

# Ingres<sup>®</sup> 2006

## System Administrator Guide

**INGRES<sup>®</sup>**

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

---

<b>Chapter 1: Introducing Ingres System Administration</b>	<b>15</b>
Audience .....	15
Special Considerations .....	15
Query Language Used in This Guide.....	15
System-specific Text in This Guide.....	16
Path Notation in This Guide .....	16
Terminology Used in This Guide .....	16
Syntax Conventions Used in This Guide .....	17
<b>Chapter 2: Understanding Ingres Components</b>	<b>19</b>
Responsibilities of a System Administrator .....	19
Components of an Ingres Installation.....	20
DBMS Server.....	21
Query Processing .....	21
The Ingres Tool or Application Process.....	21
DBMS Server Process .....	22
Query Environment.....	22
DBMS Server Facilities .....	22
Abstract Data Type Facility.....	23
Compatibility Library Facility .....	23
Data Manipulation Facility .....	23
Optimizer Facility.....	23
Parser Facility .....	24
Query Execution Facility .....	24
Query Storage Facility .....	24
System Control Facility.....	24
Relation Description Facility .....	24
General Communication Facility .....	25
Name Server .....	25
Communications Server.....	26
Data Access Server .....	26
Bridge Server.....	26
General Communications Architecture .....	26
Internet Communication .....	27
Querying and Reporting Tools .....	27
Visual Tools.....	27
Visual Manager.....	27

---

Visual Performance Monitor .....	28
Visual DBA .....	29
Remote Command Server .....	30
Configuration Manager .....	31
Network Utility .....	31
Export Assistant .....	31
Import Assistant.....	32
Journal Analyzer.....	32
Visual SQL.....	32
Visual Configuration Differences Analyzer.....	32
Visual Database Objects Differences Analyzer .....	32
Logging and Locking Systems .....	33
Lock Manager .....	33
Logging Facility .....	33
Recovery Process.....	34
Archiver Process.....	35
Transactions Recovery.....	36
Ingres Log Files.....	37
Optional Log Files .....	40

### **Chapter 3: Configuring Ingres** **43**

Configurable Components .....	43
Configuration Tools .....	43
Access Configuration Manager Utility .....	44
Access Configuration-By-Forms Utility .....	44
Access VCDA Utility.....	44
Component Configurations .....	45
Configurable Properties for a Component .....	45
Configuration Snapshot .....	46
Configuration Parameters .....	46
DBMS Server Configuration .....	49
Fast Commit Option .....	50
Write Behind Option.....	51
Data Access Server Configuration.....	52
Communications Server Configuration.....	52
Bridge Server Configuration.....	53
Star Server Configuration.....	53
Name Server Configuration.....	54
Recovery Server Configuration.....	54
Internet Communications Configuration.....	55
Locking System Configuration .....	55
Logging System Configuration.....	56

---

Multiple Log Partition Configuration .....	56
Primary Transaction Log and Dual Transaction Log .....	57
Non-Configurable Components .....	57
History of Configuration Changes.....	57

## **Chapter 4: Setting Environment Variables and Logicals** **59**

Scope of the Environment Variables.....	59
Symbol Table .....	60
Windows Environment Variables.....	60
UNIX Environment Variables .....	61
VMS Logicals .....	61
Installation-Wide Environment Variables and Logicals.....	61
View Environment Variables in Windows and UNIX .....	62
View Environment Variables in VMS.....	63
Example: Set the ING_EDIT Variable.....	64
User-Definable Environment Variables and Logicals .....	65
Customization of Query Runtime Environment.....	70
Syntax Rules for Startup Files and Environment Variables.....	71
Dayfile .....	72
DBNAME_ING.....	72
DBNAME_SQL_ INIT.....	74
II_SQL_INIT .....	76
ING_SET .....	78
ING_SET_ DBNAME.....	80
ING_SYSTEM_ SET .....	82
INIT_INGRES .....	84
Startup .....	86
Startsql.....	86
Installation Locations.....	87
When Not to Change Installation Locations.....	87
Location Variables that Do and Do Not Require Unloading (UNIX and Windows) .....	88
Change an Installation Location with Unloading (Windows) .....	89
Change an Installation Location with Unloading (UNIX) .....	90
Change Location Variables (Other than ING_ABFDIR) on VMS .....	90
Use of Symbolic Links (UNIX) .....	91

## **Chapter 5: Customizing Ingres** **93**

Ingres Character Sets.....	93
Default Character Sets .....	94
Local Collation Sequences .....	94
Supported Collation Sequences.....	94

---

Archiver Exit Script (acpexit) .....	100
Customization of Archiver Exit Script .....	101
Archiver Exit Script Parameters .....	101
How You Use a Custom Mapping File for Unicode Coercion .....	102

## **Chapter 6: Troubleshooting Ingres** **103**

Process of Troubleshooting .....	103
Troubleshooting Tools .....	103
Determine the Problem Area .....	104
Error Log Files .....	104
Optional Error Log Files .....	105
Display Value for II_CONFIG .....	105
View List of Log Files .....	106
Check the Error Log Files .....	106
Find Your Problem Category .....	107
Troubleshoot Startup, Shutdown, or Configuration Problems .....	108
Check Ingres Installation on Windows .....	110
Check Ingres Installation on UNIX .....	112
Check Ingres Installation on VMS .....	116
Detect Ingres Startup Problems on System Reboot (Windows) .....	119
Detect Ingres Startup Problems on System Reboot (UNIX) .....	120
Detect Ingres Startup Problems on System Reboot (VMS) .....	121
Check Shutdown Problems on Windows .....	122
Check Shutdown Problems on UNIX .....	124
Check Shutdown Problems on VMS .....	127
Ingres Processes on Windows .....	131
Check for Name Server Errors on Windows .....	131
Check for Communications Server Process Errors on Windows .....	132
Check for Bridge Server Errors on Windows .....	132
Check for ICE Server Errors on Windows .....	132
Check for Remote Command Process Errors on Windows .....	132
Check for Recovery Process Errors on Windows .....	133
Check for Archiver Process Errors on Windows .....	134
Check for DBMS Server Process Errors on Windows .....	135
Check for JDBC Server Errors on Windows .....	135
Check for Data Access Server Errors on Windows .....	136
Ingres Processes on UNIX .....	136
Name Server Errors on UNIX .....	136
Check for Communications Server Process Errors on UNIX .....	138
Check for Bridge Server Process Errors on UNIX .....	139
Check for ICE Server Process Errors on UNIX .....	139
Recovery Process Errors on UNIX .....	140

---

Check for Remote Command Process on UNIX .....	141
Archiver Process on UNIX .....	141
Check for DBMS Server Process Errors on UNIX .....	143
Check for JDBC Server Process Errors on UNIX .....	144
Check for Data Access Server Process on UNIX .....	144
Ingres Processes on VMS .....	145
Name Server Errors on VMS .....	145
Communications Server Process on VMS .....	146
Bridge Server Errors on VMS .....	146
Recovery Process Errors on VMS .....	147
Remote Command Process Errors on VMS .....	148
Archiver Process Errors on VMS .....	149
Check for DBMS Server Process Errors on VMS .....	151
Check for JDBC Server Process Errors on VMS .....	152
Data Access Server Errors on VMS .....	152
Check for Data Access Server Errors on VMS .....	153
Problems with Tools Startup .....	153
DBMS Server Stopped .....	154
Database Connection Problems .....	155
Inconsistent Databases and Recovery .....	157
Automatic Recovery .....	157
Inconsistent Database .....	160
Performance Problems .....	168
Flow Diagram for Troubleshooting Performance Problems .....	168
Diagnose DBMS Server Due to Logging System Problem .....	169
Resource and Maintenance Problems .....	174
Tools for Identifying Operating System Problems .....	175
What You Need Before Contacting Technical Support .....	179
Windows Installations .....	179
UNIX Installations .....	179
VMS Installations .....	180

## **Chapter 7: Using Monitoring and Tracing Tools** **183**

Supported Monitoring and Tracing Tools .....	183
System Utilities .....	184
Operating System Utilities .....	185
Windows Operating System Utilities .....	185
UNIX Operating System Utilities .....	187
VMS Operating System Utilities .....	193
Vendor Utilities .....	196
Error Messages .....	197
Error Message Format .....	197

---

Message Help Files.....	197
Fatal Errors.....	198
Non-Fatal Errors.....	198
Ingres Facility Codes.....	198
Log Files.....	202
Transaction Log File.....	202
Error Log.....	203
Archiver Log.....	203
Recovery Log.....	204
Primary Configuration Log Files.....	204
Optional Configuration Log Files.....	205
Other Optional Log Files.....	205
Trace Utilities.....	206
ODBC Call-level Interface.....	207
Set Statement.....	207
Set Options for Tracing Queries.....	211
Canceling Set Options.....	213
I/O Tracing.....	215
User-Server Communications.....	218
UNIX Trace Facilities.....	220
Bourne Shell -x Option.....	221

## **Chapter 8: Managing Your System and Monitoring Performance** **223**

Visual Manager.....	223
Functions of Visual Manager.....	224
Visual Manager Window.....	225
System and User Parameter Configuration Through IVM.....	226
Set Parameter Configuration Through IVM.....	226
Message and Notification Management Through IVM.....	227
Event Monitoring Management Through IVM.....	227
Alert Events Notification Settings in IVM.....	228
View Message Explanations in IVM.....	228
Component Monitoring Through IVM.....	228
Server Monitoring Through IVM.....	229
Logging System Monitoring Through IVM.....	229
Locking System Monitoring Through IVM.....	229
Access to Visual Tools and Documentation Through IVM.....	230
Visual Performance Monitor.....	231
Visual Performance Monitor Window.....	232
Visual Configuration Differences Analyzer.....	233
Configuration Snapshot File.....	233
Comparison of Configuration Snapshots.....	234



---

Restoration of Configuration Parameters Through VCDA .....	234
How VCDA Handles Concatenated config.dat Files.....	235
<b>Chapter 9: Analyzing and Recovering Journaled Transactions</b> .....	<b>237</b>
Journal Analyzer .....	237
Start Journal Analyzer .....	238
Journal Analyzer Window .....	238
Transaction Views .....	239
Database Level Journaled Transactions .....	239
Table Level Journaled Transactions.....	242
Select Transactions and Row Changes .....	243
Recover Transactions or Individual Row Changes .....	244
Recover "Whole" Transactions or Individual Rows .....	245
Rules, Constraints, and Integrities .....	245
Order of Transactions and Statements in Transactions .....	246
How to Check If Rows Have Changed After the Transaction.....	247
Users of Transactions .....	248
Number of Rows Affected by Each Individual "Reverse Statement" .....	249
Recover Immediately or Generate External Scripts Option .....	249
Redo Transactions and Individual Row Changes .....	251
<b>Chapter 10: Understanding Ingres Management Architecture</b> .....	<b>253</b>
Ingres Management Architecture (IMA) .....	253
IMA Audience .....	254
IMA Architecture .....	254
Management Information Base.....	255
Data and the Object Model.....	256
Components of the IMA.....	258
IMA Management Protocol .....	259
IMA Server .....	259
SQL MIB Creation.....	260
IMA Tables Registration .....	260
Register Table Examples.....	264
Removing Table Registrations.....	272
SQL Management Domains.....	272
Control Objects.....	272
IMA Dbmsinfo Constants .....	273
Query Examples to IMA Tables.....	273
Example: Controlling Management Domain.....	274
Example: What Are User Sessions Doing? .....	274
Example: Who Is Waiting for a Lock?.....	275

---

Example: Who Holds Locks? .....	275
Example: Who Is Holding a Lock Others Need? .....	276
Example: Shutting Down a Server .....	277
Example: Monitoring Performance .....	277
Restrictions IMA Use .....	278

## **Appendix A: Configuring the High Availability Option for Sun** **279**

Failover .....	279
Ingres High Availability Option .....	279
Configuration Guidelines for System and Data Files .....	280
Cluster File Systems .....	280
DBMS Server Requirements .....	280
Set Up a Logical Hostname .....	281
Configure Independent Node .....	281
Sun Cluster Software Installation Verification .....	281
Where to Install High Availability Option .....	282
Resource Registration .....	282
Register the Ingres High Availability Option Resource .....	282
Configure Ingres High Availability Option .....	283
Starting and Stopping Ingres in a Cluster Environment .....	285
Start Ingres in a Cluster Environment .....	286
Stop Ingres in a Cluster Environment .....	286
Clear a Fault in a Resource Group .....	286
Change a Resource Property .....	286
Turn on Debugging and Tracing for the High Availability Option .....	287
Remove the High Availability Option .....	288

## **Appendix B: Configuring the High Availability Option for Windows** **289**

Failover .....	289
Ingres High Availability Option .....	289
Configuration Guidelines for System and Data Files .....	290
Cluster Disk .....	290
Physical Disk Resource Type .....	290
Generic Service Resource Type .....	290
DBMS Server Requirements .....	291
Virtual Servers .....	291
Virtual Server Name .....	291
Windows Cluster Service Installation .....	292
Cluster Administrator .....	292
Verify Windows Cluster Service Installation .....	292
Cluster Command .....	292

---

Prerequisites for High Availability Option .....	293
Install High Availability Option .....	293
Functions Performed by the Ingres High Availability Option Setup Utility .....	294
Starting and Stopping Ingres in a Cluster Environment .....	294
Start Ingres in a Cluster Environment .....	295
Stop Ingres in a Cluster Environment .....	296
Remove the High Availability Option .....	297

## **Appendix C: Environment Variables and Logicals 299**

DBNAME_ING .....	299
DBNAME_SQL_INIT .....	299
DD_RServers .....	300
II_4GL_DECIMAL .....	300
II_AFD_TIMEOUT .....	301
II_APPLICATION_LANGUAGE .....	301
II_BIND_SVC_xx .....	301
II_C_COMPILER (VMS) .....	301
II_CHARSETxx .....	302
II_CHECKPOINT .....	302
II_CLIENT .....	302
II_COLLATION .....	302
II_CONFIG .....	302
II_CONNECT_RETRIES .....	302
II_DATABASE .....	303
II_DATE_FORMAT .....	303
II_DATE_CENTURY_BOUNDARY .....	304
II_DBMS_LOG .....	305
II_DBMS_SERVER .....	305
II_DECIMAL .....	305
II_DIRECT_IO (UNIX) .....	305
II_DISABLE_SYSCHECK .....	305
II_DML_DEF .....	306
II_DUMP .....	306
II_EMBED_SET .....	306
II_ERSEND (UNIX) .....	307
II_FORCE_HET .....	308
II_FRS_KEYFIND .....	308
II_GC_REMOTE .....	308
II_GCA_LOG .....	308
II_GCx_TRACE .....	309
II_GCC_LOG .....	309
II_GCD_LOG .....	309

---

II_GCN_LOG .....	310
II_GCNxx_PORT .....	310
II_HALF_DUPLEX .....	310
II_HELP_EDIT.....	310
II_INSTALLATION .....	310
II_JOURNAL .....	310
II_LANGUAGE.....	311
II_LOG_DEVICE (VMS) .....	311
II_MONEY_FORMAT.....	311
II_MONEY_PREC .....	312
II_MSGDIR .....	312
II_NULL_STRING .....	312
II_NUM_OF_PROCESSORS .....	312
II_NUM_SLAVES (UNIX).....	313
II_OO_TABLE_SIZE.....	313
II_PATTERN_MATCH.....	313
II_PF_NODE .....	314
II_POST_4GLGEN.....	314
II_PRINTSCREEN _FILE.....	314
II_SQL_INIT.....	314
II_STAR_LOG .....	314
II_SYSTEM.....	315
II_TEMPORARY .....	315
II_TERMCAP_FILE .....	315
II_TFDIR .....	315
II_THREAD_TYPE .....	315
II_TIMEZONE_NAME.....	316
II_TM_ON_ERROR.....	316
II_TUXEDO_LOC .....	316
II_TUX_SHARED .....	316
II_TUX_AS_MAX .....	316
II_TUX_XN_MAX .....	316
II_UNICODE_CONVERTER .....	317
II_UUID_MAC (Windows) .....	317
II_VNODE_PATH .....	318
II_WORK .....	318
II_XA_TRACE_FILE.....	318
IIDLDIR.....	318
ING_ABFDIR .....	319
ING_ABFOPT1 .....	319
ING_EDIT .....	320
ING_PRINT .....	320

---

ING_SET.....	321
ING_SET_DBNAME .....	321
ING_SHELL (UNIX).....	321
ING_SYSTEM_SET.....	321
INGRES_KEYS .....	321
INIT_INGRES .....	322
TERM (UNIX).....	322
TERM_INGRES.....	322

<b>Index</b>	<b>323</b>
--------------	------------



# Chapter 1: Introducing Ingres System Administration

---

The *System Administrator Guide* provides a comprehensive look at Ingres® from the perspective of the system administrator. This chapter introduces various conventions and notations followed throughout the guide.

## Audience

The *System Administrator Guide* is designed for the system administrator who has overall responsibility for the operation of the Ingres installation. The system administrator must have a solid working knowledge of operating system management or be in close contact with the operating system administrator.

## Special Considerations

Before using this guide, you should be aware of the following issues.

Ingres installations can be administered in compliance with the C2 security standard.

If you are using an Enterprise Access product, see your Enterprise Access documentation for information about syntax that may differ from that described in this guide.

Ingres is compliant with ISO Entry SQL92. In addition, numerous vendor extensions are included. For details about the settings required to operate in compliance with ISO Entry SQL92, see the *SQL Reference Guide*.

## Query Language Used in This Guide

The industry standard query language, SQL, is used as the standard query language throughout this guide.

QUEL commands and system parameters, such as environment variables or logicals, are also included in this guide. For a complete description of available QUEL statements, see the *QUEL Reference Guide*.

## System-specific Text in This Guide

This guide provides information that is specific to your operating system, as in these examples:

**Windows:** This information is specific to the Windows operation system.

**UNIX:** This information is specific to the UNIX operation system.

**VMS:** This information is specific to VMS operating system.

When necessary for clarity, the symbol ■ is used to indicate the end of the system-specific text.

For sections that pertain to one system only, the system is indicated in the section title.

## Path Notation in This Guide

The directory structure of an Ingres installation is the same regardless of operating system. Rather than showing path examples for all environments, this guide uses UNIX notation only.

For example: When describing the location of the collation sequence file, the guide shows: `$II_SYSTEM/ingres/files/collation/collation_name`.

On Windows, the location is:

`%II_SYSTEM%\ingres\files\collation\collation_name`

On VMS, the location is:

`II_SYSTEM:[INGRES.FILES.COLLATION]collation_name`

## Terminology Used in This Guide

This guide uses the following terminology:

- A *command* is an operation that you execute at the operating system level. An extended operation invoked by a command is often referred to as a *utility*.
- A *statement* is an operation that you embed within a program or execute interactively from a terminal monitor.

**Note:** A statement can be written in Ingres 4GL, a host programming language (such as C), or a database query language (SQL or QUEL).



## Syntax Conventions Used in This Guide

This guide uses the following conventions to describe syntax:

<b>Convention</b>	<b>Usage</b>
Regular fixed font	Indicates key words, symbols, or punctuation that you must enter as shown
Italics	Represent a variable name for which you must supply an actual value
[ ] (brackets)	Indicate an optional item
{ } (braces)	Indicate an optional item that you can repeat as many times as appropriate
(vertical bar)	Separates items in a list and indicates that you must choose one item



# Chapter 2: Understanding Ingres Components

---

This chapter discusses the Ingres components, utilities, and tools available to the system administrator to work with Ingres.

## Responsibilities of a System Administrator

The system administrator has all privileges and holds the primary responsibility for installing and maintaining Ingres. The system administrator (often referred to as the installation owner) owns the user ID specified during Ingres installation. Logging in as this user provides permissions that are needed for installation and maintenance.

The system administrator has the following responsibilities:

- Installing Ingres (and optional components)
- Authorizing users to access Ingres
- Defining Ingres variables such as `II_DATABASE` and `II_INSTALLATION`
- Starting, stopping, configuring, and monitoring servers
- Disconnecting or suspending a session connected to a server
- Shutting down the Ingres installation (or its components)

## Components of an Ingres Installation

The system administrator installs and maintains the following principal components or associated family of compatible components:

- DBMS Server
- Internet Communication
- General Communication Facility
- Querying and reporting tools
- Visual tools
- Logging and locking systems

An Ingres installation also has:

- Library files and utilities, provided with the product distribution
- Configuration files and error log files, created during installation and at runtime by Ingres
- Databases and their associated files: the checkpoint, journal, dump, and work files, also created by Ingres during installation and at runtime

The locations of these files are selected at installation time, together with the configuration of the Ingres DBMS Servers and the logging and locking systems. For the initial installation, default settings are provided by the system whenever possible.

After the installation is running, all Ingres servers must be monitored to ensure that they are running and that they are configured for top performance. You can also configure selected Ingres parameters.

## DBMS Server

The Ingres client-server architecture allows multiple users access to databases through connections to one or more DBMS Server processes. The DBMS Server (iidxbms) is a multi-threaded daemon process that performs asynchronous disk input/output. The number of users connected to a server is limited by hardware or the operating system.

Since it is unlikely that the default DBMS server configuration suits all users' requirements, the DBMS Server is usually configured to improve performance in a particular configuration. For a selected DBMS server definition, its buffer caches can be configured as well as other parameters and derived parameters.

DBMS servers can be configured to fit specific needs. For example, you can:

- Designate one server to access a particular database, denying other servers access to that database.
- Configure a server as a "fast commit" server to achieve higher levels of performance.
- Specify that two or more servers share a common memory buffer cache.
- Specify private buffer cache with distributed multi-cache protocol.

Some elements of the DBMS Server are described as facilities.

**Note:** Ingres preallocates most of the resources it requires. It may use hundreds of megabytes of virtual memory and allocate hundreds of thousands of handles; the use of these resources depends on the number and size of the page caches enabled.

## Query Processing

A typical interaction with Ingres consists of at least two processes:

- An Ingres tool or application
- DBMS Server process

### The Ingres Tool or Application Process

The Ingres tool or application performs the following actions:

- Takes user input and issues a query that is sent to the DBMS Server.
- Formats, optimizes, and executes the query on behalf of the user.
- Displays the resulted data returned by the DBMS Server.

## DBMS Server Process

The Ingres DBMS Server is a multi-threaded process. It can execute queries for a large number of users, each running an Ingres tool. Even though it is a single process, the DBMS Server can execute queries as multiple “sessions” on behalf of multiple users. The Visual Performance Monitor tool lets you view which sessions are running in the DBMS Server at any moment.

## Query Environment

When a thread or session executes a query inside the DBMS Server, it does so in a *query environment*. The query environment consists of:

- A quantity of resources available from the operating system for use by the session.
- The rules under which the query is executed. These rules reflect which query language is used, which locking strategy is employed, which diagnostic information is returned, which default behavior Ingres adopts for various query language statements, and so on.

## DBMS Server Facilities

The following are DBMS Server facilities, which are configured using DBMS Server configuration parameters:

- Abstract Data Type Facility
- Compatibility Library Facility
- Data Manipulation Facility
- Optimizer Facility
- Parser Facility
- Query Execution Facility
- Query Storage Facility
- System Control Facility
- Relation Description Facility
- General Communication Facility

## Abstract Data Type Facility

Abstract Data Type Facility (ADF) does all the work that involves data types. It manipulates floating-point numbers, character strings, integers, and all the conversions and comparisons between them. This facility can be executed independently from the server so that other Ingres tools can also use ADF.

This abstract data type facility is also compiled into the Ingres tools such as QBF (Query-By-Forms) and the Terminal Monitor. It is used to manipulate Ingres data types in the front-end process before sending them to the DBMS Server.

## Compatibility Library Facility

Compatibility Library Facility (CLF) provides the Ingres DBMS Server with a platform independent interface to the operating system. It contains functions that perform services such as disk I/O, memory management, Semaphore manipulation, and context switching (when Ingres threading is configured). It also provides low level system communications functions.

For example, when a thread in the DBMS is idle, it normally shows up in iimonitor, IPM, or VDBA as having current facility CLF. This indicates that the thread is executing a CLF function to read information from the communications channel that connects the thread to the system process that instigated the database connection associated with this thread.

## Data Manipulation Facility

The Data Manipulation Facility (DMF) manages the DBMS Server interface to disk storage. In addition to managing all storage structures (hash, heap, ISAM, B-tree, BLOBs, and so on), DMF uses the logging and locking systems to control transaction processing and to handle concurrency issues. Included within DMF is a buffer manager that controls access to a cache (possibly shared) of database pages.

## Optimizer Facility

The Optimizer Facility (OPF) selects the optimal plan for implementing queries. It also converts the query tree that comes out of the parser into the query execution plan (QEP). OPF uses a memory pool for its operations.

## Parser Facility

The Parser Facility (PSF) converts queries from text form to internal format. The parser adds data from the system catalogs to the query, such as information about the table structure and keys that the optimizer needs to make a useful query plan.

## Query Execution Facility

The Query Execution Facility (QEF) executes query plans and database utilities. It provides internal query services that other facilities use. QEF manages repeat queries, transactions, and cursors.

## Query Storage Facility

The Query Storage Facility (QSF) provides shared memory facilities with a temporary or permanent place to store query trees and query plans.

## System Control Facility

The System Control Facility (SCF) is the central controlling facility that manages sessions on behalf of user (client) requests. It coordinates the actions among the various facilities involved in processing a query, including thread monitoring and switching. SCF is also responsible for managing server-wide access to shared resources, such as operating system semaphores and memory.

## Relation Description Facility

The Relation Description Facility (RDF) is a central caching point for information about tables. RDF is used by Ingres Distributed Option, PSF, and OPF.



## General Communication Facility

The General Communication Facility (GCF) manages communication among all components of Ingres. GCF has five elements:

- Name Server
- Communications Server
- Data Access Server
- Bridge Server
- General Communications Architecture

### Name Server

The Name Server keeps track of all DBMS, Data Access, JDBC, Communications, Bridge, Star, Internet Communications, and Enterprise Access servers in an installation and provides user processes with the information they require to connect to local or remote installations.

#### How the Name Server Works

The Name Server provides information to user processes that enables a connection to a local DBMS Server. When a process wants to connect to a remote DBMS Server:

- The Name Server provides information that allows the process to first connect to a Communications Server.
- The Communications Server establishes communication with the remote DBMS Server.
- The Name Server checks regularly (the default is every five minutes) to ensure that all DBMS Servers on its list are functioning. If a server has shut down, the Name Server automatically deletes its registration.

#### Name Server Maintenance Utility

The Name Server Maintenance Utility (iinamu) is a part of the General Communication Facility provided by Ingres. This utility allows the system administrator to manually add servers to or delete servers from the list maintained by the Name Server, or to stop and restart the Name Server.

For more information about this utility, see the *Command Reference Guide*.

## Communications Server

The Communications Server (iigcc) is a daemon process that provides the network communication element of the Ingres Net product. It monitors outgoing communication from local applications to remote DBMS Servers and incoming communication from remote applications to local DBMS Servers. An installation can have multiple Communications Server processes.

For more information on the Communications Server, see the *Connectivity Guide*.

## Data Access Server

The Data Access Server enables communication between Java applets and applications and Ingres. The Data Access Server is the server component of the JDBC Driver, and also provides DBMS Server access for the Ingres .NET Data Provider. Data Access Server parameters include the maximum number of inbound sessions, client connections, connection pooling, pooling timeouts, network port, listen address, and protocol status. Configuration of the Data Access Server is done at installation and rarely requires reconfiguration.

## Bridge Server

The Ingres Bridge extension enables communication between Ingres clients on one network and servers on a different type of network. Reconfiguration of the Bridge Server is rarely required.

For information on the Bridge Server's role in Ingres connectivity, see the *Connectivity Guide*.

## General Communications Architecture

The General Communications Architecture (GCA) is the lowest level General Communications Facility Application Program Interface. The GCA maintains communication connections between processes on the same local Ingres installation. The GCA is a subroutine library that is a part of all Ingres tools, DBMS Servers, Star Servers, and the libraries associated with embedded SQL and EQUJEL.

## Internet Communication

The Ingres Web Deployment Option provides the foundation for internet-based electronic commerce. Using this component, users can quickly and easily build Web pages with corporate data and deploy database Web applications.

For more information, see the *Web Deployment Option User Guide*.

## Querying and Reporting Tools

Ingres provides several querying and reporting tools such as Report-Writer, Report-By-Forms (RBF), Query-By-Forms (QBF), and Vision to enable users to access Ingres databases.

When a user accesses a database through tools like Report-Writer, Report-By-Forms (RBF), Query-By-Forms (QBF), or Vision, a process is created and communication is initiated with the database through a DBMS Server. The name server, which is a part of the General Communication Facility, provides user information.

For more information on Ingres querying and reporting tools, see the *Character-based Querying and Reporting Tools User Guide*.

## Visual Tools

Ingres lets you start, stop, manage, and monitor the Ingres installation using various GUI tools.

### Visual Manager

Ingres Visual Manager (IVM) provides a global view into the Ingres installation. It serves as a system console from which you can manage Ingres components and access other utilities. This utility captures events that are occurring in the system and allows them to be filtered for emphasis, based on your preferences.

You can access all other Ingres visual tools through IVM. For an overview of IVM tasks, see the chapter "Managing your System and Monitoring Performance." For instructions on using IVM screens, see the IVM online help. For information on IVM command line tasks, see the *Command Reference Guide*.

## Visual Performance Monitor

Visual Performance Monitor can be used as a monitoring tool, a performance analysis tool, and a system management tool.

Visual Performance Monitor lets you monitor the following:

- Servers

You can view a list of the servers that are started on an Ingres installation, and information about each server. For example, you can view the sessions that are currently active for each server. You can also remove a session or stop a server (if you are a privileged user).

Visual Performance Monitor provides the following information for a particular server:

- Lock lists
- Locks
- Transactions
- Locked databases, tables, and pages
- Other locked resources
- Network traffic for Net servers

- Users and sessions

You can monitor users for whom there are open sessions. You can find out which sessions are open for a particular user, and drill down further to reveal all the related information about those sessions (lock lists, transactions locked, databases locked, and so on).

- Logging

You can view logging system summaries, transaction lists, process, and database lists. Log information can be used to monitor transaction rates, log file activity, processes, and databases in the logging system. This information is useful in determining which logging parameters need to be adjusted.

- Locking

You can monitor lock information to help determine which lock parameters need to be adjusted. Viewing locking system summaries, lock lists, and resources provides you with the information you need to spot conditions when, for example, additional locking system resources need to be added.

Viewing your lock lists is useful for locating transactions that cannot proceed because they are blocked by another transaction.

- System performance

You can view your performance information from a “database” point of view. This means you can access performance information using a database branch, as opposed to the root branches of Visual Performance Monitor.

- Replication

You can start, stop, and monitor Replicator servers that are required for the replication scheme that has been defined in a VDBA DOM window. You can set up startup parameters for these servers, send events to these servers, view and manage collisions, and display other miscellaneous replication monitor information.

Visual Performance Monitor lets you perform the following actions:

- Refresh data in the window from a server
- Shut down a database server or close a session

## Visual DBA

Visual DBA provides an alternative set of system administration tools. It allows you to perform system administrator functions including configuring, performance monitoring, backup and recovery, and remote database optimization. For an overview of Visual DBA tasks related to communications, see the *Connectivity Guide*. For instructions on using Visual DBA screens, see the Visual DBA online help. For information on Visual DBA command line tasks, see the *Command Reference Guide*.

## Remote Command Server

Some options and dialogs in Visual DBA result in the launch of an operating system level command (such as `createdb` or `destroydb`) on the server. These "remote commands" are performed by the Remote Command Server (RMCMD) process on behalf of Visual DBA.

The Remote Command Server (RMCMD) allows remote execution of operating system commands. It must be started in installations where a DBMS server is running for certain DBA tasks to be accessible remotely with Visual DBA. The Remote Command Server is started on the server side, not on the Visual DBA client.

DBA tasks are primarily those that do not have an equivalent through an SQL statement, for example, creating or dropping a database, displaying selected portions of the journal for a database, or starting a replication server remotely. By default, only the installation owner is authorized to perform such tasks remotely through Visual DBA.

The `rmcmd` process is started and stopped with Ingres, according to the settings for the Remote Command component in CBF or Configuration Manager.

## Grant Access to Remote Users

To allow a user other than the installation owner to execute remote commands, you can use either of the following methods:

- You use the Create or Alter User dialog in VDBA. Enable the Remote Command (`rmcmd`) Privileges checkbox.
- You can execute the following SQL statements, while connected to the `imadb` database as the installation owner:

```
grant select,insert,update,delete on $ingres.remotecmdinview to user
grant select,insert,update,delete on $ingres.remotecmdoutview to user
grant select,insert,update,delete on $ingres.remotecmdview to user
grant execute on procedure $ingres.launchremotecmd to user
grant execute on procedure $ingres.sendrmcmdinput to user
grant register, raise on dbevent $ingres.rmcmdcmdend to user
grant register, raise on dbevent $ingres.rmcmdnewcmd to user
grant register, raise on dbevent $ingres.rmcmdnewinputline to user
grant register, raise on dbevent $ingres.rmcmdnewoutputline to user
grant register, raise on dbevent $ingres.rmcmdstp to user
```

**Note:** Grants must not be made directly to the underlying tables.

## How Remote Commands Are Executed

Operating system commands executed by the Remote Command (rmcmd) Server are executed on the server side under the user ID that launched that server (typically the installation owner).

For commands that support the `-u` option, the rmcmd server will accept the command execution request only if it contains a `-username` argument, where *username* is the user ID of the session through which the client application (typically VDBA) does the request.

For commands that do not support the `-u` option, such as `sysmod`, `alterdb`, or `relocatedb`, the request will be accepted only if the user ID of the session is the same as the user who launched the rmcmd server on the server side.

## Configuration Manager

The Configuration Manager utility provides a GUI interface for configuring the Ingres installation. For more information, see Configuration Manager online help. For information on Configuration Manager command line tasks, see “vcbf” in the *Command Reference Guide*.

## Network Utility

The Network Utility is a stand-alone tool that allows you to view and define virtual node (vnode) definitions, which are used to connect to remote Ingres installations through Ingres Net. The Network Utility provides all the same functionality as the virtual nodes toolbar in Visual DBA. The Network Utility also enables launching of the Database Object Manager, Visual Performance Monitor, and Visual SQL utilities for such installations.

For more information, see the *Connectivity Guide* and the Network Utility online help.

## Export Assistant

The Export Assistant allows you to export Ingres or Enterprise Access table data into external file formats such as CSV, XML or DBF. It can be run as a stand-alone utility or from within Visual DBA. For instructions on using the Export Assistant, see the Export Assistant online help.

## Import Assistant

The Import Assistant allows you to import an external text file into an Ingres or Enterprise Access table. It can be run as a stand-alone utility or from within Visual DBA. For instructions on using the Import Assistant, see the Import Assistant online help.

## Journal Analyzer

The Journal Analyzer utility allows you to view and analyze journaled transactions and individual underlying statements. It also enables you to create SQL scripts, which are used to recover or redo individual row operations without redoing or rolling back the whole database or table.

For instructions on using the Journal Analyzer, see the chapter “Analyzing and Recovering Journaled Transactions” and the Journal Analyzer online help.

## Visual SQL

The Visual SQL utility allows you to enter and execute SQL queries. As part of Visual SQL, the SQL Assistant helps you build SQL queries. Visual SQL can be run as a stand-alone utility or from within Visual DBA. For instructions on using this utility, see the Visual SQL online help.

## Visual Configuration Differences Analyzer

The Visual Configuration Differences Analyzer (VCDA) allows you to compare configuration snapshots of an Ingres installation.

To learn about the benefits of this tool, see chapter “Managing Your System and Monitoring Performance”. For instructions on using this tool, see the Visual Configuration Differences Analyzer online help.

## Visual Database Objects Differences Analyzer

The Visual Database Objects Differences Analyzer (VCDA) allows you to compare the definition of Ingres database objects that are either currently installed or saved in a snapshot file. For instructions on using this tool, see the VCDA online help.



## Logging and Locking Systems

The logging and locking systems coordinate the locking, recovery, and journaling of databases. The system is composed of the following components:

- Lock manager
- Logging facility
- Recovery process
- Archiver process
- Log files

### Lock Manager

The locking component controls concurrent access to a database. Shared memory is allocated to components initially during the installation procedure. The amount of shared memory that your installation requires depends on the logging and DBMS server parameters that you select during the procedure.

**UNIX:** In UNIX, the lock manager component makes use of the shared memory segments and semaphore resources that you install when you configure the UNIX kernel.

### Logging Facility

The logging facility implements a circular “write-ahead” transaction log file for the management of transactions within the installation. It ensures that log records are written in a way that makes them accessible to the recovery and archiver processes.

## Organization of the Logging Facility

The logging facility maintains in-memory log buffers. For every in-memory log buffer, there is a corresponding disk log file. Memory log buffers and disk log file blocks are the same length. The number of log buffers is configurable, set according to the performance requirements of the system. The logging system manages multiple asynchronous writes of log buffers to the log file.

The number of disk log file blocks corresponds to the size of the log file, and is specified when the log file is created. In a properly tuned system, most log file buffers are completely full of log records when they are written to the log file. However, as all log records associated with a transaction must be forced to the log file at certain times, principally at end transaction time, a small percentage of log file blocks may contain unused space. Server group commit logic is designed to minimize the frequency of log force operations and increase log file space use.

## How Logging Facility Works

All servers in the Ingres installation, as well as the recovery and archiver processes, share the logging facility. Log records written by different servers, or written by several threads in the same server, can be combined with log records written by other servers or threads. The Logging Facility performs the following actions:

- Copies the Log records to in-memory log buffers, which are written to the log file as they fill.
- The log file block is rewritten when the file wraps around.

## Recovery Process

The recovery process performs recovery in the event of a server or system failure. The recovery process maintains a history of important actions in its own message log file, `iircp.log`.

## Data Manipulation Facility Recovery Process (dmfrcp)

Each installation of Ingres has a `dmfrcp` (data manipulation facility recovery process) that is assigned recovery responsibilities. In normal circumstances, transaction commit and rollback processing are handled within each DBMS server. In the event of a server or system failure, however, the recovery process performs the required recovery, including backing out uncommitted transactions and ensuring that committed transactions are properly reflected on disk.

## Recovery Modes

The recovery process operates in the following modes:

- Online recovery
- Offline recovery

### Online Recovery

Online recovery is performed when a server stops abnormally. In this case, users connected to other servers are generally unaffected by the recovery.

### Offline Recovery

Offline recovery is performed when the installation is brought back up after it has stopped abnormally. In this case, the installation remains unavailable until the recovery process completes all required recovery.

## Archiver Process

The archiver process (dmfacp) is responsible for copying the history of operations performed on journaled databases from the transaction log file to the journal files. The archiver process maintains a history of important actions in its own message log file, liacp.log.

The archiver process removes completed transactions from the transaction log file. For journaled tables, the archiver writes any completed transaction to the journal files for the database. Each database has its own journal files, which contain a record of all the changes made to the database since the last checkpoint was taken. The archiver process “sleeps” until sufficient portions of the transaction log file are ready to be archived or until the last user exits from a database.

### Data Manipulation Facility Archiver Process (dmfacp)

Each installation has a single archiver process called the dmfacp (data manipulation facility archiver process). The archiver process is responsible for copying the history of operations performed on journaled databases from the transaction log file to the journal files. Journal files contain the subset of transaction log file information associated with a specific database.

The name of the DMFACP process depends on the installation code. For Ingres installations with installation code II, the process name is DMFACP. For installation codes other than II, the DMFACP process name is DMFACPxx where xx is the installation code.

## Journal Files

Journal files contain the subset of transaction log file information associated with a specific database. Journal files are created (and optionally destroyed) during a checkpoint operation. In the event of a disaster, the database can be rebuilt by restoring the latest checkpoint and applying journal file information.

## Disaster Recovery Command

Disaster recovery operations are coordinated by the `rollforwarddb` command. It can be accessed through the Visual DBA Database menu, DOM window.

## Transactions Recovery

Transaction recovery involves the transaction log file that is used as a write-ahead log, plus journal files maintained on a per-database basis. Log files contain short-term recovery information regarding active databases, while the journal files contain long-term information used for auditing and disaster recovery. While the log file is circular and wraps around, journal files are of configurable length and are retained indefinitely.

Ingres employs a page-oriented recovery scheme, where changes to pages are reflected in the transaction log file.

## Types of Transaction Recovery

Recovery information is divided into two types:

- Undo (or Rollback) Operations
- Redo (or Cache Restore) Operations

Ingres performs both online and offline recovery, as described in Recovery Modes (see page 35).

## Undo Operation

Undo or transaction backout recovery is performed by the DBMS Server. For example, when a transaction is aborted, transaction log file information is used to roll back all related updates. The DBMS Server writes the Compensation Log Records (CLRs) to record a history of the actions taken during undo operations.

## Redo Operation

A Redo recovery operation is database-oriented. Redo recovery is performed after a server or an installation fails. Its main purpose is to recover the contents of the DMF cached data pages that are lost when a fast-commit server fails. Redo recovery is performed by the recovery process. Redo recovery precedes undo recovery.

## Redo Operation in a Cluster Environment

In a Linux cluster environment, or any other Ingres cluster environment where all nodes are active, the local recovery server performs transaction redo/undo for a failed DBMS server on its node, just like in the non-cluster case. The difference in a cluster installation is that if the recovery process (RCP) dies on one node, either because of an Ingres failure, or a general failure of the hardware, an RCP on another node will take responsibility for cleaning up transactions for the failed nodes.

## Ingres Log Files

Ingres maintains several log files to which it writes information about the installation activities. The various log files are:

- Transaction Log File
- Error Log
- Archiver Log
- Recovery Log
- VMS Cluster Service Process Log
- Primary Configuration Log Files
- Operational Configuration Log Files
- Optional Log Files

## Transaction Log File

Each installation has a transaction log file, and an optional dual log file. The transaction log file holds information about all open transactions and is used to recover active databases after a system failure. You have the option to change its size and number of partitions at startup.

**UNIX:** The UNIX log file can be created as a raw partition, or as a number of raw partitions if you choose. For details, see the *Getting Started* guide.

## Error Log

The Error Log (errlog.log) file is the main log file. It is a readable text file that you can use for troubleshooting. Messages about the installation are appended to this log with the date and time at which the error occurred. This is generally the first place to look when troubleshooting a problem. The error log contains the following information:

- Error messages
- Warning messages
- Server start up and shutdown messages

The system administrator maintains the error log file. This file continues to grow until manually truncated. The installation must be shut down before truncating or removing the errlog.log file.

The error log file is located at the \$II\_SYSTEM/ingres/files/ directory.

## Archiver Log

The archiver log (iiacp.log) contains information about the current archiver process. This file is appended to when the archiver process starts. The log contains the following information:

- Archiver start up
- Archiver error messages
- Archiver warning messages

The archiver log is located at \$II\_SYSTEM/ingres/files/ directory.

## Recovery Log

The recovery log (iircp.log) contains information about the current recovery process. This file is appended to when the recovery process starts. The log contains the following information:

- Current logging and locking parameter values
- Recovery process error messages
- Recovery process warning messages
- Recovery operations information

**Note:** The recovery log must be monitored if you are unable to connect to Ingres and suspect that the DBMS Server is in recovery mode.

The recovery log is located at the \$II\_SYSTEM/ingres/files/ directory.

## Primary Configuration Log Files

Ingres maintains transcripts of various configuration operations. The primary configuration log files that Ingres uses are as follows:

### **config.log**

**The config.log file contains a log of the changes made with the Configuration Manager or Configuration-By-Forms (cbf) utility.**

### **rcpconfig.log**

The rcpconfig.log file contains log and error information of the last time rcpconfig was run. For details of the rcpconfig command, see the *Command Reference Guide*.

The primary configuration log files are located in the \$II\_SYSTEM/ingres/files/ directory.

## Optional Configuration Log Files

Ingres provides optional configuration log files. Configuration log files contain a log of messages that are generated every time configuration specific commands are executed. Ingres provides the following configuration log files:

### **iilink.log**

The iilink.log file contains a log of the last time the iilink command was run. For details of the iilink command, see the *Command Reference Guide*.

### **iivdb.log**

The iivdb.log file contains a transcript of the last time the verifydb command was run. This file is created the first time when the verifydb command is run. For more information on the verifydb command see *Command Reference Guide*.

### **lartool.log**

The lartool.log file contains a transcript of the last time the lartool command was run. This file is created the first time lartool is run. For more information on the lartool command see *Command Reference Guide*.

These files are present in the II\_CONFIG directory of the Ingres installation.

**Note:** In UNIX, the II\_CONFIG file must always be \$II\_SYSTEM/ingres/files.

## Optional Log Files

Ingres provides various optional log files. These log files can be used for troubleshooting purposes. Ingres provides optional log files for:

- Logging messages for specific processes
- Generating logs for a specific Ingres facility
- Tracing logs and tracking messages at a greater level of detail

## Process Logs

Process logs can be set up for the following processes to isolate the error messages relating to that process:

- DBMS
- GCC



## Optional Log Facility

Logging can be specified for a specific Ingres facility. The following optional log and trace log files can be established by setting the associated Ingres variables.

### **DBMS Error Log File**

The DBMS Server error log is optionally defined as a separate file. This log file is established by setting the Ingres variable `II_DBMS_LOG` to a user-defined file name.

**Note:** By default, all logged messages are sent to the `errlog.log` file. If you define separate error logs, all messages are sent to both that error log file and `errlog.log`.

### **GCC Error Log File**

The GCC error log is used for installations using Ingres Net. This is in a user-defined file. This log file is established by setting the environment variable `II_GCC_LOG` to the full name of a file. The error logging level is specified by `II_GCC_LOG_LVL`.

### **GCC Trace Log File**

The GCC trace log is set up for specific troubleshooting efforts. You set `II_GCA_LOG` to a user file name. The associated Ingres environment variable `II_GCC_TRACE` defines the level of tracing.

### **Star Error Log**

The Star error log is optionally defined as a separate file. This log file is established by setting the Ingres variable `II_STAR_LOG` to a user file name.

**Note:** All Ingres Distributed Option errors and messages are sent to `errlog.log` by default.

For more information about these log files, see the chapter "Setting Environment Variables and Logicals" and the appendix "Environment Variables and Logicals."



# Chapter 3: Configuring Ingres

---

This chapter shows how to get started using the Ingres configuration tools and describes configuration considerations for each Ingres component.

## Configurable Components

You can view and set values for the following components:

- Server components
  - DBMS
  - JDBC
  - Data Access
  - Communications
  - Bridge
  - Star
  - Name
  - Recovery
- Internet and system components
  - Internet communication
  - Security
  - Locking system
  - Logging System
  - Primary and dual transaction logs

## Configuration Tools

Ingres provides the following tools to configure an Ingres installation.

- Configuration Manager (VCBF)
- Configuration-By-Forms (CBF)
- Visual Configuration Difference Analyzer (VCDA)

## Access Configuration Manager Utility

To configure Ingres using the Configuration Manager Utility, follow these steps:

**Windows:** Do **one** of the following:

- Click Start on the Windows taskbar and then choose Programs, Ingres, Configuration Manager.
- In Ingres Visual Manager (IVM), click the Configuration toolbar button, or choose File, Configure.
- Right-click the database button on the system tray and choose Configuration Manager from the popup menu.
- Enter **vcbf** on the command line.

**UNIX:** Enter **vcbf** at the command line.

## Access Configuration-By-Forms Utility

To access the Configuration-By-Forms utility, enter **cbf** at the command line.

## Access VCDA Utility

To access VCDA utility:

**Windows:** Do **one** of the following:

- Click the VCDA toolbar button in Ingres Visual Manager (IVM).
- Right-click the database button on the system tray and choose VCDA from the popup menu.
- Enter **vcda** on the command line.

**UNIX:** Enter **vcda** at the command line.

**Note:** On UNIX systems that support visual tools, you can access Configuration Manager and VCDA through IVM, as described for the Windows platform.

## Component Configurations

The Configuration Manager allows you to set up the components installed on the system using the Configure tab, which is displayed in the left pane in Configuration Manager.

Different configurations are needed for servers that perform different functions. For example, batch jobs can require a completely different DBMS server configuration than that of an online server. If so, create a DBMS server configuration called "batch" and configure it for batch processing.

The Configuration Manager lets you add, delete, or rename the various configurable server components listed in Configurable Components (see page 43). For detailed steps on performing these procedures, see the following topics in the Procedures section of the Configuration Manager online help.

- Adding a Server Component Configuration
- Deleting a Server Component Configuration
- Changing the Name of a Server Component Configuration

**Note:** You can perform these procedures for non-default components configurations only.

**Note:** You cannot delete configurations for the Name Server and Recovery Server.

## Configurable Properties for a Component

Each non-default component configuration has two properties that can be changed. They are:

- Startup Count
- Name

### Startup Count

The startup count for a component represents the number of instances to be initiated when the Ingres installation is started.

For example, if you set the startup count to five for the Communications Server, five Communications Servers start at startup time. All components that have a startup count that can be edited are daemon processes (the DBMS Server, Star, GCC, and so on). You can change the startup count for any component using Configuration Manager or the Configuration-By-Forms utility. For instructions on changing startup count, see "Changing the Startup Count" in the Configuration Manager online help.

## Name

Every non-default component can be assigned a name. The name allows you to assign a specific name to a component configuration.

For instructions on changing component names, see Setting Server Properties in Configuration Manager online help.

**Note:** You can only rename non-default server components.

## Configuration Snapshot

A configuration snapshot file contains information on the installation's parameters including:

- Configuration parameters
- vnode definitions
- System variables
- User variables

The snapshot file can be used to compare configurations changes.

For example, if you encounter problems with the installation later, take another snapshot of the configuration and compare it to the earlier snapshot to determine if any configuration changes have contributed to the problem. Also, keep an on-going record of configuration changes by taking a snapshot of the installation each time you change its configuration.

**Note:** Before configuring or editing the system parameters, you must use the VCDA tool to create a "snapshot" file of the current configuration.

## Configuration Parameters

Ingres configuration parameters are stored in the config.dat file. Some Ingres parameters are dependent on, or derived from other Ingres parameters.

## Setting Parameter and Derived Parameter Values

The Parameters page in Configuration Manager lists the parameters associated with a selected component configuration.

For most component configurations, you can edit the value of a parameter, restore the parameter to a value that has been saved in a configuration "snapshot" file, or restore the parameter to its original installation value.

When you select a particular component and click on one of the parameters in the Parameters tab, the Edit Value and Restore buttons become enabled. For instructions on editing and restoring parameter values, see Setting and Restoring Component Parameters in the Procedures section of the Configuration Manager online help.

## Using VCDA to Restore Configuration Parameter

VCDA can be used to restore configuration parameter values saved in a snapshot file. For detailed steps on using this tool, see the VCDA online help.

## Viewing and Recalculating Derived Parameters

By clicking the Derived tab, you can view a list of, as well as recalculate any of the derived parameters. Any derived parameter can be set in "protected" mode. This means that Configuration Manager or the Configuration-By-Forms utility cannot recalculate the derived value. This allows you to set a desired value for the parameter, even though the value is normally derived by the system from other parameters. If protected mode is disabled, the value of the parameter is adjusted whenever new values are set for other parameters on which the derived value is based. The new and old values are shown in the change log.

