



PSSC – IBM Customer Center Montpellier

IBM System z9 Platform Speciality engines Large Systems Update November 14, 2006


Karl-Erik Stenfors
zAtlas team member
Mainframe evangelist
IBM PSSC, Montpellier
kestenfors@fr.ibm.com

© 2006 IBM Corporation

PSSC - IBM Customer Center Montpellier

Z9 EC 12-way MCM

- **Advanced 95mm x 95mm MCM**
 - ▶ 102 Glass Ceramic layers
 - ▶ 16 chip sites, 217 capacitors
 - ▶ 0.545 km of internal wire

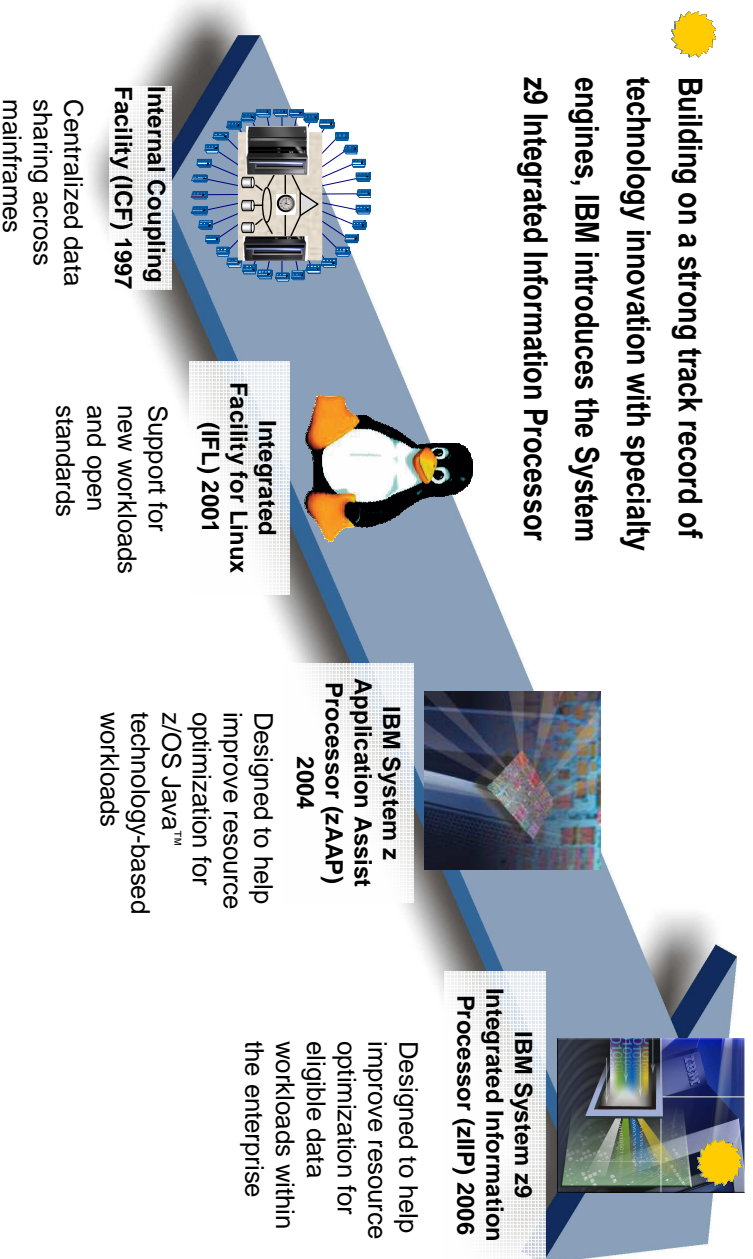


- **CMOS 10Ks0 chip Technology**
 - ▶ PU, SC, SD and MSC chips
 - ▶ Copper interconnections, 10 copper layers
 - ▶ 8 PU chips/MCM
 - 15.78 mm x 11.84 mm
 - 121 million transistors/chip
 - L1 cache/PU
 - 256 KB I-cache
 - 256 KB D-cache
 - 0.58 ns Cycle Time
 - ▶ 4 System Data (SD) cache chips/MCM
 - 15.66 mm x 15.40mm
 - L2 cache per Book
 - 660 million transistors/chip
 - 40 MB
 - ▶ One Storage Control (SC) chip
 - 16.41mm x 16.41mm
 - 162 million transistors
 - L2 cache crosspoint switch
 - L2 access rings to/from other MCMs
 - ▶ Two Memory Storage Control (MSC) chips
 - 14.31 mm x 14.31 mm
 - 24 million transistors/chip
 - Memory cards (L3) interface to L2
 - L2 access to/from MBAs (off MCM)
 - ▶ One Clock (CLK) chip - CMOS 8S
 - Clock and ETR Receiver



