

QVS Distributed Data Services/Controller Services Facility for Windows

Installation and Configuration Guide Version 3

First edition (June 2004)

This edition applies to Version 3 of the QVS Distributed Data Services/Controller Services Facility for Windows, and to all subsequent releases and modifications until otherwise indicated in new editions.

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

© Copyright QVS Corp. 2004

Updated: 8/15/07

Preface.....	6
How to Use this Manual	6
Chapter 1. System Overview	8
Nodes	9
Node IDs	9
System ID.....	10
Logical Names and Role Names.....	10
Reserved Role Names	11
Broadcast Domains	11
Distribution Roles and Domains	12
Primary Distributor	12
Backup Distributor.....	12
Subordinate	12
Mirrored domain	12
Broadcast domain.....	12
File Names and Queue Names.....	13
Components	14
Name Services	14
File Services.....	14
File System Interface	14
Interprocess Communications (IPC).....	15
Data Distribution.....	15
Node Control.....	15
Problem Determination and Analysis	16
4690 Controller Services Facility (CSF).....	16
File Distribution	16
Disk I/O Prioritization.....	16
Chapter 2. Planning for DDS/CSF Installation.....	18
System Topologies.....	18
Stand-Alone Node.....	19
Multiple Nodes on a Single LAN	19
Multiple Nodes on Multiple LANs.....	19
Multiple Servers on Multiple LANs	20
Hardware Requirements.....	21
Software Requirements.....	22
Preparing Your Workstation.....	22
Partitioning Your Hard Disk.....	23
Designating Controlled Drives	23
Controlled Drives and Disk I/O Prioritization	24
Controlled Drives and Data Distribution	25
Controlled Drive Restrictions	26
Example for Partitioning Hard Disks.....	27
Planning for TCP/IP Configuration	27
Name Resolution.....	29

Dynamic Host Control Protocol.....	30
Blocking Interfaces	30
Automatic Primary Distributor Switch-Over.....	30
IP Address Computer Name Switch-Over.....	32
Chapter 3. Installation.....	33
Choosing Attended or Unattended Installation and Configuration.....	33
Performing an Attended Installation of DDS or DDS/CSF	34
Performing an Unattended Installation	39
Transferring DDS Installation Files to a Code Server	39
Creating Response Files.....	39
Syntax Rules for Response Files.....	40
Automating Node ID Specification	41
Configuration Keywords.....	44
Starting Unattended Installation with SETUP	48
Example of Starting Unattended Installation Using SETUP	50
Upgrading an Installed Version of DDS.....	51
Resetting Controller Services Facility (CSF) Parameters.....	51
Troubleshooting Installation Errors	51
Chapter 4. Uninstalling DDS or DDS/CSF.....	53
Performing an Attended Uninstall	53
Unattended Uninstall	53
Troubleshooting Uninstall Errors	54
Installing and Removing Corrective Service	54
Chapter 5. Configuration Data.....	55
Managing Configuration Data	55
Changing Configuration Data	56
Activating Configuration Data.....	56
Chapter 6. Configuration File Utility.....	58
Starting the Configuration File Utility.....	58
Generating Configuration Files.....	60
Loading a Configuration File.....	60
Generating a New Configuration File.....	60
Working with Advanced Configuration Data	63
Generating Logical Names Files.....	66
Generating a New Logical Names File.....	66
Loading a Logical Names File.....	68
Adding Logical Names	68
Editing Logical Names	69
Removing Logical Names.....	70
Chapter 7. Configuration Keywords.....	71
Basic Configuration Keywords.....	72
AutoSwitchOver	72
AutoSwitchOverDelay	73
AutoSwitchOverForce	74
ControlledDrives.....	75
DDActive	76

DistributionRole.....	77
IPCTransport.....	77
MaxRequestors	78
PrimaryComputerName	78
PrimaryIPAddress	79
PrimaryIPAdapter	80
PrimarySubnetMask.....	81
RemoteIPC	82
WorkDirectory	82
Advanced Configuration Keywords	84
AdptrResetValue.....	84
DDSBlockedInterface	84
DistRenamedFile.....	85
IPCHeartbeatInterval	86
IPCPortCount.....	86
IPCPortStart	86
IPCTimeout.....	87
LocatePrimary.....	87
MaximumMemory	88
NetworkRequestInterval	88
NetworkRequestRetries	89
Response-File-Only Configuration Keywords	89
BackupConfig	89
ChangeLogicalName.....	90
ConfigUpdateType.....	91
DeleteLogicalName	91
Include.....	92
LogicalNamesUpdateType	93
Reboot.....	93
RestartFDS.....	94
RestoreConfig	94
Logical-Names-File-Only Configuration Keywords	95
LogicalName.....	95
Configuration Keywords Summary	96
Data Distribution Configuration Keywords.....	98
Interprocess Communication Keywords.....	98
Chapter 8. Unattended Configuration	100
Creating a Response File for Unattended Operation	100
Creating Response Files for an Unattended Initial Installation	100
Creating Response Files for an Unattended Configuration Change	101
Configuration Response File Processor (FDSCRFP)	102
Syntax for FDSCRFP.....	102
Using FDSCRFP to Generate Configuration Files and Logical Names Files	105
Chapter 9. DDS/CSF Installation and Configuration	106
Installation.....	106
Checklist of Steps: One or Two Controllers	106

Installing Windows NT, 2000, XP, or Server 2003 on Your Controller	108
Installing a Telnet Server for 4690 OS Auxiliary Consoles	109
Store Loop Adapter Installation.....	109
PCI Token-Ring Adapter Setup	110
Installing DDS and CSF.....	110
Copying Files from 4690 to the Windows Machine via FTP	114
Set the Distribution Attributes of the 4690 Files	115
Starting CSF.....	117
CSF Initialization Error Messages	118
Configuration and Customization	118
4690 CSF Configuration	119
Automatic Start and Restart of 4690 CSF.....	119
Automatic Pre-Processing, Post-Processing, Custom Processing	119
Configuring a Printer	123
Appendix A. Controller Services Facility Directory Structure.....	124

Preface

This manual describes how to install and configure IBM Distributed Data Services/Controller Services Facility (DDS) for Windows NT ®, 2000, XP, or Server 2003 operating systems. All input needed for installation and configuration is provided either through a graphical user interface or from response files. **Who Should Read This Manual** This manual is primarily for retail systems programmers who are installing and configuring the QVS Distributed Data Services/Controller Services Facility on the Windows NT, 2000, XP, or Server 2003 operating system. This manual assumes that the reader is familiar with Windows operating systems.

How to Use this Manual

Read Chapters 1, 2, and 3 before you install DDS. This book is separated into ten chapters and two appendixes:

Chapter 1. System Overview provides an overview of DDS and describes the main concepts you need to know in order to use DDS effectively.

Chapter 2. Planning for Distributed Data Services/Controller Services Facility Installation describes the available system topologies and the hardware and software requirements. This chapter also provides instructions for preparing your workstation.

Chapter 3. Installation explains how to install DDS or DDS/CSF using either attended installation or unattended installation procedures.

Chapter 4. Uninstalling DDS provides instructions for deleting DDS from your system. [Installing and Removing Corrective Service](#) provides instructions for installing and deleting corrective service.

Chapter 5. Configuration Data describes configuration data and some methods for managing it.

Chapter 6. Configuration File Utility provides instructions for configuring DDS using graphical user interface panels.

Chapter 7. Configuration Keywords provides descriptions of the keywords that are used for creating response files and configuration files.

Chapter 8. Unattended Configuration describes unattended configuration and provides instructions for using the Configuration Response File processor to create configuration files and logical names files.

Chapter 9. DDS/CSF Installation and Configuration describes the procedures for installing and using the 4690 Controller Services Facility.

Appendix A. Controller Services Facility Directory Structure describes the directory structure used by DDS.

Appendix B. Reading the Online Documentation explains how to view online the **QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide** and the **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide**. © Copyright QVS Corp. 2004

Other related publications:

- **Controller Services Facility Version 3 Related Publications** **QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide**, SC30-4041 This manual is shipped as an online document when you purchase DDS. **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide**, GC30-4042 This manual is shipped as an online document when you purchase DDS.
- **IBM VisualAge ® C++ for Windows Library** IBM VisualAge C++ for Windows, Version 2.5 - Standard Manuals This collection of manuals includes: User's Guide, Programming Guide, Open Class Library User's Guide, and Visual Builder User's Guide. IBM VisualAge C++ for Windows, Version 2.5 - Reference Manuals This collection of manuals that includes: Visual Builder Parts Reference, Building VisualAge C++ for Fun and Profit, Language Reference, C Library Reference, Open Class Library Reference Volume I, II, III, and IV, SOM Programming Guide, and SOM Programming Reference.
- **C-Related Publications** Portability Guide for IBM C x Part 1. Introduction

Chapter 1. System Overview

The QVS Distributed Data Services/Controller Services Facility for Windows NT, 2000, XP, or Server 2003 is a distributed software platform designed as a base for the development of distributed applications for the store environment.

In this manual references to the product are as follows:

- DDS or Distributed Data Services is used when the information is specific to DDS functions.
- CSF or Controller Services Facility is used when the information is specific to CSF functions.
- DDS/CSF is used when the information applies to both DDS and CSF functions.
- Windows refers to Windows NT, 2000, XP, and Server 2003 unless specified otherwise.

The main functions are:

- Interprocess communication with local and remote transparency
- Data access with local and remote transparency
- Data distribution for redundancy, performance, and availability
- Directory services
- Installation, configuration, and administration

The following sections provide an overview of the concepts that are common to all components of DDS:

- “Nodes” on page 9
- “Logical Names and Role Names” on page 10
- “Broadcast Domains” on page 11
- “Distribution Roles and Domains” on page 12
- “File Names and Queue Names” on page 13

- ““Components” on page 14
- “File Distribution” on page 16
- “Disk I/O Prioritization” on page 16

Nodes

A node is a LAN-attached machine that is running DDS. Nodes can be connected via LANs to form a system. A system is the group of nodes for which files are managed. A node can be connected to a subset of nodes via one LAN and to other subsets of nodes via different LANs. Data distribution, remote file access, and interprocess communication are supported only among nodes that are connected by a LAN. DDS does not provide a bridge or router function.

For all system topologies with more than one node, one node must be installed and configured as the configured primary distributor and another node can be installed and configured as the configured backup distributor. The backup distributor is required to add redundancy to the critical store data. Each node must be connected to the primary distributor and backup distributor by a LAN. See “Distribution Roles and Domains” on page 12 for an explanation of primary distributor, backup distributor, and subordinate.

Node IDs

Node IDs are specified for each workstation (node) when DDS is installed. Each node has a single node ID, regardless of how many LANs are used to connect it to other nodes.

The rules for node IDs are:

- Each workstation within a system must have a unique node ID. Duplicate node IDs are not detected during installation, but they are detected when DDS initializes. You must reinstall DDS to correct the problem.
- No two workstations on a LAN can have the same node ID even if they are members of different DDS systems.
- Node IDs can be from 1 to 8 characters (blanks are not allowed). Node IDs are case-sensitive.
- Do not use node IDs whose first 3 characters are "FDS". This prefix is reserved.

- Do not use node IDs whose first character is a / (forward slash).
- Greater than and less than characters (< and >), question marks (?), asterisks (*), and colons (:) are not valid characters for node IDs.
- After installation, node IDs can be changed only by reinstalling DDS.

System ID

Each system has a system ID of 1 to 4 bytes. The system ID is specified at each node during the installation of DDS, and defaults to 0000. The Name Services component uses the system ID as a qualifier when locating the node that has a particular role (see Logical Names and Role Names for a description of roles). This method ensures, for example, that each node finds the acting primary distributor for its system, and that files are distributed only within a system. The system ID must be set to a unique value for each system in environments where multiple systems are interconnected via bridges. A system ID cannot be changed without reinstalling DDS.

If you assign a different system ID to each system, the systems can be connected using a LAN or a gateway without causing messages intended for one system to be sent to another system.

If you assign the same system ID to each system, the LANs can be connected using a gateway; both LANs are considered the same system. If this method is used and you are using the Data Distribution component, there should only be one configured primary distributor and one configured backup distributor for the entire system, even though they might be on different LANs.

Note: A single asterisk (*) is not a valid system ID. Alphabetic characters in the system ID are case-sensitive.

Logical Names and Role Names

DDS provides a name-resolution capability, allowing applications to use logical names instead of hard-coded file names, interprocess communication (IPC) queue names, and node IDs. These logical names are dynamically resolved when the application runs.

The use of a logical name allows the file name to be changed without having to rebuild the application. A logical name has the following format: <name> where name is 1 to 260 characters, and the less than and greater than characters (< and >) are required delimiters.

Some names are very dynamic, such as the node ID of the primary distributor. This name changes whenever the backup distributor takes over for the primary distributor. In this case a role (the primary distributor) is assumed by a particular node. A logical name can be used to identify this role and is referred to as the role name. A role name has the following format:

<name::> where name is 1 to 8 characters, the less than and greater than characters (< and >) are required delimiters, and double colons (::) indicate that this is a role name.

The use of a role name makes it easy for an application to open a file or IPC queue on a node that provides a particular service when the service may move from node to node as conditions change. The primary distributor is an example of a service provided by DDS. Applications can be written that provide other services and a role name can be defined for them.

Reserved Role Names

The following role names are reserved by DDS and are used to identify the primary distributor and backup distributor nodes. These names are dynamically maintained by DDS and cannot be modified by the user.

Role Name : Reserved for:

FDSFDXCP : Configured primary distributor

FDSFDXCB : Configured backup distributor

FDSFDXAB : Acting backup distributor

Broadcast Domains

Broadcast domains are useful for maximizing performance and minimizing resource utilization when distributing messages or files to a large number of nodes. Within a broadcast domain, DDS exploits the LAN hardware's ability to broadcast a datagram to all nodes. The DDS Data Distribution component supports distributing files to all nodes within a broadcast domain. The Interprocess Communications component can send a message to every node within a broadcast domain.

A subset of nodes within a system can be grouped into a broadcast domain. The broadcast domain name has these rules:

- Broadcast domain names can be from 1 to 8 characters (blanks are not allowed). Greater than and less than characters (< and >), question marks (?), asterisks (*), and colons (:) are not valid characters.
- Broadcast domain names are case-sensitive.

- Do not define broadcast domain names whose first three characters are FDS. These names are reserved.
- DDS currently supports a maximum of one broadcast domain.

Distribution Roles and Domains

Each node can assume a distribution role. There are three possible distribution roles:

Primary Distributor

The primary distributor controls the primary copy of all distributed files. Only the primary copy of a file can be changed directly by an application.

Backup Distributor

The backup distributor controls the backup copy of a file and can take over for the primary distributor if the primary distributor fails or is deactivated.

Subordinate

All other nodes configured with the Data Distribution component are considered subordinates and manage image copies of distributed files.

A node can be configured with a distribution role, but under certain circumstances can assume another role. In that instance, the node is said to be acting the role. Specifically, when a node that is configured as the backup distributor assumes the role of the primary distributor that has failed, the backup distributor is said to be the acting primary distributor.

A group of nodes forms a distribution domain. There are two types of distribution domains:

Mirrored domain

A mirrored domain is defined to be the primary distributor and backup distributor. There is only one mirrored domain, so it is not named.

Broadcast domain

A broadcast domain can include zero or more nodes. Files that are distributed to a broadcast domain are distributed to each node in the domain, as well as to the acting backup distributor.

A given file can be distributed to only one domain.

By default, all files on a node are local files. A local file is a file that is not distributed. That is, it is not a primary copy, backup copy, or image copy. A file or subdirectory on the acting primary distributor can be made distributed using either the Data Distribution Utility or an application that calls the FdsSetDistribution API. (See QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide for more information about the FdsSetDistribution API.)

File Names and Queue Names

For a specific node, you can identify each file using the operating system path where the file is located and the file name. The operating system path and file name are called the file specification. Similarly, you can identify a queue using the queue name.

However, with DDS you can access files and queues on any node. Therefore, you must use a retail path specification to identify the file or queue. A retail path specification contains a node specification or broadcast domain specification as a prefix to the file specification or queue name.

Note: If you do not include a node or broadcast domain specification in the retail path specification, DDS assumes that the file or queue resides locally.

A node specification can be in one of these forms:

- A node name followed by two colons, which specifies the ID of the node directly.
For example:
NodeID::
- A role name followed by two colons and delimited with the greater-than and less-than signs. The role name will resolve to the ID of the node that is acting in the role. For example:

<RoleName::>

A broadcast domain specification is valid only to identify a queue. It must be a string, containing no blank characters, that includes a broadcast domain name followed by two colons. For example:

B_DOMAIN::

Assume that you have a file called MYFILE.DAT located in the subdirectory D:\FILES on the primary workstation with a node ID of Node1. The following retail path specifications are valid:

D:\FILES\MYFILE.DAT

Node1::D:\FILES\MYFILE.DAT
<FDSFDXAP::>D:\FILES\MYFILE.DAT

You can also use logical names for any part (or all) of a retail path specification. See “Logical Names and Role Names” on page 10 for more information about logical names.

Note: If you use a logical name for a file name or queue name in a retail path specification that contains a remote node ID or role name, the logical name is resolved on the remote node.

Components

The components of DDS are:

Name Services

Provides a name-resolution capability. This allows applications to use logical names, or aliases, instead of hard-coded file names, IPC queue names, and node IDs. These logical names are dynamically resolved at run time. Refer to the QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide for more information about the Name Services component.

File Services

Allows you to access both local and remote files. File Services can be optionally configured on zero or more nodes to share files with other nodes. The configured primary distributor and configured backup distributor share retail files by default.

Three types of files are supported:

- Keyed files
- Sequential files
- Binary files

Refer to the QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide for more information about the File Services component.

File System Interface

Provides support for distributing native operating system files, referred to as byte stream files.

This component also provides a disk I/O prioritization mechanism that overrides the standard operating system prioritization scheme. DDS prioritizes disk I/O that is

based on thread priority, allowing high priority requests, such as price lookups, to be processed ahead of lower priority requests.

Interprocess Communications (IPC)

Provides a peer-to-peer messaging service that allows application programs to send and receive messages. The messaging service is provided between processes running on a single node (intranode IPC) and on different nodes (internode IPC). The internode IPC function is a configuration option. It is required on all nodes in a system unless the system consists of a single stand-alone node.

The following LAN media are supported:

- Ethernet
- Token ring

Note: If you install the DDS 4690 Controller Services Facility, DDS also provides store loop support.

Data Distribution

Provides a distributed file capability that replicates data to multiple nodes, keeping each image synchronized during normal operations. It also performs reconciliation when failed nodes are returned to service. You can use these functions through a utility. Refer to the **IBM Distributed Data Services/Controller Services Facility for Windows Programming Guide** for more information about the Data Distribution component. Refer to the **IBM Controller Services Facility for Windows User's Guide** for more information about the utility.

Node Control

Allows you to perform administrative functions such as viewing information about the nodes within the DDS system and activating or deactivating the primary distributor. Node control is also responsible for synchronizing the time and date of all nodes within a system.

There is also an API that generates a list of all node IDs known to the DDS system, including nodes that DDS has detected as being active on the system and user-defined nodes that are not yet active.

Refer to the **QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide** for more information about using the node list API. Refer to the **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide** for more information about using the Node Control Utility.

Problem Determination and Analysis

Collects problem determination information. You can select the system message logs, system error logs, and system dump files you wish to work with. Refer to the **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide** for more information.

4690 Controller Services Facility (CSF)

You can install this facility on zero or more nodes. It supports 4680/4690 Operating System controller applications running under Windows NT, 2000, XP, or Server 2003 and the attachment of registers running the 4690 Operating System. It also supports the attachment of registers running Terminal Services feature. It provides support for DDS 4690 Controller Services Facility nodes to interact with other DDS 4690 Controller Services Facility nodes. See **Chapter 9. DDS/CSF Installation and Configuration** for more information about the 4690 Controller Services Facility.

File Distribution

DDS enables the distribution of files to other nodes in a distribution domain. When a file operation is directed to a controlled drive, DDS determines if the file has been defined as a distributed file or is in a distributed subdirectory, and then distributes the file operation as appropriate to other nodes.

File distribution is performed with no user intervention. Any operating system command (for example, COPY or ERASE) or application program statement that results in the modification of a distributed file causes DDS to distribute the operation to other nodes.

DDS must be running to detect whether a file is distributed. Until DDS is started, it assumes that all files on a controlled drive are distributed. Therefore, any attempt to change a file on a controlled drive when DDS is not running results in an error.

Note: If the **DDActive** configuration keyword is set to NO, DDS does not check to see if files on controlled drives are distributed, so this error is not returned. See the Data Distribution chapter in the **QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide** for more information about distributing files. See “Designating Controlled Drives” on page 23 for more information about controlled drives.

Disk I/O Prioritization

Processing of certain file I/O operations is time critical in the retail environment. For example, when scanning items at the point-of-sale terminal, the salesperson expects a consistent response time. If the processing load on the disk used to service price

lookup requests increases, the response time experienced by salespeople scanning items at the point-of-sale terminals must remain relatively constant.

DDS uses a disk I/O prioritization scheme to ensure that time-critical processes are given highest priority for disk access. The standard prioritization scheme of the operating system is overridden by this prioritization scheme, so that file access is granted based on the operating system priority of the process thread issuing the request.

To take advantage of disk I/O prioritization, an application program can increase the priority level of time-critical threads using the `SetThreadPriority()` API on Windows NT, 2000, XP, or Server 2003.

Disk I/O prioritization is most effective when DDS controls all partitions on a single physical disk. You should consider this fact when configuring partitions as controlled drives. See “Designating Controlled Drives” on page 23 for more information about controlled drives.

Chapter 2. Planning for DDS/CSF Installation

When you plan your DDS/CSF installation, there are several factors you must consider, including:

- System topology
- Hardware requirements and considerations
- Software requirements and considerations
- Workstation preparation
- Type of installation and configuration

This chapter provides information on all these factors.

System Topologies

You can install and use DDS on any system topology that supports the Windows NT, 2000, XP, or Server 2003 operating systems. Examples of valid system topologies include:

- “Stand-Alone Node” on page 19
- “Multiple Nodes on a Single LAN” on page 19
- “Multiple Nodes on Multiple LANs” on page 19
- “Multiple Servers on Multiple LANs” on page 20

The following notation is used in the illustrations for each topology:

- Boxes represent nodes.
- Lines connecting boxes represent LANs.
- P represents the primary distributor.
- B represents the backup distributor.
- S_x, where x is a digit, represents a server.
- Digits and lowercase letters represent all other DDS nodes.

Keep in mind that the following details when considering the topology illustrations:

- A LAN can be a single LAN segment, or multiple segments connected by one or more bridges (a bridged LAN).
- The backup distributor is optional in each of the topologies. However, a backup distributor is strongly recommended. Without one, DDS provides no recovery from loss of data due to hardware failures on the primary distributor. All DDS nodes in the same system must be located on the same LAN, or on LANs connected by bridges that are configured to forward broadcast messages.