For detailed steps on performing these procedures, see the following topics in Configuration Manager online help:

- Recalculating Derived Parameter Values
- Protecting Derived Parameters

## Components That Use Derived Parameters

The following components contain parameters that are derived from other parameters:

- DBMS Server
- Star Server
- Security
- Locking System
- Logging System
- Recovery Server

When a parameter is changed, any parameters derived from that parameter are automatically updated. Derived parameters appear on the Derived tab.



## DBMS Server Configuration

The Ingres client-server architecture allows multiple users access to databases through connections to one or more DBMS Server processes. The DBMS Server (iidxbms) is a multi-threaded daemon process that performs asynchronous disk input/output. The number of users connected to a server is limited by hardware or the operating system.

Since it is unlikely that the default DBMS server configuration suits all users' requirements, the DBMS Server is usually configured to improve performance in a particular configuration. For a selected DBMS server definition, its buffer caches can be configured as well as other parameters and derived parameters.

DBMS servers can be configured to fit specific needs. For example, you can:

- Designate one server to access a particular database, denying other servers access to that database.
- Configure a server as a "fast commit" server to achieve higher levels of performance.
- Specify that two or more servers share a common memory buffer cache.
- Specify private buffer cache with distributed multi-cache protocol.

Some elements of the DBMS Server are described as facilities.

**Note:** Ingres preallocates most of the resources it requires. It may use hundreds of megabytes of virtual memory and allocate hundreds of thousands of handles; the use of these resources depends on the number and size of the page caches enabled.

For instructions on configuring DBMS server parameters, see the following topics in Configuration Manager online help:

- Cache Parameters Page, DBMS Server Component
- Parameters Page, DBMS Server Component
- Derived Page, DBMS Server Component

## Fast Commit Option

The `fast_commit` option, which is enabled using the Derived Page of the DBMS Server Component in Configuration Manager, allows the server to execute commits without forcing data pages to the database. This reduces I/O in the server and improves response time.

Without fast commit, all update transactions must do the following when a commit is executed:

- Force all "log" records for this transaction to disk
- Force all updated data pages to disk
- Write a "commit" record to the "log" file and force it to disk

This ensures that the data is committed and the data in the database has been written to disk when a system crash occurs immediately following a commit.

With fast commit, a server can perform a commit by writing a commit record to the log file and forcing it to disk. The actual data pages are not written to the database disk until an optimal time to do so is reached. If necessary, the committed transactions can be recovered from the log file that was written to disk.

Fast commit can greatly improve transaction response time and throughput and greatly reduce direct I/O to the database. All update transactions gain some benefit from using fast commit because control is returned to the application while database writes are processed asynchronously. Applications that tend to update the same group of records many times gain a big increase in throughput when using fast commit because writes to the database are delayed until many transactions have actually been committed. This allows the system to write out many updates with few I/O operations.

However, using fast commit can also increase recovery times. If the system crashes, all committed transactions that have not yet had their data written to the disk must be recovered. To control this recovery time, adjust the consistency point frequency.

## Write Behind Option

The `dmf_write_behind` option, which is enabled using the Cache Parameters Page, DBMS Server Component in Configuration Manager, allocates server threads that asynchronously write modified data pages to disk. The write behind threads are controlled by the `dmf_wb_start` and `dmf_wb_end` thresholds that were defined at server startup time. (Both these thresholds can be changed using the Derived Cache Parameters Page, DBMS Server Component in Configuration Manager.)

By allowing transactions to commit without waiting for data to be flushed to disk, fast commit can significantly improve response time. However, the data page caches of fast commit servers tend to become filled with modified pages. The write behind threads control the number of modified pages by asynchronously writing them to disk.

If write behind threads are not used, a server quickly reaches its modified page limit (specified by `dmf_modify_limit`). When that happens, transactions begin to do synchronous writes to free up space in the cache, negating some of the benefits of using fast commit. For this reason, if you configure a server with fast commit, you must also specify write behind to get the full performance benefits of fast commit.

For two sessions to have concurrent access to a database using fast commit, each session must access the database through the same DBMS Server or through servers connected to a common buffer. To be connected to a common buffer, the server must use the `cache_sharing` parameter. When a server configured with the `fast_commit` option opens a database, it locks the database so that no other server (except a server using the same buffer cache) can access the database.

**Note:** You can use the `dmf_write_behind` option even if you are not using the `fast_commit` option; write behind threads can write to the database disk even if fast commit is not specified.

The number of write behind threads used to configure the DBMS Server depends on the system's I/O requirements. The number of writes per second each write behind thread can process depends on the speed of the disk drives. You must configure enough write behind threads to keep up with the expected I/O rate to the database.

Generally, a good rule is to allocate a write behind thread for each disk on which the databases are stored. Adding more write behind threads than disk devices does not increase the system's I/O throughput. A default installation allocates two write behind threads, which is sufficient for most application mixes.

## Data Access Server Configuration

The Data Access Server enables communication between Java applets and applications and Ingres. The Data Access Server is the server component of the JDBC Driver, and also provides DBMS Server access for the Ingres .NET Data Provider. Data Access Server parameters include the maximum number of inbound sessions, client connections, connection pooling, pooling timeouts, network port, listen address, and protocol status. Configuration of the Data Access Server is done at installation and rarely requires reconfiguration.

Data Access Server parameters include the maximum number of inbound sessions, client connections, connection pooling, pooling timeouts, network port, listen address, and protocol status. Configuration of the Data Access Server is done at installation and rarely requires reconfiguration.

For additional information on the Data Access Server's role in Ingres connectivity, see the *Connectivity Guide*. For instructions on configuring Data Access Server parameters, see the following topic in Configuration Manager online help:

- Parameters Page, Data Access Server Component

## Communications Server Configuration

The Communications Server (iigcc) is a daemon process that provides the network communication element of the Ingres Net product. It monitors outgoing communication from local applications to remote DBMS Servers and incoming communication from remote applications to local DBMS Servers. An installation can have multiple Communications Server processes.

For more information on the Communications Server, see the *Connectivity Guide*.

Communications Server parameters include the encryption mode, logging level, registry type, number of inbound and outbound sessions, and the logging and error levels. Configuration of the Communications Server is done at installation and rarely requires reconfiguration. However, changes to a protocol or the listen address are needed. For information on the Communication Server's role in Ingres connectivity, see the *Connectivity Guide*.

For instructions on configuring Communications Server parameters, see the following topics in Configuration Manager online help:

- Parameters Page, Net Server Component
- Protocols Page, Net Server Component

## Bridge Server Configuration

The Ingres Bridge extension enables communication between Ingres clients on one network and servers on a different type of network. Reconfiguration of the Bridge Server is rarely required.

For information on the Bridge Server's role in Ingres connectivity, see the *Connectivity Guide*.

For instructions on configuring Bridge Server parameters, see the following topics in Configuration Manager online help:

- Parameters Page, Bridge Server Component
- Protocols Page, Bridge Server Component

The Bridge Server parameters include the number of inbound sessions, the enabling of specific protocols, and a listen address.

## Star Server Configuration

The Star Server supports Ingres Distributed Option. Ingres Distributed Option is a distributed extension to the DBMS Server, which allows data from multiple independent databases to be accessed as if it were contained in a single, centralized relational database. In this way, Ingres Distributed Option provides client applications with transparent access to enterprise-wide data.

Star Server parameters include the number of connections supported by Star Servers, accounting data maintained, and the maximum number of tables that can be kept in the cache at one time.

Distributed recovery can be disabled, meaning that the Star Servers created with this definition do not recover distributed transactions at startup time. This feature is helpful if there is a problem with the installation, but you want it to continue working while the problem is being investigated. Like the DBMS Server, configuration is usually required to provide optimal performance in a given scenario.

For instructions on configuring Star Server parameters, see the Parameters Page, Star Server Component in the Configuration Manager online help.

## Name Server Configuration

The Name Server keeps track of all DBMS, Data Access, JDBC, Communications, Bridge, Star, Internet Communications, and Enterprise Access servers in an installation and provides user processes with the information they require to connect to local or remote installations.

The Name Server parameters include the time interval for checking for active servers, and session limits. The Name Server rarely needs to be reconfigured.

For instructions on configuring Name Server parameters, see the following topic in Configuration Manager online help:

- Parameters Page, Name Server Component

## Recovery Server Configuration

In the event of a server or system failure, the Recovery Server performs any required recovery, including backing out uncommitted transactions and ensuring that committed transactions are properly reflected on disk.

The configurable Recovery Server system parameters include connection limits, consistency points, events, gather write algorithms, server name, thread count, Name Server registration and stack memory allocation.

For instructions on configuring Recovery Server parameters, see the following topic in Configuration Manager online help:

- Parameters Page, Recovery Server Component

## Internet Communications Configuration

The Ingres Web Deployment Option provides the foundation for internet-based electronic commerce. Using this component, users can quickly and easily build Web pages with corporate data and deploy database Web applications.

For more information, see the *Web Deployment Option User Guide*.

Internet communication parameters include setting the full path to the Web server's Common Gateway Interface (CGI) directory and the full path to the primary HTML document directory. Configuration of this component is usually done at installation time, when the output directories are configured, and rarely requires reconfiguration.

For instructions on configuring Internet communications parameters, see the following topic in Configuration Manager online help:

- Parameters Page, Internet Communications Component

## Locking System Configuration

The locking system controls concurrent access to a database.

Reconfiguration of the locking system is often required. The locking system can be configured to set the default locking level (`system_lock_level`) and the maximum number of locks that can be obtained per transaction (`system_maxlocks`). For detailed information on the locking system, see the *Database Administrator Guide*.

For instructions on configuring the locking system parameters, see the following topic in Configuration Manager online help:

- Parameters Page, Locking System Component

## Logging System Configuration

The logging system implements a circular transaction log file for the management of transactions within the installation. It ensures that log records are written in a way that makes them accessible to the Recovery Server and archiver process. All DBMS and Star Servers in the installation, as well as the archiver process share the logging system. Reconfiguration of the logging system is normally required to provide optimal performance for an installation.

Configurable parameters for the transaction log files include: transfer block size, buffer count, full limit, logging memory status, location, name, and number of partitions.

For instructions on configuring logging parameters, see the following topic in Configuration Manager online help:

- Parameters Page, Logging System Component

For a detailed explanation of the logging system, see the *Database Administrator Guide*. For instructions on creating and configuring a raw log file for UNIX platforms, see the *Getting Started* guide.

**Note:** Before making changes to the transaction log, Ingres must be shut down using Ingres Visual Manager, Service Manager, or the `ingstop` command. This ensures all processes are stopped correctly and completely.

## Multiple Log Partition Configuration

Ingres allows both the primary and dual transaction log files to exist on multiple devices (up to 16). Ingres improves logging system performance by allowing it to write to multiple disk devices and, on 32-bit file systems, allowing the total log file size to be greater than 2 GB.

You can configure multiple log partitions using the `log_file_parts` parameter. When configuring multiple log partitions, keep the following in mind:

- All log partitions must be the same size
- The primary and dual transaction logs must use the same number of partitions

For instructions on configuring multiple log partitions, see the following topic in Configuration Manager online help:

- Parameters Page, Logging System Component



## Primary Transaction Log and Dual Transaction Log

The Transaction Log pages display information about the primary and dual transaction logs. If the Ingres installation is running, the information on these pages is read-only. If it is not running, and no recovery is required from a system failure, these pages allow various operations to be performed on the transaction log files. These operations include editing the log file name, setting the log size per file, and adding, modifying or deleting locations for the primary and dual transaction log files.

For instructions on configuring transaction log parameters, see the following topic in Configuration Manager online help:

- Primary Log Page, Transaction Log Component

## Non-Configurable Components

The following components cannot be configured.

- Archiver Processes
- Remote Command Server

## History of Configuration Changes

Using the History of Changes page, you can view a list of all the changes to parameters that have been made by Configuration Manager or the Configuration-By-Forms (cbf) utility. The list contains changes beginning with the least recent. Each "CHANGE" line is followed by a list of all modified parameters that were derived because of the user change. Pressing F3 or Ctrl+F (in Configuration Manager) displays a Find dialog for searching for text.



# Chapter 4: Setting Environment Variables and Logicals

---


This chapter describes the most commonly used Ingres environment variables and logicals.

## Scope of the Environment Variables


Ingres has a set of available logical names through which you can define various file names and other values. For an alphabetical list of the basic Ingres environment variables and logicals, see the appendix “Environment Variables and Logicals.”

Ingres uses the following types of operating-system dependent environment variables and logicals:

### **Windows:**


- Ingres environment variables 

### **UNIX:**

- UNIX environment variables
- Installation-wide environment variables 

### **VMS:**

- Group-level Ingres logicals
- User-defined or process-level logicals

Ingres logicals are defined in the appropriate logical name table. 

Ingres environment variables in UNIX and Windows are active only when you are using Ingres. A small subset of these environment variables can be reset by the user in the local environment.

## Symbol Table

Most Ingres environment variables in Windows and UNIX are set in the Ingres symbol table (symbol.tbl) and are visible only with the user command `ingprens`.

When setting, unsetting, or changing environment variables, a backup of the symbol.tbl file, called 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.log file contains information on the types of updates performed as well as some application information. The symbol.bak and symbol.log files are located in the II\_CONFIG location or the II\_ADMIN location (in the case of NFS Ingres instances).

## Windows Environment Variables

Windows environment variables are set at the operating system level. These take effect before, during, and after Ingres is invoked. II\_SYSTEM and PATH are the only environment variables used directly by Ingres. These are set in the user's environment using the following commands entered at the operating system prompt:

```
set environment_variable=value
```

The Windows environment variables are set using Ingres Visual Manager (IVM). For more information, see the IVM online help topic, Parameters Page, Ingres Installation branch.

Alternatively, you can set Windows environment variables through the Control Panel.

## UNIX Environment Variables

UNIX environment variables are set at the UNIX operating system level. These are in effect before, during, and after Ingres is invoked. II\_SYSTEM, PATH, and TERM are the only UNIX environment variables required and used directly by Ingres. These environment variables are set in each user's environment using the following UNIX commands entered at the operating system prompt:

C Shell:

```
setenv environment_variable value
```

Bourne Shell:

```
environment_variable= value; export environment_variable
```

The environmental variables are typically set in a user's .login or .profile file, but can also be set in operating system startup scripts such as /etc/login.

## VMS Logicals

VMS logicals are set in the system or group logical table and are visible with the system command, show logical/system or show logical/group. These are in effect before, during, and after Ingres is invoked. II\_SYSTEM is the only VMS logical required and used directly by Ingres. This logical is defined by the system administrator using the following VMS command entered at the operating system prompt:

```
define [/table_level] logical_name "value"
```

The value assigned to the Ingres logical must be enclosed in double quotes.

This definition is typically placed in the operation-system wide command procedure for defining site-specific logical names (sys\$manager:sylogicals.com).

## Installation-Wide Environment Variables and Logicals

Ingres environment variables and logicals that are defined at the Ingres system or installation-wide level affect all users in an installation. These are usually defined during the install procedure. However, the system administrator can reset some manually.

## View Environment Variables in Windows and UNIX

Use Ingres Visual Manager to view Ingres environment variables and to manually register and deregister them. Additionally, you can use it to override certain environment variables at the user level.

To view environment variables in Ingres Visual Manager:

1. Select the Ingres Installation branch in Ingres Visual Manager.
2. Select the Parameters tab and select the System subtab (if not already selected).

The Ingres parameters that are set at the installation (system) level are displayed.

3. Select the Show Unset Parameters check box.

Parameters that are not currently set are displayed.

If the User subtab is selected, those parameters that can be overridden at the user level are displayed.

4. Select the Extra subtab.

Any extra Ingres environment variables that are added or variables whose names are dynamic are displayed. You can edit these parameters from this screen.

Alternatively, to view all installation-wide environment variables from the command line, enter the following command at the operating system prompt:

```
ingprens
```

To manually register or to remove an environment variable from the Ingres symbol table, use the `ingsetenv` and `ingunset` utilities. Always use these commands to alter the symbol table. Never edit this file directly.

A similar set of system administrator commands are available to set, print, and unset installation-wide Ingres environment variables simultaneously for all clients of an Ingres installation server. These commands are `ingsetall`, `ingprall`, and `ingunsetall`.

## View Environment Variables in VMS

Ingres logicals (with the exception of II\_SYSTEM) are defined in II\_CONFIG:config.dat during the install process. These logicals can be redefined locally using the operating system command `define/job logical_name "value."` However, to update the logical definitions permanently, the II\_CONFIG:config.dat file must be manually edited and the logical definitions modified.

To facilitate user access to Ingres tools and logical definitions, add the following access commands to the users' login.com file when you log into the system.

For users, add:

```
@II_SYSTEM:[ingres]ingusrdef.com
```

For database administrators, add:

```
@II_SYSTEM:[ingres]ingdbadef.com
```

For the system administrator, add:

```
@II_SYSTEM:[ingres]ingsysdef.com
```

The commands in these files provide tools access to all users in a system-level installation and to all users with the appropriate group user identification code (UIC) in a group-level installation. The Ingres logicals contained in II\_CONFIG:config.dat are redefined at a job level for each user when one of the above commands is executed.

You can display all installation-wide Ingres logicals by typing the following command at the operating system prompt:

```
show logical [/system|/group|/job] ii*
```

## Example: Set the ING\_EDIT Variable

### Windows and UNIX:

To set ING\_EDIT for an installation using Ingres Visual Manager, follow these steps:

1. Log in as the system administrator
2. In the Ingres Visual Manager, select the Ingres Installation branch.
3. Select the Parameters tab, and select the System subtab.

The Ingres parameters that are set at the installation (system) level are displayed.

**Note:** For variables that have a dynamic name, select the Extra tab instead of the System tab.

4. Select the Show Unset Parameters check box to display parameters that are not currently set.
5. Choose ING\_EDIT from the Parameters list, and double-click in the value area.
6. Edit the value and press Enter.
7. Alternatively, click Add and scroll to choose the ING\_EDIT variable. Enter the required value, and click OK.

To set ING\_EDIT using a system command, do the following:

1. Log in as the system administrator
2. At the operating system prompt type the following:

### Windows:

```
ingsetenv ING_EDIT D:\TOOLS\VI.EXE
```

### UNIX:

```
ingsetenv ING_EDIT /usr/bin/vi
```

### VMS:

```
define [/table_level] logical_name "value"
```

**Note:** On VMS, the value assigned to the Ingres logical must be in double quotes.



## User-Definable Environment Variables and Logicals

Some Ingres environment variables and logicals can be set or reset by individual users in their local environment using operating system commands. Those set in a user's local environment supersede the Ingres environment variables and logicals set system-wide.

**UNIX:** A good place to set user-defined environment variables is the user's .login (for the C shell) or profile (for the Bourne Shell) file.

**VMS:** Individual users can define Ingres environment variables at the job or process level with DCL commands, or they can be defined in the user's login.com file.

### Example: Set an Environment Variable in the Local Environment

One environment variable usually set in the user's local environment is TERM\_INGRES. It specifies the termcap definition to be used by the forms system. It can be redefined locally by entering commands at the operating system prompt as in the following example:

**Windows:**

```
SET TERM_INGRES=IBMPC
```

**UNIX:**

C Shell:

```
setenv TERM_INGRES vt100f
```

Bourne Shell:

```
TERM_INGRES=vt100f export TERM_INGRES
```

**VMS:**

```
DEFINE TERM_INGRES "VT100F"
```

## Display Current Value for Variables

A user can display the values set in their environment with the following command entered at the operating system prompt:

**Windows:**

SET

**UNIX:**

BSD:

printenv

System V:

env

## Locally Resettable Environment Variables and Logicals

In general, only the following Ingres environment variables and logicals must be set in your local environment:

### **Windows:**

- II\_SYSTEM
- PATH

### **UNIX:**

- II\_SYSTEM
- PATH

### **VMS:**

II\_SYSTEM

The following Ingres environment variables and logicals can be reset by users in their local operating system shell:

- DBNAME\_ING
- DBNAME\_SQL\_INIT
- II\_4GL\_DECIMAL
- II\_ABF\_RUNOPT
- II\_AFD\_TIMEOUT
- II\_APPLICATION\_LANGUAGE
- II\_DATE\_CENTURY\_BOUNDARY
- II\_DATE\_FORMAT
- II\_DBMS\_SERVER
- II\_DECIMAL
- II\_DML\_DEF
- II\_EMBED\_SET
- II\_FRS\_KEYFIND
- II\_GC\_REMOTE
- II\_GCA\_LOG
- II\_GCx\_TRACE
- II\_HELP\_EDIT
- II\_LANGUAGE
- II\_MONEY\_FORMAT

- II\_MONEY\_PREC
- II\_NULL\_STRING
- II\_PATTERN\_MATCH
- II\_PF\_NODE
- II\_POST\_4GLGEN
- II\_PRINTSCREEN\_FILE
- II\_SQL\_INIT
- II\_SYSTEM
- II\_TEMPORARY
- II\_TERMCAP\_FILE
- II\_TFDIR
- II\_TIMEZONE\_NAME
- II\_TM\_ON\_ERROR
- II\_VNODE\_PATH
- II\_WORK
- IIDLDIR
- ING\_ABFDIR
- ING\_ABFOPT1
- ING\_EDIT
- ING\_PRINT
- ING\_SET
- ING\_SET\_DBNAME
- ING\_SHELL
- INGRES\_KEYS
- INIT\_INGRES
- TERM
- TERM\_INGRES

**Important!** *II\_TIMEZONE\_NAME* can be reset for client installations only. Server installations must not reset this logical because it may affect date conversions from the local system time to the internal GMT-based value.

## Non-Locally Resettable Environment Variables and Logicals

The following Ingres environment variables and logicals must not be reset by users:

- II\_BIND\_SVC\_XX
- II\_C\_COMPILER
- II\_CHARSET\_XX
- II\_CHECKPOINT
- II\_CLIENT
- II\_COLLATION
- II\_CONFIG
- II\_CONNECT\_RETRIES
- II\_DATABASE
- II\_DBMS\_LOG
- II\_DIRECT\_IO
- II\_DUMP
- II\_ERSEND
- II\_GCN\_XX\_PORT
- II\_INSTALLATION
- II\_JOURNAL
- II\_LOG\_DEVICE
- II\_MSGDIR
- II\_NUM\_SLAVES
- II\_TUXEDO\_LOC
- II\_TUX\_SHARED
- II\_TUX\_AS\_MAX
- II\_TUX\_XN\_MAX
- II\_XA\_TRACE\_FILE
- ING\_SYSTEM\_SET

**UNIX:** These environment variables and logicals must not be visible using the env or printenv command.

## Customization of Query Runtime Environment

Most aspects of the environment in which queries execute can be customized at runtime for a particular session. This can be done by executing one or more set statements to customize the environment before query execution begins.

Here are some examples:

- Reduce lock conflicts when running read-only queries, as for reports:

```
set lockmode session where readlock=nolock
```

- Display the text of the query as it is executed:

```
set printqry
```

Ingres makes customization flexible by permitting most components of the query environment to be set by means of set statements at the startup of individual queries (sessions). This can be done by using various Ingres environment variables and logicals and startup files that permit users to control both the target and duration of set statement customization.

A full list of set statements can be found in your query language reference guide.

## Syntax Rules for Startup Files and Environment Variables

Keep the following general syntax rules in mind while setting up startup files and environment variables:

- Set statements in startup files cannot contain more than one set statement per line. Each statement must be on a line by itself.
- Set statements that are read from a file, but executed only by the single-line terminal monitors, must be terminated with “\g”.
- Set statements defined directly inside Ingres environment variables and logicals (not written in a file) never contain the string “\g”.
- Set statements defined directly inside Ingres environment variables and logicals (not written in a file) can contain multiple set statements separated by a semicolon up to a total length of 64 characters.
- The terminal monitor dayfile cannot contain set statements. It contains informational startup messages only.
- Ingres environment variables and logicals set with the following commands are global for the Ingres installation and affect the target of the environment variable/logical for all users in an installation:

### **Windows and UNIX:**

ingsetenv

### **VMS:**

For a system level installation:

define/system

For a group level installation:

define/group

- Ingres environment variables and logicals set from the local user environment apply only to the user setting them. Other users are not affected. This is the case, for example, when Ingres environment variables and logicals are set interactively, or:

**Windows:** Through the System icon in the Control Panel.

**UNIX:** From a user's .login, .profile, or .cshrc file.

**VMS:** Defined at the process level by the individual user with the VMS command define/process.

## Dayfile

This name affects all users of the SQL and QUEL single-line terminal monitors.

Dayfile contains text messages only (not set statements), which appear as an informational header whenever the terminal monitors are started up. It is located in `$II_SYSTEM/ingres/files` directory. The message header can be suppressed with the `-d` command line flag.

## DBNAME\_ING

This name applies to the QUEL single-line terminal monitor only. `DBNAME_ING` points to set statements within a file that are executed whenever this Ingres environment variable takes effect.

`DBNAME_ING` affects users who connect to the database specified by `DBNAME` through the QUEL terminal monitor. If set globally with `ingsetenv`, it affects all users of `DBNAME`. If set interactively or locally in this user's `.login`, `.profile`, or `.cshrc` file it affects only this user.

Setting the `DBNAME_ING` environment variable is equivalent to a user's executing `\i filename` in the QUEL terminal monitor each time they connect to `DBNAME`.



### Example: Set DBNAME\_ING

To define this Ingres environment variable, use the following command syntax at the operating system prompt:

**Windows:**

```
set DBNAME_ING=path_to_file
```

where:

*DBNAME* is the name of the database and must be in uppercase.

For example:

```
set MYDB_ING= D:\usr\george\mystartfile
```

**UNIX:**

C Shell:

```
setenv DBNAME_ING path_to_file
```

Bourne Shell:

```
DBNAME_ING=path_to_file export DBNAME_ING
```

where:

*DBNAME* is the name of the database and must be in uppercase.

For example:

```
setenv MYDB_ING /usr/george/mystartfile
```

The file contains lines in the following format:

```
set lockmode session where readlock = no-lock \g
```

This file contains the set statements. If you have several set statements, separate the statements with a space. You must end the entire file with “\g”.

## DBNAME\_SQL\_INIT

This name applies to the SQL single-line terminal monitor only. *DBNAME\_SQL\_INIT* points to set statements within a file that are executed whenever this Ingres environment variable/logical takes effect.

*DBNAME\_SQL\_INIT* affects users who connect to the database specified by *DBNAME* through the SQL Terminal Monitor. If set globally, it affects all users of *DBNAME*. If set interactively or locally, it affects only this user.

**Windows:** Global setting is with `ingsetenv`, or locally in the user's environment.

**UNIX:** Global setting is with `ingsetenv`, or locally in the user's `.login`, `.profile`, or `.cshrc` file.

**VMS:** Global setting is with `define/system`, `define/group`, or locally with `define/process` in the user's `login.com` file.

Setting the *DBNAME\_SQL\_INIT* environment variable/logical is equivalent to a user's executing `\i filename` in the SQL Terminal Monitor each time they connect to *DBNAME*.

### Example: Set DBNAME\_SQL\_INIT

To define this Ingres environment variable, use the following command syntax at the operating system prompt:

#### **Windows:**

```
SET DBNAME_SQL_INIT = path_to_file
```

where:

*DBNAME* is the name of the database and must be in uppercase.

For example:

```
SET MYDB_SQL_INIT=c:\user\mystartfile
```

#### **UNIX:**

C Shell:

```
setenv DBNAME_SQL_INIT path_to_file
```

For example:

```
setenv MYDB_SQL_INIT /usr/george/mystartfile
```

Bourne Shell:

```
DBNAME_SQL_INIT=path_to_file export DBNAME_SQL_INIT
```

For example:

```
MYDB_SQL_INIT=/usr/george/mystartfile export MYDB_SQL_INIT
```

where:

*DBNAME* is the name of the database and **must** be in uppercase.

#### **VMS:**

```
define DBNAME_SQL_INIT path_to_file
```

For example:

```
DEFINE /PROCESS MYDB_SQL_INIT DUA1:[USER.GEORGE]MYSTART.FILE
```

where:

*DBNAME* is the name of the database and **must** be in uppercase.

The file contains lines in the following format:

```
set lockmode session where readlock = no!ock \g
```

This file contains the set statements. If you have several set statements, separate the statements with a semicolon (;). You must end the entire file with "\g".

## II\_SQL\_INIT

This name applies to the SQL single-line Terminal Monitor only. II\_SQL\_INIT points to set statements within a file that are executed whenever this Ingres environment variable/logical takes effect.

II\_SQL\_INIT affects users who connect to the SQL terminal monitor. If set globally, it affects all users. If set interactively or locally, it affects only this user.

**Windows:** Global setting is with `ingsetenv`, or locally in the user's environment.

**UNIX:** Global setting is with `ingsetenv`, or locally in the user's `.login`, `.profile` or `.cshrc` file

**VMS:** Global setting is with `define/system`, `define/group`, or locally with `define/process` in the user's `login.com` file.

Setting this environment variable/logical is equivalent to a user's executing `\i filename` in the SQL terminal monitor each time they connect to a database.

### Example: Set II\_SQL\_INIT

To define this Ingres environment variable, use the following command syntax at the operating system prompt:

**Windows:**

```
SET II_SQL_INIT=path_to_file
```

**UNIX:**

C Shell:

```
setenv II_SQL_INIT path_to_file
```

For example:

```
setenv II_SQL_INIT /usr/george/mystartfile
```

Bourne Shell:

```
II_SQL_INIT=path_to_file export II_SQL_INIT
```

For example:

```
II_SQL_INIT=/usr/george/mystartfile export II_SQL_INIT
```

**VMS:**

```
define II_SQL_INIT path_to_file
```

For example:

```
DEFINE /PROCESS II_SQL_INIT -DUA1:[USER.GEORGE]MYSTART.FILE
```

The file contains lines in the following format:

```
set lockmode session where readlock = nolock \g
```

This file contains the set statements. If you have several set statements, separate the statements with a semicolon (;). You must end the entire file with "\g".

## ING\_SET

This name applies to SQL, QUEL single-line terminal monitors as well as Ingres tools such as ABF, VIFRED, embedded SQL, QBF, IQUEL, ISQL, and so forth. ING\_SET points to set statements that are executed whenever this Ingres environment variable/logical takes effect.

ING\_SET affects users who connect to an application, an Ingres tool, or a terminal monitor. If set globally, it affects all users. If set interactively or locally, it affects only *this* user.

**Windows:** Global setting is with `ingsetenv`, or locally in the user's environment.

**UNIX:** Global setting is with `ingsetenv`, or locally in the user's `.login`, `.profile`, or `.cshrc` file

**VMS:** Global setting is with `define/system`, `define/group`, or locally with `define/process` in the user's `login.com` file.

ING\_SET is set to a string surrounded by quotes. The string must be 64 characters or less, or it is invalid. The string can contain either:

- One or more set statements totaling no more than 64 characters
- The word `include` followed by the full path name to a file containing any number of set statements

## Example: Set ING\_SET

To define this Ingres environment variable, use the following command syntax at the operating system prompt:

### Windows:

```
SET ING_SET = set-statement{; set statement}
```

Examples:

```
SET ING_SET=set lockmode show  
SET ING_SET=include C:\users\default\mystartfile
```

### UNIX:

C Shell:

```
setenv ING_SET 'set-statement {; set-statement}'
```

or

```
setenv ING_SET 'include path_to_file'
```

Examples:

```
setenv ING_SET 'set lockmode show'  
setenv ING_SET 'include /usr/george/mystartfile'
```

Bourne Shell:

```
ING_SET='set-statement {; set-statement}'
```

or:

```
ING_SET='include path_to_file' export ING_SET
```

Examples:

```
ING_SET='set lockmode show' export ING_SET  
ING_SET='include /usr/george/mystartfile' export ING_SET
```

### VMS:

```
define ING_SET "set-statement {;set-statement}"
```

or:

```
define ING_SET "include path_to_file"
```

Examples:

```
DEFINE /PROCESS ING_SET "SET LOCKMODE SHOW"

DEFINE /PROCESS ING_SET -
"INCLUDE DUA1:[USER.GEORGE]MYSTART.FILE"
```

For the include format, the file specified by *path\_to\_file* contains the set statements. If you have several set statements, separate the statements with a semicolon (;). Place each set statement on a separate line in the file. For example:

```
set autocommit on;
set lockmode session where readlock=nolock;
set result_structure cmtree
```

## ING\_SET\_DBNAME

This name applies to both SQL and QUEL single-line terminal monitors as well as Ingres tools such as ABF, VIFRED, embedded SQL, QBF, IQUEL, ISQL, and so forth. *ING\_SET\_DBNAME* points to set statements that are executed whenever this Ingres environment variable/logical takes effect.

*ING\_SET\_DBNAME* affects users who connect to the database specified by *DBNAME* through an application, an Ingres tool, or a terminal monitor. If set globally, it affects all users of *DBNAME*. If set interactively or locally, it affects only this user.

**Windows:** Global setting is with *ingsetenv*, or locally in the user's environment.

**UNIX:** Global setting is with *ingsetenv*, or locally in the user's *.login*, *.profile* or *.cshrc* file.

**VMS:** Global setting is with *define/system*, *define/group*, or locally with *define/process* in the user's *login.com* file.

*ING\_SET\_DBNAME* is set to a string surrounded by quotes. The string must be 64 characters or less, or it is invalid. The string can contain either:

- One or more set statements totaling no more than 64 characters
- The word *include* followed by the full path name to a file containing any number of set statements



**Example: Set ING\_SET\_DBNAME**

To define this Ingres environment variable/logical use the following command entered at the operating system prompt:

**Windows:**

```
SET ING_SET_DBNAME=set-statement {;set-statement}
```

For example:

```
SET ING_SET_MYDB=set lockmode show
SET ING_SET_MYDB=include C:\users\default\mystartfile
```

**UNIX:**

C Shell:

```
setenv ING_SET_DBNAME 'set-statement {; set-statement}'
```

or:

```
setenv ING_SET_DBNAME 'include path_to_file'
```

For example:

```
setenv ING_SET_MYDB 'set lockmode show'
setenv ING_SET_MYDB 'include /usr/george/mystartfile'
```

Bourne Shell:

```
ING_SET_DBNAME='set-statement {; set-statement}'
```

or:

```
ING_SET_DBNAME='include path_to_file' export ING_SET_DBNAME
```

For example:

```
ING_SET_MYDB='set lockmode show'export ING_SET_MYDB
ING_SET_MYDB='include /usr/george/mystartfile'export ING_SET_MYDB
```

**VMS:**

```
define ING_SET_DBNAME "set-statement {; set-statement}"
```

or:

```
define ING_SET_DBNAME "include path_to_file"
```

For example:

```
DEFINE /PROCESS ING_SET_MYDB - "SET LOCKMODE SHOW"

DEFINE /PROCESS ING_SET_MYDB - "INCLUDE DUA1:[USER.GEORGE]MYSTART.FILE"
```

For the include format, the file specified by *path\_to\_file* contains the set statements. If you have several set statements, separate the statements with a semicolon (;). Place each set statement on a separate line in the file.

For example:

```
set autocommit on;
set lockmode session where readlock=nolock;
set result_structure cbtree
```

## ING\_SYSTEM\_SET

This name applies to both SQL and QUEL single-line terminal monitors as well as Ingres tools such as ABF, VIFRED, embedded SQL, QBF, IQUEL, ISQL, and so forth. ING\_SYSTEM\_SET points to set statements that are executed whenever this Ingres environment variable/logical takes effect.

ING\_SYSTEM\_SET affects all users who connect to a DBMS Server through an application, an Ingres tool, or terminal monitor. It is always effective globally and can be set only by a privileged user.

ING\_SYSTEM\_SET is set to a string surrounded by quotes. The string must be 64 characters or less, or it is invalid. The string can contain either:

- One or more set statements totaling no more than 64 characters
- The word include followed by the full path name to a file containing any number of set statements

**Example: Set ING\_SYSTEM\_SET**

To define this Ingres environment variable/logical use the following command entered at the operating system prompt:

**Windows:**

```
SET ING_SYSTEM_SET = set-statement{;set-statement}
```

For example:

```
SET ING_SYSTEM_SET = set autocommit on
```

**UNIX:**

C Shell:

```
setenv ING_SYSTEM_SET 'set-statement{; set-statement}'
```

For example:

```
setenv ING_SYSTEM_SET 'set autocommit on'
setenv ING_SYSTEM_SET 'include \
/usr/ingres/ourstartfile'
```

Bourne Shell:

```
ING_SYSTEM_SET='set-statement {; set-statement}'
```

or:

```
ING_SYSTEM_SET='include path_to_file'
```

```
export ING_SYSTEM_SET
```

For example:

```
ING_SYSTEM_SET='set autocommit on'
```

```
export ING_SYSTEM_SET
```

```
ING_SYSTEM_SET='include \ /usr/ingres/ourstartfile'
```

```
export ING_SYSTEM_SET
```

**VMS:**

```
define ING_SYSTEM_SET "set-statement {; set-statement}"
```

or:

```
define ING_SYSTEM_SET "include path_to_file "
```

For example:

```
DEFINE /SYSTEM ING_SYSTEM_SET -  
"SET LOCKMODE SHOW"  
  
DEFINE /SYSTEM ING_SYSTEM_SET -  
"INCLUDE DUA1:[INGRES]OURSTART.FILE"
```

For the include format, the file specified by *path\_to\_file* contains the **set** statements. If you have several set statements, separate the statements with a semicolon (;). Place each set statement on a separate line in the file.

For example:

```
set autocommit on;  
set lockmode session where readlock=nolock;  
set result_structure cbtree;
```

## INIT\_INGRES

This name applies to the QUEL single-line terminal monitor only. INIT\_INGRES points to set statements within a file that are executed whenever this Ingres environment variable/logical takes effect.

INIT\_INGRES affects users who connect to the QUEL terminal monitor. If set globally, it affects all users. If set interactively or locally, it affects only *this* user.

**Windows:** Global setting is with `ingsetenv`, or locally in the user's environment.

**UNIX:** Global setting is with `ingsetenv`, or locally in the user's `.login`, `.profile`, or `.cshrc` file.

**VMS:** Global setting is with `define/system`, `define/group`, or locally with `define/process` in the user's `login.com` file.

Setting INIT\_INGRES variable/logical is equivalent to a user's executing `\i filename` in the QUEL terminal monitor each time they connect to a database.

### Example: Set INIT\_INGRES

To define this Ingres environment variable/logical use the following command entered at the operating system prompt:

#### **Windows:**

```
SET INIT_INGRES=path_to_file
```

For example:

```
SET INIT_INGRES= c:\user\mystartfile
```

#### **UNIX:**

C Shell:

```
setenv INIT_INGRES path_to_file
```

For example:

```
setenv INIT_INGRES /usr/george/mystartfile
```

Bourne Shell:

```
INIT_INGRES=path_to_file export INIT_INGRES
```

For example:

```
INIT_INGRES=/usr/george/mystartfile export INIT_INGRES
```

#### **VMS:**

```
define INIT_INGRES path_to_file
```

For example:

```
DEFINE /PROCESS INIT_INGRES - DUA1:[USER.GEORGE]MYSTART.FILE
```

The file contains lines in the following format:

```
set lockmode session where readlock = no-lock \g
```

This file contains the set statements. If you have several set statements, separate the statements with a semicolon (;). You must end the entire file with "\g".

## Startup

This name affects the QUEL single-line terminal monitor only.

The Ingres environment variable/logical startup affects all users of the QUEL terminal monitor. It is a customizable file (delivered with some macros already defined). It is located at: II\_SYSTEM/ingres/files/startup.

Users with permission to edit this file can add statements of the such as the following:

```
{DEFINE RET; RETRIEVE}
```

The result of using this macro startup file is that users typing "RET TABLE1.ALL" has the string "RET" substituted by the string "RETRIEVE" at query execution time. For an explanation of how macros are used and defined, see the *QUEL Reference Guide*.

The quantity of some operating system resources available to each user session is fixed at Ingres DBMS Server start up. These include resources such as the quantities of memory available for parsing, storing and optimizing a query. Default values are used for these resources at installation time. These do not need to be changed unless you are receiving explicit error messages that indicate that these resources are inadequate.

## Startsql

This name affects the SQL single-line terminal monitor only.

Startsql affects all users of the SQL terminal monitor. It is a customizable file (delivered empty on the release tape). It is located at: II\_SYSTEM/ingres/files/startsql.

Users with permission to edit this file can add statements of the following example:

```
set lockmode session where readlock = noLock \g
```

## Installation Locations

The values for the following Ingres variables are established at the time of installation:

- II\_SYSTEM
- II\_DATABASE
- II\_CHECKPOINT
- II\_JOURNAL
- II\_DUMP
- II\_WORK
- II\_TEMPLATE
- ING\_ABFDIR

Some of these values can be easily changed, and some require unloading of the installation.

### When Not to Change Installation Locations

Before changing an existing Ingres installation location, be sure that the change is really necessary: ask yourself what needs to be moved, and why.

Installation locations do not need to be changed if:

- The installation needs space for future growth.
- Existing databases are outgrowing their present disk partitions.

### Allocation of Space for Future Growth

Ingres provides the capability of creating new locations and placing any new databases, tables, or indexes in these new locations. If the requirement is space for **new** tables or databases, you do not need to change existing locations.

Instead, create the necessary new locations for these future tables and databases and use the available location flags of `createdb` (`-d`, `-c`, `-j`, `-b`, and `-w`) to create the data, checkpoint, journal, dump, and work files for the newly created database in some other location.

For more information on creating new databases, see the `createdb` command in the *Command Reference Guide*.

## Allocation of Additional Space for Existing Tables

If space is becoming limited on a disk partition containing a particularly large table, you can move the table to a new disk or partition. Use the `modify` command to relocate and `modify` to reorganize statements as described in the *Database Administrator Guide*.

## Location Variables that Do and Do Not Require Unloading (UNIX and Windows)

When you need to change the value of an installation location variable, the procedure depends to some extent on the underlying operating system. On UNIX and Windows, the procedure also depends on whether the change requires unloading the database.

You can change the `ING_ABFDIR` location value without unloading the installation.

The values of the following locations are stored in the configuration file "aaaaaaa.cnf" of each database at the time it is created and can only be changed by unloading each database, re-installing Ingres and reloading the databases (thus creating a new configuration file):

- `II_SYSTEM`
- `II_DATABASE`
- `II_CHECKPOINT`
- `II_JOURNAL`
- `II_DUMP`
- `II_WORK`

## Change the Value of the `ING_ABFDIR` Variable (UNIX and Windows)

To change the installation default for `ING_ABFDIR`:

1. Shut down Ingres with `ingstop`
2. Change the value of `ING_ABFDIR` with the `ingsetenv` command
3. Restart the installation with `ingstart`

**Note:** Application developers who do not want to use the installation default can redefine `ING_ABFDIR` in their local user environment.



## Change the Value of the ING\_ABFDIR Logical (VMS)

To change the installation default value for ING\_ABFDIR in VMS:

1. Shut down Ingres with `ingstop`.
2. Change the value with a new define command.
3. Restart the installation with `ingstart`.

ING\_ABFDIR can be redefined in the local user environment for application developers who do not want to use the installation default.

## Change an Installation Location with Unloading (Windows)

Follow these steps to change the value of II\_SYSTEM, II\_DATABASE, II\_CHECKPOINT, II\_JOURNAL, II\_DUMP or II\_WORK. "II\_SYSTEM" is used as the example location to be changed:

1. Unload each database. For details, see the *Database Administrator Guide*.
2. Remove the old installation files by removing everything under the old %II\_SYSTEM%\ingres directory and any extended locations.
3. Reinstall Ingres specifying the new value of II\_SYSTEM.
4. Recreate each of the databases. For details, see the *Database Administrator Guide*.
  - Use the `createdb dbname` command to recreate a database.
  - Reload each database by appropriately editing and running the reload script.
5. Once you have reinstalled Ingres and recreated the databases in the new location, take new checkpoints of each of the databases. Checkpoints taken in old locations are no longer valid. For details, see the *Database Administrator Guide*.

## Change an Installation Location with Unloading (UNIX)

Follow these steps in UNIX to change the value of II\_SYSTEM, II\_DATABASE, II\_CHECKPOINT, II\_JOURNAL II\_DUMP, or II\_WORK. "II\_SYSTEM" is used as the example location being changed:

1. Unload each database. For details, see the *Database Administrator Guide*.
2. Remove the old installation files by removing everything under the old II\_SYSTEM/ingres directory and any extended locations.
3. Set the value of II\_SYSTEM and PATH to the desired new values, just as if you were installing Ingres for the first time.
4. Reinstall Ingres.

Use the installation medium and follow the instructions in *Getting Started* for a new (not upgrade) installation.

As part of the installation procedure, you are asked to run the program ingbuild. This program prompts you for the new values of Ingres location variables II\_DATABASE, II\_CHECKPOINT, II\_JOURNAL, II\_DUMP, and II\_WORK.

5. Recreate each of the databases. For details, see the *Database Administrator Guide*.

Use the createdb *dbname* command to recreate a database.

Reload each database by appropriately editing and running the reload.ing script.

6. Once you have reinstalled Ingres and recreated the databases in the new location, take new checkpoints of each of the databases. Checkpoints taken in old locations are no longer valid. For details, see the *Database Administrator Guide*.

## Change Location Variables (Other than ING\_ABFDIR) on VMS

Follow these steps in VMS to change the value of II\_SYSTEM, user areas, II\_DATABASE, II\_CHECKPOINT, II\_JOURNAL, II\_DUMP, or II\_WORK:

1. Shut down Ingres with ingstop.
2. Change the value of II\_SYSTEM, II\_DATABASE, II\_CHECKPOINT, II\_JOURNAL, II\_DUMP, II\_WORK, or user areas with the define command. Remember to modify any startup or other procedures where these logicals are defined.
3. Restart the installation with ingstart.

For more information about stopping and starting the installation using ingstop and ingstart, see the *Command Reference Guide*.

## Use of Symbolic Links (UNIX)

Ingres does not support or recommend the use of symbolic links to change locations in UNIX. (This is possible with the UNIX `ln` command, which allows the UNIX system administrator to create symbolic links to new directories so that files are found in another directory even though the values of Ingres installation location variables remain unchanged.)

The following are some common problems that arise from the misuse of symbolic links:

- Any files to which Ingres writes must reside on a file system that is local to the machine on which the writing process is running. This is necessary because of the operating system's inability to guarantee network writes. If symbolic links are used, writable directories can mistakenly be mounted from remote machines. This can cause write errors and data corruption.
- Occasionally Ingres executables reside on NFS file systems that have been exported with the `- nosuid` option. The permissions appear normal when the `ls -l` command is issued, but `setuid` programs fail with permission errors.
- The `rm` command does not remove files that are symbolic links. Instead, the symbolic link itself is broken, but the original file remains. This can cause cleanup routines to fail in some circumstances.



# Chapter 5: Customizing Ingres

---

This chapter describes how to implement the following customization options that are available in Ingres:

- Changing the character set or collation sequence
- Customizing the Archiver Exit Script

## Ingres Character Sets

Ingres supports read-in character sets. This feature allows you to use a character set other than the default character set at your installation.

For example, in Thailand, the standard character set is Thai TIS. If you specify the THAI character set during the installation procedure, users at your installation will save and retrieve data using this character set.

The installation utility prompts you with a list of currently supported character sets for your platform. It sets the environment variable or logical for II\_CHARSETxx to the value you have selected, unless the default character set is chosen. If the default character set is chosen, II\_CHARSETxx is not set.

**Important!** *Choose your character set carefully. An installation can use only one character set. Do not change the character set because data in existing databases can be corrupted.*

After running the install program, check the value of II\_CHARSETxx. If II\_CHARSETxx is set incorrectly, you must run the installation program again.

For a complete list of supported character sets, see the *Getting Started* guide. For a list of the character sets available on your platform, see the online list selection during installation.

## Default Character Sets

The default character sets are as follows:

**Windows:** WIN1252

**UNIX:** ISO-8859/1

**VMS:** DEC Multinational

DEC Multinational and ISO-8859/1 are extended ASCII character sets. They have identical definitions for the 128 7-bit characters. They differ only in the defaults they assign for several of the 128 additional characters in the 8-bit character set.

## Local Collation Sequences

A collation sequence determines the sorting order of any given character set. For example, in English, the order of the alphabet is most commonly used to order a list of English words. That is, words beginning with the letter c appear before words beginning with the letter d, which appear before words beginning with the letter e and so on.

A collation sequence is associated with each database when the database is created. This sequence determines in what order sorted data is returned to users and applications, what is returned when queries use pattern matching and, in some instances, how data is stored internally.

## Supported Collation Sequences

In a computer, if no other collation sequence is enforced, the sequence derived from the machine's native character set, either ASCII or EBCDIC is used. The sorting order of these character sets derives from the internal numeric representation of each character.

In addition to the sequences derived from ASCII and EBCDIC, Ingres supports two other local collation sequences. They are:

- Multi, a sequence derived from the DEC Multinational Character Set
- Spanish, a sequence derived from the Spanish language.

If none of these sequences adequately fills your needs, you can write your own local collation sequence.

## MultiCollation Sequence

The multi collation sequence is based on the DEC Multinational Character Set. This character set adds several vowels with diacritical marks to the standard 7-bit ASCII character set.

Following are the comparison sequences for the multi sequence that differ from those of ASCII:

A < Å < Á < Â < Ã < Ä < B  
 C < Ç < D < E < È < É < Ê < Ë  
 I < Ì < Í < Î < Ï < J  
 N < Ñ < O  
 O < Ò < Ó < Ô < Õ < Ö < Œ < P  
 U < Ù < Ú < Û < Ü < V  
 Y < Ÿ < Z < Æ < Ø < Å  
 a < à < á < â < ã < ä < b  
 c < ç < d < e < è < é < ê < ë  
 i < ì < í < î < ï < j  
 n < ñ < o  
 o < ò < ó < ô < õ < ö < œ < p  
 ss < ß < st  
 u < ù < ú < û < ü < v  
 y < ÿ < z < æ < ø < å

For example:

**cote < côte < czar < cæsar**

Pattern matching rules:

- Œ, œ, ß are special when pattern match searching is used.
- Œ matches O\_ or O% as well as Œ
- œ matches o\_ or o% as well as œ
- ß matches s\_ or s% as well as ß

## Spanish Collation Sequence

The Spanish collation sequence is based on the multi sequence but contains additional support for the Spanish letters ll and ch. Listed below are the comparison sequences for the Spanish collation sequence that differ from those of ASCII. Some pattern matching rules are also described.

