

IBM Smarter Analytics 2012

Manage Information,
Gather Insights, **Outperform.**



Netezza

The True Data Warehouse and Analytics Appliance



Topics

What is Netezza?

Netezza Market Disruption

Transforming Decision Making with Embedded Analytics

IBM's Strategy: Smart Warehouse Consolidation

Summary

Netezza



- 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 and tuning
- **Scalability:** Peta-scale user data capacity
- **Smart:** High-performance advanced analytics

Appliance Simplicity

- Dedicated device
- Optimized for purpose
- Complete solution
- Fast installation
- Simple operation
- Standard interfaces
- Low cost



Appliance Family



IBM Netezza 100

Development &
Test

1 TB to 10 TB

IBM Netezza 1000

High performance
data warehouse &
analytics

1 TB to 1.5 PB

IBM Netezza High Capacity

Analytics on Big Data
Queryable Archiving

500 TB to 10+ PB

Sample Users

Digital Media



Financial Services



Government



Health & Life Sciences



Retail / Consumer Products



Telecom



Other



Implementation: Netezza Simplicity

Oracle DDL to Netezza migrated <1 day / 95%+ of Oracle DDL redundant

Oracle: 1821 lines of SQL DDL


```
CREATE TABLE CMVFACT.FCT_MKV_CUST_BAL
(
  PER_CCYYMM      NUMBER(6) NOT NULL,
  SK_CIF          NUMBER(15) NOT NULL,
  SK_PR_L5        NUMBER(8) NOT NULL,
  BAL_CD          NUMBER(5) NOT NULL,
  BUS_SEG_L3_ID   NUMBER(4) NOT NULL,
  SEG2_CD         NUMBER(2) NOT NULL,
  ISO_CRNCY_CD    CHAR(3) NOT NULL,
  CO_CST_CTR_ID   CHAR(10) NOT NULL,
  CAL_MON_DAY_NR  NUMBER(2) NOT NULL,
  CAL_MON_BUS_DAY_NR NUMBER(2) NOT NULL,
  BAL_AM          NUMBER(15,2) NOT NULL,
  BAL_ORIG_CRNCY_AM NUMBER(18,2) NOT NULL,
  SRCE_SYS_CD     NUMBER(4)   NULL
)
TABLESPACE CMVFACT
NOLOGGING
PCTFREE 0
PCTUSED 90
INITRANS 1
MAXTRANS 255
STORAGE(BUFFER_POOL DEFAULT)
NOPARALLEL
NOCACHE
PARTITION BY RANGE(PER_CCYYMM)
(PARTITION P200205 VALUES LESS THAN (200206)
  STORAGE(FREELISTS 1
    FREELIST GROUPS 1),
PARTITION P200206 VALUES LESS THAN (200207)
  STORAGE(FREELISTS 1
    FREELIST GROUPS 1),
PARTITION P200207 VALUES LESS THAN (200208)
  STORAGE(FREELISTS 1
    FREELIST GROUPS 1),
PARTITION P200208 VALUES LESS THAN (200209)
  STORAGE(FREELISTS 1
    FREELIST GROUPS 1),
```

