

**UNIVERSITÉ DU
MAINFRAME**

3 et 4 mai 2006

zIIP: un processeur pour la gestion des données

François Launay - Product manager hw System z
flaunay@fr.ibm.com



Four enterprise-wide roles of the mainframe system


Extending mainframe qualities of service to your enterprise

- Enterprise business resilience manager
- Enterprise security manager
- Enterprise workload manager
- Enterprise hub for data & SOA



Four enterprise-wide roles of the mainframe system

Extending mainframe qualities of service to your enterprise



- Enterprise hub for data & SOA

Raising the bar for Information on Demand



- Information on demand is essential for driving business strategies.
- Data is at the core of our customers' businesses, and must be leveraged for competitive advantage.
- For over four decades mainframes have been a leader in data and transaction serving. It's time to further leverage this asset.
- January 24th, 2006, IBM is announcing technology innovation for advanced data serving:
 - New DB2[®] function
 - New specialty engine
 - Planned future directions and roadmap

Companies Face Increasing Information Realities

60%+ of CEOs: Need to do a better job capturing and understanding information rapidly in order to make swift business decisions

Only **1/3** of CFOs believe that the information is easy to use, tailored, cost effective or integrated

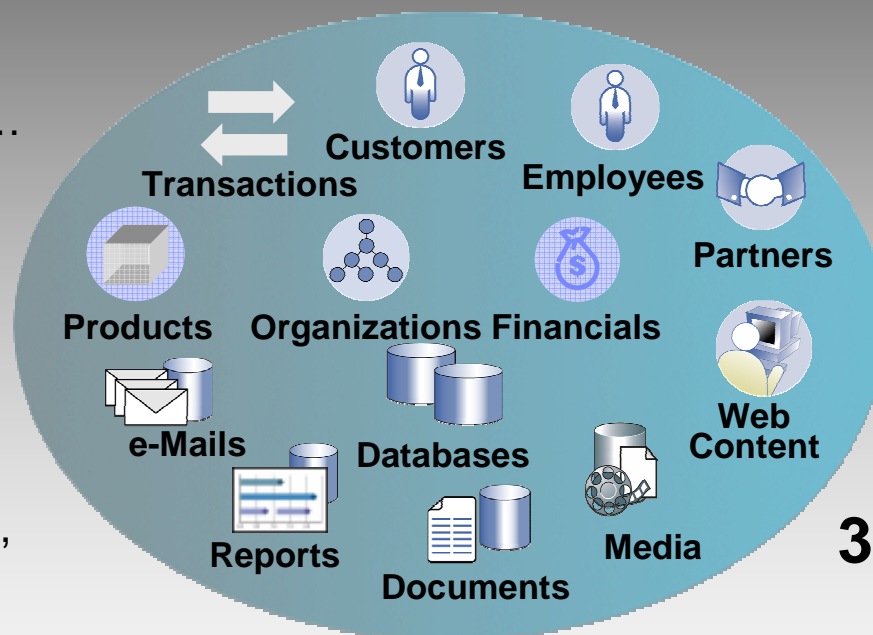
79% of companies: have 2 + repositories...
25%: have 15 +

85% of information is unstructured

17% of IT Budgets Spent on Storage HW, SW, People

37% CGR Disk Storage Growth '96-'07

122 Terabytes Disk Storage in 2005



48 disparate financial systems and 2.7 ERP systems in the average \$1B company

30-50% of design time is copy management

30% of people's time is spent searching for relevant information

40% of IT budgets may be spent on integration

Sources: IBM & Industry Studies, Customer Interviews

Where you put your Data Matters....

Confidence in System z, z/OS and DB2 for z/OS

- **Integrity**
 - z/OS® System Integrity Programming Standard – in writing
 - IBM System z™ integrity features that help protect data


- **High availability**
 - Designed with a ‘Never go down’ philosophy as opposed to a ‘rapid reboot’ philosophy
 - Capability of providing concurrent HW maintenance and upgrades and rolling changes to DB2® for z/OS (in a parallel Sysplex® cluster) can mean fewer database outages

- **Systems and database management**

- **Security**
 - Encryption, encryption, encryption – comprehensive solution
 - MLS – merge data into single server and helps preserve data isolation.
 - Helping address regulatory compliance with ability to establish centralized policies and procedures for privacy, security and audit

Technology Evolution with Mainframe Specialty Engines

- Centralized data sharing across mainframes




Internal Coupling Facility (ICF) 1997




Integrated Facility for Linux® (IFL) 2001

- Support for new workloads and open standards



System z9 Application Assist Processor (zAAP) 2004


- Incorporation of Java™ into existing mainframe solutions




IBM System z9 Integrated Information Processor (IBM zIIP)

