




IBM INFORMATION INTEGRATION & GOVERNANCE SYMPOSIUM 2012

Delivering Trusted Information for Smarter Business Decisions

High Speed Analytics

Andrew Conacher – Netezza Specialist

May 2012





Do you want to work harder or smarter?

You are already smart. You are some of the way there.

Your environment needs to be smart too!



twitter: Follow @ANZ_IM or mention #IIGS

What others are doing - where *you* can be

500+ Billion market basket records

Leveraged the same BI applications

70 times more queries on 5 times the amount of data

100,000 queries a day on Netezza systems.

Reduced Storage by almost 100TB

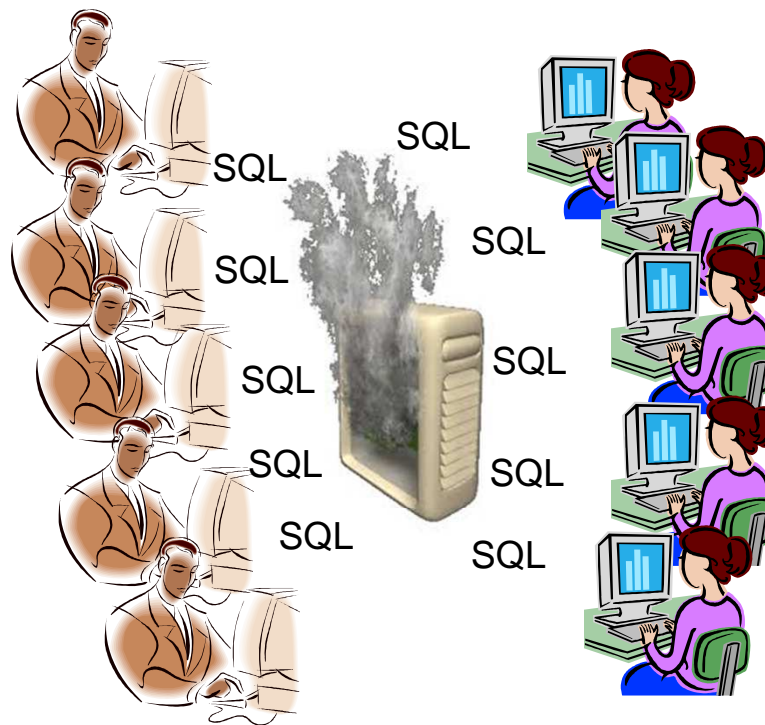
No longer needed aggregates with Netezza

'Monday mornings in hell' now a thing of the past



How we started out


Many users accessing a general purpose shared everything RDBMS server...



Crossing fingers it won't blow up. Get in early, stay back late to get your work done at quiet times when the server isn't as busy

Analytics at scale disconnected, not integrated

An effort at empowering users. Use Excel, Access and powerful tools. The DB server becomes a file Server.

