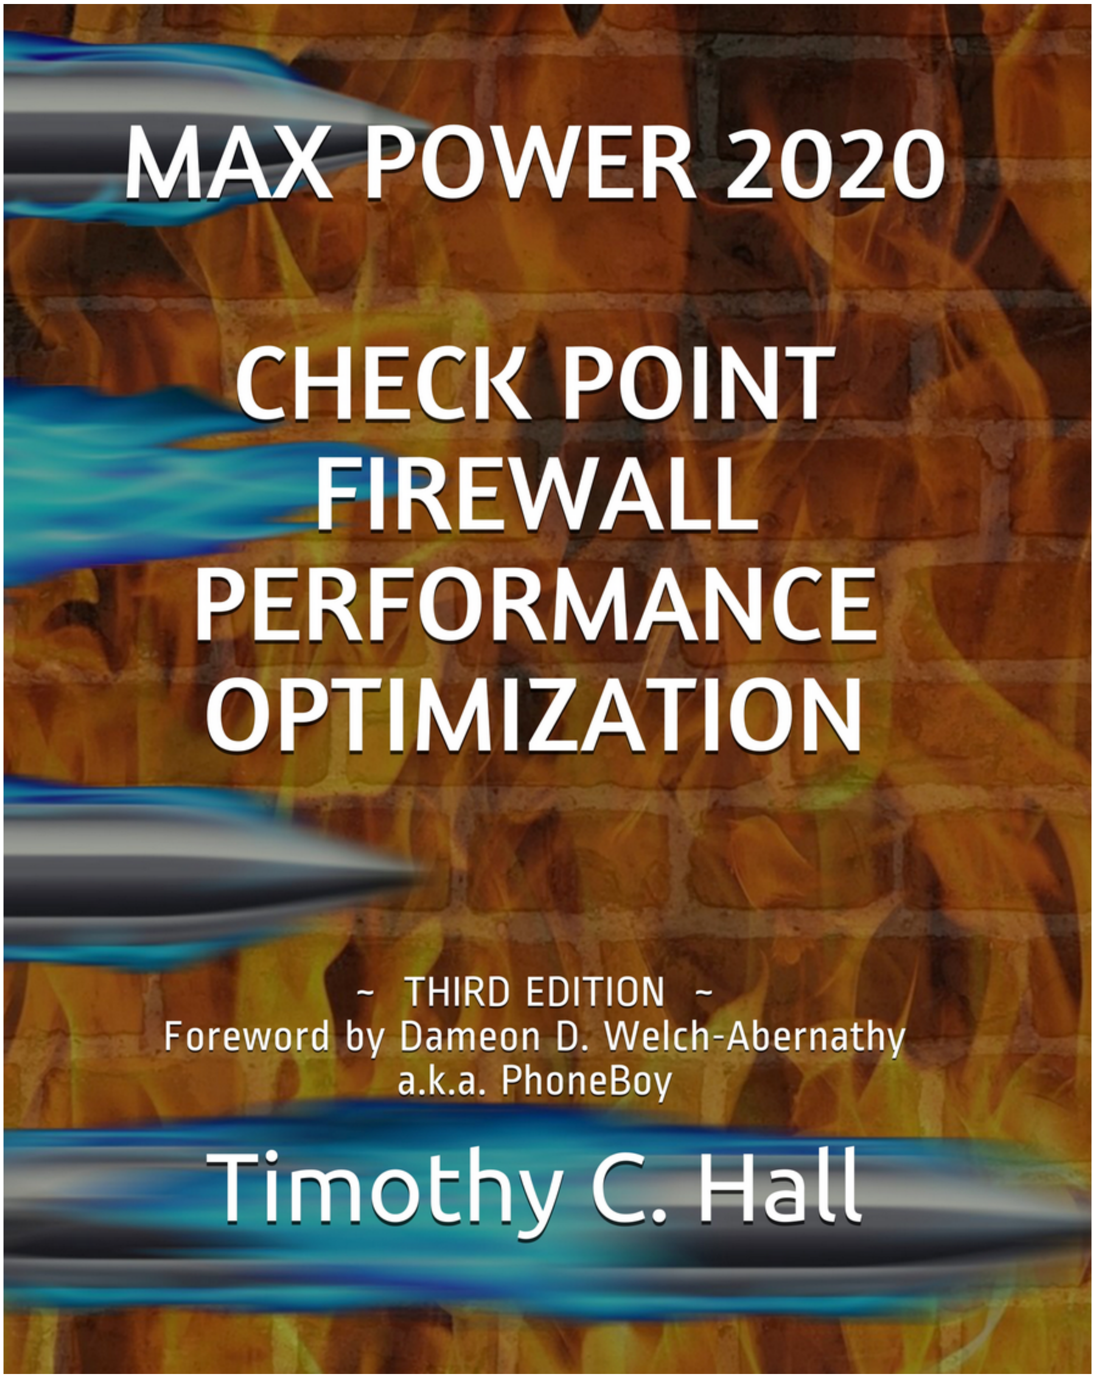
MAX POWER 2020
CHECK POINT FIREWALL PERFORMANCE OPTIMIZATION
~ THIRD EDITION ~
Foreword by Dameon D. Welch-Abernathy a.k.a. PhoneBoy
Timothy C. Hall

Max Power 2020: Check Point Firewall Performance Optimization

R80.40 Addendum 8/24/2020

Timothy C. Hall

Introduction

It has been just over seven months since the release of my book *Max Power 2020: Check Point Firewall Performance Optimization* in January 2020. The book's publication feels like it was many years ago given all that has transpired since then in our world. This 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, as well as updates for version R80.40.

Before we start off into the page number references, there are three general observations I'd like to share:

1. sk167553: Performance Investigation Procedure - How To is an outstanding, all-new SK that takes a deep dive into gateway performance optimization at both a breadth and depth not seen before in any official Check Point documentation. My gushing praise has nothing whatsoever to do with the fact that I provided some input and feedback for this SK prior to its release. Definitely worth your time to check out!
2. I'm very pleased to report that the "functionality gap" for dealing with heavy connections (a.k.a. elephant flows) on USFW-enabled firewalls has been closed in the latest R80.40 Jumbo HFAs. The various USFW-based limitations were discussed in Chapter 12 of the book and my 2020 CPX presentation Big Game Hunting: Elephant Flows (CheckMates login required). For all intents and purposes, commands such as **fw ctl multik utilize**, **fw ctl multik print_heavy_conn**, and **fw ctl multik gconn** (and even the dynamic mode 9 of Priority Queueing) now work exactly the same on both a kernel mode and user mode (USFW) firewall. Note that this functionality gap was resolved in the various R80.40 Jumbo HFAs so the initial, unpatched R80.40 release code will not have them available for USFW-enabled firewalls. These capabilities have not been back-ported via Jumbo HFAs into USFW-based R80.30 w/ Gaia 3.10

gateways other than the **fw ctl multik gconn** command. See here for the latest updates: sk164215: How to Detect and Handle Heavy Connections.

3. R80.40 Jumbo HFA Take 25+ introduced the ability for the CoreXL "split" between SND/IRQ cores and Firewall Worker cores to be automatically adjusted based on dynamic traffic loads on the various cores. This new "Dynamic Split" feature is disabled by default in R80.40, but is expected to be enabled by default in the next major release. This feature could be useful in scenarios where manual split adjustments have to made frequently on the firewall (which will still require a reboot), due to radical shifts in the amount and/or acceleration status of traffic loads through the gateway. The Dynamic Split feature also includes an anti-flap mechanism to ensure that unnecessary split adjustments are not constantly being made. There are several new **cpview** screens that can be used to monitor the status of this feature, and all split adjustments & relevant metrics are logged to the **$FWDIR/log/dsd.elg** file. Dynamic Split does not support Check Point Appliances with less than 8 CPU cores or open hardware (including virtual machines). VSX is also not supported with Dynamic Split. More reading here:
    - sk164155: Dynamic Balancing for CoreXL
    - https://community.checkpoint.com/t5/General-Topics/R80-40-Dynamic-split-of-CoreXL/m-p/82321#M16649

Supplementary Material & Corrections

p. xxv: *Correction:* The referenced hyperlink for "Check Point for Beginners" in the PDF edition of the book is incorrect, the proper link is: https://community.checkpoint.com/t5/Check-Point-for-Beginners-2-0/bg-p/check-point-for-beginners-2-0

p. 33: The location of the simkern.conf file referenced on this page has changed in R80.20+, the correct location is now $PPKDIR/conf/simkern.conf.

p. 33: The vpnkern.conf file referenced on this page has been deprecated as of R80.20 and should not be used: sk166179: In R80.20, kernel changes made using vpnkern.conf are not set after reboot.

p. 34: The following new SK nicely summarizes the available code release take numbers for all currently-supported appliances: sk166536: Software Images for Check Point Security Gateway Appliances.

p. 37: Thanks to Danny Jung, the "Super Seven" performance assessment can now be accessed as a SmartConsole extension: https://community.checkpoint.com/t5/SmartConsole-Extensions/SmartConsole-Extension-Performance-Assessment/m-p/83121

p. 38: The equivalent "Super Seven" performance assessment for SMB appliances (model numbers 600-1500) is available at CheckMates here: https://community.checkpoint.com/t5/SMB-Appliances-and-SMP/Super-Seven-Performance-Assessment-Commands-SMB-Edition/td-p/73078

p. 41: If you don't have a lot of experience taking packet captures, check out the very handy "Packet Captures for Dummies" article by Check Point employee Kyle Gordon: https://community.checkpoint.com/t5/API-CLI-Discussion-and-Samples/packet-captures-sh-Packet-Captures-for-Dummies/m-p/71820

p. 42: Note that the **cppcap** tool is bundled by default with R80.40+, and need not be downloaded from the SK and manually installed as it did for R80.30 and earlier.

p. 48: The **ifconfig** command has been deprecated, instead of **ifconfig -a** use the command **ip -s a**

p. 58: The **ifconfig** command has been deprecated, instead of **ifconfig -a** use the command **ip -s a**

p. 60: The **ifconfig** command has been deprecated, instead of **ifconfig -a** use the command **ip -s a** at the top of the page.

