IBM Software Group

zIIP-Assisted IPSec

Enterprise Network and Transformation Solutions

Jeannie Kristufek
kristufe@us.ibm.com

August 2008

© 2008 IBM Corporation

Agenda

- IPSec overview
- zSeries specialty engines
- zSeries cryptographic hardware
- What is zIIP-Assisted IPSec
- Configuring Communications Server for zIIP-Assisted IPSec
- Projecting zIIP Requirements
- Sample performance measurements

© 2008 IBM Corporation

IPSec Overview

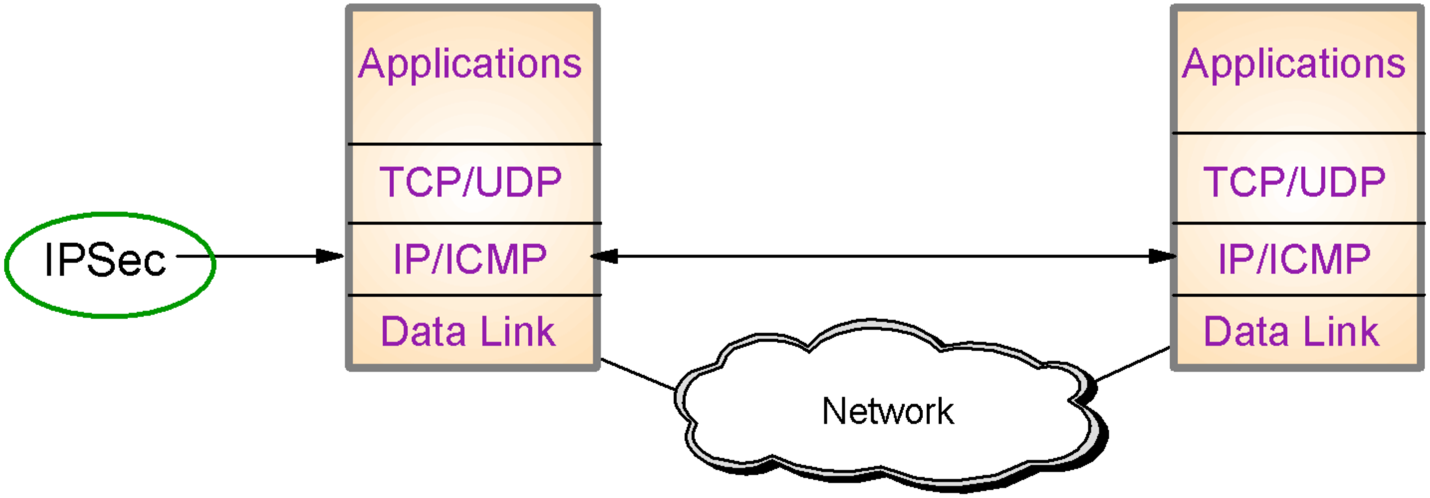

- Open network layer security protocol endorsed by IETF
- Provides authentication, integrity, and data privacy via IPSec security protocols
  - Authentication Header (AH) - provides authentication / integrity
  - Encapsulating Security Protocol (ESP) - provides data privacy with optional authentication/integrity
- Secures traffic between any two IP resources
  - Security Associations (SA)
- Management of crypto keys and security associations can be
  - manual
  - automated via key management protocol (IKE)

© 2008 IBM Corporation

IPSec Overview…

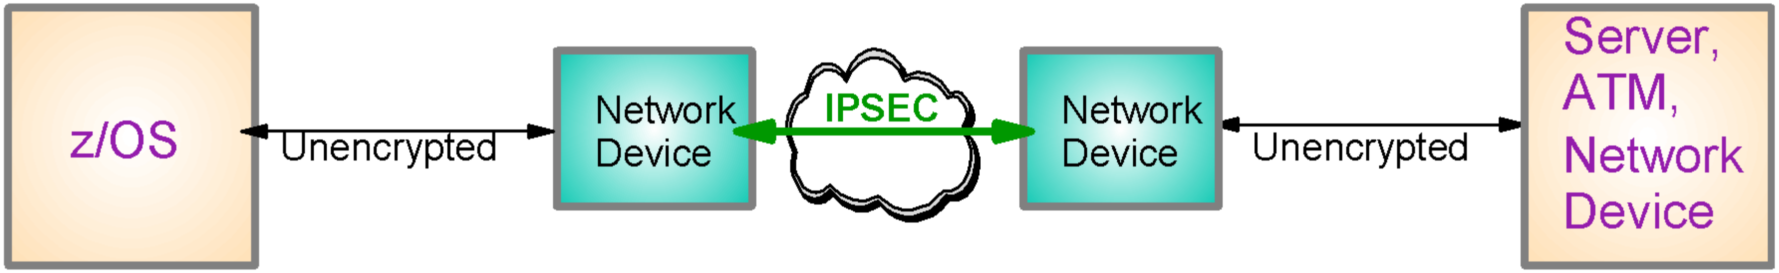

- IPSec provides end-to-end network encryption
- End-to-end network encryption is becoming more pervasive due to regulatory security policies
- End-to-end network encryption is also becoming a requirement for companies that outsource/share part of their network with business partners and need to have greater control of access to confidential data

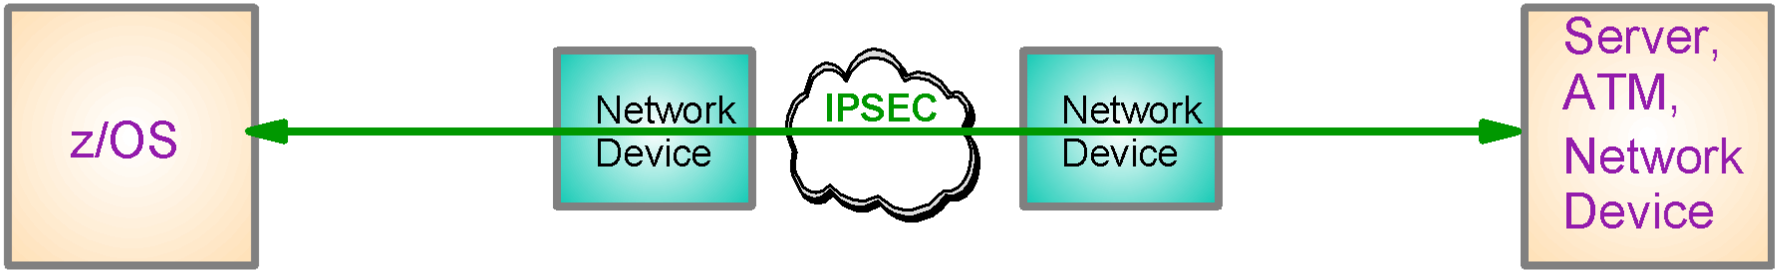

© 2008 IBM Corporation

IP Packet Filtering Overview

- Packet filtering at IP Layer
- Filter rules defined to match on inbound and outbound packets based on
  - packet information (IP address, port number, protocol)
  - network attributes (packet direction, link security)
  - time
- Possible actions
  - Permit
  - Deny
  - Permit with IPSec
  - Log (in combination with other actions)
- zIIP-Assisted IPSec only applies to “Permit with IPSec” action

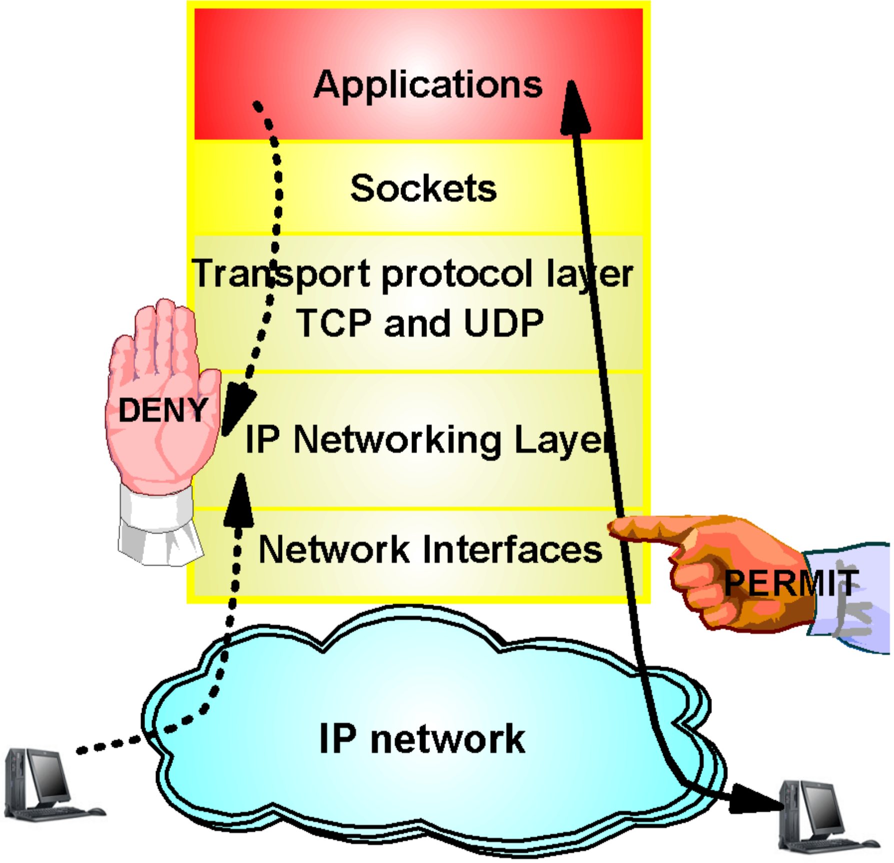

IP Packet Filtering for Enterprise Extender

- Enabling IPSEC for certain workloads requires at least a Permit filter rule for any other workloads
- Apar PK37839 addresses performance issues with enabling IP Filtering for Enterprise Extender workloads

