Information Management

IBM

Select Information from IBM where product = "*Informix*" *and version =11.70*

Sandor Szabo – Informix Development

January 30, 2011

Informix Technical Roadmap

2005

Optimized OLTP Engine

Query performance improvement

- Online index build
- Multiple page size support for better space utilization
- Enhanced buffer management
- Increased security with column encryption
- Disaster recovery with table level restore

2007 - 2008

Scale-out at lower costs

- Multi-node active cluster for high-availability (MACH) 11 with multiple remote servers and shared disk cluster
- Open source tool for administration, SQL Admin API to embed admin tasks
- Improved checkpoint performance
- Secured data encryption, LBAC, Common Criteria certification
- Enhanced application development for SOA and XML
- Data Server Client
- Text Search

2009

Business Optimization

- Informix Warehouse
- Cognos Integration
- Virtual Appliance
- Cloud Computing support
- Online Storage Optimization
- XPS features
- External Tables
- In memory support with SolidDB
- Heterogeneous Replication via CDC
- 3-D Internet collaboration
- Delay/Stop Apply Cluster
- 4GL Enhancements

2010

Application Integration

- Grid Computing
- SOA in the Database
- Enhanced Warehouse capabilities
  - **Performance/Index advisor**
  - **More XPS Features**
- Deeper Embed
  - **Automatic storage provisioning**
  - **Automatic Fragmentation**
  - **Embedability toolkit**
  - **Installation API**
- Oat Enhancements
- Warehouse
- Security
  - **Fine-grained Auditing**
  - **Trusted contex**
- 4GL Enhancements

4 Imperatives will drive our Informix strategy

- *Delight the install base: Enhance the Informix client relationship and improve client loyalty*
- *Build Strategic Differentiation: Fix the messaging and enhance the overall Informix value proposition to the market*
- *Create Proactive Sales Culture: Develop a proactive, high-touch culture, with expansion into growth areas (embedded, etc).*
- *Build New, Profitable Revenue Streams: Build new revenue streams that capitalize on the strong and growing footprint of Informix*

Philosophy Behind 11.70

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

- **Help our partners and customers WIN**

Informix 11.70 Themes

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

Information Management

IBM

Flexible Grid

January 30, 2011

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!**

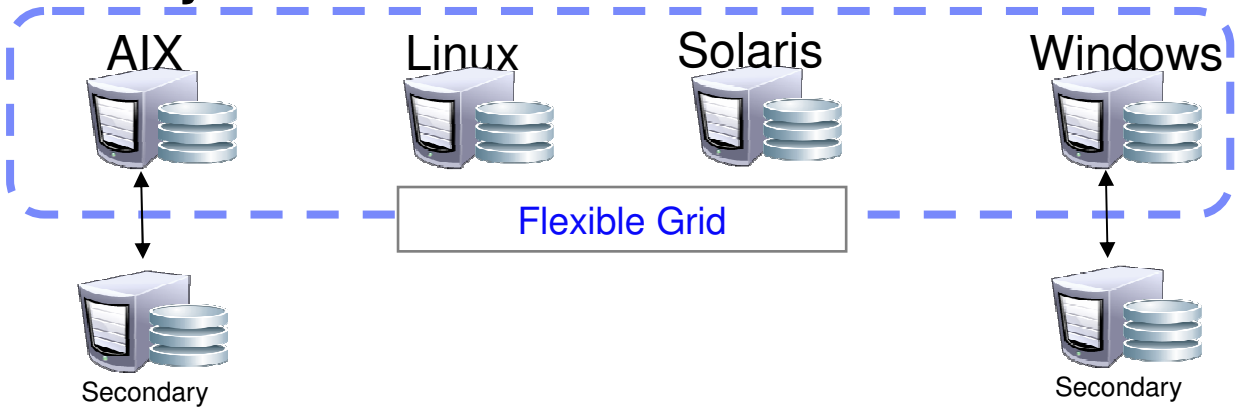

Spread your Workload Across different Hardware, Operating systems, and versions of Informix

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 (OAT):
    - 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:
  - 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.

