

IBM Big Data HW Platform

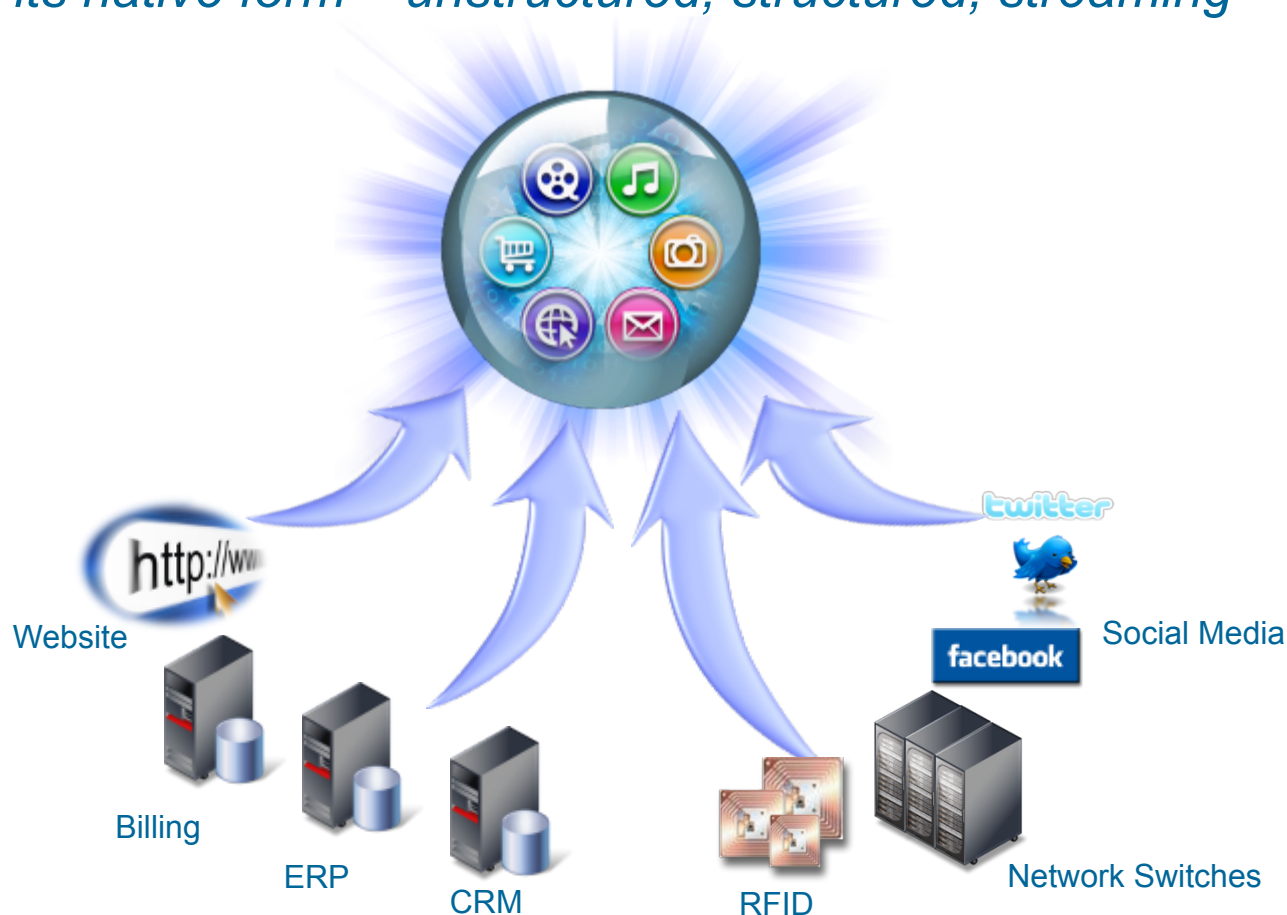
Turning big data into smarter decisions



Mujdat Timurcin
IT Architect
IBM Turk
mujdat@tr.ibm.com

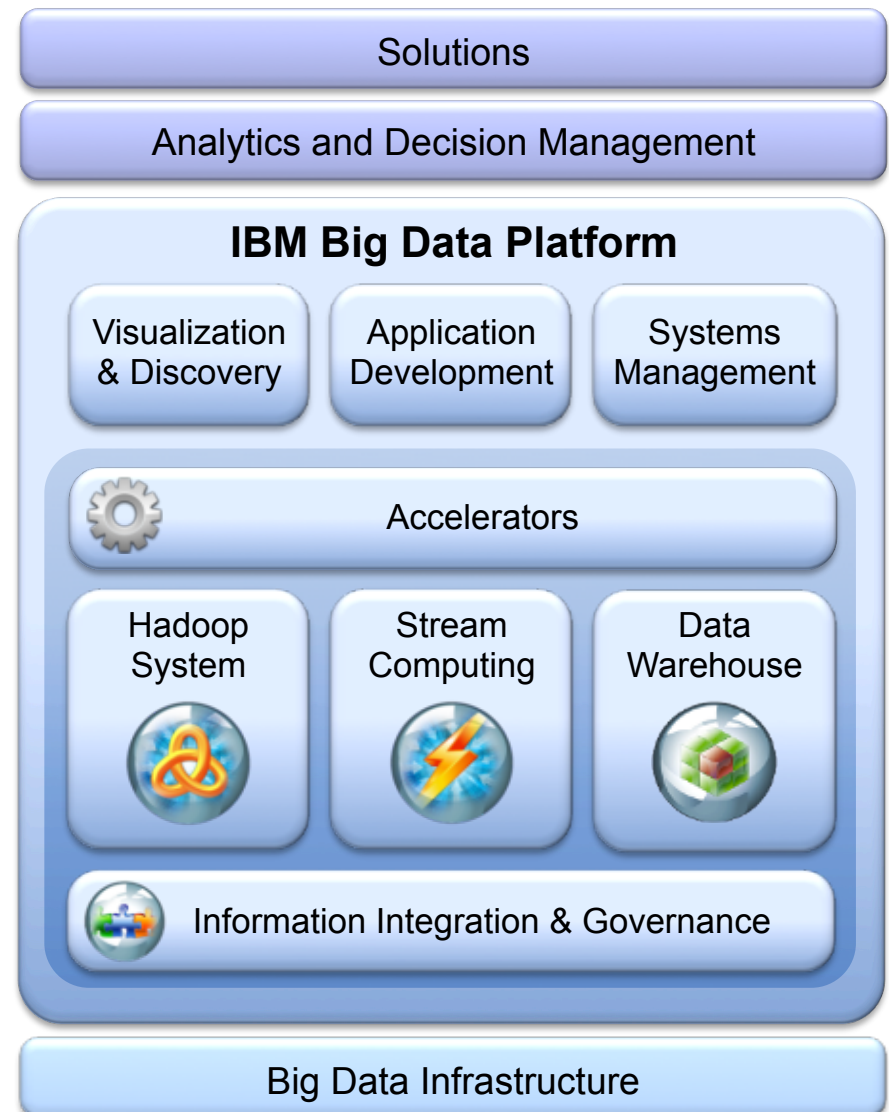
Big data is a hot topic because technology makes it possible to analyze *ALL* available data

Cost effectively manage and analyze *all available data*,
in its native form – unstructured, structured, streaming



The IBM Big Data Platform

- Process any type of data
 - Structured, unstructured, in-motion, at-rest
- Built-for-purpose engines
 - Designed to handle different requirements
- Analyze data in motion
- Manage and govern data in the ecosystem
- Enterprise data integration
- Grow and evolve on current infrastructure



