



IBM Software Group

DB2® Universal Database Version 8.2 Technical Overview

DB2 Information Management Software

DB2 Worldwide Pre-sales Support
IBM Data Management Solutions

@business on demand software

Disclaimer/Trademarks

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the pages of the presentation:

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: AIX, AS/400, DataJoiner, DataPropagator, DB2, DB2 Connect, DB2 Extenders, DB2 OLAP Server, DB2 Universal Database, Distributed Relational Database Architecture, DRDA, eServer, IBM, IMS, iSeries, MVS, Net.Data, OS/390, OS/400, PowerPC, pSeries, RS/6000, SQL/400, SQL/DS, Tivoli, VisualAge, VM/ESA, VSE/ESA, WebSphere, z/OS, zSeries

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.




Agenda

- DB2 UDB Version 8.2 Technology Overview
 - ▶ Investment Strategy
 - ▶ Product Messages
- Technology Highlights
 - ▶ SQL Enhancements
 - ▶ Security
 - ▶ Microsoft .NET
 - ▶ Business Intelligence
 - ▶ Autonomic Computing
 - ▶ Linux Enhancements
 - ▶ High Availability

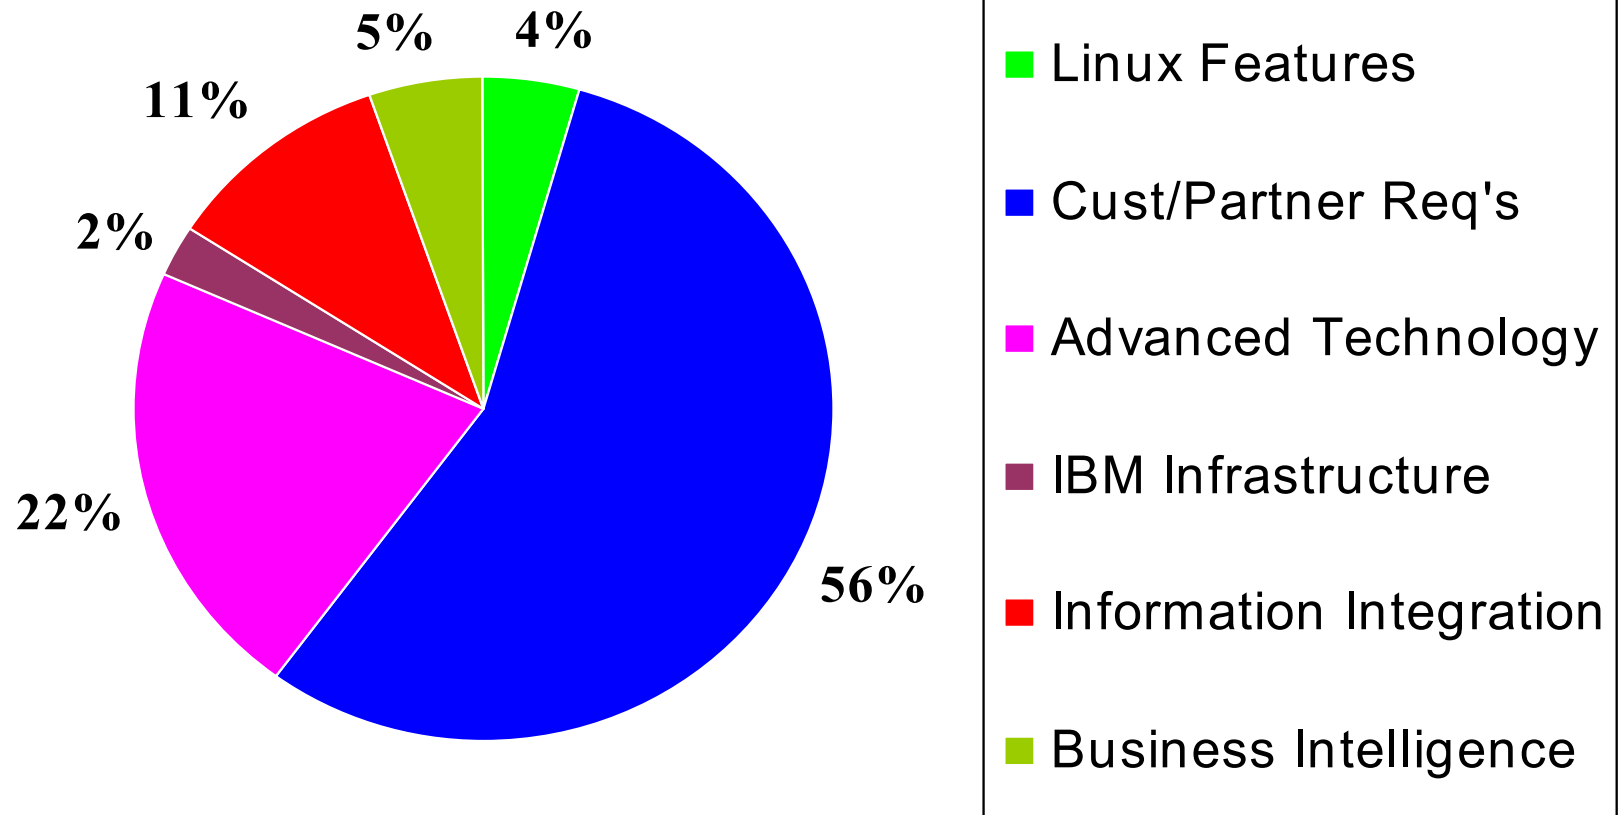


DB2 Technology Strategy Priorities

- Optimize for Multiple Workload Environments
- Deliver Optimized, Transparent Access to All Forms of Digitized Information
- Ease Application Development
- Minimize DBA Skill Requirements
- Deliver High Quality Database Services
- Be the Best ISV Partner



