

# IBM Performance 2012

Smarter Analytics. Smarter Outcomes.




## The 4 S's of Netezza *(aka "PureData Systems for Analytics")*

Falke Van Onacker

Sales Executive for Big Data and Netezza

La Hulpe - Room F - November 13, 2012





## Contact details of the speaker



*Falke Van Onacker*

*Solution Sales for Big Data and Netezza  
IBM Belgium & Luxembourg*

Mobile: [0496/57.59.13](tel:0496575913)

E-Mail: [falke.van.onacker@be.ibm.com](mailto:falke.van.onacker@be.ibm.com)

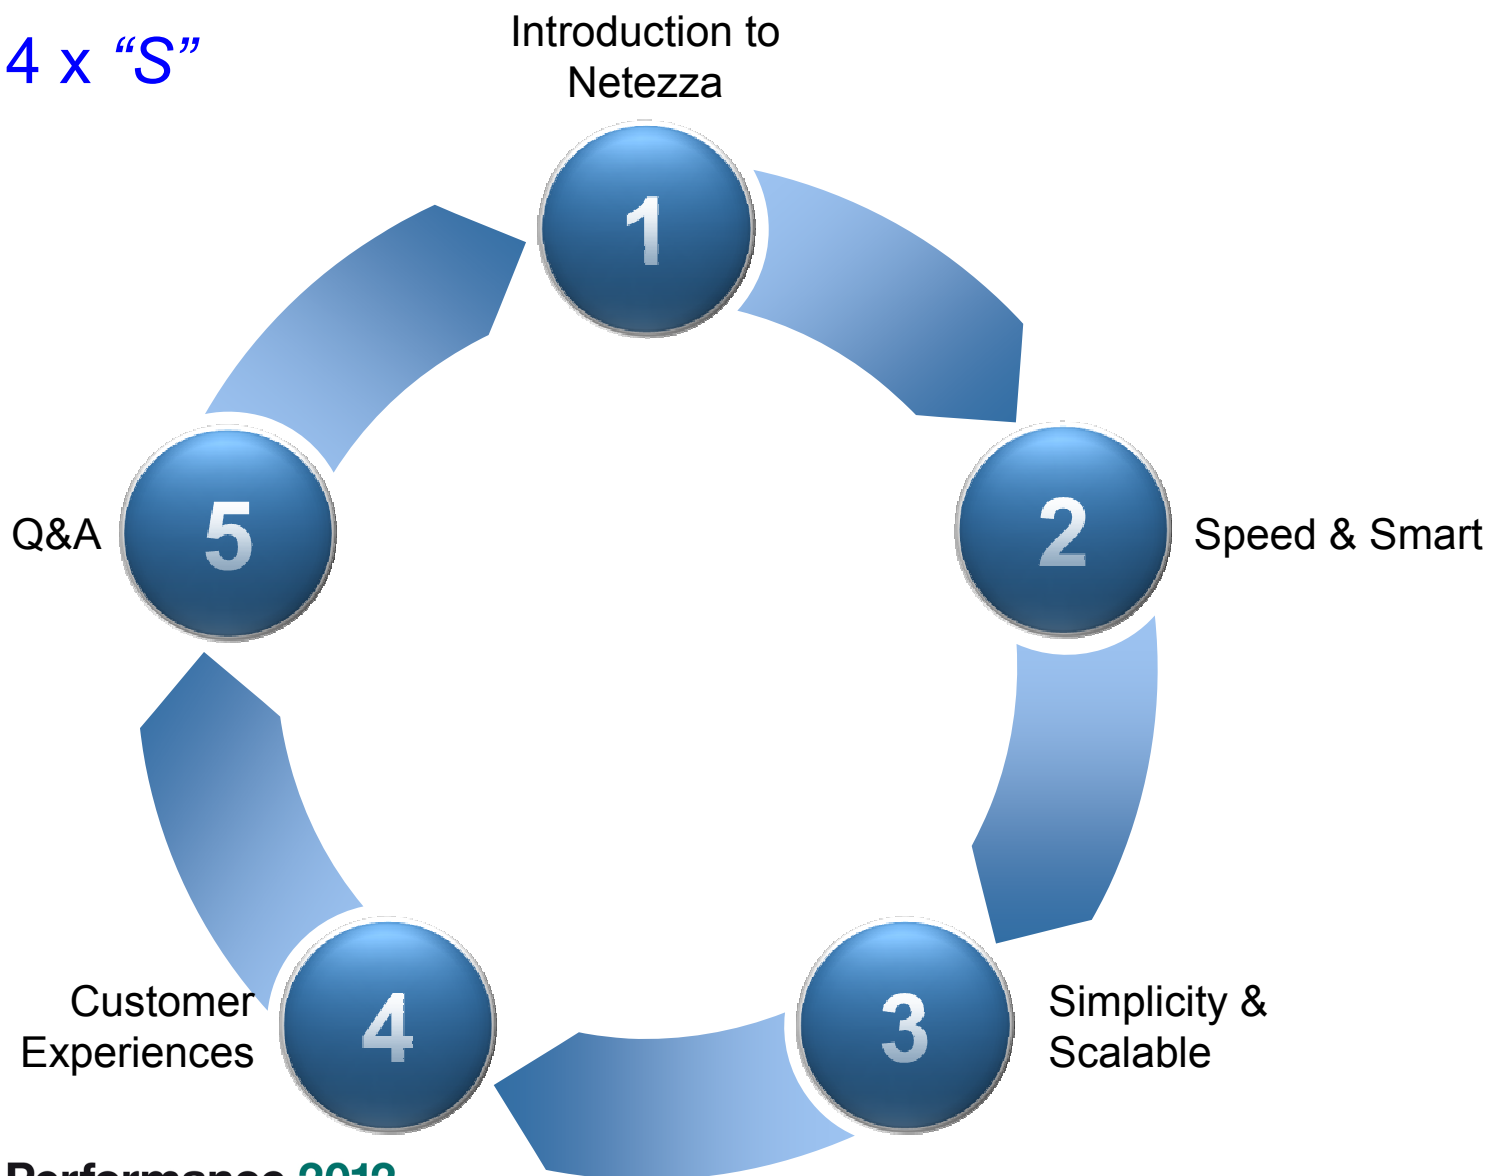
LinkedIn: <http://www.linkedin.com/in/falke>



