


# InfoPlanet

Netezza TwinFin – Leader in  
Data Warehouse Appliance



**Alessandro Pirovano,**  
*Netezza Sales Engineer, Netezza*




# Agenda



Business Perspective  
Architecture  
Enterprise Integration  
Appliance Simplicity  
Summary



# Agenda





## Business Perspective

Architecture



Enterprise Integration

Appliance Simplicity

Summary



**Would You Still Use Google  
If It Took 3 Days And 7 People  
To Get A Search Result?**



days for a single query

constant tuning

“ **Nearly 70%** of data warehouses experience performance-constrained issues of various types. ”

- Gartner 2010 Magic Quadrant


specialized resources required


months to deploy

Data continues to  
**expand exponentially.**

The right data warehouse  
**is now mission critical.**

Analytics are becoming more complex as  
**business demands faster answers.**





## Traditional data warehouses are just too complex

They are based on databases optimized for transaction processing—**NOT** to meet the demands of advanced analytics on big data.

---

- Too complex an infrastructure
- Too complicated to deploy
- Too much tuning required
- Too inefficient at analytics
- Too many people needed to maintain
- Too costly to operate

Too long to get answers





Netezza's revolutionary approach

## The Appliance

Simpler, faster, more accessible analytics



“ This is what Netezza has done in the data warehousing market: It has **totally changed the way we think about data warehousing.** ”

- Philip Howard, Bloor Research

InfoPlanet





# IBM Netezza TwinFin References



Digital Media



Financial Services



Government



Health & Life Sciences



Retail / Consumer Products



Telecom



Other



92% referencable

67% repeat business



# IBM Netezza TwinFin Client Derive Value Across Industries

## Speed

15,000 users running  
800,000+ queries per day  
50X faster than before

*"...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



## Scalability

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

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

ComputerWeekly.com



## Simplicity



Delivers out-of-the-box performance with **no tuning**

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

Steve Taff,  
Executive Dir. of IT Services



## Smart



Identifies items that shoppers are likely to buy in **future visits**

*"Using results derived from Netezza system, Catalina Marketing is able to give shoppers color coupons at point of sale for items they would like want to buy in future visits."*

Editorial Director,  
DM Review



# Appliance



- Dedicated Device
- Optimized for purpose
- Complete solution
- Fast installation
- Very easy operation
- Standard interfaces
- Low cost




**twinFin™**




- Dedicated Device
- Optimized for purpose
- Complete solution
- Fast installation
- Very easy operation
- Standard interfaces
- Low cost

- Appliance Simplicity
- Integrated database, server & storage
- Lowest TCO
- 10-100x faster than traditional systems
- Peta-scale data capacity
- Advanced analytics
- Standard interfaces



