Big Data System and Architecture

Jian Li
IBM Research in Austin
Email: jianli@us.ibm.com

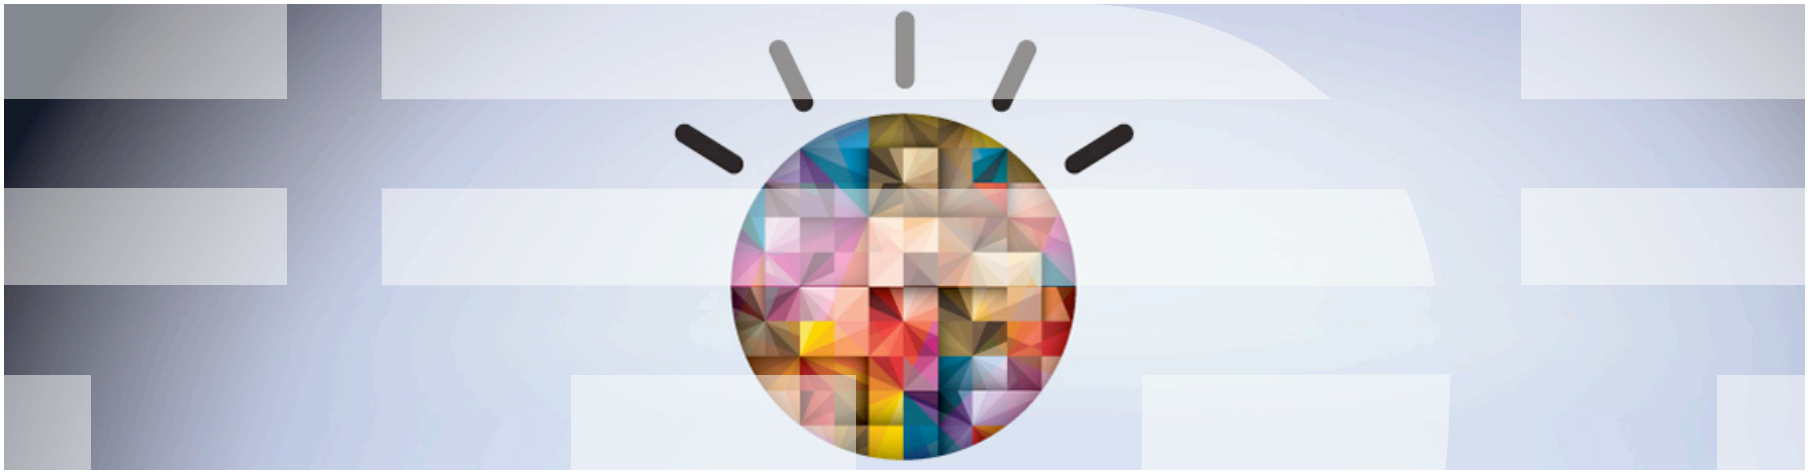

IBM Disclaimer

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

More info at: http://www.ibm.com/bigdata

Big Data → Big/Deep Insights

44x

as much Data and Content Over Coming Decade

Kilobyte (kB)	1,000 Bytes
Megabyte (MB)	1,000 Kilobytes
Gigabyte (GB)	1,000 Megabytes
Terabyte (TB)	1,000 Gigabytes
Petabyte (PB)	1,000 Terabytes
Exabyte (EB)	1,000 Petabytes
Zettabyte (ZB)	1,000 Exabytes

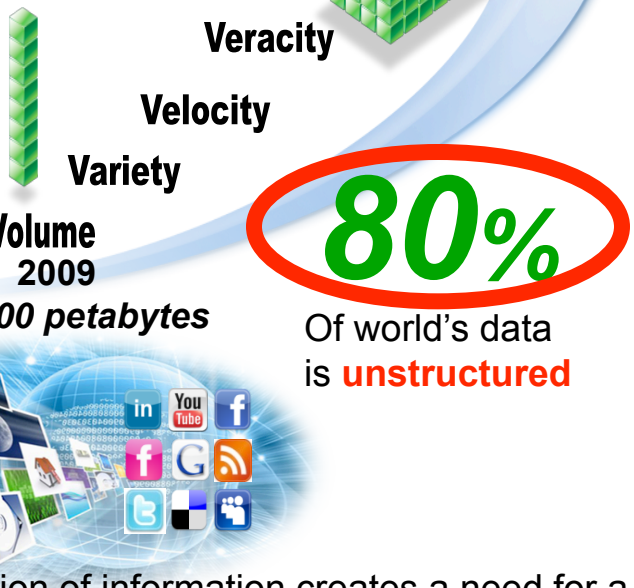

80%

Of world's data is **unstructured**

1 in 3 Business leaders frequently make decisions based on information they don't trust, or don't have

1 in 2 Business leaders say they don't have access to the information they need to do their jobs

83% of CIOs cited "Business intelligence and analytics" as part of their visionary plans to enhance competitiveness

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

The resulting explosion of information creates a need for a new kind of intelligence ….

… Build both integrated and ecosystem solutions, contribute to and leverage open source with own differentiators, open to business and research partners

Big Data Presents Big Opportunities

Extract insight from a high volume, variety, velocity and veracity of data in a timely and cost-effective manner