p. 60: The modified command to detect non-1500 byte MTUs using the **ip** command instead of **ifconfig** is: **ip -s a | grep mtu | grep -vi loopback | grep -v "mtu 1500"**

p. 63: For performing network discovery with CDP and LLDP from the firewall, there is a great CDP/LLDP daemon available for Gaia: https://github.com/oribit/cdpd-cp. Note that the direct use of LLDP between a Check Point firewall NIC and network switch is still not officially supported as mentioned in the book; LLDP must be disabled on the switch to avoid problems.

p. 63: Check Point nicely summarizes all supported transceivers for the various appliance models in the following SK, which can help avoid unsupported configurations (and subsequent interface performance and stability problems): sk92755: Compatibility of transceivers for Check Point appliances.

p. 69: The **ifconfig** command has been deprecated, instead of **ifconfig -a** use the command **ip -s a**

p. 70: Be aware that a misconfigured bond interface can cause SecureXL to silently drop traffic, which is then resolved if SecureXL is disabled. Verify the configuration any new bond interfaces with the following SK to avoid this problem: sk167613: SecureXL dropping traffic due to bond mis-configuration.

p. 70: Note that up to a maximum of 8 physical interfaces can be configured in a single bond interface.

p. 75: The **ethtool** utility was updated and picked up a bunch of new command line switches in Gaia kernel 3.10. For fiber interfaces in particular, from expert mode try **ethtool -m (interfacename)**. To obtain the low-level NIC hardware statistics, also try the **ethtool --phy-statistics (interfacename)** command. On Gaia 3.10 you should see something like this from **ethtool -m**, which is very useful for diagnosing low-level fiber issues that are causing RX-ERRs or spurious carrier transitions:

```
Identifier                    : 0x03 (SFP)
Extended identifier               : 0x04 (GBIC/SFP defined by 2-wire interface ID)
Connector                      : 0x07 (LC)
Transceiver codes                 : 0x00 0x00 0x00 0x01 0x20 0x40 0x0c 0x05
Transceiver type                 : Ethernet: 1000BASE-SX
Transceiver type                 : FC: intermediate distance (I)
Transceiver type                 : FC: Shortwave laser w/o OFC (SN)
Transceiver type                 : FC: Multimode, 62.5um (M6)
Transceiver type                 : FC: Multimode, 50um (M5)
Transceiver type                    : FC: 200 MBytes/sec
Transceiver type                    : FC: 100 MBytes/sec
Encoding                          : 0x01 (8B/10B)
BR, Nominal                         : 2100MBd
Rate identifier                    : 0x00 (unspecified)
Length (SMF,km)                       : 0km
Length (SMF)                        : 0m
Length (50um)                       : 300m
Length (62.5um)                      : 150m
Length (Copper)                      : 0m
Length (OM3)                        : 0m
Laser wavelength                     : 850nm
Vendor name                         : JDSU
Vendor OUI                          : 00:01:9c
Vendor PN                          : PLRXPL-VI-S24-22
Vendor rev                         : 1
Optical diagnostics support              : Yes
Laser bias current                    : 21.348 mA
Laser output power                      : 0.3186 mW / -4.97 dBm
Receiver signal average optical power     : 0.3195 mW / -4.96 dBm
Module temperature                        : 41.70 degrees C / 107.05 degrees F
Module voltage                       : 3.2947 V
Alarm/warning flags implemented           : Yes
Laser bias current high alarm            : Off
Laser bias current low alarm             : Off
Laser bias current high warning           : Off
Laser bias current low warning            : Off
Laser output power high alarm             : Off
Laser output power low alarm              : Off
Laser output power high warning            : Off
Laser output power low warning             : Off
Module temperature high alarm              : Off
Module temperature low alarm               : Off
Module temperature high warning             : Off
Module temperature low warning          : Off
Module voltage high alarm               : Off
Module voltage low alarm                : Off
Module voltage high warning             : Off
Module voltage low warning              : Off
Laser rx power high alarm               : Off
Laser rx power low alarm                : Off
Laser rx power high warning             : Off
Laser rx power low warning              : Off
Laser bias current high alarm threshold   : 10.000 mA
Laser bias current low alarm threshold    : 1.000 mA
Laser bias current high warning threshold : 9.000 mA
 Laser bias current low warning threshold  : 2.000 mA
Laser output power high alarm threshold   : 0.8000 mW / -0.97 dBm
Laser output power low alarm threshold    : 0.1000 mW / -10.00 dBm
Laser output power high warning threshold : 0.6000 mW / -2.22 dBm
Laser output power low warning threshold  : 0.2000 mW / -6.99 dBm
Module temperature high alarm threshold   : 90.00 degrees C / 194.00 degrees F
Module temperature low alarm threshold    : -40.00 degrees C / -40.00 degrees F
Module temperature high warning threshold : 85.00 degrees C / 185.00 degrees F
Module temperature low warning threshold  : -40.00 degrees C / -40.00 degrees F
Module voltage high alarm threshold       : 4.0000 V
Module voltage low alarm threshold        : 0.0000 V
Module voltage high warning threshold     : 3.6450 V
Module voltage low warning threshold      : 2.9550 V
Laser rx power high alarm threshold       : 1.6000 mW / 2.04 dBm
Laser rx power low alarm threshold        : 0.0100 mW / -20.00 dBm
Laser rx power high warning threshold     : 1.0000 mW / 0.00 dBm
Laser rx power low warning threshold      : 0.0200 mW / -16.99 dBm
```

p. 81: The following SK nicely summarizes the expected sensor threshold values for all Check Point appliances; very useful for detecting if reported values are on the borderline of normalcy: sk119232: Hardware sensors thresholds on Check Point appliances.

p. 109: The maximum ARP table size was increased from 16384 to 131072 on Gaia 3.10 kernel systems via Jumbo HFA; if you still see the old limit and wish to increase it beyond that value, load the latest GA Jumbo HFA. Latest updates here: sk43772: 'kernel: neighbour table overflow' appears repeatedly in /var/log/messages files.

p. 112: As mentioned in the book, where CPU time is actually expended (user/process space vs. kernel/system space) on a firewall that has USFW enabled looks much different than a firewall in the traditional kernel-based mode. Here is a great Checkmates discussion detailing all the USFW-based processes and their functions: https://community.checkpoint.com/t5/General-Topics/High-CPU-utilization-during-process-fwk0-dev-0-UMFW-vs-KMFW/m-p/70648#M14307

p. 118: If you suspect a process that is monitored by cpwd is constantly crashing, you can confirm this by looking at the #START value for the monitored process shown by **cpwd_admin list**; this value is incremented every time a process is restarted by cpwd.

p. 121: Another trick for helping identify an unknown process and its function is to examine the individual threads of execution for that process with the following command: **top -Hbn1 -p (process ID#)**

p. 137: For the **cpview** tool, the **Network...Top Connections** screen went away in R80.20, and **CPU...Top Connections** was also not present if USFW was enabled. Both screens have now returned regardless of USFW state in R80.40 Take 69+: sk167903: CPview Top Connections tab shows no data.

p. 146: A new mode for ClusterXL has been added in R80.40 called “Active-Active”:

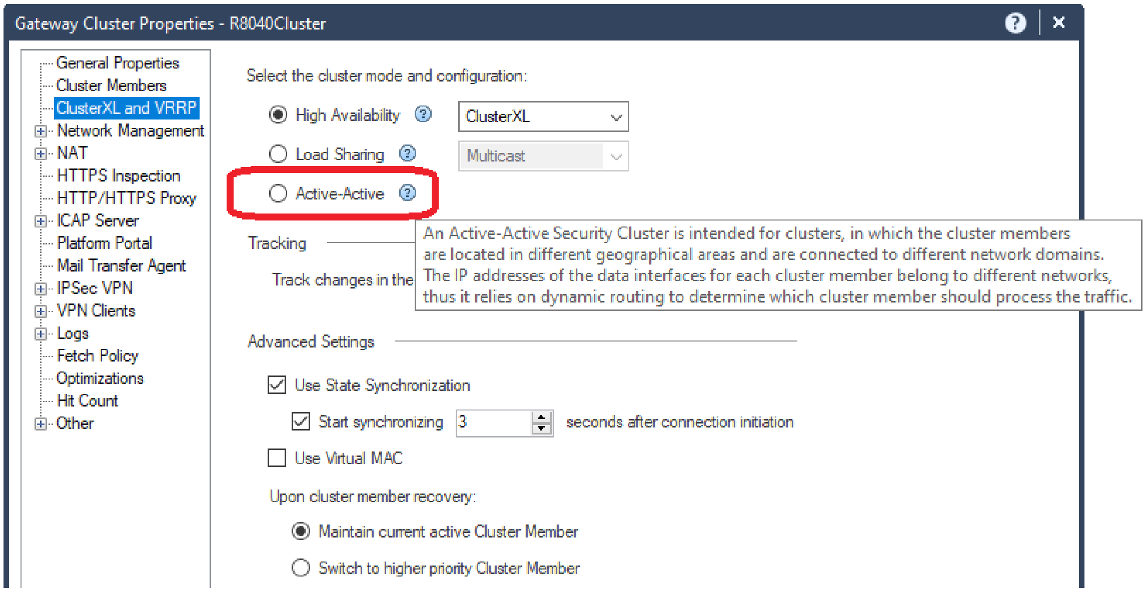

Gateway Cluster Properties - R8040Cluster

- General Properties
- Cluster Members
- ClusterXL and VRRP
- Network Management
- NAT
- HTTPS Inspection
- HTTP/HTTPS Proxy
- ICAP Server
- Platform Portal
- Mail Transfer Agent
- IPSec VPN
- VPN Clients
- Logs
- Fetch Policy
- Optimizations
- Hit Count
- Other

Select the cluster mode and configuration:

● High Availability — ClusterXL

○ Load Sharing — Multicast

○ Active-Active

Tracking

Track changes in the

An Active-Active Security Cluster is intended for clusters, in which the cluster members are located in different geographical areas and are connected to different network domains. The IP addresses of the data interfaces for each cluster member belong to different networks, thus it relies on dynamic routing to determine which cluster member should process the traffic.

