#SHAREorg

SHARE

Transforming Your Business with IBM PureData System for Analytics

Theresa Tai
IBM Corporation
ttai@us.ibm.com

Wednesday, March 12, 2013
Session 15034

Insert Custom Session QR if Desired.

SHARE in Anaheim

IBM DB2 Analytics Accelerator

Do things you could never do before!

- What is it?
  - A high performance appliance that integrates Netezza technology with zEnterprise technology, to deliver dramatically faster business analytics
- What does it do?
  - Accelerates complex queries, up to 2000x faster
  - Lowers the cost of storing, managing and processing historical data
  - Minimizes latency
  - Improves security and reduces risk
  - Complements existing investments

IBM PureData System for Analytics
Optimized Exclusively for Analytic Data Workloads

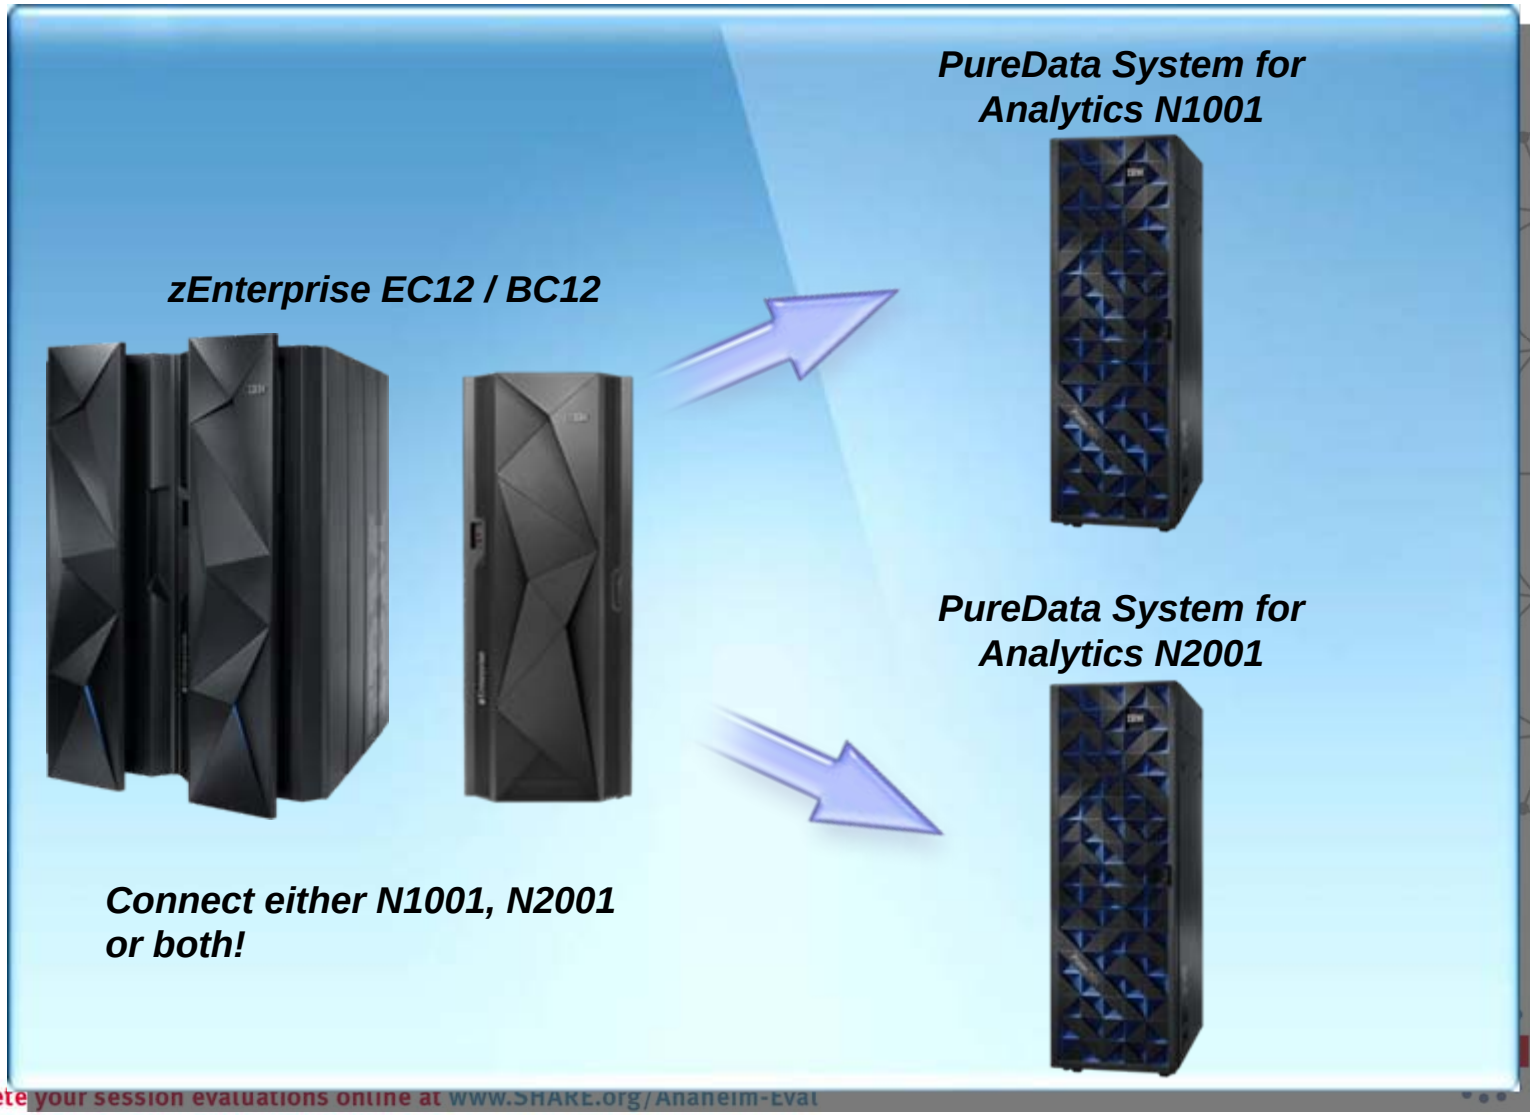

zEnterprise EC12

Extending the unmatched scalability and performance of System z

Support of System z's Newest zEnterprise

- Up to 50% system capacity performance improvement over z196 80-way
- 25% performance improvement over z196 uniprocessor
- Industry standard 8 GBps InfiniBand supporting high speed connectivity and high bandwidth
- New IBM zAware offering which monitors large quantities of message logs for smarter monitoring with high speed analytics

Coupled with the industry's newest, fastest most efficient system

DB2 for z/OS
Delivers a Single Workload-Optimized System
Foundation for Business Critical Analytics

- **Simplification**
- **High fidelity data**
- **Competitive price/performance**
- **Dynamic routing** for most efficient fit for purpose execution architecture
- Combines the strengths of both System z and Netezza technology
- Merges operational and data warehouse into a single optimized environment
- Single environment for security, logging, back-up, and recovery

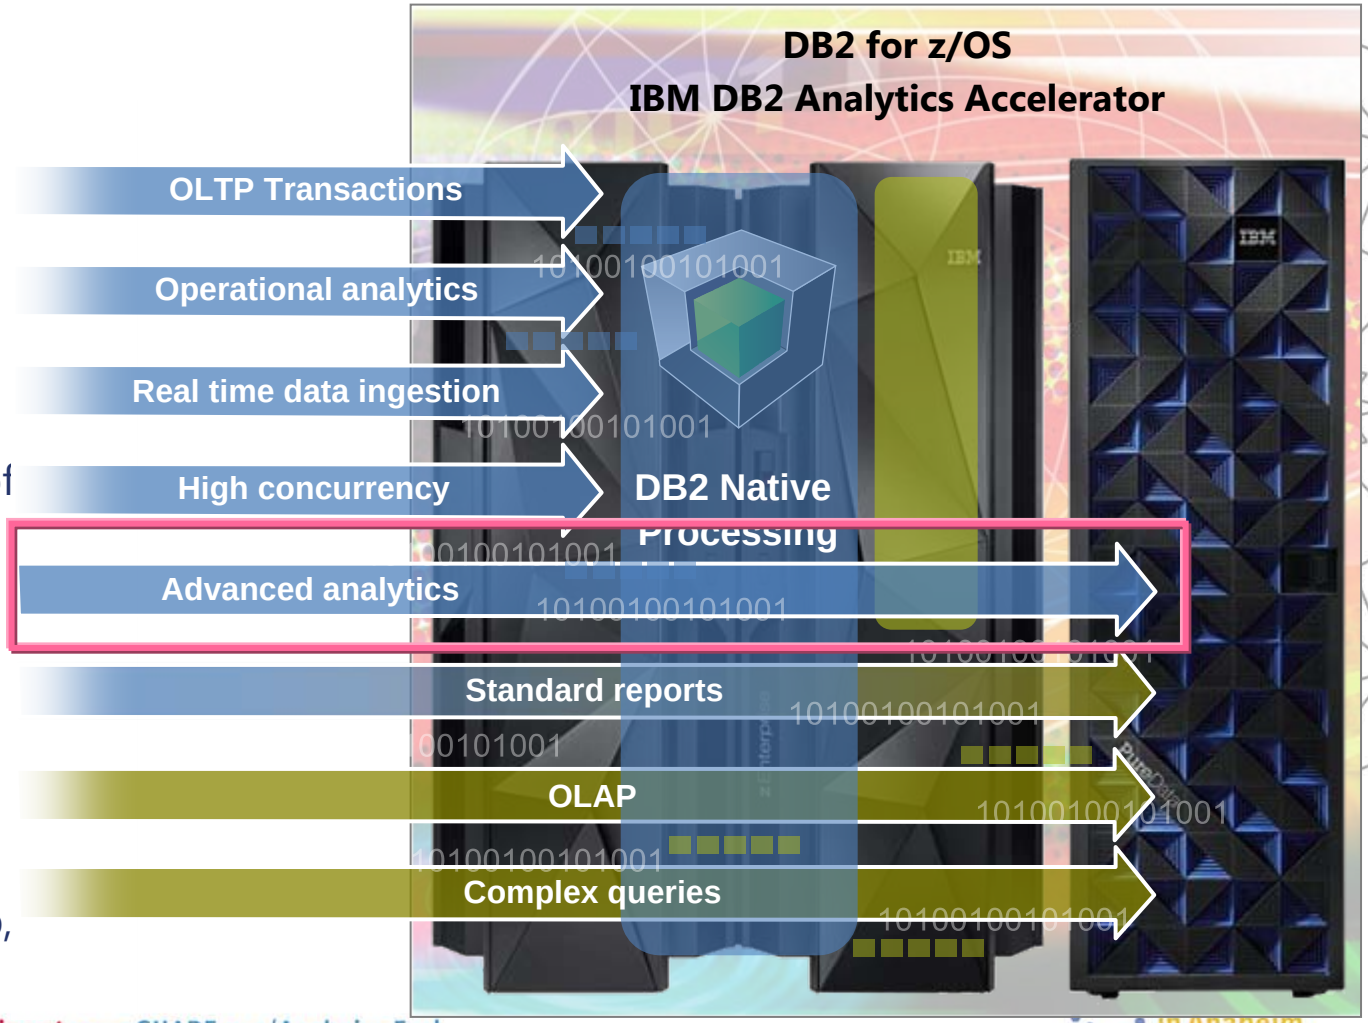

IBM DB2 Analytics Accelerator
Solution Technology at a Glance

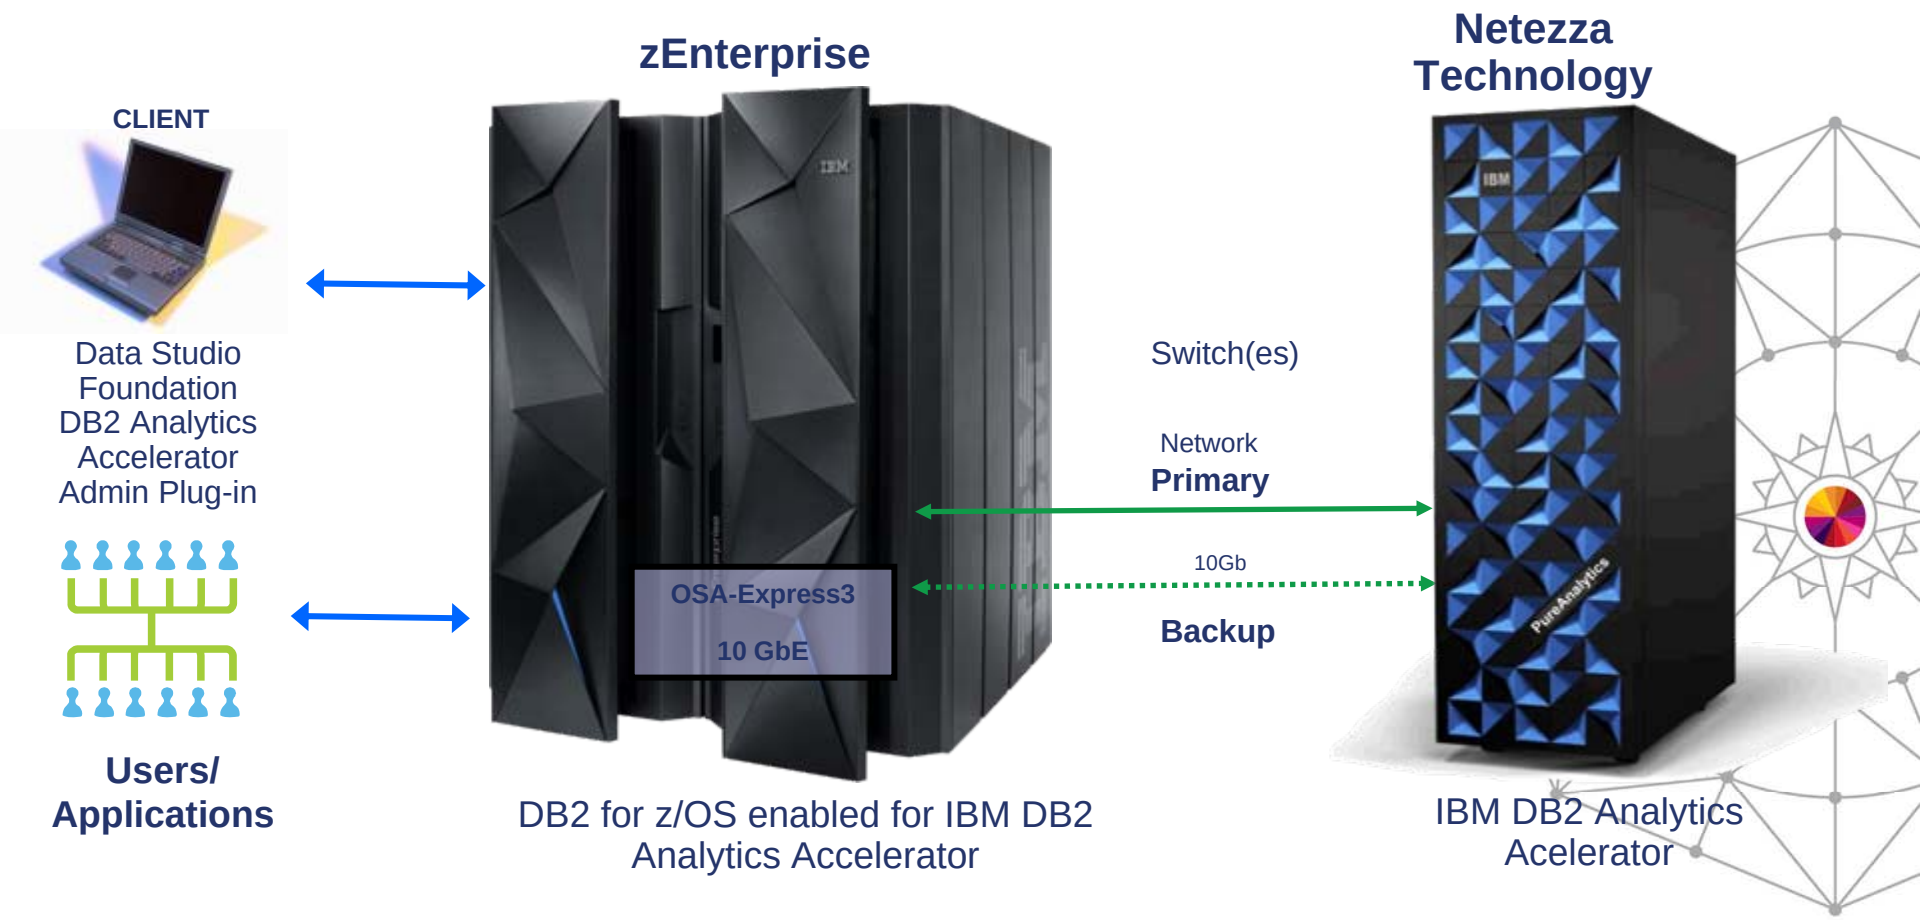

Introducing IBM PureData System N2001

The fastest performance of Netezza technology to date!

