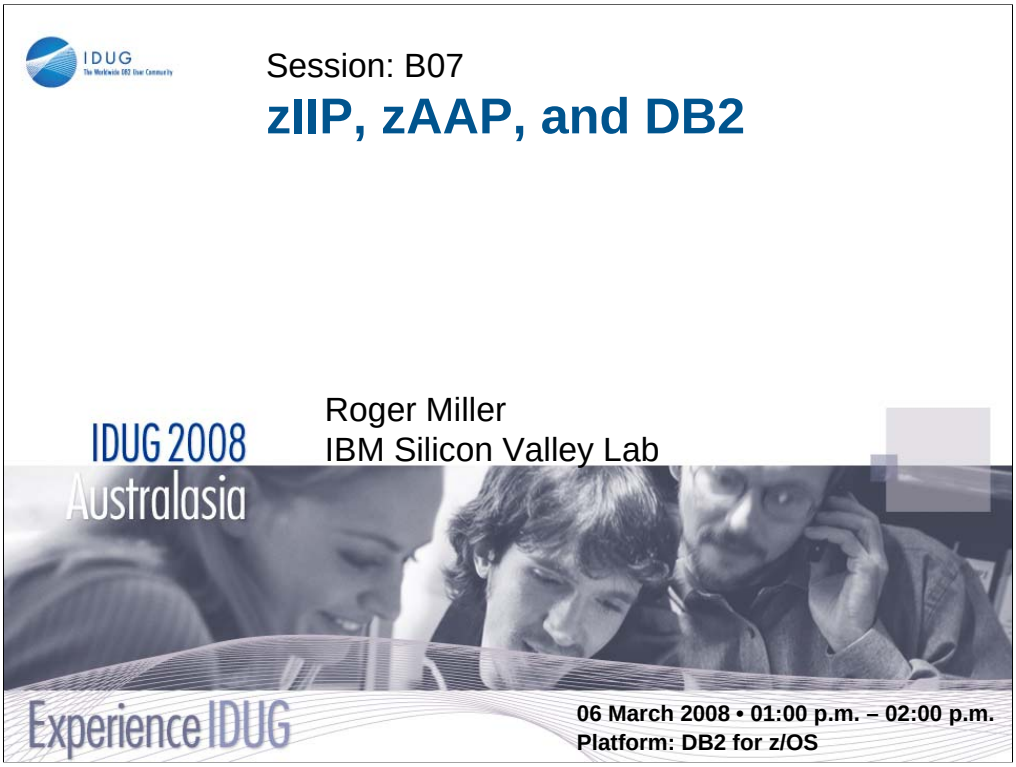

The IBM System z9 Integrated Information Processor (zIIP) is the latest customer inspired specialty engine for the IBM System z9 mainframe. The zIIP is designed to help improve resource optimization and lower the cost of portions of eligible workloads, enhancing the role of the System z9 mainframe as the data hub of the enterprise. DB2 V8 and DB2 9 use this recent improvement in hardware and z/OS to provide better performance, improved value, more resilience and better function. Deploying applications and workloads on specialty engines like the Integrated Facility for Linux (IFL), System z Application Assist Processor (zAAP), and System z9 Integrated Information Processor (zIIP) can help free up general purpose processor capacity that can be used by other work. The amount of white space and/or software savings you get from a specialty engine like a zIIP may vary and depend on the workload. In this session, we will discuss the latest improvements and news to help you understand how to implement and how much improvement the zIIP can provide for your environment.

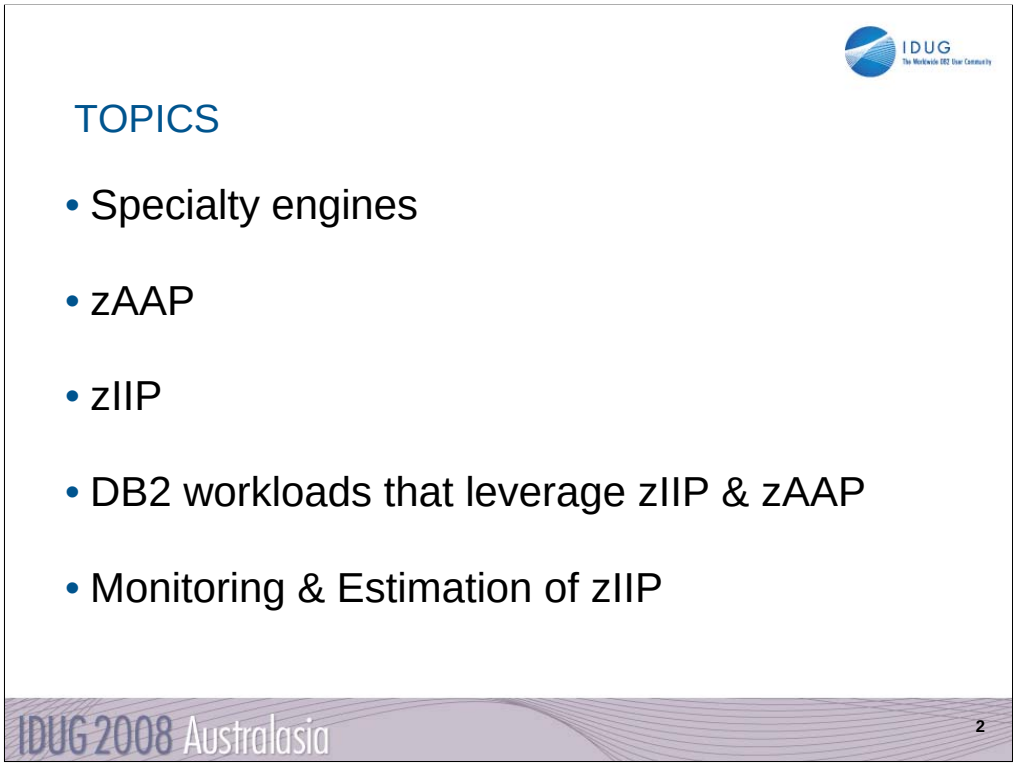

We'll start with an overview, then explain more about the work loads and what is eligible. Then we can work down to show some of the techniques for measurement and estimation, ending with a summary.

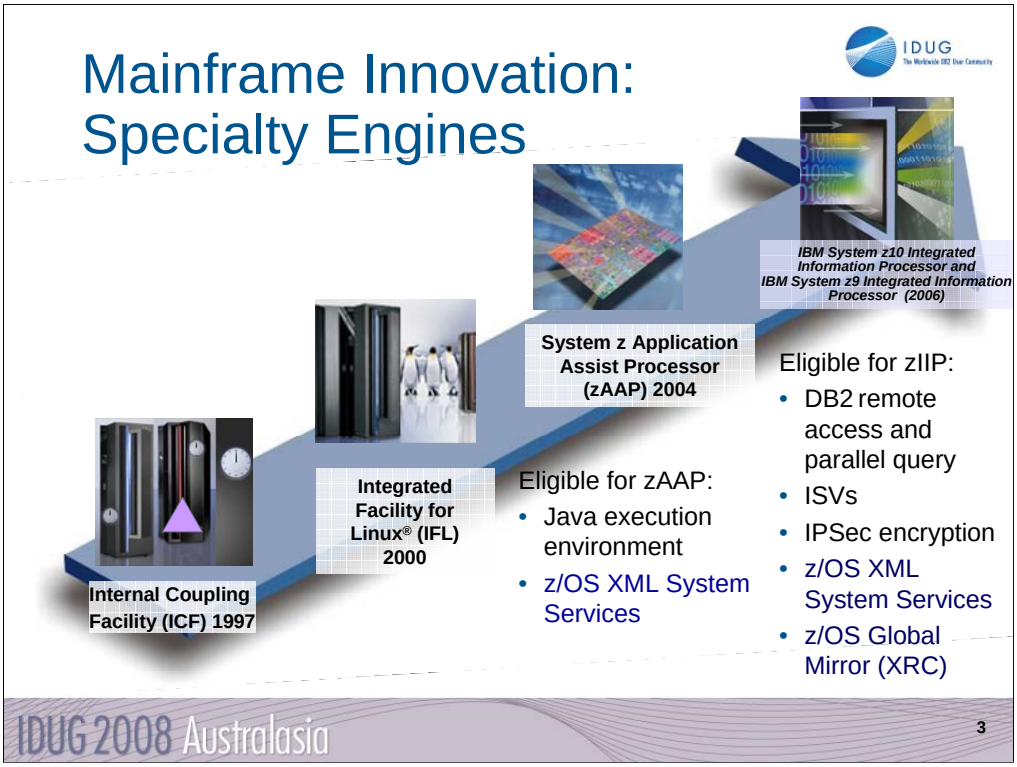

This chart shows the IBM System z family of specialty engines, with some that have been available for over ten years. The latest change to deliver is the ability to use the zIIP for IPSec encryption in August 2007. A statement of direction was provided for z/OS XML to be able to use zAAP or zIIP for XML parsing in April and August 2007. The February 2008 news is the ability to have the z/OS Global Mirror or XRC work redirected to zIIP.

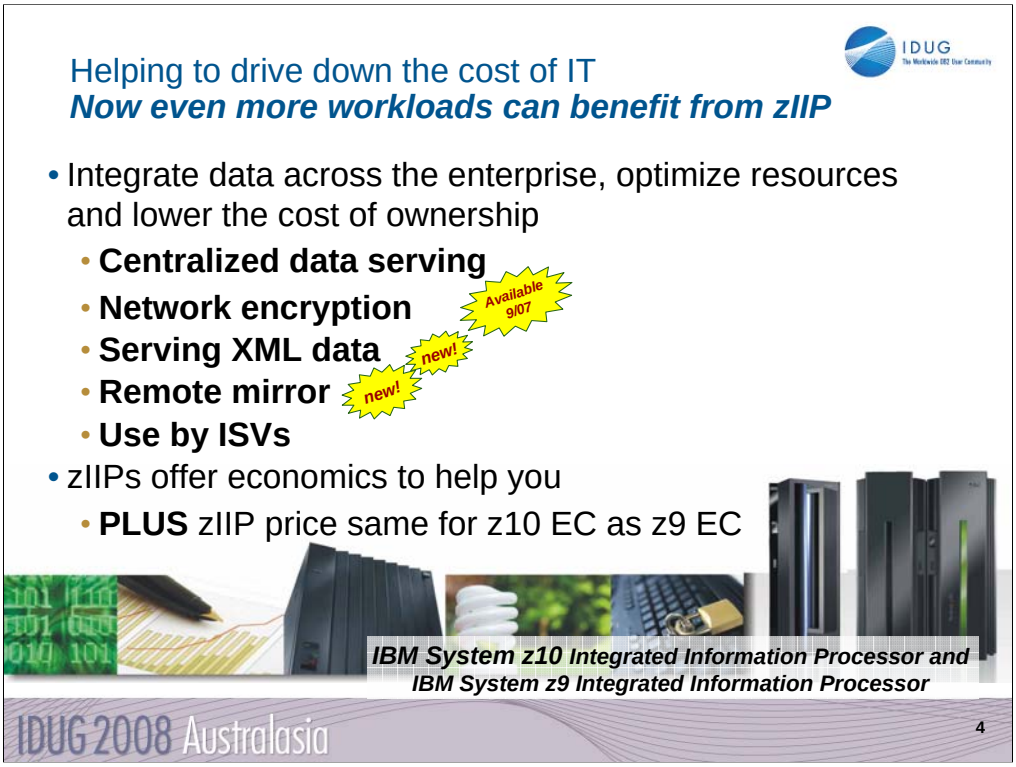

zIIP can help to integrate data across the enterprise by optimizing resources and lowering the cost of ownership for eligible data and transaction processing workloads

Centralized data serving – First to exploit zIIP were workloads such as BI, ERP, and CRM applications running on distributed servers with remote connectivity to DB2® V8

Network encryption – zIIP becomes an IPSec encryption engine helpful in creating highly secure connections in an enterprise

Serving XML data – zIIP is enabled for XML parsing, first to exploit this is inserting and saving DB2 9 XML data over DRDA®

Remote mirror – zIIP becomes a data mirroring engine with zIIP assisted z/OS Global Mirror function (zGM, formerly XRC) helpful in reducing server utilization at recovery site (with z/OS V1.8 and above)

Exploiting of zIIPs by ISVs

zIIPs offer economics to help you

PLUS zIIP price is same for z10 EC as z9 EC and we offer no charge MES

upgrades when moving to new technology

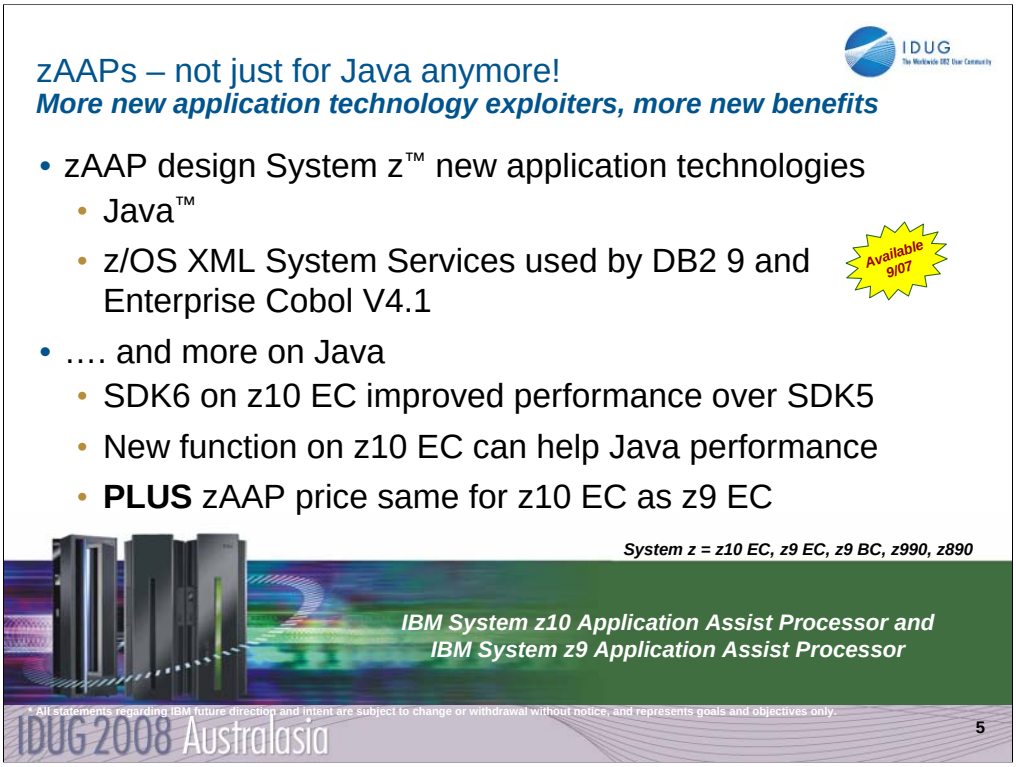

zAAPs – not just for Java anymore!

More new application technology exploiters, more new benefits

zAAP designed to help implement new application technologies on System z™

Java™ was the first - lowering the cost of computing for WebSphere® Application Server and other Java technology-based applications

z/OS XML System Services (introduced with z/OS V1.9 and rolled back to V1.8 and V1.7) helps make hosting XML data and transactions on System z more attractive. DB2 9 and Enterprise Cobol V4.1 are the first users.

…. and more on Java

SDK6 on z10 EC designed to deliver improved performance over SDK5 on z9 EC

New function on z10 EC may benefit Java performance

New z10 processor chip design and more available server memory .. plus in 2Q08 –support for new decimal floating point on z10 EC by Java BigDecimal *

PLUS zAAP price is same for z10 EC as z9 EC and we offer no charge MES upgrades when moving to new technology

Attractive pricing for new workloads
Helping Driving down the cost of IT

- Specialty engines – addressing more new workloads:
  - NEW: IPSec, XML, Global Mirror
  - Java, Data-serving
  - Linux® Internal Coupling Facility
- Specialty engines – designed to offer financial benefits now and in the future:
  - Technology investment protection
  - Lowering the cost of future growth with 'extra capacity at no extra cost'

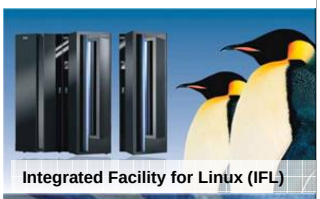
Integrated Facility for Linux (IFL)


IBM System z Application Assist Processor (zAAP)

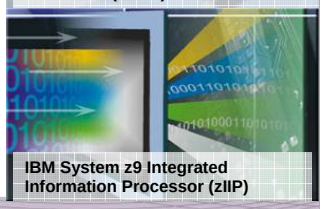
IBM System z9 Integrated Information Processor (zIIP)

The value of specialty engines continues to improve as more work becomes eligible. One new workload has delivered and two new ones announced recently.

Why do zIIPs, zAAPs and IFLs Reduce Cost?

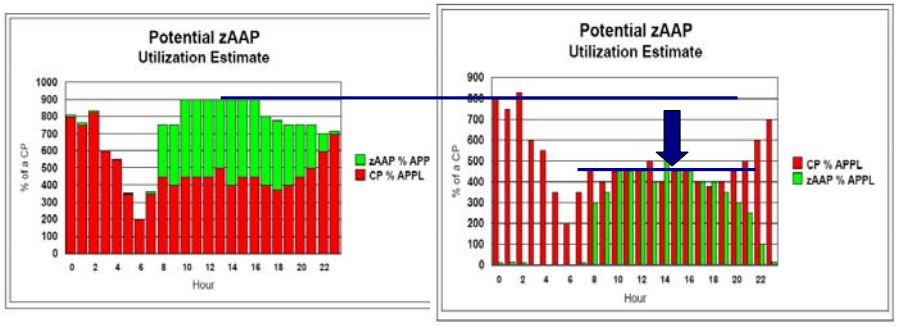

1.Hardware costs: By moving workload from general purpose processors to zIIP, zAAP and IFL processors (higher cost to lower cost processors).

2.Software Costs: license/maintenance costs based on number of and usage of general purpose central processors. Specialty engines can reduce # of CP's.

No z/OS software charges based on zIIP, zAAP and IFL processors or usage.

IDUG 2008 Australasia 7

zIIPs and zAAPs do not add functional capabilities, but they do address hardware and software costs. In addition to being lower cost processors for the specialty purposes, they allow you to reduce the license and maintenance cost for software on z/OS, as there is no z/OS software charge for processing running on zIIP, zAAP or IFL processors.

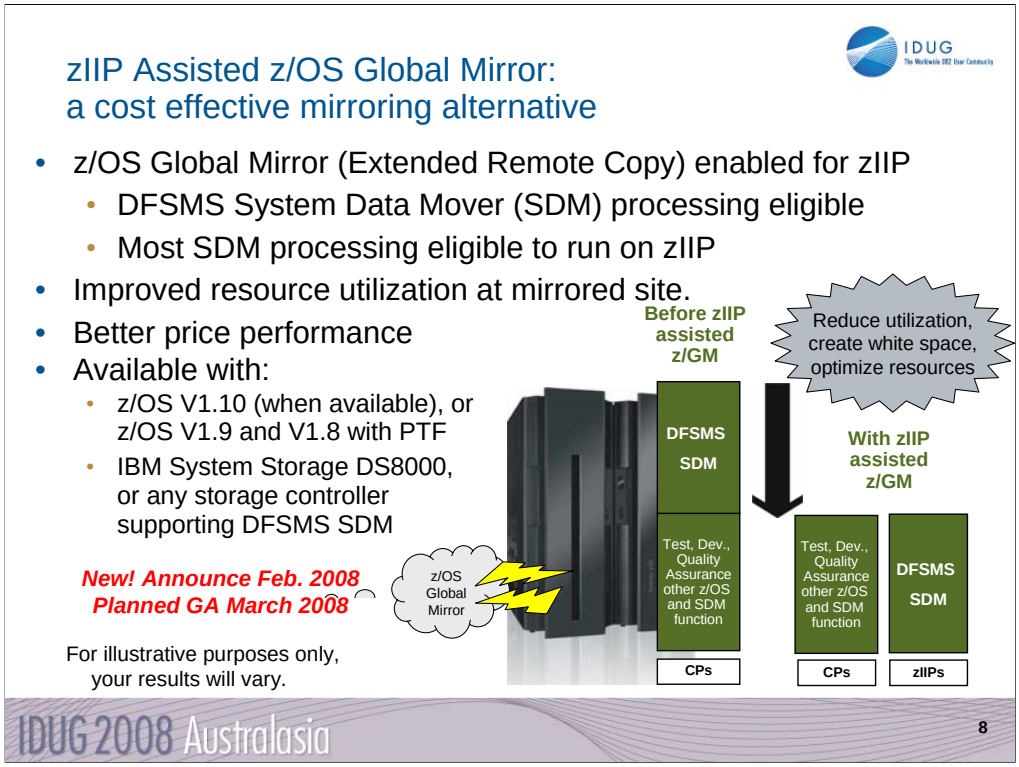

zIIP Assisted z/OS Global Mirror: a cost effective mirroring alternative
z/OS Global Mirror (formally Extended Remote Copy, XRC) is enabled for the zIIP. z/OS DFSMS allows System Data Mover (SDM) processing to be eligible for the zIIPs. Most SDM processing associated with zGM/XRC is made eligible to run on the zIIP.

zIIP assisted z/OS Global Mirror function, can help provide better price performance and improved utilization of resources at the mirrored site.

DFSMS SDM processing is redirected to a zIIP processor which can lower system utilization at the mirrored site.

Available with:

z/OS V1.10 (when available), or z/OS V1.9 and V1.8 with PTF for APAR #OA23174 (March, 2008)

IBM System Storage DS8000, or any storage controller supporting DFSMS SDM

August 2007 announcement - z/OS XML System Services*

1. NEW! z/OS XML System Services enabled to take advantage of zAAPs. Statement of Direction, at a future date:
2. IBM intends to enable z/OS XML to take additional advantage of zIIPs. (i.e. 100% zIIP redirect, greater than the current (about half) for DRDA)
3. IBM also intends to extend and expand the use of z/OS XML System Services with additional future enhancements:
   - IBM intends to enhance the XML Toolkit for z/OS so eligible workloads use z/OS XML. This allows eligible XML Toolkit processing to exploit zAAP.
   - IBM intends to add validating parsing to the z/OS XML component. This extends zAAP and zIIP exploitation to include XML validating parsing workload as well.

IDUG 2008 Australasia 9

STATEMENT OF DIRECTION: z/OS XML enabled for both zAAP and zIIP specialty processors
This SOD is broken down into 3 parts so that it is easier to explain and communicate.

In z/OS V1.8, IBM introduced a new system component of z/OS, z/OS XML System Services (z/OS XML), a system-level XML parser integrated with the base z/OS operating system and designed to deliver an optimized set of services for parsing XML documents (z/OS XML has also been made available on z/OS V1.7). The initial beneficiaries of this system component are middleware and applications requiring high performance non-validating XML parsing. z/OS XML may currently be accessed by an Assembler programming interface and one of the first exploiters, DB2 9 for z/OS, uses this Assembler interface for XML native support. IBM plans to add C/C++ support for z/OS XML with z/OS V1.9 (satisfying the Statement of Direction in Software Announcement 206-039, dated February 28, 2006).

1) IBM announced its intent to enable the z/OS XML component to take advantage of zAAPs. This future enhancement means that middleware and applications requesting z/OS XML System Services (for example DB2 processing via local connection) will have z/OS XML System Services processing execute on the zAAP. Specifically, all z/OS XML System Services parsing executing in TCB mode will be redirected to the zAAP.

