Bart Sjerps
Advisory Technology Consultant
Oracle SME - EMEA
bart.sjerps@emc.com
+31-6-27058830
Blog: http://bartsjerps.wordpress.com

Greenplum

Enabling Business Intelligence Through Virtual Enterprise Data Warehousing

EMC²
where information lives®

History: From reports to advanced analytics

- Early days: run a simple report against the OLTP Database
- Run heavy batch reports against OLTP Database
  - Dayly, weekly, monthly, year-end, ad-hoc
- Run custom queries against OLTP Database (using standard reporting tools)
  - First use of what later became Business Intelligence, getting (market) knowledge from large amounts of information
- Note: Running Batch and reporting on OLTP kills OLTP response time and performance
- Offload databases for reporting and querying only
  - Implemented as 1:1 copies, or custom designed databases (the first pure Data Warehouses)
- Need for Extract, Transform, Load tools (ETL)
- Evolved into OLAP (Online Analytical Processing); specialized methods for running Analytics
- This required special reporting tools as well

Classic vs. Next-gen business intelligence

Old-style Datawarehousing:

- Frequently run reports/batch
  - Built by programmers, optimized for performance and minimizing resource usage, requires huge developer and DBA efforts
  - This is achieved by classic tuning such as using table indexes, partitioning, SQL optimization
  - Very efficient but only for predictable queries
- Ad-hoc queries against OLTP data
  - Can kill OLTP service levels, therefore this is often offloaded against prod database copy
  - Optimizing using "tricks" such as materialized views
  - Classic tuning fails (because it's unpredictable)
- DWH misused for pieces of business process
  - Now mission-critical!
  - Consider HA / DR / Compliancy

New style Datawarehousing:

- Does not replace classic DWH!
- Get as much data from as many sources as possible
  - Web, data feeds, legacy systems, "smart" electronics, etc etc
- Clean it up and modify it for analytics using ETL tools
  - This is very resource intensive and typically requires long processing times
  - Loading in the DWH can be problematic
    - Classic DB systems again use workarounds for speeding it up
  - Data needs to be as up-to-date as possible (less than 24 hours old)
- Build multi-dimensional databases
  - That can have holes with "missing" data
- Build specialized data-marts
  - Optimized by purpose
  - Contains sub-set of all data

Classic Data Warehouse Architecture

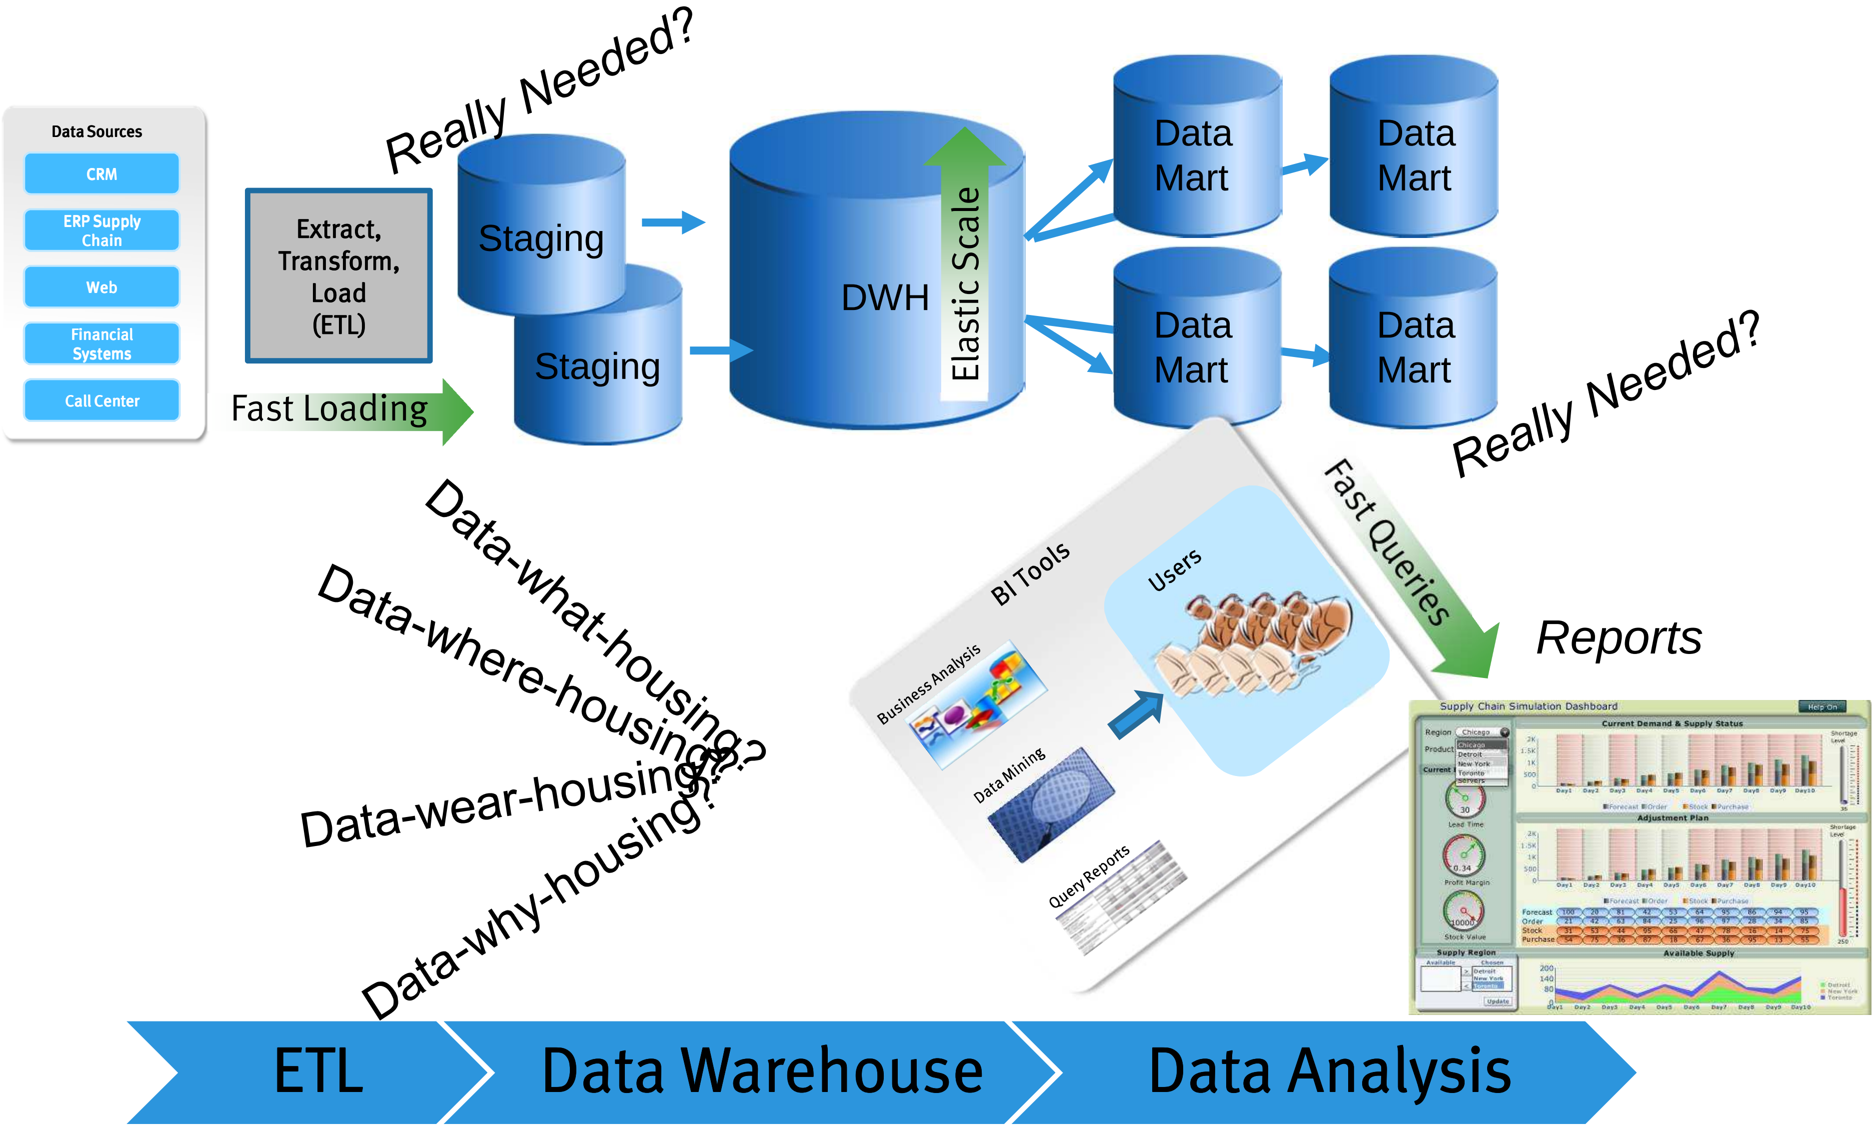

Datamart “Sprawl”

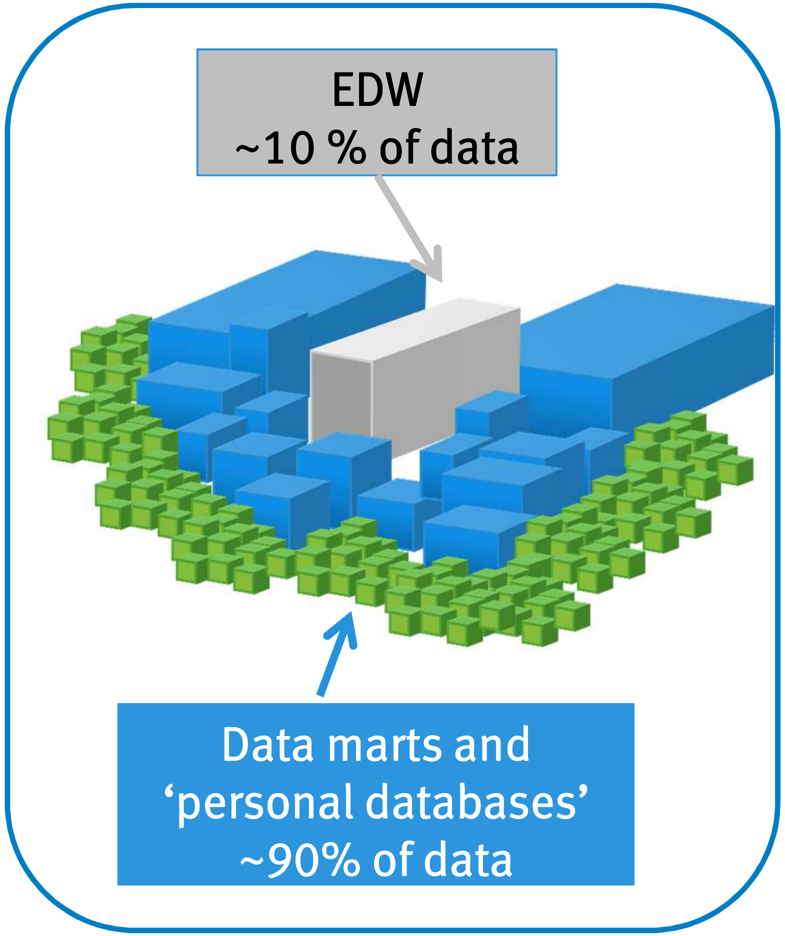

- Data is everywhere and growing
  - 44X data growth by 2020
  - 100s of data marts
  - ‘Shadow’ databases
- Critical business insight is outside EDW
- Centralized legacy systems are expensive
- System expansion is slow and process heavy
- Proprietary HW systems lag behind open systems innovation

Traditional solutions cannot scale to meet the DW/BI challenges

Business Intelligence Challenges (1)

Related to Infrastructure

- Higher service levels
  - DWH not allowed to be down for a few days
  - Need for backup/recovery/DR
  - No SPOF, high-availability architecture
  - Don't forget security, auditing, compliancy, data leakage prevention, customer privacy considerations (think Facebook and Google)
- Massive growth
  - According to research firms, unstructured data will be biggest growth factor for companies
  - Business Intelligence is #2
  - Soon we will see datawarehouses 100's of Terabytes in size (And the first Petabyte customers)
  - Business people want to store more and more in the DWH

Business Intelligence Challenges (2)

Related to Infrastructure

- Loading time
  - DWH needs to have up-to-date info
  - Load times of multiple days is simply no longer acceptable
  - 24H is max (for the whole process, not just loading)
  - Long term, drive to real-time (ouch!)
- "Scan" time (how long does it take to run a query)
  - More data
  - More impatient end users
  - More ad-hoc queries
  - Cannot optimize this anymore with classic SQL tuning and database tricks & magic

Business Intelligence Challenges (3)

Related to Infrastructure

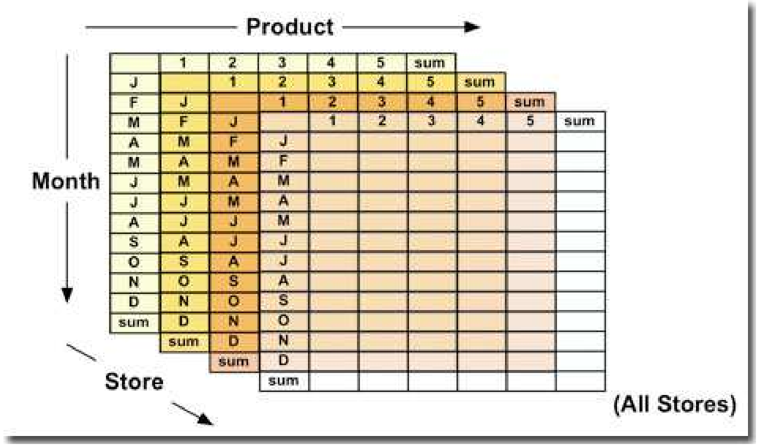

And finally… New paradigms

- Multi-dimensional OLAP databases
- In-memory statistical calculations
  - Needs to load a data subset in memory real quick
- Web users accessing BI data
  - Of course, through web applications
  - Massive scale-up in # of parallel transactions

Business and technology challenges

- Increased regulatory scrutiny and business reporting requirements
  - Insufficient data transparency across all risk exposures
  - Processing cycle taking too long
  - Lengthy reporting turn-around time
  - Need to retain more data over extended period of time

- All these need to be enabled by IT and supporting Infrastructure
  - Maintain performance amid escalating data volumes
  - Aggregate data sets from many silos
  - Ad hoc analysis and reporting occurring more frequently
  - Enable accessibility to historical raw data
  - Enable easy provisioning and expansion

- Upgrading existing infrastructure is very expensive and in many cases is cost prohibitive

An illustration of the massive technical challenge

Multiple credit servicing systems with inconsistent data of variable integrity and many manual processes

Legacy applications not taking advantage of middleware and hence not inter-operable

