

## Philosophy Behind Panther Release

- Focus on our core strengths:
  - High Availability and Reliability
  - Scale Out
  - Ease of use
  - Simple administration
  - Security
  - Performance
  
- Help our partners and customers WIN



## Panther Themes

- Flexible Grid
- Easy Embeddability
- Expand Warehouse Infrastructure
- Empower Applications Development
- Enhanced Security Management
- Increased Performance



# Flexible Grid



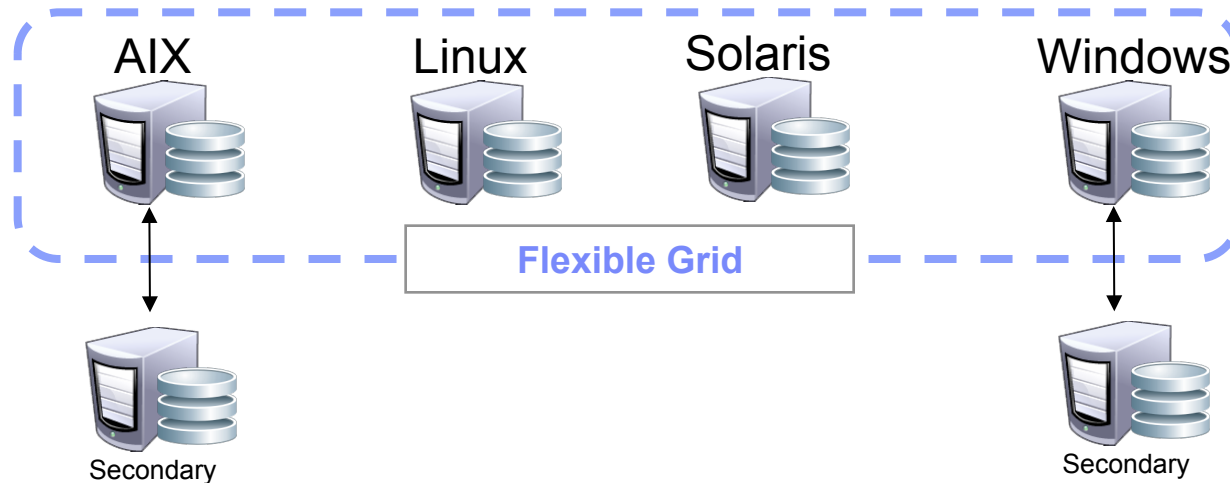
## Flexible Grid – Business Case

- I want to create a grid with a mixture of hardware, software, and Informix versions
- I want to set up my grid quickly and easily
- I want to easily administer all the servers in my grid
- I want my grid to support from 2 to 1000s of servers
- I want to synchronize my schema and data across the grid



## Flexible Grid: Simple Setup

- Install Informix on your server(s)
  - Servers may have secondary servers attached such as HDR, RSS, or SDS servers
- Define a grid to contain your servers:
  - Just give your grid a name and tell it the list of servers to use
  - You can either use the Open Admin Tool graphical interface or a command line tool to define your grid
  - Configure whether you want to replicate just schema changes or schema and data changes
- The grid is ready to use!



**Spread your Workload Across different Hardware, Operating systems, and versions of Informix**

## Flexible Grid: Simple Administration



### ■ Flexible and easy to use


- Administer all the servers in the Grid with SQL statements or the graphical Open Admin Tool
  - Define which servers are allowed to administer the grid
  - Attach to an “admin” server and administer all servers in your grid
  
- Comprehensive set of administration commands
  - Dynamically add and drop servers to/from the grid
  - Performing all DDL operations on all servers in the grid, such as Create Table, Alter Table, ...
  - Interface for targeting administration to all servers in the grid
    - Allows running a query simultaneously on all servers in the grid
    - Output of the query is kept in a local table on each server for later inspection
    - This interface can be used as a general purpose distributed query executor as well:
      - » For instance look at the inventory of each store to see what items need to be reordered
  
- It is still possible to administer individual servers in the Grid directly
  - Through OAT or command line tools



