BIG DATA – technológie a referenčné architektúry v podaní IBM

Big*/Fast* Data and Needs **Big***/Fast* Analytics

Maximum Insight Requires Combining Deep and Reactive Analytics

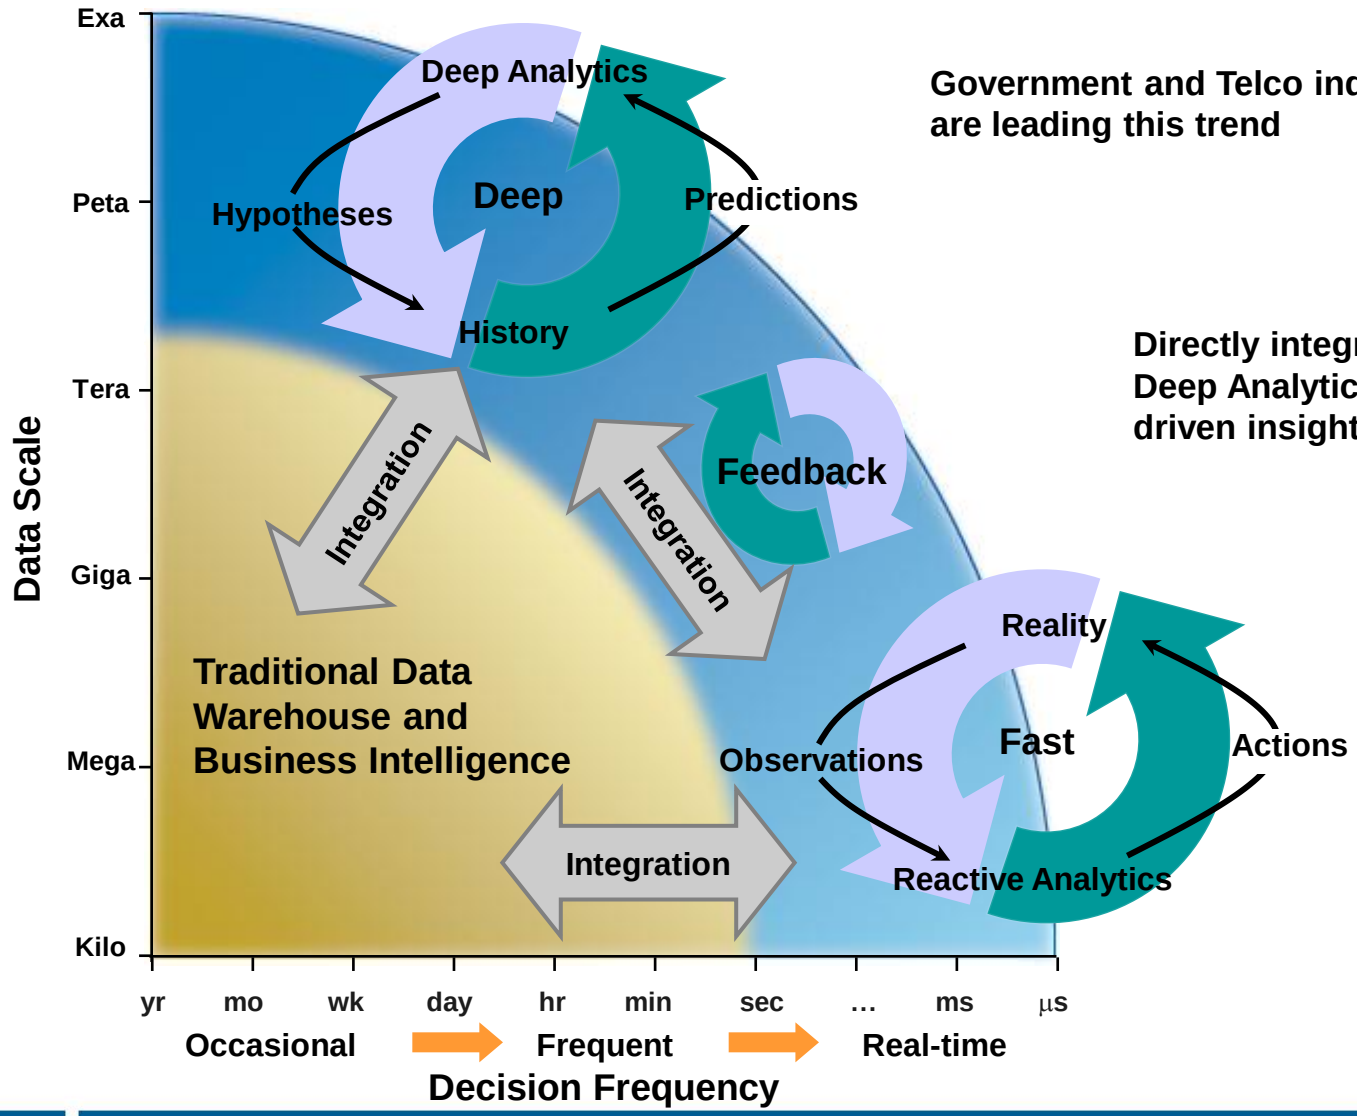

Government and Telco industries are leading this trend

Directly integrating Reactive and Deep Analytics enables feedback-driven insight optimization

IBM investments in Analytics and Big Data

- ✓ More than **$16B** in Acquisitions Since 2005
- ✓ More than **10,000** Technical Professionals
- ✓ More than **7,500** Dedicated Consultants
- ✓ **Largest** Math Department in Private Industry
- ✓ More than **27,000** Business Partner Certifications

