NETEZZA

The Power to Question Everything™

A Data Warehouse Approach to Analyzing All the Data All the Time

Bill Blake

Netezza Corporation

April 2006

Sometimes A Different Approach Is Useful

- The challenge of scaling up systems where many applications need to access large data in global parallel file systems is well documented
- At the Multi Terabyte scale, It is hard to move the data from where it is stored to where it is processed …
- But if moving data to processing is so difficult, why not try an approach where the application owns the data and processing is moved to where the data is stored?
- The application in this case is the relational database, a very useful tool for data intensive computing

Operational vs. Analytical RDBMS

Extraction /Cleansing

Trans-formation

Loading

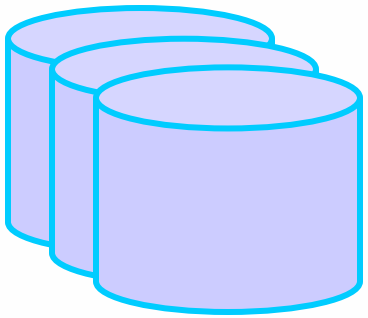

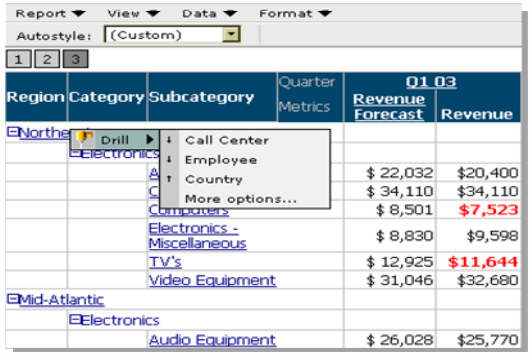

Reports

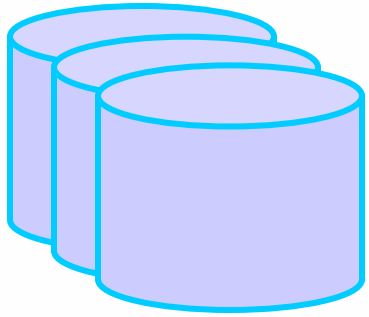

OLTP Databases

Processing Millions Of Transactions/sec

OLAP Databases aka Data Warehouse

Processing Hundreds of Terabytes/hour

NETEZZA
The Power to Question Everything™

The Legacy Focus: Transaction Processing

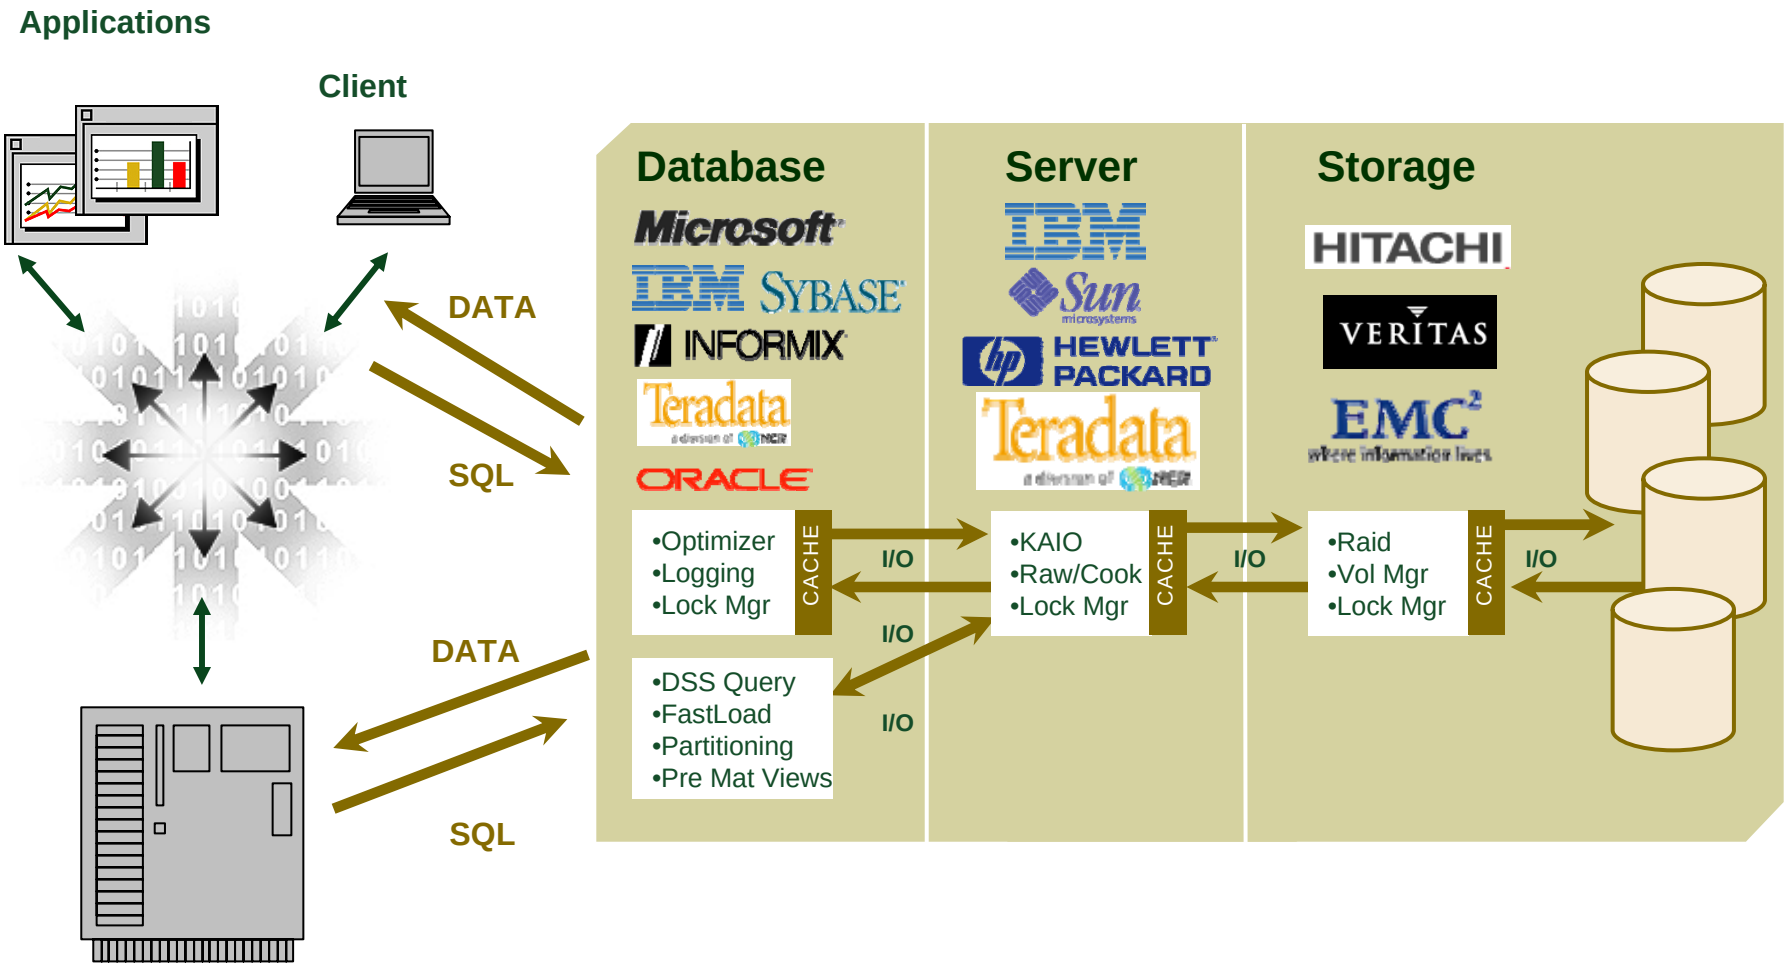

NETEZZA
The Power to Question Everything™

Netezza's Data Warehouse Appliance

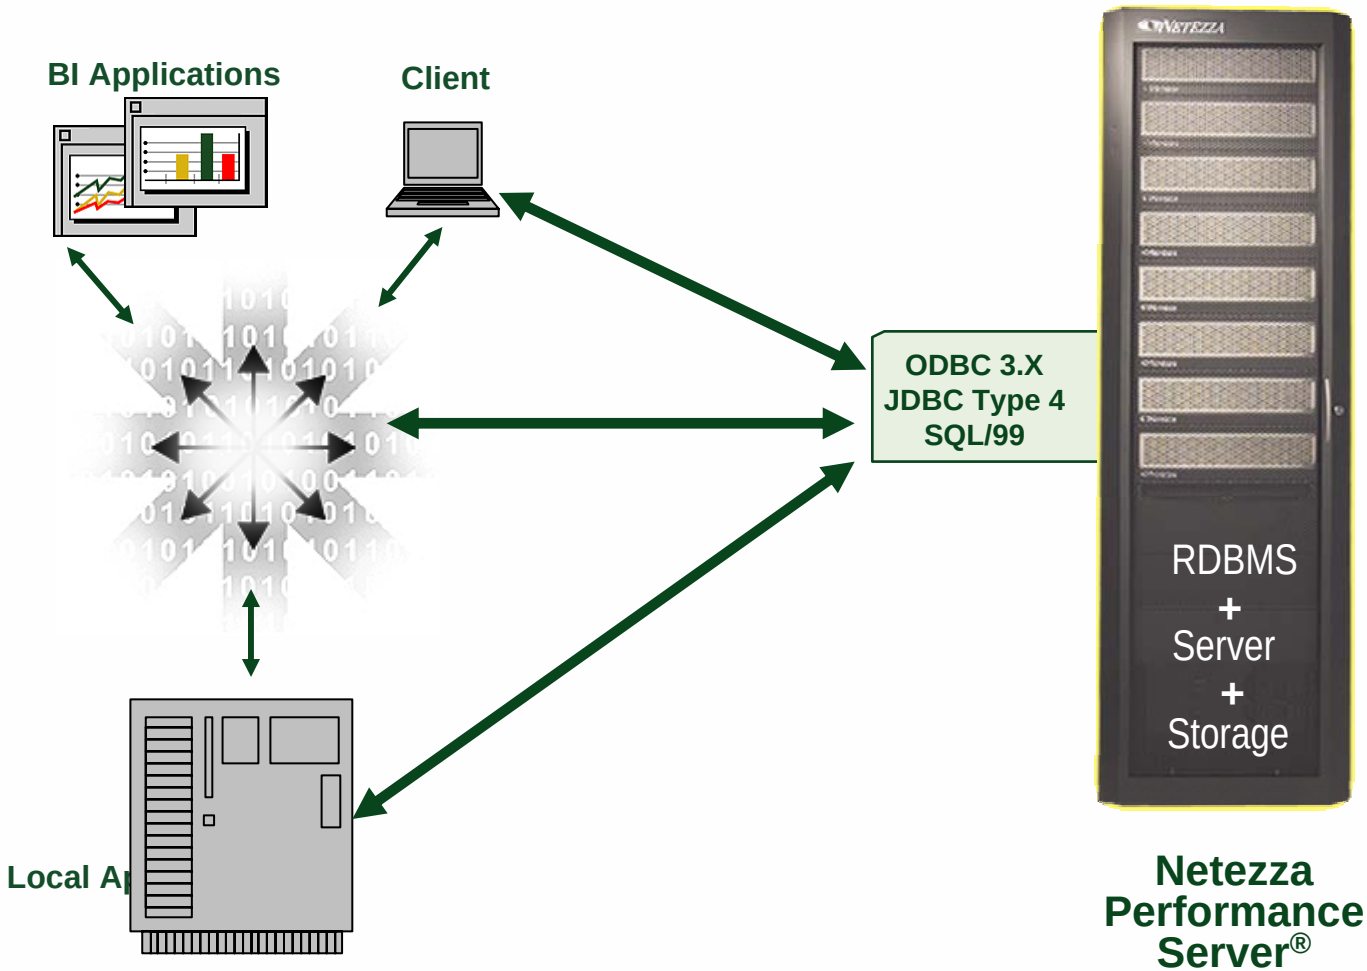

NETEZZA
The Power to Question Everything™

The Challenges Driving Us At Netezza

Forces driving disruptive change

- Sub-transactional data in a fully-connected world
- Ever-increasing need for speed
- Increasing regulatory requirements
- Market mandate for operational simplicity
- Need for actionable intelligence from unlimited data at real-time speeds

This Need Cannot be Met by Today's Systems

- ✓ ***Linux cluster scaling limited by network performance & system management complexity***
- ✓ ***Scaling with large NUMA SMP servers limited by I/O, network performance & operating system complexity***

NETEZZA
The Power to Question Everything™

Not All Computing Tasks Fit into Memory – The Analytic DB Challenge

There are benefits to scaling up analytic DBs:

- Transactional and referential integrity
- High level query language (with parallel run time optimization performed by application's query planner)
- Operation on sets of records in tables (vs sequential access of records in files)
- Database standards have matured and are now consistent across the industry
- Data volumes have grown from gigabytes to hundreds of terabytes
- Disk storage is now less than $1 per Gigabyte!

For Perspective (1980's) …

- The relational database was invented on a system that merged server, storage and database
- It was called a mainframe!