Accelerate Performance of Analytic Queries

- ***3X faster performance***[1] *for Big Data analytics*
- ***128 GB/sec effective scan rate per rack***[2] *to tackle Big Data faster*

Increase Efficiency of your Data Center

- ***50% greater data capacity per rack***[3] *helps optimize data center efficiency*
- ***More capacity and less power per rack*** *than both Oracle and Teradata*

Simplicity and Ease of Administration

- ***Improved system management and resilience*** *to spend less time managing and more time delivering value*

[1] *Based on a comparison of the IBM PureData System for Analytics N2001 to the IBM PureData System for Analytics N1001. The performance speed refers to the query times on both macro-analytic and mixed workload tests as conducted in IBM engineering lab benchmarks. The N2001 query times were an average of 3x faster than those of the N1001. Individual results may vary.*

[2] *128 GB/sec scan rate assuming an average of 4x compression across the system. Individual results may vary.*

[3] *Capacity of IBM PureData System for Analytics N2001 compared to previous generation IBM PureData System for Analytics N1001.*

N2001 Hardware Overview

12 Disk Enclosures
- 288 600 GB SAS2 Drives
  - 240 User Data, 14 S-Blade
  - 34 Spare
- RAID 1 Mirroring

2 Hosts (Active-Passive)
- 2 6-Core Intel 3.46 GHz CPUs
- 7x300 GB SAS Drives
- Red Hat Linux 6 64-bit

7 PureData for Analytics S-Blades™
- 2 Intel 8 Core 2+ GHz CPUs
- 2 8-Engine Xilinx Virtex-6 FPGAs
- 128 GB RAM + 8 GB slice buffer
- Linux 64-bit Kernel

Scales from ½ Rack to 4 Racks

- User Data Capacity: 192 TB*
- Data Scan Speed: 478 TB/hr*
- Load Speed (per system): 5+ TB/hr

- Power Requirements: 7.5 kW
- Cooling Requirements: 27,000 BTU/hr

*** 4X compression assumed**

PureData Systems for Analytics Models

	Pure Data System for Analytics **N1001**	Pure Data System for Analytics **N2001**
Blade Type	**HS22**	**HX-5**
CPU Cores / Blade	**2 x 4** Core Intel CPUs	**2 x 8** Core Intel CPUs
# Disks	96 x 3.5" / 1 TB SAS (92 Active)	288 x 2.5" / 600GB SAS2 (240 Active)
Raw Capacity	96 TB	172.8 TB
Total Disk Bandwidth	**~11 GB/s**	**~32 GB/s**
S-Blades per Rack (cores)	14 (112)	7 (112)
S-Blade Memory	**24 GB**	**128 GB**
Rack Configurations	¼, ½, 1, 1 ½, 2 – 10	½, 1, 2, 4
FPGA Cores / Blade	**8** (2 x 4 Engine Xilinx FPGA)	**16** (2 x 8 Engine Xilinx Virtex 6 FPGA)
User Data / Rack *	**128 TB**	**192 TB**