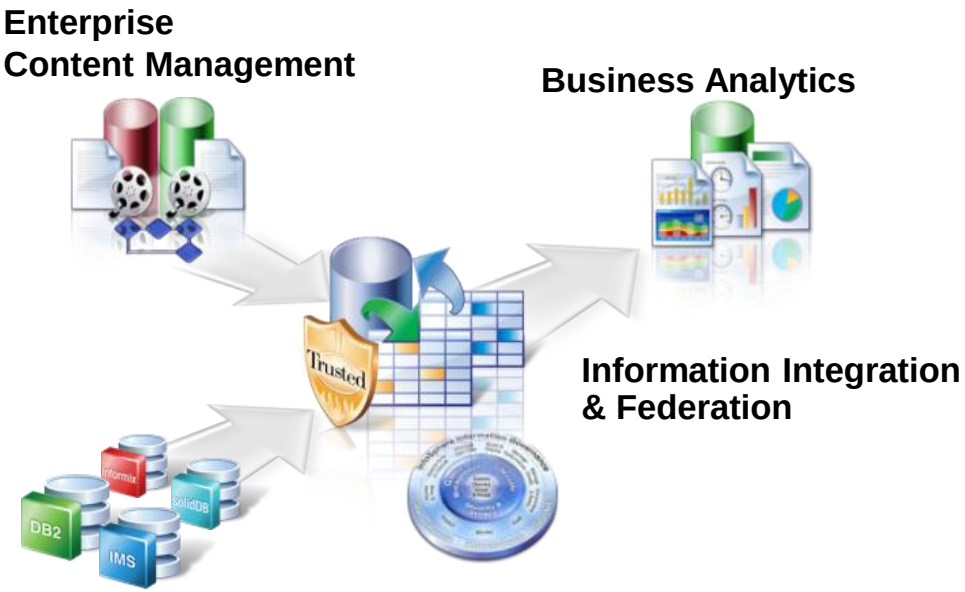

Analyze Information of any Variety

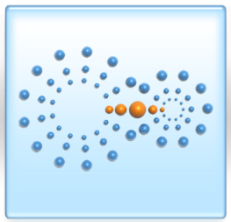

Analyze Information in Motion

Analyze Extreme Volumes of Information

Discover & Experiment

Manage & Plan

Data Warehouse and BigInsights Comparison Chart

	Data Warehouse	Hadoop / Streams
Data Types	Largely structured data	Any type of data, structured or unstructured
Data Loading	Data is cleansed/structured before going into the warehouse to maximize its utility	Raw data may be ingested as is, without any modification, as the relationships may not be understood or defined
Reliability	ACID compliant (Atomicity, Consistency, Isolation, Durability)	Not ACID compliant
Integrity	Database maintains integrity	Applications code integrity
Analytic Approach	• ***High value***, structured data • ***Repeated*** operations and processes (e.g. transactions, reports, BI, etc.) • Relatively ***stable*** sources • Well-understood requirements • Optimized for fast access and analysis	• ***Highly variable*** data and content • ***Iterative***, exploratory analysis (e.g. scientific research, behavioral modeling) • ***Volatile*** sources • Ill-defined questions and changing requirements • Optimized for flexibility
Hardware	Powerful appliance and optimized hardware	Inexpensive, commodity hardware

Background: 3 Styles of Massively Parallel Systems

Hadoop/MapReduce (BigInsights)

JAQL, Java

Data at Rest*:

High Volume

Mixed Variety

Low Velocity

Deep Analytics
Extreme Scale-out

(*pre-partitioned)

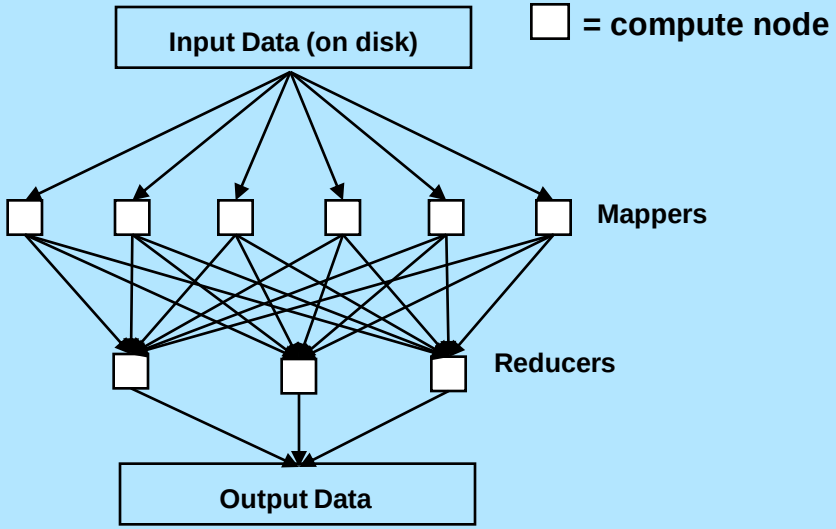

Streaming (Streams)

SPL, C, Java

Data in Motion:

High Velocity

Mixed Variety

High Volume*

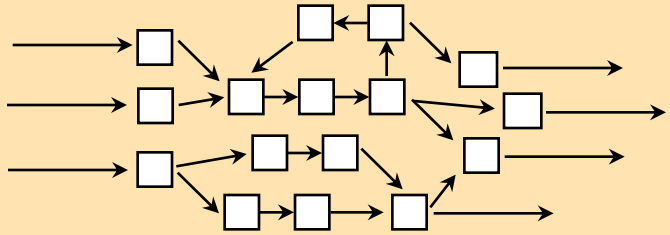

Reactive Analytics
Extreme Ingestion

(*over time)

Simulation (BlueGene)

