




Netezza 101

an introduction to the data warehouse appliance

Session Number IDW-3149B

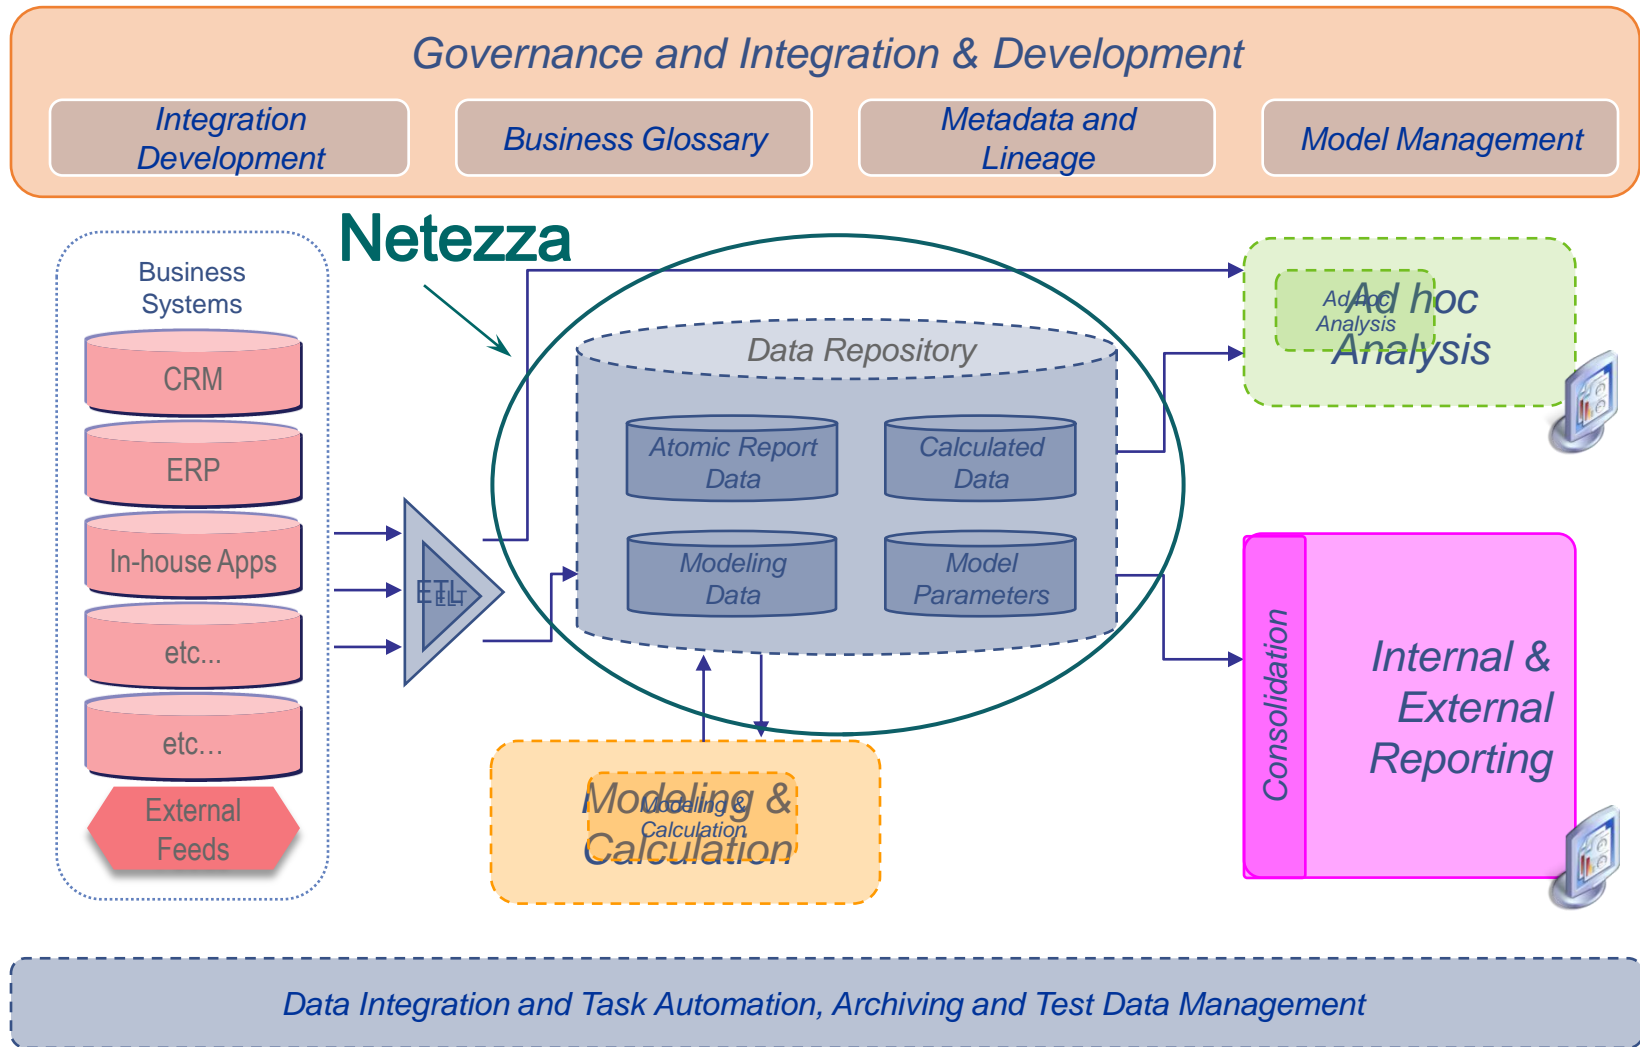
dai clegg, IBM



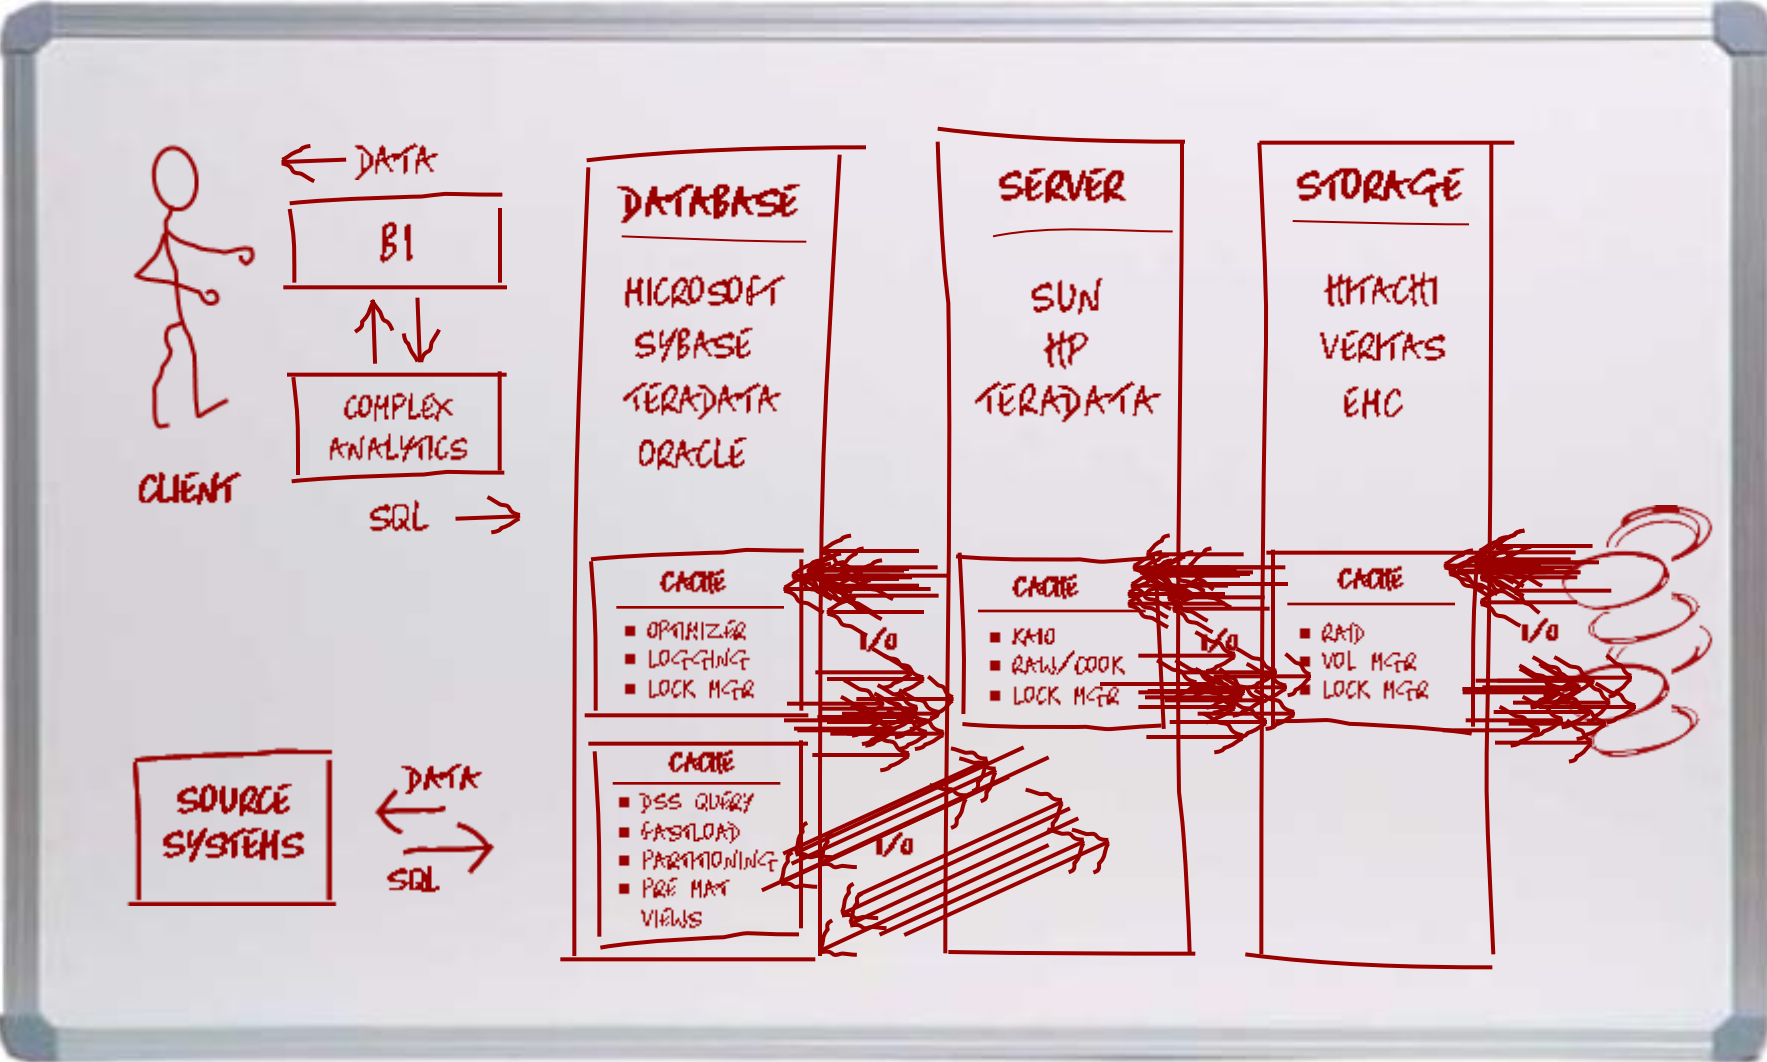
IBM Software

Information On Demand **2011**