Investment Breakdown




SQL Enhancements

```
with tx1(item, avg_sales) as
(
  select tx_item, avg(tx_quantity) from
    store_txs tablesample system(10)
  group by tx_item
),
tx2(item, avg_sales) as
(
  select tx_item, avg(tx_quantity) from
    store_txs_dup tablesample bernoulli(10)
  group by tx_item
)
select tx1.item, tx1.avg_sales, tx2.avg_sales from
  tx1, tx2
where
  tx1.avg_sales > tx2.avg_sales and
  tx1.item = tx2.item
order by tx1.item
```



Internal Stored Procedures

- Internal SQL PL eliminates the current dependency on a C compiler for the creation of SQL procedures
 - ▶ Eliminating this dependency has been the most consistent PSM-related requirement
- New PSM generates a plan and bytecode to execute in the engine
 - ▶ Performance should be comparable to existing implementation
- Bytecode stored in the catalog and loaded at execution time



Support of CALL statement in a Trigger body

- The CALL statement was restricted to external routines and applications
- In DB2 UDB Version 8.2 , the CALL statement can be used in inline SQL PL
 - ▶ CALL be used within a Trigger, a SQL Table Function, a SQL Method or in a dynamic compound statement
 - ▶ The result of the RETURN statement of the procedure can be retrieved through the GET DIAGNOSTICS RETURN_STATUS statement.



Support for REOPT Bind option

- Occasionally running dynamic SQL queries using parameter markers will cause performance problems
 - ▶ SQL queries could perform poorly during execution if the values used for the input variables (i.e. parameter markers, host variables, and special registers) are outside the predictive range of the default filter factor estimates.
- New REOPT option causes the compilation of the statement is deferred until run-time when the access plan can be optimized using the actual values for the input variables.



Larger SQL Statements

- The 64KB limit on statement size currently limits the total size of statements (e.g. 'CREATE PROCEDURE' or 'CREATE TRIGGER') which therefore limits the size of the object.
- This limit has been increased in order to allow partners who have the majority of their application logic in stored procedures or triggers, or need to use other large statements.
 - ▶ Currently partners have to break up the procedures and make them more complex in DB2 than in other RDBMS
 - ▶ In some cases, it is not possible to migrate a trigger or psm statement from other RDBMS to DB2 with DB2's 64K maximum statement size
- The limit will be increased to 2MB.



Support SET LOCK WAIT / NO WAIT

- Traditional locking that is used to implement transaction isolation level can result in applications blocking each other
 - ▶ This happens when one application must wait for another application to release the lock on a particular table, row, or index key value.
- Strategies to deal with the impact of the blocking usually provide a mechanism to specify the duration of the block.
 - ▶ the lock wait mode timeout for an application/session can be set not to wait, to wait indefinitely, to wait for a specified amount of time (up to 32767 seconds), or to reset back to the value of database configuration


```

SET LOCK MODE TO --+-- NOT WAIT-----+---|
                |                               |
                +-- WAIT -----+-----+
                               |             |
                               +-integer-constant--+
    
```



Toggle Generated Column property

- New options for changing the generated values for a column
- Alter Identity column
- Alter Generated Expression Column
- Alter Column with user-defined default value
- Drop <generated/default> SET <generated-default>
- ALTER the default value of a column to a different value
 - ▶ Existing values in the database that were set to the previous value will not be changed



Nested Savepoints

- This feature expands the support for external savepoints first introduced in DB2 UDB Version 7 by providing a number of new abilities:
 - ▶ The ability to nest savepoints within another savepoint
 - ▶ The ability to use savepoint-related statements within atomic compound SQL statements
 - ▶ The ability to use atomic compound SQL statements within an active savepoint
 - ▶ The ability to create a new savepoint level within a stored procedure invocation
 - ▶ The automatic creation of a new savepoint level within each call to a user-defined function (UDF)


```
create table t(i int);
insert into t values (1);
savepoint sp1;
insert into t values (2);
savepoint sp2;
insert into t values (3);
rollback to savepoint sp1;
insert into t values (4);
```

Security Enhancements



Data Encryption

- Encrypts user data during client/server communication over the network
- Two new authentication types
 - ▶ SQL_AUTHENTICATION_DATAENC
 - connections must use the data encryption
 - ▶ SQL_AUTHENTICATION_DATAENC_CMP
 - allows for a compatibility mode with down-level products that do not support the new authentication type
 - they will be allowed to connect with SERVER_ENCRYPT and not encrypt user data
- Encrypts all sensitive data




Common Criteria Certification

- DB2 will certify under the Common Criteria against the US Controlled Access Protection Profile (US-CAPP 1.d) at assurance level 4 (EAL4)
- Testing procedure underway with the following material being developed for certification:
 - ▶ New book focused on "security administration" (user and admin guidance)
 - ▶ Additional documentation changes to reflect the "evaluated configuration"