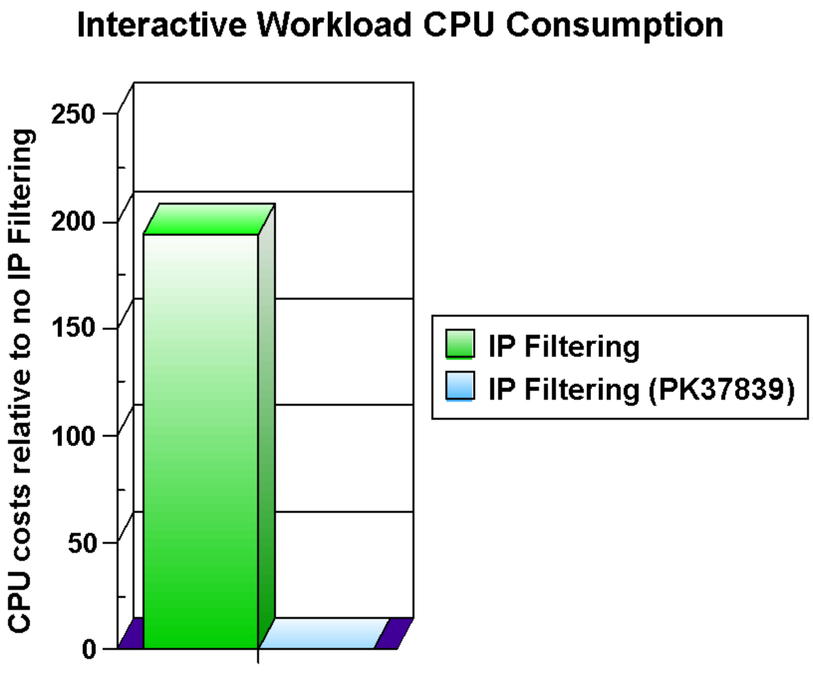

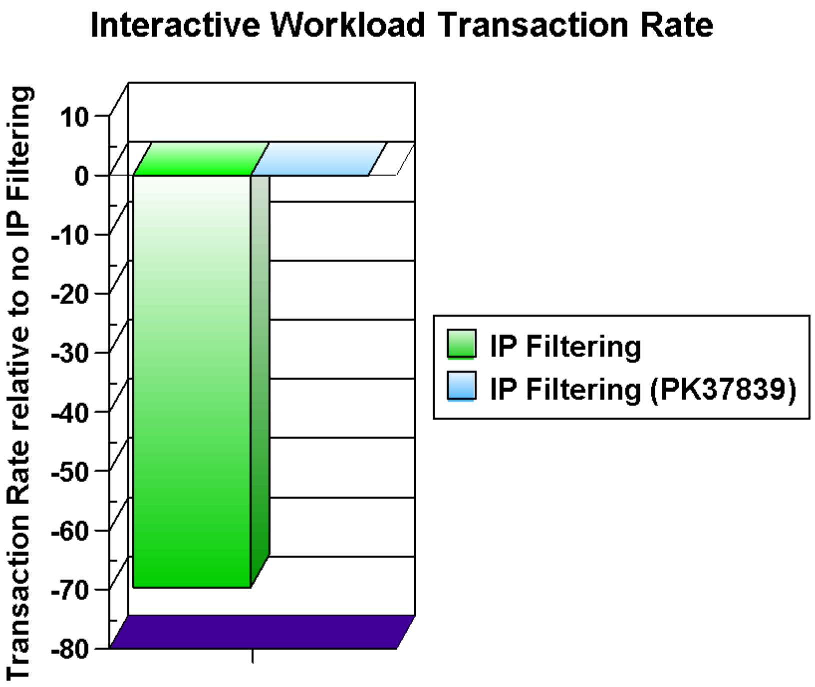

© 2008 IBM Corporation

AT-TLS Overview

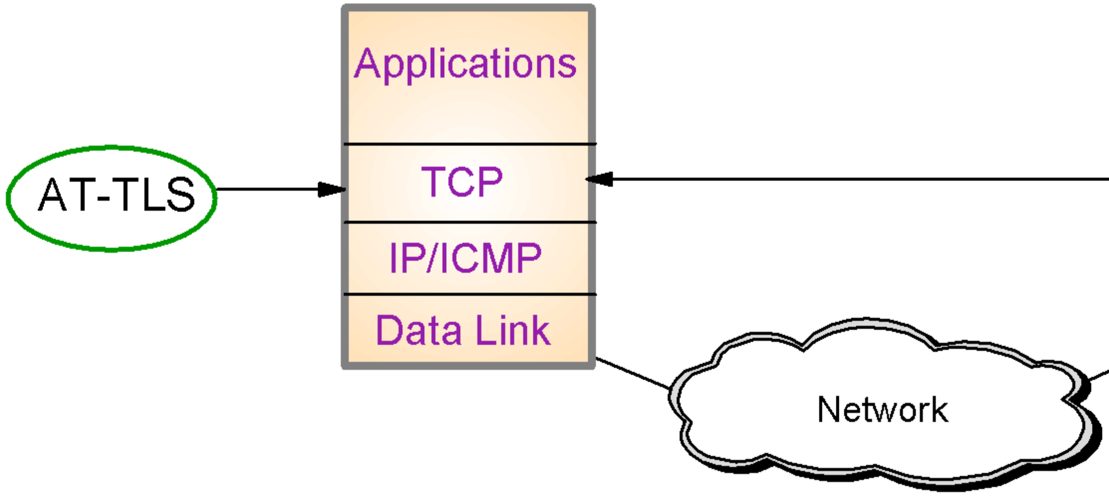

- Transport Layer Security (TLS) is defined by IETF
- Provides security services as a socket service
  - AT-TLS processing done at TCP layer without requiring any application changes
- Secures traffic between two TCP applications
  - SSL handshakes
- No zIIP assistance for AT-TLS encryption/decryption

© 2008 IBM Corporation

Specialty Engines

- Integrated Facility for Linux (IFL)
  - Provides additional processing capacity for Linux workloads without affecting IBM software charges
- System z Application Assist Processor (zAAP)
  - Provides ability to lower costs of CPU-intensive web-based applications (i.e. Java, XML)
- System z9 Integrated Information Processor (zIIP)
  - Provides ability to lower costs for select data and transaction processing workloads
  - DB2/DRDA exploits zIIPs for portions of their workloads
  - Communications Server exploits zIIPs for portions of their IPSec workloads

© 2008 IBM Corporation

Cryptographic Hardware

- Crypto coprocessors
  - Available on previous generations of zSeries
    - General CP with “built-in” crypto functions
  - Provides hardware encryption/decryption
    - Unit of work must be running on this processor
- Crypto cards
  - Available on z/990, z/9, and z/10
    - PCIX Cryptographic card (PCIXCC)
    - CryptoExpress2 card (CEX2C)
  - Provides hardware RSA signature generation/verification for peer authentication during IKE negotiations
- Hardware instructions
  - Available on z/990, z/9, and z/10
    - CP Assist for Cryptographic Function (CPACF)
  - Provides hardware encryption/decryption and authentication
    - Unit of work can be running on any general CP

© 2008 IBM Corporation

What is zIIP-Assisted IPSec?

- Even with zSeries specialized Crypto hardware, IPSec's data encryption/decryption and authentication processing can incur very heavy CPU consumption
- zIIP-Assisted IPSec allows for the movement of the bulk of Communications Server IPSec processing from general CPs to zIIPs
  - SRB-mode IPSec protocol traffic directed to zIIPs
    - Encryption/decryption, message authentication, and IPSEC header processing
    - Work is assigned to an independent WLM enclave
- Will provide CPU-busy relief on general CPs for customers already running IPSec on z/OS
- Makes z/OS IPSec deployment more attractive for customers concerned about IPSec CPU consumption
  - IBM does not impose software charges for zIIP capacity
- Does not replace CPACF
  - Simply performs CPACF instruction on zIIP rather than on a general CP

© 2008 IBM Corporation

Configuring Communications Server for zIIP-Assisted IPSec

- New GLOBALCONFIG statement in TCPIP profile
  - GLOBALCONFIG ZIIP IPSECURITY

- New IEAOPTxx statement in PARMLIB to control whether zIIP eligible IPSec work can spill over to general CPs
  - IIPHONORPRIORITY=YES
    - Allows IPSec work to run on general CPs if zIIP requests help
    - Default and recommended
  - IIPHONORPRIORITY=NO
    - All zIIP eligible work is contained on zIIPs
    - May result in throughput and/or response time degradation if zIIPs are heavily utilized
    - May be reasonable tradeoff in some environments where minimizing usage of general CPs is important

- Required apars
  - PK40178 (Communications Server enablement apar for zIIP support)
  - OA20045 (z/OS apar for IIPHONORPRIORITY support)