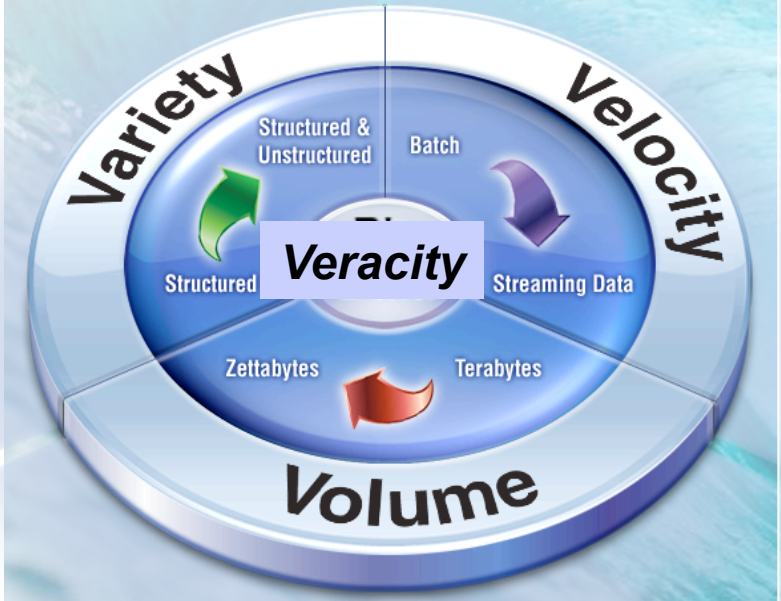

Variety: Manage and benefit from diverse data types and data structures

Velocity: Analyze streaming data and large volumes of persistent data

Volume: Scale from terabytes to zettabytes

Veracity: Establish confidence in data, information and solutions

Categories of Analytics

Learning System, e.g. Watson

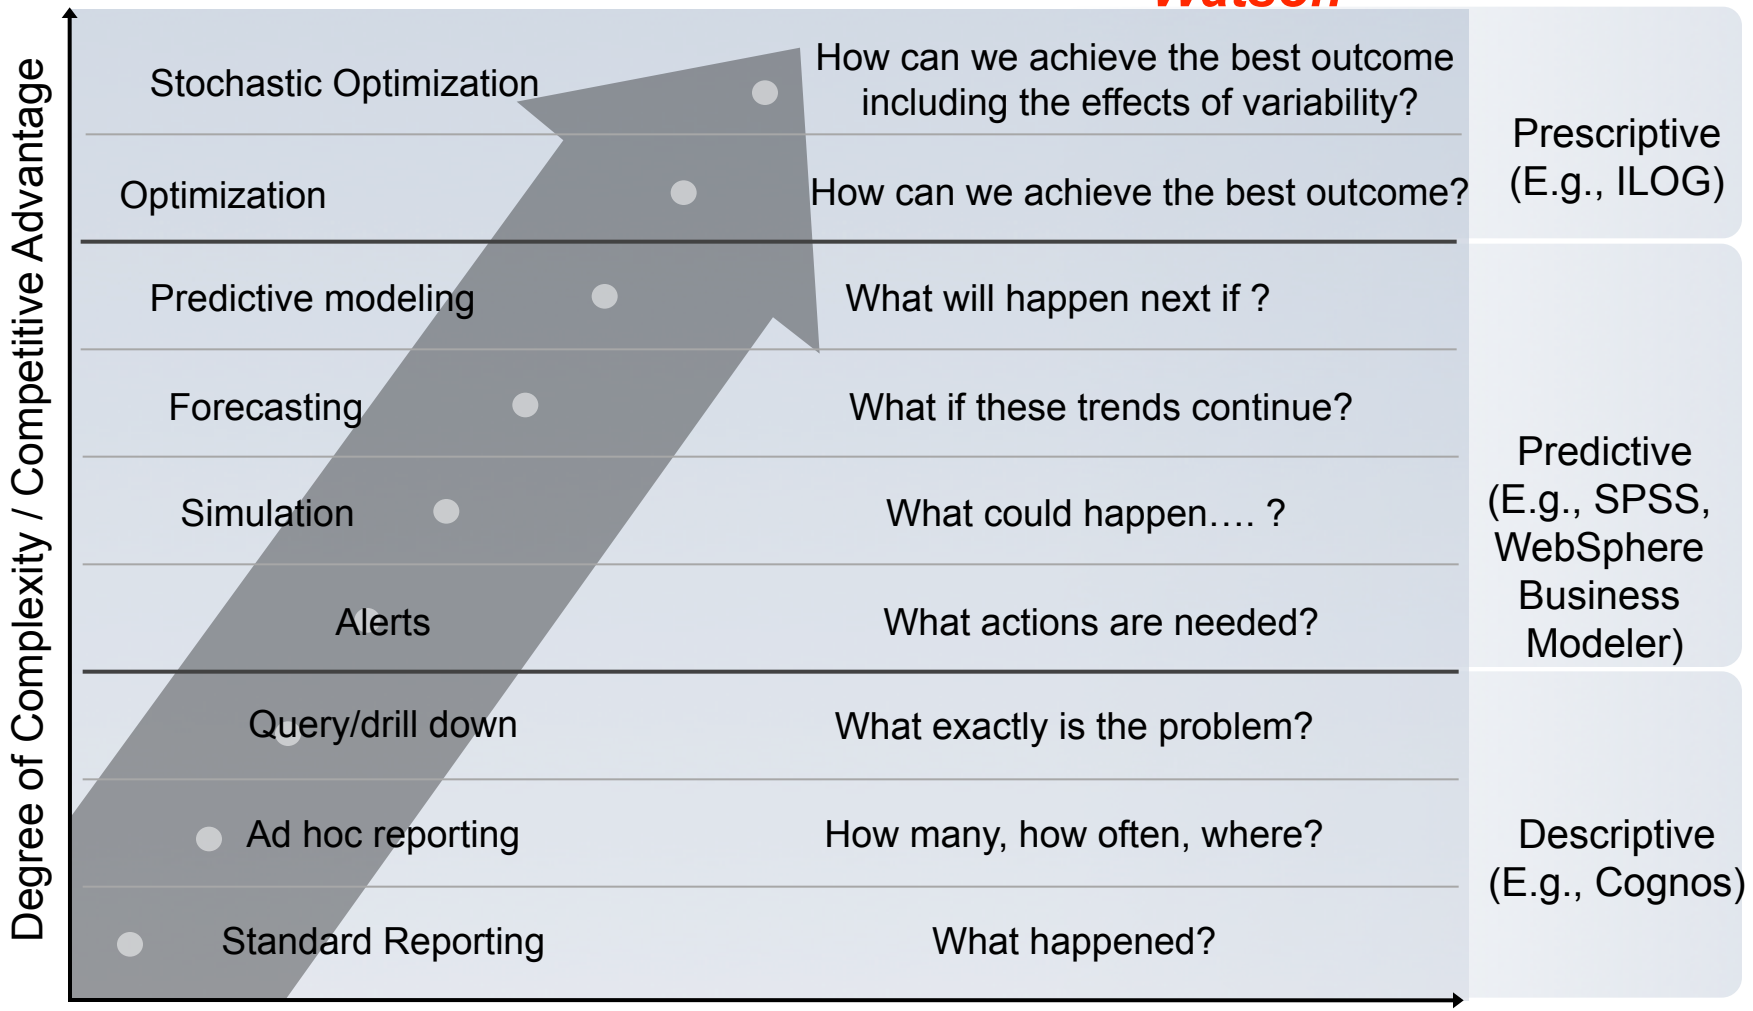

IBM Watson

IBM Watson is a breakthrough in analytic innovation, but it is only successful because of the quality of the information from which it is working.