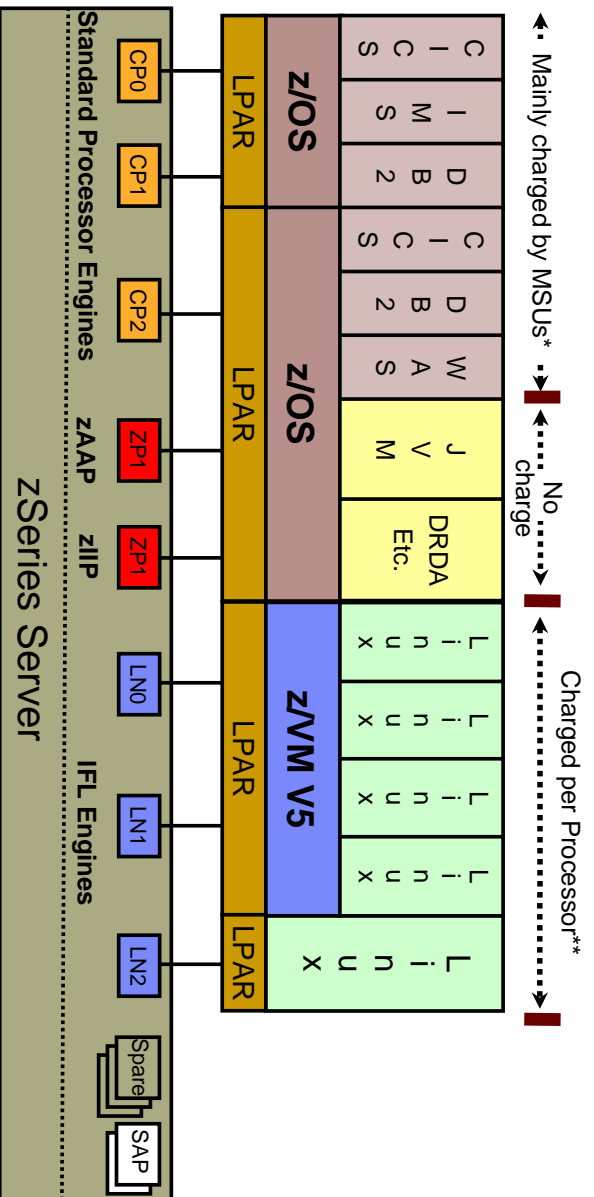
Technology evolution with speciality engines

Building on a strong track record of technology innovation with speciality engines, IBM introduces the System z9 Integrated Information Processor



3

System z9 speciality engines



* Based on MSUs
 ** Or "Engine"

4

System z9 PU Characterization

- **The type of Processor Units (PUs) that can be ordered on System z9 :**
 - ▶ Central Processor (CP)
 - Provides processing capacity for z/Architecture™ and ESA/390 instruction sets
 - Runs z/OS, z/VM, VSE/ESA, z/VSE, TPF4, Z/TPF, Linux for System z and Linux under z/VM or Coupling Facility
 - z9 EC has Capacity Marker features NOT Unassigned CP features
 - ▶ IBM System z Application Assist Processor (ZAAP)
 - Under z/OS, the Java Virtual Machine (JVM) assists with Java processing to a ZAAP
 - ▶ IBM System z9 Integrated Information Processor (zIIP) – when available
 - Provides processing capacity for selected workloads e.g., DB2 for z/OS V8 workloads executing in SRB mode
 - ▶ Integrated Facility for Linux (IFL)
 - Provides additional processing capacity for Linux workloads
 - ▶ Internal Coupling Facility (ICF)
 - Provides additional processing capacity for the execution of the Coupling Facility Control Code (CFCC) in a CF LPAR
 - ▶ Optional System Assist Processors (SAP)
 - SAP manages the start and ending of I/O operations for all Logical Partitions and all attached I/O

5

System z Application Assist Processor (ZAAP)

- **ZAAP for on demand business integration & infrastructure simplification**
 - ▶ Leveraged by workloads with Java cycles, e.g.: WebSphere®, WebServices Gateway (now WAS ESB), DB2 as well as home-grown JAVA applications
 - ▶ Can help simplify and reduce server infrastructure and improve operational efficiencies
 - ▶ Enables integration of on demand business applications with mission-critical database workloads
 - ▶ Potential operational advantages over distributed multi-tier solutions
- **Available on System z9, z990, and z890**
 - ▶ Executes Java cycles with no anticipated changes to applications
 - ▶ Enabled by the **IBM JVM**, z/OS 1.6 or higher and PR/SM™ virtualization

6



Requirements for zAAP Exploitation

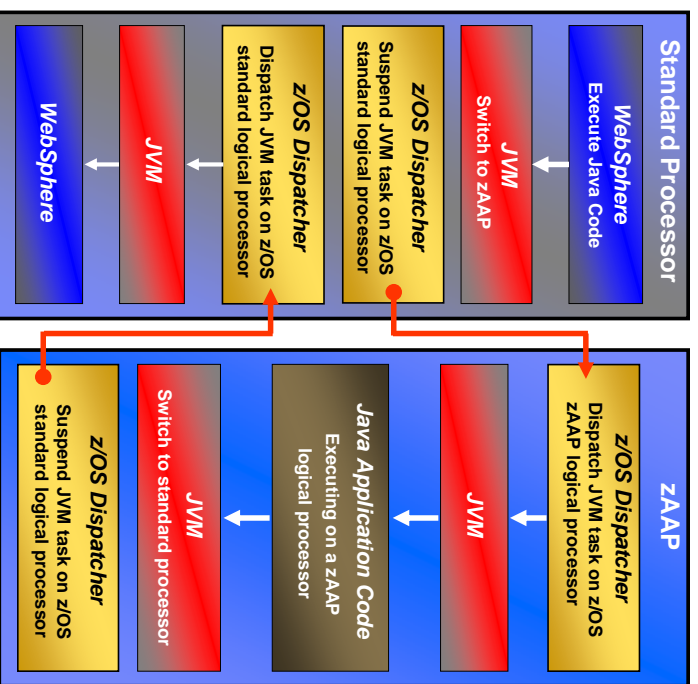
- Available on System z9, z990, and z890
- Prerequisites:
 - ▶ z/OS 1.6 (or z/OS.e 1.6 on z890)
 - ▶ IBM SDK for z/OS, Java 2 Technology Edition, V1.4 with PTF for APAR PQ86689
- Subsystems and Applications using SDK 1.4 will exploit zAAPs automatically:
 - ▶ WAS 5.1
 - ▶ CICS®/TS 2.3
 - ▶ DB2 V8
 - ▶ IMS™ V8
 - ▶ WebSphere WBI for z/OS
 - ▶ Web Services Gateway (now WAS ESB)
 - ▶ Home-grown JAV/A applications
- zAAPs must be jointly configured with general purpose processors within z/OS LPARs
 - ▶ Number of zAAPs may not exceed the number of permanently purchased CPUs (including z990 unassigned CPUs or z890 Downgrade - Record Only CPUs) on a given machine model