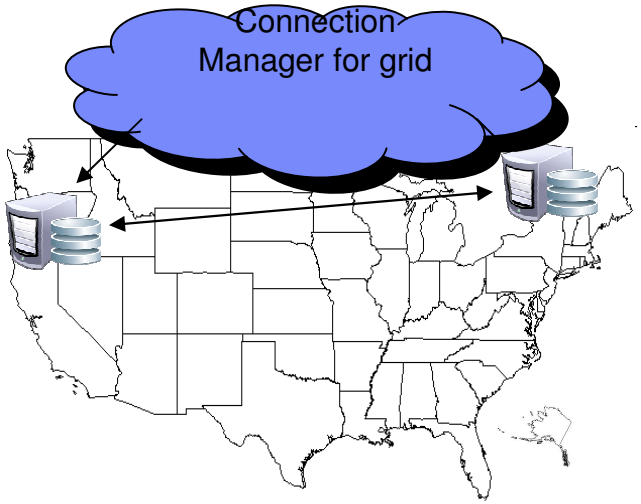

Informix Flexible Grid - Introduction

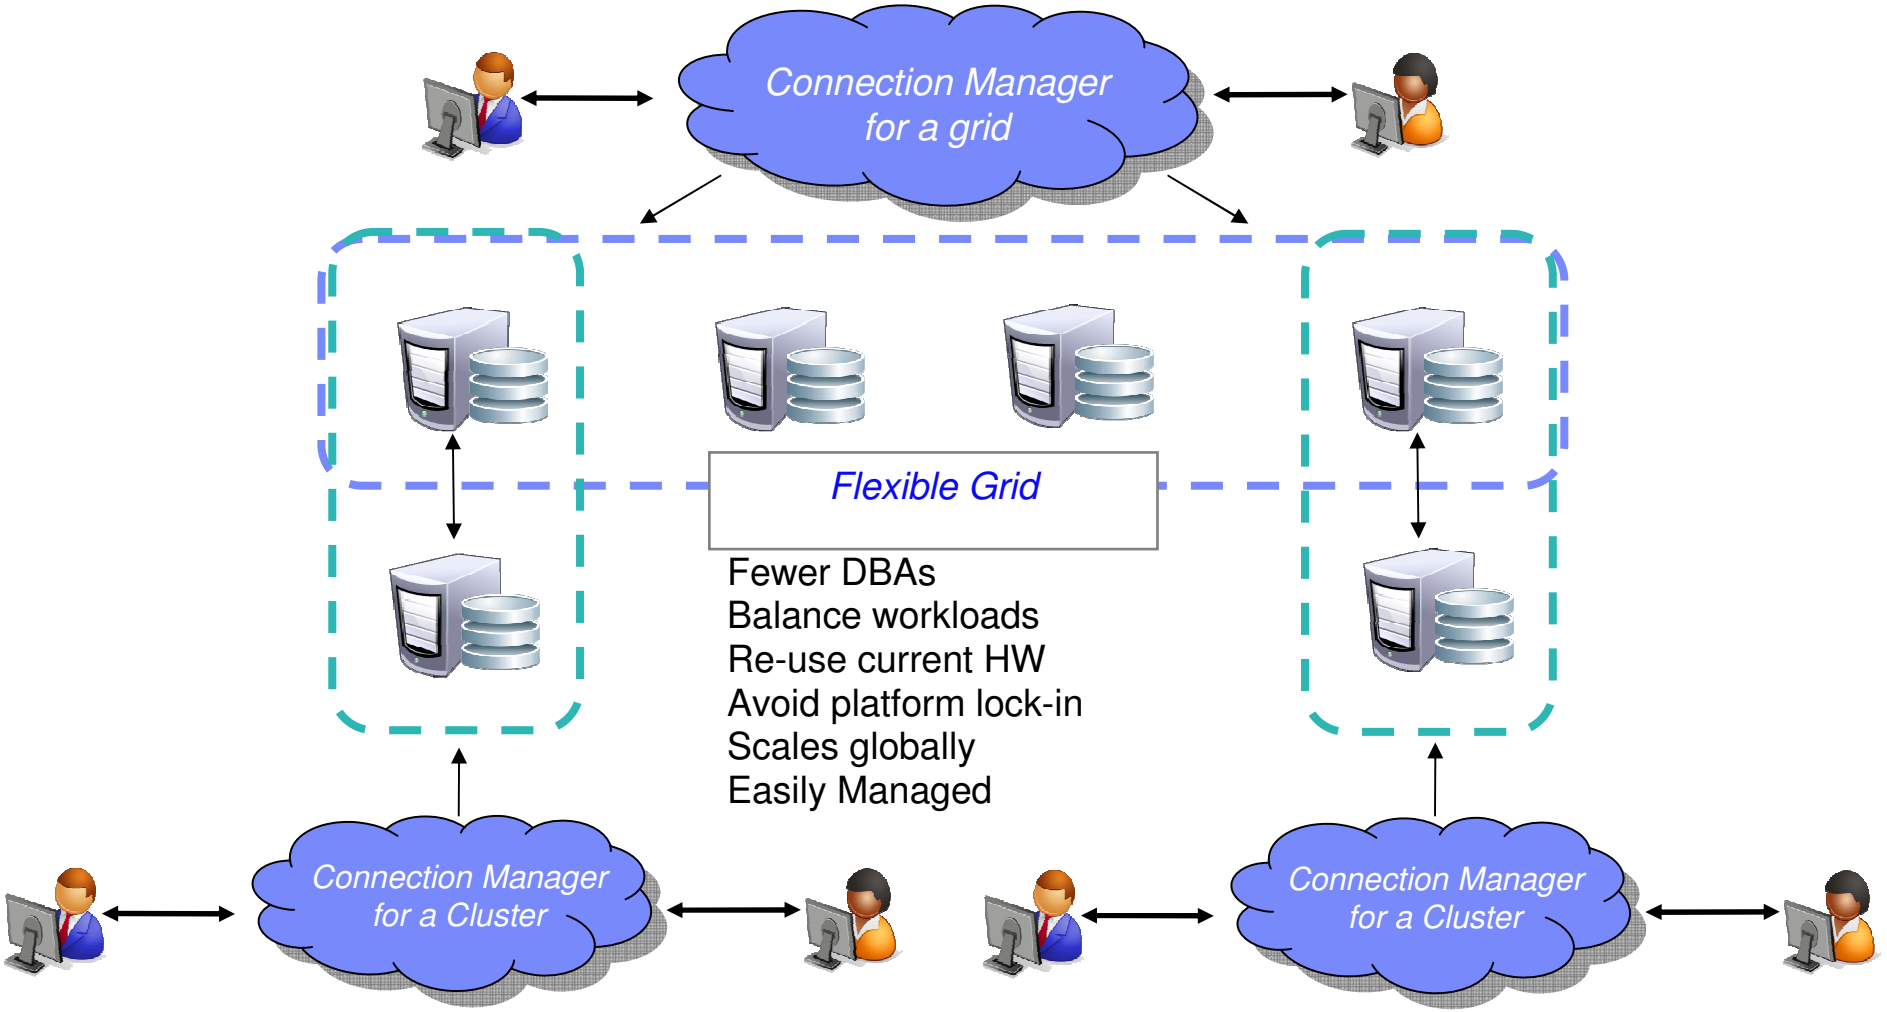

New Cluster Replication Features

- **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 (S3).

Information Management

IBM

Embeddability

January 30, 2011

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

- **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 CD's.
- **Informix Embeddability toolkit:**
  - Tutorial for creating an end to end embeddability scenario.
  - Example scripts for using Deployment Assistant/Utility.

Embeddability

- 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.

Storage Provisioning

- **Storage provisioning solves a long time problem for most IDS customers; namely, how do we dynamically and efficiently add additional capacity for IDS storage spaces, not knowing in advance when or which one of the spaces will either run out of space or cross the threshold of a space allocation requirement?**

- **Create a pool of storage, unallocated, within IDS:**
  - Configure a little in advance.

- **Optional automatic expansion of** dbspaces, temporary dbspaces, sbspaces, temporary sbspaces, **and** blobspaces:
  - "Out-of-space" errors are significantly reduced.
  - Chunks are extended and/or created as needed.
  - IDS can also expand spaces *before* they are full.

Storage Provisioning: The Power of 2

- Two available modes:
  - Manual
  - Automatic
    - mon_low_storage in sysadmin:ph_task enabled by default.
- Two available space expansion methods:
  - Chunk extension
  - Chunk creation
- Two available interfaces:
  - sysadmin task() / admin() functions (SQL interface)
  - OAT (Graphical interface)

What is the Storage Pool?

- **A list of raw devices, cooked files, and/or directories from which IDS can allocate space as needed.**

- **For example:**

  Path: /dev/rawdevice1
  Offset: 500000
  Total Size: 2000000

  Path: /ifmx/STORAGE_DIR
  Offset: 0
  Total Size: 0

- storagepool **table in the** sysadmin **database** (sysadmin:storagepool)

