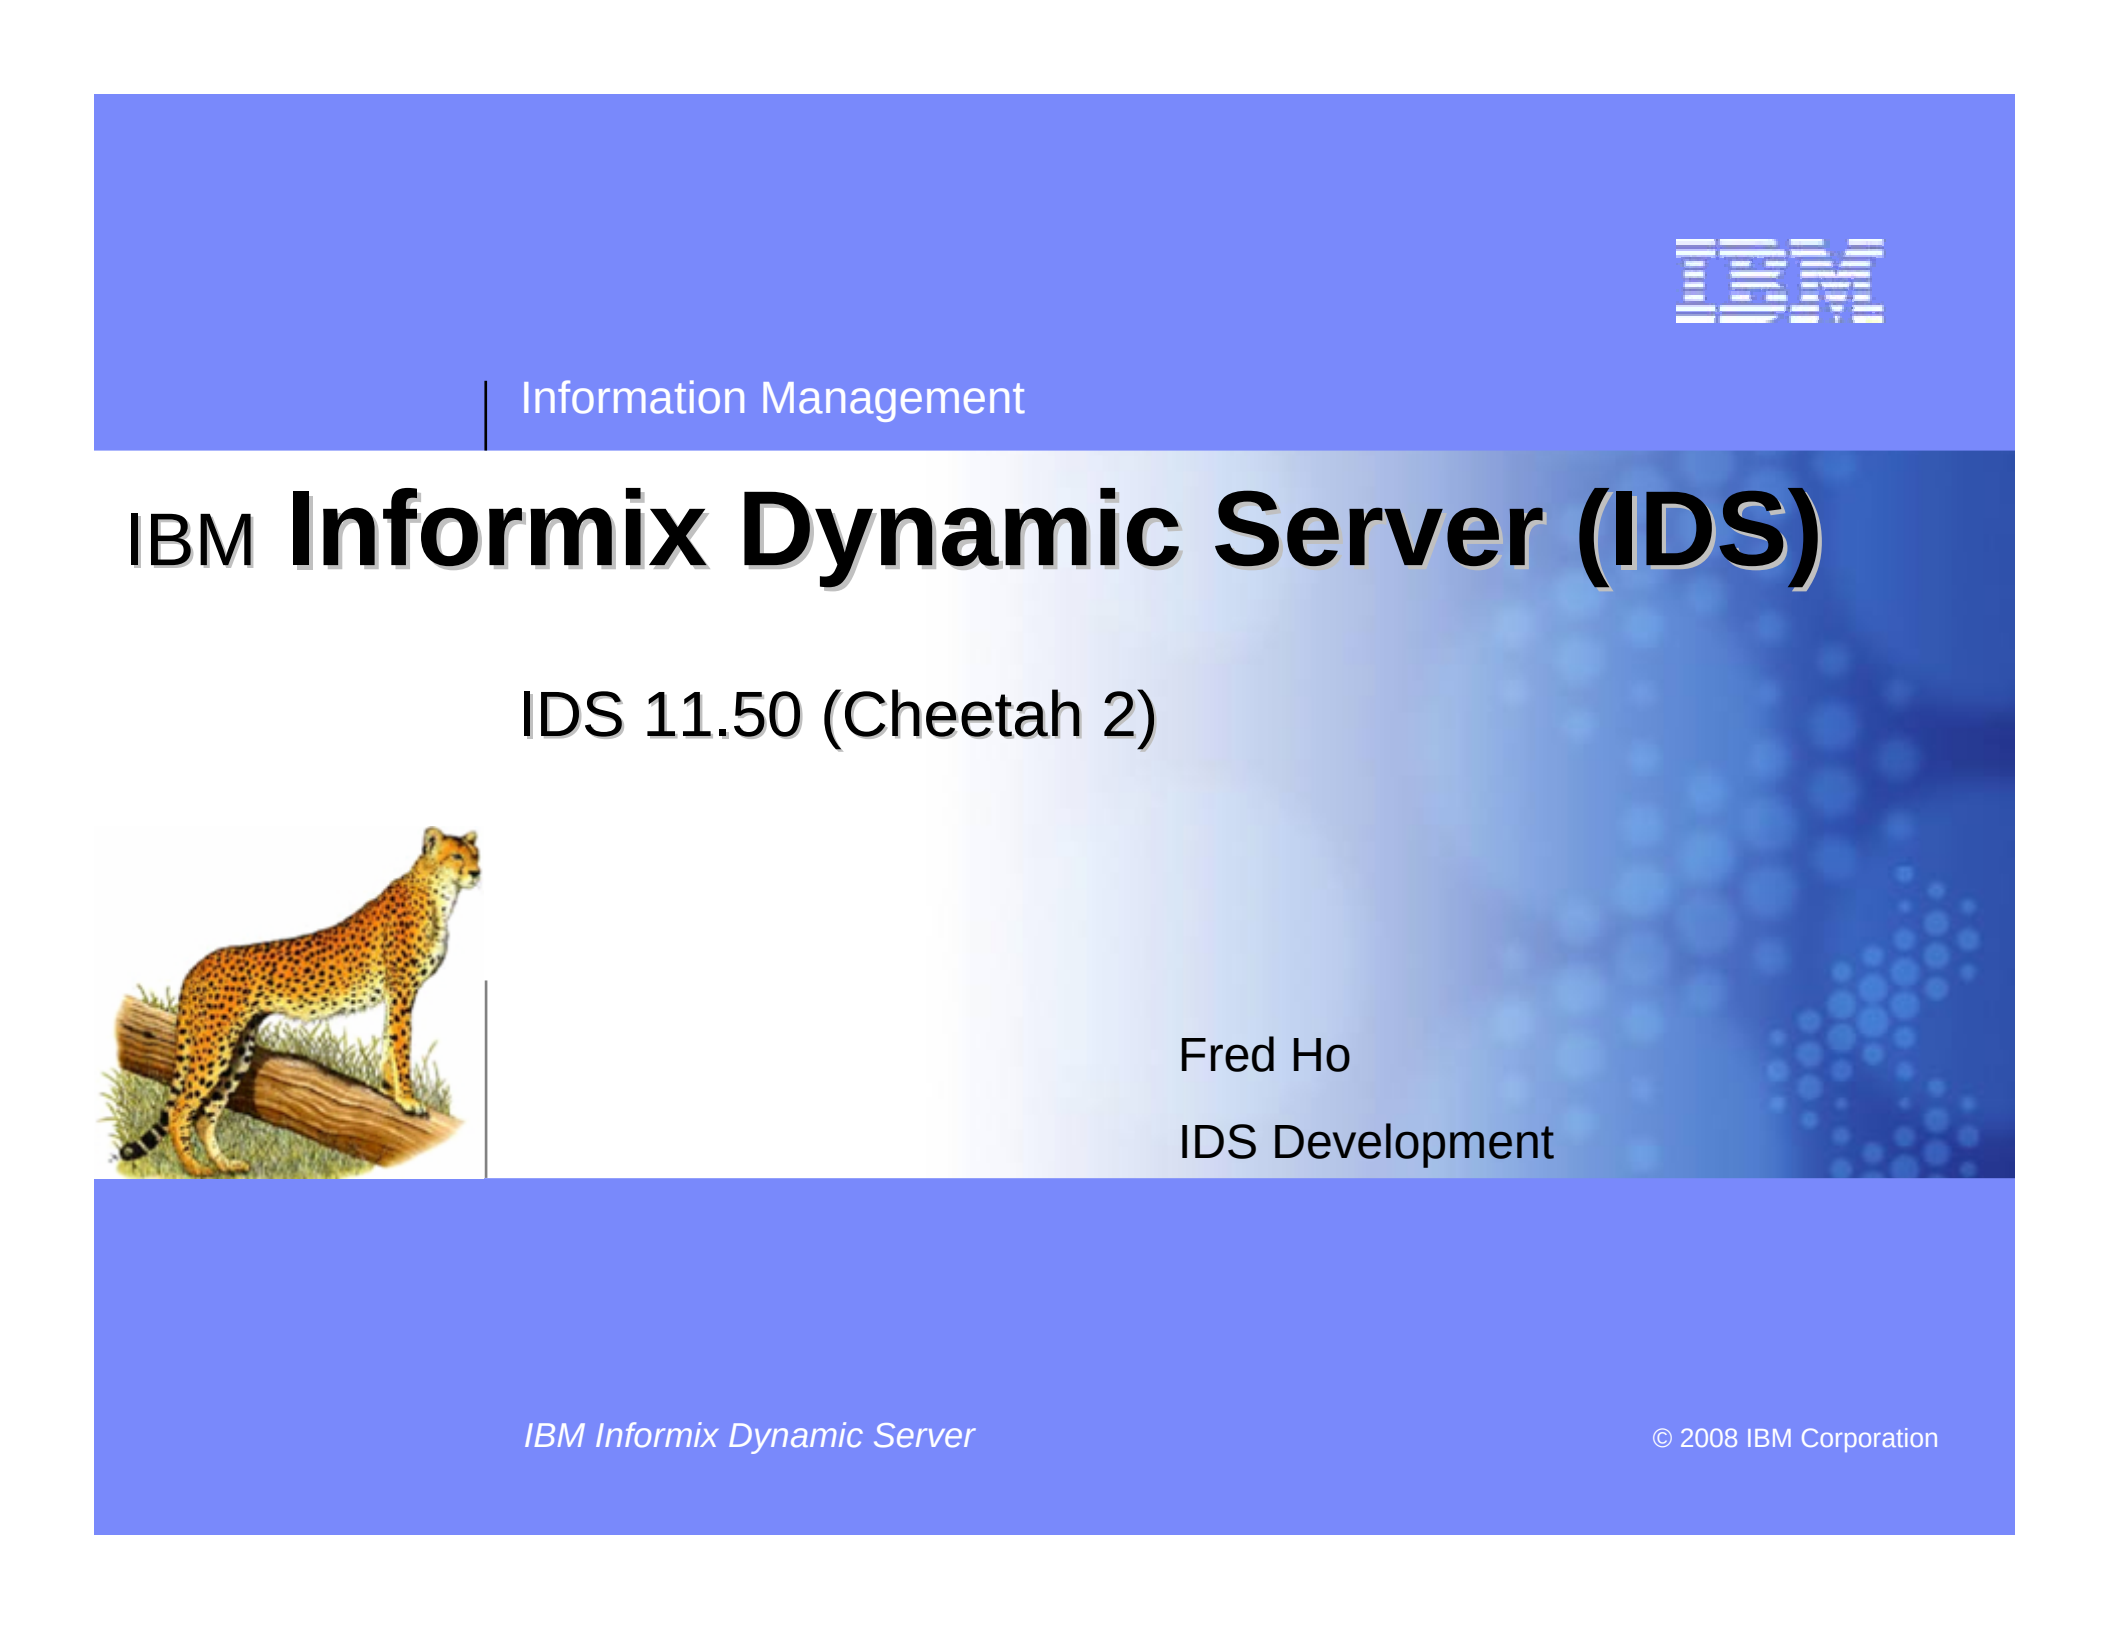

Information Management

IBM Informix Dynamic Server (IDS)

IDS 11.50 (Cheetah 2)

Fred Ho

IDS Development

Contents

- IDS Business Overview
- IDS 11 Availability
  - What is MACH-11?
  - HDR-SDS-RSS-ER
- IDS 11.5
  - Mach11 Phase 2
  - Auto Update Stats
  - Dynamic SQL in SPL
  - DRDA Support in IDS
  - Visual Explain
  - Configuration Wizard
  - Macintosh Port
  - XML Enhancements
  - BIGINT and BIGSERIAL SQL
  - Single Sign-On
- Beyond 11.5

IDS Helps Run the World
Over 100,000 Clients, Over 2,500 ISVs

"…we've gone years in some cases without having to take the database down…"

95% of all telecom service delivery providers

Informix
Resilient Agile Invisible
IBM Informix Dynamic Server 11

中国移动通信
CHINA MOBILE

" …enhances our system performance and is easy to use…"

20 of the top 25 U.S. Supermarkets

US Emergency 911 calls

All VISA transaction authorizations US

8 of the top 10 US Retailers

About: Some Facts About IDS

- 2006/2007 license double digit growth YTY
- Is a *significant* portion of IBM-IM L/U/W business
- Over 100K loyal customers world-wide
- This customer base is extremely loyal and has continued to invest with Informix
- But we are also acquiring new customers and partners WW
- IDS is the top selling IM product for SMB
- Over 40% of our business is ISV/OEM: Great technology for embedded solutions
- Popular on non-IBM platforms
  - Informix has superior support on HP, SUN, Linux, Windows platforms
- Strong commitment and investment to the future of IDS from top management
  - Launched IDS 11 in 2007 and 11.50 in 2008
  - Strong future and roadmap

IBM Data Servers

The industry's broadest range of capabilities optimized for every business need

Optimized for very high performance transactions requiring very low latency response

Optimized for rapid, cost effective development of solutions with variable and dynamic data records

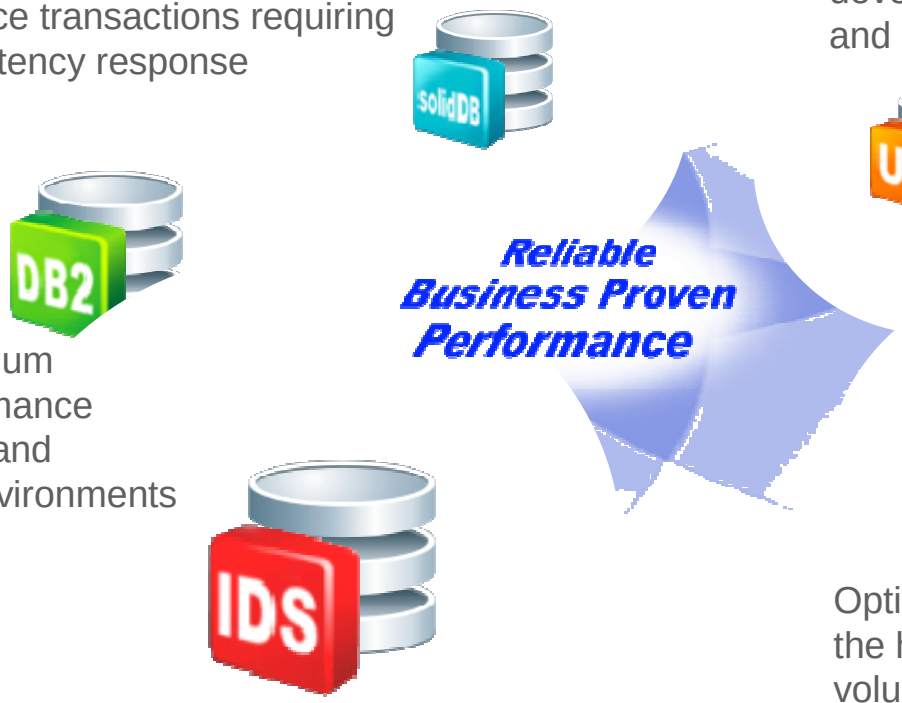

Optimized for maximum flexibility and performance in diverse workload and operating system environments

Optimized for solutions requiring the highest levels of transaction & data volume and performance

Optimized for near hands free administration for transaction intensive solutions and solution fleets

IDS Core Business Value

- Lowest cost of administration in the industry
  - ***Manage 2,000 IDS servers with 1 DBA***
- Highest performance **OLTP** technology for extreme workloads
  - ***Financial brokerage firm handles 100,000 fully committed transactions per second***
- Near zero administration
  - ***Remotely manage thousands of databases with few DBAs***
  - Inline version upgrades without taking the server down and no additional tuning
- Ideal for **embedded** environments – Applications & HW Devices
- Business continuity regardless of geographic location with **legendary reliability**
  - with **Continuous Availability Feature (CAF) – Shared Disk, HDR, Replication in One**
- Security
  - Achieve compliance with the latest regulations
  - Protect you and your customer's information
    (Column level encryption, encrypted communications, and label-based access control (LBAC))

And Industry Analysts Agree…

" …much of IBM's growth in the market may be attributed to the success of Informix Dynamic Server…"

IDC, December 2007*

"…With IDS 11, they have sent a message to IDS users that they are serious about the DBMS and that it's part of IBM's overall DBMS strategy…"

Quoted in Intelligent Enterprise, June 2007**

**IDC - Worldwide RDBMS 2006 Vendor Analysis: Top 10 Vendor License Revenue by Operating Environment and 2007 Year in Review, December 2007*
***As Quoted in Intelligent Enterprise by Doug Henschen Database Trends: Q&A with Donald Feinberg*
http://www.intelligententerprise.com/showArticle.jhtml?articleID=199903507&pgno=1

IDS Bundles

- **Express Edition**

  Self managing data server for mid-market businesses

- **Workgroup Edition**

  A low-maintenance OLTP data server for workgroup computing

- **Enterprise Edition**

  Continuous availability & blazing OLTP performance for the global enterprise

- **Developer Edition**

  Available free of charge, for development use only, with no time limits, and no support.

