particular case to ensure a complete capture. Another advantage of **tcpdump** is the capability to capture the Layer 2 headers for analysis (including MAC addresses), whereas **fw monitor** cannot. However be warned that attempting to use **tcpdump** on a SAM-accelerated interface of a 21000 series firewall may lead to a large number of packet drops.

If you are unable or unwilling to use **tcpdump** in lieu of **fw monitor**, and will need to frequently toggle the state of SecureXL while passing production traffic, consult the following SKs for hotfixes that should be installed to prevent possible (but quite rare) traffic disruptions from occurring:

- sk106934: Security Gateway might crash when disabling and re-enabling SecureXL
- sk109468: Connections are broken for short time after disabling SecureXL, or after installing a policy

Page 169: While **fwaccel stats -s** provides useful acceleration packet counters showing total number of packets processed by the SXL/PXL/F2F processing paths, you can also view live throughput numbers for each of the three paths in expressed in pps and Mbps. Run **cpview** then select Advanced...Network...Path.

Page 172: If you are utilizing 21000-series appliances equipped with a Security Acceleration Module (SAM) card, reading through the following two all-new SKs to understand the capabilities and specific optimization strategies for the SAM card is highly recommended: sk107157: ATRG: Security Acceleration Module (SAM) card and sk93036: Known Limitations of Security Acceleration Module (SAM) on 21000 Appliance.

Page 173: There are a plethora of stability fixes for the 21000 firewall units that utilize the SAM card in R77.30. If using a SAM card upgrading to R77.30 (with the latest jumbo hotfix) or at least loading the latest R77.20 jumbo hotfix is highly recommended.

Page 176-178: *Correction:* Changing the IPS Scope setting from "Perform Inspection on all Traffic" to "Protect internal hosts only" does NOT potentially make more traffic eligible for the Accelerated Path. Setting "Protect internal hosts only" has a similar effect to creating an IPS Exception in that it can save CPU time in the Medium Path (PXL). So while changing this setting does have a positive impact on performance (by potentially saving CPU time in the Medium Path), it is not for the reason originally stated in the book (that more traffic is made eligible for Accelerated Path). However there is one exception to this: on a 21000 series firewall equipped with a SAM card, traffic matched by an IPS Exception is eligible to be fully accelerated by SecureXL in the SAM card itself, assuming the traffic is not subject to inspection by another blade that requires Medium or Firewall Path processing. Configuration techniques that ensure as much traffic as possible can be accelerated by the SAM card are described in sk94484: Accelerating traffic with the Security Acceleration Module (SAM) while also using non-accelerated blades

Page 194: This section of the book spends a great deal of time trying to reduce firewall CPU load on the Firewall Worker Cores, most of which occurs in the Medium Path (PXL) on the vast majority real-world firewalls. R77.30 and later has introduced an exciting ability to view the *top connections by CPU usage*. This capability is a subset of the new R77.30 Firewall Priority Queues feature (sk105762: Firewall Priority Queues in R77.30), and the good news is that this helpful information can be obtained without having to fully enable this feature. To obtain this ability run the following command: **fw ctl multik set_mode 1** and reboot the firewall. Now when running **cpview** select CPU...Top Connections to see the top individual connections by CPU consumption. Priority Queues are enabled by default on a R80.10 gateway, regardless of whether the gateway was fresh-loaded or upgraded.

Page 195-204: R80.10 gateways now have the IPS function fully integrated into the Threat Prevention policy layer. Instead of having to configure specific IPS Exceptions, the Threat Prevention policy layer columns "Protected Scope", "Source", and "Destination" can be used to explicitly specify which IPS profile should be applied to network traffic matching these columnar criteria. Exceptions can still be added to the Threat Prevention policy layer, but they can potentially apply to the entire Threat Prevention policy, not just IPS. Some items that were previously listed as IPS Protections are now separated into the "Inspection Settings" portion of the R80/R80.10 SmartConsole. Changes to these "Inspection Settings" values require a Network Access policy reinstallation to take effect because they have been separated from the Threat Prevention policy layer.

