This presentation discusses the zIIP Assisted IPSec function in the z/OS V1R9 Communications Server.

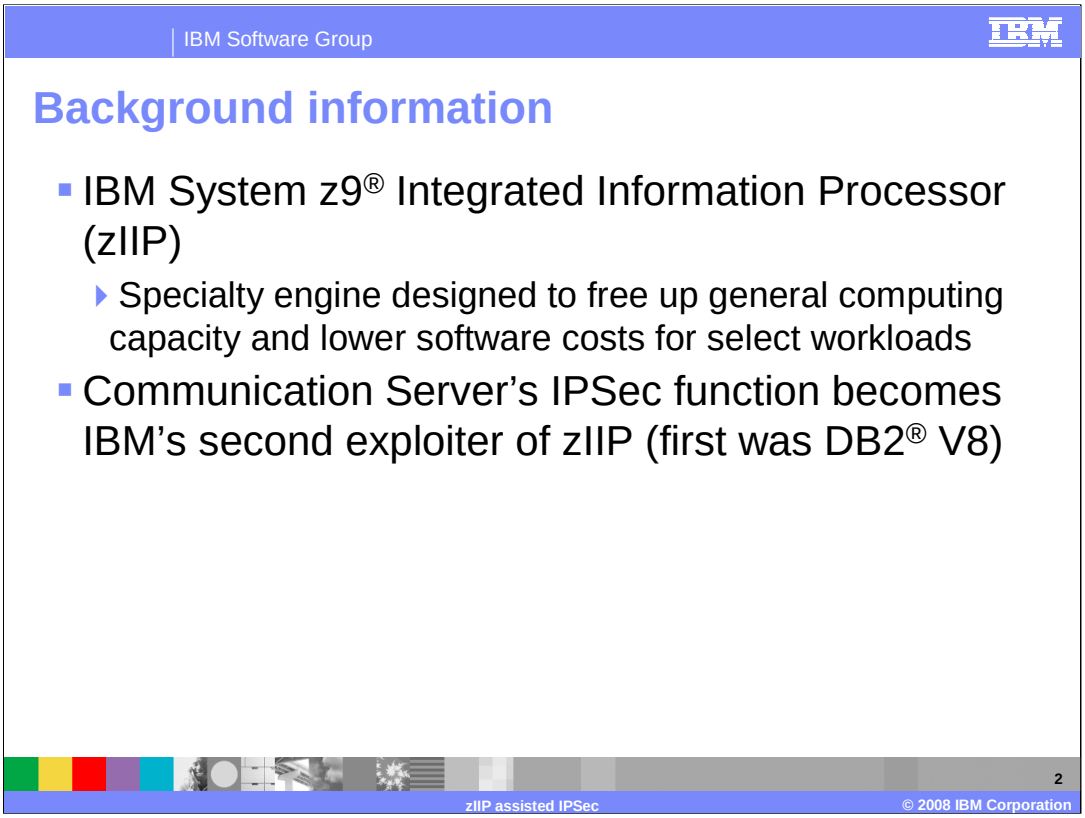

The z9 Integrated Information Processor (zIIP) was announced in 1Q2006. At that time, IBM DB2 V8 was the only zIIP exploiter.

