
In-Memory-Technologie mit IoT Daten

sandor.szabo@de.ibm.com




-
- What are the industry trends
 - How IWA works
 - Demo IoT & IWA

Database Trends

NoSQL

SQL



- Document store
- Schema less
- Indexing Aggregation
- Sharding
- Replication
- Map/reduce
- Scalability

- Relational
- Fixed schema
- Joins, DML Optimizer
- Transactions(multi-stmt)
- Analytics
- Replication
- Triggers, Stored Procedures

The Hybrid Solution

- **Relational and non-relational data in one system**
- **Invisible Row and Columnar in memory storage**
- NoSQL/MongoDB Apps and traditional application
- Distributed Queries
- Multi-statement Transactions
- Enterprise Proven Reliability
- Enterprise Scalability
- Enterprise Level Availability



**Informix provides the capability to leverage
the abilities of both relational DBMS and document store
systems.**

What Makes BLU Acceleration Different?

Unmatched Innovations from IBM Research & Development Labs

Next Generation In-Memory

In-memory columnar processing with dynamic movement of data from storage




Analyze Compressed Data

Patented compression technique that preserves order so data can be used without decompressing




CPU Acceleration


Multi-core and SIMD parallelism (Single Instruction Multiple Data)



Data Skipping

Skips unnecessary processing of irrelevant data





IBM Informix Warehouse Accelerator (IWA)

TECHNOLOGY OVERVIEW

IWA Benefits



- **Extreme performance for Analytics: 100x+** faster response times for complex BI queries
- **Leverages** existing Informix database, builds on top, to provide **instant performance** boost
- Uses **low cost commodity HW**: Linux on Intel/AMD 64-bit
- Handles **Terabytes** of data in-memory, thanks to compressed storage and query technology
- Works “behind the scenes” in Informix, **transparent** to client applications
- Very **simple** and **flexible** installation, configuration and administration
- Informix + IWA is a **hybrid database** platform which provides the best technology and performance for both **OLTP and OLAP** workloads and support of **Big Data** solutions
- **No need to keep doing all this** in order to get high performance OLAP queries:
 - Indexes
 - Aggregates / summary tables
 - Materialized query tables/views
 - Cubes
 - Decide on best data partition strategies
 - Keep different database systems for each type of workload: OLTP vs OLAP
 - Migrate data to another OLAP database
 - Change your database schema
 - Change your analytic applications
 - Tune I/O, memory and CPU for OLAP
 - Update Statistics
 - Tune Queries with Optimizer Directives

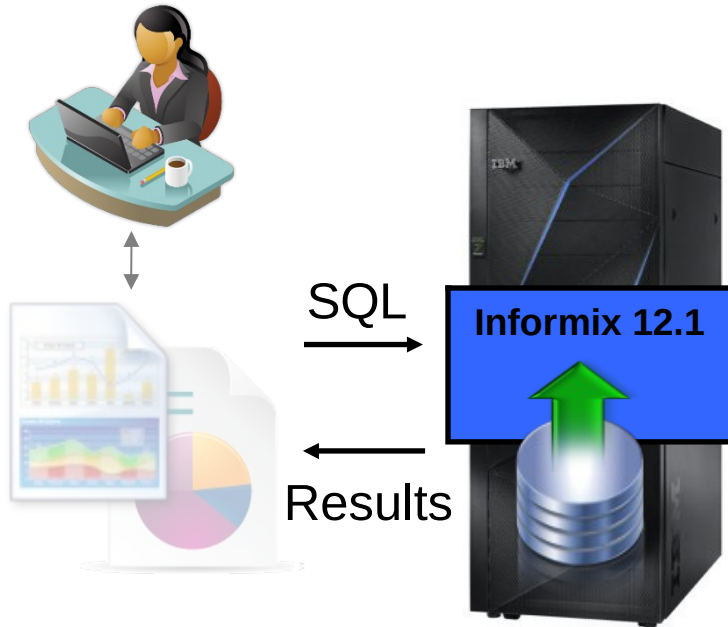
IBM Informix Warehouse Accelerator (IWA)

HOW IT WORKS



IWA Overview and Seamless Integration with Informix/IDS

- Before IWA...

