



IBM Software Group

## IBM Informix 产品发展及IDS 10.0发布



 e-business software

© 2005 IBM Corporation

IBM Software Group



## 议题

- IBM Informix 数据库系统未来的发展
- IBM Informix 数据库产品的技术特性
- IBM Informix 数据库系统10.0的新功能

IBM DB2 Information Integrator Software | February 2003

© 2003 IBM Corporation

# IBM Informix IDS的发展

**10** 专注于关键业务联机事务处理的企业级数据库系统

**9.40** 性能更上一层楼  
高达128PB的扩展能力  
更高的数据安全性  
灵活的管理能力

**9.30** 共享SQL语句缓存  
模糊检查点  
动态日志技术  
动态锁管理  
数据刀片技术

**7x** 高端数据库系统...  
- 数据分片, 并行数据库查询, 高可用数据复制和企业级数据复制, 支持大量并发用户等.

**6x** 增强的功能和性能...  
- 多线程动态可伸缩体系结构(DSA), 异步 I/O, 完善的备份管理等.

**5x** 低端数据库系统, 具有...  
- 基于成本的优化器, 两阶段提交, 参照完整性等.

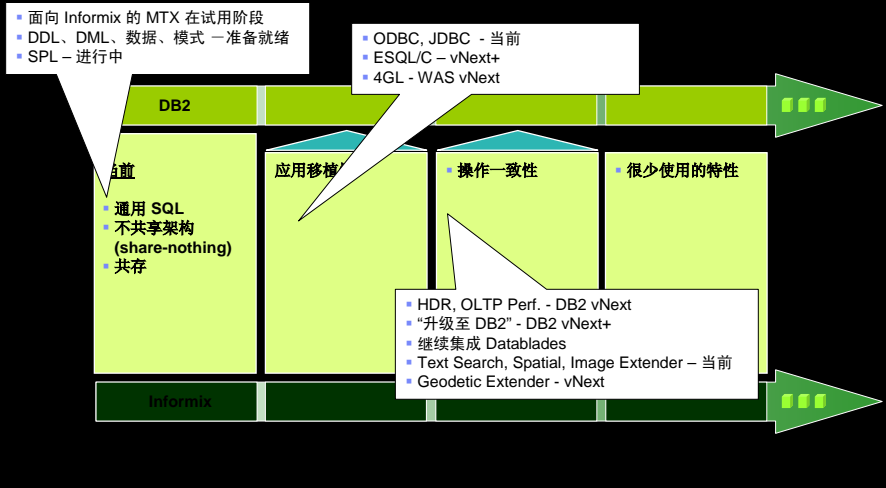
# IBM Informix数据库未来的发展

IBM公司将定期发布新的IDS版本

Year	Version	Status
2003	9.4	✓ Delivered 3Q 2003
2005	10.0	✓ Delivered 2Q 2005
2006	10.x	✓ In Concept

与DB2相互补充、融合

## IBM Informix数据库未来的发展



## IDS / DB2 产品互操作性

Informix 产品	DB2 产品	认证?
IDS 7.31/9.3/9.4/10.0	DB2 Table Editor	是
IDS 7.31/9.3/9.4/10.0	DB2 Web Query Tool	是
IDS 9.3/9.4/10.0	DB2 Content Manager	是
IDS 9.3/9.4/10.0	DB2 Information Integrator	是
IDS 9.3/9.4/10.0	DB2 Office Connect	是



## XPS 和 Red Brick / DB2 产品互操作性

Informix 产品	DB2 产品	认证?
XPS 8.4 / Red Brick 6.2	DB2 Information Integrator	是
XPS 8.4 / Red Brick 6.2	DB2 Warehouse Manager	是
XPS 8.4 / Red Brick 6.2	DB2 Office Connect	是
Red Brick 6.2	DB2 Intelligent Miner	是