**IBM Performance 2012**  
Smarter Analytics. Smarter Outcomes.



# Netezza 4 x "S"





# Netezza 4 x “S”

## Introduction to Netezza




# The IBM Netezza Appliance – Revolutionizing Analytics




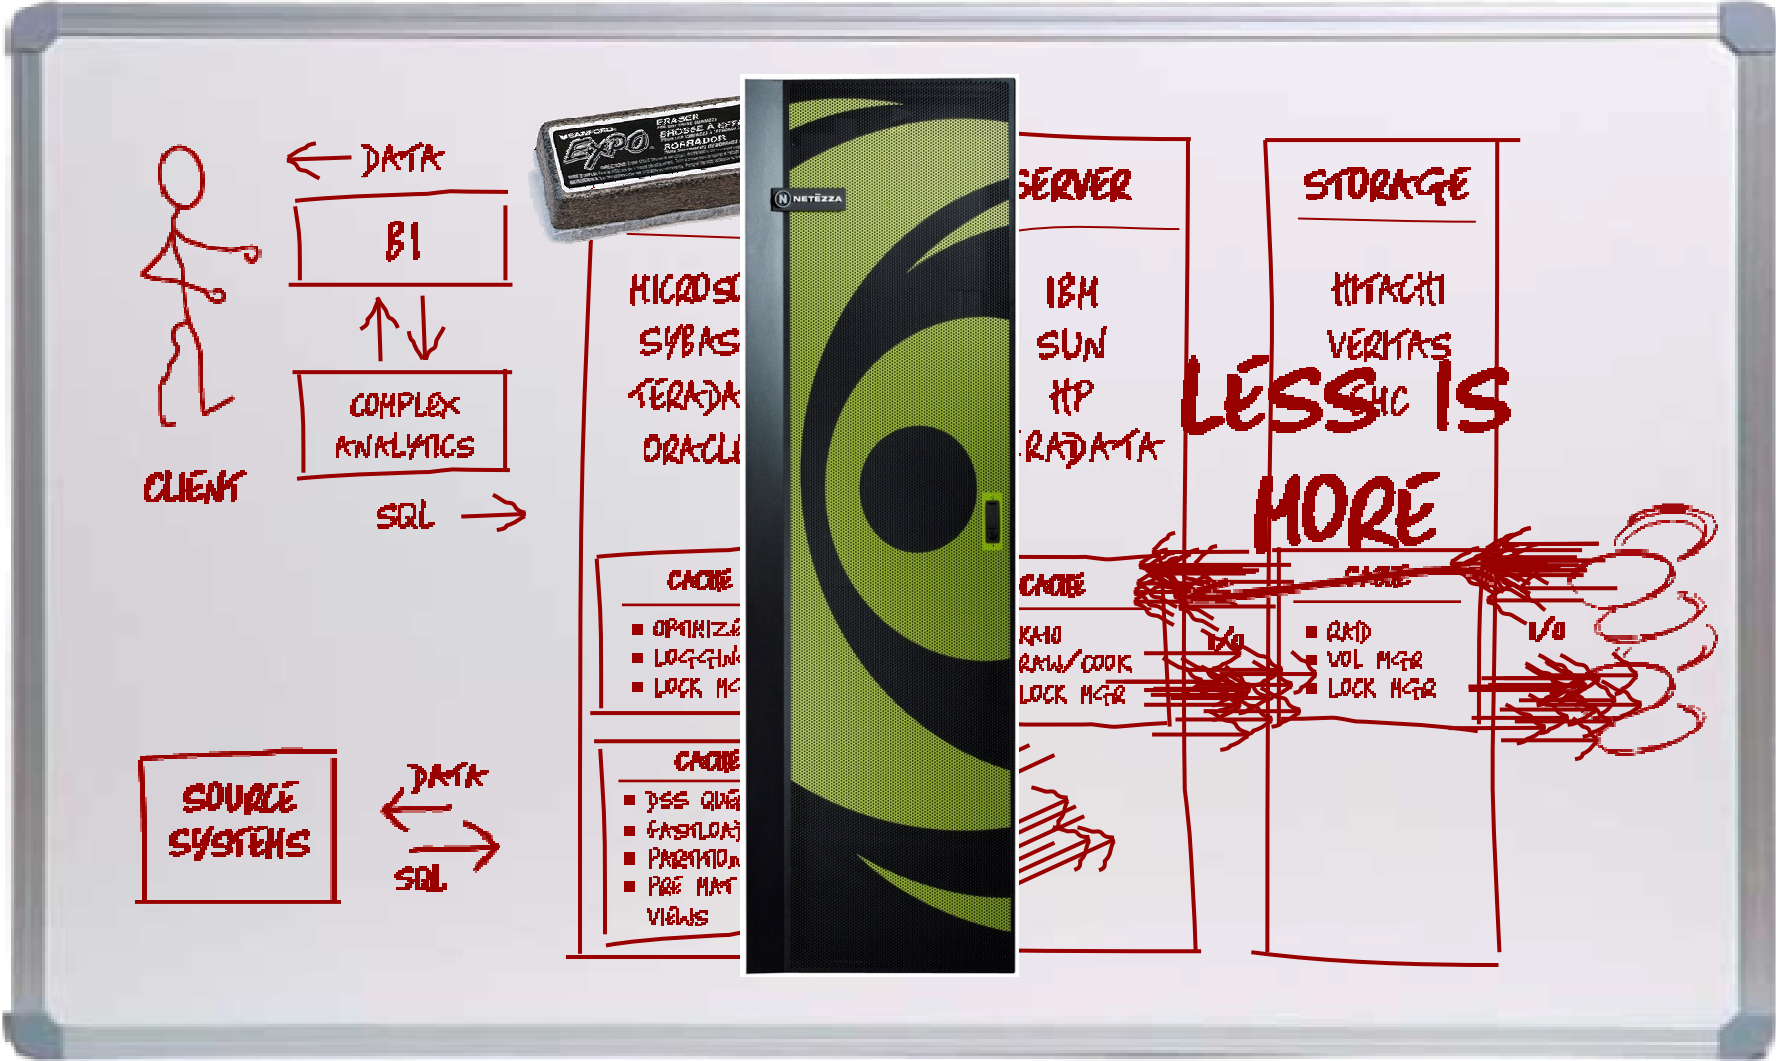
- **Purpose-Built** Analytics Appliance
- Integrated database, server & storage
- Open, thanks to standard interfaces
- Low Total Cost of Ownership
- **Speed:** 10-100x faster than traditional systems
- **Simplicity:** “real” Appliance – minimal administration & tuning
- **Scalability:** From 50 Gigabytes up to > Petabytes
- **Smart:** High-performance Advanced Analytics



# Traditional Data Warehouse Complexity



# Data Warehousing – Simplified



# Unique Value Proposition of IBM Netezza



## Operations

- *There's nothing to do ... it's an Analytics Appliance*

## BI Developers & DBAs – faster delivery times

- No configuration
- No physical modeling
- No indexes
- No tuning – out of the box performance
- *Data model agnostic*

## ETL Developers

- No aggregate tables needed – simpler ETL logic
- Faster load and transformation times - *ELT*

## Business Analysts

- *Performance that enables new types of analysis – 10 to 100x faster than legacy systems*
- *True ad hoc* queries
- Ask complex queries against large datasets
- Lower latency – *analysis on “fresh” data*



