



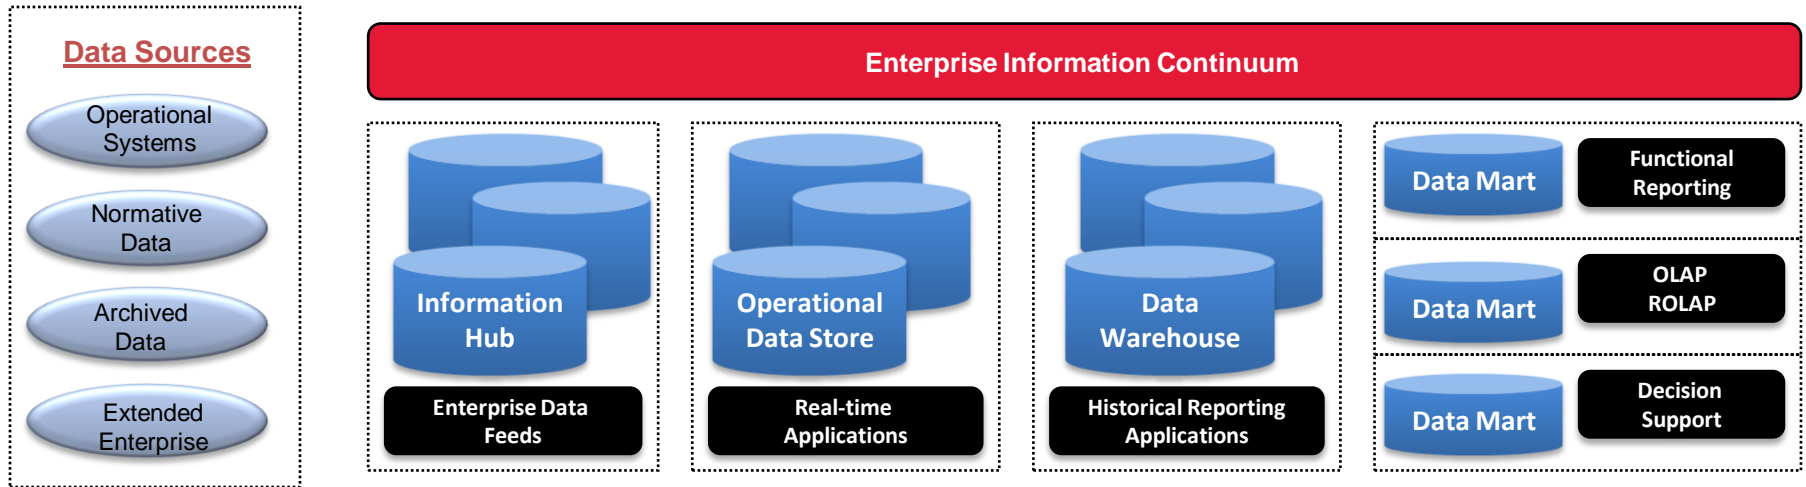
XtremeData Inc.

dbX Analytics System Introduction

Company

- Founded in 2004
- Locations: Schaumburg, IL (HQ / Engineering/ Sales); Bangalore, India (Engineering); St. Louis, MO (Engineering); South Bend, IN (Sales)
- Current business and products
 - Patented “In-Socket Accelerator” (ISA) that enables companies to create higher performance, lower power appliances. Branded as XD2000F and XD2000I in the High-Performance Computing market.
- Management:
 - Ravi Chandran; CEO and Founder
 - Faisal Shah; Chief Technologist, Database Appliance
 - Former Cofounder Knightsbridge Solutions LLC
 - Jay Desai; Sr. VP of Strategy
 - Former Cofounder Knightsbridge Solutions LLC
 - Susan Clarke; VP of Finance
 - Geno Valente; VP of Sales and Marketing

Analytic queries - an unmet need



Major Trends Driving BI

Becoming more real time & pervasive & even Mission Critical

Enterprise wide in scope with increasing data volume and data access needs

Customer/partner focused - business process differentiation; allowing outside users access to previously internal data

Consequences

Managing service levels for BI applications is increasingly becoming a challenge

Educated users want more data - want ad hoc access/unconstrained exploration of larger data sets

Pre-Defined queries no longer work; outside users can not or will not tell you how they want to use/access the data. They just want to use it.