## IDS 与 IBM 软件产品

### WebSphere software

IDS 9.3, 9.4, 10.0	WebSphere Application Server, v4.0, V5.0
IDS 9.3, 9.4, 10.0	WebSphere Portal Server, V5.0
IDS 9.3, 9.4, 10.0	WebSphere Studio Application Developer, V4.0 Web Services Object Runtime Framework (WORF) V3.5
IDS 9.3, 9.4, 10.0	WebSphere Studio Enterprise Developer, v.5.0

### Tivoli software

IDS 9.3, 9.4, 10.0	Tivoli Monitoring for Databases, V5.1.0
IDS 9.3, 9.4, 10.0	Tivoli Data Protection for Informix

### Lotus software

IDS	Lotus Enterprise Integration (LEI) 6 / Domino 6
-----	---

## Informix – 先进技术融入IBM

- Cloudscape : 零管理, 占用资源少, J2EE/JDBC 兼容可嵌入的 Java 数据库, 带完全 SQL 支持
  - 集成入**50多个** IBM 软件产品 – 包括 WebSphere 和 DB2
- Informix HDR
  - 内置于将予已发布的 DB2 UDB 8.2 版
- Informix Dynamic Server
  - 内存管理方面的改进来自于 DB2 UDB
- IDS、XPS 和 Red Brick 特性将融入 DB2 UDB 的每一个未来版本

Rational software DB2 Information Management Software

Tivoli software WebSphere software Lotus software

## IBM开发者园地– Informix开发者社区

- 为 IDS 开发应用你需要了解的一切
  - IDS
  - DataBlade 模块
  - 4GL
  - Java
  - 系统管理
  - 工具
  - 连接性
  - ...更多!

[www7b.boulder.ibm.com/dmdd/zones/informix](http://www7b.boulder.ibm.com/dmdd/zones/informix)

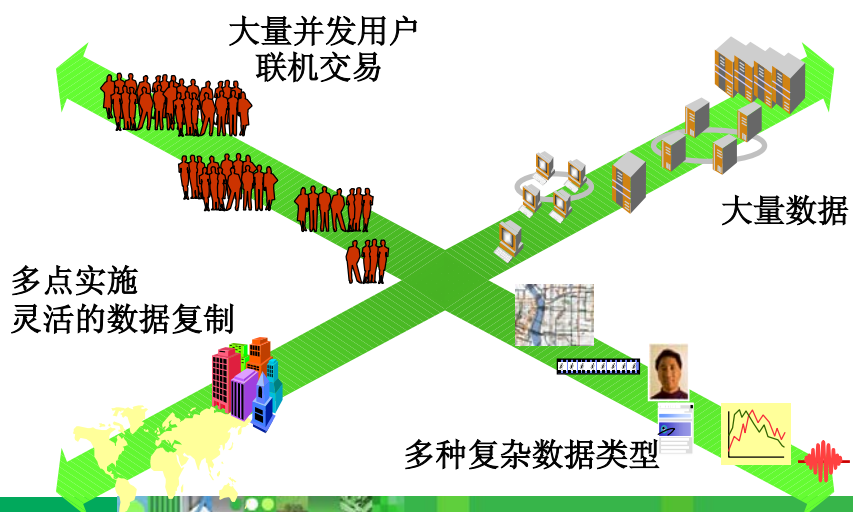
**Informix Developer Zone**  
your growing on-line community