**** Assuming 4x Compression***

The PureData System for Analytics
AMPP Architecture

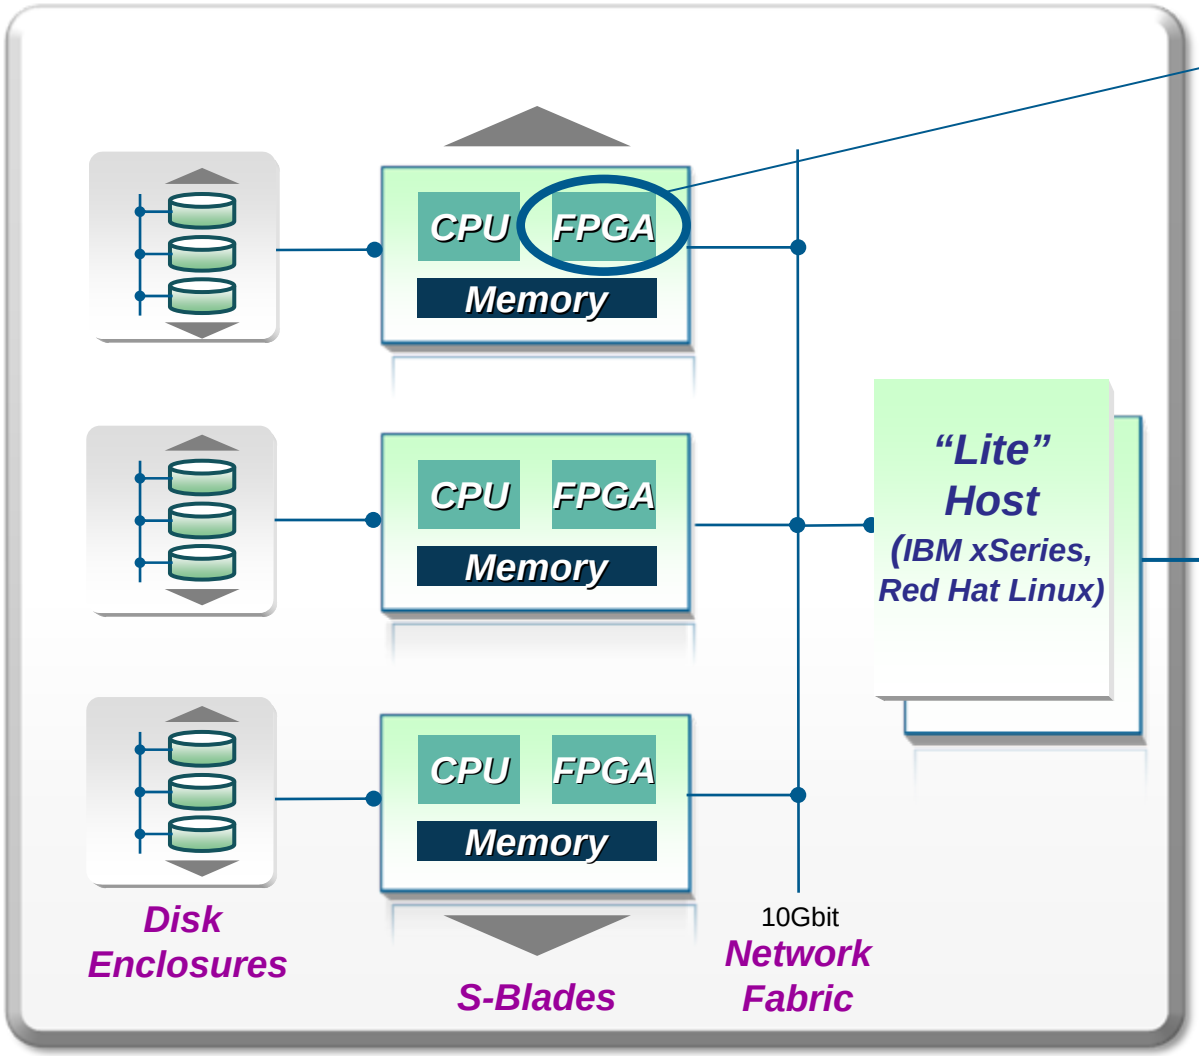

PureData System for Analytics Appliance

Field Programmable Gate Array **=**
a blank canvas until it's configured

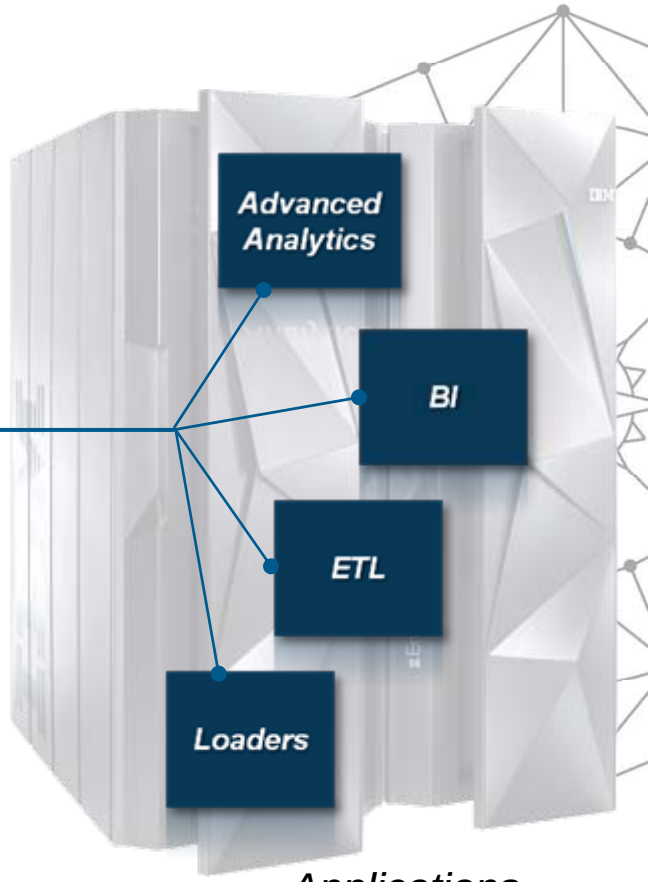

Applications

Speed Through Hardware Acceleration

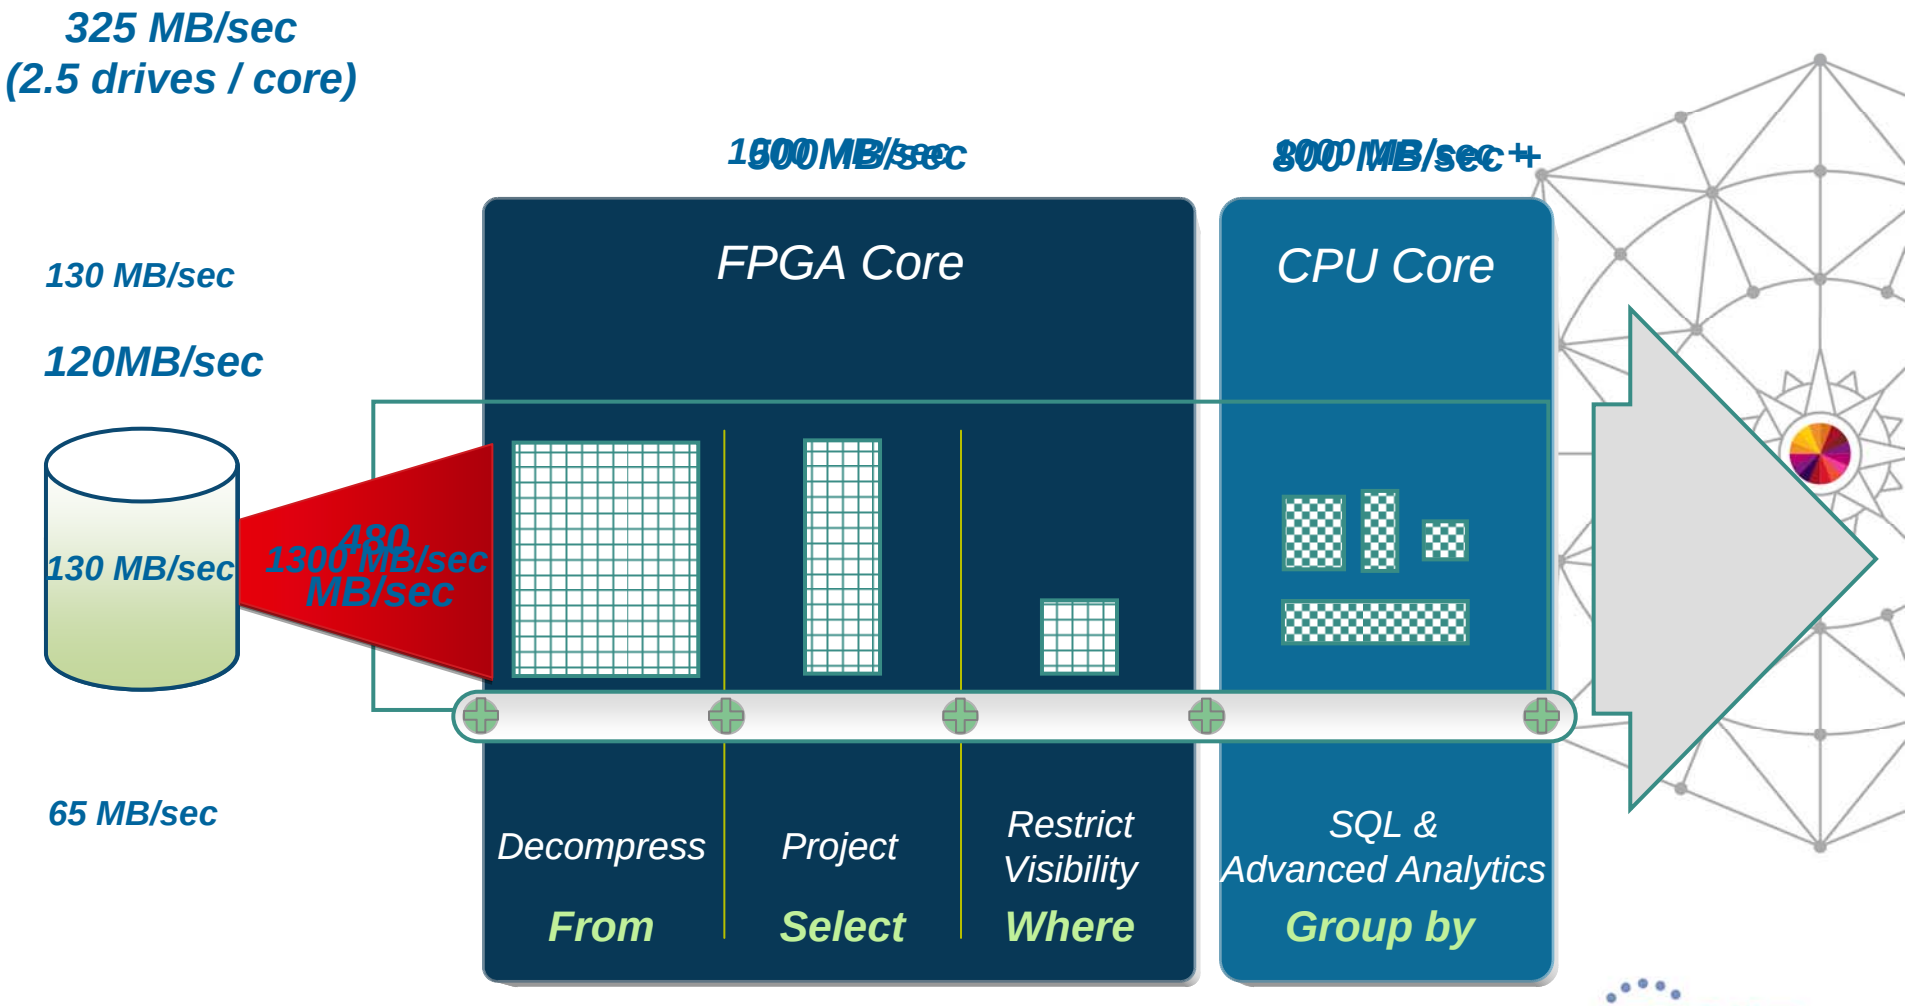

ACID (Atomicity, Consistency, Isolation and Durability) compliance

S-Blade Data Stream Processing

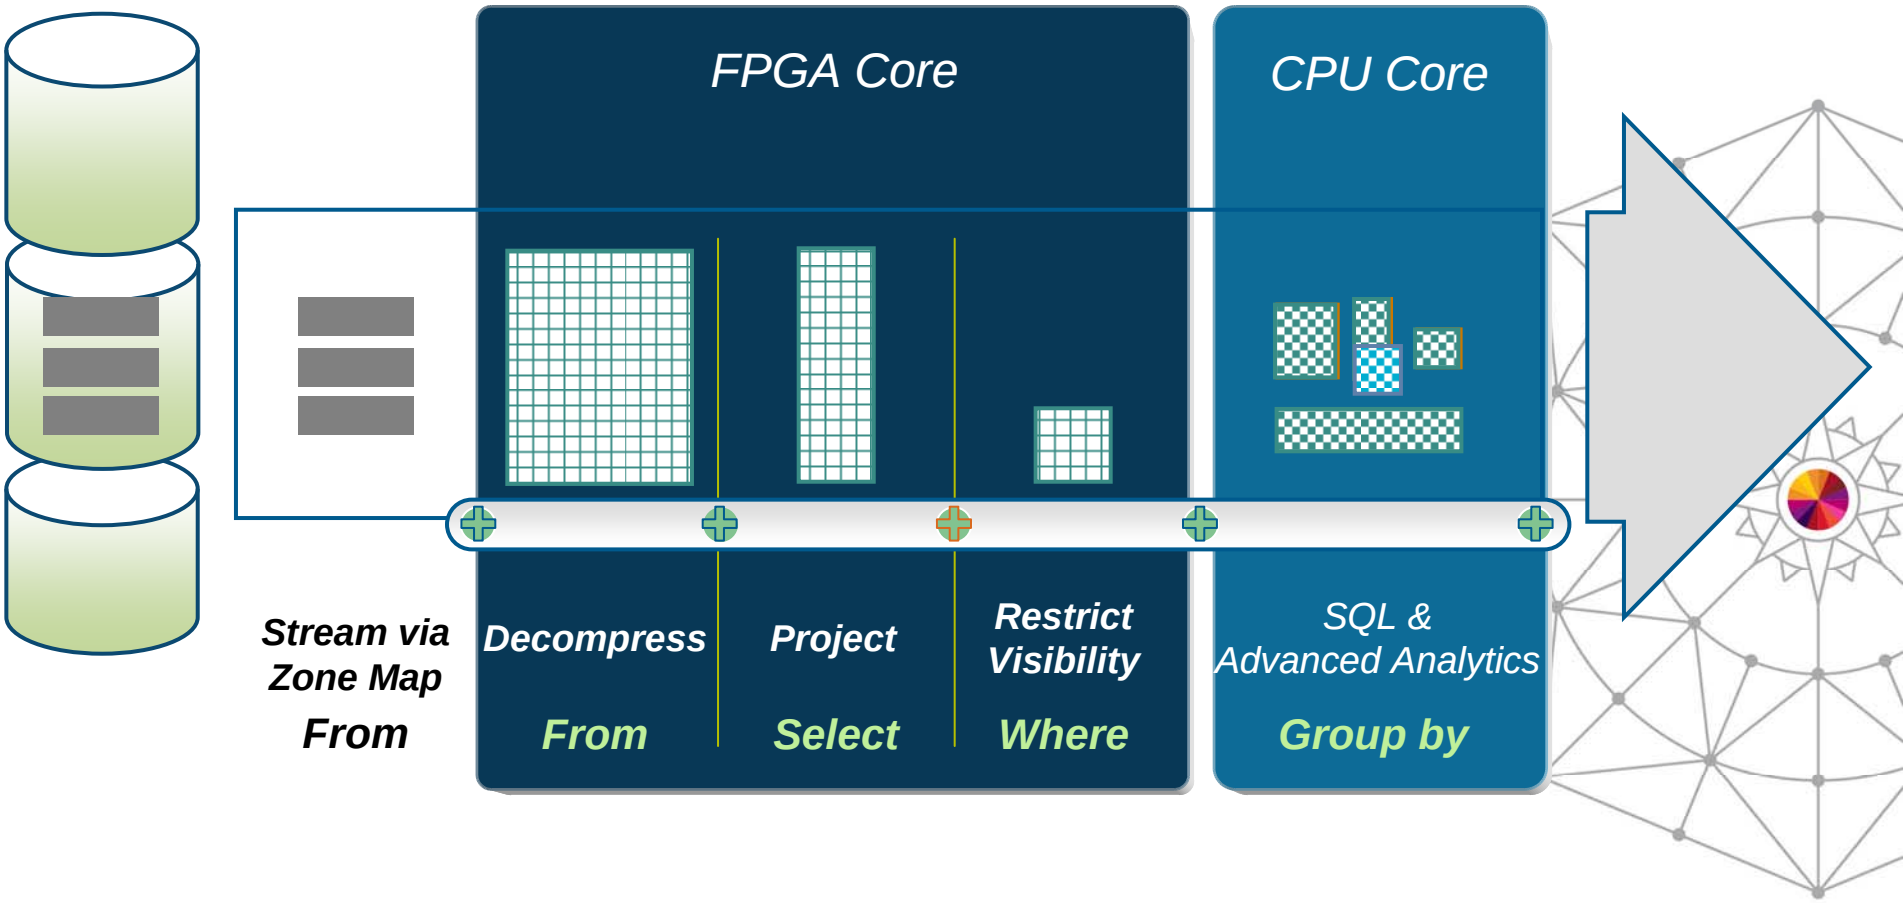

Select State, Age, Gender, count(*) From BillionRowCustomerTable Where BirthDate ‘01/01/1960’ And State in (‘FL’, ‘GA’, ‘SC’, ‘NC’) Group by State, Age, Gender Order by State, Age, Gender

Architectural Efficiency

Asymmetric Massively Parallel Processing™ (AMPP)