Advanced Settings

☑ Use State Synchronization

☑ Start synchronizing 3 seconds after connection initiation

□ Use Virtual MAC

Upon cluster member recovery:

● Maintain current active Cluster Member

○ Switch to higher priority Cluster Member

As stated in the ToolTip, this allows an external entity such as a Dynamic Routing Protocol (BGP, OSPF, etc.) or a load balancer appliance to decide how to balance traffic through the Cluster Members. ClusterXL in Load Sharing Multicast mode or Load Sharing Unicast mode uses its own internal mechanisms to balance cluster member load. This Active-Active scenario was previously one of the few remaining use cases for VRRP over ClusterXL, and opens up some new possibilities for geographically disparate cluster members.

p. 148: *Correction:* The ClusterXL Sticky Decision Function (SDF) no longer exists in R80.20+, but connection “stickyness” is still guaranteed by other methods that are compatible with SecureXL and do not impact the acceleration status of traffic. https://community.checkpoint.com/t5/General-Topics/What-does-Sticky-Decision-Function-DO/m-p/92923#M18447

p. 158: The **ifconfig** command has been deprecated, instead of **ifconfig -a** use the command **ip -s a**

p. 175: *Correction:* The **/sbin/cpuinfo** command that was used under Gaia kernel 2.6.18 to determine if SMT/Hyperthreading is enabled no longer works in Gaia kernel 3.10, use this command instead: **cat /sys/devices/system/cpu/smt/active**

1=active
0=inactive

p. 175: Some open hardware firewalls may have more than 2 execution threads per core present when SMT/Hyperthreading is enabled; run the following command to see the number of execution threads per physical core: **lscpu | grep Thread**

p. 177: If CoreXL is off and there is no menu option in **cpconfig** to modify its configuration, it is possible CoreXL was never installed in the first place. To rectify this, manually edit the file **/etc/fw.boot/boot.conf** and change "COREXL_INSTALLED 0" to "COREXL_INSTALLED 1", then reboot the firewall.

p. 190: To see an extremely high level of detail for Multi-Queue's operation (including network counters for each of the individual queues and IRQs), run the following command: **mq_mng –o –vv**

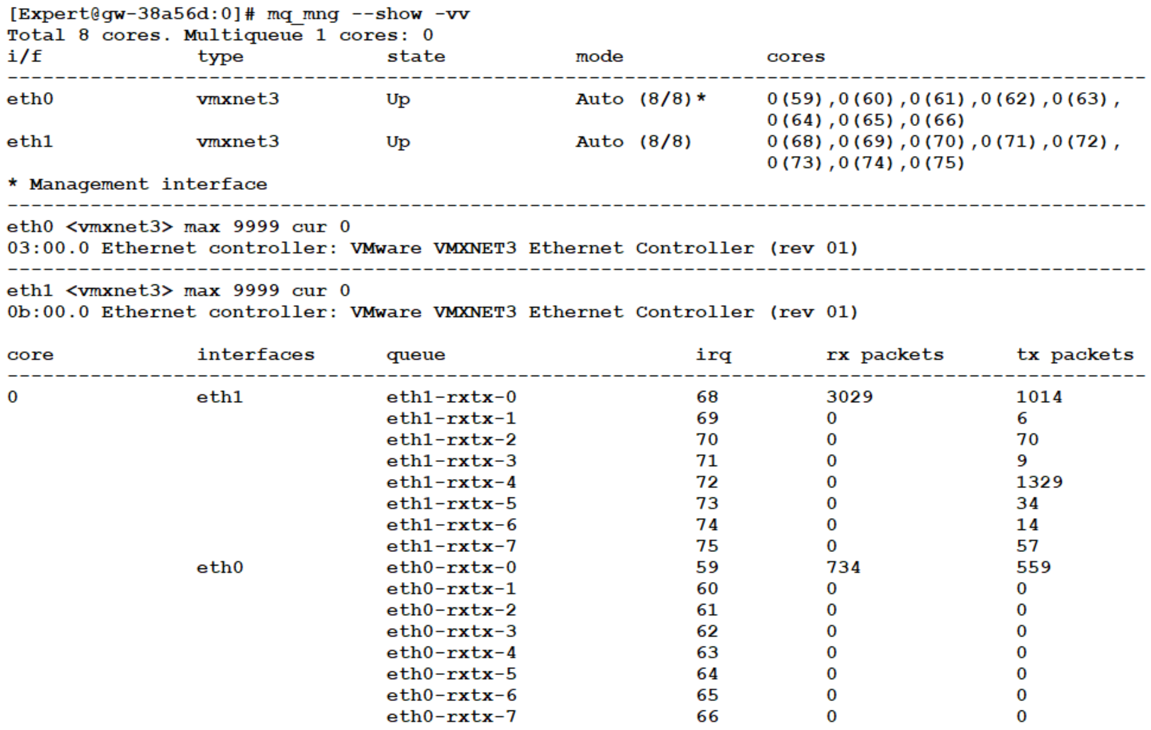

```
[Expert@gw-38a56d:0]# mq_mng --show -vv
Total 8 cores. Multiqueue 1 cores: 0
i/f              type             state            mode             cores
------------------------------------------------------------------------------------------------
eth0             vmxnet3          Up               Auto (8/8)*      0(59),0(60),0(61),0(62),0(63),
                                                                    0(64),0(65),0(66)
eth1             vmxnet3          Up               Auto (8/8)       0(68),0(69),0(70),0(71),0(72),
                                                                    0(73),0(74),0(75)
* Management interface
------------------------------------------------------------------------------------------------
eth0 <vmxnet3> max 9999 cur 0
03:00.0 Ethernet controller: VMware VMXNET3 Ethernet Controller (rev 01)
------------------------------------------------------------------------------------------------
eth1 <vmxnet3> max 9999 cur 0
0b:00.0 Ethernet controller: VMware VMXNET3 Ethernet Controller (rev 01)

core             interfaces       queue                      irq         rx packets       tx packets
------------------------------------------------------------------------------------------------
0                eth1             eth1-rxtx-0                68          3029             1014
                                  eth1-rxtx-1                69          0                6
                                  eth1-rxtx-2                70          0                70
                                  eth1-rxtx-3                71          0                9
                                  eth1-rxtx-4                72          0                1329
                                  eth1-rxtx-5                73          0                34
                                  eth1-rxtx-6                74          0                14
                                  eth1-rxtx-7                75          0                57
                 eth0             eth0-rxtx-0                59          734              559
                                  eth0-rxtx-1                60          0                0
                                  eth0-rxtx-2                61          0                0
                                  eth0-rxtx-3                62          0                0
                                  eth0-rxtx-4                63          0                0
                                  eth0-rxtx-5                64          0                0
                                  eth0-rxtx-6                65          0                0
                                  eth0-rxtx-7                66          0                0
```

p. 191: As stated in the book, Multi-Queue is enabled by default on all network interfaces except the management interface for Gaia 3.10 firewalls. Trying to manually change the state or operation of Multi-Queue on individual network interfaces is strongly discouraged under Gaia 3.10, as doing so is likely to disrupt the proper balancing of traffic across the multiple traffic queues thus causing voluminous RX-DRPs. To determine if this issue is present and how to rectify it, see this SK: sk168498: High rate of input discards after reboot when Multi-Queue is configured.

p. 192: Similar to the older 3200/5000/15000/23000 appliances mentioned in the book, the new 6500 model also uses the Intel I211 controller hardware for its on-board interfaces, which do not support more than 2 queues per interface: sk165516: Multi-Queue and Dynamic Split cannot add more than 2 CoreXL SNDs on 6500 appliance

p. 192: The vmxnet3 network driver also supports Multi-Queue operation in a virtual environment with a maximum of eight queues per interface. Note that the e1000 driver (which is normally the default network driver provided in a virtualized environment) does NOT support Multi-Queue and cannot be managed with the **mq_mng** command.

p. 205: An easier way to see the CoreXL"split" of SND/IRQ vs. Firewall Worker cores can be found on the CPU...Overview screen of **cpview** by scrolling down:

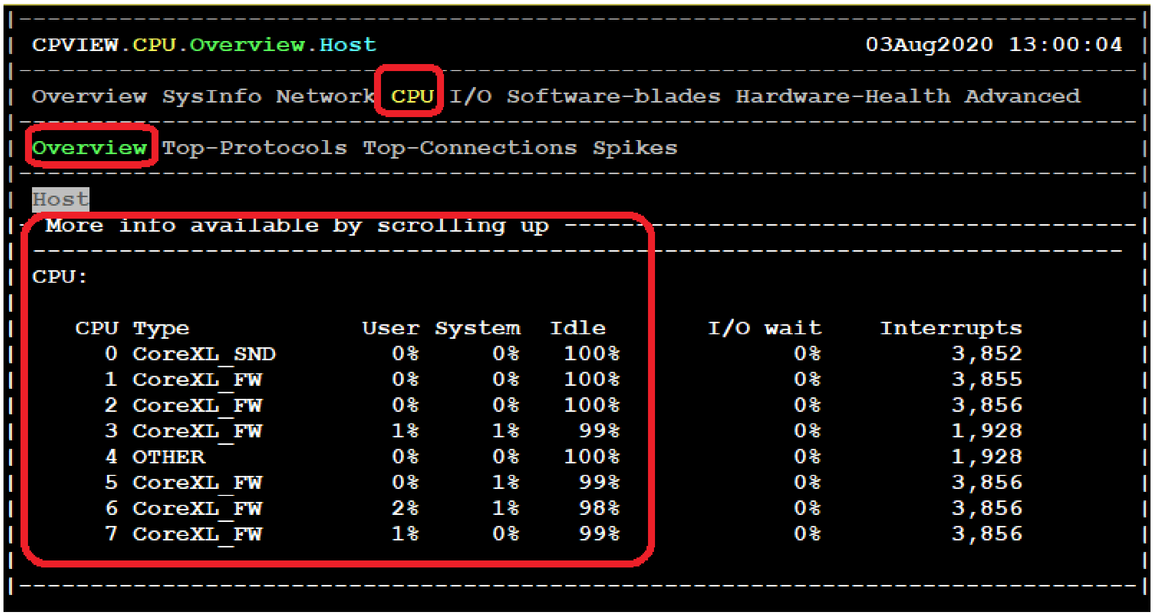

```
|-------------------------------------------------------------------------------|
| CPVIEW.CPU.Overview.Host                                    03Aug2020 13:00:04 |
|-------------------------------------------------------------------------------|
| Overview SysInfo Network CPU I/O Software-blades Hardware-Health Advanced     |
|-------------------------------------------------------------------------------|
| Overview Top-Protocols Top-Connections Spikes                                  |
|-------------------------------------------------------------------------------|
| Host                                                                           |
|- More info available by scrolling up -----------------------------------------|
| ------------------------------------------------------------------------------ |
| CPU:                                                                           |
|                                                                                |
|   CPU Type             User  System  Idle       I/O wait    Interrupts       |
|     0 CoreXL_SND         0%      0%  100%             0%         3,852       |
|     1 CoreXL_FW          0%      0%  100%             0%         3,855       |
|     2 CoreXL_FW          0%      0%  100%             0%         3,856       |
|     3 CoreXL_FW          1%      1%   99%             0%         1,928       |
|     4 OTHER              0%      0%  100%             0%         1,928       |
|     5 CoreXL_FW          0%      1%   99%             0%         3,856       |
|     6 CoreXL_FW          2%      1%   98%             0%         3,856       |
|     7 CoreXL_FW          1%      0%   99%             0%         3,856       |
|                                                                                |
|-------------------------------------------------------------------------------|
```

p. 206: *Correction:* The "Performance Optimization" screen of the Gaia web interface shown on this page is no longer available on firewalls that are utilizing the new Gaia 3.10 kernel.