## 议题

- IBM Informix 数据库系统未来的发展
- IBM Informix 数据库产品的技术特性
- IBM Informix 数据库系统10.0的新功能

## IBM Informix IDS产品技术特性

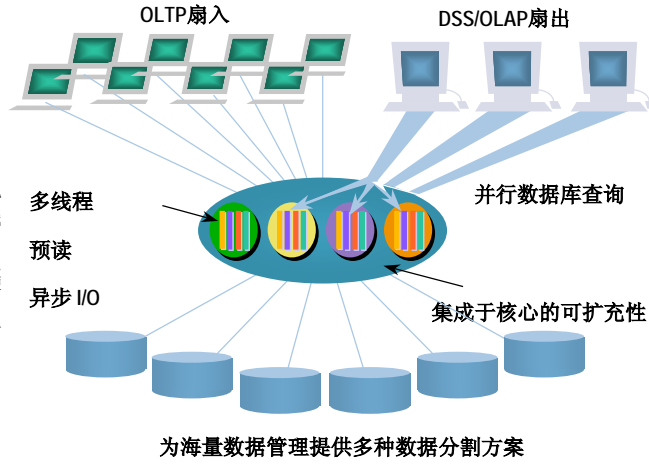
--高性能，可伸缩性，可管理性和可扩充性



## IBM Informix IDS动态可伸缩的体系结构

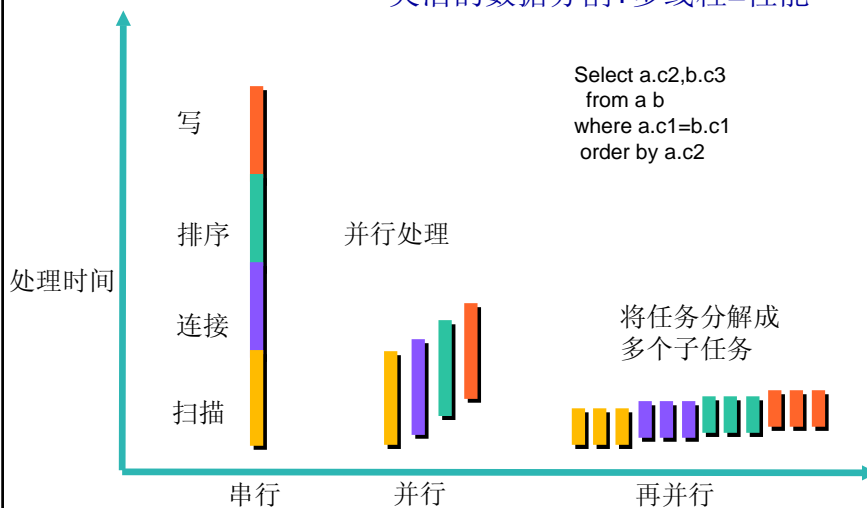
Modern parallel-everywhere base RDBMS

- 核心并行能力
- 动态调谐和管理
- 不断进行性能提升
- 不断进行功能扩展
- 多线程的进程构成虚处理机
- 分类的虚处理机组成流水线的工作方式
- 虚处理机数量可动态调整以适应系统负载及充分利用硬件资源
- 在多处理机性能扩展和并行处理方面有突出表现