DB2 Security Exit

- DB2 does not have its own mechanism for maintaining userids and passwords or userid group memberships
 - ▶ All authentication in DB2 is managed either through the underlying operating system or an external security system such as Kerberos
- This feature will allow the customers to create their own authentication mechanisms to handle:
 - ▶ group membership
 - ▶ authentication on the client side
 - ▶ authentication on the server




Windows Domain Handling and Userids

- Improved Windows Domain handling support
 - ▶ Support using cached credentials for DB2's authentication and group lookup
 - ▶ Support nested group semantics.
 - ▶ Support domain local group
 - ▶ Support implicit trusts between domains
- Userid Enhancements
 - ▶ Accepting additional special characters in userids and security mechanism group names (!%&(){}-.^~ and space)
 - ▶ Accepting security mechanism group names that are longer than 8 characters
 - ▶ Allowing two-part names on CONNECT and ATTACH that contain a Windows domain name and the userid to avoid the network traffic associated with looking up the user name in the trusted domain forest

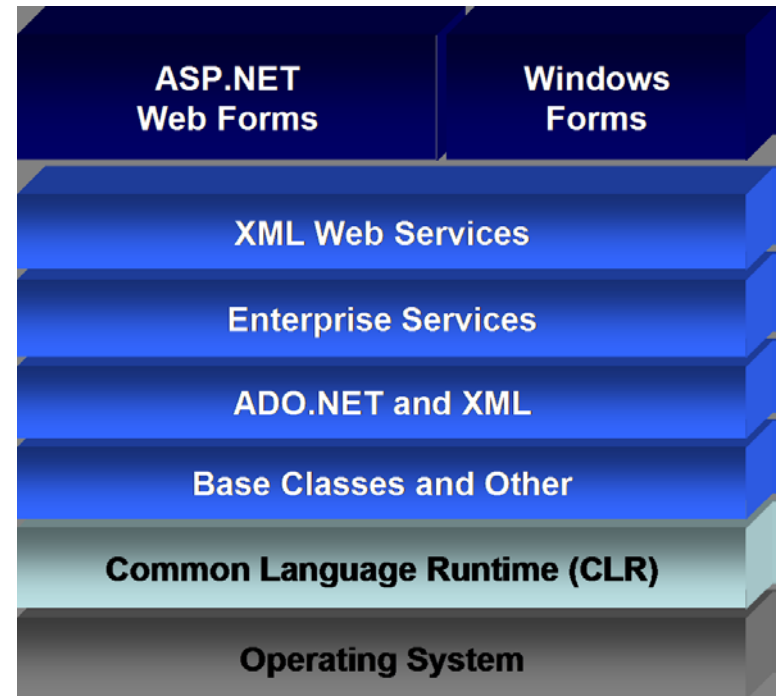


Local System Accounts (Windows)

- Support for Windows Local System Account (LSA)
 - ▶ Allows DB2 applications (including DAS scheduler, DWC, and customer applications) to run under LSA when accessing DB2 rather than forcing the user to provide a separate account
 - ▶ Ability for DB2 to authenticate customer applications running under default LSA which appear to DB2 as user "system"
 - ▶ Treat Local System Account ("system") as a Local Administrator Account -- any authorization checking done by DB2 should recognize LSA as system administrator.

