

Ingres[®] 2006

Release Summary

INGRES[®]

This documentation and related computer software program (hereinafter referred to as the "Documentation") is for the end user's informational purposes only and is subject to change or withdrawal by Ingres Corporation ("Ingres") at any time.

This Documentation may not be copied, transferred, reproduced, disclosed or duplicated, in whole or in part, without the prior written consent of Ingres. This Documentation is proprietary information of Ingres and protected by the copyright laws of the United States and international treaties.

Notwithstanding the foregoing, licensed users may print a reasonable number of copies of this Documentation for their own internal use, provided that all Ingres copyright notices and legends are affixed to each reproduced copy. Only authorized employees, consultants, or agents of the user who are bound by the confidentiality provisions of the license for the software are permitted to have access to such copies.

This right to print copies is limited to the period during which the license for the product remains in full force and effect. The user consents to Ingres obtaining injunctive relief precluding any unauthorized use of the Documentation. Should the license terminate for any reason, it shall be the user's responsibility to return to Ingres the reproduced copies or to certify to Ingres that same have been destroyed.

To the extent permitted by applicable law, INGRES PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IN NO EVENT WILL INGRES BE LIABLE TO THE END USER OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF INGRES IS EXPRESSLY ADVISED OF SUCH LOSS OR DAMAGE.

The use of any product referenced in this Documentation and this Documentation is governed by the end user's applicable license agreement.

The manufacturer of this Documentation is Ingres Corporation.

For government users, the Documentation is delivered with "Restricted Rights" as set forth in 48 C.F.R. Section 12.212, 48 C.F.R. Sections 52.227-19(c)(1) and (2) or DFARS Section 252.227-7013 or applicable successor provisions.

Copyright © 2005-2006 Ingres Corporation. All Rights Reserved.

Ingres, OpenROAD, and EDBC are registered trademarks of Ingres Corporation. All other trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

Contents

Chapter 1: New Features	5
What Is Ingres 2006?	5
Open Source	5
Features Included in Open Source	6
Features Not Included in Open Source	6
New Features for Database Administrators	7
Parallel Query	7
Key Range Table Partitioning	7
Online Modify	7
Ingres Cluster Solution for Linux	7
Ingres High Availability Option	8
Unextenddb Utility	8
Killing Queries	8
Numeric Overflow Support in Report-Writer	8
Collation Specification at the Column Level	8
System-wide Setting for Default Lock Level	9
New Features for Application Developers	9
Automatic Sequence Number Generation	10
No Wait for Lock Requests	10
Support for New Data Types	11
Additions to the Visual DBA Suite	11
Visual Database Objects Differences Analyzer	11
Visual Configuration Differences Analyzer	12
Export Assistant	13
Connectivity Enhancements	16
Support for JDBC 3.0 API	16
Data Access Server	16
JDBC Driver	17
Updateable Result Sets in JDBC	18
.NET Data Provider and Visual Studio .NET Integration	19
Ingres ODBC Administrator	20
WinSock 2.2 API TCP/IP Protocol Driver for Windows	21
ODBC Call-level Interface	21
New Features for Linux	22
RPM Packaging	22
KDE/GNOME Desktop Integration	22
Linux Cluster Support	22
Additions to User Documentation	22

Getting Started Guide for Linux	22
Connectivity Guide.....	23
Interactive Performance Monitor User Guide	24
Chapter 2: Changes to Existing Features	25
Enhanced Unicode Support	25
Unicode Coercion.....	25
Collation Sensitive Support for Wildcard Searching	25
Complex Query Optimization	25
Increased Range Table Limit.....	26
JDBC User ID Enhancements	26
Increased Column Limit	27
VDBA Enhancements	27
VDBA Architecture Split	27
Visual SQL.....	28
Visual Performance Monitor	29
Additional Properties	30
Notification of Metadata Changes	30
Ingres Visual Manager Enhancements	31
Direct Access to All Ingres Tools	31
Registering Alerts in the OS Event Log.....	32
Viewing Message Explanations.....	32
Configuration Rules System Enhancements.....	33
Shadow Copy of the Symbol Table.....	33
Additional Join Functionality.....	33
Improved Out-of-the-Box Configuration Defaults	33
Improved IMA Support.....	33
Createdb Enhancements	34
Alterdb Enhancements	34
Terminal Monitor Enhancements.....	34
Enhancements for Log Full.....	34
Extended B-tree Limits	34
Installer Enhancements.....	35
Installation as a User Other Than ingres	35
Supportability Enhancements.....	35
Help System Enhancements.....	36

Chapter 1: New Features

The *Release Summary* for Ingres® describes new features and changes to existing features for Ingres 2006.

What Is Ingres 2006?

Ingres 2006 is an enterprise-class open source database. It is a reliable, high-performance database solution that offers scalability, integration, and flexibility. Ingres scales from the laptop to the cluster in a single version and has the lowest total cost of ownership in the industry.

Its enterprise-class features are available under an Open Source license. Designed for “set-and-forget” operations on a compact footprint, Ingres is the ideal unattended database for both enterprise and embedded deployment.

Ingres 2006 was previously released by Computer Associates as Ingres r3.

Open Source

The open source model provides many business and technical advantages, such as:

- Ability to receive features more quickly
- Ability to extend Ingres to meet specific requirements
- Collaboration on projects
- Building relationships with leading open source application vendors
- More easily extending support to non-traditional or emerging platforms
- Enhancing security by providing code that is less vulnerable to attack

Features Included in Open Source

The following Ingres components are contributed to the open source community:

- Ingres DBMS and associated database administration tools
- Embedded SQL precompilers
- Character-based querying, reporting, and application development tools
- Connectivity components, including ODBC, JDBC, and the .Net Data Provider
- Ingres Distributed Option
- Ingres Replicator Option
- Ingres Web Deployment Option
- TP monitors, including CICS, Tuxedo, and Encina

Features Not Included in Open Source

Features not included in the open source edition are as follows:

- Support for spatial objects
- B1 security

The spatial object library is available for download from www.ingres.com if you have a valid technical support contract with Ingres Corporation.

