

Greenplum: Driving the Future of Data Warehousing and Analytics

Luke Lonergan, cofounder and CTO
Greenplum Software
January 20, 2010



Greenplum: What We Do

We make a high-end

MPP Database

Architected and optimized for
data warehousing and analytics

Greenplum: How We're Doing

\$100 Million

Dollars saved by customers choosing
Greenplum over Teradata

Greenplum: How We're Doing

5 Billion

Shares analyzed daily by stock exchanges
and regulatory firms using Greenplum

Greenplum: How We're Doing

6.5 Petabytes

Size of largest database in the world, Ebay,
powered By Greenplum

Greenplum: How We're Doing

300 Million

Consumers receiving more secure and personalized services from Greenplum customers

Greenplum: How We're Doing

20 Trillion

Rows of data being managed by
Greenplum databases

Greenplum: Our Customers Include...



Our Customers Do Cool Stuff



Powered by Greenplum: Xtract OEM

Powered by Greenplum:




The Offering:

- Unique Telco-focused Customer Retention and Churn Analysis
- Used by the worlds largest Telcos: Vodafone, Verizon, Sprint, Orange, etc.
- Ease of integration and ability to support existing SQL very key; expanded with embedded analytics

The Customer Value:

- Provide insight into how subscribers are using services and effectiveness of new service introductions
- Enhance visibility to how subscribers interact, and how that interaction may affect churn
- Substantially reduce cost of internal development and support burden



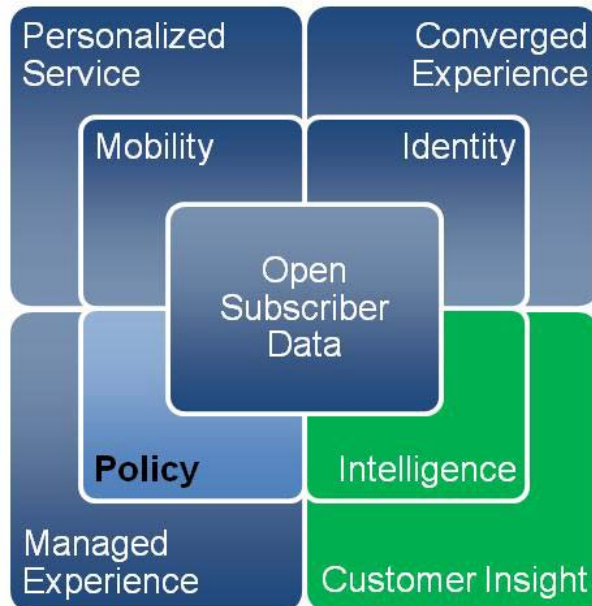
Powered by Greenplum: Nokia-Siemens Networks OEM

Powered by Greenplum:

Nokia Siemens
Networks

The Offering:

- New cornerstone offering within NSN's Subscriber Data Management line
- To date, SDM data was not analyzed due to cost and complexity
- Greenplum OEM provides core data store and analytics



The Customer Value:

- Provides framework for new class of analytic capabilities across “goldmine” of Subscriber Data for large Telco Carriers
- Understand what subscribers are doing, in what context, and across what offerings
- Substantially more insight than typical CDR or general “who talks to who” billing-centric data
- Enabling technology for more advanced Location-Based Services