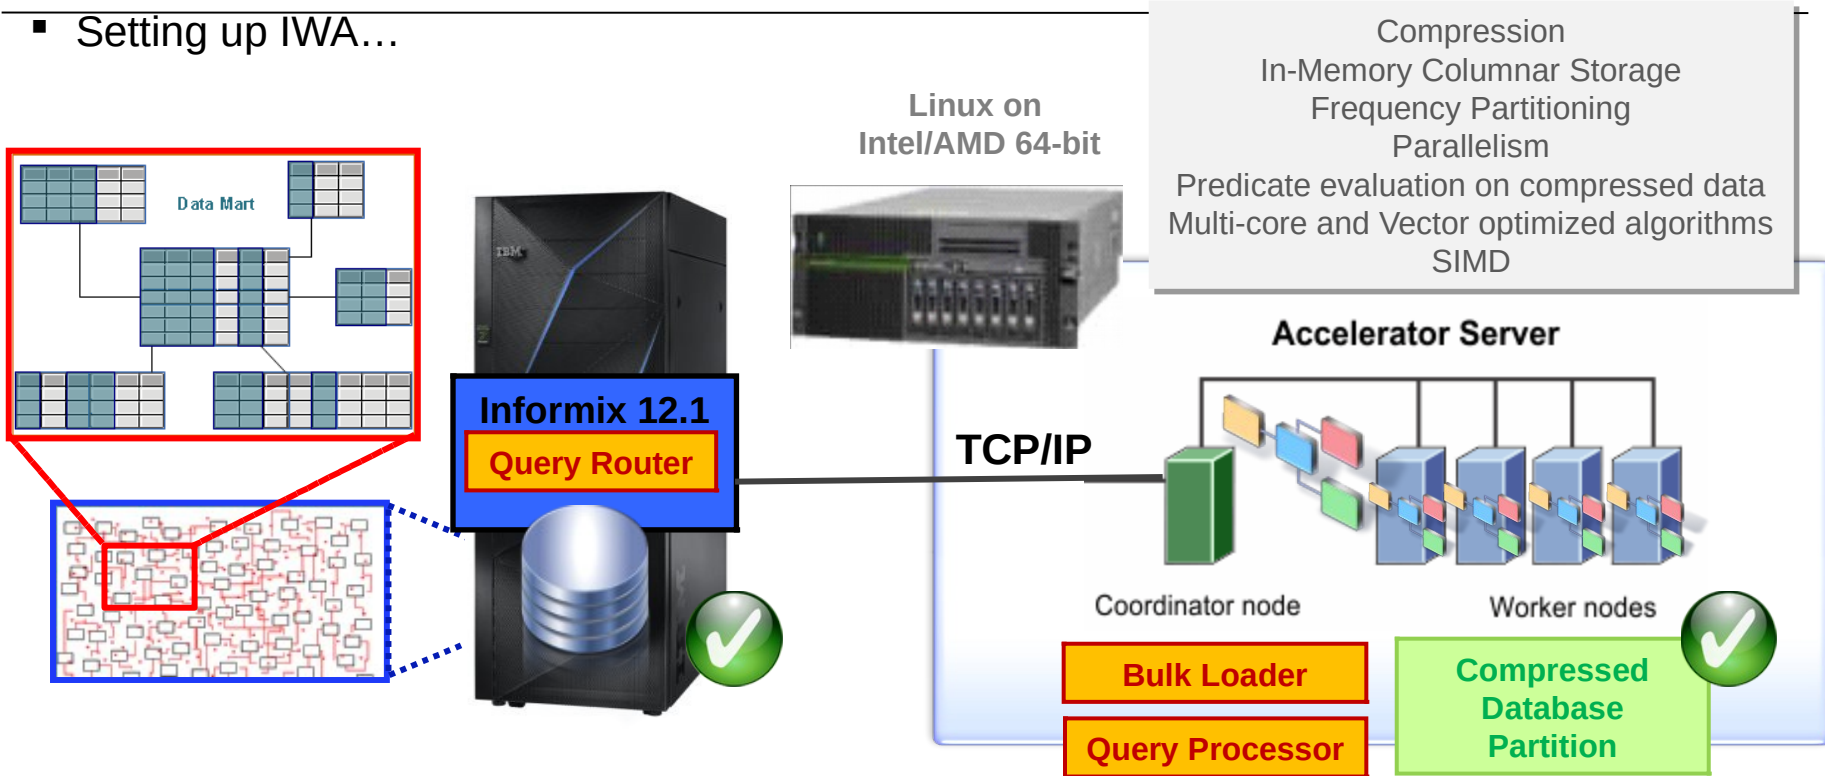


- Informix

- Receives analytic query from client
- Spends some time doing intensive I/O
- Returns results back to the client

IWA Overview and Seamless Integration with Informix/IDS

Setting up IWA...