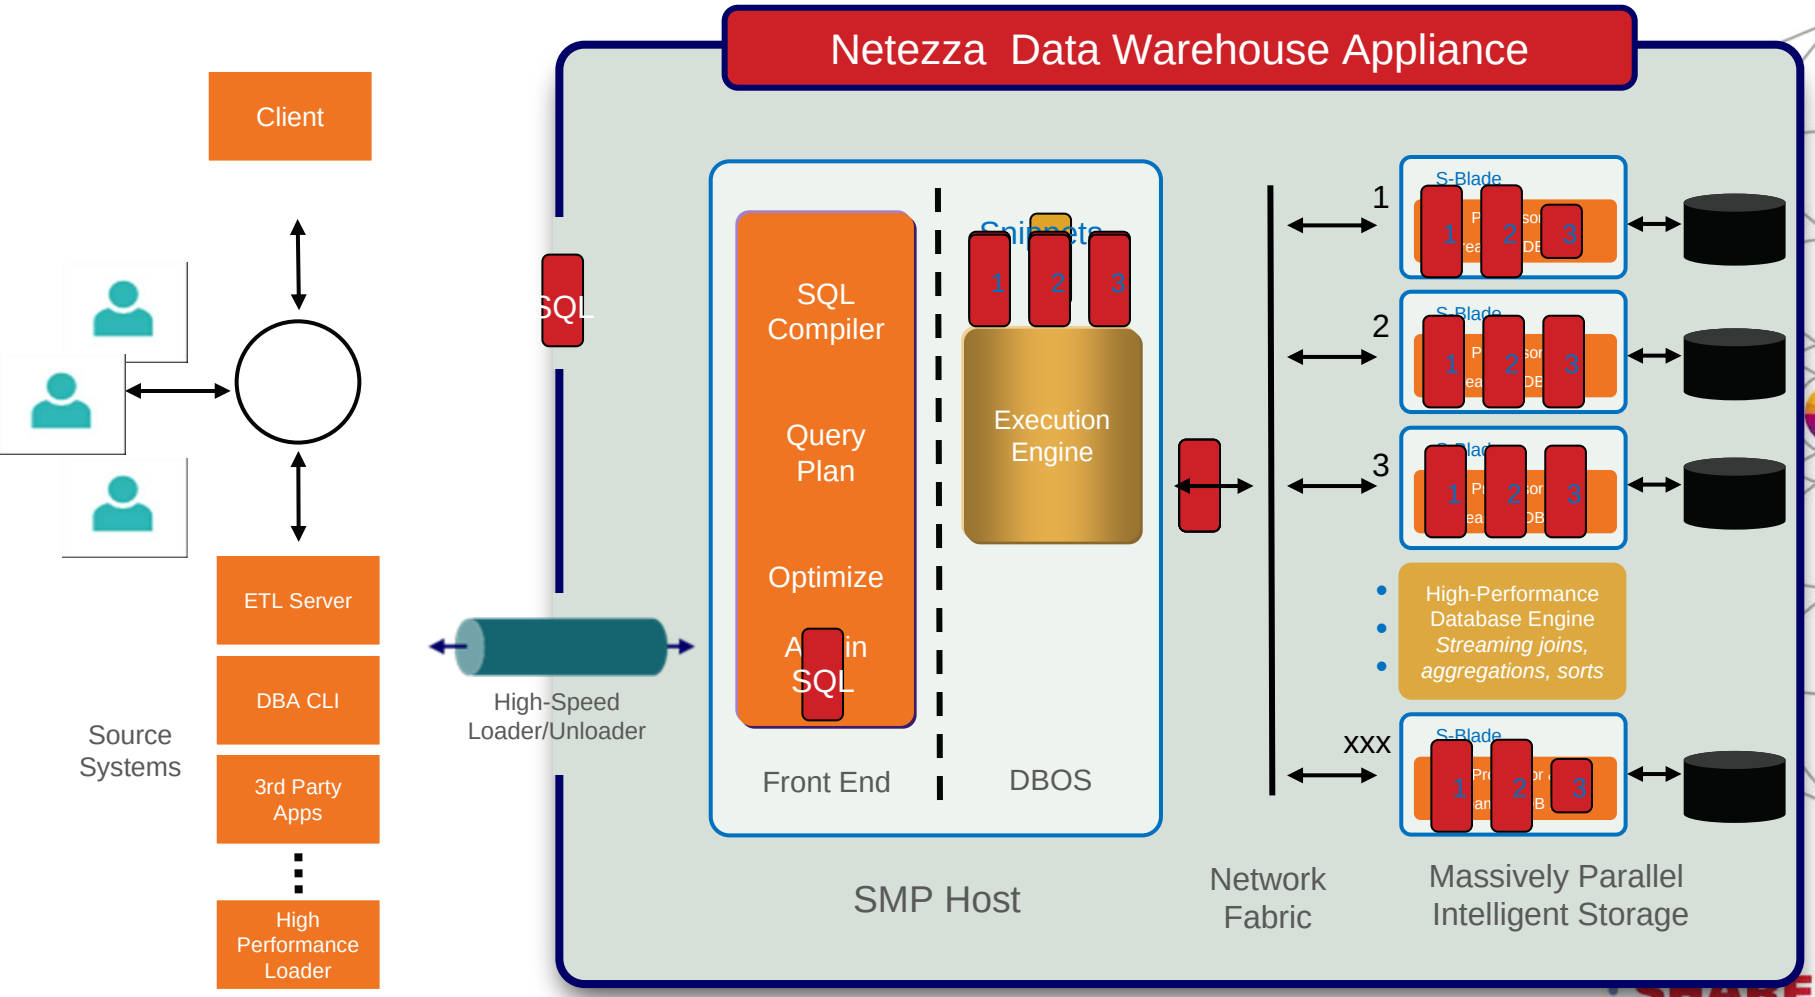

Asymmetric Massively Parallel Processing™

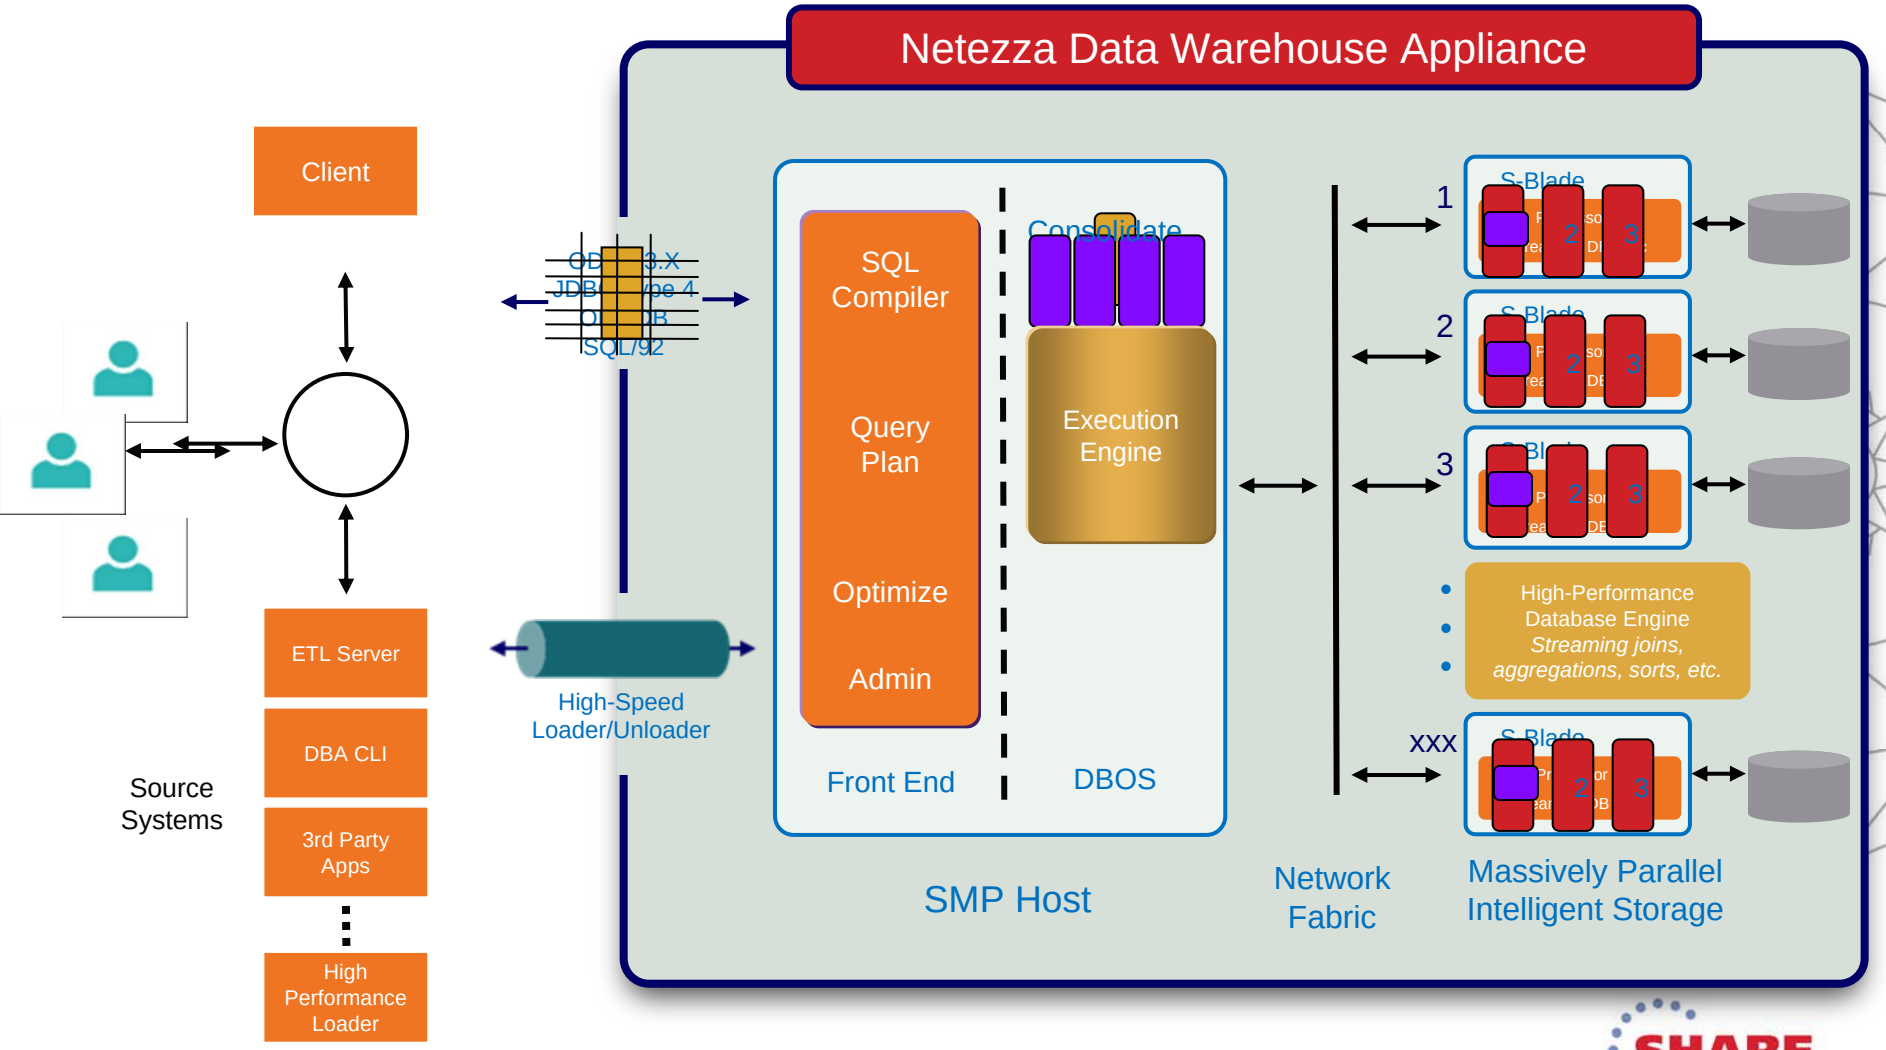

Disaster Recovery Scenario

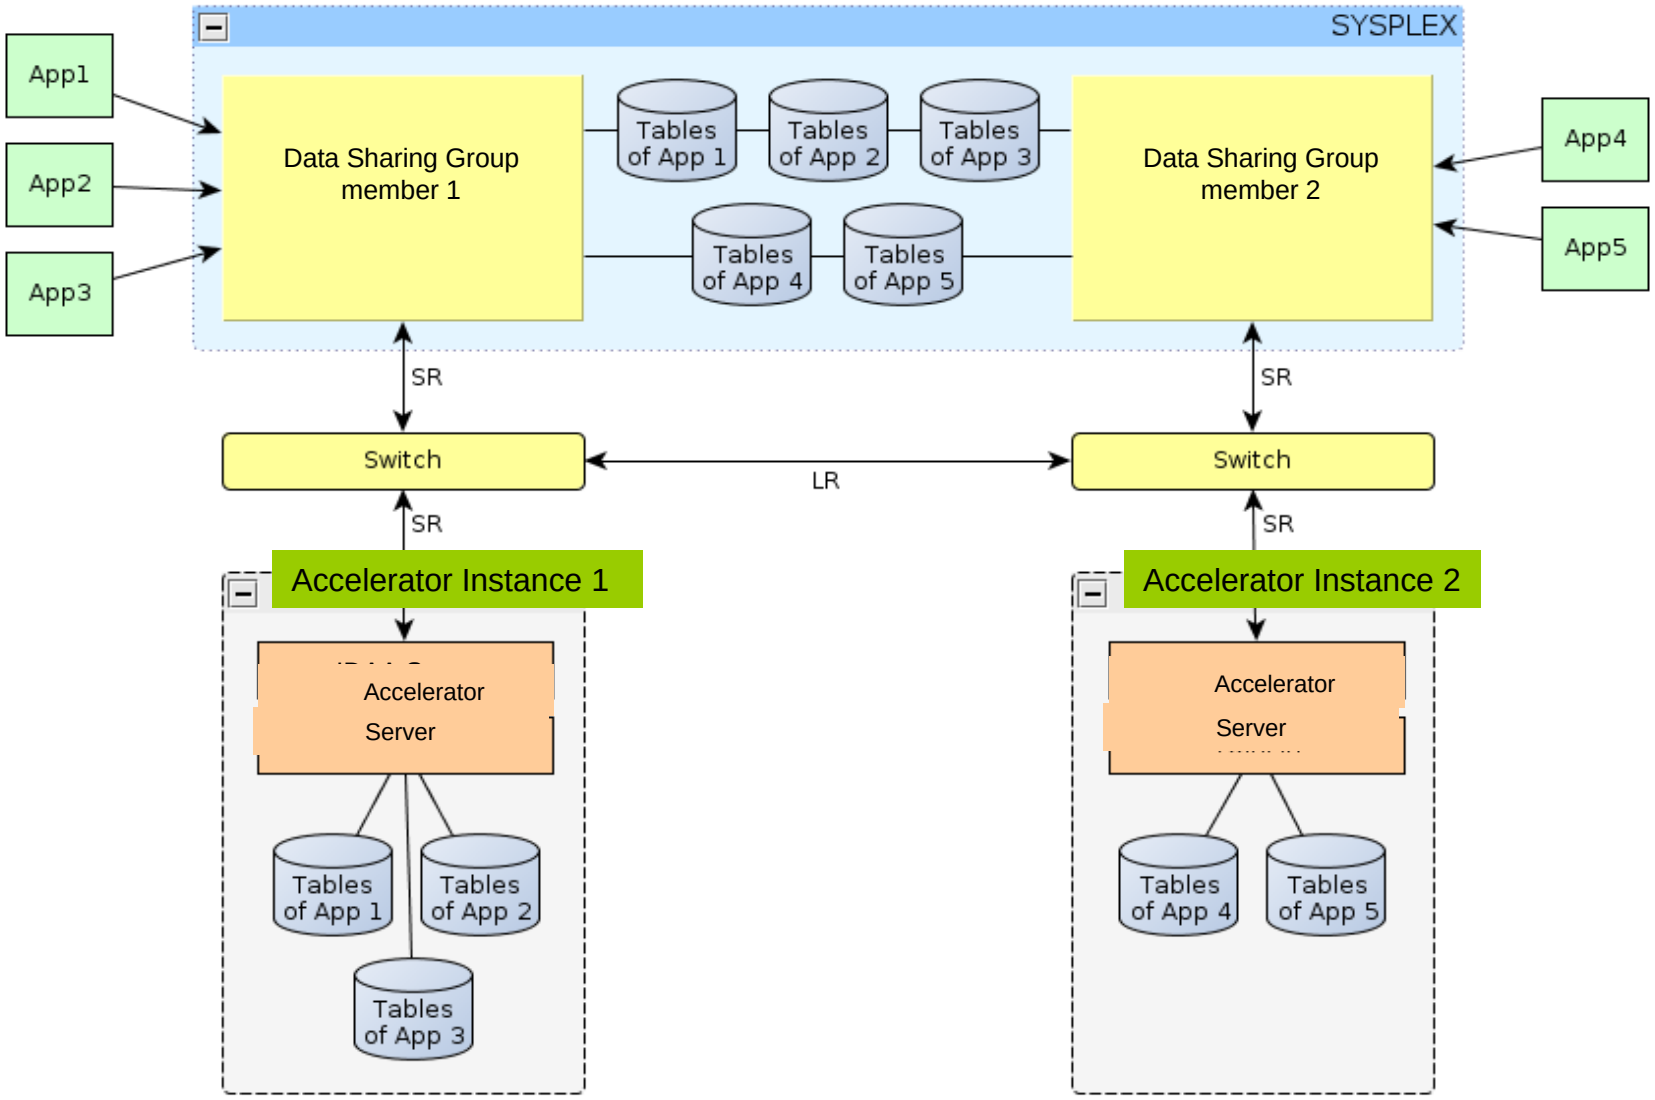

Disaster Recovery in Action

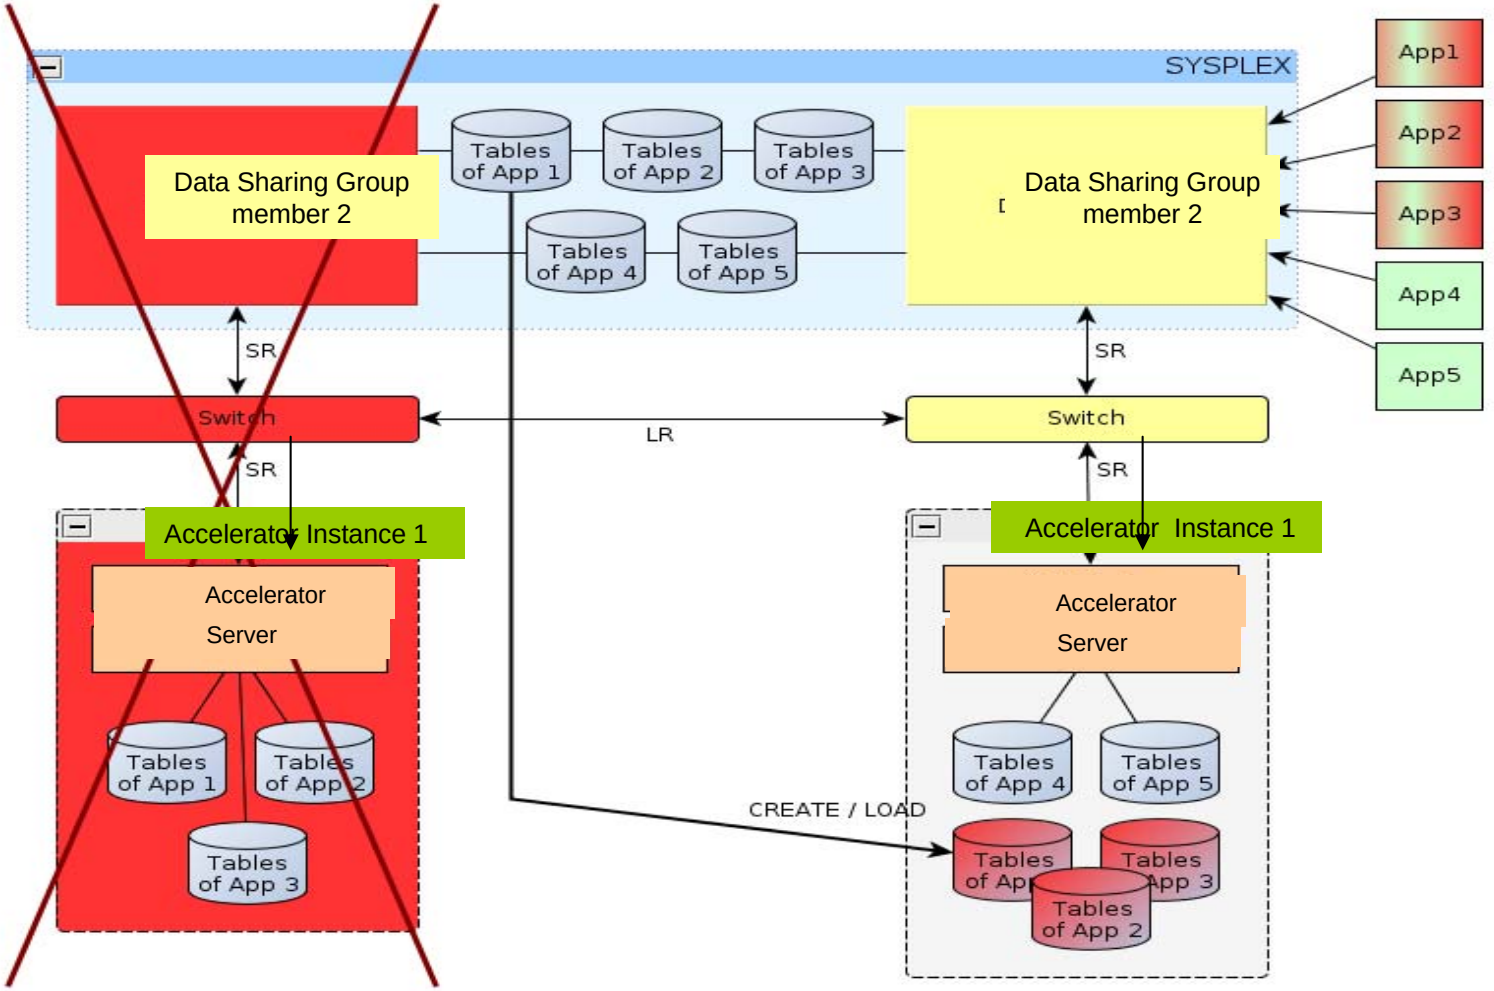

Query Execution Process Flow

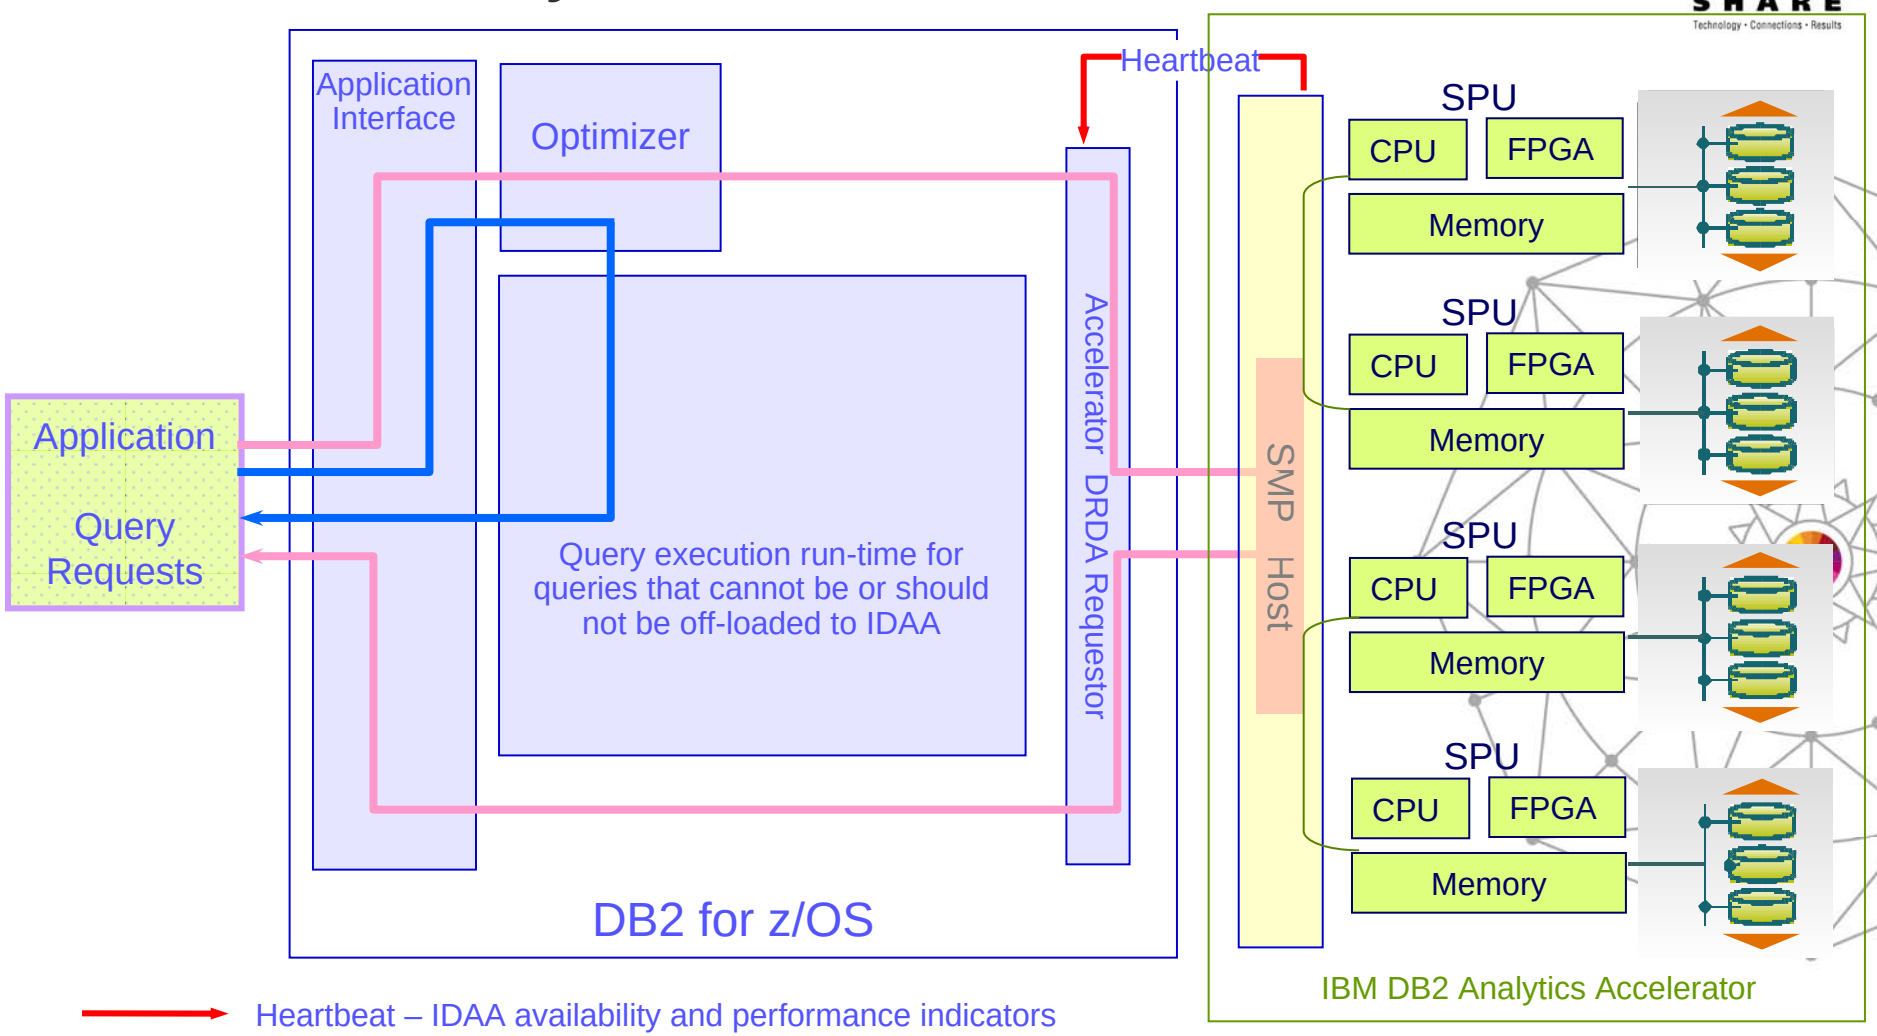

IBM DB2 Analytics Accelerator V4.1

Building on the Core Values

Helping organizations change the way they do business

Fast	Cost Saving	Appliance
Complex queries run up to 2000x faster while retaining single record lookup speed	Eliminate costly query tuning while offloading complex query processing	No applications to change, just plug it in, load the data, and gain the value

Expanding the Value of High Speed Analytics

- **Enable Acceleration of more queries**
  More SQL query constructs are now eligible for routing to the accelerator
  - ⇨ Static SQL Query & Rowset query offload and multi row fetch support
- **Enterprise Robustness & Scalability**
  More Efficient use of the accelerator appliance
  - ⇨ Support for multiple accelerators with workload balancing
  - ⇨ Incremental update support of multiple DB2 subsystems
- **Enhanced ease-of-use capabilities**
  Enhanced Storage Saver Solution
- **Increase DB2 Analytics Accelerator Transparency**
  Easier operation and maintenance
  - ⇨ Faster initial load
  - ⇨ Faster refresh and automated hardware system updates
- **Support new PureData and DB2 Technology**
  Full range of IBM PureData for Analytics models
  - ⇨ N1001 and N2001 supported
  - ⇨ DB2 11 support

IBM DB2 Analytics Accelerator V4.1 Highlights

- **More Acceleration – Static SQL Query support**
  - More SQL query constructs are now eligible for rerouting to the accelerator
  - Multi-row fetch for increased throughput
- **Enterprise Robustness – HA/DR**
  - Support for multiple Accelerators with group workload balancing
  - Workload balancing between Accelerators
  - Improved Workload Management for high priority work
- **Enhanced ease-of-use capabilities**
  - Improved High Performance Storage Saver (HPSS) Solution
    - HPSS partition read only
    - HPSS Restore function
    - HPSS Image Copy
    - Enhance Monitoring
- **Increase DB2 Analytics Accelerator Transparency**
  - Faster initial load
  - Refresh and automated hardware system updates
  - Automated Netezza Kit Installation
- **New Capacity**
  - Full range of IBM PureData for Analytics - models N1001 and N2001
  - Exploitation of new N2001 hardware features
    - *Accelerate Performance of Analytic Queries*
    - *Increase Efficiency, System Management and Resilience*

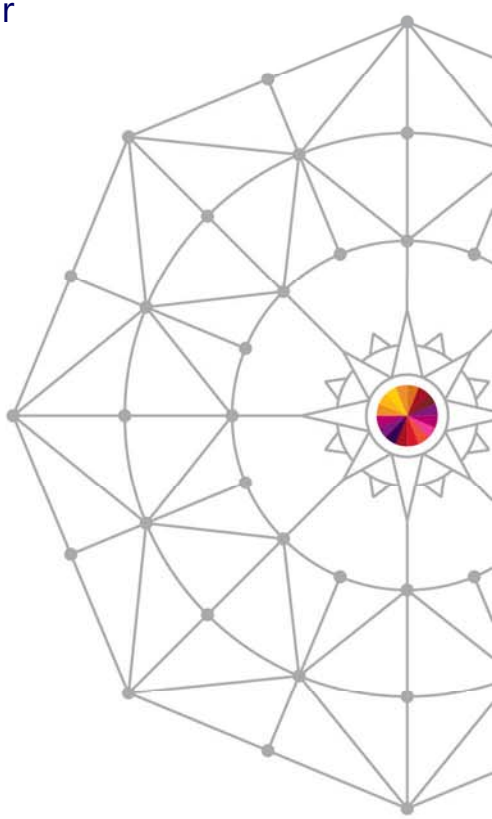

Extend Eligible SQL Query Constructs for Acceleration

