IBM Informix Dynamic Server (IDS)

IDS 11* Technical Features/Enhancements Overview

Jan Musil

jan_musil@cz.ibm.com

* Includes version 11.10 and 11.50[5]

12 January 2010

Content

- **Feature Summary**
- **Administration and Usability**
- **Business Continuity**
- **Performance Enhancements**
- **Security**
- **Application Development**
- **Integrated solutions and related offers**
- **Q & A**

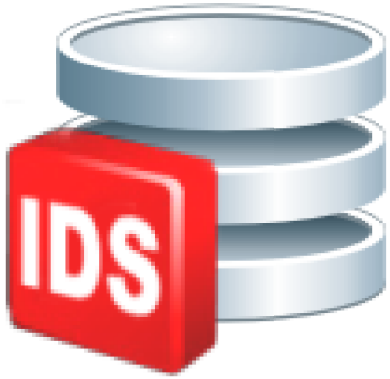

★ IDS 11 Features ★

Administration and Usability				
Recovery Time Objective (RTO)	Deployment wizard: modular installation	Better Install on Windows	Replication Management	Forcing the Database Server to Shut Down
AIO VP automatic tuning	Schedule admin tasks	Improved statistics maintenance	Automatic Update Statistics	Enhanced Support for Multibyte Character Strings
LRU MIN/MAX automatic tuning	Dynamically change ER config params and env vars	Wait on server startup	Admin mode enhancement	New default for the INFORMIXTERM environment variable for terminal set up on UNIX
interval/automatic checkpoints	Better install config values	Enhanced OAT	Monitoring ER thru SMI tables	ER: Truncate table
Monitor and analyze SQL stmt	SQL Admin API	Improvements of SQL admin API	Compression	ER Synchronization memory control
Improved onconfig.std	Explain enhancements	Limit number of sessions	Additional OAT Enhancements	Backup/restore to directories with ontape
ER: rename columns, tables, and DBs	Dbexport and dbschema enhancement (-si)	More onconfig parameters made dynamic	Enhancements to the ER plug-in for the OAT	

★ IDS 11 Features ★

Business Continuity		Security
Enhanced Enterprise Replication (ER)	HA: Configuring RS Secondary Server Latency for Disaster Recovery	Encrypted HDR communication
Continuous High Availability Data Replication (HDR)	ER: ER Stops if Memory Allocation Fails	SSL encryption
Updatable Shared Disk Secondary (SDS)	ER: Notification of an Incorrect Log Position When ER Restarts	Label-based access control (LBAC)
Updatable Remote Standalone Secondary (RSS)	ER: Workaround: Reclaiming Space from ER Paging Smart Large Objects	Backup and restore filters
Connection Manager / Failover Arbitrator	ER: Improving the Performance of Consistency Checking with an Index	Single sign-on
Dynamic reconfig of Connection Manager (CM)	ER: Specifying the Range of Data Sync Threads to Apply Replicated Transactions	Encryption Expert
Continuous Log Restore (CLR)		Improved Installation and Runtime Security

★ IDS 11 Features ★

Application Development (AD)			
Derived table in FROM clause	Basic text search Server Extension	J/Foundation JRE 5.0	BTS XML attribute indexing
Savepoint support	DRDA support	Auto-reprepare	XSLT Support
BTS support in HA clusters	DBAccess “C” style comments support	External directives control for a session	Expanded AD Technologies
Multiple trigger enhancements	SPL enhancements	32KB max return from string functions	Change Data Capture API
XML publishing and manipulation	Dynamic SQL in SPL	UPDATE, DELETE sub-query support	sysdbopen() / sysdbclose()
Embeddability – Server initialization enhancement on Windows	Enhanced data type support in distributed queries	Oracle (SPL - PL/SQL) compatibility functions	Optimizer directive for ANSI joined queries
Web feature services	IS [NOT] NULL predicate	Indexable Binary Type Server Extension	Loading Data into a Warehouse with the MERGE Statement
Indexable Hierarchical Type Server Extension	Named parameters in JDBC callable statements	Retrieving Data by Using Hierarchical Queries	

★ IDS 11 Features ★

Performance Enhancements Integrated Solution & Related Offers		Integrated Solution & Related Offers
Non-blocking Checkpoints	Modifiable first extent size	Explain output in XML format
Index self-join query plan	Improved and Enhanced concurrency	DRDA session information
Improved parallelism during backup/ restore	BIGINT and BIGSERIAL	Rapid Windows Application Development
ER Performance improvements	No temp tables in Unions	JCC and JDBC Support
VP Private memory	Row versioning	Common drivers for IDS and DB2
Direct I/O support	Automatic Ordering of DBSpaces	Data Studio
Set read for index scans	New ONCONFIG VPCLASS syntax	.Net 2.0 support
ON-Bar Performance Report	Disable IPv6 Support	ZendCore (PHP) Support
Dynamic index compression	Concurrent I/O Support on AIX	IDS Virtual Appliance/ AWS EC2 support

★ Enhancements In 11.50.xC6* Release ★

Administration and Usability	Business Continuity	Data Warehousing
Enhance OAT and Schema Manager	Mach11 XA Support	External Tables
Bundle Data Studio Common Client with CSDK	Mach11 Failover Transactional Support	Light Scan for Varchars Update/Delete Joins via Merge
IPV6 Support on HP, MacOS, and more Linux platforms	External Backup in RSS	Attach/Detach Fragments Optimization
	Connection Manager Proxy Support	

* Only key enhancements indicated here

Technical Overview

- **Administration and Usability**
- **Business Continuity**
- **Performance Enhancements**
- **Security**
- **Application Development**
- **Integrated solutions and related offers**

8

Administration and Usability

- **Deployment Wizard: Modular Installation**
- **Enhanced Open Admin Tool**
- **Compression and Storage Optimization**
- **Automatic Update Statistics**
- **SQL administration API**
- **Schedule administrative tasks**
- **Monitor and analyze recent SQL statements**
- **Backup and restore directories**
- **Recovery Time Objective (RTO)**
- **Replication management**
- **External Backup in RSS**

9

Deployment Wizard: Modular Installation

- Allows typical and custom installation of selected components
- Available in GUI and console mode
- Pre-installed components identified automatically
- Flexibility to install selective subcomponents
  - minimizes disk space for a custom installation
- Component dependency enforced
  - Provides estimated footprint for dependencies
- Maintenance flexibility for installed components
- Flexibility of hands free installation for embedding via '–silent' option

Enhanced Open Admin Tool (OAT)

- PHP-based Open Source Web-based graphical tool
- Define and manage automated tasks
- Create and display performance histograms for analysis and tuning

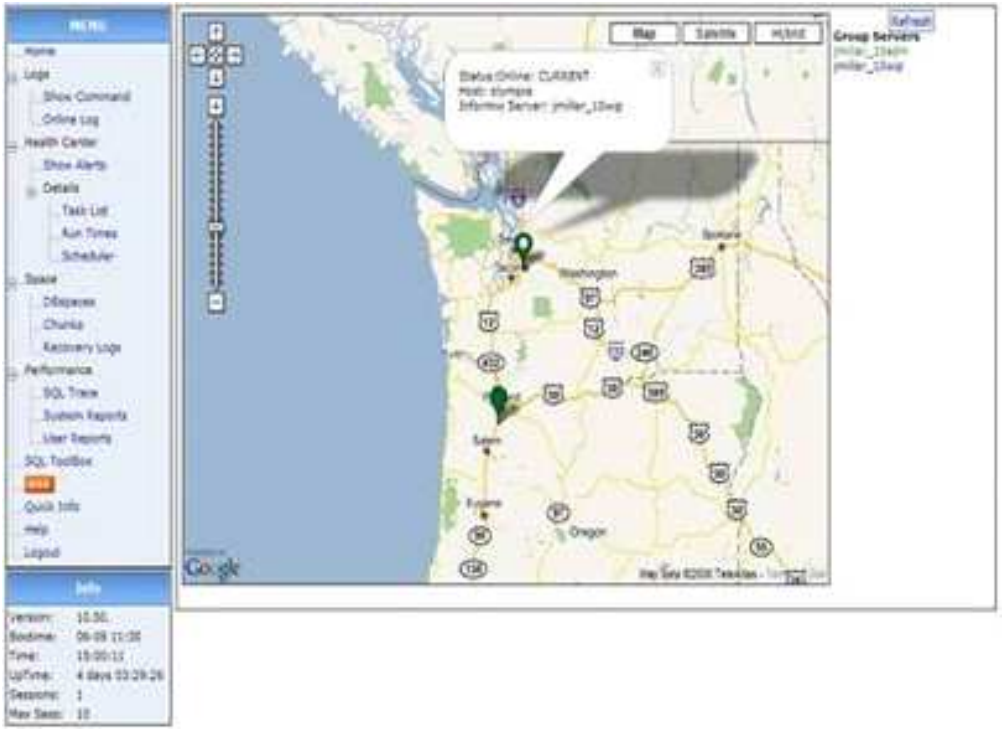

- Administer instances from a single location
- Drill down to complete details
- Easy to customize via extensions

Compression and Storage Optimization

- Row base compression
- Estimation tools helps in finding compression ratio
- Compression can save 40-50% of the db storage requirements
- For IO-bound workloads compression also improves performance
- Significant memory savings – more efficient memory utilization
- Also save on backup storage and disaster recovery storage

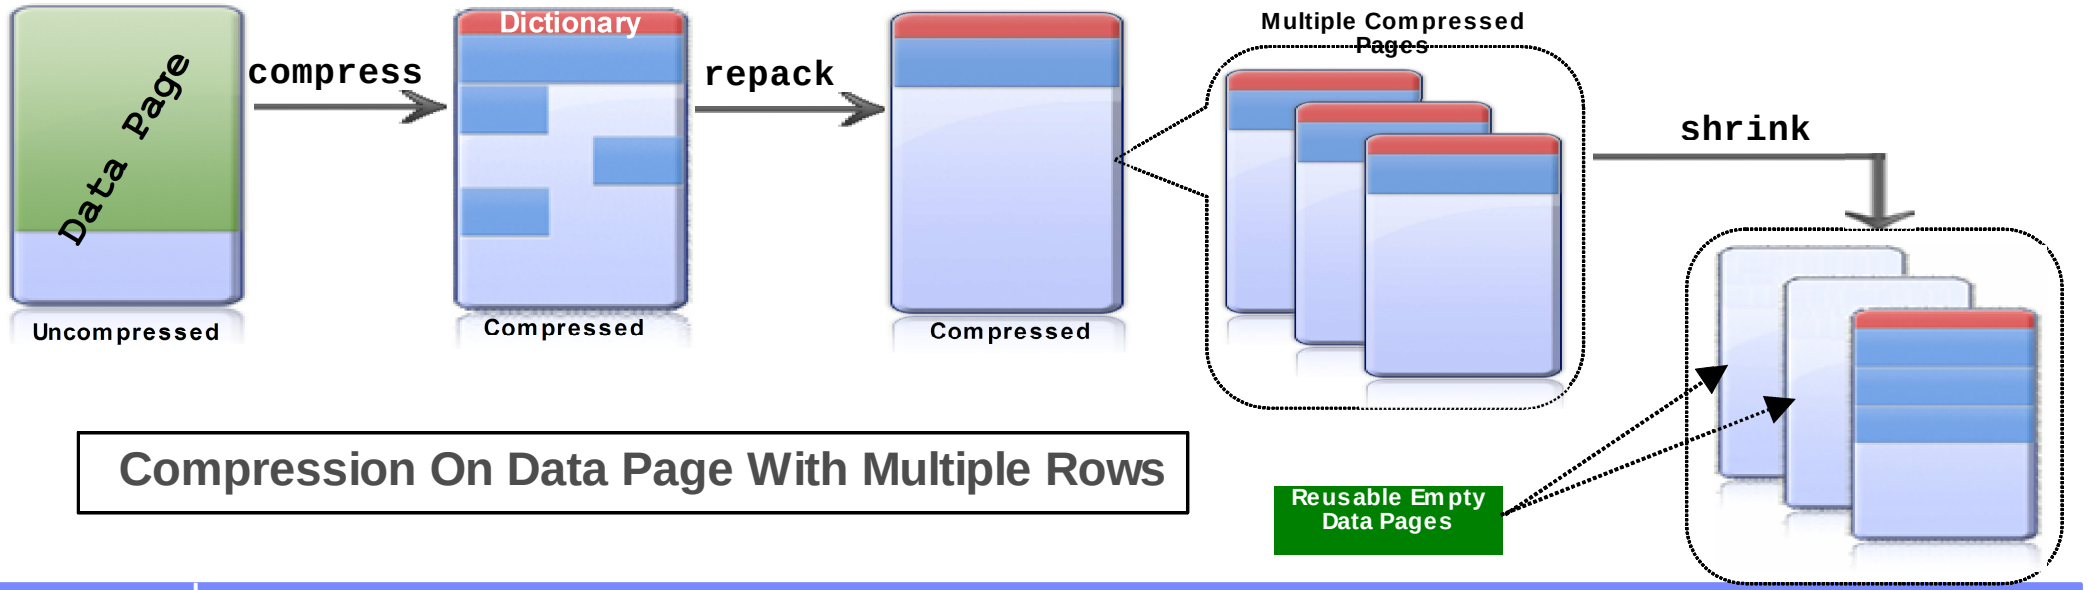

Compression On Data Page With Multiple Rows

Automatic Update Statistics (AUS)

- Automatically maintain optimizer statistics
- Simplifies the repetitive maintenance work on the database to ensure optimal performance
- Easy setup and administration
- Implemented via set of procedures
- Easy admin of AUS policies via OAT

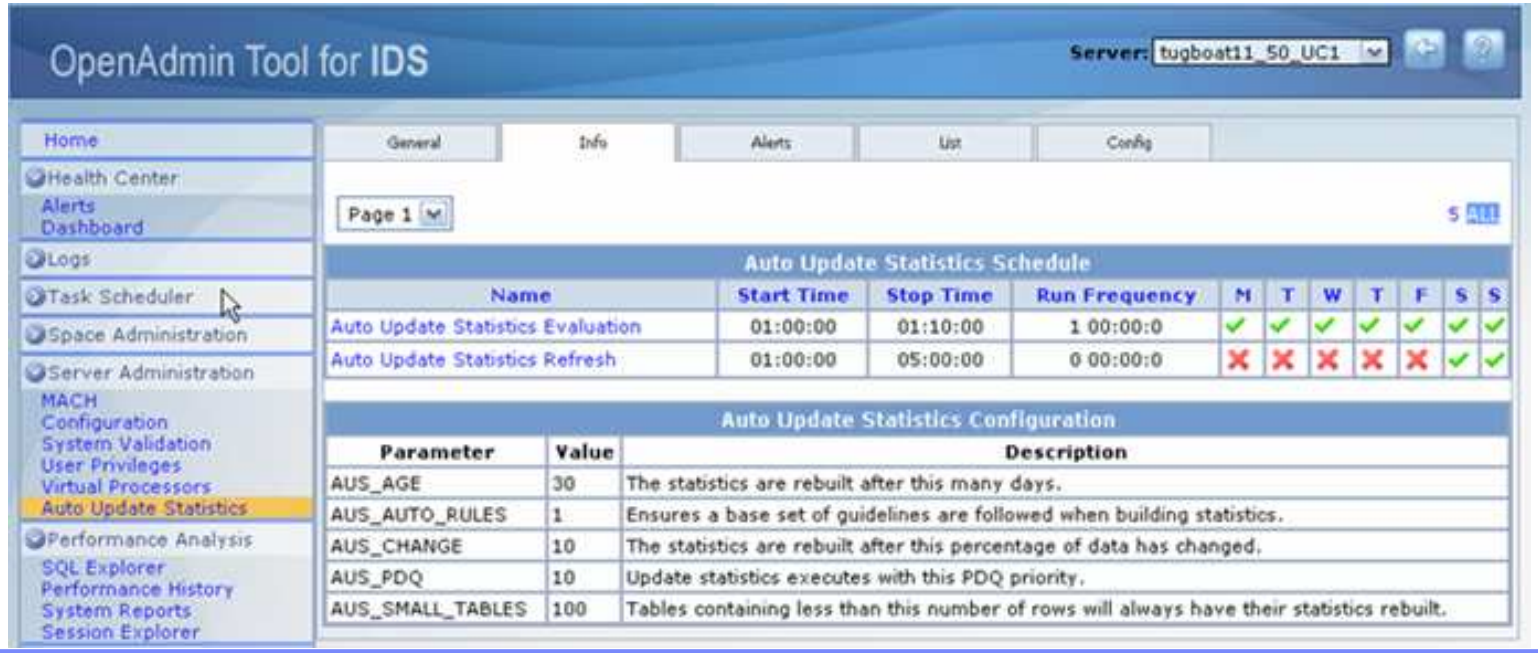

SQL Administration API

- Enables the DBSA to perform administrative tasks remotely by issuing SQL statements
- Accomplished by using built-in functions with arguments:
  - admin( )
  - task( )
- Example:

  To perform “oncheck –ce” command invoke API

```
EXECUTE FUNCTION admin('check extents');
```

Schedule administrative tasks

- Manage and run scheduled maintenance, monitoring, and administration tasks
- Tasks scheduled at predefined times or as determined internally by the server
- Monitor activities and create automatic corrective actions
- Allows DBAs to schedule common administration tasks for automatic execution

Monitor and analyze recent SQL statements

- Monitor the performance of recently executed SQL statements by configuring SQL statement tracing
- Provides statistical information about each SQL statement executed on the system
- Statistical information stored in configurable circular buffer
- By default, this feature is turned off
- Selective enablement per user can be performed
- Allows for easy analysis of SQL statements for performance and tuning purposes.