Informix

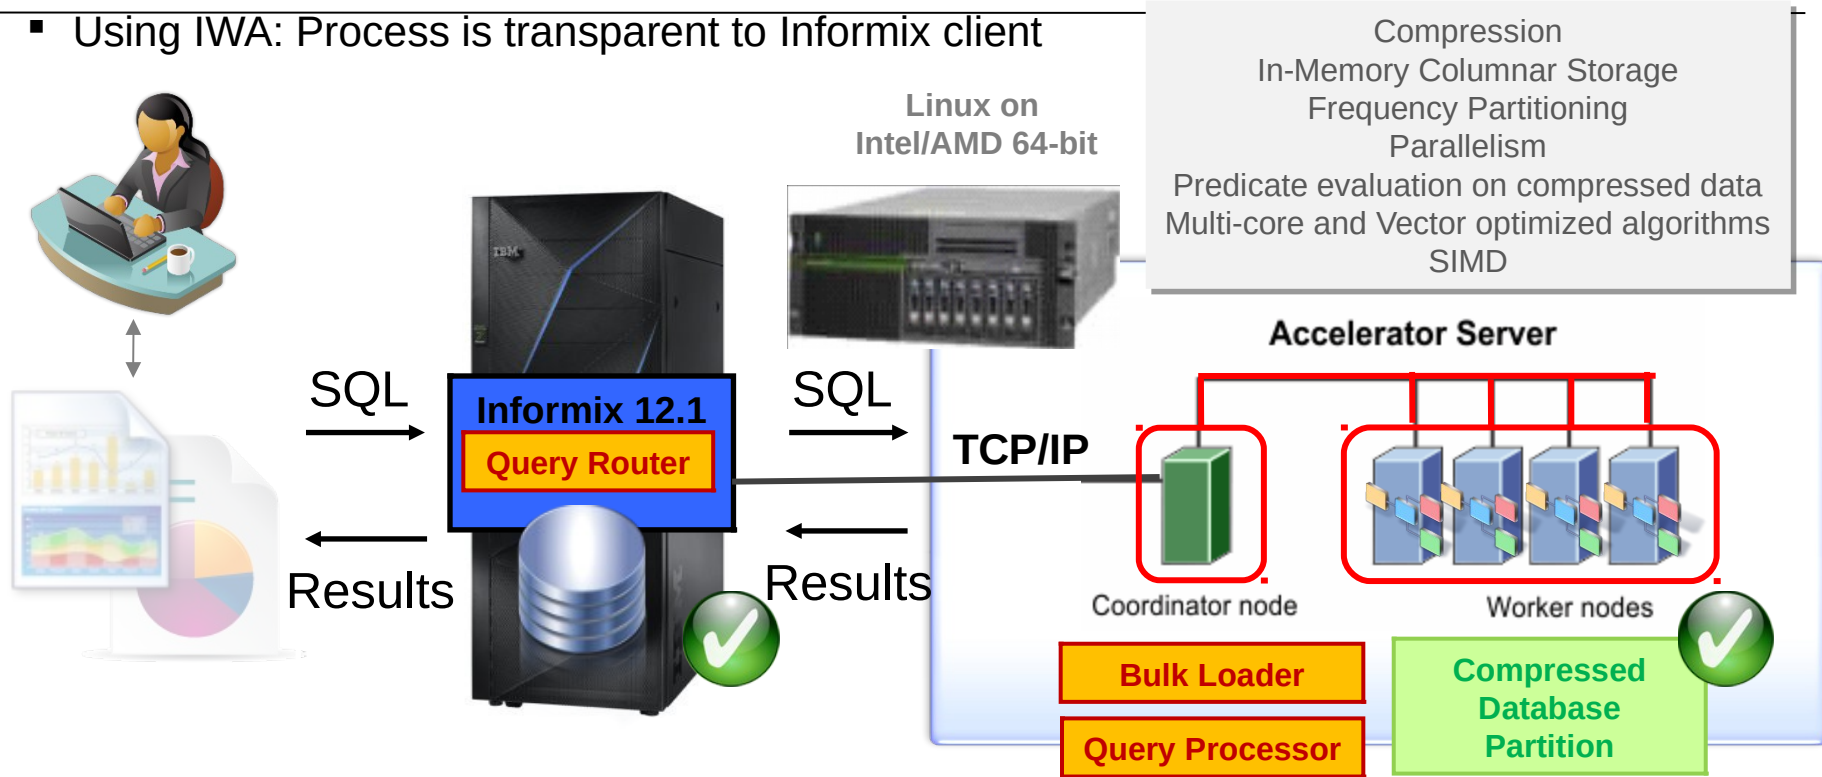
- Determine the database subset used in analytic queries to accelerate (data mart)
 - Manually or through Workload Analysis
- Deploy an IWA data mart based on DB subset
 - Stream Load the data from Informix into IWA
- Informix Optimizer is aware of IWA datamart