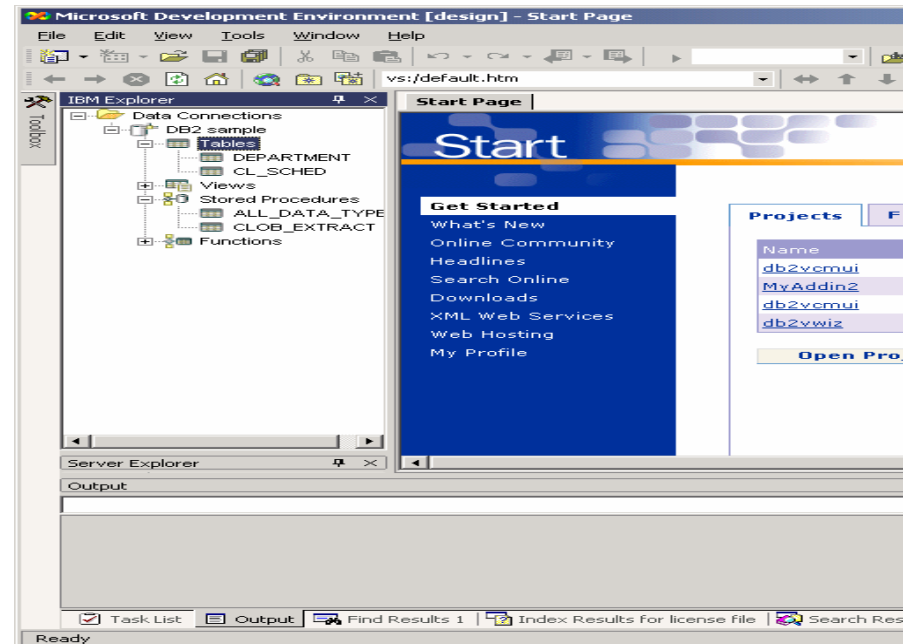


Microsoft Development Enhancements



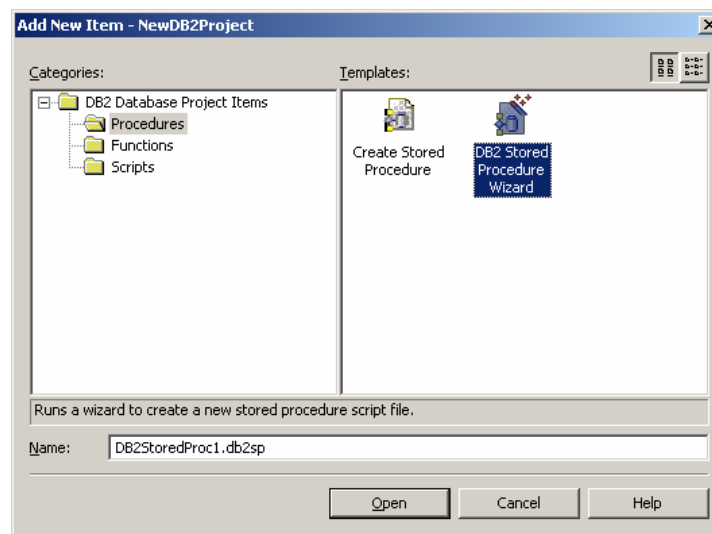
DB2 Development Add-In for VS.NET 2003

- Extended DB2 family support
 - ▶ DB2 UDB LUW V8
 - ▶ DB2 for OS/390 and z/OS V6, V7 and V8
 - ▶ DB2 for iSeries V5.2 and V5.3
- Enhanced data connections support
- Enhanced SQL Textbox control
- Simple SQL parser
- Improved managed provider tooling support
- CLR stored procedures support




DB2 Administration Add-in for VS .Net 2003

- This feature extends the IBM Explorer in Visual Studio .NET 2003 to include:
 - ▶ Adding DB2 remote data connections that are not cataloged on the local client
 - ▶ Adding 'create new' feature for supported DB2 objects from the data connection folder directly through the use of existing designers or wizards as applicable.
 - ▶ Adding 'generate DDL' feature for supported DB2 objects from the data connection folder into a new script file or through drag and drop onto a script file opened by the DB2 SQL editor.



Microsoft .Net CLR Stored Procedures

- DB2 introduces support for .NET stored procedures and UDFs
- .NET is Microsoft's vision for interoperability of disparate codebases
 - ▶ allows C, Java, VB, to be seamlessly integrated into a single application
 - ▶ .NET makes a compiler into a syntax checking device, that generates bytecode (same idea as Java), rather than a syntax checking device that generates machine specific instructions
- DB2 supports CLR (Common Library Runtime) before SQLServer!



DB2 .Net Data Provider Additional Support

- The new .NET Data Provider has been enhanced to support:
 - ▶ .NET framework 1.1 (Everett)
 - ▶ Visual Studio 7.1 (Everett)
- Supported Servers
 - ▶ DB2 UDB V8, iSeries, zSeries and SQL/DS as supported at V8 GA (i.e.. all supported servers except DB2 UDB V7)
 - server platform also includes DB2 UDB for zSeries V8 QPP level
- OLE DB .Net / ODBC .Net Data Providers
 - ▶ Tested and certified to ensure that they work correctly with DB2
 - ▶ The recommended (and only supported) data provided of choice is the DB2 .NET Data Provider




VS.NET: DB2 WebServices Support

- This feature extends the DB2 Development Add-In for Visual Studio.NET 2003 to support web services:
 - ▶ Add wizard to deploy an embedded WAS WebService for DB2 SQL statements and Stored procedures.
 - ▶ Add wizard to generate C# web methods using ADO.NET for DB2 SQL statements and Stored procedures.




Business Intelligence



Materialized Query Table Advisor

- An advisor to recommend creation (and removal) of Materialized Query Tables (formerly known as AST's in DB2) based on submitted SQL workload
 - ▶ Very difficult to choose the correct MQTs to create without the help of this wizard
 - ▶ Much more complex than deciding which indexes to create
- Includes enhancements to the index wizard
 - ▶ allow indexes and MQTs to be recommended together (including recommendations for indexes on MQTs).
 - ▶ "db2advis" command with added flags to allow index and MQT recommendations
 - ▶ This feature procedure will also be accessible using a GUI interface.