2) In addition, IBM is announcing its intent to enable the z/OS XML component to fully take advantage of zIIPs, when present. With respect to DB2, z/OS XML processing may be partially directed to zIIPs when utilized as part of a distributed request (like DB2 DRDA). The future enhancement will further benefit DB2 by directing the full amount of the z/OS XML System Services processing to zIIPs when it is utilized as part of any zIIP eligible workload (like DRDA). Specifically, all z/OS XML System Services parsing that is executed in SRB mode from zIIP-eligible enclaves will be redirected to the zIIP.

zAAPs and zIIPs are designed to help free-up general computing capacity and lower total cost of operation for select new workloads such as Java, business intelligence (BI), ERP, CRM and IPSec encryption on the mainframe. IBM does not impose software charges on zAAP and zIIP capacity. Collectively, z/OS XML System Services support of zAAP and zIIP means that you have the advantages of XML processing on z/OS with TCO benefits of either the zIIP or the zAAP processor regardless of the invocation environment.

3) As part of a comprehensive plan, IBM intends to extend and expand on the use of z/OS XML System Services enabled for zAAP specialty processors as the basis for future enhancements:

IBM intends to enhance the XML Toolkit for z/OS so that eligible workloads may exploit the z/OS XML component -- this extends zAAP exploitation to the XML Toolkit for z/OS. IBM intends to add validating parsing to the z/OS XML component -- this extends zAAP exploitation for XML validating parsing as well.

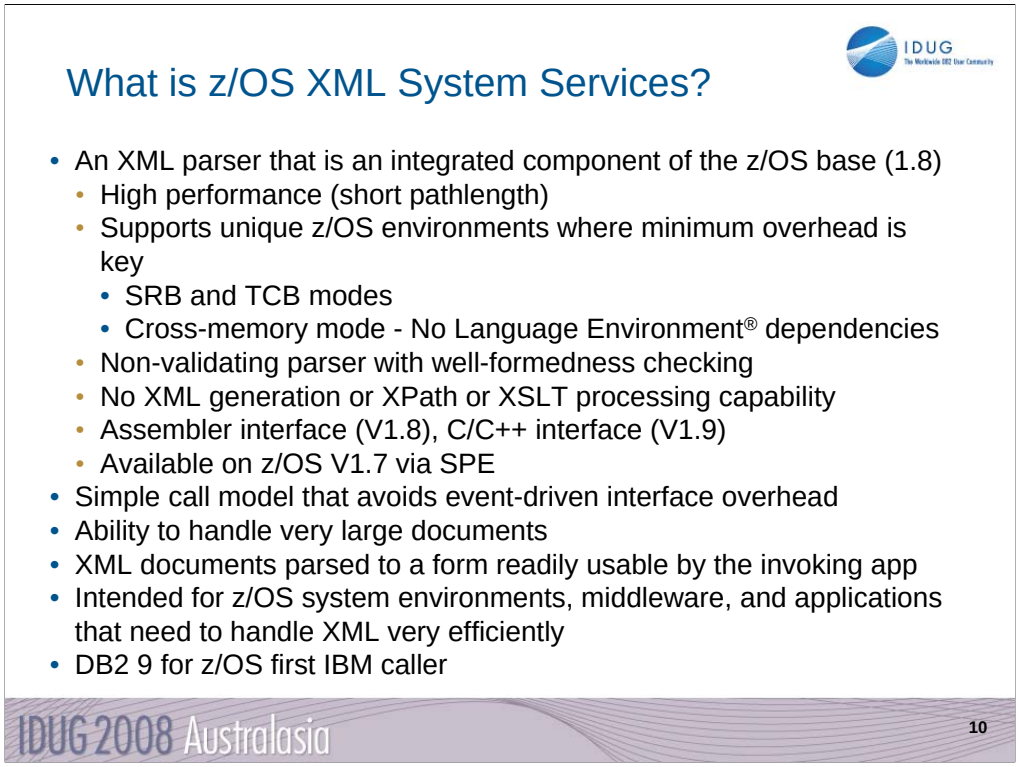

In z/OS V1.8, IBM introduced a new system component of z/OS, z/OS XML System Services (z/OS XML), a system-level XML parser integrated with the base z/OS operating system and designed to deliver an optimized set of services for parsing XML documents (z/OS XML has also been made available on z/OS V1.7). The initial beneficiaries of this system component are middleware and applications requiring high performance non-validating XML parsing. z/OS XML may currently be accessed by an Assembler programming interface and one of the first exploiters, DB2 9 for z/OS, uses this Assembler interface for XML native support. IBM plans to add C/C++ support for z/OS XML with z/OS V1.9 (satisfying the Statement of Direction in Software Announcement 206-039, dated February 28, 2006).

Simple call model that avoids event driven interface overhead

z/OS XML System Services is a lower level interface – z/OS XML parser looks at the whole document and parses everything it can thus avoiding interactive overhead.

Other parsers on parse interactively thus causing some performance overhead.

Ability to handle very large documents

a lot of parses have difficulty handling large documents – they want to bring in the whole doc but cannot... z/OS XML has large input and output buffers and can process more/ larger docs... in the event the buffer fills z/OS XML can request more memory from the application

XML documents parsed to an output buffer in a form readily usable by the invoking application

The binary form is sharable within z/OS systems and provides possible performance improvement by avoiding relocation overhead. (ie you move the XML parsing from one address space to another, no need to relocate the code)

Intended for z/OS system environments, middleware, and applications that need to handle XML very efficiently

Well- formedness checking – checks the syntax of the document (punctuation, brackets, spaces)

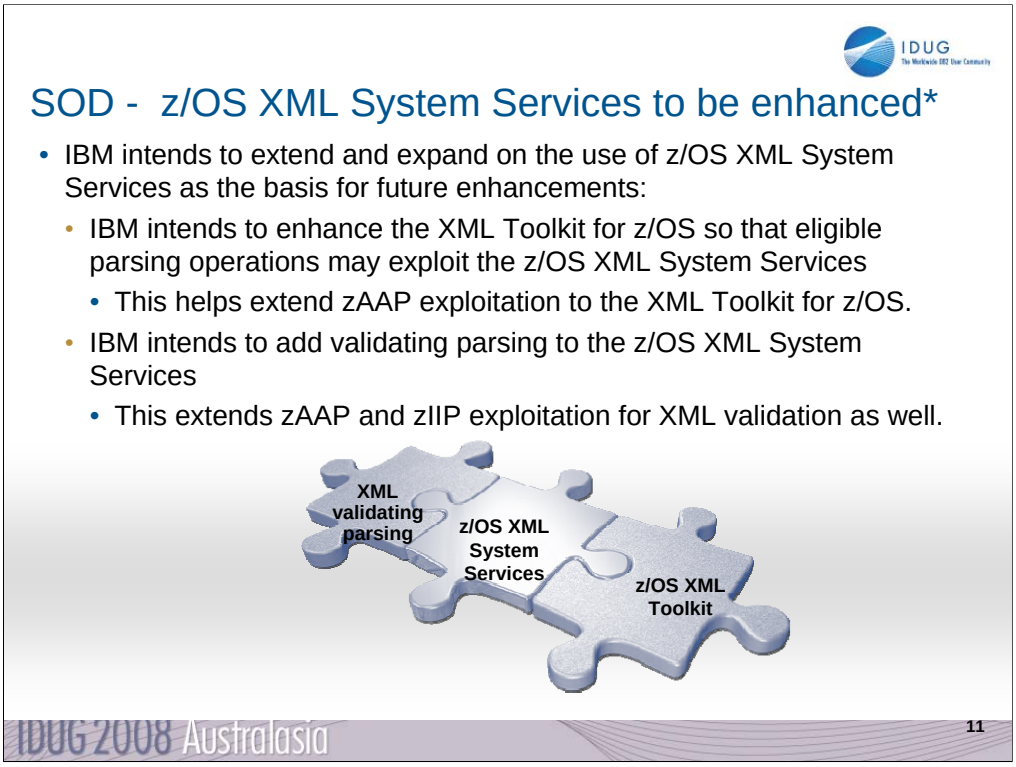

But wait, there is more ... As part of a comprehensive plan, IBM intends to extend and expand on the use of z/OS XML System Services (enabled for zAAP specialty processors) as the basis for future enhancements:

IBM intends to enhance the XML Toolkit for z/OS so that eligible parsing operations may exploit the z/OS XML System Services

This helps extend zAAP exploitation to the XML Toolkit for z/OS.

Please note, not ALL XML Toolkit workloads will be eligible for zAAP.

IBM intends to add validating parsing to the z/OS XML System Services

This extends zAAP exploitation for XML validation as well.

Therefore, in the future both XML parsing and XML validation are planned to be eligible for the zAAP.

Remember, the previous parts of the SOD stated that only z/OS XML parsing is eligible for zIIP and zAAP.

Future Directions – *Extending Encryption to IBM TotalStorage*

- Statement of Direction: To address customers' growing concern with data security, IBM is announcing a statement of direction for the development, enhancement and support of encryption capabilities within storage environments such that the capability does not require the use of host server resources.
- This includes the intent to offer, among other things, capabilities for products within the IBM TotalStorage® portfolio to support outboard encryption and to leverage the centralized key management functions planned for z/OS ICSF.

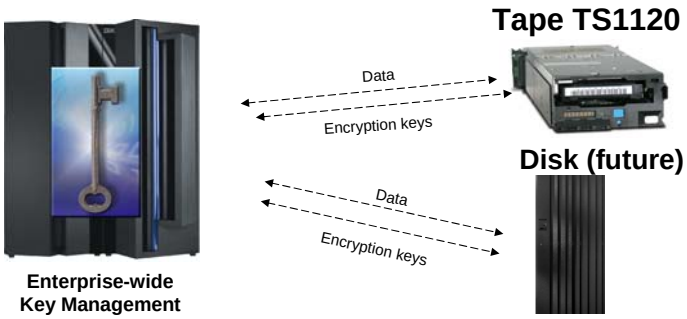

Statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only

IDUG 2008 Australasia 12

Statement of Direction: To address customers' growing concern with data security, IBM is announcing a statement of direction for the development, enhancement and support of encryption capabilities within storage environments such that the capability does not require the use of host server resources.

This includes the intent to offer, among other things, capabilities for products within the IBM TotalStorage® portfolio to support outboard encryption and to leverage the centralized key management functions planned for z/OS ICSF.

The first change comes in the TS1120 tape drive, but the rest is beyond currently announced products, including DB2 9.

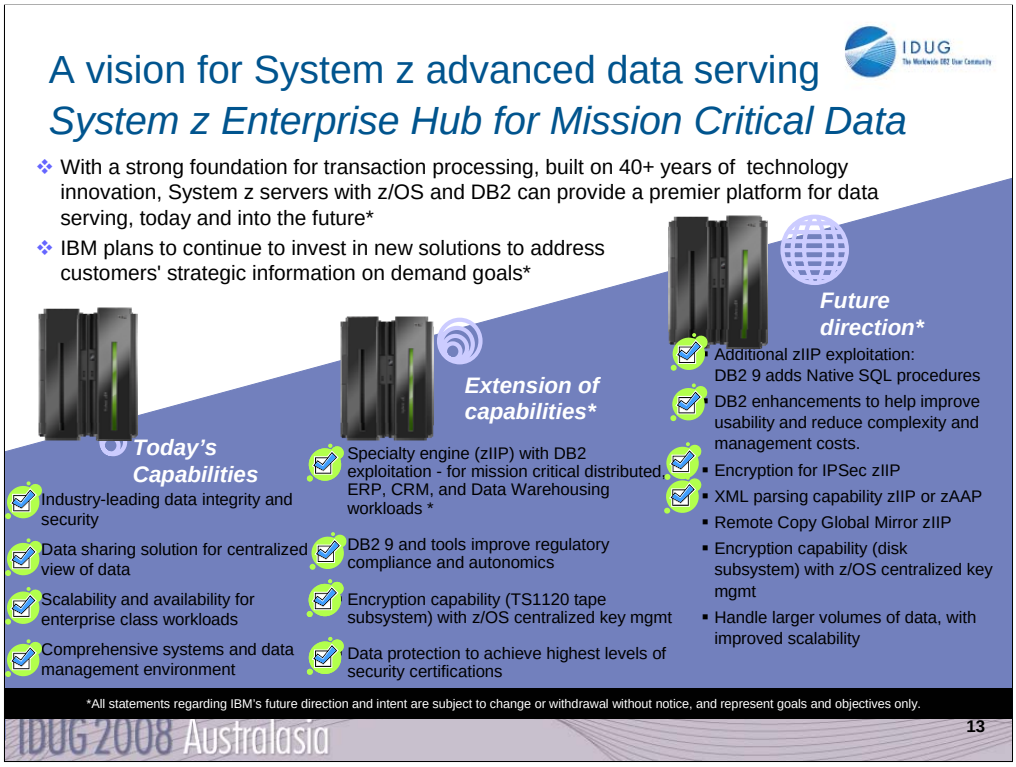

We have been working on specialty processors for a long time, and have delivered capabilities with the zIIP. The improved cost of ownership can help a lot for some customers, but not at all for others. There have been several big recent changes:

The z9 Business Class and Enterprise Class extend zIIP capabilities to many more customers. Only the largest customers needed the z9-109 processors, and the upgrade steps were very large ones. The z9 Business Class and Enterprise Class processors have a much greater range of processing power with more granular upgrade options. The entry level z9 processor now starts at under $100,000.

Query work was broadened beyond just star joins to all large parallel queries. If you have a warehouse that uses parallel processing and significant CPU time, then the zIIP may provide a benefit.

The TS1120 tape subsystem has added encryption capability with several options for centralized key management.

DB2 9 for z/OS adds three check marks. Another comes in August 2007, and some delivers with z/OS 1.9, but statements of direction note other improvements to come.

The Data Serving Roadmap provides more information about this slide: http://www.ibm.com/systems/z/ziip/data.html

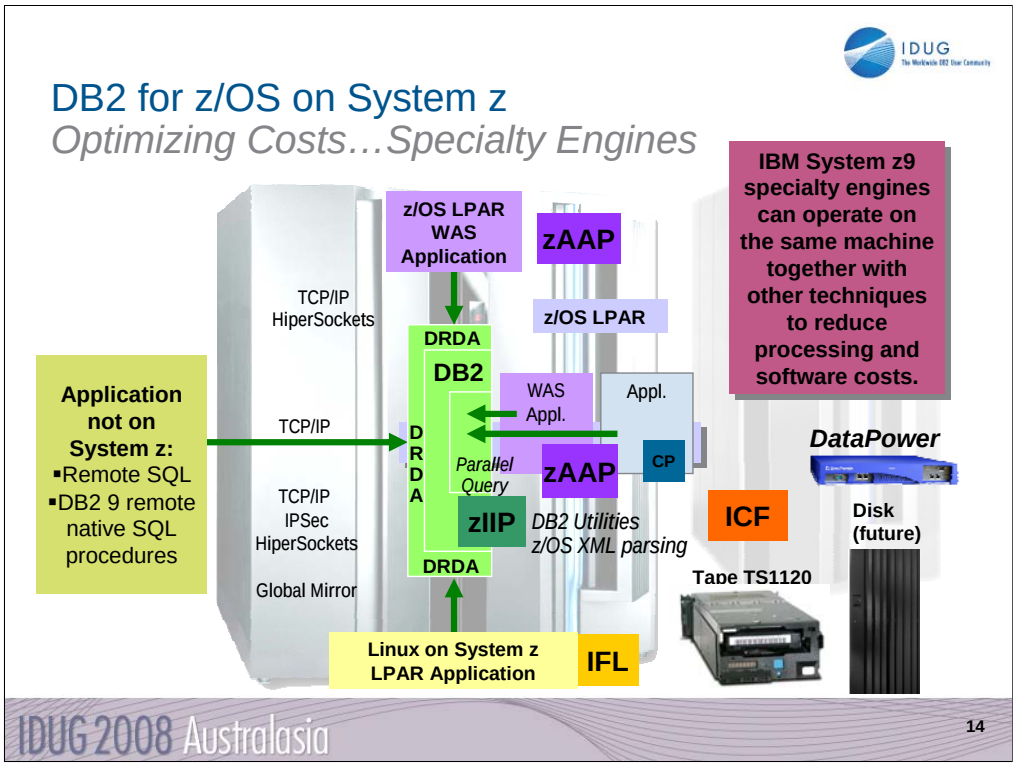

The specialty engines can be used to improve the cost of ownership, providing a low price for the hardware and not incurring software charges, since they are not general purpose processors. Customers can use all of the engines together with DB2. The ICF provides the Coupling Facility for DB2 data sharing with Parallel Sysplex for availability and scalability. The IFL can run Linux applications using DB2 Connect over a communication link or hipersockets to DB2 for z/OS. The zAAP can run Java applications, while the zIIP runs part of the DB2 work.

Announcements for zIIP:

http://www.ibm.com/common/ssi/rep_ca/7/897/ENUS106-287/ENUS106-287.PDF

http://www.ibm.com/common/ssi/rep_ca/3/897/ENUS106-293/ENUS106-293.PDF

http://www.ibm.com/systems/z/ziip/

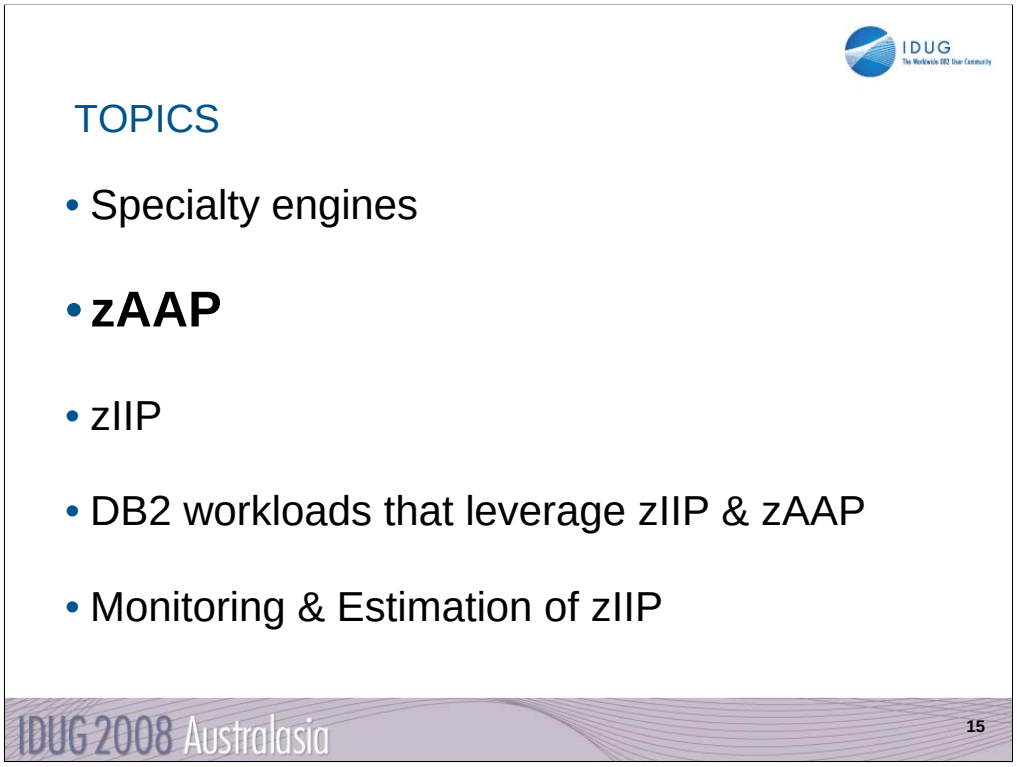

We'll start with an overview, then explain more about the work loads and what is eligible. Then we can work down to show some of the techniques for measurement and estimation, ending with a summary.

IBM System z Application Assist Processor (zAAP)

- Lower cost of computing for Java workloads on WebSphere Application Server for z/OS.
- Dedicated exclusively to Java Workloads
- Based on PR/SM™ LPAR technology, z/OS 1.6, SDK 1.4.1
- Executes Java Code with no changes to applications and no additional charges
- Lower price for zAAP than standard engine
- Up to (1) zAAP per general purpose processor
- IBM traditional System z software charges unaffected

See the web references for much more information.

Leveraging zAAPs for e-business Integration and Infrastructure Simplification

- zAAPs can help consolidate, simplify and reduce server infrastructure and improve operational efficiencies
  - **Enables strategic integration of e-business applications with mission-critical database workloads**
  - **Potential operational advantages over distributed multi-tier solutions**
- Eliminates separate tier to handle application server workload
  - **Remove one hardware tier**
  - **Remove one TCP/IP link**
- Leverage core System z strengths and manage Java Workloads automatically with z/OS
  - **System z Security**
  - **System z Workload Manager (WLM)**
  - **System z Availability**
  - **System z Scalability**
  - **System z Flexibility**

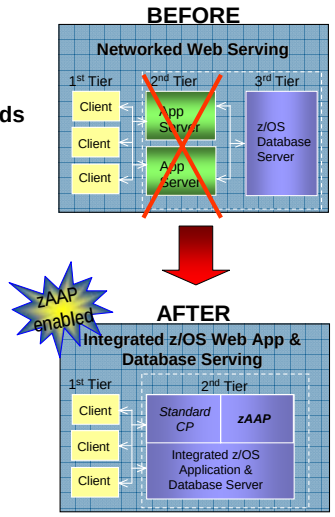

Provides a cost-competitive single tier integrated CONTAINER FOR:

Application deployment (e.g., WAS Java development platform)

Direct, integrated DB access within the same logical partition w/o all the overhead implied by a TCP/IP communications stack

Security & transaction contexts, Network stack

Provides all of the System z and z/OS QoS within the same logical partition

Takes advantage to the z990 multi-book scaling capabilities

Builds upon the z/OS increased processor scaling capabilities in z/OS 1.6

> 16 processors within a logical partition

Introduces a special purpose engine built upon existing System z specialized engine technologies

ICFs, IFLs, SAPs

Why zAAP? A Simplified Example

Consider a WebSphere Application that is transactional in nature and requires 1000 MIPS today on System z.