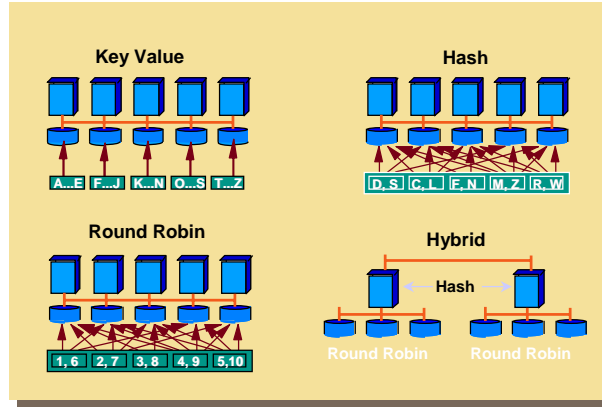
## IBM Informix IDS的技术特点

--灵活的数据分割+多线程=性能



## IBM Informix IDS 技术特点

### --灵活的数据分割技术

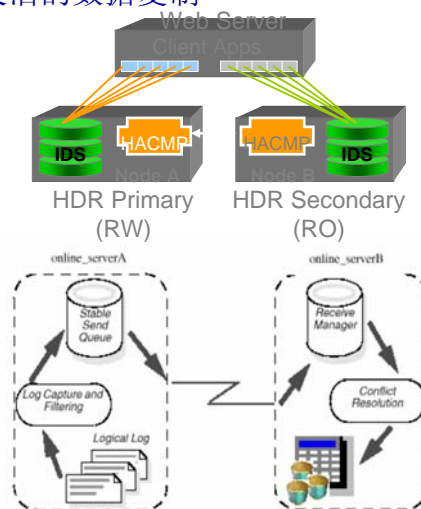


提供多种分割方案，最大限度地防止I/O倾斜

## IBM Informix IDS 技术特点

### --灵活的数据复制

- 例程级的复制
- 不同主机系统间进行复制
- 针对数据库表级的数据复制，对DML进行复制
- 基于数据库逻辑日志
- 基于事务
- 配置灵活
  - 可指定一对一
  - 一对多的复制
  - 提供级连复制
  - 可选择单向或双向复制
  - 提供多种冲突解决办法







IBM Software Group

## IBM Informix Dynamic Server 10.0的新特性



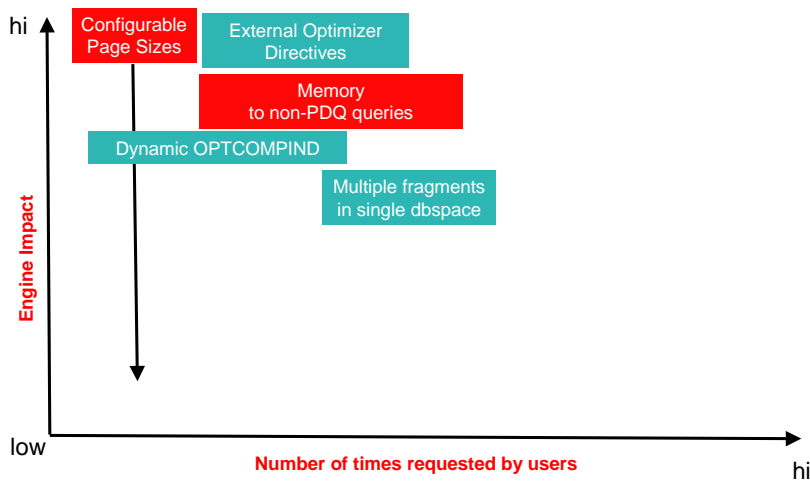
Your Name  
Here  
IBM Software  
Group  
Data

© 2005 IBM Corporation

IBM Software Group



## 高性能方面的增强



IBM DB2 Information Integrator Software | February 2003

© 2003 IBM Corporation

## 高性能: Configurable Page Size

- Space efficiency - larger pages up to **32K** bytes contiguous space
- Increased maximum key size - longer keys up to **3K\*\*** bytes
- Access efficiency- less I/O operations for data and indices

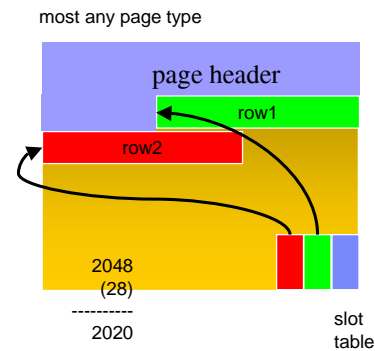
- Space efficiency - larger pages up to 32K bytes contiguous space

A row size of 1200 bytes:

- 1 row fits on a 2k page (6k every 3 rows).
- 3 rows fit on a 4k page, a savings of %33

For thirty 1200-byte rows:

- A 2k page size requires 60k.
- A 4k page size would require only 40k.
- A 6k page size requires just 36k, a 40% savings