Extending a chunk

- **Chunks are *not* extendable by default. To make a chunk extendable:**

```
EXECUTE FUNCTION task("modify chunk extendable",
      "<chunk_num>");
```

- **Now, you can extend it manually at any time:**

```
EXECUTE FUNCTION task("modify chunk extend",
      "<chunk_num>",
      "<min_size>");
```

- **You can also wait for the server to extend it automatically.**
- Restriction: **Only unmirrored dbspace and temporary dbspace chunks can be extended.**

New SQL Administration API admin() & task() Commands (1)

- storagepool add
- storagepool modify
- storagepool delete
- storagepool purge
- modify space sp_sizes
- modify space expand
- modify chunk extendable [off]
- modify chunk extend

New SQL Administration API admin() & task() Commands (2)

- create <object> from storagepool
- drop <object> to storagepool
- storagepool add
- modify space expand
- create <space_type> from storagepool
- drop <space_type> to storagepool

Storage Provisioning Configuration

- **Instance level**
  - SP_AUTOEXPAND
  - SP_THRESHOLD
  - SP_WAITTIME
- **dbspace level**
  - Create size
  - Extend size
- **Storage pool entry level**
  - Chunk size
  - Priority

Information Management

IBM

Warehouse

January 30, 2011

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 ?**

OLTP vs. Data Warehousing Workload

- Short Transactions
  - Relatively simple SQL
- Random Updates
  - Few Rows accessed
- Sub-second response time
- ER Modeling
  - Minimizes redundancy
- Normalized data (5NF)
  - Minimizes duplicates
- Few indexes
  - Avoids index maintenance
- Pre-compiled queries
  - Repeated execution of queries

- Longer Transactions
  - **Complex SQL with analytics**
- Sequential Updates
  - **Many Rows Accessed**
- Secs to Mins response time
- Dimensional Modeling
  - **OK to have redundancy**
- De-normalized data (3NF)
  - **Duplicates are OK**
- OK to have more indexes
  - **Mostly read only**
- Ad-hoc queries
  - **Unpredictable load**