The Datawarehouse Solution Space



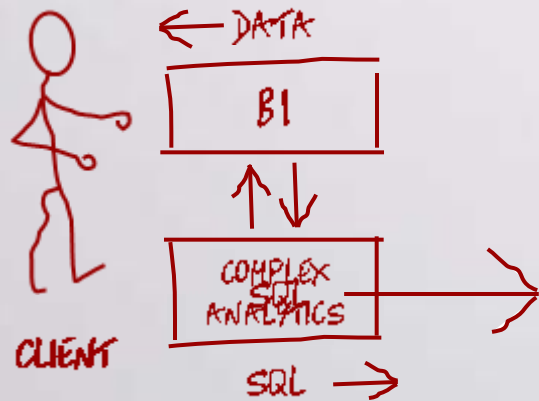
Big data overwhelms traditional data warehouses



Let's simplify this mess ...



... and bring analytics into the warehouse



LESS IS MORE

The IBM Netezza Appliance: Revolutionizing Analytics

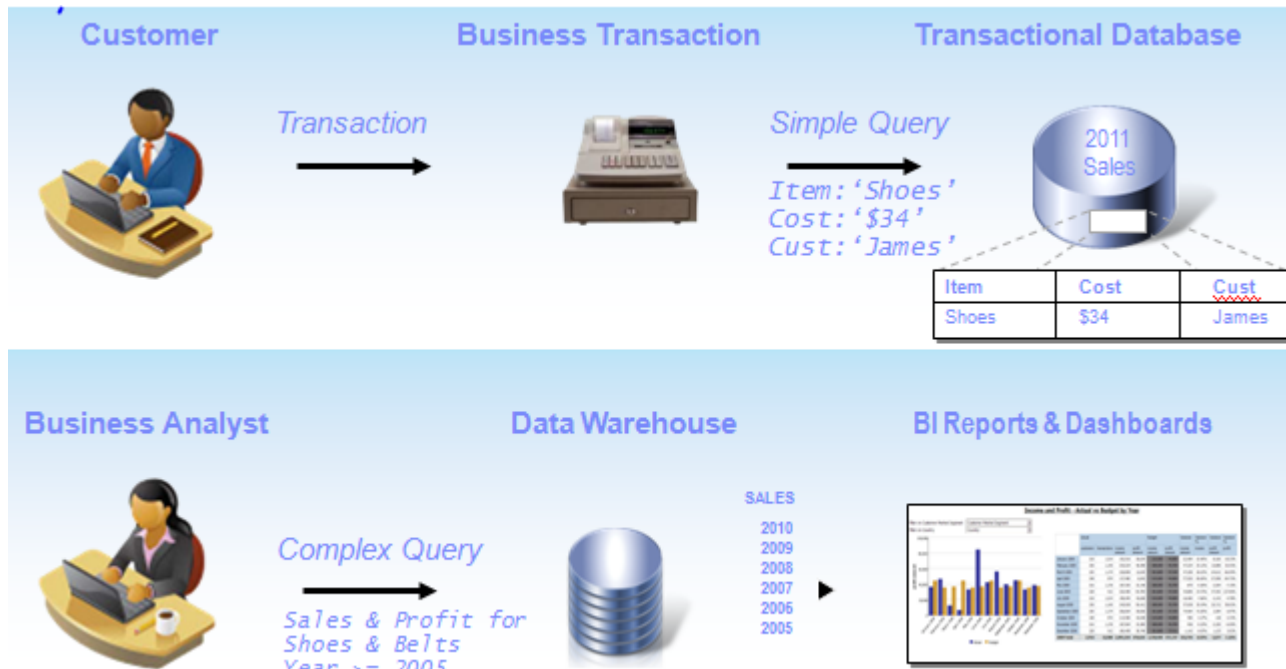


- Purpose-built analytics engine
- Integrated database, server & storage
- Standard interfaces
- Low total cost of ownership