7



zAAP Architecture and Workflow: Executing Java under IBM JVM control


- IBM JVM, parts of Language Environment® runtime, and z/OS Supervisor needed to support JVM execution can operate on zAAPs
- IBM JVM communicates to z/OS dispatcher when Java code is to be executed
 - ▶ When Java is to be executed, the work unit is "eligible" to be dispatched on a zAAP
- z/OS dispatcher attempts to dispatch zAAP eligible work on a zAAP (when present)
 - ▶ zAAP ineligible work only dispatched on standard processors
- If there is insufficient zAAP capacity available, or standard processors are idle, the dispatcher may dispatch zAAP eligible work on a standard processor
 - ▶ There is an installation control to limit the use of standard processors to execute zAAP eligible work (see Java code execution options)



8

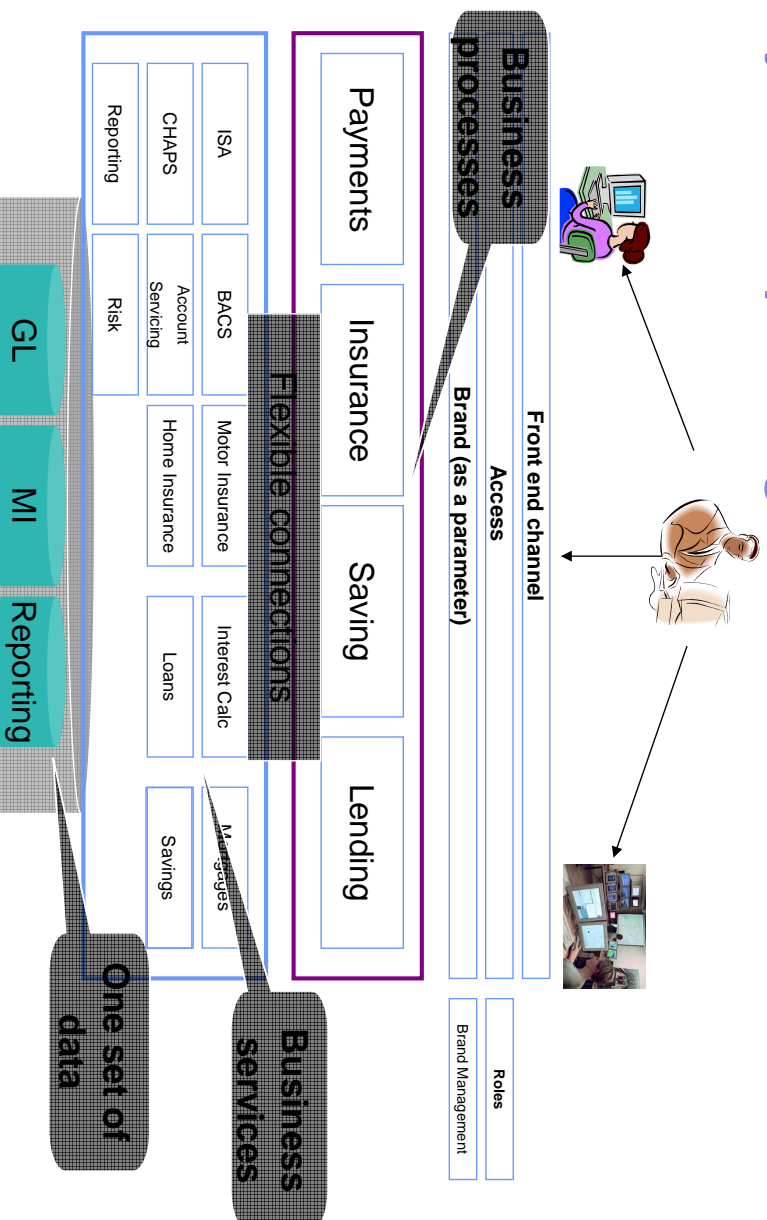
ZAAP Characteristics: How do ZAAPs differ from CPSs?

- **ZAAPs Limitations**
 - ▶ ZAAPs cannot be IPLed
 - ▶ Only executes z/Architecture mode instructions
 - ▶ ZAAPs do not support all manual operator controls
 - No: PSW Restart, LOAD or LOAD derivatives (load from file, CDROM, Server)
 - ▶ ZAAPs don't respond to SIGP requests unless enabled by a z/OS that supports ZAAPs
 - ▶ Additional architecture differences are anticipated in future implementations
 - e.g., Java specific performance enhancements
- **The z/OS design accommodates processor differences for ZAAPs:**
 - ▶ No I/O interrupts
 - ▶ No Clock Comparator interrupts
 - ▶ No affinity scheduling



