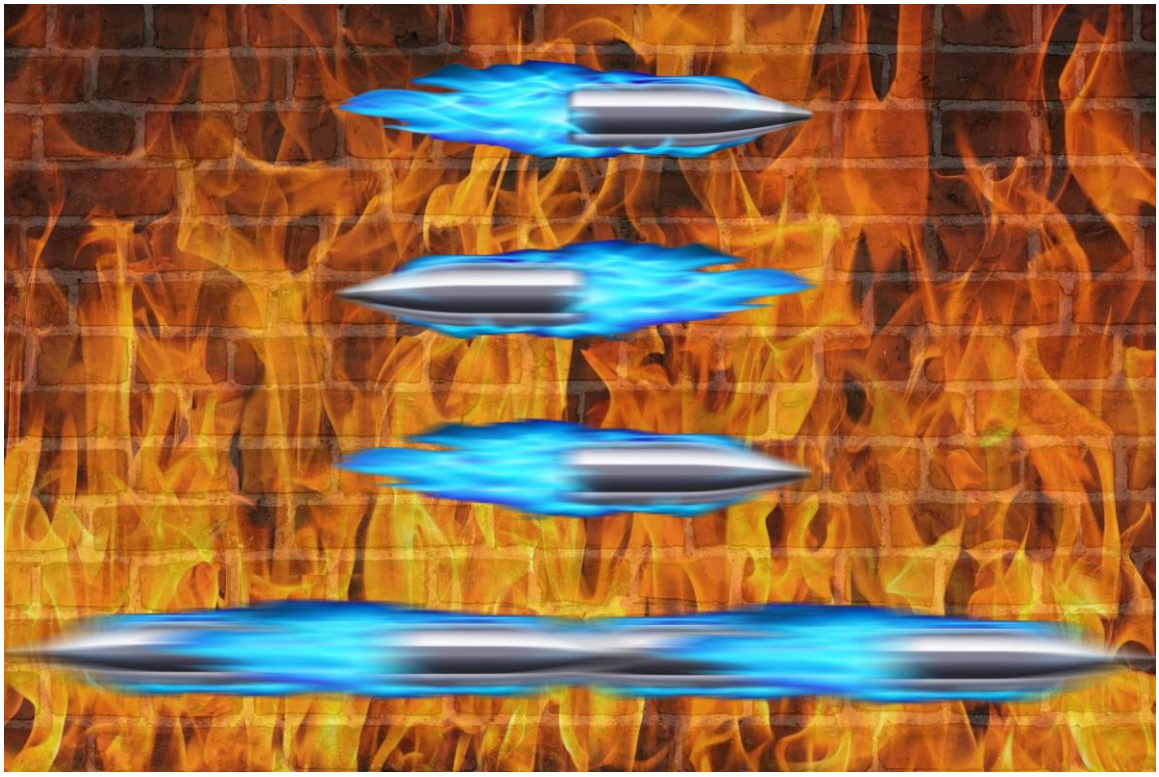

Max Power: Check Point Firewall Performance Optimization Addendum – 4/11/2016

Additional Tips, Tricks & R80 Management Supplement

Timothy C. Hall

Introduction

It has been approximately one year since the release of my book *Max Power: Check Point Firewall Performance Optimization*. This second 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 original addendum released 7/29/2015, new tips and tricks I've run into during my consulting work since then, as well as those detailed here:

https://www.cpug.org/forums/showthread.php/20631-Free-Max-Power-Tips-Tricks-R77-30-Addendum-Now-Available

In the first section of this document, **Supplementary Material by Page Number**, all added or updated content from the previous addendum dated 7/29/2015 will be highlighted like this to clearly show what has changed.

The new R80 release is management-only, and since the vast majority of performance tuning occurs directly on the firewall, all of the techniques detailed in the original book still apply under R80 Management, although some of them may look a little different in the new SmartConsole. In the second section, **R80 Management Updates**, these differences will be documented so that the optimization steps described in Max Power can be performed from the new R80 SmartConsole or equivalent GUI tool.

Also I'm pleased to announce that a new book covering Check Point R80 has been in production for 5 months now and is nearing completion. This new book is the first of a series and has a working title of "New Frontier: Check Point R80". It is a collaborative effort featuring myself and three other professionals with considerable Check Point expertise:

- Valeri Loukine – Switzerland
- Kishin Fatnani – India
- Eric Anderson – USA

More information will be available soon, stay tuned!

