DataFlux® Data Management Studio 2.7: Installation and Configuration Guide

Contents

What's New

What's New in the DataFlux Data Management Studio Installation and Configuration Guide

Overview

The June 2017 update for DataFlux Data Management Studio 2.7 provided performance and stability enhancements. Additional information and clarity was added to the *DataFlux Data Management Studio 2.7: Installation and Configuration Guide* to reflect some of those enhancements.

DataFlux Data Management Studio has the following new installation and configuration enhancements:

- new DAC option on page vi

- new web service node option on page vi

- renamed configuration and log files on page vi

- changes to existing nodes on page vi

- updated migration instructions on page vii

- restriction on saving configuration files on page vii

- updates to the configuration table on page vii

New DAC Option

A new option, DAC/DISABLE_OLD_STYLE_FS_DSN, was added to determine whether to use DataFlux Federation Server definitions from the SAS Metadata Repository or use existing DNS files.

New Web Service Node Option

A new option, WEBSERVICE/REQUEST_TIMEOUT, was added to specify the number of seconds that the **Web Service** node should wait for a response from the web service.

Renamed Files

The following files were renamed:

- Dfwfproc was renamed to batch.cfg.

- Dfwsvc.cfg renamed to service.cfg.

- Dfwfproc.log.xml renamed to batch.log.xml.

- Dfwsvc.log.xml renamed to service.log.xml.

Changes to Nodes

The **SAP Remote Function Call** node should be used only with the SAP Unicode Libraries.

Updated Migration Instructions

All DataFlux Data Management Studio or DataFlux Data Management Server migrations now use the SAS Deployment Wizard. See the "Upgrading Existing Installations of the Software" on page 2 for more information.

With the release of SAS Visual Process Orchestration 4.2, the migration table in "Migrating to SAS 9.4" on page 2 has been updated. Migration is still manual.

Restriction on Saving Configuration Files

When editing a configuration file, save it as an ASCII or UTF-8 file. If it is saved as a Unicode file, the software can fail to start and will log this error message: "Read failed due to missing file or possibly invalid content in 'c:\path\file.cfg.'"

Configuration Options

Updated the format of the configuration options and added several options.

The following BASE options were added:

- BASE/AUTH_SEARCH_DOMAIN

- BASE/AUTH_SERVER_REPOSNAME

- BASE/HIGH_PRECISION_TIMER

- BASE/MAINTAIN_GROUP

- BASE/SHARED_OUTPUT_PATH

- BASE/USE_SASSORT

The following DAC option was added:

- DAC/TKTSLOGFLAGS

The following Profile option was added:

- PROFCONFIG/MAXFREQCACHE

The following Architect options were added:

- JAVA/VMHEAP_INIT

- JAVA/VMHEAP_MAX

- VERIFY/GEO

- VERIFY/CACHESIZE

- VERIFY/USEDPV

- VERIFY/USERDI

- VERLQT/LOQATE_DATA_PATH

- VERLQT/COUNTRY_INIT_FILTER

- VERLQT/PRELOADTABLELIST

- VERLQT/MAX_REF_PAGE_CACHE_SIZE

Accessibility

Accessibility Features of DataFlux Data Management Studio

Overview

For information about the accessibility of DataFlux Data Management Studio, see the *DataFlux Data Management Studio Users Guide*.

For information about the accessibility of any of the products mentioned in this document, see the documentation for that product.

DataFlux Data Management Studio has not been tested for compliance with U.S. Section 508 standards and W3C web content accessibility guidelines. If you have specific questions about the accessibility of SAS products, contact accessibility@sas.com or SAS Technical Support.

Documentation Format

If you need this document in an alternative digital format, contact accessibility@sas.com.

1

Installing DataFlux Data Management Studio

Installing the Software for the First Time

DataFlux Data Management Studio is available through SAS delivery channels. For information about installing this product, see your SAS Software Order E-mail (SOE).

The default installation path under Windows is: `SASHome \<product_instance_name>\<version>`.

The default installation path under UNIX is: `SASHome/<product_instance_name>/ <version>`.

For configuration instructions, see Chapter 2, "Configuring DataFlux Data Management Studio," on page 7. To install add-on products, see Chapter 3, "Installing Add-On Products," on page 39.

Upgrading Existing Installations of the Software

Migrating to SAS 9.3

The following table summarizes the upgrade paths when the target system is SAS 9.3.

Existing Version	Upgrade Version	Migration Instructions
DataFlux Data Management Studio or DataFlux Data Management Server (all versions)	DataFlux Data Management Studio or DataFlux Data Management Server (all versions)	The SAS Deployment Wizard installs the products into SASHome on the target machine. Any metadata in the SAS Metadata Repository is updated. Any configuration files or DataFlux jobs that are in the source environment must be migrated manually to the target environment. See the DataFlux Migration Guide for instructions about migrating this content.
DataFlux Federation Server 2.1	DataFlux Federation Server 3.1	Migration is manual. See the administration documentation for DataFlux Federation Server 3.1.
DataFlux Authentication Server 2.1	DataFlux Authentication Server 3.1	Migration is manual. See the administration documentation for DataFlux Authentication Server 3.1.

Migrating to SAS 9.4

The following table summarizes the upgrade paths when the target system is SAS 9.4.

Existing Version	Upgrade Version	Migration Instructions
SAS 9.4 offerings that include DataFlux Data Management Studio and DataFlux Data Management Server (all versions)	SAS 9.4 offerings that include DataFlux Data Management Studio and DataFlux Data Management Server (all versions)	The SAS Deployment Wizard installs the products into SASHome on the target machine. The metadata in the SAS Metadata Repository is updated. Any configuration files or DataFlux jobs that are in the source environment must be migrated manually to the target environment. See the DataFlux Migration Guide for instructions about migrating this content.
Any SAS 9.3 installation of DataFlux Data Management Studio and DataFlux Data Management Server (all versions)	SAS 9.4 offerings that include DataFlux Data Management Studio and DataFlux Data Management Server (all versions)	The SAS Deployment Wizard installs the products into SASHome on the target machine. Any metadata in the SAS Metadata Repository is updated. Any configuration files or DataFlux jobs that are in the source environment must be migrated manually to the target environment. See the DataFlux Migration Guide for instructions about migrating this content.
DataFlux Federation Server 3.*	SAS Federation Server (all versions)	Migration is manual. See the administration documentation for SAS Federation Server 3.2.
SAS Federation Server 4.1	SAS Federation Server 4.2	Migration is automated. Use the SAS 9.4 Deployment Wizard to install the SAS offering that includes 4.2.
SAS Federation Server 4.2	SAS Federation Server 4.2	Migration is automated. Use the SAS 9.4 Deployment Wizard to install the SAS offering that includes 4.2.

Existing Version	Upgrade Version	Migration Instructions
DataFlux Authentication Server 3.*	DataFlux Authentication Server (all versions)	Migration is manual. See the administration documentation for DataFlux Authentication Server 3.2.
DataFlux Authentication Server 4.1	DataFlux Authentication Server 4.1	Migration is automated. Use the SAS 9.4 Deployment Wizard to install the SAS offering that includes 4.1.
SAS Visual Process Orchestration Server (all versions)	SAS Visual Process Orchestration Server (all versions)	Migration is manual. See the DataFlux Migration Guide for instructions about migrating this content.
Any SAS 9.3 installation of DataFlux Web Studio Server (all versions)	SAS 9.4 offerings that include DataFlux Web Studio Server (all versions)	The SAS Deployment Wizard installs the products into SASHome on the target machine. There is no metadata in the SAS Metadata Repository. Any configuration files or jobs that are in the source environment must be migrated manually to the target environment. See the DataFlux Migration Guide for instructions about migrating this content.

Existing Version	Upgrade Version	Migration Instructions
Any SAS 9.4 installation of DataFlux Web Studio Server (all versions)	SAS 9.4 offerings that include DataFlux Web Studio Server (all versions)	The SAS Deployment Wizard installs the products into SASHome on the target machine. Prior to Data Management Platform 2.7, there was no metadata in the SAS Metadata Repository. With the 2.7 release, users must manually register the DataFlux Web Studio Server in the SAS Metadata Repository. This data is automatically migrated but should be reviewed by the SAS administrator. Any configuration files or jobs that are in the source environment must be migrated manually to the target environment. See the DataFlux Migration Guide for instructions about migrating this content.

System Requirements

SAS Data Management Studio is licensed as part of SAS Data Management offerings, which include SAS Data Quality Desktop, SAS MDM Advanced, and SAS Financial Crimes Monitor. You can review the system requirements for your product at: http://support.sas.com/installcenter.

Select **SAS 9.4** under the Current Releases section to open the search window for the install center documentation. Enter `Data Management Studio` into the **Search only SAS 9.4 Install Center Documentation** text box. A results page appears with links to the system requirements by offering.

2

Configuring DataFlux Data Management Studio

About Configuration Files

When the software starts, it determines which configuration options are in effect by reading a series of configuration files, by looking in the environment, and by reading the command line. If two settings of the same name exist in different configuration settings, then the order in which the settings are read determines which value is used. The last value read is used as the configuration setting.