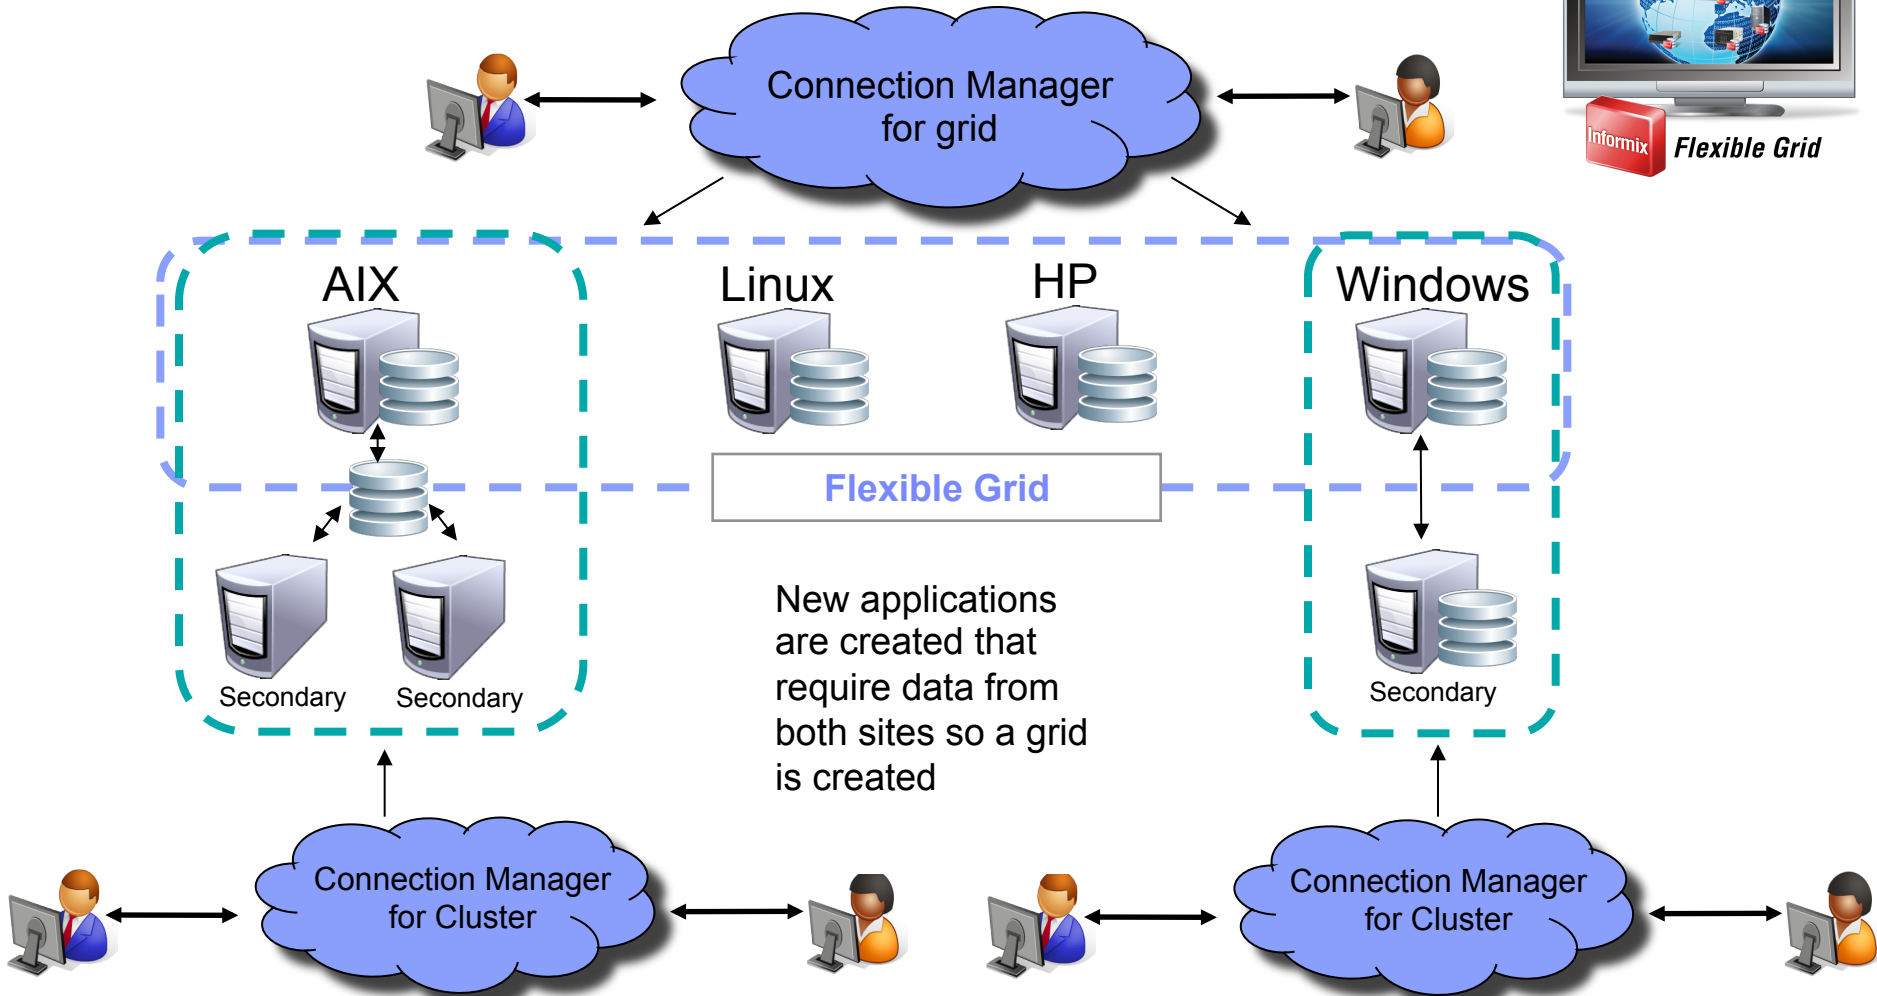
## Flexible Grid: Simple Workload Management



- Setup a connection manager for the grid
  - A simple configuration file determines what servers to manage
  - The connection managers for different sets of servers can be run simultaneously
  - Run multiple connection managers as primary/secondary in case one goes down
- The connection manager for grid distributes application connection requests to the servers in the grid
- Example:
  - Assume a grid with a server in New York and a server in San Francisco
  - The connection manager for grid can be setup for users on both coasts:
    - Use it to connect users in SF to the server in SF and the users in NY to the NY server
  - If a server goes down the connection manager for grid will reroute connections to the remaining server
    - For instance SF users would start connecting to the NY server
  - If a server goes down that has a backup secondary server then connections will be rerouted to the backup server once failover is complete




