Advanced Technical Skills (ATS) North America

CPU MF Counters Enablement Webinar

John Burg
Kathy Walsh

May 2, 2012

Announcing *CPU MF Enablement* Education

- **Two Part Series**
  - Part 1 – General Education – **Today's Agenda**
    - Why is it important to use
    - What features does it have
    - How do you enable the function
    - Q&A
    - How do you get to Part 2?
      - Need to implement CPU MF Counters and send WSC data - **by June 1, 2012**
      - WSC will provide data collection guidelines
      - When valid data is received an invitation to the 2nd webinar will be issued
    - Q&A
  - Part 2 – Specific Education – **June 14, 2012 – 2:00-3:30 PM ET**
    - Send us the data and we will send back a CPU MF report of your data
    - Session will cover the use and meaning of the different data in the report
    - Provide overall "profiles" of the data that was sent
      - No customer names will be visible
      - Provide things like: (95% of all data samples in Part 2 were AVERAGE Workload
    - Questions and Answers about the data

New Day Dawning in System z Capacity Planning

Processor Design
- **CPU**
- **Memory Hierarchy (Nest)**

Hipervisor (PR/SM)
- **Amount of virtualization**

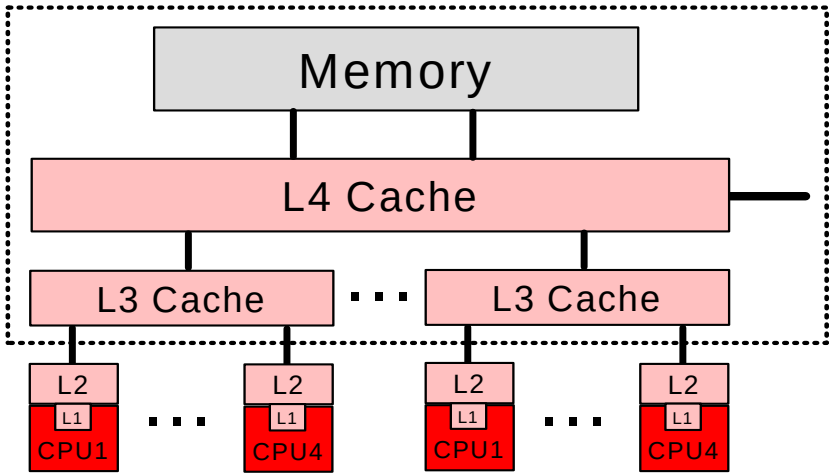

Nest

Operating System
- **Virtualization at address space level**

Workload Characteristics
- **Instructions**
- **Dispatch Profile**
- **I/O Rate**

Introduction to LSPR

- **A set of representative SCP/workload environments**
  - SCPs: z/OS, z/VM, and Linux on System z
  - Workload categories: Low ←Relative Nest Intensity→ High
  - Current LSPR workload categories: Low, Average, High
  - zPCR extends published categories
    - Low-Avg
    - Avg-High
  - A methodology focused on processor capacity
  - No significant external constraints
  - Equivalent (reasonably high, e.g. >= 90%) processor utilization
- **A metric to communicate the results**
  - ITR: Internal Throughput Rate
  - Transactions or Jobs per processor busy second
- **Information stored on the web**
  - **https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument**

LSPR Workload Categories

- Various combinations of workload primitives are measured on which the new workload categories are based
  - **Applications include CICS, DB2, IMS, OSAM, VSAM, WebSphere, COBOL, utilities**

- Low (relative nest intensity)
  - **Workload curve representing light use of the memory hierarchy**
  - **Similar to past high scaling workload primitives**

- Average (relative nest intensity)
  - **Workload curve expected to represent the majority of customer workloads**
  - **Similar to the past LoIO-mix curve**

- High (relative nest intensity)
  - **Workload curve representing heavy use of the memory hierarchy**
  - **Similar to the past DI-mix curve**