Multi-Dimensional Clustering Advisor

- The MDC advisor feature of the DB2 Design Advisor recommends MDC clustering dimensions, including coarsifications on base columns in order to improve workload performance
 - ▶ This includes potentially recommending generated columns that define coarsification of dimensions.
- Cost benefit analysis includes impact of MDC on Insert/Update/Delete activity against dimensions (which may cause records to move cells).
- Cost benefit analysis also models the potential negative effect of MDC caused by table expansion.
- The MDC advisor has a goal to select MDC solutions that result in a moderate table expansion




Partitioning Advisor

- The partitioning advisor is an extension to db2advis utility and is intended to run on a DB2 EEE/DPF environment
- The advisor takes a workload (consisting of SQL statements) as input and outputs the best partition (minimizing the cost of the workload) for each table in the workload. The tool can be used for the following applications:
 - ▶ choose partitioning for MQTs recommended by the MQT Advisor
 - ▶ decide the initial database partitioning before loading the data into the database
 - ▶ migrating from a non-partitioned DB2 to partitioned DB2 ESE
 - ▶ migrating to partitioned DB2 ESE applications from Informix or competing database systems
 - ▶ find out the right database partitioning after the environment (workload, underlying data, number of partition groups) is changed



Combined Index/MQT/MDC/Partitioning Wizard

- New graphical user interface panels in the design advisor provides more choices and controls, including:
 - ▶ the ability to display, edit or save the new recommendation set
- Implementation of the recommendations by creating:
 - ▶ indexes on the base tables
 - ▶ Materialized Query Tables
 - ▶ Indexes on the Materialized Query Tables
 - ▶ converting non Multi-Dimensional Clustering tables to Multi-Dimensional Clustering tables
 - ▶ repartitioning existing tables with a new set of partitioning key and/or in a new tablespace.



RUNSTATS Enhancements

- Sampling Support

- ▶ Boosts the performance of RUNSTATS by exploiting page level and row level sampling
- ▶ Page level sampling can improve performance because it reduces the number of table pages that need to be accessed to collect the statistics

- New RUNSTATS options flag will be added to indicate the type of sampling

```
RUNSTATS ON TABLE <table> WITH DISTRIBUTION TABLESAMPLE BERNOULLI(30)
```

```
RUNSTATS ON TABLE <table> WITH DISTRIBUTION TABLESAMPLE BERNOULLI(30)  
REPEATABLE
```

```
RUNSTATS ON TABLE <table> WITH DISTRIBUTION TABLESAMPLE SYSTEM(1.5)
```

- Support for utilities using a statistics profile

- ▶ There are some utilities such as LOAD and REORGCHK which can also specify that statistics are to be gathered.




Spatial Geodetic feature

- The DB2 Geodetic Supplement for Spatial Extender for DB2 UDB Version 8.2 will provide DB2 Customers with the ability to treat the earth as one seamless globe
- The DB2 Geodetic Supplement enables the existing Spatial Types to work with the understanding that the Earth is a closed sphere.
 - ▶ This enables the creation of trans-polar and trans-dateline seamless spatial values and enables trans-polar and trans-dateline seamless queries
- The DB2 Geodetic Supplement only adds new functionality to complement the current functionality in the Spatial Extender



Autonomic Computing Enhancements




Automated Backup w/Policy

- This feature provides a simple, automated backup that works by default "out of the box". This provides SMB users with an automated backup solution to help ensure their database is being backed up properly and regularly, without having to worry about when to back up or having any knowledge of the backup command.
- The new functions encapsulated in this feature includes:
 - ▶ Backup automation generated through a backup policy.
 - ▶ A new algorithm to identify the need for a backup on a particular database.
 - ▶ A new state-based Backup Health Indicator to surface the 'backup required' state, when the backup cannot be accomplished through the automated policy.
- The addition of a new Database Configuration Parameter that will be a switch to turn the backup automation on/off (AUTOMATIC_DB_BACKUP).




Self-tuning Backup and Restore

- Backup and Restore will automatically chose an optimal value for, # buffers, Parallelism, Buffersize whenever these values are not explicitly specified
- Any one or combination of these parameters can be explicitly specified
- For those customers using CLP, if a value for BUFFERS, BUFFER and PARALLELISM is not explicitly specified then a value of 0 will be assumed and the utility will chose optimal values.
- The BACKBUFSZ and RESTBUFSZ in the DBM CFG will be ignored. If a customer wished to use the values they must explicitly specify the value on the utility invocation




Inclusion of log files in online backup images

- Current architecture of the backup and restore utilities requires that the backup image and the required log files be shipped as separate objects
 - ▶ Possibility of losing the log files required for recovery
- Online backup images of databases and tablespaces will include the set of log files necessary to restore and recover the backup image to a consistent point in time.
 - ▶ Require to have a single object to ship to disaster recovery sites that would allow the customer to bring the database back to a consistent point in time
- New INCLUDE LOGS/EXCLUDE LOGS on ONLINE backup




Integrated and automated log file management

- DB2 introduces autonomic log management issues
- Allows DB2 to manage the use, archive, retrieval, and eventual deletion of log files with minimal user interaction
 - ▶ Many of the requirements came from ISVs
- New log manager introduced in DB2 UDB Version 8.2
- New History file entries
 - ▶ indicates when log file are archived, and where they are archived
 - ▶ Can determine which database backup, and exactly which log files are required to restored to any given point in time




Simplified Memory Configuration

- Database configuration parameters relating to heap sizes represent the maximum size for that particular heap
 - ▶ Determining the appropriate maximum size for any heap requires extensive knowledge about how that heap is used by DB2
 - ▶ No means of providing any guarantees that a particular heap will be able to obtain any memory at all
- New feature overcomes these limitations
 - ▶ simplify memory configuration by allowing the user to provide a single value for how much memory they want DB2 to use for each active database
 - ▶ The configuration size for a heap will now represent the guaranteed minimum, or reserved, size for that heap
 - ▶ Introduction of an overflow buffer