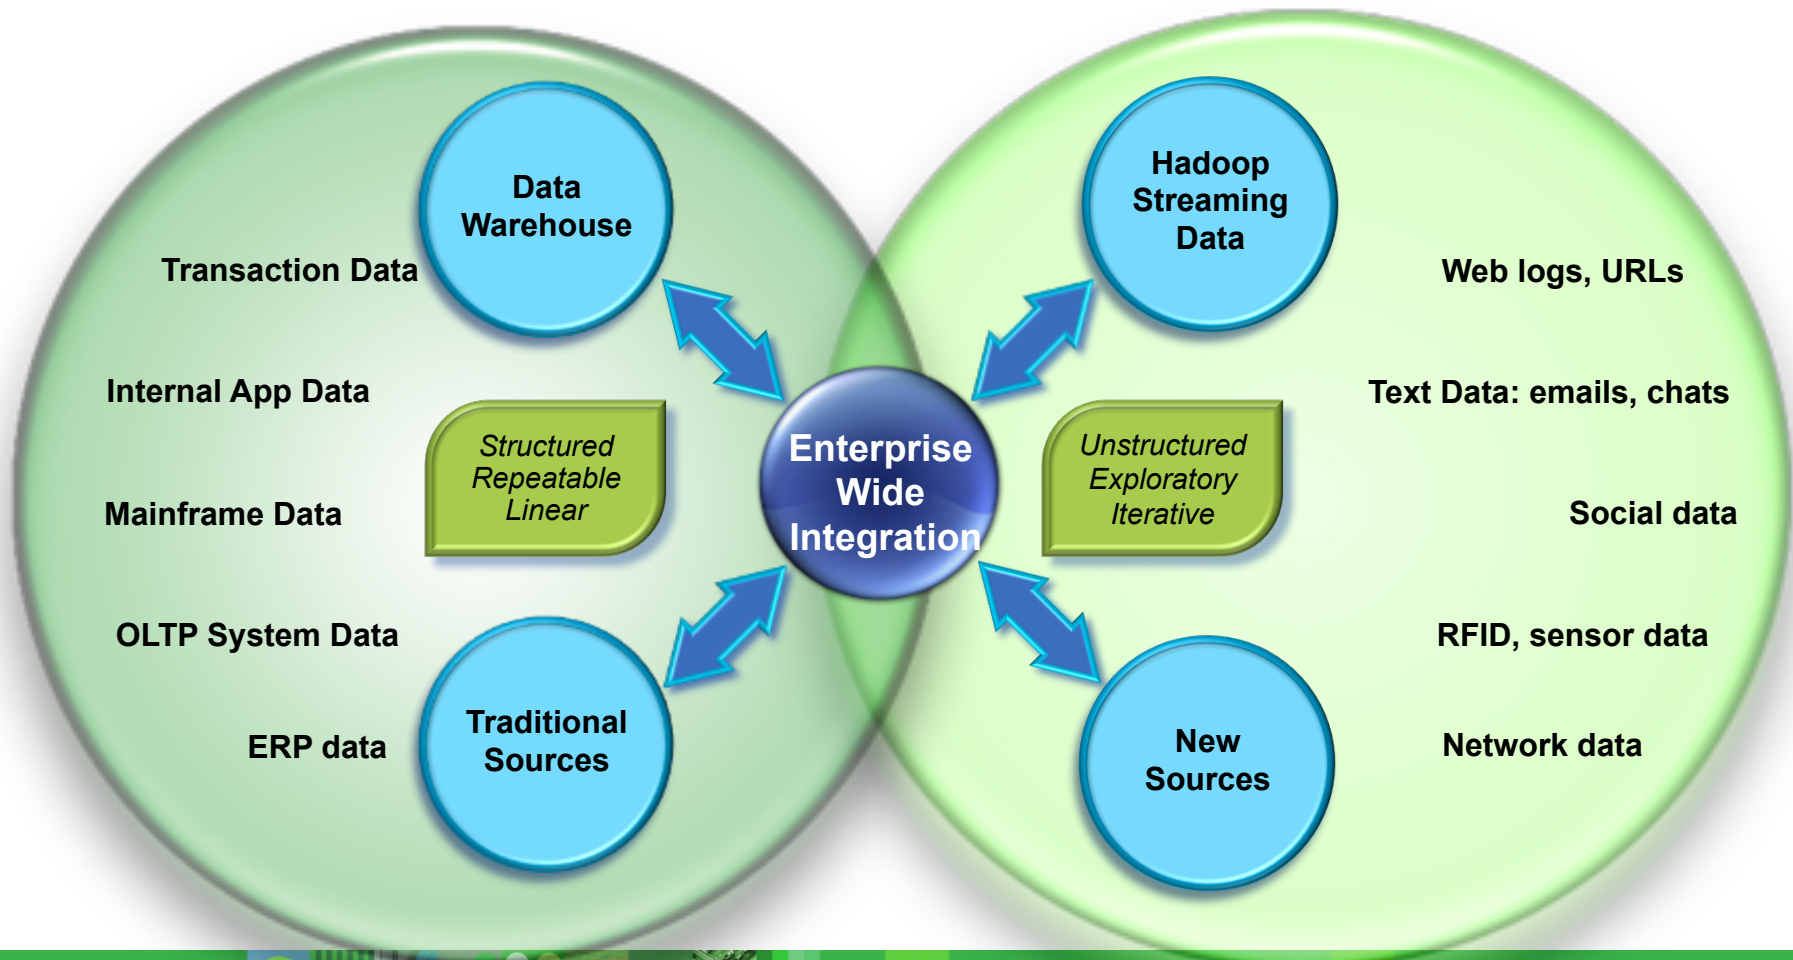
Analysis expanding from enterprise data to big data, creating new cost-effective opportunities for competitive advantage

Traditional Approach

Structured, analytical, logical

New Approach

Creative, holistic thought, intuition



A recent IBM/Oxford study highlights how organizations are adopting big data in four phases



Big data adoption



When segmented into four groups based on current levels of big data activity, respondents showed significant consistency in organizational behaviors

Total respondents n = 1061

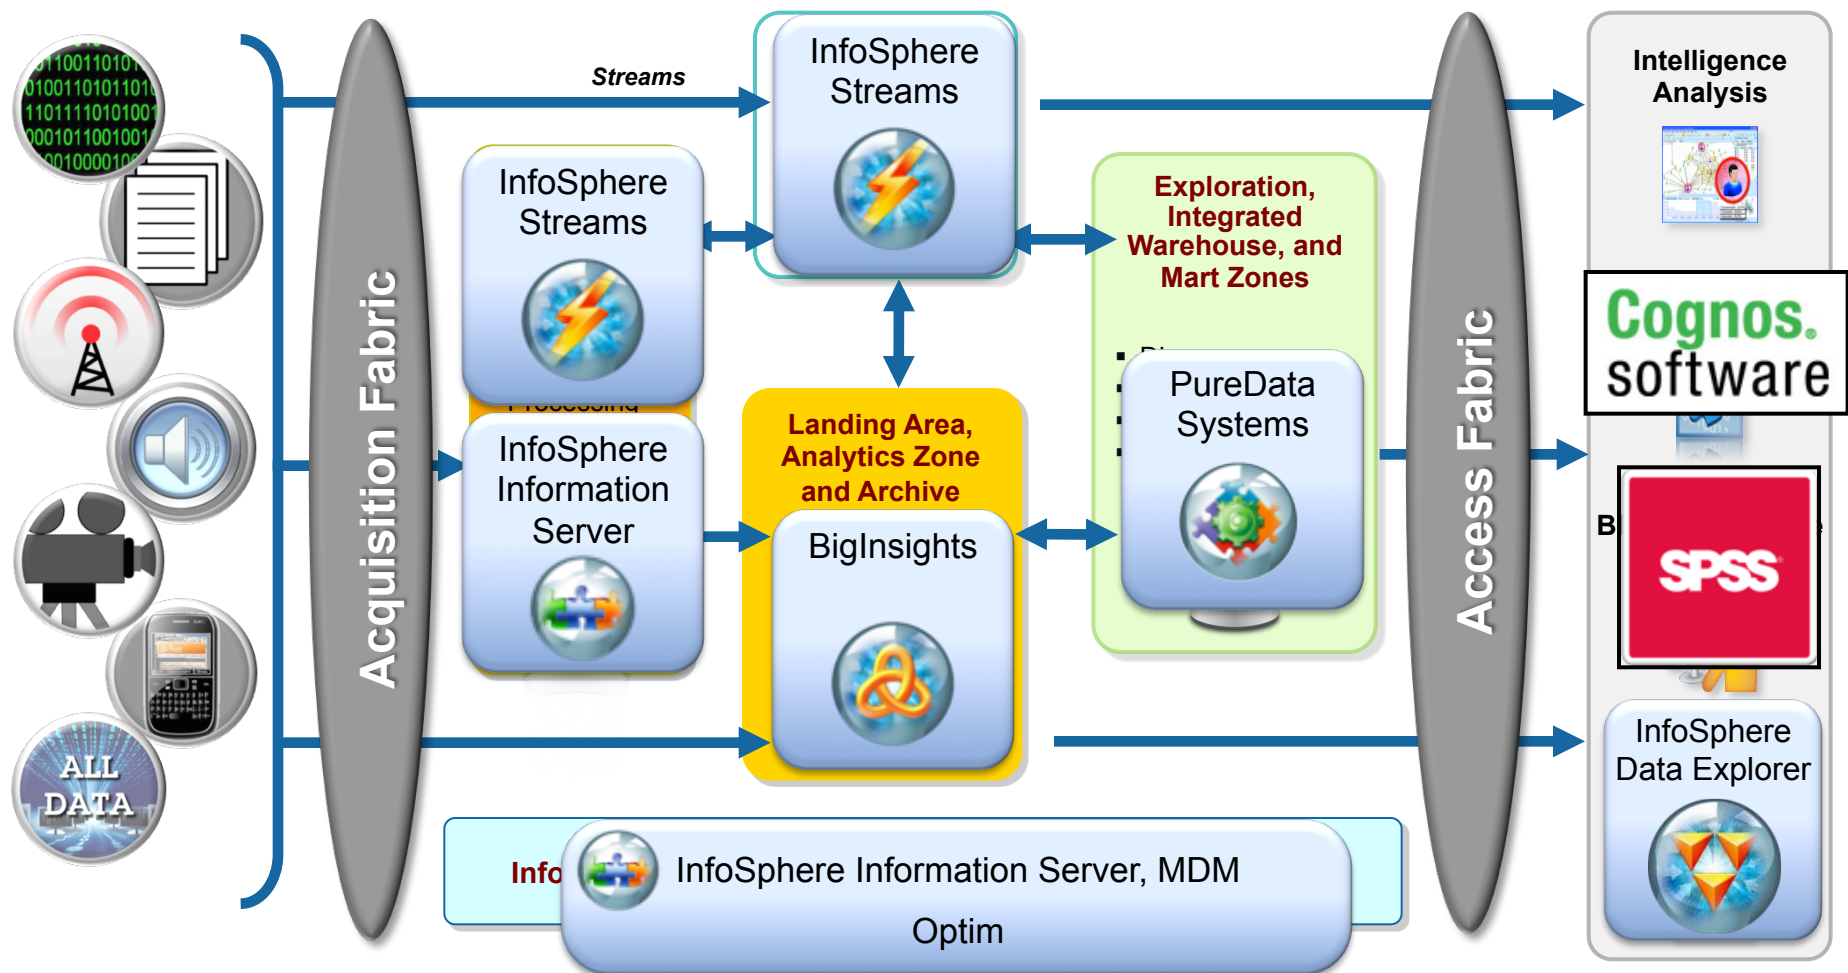
Totals do not equal 100% due to rounding