- zPCR extends published categories
  - **Low-Avg**
    - **50% Low and 50% Average**
  - **Avg-High**
    - **50% Average and 50% High**

zPCR Workload Characterization for z/OS

"Scope of Work" Definition Change

New z/OS Workload Categories Defined

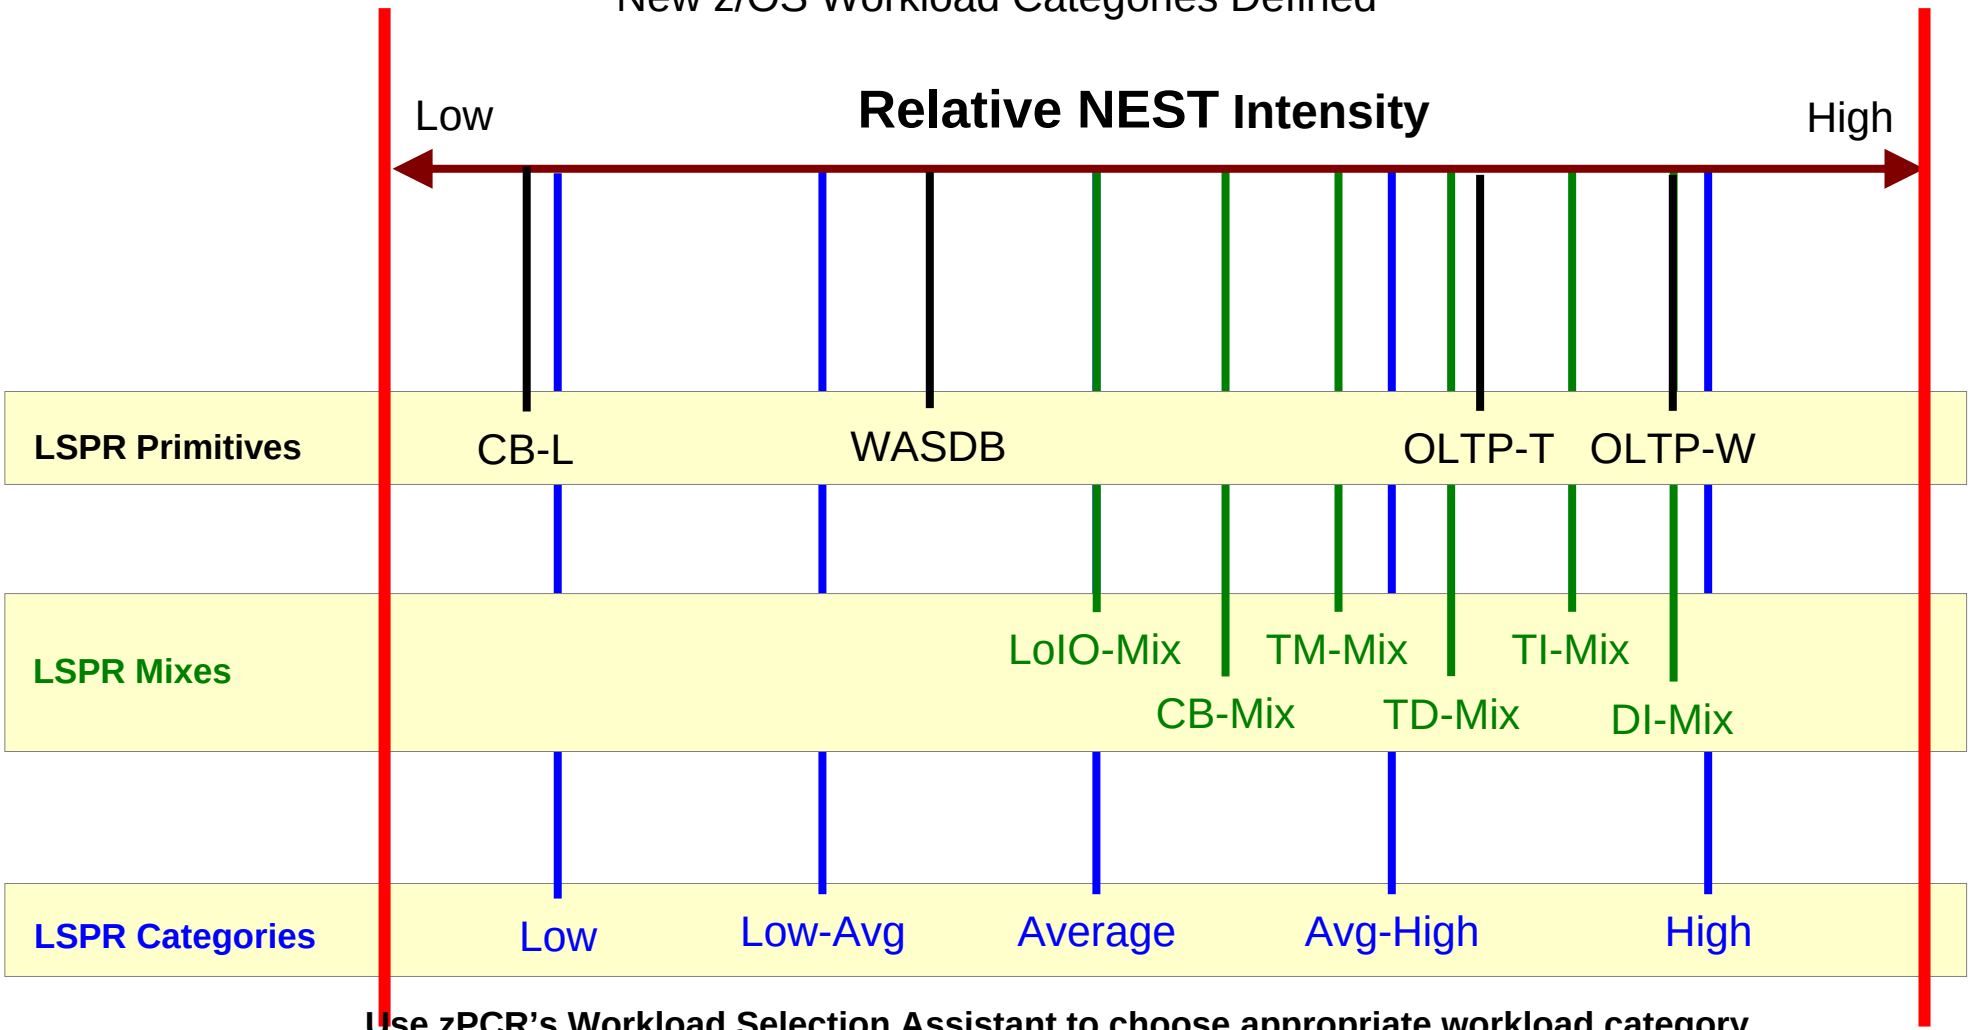

Use zPCR's Workload Selection Assistant to choose appropriate workload category

Automated with EDF input into zPCR

Note: Workload selection is automated in zCP3000

CPU Measurement Facility

- **Introduced in z10 and later processors**
- **New facility that provides hardware instrumentation data for production systems**
- **Two Major components**
  - Counters
    - Cache and memory hierarchy information
    - SCPs supported include z/OS and z/VM
  - Sampling
    - Instruction time-in-CSECT
- **New z/OS HIS started task**
  - Gathered on an LPAR basis
  - Writes SMF 113 records
- **New z/VM Monitor Records**
  - Gathered on an LPAR basis – all guests are aggregated
  - Writes new Domain 5 (Processor) Record 13 (CPU MF Counters) records
- **Minimal overhead**