While the source for the Visual DBA suite is not contributed to the open source community, the suite is included in the Ingres for Windows download.

The following members of the Ingres product family are not contributed to open source and continue to be available for purchase from Ingres Corporation:

- OpenROAD
- Enterprise Access
- EDBC products

New Features for Database Administrators

The new features for database administrators make it easier to administer an Ingres database and allow database administrators to deploy Ingres in a more scalable environment.

Parallel Query

As a multi-threaded server, Ingres has long supported symmetric multi-processing (SMP) systems by creating individual threads to handle user queries, and executing these queries in parallel across all available CPUs. Ingres 2006 introduces the ability to execute individual queries in parallel across all available CPUs in the system, which greatly improves performance. For more information, see the *SQL Reference Guide*.

Key Range Table Partitioning

With the functionality of key range table partitioning, the data in a database can be partitioned based upon the value of a given key, which significantly improves the performance of queries that require a full table scan. For more information, see the *SQL Reference Guide*.

Online Modify

The online modify functionality enables users to modify tables while working online. The DBMS performs the modify processing while allowing concurrent updates to the table. For more information, see the *SQL Reference Guide*.

Ingres Cluster Solution for Linux

Ingres Cluster Solution is an extension to the base Ingres product that allows it to operate on a cluster as a distributed application, providing transparent access to databases residing on shared storage devices. If there is an Ingres or hardware failure on one of the nodes, only those transactions that are in process on the failing node are aborted and ACID properties are maintained.

For more information, see the *Getting Started* guide for Linux.

Ingres High Availability Option

Ingres 2006 provides automatic failover support for Ingres clusters on Sun Solaris and Windows. For more information, see the *System Administrator Guide*.

Unextenddb Utility

Ingres 2006 introduces an unextenddb utility, which provides the ability to unextend a database location. For more information, see the *Database Administrator Guide* and *Command Reference Guide*.

Killing Queries

Ingres 2006 provides the ability to kill a query in another session while leaving the session in place. For more information, see the *Command Reference Guide*.

Numeric Overflow Support in Report-Writer

Ingres 2006 provides support for numeric overflow in Report-Writer. For more information, see the *Command Reference Guide*.

Collation Specification at the Column Level

This feature allows the specification of a collation sequence at the column level that differs from the database default collation sequence. A new optional "collate" clause is added to the column specification on the create table statement. The collate clause lets you specify a case-insensitive collation for columns that contain Unicode data.

In previous releases of Ingres, the DBA had the option when creating a database of defining a collation sequence to be used for non-Unicode text columns. The Unicode standard default collation sequence was provided for Unicode text columns. These collation sequences were in effect for all columns in the database and could not be changed without recreating the database.

Note: As part of this feature, the data descriptor used throughout the Ingres system changed. This data descriptor is also compiled into imaged ABF applications. After upgrading to Ingres 2006, all ABF applications should be re-imaged. Delete the contents of the ABF object directory, `$ING_ABFDIR/database-name/app-name`, and then re-image.

For details on this new feature, see the *SQL Reference Guide*.

System-wide Setting for Default Lock Level

This feature adds a system configuration parameter that allows the Ingres administrator to define the default lock level for the entire Ingres instance.

In previous releases of Ingres, to alter the default lock level the application programmer had to use the “set lockmode” statement in the application, and it affected only the current session.

The new parameter, `system_lock_level`, is available to the administrator through Configuration-By-Forms. Valid values are DEFAULT, ROW, PAGE, and TABLE. The DEFAULT value is the default and allows the system to decide the lock level. DEFAULT is the assumed value if the parameter is not present.

Note: Each of the default lock levels is subject to escalation, as in previous releases.

For more information on the default locking level, see the *Database Administrator Guide*.

New Features for Application Developers

The new features for application developers provide additional flexibility in application design and make it easier to migrate existing applications that run against non-Ingres database architectures.

Automatic Sequence Number Generation

Ingres has a sequence facility that provides the ability, through SQL, to create a column in a table that contains a sequentially incremented number for each row.

Users can define a named sequence generator by using a create sequence statement. The sequence generator can produce values in any context that requires a scalar value using the phrase next value for <sequence name> or <sequence name>.nextval. The sequence generator can be used in the values list of an insert statement, in the select list of a query, or anywhere that a scalar numeric value is required.

The sequence facility allows sequences to be defined as integer or decimal values. Ingres permits sequences to be decimal(31), which supports a range of +/- 10**32.

Use the alter sequence statement to change the parameter settings for a sequence generator and the drop sequence statement to delete a sequence generator.

For more information, see the *SQL Reference Guide*.

No Wait for Lock Requests

Ingres uses a value of "no wait" for the timeout parameter in the set lockmode statement to indicate that when a lock request that is made that cannot be granted without incurring a wait, control is immediately returned to the application that issued the request. "No wait" applies to any lock in a transaction or a lock on one or more specific tables. For more information, see the *SQL Reference Guide* and *Database Administrator Guide*.

Support for New Data Types

Ingres 2006 supports the bigint and tinyint data types.

bigint

The bigint numeric data type stores 64-bit integers. The bigint data type is an implementation of the ANSI standard bigint. Integer8, Int8, and i8 are synonyms for this data type.