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.

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:**

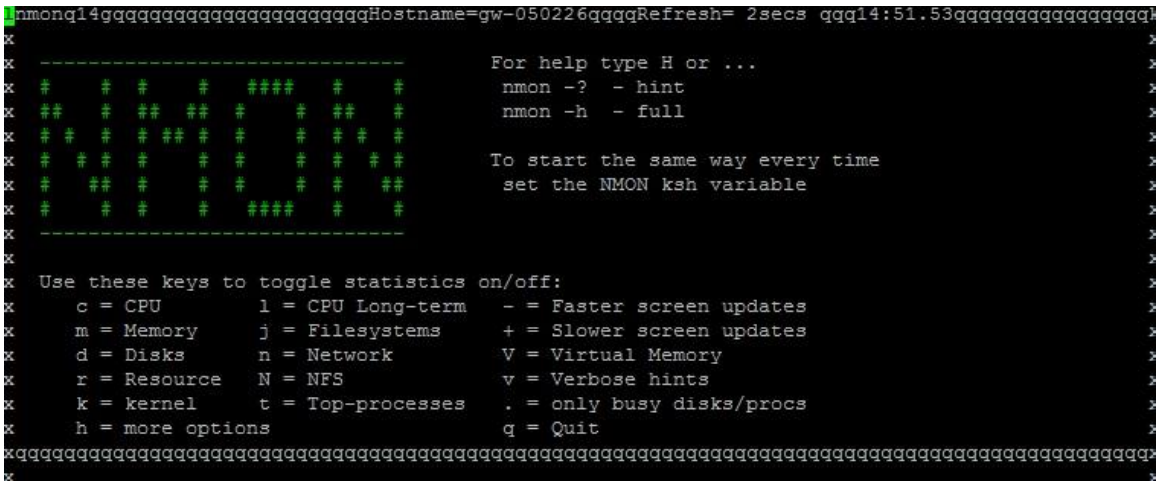

Added to Gaia in R76 it serves many of the same functions as the **top** command for monitoring CPU usage on the firewall but in a more graphical format. As an example, typing lowercase "L" provides a CPU usage graph; hitting "c" will show a graph for a for multi-core systems. While most of these CPU 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.

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":

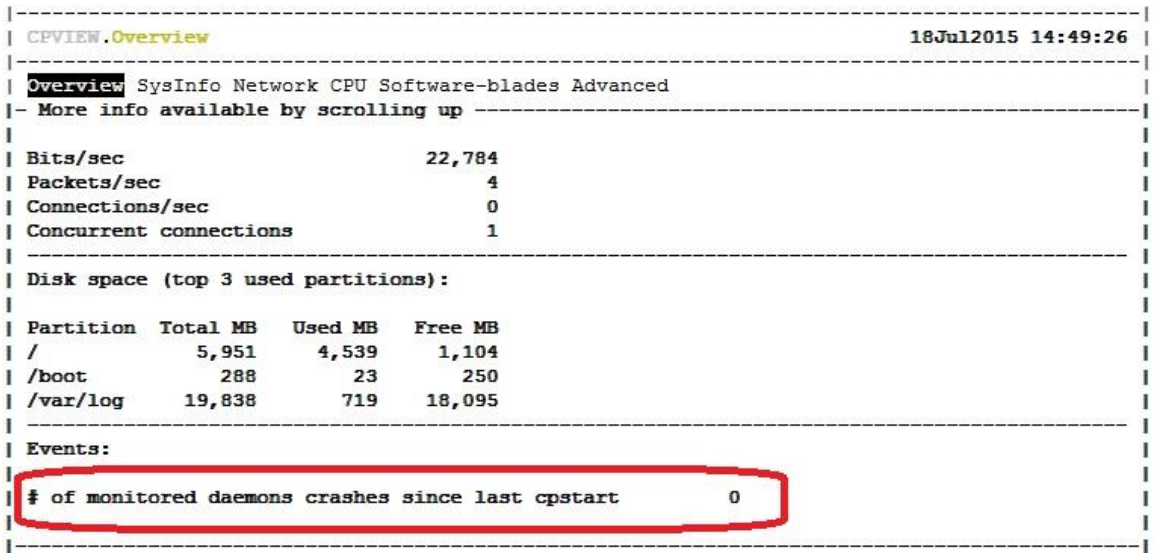

If this value is nonzero run **cpwd_admin list** to determine which daemon(s) are having a problem.