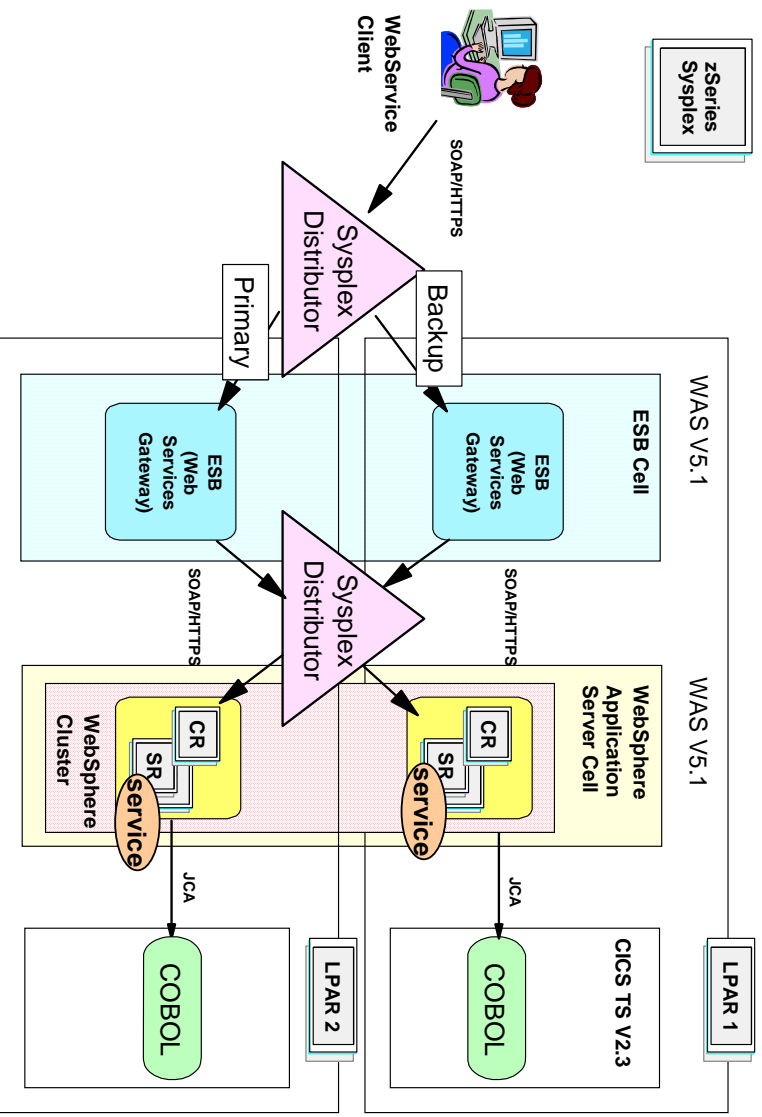
9

Project example - target SOA solution



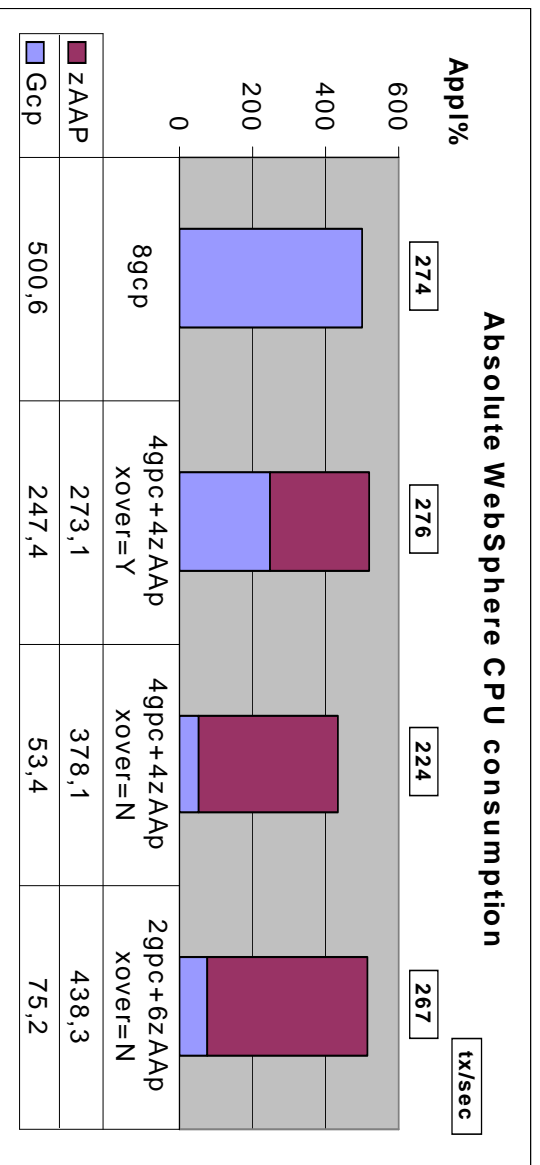
10

Tested configuration



11

ZAAP results

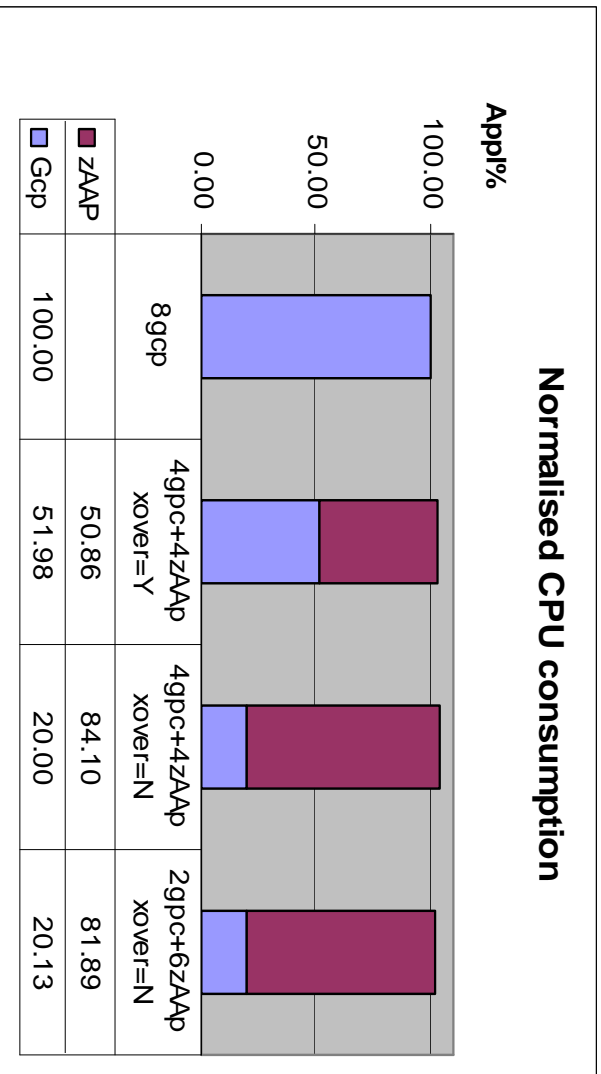


Note: Tests done prior to ZAAP IEAOPTxx parameter changes introduced in OA14131

12



ZAAP results




CPU consumption normalised for same Tx/Sec

13



What about TCO with ZAAPS ?



With ZAAP processors, System z savings would have been 37%

Source: Scorpion Study 1999 - 2004

14

DB2 V8 exploitation of zIIP

1. ERP or CRM application serving*

- For applications running on z/OS, UNIX®, Linux, Intel®, or Linux on System z that access DB2 for z/OS V8 on a System z9 EC, via DRDA® over a TCP/IP connection. DB2 gives z/OS the necessary information to have portions of these SQL requests directed to the zIIP



2. Data warehousing applications*

- Requests that utilize DB2 for z/OS V8 complex star schema parallel queries may have portions of these SQL requests directed to the zIIP when DB2 gives z/OS the necessary information

3. Some DB2 for z/OS V8 utilities*

- A portion of DB2 utility functions used to maintain index maintenance structures (LOAD, REORG, and REBUILD INDEX) typically run during batch, can be redirected to zIIP.

* The zIIP is designed so that a program can work with z/OS to have a portion of its Service Request Block (SRB) enclave work directed to the zIIP. The above types of DB2 V8 work are those executing in SRB enclaves, portions of which can be sent to the zIIP.

15

How does the zIIP work ?

The zIIP is designed so that a program can work with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to the zIIP. The types of DB2 V8 work listed below are those executing in enclave SRBs, portions of which can be sent to the zIIP.