- Static SQL Support
  - Static queries are now eligible to be routed to the Accelerator
  - New QUERYACCELERATION, GETACCELARCHIVE option
    - Same values for zparms QUERY_ACCELERATION and GET_ACCEL_ARCHIVE
  - Acceleration for static queries is determined and fixed at bind package time
  - Tables must be added and enabled or archived on accelerator before binding the package
  - Accelerator must be active and started when static query runs otherwise the execution fails
- Rowset query offload and multi-row fetch support for local applications
  - CPU reduction for retrieving data from a local accelerated query
  - For dynamic SQL in local applications:
    - Specify WITH ROWSET POSITIONING
    - To fetch use a FETCH NEXT ROWSET with FOR N ROWS clause
- More DB2 functions and data types supported
  - BITAND(), TIMESTAMPDIFF()
  - Enable comparison between different datatypes
    - e.g. VARCHAR and INTEGER

Static Query Acceleration Options

- on **active** data with ELIGIBLE or ENABLE option
- on **active** data with acceleration option ALL Static query acceleration
- on **active** data with acceleration option ENABLEWITHFAILBACK
- on **archived** data with
  - GETACCELARCHIVE=YES
  - QUERYACCELERATION<>NONE

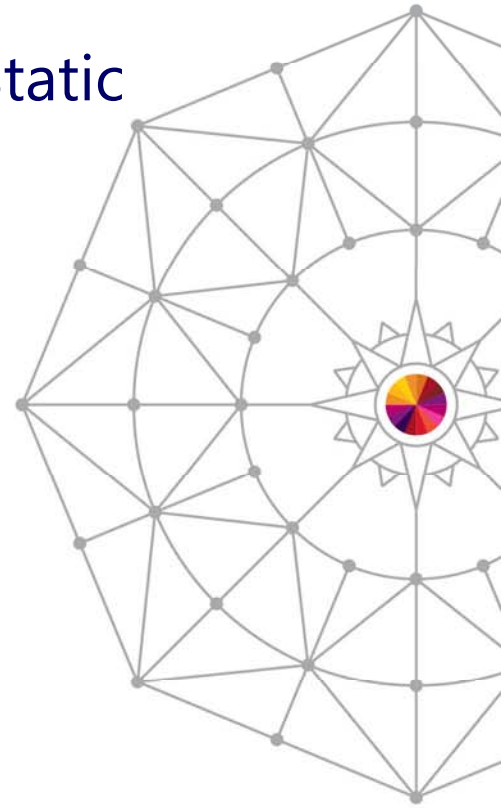

More Efficient Use of the Accelerator Appliance

- V4.1 Ensures all Defined Accelerators are well Balanced based on Utilization
  - V3.1 the query were routed to the **first Accelerator** that matched the routing criteria
  - Could lead to an unbalanced accelerator utilization
- Workload Balancing for Multiple Accelerators
  - DB2 will balance query routing between qualifying Accelerators based on utilization
  - HA workload balancing
- Multiple Accelerators are Defined and Available to a DB2 Subsystem
  - A query acceleration decision is based on **Accelerator Utilization Information**
  - Distribute workload across the lowest utilizatin value

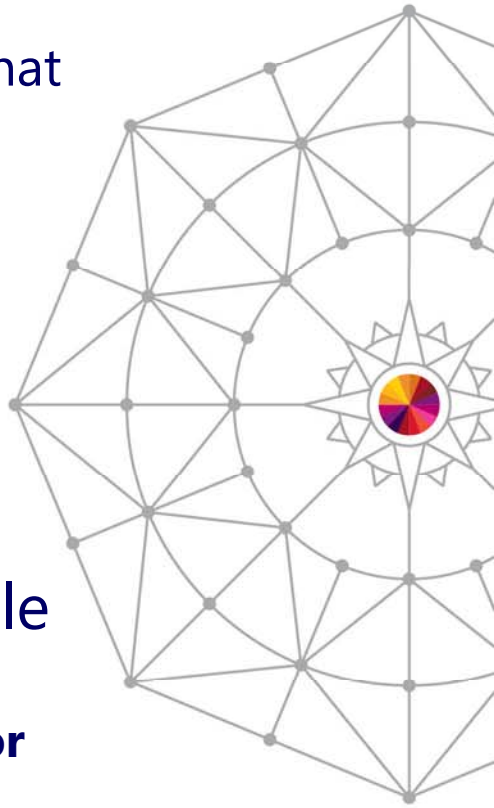

Setups Where Workload Balancing Applies

A single DB2 system connected to multiple accelerators

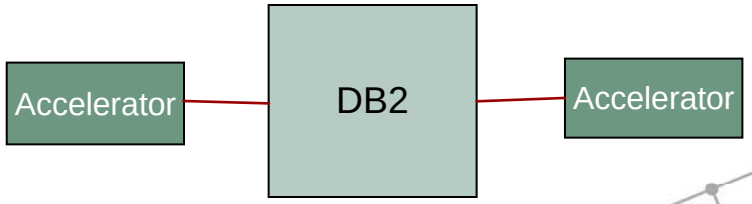

Multiple DB2 systems connected to multiple accelerators

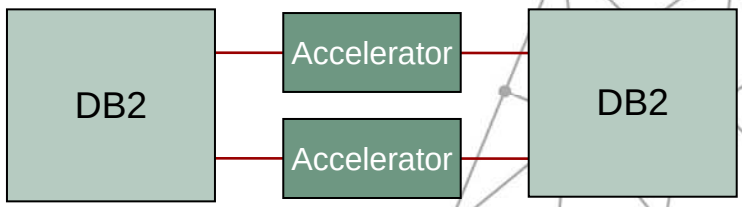

Before routing a query to an accelerator DB2 checks the utilization of each eligible accelerator to route the query

Requirements for Accelerators:

At least two Accelerators must have V4.1 installed in order that workload balancing is used.

Full Flexibility for DB2 Systems:

- residing in the same LPAR
- residing in different LPARs
- residing in different CECs
- being independent (non-data sharing)
- belonging to the same data sharing group
- belonging to different data sharing groups