Ad hoc linkages between financial and risk data

Mixture of 2 tier & 3 tier application access layers, limit inter-operability

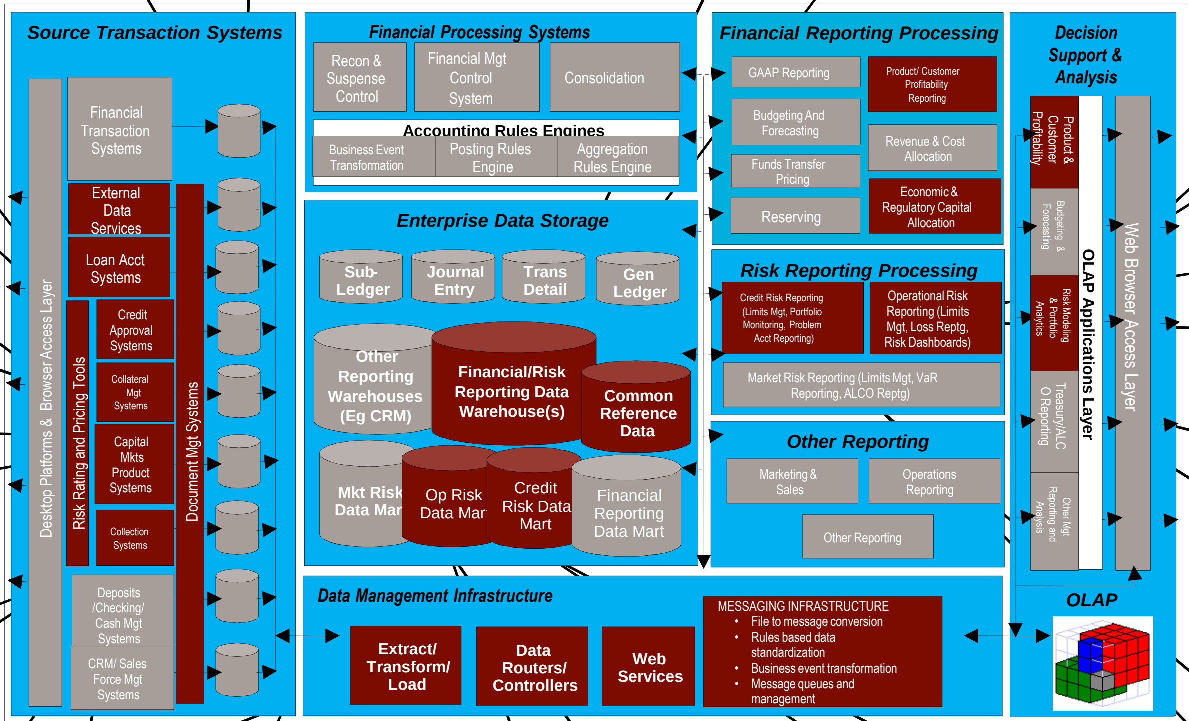

Informal & paper-based risk rating and credit processes

Ad hoc and paper-based loan documentation

Stand Alone Product Pricing Tools

Ad hoc and manual regulatory reporting processes, not transparent or readily auditable

Multiple, siloed risk analysis and modeling tools not using consistent data. Limited analytical Capability

Stand alone Recoveries systems and processes

Manual and ad hoc data loads, instead of programmed ETL

Limited operational risk data often on spreadsheets and Access Databases

Credit risk data in multiple repositories and forms, often in Access

Ad hoc and paper-based risk reporting processes not linked or using inconsistent data

Key Systems and Processes significantly impacted

EMC² where information lives®

Implementing a Data Warehouse

- In many organizations IT people want to huddle and work out a warehousing plan, but in fact
  - The purpose of a DW is decision support
  - The primary audience of a DW is therefore College decision makers
  - It is College decision makers therefore who must determine
    - Scope
    - Priority
    - Resources
- Decision makers can't make these determinations without an understanding of data warehouses
- It is therefore imperative that key decision makers first be educated about data warehouses
  - Once this occurs, it is possible to
    - Elicit requirements (a critical step that's often skipped)
    - Determine priorities/scope
    - Formulate a budget
    - Create a plan and timeline, with real milestones and deliverables!

What Takes Up the Most Time?

- You may be surprised to learn what DW step takes the most time
- Try guessing which:
  - Hardware
  - Physical database setup
  - Database design
  - ETL
  - OLAP setup

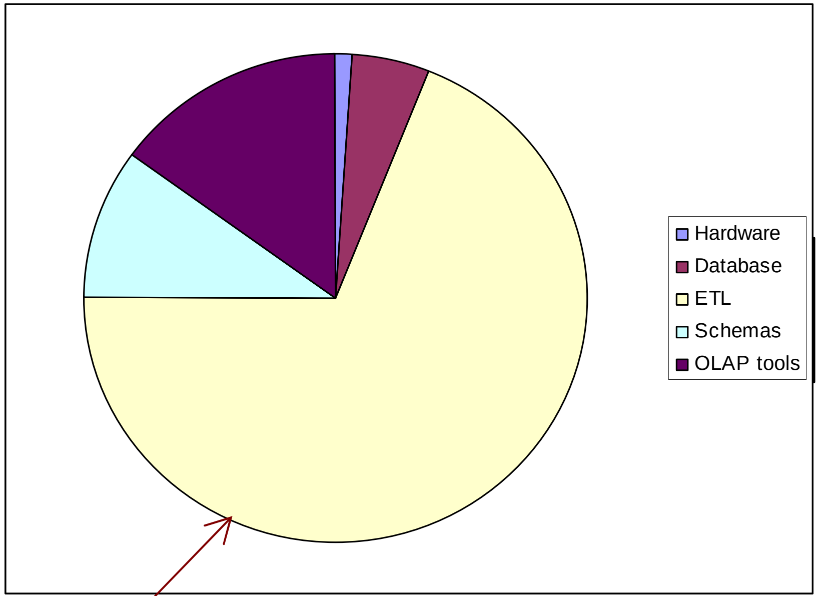

Acc. to Kimball & Caserta, **ETL** will eat up 70% of the time.
Other analysts give estimates ranging from 50% to 80%.

The most often underestimated part of the warehouse project!

Lesson…

- You cannot tune a data warehouse
  - The data growth will defeat any non-scalable infrastructure
  - The number of tables and rows will defeat any broad attempt to pre-join (i.e. denormalize) data
  - The number of users and the variety of questions will flush every cache and defeat every index scheme
  - Attempts to build specialized, redundant, data structures: pre-aggregated data or materialized views, add more and more operational complexity until you cannot build the structures in a timely manner

EMC² where information lives®

Reality Check (and rhetorical question)…
Do we see this in the real world?

Data Warehouses:

- With dozens of indexes and materialized views
- With hundreds of aggregate tables and pre-aggregated data marts
- With hand-tuned queries and no ability to support new or ad hoc business questions
- Which cannot load nightly data into all of the indexes, aggregates, marts, and materialized views in the batch window
- Which stop or slow decision-making during peak seasons
- Which constrain business growth due to the inability to expand
- Which require BI application redesign to add data, users, or attributes
- Which are fragile and require constant operational care and feeding

Scan Rate Summary: A Competitive Analysis

Example	Scan Rate (secs)
Single Node	13,320
20 Nodes	666
Row compression	266
Partition elimination	10.6
Columnar Compression	2.5
Columnar Projection	.25

Teradata and Netezza stop here

Exadata stops here

And this example only considers scanning… joins benefit from the shared-nothing architecture as well… and Exadata does not perform shared-nothing joins. It will not scale…

Differentiator	Explanation
Fast Scanning (without tuning)	Tuning a data warehouse is nearly impossible because you cannot predict tomorrow's business request. Any attempt to apply a trick to boost performance for one task will influence the others. GP does not depend on tricks like indexes, materialized views, etc. No magic tricks but easy to understand, smart "shared-nothing" architecture to run fast and scale
Fast Data Load (near real-time, parallel vs. sequential)	More accurate and recent data allows organizations to get better results faster. Loading in parallel on the segment servers instead of through a single frontend node. "External tables" use data directly from external databases and flat files – not requiring temporary files on disk. Directly update/insert back into external databases where needed saving time
Pipelining	No save to disk of intermediate results. Instead, intermediate results will be streamed directly to the next operation – fully utilizing the available CPU power, resulting in shorter queries and loads
Low administration requirement	Because no tuning is required, increasing capacity is easy and on-the-fly (no downtime), and does not require database redesign. Administrators can spend time on implementing new business functionality instead of tuning.
Polymorphic Storage	Allows data in one table to be saved with different methods, compression levels, etc. so that recent data can have highest performance and older data stored more efficiently
Integration with ETL & BI vendors	Nearly all well-known ETL and BI tools offer native integration with Greenplum
Hadoop / MapReduce integration for non-structured and semi-structured data	By connecting special Hadoop compute nodes directly on the interconnect, one platform can manage all sorts of data at extreme performance – not just structured (database) data. Think of bulk data feeds from the Internet (such as Twitter or marketing data)
Analytical and statistical functions, SAS integration	Complex statistical functions can run directly on the cluster without having to load to an application server first. SAS offers a special integration to make use of this
High Availability and Business Continuity	As BI quickly becomes mission critical, integration with classic EMC storage allows extreme availability and quick recovery in case of disasters
Based on standard Intel Architecture - low cost and high flexibility	No special hardware required – allows to benefit from improvements in server performance (Moore's law) without redesign – and customers can run Software Only versions on regular servers) and even deploy on standard virtual machines for test/dev purposes

Top-10 differentiators of EMC Greenplum

big•data \ datasets so large they break traditional IT infrastructures.