# Usage Example: Expanding Beyond a Single Cluster





## New Cluster Replication Features



- Rolling Upgrade Support
  - No down time needed to upgrade to 11.70 and beyond
- Support for ER replication of tables without primary keys
- Support DDL changes across ER replicate set and HDR/RSS/SDS servers
- Snapshot clone of an existing server
  - Resulting server can be an RSS server or ER server
- Transactions on secondaries continue to completion when primary fails over
- Backup to the cloud
  - Support for ontape to backup/restore to/from a cloud storage manager
    - Just specify the name of the cloud storage vendor
    - Initial support for Amazon Simple Storage Service



# Embeddability



## Embeddability - Business Value

- Make it even easier for Customers to Embed IDS
- Continue to enhance our core Strengths:
  - Self Maintaining
  - Self Configuration
  - Self Healing
  - Online operation
- Further reduce human & programmatic intervention required with IDS
- Simplify installation





## Embeddability – Panther Features

- Self Maintenance
  - Automated storage allocation
    - Allow existing chunks to be extended automatically
    - Add chunks automatically when space is low
    - Threshold used to determine when to add space
  
- Self Configuring
  - Deployment Assistant / Utility
    - Build a package containing
      - IDS
      - (Optional) pre-built database(s)
      - (Optional) applications
    - Compresses the package
    - Deploy, decompress, and install the package on multiple systems
    - Can be used for media distribution such as CDs
  
- Informix Embeddability toolkit
  - Tutorial for creating an end to end embeddability scenario
  - Example scripts for using Deployment Assistant/Utility



## Embeddability – Panther Features



### ■ Self Healing:

- Standardize errors and alarm codes for application exception handling
  - Out of memory, out of disk, root uninitialized, assertion failure, IDS not running, etc...
- Internally handle errors and retry when possible
  - For example, automatically adding space to a table, automatically registering blades...
- Many Datablades no longer need to be registered by hand
  - The following blades are automatically registered on first use:
    - Basic Text Search, Web Feature Service, Node, Spatial, Timeseries, MQ
- Automated DB Scheduler tasks added
  - Timeout users that have been idle for too long
  - Automatic table storage optimization (compression) based on user settable parameters
  - Automatically allocate CPU VPs to match hardware/licensing when IDS starts



# Warehouse



## Expanded Informix Warehouse – Business Case

- My business is increasingly information driven
  - What capabilities will IDS provide
    - To support this transforming business env?
- How do I manage mixed workloads on a single platform ?
  - Leverage my current investment
- How do I cost-effectively manage my growing warehouse needs?



## Warehouse Infrastructure - Features



- Query optimization
  - Star and Snowflake join optimizations
  - Support for using more than one btree index in a query qualification
  - Optimal memory management for warehouse queries
  
- Improved statistics handling
  - Fragment level statistics are saved
  - Fragment statistics are combined to form global table statistics
  - Will not rebuild statistics for fragments that haven't changed
  - Attach/detach will automatically update statistics when applicable
  
- As tables grow automatically add table fragments based on intervals
  - Intervals can be date intervals (e.g. 1 month) or numeric intervals (e.g. 1-100)
  
- “list” fragmentation
  - Will create a fragment for each item in a list, for example:  
... PARTITION p1 VALUES ('CA', 'OR') IN dbs1 ...





# Informix Ultimate Warehouse Edition



# 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

\*Informix Warehouse Accelerator requires a Linux Intel system as it relies on optimizations in that environment

