

Deep Dive on Informix

Kevin Brown, Informix Chief Architect

kbrown3@us.ibm.com



Agenda

- Informix TimeSeries
- Informix Ultimate Warehouse Edition (IUWE)
- Informix Genero Edition



TimeSeries

Overview

- Why do we need Time Series?
- A little technology
 - Background on IBM Informix relational database with time series support
- Customer examples

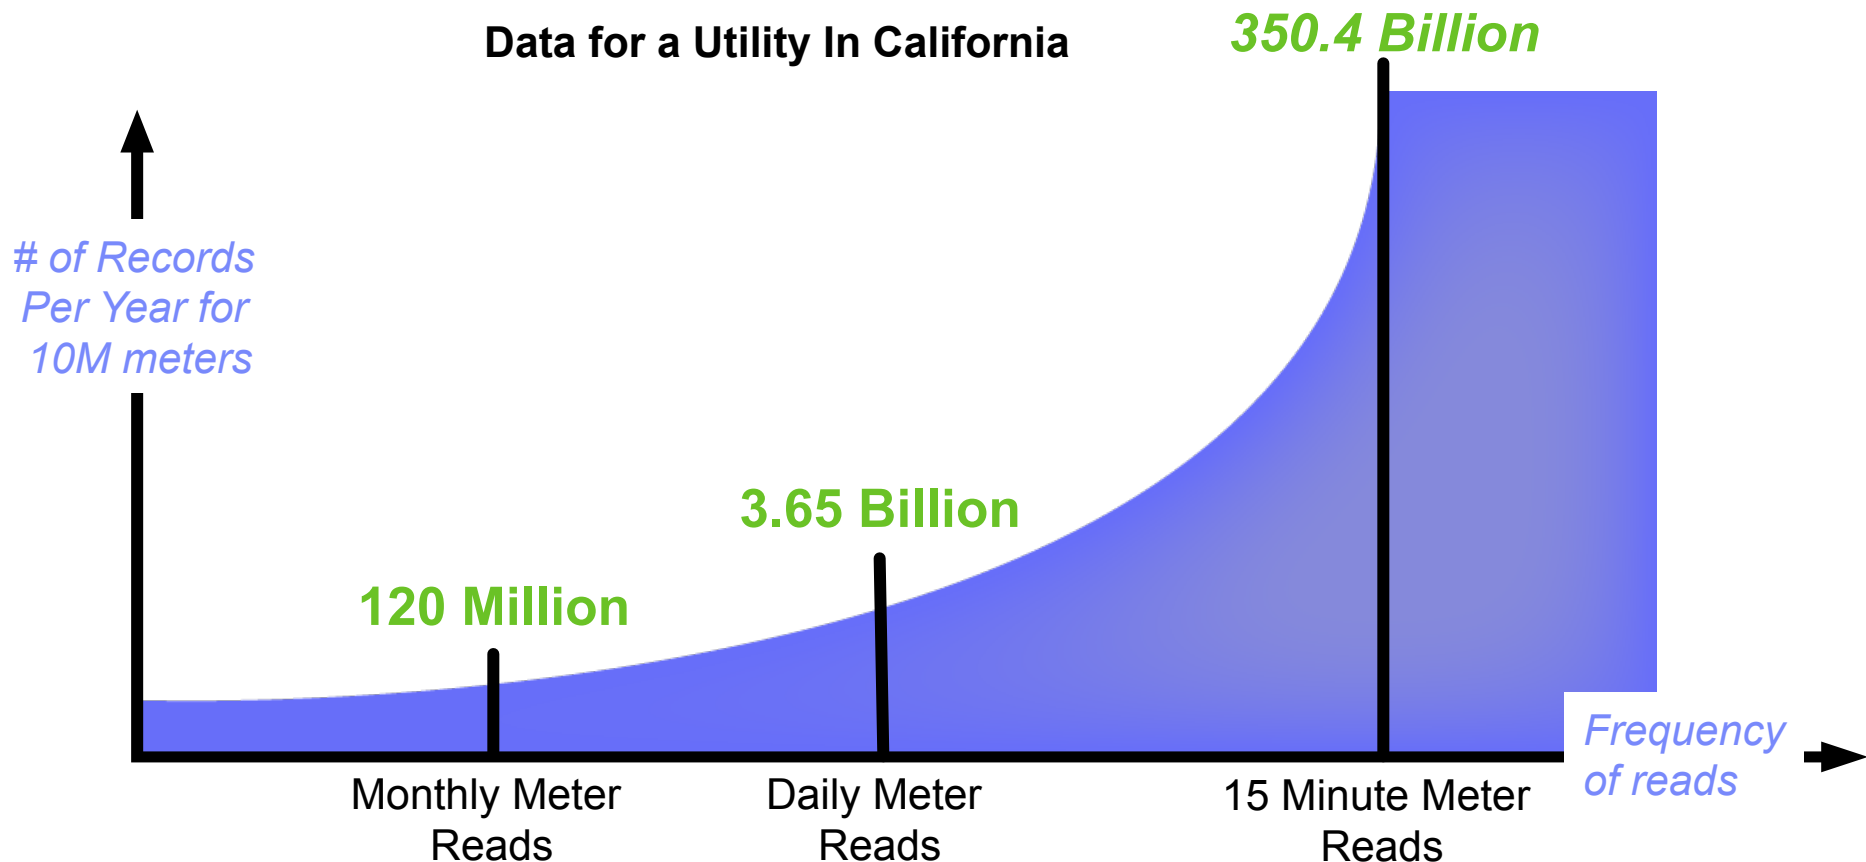
Example of Changing Storage Requirements

Changing Workloads For 10 Million Smart Meters:

Today – Each meter is read once per month

Very soon – Each meter is read once every 15 minutes

Regulations – Need to keep data on line for 3 years (PUC) and, perhaps, save for 7 years



Keeping up with Smart Meter Data

Large amounts of data causes problems in 2 areas:

1. Storage management

- Will get expensive and cumbersome to maintain

2. Query performance

- Compliance Reports must be completed before the end of each day
- Customer portal queries must be handled in a timely manner
- Customer billing

Introducing IBM Informix Relational Database

■ Why IBM Informix?

- Low cost/Low administration
 - Manage thousands of servers with one database administrator
 - Many installations have no DBA's
- High performance
 - Insert 100's of thousands of records per second
 - Analytic functions not available in other database products
- High Availability
 - Scale up, scale out, and disaster recovery
- Native time series data support
 - Load and analyze time series data

“The idea was to run the project in three cycles, and use a different type of analysis in each cycle, to see if we could draw any conclusions in terms of how to promote change in the way people view their energy consumption. With help from the IBM Hursley team, we quickly found that Informix TimeSeries could deliver spectacular results.” Clive Eisen, Chief Technology Officer at Hildebrand.