Netezza: 17 lines of SQL DDL

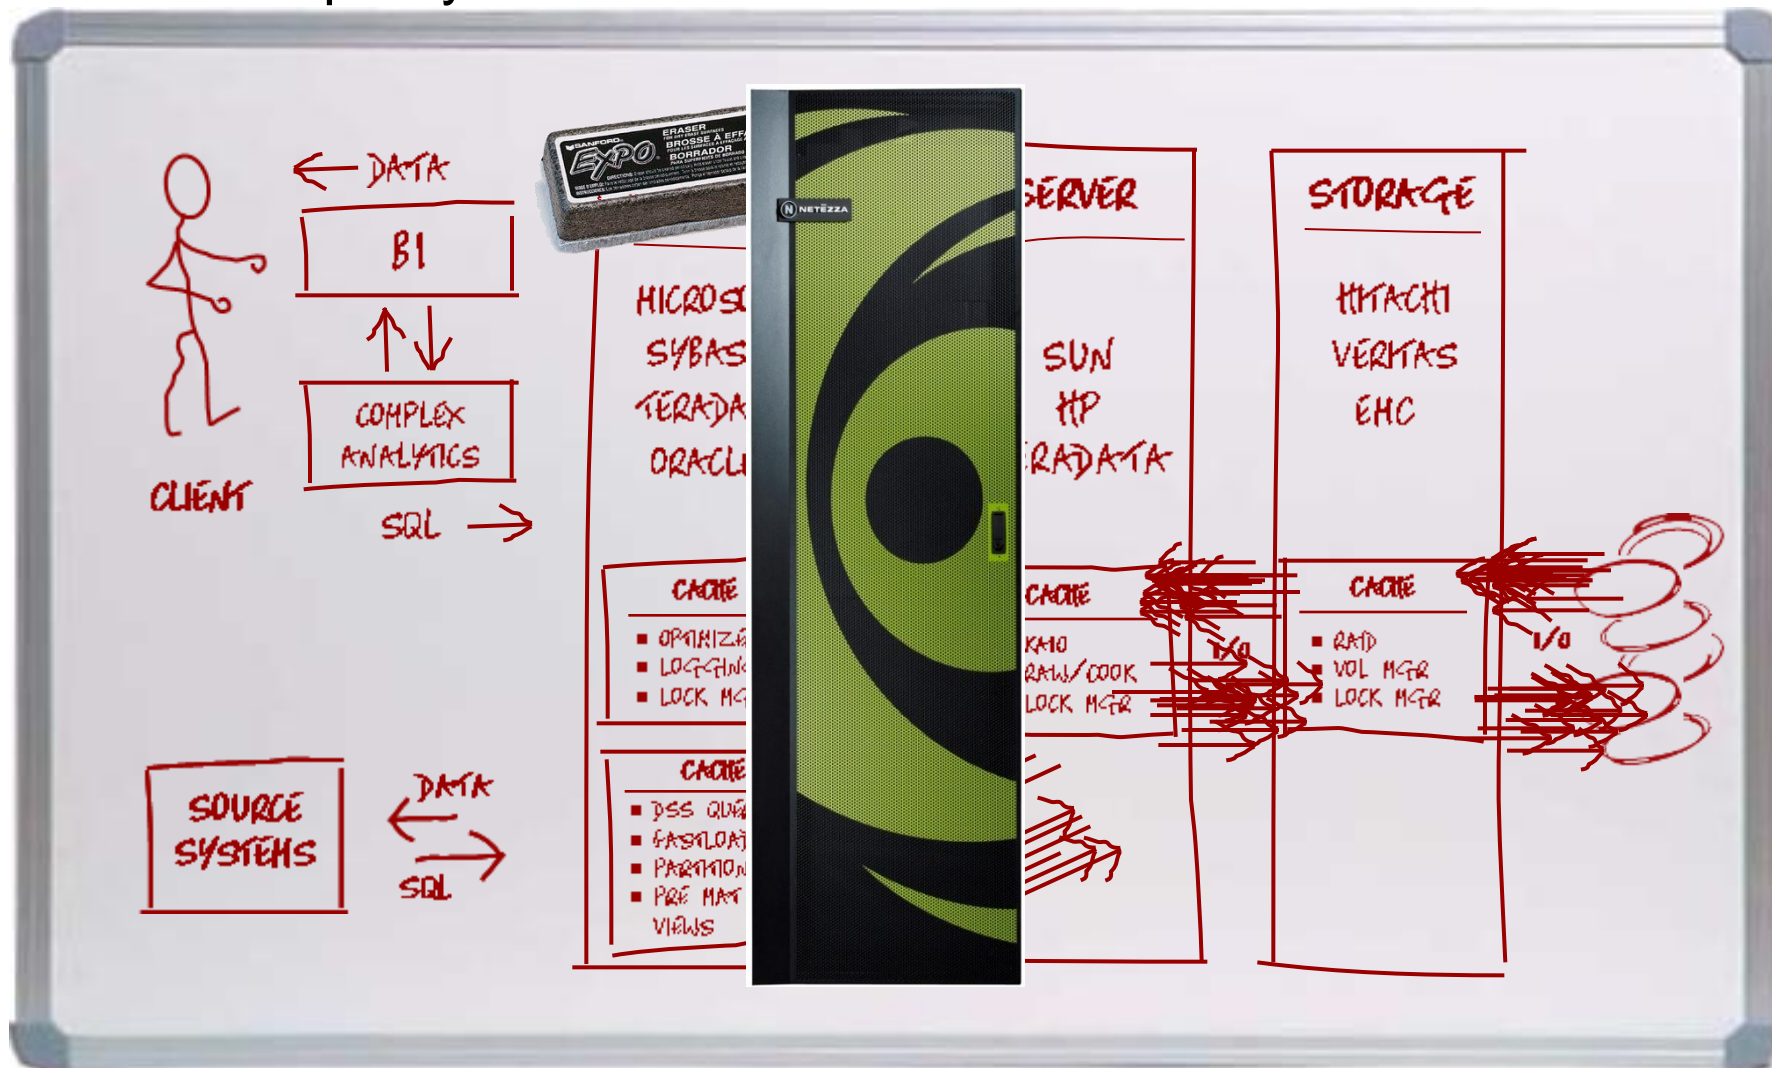
```
CREATE TABLE FCT_MKV_CUST_BAL
(
  PER_CCYYMM      integer NOT NULL,
  SK_CIF          bigint NOT NULL,
  SK_PR_L5        integer NOT NULL,
  BAL_CD          integer NOT NULL,
  BUS_SEG_L3_ID   integer NOT NULL,
  SEG2_CD         integer NOT NULL,
  ISO_CRNCY_CD    CHAR(3) NOT NULL,
  CO_CST_CTR_ID   CHAR(10) NOT NULL,
  CAL_MON_DAY_NR  integer NOT NULL,
  CAL_MON_BUS_DAY_NR integer NOT NULL,
  BAL_AM          numeric(15,2) NOT NULL,
  BAL_ORIG_CRNCY_AM numeric(18,2) NOT NULL,
  SRCE_SYS_CD     integer  NULL )
distribute on (sk_cif);
```