Page 208: The book indicates that **fw ctl zdebug drop** can be used to determine what non-logged IPS signatures are inappropriately dropping traffic. This is not completely accurate, because the reason for the drop shown by zdebug by default will be very generic, and simply indicate it had something to do with IPS enforcement.

Warning: The following procedure will substantially increase the size and memory requirements for enforcing the compiled policy on the firewall. Use with caution on production systems.

To obtain the actual IPS signature name in the zdebug output, launch the SmartConsole tool GUIdbedit and under Table...Global Properties...Properties change variable **enable_inspect_debug_compilation** from **false** to **true** and reinstall policy to the firewall. This setting will cause additional debug information to be compiled into the firewall's policy, such that the actual offending IPS signature name can be displayed in the zdebug output.

Page 213: If the **Website Categorization Mode** has been set to **Hold** as recommended in the book, and an unacceptable level of latency is encountered categorizing websites for the URL Filtering function, additional statistics can be enabled in the Resource Advisor Daemon (RAD). The RAD process handles interaction between the firewall and the Check Point cloud for dynamic lookups of content such as URLs. Note that this daemon is also used to update signatures and verify content for the Application Control, Anti-Malware, and Anti-Virus software blades; therefore statistics are available for these other three functions as well. To enable statistics for the URL filtering function specifically, execute the command **rad_admin stats on urlf**. To view URL caching and cloud interaction statistics run **cpview** and select Advanced...RAD:

```
|-------------------------------------------------------------------------------|
| CPVIEW.Advanced.RAD                                    19Jul2015 14:04:23     |
|-------------------------------------------------------------------------------|
| Overview SysInfo Network CPU Software-blades Advanced                         |
|-------------------------------------------------------------------------------|
| CPU-Profiler Memory Network SecureXL CoreXL PrioQ Streaming RAD               |
|- More info available by scrolling up ---------------------------------------|
| Name                              APPI        AB        AV       URLF       |
| Found in LDB                       N/A       N/A       N/A       N/A       |
| Sent to Site                       N/A       N/A       N/A       N/A       |
| Round Trip (ms)                    N/A       N/A       N/A       N/A       |
| Hit Count                          N/A       N/A       N/A       N/A       |
| Miss Count                         N/A       N/A       N/A       N/A       |
| Error Count                        N/A       N/A       N/A       N/A       |
| Cache Size (bytes)                 N/A       N/A       N/A       N/A       |
| Max Cache Size (bytes)             N/A       N/A       N/A       N/A       |
| Cache Total Host Records           N/A       N/A       N/A       N/A       |
| Max Cache Total Host Records       N/A       N/A       N/A       N/A       |
| Avg Family Size                    N/A       N/A       N/A       N/A       |
| Max Family Size                    N/A       N/A       N/A       N/A       |
| Expired Requests                   N/A       N/A       N/A       N/A       |
|                                                                               |
|-------------------------------------------------------------------------------|
```

Don't forget to turn off the statistics gathering with the **rad_admin stats off urlf** command when finished!

Page 220: Check Point has created a new SK for HTTPS Inspection Best Practices which includes some performance-related information, see sk108202: Best Practices - HTTPS Inspection if you are utilizing the HTTPS Inspection feature. Just as in R77.30 and earlier, the bulk of HTTPS Inspection on R80.10 gateway takes place in the Firewall Path (F2F) and is not eligible for acceleration. Also keep in mind that the initial HTTPS handshake and key calculation is performed in process space by the **wstlsd** and **pkxld** daemons on the firewall. If the firewall is experiencing extremely high CPU usage in kernel space (sy, si, hi, wa) on all cores as shown by the **top** command, manual process affinity for these daemons may need to be configured to ensure they receive adequate CPU slices.

Another example I've seen of possible process space CPU starvation on a oversubscribed firewall is the Policy Decision Point Daemon (pdpd) component of Identity Awareness, which maintains the WMI connections to Active Directory controllers and associated monitoring of Security Log events for AD Query. If the

single-threaded pdpd process cannot obtain enough CPU slices due to a busy firewall kernel, it can begin to fall behind in timely processing of Security Log events in a large Windows domain, and formation of IP to user mappings will not occur in a timely fashion. I'm pleased to report that Check Point has provided a new feature known as the Identity Collector to help avoid this situation. This is the ability to essentially "outsource" the extensive overhead of parsing AD Security Logs via WMI to a software agent running on any Windows system in the domain. The Identity Collector agent software does not need to be loaded on the Windows domain controller itself. For more information about the R80.10 Identity Collector see sk108235: **Identity Collector - Technical Overview**.