Example 1 = Distributed SQL requests (DRDA)

Queries that access DB2 for z/OS V8 via DRDA over a TCP/IP connection are dispatched within z/OS in enclave SRBs. z/OS directs a portion of this work to the zIIP.

Example 2 = Complex parallel query (BI)

Complex star schema parallel queries will now use enclave SRBs. z/OS directs a portion of this work to the zIIP.


Example 3 = DB2 utilities for index maintenance

DB2 utilities LOAD, REORG, and REBUILD will now use enclave SRBs for the portion of the processing that is related to index maintenance. z/OS directs a portion of this work to the zIIP.

16

Example for zIIP data serving workload SAP solutions

- SAP NetWeaver based SAP Solutions can exploit the benefits of the IBM zIIP




May enable growth of System z9 SAP workloads through resource optimization

17

Important technical notes for zIIP

- Utilization of the zIIP is expected to be transparent to the application
 - No anticipated changes to applications that use DB2 for z/OS V8
- The enclave SRB interface is available upon request to non-IBM vendors as well
- Pre-requisites
 - z/OS or z/OS.e 1.6 or higher
 - DB2 for z/OS V8
 - System z9 with zIIP




18


Cúram V4 Application Benchmark



- Curam V4 Application: Social Services Integrated Case Management (pension payments)
- Solution based on Websphere V6 on AIX 5.3 and DB2 UDB V8 on z/OS 1.6.
- Test objectives: validate Curam V4 performance and build sizing methodology for Curam application:
 - ✓ Online Test: 37230 Business Transactions per hour for 16,000 users
 - ✓ Generate Payment: 2.6 Million payments generated
 - ✓ Payment Batch: 2.6 Million payments processed