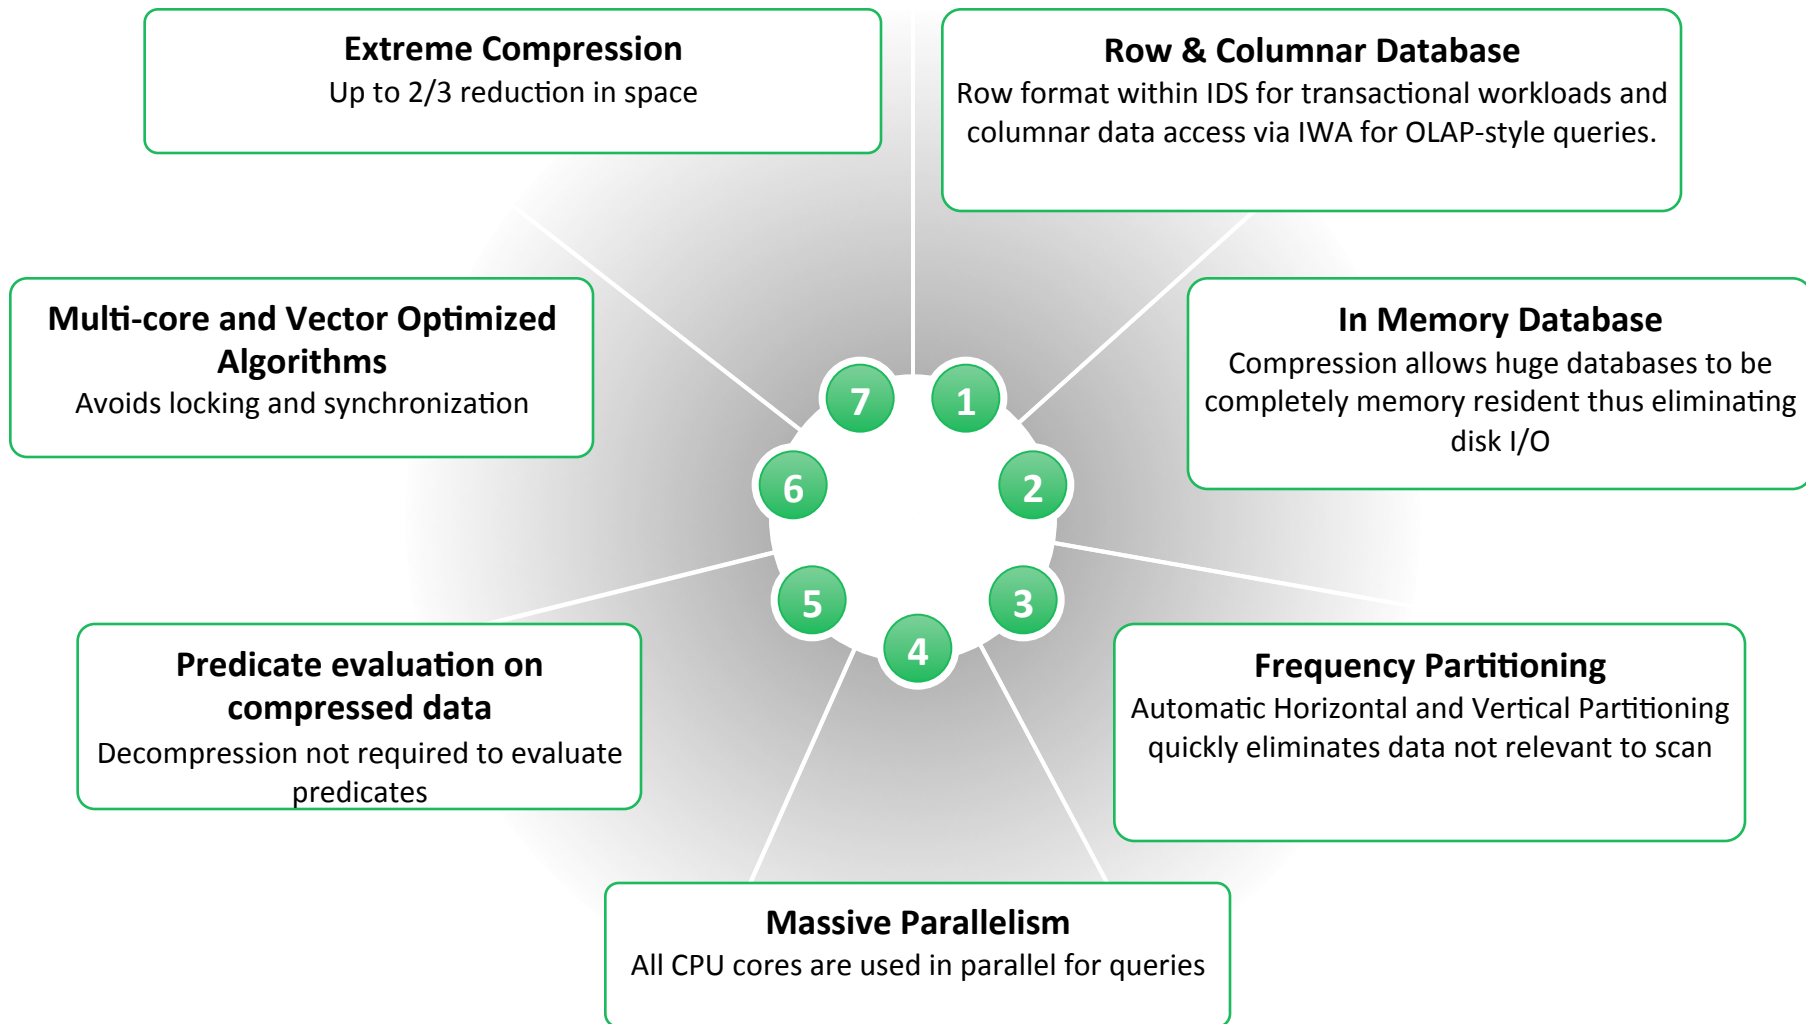
# IWA requires you to do...

- 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

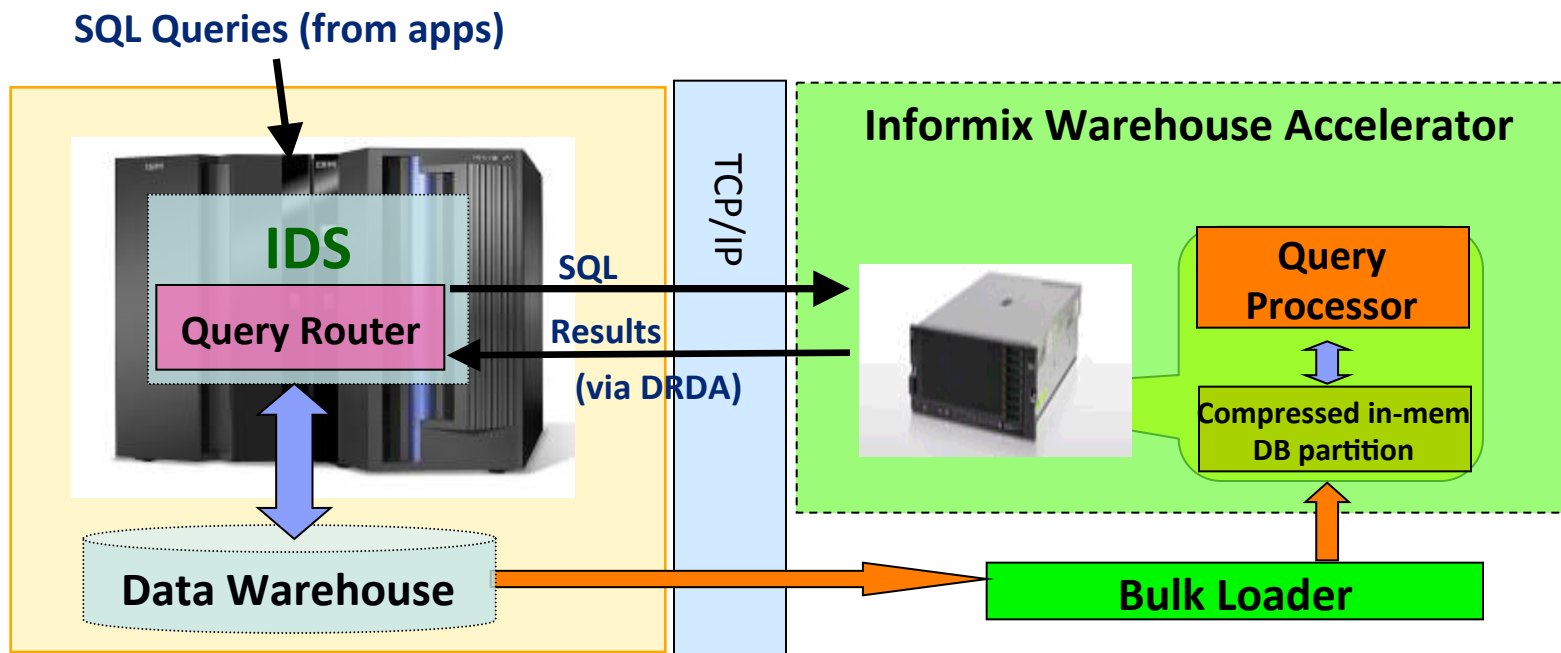
## Power of Simplicity!