All DDS nodes in the same system must be located on the same LAN or on LANs connected by bridges that are configured to forward broadcast messages.

Stand-Alone Node

This topology is the simplest. It typically represents a single-node specialty store using the Services of DDS.



Figure 1. Stand-Alone Node

Multiple Nodes on a Single LAN

This topology represents specialty stores, supermarkets, mass merchants, and small-to-moderate sized department stores.

See “System Topologies” on page 18 for information about the notation in the illustration.

All nodes in the picture, including the primary and backup, must be in the same IP subnet in order to use DDS to communicate. Multiple Nodes on Multiple LANs




Figure 2. Multiple Nodes on a Single LAN

Multiple Nodes on Multiple LANs

This topology supports up to four LANs with all nodes attached to a LAN to which both the primary distributor and backup distributor are attached.

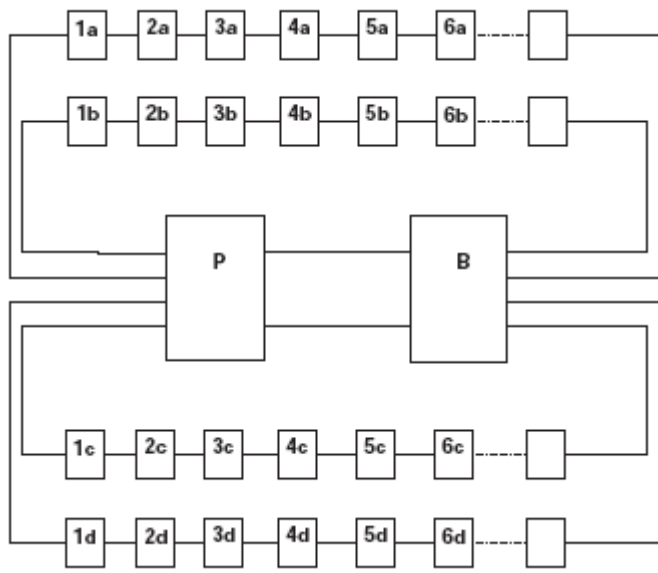


Figure 3. Multiple Nodes on Multiple LANs

See “System Topologies” on page 18 for information about the notation in the illustration.

DDS supports up to four LAN adapters. The number of adapters that a particular node can support depends on the hardware involved.

Unless the LANs shown are bridged, each LAN to which the primary and backup are attached must be located in a different IP subnet in order for the primary and backup distributors to communicate with the DDS nodes on that LAN. However, DDS does not provide IP routing between IP subnets. Therefore the nodes on each LAN will not be able to communicate with the nodes on other LANs using DDS because they are in different IP subnets. For example, node 1a will not be able to communicate with node 1b using DDS.

Multiple Servers on Multiple LANs

This topology supports up to four LANs with all nodes attached to a LAN to which both the primary distributor and backup distributor are attached. Four additional nodes are configured as servers. These servers provide a service to other nodes, such as file access if the File Services component is configured on the node to share files, or an application-defined service using the Interprocess Communications component. See the **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide** and the **QVS Controller Services for Windows Programming Guide** for more information.

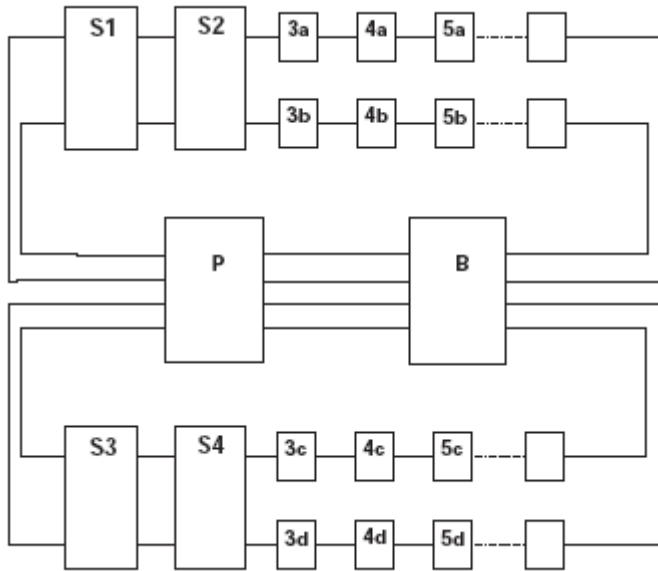


Figure 4. Multiple Servers on Multiple LANs - Software Requirements

S1 and S2 can serve all nodes that are on the same LANs as 3a and 3b. Servers S3 and S4 can serve all nodes that are on the same LANs as 3c and 3d.

See “System Topologies” on page 18 for information about the notation in the illustration.

DDS supports up to four LAN adapters. The number of adapters that a particular node can support also depends on the hardware involved.

Hardware Requirements

To install and use DDS, you need the following:

- A hardware platform that supports Windows NT, 2000, XP, or Server 2003. . On Windows NT, DDS/CSF runs on a single processor workstation. On Windows 2000, Windows XP, or Windows Server 2003, DDS/CSF can run on either a single processor or dual processor workstation. See “System Topologies” on page 18 for examples of valid system topologies
- A minimum of one hard disk on each workstation. See “Partitioning Your Hard Disk” on page 23 to determine the size requirement for the hard disk.
- An adapter that provides a network adapter driver that complies with Network Driver Interface Specification (NDIS), if you intend to communicate across a LAN.

Microsoft® Corporation and 3Com Corporation jointly developed the NDIS as a medium access control (MAC) interface for LAN adapter drivers and protocol drivers. NDIS has become an industry standard, providing a common, open interface that enables network adapters and LAN software from different manufacturers to communicate with each other.

Your communications adapter card manufacturer can tell you if your adapter cards are NDIS-compliant.

You should also consider an uninterruptible power supply for any workstation that performs a critical function. An uninterruptible power supply can ensure against failures, loss of data, and damaged files caused by temporary interruption of power. **Note:** For more information about the 4690 Controller Services Facility, see **Chapter 9. DDS/CSF Installation and Configuration**.

Software Requirements

To install DDS or DDS/CSF, you must be running one of the following operating systems on each workstation:

- Windows NT Version 4.0 Workstation with Service Pack 6
- Windows NT Version 4.0 Server with Service Pack 6
- Windows 2000 Professional with Service Pack 4
- Windows 2000 Server with Service Pack 4
- Windows XP with Service Pack 1 or 2
- Windows Server 2003

Note: To install or delete DDS, you must log on to an account with Administrator privilege.

Refer to the README.FDS file on the installation disk for information about any Windows service packs that may be required to run DDS. For information about upgrading an existing DDS/Windows NT installation to a DDS Version 3/Windows 2000 installation, see **Upgrading an Installed Version of DDS**.

Note: For more information about the 4690 Controller Services Facility, see **Chapter 9. DDS/CSF Installation and Configuration**.

Preparing Your Workstation

For each workstation you intend to use with DDS, you must:

1. Plan how to partition your hard disk. You need to partition your hard disk to set up controlled drives that will be used by DDS. See **Partitioning Your Hard Disk** for more information about hard disk partitions.

2. Format drives. You must format all disk partitions that you intend to use.

3. Install Windows NT, 2000, XP, or Server 2003. Follow the installation instructions to install the operating system. During the installation process, Windows prompts you for partitioning information. Use the partitions and sizes planned in **Partitioning Your Hard Disk**. From a command line, format all drives that have not already been formatted. If you plan to use the 4690 Controller Services Facility, read **Chapter 9. DDS/CSF Installation and Configuration** for information about Windows requirements for DDS.

4. Plan your TCP/IP configuration. You need to configure the IP address of each interface on your workstation if you plan to use DDS to communicate with other DDS workstations across a TCP/IP network. See **Planning for TCP/IP Configuration** for more information about configuring the IP addresses for your workstation.

Note: When planning your directory structure, keep in mind that DDS cannot be installed in a directory where any name in the path has embedded blanks. For example, the following directory cannot be used: d:\prod dir\fds

Partitioning Your Hard Disk

Before partitioning your hard disk for use with DDS, it is important that you understand the concept of controlled drives and how configuring a partition as a controlled drive will affect your system. The following sections describe the DDS features that are dependent on controlled drives, which partitions can and cannot be controlled drives, and the restrictions associated with defining and using a hard-disk partition as a controlled drive.

Designating Controlled Drives

During configuration you can designate one or more hard-disk partitions as controlled drives. (See **Generating a New Configuration File** and **Designating Controlled Drives** for detailed information about designating hard-disk partitions as controlled drives.) DDS controls the hard-disk partitions that are designated as controlled drives for the purpose of providing the following services:

- Disk I/O prioritization
- File (data) distribution

See **Disk I/O Prioritization** and **File Distribution** for descriptions of these services.

Note: Sequential and keyed file services can be used with data on a drive whether or not it is a controlled drive.

Depending on the version of Windows, you can choose to use the file allocation table (FAT) file system, the New Technology File System (NTFS), FAT32 on Windows 2000, or a combination of the three as the native file system for the partitions you

configure as controlled drives. However, it is recommended that you use NTFS, as it has been shown to provide improved reliability over FAT.

Any hard-disk partition can be configured as a controlled drive except the boot drive and the partitions used to store the paging file (PAGEFILE.SYS). To find out which partitions are used to store the paging file:

1. Click Start on the task bar.
2. Point to Settings.
3. Point to Control Panel.
4. Select System.
5. Take one of the following actions:
 - Select the Performance tab or
 - Select the Advanced tab then the Performance Options.
6. Select Change.

Note: When formatting controlled drives, you should allow enough space on each controlled drive for the complete 4690 image you intend to load on it plus the largest distributed file that will be a mirror image of the primary distributor's controlled drives plus the largest file that will be distributed by the primary distributor. This action allows space for temporary files that are created on non-primary workstations during distribution. Failure to allow sufficient space in this manner can result in out-of-space errors later.

Diskette drives, CD-ROM drives, RAM disk drives, and remote drives cannot be configured as controlled drives.

Controlled Drives and Disk I/O Prioritization

Disk I/O prioritization occurs automatically on any hard-disk partition that is configured as a controlled drive. You do not need to use any special commands or APIs to gain the benefits of disk I/O prioritization. All of the file system APIs from Windows and any of the DDS file services APIs can be used without modification. If you are not using the Data Distribution component, there is only one functional difference in having DDS control the partition rather than the native file system of Windows NT, 2000, XP, or Server 2003: the I/O to the disk is prioritized based on the priority of the process issuing the I/O request. If two I/O requests are issued at the same time, the I/O request from the process with the higher priority is completed before the request from the process with the lower priority.

This behavior is important in a store environment, where the I/O issued from the process performing price-lookup requests should receive priority over the I/O issued from less time-critical processes. You should therefore partition your disks so that the partitions that will be used to store the price-lookup file and other files accessed during a sales transaction are configured as controlled drives.

Disk I/O prioritization is most effective when all partitions on a single physical disk are controlled by DDS. If a workstation only has a single hard disk, it will not be possible for all the partitions on that disk to be controlled drives because the boot drive cannot be a controlled drive. However if a workstation has two or more disks, disk I/O prioritization will be maximized if one or more of the disks are partitioned such that all partitions on the disk are controlled drives.

Controlled Drives and Data Distribution

Before deciding which hard-disk partitions to configure as controlled drives, you should consider the unique operating characteristics of controlled drives, based on whether you are using the Data Distribution component.

If you are not using the Data Distribution component, you are not required to have more than one partition or to specify any partition as a controlled drive. If you do configure a partition as a controlled drive for the purpose of having the I/O on that partition prioritized, any file or directory operation on the controlled drive is allowed at any time on all nodes.

When you are using the Data Distribution component, there must be a minimum of two hard-disk partitions on each workstation. At least one of these partitions must be configured as a controlled drive. After defining which files or directories are distributed, data distribution occurs automatically on any hard-disk partition that is configured as a controlled drive. When you are using the Data Distribution component, certain types of I/O operations on controlled drives are restricted under some conditions. You should consider these restrictions when you decide which partitions to configure as controlled drives and which files to store on them.

These restrictions are:

1. File I/O that results in changes to a file or directory requires DDS to be running. Because distributed files and directories can only be modified on the acting primary distributor, the File System Interface must be able to determine if a file or directory is distributed before allowing it to be modified. DDS must be running for this determination to be made. Examples of operations that are restricted when DDS is not running include:
 - Opening a file for write access
 - Creating, renaming, or deleting a file or subdirectory
 - Writing to a file
 - Copying a file when the target file of the copy resides on a controlled drive

Note: Files and directories in the target installation and WORK directories that reside on a controlled drive are exceptions to this restriction. For these

directories, any I/O operation is allowed on any node, regardless of whether DDS is running or has stopped. (The target installation directory is the directory defined by the FSDIR System Environment Variable. The WORK directory is located in the target installation directory unless a different directory is specified during DDS configuration.) No file or directory in the target installation or WORK directories can be distributed. When these directories reside on a controlled drive, the file system interface does not require that DDS be running in order to determine if the operation should be distributed.

2. Any file handle that is obtained by opening a file for write access on the primary distributor cannot be used to modify the file after the primary distributor is deactivated or DDS is stopped. Even if the file remains open until the primary distributor is activated or DDS is restarted, the file handle cannot be used to modify the file; it must be closed and reopened. The file handle can, however, continue to be used for read or query operations both during the time that DDS is stopped and after it has been restarted.

Any file I/O that does not result in a changed file or directory is possible at all times on all nodes. Examples of operations that fall into this category include:

- Opening a file for read-only access
- Starting a program (when an .EXE or .CMD file is started, the file is opened for read access only)
- Issuing a directory command to see a list of files and directories
- Reading a file
- Copying a file when the source file of the copy resides on a controlled drive

Controlled Drive Restrictions

You can reformat a hard-disk partition specified as a controlled drive by using the format command. A restriction placed on controlled drives is that you cannot use the /FS:xxxx option of the format command. Use this option to change the file system from NTFS to FAT or FAT32 or from FAT or FAT32 to NTFS.

To change the file system using the format command:

1. Delete the drive letter from the list of letters specified as controlled drives in the DDS configuration file or by using the Configuration File Utility.
Note: Before deleting the drive letter as a controlled drive, make sure that all files and directories on the drive are local (not distributed).
2. Reboot your system.
3. Use the format command with the /FS:xxxx option to change the file system.

After reformatting the partition, you can designate it as a controlled drive again by adding the drive letter back to the DDS configuration file (or by using the Configuration File Utility) and restarting your workstation.

Example for Partitioning Hard Disks

For this example, assume the following configuration for the workstation on which you are planning to install DDS:

- The workstation will be a primary distributor or backup distributor.
- You are installing Windows NT, DDS/CSF, and some additional products.

Note: Please allow for 50MB of disk space for the installation of DDS/CSF.

The steps for partitioning the hard disks are as follows:

The first physical hard disk can be partitioned with two partitions: one for the OS and one as a controlled drive(s). The user should assure that the physical drive is large enough to contain the OS, the 4690 image and any expected code growth from either source. The C hard disk partition will not be a DDS controlled drive, and therefore cannot contain data that must be distributed or subject to I/O priority control. If your 4690 OS uses two physical hard drives (C: and D), then consider creating two additional partitions as follows:

C: - contains the Windows OS and windows applications including DDS/CSF.

D: - contains the 4690 OS image copied from the 4690 C: drive

E: - contains the 4690 OS image copied from the 4690 D: drive (if used)

This approach provides an organized separation between your product files and your mission-critical data that are managed by DDS.

Note that logical drives D: and E: will be controlled drives managed by DDS/CSF.

Optionally a user could install and configure a second or even a third physical drive to hold the C: and D: 4690 images. The advantage for this would be possibly better performance and the requirement to restore only a part of a system in case of a hard drive crash.

Planning for TCP/IP Configuration

DDS requires that a set of nodes (workstations) that are going to communicate with each other must be part of the same DDS system. A DDS system is determined by the following characteristics:

- The DDS System ID (defined during installation)
- The beginning TCP/UDP port number to be used by IPC (defined by the DDS configuration keyword **IPCPortStart**)
- The IP subnet ID (determined by the workstation IP address and subnet mask)

DDS nodes can communicate only with other DDS nodes that have been configured to use the same beginning TCP/UDP port number to be used by IPC and that are in the same IP subnet.

To determine if a set of nodes is within the same IP subnet, you must first understand TCP/IP addressing and how DDS uses it. Every network interface adapter that a node has connected to a TCP/IP network must have an IP address that is unique on that network. An IP address is a 32-bit value that is usually expressed as four integers, one for each byte of the address. The integers are separated by ".", which is referred to as a "dot". For example: 192.168.5.6

This representation of an IP address is referred to as dotted decimal notation.

In addition to its IP address, each interface has an associated subnet mask. This value is also a 32-bit number, and is usually represented in dotted decimal notation. A TCP/IP subnet (or IP subnet) is a logical grouping of nodes based on their IP addresses and subnet mask.

Nodes are within the same TCP/IP subnet only if they have the same subnet ID. The subnet ID is obtained by logically ANDing the IP address with the subnet mask. Example 1 illustrates two nodes, each having a single interface:

Example 1:

Node A

IP: 192.168.0.200 0xC0A800C8

SubnetMask: 255.255.0.0 0xFFFF0000

SubnetId : 192.168.0.0 0xC0A80000

Node B

IP: 192.168.200.5 0xC0A8C805

SubnetMask: 255.255.0.0 0xFFFF0000

SubnetId : 192.168.0.0 0xC0A80000

In Example 2, the subnet IDs of nodes A and B are the same, so both nodes are on the same IP subnet. Assuming that there is a working physical network path between these nodes, the nodes can communicate using the TCP/IP protocol.

Example 2:

Node A

IP: 192.168.0.200 0xC0A800C8

SubnetMask: 255.255.0.0 0xFFFF0000

SubnetId : 192.168.0.0 0xC0A80000

Node B

IP: 192.169.200.5 0xC0A9C805

SubnetMask: 255.255.0.0 0xFFFF0000

SubnetId : 192.169.0.0 0xC0A90000

In Example 2, the subnet IDs of nodes A and B are different. Therefore, these nodes are in different IP subnets. Even if there is a working physical network between these nodes, they will not be able to communicate using TCP/IP.

Note: If you plan to have multiple DDS systems on the same network, in addition to configuring two different system IDs, it is a good idea to configure different IP subnet IDs for each DDS system. Multiple DDS systems can be on the same IP subnet. However, you may reduce congestion by using different IP subnets for the different DDS systems, or by using a different beginning TCP/UDP port number for each DDS system.

If a DDS node has multiple network interface cards attached to different LANs, a unique IP-address/subnet-mask pair must be assigned to each network interface card. In addition, multiple IP-address/subnet-mask pairs can be assigned to a single network interface card. DDS selects one IP-address/subnet-mask pair for each network interface card and calculates the subnet ID for that IP address/subnet mask pair. As a result, DDS communicates only with other nodes in the same subnet as the IP-address/subnet-mask pair selected. A message is logged during initialization to identify each IP address and subnet mask that DDS is using. To make certain that DDS selects the IP address that will enable communication with other DDS nodes, you may have to block one or more IP addresses. See Blocking Interfaces for more information.

Name Resolution

DDS does not make Domain Name System (DNS) or Windows Internet Name Services (WINS) queries to establish connections. Name resolution is based on node

IDs and IP addresses rather than Windows computer names or TCP/IP host names. As a result, DDS needs neither a DNS name server nor a Windows WINS server.

Dynamic Host Control Protocol

DDS does not support Dynamic Host Control Protocol (DHCP) under Windows NT, 2000, XP, or Server 2003.

Blocking Interfaces

DDS uses message broadcasting as part of its internode communications. These messages are subnet-directed, which means that they will be received by all nodes on the same IP subnet as the sending interface. DDS selects one subnet ID for each network interface card. If a node has multiple network adapter cards, or has multiple IP-address/subnet-mask pairs defined for a single network adapter card, it may be necessary to block one or more addresses. Doing this will ensure that DDS uses the one address that will allow it to communicate with other DDS nodes.

If a workstation has multiple network interface cards, DDS will, by default, select one IP address/subnet mask pair for calculating the subnet ID that DDS will use for that interface. In some cases, not all of the configured interfaces will be connected to a network being used by DDS. In order to prevent DDS from sending unneeded broadcasts to non-DDS networks, you can configure the network interface as "Blocked".

If a workstation has multiple IP-address/subnet-mask pairs defined for a single network adapter card, it may be necessary to block some of the IP addresses to ensure that DDS uses the IP address/subnet mask in the same subnet as all the other DDS nodes

Automatic Primary Distributor Switch-Over

You can configure DDS to automatically activate the acting backup distributor as the primary distributor. This activation occurs after the acting backup distributor has lost communication with the node that was performing the primary distributor role. See **QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide** for a detailed description of the Automatic Primary SwitchOver feature of DDS.

Because an automatic switch-over is triggered by a loss of communication between the primary and backup distributors, the selection and configuration of your LAN hardware is critical in preventing the backup distributor from switching over when the acting primary is still running. For example, if the primary and backup distributors are connected to different wiring concentrators or LAN segments, a break in the cables connecting the concentrators or LAN segments can result in a loss of

communication between the primary and backup machines. This condition would trigger an automatic switch-over and result in two active primary distributors.

To prevent this situation, DDS requires that the configured primary and backup workstations must be connected to the same LAN wiring concentrator, such as a token-ring multistation access unit (MAU) or Ethernet hub.

The accidental removal of LAN cables is also a common occurrence. To help prevent switch-over from occurring unnecessarily when a LAN cable is removed, it is recommended that two network interface cards be installed in both the configured primary and the configured backup distributors. The second network interface card provides a redundant network connection between the two machines so that when a single cable is removed, the primary and backup distributors can still communicate. To further reduce the chances of a cable being accidentally removed, one of the two network interface cards in each machine should be an Ethernet adapter that is directly cabled to the other machine using an Ethernet crossover cable. This cabling method will eliminate the need to route the cable to a wiring closet or hub, thereby reducing the probability that the cable will be accidentally removed.

If multiple network interface cards are installed in either workstation and they are attached to multiple LAN media then one of the following choices must be true:

- The configured primary and backup distributors must be connected to all of the same LAN media. For example, if the configured primary distributor is connected to token-ring and Ethernet LANs, and DDS is using both interfaces to communicate, then the configured backup distributor must also be connected to both the token-ring and Ethernet LANs.
- The IP addresses on any network interface cards that are attached to LAN media that are not common between the configured primary and the configured backup distributors must be blocked from use by DDS. For example, if the configured primary distributor is connected to a token-ring and an Ethernet LAN, but the configured backup is connected only to the Ethernet LAN, you must block the token-ring interface so that DDS only uses the Ethernet LAN.