The Accelerator

- Install IWA on Linux x86_86 or Openpower
- Connect with Informix using custom protocol
- IWA compresses and stores a copy of the Informix DB set into data marts in-memory
- IWA data mart is fully loaded, valid and ready for Informix server to use as needed

IWA Overview and Seamless Integration with Informix/IDS

- Using IWA: Process is transparent to Informix client



Informix

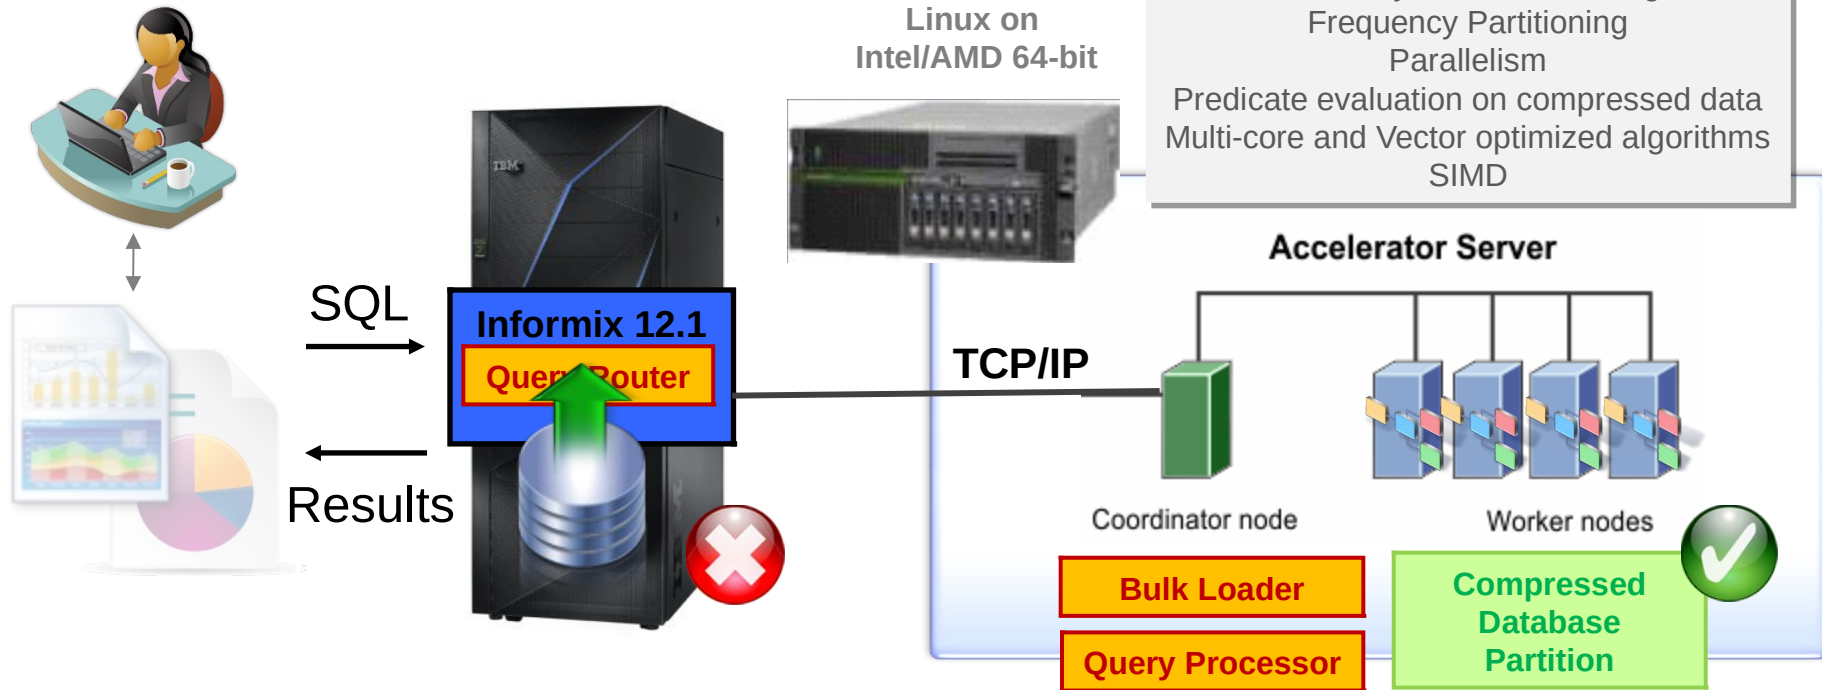
- Receives analytic query from client
- If query uses data matching an IWA datamart and can be accelerated, route/offload it to IWA
- Returns results back to the client