The int8 function converts the specified expression, which can be a c, char, varchar, nchar, nvarchar, text, float, money, decimal, integer1, smallint, to a 64-bit integer. Decimal and floating-point values are truncated. Numeric overflow occurs if the integer portion of a floating-point or decimal value is too large to be returned in the requested format.

tinyint

The tinyint numeric data type is a synonym for i1. Supported values are from -128 to +127.

For more information on new data types, see the *SQL Reference Guide*.

Additions to the Visual DBA Suite

Three new Visual DBA (VDBA) Suite tools have been added.

Visual Database Objects Differences Analyzer

The new Visual Database Objects Differences Analyzer (VDDA) tool allows you to compare groups or individual Ingres database objects, either in the current installation or saved into a snapshot file. It also allows you to visualize the differences in the database objects.

The following options are supported:

- Performing the comparison either at the installation level (that is, comparing any database objects present in an installation), or at the schema level, (that is, comparing objects owned by a given user within a database). In both cases, you can limit the comparison to certain object types.
- Saving such groups of database object definitions into snapshot files for later comparison with current or other saved database object definitions. This allows comparisons within an installation over time.

The list of differences generated by VDDA includes one line for each difference in a property of the given database object. The bottom status line indicates the number of objects with differences and the number of differences found in the comparison.

For more information, see online help for VDDA.

Visual Configuration Differences Analyzer

The new Visual Configuration Differences Analyzer (VCDA) tool allows you to compare the configuration information for two Ingres installations. VCDA enables you to take snapshots of the installation and compare these snapshots with either a snapshot taken at some point in the past or with a configuration snapshot taken on another machine.

For example, after you install Ingres and have tuned the configuration to meet your needs, take a snapshot of the configuration. If you encounter problems with the installation later on, take another snapshot of the configuration and compare it to the earlier snapshot to determine if any configuration changes have contributed to the problem. Keep an on-going record of configuration changes by taking a snapshot of the installation each time you change its configuration.

VCDA snapshots contain information taken from the config.dat file, the symbol table, and the vnode database, as well as environment variables set at the system and user level.

Specifically, VCDA allows you to:

- Save a snapshot of the current installation configuration into a file
- Compare two snapshot files, or the current installation with a saved snapshot
- Restore selective groups of configuration parameters from a saved snapshot

VCDA lists the differences with an associated icon that distinguishes between those parameters that are different and those that exist in one snapshot but not in the other.

VCDA uses information from the config.dat file, which typically contains parameters that apply to the local host name only. However, if you want to concatenate the contents of all your config.dat files into a single config.dat file to distribute across multiple environments, VCDA manages this situation as follows:

- The host name under which a snapshot is saved becomes part of the snapshot information.
- If VCDA detects that host names other than the snapshot host name are managed within the config.dat information of the snapshot, VCDA displays them.
- A host name mapping option is available in that situation, so that VCDA can compare the additional host names' configuration parameters. If this option is not used, the parameters are compared, including their host names (that is, only parameters that are identical for the same host name in the two snapshots are considered identical).

For more information, see online help for VCDA.

Export Assistant

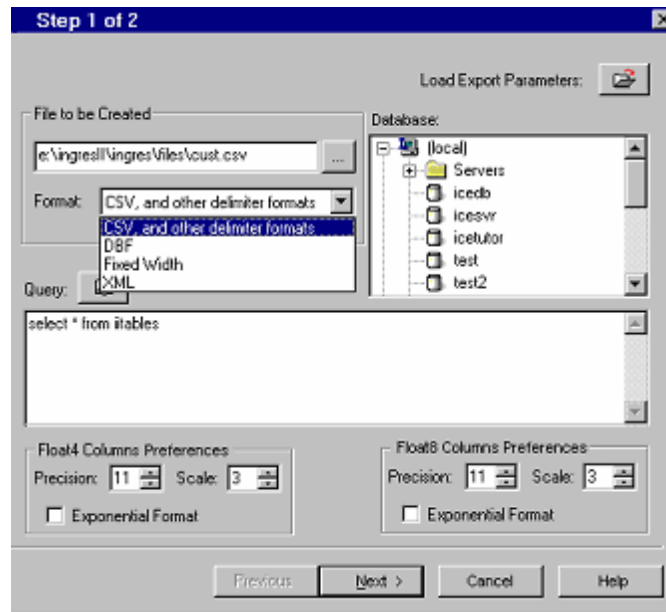
The new Export Assistant complements the existing Import Assistant in both design and function. The Export Assistant is a wizard designed to simplify the task of exporting Ingres (or Enterprise Access) data into external files. Specifically, you can export data into the following file formats:

- .csv (and other delimiter formats)
- .xml
- .dbf
- fixed widths

The Export Assistant is accessible from the Start menu and from the Ingres Visual Manager and Visual DBA tools. You can also invoke the Export Assistant from the command line.

The Export Assistant uses a two-step process to build the export file.

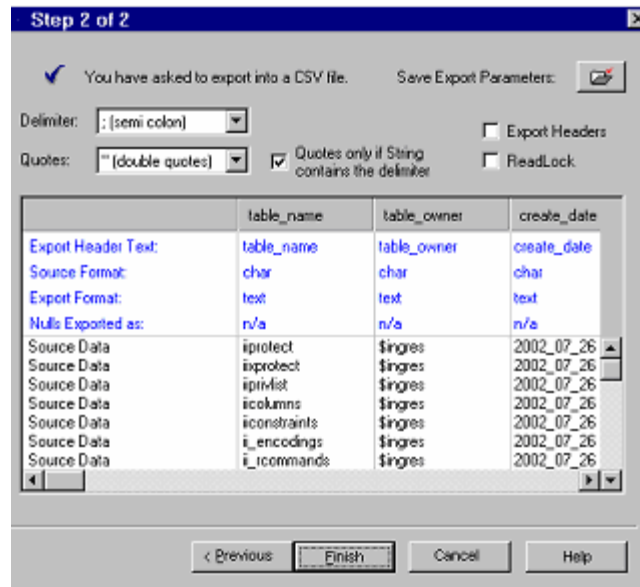
Step 1: On the first dialog, select a path and filename for the file you want to create.