Automatic Runstats

- Table maintenance involves executing RUNSTATS and REORG utilities
 - ▶ More and more systems are required to be available 7/24
 - ▶ Crucial to be able to maintain most of the tables in an "online" mode
- Autonomic Runstats Feature
 - ▶ makes process of statistic collection/maintenance completely transparent to users
 - ▶ Maintains the table statistics up-to-date, acting both reactively and preemptively, so that the optimizer can always choose the best access plan
 - ▶ a state-based collection health indicator in the health monitor to surface any runstats actions that cannot be handled automatically.
 - ▶ notification of the identification of tables requiring manual runstats (if automation is disabled) or if automation fails



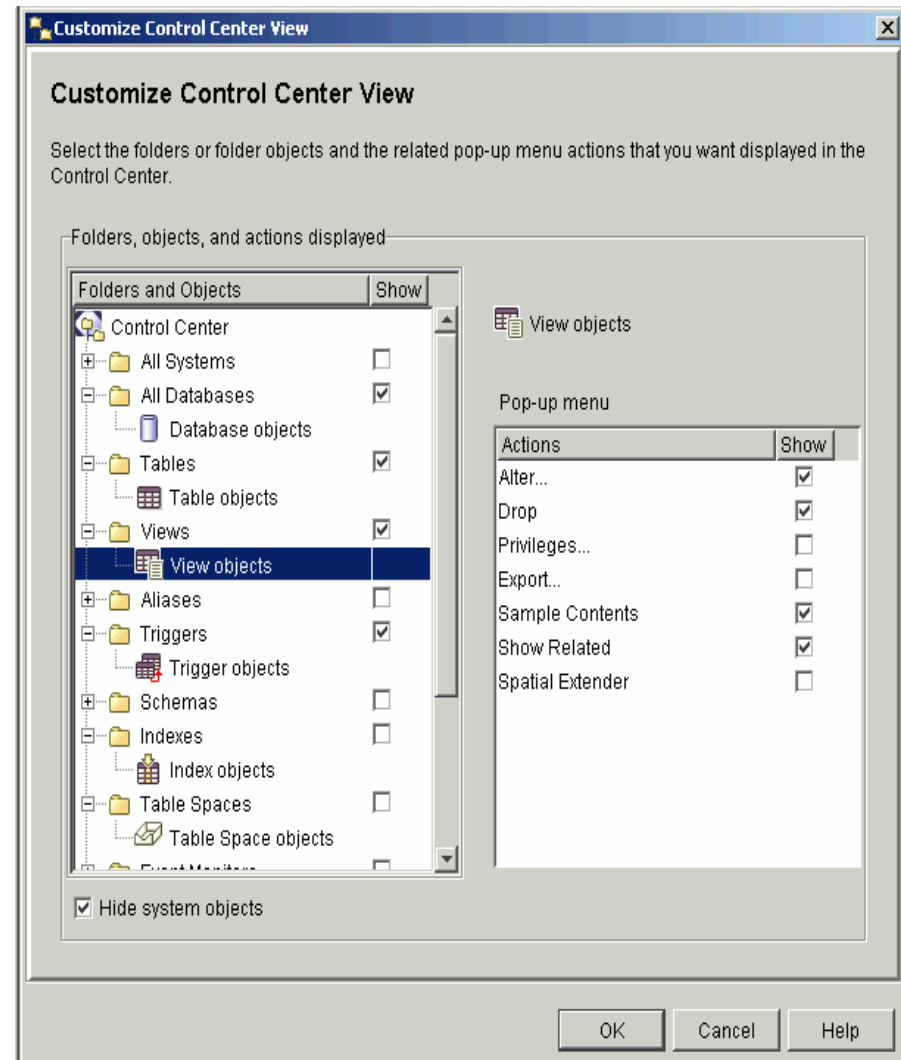
Automated Table Reorg w/ Policy

- This feature is used to manage table/index reorg without requiring manual intervention
- Accomplished through a health indicator to surface the need for table or index reorganization and policy and maintenance automation database configuration parameters to manage table reorganization automation
- The new function includes:
 - ▶ a state-based collection health indicator in the health monitor to surface any reorganization actions that cannot be handled by the policy
 - ▶ notification of the identification of tables requiring manual reorganization
 - ▶ function to control database maintenance automation
 - ▶ knowledge in DB2 to automatically identify the need for reorganization on a table
 - ▶ knowledge in DB2 to determine the type of reorganization required and when to run it
 - ▶ scheduling of reorganization actions based on automation specifications



Novice CC: Navigation

- This feature is designed to benefit developers from the aspect of providing easier and more efficient Control Center navigation
- Features Overview
 - ▶ Different Control Center views (Basic, Advanced and Custom) give users choices to pick the one that best suit their needs.
 - ▶ Simplification of the Control Center navigation
 - ▶ Basic Control Center meets majority of developer's needs



Schema Manipulation - Alter Table

- The new UI will support what is currently supported by the Alter Table statement, and go beyond the functionality of ALTER TABLE statement by offering the following new features:
 - ▶ Renaming columns
 - ▶ Removing columns
 - ▶ Altering column type and transform existing data using SQL scalar functions
 - ▶ Increasing or decreasing column size
 - ▶ Changing column default value
 - ▶ Changing column from NOT NULL to NULLABLE
 - ▶ Changing precision and scale for decimal



