Note

Before using this information and the product it supports, be sure to read the general information under “Notices” on page v.

First Edition (October 1996)

This edition applies to Version 1 of Intelligent Decision Server, 5639–IDS, and to any subsequent releases until otherwise indicated in new editions or technical newsletters. Make sure you are using the correct edition for the level of the product.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1996. All rights reserved.**
Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Notices v
Trademarks v

About This Book vii
Who Should Read This Book vii
We'd Like Your Comments vii

Chapter 1. Introducing the Database Gateway Services 1
How the Database Gateway Services Function 1
The Gateways 2
The Client/Server Dialog 3
The Database Gateway Interface 5
Data Access Tools 5
The Database Gateway Services 6
DB2 for OS/2 Gateway 6
DB2 for OS/2 DRDA Gateway 7
DB2 and MVS CAS Gateways 7
SYBASE Gateway 10
Oracle Gateway 11
Red Brick Gateway 13
DB2 for AIX and DataJoiner Gateways 14
DB2 for AIX DRDA Access 14

Chapter 2. Registering Your Workstation 15
Accessing the Asset Manager 15
Updating Your Site Licenses 16
Adding a Node 17
Adding Services 18

Chapter 3. Configuring a Service 23
Accessing the Database Gateway Service Interface 23
Allocating Licenses 24
Configuring a Database Gateway Service 25
Adding and Configuring a Gateway 26
DB2 for OS/2 and DB2 for OS/2 DRDA Gateways 29
SYBASE and Red Brick Gateways 30
Oracle Gateway 31
DB2 for AIX and DataJoiner Gateways 32
Other Database-Access Tasks 33
Configuring Data Access Tools for Database Access 33
Locating the Database Folders 33

Chapter 4. Administering a Service 35
Accessing the Gateway Administrative Functions 35
Administering the Service and the Gateway 36

© Copyright IBM Corp. 1996

Changing the Service Password 36
Updating Service Status 37
Displaying Service Information 38
Displaying Gateway Configuration 38
Displaying Traffic Information 39
Displaying Gateway Details 41
Modifying Existing Gateway Sessions 41
Sorting Functions 44
Administering Gateways at the Gateway Level 47
Modifying All Sessions on a Specific Gateway 47
Sorting Controls 48
Modifying the Configuration of an Existing Gateway 49
Deleting a Gateway 51
Performing Administration Tasks at the UNIX Interface 51
Starting a UNIX Gateway 51
Using UNIX Gateway Administration Commands 52
Copying Status Information to Other Tools 53
Related Information 53

Glossary 55

IDS Publications 59

Index 61

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594
U.S.A.

Licensees of this program who wish to have information about it for the purpose of enabling (1) the exchange of information between independently created programs and other programs (including this one) and (2) the mutual use of the information that has been exchanged, should contact:

IBM Corporation
555 Bailey Avenue, W92/H3
P.O. Box 49023
San Jose, CA 95161–9023

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks

The following terms are trademarks of IBM Corporation in the United States or other countries or both:

ACF/VTAM	DXT	PROFS
AIX	Enterprise System/9000	PS/2
Application System/400	ES/9000	QMF
APPN	IBM	RACF
AS/400	MVS/ESA	RISC System/6000
Capsule	MVS/XA	RS/6000
DATABASE 2	OfficeVision	S/390

© Copyright IBM Corp. 1996

DataJoiner	OfficeVision/VM	SQL/DS
DB2	Operating System/2	System/360
DB2 for OS/2	Operating System/400	System/370
DB2/6000	OS/2	Virtual Machine/Enterprise Systems Architecture
DFSORT	OS/400	VM/ESA
Distributed Relational Database Architecture	Personal System/2	VTAM
DRDA	Presentation Manager	Win-OS/2

C-bus is a trademark of Corollary, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Other company, product, and service names, which may be denoted by a double asterisk (**), may be trademarks or service marks of others.

About This Book

The IBM Intelligent Decision Server (IDS) *Configuring and Administering Database Gateway Services* is intended to be used as both a guide and as a reference. You can use this publication to configure and administer the database gateway services and their associated gateways. In addition, this book will help you to quickly locate the windows, buttons, and choice fields within the *Database Gateway Services* interface.

Who Should Read This Book

This book is intended for all system administrators responsible for administration and maintenance of any of the IDS Database Gateway Services. Information provided in this publication is written on the assumption that you have a working knowledge of:

- IBM Communications Manager/2 (for OS/2 database gateway services only)
- IDS system administration procedures
- IDS tools operation
- UNIX system administration (for UNIX database gateway services only)

We'd Like Your Comments

Your feedback is important to us in providing you with the most accurate and high-quality information. If you have any comments about this book or any other IDS documentation, please visit our Website at:

http://www.dss.ibm.com

There you'll find the Feedback Connection page where you can enter comments and send them to us.

© Copyright IBM Corp. 1996

Chapter 1. Introducing the Database Gateway Services

The IBM Intelligent Decision Server (IDS) provides a variety of desktop tools that can access databases and analyze data. By using the IDS database gateway services, desktop tool users can access a variety of relational databases and run Structured Query Language (SQL) queries and capsules.

This book describes how to configure the workstation on which the gateway service resides, how to configure the gateway services and their gateways, and how to perform database gateway administration tasks.

With few exceptions, the tasks and procedures described in this book apply equally to all relational database management systems (RDBMSs) and Cooperative Application Services (CAS) applications accessible through the IDS database gateway services.

This chapter describes the gateway services and their associated gateways, and contains the following topics:

"How the Database Gateway Services Function"
"Data Access Tools" on page 5
"The Database Gateway Services" on page 6

How the Database Gateway Services Function

The database gateway services provide data and protocol conversion between the IDS data access tools and an RDBMS or a CAS application. CAS is a distributed IDS facility that enables IDS desktop users on a local area network (LAN) to run command scripts. The command scripts access applications and databases on any MVS host system supported by IDS Host Services.

The IDS desktops can send and receive data through the gateways. The gateways establish *sessions* (or connections) with the RDBMS or CAS application. The sessions are allocated when a request is issued by one of the IDS data access tools. Because the interface is uniform, you can administer services for a number of different RDBMSs and CAS applications the same way.

For database access, the database gateway service accepts relational data, database management instructions, and SQL queries from the tools. Next, it converts the data into a format expected by the RDBMS. Similarly, the database gateway service accepts data from the RDBMS, converts it, then forwards it to the appropriate desktop tools.

For CAS-application access, the service accepts CAS data or application requests. The server then sends the data and requests to the host system for processing, then converts them into a format that is accepted on each end (client and server).

In either case, the IDS database gateway services convert data formats appropriately in both directions to ensure that the desktop and the RDBMS or CAS application receive data and instructions in the correct format.

© Copyright IBM Corp. 1996

The data access capabilities of the database gateway services are based on a client/server architecture. This allows data to be distributed on a variety of workstations that can be located remotely or attached directly to IDS networks. For example, the IDS desktop can act as the client, and the workstation on which the database resides can act as the server. The client/server architecture enables all users to work from the same data.

The Gateways

The IDS database gateway services are designed to support one or several database gateways simultaneously. Multiple gateway types and RDBMS connections can be configured within a single service.

If you want to access data from an MVS system, you need to use the following OS/2 gateways:

DATABASE 2 (DB2)
MVS Cooperative Application Services (CAS)

When you need to access data from workstations running UNIX-based software, you must use the following UNIX** gateways:

SYBASE**
Oracle**
Red Brick**
DB2 for AIX
DataJoiner

When you need access to a database residing on a workstation, you must use DB2 for OS/2.

When you need to access a remotely located host system, you must use the DB2 for OS/2 Distributed Relational Database Architecture (DRDA) gateway.

For DB2 for OS/2 remote gateway access, the IBM Communications Manager/2 and IBM Distributed Data Connectivity Services/2 (DDCS/2) must be also be configured to establish connectivity to a remote DRDA database. For more information about this access, see “DB2 for OS/2 and DB2 for OS/2 DRDA Gateways” on page 29.

All of these gateways perform protocol conversion and allocate resources to the administration interface for managing resources and controlling access to and from the workstation containing the database.

Table 1 on page 3 lists the IDS database gateway services, the types of database gateways each service supports, the types of workstations on which the databases reside, and the software running on the workstation containing the database.

Table 1. Types of database gateway services

Database gateway service	Gateways supported by the service	Type of computer on which the database resides	Operating system running on the computer containing the database
OS/2 Database Gateway Service	DB2 for OS/2, DRDA, DB2 for MVS	IBM PC or any host computer that can support MVS/ESA	IBM OS/2 Warp, MVS/ESA, version 4, and DB2, versions 2.3 or 3.0
AIX Database Gateway Service	SYBASE, Oracle, Red Brick, DB2 for AIX, DataJoiner	IBM RISC System/6000	IBM AIX, version 3.2 or higher
HP-UX Database Gateway Service	SYBASE, Oracle, Red Brick	HP 9000, Series 800	HP-UX, version 9.04 or higher
Solaris Database Gateway Service	SYBASE, Oracle	Sun SPARC platform	Sun Solaris operating system, version 2.3

The Client/Server Dialog

IDS tools for database access are designed for direct communication with relational database management systems. To achieve that direct communication, several levels of communication are required to implement the graphical view users of IDS desktop tools see.

In a similar scheme, the CAS transformer runs on an IDS desktop while the command script associated with the transformer runs in an MVS address space on the host. Therefore, several levels of communication are necessary to link the host components and LAN components of the CAS application.

The database gateway service and the IDS RDBMS or CAS application interface work together to provide data access capabilities for IDS users. As a result, the gateway service and the interface perform a variety of tasks, including:

Providing a uniform database interface to IDS desktop tools
Monitoring gateway traffic
Administering gateway configuration
Translating network protocols
Allocating internal resources
Interpreting SQL

Depending on the type of database you use, the database and its database gateway service might reside on the same workstation or on separate workstations. The data access and CAS application paths are similar for each IDS gateway. The paths differ only in the location of the data. Figure 1 on page 4 shows the processing layers in the data access path between the client and server.

The IDS tools present the uniform graphical representation of the database to desktop users. For RDBMS access, the database gateway service and IDS RDBMS interface make it appear to IDS users that the database is running in the desktop. Together, these components hide the location and type of database being accessed, as well as

the complexity with which database connections are made. When desktop tools send data and SQL statements to the database gateway service, the service forwards them through the IDS RDBMS interface to the RDBMS or through the CAS Application Program Interface (API) to the host application. The RDBMS or host application returns the results to the desktop tools for display or further manipulation.

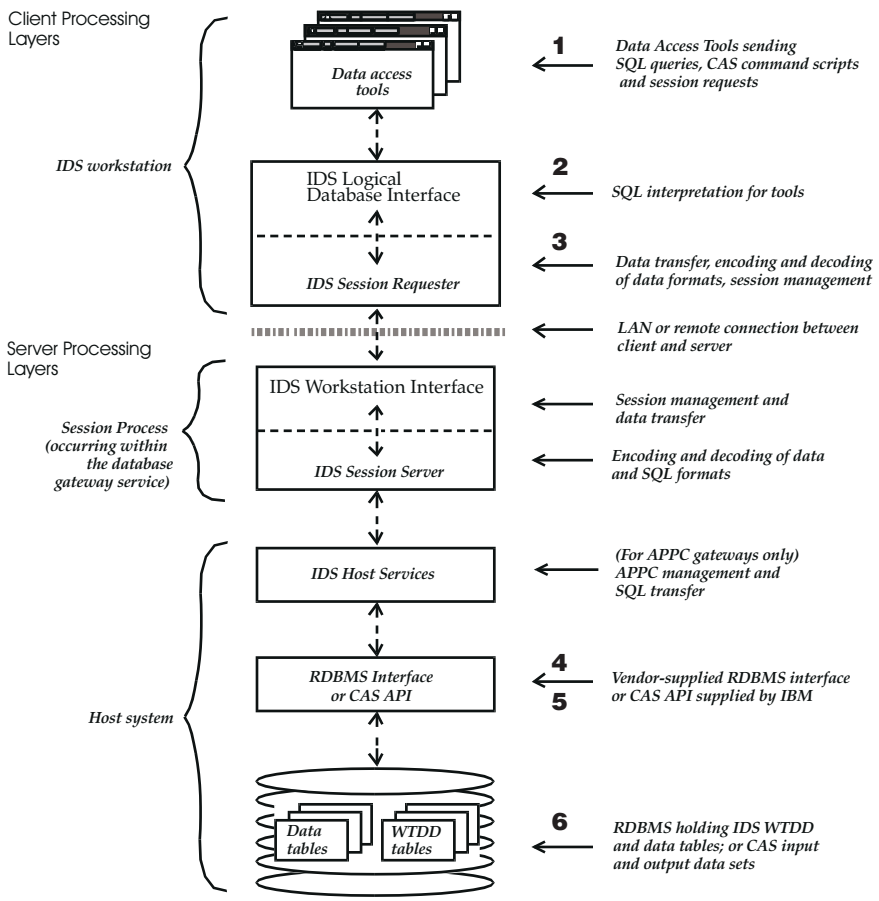

Figure 1. Client and server processing layers

The following numbered items correspond to the numbered items in Figure 1:

1. Data access tools can be run from any workstation with appropriate data access privileges.
2. The IDS logical database interface enables database gateway access.
3. The IDS session requester hides the location of the database.
4. The vendor-supplied RDBMS interface receives data and session requests from the session server, passes SQL statements, and returns data to the database.

5. The CAS API transfers data between the IDS workstation and the host application. The API must be invoked by a REXX command language procedure executed by Host Services through TSO/E. The CLIST command language can call REXX procedures to invoke the API, but cannot call the API directly.
6. An RDBMS can reside in the same workstation as its RDBMS interface or in a separate one.

The Database Gateway Interface

IDS provides a uniform interface for the IDS database gateway services that can be accessed with the System Administration (SA) tool. The interface, known as the Database Gateway interface, enables you to configure and administer the database gateway services and their associated gateways.

After you install the gateway service, you must configure it before you can use it. You'll need to specify:

The network address of the database gateway service workstation to which you want to connect.

The types of gateways.

The number of sessions for each gateway.

Whether you want to participate in a load-balancing pool.

The type of database that you want to connect to.

You can perform these configuration tasks by using the database gateway interface. See Chapter 3, “Configuring a Service” on page 23 for detailed information about configuring the gateway services and their gateways.

After you configure the gateway service and its gateways, you must prepare the data access icons, WTDD and data tables, and user accounts that you will need. For details, see these IDS books:

DBA Tool

Administering Databases with IDS

Data Access Tools

Because the IDS graphical interface to all supported RDBMSs is uniform throughout the database gateway services products, desktop users, as a rule, need not be aware of what type of database they're using. Desktop users view the database through one of the following data access tools, which provide direct access to database tables:

Query
Reporter
DBA
SQL Entry
Forms Design
SQL Input

Browser

The WTDD consists of 10 tables, which provide comprehensible object names to data access tools. The WTDD defines the *logical database*, which provides the information needed to logically map the data between the data access tools and the physical database. The rows of data and the mapping between the physical and logical databases are stored in the *physical database*, which is the actual data table. The data access tools use the WTDD to translate commands and data between the logical and the physical databases. See *Administering Databases with IDS* for information about creating the WTDD and data tables.

When you access the database through data access tools, such as the Query tool, you can see how the database tables are organized. You can design queries customized for their applications and display the results. When a query is run in a capsule, the capsule can send the results automatically to other desktop tools to create reports or display the results graphically.

Database vendors provide various ways to access remote databases. See *Administering Databases with IDS* for details about configuring IDS data access tools to reach remote databases.