Big Data and Watson

Big Data technology is used to build Watson's knowledge base

Watson used the ***Apache Hadoop*** open framework to distribute the workload for loading information into memory.

Approx. 200M pages of text
(To compete on *Jeopardy*!)

Watson's Memory
(15TB)

10 racks of P750s,
2870 processor cores

Watson technology offers great potential for advanced business analytics

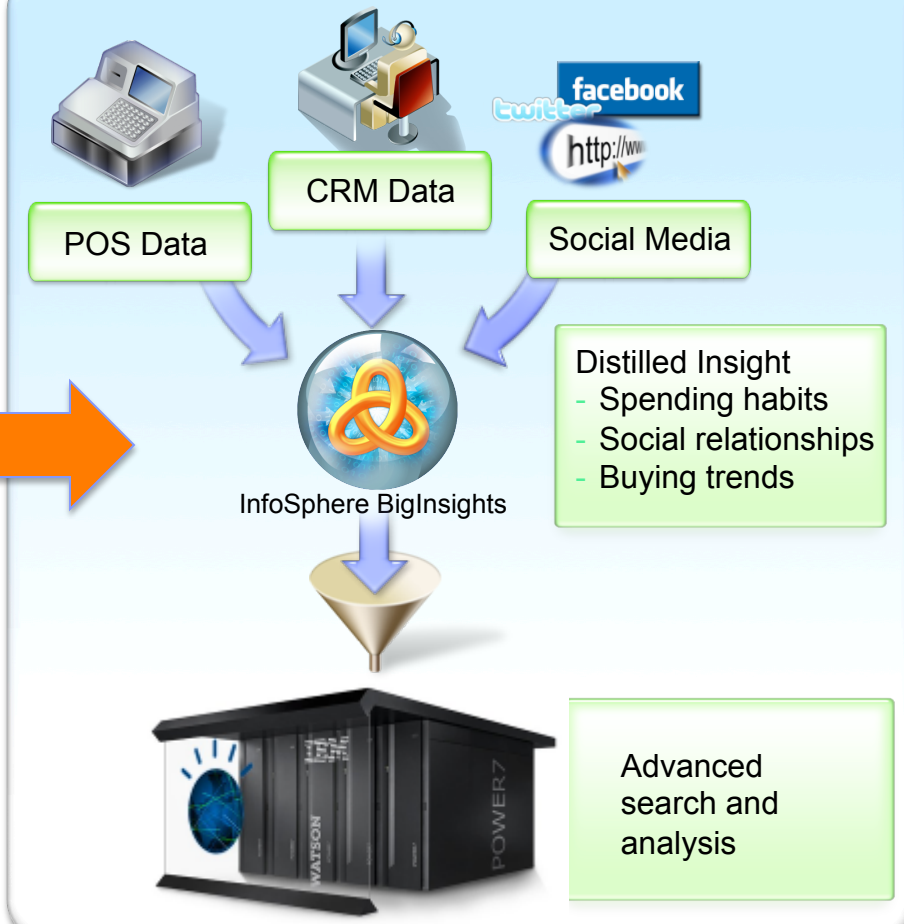

Watson Showcased Advantages of Power Linux for Big Data

Search vast amounts of unstructured information in fractions of a second

Run thousands of tasks in parallel

- 8 higher frequency cores per socket
- 4 threads per core
- Larger eDRAM on-chip cache
- Single socket to 32 sockets per system

2x the bandwidth of other commercially available systems at 500 GB per chip

Massive scale-out flexibility

- Choice of dense rack or blade nodes
- High speed, low latency interconnect
- New, highly affordable pricing options comparable to x86 rack or blade nodes