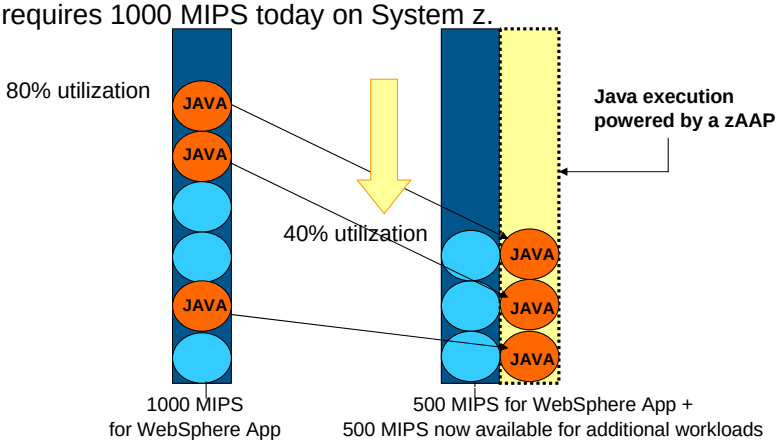

In this example, with zAAP, we can reduce the standard CP capacity requirement for the application to 500 MIPS or a 50% reduction.* * For illustrative purposes only

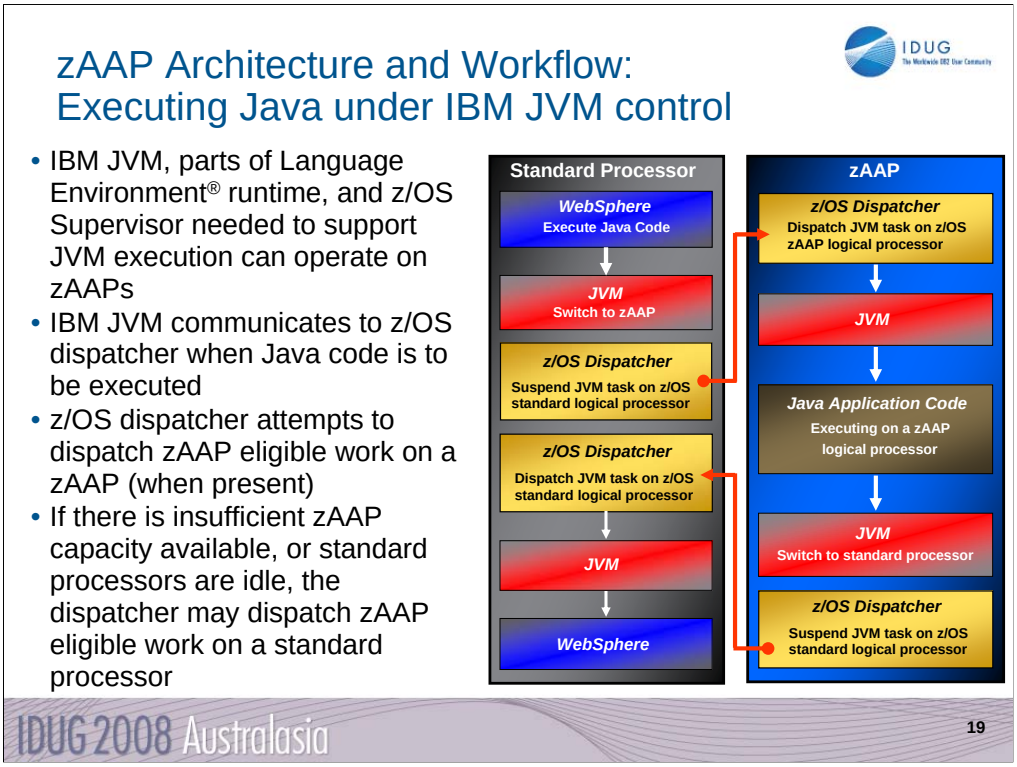

IBM JVM, parts of Language Environment® runtime, and z/OS Supervisor needed to support JVM execution can operate on zAAPs

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)

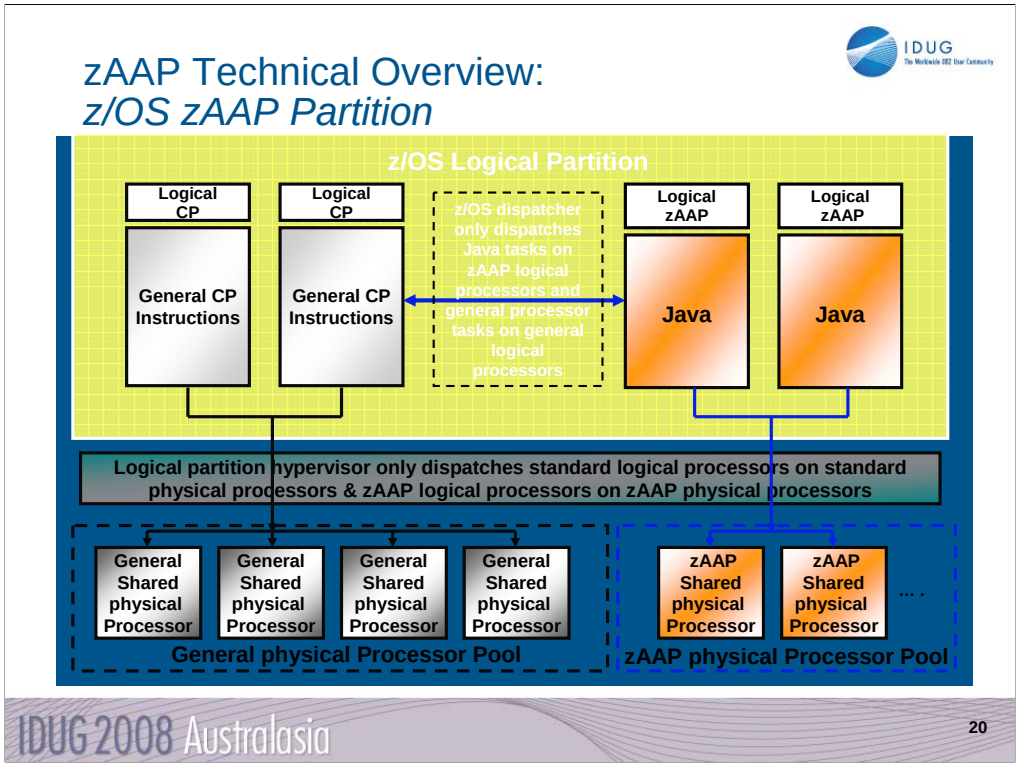

zAAP Technical Overview

When a z/OS logical partition is configured, both CPs and zAAPs are defined as necessary to support the planned Java and non-Java workloads. zAAPs may be configured as initially online or reserved for subsequent use by z/OS as necessary. Since zAAPs cannot be IPLed at least one central processor is required for each z/OS partition.

zAAPs may be defined as either shared by other logical partitions or dedicated to a specific partition; however, both the CPs and zAAPs for each partition will have the same shared or dedicated attribute. For a given partition, you cannot define shared central processors and dedicated zAAPs or vice versa. PR/SM configures shared zAAPs from the same pool of shared special purpose processors as the Internal Coupling Facility (ICF) and Integrated Facility for Linux (IFL). Collectively, all shared ICFs, IFLs and zAAPs will also dynamically share in the processing requirements for all three special purpose processor types as controlled by PR/SM.

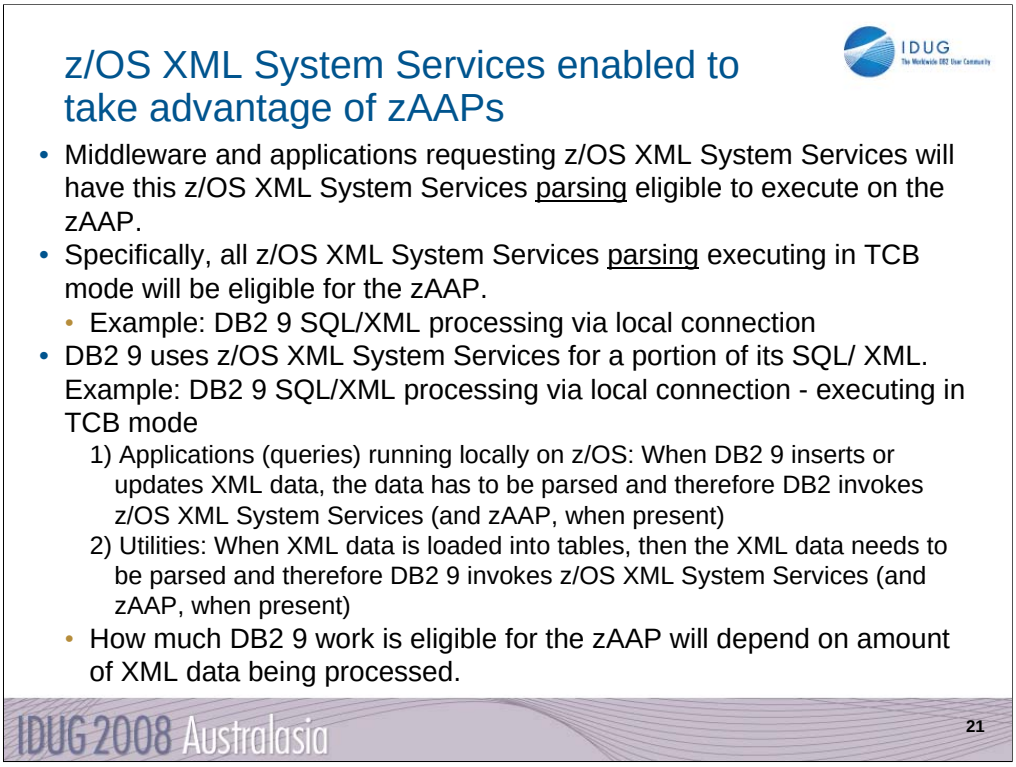

Middleware and applications requesting z/OS XML System Services will have this z/OS XML System Services parsing eligible to execute on the zAAP. Only the XML parsing being performed by z/OS XML System Services (a base element of z/OS) is eligible for zAAP.

Please note, there is a Java-based XML parser in the IBM SDK 1.3 (and above) – the technology is called XML4J. This Java-based XML parser is already fully eligible for zAAP because it is Java workload. Other C++, COBOL, PL/I and roll-your own parsers are not eligible for zAAP.

Specifically, all z/OS XML System Services parsing executing in TCB mode will be eligible for the zAAP. Example: DB2 9 SQL/XML processing via local connection

Currently DB2 9 is the only IBM exploiter of z/OS XML System Services (via an Assembler interface)

DB2 9 utilizes z/OS XML System Services for a portion of its SQL/ XML.

Please note, not all DB2 9 XML processing is done using z/OS XML System Services. XML Validation is not eligible.

Example: DB2 9 SQL/XML processing via local connection. When executing locally, DB2 9 is already executing in TCB mode

Types of DB2 9 XML parsing workloads that are eligible for zAAP would be

1) Applications (queries) running locally on z/OS

When DB2 9 inserts or updates XML data, the data has to be parsed and therefore DB2 invokes z/OS XML System Services (and zAAP, when present)

2) Utilities

When XML data is loaded into tables, then the XML data needs to be parsed and therefore DB2 9 invokes z/OS XML System Services (and zAAP, when present)

How much DB2 9 work is eligible for the zAAP will depend on amount of XML data being processed.
As a reminder.. Please note, XML validation not eligible for zAAP

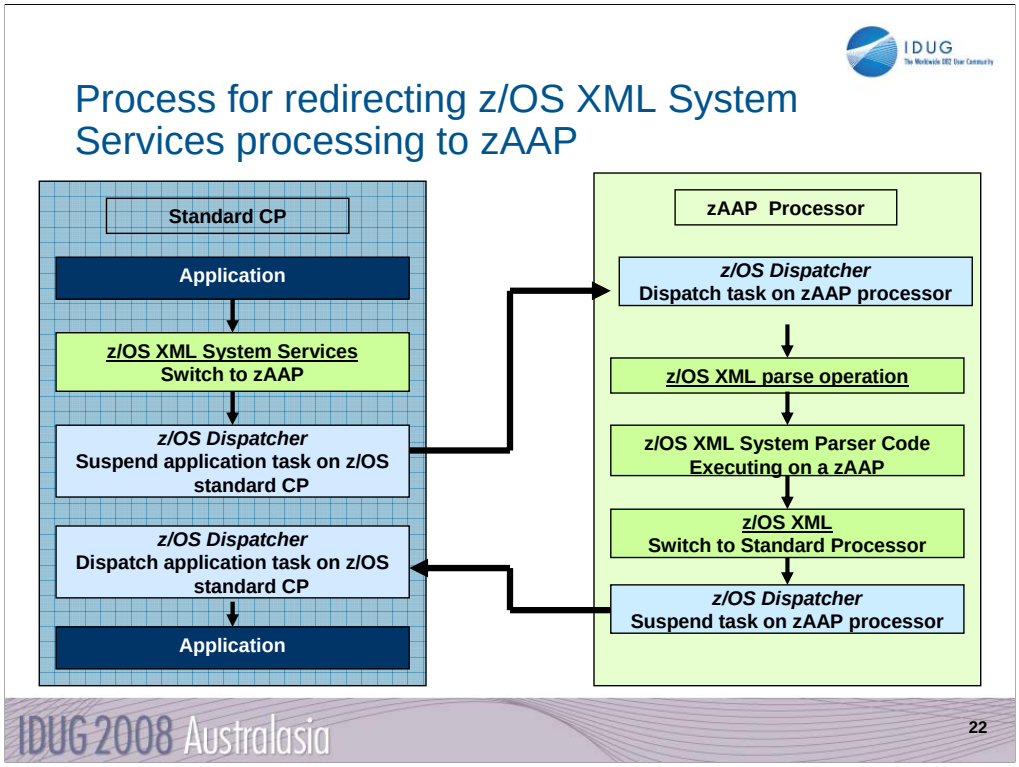

Essentially, z/OS XML System Services workload is dispatched to the zAAP the same way that Java workload is dispatched to the zAAP.

How Do I know if zAAP is Right for My Workloads?

- *zAAP Projection Tool for Java 2 Technology Edition, SDK 1.3.1*
  - **Available with Microsoft® Excel® Summary Workbook**
  - **Runs in test environment**
  - **Gathers usage information on % of Java in your workloads that could execute on zAAP**
  - **Useful in predicting number of zAAPs necessary for optimum configuration**
- *z/OS Performance: Capacity Planning Considerations for zAAP* White Paper
  - **Describes the zAAP Projection Tool**
  - **Describes the prototype measurements**
  - **Describes Capacity Planning Methodology**
- Size 390
  - **Provides special assistance for the sizing methodology described in white paper**
  - **May help with sizing consolidation of distributed Java workloads onto System z and zAAP(s)**

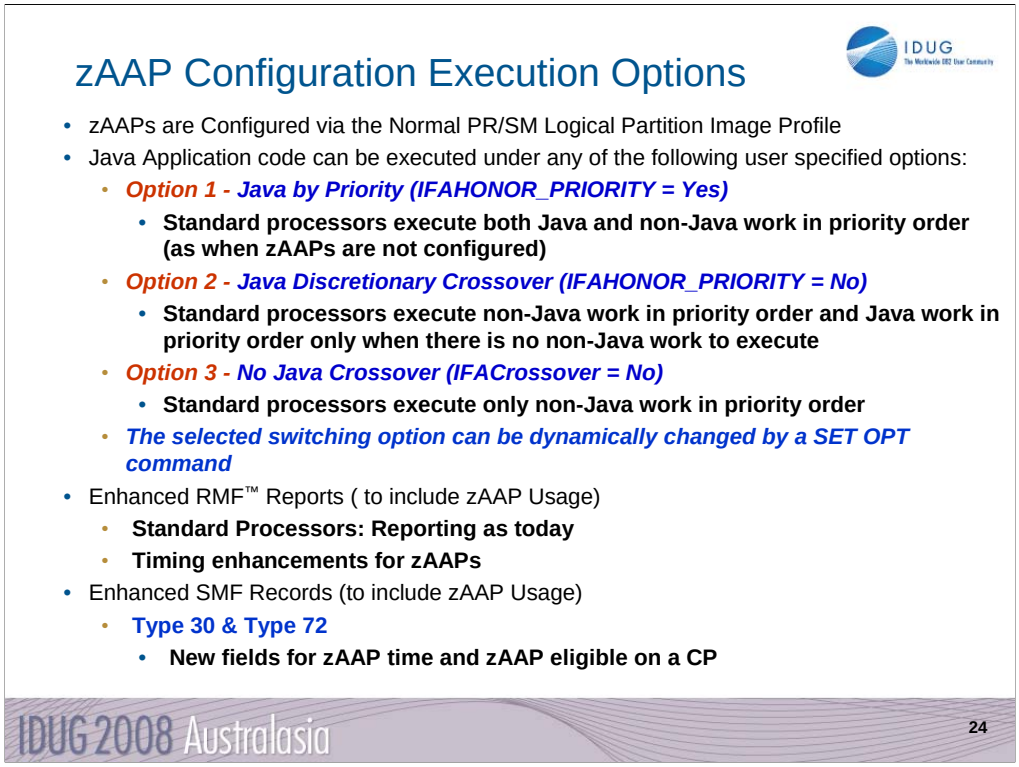

zAAPs Are Configured via the Normal PR/SM Logical Partition Image Profile. *zAAPs may be defined as shared or dedicated processors. Shared zAAPs may be PR/SM weighted and capped as per standard processors. The existing single set of PR/SM logical partition processor weights (INITIAL, MIN, MAX) are applied independently to the shared standard processors and to the shared zAAP processors configured to the logical partition*

z/OS WLM will manage standard shared CPs as today but not shared zAAPs

Java application code can be executed under any of the following user specified options

Option 1 - Java by Priority (HONOR_PRIORITY = Yes)

Standard processors execute both Java and non-Java work in priority order (as when JAFs are not configured)

JAFs execute only Java work in priority order

Java work reduced or stopped on standard processors when maximum MSUs would be exceeded

This is the z/OS default option

Option 2 - Java Discretionary Crossover (HONOR_PRIORITY = No)

Standard processors execute non-Java work in priority order and Java work, in priority order, only when there is no non-Java work to execute

JAFs execute only Java work in priority order

Might be selected when Java work would consume too many standard processor cycles due to its priority

Option 3 - No Java Crossover (CROSSOVER= No):

Standard processors execute only non-Java work in priority order

JAFs execute only Java work in priority order

For customers that don't want Java executed on more standard processors (e.g., due to a SCRT software pricing license)

Processor crossover only if last JAF fails

zAAP Summary

- Challenges:
  - **Strategic Web-based application exploitation increasing at exponential rates**
  - **Much of this technology is driven by Java technology-based applications**
  - **Java technology-based applications require up to 2-3X resources over traditional applications**
  - **Achieve e-business requirements for resilient, security-rich, simplified and cost-effective on demand computing.**
- zAAP …an industry first
  - **Only specialized processing units for Java Code today**
  - **Supported by IBM Middleware such as WebSphere, DB2…**
  - **Helps reduce demands on general purpose processors – makes them available for other work**
- zAAPs for e-business Integration and Infrastructure Simplification
  - **Integrate Java technology-based applications with mission-critical data**
  - **Helps reduce infrastructure complexity for multi-tier applications**
- zAAPs Provide Investment Flexibility
  - **Extend the value of existing System z investments and lower total cost of ownership**
  - **Cost-effective, specialized Java execution environment**
  - **Low Total Cost of Acquisition ($125K USD per zAAP)**
  - **Helps reduce Total Cost of Ownership (software and maintenance savings)**

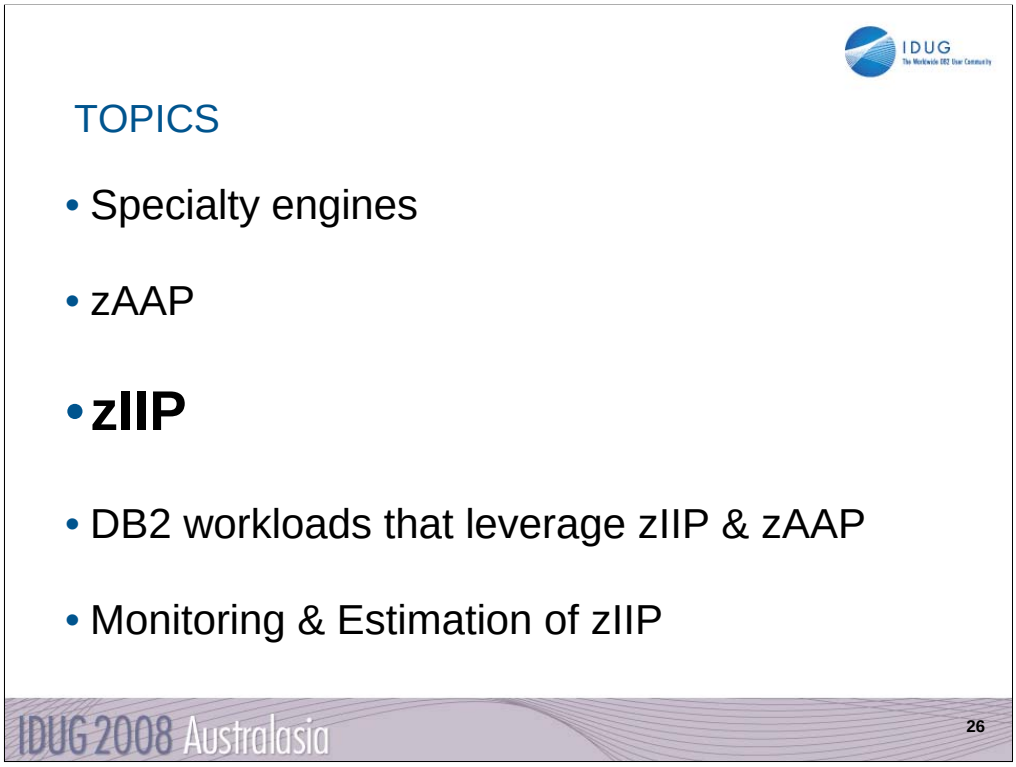

We'll start with an overview, then explain more about the work loads and what is eligible. Then we can work down to show some of the techniques for measurement and estimation, ending with a summary in the next three sections.

Mainframe Innovation: zIIP Specialty Engines

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

IBM System z9 Integrated Information Processor (IBM zIIP) 2006

Eligible for zIIP:

- DB2 remote SQL access, utilities, parallel query, remote native SQL procedures
- ISVs
- IPSec encryption
- z/OS XML
- Z/OS Global Mirror XRC SDM

One of the latest changes to deliver is the ability to use the zIIP for IPSec encryption in 2007. A statement of direction was provided for z/OS XML to be able to use zAAP or zIIP for XML parsing in April and August 2007, and that is partly delivered.

Example for zIIP data serving workload SAP solutions

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

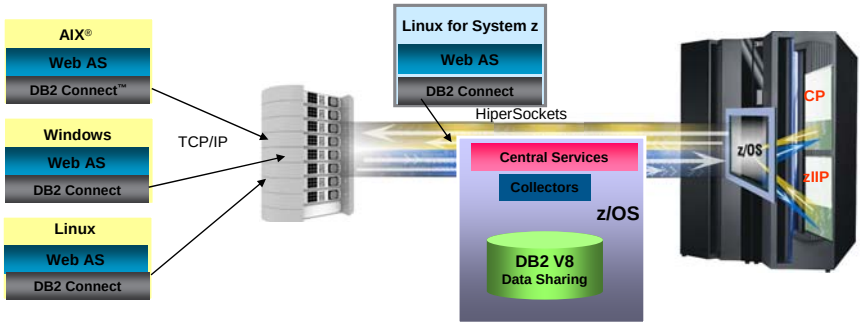

May enable growth of System z9 SAP workloads through resource optimization

The zIIP is for customers who are concerned about costs for growth. The big cost reduction is in hardware cost, which is much less than a standard processor. The biggest cost reductions are in software, as IBM does not charge for software running on the specialty processors. The zIIP will fit some customers very well, but will not apply for all. As a specialty processor, not all work can use the zIIP, which will only process work running under an enclave SRB. Most applications cannot run in SRB mode. The specifics of the software charging need to be considered. Customers must be current on hardware (System z9), current on software (z/OS 1.6 or later, DB2 V8 or later) and have a work load peak using the types of work supported by the zIIP:

Large financial services company delighted with zIIP

Issues:

- **Mission critical data on DB2 for z/OS: scalable, available, secure, central management**
  - **... Pressure to reduce cost of DB2 for ERP application**
  - **... Data warehouse outside the mainframe, multiple copies of data, challenge security & management**

Solution:

- **System z9 Integrated Information Processors (zIIPs)**
  - **Simple installation No application code changes**
  - **No DB2 configuration changes**
  - **DB2 workload dispatched flawlessly**

Results:

    - **"The overhead of zIIP was monitored and appears to be insignificant"**
    - **"The response time improved with the available capacity of the zIIP"**
    - **"zIIP utilization reached 97% (avg over 15 min) & saturated the zIIP."**
- **System z9 utilization & software costs significantly reduced**
- **Investigating Data Warehousing on System z**

Mission critical data is on DB2 for z/OS. Data is highly scalable and available, is secure, and has centralized management, backup, and recovery, but there is pressure to reduce the cost of distributed database calls (DRDA).

An ERP application currently uses DB2 on z/OS for its data store.

Data warehouse performed outside of the mainframe, causing multiple copies of disparate data, potentially compromising security, management, backup & recovery.

Five (5) System z9 Integrated Information Processors (zIIPs) were purchased.

Simple installation of DB2 and z/OS maintenance

No code changes to the application

No configuration changes to DB2

zIIP activated automatically without any tuning requirements

DB2 workload was dispatched flawlessly

Customer quotes:

"The overhead of zIIP was monitored and appears to be insignificant"

"The response time improved with the available capacity of the zIIP"

"zIIP utilization reached 97% (avg over 15 min) & saturated the zIIP."