A < Å < Á < Ä < Ä < B

CZ < CĂ < Ç < CH < Ch < D < E < È < É < Ê < Ë

I < Ì < Í < Î < Ï < J

LZ < LĂ < LL < LI < M

N < Ñ < O

O < Ò < Ó < Ô < Õ < Ö < Œ < P

U < Û < Ú < Û < Ü < V

Y < Ÿ < Z < Æ < Ø < Å

a < à < á < â < ã < ä < b

cz < câ < ç < cH < ch < d < e < è < é < ê < ë

i < ì < í < î < ï < j

lz < lâ < ll < ll < m

n < ñ < o

o < ò < ó < ô < õ < ö < œ < p

ss < ß < st

u < ù < ú < û < ü < v

y < ÿ < z < æ < ø < å

Examples:

loop < llama

cote < côte < czar < cæsar < chair

The pattern matching rules are:

- Œ, œ, ß, Ç, ç are special when pattern match searching is used.
- Œ matches O\_ or O% as well as Œ
- œ matches o\_ or o% as well as œ
- ß matches s\_ or s% as well as ß
- Ç matches C\_ or C% as well as Ç
- ç matches c\_ or c% as well as ç



## Custom Collation Sequence

If you have special needs that are not met by the available collation sequences, you can write your own. Ingres allows you to write a collation sequence that has any of the following characteristics:

- Character skipping—one or more specified characters are ignored for collation
- One-to-one mapping—a character can be substituted for another or weighted differently for collation
- Many-to-one mapping—groups of characters can be substituted for a single character or weight value for collation
- Many-to-many mapping—groups of characters can be substituted for a sequence of characters or weight values for collation

## Guidelines for Creating a Custom Collation File

Keep the following points in mind as you design and test your custom collation file:

- Never create a production database with an untested collation sequence. Always test your collation file on a sample database. Each time that you modify the collation sequence to correct any bugs, you must unload the database, destroy the old database, install the new sequence, create a database with the new sequence, and reload the database.
- Some collation sequences allow two strings that are different to compare as equal. These sequences are called information loss sequences. An example of this type of sequence is a sequence that ignores case. Problems that can result from such a sequence are:
  - If duplicates are not allowed, the DBMS drops all but one string.
  - If duplicates are not allowed, the DBMS does not allow you to add a row to a table if it appears to match an existing row.
  - The hash storage structure cannot detect when two equal but different strings are placed in a hashed relation.
  - In a query on a hash table, the “=” operator can only fetch one of the ‘equal’ strings that matches the expression.

Because of these problems, we suggest that you do not use information loss sequences and the hash storage structure together.

## How You Write a Customize Collation Sequence

To create a customized collation sequence, follow these steps:

1. Write a description file.
2. Run the description file through the `aducompile` utility.
3. Test your collation sequence with a small sample database.

## Description File—Describe Collation Sequence

To define a custom collation sequence you must create a description file, which consists of a list of “instructions” that, taken as a whole, describe the collation sequence. Each instruction must appear on a separate line in the file.

The format of each instruction is:

*value: string*

where:

### ***value***

Determines the numerical weight assigned to string. (The internal numerical weight of each character determines where a character appears in the sort order.)

The *value* can have any of the following formats:

### **char+number**

Instructs sorting of the specified string after the specified character and before the next higher-weighted character in the character set. For example, in the following instruction, *string1* is mapped as a single character that is ordered immediately after the letter H and before I in a sorted sequence:

H+1:string1

In the following instruction, *string2* sorts after *string1* and before the letter I.

H+2:string2

You can specify  $H+1$ : *string* or  $Hz+1$ : *string* and both sorts in the same manner, that is, after H and before I. However, the two examples do not behave the same when pattern matching is applied. To illustrate using an example from the Spanish language, the following instruction maps CH as a single character that exists between C and D:

```
C+1:CH
```

If you ask for a pattern match using the format C%, instances of CH are not returned. The alternative,  $Cz+1$ :CH, maps CH into two characters, C and a virtual character just after z. This causes CH to match as two characters. A pattern match using the format C% finds the instances of CH.

### **charstring**

Sorts the specified string as the equivalent of the specified charstring. For example, in the following instruction, the word tax sorts as if it were the word revenue:

```
revenue:tax
```

### **+ number**

Gives the specified string the internal numerical weight specified by given number. The number must be between 0 and 32766. The weighting of a character in this manner is less portable than giving the character a relative weight.

```
+ *
```

Causes the specified string to be ignored when collation is performed. For example, in the following instruction, the "?" is ignored whenever collation takes place.

```
+* ; ?
```

### **(empty)**

When no value is specified (the instruction takes the form: string), the collation compiler ignores the instruction. Use this format to insert comments into your collation sequence. For example:

```
:This is a comment
```

### **string**

Is any character or character string. An empty string causes a syntax error.

### The aducompile Utility

The aducompile utility compiles your description file into a binary file and installs that file as a collation sequence that can be used. You must be the installation owner to use this utility. Be sure to give your resulting collation file a unique name so that you do not overwrite any existing collation files.

Your new collation sequence is located at  
`$II_SYSTEM/ingres/files/collation/collation_name.`

**Note:** In UNIX, all system users must have rights to read the new collation file.

## Archiver Exit Script (acpexit)

The Archiver Process Exit command script is executed by the archiver before an error halt. This shell script (acpexit in UNIX or ACPEXIT.COM in VMS) resides in the utility directory. You can use the default script behavior, which sends mail to the installation owner account with the reason for the archiver failure, or, the script allows you to customize certain aspects of archiver recovery.

Because the archiver writes sequentially to the end of journal files, a common problem occurs when the archiver stops as a result of running out of journal disk space. If left unresolved, running an installation without an archiver (and with journaled databases) eventually causes the log file to fill.

To address this concern, whenever the archiver detects an error condition that forces it to stop (including running out of journal disk space) it first executes the archiver exit script.

## Customization of Archiver Exit Script

You can modify the archiver exit script to take action appropriate to your installation. For example, one possible course of action is to delete a reserved file on the journal disk to free up space and restart the archiver.

You can customize the archiver exit script in these respects:

- The default script sends mail to the installation owner account when archive processing halts, and includes the reason for the failure. You can modify this script to notify the system administrator directly of a situation that needs resolving.
- You can tailor the script to handle certain resource and environmental problems automatically. For example, if archive processing has stopped because of insufficient disk space, the script can automatically reclaim disk space and resume archive processing.
- The default script is executed with one or two arguments, described in Archiver Exit Script Parameters (see page 101). You can choose whether to incorporate these arguments if you customize the script.

The default script includes examples of processing you can use for your installation. These examples are formatted as comments, and can be included in the script by removing the comment prefixes. The script is located at `$II_SYSTEM/ingres/files/acpexit`.

## Archiver Exit Script Parameters

The archiver exit script parameters are as follows:

Parameter	Description	Status
<i>errornumber</i>	An error number that describes the reason archive processing has stopped. At the end of the default script is a list of all possible error numbers, with an explanation of the error and, in many cases, the recommended actions to correct the error.	Always passed in default script.
<i>database_name</i>	The name of the database that was being processed at the time archiving stopped. The name is provided only if the halt in archive processing was associated with a particular database problem, as opposed to a general system problem.	Sometimes passed in default script.

## How You Use a Custom Mapping File for Unicode Coercion

The DBMS Server carries out the coercion of Unicode data to local character set value based on the Unicode converter (or mapping table). Ingres automatically chooses the converter to use, based on the encoding used by the locale on your system.

You can create a custom mapping table for coercing Unicode data to local character set value if, for example, you want to map a particular character to a different Unicode code point, or create a completely new mapping table for an encoding not yet supported by Ingres. Always use the custom mapping table for your encoding once you have implemented it.

Always create a new mapping file. If you wish to modify an originally supplied mapping table for an encoding do not modify the original mapping files but copy it to a new file name. To avoid data corruption, test your mapping before implementing in a production installation.

If you are using a custom converter (mapping table) for an encoding, make sure that you always use the same converter for that encoding.

The process for creating and using a custom mapping file is as follows:

1. Create a custom mapping file in XML format.

This format should conform to `charactermapping.dtd` present in the directory `$II_SYSTEM/ingres/files/ucharmaps`. This directory also contains other converters, which may serve as examples of xml file format. For more information on the xml file format, see <http://www.unicode.org/reports/tr22/>.

2. Compile this mapping file using `unimapcompile` utility, as follows:

```
unimapcompile -o=yourmapfilename yourmapfilename.xml
```

3. Set the mapping file in the symbol table for the server using the `ingsetenv` command, as follows:

```
ingsetenv II_UNICODE_CONVERTER yourmapfilename
```

To check that the value is set correctly, run:

```
ingprenv II_UNICODE_CONVERTER
```

4. If you have remote clients connecting to this server and will be using the encoding for which you have created the custom mapping table, then copy the compiled mapping file to all clients in the directory `$II_SYSTEM\ingres\files\ucharmaps`. Also ensure that you have set `II_UNICODE_CONVERTER` to your custom converter in the symbol table for the clients.

# Chapter 6: Troubleshooting Ingres

---

This chapter describes procedures and techniques for troubleshooting the Ingres installations. Using these procedures will help you to solve the problems yourself, or, to accurately define the problem if it is still necessary to call technical support.

## Process of Troubleshooting

Troubleshooting is a process of defining and correcting a problem that occurs in an otherwise functioning Ingres installation and includes narrowing down the problem to a well-defined point, identifying the cause of failure, and eliminating it. When troubleshooting, you:

1. Determine the nature of the problem
2. Isolate the problem to a defined area
3. Eliminate the cause, using the following techniques:
  - Correcting user errors
  - Changing the Ingres installation environment
  - Changing the user environment
  - Changing the operating system environment
  - Changing the system configuration
  - Restarting the system if needed

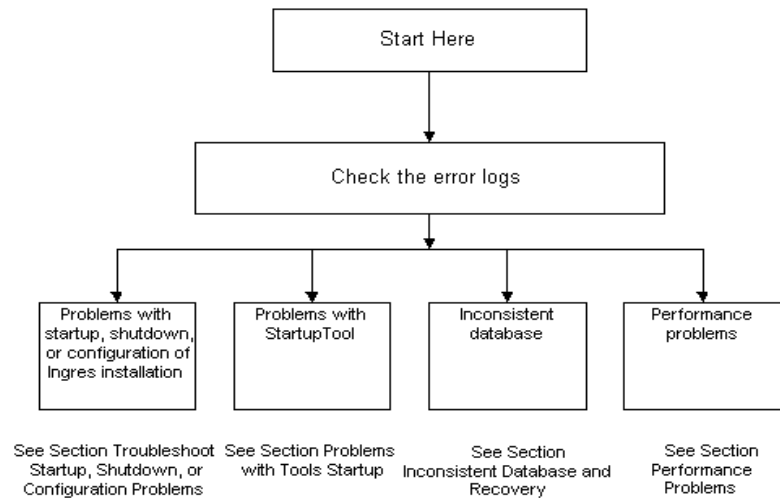
## Troubleshooting Tools

VDBA provides a set of system administration tools. It allows you to perform system administrator functions—including configuring, performance monitoring, backup and recovery, and remote database optimization—using a GUI interface. For instructions on performing these functions, see the VDBA online help.

Command line instructions are also available. For information on Ingres commands, see the Command Reference Guide.

## Determine the Problem Area

The first step in the process of troubleshooting is to determine the problem area. The following troubleshooting flow chart shows the major problem categories:



## Error Log Files

The error log files are located in the directory indicated by the Ingres environment variable/logical `II_CONFIG`. The log files are as follows:

### **errlog.log**

The main DBMS error log and the default log for most programs.

### **iiacp.log**

Archiver error log.

### **iircp.log**

Recovery error log.

**Note:** If your system is configured for Ingres Cluster Solution, each node in the Ingres cluster maintains a separate Archiver and Recovery error log. Each log is distinguished by having `_nodename` appended to the base log name, where `nodename` is the Ingres node name for the host machine as returned by `iipmhost`.

On VMS, if `SCSNODE` is set, the log files except `errlog.log` is named `*.LOG_nodename`.



## Optional Error Log Files

The names of optional log files can vary, but all log files end with the extension LOG. Optional log files include:

**II\_DBMS\_LOG**

DBMS error log

**II\_GCCxx\_LOG**

GCC error log

**II\_GC\_LOG**

GCC trace log

**II\_STAR\_LOG**

Star error log

## Display Value for II\_CONFIG

To display the value for II\_CONFIG, type the following command at the operating system prompt:

**Windows:**

```
ingprens
```

**UNIX:**

```
ingprens
```

**VMS:**

```
set def II_SYSTEM: [INGRES.FILES]
```

or

```
set def II_CONFIG.
```

## View List of Log Files

For a list of log files, type the following command at the operating system prompt from the directory indicated by \$II\_CONFIG (UNIX) or II\_CONFIG (VMS):

**Windows:**

```
dir *.log
```

**UNIX:**

```
ls *.log
```

**VMS:**

```
dir *.log
```

## Check the Error Log Files

Checking for error logs is the first step for determining the nature of the problem. To check for error logs, follow these steps:

1. Check the log files for error messages. Examine errlog.log first.
2. If indicated, check optional log files for error messages.
3. Identify the errors associated with your problem. All errors are time stamped. Find the most recent error message associated with your problem. Read back up the log file from there to the first error relating to that problem. Many error messages cascade from that initial error. This initial error is usually the most important in identifying the problem even though it is not the error that users report.

For example, the following error message merely notifies you that a DBMS server has exited for one of many possible reasons:

```
E_SC0221_SERVER_ERROR_MAX
```

You must look in the errlog.log for associated DBMS server errors such as "E\_DM9300\_DMOP\_CLOSE\_PAGE Buffer still fixed," a fatal error message. Search the error log for additional errors occurring around the time of the error that was displayed.

## Find Your Problem Category

To find the problem category, follow these steps:

1. Determine the general category in which your problem lies. The following categories relating to running the Ingres installation are described in this guide:

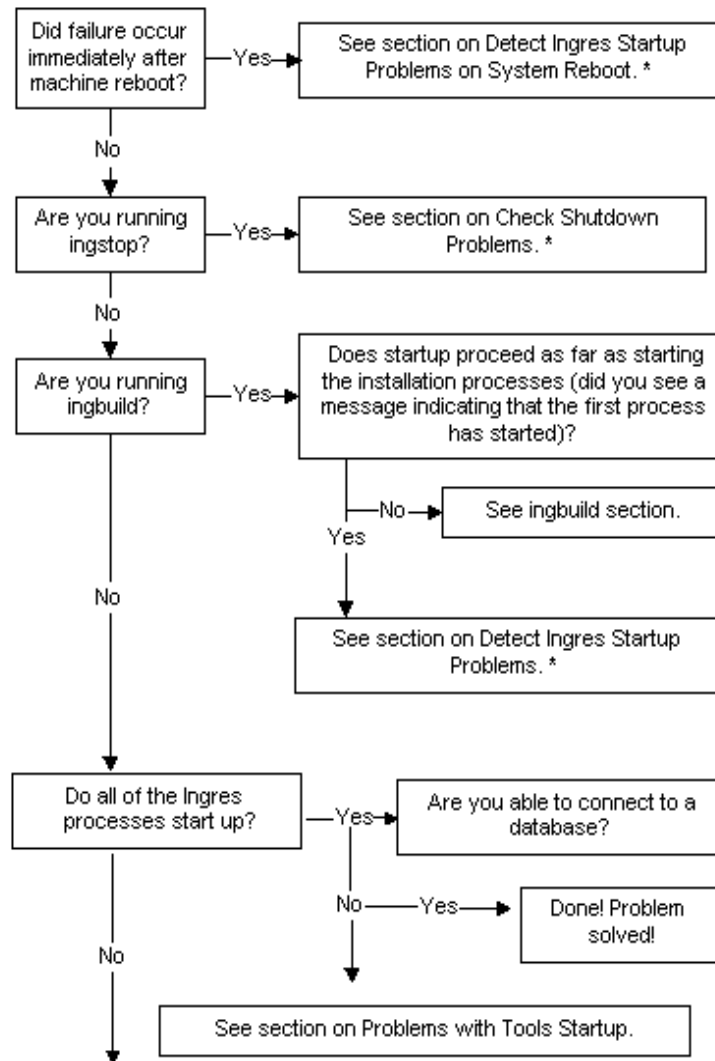
- Ingres startup and shutdown
- Ingres configuration
- Ingres tools startup
- Inconsistent database or recovery
- Operating system performance

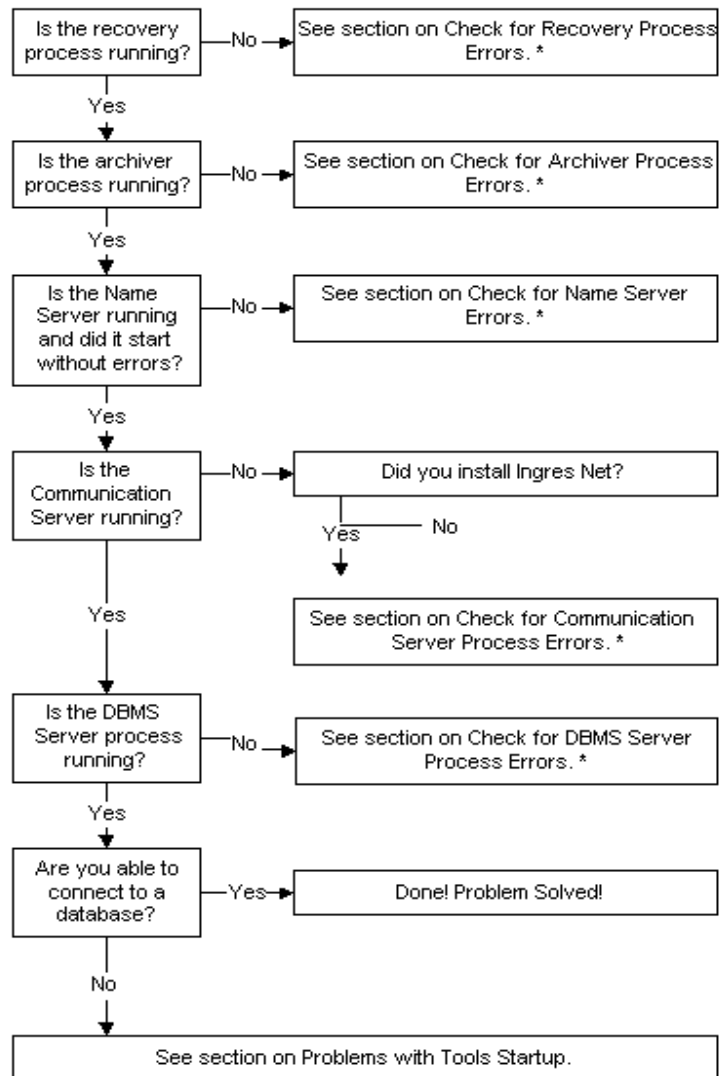
For a description of problems relating to queries and performance, see the *Database Administrator Guide*.

2. Use the information in the error logs to determine which category to check. Error messages always include the first two letters of the facility code that generated the error.

## Troubleshoot Startup, Shutdown, or Configuration Problems

Use the following flow chart to isolate a problem with startup, shutdown or configuration of your Ingres installation:





\* Note: Refer to the sections specific to the platform Ingres is running on.

## Check Ingres Installation on Windows

To check if the Ingres installation is working fine, follow these steps:

1. Check that you are logged in as the installation owner. If not, log off and log in again as this user.
2. Check that all users have II\_SYSTEM set by issuing the following command at the operating system prompt:

```
echo %II_SYSTEM%
```

All users must have Ingres executables in their path variables. Check that everyone has the full search path to %II\_SYSTEM%\ingres\bin.

The installation owner must also include %II\_SYSTEM%\ingres\utility.

3. Check that each of the Ingres variables has a valid value.

Ingres environment variables are only used and "seen" by Ingres and can be displayed with the following command entered at the operating system prompt:

```
ingprens
```

If you are in doubt about the function or legal value of a particular environment variable, see the chapter "Setting Environment Variables and Logicals" and the appendix "Environment Variables and Logicals."

Ingres environment variables denoting installation locations (II\_DATABASE, II\_CHECKPOINT, II\_DUMP, II\_JOURNAL, II\_LOG\_FILE, II\_SYSTEM) cannot be reset. To change these, you must unload the databases, rerun the installation program, and reload the databases, as described in Installation Location (see page 87).

4. Check the Ingres environment variables that have been set locally, overriding the Ingres installation-level definitions.

Only a small category of Ingres environment variables must be defined in the local user environment: those that permit you to access Ingres, and those that define values that are different for your local environment. They include TERM\_INGRES and ING\_EDIT.

If you trace the problem to an Ingres environment variable setting, correct the value. For procedures and scenarios for setting these environment variables, see the chapter "Setting Environment Variables and Logicals." If the installation does not start up, continue with this procedure.

5. Check that all Ingres processes are shut down. If there are processes that continue to run, see the section Check Shutdown Problems on Windows (see page 122).
6. Restart Ingres. Attempt once again to start up the installation by issuing the following command at the operating system prompt:

```
ingstart
```

If startup problems persist, continue diagnostics in the `ingstart` section below.

## Detect Ingres Startup Problems on Windows

To troubleshoot startup problems on Windows, follow these steps:

1. Determine which Ingres processes are running.
2. Verify that all required Ingres system processes are running. The following processes are the minimum required for a complete installation:

**iigcn**

Name Server process

**iigcc**

Communication Server process (present only on sites with Ingres Net)

**II\_IUSV\_###**

Recovery Server process

**dmfacp**

Archiver process

**iidbms**

DBMS Server process

**iigcd**

Data Access Server process

**iistar**

Distributed Database Server process (present only on sites with Ingres Distributed Option)

**icesvr**

ICE Server process

**rmcmd**

Remote command process

**iigcb**

Bridge Server process (present only on sites with Ingres Bridge)

3. If `ingstart` does not complete successfully, do the following:
  - a. Check the output from `ingstart`. This is saved in the file:  

```
%II_SYSTEM%\ingres\files\ingstart.log
```

Error log files and individual component log files are listed in Logging and Locking Systems (see page 33).
  - b. Try to determine the reason for startup failure.
  - c. A process failed to start. If a process failed to start, continue on to the detailed sections on startup problems for that specific process.

## Check Ingres Installation on UNIX

To check if the Ingres installation is working fine, follow these steps:

1. Check that you are logged in as the installation owner by issuing the following command at the operating system prompt:

```
whoami
```

If the user ID of the installation owner is not shown, log off and log in again as this user.

2. Check that all users have `II_SYSTEM` set by issuing the following command at the operating system prompt:

```
echo $II_SYSTEM
```

```
/usr/r6 (this varies by system)
```

All users must have Ingres executables in their path variables. Check that everyone has the full search path to `$II_SYSTEM/ingres/bin`.

The installation owner must also include `$II_SYSTEM/ingres/utility`.

3. Check that each of the Ingres installation variables has a valid value.

Ingres environment variables are only used and "seen" by Ingres and can be displayed with the following command entered at the operating system prompt:

```
ingprens
```

If you are in doubt about the function or legal value of a particular environment variable, see the chapter "Setting Environment Variables and Logicals" and the appendix "Environment Variables and Logicals."

Ingres environment variables denoting installation locations (`II_DATABASE`, `II_CHECKPOINT`, `II_DUMP`, `II_JOURNAL`, `II_LOG_FILE`, `II_SYSTEM`) cannot be reset. To change these, you must rerun the installation program, `ingbuild`, and on UNIX possibly unload and reload your database with `unloaddb`. More information is provided (see page 87)ed in Installation Location.



4. Check the Ingres environment variables that have been set locally, overriding the Ingres installation-level definitions. Issue the following commands at the operating system prompt:

BSD:

```
printenv | grep II  
printenv | grep ING
```

System V:

```
env | grep II  
env | grep ING
```

Only a small category of Ingres environment variables must be defined in the local user environment: those that permit you to access Ingres, and those that define values that are different for your local environment. They include TERM\_INGRES and ING\_EDIT.

If you trace the problem to an Ingres environment variable setting, correct the value. For procedures and scenarios for setting these environment variables, see "Setting Environment Variables and Logicals." If the installation does not start up, continue with this procedure.

5. Identify your installation code. If there is more than one Ingres installation on this machine, type the following command at the operating system prompt. The installation code is used to distinguish which processes belong to which installation at sites with more than one Ingres installation on the same machine:

```
ingpreenv | grep II_INSTALLATION
```

The two-letter installation code is displayed (for example, the following code R6):

```
II_INSTALLATION=R6
```

Take note of your installation code: \_\_\_\_\_.

6. Check that all Ingres processes are shut down. If there are processes that continue to run, see Check Shutdown Problems on UNIX (see page 124).
7. Restart Ingres: Attempt once again to start up the installation by issuing the following command at the operating system prompt:

```
ingstart
```

8. If startup problems persist, continue the diagnostics described in Ingbuild on UNIX (see page 114) or Detect Ingres Startup Problems on UNIX (see page 115).

## Ingbuild on UNIX

The executable script `ingbuild` performs all the steps necessary to set up an installation. It checks system resources, installs shared memory and semaphores, configures DBMS server parameters, configures the logging and locking system, and starts all the required processes.

The `ingbuild` program is located in `$II_SYSTEM/ingres/utility`. It makes use of numerous shell commands as well as the following Ingres binary and shell executables:

- `createdb`
- `iilink`
- `ingstop`
- `ingstart`
- `ingprenv`
- `ingunset`
- `sql`

One of the last things `ingbuild` does is call the `ingstart` script to actually start installation processes. When `ingstart` is called, it displays the message "Starting the Name Server process (`iigcn`).". If there are startup problems after this message has displayed, see `Ingstart on UNIX` (see page 115).

Before you can diagnose a problem with `ingbuild`, you must identify which subroutine is failing. If you know which routine is failing and it is `ingstart` or one of the main installation processes (`iigcn`, `iigcc`, `II_IUSV_nnn`, `dmfacp` or `iidbms`), see the section below that addresses that executable.

Details on tracing are described in `Bourne Shell -x Option` (see page 221).

## Detect Ingres Startup Problems on UNIX

To diagnose Ingres problems, use the following procedure.

1. Display which processes are running by using the `csreport` and operating system `ps` commands.

The `csreport` utility is described in *Operating System Utilities* (see page 185) and the `ps` command is described in *UNIX Operating System Utilities* (see page 187).

2. Verify that all required Ingres system processes are running. The following processes (in the order they are started) are the minimum required for a complete installation:

### **iigcn**

Name Server process

### **iigcc**

Communications Server process (present only on sites with Ingres Net)

### **iidbms (II\_IUSV\_\*\*\*)**

Recovery Server process

### **iigcd**

Data Access Server process

### **dmfacp**

Archiver process

### **iidbms**

DBMS Server process

### **iislave**

DBMS Server asynchronous write daemons—these processes are created at the time of the first database write operation (only present on sites where Ingres does not use OS threads or asynchronous I/O)

### **iistar**

Distributed Database Server process (present only on sites with Ingres Distributed Option)

### **icesvr**

ICE Server process

### **iigcb**

Bridge Server process (present only on sites with Ingres Bridge)

**Note:** If the command `ingprenv | grep II_CLIENT` shows “`II_CLIENT = true`”, you need to run only the Name Server and Communications Server processes.

In a few special cases Ingres can be run without the Communications Server process. This includes releases that use the streams device driver for interprocess communications. If you are not sure if your release uses the streams device driver, see your Readme file.

3. If `ingstart` does not complete successfully, try to identify the reason for startup failure. For example:

The problem is with `ingstart`. The `ingstart` script fails due to results of the checks it makes for sufficient resources and installation settings. If this is the reason for startup failure, correct the deficiency.

A process failed to start. If a process failed to start, continue on to the details sections on startup problems for that specific process.

## Check Ingres Installation on VMS

It is recommended that you first make the following basic installation checks:

1. Check that you are logged in as the ingres system administrator (ISA) user by issuing the following command at the operating system prompt:

```
show process
```

If the ISA is not shown, log off and log in again as the ISA. The ISA is typically `ingres` if there is only a single installation per machine. For multiple installations, there is a different ISA for each installation.

2. Check that all users have `II_SYSTEM` set by issuing the following command at the operating system prompt:

```
show logical II_SYSTEM
```

3. Check that each of the Ingres installation logicals has a valid value.

Ingres symbols are defined in the following files:

- `II_SYSTEM:[INGRES.UTILITY]INGSYSDEF.COM` - Defines Ingres system administration, DBA, and user symbols.
- `II_SYSTEM:[INGRES.UTILITY]INGDBADEF.COM` - Defines Ingres DBA and user symbols.
- `II_SYSTEM:[INGRES.UTILITY]INGUSRDEF.COM` - Defines Ingres user symbols.

If you are in doubt about the function or legal value of a particular logical, see the chapter “Setting Environment Variables and Logicals” and the appendix “Environment Variables and Logicals”.

Ingres logicals denoting installation locations (II\_DATABASE, II\_CHECKPOINT, II\_DUMP, II\_JOURNAL, II\_LOG\_FILE, II\_SYSTEM) can not be reset. To change these, you must edit II\_CONFIG:config.dat, update the existing definition, and execute @ii\_system: [ingres]ingsysdef.com. For additional details, see Installation Locations (see page 87).

4. Check the Ingres logicals that have been set locally, overriding the Ingres installation-level definitions. Issue the following command at the operating system prompt:

```
show logical *II*, *ING*
```

This shows which ones have been set at process, job, or group level, overriding the system-level definitions.

Only a small category of Ingres logicals are defined in the local user environment: those that permit you to access Ingres, and those that define values that are different for your local environment. They can include TERM\_INGRES and ING\_EDIT.

If you trace the problem to an Ingres logical setting, correct the value. For procedures and scenarios for setting these logicals, see the chapter "Setting Environment Variables and Logicals." If the installation does not start up, continue with this procedure.

5. Identify your installation code if there is more than one Ingres installation on this machine. To do so, type the following command at the operating system prompt:

```
show logical II_INSTALLATION
```

The two-letter installation code is displayed. For example:  
II\_INSTALLATION=R6

If this is a system-wide Ingres installation, there is no definition for II\_INSTALLATION. The installation code is used to distinguish which processes belong to which installation at sites with more than one Ingres installation on the same machine:

6. Check that all Ingres processes are shut down. If processes continue to run, see Check Shutdown Problems on VMS (see page 127).
7. Restart Ingres. Attempt once again to start up the installation by issuing the following command at the operating system prompt:

```
ingstart
```

8. If startup problems persist, continue diagnostics contained in the next section—vmsinstal or ingstart.

## VMSINSTAL on VMS

The executable script `vmsinstal` performs all the steps necessary to set up an installation. It checks system resources, installs shared images, configures DBMS server parameters, configures the logging and locking system, and starts all the required processes.

During the execution of the `vmsinstal` script, if a problem occurs, do the following:

1. Execute the `vmsinstal` program with the **L** option to log the contents of the VMSINSTAL session:

```
@sys$update:vmsinstal * device_name options L
```

2. Examine the DCL by issuing the following command at the operating system prompt:

```
set verify
```

3. Examine the `vmsinstal` log for error notifications and attempt to resolve the problem.
4. If the Installation Verification Procedure (IVP) for the DBMS Server fails with the message "ii.node\_name.syscheck.gblpages not found in config.dat" after you enter the number of concurrent users that the installation supports, check the value of the account parameter PRCLM (process limit) to ensure that it is not zero. The PRCLM must be a minimum of 15.
5. If the problem persists, contact technical support, as described in What You Need Before Contacting Technical Support (see page 179).

## Detect Ingres Startup Problems on VMS

If you are having `ingstart` problems, use the following procedure:

1. Display the processes that are running by issuing the following command at the operating system prompt:

```
show system /out = filename
```

```
search filename "II_", "DMF"
```

The search command searches the user-specified *filename* for the process names.

2. Verify that all required Ingres system processes are running. The following processes (in the order they are started) are the minimum required for a complete installation:
  - II\_IUSV\_*nn\_ppp*—Recovery process
  - DMFACP\_*nn*—Archiver process

- II\_GC $N$ \_nn—Name Server process
- II\_GC $C$ \_nn\_pppp—Communications Server process (present only on sites with Ingres Net)
- II\_DBMS\_nn\_pppp—DBMS Server process
- II\_STAR\_nn\_pppp—Distributed Database Server process (present only on sites with Ingres Distributed Option)
- RMCMD\_nn\_ppp—Remote Command Server for use with VDBA
- II\_GCB\_nn\_pppp—Bridge Server process (present only on sites with Ingres Bridge)

Where *nn* is the optional two-letter installation code, and *pppp* is the lower two bytes of the process' process ID.

**Note:** If you are checking the processes on a client node, only the Name Server and Communications Server processes are displayed when issuing the \$ SHOW SYSTEM command.

3. If ingstart does not complete successfully, do the following:
  - a. Check for error messages in the errlog.log:
 

```
II_SYSTEM: [INGRES.FILES] II_CONFIG:errlog.log
```

For a complete list of error log files and individual component log files, see Logging and Locking Systems (see page 33).
  - b. Try to identify the reason for startup failure.
    - The problem is with ingstart. If the ingstart executable fails due to results of the checks it makes for sufficient resources and installation settings, correct the deficiency.
    - A process failed to start. If a process failed to start, continue on to the detail sections on startup problems for that specific process.

## Detect Ingres Startup Problems on System Reboot (Windows)

To detect Ingres startup problems on system reboot on Windows, follow these steps:

1. Check that the "Startup Type" for the "Ingres Intelligent Database" service has been set to "automatic" in the services dialog.
2. Check that the password specified for the "Ingres Intelligent Database" service matches the installation owner user password.
3. Start the service manually.
4. If the service fails to start, run ingstart from the command prompt.
5. If the installation still does not start, contact technical support, as described in What You Need Before Contacting Technical Support (see page 179).

## Detect Ingres Startup Problems on System Reboot (UNIX)

To detect Ingres startup problems on system reboot on UNIX, follow these steps:

1. The most common cause of startup failure following a reboot is failure to include the startup command `ingstart` in the boot script for your machine. (The boot file is vendor-specific but can be named `"/etc/rc"` or `"/etc/rc.local"`.) This file contains the commands that are to be executed immediately after a reboot.

Make sure that the following line appears in the boot script:

```
su userid -c "ii_system/ingres/utility/ingstart ii_system" /dev/console
```

where:

*userid* refers to the user that owns the installation

*ii\_system* refers to the value of `II_SYSTEM` for your installation.

2. Make sure that `/dev/kmem` is readable to the user that owns the installation. If this is a security problem for your machine, you can add this user as a member of `/dev/kmem`'s group and make the `/dev/kmem` group readable.

Issue the following command at the operating system prompt:

```
chmod g+w /dev/kmem
```

The user that owns the installation *must* be able to read `/dev/kmem` or the kernel resource checks in `ingstart` fails.

3. Run `ingstart`. If the installation still does not start, contact technical support, as described in [What You Need Before Contacting Technical Support](#) (see page 179).



## Detect Ingres Startup Problems on System Reboot (VMS)

To detect Ingres startup problems on system reboot on VMS, follow these steps:

1. Make sure that a real error has occurred. If Ingres is performing startup recovery, see Automatic Recovery (see page 157).
2. Check to see that the Ingres initialization code is called from the VMS startup files. There are two main startup files:

- **SYS\$MANAGER:SYLOGIN.COM**

SYLOGIN.COM is run when any user logs in. The script to set up Ingres user symbols are defined here.

```
@II_SYSTEM: [INGRES]ingusrdef.com
```

If the DBA symbols are also to be known by all users, the ingdbadef.com script can be used instead.

If the system administrator symbols are also to be known by all users, the ingsysdef.com script can be used instead.

Only one script needs to be defined.

An alternative method is to include these in a user's particular login.com.

- **SYS\$MANAGER:SYSTARTUP\_VMS.COM**

SYSTARTUP\_VMS.COM is run when the system is booted. To have Ingres restarted automatically on reboot, add the following line:

```
@SYS$STARTUP: INGRES_STARTUP.COM
```

Edit the file SYS\$STARTUP: INGRES\_STARTUP.COM to ensure it is correct for your environment. A line is added to start *each* installation, but the line is commented out to prevent accidental errors.

Edit this file to meet your needs.

3. Be sure that you have examined all of the log files. Look in these files for possible error messages:

- II\_CONFIG:ERRLOG.LOG
- II\_CONFIG:IIACP.LOG (with *\_nodename* name suffix if configured for clusters)
- II\_CONFIG:IIIRCP.LOG (with *\_nodename* name suffix if configured for clusters)
- II\_CONFIG:II\_CSP.LOG\_*nodename* (cluster installations only)

4. Determine if any system logicals have been changed interactively since the last startup of Ingres but the command files have not been updated (use commands such as define/system/exec). Logical names become undefined when VMS is shut down and must be redefined (using the command files) at each boot.
5. Make sure that DECnet and/or TCP/IP is started before either Ingres Cluster Option or Ingres Net is started.
6. Execute the ingstart utility to start up Ingres.

If the installation still does not start, contact technical support, as described in What You Need Before Contacting Technical Support.

## Check Shutdown Problems on Windows

If Ingres has problems during shutdown on Windows, follow these steps:

1. Check environment variables in the local user environment by entering the following commands at the operating system prompt:

- a. Verify that II\_SYSTEM is set correctly.

```
echo %II_SYSTEM%
```

- b. Check that you have the full search path to %II\_SYSTEM%\ingres\bin, and %II\_SYSTEM%\ingres\utility

```
echo %PATH%
```

If you are having problems shutting down a client installation because Ingres believes a Communications Server is running locally when there is none, remove this file:

```
%II_SYSTEM%\ingres\files\name\clientname\IICOMSVR_clientname
```

2. Identify whether Ingres is recovering aborted transactions. Details of recovery delays are described in Recovery Process Monitoring (see page 159). Issue the following command at the operating system prompt and examine the output for the word "RECOVER."

```
logstat | findstr RECOVER
```

If Ingres is recovering aborted transactions, wait for this process to finish. Continue reissuing the logstat command and examining the STATUS field. When it reads "ONLINE, ECP DONE," proceed with normal shutdown.

3. Shut down the installation by typing the following command:

```
ingstop
```

You can also click the Stop toolbar button in Ingres Visual Manager to shut down the installation.

**Note:** All users must be logged out of Ingres (that is, no sessions running in the server) for `ingstop` to succeed.

4. Check that all processes are shut down. If shutdown succeeds, none of the following processes are running on your machine:

- `iigcn`
- `iigcc`
- `II_IUSV_***`
- `dmfacp`
- `iidbms`
- `ijdbc`
- `iigcd`
- `iistar`
- `icesvr`
- `iigcb`
- `rmcmd`

5. Use one of your operating system utilities to kill any of these processes that are still running.

6. Remove any shared memory segments that remain allocated to this installation. Execute the following command at the operating system prompt:

```
ipclean
```

Your installation is now shut down.

**Note:** For information about shutting down an Ingres installation that includes an ICE Server, see the chapter "Managing and Monitoring Web Deployment Option" in the *Web Deployment Option User Guide*.

## Check Shutdown Problems on UNIX

If Ingres has problems during shutdown on UNIX, follow these steps:

1. Check environment variables in the local user environment by entering the following commands at the operating system prompt.

- a. Verify that `II_SYSTEM` is set correctly:

```
echo $II_SYSTEM
```

`/usr/r6` (this varies system by system)

- b. Check that you have the full search path to `$II_SYSTEM/ingres/bin` and `$II_SYSTEM/ingres/utility`

```
echo $PATH
```

2. If there is more than one Ingres installation on this machine, identify your installation code by typing the following command at the operating system prompt:

```
ingprenv | grep II_INSTALLATION
```

The two-letter installation code is displayed (for example, the code here is JB):

```
II_INSTALLATION=JB
```

Take note of your installation code: \_\_\_\_\_.

3. Identify if you are shutting down a client installation or a full installation by issuing the following command at the operating system prompt:

```
ingprenv | grep II_CLIENT
```

If this displays "`II_CLIENT=true`", this is a client installation, and only two processes (`iigcn`, `iigcc`) are running on this node.

4. If you are having trouble shutting down a client installation because Ingres believes a Communications server is running locally when there is none, remove this file:

```
$II_SYSTEM/ingres/files/name/clientname/IICOMSVR_clientname
```

For troubleshooting details see the *Connectivity Guide*.

5. Identify whether Ingres is recovering aborted transactions. Details on recovery delays are described in Recovery Process Monitoring (see page 159). Issue the following command at the operating system prompt and examine the output for the word "RECOVER":

```
logstat | grep RECOVER
```

If Ingres is recovering aborted transactions, wait for this process to finish. Continue reissuing the `logstat` command and examining the `STATUS` field. When it says: "ONLINE, ECP DONE," proceed with normal shutdown.

6. Now shut down the installation:

```
ingstop
```

**Note:** All users must be logged out of Ingres (that is, no sessions running in the server) for the `ingstop` script to succeed.

7. Check that all processes are shut down:

- a. Display running processes by issuing the following command at the operating system prompt:

BSD:

```
ps -aux | grep ingres
```

System V:

```
ps -ef | grep ingres
```

- b. If shutdown succeeds, none of the following processes are running:

- `iigcn`
- `iigcc`
- `II_IUSV_nnn`
- `dmfacp`
- `iidbms`
- `ijdbc`
- `iigcd`
- `iistar`
- `icesvr`
- `rmcmd`
- `iigcb`

8. If any of these installation processes are not shut down, note the process ID of the running processes and do the following:

- a. Shut them down manually with the operating system command:

```
kill- QUIT process_id
```

where:

*process\_id* refers to the process ID of the process to stop.

- b. If processes are still not shut down, issue the operating system command:

```
kill -9 process_id
```

**Important!** *If your site has more than one Ingres installation, examine the installation code associated with the iidbms process to make sure you are stopping only the processes associated with the installation you need to shut down.*

9. Check that no shared memory segments remain allocated to this installation. Execute the operating system command:

```
csreport
```

The following message indicates that shared segments have been properly removed:

```
!Can't map system segment
```

10. If shared memory segments remain for this installation, interactively remove them:
  - a. To deallocate shared memory resources, issue the operating system command:

```
ipcclean
```

- b. Use the UNIX command `ipcs` to verify that the actual segments have been removed:

```
ipcs
```

- c. If they have not been properly removed, you must delete them manually.

**Important!** *If your site has more than one Ingres installation, you must take care to only remove shared memory or semaphores for the installation to be shut down. If your machine contains more than one installation, enter the following command from the environment of the installation you need to target:*

```
csreport
```

The **csreport** utility displays the shared memory and semaphore segment identifiers for this installation.

To remove the targeted segment(s), use the UNIX command:

```
ipcrm -mid
```

or

```
ipcrm -ssid
```

where *mid* is the shared memory segment identifier and *sid* is the semaphore identifier.

11. Verify that the following files are *not* present. If they are still present, you must remove them:

```
$II_SYSTEM/ingres/files/memory/lockseg.mem
```

```
$II_SYSTEM/ingres/files/memory/sysseg.mem
```

Your installation is now shut down. Instructions on how to check and restart Ingres are described in Check Ingres Installation on UNIX (see page 112).

**Note:** For information about shutting down an Ingres installation that includes an ICE Server, see the chapter "Managing and Monitoring Web Deployment Option" in the *Web Deployment Option User Guide*.

## Check Shutdown Problems on VMS

If Ingres has problems during shutdown on VMS, follow these steps:

1. Check that you are logged in as the system administrator or another privileged account by issuing the following command at the operating system prompt:

```
show process
```

Make sure you have the WORLD and CMKRNL privileges required for stopping Ingres processes by issuing the following command at the operating system prompt:

```
show process/privilege
```

A privileged account must not only have the aforementioned VMS privileges but also must have the SERVER\_CONTROL privilege authorized in II\_CONFIG:CONFIG.DAT.

2. Check that the logical II\_SYSTEM is defined by issuing the following command at the operating system prompt:

```
show logical II_SYSTEM
```

3. Check to see that users are logged out of Ingres. You can use the utility logstat to determine which databases are open and active and find out whom to notify of the impending shutdown. For additional information, see the Logstat section in the *Command Reference Guide*.

To obtain a list of running processes, use the DCL command:

```
show system
```

4. This command gives you a complete list of all processes that are running on your machine. It displays the process ID ('Pid'), Process Name, and other information. The main processes have the names as shown in Detect Ingres Startup Problems on VMS (see page 118).

If you use the form show system/full, you can also see the owner of the process. This is helpful if you have group level installations.

5. Identify your installation code if there is more than one Ingres installation on this machine.

For a group level installation the two-lettered code is displayed after the process name. In this example, the code is BE:

```
Pid      Process Name    State Pri I/O   CPU          Page flts Ph.Mem
236012AD II_DBMS_BE_1E3F HIB   6   86463 0 00:13:53.24 14401    8192
```

**Important!** *If your site has more than one Ingres installation, examine the installation code associated with the II\_DBMS\_nnn process to make sure you are only shutting down processes associated with the installation you need to shut down.*

6. Use the iinamu utility to find the names of the servers for this installation; iinamu shows you only the servers for the installation in which you currently reside. For example:

```
IINAMU> show ingres
INGRES *      II_DBMS_BE_1E3F
found.
IINAMU> show comsvr
COMSVR *      II_GCC_BE_2B77
```

7. Shut down servers. See the individual descriptions below.

**Note:** Ingres can be shut down in several ways. Following are some tips on how each shutdown procedure works to enable you to make a choice as to how to shut down the installation.

- The ingstop utility shuts down servers in the following order:
  - Remote Command Server (rmcmd)
  - Star Server
  - DBMS Server
  - Protocol Bridge Server
  - Communications Server
  - Recovery and archiver processes-IUSV and ACP by rcp.config.
  - If present, the Cluster Server is also shut down.
  - Name Server
- The following order also works. This ordering manually shuts down the installation, shutting down the Name Server first:
  - Name Server
  - Communications Server
  - Star Server



- DBMS Server
  - Recovery and archiver processes-IUSV and ACP
- a. Shutting down the Name Server

When you issue the stop command from iinamu, the Name Server process is stopped. Once this command is issued, no new servers are allowed to register with the Name Server. Therefore, it makes sense to take the Name Server down first for a clean shutdown. Current, active sessions are not affected by this command.
  - b. Shutting down the DBMS Server

Use iimonitor to stop the DBMS Server. The two commands you can use are:

```
set server shut
```

This causes the server to refuse further connections and shuts the server down gracefully, completing all current transactions first, including the iimonitor session.

```
stop server
```

This stops the server immediately without waiting for current connected sessions to complete. Use this command instead of the VMS command stop/id.
  - c. Shutting down the recovery and archiver processes

Take down the recovery and the archiver processes, either quickly or slowly, depending on the method you used to take down the DBMS Server in step b. Use either of the following commands at the operating system prompt:

```
rcpconfig/shutdown
```

This causes the processes to refuse all further connections and transaction processing, but allows current transactions to finish. This allows the archiver and recovery processes to shut down cleanly.

```
rcpconfig/imm_shutdown
```

This stops all pending transactions and immediately shuts down the recovery and archiver processes. This leaves open transactions in the log file, which requires rollback (and slow startup) when the system is restarted.
  - d. Writing a DCL command file to take Ingres down

Alternatively, you can write a DCL script that puts the above steps in a command procedure to shut down each of the Ingres processes.

8. After shutting down Ingres, check that all processes are stopped:

a. Issue the following command at the operating system prompt:

```
show system
```

Check to be sure that none of the processes listed in the Ingstart on VMS section is shown.

b. If any of these installation processes are not shut down, note the process ID (PID#) of the running process and issue either of the following commands at the operating system prompt:

```
stop process_name
```

or

```
stop/id = pid#
```

**Note:** For serious system problems that require immediate attention, the VMS stop command can be used to stop an Ingres process. This is not recommended because it does not permit the recovery process to back out open transactions or the archiver process to flush the log file to disk.

If this is a cluster installation, remember that these Ingres processes run on each node.

Your Ingres installation is now shut down. Information about checking and restarting Ingres is described in Check Ingres Installation on VMS (see page 116).

**Note:** For information about shutting down an Ingres installation that includes an ICE Server, see the chapter "Managing and Monitoring Web Deployment Option" in the Web Deployment Option User Guide.

**Note:** For information about shutting down an Ingres installation that includes an ICE Server, see the chapter "Managing and Monitoring Web Deployment Option" in the *Web Deployment Option User Guide*.

## Ingres Processes on Windows

The major Ingres processes are as follows:

- Name Server process (iigcn)
- Communications Server process (iigcc)
- Recovery process (II\_IUSV\_ *nnn*)
- Archiver process (dmfacp)
- DBMS Server process (iidbms)
- JDBC Server process (ijjdbc)
- Data Access Server process (iigcd)
- Bridge Server process (iigcb)
- ICE Server process (icesvr)
- Remote Command process (rmcmd)

### Check for Name Server Errors on Windows

The Name Server process (iigcn) is not running if either of the following occurs:

- You receive a specific error indicating that the Name Server failed to start.
- The iigcn.exe does not appear in the process list.

You can verify this by attempting to start the Name Server manually.

The most common causes of Name Server startup failure are:

- An improper installation environment exists. Be sure to verify the installation environment by performing the Check Ingres Installation on Windows (see page 110) before attempting to start the Name Server.
- The Name Server fails to start if there is already a Name Server running for this installation.

If you still have problems, set the following trace to capture additional diagnostic data before calling technical support:

```
ingsetenv II_GC_TRACE 4
ingsetenv II_GC_LOG filename
iirun iigcn
```

## Check for Communications Server Process Errors on Windows

If the Communications Server process (iigcc) does not start, verify that you have installed the Ingres Net component and that you have installed the correct protocol drivers.

If you are still having problems, set the following trace to capture additional diagnostic data before calling technical support:

```
ingsetenv II_GC_TRACE 4
ingsetenv II_GC_LOG filename
iirun iigcc
```

## Check for Bridge Server Errors on Windows

If the Bridge Server process (iigcb) does not start, verify that you have installed the Ingres Net component and that you have installed the correct protocol drivers.

If you are still having problems, set the following trace to capture additional diagnostic data before calling technical support:

```
ingsetenv II_GC_TRACE 4
ingsetenv II_GC_LOG filename
ingstart -iigcb
```

## Check for ICE Server Errors on Windows

The command to start the ICE Server process (icesvr) is `ingstart -icesvr`. To check for errors in ICE Server, verify that you have installed the Ingres Web Deployment Option component, and that you have configured it correctly using Configuration Manager. Also, that the DBMS Server process `iidbms` and/or Ingres Net are running.

For additional information, see *Ingres Processes on Windows* (see page 131).

Issue the following command at the operating system prompt to start the ICE Server:

```
ingstart -icesvr
```

## Check for Remote Command Process Errors on Windows

You can check for the presence of the Remote Command process (`rmcmd`) using the Windows Task Manager. View the Processes tab for the existence of the process.