The Database Gateway Services

The IDS database gateway services support several gateways, which can be configured in different ways. The configuration in place at your site depends upon:

The type of database gateway service you purchased
The databases you want to access
The interface you use between the database and the database gateway service
The types of workstations you use

The following gateway services are supported by IDS:

"DB2 for OS/2 Gateway"
"DB2 for OS/2 DRDA Gateway" on page 7
"DB2 and MVS CAS Gateways" on page 7
"SYBASE Gateway" on page 10
"Oracle Gateway" on page 11
"Red Brick Gateway" on page 13
"DB2 for AIX and DataJoiner Gateways" on page 14

DB2 for OS/2 Gateway

A DB2 for OS/2 gateway can be accessed through the OS/2 database gateway service. This gateway provides access to IDS desktop tools for DB2 for OS/2 databases residing on a workstation.

When an IDS LAN desktop requests access to a database on a workstation, the following process occurs:

1. The desktop sends the request to a DB2 for OS/2 gateway session, by means of the OS/2 database gateway service.
2. The gateway session forwards the request to the DB2 for OS/2 interface, which converts it into a form acceptable by the DB2 for OS/2 database.
3. The DB2 for OS/2 interface then sends the request to the DB2 for OS/2 database.
4. Returning data (if any) travels the same path, in reverse.

DB2 for OS/2 DRDA Gateway

A DB2 for OS/2 DRDA gateway can be accessed through the OS/2 database gateway service. This gateway provides access to IDS desktop tools through a workstation supporting a DB2 for OS/2 DRDA gateway connecting to an SQL/DS database or through a DB2 target database residing on a remote host system configured to support DRDA. Code is included to provide the correct database type information to the data access tools.

When you connect through DRDA links to a remote database, the Host Services software is not required on the remote DRDA workstation.

Except for the fact that the target database name is explicitly defined in the Icon Options window of the data access tool, this secondary access path (using the DRDA link) is transparent to IDS users.

When an IDS LAN desktop requests access to a database on a workstation, the following process occurs:

1. The desktop sends the request to the DB2 for OS/2 DRDA gateway session, by means of the OS/2 database gateway service.
2. The gateway session forwards the request to the DB2 for OS/2 interface, which is configured with DDCS/2. This interface converts the request into a form acceptable by a DRDA server.
3. After the request is converted, DDCS/2 sends the request to the Communications Manager/2.
4. Communications Manager/2 passes the request to the DRDA server component of the RDBMS by means of a DRDA link.
5. Returning data (if any) travels the same path, in reverse.

Because IDS is aware of the end database source, the proper translations are performed to correctly format data requests and results.

DB2 and MVS CAS Gateways

The DB2 gateway (which includes DB2-DRDA gateway access) and the MVS CAS gateway can be accessed through the OS/2 database gateway service. IDS desktop tools are designed for direct communication with other LAN-based services. The IDS Host Services is comprised of MVS address spaces that access DB2 databases on behalf of IDS data access tools. The CAS transformer runs on an IDS desktop, whereas the command script associated with the transformer runs in an MVS address

space on a host system. The widely distributed nature of these components requires several levels of communication to implement the simple view that the system presents to the IDS desktop user for either type of data access.

The Host Services components reside in MVS address spaces. They control, maintain, and provide service administration capabilities for APPC conversations with processes in the database gateway service. They accept requests from the database gateway processes and forward them through the IDS RDBMS interface and the DB2 Call Attach Facility (CAF). Each connection between Host Services and DB2 is established through CAF. When the DB2 subsystem returns results, Host Services passes them to the gateway session process.

The database gateway service uses several internal IDS processes common to both DB2 and CAS sessions to manage communications and convert data. These processes manage allocation of internal service processes and resources. A control process maintains APPC conversations to the host service manager and manages cancel requests. The database gateway service uses a session process to perform data and format conversion. A database gateway service can use multiple session processes, one for each DB2 or CAS gateway session configured for the service. The OS/2 database gateway service can accommodate several other types of gateways supporting access to different data sources and running concurrently (for example, DB2 for OS/2). The session process forwards outgoing requests to the host system and receives incoming responses that are returned to the IDS desktop tools.

When the host system is accessed through an APPC connection, the session process performs bidirectional conversion between the network protocol and APPC protocol. If the connection is interrupted, the session process restarts the connection. One session process is active for each DB2 or CAS session between the gateway service and the host system. Each DB2 or CAS gateway within the service is capable of containing multiple sessions as defined through the asset management process during installation.

The database gateway service uses Communications Manager/2 residing on the same workstation to establish SNA APPC conversations with Host Services through Virtual Telecommunications Access Method (VTAM). You can choose the actual path between the VTAM address space and the gateway service from the several SNA configurations. At the host-system entry point into the SNA network, the IDS communications interface accesses the APPC communications protocol. The IDS communications interface uses this protocol to send and receive messages between IDS Host Services and the communications network.

DB2 and CAS MVS Communications Paths

This section describes the communication paths for the DB2 and MVS CAS gateways, which use the IDS Host Services software.

DB2 Communications Path: IDS data access tools are configured to access a target database. The entry point into the DB2 subsystem can reside on the host system, while the target database can reside on a separate workstation. In these cases, the database can be accessed remotely through DRDA links.

When a user requests access to an MVS system, the following process occurs:

1. IDS desktop requests for relational data, database management instructions, and SQL queries are sent to the DB2 gateway by means of the OS/2 database gateway service.
2. The DB2 gateway session follows a conversation path to the host system automatically established by Communications Manager/2, which resides on the same workstation as the database gateway service.
3. Communications Manager/2 passes the requests through a TCP/IP or database specific SNA connection to the appropriate IDS service address space, which resides on the host system.
4. After verifying the user ID, the receiving service address space passes the request to the DB2 subsystem.
5. The DB2 subsystem accesses the data in the DB2 database.
6. Returning data (if any) travels the same path, in reverse.

For more information about DB2 and its relationship to Host Services, see *Administering Host Services for MVS: DB2 and Cooperative Application Services*.

MVS CAS Communications Path: When a user invokes an IDS capsule, the following process occurs:

1. The CAS transformer (imbedded in the capsule) sends a command procedure (such as a REXX procedure) to the MVS CAS gateway by means of the OS/2 database gateway service.

   Procedures written in other command languages must be invoked by a REXX procedure, which returns results from the host system.
2. The MVS CAS gateway session establishes conversations with the host through Communications Manager/2, which resides on the same workstation as the database gateway service.
3. CM/2 passes the requests through an SNA connection to the appropriate CAS service address space, which resides on the host system.
4. The receiving CAS service address space executes the procedure against any host application through TSO/E.
5. Returning data (if any) travels the same path, in reverse. A REXX procedure can return results to desktop tools from the host system for display or further manipulation.

For more information about CAS and its relationship to Host Services, see the following IDS books:

Cooperative Application Services: Application Development

Host Services for MVS: DB2 and Cooperative Application Services

Connecting to Other Databases from a DB2 Gateway through DRDA

If your DB2 subsystem provides DRDA requester functions, then IDS data access tools that can access a DB2 database can also be configured to access any remote DRDA server databases that are connected to the DB2 subsystem.

The remote database can be any DRDA server database, such as another DB2 subsystem. When you connect to a remote database through the DB2 gateway, Host Services is not required on the remote DRDA workstation.

When an IDS LAN desktop requests access to a host system, the following process occurs:

1. IDS desktop requests for relational data, database management instructions, and SQL queries are sent to the DB2 gateway by means of the OS/2 database gateway service.
2. A DB2 gateway session establishes conversations with the host system through Communications Manager/2, which resides on the same workstation as the database gateway service.
3. CM/2 passes the requests through an SNA connection to the appropriate IDS service address space, which resides on the host system.
4. The receiving service address space passes the request to the DB2 subsystem via the call attach facility.
5. The DB2 subsystem forwards the request to the remote RDBMS (for example, an SQL/DS database), if it is satisfying a DRDA request.
6. Returning data (if any) travels the same path, in reverse.

With DRDA, IDS data access tools are configured to access the target database, and the access path is transparent to users.

For additional information about data access through DRDA, see *Administering Databases with IDS*.

SYBASE Gateway

Depending on the UNIX platform being used, the *SYBASE* gateway is supported by:

AIX database gateway service
HP-UX database gateway service
Solaris database gateway service

See Table 1 on page 3 for the list of workstations that can be used with each of the UNIX gateways.

The software used for the SYBASE gateway requires the SYBASE Open Client programming interface. This product acts as an interface that enables the gateway to establish a communications path and pass SQL statements to the SYBASE SQL server, and to receive the results of these commands.

This section provides two example configurations for the SYBASE gateway. The data access paths for these configurations are identical, but the components reside in different locations.

When an IDS LAN desktop requests access to a database on one of the supported UNIX platforms, the following process occurs:

1. The desktop sends a request to the SYBASE gateway by means of the AIX database gateway service, HP-UX database gateway service, or Solaris database gateway service.
2. A gateway session passes the request to the SYBASE Open Client programming interface.
3. The Open Client programming interface translates the data into a form that is understood by the SYBASE SQL server.
4. When the translation is complete, the Open Client programming interface forwards the data to the SQL server.
5. Returning data (if any) travels the same path, in reverse.

When the SYBASE gateway and the SYBASE SQL server reside on separate workstations, the SYBASE gateway communicates with the SQL server over a network link.

The protocol used on that link can be any network protocol supported by the SYBASE SQL server, such as TCP/IP, Banyan VINES, or Novell's NetWare. The network link acts as an interface between the SYBASE Open Client programming interface and the SYBASE SQL server.

Oracle Gateway

Depending on the UNIX platform being used, the Oracle gateway can be accessed through:

The AIX database gateway service
The HP-UX database gateway service
The Solaris database gateway service

See Table 1 on page 3 for the list of workstations that can be used with each of the UNIX gateways.

When the database gateway and the RDBMS are located on different workstations, the Oracle gateway requires the SQL*Net programming interface on both workstations. This interface enables the gateway to establish a communications path and pass SQL statements to the Oracle server, as well as to receive the results of these commands.

When the database gateway and the Oracle RDBMS are on the same workstation, the Oracle gateway does not require the SQL*Net programming interface.

This section provides two example configurations for the Oracle gateway. The data access paths for these configurations are almost identical, except that:

When the Oracle gateway and the RDBMS reside on separate workstations, two copies of the SQL*Net programming interface are required: one for the system containing the Oracle server and one for the UNIX platform containing the database gateway service.

The components reside in different locations.

When the gateway and RDBMS reside on the same workstation, the following process occurs when an IDS LAN desktop requests access to a database on one of the supported UNIX platforms:

1. The desktop sends a request to the Oracle gateway by means of the AIX database gateway service, HP-UX database gateway service, or Solaris database gateway service.
2. The gateway session translates the data into a form that is understood by the Oracle server.
3. When the translation process is complete, the gateway session passes the data to the Oracle server.
4. Returning data (if any) travels the same path, in reverse.

When the Oracle gateway and the Oracle server reside on separate workstations, the Oracle gateway communicates with the SQL server over a network link.

The protocol used on that link can be any network protocol supported by the RDBMS, such as TCP/IP, Banyan VINES, or Novell's NetWare. The network link acts as an interface between the two Oracle SQL*Net programming interfaces.

Additionally, when the gateway and server are on separate workstations, two copies of the Oracle SQL*Net programming interface are required.

The following process occurs when an IDS LAN desktop requests access to a database on a remote UNIX platform:

1. The desktop sends a request to the Oracle gateway by means of the AIX database gateway service, HP-UX database gateway service, or Solaris database gateway service.
2. The gateway passes the request to the SQL*Net programming interface.
3. The SQL*Net programming interface translates the data into a form that is understood by the Oracle server.
4. When the translation is complete, the SQL*Net programming interface forwards the data to the other SQL*Net programming interface on the remote workstation.
5. The other SQL*Net programming interface then passes the data to the Oracle server.
6. Returning data (if any) travels the same path, in reverse.

Red Brick Gateway

The Red Brick gateway is supported by the following database gateway services, depending on the UNIX platform being used:

The AIX database gateway service
The HP-UX database gateway service

See Table 1 on page 3 for the list of workstations that can be used with each of the UNIX gateways.

The software used for the Red Brick gateway requires the SYBASE Open Client programming interface. This product acts as an interface that enables the gateway to establish a communications path and pass SQL statements to the Red Brick Open Server Gateway, which forwards the data to the Red Brick Warehouse server. The Open Client programming interface also enables the gateway to receive the results of the SQL statements.

This section provides two example configurations for the Red Brick gateway. The data access paths for these configurations are identical, but the components reside in different locations.

When an IDS LAN desktop requests access to a database on one of the supported UNIX platforms, the following process occurs:

1. The desktop sends a request to the Red Brick gateway by means of the AIX database gateway service or HP-UX database gateway service.
2. The gateway session passes the request to the SYBASE Open Client programming interface.
3. The Open Client programming interface translates the data into a form that is understood by the Red Brick Warehouse server.
4. When the translation is complete, the Open Client programming interface forwards the data to the Red Brick Open Server Gateway.
5. The Red Brick Open Server Gateway then passes the data to the Red Brick Warehouse server.
6. Returning data (if any) travels the same path, in reverse.

When the Red Brick gateway and the Red Brick Warehouse server reside on separate workstations, the SYBASE gateway communicates with the SQL server over a network link.

The protocol used on that link can be any network protocol supported by the RDBMS, such as TCP/IP, Banyan VINES, or Novell's NetWare. The network link acts as an interface between the SYBASE Open Client programming interface and the Red Brick Open Server Gateway.

DB2 for AIX and DataJoiner Gateways

The DB2 for AIX and DataJoiner gateways are supported by the AIX database gateway service. See Table 1 on page 3 for the list of workstations that can be used with each of the UNIX gateways.

The software used for the DB2 for AIX or DataJoiner gateway require DB2 for AIX, to access DB2 for AIX or DataJoiner databases residing on the same machine.

When an IDS LAN desktop requests access to a DB2 for AIX or DataJoiner database, the following process occurs:

1. The desktop sends a request to the DB2 for AIX or DataJoiner gateway by means of the AIX database gateway service.
2. A gateway session passes the request to the DB2 for AIX or DataJoiner server, which retrieves and sends the data on the same path.

When the DB2 for AIX database server and the DB2 for AIX gateway reside on separate workstations, the DB2 for AIX gateway communicates with the DB2 for AIX server over a TCP/IP or APPC network link. DataJoiner follows a similar path in this situation.

DB2 for AIX DRDA Access

A DB2 for AIX gateway acts as a requester and can connect to an SQL/DS or DB2 target database residing on a remote host system configured to support DRDA.

When you connect through DRDA links to a remote database, the Host Services software is not required on the remote DRDA workstation.

The following conditions must be met for DRDA access:

DDCS/6000, which provides DRDA application requester functions, must be installed on the machine where the gateway service is installed.

SNA links need to be established to the host.

Databases on remote hosts need to be cataloged on the AIX side.

Chapter 2. Registering Your Workstation

The IDS System Administration (SA) tool validates assets and assigns the gateway service to the workstation (or node) on which the IDS gateway service software is installed. These tasks are considered configuration tasks.

Every node has a unique network number. A node enables the services that other IDS services, including OS/2 and UNIX database gateway services, will use. OS/2 systems are limited to one node per system; UNIX systems do not have this limitation.

The asset management feature of the SA tool (called the *asset manager*) defines IDS components by licensing specific sets of them. You can access the asset manager with the SA tool, which provides the service interfaces that enable you to configure and administer specific IDS services.

This chapter describes how to update your site licenses, add your node to the list of IDS nodes (if the database gateway service was installed on a new node), and register your workstation as a database gateway service node.

The following tasks are described in this chapter:

"Accessing the Asset Manager"

"Updating Your Site Licenses" on page 16

"Adding a Node" on page 17

"Adding Services" on page 18

Accessing the Asset Manager

The SA tool provides access to all administrative interfaces. An authorized SA tool user can start it from any IDS workstation. The SA tool resides in the Administrator Icons file drawer. Within the SA tool is the asset manager interface, which enables you to add nodes, assign IDS services to nodes, and validate assets.

The asset manager must be informed of any new node or service that is installed within its realm. If the asset manager is not informed of a newly added node or service, IDS cannot use it until it is registered.

Before you add a node or service, ensure that you have a copy of the SA tool icon on your desktop. If not, obtain a copy of it from the Administrator Icons file drawer, which can be found on the system desktop. Additionally, you must know the node name and network address of the newly added node.