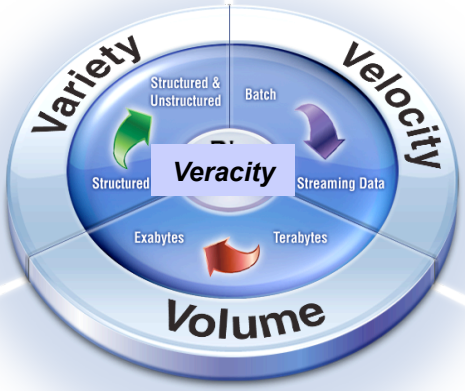

IBM Big Data Platform **Vision**

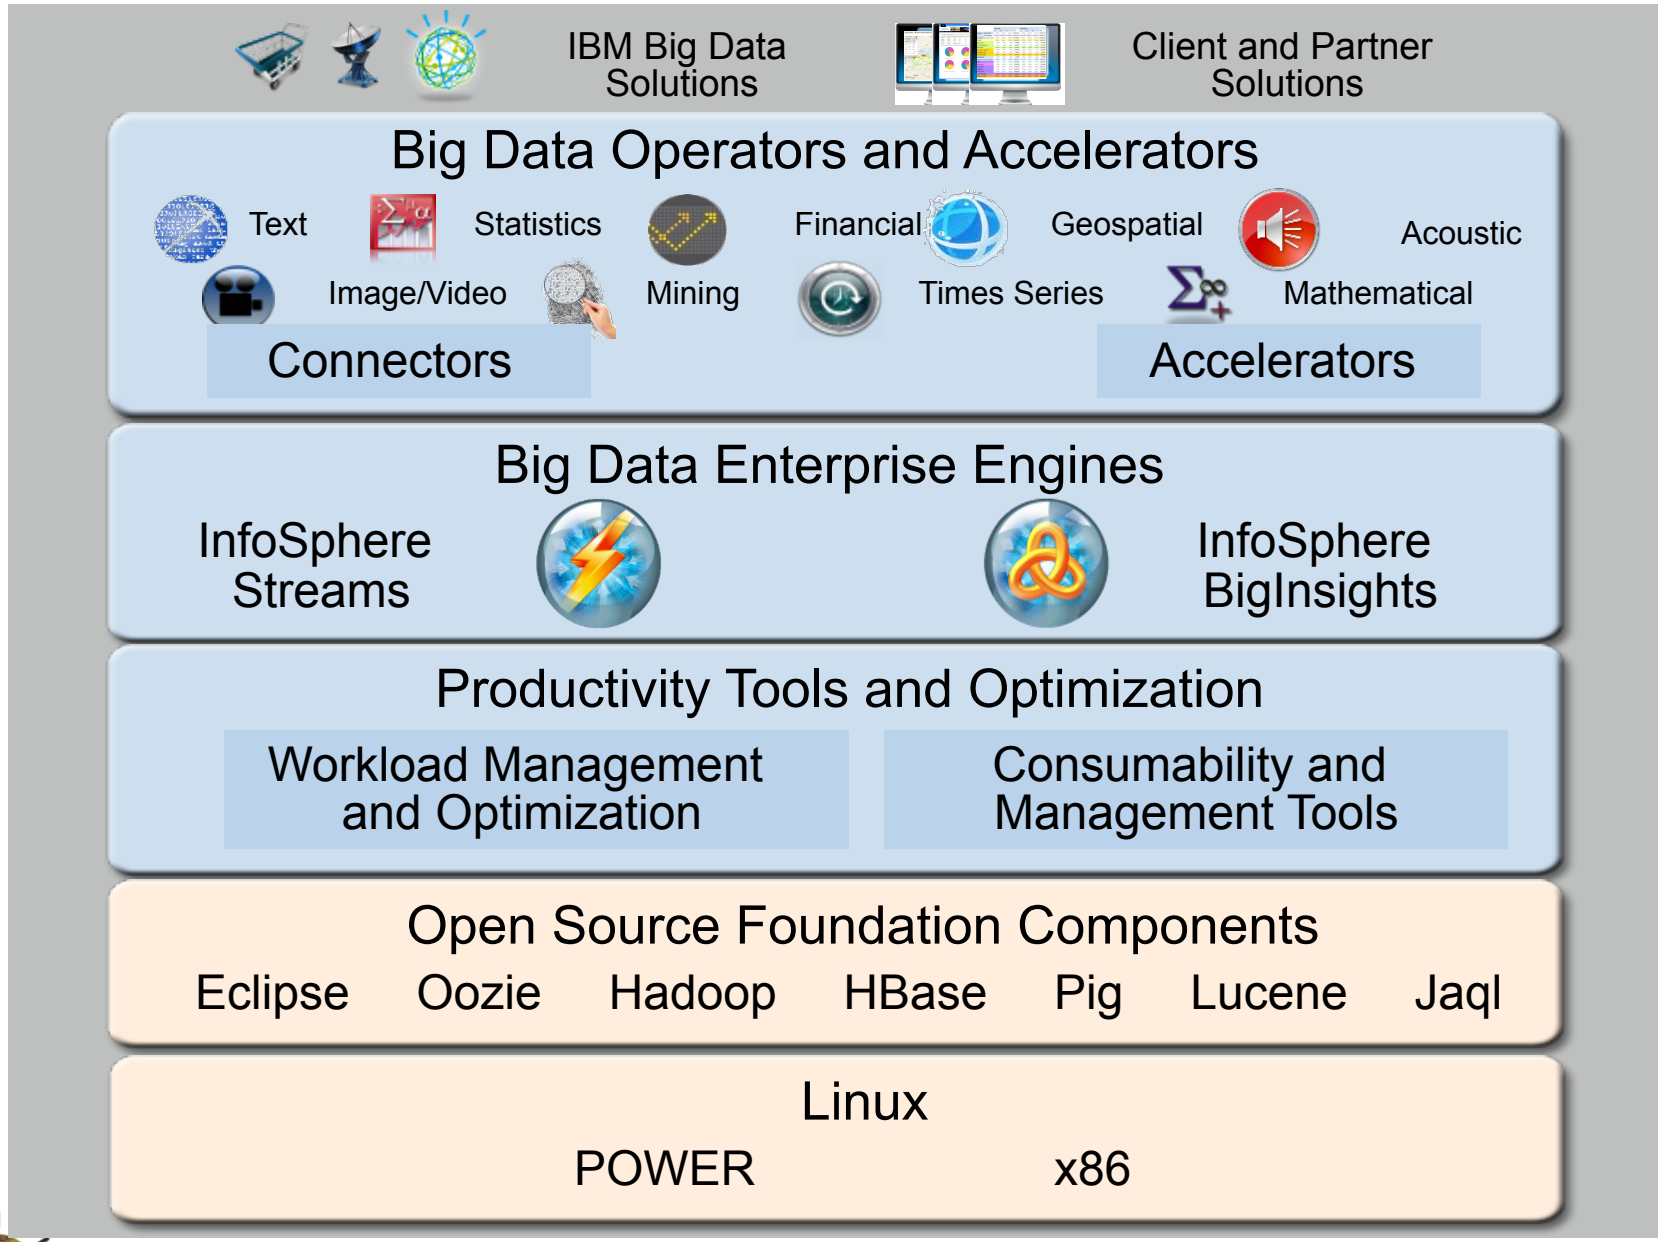

InfoSphere Streams v2.0

A Platform for Real Time Analytics on BIG Data

Volume Terabytes per second
Petabytes per day

Variety All kinds of data
All kinds of analytics

Velocity Insights in microseconds

Agility Dynamically responsive
Rapid application development

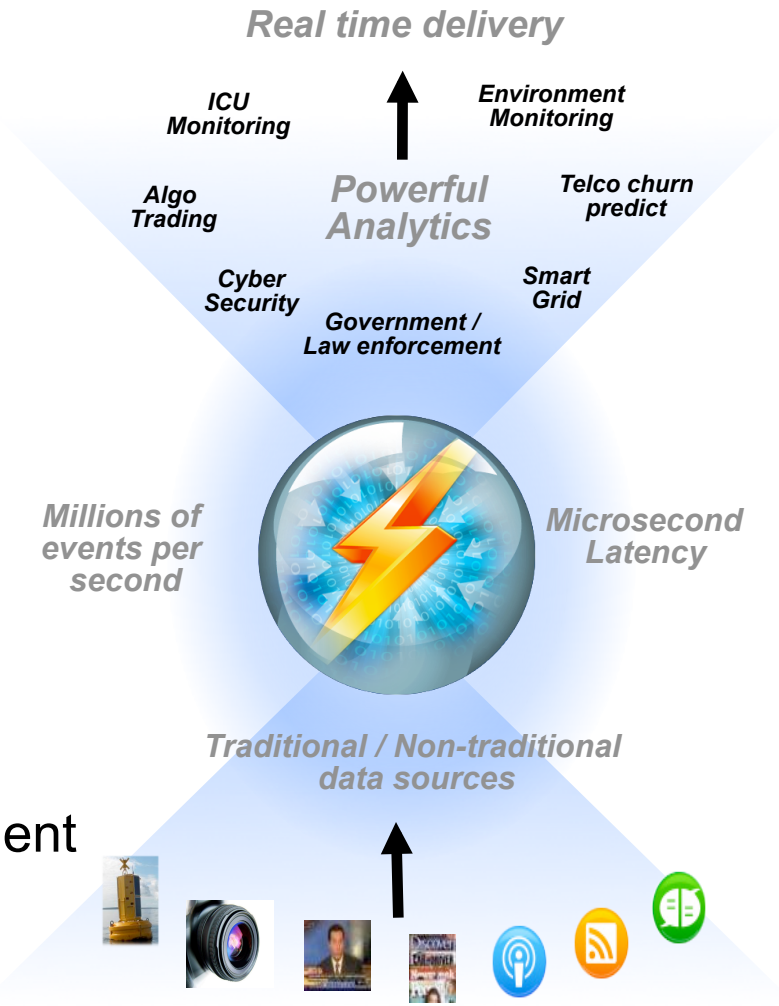

InfoSphere Streams v2.0

Development Environment

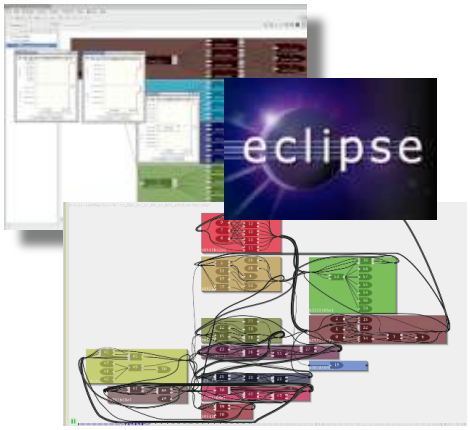

- Eclipse IDE
- Streams LiveGraph
- Streams Debugger

Runtime Environment

- Linux
- Multicore hardware
- InfiniBand and Ethernet support
- Clustered runtime for near-limitless capacity

Toolkits, Adapters & Samples

- Standard Toolkit
- Internet Toolkit
- Database Toolkit
- Mining Toolkit
- Financial Toolkit
- User defined toolkits
- Over 50 samples