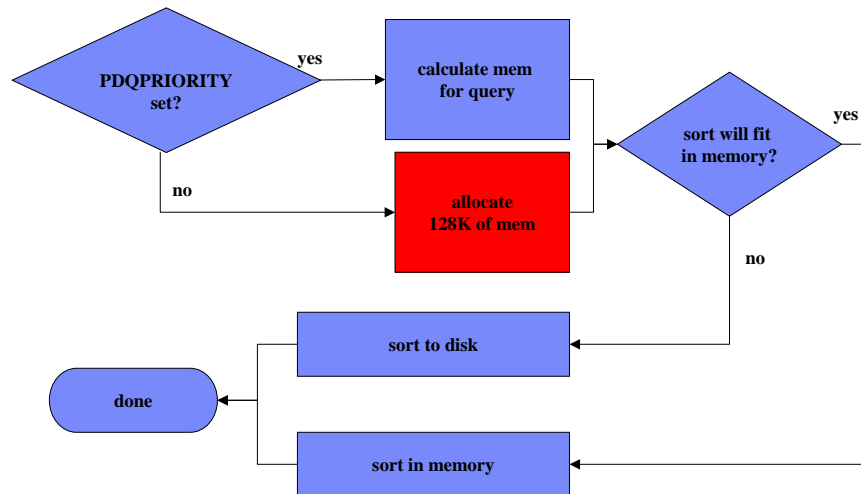


## 高性能: Configurable Page Size

### Access efficiency - less I/O operations for data and indices

- Long rows put into a single page benefits data access times by decreasing the number of pages read per row.
- Pages large enough to fit "oversize" rows eliminate the overhead of access time for remainders pages.
- More items put on a larger index page and a decreased number of levels in a btree index reduce index traversal costs.
- For DSS environments, using larger data pages that maximize disk scan rates may improve table scan performance.

## 高性能: n-PDQ queries - a problem....



## 高性能: Memory Allocation for non-PDQ Queries

- You can specify how much memory is allocated to non-PDQ queries.
  - The default of 128K can be insufficient for queries that specify ORDER BY, GROUP BY, hash joins, or other memory-intensive options.
- Use the new configuration parameter, `DS_QUERY_MEM`, to specify more memory than the 128K that is allocated to non-PDQ queries by default.
- The `onstat`, `onmode`, and `onmonitor` utilities support this configuration parameter.



## 高性能: Storing/Applying External Directives

- You can create, save, and reuse external optimizer directives.
- External optimizer directives are useful when it is not feasible to rewrite a query for a short-term solution to a problem
  - for example, when a query starts to perform poorly.
- This feature is implemented as a new SQL statement, **SAVE EXTERNAL DIRECTIVES**
  - creates and registers external optimizer directives in a new **sysdirectives** table of the system catalog.
  - Use the new **IFX\_EXTDIRECTIVES** environment variable or the **IFX\_EXTDIRECTIVES** configuration parameter to enable this feature.

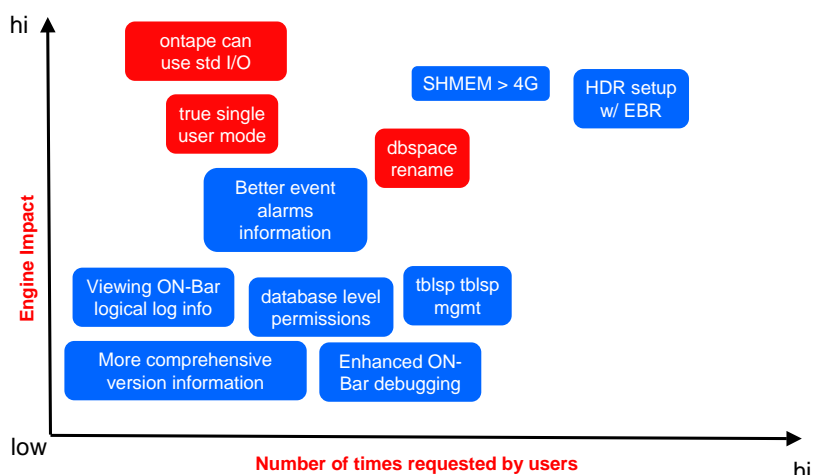
## 高性能 : Multiple Table Fragments in Single Dbspace

- You can create partitions within a dbspace that can each support a table fragment.
  - reduces the total number of dbspaces needed for a fragmented table.
- Storing multiple table fragments in a single dbspace improves query performance over storing each fragmented expression in a different dbspace.
- This feature improves performance and simplifies management of dbspaces.