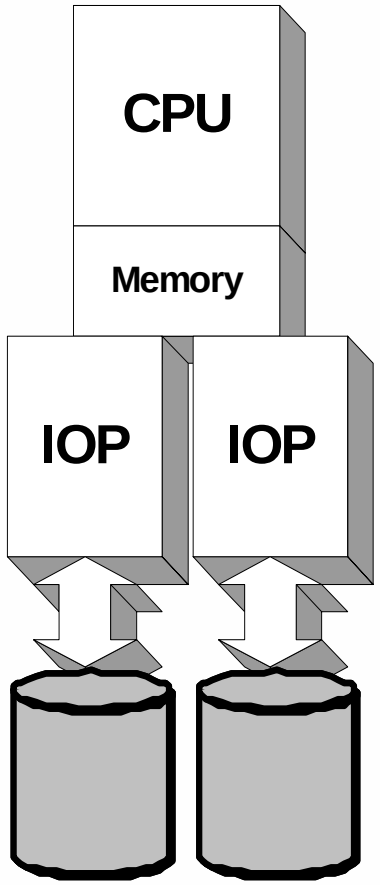

NETEZZA
The Power to Question Everything™

By The 1990's, Rules Changed

- Mainframes attacked by killer micros!
- Memory grew large
- I/O became weak
- System costs dropped
- Storage moved off to the network

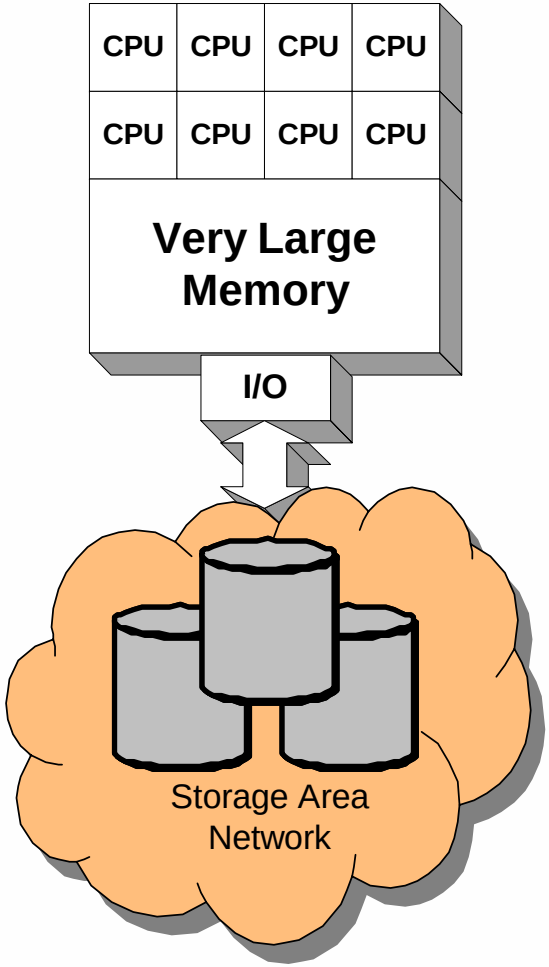

NETEZZA
The Power to Question Everything™

Capacity Was Added By Clustering

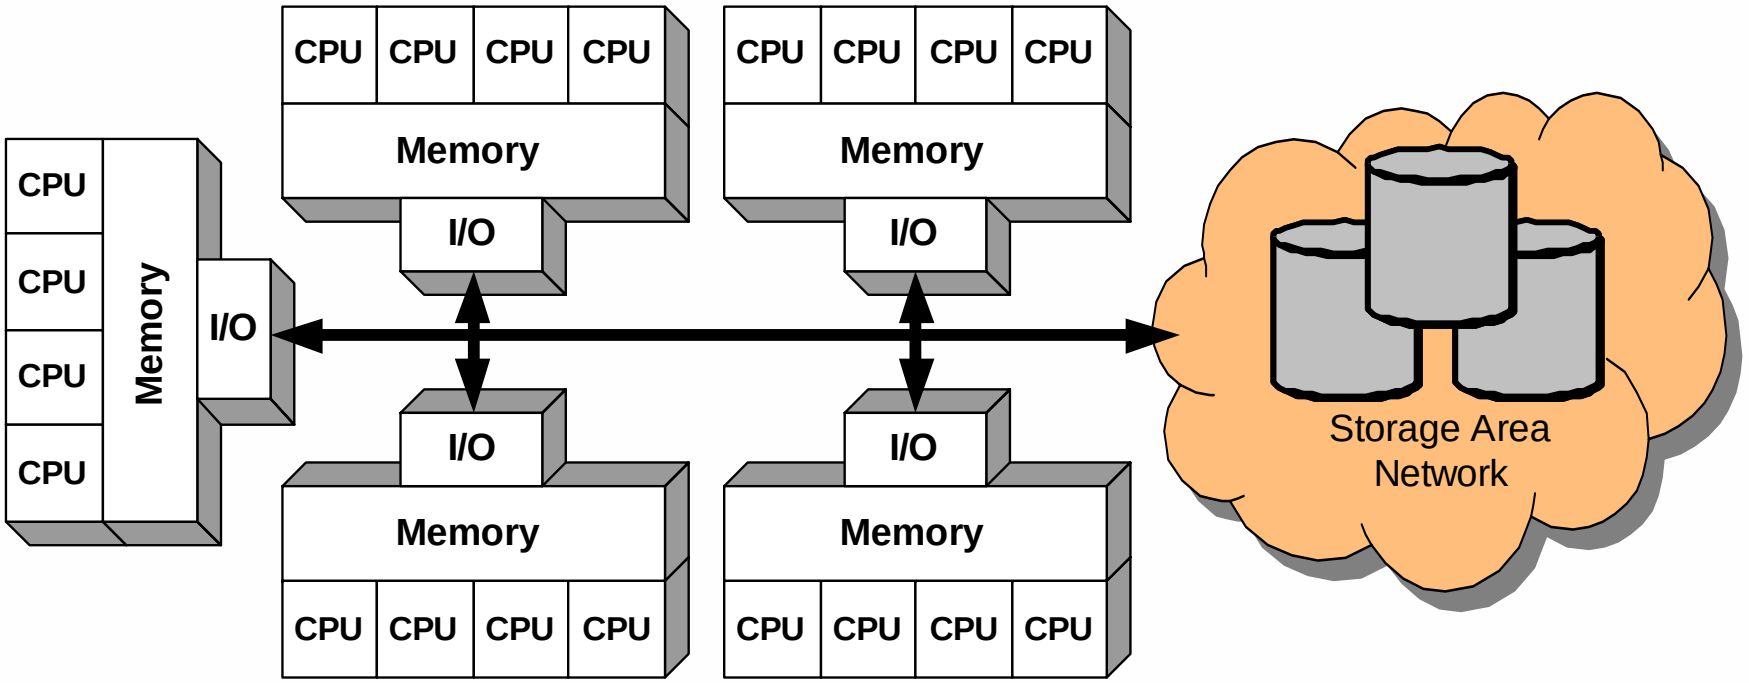

SAN limits Moving Data to the Processors

NETEZZA
The Power to Question Everything™

Data Flow – The Traditional Way

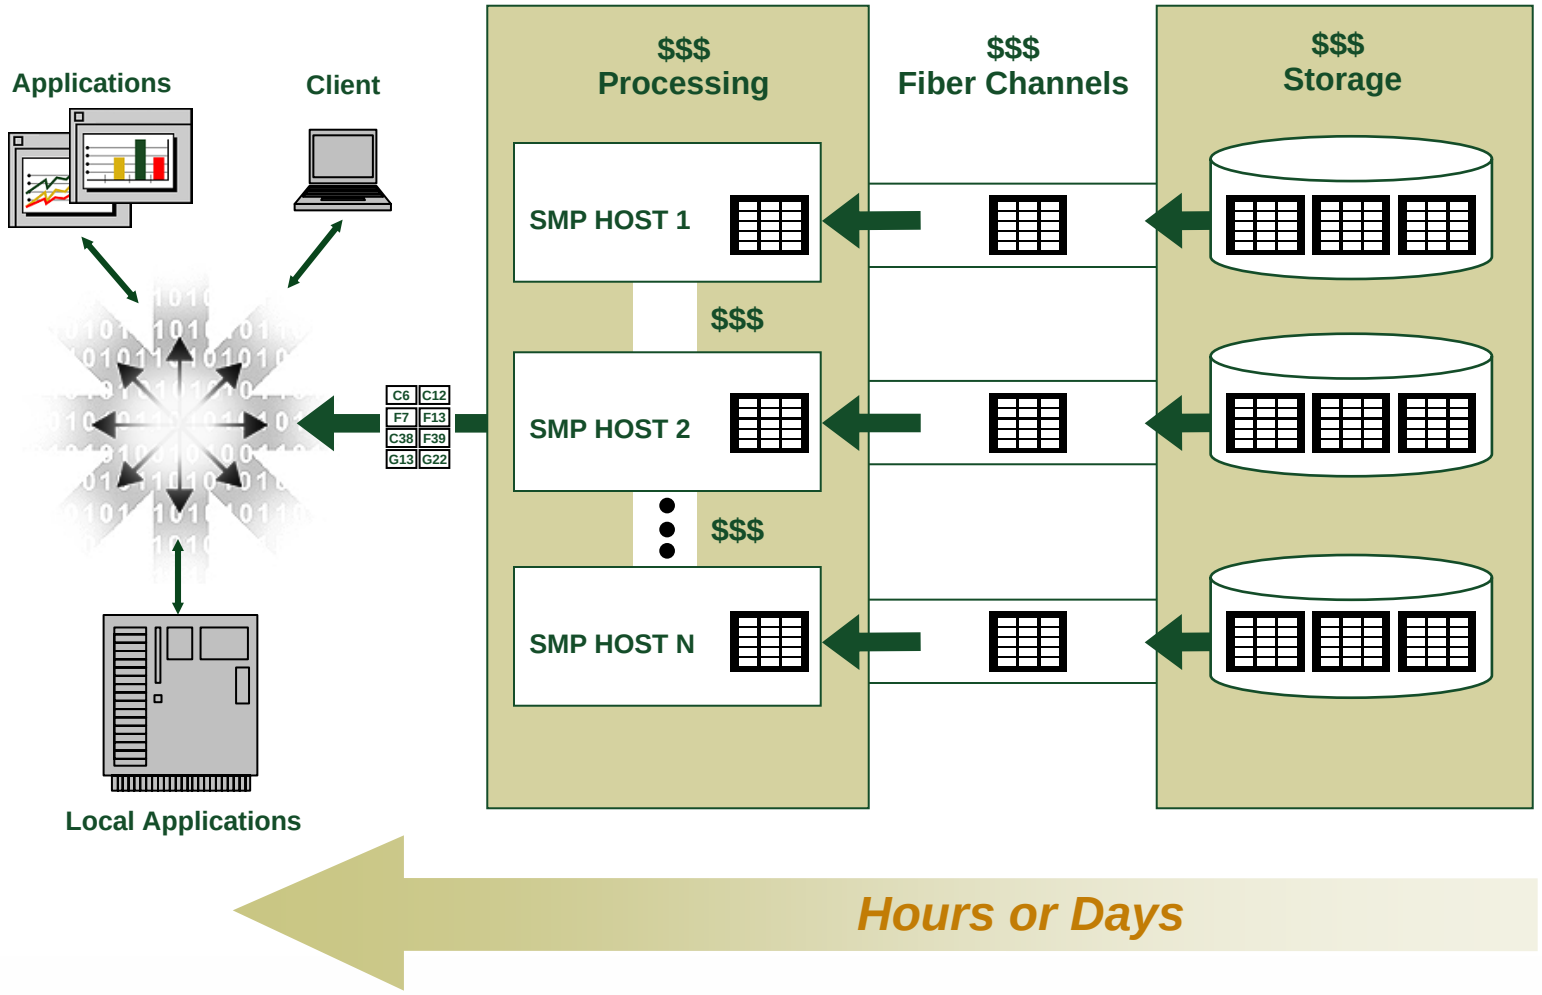

NETEZZA
The Power to Question Everything™

Moving Processing to the Data

- **Active Disk architectures**
  - > Integrated processing power and memory into disk units
  - > Scaled processing power as the dataset grew
- **Decision support algorithms offloaded to Active Disks to support key decision support tasks**
  - > Active Disk architectures use stream-based model ideal for the software architecture of relational databases

In Netezza's NPS® System: "Snippet Processing Units" take streams as inputs and generate streams as outputs

