

# Introduction to Informix Dynamic Server

**Stephen Priestley Data Server Technical Sales**

**priestl@au1.ibm.com**



## **DATA MANAGEMENT**

*DATA SERVERS, DATA WAREHOUSE  
AND DATA ARCHIVING (PRINCETON)*



## **ENTERPRISE CONTENT MANAGEMENT**

*INCLUDING FILENET PRODUCT PORTFOLIO*



## **INFORMATION INTEGRATION & MASTER DATA MANAGEMENT**

*PREVIOUSLY KNOWN AS ASCENTIAL,  
DML, SRD & TRIGO.*



## **BUSINESS INTELLIGENCE & PERFORMANCE MANAGEMENT**

*COGNOS*

# Agenda



- **Informix Business Overview**
- **IDS Architecture Overview**
- **Cheetah Features**
  - Key Features
  - Application Development
  - Replication Drill Down (MACH11)
  - Administration Drill Down
  - Optimistic Concurrency



# IBM Data Servers

## Reduce cost of deployment and management of data

- *Innovation to reduce the cost of infrastructure*
- *Innovation to manage the lifecycle of data - from modeling and design through change management and sunsetting*

## Enable rapid use of data throughout the enterprise

- *Innovation that accelerates SOA and XML initiatives*
- *Innovation that leverages Web 2.0 and situational applications*



# Information On Demand with IDS

**1**

**IDS is a strategic member of the IM portfolio that meets a spectrum of data serving needs**

**2**

**IDS is the optimal data server for integrated relational OLTP environments**

**3**

**IDS continues its rich history of technology innovation**

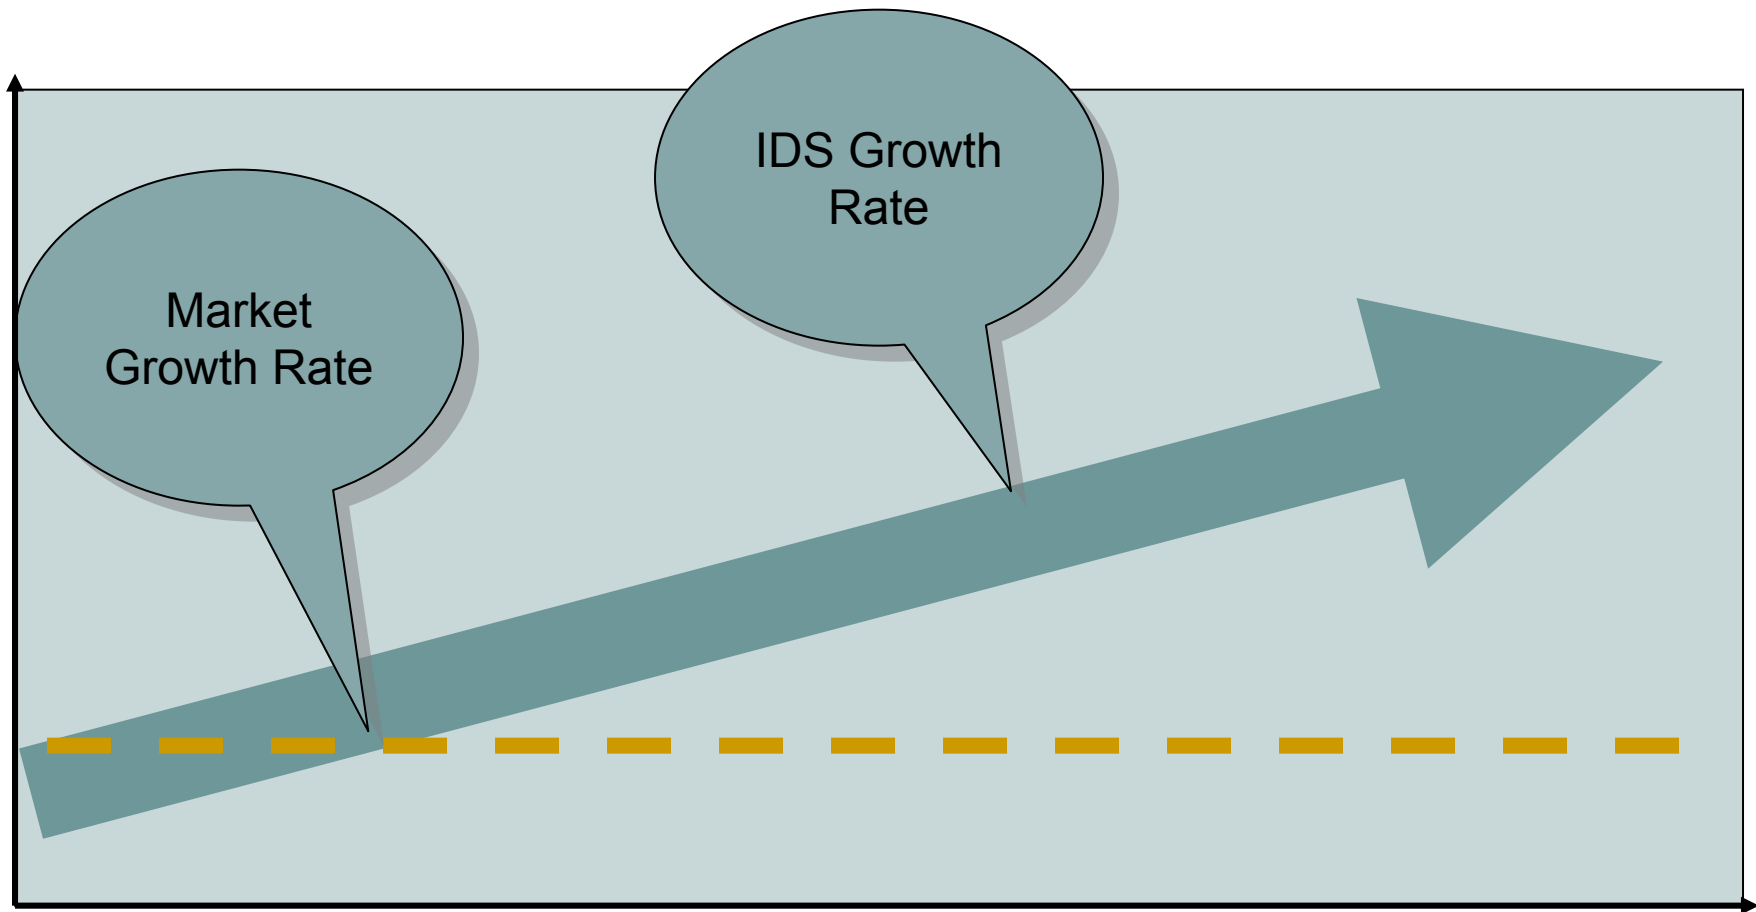
**4**

**IDS is growing faster than the data server market**



# Strong IDS Growth

*IDS is growing faster than the data server market*






# IDS in Action

## *Mission Critical Applications Relying On IDS*

✓ 8 of the top 10 U.S. Retailers




✓ 20 of the top 25 U.S. Supermarkets




✓ 95% of all telco service delivery providers




✓ Every VISA transaction authorization in U.S.



✓ United States "911" Emergency System



✓ IBM badge entry systems (and many others)



# IDS Background



- **Over 20 years in the database industry**
- **One of the strategic IBM databases**
- **Built specifically to handle OLTP workloads**
- **A standard RDBMS with many additional extensions**
  - Allows users to define custom objects and methods
  - Spatial, Timeseries, Text Search, ...
- **Key Advantages:**
  - Ease of Administration
  - Availability
  - Performance
  - Embeddability
  - Extensibility



# Significant Milestones



- **Informix 3.3 released in 1984**
- **Informix-online 5.00 released in 1991**
- **Informix Dynamic Server 7.10 release in 1994**
- **Informix acquires Illustra 1996**
- **Informix Universal Server 9.14 released in 1997**
- **IBM acquires Informix in 2001**
- **Informix releases Informix Dynamic Server 11.10  
July 2007**