Architecture Vision

All Data Sources

Big Data Ecosystem

Analytic Applications





Cisco turns to IBM big data for intelligent infrastructure management


Capabilities Utilized

- ***Streaming Analytics***
- ***Hadoop System***
- ***Business Intelligence***

Results

- Optimize building energy consumption with centralized monitoring
- Automate preventive and corrective maintenance





Vestas optimizes capital investments based on **2.5 Petabytes** of information

Capabilities Utilized
Hadoop System

Results

- Model the weather to optimize placement of turbines, maximizing power generation and longevity.
- Reduce time required to identify placement of turbine from weeks to hours.
- Incorporate 2.5 PB of structured and semi-structured information flows. Data volume expected to grow to 6 PB.

Vestas



University of Southern California Oscar Senti-meter

Capabilities Utilized

- *Stream Computing*
- *Hadoop System*
- *Visualization*
- *Consumer Sentiment Analysis application*

Results

- Measures public sentiment about the Academy Awards race
- Detects public opinion for Hollywood's awards season
- Examines trends, volume, and content of **millions** of public Twitter messages
- Analytic accelerators to understand sentiment (positive, negative, neutral)

Big data projects should be viewed as strategic initiatives, not science projects.

“The Science Experiment”



General scenario

- Start an “unbalanced configuration”
- Leverage open source only solution
- Execute a “back of the envelope” project
- Pile on more nodes to address the data volume

Soon after....

- Lack of cluster management degrades QoS
- Workload bottlenecks due to networking issues
- Issues integrating with BI/DW resources
- Users need more, but can't deliver

Client experiences suggest considering the IT architecture up-front results in fewer problems

Streams is a cluster architecture

- **Minimum Streams configuration one or three nodes**
 - One node: Entry performance, no redundancy
 - Three nodes: Redundancy, fail-over for high availability
 - Add additional compute nodes, one at a time
 - Increase performance, bandwidth, availability

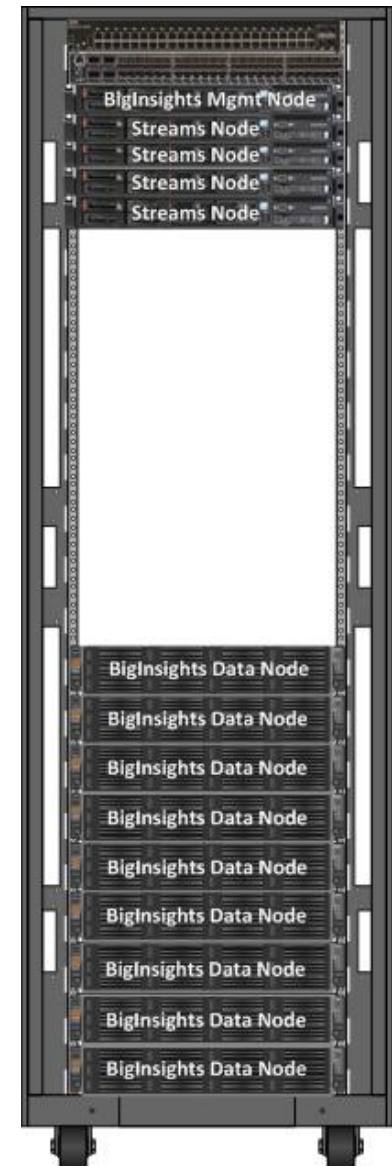


- **Co-locate Streams application /home directory**
 - Co-locate on minimum of three Streams compute nodes for high availability
 - GPFS manage data redundancy, failover across Streams nodes
 - Additional two 146 GB SAS HDDs, RAID1, 15,000rpm per node
- **Co-locate Streams recovery data base**
 - Co-locate active, failover DB2 instances on two Streams compute nodes
 - DB2 HADR manage redundancy, failover across two Streams nodes
 - Additional two 146 GB SAS HDDs, RAID1, 15,000rpm

Example Streams plus BigInsights

- Four 7R2 Streams nodes
 - Three Streams nodes include GPFS /home directory
 - Two Streams nodes include the recovery data base
- BigInsights 1/2 rack option
 - Available 1:1, 2:1 or 4:1 half rack options* depending upon data storage requirements
 - One 7R2 management node

***Please see backup slides for rack options...**



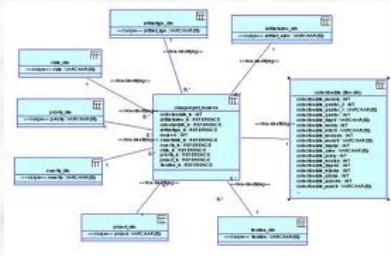


IBM Big Data Platform – Hadoop System (BigInsights)

What “schema-less” really means

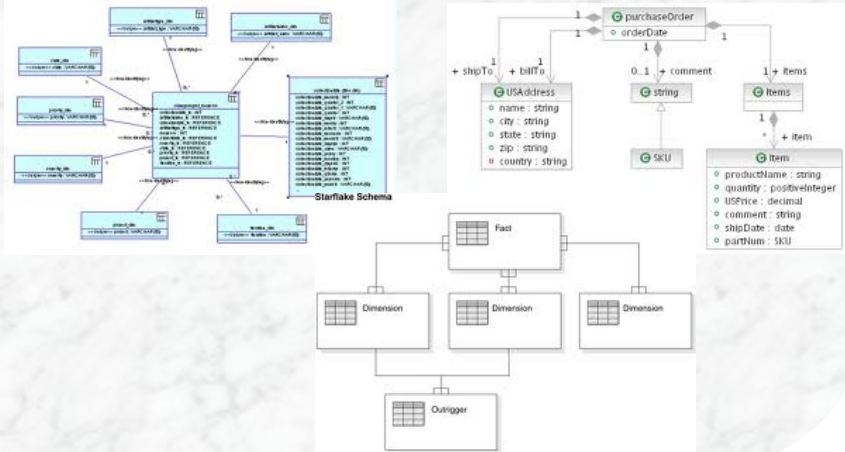
■ Relational Schema on Load

Data



■ Hadoop Schema on Run

Data



IBM Reference Architecture for Hadoop: InfoSphere BigInsights,

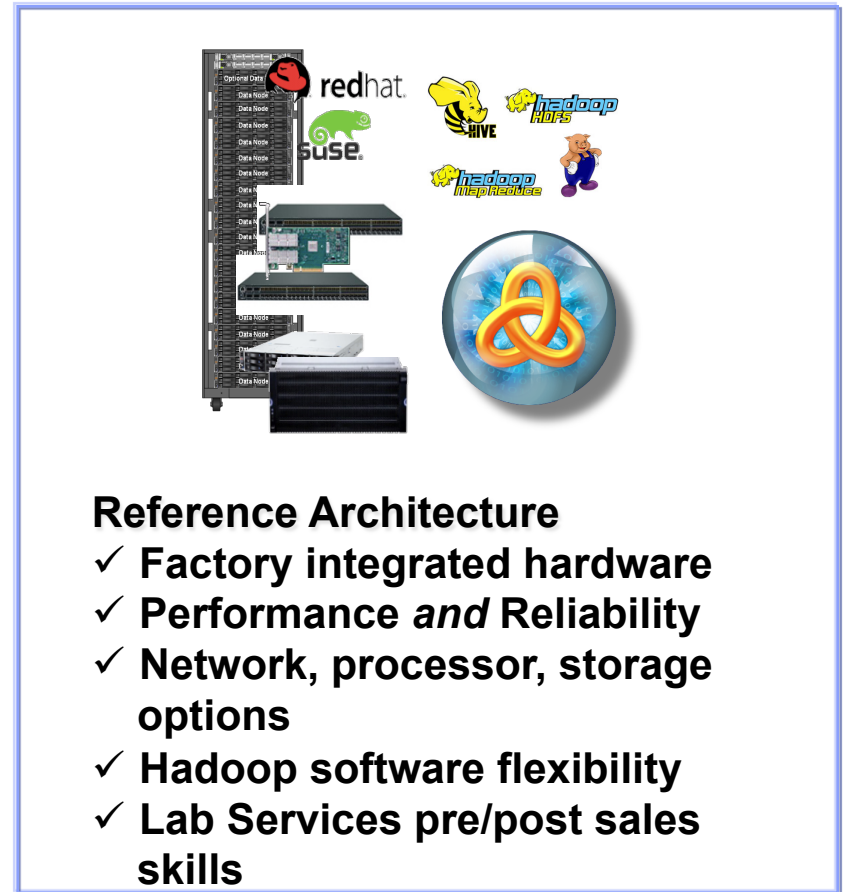
Built on System x or Power platforms

■ The Reference Architecture

- **Pre-defined Configurations:** Deep architectural and Hadoop services skills
- Big Data Hadoop stack (IBM, 3rd Party, Open Source)
- Config files on Smarter Analytics Big Data Wiki
- Leverage **TechLine** to Optimize and price.