DataFlux Data Management Studio reads configuration settings in this order:

1   The app.cfg file in the etc folder where the software is installed.

2   The app.cfg file in a user folder, such as drive:\Documents and Settings \USERNAME\Application Data\DataFlux\DataManagement\VERSION.

3   The application-specific configuration files in the etc folder, such as ui.cfg or dmserver.cfg.

4   The application-specific configuration files in a user folder.

5   The macros folder in the etc folder. The default path to the macros folder can be overridden with BASE/MACROS_PATH setting in the preceding configuration files.

6   The macros folder in a user folder.

7   The environment variables.

8   The command-line options if applicable.

DataFlux Folder Permissions for Users

Users must be able to access files that are associated with DataFlux Data Management Studio. At some sites, an administrator might have to manually set the following permissions for individual users.

Table 2.1   *Recommended File System Permissions for DataFlux Data Management Studio Users*

Folders	Recommended Permissions
Files in the DataFlux folder in a user's private storage area, under Microsoft Windows. Example path: `C:\Users\[user_ID]\DataFlux`	Full control for all files in this private folder.
`DMSTUDIO_HOME\[instance]`	A minimum of Read, Execute, and List Folder Contents for all files and folders in the [instance] folder of the DataFlux Data Management Studio installation directory.
`DMSTUDIO_HOME\[instance]\etc`	Full control for all files and folders in the **etc** folder of the DataFlux Data Management Studio installation directory.
Any folder specified in a DataFlux repository definition.	Full control for all files and folders in any **Data storage** folder and **File storage** folder specified in a DataFlux repository definition. For more information, see "Understanding Repository Definitions" in the *DataFlux Data Management Studio User Guide*.

Configuration Options

About Options

Except as noted, the following configuration options can be set for both DataFlux Data Management Studio and DataFlux Data Management Server. These options are typically set in the app.cfg file.

Note: When editing a configuration file, save it as an ASCII or UTF-8 file. If it is saved as a Unicode file, the software can fail to start and will log this error message: "Read failed due to missing file or possibly invalid content in 'c:\path\file.cfg.'"

Base or General Application Options

DEFAULT_UMASK

(Optional) On DataFlux Data Management UNIX server hosts, this specifies the default `umask` value to use when creating output files, such as the output files from jobs. The `umask` value must be numeric. If not set, the shell's `umask` value is used. Example: default_umask = 002.

ODBC_INI

(Optional) For DataFlux Data Management UNIX server hosts, this overrides the location of the odbc.ini file.

BASE/APP_CONTAINER_DOMAIN

(Optional) This identifies the authentication domain expected by the application container services. If it is not specified, the `DefaultAuth` value is used.

BASE/APP_CONTAINER_LOC

(Optional) This identifies where to locate the application container services. In most cases this is not required. If it is required, the value is typically an HTTP URI. In addition, app.cfg should always set this option to point to the metadata server (IOM://<host><port>).

BASE/APP_VER

(Optional) The application version number.

BASE/AUTH_DEFAULT_DOMAIN

(Optional) This identifies the default for resolving the identity domain.

In a metadata configuration, it is possible for the authenticated credentials to resolve the identity of a person who has multiple logins that use the same user ID. When more than one login matches the resolved user ID, the authentication domain for the presented credentials is determined in the following order:

■ BASE/AUTH_DEFAULT_DOMAIN — If the value of this option is specified and the specified value matches the authentication domain of one of the logins, then this is used. If this option is not specified, and if the specified value is not an authentication domain for one of the logins, the next rule is executed.

- DefaultAuth — If DefaultAuth matches the domain of one of the logins, the presented credentials are associated with the DefaultAuth domain. If no match is found, continue to the next rule.

- The first login — The authentication domain of the first login containing a user ID that matches the presented credentials is used.

BASE/AUTH_SEARCH_DOMAIN

(Optional) This specifies whether to search through all domain credentials. When set to `true`, the set of credentials associated with the current domain is searched to find a set that is usable by the given authenticated user. If multiple credentials are found, the first complete set is used.

BASE/AUTH_SERVER_LOC

(Optional) This specifies the location of the authenticating server and is typically specified during installation in the app.cfg file. If an authenticating server connection is attempted, but this option is not set, localhost:21030 is used as the host and port.

The basic format of the authentication server IOM URI is `iom://<host>:<port>`, where <host> is the name of the computer executing the authentication server and <port> is the port to contact the authentication server. If the authenticating server is a DataFlux Authentication Server, then the port should be specified as 21030 unless the default server has been changed.

If the authenticating server is a SAS Metadata Server, then the port should be 8561 unless the default server has been changed.

If your SAS Metadata Server is installed as a cluster of multiple hosts, then this option points to the metadata cluster profile. The pathname used by the option is the physical location that is recognized by the operating environment. The file must be accessible from the current process. The following example illustrates a typical value: `BASE/AUTH_SERVER_LOC=c:\SAS\Config \Lev1\metadataConfig.xml` The example includes the default name and location of the metadata cluster profile.

For information about valid encodings, see the *SAS National Language Support (NLS): Reference Guide*.

BASE/AUTH_SERVER_PASS

(Optional) This identifies the password to present when connecting to the authenticating server as defined by BASE/AUTH_SERVER_LOC

Typically, the BASE/AUTH_SERVER_LOC is specified during installation in the app.cfg file. Separating the user name (BASE/AUTH_SERVER_USER) and password (BASE/AUTH_SERVER_PASS) from the authenticating server location (BASE/AUTH_SERVER_LOC) enables you to run the batch command (dmpexec) with the authenticate option (-a) with individual credentials. For more information see "Running Jobs from the Command Line" in the *DataFlux Data Management Studio: User's Guide*.

(Optional) BASE/AUTH_SERVER_REPOSNAME

This identifies the name of the authenticating server metadata repository containing information to be searched when items are looked up in the metadata.

BASE/AUTH_SERVER_SSPI

(Optional) This identifies support for Security Support Provider Interface (SSPI). It enables support for single-sign on (SSO) to the SAS Metadata Server. The default is NO (no support for single-sign on). Specify `YES` to enable single-sign-on connections to the SAS Metadata Server from DataFlux Data Management Studio. Add this option to a user's app.cfg file or the default app.cfg file for DataFlux Data Management Studio. Do not set this option in the configuration files for DataFlux Data Management Servers.

BASE/AUTH_SERVER_USER

(Optional) This identifies the user name to present when connecting to the authenticating server as defined by BASE/AUTH_SERVER_LOC.

Typically, the BASE/AUTH_SERVER_LOC is specified during installation in the app.cfg file. Separating the user name (BASE/AUTH_SERVER_USER) and password (BASE/AUTH_SERVER_PASS) from the authenticating server location (BASE/AUTH_SERVER_LOC) enables you to run the batch command (dmpexec) with the authenticate option (-a) with individual credentials. For more information see "Running Jobs from the Command Line" in the *DataFlux Data Management Studio: User's Guide*.

BASE/DMSTUDIO

(Optional) This indicates whether to use the dmstudio process. If this option is set to `true` (value of 1), it indicates that the application is using the dmstudio process (not processes started by the application, like batch.cfg, for example). The user should not adjust or override this value.

BASE/DATE_FORMAT

(Optional) Specifies how to format dates. Possible values, if specified, are `iso8601` and `iso8601_no_milli`. The later does not include the millisecond portion of the ISO8601 standard.

BASE/EMAILCMD

(Required). This option specifies the command that is used to send email. The command can include %T and %B where %T is replaced with the recipient and %B is a file containing the body of the message. Monitor event, as well as the architect node, use this option.

BASE/EXE_PATH

(Optional) This is the path containing executables. The user should never need to set this.

BASE/FTPGETCMD

(Required) This is the command used for FTP Get Functionality. This should default in the install, as follows:

- %U: Replace with user name.

- %P: Replace with password.

- %S: Replace with server.

- %T: Replace with local directory.

- %F: Replace with Files to download, multiple separated by spaces.

- %L: Replace with the log file to pipe the output.

BASE/FTPPUTCMD

(Required) This specifies the command used for FTP Put Functionality. This should default in the install the same as BASE/FTPGETCMD.

BASE/HIGH_PRECISION_TIMER

(Optional) This is a DAC log setting to allow higher precision timing at the expense of forcing all threads to CPU core 1. This is available only on Windows and can decrease multithreaded performance.

BASE/JOB_LOG_ENCODING

(Optional) This specifies the encoding for the job log on a DataFlux Data Management Server.

By default, the log is written in the encoding associated with the locale of the process for the executed job. For English-speaking organizations, this might be LATIN-1 or UTF-8. If a log line contains characters that cannot be represented in the encoding, the log line is not written to the log file. This option enables you to assign the encoding of the job log.

Note: This option must be set on the DataFlux Data Management Server where jobs are executed. It has no effect on DataFlux Data Management Studio job logs.

