

IBM SolutionsConnect 2013

Turning Opportunity into Outcomes.



**Pure analytical performance
without compromises!
Ruud Lameijer**






PureSystems™

Systems with integrated expertise and built for cloud

Built-in Expertise

Capturing and automating what experts do – from the infrastructure patterns to the application patterns



Integration by Design

Deeply integrating and tuning hardware and software – in a ready-to-go workload optimized system

Simplified Experience

Making the IT lifecycle easier – with integrated management of the system and a broad open ecosystem of optimized solutions

PureFlex



MSP Editions

Infrastructure

Delivering Infrastructure Services

PureApplication



New Models

Application Platform

Delivering Platform Services

PureData



New Models

Data Platform

Delivering Data Services





IBM PureData System

Meeting Big Data Challenges – Fast and Easy!



PureData

System for Transactions

For apps like E-commerce...

Database cluster services optimized for transactional throughput and scalability

PureData

System for Analytics

Powered by Netezza technology

For apps like Customer Analysis...

Data warehouse services optimized for high-speed, peta-scale analytics and simplicity

PureData

System for Operational Analytics

For apps like Real-time Fraud Detection...

Operational data warehouse services optimized to balance high performance analytics and real-time operational throughput

Built-In Expertise Makes This as Simple as an Appliance



- Dedicated device
- Optimized for purpose
- Complete solution
- Fast installation
- Very easy operation
- Standard interfaces
- Low cost





IBM PureData System for Analytics

Optimized exclusively for analytic data workloads

PureData System for Analytics

*Delivering data services
for analytics*



Speed

- 10-100x faster than traditional custom systems
- Patented, hardware accelerated MPP (Massively Parallel Processing)

Simplicity

- Data load ready in hours
- No database indexes
- No tuning
- No storage administration

Scalability

- Peta-scale data capacity

Smart

- Designed to runs complex analytics in minutes, not hours
- Richest set of in-database analytics

Loading the PureData System for Analytics

Data Integration

- Ab Initio
- Cloudera
- Composite Software
- IBM Big Insights
- IBM Information Server
- IBM InfoSphere Streams
- Informatica
- Oracle Data Integrator
- Oracle GoldenGate
- SAP Business Objects



Querying the PureData System for Analytics

Reporting and Analysis

- IBM Cognos
- IBM SPSS
- IBM Unica
- Information Builders
- Kalido
- KXEN
- Microsoft Excel
- MicroStrategy
- Oracle OBIEE
- SAP Business Objects
- SAS
- Actuate



Inside the IBM PureData System for Analytics N2001



Optimized Hardware + Software

- Hardware accelerated AMPP
- Purpose-built for high performance analytics
- Requires no tuning

SMP Hosts


- SQL Compiler
- Query Plan
- Optimize
- Admin

Disk Enclosures

- User data, mirror, swap partitions
- High speed data streaming


Snippet Blades™

- Hardware-based query acceleration with FPGAs
- Blistering fast results
- Complex analytics executed as the data streams from disk