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 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 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.) the existing log entry is converted from a **Log** type to an **Account** type. Additional statistics are then provided for the connection including the connection duration and number of payload/data bytes sent and received by the connection. These statistics can be used to infer the connection's behavior and assist in troubleshooting.

Note that in the R80 SmartConsole the Account option is now enabled by selecting the "Accounting" checkbox in a rule's Track column. The additional connection information statistics may take up to one minute to appear in the R80 SmartConsole "Logs & Monitor" tab after the connection has terminated.

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:

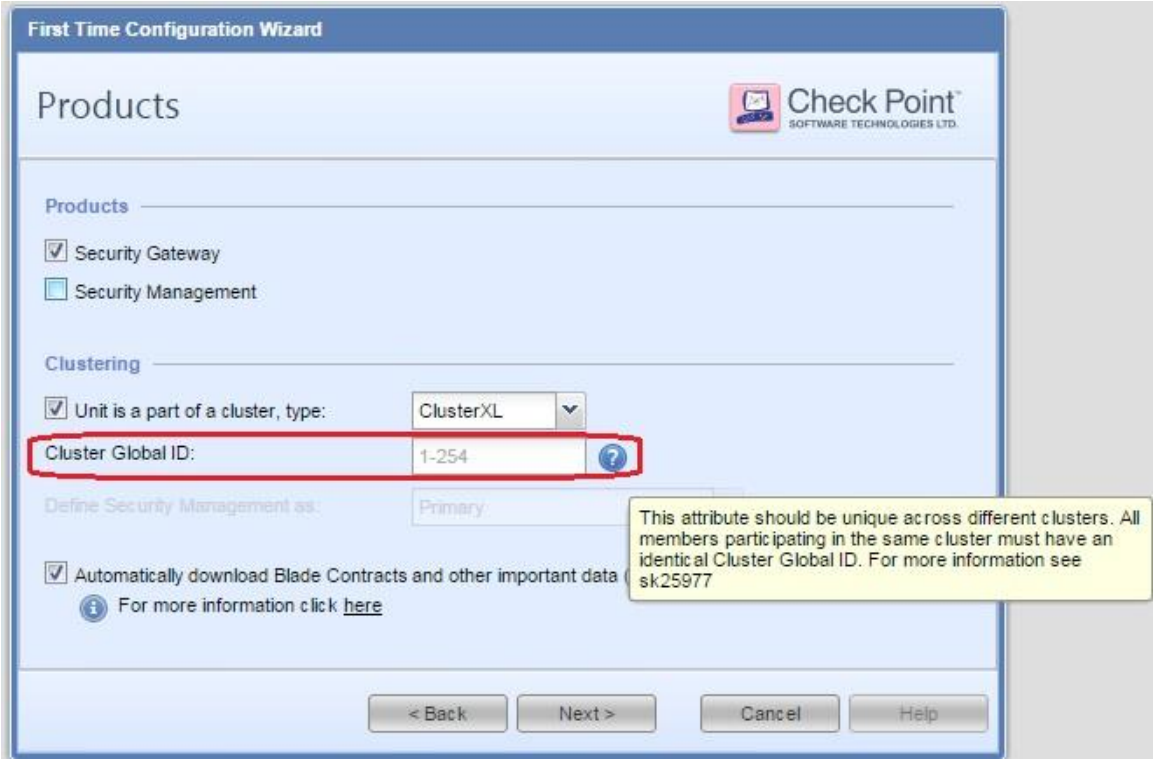

The Cluster Global ID should be set identically on all members of the same cluster, but be a different value for different clusters.

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.

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 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.

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 Check Point firewall appliances: sk98950: Slow traffic speed (high latency) when transferring files over VPN tunnel with 3DES encryption

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; as a general rule of thumb you should consider enabling this feature when the following conditions are present (but note the new warnings 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 limitation has still not been lifted in R77.30). If there is substantial IPSec and/or VoIP traffic traversing the firewall, exclude the lead Firewall Worker Core from consideration when applying the 10% rule of thumb above.