Streams and BigInsights - Integrated Analytics on Data in Motion & Data at Rest

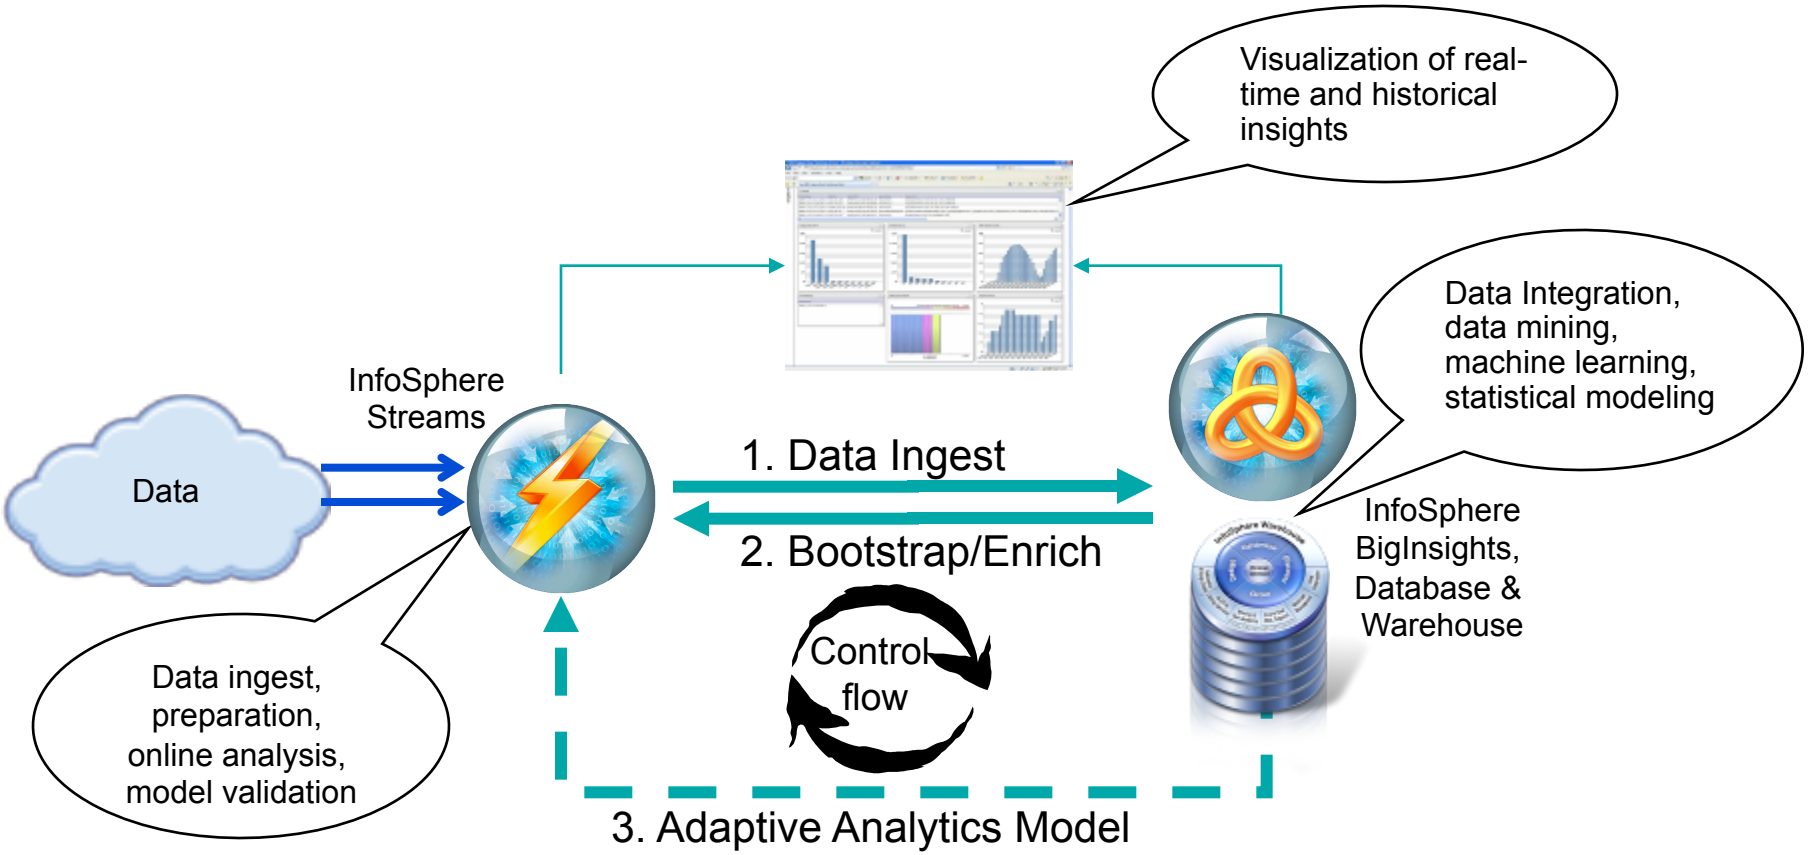

IBM: Building with the Open Source Community

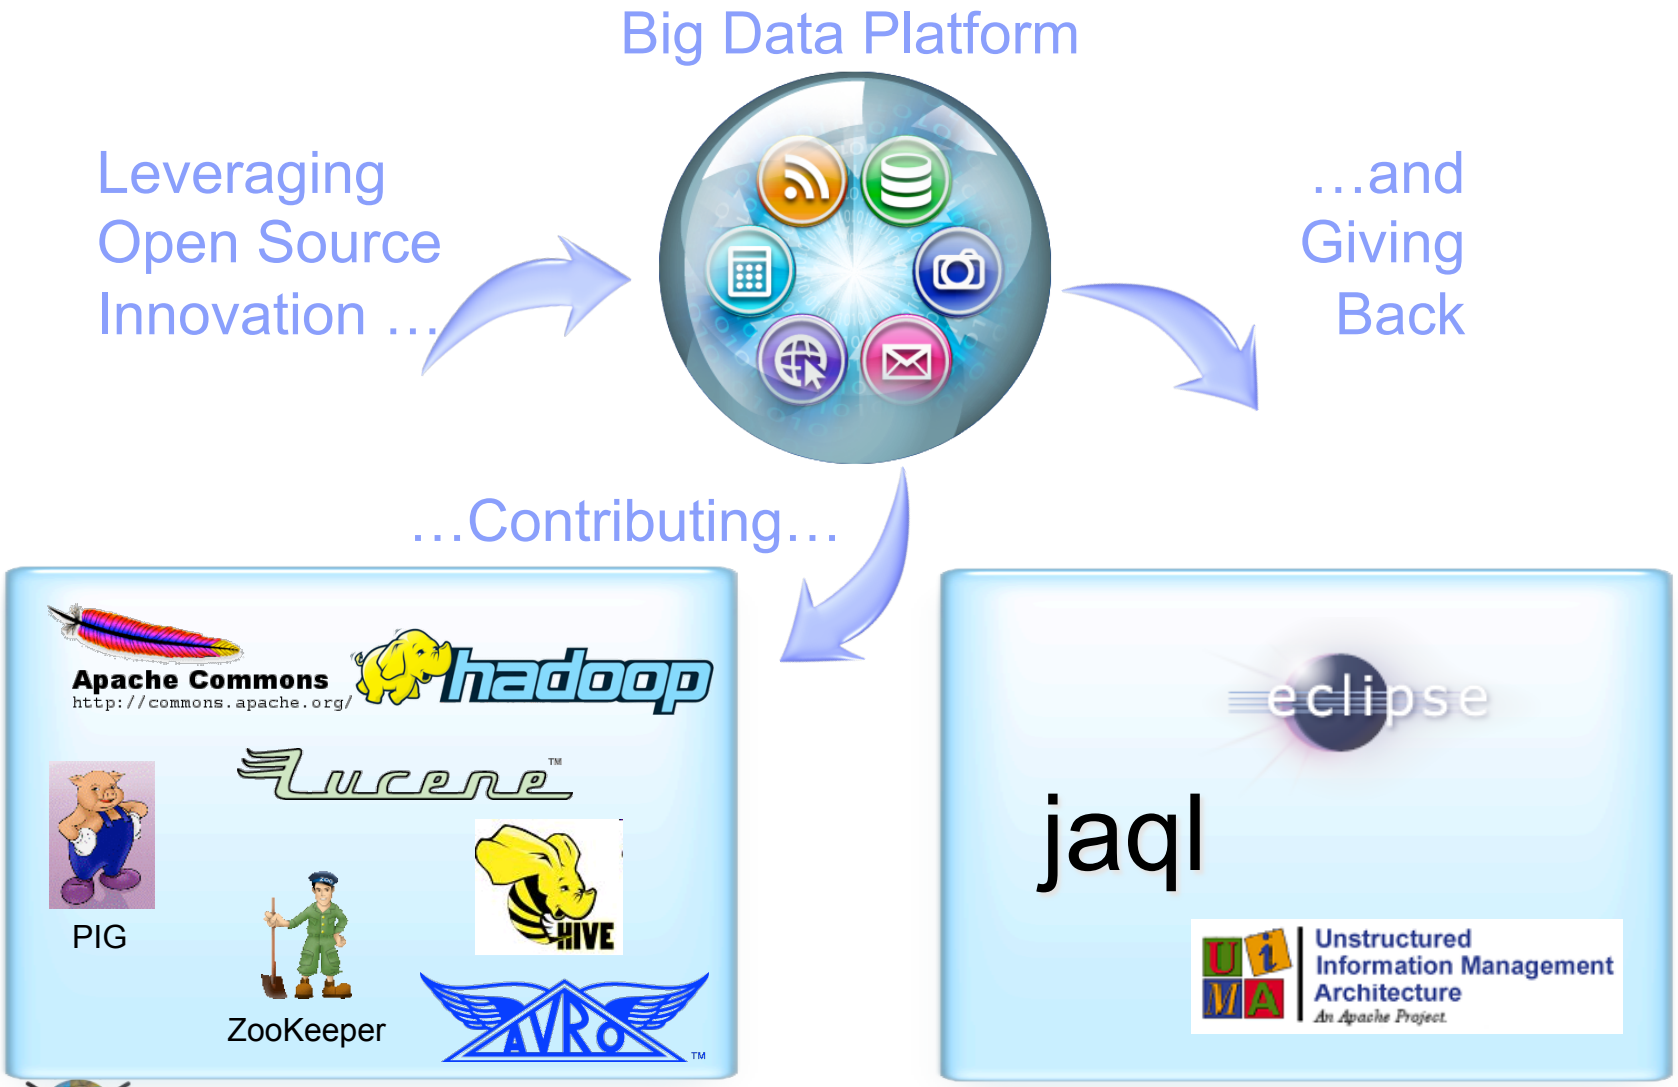

Value Beyond Open Source

- Technical differentiators
  - Built-in analytics
    - Text processing engine, annotators, Eclipse tooling
    - Interface to project R (statistical platform)
  - Enterprise software integration (DBMS, warehouse)
  - Simplified programming / query interface (Jaql)
  - Integrated installation of supported open source and IBM components
  - Web-based management console
  - Platform enrichment: additional security, job scheduling options, performance features, . . .
  - Standard IBM licensing agreement and world-class support
  - More to come in future releases!