## Check for Recovery Process Errors on Windows

The Recovery process (II\_IUSV\_*nnn*) must be running before a DBMS Server can be started. Failure of the II\_IUSV\_*nnn* starting process indicates one of the following:

- Improper installation configuration, as described in Problems with Startup, Shutdown, or Configuration.
- Problems with the log file
- Insufficient shared memory

If the Recovery process does not start, perform the following procedure:

1. Check for the existence of the transaction log by opening the Primary Transaction Log window in Configuration Manager (or the Transaction Log screen in the Configuration-By-Forms (cbf) utility), and noting the directories listed in the Log File Root Locations table in Configuration Manager (or the Primary Transaction Log Locations table in cbf).
2. In the ingres/log directory (located below all other listed directories), look for the file `ingres_log.lnn` where *nn* is an integer between 1 and 16.
3. Make the following checks on the transaction log file `ingres_log.lnn`:
  - a. Verify that `ingres_log.lnn` exists at that location by entering the following command at the operating system prompt:

```
dir
```
  - b. If it does not exist, you must create it using the `iimklog` command, and format it using the `rcpconfig -init` command.
4. If you corrected a recovery process problem, verify that Ingres starts normally by completing the following steps:
  - a. Shut down the partially started installation with the `ingstop` command.
  - b. Restart the installation with the `ingstart` command.

If you still cannot start the II\_IUSV\_*nnn*, you must erase it and recreate and reformat it. This step is done only as a last resort.

**Note:** This step reinitializes the log file, causing any outstanding transactions to be lost.

## Check for Archiver Process Errors on Windows

The archiver process (dmfacp) does not start unless the II\_IUSV\_### recovery process is running. However, an Ingres installation runs without an archiver process until the log file fills up. The user programs are suspended until any outstanding transactions in the log file are backed out. More information about Ingres in a recovery state is provided in Recovery Proces (see page 159)s Monitoring.

Archiver process (dmfacp) startup errors are likely to result from:

- Improperly shared memory resource
- Inability to read the transaction log file
- Inability to write journal files

Use the following procedure to check archiver process startup problems. Some of these checks are the same as for the recovery process in the previous section.

1. Make the following check on the transaction log file ingres.log:

Verify that ingres.log exists at that location by entering the following command at the operating system prompt:

```
dir
```

If it does not exist, you must create it using the iimklog command, and format it using the rcpcnfig –init command.

2. Check journal locations by verifying that:

The journal location name points to a valid directory containing subdirectories ingres\jnl\default\dbname.

3. Check that the disk partition containing the journal files is not full. Issue the following command from the operating system prompt:

```
dir
```

If the journal partition is 100% full, this makes it impossible to write journals and stalls the archiver. If this is the reason preventing archiver start up, you must either free space on the journal partition or temporarily disable journaling with the alterdb command. For details on alterdb, see the *Command Reference Guide*.

4. If you corrected a recovery process problem, verify that Ingres starts normally.
  - a. First, shut down the partially started installation with the ingstop command.
  - b. Restart the installation with the ingstart command.

## Check for DBMS Server Process Errors on Windows

The command to start the DBMS Server process (iibms) from the Ingres login is `ingstart -iibms`. If the DBMS Server did not start:

1. Verify that the recovery process `II_IUSV_nnn` is running, as described in [Check for Recovery Process Errors on Windows](#) (see page 133).
2. Verify that the recovery process `II_IUSV_nnn` is not in a recovery state, as described in [Recovery Process Monitoring](#) (see page 159). This is a likely situation if there was a sudden shutdown because of a power failure or some other system failure.
3. Try to start the DBMS Server. Issue the following command at the operating system prompt:  

```
ingstart -iibms
```
4. If you corrected a recovery process problem, verify that Ingres starts normally.
  - a. First, shut down the partially started installation with the `ingstop` command.
  - b. Restart the installation with the `ingstart` command.

## Check for JDBC Server Errors on Windows

If the JDBC server process (ijdbc) does not start, follow this procedure:

1. Verify that you have installed the Ingres Net component.
2. Verify that you have installed the correct protocol drivers.

If you are still having problems, set the following trace to capture additional diagnostic data before calling technical support:

```
ingsetenv II_GC_TRACE 4
ingsetenv II_GC_LOG filename
ingstart -ijdbc
```

## Check for Data Access Server Errors on Windows

If the Data Access Server process (iigcd) does not start, follow this procedure:

1. Verify that you have installed the Ingres Net component
2. Verify that you have installed the correct protocol drivers

If you are still having problems, set the following trace to capture additional diagnostic data before calling technical support:

```
ingsetenv II_GCD_TRACE
ingsetenv II_GCD_LOG filename
ingstart -iigcd
```

## Ingres Processes on UNIX

The major Ingres processes are as follows:

- Name Server process (iigcn)
- Communications Server process (iigcc)
- Recovery process (II\_IUSV\_ *nnn*)
- Archiver process (dmfacp)
- DBMS Server process (iidbms)
- JDBC Server process (ijjdbc)
- Data Access Server process (iigcd)
- Bridge Server process (iigcb)
- ICE Server process (icesvr)

Remote Command process (rmcmd)

## Name Server Errors on UNIX

The Name Server process (iigcn) is not running if either of the following occurs:

- You receive a specific error indicating that the Name Server process (iigcn) failed to start.
- The command `ps -aux` (BSD) or `ps -ef` (System V) shows that the iigcn is not running.

You can verify this by attempting to start the Name Server manually.



## Check for Name Server Errors on UNIX

If the Name Server does not start, follow these steps:

1. Verify that TCP/IP is properly installed by typing the following command at the operating system prompt:

```
telnet localhost
```

A loopback login to your machine occurs.

2. Verify that the required TCP daemon process for your operating system is running.

The specific process name is system dependent, but on many UNIX systems, the process is named "inetd" (use your process name in the command below if it is not inetd). Issue the following command at the operating system prompt, or see your operating system manual for your particular TCP/IP implementation:

BSD:

```
ps -aux | grep inetd
```

System V:

```
ps -ef | grep inetd
```

3. Check that the Ingres environment variable `II_GC_Nxx_PORT` is not set (this environment variable contains the TCP port identifier of the Name Server process):
  - a. Use the `ingprens` utility to verify that this environment variable is *not* set when the Name Server tries to start up.
  - b. If necessary, use the `ingunset` command to *unset* the `II_GC_Nxx_PORT` environment variable.
4. If you corrected a Name Server problem, verify that Ingres starts normally:
  - a. Shut down the partially started installation with the `ingstop` command.
  - b. Restart the installation with the `ingstart` command.
5. If you are still having problems, set the following trace to capture additional diagnostic data before calling technical support:

Bourne Shell:

```
II_GC_TRACE=5
```

```
II_GC_LOG = stdio (stdio or filename); export II_GC_TRACE II_GC_LOG
```

C Shell:

```
setenv II_GC_TRACE 5
```

```
setenv II_GC_LOG stdio (stdio or filename)
```

```
iirun iigcn
```

## Check for Communications Server Process Errors on UNIX

If the Communications Server process (iigcc) did not start, follow this procedure:

1. Verify that no local environment variables are set in the local user environment that contain the string "GCC". At the operating system prompt type the command:

BSD:

```
printenv | grep GCC
```

System V:

```
env | grep GCC
```

The Ingres environment variable II\_GCnxx\_PORT is set in the Ingres symbol table. It is *not* visible from the UNIX environment using the printenv or env commands, but is visible in the Ingres environment. Verify this by typing **ingprens**.

2. Check the value of the Ingres environment variable II\_RUN with the following command entered at the operating system prompt:

```
ingprens
```

- a. If this machine is running Ingres Net, the value of II\_RUN is either ",NET" or ",DBMS, NET". If the value of the Ingres environment variable II\_CLIENT is TRUE, II\_RUN is ",NET". If II\_RUN is not set correctly, reset it using the command:

```
ingsetenv II_RUN ', DBMS, NET'
```

- b. If this machine is an NFS client (that is, does not run a DBMS Server locally) the value of II\_RUN is ",NET". If II\_RUN is not set correctly, reset it using the following command entered at the operating system prompt:

```
ingsetenv II_RUN', NET'
```

For details on setting environment variables, see the chapter "Setting Environment Variables and Logicals."

3. If you are still having problems, set the following trace to capture diagnostic data and attempt to restart the Communications Server:

Bourne Shell:

```
II_GCC_TRACE=4
```

```
II_GCA_LOG=stdio
```

```
export II_GCC_TRACE II_GCA_LOG
```

```
iirun iigcc
```

C Shell:

```
setenv II_GCC_TRACE 4
setenv II_GCA_LOG stdio
iirun iigcc
```

4. If you corrected a Communications Server problem, verify that Ingres starts normally:
  - a. Shut down the partially started installation with the `ingstop` command.
  - b. Restart the installation with the `ingstart` command.

## Check for Bridge Server Process Errors on UNIX

If the Bridge Server process (`iigcb`) did not start, verify that you have installed the Ingres Net component and that you have installed the correct protocol drivers.

If you are still having problems, set the following trace to capture additional diagnostic data before calling technical support:

```
ingsetenv II_GC_TRACE 4
ingsetenv II_GC_LOG filename
ingstart -iigcb
```

## Check for ICE Server Process Errors on UNIX

The command to start the ICE Server process (`icesvr`) from the installation owner login is `ingstart -icesvr`. If the ICE Server did not start:

1. Verify that you have installed the Ingres Web Deployment Option component, and that you have configured it correctly using Configuration Manager (`vcbf`).
2. Verify that the DBMS Server process `iidbms` and Ingres Net are running, as described in [Check for DBMS Server Process Errors on UNIX](#) (see page 143)
3. Issue the following command at the operating system prompt to start the ICE Server:

```
instart -icesvr
```

## Recovery Process Errors on UNIX

The recovery process (dmfrcp) must be running before a DBMS server can be started. Failure of the dmfrcp starting process indicates one of the following:

- An improper installation configuration exists, as described in Troubleshoot Startup, Shutdown, or Configuration Problems (see page 108).
- Problems with the log file
- Insufficient or previously allocated system semaphores or shared memory
- If configured for Ingres Cluster Solution, check that required OS prerequisites are installed and operational (for example, OpenDLM in a Linux environment).

### Check for Recovery Process Errors on UNIX

If the recovery process does not start, perform the following procedure:

1. Check that the shared memory resources are properly installed. Use the `csreport` utility and check that the size, ownership, and permissions of the semaphore and shared memory segments meet the minimum requirements for your port. These are specified in the *Getting Started* guide.
2. Check for the existence of the transaction log by opening the Primary Transaction Log window in Configuration Manager (or the Transaction Log screen in the Configuration-By-Forms (cbf) utility), and noting the directories listed in the Log File Root Locations table in Configuration Manager (or the Primary Transaction Log Locations table in cbf).
3. Look for the file `ingres_log.lnn` (where `nn` is an integer between 1 and 16) in the `ingres/log` directory (located below all other listed directories).
4. Make the following checks on the transaction log file `ingres_log.lnn`:
  - a. Verify that `ingres_log.lnn` exists at that location by entering the following command at the operating system prompt:

```
ls -l
```

If it does not exist, you must recreate it using the Configuration-By-Forms (cbf) or Configuration Manager (vcbf) utility.

- b. Verify that `ingres_log.lnn` is owned by installation owner.

If it is not, issue the following command at the operating system prompt, where `userid` is the user ID of the installation owner:

```
chown userid ingres_log
```

- c. If the transaction log file was created as:

An ordinary UNIX file, make sure it has permissions 660 (that is, “-rw-rw----”). If not, issue the following command at the operating system prompt:

```
chmod 660 ingres_log
```

A raw log, permissions is “crw-----”.

Details on configuring a raw log file are in the *Getting Started* guide.

5. Verify that Ingres starts normally.
  - a. Shut down the partially started installation with the `ingstop` command.
  - b. Restart the installation with the `ingstart` command.
6. If you still cannot start the `II_IUSV_nnn` process, you need to completely shut down the Ingres installation, re-run `ingbuild`, and reconfigure the log file.

**Important!** *This step must only be done as a last resort. Keep in mind that this reinitializes the log file, and any outstanding transactions are lost.*

To shut down the system, complete the following steps:

- a. Issue the command **ingstop**.
- b. Check Shutdown Problems on UNIX (see page 124). Work through that section until your Ingres installation has been cleanly shut down.

## Check for Remote Command Process on UNIX

You can check for the presence of the remote command process by entering the following command at the operating system prompt:

```
ps -ef | grep rmcmd
```

## Archiver Process on UNIX

The archiver process (`dmfacp`) does not start unless the recovery (`dmfrcp`) process is running. However, an Ingres installation runs without an archiver process until the log file fills up. User programs are suspended as outstanding transactions in the log file are backed out. For information about the Ingres recovery state, see *Recovery Process Monitoring* (see page 159).

- Archiver process (`dmfacp`) startup errors are likely to result from:
- Improper shared memory resources
- Inability to read the transaction log file
- Inability to write journal files

## Check for Archiver Process Errors on UNIX

Use the following procedure to check archiver process startup problems. Some of these checks are the same as for the recovery process:

1. Check that the shared memory resources are properly installed: Use the `csreport` utility and check that the size, ownership and permissions of the semaphore and shared memory segments meet the minimum requirements for your port. For requirements, see the Readme file.

The command to allocate Ingres shared memory and semaphores (when logged is as the installation owner) is `csinstall`. You can display them from UNIX with the command `ipcs`.

2. Make the following checks on the Ingres transaction log file "ingres\_log":
  - a. Verify that "ingres\_log" exists at that location by entering the following command at the operating system prompt:

```
ls -l
```

If it does not exist, you must create it by running the `ingbuild` program.

- b. Verify that "ingres\_log" is owned by the user that owns the installation.

If not, issue the following command at the operating system prompt, where *userid* is the user who owns the installation:

```
chown userid ingres_log
```

- c. If the transaction log file was created as:

An ordinary UNIX file, make sure it has permissions 660 (that is, "-rw-rw----"). If not, issue the following command at the operating system prompt:

```
chmod 660 ingres_log
```

A raw log, permissions is "crw-----".

(Details on configuring a raw log file are in the *Getting Started* guide.)

Check that the `II_JOURNAL` location is a valid location. Issue the following command at the operating system prompt:

```
infodb | grep ii_journal
```

3. Check journal locations by verifying that:
  - a. The journal location name points to a valid directory containing subdirectories "ingres/jnl/default/dbname"
  - b. The permissions on these directories are:
    - 755 for the ingres directory
    - 777 for the default directory

4. Check that the disk partition containing the journal files is not 100% full. Issue the following command at the operating system prompt:

```
df
```

If a journal partition is 100% full, it is impossible to write journals and the Archiver stalls. If this is the reason preventing Archiver startup, you must either free space on the journal partition or temporarily disable journaling with the `alterdb` command. For details on `alterdb`, see the *Command Reference Guide*.

5. If you corrected an Archiver process problem, verify that Ingres starts normally:
  - a. Shut down the partially started installation with the `ingstop` command.
  - b. Restart the installation with the `ingstart` command.

## Check for DBMS Server Process Errors on UNIX

The command to start the DBMS Server process (`iidbms`) from the installation owner login is `ingstart -iidbms`. If the DBMS Server did not start:

1. Verify that the recovery process (`dmfrcp`) is running. Details are described in *Recovery Process Errors on UNIX* (see page 140).
2. Verify that the recovery process is not in a recovery state (see page 159). This is a likely situation if there was a sudden shutdown because of a power failure or some other systems failure.
3. Try to start up a DBMS server. Issue the following command at the operating system prompt:

```
ingstart -iidbms
```
4. If you corrected a DBMS server problem, verify that Ingres starts normally:
  - a. Shut down the partially started installation with the `ingstop` command.
5. Restart the installation with the `ingstart` command.

## Check for JDBC Server Process Errors on UNIX

The command to start the JDBC Server process (iijdbc) from the installation owner login is `ingstart -iijdbc`. If the JDBC server did not start:

1. Verify that the recovery process (dmfrcp) is running. More information is provided in [Recovery Process Errors on UNIX](#) (see page 140).
2. Verify that the recovery process (see page 159) is not in a recovery state. This is a likely situation if there was a sudden shutdown because of a power failure or some other systems failure.
3. Try to start up a JDBC server. Issue the following command at the operating system prompt:

```
ingstart - iijdbc
```

4. If you corrected a JDBC server problem, verify that Ingres starts normally:
  - a. Shut down the partially started installation with the `ingstop` command.
5. Restart the installation with the `ingstart` command.

## Check for Data Access Server Process on UNIX

The command to start the Data Access Server process (iigcd) from the installation owner login is `ingstart -iigcd`. If the Data Access Server did not start:

1. Verify that the recovery process (dmfrcp) is running. More information is provided in [Check for Recovery Process Errors on UNIX](#) (see page 140).
2. Verify that the recovery process (see page 159) not in a recovery state. This is a likely situation if there was a sudden shutdown because of a power failure or some other systems failure.
3. Try to start up a Data Access server. Issue the following command at the operating system prompt:

```
ingstart -iigcd
```

4. If you corrected a Data Access server problem, verify that Ingres starts normally:
  - a. Shut down the partially started installation with the `ingstop` command.
  - b. Restart the installation with the `ingstart` command.



## Ingres Processes on VMS

The major Ingres processes are as follows:

- Name Server process (iigcn)
- Communications Server process (iigcc)
- Recovery process (II\_IUSV\_*nnn*)
- Archiver process (dmfacp)
- DBMS Server process (iidxbms)
- JDBC Server process (iijdbc)
- Data Access Server process (iigcd)
- Bridge Server process (iigcb)
- ICE Server process (icesvr)

Remote Command process (rmcmd)

### Name Server Errors on VMS

The Name Server (II\_GCN) is not running if the following occurs:

- You receive a specific error indicating that the Name Server process failed to start.
- The VMS command `show system` does not list the II\_GCN process, indicating that the Name Server is not running.

**Note:** Only a privileged user can start up the Name Server.

### Check for Name Server Errors on VMS

If the Name Server does not start, follow these steps:

1. Examine the II\_CONFIG:errlog.log file for any additional information.
2. To isolate a running Name Server as the problem on failed local connections:
  - a. Define the logical II\_DBMS\_SERVER to point to the DBMS Server process name or mailbox.
  - b. If the connection now succeeds, stop and start the Name Server. For additional information about the `iinamu` command, see the *Command Reference Guide*.
3. If the Name Server continues to fail, contact technical support, as described in What You Need Before Contacting Technical Support (see page 179) .

## Communications Server Process on VMS

The Communications Server (II\_GCC\_*nnn*, where the *nnn* is a number related to the server's process ID) does not start if any of the following errors occur:

- The II\_CONFIG:errlog.log file has a notation regarding the failure of the "net server" (Communications Server) process starting.
- The show system command does not list an II\_GCC\_*nnn* process.
- The utility iinamu does not display a process name when you enter the show comsvr command at the IINAMU> prompt.

### Check for Communications Server Process Errors on VMS

If the Communications Server does not start, follow these steps:

1. Examine the II\_CONFIG:errlog.log file for any information regarding Communications Server startup.
2. Make sure that the network protocol software on which the Communications Server executes is started and functioning before attempting to start Ingres Net processes. By default, an attempt is made to open ports for the DECnet and TCP\_DEC protocols.
3. If the Communications Server continues to fail to start, contact Technical Support.

## Bridge Server Errors on VMS

The Bridge Server process (II\_GCB\_*nnn*, where the *nnn* is a number related to the server's process ID) does not start if the following errors occur:

- The II\_CONFIG:errlog.log file has a notation regarding the failure of the Bridge Server process.
- The show system command does not list a II\_GCB\_*nnn* process.
- The utility iinamu does not display a process name when you enter the show bridge command at the IINAMU> prompt.

## Check for Bridge Server Process Errors on VMS

If Bridge Server does not start, follow these steps:

1. Verify that you have installed the Ingres Net component.
2. Verify that you have installed the correct protocol drivers and they are functioning properly.
3. Verify the Bridge Server has been properly configured using the Configuration-By-Forms (cbf) utility. If you are still having problems, set the following trace to capture additional diagnostic data before calling technical support:

```
define II_GC_TRACE 4
define II_GC_LOG filename
ingstart -iigcb
```

## Recovery Process Errors on VMS

The recovery process (II\_IUSVxx\_nnn for stand-alone Ingres, or DMFRCPnnn-xx for Ingres Cluster Solution) must be running before a DBMS server can be started. Failure of the recovery process to start indicates one of the following:

- Improper installation configuration. Check Ingres Installation on VMS (see page 116).
- Problems with the log file
- Insufficient system memory (NPAGEDYN sysgen parameter)

## Check for Recovery Process Errors on VMS

If the recovery process does not start, perform the following procedure.

1. Examine the following files for specific reasons for the startup failure:

II\_SYSTEM: [INGRES]II\_CONFIG:errlog.log

II\_CONFIG: IIRCP.LOG

(If your system is configured for the Ingres Cluster Solution, the suffix *\_nodename* is added.)

2. If configured for Ingres Cluster Solution, check that DECnet is up among all nodes. Ingres uses DECnet to communicate among nodes. Check that you can set host and log into/from all nodes running Ingres.
3. Be sure that the log file location is correctly defined:

- a. Check the definition of the II\_LOG\_FILE logical.

- b. Check the II\_LOG\_FILE location for the existence of the log file:

II\_LOG\_FILE: [INGRES.LOG]INGRES\_LOG.SYSTEM\_*nodename*

4. If you still cannot start the recovery process, the problem may be a corrupted log file. Contact technical support, as described in What You Need Before Contacting Technical Support (see page 179).

## Remote Command Process Errors on VMS

The Remote Command Server (RMCMD\_ddd, where the ddd is a number related to the server's process id) does not start, if show system command does not list an RMCMD\_ddd process.

## Check for Remote Command Server on VMS

Use the following procedure if the RMCMD process did not start:

1. Ensure that the PRCLM for the system administrator is sufficient to start this additional process. This value must be a minimum of 15.
2. Check II\_SYSTEM: [ingres.bin]RMCMD.log for errors.

## Archiver Process Errors on VMS

The Archiver process (DMFACP or DCMFACPxx) does not start if any of the following conditions occur:

- Unless the recovery process is running; however, an installation runs without an archiver process until the log file fills up. User programs are suspended until the archiver is started.
- When the recovery system is backing out aborted transactions from a system failure. Recovery Process Monitoring (see page 159) can help you determine whether Ingres is in a recovery state.
- Due to improper memory resources, inability to read the transaction log file, or inability to write journal files.
- You receive an error indicating that the Archiver process failed to start.
- The VMS command show system does not list the Archiver process.

## Check for Archiver Process Errors on VMS

Use the following procedure to check Archiver process startup problems:

1. Examine the II\_CONFIG:errlog.log startup log file for information.
2. Examine the IIACP.LOG\_*nodename* file for possibly specific reasons for the startup failure. Take note of which databases are referenced in the IIACP.LOG\_*nodename* file.
3. Check that the II\_JOURNAL location is a valid location:
  - a. Check the definition of the II\_JOURNAL logical.
  - b. Check the II\_JOURNAL location. Type the following command at the operating system prompt:

```
directory II_JOURNAL: [INGRES.JNL]
```

4. Check that the disk containing the journal files is not 100% full. This makes it impossible to write journals and stops the Archiver.

If the journal disk became full as a result of the Archiver writing journal information, an error appears in the errlog.log stating a "device full" condition.

- a. After making disk space available, restart the Archiver using the `ingstart -dmfacp` command.
  - b. Optionally, modify the `acpexit` script to take whatever measures are necessary to notify you of this situation as soon as it occurs. For example, you can modify the script so that it sends you mail when the Archiver encounters an error.
5. If the journal files have been corrupted, the errlog.log reports any problems with the Archiver writing information to the journal files.

If there are corrupted files, or the Archiver process continues to fail, contact technical support, as described in *What You Need Before Contacting Technical Support* (see page 179).

## Check for DBMS Server Process Errors on VMS

If the DBMS server (II\_DBMS\_*nnn*, where the *nnn* is a number related to the server's process ID) does not start, follow these steps:

1. Examine the II\_CONFIG:errlog.log file for any information regarding DBMS Server startup.

Contact Technical Support to examine your authorization string.

2. If there is a DBMS Server startup failure reported in the II\_CONFIG:errlog.log, but no error reported in the errlog.log:
  - a. Check for the existence of any MBA*nnn* processes on the system that are linked to the IIRUNDBMS.EXE. Examine the MBA*nnn* process by issuing the following command at the operating system prompt:

```
show proc/id = pid /cont
```

where:

*pid* is the process ID of the MBA*nnn* process.

- b. If the MBA*nnn* process is linked to the IIRUNDBMS.EXE, free the IIRUNDBMS.EXE by issuing the following command at the operating system prompt:

```
stop /id = pid
```

where:

*pid* is the process ID of the MBA*nnn* process.

- c. Check the errlog.log for any new messages relating to DBMS Server startup failure.
3. Try starting a server manually, using the `ingstart -iibms` command.
  4. If the DBMS Server continues to fail to start, contact Technical Support.

## Check for JDBC Server Process Errors on VMS

The command to start the JDBC Server process (ijjdbc) from the Ingres system administrator account is `ingstart -ijjdbc`. If the JDBC server did not start:

1. Verify that the recovery process `II_IUSV_nnn` is running. If not, check for the errors described in [Recovery Process Errors on VMS](#) (see page 147).
2. Verify that the recovery process `II_IUSV_nnn` is not in a recovery state, as described in [Recovery Process Monitoring](#) (see page 159). This is a likely situation if there was a sudden shutdown because of a power failure or some other systems failure.
3. Try to start up a JDBC server. Issue the following command at the operating system prompt:  

```
ingstart - ijjdbc
```
4. If you corrected a JDBC server problem, verify that Ingres starts normally:
  - a. Shut down the partially started installation with the `ingstop` command.
  - b. Restart the installation with the `ingstart` command.

## Data Access Server Errors on VMS

The Data Access Server (`II_GCD_nnn`, where the *nnn* is a number related to the server's process ID) does not start if any of the following conditions occur:

- The `II_CONFIG:errlog.log` file has a notation regarding the failure of the Data Access Server process.
- The command **show system** does not list a `II_GCD_nnn` process.
- The utility `iinamu` does not display a process name when you enter the command `show dasrv` at the `IINAMU>` prompt.



## Check for Data Access Server Errors on VMS

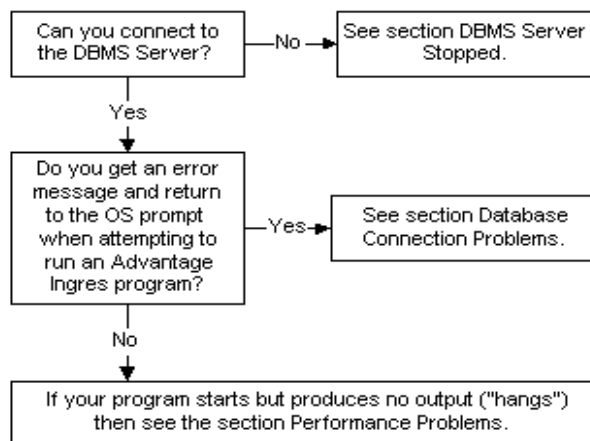
Use the following procedure if the Data Access Server process did not start:

1. Verify that you have installed the Ingres Net component.
2. Verify that you have installed the correct protocol drivers and they are functioning properly.
3. Verify the Data Access Server has been properly configured through the Configuration-By-Forms (cbf) utility.
4. If you are still having problems, set the following trace to capture additional diagnostic data before calling technical support:

```
define II_GCD_TRACE 4
define II_GCD_LOG filename
ingstart -iigcd
```

## Problems with Tools Startup

The following flow chart helps you isolate a problem when starting an Ingres tool:



## DBMS Server Stopped

Once started with `ingstart`, the DBMS Server process must continue running until the `ingstop` or `iimonitor` command is issued to stop it. If the DBMS Server stops running (“dies”) for any other reason, report it to technical support along with the associated error log messages and, if possible, the cause of the DBMS Server stopping.

1. Document error log entries associated with the process death. Details on reading the log files are described in [Check the Error Log Files](#) (see page 106). Save all errors for technical support.
2. Isolate the reason your DBMS Server process died.
  - a. Isolate which operations, application, query, and tables are needed to duplicate the problems. See the `copyapp` and `unloaddb` command descriptions in the *Command Reference Guide*.
  - b. Save this to make a test case for technical support.
3. If the immediate cause cannot be isolated, perform long-term diagnostics with `II_DBMS_LOG`. This diagnostic tool is especially valuable for fatal DMF errors.
  - a. Set `II_DBMS_LOG` to capture a “snapshot” of the DBMS Server when it stops by setting it to the full path name of a file before starting the DBMS Server. For example:

### Windows:

```
SET II_DBMS_LOG=%II_SYSTEM%\ingres\files\dbms.log
```

### UNIX:

C Shell:

```
setenv II_DBMS_LOG $II_SYSTEM/ingres/files/dbms_%p.trace
```

Bourne Shell:

```
II_DBMS_LOG = $II_SYSTEM/ingres/files/dbms_%p.trace export
```

```
II_DBMS_LOG
```

At startup, the `%p` in the `II_DBMS_LOG` specification is replaced by the Process Identifier (PID) of the server process. This prevents DBMS servers from clobbering each others logs (or the recovery process log)

### VMS:

```
define /system/exec II_DBMS_LOG -
```

```
II_SYSTEM:[INGRES.FILES]DBMS.LOG (for a system level installation)
```

```
define/group/exec II_DBMS_LOG -
```

```
II_SYSTEM:[INGRES.FILES]DBMS.LOG (for a group level installation)
```

- b. When the DBMS Server shuts down, information is dumped to the DBMS log file. You must rename the file before restarting Ingres or the new server overwrites the file. Prepare to send this file along with associated errors from the error log files to technical support for analysis.

## Database Connection Problems

Database connection problems occur in the following scenarios:

- If you cannot connect to any database
- If you encounter errors while connecting to individual databases

### No Database Connections

If you cannot connect to *any* database, follow these steps:

1. Check the `errlog.log` to see if there are any associated messages. These messages are often more informative than the message displayed on your screen and can quickly identify the source of failure. Check for associated messages in the `errlog.log` by using their timestamps.
2. Ensure that all your installation's processes are running, as described in [Check Ingres Installation on Windows](#) (see page 110), [Check Ingres Installation on UNIX](#) (see page 112), and [Check Ingres Installation on VMS](#) (see page 116).
3. Run `logstat` to insure that the logging system status is `ONLINE` and not `LOGFULL`. If the status is not `ONLINE`, see [How to Log System Issues](#) (see page 169).
4. Use `iinamu` to interrogate the Name Server. (For additional information, see `iinamu` in the *Command Reference Guide*). Type **show ingres** to verify that the DBMS Server has registered with the Name Server.
5. Use the `iimonitor` utility to see if you can connect to the DBMS Server. See `iimonitor` in the *Command Reference Guide*.

If you *can* connect, at the "`IIMONITOR>`" prompt type **show sessions** to examine DBMS server activity.

If you *cannot* connect to the DBMS Server, see [How to Log System Issues](#) (see page 169).

Check the Ingres environment variable/logical II\_DBMS\_SERVER:

**Windows:**

Use `ingpnmv` and your system's `echo` command to verify that `II_DBMS_SERVER` is not set, either in the Ingres symbol table or your local environment.

**UNIX:**

Use `ingpnmv` and your system's `env` or `printenv` command to verify that `II_DBMS_SERVER` is not set, either in the Ingres symbol table or your local environment.

**VMS:**

Use `II_DBMS_SERVER` to bypass the Name Server process. Issue the following command at the operating system prompt:

```
define II_DBMS_SERVER II_DBMS_pid
```

6. Restart the Name Server if `II_DBMS_SERVER` works.
7. If the error condition persists, contact technical support, as described in What You Need Before Contacting Technical Support (see page 179).

## Individual Database Connection Failure

If you can connect to *some* databases but not others:

1. Check the `errlog.log` for database connection error messages. These messages can quickly identify the source of the failure.
2. Check that the database is not exclusively locked by another user.
3. Check database permissions and ownership, as described in Check Ingres Installation on Windows (see page 110), Check Ingres Installation on UNIX (see page 112), and Check Ingres Installation on VMS (see page 116).

Verify that your database still exists and is not being recovered by the recovery system.

4. Type **catalogdb** and choose "catalog" / "databases" or run VDBA to check if the database is listed.

If the database does not exist, a corresponding error message is sent.

Using VDBA:

5. Select a server and click Connect DOM.
6. Select Database from the drop down menu and choose `infodb`. Make sure the database is not listed as inconsistent and that the status indicated is "VALID."

## Inconsistent Databases and Recovery

An inconsistent database occurs when administrative changes to a database in the transaction log do not agree with information maintained in the database's configuration file.

The main causes of inconsistent database errors are improper system administration procedures. These include:

- Initializing the transaction log file with the `-force_init_log` flag of `rcpconfig` while the log file still contains open transactions
- Moving or altering files or installation variables without using the appropriate utility such as `ingbuild` or `unloaddb`
- Improper procedures when recovering Ingres data from operating system backups

Inconsistent database errors can also be caused by hardware problems or software problems. For example, inconsistent database errors can be caused when the transaction log file or the configuration file has been corrupted by a hardware failure. Inconsistent database errors can also be caused by software bugs. In either case, contact technical support.

**Note:** You can also use VDBA for database backup and recovery. For more information, see the *Getting Started* guide or the online help for VDBA.

### Automatic Recovery

Ingres automatically handles the transaction failures that cause most database inconsistencies.

### Recovery During Normal Operation

If a user program exits or a transaction is aborted for some other reason, the DBMS Server automatically handles transaction rollback. This does not cause an inconsistent database.

## Recovery at Shutdown

At shutdown, all users must have exited their sessions; therefore, all transactions are committed. If users exited their sessions abnormally, the DBMS Server aborts any open transactions associated with the aborted sessions. Very long transactions take time to roll back and cause ingstop to seem to hang. The DBMS Server process cannot exit normally until it finishes recovering the aborted transactions.

If transactions are being rolled back on shutdown, allow the DBMS Server to finish this task before shutting down. If you do not, longer delays occur at startup time while the recovery process is performing rollback.

## Recovery at Startup

If transactions have been aborted and were not recovered by a normal shutdown, upon restart the recovery process performs recovery. This occurs, for example, if:

- Processes are forcibly killed from the operating system
- The machine is rebooted
- Power to the system is interrupted

The recovery process performs the following steps upon startup:

1. Reads the transaction log file. If there has not been a normal shutdown, the recovery process detects that databases are inconsistent—that is, that Ingres previously exited without completing all the transactions required for system and database consistency.
2. Proceeds through the transaction log file to back out uncommitted transactions and complete committed fast-commit transactions until the databases are again in a consistent state. While recovery is proceeding, no user interfaces can connect to a database.

Recovery actions are logged in the file `$II_SYSTEM/ingres/files/iircp.log`.

**Note:** When configured for the Ingres Cluster Solution, the log name has `_nodename` appended to the base log name. When the recovery process is complete, the message “Completed aborting transactions” is written to the RCP log file. Restart Ingres with the `ingstart` command. Users can reconnect to the databases.

## Recovery Process Monitoring

If you are monitoring Ingres startup after a machine reboot, the following messages are displayed:

```
Starting Ingres Name Server...
Starting Ingres Communications Server...
Starting Ingres Recovery Process...
```

If the transaction log contained uncommitted transactions when the machine failure occurred, the startup script pauses while the recovery process recovers transactions from the transaction log file. No messages are printed to the screen.

If you are in doubt as to whether recovery is taking place during startup, or to monitor the recovery process, use the following procedure.

Display the recovery process log file by typing the following command at the operating system prompt:

### **Windows:**

```
%II_SYSTEM%\ingres\files\iircp.log
```

### **UNIX:**

```
tail -f iircp.log
```

### **VMS:**

```
II_SYSTEM: [INGRES.FILES]IIRCP.LOG_nodename
```

If the system is recovering, the recovery actions are logged to the IIRCP.LOG file. This indicates that Ingres is automatically recovering from possible inconsistencies.

Messages are printed to the log file during recovery:

- The message at the beginning of recovering transactions indicates that transaction recovery has begun.
- Intermediate messages track recovery progress. As recovery proceeds, progress messages (for example, "Recovered 31 of 130 transactions") are displayed.
- When done, the following message is printed:  
Recovery complete.

You can also use an operating system command to determine whether the recovery process is recovering transactions by checking to see if it is accumulating CPU time.

**VMS:**

You can issue the following command at the operating system prompt several times:

```
show process/cont/id=xxxx
```

where:

xxxx is the process ID of the recovery process.

If the CPU time and disk I/O count taken by the process is increasing, the recovery process is most likely recovering transactions from the log file.

**UNIX:**

On UNIX, you can also monitor the files in the database directory of the database you suspect of being the target of the updates that are being backed out. The following command entered at the operating system prompt shows whether data files are being updated:

```
ls -lt
```

The file most recently updated is listed first along with the time of last update.

If any of the monitoring techniques above indicate that transaction recovery is taking place, continue to monitor the recovery process until recovery has completed. When the recovery process is complete, CPU time is not accumulated.

After the recovery process has finished, restart the installation with the `ingstart` command. The `ingstart` utility first shuts down and brings up all required installation processes. Programs can connect to the databases.

## Inconsistent Database

If you receive an “inconsistent database” error after recovery is complete, it means updates and modifications were not properly completed or rolled back, and the database is therefore in an inconsistent state.

Following are examples of “inconsistent database” errors that indicate your database has become inconsistent:

```
E_DM0100 DB_INCONSISTENT Database is inconsistent
E_US0026 Database is inconsistent. Please contact the ingres system manager
E_DM9327 BAD_OPEN_COUNT
```



## Diagnose an Inconsistent Database

Diagnose the cause and extent of an inconsistent database problem before you attempt to recover your database. Knowing the cause of the problem is essential to choosing the proper recovery procedures. Once a database has been rolled forward from a checkpoint, recovered from an operating system backup, or forced consistent, you cannot determine the cause of inconsistency.

To diagnose the cause and extent of an inconsistent database problem:

1. Read and save the full text of the error messages in `errlog.log` and `iircp.log`.
2. Using VDBA, choose `infodb` from the Database menu to read the database's configuration file and identify the cause of inconsistency. You can also enter the `infodb` command at the operating system prompt.

If the configuration file can be opened and read, the cause of the inconsistency is displayed. Save the output of `infodb` for technical support.

If the database's configuration file, "aaaaaaa.cnf", cannot be read, it is corrupted. You need to recover from a backup, as described in [Recover an Inconsistent Database](#) (see page 166).

3. Review the history of your Ingres installation. Look for improper system administration procedures that have caused the database to become inconsistent. See the table in [Common Causes of Inconsistent Databases](#) (see page 161).
4. Report your problem to technical support. If inconsistent database was not caused by incorrect system administration procedures, hardware failure, or known operating system software bugs, record the information, as discussed in [What You Need Before Contacting Technical Support](#) (see page 179).

## Common Causes of Inconsistent Databases

Common causes of inconsistent databases are:

- Operating system backups
- Incorrect installation paths
- Disabling of logging/recovery system
- Use of unsupported hardware configuration

## Inconsistencies Due to Operating System Backups

To recover a database from an operating system backup that was made while the installation was running, see [Recover an Inconsistent Database](#) (see page 166).

## Inconsistencies Due to Incorrect Installation Paths

Changing Ingres installation variables (such as II\_SYSTEM, II\_DATABASE, II\_CHECKPOINT, II\_JOURNAL, II\_DUMP, or II\_WORK) without using proper procedures, causes inconsistency between the information stored in the installation variables and those stored in the database configuration file "aaaaaaaa.cnf".

Database inconsistency can occur if you move a database, table, application or some other object by using operating system commands rather than the supported Ingres utilities. If the inconsistency is the result of moving a database from another location or installation without using unloaddb, you must remove the database using destroydb, recreate the database using createdb, and repopulate the database using the unloaddb utility. For details, see the *Database Administrator Guide*.

A database file can become corrupted from hardware or software failures of various kinds. A data file can be inadvertently deleted by hand, but this is rare because only the user who owns the installation can write to the database directories.

If you are in doubt about whether transactions are being recovered, run the logstat utility and examine the "Status" field. It is marked RECOVER if in the recovery state. While a recovery is taking place, for example when restarting after a system failure, the recovery process requires time to read through the transaction log file to back out uncommitted transactions and complete fast-commit transactions. To users, the system appears to hang.

## Examine Configuration File of a Database

To examine the configuration file of your database, use the VDBA Database, Infodb menu command. You can also enter the infodb command at the command prompt.

1. Compare the path information for the checkpoint, journal, data and dump locations with that defined for these environment variables and logicals as displayed by the following command:

**Windows:**

```
ingprens
```

**UNIX:**

```
ingprens
```

**VMS:**

```
show log ii
```

2. Return the installation logicals to the values displayed to by infodb, if the values have chanded. If these values are not the same, the installation logicals have been changed, or the database has been imported from some other Ingres installation.
3. If you need to change the existing values of Ingres installation variables or import a database from another site, you must use the unloaddb utility, as this creates a new, up-to-date configuration file for the database. For a discussion of Ingres environment variable and logicals that cause an inconsistent database if changed after installation is completed, see the chapter Environment Variables and Logicals.

## Recovery Rules

The following are some rules that you should keep in mind about the recovery of transactions:

- It takes at least as long to recover aborted transactions as it took to execute them originally.
- The amount of time required for recovery depends on the number of users and transactions, transaction semantics (whether autocommit is set), and the consistency point interval.
- While recovery is proceeding, all users are denied access to databases. Any attempt to connect to a database at this point returns an error such as the following:  

```
E_LQ0001_STARTUP gca protocol service request failure.
```
- Database inconsistency can occur if a user or the system administrator attempts to “force” entry into the installation by running `rcpconfig` with the `-force_init_log` flag (thus erasing the transaction log file) before the recovery system has finished rolling back the uncommitted transactions during recovery.

After a system failure, monitor recovery and always allow it to proceed until the “Recovery Complete” message appears in `iircp.log`.

## Inconsistencies Due to Disabling of Logging or Recovery System

By disabling the logging or recovery system, the DBA can temporarily turn off logging for the database to speed bulk loading of data. If logging has been turned off for this database, a NOLOGGING error message appears in the error log file. Typically this message is:

```
E_DM9050_TRANSACTION_NOLOGGING Database dbname has been updated
by a session running with SET NOLOGGING defined.
```

If the database has become inconsistent, you can check for this error message by typing the following command:

### Windows:

```
findstr DM9050\ %II_SYSTEM%\ingres\files\errlog.log
```

### UNIX:

```
grep DM9050 \$II_SYSTEM/ingres/files/errlog.log
```

### VMS:

```
search -
II_SYSTEM: [INGRES.FILES]ERRLOG.LOG -
DM9050
```

If the NOLOGGING error message appears, logging *was* disabled on this database. If the NOLOGGING message in the error log was written later than the most recent checkpoint of this database, the database must be restored from the checkpoint. To determine if this is the case, compare the timestamp on the error message in errlog.log with the timestamp in the “checkpoint history” field of the output from the command `infodb dbname`.

For details on set nologging, see the *Database Administrator Guide*.

## Database Inconsistencies Due to Use of Unsupported Hardware Configurations in UNIX

Database inconsistencies can be caused by the use of unsupported hardware configurations on NFS. In systems that include Network File System (NFS) mounts, be aware that Ingres:

- Supports NFS client installation configurations in which the DBMS Server process and data directories on one node are accessed by application programs executing on another.
- Does *not* support running DBMS servers on one node and accessing data directories on another network node using NFS. The configuration can cause undetected write errors that lead to database inconsistency.

To check your configuration, type **mount** at the operating system prompt. Make sure that the data directories (II\_DATABASE, II\_CHECKPOINT, II\_JOURNAL and II\_LOG\_FILE) are not NFS-mounted from a remote node.

For a description of the supported NFS client installation procedures, see *Getting Started*.

## Recover an Inconsistent Database

The recommended method of recovering an inconsistent database is to use rollforwarddb from VDBA. You can also enter this command from the command line. If no Ingres checkpoint exists, you can recover from an operating system backup.

## Make Inconsistent Database Consistent

The recommended way to make your inconsistent database consistent is to use rollforwarddb. It recovers the database from a previous checkpoint and, if journaling was enabled, applies the associated journals. For the full procedure, see the *Database Administrator Guide*.

## Use of Operating System Backup

Ingres does not support nor recommend the use of operating system backups as your primary means of ensuring database recoverability. The Ingres checkpoint and journaling programs provide the secure way to ensure that your data is recoverable.

**Important!** *Operating system backups must be used only as a last resort, when Ingres checkpoints have been lost or destroyed, and only under the direction of technical support.*

## When No Backup Exists

If you have an inconsistent database for which no checkpoints or operating system backups exist, you can still gain access to that database and attempt to salvage the data using the `verifydb` utility.

The `verifydb` utility can be used to unset the "inconsistent database" flag in the configuration file `aaaaaaaa.cnf`. This permits access to the database; however, it does nothing to actually make the data consistent. If `verifydb` is used to force access to a database that is inconsistent, the state of the database remains unknown. Such a database becomes unsupported by technical support. Data can be lost and problems occur weeks or months later. Technical support cannot diagnose the state of such a database because the built-in consistency checks have been overridden.

The format `verifydb -oforce_consistent` does not recover a database. It merely allows access and continued operation to a database that is in an inconsistent state.

## Gain Emergency Access to an Inconsistent Database Using `verifydb`

If you *must* use `verifydb` to gain emergency access to data in an inconsistent database, do so as follows:

1. Save all information, as outlined in What You Need Before Contacting Technical Support (see page 179).
2. Back up the database directory at the operating system level.
3. Run `verifydb` in report mode by typing the following command at the operating system prompt:

```
verifydb -mreport -sdbname "dbname" -odbms_catalogs -u$ingres
```

Verifydb output is logged in `$II_SYSTEM/ingres/files/iivdb.log`:

4. To repair inconsistencies in the system catalogs interactively and force the database consistency flag, type the following commands at the operating system prompt:

```
verifydb -mruninteractive -sdbname "dbname" -oforce_consistent -u$ingres
```

```
verifydb -mruninteractive -sdbname "dbname" -odbms_catalogs -u$ingres
```

5. Call Technical Support if additional assistance is required to resolve the inconsistency.

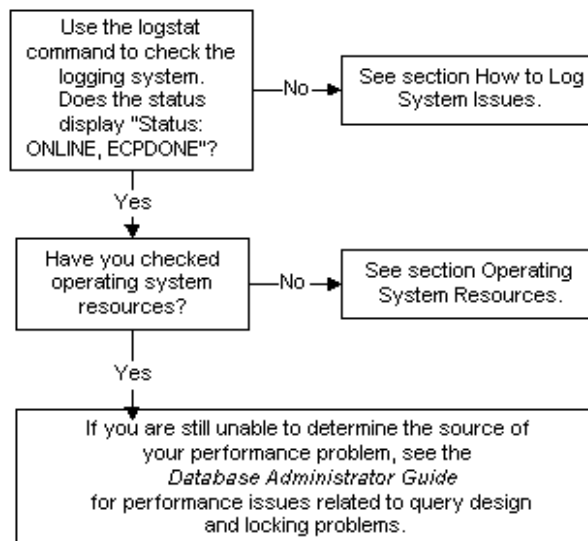
## Performance Problems

Most performance problems stem from multiple causes. A complete performance analysis must include each item in the Diagnostic Hierarchy section below. But if you need to resolve a specific performance problem quickly, focus your attention on that area. The information in this troubleshooting section is designed to assist you in resolving a performance issue after first isolating which factors are influencing performance.

Use the procedures in this section if an Ingres tool “hangs,” that is, your program seems to start but nothing happens. Other programs such as SQL display header information, but when you issue a command, nothing happens.

### Flow Diagram for Troubleshooting Performance Problems

Use the following flow chart to identify and isolate a system performance problem:





## Diagnose DBMS Server Due to Logging System Problem

Use the following procedure to diagnose a DBMS server that is not responding due to logging system problems.

1. Check the logging system by issuing the following command at the operating system prompt to invoke the logstat utility:

**Windows:**

```
logstat | more
```

**UNIX:**

```
logstat | more
```

**VMS:**

```
define sys$output filename
```

```
logstat
```

```
deassign sys$output
```

2. If you are unable to start up logstat, the recovery process (see page 159) has probably taken an exclusive lock and is in recovery.
3. If logstat starts up, check the status field, as described in Logstat Status Fields (see page 169).

### Logstat Status Fields

The problem states and resolutions of logstat fields are described in the following table:

Status Field	Description	Action
ONLINE, ECPDONE	A consistency point is completed. The system is fully functional and online.	The logging system has finished a consistency point. This status flag is present most of the time while the logging system is functioning normally. If this is your status, stop this procedure and review Operating System Resources.

<b>Status Field</b>	<b>Description</b>	<b>Action</b>
LOGFULL	The log file is full.	<p>A status of LOGFULL means that the system is suspended from processing new requests because there is no more room in the transaction log file. It remains so until the problem is corrected. Examine the other status field entries.</p> <p>If ARCHIVE is also indicated, the Archiver is actively processing the log file to free up space. The LOGFULL condition is removed when the Archiver is finished.</p>
CPNEEDED	A consistency point is needed.	CPNEEDED means the logging system is about to take a consistency point.
ARCHIVE	The Archiver is processing.	The status ARCHIVE means the Archiver process is archiving.
START_ARCHIVER	The Archiver has stopped.	A status of START_ARCHIVER means that the archiver process stopped. Restart the Archiver.
FORCE_ABORT	The force-abort- limit has been reached.	<p>If the status is FORCE_ABORT, a transaction came too close to the end of the transaction log file before the oldest was committed to disk, and reached the force-abort-limit. The oldest open transaction is aborted to free space for the new one.</p> <p>You can use logstat (or ipm) to monitor the force abort. The system frees the FORCE_ABORT status when the abort operation is completed.</p>
RECOVER	The recovery process is performing recovery.	If the status is RECOVER, the logging system is in recovery, a normal mode. Recovery requires some time because Ingres is optimized to commit rather than back out of transactions.

**Important!** *If the server is recovering a lengthy transaction, be aware that shutting down the server with iimonitor's "stop server" (or by killing processes at the operating system level) results in slower recovery. Allow the recovery to proceed normally. You can monitor the progress of the recovery with logstat and iimonitor or the ipm utility.*

Many other normal states appear on the status line. For additional information on reading logstat output, see the Logstat section in the *Command Reference Guide*.

## How to Avoid Logfull Abort

To avoid logfull aborts, examine the transaction processing strategy errors that caused the log file to fill. Consider the following points:

- If transactions do many updates or are left open for a long time, you must either commit them more frequently or increase the size of the log file. If the log file size is increased, check the "percentage of logfile written" before a consistency point is taken. For details, see the ipm, logstat, or VDBA output.
- Avoid continuous transaction errors by:
  - Setting autocommit on where applicable
  - Avoiding application errors that are caused by beginning transactions before they get all the user input required or by failing to use Ingres application timeout features. (Timeout features are discussed in the *SQL Reference Guide* and *Forms-based Application Development Tools User Guide*. Also, for a discussion of transaction processing, see the *SQL Reference Guide*.
- If you have taken the above precautions but still get log full/abort, the transaction log is sized too small for the types of transactions you need to process. Read the tips in the next section on choosing the correct transaction log file size.
- Use set log\_trace to examine the size and type of log records being written to the log file.

## Process of Resizing Transaction Logs

Resizing transaction logs involves these steps:

1. Determine whether the log requires resizing
2. Resize the transaction log file
3. Reestablish Dual Logging

## Determine Whether a Log Requires Resizing

To determine whether a log requires resizing, use **one** of the following procedures.

