BROCADE

Fabric Technology – Backbone of Big Data Solutions

Deepak Peswani

Sr. Product Manager, Data Center APAC Team

Legal Disclaimer

All or some of the products detailed in this presentation may still be under development and certain specifications, including but not limited to, release dates, prices, and product features, may change. The products may not function as intended and a production version of the products may never be released. Even if a production version is released, it may be materially different from the pre-release version discussed in this presentation.

Nothing in this presentation shall be deemed to create a warranty of any kind, either express or implied, statutory or otherwise, including but not limited to, any implied warranties of merchantability, fitness for a particular purpose, or noninfringement of third-party rights with respect to any products and services referenced herein.

Workshop For Big Data Benchmarking

Brocade

- Brocade overview
- How we understand Big Data
- Brocade's approach to Big Data
- Hopes for the future
- Considerations to the committee
- Benchmark suggestions

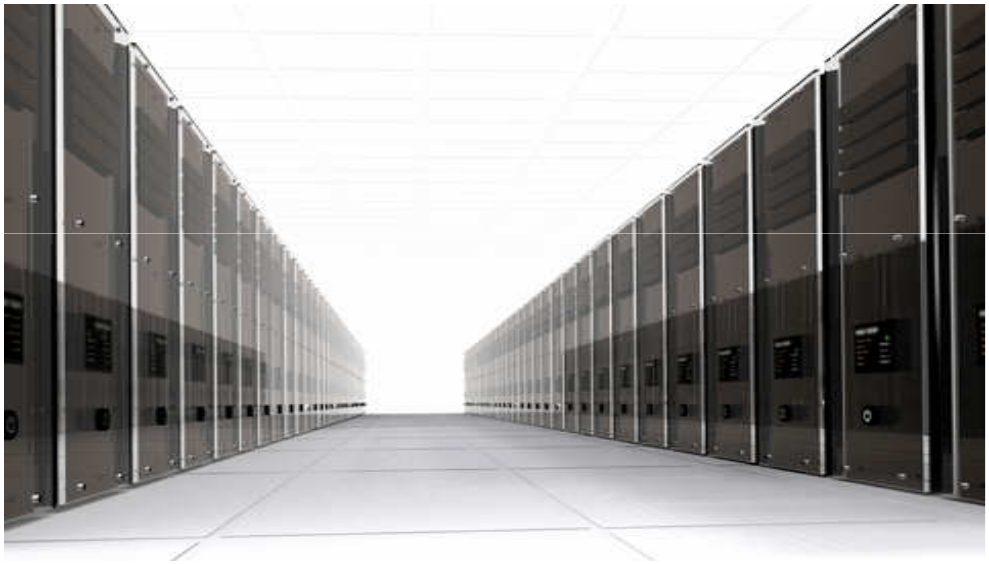

Brocade at a Glance

- Founded in 1995
- 4,500+ employees worldwide
- Headquartered in San Jose, CA
- Operating in more than 160 countries
- $2+ billion in annual revenue