Backup and restore directories with Ontape

- Enables ontape utility to backup and restore data from the file system without interactive prompts
- Enabled using TAPEDEV and LTAPEDEV configuration parameters
- A backup to a directory has the following advantages:
  - Multiple instances can simultaneously backup to the same directory
  - Use OS utilities to compress or otherwise process the data
  - Configure system for automatic log file backup when full

RTO policy to manage server restart

- RTO: Recovery Time Objective
- Set the amount of time, in seconds, that IDS has to recover from a problem after the server restart
- Set by creating RTO policy using RTO_SERVER_RESTART configuration parameter

Replication management

- Functionality allows for easier administration of ER nodes
- Dynamically change ER configuration parameters and environment variables
  - Add, change, and remove in-memory values for ER for the session
  - Dynamically rename Enterprise Replication columns, tables, and databases
  - Use the RENAME statement to rename a column, table, or database on every participant in the replicate
- Truncate replicated tables
  - Useful to application developers to clear tables before receiving updated data loads
  - Useful to administrators to more easily resynchronize environments that are not logically consistent

External Backup in RSS

- It allows the DBA to create an External backup from a RSS server, so that the primary is freed of this task.
- During the archive the primary server is totally free and can act independently about the archive (The primary server is not blocked).
- During the archive the RSS server is in STOP_APPLY mode:
  - Immediately after the archive checkpoint is processed, the RSS stops applying logical logs, however it keeps receiving them from the primary to avoid hangs in the primary.

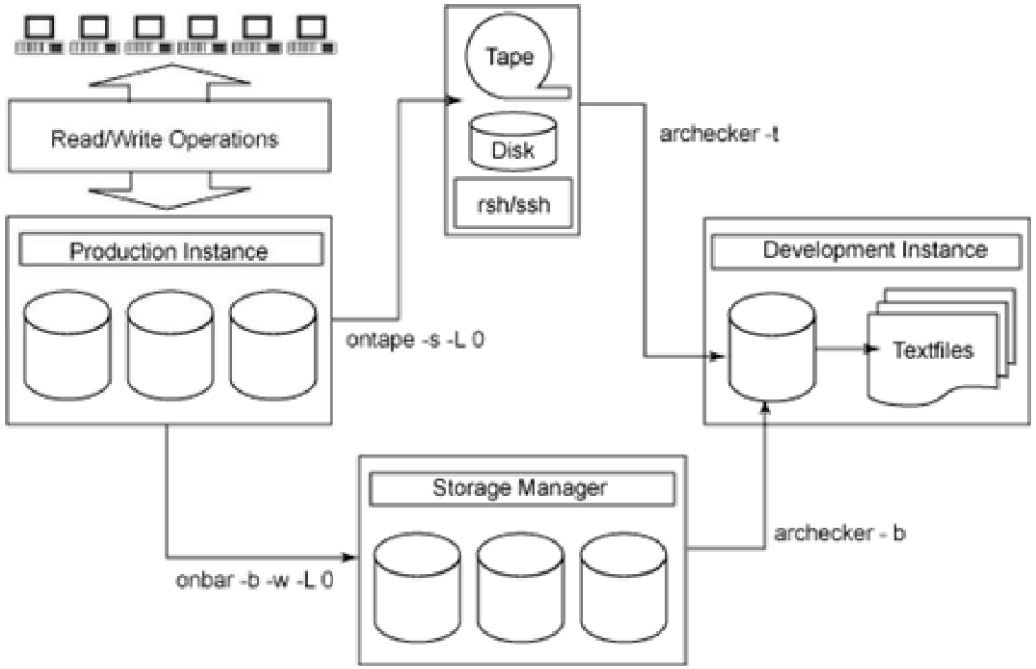

Technical Overview

- Administration and Usability
- **Business Continuity**
- Application Development
- Security
- Performance Enhancements
- Integrated solutions and related offers

21

Business Continuity

- **Continuous Availability Option (CAF)**
- **Shared Disk Secondary (SDS)**
- **High-Availability Data Replication (HDR)**
- **Remote Standalone Secondary (RSS)**
- **Enhanced Enterprise Replication (ER)**
- **Continuous Log Restore (CLR)**
- **Updatable secondary**
- **Connection Manager**
- **Failover Arbitrator**
- **XA Support in Cluster**

What is Continuous Availability Feature (CAF)?

- IDS 11 option
- Feature:
  - Shared Disk Secondary (SDS)
- Complementary features:
  - Remote Secondary Server (RSS)
  - High-availability Replication (HDR)
  - Enterprise Replication (ER)
  - Connection Manager/Failover arbitrator
- All features can be mixed and matched
- Benefits:
  - Resilience to multiple types of failures
  - Horizontal scalability

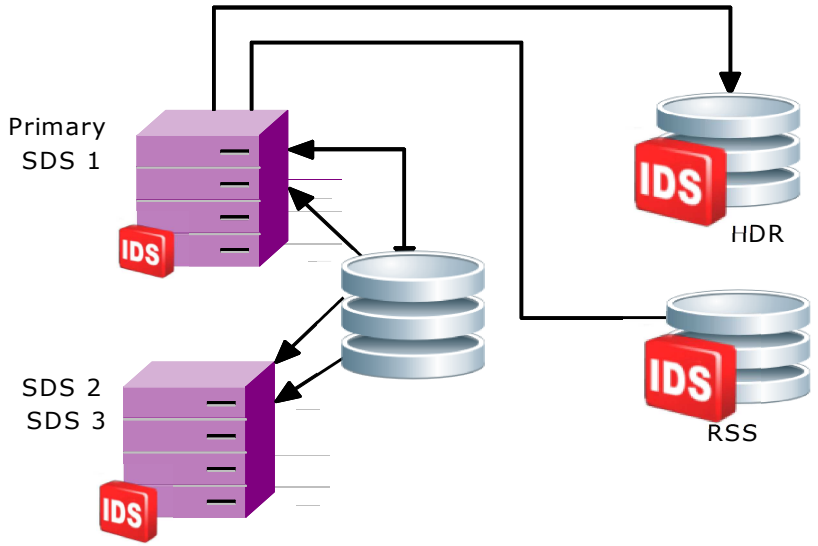

Shared Disk Secondary (SDS)

- Shared disks between multiple instances
- Support for multiple SDS in a cluster
- Encrypted communication available
- Optionally updatable (UPDATABLE_SECONDARY parameter)
- Primary transmits the current Log Sequence Number (LSN) as it is flushing logs.
- SDS instance(s) receives the LSN from the primary and reads the logs from the shared disks.
- SDS instance(s) applies log changes to its buffer cache
- SDS instance(s) resync processed LSN to primary
- Benefits:
  - Protection from machine(s) failure (failover)
  - Horizontal scalability

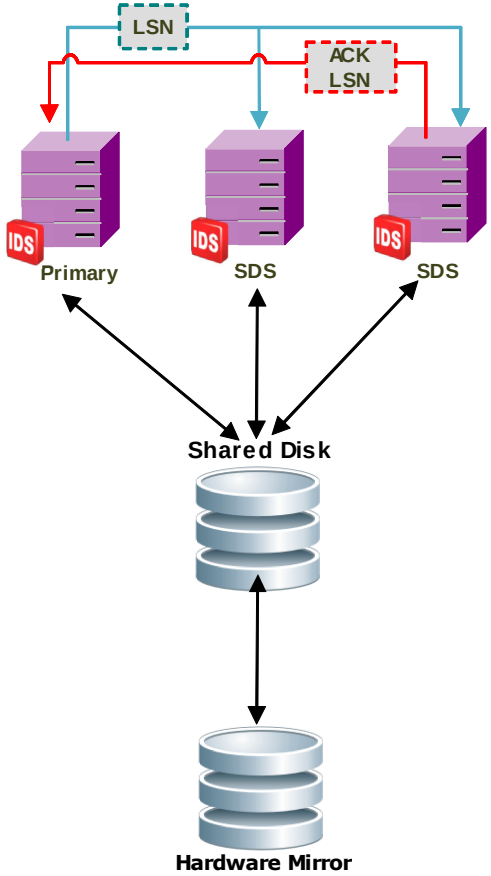

High-Availability Data Replication (HDR)

- **Two identical servers on two identical machines**
- **Optionally encrypted communication**
- **Primary server:**
  - Fully functional server (all DDLs and DMLs)
  - Sends logs to secondary server
- **Secondary server:**
  - Can be local or remote
  - Replays logs to keep in sync with primary
  - Optionally updatable server – DML statements (UPDATABLE_SECONDARY parameter on both machines)
- **Benefits:**
  - Recovery from site failure
  - Offload processing from primary

When Primary server goes down, secondary server takes over as Standard server.

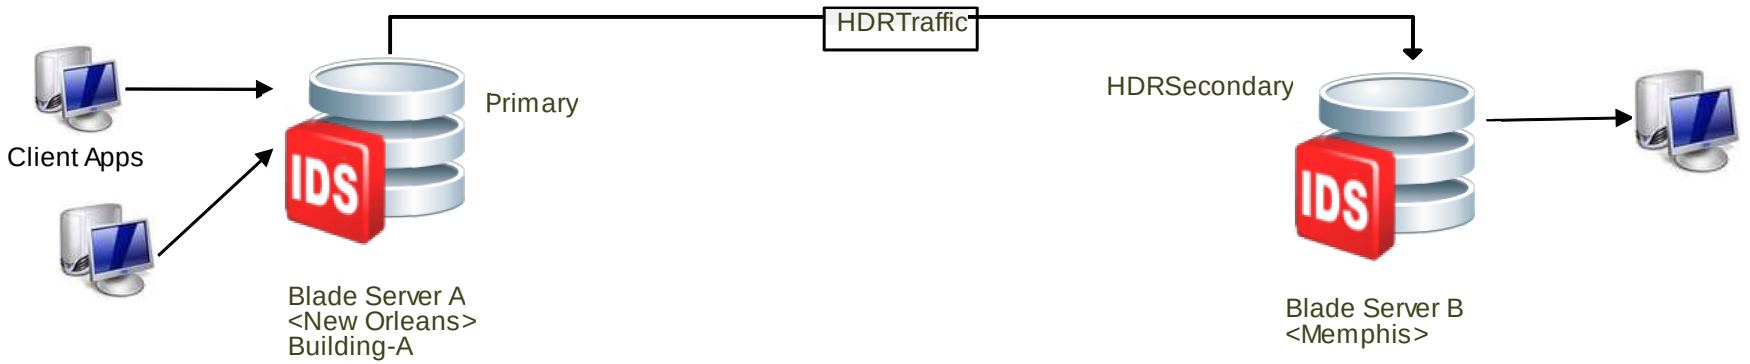

Remote Secondary Server (RSS)

- Primary can support any number of RSS nodes
- Maintains full disk copy of the database server
- Optionally updatable server
  (DELETE, INSERT, SELECT, UPDATE)
- Uses asynchronous full-duplex communication with primary
- Can be promoted to an HDR secondary
- Benefits:
  - Multi-site failure recovery (with HDR)
  - Availability over unstable networks
  - Offload processing

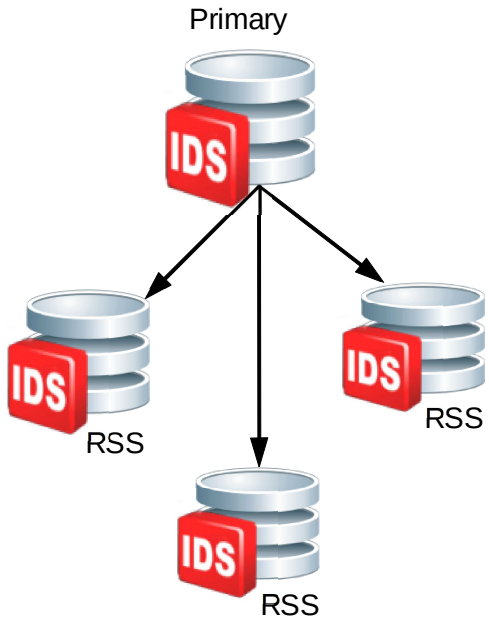

Enhanced Enterprise Replication (ER)

- Peer-to-peer, asynchronous replication
- Replicates rows and columns
- Primary-target and update anywhere support
- High performance, enhanced parallelism
- Supports multiple topologies, large number of nodes
- Supports network encryption
- Supports heterogeneous hardware
- Benefits:
  - Data distribution
  - Workload distribution
  - Data consolidation
  - High performance due to parallelism

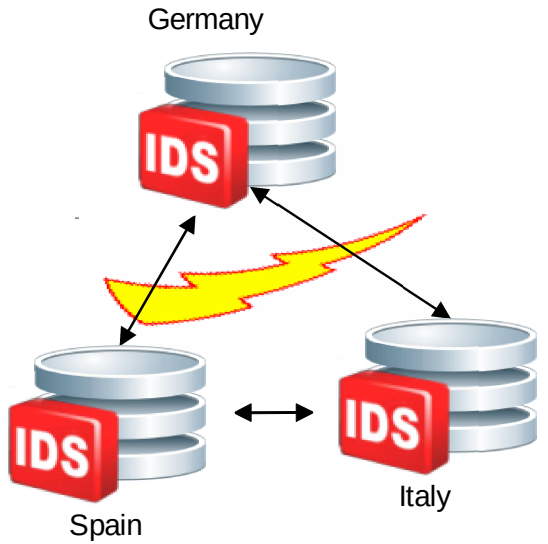

ER Topology

Enterprise Replication : Schema Examples

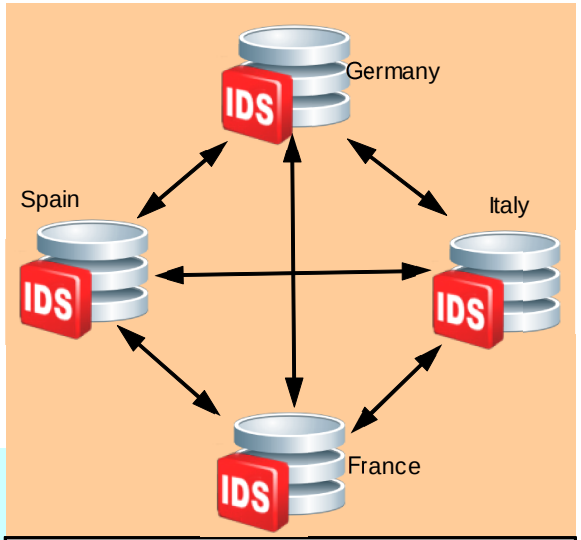

Fully Meshed–Update Anywhere

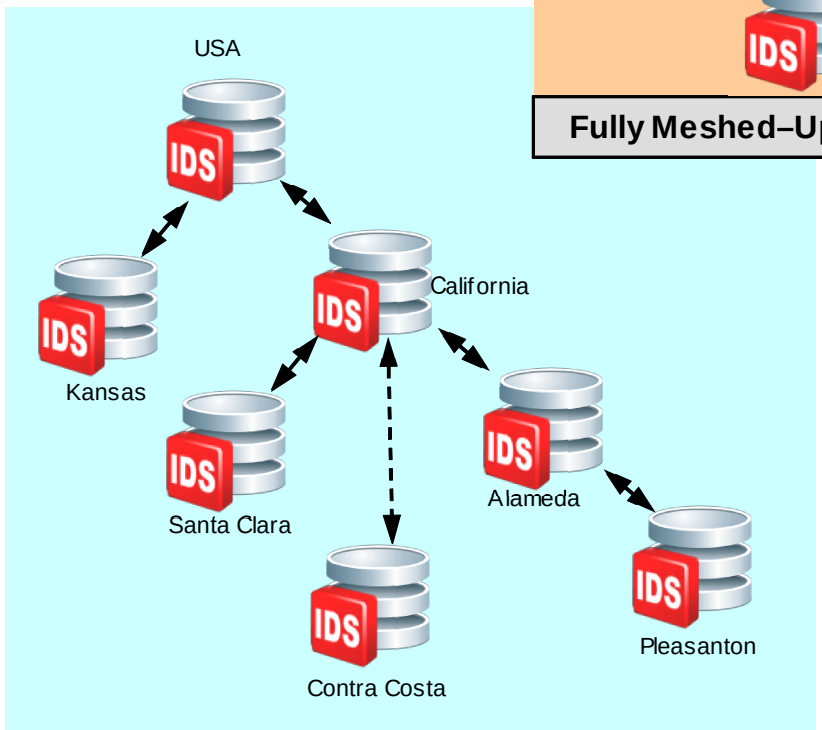

Hierarchical – Push down, Consolidate up

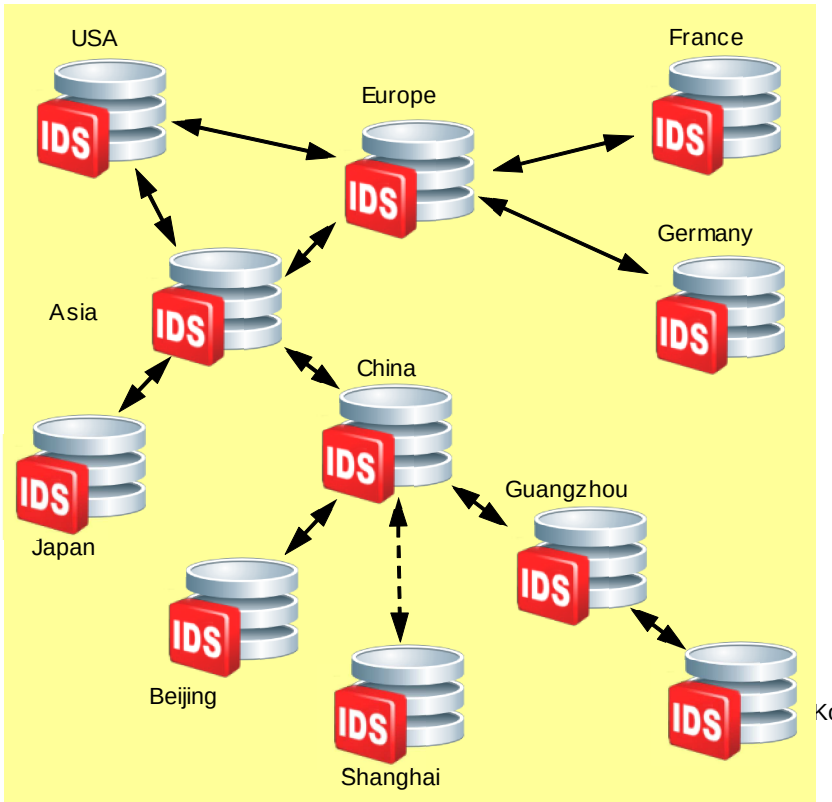

Snowflake / Forest of Trees

Nodes share more with closest nodes,
less with nodes farther away

Continuous Log Restore (CLR)