Same as Accelerator V3.1

Workload Balancing – HA Environment with Data Sharing group

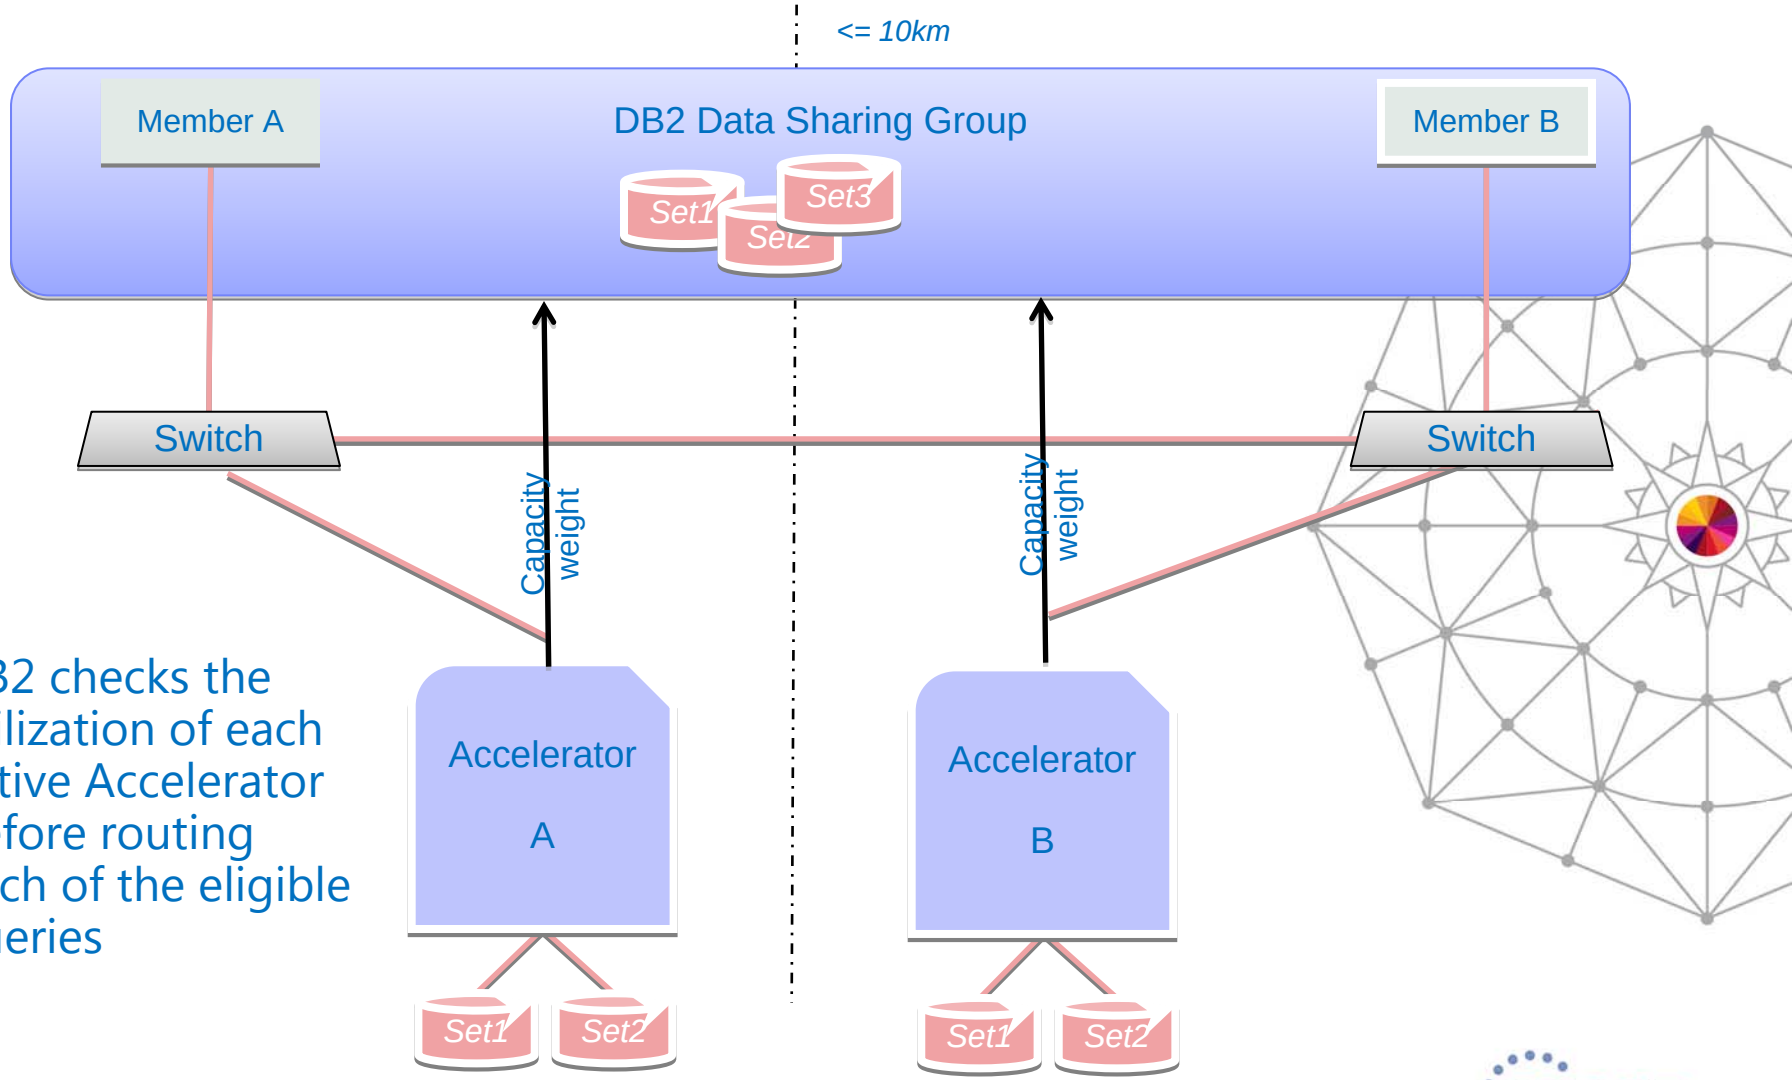

DB2 checks the utilization of each active Accelerator before routing each of the eligible queries

Workload Balancing - HA Environment with Data Sharing group

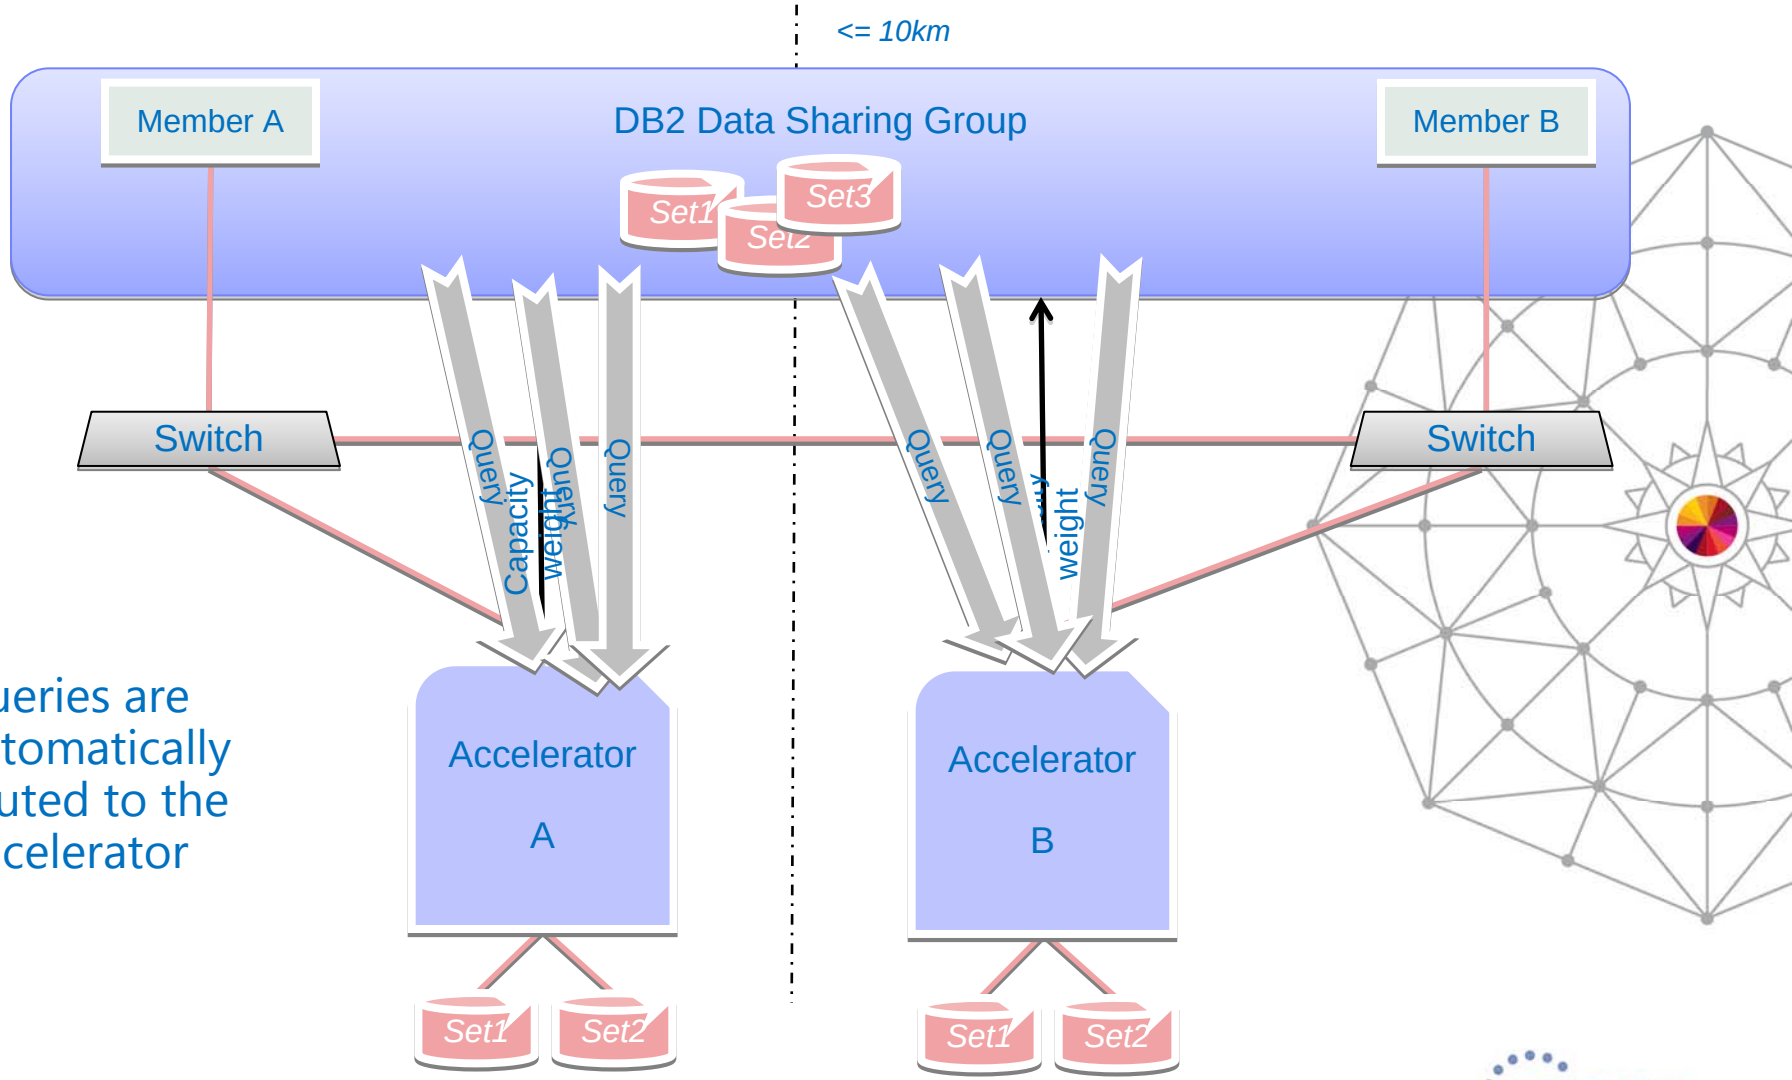

Queries are automatically routed to the Accelerator

"Accelerator Utilization Information"

Recommended HA Topology

- High availability through redundant components:
  - OSA cards, network connections, switches, and Accelerators
- Data is loaded into A1 and A2, a Load operation can run from any member of the Data Sharing Group, tables are enabled after successful load
- Incremental Update may be used to update table on both A1 and A2
- Any DB2 Data Sharing Group member can access either or both A1and A2
- Automatic selection of Accelerator based on capacity weight and availability among defined/active Accelerators

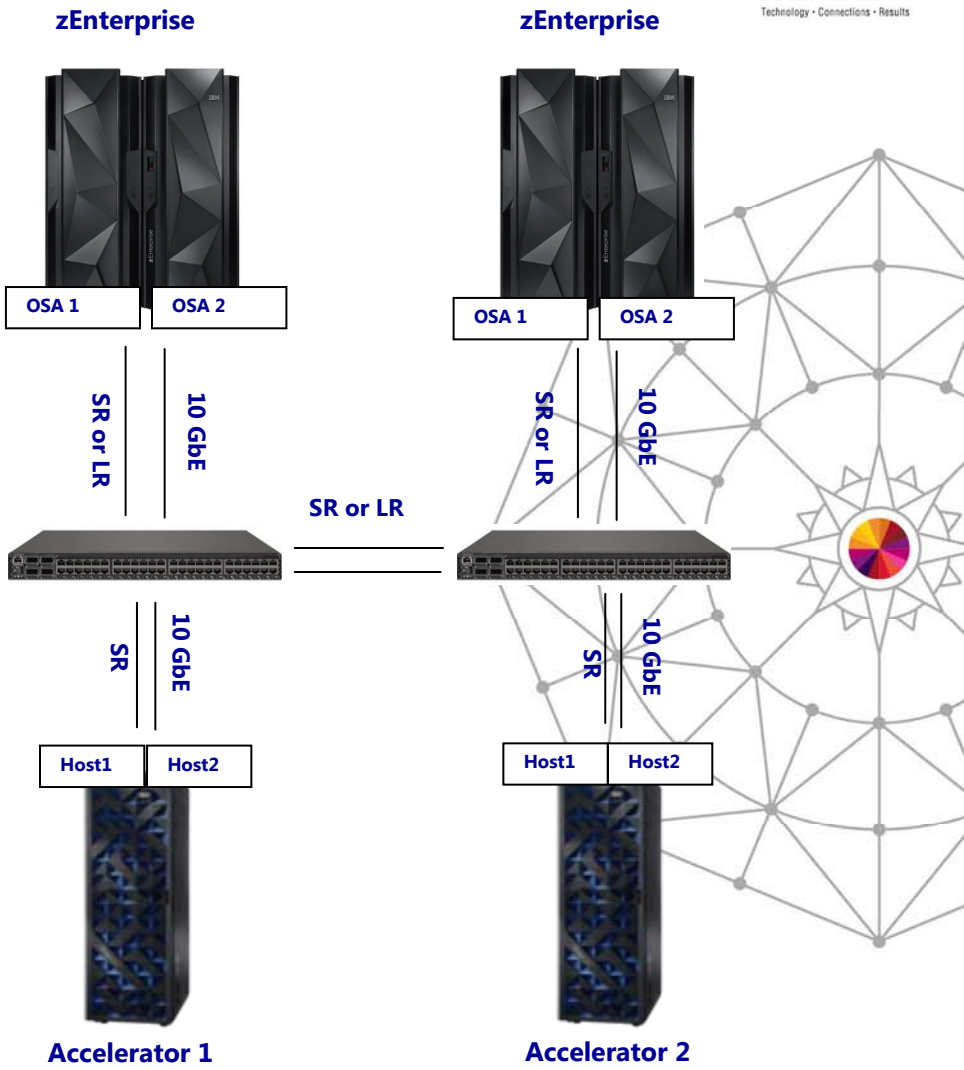

Key: A - Accelerator

End to End HA/DR

- The IBM DB2 Analytics Accelerator consists of multiple components that contribute to High Availability inside the physical machine itself
- These components are inherited from the underlying IBM PureData System for Analytics architecture and include:
  - Netezza Performance Server (NPS) hosts
  - 10 GBit dual-port Ethernet interfaces
  - Redundant S-Blades
  - Redundant Array of Inexpensive Disks (RAID) technology
- In addition to "built-in" capabilities of the appliance, there are additional software components and features to build even more extensive HA and DR concepts
- These concepts serve as the foundation to build advanced configurations and integrate them into existing HA/DR environments with System z
  - Workload Balancing
  - Data Maintenance and synchronization with multiple accelerators - High Availability setup for incremental update
  - High Performance Storage Saver (HPSS) and multiple Accelerators

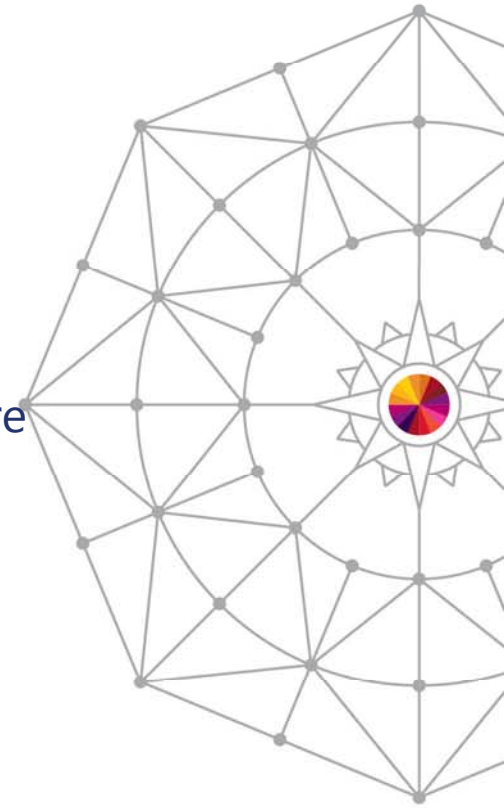

More Efficient Use of the Accelerator Appliance

- Incremental Update Support of Multiple DB2 Subsystems
  - Replicate up to **10** different subsystems into a shared Accelerator
    - IBM InfoSphere Data Replication for DB2 for z/OS v10.2.1
- Support Replication While Tables are Reloaded
  - IBM DB2 Analytics Accelerator V3.1 will stop replication for all tables if one replication enabled table needs to be reloaded
  - V4.1 allow reloading of a table while replicating other tables
- Enhanced Monitoring
  - Set of new and revised monitoring counters on **system level** exposed via Instrumentation Facility Interface (IFI) for better support of
    - Charge-back, capacity planning, monitoring and problem determination