The System z9 utilization was significantly reduced as workload was redirected to the zIIP. Software [ISV] costs were significantly reduced.

Investigating Business Intelligence/ Data Warehousing on System z

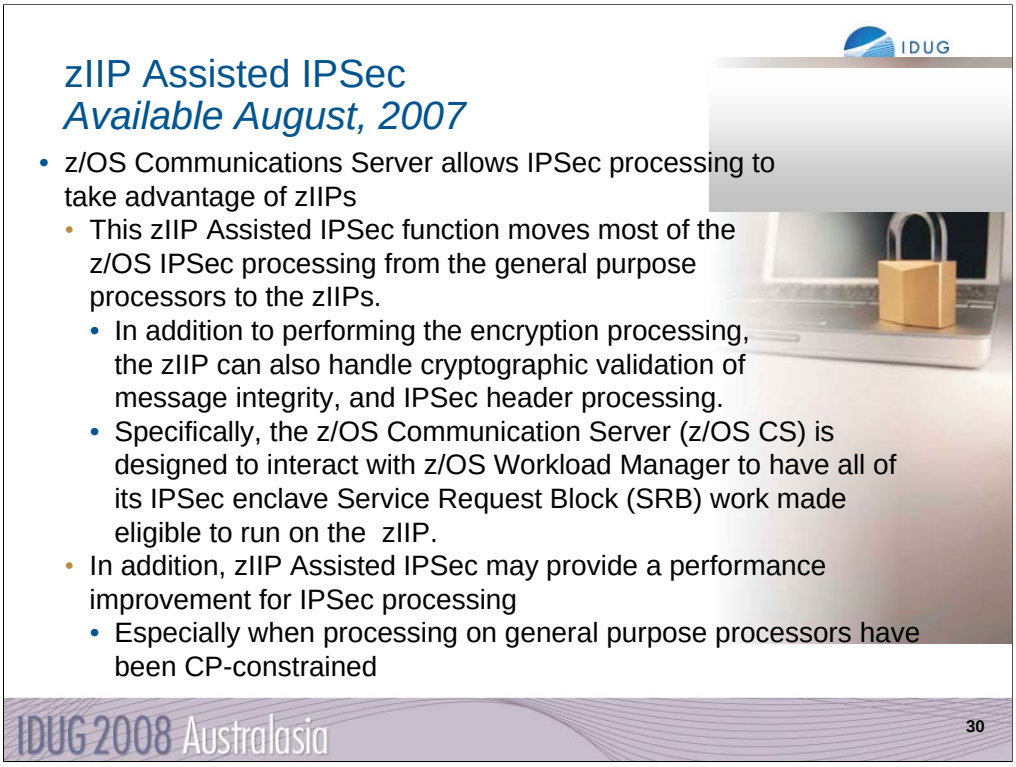

IBM previewed an enhancement to the z/OS Communications Server that allows the IPsec processing to take advantage of zIIPs. This IPSec zIIP Assist function moves most of the IPSec processing from the general purpose processors to the zIIPs. In addition to performing the encryption processing, the zIIP will also handle cryptographic validation of message integrity, and IPSec header processing. This capability is planned to be available on July 2007 via PTF for z/OS V1.8. Previewed is an enhancement to the z/OS Communications Server that allows the IPsec processing to take advantage of zIIPs. The zIIP, in effect, will be a high-speed encryption engine that is designed to provide better price performance. IPSec is an open networking standard used to create highly secure connections between two points in an enterprise - this may be server-to-server, or server to network device, as long as they support the IPSec standard. End-to-end encryption is deployed to provide a highly secure exchange of network traffic.

Business demands to protect sensitive data on the Internet have increased the requirement for users to implement end-to-end encryption on Virtual Private Networks (VPNs). At the same time as businesses are seeing an increased need for data protection, they are also increasing their use of more open networks such as the Internet. Moving secure business data and transactions from proprietary, trusted networks to more open networks creates an ever-increasing need for new technologies to protect this data.

The IPSec support was integrated into z/OS Communication Server in z/OS V1.7, and provides authentication, integrity, and data privacy from z/OS to other network endpoints that support IPSec. In addition to allowing you to run host based IPSec for secure end to end network flows, the V1.7 IPSec added IP filtering to protect your host. Since the IPSec support is implemented in the IP protocol layer, you can use it to protect a variety of network traffic types to/from any application without change. The new zIIP Assisted IPSec function that IBM is adding to z/OS V1.8 by PTF, moves most of the IPSec processing from processors to the zIIPs. In addition to performing the encryption processing, the zIIP will also handle cryptographic validation of message integrity, and IPSec header processing. This is designed to allow you to take advantage of the cost saving benefits of the zIIP when you implement IPSec to secure your valuable business transactions and to protect your host.

z/OS network encryption overview

- The z/OS Communications Server (z/OS CS) encrypts application data a number of ways:
  - #1 Application layer encryption (per session)
    - Application is coded with encryption
  - #2 Application is encrypted in network layer (also per session)
    - "common service" AT-TLS (z/OS 1.7)
    - Transparent to the application
  - #3 "Platform to platform" encryption (Virtual Private Networks using IPSec)
    - All traffic may be encrypted – transparent to all applications
- When do you use one form of encryption versus another?
  - Depends on client, application, topology, performance requirements.....
  - IPSec can be used for some or all traffic – can create a VPN

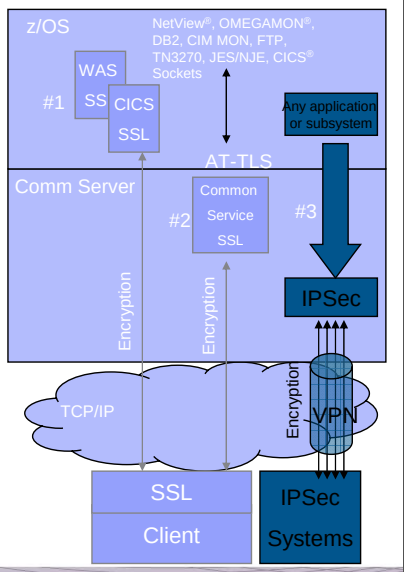

Which encryption technology one would use will depend on your requirements. All traffic can be encrypted, including the DB2 session. If the cost is key and zIIP reduces your costs, then this provides an alternative to DRDA encryption.

This is a background on encryption within the z/OS Communication Server (a component of z/OS). z/OS Comm Server can provide encryption in a number of ways. We are focusing on zIIP Assisted IPsec with is #3 VPN technology and the dark blue column on the right hand side.

Again – going forward we are focusing on those who are interested in end-to-end encryption via IPSec.

Please note, one thing that could be confusing about z/OS V1R7.0 Communications Server IPSec support is that it has been packaged together with IP filtering support and is referred to as integrated IP Security. That is because there is a very close affinity between IPSec and IP filtering in the z/OS Communications Server; while you can implement IP filtering without IPSec, you cannot implement IPSec without IP filtering

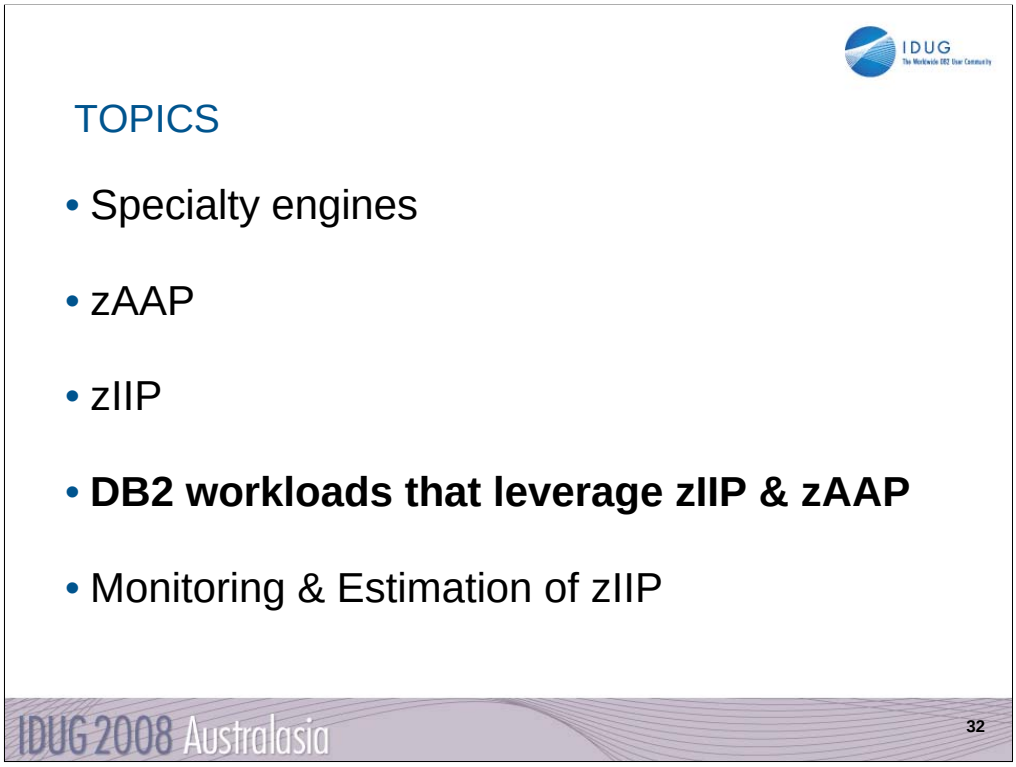

We'll start with an overview, then explain more about the work loads and what is eligible. Then we can work down to show some of the techniques for measurement and estimation, ending with a summary.

DB2 & IBM zIIP add value to database work

Portions of the following DB2 for z/OS V8 and DB2 9 workloads may benefit from zIIP (DB2 9 in blue)*:

1 - **ERP, CRM, Business Intelligence or other enterprise applications**
- **Via DRDA over a TCP/IP connection (enclave SRBs)**
- **DB2 9 for z/OS Remote native SQL procedures**
- **DB2 9 XML parsing via DRDA to fully utilize zIIP**

2 - Data warehousing applications
- **Requests that use parallel queries, including star schema**
  - **DB2 9 higher percentage of parallel queries zIIP eligible**

3 - DB2 Utilities LOAD, REORG & REBUILD
- **DB2 utility functions used to maintain index structure**

*** zIIP allows a program working with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to zIIP. Above types of DB2 work are those running in enclave SRBs, of which portions can be sent to zIIP.**

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 above types of DB2 V8 work are those executing in enclave SRBs, of which portions can be sent to the zIIP. Not all of this work will be run on zIIP. z/OS will direct the work between the general processor and the zIIP. The zIIP is designed so a software program can work with z/OS to dispatch workloads to the zIIP with no anticipated changes to the application – only changes in z/OS and DB2.

IBM DB2 for z/OS version 8 was the first IBM software able to take advantage of the zIIP. Initially, the following workloads can benefit:

• SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP), CRM (Siebel), or business intelligence and are expected to provide the primary benefit to customers. Stored procedures and UDFs run under TCBs, so they are not generally eligible, except for the call, commit and result set processing. DB2 9 remote native SQL Procedure Language is eligible for zIIP processing. BI application query processing utilizing DB2 parallel query capabilities; and functions of specified DB2 utilities that perform index maintenance.

•For more, see http://www.ibm.com/systems/z/ziip/

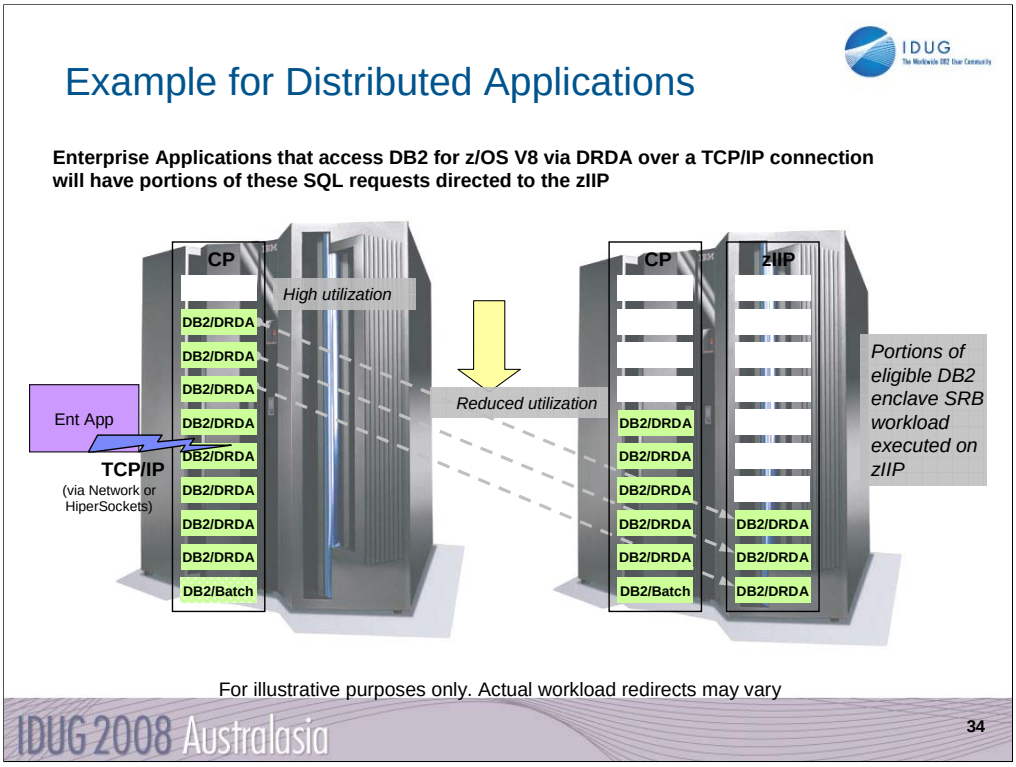

Remote SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP or PeopleSoft), CRM (Siebel), and business intelligence running on other platforms. Remote SQL is expected to provide the primary benefits to customers, as it is commonly part of the peak load. Stored procedures and UDFs run under TCBs, so they are not generally eligible for zIIP, except for the call, commit and result set processing. DB2 9 remote native SQL Procedure Language is eligible for zIIP processing.

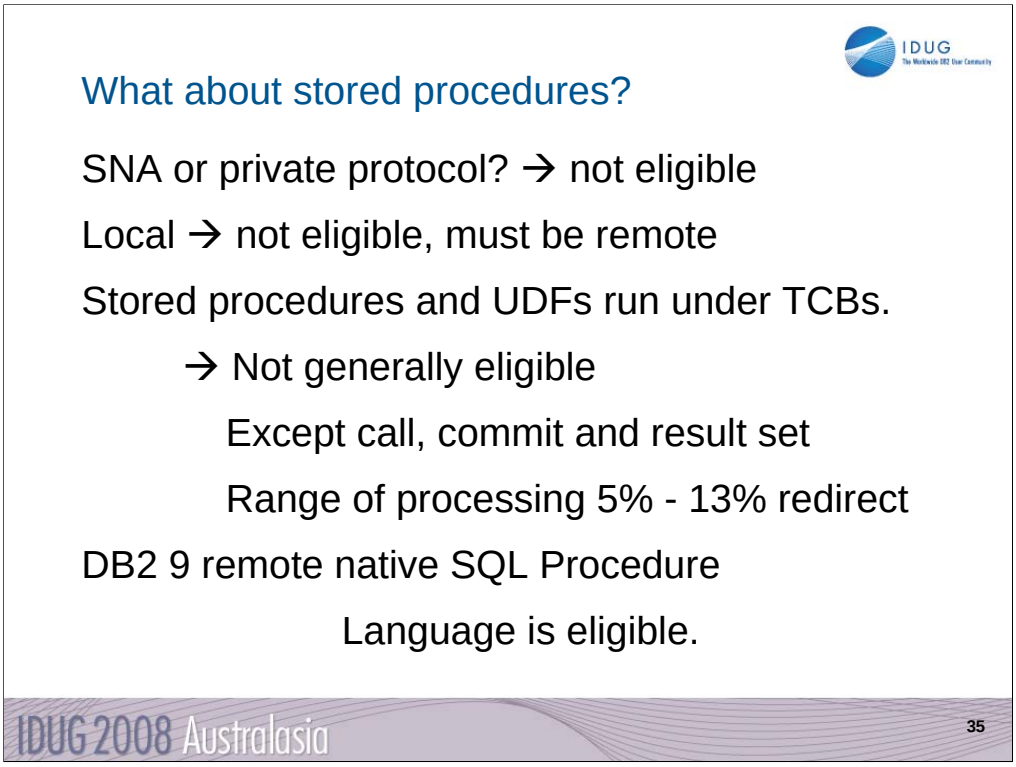

Remote SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP or PeopleSoft), CRM (Siebel), and business intelligence running on other platforms. Remote SQL is expected to provide the primary benefits to customers, as it is commonly part of the peak load. Stored procedures and UDFs run under TCBs, so they are not generally eligible for zIIP, except for the call, commit and result set processing. In a laboratory measurement, we found 13% of the stored procedure time was eligible for zIIP, but the conditions for your procedure will probably differ. DB2 9 remote native SQL

Procedure Language is eligible for zIIP processing.

Example for Business Intelligence Applications

Parallel queries via DRDA over a TCP/IP connection will have portions of this work directed to the zIIP. Large local parallel can also have portions of work directed to zIIP.

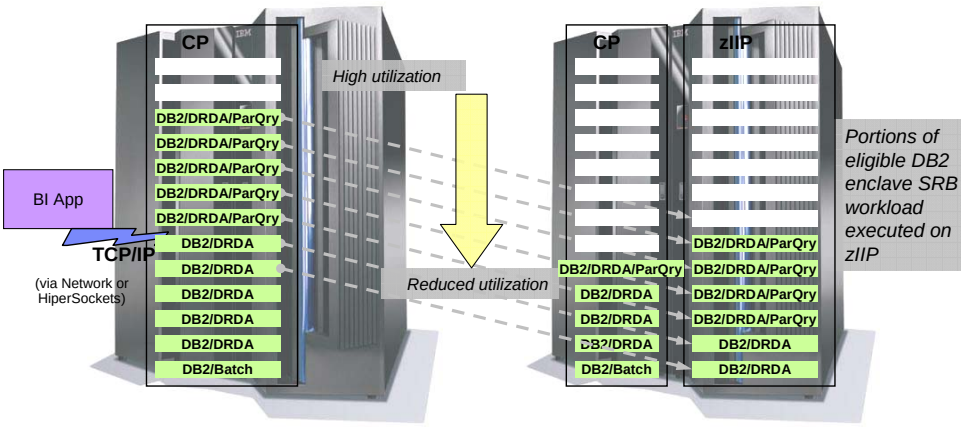

For illustrative purposes only. Actual workload redirects may vary depending on how long the queries run, how much parallelism is used, and the number of zIIPs and CPs employed

Parallel queries: If the work comes in remotely over DRDA using TCP/IP, then the initial work is eligible as remote work. After the initial time, the parallel processing threads are eligible and can use more processing on the zIIP.

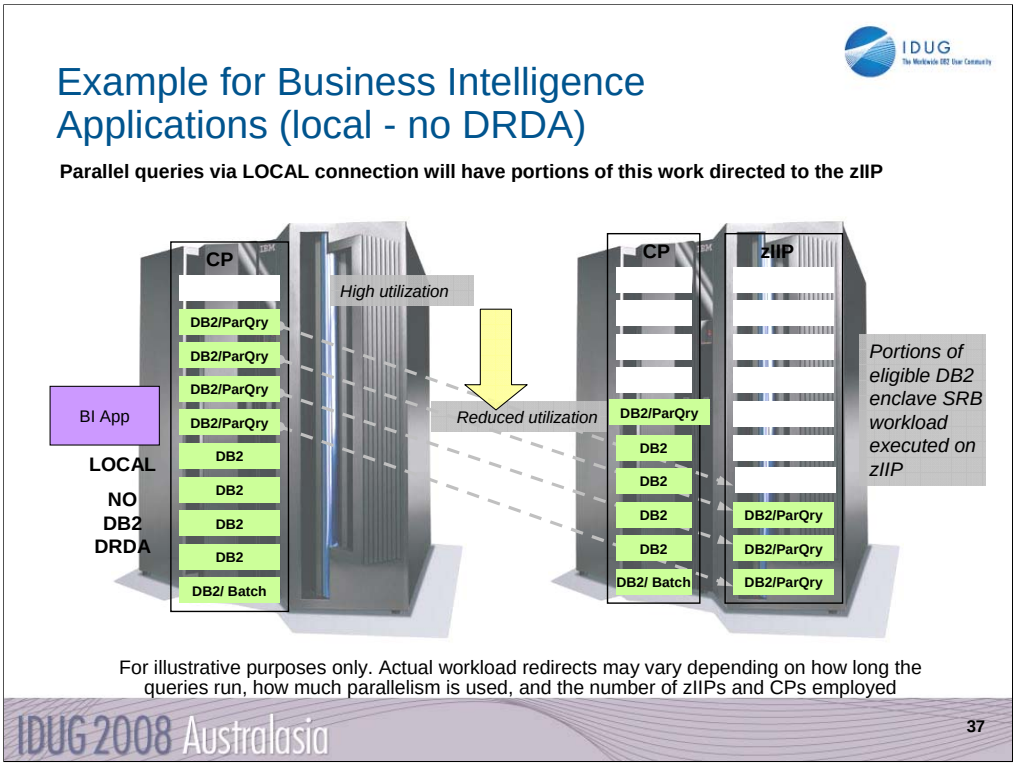

Parallel queries: If the work comes in remotely over DRDA using TCP/IP, then the initial work is eligible as remote work. After the initial time, the parallel processing threads are eligible and can use more processing on the zIIP.