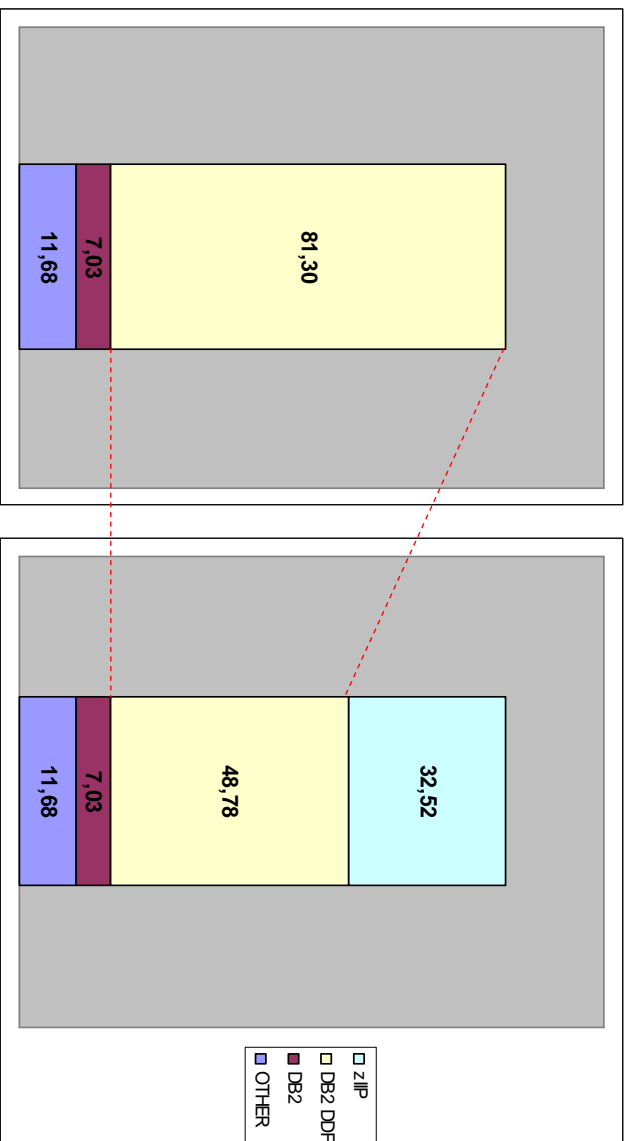


CURAM V4 – Online Workload



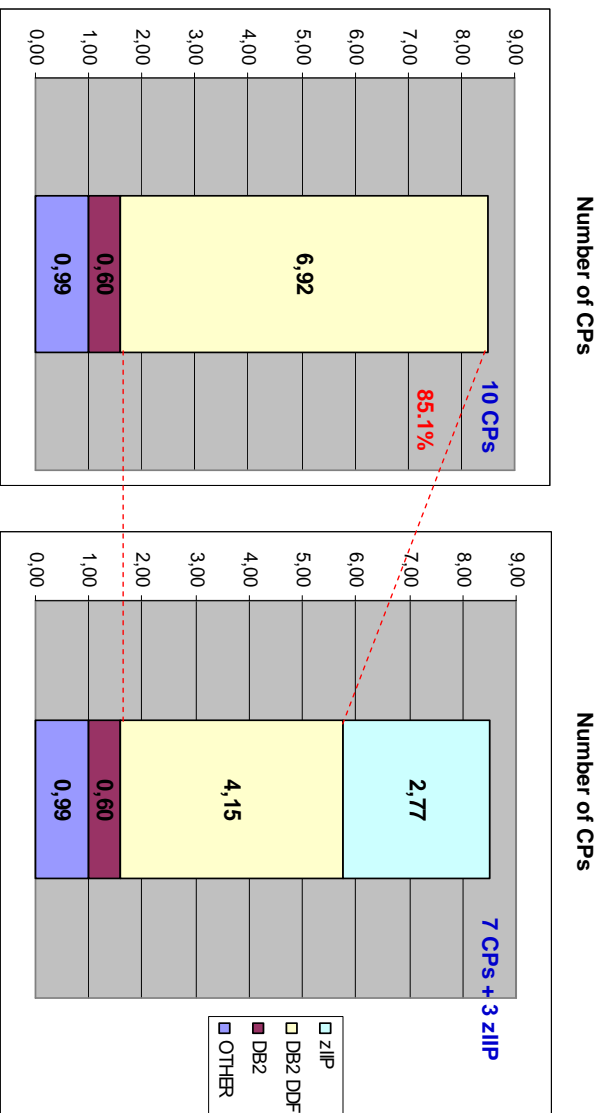
Average CPU Busy : 53.2 % - Capture Ratio : 0.92
 Average DASD Rate : 2,526 IO/sec - 16,000 active users

CURAM V4 zIIP Estimation (Online Workload)



16,000 Active Users

CURAM V4 zIIP Estimation (Online Workload)




Z9 EC MSU: 640

Z9 EC MSU: 479

16,000 Active Users


CURAM V4 – Batch Workload (Issue Payments)



Average CPU Busy : 76.1 % - Capture Ratio : 0.95
 Average DASD Rate : 1,372 IO/sec
 %DRDA for Batch workload is greater than for Online workload


23

CURAM V4 zIIP Estimation (Batch workload)



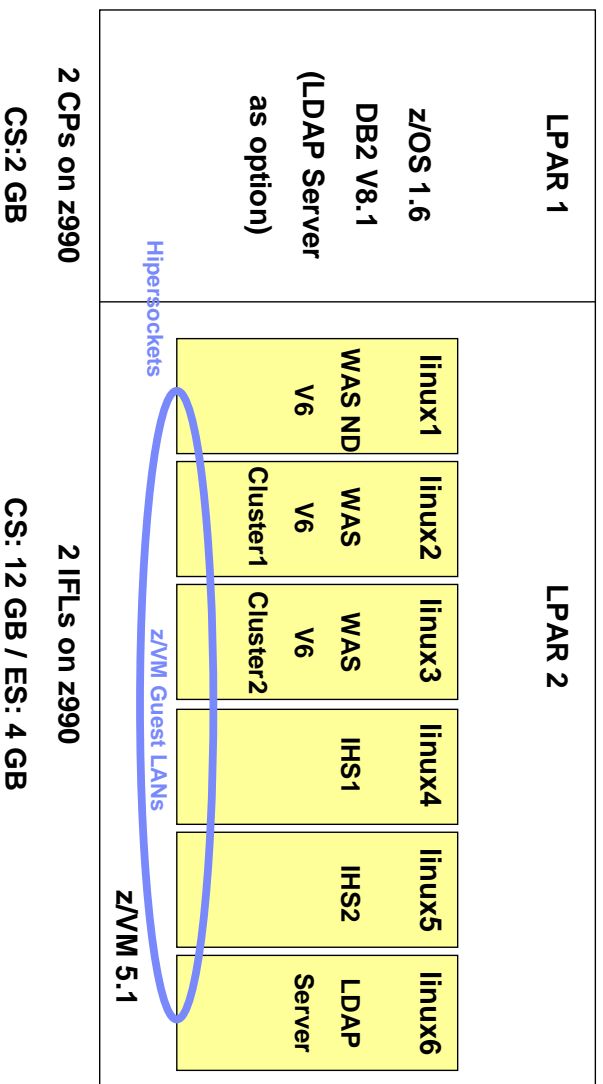
24

CURAM V4 zIIP Estimation (Batch Workload)