If you are planning to enable the Dynamic Dispatcher, be aware of the following four recently-discovered issues. In particular the first one can cause slow, creeping performance degradation on the firewall over time:

- 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:

```
CPVIEW.Advanced.Network.Path                                    18Jul2015 14:46:50
Overview SysInfo Network CPU Software-blades Advanced
CPU-Profiler Memory Network SecureXL CoreXL PrioQ Streaming RAD
Path Direction Size

Path distribution summary (available when SecureXL is on):

Totals          SXL Mbps    SXL pps    PXL Mbps    PXL pps    FW Mbps    FW pps
TCP                    0          0           0          0          0         0
UDP                    0          0           0          0          0         0
Other                  0          0           0          0          0         0

Protocol        SXL Mbps    SXL pps    PXL Mbps    PXL pps    FW Mbps    FW pps
-                      -          -           -          -          -         -
```

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 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.

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.

Pages 220-221: The HTTPS Inspection feature was significantly enhanced in R77.30. 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 be forced into the Firewall Path) 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 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 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.

Page 282: If performing lab benchmarking of Check Point firewalls, be sure to enable the following feature: sk105261: CoreXL Dynamic Dispatcher in R77.30 (but be sure to heed the warnings stated in the Pages 149-151 section above). 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.

R80 Management Updates

Page 16: The R80 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 R80 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".

Page 83: To provide the functionality for firewall thresholds as described in the book, the R80 SmartView Monitor will need to be invoked directly from the R80 SmartConsole. From the "Logs & Monitor" tab create a new tab with the "+" button, and then select "External Apps...Tunnel & User Monitoring" as shown here:

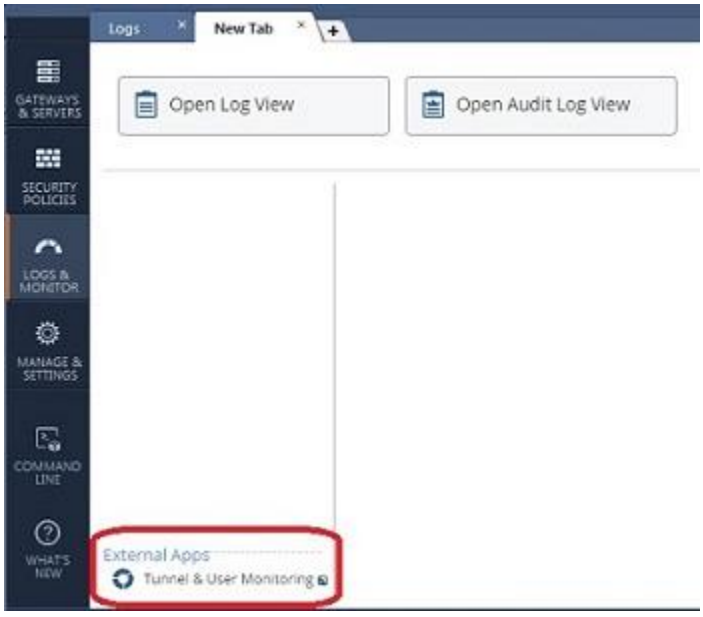

Pages 97-98: To filter for Control events on the "Logs and Monitor" tab of the R80 SmartConsole as described in the book, the proper search field syntax is "type:Control".

Page 180: In the R80 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. It is anticipated that R80.10 firewalls will 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.

Page 181: IPS Signatures are accessed in the R80 SmartConsole by clicking "IPS Protections" under Threat Tools on the "Security Policies...Threat Prevention...IPS" screen:

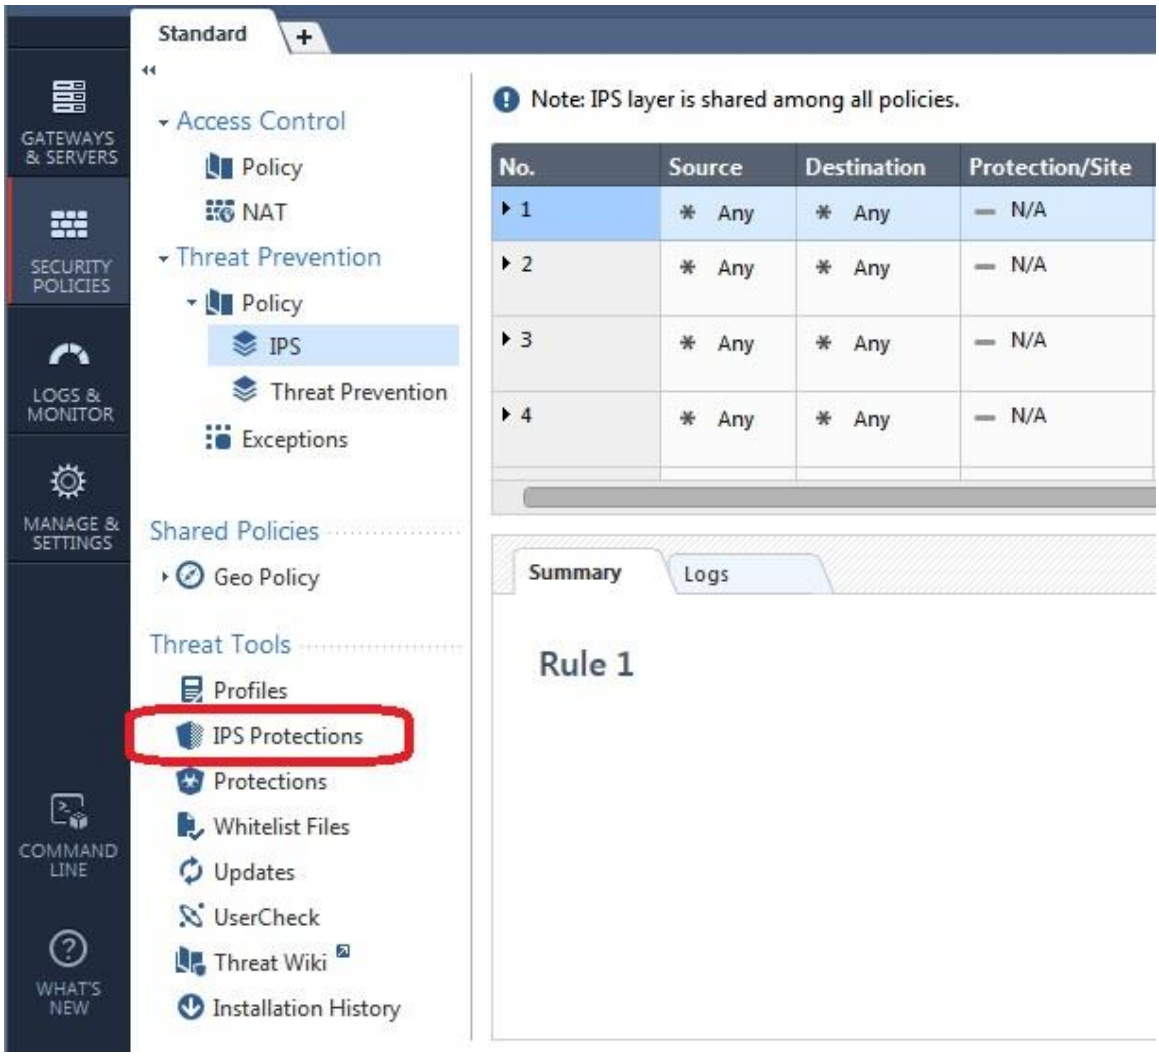

Page 186-192: APCL/URLF optimizations are now performed in the main Access Control Policy on the Security Policies tab, in a separate ACPL/URLF policy layer which is required for pre-R80 gateways.

Page 195-196: In the R80 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 R80 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 R80 SmartConsole:

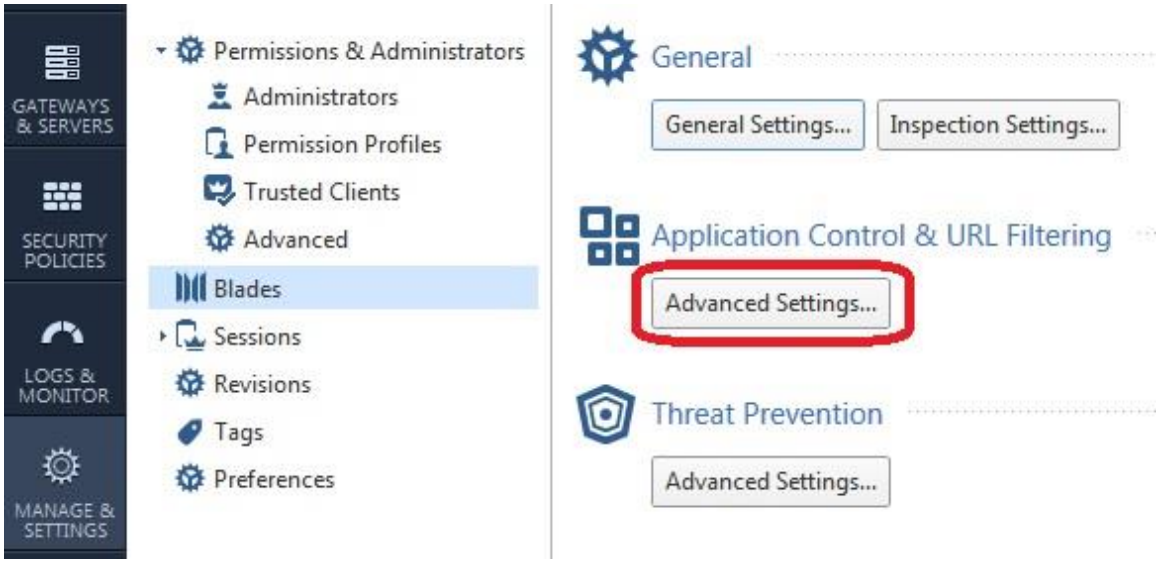

Page 214: Because the logging options for APCL/URLF are now integrated into the main Access Control Policy in the R80 SmartConsole via the APCL/URLF policy layer, the options available in the Track field of the rulebase have changed from those presented in the book:

- **Network Log**: 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 both network-level and application/category 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**: For pre-R80 gateways, this is equivalent to the Log option described above.
- **Accounting Checkbox**: Equivalent to the "Account" track option described in the "Supplementary Material by Page Number" section of this document (Page 90).
- **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 R80 SmartConsole:

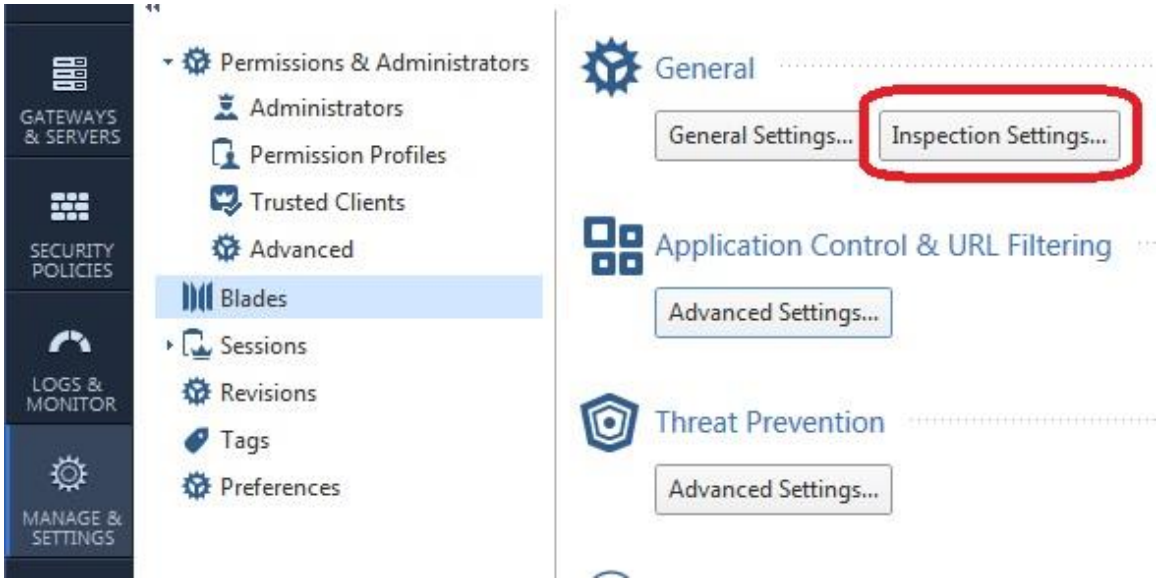

Page 234: In the R80 SmartConsole, the New Connections Rate can be accessed by right-clicking the firewall object from the "Gateways & Servers" tab, selecting "Monitor", then selecting the "Network Activity" hyperlink.

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.