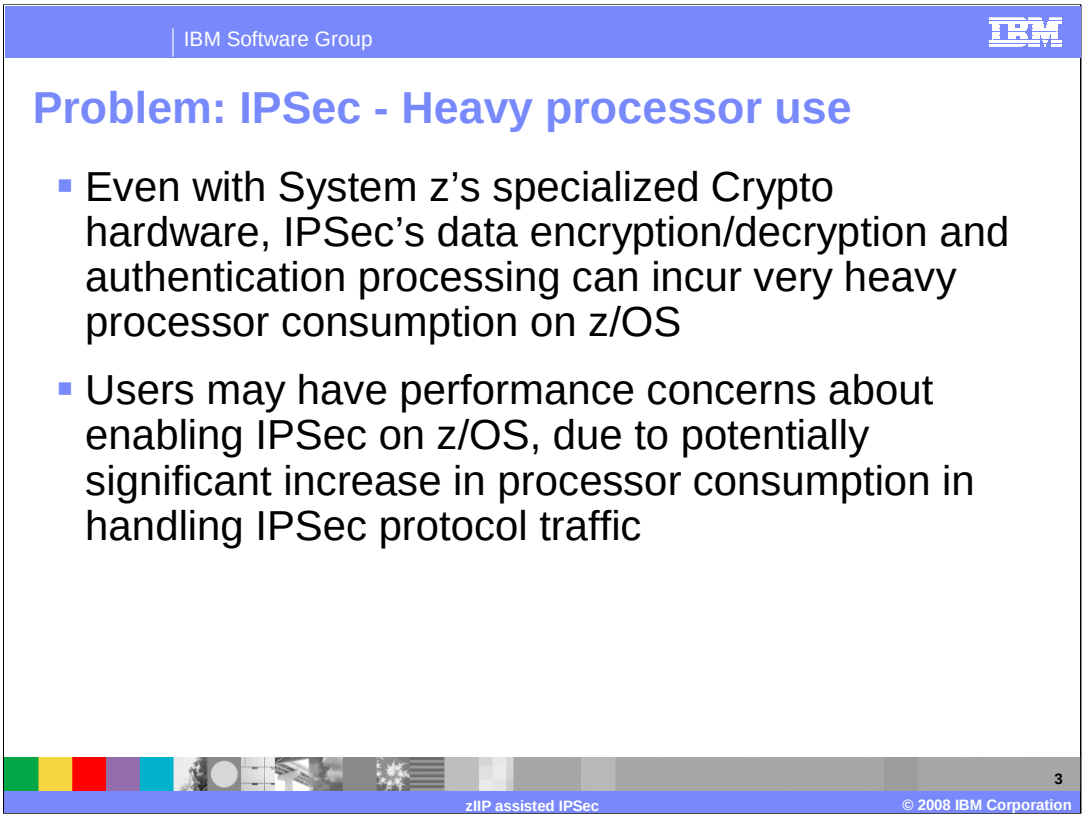

IPSec processor consumption for certain types of network traffic can be very intensive. For example, securing bulk data workloads (like FTP or TSM) with IPSec can be especially processor intensive, since IPSec processing cost is relative to the amount of data being moved. The extra cycles consumed by IPSec can be problematic If you are already running your z/OS LPARs at high utilization.