Powered by Greenplum: ClickFox OEM

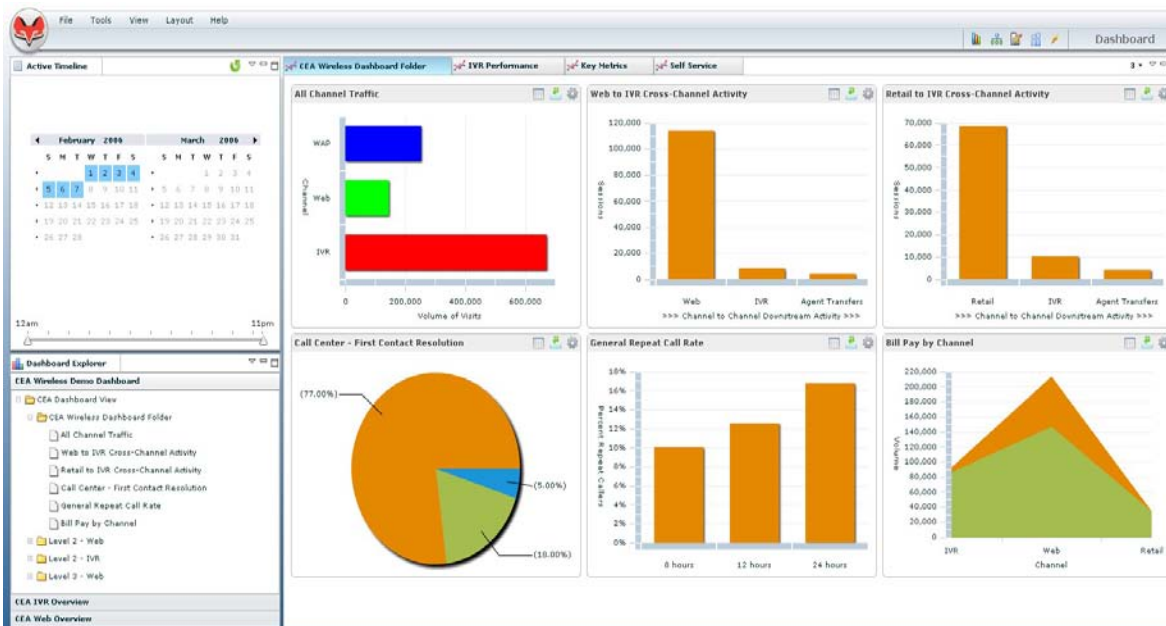


The Offering:

- Leader in Customer Experience Analytics
- GP Replacement to Oracle back-end for SAAS/Hosted and on premise Clickfox customers
- Requirement to scale larger, provide answers faster, and more easily deploy than alternatives

The Customer Value:

- Provide deeper analytics on how end-customers are interacting with the call center
- Ensure companies can maintain a technology and business intelligence advantage as the volume and complexity of customer data continues to rise
- Reduce cost and complexity of SAAS offering; lower DBA support in on-premise installations



Greenplum: What's New

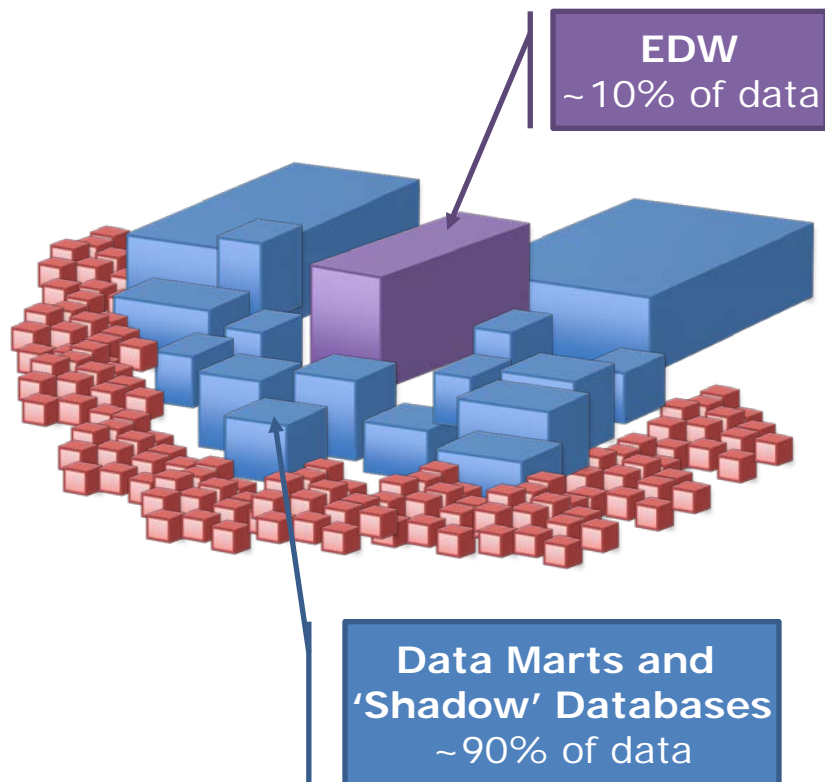


Enterprise Data Cloud™ Initiative

“Greenplum cloud plan promises a new era for data warehousing. Flexible private-cloud scalability meets massively parallel processing.”