SQL Query Flow Diagram

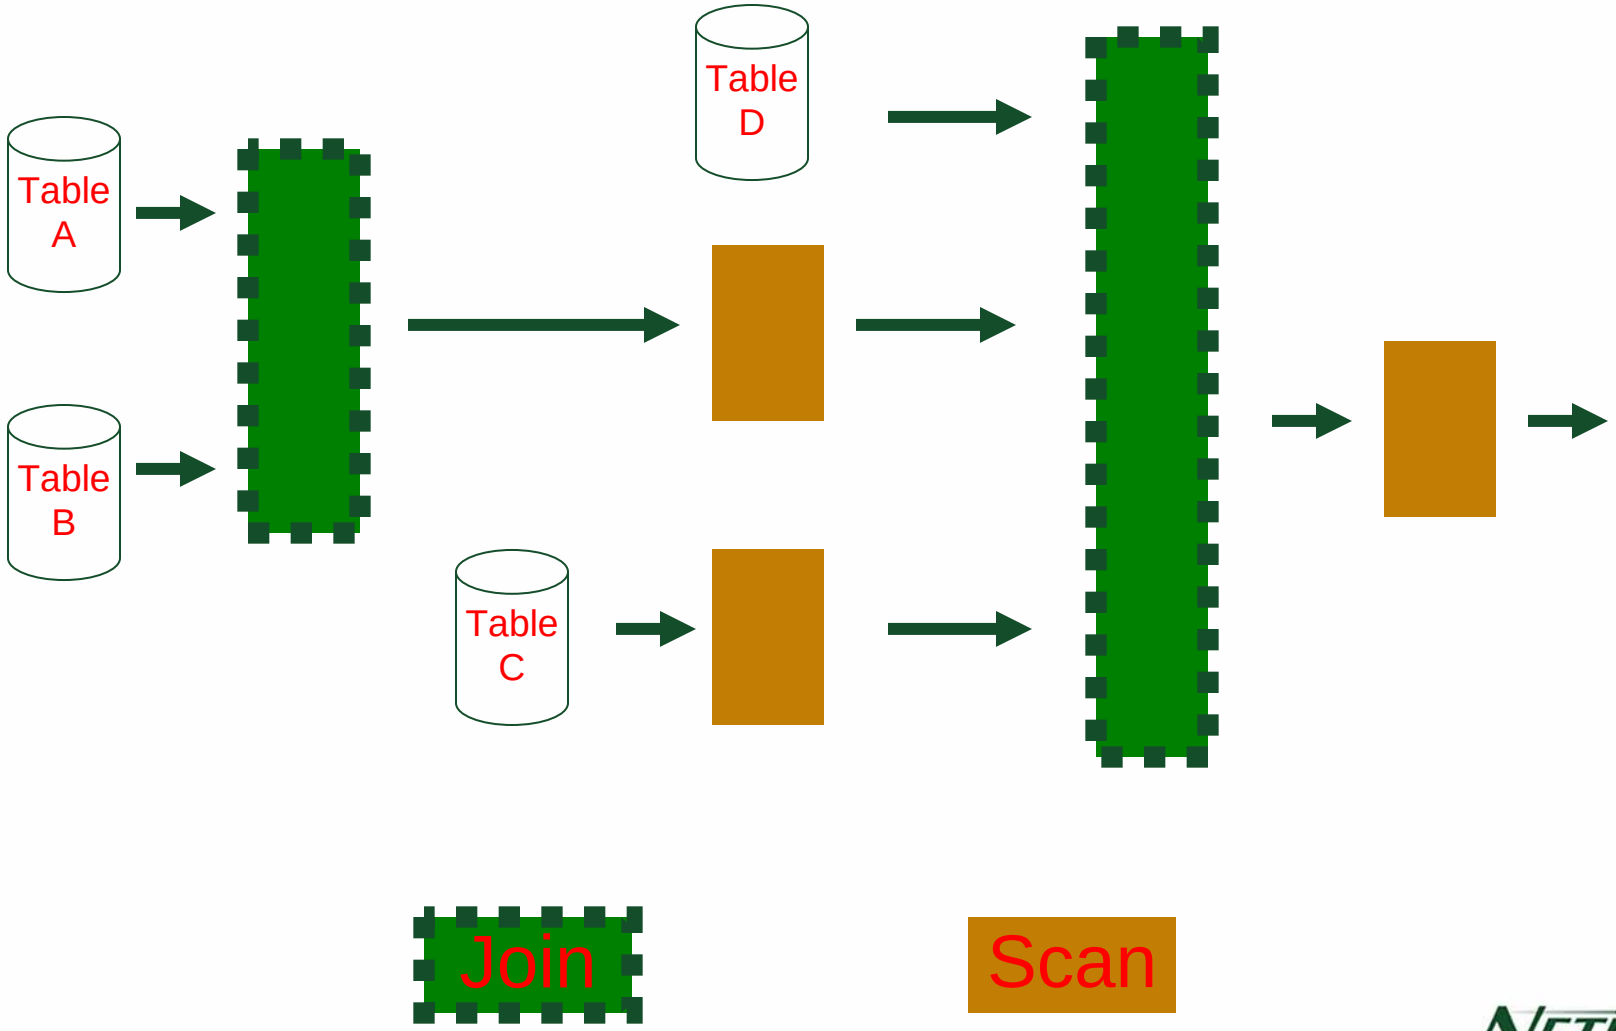

NETEZZA
The Power to Question Everything™

Streaming Data Flow

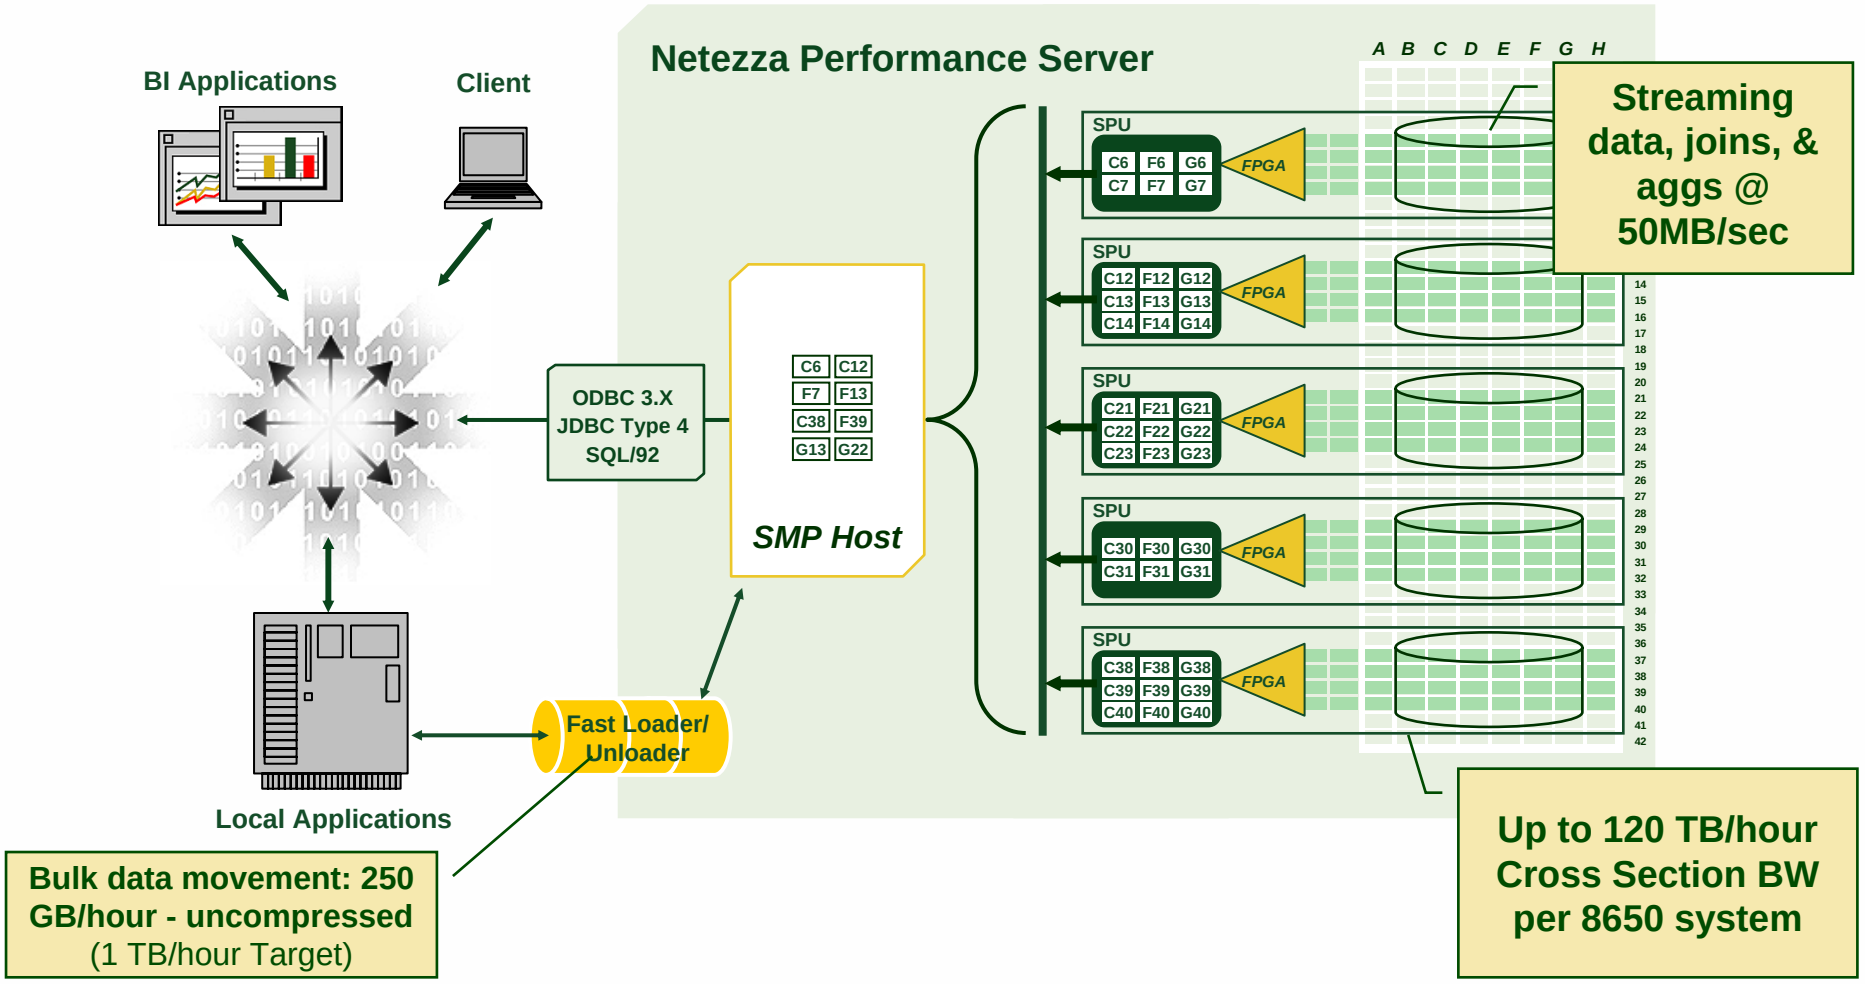

NETEZZA
The Power to Question Everything™

Active Disks as Intelligent Storage Nodes

Netezza Performance Server

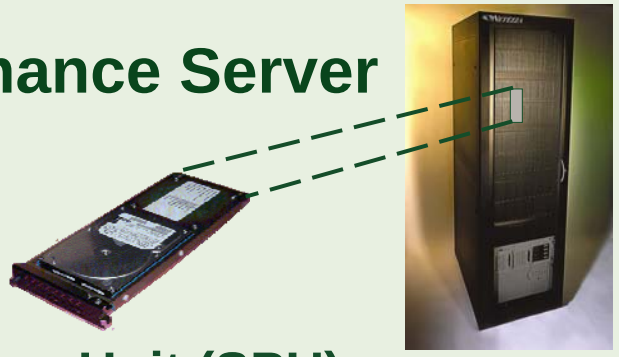

Netezza added:

- Highly optimized query planning
- Code generation
- Stream processing

Result: 10X to 100X performance speedup over existing systems

Snippet Processing Unit (SPU)