# Informix Product Family

## Data Servers

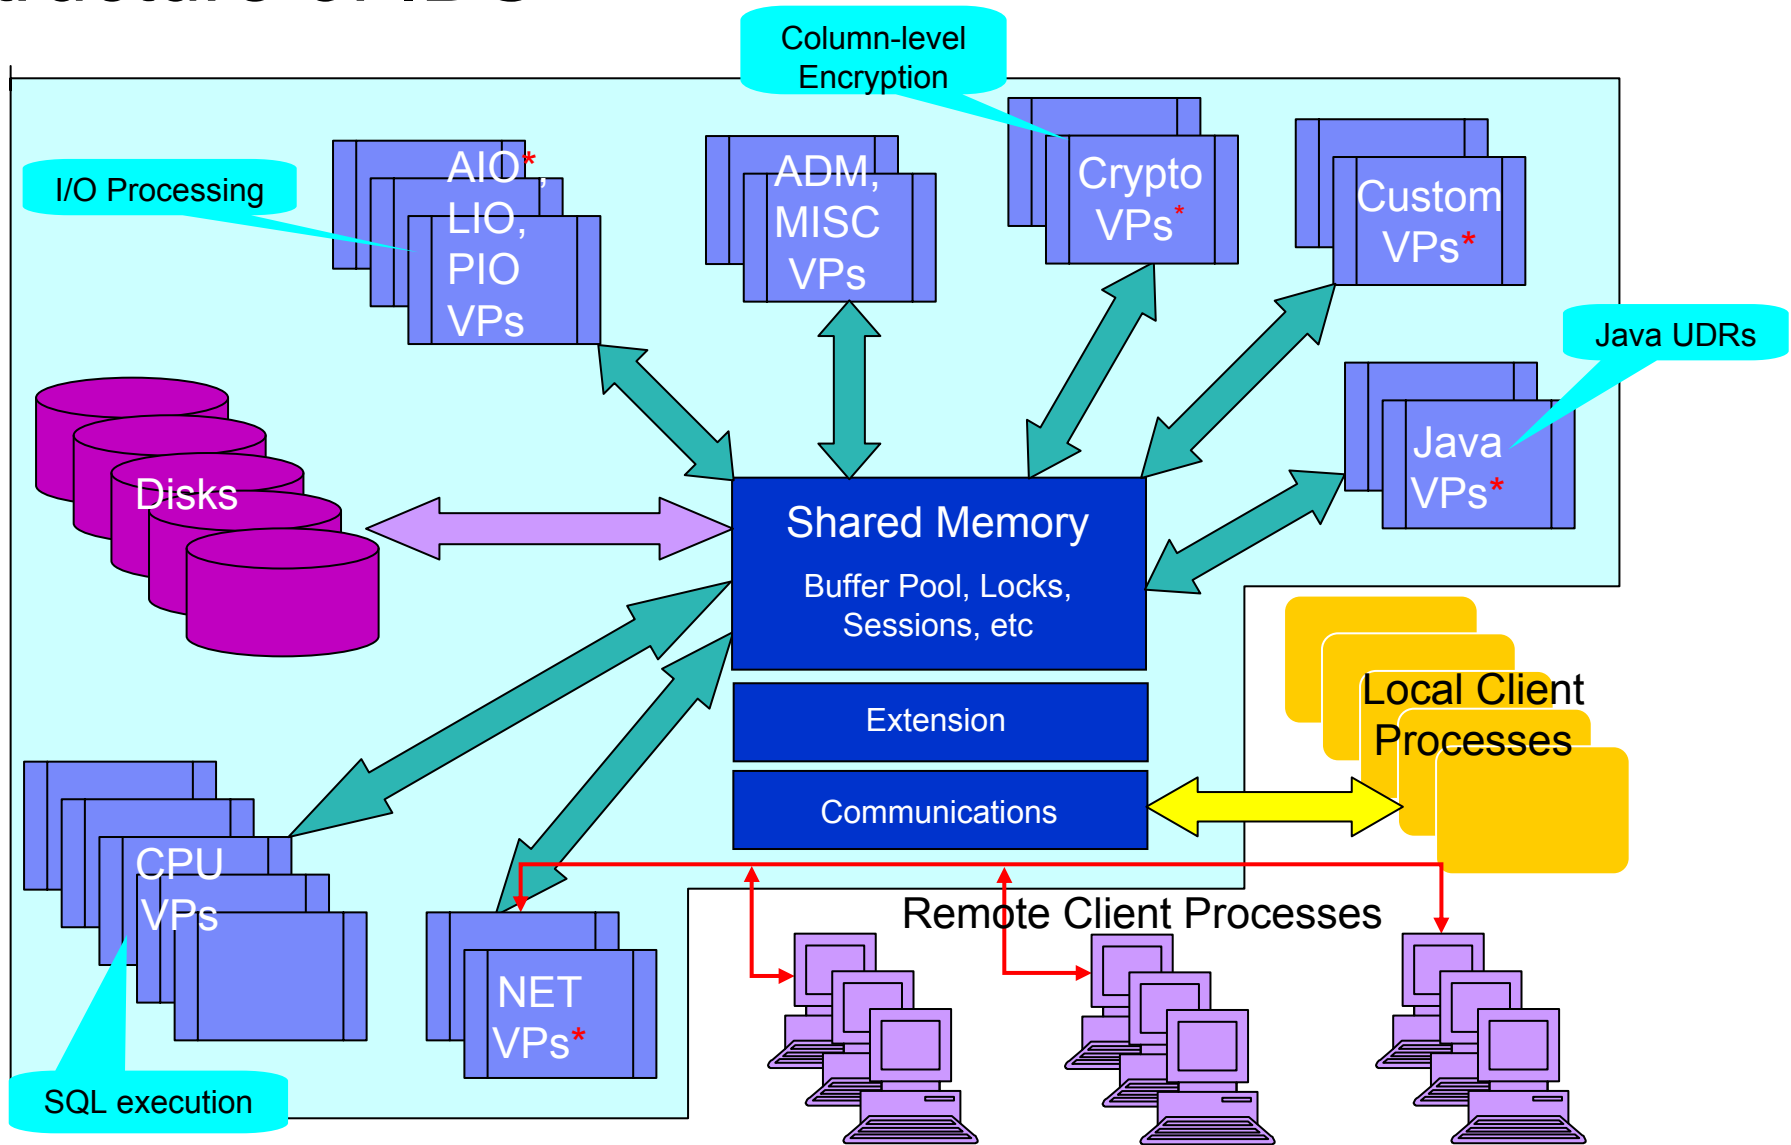
- **Informix Dynamic Server 9.x, 10.x, 11.x**
- **Informix Dynamic Server 7.x**
- **Classics**
  - Informix C-ISAM
  - Informix Standard Engine – SE
  - Informix Turbo
  - Informix OnLine – Version 5.x
- **Informix XPS 8.x**
- **Informix Red Brick 6.x**

## Tools & Connectivity

- **Informix 4GL**
- **Informix ESQL**
  - C
  - Cobol
- **Informix EGM [w/ DRDA]**



# Structure of IDS



# IDS Strengths

## Invisible

- **Simple to use**
  - “Set it and forget it”
  - Comprehensive Administration features
  - Easy to integrate with other applications

## Agile

- **High Performance, highly scalable OLTP**
  - From 1 CPU to 80 or more CPUs on an SMP box
  - Used on Wall St. for high volume market data

## Resilient

- **Continuous Availability**
  - High Availability
  - Capacity Relief
    - Dynamically add/remove additional nodes
  - Workload partitioning
    - Active/Active support



# Simple to Use

Invisible

- **Most administration can be done online**
  - Good for integrated applications
- **Simple administration**
  - New open source tool “OpenAdmin Tool for IDS”
  - Task scheduler built into the server
  - Query drill down simplifies problem determination
- **Operations are simple to setup**
  - HDR/ER
  - Adding additional table space
- **Long history of reliability = less down time**
  - Many customer report no unplanned down time for



# High Performance OLTP

Agile

- **Threaded architecture insures fast process switching**
  - Customer designed thread package
- **Designed to perform well on SMP boxes**
  - CPU VP's can be added dynamically
  - Query threads can migrate between processors
- **Storage management**
  - Non-Blocking disk I/O
  - Intelligent buffer management
- **Memory management**
  - Tunable for large or small memory sizes
- **SQL extensions for operations that most RDBMS's cannot do in the database**
  - Spatial datablade
  - Timeseries datablade



# Continuous Availability

Resilient

- **Workload partitioning**
  - Enterprise Replication: Active-Active
- **Capacity relief**
  - Nodes can be dynamically or removed as needed
- **High availability**
  - Both local and remote failover supported
  - Shared and non-shared disk available
- **Disaster recovery**
  - Log shipping supported
  - Multiple nodes support



# New Features in v11.10 (Released July 2007)



# V11.10 Data Availability



- **MACH11**
  - RSS servers
  - SDS Servers
  - Index Page Logging
  
- **Continuous Logical Log Restore**
  
- **Improved Parallelism during Backup and Restore**
  
- **Backup and Restore to Directories with Ontape**
  
- **Enterprise Replication**
  - Dynamically Change Enterprise Replication Configuration Parameters and Environment Variables
  - Truncate Replicated Tables
  - Performance Improvements for Enterprise Replication
  - Rename Enterprise Replication Columns, Tables, and Databases



# V11.10 Administration

- **Session Configuration Routines**
- **Database Administration System**
  - SQL Administration API
  - Schedule Administrative Tasks
  - Monitor and Analyze Recent SQL Statements
  - Enhanced thread wait information
  - Enhanced Sysmaster
- **Open Admin Tool for IDS**
- **Improved Statistics Maintenance**
  - Automatically Distributions & Statistics Creation
  - Real Time Temp Table Statistics
  - Improved Set Explain
  - Enhanced Support for Large Tables
- **Onconfig Improvements**
  - Change to onconfig
  - Dynamically modifying onconfig
- **Create Schema Enhancements**
- **Installation Improvements on Windows Platforms**
- **Global control of temp table logging**
- **Unbundling of Products**
- **Enhanced Btree Scanner**
- **Enhanced Variable Length Row Compression**
- **Dumping Raw Stacks**



# V11.10 Performance


- **Checkpoint Improvements**
  - Non-blocking Checkpoints
  - Restart Policy
  - Auto LRU Tuning
  - Auto AIO VP Tuning
- **Enhanced Concurrency with Committed Read Isolation**
- **Index Self-Join Query Plans**
- **Improved Concurrency with Private Memory Caches for Virtual Processors**
- **Fast Polling**
- **Direct I/O for cooked files**
- **Enhanced B-Tree Scanner**
- **Global control of temp table logging**
- **Multiple MISC VPs**
- **Faster Sorting**
- **Improved and Predicable Fast Recovery**
- **Improved prepare time with DBA cache**
- **Optimizer Improvements**
- **Enhanced Support for Large Tables**
- **Engine Algorithm Changes**
- **Set Reads for Index Scans**
- **Improved Fragment Elimination, Date Functions**



# V11.10 Application Development



- **Enhanced Data Type Support for Cross-Server Distributed Queries**
- **Named Parameters in a JDBC Callable Statement**
- **Trigger Enhancements**
- **Index Hierarchical Data**
- **Indexable Binary Data Types**
- **Improved Built-in SQL Expressions**
- **Improved Stored Procedure Looping Syntax**
- **Support for Common Clients with DRDA**
- **XML Publishing**
- **Basic Text Search**
- **Deployment Wizard**
- **Derived tables in the FROM Clause of Queries**
- **C Style Comments in dbaccess**
- **Web Feature Services for Spatial Data**
- **Expanded DBINFO functionality**
- **Optimizer Directives in ANSI-Compliant Joined Queries**




# Feature Drill-down: Application Development



# Cheetah Application Development Feature

Agile

- XML Publishing
  - IDS is taking a data-centric approach to XML
  - Functions to generate XML
  - Functions to extract part of an XML document
  - Functions to test for the existence of an element in an XML document
  - Extract, exist and validate functions use Xerces-C
    - Open source parser provided by IBM
  - Requires a new vp:  
`VPCLASS idsxmlvp,num=1`