© 2008 IBM Corporation

Configuring Communications Server for zIIP-Assisted IPSec…

- Classify the IPSec independent enclave differently from the TCPIP address space
- IPSec SRBs may have longer execution times than other work directed to zIIPs
  - Running these SRBs at high priorities could lead to significant processor delays for other work scheduled for the zIIP
- If using zIIPs exclusively for IPSec, failing to classify the independent enclave results in it being assigned the SYSOTHER class
  - Any IPSec work spilled over onto general CPs will be running at this lowest priority

© 2008 IBM Corporation

Displaying zIIP Usage

- NETSTAT STATS command displays the IPv4 and IPv6 inbound/outbound packets processed on zIIPs

```
D TCPIP,,N,STATS
EZD0101I NETSTAT CS V1R9 TCPCS

IP STATISTICS (IPV4)
  …
  …
  FRAGMENTS CREATED                  = 0
  INBOUND PACKETS HANDLED BY ZIIP    = 8197069
  OUTBOUND PACKETS HANDLED BY ZIIP   = 323874
IPV6 STATISTICS
  …
  …
  FRAGMENTS CREATED                  = 0
  INBOUND PACKETS HANDLED BY ZIIP    = 154319
  OUTBOUND PACKETS HANDLED BY ZIIP   = 33273

IP GENERAL STATISTICS
  …
  …
  …
  …
UDP STATISTICS
  DATAGRAMS RECEIVED     = 0
  NO PORT ERRORS         = 13
  RECEIVE ERRORS         = 0
  DATAGRAMS SENT         = 13
END OF THE REPORT
```

Displaying zIIP Availability

- D M=CPU command displays the zIIP online/offline status

```
D M=CPU
IEE174I 01.35.25 DISPLAY M 277
PROCESSOR STATUS
ID  CPU                  SERIAL
00  +                     029B8E2094
01  +                     029B8E2094
02  +I                    029B8E2094


CPC ND = 002094.S38.IBM.02.000000029B8E
CPC SI = 2094.730.IBM.02.0000000000029B8E
CPC ID = 00
CPC NAME = RP569
LP NAME = RALNS42    LP ID =  2
CSS ID  = 0
MIF ID  = 2

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE

I        INTEGRATED INFORMATION PROCESSOR (zIIP)
CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
CPC SI  SYSTEM INFORMATION FROM STSI INSTRUCTION
CPC ID  CENTRAL PROCESSING COMPLEX IDENTIFIER
CPC NAME CENTRAL PROCESSING COMPLEX NAME
LP NAME  LOGICAL PARTITION NAME
LP ID    LOGICAL PARTITION IDENTIFIER
```