- Speed: 10-100x faster than traditional systems
- Simplicity: Minimal administration
- Scalability: Peta-scale user data capacity
- Smart: High-performance advanced analytics

Transactional workloads vs. analytic workloads

Two very different requirements for storing and processing data



The Netezza Appliance – Loading

Data Integration

IBM Information Server
Ab Initio
Business Objects/SAP
Composite Software
Expressor Software
GoldenGate Software (Oracle)
Informatica
Sunopsis (Oracle)
WisdomForce

Data In

SQL ODBC JDBC OLE-DB



The Netezza Appliance – Querying



Digital Media



Financial Services



Government



Health & Life Sciences



Retail / Consumer Products



Telecom



Other



Netezza delivers speed

- 15,000 users
- Running 800,000+ queries per day



“

...when something took 24 hours I could only do so much with it, but when something takes 10 seconds, I may be able to completely rethink the business process ...

”



-- SVP Application Development, Nielsen

http://www.youtube.com/watch?v=yOwnX14nLrE&feature=player_embedded

Netezza delivers scalability

- 1 PB on IBM Netezza
- 7 years of historical data
- 100-200% annual data growth



“

*NYSE ... has replaced an Oracle 10 relational database with a data warehousing appliance from Netezza, allowing it to **conduct rapid searches of 650 terabytes** of data.*

”



-- ComputerWeekly.com

Netezza delivers **smarts**

- Identifies items that shoppers are likely to buy in future visits
- Coupon redemption rates as high as 24%



“

Because of (Netezza's) in-database technology, we believe we'll be able to do 600 predictive models per year (10X as many as before) with the same staff.

”



-- Eric Williams,
CIO & Executive VP

Catalina Marketing:

building loyalty one customer at a time

- Identifies its customers into segments
- Marketing to a segment of one – 195 million US loyalty program members
 - Every coupon printed is unique to the individual customer
 - Customized based on three years' worth of purchase history
- Increased efficiency – from 4 hours to 60 seconds to score a model

		Basic targeting e.g., offer dog food coupon to customer buying dog food	Using predictive models to find latent correlations
Coupon redemption rate	1%	6-10%	25%

Because of (Netezza's) in-database technology, we believe we'll be able to do 600 predictive models per year (10X as many as before) with the same staff.

”



-- Eric Williams,
CIO & Executive VP

Netezza delivers **simplicity**

- Up and running 6 months before being trained
- 200X faster than Oracle system
- ROI in less than 3 months



“

Allowing the business users access to the Netezza box was what sold it.