To open the asset manager:

1. Double-click on the SA tool icon.

   The SA tool window opens.

2. In the **Service Type** field, select **Asset Management** from the pull-down menu.

© Copyright IBM Corp. 1996

3. Click on the **Connect to Service** button. The Asset Management password window opens.
4. Type the password for the asset management interface and press enter. The Asset Management window opens.

   Because **Nodes** is selected (by default) in the **Display** field, a list of the IDS network's workstations that are known by the asset manager opens.

Updating Your Site Licenses

Update your site licenses before using a new service of any type on your network.

To update your licenses to include a database gateway service:

1. In the Asset Management window, select **Assets** from the **Display** field.

   The window displays the assets (services and tools) your site has purchased.

   This Asset Management window provides:

   The name of the vendor from which you purchased the software, in the *Vendor* column.

   A complete list of previously purchased assets, in the *Asset* column.

   The number of licenses your site purchased, in the *Number of licenses* column.

   The number of licenses currently available, in the *Available Licenses* column.
2. Click on the **Modify Asset Configuration** button to specify the assets you purchased and the quantity of each.

   The Asset Mgmt - Asset Configuration window opens.

   This window provides a complete list of IDS assets that you can purchase. The name of each asset is followed by a number enclosed in parentheses. This is the number of licenses specified in the current configuration.

   The numbers listed in the fields to the right are the number of licenses in the proposed configuration. For example, INR Service (10), means that your realm currently owns 10 licenses of INR service.
3. Type the total quantity of each gateway-related asset your site purchased (current assets plus any newly purchased assets) in the associated fields to the right.

   You must enter all of your asset quantities in one session.
4. In the **Verification String** field, type the verification string that IBM supplied when you most recently purchased your IDS assets.

   The verification string is an 8-digit, hexadecimal number that is valid for your entire configuration.
5. Click on the **Enable New Configuration** button to validate your new gateway-related assets.

If no errors are encountered, you'll receive a message at the bottom of the IDS desktop:

Apply completed successfully.

6. Close the Asset Mgmt - Asset Configuration window.
7. From the Asset Management window, select **Nodes** in the **Display** field. The Asset Management window opens, showing node information.

Adding a Node

The database gateway service can be installed on a node with or without existing IDS software. If the service was installed on a node that contains IDS software, skip to the section, “Adding Services” on page 18 to add and register the service with the asset manager.

To register a new node with the asset manager through the SA tool:

1. From the Asset Management window, click on the **Add Node** button. The Asset Mgmt - Node Details window opens.
2. For a TCP/IP network, in the **Network Address** field, enter the IP address of the Ethernet or token-ring adapter such as 143.241.2.156, which will be used for the Database Gateway services,
3. In the **Node Name** field, type the name that is to be associated with the workstation.

   You should have a systematic naming scheme for all of your nodes, based on the usual user's name, location, or some other aspect of your particular site.
4. In the **Comment** field, type any text that will help you further identify the person or group to whom the workstation belongs. The comment must be no more than 63 characters long.
5. In the **Status** field, leave the default setting of **Active**. If you want to register the node but keep it unusable until some future time, change the setting to **Inactive**.

   See the *System Administration Guide and Reference* for a thorough description of the **Active** and **Inactive** settings.
6. From the window header, click on the **Register Node Information** button to register this information.

   The address in the **Network Address** field is now preceded by **0#** and followed by **:0**. You can ignore these additions.

The workstation is now added to the list of IDS nodes, which appears in the Asset Management window.

Keep the Node Details window open so you can add the database gateway service to the newly added node.

Adding Services

If you performed the tasks described in the section “Adding a Node” on page 17, skip to step 3 to add a service to the node.

To add a service and register it with the asset manager:

1. From the list of registered IDS nodes, find the name of the node (such as Mktg_3) to which you want to add a database gateway service. Ensure that the network address in the list matches that of the actual node, then select that node.
2. Click on the **Show Details** button to display information about the node. The Asset Mgmt - Node Details window opens.

   If there is nothing displayed in the *Show Details* portion of the window, as shown in Figure 2, then no assets are assigned and registered to this node.

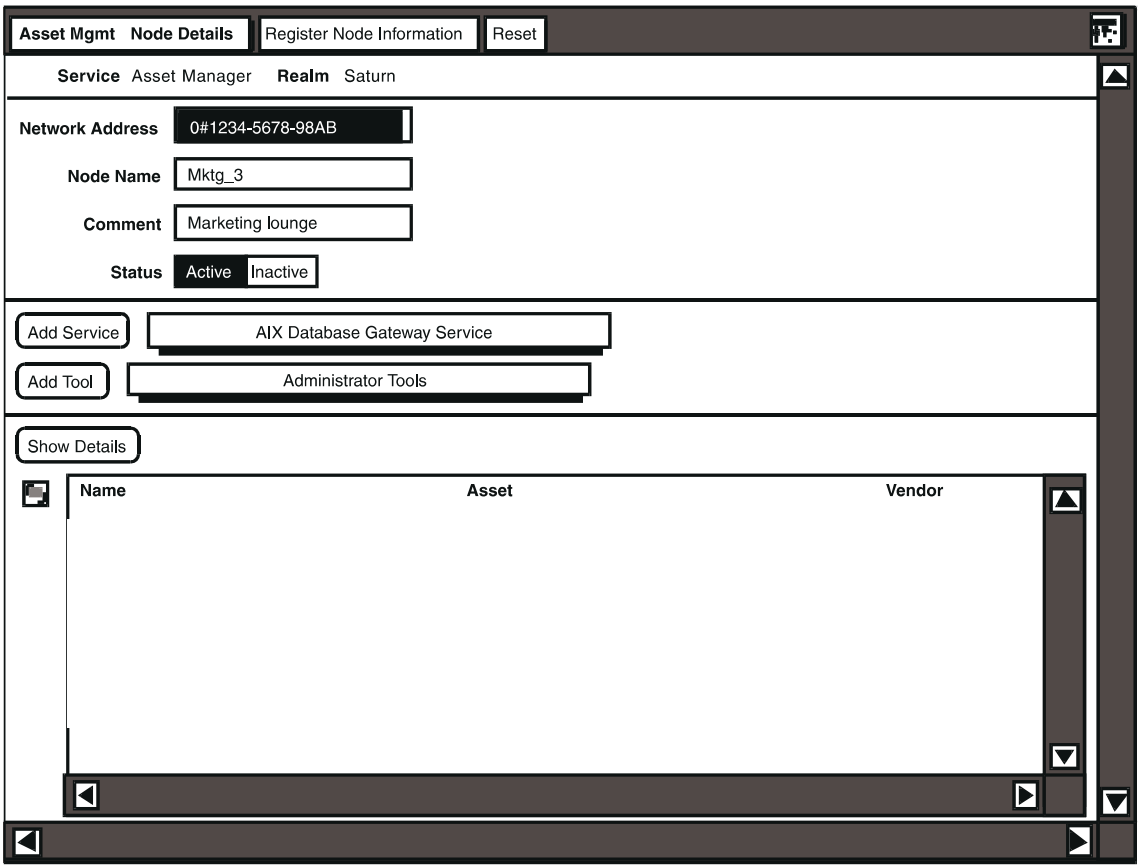

Figure 2. Asset Mgmt - Node Details window of an existing node without assets

3. Pull down the menu next to the **Add Service** button, as shown in Figure 3 on page 19.

The menu lists every type of service that was validated during IDS installation. If the window is near the bottom of your screen, you might have to scroll down to see the entire list.

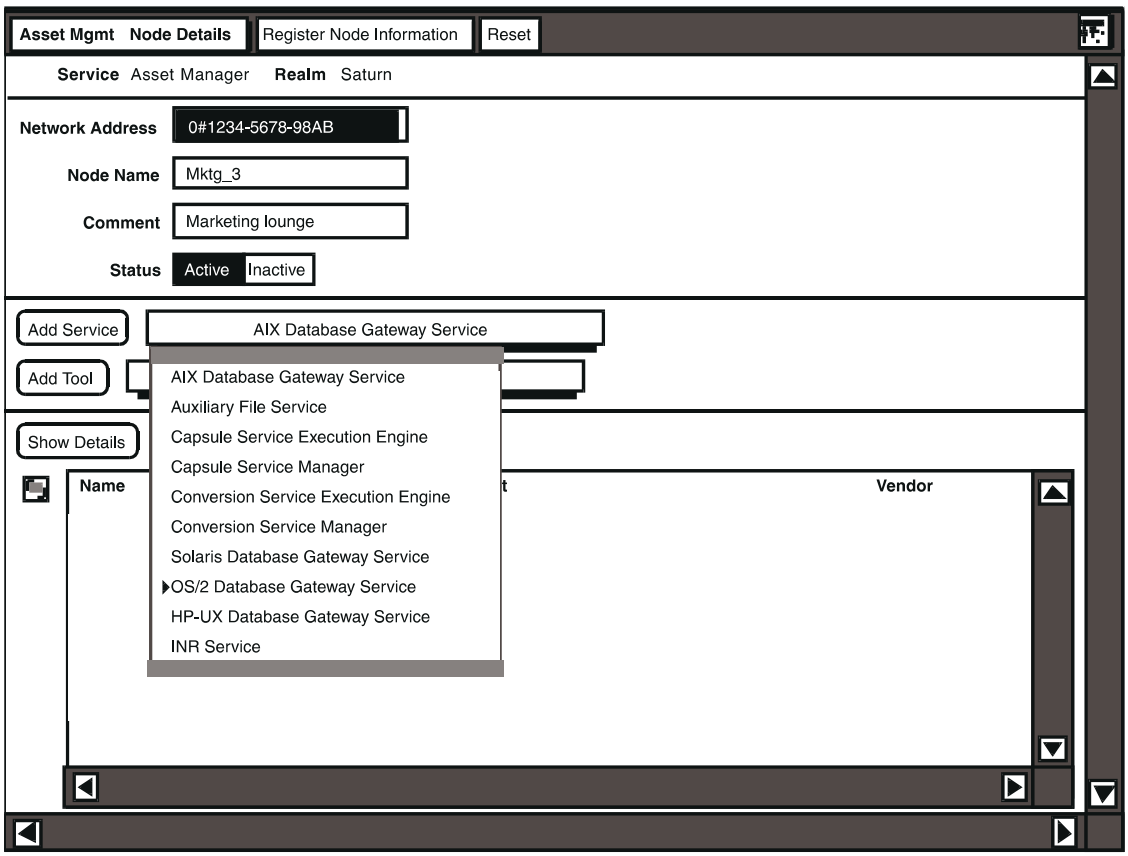

Figure 3. Asset Mgmt - Node Details window, displaying the Add Service menu

4. Select the type of database gateway service you want running on your node, such as OS/2 Database Gateway Service.

   When you choose a service, it is displayed in the **Add Service** field.

5. Click on the **Add Service** button.

   The Asset Mgmt - Service Details window opens, as shown in Figure 4 on page 20.

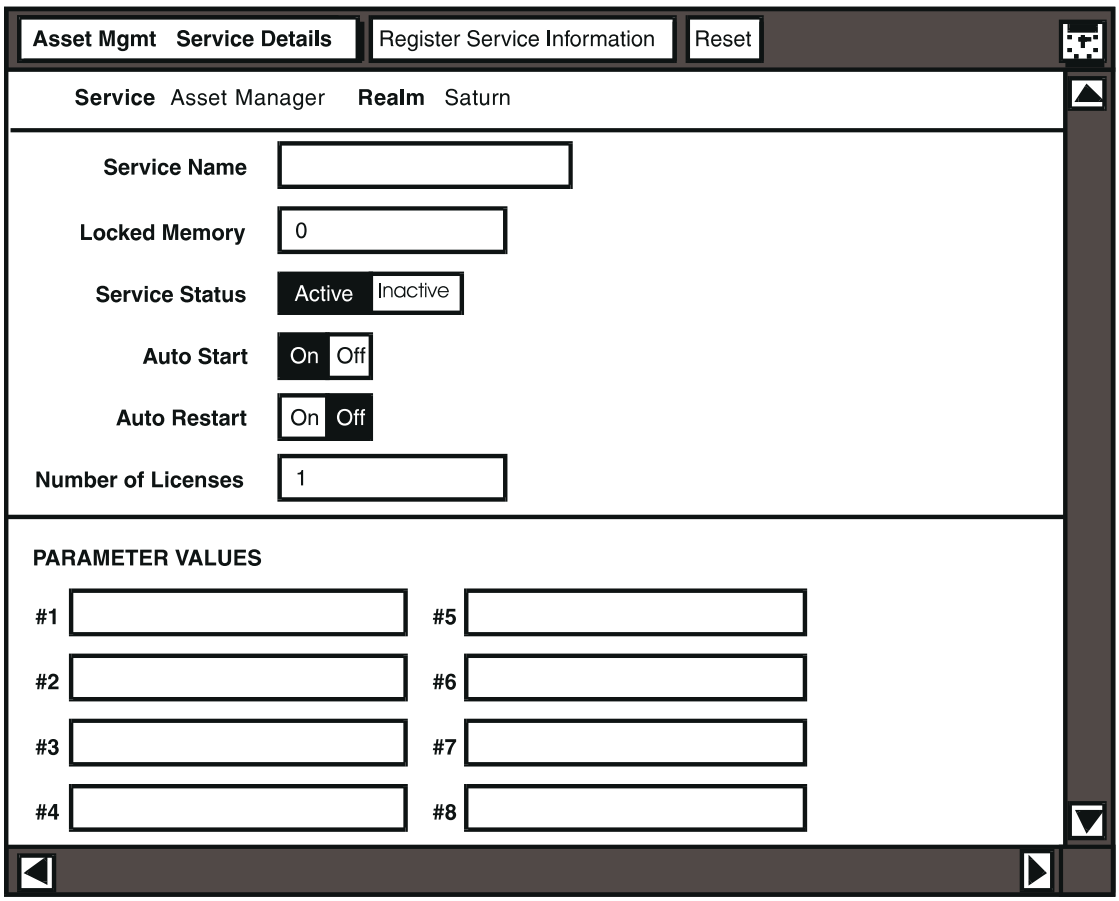

Figure 4. Asset Mgmt - Service Details window

6. In the **Service Name** field, type the name you want to assign to the database gateway service, such as Marketing DB.

   The name can include spaces and punctuation, and must be no more than 63 characters. The service name cannot be the same as any of the gateway names. An easy way to avoid this naming problem is to add the word Service to your gateway name. For example, the service name for the Marketing DB gateway would be MarketingDBService.

7. In the **Locked Memory** field, leave the default value of 0.
8. In the **Service Status** field, leave the default setting of **Active** as is.

   If you want to register the service, but not use it until later, select **Inactive**.

   The *System Administration Guide and Reference* discusses the **Active** and **Inactive** settings in more detail.

If you change the status in the **Auto Start** field to **On**, service will start automatically upon registration and whenever the server is booted.

9. In the **Auto Restart** field, change the status to **On** if you want the service to restart automatically whenever the server has stopped for any reason (including intentional shutdowns).

   Otherwise, leave it at its default value of **Off**.

10. In the **Number of Licenses** field, type the number of copies of this service you want to run.

    In most instances, you can use the default number of **1**, which is the minimum for any service.

11. From the window header, click on the **Register Service Information** button. The Asset Management window expands to show a list of gateway types for the database gateway service that you selected.

12. In the **Gateway Type** column, select a type of database to which you want the gateway to connect, such as DB2 for OS/2 Sessions.

    One database gateway service can give you access to several types of databases and applications. The only restriction is that a gateway session cannot connect to a database service running on UNIX and also connect to a database service running on OS/2.

13. With a gateway type selected, click on the **Show Details** button from the expanded portion of the window. The Asset Mgmt - Gateway Type Details window opens in the top left corner of the desktop, as shown in Figure 5.