# IWA: Breakthrough Technologies for Performance



# Informix Warehouse Accelerator Configuration



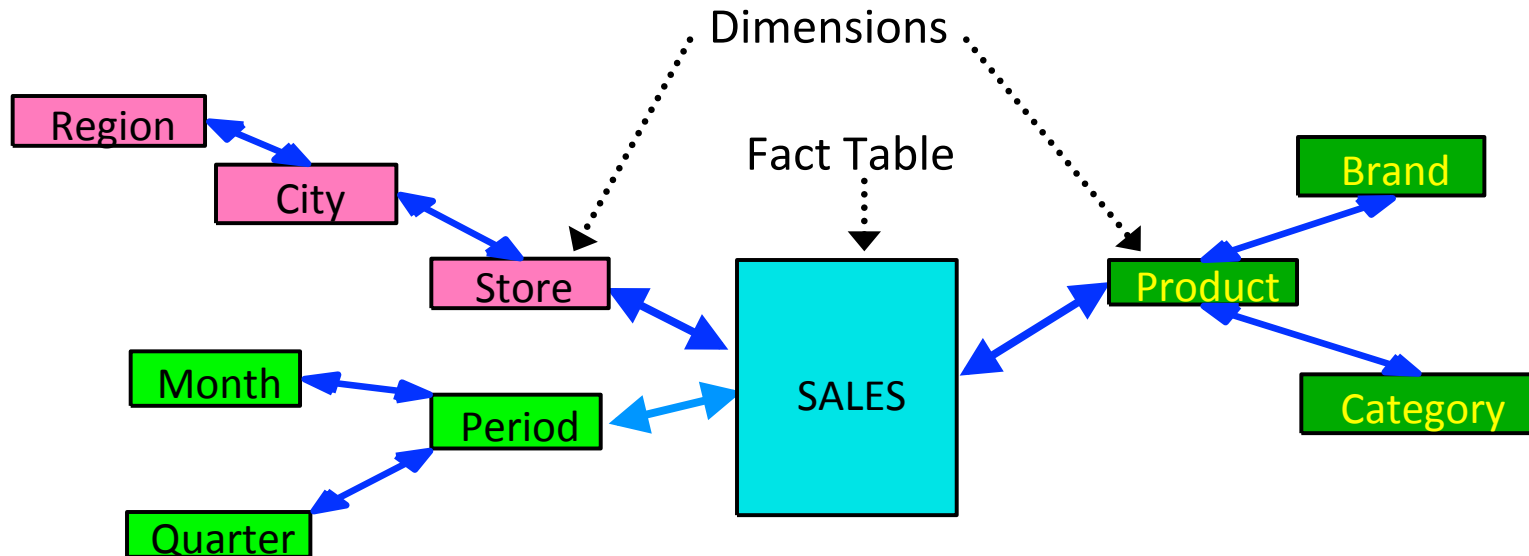
## IDS:

- 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



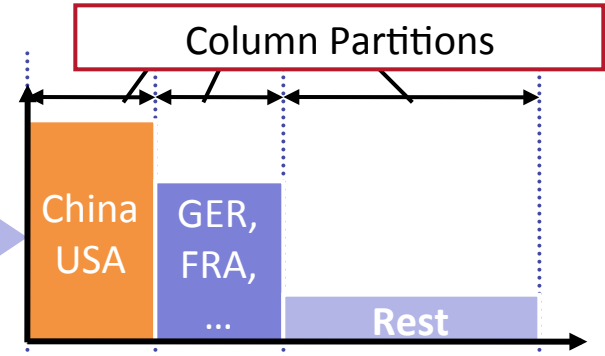




## Complex, ad hoc queries that typically

- 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**



# IBM Smart Analytics Optimizer Studio (ISAO)

The screenshot displays the IBM Smart Analytics Optimizer Studio (ISAO) interface. The main window shows a data mart diagram for 'dm\_ast\_201101' with an estimated size of 0.4 GB. The diagram includes a central table 'td\_ast\_201...' and several smaller tables like 'td\_dwh\_sta...', 'td\_dwh\_ges...', 'td\_dwh\_beru...', 'td\_dwh\_ges...', 'td\_dwh\_all er...', 'td\_dwh\_all e...', 'td\_dwh\_mo...', 'td\_dwh\_arb...', 'td\_dwh\_ber...', and 'td\_dwh\_all e...'. The interface includes a menu bar (File, Edit, Diagram, Navigate, Search, Project, Run, Data Mart, Report, Advice, Window, Help), a toolbar, and a 'Data Project Explorer' on the left. Below the diagram, there are tabs for 'SQL Results', 'Execution Plan', and 'Properties'. The 'Properties' tab is active, showing details for the 'DEMO' accelerator, including 'Name: DEMO (Fully Operational)', 'Shared Memory available: 0.3 GB of 0.4 GB', and 'Trace Profile: DEFAULT'. The bottom status bar indicates 'Cluster status (1 of 1 coordinators and 1 workers active)'.

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	<b>10.7261</b>	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	<b>1.8673</b>	9
5	Query0	Scan - Select First 1000	106m44.474	0m20.946	<b>305.7612</b>	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	<b>51.8780</b>	100
9	Query2	Scan - Select First 100	65m17.075	1m29.518	<b>43.7574</b>	100
10	Query3	Scan - Sum by Dim	87m47.842	0m17.207	<b>306.1453</b>	3107
11	Query4	Scan - Select First 100	53m57.169	0m14.162	<b>228.5813</b>	100
12	Query5	Index - by Dim	0m0.049	0m12.692	0.0039	0
13	Query6	Scan - Select First 1000	1m48.886	0m9.495	<b>11.4677</b>	1000
14	Query7	Scan - by Dim w Group by	0m15.855	0m9.171	<b>1.7288</b>	26937
15	Query8	Scan - by Dim w Group by	0m14.846	0m2.347	<b>6.3255</b>	28358
16	Query9	Scan - by Dim w Group by	79m14.253	0m42.011	<b>113.1669</b>	523499
<b>Total in hours</b>			<b>9hr 40min</b>	<b>15min 32sec</b>		
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



# Application Development



## Enhanced Application Development – Business Needs

- How do I reduce the cost of App Dev?
  - Using standard applications and tools
  - Find problems as quickly as possible
- How do I maximize my productivity?
  - Reduced time and effort for App Dev
  - Better , broader integration with IBM stack
- How can I transform my applications ?
  - Serve modern business needs





## Enhanced Application Development - Features

- Stored Procedure Debugging
  - Integration with the Optim Data Studio procedure debugger
  - Integration with Microsoft Visual Studio debugger (post panther)
  
- Improved compatibility with open source applications:
  - Drupal, Hibernate, Geronimo, iBATIS, Mediawiki, Tomcat, Xwiki
  - Available on the IIUG web site
  
- New SQL syntax added for compatibility
  - As open source applications are ported syntax changes are made to accommodate those programs:
    - NULL clause now supported
    - Relax the position of the default clause
    - Relax the position of ON DELETE CASCADE
    - Support for “if [not] exists”
    - Expressions allowed as arguments to “count” aggregate



## Enhanced Application Development - Features



- Table defragmentation (online table reorg)
  - Coalesce contiguous extents to reduce number of extents
- Provide support for MQ series server on a remote box
  - Multiple instances of IDS to attach to the same MQ server
  - Expose the messages in the queue as an IDS table
- New install framework
  - Consistent UI across all platforms
- More languages supported by OAT
- Optim Performance Manager support for IDS (post panther)
- Mashup Hub bundled with Informix



# Informix + IBM Mashup Center– UI modernization

```

CUSTOMER ORDERS
CUSTOMER
Customer Number:[ 104]      Company Name:[Play Ball! ]

ORDER NUMBER   ORDER DATE   PAID DATE   INSTRUCTIONS
1001           5/20/2008   7/22/2008   express
1003           4/22/2008   6/14/2008   express
1011           6/18/2008   8/29/2008   express
1013           6/22/2008   7/31/2008   express

Do you want to see another report? (n/y): █
    
```

## Enhancing Existing and New Reports

- Reports no longer solely tethered
- Integrating with other reports
- Expanding on 24x80 pixel screen
- Different Visualization: Map, Chart, Tiimeline
- Modern look-up techniques (AutoComplete)
- Accommodate Digital Generation

Leverage Your SQL!



Mobile

Portals

External  
3rd Party Sites

*“With the mashup capability now available within Informix, we will be able to increase customer satisfaction in line with relevant product messages akin to “AdSense” techniques while reducing support costs – all at a fraction of the cost of traditional development.. ~ John Baker, Director, Oninit Consulting Ltd*