© 2008 IBM Corporation

Potential benefit of zIIP-Assisted IPSec – Interactive

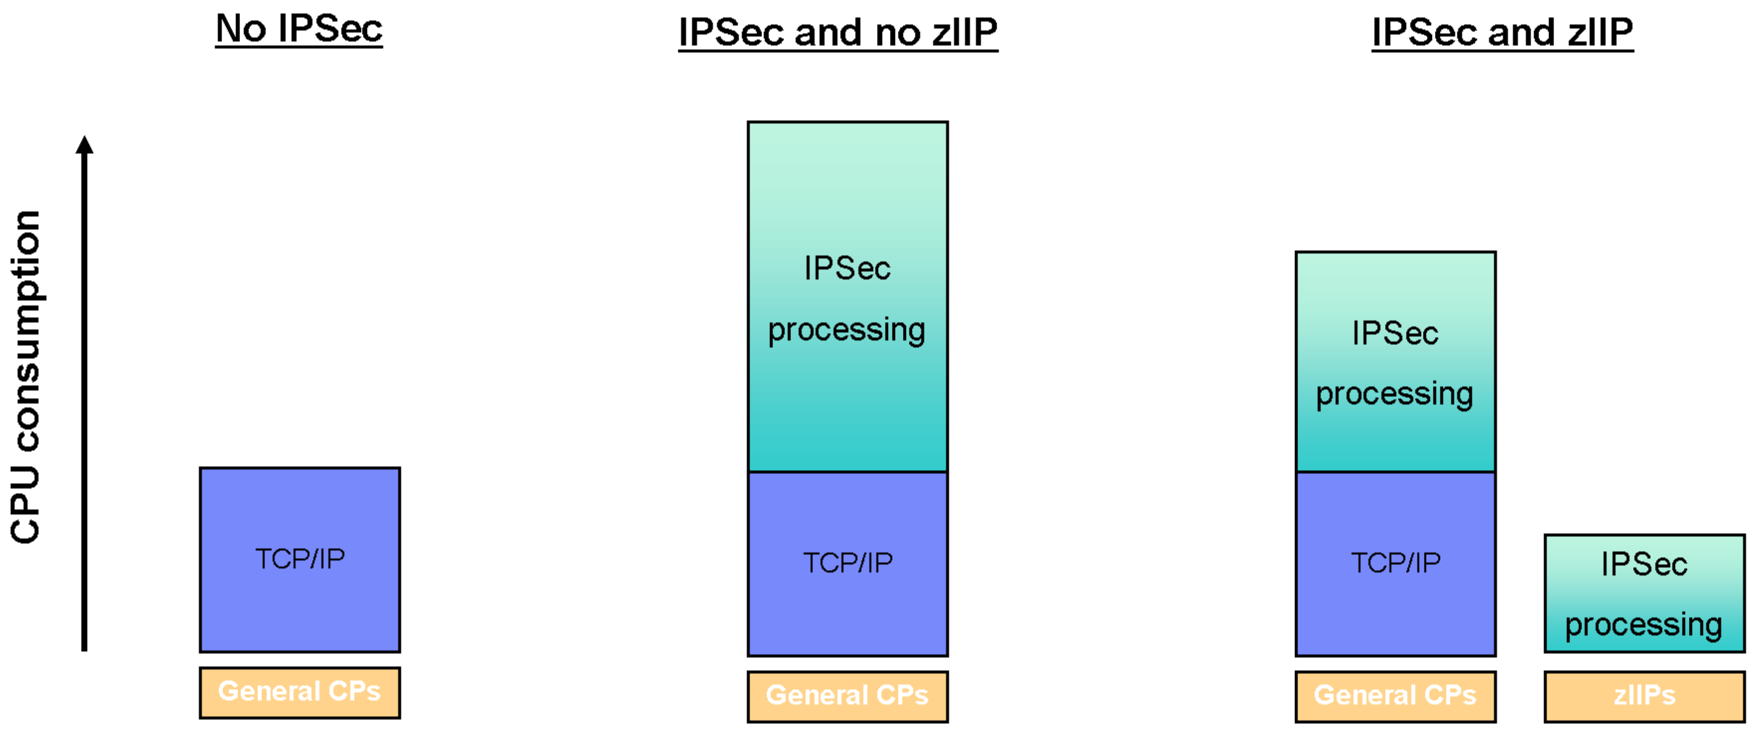

© 2008 IBM Corporation

Potential benefit of zIIP-Assisted IPSec – Inbound Bulk

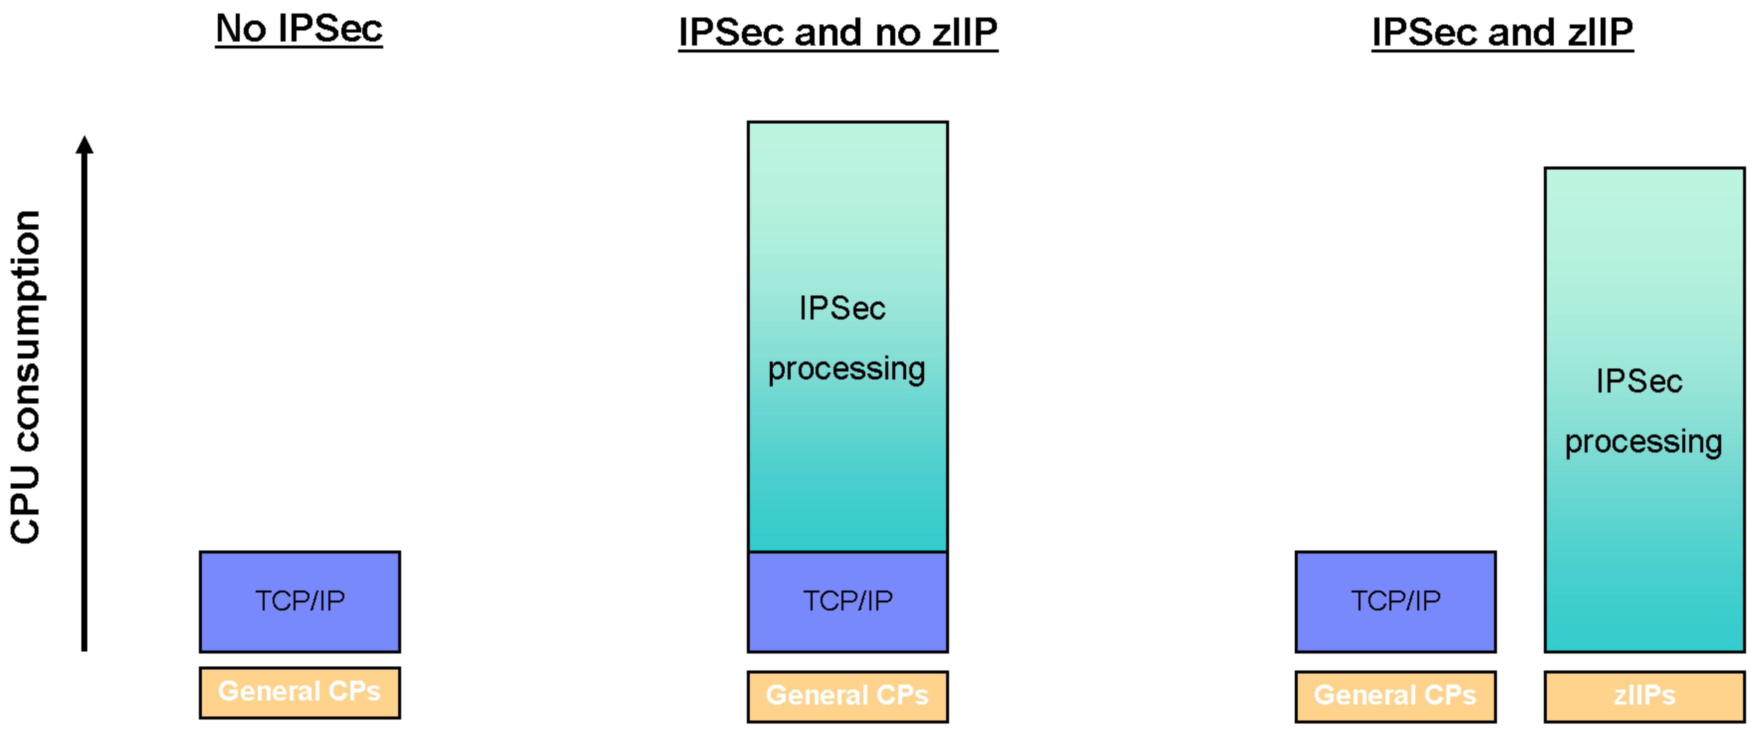

© 2008 IBM Corporation

Potential benefit of zIIP-Assisted IPSec – Outbound Bulk

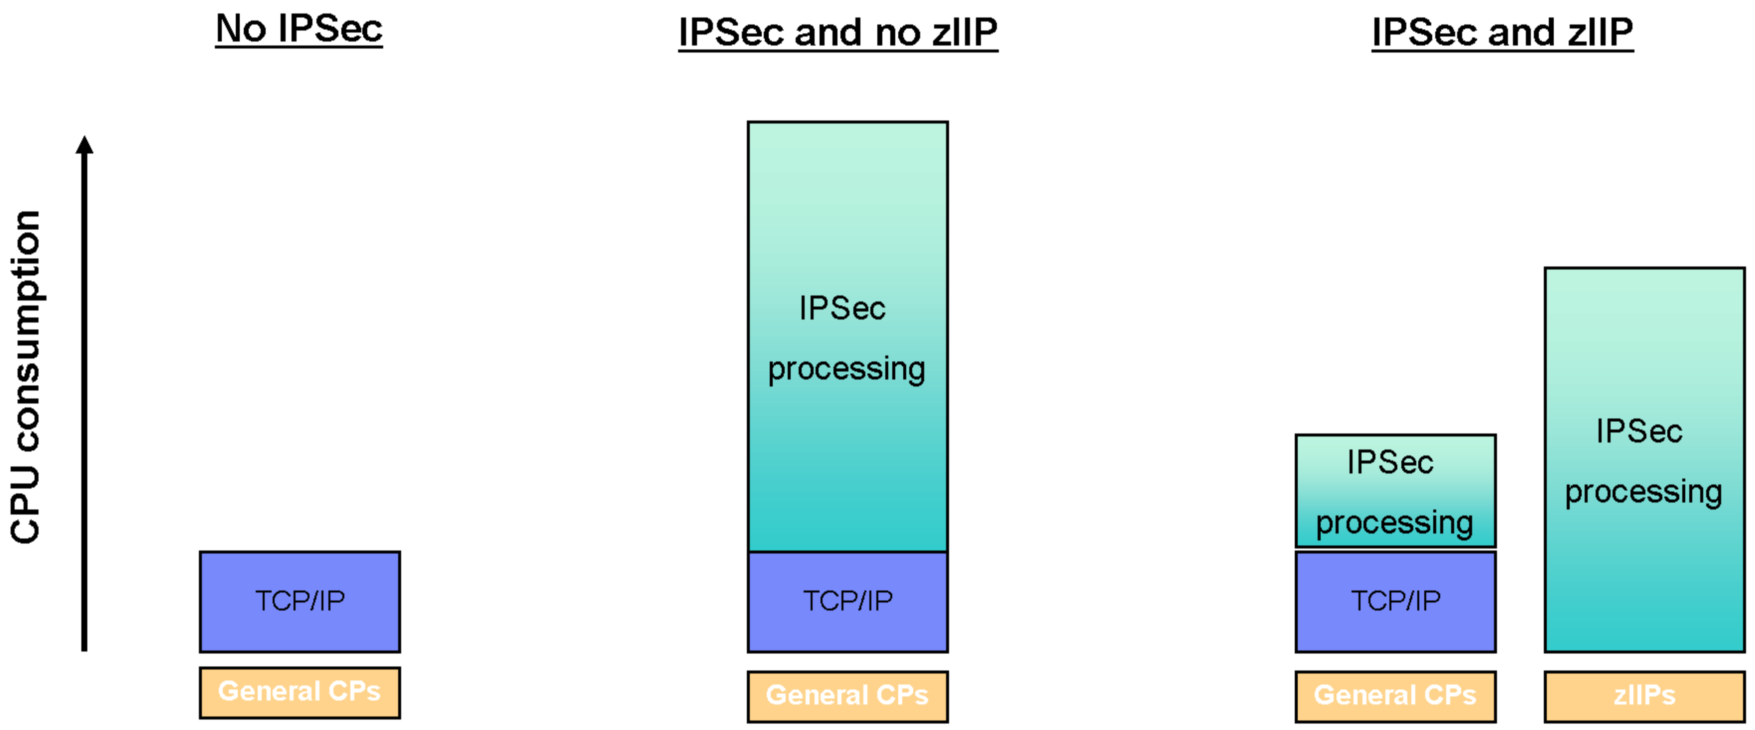

© 2008 IBM Corporation

Projecting zIIP Requirements

- How much of my existing (or future) workload is eligible to move to zIIPs?
- How many zIIPs would I need to handle my existing (or future) IPSec workload?
- Once I have zIIPs, how much CPU Busy relief can I expect on my standard CPs?

- Two general methods for projecting zIIP effectiveness:
  - If you’re already running IPSec, use PROJECTCPU function in z/OS Workload Manager
  - If you’re not yet running IPSec, some traffic modeling may be necessary – IBM’s Washington System Center will guide you through this

© 2008 IBM Corporation

Using PROJECTCPU

- Code PROJECTCPU=YES in IEAOPTxx PARMLIB member
- Code GLOBALCONFIG ZIIP IPSECURITY in TCPIP profile
- Run your IPSec workload
  - Collect RMF Workload Activity Report during representative interval