IDS Express Bundle

- Self managing data server for mid-market businesses
- Provides the power, function and reliability with simplicity in packaging, installation and deployment at a minimal investment cost
- Available on Linux, Windows (Intel 32-bit) and Mac OS X
- CPU memory restrictions: 2 CPUs / 4 GB RAM maximum

Available On

Vendor	Specification
Linux (Intel 32-bit)	RHEL, Asianux, Debian, SUSE SLES, Ubuntu LTS
Microsoft Windows (Intel 32-bit)	Windows XP, Windows 2003, Windows Vista
Apple (Intel EM64T)	Mac OS X (3Q08)

IDS Workgroup Bundle

- Low-maintenance OLTP data server for workgroup computing
- Useful for departments within large enterprises, midsized companies, ISVs for OEM use
- Includes all of the features of IDS Express plus features to handle high transaction volume
- CPU memory restrictions: 4 CPU, 8GB memory maximum

Available On

Vendor	Specification
IBM	AIX
HP	HP-UX on Itanium, HP-UX on PA-RISC
SUN	Solaris Sparc , x64
IBM POWER	RHEL, Asianux, Debian, SUSE SLES 10
IBM zSeries	RHEL 4, RHEL 5, SUSE SLES 10
Intel x86/x86_64	RHEL, Asianux, Debian, SUSE SLES, Ubuntu LTS
Microsoft	Windows XP, Windows 2003, Windows Vista
Apple (Intel EM64T)	Mac OS X (2Q08)

IDS Enterprise Bundle

- Create a failsafe, multi-site global availability plan while maximizing IT investment
- Useful for large enterprises, ISVs for OEM use
- Includes all of the features of IDS Workgroup Edition plus features required to provide the scalability to handle high user loads and provide 24x7x365 high availability

Available On

Vendor	Specification
IBM	AIX
HP	HP-UX on Itanium, HP-UX on PA-RISC
SUN	Solaris Sparc , x64
IBM POWER	RHEL, Asianux, Debian, SUSE SLES 10
IBM zSeries	RHEL 4, RHEL 5, SUSE SLES 10
Intel x86/x86_64	RHEL, Asianux, Debian, SUSE SLES, Ubuntu LTS
Microsoft	Windows XP, Windows 2003, Windows Vista
Apple (Intel EM64T**)**	Mac OS X (2Q08)

IDS Developer Bundle

- Develop reliable, high performing SOA solutions. Designed for applications development and prototyping
- Integrated with Data Studio to provide a full range of tooling for development
- Available free of charge, for development use only (no production), with no time limits, and no support.
- The following restrictions apply:
  - Maximum database size: 8 GB
  - Maximum concurrent connection: 20
  - CPU memory restrictions: 1 CPU, 1 GB memory maximum

Available On

Vendor	Specification
IBM	AIX
HP	HP-UX on Itanium, HP-UX on PA-RISC
SUN	Solaris Sparc , x64
IBM POWER	RHEL, Asianux, Debian, SUSE SLES 10
IBM zSeries	RHEL 4, RHEL 5, SUSE SLES 10
Intel x86/x86_64	RHEL, Asianux, Debian, SUSE SLES, Ubuntu LTS
Microsoft	Windows XP, Windows 2003, Windows Vista
Apple (Intel EM64T)	Mac OS X (2Q08)

Information Management

IBM IDS 11.5 Overview

Compiled By IDS-CTE
anupn@us.ibm.com

IDS 11.5 Features

Continuous Availability

- Updatable secondaries
- Connection Manager for Workload Management
- Automatic Failover Arbitrator

Admin Free Zone

- Auto Update Stats
- Admin API to Configure HA clusters
- Config options for Defaults during Install
- DRDA Config during Install
- Improved Onconfig.std
- Enhanced Open Admin Tool

Application Development

- Dynamic SQL in SPL
- Basic Text Search support for XML
- XSLT
- DRDA enhancements

Mac OS support

Security

- Single sign-on via Kerberos
- SSL Encryption

Performance

- BIGINT and BIGSERIAL

IDS 11 Availability

(Includes MACH-11 Current Functionality)

- **What is MACH-11?**
- **High-Availability Data Replication (HDR).**
- **Shared Disk Secondary (SDS).**
- **Remote Standalone Secondary (RSS).**
- **Enterprise Replication (ER).**
- **Above Combination.**

IDS Replication Solutions – The Beginning

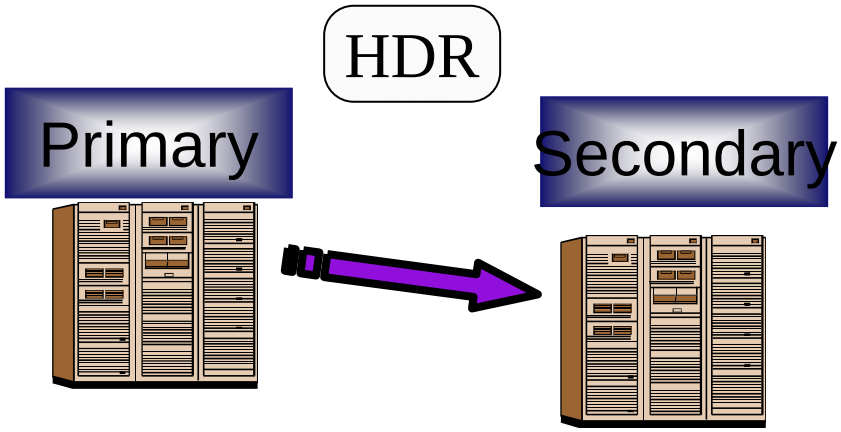

- Replicates an entire database
- Designed for a highly available hot backup
- Secondary can be used for dirty reads
- Simple to administer
- Primary to Secondary

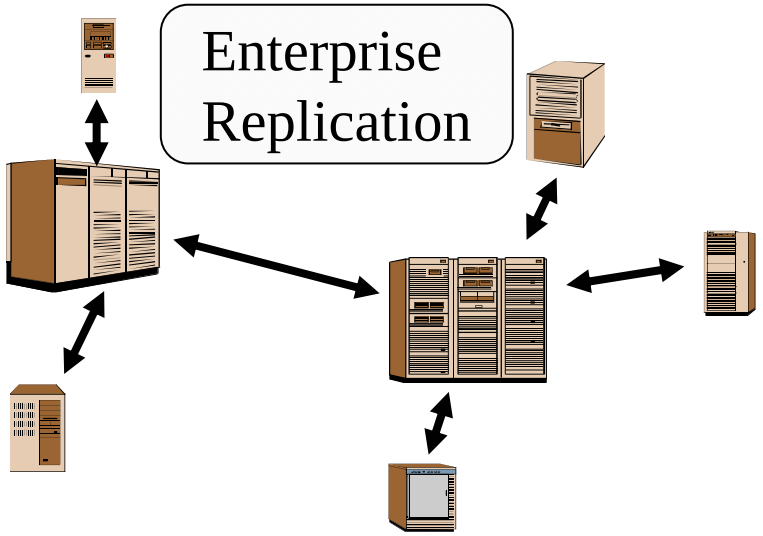

- Replicates parts of a database
- Designed for enterprise data distribution
- Supports active/active updates
- More complex to administer
- Very low latency

N.B. Any ER node can also be an HDR pair

HDR: Remote Standalone Secondaries

- Benefits:
  - Allows for simultaneous local and remote replication for higher availability
  - Capacity relief
  - Simple online setup and use
- Introduced in IDS v11.10
  - Can have 0 to N RSS nodes
  - Compatible with all other IDS availability solutions
- Uses:
  - Reporting
  - Web Applications
  - Additional backup in case primary fails
- In v11.10 RSS nodes are read-only

Replication to Multiple Remote Secondary Nodes

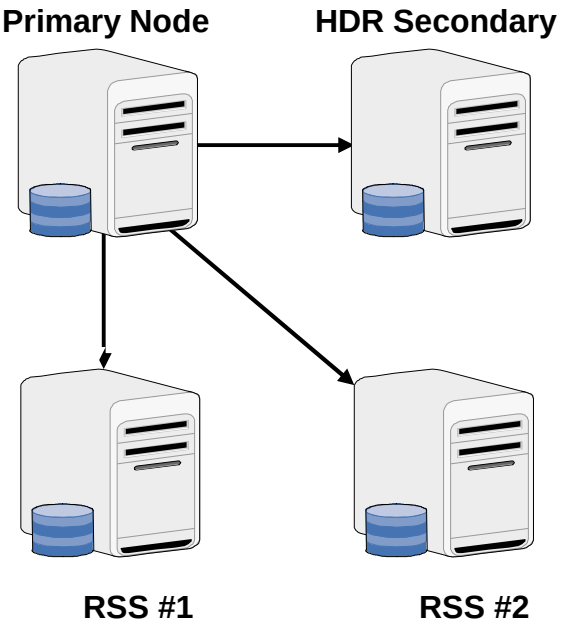

HDR: Shared Disk Secondary Nodes

- Benefits
  - Online capacity relief
  - Very easy to setup
  - Does not duplicate disk space
- Introduced in IDS v11.10
  - SDS nodes share disks with the primary
  - Can have 0 to N SDS nodes
  - Doesn't require any specialized hardware
  - Compatible with all other IDS availability solutions
- In v11.10 SDS nodes are read-only