Much more work can be processed in parallel on DB2 9, due to a fundamental change in optimization. If you want to use more parallel processing, then the two keys are

1. Provide adequate partitions for parallel processing
2. Move to DB2 9

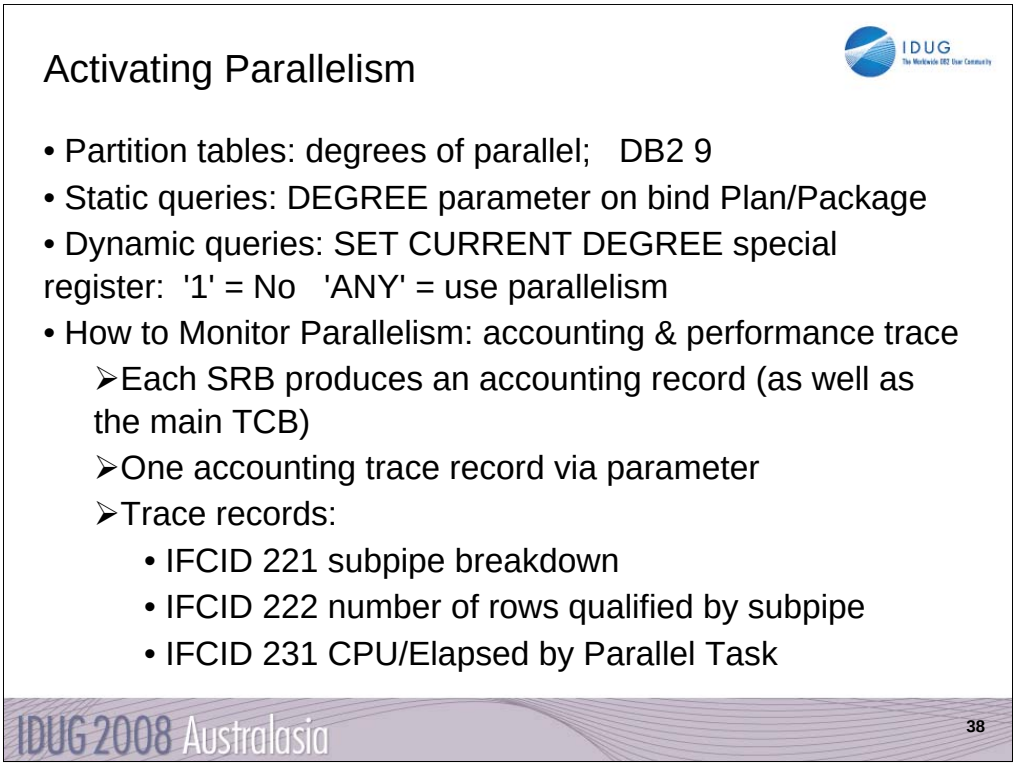

How to Activate Parallelism
• Static queries: DEGREE parameter on bind Plan/Package
• Dynamic queries: SET CURRENT DEGREE special register
– '1' -- DB2 will not consider parallelism for queries
– 'ANY' - DB2 will use parallelism for queries where possible
• DSNZPARMS
– DEFAULT CURRENT DEGREE in INSTALL panel DSNTIP4 (CDSSRDEF = ANY)
– Default CURRENT DEGREE for dynamic queries (no effect on static queries)
• To avoid query regression:
– Parallelism can be controlled with 2 ZPARMs
– PARMMDEG for the maximum degree of parallelism in DSNTIP4 (not often needed)
– Ensures single query does not consume all parallel tasks
– Minimizes runtime regression if large number of resources are not available
• SPRMPTH: Threshold to disable parallelism for short running queries

How to Monitor Parallelism

• Each SRB produces an accounting record (as well as the main TCB)
• You can tell DB2 to roll-up information into one accounting trace record via ZPARM: PTASKROL=YES
• IFCID 221 gives subpipe breakdown
• IFCID 222 gives # rows qualified by subpipe
• IFCID 231 gives CPU/Elapsed by Parallel Task

Example for DB2 V8 z/OS utilities

DB2 for z/OS utilities used to maintain index structures

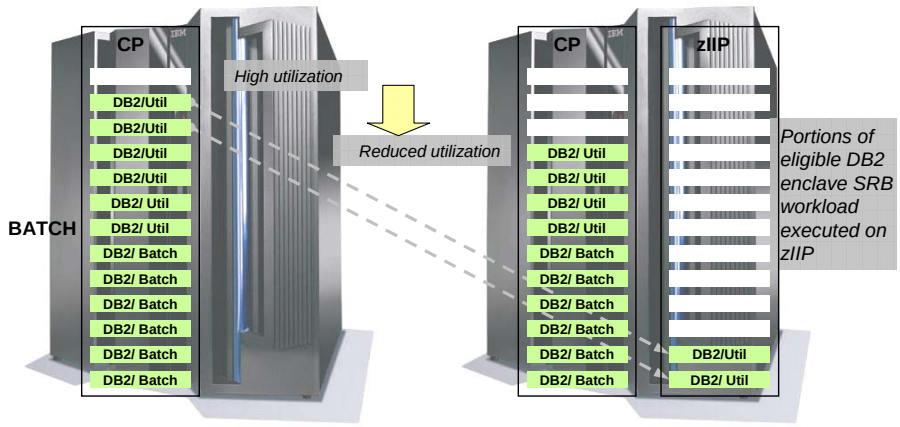

For illustrative purposes only, actual workload redirects may vary. Only the portion of the DB2 utilities used to maintain index structures (within LOAD, REORG, and REBUILD) is redirected.

DB2 utility index processing: Functions of the LOAD, REORG and REBUILD DB2 utilities that perform index maintenance are eligible for zIIP. This is not a common peak capacity constraint, but could be useful in reducing CPU charges.

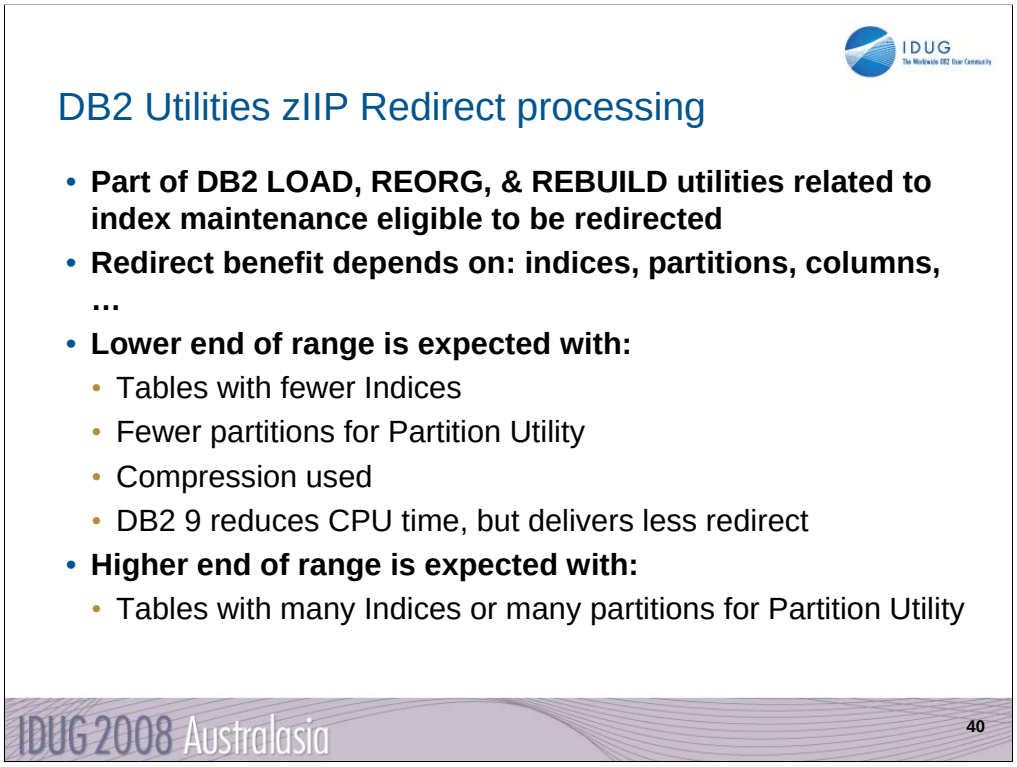

Portions of DB2 Utilities (LOAD, REORG, & REBUILD) processing related to Index maintenance are eligible to be redirected.

Redirect benefit depends on:

How many Indices are defined on the Table

How many Partitions are in the Table for Partition Utility

Number of Columns, Column data type etc.

Use of data compression

Lower end of range is expected with:

Tables with fewer Indices

Fewer partitions for Partition Utility

Compression used (more CPU overall, lower percentage redirect)

Higher end of range is expected with:

Tables with many Indices or many partitions for Partition Utility

Some instances where zIIP would not be used – local application not running parallel query

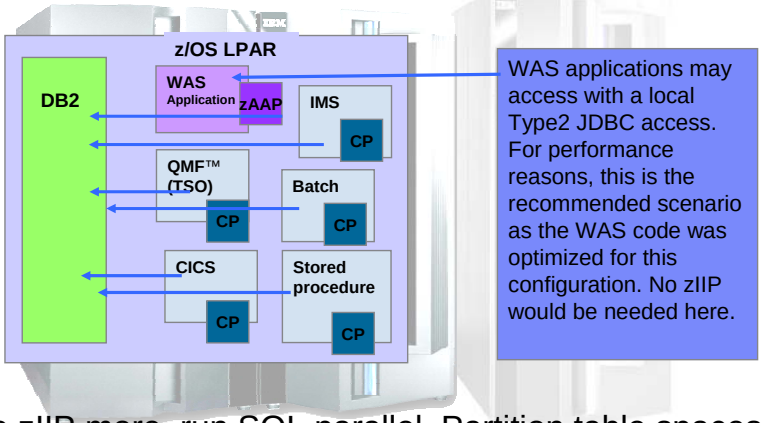

To use zIIP more, run SQL parallel, Partition table spaces, Change to DEGREE(ANY), Add join predicates to indexes for efficient cutting

If the work is local and running under a TCB, as normal applications do, then the zIIP would not be used. This is the case for a local WAS application, which could use a zAAP for the Java processing. Local QMF (TSO or CICS applications), batch programs, CICS, IMS, Stored procedures and User-Defined Functions (UDFs) do not use zIIP, unless they are running parallel queries or LOAD, REORG or REBUILD utilities.

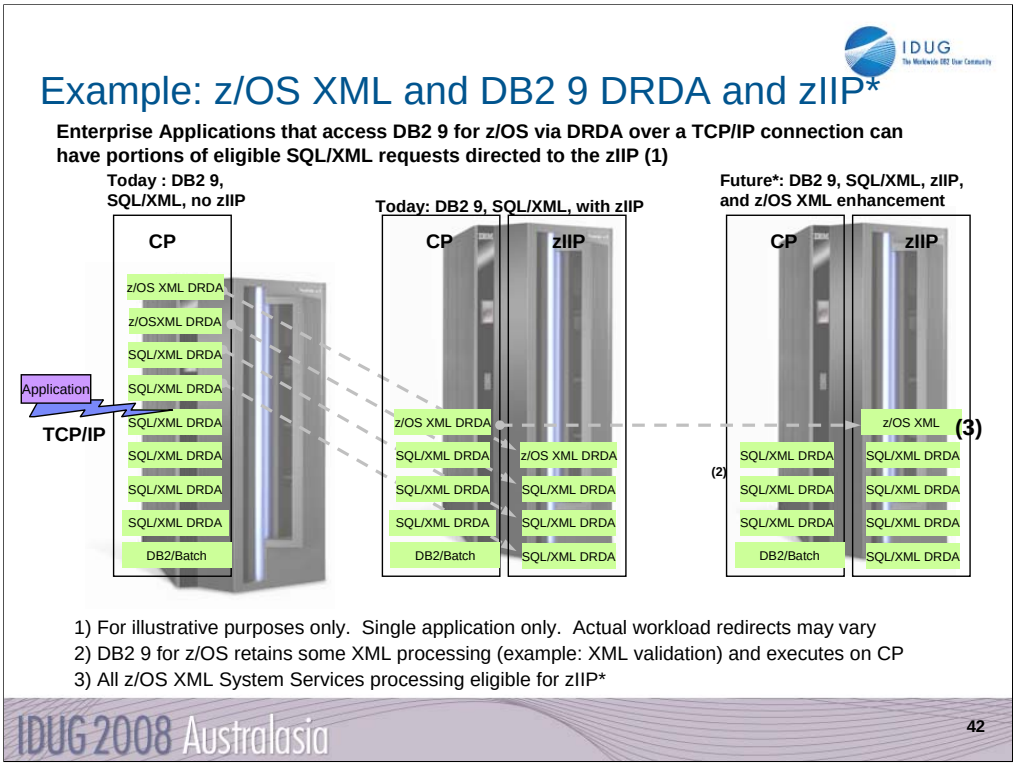

Remote SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP or PeopleSoft), CRM (Siebel), and business intelligence running on other platforms.

Remote SQL is expected to provide the primary benefits to customers, as it is commonly part of the peak load. Stored procedures and UDFs run under TCBs, so they are not generally eligible for zIIP, except for the call, commit and result set processing.

DB2 9 remote native SQL Procedure Language is eligible for zIIP processing.

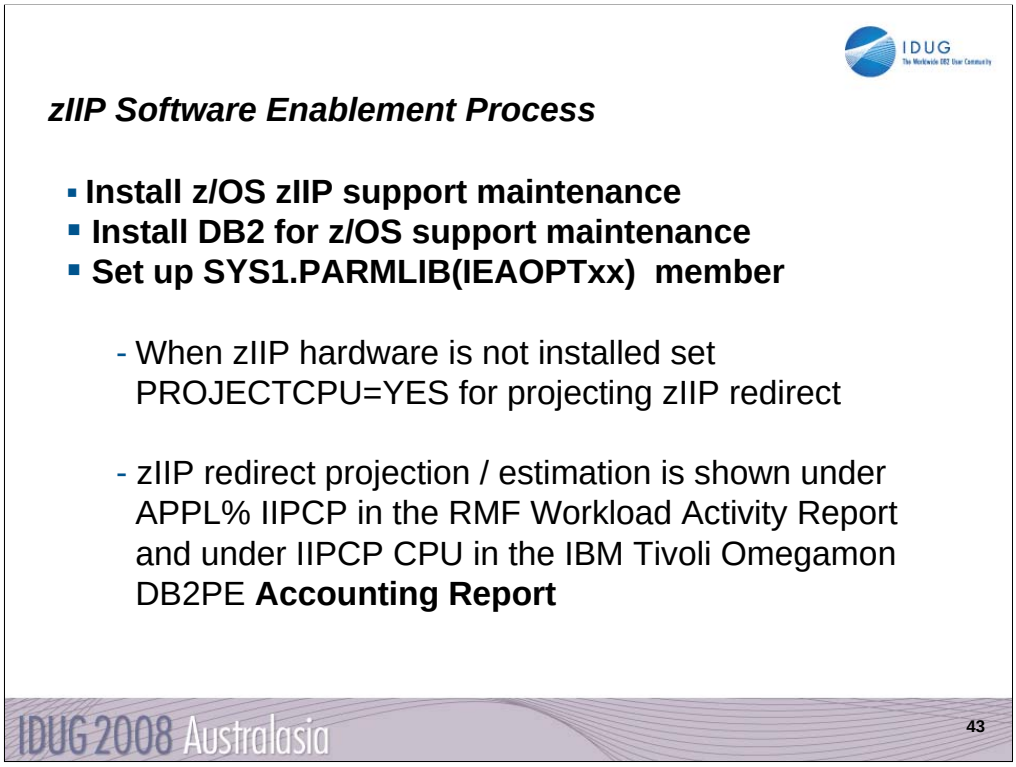

The PROJECTCPU=YES option (also available on z/OS V1R6 and z/OS V1R7 as part of the zIIP FMIDs) now also allows zAAP projection to occur, without requiring any per JVM configuration changes. Previously, each impacted JVM had to be individually configured to cause zAAP statistics to be collected in RMFand SMF.

To aid in determining the number of zIIP engines required to satisfy a specific customers usage, this parmlib option is available once all the software updates have been applied. The PROJECTCPU=YES parameter enables z/OS to collect zIIP usage as if there was one configured, when the target workload is being run. This projection capability can be run at any time, on a production environment if desired. RMF, SMF and IBM Tivoli Omegamon DB2 Performance Expert now show this calculated zIIP time so that an accurate zIIP projection can be made.

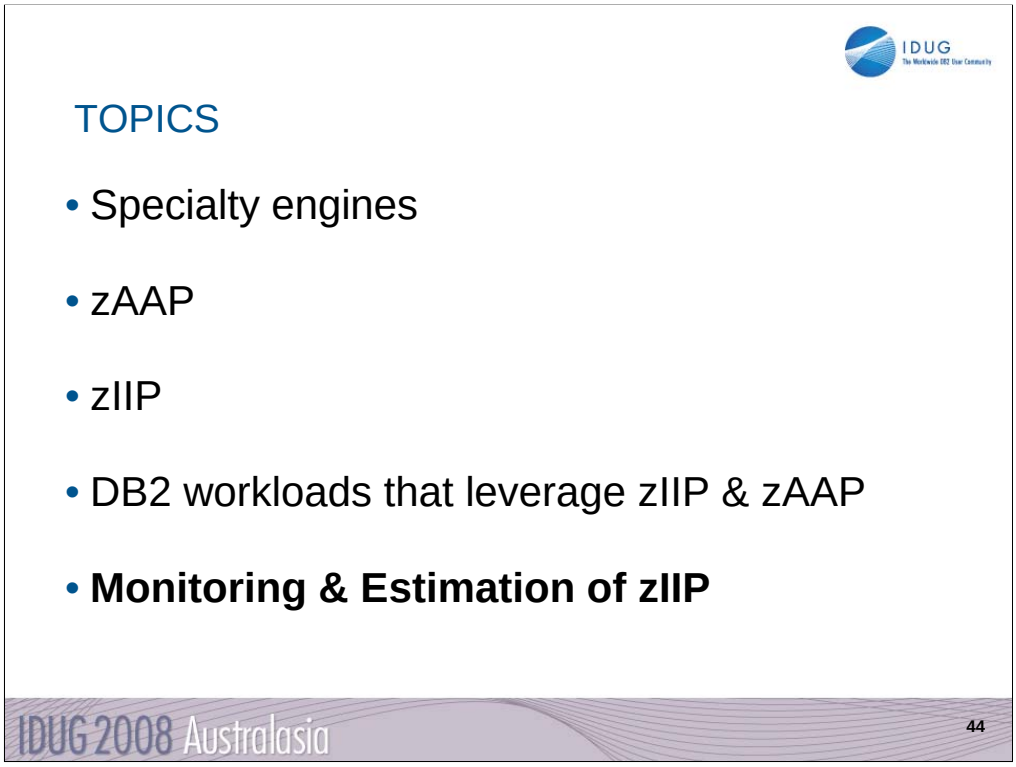

We'll start with an overview, then explain more about the work loads and what is eligible. Then we can work down to show some of the techniques for measurement and estimation, ending with a summary.

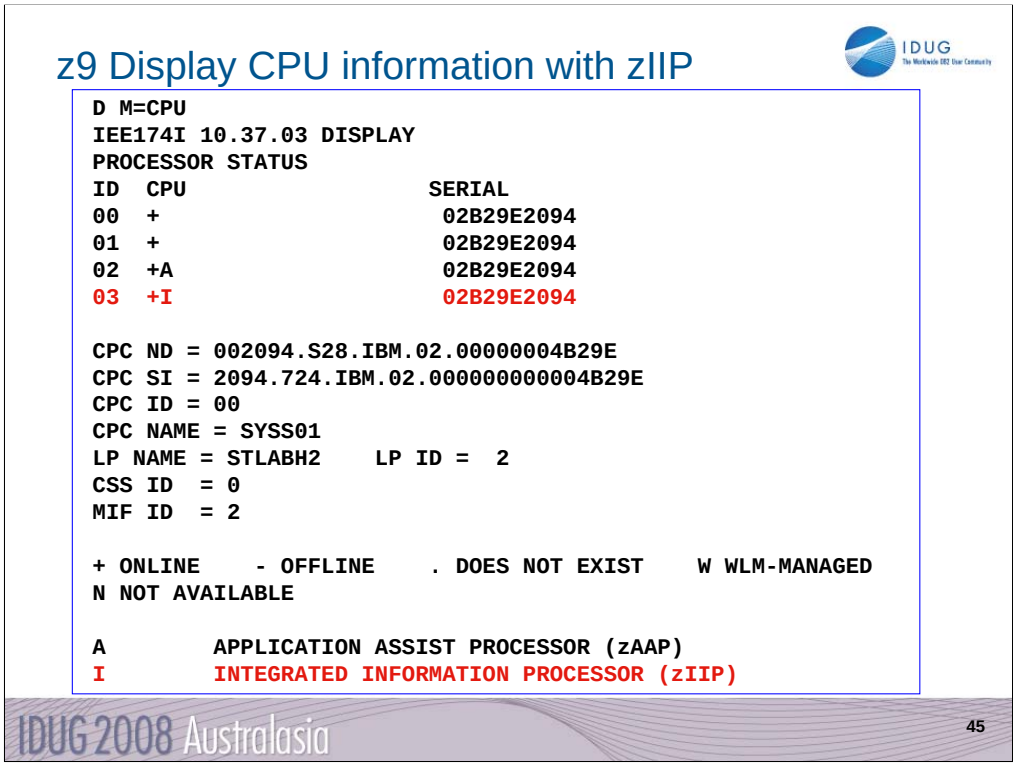

This chart shows the CPU configuration display on the z/OS console with the D M=CPU command.

This display shows a configuration with 2 CPs, 1 zAAP and 1 zIIP processor. + sign indicates that the processor is online.

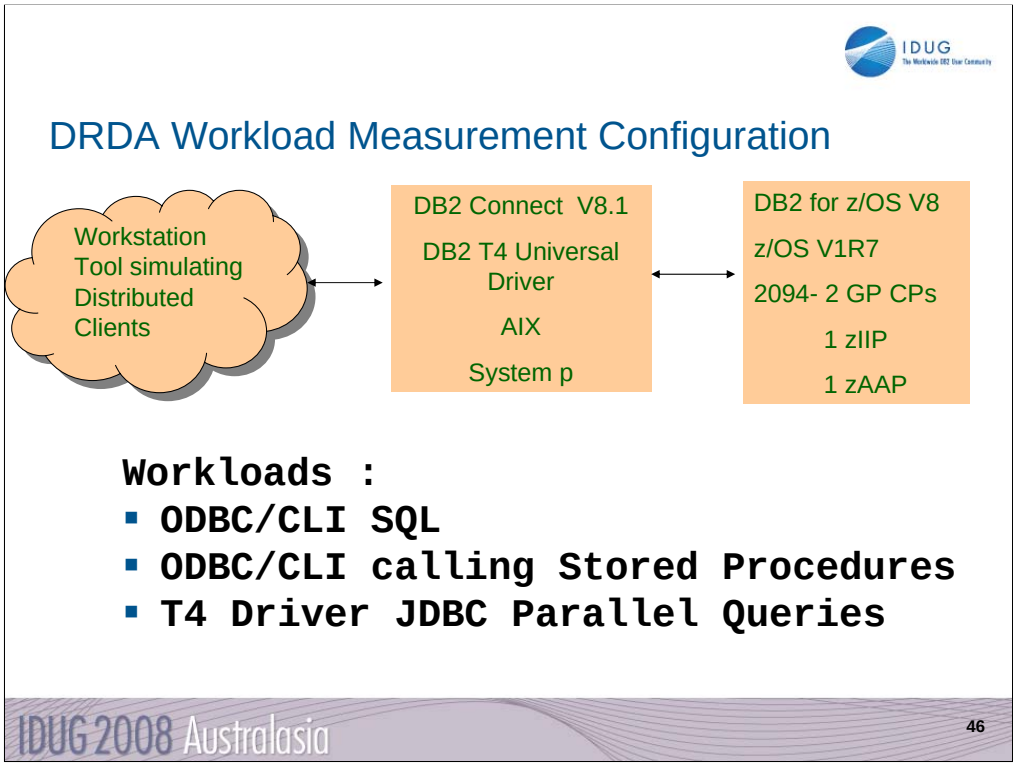

This is the configuration used for the DB2 tests and benchmarks I'll be showing you. The driving system is System p running AIX and connecting with DB2 Connect or a Java type 4 Universal Driver to DB2 for z/OS.