If DDS detects that there are two acting primary distributors, a message is logged in the Windows NT, 2000, XP, or Server 2003 application event log. The machine that automatically switched over will be deactivated and reconciled with the remaining acting primary distributor. This reconciliation can result in lost data. Any changes that were made to files on the machine that deactivated will be overlaid during reconciliation by data from the remaining acting primary distributor. If files are distributed to a broadcast domain of subordinate nodes, DDS will also attempt to synchronize the distributed files on all subordinate nodes with those on the acting primary distributor. In some cases this synchronization may not be possible, so it is

recommended that you force a redistribution of all files/directories when this situation is detected. A redistribution can be forced using either the FDSSETD command or the Data Distribution Utility.

IP Address Computer Name Switch-Over

The Automatic Switch-Over component also provides the capability to add an IP address, a computer name, or both to whichever workstation is activated as the primary distributor. This capability is independent of whether automatic primary switch-over is configured, and, therefore, is available when manually activating a primary distributor.

The IP address, or computer name, or both are added when a workstation is activated as the primary distributor and are removed when a workstation is deactivated as the primary distributor. When the acting primary distributor role moves from one workstation to another as a result of either an automatic or user-initiated activation, the IP address, or computer name, or both move to the new workstation. If no workstation is active as the primary distributor, these values will not be available to the network.

Any IP address or computer name added to the workstation performing the role of the acting primary distributor must be different from the IP address or computer name configured on any of your other workstations by Windows NT, 2000, XP, or Server 2003 network configuration. For example, if IP address 192.160.0.1 is defined on your configured primary distributor and 192.160.0.2 is defined on your configured backup distributor, then you must allocate an additional unique IP address, (for example: 192.160.0.3) to use with this feature. When the configured primary distributor is performing the acting primary distributor role, then it will be accessible by either the 192.160.0.1 or the 192.160.0.3 addresses. Similarly, when the configured backup distributor is performing the acting primary distributor role, then it will be accessible by either the 192.160.0.2 or the 192.160.0.3 addresses. When neither workstation is performing the role of acting primary distributor, the 192.160.0.3 address will not be active on either workstation. Part 3. Installing DDS

This part of the book guides you through installing your QVS DDS system. In addition, it tells you how to add DDS components to an installed system and how to uninstall (delete) DDS.

See **Chapter 5. Configuration Data** for configuration instructions and information.

Note: If you are installing 4690 Controller Services Facility, be sure to follow the instructions in **Chapter 9. DDS/CSF Installation and Configuration**.

Chapter 3. Installation

The installation program:

- Available in two flavors: One installs just DDS; the other installs both DDS and CSF.
- Creates the necessary subdirectories. For more information about the subdirectories created, see **Appendix A. Controller Services Facility Directory Structure**
- Updates the system environment and configuration information.
- Gathers permanent workstation information, which includes an identifier for the workstation and an identifier for the system to which the workstation belongs.
- Gathers initial configuration information.

Choosing Attended or Unattended Installation and Configuration

You can install and configure DDS or DDS/CSF by using one of two possible procedures:

Attended

With attended installation, you perform the installation and configuration process from each workstation where DDS or DDS/CSF will be installed. You can either:

- Install DDS or DDS/CSF by downloading the installation code to the hard drive of each workstation.
- Install DDS or DDS/CSF from each workstation by accessing a hard drive of a remote server where the installation package is located.

Typically, you would choose to perform an attended installation if you need to install DDS on a small number of workstations or if you are unfamiliar with installing products on a Windows operating system. See **Performing an Attended Installation** for instructions to perform an attended installation.

Unattended

With unattended installation, you build a code server, which can then be used to automate the installation and configuration of DDS or DDS/CSF on each workstation. Setup and response files provide installation and configuration information for each of the workstations.

Typically, you would choose to perform an unattended installation if you need to install DDS on a large number of workstations, are familiar with Windows operating system and installing products, and have a software distribution manager available.

See **Performing an Unattended Installation** for instructions about performing an unattended installation.

During initial installation of the DDS or DDS/CSF the DDS configuration process automatically starts. You can also run the DDS configuration program any time the installation program is run after the DDS or DDS/CSF has been installed. You can provide configuration information through a graphical user interface (when performing an attended installation) or by adding configuration keywords to a response file (when performing an unattended installation).

The DDS configuration program can be run independent of the installation process. You can reconfigure your workstation after DDS has been initially installed using either an attended or unattended procedure. See **Chapter 6. Configuration File Utility** for more information about performing an attended configuration. See **Chapter 8. Unattended Configuration** for more information about performing an unattended configuration.

Performing an Attended Installation of DDS or DDS/CSF

Note: You must log on to an account with administrator privilege to install or delete DDS.

1. Open a Command Prompt Window.
2. Access the drive where the DDS installation files are located.
 - a. Make sure that you are connected to the drive.
 - b. Type the drive letter and a colon (for example, X:) and press Enter.
 - c. Change directories to the subdirectory containing the installation files.
3. Type SETUP and press Enter.

Note: You can also run SETUP from the desktop. From the desktop, click Start, point to Programs, and click Windows NT Explorer or Windows 2000 Explorer. Then select the drive and folder where the installation package is located, and double-click SETUP.EXE

4. On the welcome screen, select Next. README.FDST, a file containing last-minute installation information, is displayed. Read it if you wish, then click Next.

5. Specify the directory where you want to install DDS or DDS/CSF. The Destination Directory field displays the installation directory where DDS will be installed (the target installation directory). The default target installation directory is C:\FDS.

To change the target installation directory:

- a. Select Browse. From the Choose Folder window, you can do one of two things:

- Select an existing directory or drive using the Directories and Drives graphic. After choosing an existing directory, click on OK
- Choose a new directory by typing over the C:\FDS in the Path dialog. After typing in the new directory, click on OK. **Note:** The target installation directory must be on a local hard drive. However, you cannot choose a root directory as the target installation directory.

- b. Select Next.

6. If you are installing the DDS-only version of the installation, the copying of files will proceed. Following the file copy, you will see a popup configuration screen as shown in Figure 5. Skip to step 11 to continue. If you are installing the DDS/CSF version of the installation, proceed to the next step.
7. The screen dialog now allows you to specify the 4690 store controller ID to be used by this controller. Follow the naming convention laid out by your store policies. Normally, the primary controller is set to be CC; the backup controller is set to be DD. The default value is CC. When complete, click on Next.
8. Choose the correct option regarding a Store Loop card. Store Loop is a controller-to-terminal communications protocol that has mostly been replaced by an Ethernet LAN. If you are using a Store Loop interface card to communicate with your terminals, pick the proper option. If not, choose the “No Store Loop Adapter” option. If you clicked on ISA Store Loop Adapter(s), click Next and go to the next step. Otherwise, click next and go to step 11.
9. If you clicked on the ISA Store Loop option, you must now specify how many adapters you are using. Choose either 1 or 2 and click Next.
10. If you chose either 1 or 2 Store Loop cards, you will be given the appropriate options to choose the IO addresses that each ISA card is set for. You should also choose the appropriate option for the IRQ. The IRQ must be the same for

both cards (assuming you install 2 cards). On the last option screen , click Next.

11. At this point, files will be copied to the proper directory. You will then see a configuration screen popup as shown in Figure 5.
12. Using the DDS Configuration popup screen, specify the node ID for this workstation. Node IDs are identifiers for workstations in the DDS system. The rules for node IDs are:
 - Each workstation within a system must have a unique node ID. Duplicate node IDs are not confirmed during installation, but they are detected when DDS initializes. If duplicate node IDs are detected, you must first uninstall DDS and then reinstall DDS, specifying the correct node ID.
 - No two workstations on a LAN can have the same node ID, even if they are members of different DDS systems (have different system IDs).
 - Node IDs can be from 1 to 8 characters (blanks are not allowed). Node IDs are case-sensitive.
 - Do not use node IDs whose first three characters are FDS. These names are reserved.
 - Greater than and less than characters (< and >), asterisks (*), question marks (?), and colons (:) are not valid characters for node IDs
 - After installation, node IDs can be changed only by reinstalling DDS or DDS/CSF.
13. Specify the system ID for this workstation. This identifier helps to define the set of workstations that make up a DDS system. The system ID can be up to 4 alphanumeric characters (blanks are not allowed). Do not use system IDs whose first three characters are FDS. These names are reserved.

If you assign a unique system ID to each system, the systems can be connected using a LAN or a gateway without causing messages intended for one system to be sent to another system.

If you assign the same system ID to each system, the LANs can be connected using a gateway; both LANs are considered the same system. If this type of system is used and you are using the Data Distribution component, there should be only one configured primary distributor and one configured backup distributor for the entire system, even though they might be on different LANs.

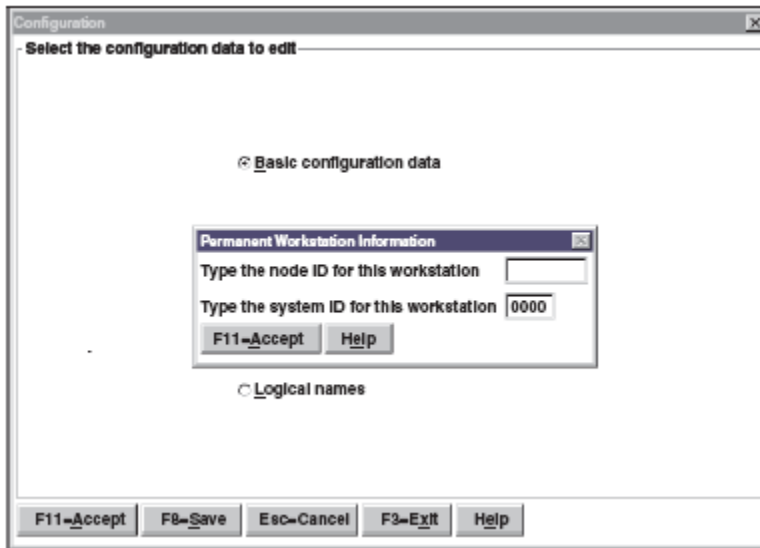


Figure 5. Permanent Workstation Information window

14. Select Accept to exit the Permanent Workstation Information window and to save changes.

When you exit the Permanent Workstation Information window, the configuration portion of the **Configuration File Utility** is started automatically.

15. You can choose to work with basic configuration data, advanced configuration data, or logical names. See **Chapter 5. Configuration Data and Logical Names and Role Names** for general information about configuration data and logical names.

Select the type of data you want to work with and go to the section indicated below for instructions. After you fill in the windows for the type of data you have selected, go to step 18.

For information about:

Basic configuration data, see **Basic Configuration Keywords**.

Advanced configuration data, see **Advanced Configuration Keywords**.

Logical names, see **Logical-Names-File-Only Configuration Keywords**.

16. Select Save. The configuration files are written to disk.

17. The Configuration File Utility-Main Window is displayed. Select Exit or press F3.
18. The Setup Complete window is displayed. From this window, you can choose to restart your computer now or later. However, you must restart your computer before you can use DDS.

Attention: Do not perform an uninstall before the computer has been restarted.

To restart now, select Yes, I want to restart my computer now and select Finish. When your computer has finished restarting, DDS will automatically be started.

Note: If you installed DDS from diskettes, you may receive error messages about drive A (or B) failing or not being ready when the computer is restarted. If you receive these messages, simply clear them. The operation of DDS will not be affected.

After you have successfully installed DDS, a folder called QVS DDS is created under Programs. To see the programs in the DDS program folder, click Start, point to Programs, and then select QVS DDS. Depending on which components are installed, you will be able to choose from the following icons:

- 4690 Console. This appears ONLY if you install the DDS/CSF version of the program.
- Configuration. See **Chapter 6. Configuration File Utility** for more information about using the Configuration utility.
- Keyed File. Refer to the **QVS Controller Services Feature for Windows User's Guide** for more information about using the Keyed File Utility.
- Node Control. Refer to the **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide** for more information about using the Node Control Utility.
- Problem Analysis. Refer to the **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide** for more information about using the Problem Analysis Diskette Utility.
- Data Distribution. Refer to the **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide** for more information about using the Data Distribution Utility.

- Readme File. Select Readme File to view the readme file for DDS.

At this point, if you are installing DDS only, your work should be complete. If you are installing DDS/CSF, please refer to **Chapter 9. DDS/CSF Installation and Configuration** for information on additional tasks you may have to complete in order to install a fully-working system.

Performing an Unattended Installation

To perform an unattended installation on a large number of workstations:

1. Create response files (see **Creating Response Files** for more information). At a minimum, you must provide a unique node ID for each workstation. In addition, you can specify the ID of the system to which the workstation belongs and specify unique installation or configuration values for your workstations.

You can create a different response file for each workstation, specifying the node ID and any unique installation and configuration values that are required.

However, you can significantly reduce the number of response files required by automating the specification of node IDs. This way, you will typically need only to create three response files: one for the primary distributor, one for the backup distributor, and one for all subordinate nodes.

2. Use the **SETUP** command from the target installation machine to initiate the unattended installation and configuration process. Parameter information for the setup command is found in **Starting Unattended Installation with SETUP**.

Note: Instead of using the **SETUP** command from the local machine, you could use a software distribution manager, such as TME 10 or Software Management Services, to manage the installation of the workstations.

Transferring DDS Installation Files to a Code Server

DDS provides installation modules which you should copy to a subdirectory on your code server.

Creating Response Files

Response files provide predefined responses to prompts that are normally answered by the user at a workstation during the installation and configuration process.

Response files literally replace the person who would respond to prompts. A response file is an ASCII text file that contains keyword=value pairs. The term keyword in this book refers to either keywords or to keyword=value pairs in response files. You can provide configuration keywords in a response file. For more information about configuration keywords, see

Chapter 7. Configuration Keywords.

Note: Any configuration keywords specified in the installation response file will be processed only if the DDS Base component is being installed for the first time or if the DDS Base component was previously installed and the keyword **RunConfigForUpdate=YES** is included in the response file.

You can use any standard ASCII editor or write a program to update the response files. You can also install DDS on the server and use the **DDS Configuration File Utility** to create response files. For more information about using the utility, see **Chapter 6. Configuration File Utility**.

During the first installation of the DDS Base component, you must specify a unique node ID and, optionally, the system ID for the system to which the workstation belongs. The node ID and system ID are only used when the DDS Base component is initially installed. Typically, you will also want to specify different installation or configuration values for the workstation, depending upon whether the workstation is a primary distributor, backup distributor, or subordinate node.

You can either create a separate response file for each workstation, or you can use the workstation's computer name, host name, or TCP/IP address to define the node ID, which allows you to reduce the number of response files you need to create. See **Automating Node ID Specification** for more information about using the workstation's computer name, host name, or TCP/IP address to define a node ID. DDS validates your response file keywords and verifies the values you provide for each keyword:

- When you use the **Configuration File Utility** to create or modify a response file
- During unattended installation as the response file is being processed

DDS also verifies configuration values when you start DDS.

Syntax Rules for Response Files

- The syntax rules for response files are as follows:
- Keywords are not case-sensitive.
- Some values are case-sensitive.
- Keywords can be in any order.
- There can only be one keyword=value pair per line.
- The maximum length for each line is 300 characters.
- Comment lines must begin with an asterisk or a semicolon, but they do not have to start in the first column of the file.
- Blank lines are considered comment lines.

- For keywords that specify more than one value, separate each value with blank spaces, a comma, or a semicolon. You can also use commas without a value as positional placeholders.
- The equal sign that separates keywords from values can be surrounded by blank characters but must be on the same line as its associated keyword and value.
- Each line must end with a carriage return (hex 0x0d), a line feed (hex 0x0a), or both. ASCII text editors normally include one or both of these bytes.

Automating Node ID Specification

When installing DDS on a large number of workstations, creating a response file for each workstation may not be practical. Therefore, DDS provides several options for automatically specifying the node ID for each workstation.

By automatically specifying node IDs, you can significantly reduce the number of response files that are needed to install DDS. You need only to specify a response file for the primary distributor, backup distributor, and subordinate nodes for each system you have. DDS provides values for the **NodeID** keyword that allow you to:

- Use the computer name of each workstation to define the node ID.
- Use the host name of each workstation to define the node ID
- Use a TCP/IP address of the workstation to define the node ID

Notes:

1. To automatically specify a node ID using either the host name or TCP/IP address, you must already have TCP/IP installed and configured.
2. You cannot specify these values using the Configuration File Utility.

If the computer names or host names defined for your workstations conform to the rules for node IDs, you can specify that either the computer name or host name be used as the node ID. To use the computer name as the node ID you can specify **nodeID=\$ComputerName** in your response files. To use the host name as the node ID you can specify **NodeID=\$HostName** in your response files.

Note: See **Configuration Keywords** for more information about the **nodeID** keyword. See ["Example Using Computer Names"](#) for an example of using the **NodeID=\$ComputerName** keyword.

If the computer names or host names defined for your workstations do not conform to the rules for node IDs, or if you plan to use a TCP/IP address to define the node ID, you must specify a valid node ID to match each computer name, host name, or TCP/IP address used. To do so, create a file, called FDSNODES.TBL, that lists each

computer name, host name, or TCP/IP address and the corresponding node ID (and optionally the system ID) to use for that workstation. Then, specify **nodeID=%ComputerName**, **nodeID=%HostName**, or **nodeID=%IPAddress** in your response files.

See **Automating Node ID Specification** for an example of using host names and the FDSNODES.TBL file to create response files.

Example Using Computer Names: Assume that you are going to install DDS on 100 workstations, that the computer names of the workstations are node_1 through node_100, and that you are going to make these nodes all part of the same system, which is 1000. Also assume that you are going to have a primary distributor, backup distributor, and 98 subordinate nodes.

You could create a separate response file for each workstation and specify the node ID for each one. However, you can automate the process by specifying three response files:

Note: The computer names must be unique even if the workstations are on different domains.

- A response file for the primary distributor (with the computer name node_1). The response file would contain this information:

```
NodeID=$ComputerName
SystemID=1000
DistributionRole=PrimaryDistributor
.....
Additional keywords, as needed
```

- A response file for the backup distributor (with the computer name node_2): The response file would contain this information:

```
NodeID=$ComputerName
SystemID=1000
DistributionRole=BackupDistributor
.....
Additional keywords, as needed
```

- A response file for all subordinate workstations (node_3 through node_100): The response file would contain this information:

Note: Subordinate workstations are not supported under DDS/CSF

```
NodeID=$ComputerName
SystemID=1000
```

DistributionRole=Subordinate

.....

Additional keywords, as needed

In each of these response files, the keyword **NodeID=\$ComputerName** is used to specify that the computer name for the workstation will be used for the node ID. After creating the response files, you would use the software distribution manager to distribute the response files and the DDS installation package to the appropriate workstations for installation.

Example Using Host Names: Assume that you are going to install DDS on 100 workstations using the host name of these workstations as the node ID. The host names for the workstations are computer1.lana.domain.com through computer100.lana.domain.com. Also assume that you are going to have a primary distributor, backup distributor, and 98 subordinate nodes.

Note: The information in this example also applies if you are specifying the TCP/IP address of the workstations as the node ID.

In this case, you cannot simply specify **NodeID=\$HostName** because the host name does not conform to the rules for node IDs. Therefore, you specify **NodeID=%HostName** in the response files and create a file called FDSNODES.TBL that will map the host names to valid node IDs.

Note: The file FDSNODES.TBL must reside in the same subdirectory as the response files.

Again, you would specify three response files:

- A response file for the primary distributor (with the host name of computer1.lana.domain.com). The response file would contain this information:

NodeID=%HostName

DistributionRole=PrimaryDistributor

.....

Additional keywords, as needed

- A response file for the backup distributor (with the host name of computer2.lana.domain.com). The response file would contain this information:

NodeID=%HostName

DistributionRole=BackupDistributor

.....

Additional keywords, as needed

- A response file for all subordinate workstations (host names of computer3.lana.domain.com through computer100.lana.domain.com). The response file would contain this information:
Note: Subordinate workstations are not supported under DDS/CSF

```
NodeID=%HostName
DistributionRole=Subordinate
.....
Additional keywords, as needed
```

After creating the response files, you would create the file called FDSNODES.TBL, which contains the host name, node ID, and system ID (optional) for each workstation.

- The file FDSNODES.TBL must be located in the same directory as the response files.
- If you have multiple host names for a workstation, you must specify all host names in FDSNODES.TBL, using the same node ID for the host names. You can use any ASCII editor to create FDSNODES.TBL, which would contain information such as:

```
computer1.lana.domain.com
node_1
1000
computer2.lana.domain.com
node_2
1000
...
computer100.lana.domain.com node_100 1000
```

In FDSNODES.TBL, specify the host name first followed by the node ID and, optionally, the system ID. Separate each of the values by one or more spaces. After creating the response files and the file FDSNODES.TBL, use the software distribution manager to distribute the response files, FDSNODES.TBL, and DDS installation package to the appropriate workstations for installation.

Configuration Keywords

This section contains a list of the keywords that are valid in an installation response file.

Include

This keyword is optional and specifies which general response files to include with a specific response file.

The physical placement of this keyword affects the resulting merged response file definition. For example, keywords from an included response file override duplicate keywords in the including response file if the physical placement of the Include keyword follows the occurrence of the original keywords. The included response file is handled in-line. If you specify an included response file, but the response file is not found, no message is recorded in the error log or history log.

Note: You cannot have more than five levels of included response files.

The following search order is used to find the general response files specified:

1. The fully-qualified file specification
2. The current directory (the directory from which the SETUP command is run).

For example:

```
Include = z:\install\rsp\genresp.rsp
```

NodeID

This keyword specifies the node ID for this workstation.

The rules for node IDs are:

- Each workstation within a system must have a unique node ID. Duplicate node IDs are not confirmed during installation, but they are detected when DDS initializes. If duplicate node IDs are detected when DDS starts, you must reinstall DDS to correct the problem.
- No two workstations on a LAN can have the same node ID, even if they are members of different DDS systems.
- Node IDs can have 1 to 8 characters (blanks are not allowed). The value for this keyword is case-sensitive.
- Do not use node IDs whose first three characters are FDS. These names are reserved.
- Greater than and less than characters (< and >), asterisks (*), question marks (?), and colons (:) are not valid characters for node IDs

- After installation, node IDs can be changed only by reinstalling DDS. In addition to specifying a node ID, you can specify one of these values:

The following values can be substituted for a valid Node ID:

\$ComputerName

Use the computer name of the workstation as the node ID. The computer name must conform to the rules for node IDs.

\$HostName

Use the host name of the workstation as the node ID. The host name must conform to the rules for node IDs.

%ComputerName

The node ID associated with the computer name (specified in the file FDSNODES.TBL) will be used as the node ID. To use this value, you must create the file FDSNODES.TBL and it must be placed in the same directory as the response files. The format of each line of the FDSNODES.TBL will be:

Computer_Name NodeID SystemID_(optional).

Each value is separated by blanks.

%HostName

The node ID associated with the host name (specified in the file FDSNODES.TBL) will be used as the node ID. To use this value, you must create the file FDSNODES.TBL and it must be placed in the same directory as the response files. The format of each line of the FDSNODES.TBL will be:

Machine_HostName NodeID SystemID_(optional).