**IBM Performance 2012**  
Smarter Analytics. Smarter Outcomes.



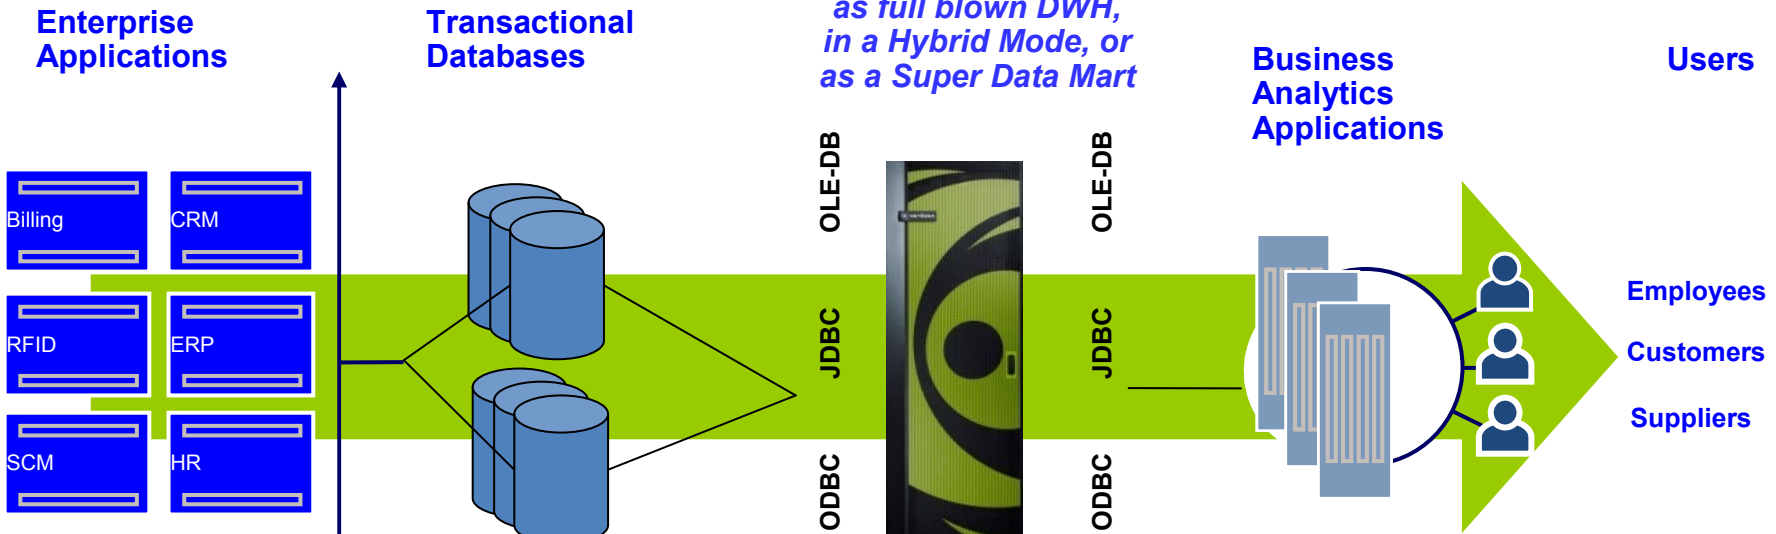
# IBM Netezza in the IT Ecosystem



*Open standards platform to enable actionable intelligence "out of the box"*

## Netezza Appliance

*as full blown DWH, in a Hybrid Mode, or as a Super Data Mart*



Representative vendors  
**ORACLE**  
**IBM**

**NETEZZA**  
 an IBM Company

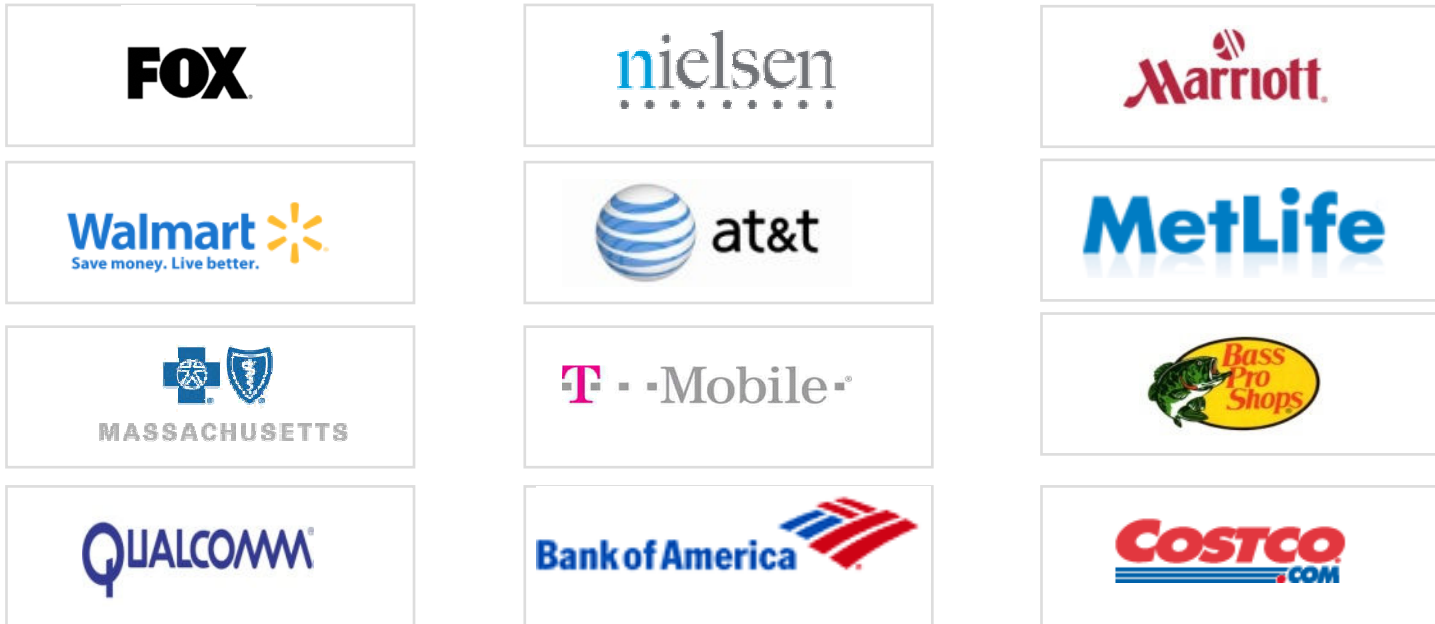
**MicroStrategy** **BUSINESS OBJECTS**  
**sas** **QlikView** **SPSS**  
**COGNOS** **Microsoft**  
**ORACLE**

- Key Advantages**
- ▶ Performance
  - ▶ Scalability
  - ▶ Low TCO




**IBM Performance 2012**  
 Smarter Analytics. Smarter Outcomes.

# More than 85 clients already use Cognos with Netezza



- 2x faster** in report execution
- At least 60% faster** in SQL execution
- 40% faster** in reports with analysis functions
- 2x faster** in drill-down
- Instantaneous** drilling up is instantaneous because data is in cache