A compute node for directly processing SQL queries on tables

NETEZZA
The Power to Question Everything™

Asymmetric Massively Parallel Processing™

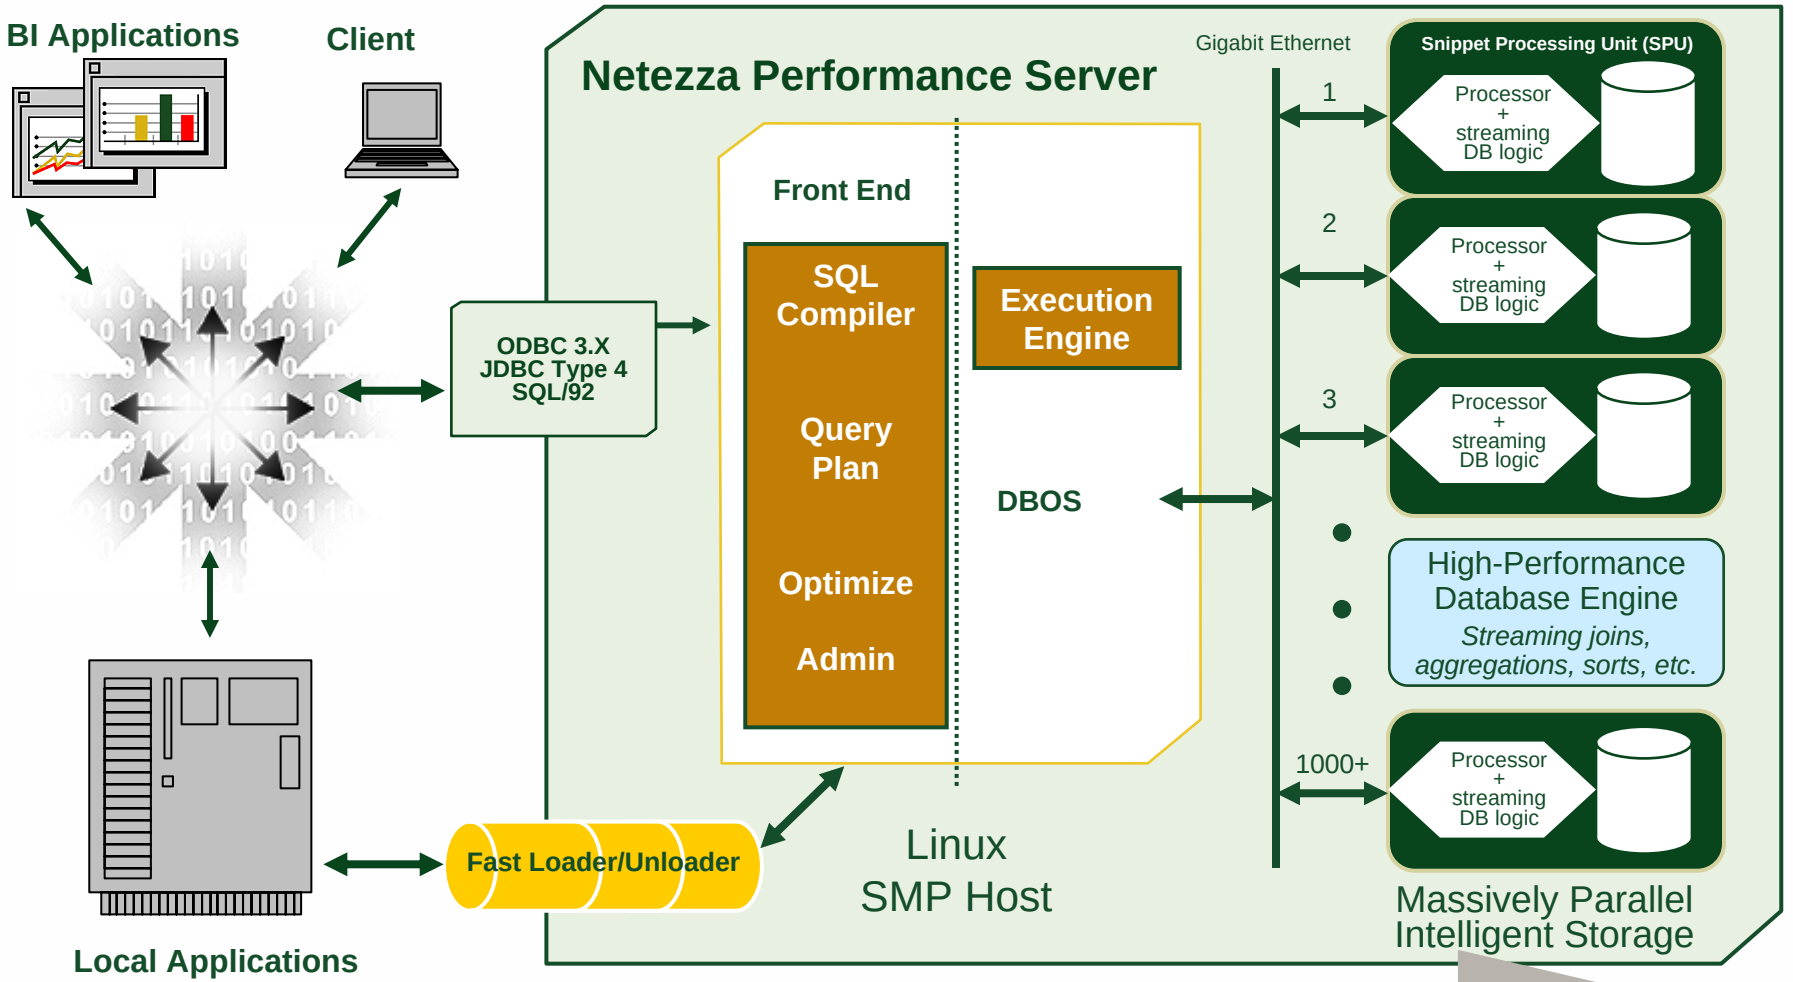

Packaging For High Density And Low Power

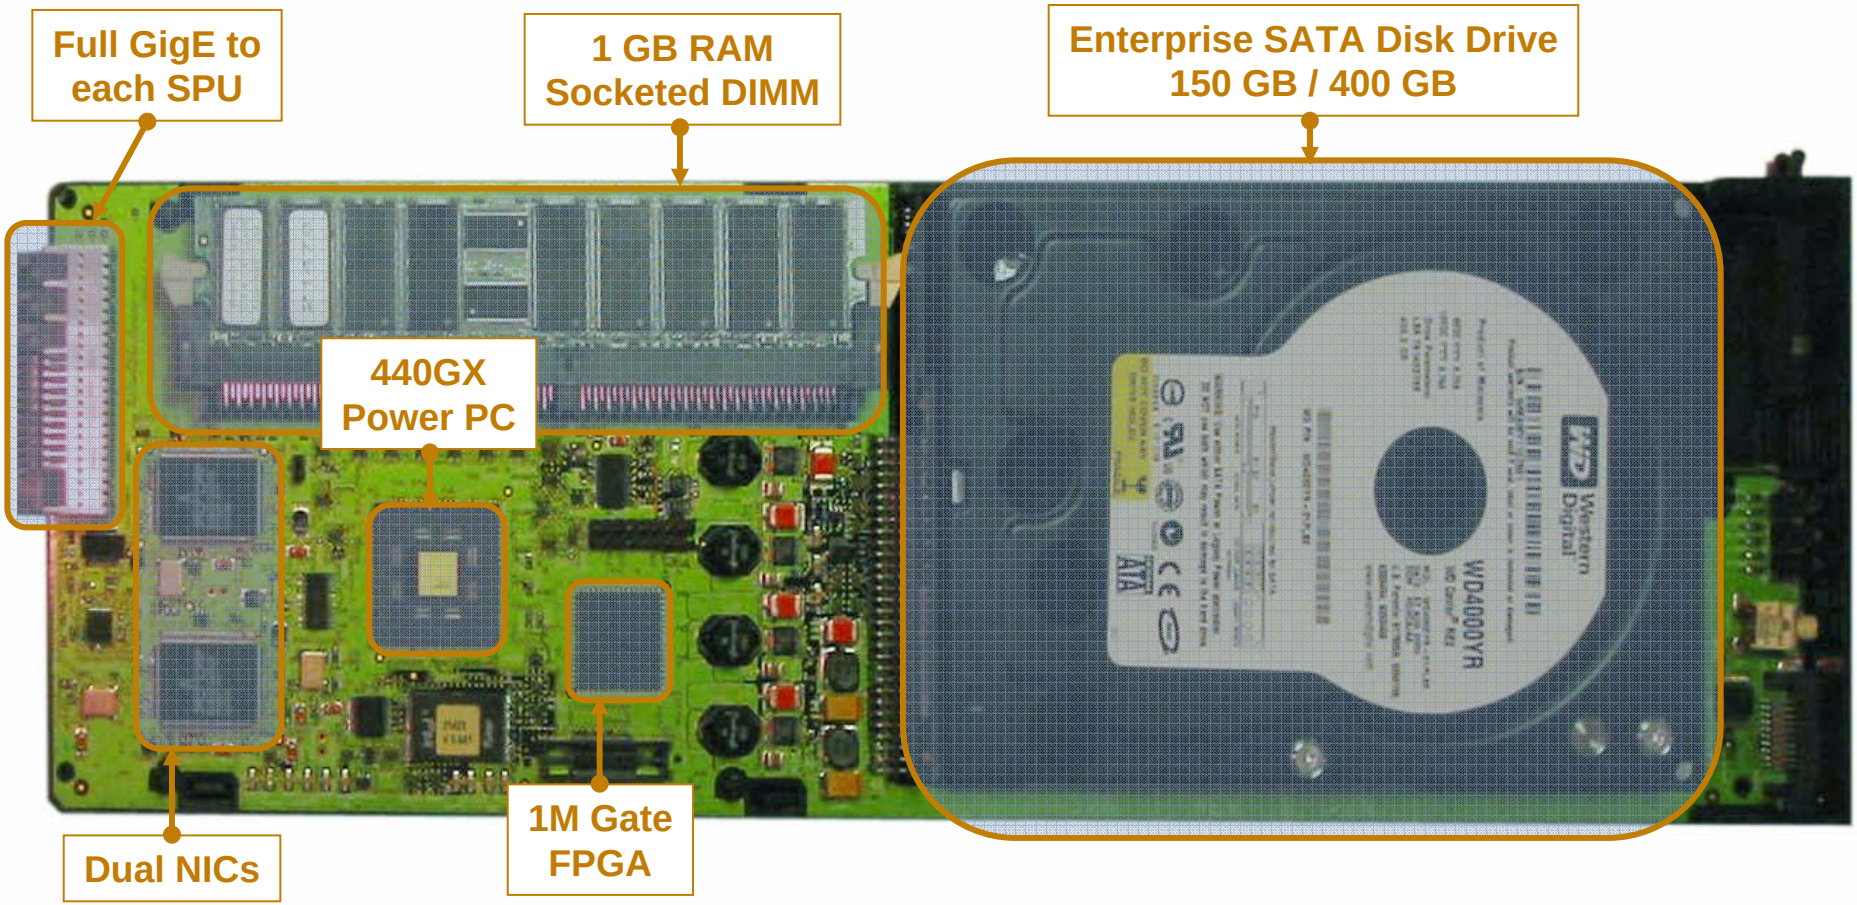

AFTER

Higher Performance – Greater Scalability – Higher Reliability

Binary Compiled Queries Executed on a Massively Parallel Grid

```
select c_name, sum(o_tot      rice) ice from customer, orders
where o_orderkey in (s      t l_orde      ey from lineitem2 where
o_orderkey=l_orderkey and l_shipda      ='01-01-1995' and
l_shipdate
c_name;" t
```

```
/********* Code *******/

void GenPlan1(CPlan *plan, char *bufStarts,char *bufEnds, bool
lastCall) {
    //
    // Setup for next loop (nodes 00..07)
    //
    // node 00 (TScanNode)
    TScanNode *node0 = (TScanNode*)plan->m_nodeArray[0];
    // For ScanNode:
                TScan0 *Scan0 = BADPTR(TScan0*);
                CTable *tScan0 = plan->m_nodeArray[0]->m_result;
    char *nullsScan0P = BADPTR(char *);
    // node 01 (TRestrictNode)
    TRestrictNode *node1 = (TRestrictNode*)plan->m_nodeArray[1];
    // node 02 (TProjectNode)
    TProjectNode *node2 = (TProjectNode*)plan->m_nodeArray[2];
    // node 03 (TSaveTempNode)
    TSaveTempNode *node3 = (TSa      pNode*)plan->m_nodeArray[3];
    // For SaveTemp Node:
    TSaveTemp3 *SaveTemp3 =      PTR(TSav      mp3*);
    CTable *tSaveTemp3 = node3->m_resul
    CRecord
    // node
…
```

```
1011010101010101010111111010101001001010101110101010010111101
0100101011111011010010101010111010101100101010101011111010010010
0101010101010101010010101001111110101010101010100101010101010010
1001011010011111111010101010011010010101010100101010101010010
0101010101001010101010011101010101010101010…
```

```
 c_name       |   price
--------------------+-----------
 Customer#000000796 | 318356.97
 Customer#000001052 | 293680.56
 Customer#000001949 | 215280.98
 Customer#000002093 | 282531.93
 Customer#000005656 | 335297.31
 Customer#000005861 | 233691.03
 Customer#000006002 | 267000.92
 Customer#000006343 | 595819.82
 Customer#000006532 | 442254.91
….
real    0m0.552s
user    0m0.010s
sys     0m0.000s
```