Select the desired file format. In the From Database tree, select the node/gateway and database on which you want to perform the query that builds your export file. You can enter this query in the bottom part of the window, or click the Query button to launch the SQL Assistant to help you build the query.

The Load Export Parameters button allows you to load a set of parameters that had been previously saved in Step 2 of the Export Assistant. These parameter files have an extension of .ii_export. Double-clicking on a file with this extension in Windows Explorer launches the Export Assistant with the corresponding Export Parameters loaded.

Step 2: Based on the query you entered in Step 1, the next dialog displays a set of proposed default values for several editable export parameters.



It also displays the first 100 rows of the export data created from the query to help you edit the parameters.

Note: The dialog shown in this example corresponds to a CSV format export file and shows the controls specific to that format. The controls vary depending on the file format chosen in Step 1.

The Save Export Parameters button allows you to save the parameters entered in Steps 1 and 2 for later reuse in Step 1.

For more information, see online help for the Export Assistant.

Connectivity Enhancements

The following new features provide connectivity enhancements.

Support for JDBC 3.0 API

Support for the JDBC 3.0 API includes three components: an installation server, a Java client driver, and an information utility. For more information, see the *Connectivity Guide*.

Data Access Server

The Data Access Server (DAS) runs as part of a standard Ingres installation. The DAS translates JDBC requests from the Ingres JDBC Driver into Ingres internal format and forwards the request to the appropriate DBMS Server. The DAS supports the same network protocols and port designations as the Communications Server.

The DAS also supports the new Ingres .NET Data Provider component that enables high-performance native .NET access to Ingres data sources and delivers Ingres data to the Microsoft .NET Framework.

Through the DAS, a JDBC client and Ingres .NET Data Provider have full access to Ingres, Enterprise Access, and EDBC databases. The DAS can also access database servers on remote machines using Ingres Net.

For more information, see the *Connectivity Guide*.

JDBC Driver

The Ingres JDBC Driver is a pure Java implementation of the JDBC 3.0 API released with the Sun Java 2 SDK, version 1.4. The driver supports application, applet, and servlet access to Ingres data sources through the Data Access Server.

The JDBC driver provided in Ingres 2.6 continues to be supported in Ingres 2006. For migration instructions related to the JDBC driver, see the *Migration Guide*.

The Ingres JDBC Driver with the DAS supports the following JDBC 3.0 features:

- Boolean data type (similar to Bit)
- Savepoints
- Named procedure parameters
- Auto-generated keys
- Connection Pool Configuration

The Ingres JDBC Driver is delivered as a single Java archive file, `ijdbc.jar`, located in the library directory (`lib`) of the Ingres installation. Access to the driver can require, depending on the Java environment used, adding the Java archive to the `CLASSPATH` environment setting or as a resource in the appropriate utility. For browser/applet access, the Java archive must be copied to the Web Server directories.

JDBC Information Utility

The JDBC 3.0 API support includes a JDBC information utility, `JdbcInfo`. This utility displays the Ingres JDBC Driver internal release information. The class files for the `JdbcInfo` utility are located in the library directory (`lib`) of the Ingres installation.

For more information, see the *Connectivity Guide*.

Updateable Result Sets in JDBC

The Ingres JDBC driver supports updateable result set features of the JDBC 2.1 API. Updateable result sets permit an application to update or delete the current row of the result set, or insert rows into the associated table using methods provided by the JDBC `ResultSet` class. A new class, `RsltUpdt`, has been added as an extension to the cursor result set class, `RsltCurs`, to support updateable result sets. The result set methods associated with the `RsltUpdt` class are listed below.

The ability to update a result set is determined by calling the following method:

```
ResultSet.getConcurrency()
```

The current row of a result set can be deleted using the following method:

```
ResultSet.deleteRow()
```

Columns values of the current row can be set using the following methods:

```
ResultSet.updateAsciiStream()
```

```
ResultSet.updateBigDecimal()
```

```
ResultSet.updateBinaryStream()
```

```
ResultSet.updateBoolean()
```

```
ResultSet.updateByte()
```

```
ResultSet.updateBytes()
```

```
ResultSet.updateCharacterStream()
```

```
ResultSet.updateDate()
```

```
ResultSet.updateDouble()
```

```
ResultSet.updateFloat()
```

```
ResultSet.updateInt()
```

```
ResultSet.updateLong()
```

```
ResultSet.updateNull()
```

```
ResultSet.updateObject()
```

```
ResultSet.updateShort()
```

```
ResultSet.updateString()
```

```
ResultSet.updateTime()
```

```
ResultSet.updateTimestamp()
```

Once column values have been set, the changes can be saved or dropped using the following methods:

```
ResultSet.updateRow()
```

```
ResultSet.cancelRowUpdates()
```

To insert a row, the result set current position must be moved to a special reserved row. The following methods control the positioning of the result set and the insertion of rows:

```
ResultSet.moveToInsertRow()
```

```
ResultSet.moveToCurrentRow()
```

```
ResultSet.insertRow()
```

The following methods can be used to determine the status of a result set row:

```
ResultSet.rowDeleted()
```

```
ResultSet.rowInserted()
```

```
ResultSet.rowUpdated()
```

For more information, see the *Connectivity Guide*.

.NET Data Provider and Visual Studio .NET Integration

This release of Ingres introduces support for the Microsoft .NET Framework and Visual Studio .NET application development tools.