Informix 11.70 Warehousing Features

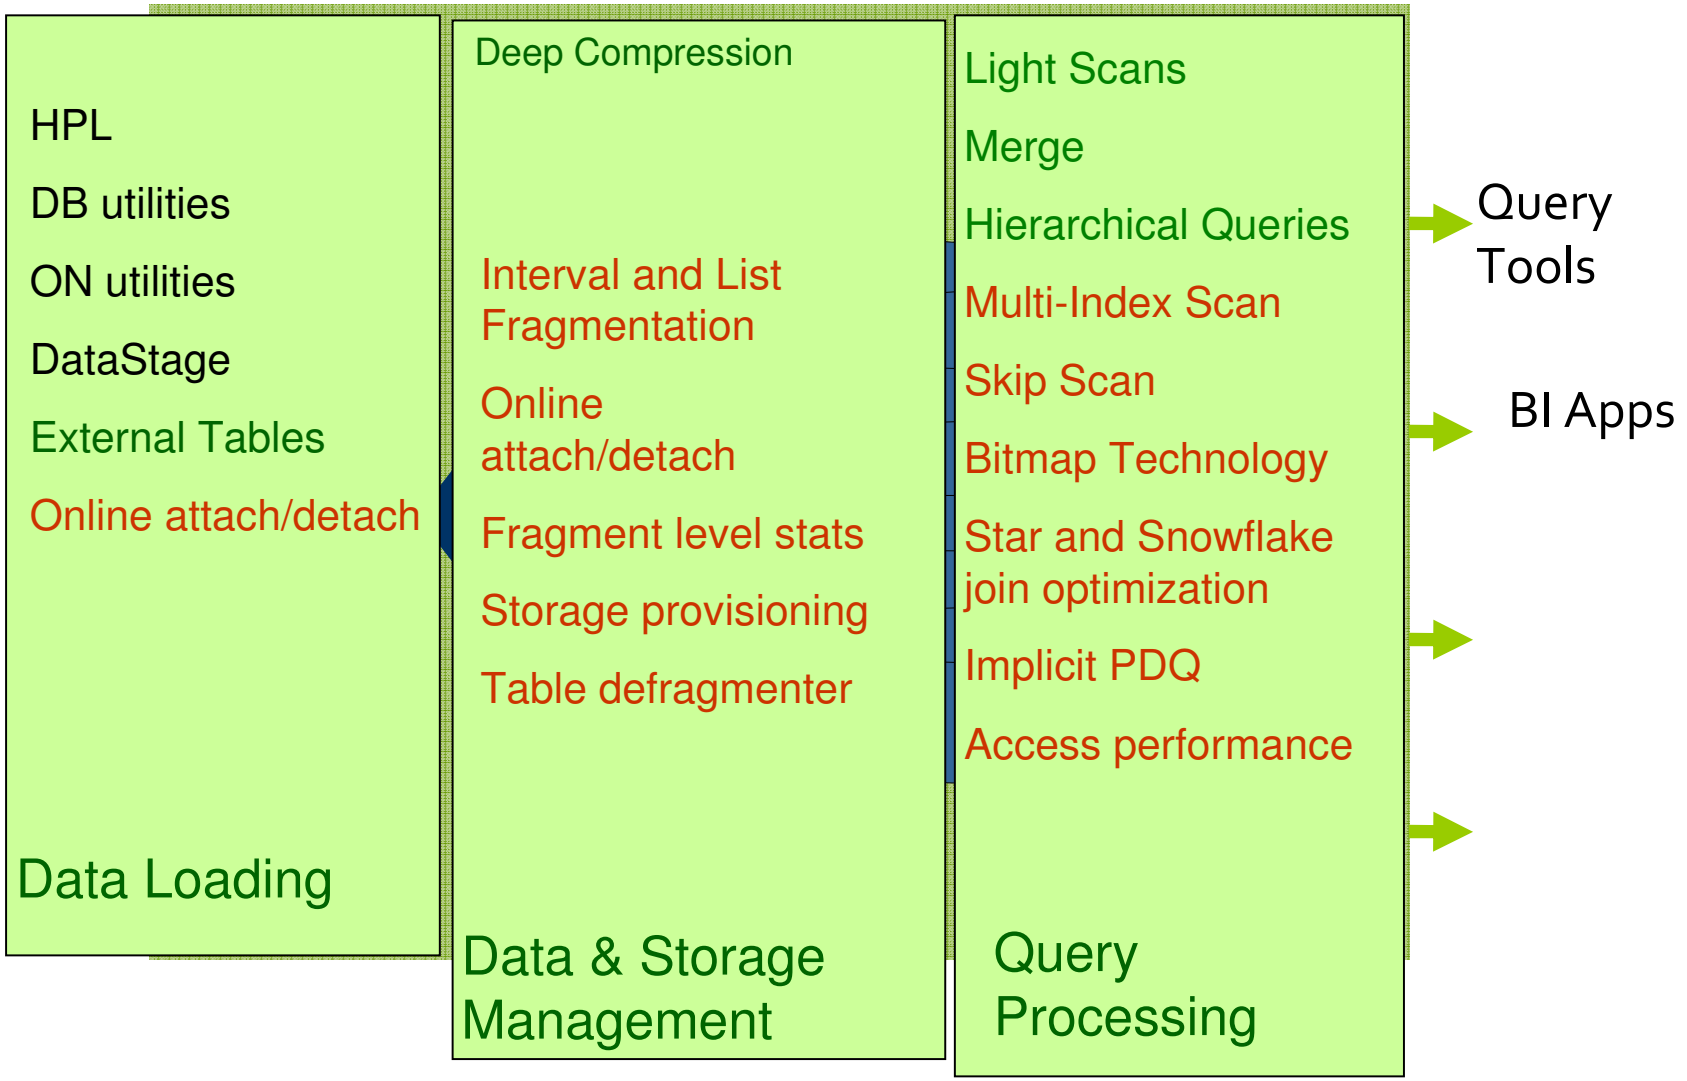

Interval Fragmentation

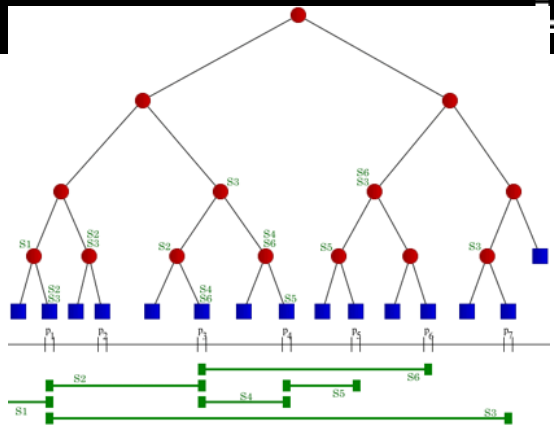

- Fragments data based on an interval value
  - E.g. fragment for every month or every million customer records
- Tables have an initial set of fragments defined by a range expression
- When a row is inserted that does not fit in the initial range fragments,
- IDS will automatically create fragment to hold the row (no DBA intervention)
- No Exclusive-lock is required for fragment addition

Fragment ONLINE operations

- ATTACH a fragment ONLINE
- DETACH a fragment ONLINE
- MODIFY transition value ONLINE
- Automatic ADDing of new fragments on insert or update
- Automatic statistics collection after attach and detach
- DBA tasks eliminated:
  - Scheduling downtime to get exclusive access for ADD, ATTACH, DETACH
  - Defining proper expressions to ensure fragment elimination
  - Running of update statistics manually after ALTER operations
  - Time taken to collect statistics is reduced as well.

Fragment Level Statistics (FLS)

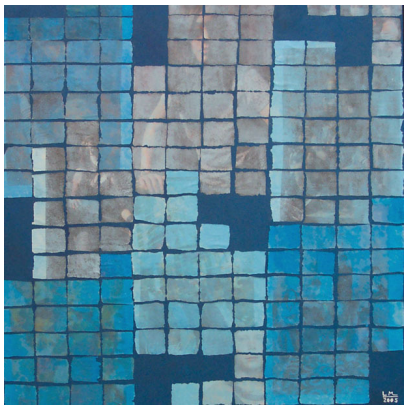

- Generate and store column distribution at fragment level
- Fragment level stats are combined to form column distribution
- System monitors UDI (Update/Delete/Insert) activities on each fragment
- Stats are refreshed only for frequently updated fragments
- Fragment level distribution is used to re-calculate column distribution
- No need to re-generate stats across entire table

Star and snowflake schema

fact table + dimension tables

typically have primary/foreign constraints (explicit or implicit)