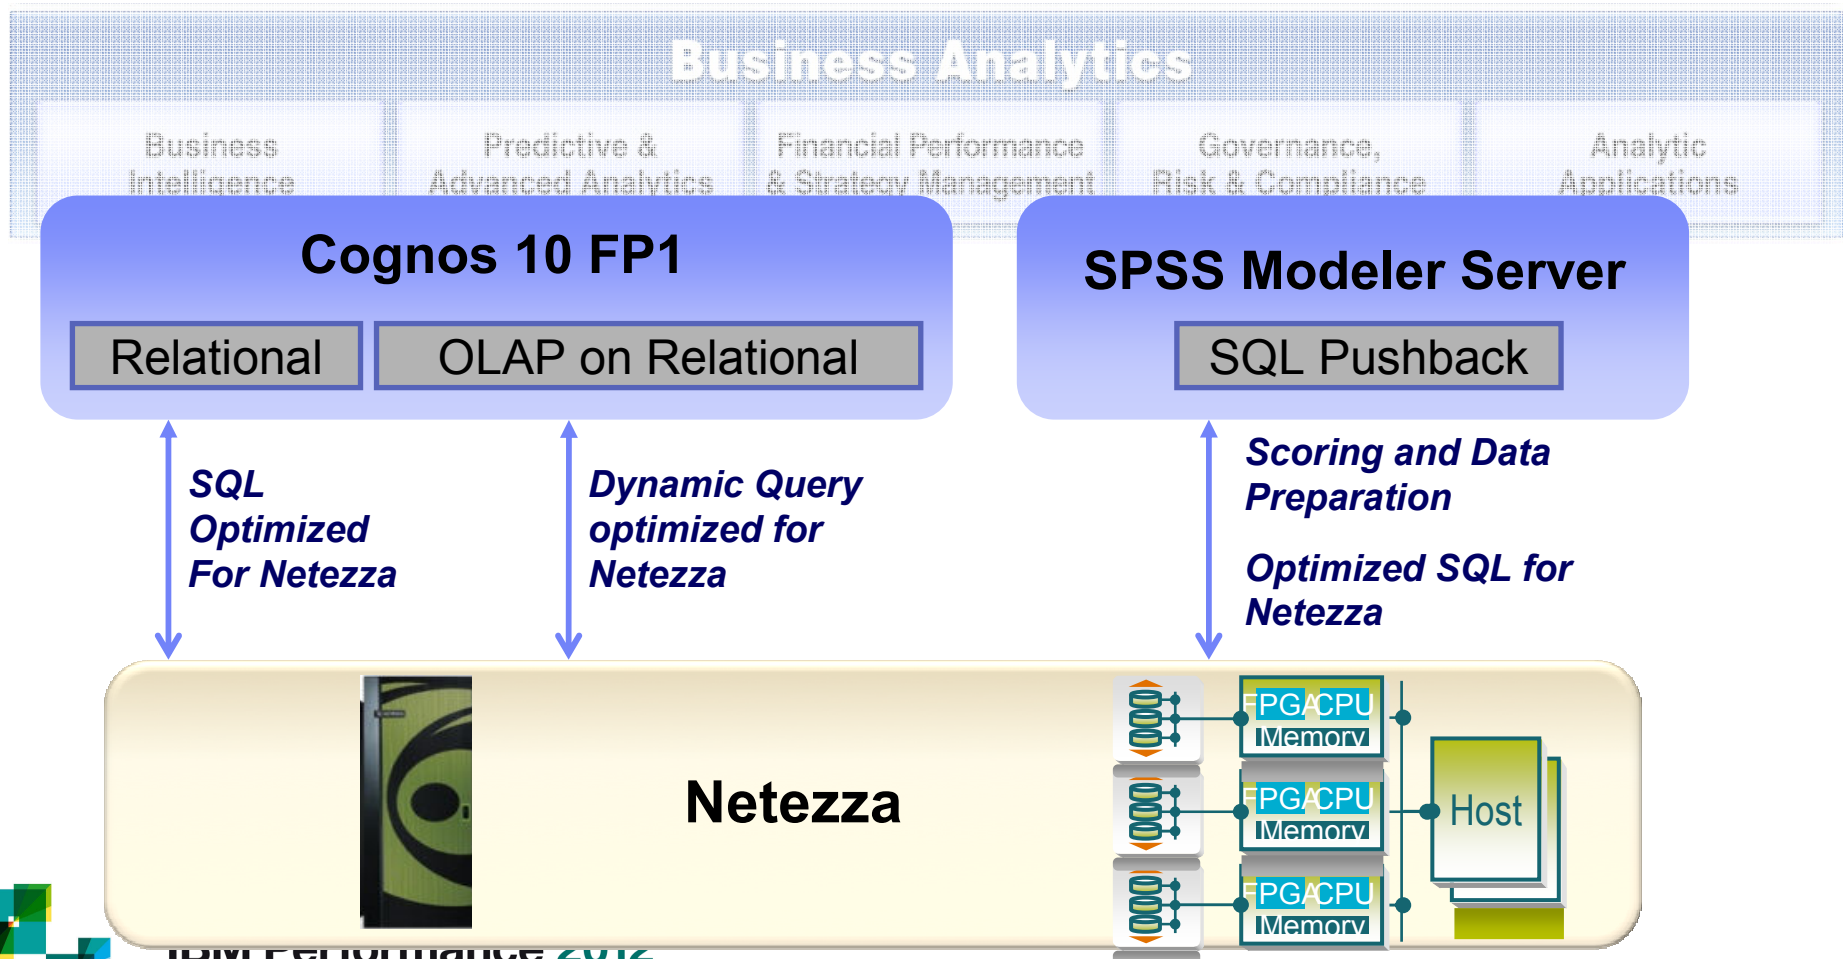


# IBM Business Analytics Optimizations for Netezza

*exploiting the MPP capabilities of Netezza*

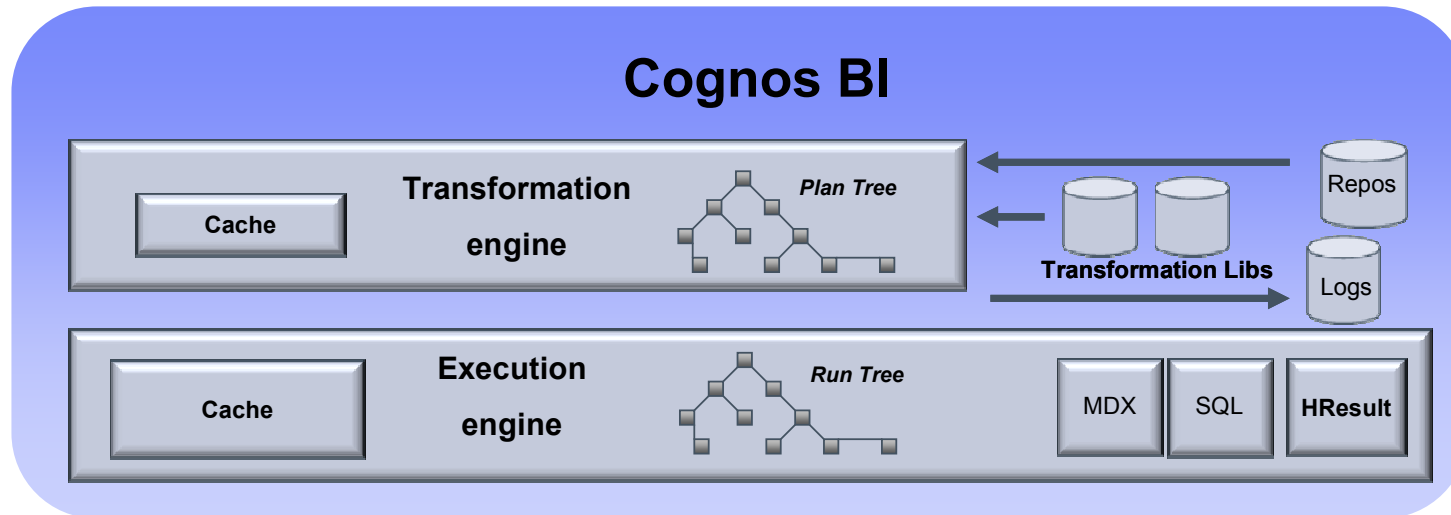


## Optimized for large scale analytics





# Netezza optimizations for Cognos BI



## Fastest Query Response

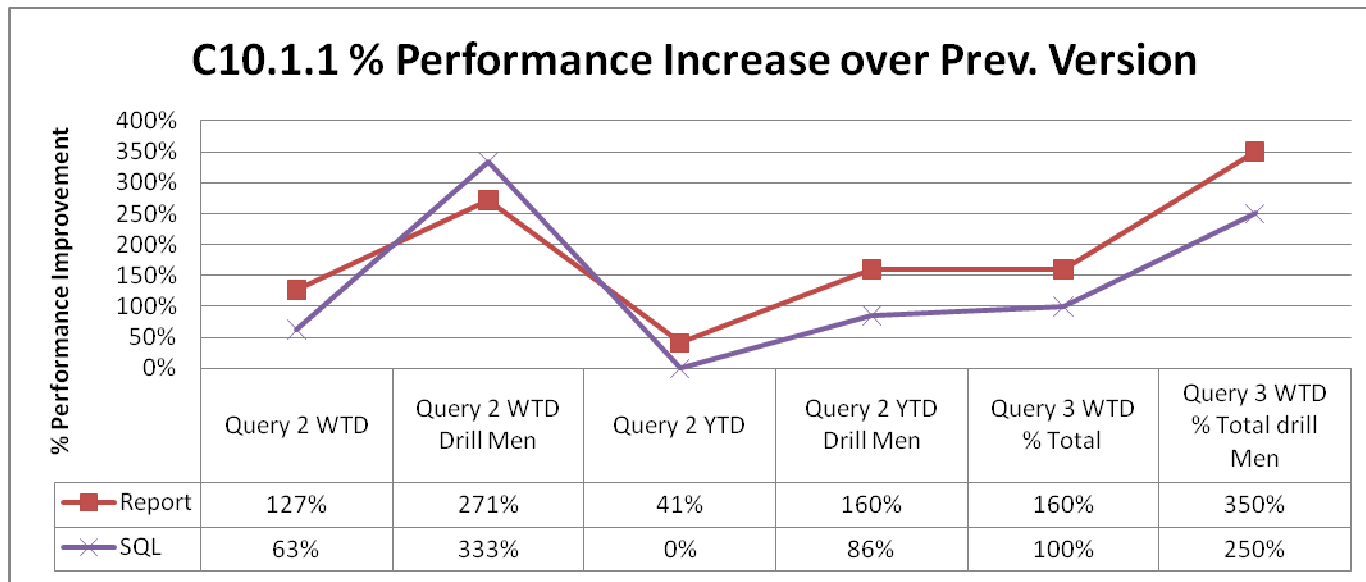
- Enhanced SQL
- Dynamic Query - fine grained secure caching
- Optimized for aggregation
- Multi-pass SQL leverages multi-parallel processing capabilities of Netezza



**IBM Performance 2012**  
Smarter Analytics. Smarter Outcomes.

# Netezza with Cognos BI 10.1.1

*A Great User Experience requires fast Performance*



*Analysis is:*

- Cognos BI V10.1.1 scales significantly better with larger data scenario
- Cognos BI V10.1.1 reduces the SQL load on the Netezza server due to caching

\* Based on findings in the lab using Cognos Business Intelligence V10.1.1 with Dynamic Query Mode compared to previous versions of Cognos Business Intelligence. TPC-DS sample data was used.



**IBM Performance 2012**

Smarter Analytics. Smarter Outcomes.



# Netezza 4 x "S"

Introduction to  
Netezza



Speed & Smart



**IBM Performance 2012**  
Smarter Analytics. Smarter Outcomes.

# IBM Netezza Appliance Approach

Out-of-the-box Speed

Built-in Statistical Libraries




## Optimized Hardware + Software

- Purpose-built for high performance analytics
- Requires no tuning

## True MPP

- All processors fully utilized
- Maximum speed and efficiency



## Streaming Data

- Hardware-based query acceleration **FPGA**
- Blistering fast results