What makes Netezza so fast?

Traditional Data Warehouse Systems




Netezza Simplicity




Netezza Simplicity



The Netezza Appliance




The Netezza S-Blade



How does Netezza go so deep?

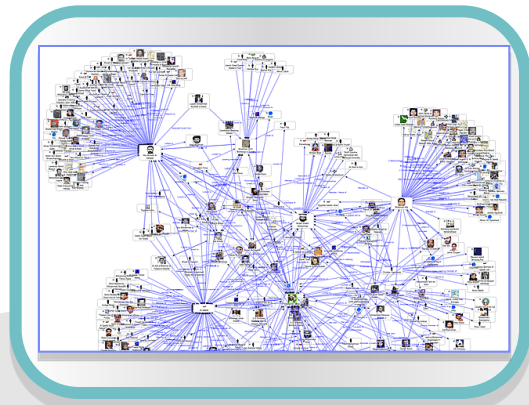
Becoming an Analytic Enterprise

Optimization



- What is the best choice?

Predictive Analytics



- What will happen?
- What will the impact be?

BI Reporting and Ad-Hoc Analysis



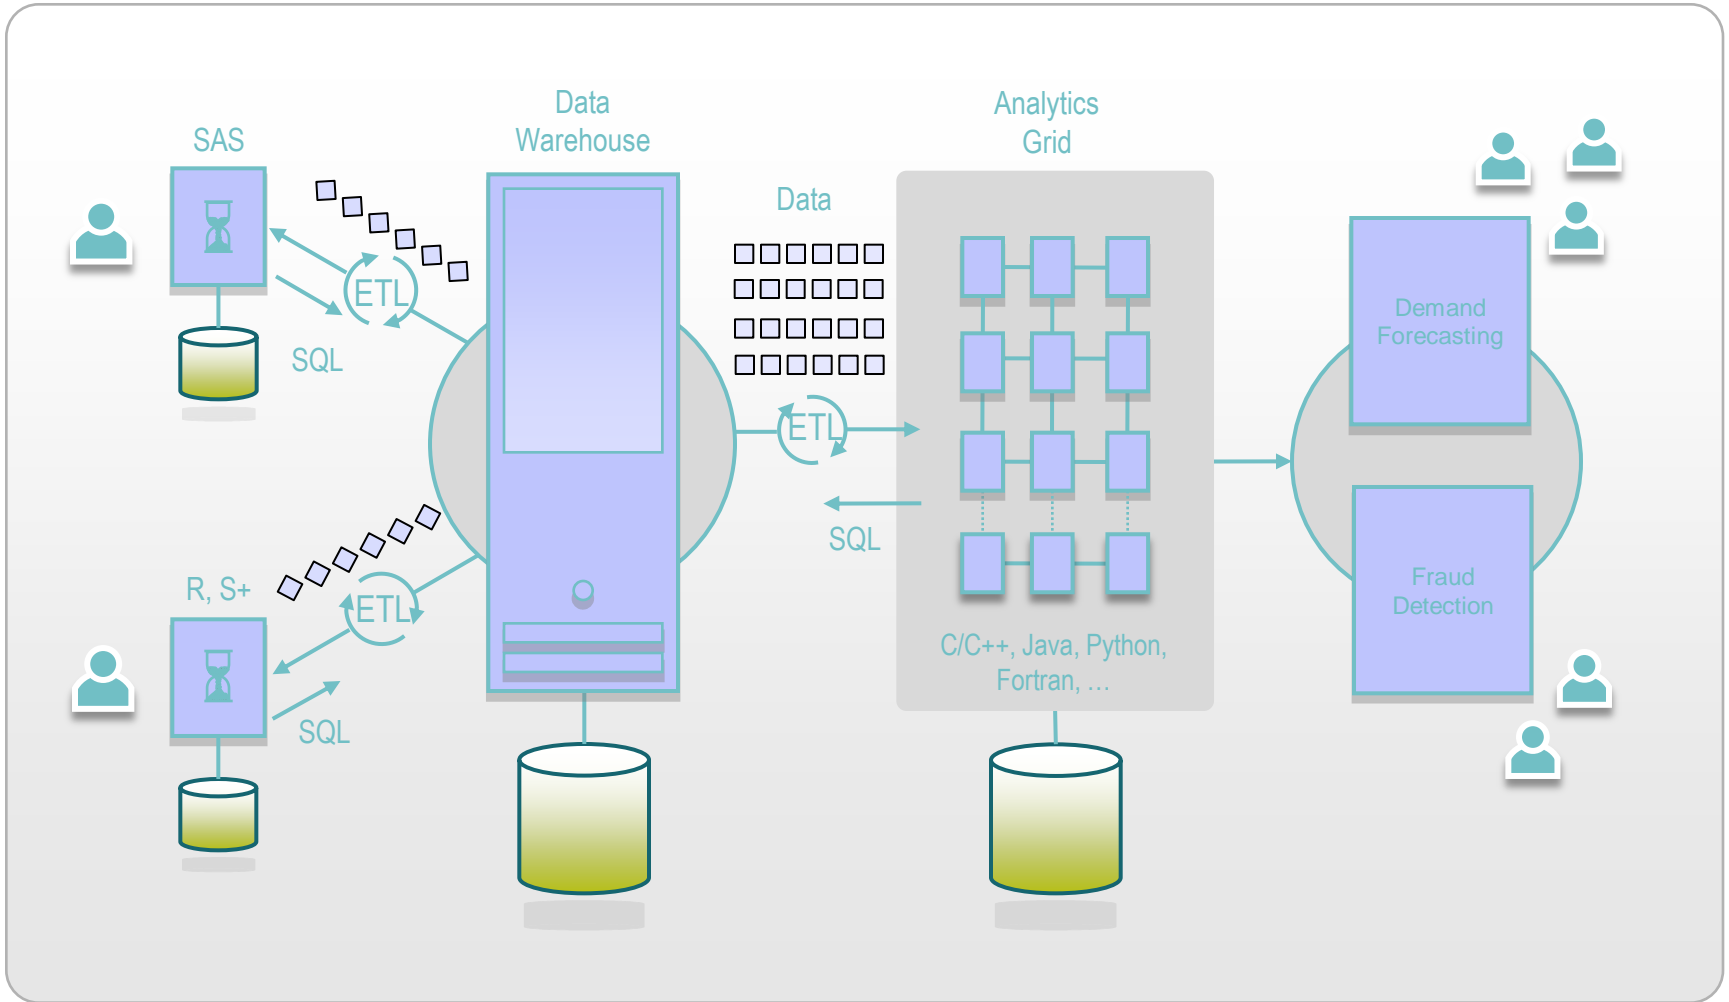
- What happened?
- When and where?
- How much?