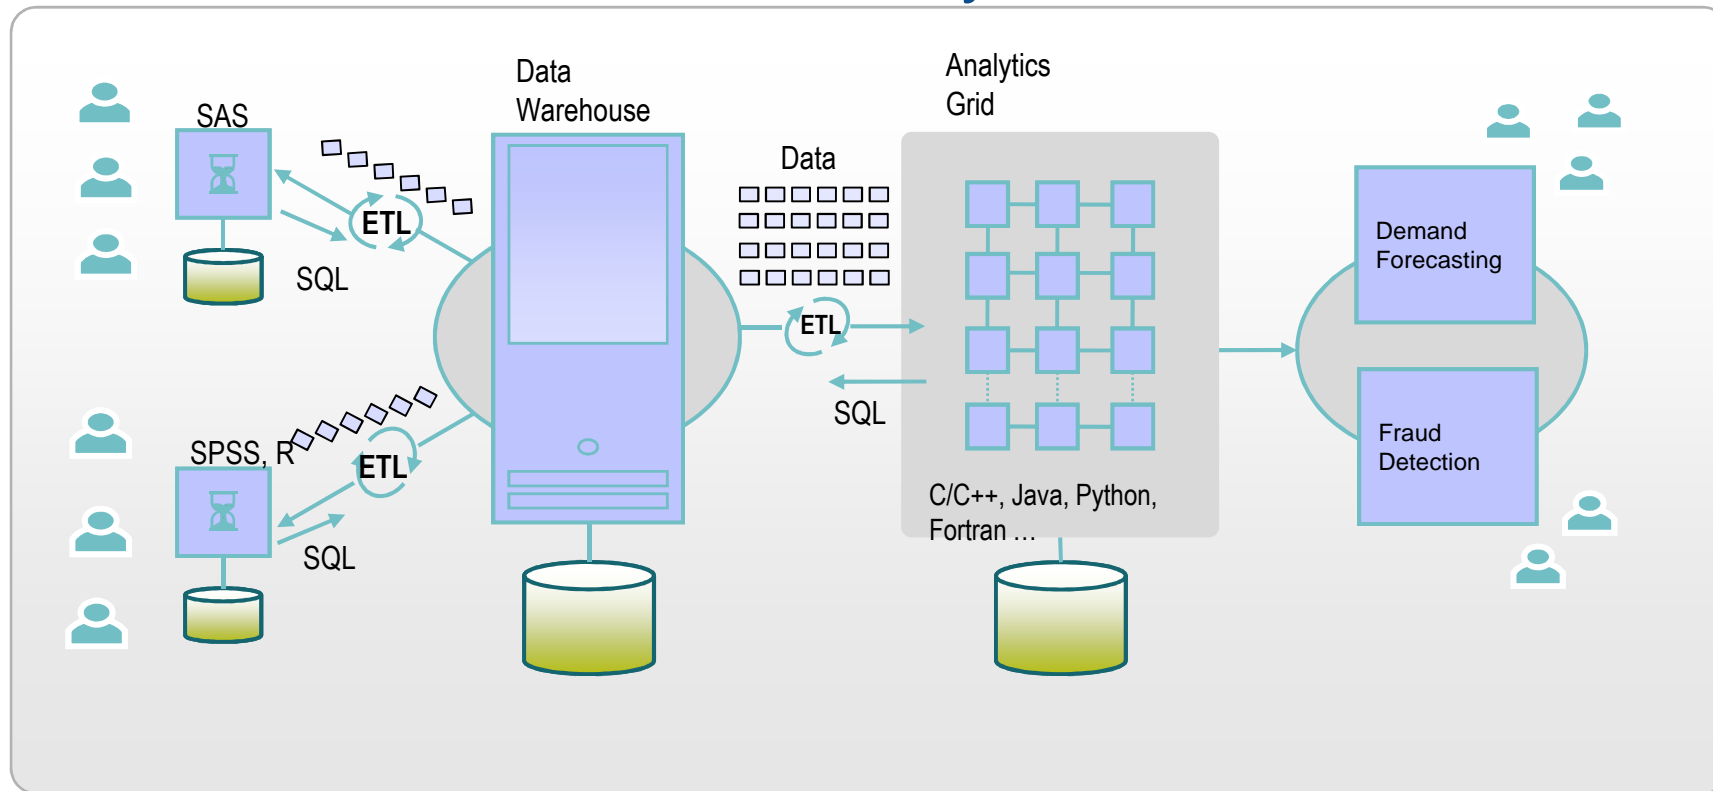


but there is gold in your code and formulas!