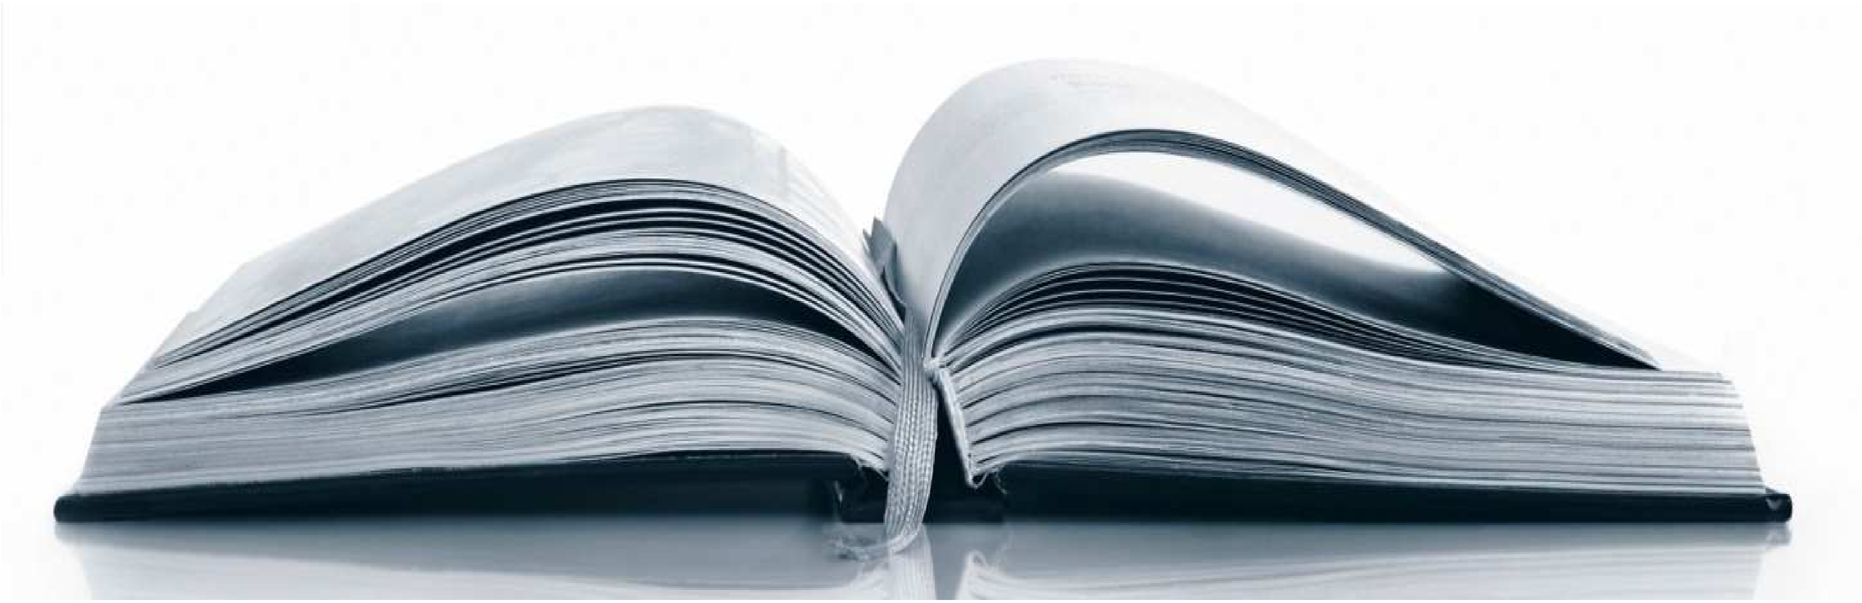

BIG DATA TRANSFORMS BUSINESS

EMC²

Most Enterprises Will Have Multiple "Big Data" Use Cases

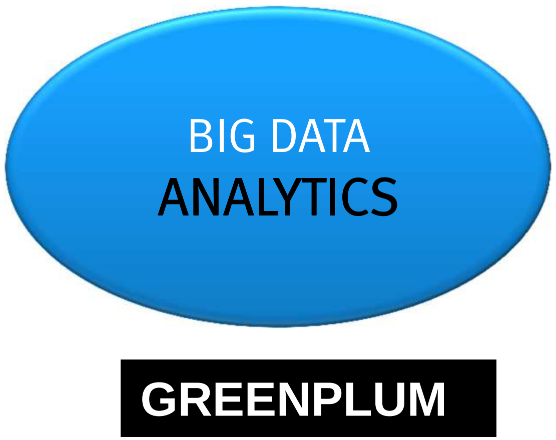

Most Enterprises Will Have Multiple “Big Data” Use Cases

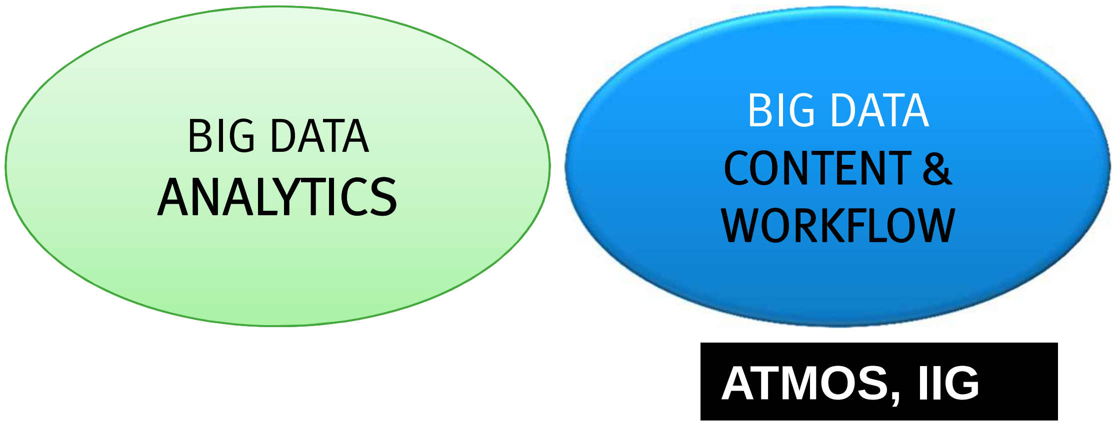

Most Enterprises Will Have Multiple "Big Data" Use Cases

BIG DATA
ANALYTICS

BIG DATA
CONTENT &
WORKFLOW

BIG DATA
ENTERPRISE

VMAX, VNX

Most Enterprises Will Have Multiple “Big Data” Use Cases

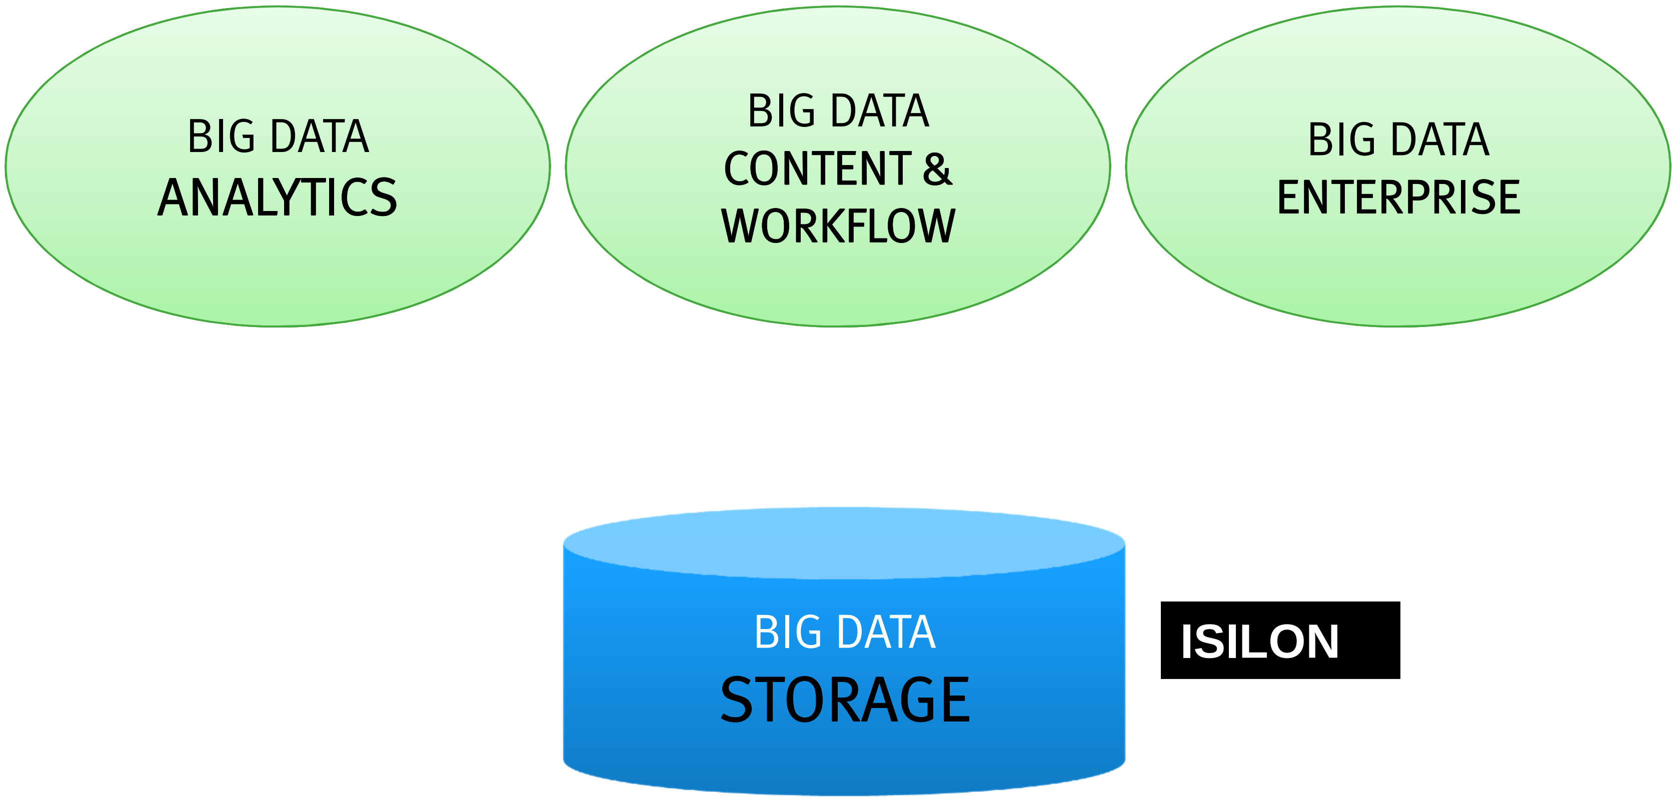