- Also known as “Log Shipping”.
- Allows logical recovery to span multiple ‘ontape/onbar’ commands.
- Provides a secondary instance with ‘log file granularity’.
- Does not impact the primary server.
- Can co-exist with “the cluster” (HDR/RSS/ SDS) as well as ER
- Can be automated by scripting the log backup alarms
- Useful when backup site is totally isolated (i.e. no network)

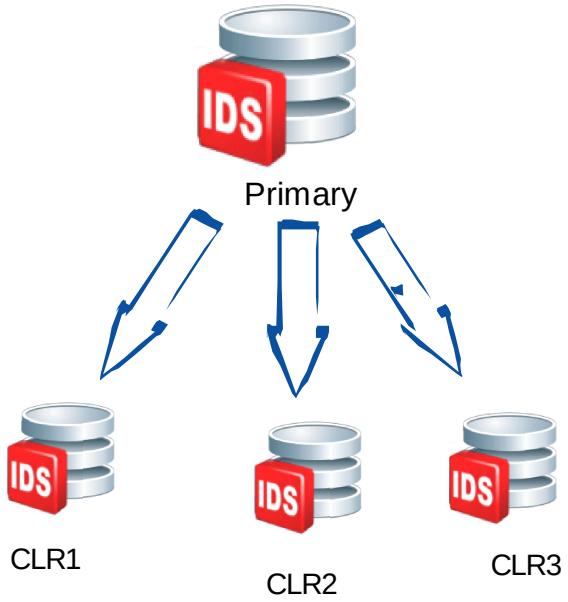

Combination of the Availability Features

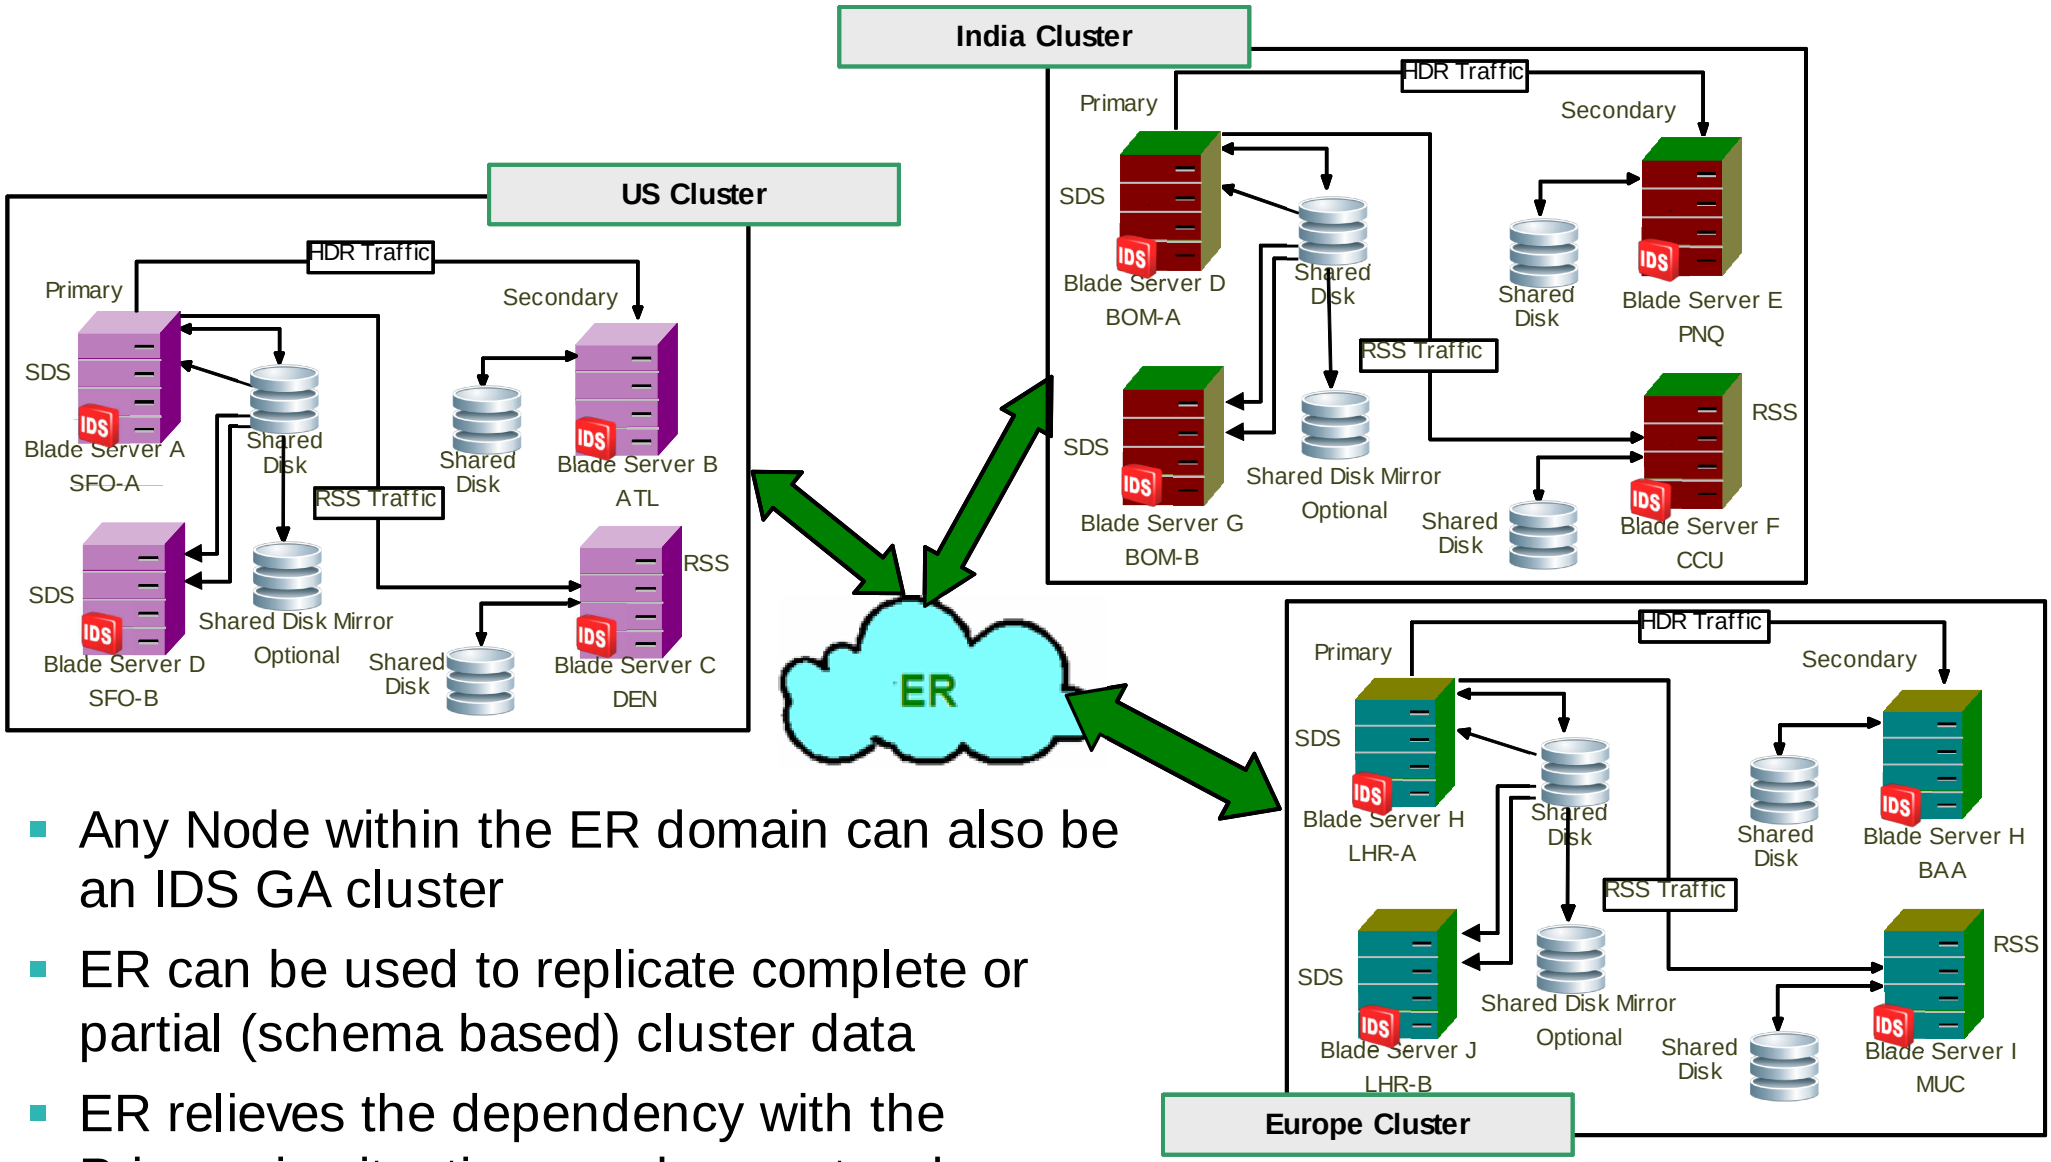

- Any Node within the ER domain can also be an IDS GA cluster
- ER can be used to replicate complete or partial (schema based) cluster data
- ER relieves the dependency with the Primary in situations such as network outages