The Accelerator

- Processes the routed SQL query extremely fast and returns answer back to Informix

IWA Overview and Seamless Integration with Informix/IDS

- Using IWA: Process is transparent to Informix client



Informix

- Receives analytic query from client
- If query uses data matching an IWA datamart and can be accelerated, route/offload it to IWA
- Returns results back to the client
- If query is not based on an IWA datamart or cannot be accelerated, Informix will resolve it

The Accelerator

- Processes the routed SQL query extremely fast and returns answer back to Informix

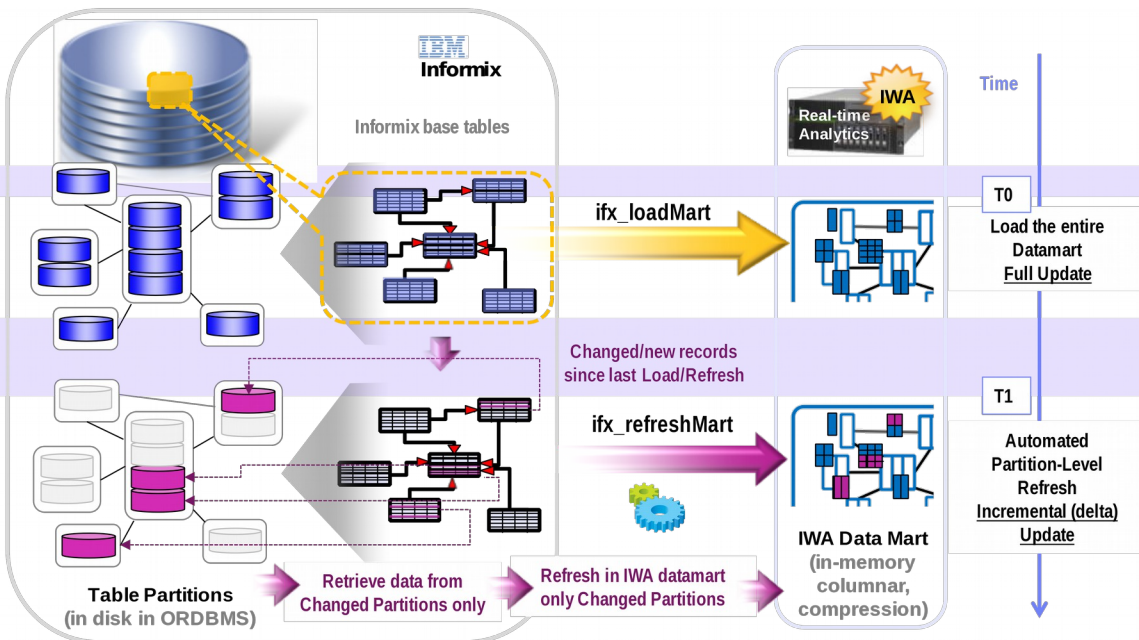
Automatic Partition Refresh

With this enhancement...

- **A single command** instructs IWA to **refresh only changed data partitions** from Informix database to IWA
- Applies to Fact and Dimension tables

Benefits...