Doug Henschen,
Editor-in-Chief,
Intelligent Enterprise

State of Play – Data in a Typical Enterprise



- Data is everywhere – corporate EDW, 100s of data marts, ‘shadow’ databases and spreadsheets
- The goal of centralizing all data in a single EDW has proven untenable
- The EDW plays an important role
 - Top-down control for the most regimented business data
 - Heavy modeling and process is essential
- What about the other 90% of the data?
 - Bottom-up generation and access
 - Usage patterns change quickly – can’t wait for EDW processes
 - Had lead to silos and shadow IT

Enterprise Data Cloud (EDC) In Action

The screenshot displays the EDC dashboard with the following sections:

- Infrastructure Pools:** A map of the United States with callouts for CA, MN, and TX. A callout for TX shows 'Dallas_Datacenter' with location 'Dallas, TX', subpools '7', and instances '3'.
- Tasks (4):** A list of tasks including 'Data Transfer to my DB/Inventory' (Scheduled 25 Sep, 2009), 'Data Transfer to test DB/Inventory' (In Progress, 30%), 'Data Transfer to Stock/Inventory' (In Progress, 70%), and 'Data Transfer to Sales/Projections' (Completed).
- EDC Statistics:** Three charts: 'Server Up Time Statistics' (127 GB loaded last 7 days, 489 TB of total data), 'Database Usage' (127 GB loaded last 7 days, 489 TB of total data), and 'Resource Usage' (127 GB loaded last 7 days, 489 TB of total data).
- Table:** A table with columns: Name, Subpool, Server Profile, Total, Allocated, Available. The first row shows 'THX' with 'Default' subpool and 'HPDL1806G' server profile.

Provision, manage, and virtualize database infrastructure

The screenshot displays the EDC dashboard with the following sections:

- Share your comments or insights...:** A text input field with a 'Share' button.
- Recent Activity:** A list of activity items including 'Scott published data object "Customer"', 'Steve created data perspective "FINA Sales"', 'Martin A. shared 5 data objects', and 'Steve: I can't find product information anywhere - does anyone have a clue?'.
- EDC Statistics:** A summary of system metrics: '127 GB loaded last 7 days, 489 TB of total data', '314 Perspectives, 12,349 Shared objects', '249 Databases, 203 Instances', '179 users', and '5 Infrastructure Pools'.
- Popular Picks:** A list of popular data objects such as 'Sales Info', 'cust_rate', 'event_log', 'Products_details', and 'Customer Data'.