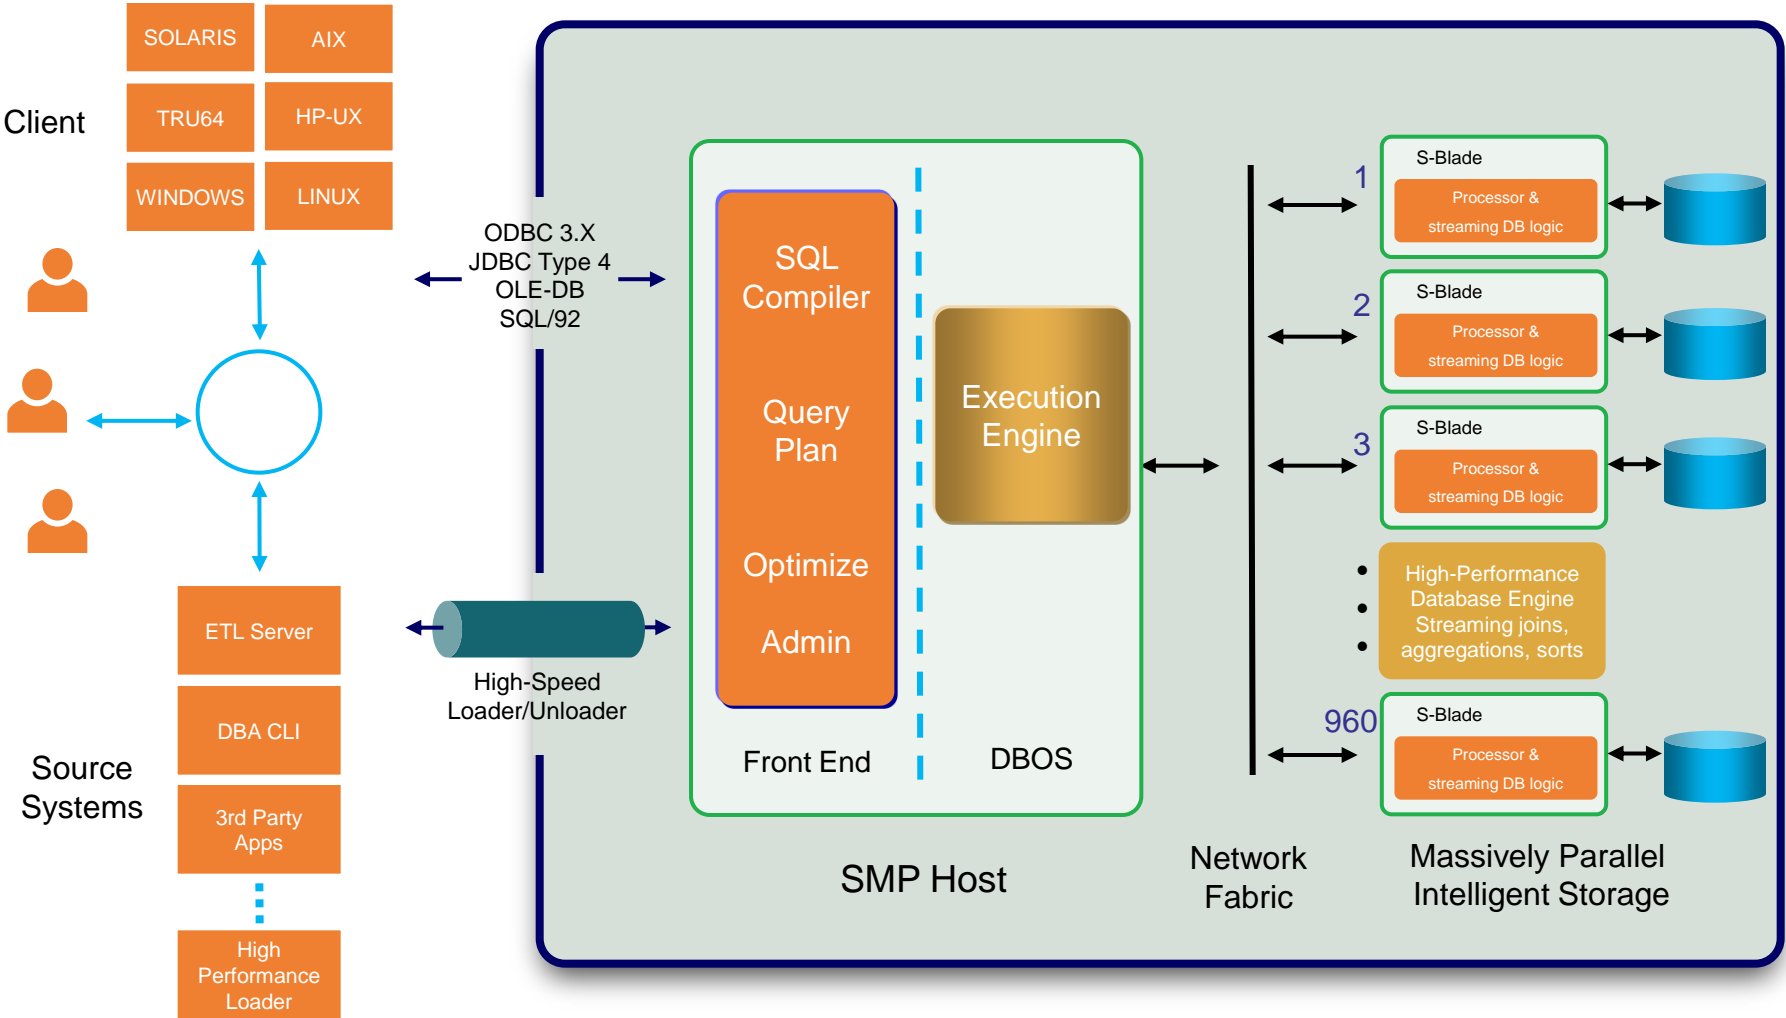
”