- It removes the potentially error-prone process for manual identification of changed partitions in Informix
- Easier administration for keeping current Informix data in IWA



Automatic IWA synchronization with Informix, on-demand

Let Informix find the changed partitions since last refresh and refresh them in IWA for you

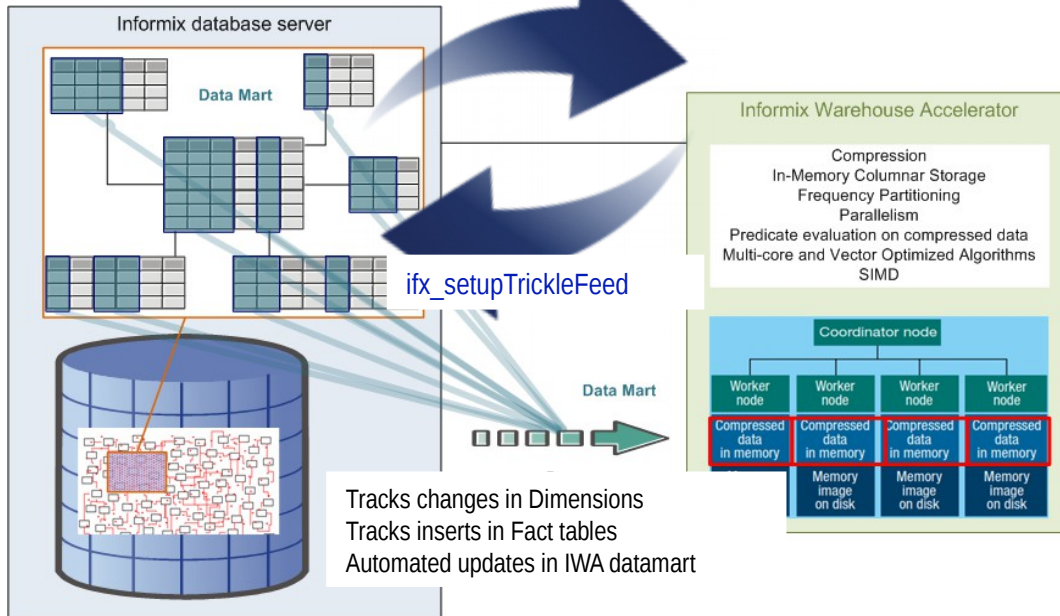
Easy adoption and maintenance of Real-Time Analytics

With this enhancement...

- **Incremental inserts** to the Fact tables and **changes** to Dimensions tables can be **continuously updated into IWA**
- Changes can be at **row level**, which is more granular than at partition level

Benefits...

- We can have “speed of thought” analytics in a real-time data warehouse or mixed workload environment
- Actionable analytics on operational data



Low Administration

Automated and fast small updates
keep IWA data current

Allows for Real/Right-Time
Analytics and Operational BI

With this enhancement...

- You can **include Time Series data** coming from smart sensors into IWA
- Data marts in IWA can be defined and loaded from an Informix' s Virtual Table Interface (VTI) of your TimeSeries data

Benefits...

- High-performance right-time analytics on big data collected from your sensors, meters, events, GPS/location, RFIDs, to anticipate and improve actions
- Combine TimeSeries and IWA for operational actionable analytics based on historic and current sensors data

Informix TimeSeries table

Meter_Id	TimeSeries column
1	[(4,160,40), (4.5,155,35), (5,165,33)...]
2	[(5,175,44), (4.5,160,35).]
--	--

Virtual Table Interface (VTI)
representation of TimeSeries table

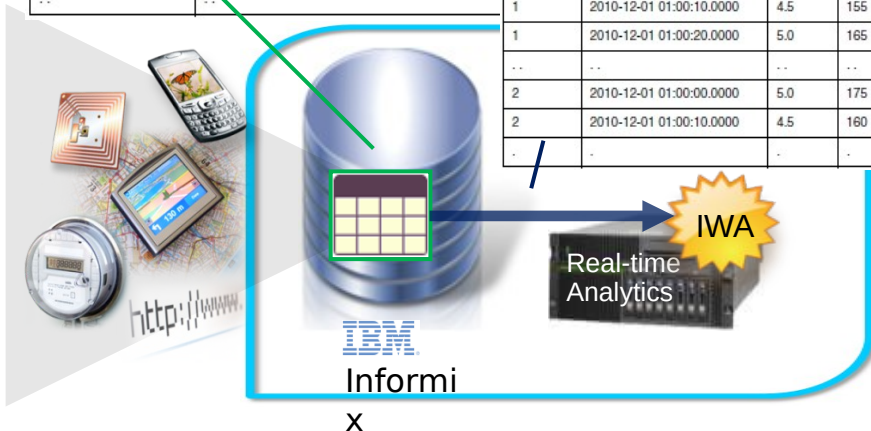
Meter_Id	Timestamp	Current	Voltage	Resistance
1	2010-12-01 01:00:00.0000	4.0	160	40
1	2010-12-01 01:00:10.0000	4.5	155	35
1	2010-12-01 01:00:20.0000	5.0	165	33
--	--	--	--	--
2	2010-12-01 01:00:00.0000	5.0	175	44
2	2010-12-01 01:00:10.0000	4.5	160	35
-	-	-	-	-