IDS Global Availability Fabric

Full Suite of High Availability Options to Lower Costs

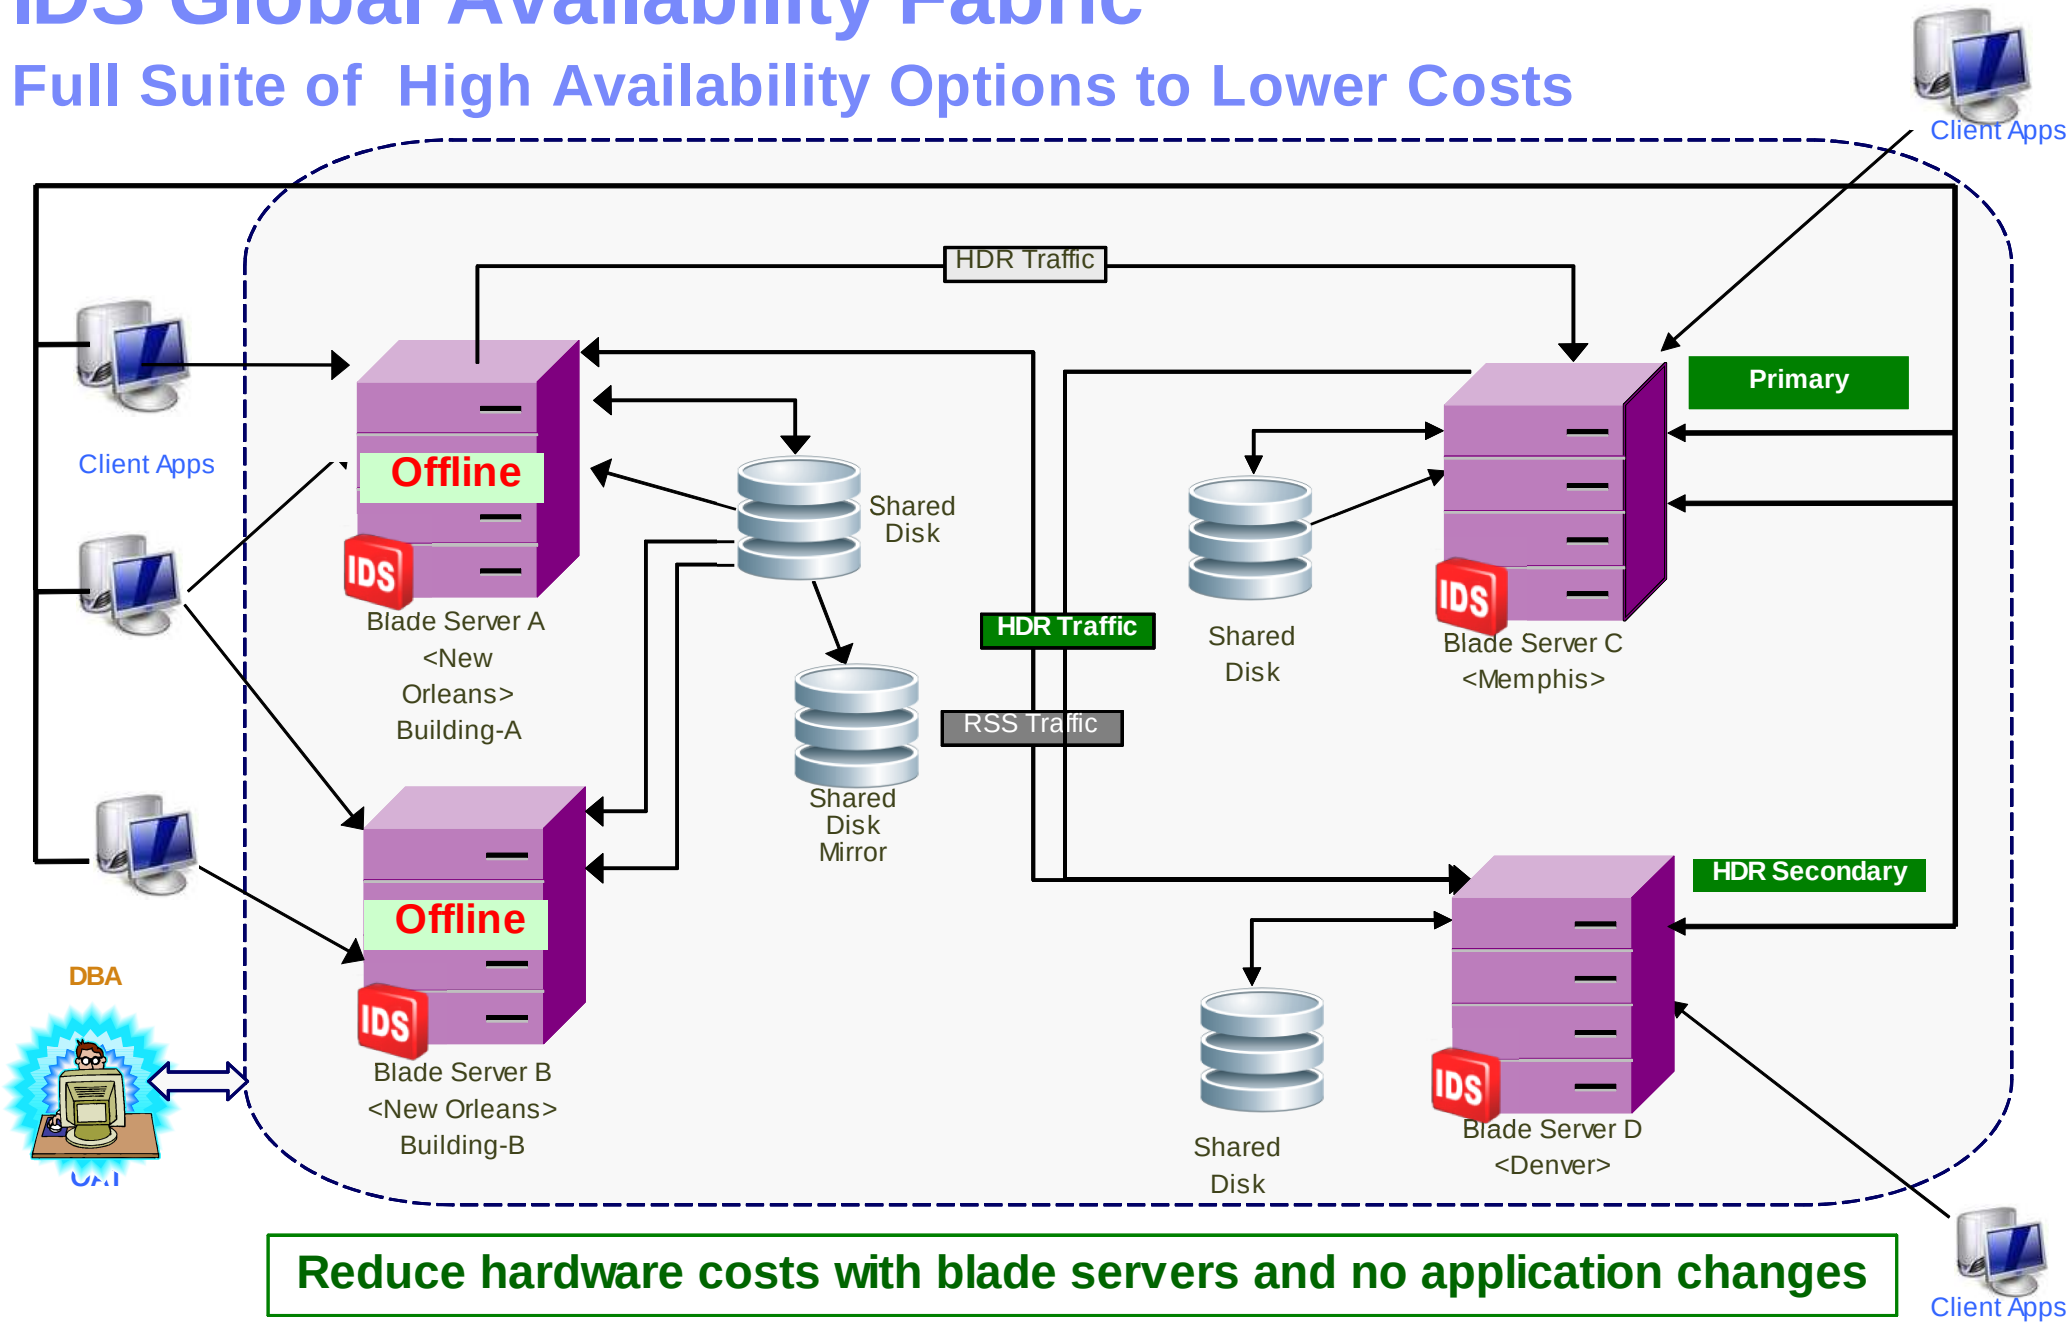

Reduce hardware costs with blade servers and no application changes

Availability Decision Tree

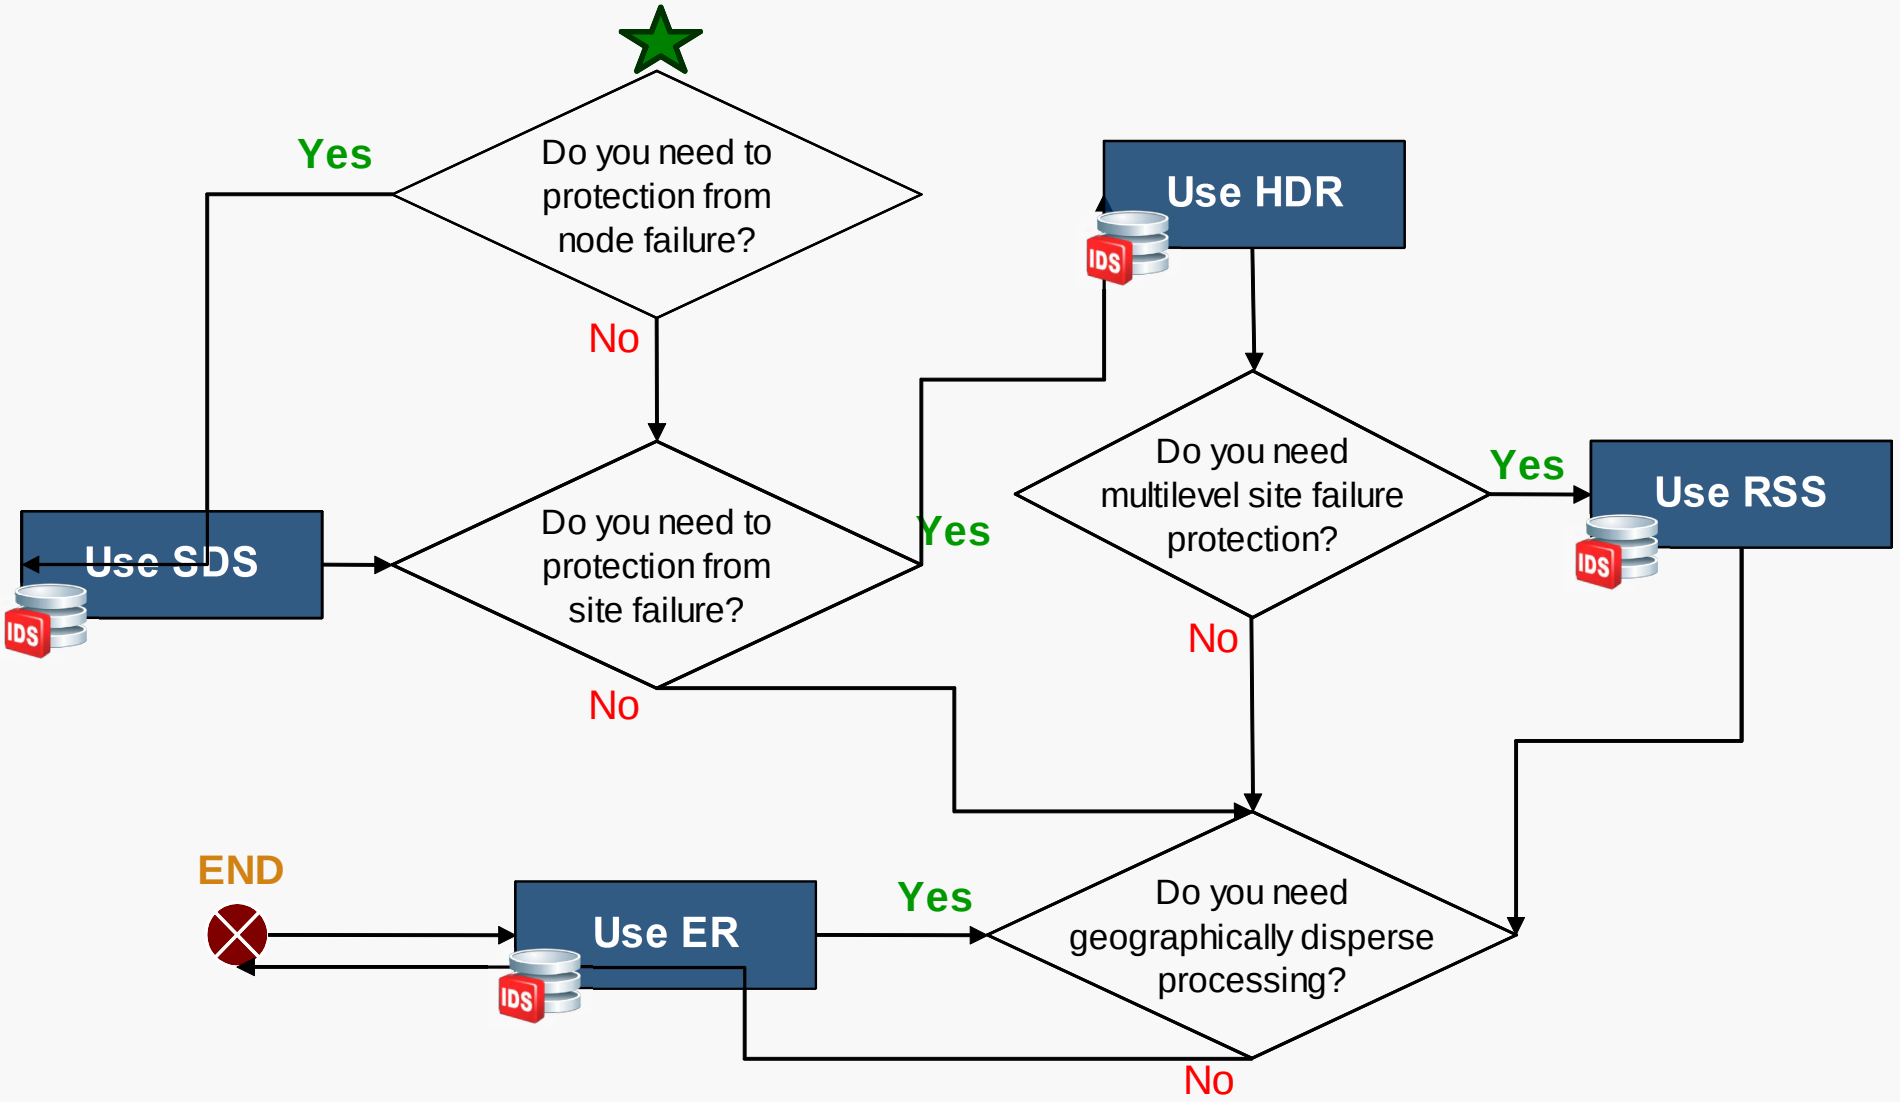

Updatable Secondary

- Secondary servers (SDS, RSS, and HDR) now support insert, update, and delete statements
- Dirty read, last committed, and committed read isolation supported on SDS node, only dirty reads on RSS and HDR
  - More isolation levels supported in future release
- Basic data types, UDTs (those that store data in the server), logged smart BLOBs, and partition BLOBs can be updated
- Only supports DML statements
- Supports temp tables - both explicit and implicit
- Continues to work with ER

Connection Manager (CM)

- Separate daemon program distributed with client bundle
- Performs node failover based on failover config parameter
- Reconfiguration of parameters can be performed dynamically
- Multiple CM can be used to avoid single point of failure
- Re-route connections to the best-fit nodes
- Monitors nodes for failure
- Benefits:
  - Automates node failover
  - Automatic connection re-routing (during node failure) based on Service Level Agreement

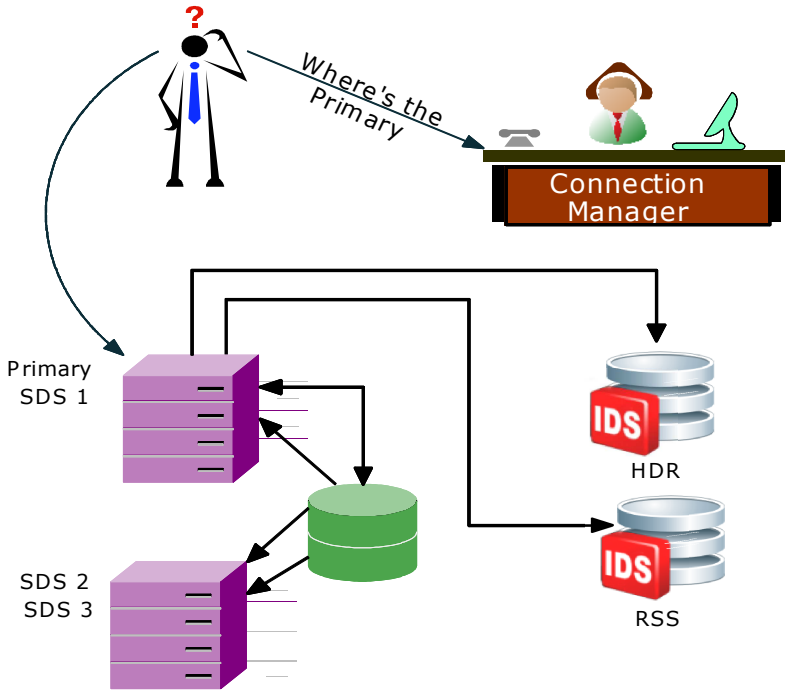

Failover Arbitrator

- Part of the Connection Manager (CM)
- Uses multiple checks to insure failure
- Performs failover when primary is down
- Order of failover defined in the CM Configuration File

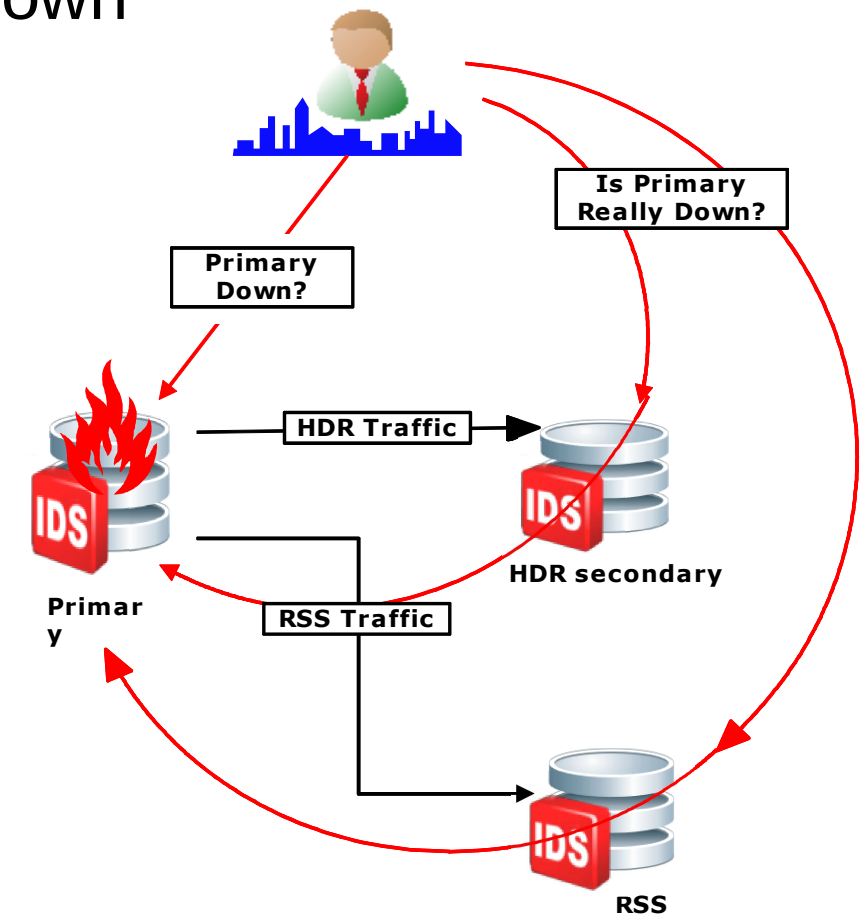

XA Support in IDS Cluster

- X/Open Distributed Transaction (Tx) Processing Model allows multiple resources (e.g. DB, AppSrvr, MsgQ) to be accessed within the same transaction.
- Enables XA interface on IDS secondary server node.
- A global Tx branch can be accessed from any updatable secondary server node in IDS cluster, not necessarily the server node that started the Tx branch.
  - Support both loosely-coupled and tightly-coupled XA Tx
  - Support XA interface via Connection Manager

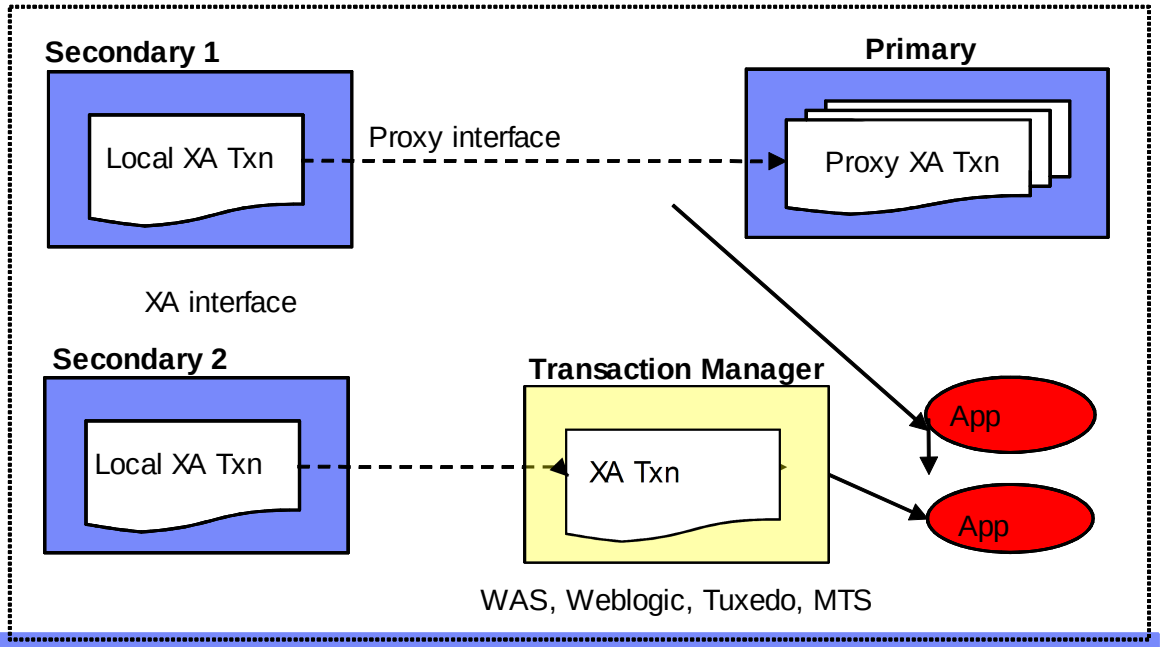

Technical Overview

- Administration and Usability
- Business Continuity
- **Performance Enhancements**
- Security
- Application Development
- Integrated solutions and related offers

37

Performance Enhancements

- **Index Self-Join query plans**
- **ON-Bar Performance Report**
- **Direct I/O for cooked files**
- **Concurrent I/O on AIX**
- **Automatic ordering of dbspaces**
- **Improved parallelism during backup and restore**
- **Non-blocking Checkpoints**
- **Improved and Enhanced concurrency**
- **BIGINT and BIGSERIAL**
- **Light Scan for VARCHAR**

38

Index self-join query plans

- Implemented in the IDS query optimizer
- By default, the optimizer considers this type of scan
- The table is logically joined to itself
- Added support for two new join-method directives:
  - INDEX_SJ : Forces an index self-join path using the specified index, or choosing the least costly index in a list of indexes
  - AVOID_INDEX_SJ : Prevents a self-join path for the specified index or indexes
- Improves query performance on tables with composite indexes.