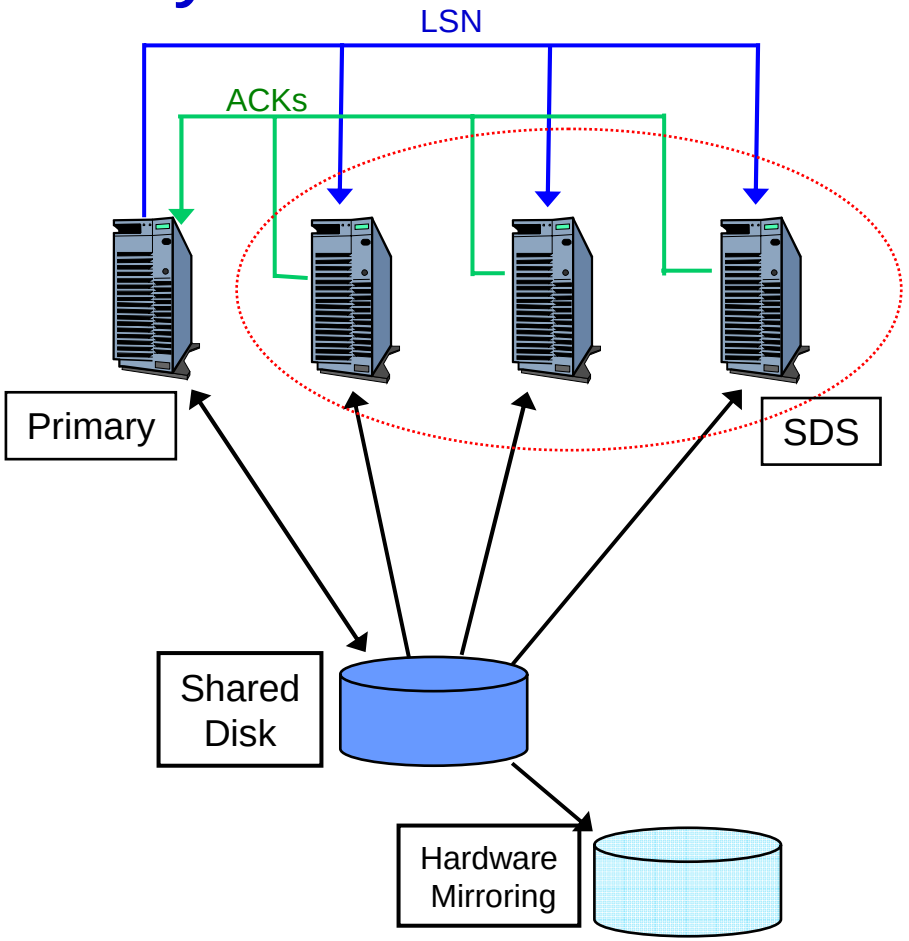

SDS Usage: Capacity as Needed

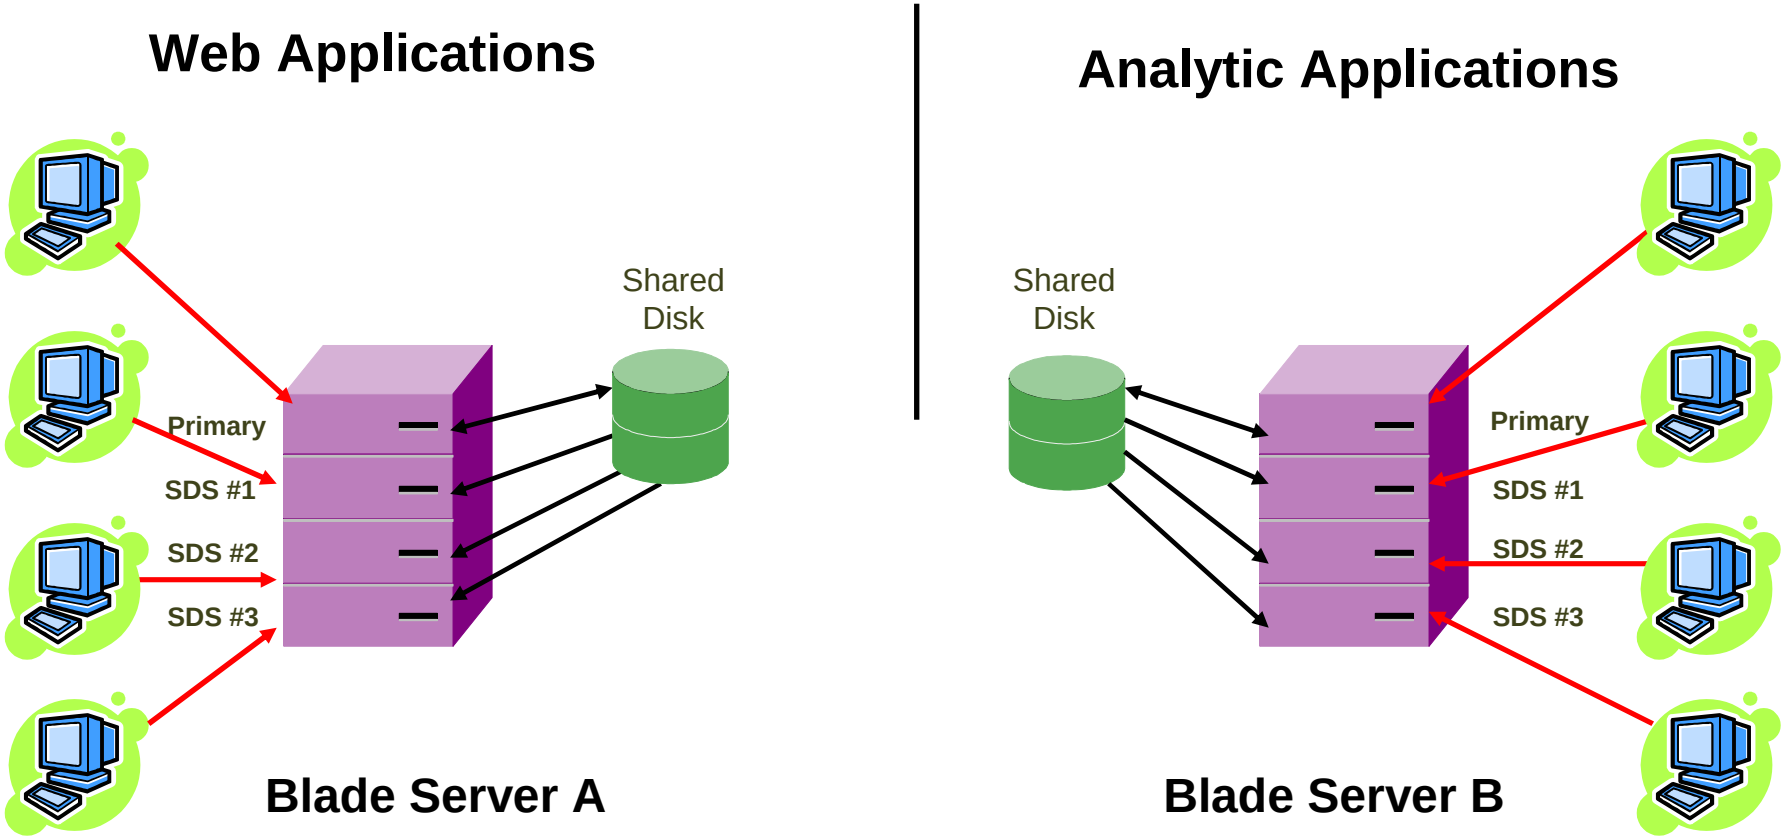

SDS Usage: Capacity as Needed

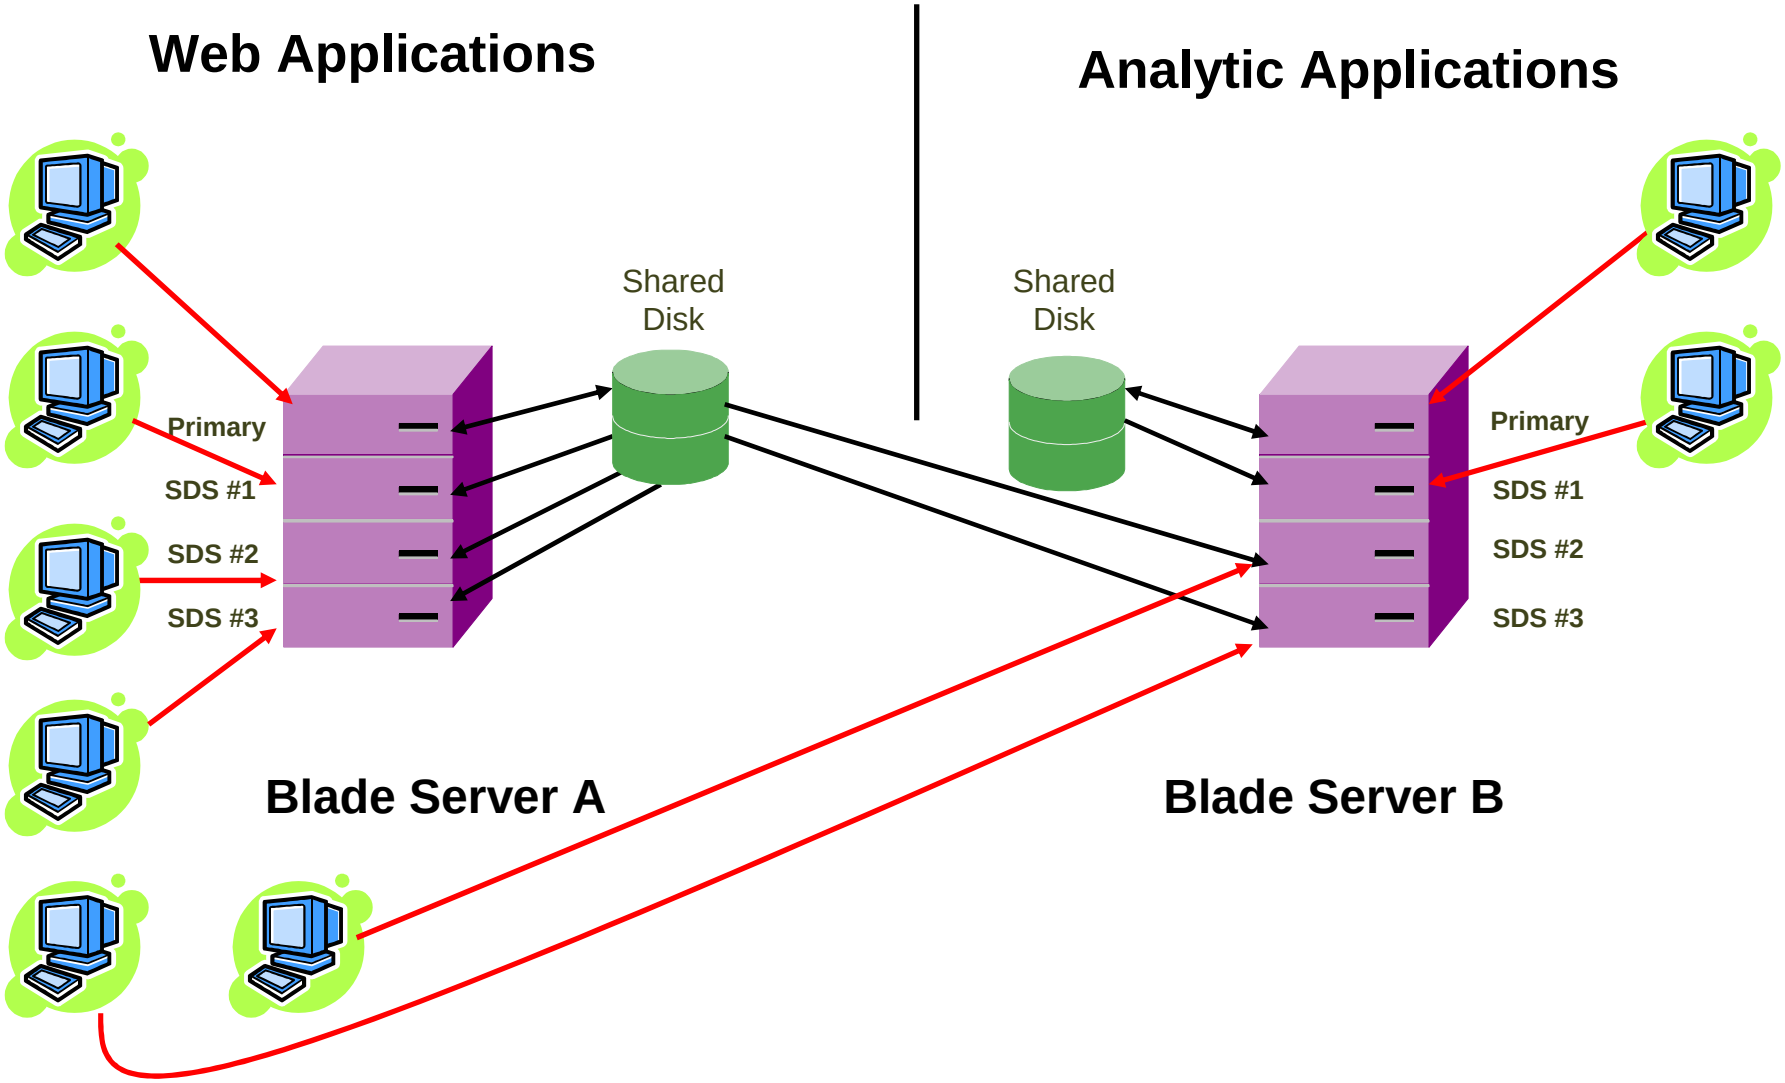

Availability – The Complete Picture

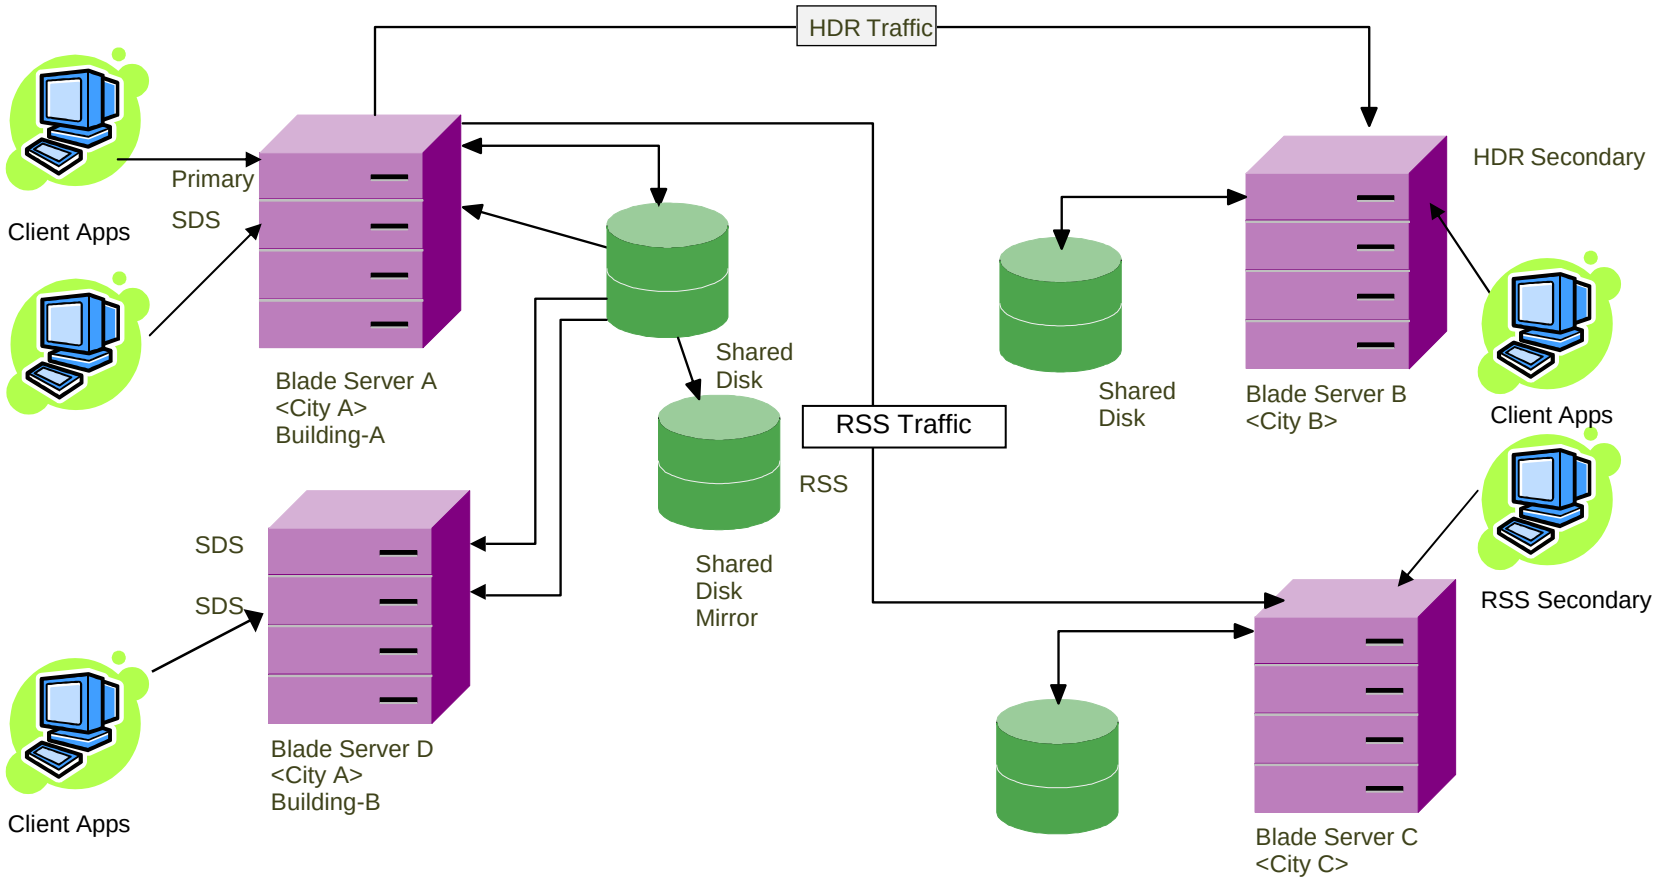

Availability – Both Buildings in City “A” Go Down

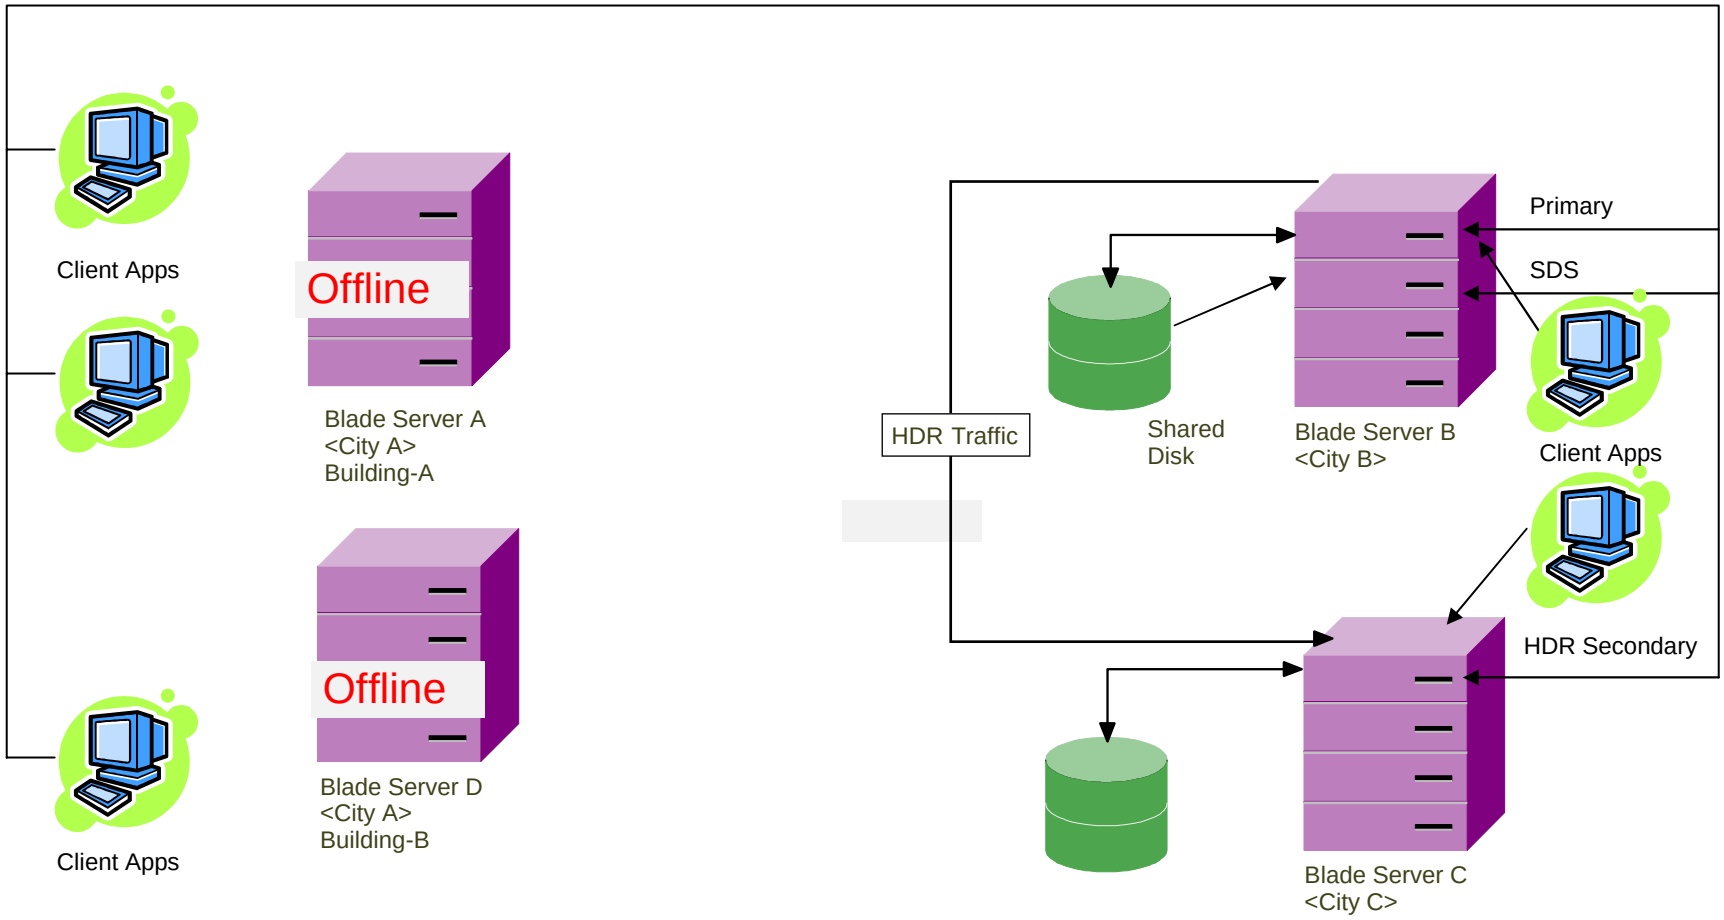

Replication – The Complete Picture

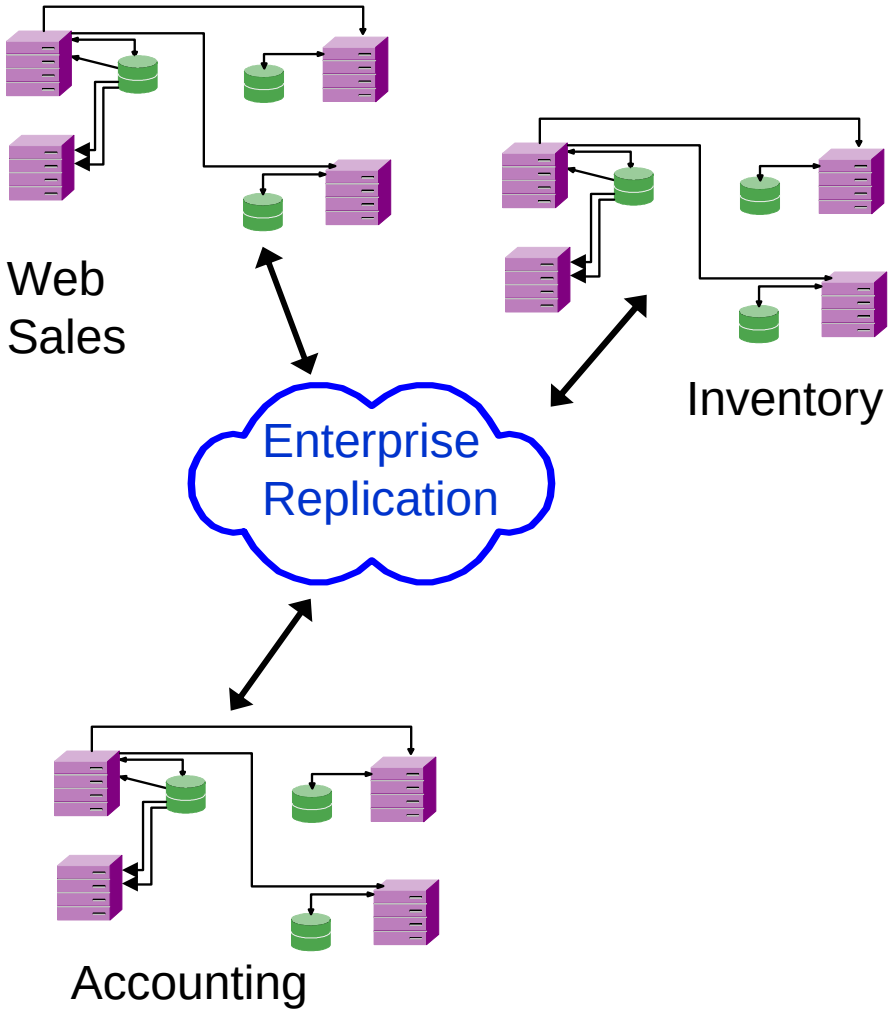

- Any node within the Enterprise Replication can also be a cluster.
- Not only do the nodes within the cluster automatically realign, but so does the ER connections.
- This provides for the ability to not only provide multiple levels of availability, but also the integration of multiple systems.