Addressing Key Market Problems

- Solutions today “restrict” access to data based on pre-defined query patterns to get performance.
- Analysts and Data researchers are demanding “unrestricted” ad hoc access to find the business value inside their data – reference “Competing on Analytics”
- Customers are looking for cost take-out and low TCO. Major vendors today are Expensive.
- Customers are looking to bring longer time-series, archived data, and low-information density data into play for analytics, at a justifiable ROI.




What's missing?

- Legacy solutions are not designed for analytic (ad hoc) access to large scale data
- Existing large-scale BI solutions are too expensive ($> \$100\text{K/TB}$)
- Cheaper solutions ($< \$50\text{K/TB}$) cannot perform analytics with acceptable execution speed



Key Difference

- Only XtremeData offers an appliance purpose built for ad hoc analysis of large data sets.



XtremeData Approach

- Start with a commodity **Linux Cluster**
 - Add direct-attached distributed storage
 - Add high-speed interconnect network (InfiniBand)
 - Add a re-engineered open-source DB engine (postgresql) that:
 - Supports a shared-nothing, parallel query execution model
 - Supports hardware acceleration via “SQL on a Chip”



XtremeData Approach

(cont.)

- ...and it morphs into an **SQL Super-Computer** delivering:
 - Scalable, shared-nothing MPP architecture
 - Hardware accelerated parallel SQL processing
 - Efficient, zero-copy, data exchange on high bandwidth network
 - Dynamic load balancing at run-time

Introducing  DB^x™

Performance Results

Sampling of performance results from recent Proof-of-Concept engagements:

- Financial institution, household credit database: large tables, N-way Joins, Aggregations.
 - **Result:** On 8xNode system (MSRP \$600,000): 33 min 42 sec. ~**16x** faster than existing solution, with estimated MSRP of \$1,500,000.
- *National Laboratory, Graph traversal to discover linkages: single narrow long table, N-way self-Joins.*
 - **Result:** On 8xNode system (MSRP \$600,000): 1 hour 25 mins. ~ **5x** faster than closest competitor, system configuration and MSRP unknown.
- *National Laboratory, Web page ranking (PageRank): rank computation, iterative traversal with updates.*
 - **Result:** On 8xNode system (MSRP \$600,000); **2x-8x** faster than closest competitor, with estimated MSRP of \$1,200,000.

Differentiation

Utility

- **Only** XtremeData's dbX cost-effectively enables wide, deep, and unrestricted querying of structured data.

Price

- **Only** XtremeData's dbX enables Petabyte-scale analytic solutions that were previously cost-prohibitive - An appliance priced at **\$20K/Usable TB**

Performance

- **Only** XtremeData's dbX delivers **4-30x faster** results than competing products at 1/3 the investment

Green

- **Only** XtremeData's dbX provides power and cooling advantages with an unmatched carbon footprint.


DW/BI Industry trends

The four components of a data analysis system are:


Hardware (3x): Storage, Computing (Servers) , Interconnect Network

Software (1x): Relational DataBase Management System (RDBMS)