## 高性能 : Dynamic OPTCOMPIND

- You can use SET ENVIRONMENT OPTCOMPIND to set **OPTCOMPIND** environment variable dynamically for the current session.
- The value that you enter using this statement takes precedence over the current setting specified in the ONCONFIG file.
- The default setting of the **OPTCOMPIND** environment variable is restored when your current session terminates.
- No other user sessions are affected by SET ENVIRONMENT OPTCOMPIND statements that you execute.

## 实用性方面的增强





## 可管理: ontape use of standard I/O

- You can now specify that **ontape** uses standard I/O instead of a tape device or disk file.
- Specifying stdout or stdin allows **ontape** to use pipes (an OS provided buffer mechanism to connect separate programs to a data stream) for archives and restores.
- Using pipes, the data can be processed by other programs without requiring that the data be saved in files or tape devices.
  - For example, you can use compression to save media space, use cloning to duplicate the archive for safety reasons, or restore the data onto another server instance.
  - This feature is **especially efficient for setting up High-Data Availability** Replication by restoring the data to the secondary server while skipping the intermediary step of saving the data to a file or disk.



## 可管理: Renaming Dbspaces

### Description

The ability to change the name of a previously defined dbspace.

### The Problem

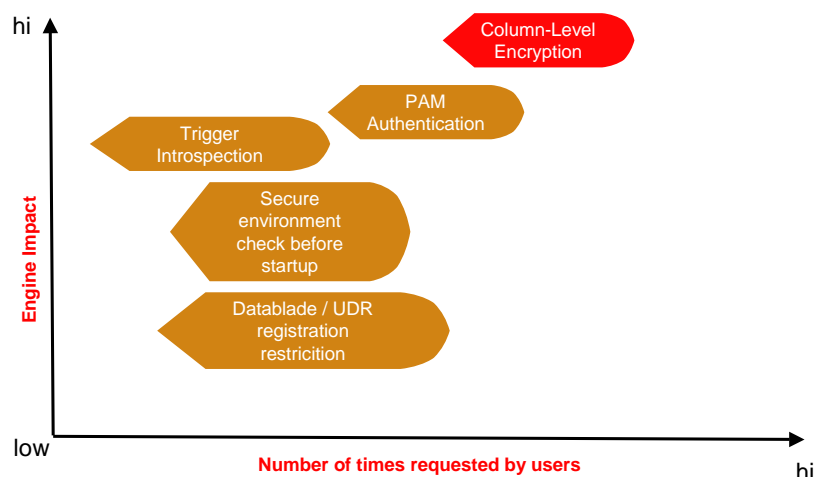
- Existing customers who wish to reorganize their data are moving their data to a new dbspace, then reloading the data back into the original dbspace to regain the original dbspace name. The last reload step could be avoided by using a rename dbspace option. Helps in Recycling of Dbspaces.
  - Time consuming operations as reorganizing the data in an existing dbspace can benefit from this feature.
- **The rename dbspace operation only changes the dbspace name; it does not reorganize data**
  - **A Level 0 archive of the renamed space and root dbspace need to be taken after renaming.**
  - **Rename done on HDR primary will propagate to the secondary.**

## 可管理 : Single User Mode

### Overview

- "Single-User Mode" feature allows the IBM Informix Dynamic Server to be put into a maintenance mode, allowing only the user 'informix' to connect the server.
- It is intended to be viewed as a mode intermediate between Quiescent mode and Online mode.
- It allows a DBA to have the server in a fully functional mode where any required maintenance may be performed.
- It is intended to be used by DBA's to perform any SQL/DDDL maintenance while preventing normal users from connecting.
- ISA also supports the Single-User Mode functionality.
- Utility Enhancements :oninit : -j, onmode -j