25

WebSphere V6 Application Benchmark (z990)



26

Websphere V6 Workload

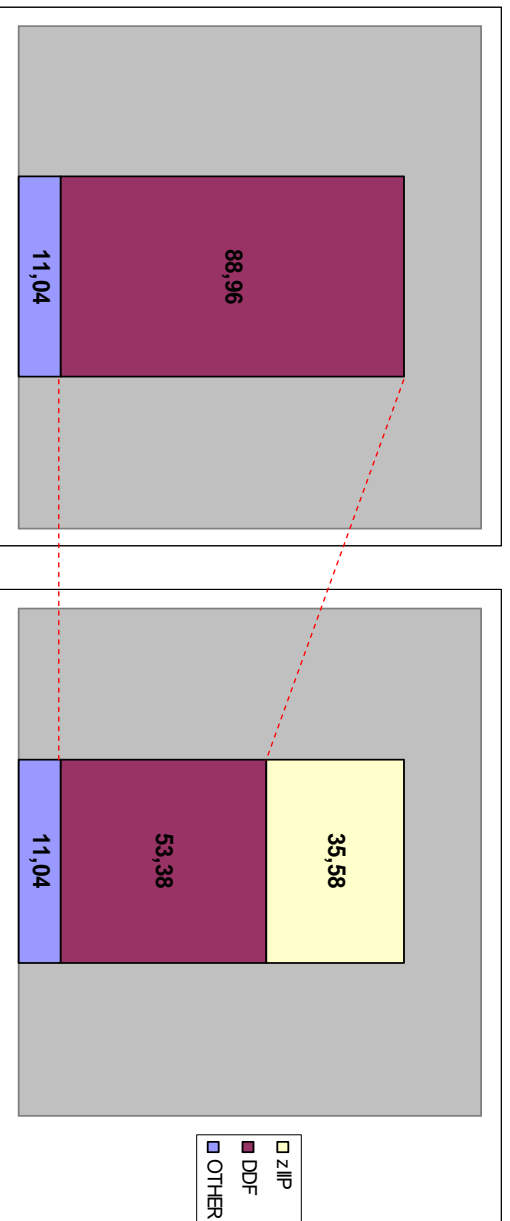


Capture Ratio : 0.95

If LDAP Server is running on z/OS, DDF % drops to 50%

27


Websphere V6 Workload – zIIP Estimation



91.5 transactions/sec

28

Websphere V6 Workload – zIIP Estimation



z9-109 MSU : 158

z9-109 MSU : 81

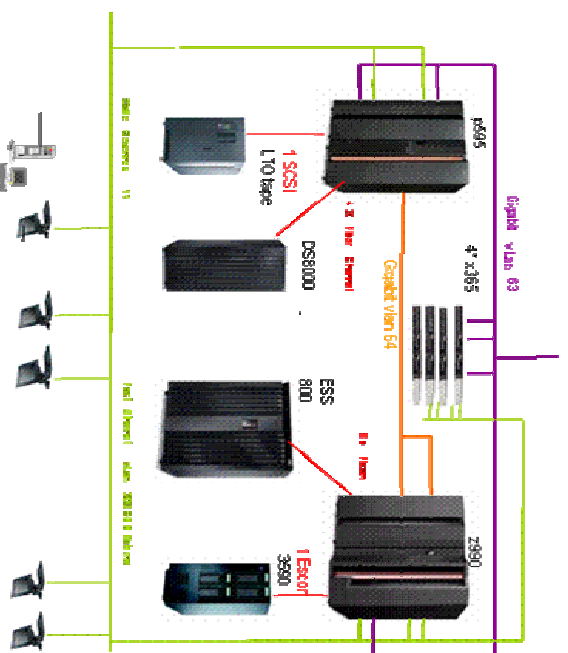
91.5 transactions/sec

29

SAP SCM 4.1 APO Benchmark




- SAP SCM 4.1 (Supply Chain Management) APO is SAP Production Planning Application.
- Solution based on SAP Application Server running on AIX and database server on DB2 UDB V8 on z/OS 1.6.
- Test objectives: provide customer with sizing information for SAP SCM APO application.
 - ✓ Batch APO R/3 integration in 3 hours
 - ✓ 580,000 purchase requisitions
 - ✓ 18,000 manufacturing orders
 - ✓ Database size is about 200 GB



30


SAP SCM 4.1 APO - Batch Workload



Capture Ratio : 0.92
 580,000 purchase requisitions / 18,000 manufacturing orders