p. 211: The **ifconfig** command has been deprecated, instead of **ifconfig -a** use the command **ip -s a**

p. 214: The permanent URL link to Tobias Lachmann's Check Point Appliance Hardware List has changed to the following: https://lwf.fink.sh/check-point-appliance-hardware-lachmann-list-permanent/

p. 216: If the **fw ctl multik stat** command shows a large imbalance of connections on the Firewall Worker cores, it is probably caused by elephant flows/heavy connections. The goal of the Dynamic Dispatcher is to balance CPU load across the Firewall Workers, not necessarily the number of concurrent connections per Firewall Worker. However it is possible to override this behavior in the Dynamic Dispatcher, and try to balance by number of connections instead. This is not recommended for most environments, but if you want to modify this behavior, add the following two variables to **$FWDIR/boot/modules/fwkern.conf** and reboot the firewall:

```
fwmultik_enable_round_robin=1
fwmultik_enable_increment_first=1
```

p. 222: *Correction:* The **/sbin/cpuinfo** command that was used under Gaia kernel 2.6.18 to determine if SMT/Hyperthreading is enabled no longer works in Gaia kernel 3.10, use this command instead: **cat /sys/devices/system/cpu/smt/active**

1=active
0=inactive

p. 225: Note that in R80.40 GA when running **fwaccel stats -s**, you will see several additional processing paths not shown in the book. Some paths such as PSLXL and CPASXL appear to have been further "sliced and diced" for statistical purposes:

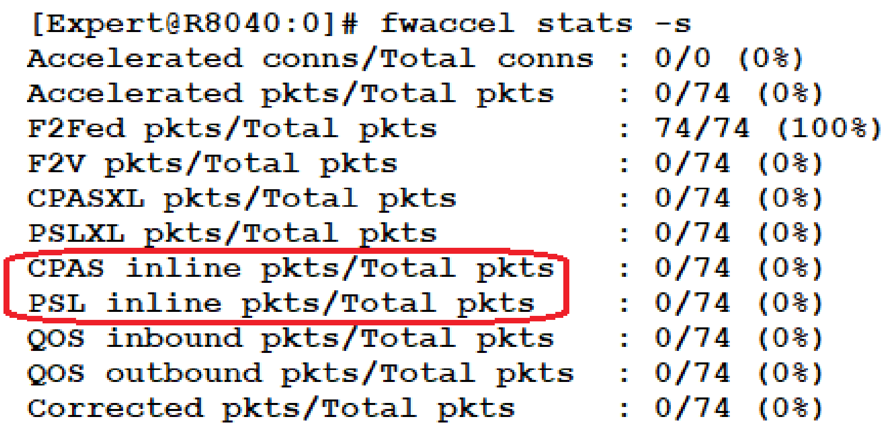

```
[Expert@R8040:0]# fwaccel stats -s
Accelerated conns/Total conns : 0/0 (0%)
Accelerated pkts/Total pkts   : 0/74 (0%)
F2Fed pkts/Total pkts         : 74/74 (100%)
F2V pkts/Total pkts           : 0/74 (0%)
CPASXL pkts/Total pkts        : 0/74 (0%)
PSLXL pkts/Total pkts         : 0/74 (0%)
CPAS inline pkts/Total pkts   : 0/74 (0%)
PSL inline pkts/Total pkts    : 0/74 (0%)
QOS inbound pkts/Total pkts   : 0/74 (0%)
QOS outbound pkts/Total pkts  : 0/74 (0%)
Corrected pkts/Total pkts     : 0/74 (0%)
```

This “dividing up” of processing paths continued through the Jumbo HFAs issued for R80.40, here is what they now look like in R80.40 Jumbo HFA Take 69:

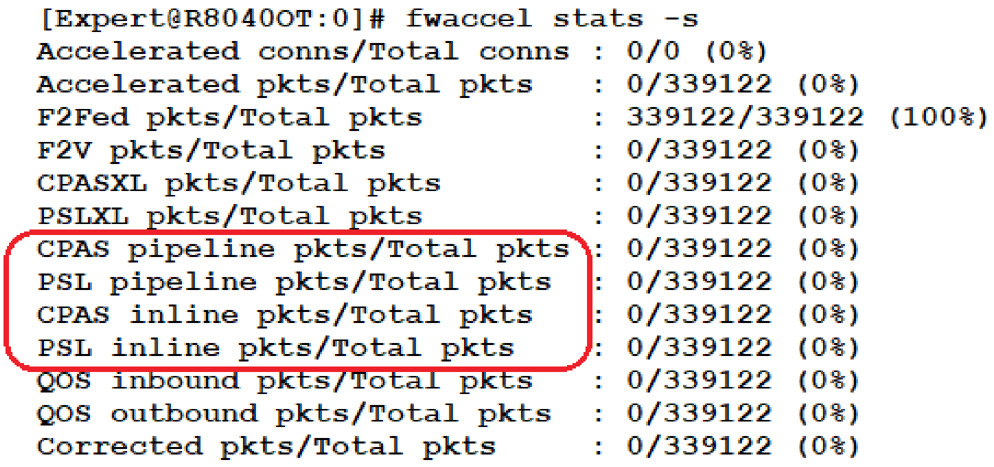

```
[Expert@R8040OT:0]# fwaccel stats -s
Accelerated conns/Total conns : 0/0 (0%)
Accelerated pkts/Total pkts   : 0/339122 (0%)
F2Fed pkts/Total pkts         : 339122/339122 (100%)
F2V pkts/Total pkts           : 0/339122 (0%)
CPASXL pkts/Total pkts        : 0/339122 (0%)
PSLXL pkts/Total pkts         : 0/339122 (0%)
CPAS pipeline pkts/Total pkts : 0/339122 (0%)
PSL pipeline pkts/Total pkts  : 0/339122 (0%)
CPAS inline pkts/Total pkts   : 0/339122 (0%)
PSL inline pkts/Total pkts    : 0/339122 (0%)
QOS inbound pkts/Total pkts   : 0/339122 (0%)
QOS outbound pkts/Total pkts  : 0/339122 (0%)
Corrected pkts/Total pkts     : 0/339122 (0%)
```

I believe the fundamental nature of the CPASXL and PSLXL paths have not changed in R80.40, and how the traffic is being handled through these two paths is being “sliced and diced” in the output for statistical purposes. In regards to firewall performance tuning, simply treat all paths with “CPAS” in their name as being part of the CPASXL path as specified in the book, and treat all paths with “PSL” in their name as being part of the PSLXL path for tuning purposes as specified in the book.

p. 234: To reiterate point #2 on this page, running **fw ctl chain** clearly shows that all inbound packets come through SecureXL first in R80.20+; this command also gives some extra insight into the "order of operations" for SecureXL under R80.20+:

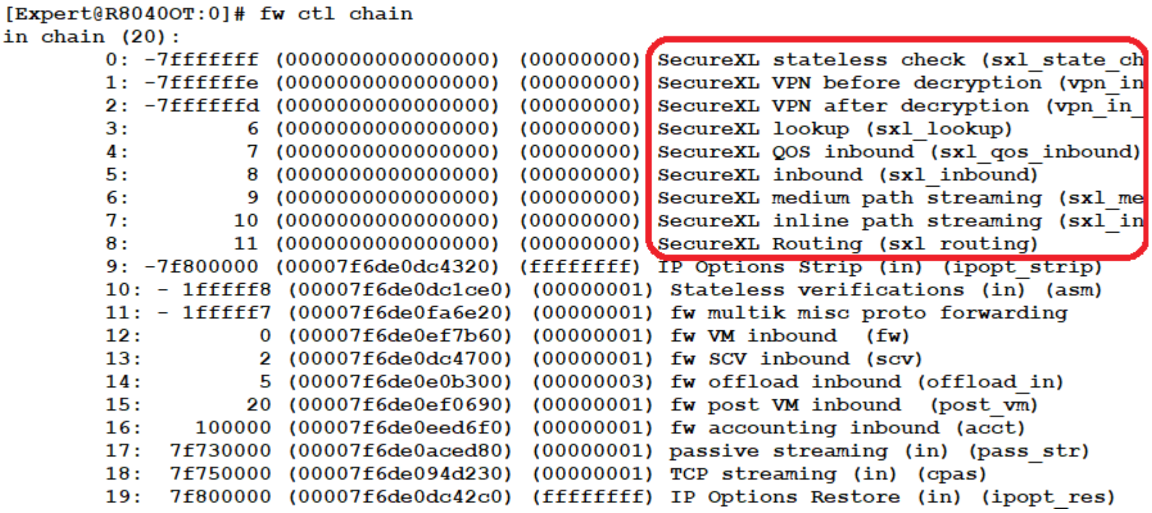

```
[Expert@R8040OT:0]# fw ctl chain
in chain (20):
        0: -7fffffff (0000000000000000) (00000000) SecureXL stateless check (sxl_state_ch
        1: -7ffffffe (0000000000000000) (00000000) SecureXL VPN before decryption (vpn_in
        2: -7ffffffd (0000000000000000) (00000000) SecureXL VPN after decryption (vpn_in_
        3:         6 (0000000000000000) (00000000) SecureXL lookup (sxl_lookup)
        4:         7 (0000000000000000) (00000000) SecureXL QOS inbound (sxl_qos_inbound)
        5:         8 (0000000000000000) (00000000) SecureXL inbound (sxl_inbound)
        6:         9 (0000000000000000) (00000000) SecureXL medium path streaming (sxl_me
        7:        10 (0000000000000000) (00000000) SecureXL inline path streaming (sxl_in
        8:        11 (0000000000000000) (00000000) SecureXL Routing (sxl routing)
        9: -7f800000 (00007f6de0dc4320) (ffffffff) IP Options Strip (in) (ipopt_strip)
        10: - 1fffff8 (00007f6de0dc1ce0) (00000001) Stateless verifications (in) (asm)
        11: - 1fffff7 (00007f6de0fa6e20) (00000001) fw multik misc proto forwarding
        12:         0 (00007f6de0ef7b60) (00000001) fw VM inbound  (fw)
        13:         2 (00007f6de0dc4700) (00000001) fw SCV inbound (scv)
        14:         5 (00007f6de0e0b300) (00000003) fw offload inbound (offload_in)
        15:        20 (00007f6de0ef0690) (00000001) fw post VM inbound  (post_vm)
        16:    100000 (00007f6de0eed6f0) (00000001) fw accounting inbound (acct)
        17:  7f730000 (00007f6de0aced80) (00000001) passive streaming (in) (pass_str)
        18:  7f750000 (00007f6de094d230) (00000001) TCP streaming (in) (cpas)
        19:  7f800000 (00007f6de0dc42c0) (ffffffff) IP Options Restore (in) (ipopt_res)
```