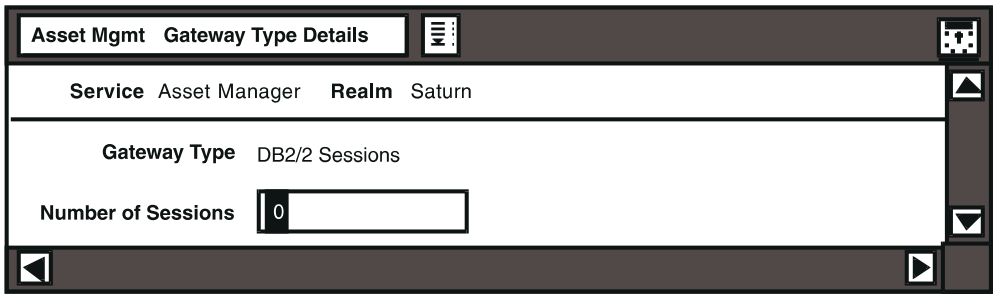

Figure 5. Asset Mgmt - Gateway Type Details window

14. In the Asset Mgmt - Gateway Type Details window, type the maximum number of simultaneous user-access sessions of the selected type to be allowed on the gateway.

15. In the window header, click on the down arrow button to display the pull-down menu.

16. Select **Register Gateway Information** from the menu.

17. Close the Asset Mgmt - Gateway Type details window.

In the expanded portion of the Asset Mgmt - Service Details window, the number of sessions (under the Number of Sessions column) for the gateway type you selected (such as DB2 for OS/2 Sessions) has changed to the number of sessions you specified in the Asset Mgmt - Gateway Type Details window.

18. If you want to allow access to another type of gateway accessible through the OS/2 database gateway service, repeat steps 13 through 18. Otherwise, close the Asset Mgmt - Service Details window.

    When you close the window, the Asset Mgmt - Node Details window reopens. The database gateway service you just added appears as an asset under the Asset column (shown in Figure 6), thereby enabling use of the service on that node.

    If you add services to a node on which a UNIX database gateway service is currently running, you must stop and restart the UNIX database gateway service.

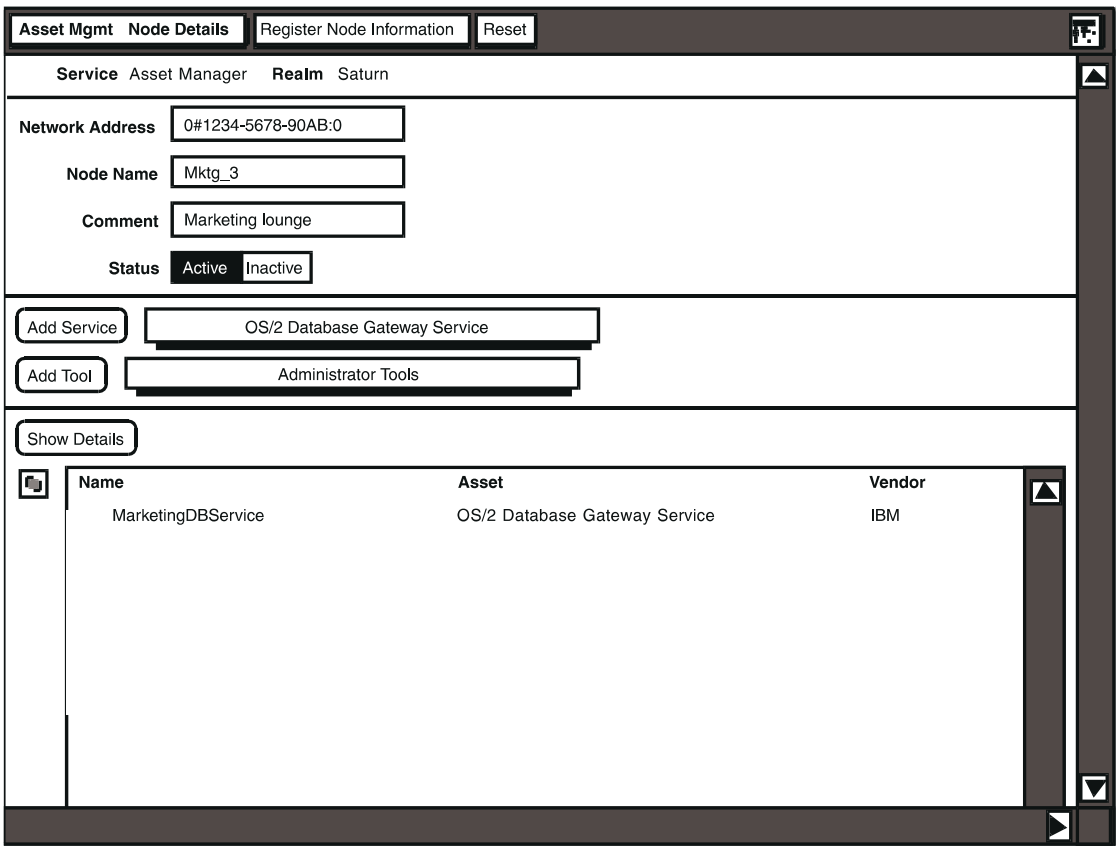

Figure 6. Asset Mgmt - Node Details window, displaying newly added database gateway service

Chapter 3. Configuring a Service

After an Intelligent Decision Server (IDS) database gateway service is assigned to your node, you can configure the gateways that are to run on your service.

The SA tool contains all of the IDS administrative functions, which are provided through a separate interface for each service. For a database gateway service, this interface is called the database gateway service interface. This chapter describes how to access this interface to configure the database gateways and their associated sessions.

The following topics are described in this chapter:

"Accessing the Database Gateway Service Interface"

"Allocating Licenses" on page 24

"Configuring a Database Gateway Service" on page 25

"Adding and Configuring a Gateway" on page 26

"Other Database-Access Tasks" on page 33

This book describes only the database gateway service interface. For instructions on general use of the SA tool, refer to the *System Administration Guide and Reference*.

Accessing the Database Gateway Service Interface

Because the SA tool contains all of the administrative interfaces, you must first access the SA tool, then select the database gateway services interface:

1. Open the SA Tool icon. The SA Tool window opens.
2. From the SA Tool window, select a realm, a service type (Database Gateway Services), and a service name.
3. To view the realm names that were detected by the asset manager, click on the **Realm** field and hold down mouse button 1. A list of realms opens.
4. To select a realm, drag the pointer to the realm of your choice and release the mouse button. The selected realm appears in the field.
5. To view the IDS service names that are registered with the asset manager residing on the realm you selected, click on the **Service Type** field and hold down button 1. A menu containing IDS services (such as the directory service, printing service, and mail service) opens.
6. Select a database gateway service by dragging the pointer to the service of your choice and releasing the mouse button.

   The chosen service appears in the field.
7. You can select the service name by either choosing it from a pull-down menu or by typing the name of the service in the service field.

© Copyright IBM Corp. 1996

To select the service name from a pull-down menu, first select **From List** in the **Select Service** field. Next, view the service names by clicking on the **Service** field and holding down the mouse button. Each database gateway service that was defined within the realm you previously specified is listed. Finally, select the appropriate entry from the Service menu.

To specify a service by entering its name yourself, first select **By Name** in the **Select Service** field. A field appears within the **Service** field. Then, type the name of the service in the **Service** field.

8. Click on the **Connect to Service** button. The Database Gateway Services password window opens.
9. Enter the password for the database gateway services interface. The Database Gateway Services window opens, as shown in Figure 7.

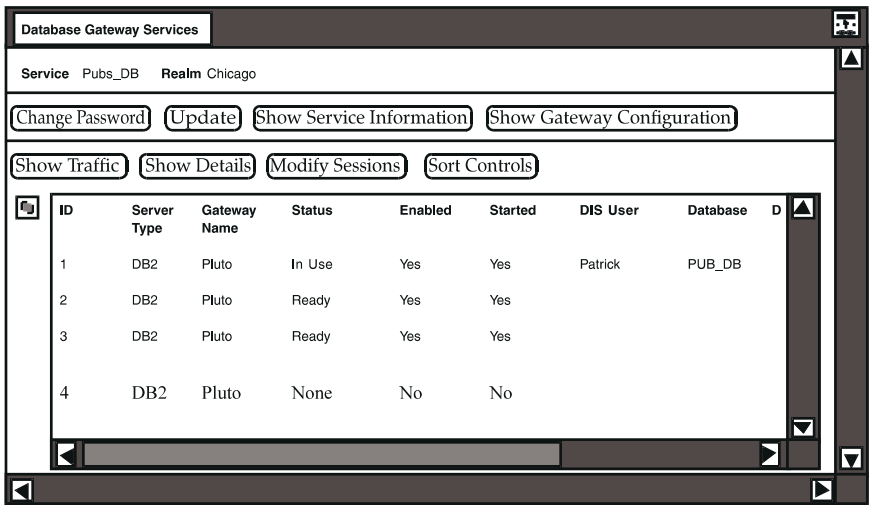

Figure 7. Database Gateway Services interface window

Allocating Licenses

All IDS database gateway service products are sold in packages called *assets*. Each asset provides connectivity to an RDBMS, from a specific operating system or hardware platform, with a set number of licenses. The number of licenses is equal to the number of sessions that are allowed to run concurrently on that particular service type.

You can distribute the gateways that run on a service and allocate the sessions that run from a gateway in any way you like. However, you can use only the number of licenses you purchased.

As part of the database gateway service configuration, you must specify the number of sessions at three levels within IDS:

The asset manager level (described in “Adding Services” on page 18).

The database gateway service level (described in “Configuring a Database Gateway Service”).

The gateway level (described in “Adding and Configuring a Gateway” on page 26).

Configuring a Database Gateway Service

Because some database gateway services run on OS/2 and require APPC connections to host workstations, configuring a gateway service can be a multi-step process. Before you configure the gateway service through the Database Gateway Service interface, you must know the following information:

The types of gateways that are registered with the asset manager.

The number of gateway sessions that are registered with the asset manager for the database gateway service.

The additional software required for each type of gateway that needs to be configured.

For example, if you are configuring a gateway that runs on OS/2 and requires APPC links to a host workstation, the IBM Communications Manager/2 must first be installed and configured. You must also know some of the parameter values that were used to configure Communications Manager/2. If you have this type of gateway and have not configured Communications Manager/2, see *Installing Host Services and Configuring Host-to-LAN Communications for MVS* .

Before you add and configure the individual gateways, specify the number of database and interface sessions for the database gateway service. This task can be performed from the DB Gateways - Service Info window, which contains general information about the service.

To configure gateway sessions for a service:

1. In the Database Gateway Services window, click on the **Show Service Information** button. The DB Gateways - Service Info window opens.

   Some fields in this window are informational, while others require that you supply or modify information. For descriptions of each field, see Table 2 on page 26.

2. Enter the appropriate information into each field requiring input.

   If you change the value in the **Maximum Database Sessions** field, the corresponding Sessions Available value does not change immediately, even if you close and reopen the DB Gateways - Service Info window. The Sessions Available value will update within a minute or so.

3. Ensure that **Auto Configuration** is set to **No**. Do not set **Auto Configuration** to **Yes** until after the configuration is complete.

4. Click on the **Set Configuration** button to save your parameters.

5. Close the window.

Table 2. Descriptions for fields in the DB Gateways - Service Info window

Field	Description
Software Version	This field indicates the version of database gateway service software currently running.
Service Up Since	This field shows the last time the service was started.
Maximum Database Sessions	In this field, specify the maximum number of database sessions the service will support. This is the number of licenses that are actually purchased. The maximum is 150.
Maximum interface Sessions	In this field, specify the maximum number of concurrent interface sessions allowed for this service. The default value is 1.
Sessions Available	This field indicates the total number of configured sessions that are not currently active.
Consoles Available	This field indicates the total number of interface sessions that are not currently active.
Auto Configuration	In this field, determine the configuration activity upon startup. If you select *Yes*, the service will be configured upon startup; if you select *No*, the service will not start up.

Adding and Configuring a Gateway

Before configuring a gateway, set the parameter values for the fields within the DB Gateways - Gateway Details window. The values you enter must be carefully coordinated with other components, including other IDS components, the target IBM Communication Manager/2, and the host software. To learn more about the process of planning connectivity for database access in IDS installations, see *Planning MVS Host-to-LAN Communications*.

Gateways that require a Systems Network Architecture (SNA) Advanced Program-to-Program Communications (APPC) connection to a host workstation have additional fields in the DB Gateways - Gateway Details window. The gateways that require APPC connections to a host system are:

DB2

MVS CAS

DB2 for OS/2 DRDA

To add and configure a gateway, click on the **Show Gateway Configuration** button in the Database Gateway Services window. The DB Gateways - Gateway Config window opens, as shown in Figure 8 on page 27.

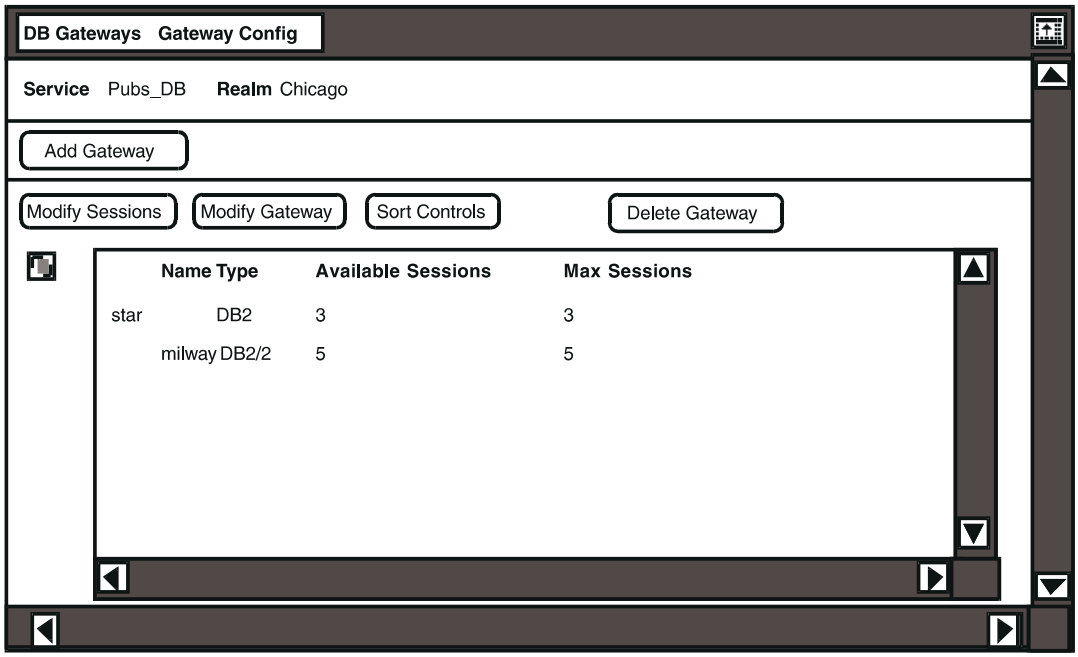

Figure 8. DB Gateways - Gateway Config window

Next, click on the **Add Gateway** button from the DB Gateways - Gateway Config window. A DB Gateways - Gateway Details window opens. The contents of this window vary depending on the type of gateway that was last selected. Table 3 describes the fields you need to fill in on the gateway details panels for each of the following gateway configurations.

Table 3 (Page 1 of 3). Gateway Details Fields

Field	Description
Gateway Name	Required for all gateway types. Specify the name of the gateway you are adding.
Gateway Type	Required for all gateway types. From a pull-down menu, select the type of gateway you are configuring. This menu contains a list of all the available gateway types for the selected service.
Pool	Displayed for all gateway types. If the service is participating in a load-balancing pool, enter the name of the pool. Each gateway can belong to only one pool.