Technology Evolution with Mainframe Specialty Engines

- Centralized data sharing across mainframes




Internal Coupling Facility (ICF) 1997



Integrated Facility for Linux® (IFL) 2001

- Support for new workloads and open standards



System z9 Application Assist Processor (zAAP) 2004

- Incorporation of Java™ into existing mainframe solutions



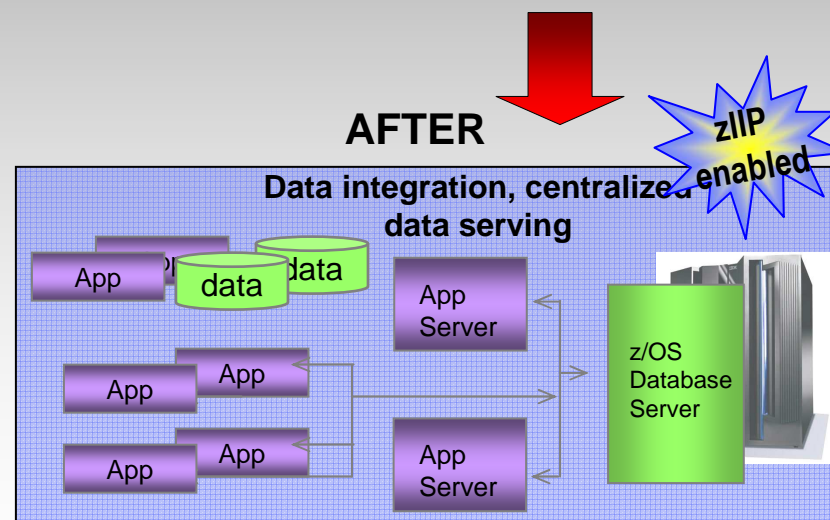
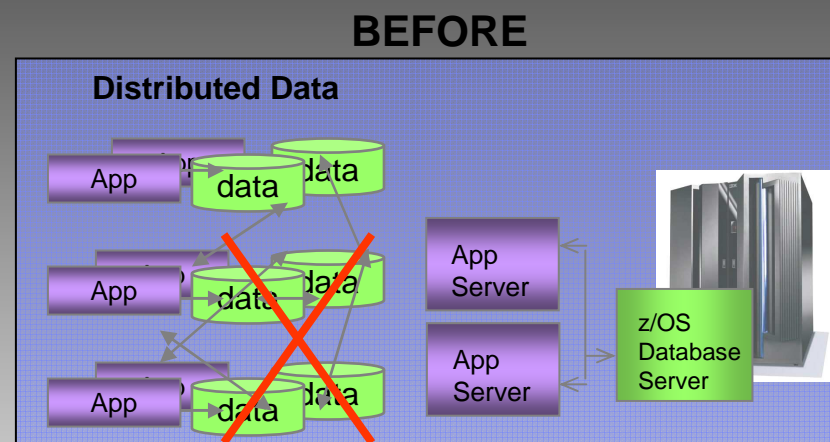
IBM System z9 Integrated Information Processor (IBM zIIP)

- Designed to help improve resource optimization for eligible data workloads within the enterprise



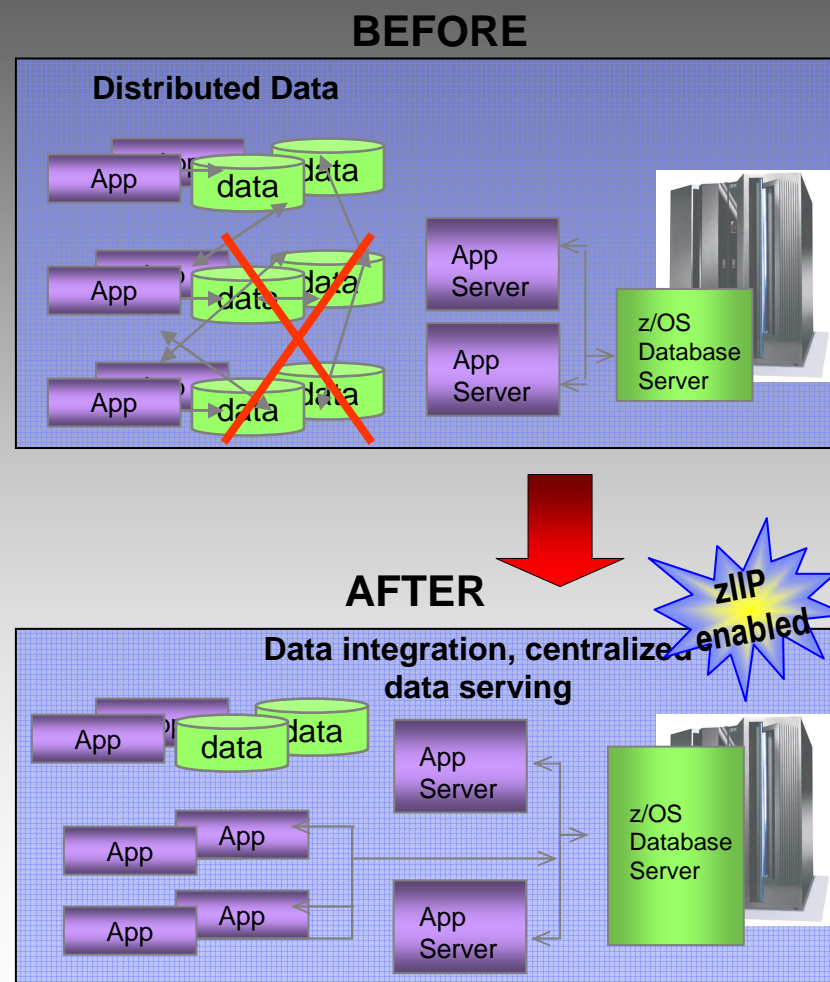
System z9 and DB2 for z/OS are an ideal data serving platform