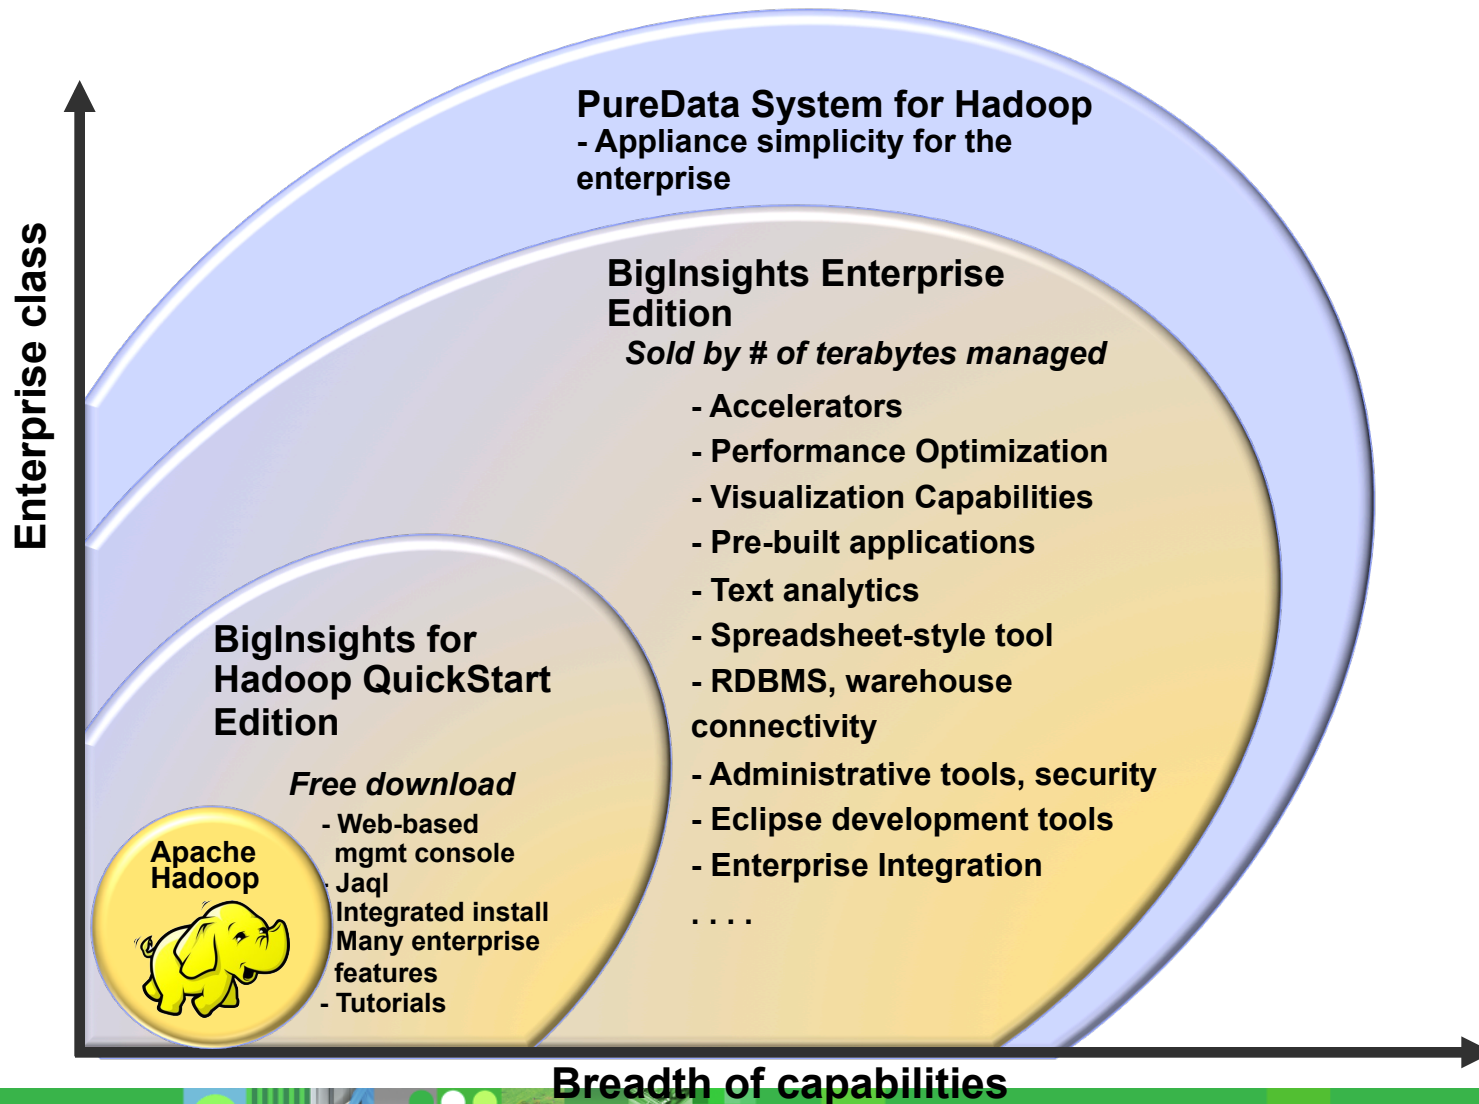
■ Addressing the needs of big data users.

- **Balanced and Flexible:** Architectural design points for compute or I/O bound analytics since **big data problems are not all the same**.
- **Easy Growth:** Building block design point to start small and grow meeting the **rapidly changing demands** of big data users.
- **Integrated Design:** High performance data ingestion to **address all big data information** requirements with options for networking and storage architectures.



From Getting Starting to Enterprise Deployment:

InfoSphere BigInsights Brings Hadoop to the Enterprise



Benefits of IBM PureData System for Hadoop

Exploring and analyzing more types of data



Built-in Expertise

*Accelerate
Big Data
Time to Value*

- **Deploy 8x faster**
than custom-built clusters¹
- **Built-in visualization**
to accelerate insight
- **Unlike big data appliances on the market, PureData System for Hadoop offers built-in analytic accelerators²**



Simplified Experience

*Simplify Big Data
Adoption & Consumption*

- **Single system console**
for full system administration
- **Rapid maintenance updates**
with automation
- **No assembly required, data load ready in hours**



Integration by Design

*Implement
Enterprise Class
Big Data*

- **Integrated Hadoop system with built-in archiving tools**
- **Delivered with more robust security**
than open source software
- **Architected for high availability**
- **Provides ability to load data at up to 14TB/hr**

¹Based on IBM internal testing and customer feedback. "Custom built clusters" refer to clusters that are not professionally pre-built, pre-tested and optimized. Individual results may vary.

²Based on current commercially available Big Data appliances product data sheets from large vendors.

What Makes PureData System for Hadoop Different?

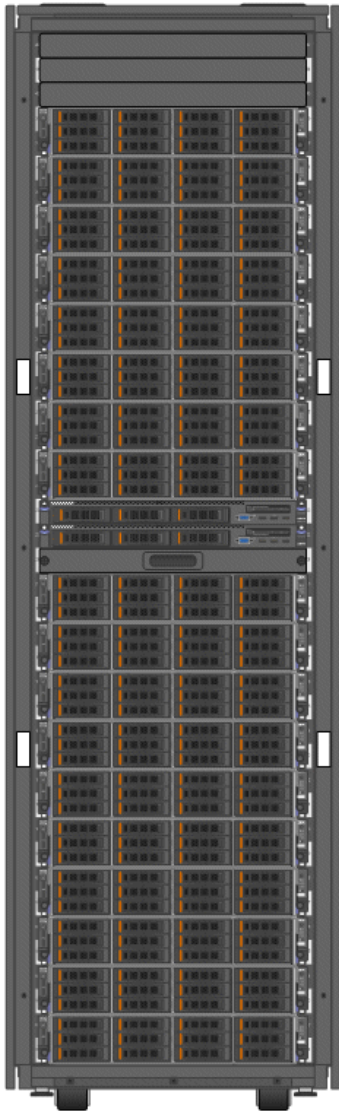
Capabilities	Benefit
Accelerate Big Data Time to Value	<ul style="list-style-type: none">– IBM Accelerators for Text Analytics, Social Data, Machine Data – start analyzing your big data much more quickly and easily– IBM BigSheets – browser-based tool enabling business users to explore data stored in BigInsights clusters and create analytic queries without writing any code– IBM Big SQL – provides the ability to perform SQL queries across your data in Hadoop
Simplify Big Data Adoption & Consumption	<ul style="list-style-type: none">– Hardware Management – advanced web-based, hardware management and monitoring tool, including capabilities turning on and off nodes, rebooting nodes, monitoring hardware health, and hardware status alerts– Deploy system updates with ease
Implement Enterprise Class Big Data	<ul style="list-style-type: none">– High Availability – hardware HA capabilities: dual master nodes, redundant networking switches, power supplies and PDUs– Security – enterprise-level security built-in, including protected master nodes and isolated data nodes– Integrated archival software to help manage data growth– BigInsights Connectors – high-speed connectors for the IBM Netezza family of data warehouse appliances, IBM DB2®, IBM InfoSphere Warehouse and IBM Smart Analytics System– Netezza SQL-side Connectors – Netezza NPS-side read and write to and from HDFS & provides archiving functionality