C/C++, Fortran, MPI, OpenMP

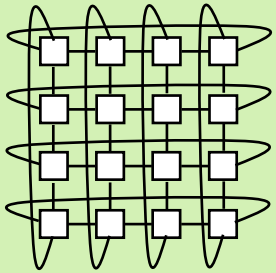

Long Running

Small Input

Massive Output

Generative Modeling
Extreme Physics

Hadoop in a Nutshell

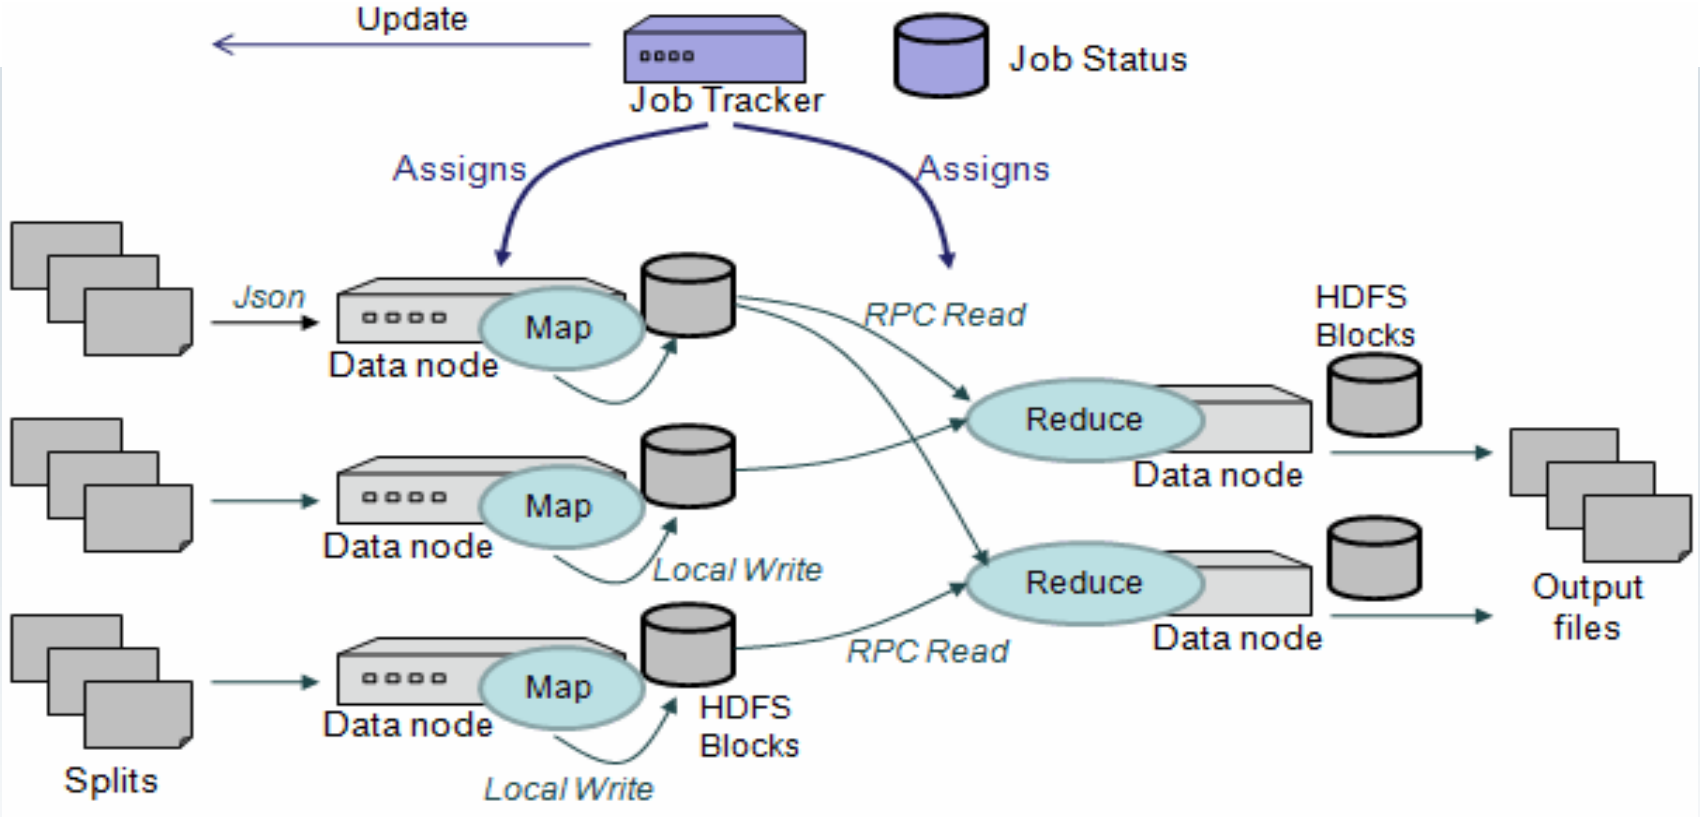

An open-source computing platform that is both distributed and redundant, handles structured and unstructured data and supports a simple and efficient programming paradigm called map-reduce.

IBM Platform for Big Data Analytics

InfoSphere BigInsights and InfoSphere Streams

- Analytics for data in-motion and at-rest
- Platform for processing large volumes of diverse data
- Complements and integrates with existing software solutions