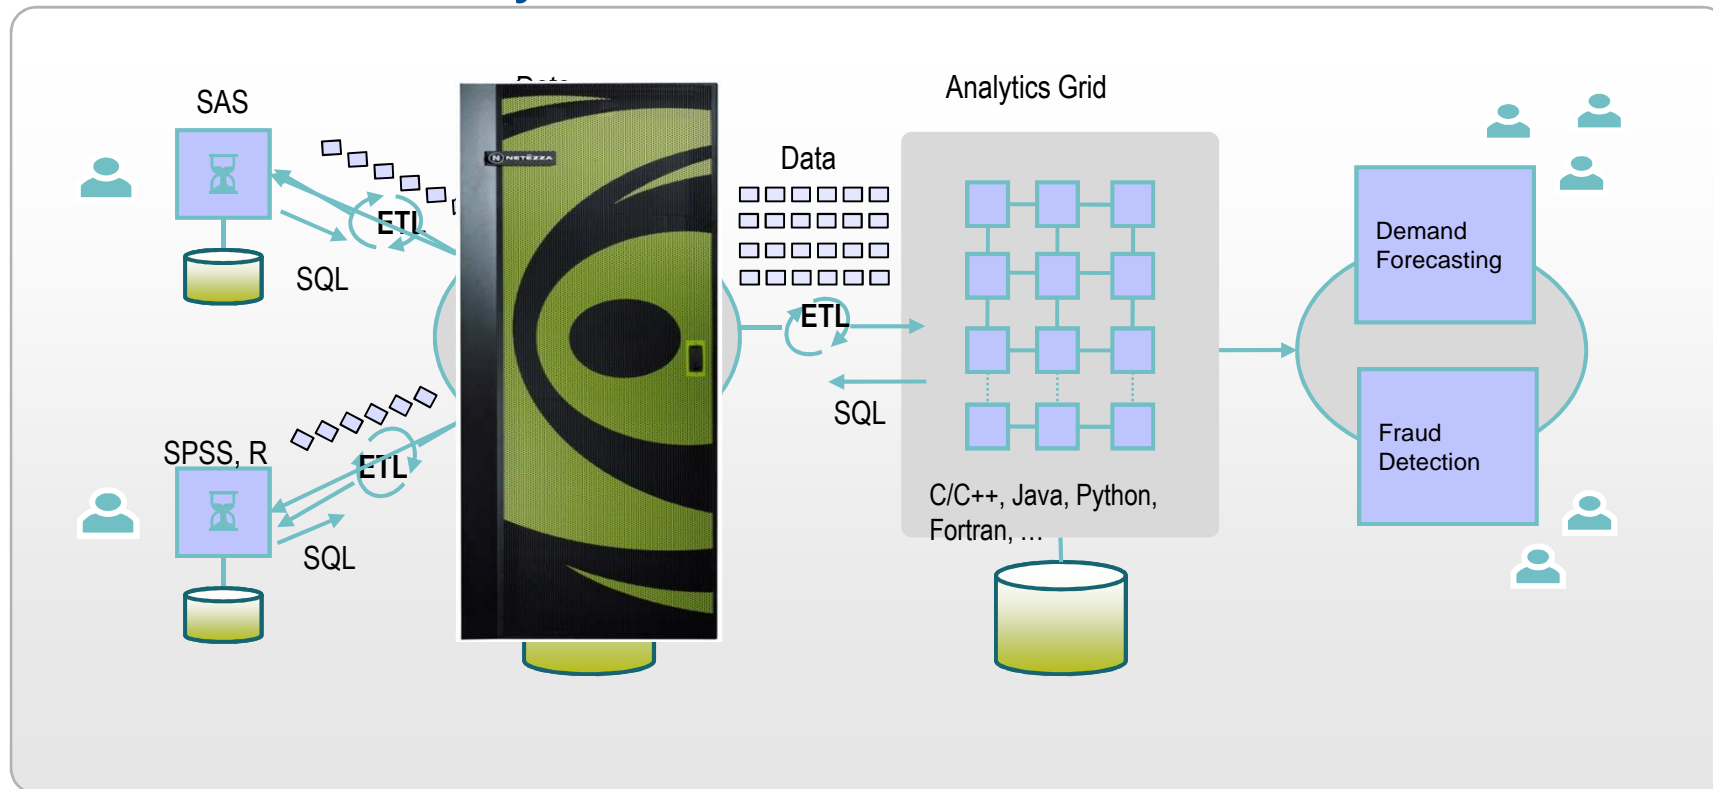


twitter: Follow @ANZ_IM or mention #IIGS

Where to next? Advanced Analytics – the traditional way

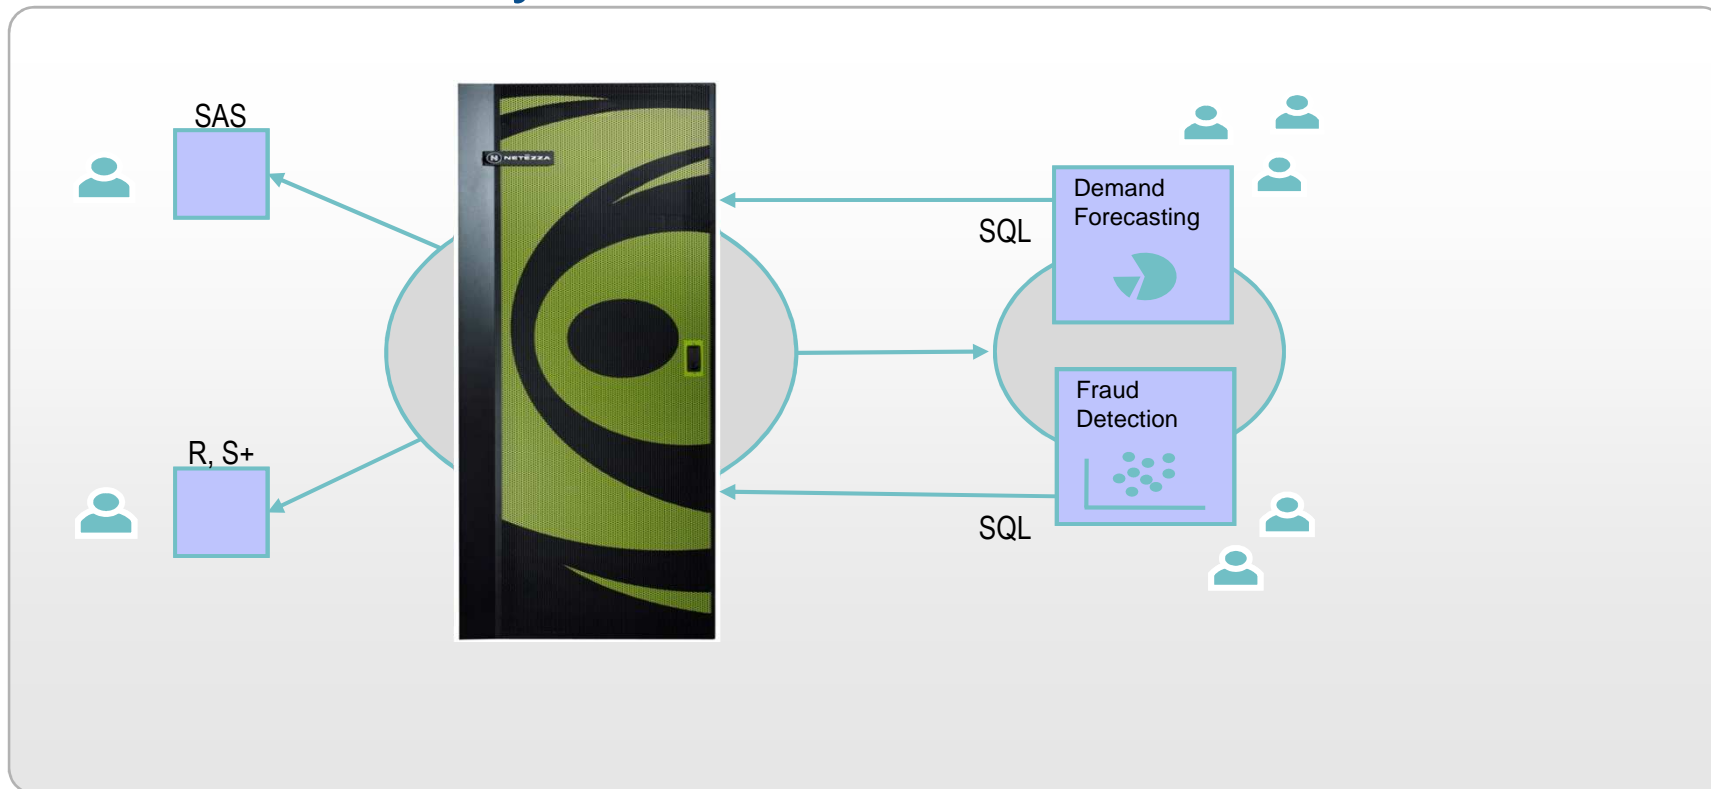


Advanced Analytics with Netezza