Using Index Self Join

Lead Keys: c1, c2
Lower Filter c1 = c1 and c2 = c2 and c3 >= 100
Upper Filter c3 <= 102

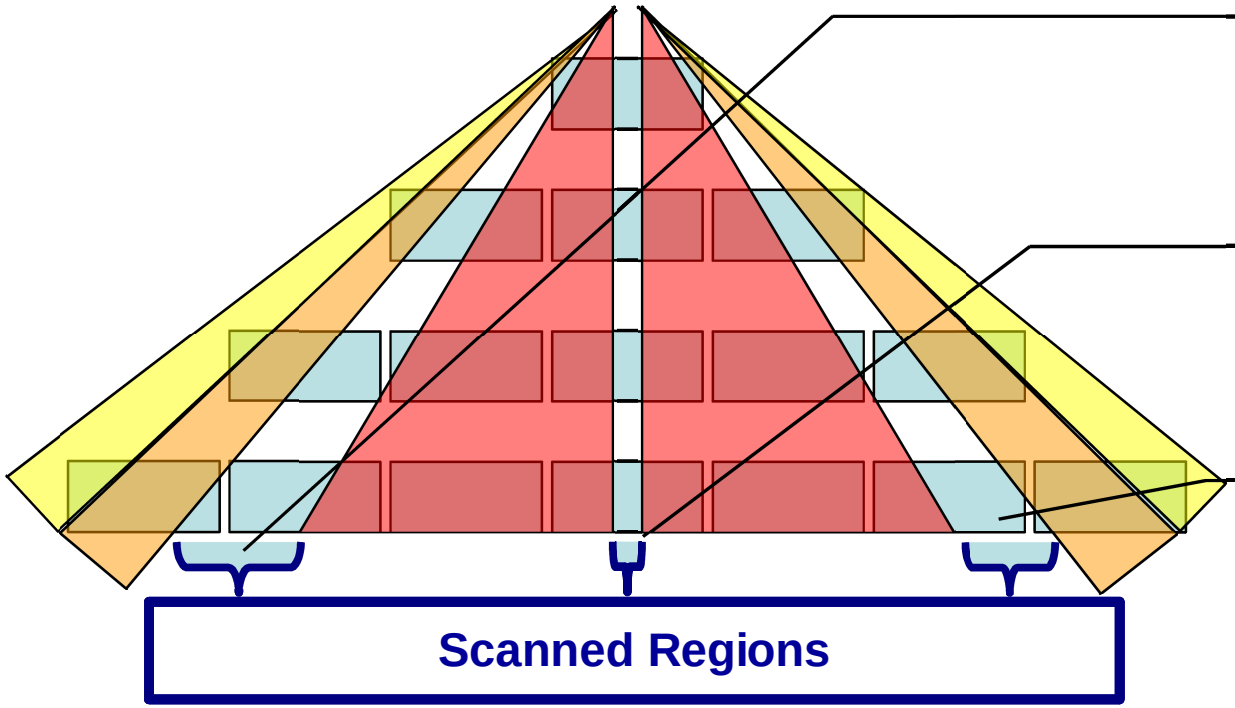

C1 = “Monday”
C2 = “Male”
C3 => 100 & c3<=102

C1 = “Monday”
C2 = “Female”
C3 => 100 & c3<=102

C1 = “Tuesday”
C2 = “Female”
C3 => 100 & c3<=102

ON-Bar Performance Report

- Provides a report of ON-Bar backup and restore performance
- Reporting available via ON-Bar activity log
- Reporting set via BAR_PERFORMANCE configuration parameter
- Granularity of the report cab be configured to contain
  - Sub-second timestamps for ON-Bar processing
  - The transfer rates between ON-Bar and the storage manager
  - The transfer rates between ON-Bar and the IDS instance

Direct I/O for cooked dbspace chunks on UNIX

- IDS allows you to use either raw devices or cooked files for dbspace chunks
- Cooked files are slower due to additional overhead and buffering provided by the file system
- Direct I/O bypasses the use of the file system buffers, hence is more efficient for reads and writes that go to disk
- Direct I/O improves the performance of cooked files used for dbspace chunks
- Direct I/O can be enabled with the DIRECT_IO configuration parameter

If your file system supports direct I/O, performance for cooked files can approach the performance of raw devices used for dbspace chunks

Concurrent I/O on AIX

- Applies only to cooked chunk files
- Benefits:
  - Uses direct I/O (no file system buffering)
  - Avoids unnecessary serialization of I/O
  - Higher I/O rates, especially when data striped across multiple disks
- Enable by setting DIRECT_IO onconfig parameter to 2
- "onstat -d" shows status
  - "C" for concurrent I/O
  - "D" for direct I/O

Automatic ordering of dbspaces during backup and restore

- Intelligent ordering of dbspaces during backup and restore to achieve maximum parallelism
- Dbspaces restored in the same order as backup
- Reduces the backup and restore time

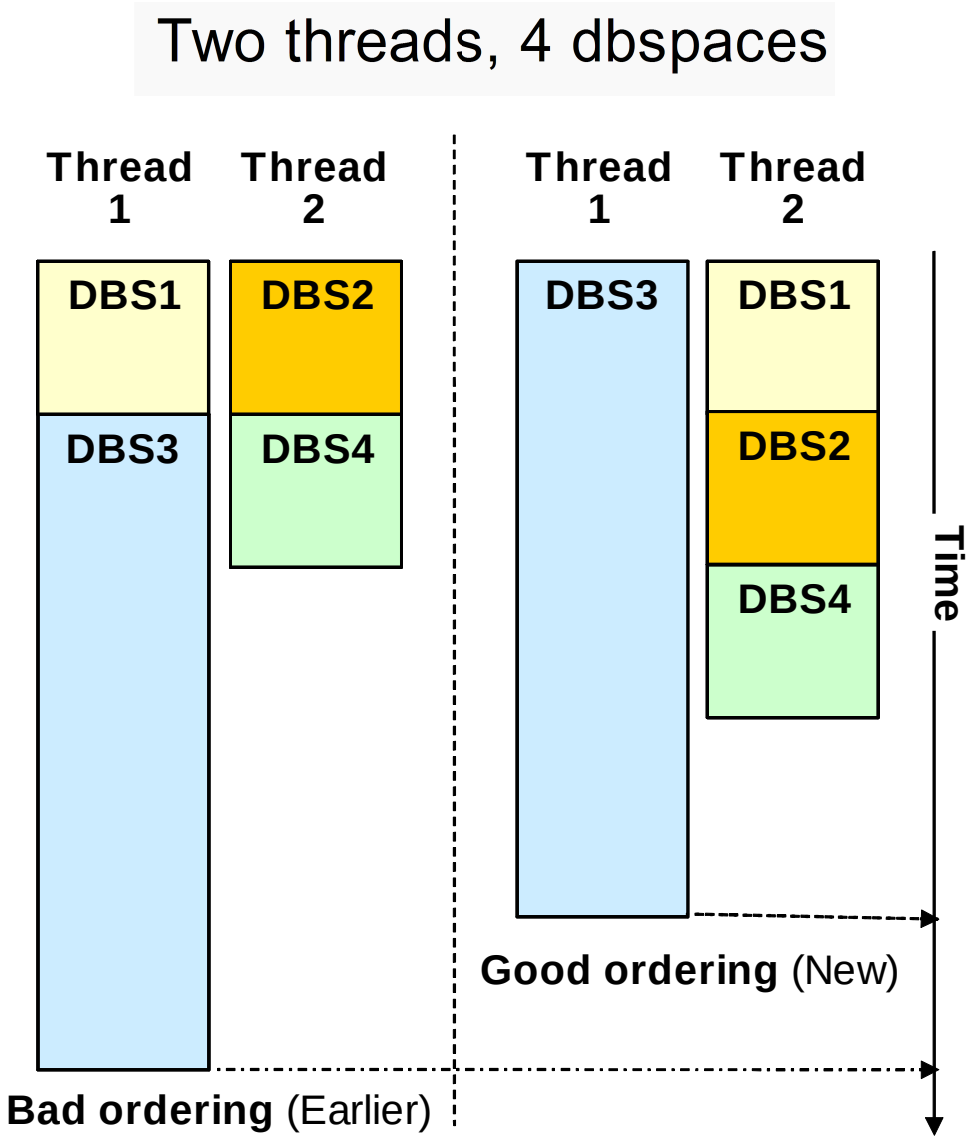

Improved parallelism during backup and restore

- ON-Bar now performs a backup and the restore of a whole system using parallel I/O
- Parallel backup is more efficient than serial backup
- Reduces the total time required for complete backup and restore
- During a parallel backup, multiple processes run simultaneously and back up data to separate dbspaces
- The maximum number of parallel processes for each onbar command controlled by BAR_MAX_BACKUP parameter
- The default value for BAR_MAX_BACKUP is 4

Improved and Enhanced concurrency

- With Committed Read (CR) isolation
  - `SET ISOLATION COMMITTED` READ uses `LAST COMMITTED`
  - Reduces the risk of locking conflicts
  - Returns the most recently committed version of the rows
  - Supports B-tree indexes and functional indexes
  - Facilitates porting from Oracle

- With private memory caches for Virtual Processors (VP)
  - Configure a private memory cache for every CPUVP
  - Decreases the time of server memory allocation on large multiprocessor computers

Non-blocking Checkpoints

- Allows transaction processing while checkpoint processing is occurring
- Allows LRU flushing to be relaxed
- Most checkpoints do not block transactions during buffer flushing except:
  - Admin and archive checkpoints
  - Checkpoints running short on resources (physical log)
- New, patented checkpoint algorithm:
  - Phase A recovery has been removed
  - Physical logging activity increases 9 to 10 times

Will need to increase size of physical log!

BIGINT and BIGSERIAL

- Support for new ANSI standard SQL data types BIGINT and BIGSERIAL
- A better performance alternative to INT8 and SERIAL8 data types
  - BIGINT/BIGSERIAL are native 8-byte integers
  - INT8/SERIAL8 are 10-byte structures made up of two 4-byte integers

Light Scan for VARCHAR

- Improves warehouse performance by
  - Enabling light scans on tables with rows that can span pages, which includes:
    - Varchar, lvarchar, nvarchar
    - Compressed tables
    - Any table with rows longer than a page
- Light Scan for fixed length rows already enabled
- ONCONFIG file:
  - BATCHEDREAD_TABLE 1
- Shell environment (set in IDS shell before starting):
  - Csh family: IFX_BATCHEDREAD_TABLE 1
  - Ksh family: export IFX_BATCHEDREAD_TABLE=1
- Single session:
  - set environment IFX_BATCHEDREAD_TABLE ‘1’;

Technical Overview

- Administration and Usability
- Business Continuity
- Performance Enhancements
- **Security**
- Application Development
- Integrated solutions and related offers

50

Security

- **Advanced Access Control feature with LBAC**
- **Encrypted HDR communication**
- **Secure Socket Layer (SSL) Encryption**
- **Kerberos Single Sign On (SSO)**
- **Encryption Expert**

51

Advanced Access Control feature with LBAC

- LBAC : Label Based Access Control
- An implementation of Mandatory Access Control (MAC)
- Offers row-level and column-level protection
- Security label assigned to protected object (row/ col)
  - Label stores data classification/sensitivity information
- User assigned security label for access restriction

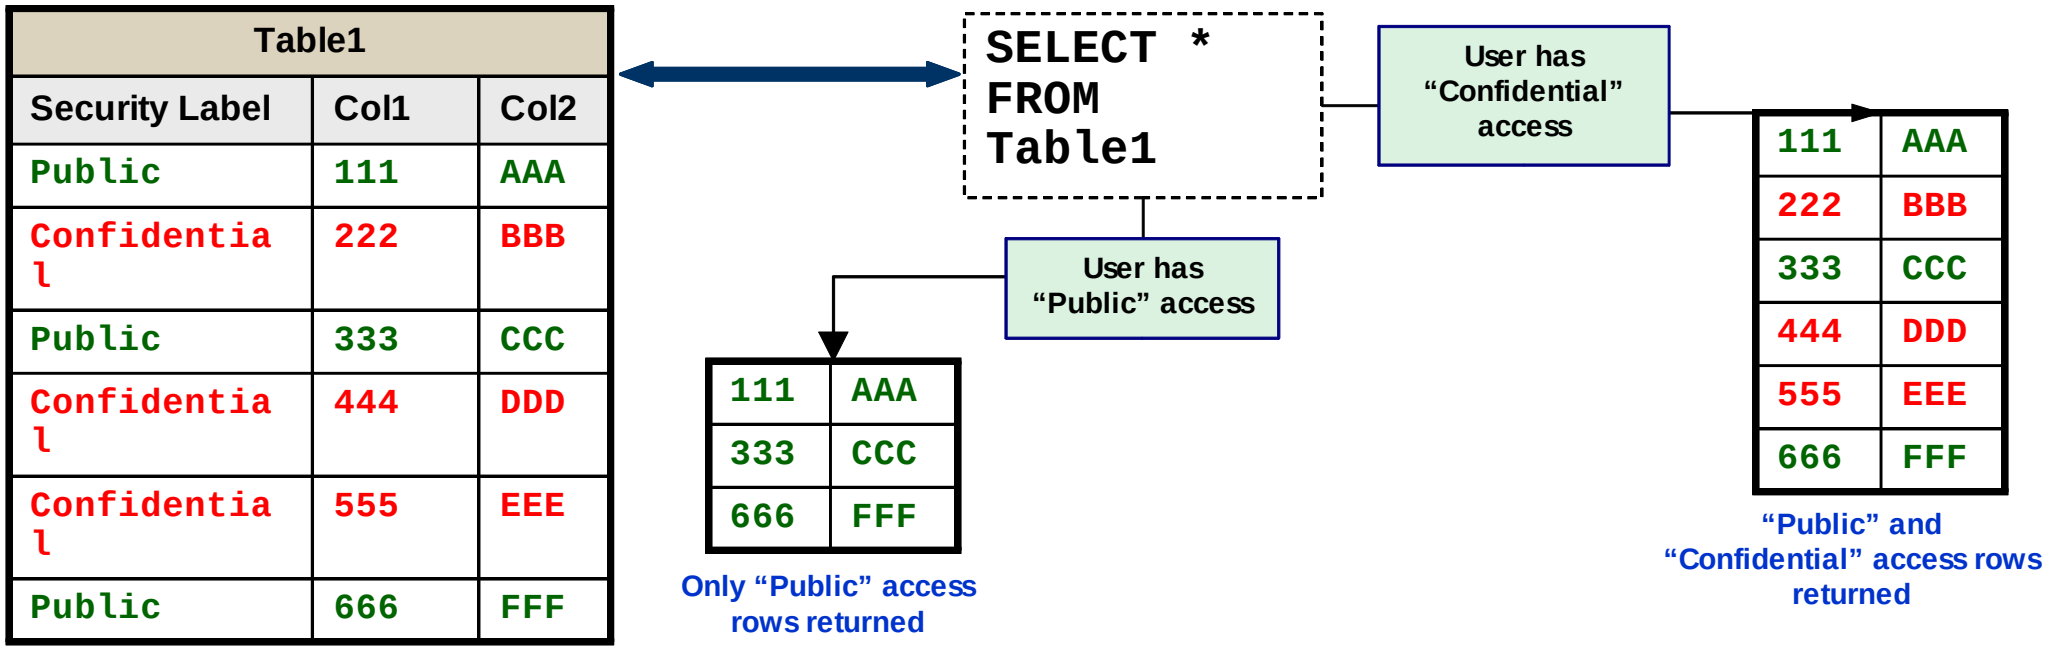

Encrypted HDR communication

- Provides secured transmission of data for HDR
- Enabled by administrator using configuration parameter ENCRYPT_HDR
- Once enabled
  - HDR primary encrypts data before sending
  - HDR secondary decrypts data after receiving

Secure Socket Layer (SSL)

- SSL is a well known security standard
- Communication protocol that provides privacy and integrity of data transmitted over the network
- Uses encryption to provide an end-to-end secure connection
- Supported with SQLI clients, DRDA clients and server to server communication .
- On SQLI: ISTAR, HDR, ER, SDS/RSS

Digital Certificates, Certificate Authority & Keystores

- SSL uses digital certificates to exchange keys for encryption and server authentication
- Digital Certificates are electronic ID cards issued by trusted parties know as Certificate Authority (e.g. VeriSign)
- Digital certificates are stored in key database (also known as keystore)

IBM's Global Security Kit bundled with IDS and CSDK provides an iKeyman utility that can be used to create keystores and manage digital certificates

Kerberos Single Sign On (SSO)

- Users enter their password once to gain access
- Password entered during login
- Password need not be reentered to authenticate in IDS
- Authentication is invisible to the user
- IDS uses Kerberos for SSO support
- Easy administration
  - User does not have to reenter password
  - Password management becomes centralized and easier
- SSO becoming ubiquitous with web apps and others
- Allows IDS to easily integrate with existing single sign on infrastructure

Encryption Expert*

- Encrypts Archived Data
- Strongly protects private and confidential data online and offline environments
- Helps comply with regulations and legislative acts example, PCI-DSS
- Has centralized policy and key management to simplify data security management
- Helps achieve key goals:
  - Minimize the risk of losing data backups that are written to portable media and stored offsite for long periods of time
  - Reduce the risk of replicated data being stolen or lost
  - Limit user access to sensitive data.

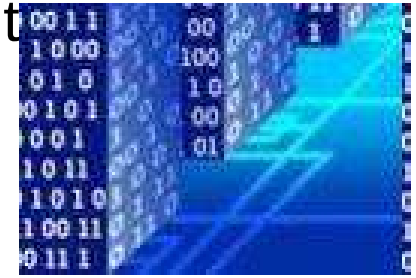

* Separately Charged Product

Technical Overview

- Administration and Usability
- Business Continuity
- Performance Enhancements
- Security
- **Application Development**
- Integrated solutions and related offers

58

Application Development

- **Automatic re-prepare**
- **Derived tables in the FROM clause**
- **Loading data with MERGE statement**
- **Update and Delete with MERGE statement**
- **External Table**
- **Optimizer directives**
- **Attach/Detach Fragment Optimization**
- **Distributed query enhancements**
- **Named parameters in JDBC Callable Statement**
- **Multiple Trigger Enhancements**
- **Sysdbopen()/Sysdbclose()**
- **Dynamic SQL in SPL**

Cont'd …

59

Application Development

- SPL Enhancements
- SPL - PL/SQL Compatibility
- XML Publishing and Manipulation
- XSLT Support
- Basic Text Search (BTS) Index
- BTS XML Attribute Indexing
- Hierarchical data type
- Retrieving data with Hierarchical queries
- Indexable binary data types
- Web Feature Service for geospatial data
- IBM Data Server Client
- Expanded Technologies

60

Automatic re-prepare

- Automatically detect changes to underlying objects and re-prepares the statement if needed
- After any DDL operation modifies the schema of a database table, the database server automatically performs the following actions:
  - Issue the UPDATE STATISTICS statement to recalculate routine statistics for all SPL routines that reference the table
  - Issue the PREPARE statement to update any prepared objects that reference the table

Derived tables in the FROM clause of queries

- The SELECT statement can now include syntax that complies with ISO/IEC 9075:1992, the SQL-92 standard
- Specifies uncorrelated sub-queries in the FROM clause as a data source for the query called derived tables or table expressions
- Supports UNION, or joined sub-queries, including OUTER joins
- Can include the ORDER BY clause
- AS correlation for columns allowed in FROM clause to declare temporary names within the query

Loading data with MERGE statement

- Transfers data to target table using UPSERT
  - UPSERT = a combination of INSERT and UPDATE in one statement
- Replaces current UPSERT logic implemented in client applications.
- Extension of ANSI/ISO 2003 standard.
- Key components of ETL (ELT)* in data warehouse environments.
- Merges rows to a target table based on a condition.
  - TRUE: Update the target table row.
  - FALSE: Insert the target table row

```
MERGE
  INTO TARGET t
  USING SOURCE s
  ON t.c1=s.c1
  WHEN MATCHED THEN
    UPDATE SET t.c1=t.c1+5
  WHEN NOT MATCHED
    THEN INSERT (t.c1, t.c2, t.c3)
    values (s.c1, 70, 70);
```