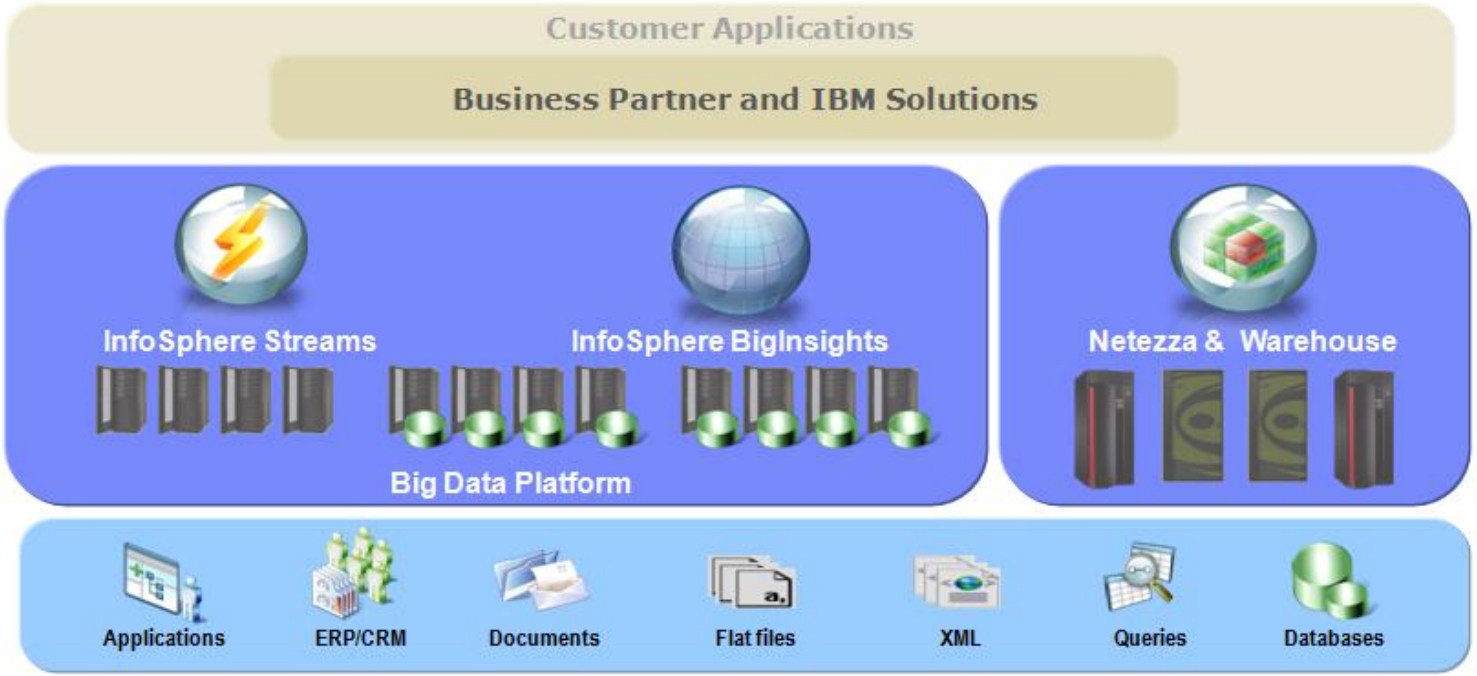

Scale-in is the New Systems Battlefield

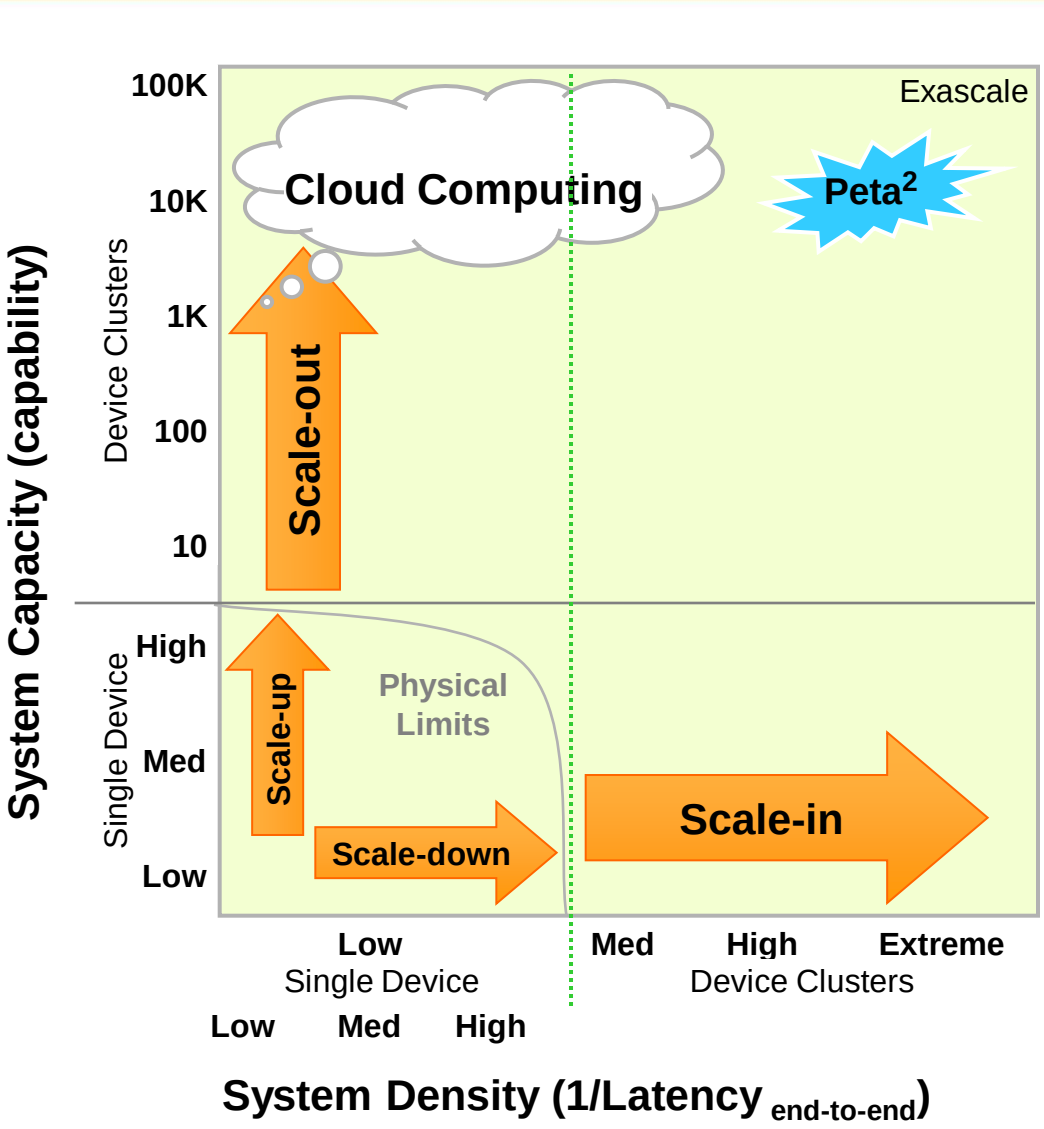

Scale-down
Maximize feature density

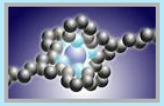
Atom Transistor


Atom Storage

Scale-up
Maximize device capacity

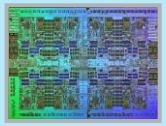
POWER 7


Terabyte HDD

Scale-out
Maximize system capacity


Blade Server


NAS

Scale-in
Maximize system density
Minimize end-to-end latency

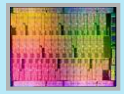
Manycore


FPGA


3D Chips


BPRAM/SCM


FLASH SSD

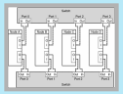
Interconnect


In-mem DB


DAS

Massive Scale Analytics

Heterogeneous Workloads (IO/CPU)