Netezza Analytics: Big Data Meets Big Math

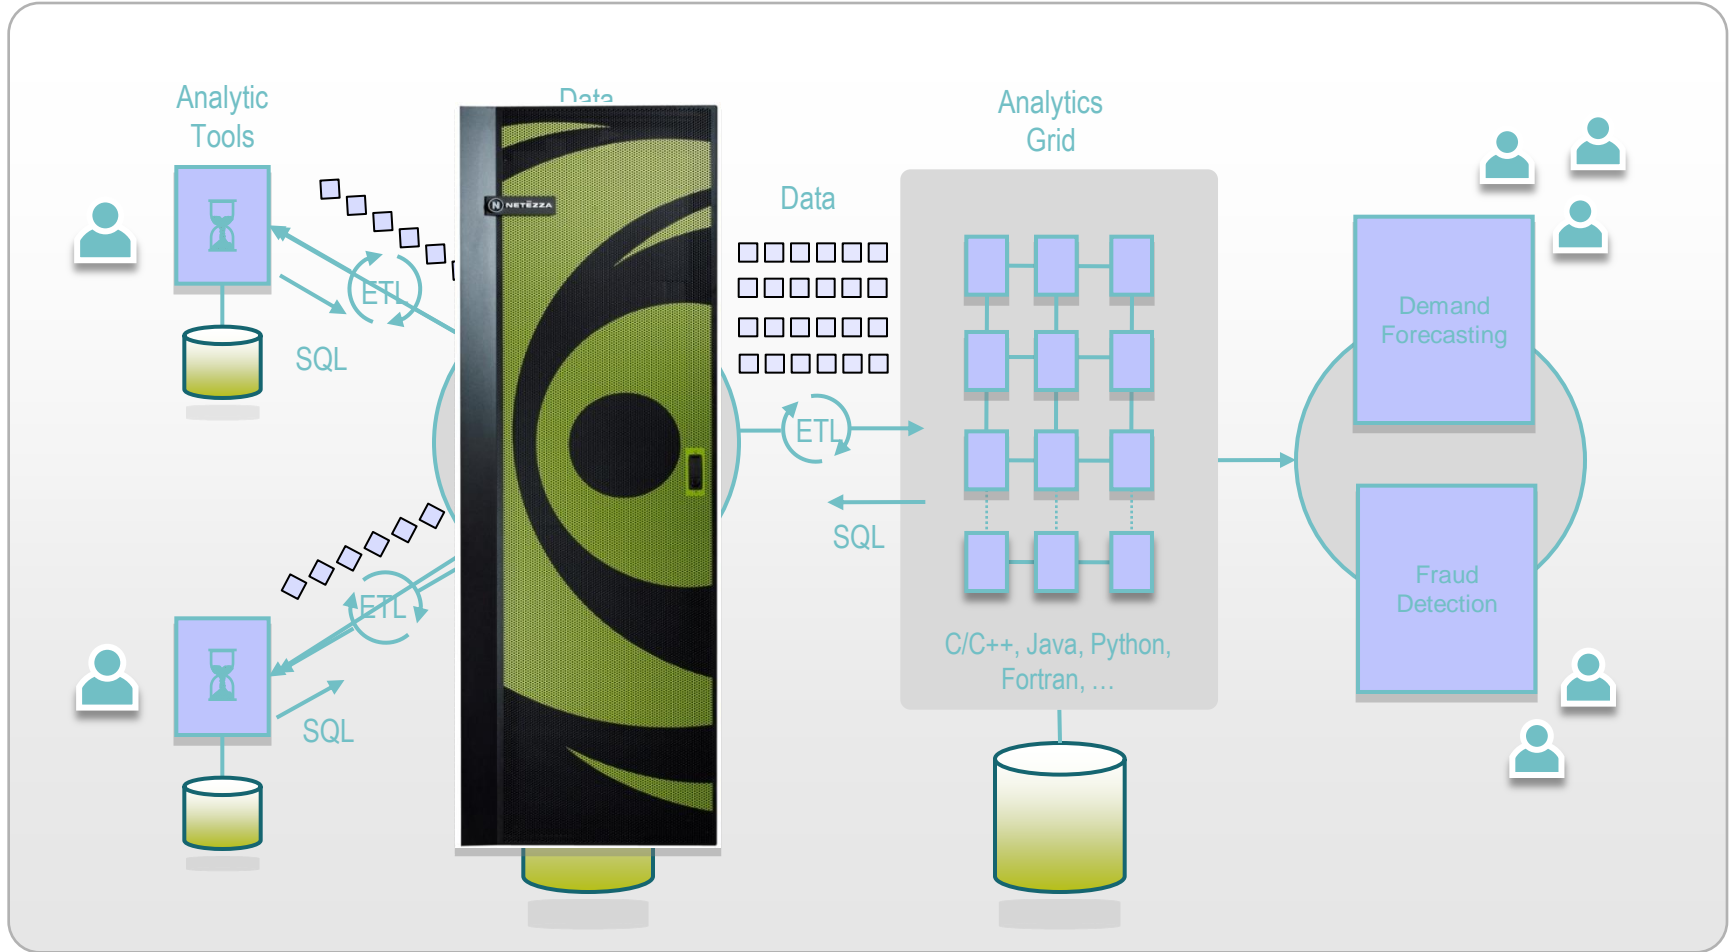


Analytics Without Constraints


Business Analytics Optimization (BAO) the Traditional Way



BAO with Netezza



BAO with Netezza



The Big Picture

IBM Netezza is Part of IBM Big Data Platform



InfoSphere BigInsights

Hadoop-based low latency analytics for variety and volume

Hadoop (NoSQL)