31

SAP APO - zIIP Estimation (Batch Workload) *



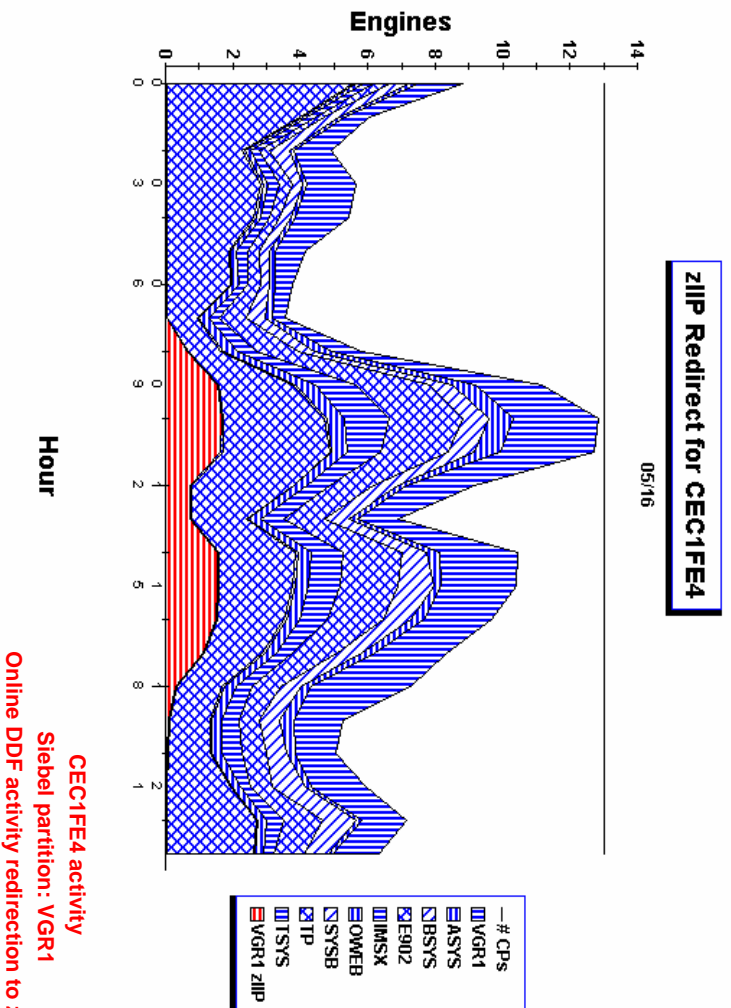
32

SAP APO – zIIP Estimation (Batch Workload) *



33

Siebel zIIP utilization for VGR1 online activity



34

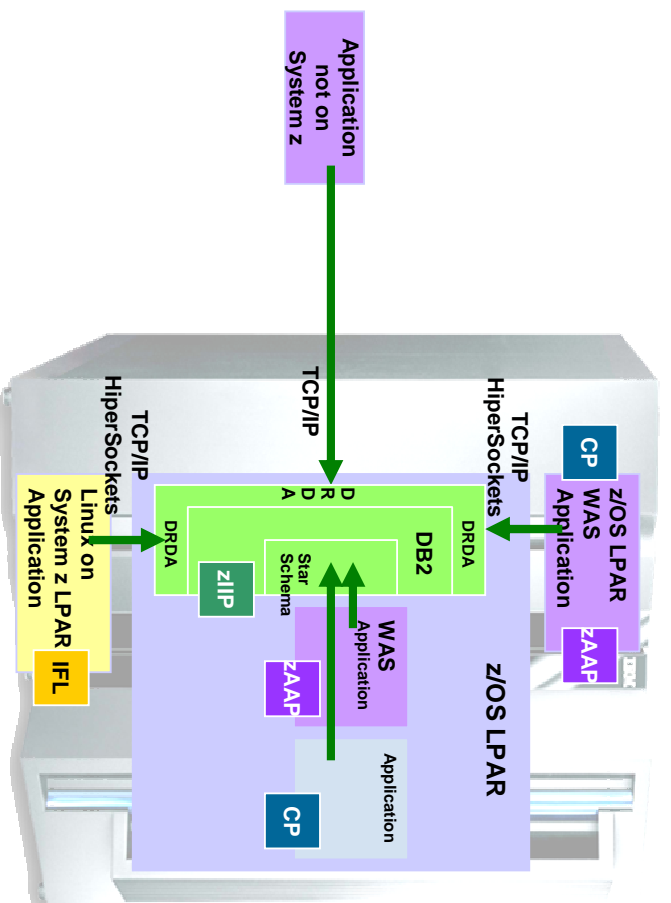
Siebel zIIP utilization for VGR1 online activity

- This analysis is based on the following DDF Services Classes DDF (DRDA TCP/IP) for Siebel online activity characterized in WLM (z/OS workload manager):
 - STCALL.DDFTP
 - STCALL.DDFTP1
 - STCALL.DDFTP2
 - STCALL.DDFTP3
 - Number of installed CPs (z9 model 713) :
 - 13 CPs loaded to 98.46 % during peak time
 - Target configuration with zIIP engines:
 - 12 CPs loaded at 93.33 % during peak time (**potential saving of 1 CP**)
 - 2 zIIPs loaded at 80.00% during peak time (**addition of 2 zIIPs**)
- Potential saving of 53 MSUs** on VGR1 logical partition

	Date	Time	# CPs	# zIIPs	Load %
Peak without zIIP	May 16	10h00	12.8	-	CP : 98.46%
CEC Peak with zIIP	May 16	10h00	11.2	1.6	CP : 93.33% zIIP : 80.00%
zIIP Peak	May 16	10h00	11.2	1.6	CP : 93.33% zIIP : 80.00%

35

Specialty engines work together



The IBM System z9 specialty engines can run independently or complement each other (shown at left, instances where specialty engines can be employed)

36

System z9 and DB2 for z/OS are an ideal data serving platform

- **Data consolidation helps reduce:**
 - ▶ Multiple copies, disparate data
 - ▶ Cost and complexity of back up and recovery
 - ▶ Network traffic
 - ▶ Amount of storage
 - ▶ DB administration and management
 - ▶ Risk associated with distributed privacy, security, and audit policies
- **Leverage System z technology**
 - ▶ Parallel Sysplex clustering for scalability AND availability AND performance
 - ▶ Data sharing = single view of the data
 - ▶ Data compression for TCO
 - ▶ Centralized backup, recovery, privacy, security and audit policies
- **New System z9 specialty engine designed to help:**
 - ▶ Customers integrate data across enterprise
 - ▶ Improve resource optimization
 - ▶ Lower the TCO for data serving workloads