Snippet-Blade™ (S-Blade) Components

- HX5 Blade
- 128 GB RAM
- 16 Intel Cores



- BPE4 Side Car
- 16 GB RAM
- 16 Virtex-6 FPGA Cores
- SAS Controller

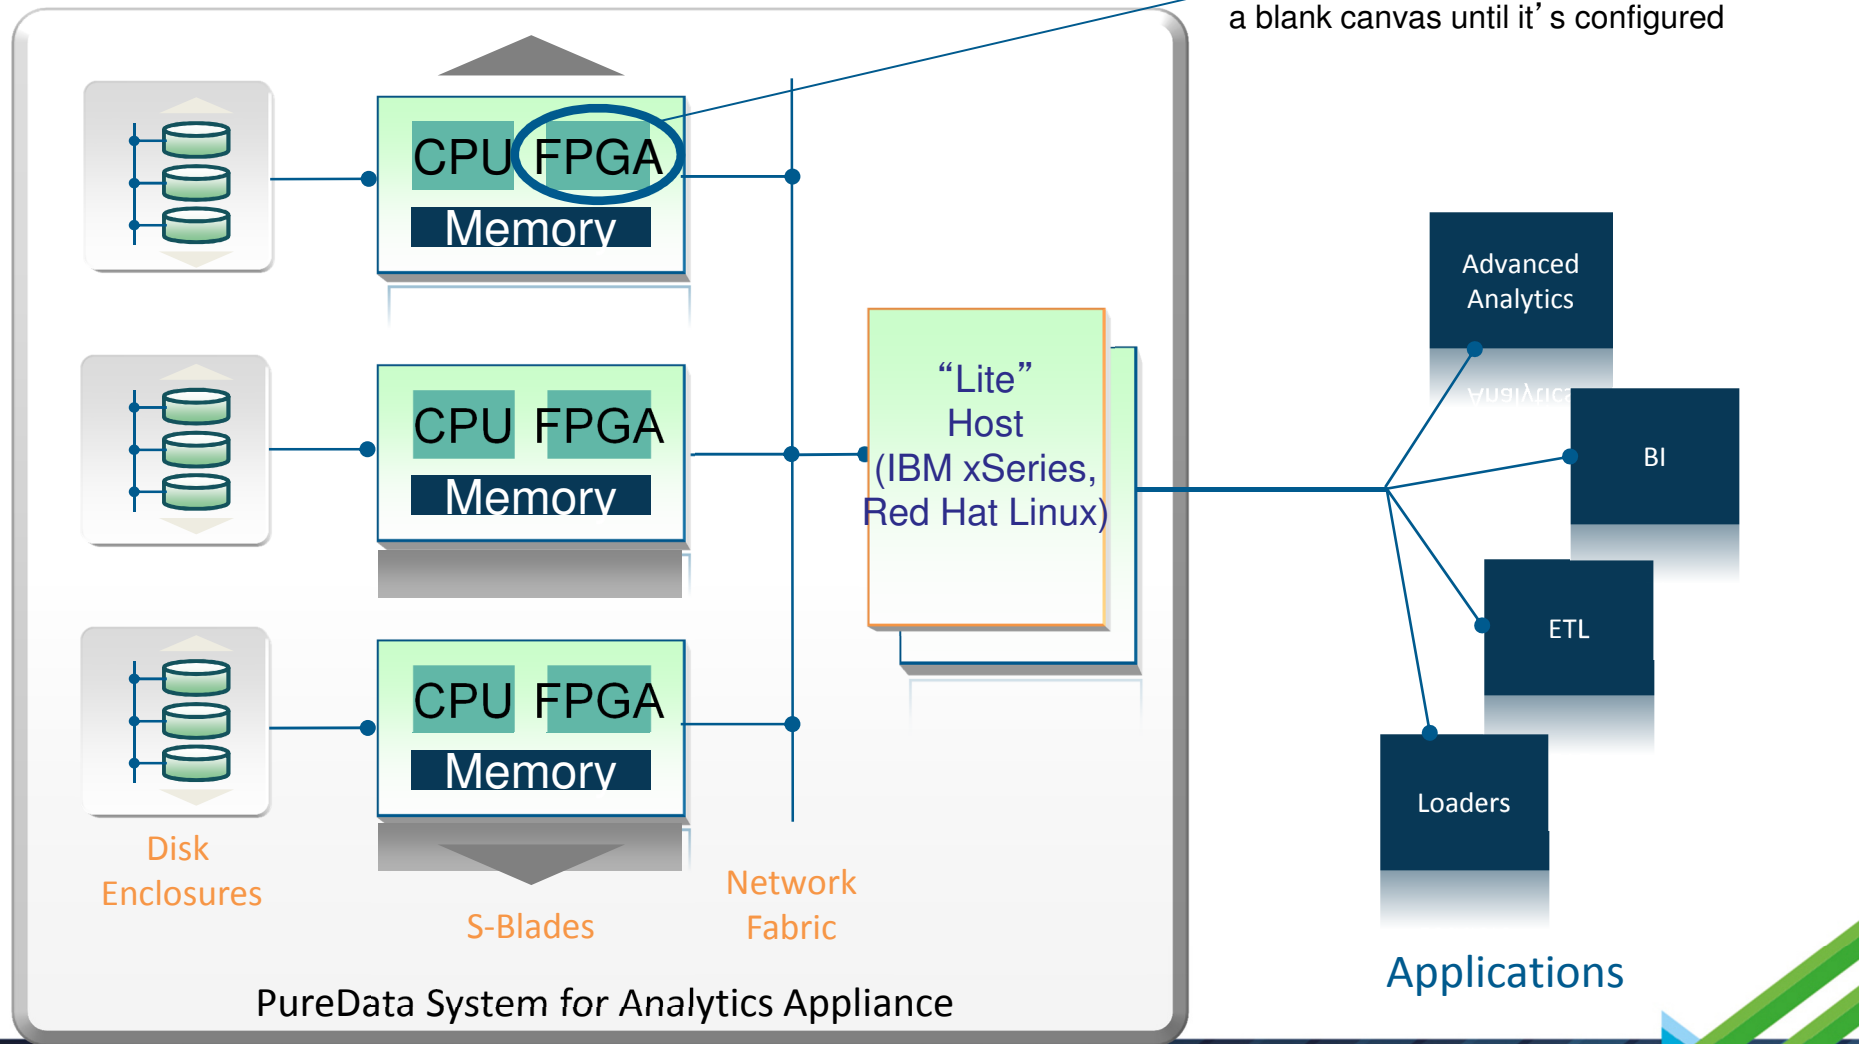
IBM BladeCenter Server