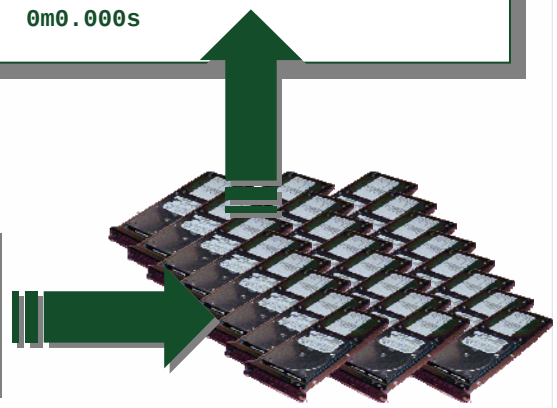

NETEZZA
The Power to Question Everything™

A Look Inside the SPU

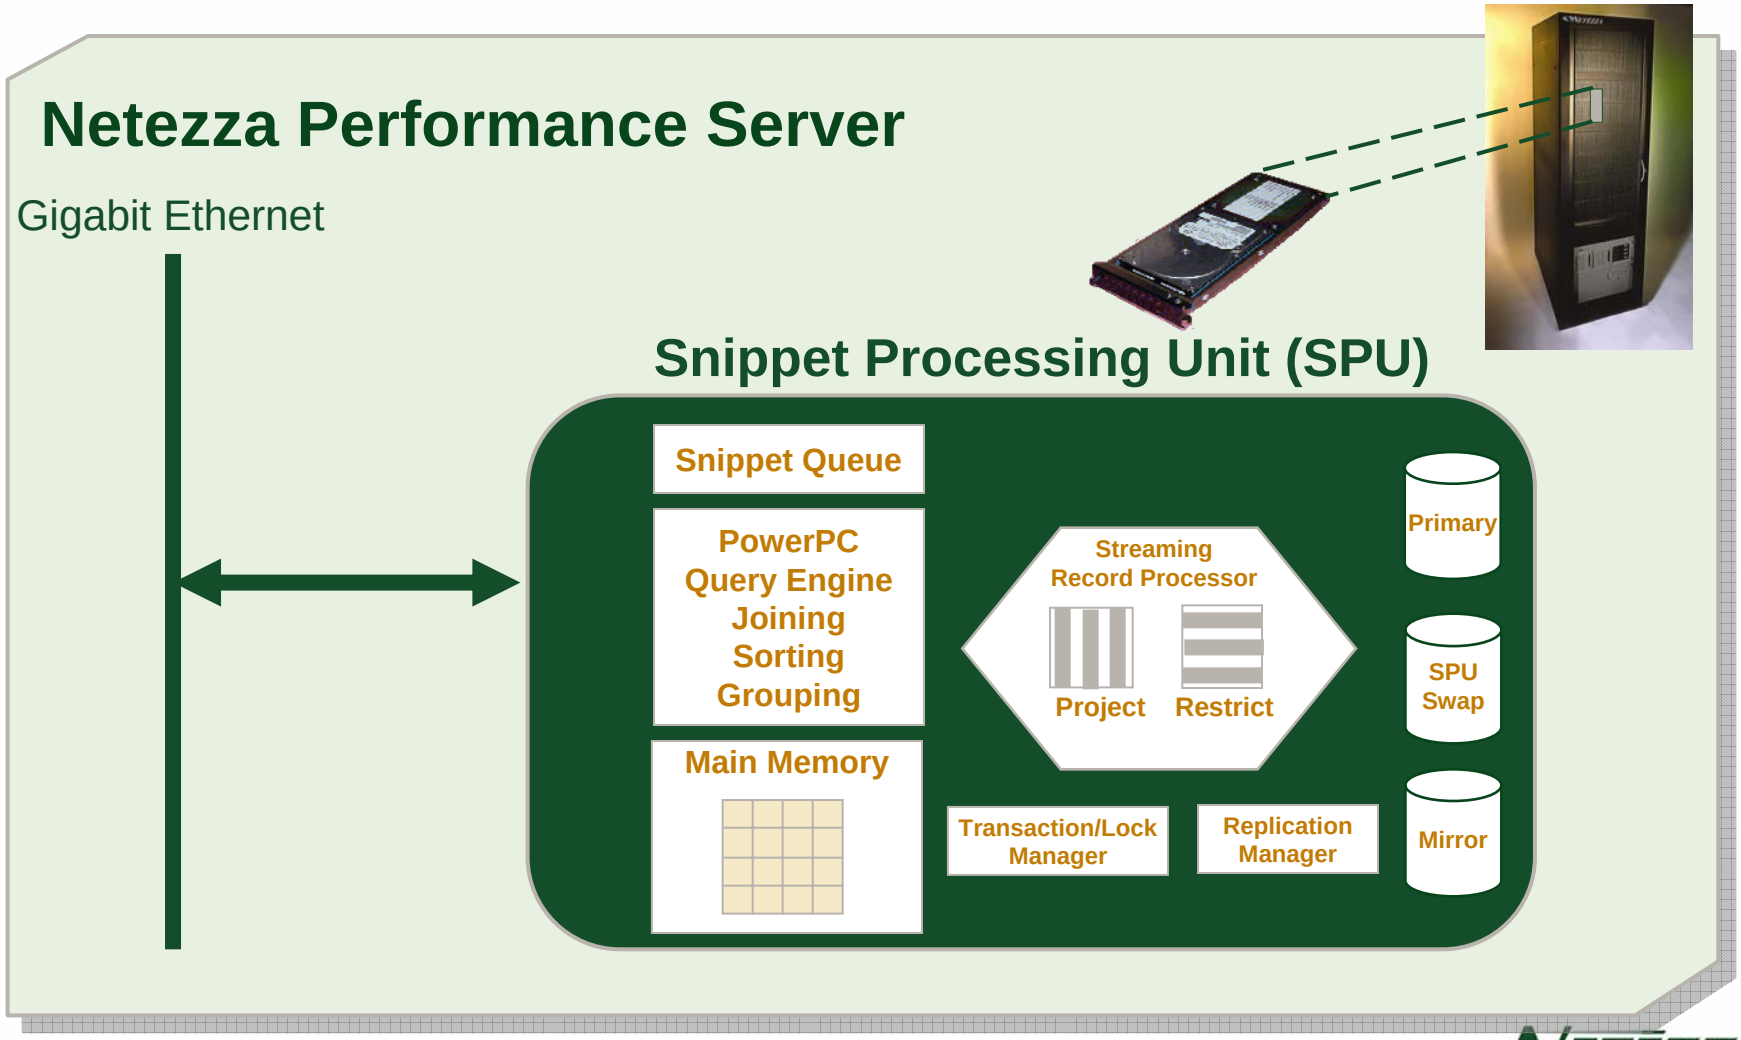

NETEZZA
The Power to Question Everything™

SELECT count (*) , sex , age FROM emp WHERE state = ‘VA’ and age > 18 GROUP BY sex , age ORDER BY age ;

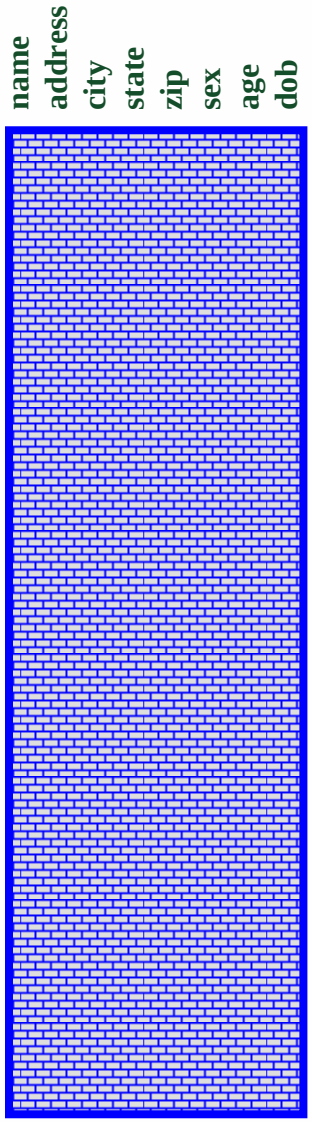

NETEZZA
The Power to Question Everything™

SELECT count (*) , sex , age FROM emp WHERE state = 'VA' and age > 18 GROUP BY sex , age ORDER BY age ;

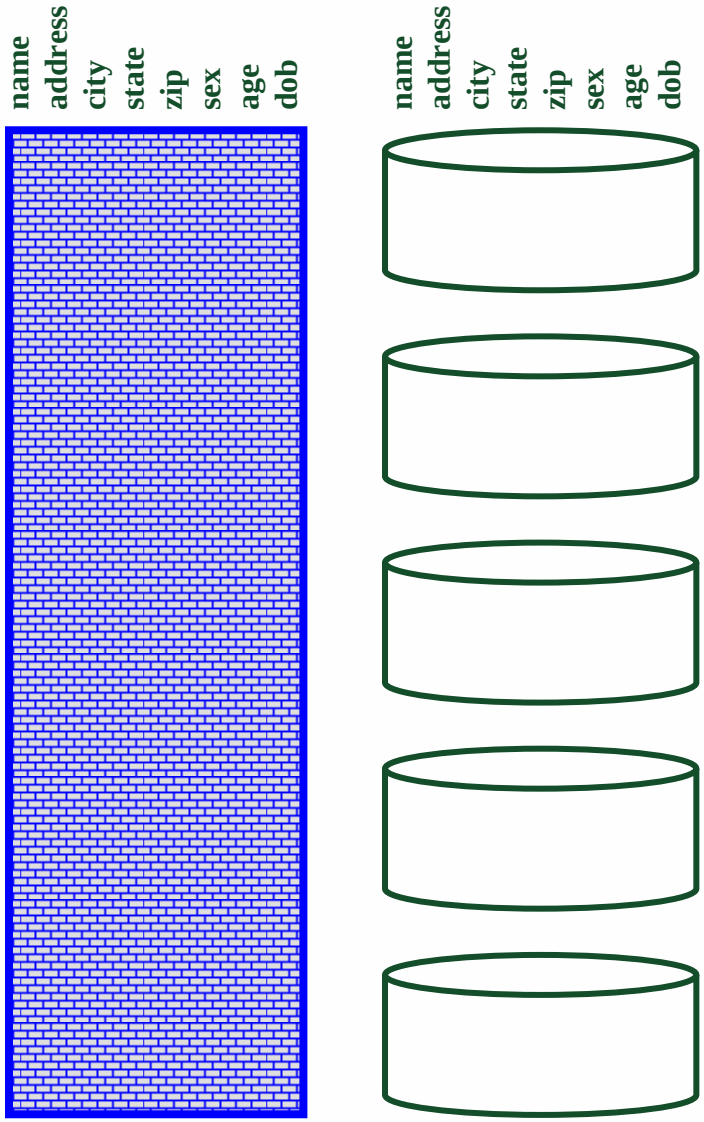

First things first. The table is distributed amongst all of the SPU's in the system so that is can be processed in parallel.

When the table is read, your scan speed is the SUM of the speed of all of the disk drives combined.

NETEZZA
The Power to Question Everything™

SELECT count (*) , **sex , age** FROM emp WHERE state = 'VA' and age > 18 GROUP BY sex , age ORDER BY age ;

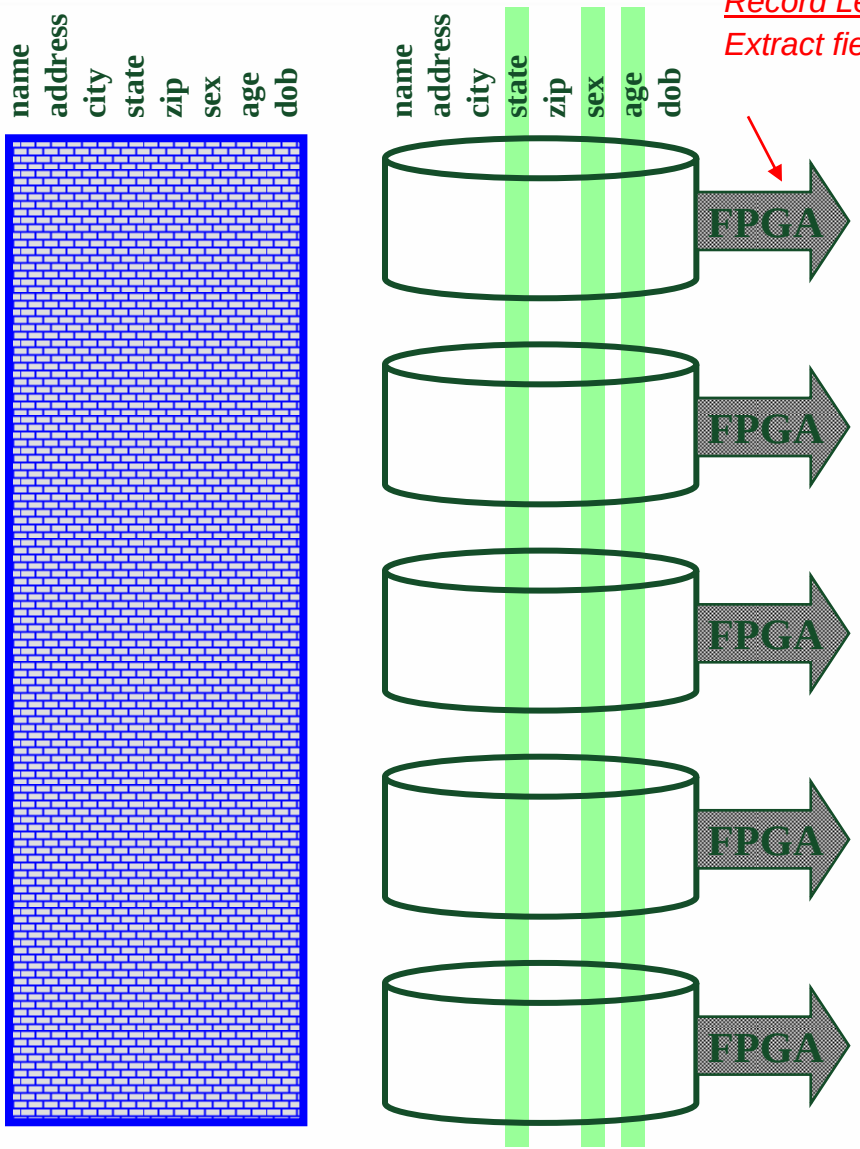

Record Level Operations

Extract fields of interest

PROJECTION

On each SPU, the FPGA / disk controller SELECTs just the columns of interest.