BASE/LIBRARY_PATH

(Optional) This is the path in which the Java JAR dependencies reside. The path is automatically determined (DFBCO_EXE_PATH/lib) by start-up code, so you should not need this setting in a configuration file.

BASE/LOGCONFIG_PATH

(Optional) This is the full path to the log configuration file. If set, this must be set in the application executable (dis.cfg, and so on) configuration file, not the app.cfg, or macros.cfg files. If it is not set, it defaults to logging.xml in the `etc` directory.

BASE/LOGEXCEPTIONS

(Optional) This is used for Exception logging and defaults to `off`. This must be set before starting the application of interest, because this setting is read only at start-up.

Set this option to either `1` or a combination of letters where: M = do not show module info, V = turn verbose on, U = install the Unhandled Exception Filter, C = install the continue Exception Filter, and f = do not install the first chance exception filter. Setting it to `1` shows all the information.

Note: This only works from DataFlux Data Management Studio, not DataFlux Data Management Server.

BASE/MACROS_PATH

(Optional) This is the path for system macros.cfg file. If not specified, this file is located in the `etc` subfolder of the installation folder.

BASE/MAINTAIN_GROUP

(Optional) This specifies whether to maintain the current execution group or job. When present in the configuration, the current process is queried to determine whether it is already in a group or job. If it is, the child processes stays in the existing group instead of being associated with a new group or job.

Note: If you use remote access software such as Citrix or Microsoft RemoteApps to run DataFlux Data Management Studio as a published application (not streamed), and jobs are failing due to a "CreateProcess() Access Denied" error, you should set the BASE/MAINTAIN_GROUP=YES option in the app.cfg file on the computer where the DataFlux Data Management Studio jobs are executed. If you do not set this option, the jobs will not execute.

BASE/MESSAGE_LOCALE

(Optional) This specifies which error message locale to use. If not specified, it is determined from the system locale.

BASE/MESSAGE_LEVEL

(Optional) This specifies the error level of messages. Possible values include `0` (or not specified) for normal messages or `1`, which includes source file and line number in messages.

BASE/MESSAGE_PATH

(Optional) This specifies the path to the message directory. This setting is not needed in a configuration file because it is automatically determined by start-up code.

BASE/MONITOR_FREQUENCY

(Optional) Disabled by default, this enables the logging of job node statistics while a job is running. If this option is disabled (or its value is `-1`), then node statistics are logged only when the job has finished executing.

If you want to log node statistics while the job is running, specify the number of milliseconds that the software should wait between logging statistics. The higher the frequency, the more run-time details about node execution are logged in to the job's log file. However, the additional collection and logging of information affects the job's performance.

Note: To enable monitor reporting, set the monitoring logger LOG_PREFIX_MONITORING. This provides monitoring information at job completion.

BASE/PLUGIN_PATH

(Optional) This is the path used by all subsystems to find plug-ins. This setting is not needed in a configuration file because it is automatically determined by start-up code.

BASE/PRIMARY_LICENSE

(Required) This is the primary licensing method. This needs to be set in the configuration file with a value of DATAFLUX or SAS.

BASE/PRIMARY_LICENSE_LOC

(Required) This is the server or file location for the primary licensing method.

BASE/REPOS_SYS_PATH

(Optional) This is the system path for repository configuration files and should be automatically determined by dfcurver.

BASE/REPOS_USER_PATH

(Optional) This is the user directory containing repository configuration files. It should be automatically determined by dfcurver.

BASE/REPOS_FILE_ROOT

(Optional) This overrides the root of the repository for URI lookups. If specified, this is used as the root for the repository when resolving the URI. The path in the URI is concatenated to this path to give the actual filename of a URI.

BASE/REPOS_EVENT_WAIT_QUERYMS

(Optional) This is the repository event processing wait time between processing queries. It specifies how frequently in milliseconds to query the repository for changes in the event table. This might need to be changed due to slow servers or IT

issues. This is an overriding value and the default is used if no value is set by the user. A setting of `-1` disables events from client.

BASE/REPOS_EVENT_CLEAN_TIMEMIN

(Optional) This sets the repository event processor that removes all events older than X minutes before start-up.

BASE/ROWSET_SIZE

(Optional) This sets the suggested RowSet size. If specified, the value is used to calculate the maximum number of rows each rowset collection should contain.

BASE/SECONDARY_LICENSE

(Required) This is the secondary licensing method. This must be set in the configuration file as either DATAFLUX or SAS.

BASE/SECONDARY_LICENSE_LOC

(Required) This is the location of the secondary license file or server and must be set in the configuration file.

BASE/SHARED_OUTPUT_PATH

(Optional) This is the directory in which to write shared output files. If set, some log and status files write files to this directory. If not set, it defaults to the value specified in "BASE/TEMP".

BASE/SORTBYTES

(Optional) This specifies the bytes to use when sorting.

BASE/SORTMERGES

(Optional) This specifies whether to merge when sorting.

BASE/SORTTEMP

(Optional) This specifies the temporary file location to be used during the sort. This can contain multiple directories separated by the appropriate host path separator character (':' on UNIX and ';' on Windows).

Note: If the SAS sort is being used, this option's value is not used.

BASE/SORTTHREADS

(Optional) This specifies the number of sort threads.

BASE/SORT_KEEPVAR

(Optional) This specifies the short version of the global temporary record width conversion option. It enables a plug-in to locally set the BASE/SORT_KEEPVAR global option.

When this Boolean value is set to `true`, the temporary data file honors the variable-width record indicator at temporary data file creation time. When the value is set to `false` (default), the temporary data file sort support converts a variable-width file to a fixed-width file, if the record does not contain any string fields, or if the lengths of the string fields in a record are within the overhead threshold necessary to sort variable-width records. Set to `true` to mimic behavior prior to the 2.4 release.

Note: This is a significantly advanced parameter that should rarely be manipulated.

BASE/TEMP

(Optional) This is the directory in which to write temporary files. If not specified, the value of the TEMP/TMP environment variable value is used.

BASE/TEXTMINE_LITI_LANG_LOCATION

(Optional) This is a document extraction node option to specify the install location of Teragram text mining LITI Language file directory. This allows the user to find the LITI language files installed and present the user with the list of provided languages and LITI files for usage in contextual extraction. This option overwrites the default location of `DFEXEC_HOME/share/liti`.

BASE/TIME_BASE

(Optional) This specifies whether to use Greenwich Mean Time (GMT). Valid values are `GMT` and `LOCAL`, which is the default. If this is set to GMT, the current date returns in GMT instead of the localized time. This affects anything that uses the current date timestamp.

BASE/UPDATE_LEVEL

(Optional) This specifies the application update level. Defaults to 0. It can be used as a minor revision number.

BASE/USE_SASSORT

(Optional) This indicates whether the SAS sort is to be used. A Boolean is set to `True` to indicate that the SAS sort is to be used. If not specified or set to `False`, then the DataFlux Data Management sort is used.

BASE/USER_PATH
> (Optional) In the GUI application, this represents the path where user configuration files are stored. This should be automatically determined by dfcurver.

Data Access Component Logging Options

DAC/DISABLESYSCATENUM
> (Optional) This sets the enumeration of SYSCAT DSNs. When set to `yes`, `1`, or `true`, this setting disables the listing of the SYSCAT type DSNs into DSNs that are on that server.

DAC/DFTKDISABLECEDA
> (Optional) This disables CEDA support. Any value turns it on.

DAC/DFTKLOGFILE
> (Optional) This sets the filename for the DFTK logging.

DAC/DFTK_PROCESS
> (Optional) This specifies whether to run DFTK out of process. The default is `off`. Any value turns it on.

DAC/DFTK_PROCESS_TKPATH
> (Optional) This specifies the TKTS path for DFTK when it is running out of process. This path defaults to a `core/sasext` directory off of the executable directory.

DAC/DSN
> (Optional) This specifies the DSN directory for TKTS. The path defaults to `DFEXEC_HOME/etc/dftkdsn`.

DAC/SAVEDCONNSYSTEM
> (Optional) This specifies the location of the system's saved connections. The path defaults to `DFEXEC_HOME/etc/dsn`.

DAC/SAVEDCONNUSER
> (Optional) This specifies the location of the user's-saved connections. The path defaults to the user's application `directory/DataFlux/dac/9.0`.

DAC/TKTSLOGFILE
> (Optional) This sets the filename for the TKTS logging.

DAC/TKTSLOGFLAGS

(Optional) This sets the TKTS logging flags. Enter a string containing a semi-colon delimited list of items that can include the following:

- PARAM_DATA

- INCOL_DATA

- OUTCOL_DATA

- TIMESTAMP

- GETDATA

- PUTDATA

- APPEND

- NOFLUSH

Monitor Options

MONITOR/DUMP_JOB_DIR

(Optional) This specifies a directory to store temporary jobs created by the business rule monitor so that they can be debugged. By default, this option is not set and the monitor does not store temporary jobs on disk.