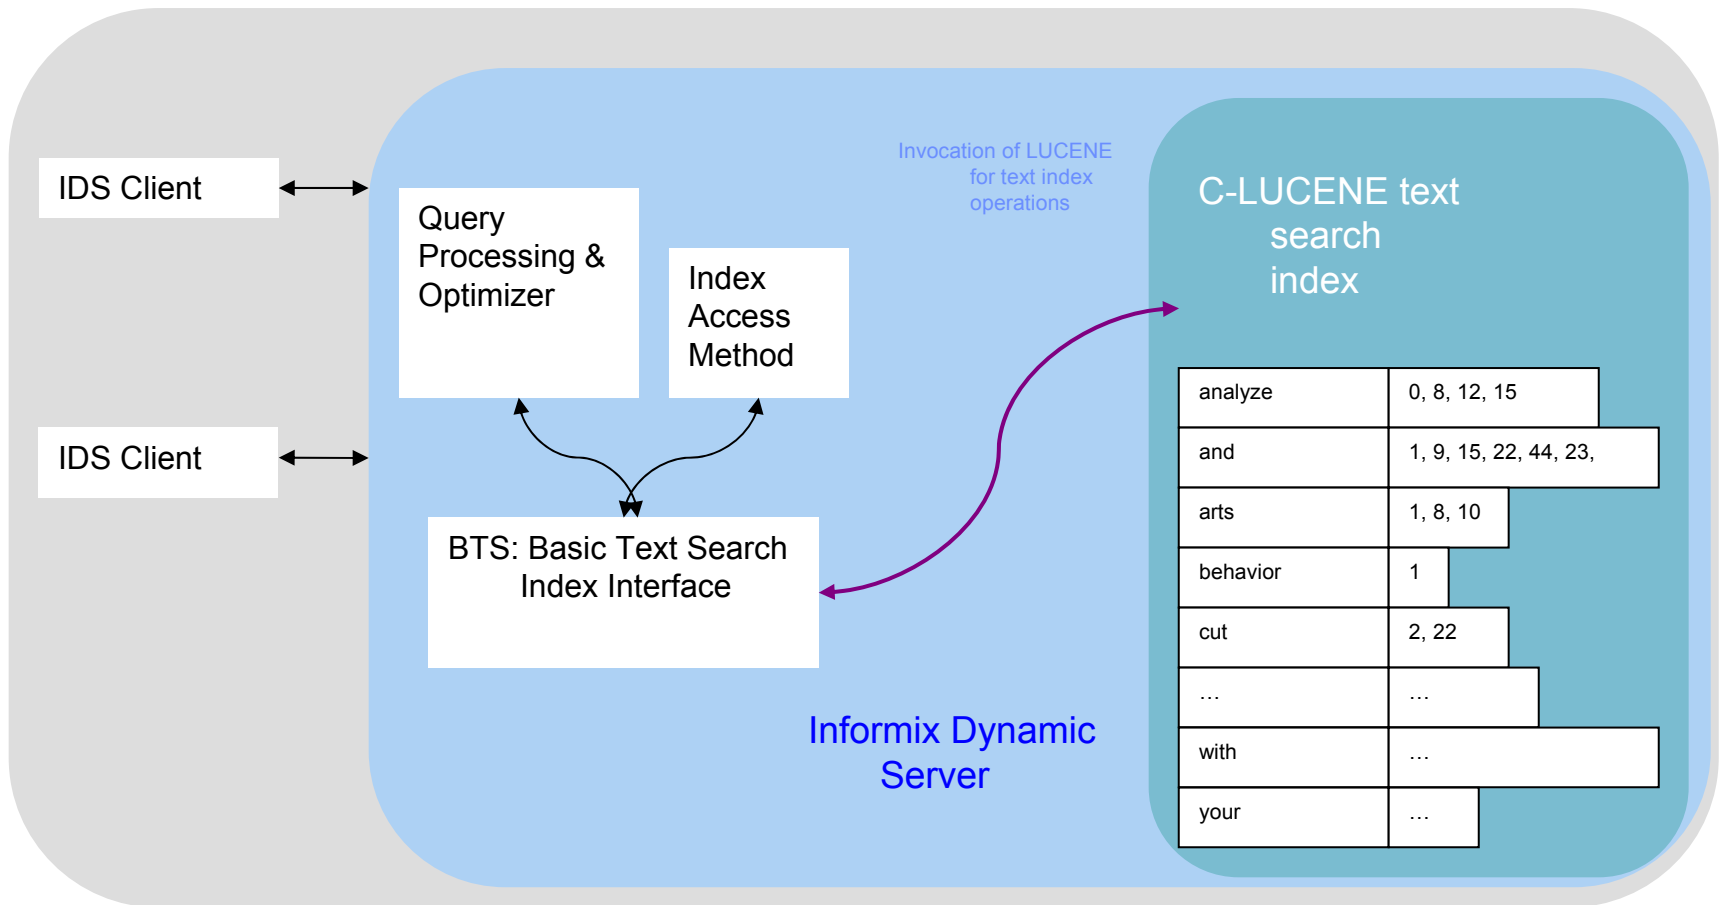


# Basic Text Search

- Text search engine built into IDS - free with IDS license.
- Provides a variety of word and phrase searching on an unstructured document repository
  - Document repository is stored in a character of smart blob column of a table
- Search engine is provided by the open source CLucene text search package
- Supports AND, OR, NOT operators, example:  
**SELECT pid FROM t**  
**WHERE bts\_contains (text\_data, 'foo AND bar') ;**
- Wildcards (“\*”, “?”) and Fuzzy Search (“~”) are supported

# Basic Text Search

Agile



# Handling Hierarchical Data


Agile

- **Indexable hierarchical data type**

- **Add new way to relate objects to one another:**

- `IsAncestor()` ,  
`IsChild()` ,  
`IsDescendant()` ,  
`IsParent()`

- **Can change processing from exponential to linear.**



# Handling MQ Data in the Database

Agile

- **Access WebSphere Message Queue (WMQ) as a table**
  - Functional interface also available
  - Full XA support
- **WMQ 5.3 supported**
- **Limitations:**
  - An instance can connect to only one WMQ manager.
  - IDS and WMQ installed on the same machine.