# Agenda





Business Perspective

**Architecture**

Enterprise Integration

Appliance Simplicity

Summary





# Massively Parallel Processing Architecture

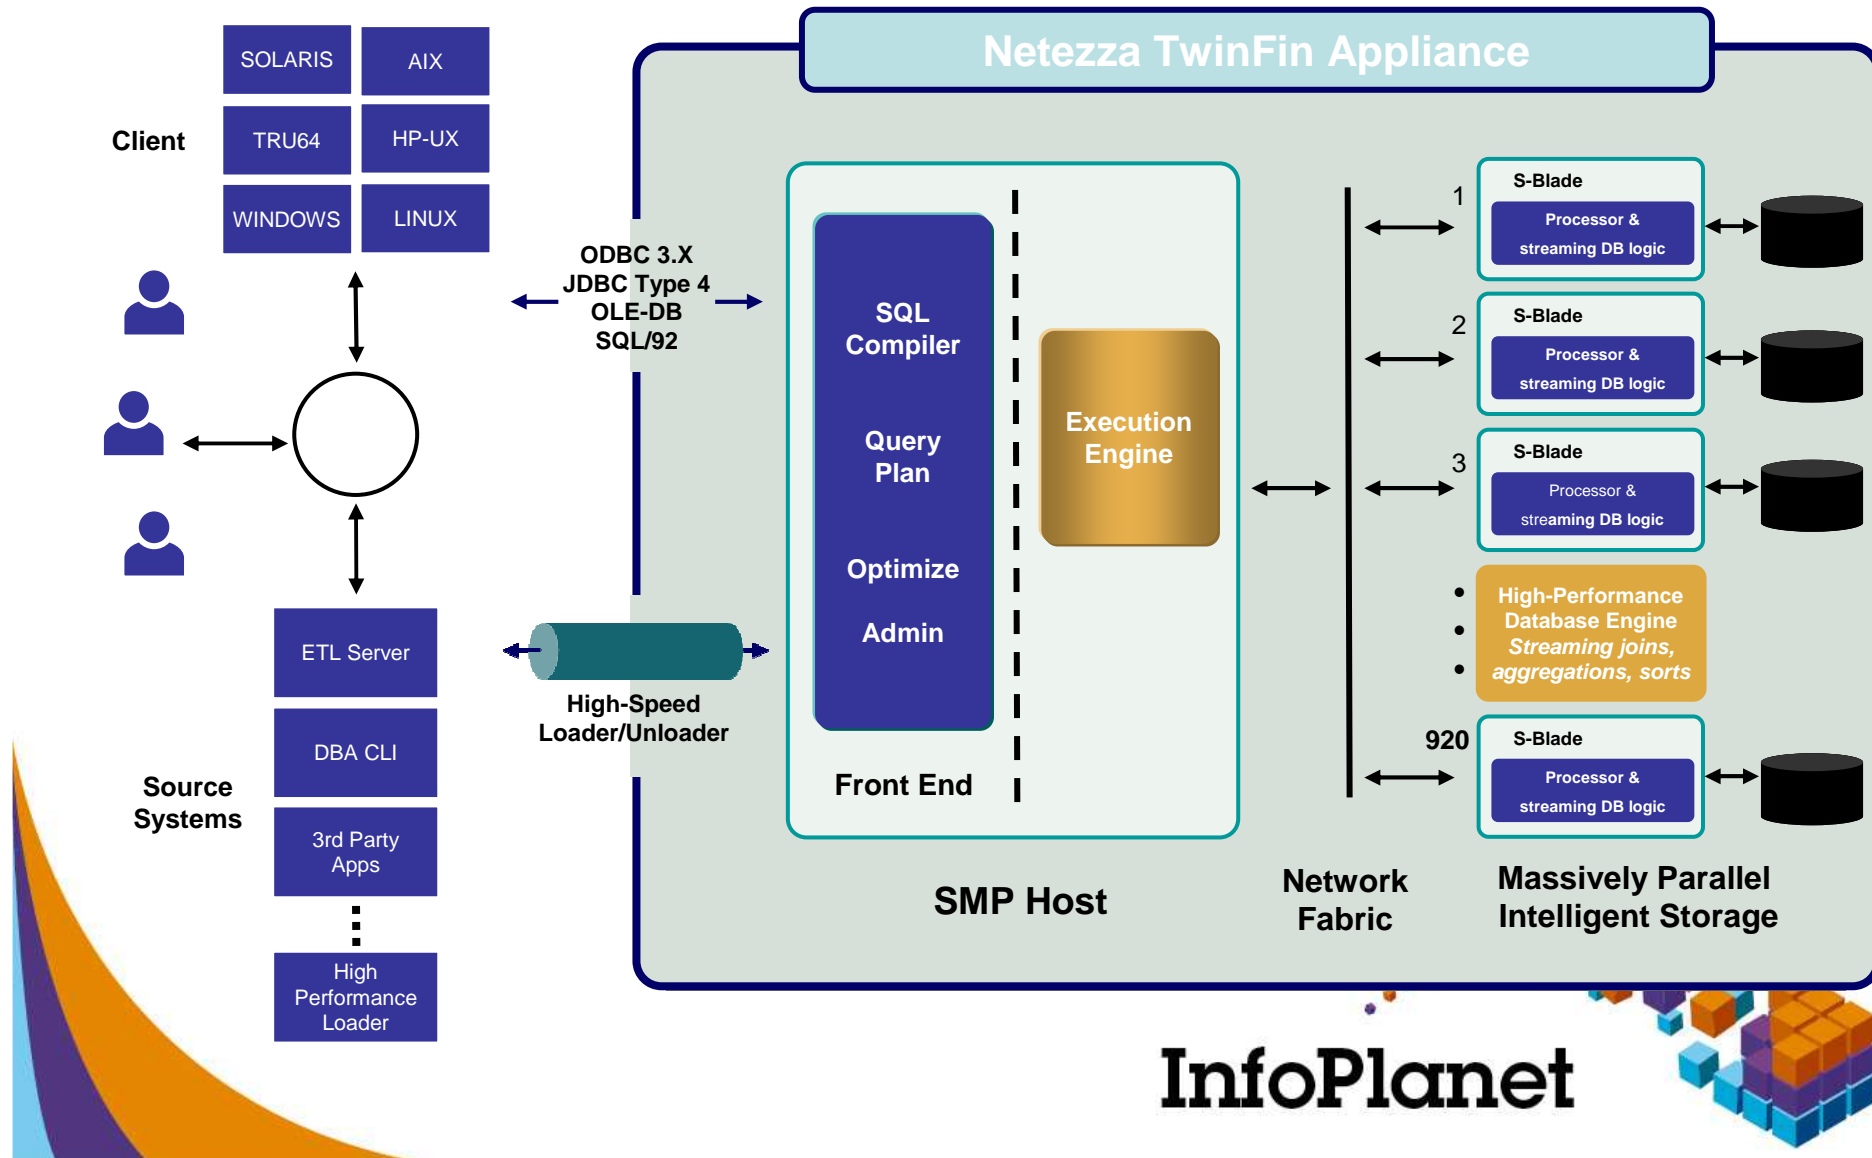

*“Divide and conquer”*



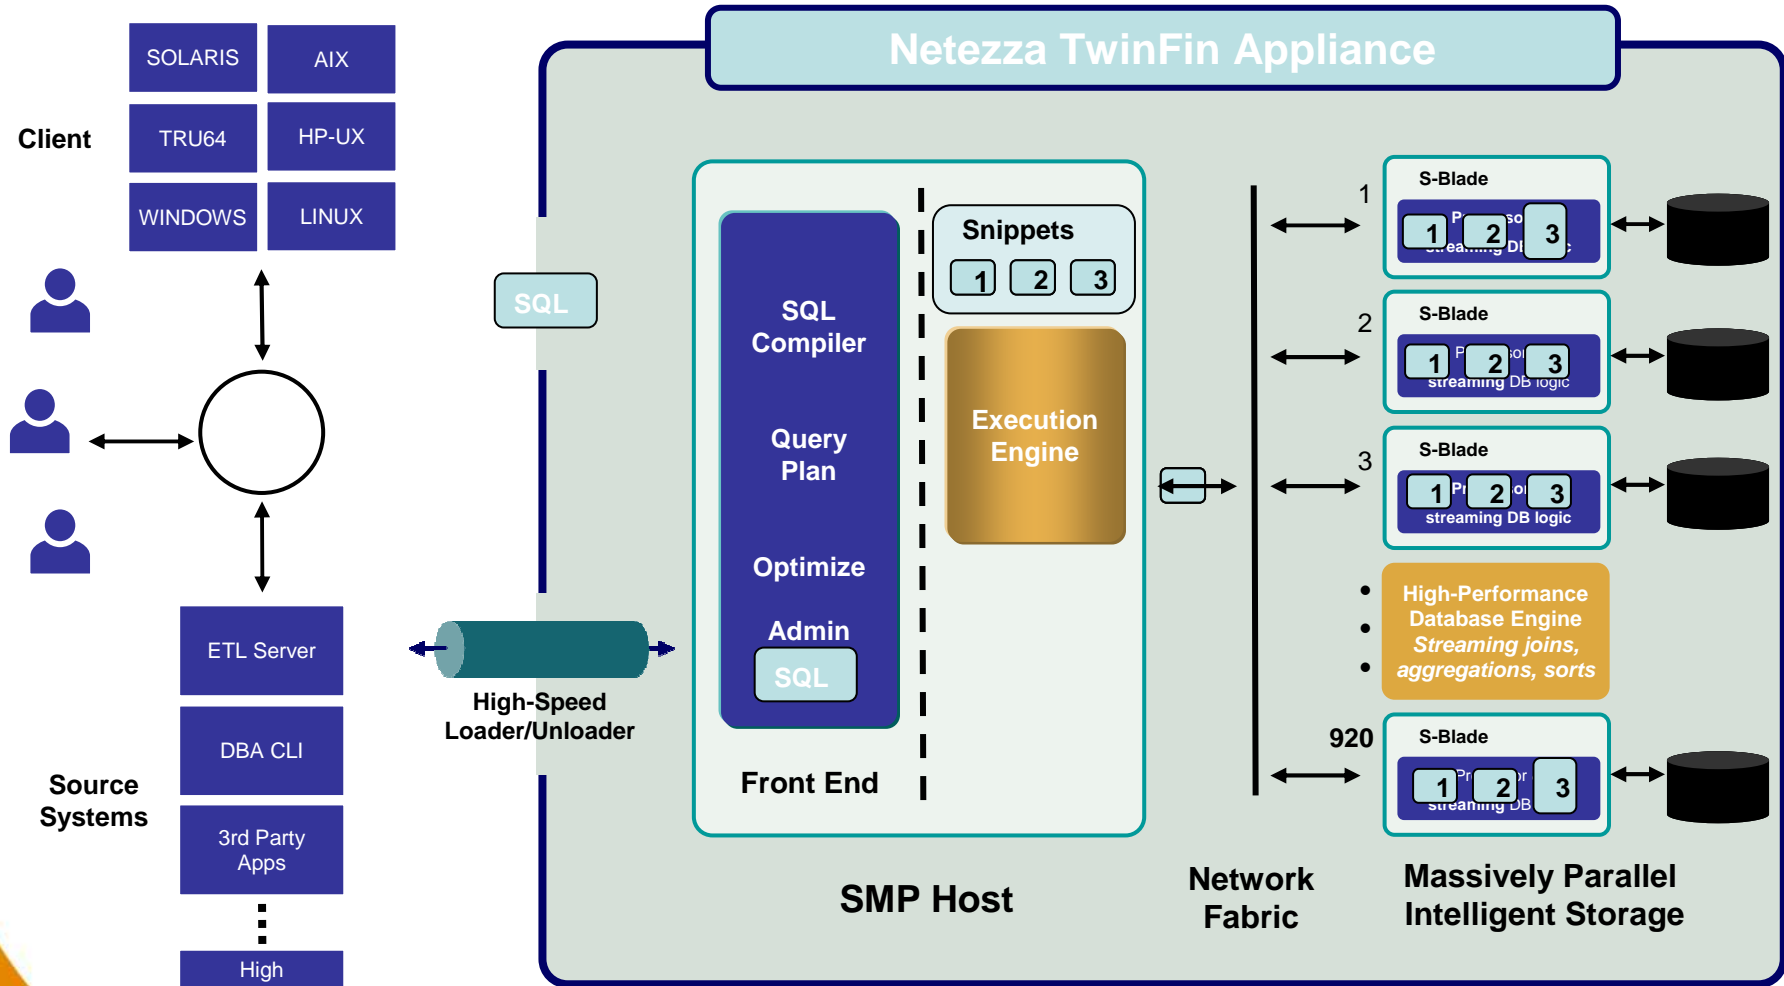
- MPP
  - “Shared Nothing” concept
  - Divides the work in smaller tasks
    - A big task is sliced vertically into a series of smaller tasks
    - The smaller tasks run independently
    - The work is automatically balanced among the tasks to minimize the time to complete
    - Each task is assigner the same amount of physical resources