Key Strengths of Informix TimeSeries




- Performance
 - Extremely fast data access
 - Data clustered and sorted by time on disk to reduce I/O
 - Handles operations hard or impossible to do in standard SQL
 - Continuous Real-Time and Batch Data Loaders

- Space Savings
 - Typically saves 50% space over standard relational layout

- Toolkit approach
 - Allows users to develop their own algorithms to run in the database
 - Algorithms running in the database leverage the buffer pool for speed

- Easier
 - Conceptually closer to how users think of time series

Typical Relational Schema for Smart Meters Data



Index Smart_Meters Table

	mtr_id	Time	KWH	Voltage	ColN
Table Grows 	1	1-1-11 12:00	Value 1	Value 2	Value N
	2	1-1-11 12:00	Value 1	Value 2	Value N
	3	1-1-11 12:00	Value 1	Value 2	Value N

	1	1-1-11 12:15	Value 1	Value 2	Value N
	2	1-1-11 12:15	Value 1	Value 2	Value N
	3	1-1-11 12:15	Value 1	Value 2	Value N

- Each row contains exactly one record = billions of rows
- Additional indexes are required for efficient lookups
- Data is appended to the end of the table as it arrives
- Meter Id' s stored in every record
- No concept of a missing row

Same Table using an Informix TimeSeries Schema

Smart_Meters Table



mtr_id	Series
1	[(1-1-11 12:00, value 1, value 2, ..., value N), (1-1-11 12:15, value 1, value 2, ..., value N), ...]
2	[(1-1-11 12:00, value 1, value 2, ..., value N), (1-1-11 12:15, value 1, value 2, ..., value N), ...]
3	[(1-1-11 12:00, value 1, value 2, ..., value N), (1-1-11 12:15, value 1, value 2, ..., value N), ...]
4	[(1-1-11 12:00, value 1, value 2, ..., value N), (1-1-11 12:15, value 1, value 2, ..., value N), ...]
...	...

Table grows 

- Each row contains a growing set of records = one row per meter
- Data append to a row rather than to the end of the table
- Meter Ids not stored in individual records
- Data is clustered by meter id and sorted by time on disk
- Missing values take no disk space, missing interval reads take 2 bytes

Real Customer Comparison TimeSeries vs. Relational (Major Texas Power Distributor)

- Simulated Environment
 - 1 million meters
 - 90 days worth of meter data
 - 15 minutes intervals
 - 200 locations
 - 500 feeders
 - 34 substations

- Hardware/OS Used
 - Power7 with 2 sockets each with 8 cores
 - 64 bit SUSE Linux v11
 - 128 GB of memory
 - Memory actually needed = 44GB
 - 6 disks dedicated to the database
 - Only 350GB disk space used by database
 - 2 additional disks for OS and input file staging

- Software Used
 - Informix Ultimate edition
 - Informix TimeSeries



Why Did they Contact Us?

- **Competitor was barely able to load and process the data in 24 hours**
 - It was taking the competitor about 5-7 hours to load and validate the data for 1 million meters
 - The *ERCOT compliance reports were each taking many hours to run
 - They were having problems running reports while data was being loaded

- **Customer was looking for ways to cut their costs**
 - Competitor used about 1.3 TB to store 90 days worth of data
 - Estimate for competitor to handle 3 years of data for 3.5 million meters was about 55 TB
 - Customer wanted to do more with the hardware they had
 - Process the data more quickly to allow other applications to run against the data

- **Customer was worried about their future**
 - In the next 1 to 2 years they expect to expand to manage 3.5 million meters
 - Their current 73 billion records stored would become 220 billion records stored
 - Even if the competitor scaled linearly there would not be enough time in the day to read and process all this data

Results of the POC with Informix TimeSeries

- Informix took about 18 minutes to validate and load a day's worth of data for 1 million meters
 - Competitor took about 5 hours
- Informix took about 6 minutes for each mandatory ERCOT compliance report and about 25 seconds if the data was already cached in memory
 - Competitor took from 2 to 3 hours depending on the report
- Disk space used by Informix was about 350GB
 - Competitor used about 1.3TB for the POC
 - Their estimate for storing 3 years worth of 3.5 million meters was 55 TB
 - Our estimate for 3 years worth of 3.5 million meters using the Informix TimeSeries was 15 TB
- Results were very linear for Informix
 - Better results if you increase cpus and storage
 - If less performance is OK then cpus and storage can be reduced

“We are committed to improving energy efficiency by adopting the latest digital technologies providing Smart Metering solutions to our customers. Our major challenges are the significant storage space and performance required to manage the data collected by the smart meters. With Informix and its native support for time series data, we believe we will meet both of these challenges. Our initial testing has shown that with an Informix based solution we are able to reduce the storage requirements by two thirds and speed the query performance up to 60 times compared to the existing solution.”

Director Technology, Strategy, & Architecture

Client Success: Hildebrand



Hildebrand is a technology consultant company on the Digital Environment Home Energy Management System (DEHEMS) project. The Hildebrand team was asked by the UK government to find a way to scale up its energy monitoring solution and enable it to monitor three million homes.



What's Smart?

- Energy management thru real time monitoring
- Ability to monitor scales to over 3 million homes
- Reduced energy

Business Benefits

- Helping people make better decisions about energy efficiency in the home.
- Collect, store and analyse up to • 50,000 data points per second
- Delivers high performance on low-cost hardware by leveraging Informix time series data management technologies.

Hildebrand: 3 Million Meters



- Test involved 3,000,000 Homes
 - Data generated every **6 seconds!**

- Up to 26 values recorded
 - Meter ID
 - Timestamp
 - 3 Electricity phases
 - 1 Gas reading
 - 20 Individual electrical sockets in the house

- Data collected every 6 seconds
 - Aggregated to per-minute readings
 - Minute readings bulk-loaded into Informix every 10 minutes
 - Average load was about 50,000 inserts per second

- **Hardware/Software Used**
 - Intel with 8 cores running:
 - 64 bit SUSE Linux v10
 - 16 GB of memory
 - Informix workgroup edition
 - Hildebrand software

Consumer Education



DEHEMS
An EU funded climate change initiative

SEVENTH FRAMEWORK PROGRAMME

Last Updated: Tue Feb 10 2009 09:17:40 GMT

Current live energy usage

Your Home	0.43 kW
Local Average	0.47 kW

Energy Saving

Compared to your last bill we currently estimate the difference in your spending is:

+£7.96
per month

Usage so far this month

kWh Cost CO₂

Your Home	114kg
Local Average	248kg

Energy Fact

Set-top boxes for the reception of digital TV are inexpensive to buy, being given away for free with subscription packages, but by 2010 they could be costing UK households £472m every year in electricity or £19 per household

Energy Usage

Your home Local Average

24 Hour Hourly Daily

Hildebrand Results/Conclusions



- People want to be “green”
- Immediate impact in behaviour change
- New behaviour seems to be sustained
- Big scope for community functions
- People want to know more
- Target of 20% savings is very achievable

AMT/SYBEX

Proving the concept for smart metering data management



- The Need:
 - Extend AMT/Sybex’s Data Transfer Solution to:
 - Load, validate, store, and provide smart meter interval and event data to external system
- The Challenge:
 - Process the enormous volume of data in a timely manner
 - Not to require a huge investment in new hardware
- The Solution:
 - Combined effort with IBM Research and IBM Informix
 - Informix TimeSeries provides the core capabilities to store and process the meter and event data:
 - VEE
 - Real-time energy monitoring
 - Analytics for developing new tariff rates
 - Help to smooth peaks in demand

“We owe a great deal to IBM, both for the Informix technology itself, and for the fantastic support from all levels of the organisation.”