Most Enterprises Will Have Multiple “Big Data” Use Cases

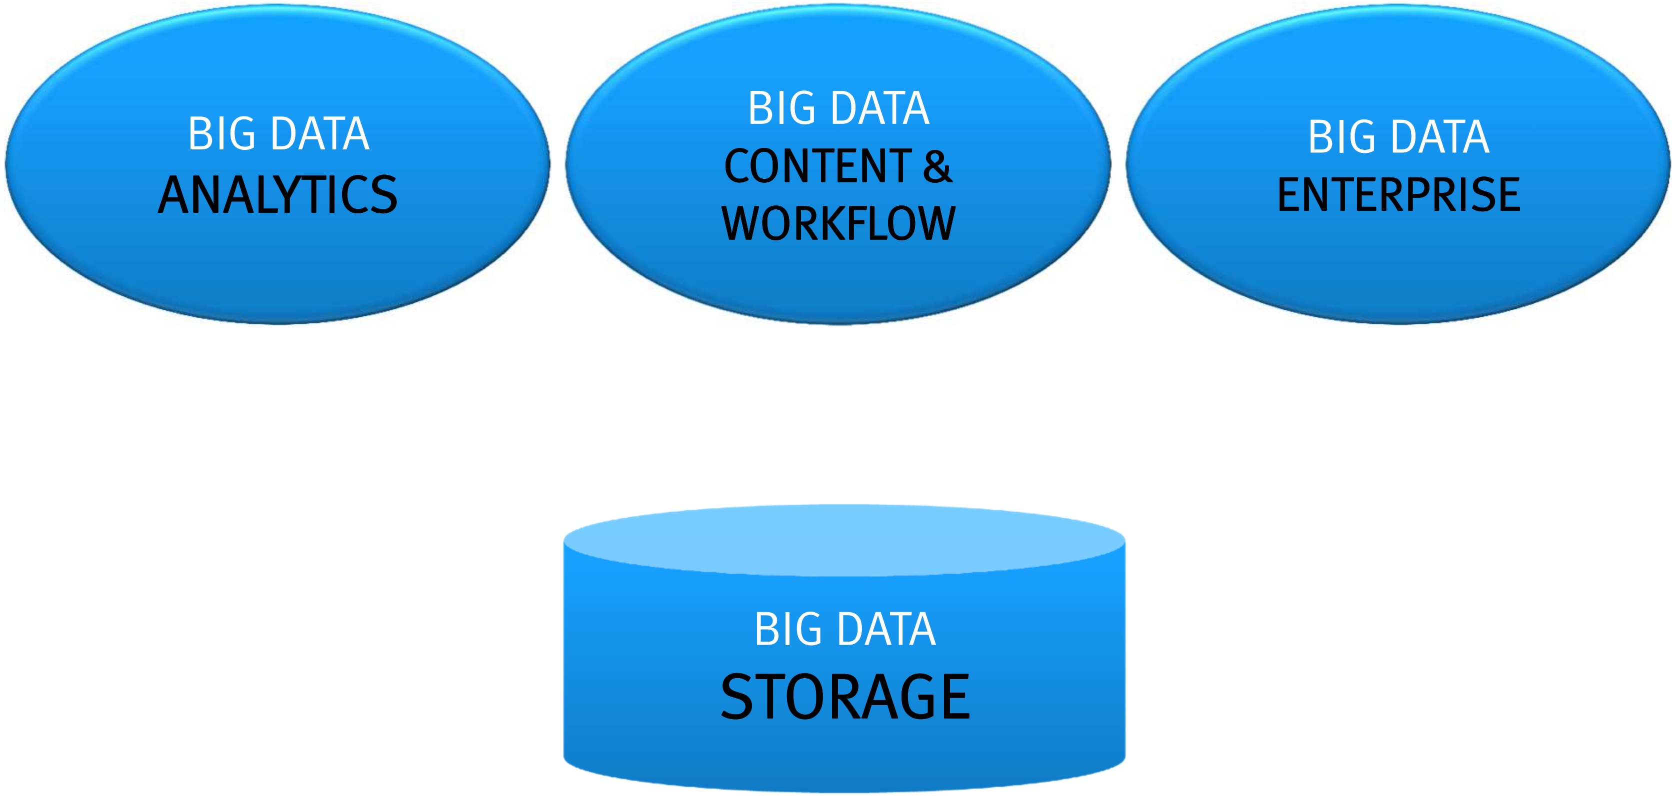

Today’s Focus – Big Data Analytics

BIG DATA
ANALYTICS

Big Data Analytics Is In Every Industry

Increase Profit Margins With Big Data

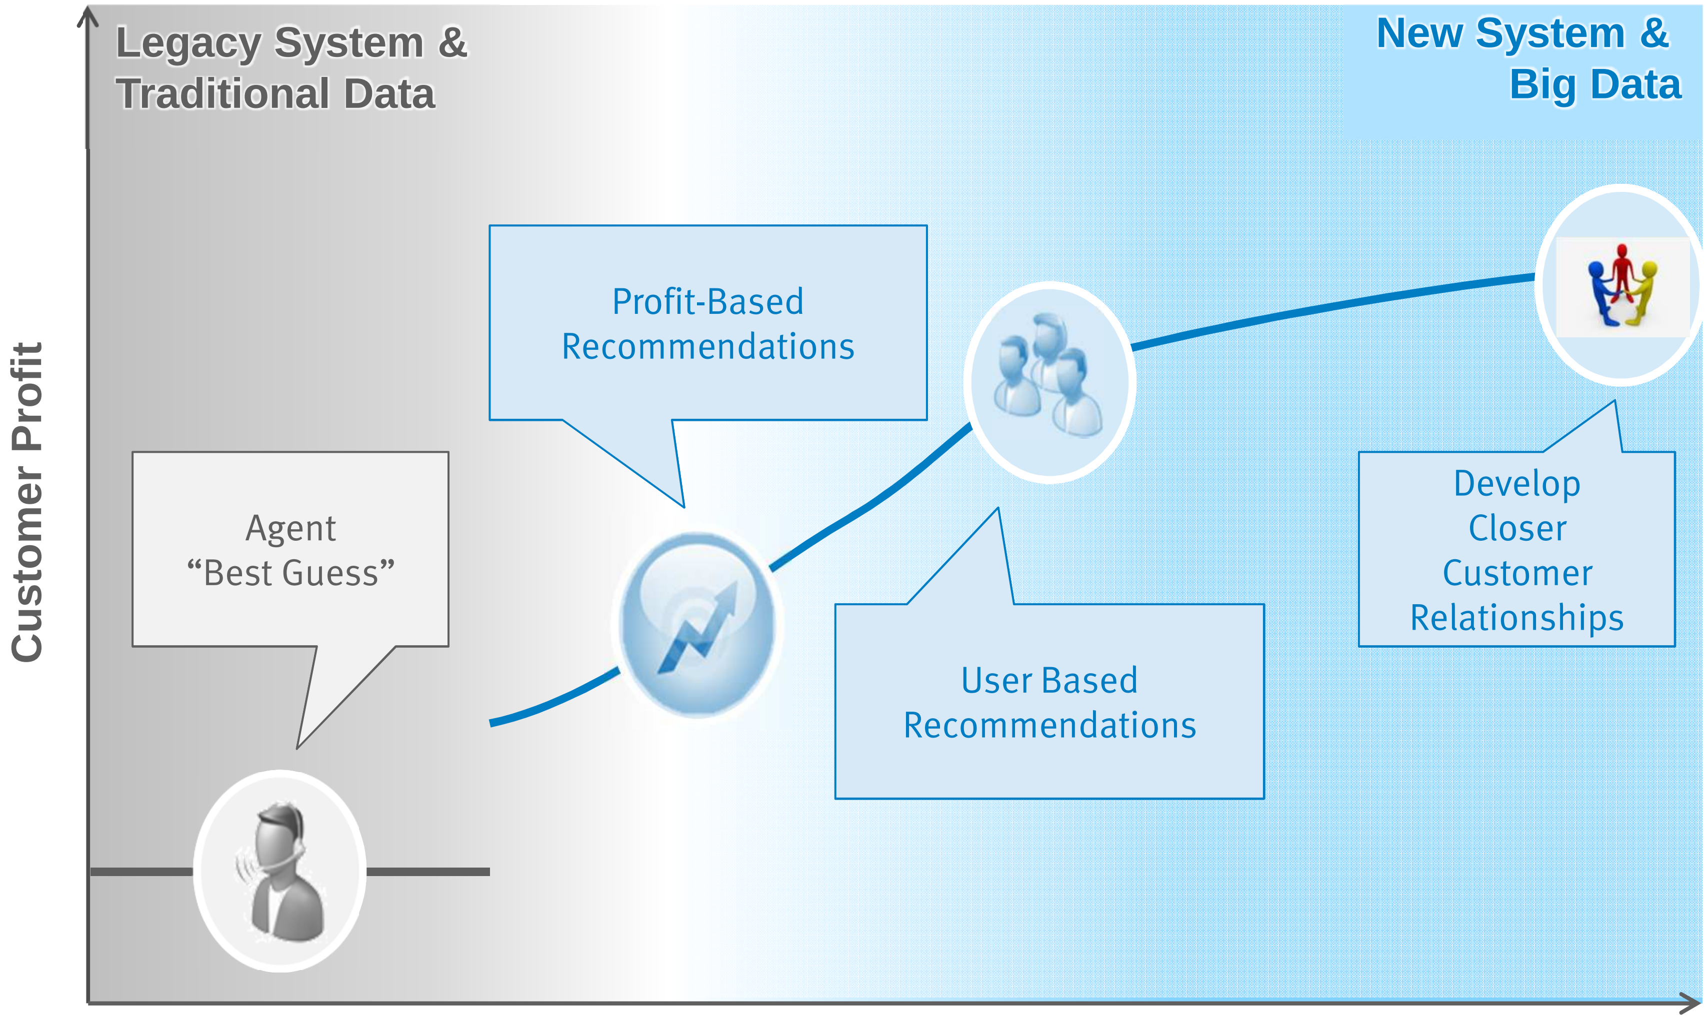

Personalized Risk Based Insurance

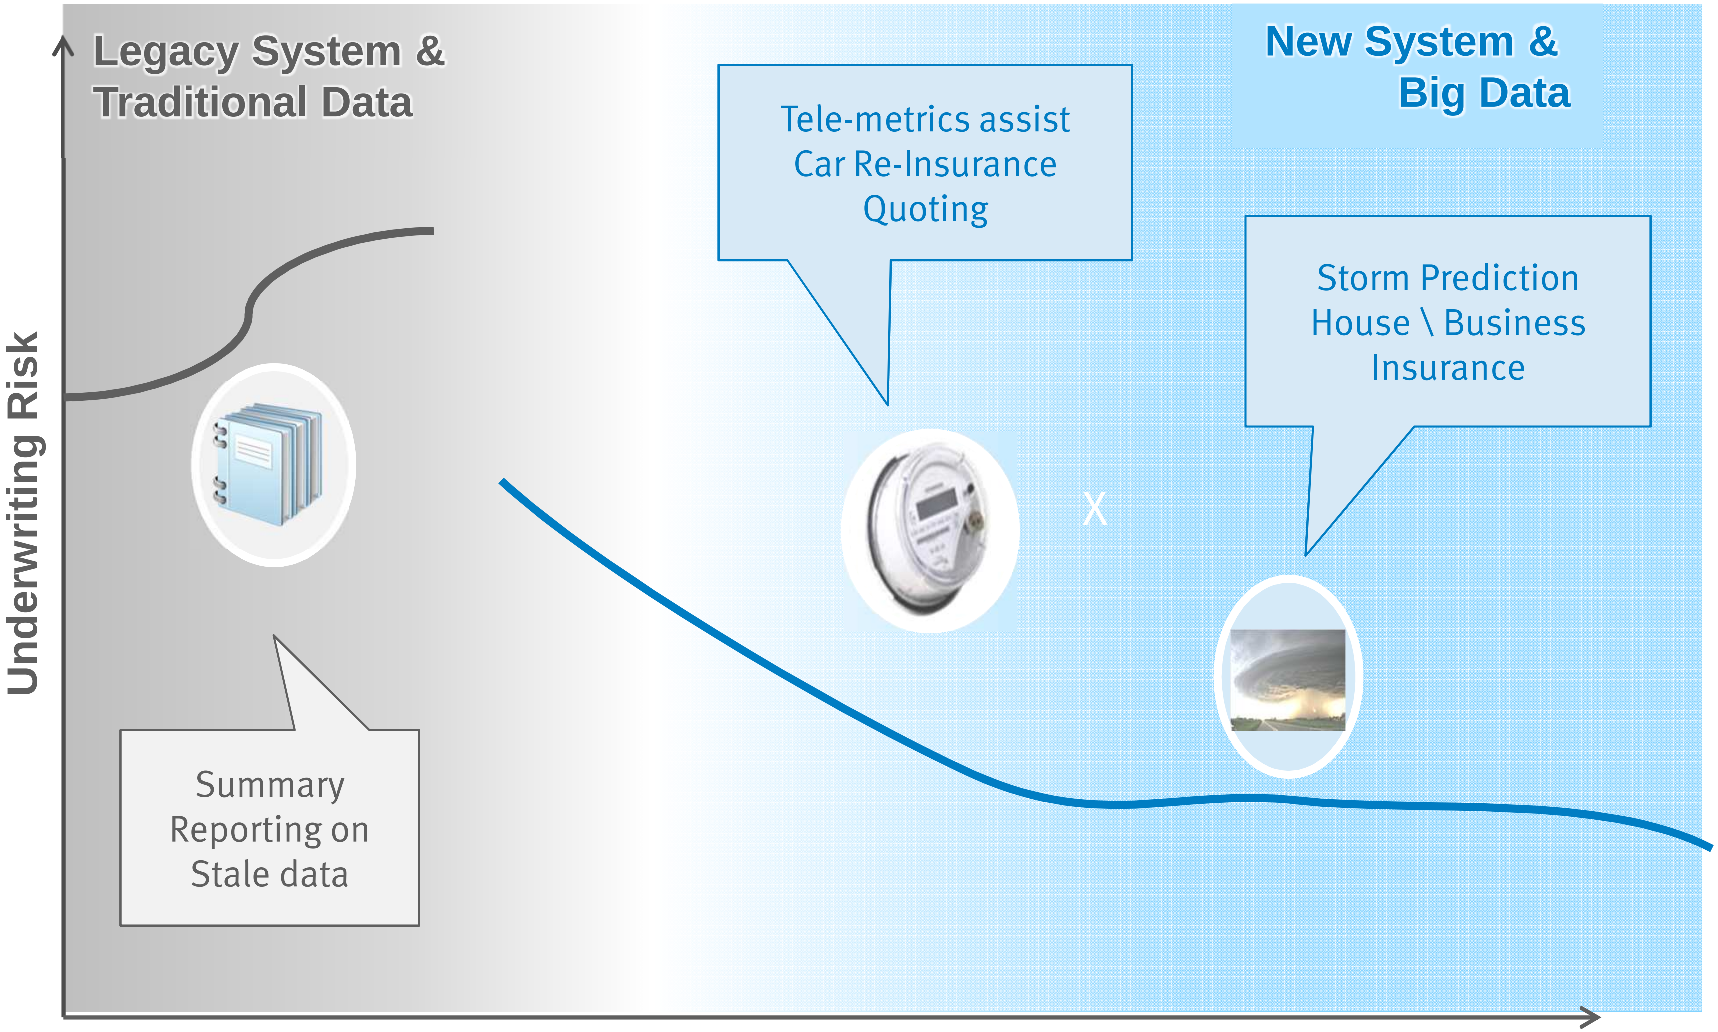

Deliver Better Healthcare With Big Data

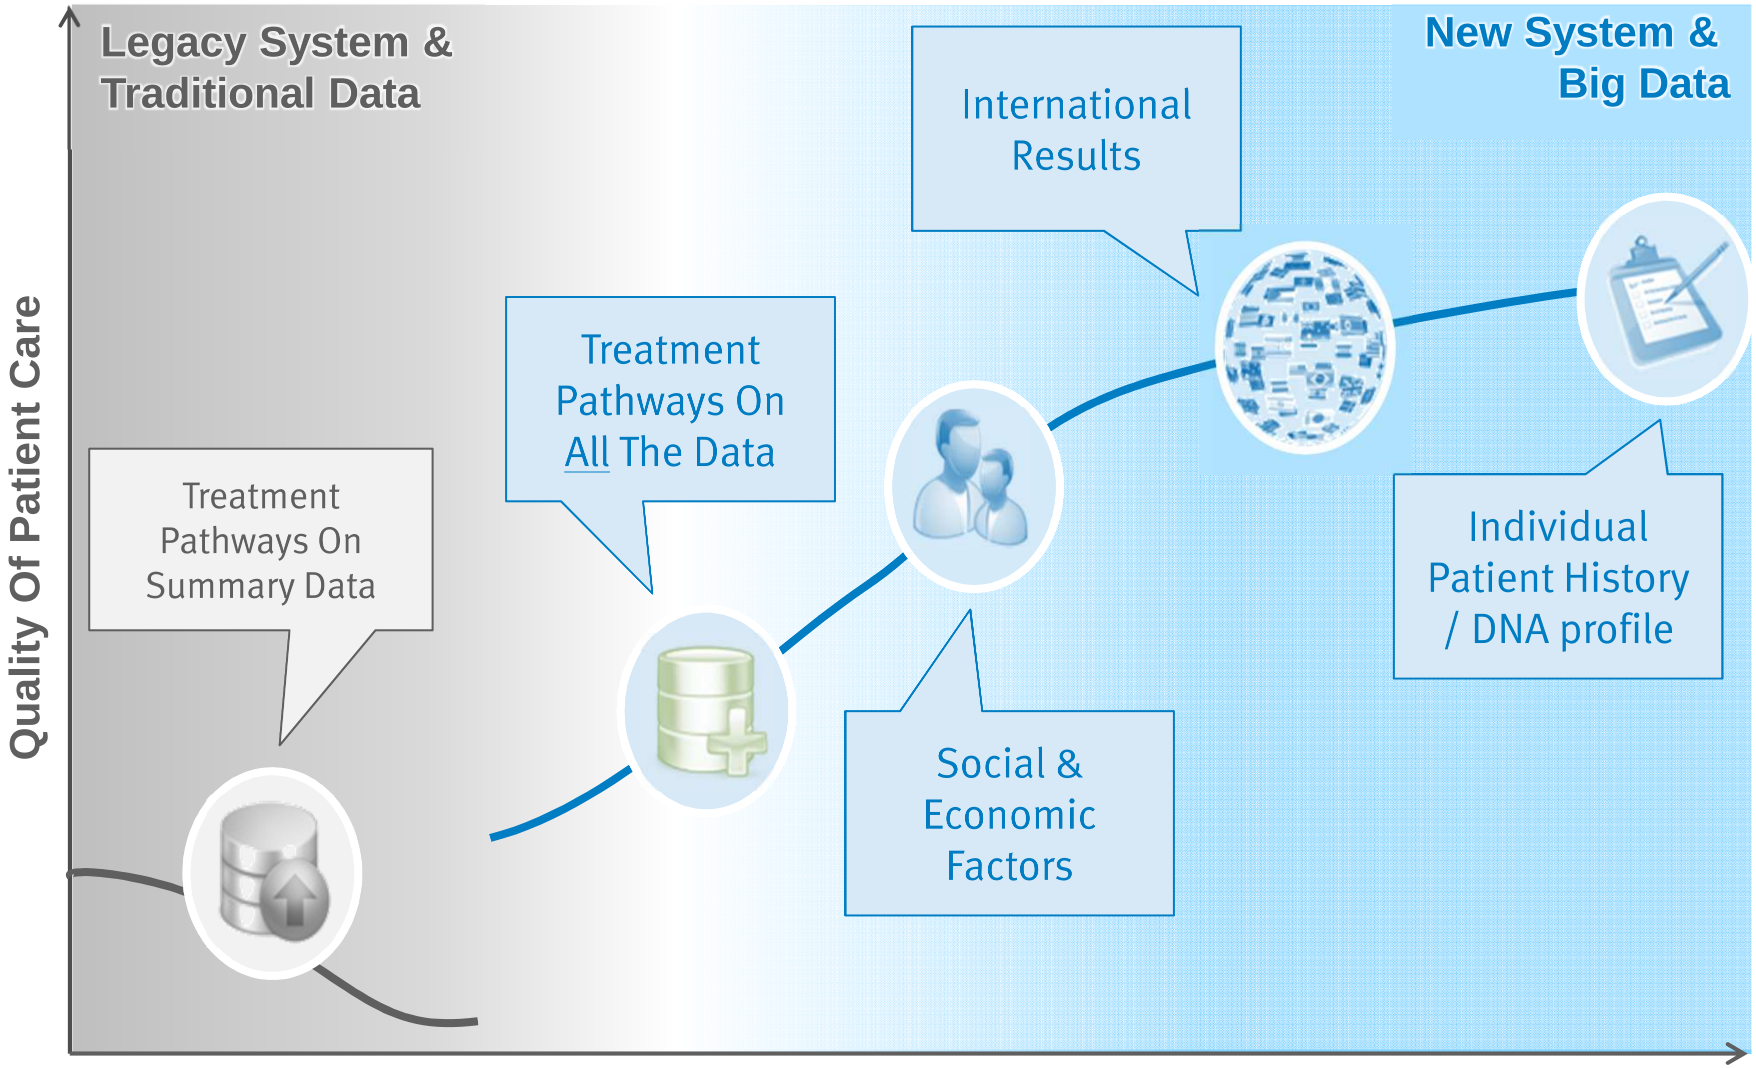

EMC IT Proven Customer Satisfaction

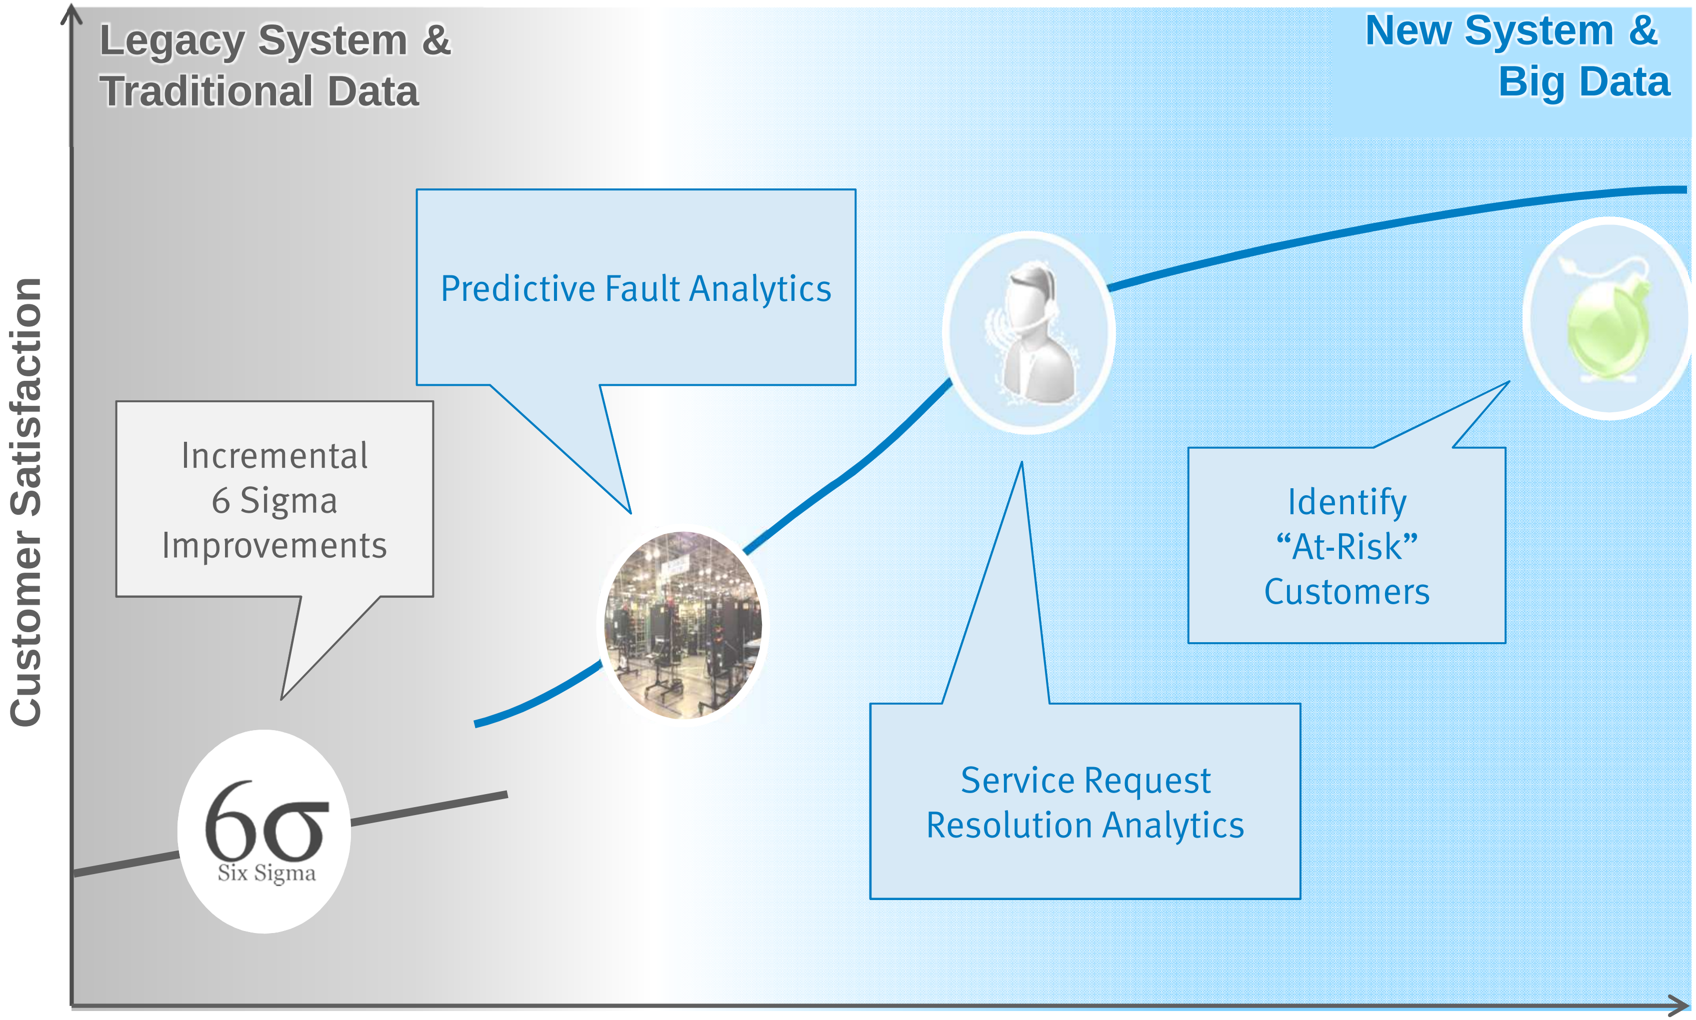

So what's the problem with Analytics on Big Data?