- Unlike typical warehouse workloads, some parts in an analytics flow are CPU intensive and some parts are IO intensive.
- Need a platform that can handle structured, semi structured, and unstructured data
- Need a platform that can handle many such workloads – custom infrastructure for each massive scale analytics problem is not viable
- Hardware optimization where possible

- Scalability and Elasticity
  - Data volumes at tens to hundreds of Terabytes, growing to Petabytes. Need to have a path for customers to scale the platform to meet growing demands.
  - Need a platform that easily supports parallel analytics – scalable, elastic, fault-tolerant MapReduce

Educate yourself on Big Data-centric Architectures for Performance

Old Compute-centric Model

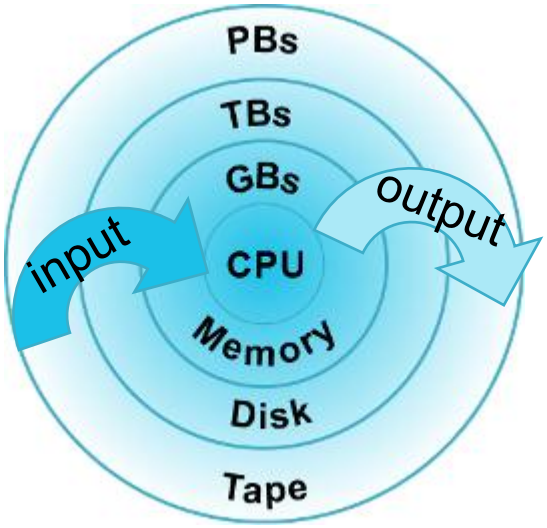

Data lives on disk and tape
Move data to CPU as needed
Deep Storage Hierarchy

Manycore **FPGA**

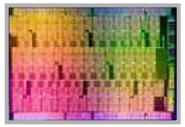

Massive Parallelism
Persistent Memory

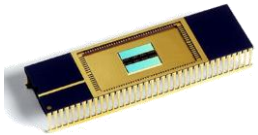

Flash **Phase Change**

New Data-centric Model

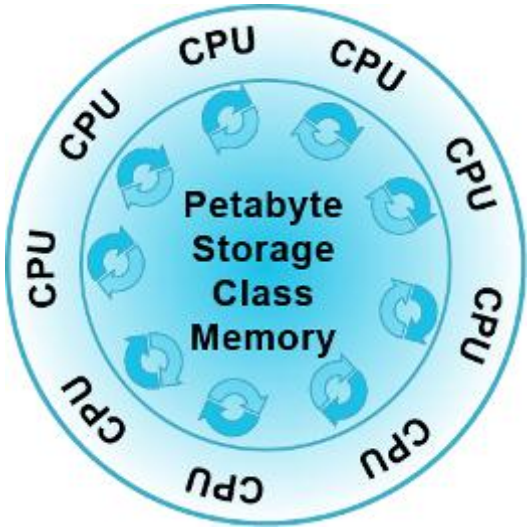

Data lives in persistent memory
Many CPU's surround and use
Shallow/Flat Storage Hierarchy

Largest change in system architecture since the System 360
Having a huge impact on hardware, systems software, and application design

Watson: A Workload Optimized System

- 90 x IBM Power 750 servers
- 2,880 POWER7 cores
- POWER7 3.55GHz chip
- 500GB per sec on-chip bandwidth
- 10Gb Ethernet network
- 15 Terabytes of memory
- 20 Terabytes of disk, clustered
- Operates at up to 80 Teraflops
- Runs IBM DeepQA software stack
- Scales out with and searches vast amounts of unstructured information with UIMA & Hadoop open source components
- SUSE Linux provides a cost-effective open platform which is performance-optimized to exploit POWER 7 systems
- 10 racks include servers, networking, shared disk system, cluster controllers