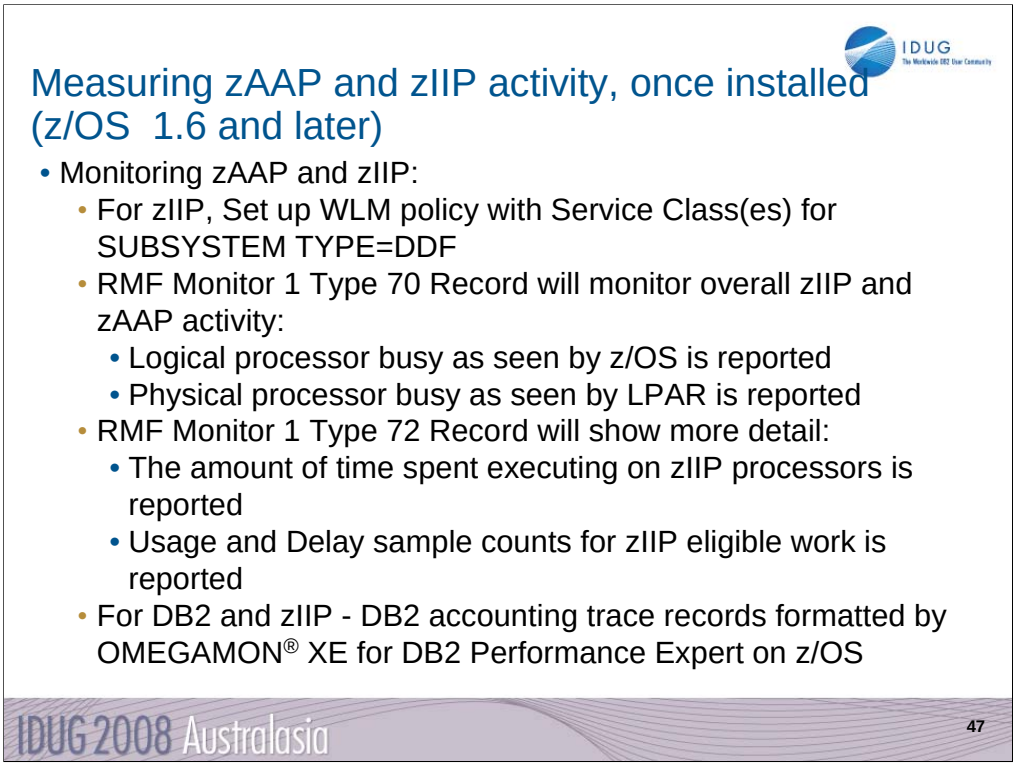

Once a zAAP or zIIP is installed (with appropriate maintenance), monitoring zAAP and zIIP activity is as follows:

For zIIP, Set up WLM policy with Service Class(es) for SUBSYSTEM TYPE=DDF

RMF Monitor 1 Type 70 Record will monitor overall zIIP and zAAP activity:

Logical processor busy as seen by z/OS is reported

Physical processor busy as seen by LPAR is reported

RMF Monitor 1 Type 72 Record will show more detail:

The amount of time spent executing on zIIP processors is reported

Usage and Delay sample counts for zIIP eligible work is reported

For DB2 and zIIP - In addition, DB2 accounting trace records can provide information on the zIIP. IBM Tivoli OMEGAMON® XE for DB2 Performance Expert on z/OS, DB2 Performance Expert or IBM Tivoli OMEGAMON XE for DB2 Performance Monitor on z/OS can be used to monitor the zIIP information.

Using PROJECTCPU

- The z/OS PROJECTCPU option in IEAOPTxx can help determine if a zIIP or a zAAP is right for you – zAAP or zIIP not needed.
- Set IEAOPTxx parmlib member option PROJECTCPU=YES. z/OS records amount of work eligible for zAAP or zIIP processors.
  - SMF Record Type 72 subtype 3 is input to the RMF post processor.
  - The Workload Activity Report lists workloads by WLM service class.
  - APPL% AAPCP indicates percentage of processor zAAP eligible
  - APPL% IIPCP indicates percentage of processor zIIP eligible
- Values as little as 25% for APPL% AAPCP or AAPL% IIPCP can be enough

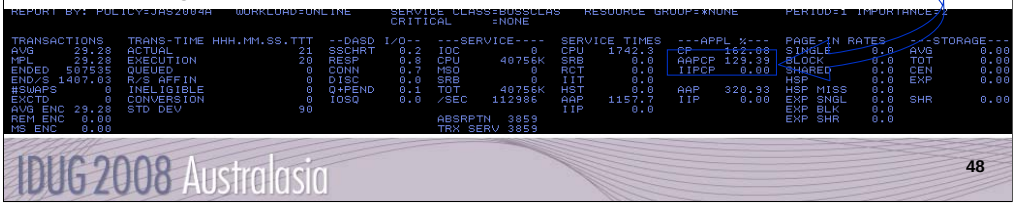

```
REPORT BY: POLICY=IAS20004   WORKLOAD=ONLINE      SERVICE CLASS=BOSSCLAS   RESOURCE GROUP=*NONE       PERIOD=1 IMPORTANCE=1
                                                  CRITICAL      =NONE

TRANSACTIONS    TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE----  SERVICE TIMES  ---APPL %---  PAGE-IN RATES    ---STORAGE---
AVG      29.28  ACTUAL                 21  SSCHRT   0.2  IOC          0  CPU    1742.3  CP     162.88  SINGLE     0.0  AVG      0.00
MPL      29.28  EXECUTION              20  RESP     0.8  CPU      40756K  SRB       0.0  AAPCP  129.39  BLOCK      0.0  TOT      0.00
ENDED   507535  QUEUED                  0  CONN     0.7  MSO          0  RCT       0.0  IIPCP    0.00  SHARED     0.0  CEN      0.00
END/S 1407.03  R/S AFFIN               0  DISC     0.0  SRB          0  IIT       0.0                HSP        0.0  EXP      0.00
#SWAPS       0  INELIGIBLE              0  Q+PEND   0.1  TOT      40756K  HST       0.0  AAP    320.93  HSP MISS   0.0
EXCTD        0  CONVERSION              0  IOSQ     0.0  /SEC     112986  AAP    1157.7  IIP      0.00  EXP SNGL   0.0  SHR      0.00
AVG ENC 29.28  STD DEV                 90                                 IIP       0.0                EXP BLK    0.0
REM ENC  0.00                                                ABSRPTN  3859                           EXP SHR    0.0
MS ENC   0.00                                                TRX SERV 3859
```


The z/OS PROJECTCPU option in IEAOPTxx (available on z/OS 1.6 with zIIP PTF) can help determine if a zIIP or a zAAP is right for you – zAAP or zIIP not needed.

Setting the IEAOPTxx parmlib member option PROJECTCPU=YES directs z/OS to record the amount of work eligible for zAAP or zIIP processors.

SMF Record Type 72 subtype 3 is input to the RMF post processor.

The Workload Activity Report lists workloads by WLM service class.

In this report the field APPL% AAPCP indicates which percentage of an processor is zAAP eligible, while the field APPL% IIPCP indicates which percentage of an processor is zIIP eligible.

Because the price of zIIP and zAAP specialty processors is less than that of general purpose processors – hardware and software - values as little as 25% for APPL% AAPCP or AAPL% IIPCP can make a zAAP or zIIP processor worth your while.

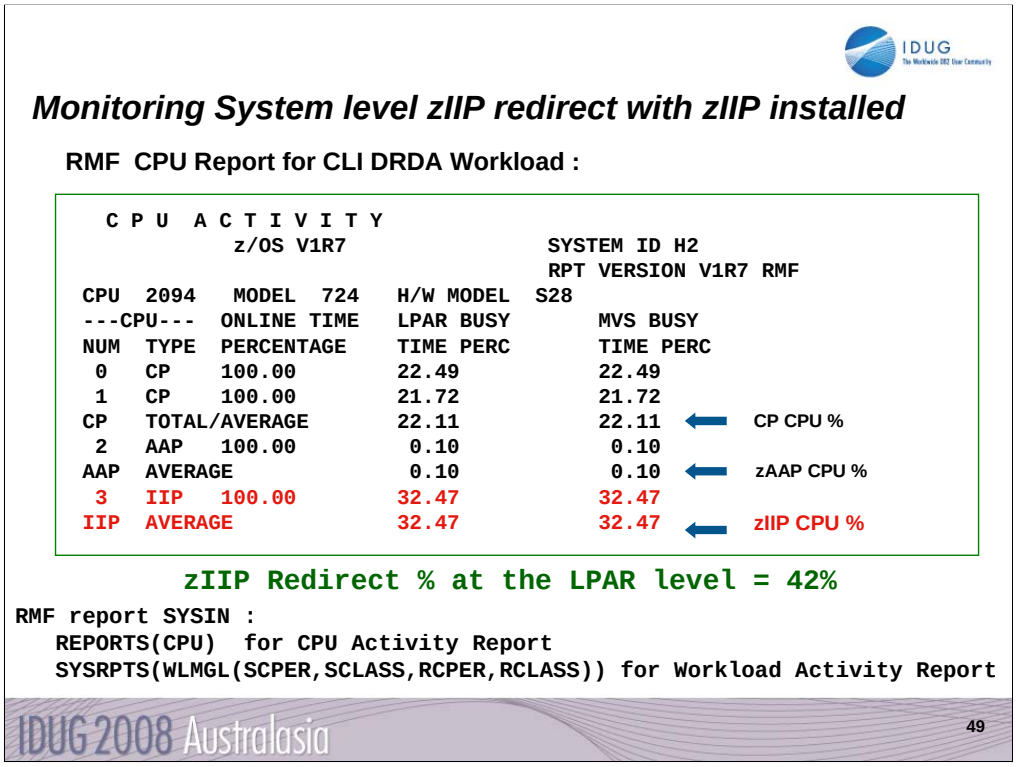

Chart shows the CPU utilization of the different processors in the LPAR generated by the RMF batch CPU report. Bottom of the chart shows the RMF batch report control cards to generate the RMF CPU report and the Workload activity reports.

The RMF CPU activity report shows 2 CPs, 1 zAAP and 1 zIIP engines.

The report is for a distributed ODBC/CLI workload showing 42% zIIP redirect at the LPAR level.

RMF Workload Activity Report Showing CLI SQL DRDA zIIP Estimate

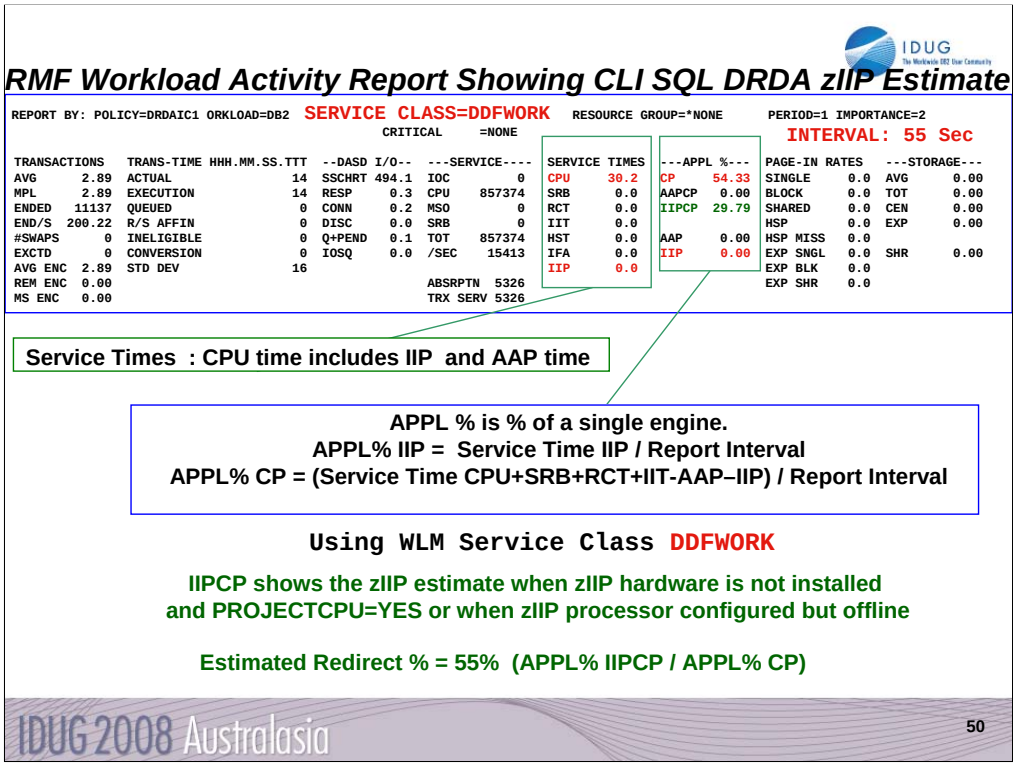

The chart shows the RMF workload activity report showing the zIIP redirect estimate for the DRDA CLI SQL workload before installing the zIIP hardware and using the SYS1.PARMLIB(IEAOPTxx) PROJECTCPU=YES parameter.

The WLM policy has been setup for Subsystem DDF and Service Class DDFWORK.

Service Times block shows the information in seconds for the interval. CPU time include IIP and AAP time. IIPCP (zIIP eligible on CP) is not shown under Service Times.

APPL % values are % of a single engine and is calculated from the Service times and dividing by the report interval. APPL% CP includes AAPCP and IIPCP.

Estimated redirect % = APPL% IIPCP / APPL% CP.

Estimate shows 55% redirect for the DB2 DRDA SQL CLI workload.

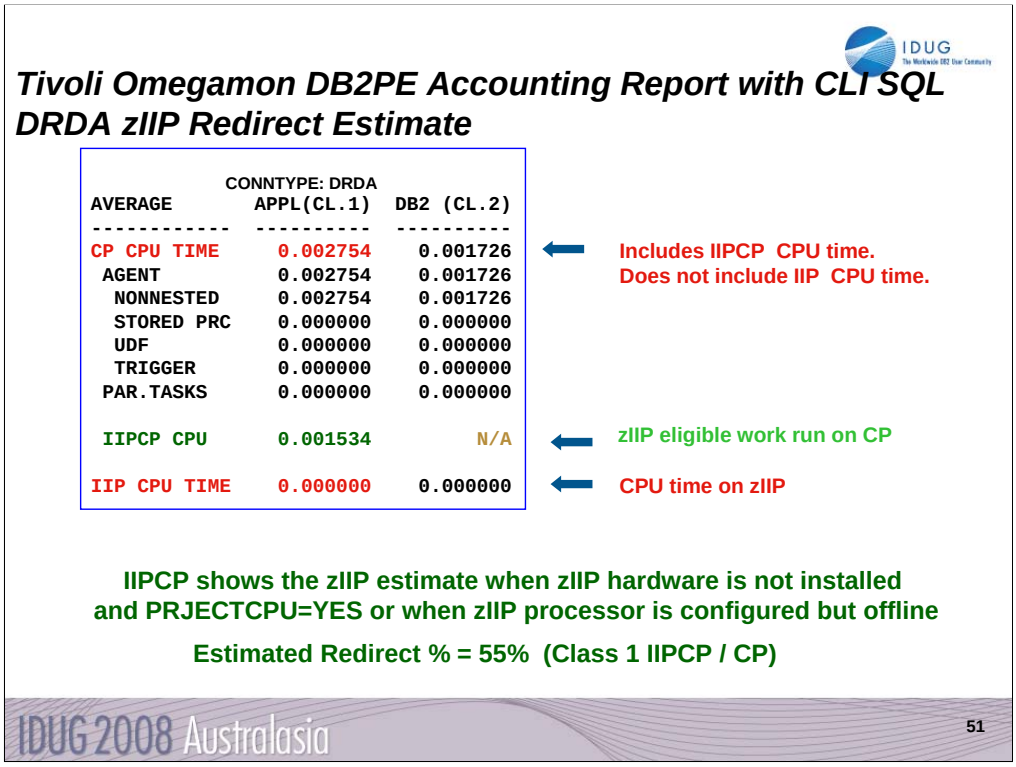

This chart shows how one can estimate the zIIP redirect using the IBM Tivoli Omegamon DB2 Performance Expert accounting report using PRJECTCPU=YES paramater when zIIP hardware is not installed. The estimation uses Class 1 CPU times. Class 1 is used because it include DDF CPU time and DB2 Class 2 CPU time.

IIP CPU time is the CPU time on zIIP. The IIP CPU time is zero since zIIP was not installed for this measurement.

IIPCP CPU time will show any zIIP eligible processing that ran on CP because zIIP was not installed or when the installed zIIP processors are too busy to handle the scheduled work. IIPCP CPU time is non-zero because zIIP is not installed and PROJECTCPU=YES.

The example shows zIIP redirect estimation of 55% at the DB2 DRDA workload level.

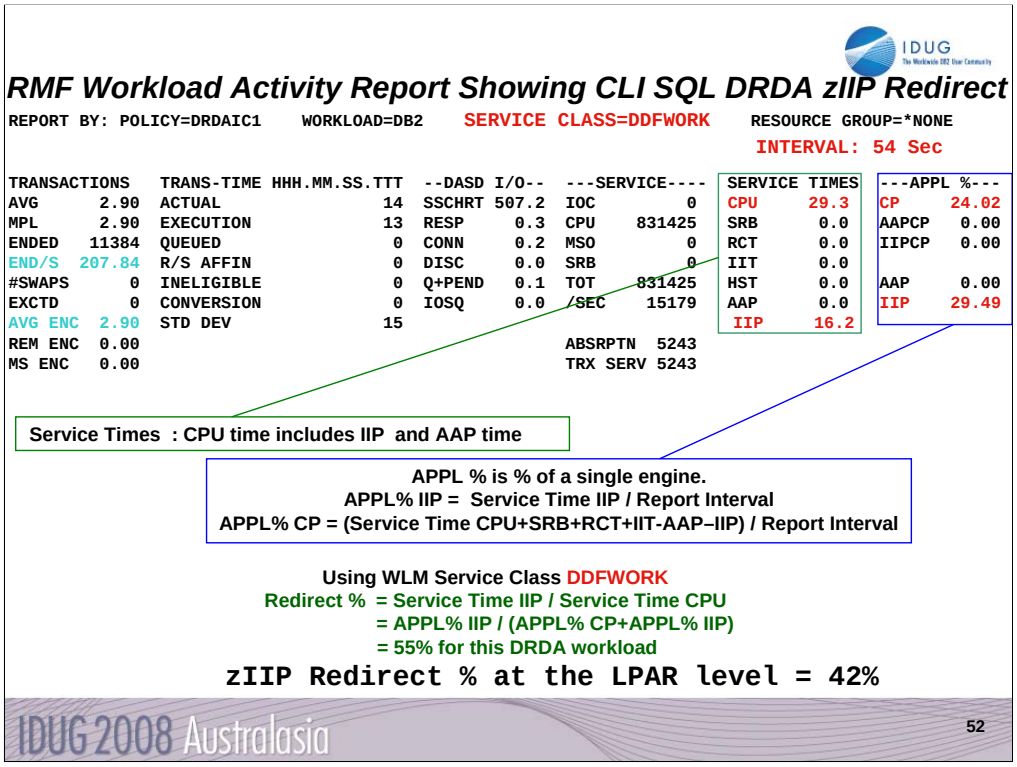

This chart shows the zIIP direct % when zIIP is being used.

It shows the redirect % of 55% at the DRDA workload level using the APPL% formula.

The effective redirect % for this workload at the LPAR level is 42% as shown in chart 27. It is lower at the LPAR level because of the CPU consumed by other non DB2 DRDA components (other DB2 address spaces, TCP/IP etc).

The DRDA redirect % can be calculated using the Service times also.

The formula is : Service Times IIP / Service Times CPU.

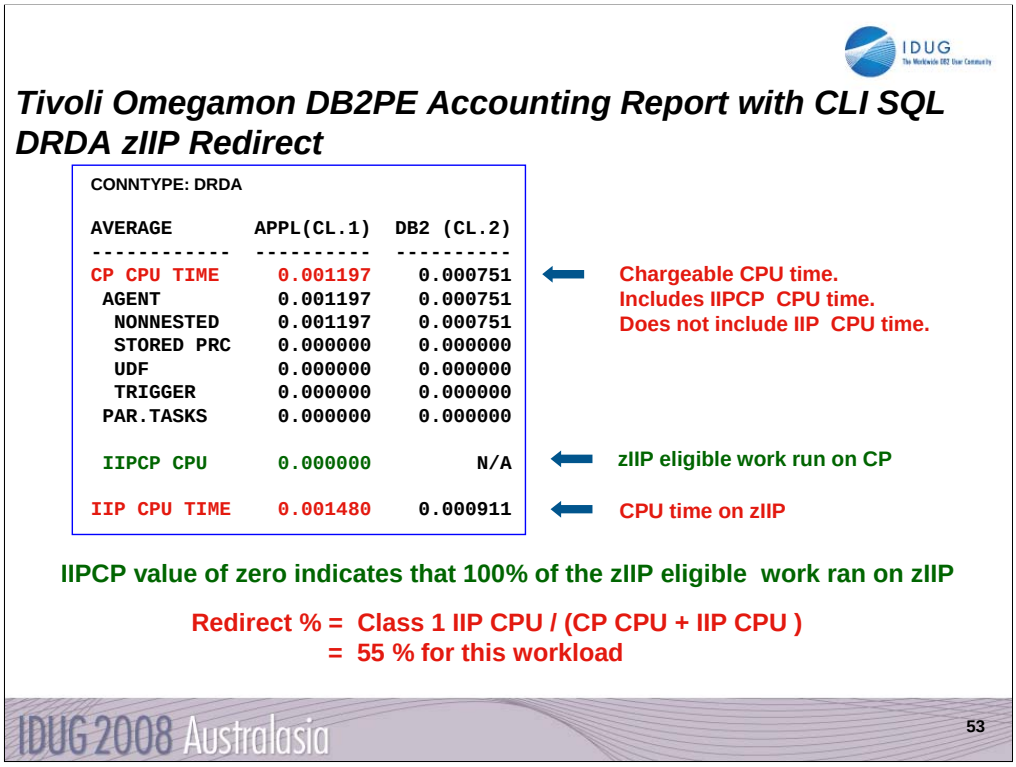

The chart shows the DRDA workload zIIP redirect % using the DB2 Performance Expert accounting report for Connect Type DRDA.

IIP CPU time is the CPU time on zIIP.

IIPCP CPU time will show any zIIP eligible processing that ran on CP because zIIP processor was busy.

High non-zero value will indicate a need to configure more zIIP processors. In this example the zero value for IIPCP CPU indicates that there is no need to configure additional zIIP processors.

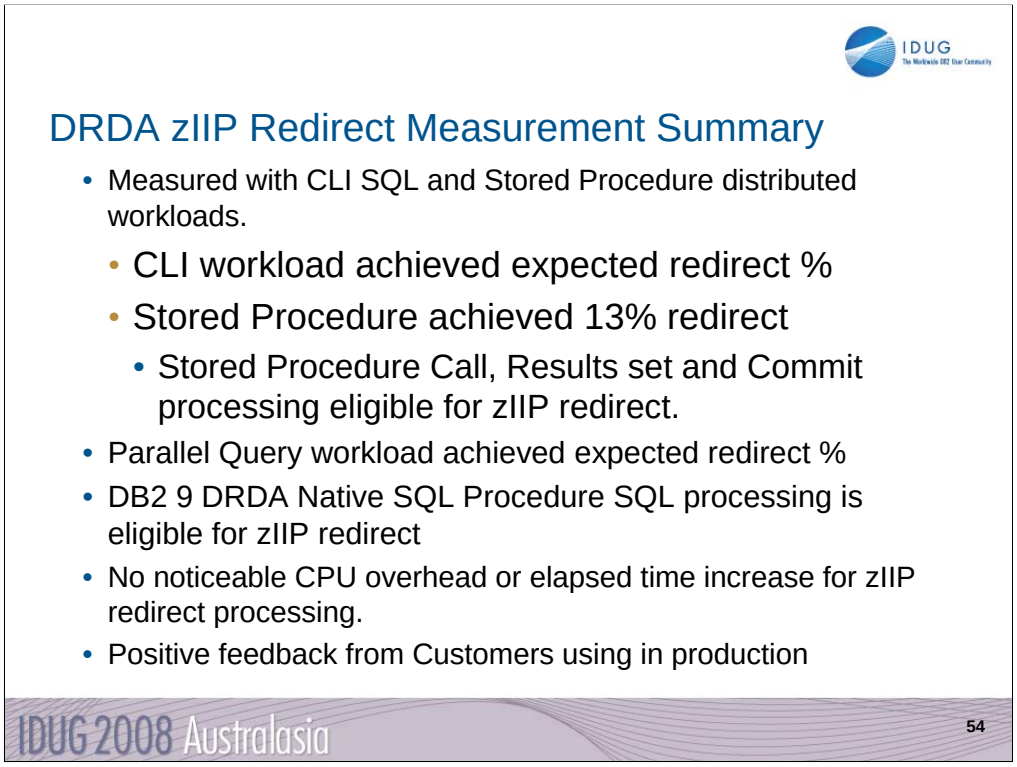

This chart summarizes the DB2 DRDA zIIP redirect measurements.