More Availability Features Via IDS 11.50

(MACH 11 Phase 2 New Functionality)

- **Writes on Secondaries**
- Connection Manager
- Connection Manager Arbitrator
- Password Manager

Writes on Secondary Nodes Now Supported!

- Secondary servers (SDS, RSS, and HDR) now support insert, update, and delete statements.
- Dirty read, last committed, and committed read iosolation supported on SDS node, only dirty reads on RSS and HDR.
  - More isolation levels supported in future release
- Redirected writes work on the basic data types, UDTs (those that store data in the server), logged smart BLOBs, and partition BLOBs.
- Only DML statements at this time.
  - DDL supported in a future release
- Supports temp tables - both explicit and implicit.
- Continues to work with ER.

How do Writes on Secondaries Work?

- Insert/Update/Delete statements executed as much as possible on the secondary
- The update portion of the statement is sent to the primary
- Threads on the primary node perform the required operation.
  - –Triggers and constraint checking is performed on the primary node.
- If the primary node fails subsequent UDI's will be sent to the new primary.

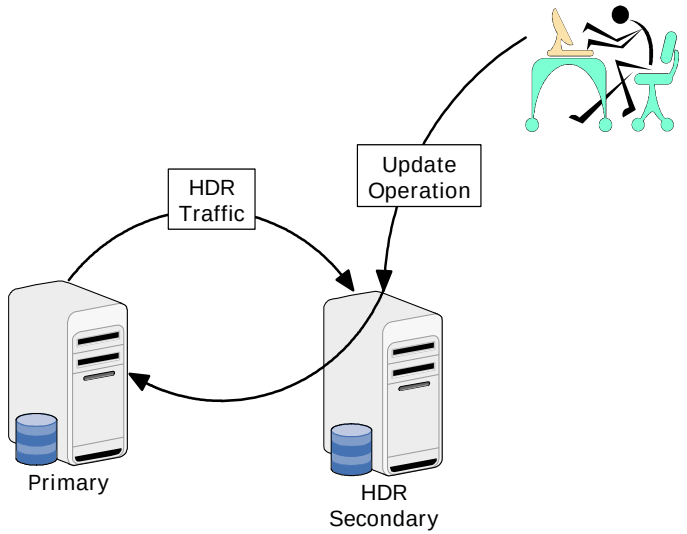

Conflict Handling Between Nodes

- Two options for detecting update conflicts between nodes:
  - Secondaries send “before” and “after” images to the primary
    - Primary compares the before image to current row
  - Secondaries send row version information along with after image
    - Row versioning is enabled when a table is created or altered by adding “with vercol” clause
    - Adds a unique version number and checksum shadow column for each row
    - Primary compares row version number and checksum to current row to detect collisions
    - Small rows may not benefit from turning on row versioning
  - Use of row versions can reduce the network traffic
  - If a collison is detected an EVERCONFLICT (-7350) error is returned.

More Availability Features Via IDS 11.50

(MACH 11 Phase 2 New Functionality)

- Writes on Secondaries
- **Connection Manager.**
- Connection Manager Arbitrator
- Password Manager

Connection Manager - oncmsm

- **ON**line **C**onnection **M**anager and **S**erver **M**onitor
- Standalone program – runs outside of IDS
- Monitors the workload of the nodes in a MACH11 cluster
- Monitors the nodes in a MACH11 cluster for failures
- Does not have to run on any of the cluster nodes
- Can run multiple connection managers for high availability

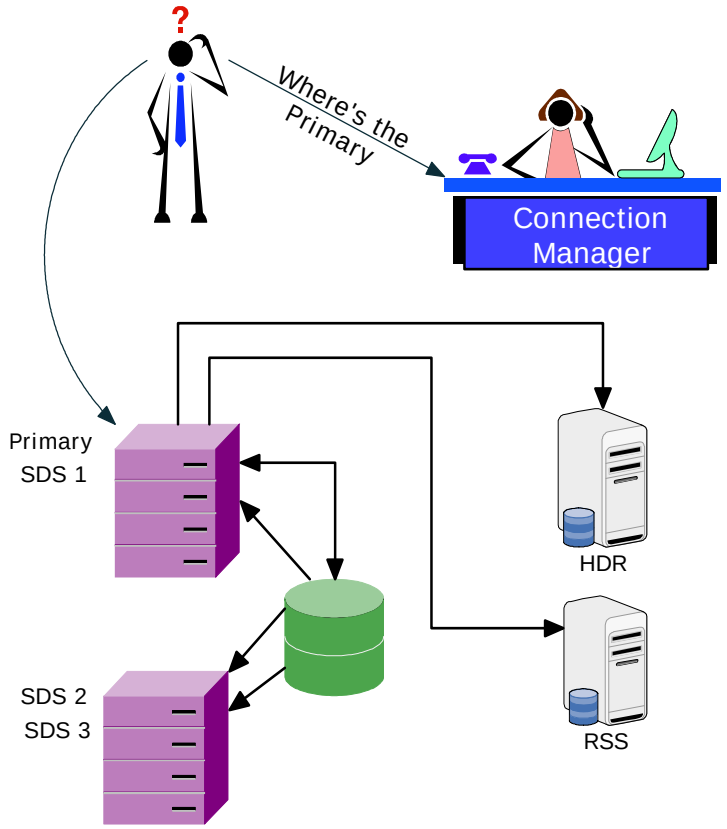

Connection Manager

- Works with CSDK, JDBC and JCC:
- Resolves connection request based on Service Level Agreement (SLA):
  - Connect to primary.
  - Connect to best secondary.
  - Connect to primary or SDS based on workload.

Managing Client Connections – an Example

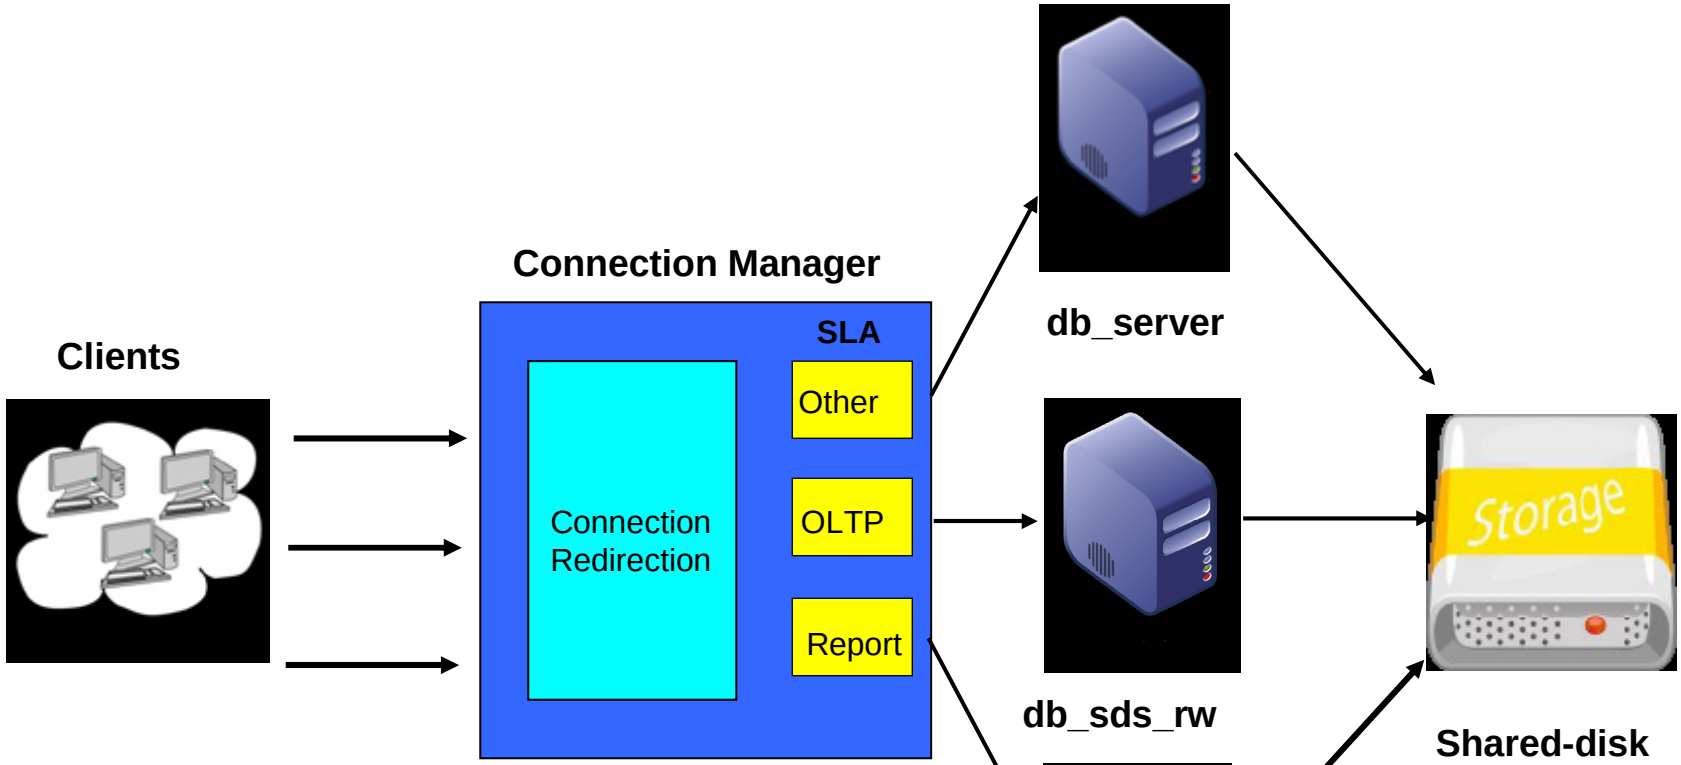

Config file

```
SLA:Report=db_sds_ro+db_sds_rw
SLA:OLTP=(db_server+db_sds_rw)
SLA:Other=db_sds_rw+db_server
```

Connection Manager reroute chart

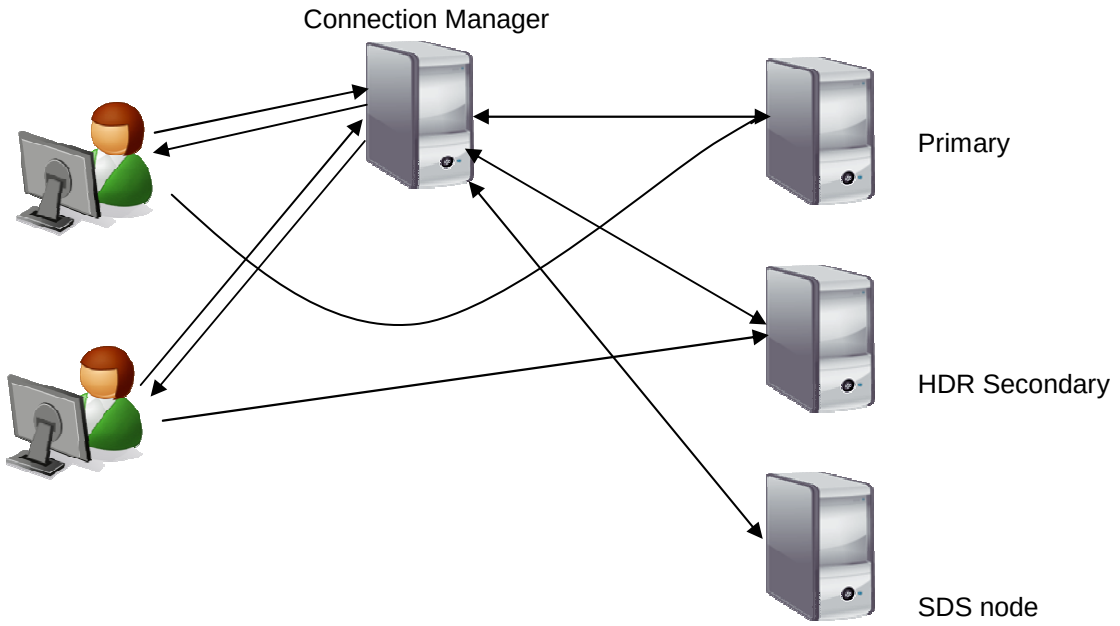

- Clients have the option of connecting to the servers directly or via the connection manager.
- Clients connected via the connection manager can automaticly get assigned to a new node in case of failures.

More Availability Features Via IDS 11.50

(MACH 11 Phase 2 New Functionality)

- Redirected Writes
- Connection Manager
- **Connection Manager Arbitrator.**
- Password Manager

Connection Manager Arbitrator

- The arbitrator function of “CM” provides automatic failover logic for high-availability clusters.
- Also known as the Failover Arbitrator.
- Continuously monitors all nodes in the cluster for failures.
- The Arbitrator decides when a node has failed and which node should become primary
- Will support failover to any of the secondary node – RSS, SDS, or HDR.
- Priority for failover is part of configuration

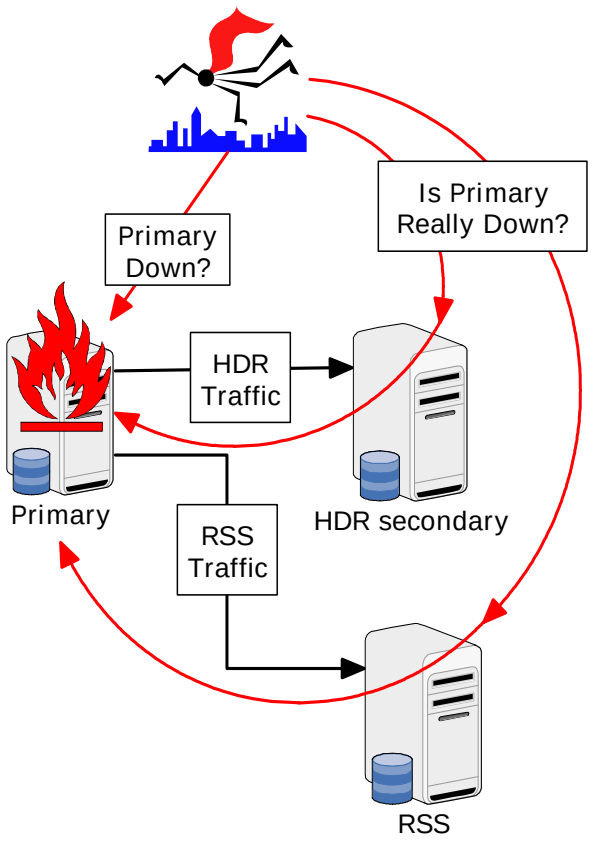

Fail Over Configuration (FOC) Parameter

- Order of failover defined in the parameter in the Connection Manager Configuration File (`$INFORMIXDIR/etc/cmsm.cfg`).

- FOC parameter format:

  `FOC` ***`failover_configuration,timeout_value`***

  Where:

  failover_configuration: One or more of Primary, SDS, HDR, RSS, or a group of server types separated by a plus (+) enclosed in parentheses.

  `timeout_value`*:* Amount of time (in seconds) to wait before attempting a failover.

- Default of the FOC parameter:

  `FOC SDS+HDR+RSS,0`

FOC Parameter

- Generic Key Words: `primary, HDR, RSS, SDS`

  Example:

  `FOC  serv1+(serv2+SDS)+HDR+RSS,10`

  Failover first to serv1, then to the best choice of serv2 or any running SDS node, then to the HDR secondary, then to any RSS node. Allow 10 seconds before performing the failover process.

- The FOC can be passed to the connection manager (oncmsm) on the command line as "`-f serv1+(serv2+SDS)+HDR+RSS,10`".

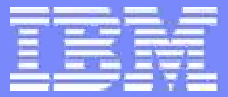

Information Management

Auto Update Statistics

IBM Informix Dynamic Server

What is Auto Update Statistics (AUS)?

- The ability to automate the maintenance of optimizer statistics
- It accomplishes this by:
  - Identifying tables based on defined set of policies which require new or updated optimizer statistics
  - Runs update statistics on the required tables in a priority order within a defined window of time

Why Auto-Update Statistics (AUS)?

- **To simplify the repetitive work DBA's are required to undertake to ensure optimal performance.**
- **Update statistics provides the information the query optimizer and database engine needs to run queries quickly and efficiently.**
  - Many customers new to Informix do not realize they must run update statistics and see poor performance due to the lack of optimizer statistics.
- **The effort it takes DBA to develop an optimal plan to run update statistics is complex and time consuming.**
  - DBA settle for a non-optimal plan or lack the understanding to develop an optimal plan.
  - Each customer must re-invent their own plan and procedures

OAT – Auto Update Statistics (AUS)

- Two pieces comprise the AUS solutions:
  - A set of procedures which will be invoked by the database scheduler to automate update statistics.
    - Installed as part of the server
    - Users can write SQL statements to manipulate the AUS policies.
    - Allows for partners to embed programmatic solutions
  - A OAT graphical interface allowing easier control of the different AUS policies.
    - Open source download available from iiug.org and IBM websites
    - Display AUS information in easy to read format
    - Simple point and click interface

AUS Database Tasks

- **AUS has two database tasks in the database scheduler**
  - Auto Update Statistics Evaluation
    - Analyzes all the tables in all logged databases
    - Locates tables which exceed the predefined polices
    - Builds the update statistics commands
    - Inserts the commands into the table aus_cmd_list
  - Auto Update Statistics Refresh
    - Execute the update statistics commands in the aus_cmd_list table
    - After the command completes, it is moved to the aus_cmd_comp table

AUS also utilizes the information produced by the mon_table_profile task in the database scheduler to determine how much the table has changed. Do not disable this task!!