– Gordon Brown, DTS Product Owner, AMT-SYBEX

Solution components:

- IBM® Informix® TimeSeries™
- IBM Research



AMT/Sybex Tests and Performance Measurements


The main tests for AMT-SYBEX were to prove key business functions at high volume:

- Can high volumes of data be loaded into Smart DTS?
- Can data be analyzed and processed in Smart DTS in a timely manner?

To facilitate these requirements, the following tests were carried out:

- These results are for 10 million meters storing data every ½ hour data for one day.

Module	Time	Readings/intervals per second
Technical Validation and Transformation	13 minutes	>600,000
High Speed TimeSeries database load (with full logging)	50 minutes	>150,000
Validation and Estimation	33 minutes	>240,000



Full processing end to end for 10 million meter points with half-hourly interval data on a single mid-range P-series 8 CPU server is just over 1 ½ hours

AMTSybex[®]
Enabling essential industries

IUWE

Informix Ultimate Warehouse Edition

Informix Warehouse: Then and Now



Informix Warehouse Feature

- SQW
- Data Modeling
- ELT/ETL

Informix Warehouse with Storage Optimization/Compression

Cognos integration
- Native Content Store on IDS
SQL Merge

External Tables

Star Join Optimization
Multi-index Scan
New Fragmentation
Fragment Level Stats
Storage Provisioning

Warehouse Accelerator



Now Introducing... Informix Ultimate Warehouse Edition

What it is

- New Informix edition for warehouse
- Package contains:
 - Informix Ultimate Edition
 - Informix Warehouse Accelerator (IWA)*
 - IBM Smart Analytics Optimizer Studio
 - Compression Feature included

What IWA does

- Compresses and processes data mart entirely in memory
- Uses a deep columnar approach for optimizing performance
- Runs as an add-on to Informix

Benefits

- Order of magnitude or more performance gains
- Transparent to BI apps – queries continue to run against Informix
- No maintenance

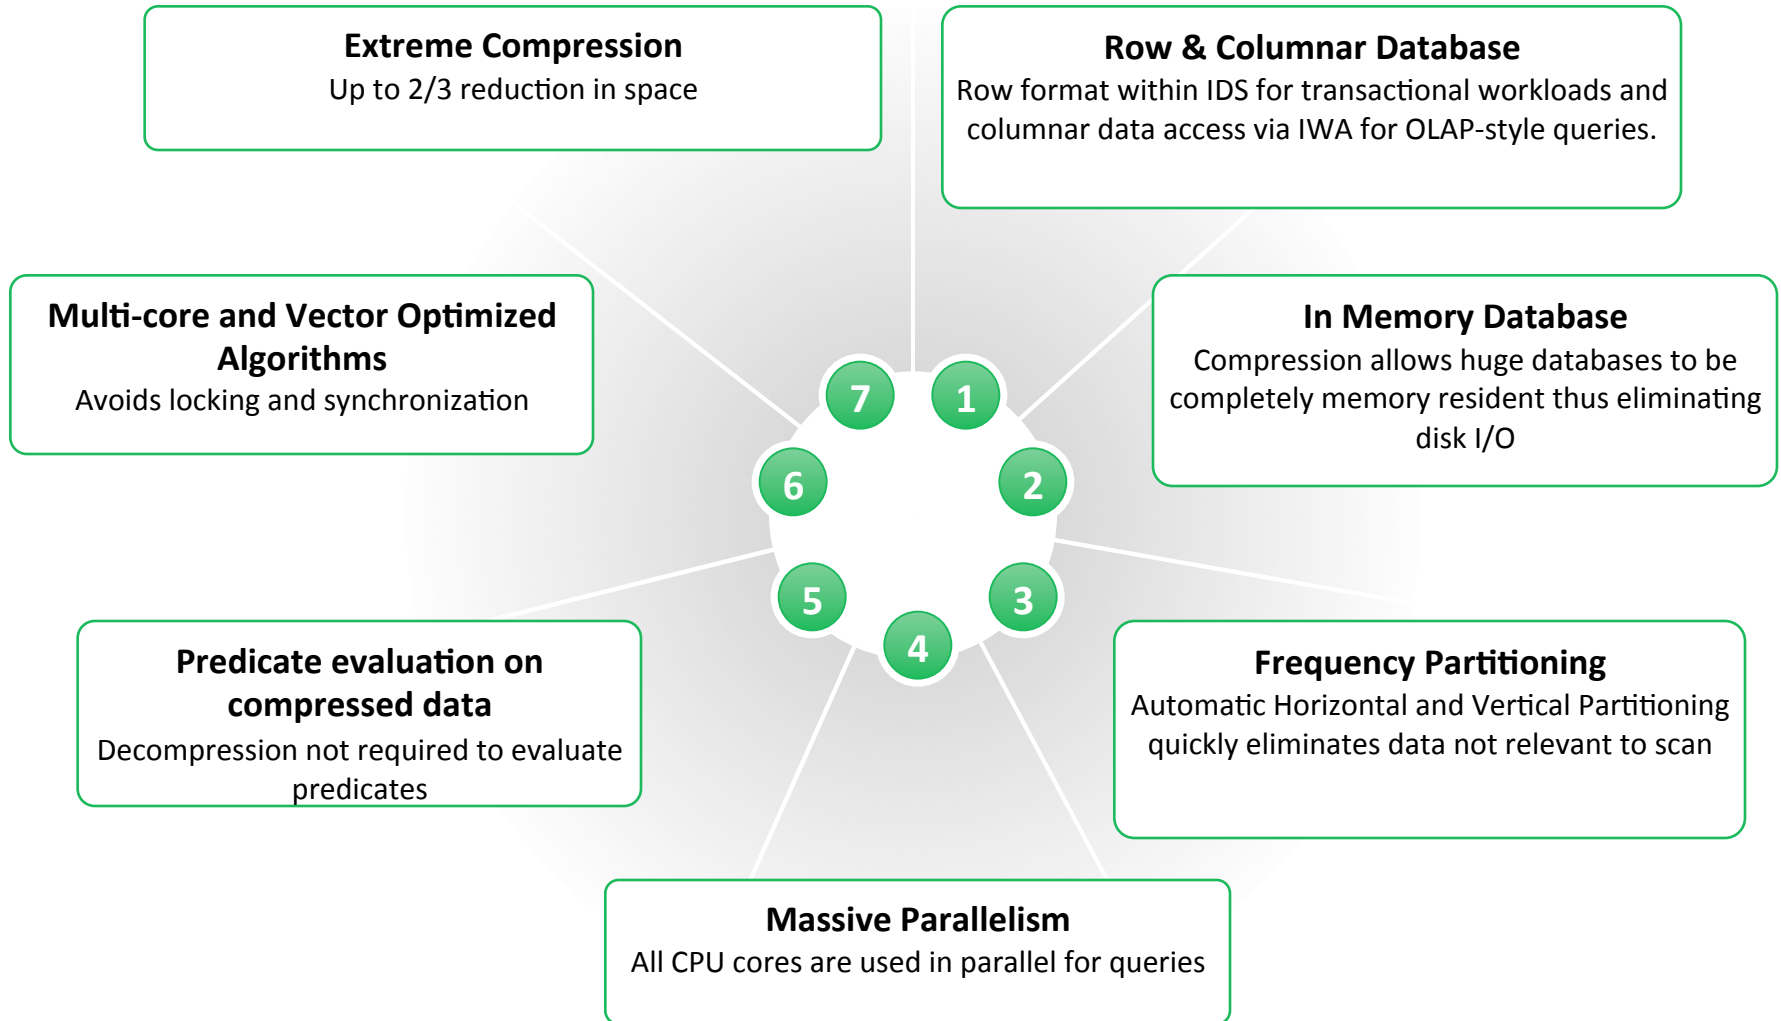
IWA requires you to do...