MONITOR/BULK_ROW_SIZE

(Optional) This specifies the size of the bulk packet. The default value is `1000`. This affects row-logging events, as well as historic custom and standard metrics. This value can be changed to enhance the performance of jobs that monitor business rules when those jobs include row-logging events.

Profile Options

PROF/DEBUG_MODE

(Optional) This sets the value for the frequency distribution (FRED) engine debug mode. Possible values include `0` (not debug mode) or `1` (debug mode). The default is `0`.

PROF/LOCK_RETRIES

(Optional) This specifies the number of times to retry SQLite repository connection when a connect attempt times out. A value of `-1` retries until a connection is established.

PROF/PER_TABLE_BYTES

(Optional) This sets the value for the frequency distribution engine per table bytes. You can use any numeric value. Default is `-1` (frequency distribution engine default).

PROFCONFIG/MAXFREQCACHE

(Optional) This specifies the amount of memory in kilobytes that the frequency distribution engine uses. If not specified, the default is `128kb`.

QKB Options

QKB/ALLOW_INCOMPAT

(Optional) This allows data jobs to run even when the software detects that a QKB definition invoked by the job was saved by a version of the software that is later than the current version of the software. Default is `NO`. The default behavior is for these definitions to fail to load. Results obtained when this option is turned on are undefined.

QKB/COMPATVER

(Optional) This tells DataFlux Data Management Studio which version of the QKB to use when running a data job. Possible values include `dfpower82`, `unity21`, `unity22`, or `dmp21` through `dmp26`. If not specified, it defaults to the latest version. You can use this option if you want to use the latest version of DataFlux Data Management Studio, but want the outputs of your QKB-related Data Job nodes (for example, matchcodes) to be exactly the same as the outputs for earlier versions.

QKB/ON_DEMAND

(Optional) This loads QKB definitions on demand. Default is `YES`. The application start-up creates a QKB pool that sets the option for all consumers (Profile, Explorer, and Nodes) with the exception of the Expression Engine, which has its own initialization. Set this option to `No` to find errors within definitions and to see error messages specific to Pattern Logic nodes.

If used for a data job, this option allows QKBs to load definitions as needed instead of loading all of them at the beginning. These definitions are kept in memory for future runs that reuse that process.

If used for a service, each loaded service loads its own QKB. Once a service is loaded into memory and runs at least once, it keeps the QKB loaded from the previous run and does not have to be loaded again. Note that a QKB in memory is not shared across different services or threads, so each initiation of either a new service or a new thread for an existing service causes the QKB to be loaded. This could have an implication on memory and performance.

QKB/PATH

(Required by QKB products.) This is the path to QKB. The path is set to the default QKB as defined in the application.

QKB/SURFACEALL

(Optional) This surfaces all parse definitions in the Data Job interface, even definitions for which the "Surface" flag is unchecked in Customize. The default is `NO`.

Note: The application start-up creates a QKB pool that sets the option for all consumers (Profile, Explorer, and Nodes) with the exception of the Expression Engine that continues to have its own initialization.

Architect Client (UI) Settings Option

ARCHITECT/AutoPassThru

(Optional) This provides client options to set mappings. Maintained by the client, choices are `0` (target), `1` (Source and Target), and `2` (All). This determines whether data flow nodes should automatically fill their pass-through list from their parent nodes.

Architect and Miscellaneous Node Options (defined in ids.h)

CLUSTER/BYTES

(Optional) This specifies the bytes to use when clustering.

CLUSTER/LOG
   (Optional) This specifies whether a clustering log is needed.

CLUSTER/TEMP
   (Optional) This specifies the cluster temporary path.

FRED/LOG
   (Optional) This specifies whether a FRED log is needed.

JAVA/CLASSPATHS
   (Optional) This specifies the Java classpath. This overrides the Java class path for the Java plug-in node. You generally should not need to set this.

JAVA/DEBUG
   (Optional) This specifies the Java debug options.

JAVA/DEBUGPORT
   (Optional) This specifies the port to remotely debug Java.

JAVA/VMHEAP_INIT
   (Optional) This specifies the initial size of the JVM in megabytes. This should be used if you are using JVM 1.6. Start at 16 and decrease if you keep getting errors.

JAVA/VMHEAP_MAX
   (Optional) this specifies the maximum amount of heap space the JVM uses in megabytes. This should be used if you are using JVM 1.6. You can decrease this along with JAVA/VMHEAP_INIT.

VERIFY/CACHESIZE
   (Optional) This specifies the cache size by percentage.

VERIFY/CANADA
   (Required by SERP nodes) This specifies the path to the Canadian data. The Canada installation should maintain this.

VERIFY/GEO
   (Required by Geocode) This specifies the location of the Geocode/Phone data. The Geocode installation should maintain this.

VERIFY/PRELOAD

(Optional) This preloads defined state data into memory for address verification. It is set in the app.cfg file. Valid value can be `ALL`.

VERIFY/USEDPV

(Optional) This enables or disables Delivery Point Value (DPV) processing to confirm that the address actually exists.

VERIFY/USEELOT

(Optional) This enables or disables Enhanced Line of Travel (eLOT) processing, which enables mailers to sort their mailings in approximate carrier-casing sequence using an eLOT sequence number and to qualify for enhanced carrier route presort discounts.

VERIFY/USELACS

(Optional) This enables or disables Locatable Address Conversion System (LACS) processing that enables mailers to identify and convert a rural-route address to a city-style address.

VERIFY/USERDI

(Optional) This enables or disables Residential Delivery Indicator (RDI) processing that confirms whether a particular address is residential or business.

VERIFY/USPS

(Required by the USPS) This specifies the USPS data path. This is maintained by the USPS installation.

VERIFY/USPSINST

(Required) This determines whether the USPS data is installed or if sample data is being used. This is maintained by the USPS installation.

VERIFYINTL/CFG

(Required by international verification) This is the address verification configuration file. The component's installation should maintain the path.

VERIFYWORLD/DB

(Required for Platon) This specifies the Platon library universal unlock code. The component's installation should maintain the path.

VERIFYWORLD/UNLK

(Required for Platon) This specifies the Platon data path. The component's installation should maintain the path.

VERLQT/COUNTRY_INIT_FILTER

(Optional) This is a comma-separated list of two-letter country codes. This limits the countries from which the user can choose and for which Loqate can load data.

VERLQT/LOQATE_DATA_PATH

(Required) This is the path to the Loqate data installation folder.

VERLQT/LOQATE_LOG_PATH

(Optional) This is the path for the Loqate log file. This is the log generated by the Loqate engine, not by the application code. It can be useful for debugging.

VERLQT/MAX_REF_PAGE_CACHE_SIZE

(Optional) This specifies the maximum reference data cache size. Default is `300000`. Larger values allow for more referenced data to be cached, which can improve performance. However, this also increases the amount of memory your process attempts to consume. If you encounter out-of-memory errors, try lowering this value.

VERLQT/PRELOADTABLELIST

(Optional) This is a comma-delimited list of country data table names that can be found in the Loqate data folder. Use "*" to load all countries. This is intended to improve performance by pre-loading data for specific countries. Note that this can use a very large amount of memory.

Address Update (NCOA) Options (in dfncoa_appcfg.h)

NCOA/DFAV_CACHE_SIZE

(Optional) This sets the verify cache percentage for the USPS data. The higher the value, the more data is cached. The faster the processing, the more memory is used. The default is `0`. It resides in macros/ncoa.cfg.

NCOA/DFAV_PRELOAD

(Optional) This sets the verify preload options for the USPS data. Valid values are `ALL` or an empty string. Using `ALL` requires a large amount of memory. It resides in macros/ncoa.cfg.

NCOA/DVDPATH

(Required) This is the path to the unpacked and unzipped NCOA data. It resides in macros/ncoa.cfg.

NCOA/QKBPARSEDEFN

(Optional) This is the path to the QKB parse definition used for address updates. The default is `Name (Address Update)`, and it resides in macros/ncoa.cfg.

NCOA/QKBPATH

(Required) This is the path to the QKB used for Address Update name parsing. It resides in macros/ncoa.cfg.

NCOA/REPOSCONNECTION

(Required) This specifies the connection string used to connect to the Address Update repository. This overrides NCOA/REPOSDSN, and one or the other is required. This is typically set by the Address Update Admin utility and resides in the app.cfg.

NCOA/REPOSDSN

(Required) This specifies the DSN used to connect to the Address Update repository. It is overridden by NCOA/REPOSCONNECTION, and one or the other is required. This is typically set by the Address Update Administration utility and resides in the app.cfg.

NCOA/REPOSPREFIX

(Required) This is the table prefix used on the Address Update tables. This is typically set by the Address Update Administration utility and resides in the app.cfg.

NCOA/REPOSTYPE

(Required) This specifies the repository type. Valid values are: `0` (Guess), `1` (ODBC), or `2` (DFTK). If the value is `0`, the node attempts to determine the type from the connection string. This is typically set by the Address Update Administration utility and resides in the app.cfg.

NCOA/USPSPATH

(Required) This is the path to the USPS `CASS/DPV/etc` data. It resides in macros/ncoa.cfg.