Hardware

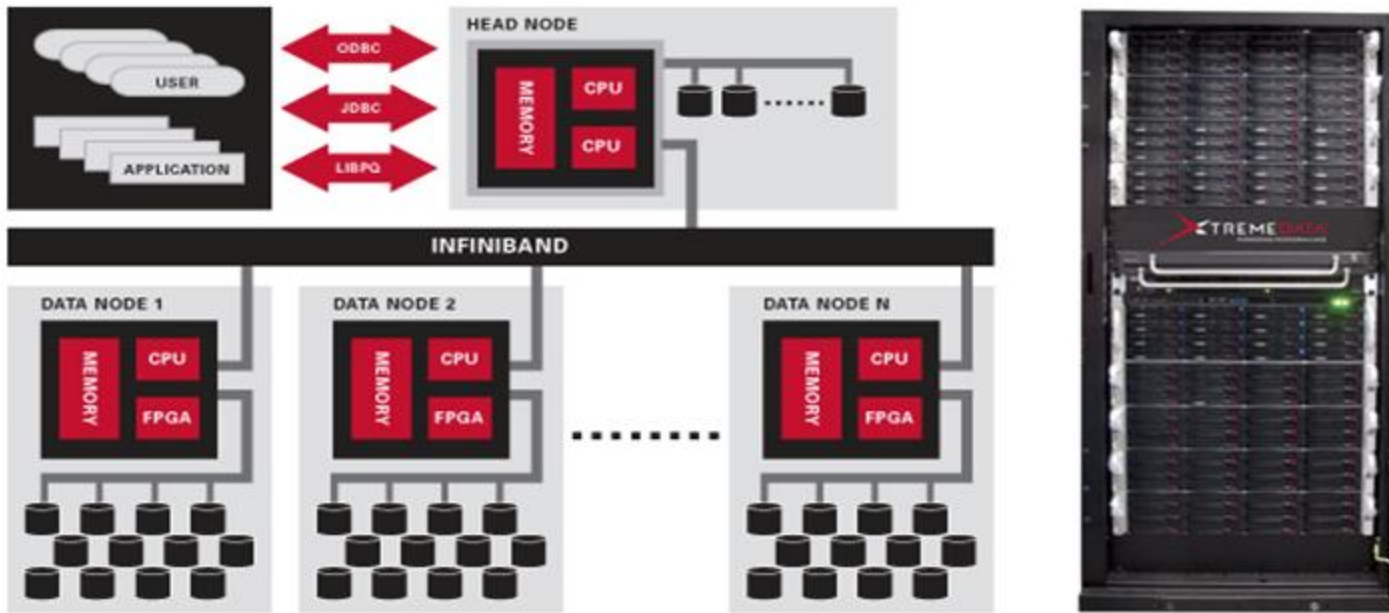


Software



Unique among all vendors in the market today: the dbX architecture leverages the best of all industry trends...We can sustain our price/performance advantage into the future


Architecture



- Engineered from first principles for large database analytics
- Breakthrough price/performance (16xNodes): \$20K/TB at 0.4TB/min query processing
- Provides clear value vs. competitive solutions (Teradata, Netezza, Oracle, GreenPlum)
 - Query performance (SQL processing)
 - Load performance (1-4 TB/hour)
 - Processing density (storage & performance/sq. foot)
 - Carbon footprint (performance/watt)


Under the hood

| dbX System | Data Node Features | SQL on a Chip |
|------------|--------------------------------|--------------------------|
| Head Node | HP DL385 with SAS Storage | Approved HP Accelerator |
| IB Switch | Quad-Core CPU; IB Connection | Contains SQL in Hardware |
| Data Nodes | Patented In-Socket Accelerator | Plugs into CPU Socket |



Patented “SQL on a Chip”

- Replaces CPU in COTS Tier 1 Servers... Qualified by HP (Accelerator Program)
- SQL operations accelerated “under the hood” in silicon
 - Big table data movement: Loads/Unloads/Scans
 - Big operators: Joins, Sorts, GroupBy, OrderBy,...
- 10x the Performance at 1/3rd the power of CPU
- Leverages rapidly advancing cutting edge FPGA technology in patented way