- Business benefits
  - Quicker time-to-value due to IBM technology and support
  - Reduced operational risk
  - Enhanced business knowledge with flexible analytical platform
  - Leverages and complements existing software assets

Performance Enhancement Examples

- Adaptive MapReduce
  - Speeds up a class of jobs
    - Example: Jobs that process many small files
    - No changes required to jobs (applications)
  - Accomplished by changing how certain MapReduce tasks executed
    - Map tasks can make runtime decisions based on environment and status of other tasks.
    - Requires communication through ZooKeeper.
  - Enabled through Jaql option, MapReduce job property setting

- Efficient processing of compressed text data
  - Use multiple Map tasks
    - Hadoop default: 1 Map task assigned to process compressed text files
  - Enabled through BigInsights LZO-based compression technology
    - Performance, compression ratios generally consistent with other LZO-based technologies
  - Automatic with Jaql; programming option with Java MapReduce

Technology Example: GPFS-SNC for better Performance, Availability, Integrity and Manageability

- Query languages like Pig and JAQL need good random I/O performance
- Sort requires better sequential throughput
  - GPFS is twice HDFS for both of the above
- For document index lookups, client side caching is a big win
  - 17x throughput speedup

GPFS-SNC Key technology
- Locality awareness
- Write Affinity
- Metablocks
- Pipelined replication
- Distributed recovery

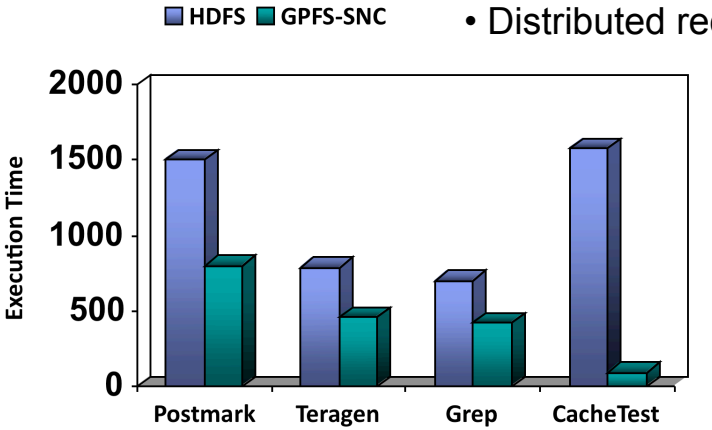

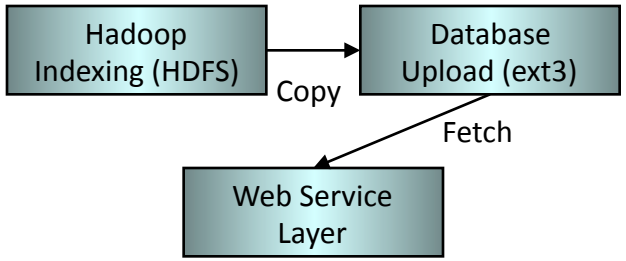

HDFS:
Extra copy overhead and network fetch, separate clusters for analytics and database

Workload Isolation

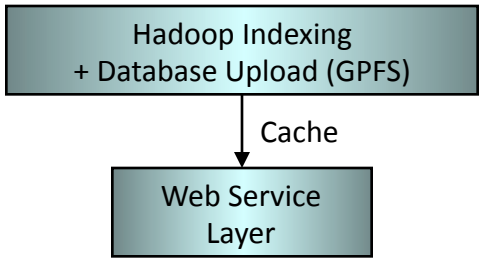

GPFS:
Single cluster for analytics and database, no copying required, caching for web layer

- Proven data integrity
- Replicated metadata services
  - Yahoo keeps 3 copies of 3 versions of HDFS because of unknown data integrity [1]
  - Quantcast deletes files once HDFS is 50% full [2]

[1] Care and Feeding of Hadoop Clusters, Marc Nicosia, Usenix 2009
[2] The Komos Distributed File System, Sriram Rao, Quantcast Inc.

Technology Example: OpenCL

- **OpenCL (Open Computing Language)** is *an open standard for cross-platform, parallel programming of modern processors found in personal computers, servers and handheld/ embedded* devices.
- Highly flexible: supports computation on CPUs, GPUs, accelerators (SIMD, FPGAs, DSPs)
  - **Research contributions**
  - 2.8 X acceleration factor for the sparse coding phase (considering the best timings for each implementation).
  - 2.0 X overall algorithm improvement factor, including the preprocessing costs.
- IBM has released "technology preview"
  - http://www.alphaworks.ibm.com/tech/opencl