IntelliServer Option

DFCLIENT/CFG

(Required) This is used for IntelliServer. It is maintained by the dfIntelliServer installation, which is typically located at `C:\Program Files\DataFlux \dfIntelliServer\etc\dfclient.cfg`. Modify the dfclient.cfg file to point to the server and port.

Pooling Options

For puddle options, the name of the puddle is placed after 'POOLING/' (for example, POOLING/WFEJOB/MAXIMUM_PROCESSES). If no puddle name is specified, it is globally applied to all puddles. Here are some examples of puddles: WFEJOB - batch jobs on DMServer, WFESVC - Process services on DMSERVER, and APISVC - DFAPI services.

POOLING/CHILD_MAXIMUM_LAUNCHES

(Optional) This specifies the throttling for launches. When specified, the number of concurrent child process launches is limited by this value. If the current child launch request exceeds the specified value, the launch waits until the number of launching processes is below the specified value. If `zero` or `not` is specified, no limit exists for concurrent child launches.

POOLING/GET_PROCESS_TIMEOUT

(Optional) This specifies the length of time, in seconds, the process requester should wait for a process to become available before it times out. If zero, the requester waits indefinitely. The acquire process time-out is in terms of the acquisition of a process and the process pooling handshaking. It does not consider the time required by the requester to complete application-level initialization. Default: `no time-out`.

POOLING/IDLE_TIMEOUT
> (Optional) This specifies the length of time, in seconds, the process remains idle before it is terminated. If zero, idle processes are not terminated. Default: `0`. This is a puddle option.

POOLING/MAXIMUM_ERRORS
> (Optional) This specifies the maximum number of pooled process errors before the process is terminated. The default is `0` (never terminate it).

POOLING/MAXIMUM_PROCESSES
> (Optional) This specifies the maximum number of concurrent pooled processes. If `0`, the number of concurrent pooled processes is unlimited. Default is unlimited after `2.2` (previously was `10`). If POOLING/GET_PROCESS_TIMEOUT is set, it waits for that amount of time to get a new process if it needs one.

POOLING/MAXIMUM_USE
> (Optional) This specifies the maximum number of pooled process uses before the process is terminated. The default is `0` (unlimited).

Repository Option

REPOS/CREATE_SPEC_PATH
> (Optional) This specifies the path to the RepositoryTableCreationXML file containing the custom SQL that configures the commands to create tables and indexes in the repository.

REPOS/FORCE_FILE_BASED
> (Optional) This specifies how to create the repository table or index. It provides a means of configuring the commands that are used to create tables and indexes in the repository.

REPOS/TABLE_LIST_PATH
> (Optional) This defines the repository XML table. The directory should contain XML files for any tables that the repository library should add during creation or update. If set, this directory contains the XML files; if not set, look in DFEXEC_HOME/etc/reposcreate.

Other Options

EXPRESS_MAX_STRING_LENGTH

(Optional) This specifies the maximum size of strings declared in the expression nodes. The default maximum length of any string in this node is 5,242,880 bytes (5MB). This enables specifying a larger value in bytes. If performance issues arise, then the suggested setting is 65536 bytes.

EXPRESSION/UDFDIR

(Optional) This specifies where to look for UDF files. If not specified, look for UDF files in `installationdir/etc/udf`.

JAVA/COMMAND

(Optional) This is the command used to launch the Java proxy process. The default is Java. The Java command must be compatible with launching from the command line. Here are some examples:

- JAVA/COMMAND = Java

- JAVA/COMMAND = Java -Djavax.net.ssl.trustStore=C:\Store\jssecacerts

- JAVA/COMMAND = Java -Djavax.net.ssl.trustStore="C:\Cert Store\jssecacerts"

- JAVA/COMMAND = "C:\Program Files\Java\jre6\bin\java"

MDM/REPOSITORY_ROOT_FOLDER

(Optional) This allows the `foundations/master_data` to be overwritten by the user when putting the contents of `[INSTALL_ROOT]/share/mdm` into a repository.

Data Access Component Directives

The Data Access Component (DAC) enables you to connect to data using Open Database Connectivity (ODBC) and Threaded Kernel Table Services (TKTS). ODBC database source names (DSNs) are not managed by the DAC, but by the Microsoft ODBC Administrator. However, TKTS DSNs are managed by the DAC, and TKTS connections are stored in a TKTS DSN directory. The DAC reads a list of settings to

determine what actions to take, where to find various things that it needs, and how to make the proper connections.

Note: The default DAC directives should be satisfactory for most sites. Change these settings only if you have special needs.

There are three types of DAC Configuration Directives:

configuration options
> Most DAC settings come from configuration values specified in app.cfg or other configuration files that are read in a specific sequence. It is recommended to put these settings in the app.cfg file, but they can be put in macro.cfg or another application-specific configuration file. See "Configuration Options" on page 9 for a more comprehensive list of possible configuration values.

registry
> The DAC looks in the registry for DSN-specific items. These are found in either or both HKEY_CURRENT_USER\Software\DataFlux\dac\{version}\{dsn} or HKEY_LOCAL_MACHINE\Software\DataFlux\dac\{version}\{dsn}. In this case, "dsn" is a string with the same name as the DSN to be affected.

environment variables
> You can identify the environment variables with $NAME in the following table. The application using the DAC typically sets environment variables to appropriate locations. On Windows, $DFEXEC_HOME is set to the application's root directory.

Setting	Description	Windows	UNIX
Command file execution	Specifies a text file with SQL commands (one per line). These commands run in turn on any new connection that is made. For example, they can be used to set session settings. This is implemented only for the ODBC driver.	The DAC/ SAVEDCONNSYSTEM configuration value specifies the path to the saved connections. The DAC checks for files with the same filename as the DSN and a .sql extension.	The same as Windows.

Setting	Description	Windows	UNIX
Connection and login timeout	Enables you to specify the time in seconds for a connection time-out and a login time-out.	The string type value is located in USER *dsn_name* \logintimeout (or connecttimeout) where dsn_name is the name of the DSN. The string value is the number of seconds for the time out. If the string type value is not found in USER\. then the DAC looks in the SYSTEM *dsn_name* \logintimeout.	Not supported.
DAC logging	Determines whether to create a log file for DAC operations.	This is controlled through the effective logging configuration files, which use the SAS logging facility. The logger name "DAC" should be used.	The same as Windows.
DFTK log file	Specifies the log file that interacts with the DFTKSRV layer and is useful only for debugging issues specific to dftksrv. This setting is applicable only if you are running DFTK out of process.	Look for config value DAC/DFTKLOGFILE.	The same as Windows.
Disable CEDA	Specifies whether to disable CEDA. This setting is applicable only to tkts connections.	Look for DAC/ DFTKDISABLECEDA configuration value, which should specify any non-null value (for example, yes).	The same as Windows.

Setting	Description	Windows	UNIX
Multiple active results sets (MARS)	Specifies whether to allow a connection to open multiple result sets at the same time. This is applicable only when using the SQL Server Native Client driver on Windows against a SQL Server 2005 (or later).	Look in USER *dsn_name*\mars for dword value of 1 where *dsn_name* is the name of the DSN. If the DAC does not find it, it looks in the SYSTEM*dsn_name*\mars.	Not applicable.
Oracle NUMBER(38) handling	If connected to Oracle (only), NUMBER(38) columns are treated as INTEGER by default. To override that functionality, treat them as REAL.	Look in USER \oranum38real for dword value of 1. If it is not found in USER, look in SYSTEM \oranum38real.	Look in $HOME/.dfpower/ dsn.cfg for a line dsn_name = oranum38real where *dsn_name* is the name of the DSN.
Read uncommitted	For SQL Server only, specifies that reading data from cursors is allowed to read uncommitted rows. This can alleviate lockups in SQL Server (applies to DMP version 2.3 and later).	Look in USER *dsn_name* \readuncommitted for dword value of 1. Where *dsn_name* is the name of the DSN. After looking in USER, look in SYSTEM.	Look in $HOME/.dfpower/ dsn.cfg for a line dsn_name = readuncommitted where *dsn_name* is the name of the DSN.
Run DFTK out of process	Specifies whether to run TKTS out of process, which enables you to perform troubleshooting.	Look for config value DAC/ DFTK_PROCESS. This should specify any non-null value (for example, yes).	The same as Windows.