NETEZZA
The Power to Question Everything™

SELECT count (*) , sex , age FROM emp **WHERE state = 'VA' and age > 18** GROUP BY sex , age ORDER BY age ;

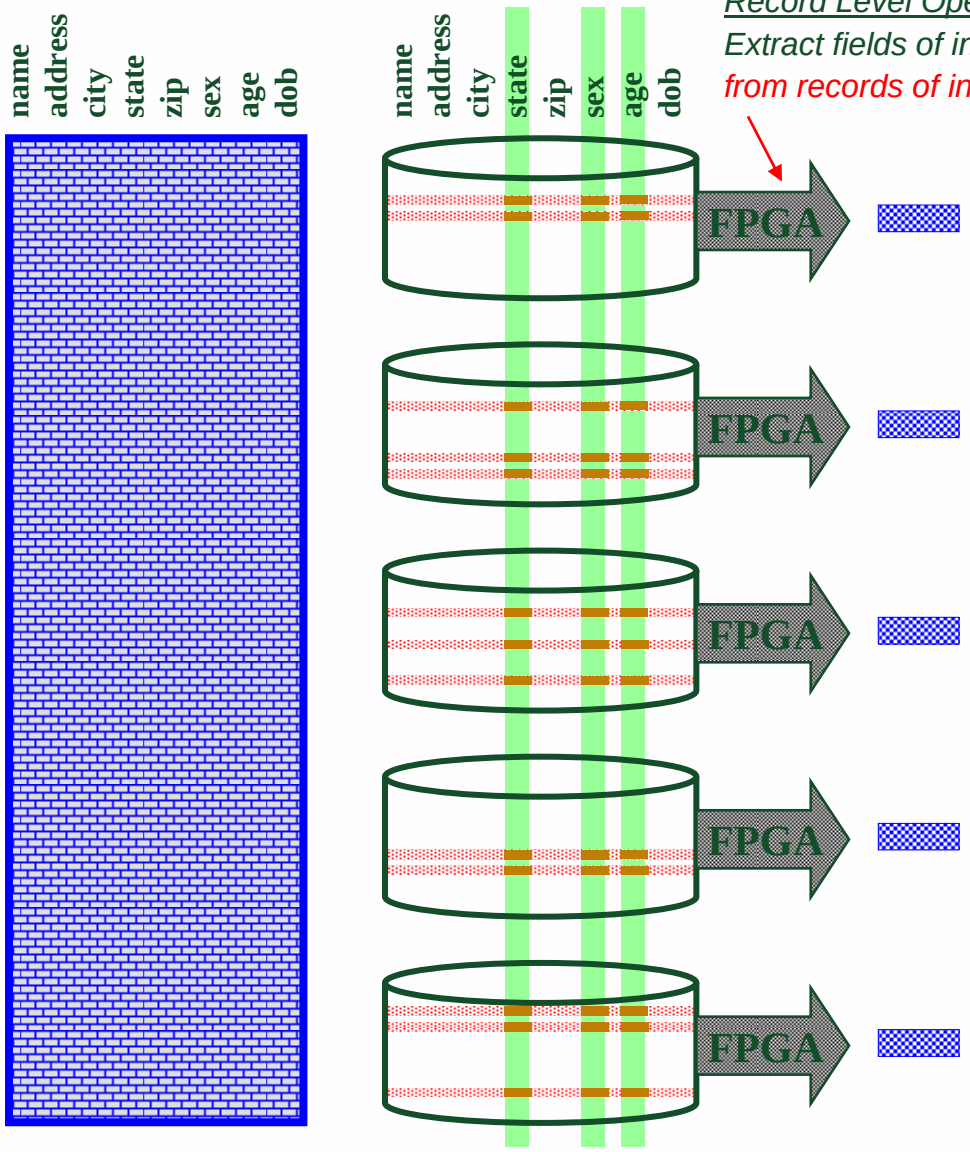

RESTRICTION

The FPGA is also responsible for choosing the records of interest – applying the conditions of the WHERE clause.

NETEZZA
The Power to Question Everything™

SELECT count (*) , sex , age FROM emp WHERE state = 'VA' and age > 18 GROUP BY sex , age ORDER BY age ;

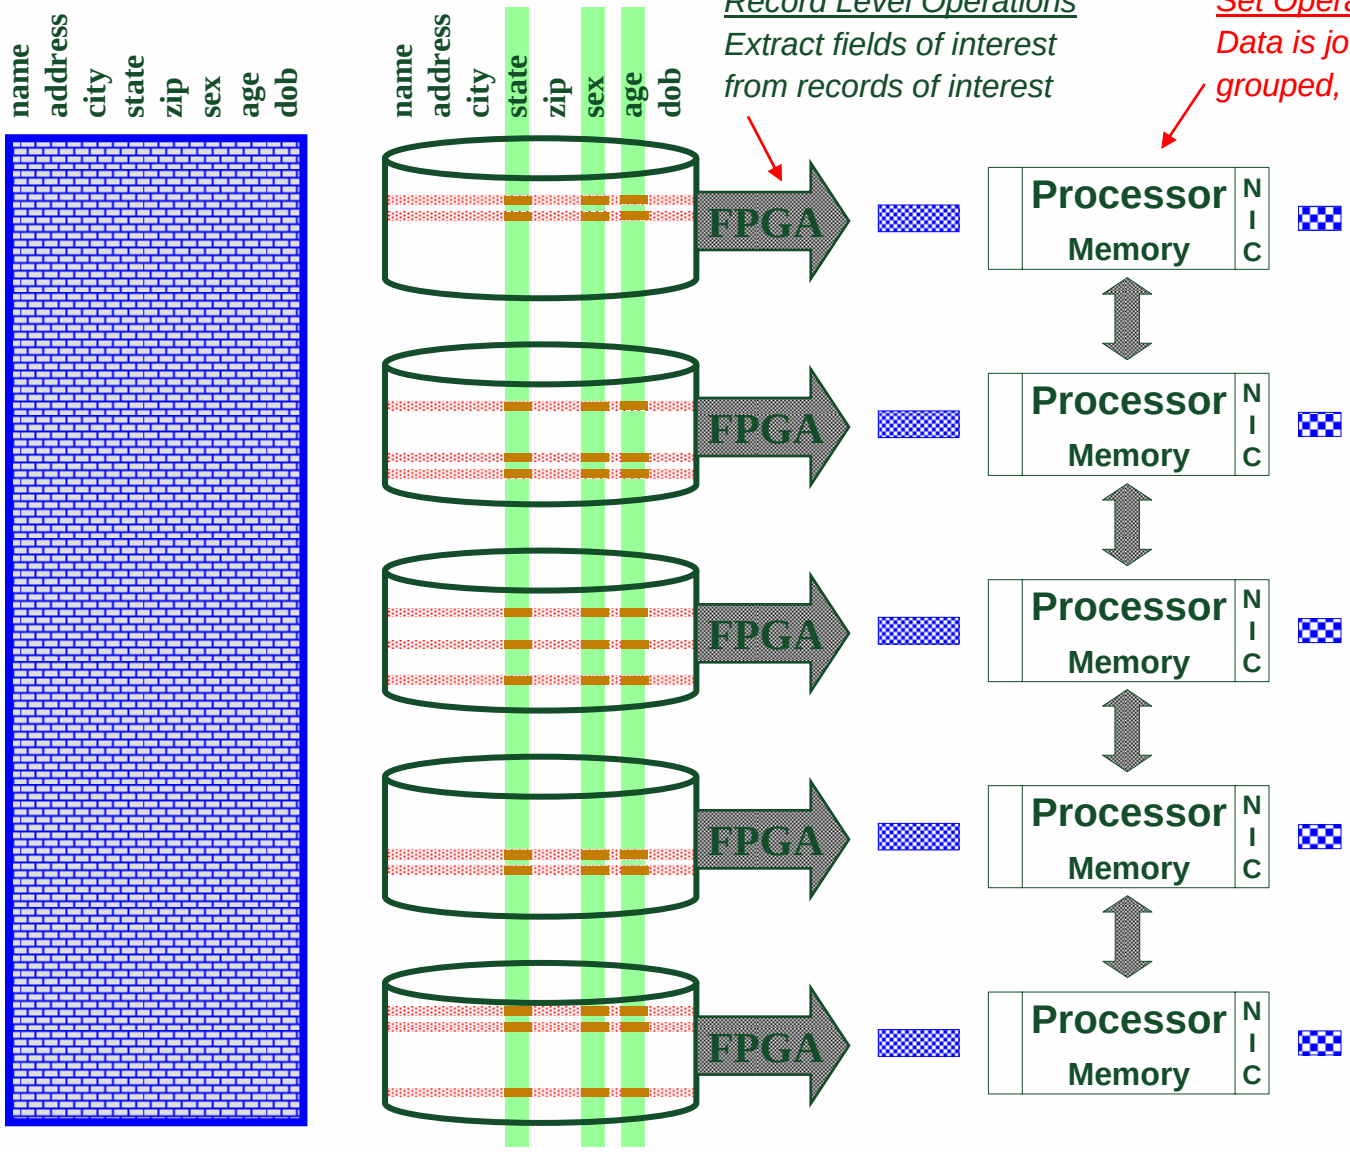

NETEZZA
The Power to Question Everything™

SELECT count (*) , sex , age FROM emp WHERE state = 'VA' and age > 18 GROUP BY sex , age ORDER BY age ;

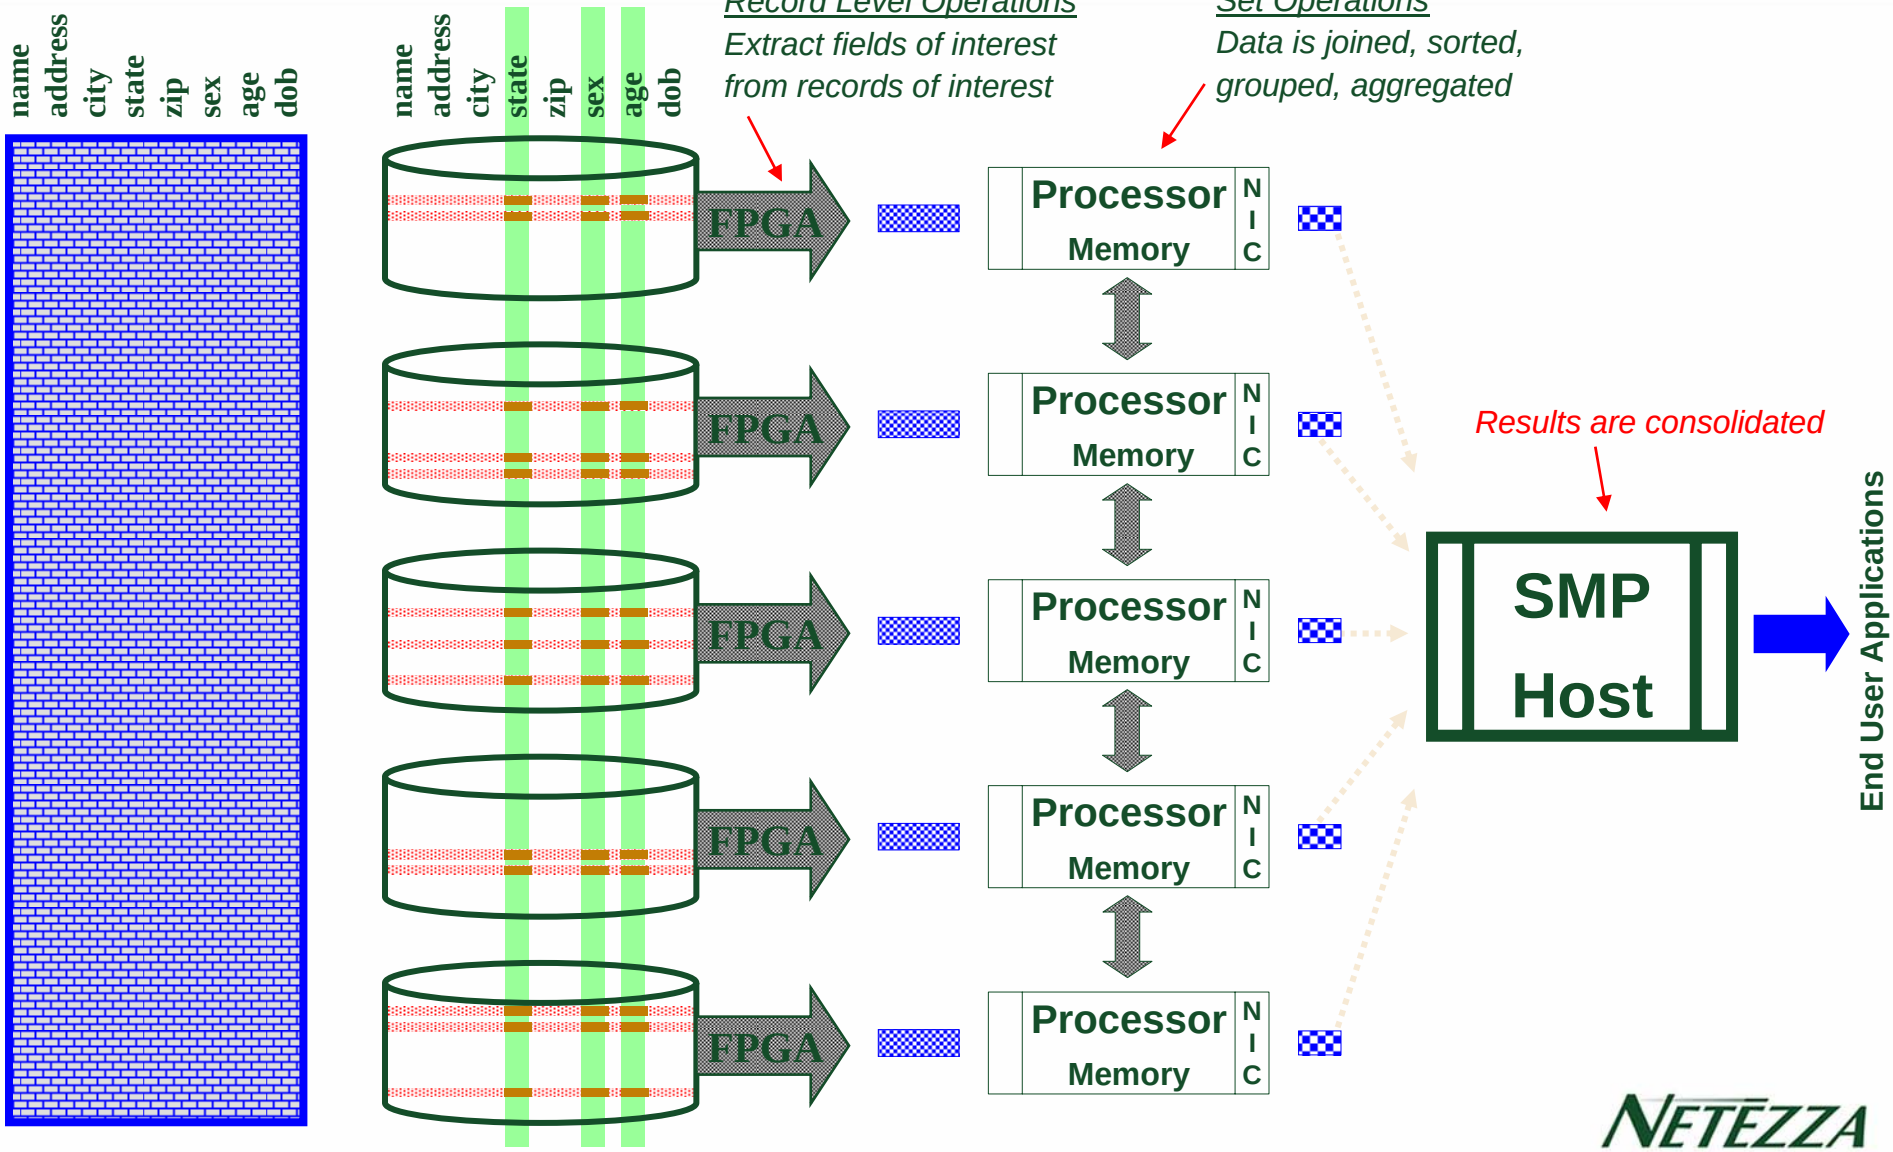

NETEZZA
The Power to Question Everything™

What about scientific data and non-SQL heuristics?