- **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



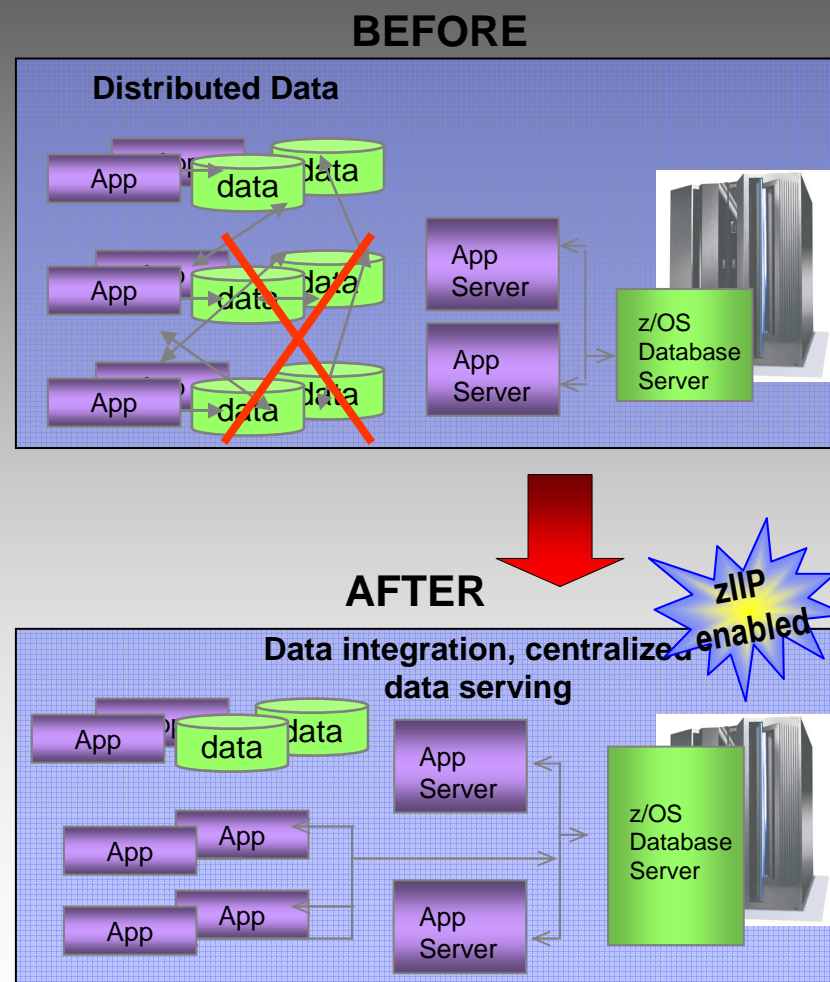
System z9 and DB2 for z/OS are an ideal data serving platform

- **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



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



New IBM System z9 Integrated Information Processor (IBM zIIP)

- New specialty engine for the System z9 mainframe designed to help:
 - Customers integrate data across the enterprise
 - Improve resource optimization and lower the cost of ownership for eligible data serving workloads
- z/OS manages and directs work between the general purpose processor and the zIIP
 - Number of zIIPs per z9-109 not to exceed number of standard processors
 - No changes anticipated to DB2 for z/OS V8 applications
- DB2 for z/OS V8 will be first IBM exploiter of the zIIP with
 - System z9 109
 - z/OS 1.6 or later
 - DB2 for z/OS V8



Types of DB2 for z/OS workloads that may benefit from zIIP



1. ERP or CRM application serving
2. Data warehousing applications
3. Some DB2 for z/OS V8 utilities

Types of DB2 for z/OS workloads that may benefit from 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 109, 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 all or 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.

What is DRDA ?

DRDA = Distributed Relational Database Architecture

–Developed by IBM

–Enables relational data to be distributed among multiple platforms

–‘any app to any db and any db to any db’

DRDA is implemented in DB2 for z/OS and reduces the need for additional gateway products that may affect performance and availability. The Open Group adopted DRDA in 1998 as the open standard for database access interoperability. An application uses DRDA application requestor or server to access a remote database. DB2 Connect is an example of a DRDA application server. The universal driver is an example of a DRDA application requestor

DRDA is network independent. It can use TCP/IP or SNA as a network protocol to flow DRDA commands. Connections using SNA are not eligible for zIIP.