Parallel query via TCP/IP DRDA will get the redirect benefit from DRDA zIIP redirect for the main task and parallel query zIIP redirect for the parallel query child tasks.

Native SQL Stored Procedure calls from TCP/IP DRDA applications will under enclaves in DBM1 (instead of under WLM) and hence eligible for zIIP redirect.

Several customers are running successfully in production (Acxium / Trans Union, AEGON, Caterpillar, Baldor Electric, Citigroup, Mayo Clinic…).

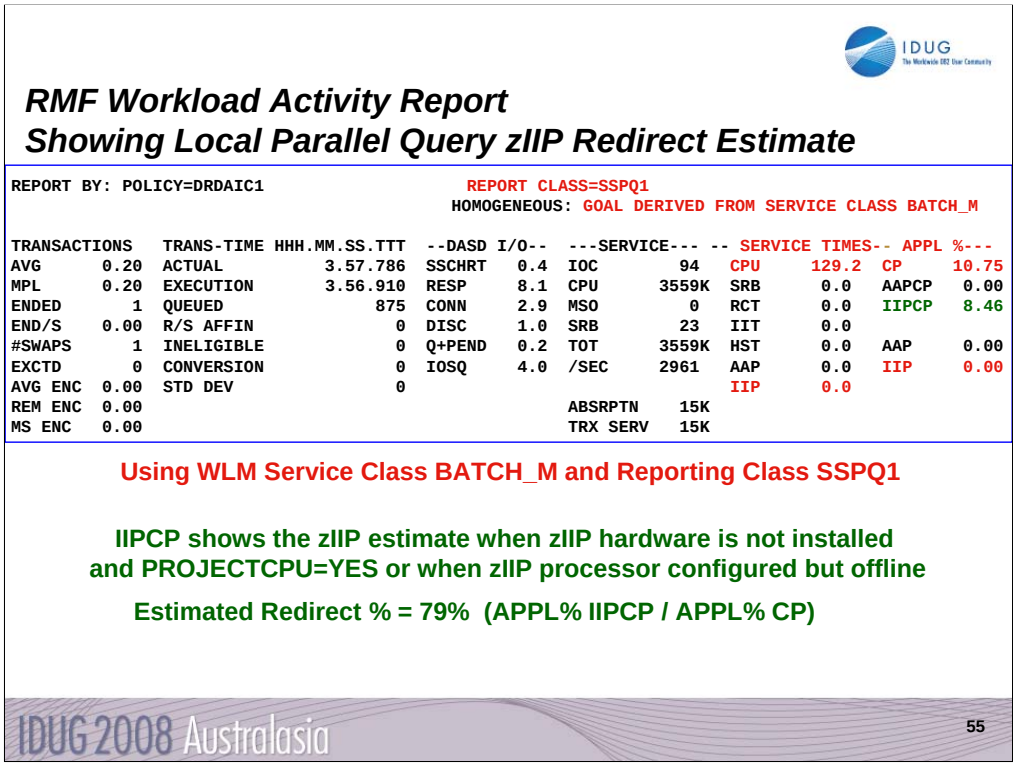

This chart is showing the zIIP redirect estimate for batch parallel query workload.

The WLM policy was setup for Subsystem JES, Service class BATCH_M and reporting class SSPQ1 for the parallel query Job name.

The example shows zIIP redirect estimate of 79% for the parallel query.

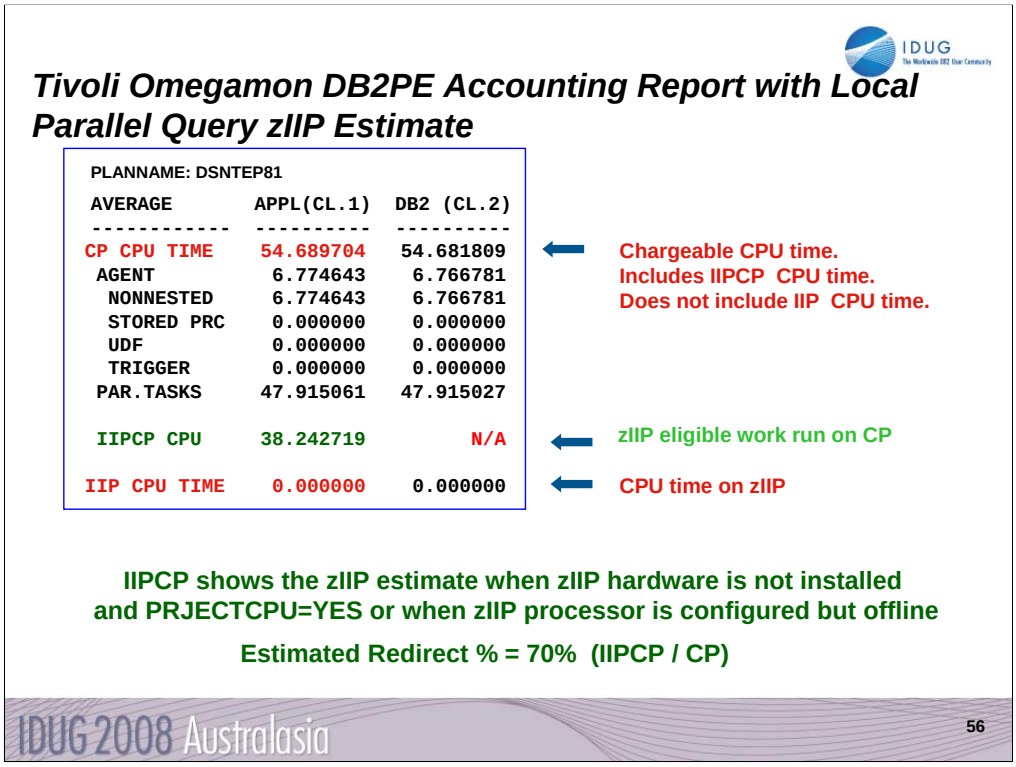

This chart shows the zIIP redirect estimate for all the parallel queries run under the Plan name DSNTEP81.

The estimated redirect % using the accounting report is 70%/

RMF Workload Activity Report
Showing Local Parallel Query zIIP Redirect

```
REPORT BY: POLICY=DRDAIC1                          REPORT CLASS=SSPQ1
                                  HOMOGENEOUS: GOAL DERIVED FROM SERVICE CLASS BATCH_M

TRANSACTIONS    TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE----  SERVICE TIMES  ---APPL %---
AVG      0.19   ACTUAL          3.52.930  SSCHRT   0.4  IOC         94  CPU     129.1  CP       2.23
MPL      0.19   EXECUTION       3.52.074  RESP     8.9  CPU       3556K SRB       0.0  AAPCP    0.00
ENDED       1   QUEUED                856  CONN     3.1  MSO          0  RCT       0.0  IIPCP    0.01
END/S    0.00   R/S AFFIN               0  DISC     1.5  SRB         28  IIT       0.0
#SWAPS      1   INELIGIBLE              0  Q+PEND   0.1  TOT       3556K  HST       0.0  AAP      0.00
EXCTD       0   CONVERSION              0  IOSQ     4.2  /SEC      2845  AAP       0.0  IIP      8.11
AVG ENC  0.00   STD DEV                 0                               IIP     101.3
REM ENC  0.00                                             ABSRPTN    15K
MS ENC   0.00                                             TRX SERV   15K
```

Using WLM Service Class BATCH_M and Reporting Class SSPQ1

Redirect % = Service Time IIP / Service Time CPU (more accurate)
= APPL% IIP / (APPL% CP+APPL%IIP)
= 78 % for this Query

This chart shows the actual parallel query zIIP offload for the WLM reporting class SSPQ1 when the zIIP is installed .

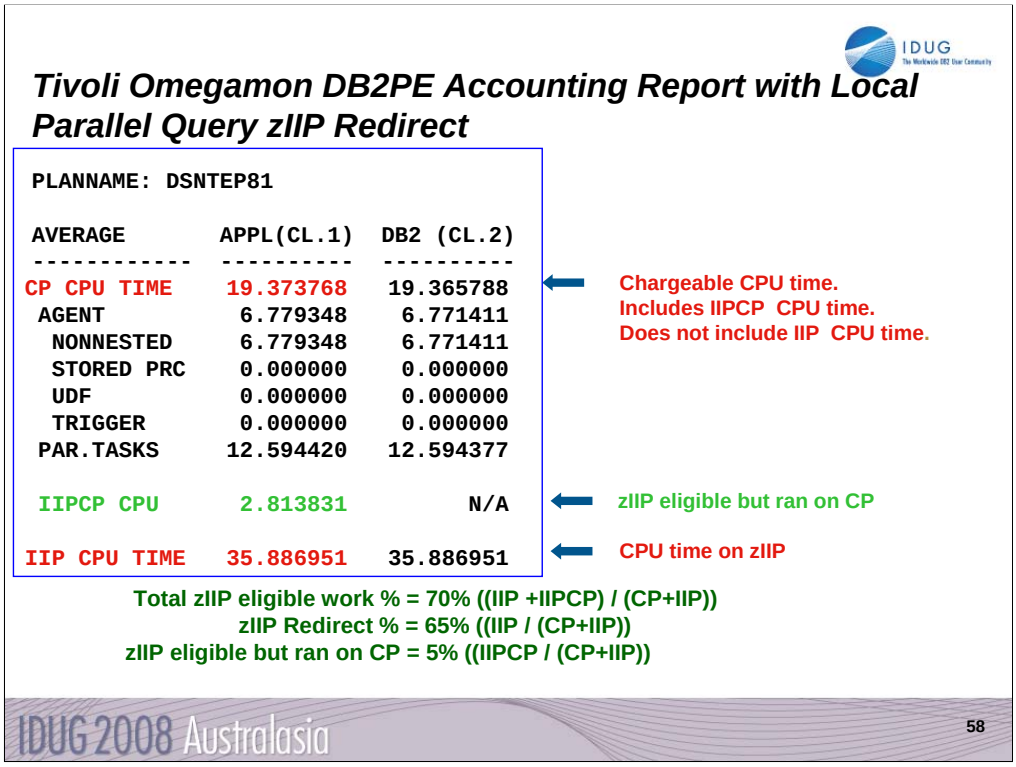

This chart shows the actual zIIP direct for all the parallel queries run under the Plan name DSNTEP81 when the zIIP is installed.

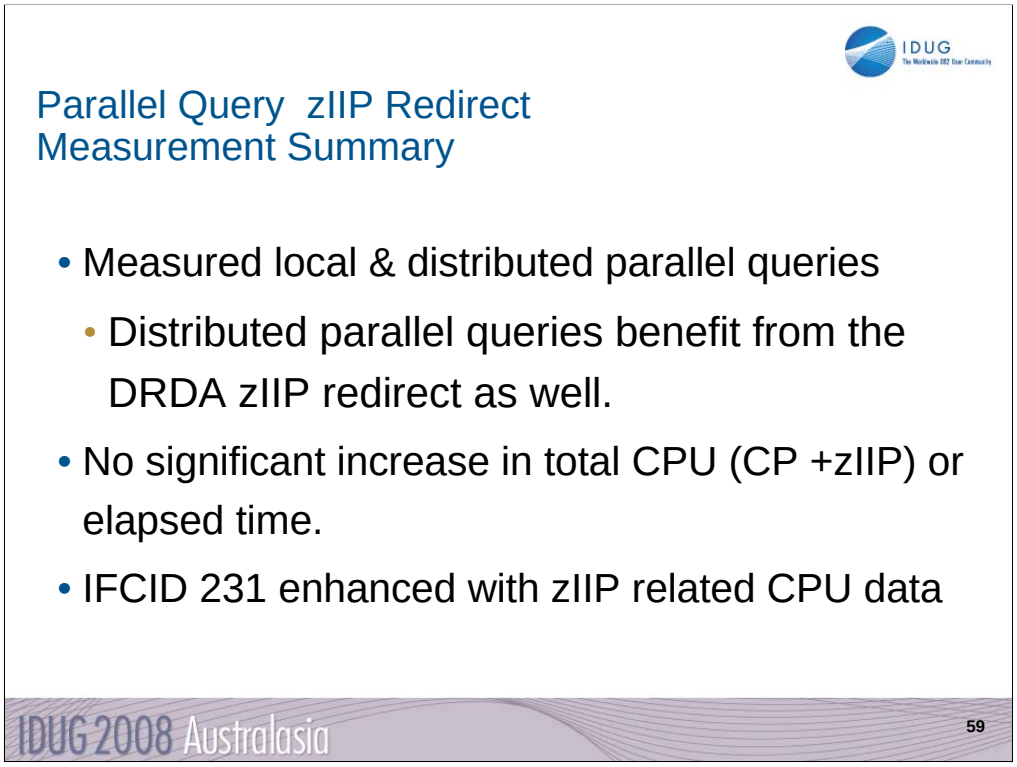

This chart summarizes the zIIP measurements for the parallel query.

CPU intensive parallel queries after their parallel group CPU consumptions exceeds certain threshold (100 ms) then subsequent child task processing will be scheduled to run under enclave SRB and a portion of it will be redirected to zIIP.

Parallel queries coming in via TCP/IP DRDA will get the DRDA zIIP redirect for the main task and also the parallel child task zIIP redirect after a certain parallel group CPU threshold is reached.

IFCID 231 has been enhanced to show the CP and IIP CPU time for the parallel child tasks.

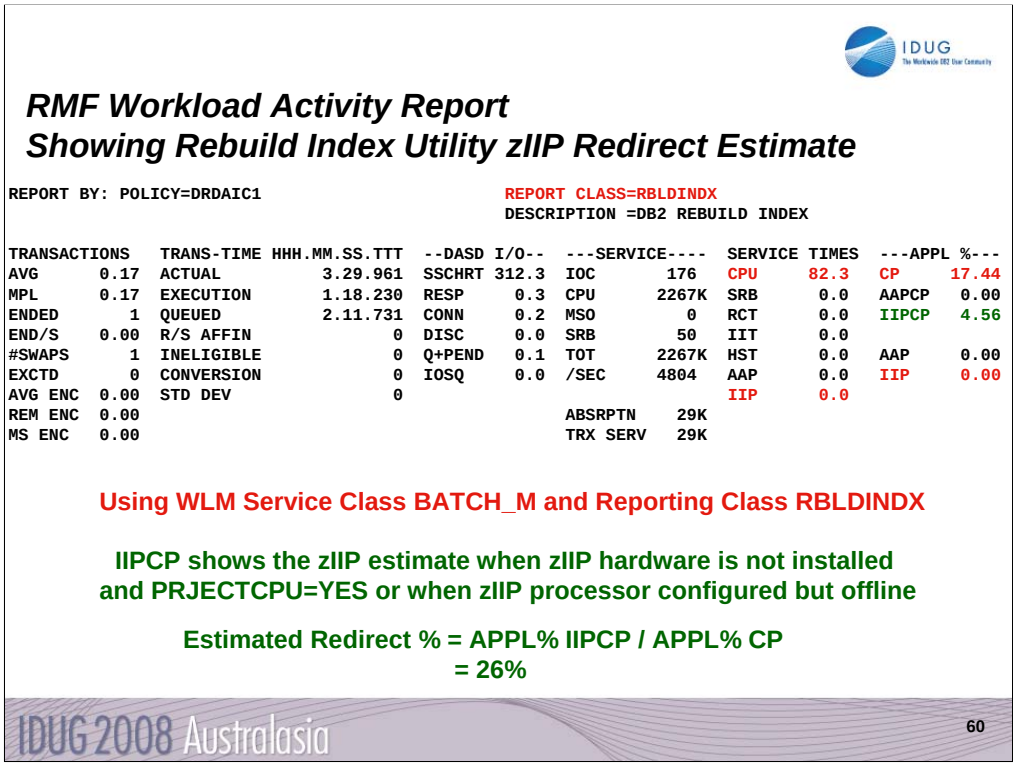

This chart shows the zIIP redirect estimate formula for the DB2 Rebuild Utility when the zIIP processor is not installed.

The WLM policy has been setup for Subsystem JES, Service Class BATCH_M and reporting class RBLDINDX for the DB2 Rebuild Utilities.

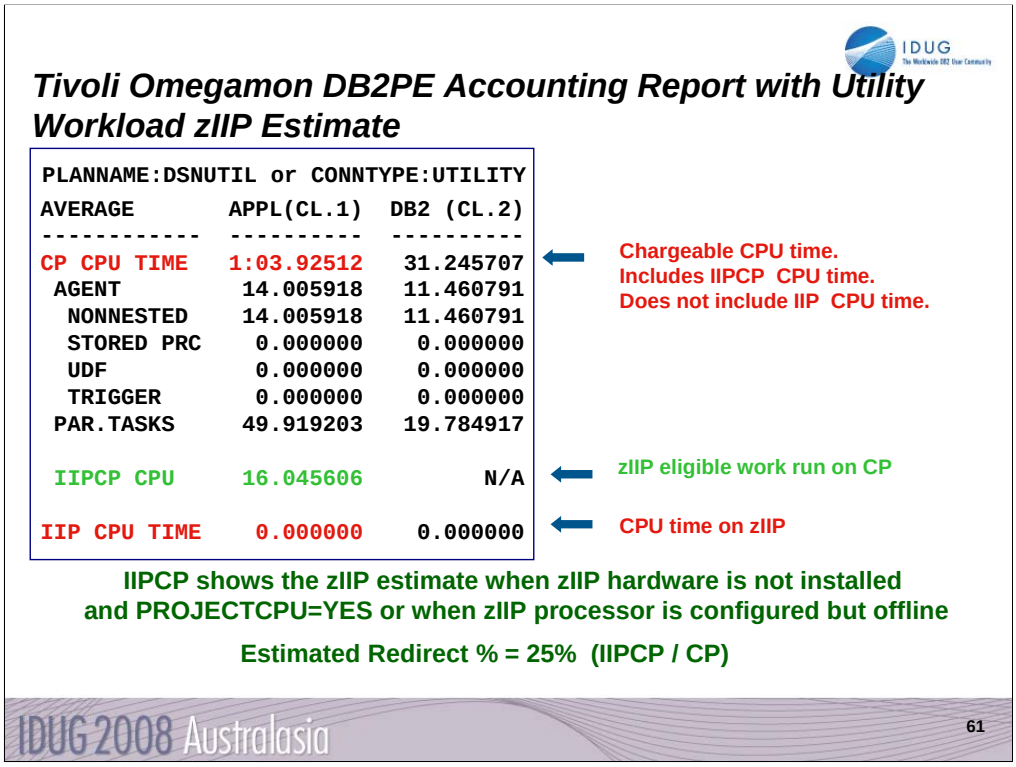

This chart shows the zIIP redirect estimate for all the DB2 Utilities that were run under Plan name DSNUTIL or Connect type Utility, when the zIIP hardware is not installed.

DB2 Class 1 IIPCP and CP CPU time is used to estimate the zIIP redirect.

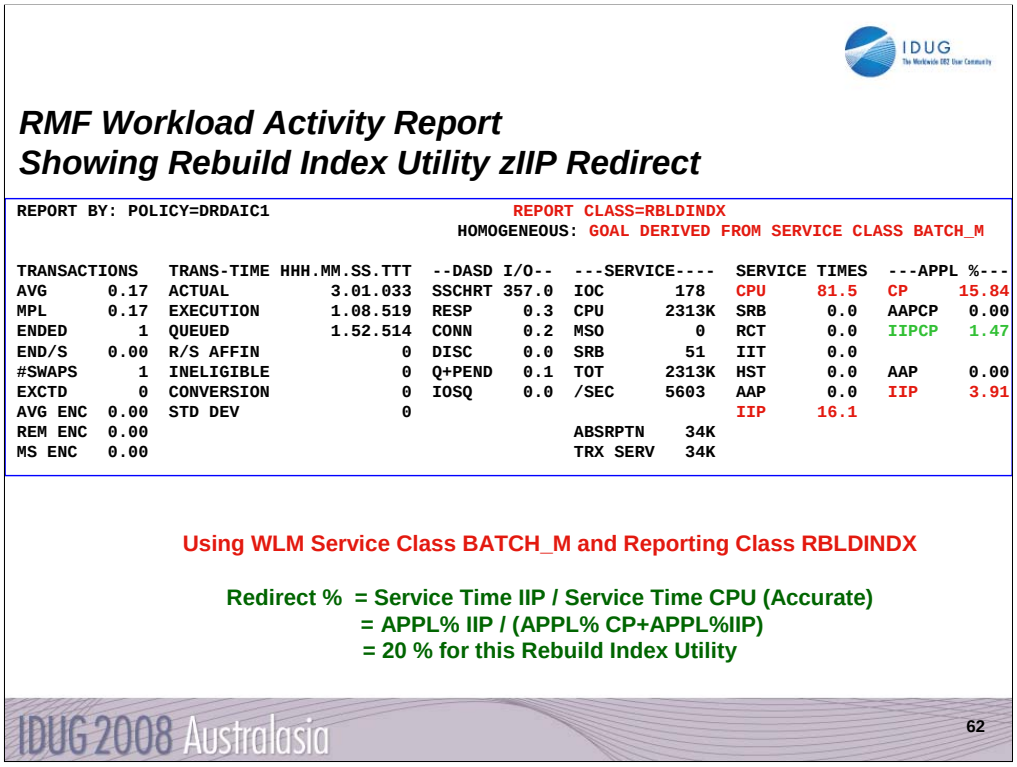

RMF Workload Activity Report
Showing Rebuild Index Utility zIIP Redirect

REPORT BY: POLICY=DRDAIC1 REPORT CLASS=RBLDINDX
HOMOGENEOUS: GOAL DERIVED FROM SERVICE CLASS BATCH_M

TRANSACTIONS		TRANS-TIME HHH.MM.SS.TTT		--DASD I/O--		---SERVICE----		SERVICE TIMES		---APPL %---	
AVG	0.17	ACTUAL	3.01.033	SSCHRT	357.0	IOC	178	CPU	81.5	CP	15.84
MPL	0.17	EXECUTION	1.08.519	RESP	0.3	CPU	2313K	SRB	0.0	AAPCP	0.00
ENDED	1	QUEUED	1.52.514	CONN	0.2	MSO	0	RCT	0.0	IIPCP	1.47
END/S	0.00	R/S AFFIN	0	DISC	0.0	SRB	51	IIT	0.0		
#SWAPS	1	INELIGIBLE	0	Q+PEND	0.1	TOT	2313K	HST	0.0	AAP	0.00
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	5603	AAP	0.0	IIP	3.91
AVG ENC	0.00	STD DEV	0					IIP	16.1		
REM ENC	0.00					ABSRPTN	34K				
MS ENC	0.00					TRX SERV	34K				

Using WLM Service Class BATCH_M and Reporting Class RBLDINDX

Redirect % = Service Time IIP / Service Time CPU (Accurate)
= APPL% IIP / (APPL% CP+APPL%IIP)
= 20 % for this Rebuild Index Utility

This chart shows the actual zIIP redirect % for the DB2 Rebuild Index Utilities under the WLM reporting class RBLDINDX.