p. 239: There is a known issue with the TLS parser that will cause an excessive amount of traffic to take the Medium Path (PSLXL) when it should instead be fully accelerated in the Accelerated Path (SXL). This only occurs when the following combinations of blades (ONLY these blades and no others) are enabled on the firewall: **Firewall** and/or **Identity Awareness** and/or **Monitoring** and/or **IPSec VPN**. With any combination of only these blades enabled, almost all traffic should be eligible for full acceleration by SecureXL. However in this specific case the TLS parser is improperly invoked by the firewall when it is not actually needed, thus forcing large amounts of traffic into the PSLXL path. If you have only these blades enabled and high PSLXL, you can verify the status of the TLS parser with the **fw ctl get int tls_parser_enable** command, 1 means the TLS parser is enabled. *Note that a fix for this issue is currently under development and will be distributed via Jumbo HFA for R80.20 and all later versions. It is NOT RECOMMENDED to manually tamper with the state of the TLS parser except under the guidance of Check Point TAC, as this can cause further problems when new blades are* *enabled for the first time. This issue is presented as a corner case where there is perpetually high PSLXL path utilization for no apparent reason.*

More info and how to disable the TLS parser is here: sk166700: High CPU after upgrade from R77.x to R80.x when running only Firewall and Monitoring blades. Note that even if the TLS parser is disabled as described in the SK, TCP 443 (https) and TCP 445 (microsoft-ds) traffic will still go PSLXL anyway as revealed in this related CheckMates thread: https://community.checkpoint.com/t5/General-Topics/First-impressions-R80-30-on-gateway-one-step-forward-one-or-two/m-p/72593

p. 239: One topic that was not covered in the book due to its rarity is a sizable CPU utilization imbalance among the SND/IRQ cores (instead of Firewall Workers). Normally this should not happen if Multi-Queue is enabled on all interfaces (which is the case under Gaia 3.10 by default except for the management interface). However an issue with distribution of Remote Access VPN connections can cause a problematic SND/IRQ CPU imbalance that will seriously impact overall firewall performance: sk165853: High CPU usage on one CPU core when the number of Remote Access users is high.

p. 240: *Correction:* The **/sbin/cpuinfo** command that was used under Gaia kernel 2.6.18 to determine if SMT/Hyperthreading is enabled no longer works in Gaia kernel 3.10, use this command instead: **cat /sys/devices/system/cpu/smt/active**

1=active

0=inactive

p. 241: SMT/Hyperthreading is now supported on open hardware (i.e. not Check Point firewall appliances) using the Gaia 3.10 kernel for the first time starting in R80.40 Jumbo HFA 48+. Note however that from a licensing perspective on open hardware, each logical core (of which there are usually two for each physical core) *will be considered as* *another physical core that must be separately licensed.* The “container” portion of a firewall license specifies the number of cores that a firewall is allowed to used for traffic processing. Example: a 5900 series appliance has 8 physical cores and the included license container for an appliance permits the use of all logical cores even if SMT/Hyperthreading is enabled. That is NOT how it works on an open hardware firewall. If SMT/Hyperthreading is enabled on an 8-core open hardware firewall there will now be 16 logical cores, *and the open hardware firewall must upgrade its container license from 8 cores to 16 cores to use all of them.* Considering that enabling SMT/Hyperthreading grants a roughly 30% performance increase, with an open hardware firewall in this scenario you would be paying for 8 more physical cores yet only really getting about 30% of that performance. If at all possible on open hardware firewalls, add more *physical* cores first instead of logical ones via SMT/Hyperthreading! Further information can be found in this informative CheckMates post by Dr. Dorit Dor: https://community.checkpoint.com/t5/Enterprise-Appliances-and-Gaia/Does-R80-40-support-HP-DL380-G10/m-p/74952/highlight/true#M5794

p. 244: On some of the newer appliances, the menu option in **cpconfig** to control the state of SMT/Hyperthreading will be missing, and it appears SMT/Hyperthreading cannot be manually disabled. This is due to a BIOS limitation for that particular appliance; contact the Check Point TAC if you wish to disable SMT/Hyperthreading.

p. 247: *Clarification:* The **fast_accel** feature is available “out of the box” in version R80.40, and no Jumbo HFAs are required to use it.

p. 251: If you wish to prohibit fragments altogether for performance reasons, you can do so directly (and more efficiently) in SecureXL with this command (but all the caveats to doing so mentioned in the book still apply – careful!): **fwaccel dos config set --enable-drop-frags**

p. 253: *Correction:* The F/f flags no longer appear in the output of **fwaccel conns** starting in version R80.30; if a connection is being handled in the F2F path it does not appear at all in the output of **fwaccel conns**, so commands such as **fw ctl conntab** and **fw tab -u -t connections -f** must be used to verify the presence of an F2F-handled connection.

p. 266: To further clarify the proper commands to disable antispoofing "on the fly" based on code version and Jumbo HFA level:

R80.10 and earlier:

```
fw ctl set int fw_antispoofing_enabled 0
sim feature anti_spoofing off; fwaccel off; fwaccel on
```

R80.20 GA through Jumbo HFA Take 102: Impossible

R80.30 (kernel 2.6.18) GA through Jumbo Take 75: Impossible

R80.20 Jumbo HFA Take 103+, R80.30 (kernel 3.10) GA, R80.30 (kernel 2.6.18) Jumbo HFA 76+, R80.40 GA:

```
fw ctl set int fw_antispoofing_enabled 0
fw ctl set int sim_anti_spoofing_enabled 0 -a
```

p. 269: Note that if an exception is created for one of the 39 IPS "Core" Protections, it may interfere with the proper enforcement of Geo Policy. R80.20+ Geo Updatable Objects are not impacted by this issue – another great reason to use them instead of Geo Policy! sk164916: Geo Protection does not block countries.

p. 281: Note that the Dynamic Block Lists mentioned in the book are heavily dependent on DNS working properly. In *ALL* rules permitting DNS through your firewall, make sure you include both the **domain-udp** (UDP port 53) *AND* the **domain-tcp** (TCP port 53) services! The **domain-tcp** service is used in cases where the original DNS response is too large for one UDP port 53 packet. If your firewall policy does not permit the **domain-tcp** service alongside **domain-udp**, it can cause strange-looking traffic stalls and other "random" performance issues that are very difficult to diagnose! Further reading: https://community.checkpoint.com/t5/General-Management-Topics/Domain-Objects-FQDN-An-Unofficial-ATRG/m-p/72958

p. 287: If when running the **fwaccel stat** command you see "disabled by Firewall Layer ---" with no rule number shown, the name of the Network/Firewall policy layer as defined in the SmartConsole is too long. If it is shortened to under 32 characters the **fwaccel stat** display should work properly: sk145533: "Layer ---" is displayed instead of specific layer name and rule number in output of 'fwaccel stat'.

p. 291: Note that specifying an explicit SOURCE port number in a service's properties (this is not common) can also halt rule templating by SecureXL:

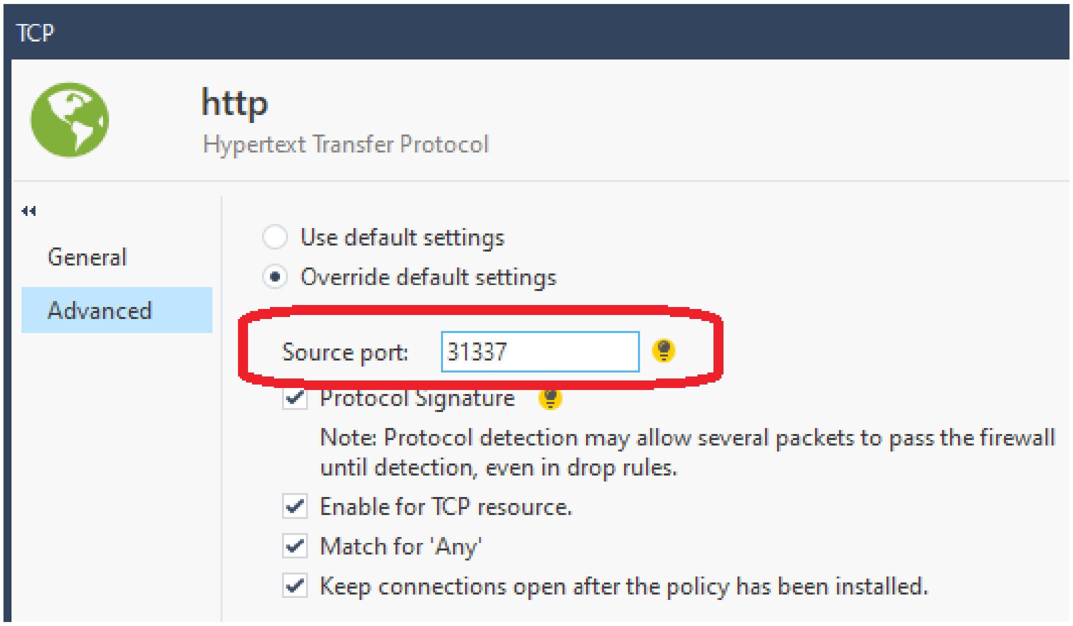

TCP

http

Hypertext Transfer Protocol

○ Use default settings

● Override default settings

Source port: 31337

☑ Protocol Signature

Note: Protocol detection may allow several packets to pass the firewall until detection, even in drop rules.