Setting	Description	Windows	UNIX
Suffix for CREATE TABLE statements	This enables you to provide a string that is appended to every CREATE TABLE statement. If you include %T in this string, it is substituted with the table name.	Look in USER *dsn_name* \postcreate for a string where *dsn_name* is the name of the DSN. If it is not found in USER, look in SYSTEM.	Not supported.
System saved connection	Specifies where to find system-saved connections.	Look in DAC/ SAVEDCONNSYSTE M config value for a path. If it is not found, the DAC checks for the DFEXEC_HOME environment variable in the $DFEXEC_HOME \etc\dsn subdirectory, which is in the installation directory.	Look for the same config value. If not found, look in the $DFEXEC_HOME/et c/dsn subdirectory, which is in the installation directory.
TK Path	Specifies where threaded kernel files are located. This setting is applicable only if you are running Data Factory Took Kit (DFTK) out of process. The dftksrv path and core directory should be specified.	Look for config value DAC/ DFTK_PROCESS_P ATH. If still not found, set it to $DFEXEC_HOME \bin; $DFEXEC_HOME \bin\core\sasext.	Same as for Windows, but sets it to $DFEXEC_HOME/lib/ tkts if it is not found.
TKTS DSN directory	Specifies the path where TKTS DSNs are stored in XML files.	Look in config value DAC/DSN. If it is not found, look in $DFEXEC_HOME \etc\dftkdsn\.	Look in the same config value. If it is not found, look in $DFEXEC_HOME/et c/dftkdsn/.

Setting	Description	Windows	UNIX
TKTS log file	Specifies the log file that is produced by the TKTS layer and is useful for debugging tkts issues.	Look for config value DAC/TKTSLOGFILE.	The same as Windows.
TKTS startup sleep	Specifies how much time in seconds to delay between the start of the dfktsrv program and the booting of threaded kernel. This setting is applicable only if you are running DFTK out of process.	The DAC checks the registry for a tktssleep value.	Not supported.
User saved connection	Specifies where to find user-saved connections.	Look in DAC/ SAVEDCONNUSER config value for a path. If not found, look in the user's application settings directory, which is usually found in the %APPDATA% directory, in the DataFlux\dac*version* subdirectory. The %APPDATA% location is C:\Users \userid\AppData\ for Windows 7 and C: \Documents and Settings\userid \Application Data\ for Windows XP.	Look for the same config value. If not found, look in the $HOME/.dfpower/dsn directory.

Setting	Description	Windows	UNIX
Use braces	Specifies whether to enclose DSN items in braces when they contain reserved characters. Most drivers do not support use of braces. However, this might be an option if problems occur with saving credentials.	The DAC looks in USER\usebraces for dword value of 1. If it is not found there, the DAC looks in SYSTEM\usebraces.	Look in $HOME/.dfpower/ dsn.cfg for a line dsn_name = usebraces where dsn_name is the name of the DSN.

Logging Options

The following log files are provided for DataFlux Data Management Studio:

- Studio log

- Platform log

- DAC log

- TKTS log

If enabled, the TKTS log captures TKTS events such as BASE data sets and data sent across the wire to SAS Federation Server. The TKTS log provides a default set of information that cannot be customized.

The Studio log, the Platform log, and the DAC log are enabled by default. These logs use the SAS Logging Facility, a flexible and configurable framework that you can use to collect, categorize, and filter events. Then you can write them to a variety of output devices. The logging facility supports problem diagnosis and resolution, performance and capacity management, and auditing and regulatory compliance.

The logging facility framework categorizes and filters log messages in SAS server and SAS programming environments. It also writes log messages to various output devices.

In the server environment, the logging facility logs messages based on predefined message categories such as the following:

- Admin for administrative messages

- App for application messages

- Perf for performance messages

Messages for a category can be written to files, consoles, and other system destinations simultaneously. The logging facility also enables messages to be filtered based on the following thresholds, in order from most inclusive to least inclusive: TRACE, DEBUG, INFO, WARN, ERROR, and FATAL. The configuration files for the Studio log, the Platform log, and the DAC log contain a template for logging options. You need only to uncomment certain lines to turn on the corresponding options. For a full description of these options, see the *SAS Logging: Configuration and Programming Reference*. There are also some DataFlux options for logs. See "Data Access Component Logging Options" on page 19.

The following table specifies the scope, level thresholds, and configuration location for each of these logs:

Name	Scope	Level Thresholds	Configuration Location
Studio	DataFlux Data Management Studio events	TRACE, DEBUG, INFO, WARN, ERROR, and FATAL	By default, these are enabled. To disable, see note below.
Platform	Engine events	TRACE, DEBUG, INFO, WARN, ERROR, and FATAL	By default, these are enabled. To disable, see note below.
DAC	Data access events	TRACE, DEBUG, INFO, WARN, ERROR, and FATAL	By default, these are enabled. To disable, see note below.
TKTS	TKTS events such as BASE data sets and data sent across the wire to SAS Federation Server	Not configurable	`drive:\Program Files\DataFlux \DMStudio \studio1\etc \app.cfg`

Name	Scope	Level Thresholds	Configuration Location

Note: By default, Studio, Platform, and DAC are enabled. To turn them off, comment out the BASE/LOGCONFIGPATH line in ui.cfg and in batch.cfg and restart the application. For configuration, go to `drive:\Program Files\DataFlux\DMStudio\studio1\etc\platform.log.xml` and `drive:\Program Files\DataFlux\DMStudio\studio1\etc\batch.log.xml`.

By default, all Windows XP log files are written to `drive:\Documents and Settings\USERNAME\Application Data\DataFlux\DMStudio\[version]\Logs`. The default location in Windows 7 is `drive:\Users\[username]\AppData\Roaming\DataFlux\DMStudio\[version]\Logs`.

All Studio, Platform, and DAC logs by default are all logged to the platform log file. This log file is prefixed with platform_. The DAC logging can be redirected to its own log file by uncommenting the appender-ref setting in the logger section for DF.DAC and by uncommenting the RollingFileAppender section for DacFile. The latter contains the filename pattern that controls where and what name the DAC log file takes. The Studio logging can be redirected to its own log file by uncommenting the appender-ref setting in the logger section for DF.DMStudio and by uncommenting the RollingFileAppender section for DMStudioFile. The latter contains the filename pattern that controls where and what name the Studio log file takes.

The Studio, Platform, and DAC logging all use a logging level of INFO by default. This level is set in the root logger section of the platform.log.xml configuration file. The level value can be changed here to change the default for all three logs. In addition, the logging level for the DAC and Studio logs can be set separately in the DF.DAC and DF.DMStudio logger sections respectively by uncommenting (and possibly updating) the level values in each.

The platform.log.xml file controls all of Studio, Platform, and DAC logging within the DataFlux Data Management Studio application except for job runs. When you run a job, the application creates a separate batch process for that job run. Therefore, you should modify the batch.log.xml to configure logging for jobs.

By default, the Platform log contains run-time statistics for job nodes. See option for details about these statistics. To

disable all job node statistics from the Platform log, add the following block immediately before the <root> section at the bottom of *.log.xml file:

```
<logger name="DF.Monitor">

<level value="OFF"/>

</logger>
```

3

Installing Add-On Products

Installing Data Packs

Data Packs are collections of third-party enrichment data, such as Loqate international address data, and are not installed as part of the SAS Deployment Wizard process. If you have licensed Data Pack, go to SAS support site to download and install these products. Use the Release Notes to verify that they are compatible with your version of SAS software.

Download Datapacks

After you install a data pack, see the topic for that feature in the *DataFlux Data Management Studio: User's Guide* to enable access to that data pack.

Installing Supplemental Language Support

If you plan to use DataFlux Data Management Studio for data that includes East Asian languages or right-to-left languages, you must install additional language support. To install these packages:

1 Select **Start ▸ Settings ▸ Control Panel**.

2 Double-click **Regional and Language Options**.

3 In the Regional and Language Options dialog box, click the **Languages** tab.

4 Under Supplemental Language Support, select the check boxes for **Install Files for complex script and right-to-left languages (including Thai)** and **Install files for East Asian languages**.

5 The Microsoft Windows installer guides you through the installation of these language packages.

Setup for the Address Update Add-On

Overview

The DataFlux® Address Update add-on enables you to use the United States Postal Service (USPS) NCOALink® system to identify and update customer address information. This section describes how to perform some of the setup tasks for the Address Update Add-on to DataFlux Data Management Studio. The setup tasks

described here are performed outside of the application. Also see the Address Update Add-On chapter in the *DataFlux Data Management Studio User's Guide*, which describes setup tasks required in the interface.

Obtain a License for DataFlux Products and USPS Components

If you are receiving this DataFlux product as part of a SAS® license, the DataFlux license file is provided with your SAS installation media. If you need assistance in obtaining a license, contact your SAS account representative or access the DataFlux Unlock Code page at support.sas.com/contact/additional/dfunlockreq.html.

Contact the USPS to obtain the required licensing and certification.

If you are using the DataFlux Data Management Server to execute address update jobs, you must also have the appropriate DataFlux licenses and USPS licenses on that server.

Install DataFlux Software

Install the following DataFlux software if this has not been done:

- DataFlux Data Management Studio

- DataFlux Quality Knowledge Base (QKB) for Contact Information 2011A

- Address Update Add-on installer for DataFlux Data Management Studio

If you are using the DataFlux Data Management Server to execute address update jobs, you must also install the following when you are ready to use the server.

- DataFlux Data Management Server: See the DataFlux Data Management Server Administrator's Guide.

- Address Update Add-on installer for Data Management Server, which updates Data Management Server to support the Address Update Add-on.