So lets take your standard analytic equation…☺

$$
\begin{aligned}
&\propto \frac{\prod_{i\neq k}\Gamma(n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i)}{\Gamma\big((\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i)+1\big)} \prod_{i\neq k}\frac{\Gamma(n_{(\cdot),v}^{i,-(m,n)}+\beta_v)}{\Gamma\big(\sum_{r=1}^{V} n_{(\cdot),r}^{i,-(m,n)}+\beta_r\big)} \\
&\times\Gamma(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k+1)\frac{\Gamma(n_{(\cdot),v}^{k,-(m,n)}+\beta_v+1)}{\Gamma\big((\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r)+1\big)} \\
&\propto \frac{\Gamma(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k+1)}{\Gamma\big((\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i)+1\big)}\frac{\Gamma(n_{(\cdot),v}^{k,-(m,n)}+\beta_v+1)}{\Gamma\big((\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r)+1\big)} \\
&= \frac{\Gamma(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k)\big(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k\big)}{\Gamma\big(\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i\big)\big(\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i\big)}\frac{\Gamma\big(n_{(\cdot),v}^{k,-(m,n)}+\beta_v\big)\big(n_{(\cdot),v}^{k,-(m,n)}+\beta_v\big)}{\Gamma\big(\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r\big)\big(\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r\big)} \\
&\propto \frac{\big(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k\big)}{\big(\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i\big)}\frac{\big(n_{(\cdot),v}^{k,-(m,n)}+\beta_v\big)}{\big(\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r\big)} \\
&\propto \big(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k\big)\frac{\big(n_{(\cdot),v}^{k,-(m,n)}+\beta_v\big)}{\big(\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r\big)}.
\end{aligned}
$$

What are the main mathematical functions??

$$\times\frac{\prod_{i\neq k}\Gamma(n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i)}{\Gamma((\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i)+1)}\prod_{i\neq k}\frac{\Gamma(n_{(\cdot),v}^{i,-(m,n)}+\beta_v)}{\Gamma(\sum_{r=1}^{V} n_{(\cdot),r}^{i,-(m,n)}+\beta_r)}$$

MANY MULTIPLICATIONS

$$\times\Gamma(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k+1)\frac{\Gamma(n_{(\cdot),v}^{k,-(m,n)}+\beta_v+1)}{\Gamma((\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r)+1)}$$

MANY DIVISIONS

$$\times\frac{\Gamma(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k+1)}{\Gamma((\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i)+1)}\frac{\Gamma(n_{(\cdot),v}^{k,-(m,n)}+\beta_v+1)}{\Gamma((\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r)+1)}$$

MANY ADDITIONS

$$=\frac{\Gamma(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k)(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k)}{\Gamma(\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i)(\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i)}\frac{\Gamma(n_{(\cdot),v}^{k,-(m,n)}+\beta_v)(n_{(\cdot),v}^{k,-(m,n)}+\beta_v)}{\Gamma(\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r)(\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r)}$$

$$\propto\frac{(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k)}{(\sum_{i=1}^{K} n_{m,(\cdot)}^{i,-(m,n)}+\alpha_i)}\frac{(n_{(\cdot),v}^{k,-(m,n)}+\beta_v)}{(\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r)}$$

WITH BIG DATA => PERFORMANCE PROBLEMS!!

$$\propto(n_{m,(\cdot)}^{k,-(m,n)}+\alpha_k)\frac{(n_{(\cdot),v}^{k,-(m,n)}+\beta_v)}{(\sum_{r=1}^{V} n_{(\cdot),r}^{k,-(m,n)}+\beta_r)}.$$

EMC² where information lives®

So your Problem is?

- Your data is getting too large to run complex queries?
- You have a shortage of resources with the required analytical skills??
- You would like your analytics to be part of every business process?

THE PLATFORM

Greenplum UAP

- Unified Analytics Platform for Big Data
  - Greenplum Database for structured data
  - Greenplum HD, Enterprise-ready Hadoop for unstructured data
  - Greenplum Chorus, the social platform for data science