Table 3 (Page 2 of 3). Gateway Details Fields

Field	Description
Macro Folder	Required for all gateway types. Enter the name of the folder containing the macros for this gateway type, such as DBS-REDBRICK or DBS-SYBASE. If this entry field is left blank, the SA tool provides a default name based on the gateway type. If you are configuring a DB2 for DRDA gateway, provide the name of the folder for the DRDA server type, such as DBS-DB2, DBS-OS400, or DBS-SQLDS. If you are configuring a DB2 for AIX gateway, provide the name of the folder for the gateway type, such as the DB2 for AIX PE. This field is not required for MVS-CAS.
Maximum Number of Sessions	Required for all gateway types. Specify the maximum number of concurrently active sessions this gateway supports.
Download Commit Count	Required for all gateway types. Enter a number indicating how many rows of data are to be downloaded to the database before the data is actually recorded. For example, if you enter *60* as the Download Commit Count, the database updates will be committed after every 60th row of data is downloaded. The number you specify depends upon the type of database you are using. For instance, you might download more records before a save to a DB2 database than to another type of database. This field does not apply to an MVS-CAS database gateway service. The value of *0* defers the commit until all of the data has been transferred.
Oracle Home Directory	Required only for an Oracle gateway. The directory specified in this field corresponds to the Oracle_HOME environment variable.
Default SID	Required only for an Oracle gateway. The SID specified in this field corresponds to the Oracle_SID environment variable. Together, the *Default SID* variable and the *Oracle Home Directory* variable define a default database for the Oracle gateway. If they are specified, then the **Database Name** field in the options sheet for data access tools can be left blank, and the default database will be accessed automatically.
DB2INSTANCE	Required only for DB2 for AIX or DataJoiner databases. Specifies the instance that is active by default. This field corresponds to the *DB2INSTANCE* environment variable. To use CAE to connect to a remote database, specify the CAE instance, not the remote instance.
Interfaces File Directory	Required only for SYBASE and Red Brick gateways. Enter the name of the directory that contains the interfaces file, which is accessed by the SYBASE Open Client interface. The directory specified in this field corresponds to the SYBASE environment variable.
Target Server Field	Required only for SYBASE and Red Brick gateways. Enter the name of the default SYBASE or Red Brick server. The server name entered in this field corresponds to the *DSQUERY* environment variable and must be present in the interfaces file.

Table 3 (Page 3 of 3). Gateway Details Fields

Field	Description
LU Alias	Required for DB2 and MVS CAS gateways. Enter the gateway's LU alias. This is the same value supplied earlier in the Communications Manager/2 configuration. For more information on the LU alias, see *Planning MVS Host-to-Lan Connectivity*.
Host Library	Specify the name of the host library associated with the Host Services software for the host; it is needed to start jobs.
Gateway ID	Required for DB2 and MVS CAS gateways. Enter the three-character identifier for the gateway. If you are using the IDS default values, the identifier will be something like *GWA* or *GWB*.
Mode Name	Required for DB2 and MVS CAS gateways. Enter the mode name defined earlier for this service. The mode name must match that specified in the Communications Manager/2 configuration.
Pre-Start	Required for DB2 and MVS CAS gateways. Pre-Start choices *On* and *Off* are used to indicate whether database access sessions are to be automatically started when the database gateway service is initialized. *On* indicates sessions are prestarted. *Off* indicates sessions start in response to a connect issued from a data access tool. The default for this field is *Off*. It is generally more efficient to prestart APPC connections.

DB2 for OS/2 and DB2 for OS/2 DRDA Gateways

When you select a DB2 for OS/2 or DB2 for OS/2 DRDA gateway, a Gateway Details window opens.

To add one of these gateways:

1. Type the required information in each field.

   See Table 3 on page 27 for descriptions of the fields within this window and the information they require.

2. Click on the **Add Gateway** button at the top of the DB Gateways - Gateway Details window to save your configuration information.
3. Repeat steps 1 and 2 for each gateway you create.
4. Close the DB Gateways - Gateway Details window.

To connect to a DB2 and SQL/DS server by means of a DB2 for OS/2 and Distributed Data Connectivity Services/2 (DDCS/2) connection, you must:

Set up a DB2 for OS/2 database.

Set up a User Profile Management (UPM) user ID that can be configured into the Icon Options windows of the various data access tools.

DDCS/2 and UPM require that an application (such as an IDS gateway session) first connect to a DB2 for OS/2 database before it can issue the CONNECT command for the DRDA server.

To connect to a DB2 for OS/2 database and support the UPM user ID, enter two-part values into the **Database Name** field, the **User Name** field, and the **Password** field, as follows:

local.remote

where:

local	is the DB2 for OS/2-specific value
remote	is the DRDA service-specific value

You can use a one-part user ID and password if your user ID and password match your host user ID. Otherwise, always specify both values.

To simplify UPM administration, you can create a single local user ID to use in all DDCS/2-connected data access icons. Create a local ID of two characters so that you'll have up to eight characters for the remote ID, plus one character for the period.

SYBASE and Red Brick Gateways

When you select a SYBASE or Red Brick gateway, the DB Gateways - Gateway Details window opens, as shown in Figure 9.

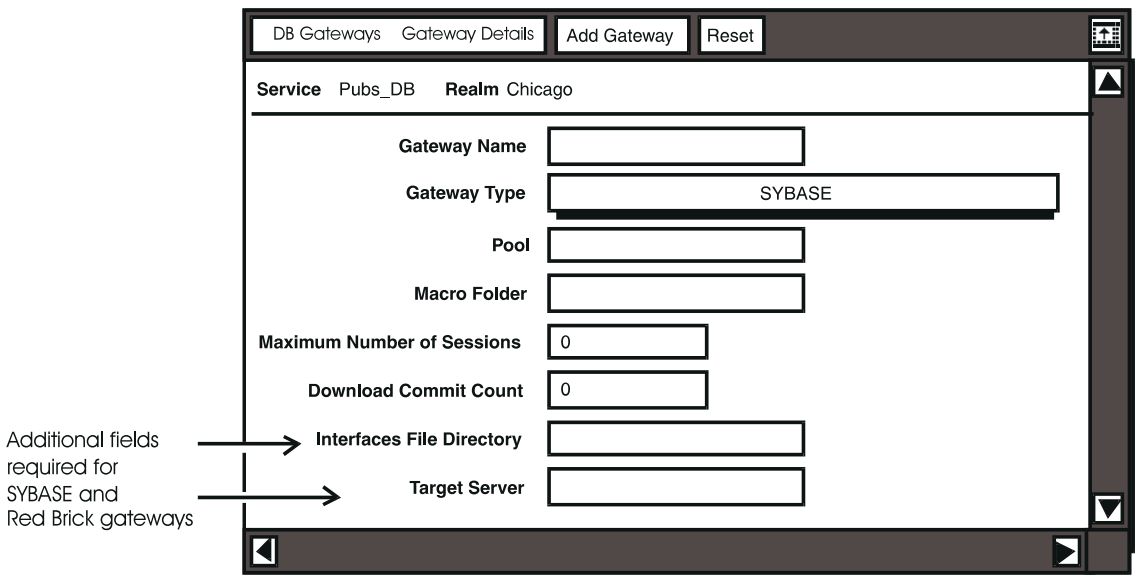

Figure 9. DB Gateways - Gateway Details window for the SYBASE and Red Brick gateways

There are two fields specific to SYBASE and Red Brick gateways, the **Interface File Directory** field and the **Target Server** field. When you configure these gateways, you

must provide the appropriate information in the additional fields. For example, you use the **Interface File Directory** field to identify the location of the SYBASE or Red Brick server you want to connect to.

To add this gateway:

1. Type the appropriate information into each field. See Table 3 on page 27 for descriptions of the fields within this window and the information they require.
2. Click on the **Add Gateway** button at the top of the DB Gateways - Gateway Details window to save your configuration information.
3. Repeat steps 1 and 2 for each gateway you create.
4. Close the DB Gateways - Gateway Details window.

Oracle Gateway

When you select an Oracle gateway, the DB Gateways - Gateway Details window opens, as shown in Figure 10.

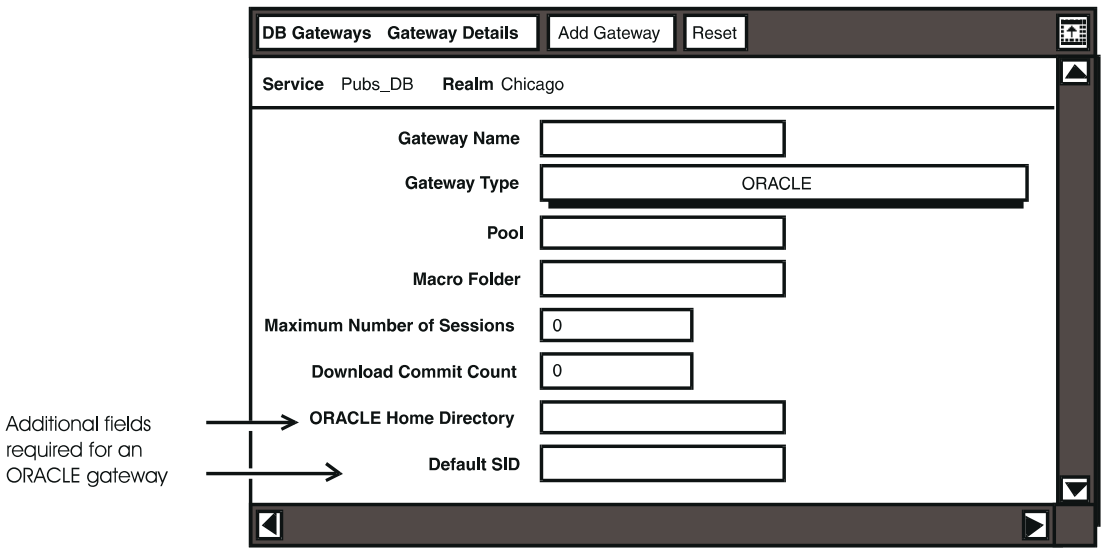

Figure 10. DB Gateways - Gateway Details window for the Oracle gateway

There are two fields specific to an Oracle gateway. When you configure this gateway, you must provide the appropriate information in the additional fields.

To add this gateway:

1. Type the appropriate information into each field. See Table 3 on page 27 for descriptions of the fields within this window and the information they require.
2. Click on the **Add Gateway** button at the top of the DB Gateways - Gateway Details window to save your configuration information.

3. Repeat steps 1 and 2 for each gateway you create.
4. Close the DB Gateways - Gateway Details window.

DB2 for AIX and DataJoiner Gateways

When you select a DB2 for AIX or a DataJoiner gateway, the DB Gateways - Gateway Details window opens, as shown in Figure 11.

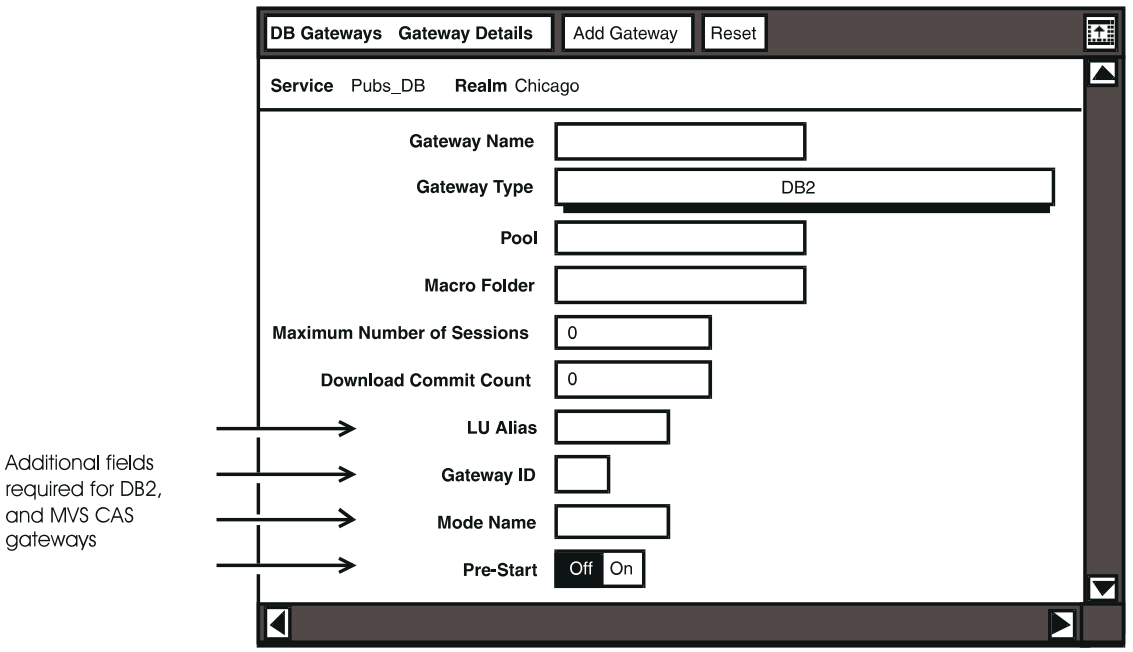

Figure 11. DB Gateways - Gateway Details window for the DB2 for AIX gateway

There is one field specific to a DB2 for AIX gateway. When you configure this gateway, you must provide the appropriate information in this field.

To add this gateway:

1. Type the appropriate information into each field. See Table 3 on page 27 for descriptions of the fields within this window and the information they require.
2. Click on the **Add Gateway** button at the top of the DB Gateways - Gateway Details window to save your configuration information.
3. Repeat steps 1 and 2 for each gateway you create.
4. Close the DB Gateways - Gateway Details window.

Other Database-Access Tasks

When a database gateway service and its gateways are configured, you must perform several other tasks related to setting up database access. This section provides only a brief overview of the tasks. They are described in detail in *Administering Databases with IDS*.

Configuring Data Access Tools for Database Access

For IDS users to run queries on databases and create reports through data access tools, the icons for the IDS data access tools must be configured to access specific databases. The data access tools include:

Query
Reporter
DBA
SQL Entry
Forms Design
SQL Input
Browser

The procedure for configuring the data access tools is consistent for all tools, as described in each tool guide. Also see *Administering Databases with IDS* for information on how to use the Reporter tool.

The data access tools must refer to the requested database gateway service by the gateway name, not the service name.

Locating the Database Folders

IDS provides several database folders, which pertain to gateway types. In some cases, there is more than one folder per gateway because different RDBMS versions can be accessed by a gateway. Each of the database folders contain macros that are called from the SA tool to create:

The Workstation Tools Data Dictionary (WTDD) and data tables
Customized macros that define a database to be used by IDS
User IDs for the database

These folders also contain the install macro and the DBA Tool Query file. The install macro and the DBA Tool Query file contain enhanced SQL functions that the DBA tool uses to perform database-specific operations.

To access the database gateway folders:

1. From the system desktop, select the System disk icon.
2. Open the File Drawers file drawer.
3. Open the DBA file drawer. A list of database gateway folders appears.

You will find the following macro folders in the DBA file drawer:

DBS-DB2

DBS-Oracle7

DBS-OS2

DBS-OS400

DBS-REDBRICK

DBS-SQLDS

DBS-SYBASE

DBS-Teradata

DBS-DB2 for AIX

DBS-DB2 for AIX PE

DBS-DataJoiner

Administering Databases with IDS describes the macros contained within each of these folders.

Chapter 4. Administering a Service

This chapter describes the processes available to system administrators for administering a database gateway service with the SA tool.

The administration interface in the SA tool enables you to perform routine administrative tasks, such as checking status, modifying the data access service and its gateways, and changing passwords. It also includes functions for monitoring network connections and isolating problems.

The following tasks are described in this chapter:

"Accessing the Gateway Administrative Functions"

"Administering the Service and the Gateway" on page 36

"Administering Gateways at the Gateway Level" on page 47

"Copying Status Information to Other Tools" on page 53

"Related Information" on page 53

Accessing the Gateway Administrative Functions

Access to all administrative functions for a database gateway service is provided with the SA tool. You perform all administrative tasks in the windows of the administration interface. For certain tasks, you will be required to supply information about the database or data access service.