Ingres .NET Data Provider

The Ingres .NET Data Provider is a .NET component that enables high-performance native .NET access to Ingres data sources and delivers Ingres data to the Microsoft .NET Framework.

The Ingres .NET Data Provider offers a series of .NET types to describe the user's data, .NET provider classes to manipulate the data, and connection pooling to efficiently manage data connections.

The design and naming conventions of the Ingres .NET Data Provider's data types, classes, properties, and methods follow the same pattern as the Microsoft .NET Data Providers. Consequently, developers who are familiar with the Microsoft providers can easily develop or convert existing code from Microsoft databases to Ingres databases.

All Ingres .NET Data Provider modules are written in C#, a managed .NET language with full access to every .NET Framework capability. Even though the data provider is written in C#, any managed language such as VB.NET or J# can use the data provider because of .NET's language interoperability feature.

For more information, see the *Connectivity Guide*.

Visual Studio .NET Integration

The .NET Framework was written with design-time support. Integration with the Visual Studio .NET visual tools allows programmers to drag-and-drop the Ingres .NET Data Provider design component onto a design surface such as a the Windows Form Control (WinForm). Integration also includes the following design components:

- Data Adapter Configuration Wizard: Enables programmers to specify the design properties of the Ingres DataAdapter object.
- Query Builder: Enables programmers to build SQL statements that the Ingres .NET Data Provider uses to retrieve and modify database information.
- Parameter Collection Editor: Enables programmers to add parameters to the Command component.

For more information, see the *Connectivity Guide*.

Ingres ODBC Administrator

Prior to this release, ODBC users on non-Windows platforms were required to manually edit the `odbc.ini` configuration file to define a data source. This release introduces an Ingres ODBC Administrator utility for non-Windows platforms. Supported platforms include UNIX, Linux, and VMS. The new ODBC Administrator enables users to:

- Create, edit, and delete data source definitions
- View configuration details about a particular data source
- Display a list of installed drivers and view configuration details for a selected driver
- Define an alternate path for accessing driver definitions
- Define an alternate path for accessing data source definitions
- Turn ODBC tracing on or off
- Test a data source connection

For more information, see the *Connectivity Guide*.

WinSock 2.2 API TCP/IP Protocol Driver for Windows

This release provides a new Windows TCP/IP protocol driver that takes advantage of the latest Windows Winsock 2.2 API. The new driver (tcp_ip) removes an architectural limitation of the previous implementation (wintcp), which resulted in a performance problem when a client application used many INSERT statements or a large number of single SELECT statements. Also, connection attempts sometimes failed when multiple connects were attempted simultaneously.

To allow a smooth migration from the existing protocol driver to the new one, the existing wintcp protocol driver is included in this release, but will be removed in the future. We recommend that you use the tcp_ip protocol driver.

For more information, see the *Connectivity Guide*.

ODBC Call-level Interface

The Ingres ODBC Call-level Interface (CLI) provides access to the ODBC application environment without the need to use third-party software. It is installed when you install the Ingres ODBC Driver and is supported on all platforms on which Ingres runs.

The Ingres ODBC CLI performs the following functions:

- Optionally determines driver characteristics from ODBC configuration files
- Loads and unloads the ODBC driver into and from application memory
- Maps the driver manager API to the driver API
- Performs basic error checking
- Provides thread safety
- Provides ODBC tracing
- Provides function templates, type definitions, and constant definitions for ODBC applications

For more information, see the *Connectivity Guide* and the *System Administrator Guide*.

New Features for Linux

The following features are new for the Linux platform.

RPM Packaging

Ingres 2006 for Linux is packaged using the Red Hat Package Manager. For more information, see the *Getting Started* guide for Linux.

KDE/GNOME Desktop Integration

When Ingres 2006 is installed on Linux, it installs the Getting Started guide for Ingres for Linux on the desktop. It also creates two program groups, one for the Visual DBA suite and the second for the complete set of Ingres documentation. The user has the information needed to start and use Ingres within minutes. For more information, see the *Getting Started* guide for Linux.

Linux Cluster Support

For information about the Ingres Cluster Solution on Linux, see Ingres Cluster Solution for Linux (see page 7).

Additions to User Documentation

The following new guides have been added to the Ingres documentation set for Ingres 2006.

Getting Started Guide for Linux

The *Getting Started* guide for Linux provides installation and configuration information for Ingres users who are installing and configuring Ingres on the Linux platform.

Connectivity Guide

The *Connectivity Guide* includes information previously located in the *Getting Started* guide and the *System Administrator Guide*. This central resource makes it easier for users to locate connectivity-related information for supported Ingres APIs. Specifically, the *Connectivity Guide* provides the following information:

- Instructions for installing, configuring, and using Ingres Net and Ingres Protocol Bridge.
- An introduction to the JDBC components that enable JDBC connectivity to Ingres data sources. It provides a description of each component, a list of supported API features, driver and server configuration information, and guidelines for implementing Java applications in the Ingres environment.
- An introduction to the Ingres ODBC components that enable ODBC connectivity to Ingres data sources. This information includes a description of each component, a list of supported API features, data source configuration instructions, connection string keyword definitions, and guidelines for implementing ODBC-enabled applications in the Ingres environment.
- An introduction to the Ingres .NET Data Provider components that enable high-performance native .NET access to Ingres data sources and deliver Ingres data to the Microsoft .NET Framework. This information explores how components and wizards in the provider objects help integrate the Ingres .NET Data Provider with MS Visual Studio .NET to aid in the development of .NET applications that access Ingres data.
- Troubleshooting tips and specific configuration information for each of the network protocols supported by Ingres.