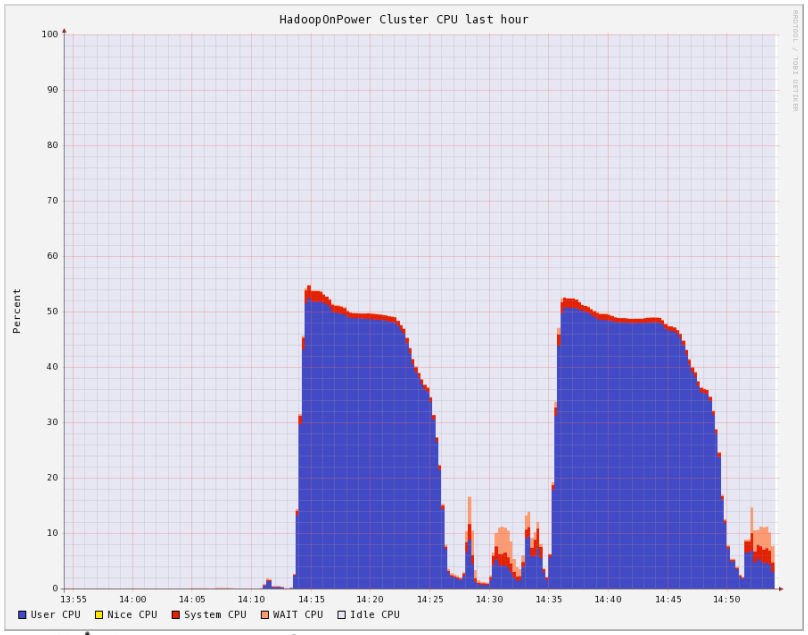

Java Cluster CPU Usage

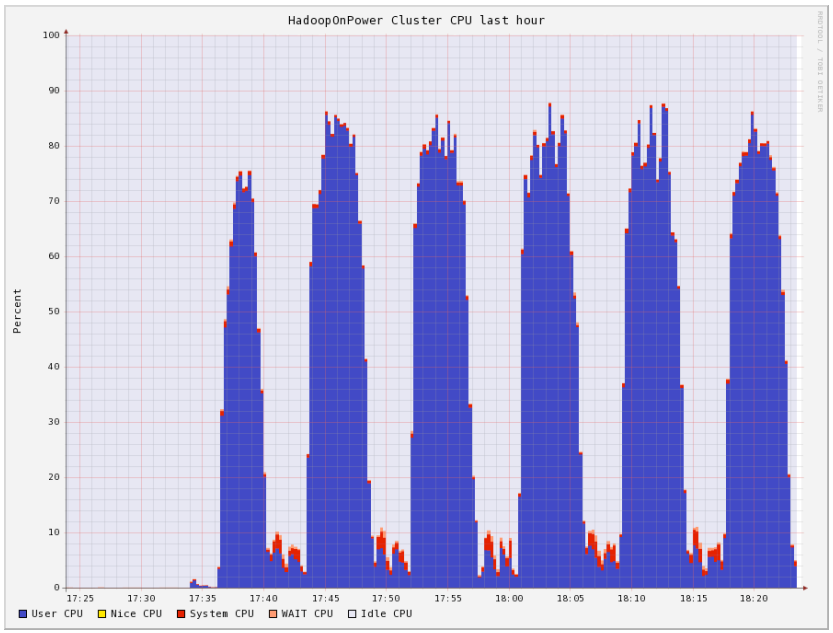

JOCL Cluster CPU Usage

Technology Example: Reconfigurable FPGA Acceleration

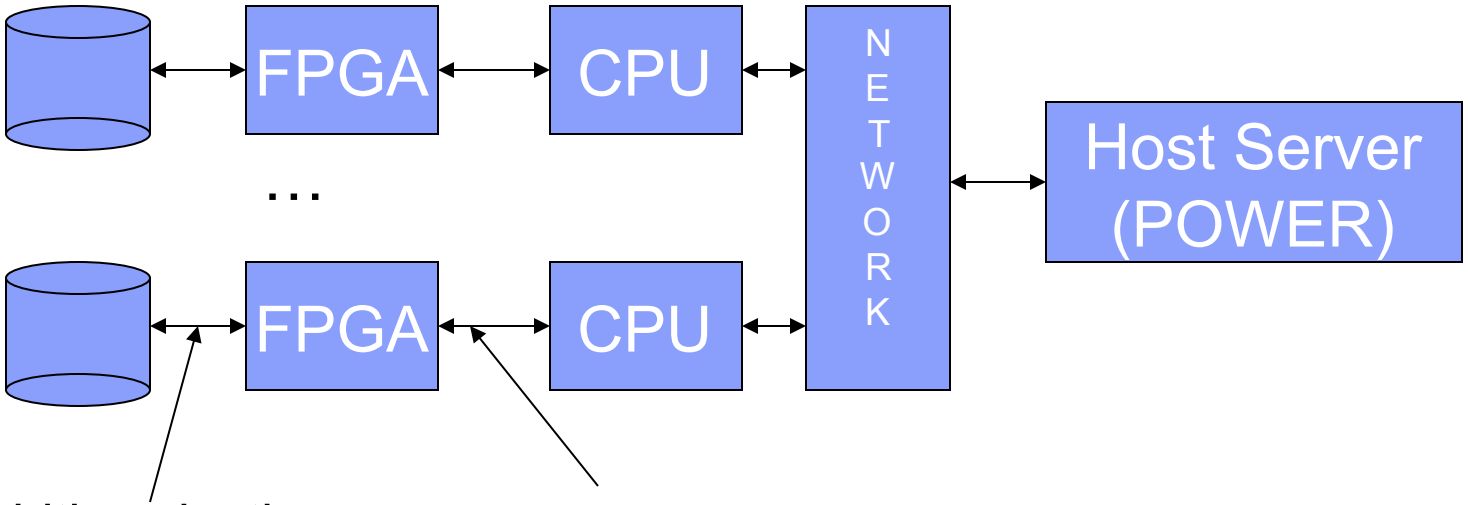

Bandwidth reduction
(& capacity increase)
Through (De)Compression

World's fastest gzip
(Research Contributions)

Bandwidth reduction
Through Filtering

Big Data: Dictionary & Regexp
based filtering

Net result: Significant increase in capacity and throughput in place

Example System Research Issues

- Scalable system and network designs for capturing large numbers of concurrent data streams or high bandwidth data streaming
- Data management for vast amounts of unstructured data
- OS, distributed systems and system management support for very large-scale analytics
- Debugging and performance analysis tools for analytics and data-intensive computing
- Programming systems and language support for deep analytics
- Mapreduce and other processing paradigms for analytics
- Processor, memory and system architectures for data analytics
- Benchmarks, metrics and workload characterization for analytics and data-intensive computing
- Accelerators for analytics and data-intensive computing
- Implications of data analytics to cloud computing
- Implications of data analytics to mobile and embedded systems
- Energy-efficiency and energy-efficient designs for analytics
- Availability, fault tolerance and data recovery in large-scale data-oriented environments

Conclusions


- Significant IBM investment in “Big Data” solutions

- IBM BigInsights+Streams: strategic platform for Big Data analytics
  - Leverage and extend open source
  - Enable firms to exploit growing variety, velocity, volume and veracity of data
  - Deliver diverse range of analytics: descriptive, predictive, prescriptive, learning
  - Complement existing software investments and commercial offerings
  - Provide enterprise-class infrastructure and supporting services

- IBM advantage
  - Combination of software, hardware, services and advanced research
  - Both integrated and ecosystem solutions

- Open to research collaboration and business partnership

Big Data System and Architecture

Jian Li
IBM Research in Austin
Email: jianli@us.ibm.com