Auto Update Statistics Policies

- **Threshold Policies**
  - A set of thresholds and rules which apply to different attributes of the optimizer statistics.
  - These policies are aimed at determining which tables require updated statistics.

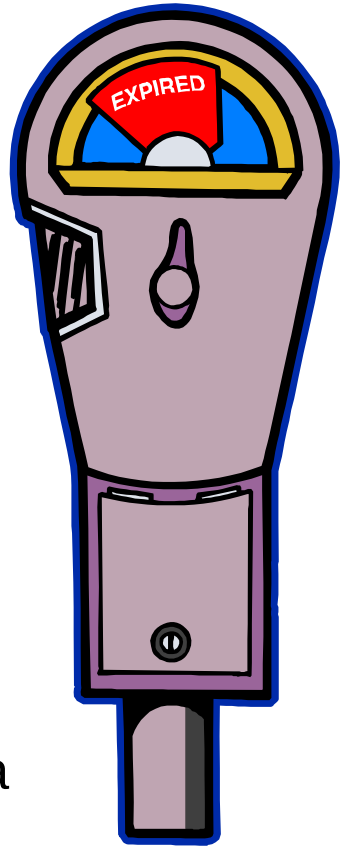

- **Two kinds of Thresholds:**
  - **Time Based**
    - When the table statistics or distributions exceed a specified age, they reached the threshold.
  - **Modification Based**
    - When the number of inserts, updates or deletes to a table exceeds a specific percentage, the threshold for statistics and distributions are reached.

Configuration Parameters

- **Configuration Parameters can be updated two ways**
  - Modified in OAT's AUS configuration page
  - Directly by updating the table sysadmin:ph_threshold

Ph_Threshold Table Parameter	Default Value	Description
AUS_AGE	30 days	How old the statistics or distributions can be before they will be updated; even if there are no changes.
AUS_CHANGE	10	If the table has changed by more than this percentage then the statistics an distributions will be updated.
AUS_AUTO_RULES	1	Turning this on will ensure statistics and distributions are updated to a Informix minimum suggested guidelines. If the current statistics or distributions exceed the minimum suggested guidelines then the current setting will be used.
AUS_SMALL_TABLES	100	Tables containing less than this number of rows will always have their statistics rebuilt.
AUS_PDQ	10	Auto Update statistics executes with this PDQ priority.

OAT's Graphical Interface for AUS

- **General overview of the state of Statistics**

Graph showing the state of the entire data server

Show state of each databases statistics by table

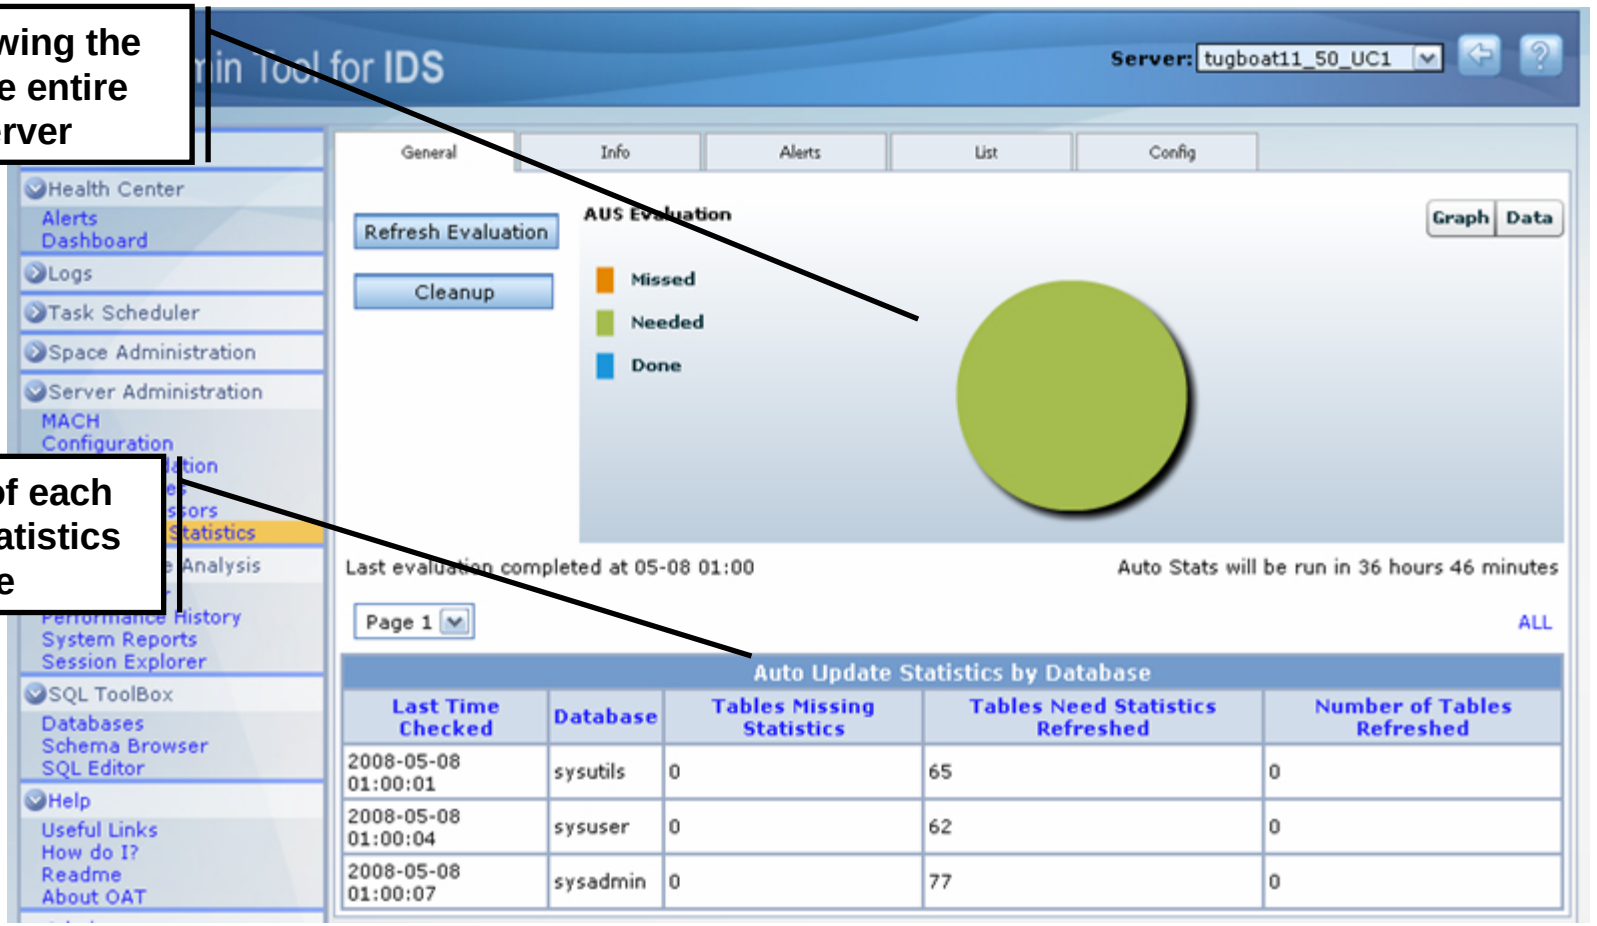

OAT's Graphical Interface for AUS

- Overview of AUS's schedule and policy settings

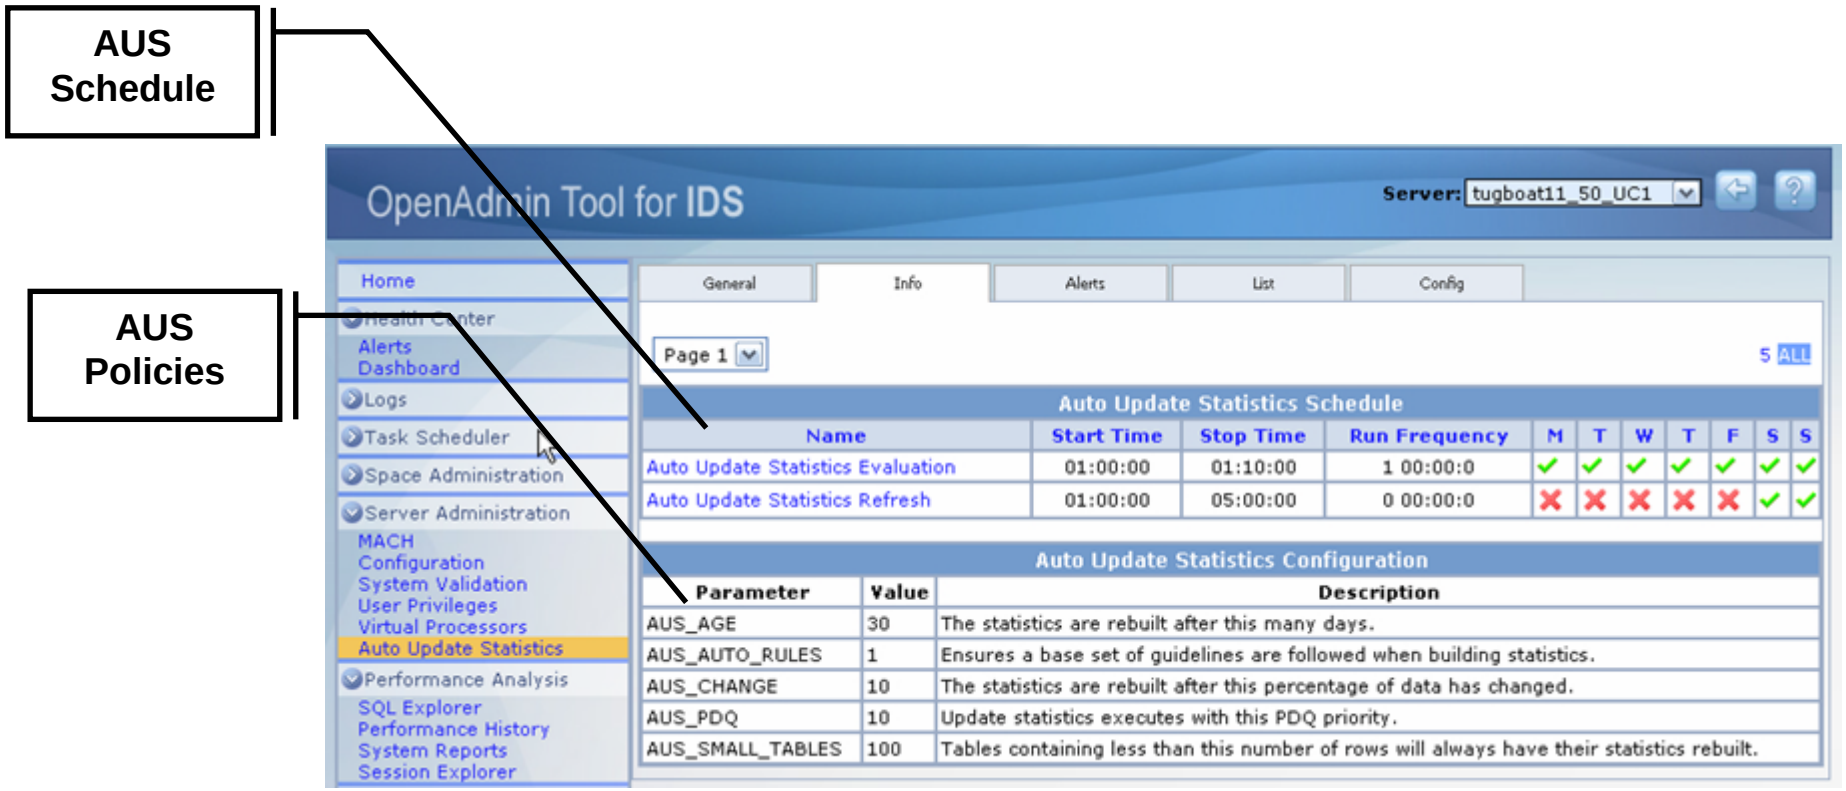

OAT's Graphical Interface for AUS

- The AUS evaluator will post alerts about databases which require new or refreshed statistics