© 2008 IBM Corporation

Example: RMF report for bulk data and no zIIPs

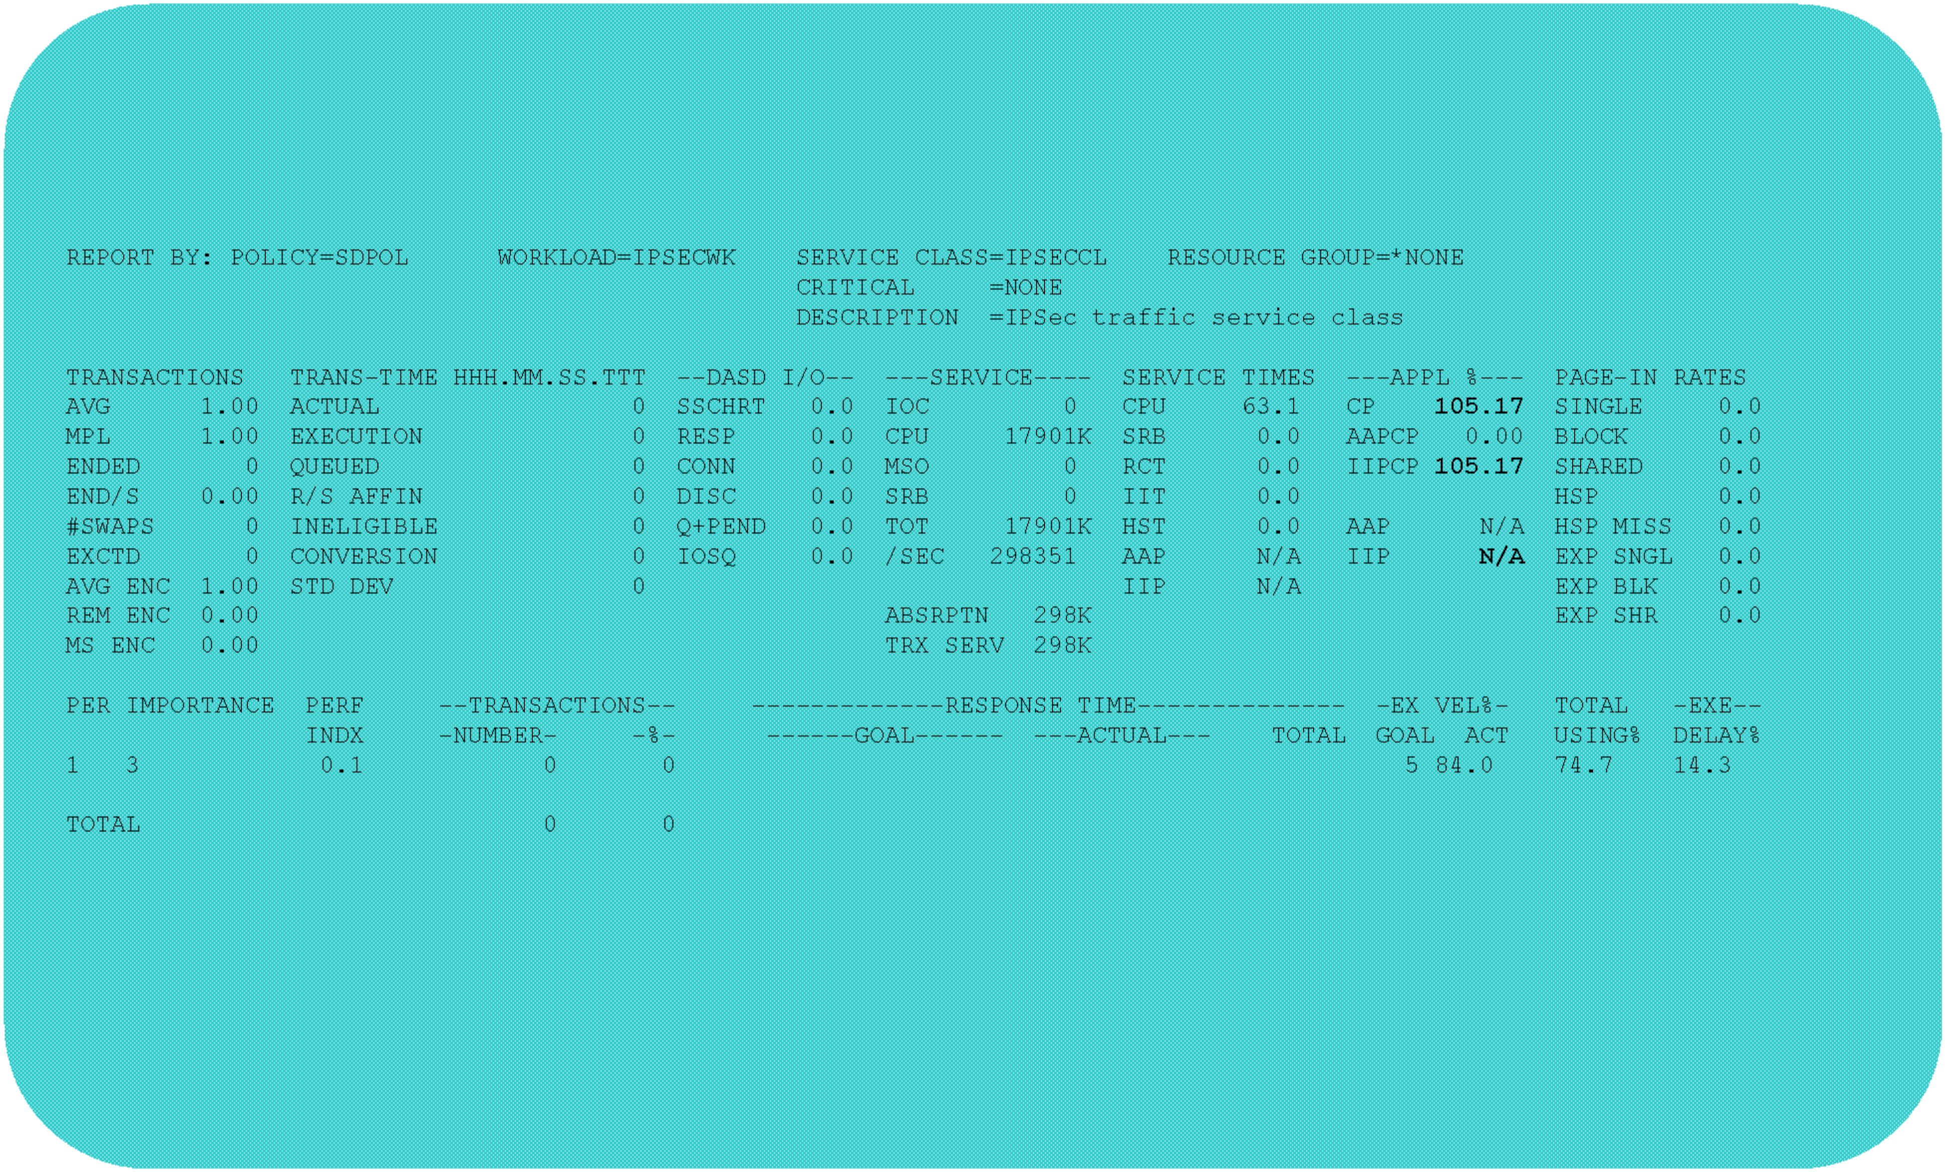

```
REPORT BY: POLICY=SDPOL      WORKLOAD=IPSECWK    SERVICE CLASS=IPSECCL    RESOURCE GROUP=*NONE
                                                 CRITICAL     =NONE
                                                 DESCRIPTION  =IPSec traffic service class

TRANSACTIONS   TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE----  SERVICE TIMES  ---APPL %---  PAGE-IN RATES
AVG      1.00  ACTUAL                 0  SSCHRT   0.0  IOC          0  CPU      63.1  CP    105.17  SINGLE     0.0
MPL      1.00  EXECUTION              0  RESP     0.0  CPU     17901K  SRB       0.0  AAPCP   0.00  BLOCK      0.0
ENDED       0  QUEUED                 0  CONN     0.0  MSO          0  RCT       0.0  IIPCP 105.17  SHARED     0.0
END/S    0.00  R/S AFFIN              0  DISC     0.0  SRB          0  IIT       0.0                HSP        0.0
#SWAPS      0  INELIGIBLE             0  Q+PEND   0.0  TOT     17901K  HST       0.0  AAP      N/A  HSP MISS   0.0
EXCTD       0  CONVERSION             0  IOSQ     0.0  /SEC   298351   AAP       N/A  IIP      N/A  EXP SNGL   0.0
AVG ENC  1.00  STD DEV                0                                IIP       N/A                EXP BLK    0.0
REM ENC  0.00                                          ABSRPTN   298K                               EXP SHR    0.0
MS ENC   0.00                                          TRX SERV  298K

PER IMPORTANCE  PERF     --TRANSACTIONS--     -------------RESPONSE TIME--------------  -EX VEL%-  TOTAL   -EXE--
                INDX     -NUMBER-     -%-      ------GOAL------  ---ACTUAL---    TOTAL  GOAL  ACT  USING%  DELAY%
1   3            0.1            0       0                                                 5 84.0   74.7    14.3

TOTAL                           0       0
```

© 2008 IBM Corporation

Example: RMF report for bulk data and no zIIPs …

- We're interested in the Workload Activity Report for the IPSECCL service class, since IPSec traffic that can be processed on available zIIP processors will run in this WLM Service class

- IIP N/A
  - Since zIIP is not configured

- IIPCP 105.17
  - Percentage of CPU time used by zIIP-eligible work running on general CPs
    - This workload would saturate a single zIIP, with at least 5% spilling over to CPs

- CP 105.17
  - The two CPs are each averaging 105.17/2 = ~52.59% busy handling this IP workload.
  - On CPU activity report, each of the two standard CPs are averaging 56.88% busy (normalizing to the capacity of a single standard CP this becomes 113.76%)
  - Therefore the percentage of CPU time that was not zIIP eligible in this benchmark is 113.76 – 105.17 = 8.59%.
    - If two zIIPs were added to this configuration, 113.76/105.17 = approximately 92% of the CPU consumption for this IP workload would move to zIIP.