- No Query Tuning and Optimizer Hints
- No Database tuning
- No Index creation, reorganization
- No Update Statistics
- No Partitioning/Fragmentation
- No Storage Management/Page size configuration
- No Database/schema Changes
- No Application Changes
- No Summary Tables/Materialized Views
- No Buying more expensive Hardware
- No Change of expectations

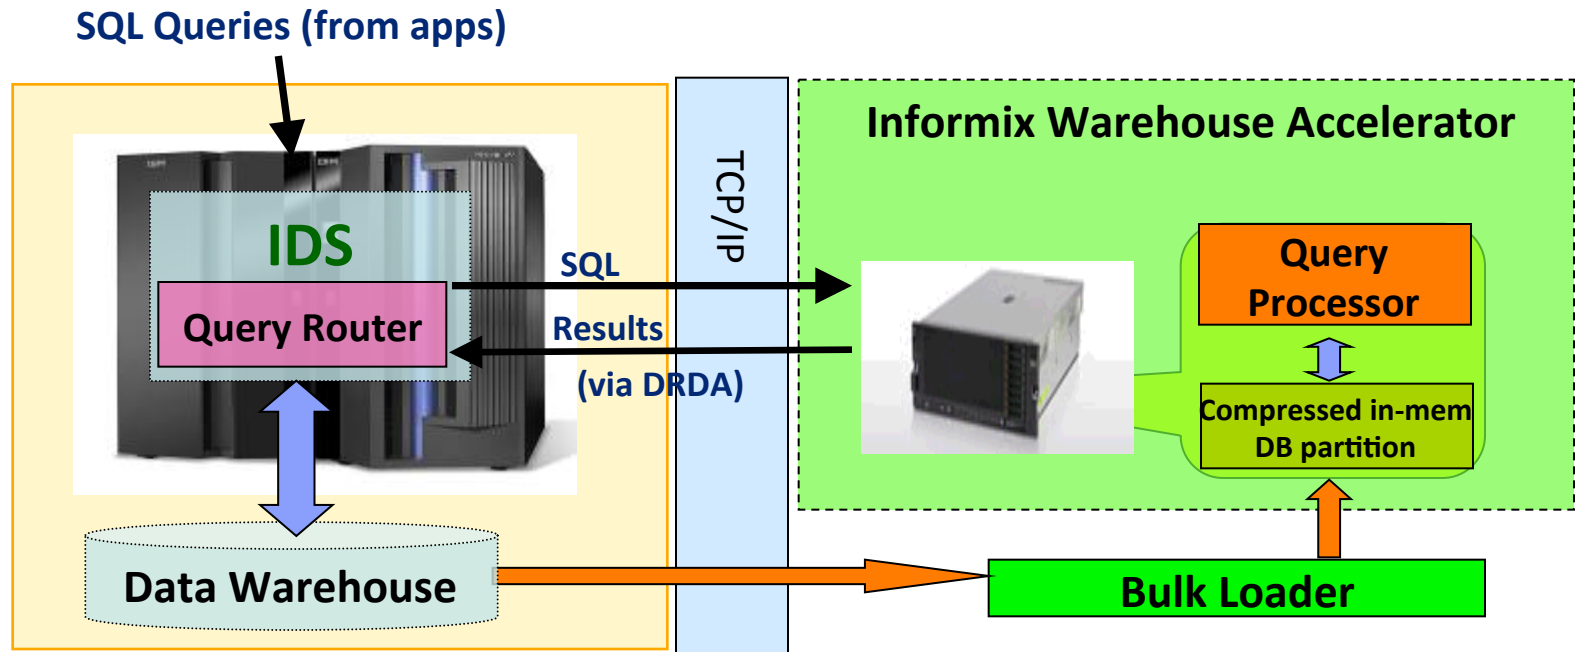
Power of Simplicity!



IWA: Breakthrough Technologies for Performance



Informix Warehouse Accelerator Configuration



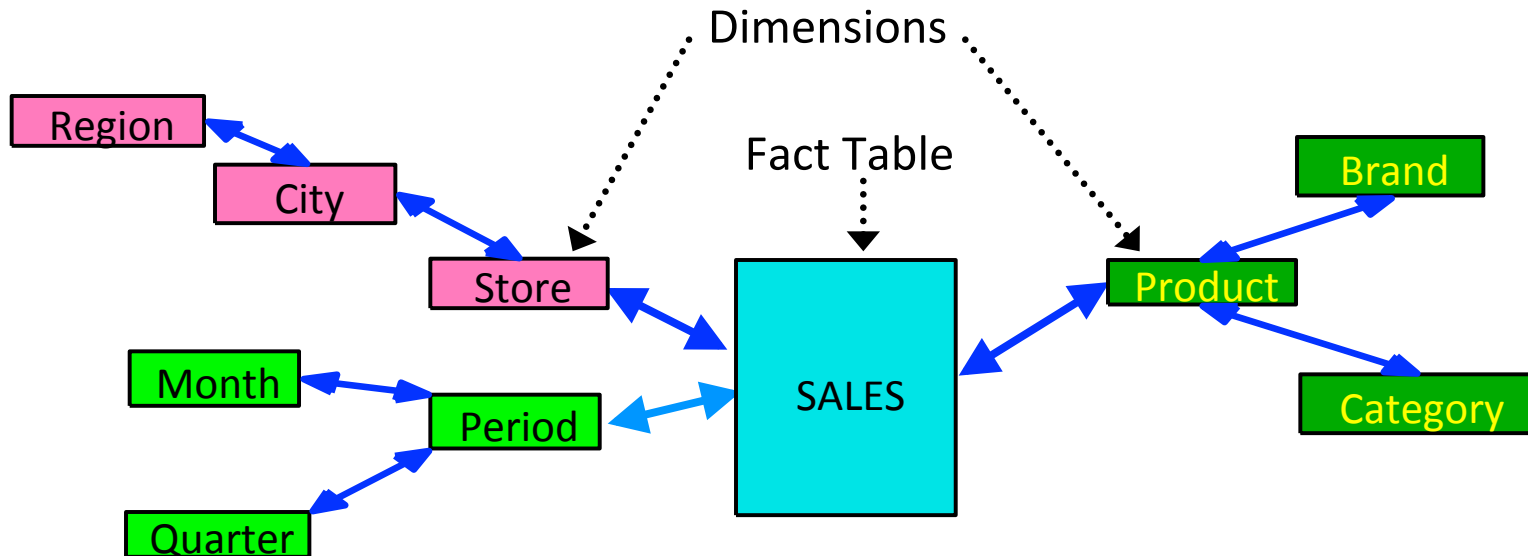
IDS:

- User need not change SQL or apps.
- Routes SQL queries to accelerator as needed and reads back results
- Can always run query in IDS, e.g., if estimated execution time is very short

Informix Warehouse Accelerator:

- Communicates with IDS via TCP/IP & DRDA
- Analyzes, compresses, and loads a portion of the Informix warehouse into IWA
- Processes routed SQL query and returns answer to IDS

IWA Sweet Spot: Queries against Star and Snowflake Schemas

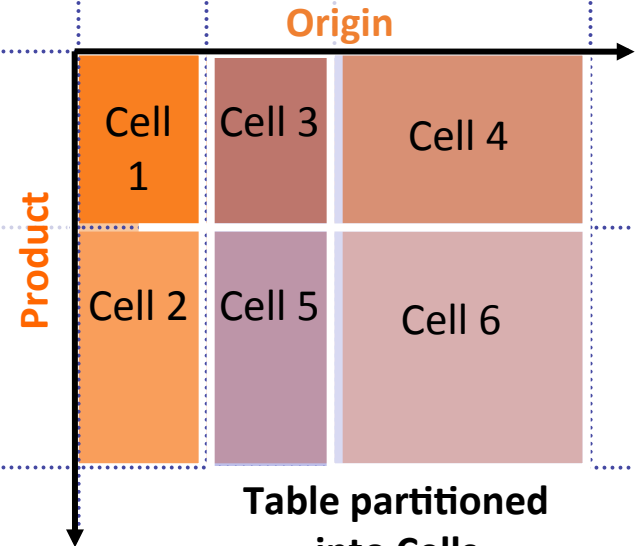
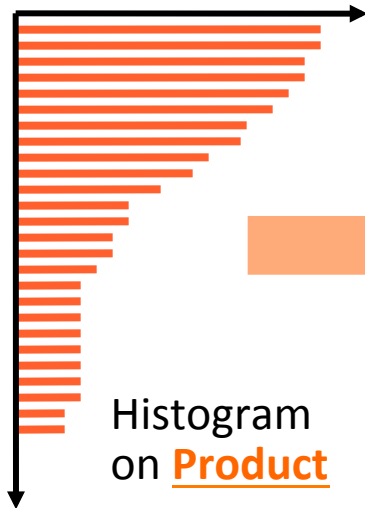
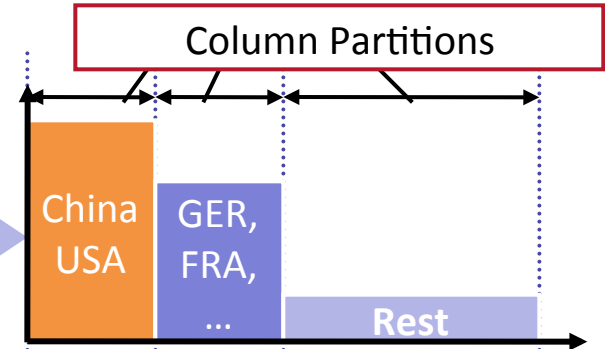
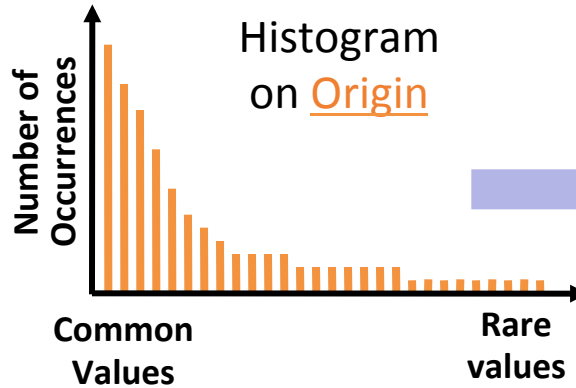



Complex, ad hoc queries that typically

- Look for trends, exceptions to make actionable business decisions
- Touch large subset of the database (unlike OLTP)
- Involve aggregation functions (e.g., COUNT, SUM, AVG,...)
- **The “Sweet Spot” for the IWA!**

Compression: Frequency Partitioning

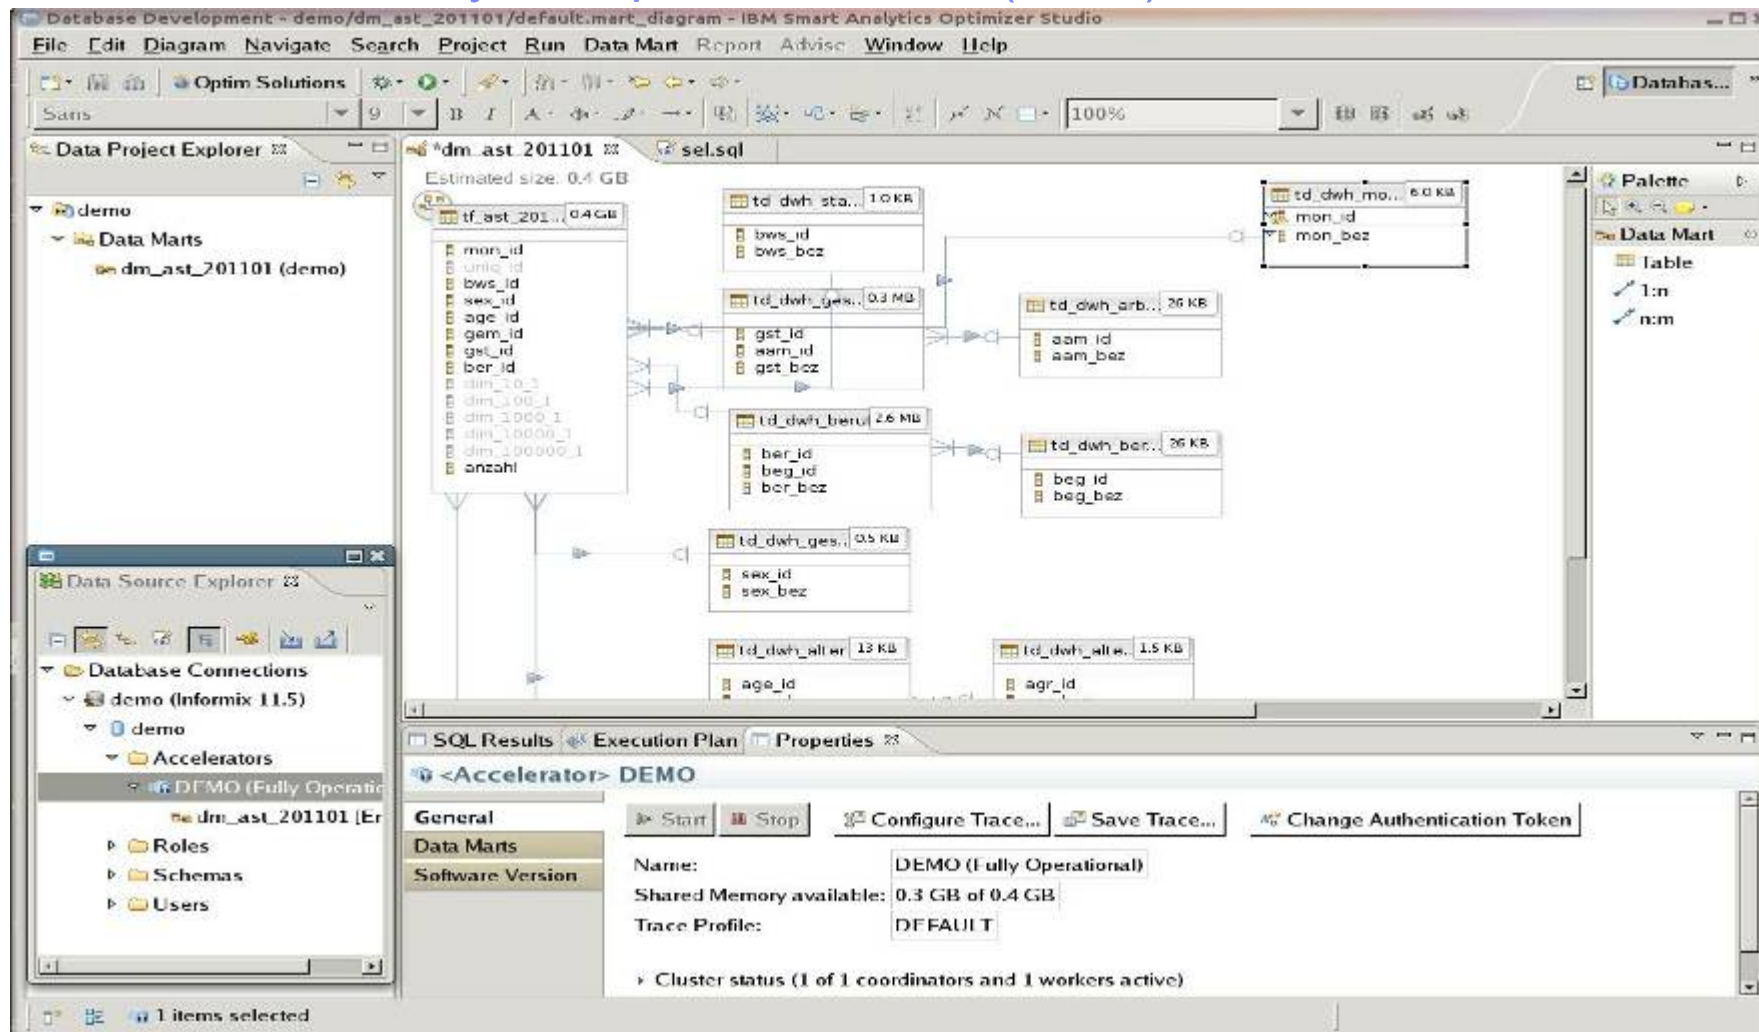
Trade Info (volume, product, origin country)