Note: The connectivity information presented in this *Connectivity Guide* for accessing Ingres 2006 databases also applies to Enterprise Access and EDBC products and the databases they support.

Interactive Performance Monitor User Guide

The Interactive Performance Monitor (IPM) tool provides an interactive, forms-based interface for displaying information about servers, sessions, and locking and logging activity in an active Ingres installation. The *Interactive Performance Monitor User Guide* explains how to use this tool to perform the following tasks:

- Monitoring the state of server processes, logging and locking systems, database servers, log files, and user sessions
- Analyzing performance problems, including locked tables and blocked or deadlocked transactions
- Maintaining and tuning the system using IPM information such as identifying transactions that are neither committed nor rolled back over a long period of time

Chapter 2: Changes to Existing Features

This chapter describes changes made to existing features of Ingres.

Enhanced Unicode Support

Unicode is a standard method of storing character data for multinational situations. Many applications require the ability to store and retrieve Unicode data. Ingres 2006 supports the second phase of a phased implementation of full Unicode support and extends support to the UTF-16 encoding scheme.

For details on this feature, see the *SQL Reference Guide*, the *OpenSQL Reference Guide*, and the *Command Reference Guide*.

Unicode Coercion

Ingres coerces between Unicode data types and non-Unicode data types, including nchar/nvarchar and char/varchar types. Unicode coercion includes the following features:

- Unicode and non-Unicode columns can be joined together in a query.
- The Ingres copy statement (and by inference, copydb and unloaddb) can be executed to generate or load non-binary data files where a table contains at least one Unicode column.
- Query strings containing non-Unicode literals that reference Unicode columns are coerced to Unicode before the DBMS executes the query.

Collation Sensitive Support for Wildcard Searching

Ingres performs wildcard searches involving Unicode strings by making use of a case insensitive collation table.

Complex Query Optimization

The Ingres query optimizer has been enhanced to handle queries that reference large numbers of tables and indexes with shorter response time. For more information, see the *SQL Reference Guide*.

Increased Range Table Limit

The limit for the number of table references in a query has increased from 30 to 126.

This limit refers to the sum of explicit table references (even if the same table appears several times in a query), explicit view references, and tables and views included by the expansion of a view definition. It also refers to the sum of such references in all subselects and unioned selects of a query.

For more information, see the *SQL Reference Guide*.

JDBC User ID Enhancements

In previous releases of Ingres, the JDBC driver required a user ID and password that was valid in the context of the DBMS platform, even for local connections where the local user ID is sufficient. The JDBC driver also failed to use virtual node (vnode) login information when making remote connections, and was not able to access private vnode definitions.

In Ingres 2006, the JDBC driver no longer requires a user ID and password when the Data Access Server is running on the same platform as the Java client. Instead, for local connections, the local user ID is sufficient to establish the DBMS Server connection.

When the Data Access Server and DBMS Server are on separate platforms, a vnode is required in the target database specification that defines how the connection between the Data Access Server and DBMS Server is to be established. The vnode provides the connection information while the JDBC user ID and password are used to access the remote DBMS Server.

A new driver connection property/URL attribute allows the JDBC application to control the context (Data Access Server (local) or DBMS Server (remote)) in which the user ID is used.

When used in the Data Access Server context, the user ID and password allow access to the private vnode information for the user ID provided, and both the login and connection information from the vnode is used to access the remote DBMS Server.

When used in the DBMS Server context, global vnode definitions are used for (nonsensitive) connection information and the provided user ID and password are used to access the remote DBMS Server.

For more information, see the *Connectivity Guide*.

Increased Column Limit

The limit on the number of columns per table has increased from 300 to 1024. This feature is upwardly compatible; however, programs written to take advantage of the new limit cannot be used with earlier Ingres releases that included the 300-column limit. For more information, see the *SQL Reference Guide*.

VDBA Enhancements

The following VDBA enhancements have been implemented for Ingres 2006. These enhancements are available only on Windows, Solaris, HP-UX, and AIX platforms.

VDBA Architecture Split

Two VDBA utilities, SQL/Test and Performance Monitor, have been split from the VDBA architecture and are now called Visual SQL and Visual Performance Monitor. By isolating these components as smaller, stand-alone executables, users benefit from increased response time because they no longer need to launch the whole VDBA executable.

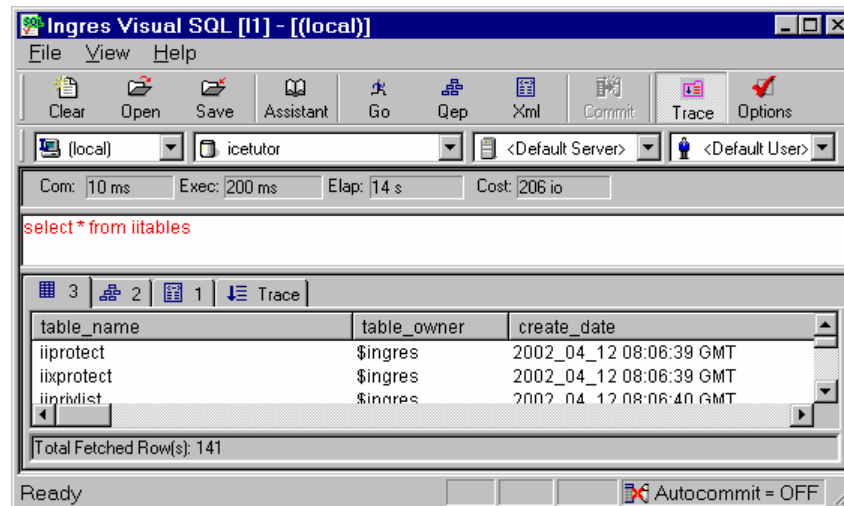
When launched in VDBA, these new components also help avoid VDBA locking itself. This can happen, for example, when there is an uncommitted query on a table in an SQL Test window, and an operation is performed on the same table in a DOM window opened in the same VDBA instance.