Netezza DB Accelerator

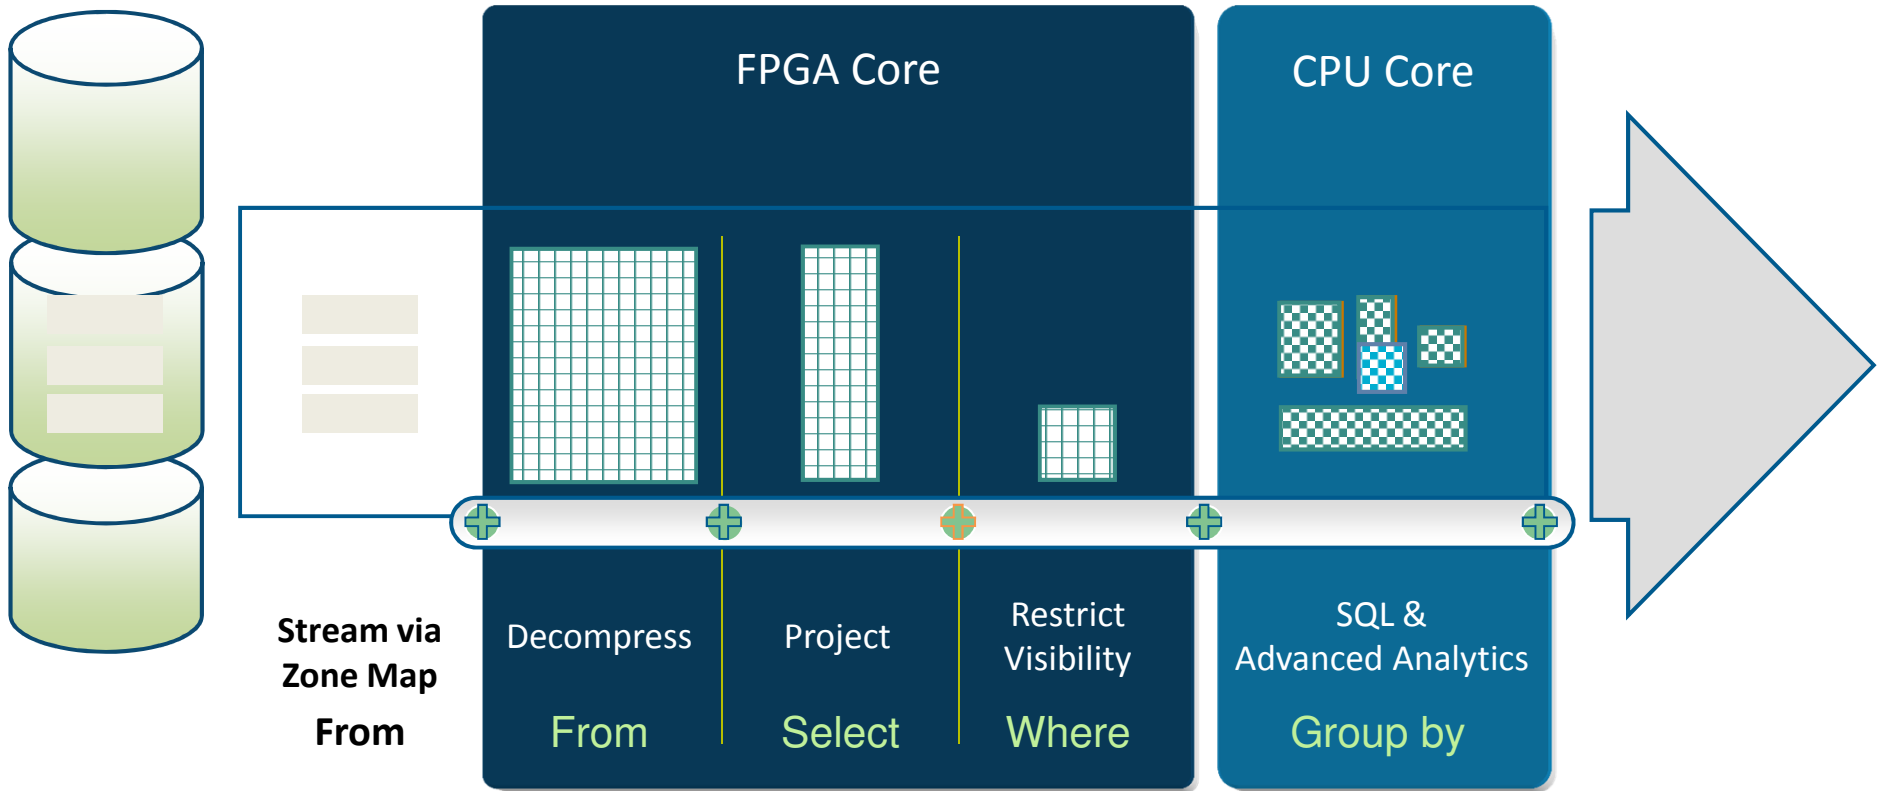
The PureData System for Analytics AMPP Architecture