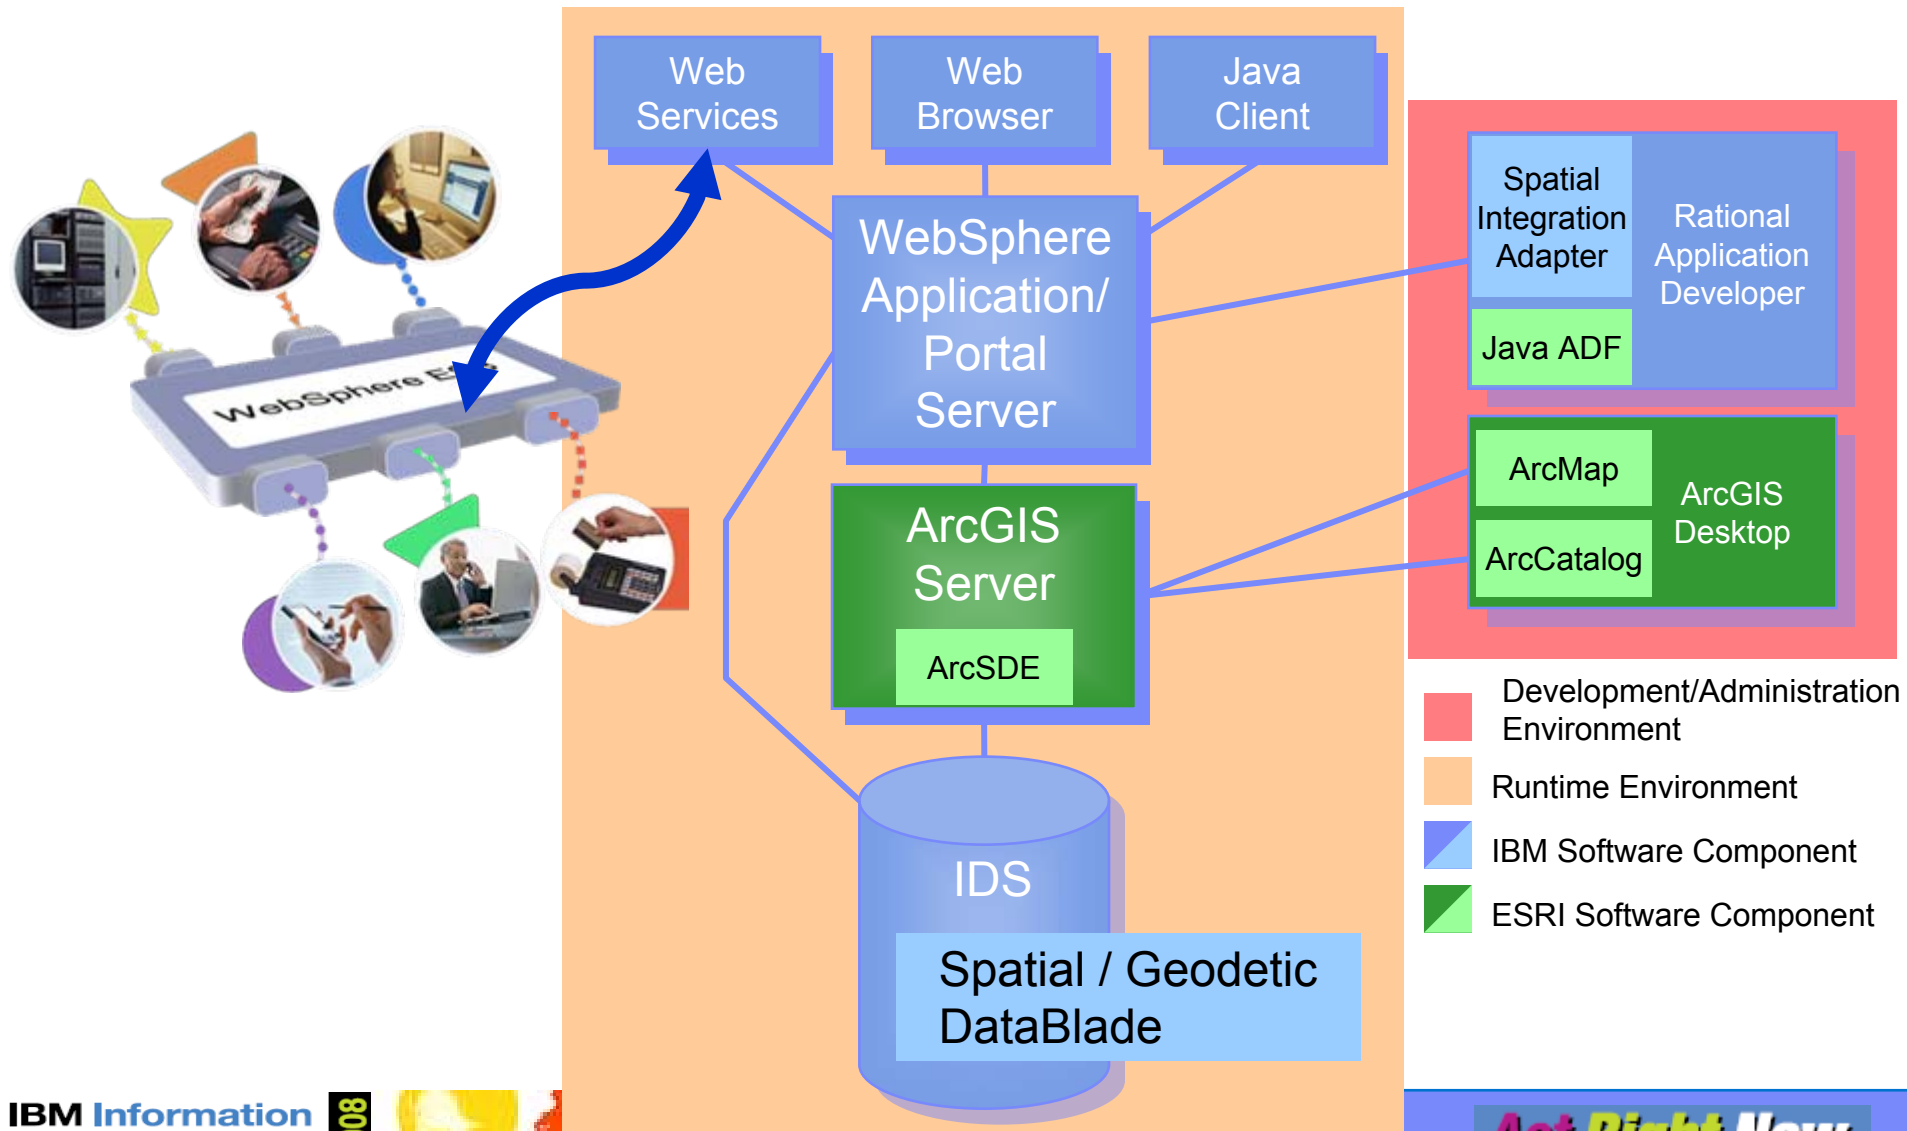
# IDS Web Services Interface in Cheetah

Agile

- **Based on Web Feature Service (WFS) Specification 1.1 from OGC**
  - Uses XML (with GML extensions) to query, insert, update and delete features from a spatial data store
- **Works with Spatial and Geodetic datatypes**
- **Opens IDS spatio-temporal technology to more vendors**
- **SOA Enabled**
  - WFS data can be consumed by many different applications (Java, .NET, PHP, traditional 3<sup>rd</sup> generation languages)



# Spatial Architecture Example





# IDS Install – Deployment Wizard – Why?

Invisible

- **No way to install just a part of IDS**
- **The IDS installation may be too big for some embedded solutions**
- **Manually removing files is not supported by Tech Support currently**
- **Many customers do not use all IDS components**
- **Large footprint of IDS can be a problem for install over the net**



**Installer**

Select the features to be installed.

	Description
<ul style="list-style-type: none"> <li>IBM Informix Dynamic Server               <ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> Base server (Required)</li> <li><input checked="" type="checkbox"/> Database Server Extensions</li> <li><input checked="" type="checkbox"/> Global Language Support</li> <li><input checked="" type="checkbox"/> Backup and Restore</li> <li><input checked="" type="checkbox"/> Demos</li> <li><input checked="" type="checkbox"/> Data loading utilities                   <ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> onunload and onload utilities</li> <li><input checked="" type="checkbox"/> dbload utility</li> <li><input checked="" type="checkbox"/> High-Performance Loader</li> </ul> </li> <li><input checked="" type="checkbox"/> Enterprise Replication</li> <li><input checked="" type="checkbox"/> Administrative Utilities                   <ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> Performance monitoring utilities</li> <li><input checked="" type="checkbox"/> Misc. monitoring utilities</li> <li><input checked="" type="checkbox"/> Auditing utilities</li> </ul> </li> </ul> </li> </ul>	<p>The database server without optional extensions, libraries, or utilities.</p> <p>It requires 82.8 MB of disk space.</p>

InstallShield

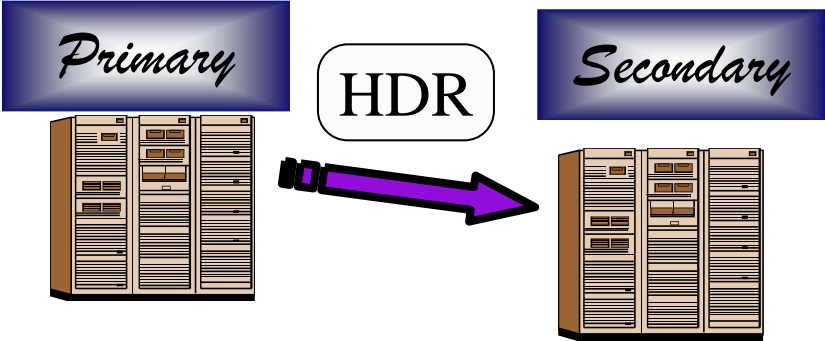
< Back
Next >
Cancel




# Feature Drill-Down: Replication



# IDS Replication Solutions – What we had (pre-v11.10)




- 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***



# New for v11.10 (Code Name MACH-11)




- **Remote Standalone Secondary Nodes (RSS)**
  - Allows multiple HDR-like Secondary Nodes
  - Coordinated with the HDR secondary
  - Functions as a secondary to the HDR secondary
  
- **Shared Disk Secondary Nodes (SDS)**
  - Maintains the buffer pool rather than the disk
  - Secondary uses same disks as the primary
  - Uses LSN advancement rather than locks to coordinate with primary



# 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
- Currently RSS nodes are read-only


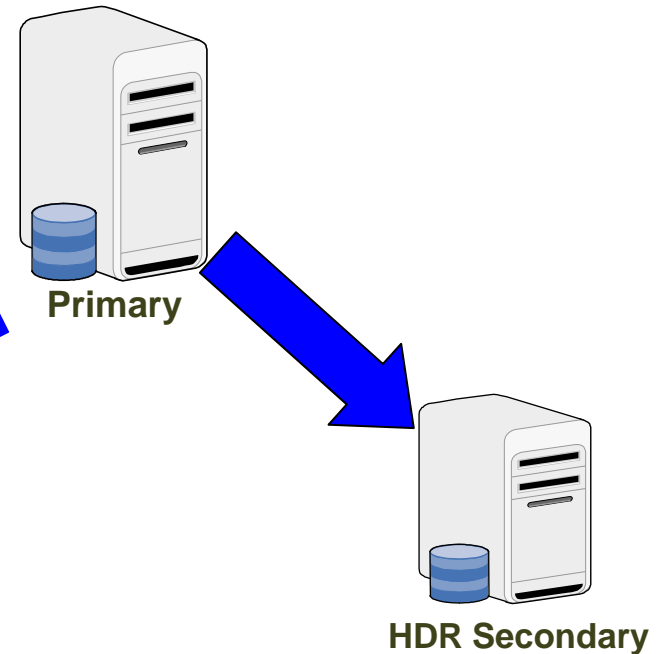
## Replication to Multiple Remote Secondary Nodes



# Usage of RSS – Bunker Backup

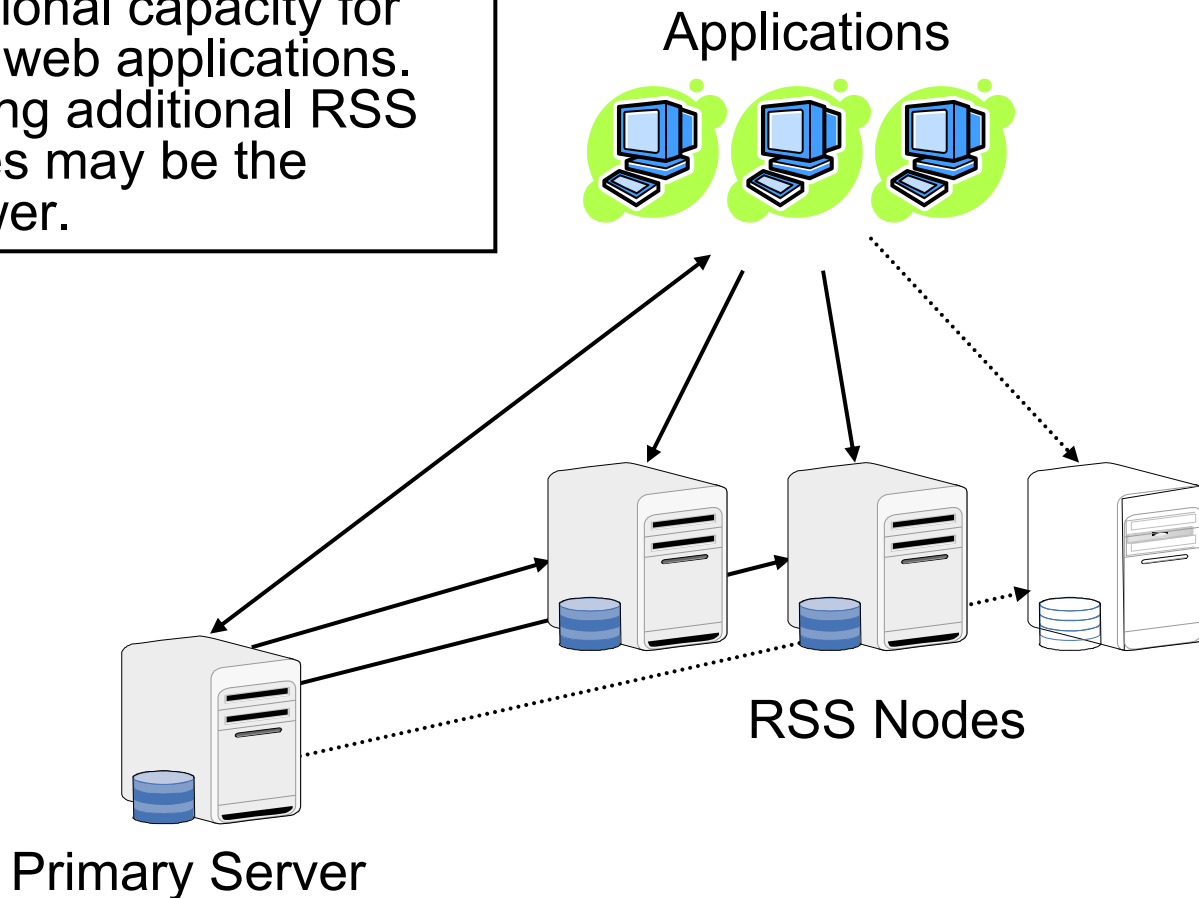
Customer currently has their primary and secondary in the same location and is worried about losing them in a disaster. They would like to have an additional backup of their system available in a remote location for disaster recovery.

Using HDR to provide High Availability is a proven choice. Additional disaster availability is provided by using RSS to replicate to a secure 'bunker'.



# Usage of RSS: Additional Capacity

A customer needs to add additional capacity for their web applications. Adding additional RSS nodes may be the answer.