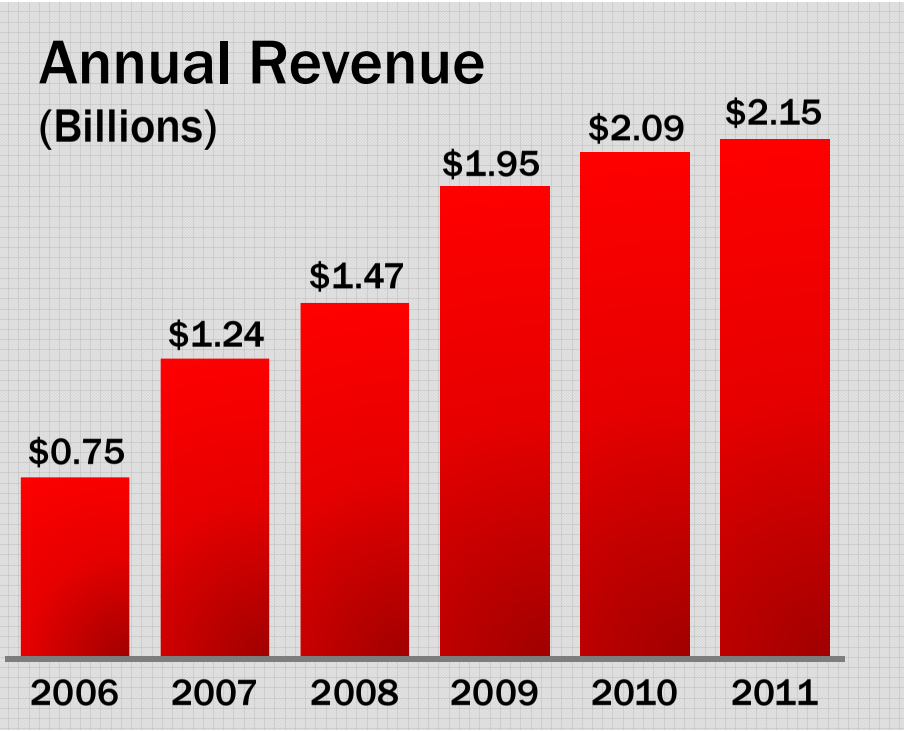

A Difference in Data

New Sources of Data

Customer

- Behavior
- Sentiment
- Social Networks

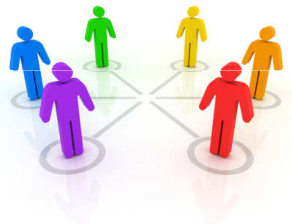

Newspaper
Magazine

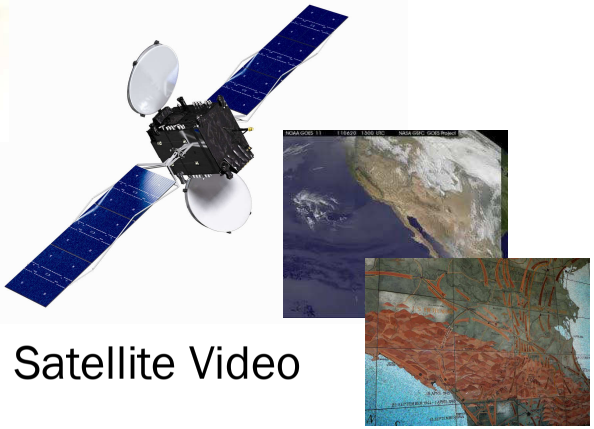

Satellite Video

"Unstructured Data"