## Deep Analytics

- Complex analytics executed in-database
- Deeper insights



**IBM Performance 2012**

Smarter Analytics. Smarter Outcomes.

© 2012 IBM Corporation

# Example of out-of-the box Speed



## Unica on Netezza: PoC results of a Belgian Client

### NetInsight response times on Current environment with Oracle

| Reports                            | dataset | start time | end time | delta    |
|------------------------------------|---------|------------|----------|----------|
| [redacted]                         | January | 10h00      | 12h45    | 02h45    |
| [redacted]                         | January | 10h01      | 13h04    | 03h03    |
| [redacted] Trended per day         | January | 10h02      | 10h57    | 00h55    |
| Campaign Summary per Channel       | January | 10h03      | 10h54    | 00h51    |
| Traffic per Day for [redacted]     | January | 10:01:21   | 10:53:23 | 00:52:02 |
| Traffic Sources for [redacted]     | January | 10:01:36   | 10:51:10 | 00:49:34 |
| Traffic per Month for Support Site | January | 10:01:46   | 10:53:23 | 00:51:37 |
| Events for [redacted]              | January | 10:01:54   | 10:59:26 | 00:57:32 |

#### Explanation of above chart:

To map the setup used by [redacted] for the Unica / Oracle test, the 8 reports will be executed simultaneously using 2 users (each user will issue 4 reports). Timings will be derived from the NetInsight application (start time = time of report launch, end time = report presented in application)



**IBM Performance 2012**

Smarter Analytics. Smarter Outcomes.