Incremental Update with DB2 Analytics Accelerator V3.1

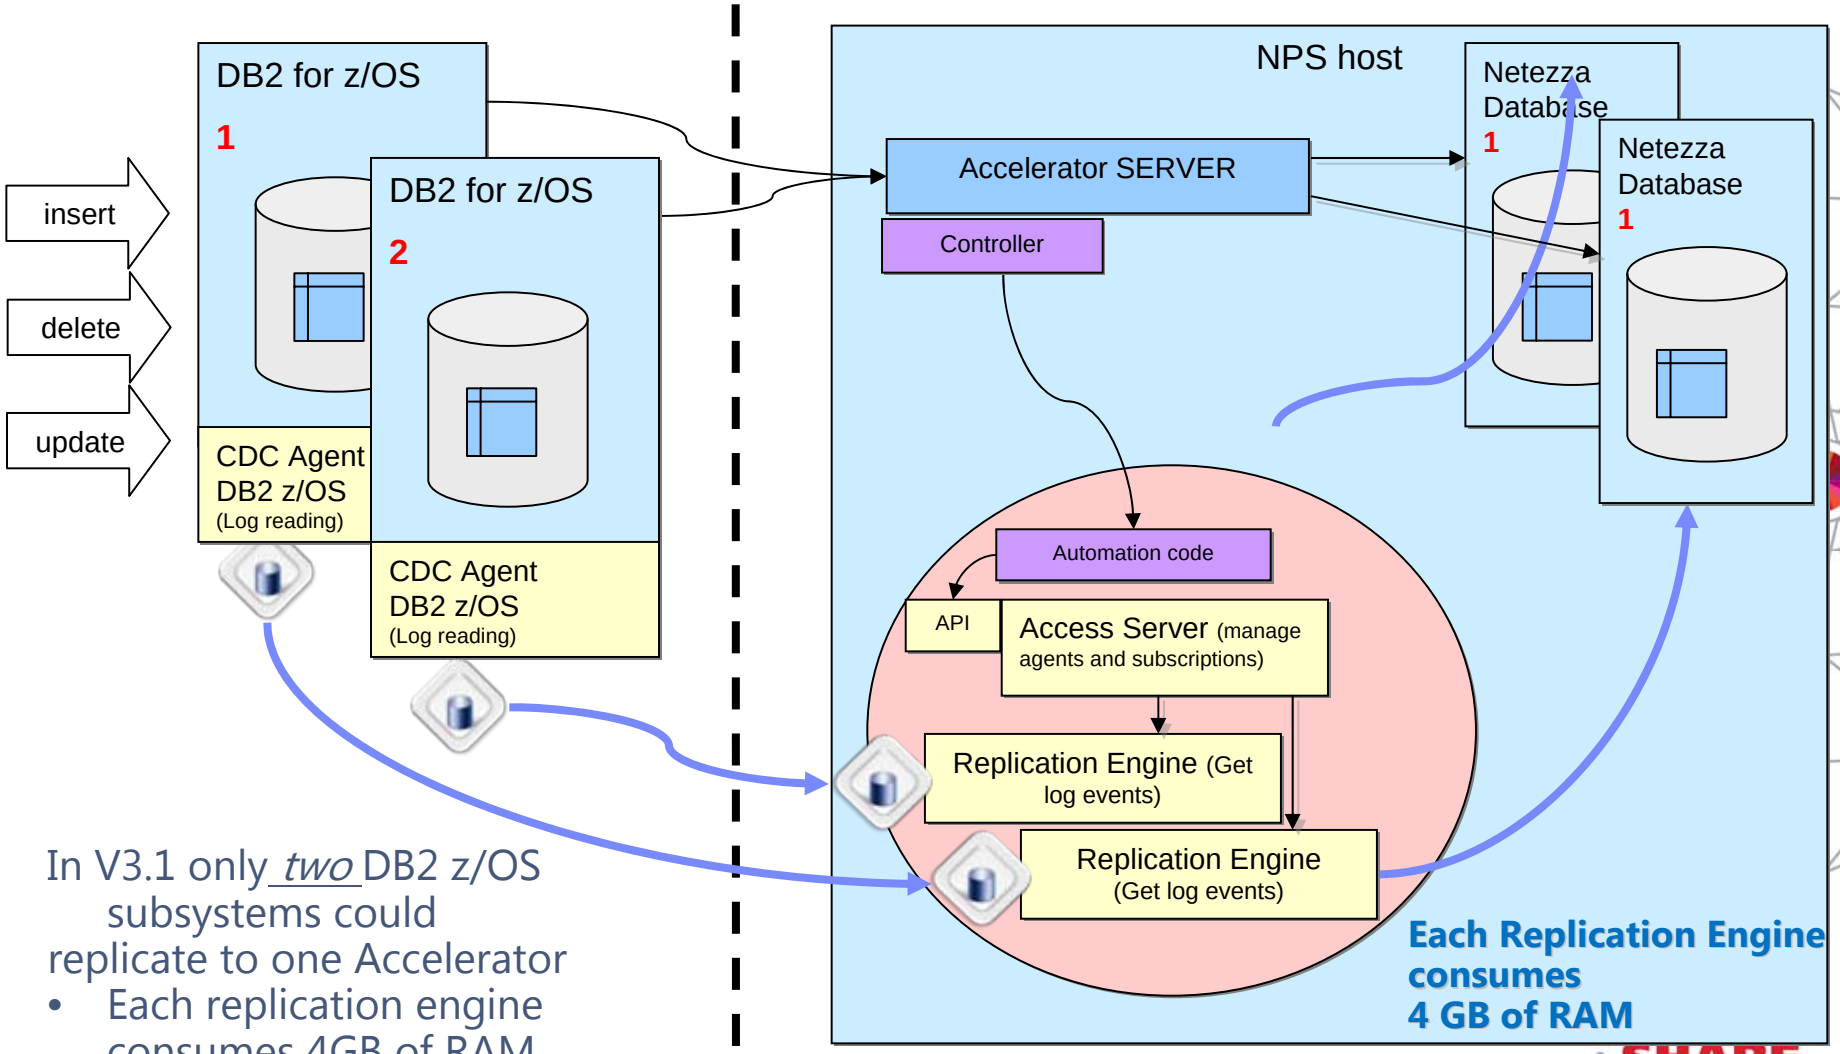

In V3.1 only *two* DB2 z/OS subsystems could replicate to one Accelerator

- Each replication engine consumes 4GB of RAM

Incremental Update with IBM DB2 Analytics Accelerator V4.1

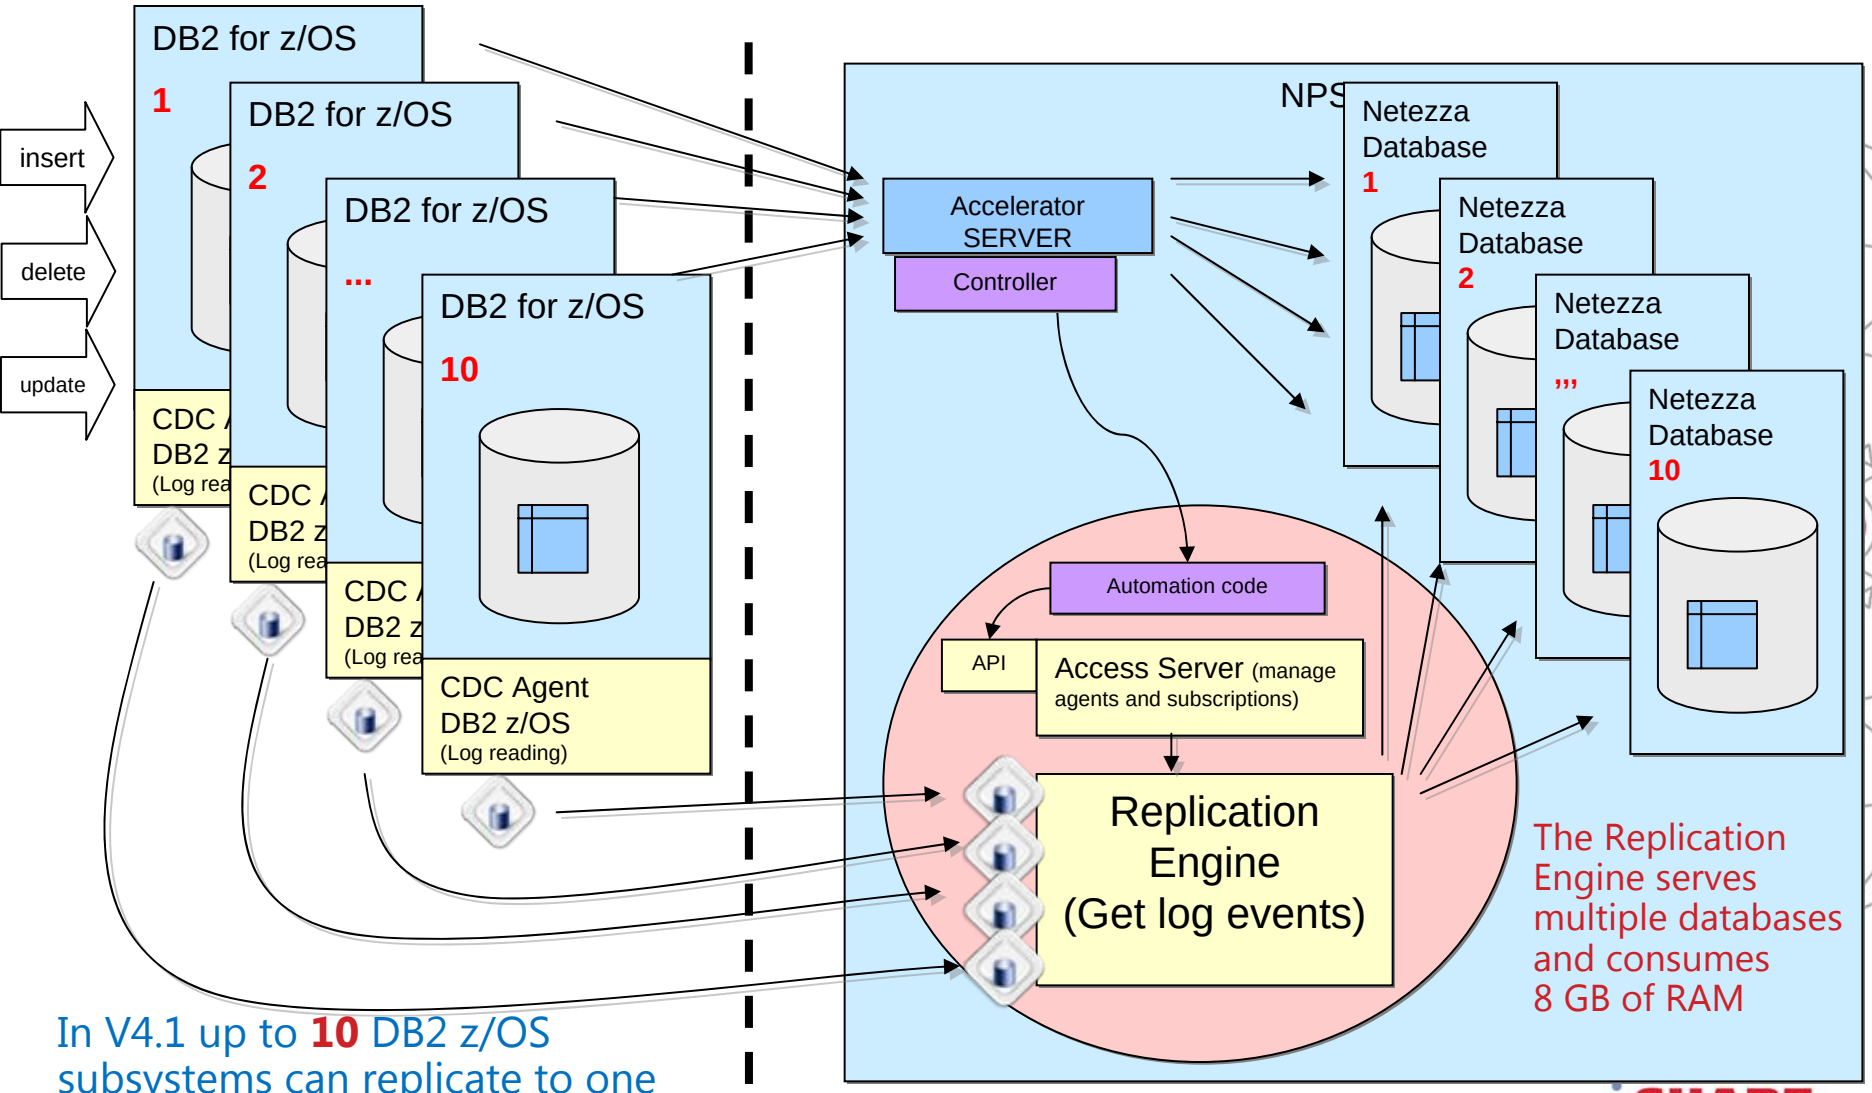

High Performance Storage Saver (HPSS)
Reducing the cost of high speed storage

Store historic partitions on the Accelerator only

- When partitions no longer require updating, reclaim the DB2 storage
- Continue to search the entire table at Accelerator speeds
- Direct mixed workload queries to hotter temperature partitions while focusing aggregate type querying to the warm or cold temperature data

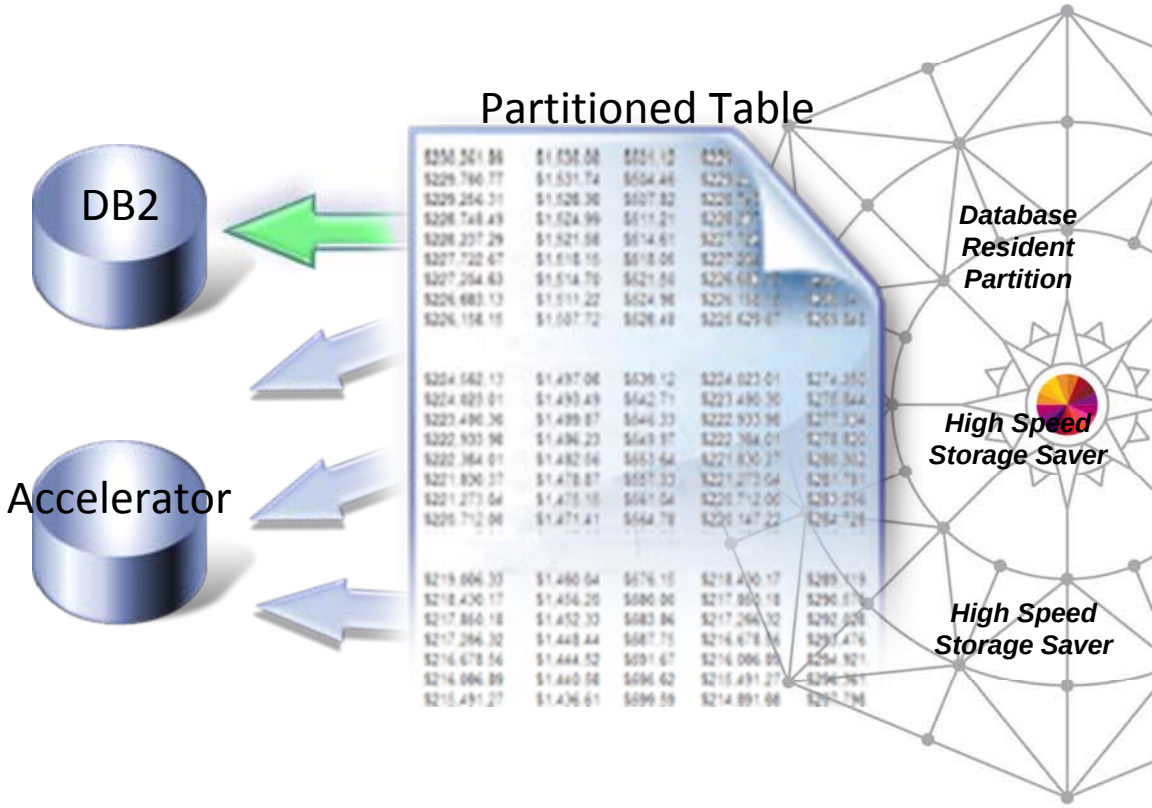

HPSS can be implemented for a set of partitions or all partitions

Special Registers - Managed by zParms

CURRENT QUERY ACCELERATION

CURRENT GET_ACCEL_ARCHIVE

Enhanced Storage Saver Solution (HPSS)