Address Update jobs must be able to access the DataFlux USPS Verify data from the computer where the jobs are executed. The Verify data is available for Microsoft Windows and UNIX systems. The path to this data is specified in the ncoa.cfg file, which

is described later. Be sure to record the installation path to this data. The default installation path for Windows is `C:\Program Files\DataFlux\VerifyData\`.

Install Address Update Data from USPS

Address Update jobs must be able to access the USPS NCOALink® data from the computer where the jobs are executed. If you are using the Data Management Server to execute the address update jobs, you must obtain and license this data from the USPS and install the data so that it is accessible to both the DataFlux Data Management Studio and the Data Management Server. You can copy this information to both computers, or make the information available on a shared file system, as long as you follow the terms of the license. The path to this data is specified in the ncoa.cfg file, which is described later. This data is delivered either on a DVD or electronically from USPS.

Address Update processing does not take place unless you have specific data folders in the expected locations with the correct data files. To ensure that this folder structure is created:

1 Create a folder called "NCOA Data" on your drive.

2 Add subfolders called "DVD 1" and "DVD 2" under the NCOA Data folder.

3 Add subfolders called "CLK" and "RV9" under the DVD 2 folder.

This structure yields the following layout:

> *C:\NCOA Data\DVD 1*
> *C:\NCOA Data\DVD 2\ CLK*
> *C:\NCOA Data\DVD 2\ RV9*

To install Address Update Data:

1 Copy all of the files on DVD 1 into the DVD 1\ directory.

2 Unzip all of the ZIP files on DVD 1 into the DVD 1\ directory.

3 Copy all of the files on DVD 2 into the DVD 2\ directory.

4 Unzip the file CLK.ZIP into the CLK\ directory under the DVD 2\ directory.

5 Unzip the file RV9.ZIP into the RV9\ directory under the DVD 2\ directory.

If you are processing flat files, a sample layout looks like this:

Subfolder	Files	Source
C:\NCOA Data\ DVD 1\	clk.flt.1 clk.flt.2 and so on	
C:\NCOA Data\DVD 2\	clk.a.hf0 clk.b.hf0 and so on	
C:\NCOA Data\DVD 2\ CLK\	business.lst clk.flt.3 and so on	Folder unzipped from CLK.ZIP
C:\NCOA Data\DVD 2\ RV9\	cln.dat cln.txt crd.dat and so on	Folder unzipped from RV9.ZIP

If you are processing Hash files, a sample layout looks like this:

Subfolder	Files	Source
C:\NCOA Data\ DVD 1\	clk.a.hs1	Unpacked from CLK.A.HS1.SZP
	clk.a.hs2	Use the dfAddrHu utility
	clk.a.hs3	
	...	
	clk.z.hs8	hs8 files are unzipped from CLK.*.HS8.ZIP files
	and so on	
C:\NCOA Data\ DVD 2\	clk.a.hs0	
	clk.b.hs0	
	and so on	
C:\NCOA Data\ DVD 2\CLK\		Folder unzipped from CLK.ZIP
	business.lst	
	clk.flt.3	
	and so on	
C:\NCOA Data\ DVD 2\RV9\		Folder unzipped from RV9.ZIP
	cln.dat	
	cln.txt	
	crd.dat	
	and so on	

Note: You must unpack Hash files (.szp). DataFlux has provided a utility called dfAddrHu .

You can access the dfAddrHu utility in the `C:\Program Files > DataFlux > DM Studio > [version number] bin` directory, with the following syntax: *dfAddrHu [<options>] −i [<input>] −o [<output>].*

To run the utility, specify the path location to the DVD 1 folder as both the input and the output, as follows: `dfAddrHu <path to DVD 1 folder> <path to DVD1 folder>`.

For example, you can specify the paths as `dfAddrHu "C:\NCOA Data\DVD 1" "C:\NCOA Data\DVD 1"`.

The utility reads through the DVD 1 folder and unpacks all of the .szp files into the same folder.

Update the ncoa.cfg File

The physical paths for Address Update data, the DataFlux QKB, and USPS data for Verify (CASS, DPV, and LACS) must be specified in the ncoa.cfg file on any computer that executes address update jobs. You should configure the ncoa.cfg file for DataFlux Data Management Studio and confirm that it is working correctly. Then you can copy the ncoa.cfg file to Data Management Server and edit paths if necessary.

For DataFlux Data Management Studio, add the following options to the ncoa.cfg file, located at `C:\Program Files\DataFlux\DMStudio\<version number>\etc \macros`:

NCOA/DVDPATH

Location where the USPS NCOALINK data is installed

Example: ncoa/dvdpath = "C:\NCOA Data\"

NCOA/QKBPATH

The location of the QKB (minimum of QKB for Contact Information 2011A is required)

Example: ncoa/qkbpath = "C:\Program Files\DataFlux\QKB\CI\2011A\"

NCOA/USPSPATH

Location of the DataFlux USPS Verify data (CASS, DPV, LACS)

Example: ncoa/uspspath = "C:\Program Files\DataFlux\VerifyData\data\1105\usps\"

The following optional settings are available for the United States Postal Service (USPS), CASS, and LACS configuration. Set this option for the Address Update node in the ncoa.cfg file. Note that these options do not affect the operation of the Address Verify nodes.

NCOA/DFAV_CACHE_SIZE

Range: 0 through 100 and indicates how much USPS data to cache. The higher the value, the more data is cached. The faster the processing, the more memory used. The default is 0.

NCOA/DFAV_PRELOAD

Set dfav/verify preload options for the USPS data. Provide the names of states to preload, to accelerate address verification for those states. Valid values include:

> "." - do not preload any states (default)
>
> "ALL" - preload all states
>
> "MIL" - preload military addresses only
>
> "CONUSA" - preload the 48 contiguous states
>
> "TX FL" - preload Texas and Florida

Note: Monitor the impact of the NCOA/DFAV_CACHE_SIZE and NCOA/DFAV_PRELOAD options on the available memory on your system. If you run out of memory, you might consider decreasing the amount of memory that you have allocated to these options.

After you have run a sample address update job and verified that the ncoa.cfg file is working properly, copy the ncoa.cfg file to DataFlux Data Management Server and edit paths if necessary (for example, `C:\Program Files\DataFlux\DMServer \<version number>\etc\macros`).

Update the app.cfg File

DataFlux Data Management Studio has a PAF administration dialog box. For instructions about using this dialog box, see the Online Help under "Administering the Address Update Add-On". The PAF administration dialog box updates the app.cfg file.

Descriptions of the options added by the PAF administration dialog box are as follows:

NCOA/REPOSDSN
    DSN connection for the address update repository

NCOA/REPOSPREFIX
    Table prefix for the tables in this repository, if a prefix has been specified

NCOA/REPOSTYPE
    Value that indicates the type of repository:

> 0 = NULL (DataFlux Data Access Component (DAC) attempts to determine the repository type from the connect string)
>
> 1 = ODBC DSN
>
> 2 = Custom DSN

The type depends on how the address update repository was configured in the PAF administration dialog box.

If you are using the DataFlux Data Management Server to execute address update jobs and have run a sample address update job and verified that the app.cfg file is working properly, then copy the app.cfg file to DataFlux Data Management Server and edit paths if necessary (for example, `C:\Program Files\DataFlux\DMServer\<version number>\etc`).

Setup Continues in DataFlux Data Management Studio

See the DataFlux Data Management Studio Online Help Address Update Add-On section for additional information about setup tasks that must be performed.

USPS Copyright Disclosure

DataFlux is a non-exclusive interface distributor of the United States Postal Service and holds a non-exclusive license from the United States Postal Service to publish and sell USPS CASS, DPV, and RDI information. This information is confidential and proprietary to the United States Postal Service. The price of these products is neither established, controlled, or approved by the United States Postal Service.

Installing Support for the SAP RFC Node

If you open a data job in DataFlux Data Management Studio, the **SAP Remote Function Call** node is available in the **Data Integration** folder. This node enables you to read and write data using the SAP RFC and Business Application Programming Interfaces (BAPI) provided by SAP when using the SAP Unicode libraries. This functionality is useful for the following:

■ reading information from SAP that is not readily available in tables

■ using the functionality that SAP exposes as BAPIs

■ loading data to SAP by using functions provided by SAP that are a standard way of updating objects

■ triggering an upload to SAP by using functions that take (for example, an external file input)

In order to use the SAP Remote Function Call node, you must install SAP Unicode libraries (DLLs) on the computer where the job that contains the node is executed. For example, if you use DataFlux Data Management Studio to execute a job with an **SAP Remote Function Call** node, then you must install the SAP libraries on the computer where the software is running. If you execute a job with an SAP Remote Function Call node, then you must install the SAP libraries on the computer where Data Management Server is running.

DataFlux Data Management Studio is a 32-bit Windows application and requires 32-bit SAP libraries. See the SAP documentation for details about installing these libraries.