☑ Enable for TCP resource.

☑ Match for 'Any'

☑ Keep connections open after the policy has been installed.

p. 304: The Hide NAT port allocation mechanism was enhanced in version R80.40 for firewalls with 6+ cores, dynamically pooling available source ports among the Firewall Workers thus reducing the chance of running out of source ports on a particular Firewall Worker core. Further information: sk156852: "NAT Hide Failure" error in SmartLog / SmartView Tracker. In addition, new screens to monitor NAT state were added to **cpview** in R80.40:

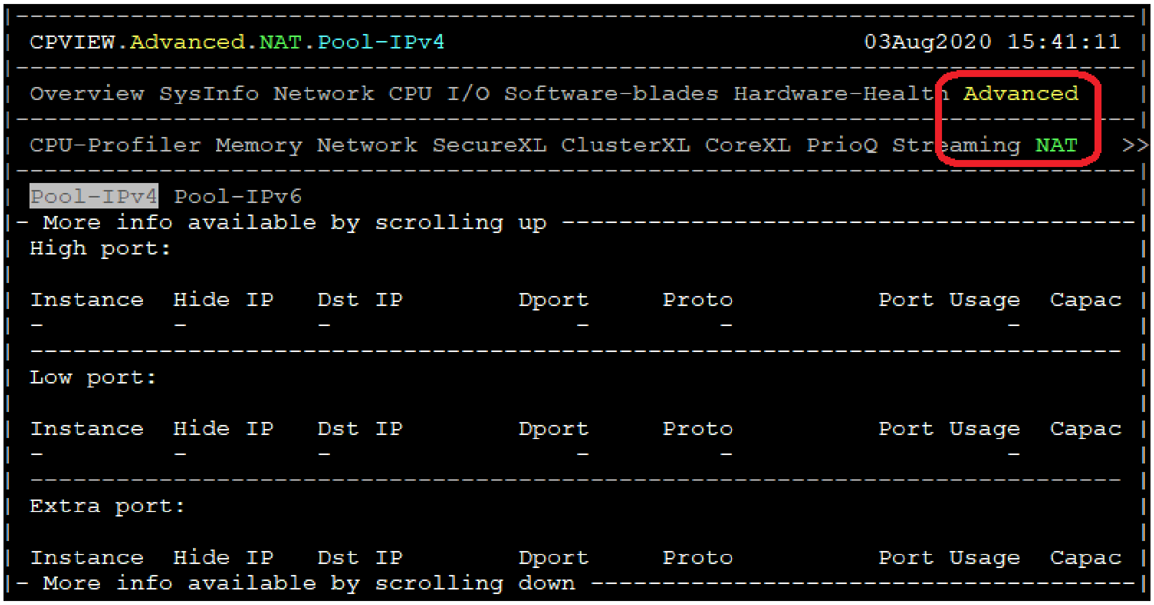

```
|---------------------------------------------------------------------------|
| CPVIEW.Advanced.NAT.Pool-IPv4                          03Aug2020 15:41:11 |
|---------------------------------------------------------------------------|
| Overview SysInfo Network CPU I/O Software-blades Hardware-Health Advanced |
|---------------------------------------------------------------------------|
| CPU-Profiler Memory Network SecureXL ClusterXL CoreXL PrioQ Streaming NAT  >>
|---------------------------------------------------------------------------|
| Pool-IPv4 Pool-IPv6                                                       |
|- More info available by scrolling up -------------------------------------|
| High port:                                                                |
|                                                                           |
| Instance  Hide IP   Dst IP        Dport     Proto          Port Usage  Capac |
| -         -         -                 -         -                   -       |
|                                                                           |
| ------------------------------------------------------------------------- |
| Low port:                                                                 |
|                                                                           |
| Instance  Hide IP   Dst IP        Dport     Proto          Port Usage  Capac |
| -         -         -                 -         -                   -       |
|                                                                           |
| ------------------------------------------------------------------------- |
| Extra port:                                                               |
|                                                                           |
| Instance  Hide IP   Dst IP        Dport     Proto          Port Usage  Capac |
|- More info available by scrolling down -----------------------------------|
```

There are also some Checkmates community-developed tools to provide even more NAT troubleshooting information such as the “NAT table (fwx_alloc) top users“ and “NAT table (fwx_alloc) specific NAT IP address analyses“ utilities by Kaspars Zibarts:

https://community.checkpoint.com/t5/API-CLI-Discussion-and-Samples/NAT-table-fwx-alloc-top-users/m-p/78543

https://community.checkpoint.com/t5/API-CLI-Discussion-and-Samples/NAT-table-fwx-alloc-specific-NAT-IP-address-analyses/m-p/89644

p. 306: Check Point has created an official SK detailing the "many-to-fewer" Hide NAT setup shown in the book: sk142833: How to create manual NAT rules in Many-To-Few mode.

p. 318: If possible, try to avoid overriding the default advanced settings for a service as shown here:

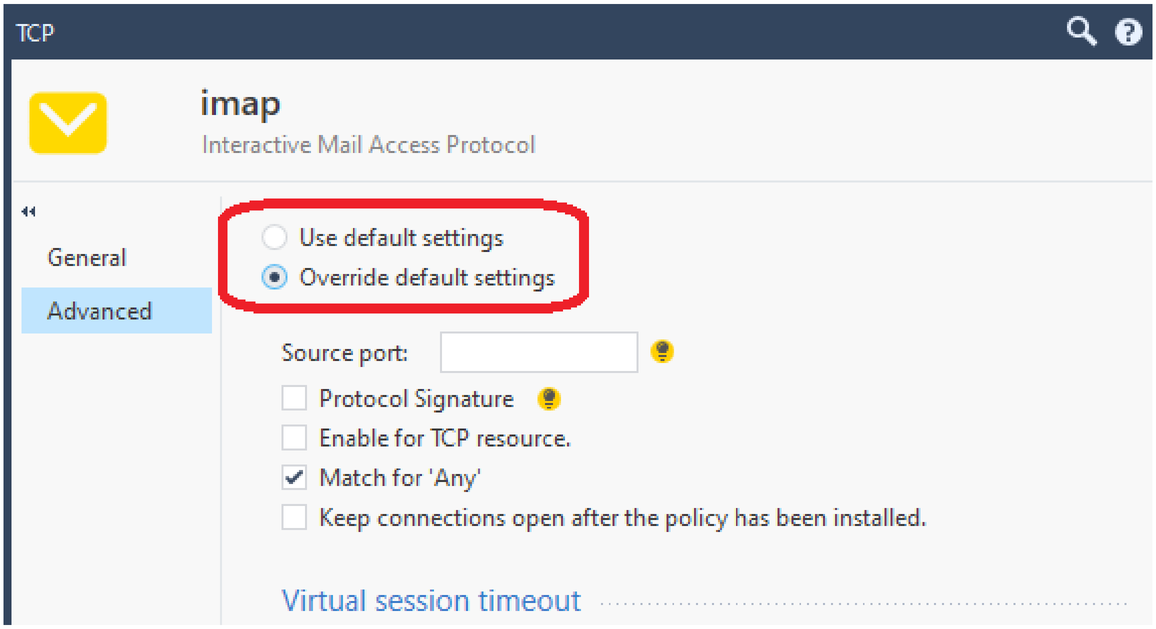

TCP

imap

Interactive Mail Access Protocol

General

Advanced

○ Use default settings
● Override default settings

Source port:

□ Protocol Signature
□ Enable for TCP resource.
☑ Match for 'Any'
□ Keep connections open after the policy has been installed.

Virtual session timeout

Overriding the default settings of a service has been reported by CheckMates users as adversely impacting the ability for SecureXL to fully accelerate traffic in some circumstances. This effect seems to vary depending on the service in question and firewall code version, so if traffic that should be fully accelerated is not, check the Advanced screen of the matching service object. If an override is present try selecting "Use default settings", reinstall policy, start NEW connections matching the service, then check acceleration status of those new connections with **fwaccel conns**. This issue is somewhat alluded to in this SK: sk101232: Connections are dropped as Out-of-State after some idle time when SecureXL is enabled. This all assumes of course that there was not some other factor that would have kept the connection from being fully accelerated anyway, such as requiring passive or active steaming inspection on a Firewall Worker core. Note that TCP/443 and TCP/445 traffic can never be fully accelerated by SecureXL and will always go at least PSLXL (see the first **p. 231** note earlier in this addendum).

p. 332: If AES-NI is not showing up when running the **dmesg** command on some of the newer Check Point appliances, AES-NI may have been be rolled up under another processor extension, so try this alternate command: **dmesg | grep -i sha**

```
[Wed Apr 15 22:52:55 2020] sha1_ssse3: Using SHA-NI optimized SHA-1 implementation
[Wed Apr 15 22:52:55 2020] sha256_ssse3: Using SHA-256-NI optimized SHA-256 implementation
```

p. 333: *Correction:* There is a way to see if NAT-T is in use with IPSec VPNs without taking a packet capture and looking for UDP port 4500, and it is via the SmartView Monitor "Tunnels on Gateway" screen:

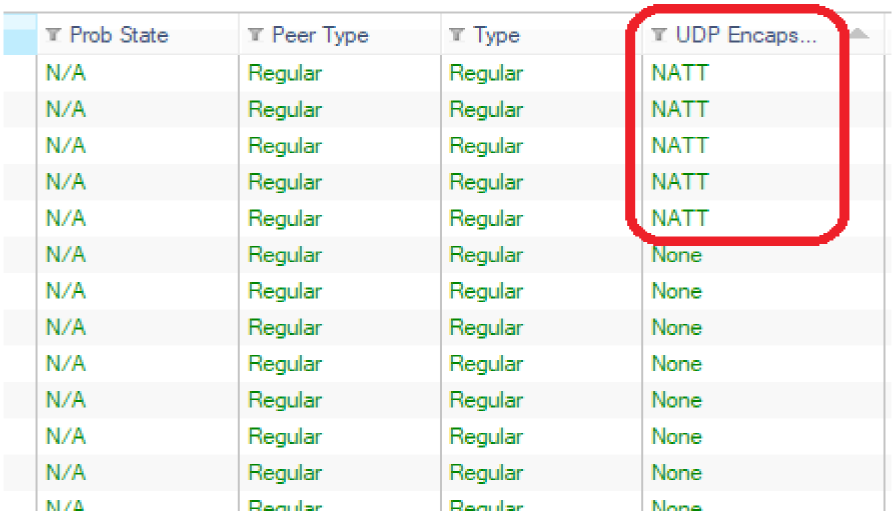

Prob State	Peer Type	Type	UDP Encaps...
N/A	Regular	Regular	NATT
N/A	Regular	Regular	NATT
N/A	Regular	Regular	NATT
N/A	Regular	Regular	NATT
N/A	Regular	Regular	NATT
N/A	Regular	Regular	None
N/A	Regular	Regular	None
N/A	Regular	Regular	None
N/A	Regular	Regular	None
N/A	Regular	Regular	None
N/A	Regular	Regular	None
N/A	Regular	Regular	None
N/A	Regular	Regular	None