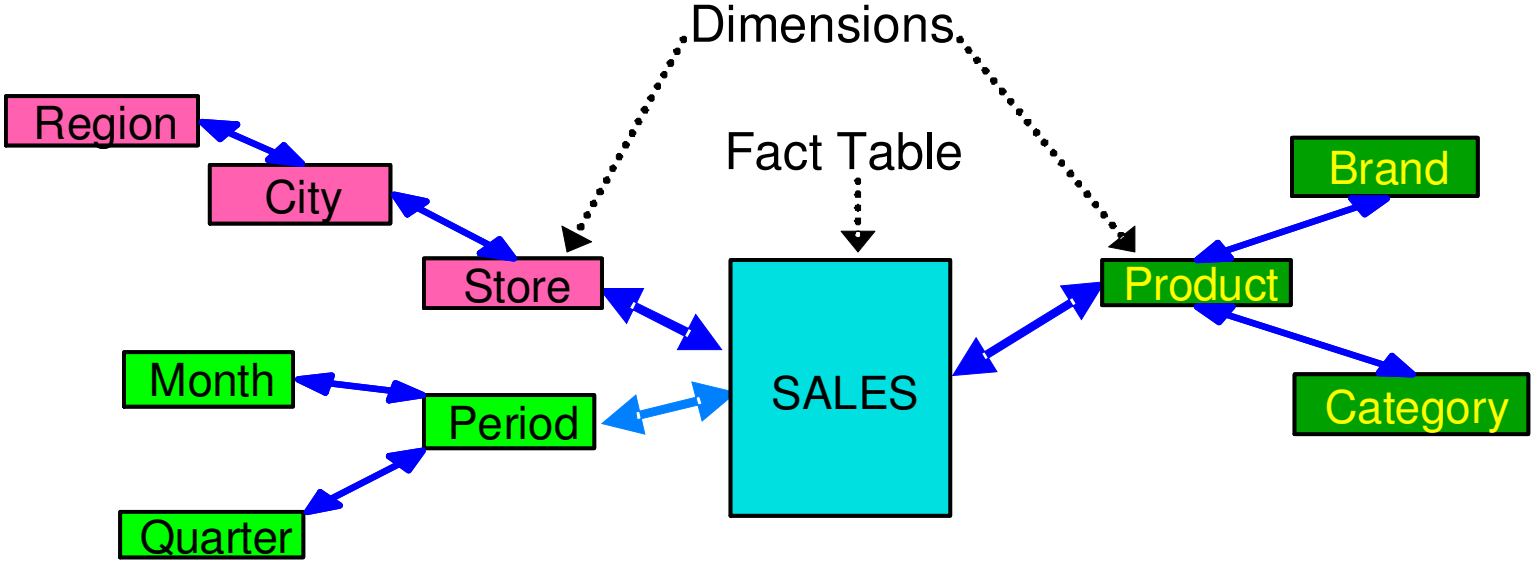

- Large Central “Fact” table, Smaller “Dimension” tables

Small fraction of Fact table in result

Star-Join - Requirements

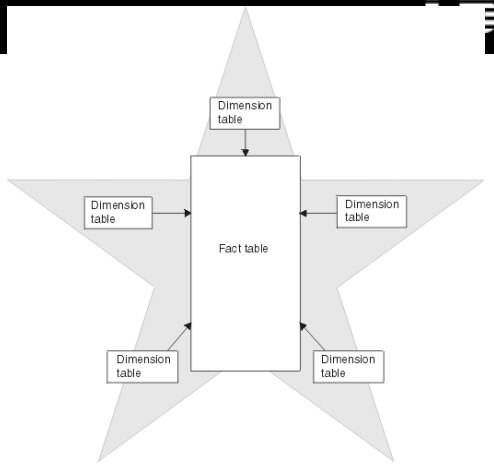

Query qualifies for a Star-Join plan if:

- Has Update Stats low on each table
- PDQ is ON
- Estimated number of rows from the identified fact table has to be greater than all other tables.
- Fact table itself is identified with equijoins to all other dimensions.
  - Same method repeated to identify branches from each dimension resulting in snowflake identification).

Optimization environment

- New syntax for SET OPTIMIZATION statement:
- set optimization environment star_join 'enabled/disabled/forced';
  - 'enabled' – default, star join optimization is considered
  - 'disabled' – star join optimization is not considered
  - 'forced' – star join execution plan is chosen if available
- set optimization environment fact 'tab1,tab2,...tabN';
  - only consider tables listed as fact table
- set optimization environment avoid_fact 'tab1,tab2,..tabN';
  - does not consider tables listed as fact table
- set optimization environment non_dim 'tab1,tab2,...tabN';
  - does not consider tables listed as dim

Can be set in sysdbopen

Information Management

IBM

Application Development

January 30, 2011

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.

Informix Timeseries Included for Free

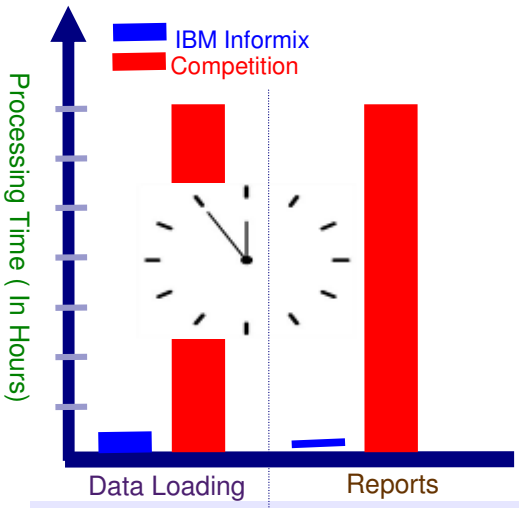

Performance Comparison for Data Loads and Reports for 1 Million Meters

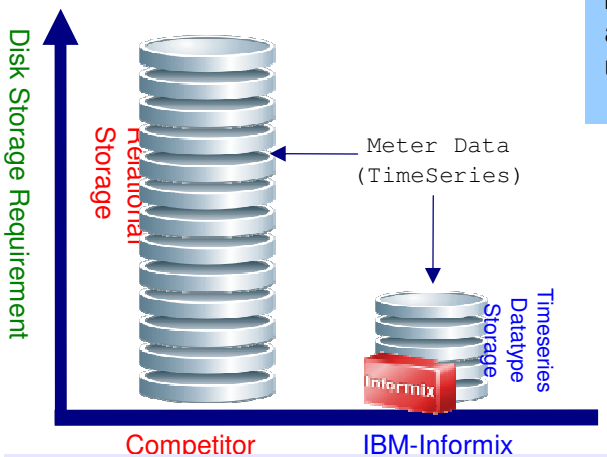

Data Storage Comparison for 1 million meters

Based on actual benchmark test run for a large US electrical utility company

- **Informix took ~ 18 minutes to load 1 day of data for 1 million meters;** *The competitor took ~ 7 hours.*
- **Informix took from 25 seconds to 6 minutes to run the utility commission reports;** *The competitor took from 2 - 7 hours depending on the report.*
- **Data space used by Informix is ~ 350GB;** *The competitor uses about 1.3TB.*
- **Informix performance and storage comparison is linear with more meter data.**
- **Informix gives even better results if you increase cpu's and storage.**

Information Management

IBM

Security

January 30, 2011

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 for 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.
  - *Vormetric* currently supports only *HP*, but more available shortly.
  - Most Informix customers use raw disks.

Information Management

IBM

Performance

January 30, 2011

Performance and Resource Agenda

- Implicit PDQ
- Pre Load a C-UDR Library
- Network and Connection Performance
- Linux Large Page Support
- Prevent Accidental Reinitialization.