Trading example of Real-Time Stream Computing

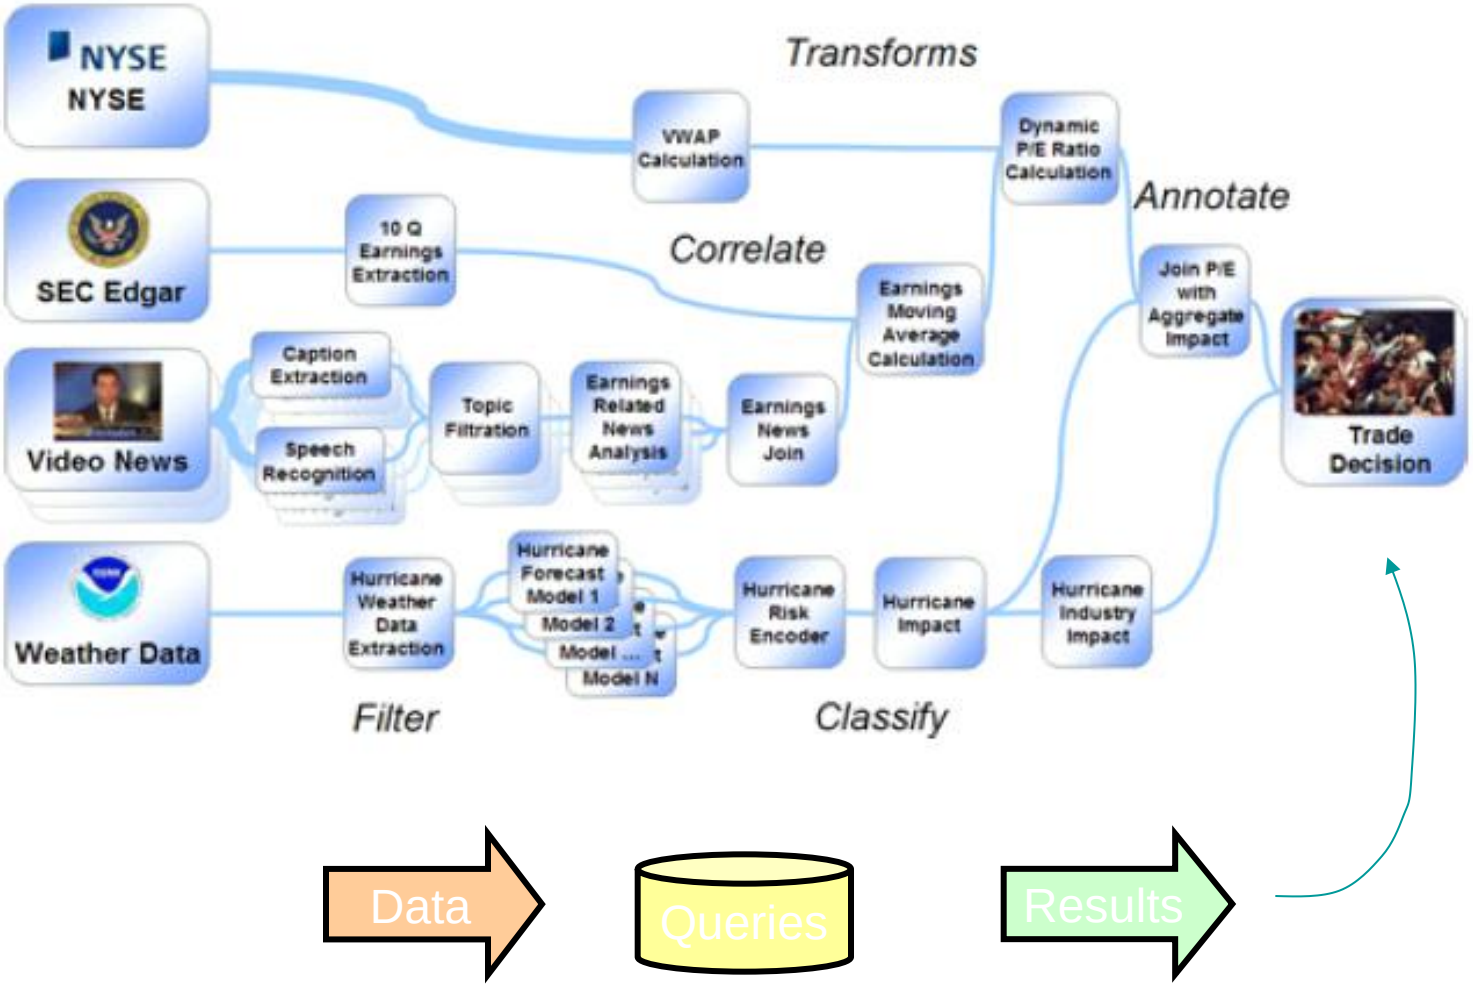

What is a typical Big Data system ?

Emerging “Hot” Analytics area

- The Hadoop EcoSystem is rapidly expanding
- Environment is maturing, moving to more general acceptance

Typical Deployments

- Average cluster size is 120 nodes
- 44% of clusters are between 10 and 100 nodes
- 52% of clusters are between 100 and 1,000
- Largest cluster is over 20 PB
- 13.1% have a cluster bigger than 100 TB
- 12.8% have a cluster bigger than 1 PB

IBM Smart Analytics System *Design Principles*

The Smart Analytics System architecture has been guided by these key design principles which were validated with the field and our largest BI customers:

- Balanced performance
- Stability
- Price to performance ratio
- Scalability
- Fault tolerance
- High availability
- Ease of installation and implementation
- Packaging density
- Industry-standard components

Simplicity of solution portfolio elements

Value Configuration – Switch HA using vLAG

VLAN tagging of Data and Admin network using vLAG of 2x 1Gb links over 2 G8052s in 2 racks