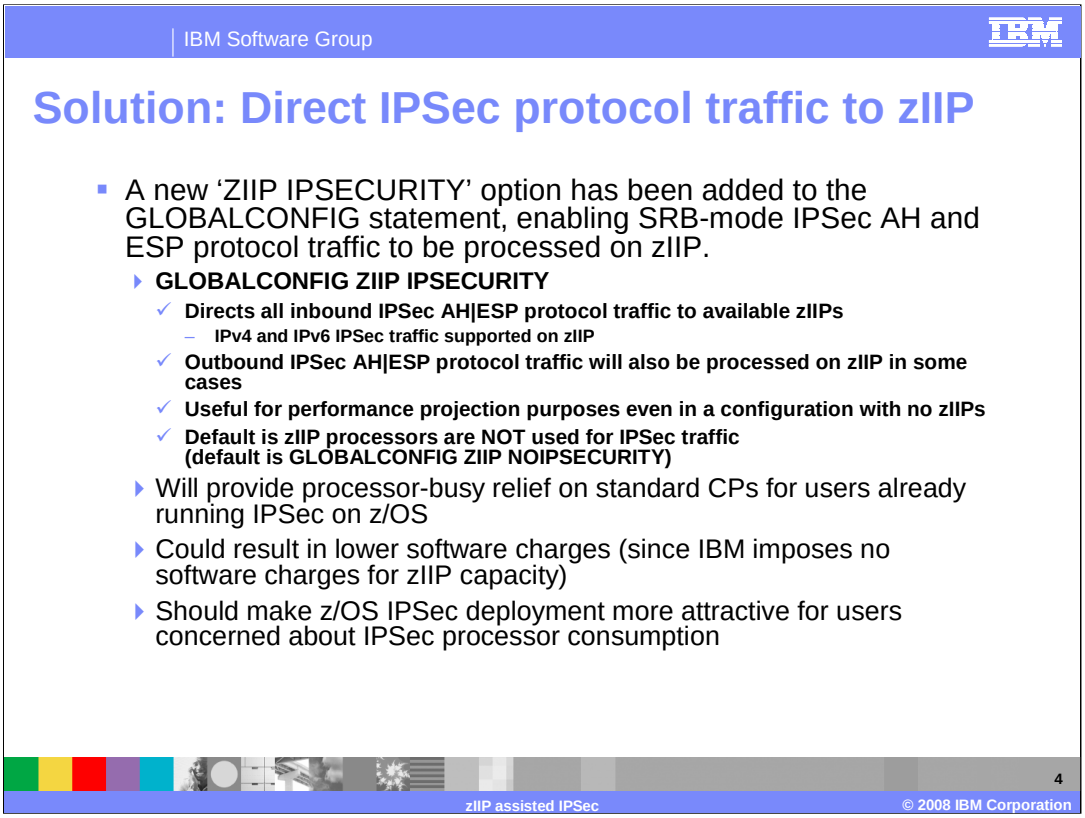

The zIIP IPSECURITY feature helps position IBM System z9 as a cost-effective server in environments requiring end-to-end security for IP network traffic. By directing IPSec's Authentication Header (AH) and Encapsulating Security Payload (ESP) protocol traffic to zIIP, your standard CPs will run less busy, and this can result in reduced software charges (since IBM imposes no software charges for zIIP capacity). If you have decided against IPSec deployment on z/OS (due to processor consumption issues), you may find the zIIP IPSECURITY feature now makes such deployment feasible.

Configuring GLOBALCONFIG ZIIP IPSECURITY causes *inbound* ESP and AH Protocol traffic to be processed in Enclave SRBs, and targeted to available zIIPs. *Outbound* ESP and AH protocol traffic may also be processed on available zIIPs when either the application invoking the send() function is already running on a zIIP, or when the data to be transmitted is in response to normal TCP flow control (for example, data transmitted in response to a received TCP acknowledgement or window update).

If you have no zIIPs, you can also use GLOBALCONFIG ZIIP IPSECURITY in conjunction with the MVS PROJECTCPU function, to obtain RMF projection data on the percentage of workload that is eligible to be run on zIIP.

The default setting is to leave IPSec processing on standard CPs, so if you do want to direct your IPSec processing to zIIP, you need to code GLOBALCONFIG ZIIP IPSECURITY.