To access the administrative functions for a database gateway service:

1. Open the SA Tool icon. The SA Tool window opens.
2. From the pull-down menu in the **Realm** field, select an IDS realm to connect to.
3. In the **Service Type** field, select **Database Gateway Services** from the pull-down menu.
4. From the pull-down menu in the **Service** field, select or type a service name to connect to.
5. Click on the **Connect to Service** button. The Database Gateway Services password window opens.
6. Type the password for the database gateway services interface in the entry field. If you type a valid password, the Database Gateway Services window opens.

All database gateway windows show the name of the service and the realm to which the service belongs.

Only started sessions are listed in the window. DB2 for OS/2 and UNIX gateways will not show any sessions if they have no open tools.

© Copyright IBM Corp. 1996

Administering the Service and the Gateway

You can perform administrative tasks for the database gateway service through the Database Gateway Services window. By using the buttons in this window, you can:

Change the service password
View the current status
Show service information
Show gateway configurations
Show traffic
Show session details
Modify gateway sessions
Sort controls

If you have a DB2 or MVS CAS gateway, you can also monitor traffic on the link between the gateway and host system.

The session status display also appears when the Database Gateway Services window is open. An example of a session status display window is shown in Figure 12.

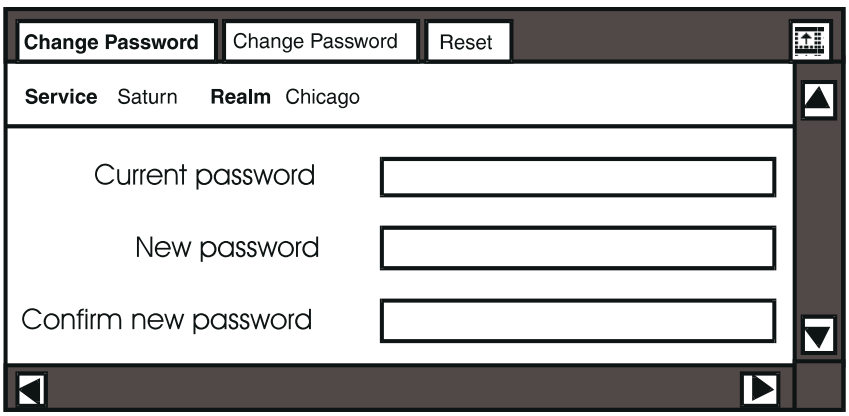

Figure 12. Example of a session status display window

The status display is used with the buttons to perform routine administration tasks and to modify the gateway service environment. Some of the buttons change the status of the service environment, while others produce windows containing additional functions. The remainder of this section describes typical administrative tasks that can be performed by using the buttons.

Changing the Service Password

All of the administrative functions in the Database Gateway Services interface are protected by the service password, which can be changed from within this interface.

To change the current password:

1. Click on the **Change Password** button in the Database Gateway Service window. The Change Password window opens.
2. In the **Current Password** entry field, type the current password that you want to change.
3. In the **New password** entry field, type a new password for the service.
4. In the **Confirm new password** entry field, type the new password again to verify that you entered it correctly the first time.
5. Click on the **Change Password** button in the window header so the new password can take effect.

After you perform this procedure, you will need to enter the new password to access the service.

If you decide not to continue with the password change after you have started, you can click on the **Reset** button in the window header to erase anything you entered, then close the window.

Updating Service Status

When a status display is produced, it reflects a snapshot view of the service at the time the display is requested. Status displays are not dynamically updated as the service runs. In windows where you can manipulate the database gateway environment or make changes that affect a status display, you need to refresh status displays with current information.

In the Database Gateway Services window, you can click on the **Update** button to refresh the status display with the latest activity occurring on the service.

At the service level, you can see:

The status of the service (such as Ready, None, or In Use).

If the service is enabled or disabled.

Whether the service has been started.

If the service is enabled and active, the name of the IDS user who is connected to the service.

If the service is enabled and active, the name of the database to which the gateway is connected.

If the service is enabled and active, the database user IDs of the people using the gateway.

If the logging feature is on or off.

The tracing level to which the service is set.

The appearance of the **Update** button does not change when you select it, but the change is reflected in the status display.

An **Update** button appears in the Database Gateway Services and the Traffic windows.

Displaying Service Information

The **Show Service Information** button in the Database Gateway Services window enables you to see the current configuration settings for the database gateway service. When you click on that button, the DB Gateways - Service Info window opens:

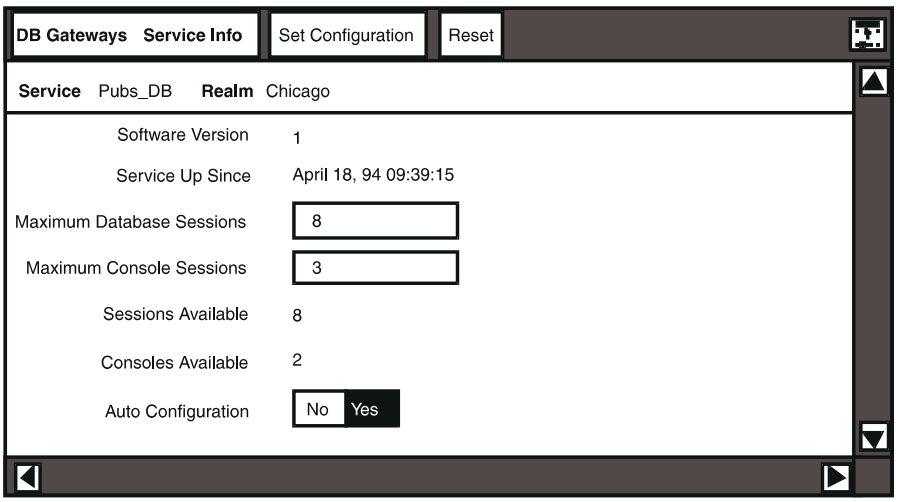

Figure 13. DB Gateways - Service Info window

From this window, you can see:

- The software version of the database gateway service
- The date and time the service was last started
- The maximum number of database sessions currently assigned
- The maximum number of interface sessions currently assigned
- The setting for the Auto Configuration function

If you change the value in the **Maximum Database Sessions** entry field, the corresponding Sessions Available value does not change immediately, even if you close and reopen the DB Gateways - Service Info window. The Sessions Available value is updated within a minute or so.

Displaying Gateway Configuration

You can see the types of gateways that are configured for your database gateway service, and other information, by clicking on the **Show Gateway Configuration** button from within the Database Gateway Services window. When you click on the button, the DB Gateways - Gateway Config window opens, as shown in Figure 14 on page 39.

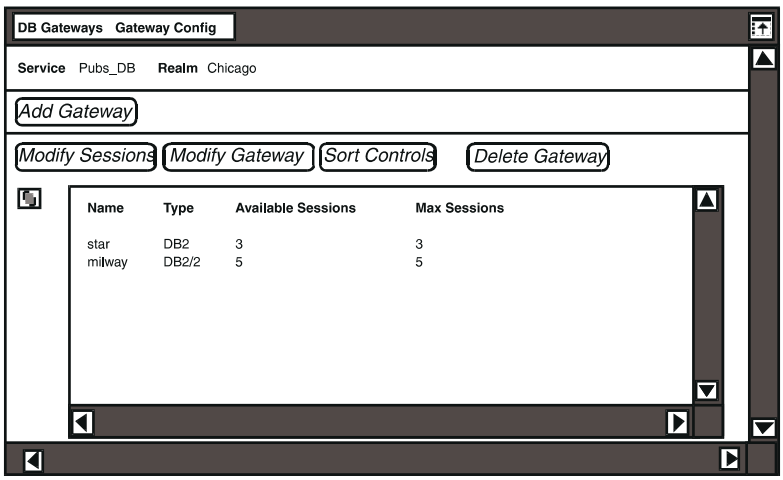

Figure 14. DB Gateways - Gateway Config window

Some of the buttons provided in this window are also contained in the Database Gateway Services window.

However, these buttons display different information, depending on the level (database gateway services or gateway) from which they are invoked.

If you want to add and configure a new gateway and establish sessions for that gateway, see Chapter 3, “Configuring a Service” on page 23.

For descriptions of the buttons within this window, see “Administering Gateways at the Gateway Level” on page 47.

Displaying Traffic Information

For gateways accessed through APPC connections, you can monitor packet traffic over the link. You should monitor traffic over the APPC links between the database gateway service and the host system when a problem is suspected. Indications that there might be a problem with the link include slow or no response from the database.

Only the gateways accessed through APPC connections display conversation status, which means information is displayed in the Database and DB User columns in the Database Gateway Services window.

To display traffic information for the APPC-configured gateways, first select a gateway for which you want traffic information. When the gateway is highlighted, click on the **Show Traffic** button from the Database Gateway Services window. The DB Gateways - Traffic window opens, as shown in Figure 15 on page 40.

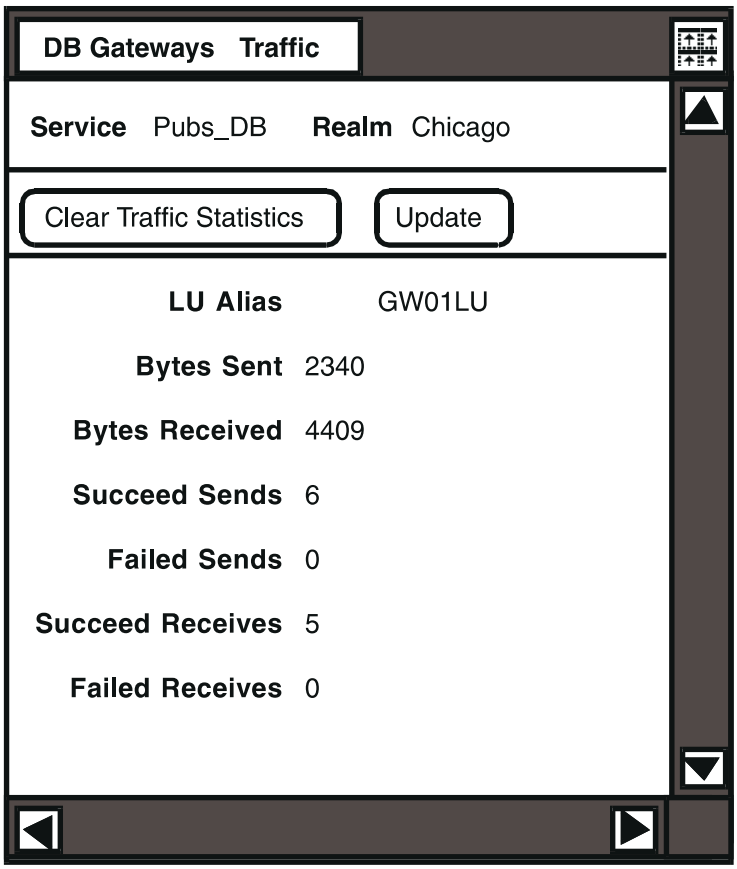

Figure 15. DB Gateways - Traffic window

The traffic statistics show link activity by counting packets that traverse the network through the database gateway service. This window displays information such as:

- The alias assigned to the local LU
- The number of bytes received and sent
- The number of packets sent and the number that were not sent
- The number of packets received and the number that were not received

If you suspect a problem with the APPC connection to the host system, you can reset traffic statistic counters at regular intervals and record the number of packets sent and received during each interval. If you are familiar with traffic statistics under normal operating conditions, you may be able to recognize a problem when you suspect poor performance.

To reset the statistics counters to zero, click on the **Clear Traffic Statistics** button. When you reset the statistics, the counts immediately return to zero.

To update the statistics counters with the latest information, click on the **Update** button.

Displaying Gateway Details

To display details about a gateway, select a gateway from the Database Gateway Services window. When the gateway of your choice is highlighted, click on the **Show Details** button. The DB Gateways - Gateway Details window opens.

This window displays the following information:

- Alias assigned to the local LU
- Alias assigned to the partner LU (PLU), the remote link
- Host system library associated with the database
- Name of the transaction program
- Date and time the transaction program was last started
- State of the conversation, such as *pending*
- Primary return code number
- Secondary return code number

For more information on the LU, see *Planning MVS Host-to-Lan Connectivity*.

Modifying Existing Gateway Sessions

The Database Gateway Services window provides a subset of functions that allow you to monitor system performance and help you locate problem areas related to the database gateway service.

To access these functions, first select a session from the Database Gateway Services window. When the session is highlighted, click on the **Modify Sessions** button. The DB Gateways - Sessions window opens.

This window contains settings for:

- Enabling and disabling sessions between the gateway and host systems
- Starting and stopping database gateway sessions
- Starting and stopping logging
- Setting trace levels

For UNIX, the DB Gateways - Sessions window settings take effect after the sessions are established.

The changes that were made in the status display that was last modified (such as the Database Gateway Services window) take priority over the settings in the DB Gateways - Gateway Config window

Disabling Sessions for a Gateway

For non-UNIX gateways, the Disable function prevents tools that are opening from using a particular function.

The Disable function is useful when you are concerned about a heavy load on the gateway service causing a degradation in performance. For example, if your database gateway service is configured to support more sessions than are currently active, but you notice that performance is not quite what you normally expect, you might want to disable unused sessions to prevent further degradation. When performance reaches normal expectations, you can use the Enable function to allow these sessions to be used as needed.

To disable gateway sessions, select **Disable** from the DB Gateways - Sessions window. This renders all of the sessions for the selected gateway inactive and prevents new sessions from starting. Sessions active at the time you select the **Disable** choice continue, until the tool using the session is closed. For UNIX gateways, the Disable function will prevent new sessions from starting.

Enabling a Session or Gateway

For non-UNIX gateways, select **Enable** from the DB Gateways - Sessions window to reverse the effect of Disable by placing the sessions in active status and allowing IDS tools to use them. The Enable choice also re-enables automatic startup of the host services link. In UNIX gateways, the Enable function allows you to enable and disable sessions.

Starting and Stopping Data Access Sessions

The **Start** and **Stop** fields located in the DB Gateways - Sessions window allow you to start or stop all sessions assigned to the selected gateway.

Using the Stop Function: Use the **Stop** function if you suspect the gateway service is malfunctioning or causing a performance problem with data access. This will immediately terminate an individual session or all sessions defined for a particular gateway.

To stop the sessions within the specified gateway, select **Stop** from the DB Gateways - Sessions window. This renders the sessions inactive and prevents new sessions from starting. Sessions that are active at the time you select the **Stop** button will stop immediately.

You can use the **Stop** function as part of an orderly shutdown procedure for the gateway. First, use the **Disable** function to prevent new sessions from starting. When there are no more active sessions, use the **Stop** function to terminate the sessions.

Using the Start Function: For non-UNIX gateways, the **Start** function initiates startup of all of the selected gateway sessions. This function is normally used after you have previously used the **Stop** function to terminate sessions with the RDBMS or CAS application.

To start the sessions within the specified gateway, select the **Start** choice. This starts the sessions for the selected gateway.

Recording Diagnostics Information in a Log File

To help you isolate problem areas, you can record diagnostic messages in a log file. This logging process uses the standard system administration log.

For OS/2, the name of the log file is:

system|File Drawers|Error Logs|service_name|DBServiceLog

where *service_name* is the name of the database gateway service on which you are running a log.

Details for using logs for problem diagnosis are contained in the *System Administration Guide and Reference*.

Establishing Gateway and Session Logging for UNIX Gateways

One of the advantages of the database gateway services is the ability to create logs at either the gateway or session level. Because this feature stores a large amount of information, it will affect your system's performance. Do not use this feature without recommendation from the IBM Support Center.

Before you use this logging feature, set the *M4SQLLOG* environment variable. First, you must issue the **startgateway** command before starting the service.

To set the *M4SQLLOG* variable, use one of the following commands, depending on the type of shell you are using:

If you are using *sh*, enter:

```
M4SQLLOG=1; export M4SQLLOG
```

If you are using *csh*, type:

```
setenv M4SQLLOG 1
```

You can also modify the startgateway script by inserting `M4SQLLOG=1; export M4SQLLOG` before the `echo ' '` line.

After you change the script, the *M4SQLLOG* variable will automatically be set whenever you issue the *startgateway* command.