Right-Time Analytics on time-stamped data

Big Data solutions on Sensor data

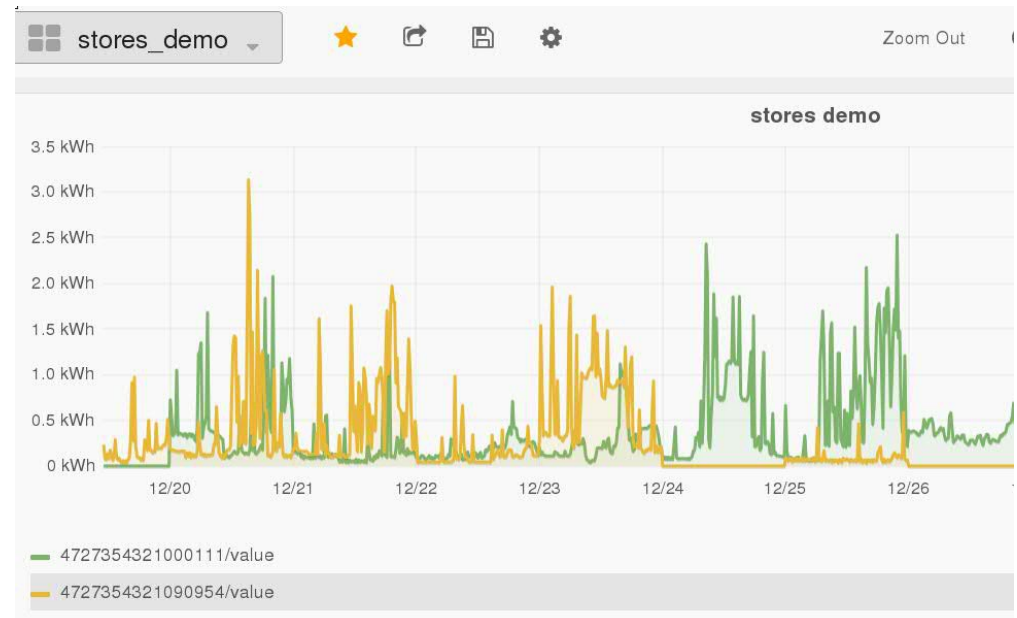
Operational Intelligence

- Unique platform, flexible, fast and scalable, for the most challenging Big Data and Smart Planet solutions



Demo with Grafana

- Grafana is a dashboard builder for visualizing time series metrics
- Open Source
- Grafana is most commonly used for visualizing time series data for IoT
- Visualize industrial sensors, home automation, weather, and process control
- Informix Timeseries plugin for Grafana.



Demo with Grafana

- Shell script to setup a mart for Grafana demo

```
DWA=DWA
DB=stores_demo
MART=iiug
Q='select * from ts_data_v;'

dbaccess -e ${DB} - <<!!!
set environment use_dwa 'probe start';
set explain on avoid_execute;
${Q}
set explain off;
set environment use_dwa 'probe stop';
execute procedure ifx_probe2mart('${DB}','${MART}');
execute function ifx_createmart('${DWA}','${MART}');
execute function ifx_loadmart('${DWA}','${MART}','NONE');
```

IWA data mart supports External Tables

With this enhancement...

- We can **load data directly from external tables into IWA data marts** without having to load it into Informix database first
- **Large external data** in ASCII / binary files or network devices –ex: through named pipes– can be used to quickly populate an IWA data mart

Direct, fast and flexible way to leverage external data for in-memory analytics

Benefits...

- Run extremely fast in-memory analytic queries on operational data from non-Informix external files and devices.
- Large amount of external data is quickly loaded and made available in IWA, thanks to high performance reads of Informix External Tables and the no need for the external data to be loaded into Informix database first
- Storage savings and flexibility to do in-memory analytics on large data in file systems or devices and integrate it with other SQL and NoSQL data in Informix.