New Hardware Capabilities to Size z/OS Workloads

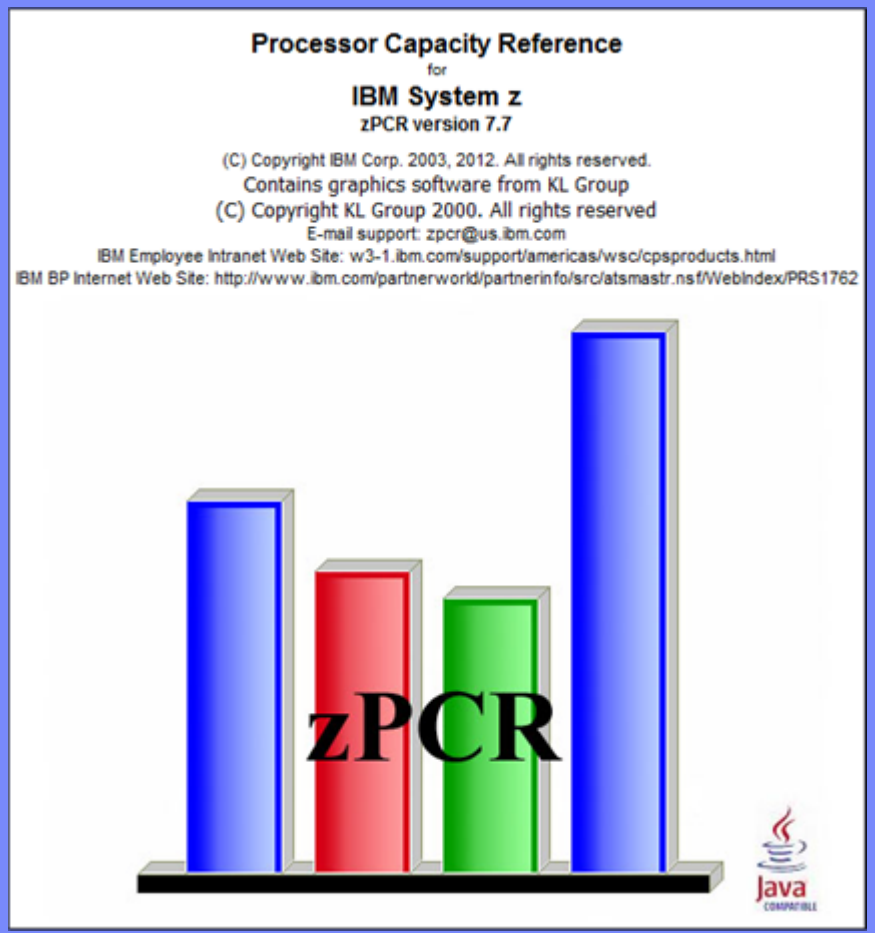

LSPR BENCHMARKS

Manual Input

RMF Report

CP3KEXTR

SMF 70
- **LPAR Data**

SMF 113
- **Counters cache / memory hierarchy**

Importance of using CPU MF Counters

- **New CPU MF Counters provide better information to do more successful capacity planning**
- **Same data used to validate the LSPR workloads can now be obtained from production systems**
  - Matches your production workload to the LSPR workloads
    - zPCR automatically processes CPU MF data to provide a match
      - Based on Relative Nest Intensity (RNI)
- **CPU MF Counters also useful for performance analysis**