© 2012 IBM Corporation



# Example of out-of-the box Speed



## Unica on Netezza: PoC results of a Belgian Client

### NetInsight Response Times on Netezza

| Reports                            | dataset                  | Test on Dec. 5 at 15:26 | Test on Dec. 5 at 16:00 | Test on Dec. 5 at 16:05 | Test on Dec. 6 at 9:37 | Test on Dec. 6 at 17:40 |
|------------------------------------|--------------------------|-------------------------|-------------------------|-------------------------|------------------------|-------------------------|
| [redacted] Traffic                 | January                  | <1min                   | <1min                   |                         | <1min                  | <1min                   |
| [redacted] Monthly Traffic         | January                  | <1min                   | <1min                   |                         | <1min                  | <1min                   |
| [redacted] Trended per day         | January                  | <1min                   | <1min                   |                         | <1min                  | <1min                   |
| Scenario Funnel Example            | January                  | <1min                   | <1min                   |                         | <1min                  | <1min                   |
| Traffic per Day for [redacted]     | January                  | <1min                   | <1min                   |                         | <1min                  | <1min                   |
| Traffic Sources for [redacted]     | January                  | <1min                   | <1min                   |                         | <1min                  | <1min                   |
| Traffic per Month for Support Site | January                  | <1min                   | <1min                   |                         | <1min                  | <1min                   |
| Events for [redacted]              | January                  | <1min                   | <1min                   |                         | <1min                  | <1min                   |
| All the above reports              | 01-12-2010 to 31-05-2011 |                         |                         | <2 min                  |                        |                         |





# IBM Netezza Built-in Statistical Library

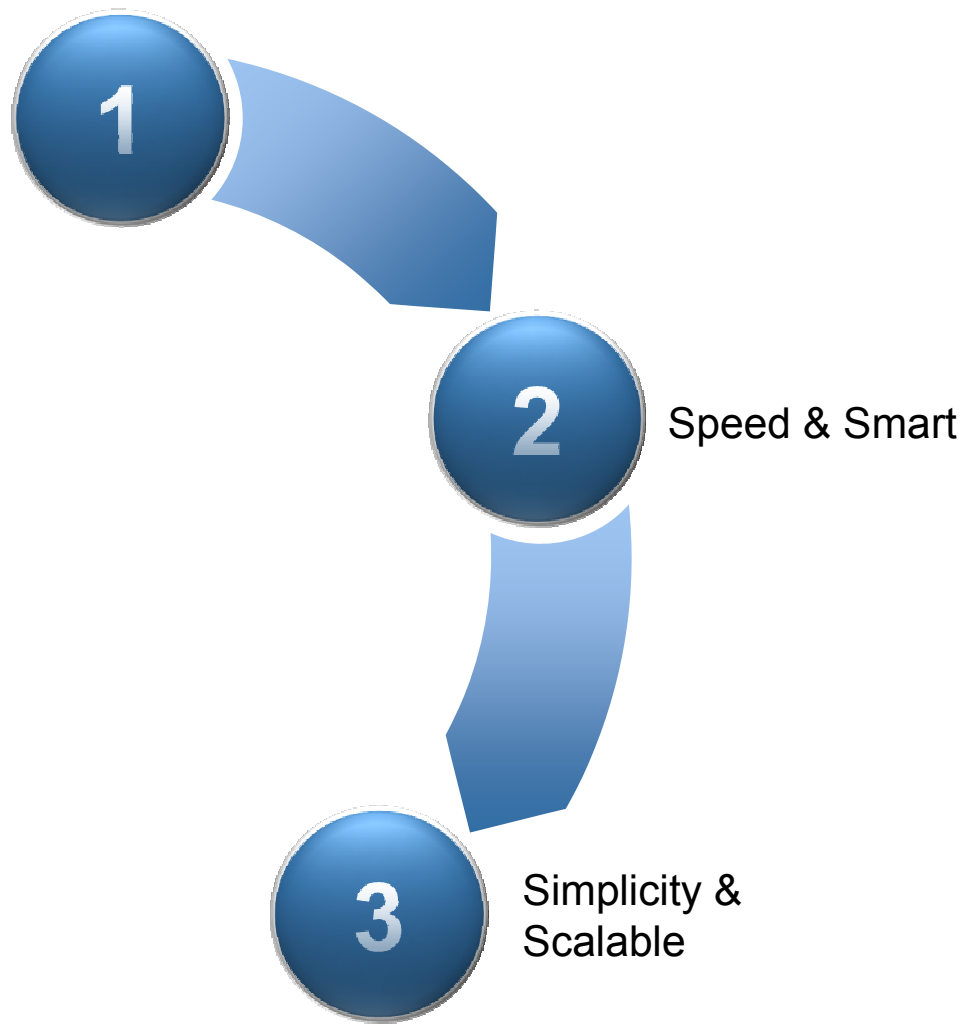
## Extensive In-Database Analytics Library for Big Data Crunching





# Netezza 4 x “S”

Introduction to  
Netezza



**IBM Performance 2012**  
Smarter Analytics. Smarter Outcomes.




## Appliance Simplicity

- **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
- **No software installation**

Resources become Data Managers instead of Database Administrators

Become a Netezza Expert with 2 trainings, in 1 week.





# Traditional Complexity

...

# Netezza Simplicity

## ORACLE Indexes

```
CREATE INDEX "MRDWDDM"."RDWF_DDM_ROOMS_SOLD_IDX1" ON "RDWF_DDM_ROOMS_SOLD"
("ID_PROPERTY" , "ID_DATE_STAY" , "CD_ROOM_POOL" , "CD_RATE_PGM" ,
"CD_RATE_TYPE" , "CD_MARKET_SEGMENT" ) PCTFREE 10 INITRANS 6 MAXTRANS 255
STORAGE( FREELISTS 10) TABLESPACE "DDM_DATAMART_INDEX_L" NOLOGGING
PARALLEL ( DEGREE 4 INSTANCES 1) LOCAL(PARTITION "PART1" PCTFREE 10
INITRANS 6 MAXTRANS 255 STORAGE(INITIAL 4194304 NEXT 4259840 MINEXTENTS 1
MAXEXTENTS 100000 PCTINCREASE 0 FREELISTS 10 FREELIST GROUPS 1 BUFFER_POOL
DEFAULT) TABLESPACE "DDM_DATAMART_INDEX_L" NOLOGGING, PARTITION "PART2"
PCTFREE 10 INITRANS 6 MAXTRANS 255 STORAGE(INITIAL 4194304 NEXT 4259840
MINEXTENTS 1 MAXEXTENTS 100000 PCTINCREASE 0 FREELISTS 10 FREELIST GROUPS
1 BUFFER_POOL DE
```

## ORACLE Bitmap index

```
CREATE BITMAP INDEX "CRDBO"."SNAPSHOT_MONTH_IDX13" ON
"SNAPSHOT_OPPTY_MONTH_HIST" ("SNAPSHOT_YEAR" ) PCTFREE 10 INITRANS 2
MAXTRANS 255 STORAGE(INITIAL 4194304 NEXT 4194304 MINEXTENTS 2 MAXEXTENTS
2147483645 PCTINCREASE 0 FREELISTS 1 FREELIST GROUPS 1 BUFFER_POOL
DEFAULT) TABLESPACE "SFA_DATAMART_INDEX" NOLOGGING ;
```

## ORACLE Table Clusters

```
CREATE CLUSTER "MRDW"."CT_INTRMDRY_CAL" ("ID_YEAR_CAL" NUMBER(4, 0),
"ID_MONTH_CAL" NUMBER(2, 0), "ID_PROPERTY" NUMBER(5, 0)) SIZE 16384
PCTFREE 10 PCTUSED 90 INITRANS 3 MAXTRANS 255 STORAGE(INITIAL
83886080 NEXT 41943040 MINEXTENTS 1 MAXEXTENTS 1017 PCTINCREASE 0
FREELISTS 4 FREELIST GROUPS 1 BUFFER_POOL RECYCLE) TABLESPACE
"TSS_FACT" ;
```