Control Center Support Throttle Utilities

- Allows DBA to reduce impact of running resource-intensive utilities on operational workload
 - ▶ Backup, Rebalance, Runstats
- Support setting and display for utility execution priority for the supported utilities:
 - ▶ backup, rebalance, runstats
 - ▶ support new option to set priority in utility dialogs
 - ▶ support new option to set priority in SHOW COMMAND
- Display priority of executing utilities in Task Center, and the current priority set for the session




Health Center Recommendation Advisor

- The Recommendation Advisor guides the user to take action to resolve the alert condition
- Recommendation Advisor will involve the following steps:
 - ▶ What is this Health Indicator?
 - what is this aspect of health for which the alert was raised?
 - what is the alert condition/severity
 - ▶ Why is this health condition important to me?
 - ▶ What are the possible courses of actions (recommendations) to resolve this problem.
 - ▶ What is actually going to be done based on the recommendation I selected? what, if any additional steps must I take?



Linux Enhancements



DIRECT I/O for Linux

- Currently, the file system caching policy of all I/O operations in DB2 are performed in buffered mode by default
- While this caching policy is extremely effective when the cache hit ratio is high, it has an overhead of making an extra copy of the buffer from the disk to file cache (in case of read) or from file cache to disk (in case of write).
 - ▶ Since the buffer is already cached in DB2's bufferpool layer, this dual level of caching proves to be unnecessary in situations where the cache hit ratio is low and a lot of I/O is performed.
- Direct I/O (DIO) is an alternate caching policy that reduces CPU utilization for reads and writes by eliminating the copy from file cache to user buffer.
- This feature extends DIO support on Linux distribution which are based on a 2.6 kernel (this was not ported back to previous kernel levels)
- Lab tests indicate a 12% performance improvement (OLTP workload)
- Specified by using the NO FILE SYSTEM CACHING parameter on the CREATE/ALTER TABLESPACE or CREATE DATABASE statement



CPU pinning on Linux

- NUMA technology allows for multiple cpus to run workloads in Linux
 - ▶ Processes can move from cpu to cpu
 - ▶ Penalty for switching CPUs in the middle of a task can be quite high (L1, L2 cached data)
- This feature allows db2 processes to be pinned to a CPU once assigned by the OS to a specific CPU.
 - ▶ This is not the same as having to decide which CPU to pin a process to, but rather whether or not to pin once a process is on a certain CPU.
 - ▶ In this case, the OS still decides which CPU to run the task on, but DB2 decides if the task should stay on that CPU until the task is finished.
- Enabled by specifying an XML format file (called the resource affinity configuration file)
- Can cause performance degradation
- Properly tuned can result in a 1%-2% improvement



Asynchronous I/O on Linux

- Allows DB2 to take advantage of the kernel-based asynchronous I/O capabilities of the 2.6 Linux kernel
- AIO is not enabled by default in DB2 UDB Version 8.2
- use `db2set` and the `DB2NOLIOAIO` variable to enable AIO
 - ▶ If false, AIO will be enabled in DB2
 - ▶ If true, AIO will be disabled.
- DB2 uses the Linux AIO interface for its page cleaners
- AIO can noticeably improve performance on Red Hat AS 2.1. SuSE SLES 8 (2.4 kernel) also supports it, so customers can manually enable and run AIO with DB2, even with existing code.
- Lab tests have shown a 14% performance improvement when running OLTP workload on a file system. 5% on Raw Devices




Large Page Support on Linux

- The Linux 2.6 kernel allows for large page sizes. Reduces overhead of mapping large memory in the system..
 - ▶ This support has also been back ported to the 2.4 kernel in certain distributions (Red Hat AS 2.1 and SuSE SLES 8)
 - ▶ This feature is necessary for improving performance on fast CPUs today that require lots of memory to be fully utilized.
- Must enable both Linux and DB2!!
- Internal lab tests indicate potential 9% improvement when DB requires large memory




High Availability Enhancements




High Availability Disaster Recovery (HADR)

- Target Market
 - ▶ Online commercial applications
- Challenge
 - ▶ 24 x 7 Availability
- Solution
 - ▶ Offsite Disaster Recovery:
 - Failover to disaster site without losing work
 - ▶ Onsite Standby:
 - Apply a security patch without taking down the database
- Value
 - ▶ Business continuation
 - ▶ Ease-of-use
 - ▶ Automatic Client Reroute



Offsite Disaster Recovery



Onsite Standby


High Availability Disaster Recovery

- DB2 High Availability Disaster Recovery (HADR) is an easy to use data replication feature that provides a high availability (HA) solution for both partial and complete site failures
 - ▶ HADR is intended to be the primary HA solution for complete site failures as well as the choice for applications demanding ultra-fast failover for partial site failures.
- HADR replicates data changes from a source database (called the primary) to a target database (called the standby)
 - ▶ Using synchronous mode, HADR can guarantee that any transaction committed on the primary is also committed on the standby
 - ▶ HADR allows failover and failback between the two systems
 - ▶ HADR requires the same hardware, OS and DB2 software on the two systems (except some minor differences during rolling upgrade).
- HADR propagates data changes by shipping database log records from the primary to the standby
 - ▶ It is tightly coupled with DB2 logging and recovery.