## 安全性



## 安全性: Column Level Encryption

### ▪ Highlights

- Built-in SQL encryption functions to support data encryption
- Latest cryptographic standards
- 128 bit AES and Triple-DES
- Passwords up to 128 characters
- Session wide password management for easy programming
- View/trigger/SPL support
- **Consistent with DB2 syntax**

## 安全性: Column Level Encryption

### - Usage examples

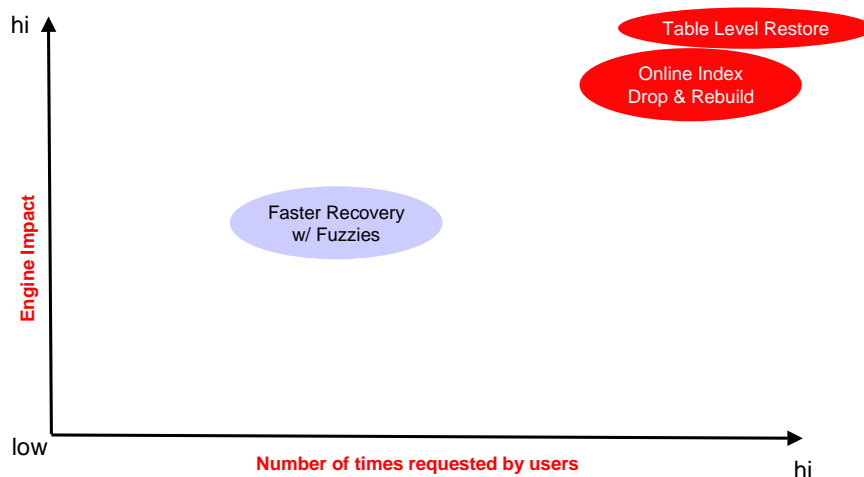
```
INSERT INTO table VALUES (1, ENCRYPT_AES(data, password));  
UPDATE table SET column = ENCRYPT_TDES(data, password)  
WHERE ...;  
SELECT DECRYPT_CHAR(column, password) FROM table;  
EXECUTE FUNCTION ENCRYPT_AES(data, password, hint);
```



## 安全性: Secure Environment Check

- Server utilities on UNIX now check if the environment is secure by testing for the following:
  - The permissions on **\$INFORMIXDIR** and some directories under it are correct. For each directory, check that the directory exists, is owned by user **informix** and the correct group, and that its permissions do not include write permissions for the group or other users.
  - The permissions on the ONCONFIG file are correct. The file must belong to the DBSA group. If the DBSA group is group **informix** (default), then the ONCONFIG file should be owned by user **informix** too; otherwise, the ownership is not constrained. The file must not have write permissions for others.
  - The permissions on the **sqlhosts** file are correct. Under a default configuration, the **sqlhosts** file is **\$INFORMIXDIR/etc/sqlhosts**; the owner should be user **informix**, the group should be either the **informix** group or the DBSA group, and there should be no public write permissions. If the file is specified by setting the **INFORMIXSQLHOSTS** environment variable, then the owner and group are not checked, but public write permissions are not permitted.
  - The length of both the filenames **\$INFORMIXDIR/etc/onconfig.std** and **\$INFORMIXDIR/etc/\$ONCONFIG** must be less than 256 characters.

## 实用性的增强



## 实用性: Index Changes

- CREATE INDEX and DROP INDEX now supports DDL operations that apply no exclusive lock to the table on which the specified index is defined.
  - If you use this syntax to create an index on a table that other users are accessing, the index is not available until no user is updating the table.
  - After you issue the new syntax to drop an index, no one can reference the index, but concurrent DML operations can use the index until they terminate.
  - Dropping the index is deferred until no user is using the index.
  - This feature maintains the availability of the table within a production environment after an existing index has ceased to be efficient.