# HDR: Shared Disk Secondary Nodes

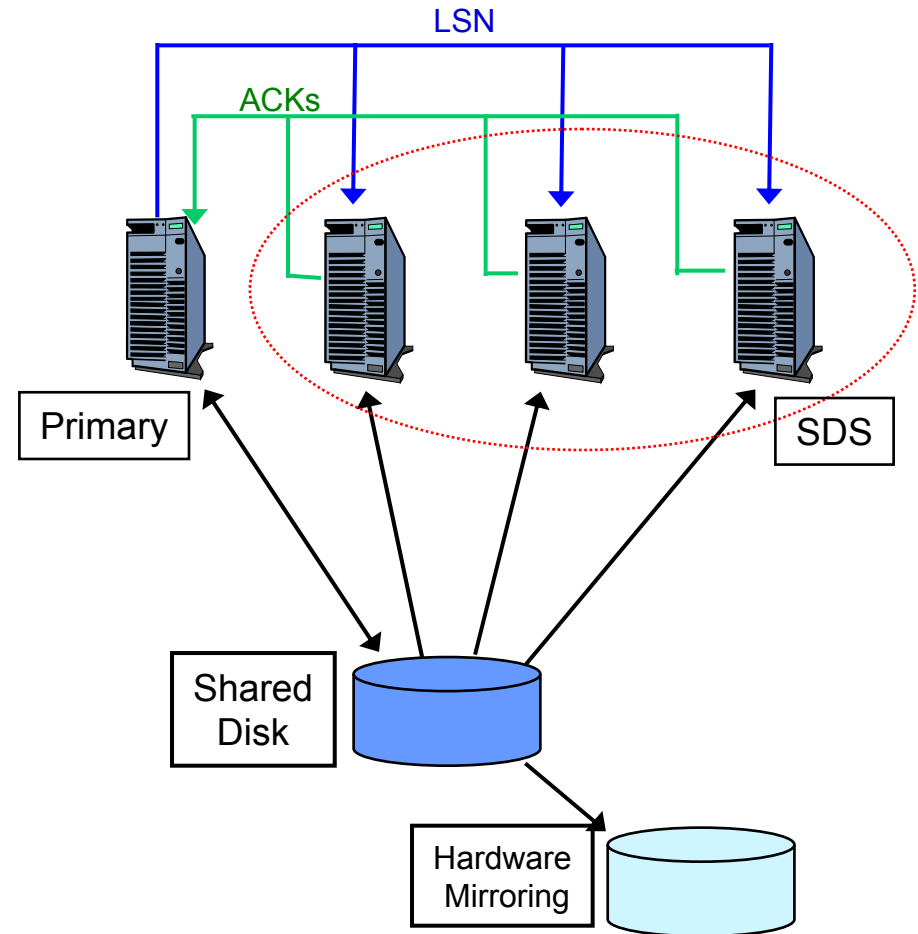
## ■ Benefits

- Online capacity relief
- Very easy to setup
- Does not duplicate disk space

## ■ Introduced in IDS v11.10

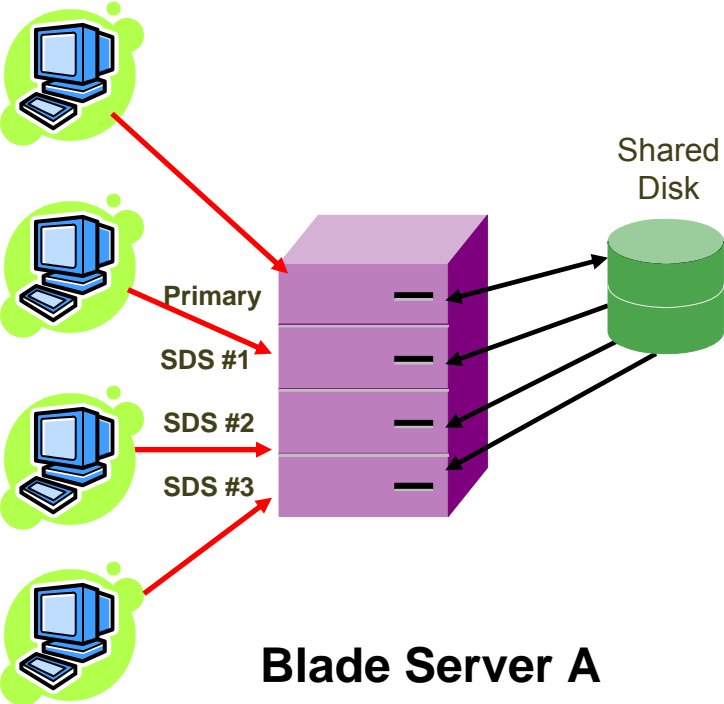
- SDS nodes share disks with the primary
- Can have 0 to N SDS nodes
- Doesn't require any specialized hardware
- Compatible with all other IDS availability solutions