Each value is separated by blanks.

%IPAddress

The node ID associated with the IP address (specified in the file FDSNODES.TBL) will be used as the node ID. To use this value, you must create the file FDSNODES.TBL and it must be placed in the same directory as the response files. The format of each line of the FDSNODES.TBL will be:

Machine_IP_Address NodeID SystemID_(optional).

Each value is separated by blanks.

For example:

NodeID = reg1

RunConfigForUpdate

Set this keyword to Yes when you want to run the configuration program during the installation program. The configuration program is run automatically during an initial installation. Otherwise, the configuration program is not run unless

RunConfigForUpdate is set to YES.

SystemID

This keyword specifies the system ID for this workstation. The value for this keyword is case-sensitive.

The **SystemID** value can be up to 4 bytes in length. System IDs cannot start with the characters FDS and cannot contain blanks or asterisks. The default value for **SystemID** is 0000.

This identifier helps to define the set of workstations that make up a DDS system. If you assign a unique system ID to each system, the systems can be connected using a LAN or gateway without causing messages intended for one system to be sent to another system.

Note: You do not need to specify a system ID keyword in the response file if you are using FDSNODES.TBL and the system ID is specified in that file. See the discussion above of how to define a NodeID using the FDSNODES.TBL. The SystemID can be optionally appended to each line in that file.

If you assign the same system ID to each system, the LANs can be connected using a gateway and the LANs will be considered the same system. If this type of system is used and if you are using the Data Distribution component, there should be only one configured primary distributor and one configured backup distributor for the entire system, even though they might be on different LANs.

A valid example of a system ID is: **SystemID** = st01

Common Processing Errors

The following common conditions produce errors:

- If a case-sensitive value is in the wrong case, it is not valid.
- If a required keyword is not provided or is not recognized, the requested action fails with an error.
- If an incorrect value is specified for an optional keyword that requires a value, the requested action fails with an error.
- If an incorrect value is specified for a required keyword, the requested action fails with an error.
- If a keyword is specified more than once, the value of the last occurrence is used unless multiple occurrences of the same keyword are specifically allowed.

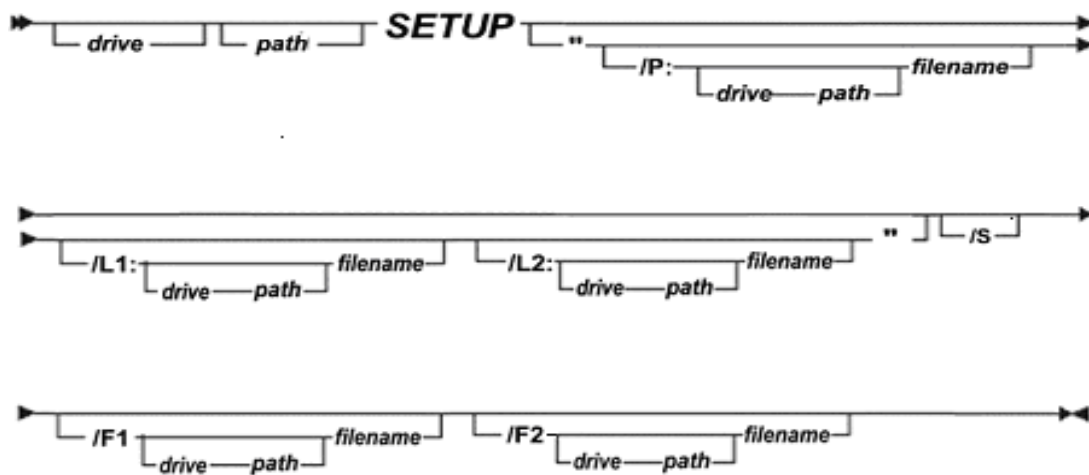
- If a keyword is not recognized, the installation and configuration program fails with an error.

Starting Unattended Installation with SETUP

In order to perform an unattended installation, you must first perform an attended installation using the “/R” parameter. This parameter causes the Windows installer to create an installation script that will be used later in executing the installation in unattended (silent) mode. This script will be located in the WINNT directory in a file: SETUP.ISS. Retain that file for use in running the silent installation.

The SETUP command performs an unattended installation by providing parameters to control DDS installation processing attributes. These attributes include the locations of the response files, history logs, and error logs. The syntax for this command is shown below:

Syntax for the SETUP Command



The parameters for this command are not case-sensitive. If the value of any of the parameters contains spaces, you must use double quotes (" ") to enclose the value. For example, you would specify a response file in the directory C:\RESPONSE FILE\FDSSUB.RSP as **/P:C:\RESPONSE FILE\FDSSUB.RSP**.

Note: Parameters may be specified in any order with one exception. The series of parameters enclosed in double quotes, if used, **MUST** be the first parameter(s) following the SETUP command. The double quotes allow the enclosed parameters to be treated as a “user” parameter. These parameters are passed intact to the FDSCR

/P:drive path filename

Specifies the drive, path, and file name of the specific response file. A response file must be specified for unattended installation.

/L1:drive path filename

Specifies the file name of the error log file. The drive and path are optional.

If you specify the log file, it will contain any history and errors associated with DDS or DDS/CSF installations.

These log files are also written to the client workstation, replacing any existing log files:

- The default installation log file is C:\WINNT\FDSSETUP.LOG. The subdirectory path (C:\WINNT) is determined by the WINDIR system environment variable.
- The default configuration error log file is target_install_directory\CONFIG\FDSCERR.LOG. This file is created only the first time you invoke SETUP.

/L2:drive path filename

Specifies the file name of the history log file. The drive and path are optional.

Notes: The value for this parameter may not contain hyphen (-) or slash (/). This parameter is processed during an initial installation only if you are also configuring DDS.

A history log file is also written to the client workstation, replacing any existing history log file. It is named target_install_directory\CONFIG\FDSCHIST.LOG.

/S

Specifies a silent installation and is required for unattended installation. Using this parameter ensures that no panels or user prompts are displayed to the screen during installation. The machine on which DDS is installed will automatically restart when SETUP is complete only if the installation response file specifies a restart.

/F1drive\path\filename

Specifies the fully qualified file name of the file created by the Windows Installer during the attended installation. The standard file name is SETUP.ISS. Note there are no spaces between the parameter and the fully-qualified file name.

Example: /F1c:\ddsinstall\setup.iss

/F2drive\path\filename

Specifies the fully qualified file name to use to write a log file created by the Windows installer during the unattended installation. This file should be

examined after the installation has been run in order to determine if it was successful or if a problem of some type has occurred.

Example of Starting Unattended Installation Using SETUP

To start an unattended installation using SETUP:

1. Set up the code server on drive D and type:

```
D:\INSTALL\DISK1\SETUP "/L1:D:\INSTALL\LOG\FDS\ERROR.LOG
                        /L2:D:\INSTALL\LOG\FDS\HISTORY.LOG
                        /P:D:\INSTALL\RSP\FDS\FDSDFLT.RSP"
/S
/F1D:\INSTALL\SETUP.ISS
/F2D:\INSTALL\SETUP.LOG
```

Note: The parameters in this example are shown on separate lines for clarity. Note also that the /L1:, /L2:, and /P: parms all together have to be enclosed in double quotes AND have to be the first parms following the SETUP command. Note also that the /F1 and /F2 parms are not immediately followed by a colon as are the /L1:, /L2:, and /P: parms.

In this example:

D:\INSTALL\DISK1\SETUP

Specifies that the DDS setup program is located in the D:\INSTALL\DISK1 path.

/L1:D:\INSTALL\LOG\FDS\ERROR.LOG

Specifies that all error information will be stored in ERROR.LOG, which is located in the INSTALL\LOG path.

/L2:D:\INSTALL\LOG\FDS\HISTORY.LOG

Specifies that all status information from the configuration will be stored in HISTORY.LOG, which is located in the INSTALL\LOG path.

/P:D:\INSTALL\RSP\FDS\FDSDFLT.RSP

Specifies that the response file, D:\INSTALL\RSP\FDS\FDSDFLT.RSP, will be used for the installation configuration commands.

/S

Specifies a silent installation.

/F1D:\INSTALL\SETUP.ISS

Specifies the file which contains the Windows Installer (InstallShield) installation script.

/F2D:\INSTALL\SETUP.LOG

Specifies the location and file name of a log file which will be created during the unattended installation on the target machine.

2. Invoke the software distribution manager you are using to start and manage the installation.

Upgrading an Installed Version of DDS

If you are upgrading from Windows NT to Windows 2000, XP, or Server 2003 on a workstation that already has DDS installed, you must first upgrade DDS to DDS Version 3, which supports Windows 2000, XP, or Server 2003 before upgrading Windows.

If DDS is already installed on your computer, run the installation program to upgrade to the current release. The upgrade will bring all installable DDS files up to the current release, with the following exceptions:

- The contents of the FDS\CONFIG and the FDS\WORK directories
- The user-modifiable DDS Base .BAT files, FDSBEGIN.BAT and FDSEND.BAT
- The user-modifiable Data Distribution .BAT files, FDSAP.BAT and FDSDP.BAT

Follow the instructions for either attended or unattended installation, and the upgrade will be performed automatically if the installation program detects that an earlier version of DDS is already installed. All installed components of DDS and DDS will be stopped before the upgrade is performed. After the upgrade is complete, you must restart your computer before you can restart DDS and DDS.

Note: The unattended installation would be run by recording a Windows Installation script using the SETUP /R sequence to create a script in WINNT\SETUP.ISS. The unattended installation would use that script file to install the unattended upgrade.

Resetting Controller Services Facility (CSF) Parameters

You can reset CSF parameters through either an attended or unattended process. The CSF-related parameters that can be changed are Controller ID, ISA Loop IO addresses, and ISA Loop IRQ settings. In order to change any of these values, you must run the command, **FDSCSFBH**. Please refer to that section for details on how to do this. This command can be run remotely.

Troubleshooting Installation Errors

If you encounter error messages while installing DDS/CSF:

- Restart the computer.
- If DDS/CSF is listed as an installed program, then uninstall it and restart the computer
- Begin the installation process again after correcting the installation error.

Chapter 4. Uninstalling DDS or DDS/CSF

Uninstalling deletes DDS or DDS/CSF from your workstation. You can perform either an attended uninstall or an unattended uninstall.

When you delete DDS, all system environment and configuration information added to your workstation by DDS or DDS/CSF is deleted. In addition, most files and subdirectories in the DDS installation directory (default is C:\FDS) are deleted. If you have stored any of your own files in these directories, these files may also be deleted.

To ensure that you do not adversely affect running applications that are using DDS when you attempt to uninstall the product, stop DDS before beginning the uninstall process. Some files may not be deleted immediately because they are in use. These files will be deleted when you reboot your workstation.

Performing an Attended Uninstall

You can perform an attended deletion of DDS or DDS/CSF from the Add/Remove Programs option of Windows NT, 2000, XP, or Server 2003. Also, the user can run the SETUP command initially used to install the program. When Setup is started and it senses a valid installation, an option is provided to uninstall. If proper prompts are followed, DDS or DDS/CSF may be uninstalled in that manner.

Unattended Uninstall

The uninstall command uses standard Windows protocol to remove this program in unattended mode. The user must create a file, “\winnt\setup.iss” to use in running an unattended uninstall. This file is created while running an attended uninstall using the SETUP command. Setup must be run as follows:

```
Setup /r
```

This command causes a recording to be made of user choices made during the uninstallation sequence. The file, “\winnt\setup.iss”, is saved for use later in doing an unattended uninstll. The user should be aware that all of the choices made during the uninstall will be recorded in a script including the choice of rebooting the machine following the uninstallation.

When the setup command is run later on in unattended mode, the command is run similar to what follows:

```
SETUP /S /F1[path]setup.iss /F2[path]uninstall.log
```

Note: In this example, “setup.iss” contains the Windows uninstallation script and “uninstall.log” contains the log created when the uninstall is run. [path] is simply the fully-qualified path to either of these files.

Troubleshooting Uninstall Errors

Check the log information created during the uninstallation process to determine if the uninstallation was successful and, if not, why not.

Installing and Removing Corrective Service

To apply or remove corrective service, refer to the README.FDS file on the Install diskette for the specific version of corrective service you are installing. Installation details may vary slightly from version to version but in all cases the process begins with the Install diskette and the setup command.

Chapter 5. Configuration Data

Each DDS workstation has its own configuration data, which is stored in two files on the local hard drive.

Logical names file (FDSL.N.LN)

The logical names file contains all of the logical names that you have defined for the DDS system. See “Logical Names and Role Names” on page 6 and refer to Name Services in **QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide** for more information about logical names.

Configuration file (FDSCONFIG.CFG)

The configuration file contains the remainder of the DDS configuration data. See

Chapter 7. Configuration Keywords for more information about the keywords that are located in the configuration file.

DDS classifies configuration data into two categories:

Basic

Affects most aspects of DDS operation and should be considered by every user.

Advanced

Affects specific aspects of DDS operation. It does not replace the basic data, but allows you to further control DDS.

Managing Configuration Data

To manage configuration data for the workstations in your DDS system, you can use tools and techniques supplied by DDS or provide your own. In fact, with the exception of the node ID and the system ID, you can choose not to manage configuration data at all. DDS uses defaults for any configuration item you do not specify.

Note: DDS provides a default configuration file that contains comment lines. Even if you use the default configuration values supplied with DDS, this file is required for DDS to initialize correctly.

DDS provides these tools and techniques for managing your configuration data:

- When you install the base component of DDS on a workstation for the first time, you also have the option to configure the workstation. If you are performing an attended installation, a series of configuration panels are provided for you at each workstation.
- The first time you perform an unattended installation, the configuration program (FDSCRFP) is automatically invoked and configures the workstation based on the configuration response file you provide (or using defaults if no file is provided).
- After you have configured the workstation, you must use the **Configuration File Utility** to modify the configuration. You can use the Configuration Response File Utility either to perform an attended configuration or to create response files for use with the Configuration Response File Processor. See **Chapter 6. Configuration File Utility** for information about the utility.
- You can also use the Configuration Response File Processor to manage configuration data from a central location. Then you can use a software

distribution manager to configure DDS on the workstation. You can also use tools or techniques not supplied with DDS to manage configuration data. The DDS configuration and logical names files are standard ASCII files.

Therefore, you can:

- Write a program that generates the configuration data file automatically.
- Use any standard ASCII editor to maintain the configuration data file.

Changing Configuration Data

Configuration data is easily changed to reflect any modifications in your store operation. When DDS is started, the local configuration data is read from `FDSCONFG.CFG`, which is located in the directory path `target_install_directory\CONFIG`. All configuration values are verified, and the configuration data is stored in memory.

DDS uses the configuration data in memory until DDS is stopped. Therefore, if you change configuration data, you must restart DDS for the changes to become effective. If you change the controlled drives for the workstation, or change whether data distribution is active, you must reboot the workstation for the changes to become effective.

Some configuration data, such as changes to the controlled drives used on the workstation or whether data distribution is used, requires the workstation to be rebooted before they become effective.

Activating Configuration Data

DDS uses the following guidelines for activating configuration data:

- If the configuration data file `target_install_directory\CONFIG\FDSCONFG.CFG` is not found, the initialization will fail. If you want to use all the default configuration values, you can create this file with no keywords in it. The default configuration file that is shipped with DDS contains a configuration data file with a comment line. However, comment lines are not required.
- Initialization fails if keywords or keyword values are detected that are not valid.
- Initialization fails if inconsistencies between keywords are detected. For example, IPC must be enabled if you intend to activate data distribution within the system. Therefore, specifying **RemoteIPC=No** and **DDActive=Yes** will cause initialization to fail.

- For configuration keywords that are not specified in the configuration file, DDS uses a default value and continues initialization.
- If the node ID or system ID is missing or not valid, initialization fails. The node ID and system ID are set during installation and validated during initialization. You cannot set these values via configuration keywords. See **Chapter 1. System Overview** and

Chapter 3. Installation for information about these values and how to set them.

- If a keyword is specified more than once, the value of the last occurrence is used unless multiple occurrences of the same keyword are specifically allowed.
- If no errors are encountered while processing the configuration data file, DDS writes out all the configuration keywords and values to a file that has the same format as a configuration file. The name of this file is `target_install_directory\CONFIG\FDSUSED.CFG`. It contains the list of values for all basic and advanced configuration keywords used for initialization. Default values, as well as explicitly set values, are written to this file. The file is overwritten each time that DDS is initialized.

Chapter 6. Configuration File Utility

DDS provides a window-driven utility to create or change your configuration files and logical names files. You can run the **Configuration File Utility** on your local workstation. The utility provides panels that prompt you for the desired configuration data. Default values are provided for every field.

When using the Configuration File Utility, you also have the option to specify an existing file (either a response file or a configuration data file) that should be used to provide the default values. When you specify this file, the installation and configuration values in the file are automatically supplied on the panels of the Configuration File Utility. You can also specify an existing logical names file to provide initial settings for your logical names.

Each window within the **Configuration File Utility** has an option for restoring all fields on the window to their default values. The default values are stored within DDS. You do not have to specify them.

You can exit the **Configuration File Utility** at any time. Upon exiting, you are asked whether you want to save or discard the changes you made. If you choose to save your changes, the file you specified as your output file is written with the values from your recent **Configuration File Utility** session. If you choose to discard your changes, none of the work from your current **Configuration File Utility** session is written to disk.

Starting the Configuration File Utility

You can start the utility from a command line or from the desktop:

- To start the utility from a command line, type FDSCFGUT and press Enter.
- To start the utility from the desktop, click Start on the task bar, point to Programs, point to QVS DDS or QVS 4690 Controller Services, and then click the Configuration icon.




Figure 6. **Configuration File Utility** -- Main window Generating Configuration Files

The main window of the **Configuration File Utility** is displayed with three menu options:

File

From the File menu option, you can choose one of the following actions: **Note:** You cannot select this menu option until you have initiated actions using the Generate menu option.

Edit

Allows you to work with a file. After you select the kind of file you want to work with (using the Generate menu option), select Edit to work with the file.

Save

Saves a file.

Save As...

Saves a file under a different name or allows you to provide a name for a newly generated file.

Generate

You can choose one of these actions from the Generate menu option:

Generate Configuration File

Create a new configuration file or update an existing file. See Generating Configuration Files on page 66 for detailed information.

Generate Logical Names File

Create a new logical names file or update an existing file. See **Generating a New Logical Names File** below for detailed information.

Help

Provides access to the online help for the Configuration File Utility.

Generating Configuration Files

You can load an existing configuration file or generate a new configuration file. To load an existing configuration file, see **Loading a Configuration File**. To generate a new configuration file, see **Generating a New Configuration File**.

Loading a Configuration File

To work with an existing configuration file:

1. Select **Generate** from the **Configuration File Utility** menu bar.
2. Select **Configuration File**.
3. Select **Load Configuration File**.
4. Type the path and name of the configuration file you want to work with, or select a drive, directory, and name from the list. The file you specify must already exist as a valid DDS configuration file. It should have a file type of “cfg”. The standard name is “fdsconfig.cfg”.
5. Select **Accept**.
6. The **Configuration File Utility** menu is displayed. Select **File** from the menu bar.
7. Select **Edit**. You can choose to work with basic configuration data, advanced configuration data, or both. If you select both, the basic configuration panels are displayed first, followed by the advanced configuration panels.

To work with basic configuration data, see **Generating a New Configuration File**.
To work with advanced configuration data, see **Working with Advanced Configuration Data** below.

Generating a New Configuration File

To create a new configuration file:

1. Select **Generate** from the **Configuration File Utility** menu bar.
2. Select **Configuration File**.
3. Select **New Configuration File**.
4. The **Configuration File Utility** menu is displayed. Select **File** from the menu bar.
5. Select **Edit**. You can choose to work with basic configuration data, advanced configuration data, or both. If you select both, the basic configuration panels are displayed first, followed by the advanced configuration panels.

Working with Basic Configuration Data

After you choose basic configuration data, the Basic Configuration Data window is displayed.

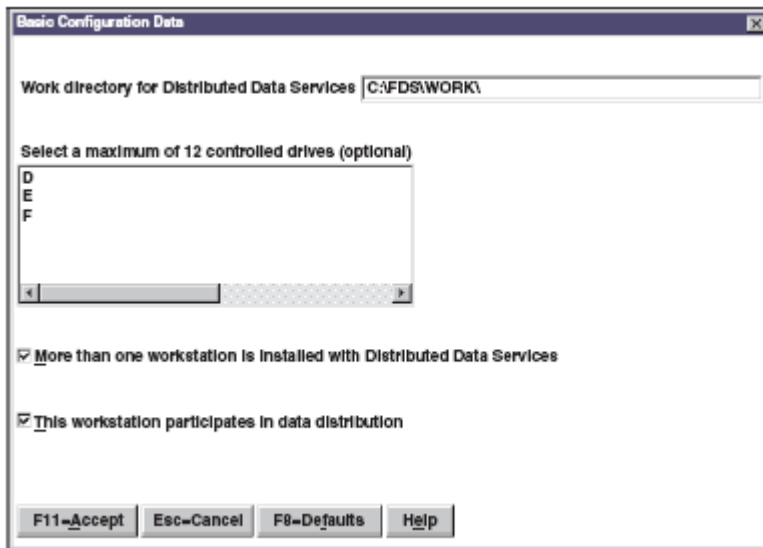


Figure 7. Basic Configuration Data

To complete the basic configuration data:

1. Type the complete path specification for the work directory you want to use for working disk space during normal operations. The default work directory is target_install_directory\WORK.

The directory you specify must be a local hard disk with at least 40 MB of available space. This directory should not be placed on a controlled drive. **Note:** This directory should be used only for working space. Do not work with or delete files from this directory. When you uninstall DDS, all files in this directory are deleted, unless you use the fdsunist command and specify the /K parameter.

2. Select the controlled drives for the DDS system, if you are using them. You can specify up to 12 controlled drives. The default is to have no controlled drives. However, if you select This workstation participates in data distribution, you must select at least one controlled drive.

Each drive you specify must be a local hard drive. However, you cannot specify as a controlled drive either the boot drive or the drive used to store the operating system paging file. To determine the drive where the paging file is located:

- a. Click **Start** from the task bar.
- b. Point to **Settings**.
- c. Select **Control Panel**.
- d. Select **System**.

- e. Take one of the following actions:
 - Select the **Performance** tab on Windows NT.
 - Select the **Advanced** tab and Performance Options on Windows 2000.
 - f. Select **Change**.
3. If you intend your local workstation to use the data distribution function as the primary distributor, the backup distributor, or a subordinate workstation, put a check in the box for This workstation participates in data distribution.

If this workstation does not participate in data distribution, proceed to step 13 in the next section, **Working with Advanced Configuration**.

4. The Distribution Role dialog box is displayed. Select the distribution role of this workstation and select **Accept**. The distribution roles are:

PrimaryDistributor

Controls the prime copy of all distributed files. Only the prime copy of a file can be modified directly by an application.