Example: RMF report for bulk data and one zIIP

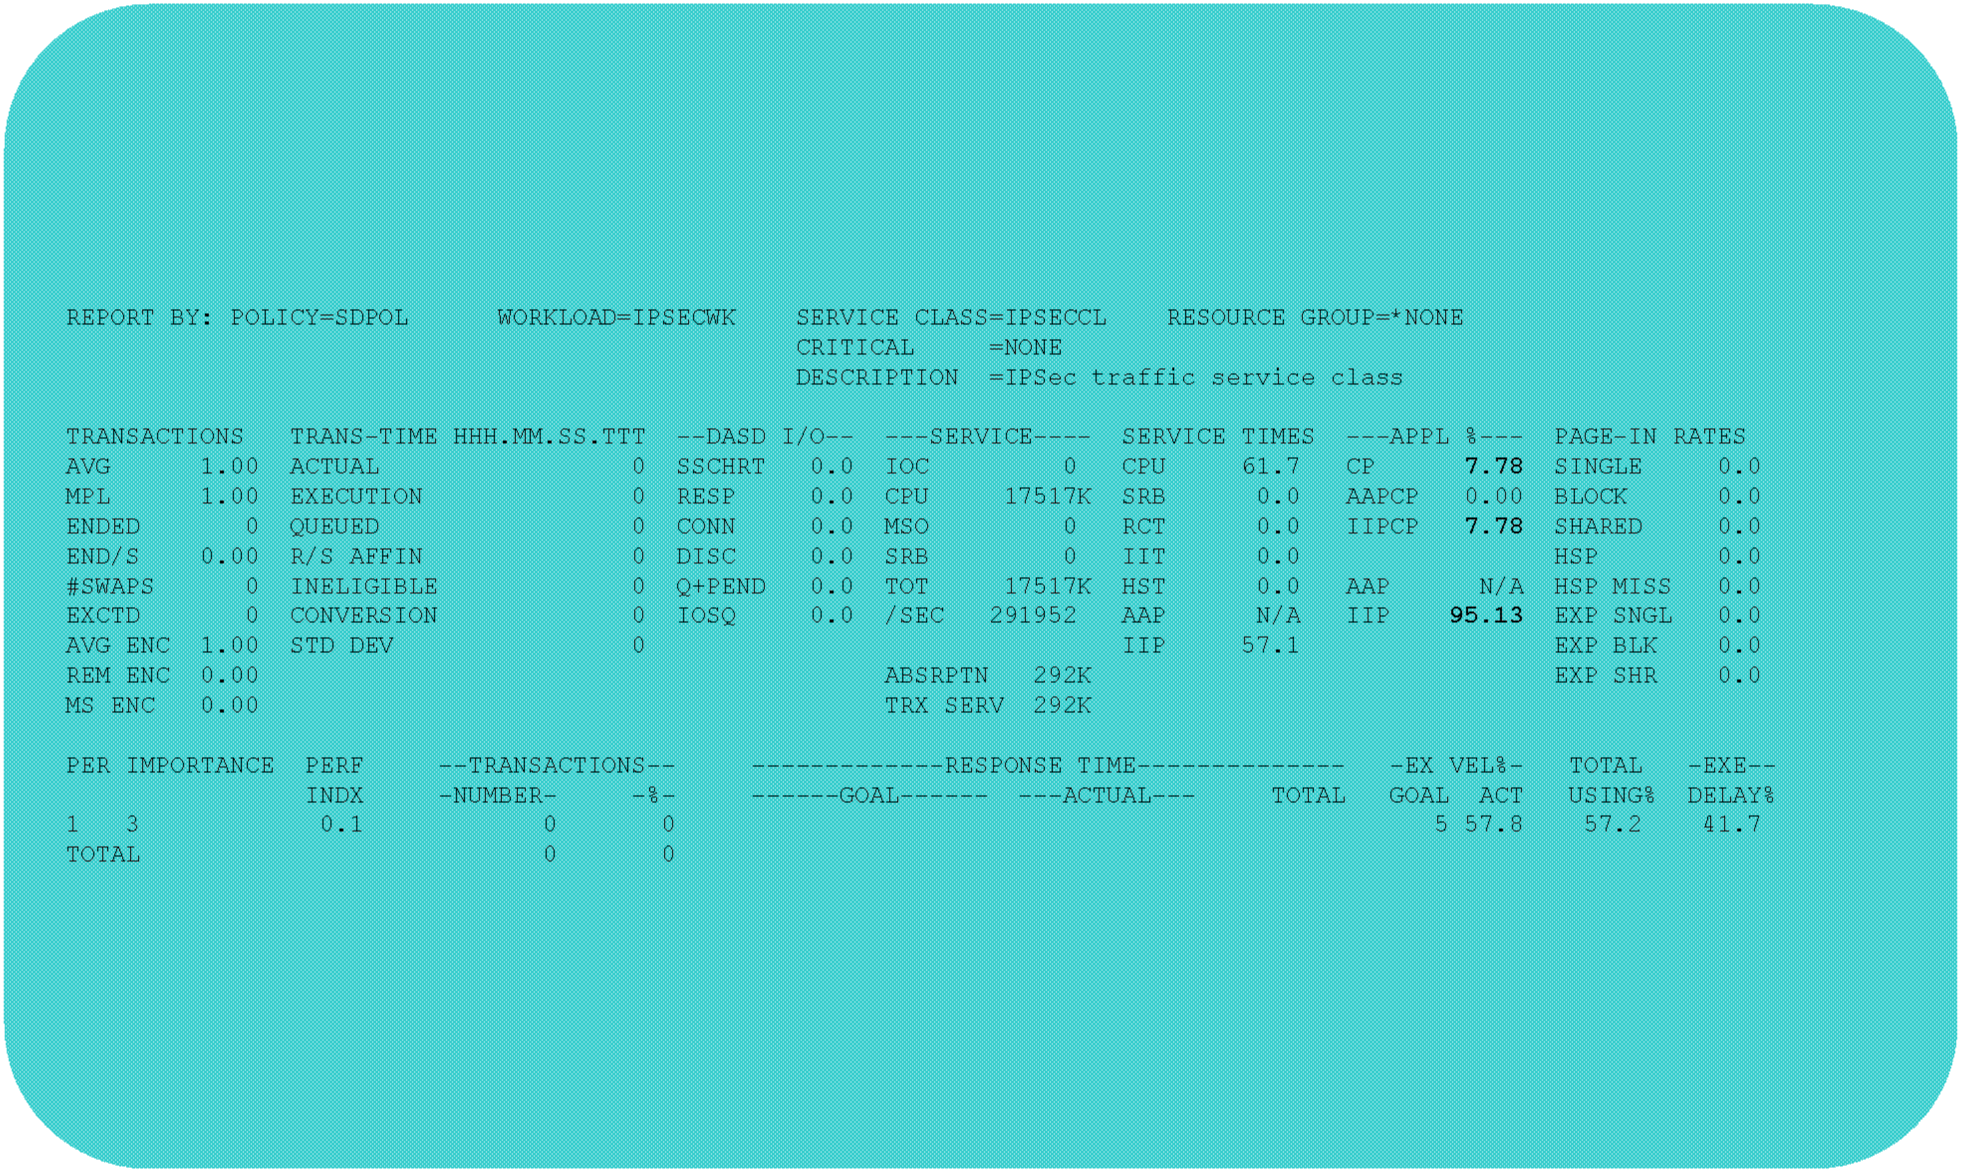

```
REPORT BY: POLICY=SDPOL      WORKLOAD=IPSECWK    SERVICE CLASS=IPSECCL    RESOURCE GROUP=*NONE
                                                 CRITICAL     =NONE
                                                 DESCRIPTION  =IPSec traffic service class

TRANSACTIONS   TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE----  SERVICE TIMES  ---APPL %---  PAGE-IN RATES
AVG      1.00  ACTUAL                 0  SSCHRT   0.0  IOC         0   CPU     61.7   CP      7.78  SINGLE     0.0
MPL      1.00  EXECUTION              0  RESP     0.0  CPU     17517K  SRB      0.0   AAPCP   0.00  BLOCK      0.0
ENDED       0  QUEUED                 0  CONN     0.0  MSO         0   RCT      0.0   IIPCP   7.78  SHARED     0.0
END/S    0.00  R/S AFFIN              0  DISC     0.0  SRB         0   IIT      0.0                 HSP        0.0
#SWAPS      0  INELIGIBLE             0  Q+PEND   0.0  TOT     17517K  HST      0.0   AAP      N/A  HSP MISS   0.0
EXCTD       0  CONVERSION             0  IOSQ     0.0  /SEC   291952   AAP      N/A   IIP    95.13  EXP SNGL   0.0
AVG ENC  1.00  STD DEV                0                                IIP     57.1                 EXP BLK    0.0
REM ENC  0.00                                          ABSRPTN   292K                               EXP SHR    0.0
MS ENC   0.00                                          TRX SERV  292K

PER IMPORTANCE  PERF     --TRANSACTIONS--     -------------RESPONSE TIME--------------   -EX VEL%-   TOTAL   -EXE--
                INDX     -NUMBER-     -%-     ------GOAL------  ---ACTUAL---     TOTAL   GOAL  ACT   USING%  DELAY%
1   3            0.1            0       0                                                   5 57.8    57.2    41.7
TOTAL                           0       0
```

Example: RMF report for bulk data and one zIIP…

- Same workload was run as previous chart, except now running with one zIIP configured

- IIP 95.13
  - The zIIP is 95.13% busy handling this IP workload
    - The remaining 4.87% of single-zIIP capacity is uncaptured time (z/OS base functions such as interrupt handling, dispatching, etc).