Using VDBA's Performance Monitor, start a monitoring session on the desired node:

1. Select the node in the left pane and click the Monitor toolbar button.
2. Select Log Information in the branch.
3. Click the Header tab in the right pane to display log usage.

Start the Interactive Performance Monitor utility (ipm).

1. Highlight the Log\_info menu item and select the Select option.
2. Highlight the Header menu item and select the Select option.
3. The Log file Diagram shows the percentage use of the Transaction Log.
4. Exit the ipm program.

Find the current peak transaction load by using the logstat utility to monitor the amount of log file that is in use during your peak hours. This allows you to choose the correct size for your log file. Because requirements change, you must monitor the log file regularly.

1. Enter the logstat command.
2. Find the value under % of log file in use.
3. If the log file needs to be resized, select the new size.

### **UNIX:**

Use mkrawlog to reconfigure a raw log file.

4. Shut down the Ingres installation by entering the ingstop command.

## Resize the Transaction Log File

To resize the transaction file, follow these steps:

1. Shut down the Ingres installation by entering the `ingstop` command.
2. Start the Configuration-By-Forms (cbf) or Configuration Manager (vcbf) utility.

- a. Choose Transaction Log from the menu.

The file name is displayed along with the current size of the Transaction Log.

- b. Choose Destroy.

When destroyed, the Transaction Log information table is emptied.

- c. Select the Create option.

- d. Enter the required Transaction Log file size in megabytes.

The transaction log is created.

- c. Exit cbf or vcbf.

4. Restart the Ingres installation by entering the command `ingstart`.

For details, see the *Getting Started* guide.

**Important!** *Reconfiguring your log file destroys the current contents of the file and leaves an empty, reinitialized log file after the reconfiguration is complete. When Ingres is shut down, all transactions are written to disk; therefore, to prevent inconsistent database problems, reconfigure the log file only after a successful Ingres installation shutdown procedure.*

## Reestablish Dual Logging

A failure in one of the two dual log files is most often symptomatic of hardware problems. In the event of such a failure, you must reestablish dual logging at the first opportunity.

To re-establish dual logging after a failure of one of the log files, follow these steps:

1. Shut down the installation.
2. Create a new log file in the location of the failed log.
3. Access the Configuration-By-Forms utility

**Windows:** Click Start on the taskbar, and chose Programs, Ingres, Configuration Manager.

**UNIX:** Type **cbf** or **vcbf** on the command line to access the Configuration-By-Forms and Configuration Manager Utility.

**VMS:** Type **cbf** or **vcbf** on the command line to access the Configuration-By-Forms and Configuration Manager Utility.

4. Select dual transaction log.
5. Select Reformat to reestablish dual logging.

This automatically determines which of the log files is valid and copies the contents of the valid log to the newly created log file. After the copy is complete, both log files are marked valid.

6. Restart the installation with the `ingstart` command.
7. Check that both log files are enabled. You can check the enabled status using either the Configuration-By-Forms (`cbf`) or Configuration Manager (`vcbf`) utility.

## Resource and Maintenance Problems

Good performance requires planning and regular maintenance. Make sure your operating system is configured with sufficient resources for Ingres. Insufficient system resources cause poor performance or prevent Ingres from starting.

## Tools for Identifying Operating System Problems

The following tools can help you identify operating system resource problems:

- Review the minimum requirements for a basic Ingres installation given in the Readme file. If your particular environment requires more resources, use the Ingres utilities to verify that there are enough resources.
- System resources can be monitored by some operating system utilities. Syntax details are described in *Unix Operating System Utilities* (see page 187).

### UNIX:

BSD:

pstat utility—to display the status of UNIX system tables and system swap space

vmstat utility—to display virtual memory status

System V:

- sar utility—to display activity of various system resources such as CPU utilization, swapping activity, and disk activity.
- show memory—displays the system memory resources and the amount of non-paged dynamic memory (total, free and in use).
- show process/id=pid/continuous—displays the amount of page faulting, working set, buffered I/O, and direct I/O the server is doing.
- show device—indicates if a particular disk drive is out of disk space.
- show device /files—if there is a problem starting an installation, this command can be used to make sure that an Ingres process is not holding on to a mailbox.

The installation utility allows the examination of all Ingres installed images, showing the amount of global pages and sections available and used.

### VMS:

The following VMS tools are very useful for tracing script problems:

- set verify—allows you to see the commands used in the installation script or any other VMS command procedure.
- set watch file/class=major (requires CMEXEC VMS privilege)—allows all the files that are being accessed to be displayed. This assists in diagnosing if a location of a file or library is being incorrectly referenced. Use this command when debugging Ingres processes interactively.
- The displays can be turned off with the following command entered at the operating system prompt:

**Important!** *This is an unsupported VMS command. Use at your own risk.*

```
set watch file/class=none.
```

Syntax details are described in VMS Operating System Utilities (see page 193).

## Diagnose Poor Performance Due to Insufficient System Resources

If Ingres seems slow or unresponsive for no apparent reason, follow these steps to diagnose the problem. Write down any error messages you receive when performing these steps:

1. Connect to your DBMS Server through Ingres monitors:

- a. First display the *server\_number* of your DBMS Server using the *iinamu* utility:

```
iinamu  
IINAMU> show ingres
```

- b. Connect to the DBMS Server monitor by typing the command:

```
iimonitor server_number
```

- c. To see the DBMS Server sessions, at the *iimonitor* prompt type:

```
IIMONITOR> show sessions
```

- d. Check the status of the sessions to determine which one is making excessive use of the server. (You can use VDBA to check session status.)

For syntax details, see the sections *iimonitor* and *iinamu* in the *Command Reference Guide*.

2. If repeated "show sessions" commands in *iimonitor* show that the query session is continually in a *CS\_EVENT\_WAIT* (LOCK) state, the problem involves concurrency and locking.

Alternatively, you can use the VDBA Performance Monitor to check for this problem.

- a. Select Servers in the left pane of the Performance Monitor.
- b. Select INGRES.
- c. Select Sessions in the Servers.



3. If the session alternates between CS\_EVENT\_WAIT and CS\_COMPUTABLE, this indicates that the query is processing. However, if the query is taking an excessive amount of time, set up a trace on it, as described in Trace Utilities (see page 206).

- a. Interrupt the query that is running:

- *Interactively*, use Ctrl+C and wait.
- *In batch or background mode*, use the following command to terminate:

**Windows:** In the Task Manager, highlight iidbmst and click on End Process

**UNIX:**

```
kill pid
```

**VMS:**

```
stop process/id = pid
```

where *pid* is the process ID of the query.

**Important!** *The command format "stop proc" must be used only as a last resort. Use of this option can cause more problems than it solves.*

- b. Issue the command **set gep**.

- Rerun the query. This outputs a query execution plan.
- Alternatively, start an SQL window on the database in VDBA and click the Display Query Execution Plan button to graphically display the query plan.
- It is important to note that interrupting a query requires some time because Ingres is optimized to commit rather than back out of transactions. It takes at least as long to back out of a transaction as to process the transaction normally. The transaction must be fully backed out before sessions can resume and locks are freed.

4. It is useful to note whether the query runs differently when called from other Ingres tools. For example, try issuing the same query from Interactive SQL, and Embedded SQL
5. Determine if you can access all of the data in the tables in all components of the query.

- a. From the Terminal Monitor type:

```
select count(*) from tablename
```

This verifies that Ingres can sequentially access every row in the table and indicates that other access paths (secondary indexes, hash pointers, B-Tree page pointers, and so on) can cause the problem. Queries using restrictive where clauses probably are using these secondary access methods.

- b. Check for permits that apply to this data by typing the following command from the Terminal Monitor:

```
help permit tablename
```

- c. Check for restrictions that you have placed on this data through the Knowledge Management Extension features, such as database rules. From the Terminal Monitor enter the following query to determine if a table is subject to restrictions imposed by a database rule:

```
select * from iirules where table_name = 'tablename'
```

Alternatively, in the right pane of VDBA's Database Object Manager (DOM) window:

- Select the Databases branch and the desired database
- Select Tables and the desired *table\_name*
- Select Rules

## What You Need Before Contacting Technical Support

Before contacting technical support, gather as much information as possible and have the following information available.

### Windows Installations

Save any relevant errors in each of the following logs:

- `errlog.log`
- `iiacp.log`
- `iircp.log`

Have the following information available:

- Your contact ID.
- Your exact Ingres release.
- Your exact operating system release. Obtain this with the `VER` command.
- The current Ingres installation environment. Use the following command at the operating system prompt:

```
ingprenv > filename
```

- The current user environment. Use the following command at the operating system prompt:

```
SET > filename
```

- A clear description of what you are trying to do.
- An indication of whether the failure occurs reproducibly.

### UNIX Installations

Save any relevant errors in each of the following logs:

- `errlog.log`
- `iiacp.log`
- `iircp.log`

**Note:** If your system is configured for Ingres Cluster Solution, separate `iiacp.log` and `iircp.log` files will exist for each node with `_nodename` appended to the log names. All logs for all nodes must be checked for pertinent errors.

Have the following information available:

- Your Contact ID.
- Your exact Ingres release. Obtain this with the following command at the operating system prompt:

```
cat $II_SYSTEM/ingres/version.rel
```

- Your exact operating system release. Obtain this with the following command at the operating system prompt:

```
uname -a
```

Or:

```
cat /etc/motd
```

- The current Ingres installation environment. Use the following command at the operating system prompt:

```
ingpreenv > filename
```

- The current user environment. Use the command:

BSD:

```
printenv > filename
```

System V:

```
env > filename
```

- A clear description of what you are trying to do
- An indication of whether the failure occurs reproducibly

## VMS Installations

Have the following information available:

- Your Contact ID.
- Your exact Ingres release including any Ingres patches installed. The following header shows the release number in the format *nn/nn*:  
Ingres Alpha Version *nn/nnnn* (vax.vms/00)
- Your exact operating system release including any VMS patches installed.
- The current Ingres installation environment. To create a file type:

```
define sys$output filename  
show logical II*  
show logical *ING*  
deassign sys$output
```

- If you have been unsuccessful in recovering an inconsistent database, it is important to save this additional diagnostic information.

- For the transactions in the transaction log file, type **logstat**:
- For the state of the database configuration file, type **infodb dbname**
- For copies of the relevant error log files type:

```
search ERRLOG.LOG date
type IIRCP.LOG
```

where:

*date* is the date the database became inconsistent.

IIRCP.LOG\_*NODENAME* for each node in the cluster, if configured for Ingres Cluster Solution.

After retrieving all the diagnostic information, type **deassign sys\$output**.

- A clear description of what you are trying to do when the problem occurs. This includes commands, queries, and so on.
- An indication of whether the failure occurs reproducibly.
- Save the information in your error logs. Save any relevant errors in:
  - ERRLOG.LOG
  - IIACP.LOG (IIACP.LOG\_*NODENAME* for each node in the cluster, if configured for Ingres Cluster Solution)
  - IIRCP.LOG\_*NODENAME* for Ingres Cluster Solution
- Save DBMS server parameter settings:  
II\_CONFIG:CONFIG.DAT
- A copy of any changes that have been made to II\_config:config.dat.  
If the Configuration-By-Forms (cbf) utility is used to update the configuration file, a copy of all changes are written to II\_CONFIG:CONFIG.LOG.
- Have Dial-In information available (if allowed).



# Chapter 7: Using Monitoring and Tracing Tools

---

This chapter describes monitoring and tracing tools that help you troubleshoot the Ingres installation.

## Supported Monitoring and Tracing Tools

A number of tools are available to help you troubleshoot and maintain your installation. In addition to the standard commands and utilities, the tools include:




- Special purpose utilities:
  - Ingres system utilities
  - Operating system utilities
- Trace utilities:
  - Ingres trace utilities
  - Operating system trace utilities

**Note:** You can use Ingres Visual Manager, Visual Performance Monitor, and VCDA to visually monitor and maintain your installation. For additional information about these tools, see the chapter “Managing Your System and Monitoring Performance” and online help.

## System Utilities

Ingres system utilities monitor aspects of the installation for troubleshooting or other special purposes. The system utilities are not needed for normal operation and are generally used with the help and advice of Technical Support. For more information about system utilities see the *Command Reference Guide*.

Day-to-day Ingres commands follow:

	<b>Command</b>	<b>Name</b>	<b>Description</b>
	cacheutil	Shared Buffer Cache Utility	Lists status information on active shared memory segments; can also be used to deallocate an inactive shared memory buffer cache.
<b>Windows:</b>	ipcclean	Shared Memory Allocation	Deallocates shared memory. 
<b>UNIX:</b>	csinstall, cscleanup, ipcclean csreport	Shared Memory Allocation	Allocates, deallocates, and lists status of UNIX shared memory and semaphore resources. Allocation is based on the value assigned II_LG_MEMSIZE. 
	iimonitor	Server Maintenance Utility	Lists DBMS server information; also allows a privileged user to stop the server or an active session.
	iinamu	Name Server Utility	Lists current servers; also allows a privileged user to add and delete servers.
<b>VMS:</b>	iishowres	Show Resources	Lists shared memory used by the locking and logging system. 
	lartool	Logging, archiving and recovery tool	Allows the system administrator to manually abort or commit a distributed transaction involved in two-phase commit protocol.
	lockstat, logstat	Locking, Logging Status	Display information from the locking or logging system. These functions are replaced by the forms-based ipm utility.



Command	Name	Description
rcpconfig	Logging/ Locking Configuration	Manipulate log files or shut down the logging/locking systems (RCP module).
ija	Logging, archiving, and recovery tool	Allows the system administrator to manually abort or commit a transaction involved in two-phase commit protocol.

## Operating System Utilities

Many operating systems provide utilities to administer and troubleshoot. Most of the operating systems provide context sensitive help to use the tools.

### Windows Operating System Utilities

You can use the following Windows operating system utilities to monitor Ingres:

- Windows Diagnostics
- Windows Performance Monitor
- Windows Event Viewer
- Windows Registry Editor
- Windows Task Manager

For a full description of the Windows utilities, see your Windows documentation.

### Windows Diagnostics

The Windows Diagnostics program can help you determine your operating system's configuration. This tool can be found in Settings, Control Panel, Administrative Tools, Computer Management.

## Windows Performance Monitor

Performance Monitor is a Windows graphical tool for measuring the performance of your own computer or other computers on a network. On each computer, you can view the behavior of objects such as processors, memory, cache, threads, and processes. Each of these objects has an associated set of counters that provide information on such things as device usage, queue lengths, and delays, as well as information used for throughput and internal congestion measurements.

It provides charting, alerting, and reporting capabilities that reflect current activity along with ongoing logging. You can also open log files at a later time for browsing and charting as if they were reflecting current activity. To monitor performance on Windows, see the operating system documentation.

## Windows Event Viewer

Event Viewer is a tool for monitoring events in your system. You can use Event Viewer to view and manage System, Security, and Application event logs. To access the Event Viewer, right-click on the My Computer icon and select Manage. The Computer Management Window is displayed. The Event Viewer is available under the System Tools.

## Windows Registry Editor

This program can be used to view the system configuration and environment. For a description of how the information is presented and the capabilities of the utility, see the online help.

## Use REGEDIT32

To start the Registry Editor, run REGEDT32.EXE from File Manager or Program Manager, or type **start REGEDT32** in a command window. The Registry Editor program has a similar view to the File Manager program.

## Windows Task Manager

The Task Manager enables you to monitor and control your computer and what is running on it. It shows you programs and processes that are running as well as performance. To access the Task Manager, right click an empty area in the task bar and click Task Manager.

## UNIX Operating System Utilities

You can use the following UNIX operating system utilities to monitor Ingres:

- ps
- iostat
- pstat
- sar
- sdsydef
- vmstat

For a full description of the UNIX options, see your UNIX documentation (or online help). A detailed description of the utilities can be found in the *Command Reference Guide*.

**Note:** Not all of the utilities are present on every UNIX system. Some are present only in a BSD or System V environment but not both.

### ps

This command provides virtual memory and cpu information on each active process submitted from your account. Here is a sample ps output:

```
PID  TT  STAT  TIME  SL  RE  PAGEIN  SIZE  RSS  LIM  %CPU  %MEM  COMMAND
xx06 p3  s    28:50  13  99  45886  3696  2856  xx   0.0  39.3  iidbms
xx94 p3  s     4:24   0  99  2899   720   344  xx   0.0   4.7  dmfrcp
xx09 p3  I     0.51  99  99  4488   64   184  xx   0.0   2.5  iislave
xx19 p3  I     0.57  99  99  5764   64    17  xx   0.0   2.4  iislave
xx96 p3  I     0.04  99  99  1852   696   160  xx   0.0   2.2  dmfacp
```

The display fields are as follows:

Field	Description
PID	Process ID field
TT	Controlling terminal
STAT	Process status  Runnability of the process: Runnable (r), Stopped (t), Disk or other short-term wait (d), Sleeping (s), or Idle (i)  Swap status: Swapped out (w) or Loaded in core (blank)  Process priority change: Reduced (n), Increased (>) or No change (blank)
SL	Sleep time (seconds blocked)
RE	Residency time (seconds in core)

Field	Description
PAGEIN	Number of disk I/Os resulting from page references not in core
SIZE	Virtual process size
RSS	Resident set size
LIM	Soft memory limit (setrlimit), else "xx"
%CPU	CPU utilization (1 minute decaying average)
%MEM	Memory utilization
COMMAND	Process name

## iostat

The `iostat` command returns information about I/O status. It lists statistics on current I/O activity for each disk device and system CPU utilization percentages. Here is a sample `iostat` output:

```

tty  cpu
tin  tout   us   ni   sy   id
  1    18   19    0    3   78

/dev/*dsk/c0d*s*/dev/*dsk/c1d*s*/dev/*dsk/c2d*s*/dev/*dsk/c4d*s*
bps  sps  msp  bps  sps  msp  bps  sps  msp  bps  sps  msp
  2   0.1 61.7  1   0.0 95.5  1   0.0 60.4  2   0.2 44.3

```

- The `tin` and `tout` display fields show the number of characters written to and from terminal devices.
- CPU information includes the % time spent in user mode (`us`), "niced" user mode (`ni`), system mode (`sy`), and idle (`id`).
- The disk I/O for each disk device shows the average number of blocks transferred per second (`bps`), average number of seeks per second (`sps`), and average time per seek in ms (`msp`).

## pstat (BSD)

Print Statistics utility is useful for examining system tables and system swap space.

The command `/etc/pstat -s`, shows system swap space activity. An example output:

```
589944 used, 667364 free, 61864 wasted
avail (num*size): 316*2048 5*1024 6*512 13*256 21*128 38*64 54*32
```

**Note:** For a description of the output units, see your operating system guide or online MAN pages. On some machines output is in bytes (as in this example). On others it is in units equal to the machine's page size.

The command `/etc/pstat -T` gives a list of used and free slots in various system tables. Here is a sample output:

```
1717/3216 files
937/2384 inodes
333/1300 processes
266/325 mfiles
595176/1256780 swap
```

The `pstat` utility can be used to give other useful information if the user has knowledge of the UNIX kernel. It can tell you, for example, the state of a running process, which files are open, which operating system locks apply and which processes have which files open.

## sar (System V)

The sar (System Activity Reporter) command can be used to list CPU utilization statistics, disk and swapping activity, and other status options. The UNIX sar command samples cumulative activity counters in the operating system and reports on various system activities. Just a few of the reports available that are of use for the system administrator are CPU utilization, disk activity, and swapping activity. See examples below.

There are many other options available, including reports on unused memory pages and disk blocks (swap space), and message and semaphore activity. For the system administrator, sar is an indispensable tool.

- The command **sar -u 5**, shows CPU utilization output as in the following sample output:

```
06:38:26  %usr  %sys  %wio  %idle
06:38:31      1    3    96    0
```

This output shows the portion of CPU time spent running in user mode, system mode, idle with some process waiting for block I/O, or otherwise idle.

- The command **sar -d 5**, shows disk activity output as in the following sample output:

```
06:39:53  device  %busy  avque  r+w /s  blks/s  await  avserv
06:39:53  disc0-0   16    13.6    6      17    349.6   29.7
           disc0-5    0     1.0     0       2     0.0    20.0
```

The display shows the portion of time the device was busy servicing requests (%busy), the average number of outstanding requests (avque), the number of data transfers (r+w/s), the number of bytes transferred in 512-byte units (blk/s), the average time in milliseconds that requests wait idly on queue (await), and the average time for requests to be serviced (avserve).

- The command **sar -w 5** shows swapping activity output as in the following sample output:

```
06:40:45  swpin/s  bswin/s  swpot/s  bswot/s  pswch/s
06:40:50    0.00    0.0     0.00    0.0     19
```

The fields swpin/s, swpot/s are the number of swap transfers. The fields bswin/s, bswot/s are the number of 512-byte units transferred in/out. The field pswch/s is the number of process switches.

## sysdef (System V)

The `/etc/sysdef` (System Definition) command (implemented only on some System V operating systems) displays a list of all tunable kernel parameters and their configuration in the currently active kernel. Here is a sample `sysdef` output:

Tunable Kernel Parameters:

```

250 buffers in buffer cache (NBUF)
 30 entries in callout table (NCALL)
150 inodes (NINODE)
150 s5inodes (NSSINODE)
150 entries in file table (NFILE)
 25 entries in mount table (NMOUNT)
100 entries in proc table (NPROC)
210 entries in shared region table (NREGION)
120 clist buffers (NCLIST)
 25 processes per user id (MAXUP)
 64 hash slots for buffer cache (NHBUF)
 20 buffers for physical I/O (NPBUF)
 50 size of system virtual space map (SPTMAP)
 16 fraction of memory for vhandlow (VHNDFRAC)
  0 maximum physical memory to use (MAXPMEM)
 10 auto update time limit in seconds (NAUTOUP)
 60 maximum number of open files per process (NOFILES)
256 number of streams queues (NQUEUE)
 48 number of streams head structures (NSTREAM)
  4 number of 4096 bytes stream buffers (NBLK4096)
 32 number of 2048 bytes stream buffers (NBLK2048)
 32 number of 1024 bytes stream buffers (NBLK1024)
 32 number of 512 bytes stream buffers (NBLK512)
 64 number of 256 bytes stream buffers (NBLK256)
128 number of 128 bytes stream buffers (NBLK128)
256 number of 64 bytes stream buffers (NBLK64)
256 number of 16 bytes stream buffers (NBLK16)
128 number of 4 bytes stream buffers (NBLK4)

2560 maximum size of user's virtual address space in pages (MAXUMEM)
2560 for package compatibility equal to MAXUMEM (MAXMEM)
 25 page stealing low water mark (GPGSLO)
 40 page stealing high water mark (GPGSHI)
  9 page aging interval (AGEINTERVAL)
  1 bdflush run rate (BDFLUSHR)
 25 minimum resident memory for avoiding deadlock (MINARMEM)
 25 minimum swappable memory for avoiding deadlock (MINASMEM)
  1 maximum number of pages swapped out (MAXSC)
  1 maximum number of pages saved (MAXFC)

*
* Utsname Tunables
*
 3.2 release (REL)
chipmunk node name (NODE)
chipmunk system name (SYS)
      2 version (VER)

```

```
*
* Streams Tunables
*
  87 number of multiplexor links (NMUXLINK)
  9 maximum number of pushes allowed (NSTRPUSH)
 256 initial number of stream event calls (NSTREVENT)
  1 page limit for event cell allocation (MAXSEPGCNT)
4096 maximum stream message size (STRMSGSZ)
1024 max size of ctl part of message (STRCTLSZ)
  80 max low priority block usage (STRLOFRAC)
  90 max medium priority block usage (STRMEDFRAC)

*
* RFS Tunables
*
 10 entries in server mount table (NSRMOUNT)

*
* IPC Messages
*
 100 entries in msg map (MSGMAP)
2048 max message size (MSGMAX)
4096 max bytes on queue (MSGMNB)
  50 message queue identifiers (MSGMNI)
   8 message segment size (MSGSSZ)
  40 system message headers (MSGTQL)
1024 message segments (MSGSEG)
*
* IPC Semaphores
*
 10 entries in semaphore map (SEMAP)
 10 semaphore identifiers (SEMMNI)
 60 semaphores in system (SEMMNS)
 30 undo structures in system (SEMMNU)
 25 max semaphores per id (SEMMSL)
 10 max operations per semop call (SEMOPM)
 10 max undo entries per process (SEMUME)
32767 semaphore maximum value (SEMVMX)
16384 adjust on exit max value (SEMAEM)

*
* IPC Shared Memory
*
1048576 max shared memory segment size (SHMMAX)
  1 min shared memory segment size (SHMMIN)
 100 shared memory identifiers (SHMMNI)
  6 max attached shm segments per process (SHMSEG)
 512 max in use shared memory (SHMALL)
*
* File and Record Locking
*
 100 records configured on system (FLCKREC)
```



## vmstat

The `vmstat` (Virtual Memory Statistics) command returns virtual memory status information, including process states and paging activity. Here is a sample `vmstat` output:

```
procs      memory      page      disk      faults      cpu
r b w avm  fre di  re rd pi po de z0 z1 z2 z3 in sy cs us sy id
1 0 0 2536 456 24 2 1 4 0 0 1 1 0 1 24 475 23 4 6 91
0 0 0 2748 356 24 0 0 0 0 0 2 0 0 3 26 323 29 1 5 95
0 0 0 1044 344 24 0 0 0 0 0 0 0 0 0 16 216 18 0 3 97
0 0 0 2288 344 24 0 0 0 0 0 0 0 0 0 19 334 27 1 4 95
1 0 0 2372 332 24 0 0 4 0 0 0 0 0 1 28 552 40 1 6 93
```

The `procs` columns define the process states: in run queue (`r`), blocked for resources (`b`), and runnable or short sleeper (`w`). The memory columns show virtual and real memory status: `avm` is active virtual pages (belonging to processes active in approximately the last 20 seconds), `fre` is size of the free list, and `di` is the number of dirty pages. The page columns show page faults and paging activity. These are expressed in units per second, averaged over 5 seconds: `re` are page reclaims, `rd` are page reclaims from the dirty list, and `pi`, `po` are pages paged in/out. The `de` field is anticipated short-term memory shortfall. The disk columns list disk activity, showing the device name and operations per second.

## VMS Operating System Utilities

VMS utilities vary depending on the operating system release. For detailed information on utilities supported by your system, see your operating system guides (or online help).

## help

The VMS online help facility (`help`) provides information on the specified command. The syntax is `help command`.

## monitor

The VMS monitor utility returns dynamic information about VMS system performance for a specified component. Some of the most useful components to monitor are:

- monitor system — monitors system statistics
- monitor process/topcpu—monitors the process currently using the most CPU
- monitor process/topdio—monitors the processes that use the most disk I/O
- monitor process/topbio—monitors the processes that use the most buffered I/O
- monitor process/topfault— monitors the processes that cause the most page faults
- monitor pool—monitors non-paged pool statistics

## show

The VMS show utility returns process and other VMS system status information on a specified component. Qualifiers (/qualifier) are used on certain components, as shown in the examples that follow. Some useful forms of the VMS show command appear in the following list.

- **show system or show system/full**—displays status information about current processes. Here is a sample output:

```
date and time stamp
process ID number
process name and identification with UIC
processing state
priority
total process I/O
cumulative cpu time used
cumulative page faults
amount of physical memory being used
type of process
```

```

VAX/VMS V6.0 on node F00 9-FEB-2001 09:34:17.48 Uptime 13 13:46:12
Pid      Process Name      State  Pri  I/O    CPU           Page flts
Ph.Mem
00205374 DMFRCPBE           HIB    6    6270   0 00:00:15.71  2756
4280
00202F75 DMFACPBE           HIB    5    934    0 00:00:57.09  1115
2302
00204376 II_GCNBE           HIB    6    667    0 00:00:03.15  782
1363
00202B77 II_GCC_BE_2B77     HIB    5    142    0 00:00:02.15  664
1549
00205578 II_DBMS_BE_5578    HIB    5   38540   0 00:04:31.58  9390
13903
002002C3 DMFCSPBE           HIB    5   71071   0 00:00:22.27  399
173

```

- **show system /out = filename**—outputs to a file. This provides a record of the current status.
- **show process**—displays information about a process and subprocesses. This command requires the GROUP privilege to show other processes in the same group or the WORLD privilege to show processes outside your group. If no qualifier is entered, only a basic subset of information is displayed:

```

date and time stamp
process terminal
user name and UIC
node name
process name and process identification
priority
default directory
allocated devices

```

Some of the show process qualifiers you can use are:

**/memory**—displays the process' use of dynamic memory areas. This qualifier is allowed only for the current process.

**/quotas**—displays, for each resource, a quota or a limit. The quota values reflect any quota reductions resulting from subprocess creation. Limit values reflect the resources available to a process at creation.

**/priv**—displays current privileges for the process.

- **show cluster**—monitors and displays cluster activity and performance.
- **show cpu**—displays the current state of the processors in a VMS multiprocessing system. This command requires a change of mode to the kernel (CMKRNL) privilege. The show cpu qualifiers are:
  - /all**—selects all configured processors, active or inactive, as the subject of the display.

`/full`—produces information from the summary display and also lists the current CPU state, current process (if any), revision levels, and capabilities for each configured processor. The display indicates which processes can execute only on certain processors in the configuration.

- **show device**—displays the status of a device on the system. You can use the following show device qualifiers:

`/allocate`—displays all devices currently allocated to processes.

`/file`—shows all open files on the device, together with the process name and the process ID that opened the file.

`/mounted`—displays all devices currently having volumes mounted.

- **show logical /full**—displays all logical names in one or more logical name tables, or displays the current equivalence strings assigned to specified logical names, with iterative translations.
- **show memory /full**—displays status information on VMS system tables and system swap space.
- **show symbol**—displays the current value of a local or global symbol. Ingres symbols are defined by running the following:  
  
II\_SYSTEM: [INGRES.UTILITY]INGSYSDEF.COM – Defines Ingres system administration, DBA, and user symbols.  
  
II\_SYSTEM: [INGRES.UTILITY]INGDBADEF.COM – Defines Ingres DBA and user symbols.  
  
II\_SYSTEM: [INGRES.UTILITY]INGUSRDEF.COM – Defines Ingres user symbols.
- **show users**—displays status information on current users.
- **show working\_set**—displays the working set limit, quota, and extent assigned to the current process.

## sysgen

The VMS sysgen utility displays system configuration parameters. To run this utility, execute the following command at the operating system prompt:

```
run sys$system:sysgen
```

For a full description of the system definitions, see your VMS System manual.

## Vendor Utilities

Many vendors also provide monitor utilities. These often present a more user-friendly interface than the generic utilities mentioned above. Check your operating system documentation for availability of these useful vendor-specific utilities.

## Error Messages

Errors in Ingres can be classified into Fatal and Non-Fatal errors. Ingres displays error messages in a particular format. This format helps you find the specific problem that Ingres has encountered. Error messages and their formats are explained in the following sections.

### Error Message Format

Error messages in Ingres consist of a three-part error code and accompanying message text. The code segment has the format:

`E_FCxxx_msg message text`

The message codes are as follows:

#### **E\_, I\_, or W\_**

E\_ (error message), I\_(informational message), or W\_ (warning message).

#### **FC**

Two-letter Facility Code (see page 198).

#### **xxxx**

4-digit hexadecimal number that represents the value returned to the application for use in error handling. The 2-letter facility code combined with the 4-digit hexadecimal number uniquely identifies this error code.

#### **\_msg**

"msg" is an abbreviated message text or decimal number. As a decimal, it is always equivalent in value to the preceding hexadecimal number.

### Message Help Files

Help is available online to aid in reading error messages. The error messages are stored in a single file `messages.txt` in the directory `$II_SYSTEM/ingres/files/english/messages/`.

## Fatal Errors

Fatal errors are those errors that require correction before Ingres can proceed with the program. All errors, including fatal errors, are recoverable.

Fatal errors occur in the following categories:

- Ingres tools, such as QBF or ABF
- Ingres server errors
  - Environmental OS/hardware limits and problems
  - A command limit is reached, for example, exceeding the maximum number of aggregates in a query
  - System software error

To remove the cause of a fatal error, use the troubleshooting techniques in the chapter Troubleshooting Ingres.

## Non-Fatal Errors

Non-fatal errors are diagnostic indications of user errors, usually in application development. For example, a non-fatal error code is returned when a user attempts to destroy a non-existent table. These errors are not severe and do not halt the user program or Ingres. User applications can choose to print or suppress these errors.

## Ingres Facility Codes

Ingres facility codes consist of a two-character code that is generated by a particular facility. For example, the Abstract Data Type Facility has a facility code of AD so error messages generated by this facility have an error message format of:

*E\_ADxxxx\_msg*

Error messages often include an explicit reference to which internal Ingres function returned the error. A list of these modules and their functions can help you identify the source of errors requiring troubleshooting.

## Facility Codes for Primary Components

The facility codes for the primary Ingres components, such as the DBMS Server and General Communications Facility, are shown in the following table:

Facility Code	Facility
AD	Abstract Data Type Facility (ADF)
DM	Data Manipulation Facility (DMF)
DU	Database Utility Facility (DUF)
GC	All GCF General Communication Facilities, including: Communications Server (GCC) General Communications Area (GCA) Name Server (GCN)
OP	Optimizer Facility (OPF)
PS	Parser Facility (PSF)
QE	Query Execution Facility (QEF)
QS	Query Storage Facility (QSF)
RD	Relation Description Facility (RDF)
SC	System Control Facility (SCF)
US	User Errors

## Server-Only Components

The server-only components are mostly DMF routines. The facilities listed in the following table are all subcomponents for Compatibility Library (CL) and do not appear in the facility-code part of an error message. All messages involving these facilities start with E\_CL, and usually contain the two-letter facility code in the abbreviated message part of the error text.

Facility Code	Name	Description	Used By
CS	Central System	Control of server	SCF, Utilities
DI	Database I/O	Access to Database files (different from SI)	DMF, Utilities
JF	Journal Files	Access to Journal files (different from DI)	DMF, Utilities

Facility Code	Name	Description	Used By
LG	Logging	Server logging (transaction log)	DMF
LK	Locking	DBMS data locking	DMF, Utilities
SR	Sorting	Sorting for queries	DMF
TR	Tracing	DBMS event tracing	Server components

### Tools-Only Components

The Ingres Tools-only components are shown in the following table:

Facility Code	Name	Description	Used By
AR	ABF Runtime	Application-By-Forms runtime component	ABF
DY	Dynamic Utilities	Dynamic linking	ABF
OL	4GL Support	Produce Ingres 4GL object code	ABF
PC	Process Control	Fork, sleep, exit, and so forth	Tools components
PE	Permissions	Object access authority	Tools components
TE	Terminal Driver	Terminal I/O	Tools terminal driver
UT	Utility Invocation	Compile, link, call, print, and so forth	Tools components

### Components for Stream File Management

Ingres components for stream file management are shown in the following table. The facilities listed below are all subcomponents for Compatibility Library (CL ) and do not appear in the facility-code part of an error message. All messages involving these facilities start with E\_CL, and usually contain the two-letter facility code in the abbreviated message part of the error text.

Facility Code	Name	Description	Used By
CP	Switch protections	Special-purpose file permissions handling	auditdb



Facility Code	Name	Description	Used By
LO	Locations	Abstract file system support (hierarchical)	Tools components
SI	Stream I/O	Stream functions	Tools components

### Components for Utility Routines

Ingres components for utility routines are shown in the following table:

Facility Code	Name	Description	Used By
CI	Authorization	Authorization string interpretation	Server and Tools components
ER	Error Messages	Error Message Handling/supports international	Server and Tools components
EX	Exception Handling	Establish routine chaining	All
GC	GCA CL	GCA (IPC) support	All
GV	Global Variables	Global symbols (Release ID, date, and so forth)	All
ID	Object IDs	User identifications	Usually Tools components
ME	Memory Allocation	Memory management	Usually Tools components
NM	Logical Symbols	Logical symbol management (or simulation)	All
QU	Queue Manipulation	Management of queues	All
TM	Timing	Timing services	All; some timing is server only

## Miscellaneous Components

Miscellaneous components are shown in the following table:

Facility Code	Name	Description	Used By
BT	Bit Manipulation	Utility routines to test, set, clear bits	All
CM	Character Manipulation	Character handling and for Kanji support	All
CV	Conversion	Convert internal representations	All
MH	Math	Functions (tan, ln, log, exp, and so forth)	All
ST	String Handling	Utilities for string operations	All

## Log Files

Ingres maintains log files to which it writes information about the installation activities.

### Transaction Log File

Each installation has a transaction log file and an optional dual log file. The log file holds information about all open transactions and is used to recover active databases after a system failure. Optionally, you can change its size and number of partitions at startup.

**UNIX:** The UNIX log file can be created as a raw partition or a number of raw partitions.

## Error Log

The main error log is a readable text file that you can use for troubleshooting. The error log file (errlog.log) is maintained in \$II\_SYSTEM/ingres/files/ directory.

Messages about the installation are appended to this log with the date and time at which the error occurred. This is generally the first place to look when troubleshooting a problem.

The error log contains the following information:

- Error messages
- Warning messages
- Archiver shutdown
- DBMS server start up and shutdown

The system administrator maintains the error log file. The file continues to grow until manually truncated. The installation must be shut down before truncating or removing the errlog.log file.

## Archiver Log

The archiver log contains information about the current archiver process. The archiver log file (liacp.log) is located in the \$II\_SYSTEM/ingres/files directory.

This file is appended to when the archiver process starts. The log contains the following information:

- Archiver start up
- Archiver error messages
- Archiver warning messages

## Recovery Log

The recovery log (dmfrcp) contains information about the current recovery process. The recovery log file (iircp.log) is located in the \$II\_SYSTEM/ingres/files/ directory.

This file is appended to when the recovery process starts. The log contains the following information:

- Current logging and locking parameter values
- Recovery process error messages
- Recovery process warning messages
- Recovery operations information

The recovery log must be monitored if you are unable to connect to Ingres and suspect that the DBMS Server is in recovery mode.

## Primary Configuration Log Files

Ingres maintains transcripts of various configuration operations. The primary configuration log files that Ingres uses are as follows:

### **config.log**

**The config.log file contains a log of the changes made with the Configuration Manager or Configuration-By-Forms (cbf) utility.**

### **rcpconfig.log**

The rcpconfig.log file contains log and error information of the last time rcpconfig was run. For details of the rcpconfig command, see the *Command Reference Guide*.

The primary configuration log files are located in the \$II\_SYSTEM/ingres/files/ directory.

## Optional Configuration Log Files

Configuration log files that are optionally maintained are follows. These files, if present, are in the directory indicated by II\_CONFIG:

### **iilink.log**

This log is used only in conjunction with the Object Management Extension. The file contains a log of the last time the iilink command was run. For details on the iilink command, see the *Object Management Extension User Guide*.

### **iivdb.log**

Contains a transcript of the last time verifydb was used to diagnose or attempt recovery of a damaged or inconsistent database. This file is created the first time verifydb is run

### **lartool.log**

Contains a transcript, with any errors, of the last time lartool was used to manually abort or commit a running transaction. This file is created the first time lartool is run.

## Other Optional Log Files

The following types of optional log files are present on your installation:

- Individual process logs. For the following processes you can set up a separate log file to isolate the error messages relating to that process:
  - DBMS
  - GCC

You normally do this only for a specific troubleshooting purpose. The default is that these messages are sent to the errlog.log. If you define one of these separate error logs, all messages are sent to *both* that file and errlog.log.

- Trace log for tracking messages at a greater level of detail
- Log for an optional Ingres facility

These optional log and trace log files can be established by setting the associated Ingres variables:

#### **DBMS Error log**

The DBMS error log, optionally defined as a separate file.

This log file is established by setting the Ingres variable `II_DBMS_LOG` to a user file name. For additional information, see the appendix "Environment Variables and Logicals." The default is all DBMS server errors and messages are sent to `errlog.log`.

#### **GCC error log**

The GCC error log, for installations using Ingres Net. This is in a user-defined file.

This log file is established by setting `II_GCCxx_LOG` to the full name of a file. The error logging level is specified by `II_GCCxx_LOG_LVL`. These Ingres environment variables/logicals are listed in the appendix "Environment Variables and Logicals."

#### **GCC trace log**

The GCC trace log is set up for specific troubleshooting efforts. You set `II_GCA_LOG` to a user file name. The associated Ingres variable `II_GCC_TRACE` defines the level of tracing. These Ingres environment variables/logicals are listed in the appendix "Environment Variables/Logicals." Details on using a GCC trace are described in User-Server Communications (see page 218).

#### **Star error log**

The Star error log for installations running Ingres Distributed Option, if defined.

This log file is established by setting the Ingres variable `II_STAR_LOG` to a user file name. For additional information, see the appendix "Environment Variables and Logicals." The default is all Distributed Option errors and messages are sent to `errlog.log`.

## Trace Utilities

Commands are available to trace such things as disk file access, locks, user interface-to-DBMS server communications, query plans of the Ingres Query Optimizer and various types of memory usage. For more information on tracing, see the *Database Administrator Guide*.

DBMS server tracing is generally enabled by some form of the set command, although some forms of tracing make use of Ingres environment variables and logicals. Set commands used specifically for debugging and troubleshooting are discussed in this document. More general set commands are documented in the *SQL Reference Guide*.

---

## ODBC Call-level Interface

The Ingres ODBC Call-level Interface (CLI) follows the existing convention of ODBC tracing as performed by the UNIX ODBC Driver Manager and Microsoft ODBC Driver Manager. Registry or configuration files are scanned for trace and trace log settings. Tracing output is similar to what is currently provided by the Ingres ODBC tracing DLL on Windows. Optionally, an application can set standard ODBC tracing using the following Ingres environment variables:

### **II\_ODBC\_LOG**

A string indicating the path and file name of an ODBC trace file. For example, the path and file name of an ODBC trace file for UNIX and Linux is `\tmp\odbc.log` and for Windows is `c:\temp\odbc.log`.

### **II\_ODBC\_TRACE**

A positive integer with a value of 1, 3, or 5. A setting of 1 provides standard ODBC tracing and is the most useful for debugging ODBC CLI applications. A setting of 3 includes ODBC function entry calls in the Ingres ODBC driver (as opposed to the ODBC CLI driver manager). A setting of 5 displays information about internal function entries in the Ingres ODBC driver.

If an ODBC application does not use the Ingres ODBC CLI, the ODBC trace settings of 3 and 5 are still recognized, but the setting of 1, which belongs to the driver manager component, is not. However, the existing tracing capability for third-party driver managers is unaffected. In this case, the driver manager tracing is written to the ODBC trace log as specified in the registry (Windows) or `odbcinst.ini` file (UNIX and Linux), but the detailed tracing information is written to the log file as specified by the `II_ODBC_LOG` environment variable.

## Set Statement

The set statement is used to specify a number of runtime options and has the general format:

```
set [no]option [additional parameters];
```

The *option* is set on or active by `set option` or `set option on` (depending on the option). It is turned off by a corresponding `set nooption` or `set option off`. Additional parameters are required, depending on the option. For a detailed syntax description of the set statement, see the *SQL Reference Guide*.

## Environment Variables and Logicals Commonly Used with Set

Set statements are executed as part of query language startup procedures at selected levels of scope. A set statement is entered directly at a terminal for temporary settings or executed using an Ingres environment variable/logical to establish a default option setting for all users. Several levels are shown in the following table:

<b>Where Specified</b>	<b>Scope of Effect</b>
ING_SYSTEM_SET	Executed whenever Ingres is started up, affecting all users
ING_SET	Executed whenever any Ingres tool connects to a server
ING_SET_DBNAME	Affects only Ingres tools starting up on the specified database
II_EMBED_SET	Similar to ING_SET, but diagnostic output is to a file
set statement entered directly	In effect until changed by a subsequent set statement for the option. This occurs by another direct set statement or using any of the Ingres environment variables and logicals above reestablishing the default.



### Example: Set Statement Entered Directly

One or more set options can be specified using the `_SET` environment Ingres environment variables and logicals. They can be assigned directly from the operating system shell. Separate multiple set statements by semicolons (up to a limit of 64 characters), as shown in the following examples:

**Windows:**

```
SET ING_SET=set nojournaling;set printqry
```

**UNIX:**

C Shell:

```
setenv ING_SET "set nojournaling; set printqry"
```

Bourne Shell:

```
ING_SET = "set nojournaling; set printqry" export ING_SET
```

**VMS:**

```
DEFINE ING_SET "SET NOJOURNALING; SET PRINTQRY"
```

## Example: Set Statements in an Include File

Set statements can also be implemented by means of an include statement. The include statement allows you to place the set statements in a file and specify the file name in the setenv (or equivalent) command. Use of the include option also avoids the 64-character line limit.

For example:

### **Windows:**

```
set ING_SET=include c:\extra\ingres\set.ing
```

### **UNIX:**

C Shell:

```
setenv ING_SET "include /extra/ingres/set.ing"
```

Bourne Shell:

```
ING_SET = "include /extra/ingres/set.ing"  
export ING_SET
```

### **VMS:**

```
DEFINE ING_SET -"INCLUDE DUA0:[EXTRA.INGRES]SET.ING"
```

The file *set.ing* includes the following statements:

```
set autocommit on;  
  set lockmode on mastertable  
  where level = page,  
  maxlocks = 10;
```

## Set Statements in Startup Files

Set statements are automatically executed as part of query processor start up. Various startup defaults are enabled either by setting an Ingres environment variable/logical (such as ING\_SET or II\_EMBED\_SET described above) or by including these set statements in one of the query processor startup files.

For an alphabetical listing of Ingres environment variables/logicals, see the appendix "Environment Variables and Logicals."

---

## Set Options for Tracing Queries

Set options of special interest in query troubleshooting include:

- [no]printqry: Set query display on or off.
- [no]qep: Set query execution plan display on or off.
- joinop [no]timeout: Set optimizer timeout on or off.

### The set printqry Option

The command **set printqry** prints out the query before it is optimized and executed. This is especially useful when evaluating queries for performance tuning or troubleshooting.

### Example: Turn Query Display On Using ING\_SET

The following command turns query display on:

#### Windows:

```
SET ING_SET=set printqry
```

#### UNIX:

C Shell:

```
setenv ING_SET "set printqry"
```

Bourne Shell:

```
ING_SET = "set printqry"  
export ING_SET
```

#### VMS:

```
DEFINE ING_SET "SET PRINTQRY"
```

As an example, assume that QBF is started up. The following is displayed:

```
Query Buffer:  
set Autocommit on  
Query Parameters:
```

```
Query Buffer:  
select cap_value, cap_capability from iidbcapabilities  
Query Parameters:
```

```
Query Buffer:  
select user_name, dba_name from iidbconstants  
Query Parameters:
```

### Example: Turn Query Display On Using II\_EMBED\_SET

The following command turns query display on through II\_EMBED\_SET. This is similar to the ING\_SET environment variable/logical. All information is gathered in the client application and printed to the log by the library routines linked to the client application:

**Windows:**

```
SET II_EMBED_SET=printqry
```

**UNIX:**

C Shell:

```
setenv II_EMBED_SET "printqry"
```

Bourne Shell:

```
II_EMBED_SET = "printqry"  
export II_EMBED_SET
```

**VMS:**

```
DEFINE II_EMBED_SET "PRINTQRY"
```

The query output with query timings is placed into a file called iiprtqry.log. Here is a sample output:

```
Query text:  
select * from iirelation  
Query Send Time: Thu Mar 26 17:20:43 2001  
Query Response Time: Thu Mar 26 17:20:44 2001  
Response CPU Time: 2630  
Qry End CPU Time: 3370
```

This type of query monitoring can be useful for spotting slow-running queries in applications.

## The set qep Option

The **set qep** option provides a display of the optimizer's query execution plan (QEP) for the query after it has been optimized but before it is executed.

The following command turns QEP display on as the default ING\_SET level:

### Windows:

```
SET ING_SET=set qep
```

### UNIX:

C Shell:

```
setenv ING_SET "set qep"
```

Bourne Shell:

```
ING_SET = "set qep"  
export ING_SET
```

### VMS:

```
DEFINE ING_SET "SET QEP"
```

The Ingres Query Optimizer and QEPs are described in detail in the *Database Administrator Guide*.

## The set joinop notimeout Option

The **set joinop notimeout** option can be used in tracing query performance. This statement turns the optimizer timeout feature off. With timeout on, the optimizer stops checking for further query execution plans when it believes that the best plan it has found takes less time to execute than the amount of time already spent searching for a plan. If you issue a **set joinop notimeout** statement, the optimizer continues searching query plans. This option is often used with the **set qep** option to ensure that the optimizer is picking the best possible query plan.

## Canceling Set Options

To cancel any of these options that have been set, you issue the opposite set statement (set nooption to turn an option off or set *option* in the case of joinop notimeout to restore the default behavior).

## The set session with on\_error Option

The following set statement enables you to specify how transaction errors are handled in the current session:

```
set session with on_error = rollback  
statement| transaction
```

To direct Ingres to roll back the effects of the entire current transaction if an error occurs, specify `rollback transaction`. To direct Ingres to roll back only the current statement (the default), specify `rollback statement`. To determine the current status of transaction error handling, issue the `select dbmsinfo('on_error_state')` statement.

Specifying `rollback transaction` reduces logging overhead, and help performance; the performance gain is offset by the fact that, if an error occurs, the entire transaction is rolled back, not the single statement that caused the error.

The following errors always roll back the current transaction, regardless of the current transaction error-handling setting:

- Deadlock
- Forced abort
- Lock quota exceeded

To determine if a transaction was aborted as the result of a database statement error, issue the statement `select dbmsinfo('on_error_state')`. If the error aborted the transaction, this statement returns 0, indicating that the application is currently not in a transaction.

You cannot issue the `set session with on_error` statement from within a database procedure or multi-statement transaction.

## I/O Tracing

The set option `io_trace` prints out information about disk I/O during the life of each query.

The following command turns I/O trace on as the default `ING_SET` level:

### Windows:

```
SET ING_SET=set io_trace
```

### UNIX:

C Shell:

```
setenv ING_SET "set io_trace"
```

Bourne Shell:

```
ING_SET = "set io_trace"
export ING_SET
```

### VMS:

```
DEFINE ING_SET "SET IO_TRACE"
```

For example, given the query:

```
select * from iirelation;
```

Here is a sample output from the I/O trace. The counts are the number of pages read/written:

```
*****
I/O READ File: aaaaaaac.t00 (iidbdb, iirel_idx, 13) count:1
*****
I/O READ File: aaaaaaab.t00 (iidbdb, iirelation, 0) count:8
*****
I/O READ File: aaaaaaab.t00 (iidbdb, iirelation, 8) count:8
*****
I/O READ File: aaaaaaab.t00 (iidbdb, iirelation, 16) count:7
```

**Note:** When tracing the I/O or the locks of a parallel query (using `set io_trace` or `set lock_trace` with `set parallel n`), the trace messages from child threads of the QEP are logged to the `II_DBMS_LOG`. The trace messages for the main thread are sent to the user session in the normal manner.