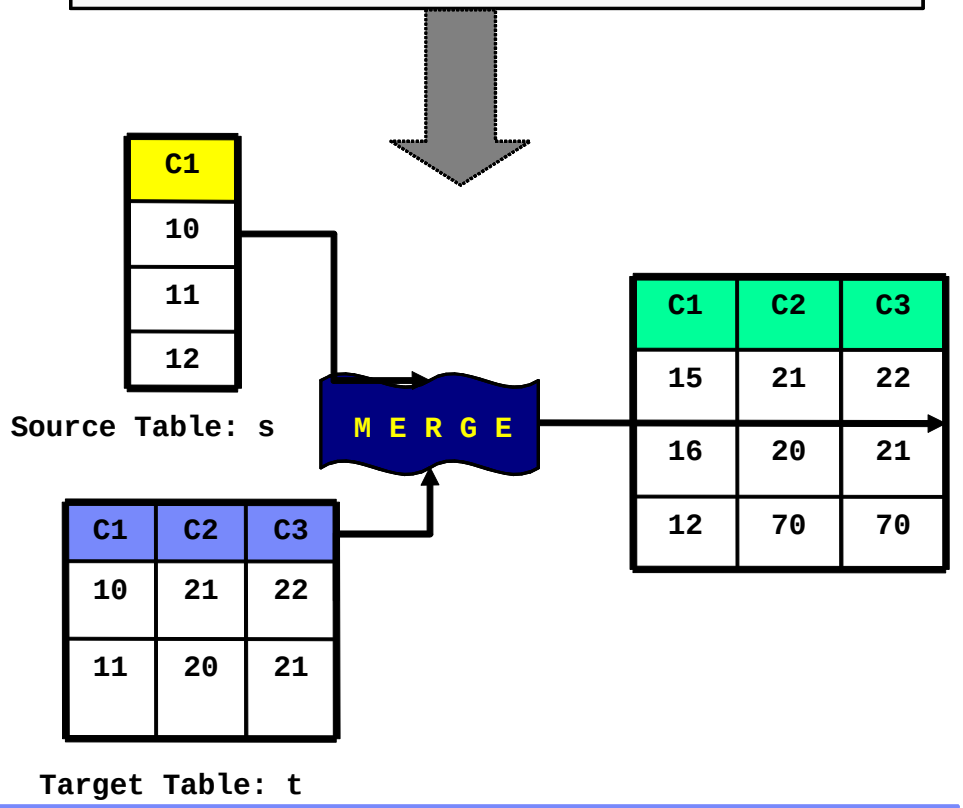

Update and Delete Joins with MERGE statement

```
MERGE INTO bonuses D
USING employee S
ON D.employee_id = S.employee_id
WHEN MATCHED THEN DELETE;

Optional: "WHEN NOT MATCHED"
```

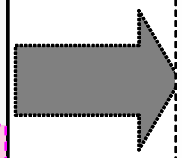

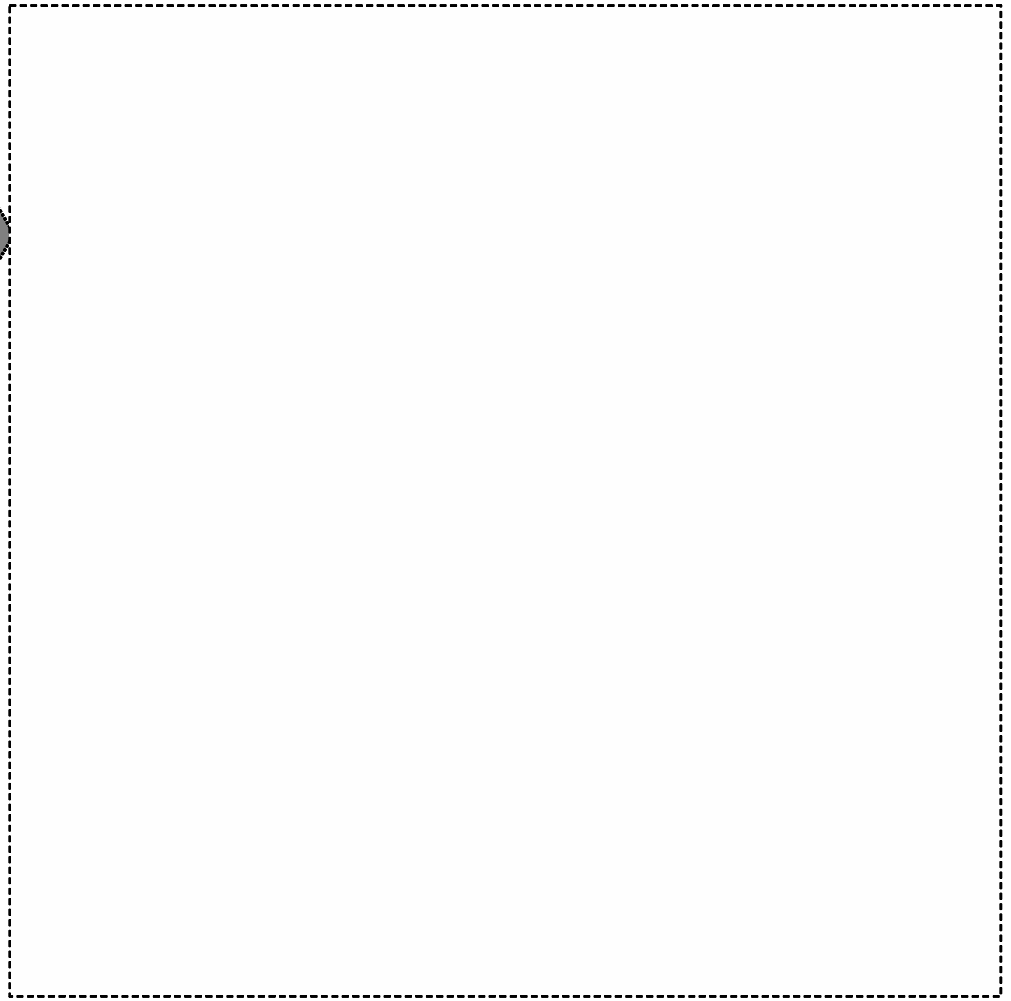

Source table may be a table expression consisting of joins

```
* C1 -> employee_id, C2 -> salary
```

External Tables

- **Performance**
  - Faster load and unload of massive data sets.
    - Internal tests shows 2x for unload and 3x for load of large tables over HPL.

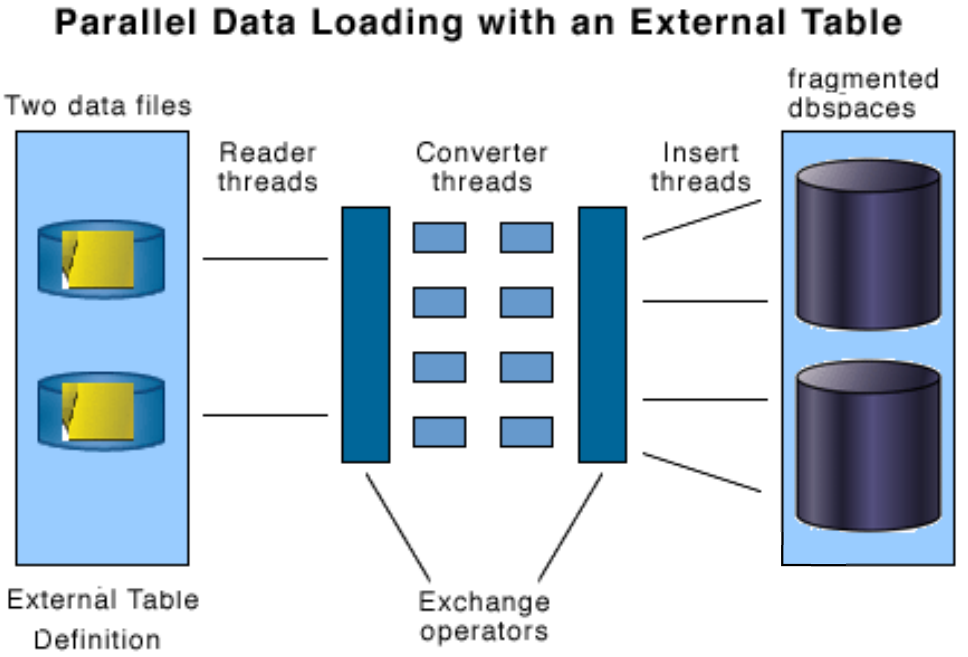

- **Ease of Use**
  - Syntactically, an external table can be used in an SQL statement in place of a regular table.
    - No coding needed compared to use of VTI.
  - External table can be used in Stored procedure for load and unload
  - No need of DBA privilege to do Load/Unload using external tables
- **Most requested feature for XPS users**

Optimizer directives in ANSI-compliant joined queries

- Extends support in ANSI/ISO joined queries to the following classes of optimizer directives:

Class	Optimizer Directives
Access-method	FULL, AVOID_FULL, INDEX, AVOID_INDEX, INDEX_SJ, AVOID_INDEX_SJ
Explain-mode	EXPLAIN, AVOID_EXECUTE
Optimization-Goal	ALL_ROWS, FIRST_ROWS
Join-order	ORDERED

Attach/Detach Fragment Optimization

- Currently, an exclusive lock on the base table is required to add/drop a fragment
- Optimization is for the system to force users off by aborting transactions
- Connections are still intact
- Lock mode wait can be set to wait certain amount of time before “Alter Fragment” statement
- An ALTER FRAGMENT operation can now force out transactions to get exclusive access to tables
- Allows programmatic attaching/detaching fragments

Enhancements to distributed queries

- Extends support for:
  - UDRs in cross-database and cross-server distributed operations
  - External routines written in the C or Java languages.
  - Additional data types that can be used in cross-server distributed queries to built-in non-opaque SQL data types
    - BOOLEAN
    - LVARCHAR
    - DISTINCT of non-opaque built-in types
    - DISTINCT of BOOLEAN
    - DISTINCT of LVARCHAR
    - DISTINCT of the DISTINCT types that are listed above

Named parameters in a JDBC Callable Statement

- Adds the convenience of identifying parameters by name instead of by ordinal position
- For unique stored procedure
  - You can omit parameters that have default values
  - You can enter the parameters in any order
- Useful for calling stored procedures that have many arguments

Multiple Triggers Enhancements

- Several new trigger features and enhancements added
- Expands the syntax and functionality of triggers on tables and on views
  - Can define multiple INSERT, DELETE, UPDATE, and SELECT triggers on a table
  - Can define multiple INSTEAD OF triggers for the view
- Allows more flexibility and performance improvements

Sysdbopen()/Sysdbclose() UDRs

- Executed after a successful open/close db or connect/ disconnect
  - Set execution parameters
  - Send alerts
  - Begin auditing
- Different versions per user and for PUBLIC
- Does not execute on remote UDRs or distributed DMLs
- Can be defined by DBA and user Informix

Dynamic SQL in SPL

- In static SQL statement all the information is known at compile time

```
CREATE PROCEDURE foo_static (vpkid INT,vname VARCHAR(32))
  RETURNS INT;
  UPDATE mytab SET myname = vname WHERE pkid = vpkid;
END PROCEDURE;
```

- Dynamic SQL allows an client program to build an SQL statement at runtime

- Contents of the Dynamic SQL statement is determined by user input

```
CREATE PROCEDURE foo_dynamic(vpkid INT,vname VARCHAR(32),
  tabname VARCHAR(128)) RETURNS INT;
  DEFINE query LVARCHAR(512);
  LET query = "UPDATE " || tabname || "SET myname = '" || vname
  || "' WHERE pkid = " || vpkid;
  EXECUTE IMMEDIATE query;
  RETURN SQLCODE;
END PROCEDURE;
EXECUTE PROCEDURE foo_dynamic (123, "John Doe", "customer");
```

Dynamic SQL in SPL

- Using procedures with dynamic statements within other DML statements is disallowed
- Cursor support in dynamic SQL Only string values or variables are allowed. Statement length is limited to 32K
- Following Cursory statements are allowed
  - `PREPARE, DECLARE, OPEN, FETCH, CLOSE, FREE`
- Following non cursory statements are allowed
  - `EXECUTE IMMEDIATE`
- Use the new SQLCODE function to check for any error. It should return zero when the statement is successful

SPL Enhancements

- GOTO and LOOP statements provide greater flexibility
- Facilitate easier migration from other database
- The GOTO label statement
  - Can unconditionally exit from a loop
  - Transfer control to the specified statement block
- The LOOP statement
  - Executes a statement block for an unspecified number of iterations. Used within a
    - WHILE condition LOOP
    - FOR condition LOOP
    - Independent of any FOR or WHILE
  - LOOP statements can also be nested
  - EXIT and CONTINUE options in IF statement for control

SPL - PL/SQL Compatibility

- Goal: Facilitate migration from Oracle to IDS.
- New loop syntax:
  - `LOOP…END LOOP`
  - `WHILE LOOP… END LOOP`
  - `FOR LOOP… END LOOP`
- Loop label:
  - `<<lp_label>> LOOP… END LOOP <<lp_label>>;`
  - `<<lp_label>> WHILE LOOP… END LOOP <<lp_label>>`
  - `<<lp_label>>  FOR LOOP… END LOOP <<lp_label>>`
- Loop exit statement support:
  - `EXIT WHEN <expr>;`
  - `EXIT << lp_label>> WHEN <expr>;`
- GOTO and label syntax support:
  - `GOTO <<label>>`
  - `<< lp_label>>`

XML Publishing and Manipulation

- IDS supports XML publishing functions
- Transforms query results to XML for use in XML applications
- Can be used in heterogeneous database environment
- Supports built-in functions like XPATH
- Extracting XML portions or values
  - `extract()/extractxmlclob()` → extract an element and its children
  - `extractvalue()/extractxmlclobvalue()` → extract the value of one element
- Condition on XML
  - `existsnode()` → test for the existence of a node

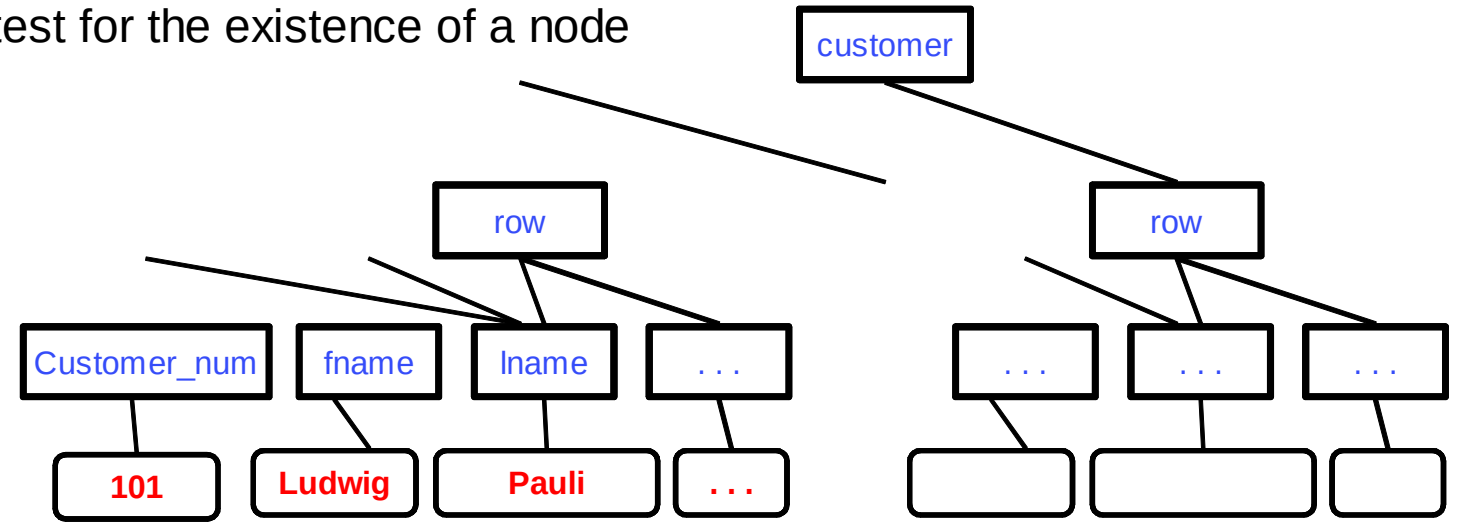

XSLT Support

- XSLT: eXtensible Stylesheet Language Transformations
- XSL and XSLT is part of W3C standard
- XSLT is used to transform XML documents
  - XML to XML (transform to confirm different schema/standard)
  - XML to HTML
  - XML to PDF
- Provides Transform Functions based on XSLT
- XSLT helps you to transform to customize display
- XSLT helps you to transform from one XML schema into another.
- XSLT enables loose integration thru transformation
- XSLT enables same data to be publishable to multiple targets with unique requirements

Basic Text Search (BTS) Index

- Available as a bundled free server extension
- Allows search of words and phrases
- Fast retrieval and automatic indexing of text data
- Search strategies
  - Single and multiple character wildcard searches
  - fuzzy and proximity searches
  - AND, OR and NOT Boolean operations
- Uses open source CLucene text search package
- Example
  - Fuzzy search: type soarness matches soreness

```
SELECT cat_advert, score FROM catalog
WHERE bts_contains(cat_advert, 'soarness~ and
classic', score # real);
```

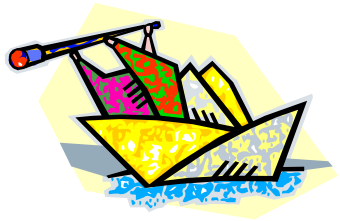

BTS XML Attribute Indexing

- Customizable stopword lists
- Flexibility of choosing xml tags for index creation
- Allows restricted XML tags searching
- “XPATH" queries supported in search
- Optional namespace use in searches
- Flexibility of ignoring XML Markup/tags in sentences
- Additional function returns unique tags in XML documents
- Example
  - XML search using XPATH

```
SELECT xml FROM xmlcustomer
WHERE bts_contains(xml,'/customer_record/row/
city:"Denver"');
```

Hierarchical data type : The Node Type

- "Hierarchically" aware type: Node
- Pre-process the hierarchical relationships
  - `ex: Chapter 11,section 7,para 3: 11.7.3`
- Add new way to relate objects to one another
  - `IsAncestor()`
  - `IsChild()`
  - `IsDescendant()`
  - `IsParent()`
- example:
  - Geo hierarchy: Country > State > Metro > City