So.... regarding the zIIP: If DB2 for z/OS V8 work load comes over TCP/IP and is DRDA compliant, a portion of that DB2 workload is eligible to be redirected to the zIIP –you need BOTH TCP/IP and DRDA.

What are enclave SRBs?

- z/OS dispatches work in either TCB (Task Control Block) mode or SRB (Service Request Block) mode. DB2 parallel tasks use SRB mode and are assigned the same importance as the originating address space.
- Preemptable enclaves are used to do the work on behalf of the originating TCB or SRB address space. Enclaves are grouped by common characteristics and service requests and since they are preemptable, the z/OS dispatcher (and WLM) can interrupt these tasks for more important ones (i.e. manage a transaction end-to-end). There are two types of preemptable SRBs: client SRBs and enclave SRBs.
- If the DB2 for z/OS V8 request is coming in over distributed (i.e. DRDA over TCP/IP) then most of that work is executed in enclave SRBs.
- If the request is coming over local / native connection, then that work is dispatched between TCBs, client SRBs, and enclave SRBs (star schema parallel queries and some utility index maintenance now use enclave SRBs)

So regarding the zIIP, only the enclave SRB work (not the client SRB work or non-preemptable SRB work or TCB work) is eligible to be redirected to the zIIP. DB2 V8 knows how its work is dispatched and directs z/OS 1.6 or later to dispatch (redirect) a portion of the eligible work to the zIIP.

What is Index Maintenance?

- **An index allows quick access to the rows in a table. Indexes are created using one or more columns of a table.**
- **Over time, as data in a large database is manipulated indexes can become less efficient. They need to be updated and maintained. This can be a very big task.**

–LOAD –loads your tables

–REORG –improves your index performance

–REBUILD INDEX –creates or rebuilds your indexes

So, regarding the zIIP: The BUILD phase of LOAD, REORG, and REBUILD utilities performs index maintenance. Most of the BUILD phase will be eligible to be redirected to the zIIP.

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.



DB2 DRDA zIIP offload estimation methodologies

Estimation Assumptions:

- DRDA workloads being estimated use TCP/IP connectivity
- Workloads using SNA connectivity are not eligible for offload
- Exception: star schema query parallelism child tasks
- DB2 trace instrumentation does not allow differentiation between TCP/IP and SNA connectivity

• **Methodologies for estimation :**

Choose the DB2 subsystems and the peak periods of interest for estimation

–Methodology 1:

- Using RMF Workload Activity report for environments that do not use or minimally use Stored Procedures and/or UDFs(User Defined Function)

–Methodology 2:

- Using RMF Workload Activity and DB2 Accounting reports for environments that significantly use Stored Procedures and/or UDFs

Methodology 1 : Using RMF Workload Activity Report

This methodology is applicable for environments that do not use or minimally use Stored Procedures and/or UDFs

–Because Stored Procedure and UDF run under TCB and are captured under the DDF work Service Class and they are not eligible for zIIP offload

Ensure that WLM policy is setup with Service Class(es) for SUBSYSTEM TYPE=DDF

Run the RMF Workload activity report (SYSRPTS) for the peak periods of interest showing DB2 DDF work related Service Classes

The APPL% CP under the DDF work Service Class(es) can be used to determine the DB2 enclave SRB CPU eligible to be offloaded

–APPL% represents CPU % usage relative to one CP

–Add APPL % from all the DDF Service Classes for the DB2 subsystem

–Approximately x% of this eligible CPU will be offloaded to zIIP


–Example:

•Total APPL% CP for all DDF Service Classes = 115.2%

•40% zIIP offload = $115.2 * 0.40 = 46\%$ (approximately 46% of one CP will be offloaded to a zIIP processor)

Example 1: Enterprise Applications

Enterprise Applications that access DB2 for z/OS V8 via DRDA over a TCP/IP connection will have portions of these SQL requests directed to the zIIP