## ■ Currently SDS nodes are read-only




# SDS Usage: Capacity as Needed

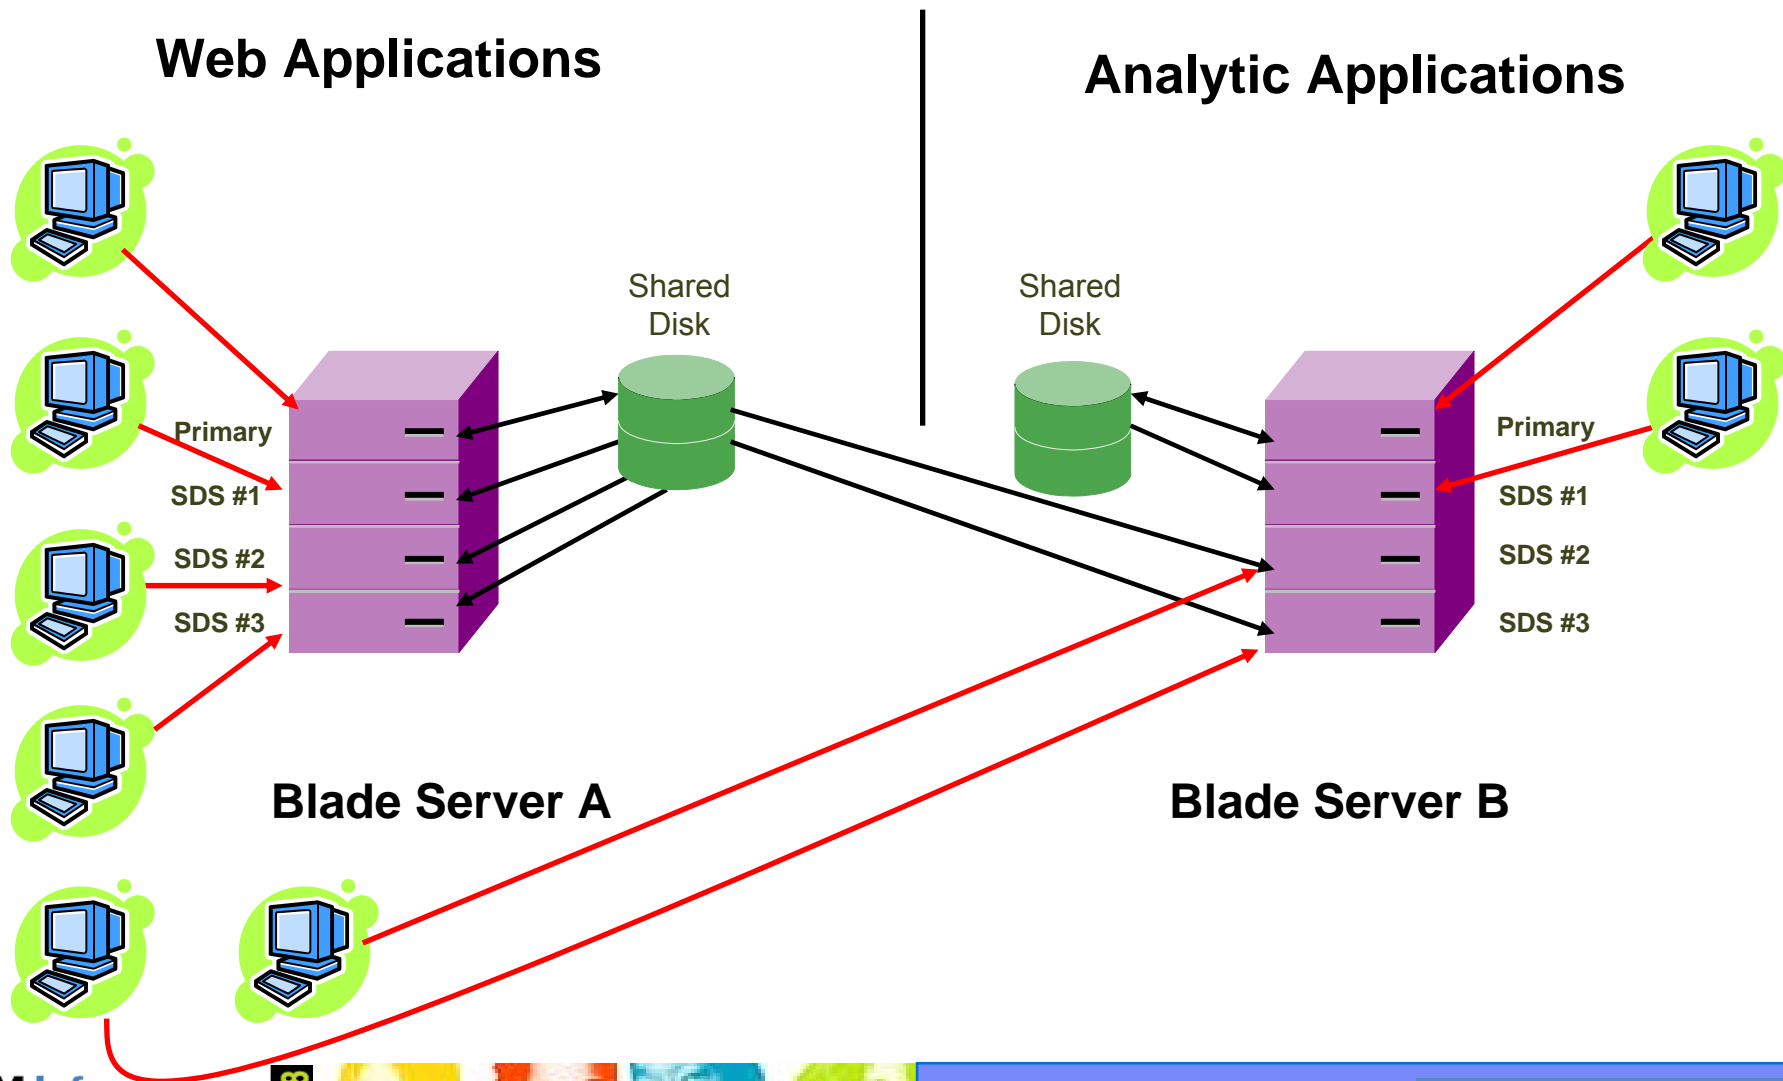
## Web Applications



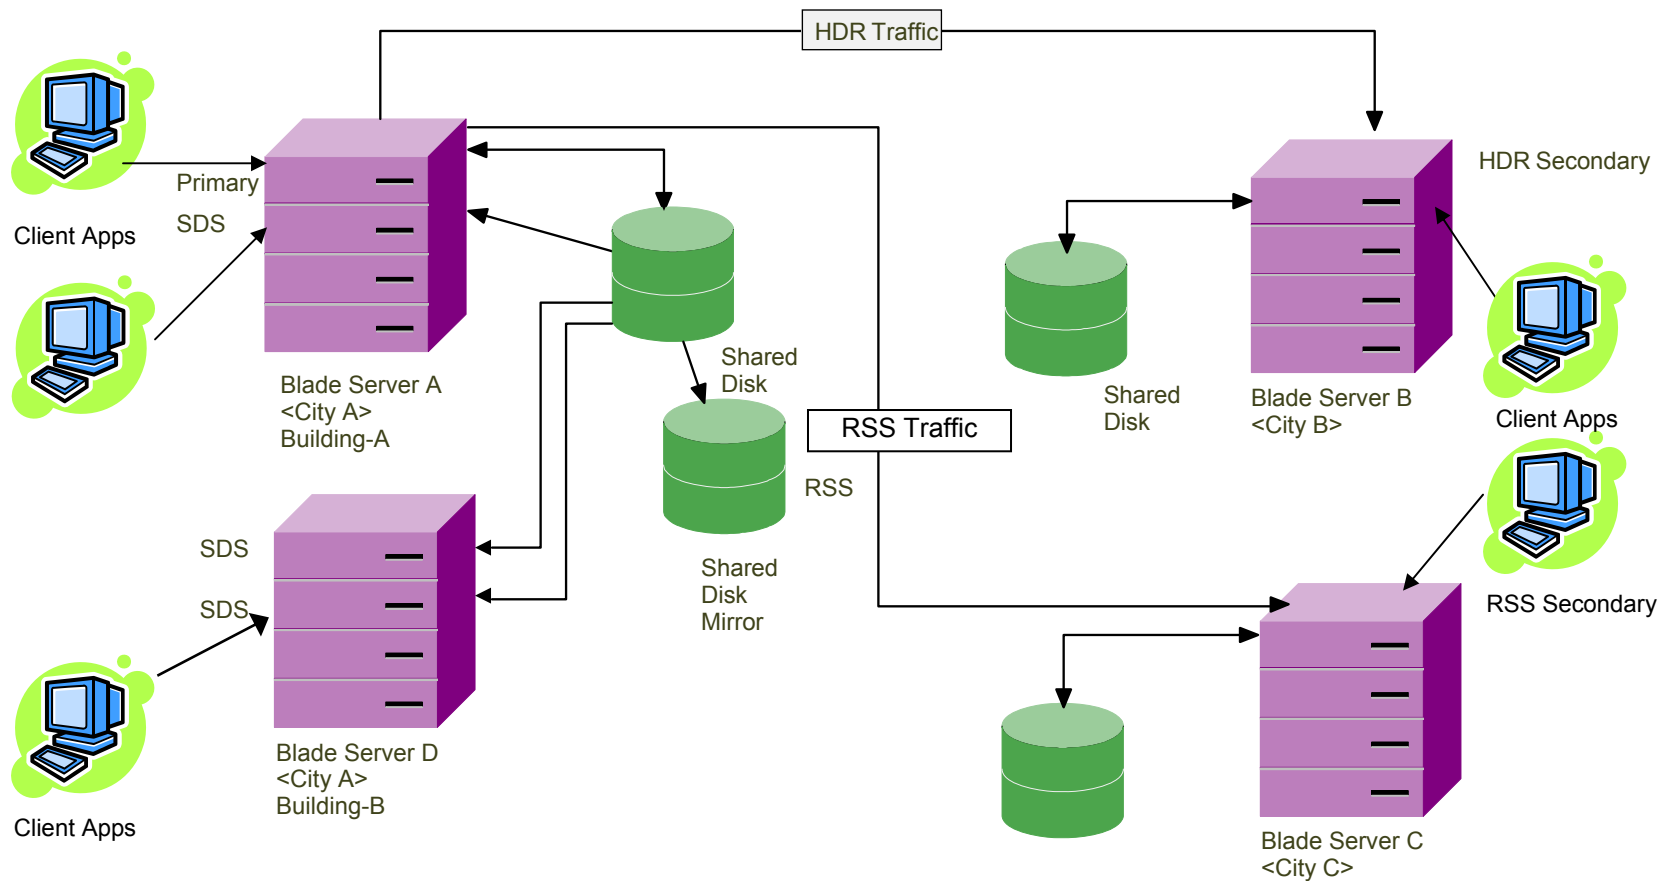
## Analytic Applications



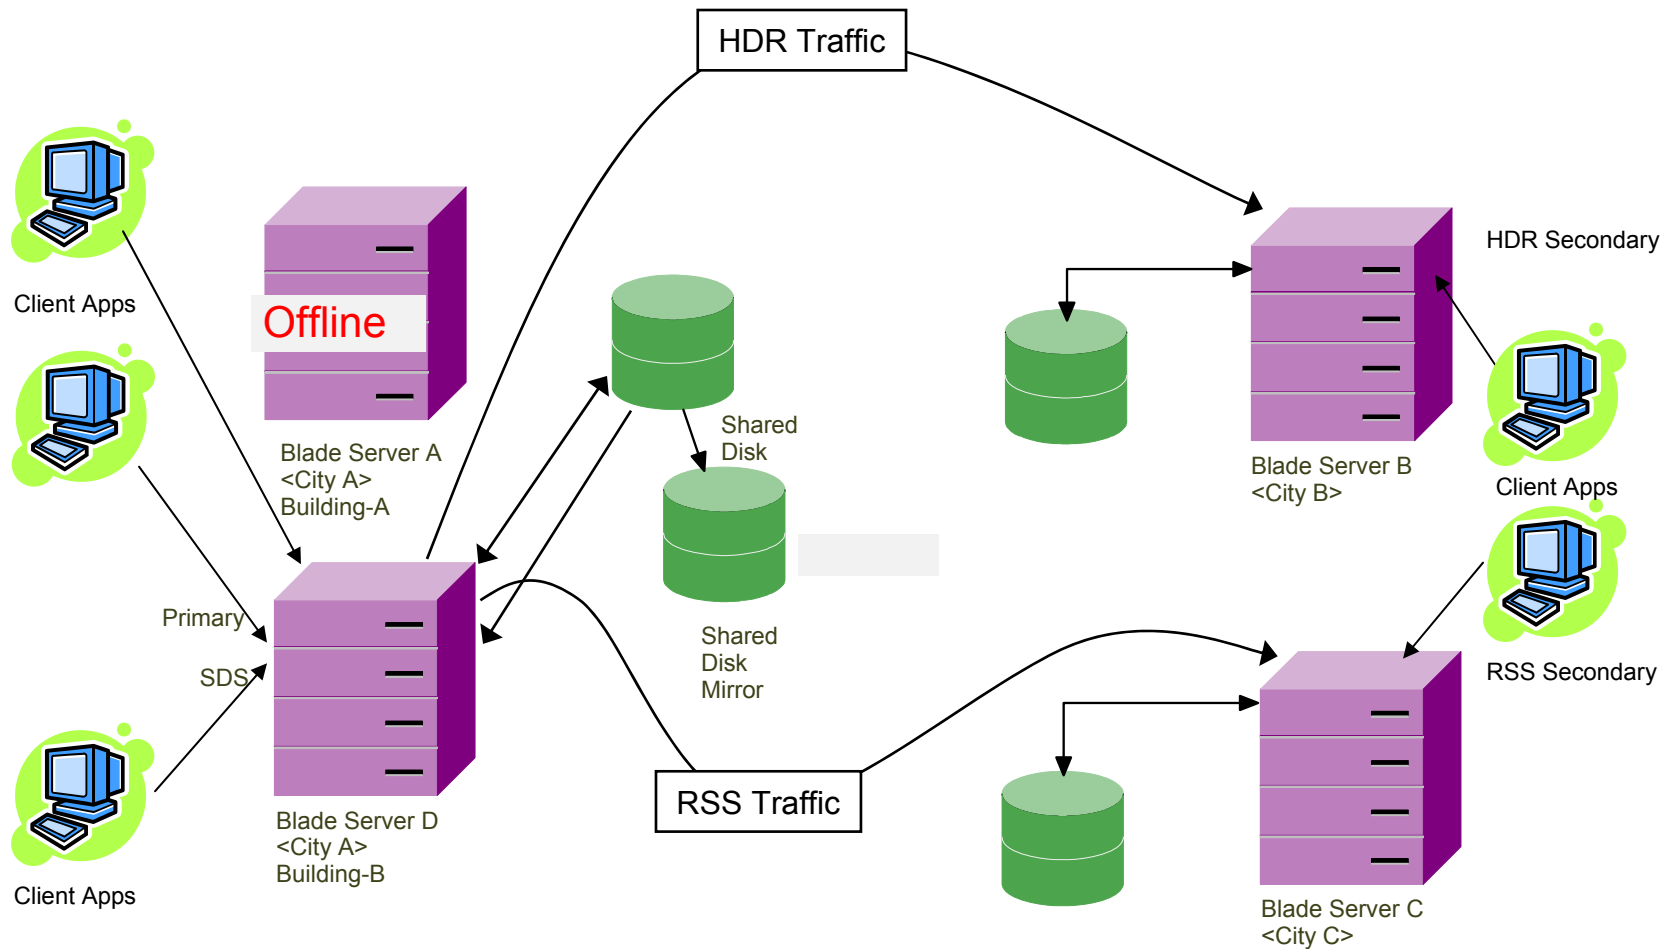