Robust Enterprise Readiness

- **High Availability**
 - **Hardware: Dual Master Nodes, Redundant Network Switches, Power Failover**
 - **Software: Customized Linux HA**
- **Security**
 - **Protected Master Node with secure port access**
 - **Data Node Isolation with secure port access**
 - **Integration with IBM Guardium - best in class security**
- **Integrated Cluster Management**
 - **Hardware monitoring and management via BigInsights Web Console**



Key Features and Specifications



Key Features

Hadoop Distribution	– InfoSphere BigInsights V2.1
Built-in Analytics/Accelerators	<ul style="list-style-type: none"> – IBM BigSheets – IBM Accelerator for Text Analytics – IBM Accelerator for Social Data – IBM Accelerator for Machine Data – IBM Big SQL
Development / Administration	<ul style="list-style-type: none"> – Eclipse-based Development Environment – Exposed Node Management
Enterprise Readiness	<ul style="list-style-type: none"> – Security – High Availability SW & HW – Hardware management & monitoring
Data Warehouse Integration	<ul style="list-style-type: none"> – Enterprise data warehouse connectors – Archival capabilities

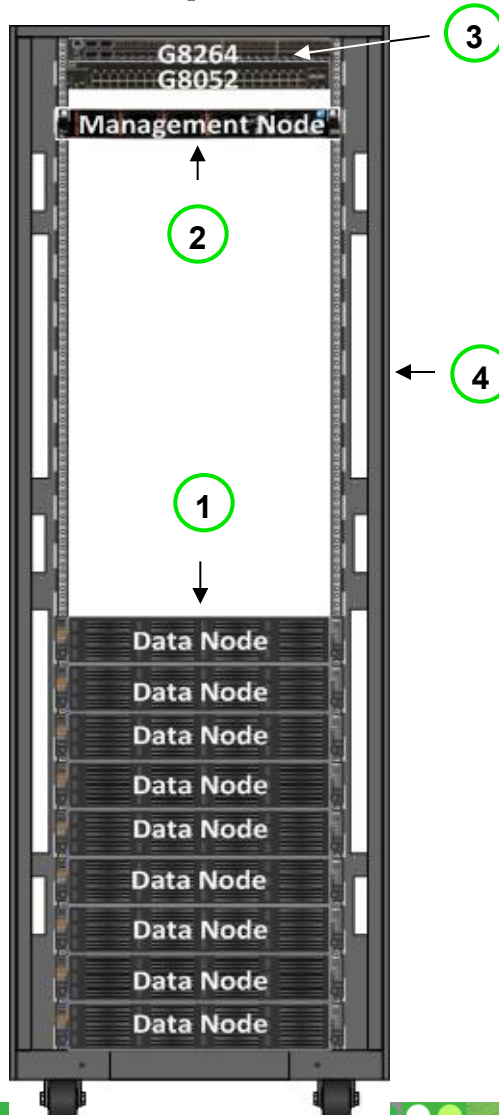
Specifications

Full Rack

Management Nodes	1 primary, 1 standby (x3550 M4)
Data Nodes	18 (x3630 M4)
CPU Cores	216
Memory	96 GB per node, 1728 GB total
Raw Storage	216 drives, 3 TB each. 648 TB total
User Space	216 TB

IBM Reference Architecture

4 main components:



① **Data Node:** the Data node is where the data is kept in a Big Data System. It is typically comprised of CPUs and lots of Storage.

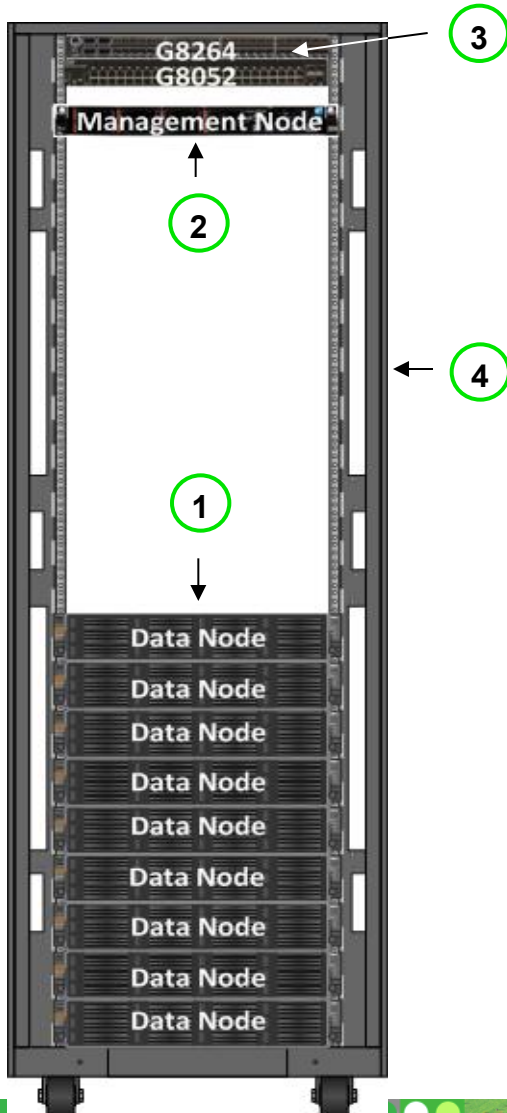
② **Management Node:** the Management node controls the system and keeps track of where the data is and where jobs are running. (Job Tracker, Name node and console functions).

③ **Rack Switches:** the Rack switches control the movement of data within the Rack .. We typically include 2 switches, one for the data traffic and one for the management traffic.

④ **Rack:** All the components are installed in a 42U rack .. As part of the ordering process, all components can be racked by manufacturing.

IBM Reference Architecture

for System x and Power:



System x:

- 1 Data Node: 3630 m4 w/ 12 cores, 48gb memory, 12 -1,2,3tb drives for data, JBOD, 2 drives for OS
- 2 Management Node: 3550 m4 w/ 16 cores, 96gm memory, 4 - 600gb drives
- 3 Rack Switches: G8264 10GbE switch for Data, G8052 1GbE switch for Mgmt
- 4 Rack: Standard 42U rack

Power:

- 1 Data Node: Power 7R2+ w/16 cores, 128gb memory, I/O Exp Unit 24 -900 gb drives for data, 6 drives for OS
- 2 Management Node: Power 7R2+ w/ 16 cores, 96gm memory, 2 - 600gb drives
- 3 Rack Switches: G8264 10GbE switch for Data, G8052 1GbE switch for Mgmt
- 4 Rack: Standard 42U rack

Hadoop System x M4 Pre-defined configurations

Incorporating a balance of value, enterprise and Performance options

Management Node

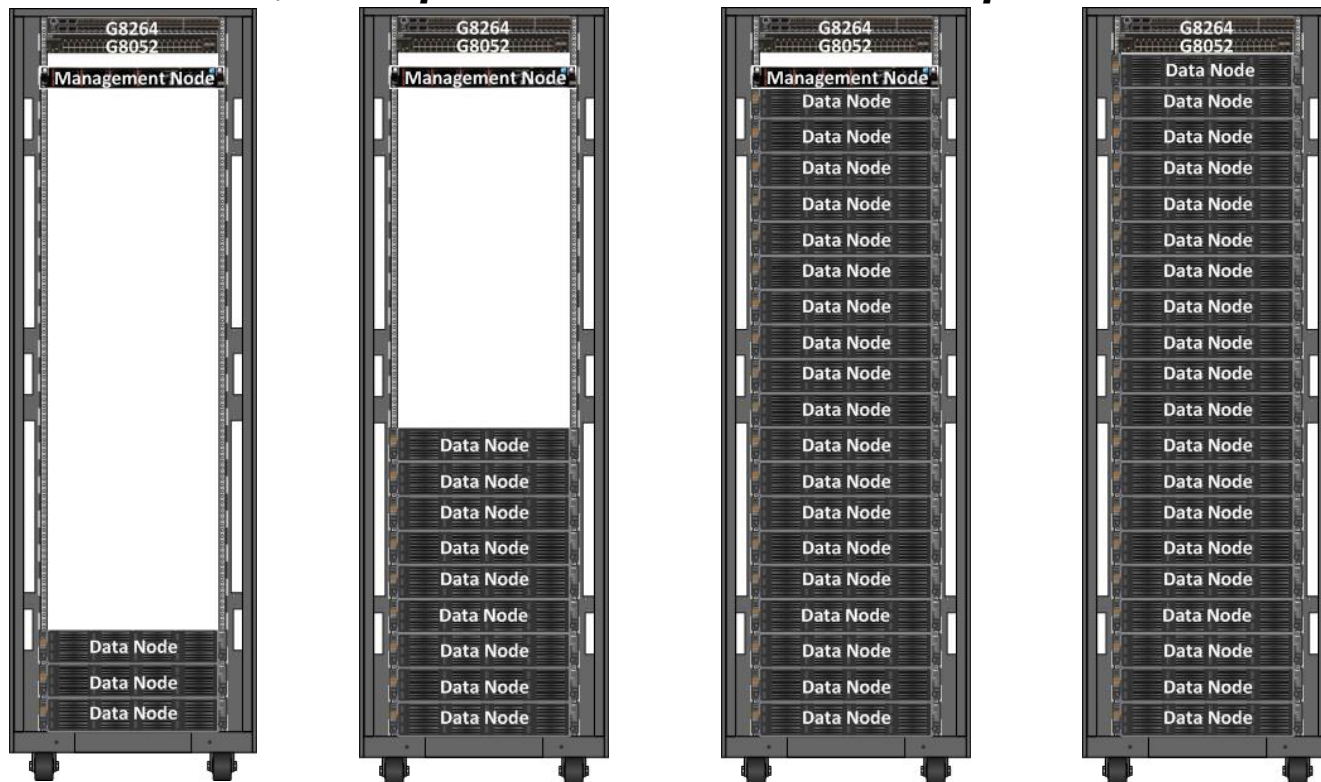
x3550 M4 with

- ✓Two E5-2650 2GHz 8-core CPU
- ✓128GB RAM,
- 16x 8GB 1600MHz RDIMM
- ✓Four 600GB 2.5" HDD (OS)
- ✓Two Dual-port 10GbE (data)
- ✓Dual-port 1GbE (mgmt)