Implicit PDQ

- Implicit Parallel Data Query (PDQ) support enables IDS server to use resources based on query's need.
- Smarter utilization of resources.
- Explicit PDQ:
  - User setting (SET PDQPRIORITY statement).
  - All queries in current session use same setting.
- Implicit PDQ:
  - IDS determines resource requirement based on optimizer's estimates.
  - Each query can have a different PDQ setting.
- Requires at least low-level statistics and distribution statistics on all tables involved in a query if IMPLICIT_PDQ is turned on for IMPLICIT_PDQ to work.
- Does not work with external tables at this time.
- Can use in sysdbopen()

Pre-load a C-UDR Shared Library

- Allows IDS threads to migrate from one cpuvp to another during the execution of the C-UDR.
- Without this feature, the thread executing the UDR is bound to the cpuvp for the duration of the C-UDR execution.
- Allowing thread migration during C-UDR execution can increase performance.
- PRELOAD_DLL_FILE **-** onconfig **parameter:**
  - This parameter is repeatable, once for each shared library.
  - Directory location of the library is its value.
- Sample:

```
PRELOAD_DLL_FILE   $INFORMIXDIR/extend/test.udr
PRELOAD_DLL_FILE   /work4/jmcmahon/foo.so
```

Connectivity Performance Enhancements

- There is improved connection time performance by looking for O/S connectivity data in server memory instead of server disk.

- New onconfig parameter NS_CACHE :

  - Used to define an expiration time measured in seconds for the database server to search for and then get host, service, user, and group information from the O/S Name Service Cache.

  - Search the memory cache first for the data before using the operating system:

    - Its much quicker than disk ……
    - Example:

    - NS_CACHE host=900, service=900, user=900, group=900 (default)

Connectivity Performance Enhancements

- The onconfig parameters DBSERVERNAME and DBSERVERALIASES have been enhanced to allocate multiple listener threads to an instance, resulting in improved connectivity performance as well.

- Syntax:
  - DBSERVERNAME name[-n]
  - DBSERVERALIASES alias[-n],...
    - where ‘n’ is the number of threads to allocate.

- Example:
  - DBSERVERNAME ids-8

- onstat –g ath | grep “lst” **to monitor.**

- Both the listener thread and DNS memory search together result in about 90% connectivity performance improvement.

Linux Large Page Support

- Large Page Support for Linux - Large pages for non-message shared memory segments that reside in physical memory are now enabled by default on Linux platforms and previously, on AIX and Solaris systems:
  - On Linux x86_64 the large pages used by Informix are defined by the
  - Hugepagesize entry in the /proc/meminfo file.
- The use of large pages can provide huge performance benefits in large memory configurations. To enable or disable support for large pages, use the IFX_LARGE_PAGES environment variable.
- The system administrator or DBSA must activate large pages in the operating system.
- setenv IFX_LARGE_PAGES 1 – To activate (0 is default).

Preventing Accidental Instance Re-initialization

- You seen this slide before …… where we warn you not to execute oninit –i when bringing up an existing instance, because of the total instance destruction that will occur to the existing instance.
- Wouldn't it be nice if you could prevent that possibility……
- New onconfig parameter: FULL_DISK_INIT – two possible values:
  - Default value is 0 – oninit –i will successfully execute only if there is not a page-0 in the root path location (the instance has not been previously initialized and does not exist).
  - Value of 1 – oninit –i will successfully execute in all circumstances. Once the instance is successfully initialized the first time, this value is then changed to 0, to prevent accidental initialization of the new instance.
- A message is written to online.log reporting unsuccessful re-initialization attempts.

Forest of Trees Indexes

- A larger B-Tree index divided into smaller subtrees called “buckets”.
- Programmer defines which columns are used to hash to a “bucket” and the number of “buckets”.
- Each “bucket” acts as sort of a mini index.
- Traditional Btree Index:
  - create unique index security_idx on
    SECURITY(S_SYMB, S_CO_ID) in dbs;
- Forest of Trees Index:
  - create unique index security_idx on
    SECURITY(S_SYMB, S_CO_ID) in dbs
    hash on (S_SYMB) with 1000 buckets;
- Reduce rootnode contention on small to medium sized tables.

Autonomics

- **Idle User Time Out:**
  - A scheduled task that terminates idle database sessions longer than a configurable time threshold.
- **Bad Index Alert:**
  - A daily scheduled task that looks for indexes which have been marked as unusable and creates an entry in the administration database for each index it finds.
- **Auto CPU VP Configuration:**
  - A system startup task will add additional Virtual Processors based on the number of system cpus.
- **Auto In-Place-Alter Completion:**
  - A weekly scheduled task checks for tables that have an outstanding in place alter.
- **Auto Compress/Repack/Shrink/Defrag on Tables:**
  - Each operation meeting threshold criteria can also be separately enabled, disabled and configured via specific entries in the administration database.