Challenge to Use SMF to Select a LSPR Workload Mix

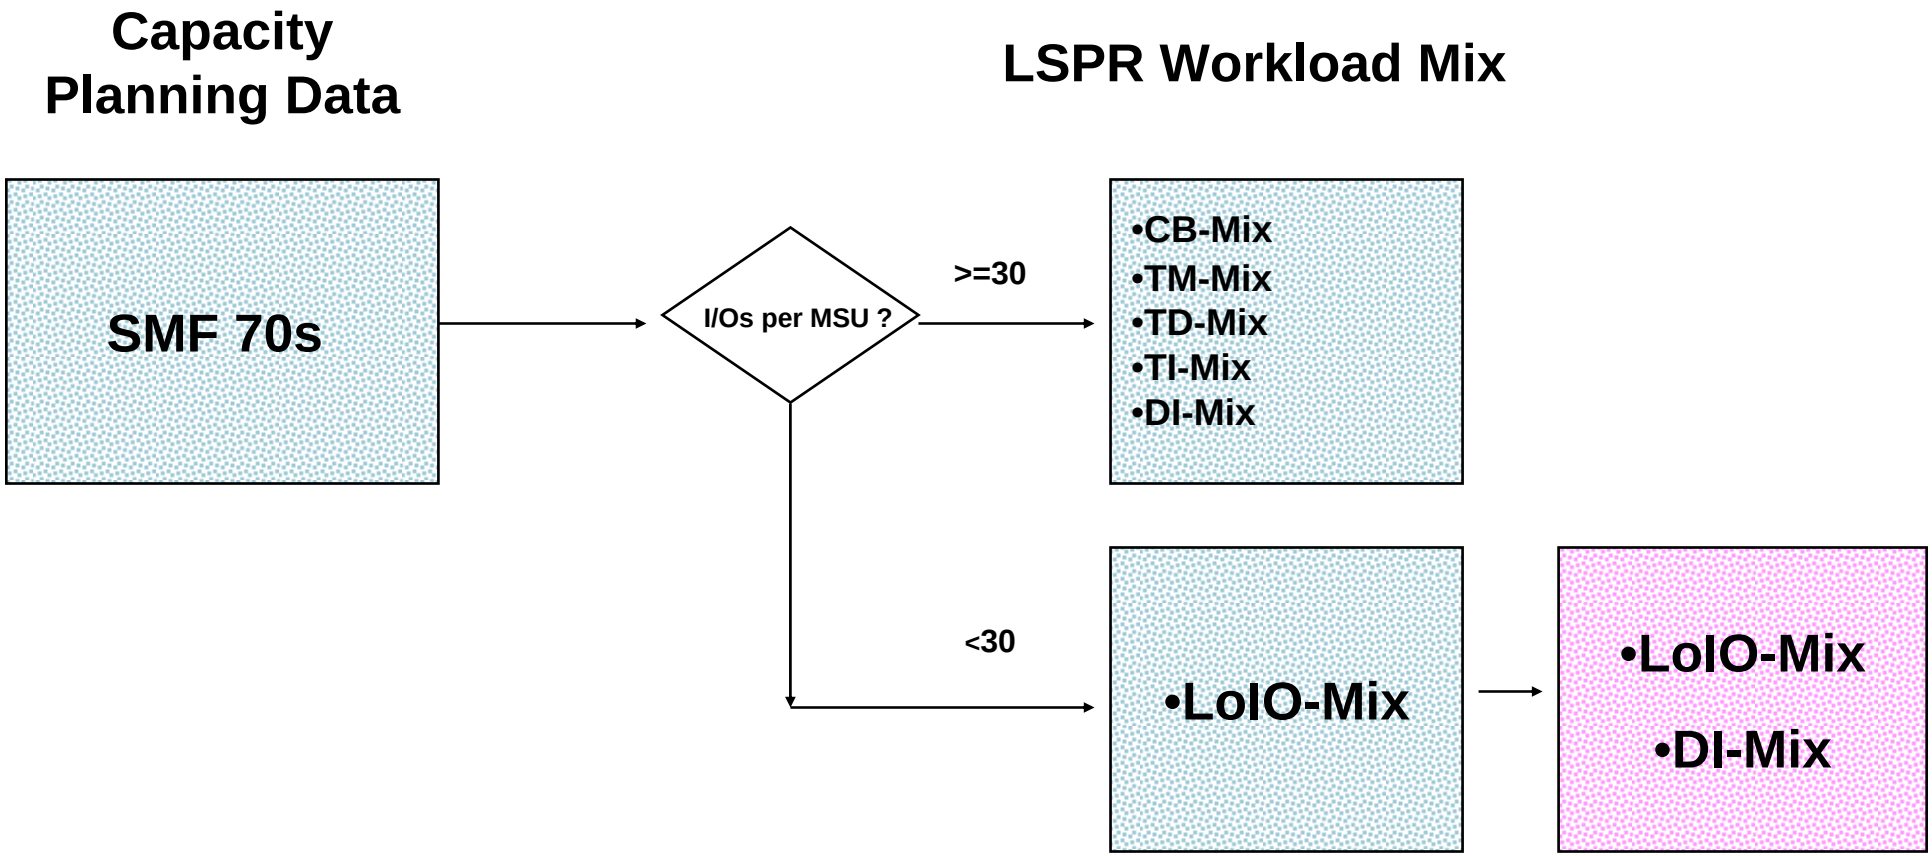

SMF 113s Provide Better LSPR Workload Selection

Capacity Planning Data

LSPR Workload Category

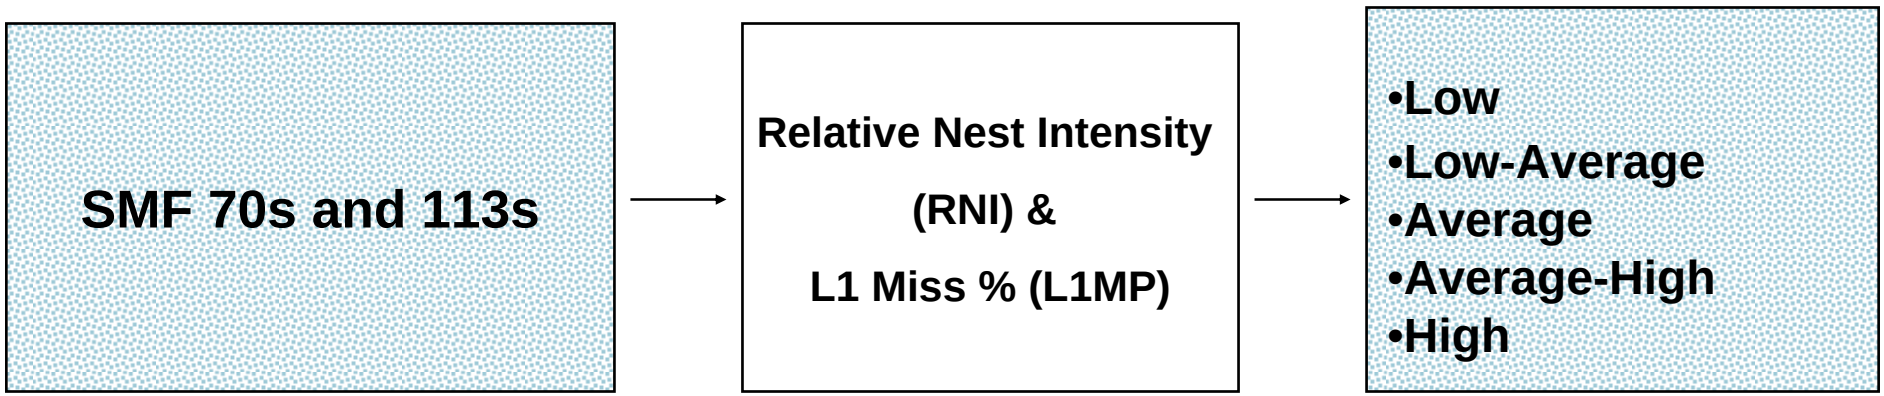

Validation and Enablement Details

Workload Category Validation Process

- **Analyzed customer data**

- **Initial z10 Study**
  - 100 z10 customer LPARs
  - SMF 70s and SMF 113s
  - Built relationship between performance and Nest
  - Built new metric to more precisely match workloads
    - Relative Nest Intensity (RNI)
  - LSPR workload categories defined as Low, Average, and High RNI

- **Validated RNI metrics with review of z10 to z196 Migrations**
  - 75 z10 to z196 customer LPAR migrations
  - Validated RNI based workload match