These are not databases nor spreadsheets

Social Media

- Blogs
- Posts
- RSS Feeds
- Videos

Defining Big Data

Data is useless if it doesn't add business value

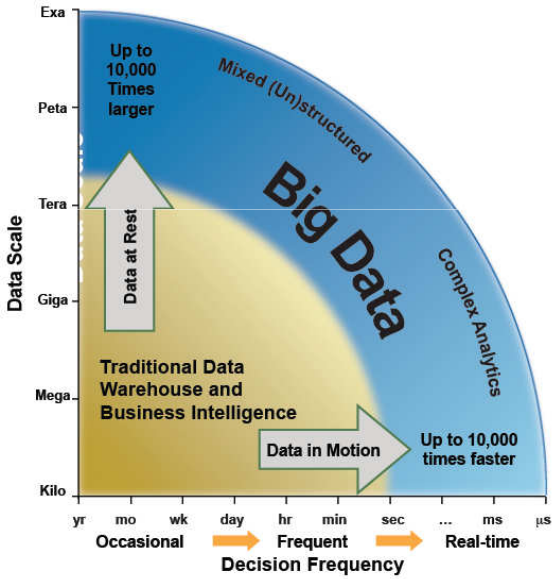

Defined by: Volume, Velocity, Variety

Data Analytics

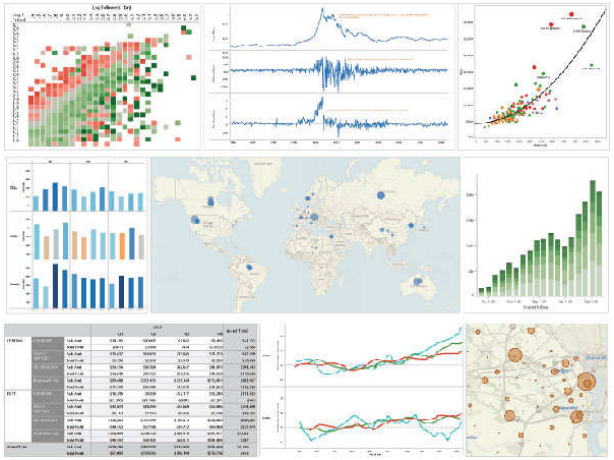

The Process of: "Discovery and communication of meaningful patterns in data"

Big Data Analytics

Turning disparate data into business insights

Modeling
Data Mining

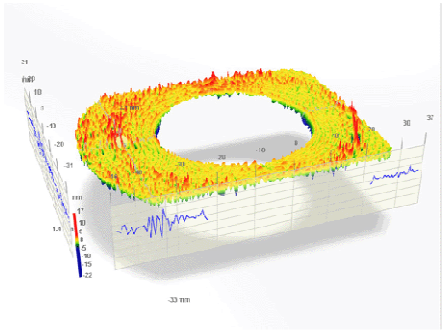

Statistical regression
Decision-Tree
Classless learning

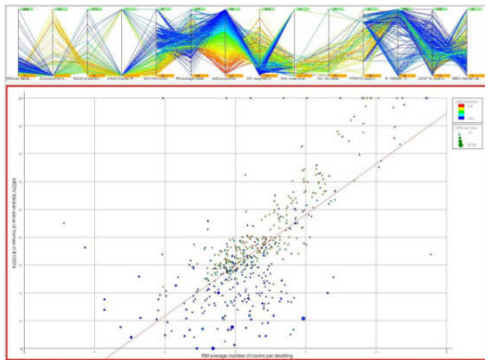

Heat displays
Sentiment Analysis
Predictive Analytics

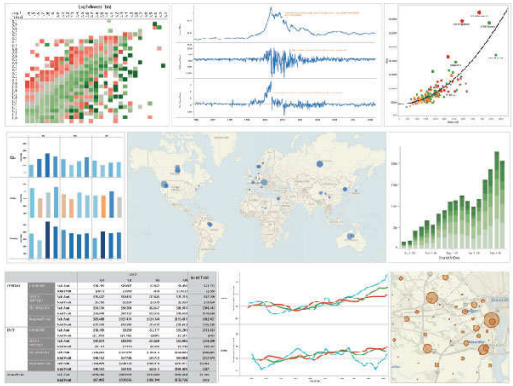

Orchestration & Performance Tuning

Data Visualization & Analytical Engine

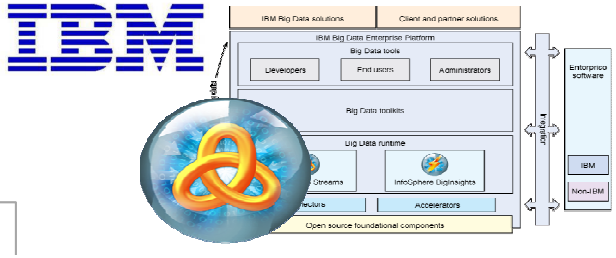

Distributed File System

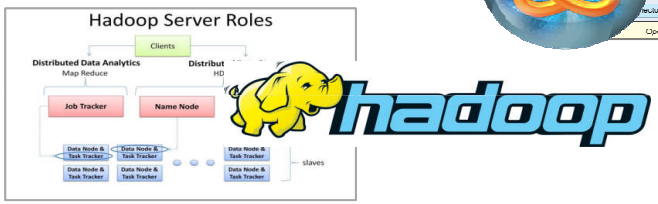

Server Hardware

Network Fabric

Brocade Big Data Solution

- Tested
- Certified
- Configurable
- Customizable
- Supportable
- Scalable
- Built for delivery