Informix 11.7 – *Creating Growth Opportunities for Clients and Partners*

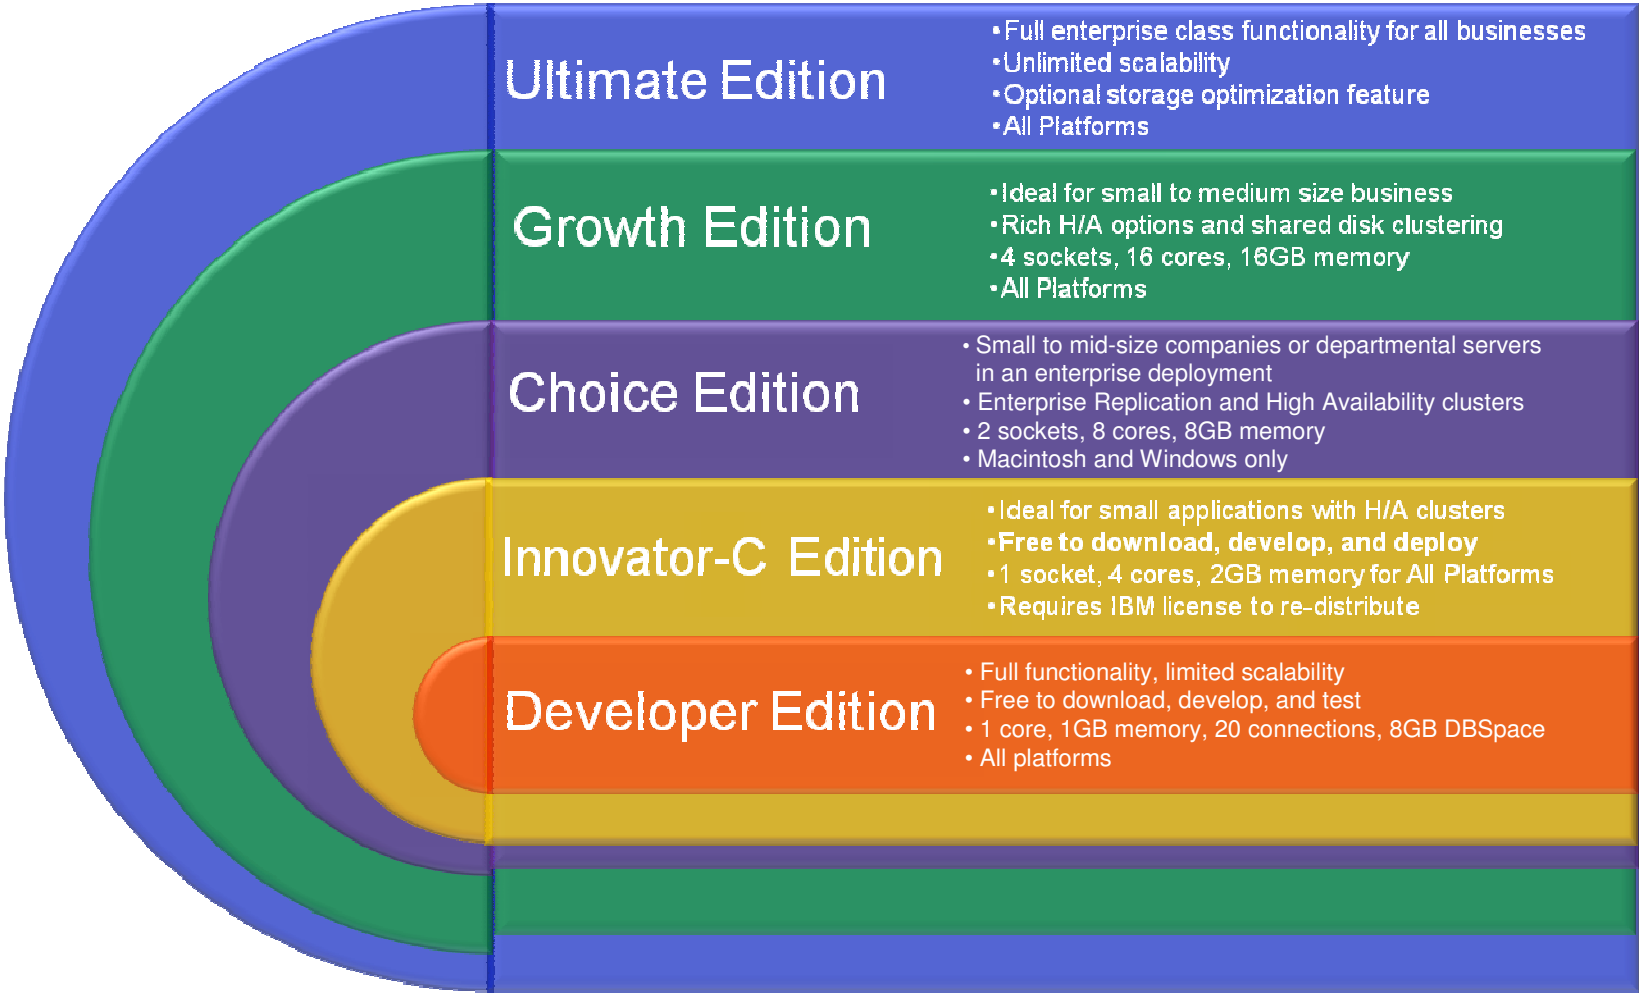

Informix Technical Roadmap

Development

Flexible Grid

Performance

Applications

Availability

Security

***Informix vNext+ 2012**

- Warehouse/BI improvements
- Benchmarks
- Embed SolidDB into IDS
- Informix for Hadoop/Cloud
- Informix for Handheld devices
- 4GL stored procedures
- Support for multi-temperature data
- Enhance Deep Embed
- Enhance Industry Offering (Utilities)

***Informix vNext++ 2014**

- Distributed query processing
- Materialized Query Tables
- Extend Grid to non-database sources

***Informix vNext+++ 2016…**

Continue to invest in Informix to increase its strengths and to open new markets

** Features Subject to change*

Logo