After you set the variable, logging is turned on for each session. You can enable and disable the logging feature from the Database Gateway Services window.

All information will be written into the m4sql.log file.

If you do not want logging turned on for a particular session, you can disable the logging feature at the session level after the session is started by setting the logging feature to **Off**.

You can also disable all of the sessions for a particular gateway by setting the logging feature to **Off** at the gateway level. However, after logging is disabled at the gateway level, all subsequent sessions that are started from that gateway will still have logging turned on.

Running a Trace

If you call the IBM Support Center for assistance, you might be asked to run a trace on the gateway. Within the database gateway service interface is a trace function that can run traces on active sessions to help isolate suspected problems.

Do not run traces unless the IBM Support Center requests that you do so. Tracing some events may affect the performance of the gateway and increase network traffic.

The trace function records specified events in a circular buffer maintained by the Mettrace utility. It will run on all current sessions controlled by the gateway selected in the status display.

The tracing feature provides four levels of support:

Off	Indicates the trace is not active for the specified session
0	Records the gateway status and error messages
1	Records gateway status, error messages, and debug messages
2	Records buffer contents, repeated messages, and warnings

To initiate a trace, select one of the numbered choices in the **Trace** field. The status window will show that a trace is active on the gateway. To discontinue the collection of trace information, set the Trace function to Off.

The IBM Support Center may ask you to print the trace file, copy it to a diskette, or send it by electronic mail for diagnosis.

Sorting Functions

For your convenience, you can specify the order in which individual entries appear in their associated status displays. To do this, modify the parameters that appear in the Sort Controls window. These parameters do not automatically reset; they will remain as you modify them until you change them again. You can open this window, in slightly different forms, from either the Database Gateway Services window or the DB Gateways - Gateway Configuration window.

To access the Sort Controls window, click on the **Sort Controls** button from the Database Gateway Services window. The Sort Controls window opens as shown in Figure 16 on page 45.

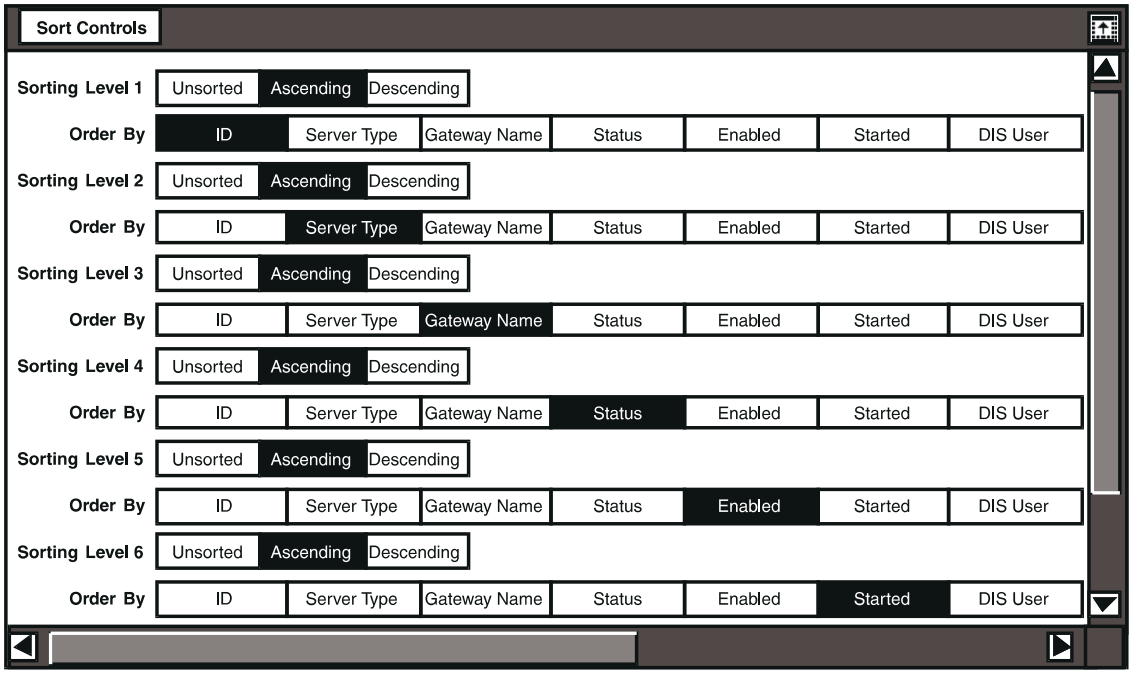

Figure 16. Sort Controls window

From this window, you can:

Choose one of three sorting methods (Unsorted, Ascending, and Descending).

Sort by one or more of the columns listed in the Order By row.

Specify the level of sorting.

The number of Sorting Level n and Order By rows in the Sort Controls window directly corresponds to the number of columns in the status display.

Unsorted Level

To make all status display columns appear in random order, select **Unsorted** at Sorting Level 1 of the Sort Controls window. If you select **Unsorted** in the Sorting Level 2 field, you specify that entries after the first sorting are to appear in random order. Subsequent sorting levels work similarly.

When you select **Unsorted** at any level, any lower-level Sorting Level and Order By rows immediately disappear from the Sort Controls window.

Ascending Level

To specify that the list of services is to appear in alphanumeric order, with numbers listed *0* through *9*, followed by letters listed *A* through *Z*, select Ascending. The sorting can be based on any column.

Descending Level

To specify that the list of services is to be sorted in alphanumeric order, with letters listed *Z* through *A*, followed by numbers listed *9* through *0*, select Descending. The sorting can be based on any column.

Selection List

Because there are 11 columns of information in the Database Gateway Services window, there are 11 Sorting Level n rows and 11 listings in the Order By selection list. For each column in the display, set the sorting order by selecting a choice in the Sorting Level n row. You can leave the field unsorted, sort it in ascending order, or sort it in descending order.

The following example shows the results of the sort control fields shown in Figure 16 on page 45.

	ID	Server Name	Gateway Name	Status	Enabled	Started	DIS User	Database	DB User	Logging	Tracing
1	DB2		Pluto	In Use	Yes	Yes	Patrick	PUBS_DB	Patrick	On	0
2	DB2		Pluto	Ready	Yes	Yes				On	0
3	DB2		Pluto	Ready	Yes	Yes				On	0
4	DB2		Pluto	None	No	No				On	0

The next example shows that the *ID* column is the first sorting level and is set to *Descending*:

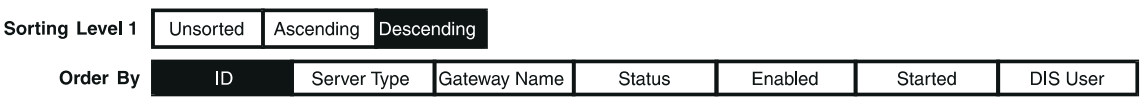

The results of this setting display all of the columns in the status display area in descending order, as shown in the following example:

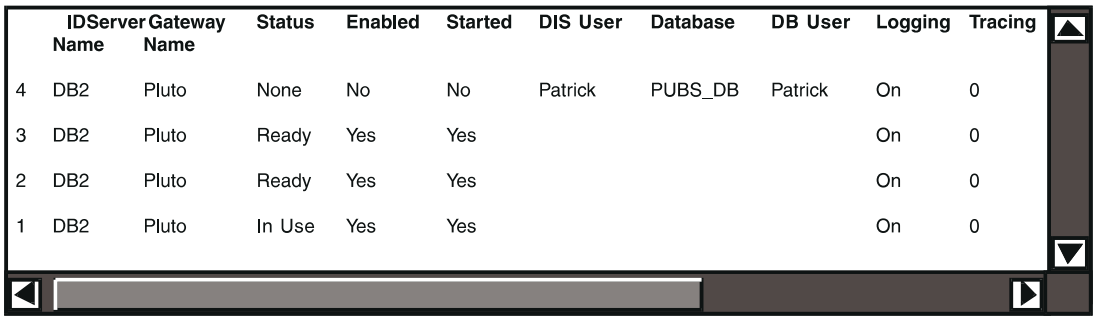

	IDServer Name	Gateway Name	Status	Enabled	Started	DIS User	Database	DB User	Logging	Tracing
4	DB2	Pluto	None	No	No	Patrick	PUBS_DB	Patrick	On	0
3	DB2	Pluto	Ready	Yes	Yes				On	0
2	DB2	Pluto	Ready	Yes	Yes				On	0
1	DB2	Pluto	In Use	Yes	Yes				On	0

Administering Gateways at the Gateway Level

At the gateway level, you can perform additional administrative tasks that pertain only to the gateway. To do this, click on the **Show Gateway Configuration** button from within the Database Gateway Services window. The DB Gateways - Gateway Config window opens.

The functions in this window allow you to:

Add and configure a new gateway.

Establish or modify sessions for a specified gateway.

Sort gateway controls.

Modify the configuration of an existing gateway.

Delete an existing gateway.

The functions that can be accessed only from the gateway level are described in the following sections.

Modifying All Sessions on a Specific Gateway

For non-UNIX gateways, the procedure for modifying all sessions on a specific gateway works identically from both the Database Gateway Services window and the DB Gateways - Gateway Config window. For UNIX gateways, the buttons have slightly different functions.

To modify sessions from the DB Gateways - Gateway Config window, first select a gateway from the status display window. Next, click on the **Modify Sessions** button. The DB Gateways - Sessions window opens, as shown in Figure 17 on page 48.

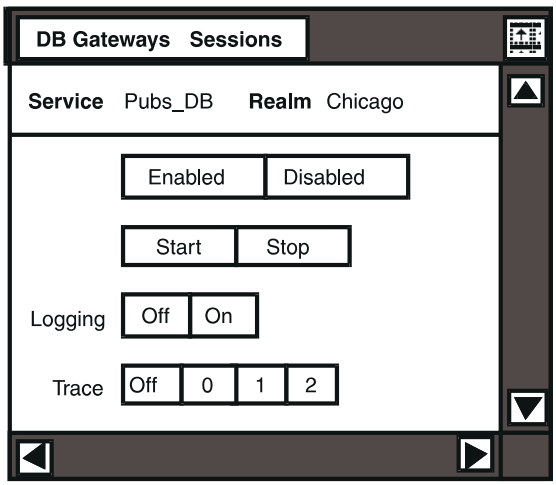

Figure 17. DB Gateways - Sessions window

The changes that were made in the status display that was last modified (such as the DB Gateways - Gateway Config window) take priority over the settings in the Database Gateway Services window.

For UNIX gateways:

Enable/Disable	The default is **Enabled**. This allows sessions to start. **Disabled** means that no new sessions can start. Existing ones are kept online.
Start/Stop	The **Start** function has no meaning. **Stop** means end all existing sessions.

For additional information about modifying existing gateway sessions, see “Modifying Existing Gateway Sessions” on page 41.

Sorting Controls

As with the status display entries in the Database Gateway Services window, you can specify the order in which individual entries appear in the DB Gateways - Gateway Config window by using the **Sort Controls** button.

To access the Sort Controls window, click on the **Sort Controls** button from the DB Gateways - Gateway Config window. The Sort Controls window opens, as shown in Figure 18 on page 49.

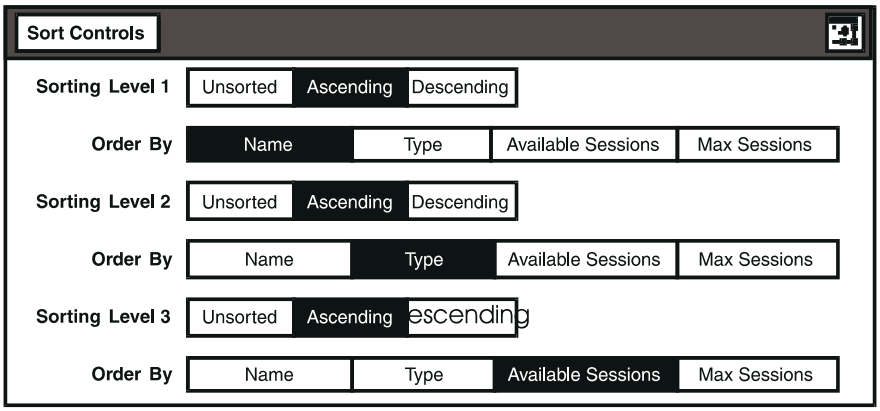

Figure 18. Sort Controls window accessible from the DB Gateways - Gateway Config window

The only difference between this window and the Sort Controls window that you can access from the Database Gateway Services window is in the number of columns in the Order By row. This window contains four columns:

Name
Type
Available Sessions
Max Sessions

You have three sorting methods to choose from: Unsorted, Ascending, and Descending. Also, you can sort on any one or more of the columns in the Order By row. Further, you can specify the levels of sorting, each based on one of the columns in the list.

The number of Sorting Level n and Order By rows in the Sort Controls window directly corresponds to the number of columns in the status display.

For additional information about using sort controls, see “Sorting Functions” on page 44.

Modifying the Configuration of an Existing Gateway

To modify the configuration of a gateway that has been configured in a service, first select a gateway from the status display list in the DB Gateways - Gateway Config window. Next, click on the **Modify Gateway** button. The DB Gateways - Gateway Details window opens. Fields in this window can vary depending on the type of gateway you select.

DB2 for OS/2 and DB2 for OS/2 DRDA Gateways

When you select a DB2 for OS/2 or DB2 for OS/2 DRDA gateway, a gateway details window opens.

To modify one of these gateways:

1. Enter the appropriate information into each field. See Table 3 on page 27 for descriptions of each field within this window.
2. Click on the **Set Configuration** button at the top of the DB Gateways - Gateway Details window to save the new information.
3. Repeat steps 1 and 2 for each gateway you modify.
4. Close the DB Gateways - Gateway Details window.

SYBASE and Red Brick Gateways

A gateway configuration can be modified only if it contains no active sessions and you disable it. Ensure that these two conditions are met before attempting to modify a gateway configuration.

When you select a SYBASE or Red Brick gateway, the DB Gateways - Gateway Details window opens.

There are two fields specific to a SYBASE or Red Brick gateway. You must provide the appropriate information in the additional fields when modifying these gateways.

To modify one of these gateways:

1. Type the appropriate information into each field. See Table 3 on page 27 for descriptions of each field within this window.
2. Click on the **Set Configuration** button at the top of the DB Gateways - Gateway Details window to save the new information.
3. Repeat steps 1 and 2 for each gateway you modify.
4. Close the DB Gateways - Gateway Details window.

Oracle Gateway

A gateway configuration can be modified only if it contains no active sessions and you disable it. Ensure that these two conditions are met before attempting to modify a gateway configuration.

When you select an Oracle gateway, the DB Gateways - Gateway Details window opens.

There are two fields specific to an Oracle gateway. When you configure this gateway, you must provide the appropriate information in the additional fields.

To modify this gateway:

1. Type the appropriate information into each field. See Table 3 on page 27 for descriptions of each field within this window.
2. Click on the **Set Configuration** button at the top of the DB Gateways - Gateway Details window to save the new information.
3. Repeat steps 1 and 2 for each gateway you modify.
4. Close the DB Gateways - Gateway Details window.

DB2 for AIX and DataJoiner Gateways

A gateway configuration can be modified only if it contains no active sessions and you disable it. Ensure that these two conditions are met before attempting to modify a gateway configuration.

When you select a DB2 for AIX or DataJoiner gateway, the DB Gateways - Gateway Details window opens.

There is one field specific to DB2 for AIX and DataJoiner gateways. When you configure one of these gateways, you must provide the appropriate information in this field.

To modify this gateway:

1. Type the appropriate information into each field. See Table 3 on page 27 for descriptions of each field within this window.
2. Click on the **Set Configuration** button at the top of the DB Gateways - Gateway Details window to save the new information.
3. Repeat steps 1 and 2 for each gateway you modify.
4. Close the DB Gateways - Gateway Details window.

Deleting a Gateway

To delete a configured gateway from this window, select it by clicking on its name.