## Netezza

```
CREATE TABLE MRDWDDM.RDWF_DDM_ROOMS_SOLD (
ID_PROPERTY numeric(5, 0) NOT NULL ,
ID_DATE_STAY integer NOT NULL ,
CD_ROOM_POOL CHAR(4) NOT NULL ,
CD_RATE_PGM CHAR(4) NOT NULL ,
CD_RATE_TYPE CHAR(1) NOT NULL ,
CD_MARKET_SEGMENT CHAR(2) NOT NULL ,
ID_CONFO_NUM_ORIG integer NOT NULL ,
ID_CONFO_NUM_CUR integer NOT NULL ,
ID_DATE_CREATE integer NOT NULL ,
ID_DATE_ARRIVAL integer NOT NULL ,
ID_DATE_DEPART integer NOT NULL ,
QY_ROOMS integer NOT NULL ,
CU_REV_PROJ_NET_LOCAL numeric(21, 3) NOT NULL ,
CU_REV_PROJ_NET_USD numeric(21, 3) NOT NULL ,
QY_DAYS_STAY_CUR smallint NOT NULL ,
CD_BOOK_SOURCE CHAR(1) NOT NULL)
distribute on random;
```

No indexes

No Physical Tuning/Admin

Stripe data randomly, or by Columns





# IBM Netezza Models



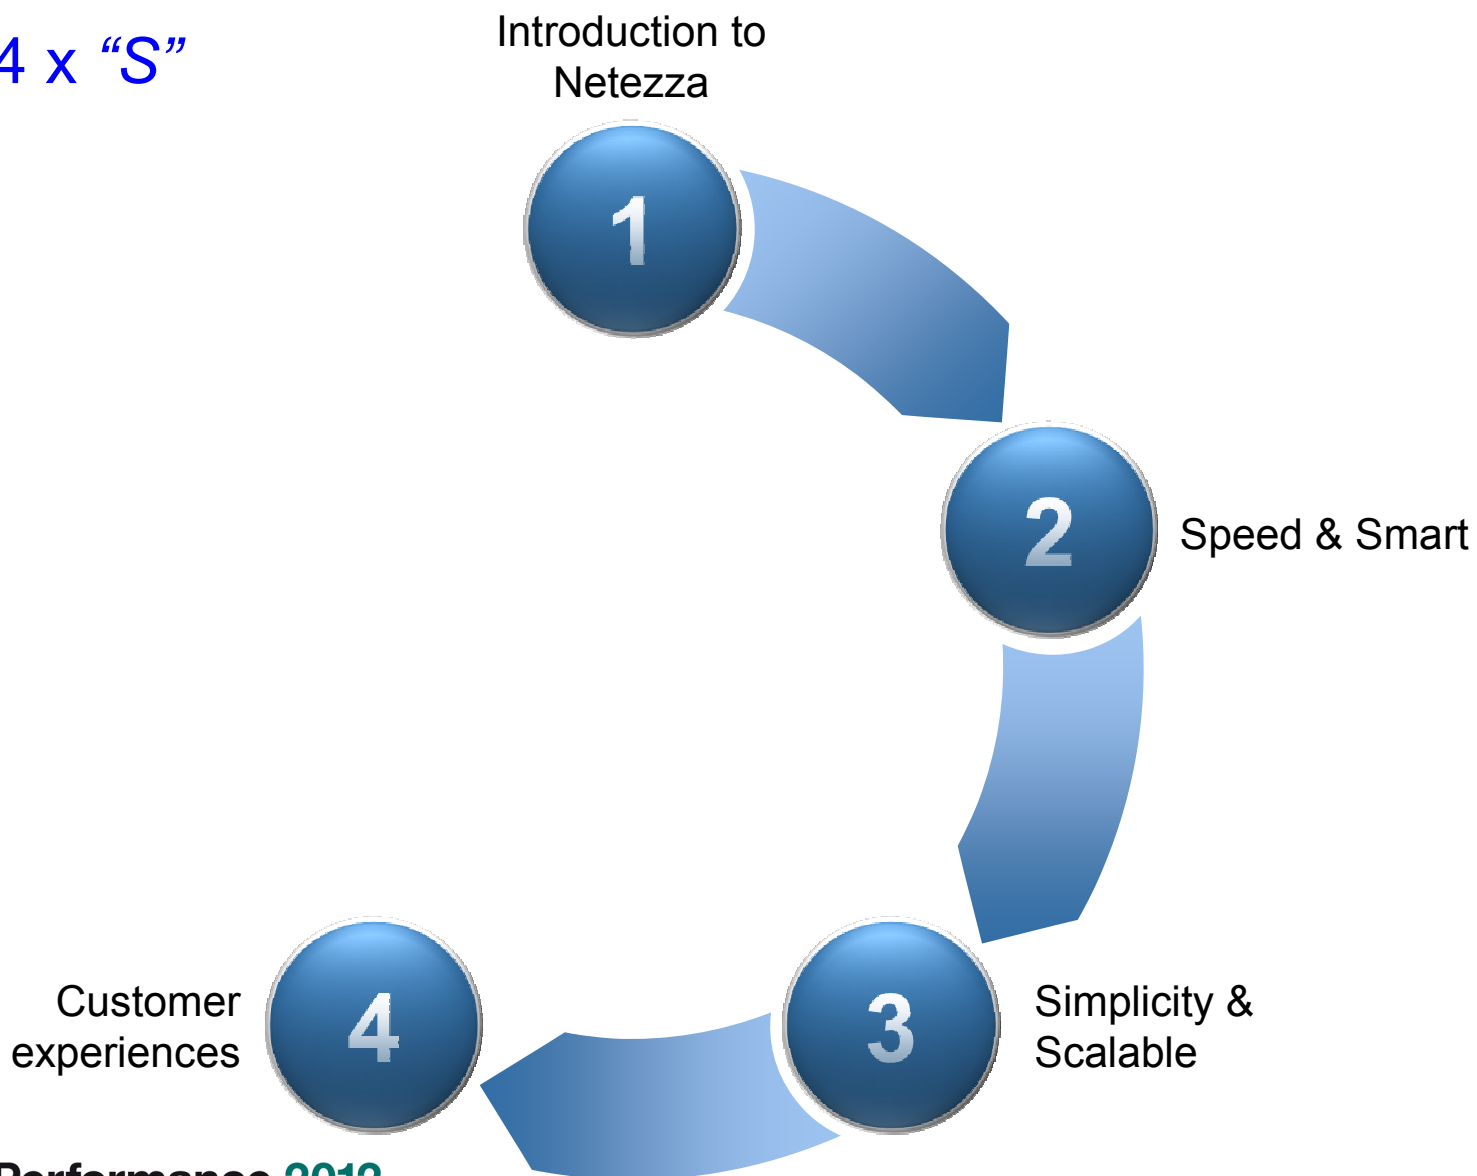
|                    | Quarter Rack | Half Rack | Full Rack | Two Racks | ...More Racks | ... 10 Clustered Racks |
|--------------------|--------------|-----------|-----------|-----------|---------------|------------------------|
| Snippet Processors | 24           | 48        | 96        | 192       | 384           | 960                    |
| Capacity (TB)      | 8            | 16        | 32        | 64        | 128           | 320                    |
| Compression (TB)   | 32           | 64        | 128       | 256       | 512           | 1280                   |