Each x3630 M3 has 2 integrated 1Gb ports and one IMM 100Mb port

One G8052 switch per rack, each with 48 x 1Gb ports & 4 Std. 10Gb uplinks

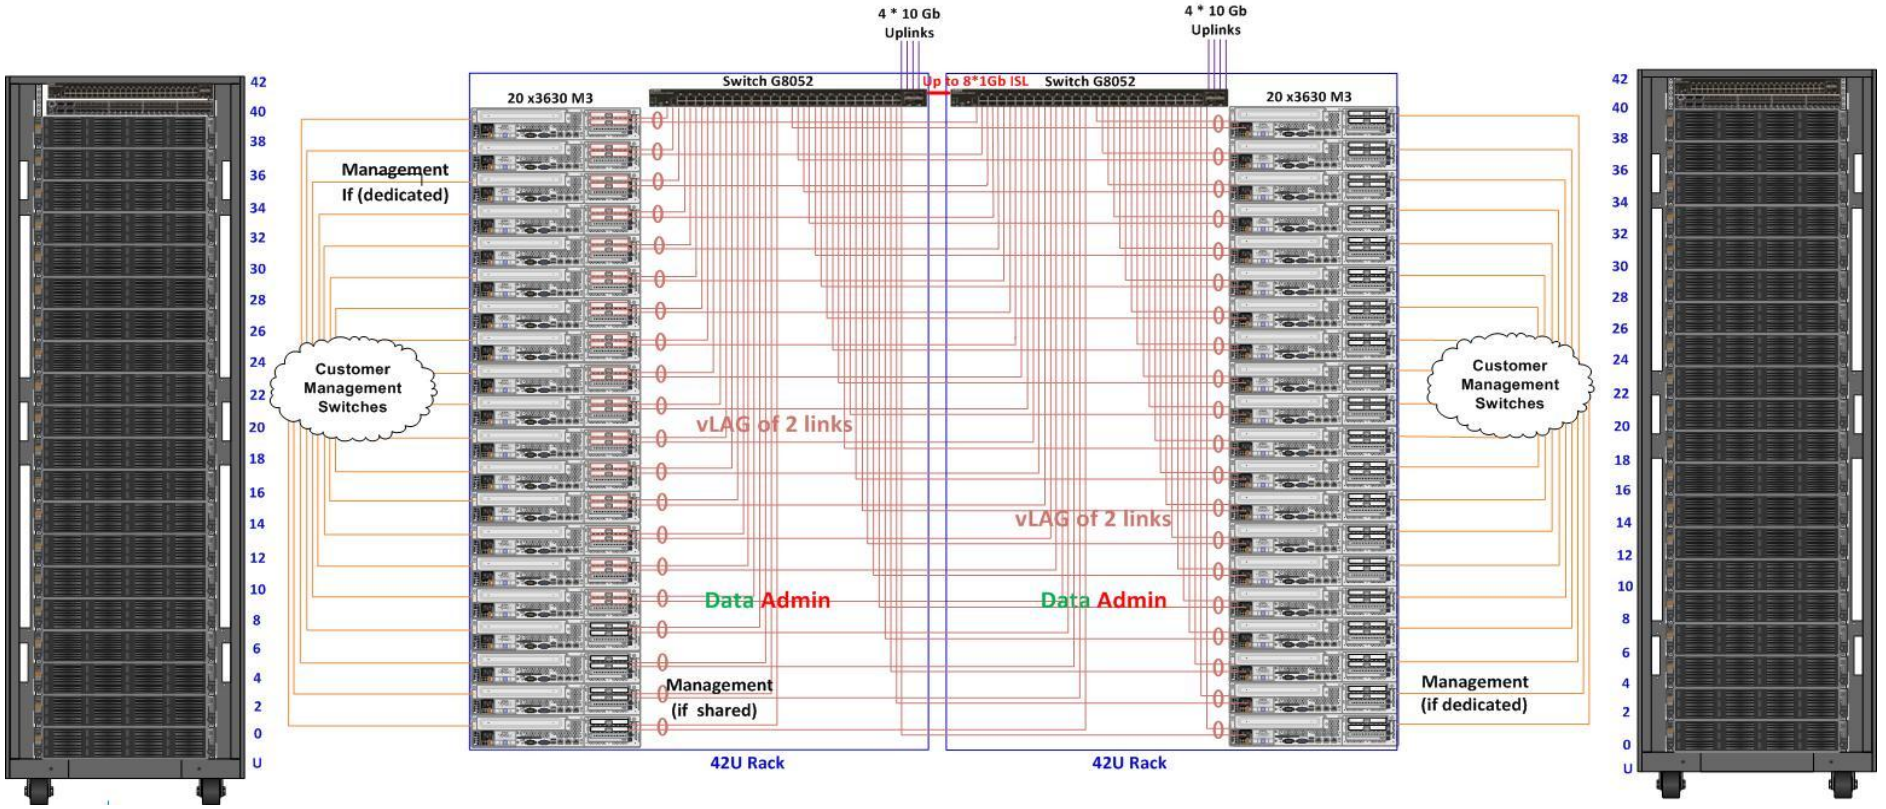

Pre-defined configurations

Incorporating a balance of value, enterprise and Performance options

Management Node

x3550 M3 with
üTwo 3.06GHz 6-core CPUs
ü96GB RAM
üTwo 900GB 2.5” HDD (OS)
üTwo 900GB 2.5” HDD (app)
üDVD drive
üDual-port 10GbE (data)
üDual-port 1GbE (mgmt)

Data Node

x3630 M3 with
üTwo 3.06GHz 6-core CPUs
ü48GB RAM
üTwo 2TB 3.5” HDD (OS/app)
üTwelve 2TB 3.5” HDD (data)
üOptional 3TB HDD upgrade
üDual-port 10GbE (data)
üDual-port 1GbE (mgmt)

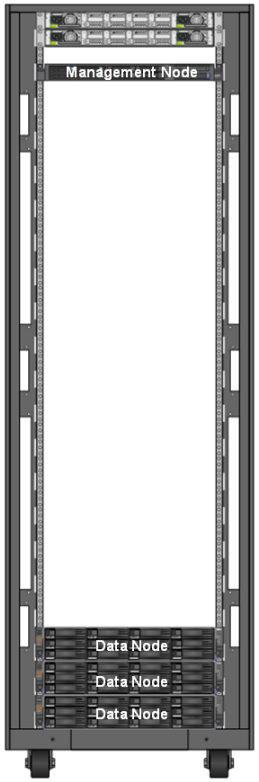

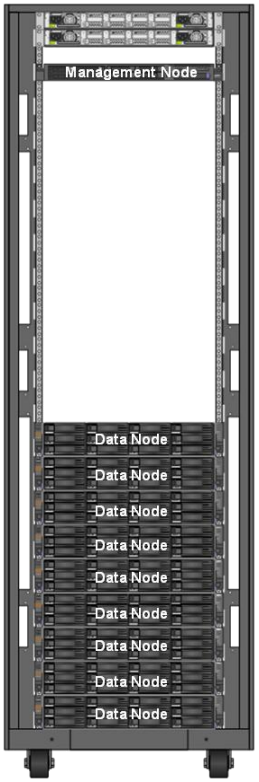

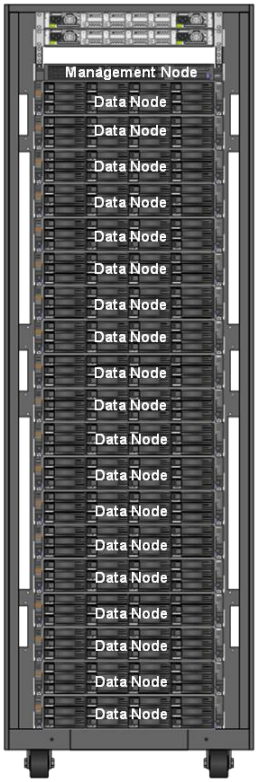

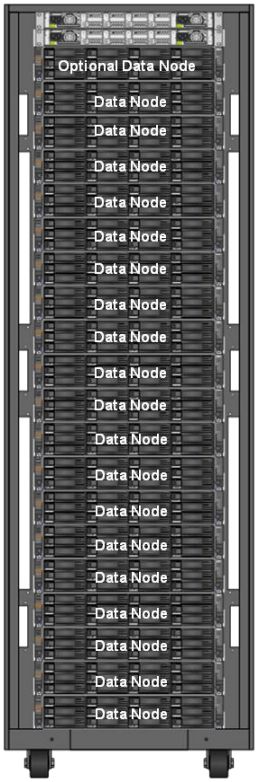

Configuration	Starter	Half Rack	1st Thru 4th Full Rack*	Additional Full Racks
Usable Storage (2TB / 3TB)	72TB / 108TB	216TB / 324TB	456TB / 684TB	456(480)TB / 684(720)TB
User space (2TB / 3TB)	24TB / 36TB	72TB / 108TB	152TB / 228TB	152(160)TB / 228(240)TB
Mgmt Nodes / Data Nodes	1 Mgmt / 3 Data	1 Mgmt / 9 Data	1 Mgmt / 19 Data	0 Mgmt / 19 (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