    - Communication between is made only at the beginning and end of the task
- Benefits
  - A large task completes in a short elapsed time
  - Maximizes use of resources
- Points of Attention
  - Complexity on administration and management
  - Communication bottlenecks



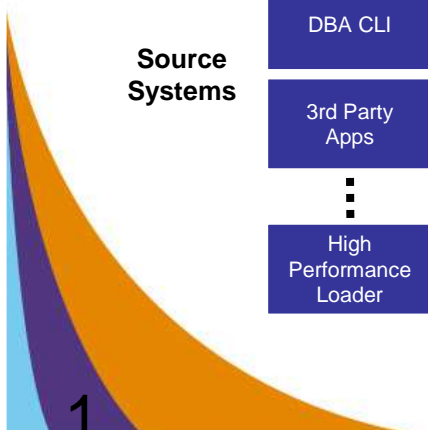
# Asymmetric Massively Parallel Processing™



# Asymmetric Massively Parallel Processing™




InfoPlanet






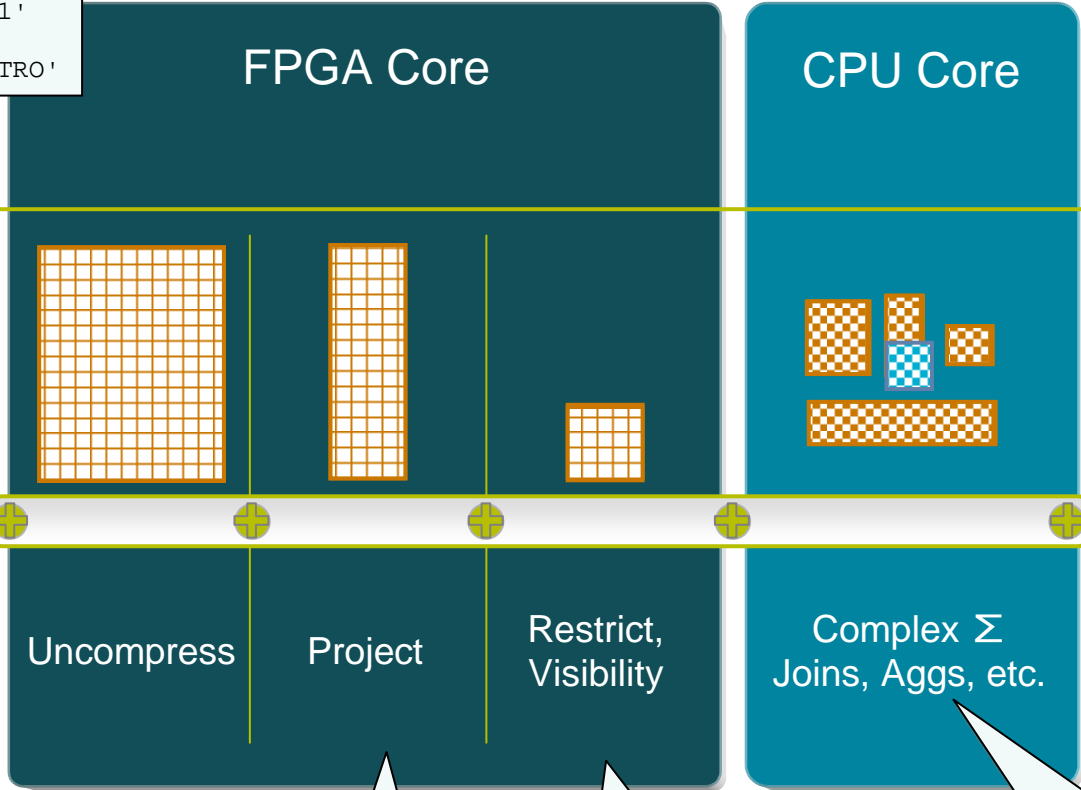
# S-Blade Data Stream Processing



```
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)

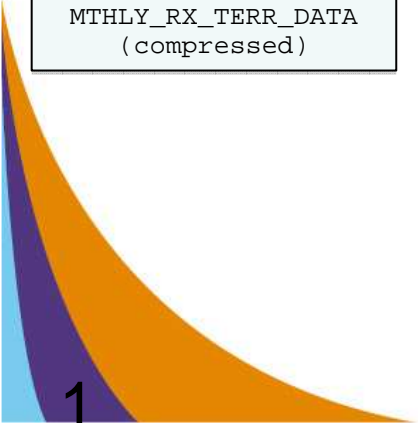