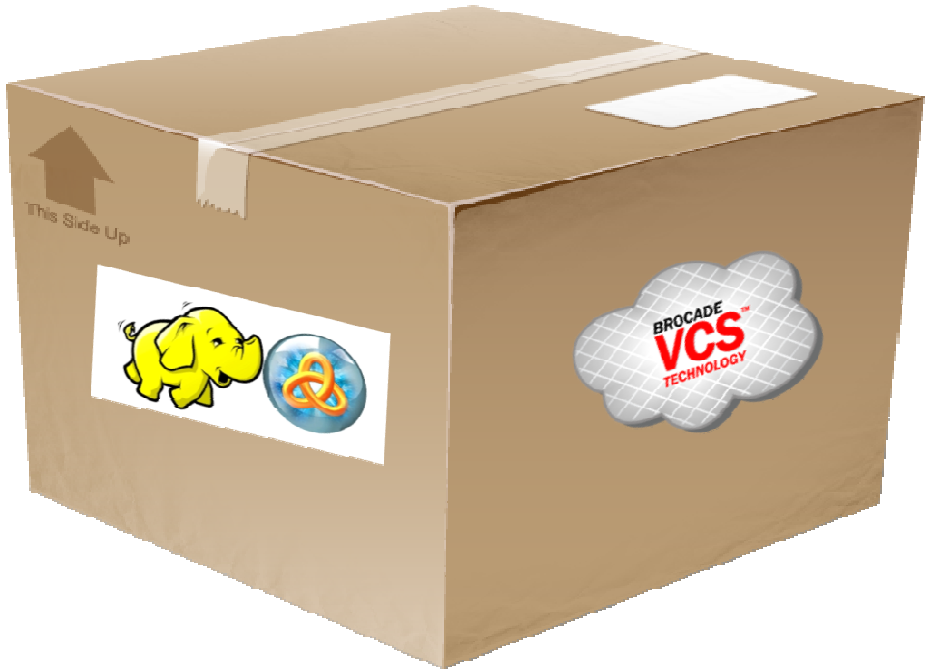

Consideration #1: Multi-Speed support

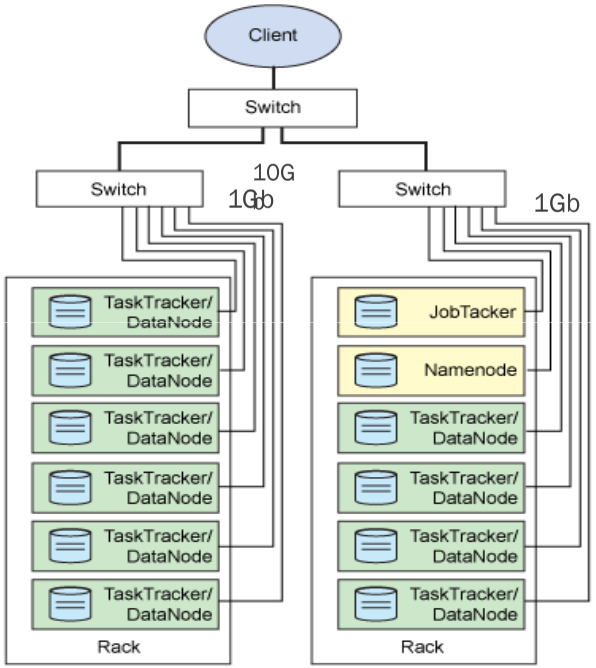

Bandwidth options

- 1/10GbE Nodes
- 10/40GbE Spine
- 40GbE to Intermediate Nodes
- 40GbE to NameNode/Jobtracker
- 40GbE to Enterprise
- 100GbE to Source Data
- 100GbE to Warehouse/Mart
- 100GbE Mirrored Storage

Consideration #2: The fabric needs to be as reliable as HDFS or better.