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



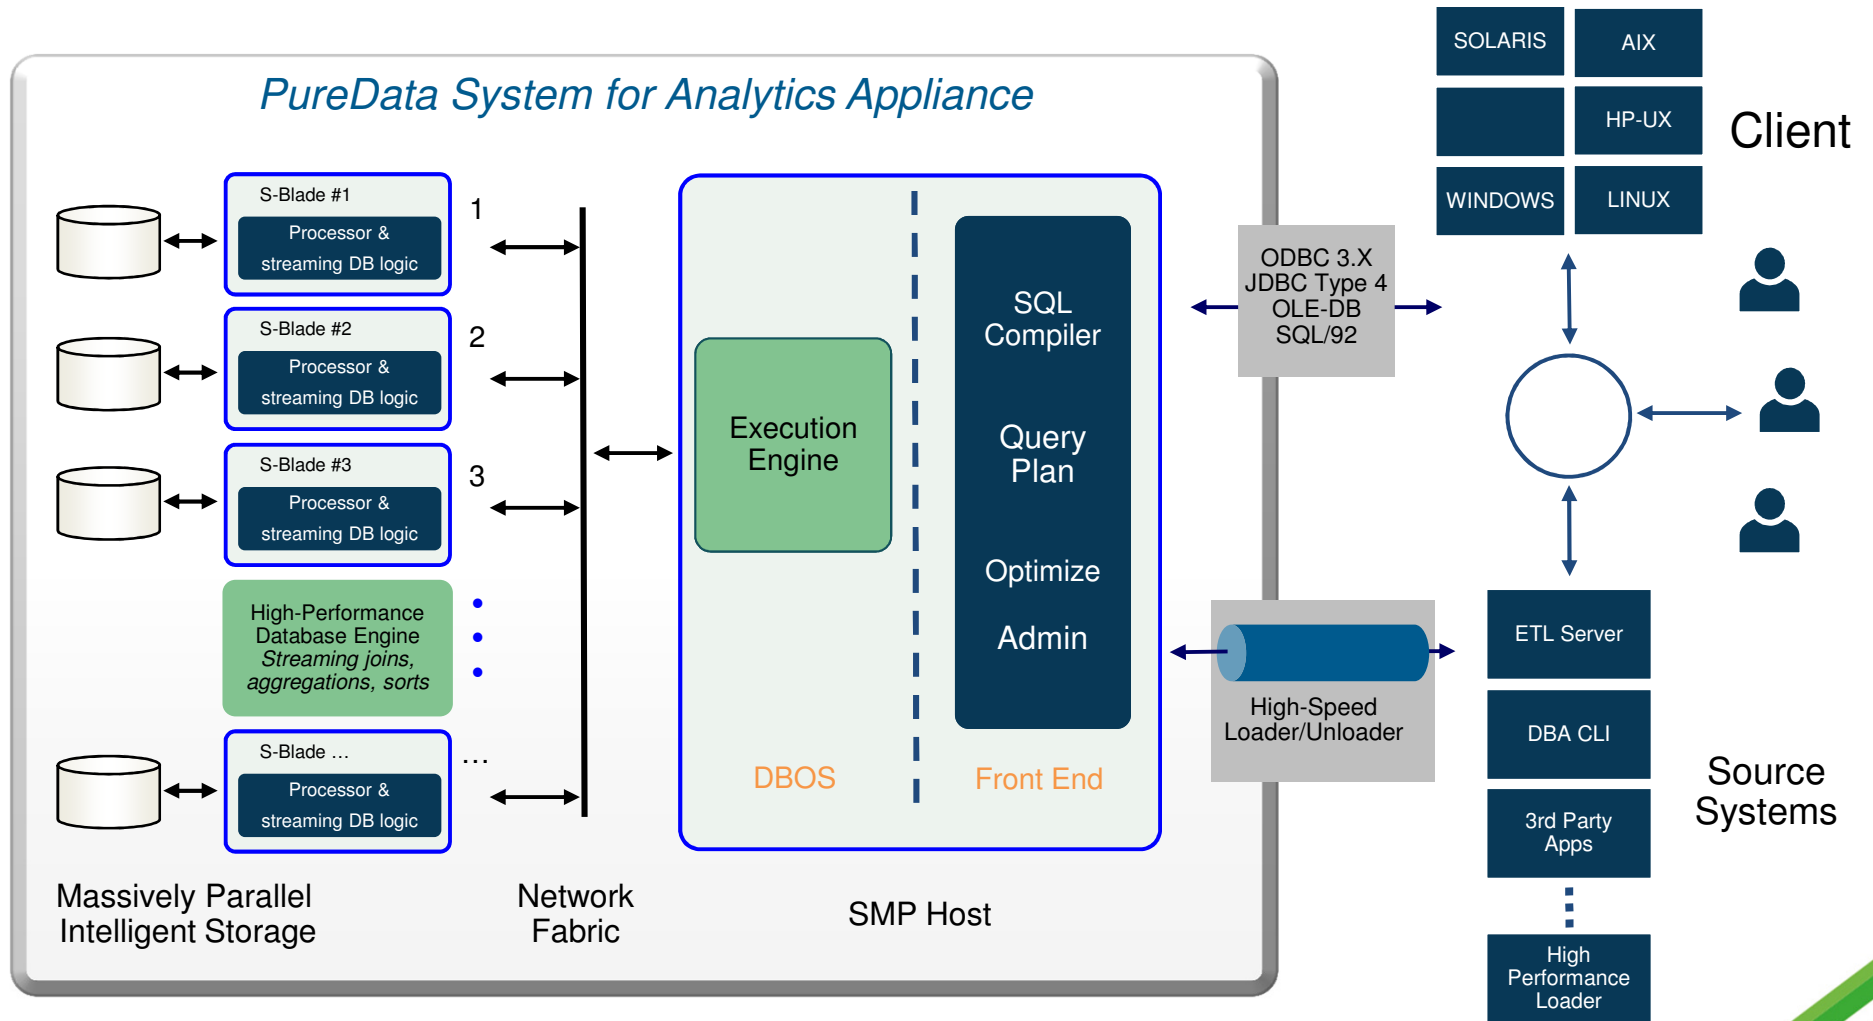
S-Blade Data Stream Processing



```

Select State, Age, Gender, count(*) From MultipleBillRowCustomTable Where BirthDate >= '1/01/1960'
And State = 'CA' in ('SEL', 'NGA') Group By State, Age, Gender Order By State, Age, Gender
State, Age, Gender
    
```