```
select DISTRICT,
       PRODUCTGRP,
       sum(NRX)
```

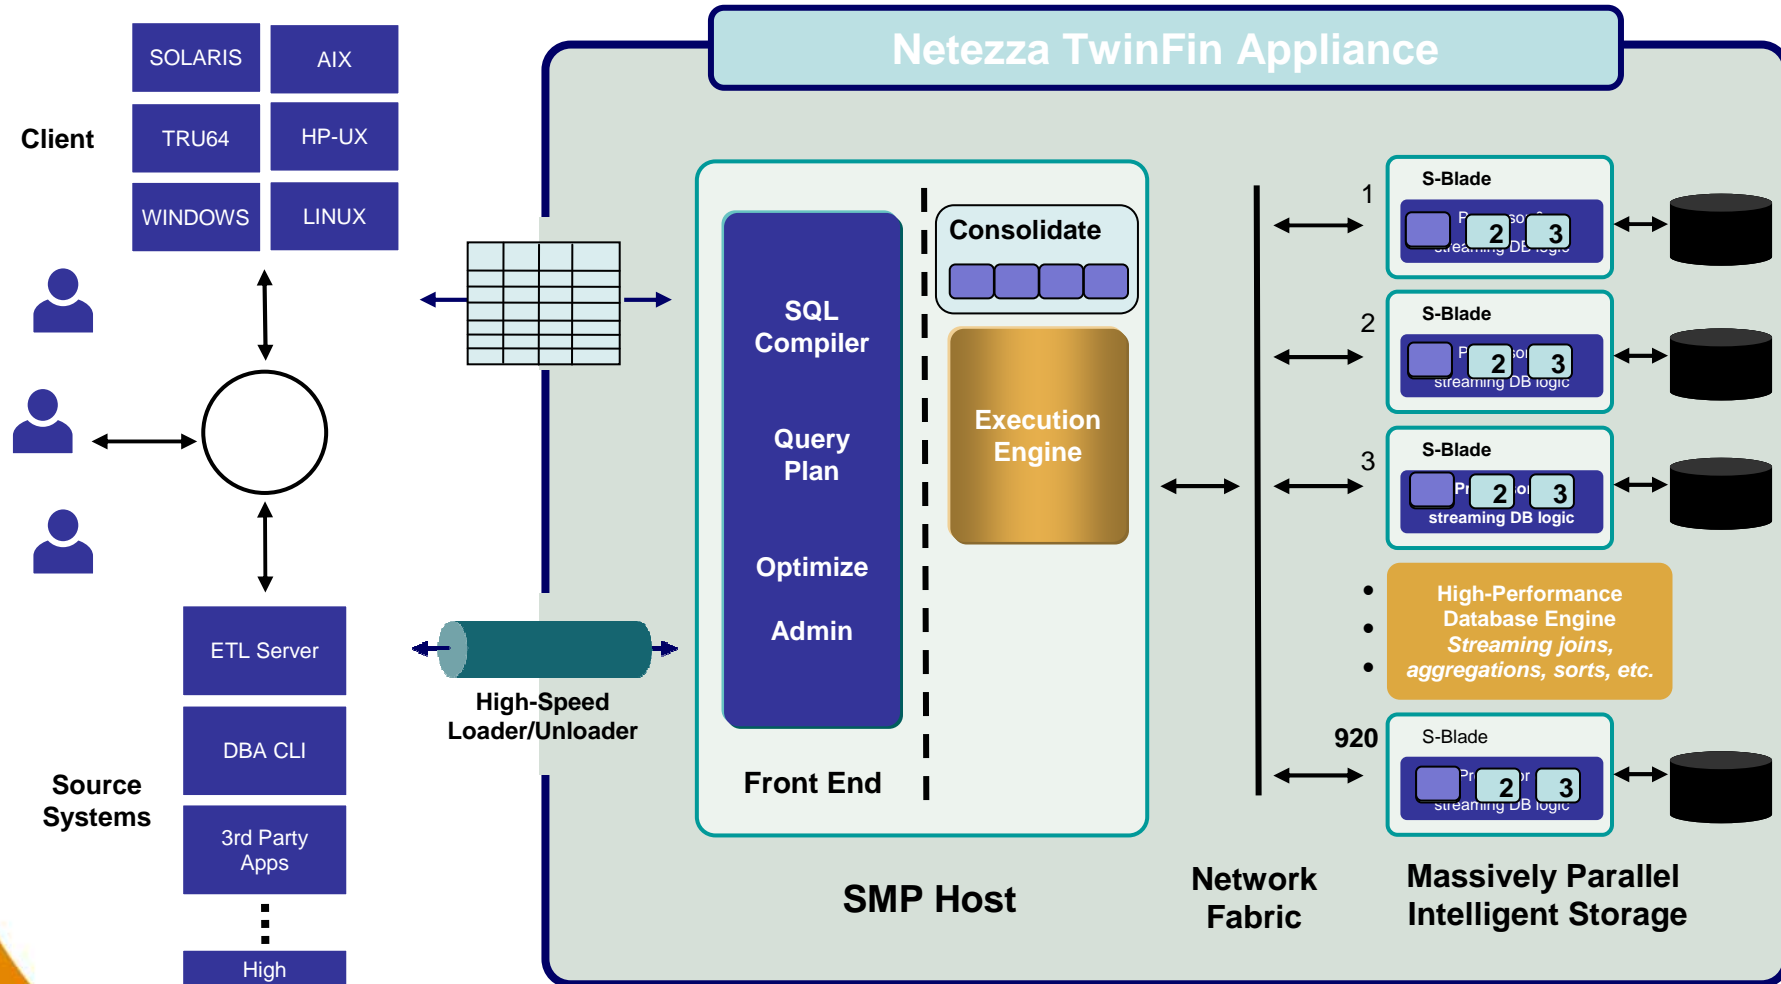