Just to be clear, firewall operations that are executed by a process instead of in the kernel are not necessarily a bad thing. Most typically these process-based operations are handling what I would term "potentially unsafe" situations. A classic example is IKE Phase 1 and Phase 2 negotiations being performed by the **vpnd** daemon on the firewall, while IPSec encryption and decryption occurs in the kernel via the SXL or F2F paths. If a hostile IKE peer tries to crash or flood the firewall with malformed IKE negotiations, the crash occurs in the **vpnd** daemon which is immediately and gracefully restarted by the **fwd** parent process, and system stability is ensured. If IKE negotiation operations were handled directly in the firewall kernel, a hostile IKE peer could potentially crash or significantly impair the entire firewall. This IKE scenario is a particularly good example since IKE peer authentication via a pre-shared secret or certificates happens very late in IKE Phase 1, after protocols have been negotiated and a computationally expensive Diffie-Hellman key calculation has been performed.

Pages 220-221: The HTTPS Inspection feature was significantly enhanced in R77.30 and later. While many of the relevant fixes are included in the R77.20 jumbo hotfix, it appears that there are many enhancements exclusive to R77.30 that can improve the functionality and performance of the HTTPS Inspection feature. While the bulk of HTTPS Inspection operations appear to still occur in the Firewall Path, the firewall

performance impact of Bypass actions and SSL negotiation have been substantially improved.

Page 224: The ISP Redundancy feature is well-known for forcing almost all traffic into the Firewall Path, even traffic that is not involved with the External interfaces leading to the redundant ISPs. However loading the R77.30 Jumbo Hotfix (take 15+) will permit acceleration of firewall traffic when ISP Redundancy is used in Primary/Backup mode. (Note that if ISP Redundancy's Load Sharing mode is selected, almost all traffic will still go F2F) See: sk104679: SecureXL Accept Templates not created when ISP Redundancy is enabled in Primary/Backup mode.

Page 234: Alternatively, to view the firewall's New Connection Rate (Connections/sec) from the CLI, run the **cpview** command and select **Network**.

Page 239-240: I'm pleased to announce that you will be FAR more likely to see your entire rulebase become eligible for SecureXL templating in R80.10 than in earlier firewall versions as shown by **fwaccel stat**, due to the new R80.10 NMR and NMT templates. On R80.10 gateway the following objects being present in a rule will no longer halt templating eligibility for the remaining rules:

- Domain objects
- Time objects
- Dynamic objects
- **traceroute** service
- **dhcp-request** service
- **dhcp-reply** service

The only object types that seem to still halt templating eligibility in R80.10 gateway policies are listed below, so try to move rules utilizing these services towards the bottom of your rulebase if possible:

- RPC/DCOM/DCE-RPC services
- Services of type Other with a custom matching condition
- Certain rare complex services (i.e. **http_mapped**, **ftp_mapped**)

The best part is that these new templating abilities happen by default on R80.10 gateway, and you don't need to do anything whatsoever to enable them. However just because your entire rulebase is much more likely to be eligible for templating in R80.10 gateway, the rulebase optimization techniques described on pages 240-244 should still be followed for best performance, and to reduce CPU load caused by full policy lookups on the Firewall Worker Cores for non-templated connections.

Page 256: To quickly check if the IPS Aggressive Aging feature is currently expiring connections early due to excessive firewall memory or connection table utilization, run

the command **fw ctl pstat** on the firewall and look under the **System Capacity Summary** section of the output.

Page 258-262: A very nice complement to the SecureXL-friendly blocking capabilities of the **fw samp/sim_dos** commands described in the book is the ability to dynamically receive a real-time list of blacklisted IP addresses from the Check Point cloud, and have your firewall efficiently block them in the SecureXL Accelerated Path. This is quite similar to the old Dshield.org "Storm Center" capability but imposes much less performance overhead. On R77.30 this feature is accessed via the **ip_block.sh** script. See sk103154: How to block traffic coming from known malicious IPs for more details about this little-known feature.