- BLAST is a widely used tool for finding similar sequences in large databases of sequences
- Netezza has integrated the BLAST heuristic algorithms into a new type of SQL Join:

The syntax is an extension of the SQL92 generalized join syntax:
SQL92: SELECT <cols> FROM <t1> <jointype> <t2> ON <join-condition>

The blast join syntax where the controls is a literal string is:

SELECT <cols>
FROM <haystack> [ALIGN <needles>] [WITH <controls>]
ON BLASTX(<haystack.seq>,<needles.seq>,<controls.args>)

Thus a simple literal protein blast looks like:
SELECT <cols> FROM haystack ON BLASTP(haystack.seq, 'ZZAADEDAAM', '-e.001')

Netezza Performance Server® Family
NPS 8000z-Series High-Performance Products

8000z Series:
1-33 TB

	8050z	8150z	8250z	8450z	8650z
Processors	56	112	224	448	672
User Space	2.75 TB	5.5 TB	5.5 – 11 TB	5.5 – 22 TB	5.5 – 33 TB

Continued Innovation in the 8000 Family

- Built & Priced for PERFORMANCE
- Enhanced performance, reliability and system capacity
- Simple, scalable capacity expansion across the product range

NETEZZA
The Power to Question Everything™

Netezza Performance Server® Family
NPS 10000-Series High-Density Products

10000 Series:
Up to 100 TB

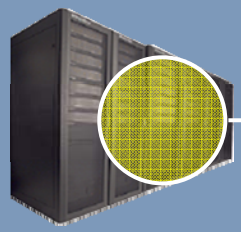

0.9 Terabyte Total DRAM
0.3 Petabyte Total Storage

	10400 HD	10800 HD
Processors	448	896
User Space	10 – 50 TB	10 – 100 TB

Introducing the 10000 Product Family

- HIGH PERFORMANCE & HIGH DATA DENSITY in a single NPS appliance
- Simple, cost-effective, scalable capacity expansion
- Up to 12.5 TB of user space per rack

NETEZZA
The Power to Question Everything™

Some of Our Customers by Vertical Market…

Retail

Ahold USA

Albertsons

Catalina Marketing

Debenhams

Neiman Marcus

Restoration Hardware

Ross Dress for Less

Ryder

Shoppers Drug Mart

Sysco

Telecom

at&t

The Carphone Warehouse …for a better mobile life

caudwell communications

cingular raising the bar

iBasis

Sprint

orange

Rogers

Telus

Online

Advertising.com

amazon.com

America Online

CNET Networks

Financial Services

Bank of America

Capital One

Loan Performance

Nationwide

Option One Mortgage

Analytic Services

ACNielsen

Acxiom

Clarity Blue Customer Intelligence

epsilon

Merkle

Healthcare

American Red Cross

NDCHEALTH

Premier

Other

CenterParcs

Proven Results:
Leading Food & Drug Retailer

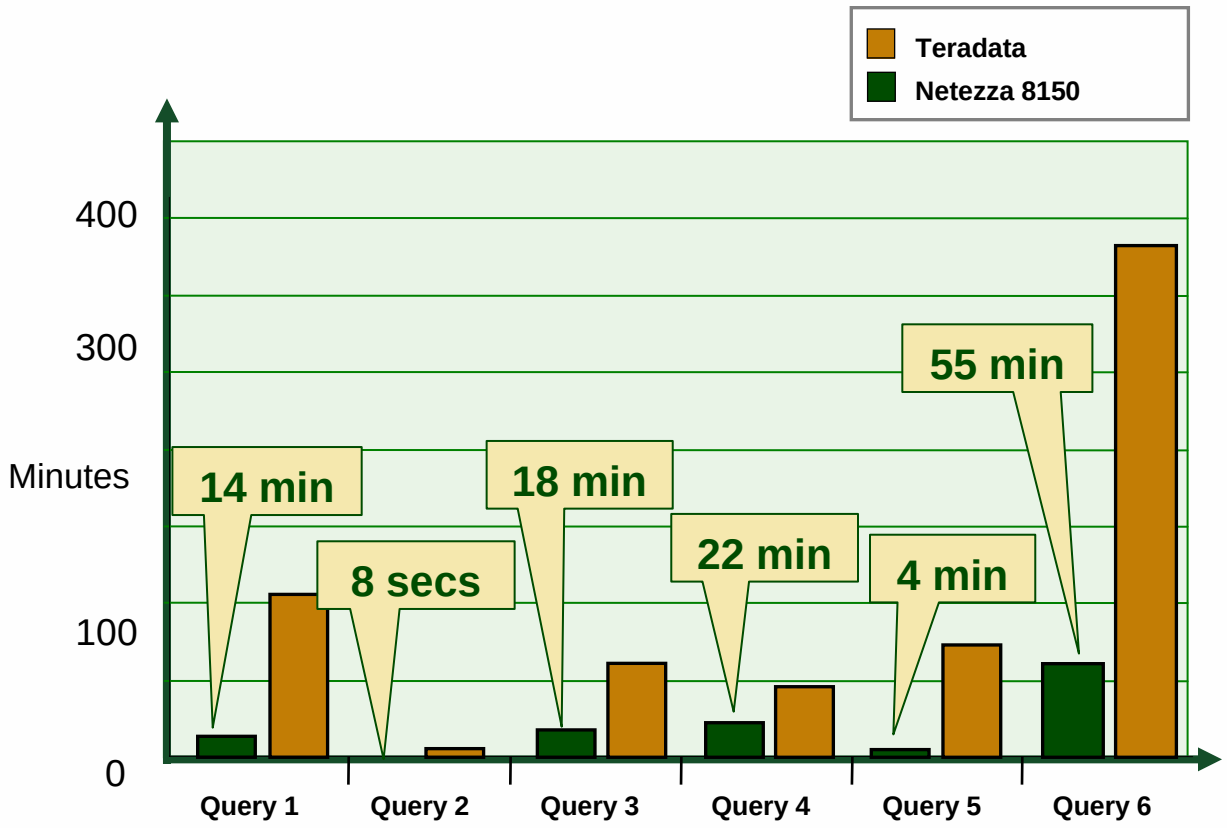

**Netezza results based on an NPS 8150.*
Teradata queries run on a 40-node system (5200 &5300)

Situation

- Competed against a **Teradata system 25x more expensive**
- Total amt of data loaded: 3.5TB
- Time to load: 28 hours
- 118GB/hour
- Queries included market basket penetration, Y/Y comparisons, top UPC by movements, price optimization tracking, etc.
- Running SQL

Results

- NPS system went from loading dock to installed, configured and running in three hours
- Queries were run substantially faster, including one that was over five times faster

NETEZZA
The Power to Question Everything™

Proven Results:
Report Execution from a Government POC

Situation

- Competed against a very large Teradata system (96 nodes)
- Total amount of data loaded: 2.5 TB
- 200+GB/hour load rates (single stream)
- Representative set of resource-intensive production MicroStrategy reports

Impact

- NPS system went from loading dock to installed, configured and running in five hours
- Queries showed substantial improvement on Netezza – **15 times faster** on average!
- Total execution time (13 reports) was ~7 ½ hours on TD vs. only 47 min on Netezza

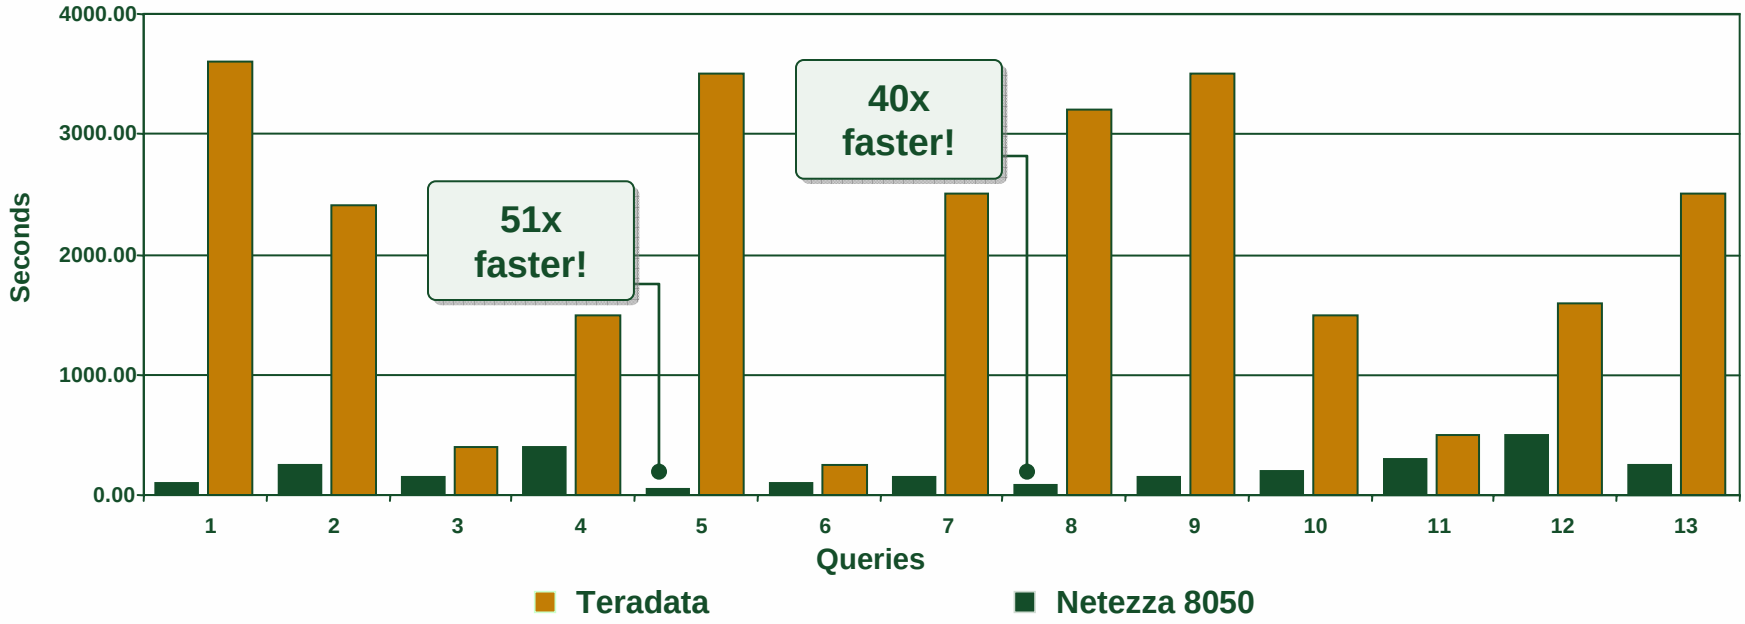

* *Netezza results based on an NPS 8250.*
Teradata queries run on a 96-node system (52xx and 53xx)

NETEZZA
The Power to Question Everything™

Proven Results:
Analytic Service Provider

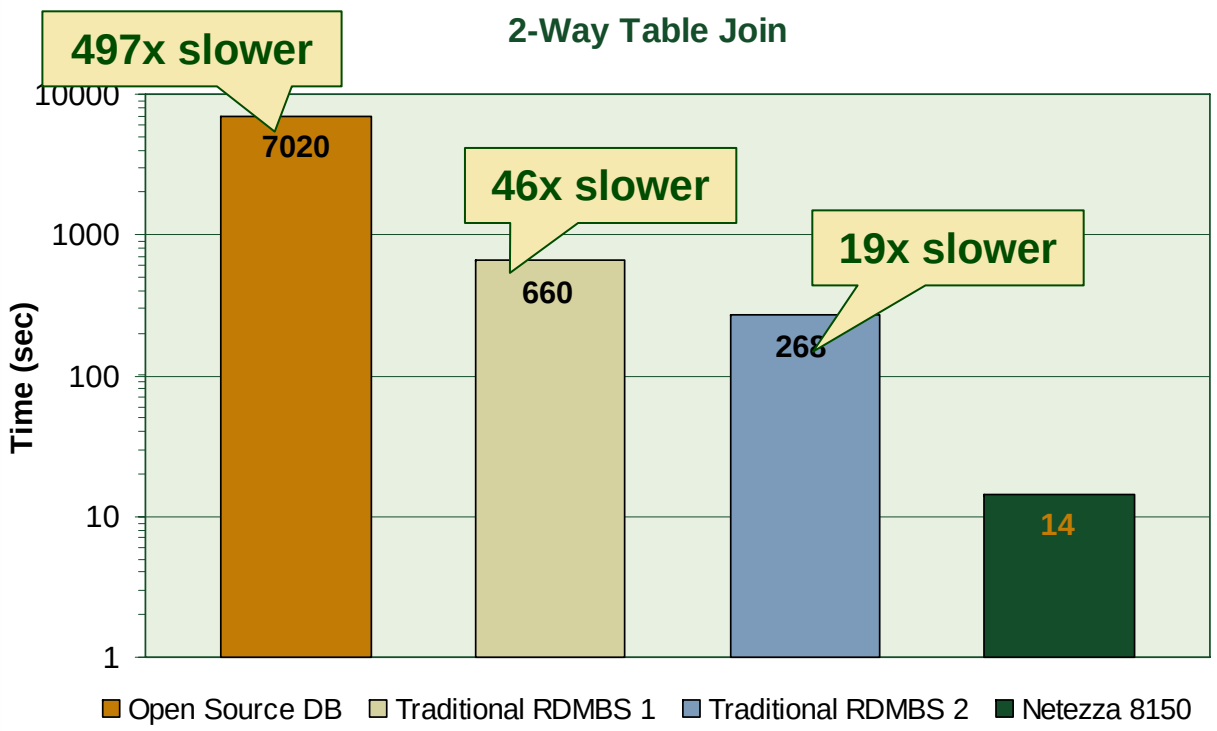

POC Performance

- 2-way Cartesian Join Mixed Read/Write Test
- 37.6M rows with 122.9M rows
- Performance Improvement with Netezza
  - 497x v. Open Source DB
  - 46x v. Traditional RDBMS 1
  - 19x v. Traditional RDBMS 2