```
where MONTH = '20091201'
and    MARKET = 509123
and    SPECIALTY = 'GASTRO'
```

sum(NRX)

inforianet



# Asymmetric Massively Parallel Processing™

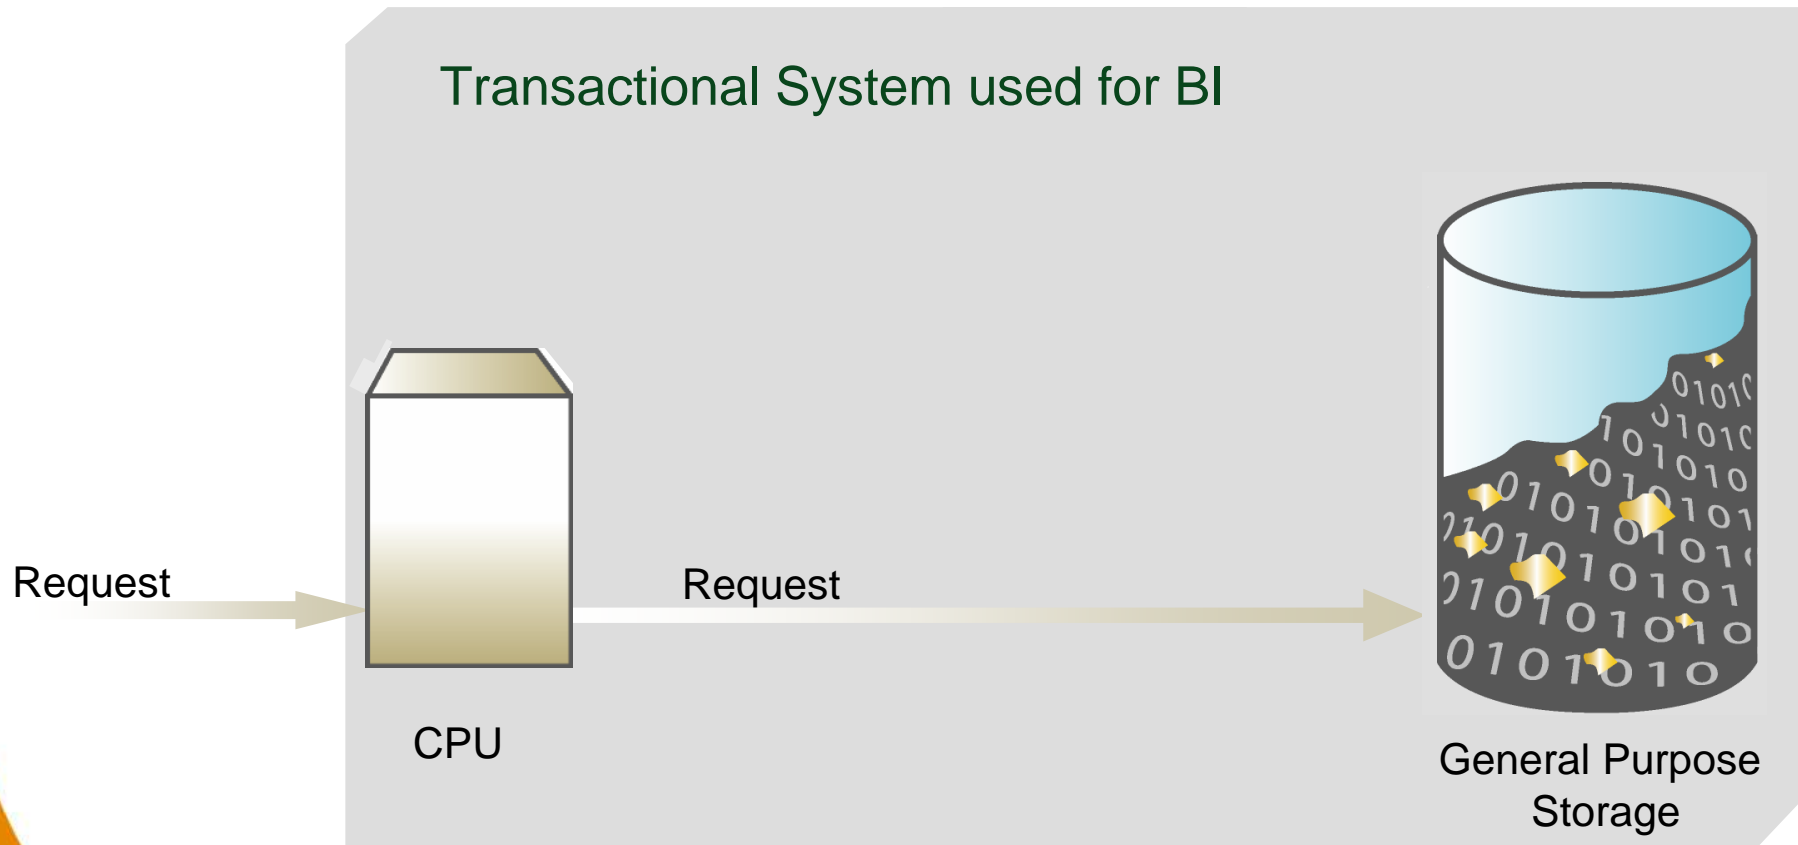


InfoPlanet




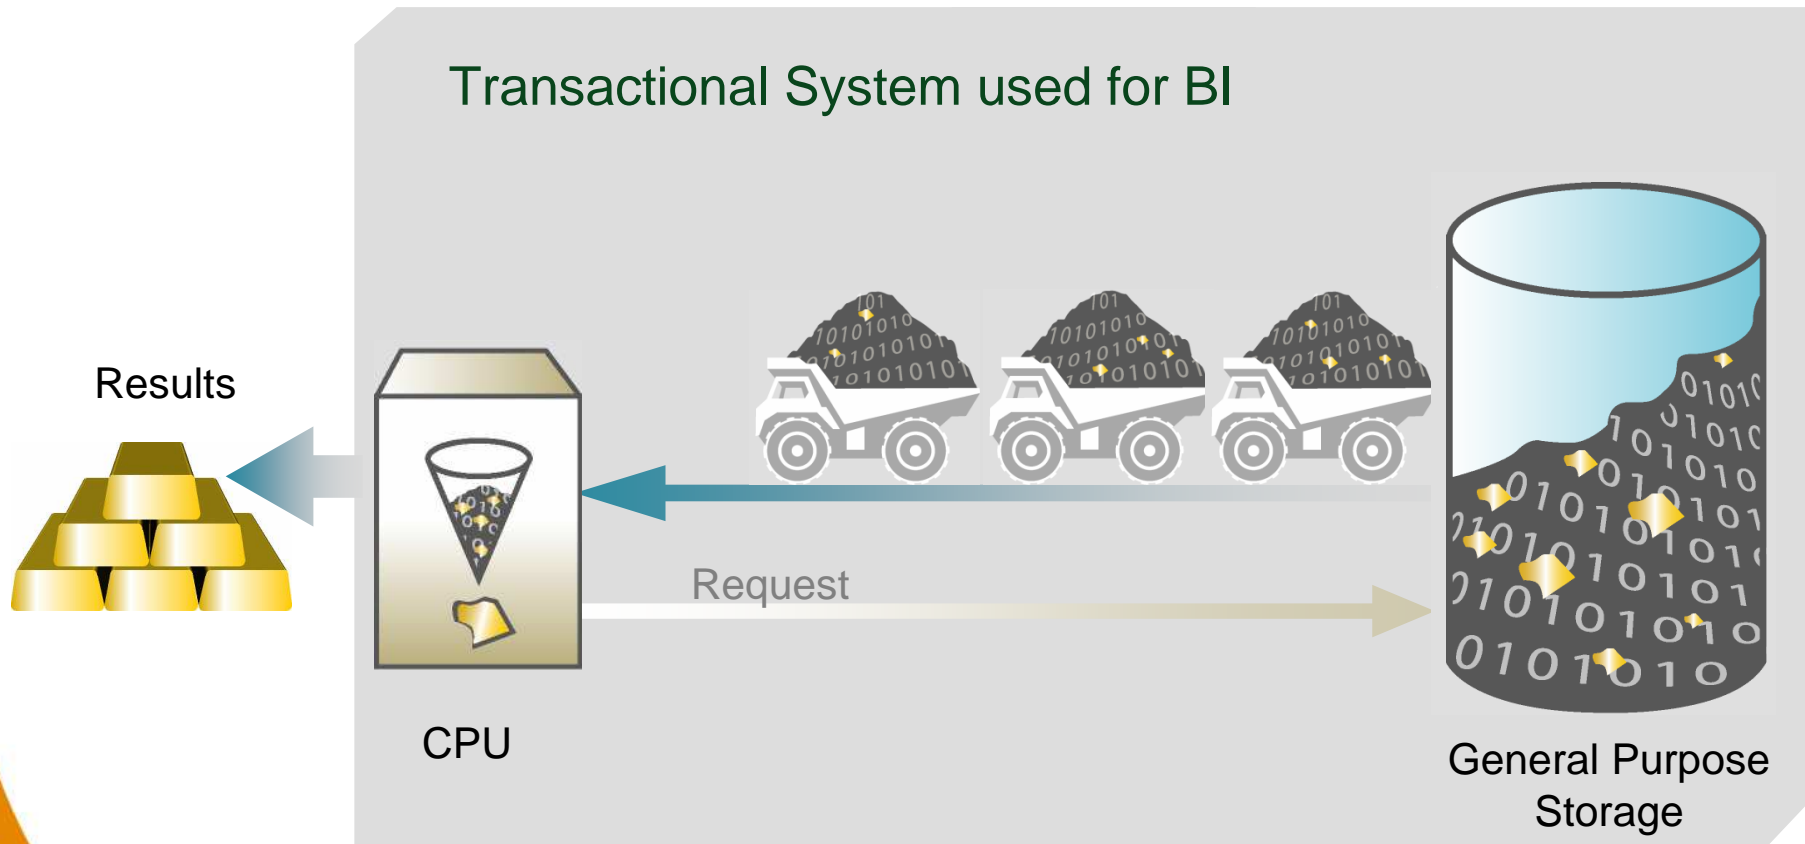
# Data Warehouse Workload

*Fewer requests, lots of data manipulation*



# Data Warehouse Workload

*Transaction systems are inefficient for data shuffling*






# Data Warehouse Blades


*Designed for Tera-scale Business Intelligence*

## Netezza Performance Server™ System



Asymmetric Massively Parallel Processing

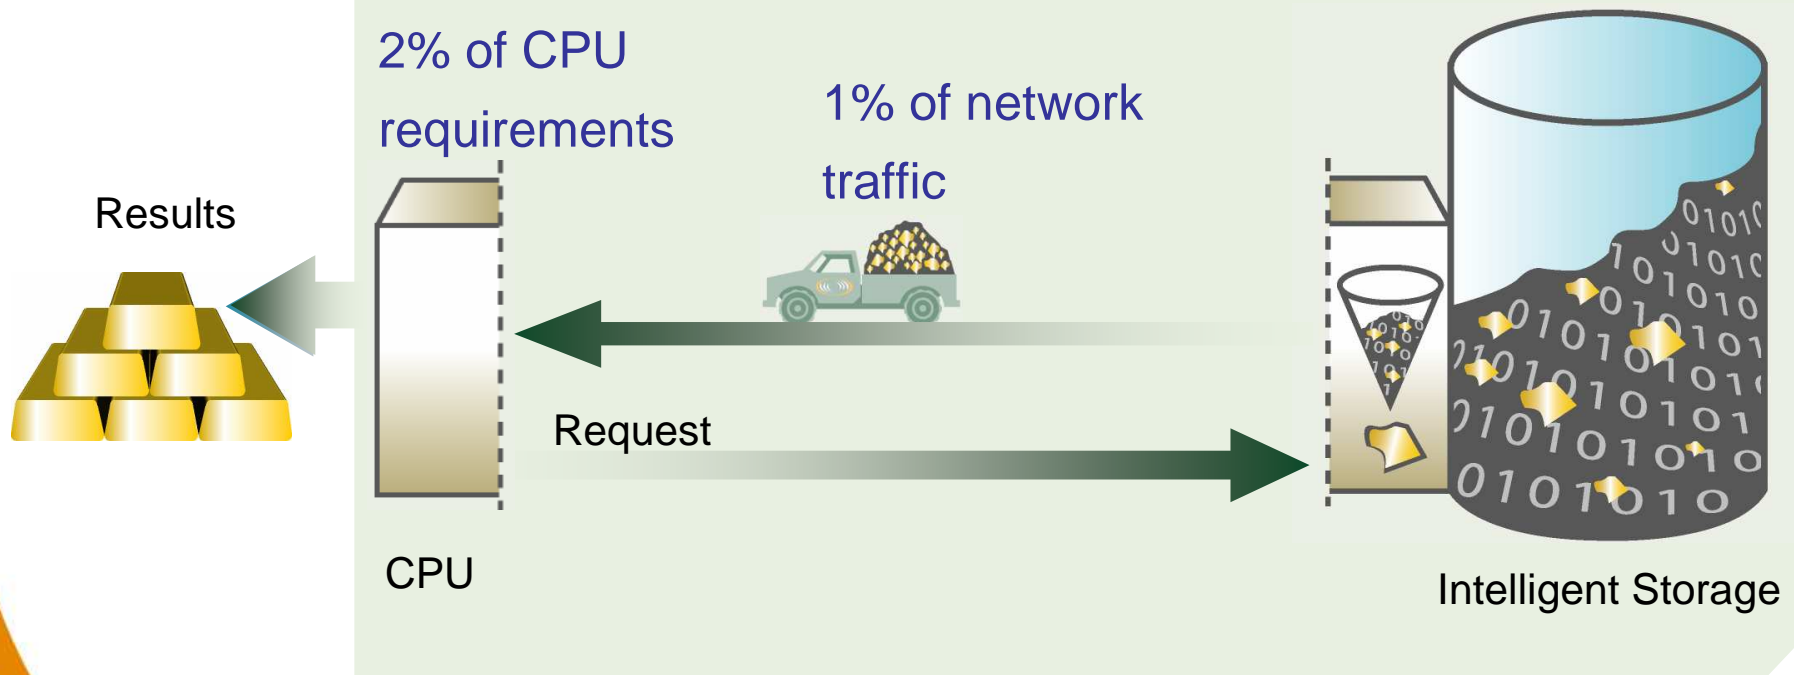
**InfoPlanet**



# Data Warehouse Blades


*Highly efficient data movement*

## Netezza Performance Server™ System



Asymmetric Massively Parallel Processing

**InfoPlanet**



# The IBM Netezza TwinFin™ Appliance



**Disk Enclosures**

*Slice of User Data  
Swap and Mirror partitions  
High speed data streaming*

**SMP Hosts**

*SQL Compiler  
Query Plan  
Optimize  
Admin*

**S-Blades™  
(with FPGA-based  
Database Accelerator)**



*Processor &  
streaming DB logic  
High-performance database  
engine streaming joins,  
aggregations, sorts, etc.*



**InfoPlanet**



# The IBM Netezza S-Blade™




**InfoPlanet**






# S-Blade™ Components




SAS Expander Module

DRAM



SAS Expander Module

Dual-Core FPGA



Intel Quad-Core



IBM BladeCenter Server

Netezza DB Accelerator


**InfoPlanet**




# Disk Mirroring and Failover




- All user data and temp space mirrored
- Disk failures transparent to queries and transactions