Page 270: In R80.10, Multi-Queue supports a new network driver in addition to igb and ixgbe: the Mellanox mlx5_core 40Gb driver which is used in the 15000 and 23000 series of appliance hardware. The Intel l211 driver used for the on-board NIC interfaces on the 3200/5000/15000/23000 appliance models are supported as well, but watch out for this known issue: sk114625: Multi-Queue does not work on 3200 / **5000** / 15000 / 23000 appliances when it is enabled for on-board interfaces.

Page 275-276: I'm pleased to report that R77.30 has an available built-in fix for the Hide NAT port allocation failures that are much more likely to occur when Hyperspect is enabled as discussed in #8. Ports used for Hide NAT source port reallocation can be dynamically pooled among the Firewall Worker Cores, instead of being statically assigned. This new feature is not enabled by default. It involves setting the **fwx_nat_dynamic_port_allocation** variable from 0 to 1. There is a separate hotfix available for R77.20 to add this functionality, however it does not appear to be a part of the R77.20 jumbo hotfix. See sk103656: Dynamic NAT port allocation feature for more details. In R80.10 gateway if there are 6 or more Firewall Worker Cores configured, by default the **fwx_nat_dynamic_port_allocation** variable will be automatically set to 1, otherwise it will still be 0.

Page 282: If performing lab benchmarking of Check Point firewalls on R77.30, be sure to enable the following feature: sk105261: CoreXL Dynamic Dispatcher in R77.30 (howver heed the warnings stated in the Pages 149-151 section above). The Dynamic Dispatcher is enabled by default on R80.10 gateway. Network load-testing traffic is infamous for its non-uniqueness, which can cause an imbalance of Firewall Worker Core loading and severely crimp firewall throughput results. Also if performing benchmarking of HTTPS Inspection on a Check Point firewall, be sure to enable HTTPS Inspection in "Test Mode" as detailed here: sk104717: HTTPS Inspection Enhancements in R77.30. HTTPS Inspection Test Mode compensates for similar quirks in HTTPS load-testing traffic and ensures accurate performance results.

Page 283: If you've reached this section of the book and can't obtain acceptable performance from your firewall despite following all the tuning recommendations, and no immediate relief is in sight in the form of newer faster hardware, consider employing this new R77.30 feature discussed in the Introduction to make the most of what you do have: sk105762: Firewall Priority Queues in R77.30. Priority Queues is enabled by default on R80.10 gateway.

R80/R80.10 Management Updates

Page 16: The SmartConsole can directly provide the same Traffic/System Counters functionality as SmartView Monitor, by right-clicking a firewall object on the "Gateways and Servers" tab and selecting "Monitor".

Pages 77-82: Virtual Links can only be configured from the legacy SmartDashboard which is accessible from the SmartConsole by selecting "Manage and Settings...Blades...Configure in SmartDashboard". The "Virtual Link" SmartView Monitor report is still available in the R80 SmartConsole by right-clicking the firewall object and selecting "Manage...Traffic".

A reader pointed out that the Virtual Links feature covered on these pages is incompatible with ClusterXL, in that a cluster object will never be selectable as a valid End Point when creating a new Virtual Link object. Also any gateways with a version of R77 or higher (whether clustered or not) will also not be selectable as a valid End Point when creating a new Virtual Link object, due to an issue with the R77.30 SmartDashboard. For the fix see: sk106085: When configuring Virtual Link in SmartDashboard, the End point combo boxes are blank.