CPU MF
z10 Customer Workload Characterization Summary

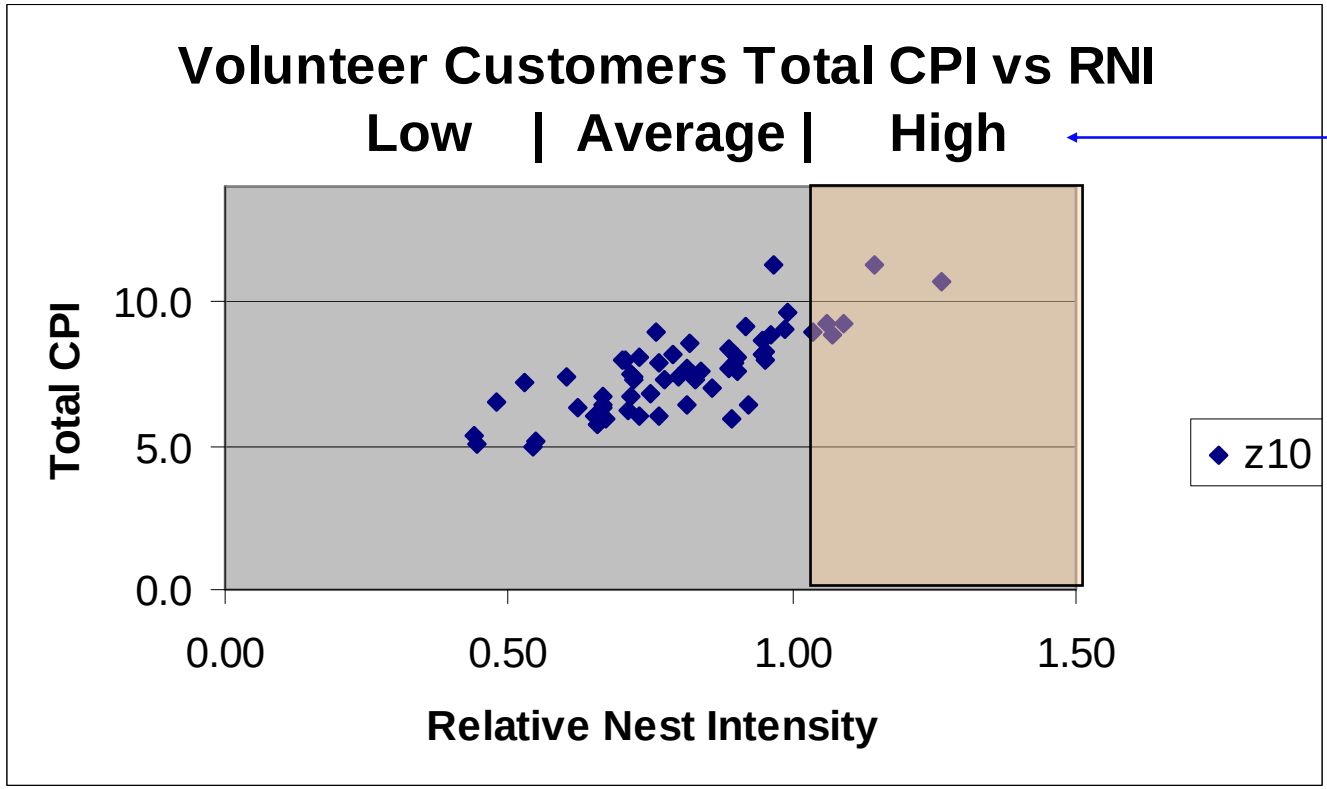

1) Customer CPI measurements

2) Created new RNI metric

3) Created new LSPR Workload Categories

DI-Mix

Requirements to Enable CPU MF Counters

- **Processor requirements**
  - z10, z196, or z114
- **z/OS requirements**
  - z/OS 1.10 or higher
    - With APAR OA30486
- **z/VM requirements**
  - z/VM at 5.4 or higher
    - With APAR VM64961

z/OS Steps to Enable CPU MF Counters

- **1 - Configure the processor to collect CPU MF**

  ___ Update the LPAR Security Tabs, can be done dynamically

- **2 - Set up HIS and z/OS to collect CPU MF**

  ___ Set up HIS Proc

  ___ Set up OMVS Directory - required

  ___ Collect SMF 113s via SMFPRMxx

- **3 - Collect CPU MF COUNTERs**

  ___ Start HIS

  ___ Modify HIS: "F HIS,B,TT='Text',PATH='/his/',CTRONLY,CTR=(B,E),SI=SYNC"

  – **Recommend to start HIS, Modify for Counters, and continuously run**

SMF 113s Space Requirements

- The SMF 113 record puts minimal pressure on SMF
  - 452 bytes for each logical processor per interval
- Example below is from 3 z196s processors
  - 713, 716 and 718
  - 10 Systems
  - 5 Days, 24 hours
- SMF 113s were 1.2% of the space compared to SMF 70s & 72s