- Field lengths **vary between cells**
 - Higher Frequencies → Shorter Codes (Approximate Huffman)
- Field lengths **fixed within cells**

Table partitioned into Cells

IBM Smart Analytics Optimizer Studio (ISAO)



Easy to use graphical tool for data mart configuration for Linux and Windows (included with IUWE)

Case Study #1: Major U.S. Shoe Retailer

- Top 7 time-consuming queries in Retail BI and Warehouse: (Against 1 Billion row Fact Table)

Query	IDS 11.5	IDS 11.7 IWA
1	22 mins	4 secs
2	1 min 3 secs	2 secs
3	3 mins 40 secs	2 secs
4	30 mins & up	4 secs
5	2 mins	2 secs
6	30 mins	2 secs
7	45 mins & up	2 secs

“Our Retail users will be really happy to see such a huge improvement in the queries processing timings.”

“This IWA extension to IDS will really bring value to the Retail BI environment.”

Advanced Data Tools Tests

Run	Query	Type of Query	IDS Only	IDS & IWA	Improvement	Rows Return
1	Query12	Scan - Select First 1000	105m7.730	9m48.071	10.7261	1000
2	Query13	Index - by Dim	0m0.116	0m11.605	0.0100	9
3	Query14	Index - by Dim	0m0.055	0m11.971	0.0046	9
4	Query15	Index - by Dim	0m8.345	0m4.469	1.8673	9
5	Query0	Scan - Select First 1000	106m44.474	0m20.946	305.7612	1000
6	Query10	Index - by Dim	0m4.817	0m6.062	0.7946	1
7	Query11	Index - by Dim	0m3.430	0m6.068	0.5653	1
8	Query1	Scan - Select First 100	63m33.083	1m13.501	51.8780	100
9	Query2	Scan - Select First 100	65m17.075	1m29.518	43.7574	100
10	Query3	Scan - Sum by Dim	87m47.842	0m17.207	306.1453	3107
11	Query4	Scan - Select First 100	53m57.169	0m14.162	228.5813	100
12	Query5	Index - by Dim	0m0.049	0m12.692	0.0039	0
13	Query6	Scan - Select First 1000	1m48.886	0m9.495	11.4677	1000
14	Query7	Scan - by Dim w Group by	0m15.855	0m9.171	1.7288	26937
15	Query8	Scan - by Dim w Group by	0m14.846	0m2.347	6.3255	28358
16	Query9	Scan - by Dim w Group by	79m14.253	0m42.011	113.1669	523499
Total in hours			9hr 40min	15min 32sec		
Total in minutes						

IWA Referenced Hardware Configuration

Intel(R) Xeon(R) CPU	X7560 @ 2.27GH 4 X 8
Memory	512G
6 disks	300 GB SAS hard disk drives each



Options:

- 4-processor, 4U rack-optimized enterprise server with Intel® Xeon® processors
- 8-core, 6-core and 4-core processor options with up to 2.26 GHz (8-core), 2.66 GHz (six-core) and 1.86 GHz (four-core) speeds with up to 16 MB L3 cache
- Scalable from 4 sockets and 64 DIMMs to 8 sockets and 128 DIMMs
- Optional MAX5 32-DIMM memory expansion (up to a total of 1.5 TB of memory)
- 16x 1.8" SAS SSDs with eXFlash or 8x 2.5" SAS HDDs

IBM Genero

IBM Informix Genero - What is the Offering?

- The offering is driven through a partnership between:

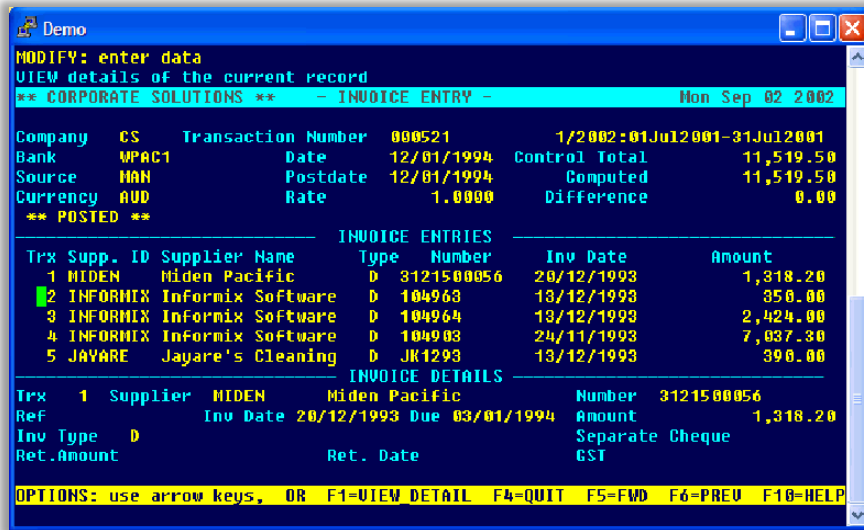


- Allows Informix 4GL users to deploy existing and new applications in:
 - Browser-based environments
 - GUI desktop environments
 - Mobile environments
 - Cloud environments
- Leverages the Four J's Genero product family:
 - Genero BDL with Web Services
 - Genero Application Server
 - Genero Desktop Client
 - Genero Report Designer and engine
 - Genero Studio server and client