DataFlux Data Management Servers can be either 32-bit or 64-bit. See the SAP documentation for details about installing the appropriate libraries. Note that for DataFlux Data Management Servers running on UNIX systems, you can set the SAP_LIBPATH option in the app.cfg file to point to the SAP DLLs.

Appendix 1

Legal Notices

Apache Portable Runtime License Disclosure

Copyright © 2008 DataFlux Corporation LLC, Cary, NC USA.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Apache/Xerces Copyright Disclosure

The Apache Software License, Version 3.1

Copyright © 1999-2003 The Apache Software Foundation. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1 Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2   Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3   The end-user documentation included with the redistribution, if any, must include the following acknowledgment: "This product includes software developed by the Apache Software Foundation (http://www.apache.org)." Alternately, this acknowledgment may appear in the software itself, if and wherever such third-party acknowledgments normally appear.

4   The names "Xerces" and "Apache Software Foundation" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact apache@apache.org.

5   Products derived from this software may not be called "Apache", nor may "Apache" appear in their name, without prior written permission of the Apache Software Foundation.

THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the Apache Software Foundation and was originally based on software copyright (c) 1999, International Business Machines, Inc., http://www.ibm.com. For more information on the Apache Software Foundation, please see http://www.apache.org.

Boost Software License Disclosure

Boost Software License - Version 1.0 - August 17, 2003

Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the "Software") to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following:

The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Canada Post Copyright Disclosure

The Data for areas of Canada includes information taken with permission from Canadian authorities, including: © Her Majesty the Queen in Right of Canada, © Queen's Printer for Ontario, © Canada Post Corporation, GeoBase®, © Department of Natural Resources Canada. All rights reserved.

DataDirect Copyright Disclosure

Portions of this software are copyrighted by DataDirect Technologies Corp., 1991 - 2008.

Expat Copyright Disclosure

Part of the software embedded in this product is Expat software.

Copyright © 1998, 1999, 2000 Thai Open Source Software Center Ltd.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

gSOAP Copyright Disclosure

Part of the software embedded in this product is gSOAP software.

Portions created by gSOAP are Copyright © 2001-2004 Robert A. van Engelen, Genivia inc. All Rights Reserved.

THE SOFTWARE IN THIS PRODUCT WAS IN PART PROVIDED BY GENIVIA INC AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,

PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

IBM Copyright Disclosure

ICU License - ICU 1.8.1 and later [as used in DataFlux clients and servers.]

COPYRIGHT AND PERMISSION NOTICE

Copyright © 1995-2005 International Business Machines Corporation and others. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

Informatica Address Doctor Copyright Disclosure

AddressDoctor® Software, © 1994-2015 Platon Data Technology GmbH

Loqate Copyright Disclosure

The Customer hereby acknowledges the following Copyright notices may apply to reference data.

Australia: Copyright. Based on data provided under license from PSMA Australia Limited (www.psma.corn.au)

Austria: © Bundesamt für Eich- und Vermessungswesen

Brazil: Conteudo firnecido por MapLink. Brazil POIs may not be used in publically accessible, internet-based web sites whereby consumers obtain POI data for their personal use.

Canada:

Copyright Notice: This data includes information taken with permission from Canadian authorities, including © Her Majesty, © Queen's Printer for Ontario, © Canada Post, GeoBase ®.

End User Terms: The Data may include or reflect data of licensors including Her Majesty and Canada Post. Such data is licensed on an "as is" basis. The licensors, including Her Majesty and Canada Post, make no guarantees, representation, or warranties respecting such data, either express or implied, arising by law or otherwise, including but not limited to, effectiveness, completeness, accuracy, or fitness for a purpose. The licensors, including Her Majesty and Canada Post, shall not be liable in respect of any claim, demand or action, irrespective of the nature of the cause of the claim, demand or action alleging any loss, injury or damages, direct or indirect, which may result from the use or possession of the data or the Data.

The licensors, including Her Majesty and Canada Post, shall not be liable in any way for loss of revenues or contracts, or any other consequential loss of any kind resulting from any defect in the data or in the Data.

End User shall indemnify and save harmless the licensors, including Her Majesty the Queen, the Minister of Natural Resources of Canada and Canada Post, and their officers, employees and agents from and against any claim, demand or action, irrespective of the nature of the cause of the claim, demand or action, alleging loss, costs, expenses, damages, or injuries (including injuries resulting in death) arising out of the use of possession of the data or the Data.

Croatia, Cyprus, Estonia, Latvia, Lithuania, Moldova, Poland, Slovenia, and/or Ukraine: © EuroGeographics

France: source: Géoroute® IGN France & BD Carto® IGN France

Germany: Die Grundlagendaten wurden mit Genehmigung der zuständigen Behörden entnommen

Great Britain: Based upon Crown Copyright material.

Greece: Copyright Geomatics Ltd. Hungary: Copyright © 2003; Top-Map Ltd.

Italy: La Banca Dati Italiana è stata prodotta usando quale riferimento anche cartografia numerica ed al tratto prodotta e fornita dalla Regione Toscana.

Norway: Copyright © 2000; Norwegian Mapping Authority

Portugal: Source: IgeoE – Portugal

Spain: Información geográfica propiedad del CNIG

Sweden: Based upon electronic data © National Land Survey Sweden.

Switzerland: Topografische Grundlage © Bundesamt für Landestopographie.

Microsoft Copyright Disclosure

Microsoft®, Windows, NT, SQL Server, and Access, are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Oracle Copyright Disclosure

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates.

PCRE Copyright Disclosure

A modified version of the open source software PCRE library package, written by Philip Hazel and copyrighted by the University of Cambridge, England, has been used by DataFlux for regular expression support. More information on this library can be found at: ftp://ftp.csx.cam.ac.uk/pub/software/programming/pcre/.

Copyright © 1997-2005 University of Cambridge. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- Neither the name of the University of Cambridge nor the name of Google Inc. nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Red Hat Copyright Disclosure

Red Hat® Enterprise Linux®, and Red Hat Fedora™ are registered trademarks of Red Hat, Inc. in the United States and other countries.

SAS Copyright Disclosure

Portions of this software and documentation are copyrighted by SAS® Institute Inc., Cary, NC, USA, 2009. All Rights Reserved.

SQLite Copyright Disclosure

The original author of SQLite has dedicated the code to the public domain. Anyone is free to copy, modify, publish, use, compile, sell, or distribute the original SQLite code, either in source code form or as a compiled binary, for any purpose, commercial or non-commercial, and by any means.

Sun Microsystems Copyright Disclosure

Java™ is a trademark of Sun Microsystems, Inc. in the U.S. or other countries.

TomTom Copyright Disclosure

© 2006-2015 TomTom. All rights reserved. This material is proprietary and the subject of copyright protection, database right protection, and other intellectual property rights owned by TomTom or its suppliers. The use of this material is subject to the terms of a license agreement. Any unauthorized copying or disclosure of this material will lead to criminal and civil liabilities.

USPS Copyright Disclosure

National ZIP®, ZIP+4®, Delivery Point Barcode Information, DPV, RDI, and NCOALink®. © United States Postal Service 2005. ZIP Code® and ZIP+4® are registered trademarks of the U.S. Postal Service.

DataFlux is a non-exclusive interface distributor of the United States Postal Service and holds a non-exclusive license from the United States Postal Service to publish and sell USPS CASS, DPV, and RDI information. This information is confidential and proprietary to the United States Postal Service. The price of these products is neither established, controlled, or approved by the United States Postal Service.

VMware Copyright Disclosure

VMware® virtual environment provided those products faithfully replicate the native hardware and provided the native hardware is one supported in the applicable DataFlux product documentation. All DataFlux technical support is provided under the terms of a written license agreement signed by the DataFlux customer.

The VMware virtual environment may affect certain functions in DataFlux products (for example, sizing and recommendations), and it may not be possible to fix all problems.

If DataFlux believes the virtualization layer is the root cause of an incident; the customer will be directed to contact the appropriate VMware support provider to resolve the VMware issue and DataFlux shall have no further obligation for the issue.

Recommended Reading

- *DataFlux Data Management Studio Users Guide*

- *DataFlux Web Studio User's Guide*

- *DataFlux Web Studio Installation and Configuration Guide*

- *DataFlux Authentication Server Administrator's Guide*

- *DataFlux Authentication Server User's Guide*

- *DataFlux Data Management Server Administrator's Guide*

- *DataFlux Data Management Server User's Guide*

- *DataFlux Expression Language Reference Guide for Data Management Studio*

- *SAS Federation Server Administrator's Guide*

- *SAS Federation Server User's Guide*

- *DataFlux Migration Guide*

- DataFlux Quality Knowledge Base Online Help and SAS Quality Knowledge Base Online Help

- *The New IT: How Technology Leaders Are Enabling Business Strategy in the Digital Age*

For a complete list of SAS publications, go to sas.com/store/books. If you have questions about which titles you need, please contact a SAS Representative:

SAS Books
SAS Campus Drive

Cary, NC 27513-2414

Phone: 1-800-727-0025

Fax: 1-919-677-4444

Email: sasbook@sas.com

Web address: sas.com/store/books