BackupDistributor

Controls the backup copy of a file and can take over for the primary distributor if the primary distributor fails or is deactivated.

Subordinate

Manages image copies of files.

Note: This option (subordinate) is NOT supported if you install the DDS/CSF version of the product.

5. If the configured distribution role of this workstation is either PrimaryDistributor or BackupDistributor, the Activation of Primary Distributor dialog box is displayed.
 - Specify an IP address if you want a specific IP address to be added to the node that is activated as the primary distributor.
 - Specify a subnet mask that will be associated with the **PrimaryIPAddress**.
 - Specify the network adapter number for which the **PrimaryIPAddress** will be added when this workstation is activated as the primary distributor.
 - Specify a computer name if you want a specific NetBIOS name to be added to the node that is activated as the primary distributor.
 - To indicate whether to enable the automatic switch-over feature, put a check in the box for Enable automatic switch-over of the primary distributor.

Note: All the parameters in this panel except **PrimaryIPAdapter** must be configured the same for both the configured primary distributor and the configured backup distributor.

Select **Accept**. If you enabled **Enable automatic switch-over** of the primary distributor, the **Automatic Switch-Over Configuration** dialog box is displayed.

Otherwise, skip the next step.

6. Specify the time in minutes to wait before automatically activating the acting backup distributor after it has lost communication with the acting primary distributor. Check the box for **Force automatic switch-over when the acting backup distributor is not fully reconciled** if you want to force the activation when the backup distributor is not fully reconciled and risk the possible loss of data.
7. If more than one workstation has DDS installed, select the box for **More than one workstation is installed with DDS**. Otherwise, go to step 13 in the **Working with Advanced Configuration Data** section below. If you select **More than one workstation is installed with DDS**, the **Maximum Number of Clients** dialog box appears when you select **Accept**. Enter the total number of clients attached to this workstation if this workstation is a file server.
8. Select **Accept**. The **Configuration File Utility** main window is displayed.
9. To save the file:
 - a. Select **File** from the menu bar.
 - b. Select **Save or Save As...** and specify the path and name for your configuration file, or select a name from the list.
 - c. Select **Accept**. The file is saved and the main window is displayed.

Working with Advanced Configuration Data

If you choose to work with advanced configuration data, the Advanced Configuration window, illustrated in **Figure 8** is displayed. The fields that are displayed on the Advanced Configuration Data window depend on the values you specified when you worked with basic configuration data. All possible fields are discussed in this section. If the field being discussed is not displayed on your window, skip the instruction for that field. For demonstration purposes, **Figure 8** illustrates a sample window.

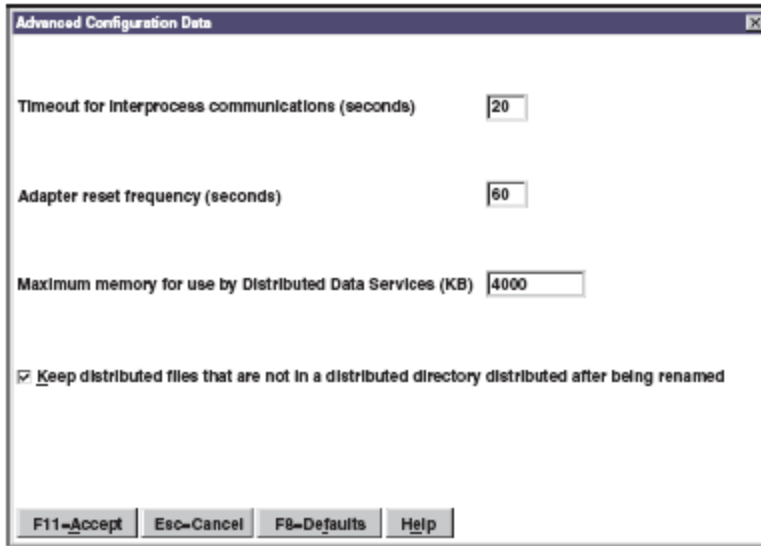


Figure 8. Advanced Configuration Data

1. Specify the time in seconds that the IPC component should wait before notifying your application that a receiving workstation is unavailable.

The IPC component manages communications between workstations in the DDS system on your LAN. When the IPC component transmits a point-to-point message on behalf of your application, it waits the specified amount of time to receive an IPC confirmation from the receiving workstation. If the confirmation is not received, the IPC component returns an error code to your application indicating that the receiving workstation is not available. (See **IPCTimeout** for detailed information.)

2. Specify the time in seconds that you want DDS to wait before attempting to reset the adapter card.

Certain physical occurrences can disrupt the IPC component from your workstation; for example, a disconnected cable from the adapter card used on a hardwired LAN media. When a LAN outage is detected, DDS attempts to reset the adapter after a short wait. If the first reset is unsuccessful, DDS waits the specified amount of time and retries. The DDS continues this retry process until communications are restored or until the DDS system is stopped. See **AdptrResetValue** for detailed information.

3. Specify the maximum amount of memory in kilobytes that you want DDS to use. The default value that you see is based on the distribution role you selected. If you allocate too little memory, DDS could slow down or could fail to initialize. See **MaximumMemory** for detailed information.

4. Indicate whether distributed files (that are not within a distributed directory) should remain distributed if they are renamed. Leave the check in the box if you want to have distributed files that are not within a distributed directory remain distributed if they are renamed. Remove the check in the box if you want the files to become local files (they will no longer be distributed) if they are renamed. See **DistRenamedFile** for detailed information.
5. Indicate whether data distribution exists in your DDS system. Check the box for Data distribution exists in the system if data distribution is active anywhere in your system.
6. Select **Accept**. The Advanced Configuration Data for IPC window is displayed, as illustrated in **Figure 9**.

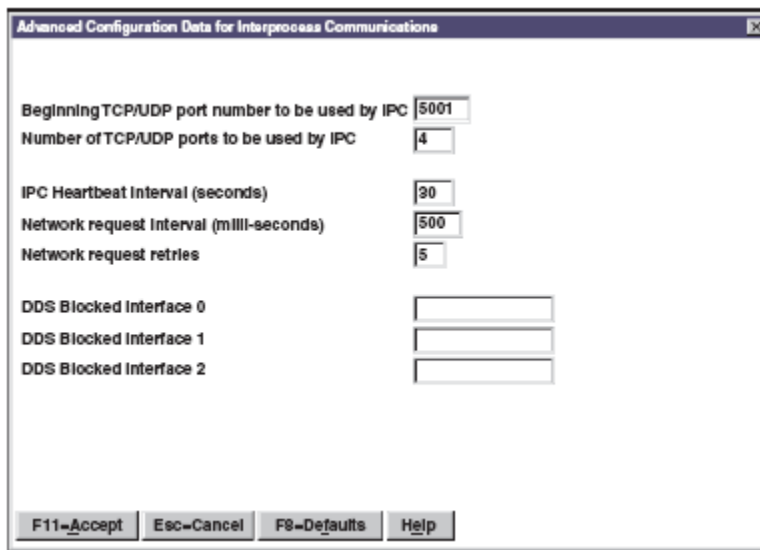


Figure 9. Advanced Configuration Data for Interprocess Communications Window

7. Enter the beginning TCP/UDP port number to be used by IPC.
8. Enter the number of TCP/UDP ports to be used by IPC.
9. Enter the IPC heartbeat interval, in seconds.
10. Enter the network request interval, in milliseconds.
11. Enter the number of network request retries.

12. If you have multiple communications adapters in your workstation, you can specify that DDS be blocked from using up to three interfaces (DDS Blocked Interface 1 through DDS Blocked Interface 3).
13. Select **Accept**. The **Configuration File Utility** main window is displayed.
14. To save the file:
 - a. Select File from the menu bar.
 - b. Select **Save or Save As...** and specify the path and name for your configuration file, or select a name from the list.
 - c. Select **Accept**. The file is saved and the main window is displayed.

Generating Logical Names Files

You can generate a new logical names file or load an existing logical names file. See **Generating a New Logical Names File** for information about generating a new configuration file. See **Loading a Logical Names File** for information about working with an existing logical names file.

See **Logical Names and Role Names** for more information about logical names.

Generating a New Logical Names File

To generate a new logical names file:

1. Select Generate from the **Configuration File Utility** menu bar.
2. Select Logical Names File.
3. Select New Logical Names File.
4. From the **Configuration File Utility** menu bar, select File.
5. Select Edit. The Logical Names window illustrated in Figure21 is displayed.

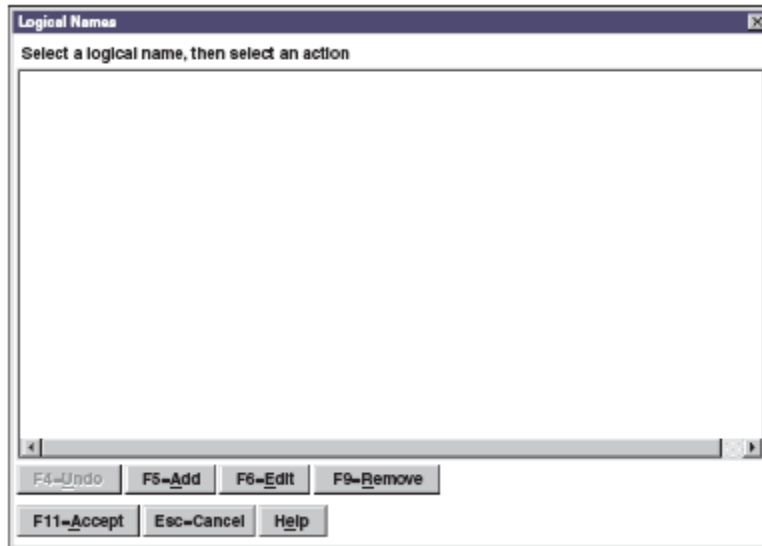


Figure 10. Logical Names

6. Select Add to create logical names. A dialog box is displayed.
7. Type the logical name you want to add, as well as the resolved name for the logical name.
8. Select **Accept**. The Logical Names window refreshes with the new logical name added to the list.

To reverse an addition, select Undo.

You can repeat this step as many times as needed.

You can also edit and remove logical names you have added. See **Editing Logical Names** and **Removing Logical Names** for more information.

9. When you have added all the names you want, select **Accept**. The **Configuration File Utility** main window is displayed.
10. To save the new file:
 - a. Select File from the **Configuration File Utility** menu bar.
 - b. Select **Save or Save As...** and specify the path and name for your logical names file, or select a name from the list. The standard filename for this file is FDSL.N.LN.
 - c. Select **Accept**. The logical names file is saved and the main window is displayed.

Loading a Logical Names File

To work with a previously existing logical names file:

1. Select Generate from the **Configuration File Utility** menu bar.
2. Select Logical Names File.
3. Select Load Logical Names File.
4. Type the path and name of the logical names file you want to load, or select a drive, directory, and name from the list. The file you specify must already exist as a valid DDS logical names file (normally FDSL.N.LN).
5. Select **Accept**. The **Configuration File Utility** main window is displayed.
6. Select File from the **Configuration File Utility** menu option.
7. Select Edit.
8. The Logical Names window, illustrated in Figure 10, is displayed. From this window, you can:
 - Add a logical name
 - Edit a logical name
 - Remove a logical name

Adding Logical Names

To add a logical name:

1. For each name you want to add:
 - a. Select Add. A dialog box is displayed.
 - b. Type the logical name you want to add, as well as the resolved name for the logical name, if requested.
 - c. Select **Accept**. The Logical Names window is displayed with the new logical name added to the list.

To reverse the addition, select Undo.

Note: If you are installing the DDS/CSF version of this product, be sure to define at least one logical name. The name should be “adxcsfh0”. Use lower case only. The value should be the drive name on which the 4690 Controller code resides (normally “d:”).

2. After adding all the names you want, select **Accept**. The **Configuration File Utility** main window is displayed.
3. To save the changes you have made:
 - a. Select File from the **Configuration File Utility** menu bar.
 - b. Select **Save or Save As....**
 - c. Specify the path and name for your logical names file, or select a name from the list.
 - d. Select **Accept**. The logical names file is saved and the **Configuration File Utility** main window is displayed.

Editing Logical Names

To edit a logical name:

1. For each name you want to edit:
 - a. Select the logical name you want to edit.
 - b. Select Edit. A dialog box is displayed with the current logical name and resolved name, if applicable.
 - c. Type a new logical name or resolved name over the current logical name and resolved name, if applicable.
 - d. Select **Accept**. The Logical Names window is displayed with the new logical name displayed in place of the original name and the new resolved name in place of the original resolved name, depending on the changes you made.

To reverse the change, select Undo.

2. After changing all the names you want, select **Accept**. The **Configuration File Utility** main window is displayed.
3. To save your changed logical names:
 - a. Select File from the **Configuration File Utility** menu bar.
 - b. Select **Save or Save As...** and specify the path and name for your logical names file, or select a name from the list.

- c. Select **Accept**. The logical names file is saved and the main window is displayed.

Removing Logical Names

To remove a logical name:

1. For each name you want to remove:
 - a. Select a logical name you want to remove from the list.
 - b. Select **Remove**. The logical name is removed from the list. Select **Undo** to reverse the removal.
2. After removing all the names you want to, select **Accept**. The **Configuration File Utility** main window is displayed.
3. To save the changes you have made:
 - a. Select **File** from the **Configuration File Utility** menu bar.
 - b. Select **Save or Save As...** and specify the path and name for your logical names file, or select a name from the list.
 - c. Select **Accept**. The logical names file is saved and the **Configuration File Utility** main window is displayed.

Chapter 7. Configuration Keywords

This section contains a list of the configuration keywords that can be used in the configuration data file and configuration response files. For information about how to obtain configuration data from workstations using the DDS APIs, refer to FdsQueryConfig() API in the **QVS Distributed Data Services/Controller Services Facility for Windows Programming Guide**.

You can also determine the setting of these values by running the Configuration File Utility, or by editing target_install_directory\CONFIG\FDSUSED.CFG after DDS has initialized. See **Activating Configuration Data** for more information about FDSUSED.CFG. **Chapter 6. Configuration File Utility** for more information about the Configuration File Utility.

Not all keywords will apply to the type of workstation you are configuring. For example, if you are configuring a stand-alone workstation, you do not require remote IPC, so any keywords dealing with communications between workstations do not apply.

Other keywords are only valid based on what you are specifying. For example, the keyword related to restarting DDS on a workstation after its configuration has been changed (**RestartFDS**) is valid only when generating response files. The following information is provided for each keyword:

Unit of Measure

Qualifies the value you provide. For example, it can be seconds of time, KB of memory, or adapter card number. Some keywords do not require a **Unit of Measure**.

Valid Values

Specifies the upper and lower bounds for the specific keyword. These bounds are verified when DDS is started.

Unless specifically stated, values are not case-sensitive.

Default

Provides the default value if one exists. This information is provided only for basic keywords and response-file only keywords. Default values for advanced configuration keywords are based on values provided for basic configuration keywords.

Notes

Further elaborates on specific aspects of the configuration data. This section helps you to determine whether you need to change the value, or gives you further information about how you can use the configuration value to affect operation of DDS.

Basic Configuration Keywords

These keywords are valid in both response files (installation and configuration) and configuration files.

AutoSwitchOver

If your workstation is configured with a **DistributionRole** of PrimaryDistributor or BackupDistributor, this configuration value is used to specify whether the acting backup distributor should automatically be activated as the primary distributor when the following conditions have been met:

- The acting backup distributor has lost communication with the acting primary distributor for the length of time specified by the **AutoSwitchOverDelay** keyword.
- DDS is not in the process of initializing on the acting backup distributor. That is, DDS is not restarting or rebooting.
- Either the acting backup distributor has determined that it is fully reconciled (all of its distribute-on-update files are an accurate image of those on the acting primary distributor) or the **AutoSwitchOverForce** keyword is set to Yes.
- The acting backup distributor can accurately determine if it is connected to the LAN and is able to successfully communicate with the LAN.
- The acting backup distributor has communicated with an acting primary distributor at least once since DDS was started.

Unit of Measure

Not applicable

Valid Values

Yes

Whenever the conditions for an automatic switch-over are met, the acting backup distributor will be activated as the acting primary distributor.

No

User intervention is required to activate the acting backup distributor as the primary distributor. See **IBM Distributed Data Services/Controller Services Facility for Windows User's Guide** for information about activating the primary using FDSADMIN or FDSNCU Node Control Utilities and **IBM Distributed**

Data Services/Controller Services Facility for Windows Programming Guide
for information about activating the primary using FdsActivateAsPrimary() API.

Default

No

Notes

1. When DDS is configured for automatic switch-over, automatic deactivation is also enabled. Automatic deactivation entails the deactivation of the acting primary distributor whenever both of the following conditions are met:
 - a. The acting primary distributor can accurately determine when it is disconnected from the LAN or when it is unable to connect to the LAN.
 - b. The acting primary distributor has communicated with an acting backup distributor since the acting primary distributor completed initializing.
2. To set the value of this keyword to Yes, installation of the Automatic Switch-Over component of DDS is required.
3. See the **AutoSwitchOverDelay** and **AutoSwitchOverForce** keywords and **Automatic Primary Distributor Switch-Over** or more information.
4. The configured primary distributor and the configured backup distributor must have the same value for this keyword.

AutoSwitchOverDelay

This value specifies the amount of time the acting backup distributor tries to reestablish communication with the acting primary distributor before automatically activating as the acting primary distributor. The timer used for this delay is started after all other conditions for automatic switch-over have been met.

Unit of Measure

Minutes

Valid Values

0 - 1440

Default

5

Notes

1. This keyword is ignored if the value of AutoSwitch-Over is set to **No**

2. See the **AutoSwitchOverDelay** and **AutoSwitchOverForce** keywords and **Automatic Primary Distributor Switch-Over** for more information.

AutoSwitchOverForce

Indicates whether the acting backup distributor will be activated as the primary distributor when all conditions for automatic switch-over have been met, excluding determination of whether the acting backup distributor is fully reconciled.

Unit of Measure

Not applicable

Valid Values

Yes

The acting backup distributor will be activated as the primary distributor whenever all of the conditions for an automatic switch-over are met, excluding the determination of whether the acting backup distributor is fully reconciled.

No

The acting backup distributor will never be automatically activated as the primary distributor if the acting backup distributor has determined that it is not fully reconciled.

Default

No

Notes

1. This keyword is ignored if the value of **AutoSwitchOver** is set to **No**
2. See the **AutoSwitchOver** and **AutoSwitchOverDelay** keywords and **Automatic Primary Distributor Switch-Over** for more information.
3. The user should carefully consider the side effects of setting this value to **Yes**. Because all updates to distributed files are initially made on the acting primary distributor and then sent to the backup distributor, activating the backup distributor when it is not fully reconciled means that any updates that have not been received and applied by the backup will be lost. The backup distributor could have mismatched, additional or missing files when compared to the primary.

ControlledDrives

This configuration value is used to specify which drives you want DDS to control. DDS performs disk I/O prioritization and distributes files only for disks that you specify here. If you do not want a drive partition to be controlled by DDS, do not include that partition in this configuration value. See **Designating Controlled Drives** for more information.

Unit of Measure

Not applicable

Valid Values

A string of drive letters. Each drive letter must be a single alphabetic character C through Z. Duplicate values are ignored.

Default

None

Notes

1. If you do not specify a value for this keyword, it means that you do not want any drives to be controlled.
2. To have DDS control your D, F, J, and K drive partitions, specify **ControlledDrives=DFJK**.
3. You can specify up to 12 drive partitions to be controlled.
4. You cannot specify that the boot drive be controlled by DDS. The boot drive is verified when you initialize DDS.
5. The drive used to store the operating system paging file cannot be controlled. To determine where the paging file is located:
 - a. Click Start from the task bar.
 - b. Point to Settings.
 - c. Select Control Panel.
 - d. Select System.
 - e. Take one of the following actions:
 - Select the Performance tab on Windows NT.

- Select the Advanced tab and Performance Options on Windows 2000.

f. Select Change.

This drive is verified when you initialize DDS.

6. Each drive specified must be a local hard drive.
7. For this configuration keyword to take effect, you must reboot the local workstation. See **Reboot** for more information.

DDActive

This configuration value is used to indicate whether your local workstation is enabled for the Data Distribution component.

Unit of Measure

Not applicable

Valid Values

Yes

Specifies that the local workstation is enabled for data distribution. Specify this value if you intend for your local workstation to use the Data Distribution function as the primary distributor, the backup distributor, or a subordinate workstation. **No**

Specifies that the local workstation is not enabled for data distribution. This selection reduces use of system resources.

Default

No

Notes

1. For more information about the Data Distribution component, see **Chapter 1. System Overview** and the Data Distribution Utility chapter in the **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide**. For this configuration keyword to take effect, you must reboot the local workstation. See **Reboot** for more information.
2. At least one drive must be specified as a controlled drive if you specify **DDActive=Yes**.

DistributionRole

This value configures the workstation with a role. This configuration value is used to set the configured primary distributor, configured backup distributor, and subordinates.

The primary distributor controls the prime copy of all distributed files. Only the prime copy of a file can be modified directly by an application. The backup distributor controls the backup copy of all distribution files and can take over for the primary distributor if the primary distributor fails or is deactivated.

All other workstations that are members of a distribution domain are considered subordinates. They manage image copies of the files.

Unit of Measure

Not applicable

Valid Values

PrimaryDistributor, BackupDistributor, Subordinate, or None

Default

If **DDActive** is set to No, the default for this keyword is None. If **DDActive** is set to Yes, the default for this keyword is Subordinate.

Note

This keyword applies only if this workstation participates in data distribution (if you specified **DDActive=Yes**). If this workstation does not participate in data distribution, specify **DistributionRole=None**.

IPCTransport

This keyword specifies the transport layer protocol to be used on this workstation by DDS IPC.

Unit of Measure

Not applicable

Valid Values

TCPIP

Default

TCPIP

Note

This keyword applies only if more than one workstation is installed with DDS (if you specified **RemoteIPC=Yes**).

MaxRequestors

If this workstation is a file server, specify the maximum number of requestors (workstations) that this workstation will serve. If this workstation is not a file server, specify 0 (zero).

Unit of Measure

Not applicable

Valid Values

0-65535

Default

0 for subordinate workstation

20 for primary or backup distributor

PrimaryComputerName

Many application programs are designed to connect (NET USE) to a file server using a specific computer name (NetBIOS name) for the purpose of reading and writing to files. If DDS is installed on that server and distributing file updates made by the application, it is important that the application is always connected to the DDS workstation that is performing the primary distributor role. The DDS workstation is the only workstation on which distributed files can be modified.

Using this keyword, a unique computer name can be specified that will be added to a workstation when it becomes the acting primary distributor. The name will not be associated with a single workstation, but will be added to whichever workstation is performing the role of acting primary distributor. Therefore, after a workstation has been activated as the primary distributor, it will have two active NetBIOS names.