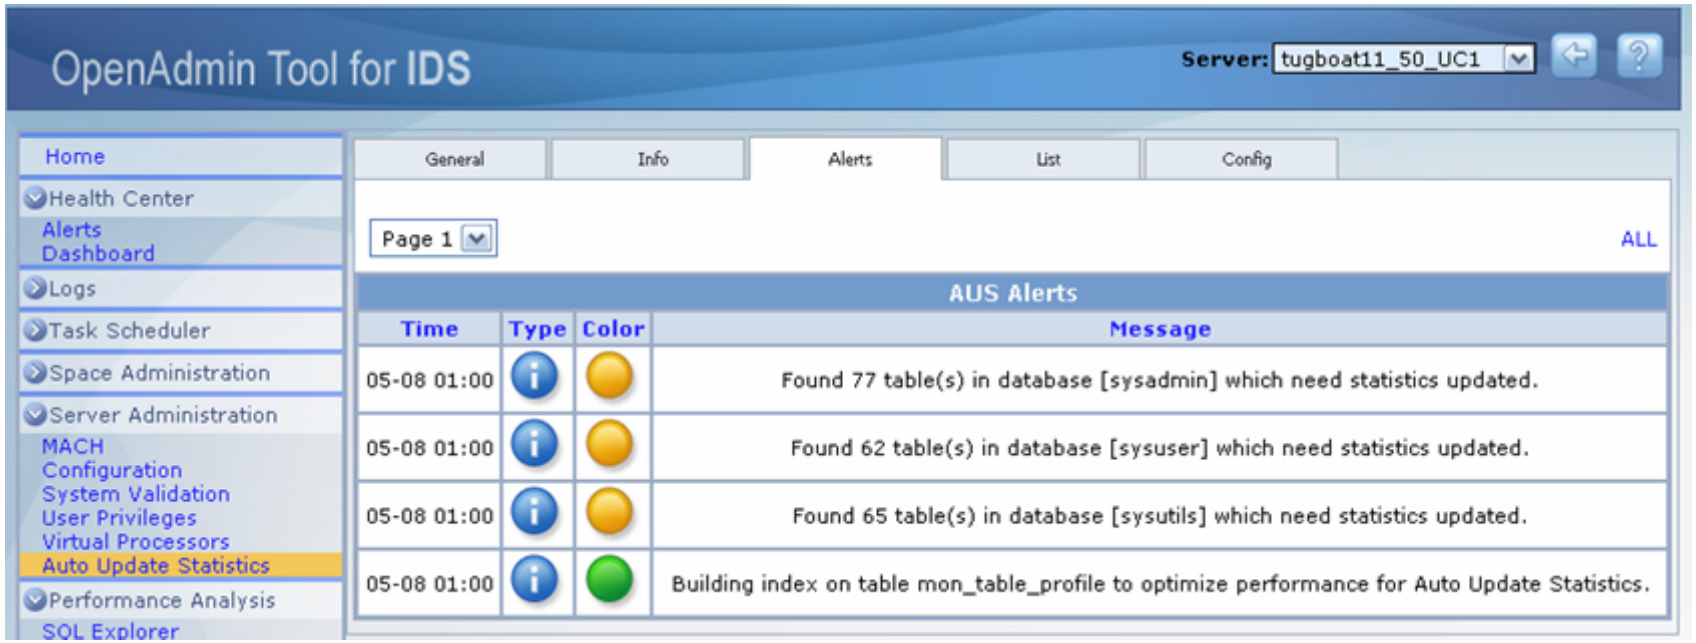

OAT's Graphical Interface for AUS

- The list of Update Statistics command created by the AUS Evaluator Task

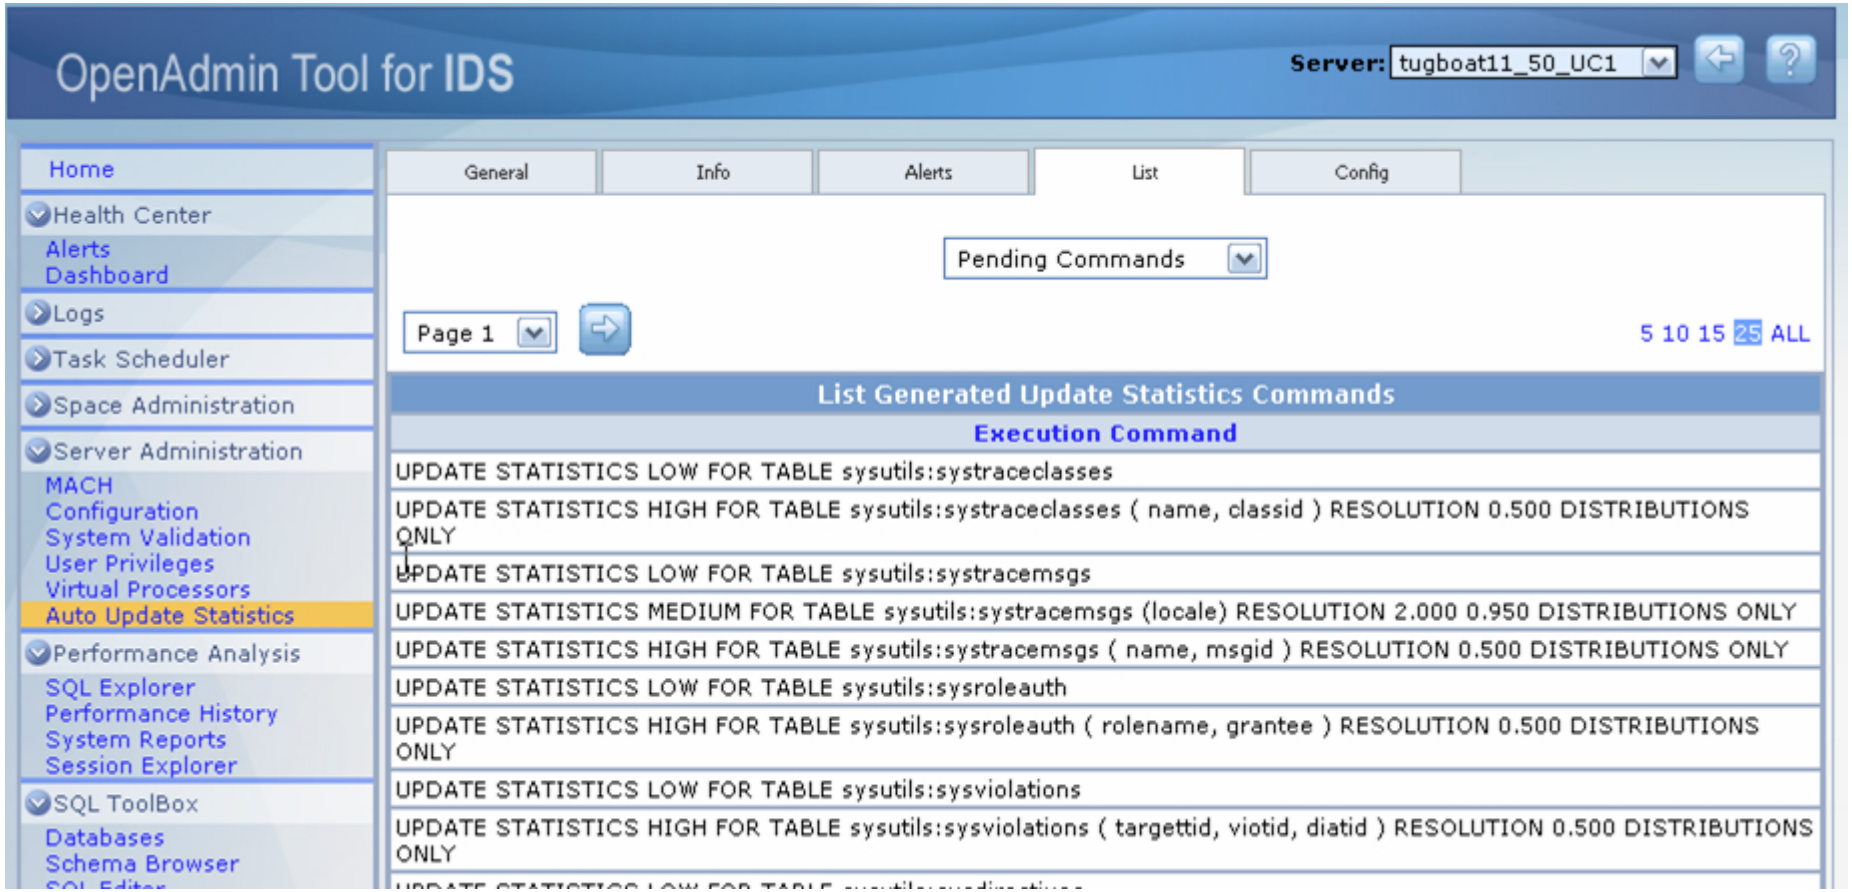

OAT's Graphical Interface for AUS

- Setting the configuration parameter of AUS

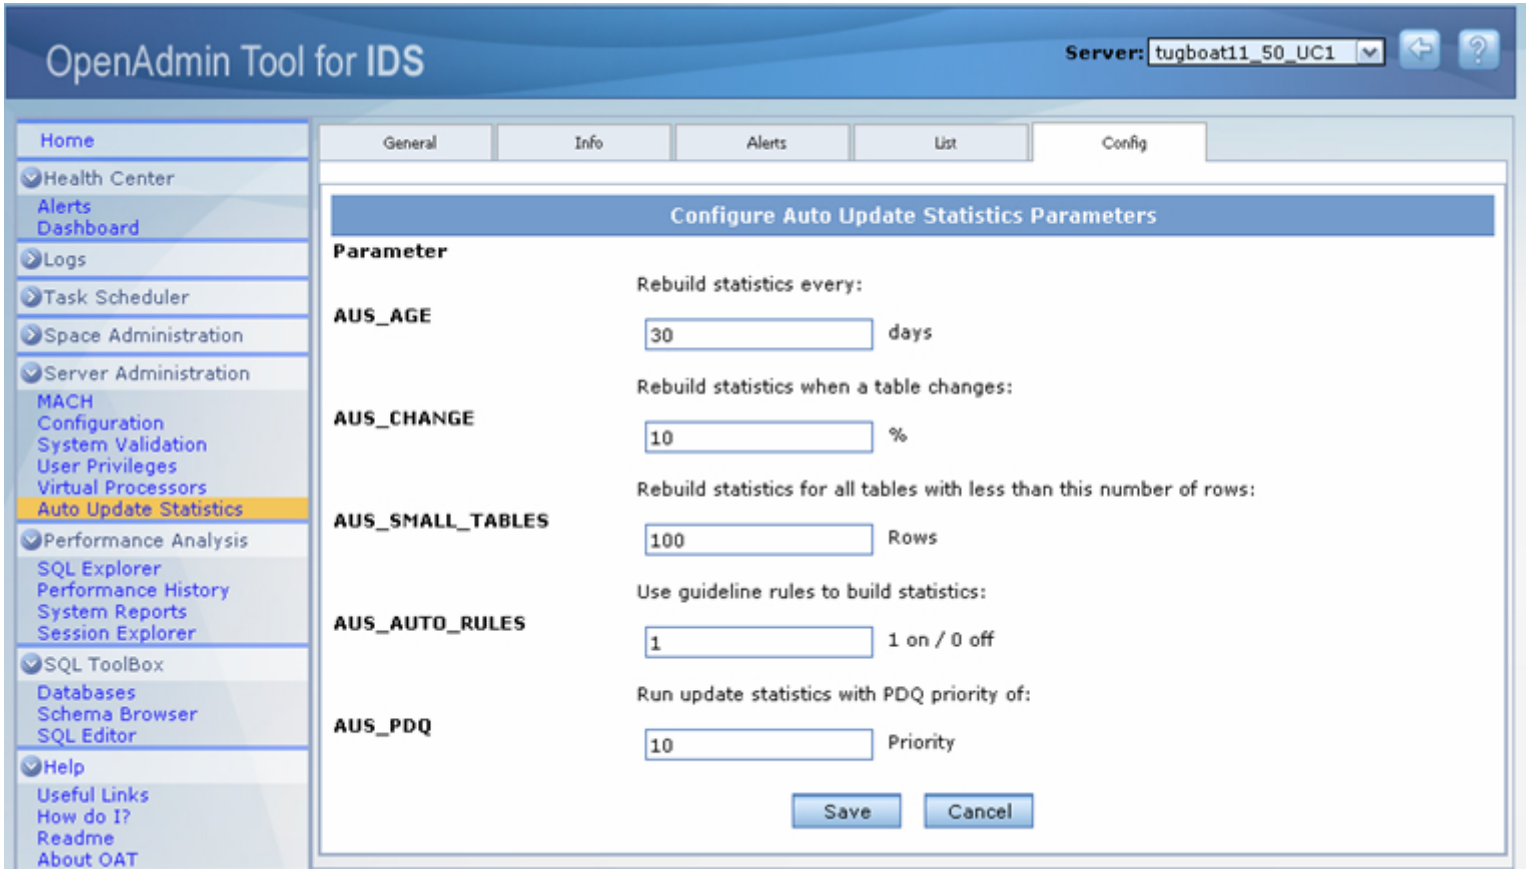

Minimum Auto Update Statistics Guidelines

- Low on all indexes (or on the table if no indexes exist).
- High on all lead keys of indexes.
- Medium on all non-lead keys:
  - The minimum resolution used for Medium is 2.0
  - The minimum confidence used for Medium is 0.95
  - The minimum resolution used for High is 0.5
- When AUTO_RULES is set to 1, for each table, the suggested update statistics guidelines will be followed. In the case of refreshing a user created distribution, we will never downgrade its level; as an example, we suggest medium mode distributions on the secondary columns of an index. If a user had created the distribution with H mode, H will be used. This is true not only for the level of distributions, but also the resolution, confidence and sampling size.
- Statistics for tables with rows less than AUS_SMALL_TABLES parameter will be unconditionally refreshed.

Scheduling Information

- **Run time window**

- Since AUS is a resource intensive operation, a specific run window can be defined. The run window definition includes start time, the end time and the days of the week to run the job.
- Configured from the Open Admin tool.
- Implemented by setting the start and end times of the aus_refresh_stats task in ph_task table in sysadmin.
- Default runtime window is 1:00 AM – 5:00 AM daily.

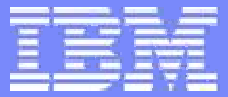

Information Management

Dynamic SQL in SPL

Dynamic SQL in Stored Procedures

- Statements can now be dynamically constructed and executed
- New Dynamic SQL Statement in SPL - EXECUTE IMMEDIATE

Syntax:

EXECUTE IMMEDIATE { *SQL_quoted_string* | *Str_variable* } ;

SQL_quoted_string: *A string containing a single SQL statement*

Str_variable: *A character variable containing the SQL statement*

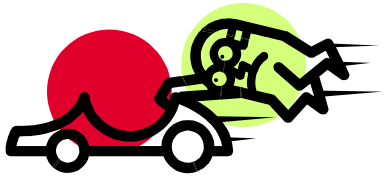

Dynamic SQL Example

```
CREATE PROCEDURE MYPROC()
       RETURNING INT;
       DEFINE A0 VARCHAR(30);
       DEFINE A1 VARCHAR(5);
       DEFINE A2 INT;
       DEFINE A3 VARCHAR(60);
       DEFINE A4 INT;
       LET A0 = "INSERT INTO DYN_TAB VALUES (";
       LET A1 = ")";
       FOR A2 = 1 TO 100
           LET A3 = A0 || A2 || A1;
           EXECUTE IMMEDIATE A3 ;
       END FOR;
       SELECT COUNT(DISTINCT C1) INTO A4 FROM T1;
       RETURN A4;
END PROCEDURE;
-- should return 100 as 100 unique values got
-- inserted by the EXECUTE IMMDIATE in loop
EXECUTE PROCEDURE MYPROC();
```

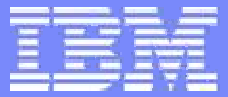

Information Management

Changes to the Open Administration Tool (OAT)

OpenAdmin Tool for IDS: New Features

- **Managing High Availability Clusters (MACH)**
  - Configuring the IDSD daemon
  - Starting and Stopping Servers in the Cluster
  - Creating an SD Secondary Server
  - Managing Service Level Agreements and Failover Configurations with the Connection Manager Wizard
- **Generic Server Administration Features**
  - Automated Update Statistics
  - Server Configuration
  - Historical Data Graphs
  - Read-Only Access
  - Integrity Checks
  - Server Switch
  - VP Administration
  - Privileges Manager
  - Connection Environment Variable Support
  - Task Scheduler Wizard

Administer Anywhere with Open Admin Tool for IDS

- Web Access
- Graphical Interface
- Drill down to complete details
- Administer multiple remote servers