    What policy applies to the Hyatt in Denver? State policy, City Resource

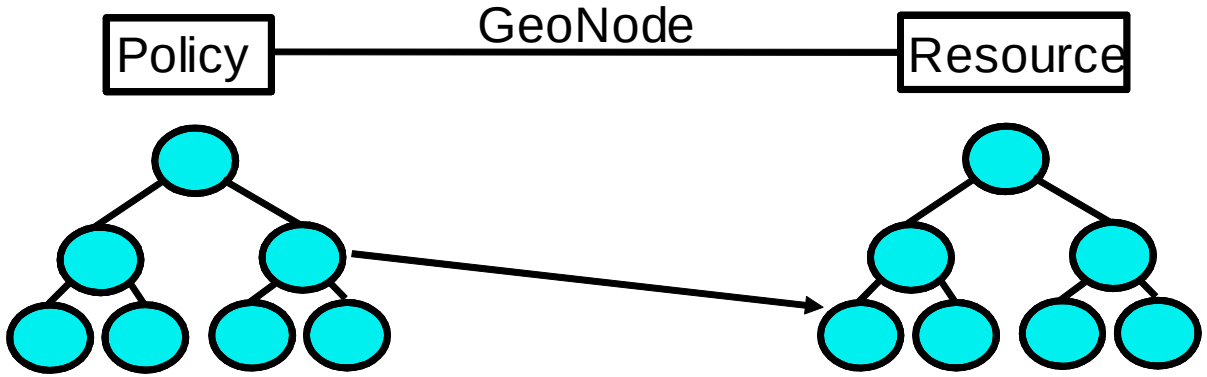

Retrieving data with Hierarchical queries

- Retrieves relational hierarchal data with parent-child dependencies
- Implement using “START WITH ... CONNECT BY” directive in SELECT clause
- Syntax:

- Comparison with Node Server Extension

Node Server Extension	CONNECT BY
Depth for hierarchy is limited to 16	Depth is (almost) unlimited (as much as fragmented temp table can hold)
Queries use normal relational operators to query hierarchy	Queries use recursive operators to query hierarchy

Indexable binary data types

- Two new data types
  - BINARYVAR
    - Variable-length opaque type with a max length of 255 bytes
  - BINARY18
    - Same as BINARYVAR data type except it holds a fixed value of 18 bytes
- Allows storing data in binary-encoded strings that can be indexed for quick retrieval
- Supports string manipulation functions
- Support bitwise logical AND, OR, XOR, and NOT operations
- Example
  - UUID `2ac07282-184e-4103-ad2a-ad8abdab8f0b` would take 36 bytes if stored as string
  - The following code inserts the above UUID string as binary-encoded string :

```
CREATE TABLE mytab (srlno INTEGER, bindata BINARY18)
INSERT INTO mytab VALUES (1,'2ac07282184e4103ad2aad8abdab8f0b')
```

Web Feature Service for Geospatial data

- Implemented as a Server Extension
- Web Feature Service (WFS) Server Extension module implements an Open Geospatial Consortium(R) (OGC) Web Feature
- OGC WFS in IDS to acts as a presentation layer for the Spatial and Geodetic Server Extension modules
- The OGC WFS interface allows requests for geographical features across the web using platform-independent calls
- The XML-based Geography Markup Language (GML) is used as the encoding for transporting the geographic features

IBM Data Server Client

- **Supports DB2, Derby, IDS 11.x, Unidata/Universe**
  - Uses a DRDA interface
- **APIs covered:**
  - Java, PHP, Perl, Ruby/Rails, CLI, .NET
  - Visual Studio Add-in
- **Some limitations in IDS support (extended types)**

See: http://www.ibm.com/software/data/db2/ad

Expanded Application Development Technologies

- Developer communities
  - C/C++
  - Java (JDBC / SQLJ)
  - .NET (C#, VB .NET)
  - EGL
  - Open Source
    - PHP/Zend FW
    - Ruby/Rails
    - Perl
    - Python/Django

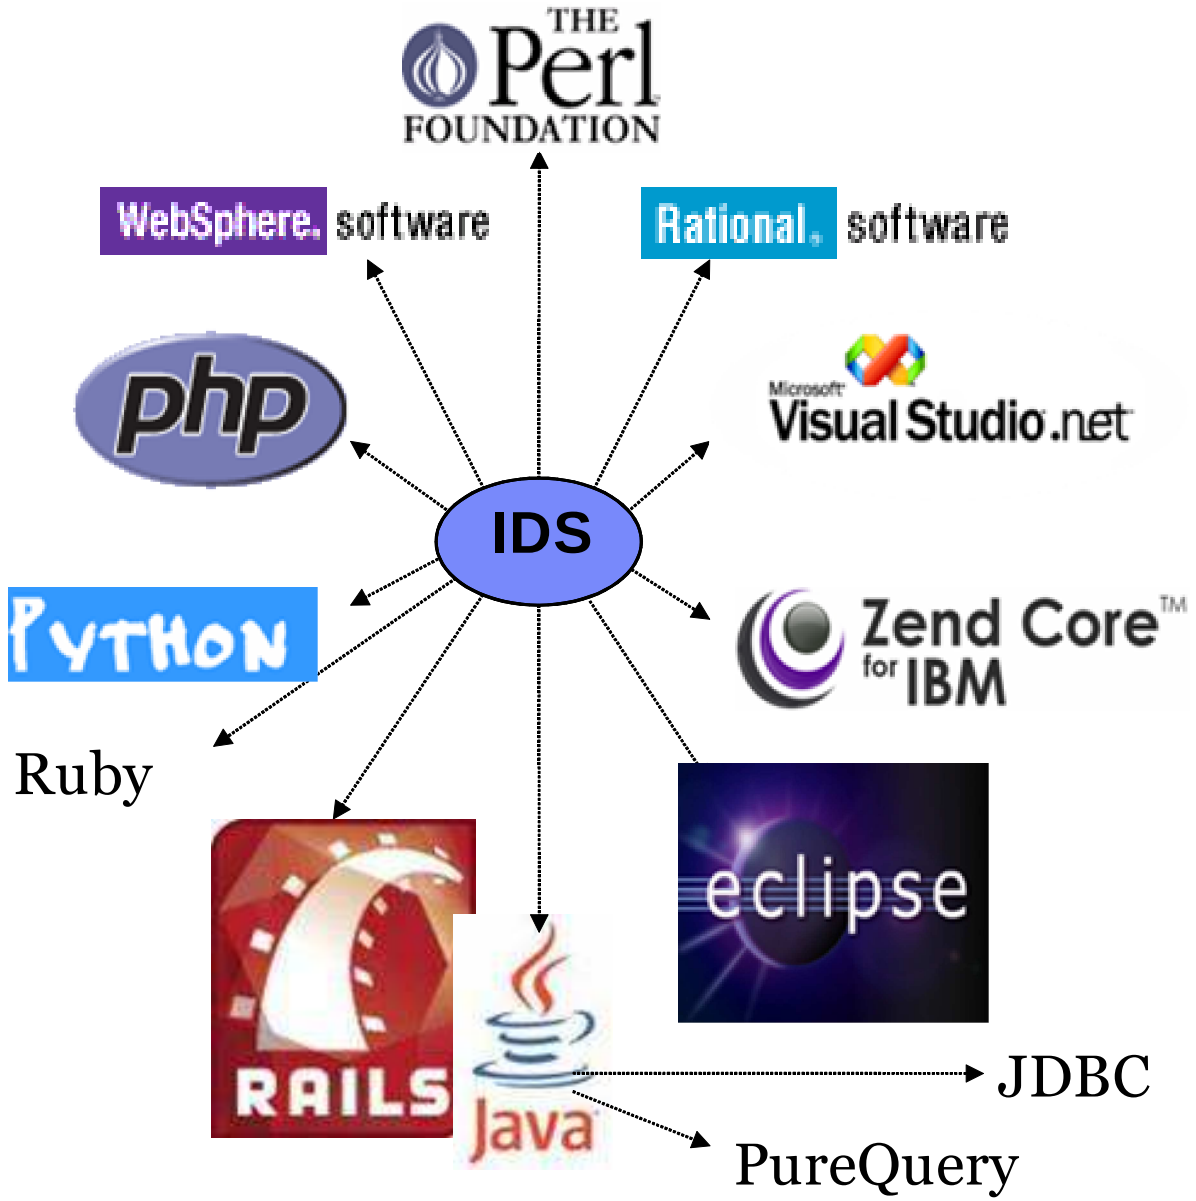

Technical Overview

- Administration and Usability
- Business Continuity
- Performance Enhancements
- Security
- Application Development
- **Integrated solutions and related offers**

86

Integrated solutions and related offers

- **IDS Virtualization**
- **IDS AWS EC2 support**
- **Visual Studio Support**
- **Java Common Client (JCC) and JDBC Support**
- **Support for common clients**
- **Data Studio**

87

IDS Virtualization

Optimize your IT investments, infrastructure and resources

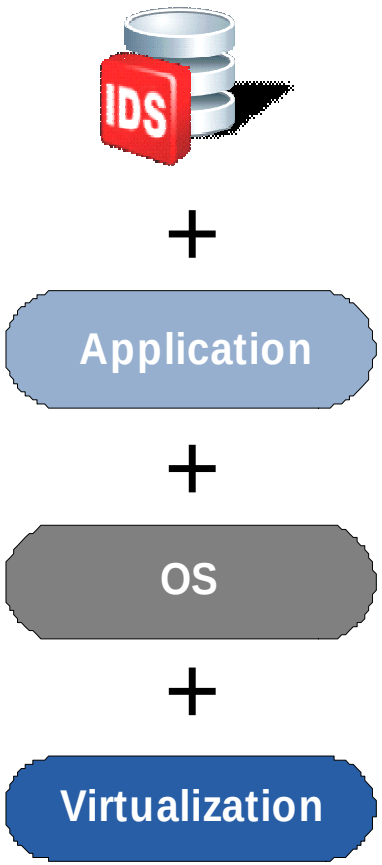

Investments

- Reduce power, cooling and real-estate expenses
- Reduce hardware expenses through consolidation and improved utilization

Infrastructure

- Simplify server configuration and deployment
- Enhance manageability and responsiveness
- Improve server utilization and control
- Scale compute resources to meet demand

Resources

- Leverage IDS high availability features
- Reduce software and maintenance complexity

Available now:

- IDS Virtual Appliance: A SUSE Linux Enterprise Server 10 SP2 virtual machine with Pre configured online IDS instance
- Developer Edition virtual appliance edition for exploration and development
- Other editions through ISVs and OEM partners
- Repackaging Guide
- IDS instances in Amazon Cloud

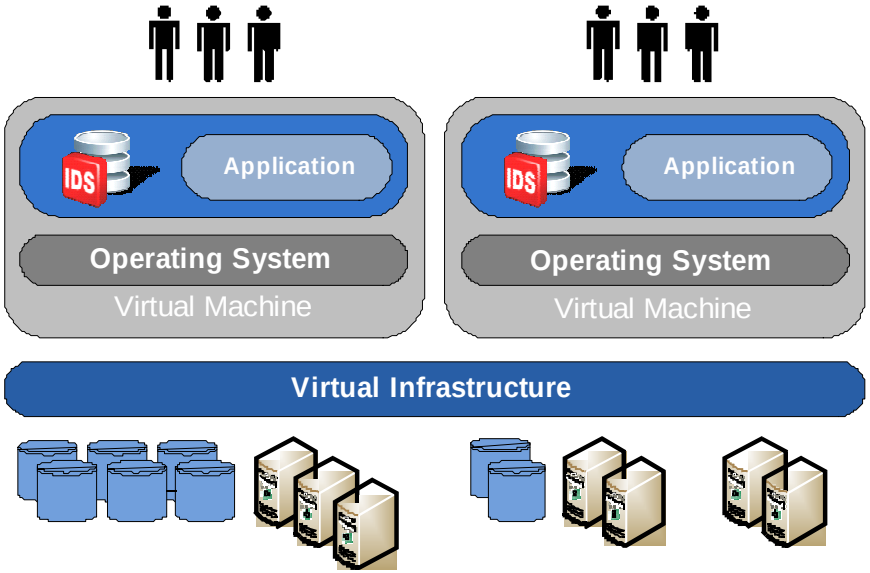

IDS AWS EC2 support

- AWS EC2 support : Amazon Web Services Elastic Compute Cloud
- Delivers production applications on an pay hourly as per use
- A key component in “Cloud Computing”
- IDS is certified for EC2
- IDS virtual appliance can be used with EC2

Visual Studio Support

- IDS supports the IBM Database Add-ins for Visual Studio
- Enables IDS application developers to use the Visual Studio functionalities
- Reduces development time
- Develop and integrate .NET applications for a wide range of IDS server families

Optimized for DB2 Information Management Software

Microsoft Visual Studio .net

Java Common Client (JCC) and JDBC Support

- Single JDBC driver for ALL IBM data servers
- Java Common Client (JCC) support IDS IDS 11
  - Support for Informix connection URLs, environment variables, data source
  - Support for Informix basic data types
  - Support for all database modes
  - Connection using DRDA protocol
  - Supports JDBC 4.0 Specification
- Direct IDS access using JDBC or JCC via **Groovy** scripting

Common Client for DRDA Enhancements

- IDS supports DRDA, the communications protocol used by DB2
- Implements IBM common client APIs to communicate with IDS as well as DB2
- Gives flexibility to customers to deploy solutions on any IBM data server they prefer to use
- Other enhancements:
  - Support MACH11 features
    - Support Connection Manager
    - Automatic client rerouting
    - Load balancing at transaction boundary
  - Support SSL for client-server communication [JCC]
  - Support for database names > 18 characters
  - Retrieval of ISAM error codes
  - Support statement batching
  - Support for CLIENT WORKSTATION Info APIs

IBM Data Studio 2.1

- Eclipse-based Integrated Development Environment (IDE)
- Available free version with limited capabilities
  - All capabilities included for a trial period
- Major features:
  - Connect to data sources and browse data objects and its properties
  - Editors and wizards to create and alter data objects
  - Modify privileges of data objects
  - Diagrams to visualize relationship between objects
  - Create web services based on SQL statements or procedure calls
  - Speed up database development with PureQuery

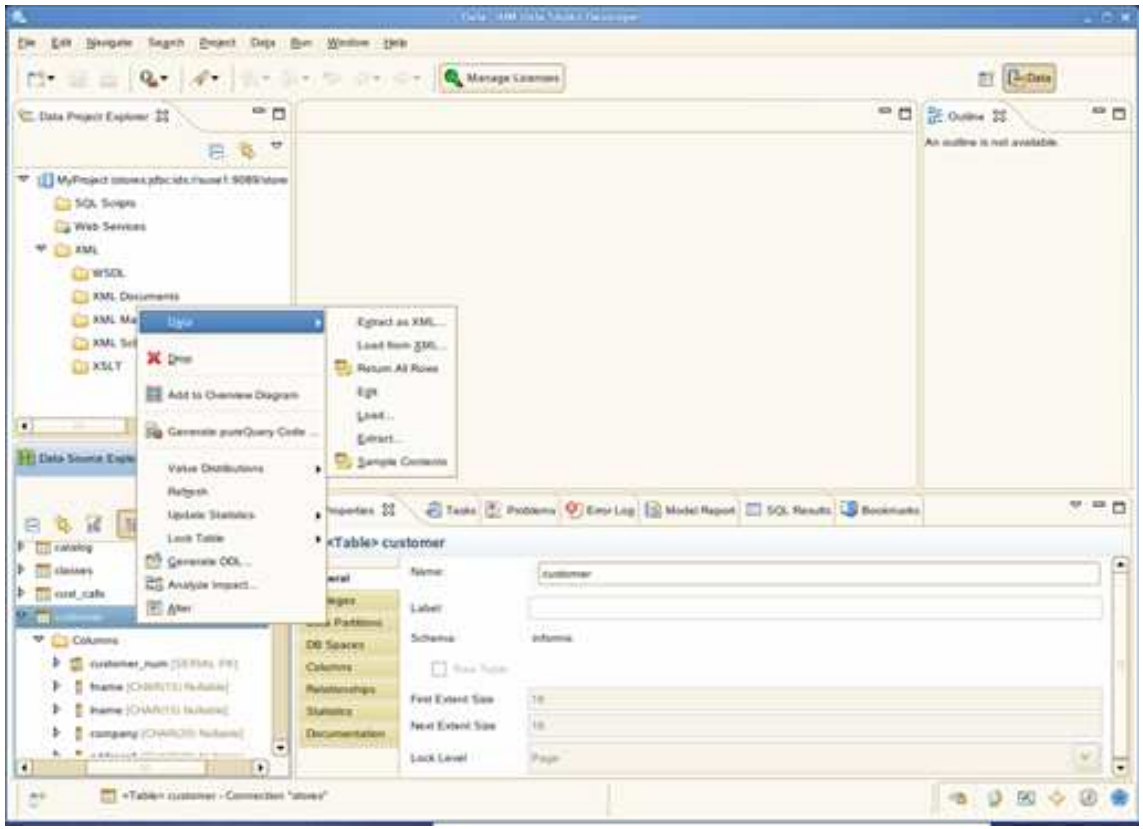

Value Additions

And

Backup Slides

The Following Slides Are Hidden.