# TimeSeries



# Example of Changing Storage Requirements

## Changing Workloads For 10 Million Smart Meters:

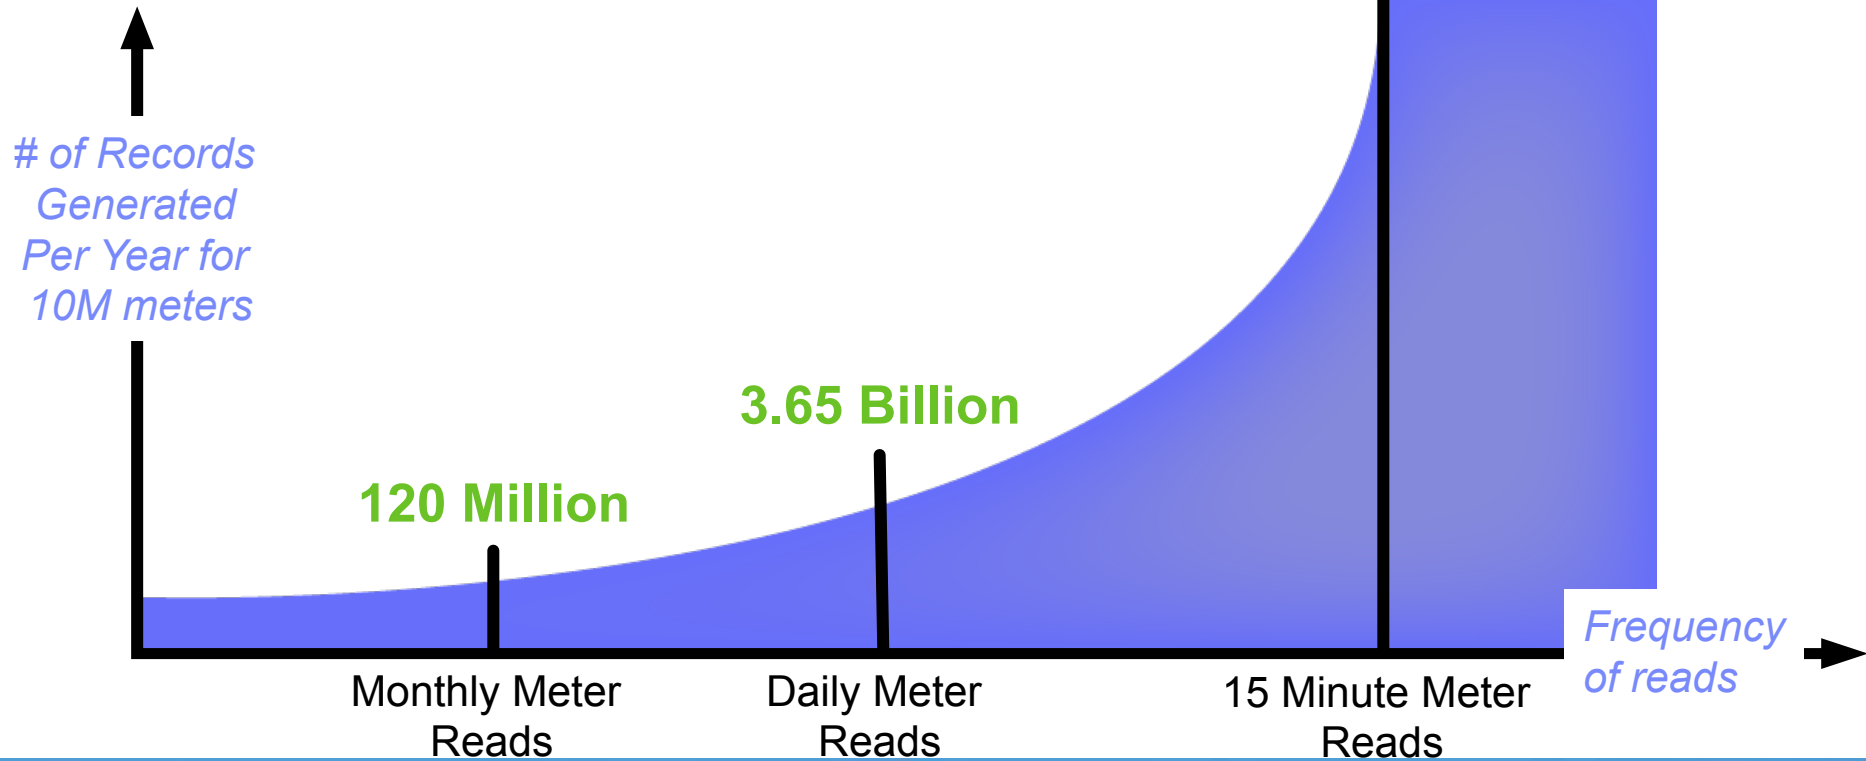
**Today** – Each meter is read once per month

**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

Data for a Utility In California

350.4 Billion



## 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




## 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

Meter_Id	Time	KWH	Voltage	.....	CoIN
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
...	...	...	...	.....	...

Table Grows

- 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



Meter_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*

## 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**
- **No problem writing the data at this rate!**

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





# Consumer Education



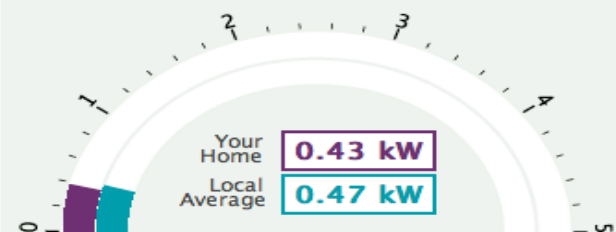


An EU funded climate change initiative

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

### Current live energy usage



### Energy Saving

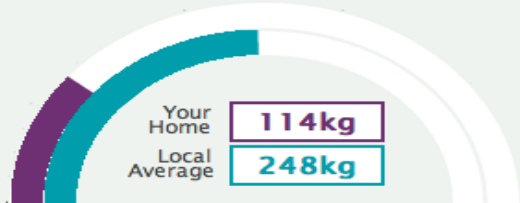
Compared to your last bill we currently estimate the difference in your spending is:

**+£7.96**

per month

### Usage so far this month

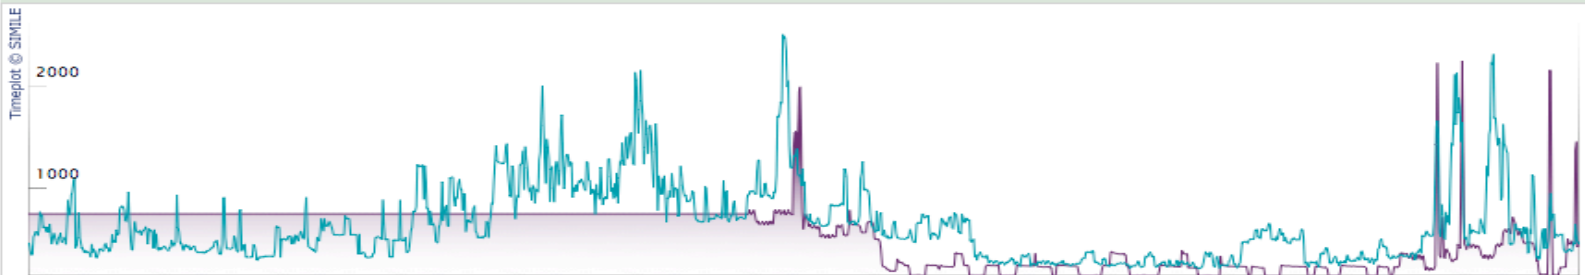
kWh
Cost
CO<sub>2</sub>



### 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

# 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/Sybox 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



# Security



## Enhanced Security – Business Needs

- How can I track data access for audit purposes?
  - Define and collect only the data I need
  - Track web users
- How can I simplify user administration?
  - Support central admin
  - Support admin of Non-OS users
- How can reduce time and effort to compliance ?



## Enhanced Security – Features

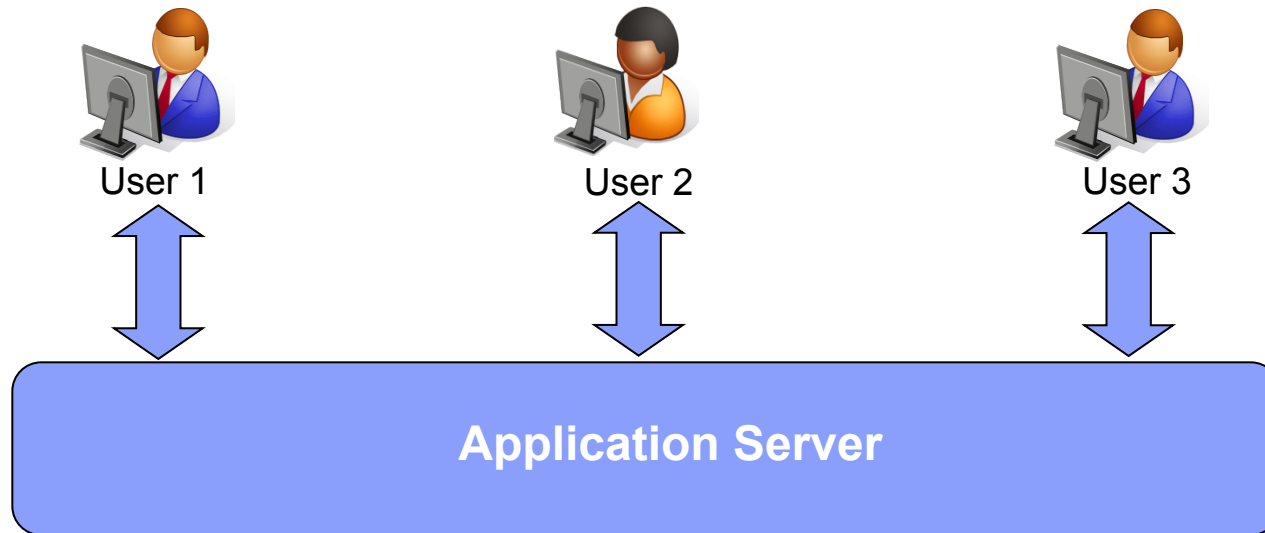


- Selective row-level (SRL) Auditing
  - Allow customers to pick which tables and operations to audit
  
- Trusted Context
  - Trust user authentication done by middle tier applications
  - Allow these middle tier applications to switch users over an existing connection without shutting down the connection
  
- Support non-OS users
  - IDS users will no longer need to have a login on the host OS
  
- Encryption of raw disks via Encryption Expert
  - Encryption for regular file systems already supported
  - Vormetrix currently supports only HP, but more available shortly
  - Most Informix customers use raw disks






# Trusted Context



- Benefit:

- No need to drop the connection and reconnect with each new user in order to maintain security
- Security is maintained without a hit in performance

- Informix trusts the Application Server to authenticate the user
- It then sends a message to Informix when a new user sends a query
- No need to drop and reconnect with each user
- Each user maintains their own identity



# Performance



## OLTP Performance - Features




- Forest of Trees (FOT)
  - New index composed of many subtrees (buckets)
- C-UDR preload feature
  - Improve C user defined function performance



# Forest of Trees Index

Hash on key value to pick a bucket / rootnode



Key Value	Bucket
1	2
2	1
3	2
...	...
47	3
221	1

## Syntax Example:

create unique *my\_index* on *my\_table* (product\_code, sold\_date, item\_number) hash on (product\_code) with 10000 buckets;

- Hash on product\_code
- Range on sold\_date
- select \* from my\_table where product\_code = ? and sold\_date between ? and ?;



# Questions

(Kevin Brown – [kbrown3@us.ibm.com](mailto:kbrown3@us.ibm.com))



Thank  
YOU