- Easy to Customize
- Open Source
- PHP-based
- Uses IDS SQL Admin API

Open Admin Tool For IDS – Replication

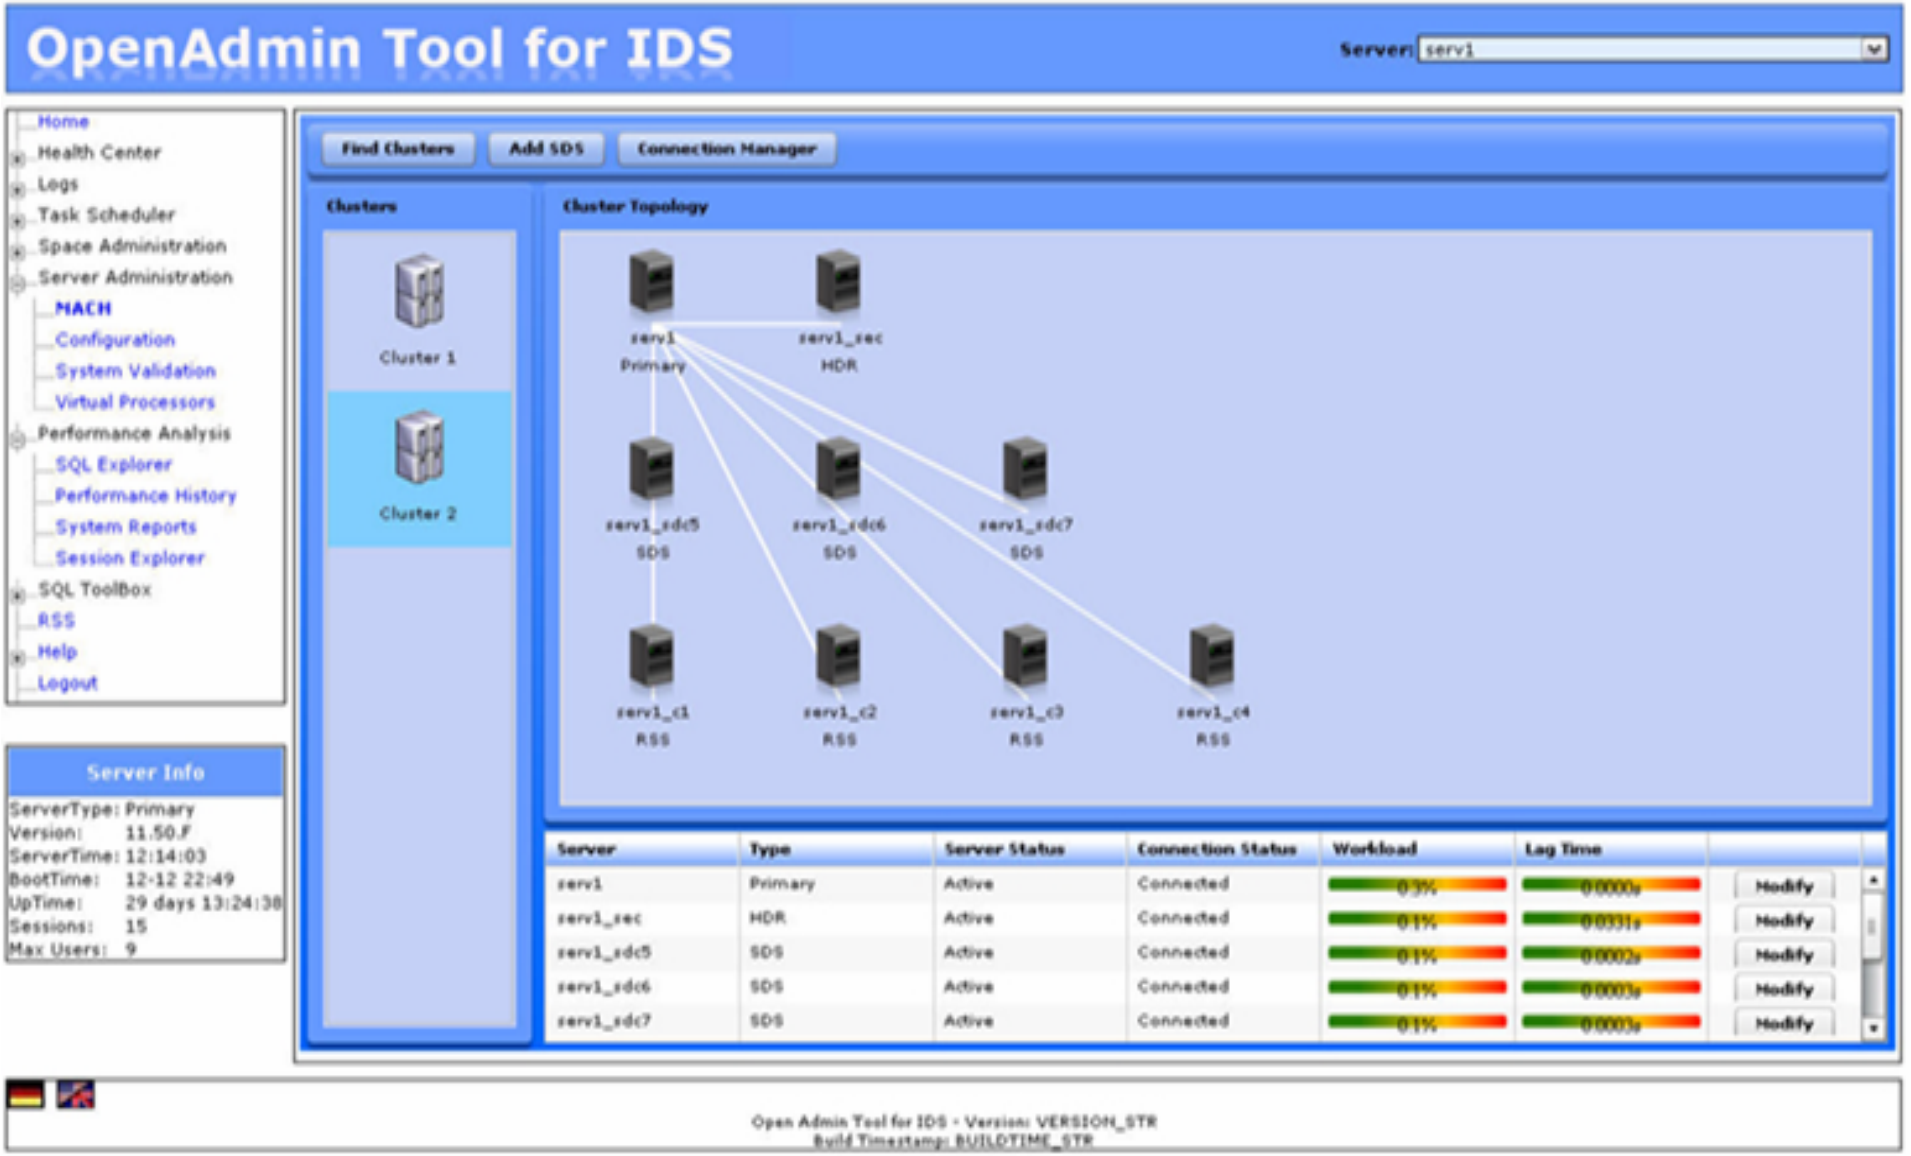

OAT For IDS – Connection Manager

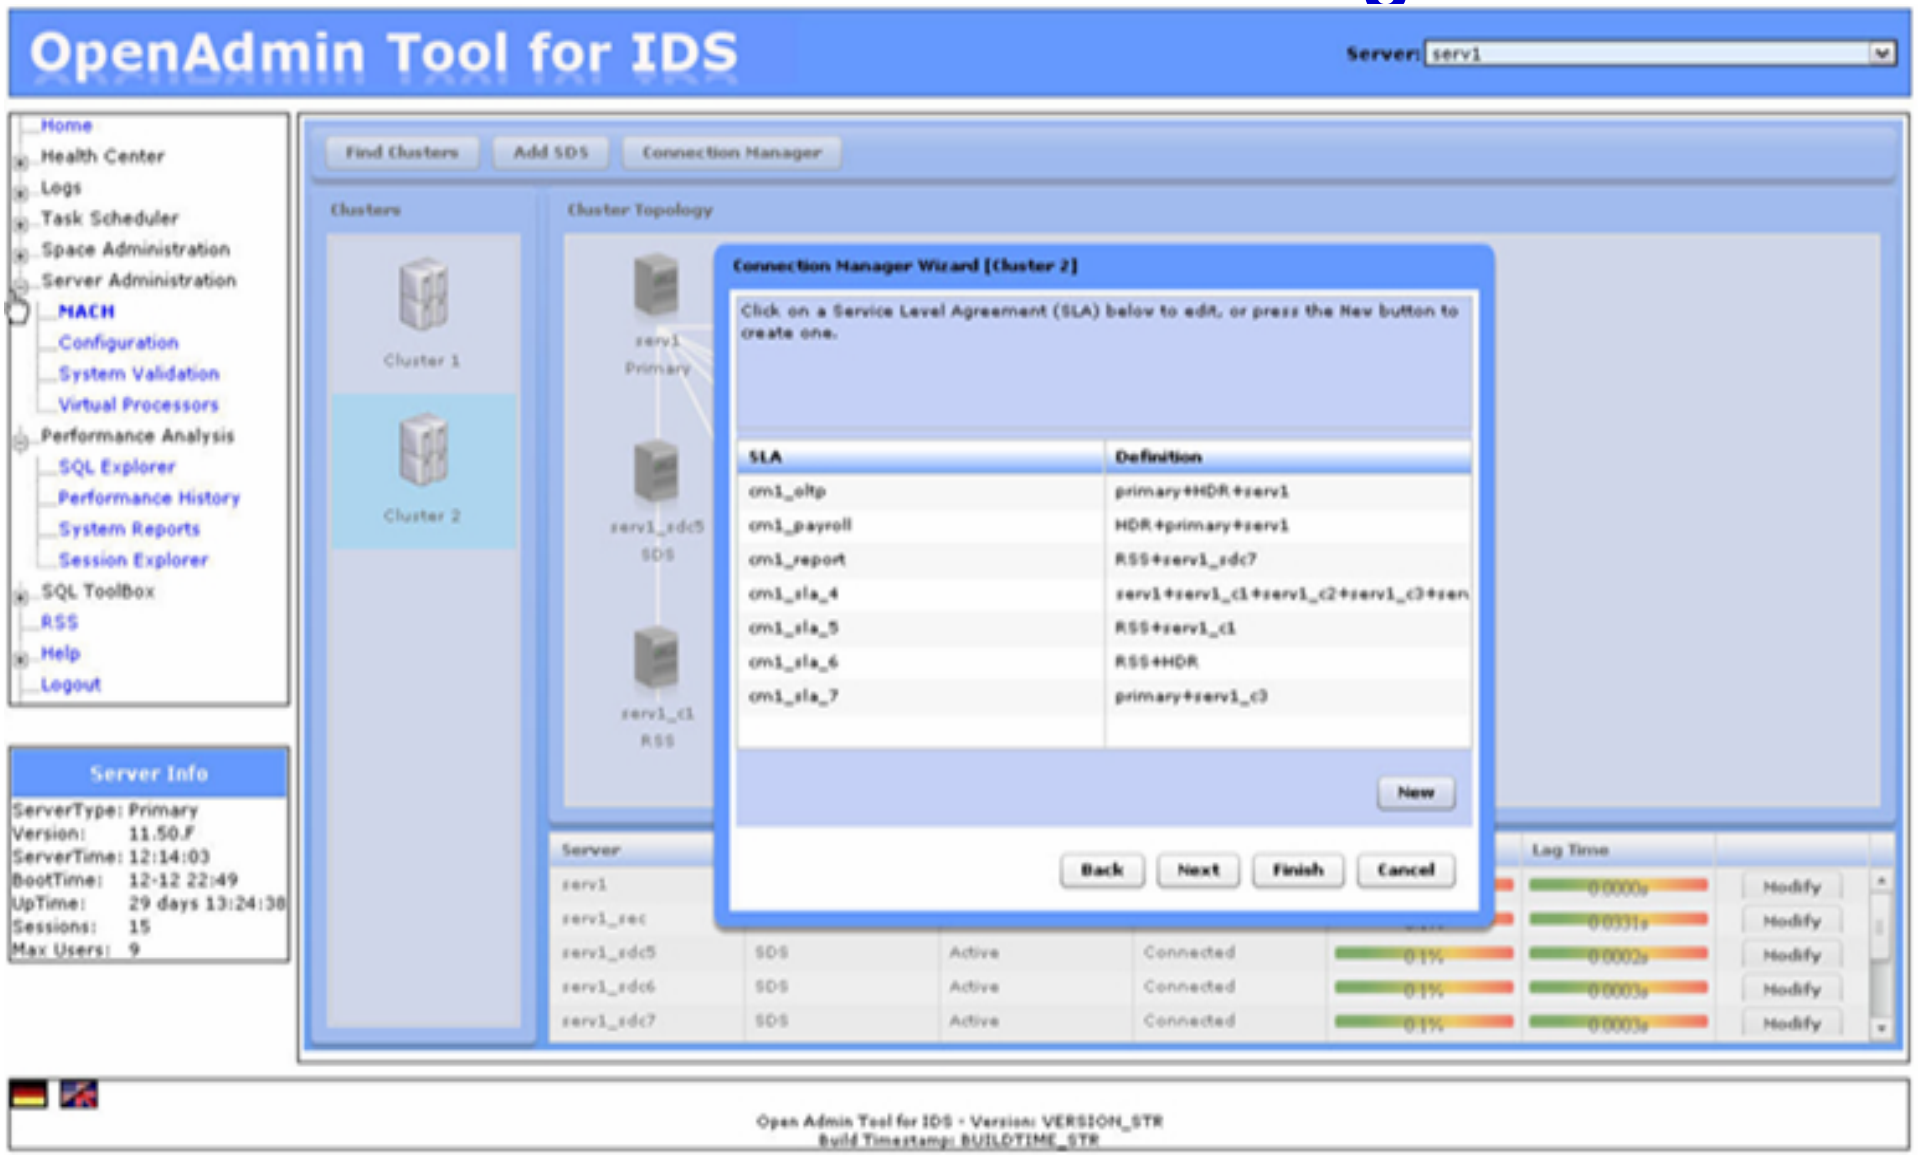

New connection daemon - IDSD

- Allows users to start or stop an IDS server from OAT - new IDS Daemon to start IDS
  - Runs on same host as IDS server
  - Started by inetd or xinetd
  - Simple configuration
  - Included with IDS distribution
- Can also start/stop nodes in a MACH11 cluster

Wizards to Add SDS and Connection Manager

- Adds an Shared Disk Secondary server to a cluster
  - Takes configuration and environment from the primary server
  - Prompts for values that need to be modified for new SDS
  - Creates ONCONFIG and SQLHOSTS files
  - Uses IDSD to transfer configuration files to remote host
  - IDSD starts starts IDS on SDS node

- Allows user to create or edit the connection manager configuration file
  - Allows user to modify Failover Configuration (FOC)

Server Configuration

The Server Configuration Feature consists of two main parts:

- Dynamic Reconfiguration
  - Updates the config variables that can be changed during run-time
- Onconfig Parameter Recommendations
  - Provides recommendations for many config parameters, based on current system information.
  - DBAs can use these recommendations to better tune their IDS database servers.
  - Recommended changes are highlighted in yellow on the System > Configuration Page; use the Show Recommendations filter (Onconfig Options field) to view all recommendations.

Historical Data Graphs

- Server profile information, such as disk reads or writes, can be captured on a regular basis.
- Historical Data Graphs provides a graphical tool for analyzing the server profile information over a period of time.
- This information can be useful in the performance analysis of the data server.

Read Only Access

- Can now specify an IDS server group as ‘read only’
  - Disables any administration privileges (such as for adding a chunk) in OpenAdmin Tool, for all servers defined within that group.
  - Specify for existing groups in the ‘Admin’ section of OAT under Manage Connections > Connection Admin.
- You can also specify this option when creating new groups.

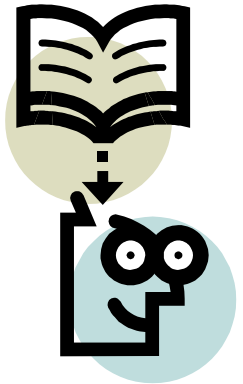

Integrity Checks

The OpenAdmin Tool now provides an interface to the following integrity checks:

- Check the integrity of all the available databases and tables on a server
- Checks all the available tables in a given database
- Checks individual tables
- Checks the integrity of one or all dbspace extents that are available on a particular server
- Checking integrity of tables and extents for the current dbspace

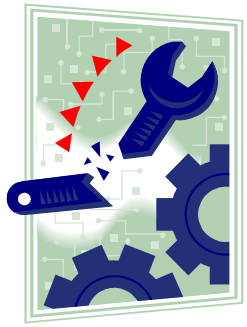

Virtual Processor Administration

OAT can now view and manage virtual processors:

- View graphical information about VP classes
- View statistics about each VP class
- Add or drop a virtual processor

Privileges Manager

- **Database-Level Privileges**
  - View all existing privileges
  - Grant, modify, or revoke database-level privileges for a user or role
- **Table-Level Privileges**
  - View all existing privileges
  - Grant, modify, or revoke table-level privileges for a user or role
- **Roles**
  - Create and view roles

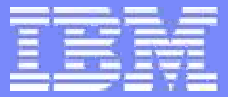

Information Management

Visual Explain

Visual Explain

- A new function called EXPLAIN_SQL converts “explain” output into XML format
- IBM DataStudio can accept this XML output and show the query plan graphically to the user.
- Signature of this function:

```
informix.explain_sql(INOUT major_version int,
                     INOUT minor_version int,
                     requested_locale varchar(33),
                     xml_input blob,
                     xml_filter blob,
                     OUT xml_output blob,
                     OUT xml_message blob)
returns blob as xml_plan
```

Visual Explain

- Sample visual explain in DataStudio

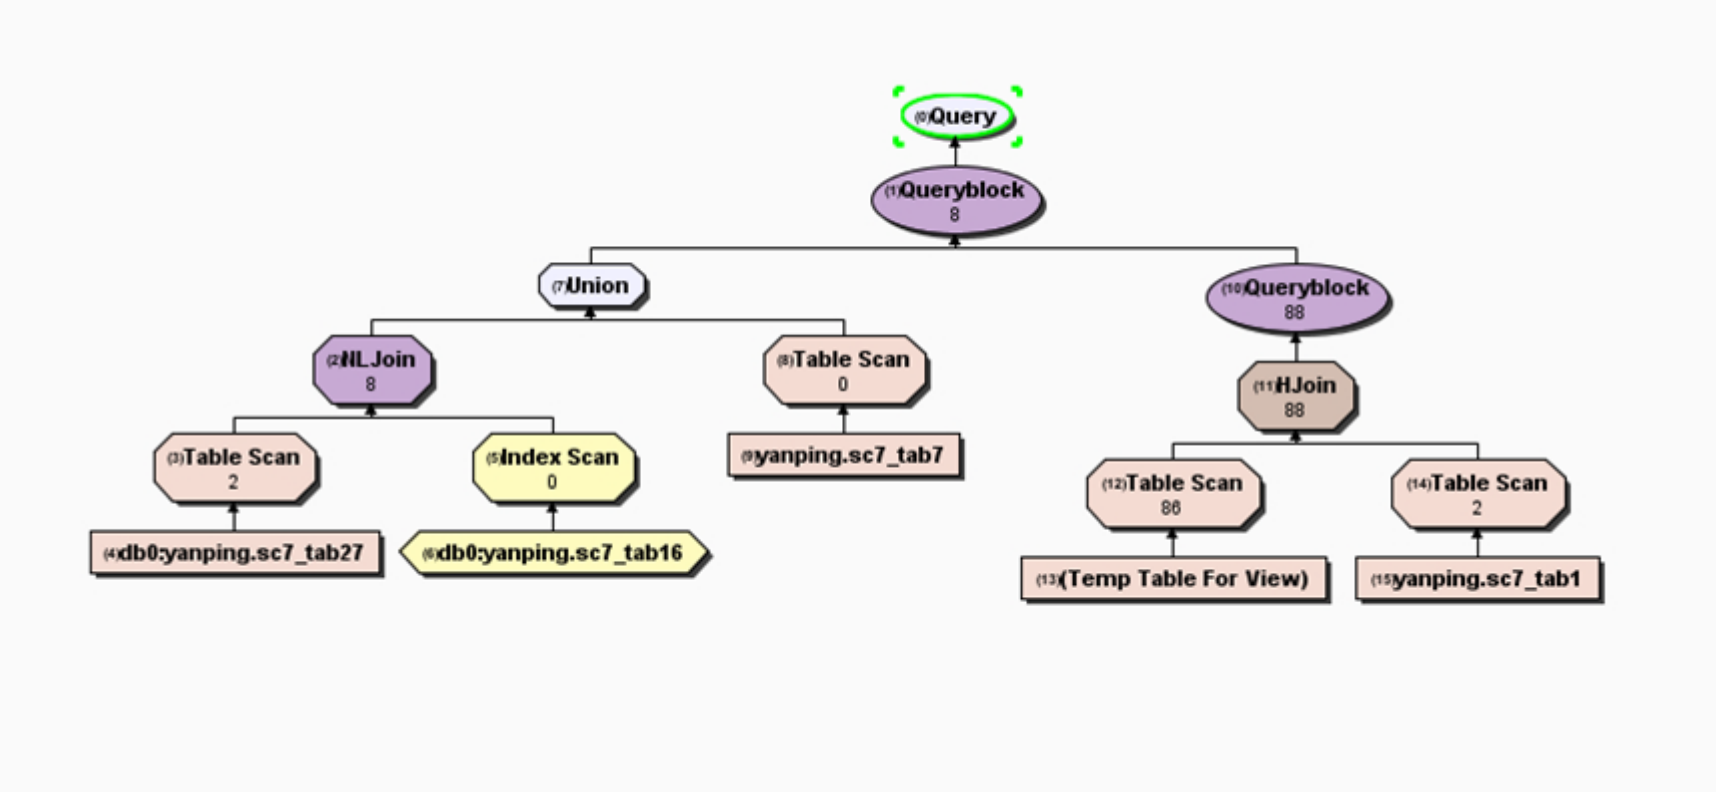

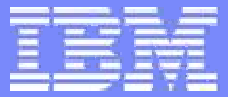

Information Management

Configuration Wizard

Configuration Wizard

- A new option to invoke the configuration wizard added to IDS installation tool
- First Configuration Wizard asks a few questions such as:
  - Number of CPU(s) (prefilled with max found)
  - Amount of memory to use (prefilled with max found)
  - Number of connections
  - Number of OLTP clients to support
  - Number of DSS clients to support
- From these answers an onconfig file is generated and installed in $INFORMIXDIR/etc
- Underlying utility is found in $INFORMIX/bin/genonconfig

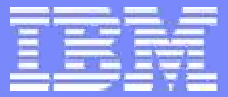

Information Management

Macintosh Port

Macintosh Port

- Runs on Leopard
- Intel 64 bit supported on client and server
- Strong interest in publishing & education, gaming market.
- Some functionality will come later
  - JDBC
  - Geodetic support
  - Excalibur Text
  - Support for FASTPOLL

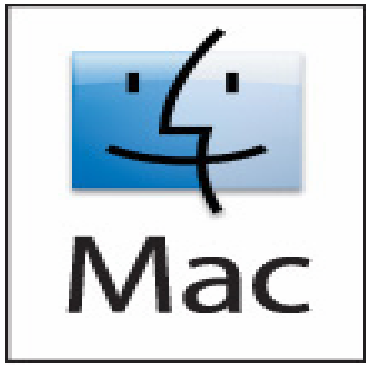

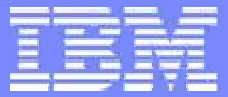

Information Management

XML Enhancements

XSLT

- eXtensible Stylesheet Language Transformation
- XSLT is used to transform XML documents
  - XML to XML (to conform to different schema/standards)
  - XML to HTML
  - XML to PDF
- 3 XSLT functions
  - XSLT Transform()
  - XSLT TransformAsCLOB()
  - XSLT TransformAsBLOB()

Searching XML Documents as part of Basic Text Search

- Customer can now generate their own stopword lists
- Can indicate which xml tags to index at index create time.
  - Values get indexed, not tags.
  - Also options to index all tags as well as treat tags as text.
- Can now limit searches to only certain XML tags
- “XPath" queries supported in search.
  - Search string can contain a path to search /books/title/chapter/....
  - Path can be relative or fixed.
- tags can be completely ignored
  - XML doc treated as a set of values
  - Looses the ability to search by tags
- You can specify whether the namespace should be required in searches
- Markup in sentences can be ignored
  - "foo <bold> bar </bold>" can either be indexed as the single "foo bar" or as "foo" and "bar". This matters for proximity searches for instance.
- New function provided which returns all the unique tags in the XML doc.