# SDS Usage: Capacity as Needed



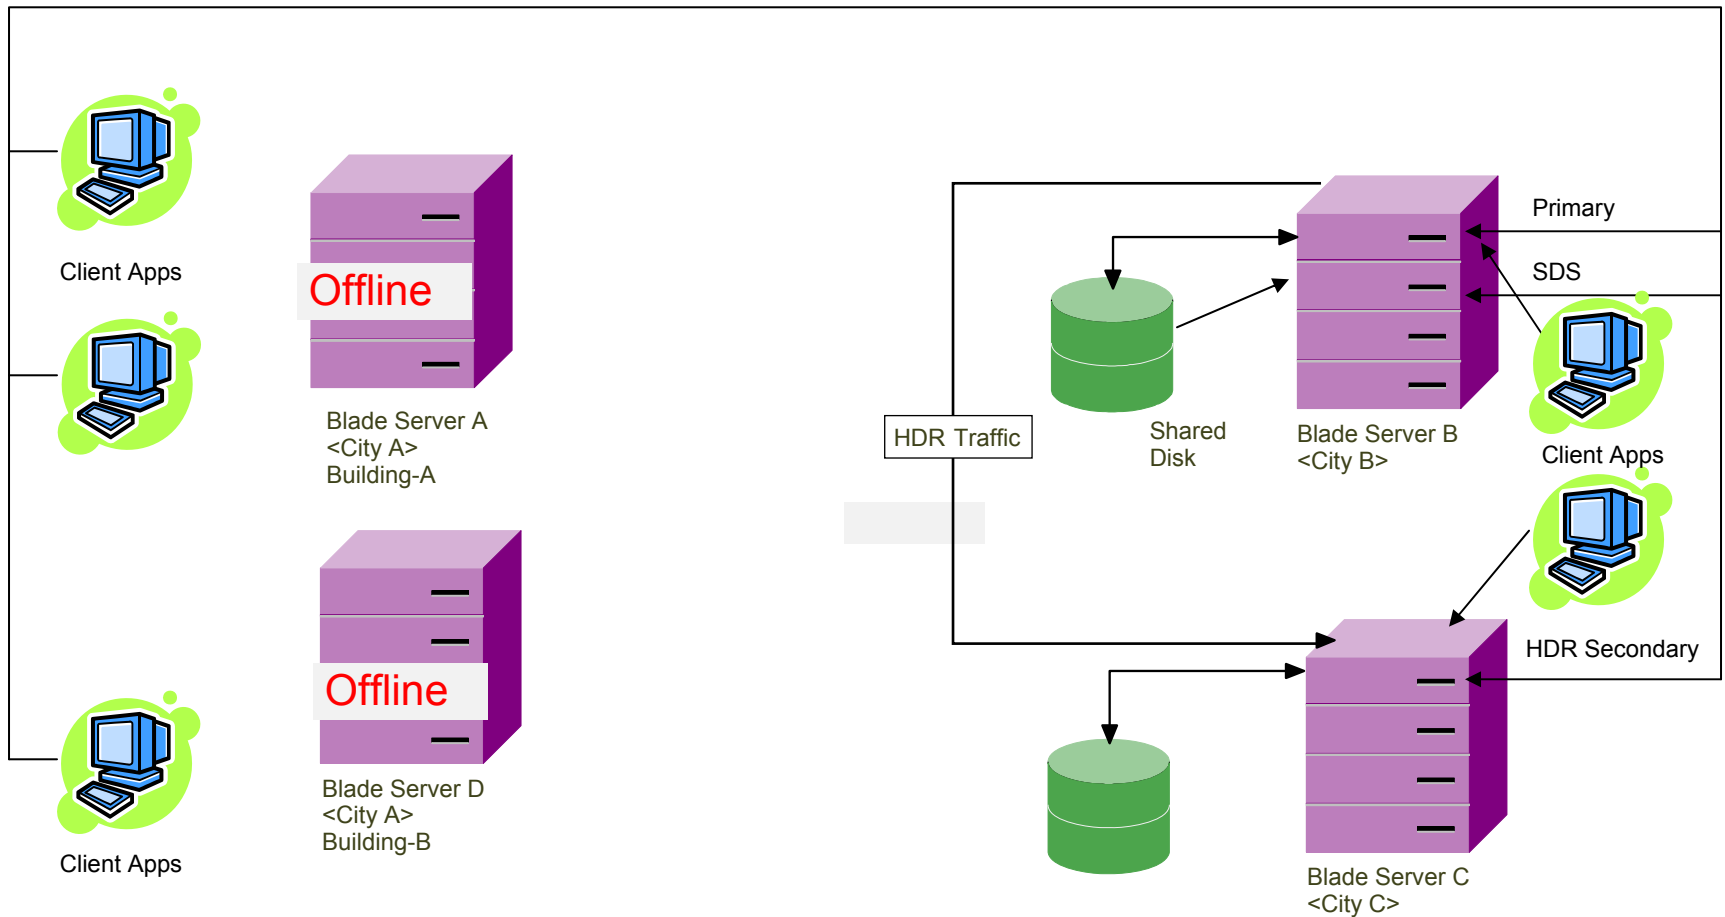
# Availability – The Complete Picture



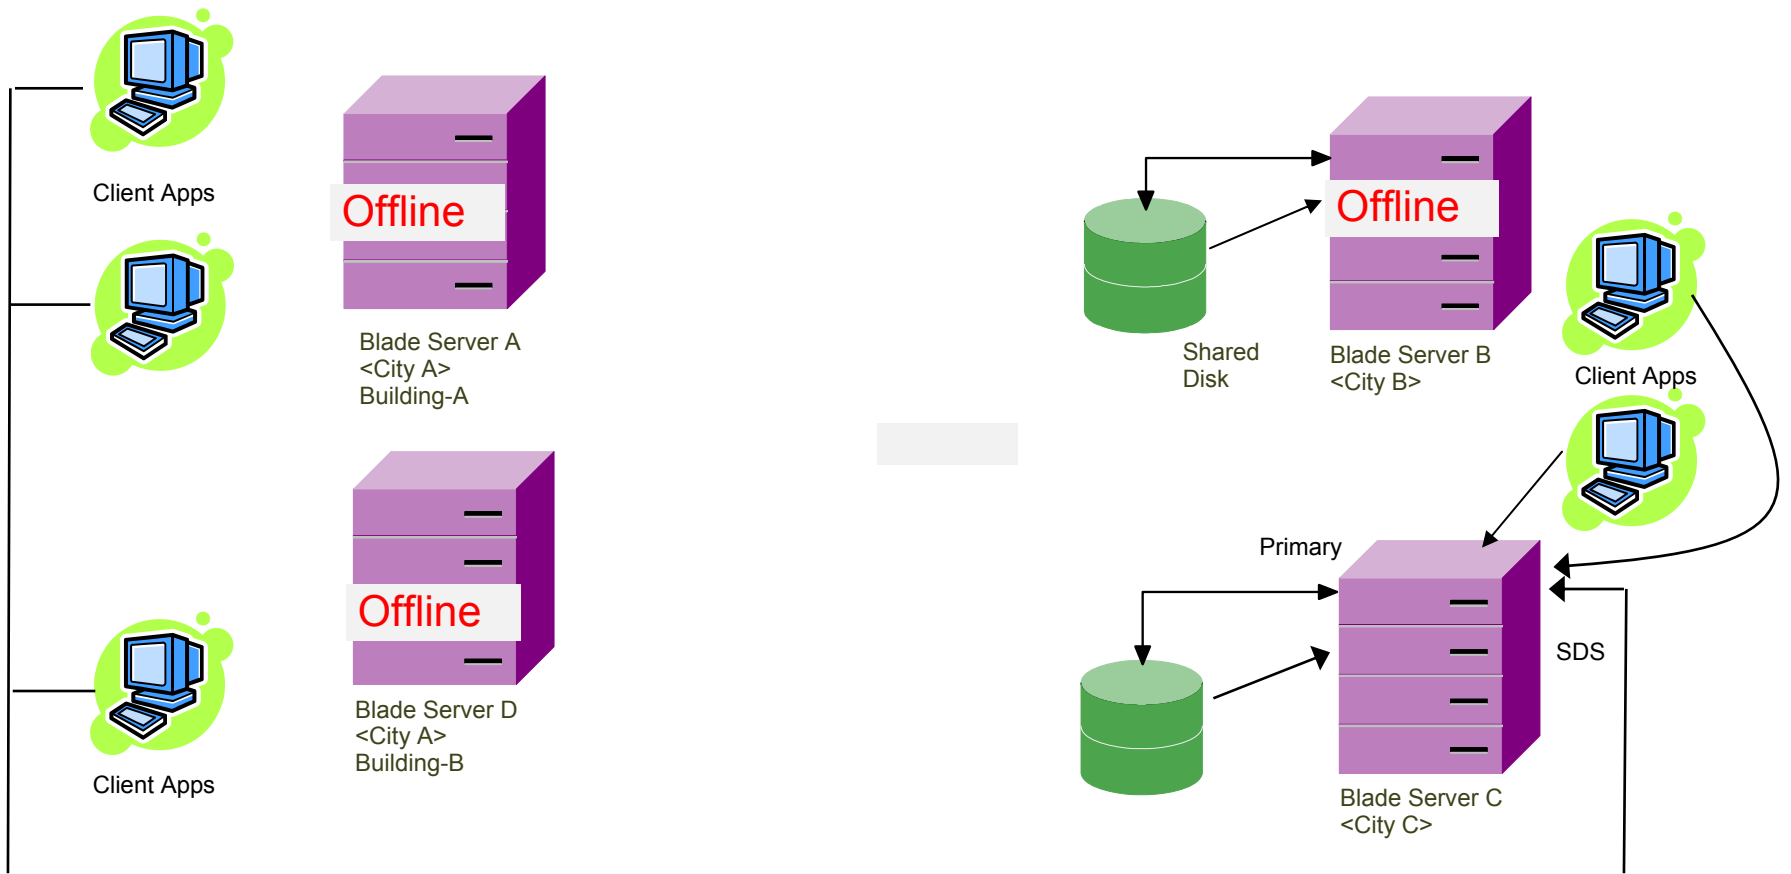
# Availability – After Fire in Building-A