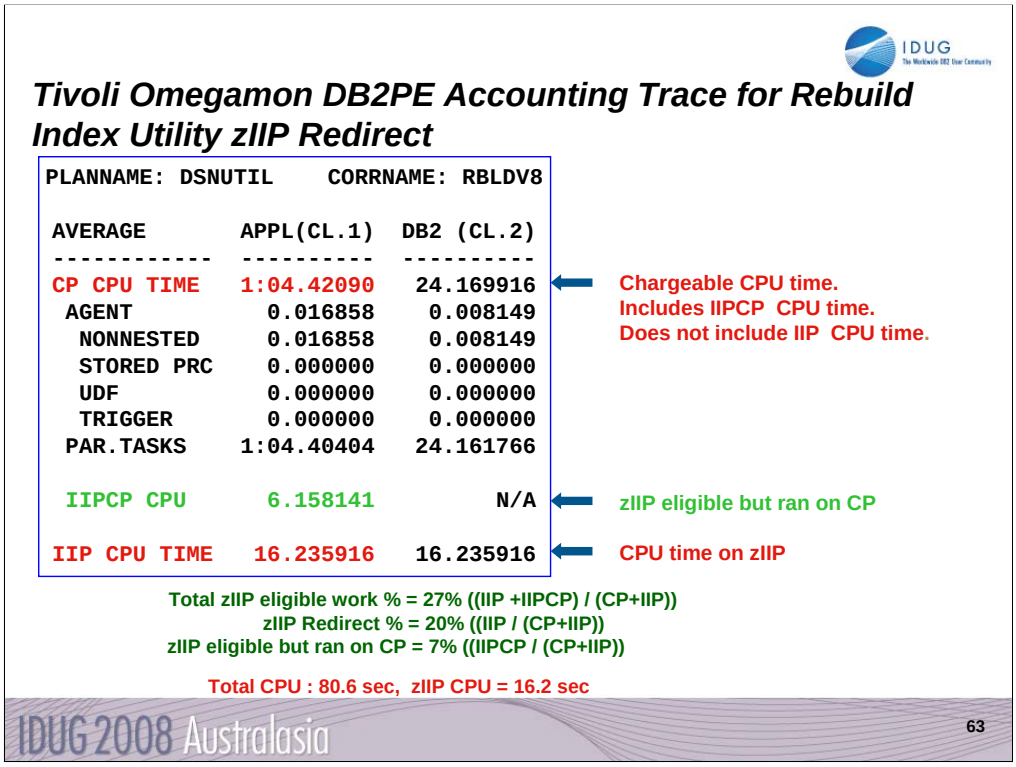

This chart shows the actual zIIP redirect % for the DB2 Rebuild Utility job RBLDV8 when the zIIP is installed.

Notice non zero value for IIPCP CPU indicating some of the zIIP eligible processing was redirected back to run on CP since the zIIP processor was busy. High number for IIPCP CPU value is an indication that additional zIIP processors could be added.

Note that the Total CPU seconds and IIP CPU seconds shown in accounting trace at the bottom matches the WLM Service times CPU and IIP seconds in chart 43.

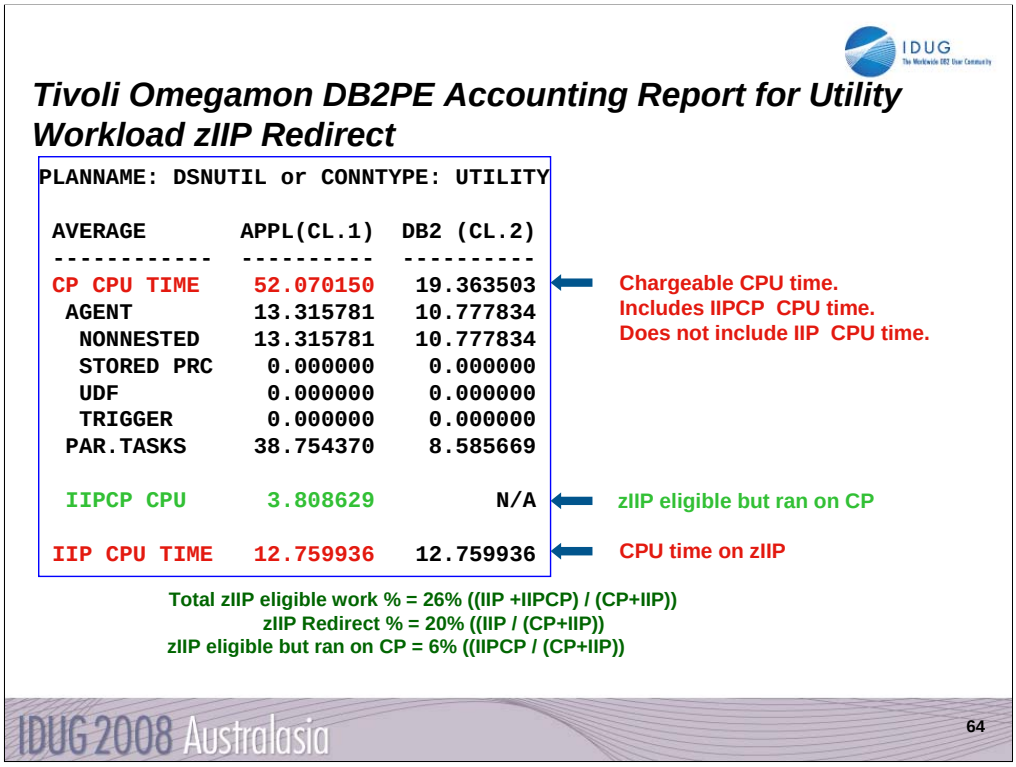

This chart shows the actual zIIP redirect for all the DB2 Utilities that were run (Plan name DSNUTIL or Connect Type UTILITY) when the zIIP processor is installed.

Notice non zero value for IIPCP CPU indicating zIIP was processor was unable to process all the zIIP eligible work.

Utility zIIP Redirect Measurement Summary

- Measured LOAD, REBUILD INDEX and REORG Utilities.
- zIIP redirect % depends on % CPU consumed by the Build Index phase of the Utility.
- Observed Class 1 CPU reduction for configuration with 4 CPs and 2 zIIPs with fixed length Index key :
  - 5 to 20% for Rebuild Index
  - 10 to 20% for Load or Reorg of a Partition with one Index only, or Load of entire Table, or Reorg of entire Tablespace
  - 40% for Rebuild Index of logical Partition of Non Partitioning Index
  - 40 to 50% for Reorg Index
  - 30 to 60% for Load or Reorg of a Partition with more than one Index
- CPU overhead incurred during execution unit switch from TCB to enclave SRB during Index Rebuild phase
  - Typically less than 10% eligible for redirect

This chart summarizes the zIIP redirect measurements for the DB2 Utilities.

The zIIP redirect % is proportional to amount of build index processing.

CPU cost associated with SORT and Compression are not eligible for zIIP redirect.

To support zIIP redirect the build index processing was changed from TCB processing to enclave SRB processing which introduces execution unit switch CPU overhead. The zIIP redirect % shown in this chart have been adjusted for the CPU overhead. The measurement were done with a Tablespace with 10 partitions and varying number of indices up to 6.

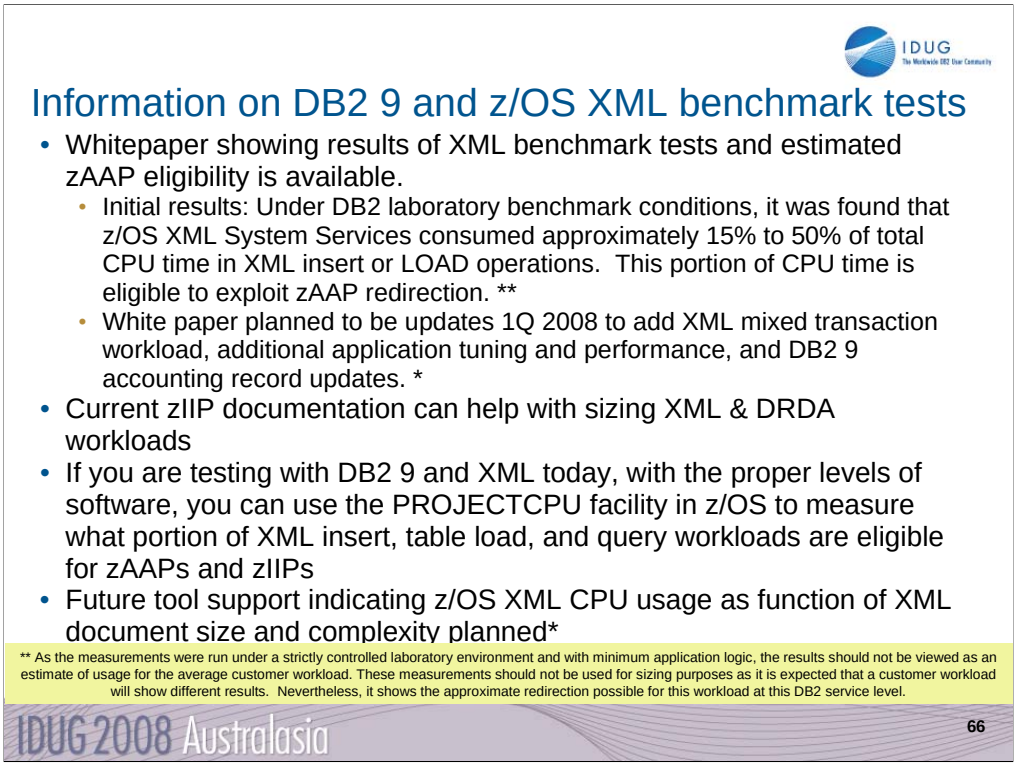

In laboratory DB2 benchmark performance runs, z/OS XML System Services consumes approximately 15% to 50% of total CPU time in XML insert or LOAD operations. This portion of CPU time is eligible for zAAP redirection.

The amount of CPU time for z/OS XML System Services will vary widely for other applications, based on the document size, its complexity, and number of indexes defined on XML tables.

When the size and number of nodes of the documents increase, the redirection percentage could increase accordingly, up to about 60%, according to laboratory measurements. Beyond that, there is no additional redirection because the corresponding DB2 processing to construct the internal representation of the XML document increases at the same time.

As more XML indexes are defined in XML tables, the percentage of redirection will be lower, as there is then more index processing within DB2 that is not related to z/OS XML System Services.

When encoding conversion is necessary, the percentage of redirection will be lower. DB2 always uses UTF-8 to process XML data. If the source XML that is inserted/updated or loaded is not in UTF-8, z/OS conversion services will be invoked to convert it to UTF-8.

As the measurements were run under a strictly controlled laboratory environment and with minimum application logic, the results should not be viewed as an estimate of usage for the average customer workload. These measurements should not be used for sizing purposes as it is expected that a customer workload will show different results. Nevertheless, it shows the approximate redirection possible for this workload at this DB2 service level.

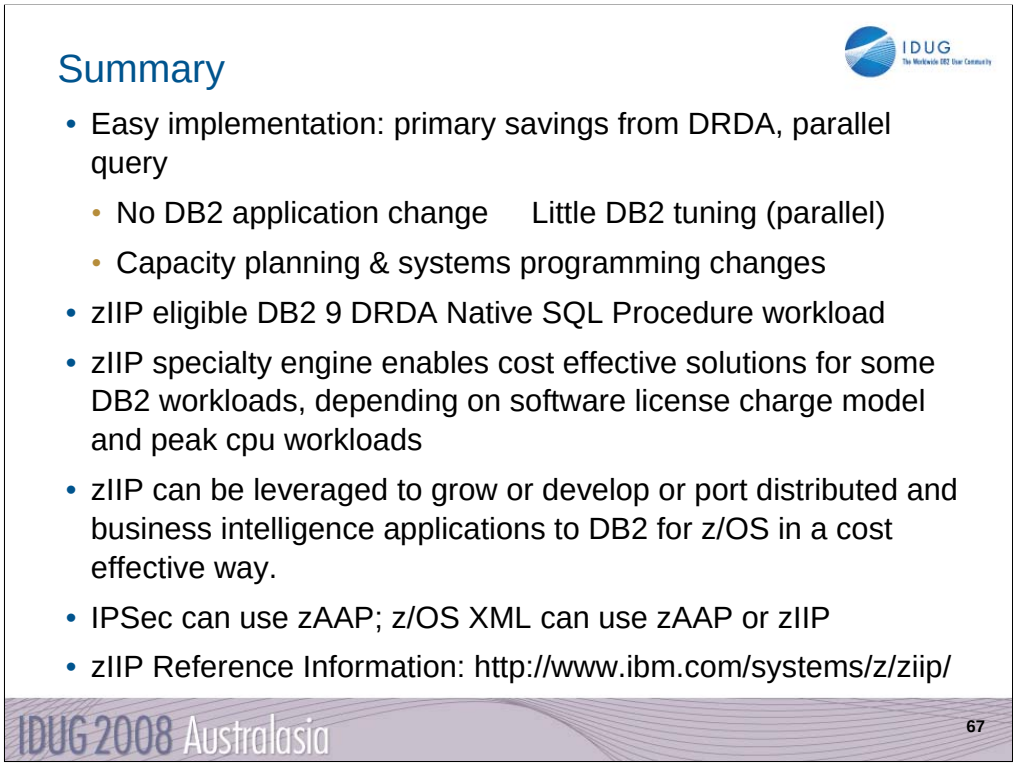

While the applications do not change and little DB2 tuning is needed, systems programming is needed to get the appropriate service and settings. Capacity planning is needed to have the right number of zIIPs and zAAPs. DB2 tuning may be desired to change some applications to use parallel processing. The primary guidance for those who want to use zIIP more is to

- Partition the tables to an appropriate number of parallel tasks
- Move to DB2 9

Resources see notes below

- System z Specialty engines
  - ibm.com/systems/z/specialtyengines/
- z/OS 1.9 Announcement Letter
- z/OS XML System Services site
  - ibm.com/servers/eserver/zseries/zos/xml/
- DB2 9 for z/OS
  - ibm.com/software/db2zos/
  - DB2 9 performance benchmark paper
    ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101088
  - Detailed XML introduction presentation:
    ftp.software.ibm.com/software/data/db2zos/db29zospureXMLgzhang.pdf
  - DB2 z/OS XML Guide:
    publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db29.doc.xml/bknxspsh.htm

Please see the web for much more information.

Specialty processors http://www.research.ibm.com/journal/rd/511/wyman.html

zAAP http://www.ibm.com/systems/z/advantages/zaap/index.html
http://www.redbooks.ibm.com/abstracts/sg246386.html
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100417
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1076

zIIP http://www.ibm.com/systems/z/advantages/ziip/
http://www.ibmsystemsmag.com/mainframe/novemberdecember07/coverstory/18227p1.aspx
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100988
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101088
http://www.ibm.com/support/docview.wss?rs=852&uid=swg27009459
http://www.redbooks.ibm.com/abstracts/sg247473.html
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD103516

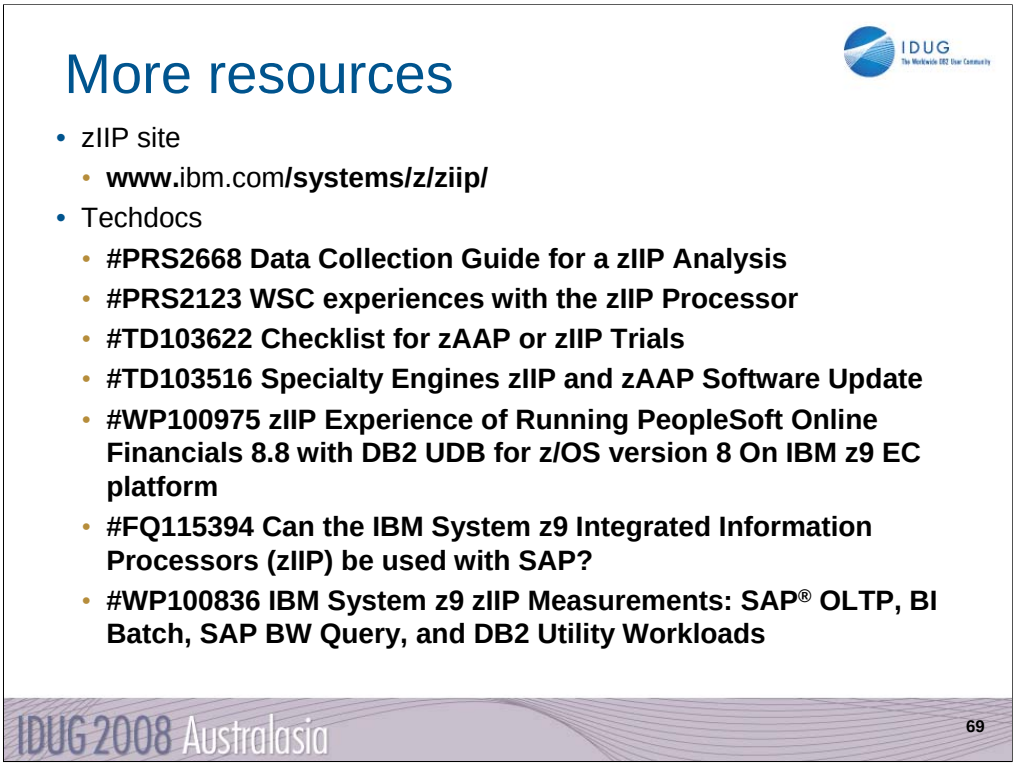

WSC papers

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2123
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2745
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2545
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS2668
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100836
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100975
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100988
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101088
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD103516
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100417
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1076

Additional resources

- Many presentations are available on the web
  - See notes below for details
- IBM Journal of Research & Development
  http://researchweb.watson.ibm.com/journal/rd/511/wyman.pdf
- zSeries Application Assist Processor (zAAP) Implementation
  http://www.redbooks.ibm.com/abstracts/sg246386.html?Open
- DB2 9 Performance Topics
  http://www.redbooks.ibm.com/abstracts/sg247473.html?Open
- Enhance Query Parallelism With zIIP Processors
  http://www.ibmsystemsmag.com/mainframe/enewsletterexclusive/18822p1.aspx

Road show handouts
http://www.ibm.com/software/os/systemz/roadshows/handouts/
ftp://ftp.software.ibm.com/software/zseries/pdf/roadshows_handouts/Session_3-1_Specialty_Engines_handout_v4.pdf
ftp://ftp.software.ibm.com/software/systemz/pdf/roadshows/handouts/06_-_Consolidating_Data_on_System_z_v2.0_-_Americas.pdf
ftp://ftp.software.ibm.com/software/systemz/pdf/roadshows/handouts/10_-_Delivering_Next_Gen_Solutions_at_the_Lowest_Cost_v2.0_-_Americas.pdf
ftp://ftp.software.ibm.com/software/zseries/pdf/roadshows_handouts/zIIPS_and_zAAPS_Capacity_Planning_Z103.pdf
ftp://ftp.software.ibm.com/software/zseries/pdf/roadshows_handouts/Specialty_Engine_WSC_Experiences_zaAPS_and_ZIIPs_Z104.pdf
ftp://ftp.software.ibm.com/software/zseries/pdf/roadshows_handouts/IBM_Z_Power_White_Paper-_Feb2007.pdf
ftp://ftp.software.ibm.com/software/zseries/pdf/roadshows_handouts/Session_3-2_Capacity_Planning_for_SOA_with_z_handout_V4.pdf
other zIIP & performance presentations
ftp://ftp.software.ibm.com/software/data/db2zos/S2_zIIP_DB2_workloads_2007Aug08.pdf
ftp://ftp.software.ibm.com/software/data/db2zos/IMTechzOS21zIIP.pdf
ftp://ftp.software.ibm.com/software/data/db2zos/v9sysperf2007oct.pdf
Primary System z web pages
specialty engines http://www.ibm.com/systems/z/specialtyengines/
zIIP http://www.ibm.com/systems/z/ziip/
zAAP http://www.ibm.com/systems/z/zaap/
z/OS 1.9 delivery http://www.ibm.com/systems/z/os/zos/overview/zosnew_summary.html
http://publibz.boulder.ibm.com/cgi-bin/bookmgr_OS390/BOOKS/E0Z2A119/1.7.5
Capacity Planning for zIIP Assisted IPSec
http://www.ibm.com/support/docview.wss?uid=swg27009459&aid=1
http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100988
zAAP capacity planning http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100417
Large Systems Performance Reference
http://www.ibm.com/servers/eserver/zseries/lspr/pdf/SC28118711.pdf
John Shedletsky http://www.ibm.com/software/swnews/swnews.nsf/n/cres78cr4r
ftp://ftp.software.ibm.com/software/systemz/pdf/TCOStamford_Mainstream.pdf
Gopal Krishnan IOD 1782 Leveraging zIIP with DB2 for z/OS
Terry Purcell IOD 1308 Tuning your SQL to get the most out of zIIPs

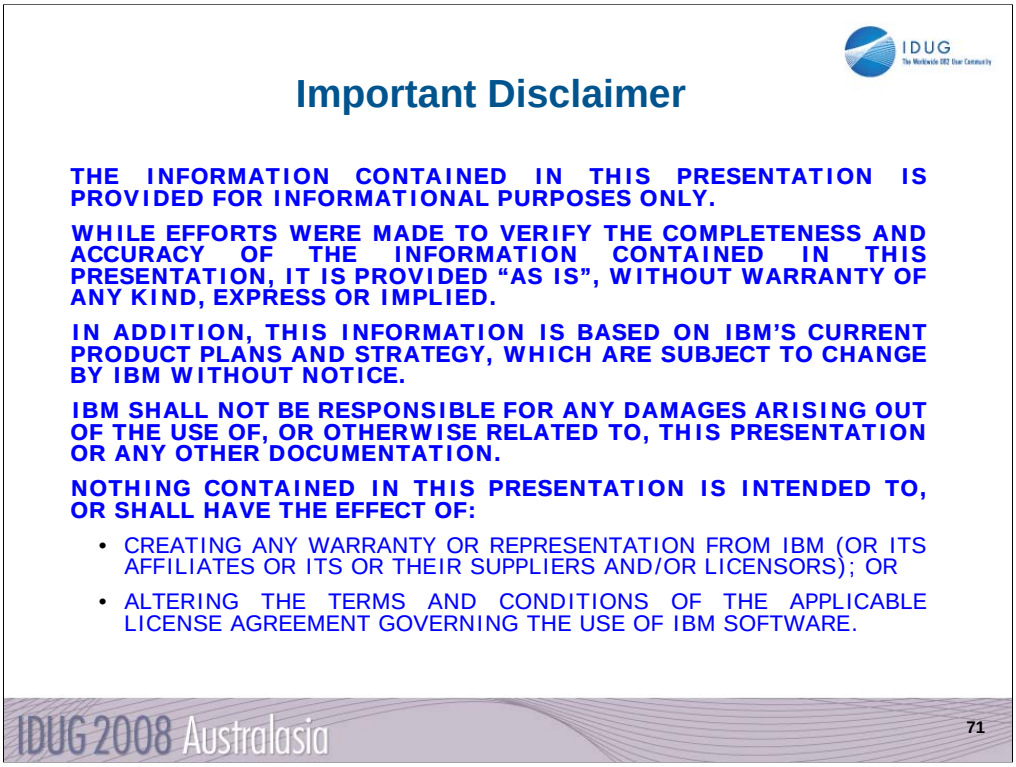

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice. Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in the operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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 and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only. This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

IBM MAY HAVE PATENTS OR PENDING PATENT APPLICATIONS COVERING SUBJECT MATTER IN THIS DOCUMENT. THE FURNISHING OF THIS DOCUMENT DOES NOT IMPLY GIVING LICENSE TO THESE PATENTS.

TRADEMARKS: THE FOLLOWING TERMS ARE TRADEMARKS OR ® REGISTERED TRADEMARKS OF THE IBM CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: AIX, AS/400, DATABASE 2, DB2, e-business logo, Enterprise Storage Server, ESCON, FICON, OS/400, Netfinity, RISC, RISC SYSTEM/6000, System i, System p, System x, IBM, Lotus, NOTES, WebSphere, z/Architecture, z/OS, System z,

The FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF THE MICROSOFT CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: MICROSOFT, WINDOWS, ODBC

Trademarks

AIX*	IBM eServer	z/VM*
CICS*	IBM logo*	zSeries*
DB2*	IMS	
DB2 Connect	On Demand Business logo	
DB2 Universal Database	Parallel Sysplex*	
DRDA*	System z	
FICON*	System z9	
GDPS*	WebSphere*	
HiperSockets	z/Architecture	
IBM*	z/OS*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of the Intel Corporation in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries.
Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.
UNIX is a registered trademark of The Open Group in the United States and other countries.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput 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 improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

Thanks for coming.