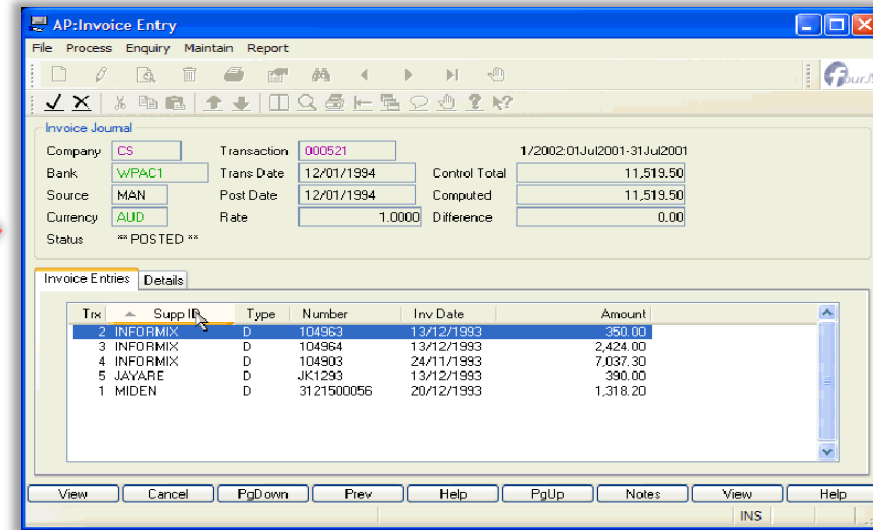


What is the Offering?

Informix 4GL



Informix Genero



Current Informix 4GL offerings:

- 4GL Compiler Development
- 4GL Interactive Debugger
- 4GL Compiler Runtime
- 4GL RDS Development
- 4GL RDS Runtime


New Informix Genero offerings:

- Genero Developer Suite
- Genero Runtime Suite

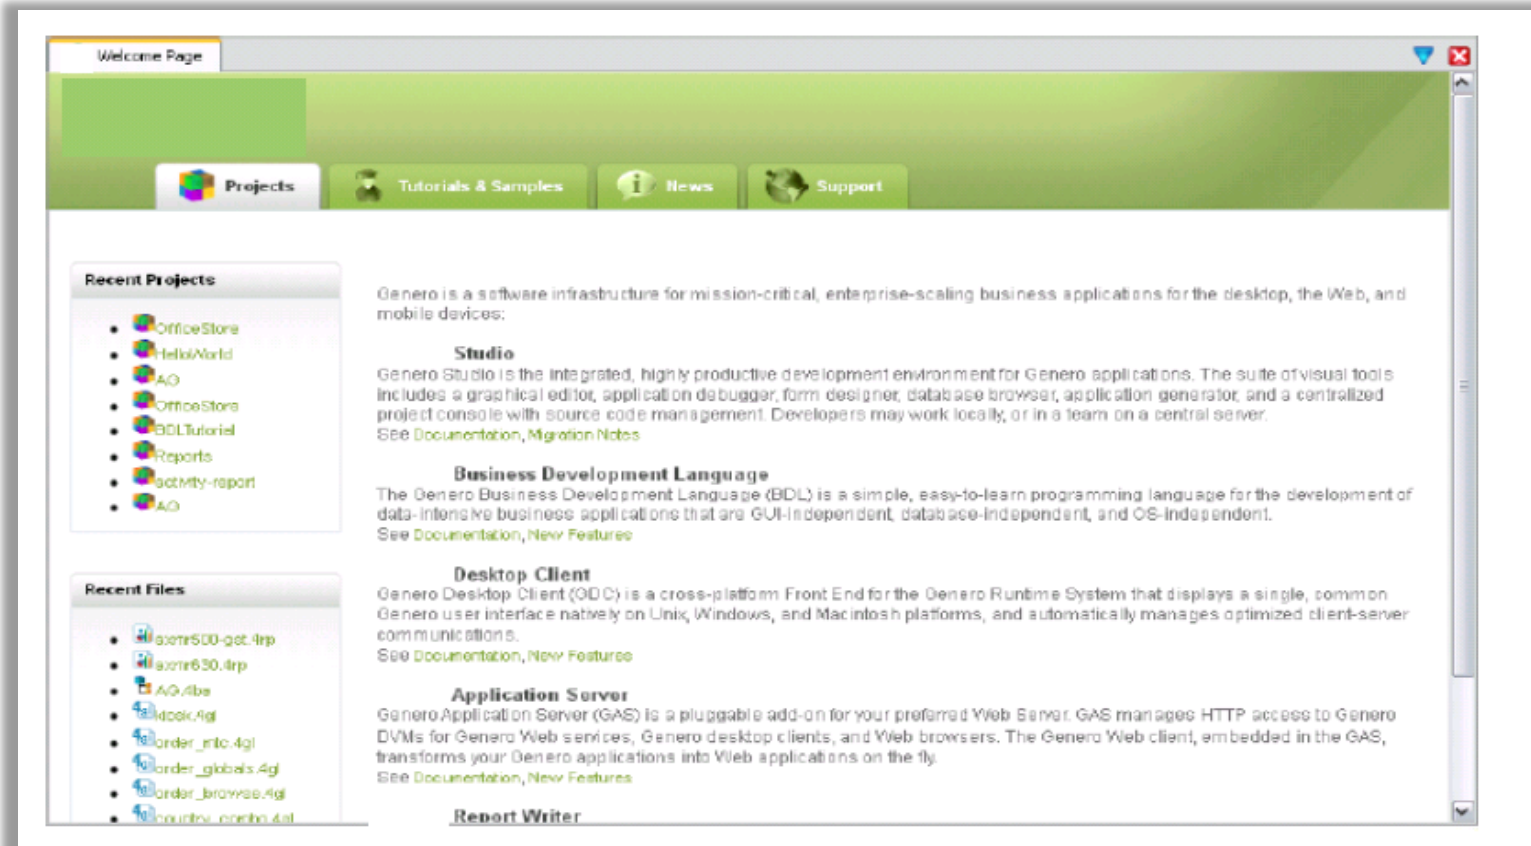
➔ 4GL continues to be available

What is in the Genero Offering?