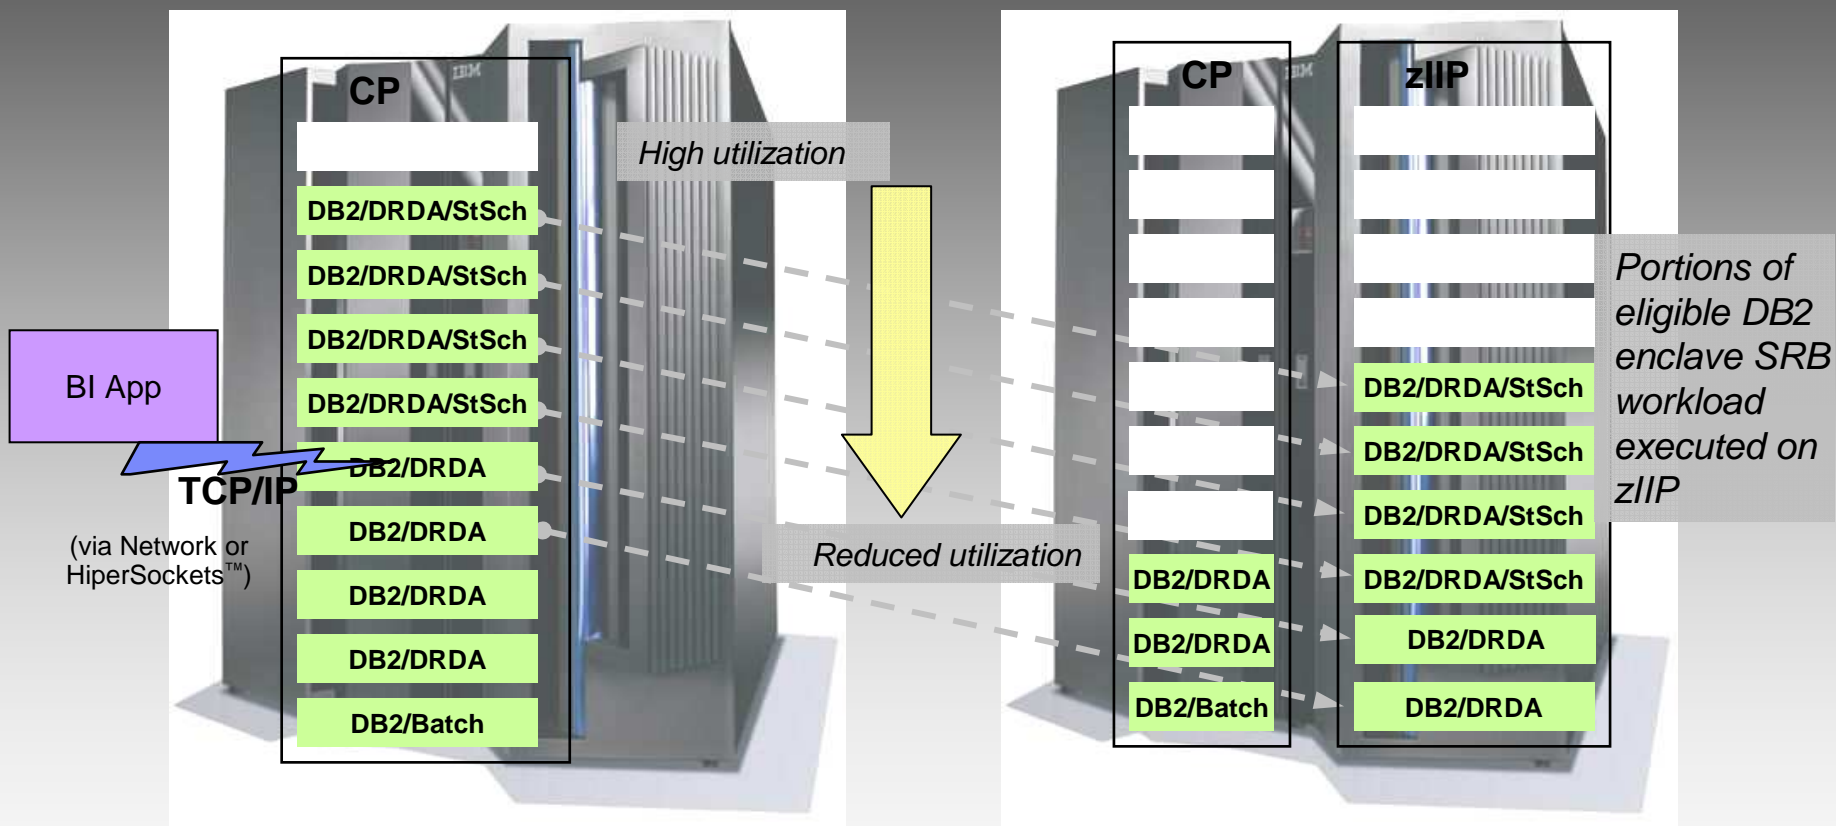


For illustrative purposes only
Actual workload redirects may vary



Example 2: Business Intelligence Applications

Complex star schema parallel queries via DRDA over a TCP/IP connection will have portions of this work directed to the 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



Example 2.5: Business Intelligence Applications (local - no DRDA)

Complex star schema parallel queries via LOCAL connection will have portions of this work directed to the zIIP




For illustrative purposes only

Actual workload redirects may vary depending on how long the queries run and how much parallelism is used


Example 3: DB2 for 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 for z/OS vNext – Addressing corporate data goals

- **Improved IT Infrastructure In Support of Compliance Efforts**
 - Trusted security context
 - Database roles
 - Auditing capabilities
 - Encryption improved
- **Simplify development and porting**
 - Many SQL improvements that simplify porting
 - Native SQL stored procedures
 - Default databases and table spaces
 - Automatic unique indexes to support primary keys
- **Decrease Complexity and Cost**
 - Fast table replacement
 - Partition by growth
 - Table append
 - Volume-based COPY/RECOVER
 - Optimization Service Center
- **Evolve Your Environment & SOA**
 - Integrated XML
 - WebSphere® integration