Data Node

x3630 M4 with

- ✓Two E5-2430 2.2GHz 6-core CPU
- ✓48GB RAM,
- 6x 8GB 1333MHz RDIMM
- ✓Two 2TB 3.5" HDD (OS/app)
- ✓Twelve 2TB 3.5" HDD (data)
- ✓Optional 3TB HDD upgrade
- ✓Dual-port 10GbE (data)
- ✓Dual-port 1GbE (mgmt)



Configuration	Starter	Half Rack	Full Rack w/Mgmt Node*	Full Data Node Rack*
Available storage (2TB/3TB)	72TB / 108TB	216TB / 324TB	456TB / 684TB	480TB / 720TB
Raw data space (2TB/3TB)	18TB / 27TB	54TB / 81TB	114TB / 171TB	120TB / 180TB
Mgmt Nodes / Data Nodes	1 Mgmt / 3 Data	1 Mgmt / 9 Data	1 Mgmt / 19 Data	0 Mgmt / 20 Data
Switches	1 x 10GbE / 1 x 1GbE	1 x 10GbE / 1 x 1GbE	1 x 10GbE / 1 x 1GbE	1 x 10GbE / 1 x 1GbE

* Number of management nodes required varies with cluster size and workload; for multi-rack configs, select combination of these racks as needed

22 Assumes Hadoop replication with 3 copies of data, no compression, 25% capacity reserved for map/reduce "shuffle" files; because compression and shuffle space can vary widely by workload, this estimate provided as a starting point only

Pre-defined PowerLinux configurations

Management Node

PowerLinux 7R2+

- ✓Sixteen 3.6 or 4.2 GHz cores
- ✓96GB RAM
- ✓Two 600GB 2.5" HDD (OS)
- ✓DVD drive
- ✓Dual-port 10GbE (data)
- ✓Dual-port 1GbE (mgmt)

Data Node

PowerLinux 7R2+

- ✓Sixteen 3.6 or 4.2 GHz cores
- ✓128GB RAM
- ✓Two 600GB 2.5" HDD (OS)
- ✓Four 900GB 2.5" HDD (app)
- ✓2.5" I/O Exp 24 900GB (data)
- ✓Dual-port 10GbE (data)
- ✓Dual-port 1GbE (mgmt)



Configuration	Full Rack 1:1	Full Rack 2:1	Full Rack 4:1
Avail Storage	255 TB	172TB	144TB
Raw data **	42TB	29TB	24 TB
Data Nodes, cores, threads	10, 160, 640	12, 192, 768	16, 256, 1024
Disk Expansion Units	10	6	4
Management Nodes	1 (displaces 1 Data/Exp pair)	1 (uses avail free 4U space)	1 – (displaces 1 data node)
Switches	1 x 10GbE / 1 x 1GbE	1 x 10GbE / 1 x 1GbE	1 x 10GbE / 1 x 1GbE

* Number of management nodes required varies with cluster size and workload; for multi-rack configs, select combination of these racks

**as needed
Assumes Hadoop replication with 3 copies of data, no compression, 25% capacity reserved for map/reduce “shuffle” files; because compression and shuffle space can vary widely by workload, this estimate provided as a starting point only

IBM Intelligent Cluster: Improving Time to Value

Which would you rather deal with . . .



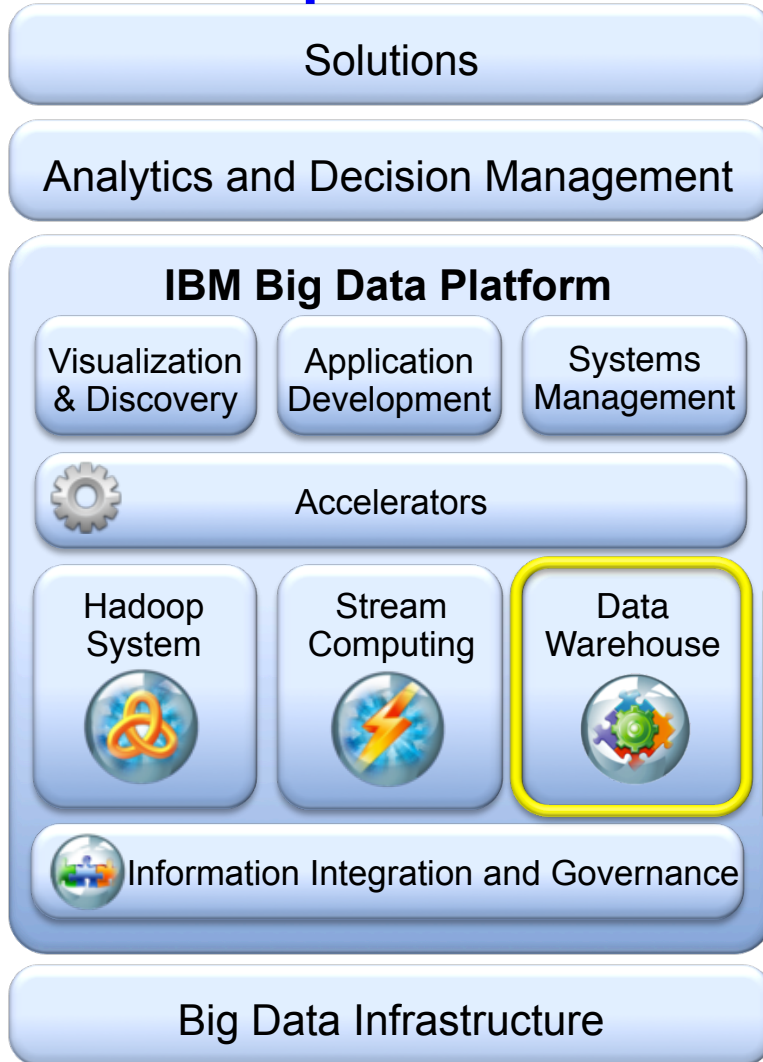
or



Intelligent cluster

Integrated solutions that include
Servers, switches, cabling, software*;
Ready to roll into your data center, apply power
and networking – and you're ready to begin
working

Part of the IBM Big Data Platform workload optimized solutions for your analytic needs



*PureData System
for Analytics*

New



*PureData System for
Operational Analytics*

New



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

PureData System for Operational Analytics hardware



- IBM POWER7 P740 & P730 16 Core servers @ 3.55GHz

- IBM Storwize® V7000 with 900GB drives
- Ultra SSD I/O Drawers, each with six 387GB SSD

- Blade Network Technologies 10G and 1G Ethernet switches
- Brocade SAN switches (SAN48B-5)

Scales to PB
+ capacity*

DB2 Analytics Accelerator for z/OS

PureData System for Analytics connected to System z only accessible through DB2

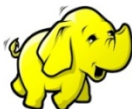
Blending System z and Netezza technologies to deliver unparalleled, mixed workload performance for complex analytic business needs.



What is the value?