- Failed drives automatically regenerated
- Bad sectors automatically rewritten or relocated



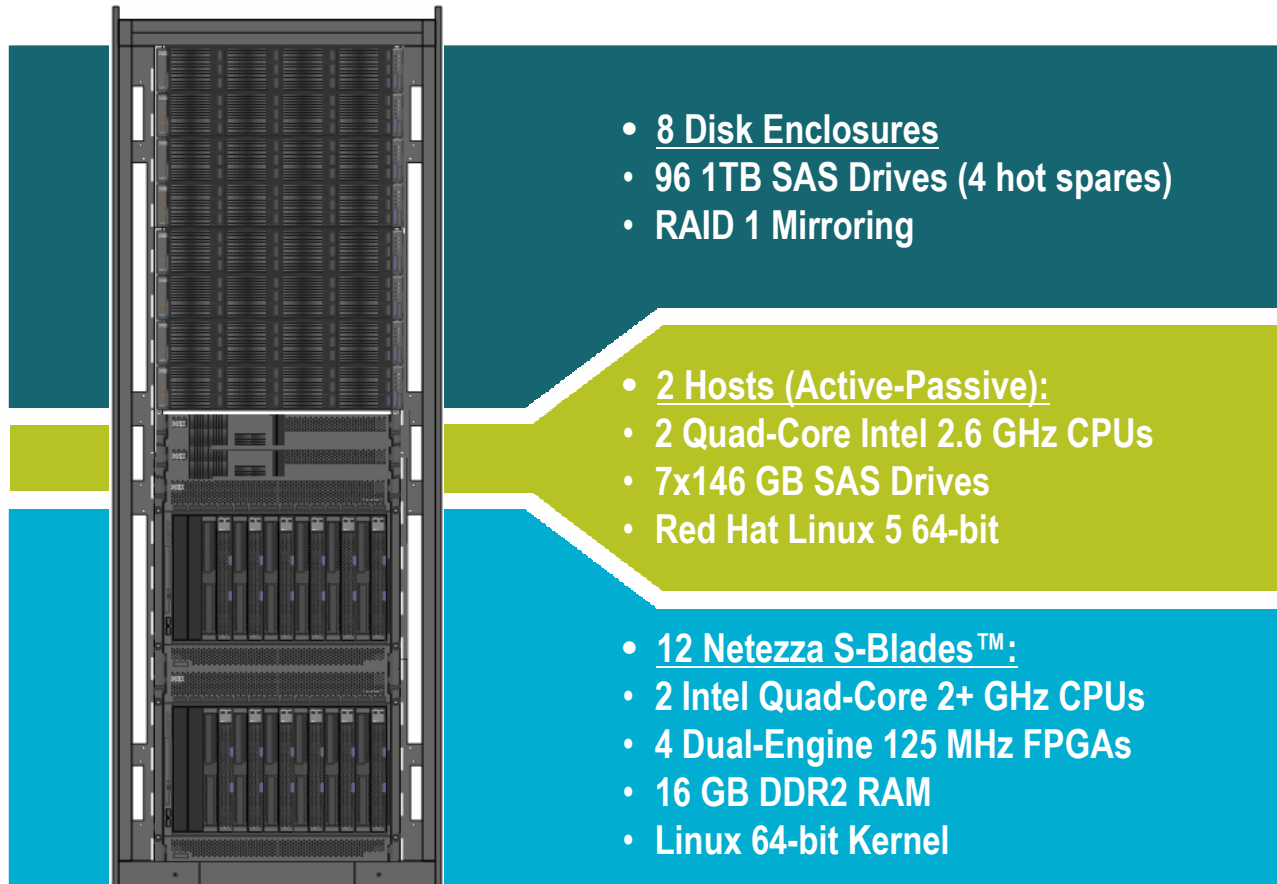
# S-Blade™ Failover



- Drives automatically reassigned to active S-Blades within a chassis
- Read-only queries (that have not returned data yet) automatically restarted
- Transactions and loads interrupted
- Loads automatically restarted from last successful checkpoint



# IBM Netezza TwinFin™ 12 Specification



- **User Data Capacity:** 128 TB\*\*
- **Data Scan Speed:** 144 TB/hr\*\*
- **Load Speed (per system):** 1.5+ TB/hr


- **Power Requirements:** 7.6 kW
- **Cooling Requirements:** 7.8 kW
- **Overall system availability:** 99.99%

\*\* : 4X compression assumed



**InfoPlanet**




# Agenda



Business Perspective  
Architecture  
**Enterprise Integration**  
Appliance Simplicity  
Summary





# Netezza **twinFin**™ Models


					...		...	
	TF3	TF6	TF12	TF24	...	TF48	...	TF120
Snippet Processors	24	48	96	192	...	384	...	960
Capacity (TB)	8	16	32	64	...	128	...	320
Compression (TB)	32	64	128	256	...	512	...	1280




## Skimmer

Snippet Processors	8
Capacity (TB)	2.5
Compression (TB)	10

Capacity = User Data space  
 Compression = Effective User Data Space



# Seamless Integration with Enterprise Ecosystem



Certified inter-operability with leading applications

Data Integration



Business Intelligence and Analytics



Advanced Analytics and Data Mining




InfoPlanet



# IBM and Netezza Integration Story



Product Name	Category	Description
Cognos 8	Reporting	Certified
Cognos Powerplay	Reporting	Certified