DB2 for z/OS vNext

Schema Evolution – Database Definition On Demand

- Fast replacement of one table with another
- Rename column and index
- Alter index to change page size
- Table space that can add partitions, for growth
- Improve ability to rebuild an index online
- Online reorganization with no BUILD2 phase
- Modify early code without requiring an IPL
- Alter table space and index logging

Performance Improvements

- CPU reductions in LOAD and REORG
- Online REBUILD INDEX
- Insert performance:
APPEND INDEX LOG
- LOB performance, function, scalability
- INDEX on expression
- SQL and optimization improvements



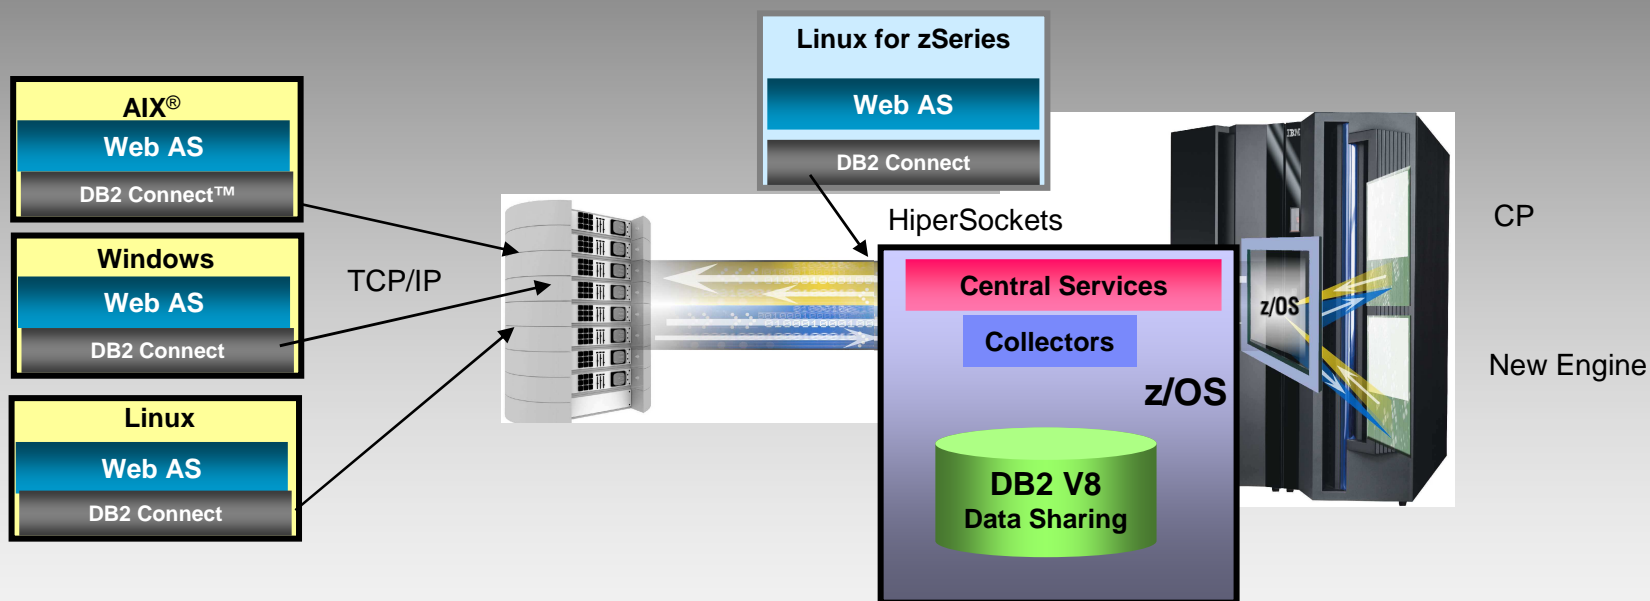
IBM System z9, z/OS & DB2 for z/OS

- IBM System z9™ 109 (z9-109) Integrated Information Processor (zIIP)
- Enhanced Cryptography
- Enhanced channels
- Faster Processors, up to 54
- More memory, better value, 64-bit virtual storage
- New backup and restore
- Multilevel Security
- Unicode conversion
- Compression
- zSeries Application Assist Processor
- z/Architecture™ new instructions
- WLM enhanced ...



An IBM zIIP data serving workload example, with SAP solutions


- **SAP Solutions** can exploit the benefits of the IBM zIIP.