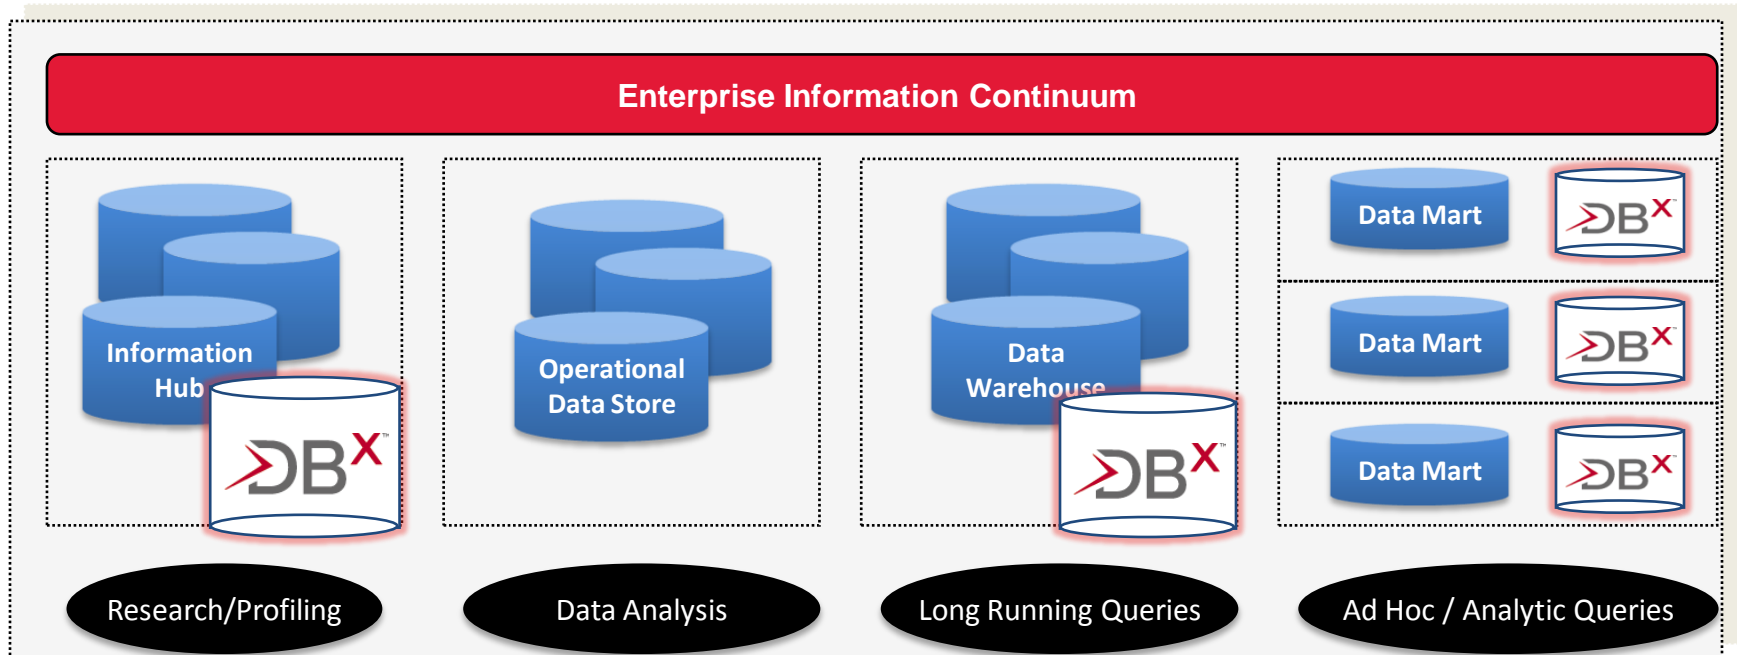
Configurations

| Model | dbX 1008 | dbX 1016 | dbX 1028 | dbX1060 |
|---|----------|----------|----------|---------|
| Rack (Standard 24x40) | 1 (50%) | 1 | 2 | 4 |
| Data Nodes (Opteron + SQL on a Chip) | 8 | 16 | 28 | 60 |
| User Data (1TB drives) | 30TB | 60TB | 105TB | 225TB |



**Scalable to 1024 Nodes
Over 5PB of User Data
(15PB total disk space)**

Summary: Applications for dbX



- Adjunct to existing environments
- Complement legacy investments in DW infrastructure
- Shift analytic query workload away from the current environment - better leverage legacy environment for operational BI applications
- “Sand box” for research, analysis, ad hoc.....
- SQL based analytics

Filling clear market needs for

Usability

XtremeData's dbX cost-effectively enables wide, deep, and unrestricted querying of structured data.

Price

An Appliance priced at \$20TB/Usable TB, only XtremeData's dbX enables Petabyte-scale analytic solutions that were previously cost-prohibitive

Performance

XtremeData's dbX delivers 10x faster results than competing products at 1/3 the investment

Environment (energy-efficiency)

XtremeData's dbX provides power and cooling advantages with an unmatched competitive carbon footprint.