Cognos Data Manager	ETL	Certified
Cognos Data Transformer	Data Modeling	Certified
DataMirror	Data Integration	Certified
DataStage SE/EE	ETL	Certified
Exeros Discovery	Data Analysis	Pending
Exeros Discovery Validator	MDM & ETL	Validated
Guardium	DB Security/Audit	Certified
SPSS Clementine	Statistical Analysis	Certified
SPSS Modeler	Statistical Analysis	Certified
Tivoli Storage Manager	Data Backup	Certified
Unica Affinium	CRM	Certified




**InfoPlanet**






# Agenda



Business Perspective  
Architecture  
Enterprise Integration  
**Appliance Simplicity**  
Summary



# The Netezza Appliance Approach to Administration




No indexes and tuning

No storage administration

- No dbspace/tablespace sizing and configuration
- No redo/physical/Logical log sizing and configuration
- No page/block sizing and configuration for tables
- No extent sizing and configuration for tables
- No Temp space allocation and monitoring
- No RAID level decisions for dbspaces
- No logical volume creations of files
- No integration of OS kernel recommendations
- No maintenance of OS recommended patch levels
- No JAD sessions to configure host/network/storage
- Simple data distribution, HASH or ROUND ROBIN

No software installation

**Resources become Data Managers instead of Database Administrators**



# NPS Simplicity & Ramifications on TCO


Telecom Retailer and Service Provider

<b>Telecom Call Detail Record FACT (6 billion rows)</b>	<b>Oracle Object Count *</b>	<b>Netezza Object Count</b>
Tables	1	1
Indexes	12	
Table Partitions	47	