Hadoop is “Rack Aware”

- Copies inside the rack
- Copies outside the rack
- 10-Rack solution, a full rack lost is only 10% hit to performance

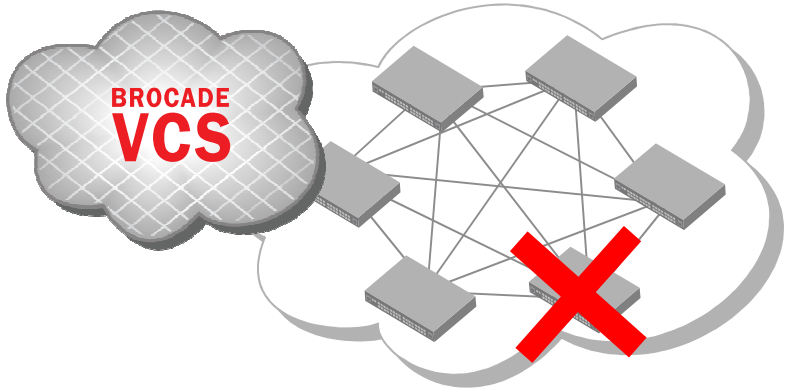

Fabrics need to be “Rack Aware”

- Any loss of a device is instantly rerouted
- Resources are instantly available to the file system
- Distributed management

Consideration #3: Automatic additions and convergance

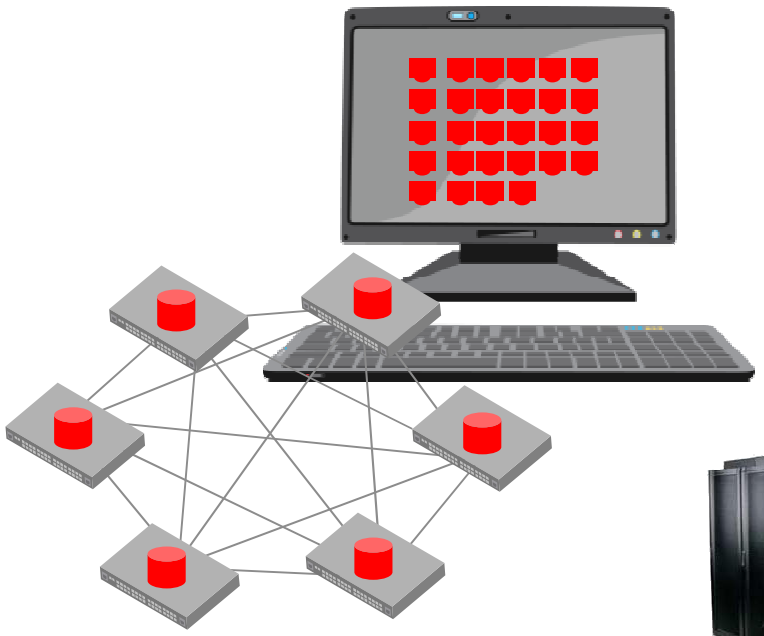

The fabric should grow with requirements

- No manual configuring needed to scale up
- Neighbors configure new boxes
- New resources and expanded footprint automagically appears on screen
- New resources are ready to be productive

Consideration #4: Reliability

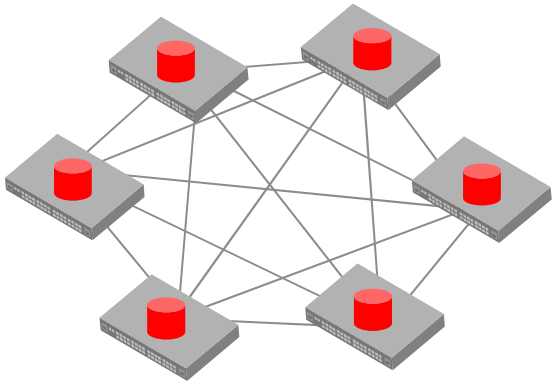

Big Data is a warehouse and a data mart in one. Data needs to be guaranteed delivery.