Advanced Analytics with Netezza



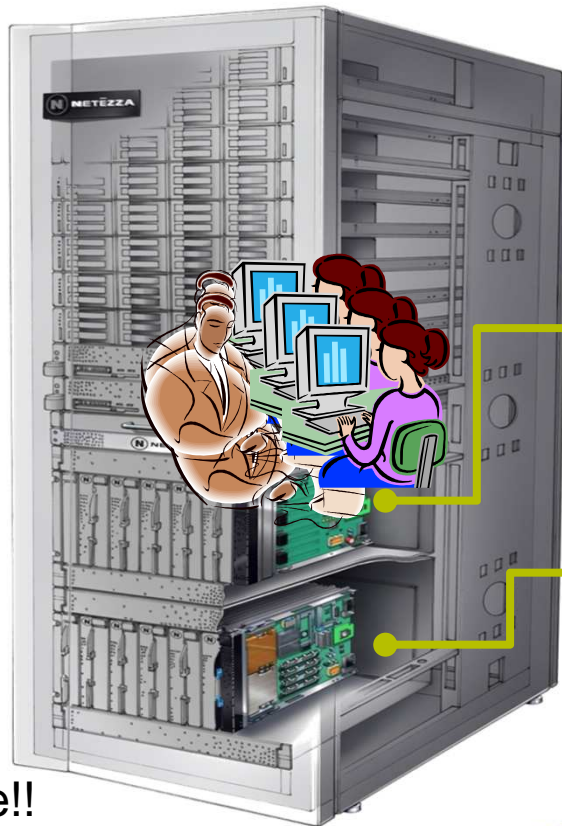
twitter: Follow @ANZ_IM or mention #IIGS



The IBM Netezza 1000 Architecture

Optimized Hardware+Software
Purpose-built for high performance analytics; requires no tuning