Index Partitions	564	
Table Partitions tablespaces	47	
Index Partitions tablespaces	47	
Table Data Files	170	
Index Data Files	122	
<b>TOTAL</b>	<b>1,010</b>	<b>1</b>



*“Look at all the weeks/months worth of effort, DBA design and maintenance that we don't have with Netezza. The appliance claims are true.”*

\*: Oracle data does not account for ADDITIONAL effort required in configuring and engineering the file system design to accommodate this index management scheme.

# Agenda



Business Perspective  
Architecture  
Enterprise Integration  
Appliance Simplicity  
Summary





# TwinFin™

The true data warehousing appliance.

---



Purpose-built analytics engine

Integrated database, server and storage

Standard interfaces

Low total cost of ownership

Speed: 10-100x faster than traditional system


Simplicity: Minimal administration and tuning

Scalability: Peta-scale user data capacity

Smart: High-performance advanced analytics

**InfoPlanet**





## Bold Claims are Made, But...We Prove It!

We prove we are simpler

We prove we deliver performance

We prove we work within your environment

We prove we integrate with your 3rd party tools

We prove we are “easy to do business with”

We prove we have the lowest TCO

We prove Business Value



**InfoPlanet**



Take the  
TwinFin™  
TestDrive

---

Out of the box thinking,  
right out of the box.

---

We're accelerating the Proof of Concept process.

We'll bring you a TwinFin™. Install it in your environment.

Then let it loose on your toughest problem.

Your data

Your site

Our technology