InfoSphere Information Server

High volume data integration and transformation

Information Integration



InfoSphere Streams

Low Latency Analytics for streaming data

Stream Computing



MPP Data Warehouse



IBM InfoSphere Warehouse

Large volume structured data analytics



IBM Netezza High Capacity Appliance

Queryable Archive Structured Data



IBM Netezza 1000

BI+Ad Hoc Analytics Structured Data

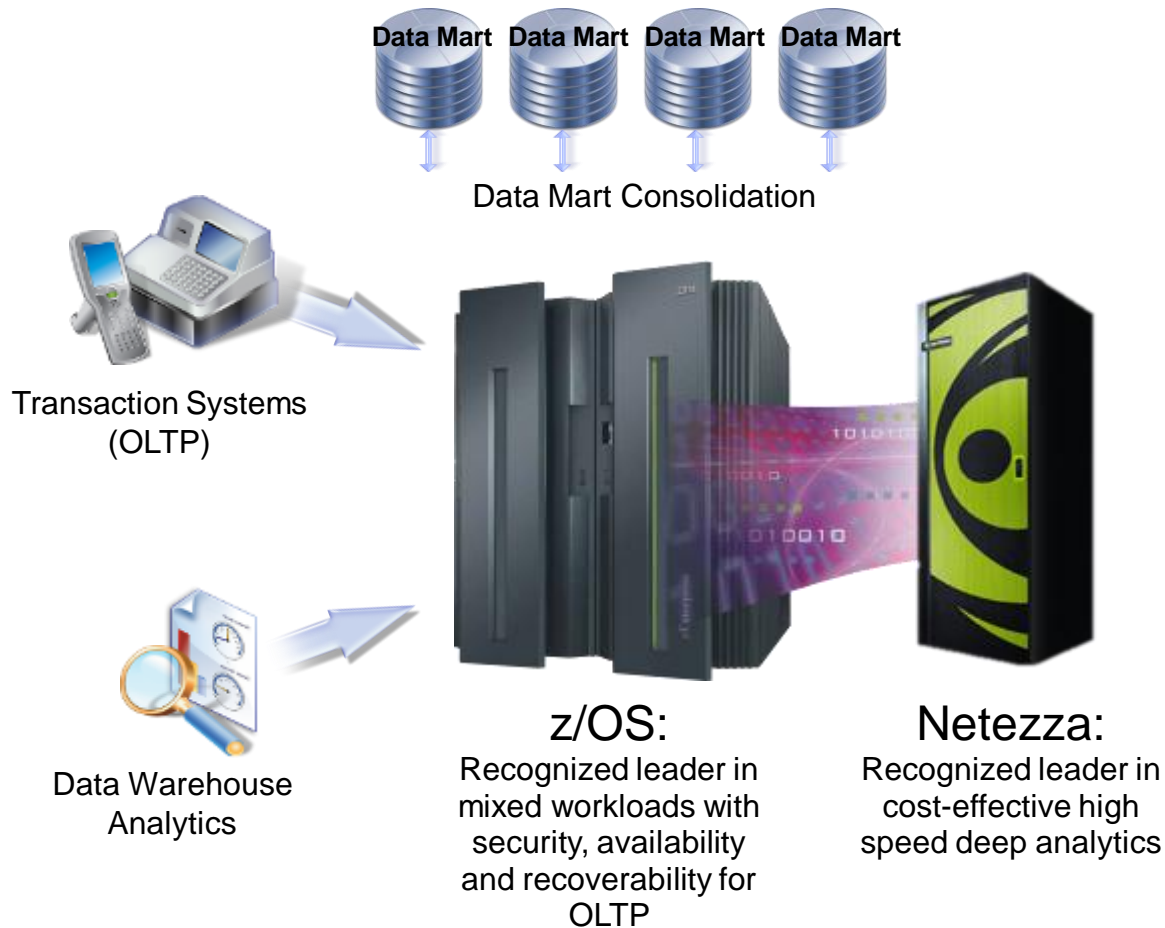


IBM Smart Analytics System

Operational Analytics on Structured Data

DB2 Analytics Accelerator for zOS

Combining the best transaction system with the best analytics system





- Manage more data
- Analyse more data, faster
- Go deeper with your analytics
- Reduce costs
- Supported by IBM
- Part of IBM Big Data vision

Pop Quiz

Netezza is a data warehouse appliance that combines which three technologies in an architecture optimized for high performance analytics?