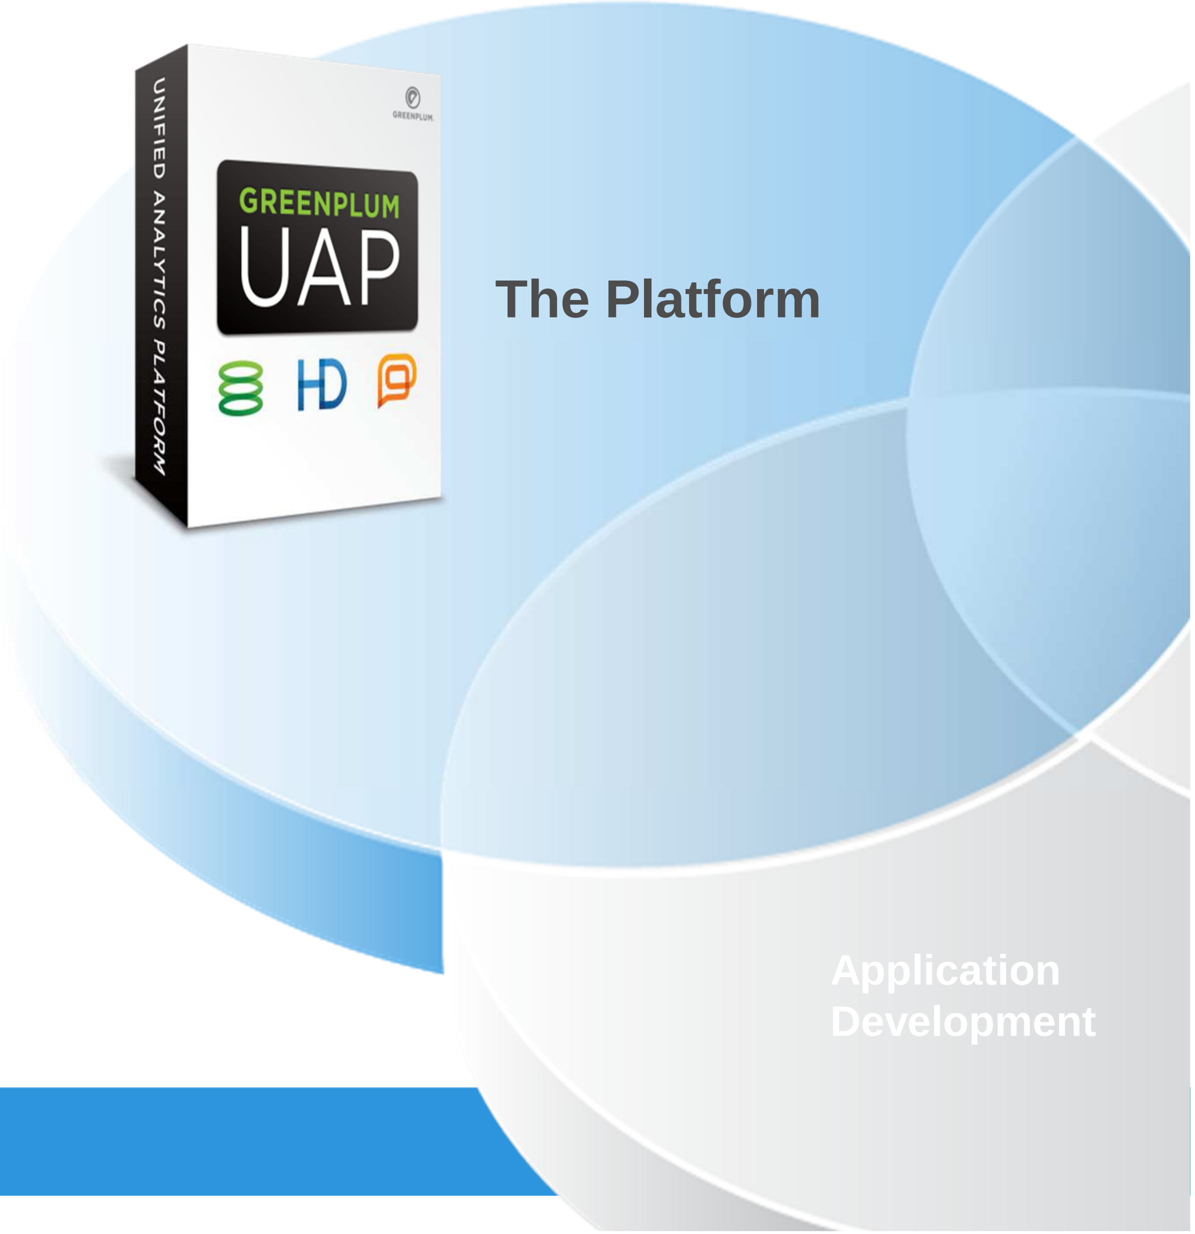

Greenplum Unified Analytic Platform

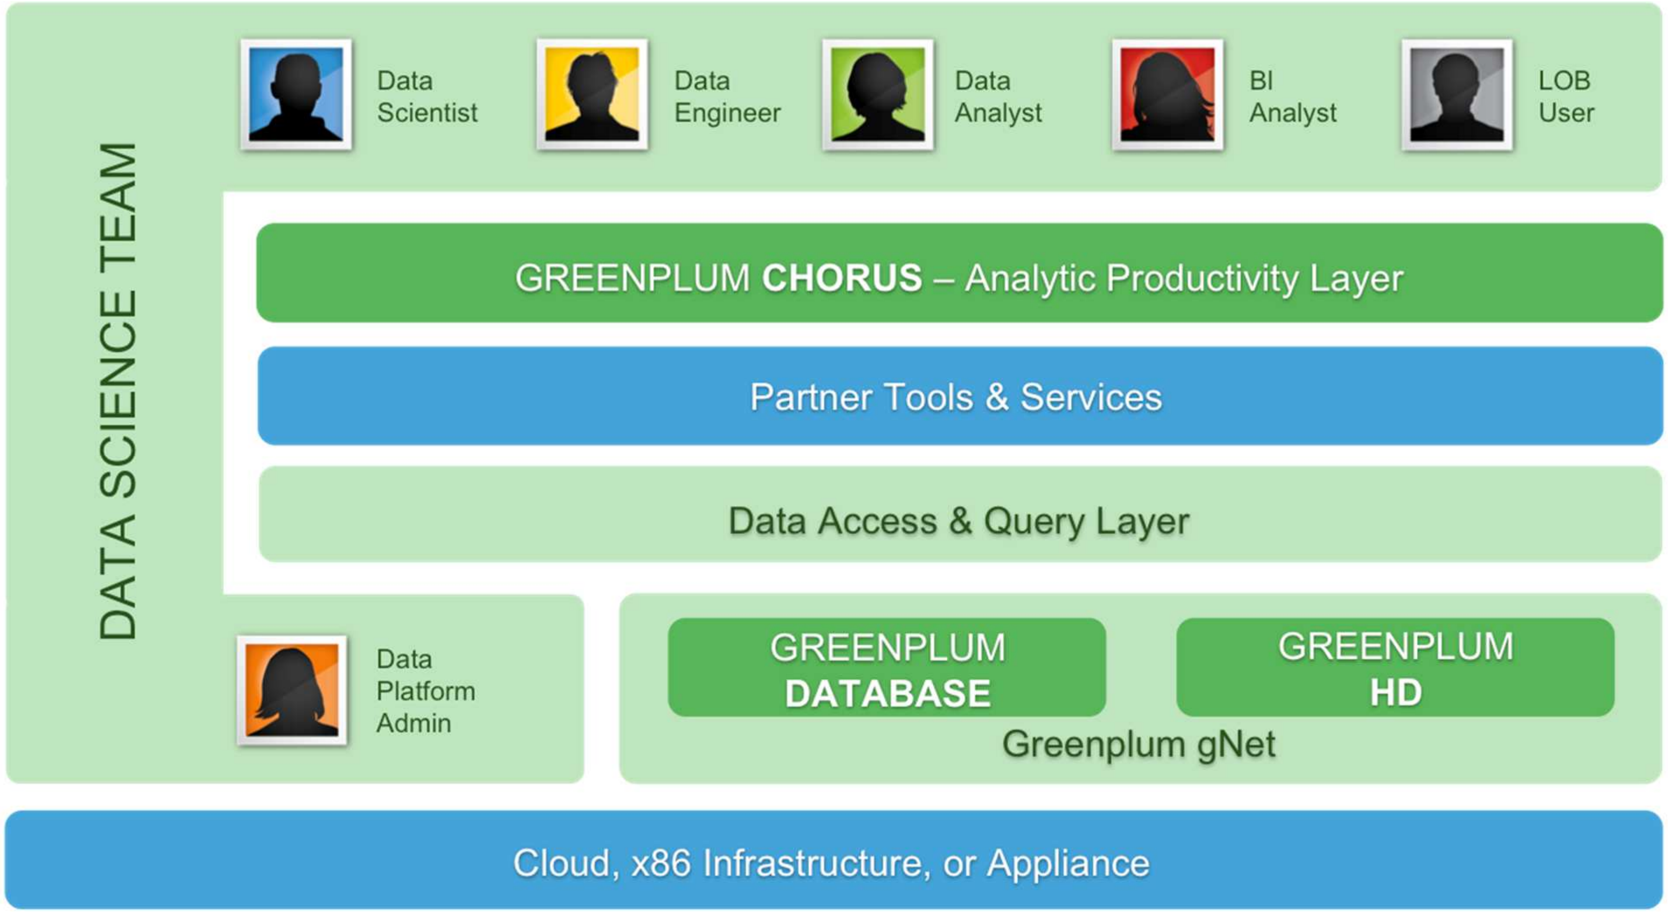

Analytics With The Greenplum Database

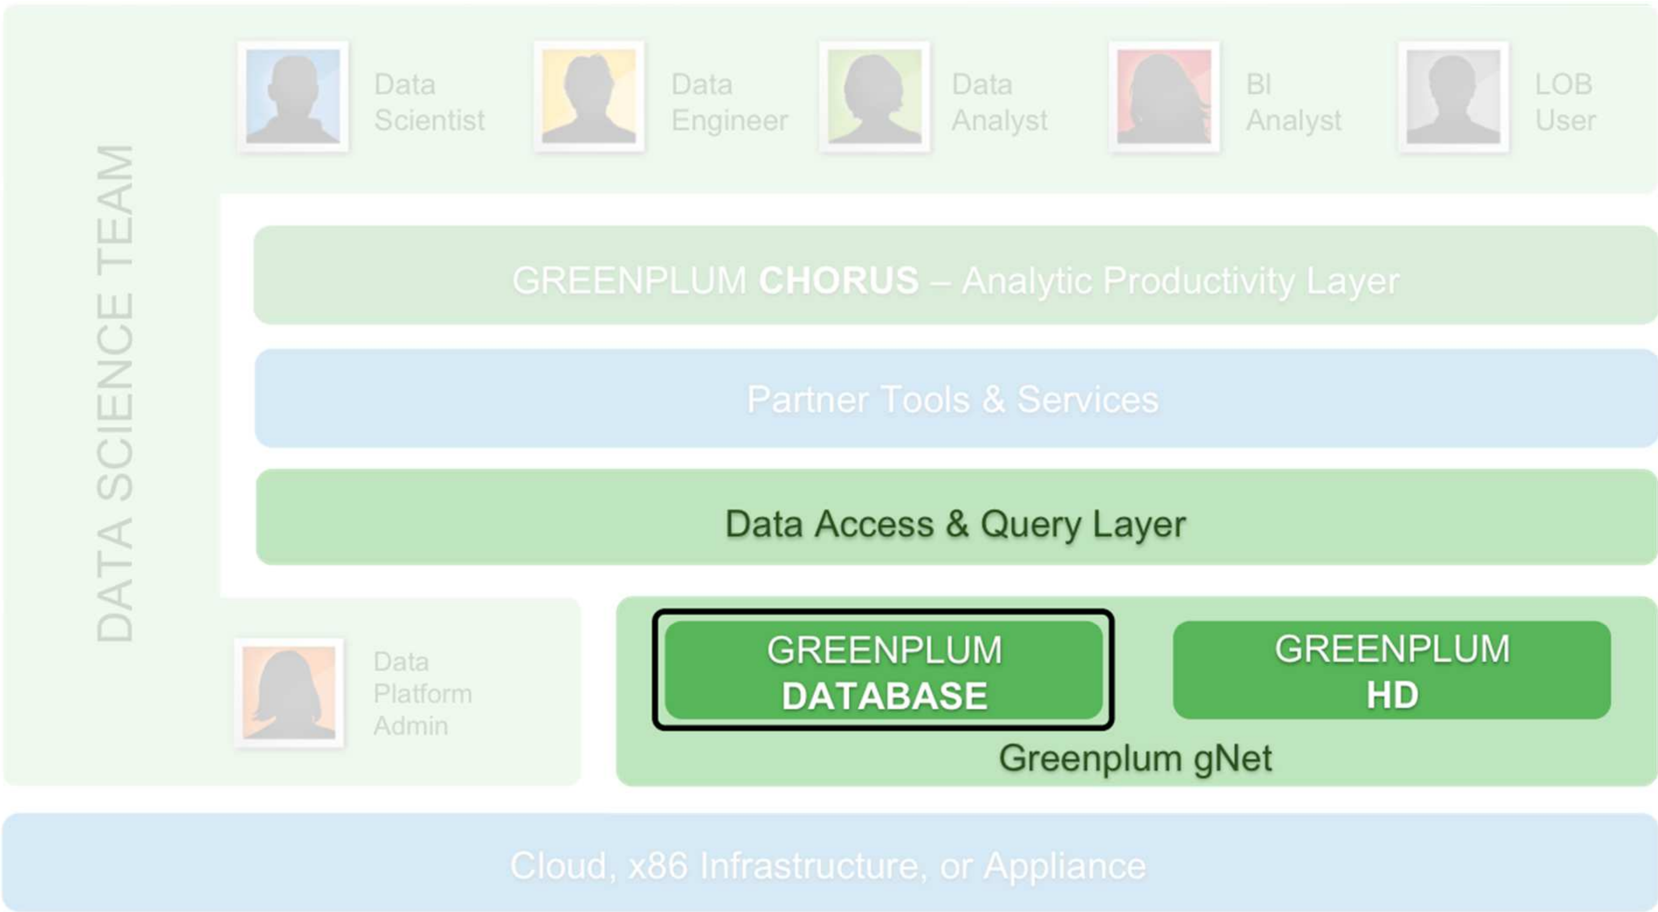

Architecture: Capacity & Performamce Via Parallelism

- MPP Scale-out architecture
  - Based on commodity servers & interconnects
- Automatic parallelization
  - Load and query like any database
  - Automatically distributed information and computation across all nodes
  - No need for manual partitioning or tuning
- Shared-nothing architecture
  - All nodes can scan and process in parallel
  - Independent storage to computation links
  - Linear scalability by adding nodes
  - On-line expansion when adding nodes