Enables growth of System z9 SAP workloads through resource optimization

System z9 109 : delivering advanced data serving

- DB2 for z/OS provides a **robust database environment for advanced data serving and information on demand**
 - IBM zIIP exploitation
 - Enhancements in XML, Data Warehousing
- The new IBM System z9 Integrated Information Processor (IBM zIIP) is designed to strengthen the **System z9 mainframe as the data serving hub**, helping customers to more fully leverage their valuable assets
 - Can improve resource optimization
 - Can help lower cost of computing for eligible workloads
- In the future IBM intends to continue System z9 technology innovation with planned investments in key areas to further strengthen **System z role as the Data Hub across the enterprise for mission critical data**



A vision for System z advanced data serving

System z Enterprise Hub for Mission Critical Data

- ❖ IBM plans to continue to invest in new solutions to address customers' strategic information on demand goals*



Today's Capabilities

- Industry-leading data integrity and security
- Data sharing solution for centralized view of data
- Scalability and availability for enterprise class workloads
- Comprehensive systems and data management environment



Extension of capabilities*

- New specialty engine (zIIP) with DB2 exploitation - for mission critical ERP, CRM, and Data Warehousing workloads
- Database support improves regulatory compliance and autonomies
- Support of encryption capability (tape subsystem) with z/OS centralized key mgmt
- Data protection to achieve highest levels of security certifications



Future direction*

- Additional zIIP exploitation
- DB2 enhancements to help improve usability and reduce complexity and management costs.
- DB2 table scan acceleration via DS8000
- Support of encryption capability (disk subsystem) with z/OS centralized key mgmt
- Handle larger volumes of data, with improved scalability

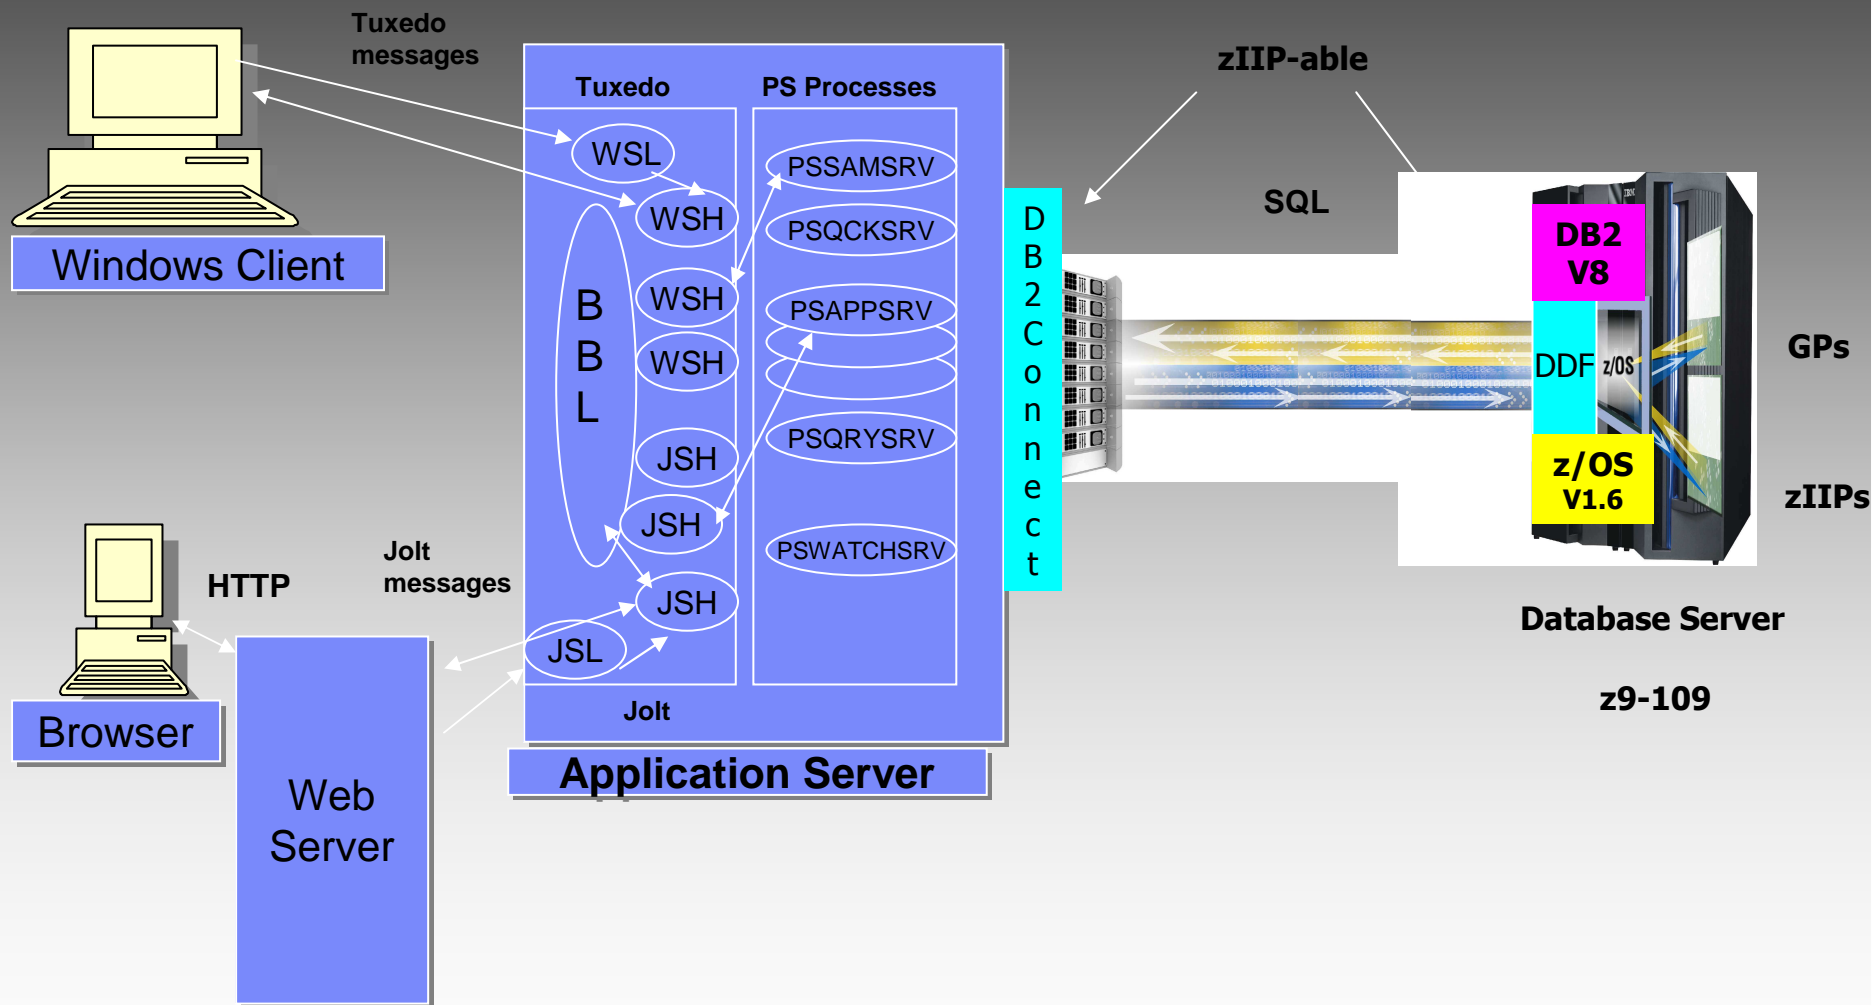
*All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