Next, click on the **Delete Gateway** button. The name of the gateway is removed from the window. Deleting a gateway removes the gateway and its associated sessions from the service. If you attempt to delete a gateway that has active sessions, a warning message will appear, informing you that you must first stop the active sessions.

Performing Administration Tasks at the UNIX Interface

This section provides some administration tasks you must perform at the UNIX interface for the UNIX database gateway services, which include SYBASE, Oracle, Red Brick, and DB2 for AIX. Some troubleshooting information is also provided.

Starting a UNIX Gateway

After gateway software is installed on a node, you can use the **startgateway** script to start the gateway and verify the service connection. See *Installing Database Gateways for UNIX Environment* for more details.

To run the **script:**

1. Log on to the UNIX interface using the database gateway user name.
2. Append the following directory to the end of the *PATH:* line:

   /u/idsgw/bin

3. Change to the directory containing the files that you loaded from the installation media. For example:

cd /u/idsgw/bin

4. To start the gateway, enter:

   startgateway

If the database gateway is having a problem starting up, the workstation will either hang or display a series of error messages after the last line of the previous display.

If a problem occurs, use the **abortgateway** command to stop the process. After you fix the problem, you can reissue the **startgateway** command.

After the gateway completes its startup routine, it does not attempt to reconnect to the database server until some action on IDS makes a connection necessary. Such an action could be a user running a query against the data on the server. If you experience a connection failure, it is possible that subsequent connection attempts will succeed. This depends on the reason for the original failure and the status of the database server.

Using UNIX Gateway Administration Commands

You can issue the commands listed in Table 4 at the UNIX interface to administer a UNIX gateway.

Table 4. UNIX gateway administration commands

Command	Description
startgateway	This command starts the gateway. A banner appears on the interface if the command is successful. If the banner does not appear, or if an error message is received, an error has occurred.
stopgateway	This command stops the gateway.
showgateway	This command displays the status of the running gateway processes (similar to the UNIX *ps* command).
abortgateway	This command terminates the gateway in case certain errors leave it in an inconsistent state, where it cannot be stopped or started. For example, if you issue a stopgateway command and a subsequent startgateway command immediately fails with a system error, issue the abortgateway command. **Attention:** The abortgateway command removes any operating system message queues, semaphores, and shared memory owned by the user ID issuing the command.

Copying Status Information to Other Tools

Status information appears automatically when you open the Database Gateway Services window and the DB Gateways - Gateway Config window. Each of the windows that displays a status report contains a **Select All** button.

To copy status information to another tool, click on the **Select All** button to highlight the entries in the status display. Using either the Copy function key or the copy function in the pop-up function pad, copy the highlighted entries to an open spreadsheet or to a Text icon.

Related Information

For information about setting up data access, see *Administering Databases with IDS*. It contains information related to:

Creating WTDD tables
Database security
Other relevant data access information unique to the particular RDBMS
SQL reserved words

Glossary

address space. (1) The range of addresses available to a program. (2) A range of virtual storage addresses created and maintained by the MVS operating system for data and program residence, and for program execution; also, the program executing in an address space. MVS allocates real storage, processor, and input/output resources to an address space as needed to execute a program in that address space. MVS can maintain and control a large number of address spaces concurrently.

Advanced Program-to-Program Communications (APPC). An application programming interface that provides workstations with the ability to share and exchange programs, files, and other data with remote nodes.

AIX operating system. IBM's implementation of the UNIX operating system. The RISC System/6000 system, among others, runs on the AIX operating system. AIX is an acronym for Advanced Interactive Executive.

APPC. Advanced Program-to-Program Communications

application program interface (API). The interface through which an application program interacts with a service. An API usually consists of a set of functions or commands.

Application System (AS). (1) (*Cooperative Application Services*) An IBM decision support product used for producing graphics and reports, performing statistical analyses, and tabulating data. (2)See Application System/400 (AS/400).

AS/400. Application System/400

Application System/400 (AS/400). The IBM Application System/400 computer, which runs the OS/400 operating system. The AS/400 is classified as a midrange system.

asset management. Identification and distribution of purchased IDS features, services, and tools licensed for use.

capsule. A tool for building an automated processing sequence using other icons.

catalog. Tables maintained by the database administrator that contain descriptions of objects such as tables, views, and indexes.

client. A process, such as a program, that requests a service from another program.

communications controller. In Systems Network Architecture (SNA) communications, a subsystem or a front-end processor designed to manage data communications. Sometimes communications services are integrated with the host computer instead of being provided by a separate processor.

Communications Manager/2. An IBM product running on OS/2 systems that provides SNA connectivity to other SNA nodes, including host systems and other OS/2 systems.

conversation. In Systems Network Architecture (SNA) communications, a connection between two APPC transaction programs (TPs) over an LU to LU session that allows them to communicate with each other.

Cooperative Application Services. A combination of workstation software (a Cooperative Application Services transformer that runs within a capsule), a Cooperative Application Services gateway, and Host Services software running on an MVS host system. These three components allow IDS users to execute command scripts (such as those written in REXX) on the host and retrieve the results for further processing on their desktops.

database administrator (DBA). The person who issues administrative instructions to a database service. The DBA defines the data tables used by data access tools, monitors database performance, and assigns user access.

DB2. DATABASE 2, an IBM relational database management system that runs on an MVS system.

DB2 for OS/2. An IBM *relational database management system* running on OS/2 systems.

DBA. Database administrator.

Distributed Relational Database Architecture (DRDA). An architecture that provides connection between a database client on a workstation and a database server on a host system, which enables a database client to access and update information in a host database.

Ethernet. A type of local area network used for IDS connections. An Ethernet network uses a bus topology

© Copyright IBM Corp. 1996

(that is, all devices are connected to a single cable). An Ethernet network can include any of several kinds of physical transmission media.

file service. In IDS, a software process that manages disk storage of the system software and user data.

file server. The computer on which the file service runs.

front-end processor. In a computer network, a processor that relieves a host computer of processing tasks such as line control, message handling, code conversion, and error control.

gateway. A software component that establishes a communications path between a workstation and a relational database.

gateway service. A software process that translates protocols and access methods between two dissimilar systems.

gateway service identifier. A unique 3-character string, such as *GWA* or *GWB*, that identifies a gateway service as used in various configuration parameters.

gateway service name. The name by which a gateway is known to other components on the local area network (LAN).

Host Services. The IDS program product residing on a host computer. Host Services provides interface services between IDS on the LAN and the database on the host computer.

host system. A data-processing system with which other network devices can communicate, through one or more network connections.

IDS. The IBM Intelligent Decision Server.

internetwork. Two or more local area networks (LANs) interconnected for exchange of data.

LAN. Local area network.

link. In communications networks, a serial-by-bit connection under the control of Synchronous Data Link Control (SDLC).

load balancing. A feature that permits the service to distribute user sessions automatically across individual gateways grouped into a *pool*. With load balancing, the service assigns a user session to the least-loaded gateway in the pool, making maximum use of load capacity.

local area network (LAN). A data communications network, spanning a limited area, in which *serial transmission* is used for direct communication among computers or other devices.

local network. The network to which a specified user or resource is directly connected. See remote network. Contrast with network.

logical database. A database representation that includes print names for tables and columns, and that groups the tables into useful categories. The logical database also contains definitions for default joins between tables. The workstation tools data dictionary (WTDD) stores the description of the logical database, the mapping of print names to internal names in the physical database, and other logical attributes, such as the type and length of data.

logical unit (LU). In Systems Network Architecture (SNA) communications, a port through which a user or program accesses the SNA network for communication with another user or program. Contrast with physical unit.

logical unit 6.2. LU 6.2

LU. Logical unit.

LU 6.2. A type of communications protocol, which is used when connecting workstations to host systems.

macro. (1) On an IBM Intelligent Decision Server LAN, a text document containing parameterized SQL and administrative commands and stored in the file service. You can use these macros through either the host console or the database administration console in the SA tool on a workstation. (2) In VTAM and NCP, a statement that defines various network resources and their attributes.

Multiple Virtual Storage/Enterprise Systems Architecture. MVS/ESA.

MVS/ESA. An IBM MVS operating system environment that supports IBM S/390, ES/9000. MVS/ESA allows an address space access to one or more data spaces. A data space is, like an address space, a virtual storage area. While an address space contains system programs, user programs, and data, a data space can contain only data. Each data space can provide up to two gigabytes of virtual storage for data. Data spaces

can save time, prevent unintentional changes to data, and permit the sharing of data among selected users. So, in addition to the two gigabytes a job has available to it for holding instructions and data in an address space, each job can also have access to many data spaces.

NetBIOS. An application program interface (API) between a LAN adapter and programs. Nodes are identified on the LAN by NetBIOS names. In the OS/2 system, NetBIOS is part of LAN Adapter and Protocol Support.

network. A set of workstations, servers, and other devices connected together for the interchange of information. Contrast with local network.

OS/400. Operating System/400. The operating system for the IBM AS/400 computer.

OS/2. Operating System/2. A single-user, multitasking, IBM operating system that runs on microcomputers with an Intel 80386 or later microprocessor.

packet. In communications networks, a sequence of binary digits, including data and control signals, that is transmitted and switched as a unit.

parameter. Data passed by a user to a program to direct or limit processing in some way.

partner logical unit (PLU). In Systems Network Architecture (SNA) communications, the remote, corresponding participant in a session in an SNA network.

physical unit (PU). In Systems Network Architecture (SNA), a network-addressable unit that manages and monitors the resources of a node. Contrast with logical unit.

PLU. Partner logical unit.

pool. A collection of similar gateways or services that access a single database or host. Based on current loads on the individual gateway service, the service assigns database-access requests to pool members dynamically.

protocol. A set of rules and conventions that governs how two entities exchange data.

PU. Physical unit.

RDBMS. Relational database management system.

realm. In IDS networks, a collection of IDS resources that are maintained as a logical group (that is, a group that is determined not by physical boundaries, but by virtue of being associated with one another). A realm must contain at least a primary file service and a directory service. Other components such as workstations, internetwork routing service, and data access services can be added to form a functional IDS local area network (LAN).

relational data. In a relational database, information that is related to other information and whose relationships can be defined in the form of tables consisting of rows and columns of data.

relational database. A database that is organized and accessed according to relationships between data items. These relationships are expressed as tables related by data values, rather than by pointers or memory locations.

relational database management system (RDBMS). The software that makes it possible to create, use, and maintain relational data in the form of tables consisting of rows and columns of data.

remote network. A network other than the local network, which is accessible to the local network by some means. Contrast with local network.

RISC System/6000. The IBM RISC System/6000 computer. This computer uses the AIX operating system. See also AIX operating system.

serial transmission. The sequential transmission of one data bit at a time.

server. The processor in which a service runs.

service. A software process that provides facilities to other parts of IDS, such as a file service, a mail service, a mail gateway service, and a database gateway service.

service manager. An MVS Host Services address space that coordinates one or more sets of subsystem services.

session. In the OS/2 environment, one instance of a running program or command prompt. IDS runs in an OS/2 full-screen session.

SNA. Systems Network Architecture.

SQL. Structured Query Language

SQL/DS. The IBM Structured Query Language/Data System relational database management system that runs on a VM system.

Structured Query Language. Structured Query Language, a command language for use with relational databases. The language consists of statements to insert, update, delete, query, and protect data. IDS tools use SQL to communicate with database servers.

Synchronous Data Link Control (SDLC). A discipline conforming to subsets of the Advanced Data Communication Control Procedures (ADCCP) of the American National Standards Institute (ANSI) and High-level Data Link Control (HDLC) of the International Organization for Standardization, for managing synchronous, code-transparent, serial-by-bit information transfer over a link connection. Transmission exchanges may be duplex or half-duplex over switched or non-switched links. The configuration of the link connection may be point-to-point, multipoint, or loop.

Systems Network Architecture (SNA). The IBM description of logical structure, formats, protocols, and operational sequences for network configuration, operation, and transmission.

table. In a database, a single file consisting of data organized in columns and rows.

TCP/IP. Transmission Control Protocol/Internet Protocol.

token-ring network. A network that uses a ring topology, in which tokens are passed in a circuit from node to node. A node that is ready to send can capture the token and insert data for transmission.

tool. A IDS workstation program, such as Text, Reporter, or Query. Every tool is represented by an icon.

TP. Transaction program.

transaction program (TP). A program that uses an APPC application programming interface to communicate with a partner program on a remote node.

Transmission Control Protocol/Internet Protocol (TCP/IP). A transport protocol that packages information for transmission, reassembles received packets, and handles the routing and transmission of data between networks.

user ID. A name used to access a desktop, system, gateway, or database service.

view. A copy of a query, which can be named, stored, edited and redisplayed for later use.

Virtual Telecommunications Access Method (VTAM). A set of programs that maintain control of the communication between terminals and application programs running under DOS/VS, VM, and OS/VS2 operating systems.

VM (Virtual Machine). (1) A functional simulation of a computer and its associated devices. Each virtual machine is controlled by a suitable operating system. (2) In VM, a functional equivalent of either a System/370 computing system or a System/370-Extended Architecture computing system. Each virtual machine is controlled by an operating system. VM controls concurrent execution of multiple virtual machines on a single system.

VM/ESA. Virtual Machine/Enterprise Systems Architecture operating system.

VTAM. See Virtual Telecommunications Access Method.

workstation. A computer at which the user can use IDS to perform work. Workstation components include the monitor, keyboard, mouse or trackball, and system unit, which may house a disk drive.

Workstation Tools Data Dictionary (WTDD). See WTDD.

WTDD. A logical data dictionary that standardizes the IDS workstation view of data across many relational database management systems. The WTDD stores the print names of database items, defines how the items are to be displayed, and includes other logical attributes of data. The WTDD is implemented as a set of standard relational tables that can be managed through the DBA tool.

IDS Publications

This section lists IDS publications, including order numbers. To order copies of the books listed here, or to get more information about a book, see your IBM representative.

General Publications

Getting Started with the Intelligent Decision Server (SC26-8892)

Working with the Intelligent Decision Server (SC26-8893)

Installation, Configuration, and Administration Publications

Installing IDS LAN Components (GC26-8877)

System Administration Guide and Reference (SC26-8890)

Messages and Codes (GC26-8878)

Database Gateway Services Publications

Installing Database Gateway Services for UNIX Environments (GC26-8884)

Configuring and Administering Database Gateway Services (GC26-8885)

Administering Databases with IDS (SC26-8891)

Host Services Publications

Planning MVS Host-to-LAN Connectivity (GC26-8880)

Installing Host Services and Configuring Host-to-LAN Communications for MVS (GC26-8881)

Administering Host Services for MVS: DB2 and Cooperative Application Services (GC26-8882)

Developing Applications with Cooperative Application Services (GC26-8883)

Configuring and Administering Mail Gateway Services (GC26-8923)

Open Data Access Service Publications

Developing Applications with Open Data Access Service (SC26-8888)

Data Access Tools Publications

Browser User's Guide (SC26-8900)

Data Entry User's Guide (SC26-8897)

DBA Tool User's Guide (SC26-8894)

SQL Entry User's Guide (SC26-8898)

Data Analysis Tools Publications

BASIC Tool User's Guide (SC26-8908)

Query User's Guide (SC26-8901)

Spreadsheet User's Guide (SC26-8904)

Presentation Tools Publications

Layout User's Guide (SC26-8902)

Plot User's Guide (SC26-8905)

Reporter User's Guide (SC26-8903)

Text User's Guide (SC26-8896)

Application Development Tools Publications

Transformers Guide (SC26-8907)

Capsule User's Guide (SC26-8906)

Forms Design User's Guide (SC26-8899)

Interacting with PC Applications (SC26-8895)

Installing and Configuring IDS Open Clients (GC26-8889)

© Copyright IBM Corp. 1996

Index

© Copyright IBM Corp. 1996

DSMSCA302W INVALID COMMAND OPTION - 'NOBOX'.
DSMSCA302W INVALID COMMAND OPTION - 'SPINE'.
DSMBEG323I STARTING PASS 2 OF 4.
DSMBEG323I STARTING PASS 3 OF 4.
DSMBEG323I STARTING PASS 4 OF 4.