- Fast, predictable response times for “right-time” analysis
- Fast Time to Value. (Deployment in days, not months.)
- Improve price/performance for analytic workloads
- Minimize the need to create data marts for performance
- Highly secure environment for sensitive data analysis
- Transparent to the application and user

A smarter consolidated infrastructure



Workload Manager(s)

Multi-tenant shared service environment



Resource Orchestration

Provisioning & Management



Enterprise Storage

IBM PLATFORM SYMPHONY

Enhanced workload scheduling and
resource sharing for big data

IBM Platform Symphony



- A heterogeneous grid management platform
- A high-performance SOA middleware environment
- Supports diverse compute & data intensive applications
 - Compute and Data intensive ISV analytic applications
 - In-house analytic applications (C/C++, C#/ .NET, Java, Excel, R etc)
 - Optimized low-latency Hadoop compatible run-time
 - Can be used to launch, persist and manage non-grid aware application services
- React instantly to time critical-requirements
- Production proven multi-tenancy with resource sharing capabilities
- Embedded single-tenant license in InfoSphere BigInsights 2.1

IBM Platform Symphony for big data

SLA management

Resource Plan

Resource Group: ComputeHosts Time Intervals and Settings

Slot allocation policy

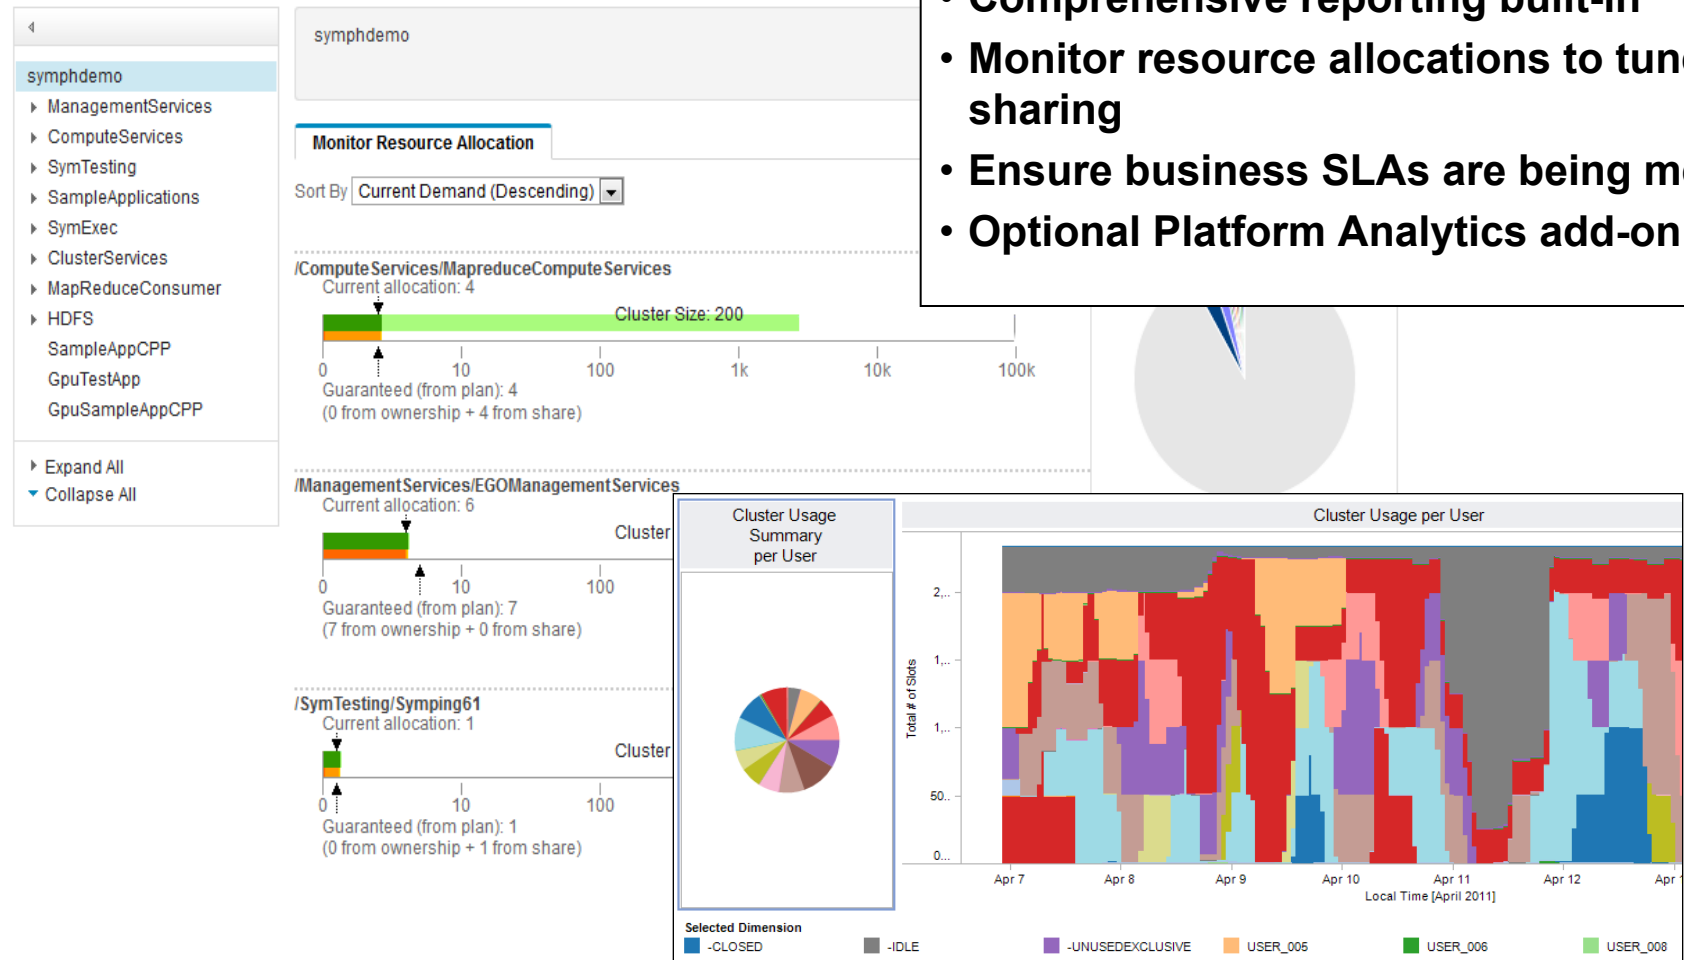
Consumer	Owned Slots	Consumer Rank	Lend Limit	Borrow Limit	Share
Model type: Ownership	00:00:00 24:00				
▼ symphdemo	108				
▼ SymTesting	0	0			<input checked="" type="checkbox"/>
■ Sympling61	0	0	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Total	0	-	-	-	-
Balance	0	-	-	-	-
▶ SampleApplications	0	0			<input checked="" type="checkbox"/>
▼ SymExec	40	0			<input checked="" type="checkbox"/>
■ SymExec61	0	0	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
■ AlgoConsumer	40	0	<input checked="" type="checkbox"/> 20 Details	<input type="checkbox"/>	<input checked="" type="checkbox"/>
■ AlgoRiskwatch	0	0	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
■ HedgeAnalysis	0	0	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Total	40	-	-	-	-
Balance	0	-	-	-	-
▼ MapReduceConsumer	20	0			<input checked="" type="checkbox"/> 1
■ MapReduce61	4	0	<input checked="" type="checkbox"/> Details	<input checked="" type="checkbox"/> Details	<input checked="" type="checkbox"/> 1
■ MapReduceHighPriority	8	0	<input checked="" type="checkbox"/> Details	<input checked="" type="checkbox"/> Details	<input checked="" type="checkbox"/> 1
■ MapReduceDefault	8	0	<input type="checkbox"/>	<input checked="" type="checkbox"/> Details	<input checked="" type="checkbox"/> 1
■ mrsleep	0	0	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/> 1
Total	20	-	-	-	-
Balance	0	-	-	-	-
■ SampleAppCPP	0	0	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/> 1
Total	60	-	-	-	-
Balance	48	-	-	-	-

- Guarantee minimum quality of service
- Time-variant sharing policies
- Multiple resource sharing models
- Granular, directed sharing
- Configurable pre-emption policies
- Maintain multiple versions of application services to simplify life-cycle management
- Share resources between Dev, Test, Production & QA application instances