- IIPCP 7.78
  - Percentage of CPU time used by zIIP-eligible work running on general CPs
    - Previous analysis indicated about 5% spillover, the extra 2% is attributable to uncaptured time

- CP 7.78
  - The 2 CPs are each averaging 7.78/2 = 3.89% busy handling this workload
    - Previous analysis indicated the 2 CPs averaging 52.59% busy
    - With single zIIP configured, IPSECCL-related CP utilization has dropped by over 48 percentage points on each of the general CPs
    - The IPSECCL-related work remaining on the standard CPs here is work that "spilled over" from the single zIIP

© 2008 IBM Corporation

Performance Measurements

- Disclaimer
  - The performance data discussed in this presentation was collected using a dedicated system environment, so the results obtained in other configurations or operating system environments may vary
- The benchmarks used in this presentation were obtained using the Application Workload Modeler (AWM) for z/OS
  - For more information, visit the Application Workload Modeler website at http://www.ibm.com/software/network/awm/index.html
- All CPU consumption measurements are for networking CPU
  - Refers to CPU used in the TCPIP stack, Unix System Services, and MVS IOS, Scheduling, and Dispatcher cycles involved in networking flows
    - Typically contributes 8% of total CPU for interactive workloads
    - Typically contributes 30% of total CPU for bulk workloads
- All measurements collected on z/10 model 2097-752 LPARs with 2 dedicated general CPs and 0 - 1 dedicated zIIPs per LPAR running Communications Server V1R10
  - IPSec configuration utilized Triple-DES encryption with SHA authentication
  - V1R10 provides optimizations for IPSEC versus V1R9
  - zIIP-assisted IPSEC provides TCP flow-control changes for bulk data transfers

© 2008 IBM Corporation

Interactive workload Measurements

- 10 concurrent interactive sessions sending/receiving 100 bytes

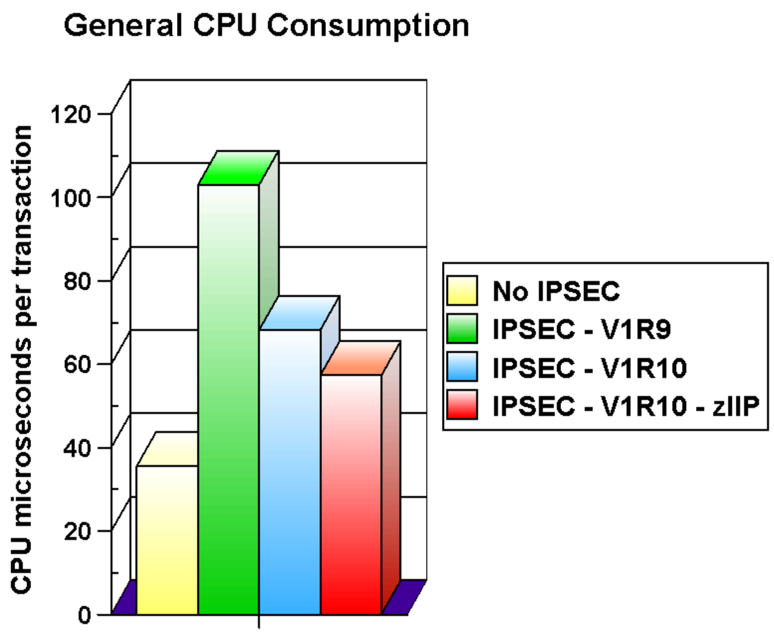

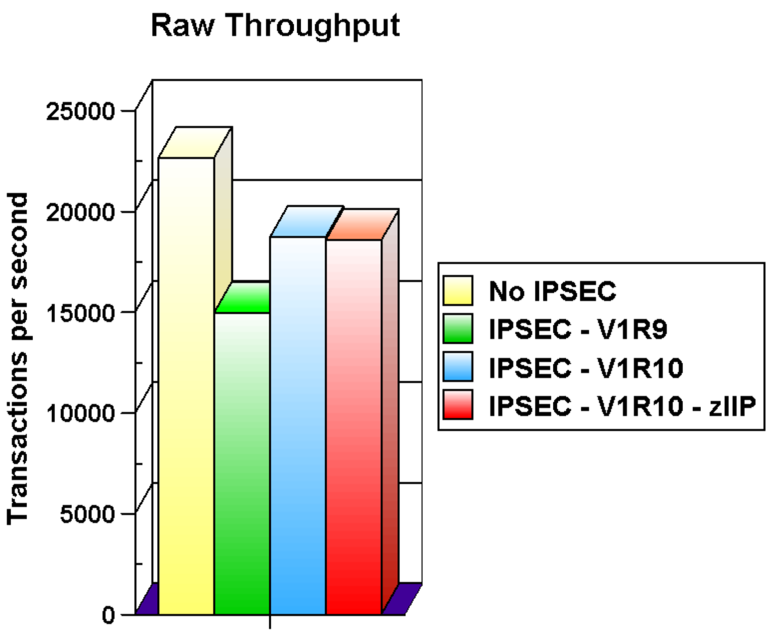

- With zIIPs and V1R10, networking CPU for IPSec drops from a 189% increase to a 61% increase compared to non-secured

- Regardless of zIIPs, overhead of IPSec processing adds latency which results in 17% lower transaction rate compared to non-secure

“Normalized” Interactive workload Measurements

- IPSEC impact to CPU consumption
  (based on 8% networking costs)

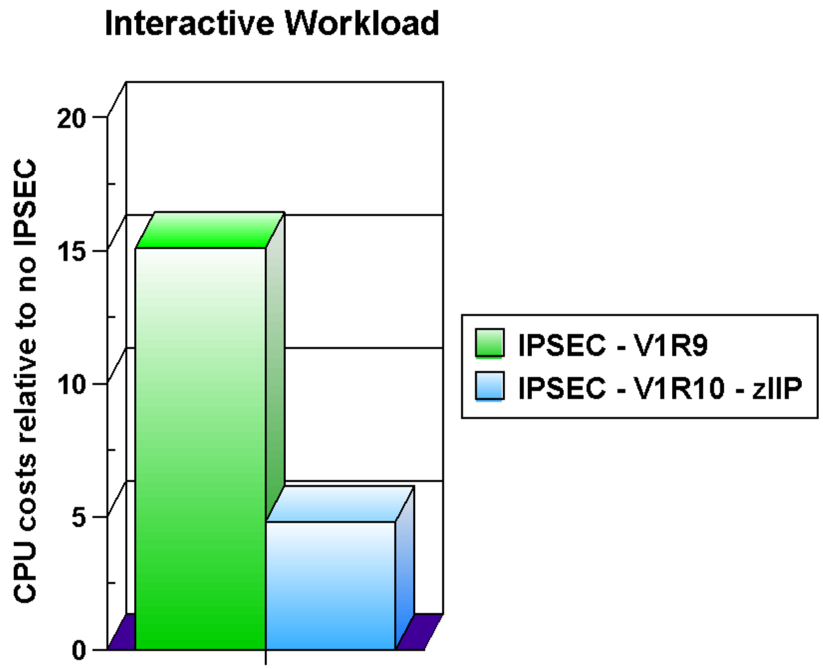

- For interactive traffic, enabling IPSEC results in a 5% increase in CPU per transaction when using zIIPs on V1R10

© 2008 IBM Corporation

Inbound Bulk workload Measurements

- 5 concurrent streaming sessions sending 1 byte and receiving 20 Mbytes

General CPU Consumption

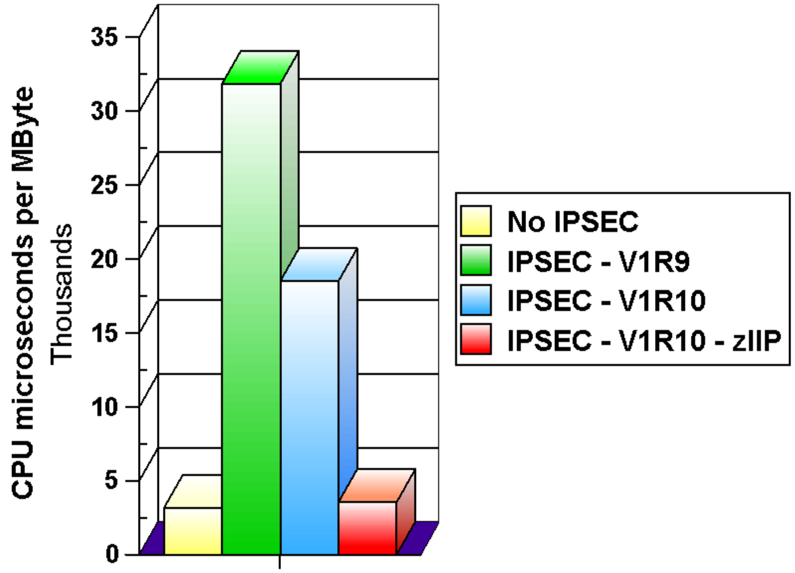

Raw Throughput

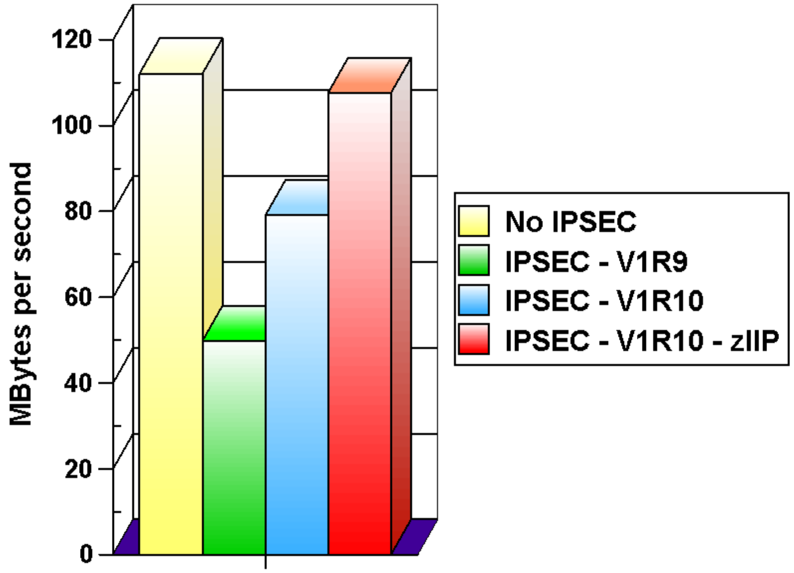

- With zIIPs and V1R10, networking CPU for IPSec drops from a 912% increase to just 13% more compared to non-secured
- Algorithm changes with zIIP and V1R10enabled on client reduces throughput gap from 56% to 4% compared to non-secure