Use These As Supplement To The Above Slides

Macintosh OS Support

- Runs on Leopard
- Both client and server are available
- Intel 64 supported
- Strong interest in publishing & education, gaming market
- Future functionality coming soon
  - JDBC and JCC
  - Geodetic support
  - Excalibur Text
  - Support for FASTPOLL

IDS – Winning Solution in All Industries

Public Sector

- Civic Administration
- Customs
- Education
- Health
- Judicial
- Law Enforcement
- Military
- National Security
- Tax Administration

Electronics

- Integrated Applications:
- Text messaging
- Pay-per-View
- Entry Access badges
- Cell phone access
- RFID

Telecommunications

95% of all telecommunications delivery providers

Leisure

Hotel and
Airline Reservations
Online Gaming

Blazing-fast OLTP
Nearly Hands-free administration
Legendary Reliability

Retail

IDS Powers **8 of the top U.S.** largest retailers

Banking and Finance

Every credit & bankcard transaction authorization in US goes through IDS

Extending the Value…

Strong IDS Growth

IDS is growing faster than the data server market

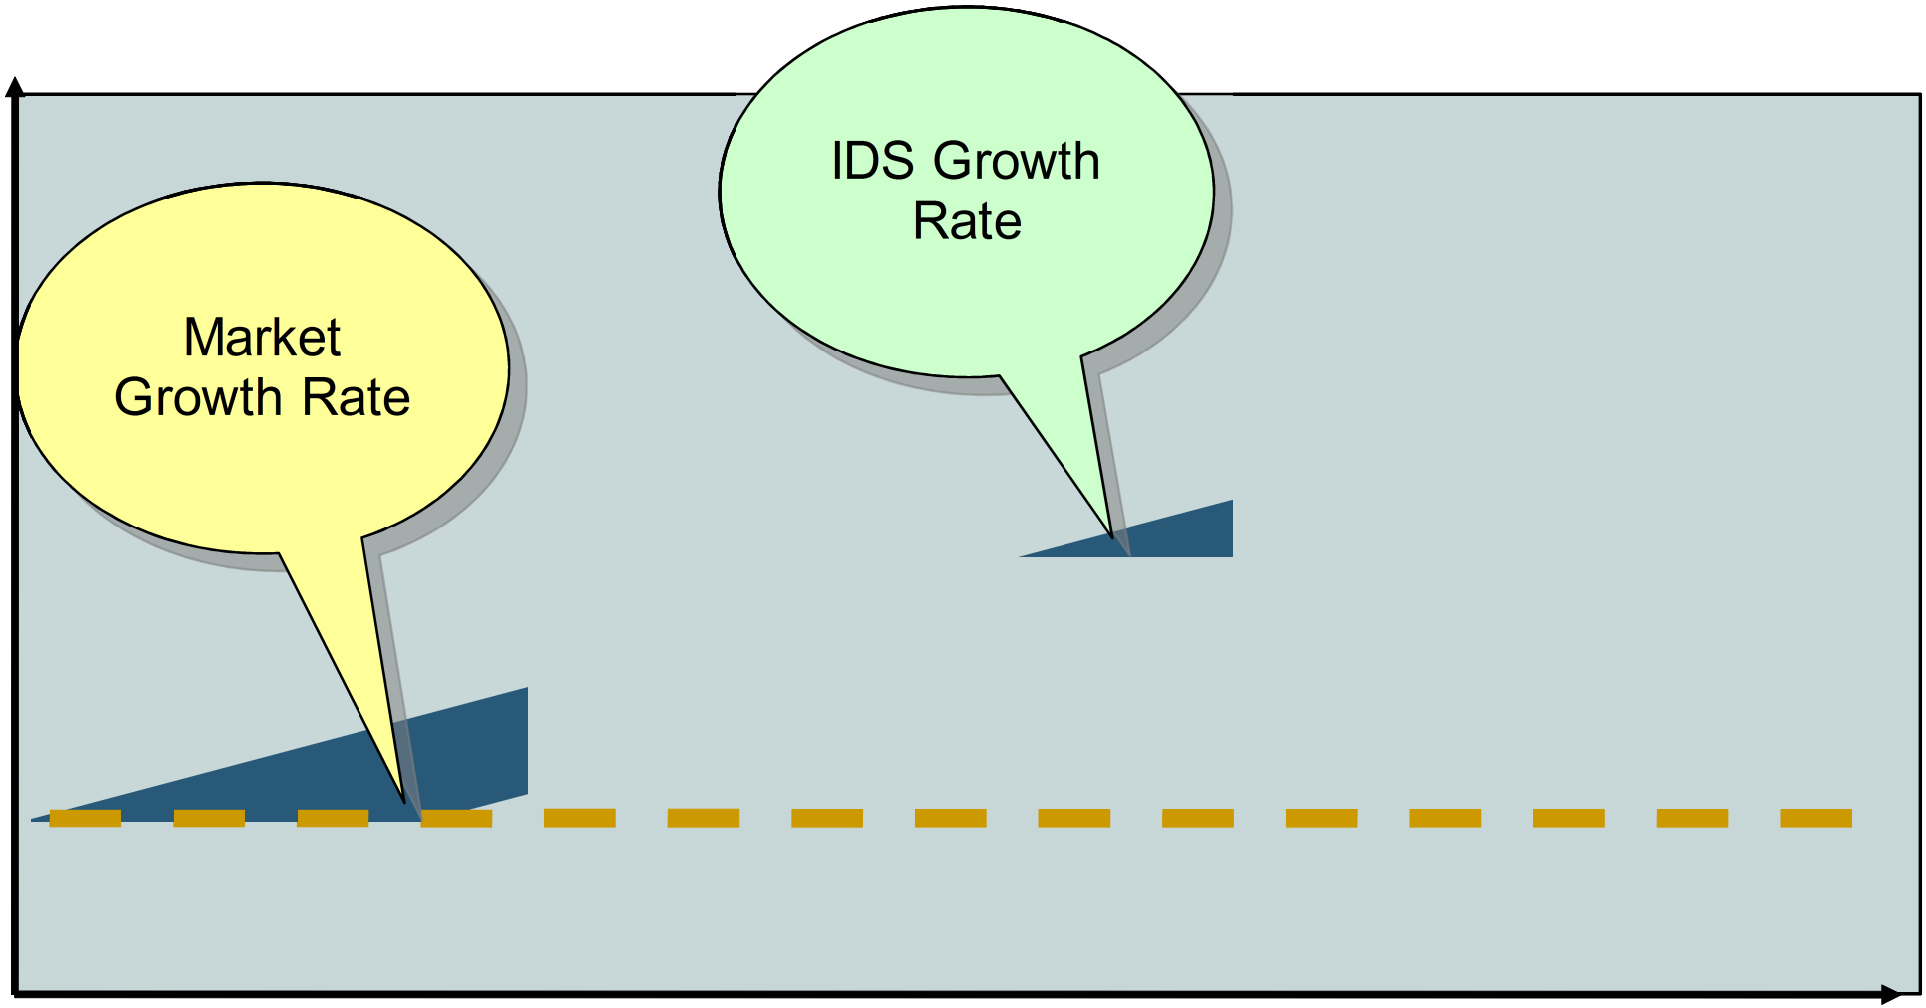

IDS Positioning

Capitalize on market strengths of Informix Dynamic Server (IDS)

On line Transaction Processing

IDS 11 new features increase the performance of IDS for OLTP.

Integrated Solutions

IDS is IBM's leading data server for industrial-strength integrated and/or embedded solutions.

Location Based Services

RFID, Spatial, Geo-Spatial, Time-Series Technology is Special in IDS

Small & Medium Businesses

Low DBA requirements and small Footprint are strong selling points.

Key Industries

IDS key industries including Telco/Communications, Retail/Distribution, Government/ Public Sector, Gaming/Entertainment, Travel/Hospitality, Banking/Financial, Electronics, Health Care, Energy & Manufacturing

Analysts are saying… IDC

"IBM remains the second largest RDBMS vendor mainly due to its substantial mainframe DB2 sales. Despite various "experts" prognosticating to the contrary, IBM has not only maintained this business but steadily grown it. What is less well known among RDBMS market observers is that much of IBM's growth in this market may be attributed to the success of Informix Dynamic Server (IDS). This RDBMS, which many had thought just a stepping stone to IBM's acquiring customers for the distributed system version of DB2, has turned out be quite a growth engine, despite very little promotion by IBM itself. IDS customers, the ISV community that offers products for it, and the many consultants who support it, have been "under the radar" of evangelists for the product, not just ensuring a stable customer base but winning over new customers as well."

Carl W. Olofson

IDC - Worldwide RDBMS 2006 Vendor Analysis: Top 10 Vendor License Revenue by Operating Environment and 2007 Year in Review, Dec 2007

http://www.idc.com/getdoc.jsp?containerId=209965

IDC Analyst White Paper on IDS

URL for customer registration and download

http://www-306.ibm.com/software/data/information/idc-ids.html

A visionary data server for forward-thinking businesses

IBM Informix Dynamic Server (IDS)

- **Continuous Availability** functionality and the IDS legendary reliability reduces the cost of downtime
- **Extreme OLTP** performance and efficient use of resources optimize IT investments
- **Near zero administration** simplifies deployment and management of data to lower cost

IDS evolves to keep up with your Information Challenges

Information needs for your business are changing rapidly…

- Regulations and industry standards are driving an increase in data volume
- The competitive landscape requires increased responsiveness, leading to the adoption of new business applications
  and SOA
- The need for insight and integration requires support for an increasing diversity of data formats
- Growing infrastructures increase costs of personnel and hardware

Manage Your Data, Not Your Database

Access your IDS data no matter where it resides

Availability - Business Impact

Loss of data access can stop business from running

- Loss of Customer Confidence
- Loss of Employee Productivity
- Loss of Company or Share Value

Availability	Downtime Minute per Year
99.999%	5 minutes
99.99%	50 minutes
99.9%	8 hours, 20 minutes
99%	3 days, 11 hours, 18 minutes
95%	18 days, 6 hours
90%	34 d… hours, 17 m…
85%	54 days, 18 h…

- Loss of Market Share and Revenue
- Penalties, Fines and Regulatory Fee

Application Segment	Average Cost of Downtime/Hour
Shipping - Distribution	$28,000 per hour
Tele-Ticket Sales	$69,000 per hour
Airline Reservations	$89,000 per hour
Home Shopping	$113,000 per hour
Pay Per View - Television	$150,000 per hour
Credit Card Sales	$2,650,000 per hour
Financial Market	$6,450,000 per hour

Source: Giga Group 2004

At 99% Uptime, a financial market would lose around $540 million per year

Lower cost of infrastructure for customers

React immediately to help ensure availability, optimal performance and lower cost

- Reliable and secure data to meet compliance needs
- Integrated automated failover and backups
- Automated Administration disappears within an application
- GUI development and administration tools

"IDS is the best choice"

"*We looked at a lot of database server vendors, and we found that IDS had the best performance, the most reliability and even the best cost.*"

- Andrew Ford, Senior Database Administrator, NetworkIP

Easy, Scalable Administration

- **Large retailer runs multiple instances in nearly every store**
  - Separate Retail, Pharmacy and other systems
  - This amounts to > 10,000 production installations worldwide
  - 24 X 7 X 365
  - Remote management using only 8 DBAs
- **Customer reports 99.998% uptime (downtime 5 minutes/year)**
- **Distributed queries and proprietary server extensions used**

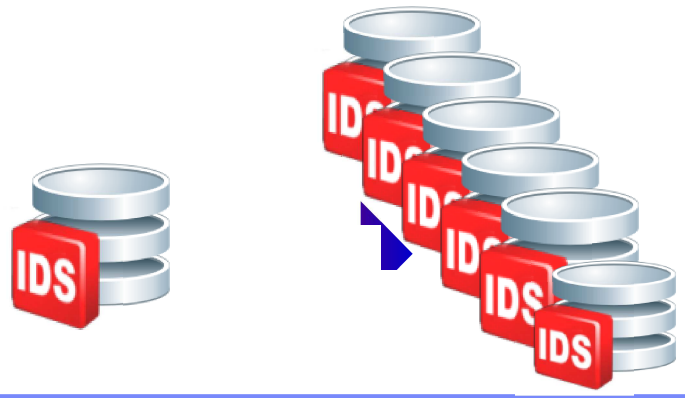

IDS 11 data server: The Innovation Continues

IDS 11

- **Resilient**
- **Agile**
- **Invisible**

- **Continuous Availability and Scalability**
  - Legendary reliability and highly available
  - Blazing fast OLTP processing
  - Scales for rapid response to changing needs
- **Confident compliance**
  - Enhances data governance and security
- **Enhanced development tools and platforms**
  - Easily integrated into distributed applications for nearly invisible, hands free administration
  - Rapid application development with IBM Data Studio

How to move to IDS 11

Current Platform	Step 1	Step 2
7.3x, 9.21, 9.30, 9.40, 10.x	**IDS 11**	**N/A**
9.1x	**9.30**	**IDS 11**
7.22, 7.23	**7.31**	**IDS 11**
Online 5.1x	**7.31**	**IDS 11**

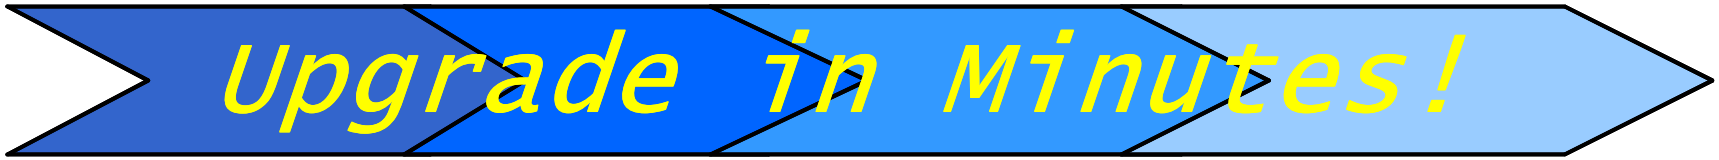

- **Recompilation of applications is not necessary in most cases**
- **If customers are changing operating systems (say HP-UX to Linux) or hardware (say SUN SPARC to IBM pSeries) they must use dbexport/dbimport**
- **Customers upgrading from Standard Engine must use dbexport/dbimport**
- **Reversion, if necessary, but highly unlikely, is fully automated**

IBM Informix Dynamic Server (IDS)

A Strategic Data Server in the Information Management Portfolio

Financial | Gaming | Government | Healthcare | Retail | Telco

Double Digit Growth!

Extensive, Active, Partner Network

Loyal Customer Base

Enter-prise

Mid-Mkt

SMB

Integrated Application Development Environments and APIs

Wide Breadth Of IBM and Partner Solutions

Java / JDBC | .Net | PHP / Ruby Rails | ODBC / Others

Open, Common Standard APIs

Robust SQL (SQLI and DRDA)

Rich Connectivity to Applications, Data, and Content

Informix Dynamic Server (IDS)

Agile | Invisible | Resilient

High Performance OLTP with Legendary Reliability

Self Tuning, Memory Management with Low Cost of Administration

Scalable, Multi-Threaded, Secure Parallel Processing Architecture

Gain Competitive Advantage with IDS

Productivity

Little or no Administration, lower TCO!

Business Continuity

Highly Available and Reliable Data Access

Easy Deployment

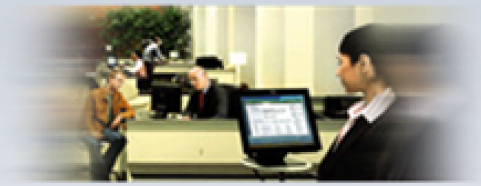

Very small footprint, customizable

Customer Satisfaction

"It Just Works!"

High Performance

Faster Transactions, Increased revenue

Privacy Compliance

Minimize risk, highly secure

Faster time-to-value

Minimal need for Training, Hand Holding …

IDS Simplifies your global business

- **Simply global**
  - Distributed data – anywhere – real time
- **Simply dependable**
  - Legendary reliability
- **Simply manageable**
  - Remotely support thousands of locations
  - Tailor your environment for push button deployment
- **Simply available**
  - Data in synch for your global business
  - Workload management to maximize resources
- **Simply scalable**
  - Dynamic scalability on demand
  - Scale up. Scale out. Scale around the world – your choice
- **Simply the most cost effective data management**
  - Minimal administration, low cost of ownership
  - Small footprint – lower cost of hardware – lower cost to the planet.

IDS simply is the best choice

IDS 11
Reference Books

- Embedded Informix Dynamic Server: A Handbook
- Extending Availability and replication
- Customizing IDS for your environment
- IDS 11: Advanced Functionality for Modern Business
- Security and Compliance Solutions for IBM Informix Dynamic Server
- Informix Dynamic Server 11: Extending Availability and Replication
- Many more…

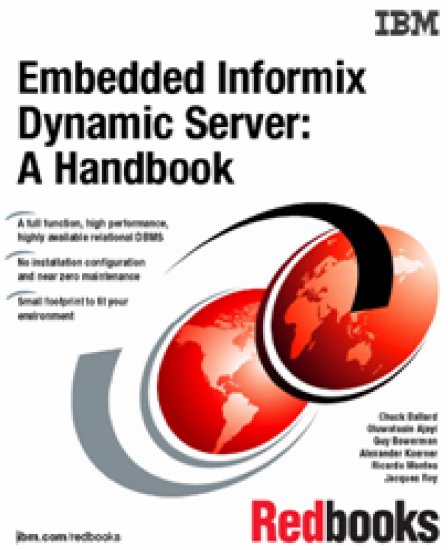

Other Resources

- **Product Family Homepage:** http://www.ibm.com/software/data/informix/
- **Platform Roadmap:** http://www.ibm.com/software/data/informix/pubs/roadmaps.html
- **Product Life Cycle (PLC) Roadmap:** http://www.ibm.com/software/data/informix/support/plc/
- **Developer's Homepage:** http://www.ibm.com/developerworks/db2/zones/informix/
- **Interoperability:** http://www.ibm.com/software/data/informix/ids/interop/

Cont'd …

Other Resources

- **Documentation/Manuals:** http://www.ibm.com/software/data/informix/pubs/library/
- **Product Webcasts:** http://www.ibm.com/software/data/informix/webcasts/
- **IDS Redbooks:** http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?Query=Informix&SearchMax=250&SearchOrder=4
- **Education/Training:** http://www.ibm.com/software/data/education/
- **IDS Success Stories :**
  - http://www.ibm.com/software/success/cssdb.nsf/softwareL2VW?OpenView&Count=30&RestrictToCategory=dmmain_InformixDynamicServer
  - http://www.ibm.com/press/us/en/pressrelease/23545.wss

Non-IBM Communities and Forums

- **IIUG (International Informix User Group)**
  - http://www.iiug.org
  - Very informative, lists ongoing, upcoming IDS events
  - Software repository of useful tools developed by IDS users
- The Informix Zone
  - http://informix-zone.com/
  - Very informative, resourceful site
  - Provides ongoing updates and information on IBM IDS
  - Owned and maintained by long term IDS enthusiast, expert and technologist
- IDS Developer Edition Wiki
  - http://www.informix-zone.com/idswiki/doku.php