- Modern Development Environment for 4GL
 - Graphical Development Environment (Genero Studio):
 - Compiler, debugger, 4GL-aware editor
 - Graphical display form builder
 - Application Generator
 - Source code control
 - Graphical Report Writer (Genero Report Writer):
 - Graphical report layout designer
 - Deployment Options (Genero display clients):
 - Windows (native), Web, Mac, iPhone, Microsoft Silverlight
 - Web Services
 - Enhancements to the 4GL Language (Genero BDL)
 - Rich O/S platform availability (including Windows and Mac)



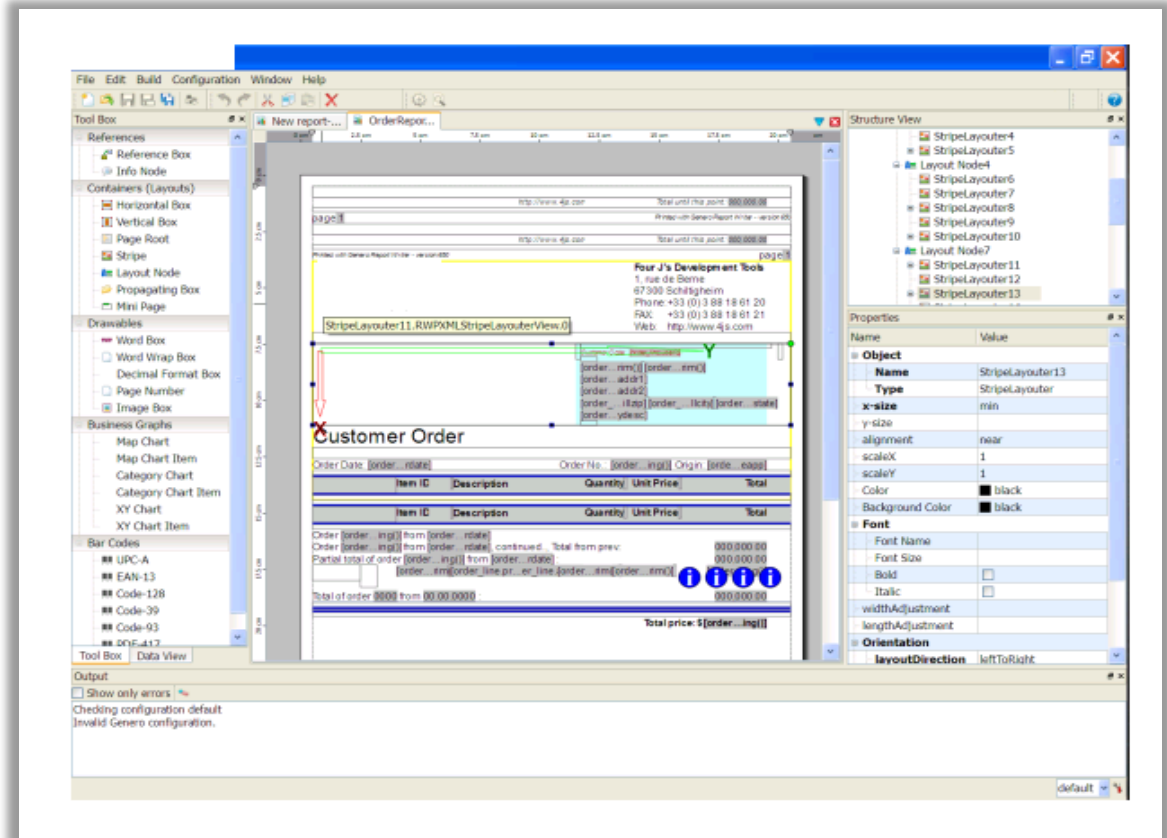
Genero Studio – Graphical IDE for 4GL



- Edit Code
- Create Reports
- Manage Projects
- Diff programs
- Application Generator
- Design Forms
- Debug applications
- Team Work (version control)
- SOA and Web Services
- Code Profiler

Genero Report Writer – Fast, Streaming, and Flexible

- Enterprise reports
- Streaming architecture
- Scalable
- Dynamic layouts
- Formats supported:
 - PDF
 - HTML
 - Excel
 - Others



Flexible Development across many Platforms

One 4GL code base can run across many display environments!



Pocket PC

Mac

Java

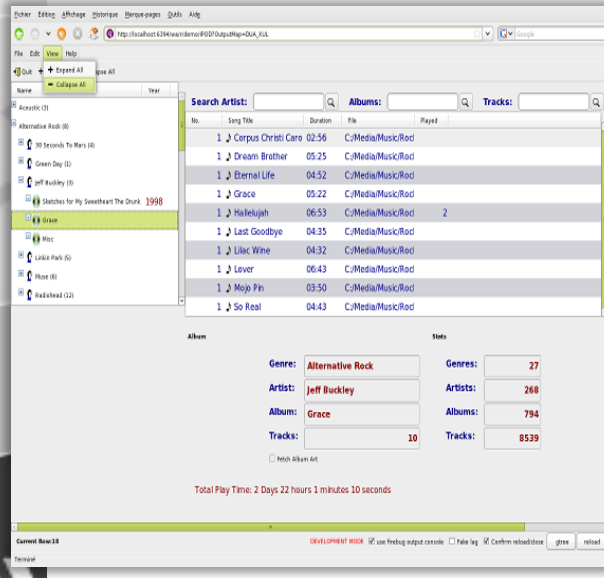
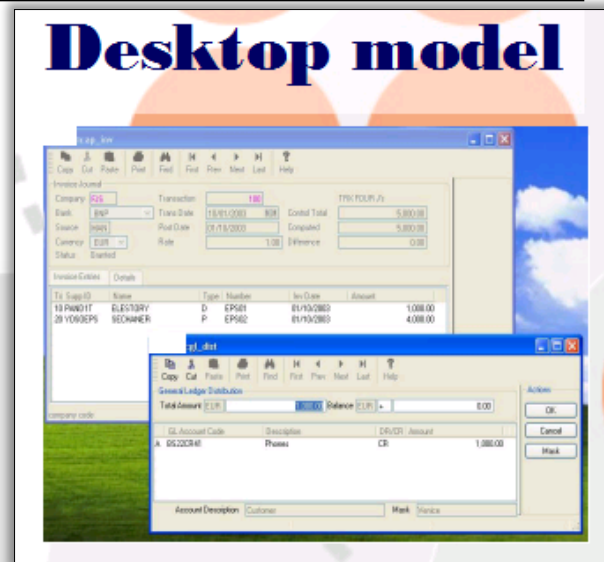
Linux

Windows

HTML

User Interface

XML



4GL to Genero - Conversion Case Studies

What is the process?

1. Re-use existing 4GL source code
2. Simply recompile existing code
3. Perform customization as needed
4. Deploy anywhere!



Project A:

- ◆ Perform 1:1 conversion
- ◆ Very short term - **three weeks to convert**
- ◆ Conversion 100% executed by Adaptris
- ◆ Future development performed by customer

Key metrics:

- ◆ 354 programs
- ◆ 3.023 modules
- ◆ 625 screens
- ◆ 1.174.100 lines of source code

Project B:

- ◆ Conversion with significant enhancements
- ◆ Medium term – **two man months to convert**
- ◆ Conversion executed in tight collaboration at customer site
- ◆ Future development performed by customer

Key metrics:

- ◆ 242 programs
- ◆ 1.522 modules
- ◆ 596 screens
- ◆ 574.684 lines of source code

Questions

Thank
YOU

kbrown3@us.ibm.com