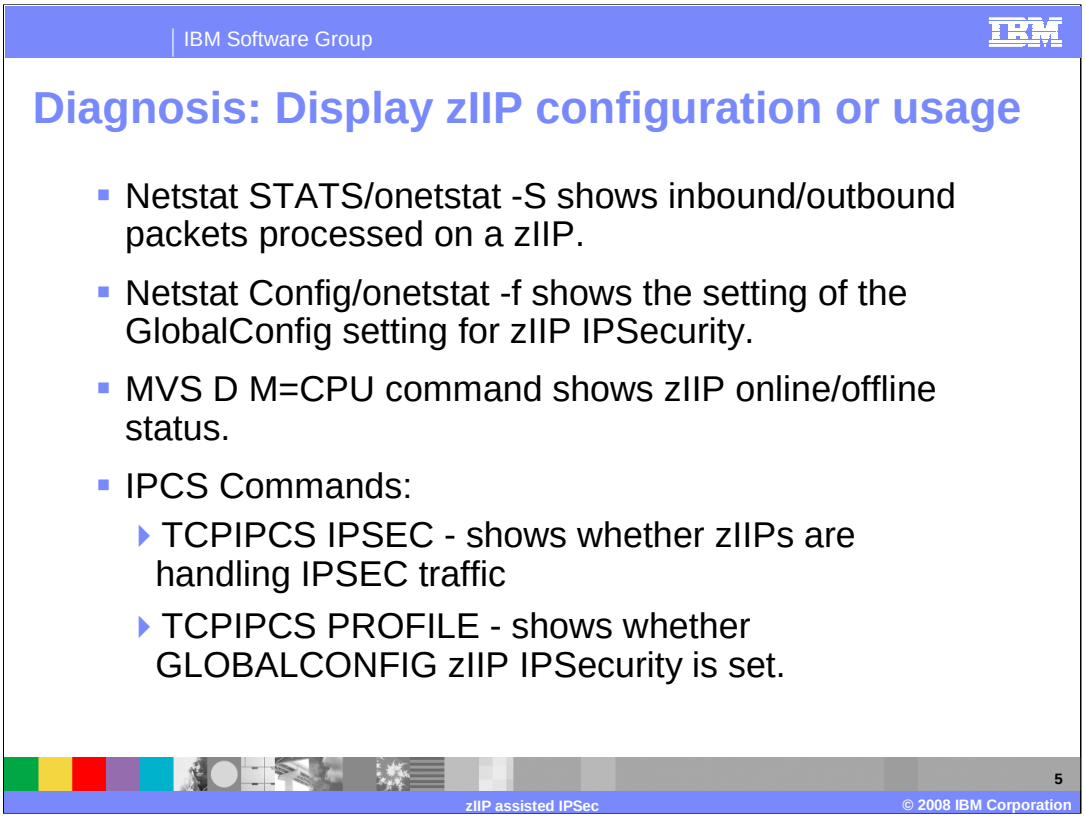

Various Netstat options are available for viewing zIIP IPSec behavior and configuration. If in doubt about zIIP online|offline status, use the

MVS D M=CPU command. If viewing a dump and you aC:\iea\content\stgv1r0\commserv\1.9z\security\zIIP_Assist_IPSec.pptre interested in zIIP IPSec configuration, you can use the TCPIPCS IPSEC or TCPIPCS PROFILE commands.

D M=CPU command example

D M=CPU shows 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
```


Once you have a zIIP configured to your LPAR, you can use the MVS D M=CPU command to display zIIP Status. In this example, there are two standard CPs and one zIIP online. The zIIP is identified by the “I” character next to the Online|Offline status indicator.

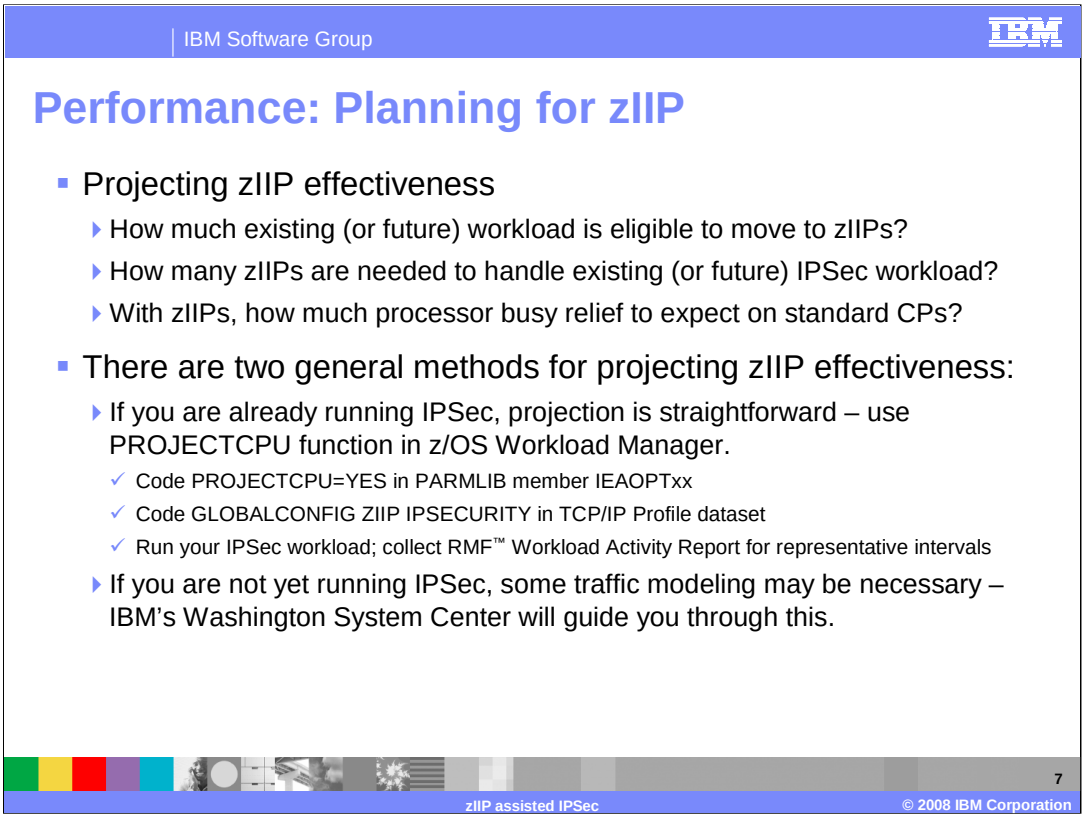

If you are running IPSec, zIIP may significantly reduce the processor utilization of your standard CPs. In planning for zIIP, you need to determine (a) how much of your workload is eligible to move to zIIP, (b) how many zIIPs are required to fully handle that load, then (c) once you do have zIIPs, how much processor busy relief you can expect on your standard CPs.

If you are already running your representative IPSec workload, performing zIIP Projection analysis is pretty simple. Function exists within z/OS that will allow users already running IPSec (but not currently using zIIPs) to accurately project the amount of their existing workload that is eligible to move the zIIPs. This function builds upon the PROJECTCPU service present in z/OS. PROJECTCPU gives a very precise accounting of workload that is zIIP-eligible. Using PROJECTCPU for zIIP capacity planning purposes is therefore very accurate and simple, since no extra analysis of network traffic is required.

If you are not yet running IPSec, some complex traffic modeling may be necessary to derive accurate estimates of zIIP effectiveness in your future IPSec configuration. System z sales personnel will engage the Washington System Center to perform this modeling, when necessary.

PROJECTCPU example: RMF workload activity report – Bulk data with no zIIP configured