Thanks to reader Tobias Moritz for providing this correction and the screenshot.

p. 334: R80.40 introduced the ability to specify custom VPN domains on a per-VPN Community basis. This new feature can be leveraged to maximize VPN tunnel sharing, improve performance, and prevent the creation of an excessive number of IPSec tunnels. As long as the SMS/MDS is version R80.40, it has been reported at CheckMates that this per-VPN Community VPN domain feature works with older gateways such as R80.30.

p. 352: To clarify, here are the proper log filters for viewing logs created by the four "classes" of IPS Protections in SmartConsole:

- IPS ThreatCloud Protection logs are matched by filter **blade:ips**
- Core Activations logs are matched by filter **blade:ips**
- Inspection Settings logs are matched by filter **blade:firewall**, but the Protection Type is "IPS"
- Geo Policy logs are matched by filter **blade:firewall**

p. 362: Note that adding exceptions for IPS Core Protections to potentially improve performance (and getting these exceptions to work correctly) can be a bit tricky, so see this SK for the correct procedure: sk162493: Adding an exception to IPS Core Protection does not take effect.

p. 365: There is a great new tool called "TailoredSafe" that can be used to analyze all IPS protections currently in Detect mode, and move them to Prevent or Inactive in one fell swoop. There is an instructional video called "Moving from Detect to Prevent" as well. While manually slogging through thousands of protections set for Detect as detailed in the book was certainly a great deterrent, there is now NO EXCUSE for leaving IPS protections in Detect mode with the resulting heavy firewall performance impact:

- https://community.checkpoint.com/t5/IPS-Anti-Virus-Anti-Bot-Anti/Moving-from-Detect-to-Prevent-TechTalk-Video-Slides-and-Q-amp-A/m-p/71878
- sk164812: Tailored Safe Extension: Tailor-made IPS profile

p. 385: Here are some other good Threat Prevention testing sites in addition to CheckMe:

- http://www.threat-cloud.com/test/files/MediumConfidenceBot.html
- http://www.threat-cloud.com/test/files/HighConfidenceBot.html
- http://sc1.checkpoint.com/za/images/threatwiki/pages/TestAntiBotBlade.html
- http://poc-files.threat-cloud.com/demo/demo.doc
- http://www.eicar.org/download/eicar_com.zip

p. 396: To see detailed statistics about packet handling in the CPASXL path (great for troubleshooting networking or other performance issues for traffic being handled in the CPASXL path), run the new **fw ctl cpasstat** command:

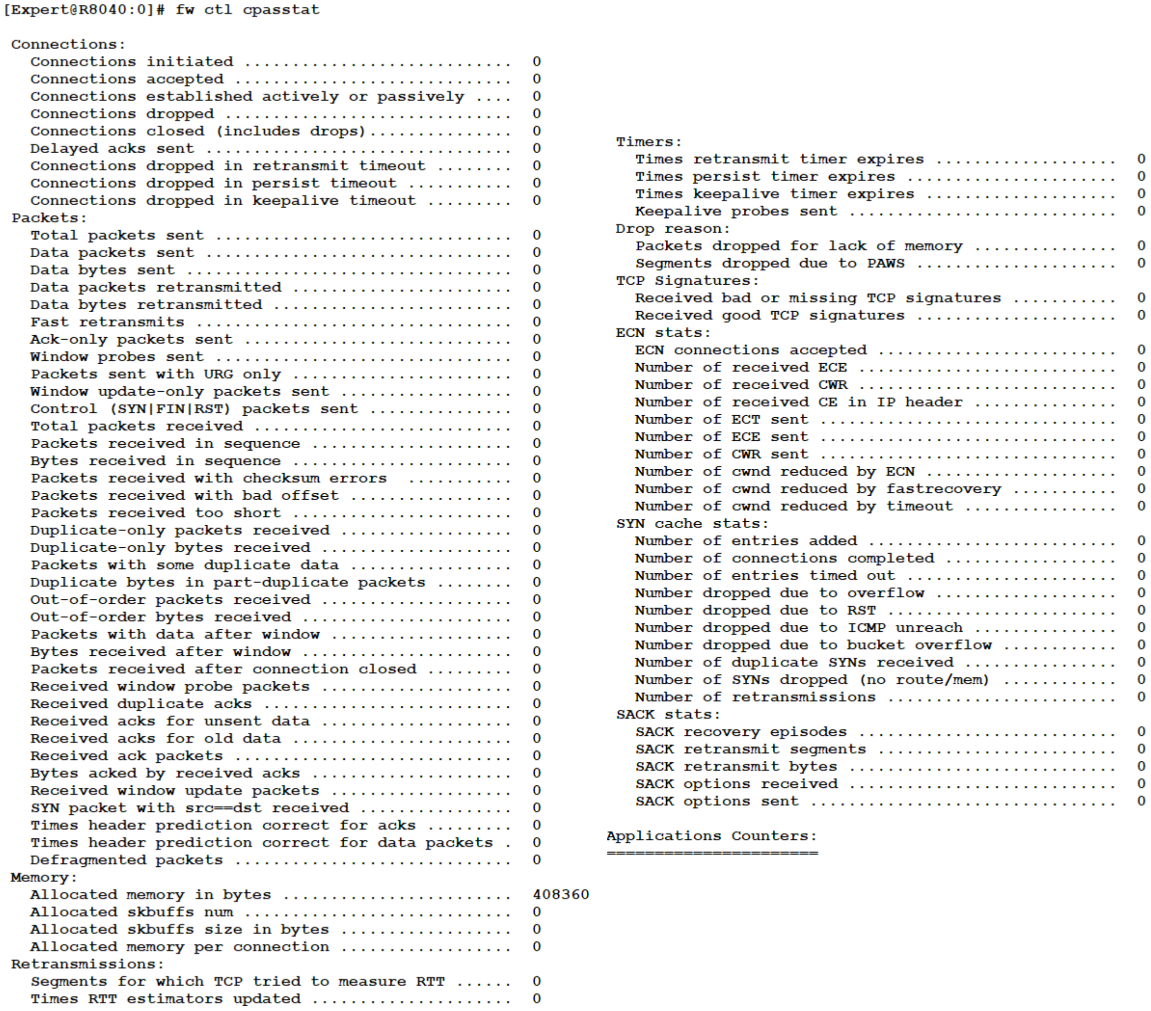

```
[Expert@R8040:0]# fw ctl cpasstat

 Connections:
   Connections initiated ............................  0
   Connections accepted .............................  0
   Connections established actively or passively ....  0
   Connections dropped ..............................  0
   Connections closed (includes drops)...............  0
   Delayed acks sent ................................  0
   Connections dropped in retransmit timeout ........  0
   Connections dropped in persist timeout ...........  0
   Connections dropped in keepalive timeout .........  0
 Packets:
   Total packets sent ...............................  0
   Data packets sent ................................  0
   Data bytes sent ..................................  0
   Data packets retransmitted .......................  0
   Data bytes retransmitted .........................  0
   Fast retransmits .................................  0
   Ack-only packets sent ............................  0
   Window probes sent ...............................  0
   Packets sent with URG only .......................  0
   Window update-only packets sent ..................  0
   Control (SYN|FIN|RST) packets sent ...............  0
   Total packets received ...........................  0
   Packets received in sequence .....................  0
   Bytes received in sequence .......................  0
   Packets received with checksum errors  ...........  0
   Packets received with bad offset .................  0
   Packets received too short .......................  0
   Duplicate-only packets received ..................  0
   Duplicate-only bytes received ....................  0
   Packets with some duplicate data .................  0
   Duplicate bytes in part-duplicate packets ........  0
   Out-of-order packets received ....................  0
   Out-of-order bytes received ......................  0
   Packets with data after window ...................  0
   Bytes received after window ......................  0
   Packets received after connection closed .........  0
   Received window probe packets ....................  0
   Received duplicate acks ..........................  0
   Received acks for unsent data ....................  0
   Received acks for old data .......................  0
   Received ack packets .............................  0
   Bytes acked by received acks .....................  0
   Received window update packets ...................  0
   SYN packet with src==dst received ................  0
   Times header prediction correct for acks .........  0
   Times header prediction correct for data packets .  0
   Defragmented packets .............................  0
 Memory:
   Allocated memory in bytes ........................  408360
   Allocated skbuffs num ............................  0
   Allocated skbuffs size in bytes ..................  0
   Allocated memory per connection ..................  0
 Retransmissions:
   Segments for which TCP tried to measure RTT ......  0
   Times RTT estimators updated .....................  0
```

```
 Timers:
   Times retransmit timer expires ...................  0
   Times persist timer expires ......................  0
   Times keepalive timer expires ....................  0
   Keepalive probes sent ............................  0
 Drop reason:
   Packets dropped for lack of memory ...............  0
   Segments dropped due to PAWS .....................  0
 TCP Signatures:
   Received bad or missing TCP signatures ...........  0
   Received good TCP signatures .....................  0
 ECN stats:
   ECN connections accepted .........................  0
   Number of received ECE ...........................  0
   Number of received CWR ...........................  0
   Number of received CE in IP header ...............  0
   Number of ECT sent ...............................  0
   Number of ECE sent ...............................  0
   Number of CWR sent ...............................  0
   Number of cwnd reduced by ECN ....................  0
   Number of cwnd reduced by fastrecovery ...........  0
   Number of cwnd reduced by timeout ................  0
 SYN cache stats:
   Number of entries added ..........................  0
   Number of connections completed ..................  0
   Number of entries timed out ......................  0
   Number dropped due to overflow ...................  0
   Number dropped due to RST ........................  0
   Number dropped due to ICMP unreach ...............  0
   Number dropped due to bucket overflow ............  0
   Number of duplicate SYNs received ................  0
   Number of SYNs dropped (no route/mem) ............  0
   Number of retransmissions ........................  0
 SACK stats:
   SACK recovery episodes ...........................  0
   SACK retransmit segments .........................  0
   SACK retransmit bytes ............................  0
   SACK options received ............................  0
   SACK options sent ................................  0

Applications Counters:
======================
```

p. 397: As of R80.40 Jumbo HFA Take 53+, the Visitor Mode implementation was moved from the vpnd process (which caused a performance-impacting "trip" to process space for this type of traffic) down into the Firewall Workers. This change is enabled by default, permits a much larger numbers of Remote Access VPN users to utilize Visitor Mode, and also substantially reduces the performance impact of Visitor Mode VPN traffic on the firewall. See the following SK for more details: sk168297: Remote Access VPN Visitor Mode in Kernel.