Page 83: To provide the functionality for firewall thresholds as described in the book, the SmartView Monitor will need to be invoked directly from the SmartConsole. From the "Logs & Monitor" tab create a new tab with the "+" button, and then select "External Apps...Tunnel & User Monitoring" as shown here:

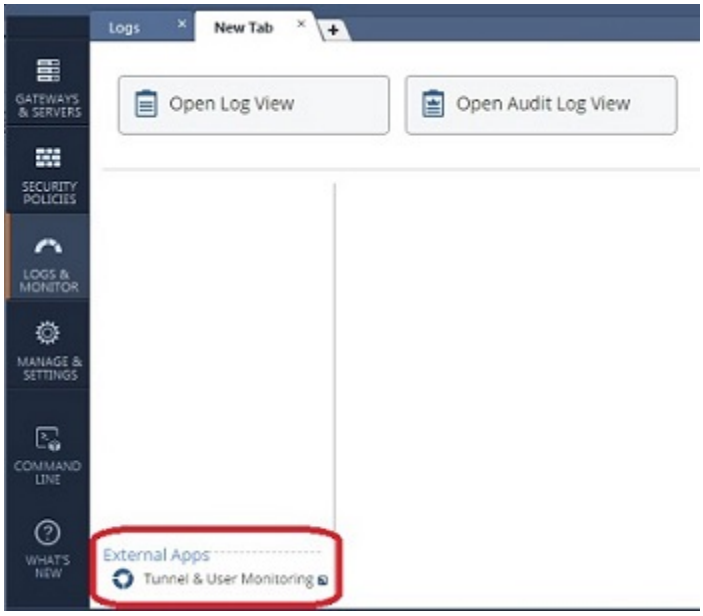

Pages 97-98: To filter for Control events on the "Logs and Monitor" tab of the SmartConsole as described in the book, the proper search field syntax is "type:Control".

Page 180: In the SmartConsole, IPS Profiles are no longer directly assigned on the firewall object; this action is performed in the Threat Prevention policy alongside Anti-virus, Anti-bot and Threat Emulation. However for pre-R80 firewalls, IPS Profiles are assigned in a special IPS "rulebase" that becomes visible if you have at least one pre-R80 firewall defined:

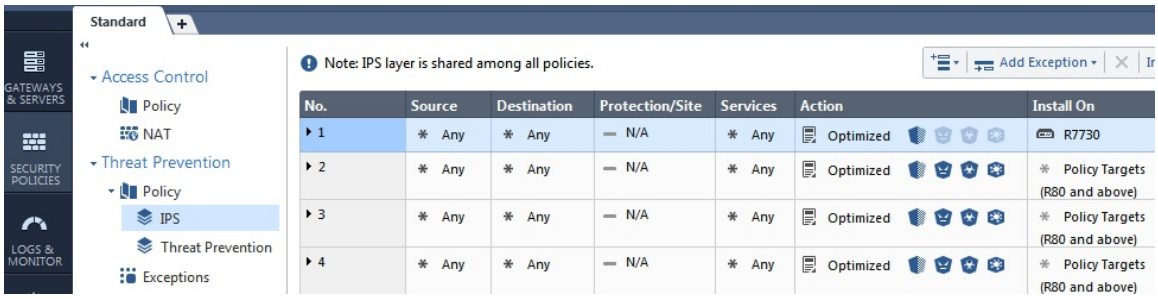

While this screen may look like a typical security policy, you'll rapidly find that it is not once you start working with it. For example, the Source, Destination, Protection/Site, and Service fields cannot be edited at all for pre-R80 gateways! Should you wish to take

more granular control of how IPS protections are applied to network traffic as described in the book, you are limited to the following:

- In the properties of the gateway object, IPS screen, you can select "Protect internal networks only" or "Perform inspection on all IPS traffic". As stated in the book, the former setting will only apply IPS protections against traffic that is heading to a non-External (i.e. Internal) interface of the gateway as defined in the gateway's topology. The latter setting will apply IPS Protections to all traffic regardless of where it is headed. Which direction the inspected connection was originally initiated (i.e. inbound or outbound) does not impact how this setting is applied.
- You can define IPS Exceptions as described in the book for a single IPS Profile or multiple profiles. With an IPS Exception rule, traffic matched by Source, Destination, and/or Service can be excluded from all IPS enforcement or a subset of IPS enforcement. The subset could be a single protection or multiple protections.

For pre-R80 gateways the IPS "policy" in the R80 SmartConsole is really just a place to define which IPS Profile is assigned to a gateway, and to create IPS Exceptions. R80.10 firewalls have their IPS settings consolidated in the much more flexible main Threat Prevention policy, along with all the other Threat Prevention features Anti-Virus, Anti-bot, and Threat Emulation.