The PureData System for Analytics AMPP Architecture in Action

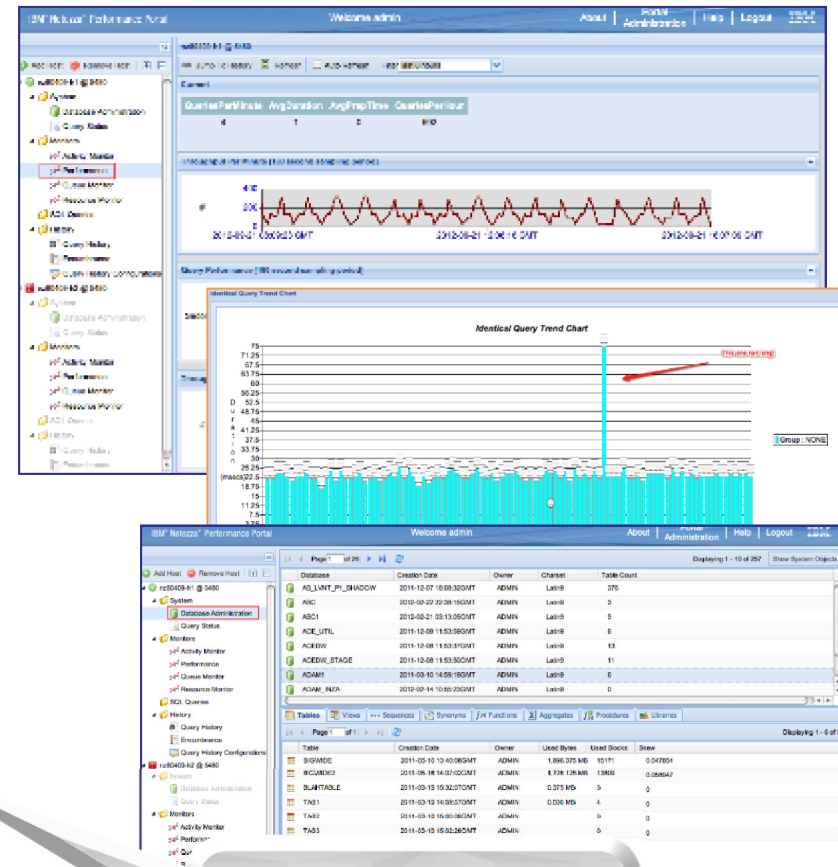


Spend Less Time Managing and More Time Innovating



- ✓ Easy Administration Portal
- ✓ No software installation
- ✓ No indexes and tuning
- ✓ No storage administration

- No dbspace/tablespace sizing and configuration
- No redo/physical/Logical log sizing and configuration
- No page/block sizing and configuration for tables
- No extent sizing and configuration for tables
- No Temp space allocation and monitoring
- No RAID level decisions for dbspaces
- No logical volume creations of files
- No integration of OS kernel recommendations
- No maintenance of OS recommended patch levels
- No JAD sessions to configure host/network/storage