* 1st and 2nd rack only if using GPFS-SNC instead of HDFS

BigInsights Hardware Foundation

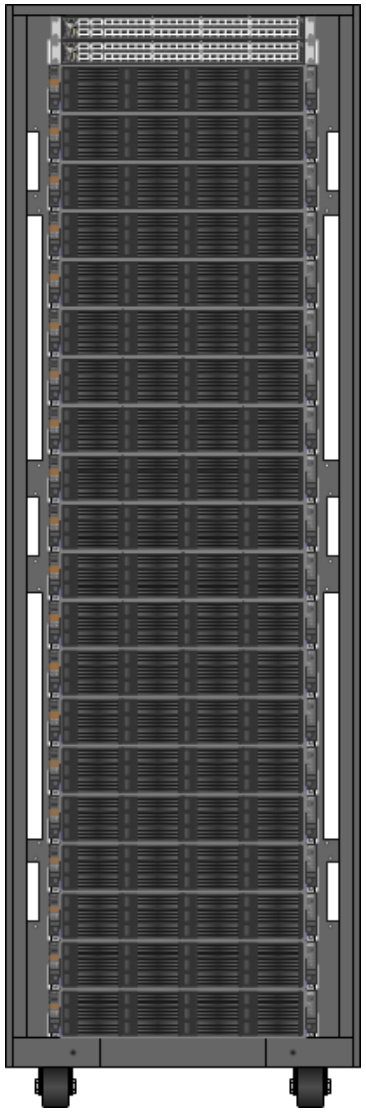

Rack-Level Features

- ✓Up to 20 System x3630 M3 nodes
- ✓Up to 840TB storage
- ✓Up to 240 cores
- ✓Up to 3,840GB memory
- ✓Up to two 10Gb Ethernet (IBM G8264) or 40Gb InfiniBand switches
- ✓Scalable to multi-rack configurations

Available Enterprise and Performance Features

- ✓Redundant storage
- ✓Redundant networking
- ✓High performance cores
- ✓Increased memory
- ✓High performance networking

Backup Slides

BigInsights Hardware Foundation - Key Components

Value Data Node

- ✓IBM System x3630 M3
- ✓Two Intel Xeon E5620 CPUs
- ✓Data: 12 x 2TB NL SAS HDDs, JBOD
- ✓OS: 1 x 2TB NL SAS HDD
- ✓48GB DDR3 RDIMMs

Value Management Node

(JobTracker, NameNode, Console)

- ✓IBM System x3550 M3
- ✓Two Intel Xeon X5675 CPUs
- ✓OS: 4 x 900GB NL SAS HDD, mirrored
- ✓DVD drive
- ✓96GB DDR3 RDIMMs

1GbE Switch

- ✓IBM RackSwitch G8052
- ✓48 × 1 GbE RJ45 ports and four standard 10 GbE SFP+ ports
- ✓Low 130 W power rating and variable speed fans to reduce power consumption

10GbE Switch (optional)

- ✓IBM RackSwitch G8264
- ✓Optimized for applications requiring high bandwidth and low latency
- ✓Up to 64 1 Gb/10 Gb SFP+ ports, four 40 Gb QSFP+ports
- ✓1.28 Tbps non-blocking throughput
- ✓ Energy efficient cost effective design

Big Data on IBM PowerLinux

The performance of Power at an x86 price point

IBM Big Data offerings being ported to run natively on PowerLinux

- InfoSphere Streams will GA March 30th 2012
- BigInsights beta in April, GA in June 2012

Positioning

- Existing Power Customers
- Competitive x86 installs
- Particularly well-suited to CPU-intensive workloads

Key Advantages for Big Data

- Potential to leverage Power's multi-thread capability (up to 64 threads of simultaneous execution vs up to 24 on comparable x86) and greater memory bandwidth
- A cluster built on 7R2 requires up to 50% fewer nodes

Successes

- Watson integrated POWER7 processors, IBM DeepQA software technology + Hadoop on Linux
- Commercializing Watson technology
  - WellPoint will use Watson's data-crunching on Power 750s running Linux to help suggest treatment options and diagnoses to doctors
  - IBM and Nuance to combine IBM's DeepQA with Nuance's speech recognition and clinical solutions for patient diagnosis and treatment

2012

PowerLinux 7R2

8246-L2C

2 socket rack, 16-core

3.55 GHz or 3.3 GHz

2x memory and 2x I/O bandwidth

Early Adopter Sales Program

- Limited quantities prior to May 4, 2012 GA
- Requires L112 and approval via RPQ # 8A1985

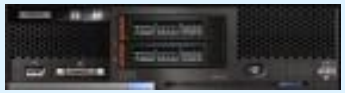

PowerLinux Blade

1457-7FL

2 socket compute node

16 core, 3.55 GHz or 3.22 GHz

12 core, 3.72 GHz

IBM Data Warehouse and Analytics Solutions

IBM Data Warehouse & Analytics Solutions

IBM Netezza	IBM Smart Analytics System	IBM InfoSphere Warehouse
		
True Appliance	**Flexible Integrated System**	**Custom Solution**
Simplicity and fastest time to value	5600S – when TCA is important 7700 – lowest cost per workload 9600 – when data resides on z	Unique needs

Simplicity ← *The right mix of simplicity and flexibility* → **Flexibility**