IWA data mart supports External Tables JOINS!!

```
create table f(f1 int); insert into f values(2508);
create external table ext_f sameas f using (datafiles('disk:/tmp/f.data'))
insert into ext_f select * from f;
create table d(f1 int);insert into d values(2508);
create external table ext_d sameas d using (datafiles('disk:/tmp/d.data'));
insert into ext_d select * from d;
```

```
set environment use_dwa 'probe cleanup';
set environment use_dwa 'probe start';
select {+ avoid_execute} * from ext_f, ext_d where ext_f.f1=ext_d.f1;
set environment use_dwa 'probe stop';
execute procedure ifx_probe2mart('$DB','$SMART');
execute function ifx_createmart('$DWA','$SMART');
execute function ifx_loadmart('$DWA','$SMART','NONE');
```

Direct, fast and flexible way to leverage external data for in-memory analytics



IWA data mart supports synonyms and views

With this enhancement...

- Until now, an IWA datamart could only contain regular local tables, all in the same database
- We can now create an IWA datamart that **uses remote tables** in another Informix DB and accelerate queries using those remote tables
 - By having a local **synonym** in the Informix DB of the datamart, which points to the remote Informix table
- We can include **views** as part of an IWA datamart definition, and accelerate queries that use views
 - Views could map to a subset of another **Informix table** or also to **NoSQL data**
 - Use views to accelerate **self-joins**

Benefits...

- Ability to combine and accelerate queries on local with remote tables, no need to make all tables local
- Allows to be able to do accelerate self-joins queries by using views
- Allows to accelerate data in JSON collections by using views
- Fast analytic queries on views, typically slow in SQL DBs due to on-the-fly view materialization

Fast analytic queries on SQL and NoSQL data
Analytics on data from different sources

More queries and datatypes can be accelerated



IWA data mart supports of views and usage in NoSQL query (1)

Example: Accelerating NoSQL data (in a JSON collection)

▪ From MongoDB shell:

- Create two collections (JSON): comments and users

```
$ mongo demo_database
MongoDB shell version: 2.4.9
connecting to: demo_database

mongos> db.comments.insert( [
  { uid:12345, pid:444, comment:"first" },
  { uid:12345, pid:888, comment:"second" },
  { uid:99999, pid:444, comment:"third" }
] )

mongos> db.users.insert( [
  { uid:12345, name:"john" },
  { uid:99999, name:"mia" }
] )

mongos> exit
```

▪ From Informix:

- Create a view on each JSON collection (comments, users)

```
$ dbaccess demo_database -
> create view vcomments(uid,pid,comment) as
select
  bson_value_int(data,'uid'),
  bson_value_int(data,'pid'),
  bson_value_varchar(data,'comment')
from comments;
> create view vusers(uid,name) as select
  bson_value_int(data,'uid'),
  bson_value_varchar(data,'name')
from users;
```

- Deploy an IWA data mart by probing a join between them:

```
set environment use_dwa 'probe cleanup';
set environment use_dwa 'probe start';

select {+ avoid_execute} * from vcomments c,
vusers u where c.uid=u.uid;

set environment use_dwa 'probe stop';

execute procedure
ifx_probe2mart('demo_database','noSQL_mart');
execute function
ifx_createmart('demo_dwa','noSQL_mart');
execute function
ifx_loadmart('demo_dwa','noSQL_mart','NONE');
```


The banner features a blue background with a bright light source on the left, creating a lens flare effect. The light illuminates a series of glowing blue circles and lines that recede into the distance, suggesting a digital or data flow. Various terms related to technology and databases are scattered in the background, including 'reduce cost', 'cloud', 'mobile', 'embeddable', 'applications', and 'PureSystems'.

reduce cost
cloud
mobile
embeddable
applications
PureSystems

IBM Informix 12.1
Simply Powerful

IBM Informix database 12.1

Relational, Embeddable, Hybrid-Database
InMemory (columnar)
Fast, Always-on Transactions
NoSQL capability
Multi-Tenancy
Sensor data management
High availability
Easy to Use

A decorative footer bar consisting of a series of overlapping, semi-circular shapes in shades of blue, orange, and red.

Conclusion

- IWA and Timeseries provides fast in-memory analytics for IoT data
- Can handle all kind of SQL not only Warehouse!
- you can kill runtime performance challenges with HW



THANK
YOU