- Based on the DCB protocols
- Fabric should be rated for storage
- Lossless data transfer
- TRILL L2 multi-path technology
- The fabric should ensure that data has the best path, plenty of bandwidth, and complete delivery throughout the fabric without the file system getting involved.

Consideration #5: Metadata

Complete and comprehensive, real-time information about the health of your solution.

Examples:

- Rebalancing file system
- Hot-Spot detection
- Total System status
- Traffic load analysis
- Benchmarking
- Capacity planning
- Outage reporting
- Integration with orchestration software
- Real-Time traffic alaysis

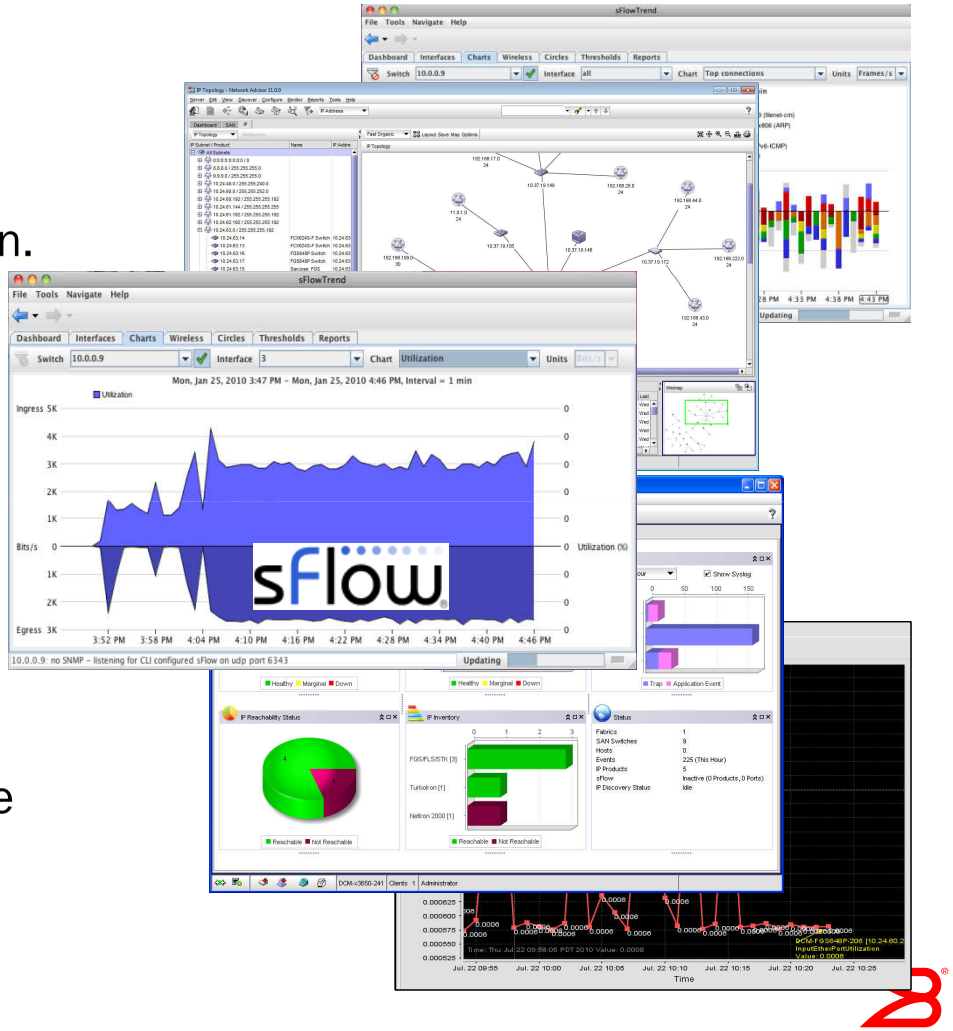

The Big Data 100…

- Be based on business productivity
- Include considerations ease of implementation, overall complexity, manageability, flexibility, upgradability, adherence to standards, etc…






…is what business around the world will use to build the systems to run their enterprise. It is our responsibility to ensure it reflects reality.

BROCADE

Thank You