## Lock Tracing

Set lock\_trace prints out information about locks within a transaction.

The following command turns lock tracing on as the default ING\_SET level:

### **Windows:**

```
set ING_SET=set lock_trace
```

### **UNIX:**

C Shell:

```
setenv ING_SET "set lock_trace"
```

Bourne Shell:

```
ING_SET = "set lock_trace"export ING_SET
```

### **VMS:**

```
DEFINE ING_SET "SET LOCK_TRACE"
```



## Event Tracing

The following features enable your application to display and trace events:

- To enable or disable the display of event trace information for an application when it raises an event, use the following command:
- To enable the display of events as they are raised by the application, specify `set printdbevents`. To disable the display of events, specify `set noprintdbevents`.
- To enable or disable the logging of raised events to the installation log file, use the following statement:

```
set [no]printdbevents
```

```
set [no]logdbevents
```

- To enable the logging of events as they are raised by the application, specify `set logdbevents`. To disable the logging of events, specify `set nologdbevents`.
- To enable or disable the display of events as they are received by an application, use the following statement:

```
set_sql(dbeventdisplay = 1 | 0)
```

Specify a value of 1 to enable the display of received events, or 0 to disable the display of received events. You can also enable this feature using `II_EMBED_SET`.

You can create a routine that traps all events returned to an embedded SQL application. To enable/disable an event-handling routine or function, your embedded SQL application must issue the following `set_sql` statement:

```
exec sql set_sql(dbeventhandler = event_routine | 0)
```

To trap events to your event-handling routine, specify *event\_routine* as a pointer to your error-handling function. For information about specifying pointers to functions, see your host language companion guide. In addition to issuing the `set_sql` statement, you must create the event-handling routine, declare it, and link it with your application.

## User-Server Communications

The General Communication Facility (GCF) is composed of the GCA protocol and three communications programs: Name Server, Communications Server, and Data Access Server. These are separate processes. You can listen in on the communications between these GCF programs using the following Ingres environment variables and logicals:

- II\_EMBED\_SET "printgca"
- II\_GCA\_TRACE
- II\_GCN\_TRACE
- II\_GCC\_TRACE
- II\_GCD\_TRACE
- II\_GC\_TRACE
- II\_GCCCL\_TRACE
- II\_GCA\_LOG

### Trace Communication Using printgca

You can trace communications occurring in the GCA with the following command:

**Windows:**

```
set II_EMBED_SET=printgca
```

**UNIX:**

C Shell:

```
setenv II_EMBED_SET "printgca"
```

Bourne Shell:

```
II_EMBED_SET = "printgca"  
export II_EMBED_SET
```

**VMS:**

```
DEFINE II_EMBED_SET "PRINTGCA"
```

This traces all GCA messages passed between Ingres tools and the server. The output is placed into a file called `iiprtgca.log`.

## Output of GCA Trace

Here is a sample output:

```
printgca = on session -2 (Thu Mar 26 15:50:32 2001)
```

```
GCA Service GCA_SEND
gca_association_id: 0
gca_message_type: GCA_QUERY
gca_data_length: 49
gca_end_of_data: TRUE
gca_data_area: GCA_Q_DATA
gca_language_id: 2
gca_query_modifier: 0x01
gca_qddata [0]:
gca_type: 51: DB_QTXT_TYPE
gca_precision: 0
gca_i_value: 29
gca_value: select name, gid from iuser
```

End of GCA Message

```
GCA Service: GCA_RECEIVE
gca_association_id: 0
gca_message_type: GCA_TDESCR
gca_data_length: 55
gca_end_of_data: TRUE
gca_data_area: GCA_TD_DATA
gca_tsize: 26
gca_result_modifier: 0x01
gca_id_tdscr: 8
gca_i_col_desc: 2
gca_col_desc[0]:
gca_attdbv:
db_data: 0x00
```

```
db_length: 24
db_datatype: 20: DB_CHA_TYPE
db_prec: 0
gca_i_attname: 4
gca_attname: (4) name
gca_col_desc[1]
gca_attdbv:
db_data: 0x00
db_length: 2
db_datatype: 30: DB_INT_TYPE
db_prec: 0
gca_i_attname: 3
gca_attname: (3) gid
```

End of GCA Message

```
GCA Service: GCA_RECEIVE
gca_association_id: 0
gca_message_type: GCA_TUPLES
gca_data_length: 338
gca_end_of_data: TRUE
gca_data_area: GCA_TU_DATA not traced
```

End of GCA Message

### Example: Level 2 Tracing Using GCA Trace Utility

The syntax of GCA trace utility are as follows:

```
define II_GCA_TRACE [n]  
define II_GCA_LOG [filename]
```

II\_GC*x*\_TRACE defines the level of GCA tracing, with  $n = 1$  (lowest) through 4 (most detailed). II\_GCA\_LOG names the destination file of the trace output. If not specified, it defaults to standard output.

Here is an example of level 2 tracing:

```
!GCcm: target_id: 9062  
!GCrequest 0: connecting on 9062  
!GC_exchange 0: status 00000000  
!GC_exchange_sm 0: status 00000000 state 0  
!GCsend 0: send 528  
!GC_send_comp 0: sent 528 status 00000000  
!GC_exchange_sm 0: status 00000000 state 1  
!GCreceive 0: want 528  
!GC_recv_colmp 0: status 000000 state 0  
!GCsend 0: send 65  
!GC_send_comp 0: sent 65 status 00000000  
!GCreceive 0: want 24
```

## UNIX Trace Facilities

The UNIX operating system has various trace facilities that are useful in troubleshooting Ingres.

## Bourne Shell -x Option

Invoking a Bourne shell with the `-x` option causes each shell command to be printed before it is executed. This is especially useful for diagnosing problems with installation shell scripts.

To use this trace on the Ingres shutdown script, call the Bourne shell (or Korn shell) with the `-x` option, passing the name `ingstop` as the script to be executed:

```
sh -x ingstop
```

This option can only be used with UNIX shell scripts, not with binary executables.

If you are in doubt as to whether a program or command is a binary or shell script, type:

```
file program_name
```

UNIX replies with "commands text," "ASCII text," or "executable shell script" if it is an executable shell script. Otherwise the file command returns a message specific to your machine—something like demand paged executable and possibly the designated chip on which it was compiled to run.

## UNIX Trace Command

Some UNIX platforms (such as Sun) provide a facility to trace system calls as they are made by a process. To invoke it, enter `trace` at the operating system prompt.

## Create an Ingres Console for Error Monitoring in UNIX

To ensure that any new errors written to the Ingres error log file are displayed on the terminal screen, issue the following command at the terminal.

```
tail -f $II_SYSTEM/ingres/files/errlog.log
```

To continuously monitor the activities of the Ingres recovery process (`dmfrcp`), issue this command:

```
tail -f $II_SYSTEM/ingres/files/iircp.log
```



# Chapter 8: Managing Your System and Monitoring Performance

---

This chapter discusses the visual tools that Ingres provides to manage and monitor Ingres performance, including:

- Visual Manager
- Visual Performance Manager
- Visual Configuration Difference Analyzer

**Note:** These GUI tools can be used on Windows platforms, or on UNIX platforms with the MainWin application by Mainsoft. To find out which UNIX platforms are supported, see the Ingres Corporation web site.

## Visual Manager

Ingres Visual Manager (IVM) provides a global view into the Ingres installation. It serves as a system console from which you can manage Ingres components and access other utilities. This utility captures events that are occurring in the system and allows them to be filtered for emphasis, based on your preferences.

IVM allows you to monitor (and start and stop) the different servers in the installation (DBMS, Name, Communications, JDBC, Data Access, Bridge, Star, ICE, and Recovery) as well as Remote Command, the Logging and Locking systems, and the Primary Transaction Log and Archiver process. It shows events occurring in the system both at the installation level and the level of each server. Using IVM, you can also edit the Ingres system and user environment variable settings.

## Functions of Visual Manager

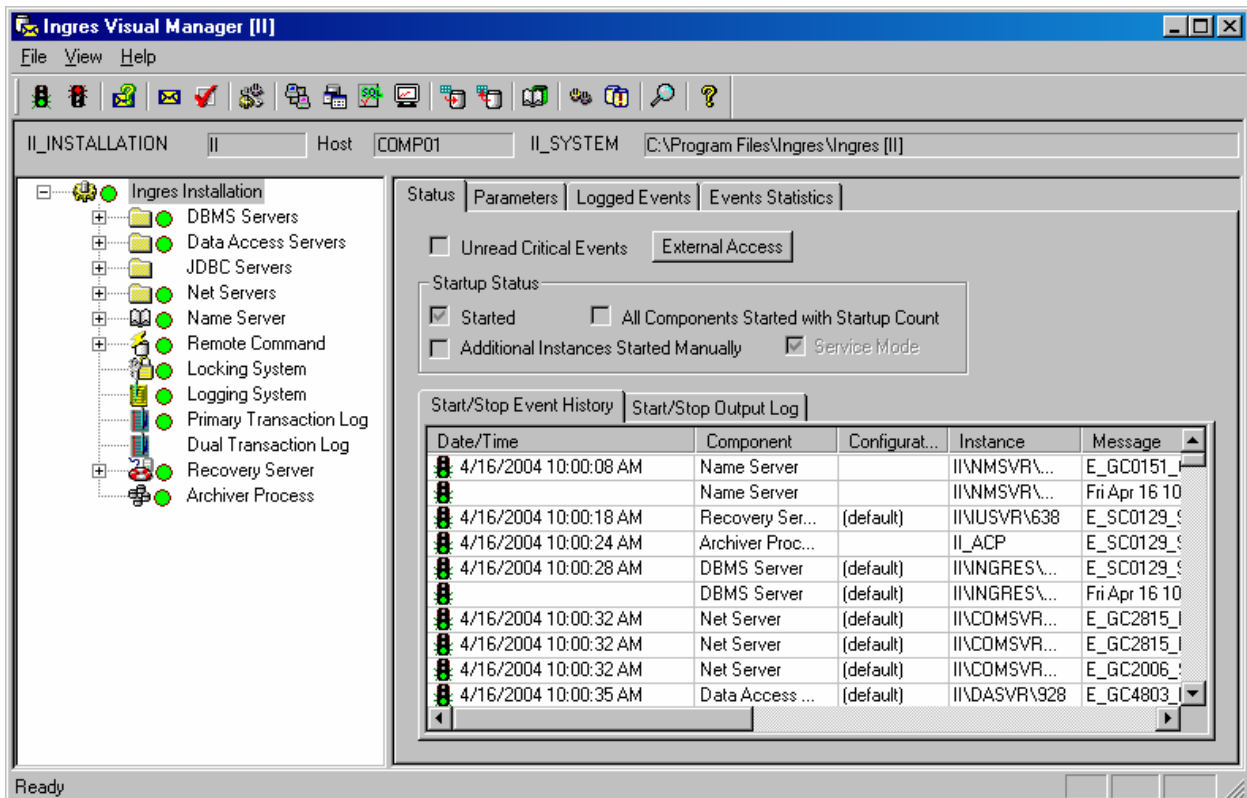
Ingres Visual Manager enables you to:

- View the started and stopped components and servers in the installation as well as the history of start/stop events
- Start or stop the entire installation or individual components
- View and edit Ingres system and user environment variables
- Control how various types of Ingres error log messages are handled
- Set preferences for event monitoring and notification
- View and filter events that occurred both at the installation level and the component or server level
- View statistics on these events
- View explanations for error log messages
- Access all Ingres visual tools and online documentation



## Visual Manager Window

The Ingres Visual Manager window contains a list of Ingres components. By selecting any branch of the corresponding tree, you can access information related to the branch, which correspond either to the entire Ingres installation (the root branch), an individual component configuration, or an individual instance of a given configuration.



The Visual Manager Window contains the following tabs:

### Status

Displays the summary of startup status of the components of the installation, the start and stop history of any component in the configuration, and related information such as the output of the last start or stop operation.

### Parameters

Displays system, user, or extra parameters.

### **Logged Events**

Displays a list of logged events for the Ingres installation and any of its sub-branches. You can also select an event to view its corresponding message text.

### **Event Statistics**

Displays statistics on events that have occurred for the selected component, grouped per event ID, and presented as a numerical count and in graph form.

For more information about these pages, see online help for Ingres Visual Manager.

## **System and User Parameter Configuration Through IVM**

The Parameters page under the Ingres Installation branch contains lists of parameters for system, user, and extra parameters along with their values and descriptions. You can add or unset a parameter and edit its value. You can also choose whether you want to view Ingres parameters that are not currently set.

For more information, see the online help topic, Parameters Page, Ingres Installation branch (Ingres Visual Manager window).

## **Set Parameter Configuration Through IVM**

From IVM, you can access the Configuration Manager utility to configure parameters for each server component in the Ingres installation. To access Configuration Manager, select the desired server component in the left pane of the Ingres Visual Manager window and click the Configuration Manager toolbar button, or choose File, Configure.

For more information, see the online help topic, Configuring Ingres Components in the Ingres Visual Manager online help.

## Message and Notification Management Through IVM

Using IVM, you can control how various types of Ingres error log messages are handled. The Define Message Categories and Notification Levels dialog lets you specify which messages in the `errlog.log` must be discarded, which are displayed normally in IVM, and which trigger a special alert. This dialog also allows you to group messages according to user-defined categories, which can be handled as a group for discard, display or alert.

To access the Define Message Categories and Notification Levels dialog, click the Categories and Notification Levels toolbar button, or choose File, Categories and Notification Levels. For information on using this dialog, see the Ingres Visual Manager online help.

## Event Monitoring Management Through IVM

Using the Preferences dialog in IVM, you can specify the following settings related to event monitoring.

- Maximum amount of memory used to store events (in megabytes)
- Maximum number of events stored
- Maximum number of days that the event is stored before being deleted
- Only messages that arrived after the last Name Server startup are stored

To access the Preferences dialog, click the Preferences toolbar button, or choose File, Preferences.

## Alert Events Notification Settings in IVM

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.

### **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. If this option is selected, you can also set preferences for generating specific Operating System events.

To access the Preferences dialog, click the Preferences toolbar button, or choose File, Preferences. For instructions on setting these additional preferences, see the Preferences dialog topic in the Ingres Visual Manager online help.

## View Message Explanations in IVM

IVM allows you to view an explanation for any Ingres message in the errlog.log file. To view an explanation for a message in the errlog.log file:

1. Select the desired message on any page or window in IVM that allows you to select messages.
2. View the corresponding explanation by clicking the Message Explanation toolbar button.

For details about the Message Explanation window, see the Ingres Visual Manager online help.

## Component Monitoring Through IVM

You can monitor the following components through IVM

- Servers
- Logging System
- Locking System

## Server Monitoring Through IVM

For each of the following server components as well as their instances, you can use Ingres Visual Manager to monitor status, logged events, and event statistics:

- Bridge Servers
- DBMS Servers
- JDBC Servers
- Data Access Servers
- Communications Servers
- Star Servers
- ICE Servers
- Name Servers
- Remote Command
- Recovery Server—a Recovery Log File page appears in the IVM window
- Archiver Process—an Archiver Log File page appears in the IVM window

For more information on monitoring server components, see online help for Ingres Visual Manager.

## Logging System Monitoring Through IVM

The Logged Events page displays a list of the logged events for the Logging System branch, and any of its sub-branches. However, the messages are filtered according to certain preferences you define in the Preferences dialog.

For more information, see online help for Ingres Visual Manager.

## Locking System Monitoring Through IVM

The Logged Events page displays a list of the logged events for the Locking System. However, the messages are filtered according to certain preferences you define in the Preferences dialog. For more information, see online help for Ingres Visual Manager.

## Access to Visual Tools and Documentation Through IVM


IVM provides direct access to all Ingres visual tools and online documentation. (The only exception is the Ingres Service Manager, whose functionality is already provided by IVM.)

You can access Ingres tools and documentation in the following ways:

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



- Select a tool from the IVM toolbar:



**Note:** Only the most frequently used tools are accessible from this toolbar.

## Visual Performance Monitor

Visual Performance Monitor can be used as a monitoring tool, a performance analysis tool, and a system management tool.

Visual Performance Monitor lets you monitor the following:

- Servers

You can view a list of the servers that are started on an Ingres installation, and information about each server. For example, you can view the sessions that are currently active for each server. You can also remove a session or stop a server (if you are a privileged user).

Visual Performance Monitor provides the following information for a particular server:

- Lock lists
- Locks
- Transactions
- Locked databases, tables, and pages
- Other locked resources
- Network traffic for Net servers

- Users and sessions

You can monitor users for whom there are open sessions. You can find out which sessions are open for a particular user, and drill down further to reveal all the related information about those sessions (lock lists, transactions locked, databases locked, and so on).

- Logging

You can view logging system summaries, transaction lists, process, and database lists. Log information can be used to monitor transaction rates, log file activity, processes, and databases in the logging system. This information is useful in determining which logging parameters need to be adjusted.

- Locking

You can monitor lock information to help determine which lock parameters need to be adjusted. Viewing locking system summaries, lock lists, and resources provides you with the information you need to spot conditions when, for example, additional locking system resources need to be added.

Viewing your lock lists is useful for locating transactions that cannot proceed because they are blocked by another transaction.

- System performance

You can view your performance information from a “database” point of view. This means you can access performance information using a database branch, as opposed to the root branches of Visual Performance Monitor.

- Replication

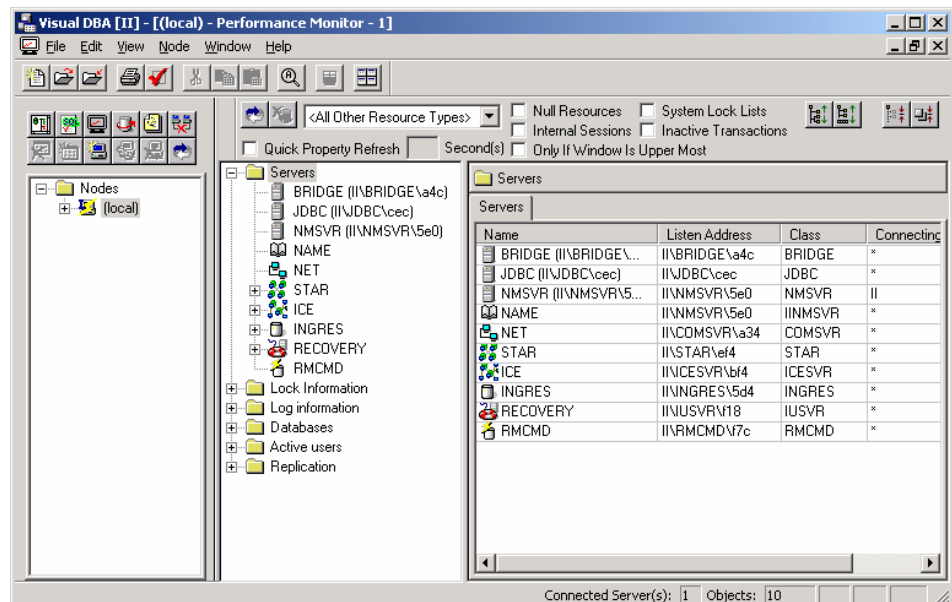
You can start, stop, and monitor Replicator servers that are required for the replication scheme that has been defined in a VDBA DOM window. You can set up startup parameters for these servers, send events to these servers, view and manage collisions, and display other miscellaneous replication monitor information.

Visual Performance Monitor lets you perform the following actions:

- Refresh data in the window from a server
- Shut down a database server or close a session

## Visual Performance Monitor Window

The Visual Performance Monitor window contains a list of categories under which various types of Ingres performance information appears. Under each root object category are branches representing each performance entity. Under each branch are one or more sub-branches that pertain to the particular type of performance entity.



For more information, see the Visual Performance Monitor online help. For context-sensitive help on any active dialog or window, press F1.



## Visual Configuration Differences Analyzer

The Visual Configuration Differences Analyzer (VCDA) allows you to compare the configurations for an Ingres local installation by taking “snapshots” of its configuration. For example, after you install Ingres and have set up the database schemas, take a snapshot of the configuration. If you encounter problems later on, take another snapshot of the configuration and compare it to the earlier snapshot to see if any configuration changes have contributed to the problem. Keep an on-going record of configuration changes by starting with a snapshot of a new installation that shows all system defaults, and taking a snapshot of the installation each time the configuration is changed.

**Note:** VCDA compares the configurations of Ingres installations, but not the database object definitions within these installations. To compare database objects, see the Visual Database Objects Differences Analyzer (VDDA), in the *Database Administrator Guide* and the Ingres online help.

Specifically, VCDA allows you to:

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

For specific instructions on performing these tasks, refer to the VCDA online help.

### Configuration Snapshot File

When you take a snapshot of the current configuration for the local installation, VCDA writes the following information to a text file:

- Main Parameters: The value of II\_SYSTEM, II\_CLUSTER, II\_CONFIG\_LOCAL, II\_NETBIOS\_NAME, hostname, and username
- All information included in the config.dat file, and in the additional config.dat files pointed to by II\_CONFIG\_LOCAL (if any)
- All vnode information, except for password data
- All Ingres system variables
- All Ingres user variables

This snapshot can be compared to other snapshots as explained below.

## Comparison of Configuration Snapshots

VCDA allows you to compare the current configuration of the local installation with a previously saved configuration snapshot. It also allows you to compare two saved snapshots that do not include the current configuration. VCDA also indicates, using an icon color, if the difference is due to a parameter that is different in both snapshots, or a parameter that exists in one snapshot but not in the other.

If there are no differences in the main parameters, VCDA displays a green icon and a message. If there are no differences in the other information, VCDA displays only a green icon.

## Comparison of Remote Installation Snapshots

Although you can only take snapshots of the local installation, you can compare configurations on remote installations (even on different platforms) by making the snapshot file from one machine available to another machine that is running VCDA. The snapshot file is a simple text file that is readable across different platforms.

## Restoration of Configuration Parameters Through VCDA

In addition to comparing configuration snapshots, VCDA allows you to restore selected groups of configuration parameters from a saved snapshot. The only parameters that VCDA does not restore are the values of `II_SYSTEM`, the hostname and username.

You restore parameters by first selecting the snapshot file that contains the desired parameters. You can specify the specific parameters you want to restore. If the snapshot file you choose corresponds to a networked installation, you can also specify the value of the corresponding Ingres variables and the content of the relevant additional `config.dat` file, with an option to create the tree structure if it does not exist.

VCDA automatically creates a backup snapshot file of the current configuration before it restores parameters. If for some reason you do not want to keep the restored parameters, you can select this backup snapshot file to restore the parameters to their previous state.

## How VCDA Handles Concatenated config.dat Files

VCDA uses information from the config.dat file, which normally contains parameters that apply to the local hostname 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 hostname under which a snapshot is saved shall be saved within the snapshot.
- If VCDA detects that hostnames other than the snapshot one are managed within the config.dat information of the snapshot, VCDA shall display it.

A “hostname mapping” option is available in that situation, so that VCDA can compare the additional hostnames configuration parameters. If this option is not used, the parameters are compared including their “hostname” part, (that is, only parameters that are identical for the same hostname within the two snapshots are considered as identical).



# Chapter 9: Analyzing and Recovering Journalled Transactions

---

This chapter discusses journalled transactions and how to analyze, recover, and redo them using the Journal Analyzer visual tool.

## Journal Analyzer

The Journal Analyzer utility allows you to view and analyze journalled transactions and individual underlying statements. It also enables you to create SQL scripts, which are used to recover or redo individual row operations without redoing or rolling back the whole database or table.

For the releases of Ingres that support the Journal Analyzer, see <http://ca.com>.

In the Journal Analyzer, you can:

- View past transactions and their begin/end timestamps, duration, LSN, and user, and the number of insert, update, and delete statements within the transaction
- View tables and rows involved in the transaction
- View before and after images and corresponding LSN for row changes within the transaction
- View rolled back transactions
- View some types of statements, such as create table or create index, inserts, updates, and deletes, and recover or redo journalled insert, update, and delete statements
- Recover and redo transactions immediately or generate scripts to do it at another time or in another database
- Recover and redo multiple transactions at one time, allowing, for example, the replaying of whole sequences of journal files on a remote node/database
- Recover and redo individual row changes

For detailed instructions on using the Journal Analyzer, see the Journal Analyzer online help.

## Start Journal Analyzer

To start the Journal Analyzer:

### Windows:

Do **one** of the following:

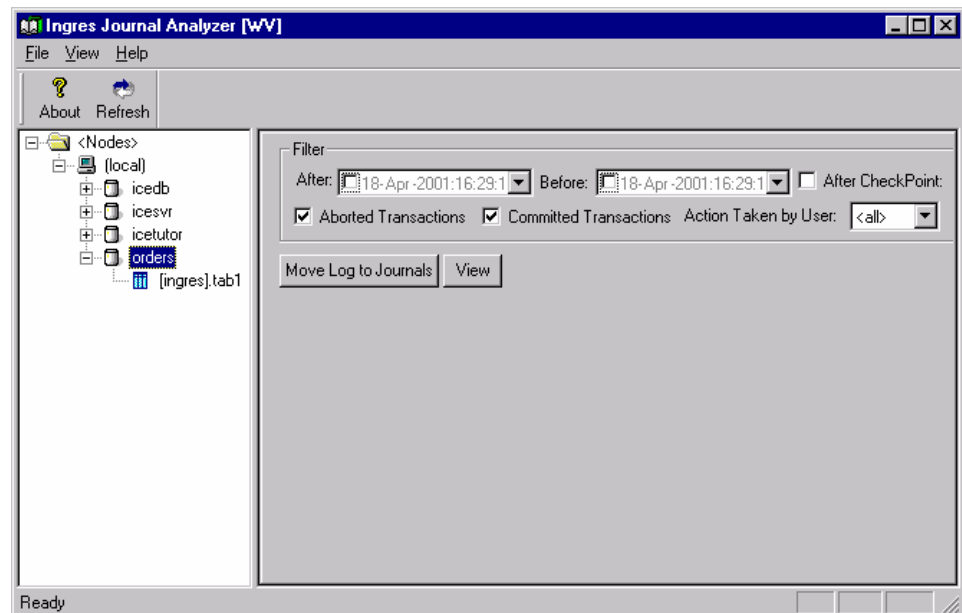
- Click Start on the taskbar and choose Programs, Ingres, Ingres Journal Analyzer.
- From within Ingres Visual Manager's right-click popup menu, select Journal Analyzer, or click the Journal Analyzer toolbar button.
- Type **ija** at the command prompt.

### UNIX:

- Type **ija** at the command prompt.

## Journal Analyzer Window

In the Journal Analyzer window, the left pane consists of the vnode, databases, and tables. In the left pane, you can drill down to the desired vnode, database, and (optionally) table. When a database or table has been selected, the right pane appears as follows:



Controls in the right pane are as follows:

**Filter**

Use the Filter options to control the quantity and type of transactions that are displayed, based on different criteria such as the transaction date and time, checkpoint number, user, and so on.

**View**

Clicking the View button displays transaction information that is currently in the journals. You must click the View button to refresh the data when you change the filter criteria.

**Move Logs to Journals**

To see any recent transactions that are in the transaction log file—but have not yet moved to the journals—click Move Log to Journals.

## Transaction Views

Based on whether you select a database or table in the navigational pane of the Journal Analyzer, and whether the Group by Transaction option is selected, you see a slightly different interface in the right pane.

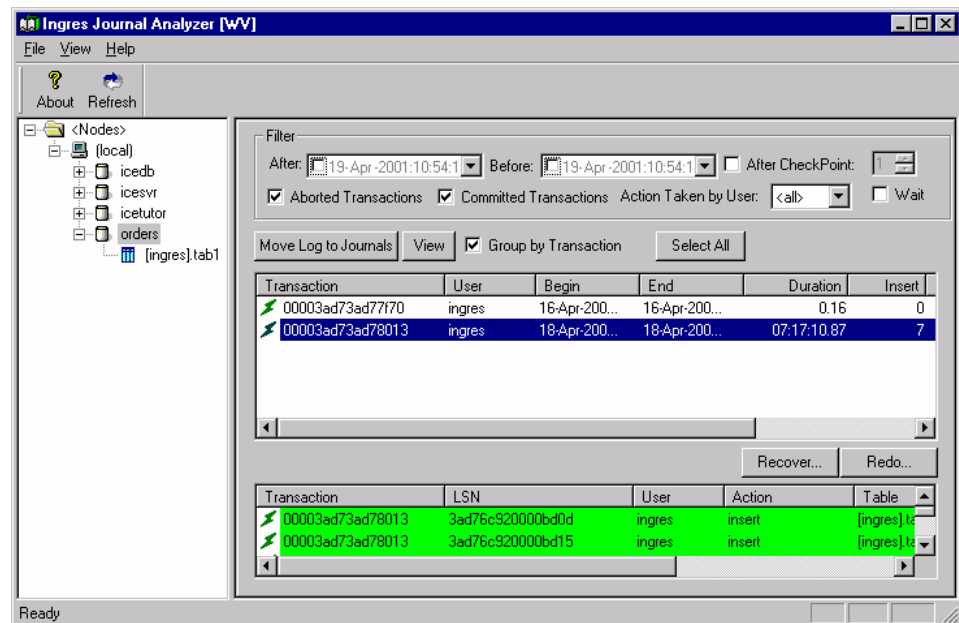
### Database Level Journalized Transactions

When you select a database, you see the journalized transactions within the database that match the criteria defined in the top of the pane. You can then determine whether you want to group transactions and view the details for only selected transactions.

## Group Changes by Transaction

When the Group by Transaction option is selected, a single line of “summary” information is displayed for each transaction in the uppermost list, sorted by transaction number. When one or more of these lines are selected, the detailed information for all selected lines is shown in the lower list. The detail information corresponds to each row of journaled tables altered in the transaction(s), in addition to statements such as create table or alter table that were executed within the transaction.

For example, in the following window, a transaction is selected and the corresponding detail lines are displayed beneath it:

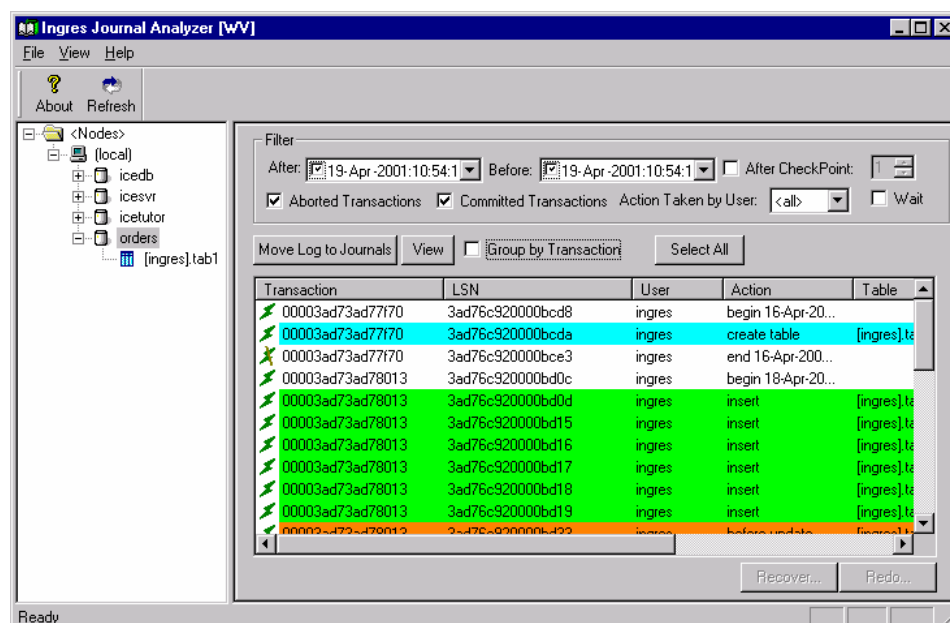


Although the rows are sorted by transaction number initially, look at (and sort by) transaction end (End column), because it is the time when all changes for the transaction become available to other users. A transaction with a higher number (that is, having started later) than another transaction, can be committed first, in which case the underlying changes become available to other users before those of the other transaction (which had a smaller number because it had started first). For this reason, sort on transaction end (by clicking on the corresponding header).



## View Transaction Details

If the Group by Transaction option is not selected, a single list appears, displaying all the individual actions for each transaction, in transaction order:



Even these detail lines provide a summary of information on the individual actions. This list cannot display column names because different tables have different column names. To get the full detail of individual rows that have changed, you must choose a table in the navigational pane. See Table Level Journalled Transactions (see page 242).

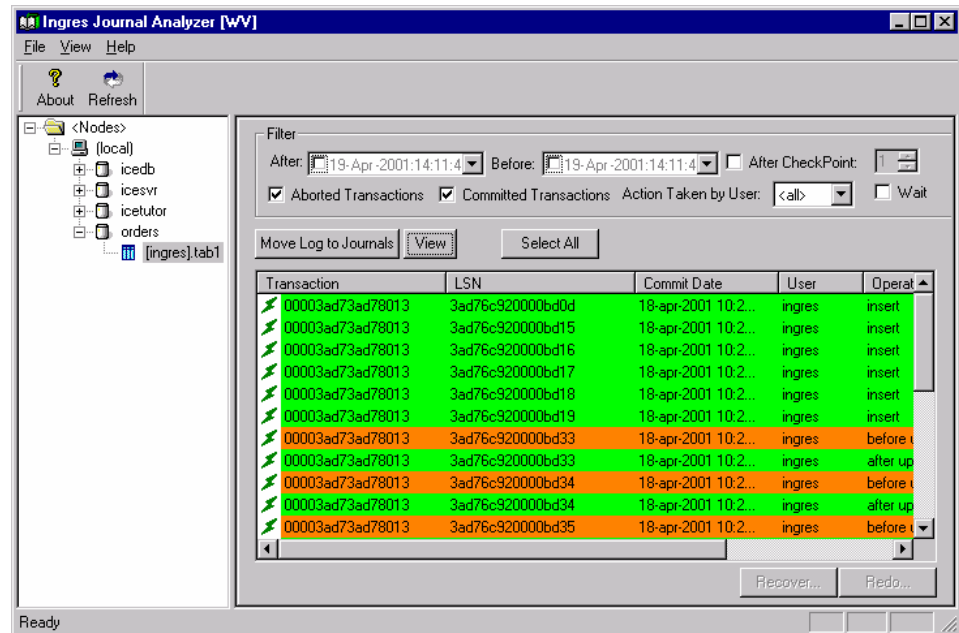
For update statements, the “before update” and “after update” images are shown. For inserts and deletes, only one row (the inserted or deleted one) appears. Different colored rows are used to denote the type of statement (in addition to the Action column that displays strings such as begin, end, insert, update, delete, and so on).

An icon preceding each line shows you whether the transaction it belongs to was committed or rolled back, included a partial rollback, or involved changes in non-journalled tables (an exclamation point next to the icon appears in this case).

You can sort the data by any criteria by simply clicking on the corresponding column header.

## Table Level Journalled Transactions

On the table level, the right pane looks and functions similar to the one at the database level, but only the row changes corresponding to the selected table appear. There is no Group by Transaction check box for the table level. At the table level, you can view a history of journalled changes in that table.



As opposed to the database level panes, where more than one table was involved, it is possible here to have the real column names become headers of the results list, which was not possible at the database level because the columns are normally different in each table that was involved in the transaction.

You can also sort the data on the table level by clicking on a column header. In particular, the following sorts are useful:

- Commit date—this is the most meaningful criteria for sorting the changes
- Transaction number—be aware that the transactions are not always committed in the same order
- LSN—shows you the exact sequence of row changes

For more information on the interface for this utility, see the Journal Analyzer online help topics.

## Select Transactions and Row Changes

After browsing for the transactions or row changes that you want to recover or redo, you can select one or more transactions or individual row changes. If the Group by Transaction check box is selected, and multiple transaction lines are selected, the detail list shows the row changes corresponding to all selected transactions. If this option is not selected, you are selecting lines at the table level, which correspond to row changes, not full transactions. The Recover and Redo dialogs proposes (but does not pre-select) the other changes in those transactions of the initial row changes that were selected.

**Tip:** Select whole transactions associated with the changes to be recovered to ensure transaction integrity. This is the preferred and most common way to recover row changes.

After making your selections, the Recover and Redo buttons are available. If you do select lines corresponding to only *part* of a transaction, there are other factors to consider, as described in Recover Transactions or Individual Row Changes (see page 244) and Redo Transactions and Individual Row Changes Option (see page 251).

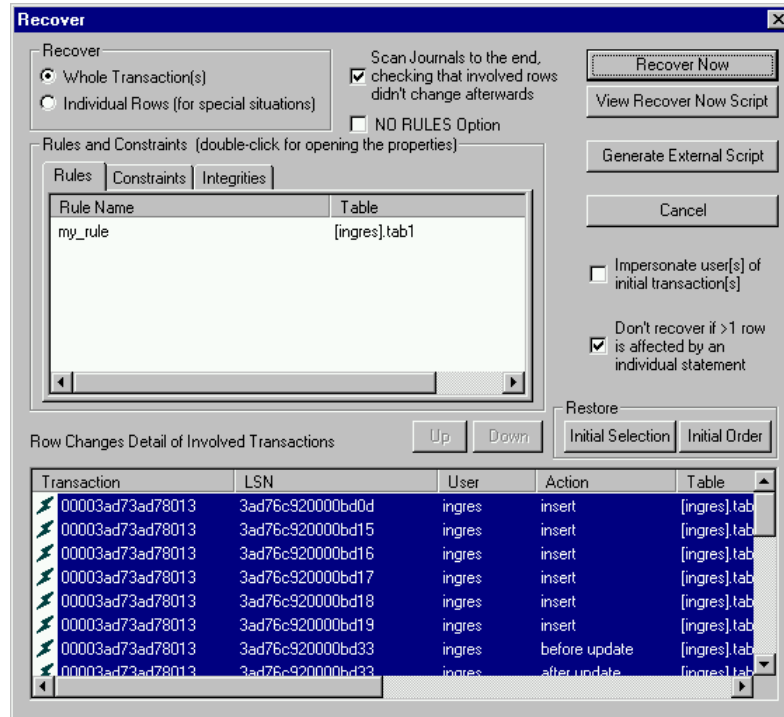
The Select All button allows you to select all transactions currently displayed in the pane for the selected criteria. This can be useful for example for replaying whole portions of journals on another node/database, through the Redo functionality described in the Redo Transactions or Individual Row Changes (see page 251).

## Recover Transactions or Individual Row Changes

The Recover dialog is used to create one of two types of “undo scripts”—an internal script that is run immediately or an external script that can be saved and run later.

To recover transactions, follow these steps:

1. When you click Recover, the Recover dialog appears as follows:



Initially, the row changes that you selected in the Journal Analyzer are displayed in the lower portion of the Recover dialog.

2. Click "Recover Now" to recover immediately, or "Generate External Scripts" to generate scripts that can be run later, as described in Recover Immediately or Generate External Scripts Option (see page 249).

## Recover “Whole” Transactions or Individual Rows

If you chose whole transactions in the Journal Analyzer, the Whole Transaction(s) option is selected automatically in the Recover dialog.

If you selected individual row changes that make up only part of a transaction (in the lower portion of the dialog), the Individual Rows option is selected. In this case, you are prompted for confirmation as to whether you want to recover only parts of transactions, and lose transaction integrity. Although the recovery of individual rows is not recommended, it can be useful, for example, if a table has Ingres rules associated with it.

## Rules, Constraints, and Integrities

If a table has Ingres rules associated with it that results in additional row changes within the same transaction, do not generate undo operations for such dependent changes because it leads to undesirable results. Also, the rules can change between the time of the initial transaction and the time when you want to undo that transaction. Instead, undo only the change in the base table, and let the rules generate the dependent changes.

Journal Analyzer detects the rules, constraints, and integrities that exist in the tables involved in the transaction(s) and displays them on the tabbed pages in the Recover dialog. From this display, you can determine whether some changes in a transaction resulted from rules and if certain constraints cannot be fulfilled. This information helps you determine if you want to use the No Rules option, and possibly disable certain rows within the transaction. You can view the properties for a particular rule, constraint, or integrity by double clicking it. Examine these entities to determine their impact on a restore.

**Note:** Integrity constraints that are violated are not specifically reported: updates and inserts that violate any integrity constraints are simply not performed.

## Scan Journals to End Option

The Scan Journals to the End option provides a warning if the tables have been altered since the time of the initial transaction, in which case the rules for these tables have changed.

If there are rules on a table, you have two choices:

- Disable individual row changes in the undo scripts, by deselecting the corresponding row changes in the Recover dialog. This can be useful, for example, if the current rules cause the desired operation in dependent tables, in which case you must disable the dependent row changes.
- Use the No Rules option to allow you to disable rules while executing the queries. This option allows the reversal of the exact row changes of the initial transaction, regardless of what the rules were at that time, and what they are now. The side effect of this option is that if there are constraints on the corresponding tables, *they are not checked if this option is used*. This operation must be used very carefully. If the Journal Analyzer previously detected current constraints on those tables, it provides a warning and asks for confirmation before starting the operation.

## Order of Transactions and Statements in Transactions

The order of the undoing of transactions is reverse of the order at commit time of these transactions and cannot be changed. If for some reason you need the order to be different than the usual order, the only way to do this is to recover each transaction separately (by launching the Recover dialog for each transaction individually) in the desired order.

In each transaction to be recovered, the order of the individual statements can be changed, using the Up and Down buttons in the Recover dialog.

## How to Check If Rows Have Changed After the Transaction

Before proposing to reverse row changes resulting from a transaction, you must take into account the fact that some of these rows have changed after the transaction. The possible cases are as follows:

- None of these rows have changed and there is no problem.

Or:

- If a row has changed, usually no row corresponding to the “after image” of the initial change can be found, and therefore the “reverse” statement, which includes a where clause that searches for a row corresponding to such “after image,” fails because no such row is found.

There is an “advanced” exception, in the case where a given row (the same or another one) has the exact same content afterwards as the after image of the initial change: in this case, it is the user’s responsibility to make a decision about taking additional actions.

For this reason, the following features are available:

- The changes that you have asked to reverse are displayed in the dialog for you to view.
- By default, the Scan Journals to the End option is selected. It indicates if later changes affected the rows involved in the transaction. If this problem is not an issue, disable the option.
- If rows of the transaction to be undone have changed since the transaction, Journal Analyzer prompts you with a warning, asking whether you want to continue.

## Users of Transactions

By default, Journal Analyzer proposes to recover transactions by executing reverse transactions under the currently connected user. However, the Impersonate User of Initial Transaction option allows you to impersonate the user under which each transaction had initially been executed.

**Note:** This option requires that the connected user is the database administrator or has appropriate privileges.

The Impersonate User of Initial Transaction option affects the way the reverse transactions are executed: that is, normally Journal Analyzer recovers multiple transactions by executing the recover statements of all transactions to be undone in a single, global “recover” transaction. This allows, in case of failure at any point in the recover operation, to cancel the whole operation.

However, given that it is not possible to execute in the same transaction different statements under different users, if this option is used, and the initial transactions had not all be executed under the same user, Journal Analyzer must commit (and change the user) before each reverse statement to be executed under a different user than the previous one.

A possible problem in this case is when a reverse statement fails, once other recover statements have already been committed. In this situation, Journal Analyzer proposes to “undo the undo” of the transactions previously undone. However in case of failure of this last operation, you need to manually “repair” this situation (typically by using Journal Analyzer) by completing or canceling the uncompleted operation.



## Number of Rows Affected by Each Individual “Reverse Statement”

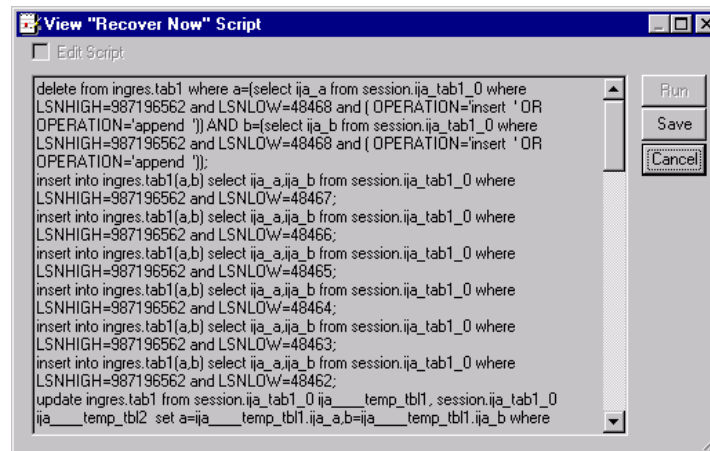
Journal Analyzer-generated reverse statements for insert and update statements are delete and update statements, with a “where” clause on all columns that can be accepted in a where clause (column types such as “long varchar” cannot).

If you need two identical rows in which all columns can be part of a where clause (SQL does not normally distinguished such rows), the Recover dialog has an option (selected by default) called “Don’t Recover if >1 Row is Affected by an Individual Statement”. Only when this option is explicitly disabled does the recover process work even if more than one row is affected by an individual recover statement.

In the case where no row is affected by an individual “recover” statement, Journal Analyzer provides an error anyway, and the recover operation fails. If the Scan Journals to End option has not been disabled, Journal Analyzer normally warns you that rows involved already have changed since the initial transaction (different warnings appear if journaling has been disabled after the initial transaction or in similar situations).

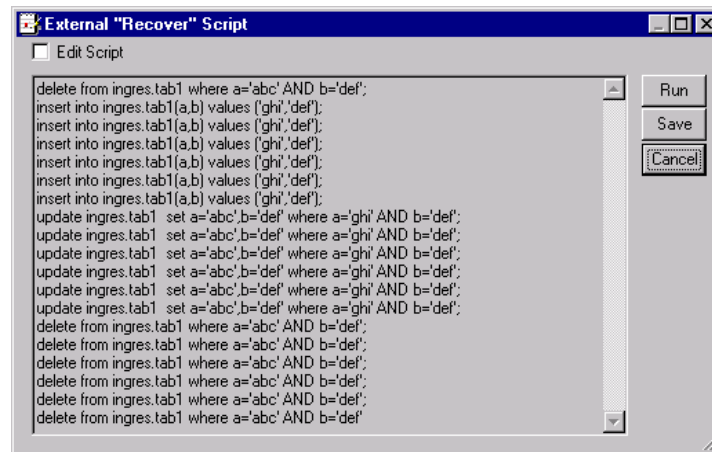
## Recover Immediately or Generate External Scripts Option

If you want to view the script that is run based on your Recover dialog selections, click the View Recover Script Now button. The following dialog appears:



This dialog allows you to see the SQL statements that are executed if you click the Recover Now button. The script only works when executed within Journal Analyzer and must not be used externally. If desired, you can also save the script by clicking Save.

If you want to generate an external undo SQL script that you can run at another time (possibly on another database), you can do so after you have made your selections in the Recover dialog. By clicking the Generate External Script button, the following dialog appears:



You can make any necessary changes to the script by selecting the Edit Script option. To save or run the script, clicking the corresponding button.

Although external scripts can be very useful, you should be aware of some issues:

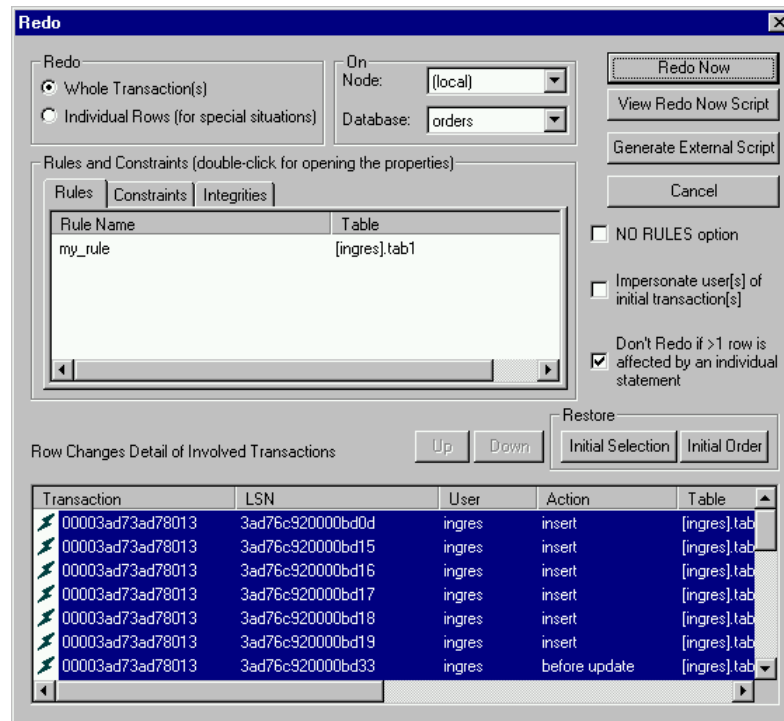
- Binary data cannot be conveniently inserted into an SQL script
- Floating point numbers can be rounded if inserted into a script in text format

For this reason, the Generate External Script options must be restricted for the Recover and Redo features to when absolutely necessary (that is, typically when executing the scripts on another node and database). In this case:

- There is an error message if some data types included in rows of the transaction cannot be inserted in the script
- The user must be aware of the context of possible rules and constraints existing on the target database at the time when the script is executed and accordingly update the script.

## Redo Transactions and Individual Row Changes

From the Journal Analyzer window, you can also launch the Redo dialog, as follows:



This dialog is useful for the following purposes:

- For redoing rolled back transactions (that had been rolled back for an unwanted reason for example)
- For regenerating similar, committed transactions in another database

This dialog is similar to the Recover dialog; however, two additional controls allow you to specify a different node and database where to redo the transaction or statement. All the features described in the Recovering Transactions section for the Recover dialog apply to this dialog and function in the same way—only for redoing transactions. For more information, see that section and the help system.



# Chapter 10: Understanding Ingres Management Architecture

---

This chapter describes the overall features of Ingres Management Architecture (IMA), shows examples of how it is used from SQL, and notes restrictions that application developers should know about.

## Ingres Management Architecture (IMA)

The Ingres Management Architecture (IMA) provides the framework to access system data for monitoring and managing installations:

### **Monitoring**

The monitoring facilities include the system data currently reported through specialized utilities such as the Ingres Visual Manager, Visual Performance Monitor, Interactive Performance Monitor (ipm), lockstat, logstat, and iimonitor.

### **Managing**

Active management operations include stopping servers and manipulating sessions in DBMS and Star servers.

With IMA you can create system management applications using standard SQL tools. With Ingres Net, the IMA provides access to multiple installations, allowing management of complex environments with a single SQL application. The IMA is also designed to be compatible with standard management protocols, such as SNMP (Simple Network Management Protocol), DME (Distributed Management Environment), and CMIP (Common Management Interface Protocol).

IMA provides its services without affecting the underlying operation of the product.

## IMA Audience

IMA provides the means of creating sophisticated DBMS management applications for users and third-party vendors. These capabilities are of interest to advanced DBAs, system administrators, and developers producing tools for use by DBAs and system administrators.


The user of this product must be familiar with Ingres tools such as OpenROAD, Ingres Terminal Monitor, and VDBA and have an overall understanding of system monitoring capacities and purposes.