True MPP
All processors fully utilised for maximum speed and efficiency



Streaming Data
Hardware-based query acceleration for blistering fast results

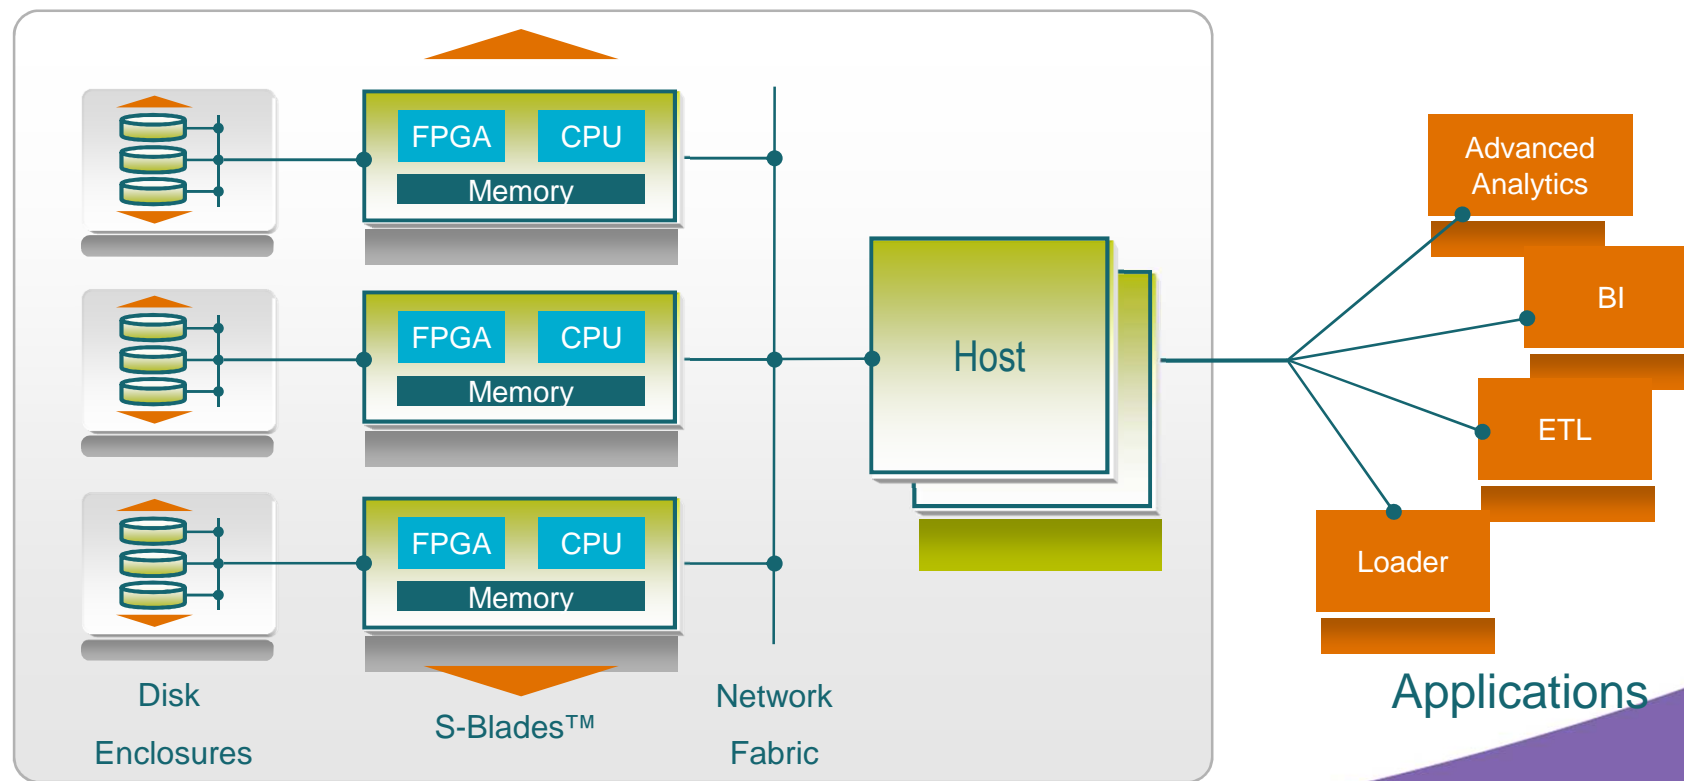
Deep Analytics
Complex analytics executed in-database for deeper insights

96 Cores working for you here!!





The Netezza AMPP™ Architecture




Netezza Appliance



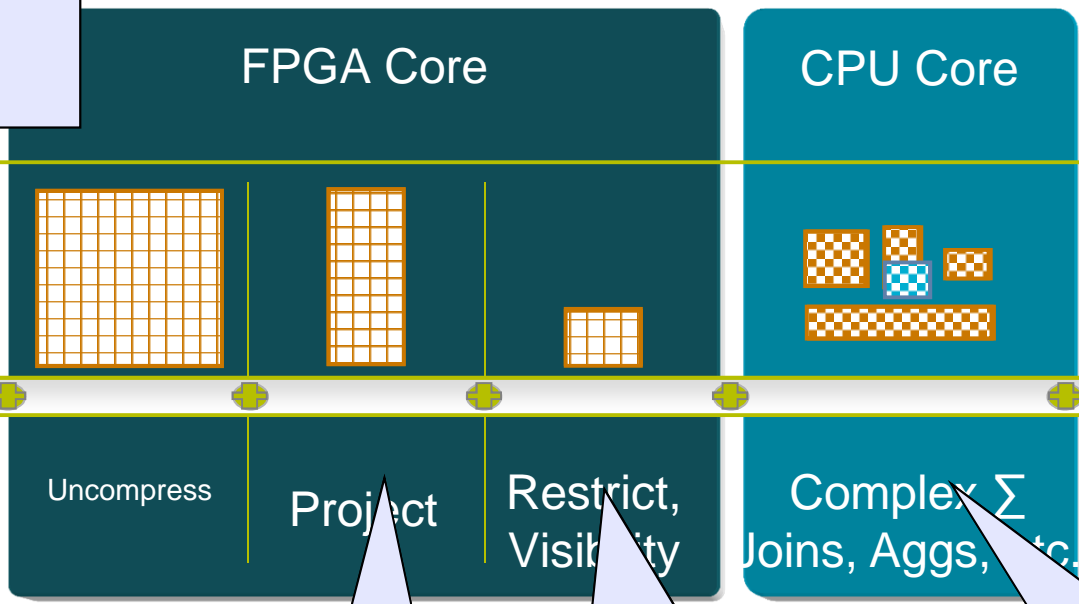
twitter: Follow @ANZ_IM or mention #IIGS