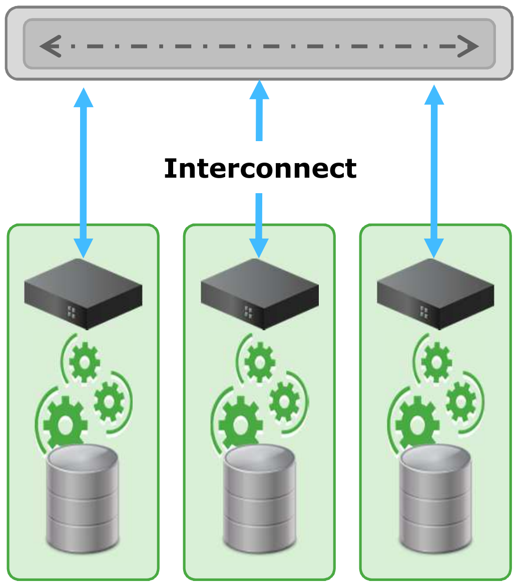

Greenplum Unified Analytic Platform

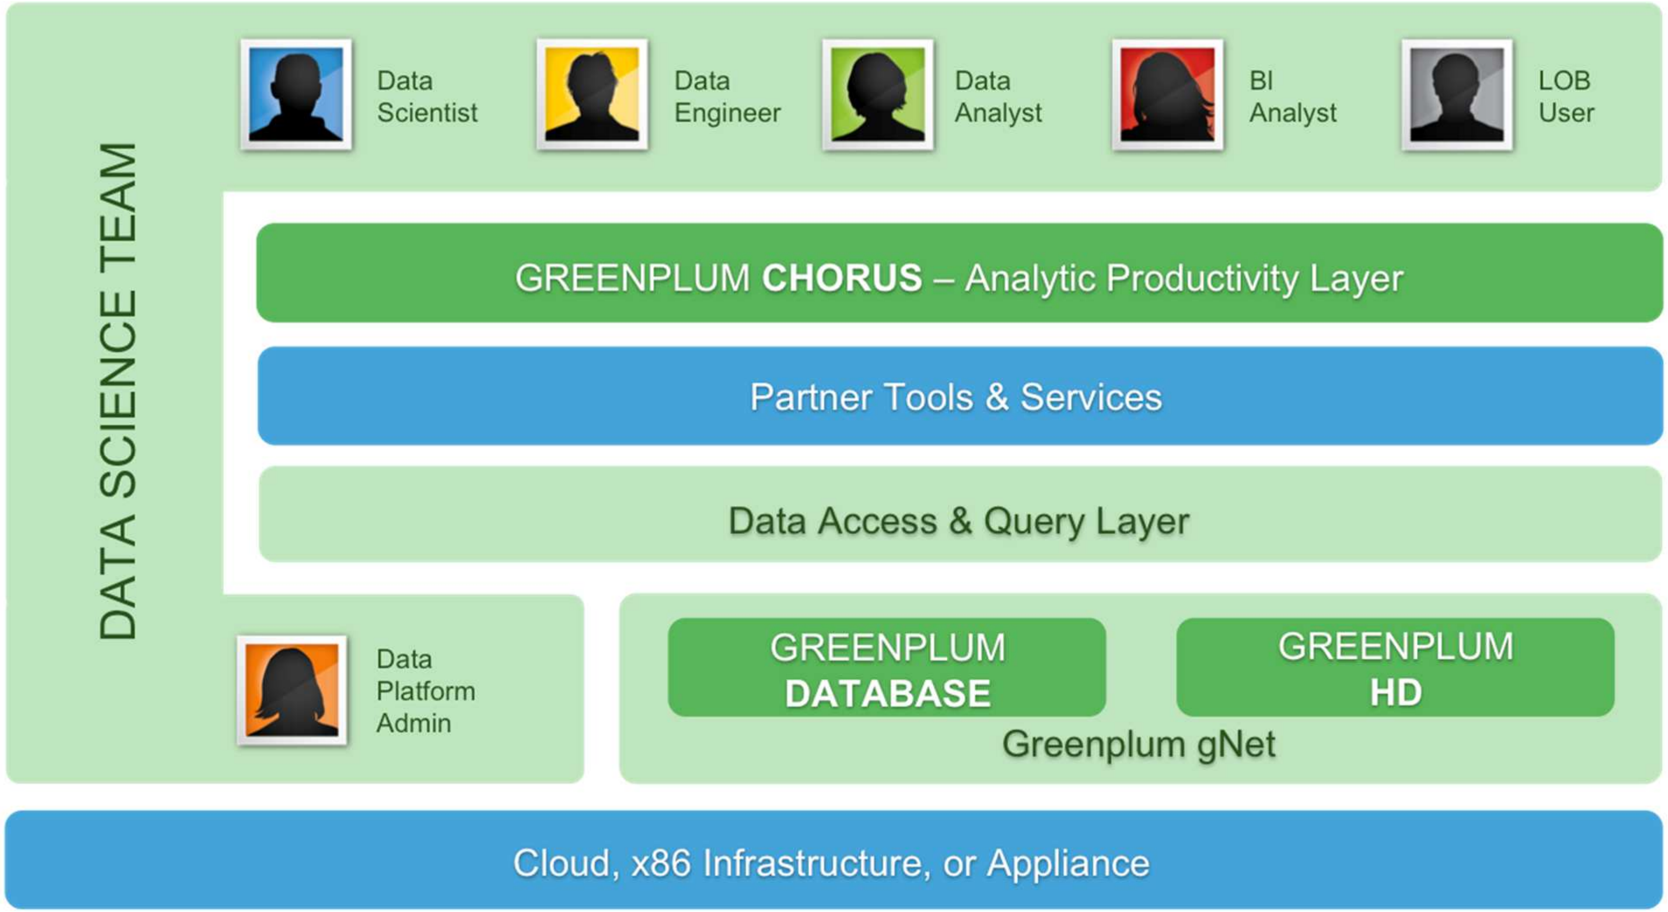

Greenplum Unified Analytic Platform

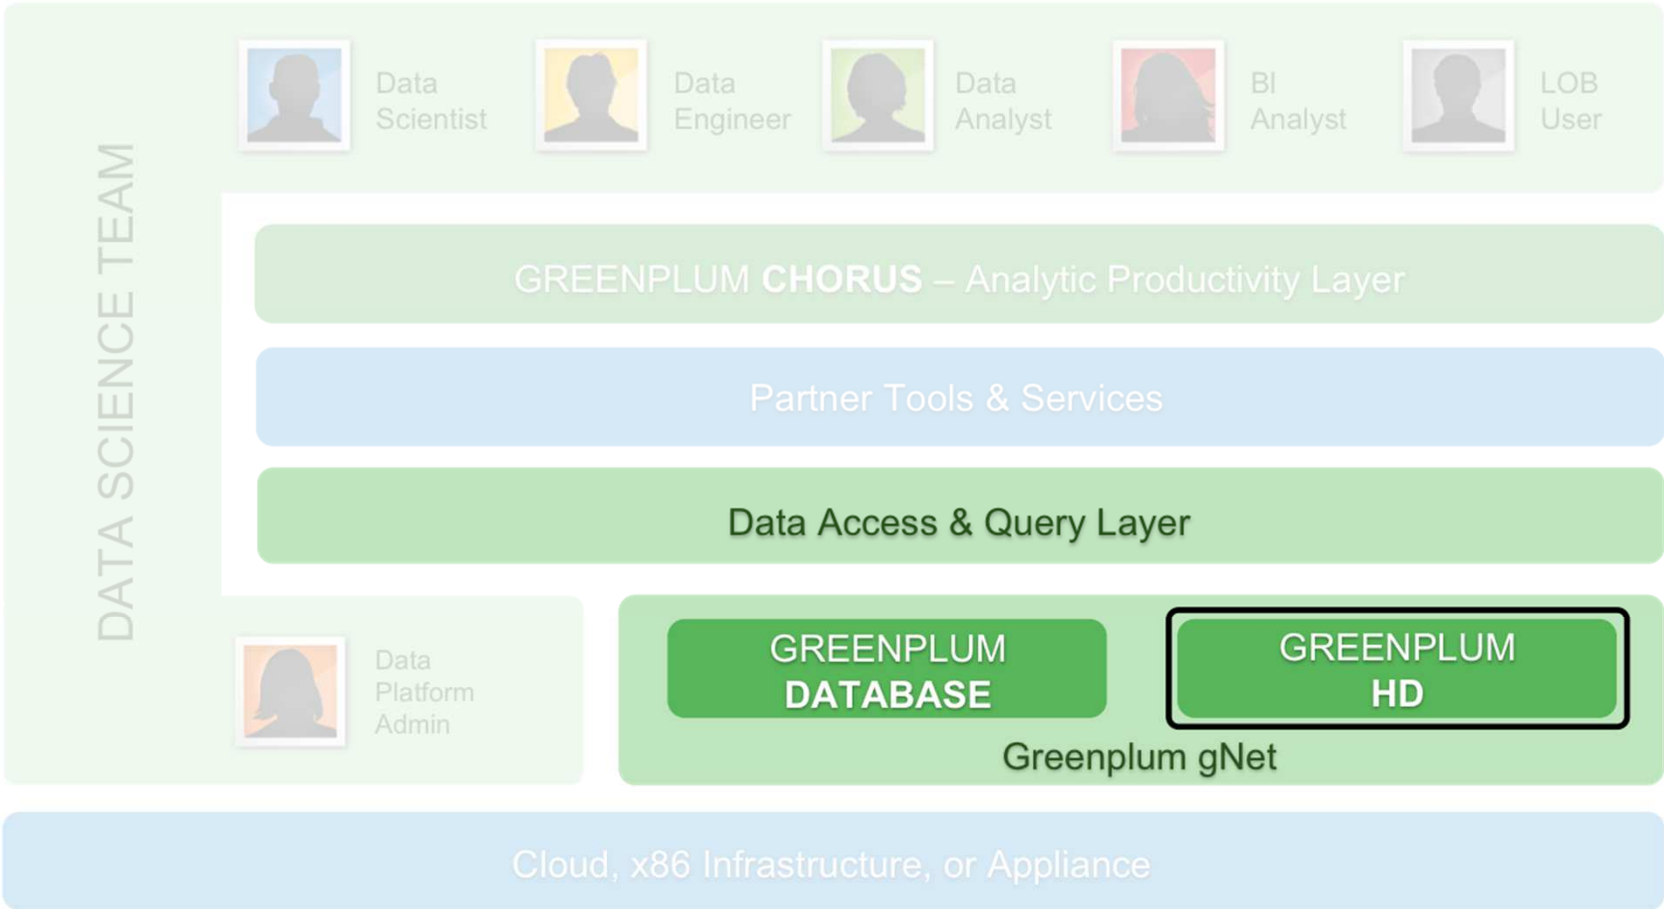

A Detailed Look At Greenplum HD

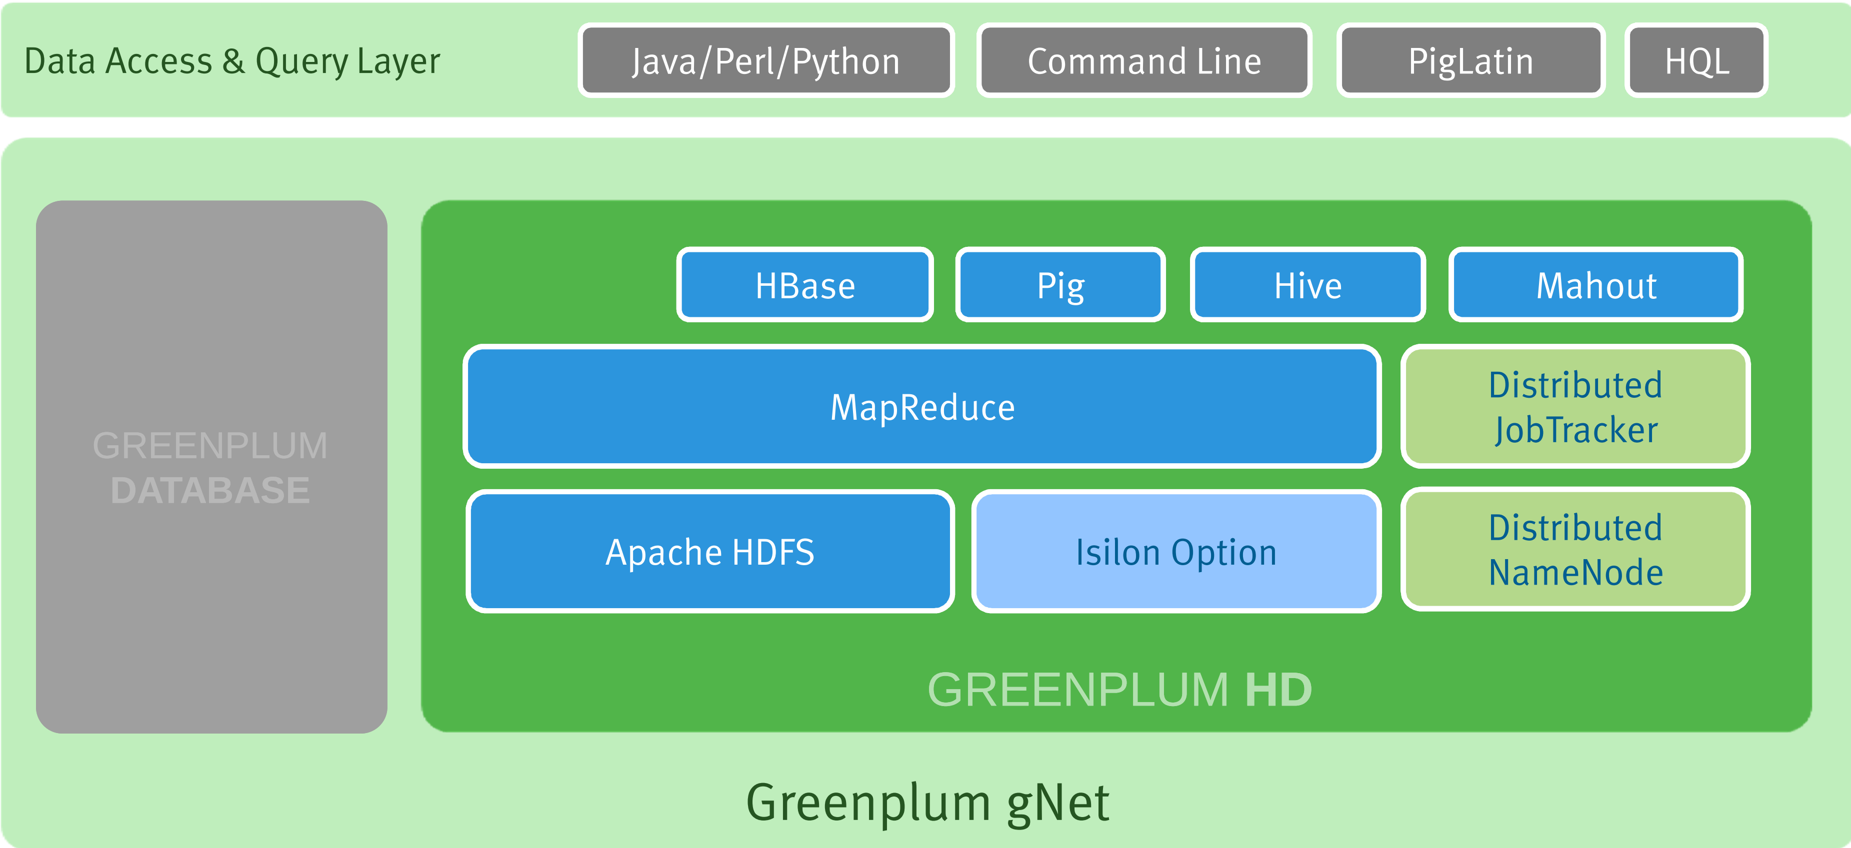

GREENPLUM gNET

UAP Unifies RDBMS and Hadoop

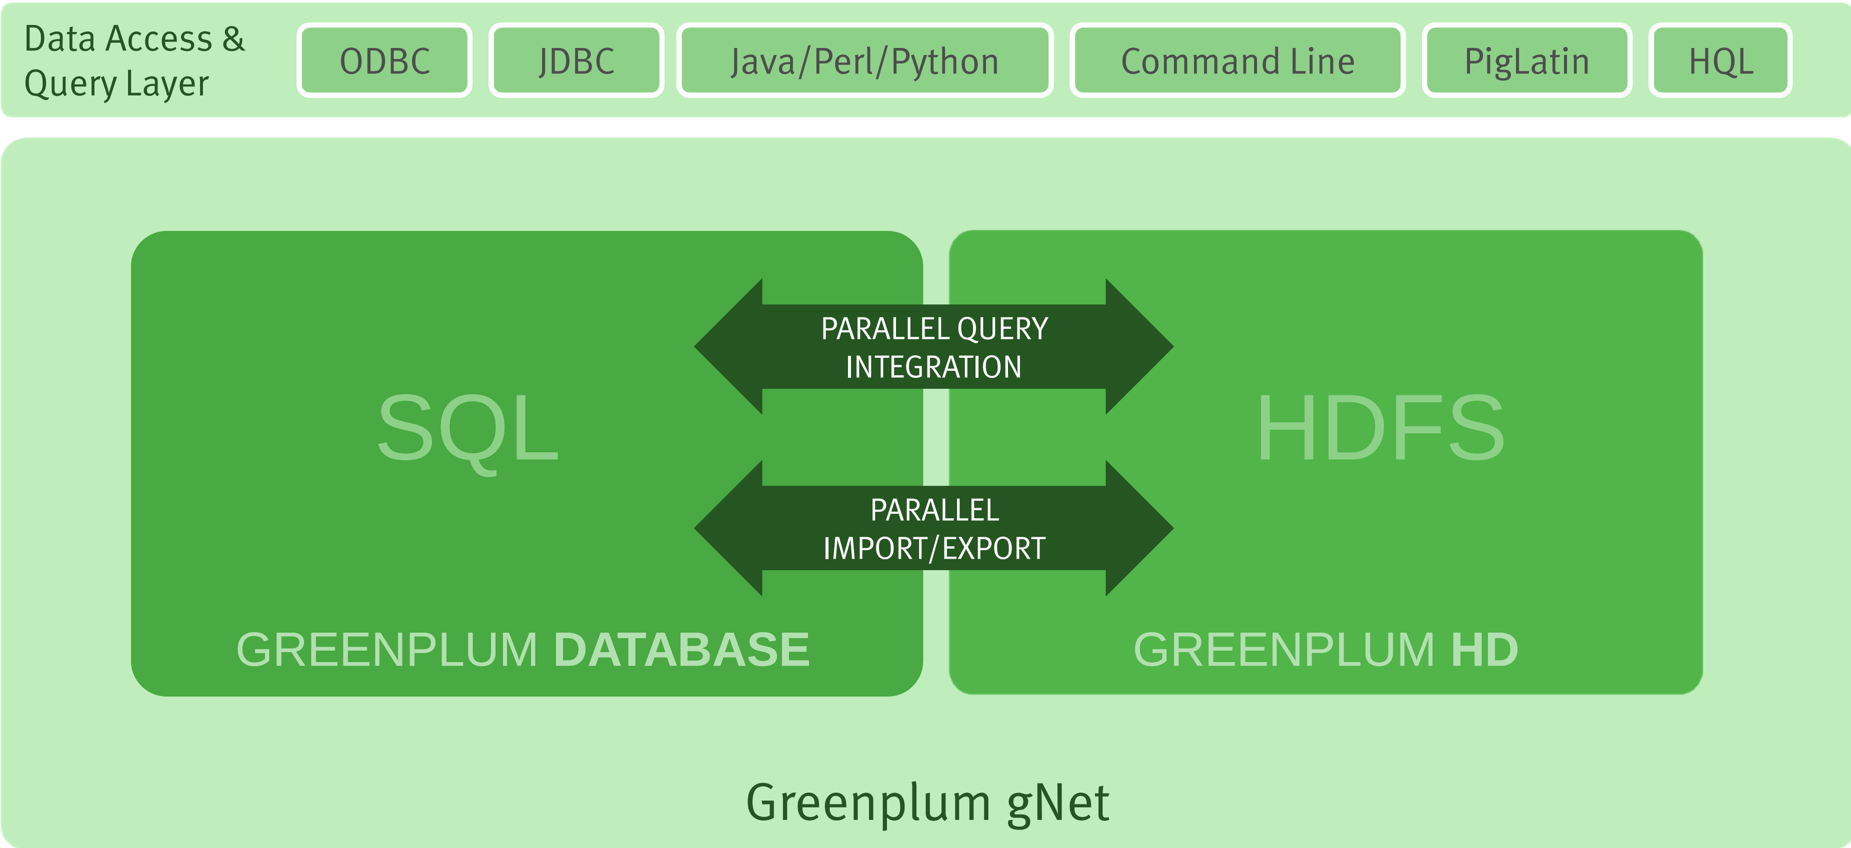

Hadoop

- A fully-compliant Hadoop implementation co-located with Greenplum in an Appliance, sharing a fast inter-connect, to provide business intelligence based on both structured and unstructured data…

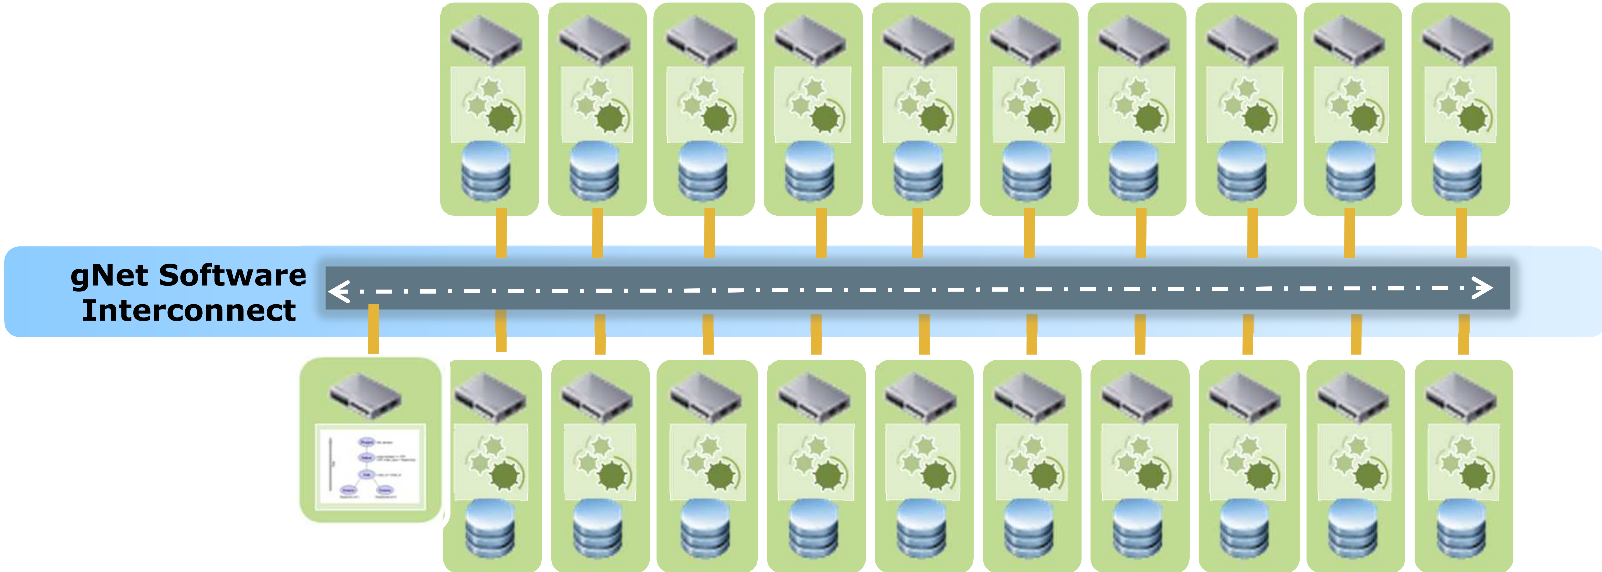

Challenges with Traditional Big Data and Hadoop Environments

- Poor utilization of storage and CPU resources in Hadoop clusters
- Inefficient data staging and loading processes
- Backup and disaster recovery missing