NETEZZA
The Power to Question Everything™

Proven Results: E-Business Customer

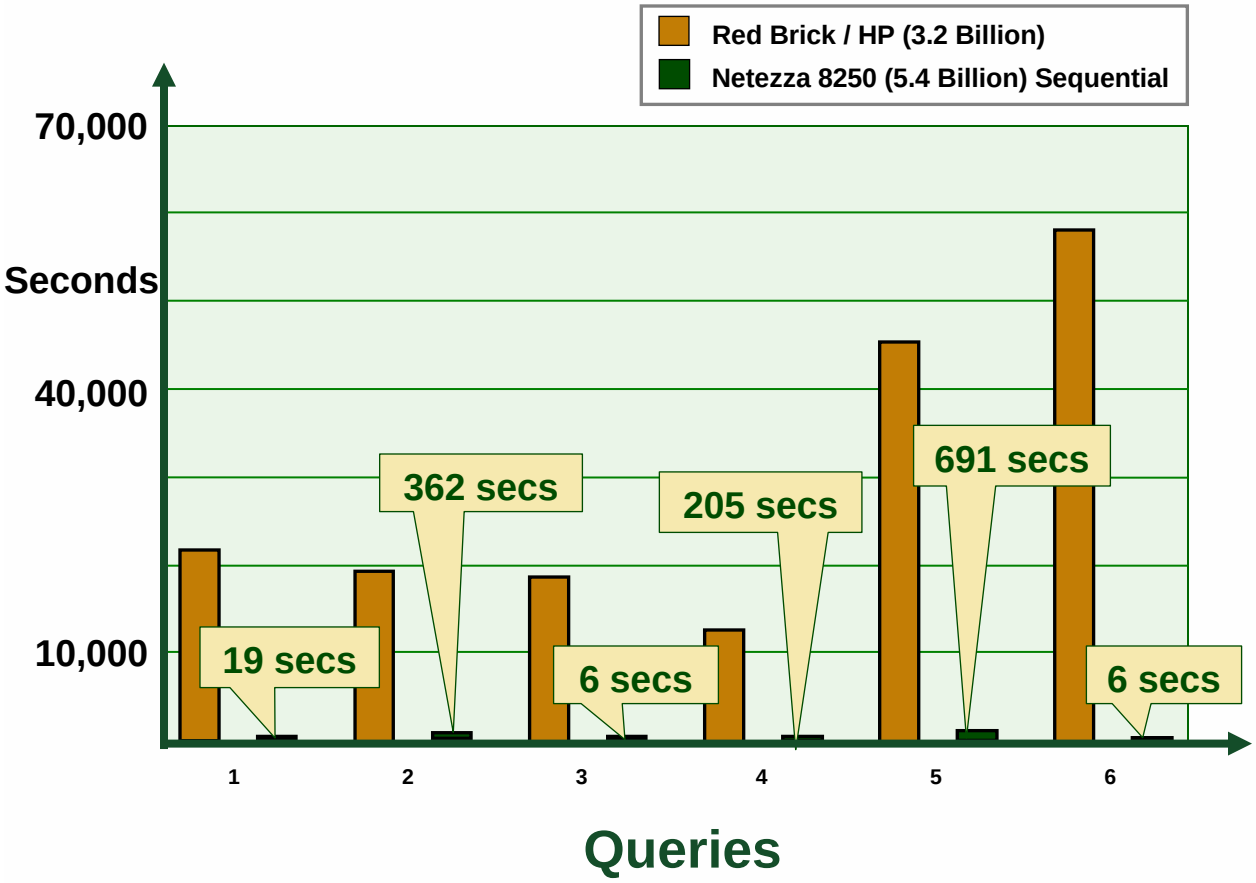

Situation

- Red Brick: 3.2 billion rows
- NPS: 5.4 billion rows
- 6 queries–load, expansion and test
- Business Objects and SQL

Query Performance

- NPS system handled 69% more data volume but was able to complete the total workload in 21 minutes vs. 50 hours, 143x faster!

Load Performance

- 140+ GB/hr

**Netezza results based on an NPS 8250.*
Red Brick results on HP SuperDome 32 CPU/32GB RAM and EMC SAN

NETEZZA
The Power to Question Everything™

Netezza's DW System Architecture Evolution

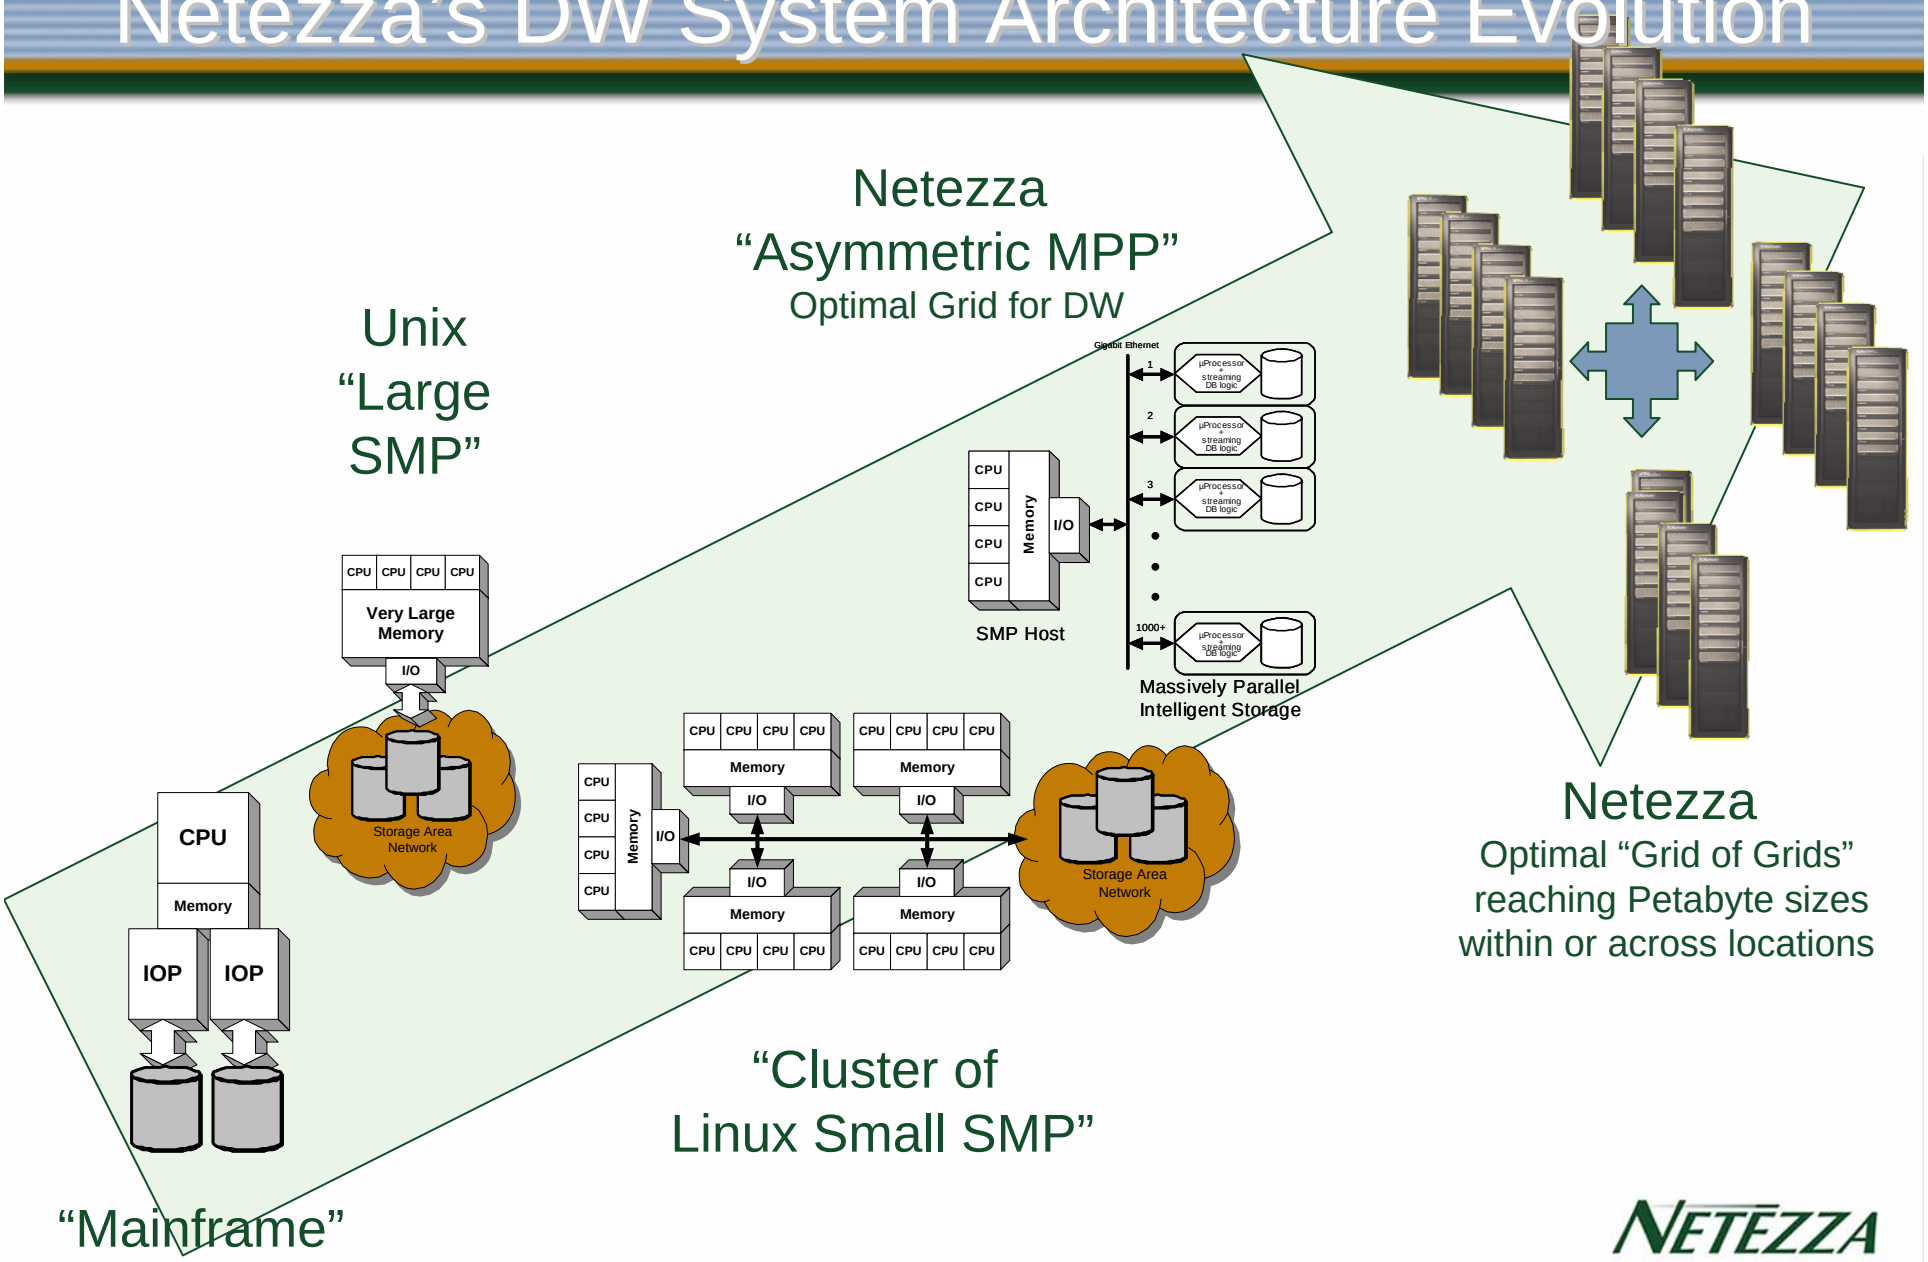

NETEZZA
The Power to Question Everything™

NETEZZA

The Power to Question Everything™

Thank You!