## 实用性: Point-in-Time Table Restore

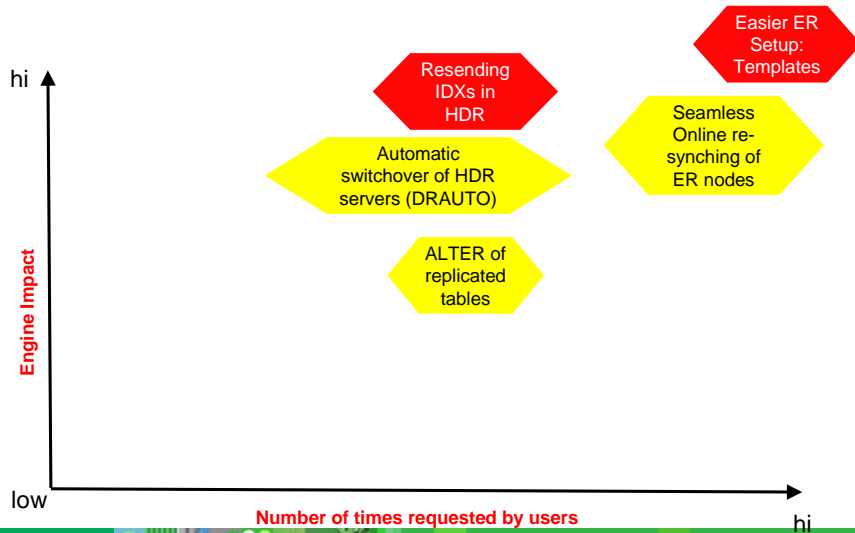
- Purpose: to provide the customer with the ability to easily extract a set of tables, a table or a portion of a table from a level 0 archive to a user specified point in time.
- The extracted data can be placed in an external table or on a table on the server of the user's choice regardless of server version or machine type as long as the database server is listed in the sqlhost file.

### Benefits to the DBA

- SQL Driven Distributed Restore (SDDR)
- Extract a table or set of tables
- Data may be placed in the same version database or on a different database version with a different machine architecture
- You may apply a filter to the retrieved data
- Retrieve just a subset of the columns
- Repartitioning of the data



## 高可用的增强



## 高可用: Resending Indexes in HDR

- You can resend an index that became corrupt on the secondary server in an HDR pair.
- Resending an index is quicker than dropping and then rebuilding the index on the primary server.
- This feature increases the availability of the HDR primary server.

## 高可用: DRAUTO

- You can automate switching servers for High-Availability Data Replication if the primary server fails by using the DRAUTO configuration parameter.
- If DRAUTO is set to either
  - RETAIN\_TYPE or REVERSE\_TYPE, the secondary database server switches to type standard automatically when an HDR failure is detected.  
RETAIN\_TYPE, the original secondary database server switches back to type secondary when the HDR connection is restored.  
REVERSE\_TYPE, the original secondary database server switches to type primary when the HDR connection is restored, and the original primary switches to type secondary.

## 高可用: Replicate Templates

### ▪ Overview

- Ease of ER administration and setup
  - The entire Enterprise Replication domain can be setup using simple commands with options such as defining replicates on all tables within database, specified on the command line or using a input file.
  - A template can perform an initial data synchronization on new servers being added to a template..
  - A template can optionally create tables during realization if they do not exist on target servers during template realization.
- Eliminates most of the table Schema related errors
  - Templates use the master dictionary from the Master node to create these tables to ensure consistent schemas between the nodes.



## 高可用: Replicate Resynch

### Overview

- Meant to be used:
  - Bring a newly participating table up-to-date with the ongoing replication
  - Repair a replicated table if replication was stopped or failed for some reason.
- Two ways to repair a table:
  - Quick way: By processing the ats/ris files.
  - Other way: By defining and running a 'resynch job'.



## 总结

### IBM Informix数据库系统将持续稳定的发展

- 领先的数据库体系结构--动态可伸缩的体系结构
- 数据库系统资源动态可调整--支持7x24关键业务系统
- 灵活的多库多表结构--易于业务的扩展
- 基于核心的并行处理--parallel every thing!
- 支持超大数据库容量- 128 PB
- 断点续恢复能力
- 内置的灵活复制技术

.....



IBM Software Group

谢谢!



© 2005 IBM Corporation