# Availability – Both Buildings in City “A” Go Down



# Availability – City “B” Goes 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.



# Feature Drill-Down: Administration

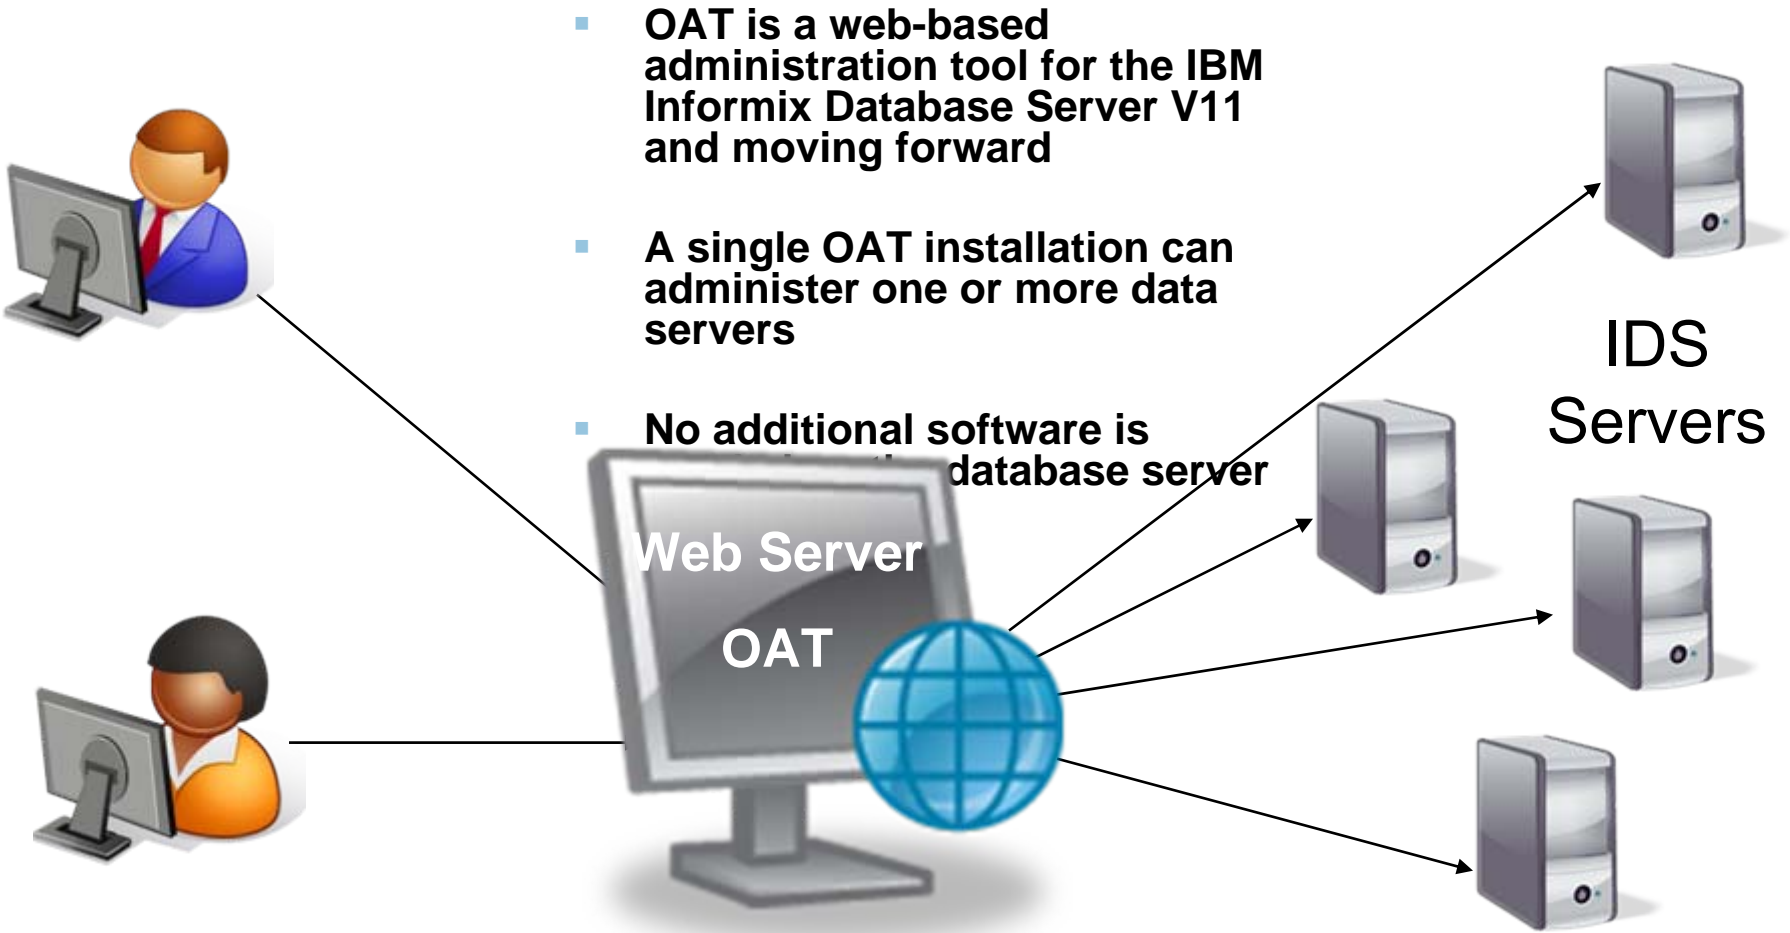


# Solutions Developed

- **New administration tool – OpenAdmin Tool for IDS (OAT)**
- **SQL based monitoring and administration APIs**
- **Database Scheduler**
- **SQL Query Tracing and Profiling**



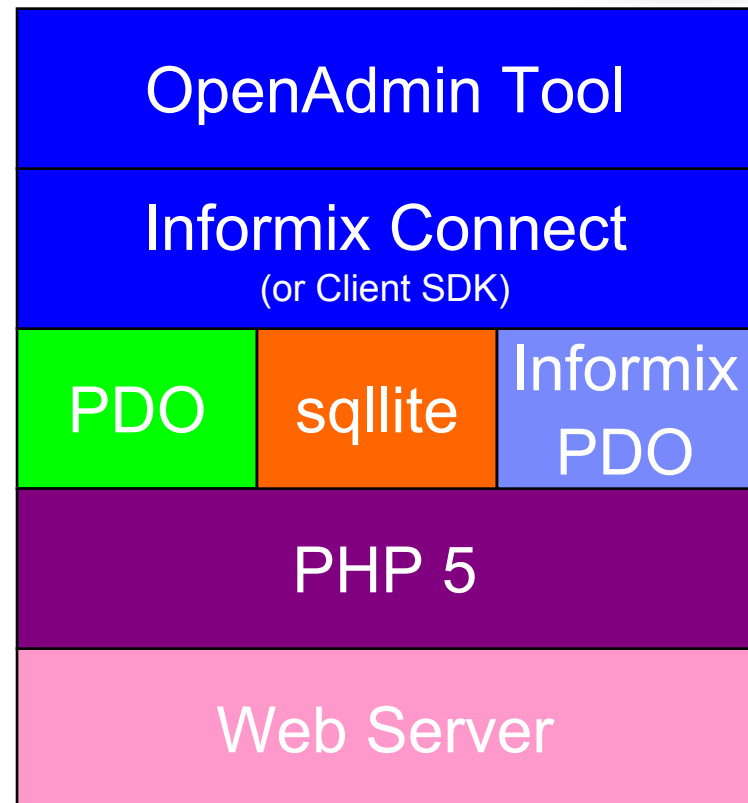
# What is OpenAdmin Tool for IDS (OAT)?



# The OAT Stack



- **Components which are required to run OpenAdmin Tool for IDS**
- **100% Free**
- **Available from various sources as individual pieces**




# SQL Administration API

- A set of User Defined Routines (UDRs) to administer the Informix database server.
- **Feature Benefits**
  - SQL Based Administration
  - Remote Administration
  - Tracking of command execution and results in a system table
- **The major categories of administration include:**
  - Space Management
  - Configuration Management
  - Routine task maintenance
  - System Validation



# Solutions – SQL Query Drill Down Feature

- Provide consolidated detail information about SQL statements through all layers
- Information available through onstat or sysmaster database
- Dynamically configurable
- Global and User Tracing modes



# Find the Slowest Three SQL statements

- Find the 3 slowest SQL statements which were executed against the

```
SELECT FIRST 3 sql_runtime, sql_statement
FROM syssqltrace
WHERE sql_database = "demo"
ORDER BY sql_runtime DESC;
```

```
0.048788974408 insert into mon_table_profile select 6
0.041202886097 INSERT INTO ph_run ( run_id, run_task_i
0.039722568015 INSERT INTO ph_run ( run_id, run_task_i
```



# Built-in Database Scheduler

Problem – Customer have to write their own scripts to perform periodic administration

Solution – A new function scheduler is now available in IDS

- **Can schedule SQL, Stored procedures or UDR's at specific times**
  - Much like a unix “cron” job
- **Two types of jobs can be scheduled:**
  - Tasks
    - Perform actions (check disk space and add chunk if low)
  - Sensors
    - Collect information (record how much disk space is used)
- **System is preconfigured with some tasks and sensors**
  - Users can easily add their own



# Feature Drill-Down: Optimistic Concurrency



# Concurrency: Writers Blocking Readers

- The default isolation level of many DBMS's requires writers to block readers
- Application performance degrades when users must wait for locks before reading data
- Dirty read helps, but applications may get unexpected results
- Deadlocks may also occur which reduce overall throughput



# Cheetah Introduces “committed read last committed”



- **New isolation level introduced**
  - Called “committed read last committed”
  - Designed to increase concurrency & throughput
- **Returns the last committed row**
- **Writers don't block readers, readers don't block writers**
  - Eliminates most deadlocks
- **Simplifies migrating applications to IDS**
  - Default isolation level can be set in the IDS config file so applications don't have to change



# IDS: Then, now, and what's to come

## IDS 10

- Fastest, most-available Informix OLTP engine ever
- Continuous advances to solidify market position
- IDS Express
- PHP driver
- Baan benchmarks

## Business Continuity with Security & Performance IDS 11 "Cheetah" 2007

- High Availability enhancements
- Significant security and encryption enhancements, LBAC and Common Criteria certification
- Enhanced spatial & geodetic web-services for location-based services
- Further reduction in TCO with improved administration functions,
- Advance application development, XML
- Enhanced solutions integration, Admin API, customizable footprint

## Continuous Availability & Scalability 2008\*

- High power scalability and availability solutions
- Enhanced data encryption
- Data compression
- Further integration with common client/tools
- Integrated servers administration
- Macintosh port

## Dynamic Enterprise OLTP Data Server Vnext\*

- Improvements in out-of-box experience
- Online Gaming Support
- Carrier Grade Support
- Additional security management
- Application migration improvements

\* Subject to change



# Questions

<http://www.ibm.com/software/data/informix>

**IBM Information**  
On Demand

2008

>>> Comes To You  
13 MAY - SYD, 15 MAY - MEL



**Act.Right.Now.**

# Thank YOU

<http://www.ibm.com/software/data/informix>

**Stephen Priestley: [priestl@au1.ibm.com](mailto:priestl@au1.ibm.com)**