-- Steve Taff,
Executive Dir. of IT Services

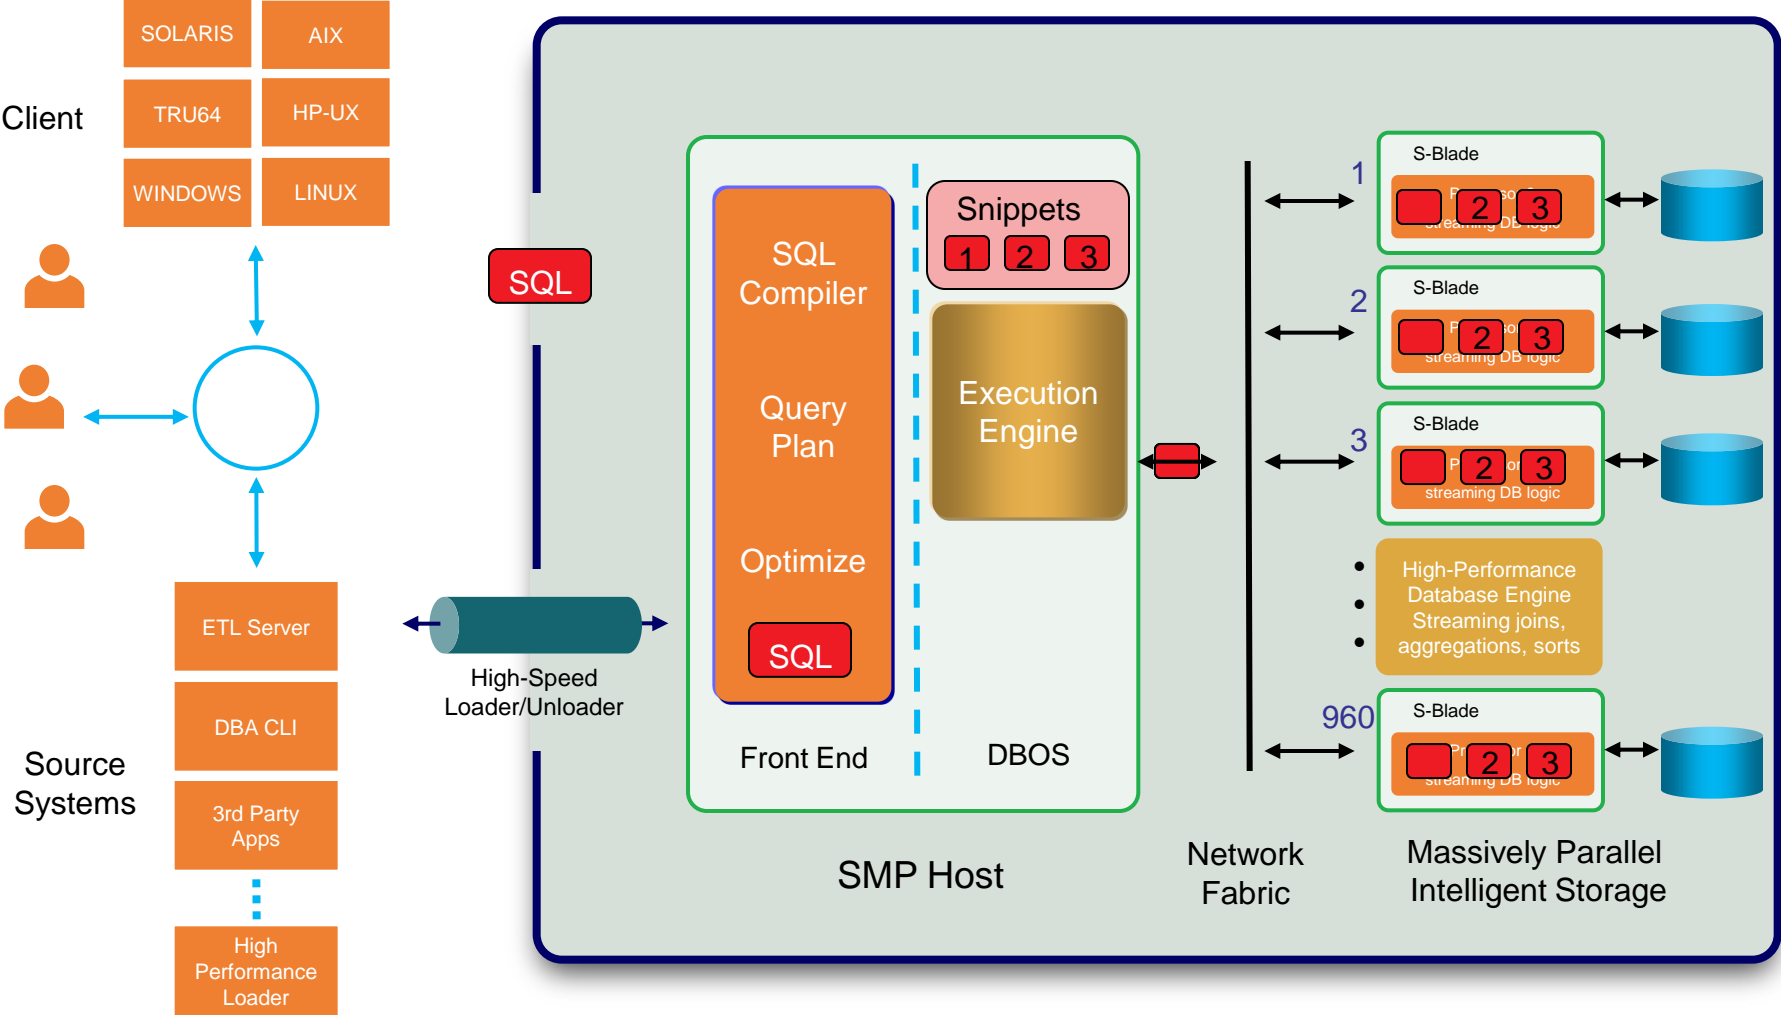
IBM Netezza:

True Appliance, Massively-Parallel Processing




IBM Netezza:

True Appliance, Massively-Parallel Processing

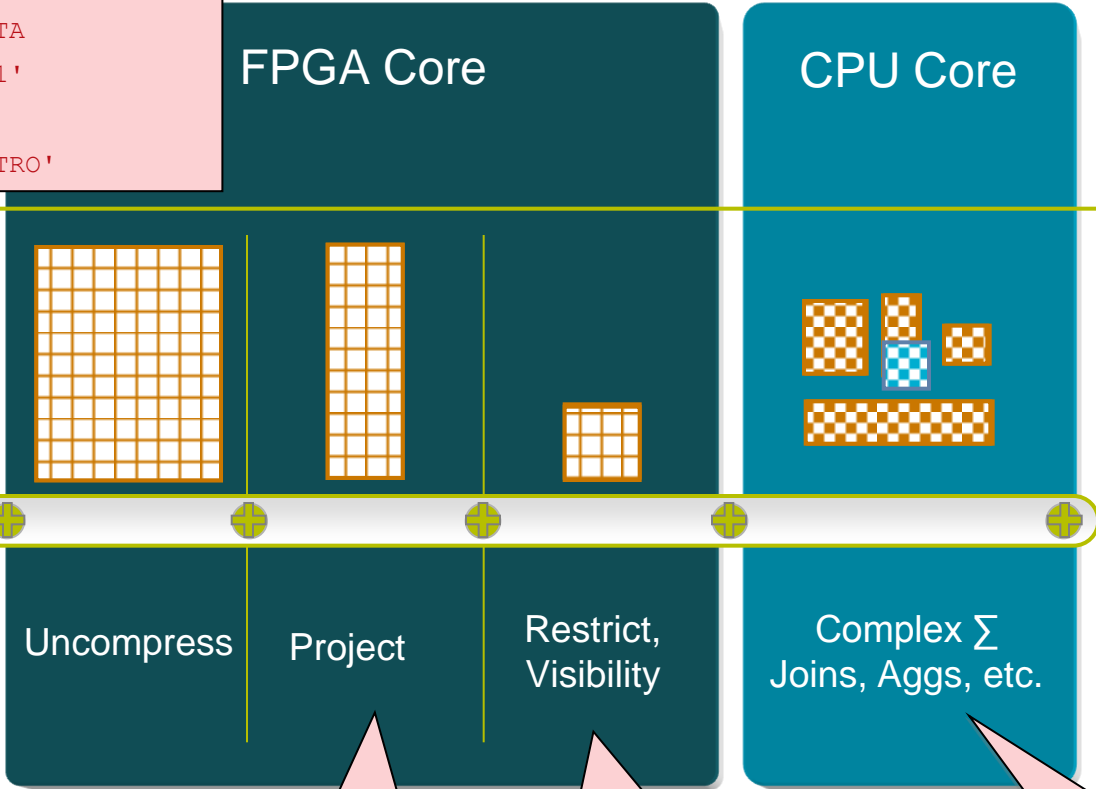


Our Secret Sauce

```
select DISTRICT,
       PRODUCTGRP,
       sum(NRX)
from   MTHLY_RX_TERR_DATA
where  MONTH = '20091201'
and    MARKET = 509123
and    SPECIALTY = 'GASTRO'
```



Slice of table
MTHLY_RX_TERR_DATA
(compressed)