The SQL Assistant has also been split into a stand-alone DOM component. The SQL Assistant can now be accessed from the Ingres Export Assistant, in addition to Ingres Network Utility and VDBA, to help build SQL queries.

Visual SQL and Visual Performance Monitor are accessible from the Start menu, Ingres Visual Manager, Network Utility, VDBA, and the command line.

Visual SQL

The new stand-alone executable (vdbasql) for Visual SQL displays the following window.



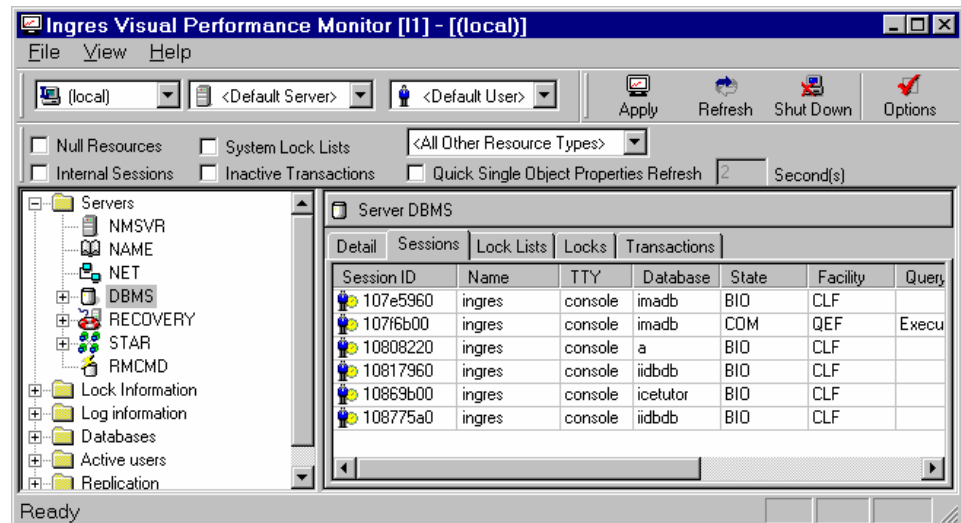
This window allows all the functionality of the previous SQL/Test window in VDBA, but has the following additional toolbar controls:

- New combo-boxes for choosing a node, server class, and impersonated user. (The equivalent function is provided in VDBA through the branch chosen in the nodes tree when launching an SQL/Test window.)
- A Commit button that allows you to directly commit the current transaction without the need to type commit in the query area.
- An Options button that provides access to the properties of the SQL Test control. (This is equivalent to the SQL/Test Preferences in previous VDBA releases, plus the session timeout parameter that was global to VDBA.)

In addition, the status bar now provides the autocommit state of the transaction.

Visual Performance Monitor

The new stand-alone executable (vdbamon) for Visual Performance Monitor displays the following window.



This window allows all the functionality of the previous Performance Monitor window in VDBA, but has the following additional toolbar controls:

- New combo-boxes for choosing a node, server class, and impersonated user. (The equivalent function is provided in VDBA through the branch chosen in the nodes tree when launching a Performance Monitor window.)
- An Apply button that displays the monitor information corresponding to the chosen node (and, if applicable, server class and impersonated user).
- An Options button that provides access to the properties of the Performance Monitor control. (This is equivalent to the Monitor Preferences in previous VDBA releases, plus the background refresh preference, the Grid In List option, and the session timeout and max number of sessions parameters that were global to VDBA.)

Additional Properties

The Properties dialog is accessible from the new Visual SQL and Visual Performance Monitor executables, as well as from VDBA. (This was previously called the Preferences dialog.)

The SQL/Test Priorities dialog now includes a session timeout parameter that was previously global to all VDBA sessions.

The Performance Monitor Properties dialog includes the following additional parameters:

- Background refresh for monitor windows (includes the choice of having background refresh active or not)
- Session timeout and the maximum number of sessions used within the Performance Monitor control's internal session cache
- Grid In List option that displays a grid in the Detail Information pane that allows for easier viewing of data

In addition, Visual SQL and Visual Performance Monitor provide a new menu option, Save Preferences As Default. If this option is selected, which is the installation default, the properties are permanently stored and used every time the utility is executed. If this option is unselected, the selected properties are only used for the utility's current instance.

Notification of Metadata Changes

Notification of changes to Ingres metadata includes the following features in VDBA:

- VDBA is notified by the DBMS of any metadata changes and automatically refreshes itself.
- The refresh preferences are no longer available.
- A new preference (called other servers) provides background refresh for those servers that do not provide metadata change notification (for example, Ingres DBMS with releases older than Ingres 2006 and gateway servers).
- The background refresh of monitor data and node definitions are placed in the Monitor and Nodes preferences.

For more information, see the online help for VDBA.

Ingres Visual Manager Enhancements

The Ingres Visual Manager contains the following enhancements:

- Direct access to all Ingres tools
- Can register alerts in the OS event log
- Message explanations

Note: These enhancements are only available on Windows, Solaris, HP-UX, and AIX platforms.

Direct Access to All Ingres Tools


Ingres Visual Manager provides direct access to all Ingres visual tools. The only exception is the Ingres Service Manager; its full functionality is already provided by Ingres Visual Manager. You can access these tools in the following ways:

- Right-click the tray toolbar to display a menu:



The Ingres 2006 documentation is accessible from the menu.

- Select a tool from the IVM toolbar.



Only the most frequently used tools are accessible from this toolbar. Use the above menu to access all other tools.