The computer name specified must be different from the computer names configured for any other workstations through Windows installation or network configuration.

If this keyword has a value, it indicates that the user wants to add the computer name to the acting primary distributor during either a user-initiated or automatic activation of the primary distributor and to remove it during deactivation.

Unit of Measure

Not applicable

Valid Values

Valid NetBIOS name

Default

No NetBIOS name

Notes

1. The computer name can be up to 15 characters long with no blank spaces. It can contain only alphanumeric characters plus the following special characters: ! @ # \$ % & () - _ ' { } . The names of these characters are exclamation mark, at or each symbol, number sign, dollar sign, percent sign, and sign, ampersand, right and left parenthesis, hyphen or dash, underscore, apostrophe, left and right curly braces, and period.
2. The computer name specified must be different from the computer name configured on this or any other workstation attached to your network.
3. To specify a transferable computer name for the acting primary distributor requires the installation of the Automatic Switch-Over component of DDS.
4. The configured primary distributor and the configured backup distributor must have the same value for this keyword.

PrimaryIPAddress

Many application programs are designed to connect to a file server using a specific IP address for the purpose of reading and writing to files. If DDS is installed on that server and distributing file updates made by the application, it is important that the application always connect to whichever DDS workstation is the acting primary distributor. This workstation is the only workstation on which distributed files can be modified. Using this keyword, a unique IP address can be specified that will be added to a workstation when it becomes the acting primary distributor. This address will not be associated with a single workstation, but will be added to whichever workstation is performing the role of acting primary distributor. After a workstation has been activated as the primary distributor, it will have two active IP addresses.

The IP address specified must be different from any address configured for your other workstations with Windows NT, 2000, XP, or Server 2003 network configuration. If this keyword has a value, it indicates that the user wants to add the IP address to the new acting primary distributor during either a user- initiated or automatic activation of the primary distributor and to remove it during deactivation.

Unit of Measure

Not applicable

Valid Values

Valid TCP/IP address that is not already configured on this or any other workstation

Default

No IP address

Notes

1. The IP address specified must be different from the IP address configured on this or any other workstation attached to your network.
2. To specify a transferable IP address for the acting primary distributor requires the installation of the Automatic Switch-Over component of DDS.
3. The configured primary distributor and the configured backup distributor must have the same value for this keyword.
4. If you specify a **PrimaryIPAddress**, you must also specify a **PrimarySubnetMask**.
5. See **PrimaryIPAdapter** below for more information.

PrimaryIPAdapter

This keyword is used to specify the network adapter on which the **PrimaryIPAddress** will be added. The procedures for determining the keyword value are different for Windows NT and Windows 2000.

Windows NT

On Windows NT, the adapter number can be found using the network configuration program in this manner:

1. Click the Start button on the task bar.
2. Choose Settings.
3. Click Control Panel.
4. Double-click Network.
5. Click the Bindings tab.
6. Select all protocols from the Show Bindings for: list box.
7. Click the '+' next to the TCP/IP Protocol entry.

If you want the **PrimaryIPAddress** to be added to the first adapter in the list, specify 1 for this keyword, or for the second adapter specify 2, and so on. Any adapters that are disabled for TCP/IP, as indicated by a circle with a line through it to the left of the

adapter name, should be ignored. For example, if there are two adapters listed, but the TCP/IP protocol has been disabled on the first adapter in the list, then the second adapter in the list is adapter 1 for purposes of specifying an adapter number for the **PrimaryIPAdapter** keyword.

Windows 2000

On Windows 2000, the command line utility fdsipcad.exe, which is shipped as part of the Controller Services base product, must be executed to determine the adapter number.

1. Run fdsipcad.exe.
2. Specify the **PrimaryIPAdapter** value shown for the adapter you want to associate with the **PrimaryIPAddress**.

Note: On Windows 2000, this value may be any non-zero integer.

Unit of Measure

Not applicable

Valid Values

Any non-zero integer

Type 1-Basic

Default

1

Notes

1. This keyword is ignored if the value of **PrimaryIPAddress** is not specified.
2. The configured primary distributor and the configured backup distributor need not have the same value for this keyword.

PrimarySubnetMask

The subnet mask for the **PrimaryIPAddress** is specified using this keyword.

Unit of Measure

Not applicable

Valid Values

Valid TCP/IP subnet mask

Default

No subnet mask

Notes

1. The configured primary distributor and the configured backup distributor must have the same value for this keyword.
2. The Automatic Switch-Over component of DDS must be installed before a value may be specified for this keyword.
3. If you specify a **PrimarySubnetMask**, you must specify a

RemoteIPC

The IPC component is a peer-to-peer messaging service that allows application programs to send and receive messages. The messaging service is provided between processes running on a single workstation and between those running on different workstations. This configuration option is used to specify whether inter-workstation IPC is enabled.

Unit of Measure

Not applicable

Valid Values

Yes or **No**

Default

Yes

Note

Specify Yes for all workstations in a system unless the system consists of a single, stand-alone workstation.

WorkDirectory

DDS uses temporary or working disk space during normal operations. For example, the Data Distribution component uses the disk space to keep status information on distributed files, and the File Services component uses the disk space to ensure that multiple write operations are processed correctly. Use this value to designate which local hard drive and directory should be used for this purpose.

Unit of Measure

Not applicable

Valid Values

A fully qualified path name including drive specification.

Default

target_install_directory\WORK

Notes

1. The directory must be specified on a local hard drive.
2. In general, this keyword should never be changed after DDS has successfully been started the first time. Should it be necessary to change it, the following steps must be followed:
 - a. Stop DDS.
 - b. Change this keyword to the desired value.
 - c. Create the subdirectory specified by the new value for this keyword.
 - d. Copy the entire contents (including all subdirectories) of the directory specified by the old value for this keyword to the directory specified by the new value for this keyword.
 - e. Start DDS.
 - f. After DDS has been successfully started, you can delete the contents of the old directory and remove the directory. g. The designated working directory will be created upon initialization of DDS if it does not exist.
3. If you specify **WorkDirectory**=D:\FDS\WORK, DDS uses the specified drive and directory path as its working directory.
4. The directory you designate should have at least 40 MB of available space and should not be used for any other purpose. It should not be placed on a controlled drive.
5. The data in this directory must be persistent across reboots. Therefore, you should not work with or delete files from this directory.
6. Important: All files in the directory you specified as the work directory will be deleted when you remove DDS.

Advanced Configuration Keywords

These keywords are valid in both response files (installation and configuration) and configuration files.

AdptrResetValue

Certain physical occurrences can disrupt interprocess communication; for example, someone disconnecting the cable from the adapter used on the hardwired LAN media. When a LAN outage is detected, DDS attempts to reset the adapter card after a specified period of time. Use this configuration value to specify how long DDS should wait before attempting to reset the adapter card. If the first reset is unsuccessful, DDS waits the specified amount of time and retries. The DDS continues this retry process until communications are restored or until the DDS system is stopped.

The user may notice the impact of resetting the adapter for several seconds.

Large values for this keyword will result in less frequent impact. Small values for this keyword allow communication with other workstations to be restored more quickly after the LAN outage is resolved.

Unit of Measure

Seconds

Valid Values

1600

Default

60

DDSBlockedInterface

In workstations that have multiple communications interfaces, this keyword can be used to stop DDS from using all of the interfaces. You can specify that DDS be blocked from using up to three interfaces. See **Blocking Interfaces** for general information about blocked interfaces.

Unit of Measure

Not applicable

Valid Values

0-2, Valid TCP/IP address

Default

All interfaces will be used by DDS.

Note

For each interface that you want to block DDS from using, specify this keyword followed by 0, 1, or 2 and a valid TCP/IP address. You must provide a TCP/IP address that currently exists in the workstation. For example, to specify that the TCP/IP address 9.88.154.134 be blocked from use by DDS and that this is the first address to be blocked, specify:

DDSBlockedInterface = 0,9.88.154.134

DistRenamedFile

Within DDS, there are two ways to make a file a distributed file. You can either set the file itself to be a distributed file or you can include the file in a directory that has been set to be distributed. If you include a file in a directory that is distributed, and then you rename the file, it will continue to be distributed, but with the new name.

When you set a file itself to be distributed (it is not located in a directory to be distributed), you can use the **DistRenamedFile** keyword to control whether or not the file remains a distributed file if it is renamed.

Unit of Measure

Not applicable

Valid Values

Yes

Any distributed file that is renamed remains distributed but with the new name. The file remains distributed regardless of whether or not the distributed file is in a distributed directory.

No

Any distributed files that are not in a distributed directory will become local files (they will no longer be distributed) if they are renamed.

Default

Yes

Notes

1. This keyword has no effect on files that are located in distributed directories. They will remain distributed as long as they are in a distributed directory, regardless of whether or not they are renamed.
2. This keyword has no effect on files that are renamed to a different drive.

IPCHeartbeatInterval

When DDS IPC uses TCP/IP as the transport layer protocol, it periodically checks every point-to-point connection established with other workstations to see if it has failed.

This keyword specifies how frequently DDS IPC should verify the communications connection with another node has failed. The value of this keyword specifies the time interval between checks.

Unit of Measure

Seconds

Valid Values

1240

Default

30 seconds

IPCPortCount

The number of TCP/UDP ports used by the IPC component when you are using TCP/IP as the transport layer protocol. The minimum value required for DDS to communicate with other workstations is 3. Increase this number by 1 for each DDS broadcast domain that includes this workstation. This keyword is used along with **IPCPortStart** to specify the range of well-known ports that DDS may use for TCP/IP communication.

Unit of Measure

Not applicable

Valid Values

3-254

Default

6

IPCPortStart

The first TCP/UDP port number to be used by the IPC when you are using TCP/IP as the transport layer protocol. This keyword is used along with **IPCPortCount** to specify the range of well-known port numbers that DDS may use for TCP/IP communication.

Unit of Measure

Not applicable

Valid Values

5001-65000

Default

5001

Note

The value of this keyword must be the same for all DDS workstations within a system.

IPCTimeout

The IPC component manages communications among workstations in the DDS system on your LAN. When the IPC component transmits a point-to-point message on behalf of your application, it waits a certain amount of time to receive an IPC confirmation from the receiving workstation. If the confirmation is not received, the IPC component returns an error code to your application indicating that the receiving workstation is not available.

Use this configuration data to specify how long the IPC component should wait before notifying your application that the receiving workstation is unavailable.

Unit of Measure

Seconds

Valid Values

5-600

Default

20

LocatePrimary

This keyword tells DDS that a primary distributor exists in the system. Specify Yes for this keyword if data distribution is active anywhere in your system. Specify No if data distribution is not active anywhere in the system. Data distribution must be active on at least one node unless the system consists of a single node.

Unit of Measure

Not applicable

Valid Values

Yes or No

MaximumMemory

Specify the maximum memory in kilobytes for Controller Services to use. **Valid Values** are in the range of 24000 to 2147483647.

Unit of Measure

1024 bytes (1K)

Valid Values

24000 - 2147483647

Type

Advanced

Notes

1. The default value for this keyword is 24000 KB. This value allows filesets of up to 10,000 files with 8.3 format file names to be handled. This memory is allocated from the Windows swap file and will not affect real memory usage on your system unless the DDS workload requires it. Distributed files with long file names should be counted as two files when determining the **MaximumMemory** value.
2. Increase the minimum value by at least 2000 KB for each additional 1000 distributed files in the system. For example, up to 10,000 distributed files would require at least 24000 KB, and 10,000-10,999 would require at least 26000 KB.

NetworkRequestInterval

The IPC component supports broadcast domains. Broadcast domains support communication among multiple workstations.

To manage this type of communication when using the TCP/IP transport layer protocol, DDS IPC ensures that all workstations that are part of a specific broadcast domain use the same TCP/IP port to represent that broadcast domain. When a broadcast domain is first used, the associated port must be determined. The workstation queries the LAN to obtain this information. This keyword specifies the time interval between queries used by the DDS IPC component when attempting to determine the TCP/IP port associated with a broadcast domain. This value is determined by the amount of congestion on the LAN, and is proportional to the level of network congestion.

Unit of Measure

Milliseconds

Valid Values

100 - 5000

Default

500 milliseconds

NetworkRequestRetries

The IPC component supports broadcast domains. Broadcast domains support communication with multiple workstations.

To manage this type of communication when using the TCP/IP transport layer protocol, IPC ensures that all workstations that are part of a specific broadcast domain use the same TCP/IP port to represent that broadcast domain.

When a broadcast domain is first used, the associated port must be determined. The workstation queries the LAN for this information.

This keyword specifies the maximum number of times the IPC component will query the LAN to obtain the TCP/IP port associated with a broadcast domain. This value is determined by the amount of congestion on the LAN, and is proportional to the level of network congestion.

Unit of Measure

Not applicable

Valid Values

Default

Response-File-Only Configuration Keywords

These keywords are valid only in a response file and are not used elsewhere.

BackupConfig

This configuration value is used to indicate that your current configuration file should be written to the specified file prior to applying any changes from the response file. The logical names file is not backed up. Current configuration data is in target_install_directory\FDS\CONFIG\FDSCONFIG.CFG.

The current configuration data will be copied to the specified file. Then the rest of the configuration values in the response file are processed. If the **ConfigUpdateType** keyword is not set to update the configuration files, the current configuration is backed up and the new configuration is the same as the old configuration.

Unit of Measure

Not applicable

Valid Values

Fully qualified file name

Notes

1. If you do not want a backup of your configuration data file, specify **BackupConfig=** in your response file, with no file name specified.
2. When reinstalling or reconfiguring DDS, you can use this configuration keyword to save the current configuration data before reinstalling. If you encounter any problems, you can use the **RestoreConfig** configuration keyword to return to your previous configuration.

ChangeLogicalName

A logical name is a name that indirectly refers to a physical entity. For example, the logical name PLU1 can refer to a specific price lookup file. The resolved name specifies the physical entity itself. The Name Services component provides the resolution services for your application program. If the logical name already exists in the configuration logical names file, the resolution of the logical name is specified with this configuration keyword. If the logical name does not exist in the configuration logical names file, this configuration data is used to establish a logical name as well as the resolved name for the logical name.

Unit of Measure

Not applicable

Valid Values

A logical-name/resolved-name pair

Notes

1. The logical name is case-sensitive.
2. You can specify as many logical names as you want.
3. Order is important for this keyword. If you specify **ChangeLogicalName** for a logical name and later specify **DeleteLogicalName** for the same logical name, the logical name will be deleted.
4. All logical names are stored in a single file that is used only for logical-name configuration data.
5. Logical names cannot be a null string and cannot start with the prefix FDS.
6. Logical names cannot exceed 260 characters.

7. Logical names cannot contain two colons (::) in a row.
8. The logical name must be delimited by the less-than and greater-than characters (< and >). These delimiters (< and >) cannot be imbedded within the logical name. The resolved name must immediately follow the logical name with no space between them. All characters, including spaces, that follow the greater than symbol are treated as part of the resolved name.
9. Resolved names cannot be a null string. The resolve string can have other logical names imbedded. If it does, the logical name must be delimited by the less-than and greater-than characters (< and >).

For example:

ChangeLogicalName=<adxcsfh0>d:

Note: In this example, the logical name, “adxcsfh0”, represents the value “d:”

ConfigUpdateType

This keyword is used to specify whether the current configuration data will be updated to create the new configuration files or whether the new configuration files will be created based entirely on the information in the response file.

Unit of Measure

Not applicable

Valid Values

NoUpdates

No changes will be made to your configuration file. This option allows you to update the installation files or logical names file without making any changes to your configuration file.

New

Your current configuration file will be deleted and a new configuration file will be created, based on the information in the response file.

Update

Your existing configuration file will be updated based on the information in the response file. If there is no existing configuration file, a warning message will be logged and the file will be created using the information from the response file.

DeleteLogicalName

A logical name is a name that indirectly refers to a physical entity. For example, the name PLU1 can refer to a specific price lookup file. The resolved name specifies the

physical entity itself. The Name Services component provides the resolution services for your application program.

This keyword is optional. Use it to delete a logical name that is already specified in the configuration logical-names file. If the logical name does not exist in the configuration logical-names file, this configuration keyword has no effect and a warning message is logged.

Unit of Measure

Not applicable

Valid Values

A logical name

Notes

1. This value is case-sensitive.
2. The logical name must be delimited by the less-than and greater-than characters (< and >).
3. The resolved name is not required and is ignored if present.
4. There is no verification to ensure that a deleted logical name is not embedded in other logical name resolutions.
5. The Name Services component fails if it tries to resolve a logical name that is not fully resolvable.

For example, to delete the logical name <plu1> from the configuration logical names file:

```
DeleteLogicalName=<plu1>
```

Include

This keyword is optional and specifies which general response files to include with a specific response file.

The physical placement of this keyword affects the resulting merged response file definition. For example, keywords from an included response file override duplicate keywords in the including response file if the physical placement of the Include keyword follows the occurrence of the original keywords. The included response file is handled in-line. If you specify an included response file but the response file is not found, no message is recorded in the error log or history log.

You cannot have more than five levels of included response files.

The following search order is used to find the specified general response files:

1. The fully-qualified file specification
2. The current directory

For example:

Include = z:\cid\rsp\genresp.rsp

LogicalNamesUpdateType

This keyword is used to specify whether the current logical-names file will be updated to create the new logical-names file, or whether a new or replacement logical-names file will be created based entirely on the information in the response file.

Unit of Measure

Not applicable

Valid Values

NoUpdates

No changes will be made to your logical-names file. This option allows you to update the installation files or the configuration file without making any changes to your logical names file.

New

Your current logical-names file will be deleted and a new logical-names file will be created, based on the information in the response file.

Update

Your existing logical-names file will be updated based on the information in the response file. If there is no existing logical-names file, a warning message is logged and the file is created.

Reboot

This configuration keyword is used to reboot the local workstation if there were any changes to the **ControlledDrives** or **DDActive** keywords. The Configuration Response File Processor does not actually initiate a reboot; it sends a return code to indicate that a reboot is required. See **Configuration Response File Processor (FDSCRFP)** for details.

Unit of Measure

Not applicable

Valid Values

Yes

If the values for the **ControlledDrives** or **DDActive** keywords have been changed, the response file is processed, DDS is stopped on the local workstation, and the local workstation is reinitialized to pick up the new set of controlled drives.

No

DDS is not stopped and the local workstation is not reinitialized.

RestartFDS

This configuration keyword is used to restart the DDS system to apply most of the configuration changes you have made using the response file. To apply changes to the **ControlledDrives** or **DDActive** keyword, you must reboot the local workstation. See **Reboot** for more information.

Unit of Measure

Not applicable

Valid Values

Yes

The response file is processed. Then the DDS system is shut down and restarted to pick up the new configuration data.

No

The new configuration files are created or updated, but DDS is not restarted to pick up these new configuration values.

Note

This keyword restarts DDS only if it was running. If it was not running prior to this reconfiguration, it is not started after the response file is processed.

RestoreConfig

This configuration value is used to indicate that your configuration file should be restored from the specified file. The logical names file is not restored. The configuration file is copied from the specified file to target_install_directory \CONFIG\FDSCONFIG.CFG. After the configuration file is restored, the configuration data in the response file is processed. If the **ConfigUpdateType** keyword is specified and is not set to Update, this keyword is ignored and the configuration file is not restored.

Unit of Measure

Not applicable

Valid Values

Fully-qualified file name

Note

You can use this keyword to restore a previous version of your configuration file.

Logical-Names-File-Only Configuration Keywords

This keyword is only valid in a logical names file.

LogicalName

A logical name is a name that indirectly refers to a physical entity. For example, the name PLU1 can refer to a specific price lookup file. The resolved name specifies the physical entity itself. The Name Services component provides the resolution services for your application program.

Use this configuration keyword to establish a logical name as well as the resolution name for the logical name.

Unit of Measure

Not applicable

Valid Values

A logical-name/resolved-name pair

Notes

1. This value is case-sensitive.
2. This keyword does not have a default.
3. You can specify as many logical names as you want.
4. All logical names are stored in a single file that is used only for logical name configuration data.
5. Logical names cannot be a null string and cannot start with the prefix FDS. 7. Logical names cannot exceed 260 characters.
6. Logical names cannot contain two colons (::) in a row.
7. The logical name must be delimited by the less-than and greater-than characters (< and >). These delimiters (< and >) cannot be imbedded within the logical name. The resolved name must immediately follow the logical name with no space between them. All characters, including spaces, that follow the greater than symbol are treated as part of the resolved name.

8. Resolved names cannot be a null string. The resolve string can have other logical names imbedded. If it does, the logical name must be delimited by the less-than and greater than characters (< and >).

For example:

LogicalName=<file1>d:\plu\file1

Configuration Keywords Summary

This table provides a summary of all configuration keywords, including a description, the type of keyword, and the page where the keyword is described.

Table 1. Configuration Keywords

Keyword	Description	Type	See Page
AdptrResetValue	Specifies the adapter reset frequency	Advanced	See page 84
AutoSwitchOver	Automatically activates the backup distributor	Basic	See page 72
AutoSwitchOverDelay	Specifies length of time for backup distributor to attempt to reestablish communication	Basic	See page 73
AutoSwitchOverForce	Indicates whether the acting backup distributor will take over as primary when applicable conditions are met	Basic	See page 74
BackupConfig	Backs up the configuration data to the specified file	Response File only	See page 89
ChangeLogicalName	ADDS or updates a logical name	Response File only	See page 90
ConfigUpdateType	Specifies the type of update for the configuration file	Response File only	See page 91
ControlledDrives	Specifies which drives are controlled (maximum of 12)	Basic	See page 75
DDActive	Indicates whether the workstation participates in data distribution	Basic	See page 76
DDSBlockedInterface	Specifies TCP/IP address that DDS should not use	Advanced	See page 84
DeleteLogicalName	Deletes a logical name	Response File only	See page 91
DistRenamedFile	Indicates whether distributed files (that are not within a distributed directory) should remain distributed if they are	Advanced	See page 85

	renamed. The default value is Yes.		
DistributionRole	Specifies the distribution role	Basic	See page 77
Include	Includes another group response file	Response File Only	See page 92
IPCHeartbeatInterval	Specifies how often DDS IPC should check to see if the communications connection with another node has failed	Advanced	See page 86
IPCPortCount	Specifies the number of TCP/UDP ports to be used by IPC when using TCP/IP as the transport layer protocol	Advanced	See page 86
IPCPortStart	Specifies the first TCP/UDP port number to be used by IPC when using TCP/IP as the transport layer protocol	Advanced	See page 86
IPCTimeout	Specifies the IPC timeout	Advanced	See page 87
IPCTransport	Specifies the transport layer protocol to be used by IPC	Basic	See page 77
LocatePrimary	Specifies whether data distribution is active anywhere in the system	Advanced	See page 87
LogicalName	Specifies the logical name and resolution for that logical name	Logical Names File only	See page 95
LogicalNamesUpdateType	Specifies the type of update for logical names file	Response File only	See page 93
MaximumMemory	Specifies the maximum memory for use by DDS	Advanced	See page 88
MaxRequestors	Specifies the number of requestors if the workstation is a file server	Basic	See page 78
NetworkRequestInterval	Specifies the time interval between queries used by the IPC component when attempting to determine the TCP/IP port associated with a broadcast domain	Advanced	See page 88
NetworkRequestRetries	Specifies the maximum number of times the IPC component will query the LAN to obtain the TCP/IP port associated with a broadcast	Advanced	See page 89

	domain		
PrimaryComputerName	ADDS a name to the new acting primary distributor	Basic	See page 78
PrimaryIPAddress	Specifies the IP address to be added when a workstation becomes the acting primary distributor	Basic	See page 78
PrimaryIPAdapter	Specifies the network adapter to which the PrimaryIPAddress will be added	Basic	See page 80
Reboot	Reboots the system to apply ControlledDrives and DDActive changes	Response File only	See page 93
RemoteIPC	Indicates if the IPC component is enabled	Basic	See page 82
RestartFDS	Stops and restarts DDS to apply configuration changes	Response File only	See page 94
RestoreConfig	Restores the configuration data from the specified file	Response File only	See page 94
WorkDirectory	Specifies the working directory to be used by DDS	Basic	See page 82

Data Distribution Configuration Keywords

The following table summarizes the configuration keywords relating to data distribution. These keywords apply only if this workstation is participating in data distribution.