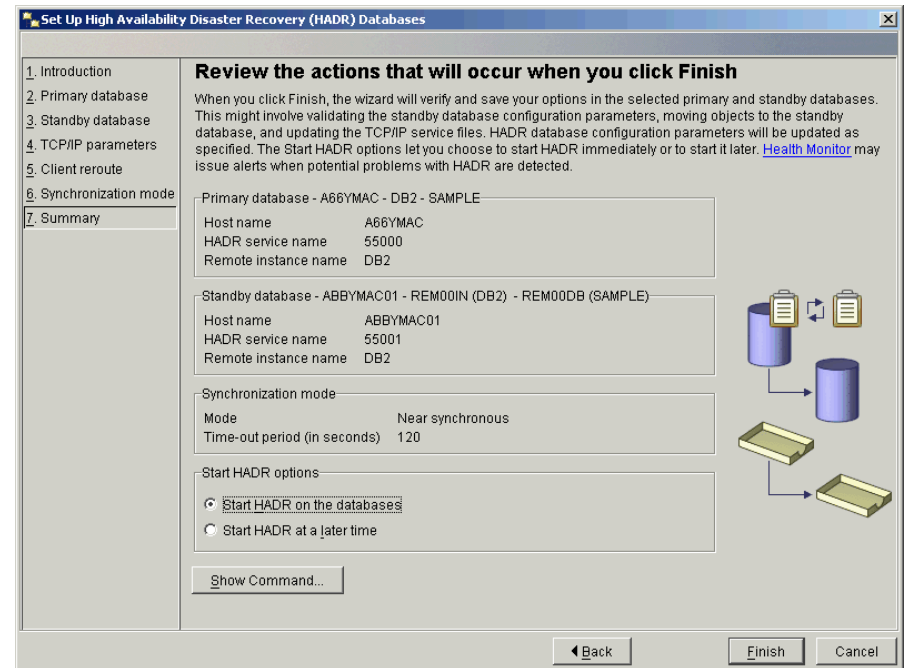
Client Reroute

- Whenever a server crashes, each client that is connected to that server gets a communication error which terminates the connection leading to an application error.
 - ▶ The DB2 UDB client code attempts to re-establish the connection to either the original server or to a new server.
 - ▶ When the connection is re-established, the application will receive an error that informs it of the transaction failure, but the application can continue with the next transaction.
- The Automatic Client Reroute feature could be used in following configuration environments:
 - ▶ Enterprise Server Edition (ESE) / Data Partitioning Feature (DPF)
 - ▶ Dpropr-style Replication
 - ▶ High Availability Cluster Multiprocessor (HACMP)
 - ▶ High Availability Disaster Recovery (HADR) environment
- Also available with LDAP Clients



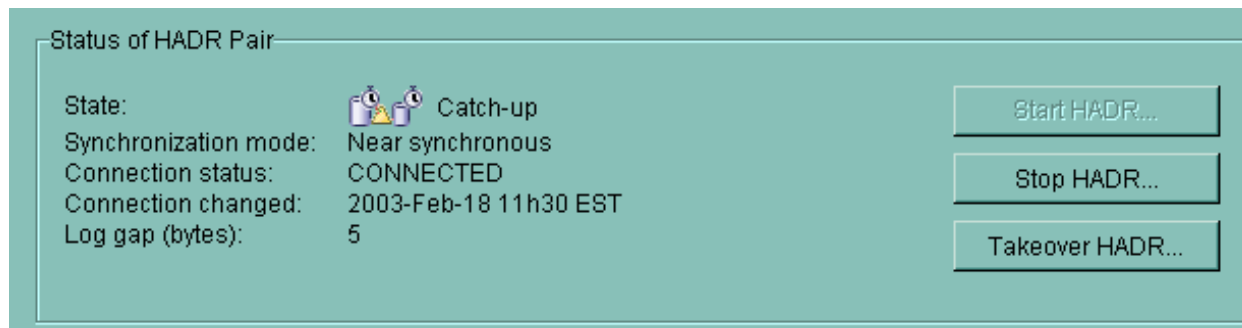
HADR: Set-up Standby Database Wizard

- The Configure HDR Databases Wizard facilitates the process to set up the primary and standby databases for HADR
- The wizard will guide the user to perform the following tasks:
 - ▶ Identify the HDR pair
 - ▶ Prepare the primary database for log shipping
 - ▶ Perform database backup
 - ▶ Copy backup image to secondary server
 - ▶ Perform restore on the selected standby database
 - ▶ Move any database objects not included in the backup image
 - ▶ Update service files
 - ▶ Update HDR related configuration parameters on both databases
 - ▶ Provide an option to Start HADR



HADR: Action windows

- A new set of tools are available to view the status of an HADR pair. From this view the following actions can be performed:
 - ▶ Refresh the status information
 - ▶ Stop and Start HADR
 - ▶ Takeover HADR
- The GUI will enhance the CLP by allowing the current role and state of a database to be used to determine which commands and options are valid before submitting the command to the engine





IBM Software Group

DB2® Universal Database Version 8.2 Technical Overview

DB2 Information Management Software

DB2 Worldwide Pre-sales Support
IBM Data Management Solutions

@business on demand software