p. 398: When the book was originally published, under what circumstances User Space Firewall (USFW) was enabled by default “out of the box” was still evolving and there was large amounts of conflicting information being reported by users on CheckMates. As it turns out, whether USFW will be enabled by default is actually much more dependent on specific hardware/appliance type than Gaia kernel or number of cores. So assuming at least version R80.30:

- USFW is enabled by default on Check Point Appliance 2019 series (3600 [4 core], 3800 [8 core], 6XXX, 7XXX, 16XXX, 26XXX, 28XXX)
- USFW is enabled by default in any kind of virtualized environment like VMWare, regardless of the number of cores. Only 2 cores present in VMWare? USFW will be enabled, except for the next bullet point.
- A Check Point appliance or Virtual Machine deployed in standalone mode (SMS & firewall on the same box) will always have USFW *disabled*.
- “Bare Metal” Open Hardware Server (not VMWare) - Depends on number of cores, less than 35 cores USFW disabled by default, more than 35 cores USFW enabled by default.
- USFW is not enabled by default on Check Point appliance 2016 series (3100, 3200, 5XXX, 15XXX, 23XXX), except for model 23900 which has USFW enabled by default.
- USFW is not enabled by default on Check Point appliance 2012 series (2200, 4XXX, 12XXX, 13XXX, 21XXX). Note that most if not all of these 2012 series appliances reach end of support in 2022.

This criteria for whether USFW is enabled by default seems to have changed over time, which may explain why some early 16000 models didn’t have USFW enabled by default, along with some other apparent inconsistencies reported by CheckMates users. Check this recently-created SK for the latest updates, limitations, and under what circumstances you may want to manually change the default state of USFW: sk167052: Check Point User-Space **firewall** support for R80.30 3.10 and above.

p. 415: In R80.40 the HTTPS Inspection Policy was moved from the legacy SmartDashboard into the main SmartConsole, and the ability to specify that only certain blades/features apply to traffic matching a particular HTTPS Inspection rule was added. This opens up some interesting optimization possibilities by potentially excluding certain decrypted HTTPS traffic from inspection by some blades/features in R80.40+. In R80.30 and earlier all traffic matching an HTTPS Inspection rule with an action of "Inspect" would be inspected by all relevant blades thus incurring more overhead. This overhead can be avoided in R80.40 via the new "Blades" column if desired.

p. 417: Further clarifications have emerged concerning HTTPS Inspection Policy optimization. The HTTPS Inspection Policy is evaluated top-down, so try to order the rules as follows:

- Bypass by IP addresses only (applications/categories set to Any)
- Bypass with IP and/or applications/categories
- Rules with Inspect actions (blades can be selected individually or left at Any)
- Clean Up Rule (any any bypass)

The following screenshot shows these principles in action (as well as the new R80.40 "Blade" column):

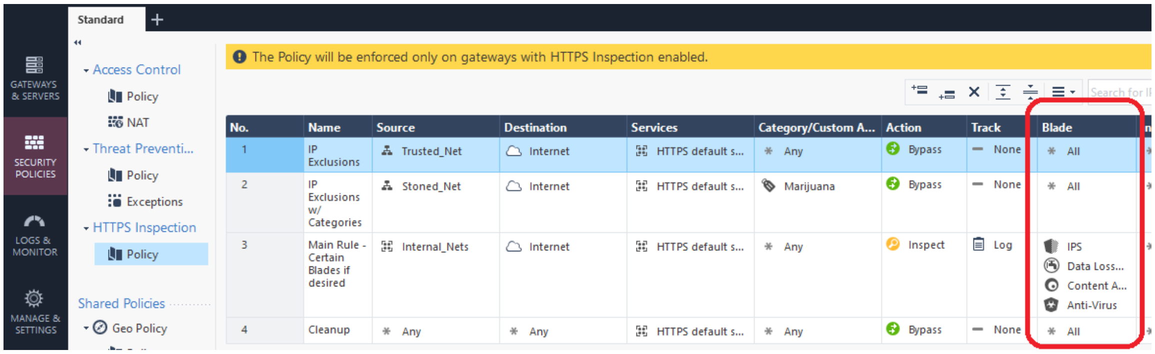

The Policy will be enforced only on gateways with HTTPS Inspection enabled.

No.	Name	Source	Destination	Services	Category/Custom A...	Action	Track	Blade
1	IP Exclusions	Trusted_Net	Internet	HTTPS default s...	Any	Bypass	None	All
2	IP Exclusions w/ Categories	Stoned_Net	Internet	HTTPS default s...	Marijuana	Bypass	None	All
3	Main Rule - Certain Blades if desired	Internal_Nets	Internet	HTTPS default s...	Any	Inspect	Log	IPS Data Loss... Content A... Anti-Virus
4	Cleanup	Any	Any	HTTPS default s...	Any	Bypass	None	All

Source: https://community.checkpoint.com/t5/General-Management-Topics/HTTPS-Inspection-Setup/td-p/83504

p. 417: After making a change with the **cipher_util** tool, you may need to perform a gateway policy installation to make the changes take effect.

p. 422: Note item #2 in the Introduction section of this addendum document; the “functionality gap” when dealing with elephant flows/heavy connections on a kernel mode vs. USFW-enabled firewall has been closed in the most recent R80.40 Jumbo HFAs.

p. 432: Note that the **fast_accel** feature will not work for connections that must go F2F for some other reason. If whatever condition is causing the traffic to go F2F can be rectified, the **fast_accel** feature will work for that matching traffic as expected.

p. 432: The Check Point Solutions Center (accessible through your local Check Point Sales Engineer) has a software feature available for certain code levels that can spread the processing of elephant flows/heavy connections across multiple Firewall Workers, but this capability is not in the main train of R80.40 (or earlier) code yet.

p. 440: Check Point employee Matan Ben David has created an excellent reference document at CheckMates called "How to Identify DDoS attack on Check Point Gear" that is absolutely invaluable reading even if you aren't currently suffering a DDoS attack: https://community.checkpoint.com/t5/Incident-Response/How-to-Identify-DDoS-attack-on-Check-Point-Gear/ba-p/85052

p. 441: One detail not mentioned in the book is how Aggressive Aging is handled when the Concurrent Connections Limit on the firewall is set to "Automatically" and not a fixed value such as the old default of 25,000 connections. Because the limit is almost always set to "Automatically" on today's firewalls (excluding VSX), Aggressive Aging only looks at memory utilization thresholds when deciding whether to activate.
sk122154: How is Aggressive Aging enforced when Concurrent Connections Capacity Limit is calculated automatically?

p. 442: Although I haven't noticed it myself, there are reports that if a very high level of firewall traffic is fully-accelerated (Accelerated pkts/sec in **fwaccel stats -s**), Aggressive Aging if enabled will cause much higher CPU loads on the SND/IRQ cores. Since the Aggressive Aging function is implemented by the Firewall Worker cores, it is speculated that there are a very large number of aging notifications being constantly sent into the SND/IRQ cores, since that is where almost all traffic is currently being handled in this case due to the high rate of fully-accelerated traffic. This effect seems more likely to happen in versions R80.10 and lower, but can also happen in versions R80.20+ depending on traffic type and delay. The most up-to-date information can be found in this CheckMates discussion:
https://community.checkpoint.com/t5/Next-Generation-Firewall/High-CPU-use-on-SND-cores-and-Aggressive-Aging/td-p/88761

p. 449: If using WinSCP to transfer files to and from a firewall or SMS/MDS using the Gaia 3.10 kernel, you may notice that file transfers are extremely slow. This is due to a problematic buffer optimization strategy in WinSCP that is enabled by default, and causes extremely slow transfers to and from the Gaia OS. sk163940: SCP/SFTP file transfer is very slow after upgrade of SSH. This issue only manifests itself when transferring large files to and from the gateway itself with SCP; this does not occur with WinSCP file transfers that are just transiting the gateway. Apparently this is due to an upgraded Gaia version of SSHD that does not care one bit for WinSCP's buffer size optimization strategy.

p. 450: Due to a recent change, Management Data Plane Separation can now be used on firewalls that have as few as 4 cores, not 8 as stated in the book.

p. 457: R80.40 introduced a new feature called "Revert to Version" that can be employed to undo all changes back to a specific date or session publication. This is similar to performing a full revert on an R77.30 or earlier SMS via Database Revision Control. The details of this new feature are covered in my updated CheckMates article "R80+ Change Control: A Visual Guide":
https://community.checkpoint.com/t5/Policy-Management/R80-Change-Control-A-Visual-Guide/m-p/39702

p. 459: Check Point has introduced a great new tool in R80.40 for tracking down excessive CPU utilization, especially if it is intermittent. It is called "CPU Spike Detective" and it is documented here: sk166454: CPU Spike Detective. Spike results can be viewed directly in **cpview** or in a variety of log files.

p. 473: R80.40 introduced a new feature called SmartTasks, that can be configured to send emails upon policy installation similarly to the SmartView Monitor-based method shown in the book.

p. 507: *Correction:* The **/sbin/cpuinfo** command that was used under Gaia kernel 2.6.18 to determine if SMT/Hyperthreading is enabled no longer works in Gaia kernel 3.10, use this command instead: **cat /sys/devices/system/cpu/smt/active**

1=active

0=inactive