R80.10 gateways now have the IPS function fully integrated into the Threat Prevention policy layer. Instead of having to configure the IPS Protection Scope and/or specific IPS exceptions, the Threat Prevention policy layer columns "Protected Scope", "Source", and "Destination" can be used to explicitly specify which IPS profile should be applied to network traffic matching these columnar criteria. Some items that were previously listed as IPS Protections are now separated into the "Inspection Settings" portion of the R80/R80.10 SmartConsole. Changes to these "Inspection Settings" values require a Network Access policy reinstallation to take effect because they have been separated from the Threat Prevention policy layer.

Page 181: IPS Signatures are accessed in the SmartConsole by clicking "IPS Protections" under Threat Tools on the "Security Policies...Threat Prevention...IPS" screen:

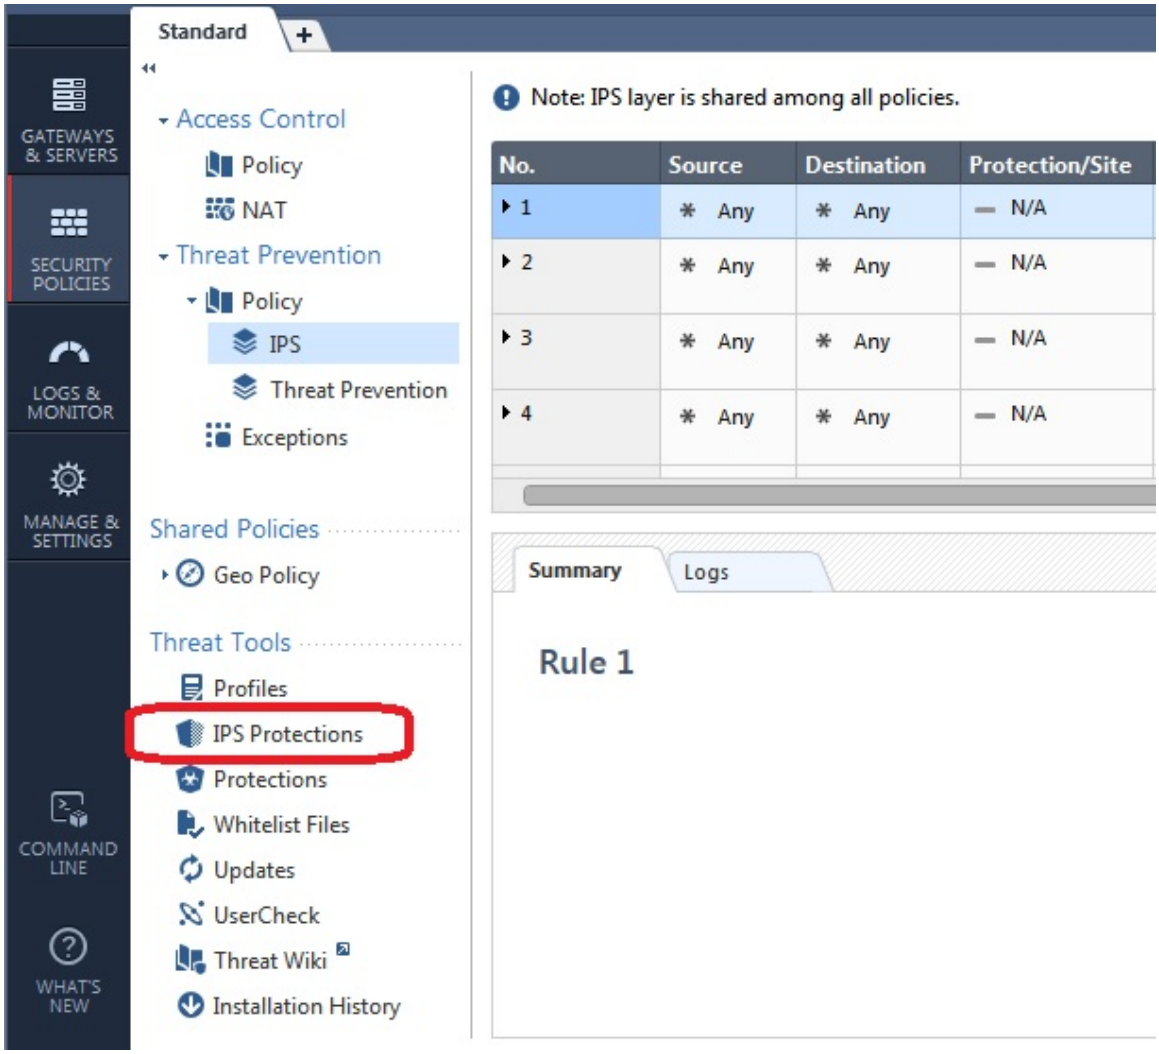

Page 186-192: APCL/URLF optimizations are now performed in the main Access Control Policy on the Security Policies tab, in a separate ACPL/URLF policy layer which is required for pre-R80 gateways. For R80.10 gateways, APCL/URLF enforcement can now take place in a single policy layer that is unified instead of in a separate policy layer. All the APCL/URLF tuning principles described on these pages still apply.