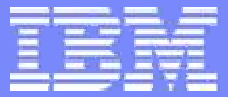

Information Management

Single Sign On In IDS

What is Single Sign On (SSO) ?

- Users enter their password once to gain access to resources
- Password entered during login
- Password need not be entered again to authenticate to the IDS
- Authentication is invisible to the user
- IDS uses Kerberos for SSO support

So what does it mean?

- 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.

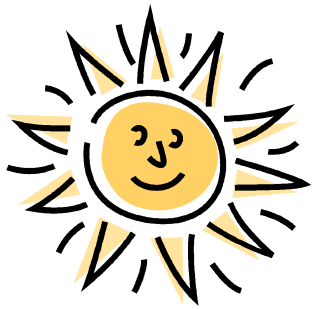

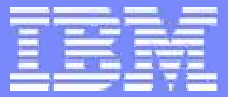

Information Management

BigInt/BigSerial

BigInt and BigSerial Support

- New ANSI standard SQL data types BIGINT and BIGSERIAL introduced in IDS.
- A better performance alternative to existing INT8 and SERIAL8 data types.
- IDS will continue its support of INT8 and SERIAL8 data types.
- Like INT8, BIGINT will store numbers in the range:
  - From -(2^63-1) to 2^63-1.
  - –9,223,372,036,854,775,808 is reserved to indicate NULL
- Like SERIAL8, BIGSERIAL will store numbers range:
  - From 1 to 9,223,372,036,854,775,807 [or 1 to 2^63-1]

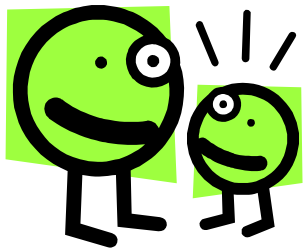

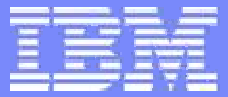

Information Management

IDS 11.5 – Enhancements to DRDA Support

IDS and DRDA

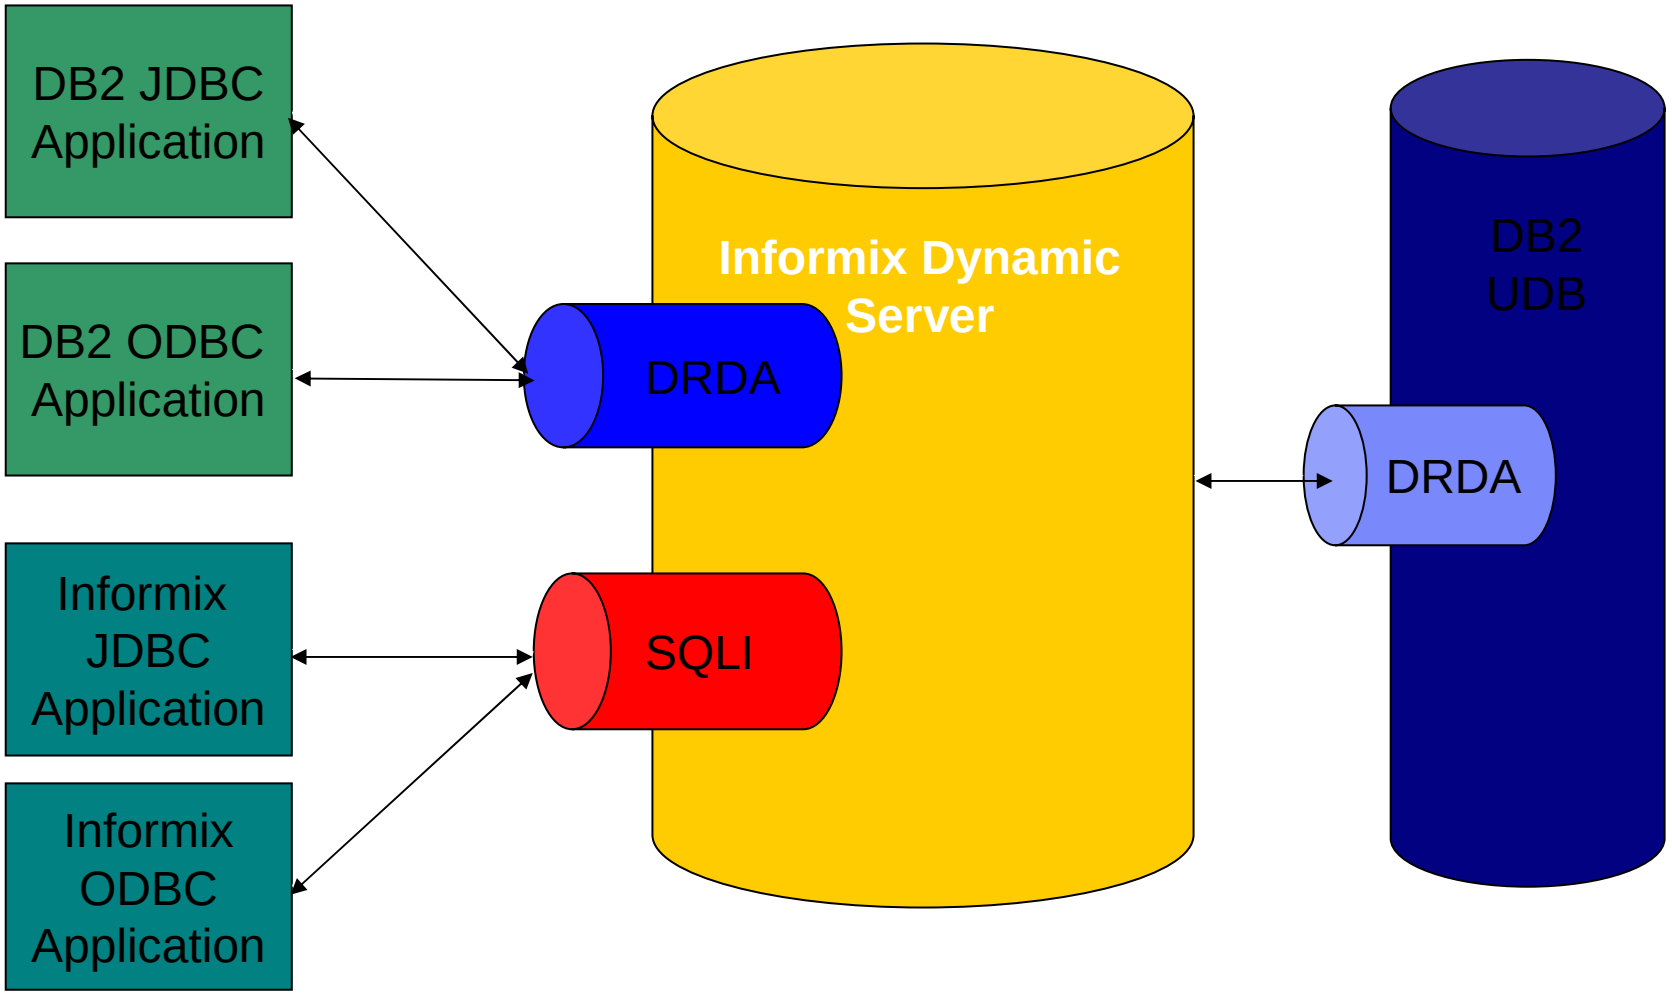

New DRDA Enhancements for Informix in Cheetah 2

- Chaining of multiple statements together supported
- Database names up to 128 bytes
- “is NULL” supported in more places

  Examples :

  #1. WHERE ... (EmployeeID + 100 IS NULL) AND ....
  #2. WHERE ... (my_UDR(c1) IS NULL) AND ...
  #3. WHERE ... <a sub-query> IS NULL AND ...
  #4 WHERE ... (CASE
  WHEN (OrderID IS NULL) THEN CAST(NULL
  AS int) ELSE 3
  END IS NULL) AND …

- Better support for LOBs of all sizes

Benefits of DRDA Support in IDS

- Supports a whole new class of tools and application that previously only worked via DRDA (for example Data Studio)
- Application compatibility for accessing both IDS and DB2 database
- Common Clients, e.g. JCC, .NET, PHP, Ruby on Rails, reduces cost for development = quicker delivery of client features
  - JCC is JDBC 4.0 compliant

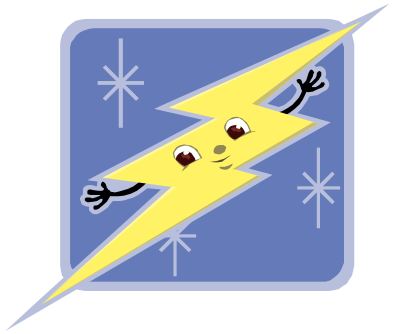

IDS 11.5 Features

Continuous Availability

Updatable secondaries
Connection Manager for Workload Management
Automatic Failover Arbitrator

Admin Free Zone

Auto Update Stats

Admin API to Configure HA clusters

Config options for Defaults during Install

DRDA Config during Install

Improved Onconfig.std

Enhanced Open Admin Tool

Application Development

Dynamic SQL in SPL

Basic Text Search support for XML

XSLT

DRDA enhancements

Mac OS support

Security

Single sign-on via Kerberos

SSL Encryption

Performance

BIGINT and BIGSERIAL

Beyond Cheetah 2

Themes

- Product Integration
- Ease of Use/Application Integration
- Simplify Application Migration
- High Performance OLTP
- Global Availability
- Data Governance

Product Integration

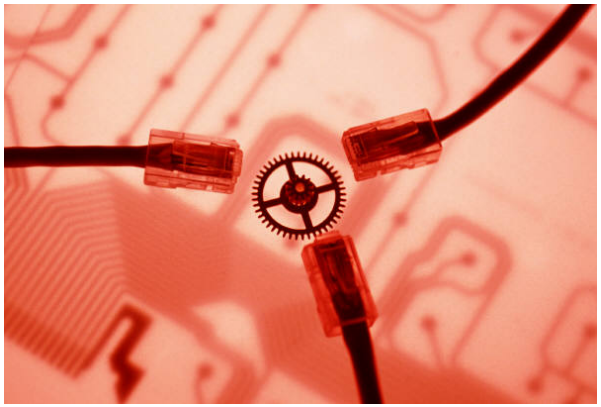

- Solid
- DataMirror
- Princeton Softech – Optim
- Cognos

IDS Business Solutions

- IDS + Solid: The In-Memory advantage
- IDS + Cognos: Enhance Your Business Intelligence
- IDS + DataMirror: Heterogeneous Replication
- IDS + Princeton Softech: IT Governance and Risk Management

IDS + solidDB: The In-Memory Advantage

Provides applications with real-time, all-the-time access to data.

- Extreme Speed
  - Low Latency: Measured in microseconds
  - Volume: Tens of thousands concurrent transactions
- Extreme Availability
  - Uninterrupted data access - 99.999% availability
  - Failover to hot standby in milliseconds
- Low Cost
  - Easy to deploy and administer within applications
  - Embedded database
  - Small footprint – 10MB or less

Cont'd …

IDS + solidDB: The In-Memory Advantage

Accelerates IDS

- Integration of SolidDB and IDS
  - New class of Extreme Transaction Processing
  - SolidDB as in-memory cache/IDS as data store
  - Concurrent execution of low latency, high volume transactions and complex, massive data volume operations
- TM1 Performance Benchmark
  - Baselines established between SolidDB and IDS
- Customer Benchmark
  - Major Telco billing application with subscriber data
- First release of integrated solution in June 2008

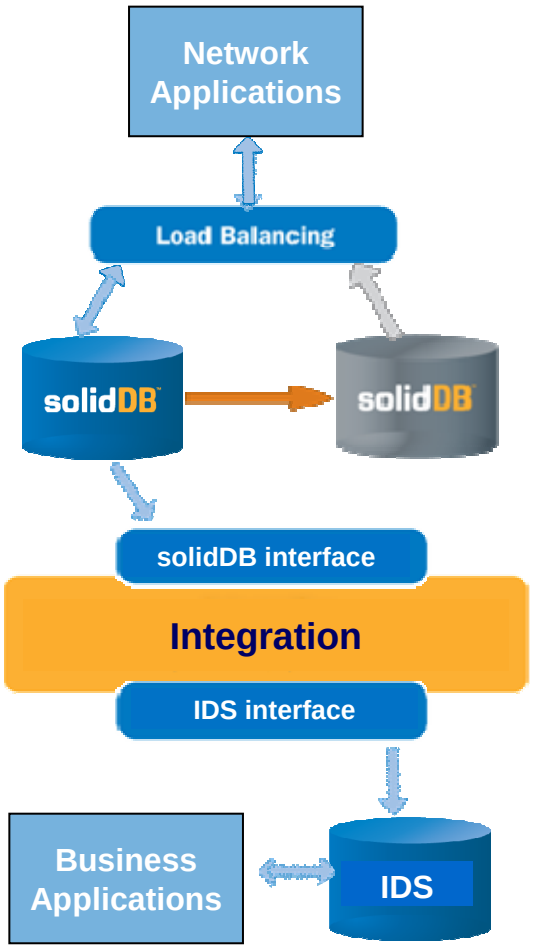

IDS + COGNOS: Enhance Your Business Intelligence (BI)

Complete solution for BI & Performance Management

Industry and Functional Solutions

Complete Coverage of all capabilities

Enterprise-Class SOA Platform

One Platform, One Architecture

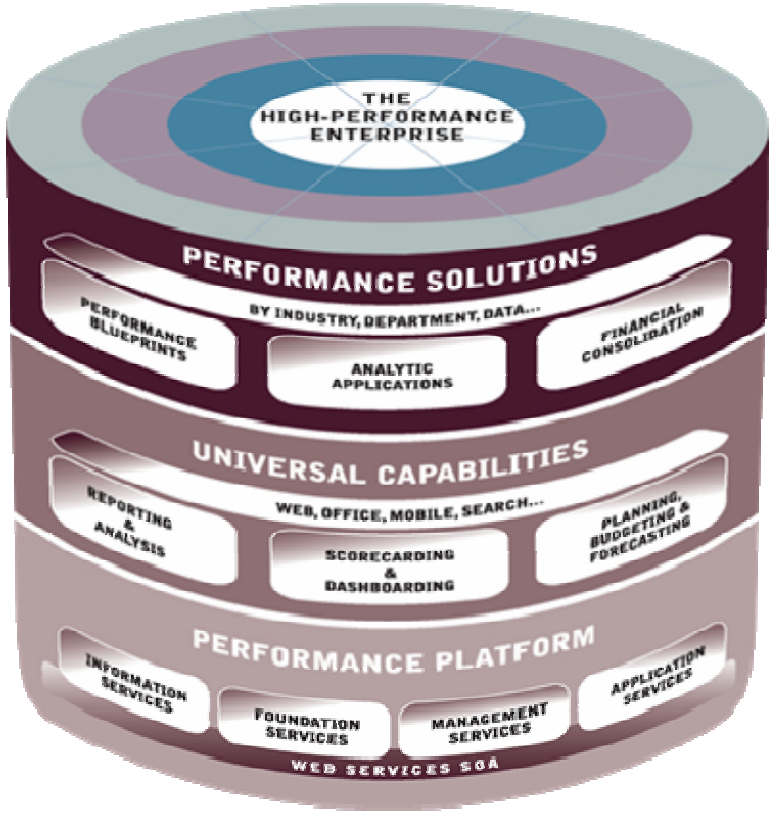

IDS + DataMirror : Heterogeneous Replication

Provides applications with real-time, all-the-time access to data.

- **Transformation Server**
  - **Captures and delivers changed data across diverse data stores in real time without impacting performance**
- **iCluster**
  - **High availability and disaster recovery solution for IBM System-i environments**
- **Changed Data Capture (CDC) can be used for heterogeneous replication**
- **Could be used for Solid integration to IDS for bi-directional replication**

Cont'd ...

IDS + Princeton Softech: IT Governance and Risk Management

Business Parameter Driven Framework to define, extract, access and selectively restore referentially intact subsets of related data from cross-platform, cross-application, RDBMS.