- Archiving Enhancements
  - DB2 partitions are set to **persistent read-only (PRO) status** after the partition was moved to the Accelerator to prevent future INSERTs or UPDATEs to the partitions
  - Create **image copies are protected** based on the PRO status of the DB2 partitions
    - No further image copies can be created
  - Up to **4 image copies** per partition are possible
    - Improves Disaster Recovery (DR) scenarios
    - Cataloged in SYSIBM.SYSCOPY
  - A **flexible naming schema** for the image copies
    - Simplifies the Restart process or rearchiving since the image copies must no longer be deleted manually beforehand
  - Ability **to archive** a table **on multiple accelerators**

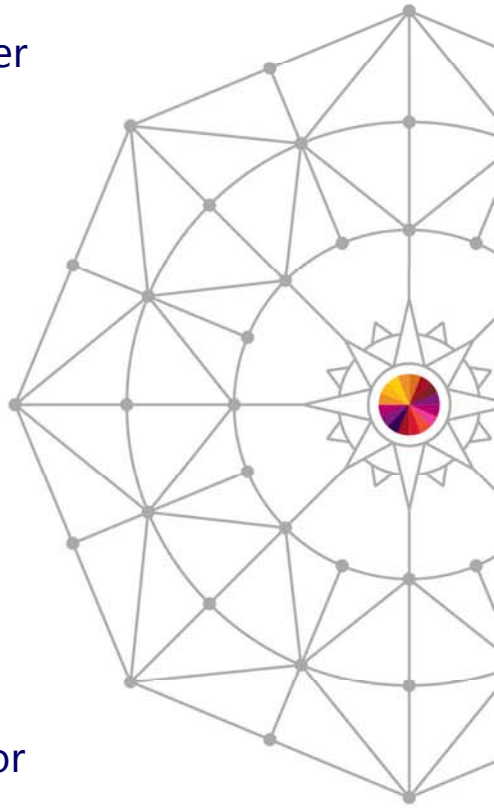

- Restore Archived Partitions
  - New stored procedure to automate the process to change one or more partitions on archived partitions via stored procedure
    - Restore is done from image copy data
    - No longer require to remove the table from the Accelerator

WLM Query Prioritization for Local Applications

- Extend WLM query prioritization to local queries
- DB2 propagates the Service Class and Importance Level for local applications to the Accelerator via a special register prior to PREPARE
  - Service Class is detected based on application's address space
    - TSO subsystem type – e.g. SPUFI
    - JES subsystem type – e.g. TEP3, host language apps run via TSOBATCH
    - STC subsystem type – any apps kicked off via START command
  - For CICS/IMS applications the Service Class is detected based on performance block
- Accelerator will map the Importance Level to the Netezza priority prior to query execution

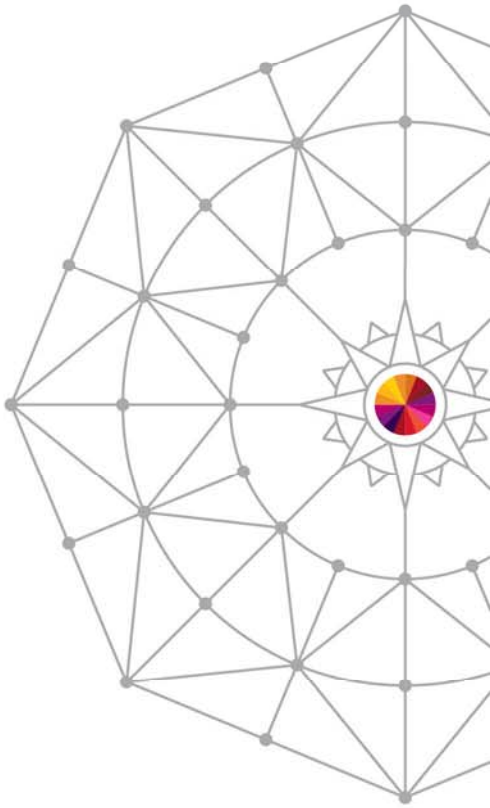

WLM Importance Level Priority Mapping

- Netezza supports 4 different priority level
  - Critical | High | Normal | Low

- The Accelerator maps the WLM Importance Level to a corresponding Netezza priority before query execution

WLM Importance Level	Netezza Priority
System	**Critical**
Importance 1	**Critical**
Importance 2	**Critical**
Importance 3	**High**
Importance 4	**Normal**
Importance 5	**Low**
Discretionary	**Low**

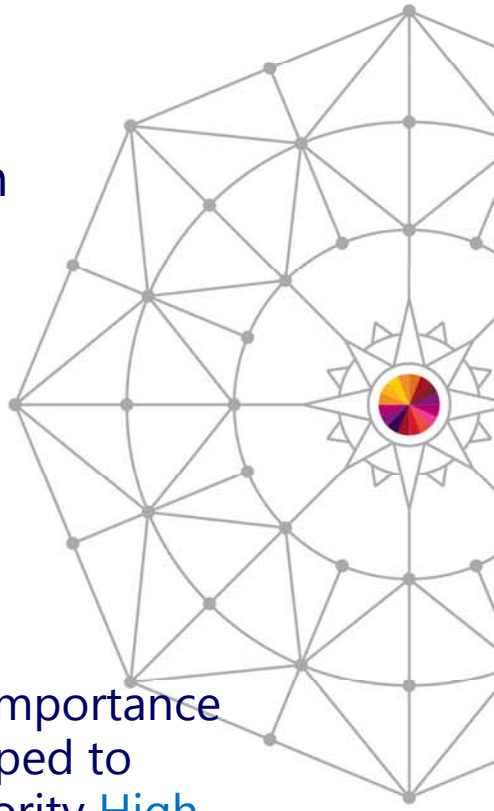

In V3.1 WLM Importance level 2 mapped to Netezza Priority High and level 3,4,5 mapped to Normal.

Sub-Capacity Licensing Growth on Demand

- Buy a full rack N2001 but pay only for 50 %, 62.5%, 75 % or 87.5 %
- License must not use more then defined capacity
  - Customer responsibility to monitor and document the actual usage
  - System provides interfaces to set/monitor resource limits
- Capacity is a combination of
  - data storage utilized as a percent of the total available physical storage
  - performance settings as a percentage of maximum resource allocation
    - Can be specified using the Set resource limits for DB2 subsystems menu item on the accelerator

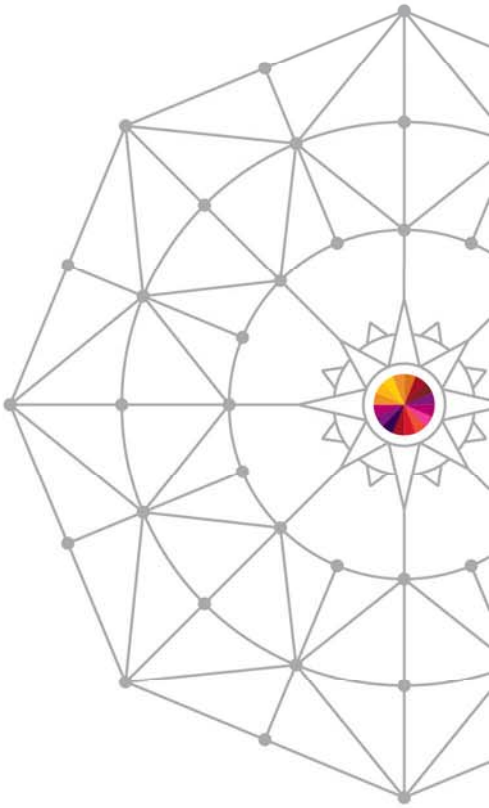

Miscellaneous Enhancements

- Multiple Codepage Support
  - Mixed EBCDIC and UNICODE tables are now allowed on the accelerator for the same DB2 subsystem
    - Queries that combine both EBCDIC and UNICODE tables can not be routed
- Fine Grained Access Control for Stored Procedure ACCEL_CONTROL_ACCELERATOR
  - SP offers several function to control an accelerator
    - e.g. Cancel task, starting replication, collecting trace
- IBM DB2 Analytic Accelerator Studio
  - Provided as a full install image based on Data Studio V4.1
  - New Restore Partition wizard
  - Usability improvements

Operation and Maintenance Enhancements

- Automated NZKit Install
  - The Netezza software update (NPS) can now be done through a stored procedure or Data Studio in the same way as Accelerator Software Updates
- Load Operation Improvements
  - DFSMS enhancements reduce zCP time up to 30% during Accelerator data loads
    - z/OS enhancement only
      - z/OS V1.12 : PTF UA68971
      - z/OS V1.13 : PTF UA68972
      - z/OS V2.1 : PTF UA68973

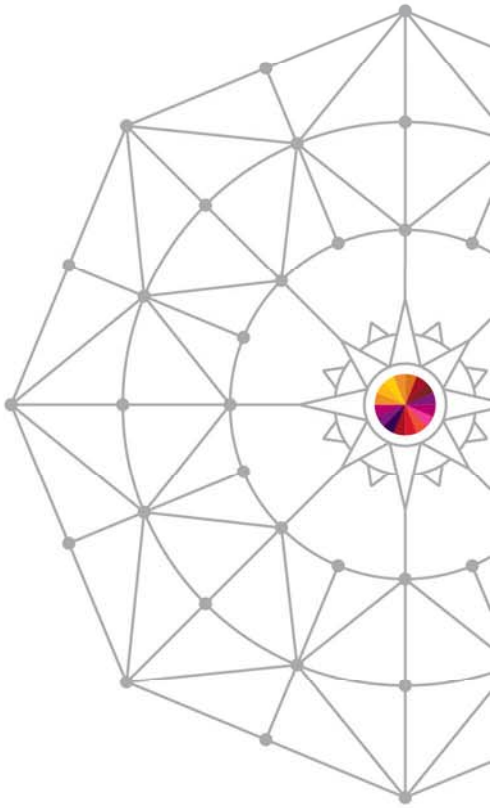

zEnterprise Analytics System (IZAS) 9700 and 9710

A cost-competitive, integrated combination of hardware, software and services to deliver business reporting and business critical analytics

zEnterprise Analytics System 9700

zEnterprise Analytics System 9710

Preselected

To deliver a comprehensive, end-to-end, flexible solution

Pretested

To meet business critical analytic demands

Solution Priced

For deployment as an add-on to an existing zEnterprise system or as a new system

IBM DB2 Analytics Accelerator and zEnterprise

Minimize latency. Improve performance. Drive innovation.

- **Bring Analytics to the Data**
  - Reduced Latency
  - Reduced Complexity
  - Reduce Cost
- **Deliver Business Critical Analytics**
  - Timely, Accurate, Secure Data
  - Rapid Deployment and Expansion
- **Evolve with the Business**
  - Start with Your Top Analytics Requirements
  - Grow without Changing Your Existing IT Environment

Transactions in

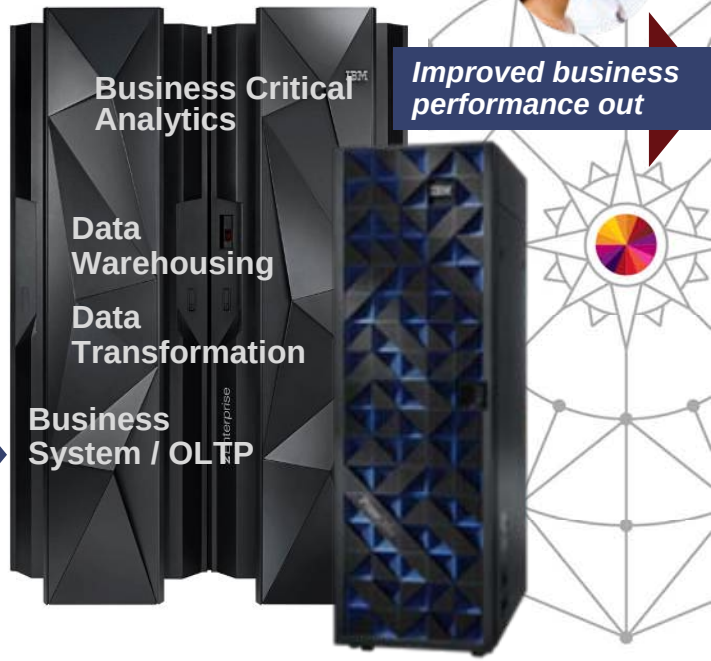

Shameless Plug

- Session 15036: Enabling Best-of-bread Analytics with zEnterprise
  - Speaker: Carl Parris, IBM, STSM
  - Wednesday, March 12, at 3:00pm (Platinum Ballroom Salon 1)
  - SHARE Analytics Spotlight session
- Session 14936: Information Governance: Key to Business Optimization with zEnterprise
  - Speaker: Theresa Tai, IBM, Executive IT Specialist
  - Thursday, March 13, at 4:30pm (Platinum Ballroom Salon 3)
- 2013 ITSO Redbooks Projects
  - Point of View (PoV)
    - Using IBM System z as a Hub for Business Analytics
      - *http://www.redbooks.ibm.com/abstracts/redp5062.html?Open*
  - Redbook
    - Hybrid Analytics Solution using IBM DB2 Analytics Accelerator V3.1
      - *http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg248151.html?Open#!*
- 2014 Deliverables
  - IBM DB2 Analytics Accelerator High Availability and Disaster Recovery Redguide, 1Q
  - IBM DB2 Analytics Accelerator V4.1 Redbook, target 2Q

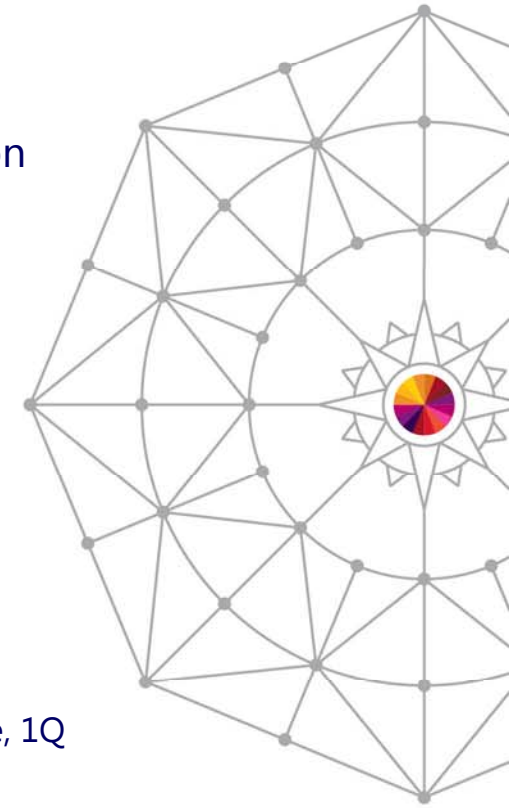

Reference Links

- IBM DB2 Analytics Accelerator
  - http://www-01.ibm.com/software/data/db2/zos/analytics-accelerator/
- IBM DB2 Analytics Accelerator Solution Video
  - http://www.youtube.com/watch?v=dUFKKqcA6kU
- Incremental Update
  - http://www-01.ibm.com/support/docview.wss?uid=swg27037912
- Business Analytics Buying Criteria
  - http://public.dhe.ibm.com/common/ssi/ecm/en/xbl03023usen/XBL03023USEN.PDF
- High Availability Network Setups
  - http://www-01.ibm.com/support/docview.wss?uid=swg27028171
- Business Partner Information
  - https://www.ibm.com/partnerworld/mem/sla.jsp?num=212-396
- Youtube videos on Client Success Stories
  - http://youtube.com

Thank You!