IBM Platform Symphony for big data

Reporting and Analytics

Resource Allocation

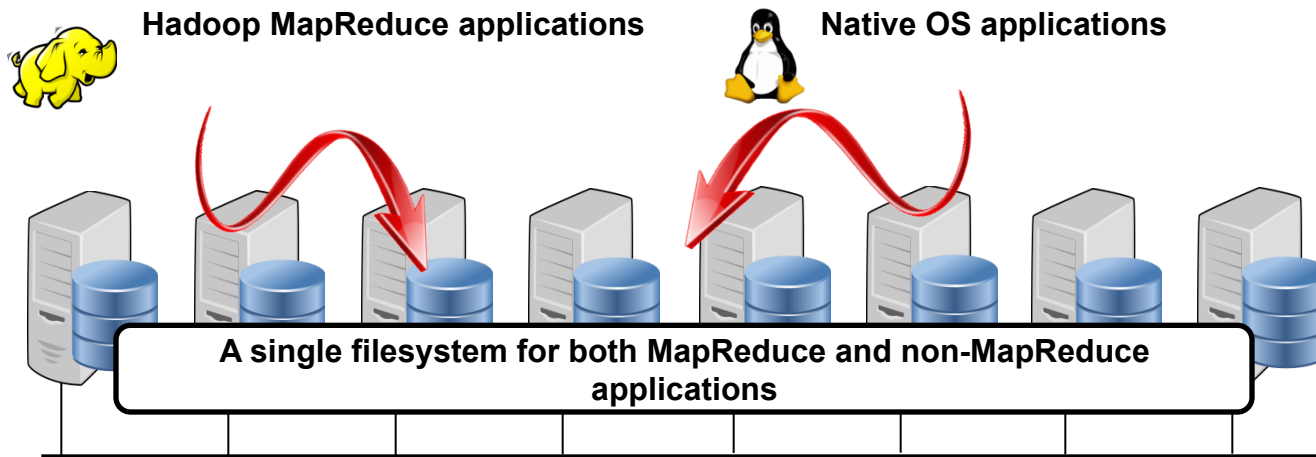


IBM GPFS

Bringing enterprise file system capabilities to big data

GPFS – bringing new capabilities to **big data**

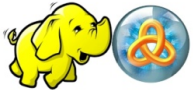
POSIX file system



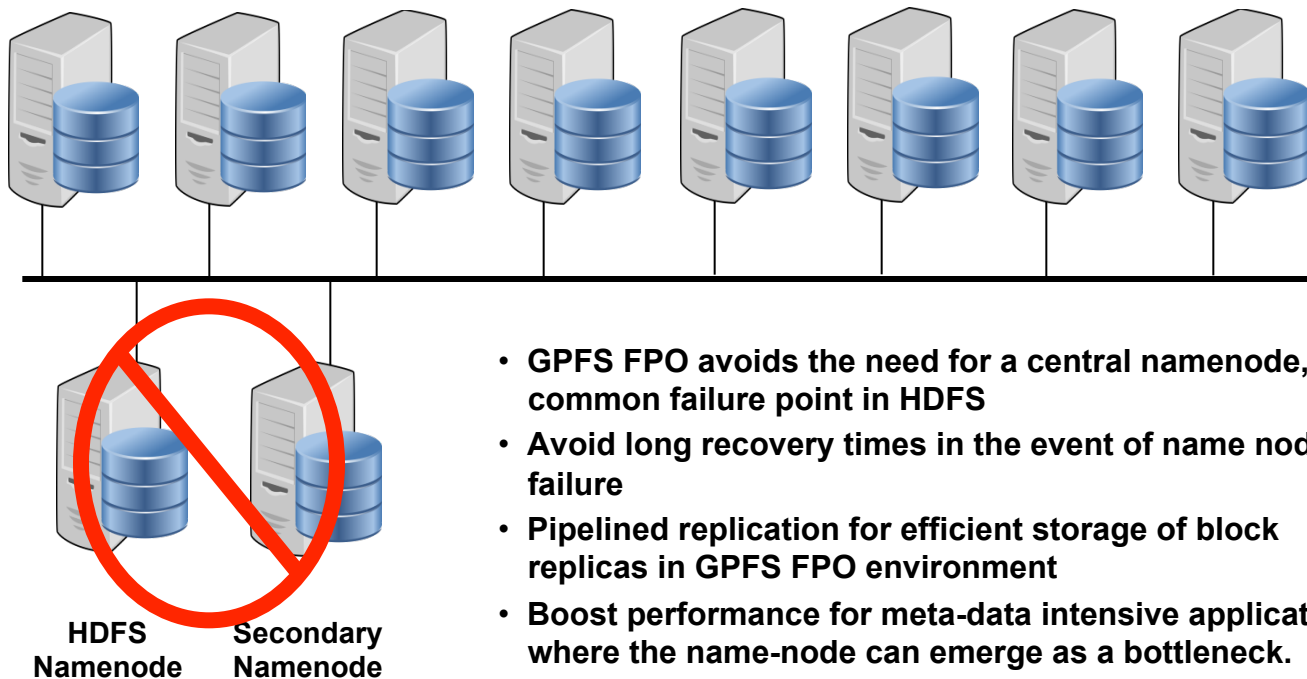
- Native POSIX file system
- Avoid needless data movement and replication
- Variable block-sizes provide good performance across a range of workloads

GPFS – bringing new capabilities to big data

File system reliability



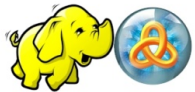
IBM BigInsights cluster with GPFS FPO



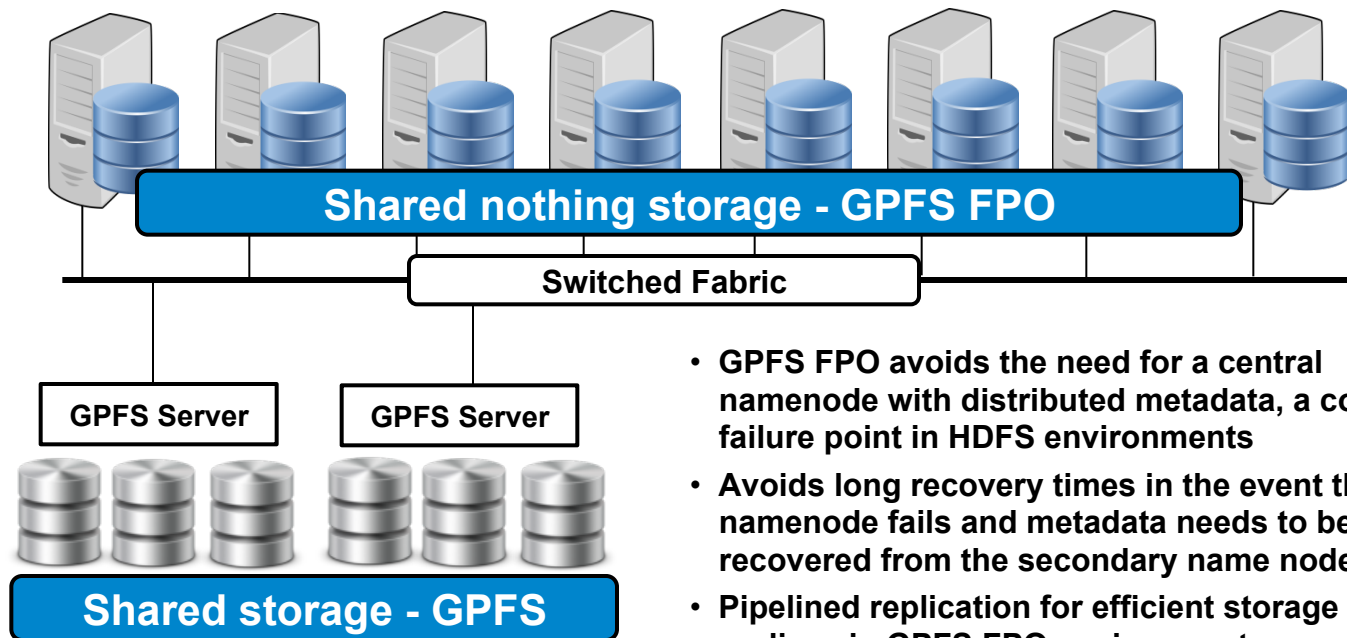
- GPFS FPO avoids the need for a central namenode, a common failure point in HDFS
- Avoid long recovery times in the event of name node failure
- Pipelined replication for efficient storage of block replicas in GPFS FPO environment
- Boost performance for meta-data intensive applications where the name-node can emerge as a bottleneck.

GPFS – bringing new capabilities to big data

Flexible storage configuration

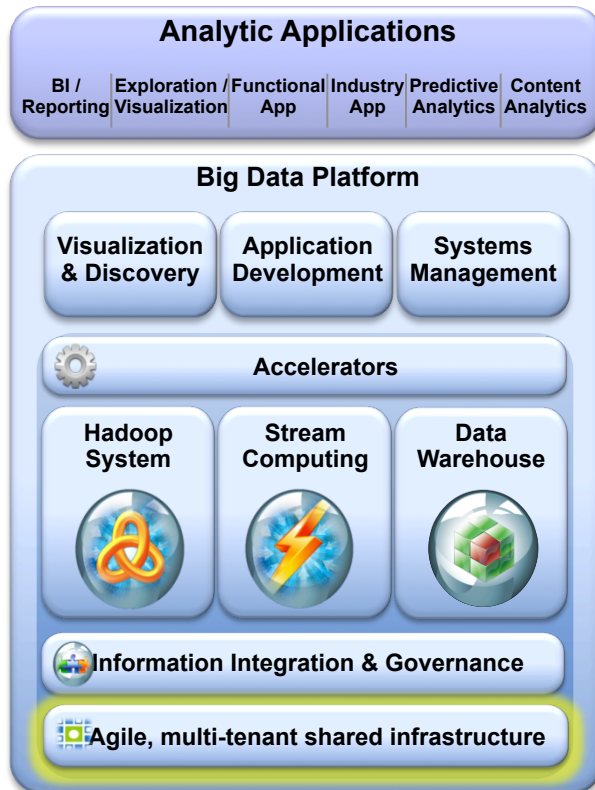


IBM BigInsights cluster with GPFS FPO



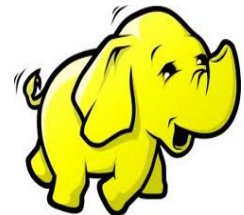
- GPFS FPO avoids the need for a central namenode with distributed metadata, a common failure point in HDFS environments
- Avoids long recovery times in the event that the namenode fails and metadata needs to be recovered from the secondary name node
- Pipelined replication for efficient storage of block replicas in GPFS FPO environment

IBM – a better infrastructure for big data



Technical computing for Big Data

- ✓ Improved performance
- ✓ More efficient use of infrastructure
- ✓ Diverse, concurrent workloads
- ✓ Dynamic Ease of management
- ✓ Guaranteed service levels
- ✓ Proven reference architectures
- ✓ State-of-the-art hardware & software
- ✓ Proven capabilities



Get Started on Your Big Data Journey Today

Get Educated

- IBM Big Data: ibm.com/bigdata
- IBMBigDataHub.com
- BigDataUniversity.com
- IBV study on big data
- Books / analyst papers

Schedule a Big Data Workshop

- Free of charge
- Best practices
- Industry use cases
- Business uses
- Business value assessment



THINK

BIG

BIG