Capacity = User Data space  
 Compression = Effective User Data Space





# Netezza 4 x "S"



# Over 600 live IBM Netezza Clients

## 90% acts as a reference



### Digital Media



### Financial Services



### Government



### Health & Life Sciences



### Retail / Consumer Products



### Telecom



### Other





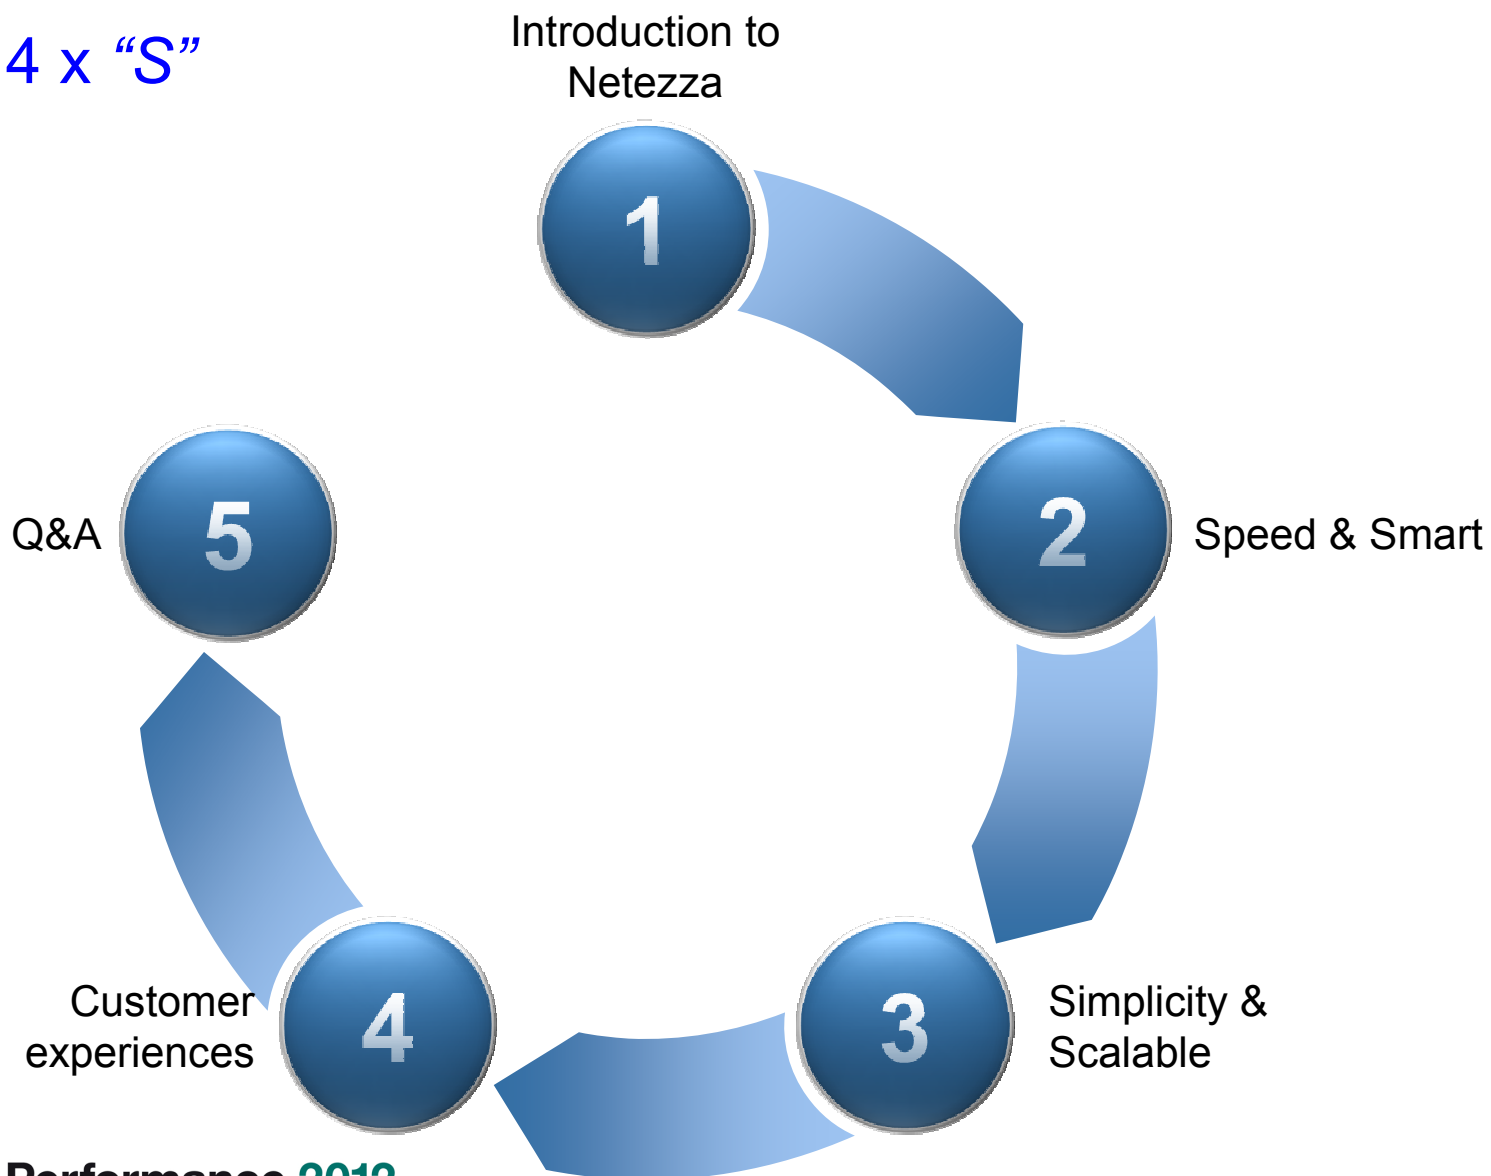
# IBM Netezza – Some Local Clients in BeNeLux



**IBM Performance 2012**  
Smarter Analytics. Smarter Outcomes.



# Netezza 4 x "S"



**Get Proof. Not Promises.**



Take the  
**Netezza**  
**TestDrive**

- ✓ Your Site
- ✓ Your Data
- ✓ Our Technology

A Proof of Concept in your own data center.

Netezza's data warehouse appliances are purpose-built to make advanced analytics simpler, faster and more accessible.

We'll bring it in.  
You bring it on.





# Netezza 4 x "S"