Hadoop Architecture

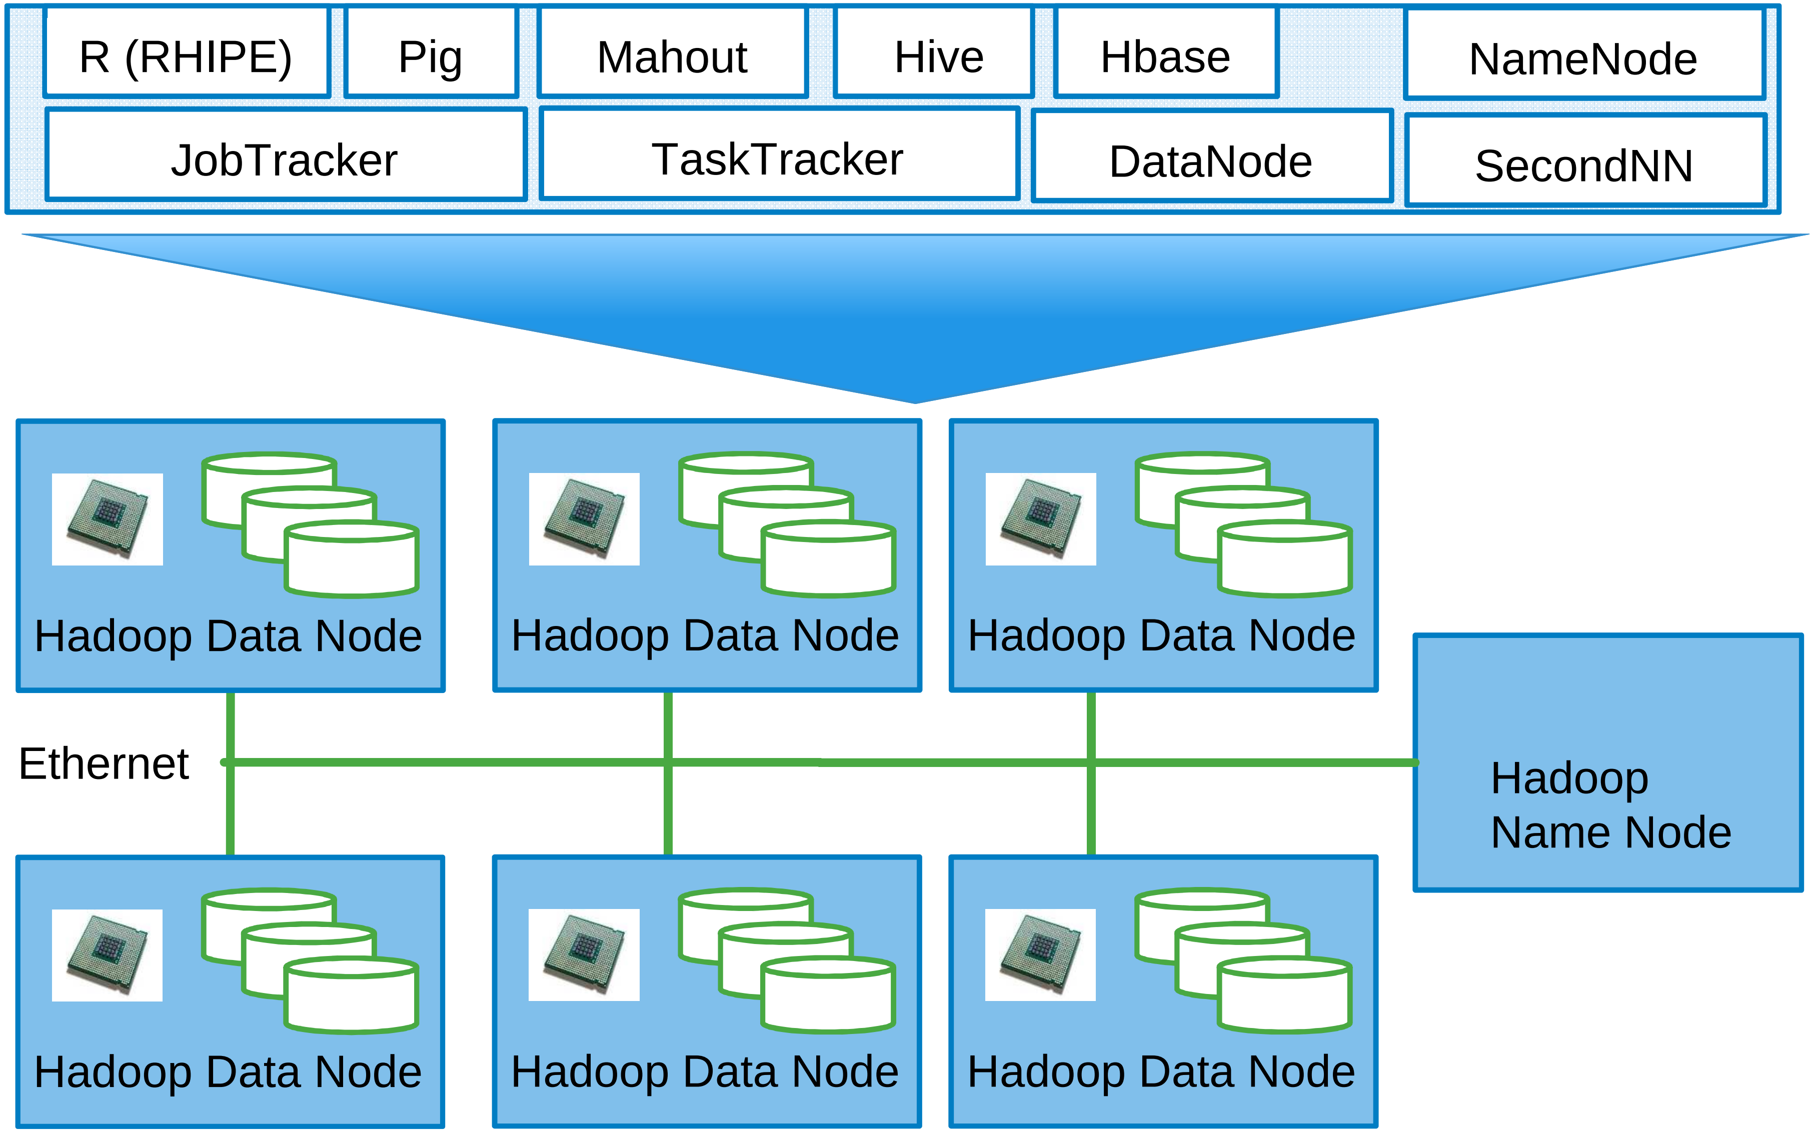

Hadoop Data Flow

1. Data is ingested into the Hadoop File System (HDFS)
2. Computation occurs inside Hadoop (MapReduce)
3. Results are exported from HDFS for use

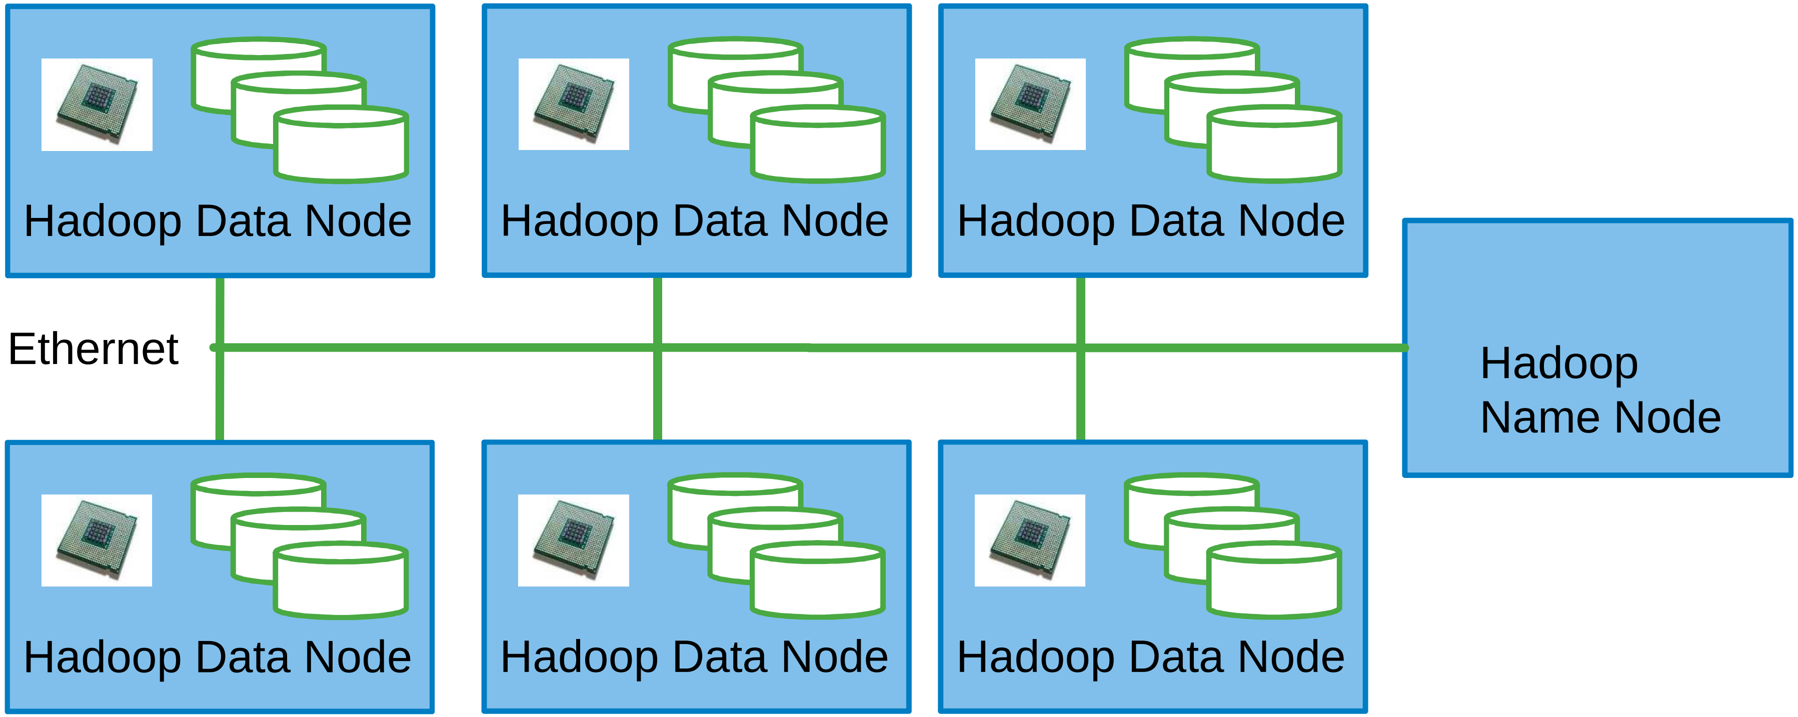

Traditional Hadoop Environment

- Servers with Direct Attached Storage (DAS)
- Data Protection: 3x mirror of all data
- Data Ingest: Tools dependant (No CIFS/NFS access)
- Scaling: Add more servers with DAS
- Single Points of Failure: NameNode
- Replication: No geographic data protection
- Data Recovery: Recreate data from other sources

Challenges with Traditional DAS approach with Hadoop

- Poor utilization of storage and CPU resources in Hadoop clusters
- Managing complexities of data and storage environments → especially DAS
- Inefficient data staging and loading processes
- Backup and disaster recovery missing
- Management @ Scale

Hadoop Architecture with Isilon

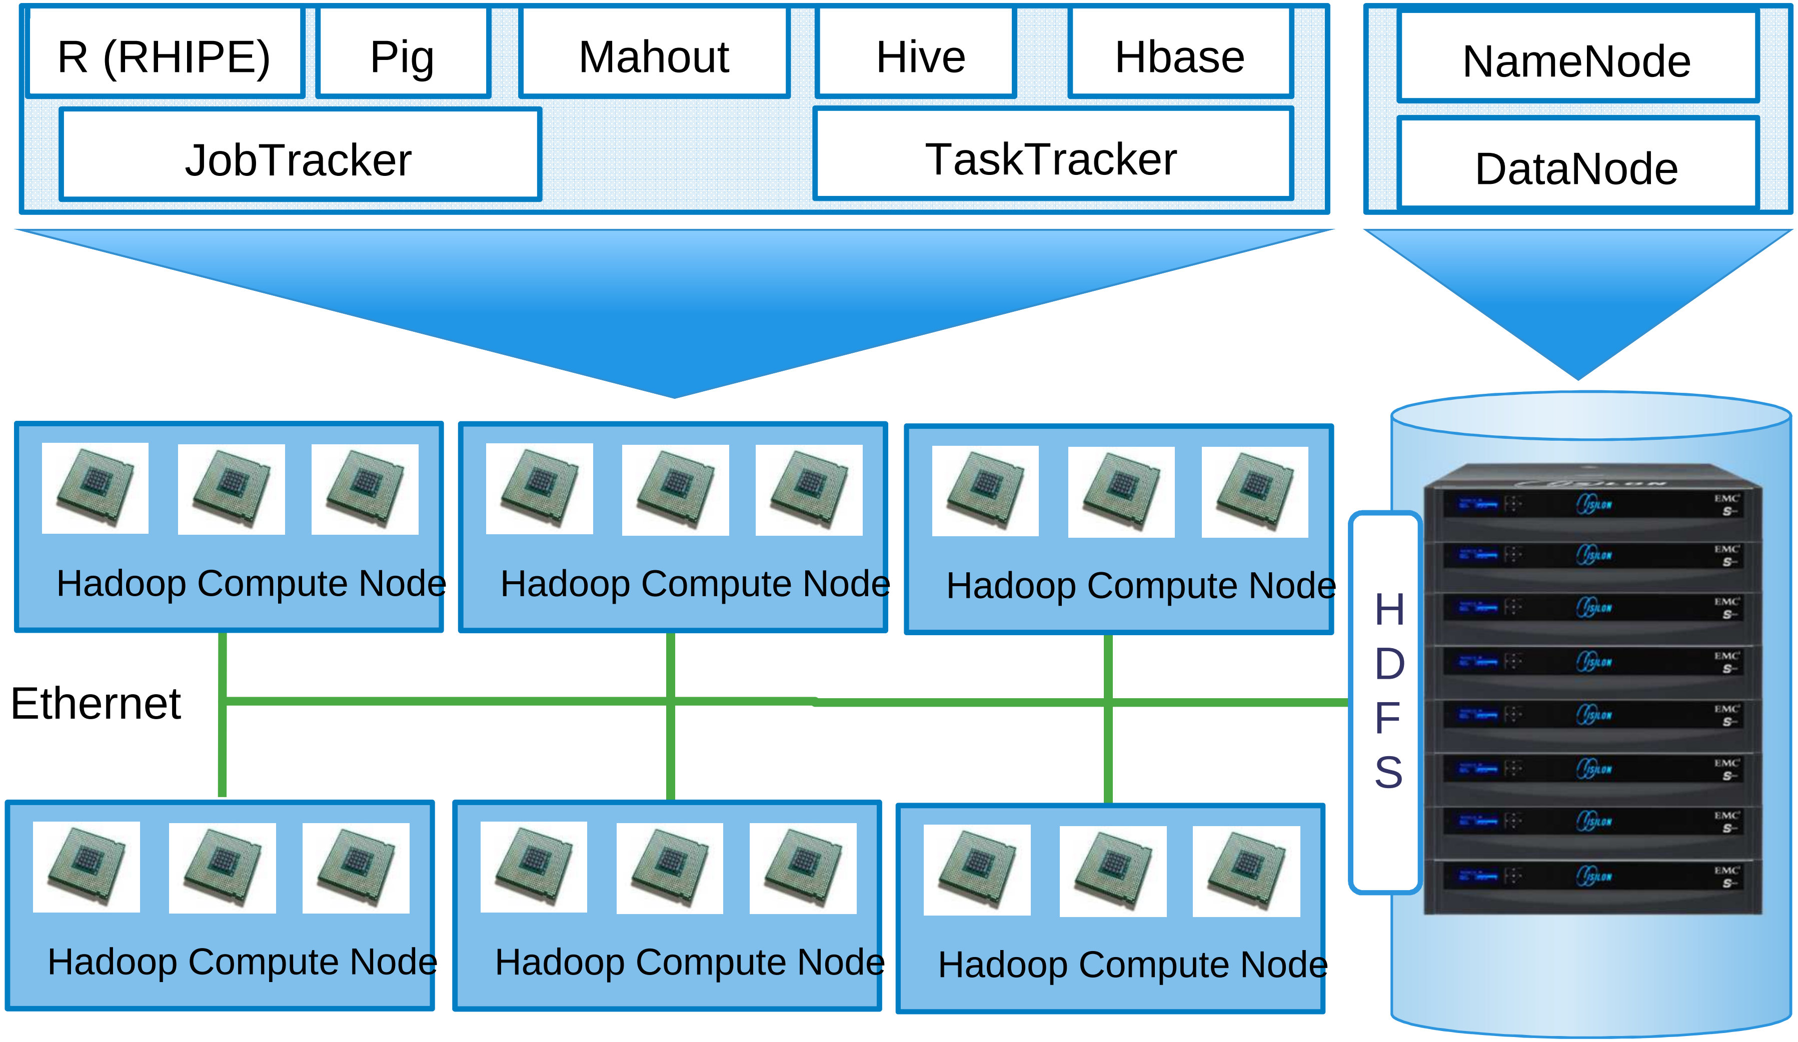

THE ANSWER MACHINE

DATA IN. DECISIONS OUT.

Introducing the
Greenplum Data Computing Appliance

Delivering the fastest data loading and best price/performance ratio in the data warehousing industry.

EMC²
where information lives®