Table 2. Data Distribution Configuration Keywords

Keyword	Description	Type	See Page
DDActive	Indicates whether the workstation participates in data distribution	Basic	See page 76
DistributionRole	Specifies the distribution role	Basic	See page 77
DistRenamedFile	Indicates whether distributed files (that are not within a distributed directory) should remain distributed if they are renamed. The default value is Yes.	Advanced	See page 85

Interprocess Communication Keywords

The following table summarizes the configuration keywords relating to the IPC component. These keywords apply only if there is more than one DDS workstation in your system.

Table 3. Interprocess Communication Configuration Keywords

Keyword	Description	Type	See Page
AdptrResetValue	Specifies the adapter reset frequency	Advanced	See page 84
DDSBlockedInterface	Specifies TCP/IP address that Distributed Data Services/Controller Services Facility should not use	Advanced	See page 84
IPCHeartbeatInterval	Specifies how often DDS IPC should check to See page if the communications connection with another node has failed	Advanced	See page 86
IPCPortCount	Specifies the number of TCP/UDP ports to be used by IPC when using TCP/IP as the transport layer protocol	Advanced	See page 86
IPCPortStart	Specifies the first TCP/UDP port number to be used by IPC when using TCP/IP as the transport layer protocol	Advanced	See page 86
IPCTimeout	Specifies the IPC timeout	Advanced	See page 87
IPCTransport	Specifies the transport layer protocol to be used by IPC	Basic	See page 77
LocatePrimary	Specifies whether data distribution is active anywhere in the system	Advanced	See page 87
MaxRequestors	Specifies the number of requestors if the workstation is a file server	Basic	See page 78
RemoteIPC	Indicates if the IPC component is enabled	Basic	See page 82
NetworkRequestInterval	Specifies the time interval between queries used by the IPC component when attempting to determine the TCP/IP port associated with a broadcast domain	Advanced	See page 88
NetworkRequestRetries	Specifies the maximum number of times the IPC component will query the LAN to obtain the TCP/IP port associated with a broadcast domain	Advanced	See page 89

Chapter 8. Unattended Configuration

You can perform the unattended configuration along with the initial installation of DDS, or you can perform configuration separately. To perform an unattended configuration of DDS:

1. Create the configuration response files. To do this, you should use a standard editor.
2. Use a software distribution manager to initiate the processing of the response files at the local workstation.
 - If you are configuring your system and performing an initial installation at the same time, see “**Starting Unattended Installation with SETUP**” for information about performing an unattended installation.
 - If you have already installed DDS or DDS/CSF and you want to reconfigure workstations, see the **FDSCRFP** command description below for more information about how this is done.

Creating a Response File for Unattended Operation

Response files can be created to be used to configure DDS or DDS/CSF at initial installation, program upgrade, or to simply modify the configuration at some later point in time.

Creating Response Files for an Unattended Initial Installation

Response files should be typically created using configuration files that have been used on a test machine in your test environment. Once they have been thoroughly verified and, assuming that similar environments exist in your stores, you can create the response files in the following manner:

1. Run an initial installation on a primary, secondary, or subordinate lab machine (you must do all three types, if applicable) using the “/R” parameter on the SETUP command.
2. Choose the appropriate options for the type of machine you are using (primary, secondary, or subordinate). Go through a complete and satisfactory setting of DDS or DDS/CSF configuration options. If you want the remote (unattended) installation to reboot at the end, click YES, when prompted or NO when prompted.

3. Copy a file, SETUP.ISS, that was created during setup. It will be found in the main directory of your Windows OS (WINNT or WINDOWS, normally). This file contains Windows SETUP commands that are used to do part of the unattended installation of DDS or DDS/CSF. Copy this file to a media of your choice.
4. Copy the files, FDSCONFIG.CFG, FDSWS.CFG, and FDSL.N.LN (normally found in the C:\FDS\CONFIG directory) onto a media of your choice. Using an editor, merge these three files. Save the resulting file as any file name. Example: RESPONSE.RSP
5. Edit the resulting file to replace the command “LogicalNames” with “ChangeLogicalNames”. Also, make any changes you need to make to the NodeID or SystemID keywords since at least NodeID must be unique for each controller in your system. Please refer to the NodeID and SystemID command descriptions to understand how you can set up your NodeID and SystemID to be unique across your store chain.
6. Create an FDSNODES.TBL file if required. See the NodeID command description for more information on this.
7. Using the software utility you will use to execute the unattended installation, you must use the response file, possibly the FDSNODES.TBL file, and the setup.iss file (along with the installation files) to execute the remote installation. Please refer to the SETUP command to understand the format of this remote setup command.

Creating Response Files for an Unattended Configuration Change

If you need to simply change an existing DDS or DDS/CSF configuration in a remote controller, you should use the FDSCRFP command to do so. This command is normally found in the C:\FDS\BIN directory on the machine to be modified. In this case, you would have in mind the configuration changes you intend to make and you must create a response file that contains just those commands. If a similar command already exists in a configuration file, you can override that command by doing an unattended update. For this situation, you must add the **ConfigUpdateType=Update** keyword in the response file. With this command, all existing configuration information is left intact and only the new commands you send are added or modified. In addition, refer to the section, **Response-File-Only Configuration Keywords**, for a discussion of other commands you might need when doing an update of this type. You can then issue the FDSCRFP command using the response file you have created. See the description of this command below.

Configuration Response File Processor (FDSCRFP)

The FDSCRFP command creates and updates a logical-names file and/or a configuration file based on information in the specified response file.

When you use the SETUP command to perform an unattended installation of the DDS base component for the first time, FDSCRFP is automatically called if its parameters are specified. It uses the response file specified by the /P: parameter and processes the configuration keywords in that file.


For information about configuration-specific keywords, see “**Chapter 7. Configuration Keywords**”. For more information about the SETUP command, see “**Starting Unattended Installation with SETUP**”.

You can use FDSCRFP to perform an unattended reconfiguration of a DDS system that is already installed. Only the configuration of the workstations is affected; no DDS product files are changed.

DDS must be installed on the client workstation from which this command is called. The target directory for configuration files and logs created by this command is the same as the target installation directory for DDS. If DDS is not installed, this command will fail.

When you use FDSCRFP, an error log and history log is maintained on the client workstation. When configuration is complete, the log files are copied to the files specified by the /L1 and /L2 parameters, appending to those files if they already exist.

Syntax for FDSCRFP



/P:

Specifies the drive, path, and file name of the specific response file. The drive and path are optional.

The following search order is used to find the response files:

1. The fully qualified file specification
2. The current directory from which FDSCRFP is called
3. Each directory in the PATH environment variable (on the server where FDSCRFP is called)

/L1:

Specifies the drive, path, and file name of the error log file. If the drive and path are not specified, the drive and path where the FDSCRFP command is running from are used. The path typically points to a location on the code server so that an administrator would have access to the logs in case of failure. The log file is also written to the client workstation, replacing any existing log. It is named target_install_directory\CONFIG\FDSERR.LOG.

Notes:

1. If you point the log file to a location on the code server, make sure that the current logon has write access to the location.
2. The value for the /L1: parameter may not contain hyphen (-) or slash (/).

/L2:

Specifies the drive, path, and file name of the history log file. If the drive and path are not specified, the drive and path where the FDSCRFP command is running from are used. The path typically points to a location on the code server so that an administrator would have access to the logs in case of failure. The log file is also written to the client workstation, replacing any existing log. It is named target_install_directory\CONFIG\FDSCHIST.LOG.

Notes:

1. If you point the log file to a location on the code server, make sure that the current logon has write access to the location.
2. The value for the /L2: parameter may not contain hyphen (-) or slash (/).

Return Codes

The following return codes are sent to the software distribution manager.

Return Code -- Description and Action

0	A successful program stop occurred but warning messages might be logged.
2048	A data resource was not found. Configuration fails. This code is returned if the value for /L1 or /L2 is not valid.
2070	The product is not configured. Errors are logged and configuration fails. This code is returned if: The system is out of memory The value for /P is not valid Keywords or values are found that are not valid, or if inconsistencies are found between keywords
5632	The program syntax was not used correctly or a disallowed parameter

- was passed. Configuration fails.
- 5636 An unexpected condition has occurred. Messages are logged and configuration fails. This code is returned if any internal errors are encountered. If, during your initial installation and configuration of DDS, a configuration error is detected (for example, you do not provide a local node ID), the initial installation or configuration is stopped and return code 0x1604 is returned to your software distribution manager. The error log and history log you specified will contain detailed information on the error and how to correct it. Correct the problem and retry the initial installation or configuration.
- 65024 The program stopped successfully. Reboot, but do not start FDSCRFP.EXE again. This code is returned if **Reboot=YES** and if the values for the **ControlledDrives** or **DDActive** keywords are changed.
- 65028 The program stopped successfully, but warning messages were logged. Reboot, but do not start FDSCRFP again. This code is returned if **Reboot=yes**, if the values for **ControlledDrives** or **DDActive** are changed, and if warnings have been logged.

Notes:

1. An error results if incorrect keywords or values are detected in the configuration file or logical-names file after the response file is processed. This error condition usually indicates that there were incorrect values in the input configuration or logical names file that were not corrected in the response file.
2. If you specified the /L1 or /L2 parameter and you do not see the error or history log file, inspect the logs on the local workstation. The local log files are copied to the file names specified by the /L1 and /L2 parameters. A failure when copying to the /L1 or /L2 file names is a warning, so you should inspect the local log files for other warning messages.

If local log files do not exist or are empty, determine whether the disk is full or cannot be written to for another reason. Correct the problem and run FDSCRFP again.

3. Back up the logical-names file before using this command-line utility if you are updating the logical-names file. Then, you can recover if the update fails due to insufficient disk space.

Using FDSCRFP to Generate Configuration Files and Logical Names Files

This example generates DDS configuration and logical names files. The response file is taken from drive X and the log files are copied to the specified file names, which are also on drive X. This example assumes that you have installed DDS in the C:\FDS directory.

```
C:\FDS\BIN\FDSCRFP      /P:X:\RSP\CLIENT5.RSP  
                        /L1:X:\LOG\CLIENT5E.LOG  
                        /L2:X:\LOG\CLIENT5H.LOG
```

Note: The parameters in this example are shown on separate lines for clarity.

Chapter 9. DDS/CSF Installation and Configuration

This section assumes a basic working knowledge of Windows operating systems, 4690 OS, and DDS. Earlier sections of this manual provide the required DDS knowledge.

DDS provides two optional features for customers who are currently using the 4690 operating system and want to preserve the investment in 4690 applications, terminal hardware, and terminal wiring while moving to an open platform for the back room:

- **DDS 4690 Controller Services Facility (CSF)**
DDS provides support for running 4690 Operating System controller functions and applications on Windows NT, 2000, XP, or Server 2003. It provides many functions of the 4690 Operating System including support for multiple controllers.

Installation

The installation process for CSF fits into the 4690 installation process. You build a system at your central site with all the required software and then package the newly built system for delivery to the stores.

Two installation checklists follow: one for a single controller environment and one for a multiple controller environment.

Checklist of Steps: One or Two Controllers

Note: More details about each of these steps follow this checklist.

User's should be aware that, if you are using a 4690 image that is configured with the Multiple Controller Feature (MCF), you should install that image onto two DDS/CSF controllers which are configured to run "data distribution". If you are using an image from a 4690 Controller that is not configured to run MCF, then you should install that image onto one DDS/CSF controller which is NOT configured to run data distribution.

1. Configure a native 4690 native system for installation onto your Windows machine.
2. Determine disk drive requirements: Normally, the Windows OS will be installed on the C: drive. Logical partitions can be created on the same physical drive or on another drive that can be used to hold the 4690 C: / D: images.

3. Install Windows NT, 2000, XP, or Server 2003 including the DLC protocol on the controller. See **Installing Windows NT, 2000, XP, or Server 2003 on Your Controller**.
4. Be aware that CSF requires a JAVA JVM in order to be fully functional. This JVM is automatically installed as part of the DDS/CSF installation.
5. Install TelnetServer 2000 from Pragma Systems, Inc., if required. See **Installing a Telnet Server for 4690 OS Auxiliary Consoles** for more information.
6. If you have one or more of the following adapters to be installed, do it now, following the instructions provided with the adapter. In addition, install any drivers for the adapters that are not provided by CSF. Refer to **“Store Loop Adapter Installation” on page 109** for more information.
 - Store Loop Adapter
 - PCI Store Loop Adapter
7. Install DDS/CSF. See **Installing DDS and CSF**.
 - Set the Store Controller ID during installation. The default value is CC. See **“Controller IDs” on page 110**.
 - If you are configuring two DDS/CSF controllers to run a system with the Multiple Controller Feature installed, you should specify “data distribution” as an option. You must also specify one of the controllers as a Primary Distributor and the other as a Backup Distributor.

If you have store loop adapters, set their parameters during installation. For more information, see **Store Loop Adapter Setup on page 111**

- Create CSF logical names for the 4690-like disk drives during installation. See **Installing DDS and CSF** for more information about the logical drives.
 - NOTE: Do not specify data distribution during installation if you are going to install DDS/CSF on only one controller.
8. Copy the 4690 OS files and the application files from the 4690 OS library machine to the Windows NT, 2000, XP, or Server 2003 target machine. See **Copying Files from 4690 to the Windows Machine via FTP** on page 114 for more details.
 9. Obtain the 4690 file distribution attributes by running FDSQDIST.286. Run FDSQDIST.BAT on the DDS primary distributor. See **Set the Distribution Attributes of the 4690 Files**. This step is not necessary if you will be running a controller without the Multiple Controller Feature enabled.

10. Boot your controller(s) to start 4690 CSF.
11. Start a 4690 console in one of two ways:
 - by entering FDSECNAT at a Windows command prompt
 - by clicking **Start**, then **Programs**, then **QVS DDS**, then **4690 CSF Console**
1. **Note:** If the store controller has more than one LAN adapter configured with the DLC protocol, then CSF needs to know which LAN adapter to use for Terminal to Controller Communications (TCC). This is done with the TCC4690_LANA Windows environment variable. You can obtain the value to set this variable to from the output of the FDS DLC command. See **Installing DDS and CSF** for more information on TCC4690_LANA and **Automatic Pre-Processing, Post-Processing, Custom Processing** for information about the FDS DLC command.

Installing Windows NT, 2000, XP, or Server 2003 on Your Controller

Determine which version of Windows you want to run. Be aware that all features of DDS/CSF will run in any Windows environments.

You may have other applications that you want to run on the controller besides DDS and CSF. Check the requirements for these applications to determine if they require a particular Windows version.

Windows must be configured for the DLC protocol. Follow the appropriate Microsoft installation instructions for Windows NT or 2000. Note that DLC is automatically installed by the DDS/CSF installation package on either Windows XP or Server 2003. In these cases, if your machine has more than one network card installed, you **MUST** assure that DLC is installed only on the card you intend to use to communicate with 4690 Terminals, either native or emulated. See **Installing DDS and CSF** for more information on TCC4690_LANA and **Automatic Pre-Processing, Post-Processing, Custom Processing** for information about the FDS DLC command.

If you intend to run under Windows XP with Service Pack 2 installed and the Windows Firewall enabled, you must do one of two things in order to run successfully with DDS.

1. Disable the Windows Firewall
2. Enable the Windows Firewall, but configure an exclusion for the following programs:
 - c:\fds\bin\fdsiprct.exe

- c:\fds\bin\fdst286.exe
- c:\fds\bin\dhcpd.exe

Note: The path of the above exclusions could be different if you have installed the DDS/CSF program in another directory. If so, you should adjust the path accordingly. You may discover that other functions on your XP platform may also have to be granted exceptions in order to run with the Windows Firewall enabled. The list above may not prove to be comprehensive.

Installing a Telnet Server for 4690 OS Auxiliary Consoles

CSF requires a Telnet server for LAN-attached 4690 OS consoles. The TelnetServer 2000 from Pragma Systems, Inc. is supported.

After installation of TelnetServer 2000, remember to configure the Graceful Termination option. For more information, see Auxiliary Console Support in the **QVS Distributed Data Services/Controller Services Facility for Windows User's Guide**.

Store Loop Adapter Installation

If you currently have an ISA store loop card installed and you want to change to a PCI store loop card, you can replace the ISA card with the PCI card. No other installation or configuration changes are necessary in order for the PCI store loop card to operate.

If you currently have a PCI store loop card installed and you want to change to an ISA store loop card, you must uninstall CSF and reinstall CSF in order to be able to fill in the settings for address and interrupt level required for the ISA store loop card.

CSF supports the use of one or two ISA or PCI bus store loop adapters. A mix of loop adapter types (PCI and ISA) is not supported. Install the adapters as specified in the instructions that come with the adapters.

Before installing the ISA store loop adapters, you may have to set jumpers to indicate the memory addresses and interrupt levels used by the adapters. You may have to work out memory address conflicts with other adapters or devices in your computer. The acceptable hex addresses for each loop adapter are:

- hex C8800
- hex CD800
- hex D0800
- hex D8800

No address may be used for more than one device.

The acceptable interrupt levels are 7 and 11. If you have two ISA store loop adapters, they must both use the same interrupt level.

Make note of the addresses and interrupt levels. You will need this information during installation of CSF.

PCI Token-Ring Adapter Setup

In a multiple controller environment with CSF, the PCI token-ring speed on both machines must be the same. Do not use the option to automatically determine ring speed.

It is strongly recommended that the NO AUTONSENSE option be used when setting up the adapter. Set all controllers to the same token-ring speed, matching the speed of the LAN.

Installing DDS and CSF

DDS has several functions:

- DDS Distribution
- Automatic Switch-Over component,

DDS/CSF has several functions:

- DDS Distribution
- Automatic Switch-Over component,
- CSF with Multiple Controller Support

If you are building a single-controller store environment, you will need the DDS/CSF version of the installation.

Chapter 3. Installation tells you how to perform the installation.

Controller IDs

Each 4690 CSF controller must have its own controller ID. The controller ID is set during initial installation of 4690 CSF. The default value is CC.

Notes:

1. You must make this ID the same as the ID for the source 4690 controller.
2. Any time a 4690 controller ID is changed on one machine in a pair of controllers, you must reboot both controllers for the change to take effect and work properly.

Store Loop Adapter Setup

CSF supports the use of one or two ISA or PCI bus loop adapters. If you are using one or more ISA store loop adapters with CSF, you will need to supply the memory addresses and interrupt levels to the CSF installation program. The memory addresses and interrupt levels are determined by the jumper settings on the adapters.

If you are not using a store loop adapter, or if you are using a PCI store loop adapter with CSF, you will not need to supply store loop adapter information.

If you are installing CSF by performing attended installation, you will be asked to indicate whether you have 0, 1, or 2 store loop adapters installed. You will need to supply memory addresses and interrupt levels for the loop adapters. See **Store Loop Adapter Installation** for more information about store loop adapter memory addresses and interrupt levels. If you add a loop adapter or change loop adapter jumper settings, you must uninstall and re-install DDS/CSF.

Defining Your Logical Drive Names

You need to consider the types of applications you will be running and the types of Application Program Interfaces (APIs) you will be using in these applications before you decide how your Windows NT, 2000, XP, or Server 2003 system should be set up.

If you will be running applications in the Windows NT, 2000, XP, or Server 2003 environment that will use the 4690 C language APIs or the 4690 Java APIs, you may wish to install your Windows NT, 2000, XP, or Server 2003 operating system on the E: drive of your machine. Designate the 4690 C: drive (adxcsfh0 DDS logical name value) as C: and designate the 4690 D: drive (adxcsfh1 DDS logical name value) as D:. This will let your application writer know where a file is located. The C: drive will be the C: drive regardless of which API is being used.

The following examples illustrate the confusion that can occur if you do not choose your 4690 C: and 4690 D: drive locations carefully.

Example 1: Running a Windows NT, 2000, XP, or Server 2003 application that uses the 4690 C APIs to access the 4690 keyed files

Since this is a Windows application, you must consider the different file systems being used when accessing files-- the Windows NT file system (FAT or NTFS), the Windows 2000 file system (NTFS, FAT32, or FAT), and the 4690 OS FAT file system. When you use C language based I/O or Windows APIs for file access, you are using the Windows NT, 2000, XP, or Server 2003 file system. When your application references a file with the name D:\MYDIR\MYFILE.DAT, Windows will look on the D: drive for it. When you use any of the 4690 C language APIs, for example, ADX_COPEN_KFILE, you are using the 4690 file system, which runs in the 4690 environment provided by CSF. When your application references a file with the name D:\MYDIR\MYFILE.DAT in the 4690 C language APIs, CSF will interpret the drive on which to find the file by using the adxcsfh1 DDS logical name. You set up this DDS logical name to tell CSF where it can expect to find the 4690 D: drive, and this may or may not be the Windows D: drive. Note: The logical names, adxcsfh0 and adxcsfh1 are case sensitive and should be specified only in lower case.

Example 2: Running a Java application that uses the 4690 Java APIs, such as KeyedFile on your controller machine.

Since this is a Windows machine, the Java Virtual Machine (JVM) will be a Windows version of the JVM, not the 4690 OS version of the JVM. You must consider the different file systems being used when accessing files-- the Windows NT file system (FAT or NTFS), the Windows 2000 file system (NTFS, FAT32, or FAT), and the 4690 OS FAT file system. When you use the pure Java APIs to access files, such as FILE, the application will be using the Windows NT, 2000, XP, or Server 2003 file system. When your Java application references a file with the name C:\JAVAPPL\FILE1.XXX, the JVM will go to the C: drive to find the file. When you use the 4690 Java API, such as KeyedFile, the application makes a call into the 4690 environment and CSF will interpret the file name using the DDS logical names you set up. When your Java application references the file with the name C:\JAVAPPL\KEYDFILE.DAT, CSF will use the DDS logical name adxcsfh0 you set up to translate the C: drive into the 4690 C: drive location on the Windows NT, 2000, XP, or Server 2003 machine, which could be a drive different from the C: drive.