Registering Alerts in the OS Event Log

When an Alert message is written to the Ingres errlog.log file, IVM indicates the alert through a special icon change in the tray toolbar and in the IVM window tree. IVM also lets you set additional preferences for alert notification using the Preferences dialog. These additional preferences include:

- **Sound** – When selected, you are alerted of an event by a sound (beep).
- **Message Box** – When selected, you are alerted of an event by a message box if there are unread Alert messages.

A new preference has been added to the Preferences dialog:

- **OS Event** (for new Events only) – When selected, the full text of all Alert messages resulting from new events are logged in the Operating System [Application] Event log in addition to the Ingres errlog.log file. This new preference allows third-party tools to be signaled by Ingres Alert messages.

In addition, if the operating system supports it (as with Windows), the message category (class) and number is stored as information belonging to the event. External applications that are monitoring error messages for the Ingres installation are given the error category and number without having to parse the text of the message.

If the OS Event option is selected, you can also set preferences for generating the following specific operating system events:

- One event is generated whenever the OS Event option checkbox has been selected or deselected and the Preferences dialog has been validated.
- One or more events are generated when the Alert States have been changed (for example, from “alert” to “non-alert”) in the Messages Categories and Notification Levels dialog.
- One event is generated if IVM detects that the errlog.log file has been manually changed.

Viewing Message Explanations

To help you respond to Ingres messages, IVM now displays an explanation for any message in the errlog.log file.

Select the desired message on any page or window in IVM that allows message selections and view its explanation by clicking the Message Explanation toolbar button.

Configuration Rules System Enhancements

The Configuration Rules File System has been modified to handle negative values and decimal values.

Shadow Copy of the Symbol Table

When setting, unsetting, or changing environment variables, a backup of the symbol.tbl file (symbol.bak) is maintained. If the original symbol.tbl file becomes corrupted, you can use the backup symbol table file to restore it. A history of updates to the symbol.tbl file is maintained in the symbol.log file. The symbol.bak and symbol.log files are located in the same location as the original symbol.tbl file. For more information, see the *System Administrator Guide*.

Additional Join Functionality

In addition to inner, left, right, and full joins, users can request cross joins (effectively inner joins without an clause) and natural joins. Also, the on clause can be replaced by a using clause that contains a list of columns, each of which appears in both tables being joined. Instead of the explicitly coded join qualification of the on clause, the using clause applies one equijoin predicate for each column pair in the list of columns. For example, "... a left join b using c1, c2 ..." is identical to coding "... a left join b on a.c1 = b.c1 and a.c2 = b.c2" For more information, see the *SQL Reference Guide*.

Improved Out-of-the-Box Configuration Defaults

Ingres configuration defaults have been updated to reflect the current hardware environments.

Improved IMA Support

Ingres 2006 introduces an additional Ingres Management Architecture (IMA) component, IMP.

IMA provides the framework for accessing system data for monitoring and managing installations through SQL without affecting the underlying operation of the product. For more information, see the *System Administrator Guide*.

Createdb Enhancements

Ingres 2006 adds a flag to createdb, which allows the specification of a nondefault page size for catalogs. For more information, see the *Command Reference Guide*.

Alterdb Enhancements

Ingres 2006 adds three new flags to the alterdb utility, which allow for the deletion of invalid checkpoints, the deletion of a specific checkpoint, and a non-Unicode enabled database to be Unicode enabled. For more information, see the *Command Reference Guide* and the *Database Administrator Guide*.

Terminal Monitor Enhancements

Ingres 2006 adds support for the `-p` flag to the terminal monitor (that allows a password to be specified) and support for the `-r` flag (that allows a role name and optional role password to be added).

Also, command completion and command history recall are supported for the Linux operating system.

For more information, see the *SQL Reference Guide* and the *Command Reference Guide*.

Enhancements for Log Full

Ingres 2006 allows the user to dictate behavior when a log full situation occurs. The options provided are `commit`, `abort`, or `continue`. For more information, see the *SQL Reference Guide*.

Extended B-tree Limits

Ingres 2006 extends the 440-byte limit on the maximum width of a B-tree key.

Installer Enhancements

Ingres on Windows environments is delivered in Microsoft Windows Installer format. Microsoft Windows Installer version 2.0 is required on the machine on which you are installing Ingres components.

The installer on Windows presents a setup wizard to guide you through the installation process. The setup wizard lets you select a Complete or Custom install. The Custom installation lets you choose individual components for installation.

For more information, see the *Getting Started* guide.

Installation as a User Other Than ingres

With Ingres 2006, the product can be installed as a user other than ingres. Ingres security is retained. When installing Ingres, you can specify a user ID (and its associated group ID) that owns the Ingres installation. During installation, this system administrator ID is automatically created and given the required permissions. If they are not previously defined, the user ID and group ID are added to your UNIX or Windows systems.

For more information, see the *Getting Started* guide.

Supportability Enhancements

Product enhancements that improve the supportability of Ingres include dumping queries to the error log file upon certain error conditions (for example, an optimizer time-out or an exhaustion of resources), and modifying the ingstart process so that the exact version and patch information is written to the error log file on startup.

Help System Enhancements

Each Ingres visual tool now has its own Help system that is independent of other visual tools. This separation of Help systems makes it easier and quicker for users to navigate through the list of Help system topics. In addition, search results now display only those topics that are applicable to the visual tool in use. Standalone Help systems are now provided for the following visual tools:

- VDBA
- Configuration Manager
- Ingres Visual Manager
- Import Assistant
- Export Assistant
- Journal Analyzer
- Network Utility
- Visual SQL
- Visual Performance Monitor
- Visual Database Objects Differences Analyzer
- Visual Configuration Differences Analyzer