RECORD TYPE	RECORDS READ	PERCENT OF TOTAL	**AVG. RECORD LENGTH**	MIN. RECORD LENGTH	MAX. RECORD LENGTH	RECORDS WRITTEN	Total Size (with AVG. Record Size)	% Total Size (with AVG. Record Size)
70	14,250	1.8%	14,236	640	32,736	14,250	202,865,850	15.1%
72	744,014	93.5%	1,516	1,104	20,316	744,014	1,128,252,590	83.7%
113	37,098	4.7%	**452**	**452**	**452**	37,098	16,768,296	**1.2%**
TOTAL	795,362	100.0%	1,695	18	32,736	795,362	1,347,886,736	100.0%

Operations – Display Command

```
F HIS,B,TT='BE Counters',PATH='/his/',CTRONLY,CTR=(B,E),SI=SYNC
```

D HIS

```
RESPONSE=SYSD
 HIS015I 10.15.54 DISPLAY HIS 286
  HIS       0025 ACTIVE
  COMMAND: MODIFY HIS,B,TT='BE Counters',PATH='/his/',CTRONLY,CTR=(B,E),
           SI=SYNC
  START TIME:  2012/04/12 10:15:45
  END TIME:    ----/--/-- --:--:--
  COMPLETION STATUS: --------
  FILE PREFIX: SYSHIS20120412.101545.
  COUNTER VERSION NUMBER 1: 1   COUNTER VERSION NUMBER 2: 2
  COMMAND PARAMETER VALUES USED:
   TITLE=  BE Counters
   PATH=   /his/
   COUNTER SET= BASIC, EXTENDED
   DURATION= NOLIMIT
   CTRONLY
   DATALOSS= IGNORE
   STATECHANGE= SAVE
   SMFINTVAL= SYNC
```

Use CPU MF Counters for Performance Analysis

- **Counters can be used as a secondary source to:**
  - Supplement current performance data from SMF, RMF, DB2, CICS, etc.
  - Help understand why performance may have changed

- **Some examples of usage include:**
  - HiperDispatch Impact
  - Configuration changes (Additional LPARs)
  - 1 MB Page implementation
  - Application Changes (e.g. CICS Threadsafe Vs QR)
  - Estimating Utilization Effect for capacity planning
  - z196 GHz change in Power Saving Mode
  - Crypto CPACF usage

Questions ?

Part 2 Enrollment Information

- **Part 2 – Specific Education – June 14, 2012 – 2:00-3:30 PM ET**
- **Data Collection Requirement**
  - Send SMF Data from 1 LPAR (your choice)
    - One LPAR per Enterprise - production preferred
    - SMF 70s (CPU) and SMF 113s (CPU MF Counters)
    - 4 hours from 1 day
- **Process**
  - Review TECHDOC for Detailed CPU MF Implementation and FTP Directions
    - http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS4922
  - Once you've enabled CPU MF Counters and decided to send SMF data to IBM
    - FTP the SMF data to IBM
    - Email John Burg (jpburg@us.ibm.com) and cc Kathy Walsh (walshk@us.ibm.com) Subject: "CPU MF Data for Part 2" and include
      - > Your name and company name.
      - > The file name that was successfully FTP'd to IBM
      - > Supply email IDs (up to 5) for Reports and Part 2 Webinar invitations
    - We'll' confirm back to you:
      - > Successfully downloaded your SMF data (target within 1 week of your email)
      - > By June 13th, we'll send your CPU MF Report and the Part 2 Webinar call in information
  - **We need the SMF data successfully received by Friday June 1st**

Questions ?

Summary

- **CPU MF Counters provide better information for more successful capacity planning**
- **Same data used to validate the LSPR workloads can now be obtained from production systems**
- **CPU MF Counters can also be useful for performance analysis**
- **Enable CPU MF Counters Today!**
  - Continuously collect SMF 113s for your production systems

Thank You for Participating!