```
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
```


This configuration has no zIIPs online, and it shows using the PROJECTCPU function to gather projection data.

Note 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. Interpretation of this workload activity report is as follows:

IIP N/A - since zIIP is not configured.

IIPCP 105.17 - this is the percentage of processor time used by zIIP-eligible work (in the IPSECCL Service Class) running on standard CPs. This statistic is normalized to the capacity of a single standard CP, so the interpretation is: *This workload would fully saturate a single zIIP, with at least 5% spilling over to CPs.* Two zIIPs would be required to handle all of the zIIP eligible workload from the IPSECCL Service Class.

CP 105.17 - The two CPs are each averaging 105.17/2 = ~52.59% busy handling this IP workload. Not shown here, the processor activity report indicates that 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 processor 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 processor consumption for this IP workload would move to zIIP, off of the standard CPs. This workload is especially well-suited to zIIP, due to the high degree of enclave SRB execution within CommServer for inbound bulk-data workload.

PROJECTCPU example: RMF workload activity report – Bulk data with ONE zIIP configured

```
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
```


This configuration has one zIIP online, running the same workload as in the previous chart.

Interpretation of this workload activity report is as follows:

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 - In the earlier analysis (Workload Activity Report and processor activity report with NO zIIP), at least 5% of the zIIP-eligible workload spills over to standard CPs if a single zIIP is added to the configuration. In this case you are seeing 7.78% is spilling over. The portion greater than 5% is attributable to uncaptured time.

CP 7.78 - In the earlier analysis (Workload Activity Report with NO zIIP), the IPSECCL Service Class accounted for 52.5% busy on each of the standard CPs. With the zIIP now configured, IPSECCL-related CP utilization (averaged on each CP) has dropped to 7.78/2 = 3.89% busy; *a utilization reduction of over 48 percentage points on each of the standard CPs.* The IPSECCL service class work remaining on the standard CPs here is work that "spilled over" from the single zIIP.

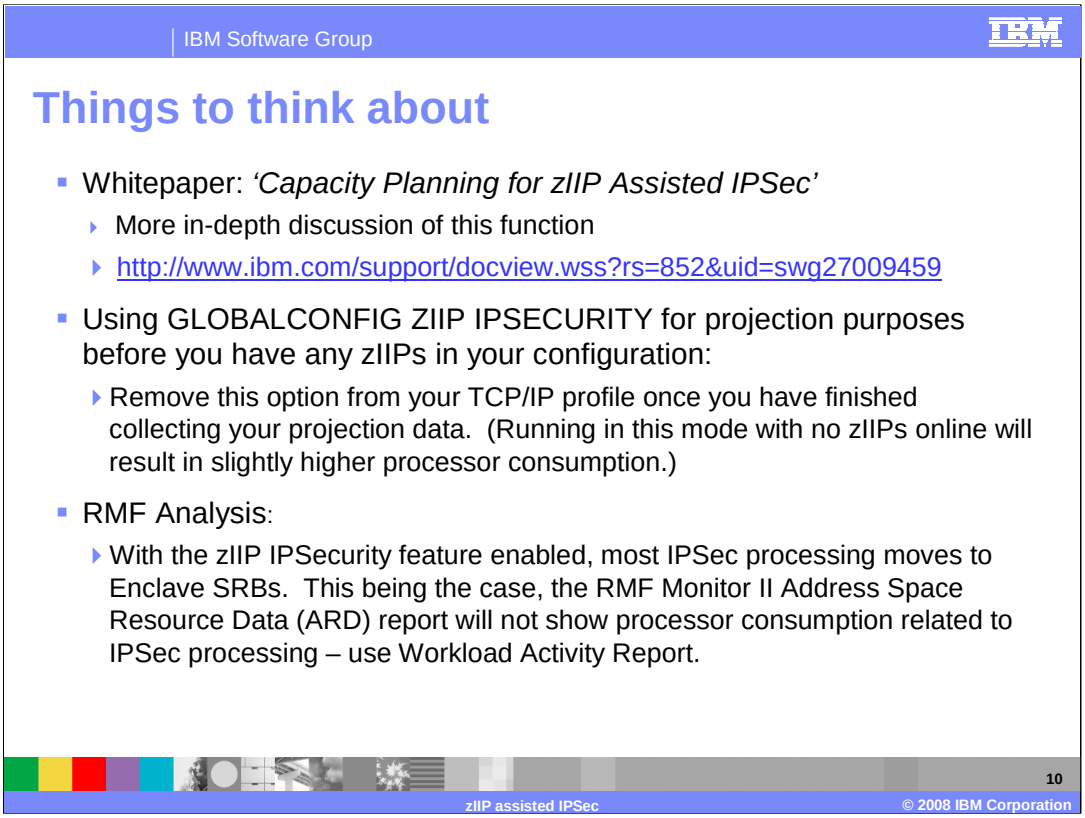

There is a whitepaper that covers zIIP IPSec projection modeling in depth, and presents some of the early zIIP IPSec performance data collected in IBM labs. It can be found on ibm.com.

Specify GLOBALCONFIG ZIIP IPSECURITY only if (a) you already have zIIPs online to your LPAR or (b) you are running performance runs to obtain RMF data to project zIIP effectiveness.

Things to think about (continued)

- Controlling “spillover” of work back to standard CPs - PARMLIB Member IEAOPTXX: IIPHONORPRIORITY statement
  - Controls whether zIIP eligible work is allowed to run on standard CPs
  - IIPHONORPRIORITY=YES is recommended and default value (zIIP eligible work is allowed to run on CPs if zIIP requests help)
  - IIPHONORPRIORITY=NO means z/OS will try to contain all zIIP-eligible workload on zIIPs; this can lead to throughput and response time degradation when zIIP is highly used

Specifying IIPHONORPRIORITY=YES allows zIIP eligible workload to run on standard CPs, if zIIP work is not completed in a reasonable time period. This is the default and recommended value.

Specifying IIPHONORPRIORITY=NO disallows any zIIP eligible work from running on CPs (unless no zIIPs are online, or zIIP work is holding system locks or other resources impeding non zIIP work). When the NO value is set and zIIPs are present in the configuration, zIIP eligible work is contained on the zIIPs. During periods of very high zIIP utilization, throughput and response time may suffer. It may be reasonable to tradeoff throughput/response time in some environments, where minimizing utilization of the standard CPs is paramount

Trademarks, copyrights, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DB2 IBM RMF System z9 z/OS

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.