Page 195-196: In the SmartConsole, IPS Exceptions are added in an Exceptions policy located under "Security Policies...Threat Prevention...Exceptions".

Page 207: To view all IPS events unfiltered in the SmartConsole, from the "Logs & Monitor" tab click "Queries" then "Threat Prevention...By Blade...IPS Blade...All".

Page 212: The Application Control & URL Filtering Engine Settings described in the book can be accessed from the "Advanced Settings" button on the "Manage & Settings...Blades" screen in the SmartConsole:

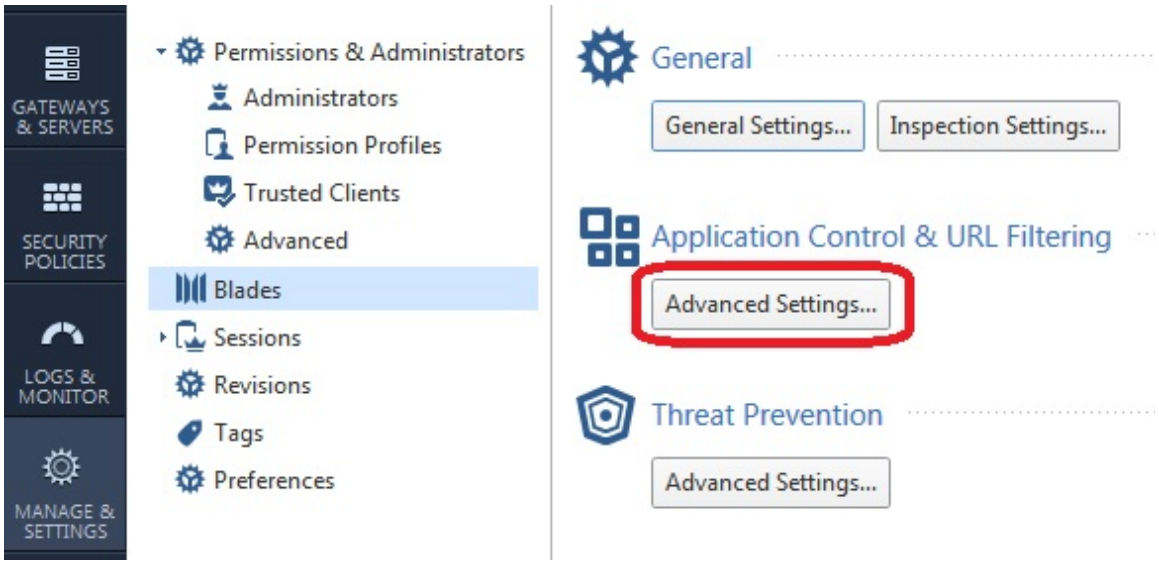

Page 214: Because the logging options for APCL/URLF are now integrated into the main Access Control Policy in the SmartConsole via the APCL/URLF policy layer, the options available in the Track field of the rulebase have changed from those presented in the book:

- **Network Log**: (Replaced in R80.10 with "Log") Generate a log with only the basic network information such as IP addresses and ports (application/category information will NOT be included). On a pre-R80 SMS, this setting is equivalent to setting the Track column to "Log" in the main rulebase (Firewall tab..Policy), but setting the Track column in the APCL/URLF policy rule to "None".