zIIP Estimation Benchmarks

- Following zIIP estimations have been made assuming:
 - Only DRDA over TCP/IP connectivity to DB2 is supported for zIIP offload
 - Application does no or small usage of Stored Procedures and /or UDF DB2 functions
- The estimation methodologies outlined here have been validated only with a limited set of Lab and customer DB2 workloads. The **methodologies may be revised** as we validate with additional DB2 customer workloads
- Because the estimation has been performed on benchmark environments (not real production systems), it represents **a maximum of what can be reached** in a customer production environment.

PeopleSoft Enterprise Example

PeopleSoft Enterprise Application Architecture V8 (Online)



PeopleSoft Enterprise Example

Customer Case 1

- Current: z890 - 450
 - PE Fin 8.9 & HR 8.3, DB2 V8, DB2 Connect + legacy CICS
 - Upgrading to HR 8.9 and adding Self Service

	System(s)	MSUs	“MIPS”
Current	z890-450	97	646
Future w/o zIIPs	Z890-460	119	789
Future with zIIPs	z9-701 w/ 1 zIIPs	81	580

PeopleSoft Enterprise Specific Information

Customer Case 1

- New configurations:
 - Without zIIPs
 - Increase MSUs to 119
 - Increase MIPS to 789
 - With zIIPS
 - 17% savings in MSUs – 32% savings from non zIIP solution
 - 10% savings in MIPS – 27% savings from non zIIP solution
 - zIIPs helped accomplish the needed upgrade with an overall **DECREASE** in MIPS and MSU

PeopleSoft Enterprise Example

Customer Case 2

- Current: z990 - 307
 - PE Fin 8.9, DB2 V8, DB2 Connect + Legacy Apps
 - Adding HR 8.9 which requires 750 Mips

	System(s)	MSUs	“MIPS”
Current	z990-307	451	2709
Future w/o zIIP	Z990-310	601	3662
Future with zIIP	z9-705 w/ 2 zIIPs	363	2633

PeopleSoft Enterprise Specific Information

Customer Case 2

- New configuration:
 - Without zIIPs
 - Increase MSUs to 601
 - Increase MIPS to 3662
 - With zIIPs
 - 20% savings in MSUs – 40% over non zIIP solution
 - 3% savings in MIPS – 28% over the non zIIP solution
 - zIIPS allowed the addition of the new module with overall **DECREASE** in MIPS and MSUs

More information:

System z web site : <http://www-03.ibm.com/systems/z/feature012406/>

zIIP FAQ: http://www.ibm.com/systems/z/faq/pdf/ziip_faq.pdf

White paper : Why Data Serving on the Mainframe:

<http://www-03.ibm.com/systems/z/feature012406/whitepaper.html>