You also need to decide where the 4690 subdirectories and files will reside on your Windows NT, 2000, XP, or Server 2003 controller. They should be on a drive controlled by DDS. A DDS logical name must be defined to identify the drive.

Determine if your 4690 environment and applications require one or two drives. 4690 OS supports one or two drives: C and D. The drives referred to by 4690 applications running in the CSF environment are logical drive references. DDS uses two logical names to refer to the corresponding 4690 drives: `adxcsfh0` corresponds to 4690 drive C: and `adxcsfh1` corresponds to 4690 drive D:. These values are defined in your DDS logical names file (`FDSL.N.LN`). This file can be modified using the Configuration File Utility, which is started at the end of DDS installation. More information can be found in **Chapter 6. Configuration File Utility**.

For example, assume that your 4690 application uses the C and D drives, and references those drives in the path for their files. Suppose that you have set up physical drive D to contain the files normally found on your 4690 native machine's C drive, and the physical drive E: to contain the files normally found on your 4690's D: drive. For this example you would add the following lines to the DDS logical names file (`FDSL.N.LN`):

Logical Name = <`adxcsfh0`> D:

Logical Name = <`adxcsfh1`> E:

Note: Drive designations are always in the form of the drive letter followed by a colon (for example, D:). The logical names, `adxcsfh0` and `adxcsfh1` are case sensitive and must be lower case.

4690 OS applications may request a file name as input. If you provide a full path for the file, you must enter the drive letter from the perspective of the 4690 application. For example, to compress a file using the File Compress/Decompress utility, you must provide file names for both input and output files. Suppose that the file you want to compress is the terminal dump file `ADXCSLTF.DAT` in the `ADX_SDT1` subdirectory. The output file is to reside in the root of the drive containing the `ADX_SPGM` subdirectory, which is the 4690 C drive. From the value for `adxcsfh0`, the 4690 C drive corresponds to the Windows D drive. When you enter the name of the output file for File Compress/Decompress, you must enter `C:ADXCSLTF.CMP`. Likewise, when you enter the name of the input file, you can enter either `ADX_SDT1:ADXCSLTF.DAT` or `C:\ADX_SDT1\ADXCSLTF.DAT`.

Multiple LAN Adapters

If your store controller will have more than one LAN adapter with the DLC protocol, then you need to tell CSF which LAN adapter to use for Terminal to Controller Communications (TCC). You can do this with the **TCC4690_LANA** Windows environment variable.

You can use the `FSDS DLC` command to determine the value of the `TCC4690_LANA` variable. This command will display information about the

LAN adapters in your store controller. You can use either the adapter number or the adapter title to set TCC4690_LANA. However, it is recommended that you use the adapter title.

You can use the complete adapter title or part of the adapter title to set TCC4690_LANA. If you use only part of the adapter title, you should use enough of the title to uniquely identify the adapter for CSF. Also, when using either the complete title or part of the title, you must start with the left part of the title. For example, if the adapter title is [1] IBM Shared RAM Token Ring Adapter, you can enter:

[1] IBM Shared RAM

This example uses only part of the title, but it starts with [1], which is the left part of the title.

If you have multiple LAN adapters and do not create and set TCC4690_LANA, then CSF will, by default, use adapter number 0 (zero). Also, if you create and set TCC4690_LANA but CSF can not match it to an adapter in your store controller, then CSF will, by default, use adapter number 0 (zero).

If CSF unsuccessfully attempts to use a LAN adapter, the following message will be displayed when CSF is initializing:

W510: LAPS 802.2

CSF creates a log file named laninfo.log in the fds\log directory. After CSF is started, this file will contain the value of **TCC4690_LANA** and the information about the LAN adapters in your store controller. This is the same information that is displayed by the **FSDSLC** command.

Copying Files from 4690 to the Windows Machine via FTP

The native 4690 OS machine must have TCP/IP installed and configured. Refer to IBM Transmission Control Protocol/ Internet Protocol for 4690 Application Interface Guide for details about setting up TCP/IP on 4690. Define a trusted user that can access at least the subdirectories from which you want to copy files (ADX_SPGM, ADX_SDT1, ADX_IPGM, ADX_IDT1, ADX_IDT4, and any other subdirectories your applications require). The easiest way to do this is to give a user access to the root of the C: drive on the 4690 machine.

Both the 4690 native machine and the Windows 2000 or XP machine must be on the same LAN. Be sure to stop all applications running on your 4690 OS machine, both foreground and background applications. Also turn off any terminals you have attached to your 4690 OS machine, or disconnect the registers from the 4690 OS controller. This ensures that the files that the applications would normally have open can be copied to the Windows machine. There are a few 4690 files that

cannot be copied due to the 4690 OS having the files open and locked. Instructions for dealing with these files are included in this section.

These instructions assume that your user ID has access to the root of the C: drive on the 4690 system and that you start FTP from the root that contains the 4690 subdirectories. From your Windows machine, enter the following commands to accomplish the file transfer:

FTP n

(where n is the name of your 4690 machine or the IP address of the 4690 machine)

Binary

(to have the files transferred as binary files)

Prompt

(to turn off prompting)

CD ADX_SPGM

Lcd ADX_SPGM (to change to the ADX_SPGM on your Windows 2000 machine)

Mget *.*

(to get all the files in the subdirectory and place them on your machine)

Cd ..

(to get back to the root of the source machine)

Lcd ..

(to get back to the root of the Windows machine)

Repeat these five commands for each subdirectory you want to copy to your Windows machine. After all the files are copied, enter the FTP QUIT command to end the FTP session.

The 4690 event log data files cannot be copied because the 4690 OS has these files opened and locked while 4690 is running. These files reside in ADX_SDT1 and are named ADXCISOAF.DAT , ADXCISOBF.DAT, ADXCISO CF.DAT, ADXCISO DF.DAT, ADXCISO EF.DAT, and ADXCISO FF.DAT. You can copy ADX_SDT1\ADXCISONF.DAT into each of the files named here. It is important that the files exist, but it is not important that they contain the same data that the 4690 OS system contained.

Set the Distribution Attributes of the 4690 Files

If you have a multiple controller environment, you should set up two controllers running DDS/CSF. 4690 OS files should initially be copied to both machines as described in previous sections. You should then use your 4690 library machine to get the 4690 distribution attributes of your 4690 files. CSF provides a program called FDSQDIST.286 to accomplish this task. It is a 4690 application that gathers the attributes of all the files in the 4690 system. It produces an FDSQDIST.BAT file, which runs on the Windows controller. To use FDSQDIST, follow these steps:

1. Copy the FDSQDIST.286 file from \FDS\BIN to a diskette.
2. Put the diskette in the 4690 native library machine and run CHKDSK F on the diskette in order to clear a file attribute mismatch between Windows and 4690 OS.
2. Copy the file, FDSQDIST.286, to the root of the C drive.
3. Execute FDSQDIST.286 from the command line.
4. Copy the resulting FDSQDIST.BAT file from the root of the C: drive to a diskette. This file contains the DDS commands to set the distribution attributes of all the 4690 files.
5. Put the diskette in the Windows machine where DDS and CSF are installed. Copy the FDSQDIST.BAT file from the diskette to the \FDS\BIN subdirectory on the hard disk of the DDS primary distributor.
6. Edit the FDSQDIST.BAT file and change the lines as required. In the following example, in spite of the listed instructions in the bat file, there is no need to modify either **nodexxx1** or **nodexxx2**. The environment variables **csfh0** and **csfh1** must match the drives that you identified for DDS logical names **adxcsfh0** and **adxcsfh1**, respectively. Also, please note that these two logical names must be defined the same way on both the primary controller and the backup controller.

Example:

```
REM Modify the nodexxxx in the following two lines:
echo FDSBDDOMN CSF2MCF /CREATE nodexxx1 nodexxx2 >>fdsqdist.log
FDSBDDOMN CSF2MCF /CREATE nodexxx1 nodexxx2 >>fdsqdist.log
REM Modify the following two lines if necessary:
set csfh0=d:
set cdsh1=e:
REM There will be a pair of lines like the next two for every file that you copied from the 4690 native machine.
echo FDSSETD %csfh0%:\ADX_SPGM\ADXNF.DAT /FILE /BROADCAST CSF2MCF /DOC >>fdsqdist.log
FDSSETD %csfh0%:\ADX_SPGM\ADXNF.DAT /FILE /BROADCAST CSF2MCF /DOC >>fdsqdist.log
echo FDSQDIST.BAT Completed >>fdsqdist.log
```

7. Stop the 4690 CSF service on both machines. This can be done in either of two ways.

a. Use the control panel. Perform the following steps:

1) Click Start from the Windows taskbar.

2) Point to Settings.

3) Take one of the following actions:

- On Windows NT, click Control Panel.
- On Windows 2000, XP, or Server 2003, click Control Panel, and then click Administrative Tools.

4) Double-click Services.

5) Select QVS Controller Services Facility.

6) Click Stop.

7) Click Yes.

b. Enter the NET STOP CSF command.

Make sure DDS is active on both machines using the FDSQUERY command and that the two machines are communicating with each other. You can verify communication with the PING command.

8. Run the FDSQDIST.BAT file to have the distribution attributes of all the 4690 files set on the DDS system. The .BAT file needs to be run on the DDS acting primary distributor only.

NOTE: After running the distribution program, examine the **FDSQDIST.LOG** file contained in the same directory as the file **FDSQDIST.BAT**. Verify that no errors are shown. Look for any errors in execution of the **FDSBDMN** command. If it shows an error, stop DDS, erase the **C:\FDS\WORK** directory, restart DDS and rerun **FDSQDIST.BAT**. Examine closely a few sample **FDSSETD** lines to see that no double backslashes are used in the path. If so, check the logical name, **csfh0** and **csfh1** to verify that neither end in a backslash.

Starting CSF

CSF has one Windows service and several Windows drivers.

The service is called CSF. The service and the drivers are set up to start when the machine on which they are installed is booted.

After your system boots, attach a console to CSF as follows: click Start, then Programs, then QVS Controller Services, then 4690 CSF Console. This action allows you to see the 4690 windows. If you want to attach an auxiliary console to CSF, use the command: FDSECNAT -a.

Note: Closing the CSF window does not log off the currently signed-on user. Also, CSF does not end when a user closes a console window. Closing the CSF window is not the same as closing a Windows window.

Stop the service and drivers using the NET STOP at a command prompt:

- NET STOP CSF stops the service.

The service and drivers can be started with the NET START command.

CSF Initialization Error Messages

During initialization of CSF you may see the following error messages:

W510: LAPS 802.2

This message indicates that the DLC protocol has not been installed or that the LAN adapter is not accessible to CSF. Verify that the DLC protocol has been installed for your LAN adapter or connection and that the adapter is functioning properly. Also, verify that CSF is trying to use the LAN adapter that will allow CSF to communicate with your terminals. See **Installing DDS and CSF** for more information.

W555: 001A

This message indicates one of at least three possibilities:

1. the DDS logical name adxcsfh0 has not been defined correctly (see **Installing DDS and CSF** for more information)
2. the native 4690 OS files are not at a supported level. See the DDS/CSF readme file, README.FDS, for details on which levels of 4690 OS are supported by CSF. The unsupported 4 digit release level is shown at the end of the **W555:001A** error message. You can determine the level of 4690 OS you are running by examining the file ADXCSTSD.DAT in the ADX_SPGM directory. Contact QVS personnel if you wish to run with a currently-unsupported DDS/CSF release or version to see how that may be accomplished.
3. You have installed a CSF configuration that is mismatched with the native 4690 image you are trying to run. If you use an image which supports the 4690 Multiple Controller Feature (MCF), you must install the Data Distribution option when you install DDS/CSF. Otherwise, there is a

mismatch. If you have a 4690 image and want to disable MCF in that image, you can rename the file ADX_SPGM\adxzm??f.dat to use a different file extension (example: ADX_SPGM\adxzm??f.sav). The ?? characters represent the node of the store controller (normally CC or DD).

Configuration and Customization

After you have installed 4690 CSF, you must determine whether you want custom processing performed before 4690 CSF starts or after 4690 CSF ends. The following sections help you configure 4690 CSF and set up any custom processing that you require.

4690 CSF Configuration

Configuration of 4690 CSF consists of 4690 OS configuration and configuration of any installed 4690 OS applications, such as the IBM 4690 Supermarket application.

This configuration can be done in whole or in part on the 4690 OS library machine at your central site. Refer to your 4690 OS documentation for details about performing this configuration.

The following configuration actions must be performed.

- Terminal configuration
- Controller configuration
- System configuration

Configuration details that are unique to each store must be configured on your Windows controller using 4690 CSF. These details might include store number, terminal numbers, and operator IDs and passwords.

Automatic Start and Restart of 4690 CSF

There may be situations when CSF will end. In the store, you will want it to start again automatically to keep the front end operational. When CSF is installed, an entry is made in the Windows registry, identifying the service used to start CSF. This registry entry indicates that the service is to start automatically. The CSF service uses the parameters field of the registry entry as the source of the name of the .BAT file to start. This field identifies FDS_CLOG.BAT, residing in the FDS\BIN subdirectory.

The start process is as follows:

1. Windows boots, starting the CSF service.
2. CSF starts the \FDS\BIN\FDS_CLOG.BAT file.

3. FDSCLOG.BAT starts FDSEPROC.EXE.

If CSF stops for any reason, the service will restart the FDSCLOG.BAT file. (This includes the times when you perform "Load Controller Storage" and "Dump Controller Storage" functions.) Up to four attempts to restart are made in three minutes. If the attempts are not successful, the service stops. This condition indicates a serious problem that must be addressed.

Automatic Pre-Processing, Post-Processing, Custom Processing

FDSCLOG.BAT queries the existence of two files: PREEPROC.BAT and POSTEPRC.BAT. If PREEPROC.BAT exists, then it is run before starting CSF. If POSTEPRC.BAT exists, it is run after CSF ends. If these files are used, they must be in the FDSCUST subdirectory, in the root of the drive where DDS is installed. This subdirectory is not created when DDS/CSF is installed.

Pre-CSF Processing

To have processing performed before CSF is started, you must have the file PREEPROC.BAT in the FDSCUST subdirectory. The FDSCUST subdirectory must be placed in the root of the drive where DDS is installed.

FDSCLOG.BAT queries the existence of the PREEPROC.BAT file. If it exists, it is called each time FDSCLOG.BAT is run. After PREEPROC.BAT ends, FDSCLOG.BAT starts CSF. This process provides much flexibility, but it should be used with caution.

For example, assume that a special file must be copied to a particular subdirectory before FDSEPROC.EXE is started. To accomplish this file copy, your PREEPROC.BAT file would contain the following line:

```
COPY C:\SPECIAL\DATA\STASH\UNIQUE.DAT D:\ADX_UDT1\UNIQUE.DAT
```

Each time FDSCLOG.BAT is run, PREEPROC.BAT will run before FDSEPROC.EXE is started.

Note: Your PREEPROC.BAT file must not contain the EXIT batch file command. If this command exists in your batch file, FDSLOG.BAT will end after PREEPROC.BAT is run and prevents CSF from starting.

Post-CSF Processing

To have processing performed after CSF ends, you must have a file named POSTEPRC.BAT in the FDSCUST subdirectory. The FDSCUST subdirectory must be in the root of the drive where DDS is installed.

FDSCLOG.BAT queries the existence of the POSTEPRC.BAT file. If it exists, POSTEPRC.BAT is called each time CSF ends. FDSCLOG.BAT waits for

POSTEPRC.BAT to end before continuing. If POSTEPRC does not end, CSF will not start FDSLOG.BAT again.

For example, suppose that your application writes a Processing Results file. You want to save at least eight generations of the data. Your holding place for this data is in D:\HOLDING\PLACE. The processing results file is D:\ADX_UDT1\PROCRSLT.DAT. You also want the data compressed to conserve disk space. Your POSTEPRC.BAT file contains the following lines:

```
REN D:\HOLDING\PLACE\RSLTDATA.007 RSLTDATA.008
REN D:\HOLDING\PLACE\RSLTDATA.006 RSLTDATA.007
REN D:\HOLDING\PLACE\RSLTDATA.005 RSLTDATA.006
REN D:\HOLDING\PLACE\RSLTDATA.004 RSLTDATA.005
REN D:\HOLDING\PLACE\RSLTDATA.003 RSLTDATA.004
REN D:\HOLDING\PLACE\RSLTDATA.002 RSLTDATA.003
REN D:\HOLDING\PLACE\RSLTDATA.001 RSLTDATA.002
FDSCOMP C D:\HOLDING\PLACE\RSLTDATA.001
D:\ADX_UDT1\PROCRSLT.DAT
```

(FDSCOMP is a utility provided by DDS for compressing files.)

Note: Your POSTEPRC.BAT file must not contain the EXIT batch file command. If this command exists in your batch file, FDSLOG.BAT will end after POSTEPRC.BAT is run, thus preventing CSF from automatically restarting after the previous invocation of CSF ends.

Custom Processing

Instead of having FDSLOG.BAT start CSF, you might want to have customized processing take place, including the starting of CSF. For custom processing, you must have a file named EPROCCMD.BAT in the FDSCUST subdirectory. The FDSCUST subdirectory must be in the root of the drive where DDS is installed. FDSLOG.BAT queries the existence of EPROCCMD.BAT. If it is present, it is run instead of starting FDSEPROC.EXE.

Part of the processing in the EPROCCMD.BAT file must be a call to FDSEPROC.EXE. This lets you control when and how FDSEPROC is called. This call must be as follows:

```
FDSEPROC
```

Note: Your EPROCCMD.BAT file must not contain the EXIT batch file command. If this batch file command exists in your batch file, FDSLOG.BAT will end after EPROCCMD.BAT is run, thus preventing CSF from automatically restarting after the previous invocation of CSF ends.

You can add tracing parameters if you wish. See **Automatic Pre-Processing, Post-Processing, Custom Processing** for details.

Example:

Suppose that you have a program you want to start at the same time that FDSEPROC is started. Your EPROCCMD.BAT file in the FDS\USER subdirectory should contain the following lines:

```
START C:\MYPROC\SIDEPROC.EXE  
FDSEPROC
```

Command Syntax for FDSEPROC and FDSECNAT

FDSEPROC is the command used to bring up CSF. In certain cases, it can be used to change runtime parameters, as well.

The syntax and parameters for FDSEPROC follow. All the parameters are optional.

FDSEPROC <-k> <-t> <-v>

- **-k** indicates that CSF will end.
- **-k1** indicates that CSF will end and produce a memory dump.
- **-t** indicates binary logging. This parameter is most useful in helping to debug problems with 4690 CSF. It generates the same log data as the **-v** option, but in a compressed binary format, which consumes less disk space.
- **-v** indicates verbose logging. This parameter is most useful in helping to debug problems, however it generates large log files. You can turn on this feature by using the FDSEPROC command with the **-v** parameter after 4690 CSF is running. CSF will always generate minimal log files if the **-v** flag is not used on the FDSEPROC command.

FDSECNAT allows you to bring up the 4690 OS screens by way of CSF. FDSECNAT has one optional parameter. The syntax and parameter for FDSECNAT are:

FDSECNAT <-a>

-a indicates that this is an auxiliary console.

FSDS DLC

You can use FSDS DLC command to display information about the LAN adapters in your store controller. There are no parameters on this command. When you run FSDS DLC from a Windows command prompt you will see a display of information about the LAN adapters in your store controller. Each adapter will have an adapter number and an adapter title. You can use the information from this display to set the TCC4690_LANA Windows environment variable. After

CSF is started, this information will also be contained in a log file named laninfo.log in the fds\log directory.

FDSCSFBH

During installation, the user is prompted to enter the 4690 Controller ID and the store loop adapter configuration values. After installation is complete, these values can be changed by either rerunning the installation and entering new values or by using the FDSCSFBH command.

The FDSCSFBH command accepts one or more optional parameters specified using a keyword, followed by a value. The syntax is as follows:

FDSCSFBH <parameters>

Keyword	Valid Values
/CTRLR	Controller ID alphabetic characters CC through ZZ
/LOOP1	Loop1Address hex 0 and C8800
/LOOP2	Loop2Address hex 0, C8800, D0800, D8800, and DC800
/LOOPIRQ	LoopIRQ decimal 7 and 11

For example, to change the interrupt used by the Loop 1 adapter to IRQ 11, and the memory address to hex C8800, first change the jumpers on the adapter to configure the hardware to those values. Then issue the FDSCSFBH command using the following syntax:

FDSCSFBH /LOOPIRQ 11 /LOOP1 hex C8800

As shown in this example, you are not required to use all the keywords and the keywords may be entered in any order. Each keyword must be followed immediately by the value for that keyword.

Configuring a Printer

4690 printer devices are mapped to Windows NT, Windows 2000, XP, or Server 2003 printer devices as follows:

- PRN1:, LPT1: maps to LPT1
- PRN2:, LPT2: maps to LPT2
- PRN3:, LPT3: maps to LPT3

4690 device drivers can be defined as serial or parallel.

There are two methods in Windows to enable device LPT1 to print, other than using a physical, parallel port:

- a. **Printer Pooling method:** Windows supports a concept called pooled printers. With this method, Windows spreads print jobs across the pool


of printers. You can define LPT1 and a network printer to be in the same pool. If no printer is attached to LPT1, Windows sends the print job to the network printer because LPT1 is in an error state and cannot be used.

- b. Net Use method: If you want LPT1 to go to a network printer, use the following command. “NET USE LPT1: \\srv1\printer”, where srv1 is the name of the server to which the printer is attached and printer is the name of the printer.

Appendix A. Controller Services Facility Directory Structure

The installation program automatically creates the FDS (or the target installation) directory. It contains all the other subdirectories and files for DDS.

The directory structure for DDS is as follows:



The subdirectories and their contents are:

BIN

Programs used by DDS.

CONFIG

Configuration files.

JAVA

Java files.

HELP

Help files and message files.

WORK

The working directory. DDS creates and manages several temporary files in this directory. DDS also uses this directory to store information when DDS stops; this information is then available the next time it starts.

Note: The location of the work directory is x:\FDS\WORK, where x is the drive where DDS is installed. This drive is specified at installation. You can specify that it reside in a different location using the **WorkDirectory** configuration keyword. See **WorkDirectory on page 82** for more information.

You should not work with or delete any files from this directory. You can change the drive and the name for this directory using the Configuration File Utility. See **Chapter 6. Configuration File Utility on page 58** for more information about the Configuration File Utility.

Note the following directories:

LOG

Contains the 4690 Controller Services Facility log files. This subdirectory exists only when the 4690 Controller Services Facility is installed.

4690RAM1 -- 4690RAM4

Contains the 4690 Controller RAM disks.