## IMA Architecture

With IMA, all the installations in an administrative domain can be managed by a collection of pure SQL applications that can be written mainly with OpenROAD:

- **OpenROADApplication**—the monitoring program written in standard SQL.
- **Tables in DBMS**—catalog tables representing the run state of various servers and global parameters from all of the installations of interest.
- **IMA**—provides access to objects defined in the DBMS Server and in other processes.
- **Server Processes**—the DBMS Server and other processes being monitored. With Ingres Net, these can be on different vnode installations.

The following illustration shows the IMA gateway and its connection to an OpenROAD application:



## Management Information Base

The data in the Management Information Base catalog tables form a Management Information Base (MIB). The available data is defined using a simple object model. The MIB is made visible to management applications through the IMA facility, which provides SQL access to the data in the MIB.

## Data and the Object Model

The MIB is defined as an abstract entity consisting of simply typed management objects:

### Simply Typed

Simply Typed means that objects have values that are strings or strings that can be interpreted as integers, and do not represent aggregates, such as structures.

### Aggregate Objects

*Aggregate* objects, such as a C language structure, are represented in the MIB by having a separate object type for each element. (This can be compared to having a single table for each potential column in a relational database, with the entire table made visible with a join of all the column tables.)

In the model each object has a name (*classid*), location, and one or more instances.

## Classid

Each data object in the MIB has a name, called the *classid*, which defines its type and semantics. Classids are strings that form a hierarchical name space by the convention of using periods (.) to separate levels of the tree.

For example, consider the following classid:

```
exp.scf.scs.scb_query
```

This object name identifies an experimental object, owned by scf, that is scs related, and that is the scb\_query. In the current release all objects are under the exp experimental branch of the name space tree.



## Location of Data Objects

Each data object is located in some place. If it is:

- Process-specific, the place is a particular server
- Visible in several different processes in an installation, its place is the vnode in which the installation is located

The place must be local, or remotely accessed using Ingres Net. Installation passwords must be working. The general format of the place string is:

```
[vnode:] [/@iigcn | /@gca_address]
```

where:

*gca\_address* is the listen address of the server as shown in Configuration Manager (vcbf), Interactive Performance Monitor (ipm) or Name Server Maintenance Utility (iinamu).

## Number of Instances of a Classid

None, one, or more objects of a particular classid can exist in a process or a vnode. Each has a corresponding instance value. This value is always a string, although it is possible in many instances to interpret the string as an integer value.

A table is modeled by having objects with instance values that form a key index.

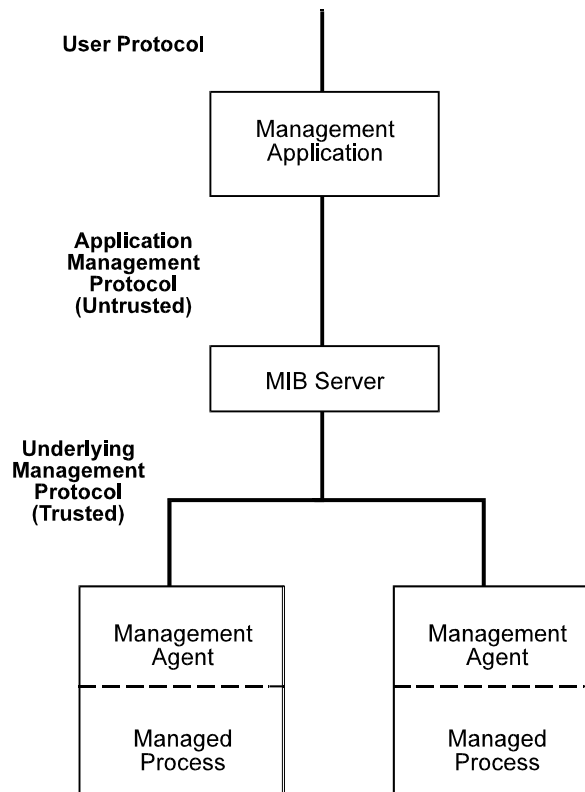
## Data Keys

All data objects in the MIB have simple data types, can have many different instances, and are located either in a server process or in a vnode. Thus, they are always addressed with a three-part key:

```
{ [vnode_or_server,] classid, instance }
```

## Components of the IMA

This following is a reference model for the IMA.



The components of the model are described below, with the current IMA implementation noted:

### User Protocol

The *User Protocol* is forms or other appropriate interface. A Management Application program provides the user protocol.

### MIB Server

The application obtains MIB data through an *Application Management Protocol*, such as SQL (or alternative protocol such as SNMP). The current IMA implementation is SQL-based.

- A *MIB Server* process accepts the application management protocol.
- The MIB server communicates with the managed processes by an *Underlying Management Protocol*.

Thus, managed processes need to talk to only a single protocol, while MIB servers can be provided to support management applications using a variety of protocols.

## IMA Management Protocol

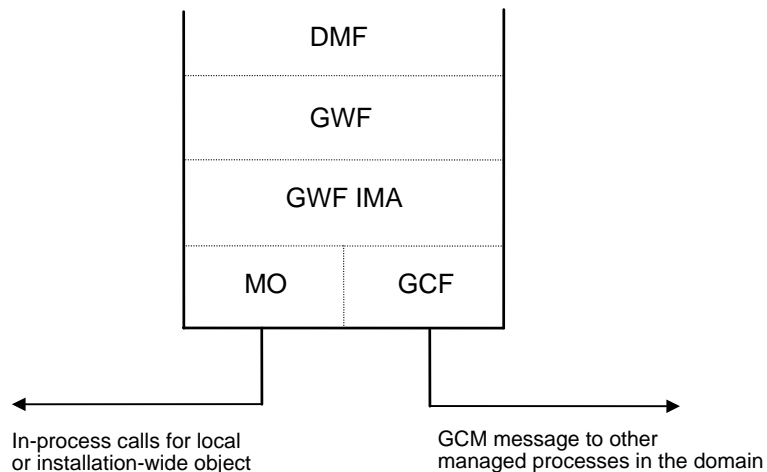
The IMA defines its own underlying management protocol for use between servers. The IMA protocol is built on top of the regular GCA protocol. Ingres Net provides IMA management of other installations transparently.

## IMA Server

The IMA Server maintains the SQL MIB data. It uses a non-relational gateway to present the management objects in the domain as SQL tables:

- For objects internal to the DBMS Server, the IMA Server reads and writes objects directly.
- Operations on objects that are in other processes are handled using the GCA management protocol.

The following illustration shows the IMA component in the DBMS Server, including the Data Management Facility (DMF), the Enterprise Access facility (GWF), the IMA Facility (IMA GWF), and the IMA's Management Object (MO) module for handling local objects, as well as the GCF facility for dealing with other managed processes.



An Ingres DBMS Server is needed for users to access the MIB data with an SQL application. Without an Ingres DBMS Server, no management data is available to non-privileged users. It is possible for privileged applications to communicate with servers using the GCM directly.

The architecture permits Enterprise Access products from other management protocols to be written using either SQL or the underlying GCA management protocol directly.

## SQL MIB Creation

The IMA's SQL MIB is the binding of the MIB object space into an SQL name space. Tables containing MIB objects can be registered into any database owned by the \$ingres user. The binding itself is done by the IMA, which converts between rows and individual MIB objects.

## IMA Tables Registration

IMA tables must be added explicitly by the \$ingres user in any database owned by \$ingres to define an SQL schema for a MIB database. The SQL register table as import statement is used to register IMA tables. IMA registered tables can be opened only in databases owned by \$ingres.

IMA catalogs must be created, if they do not already exist. Only the iidbdb has these catalogs.

## Register Table Statement

The register table as import statement has the following syntax:

```
register table tablename
(col_name format [is 'ext_format'] {, col_name format [is 'ext_format']})
as import from 'source'
with dbms = ima
[, structure = none | [unique] sortkeyed]
[, key = (col_name [asc] {, col_name})]
[, rows = nnnnnnnn] [, [no]duplicates] [, [no]update]
[, [no]journaling]
;
```

where:

### ***tablename***

Is the name of the table. This must follow the standard naming conventions. This is a required parameter.

### ***col\_name***

Is the name of the column. This must follow the standard naming conventions. This is a required parameter.

### **format**

Is the data type of the column. This is a required parameter.

### **is '*ext\_format*'**

Specifies the management object being referenced, by extended format string name. This is either:

- Classid: A reserved name identifying metadata:
- PLACE: Column is any kind of place, char type.
- VNODE: Column is vnode places only char type.
- SERVER: Column is server places only char type.
- CLASSID: Column is the classid name char type.
- INSTANCE: Column is the instance value char type.
- VALUE: Column is the char value of the {place, classid, instance} object.
- PERMISSIONS: Column is the integer management object permission mask for the classid.

### **as import**

Indicates that you are registering a gateway product. This is a required parameter.

**from 'source'**

Indicates the source is objects or tables. This parameter is ignored by the IMA. You must enclose the value of *source* in single quotation marks. This is a required field.

**with**

Introduces additional information about the table being imported. This is a required field.

**dbms= ima**

Indicates that you are registering with the IMA. This is a required field.

**structure=**

Specifies whether the table is keyed or not. If omitted, no key is assumed. For IMA tables the only valid structure is sortkeyed or unique sortkeyed, which must be specified. This is a required field.

**key=**

Specifies the column name(s) to use as keys. For the IMA flat table, the key must be : ([place,] classid, instance). For IMA cross-tab tables, the key must contain a column with the index classid, which can be preceded with a column defining the place. This parameter is required if you specified sortkeyed as the key type.

**rows=nnnnnnnn**

Specifies the approximate number of rows in the table. If omitted, a default of 1000 is used. This is an optional parameter.

**[no]duplicates**

Indicates whether the table can contain duplicate rows. If this parameter is omitted, duplicates will be considered. This parameter is optional.

**[no]update**

Indicates whether the table can be updated. If this parameter is omitted, the default is noupdate (read-only). For IMA tables, the flat object table can be updateable. All cross-tab tables must be non-updateable. This is an optional parameter.

**[no]journaling**

Indicates whether journaling is allowed. For IMA tables, only nojournaling is allowed. This is an optional parameter.

## IMA Table Types

There are two types of IMA tables:

- Flat object Tables
- Cross Tab Tables

## Flat Tables

Flat tables have one object instance per row. Only flat object tables can be updated. This is a sample flat object table:

<b>server</b>	<b>classid</b>	<b>instance</b>	<b>value</b>
here::@/1001	exp.table.index	01	01
here::@/1001	exp.table.index	02	02
here::@/1001	exp.table.user	01	gammerl
here::@/1001	exp.table.user	02	dommage
here::@/1001	exp.table.value	01	88
here::@/1001	exp.table.value	01	99

Flat tables must be registered as in the following example:

```
register table ima_mib_objects (
    server varchar(64) not null not default is 'SERVER',
    classid varchar(64) not null not default is 'CLASSID',
    instance varchar(64) not null not default is 'INSTANCE',value
    varchar(64) not null not default is 'VALUE',
    perms integer2 not null not default is 'PERMISSIONS'
)
as import from 'objects'
with update,
dbms = IMA,
structure = unique sortkeyed,
key = (server, classid, instance);
```

## Cross-tab Tables

Cross-tab tables have one object per column, per row. A cross-tab version of the same data as above is much more convenient to query, as shown below:

<b>index</b>	<b>user</b>	<b>value</b>
01	gammerl	88
02	dommage	99

A cross-tab table has 1... $n$  columns, including at least one that is the index. A place column is optional. If present, it is identified by having an "is clause" of VNODE or SERVER. If it is omitted, the place defaults to the current server. All other columns are assumed to be management object classids, sharing common index instance values.

Cross-tab tables must be registered in the following format. The columns in the registration can be in any order, but must be sortkeyed on the place, if it is present, and the index object. Here is an example of IMA cross-tab table registration syntax:

```
register table ima_dbms_servers (  
    server varchar(64) not null not default is  
        'SERVER',  
    listen_address varchar(64) not null not default is  
        'exp.gcf.gca.client.listen_address',  
    max_connections integer4 not null not default is  
        'exp.scf.scd.server.max_connections',  
    num_connections integer4 not null not default is  
        'exp.scf.scd.server.current_connections',  
    server_pid integer4 not null not default is  
        'exp.clf  
  
.unix  
  
.cs.srv_block.pid')  
as import from 'tables'  
with dbms = IMA,  
structure = sortkeyed,  
key = (server);
```

## Register Table Examples

The subsections below give examples of table registrations for the IMA MIB. These tables are used in Query Examples to IMA Tables in this chapter that show examples of SQL access to IMA.

The objects and contents of all registered tables are subject to enforcement of appropriate access restrictions.



### Example: Flat Object Table

A flat object table is required to be able to update any object. Objects that can be updated are control objects.

```
register table ima_objects (  
  server varchar(64) not null not default  
    is 'SERVER',  
  classid varchar(64) not null not default  
    is 'CLASSID',  
  instance varchar(64) not null not default  
    is 'INSTANCE',  
  value varchar(64) not null not default  
    is 'VALUE'  
)  
as import from 'objects'  
with dbms = IMA,  
  update,  
  structure = unique sortkeyed,  
  key = (server, classid, instance);
```

### Example: Object Metadata Table

This table provides metadata about the known objects in the domain. Objects appear to be supported by many servers:

```
register table ima_meta (  
  server varchar(64) not null not default  
    is 'SERVER',  
  classid varchar(64) not null not default  
    is 'exp.glf.mo.meta.classid',  
  class varchar(64) not null not default  
    is 'exp.glf.mo.meta.class',  
  perms i2 not null not default  
    is 'exp.glf.mo.meta.perms',  
  size i2 not null not default  
    is 'exp.glf.mo.meta.size',  
  xindex varchar(64) not null not default  
    is 'exp.glf.mo.meta.index'  
)  
as import from 'tables'  
with dbms = IMA,  
  structure = unique sortkeyed,  
  key = (server, classid);
```

## Examples: Domain Tables

This table shows the domain of servers and vnodes visible to the session:

```
register table ima_session_domain (  
  domplace varchar(64) not null not default  
    is 'exp.gwf.gwm.session.dom.index'  
)  
as import from 'tables'  
with dbms = IMA,  
  structure = unique sortkeyed,  
  key = (domplace);
```

The following example shows all the servers known to the DBMS in all the vnodes in all the user domains:

```
register table ima_places (  
  place varchar(64) not null not default  
    is 'exp.gwf.gwm.places.index',  
  type i4 not null not default  
    is 'exp.gwf.gwm.places.type',  
  class varchar(64) not null not default  
    is 'exp.gwf.gwm.servers.class',  
  flags i4 not null not default  
    is 'exp.gwf.gwm.servers.flags',  
)  
as import from 'tables'  
with dbms = IMA,  
  structure = unique sortkeyed,  
  key = (place);
```

## Example: Table of SCF Sessions

The following table shows information about SCF-level sessions:

```
register table ima_server_sessions (  
    server varchar(64) not null not default is  
        'SERVER',  
    session_id varchar(32) not null not default is  
        'exp.scf.scs.scb_index',  
    effective_user varchar(20) not null not default is  
        'exp.scf.scs.scb_euser',  
    real_user varchar(20) not null not default is  
        'exp.scf.scs.scb_ruser',  
    db_name varchar(12) not null not default is  
        'exp.scf.scs.scb_database',  
    db_owner varchar(8) not null not default is  
        'exp.scf.scs.scb_downer',  
    db_lock varchar(9) not null not default is  
        'exp.scf.scs.scb_dblockmode',  
    server_facility varchar(10) not null not default is  
        'exp.scf.scs.scb_facility_name',  
    session_activity varchar(50) not null not default is  
        'exp.scf.scs.scb_activity',  
    activity_detail varchar(50) not null not default is  
        'exp.scf.scs.scb_act_detail',  
    session_query varchar(1000) not null not default is  
        'exp.scf.scs.scb_query',  
    session_terminal varchar(12) not null not default is  
        'exp.scf.scs.scb_terminal',  
    session_group varchar(8) not null not default is  
        'exp.scf.scs.scb_group',  
    session_role varchar(8) not null not default is  
        'exp.scf.scs.scb_role',  
    server_pid integer4 not null not default is  
        'exp.scf.scs.scb_pid'  
)  
as import from 'tables'  
with dbms = IMA, structure = unique sortkeyed,  
key = (server, session_id)
```

## Example: Lock Tables

The following tables allow you to determine which sessions are waiting for locks and which locks they are.

**Note:** Because these lock tables are rapidly changing, they are particularly susceptible to dirty read problems. Fortunately, you are more often going to be looking for the one that is not changing, because a session is stuck on that particular lock.

```
register table ima_locklists (  
    vnode varchar(64) not null not default is  
        'VNODE',  
    locklist_id_id integer4 not null not default is  
        'exp.dmf.lk.llb_id.id_id',  
    locklist_lkb_count integer4 not null not default is  
        'exp.dmf.lk.llb_lkb_count',  
    locklist_status varchar(50) not null not default is  
        'exp.dmf.lk.llb_status',  
    locklist_lock_count integer4 not null not default is  
        'exp.dmf.lk.llb_llkb_count',  
    locklist_max_locks integer4 not null not default is  
        'exp.dmf.lk.llb_max_lkb',  
    locklist_wait_id_id integer4 not null not default is  
        'exp.dmf.lk.llb_wait_id_id',  
    locklist_type integer4 not null not default is  
        'exp.dmf.lk.llb_name0',  
    locklist_database integer4 not null not default is  
        'exp.dmf.lk.llb_name1',  
    locklist_server_pid integer4 not null not default is  
        'exp.dmf.lk.llb_pid',  
    locklist_session_id varchar(32) not null not default is  
        'exp.dmf.lk.llb_sid'  
  
)  
as import from 'tables'  
with dbms = IMA,  
structure = unique sortkeyed,  
key = (vnode, locklist_id_id);  
  
register table ima_locks (  
    vnode varchar(64) not null not default is  
        'VNODE',  
    lock_id_id integer4 not null not default is  
        'exp.dmf.lk.lkb_id.id_id',  
    lock_request_mode varchar(10) not null not default is  
        'exp.dmf.lk.lkb_request_mode',  
    lock_grant_mode varchar(10) not null not default is  
        'exp.dmf.lk.lkb_grant_mode',  
    lock_state varchar(100) not null not default is  
        'exp.dmf.lk.lkb_state',  
    lock_attributes varchar(100) not null not default is  
        'exp.dmf.lk.lkb_attribute',  
    resource_id_id integer4 not null not default is  
        'exp.dmf.lk.lkb_rsb_id_id',  
    locklist_id_id integer4 not null not default is  
        'exp.dmf.lk.lkb_llb.id_id'  
  
)  
as import from 'tables'  
with dbms = IMA,  
structure = unique sortkeyed,  
key = (vnode, lock_id_id);
```

In the lock table, *state* 1 == granted, 2 == convert, 3 == waiting:

```
register table ima_resources (
  vnode varchar(64) not null not default is
    'VNODE',
  resource_id_id integer4 not null not default is
    'exp.dmf.lk.rsb_id.id_id',
  resource_grant_mode varchar(4) not null not default is
    'exp.dmf.lk.rsb_grant_mode',
  resource_key varchar(50) not null not default is
    'exp.dmf.lk.rsb_name',
  resource_type integer4 not null not default is
    'exp.dmf.lk.rsb_name0',
  resource_database_id integer4 not null not default is
    'exp.dmf.lk.rsb_name1',
  resource_table_id integer4 not null not default is
    'exp.dmf.lk.rsb_name2',
  resource_index_id integer4 not null not default is
    'exp.dmf.lk.rsb_name3',
  resource_page_number integer4 not null not default is
    'exp.dmf.lk.rsb_name4'
)
as import from 'tables'
with dbms = IMA,
structure = unique sortkeyed,
key = (vnode, resource_id_id);
```

## Example: Logging Tables

The following logging tables give a view into the active transactions on the system. Like the locking tables, they are subject to dirty read inconsistency:

```
register table ima_log_processes (  
    vnode varchar(64) not null not default is  
        'VNODE',  
    process_id_id integer4 not null not default is  
        'exp.dmf.lg.lpb_id.id_id',  
    process_id_instance integer4 not null not default is  
        'exp.dmf.lg.lpb_id.id_instance',  
    process_status varchar(100) not null not default is  
        'exp.dmf.lg.lpb_status',  
    process_pid integer4 not null not default is  
        'exp.dmf.lg.lpb_pid'  
)  
as import from 'tables'  
with dbms = IMA,  
structure = unique sortkeyed,  
key = (vnode, process_id_id);  
  
register table ima_log_databases (  
    vnode varchar(64) is  
        'VNODE',  
    db_id_id integer4 not null not default is  
        'exp.dmf.lg.ldb_id.id_id',  
    db_id_instance integer4 not null not default is  
        'exp.dmf.lg.ldb_id.id_instance',  
    db_status varchar(100) not null not default is  
        'exp.dmf.lg.ldb_status',  
    db_database_id integer4 not null not default is  
        'exp.dmf.lg.ldb_database_id',  
    db_name varchar(32) not null not default is  
        'exp.dmf.lg.ldb_db_name',  
    db_owner varchar(32) not null not default is  
        'exp.dmf.lg.ldb_db_owner'  
)  
as import from 'tables'  
with dbms = IMA,  
structure = unique sortkeyed,  
key = (vnode, db_id_id);  
  
register table ima_log_transactions (  
    vnode varchar(64) not null not default is  
        'VNODE',  
    tx_id_id integer4 not null not default is  
        'exp.dmf.lg.lxb_id.id_id',  
    tx_id_instance integer4 not null not default is  
        'exp.dmf.lg.lxb_id.id_instance',
```

```
tx_status varchar(100) not null not default is
'exp.dmf.lg.lxb_status',
tx_db_id_id integer4 not null not default is
tx_db_id_id integer4 not null not default is
'exp.dmf.lg.lxb_db.id_id',
tx_db_name varchar(32) not null not default is
'exp.dmf.lg.lxb_db_name',
tx_db_owner varchar(32) not null not default is
'exp.dmf.lg.lxb_db_owner',
tx_pr_id_id integer4 not null not default is
'exp.dmf.lg.lxb_pr_id_id',
tx_wait_reason varchar(30) not null not default is
'exp.dmf.lg.lxb_wait_reason',
tx_first_log_address varchar (30) not null not default is
'exp.dmf.lg.lxb_first_lga',
tx_last_log_address varchar (30) not null not default is
'exp.dmf.lg.lxb_last_lga',
tx_transaction_id varchar(16) not null not default is
'exp.dmf.lg.lxb_tran_id',
'exp.dmf.lg.lxb_tran_id',
tx_transaction_high integer4 not null not default is
'exp.dmf.lg.lxb_tran_id.db_high_tran',
tx_transaction_low integer4 not null not default is
'exp.dmf.lg.lxb_tran_id.db_low_tran',
tx_server_pid integer4 not null not default is
'exp.dmf.lg.lxb_pid',
tx_session_id varchar(32) not null not default is
'exp.dmf.lg.lxb_sid',
tx_user_name varchar(32) not null not default is
'exp.dmf.lg.lxb_user_name'
)
as import from 'tables'
with dbms = IMA,
structure = unique sortkeyed,
key = (vnode, tx_id_id);
```

## Removing Table Registrations

Use the SQL remove table statement to remove the definition of an IMA table. This statement removes all catalog entries for the registered table, including any related views, integrities, and permits.

This statement has the following syntax:

```
remove table tablename;
```

where:

***tablename***

Is the table name for a registered IMA file.

## SQL Management Domains

A *management domain* consists of the objects in the servers and installations visible to an SQL application through the DBMS Enterprise Access. This domain is session-specific.

The domain is initially defined through a table registration, as described in Examples: Domain Tables (see page 266). Update queries can add or delete vnodes or servers to or from the domain.

Each vnode or server in the domain must be directly reachable through Ingres Net by the \$ingres user in the vnode processing the query.

## Control Objects

The MIB defines *control objects* that can be updated, using the SQL update statement, to produce known effects. The only rows that exist are those objects defined in the MIB, in a flat object table (see page 265). They cannot be created or deleted through direct SQL manipulation. Because SQL is providing a *view* of an abstract MIB, the tables are not granted permissions for inserts and deletes.

For example, to add a new write-behind thread, update the "number of write-behind threads" object, rather than inserting a new row in a table of write-behind threads. Updating the "number of write-behind threads" object adds new objects that are related to the new thread.



## IMA Dbmsinfo Constants

There are three new dbmsinfo constants available for use in queries involving IMA tables. Dbmsinfo is a function that returns a string containing information about the current session. This statement has the following syntax:

```
dbmsinfo('request_name')
```

where:

*request\_name* refers to one of those listed in the following table:

<b>Request Name</b>	<b>Response Description</b>
ima_server	Returns the IMA current server
ima_session	Returns the IMA current session
ima_vnode	Returns the IMA current vnode

For further information about using dbmsinfo, see the *SQL Reference Guide*.

## Query Examples to IMA Tables

The following sections provide examples of possible queries to IMA tables, based on the table registrations in the previous sections.

**Note:** Your actual queries must be based on the IMA models supplied on line.

## Example: Controlling Management Domain

Seeing the current domain is the following simple query:

```
select * from ima_session_domain;
```

The following query produces the ima\_objects table that contains all objects that can be updated.

```
select * from ima_objects
  where classid like 'exp.gcf.gca.%';
```

Rows are added by updating (using the update statement) the 'exp.gwf.gwm.session.control.add\_vnode' object. The following query adds all the servers in this vnode, expanding from the default of only this DBMS Server:

```
update ima_objects
  set value = dbmsinfo('ima_vnode')
  where
    classid = 'exp.gwf.gwm.session.control.add_vnode'
    and instance = '0'
    and server = dbmsinfo('ima_server');
```

The 'exp.gwf.gwm.session.control.del\_vnode' object can be used to delete rows, and an update with any value to the 'exp.gwf.gwm.session.control.reset\_domain' object removes all the rows in the domain.

## Example: What Are User Sessions Doing?

The following query shows what user sessions are currently doing in all servers in the domain:

```
select * from ima_server_sessions
  where db_owner != '';
```

This query includes sessions on other nodes of a cluster in a cluster installation.

## Example: Who Is Waiting for a Lock?

The following query shows which sessions are waiting for locks:

```
select distinct
    resource_id_id,
    lock_id_id,
    lock_state,
    ima_locklists.locklist_id_id,
    locklist_server_pid,
    locklist_session_id,
    effective_user,
    db_name,
    session_terminal,
    session_query
from
    ima_locks,
    ima_locklists,
    ima_server_sessions
where
    lock_state != 'GRANTED'
and
    ima_locks.locklist_id_id =
        ima_locklists.locklist_id_id
and
    ima_locklists.locklist_server_pid =
        ima_server_sessions.server_pid
and
    ima_locklists.locklist_session_id =
        ima_server_sessions.session_id;
```

## Example: Who Holds Locks?

The following query shows which sessions are holding locks:

```
select distinct
    lock_id_id,
    lock_state,
    ima_locklists.locklist_id_id,
    locklist_server_pid,
    locklist_session_id,
    effective_user,
    db_name,
    session_terminal,
    session_query
from
    ima_locks,
    ima_locklists,
    ima_resources,
    ima_server_sessions
where
    lock_state != 'WAITING'
and
    ima_locks.locklist_id_id =
        ima_locklists.locklist_id_id
and
    ima_locklists.locklist_server_pid =
        ima_server_sessions.server_pid
and
    ima_locklists.locklist_session_id =
        ima_server_sessions.session_id;
```

## Example: Who Is Holding a Lock Others Need?

The following query shows which sessions are holding locks for which other sessions are waiting:

```
select distinct
    resource_id_id,
    lock_id_id,
    lock_state,
    ima_locklists.locklist_id_id,
    locklist_server_pid,
    locklist_session_id,
    effective_user,
    db_name,
    session_terminal,
    session_query
from
    ima_locks,
    ima_locklists,
    ima_server_sessions
where
    resource_id_id in (
        select distinct
            ima_resources.resource_id_id
        from
            ima_resources
        where
            lock_state != 'GRANTED'
        and
            ima_locks.locks_id_id
                = ima_resources.resource_id_id
    )

and
    lock_state != 'WAITING'

and
    ima_locks.locklist_id_id =
        ima_locklists.locklist_id_id

and
    ima_locklists.locklist_server_pid =
        ima_server_sessions.server_pid

and
    ima_locklists.locklist_session_id =
        ima_server_sessions.session_id;
```

## Example: Shutting Down a Server

This query marks a server for shutdown when all sessions have exited, including this session (a Windows example is shown):

```
update ima_mib_objects
  set value = '1'
where classid = 'exp.clf.nt.cs.mon.shutserver'
  and instance = '0'
  and server = dbmsinfo('ima_server');
```

## Example: Monitoring Performance

First the following query is issued for an HP 9000/730 (a quick workstation) in various domains:

```
select count(ima_meta.classid)
  from ima_meta;
```

Assume that the following rates are observed for this query:

Domain	Rows	Seconds	Rows/sec
This DBMS	563	2	284
This instance, 1 DBMS	1257	5	251
Other instance, 1 DBMS	1193	28	42

Also, the following query is issued for the ima\_objects table:

```
select count (ima_objects.classid)
  from ima_objects;
```

Assume that the following results are observed for this second query:

Domain	Rows	Seconds	Rows/sec
This DBMS	8638	3	1266
This instance, 1 DBMS	17708	14	1265
Other instance, 1 DBMS	17410	54	322

Observe that operations in the local server are reasonably fast, as are operations to other servers in the local vnode. However, operations to other vnodes are notably slower.

This monitoring leads to the conclusion that queries involving scans, especially remote scans, must be avoided whenever possible. In particular, an update operation must always be done with a fully qualified server key or it scans the object space in the domain, which can be noticeably slow.

## Restrictions IMA Use

Following is a summary of restrictions on IMA use for the current release:

- Support for data types is limited in table registrations.
  - Only integer types, varchar, and char can be specified as data types in the register table statement.
  - All columns must be not null not default.
- The IMA supports full SQL but with the following current restrictions:
  - There are no transactions (as there is no locking or lock ownership of the underlying objects). Updates take immediate action and cannot be rolled back.
  - All reads are “dirty.” This means that querying the same record can return different (inconsistent) values as the data underlying the record are changed. You are not guaranteed to get consistent results if you select the same object repeatedly.
  - Rule triggering is unreliable. Rules are not triggered when the value of an object changes in a way unknown through SQL. This occurs to IMA-related data as the installation operates. Thus, while updates that are performed through SQL can trigger rules, rule integrity cannot be guaranteed for IMA data.
- The group of defined objects supplied with the IMA initial release is subject to change in subsequent releases.
- A cluster installation is visible to IMA as the installations on each node within a domain. No other special treatment is provided in the current IMA release.

# Appendix A: Configuring the High Availability Option for Sun

---

This appendix describes how to install and configure the Ingres High Availability Option in Sun cluster environments.

The information in this appendix is intended for engineers with extensive knowledge of software and hardware in the cluster environment, including cluster concepts such as high availability (HA), global file systems, the resource management object model, resource properties, logical host names, shared addresses, failover, and switchover.

**Note:** This appendix refers to specific Sun Cluster administration commands. For additional information about these commands, consult your Sun Cluster documentation.

## Failover

Failover is the process by which a cluster automatically relocates a data service from a failed primary node to a designated secondary node. If a failover occurs, a brief interruption in service can occur, requiring you to reconnect after the failover has completed. Once reconnected, the physical server on which Ingres is running is transparent to the end user.

The automatic failover feature ensures that transactions are handled by an active node in the cluster if a hardware or software failure occurs.

**Note:** Any transactions that are in flight at the time of failure are lost.

## Ingres High Availability Option

The Ingres High Availability Option for clusters runs application instances only on a single node. Depending on how the data service has been configured, when the fault monitor detects an error, it either attempts to restart the instance on the same node or to start the instance on another node for failover.

**Note:** The Ingres High Availability Option is not scalable; that is, it does not provide active instances on multiple nodes.

## Configuration Guidelines for System and Data Files

Ingres system files can exist on either the local disks of each cluster node or on the cluster file system. Ingres data and related files must exist on the shared cluster file system. It is recommend that all system, data, and data related files be stored in the cluster file system.

### Cluster File Systems

Access to Ingres data does not depend on the type of cluster file system. The cluster file systems must be configured with either Solstice DiskSuite or the VERITAS Volume Manager.

For instructions on creating a global disk group, see the Sun DiskSuite documentation or the VERITAS Volume Manager documentation.

Register the disk group with the cluster using the Sun Cluster administration command `scconf`.

Configure the disk group as a `SUNW.HAStorage` Resource Type. This synchronizes actions between HA storage and Ingres High Availability Option. For more information, see the `SUNW.HAStorage(5)` man page.

## DBMS Server Requirements

When using the Ingres High Availability Option, the Ingres DBMS server is installed and configured as a typical installation, provided that the database-related files are installed on the cluster file system.

At the time of installation, the installer checks if the database, checkpoint, journal, dump, work, and transaction log files are on cluster drives. For a clustered DBMS installation, the database, checkpoint, journal, dump, work, and transaction log files must be on a cluster drive. The `II_SYSTEM` (Ingres home directory) must be installed on a non-cluster local drive. This configuration should be done while at the Ingres DBMS Configuration screen during installation.

**Note:** An error message is displayed if the installer finds that the database files are not on a cluster drive.

Using a logical hostname simplifies the installation and configuration of the Ingres DBMS Server in the cluster environment. Conversely, each node can be independently configured.



## Set Up a Logical Hostname

To set up a logical hostname, follow these steps:

1. Run the script `iisuscha`. Answer yes to “configure a logical hostname” and enter the logical hostname to be used in the HA cluster.

This changes the `config.dat` file by replacing the server name with the logical hostname and sets `II_HOSTNAME` to the logical hostname.

2. Use this same logical host name when registering the HA option with the Sun cluster.

## Configure Independent Node

If the variable `II_HOSTNAME` is not to be used for a logical hostname, each host configuration must be set up in the `config.dat` file. The first host on which Ingres is installed does not need additional setup.

To add additional hostnames that can be used in the cluster, run the script `iisuscha`. Answer no to “configure a logical hostname.” Answer yes to “add a HA host” and enter the additional host to be added. This duplicates the entries in `config.dat` to match the primary server configuration settings.

Each node then has a different DBMS server configuration, which is helpful in an unbalanced cluster. Each time a change is made to the Ingres configuration, the same change must be completed on each individual node.

## Sun Cluster Software Installation Verification

The Sun Cluster packages must be installed and configured on the cluster nodes before installing the Ingres High Availability Option. The installation must be verified by a qualified Sun Cluster Engineer using the commands shown in the following table:

Command	Description
<code>scinstall -pv</code>	Displays the Sun cluster release and version information.
<code>scrgadm -p</code>	Displays the Sun cluster configured resource types, resource groups, and resources.
<code>scstat -p</code>	Displays the status of the Sun cluster components.
<code>scconf -p</code>	Displays the cluster configuration.

## Where to Install High Availability Option

The Ingres High Availability Option must be installed on all nodes in the cluster where the Ingres DBMS Server can failover.

## Resource Registration

Once the data service is installed on all nodes in the cluster, the resource must be registered. All administration must be done as a super user, normally the "root" user.

During the installation of the Ingres High Availability Option, the resource type (Ingres.ingres\_server) for the data service is automatically registered.

## Register the Ingres High Availability Option Resource

If the resource type is not registered, use the following command to register it:

```
# scrgadm -a -t Ingres.ingres_server -f  
/opt/Ingres/IngresSCAgent/etc/Ingres.ingres_server
```

where:

**-t**

Specifies the predefined resource type name for the Ingres.ingres\_server DBMS High Availability Option

**-f**

Specifies the pathname of the resource type registration file and is required if the file is not in the well-known directory(usually /usr/cluster/lib/rgm/rtreg).

## Configure Ingres High Availability Option

After the Ingres High Availability Option package has been installed, perform these steps:

1. Become the super user on a cluster node by using the following command:

```
% su
```

2. Create a failover resource group to hold the network and application resources by using the following command:

```
# scrgadm -a -g <resource-group> [-h nodelist]
```

where:

**-a**

Adds the data service resource type.

**-g**

Specifies the name of the resource group. This name can be of your choice but must be unique for the resource groups within the cluster.

**-h**

Specifies an optional comma-separated list of physical node names or IDs that identify potential masters. The order here determines the order in which the nodes are considered as primary during failover

3. Create a Logical Hostname resource in the group using the following command:

```
# scrgadm -a -L -g <resource-group> -l <logical-hostname>
```

where:

**-a**

Adds the data service resource type.

**-L**

Adds a logical Hostname resource type.

**-g**

Specifies the name of the resource group. This name can be your choice but must be unique for the resource groups with the cluster.

**-l**

Specifies a network resource. The network resource is the logical hostname or shared IP address that clients use to access Ingres for Sun Cluster HA

4. Create an Ingres resource in the group by using the following Sun Cluster administration command:

```
# scrgadm -a -j <resource-name> -g <resource-group> \  
-t Ingres.ingres_server \  
-x II_SYSTEM=<$II_SYSTEM> \  
-x II_OWNER=<ingres> \  
-x II_HOSTNAME=<logical_host> \  
[-x User=<user> -x Pword=<password> \  
-x Database=<database>] \  
[-x EnvFile=<environmentfile>] \  
[-x II_SCHA_TRACE=<level> -x II_SCHA_LOG=<filename>]
```

where:

**-a**

Adds the data service resource type.

**-j**

Specifies the name of the resource to be added.

**-g**

Specifies the name of the resource group. This name can be your choice but must be unique for resource groups with the cluster.

**-t**

Specifies the predefined resource type name for the Ingres.ingres\_server DBMS High Availability Option.

**-x II\_SYSTEM**

Sets the path for the Ingres system directory.

**-x II\_OWNER**

Sets the Ingres owner.

**-x II\_HOSTNAME**

Sets the logical host name.

**-x User (Optional)**

Sets the users for secondary monitoring.

**-x Pword (Optional)**

Sets the password for secondary monitoring.

**-x Database (Optional)**

Sets the database for secondary monitoring.

**-x EnvFile (Optional)**

Sets the file used to load additional user specified environment variables.

**-x II\_SCHA\_TRACE (Optional)**

Sets the trace level for debugging. Valid levels are 0 for no tracing and 1-9 for debugging information.

**-x II\_SCHA\_LOG (Optional)**

Sets the filename for the trace output.

Notes:

- When a fault occurs in the Ingres services resource that causes a restart, the whole resource group is restarted. Other resources in the resource group are restarted, even if they did not have a fault. To prevent other resources from being restarted along with the Ingres services, put them in a separate resource group.
- To override the default values, you can set additional extension properties for the Ingres High Availability Option.

5. Bring the resource group online by using the following command:

```
# scrgadm -Z -g <resource-group>
```

**-Z**

Enables all resources of the specified resource group.

**-g**

Specifies the name of the resource group. This name can be your choice but must be unique for resource groups within the cluster

6. Verify that the Ingres High Availability Option resource has come online on one node of the cluster by using the following command:

```
# scstat -g
```

## Starting and Stopping Ingres in a Cluster Environment

When using Ingres in a cluster environment, it must be started and stopped through the cluster service.

If Ingres is shut down while under the control of the High Availability Option, the cluster service sees this as a failure and either tries to restart Ingres on the same node or fail over the service to the failover node.

## Start Ingres in a Cluster Environment

To bring Ingres online in a cluster environment, use the following Sun Cluster command:

```
# scrgadm -Z -g <ingres_resource_group>
```

**Note:** Do **not** use the `ingstart` command to start Ingres in a cluster environment.

## Stop Ingres in a Cluster Environment

To bring Ingres offline, use the following Sun Cluster command:

```
# scrgadm -F -g <ingres_resource_group>
```

**Note:** Do **not** use the `ingstop` command to stop Ingres in a cluster environment.

## Clear a Fault in a Resource Group

If the resource group has a fault, Ingres cannot start. To clear the failure, use the following Sun Cluster command:

```
# scswitch -c -h <host>-j <ingres_resource_name> -f STOP_FAILED
```

To make sure Ingres is shutdown, use the command:

```
# scrgadm -F -g <ingres_resource_group>
```

## Change a Resource Property

To change a resource property, use the following command:

```
# scrgadm -c -j <ingres_resource_name> -x <extension_property>=<value>
```

## Turn on Debugging and Tracing for the High Availability Option

To turn on debugging:

1. Set the II\_SCHA\_TRACE level (0-9) and II\_SCHA\_LOG file name by following these steps:

2. Set the debug out file name:

```
“scgradm -c -j <ingres_resource_name> -x  
II_SCHA_TRACE=<path_to_debug_output>”
```

3. Set the trace level using the following Sun Cluster command:

```
# scrgadm -c -j <ingres_resource_name> -x II_SCHA_TRACE=<0-9>
```

where:

0=none, 1=minimal, and 9=maximum.

4. Set the debug out file name:

```
scgradm -c -j <ingres_resource_name>-x II_SCHA_TRACE=<path_to_debug_output>
```

To turn off debugging, set the level to 0 and the path to "".

## Remove the High Availability Option

Follow these steps to remove the Ingres High Availability Option:

1. Bring Ingres offline using the following Sun Cluster command:  

```
# scrgadm -F -g <ingres_resource_group>
```
2. Disable the resource using the following Sun Cluster command:  

```
# scswitch -n -j <ingres_resource_name>
```
3. Remove the resource using the following Sun Cluster command:  

```
# scrgadm -r -j <ingres_resource_name>
```
4. Disable the logical host using the following Sun Cluster command:  

```
# scswitch -n -j <logical_hostname>
```
5. Remove the logical host using the following Sun Cluster command:  

```
# scrgadm -r -j <logical_hostname>
```
6. Remove the resource group using the following Sun Cluster command:  

```
# scrgadm -r -g <ingres_resource_group>
```
7. Remove the resource type using the following Sun Cluster command:  

```
# scrgadm -r -t Ingres.ingres_server
```

**Note:** Do not remove this file if a new Ingres High Availability Option resource group is to be created.



# Appendix B: Configuring the High Availability Option for Windows

---

This appendix describes how to install and configure the Ingres High Availability Option in Windows cluster environments.

The information in this appendix is intended for engineers with extensive knowledge of software and hardware in the cluster environment, including cluster concepts such as high availability (HA), cluster disk, resource types, virtual servers, cluster administrator, and failover.

**Note:** This appendix assumes knowledge of Microsoft Windows 2000 Advanced Server or Datacenter Server, Windows Clustering, and Windows 2000 Cluster Service. It refers to specific Windows Cluster Service commands and procedures. For additional information about these commands, see your Windows Cluster Service documentation.

## Failover

Failover is the process by which a cluster automatically relocates a data service from a failed primary node to a designated secondary node. If a failover occurs, a brief interruption in service can occur, requiring you to reconnect after the failover has completed. Once reconnected, the physical server on which Ingres is running is transparent to the end user.

The automatic failover feature ensures that transactions are handled by an active node in the cluster if a hardware or software failure occurs.

**Note:** Any transactions that are in flight at the time of failure are lost.

## Ingres High Availability Option

The Ingres High Availability Option for clusters runs application instances only on a single node. Depending on how the data service has been configured, when the fault monitor detects an error, it either attempts to restart the instance on the same node or to start the instance on another node for failover.

**Note:** The Ingres High Availability Option is not scalable; that is, it does not provide active instances on multiple nodes.

## Configuration Guidelines for System and Data Files

Ingres system files, data, and related files must exist on a cluster disk.

### Cluster Disk

A cluster disk is a disk on a shared bus connected to the cluster nodes, which all the cluster nodes can access (though not at the same time).

Access to Ingres data does not depend on the type of cluster disk system (shared external SCSI, SCSI RAID, or SCSI over Fiber Channel).

### Physical Disk Resource Type

The cluster service uses the Physical Disk resource type to manage disks that are on a cluster storage device. Disks must have the same drive letters on all cluster nodes.

To prevent data corruption, only one node at a time can use a cluster disk. Cluster Service normally prevents more than one node from using a disk at the same time. However, there are some circumstances not protected by ordinary cluster safeguards: when you install a new disk, and when you remove a disk from a cluster. For more information, see the Windows 2000 Server documentation.

The Ingres system, data, and related files must exist on a Physical Disk resource type.

### Generic Service Resource Type

You use the Generic Service resource to manage Windows 2000 services as cluster resources.

The Ingres High Availability Option for Windows installs the Ingres Service as a Generic Service resource type. You must specify, during setup, the Physical Disk resources that contain the Ingres data and related files as dependencies of the Ingres Service cluster resource.

## DBMS Server Requirements

When using the Ingres High Availability Option, the Ingres DBMS server is installed and configured as a typical installation, provided that the database-related files are installed on the cluster file system.

At the time of installation, the installer checks if the database, checkpoint, journal, dump, work, and transaction log files are on cluster drives. For a clustered DBMS installation, the database, checkpoint, journal, dump, work, and transaction log files must be on a cluster drive. The II\_SYSTEM (Ingres home directory) must be installed on a non-cluster local drive. This configuration should be done while at the Ingres DBMS Configuration screen during installation.

**Note:** An error message is displayed if the installer finds that the database files are not on a cluster drive.

Using a logical hostname simplifies the installation and configuration of the Ingres DBMS Server in the cluster environment. Conversely, each node can be independently configured.

## Virtual Servers

To access a DBMS server in a non-cluster environment, network clients must connect to a physical server, identified by a network name and IP address. If that server fails, access to the DBMS is impossible.

Windows 2000 Clustering enables the creation of virtual servers, identified by a unique network name and IP address (different from the underlying physical server's network name and IP address). Unlike a physical server, a virtual server is not associated with a specific computer. If the node hosting the virtual server fails, clients can still access the DBMS using the same server name.

## Virtual Server Name

Installing the Ingres High Availability Option modifies the Ingres installation to use the virtual server name. This changes the config.dat file by replacing the server name with the virtual server name and sets II\_HOSTNAME to the virtual server name. To automatically reconnect to the server after a failover, network clients must use this virtual server name when connecting to the DBMS Server.

## Windows Cluster Service Installation

The Windows Cluster Service must be installed and configured on the cluster nodes before installing the Ingres High Availability Option. The Cluster Service is a component of Windows Clustering that creates a server cluster and controls all aspects of its operation.

### Cluster Administrator

Cluster Administrator is an application used to configure a cluster. It can run on any member of the trusted domain, whether the computer is a node of the cluster or not.

### Verify Windows Cluster Service Installation

To verify the Windows Cluster Service installation, a system administrator can start Cluster Administrator from the Windows Start menu. Choose Start, Programs, Administrative Tools, Cluster Administrator.

### Cluster Command

As an alternative to using Cluster Administrator to verify the Windows Cluster Service installation, you can use the cluster command from the Windows command prompt. Also, the command (cluster.exe) can be called from command scripts.

To verify the installation, a system administrator can use the following commands:

<b>Command</b>	<b>Description</b>
cluster /version	Displays the cluster version and service pack information.
cluster node	Displays the status of all nodes in the cluster.
cluster resource	Displays the status of all cluster resources.

## Prerequisites for High Availability Option

The following are the prerequisites for installing Ingres High Availability Option on Windows:

- Ingres High Availability option supports Windows 2000 Advanced Server or Windows 2003 Enterprise Edition.
- The Ingres directory (II\_SYSTEM) must be installed on a local non-cluster drive.
- For a clustered DBMS, the following DBMS related files must be available on a cluster file system:
  - Database
  - Checkpoint
  - Journal
  - Dump
  - Work
- Ingres must be installed as a service.

## Install High Availability Option

Installation of Ingres High Availability Option is an intermediary step during the Ingres installation. In the installation wizard, the Ingres Configuration dialog contains a Cluster Configuration section. The Cluster Configuration section consists of a checkbox for Install Ingres High Availability option and a text box for Ingres Cluster Resource Name. These options are only available on machines that are part of the cluster. They are disabled on machines that are not part of the cluster.

## Functions Performed by the Ingres High Availability Option Setup Utility

The Ingres High Availability Option setup utility performs the following functions:

- Ensures that Ingres is installed as a service on the cluster node
- Registers the Ingres Service as a Generic Service cluster resource
- Creates or updates the cluster group for the DBMS Server or Communications Server (client-only)
- Establishes dependencies of the Ingres Service on the required cluster resources (network name and disk)
- Modifies the existing Ingres installation for Windows Clustering

All administration must be done as a system administrator.

## Starting and Stopping Ingres in a Cluster Environment

When using Ingres in a cluster environment, it must be started and stopped through the cluster service.

If Ingres is shut down while under the control of the High Availability Option, the cluster service sees this as a failure and either tries to restart Ingres on the same node or fail over the service to the failover node.

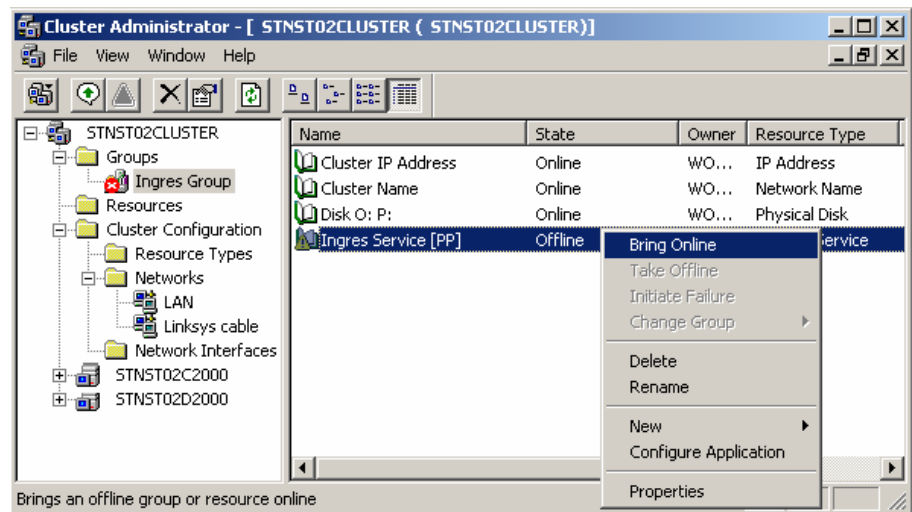
## Start Ingres in a Cluster Environment

To bring Ingres online in a cluster environment, use the Windows Cluster Administrator, and follow these steps:

1. Click Start, Settings, Control Panel, Administrative Tools, Cluster Administrator.

The Cluster Administrator window is displayed.

2. Expand the Groups folder and select the Ingres Group.
3. Right-click the Ingres Service in the right pane, and choose Bring Online.



Alternatively, you can use the following Cluster command:

```
C:> cluster resource <ingres_resource_name> /online /wait:300
```

To bring the entire Ingres group online, use the following command:

```
C:> cluster group <ingres_resource_group> /online /wait:300
```

If you omit `/wait:300,` you may see the following message, which indicates that the Ingres service did not reply within the default time period. However, the cluster command executed properly.

```
System error 997 has occurred.
Overlapped I/O operation is in progress.
```

**Note:** Do **not** use the `ingstart` command to start Ingres in a cluster environment.

## Stop Ingres in a Cluster Environment

To take Ingres offline in a cluster environment, use the Windows Cluster Administrator, as follows:

1. Click Start, Settings, Control Panel, Administrative Tools, Cluster Administrator.

The Cluster Administrator window is displayed.