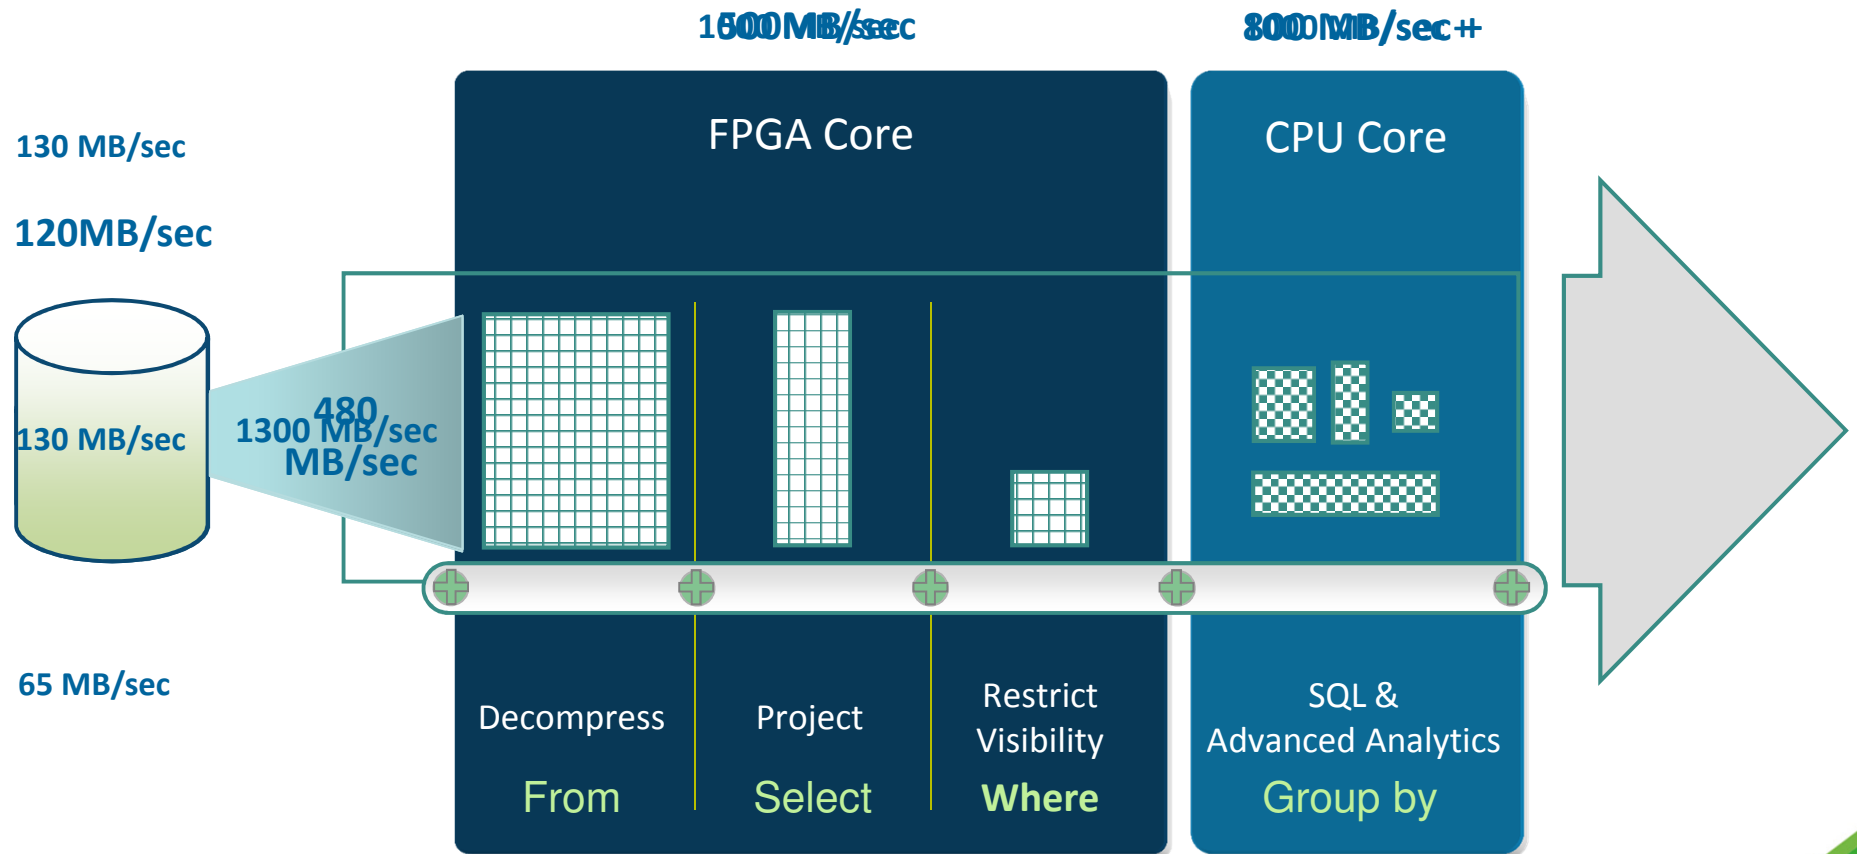
Data Experts, not Database Experts

How we did it



325 MB/sec
(2.5 drives / core)

PureData System for Analytics N2001



Benefits of the IBM PureData System for Analytics



Accelerate Performance of Analytic Queries

- **3X faster performance¹**
for Big Data analytics
- **128 GB/sec effective scan rate per rack²**
to tackle Big Data faster

Increase Efficiency of your Data Center

- **50% greater data capacity per rack³**
helps optimize data center efficiency
- **More capacity and less power per rack**

Simplicity and Ease of Administration

- **Improved system management and resilience**
*to spend less time managing and more time
delivering value*
- **70% FEWER service calls**
with more spare drives and faster disk regeneration⁴

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.

⁴ Each N2001 rack contains 34 hot spare drives and 240 active drives for a ratio of 1 spare per 7 drives. Each N1001 rack contains 4 hot spare drives and 92 active drives for a ratio of 1 spare per 23 drives. The N2001 has 3.3x more spares per active drive. Frequency of disk related service calls expected to decrease by 70% assuming the same drive failure rates.

PureData System for Analytics Hardware

Overview: Model N2001



12 Disk Enclosures

- 288 600 GB SAS2 Drives
 - 240 for User Data
 - 14 for S-Blades
 - 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
1/2 Rack to 4 Racks

- User Data Capacity: 192 TB*
- Data Scan Speed: 450 TB/hr*
- Load Speed (per system): 5+ TB/hr