© 2008 IBM Corporation

Outbound Bulk workload Measurements

- 5 concurrent streaming sessions sending 20 Mbytes and receiving 1 byte

General CPU Consumption

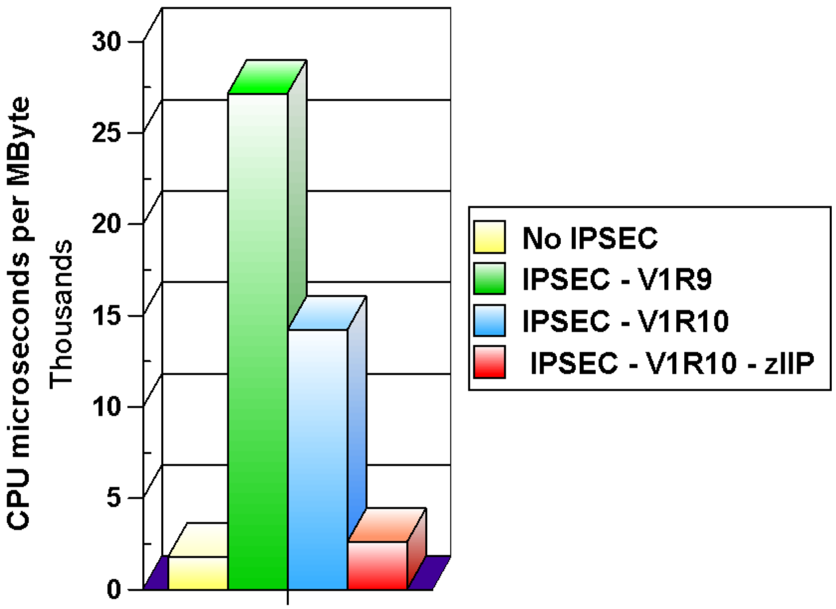

Raw Throughput

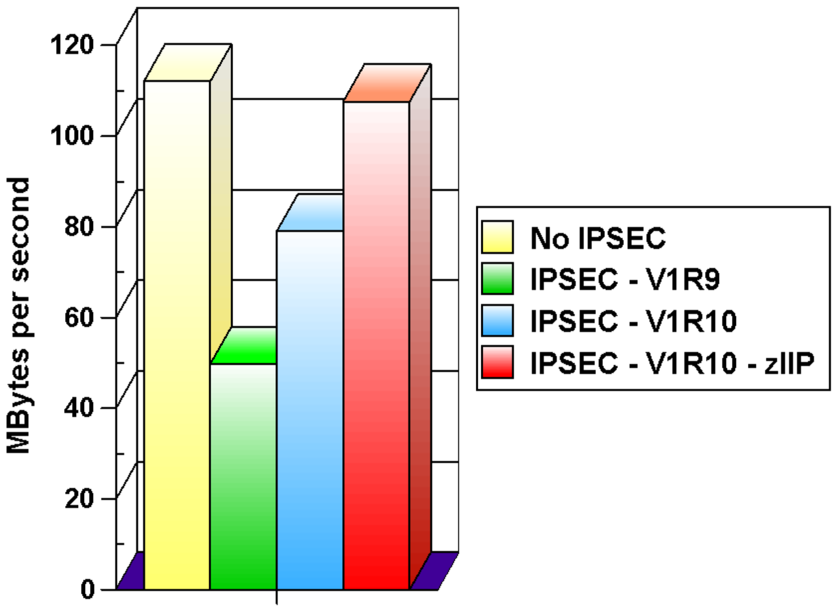

- With zIIPs and V1R10, networking CPU for IPSec drops from a 1405% increase to 45% more compared to non-secured

- Algorithm changes with zIIP enabled and V1R10 optimizations reduces throughput gap from 56% to 4% compared to non-secure

© 2008 IBM Corporation

“Normalized” Bulk workload Measurements

- IPSEC impact to CPU consumption (based on 30% networking costs)

Inbound Bulk Workload

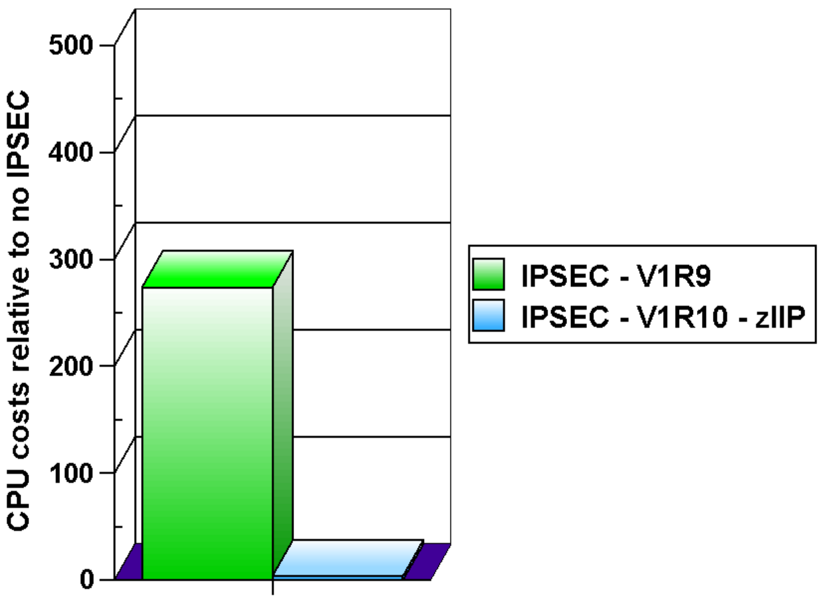

Outbound Bulk Workload

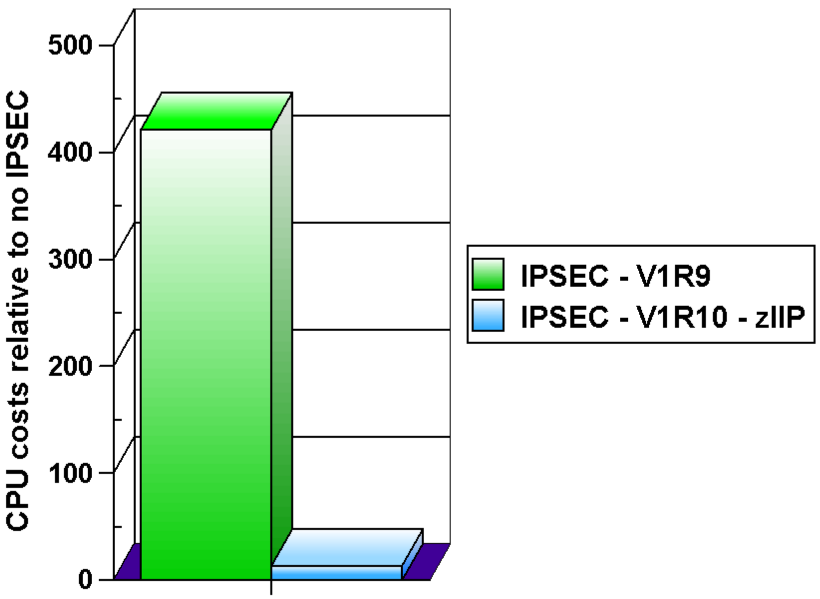

- For inbound bulk transfers, enabling IPSEC results in a 4% increase in CPU per Mbyte when using zIIPs on V1R10
- For outbound bulk transfers, enabling IPSEC results in a 14% increase in CPU per Mbyte when using zIIPs on V1R10

Final Thoughts

- Further reading
  - WP100988 – Capacity Planning for zIIP-Assisted IPSec
    - http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100988
- Questions?

© 2008 IBM Corporation

For More Information….

URL	Content
http://www.ibm.com/software/ipv6	IBM's IPv6 web page
http://www.ibm.com/systems/z/	IBM System z
http://www.ibm.com/systems/z/hardware/networking.index.html	IBM System z Networking
http://www.ibm.com/servers/eserver/zseries/networking/technology.html	IBM Enterprise Servers: Networking Technologies
http://www.ibm.com/software/network	Networking & Communications Software
http://www.ibm.com/software/network/commserver/zos	IBM z/OS Communications Server
http://www.ibm.com/software/network/commserver/library	CS White Papers, Product Doc, etc.
http://www.redbooks.ibm.com	IBM Redbooks
http://www.ibm.com/software/network/commserver/support	Communications Server Technical Support
http://publib.boulder.ibm.com/infocenter/ieduasst/stgv1r0/index.jsp	IBM Education Assistant
http://www.ibm.com/support/techdocs/	Advanced Technical Support (Flashes, Presentations, White Papers, etc.)

© 2008 IBM Corporation