- **Log:** Includes network-level, application/category, and Content Awareness logging. This setting is equivalent to setting the Track column to "Log" in the main rulebase, and the Track column to "Log" in the APCL/URLF policy on a pre-R80 SMS.
- **Full Log**: (in R80.10 this is called "Detailed Log") For pre-R80 gateways, this is equivalent to the Log option described above. For R80.10 and later gateways, this option provides additional logging for application/category, even if an explicit application/category was not specified in the policy rule.
- **Extended Log**: (R80.10 only) Provides all individual URLs visited for a matching rule, and is the equivalent of setting "Complete Log" on a pre-R80 APCL/URLF policy rule. This logging option is likely to impact firewall performance and should be used sparingly.
- **Accounting Checkbox**: Equivalent to the "Account" track option described in the "Supplementary Material by Page Number" section of this document. While the Accounting option can be used to infer connection and session behavior as described in the book, new to R80.10 management is a feature called "Session Logging". When enabled this feature will correlate multiple individual connections into a "session" that has additional logging information. See the following for more information: Infinity R80.10 "Cool Feature of the Day" - *Session logging*
- **Suppression Checkbox:** This option will consolidate numerous identical logs matching the rule over a period of 3 hours into a single log entry.

Page 219-220: IP Fragmentation settings can be accessed from the "Inspection Settings" button located on the "Manage & Settings...Blades...General" screen in the SmartConsole:

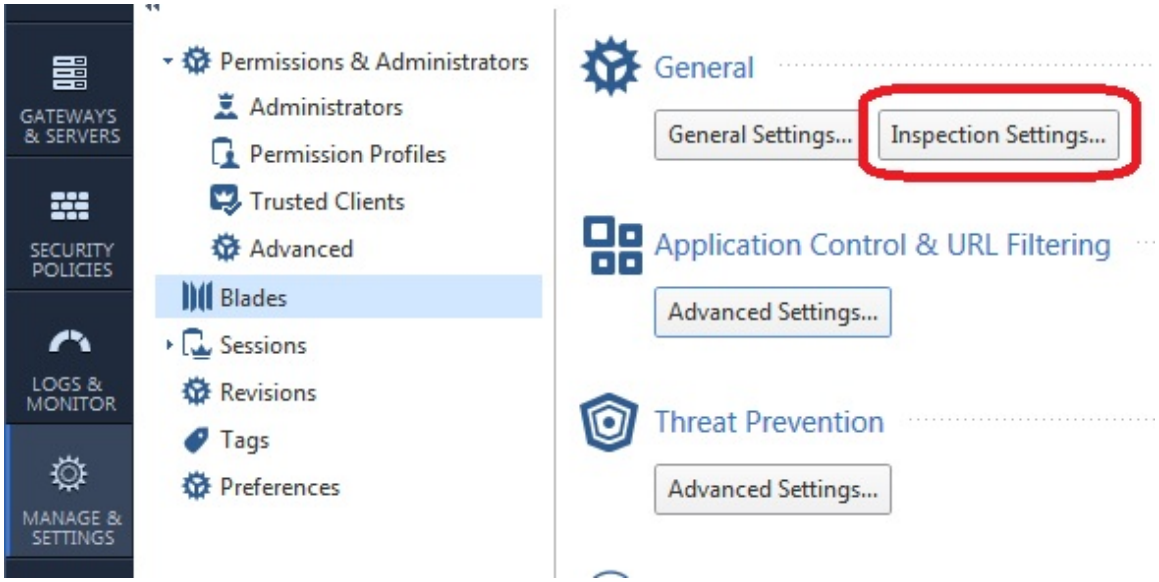

Page 234: In the SmartConsole, the New Connections Rate can be accessed by right-clicking the firewall object from the "Gateways & Servers" tab, selecting "Monitor", then selecting the "Network Activity" hyperlink.

Page 255: The Aggressive Aging settings are accessed from the "Inspection Settings" button located on the "Manage & Settings...Blades...General" screen in the R80 SmartConsole.