- Power Requirements: 7.5 kW
 - Cooling Requirements: 27,000 BTU/hr
- * Assuming 4X compression



IBM DB2 Analytics Accelerator

Now even faster with PureData System for Analytics N2001

- The PureData System for Analytics N2001 is also the next generation DB2 Analytics Accelerator
- Providing the same improvements to our DB2 for zOS customers





Listen to What Some of Our Customers Have to Say

CATALINA®

<http://www.youtube.com/watch?v=uwn8HX0IM8o>



<http://www.youtube.com/watch?v=yKIGQuSYUd4>



<http://www.youtube.com/watch?v=Y0j1XSdMSDE>



<http://www.youtube.com/watch?v=-6UBeGllc98>




Make Every Interaction Count

<http://www.youtube.com/watch?v=ySLiDYgObFc>

Take the next step....

Discover the value and begin your journey with IBM PureSystems:

- Visit ibm.com/puresystems to learn more
- Join the conversation about this new category of computing:
 - Twitter: @IBM PureSystems
 - Hashtag: #IBMPureSystems or #expertintsys
 - YouTube Channel: expertintegratedsys
 - Blog: expertintegratedsystemsblog.com
- Developers – Get started today with our no charge trial offerings!
ibm.com/developerworks/puresystems/try
- Explore PureSystem partner solutions
ibm.com/puresystems/centre





© International Business Machines Corporation 2012

International Business Machines Corporation New Orchard Road Armonk, NY 10504

IBM, the IBM logo, PureSystems, PureFlex, PureApplication, PureData and ibm.com are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide.

A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml

All rights reserved.