The Netezza Secret Sauce (some of it)

```
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,
       NRX
```

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

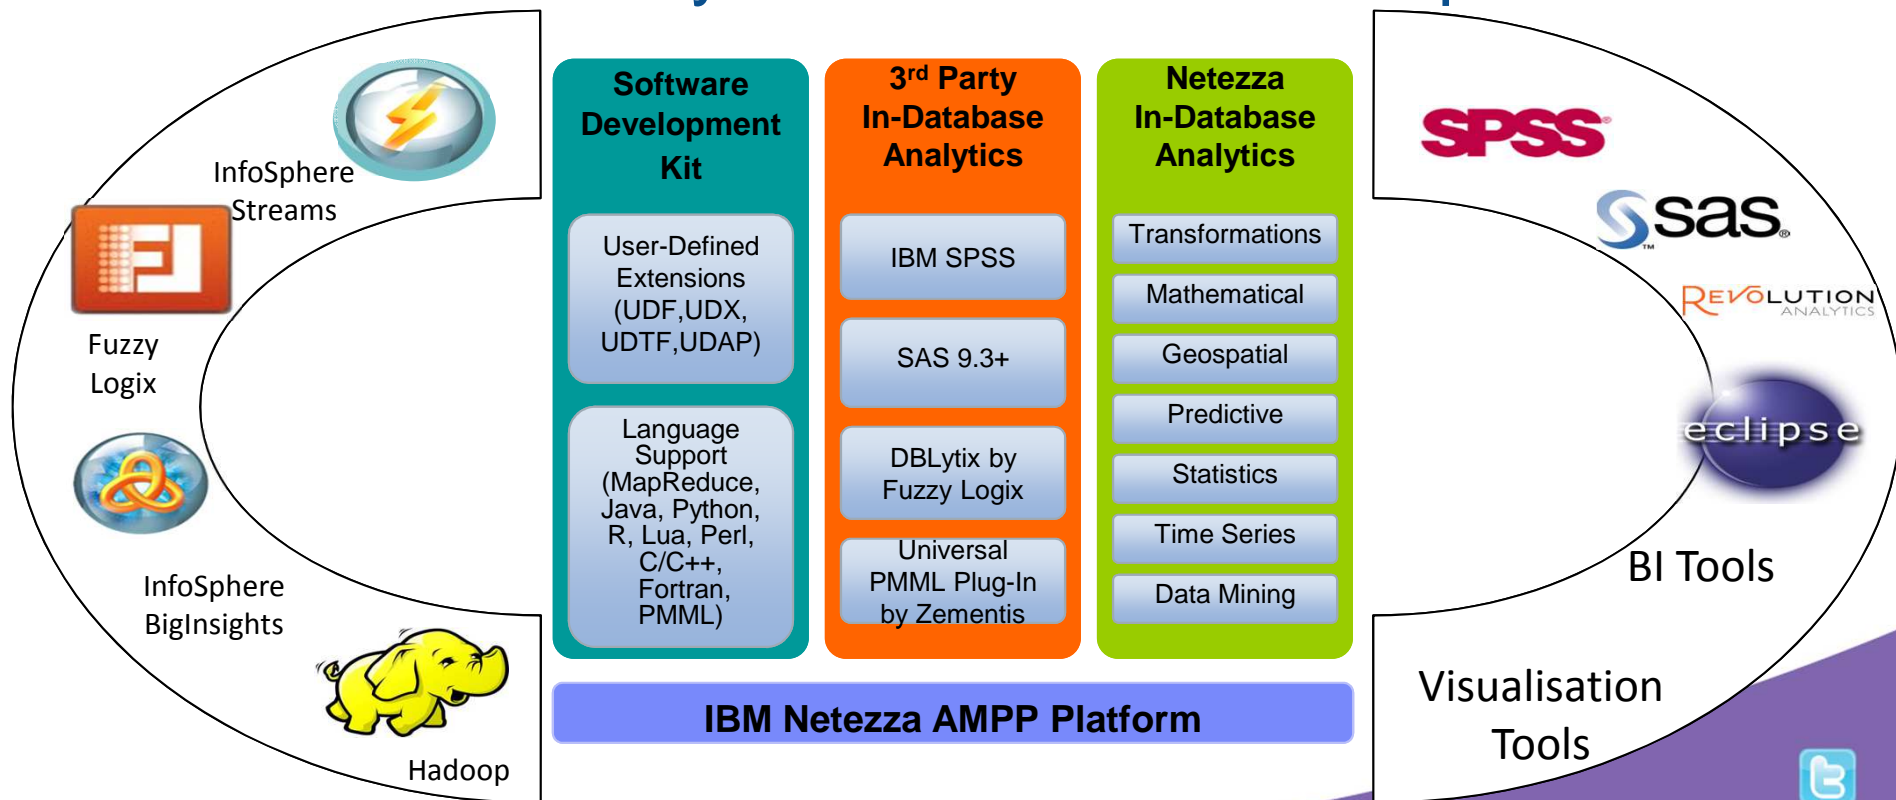
```
sum(NRX)
```



Follow @ANZ_IM or mention #IIGS

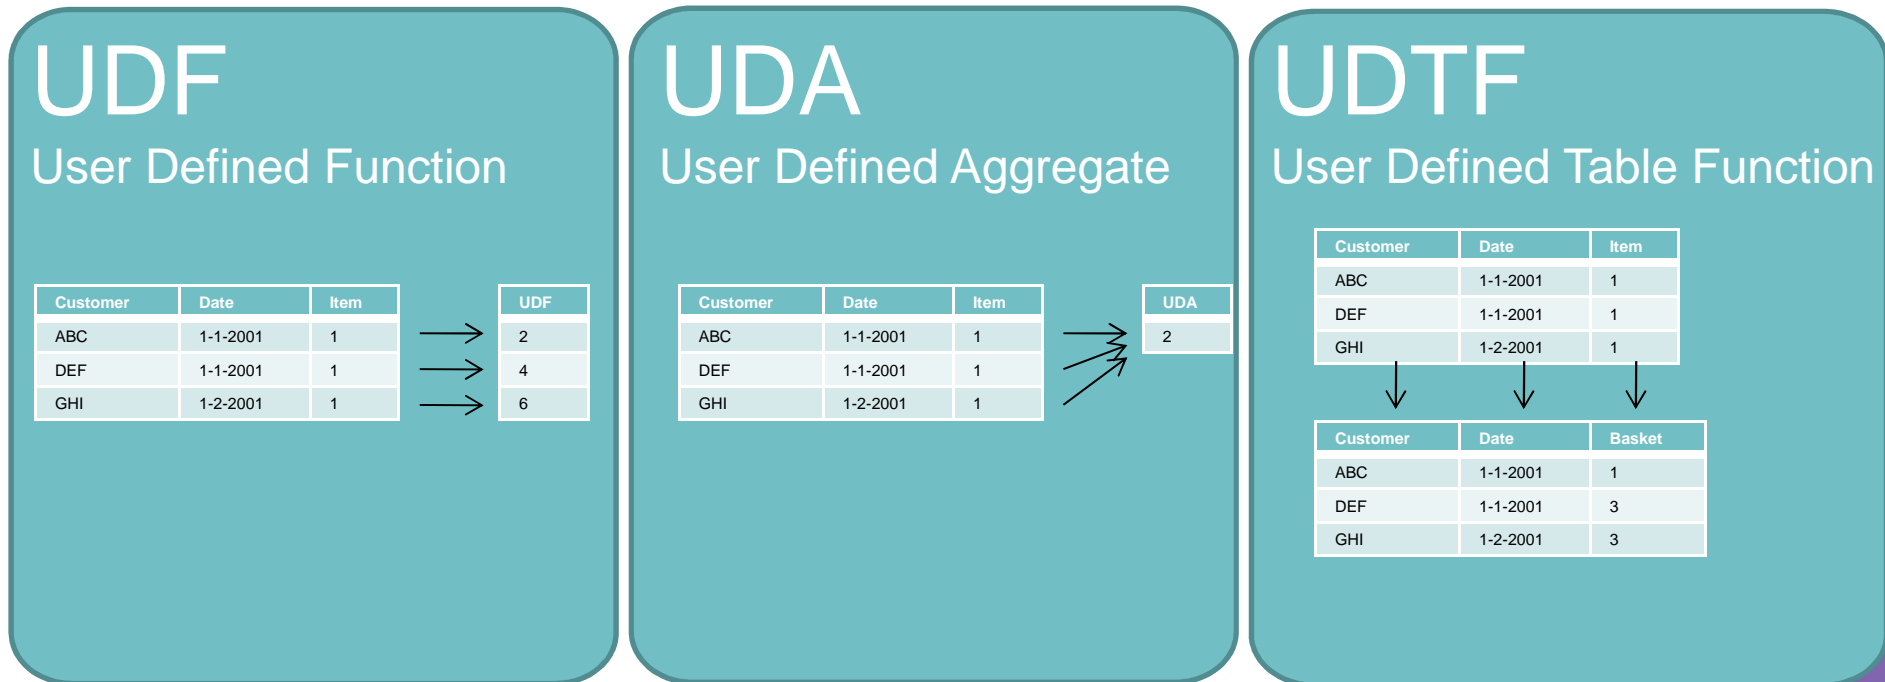


IBM Netezza Analytics: Features and Capabilities



twitter: Follow @ANZ_IM or mention #IIGS

Netezza Analytics implementations

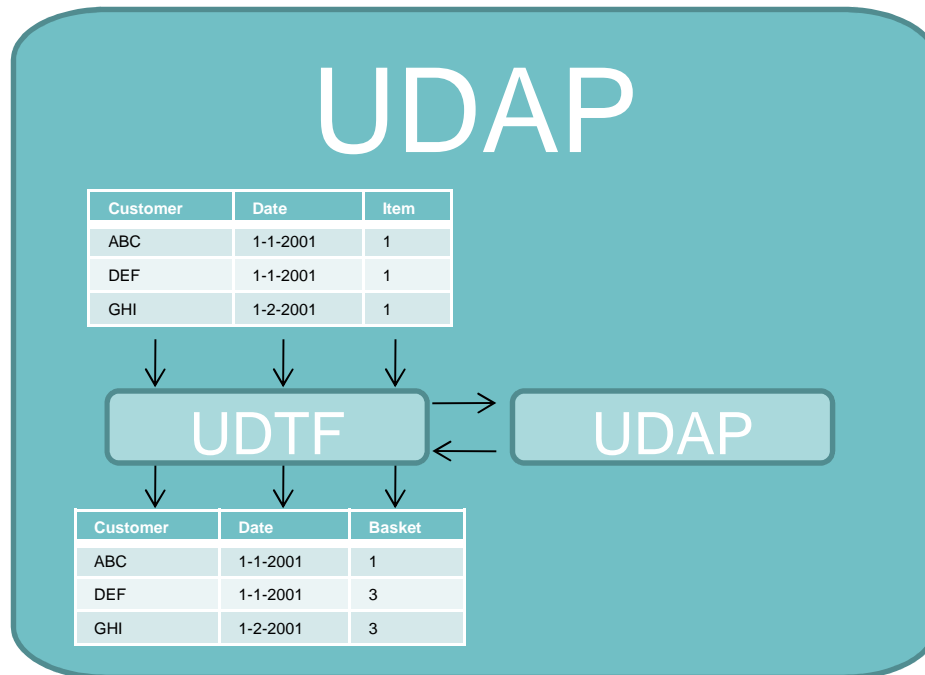


User Defined Extensions API consists of C/C++ classes Called via SQL



twitter: Follow @ANZ_IM or mention #IIGS

Newest User Defined Extension User-Defined Analytic Process



UDAP can be written
in C/C++, Java,
Python, Fortran, R,