2. Expand the Groups folder and select the Ingres Group.
3. In the right pane, right-click the Ingres Service, and choose Take Offline.

Alternatively, you can use the following Cluster command:

```
C:> cluster resource <ingres_resource_name> /offline /wait:300
```

To take the entire Ingres group offline, use the following command:

```
C:> cluster group <ingres_resource_group> /offline /wait:300
```

If you omit `/wait:300,` you may see the following message, which indicates that the Ingres service did not reply within the default time period. However, the cluster command executed properly.

```
System error 997 has occurred.  
Overlapped I/O operation is in progress.
```

**Note:** Do **not** use the `ingstop` command to stop Ingres in a cluster environment.



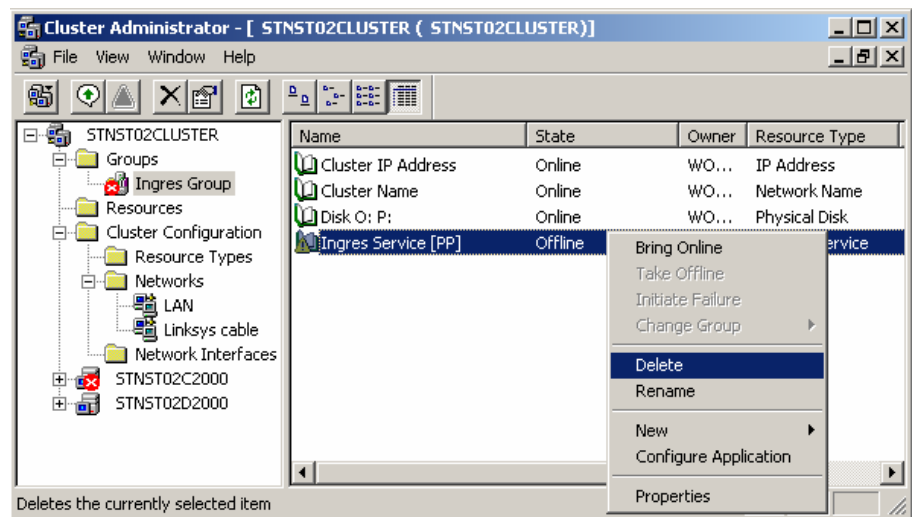
## Remove the High Availability Option

Follow these steps to remove the Ingres High Availability Option using the Windows Cluster Administrator to take the Ingres resource offline.

In the Cluster Administrator window, expand the Groups folder and select the Ingres Group.

1. Right-click the Ingres Service in the resources window, and choose Take Offline.
2. Move the resource group containing the Ingres service to the node where you originally installed Ingres. In the Cluster Administrator window, expand the Groups folder. Right-click the Ingres Group, and choose Move Group.
3. Uninstall Ingres on the node where you originally installed it. Delete the Ingres resource.

In the Cluster Administrator window, expand the Groups folder and select the Ingres Group. Right-click the Ingres Service in the resources window, and choose Delete.



A message box appears, prompting you to delete the Ingres service from the cluster. Click Yes to continue or No to cancel.

Alternatively, follow these steps to remove the Ingres High Availability Option using the Cluster command:

1. Using the following Cluster command, to take Ingres offline:

```
C:> cluster resource <ingres_resource_name> /offline /wait:300
```

If you omit “/wait: 300,” you may see the following message, which indicates that the Ingres service did not reply within the default time period. However, the cluster command executed properly.

```
System error 997 has occurred.  
Overlapped I/O operation is in progress.
```

2. Move the resource group containing the Ingres service to the node where you originally installed Ingres.

```
C:> cluster group <ingres_group_name> /move:<node_name>
```

3. Uninstall Ingres on the node where you originally installed it.
4. Delete the resource using the following Cluster command:

```
C:> cluster resource <ingres_resource_name> /delete
```

# Appendix C: Environment Variables and Logicals

---

This appendix alphabetically lists the basic Ingres environment variables and logicals and a description of each. Use this appendix as a reference in conjunction with the chapter “Setting Environment Variables and Logicals,” which describes the Ingres environment variables and logicals most commonly defined and the system level at which they are usually set.

Additional special purpose Ingres environment variables and logicals can be used in your installation, depending on the equipment and configuration.

## DBNAME\_ING

DBNAME\_ING specifies a startup file of set commands to be executed whenever a user opens the specified database (*DBNAME*). DBNAME\_ING can be defined installation-wide or locally.

**Windows:** DBNAME\_ING affects any user who connects to the database specified by DBNAME using the QUEL Terminal Monitor.

**UNIX:** DBNAME\_ING is effective in all user interfaces, including a terminal monitor.

**VMS:** The syntax for defining this logical is:

```
define DBNAME_ING path_to_file
```

## DBNAME\_SQL\_INIT

DBNAME\_SQL\_INIT specifies a startup file of set commands to be executed whenever a user connects to the database specified by DBNAME using the SQL single-line Terminal Monitor. It can be set installation-wide or locally. For details on using DBNAME\_SQL\_INIT, see the chapter “Setting Environment Variables and Logicals.”

## DD\_RServers

DD\_RServers points to the directory hierarchy that is used by the Replicator servers.

**Note:** DD\_RServers is *not* a required variable. If it is not set, it defaults to the "rep" subdirectory under the Ingres installation directory.

**Windows:** By default, DD\_RServers is defined to be %II\_SYSTEM%\ingres\rep. To define a different location, use, for example:

```
set DD_RServers c:\rep2
```

**UNIX:** By default, DD\_RServers is defined to be the \$II\_SYSTEM/ingres/rep. To define a different location, use, for example:

C Shell:

```
setenv DD_SERVERS /usr/rep2
```

Bourne Shell:

```
DD_RServers=/usr/rep2; export DD_RServers
```

**VMS:** By default, DD\_RServers is defined to be the II\_SYSTEM:[ingres.rep]. To define a different location, use, for example:

```
define DD_RServers user_disk:[rep2]
```

The \$DD\_RServers/servers (DD\_SERVERS:[servers] for VMS) directory contains server1 through server10. These directories are used by the Replicator servers during runtime.

If you need your installation to support multiple Replicator configurations and/or multiple source databases, use DD\_RServers to access each Replicator environment. For each environment, create a DD\_RServers directory with the subdirectory servers. In the servers directory, create subdirectories server1 through server*N* where *N* is the number assigned to the Replicator Server.

## II\_4GL\_DECIMAL

II\_4GL\_DECIMAL determines the decimal point character for 4GL source input. For example, in France, the decimal point character is a comma. The default value is ".". Valid values are "." and ",".

## II\_AFD\_TIMEOUT

II\_AFD\_TIMEOUT is used to set the number of seconds before timeout in Vision's Application Flow Diagram.

When you are using Vision's Application Flow Diagram, Vision periodically checks for changes made by concurrent developers that affect the diagram. If such a change has been made, Vision notifies you and gives you a chance to update the display.

You can set II\_AFD\_TIMEOUT to any value greater than 20, or 0 to disable the timeouts. Setting the value to much less than the default of 300 seconds causes Vision to time out and query the database often enough to visibly affect performance.

## II\_APPLICATION\_LANGUAGE

II\_APPLICATION\_LANGUAGE is set to the language Application-By-Forms uses to translate string constants.

ABF allows you to internationalize your applications by translating the string constants you define for them. When you run the application, it normally uses the constants defined for your default language. You can override this by defining II\_APPLICATION\_LANGUAGE to the name of another language. For more information, see *Forms-based Application Development Tools User Guide*.

## II\_BIND\_SVC\_xx

II\_BIND\_SVC\_xx is set to the bind address for Ingres tools connections to a named Ingres or Enterprise Access installation.

## II\_C\_COMPILER (VMS)

II\_C\_COMPILER is set to the value of the C compiler in use for ABF. Valid entries are VAX11(VAX/VMS only), DECC, or NONE. II\_C\_COMPILER is defined at the system or group level. (Call technical support for information about how to change the C compiler used by ABF).

## II\_CHARSETxx

II\_CHARSETxx is used to set the type of character set used for string operations. This is set during ingbuild and cannot be changed without corrupting data.

## II\_CHECKPOINT

II\_CHECKPOINT is an area set to the full file specification for the default checkpoint location. It is set by the system administrator during the ingbuild procedure and cannot be changed, even during installation updates. Specific databases can designate alternate locations for checkpoints as a parameter to the createdb command.

## II\_CLIENT

II\_CLIENT is set to "true" for client installations. Do not change II\_CLIENT if it is set.

## II\_COLLATION

II\_COLLATION determines the collating sequence for a database when the database is created. Valid entries are "multi" to specify the DEC Multinational Character Sequence or "spanish" to specify the Spanish alphabet's character sequence. More on collation sequences is described in Local Collation Sequences (see page 94). II\_COLLATION can be defined at the system level or redefined by the individual user at the process or job level.

## II\_CONFIG

II\_CONFIG sets the full file specification of the Ingres files directory during the installation procedure. It cannot be changed.

## II\_CONNECT\_RETRIES

II\_CONNECT\_RETRIES sets the number of times an application retries a connection before returning an error. The default is 1. An optional flag (d or D) causes a delay before the retry. For example: II\_CONNECT\_RETRIES=1d.

## II\_DATABASE

II\_DATABASE is an area set to the full file specification for the default database location. It is set by the system administrator during the ingbuild procedure and cannot be changed, even during installation updates. Specific databases can designate alternate locations as a parameter to the createdb command.

## II\_DATE\_FORMAT

II\_DATE\_FORMAT defines the format style for date values. It is defined installation-wide. Default is the US value with an output format of *dd-mmm-yyyy*. Default legal input formats for US dates include:

```
mm/dd/yy
mm/dd/yyyy
dd-mmm-yyyy
mm-dd-yyyy
yyyy.mm.dd
yyyy_mm_dd
mmdyy
mm-dd
mm/dd
```

If set, it replaces the default format with an alternative format. The following are valid II\_DATE\_FORMAT settings, showing the alternative output format:

Setting	Output Format
US (default)	dd-mm-yy
ISO	yymmdd
MULTINATIONAL	dd/mm/yy
SWEDEN or FINLAND	yyyy-mm-dd
GERMAN	dd.mm.yy
DMY	dd-mmm-yyyy
MDY	mmmm-dd-yyyy
YMD	yyyy-mmm-dd

For example, if the logical's setting is ISO, the dates are output as *yymmdd*. If you enter a date using a similar, six-character format, such as *mmddy*, the date is interpreted on output as *yymmdd*. For example, if you enter the date March 9, 1912, as 030912 (*mmddy*) and II\_DATE\_FORMAT is set to ISO, this date is interpreted as Sept. 12, 1903 (030912 as *yymmdd*). You receive an error message if the date you entered cannot be interpreted as a valid date in the output format.

**Note:** When using the `_date4()` function (MULTINATIONAL4), the output format for the year value always returns 'yyyy'.

For more information about valid date input and output formats, see the *SQL Reference Guide*.

## II\_DATE\_CENTURY\_BOUNDARY

II\_DATE\_CENTURY\_BOUNDARY determines the inferred century for a date that is missing the century on input. It can be set to an integer value in the range of  $0 < n < = 100$ . If this environment variable is not set in the user's environment or if it is set to 0 or 100, the current century is used.

The century is determined by the following calculation:

```
if ( date_year < II_DATE_CENTURY_BOUNDARY )
  then
    (date_year += 2000)
  else
    (date_year += 1900)
endif
```

For example, if II\_DATE\_CENTURY\_BOUNDARY is 50 and the current year is 1999, an input date of 3/17/51 is treated as March 17, 1951, but a date of 03/17/49 is treated as March 17, 2049.

If the user enters the full four digits for the year into a four-digit year field in the application, the year is accepted as entered, regardless of the II\_DATE\_CENTURY\_BOUNDARY environment variable setting.



## II\_DBMS\_LOG

II\_DBMS\_LOG points to the location of the log file containing DBMS server-specific error messages. It is set installation-wide. The server reads it at startup time, so changing it after startup has no effect on an existing server. Most error messages written to this file are also logged to the installation-wide error log (errlog.log). Some server trace messages and statistics information are written to the II\_DBMS\_LOG file that are not written to errlog.log. The default, if not set, is \$II\_SYSTEM/ingres/files/errlog.log.

## II\_DBMS\_SERVER

II\_DBMS\_SERVER points to the DBMS Server to which the client process automatically connects. It can be set installation-wide or locally. When II\_DBMS\_SERVER is set, application processes bypass the Name Server. Using II\_DBMS\_SERVER requires setting II\_GC\_REMOTE (see page 308) in the DBMS Server installation.

## II\_DECIMAL

II\_DECIMAL specifies the one character used to separate fractional and non-fractional parts of a number. Default value is the period (.), as in 12.34. Alternatively, the comma (,) can be used, as in 12,34. Only (.) and (,) are allowed. It is set installation-wide.

## II\_DIRECT\_IO (UNIX)

II\_DIRECT\_IO is specific to UNIX environment. If set, it causes disk reads and writes to be done directly to disk, circumventing the UNIX file buffer system. It is only available on some platforms.

## II\_DISABLE\_SYSCHECK

II\_DISABLE\_SYSCHECK controls the execution of the **syscheck** command. The syscheck command checks the system to confirm the availability of the required resources for running Ingres. If II\_DISABLE\_SYSCHECK is set, syscheck is still called but will immediately exit without performing any checks and with a success return status. For more information on the syscheck command, see *Command Reference Guide*.

## II\_DML\_DEF

II\_DML\_DEF defines the default query language for an installation as SQL or QUEL. It is used in installations that support both the SQL and QUEL database languages, in which case one such language can be designated as the default. Applications-By-Forms, Ingres 4GL, and Report-By-Forms are parts of the Ingres tools that use II\_DML\_DEF. SQL is the default if this is not set. It can be set installation-wide or locally.

## II\_DUMP

II\_DUMP defines an “area” that specifies the default location of the dump log files used in online backup. (For more information on online backup, see the ckpdb description in the *Command Reference Guide*.) It is defined installation-wide in the ingbuild procedure and cannot be changed.

## II\_EMBED\_SET

II\_EMBED\_SET allows application programs to set a variety of features. It can be set installation-wide or locally. The acceptable values and the actions they specify are:

**errorqry**

Prints query that caused error.

**savequery**

Enables/disables saving of the text of the last query issued. Specify 1 to enable or 0 to disable. To obtain the text of the last query, you must issue the `inquire_ingres(:query= querytext)` statement. To determine whether saving is enabled, use the `inquire_ingres(:status=savequery)` statement.

**printqry**

Prints all query text and timing information to `iiprtqry.log` file in the current directory.

**qryfile**

Specifies an alternate text file to which query information is written. The default file name is `iiprtqry.log`.

**printgca**

Traces GCA messages to the `iiprtgca.log` file in the current directory.

**gcafile**

Specifies an alternate text file to which GCA information is written. The default file name is `iiprtgca.log`.

**printtrace**

Enables or disables saving of DBMS trace messages to a text file (specified by `tracefile`). Specify 1 to enable saving of trace output, or 0 to disable saving.

**tracefile**

Specifies an alternate text file to which tracepoint information is written. The default file name is `iiprttrc.log`.

**sqlprint**

Prints all errors and database procedure messages arising from embedded SQL statement.

**eventdisplay**

Print events as they return from DBMS.

**dbmserror**

Makes local (DBMS) error numbers the default error numbers returned to the application in `errno` and the SQLCA's `sqlcode`.

**genericerror**

Makes generic error numbers the default error numbers returned to the application in `errno` and the SQLCA's `sqlcode`.

**prefetchrows**

Specifies the number of rows to be "prefetched" when retrieving data using cursors. Valid arguments are:

**0 (default)**

Ingres determines the number of rows to prefetch.

**1**

Disables prefetching. Each row is fetched individually.

**N (positive integer)**

Specifies the number of rows to be prefetched.

## II\_ERSEND (UNIX)

II\_ERSEND is set during the installation procedure to the full path name of the `errlog.log` file.

## II\_FORCE\_HET

II\_FORCE\_HET must be set to TRUE to allow seamless communication between the local client and remote databases on 64-bit HP Tru64 Alpha (formerly DEC Alpha) systems. II\_FORCE\_HET must be set before the client program attempts communication. For the change to take effect, the Communication Server and Name Server must be stopped and restarted.

## II\_FRS\_KEYFIND

II\_FRS\_KEYFIND allows users to turn off the keyfind frs function on table fields in read mode.

## II\_GC\_REMOTE

II\_GC\_REMOTE allows or disallows a direct connection between the application process on the client machine and the Ingres DBMS process on the server machine, thus bypassing Name Server processing. The vnode attribute, `connection_type`, must be set to "direct" to enable direct connections. For additional information on this vnode attribute, see the *Connectivity Guide*.

II\_GC\_REMOTE only needs to be set in the DBMS Server installation.

Valid II\_GC\_REMOTE values are:

### **ON or ENABLE**

On or Enable allows direct connections that use the vnode attribute as described above, and rejects direct connections that are attempted using the II\_DBMS\_SERVER environment variable.

### **FULL or ALL**

Full or All allows both vnode controlled and II\_DBMS\_SERVER controlled direct connections. You must avoid using this setting because it compromises network security. (No validation occurs to determine whether the client installation has permission to access the server installation.)

If II\_GC\_REMOTE is not set, or is not set to one of the above values, direct connections are disallowed.

## II\_GCA\_LOG

II\_GCA\_LOG denotes the output destination for II\_GCx\_TRACE trace information for all Ingres components. Can be set to `stdio` or to the full path to a file.

## II\_GCx\_TRACE

II\_GCx\_TRACE specifies environment variables and logicals for tracing:

### **II\_GCA\_TRACE**

Sets the trace level for General Communications Architecture (GCA)

### **II\_GCN\_TRACE**

Sets the trace level for Name Server (GCN)

### **II\_GCC\_TRACE**

Sets the trace level for Communication Server (GCC)

### **II\_GCD\_TRACE**

Sets the trace level for Data Access Server (GCD)

### **II\_GC\_TRACE**

Sets the trace level for GCA Compatibility Library

### **II\_GCCCL\_TRACE**

Sets the trace level for GCC Compatibility Library

Each variable can be set separately to a number from 0 to 4 specifying trace level, with levels 1 through 4 increasing in amount of information traced.

If any of these values are set to non-zero, trace statements for the appropriate items are logged in the II\_GCA\_LOG file.

## II\_GCC\_LOG

II\_GCC\_LOG denotes the output destination for the Communications Server (GCC) trace information. Setting this environment variable is preferable to using II\_GCA\_LOG because it is dedicated to the Communications Server process.

## II\_GCD\_LOG

II\_GCD\_LOG denotes the output destination for the Data Access Server (GCD) trace information. Setting this environment variable is preferable to using II\_GCA\_LOG because it is dedicated to the Data Access Server process.

## II\_GCN\_LOG

II\_GCN\_LOG denotes the output destination for the Name Server (GCN) trace information. Setting this environment variable is preferable to using II\_GCA\_LOG because it is dedicated to the Name Server process.

## II\_GCNxx\_PORT

II\_GCNxx\_PORT contains the connect address of a local installation's Name Server. This variable must not be reset.

xx = 2-character code defined by II\_INSTALLATION

## II\_HALF\_DUPLEX

II\_HALF\_DUPLEX, if set to 0 or undefined, enables Net duplex operation for SNA LU6.2 conversations. If set to 1, the Net half-duplex operation is enabled.

## II\_HELP\_EDIT

II\_HELP\_EDIT, if set to any value, adds an extra operation, Edit, to menus encountered in help text screens in the forms systems. The Edit operation enables users to edit the help text in the screen with the text editor defined by ING\_EDIT. II\_HELP\_EDIT is typically defined locally by the individual user.

## II\_INSTALLATION

II\_INSTALLATION is a two-character code used to define a particular installation. This code is set during installation. If you want to change it, rerun ingbuild.

## II\_JOURNAL

II\_JOURNAL is set to the full file specification for the default journal location. The system administrator during the ingbuild procedure sets it and it cannot be changed, even during installation updates. Specific databases can designate alternate locations for journals as a parameter to the createdb command.

## II\_LANGUAGE

II\_LANGUAGE determines what language is used for screen messages, menu items, and prompts. It is defined installation-wide. Valid values are as follows:

ENGLISH	English (Default)
DEU	German
ESN	Spanish
FRA	French
ITA	Italian
JPN	Japanese
PTB	Brazilian Portuguese
SCH	Simplified Chinese

## II\_LOG\_DEVICE (VMS)

II\_LOG\_DEVICE specifies the name of the logging/locking pseudo device to which the logging/locking driver is connected. This value is QAA0: for production installations. For test installations, the value is QxA0:, where *x* is the first character of the group installation identifier. This logical is defined at the system or group level.

## II\_MONEY\_FORMAT

II\_MONEY\_FORMAT defines the format of money output. It is set installation-wide. You can change the format by setting II\_MONEY\_FORMAT to a string with two symbols separated by a colon (:). The symbol to the left of the colon indicates the location of the currency symbol. It must be "L" for a leading currency symbol or a "T" for a trailing currency symbol. The symbol to the right of the colon is the currency symbol you want displayed. Currency symbols can contain up to 4 physical characters. For example:

Logical Definition	Result
L:\$	\$100
T:DM	100DM
T:F	100F

## II\_MONEY\_PREC

II\_MONEY\_PREC specifies the number of decimal places to be displayed for money values. Valid values are 0, 1, and 2. The default is 2 (for decimal currency). It is set installation-wide. II\_MONEY\_PREC is applied when a money value is converted to a character string (for printing, for example). Extra decimal places are rounded. For example, if II\_MONEY\_PREC is set to 0, 9.50 is rounded to 10.

## II\_MSGDIR

II\_MSGDIR sets the directory containing error and information messages. Users must not change this setting. The default file is the II\_LANGUAGE setting, located at \$II\_SYSTEM/ingres/files/english.

## II\_NULL\_STRING

II\_NULL\_STRING specifies the string used by a terminal monitor and other Ingres tools to represent the null value. The string consists of one to three characters. The default value is three blanks.

## II\_NUM\_OF\_PROCESSORS

II\_NUM\_OF\_PROCESSORS is set to the number of processors on the machine. The variable affects processing in the Ingres DBMS, particularly when waiting for mutually exclusive resources; a value larger than 1 causes a more compute-intensive strategy that may result in getting the resource faster; however, this strategy is not generally advisable on single processor machines.



## II\_NUM\_SLAVES (UNIX)

II\_NUM\_SLAVES specifies the number of slave processes that a DBMS server creates to do disk operations. The default is two, but systems having applications with high I/O throughput requirements, many disks, or faster disk drives can increase the number of slave processes. The following example allows each new server that is started to have four slaves:

```
ingsetenv II_NUM_SLAVES 4
```

Stop and restart DBMS servers to reset the number of slaves. The maximum supported value is 30 and the minimum is 0. Setting this environment variable/logical to 0 degrades performance substantially.

**Note:** On UNIX implementations that use O/S threads, II\_NUM\_SLAVES is not set and there are no I/O slaves.

## II\_OO\_TABLE\_SIZE

II\_OO\_TABLE\_SIZE allows you to increase the size of an internal COPYAPP table to support the copying of more application objects than is allowed by the default (4096 objects).

## II\_PATTERN\_MATCH

II\_PATTERN\_MATCH determines the set of pattern matching characters for QBF and the ABF qualification function. These can be SQL-like, QUEL-like, or not set.

### **sql**

The pattern matching characters used by QBF and the ABF qualification function are % \_ [ and ]. The % is equivalent to the QUEL "\*" and \_ is equivalent to the QUEL "?".

### **quel**

The pattern matching characters used by QBF and the ABF qualification function are \* ? [ and ].

### **<not set>**

If not set, or if it is set to anything other than sql or quel, the QUEL pattern matching characters are used.

II\_PATTERN\_MATCH can be set installation-wide or locally. The complete word, either sql or quel, must be entered. The setting for II\_PATTERN\_MATCH is case-sensitive.

## II\_PF\_NODE

II\_PF\_NODE specifies the vnode name of a remote node to be managed through the Microsoft Windows Performance Monitor.

## II\_POST\_4GLGEN

II\_POST\_4GLGEN specifies a command to be translated and executed by the Vision component following successful generation of Ingres 4GL code. The command is passed the full path name of the Ingres 4GL source file just generated. This logical is used only by sites that have installed the Vision product. For more information on this logical, see the *Forms-based Application Development Tools User Guide*.

## II\_PRINTSCREEN\_FILE

II\_PRINTSCREEN\_FILE is used with the printscreen function. It specifies a default file name for the output file of the printscreen function. This environment variable/logical is usually defined locally by the individual user. If this environment variable/logical is set and no file name is sent to the printscreen function, no prompt is given and this file name is used. If not set, the user is prompted for a file name. If the file name "printer" is specified for II\_PRINTSCREEN\_FILE, the screen depiction is sent directly to the line printer. For more information, see the *Character-based Querying and Reporting Tools User Guide*.

## II\_SQL\_INIT

II\_SQL\_INIT (see page 76) can be set locally to the full path name of a file containing SQL commands. Typically, it is defined locally by the individual user. When a user with II\_SQL\_INIT set connects to a terminal monitor, the commands in the named file are processed.

## II\_STAR\_LOG

II\_STAR\_LOG points to the location of the log file containing distributed DBMS server-specific error messages. It is set installation-wide. It is read by the server at startup time, so changing it after startup has no effect on an existing server. Most error messages written to this file are also logged to the installation-wide error log (errlog.log). Some server trace messages and statistics information are written to the II\_STAR\_LOG file that are not written to errlog.log. The default, if not set, is \$II\_SYSTEM/ingres/files/errlog.log.

## II\_SYSTEM

II\_SYSTEM specifies an “area” that specifies the parent directory of the ingres directory, where many components of your installation are located. It must not be changed without reinstalling Ingres.

## II\_TEMPORARY

II\_TEMPORARY specifies the directory where temporary files used by the Ingres tools are created. By default, these files are created in the user’s current directory. II\_TEMPORARY can be redefined installation-wide or locally.

## II\_TERMCAP\_FILE

II\_TERMCAP\_FILE specifies an alternative termcap file to use. This file must be in Ingres termcap file format. It can be redefined installation-wide or locally.

## II\_TFDIR

II\_TFDIR specifies the directory in which the Vision code generator searches for Ingres 4GL template files. This logical is used only by sites that have the Vision component installed.

## II\_THREAD\_TYPE

II\_THREAD\_TYPE specifies whether Ingres uses internal threads or OS threads. Valid settings are:

**OS (default)**

Operating System threads

**INTERNAL**

Internal threads

This variable can be set during installation. For additional information on setting this variable for your platform, see the Readme file.

## II\_TIMEZONE\_NAME

II\_TIMEZONE\_NAME specifies the world time zone that determines the installation's location for timing purposes. It is defined during the ingbuild procedure.

## II\_TM\_ON\_ERROR

II\_TM\_ON\_ERROR defines the action taken when an error occurs in a terminal monitor. The valid settings are continue (on error) or terminate (on error). If this environment variable/logical is not set, the default action for the (non-forms based) terminal monitor is to continue; the default for Interactive SQL or QUEL (the forms-based terminal monitor) is to terminate. It can be defined installation-wide or locally by the individual user.

## II\_TUXEDO\_LOC

II\_TUXEDO\_LOC supports Ingres DTP for Tuxedo. Specifies the directory where the shared memory file is to be created. It must be the same for all servers in a group. If II\_TUXEDO\_LOC is not set, II\_TEMPORARY is used.

## II\_TUX\_SHARED

If II\_TUX\_SHARED is set to USER, the name of the shared memory segment used by Ingres is t<username>.tux; otherwise the name is t1.tux.

## II\_TUX\_AS\_MAX

II\_TUX\_AS\_MAX is the maximum number of application and TMS servers that can be started. The default value is 32. A maximum of II\_TUX\_AS\_MAX servers are permitted to attach to the Ingres shared memory segment.

## II\_TUX\_XN\_MAX

II\_TUX\_XN\_MAX is the total number of transaction entries allocated in the shared memory segment. Each server that attaches the shared memory segment reserves II\_XA\_SESSION\_CACHE\_LIMIT transaction entries for its own use. The default value is 1024.

---

## II\_UNICODE\_CONVERTER

II\_UNICODE\_CONVERTER sets the converter to be used for Unicode coercion. The DBMS Server carries out the coercion of Unicode data to local character set value based on the Unicode converter (or mapping table). Several converters are available in compiled and xml format, located in the directory \$II\_SYSTEM/ingres/files/ucharmaps. By default, Ingres automatically chooses the converter to use, based on the encoding used by the locale on your system.

You can manually set II\_UNICODE\_CONVERTER if Ingres is unable to set the converter correctly for your locale for coercing Unicode data to local character set value.

### To manually set II\_UNICODE\_CONVERTER on UNIX based on locale

1. Type the following at the UNIX command prompt:

```
locale
```

The locale on your system is returned.

2. Use the locale value to find the character set on your system. For example, if locale is ja\_JP.sjis, issue the following command:

```
LC_CTYPE=ja_JP.sjis locale charmap
```

The character set used by the system is returned.

3. Look up the character set in the file \$II\_SYSTEM/ingres/files/ucharmaps/charmapalias.xml. This file lists mapping IDs and their corresponding aliases. Find the mapping ID by finding the alias name value that matches the alias name. This mapping ID is the value of the converter.
4. Set the value of II\_UNICODE\_CONVERTER in the symbol table, as follows:

```
ingsetenv II_UNICODE_CONVERTER mapping_id
```

Recycle the server by issuing the `ingstop` and the `ingstart` commands.

## II\_UUID\_MAC (Windows)

II\_UUID\_MAC, when set to TRUE, uses an older algorithm for generating UUIDs on Windows.

For security reasons, the `uuid_create` function on Windows no longer uses a machine's MAC address to generate UUIDs. You can use the II\_UUID\_MAC environment variable if you do not need the level of security provided by the new `uuid_create` algorithm and for the sake of compatibility want to continue to use the MAC address of your machine when creating UUIDs.

## II\_VNODE\_PATH

II\_VNODE\_PATH specifies a series of remote virtual nodes to be tried each time a user attempts to connect to the specified *vnode*. Node names are specified as strings to be appended to *vnode*. Commas must separate the strings.

For example, the following specifies that the remote virtual nodes *mynode0*, *mynode1*, *mynode2*, and *mynode3* be tried each time a user attempts to connect to *mynode*:

```
II_MYNODE_PATH=0, 1, 2, 3
```

The names must be valid remote virtual nodes defined by the Ingres Network Utility or the Net Management Utility (*netutil*).

A single connection attempt begins with one of the nodes, chosen at random, and continues until all of the specified nodes have been tried. The number of connection attempts made before an error is returned to the user is determined by the value of `II_CONNECT_RETRIES`.

## II\_WORK

II\_WORK specifies a directory that contains temporary files used by Ingres during the execution of certain queries such as sorts, copies, and modifies.

## II\_XA\_TRACE\_FILE

II\_XA\_TRACE\_FILE specifies a file in which Ingres DTP (Distributed Transaction Processing) logs the events occurring through the TMXA interface, as well as any SQL performed against the Ingres DBMS. The user who starts the application servers must have write access to the file. For additional information about this file, see the *Distributed Transaction Processing User Guide*.

## IIDLDIR

IIDLDIR specifies a directory that contains dynamically loaded libraries.

## ING\_ABFDIR

ING\_ABFDIR defines the directory name for the top-level directory to contain Application-By-Forms object libraries and work files. Under this directory, a subdirectory is created for each database that contains ABF applications and, within that subdirectory, another subdirectory is created for each ABF application. The ING\_ABFDIR directory level must have read, write, and execute permissions for group, owner, and world. If, however, the application designer redefines ING\_ABFDIR to point to a directory of his own, he can optionally restrict permissions on the directory to read, write, execute to owner. It can be redefined installation-wide or locally.

## ING\_ABFOPT1

ING\_ABFOPT1 specifies the name of a file, used by ABF, in which additional linker options can be placed. These options are not required to run ABF applications but can be required to include libraries, compiled objects, and other capabilities not available in ABF. Users can modify this file. For more information, see the *Forms-based Application Development Tools User Guide*. The individual user usually defines it at the local level.

## ING\_EDIT

ING\_EDIT specifies the default editor invoked by various editor commands. The default for the entire installation is set during the installation procedure. Users can also set this in their local environment.

**Windows:** You must set ING\_EDIT to the full path name of the editor. For example:

```
c:\tools\vi.exe
```

**UNIX:**

```
/usr/ucb/vi
```

**VMS:** If undefined, the callable EDT editor is called. ING\_EDIT can be defined to be:

**+EDT**

Call the EDT editor

**+TPU**

Call the TPU editor

**+LSE**

Call the language sensitive editor.

The "+" symbol indicates that the callable version of the editor is to be used. This avoids creating a new subprocess in which to run the editor and results in much faster start up of the editor when called from a terminal monitor, ABF and other Ingres tools programs. Example:

```
DEFINE ING_EDIT "+TPU"
```

No qualifiers can be used. You must specify any non-default editor behavior in an editor startup file. For example, to use the Language Sensitive Editor in FORTRAN, define the following:

```
DEFINE ING_EDIT "+LSE" DEFINE LSE$COMMAND - "SYS$LOGIN:LSE_STARTUP.LSE"
```

## ING\_PRINT

ING\_PRINT specifies the default printer command issued by the Print function. The default is PRINT. It is usually defined locally by the individual user.



## ING\_SET

ING\_SET specifies a quoted string or startup file of set commands to be executed whenever a user connects to the DBMS Server through an application, an Ingres tool, or a single-line terminal monitor. It can be set installation-wide or locally. For details on using ING\_SET, see the chapter “Setting Environment Variables and Logicals.”

## ING\_SET\_DBNAME

ING\_SET\_DBNAME specifies a quoted string or startup file of set commands to be executed whenever a user connects to the database specified by *DBNAME*. Affects any user who connects to the specified database through an application, an Ingres tool, or a single-line terminal monitor. It can be set installation-wide or locally. For details on using ING\_SET\_DBNAME, see the chapter “Setting Environment Variables and Logicals.”

## ING\_SHELL (UNIX)

ING\_SHELL, If defined, contains the path name of the shell used by Ingres Menu and a terminal monitor when the shell operation is invoked.

## ING\_SYSTEM\_SET

ING\_SYSTEM\_SET specifies a quoted string or startup file of set commands to be executed whenever a user connects to the DBMS Server through an application, an Ingres tool, or a terminal monitor. It is always global. For details on using ING\_SYSTEM\_SET, see the chapter “Setting Environment Variables and Logicals.”

## INGRES\_KEYS

INGRES\_KEYS specifies a mapping file to customize a user’s keyboard layout for use with forms based programs. For a complete description of the mapping file syntax, see *Character-based Querying and Reporting Tools User Guide*. This environment variable/logical is usually set in the user’s local environment.

## INIT\_INGRES

INIT\_INGRES can be set to the full file specification of a file containing QUEL commands. It is usually defined locally by the individual user. When a user with INIT\_INGRES set connects to the QUEL Terminal Monitor, the commands in the named file are processed. For details on using INIT\_INGRES, see the chapter “Setting Environment Variables and Logicals.”

## TERM (UNIX)

TERM is the terminal description for the terminal upon which you are executing UNIX programs. The TERM environment variable is not used directly, but it becomes the default for TERM\_INGRES, if TERM\_INGRES is not explicitly set. TERM is used by UNIX programs like vi, which can be called as subsystems from Ingres programs.

## TERM\_INGRES

TERM\_INGRES contains the terminal designation for the terminal upon which you are executing one of the Ingres forms-based products, such as QBF or VIFRED. For a list of supported values, see *Character-based Querying and Reporting Tools User Guide*. It is usually defined locally by the individual user.

**Windows:** TERM\_INGRES is most conveniently set in the symbol table to IBMPC.

**UNIX:** TERM\_INGRES is most conveniently set in a .login or .profile file. TERM\_INGRES takes precedence over TERM in defining the terminal type. It allows TERM to be defined differently for use by other UNIX programs such as vi.

**VMS:** TERM\_INGRES is most conveniently set in a LOGIN.COM file.

# Index

---

## [

[ ] (brackets) • 17

## {

{ } (braces) • 17

## |

| (vertical bar) • 17

## A

access commands, adding to login.com file • 61

acpexit • 100

ACPEXIT.COM • 100

acpexit.com, arguments • 100

ADF (Abstract Data Type Facility) • 22, 199

aducompile (utility) • 97, 100

alert events, setting • 228

OS events • 228

archiver process • 100

component description • 35

described • 36

error log • 38, 203

archiving • 100

## B

bold typeface • 17

Bridge Server

configuring the • 26

described • 25

## C

classid • 256

collation sequence

aducompile (utility) • 100

creating customized • 97

described • 94

description file • 98

local • 94

multi • 95

pattern matching • 95, 96, 98

spanish • 96

Communications Server

described • 25

facility code • 199

components, installation • 20

config.dat • 61

config.log • 39

configuration

logs • 39

name, changing • 46

parameters, restoring • 234

Configuration Manager

accessing • 43

changing the configuration name • 46

configuring Bridge Server parameters • 26

configuring Data Access Server parameters • 26

configuring DBMS server parameters • 49

configuring Internet communications • 55

configuring locking system parameters • 55

configuring logging system parameters • 56

configuring multiple log partitions • 56

configuring Name Server parameters • 25

configuring Recovery Server parameters • 54

configuring Star Server parameters • 53

configuring system components • 45

configuring transaction log parameters • 57

setting parameter values • 47

Configuration-By-Forms, accessing • 43

configurations, comparing • 234

configuring

Bridge Server • 26

changing parameter values • 47

Data Access Server • 26

DBMS Server • 21

Internet communications • 55

locking system • 55

logging system • 56

multiple log partitions • 56

Name Server • 25

parameters, using IVM • 226

protected mode • 47

Recovery Server • 54

Star Server • 53

conventions for describing syntax • 17

cscleanup (utility) • 184

csinstall (utility) • 184

csreport (utility) • 184

---

## D

- Data Access Server
  - configuring the • 26
  - described • 25
- database events, tracing • 217
- databases
  - new • 87
- dayfile • 71, 72
- DBMS Server
  - configuring the • 21
  - described • 21
  - facilities • 22
- DBNAME\_ING • 72, 299
- DBNAME\_SQL\_INIT • 74, 299
- DD\_RServers • 300
- defining message categories and notification levels • 227
- diagnostics, Windows • 185
- direct connections • 308
- DMF (Data Manipulation Facility) • 22, 199
- dmf\_write\_behind option • 51
- dmfapSee archiver process • 36
- dmfrcp Seealso recovery process • 204
- dmfrcp Seerecovery process • 38
- dmfrcpSee also recovery process • 35
- dual transaction log
  - editing parameters • 57

## E

- environment variables
  - UNIX • 59
  - Windows • 60
- environment variables/logicals • 59
  - global • 71
  - installation-wide • 61
  - local • 71
  - not resettable • 69
  - resettable • 67
  - scope • 59
  - types • 59
  - user-definable • 65
  - Windows • 59
- error messages
  - facility code • 198
  - fatal • 198
  - help • 197
  - Ingres • 197
  - non-fatal • 198
- errors

- DBMS error log • 40, 205
- errlog.log • 38, 203
- GCC error log • 40, 205
- monitoring with Ingres console • 221
- Star error log • 40

- events
  - logging Operating System events • 228

## F

- facility
  - Abstract Data Type (ADF) • 22
  - Data Manipulation (DMF) • 22
  - General Communication (GCF) • 25
  - Optimizer (OPF) • 22
  - Parser (PSF) • 22
  - Query Execution (QEF) • 22
  - Query Storage (QSF) • 22
  - Relation Description (RDF) • 22
  - System Control (SCF) • 22
- facility codes • 198

## G

- GCA (General Communications Architecture)
  - described • 25
- GCA (General Communications Area)
  - facility code • 199
- GCF (General Communication Facility) • 218
  - Bridge Server • 25
  - Communications Server • 25
  - Data Access Server • 25
  - described • 25
  - facility code • 199
  - General Communications Architecture • 25
  - Name Server • 25

## H

- help (VMS utility) • 193
- High Availability Option for Sun Clusters • 279
- High Availability Option for Windows • 289
- History of Changes page • 57

## I

- II\_4GL\_DECIMAL • 300
- II\_ABF\_RUNOPT • 67
- II\_AFD\_TIMEOUT • 301
- II\_APPLICATION\_LANGUAGE • 301
- II\_BIND\_SVC\_xx • 301
- II\_C\_COMPILER • 301
- II\_CHARSETxx • 302

---

II\_CHECKPOINT • 302  
 II\_CLIENT • 302  
 II\_COLLATION • 302  
 II\_CONFIG • 302  
 II\_CONNECT\_RETRIES • 302  
 II\_DATABASE • 303  
 II\_DATE\_CENTURY\_BOUNDARY • 304  
 II\_DATE\_FORMAT • 303  
 II\_DBMS\_LOG • 305  
 II\_DBMS\_SERVER • 67, 305  
 II\_DECIMAL • 305  
 II\_DIRECT\_IO • 305  
 II\_DISABLE\_SYSCHECK • 305  
 II\_DML\_DEF • 67, 306  
 II\_DUMP • 69, 90, 306  
 II\_EMBED\_SET • 67, 208, 212, 306  
     dbmserror • 306  
     errorqry • 306  
     eventdisplay • 306  
     gcafile • 306  
     genericerror • 306  
     prefetchrows • 306  
     printgca • 218, 306  
     printqry • 306  
     printtrace • 306  
     qryfile • 306  
     savequery • 306  
     sqlprint • 306  
     tracefile • 306  
 II\_ERSEND • 69, 307  
 II\_FORCE\_HET • 308  
 II\_FRS\_KEYFIND • 67, 308  
 II\_GC\_REMOTE • 67, 308  
 II\_GCA\_LOG • 67, 308  
 II\_GCC\_LOG • 309  
 II\_GCD\_LOG • 309  
 II\_GCN\_LOG • 310  
 II\_GCNxx\_PORT • 69, 310  
 II\_GCx\_TRACE • 67, 309  
 II\_HALF\_DUPLEX • 310  
 II\_HELP\_EDIT • 67, 310  
 II\_INSTALLATION • 69, 310  
 II\_JOURNAL • 69, 90, 310  
 II\_LANGUAGE • 67, 311  
 II\_LOG\_DEVICE • 69, 311  
 II\_MONEY\_FORMAT • 67, 311  
 II\_MONEY\_PREC • 67, 312  
 II\_MSGDIR • 69, 312  
 II\_NULL\_STRING • 67, 312  
 II\_NUM\_OF\_PROCESSORS • 312  
 II\_NUM\_SLAVES • 69, 313  
 II\_OO\_TABLE\_SIZE • 313  
 II\_PATTERN\_MATCH • 67, 313  
 II\_PF\_NODE • 67, 314  
 II\_POST\_4GLGEN • 67, 314  
 II\_PRINTSCREEN\_FILE • 67, 314  
 II\_SQL\_INIT • 67, 76, 314  
 II\_STAR\_LOG • 314  
 II\_SYSTEM • 67, 90, 315  
 II\_TEMPORARY • 67, 315  
 II\_TERMCAP\_FILE • 67, 315  
 II\_TFDIR • 67, 315  
 II\_THREAD\_TYPE • 315  
 II\_TIMEZONE\_NAME • 67, 69, 316  
 II\_TM\_ON\_ERROR • 67, 316  
 II\_TUX\_\_XN\_MAX • 316  
 II\_TUX\_AS\_MAX • 69, 316  
 II\_TUX\_SHARED • 69, 316  
 II\_TUX\_XN\_MAX • 69, 316  
 II\_TUXEDO\_LOC • 69, 316  
 II\_UNICODE\_CONVERTER • 102, 317  
 II\_VNODE\_PATH • 67, 318  
 II\_WORK • 67, 90, 318  
 II\_XA\_TRACE\_FILE • 318  
 IIACP.LOG • 38  
 IIDLDIR • 67, 318  
 iigcb process • 25  
 iigcc process • 25  
 iigcd process • 25  
 iigcn process • 25  
 iilink.log • 39, 205  
 iinamu (utility) • 25  
 iircp.log • 35, 38, 204  
 iivdb.log • 39, 205  
 IMA (Ingres Management Architecture) • 253  
     and OpenROAD • 254  
     audience • 254  
     dbmsinfo constants • 273  
     how it works • 258  
     Management Information Base (MIB) • 255  
     management protocol • 259  
     registering IMA tables • 260  
     removing table registrations • 272  
     restrictions • 278  
     SQL Management Information Base • 260  
     table types • 262  
 include file • 210  
 ING\_ABFDIR • 67, 88, 319

---

---

ING\_ABFOPT1 • 67, 319  
ING\_EDIT • 67, 320  
ING\_PRINT • 67, 320  
ING\_SET • 67, 78, 208, 211, 215, 216, 321  
ING\_SET\_DBNAME • 67, 80, 321  
ING\_SHELL • 67, 321  
ING\_SYSTEM\_SET • 69, 82, 208, 321  
ingpreenv (command) • 61  
Ingres  
    components • 20  
    facility components • 198  
    querying and reporting tools • 27  
Ingres Visual Manager  
    configuring parameters • 226  
    monitoring performance • 223  
INGRES\_KEYS • 321  
ingsetenv (command), global effect • 71  
INIT\_INGRES • 67, 84, 322  
Internet communications • 27  
    configuration parameters for • 55  
iostat (UNIX utility) • 188  
ipcclean (utility) • 184

## J

Journal Analyzer  
    filtering transactions • 238  
    moving transactions to journals • 238  
    recovering individual rows • 245  
    recovering transactions or row changes •  
        244  
    recovering whole transactions • 245  
    redoing individual rows • 245  
    redoing transactions or row changes • 244,  
        251  
    redoing whole transactions • 245  
    selecting transactions or row changes • 243  
    starting • 238  
    viewing transactions • 238, 239  
    viewing transactions on database level •  
        239  
    viewing transactions on table level • 242

## L

lartool.log • 39, 205  
locations  
    changing installation • 87  
    new • 87  
lock information, viewing • 28  
locking system

    configuration parameters • 55  
    described • 33  
lockstat utility • 184  
log files  
    archiver • 203  
    error • 203  
    individual process logs • 205  
    recovery process • 204  
log partitions, configuring multiple • 56  
log\_file\_parts (parameter) • 56  
logging  
    configuration logs • 39  
    individual process logs • 40  
    log files • 37, 202  
    transaction log file • 37, 202  
logging system  
    configuration parameters • 56  
    described • 33  
    viewing summaries • 28  
logicals  
    VMS • 59  
    Windows • 59  
login.com file, adding access commands to •  
    61

## M

matching patterns • 95, 96, 98  
memory cache • 186  
monitor (VMS utility) • 194  
monitoring  
    active user information • 28  
    logging system • 28  
multi collation sequence • 95  
multiple log partitions, configuring • 56

## N

Name Server  
    configuring the • 25  
    described • 25  
    facility code • 199

## O

operating system  
    conventions for different • 16  
Operating System events, setting • 228  
OPF (Optimizer Facility) • 22  
    facility code • 199

---

## P

- parameters
  - restoring • 234
  - viewing a history of changes • 57
- patterns, matching • 95, 96, 98
- permissions, granting for remote commands • 30
- primary transaction log, editing parameters • 57
- printgca (command) • 218
- protected mode, configuring • 47
- ps (UNIX utility) • 187
- PSF (Parser Facility) • 22
  - facility code • 199
- pstat (UNIX utility) • 189

## Q

- QEF (Query Execution Facility) • 22
  - facility code • 199
- QSF (Query Storage Facility) • 22
  - facility code • 199
- query environment • 22

## R

- rcpconfig.log • 39
- RDF (Relation Description Facility) • 22
  - facility code • 199
- recovering
  - immediately or generating scripts • 249
  - individual row changes • 245
  - whole transactions • 245
- recovery
  - dmfrcp process • 35
- recovery process
  - error log • 38, 204
- Recovery Server
  - configuring the • 54
- redoing
  - immediately or generating scripts • 249
  - individual row changes • 245, 251
  - transactions • 251
  - whole transactions • 245
- remote command server (RMCMD) • 30
- remote commands, granting permissions for • 30
- Replicator servers, monitoring • 28
- run time
  - customization • 70

## S

- sar (UNIX utility) • 190
- SCF (System Control Facility) • 22
  - facility code • 199
- server
  - Bridge • 25
  - Bridge Server, configuring the • 26
  - Communications • 25
  - Data Access • 25
  - Data Access Server, configuring the • 26
  - DBMS • 21
  - DBMS Server, configuring the • 49
  - Name • 25
  - Name Server, configuring the • 25
  - Recovery Server, configuring the • 54
  - Star Server, configuring • 53
- set (statement)
  - cancelling options • 213
  - direct set • 209
  - format • 207
  - in startup files • 71, 210
  - in tracing • 206
  - include file • 210
  - io\_trace • 215
  - joinop notimeout • 213
  - printqry • 211
  - qep • 213
- set lock\_trace (statement) • 216
- shell script, tracing • 221
- show (VMS utility) • 194
- spanish collation sequence • 96
- Star Server, configuring • 53
- startsql • 86
- startup • 86
- statement syntax • 17
- symbolic links • 91
- syntax, conventions for describing • 17
- sysdef (UNIX utility) • 191
- sysgen (VMS utility) • 196
- system administrator responsibilities • 19
- system utilities • 184
- system-level commands • 16

## T

- tables
  - moving • 88
  - new • 87
- TERM • 67, 322
- TERM\_INGRES • 67, 322

---

- trace (UNIX command) • 221
- traces
  - GCC trace log • 40, 205
  - I/O • 215
  - locks • 216
  - log • 40, 205
  - queries • 211
  - user-server communications • 218
  - utilities • 206
- transaction log file
  - described • 37, 202
- transaction logs
  - primary and dual • 57
- transactions
  - logging • 37, 202
  - recovery • 36
- transactions, identifying users of • 248
- troubleshooting
  - symbolic links caveats • 91
  - tools • 205

## U

- Unicode • 102, 317
- UNIX utilities
  - I/O Statistics • 188
  - Print Statistics • 189
  - Process Status • 187
  - System Activity Reporter • 190
  - System Definitions • 191
  - Virtual Memory Statistics • 193
- User Errors
  - facility code • 199
- users, monitoring active • 28
- utilities
  - trace • 206

## V

- viewing a history of parameter changes • 57
- Visual Configuration Differences Analyzer
  - accessing • 43
  - and concatenated config.dat files • 235
  - comparing configurations • 234
  - described • 233
  - restoring parameters • 234
  - user tasks • 233
- Visual Performance Monitor
  - monitoring active users • 28
  - user tasks • 28
  - viewing lock information • 28

- viewing log information • 28
- VMS utilities
  - help • 193
  - monitor • 194
  - show • 194
  - sysgen • 196
- vmstat (UNIX utility) • 193

## W

- Windows Diagnostics program • 185
- Windows Performance Monitor program • 186
- write behind • 51