Discover, publish and share new data and insight

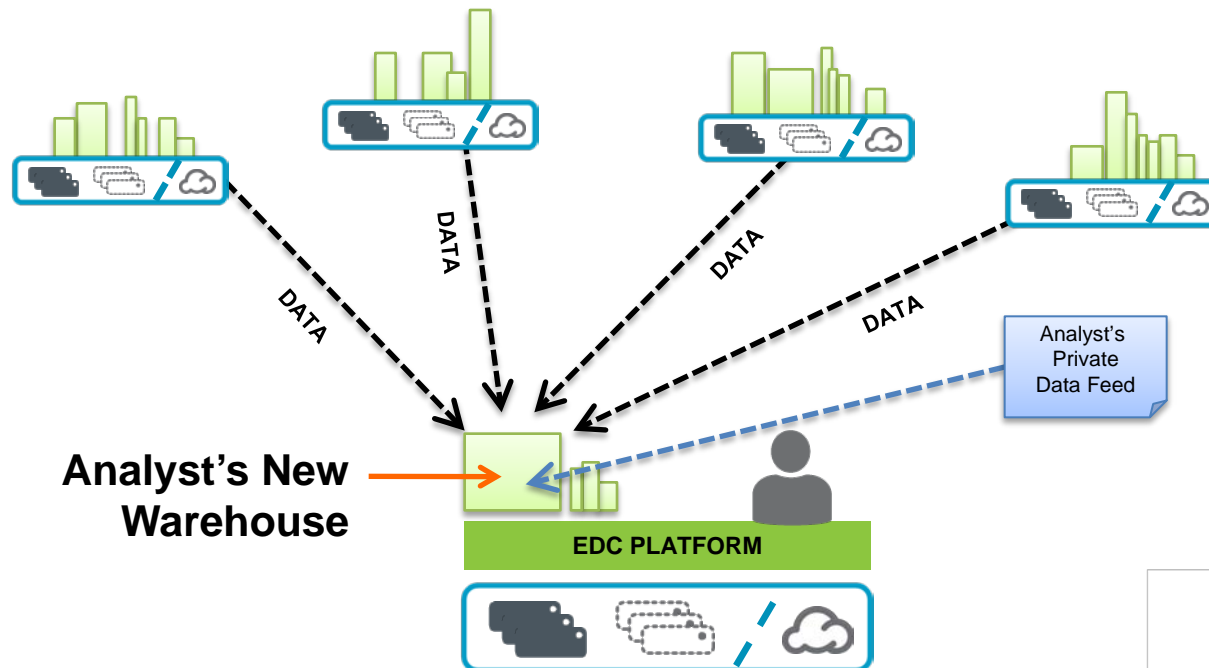
Large Global Telco

Goals:

- Remove IT barriers to analyst productivity and value creation
- Dramatically reduce IT resource constraints and delays – i.e. realize ideas sooner
- Combine centralized 'EDW' data with freshly discovered feeds and other useful sources

Approach:

- Self-serve creation of project warehouses in minutes – and elastically expand as needed
- Load new data feeds without requiring formal modeling
- Implement on full VCE stack at 1.3PB of usable storage




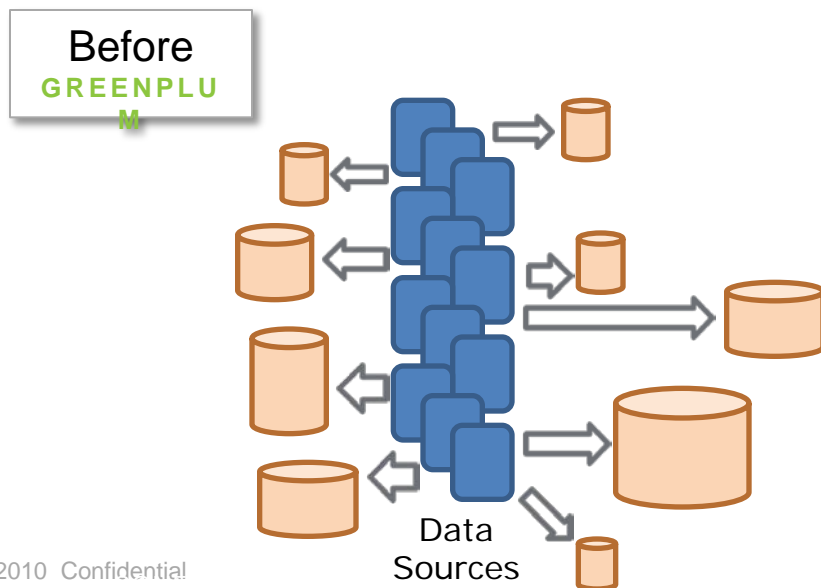
Large Global Bank

Goals:

- Transform DB IT into a virtualized, self-service infrastructure platform
- Dramatically lower costs, and increase agility to support business requirements
- Break down silo walls - provide a unified way to find and access all data

Approach:

- Provide 'self serve' model to bring shadow IT into the light
- Enable database as a service – and encourage deployment of data marts and DWs onto DB's private cloud
- Allow unified data access and pragmatic 'logical' data model unification incrementally



The Greenplum Generation Begins



Enterprise Data Cloud™ Initiative

The new architecture and methodology for modern enterprise data warehousing and analytics directly addressing customers' business issues and opportunities.

Enterprise Data Cloud: The Future of Data Warehousing and Analytics

1

Extreme Scale and Elastic Expansion

2

Self-Service Provisioning

3

Unified Data Access

Enterprise Data Cloud: The Future of Data Warehousing and Analytics

1

Extreme Scale and Elastic Expansion

The ability to handle massive scale and dynamic change of data volume, leveraging commodity hardware

Enterprise Data Cloud: The Future of Data Warehousing and Analytics

2

Self-Service Provisioning

Providing business analysts the ability to create
data marts and warehouses in minutes

Enterprise Data Cloud: The Future of Data Warehousing and Analytics

3

Unified Data Access

Enabling ease and efficiency in the enterprise-wide publishing, discovering, and sharing of data

Thank you

- Luke Lonergan
llonergan@greenplum.com