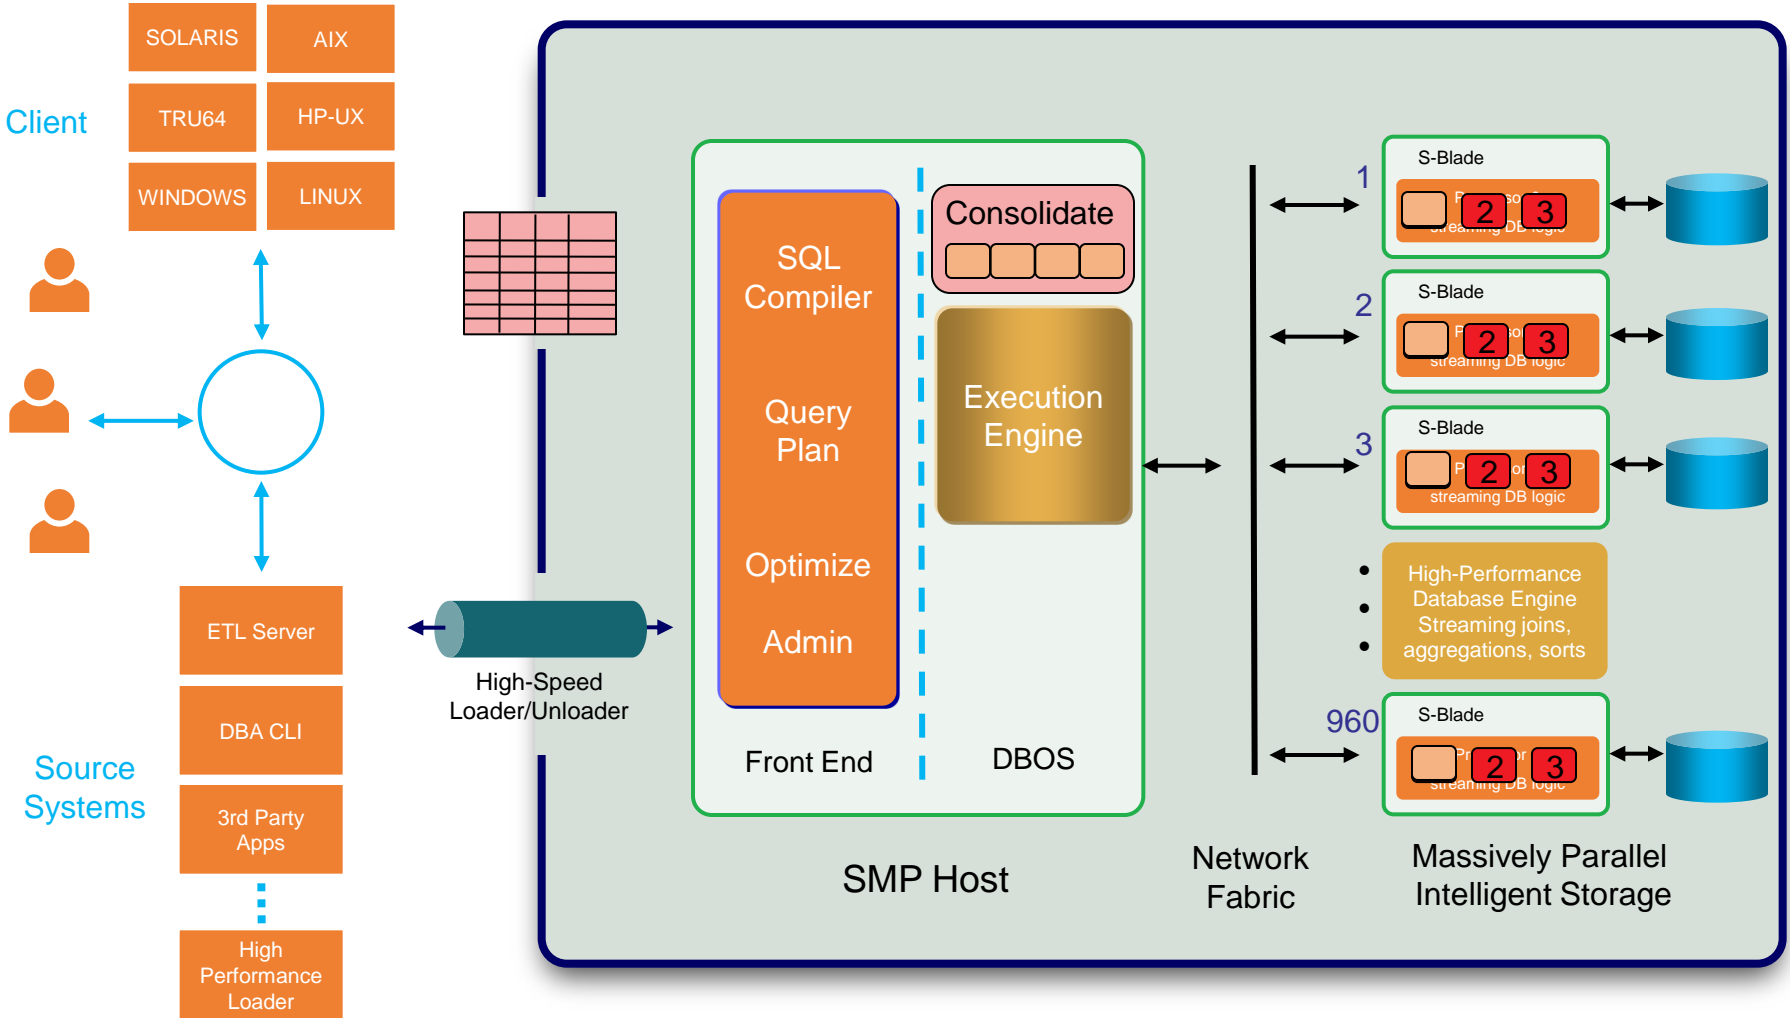
```
select DISTRICT,
       PRODUCTGRP,
       sum(NRX)
```

```
where MONTH = '20091201'
and    MARKET = 509123
and    SPECIALTY = 'GASTRO'
```

```
sum(NRX)
```

IBM Netezza:

True Appliance, Massively-Parallel Processing



IBM Netezza data warehouse appliance family

for data life-cycle management



Development & Testing



DWH, High Performance Analytics



Queryable Archive. Backup/DR

From a few terabytes to 10s of petabytes



The IBM Netezza Appliance: Revolutionizing Analytics



Digital Media	
Financial Services	
Government	
Health & Life Sciences	
Retail / Consumer Products	
Telecom	
Other	

- Speed: 10-
- Simplicity: 1
- Scalability: Peta-scale user data capacity
- Smart: High-performance advanced analytics



Simplicity Drives Performance
Visit us at Booth 509