LUA, Perl





Typical Netezza/Hadoop Use Cases



Unstructured data



Create context
(classification, text mining)



Analyse



Semi-structured data



Parse, aggregate



Analyse, report



Structured data



Analyse, report



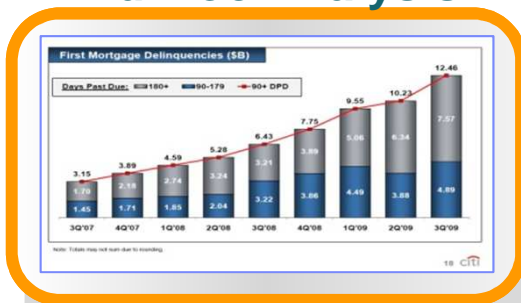
Active archival
Long running queries



twitter: Follow @ANZ_IM or mention #IIGS

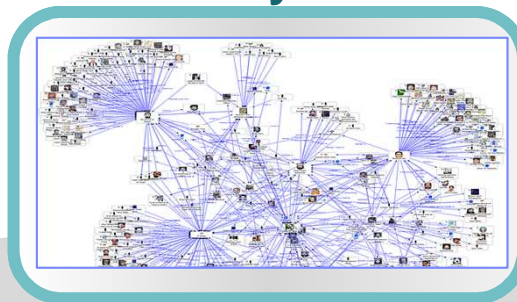
The Analytic Enterprise

BI Reporting and Ad-Hoc Analysis




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

Predictive Analytics



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

Optimisation



- What is the best choice?



twitter: Follow @ANZ_IM or mention #IIGS



Netezza at T-Mobile

- Originally a traditional RDBMS Data Warehouse – 40 Terabytes
- Managed by 3 FTEs
- Moved to Netezza
- Scaled from 40TB – 100TB – 500TB – 2 Petabytes today
- Same 3 FTEs
 - Principal Engineer
 - SQL Developer
 - Business conduit

[T-Mobile: IBM Netezza Client Success Video - YouTube](#)



www.youtube.com/watch?v=IBZYJpDT0vQ

14 Mar 2012 - 3 min - Uploaded by IBMCustomerReference
Network engineering success at scale with IBM **Netezza** Christine Twiford, T- Mobile's Manager of Network ...

- Used for Network Analytics
- 60 Applications, 1200 users, 400 data feeds
- Loads 17 billion records per day

<http://www.youtube.com/watch?v=IBZYJpDT0vQ>



twitter: Follow @ANZ_IM or mention #IIGS