
Welcome to Pivotal Greenplum Database 4.3.1

Greenplum Database is a massively parallel processing (MPP) database server that supports next generation data warehousing and large-scale analytics processing. By automatically partitioning data and running parallel queries, it allows a cluster of servers to operate as a single database supercomputer performing tens or hundreds times faster than a traditional database. It supports SQL, MapReduce parallel processing, and data volumes ranging from hundreds of gigabytes, to hundreds of terabytes.

Note: This document contains pertinent release information about Greenplum Database 4.3.1. For previous versions of the release notes for Greenplum Database, go to [Pivotal Documentation](#) or [EMC Support Zone](#).

About Greenplum Database 4.3.1

Greenplum Database 4.3.1 is a maintenance release that introduces a number of significant new features, as well as performance and stability enhancements. Please refer to the following sections for more information about this release.

- [Product Enhancements](#)
- [Changed and Deprecated Features](#)
- [Downloading Greenplum Database](#)
- [Supported Platforms](#)
- [Resolved Issues in Greenplum Database 4.3.1](#)
- [Known Issues in Greenplum Database 4.3.1](#)
- [Upgrading to Greenplum Database 4.3.1](#)
- [Greenplum Database Tools Compatibility](#)
- [Greenplum Database Extensions Compatibility](#)
- [Hadoop Distribution Compatibility](#)
- [Greenplum Database 4.3.1 Documentation](#)

Product Enhancements

Greenplum Database 4.3.1 includes these enhancements:

- [Simplified Configuration to Access HDFS Data with gphdfs](#)
- [Enhancements for Accessing External Data with Greenplum Database](#)
- [Enhancement for Restoring Data from a Greenplum Database Backup](#)
- [Query Memory Accounting Framework](#)
- [Organization of Greenplum Database Administration Documentation](#)

Simplified Configuration to Access HDFS Data with gphdfs

Starting with Greenplum Database 4.3.1, installing a separate gNet package is not required to use the gphdfs protocol. The jar files for the gphdfs extensions, the libraries, and the documentation for the gphdfs extensions are bundled with Greenplum Database. The files are installed in `$GPHOME/lib/hadoop`. The gphdfs protocol is used with external tables to access data from Hadoop file systems.

Starting with Greenplum Database 4.3.1, to upgrade to a different version of gphdfs, you must install the version of Greenplum Database that contains the version of gphdfs that you wish to use.

For information about the gphdfs protocol, see the *Greenplum Database Administrator Guide*.

Enhancements for Accessing External Data with Greenplum Database

Greenplum Database 4.3.1 includes these enhancements:

- The Greenplum Database `gpfdist` utility has been enhanced. The `gpfdist` utility is the Greenplum Database parallel file distribution program. The utility serves data files to or writes data files out from Greenplum Database segments. The new `gpfdist` optional `-w` option that specifies the number of seconds that Greenplum Database delays before closing a target file such as a named pipe.

For Greenplum Database with multiple segments, there might be a delay between segments when writing data from different segments to the file. You can specify a time to wait before Greenplum Database closes the file to ensure all the data is written to the file.

- The Greenplum Database `gpload` utility has been enhanced. The `gpload` utility runs a job that loads data into a Greenplum Database table. You control the load job with a YAML control file. The `INSERT`, `UPDATE` or `MERGE` operation that `gpload` performs, including any SQL commands specified in the SQL collection of the YAML file, are performed as a single transaction. Performing the operation as a single transaction prevents inconsistent data when performing multiple, simultaneous load operations on a target table.

The new `gpload` option `--no_auto_trans` disables performing the operation as a single transaction.

For information about the `gpfdist` and `gpload` utilities, see the *Greenplum Database Utility Guide*.

Enhancement for Restoring Data from a Greenplum Database Backup

Greenplum Database 4.3.1 includes this enhancement:

- The new option `--noanalyze` for the Greenplum Database `gpdbrestore` utility disables `ANALYZE` of tables during restore.
The default action is to run the `ANALYZE` command after a restore. This option is useful if running `ANALYZE` on tables in your database requires a significant amount of time. If you specify this option, you should run `ANALYZE` manually on restored tables. Failure to run `ANALYZE` following a restore might result in poor database performance.

For information about the `gpdbrestore`, `gpfdist` and `gpload` utilities, see the *Greenplum Database Utility Guide*.

Query Memory Accounting Framework

When an out of memory event occurs, the Greenplum Database memory accounting framework reports detailed memory consumption of every query running at the time of the event. The information is written to the Greenplum Database segment logs.

Organization of Greenplum Database Administration Documentation

To make managing Greenplum Database easier, the information to manage and use Greenplum Database systems and database instances have been consolidated into a single book for Greenplum Database 4.3.1 and has been reorganized into five sections.

The information in *Greenplum Database System Administrator Guide* and the *Greenplum Database Database Administrator Guide* have been combined into the *Greenplum Database Administrator Guide*.

See the Preface for a description of the guide sections.

For a list of Greenplum Database documents, see [Greenplum Database 4.3.1 Documentation](#). For a description of the documentation, go to the Pivotal documentation web page <http://docs.gopivotal.com/gpdb/about-docs.html>.

Changed and Deprecated Features

The following are Greenplum Database 4.3.1 changed and deprecated features.

Changed Feature

These are the updates to the Greenplum Database supported platforms:

- Greenplum Database 4.3.1 supports RHEL 6.5 and SUSE Linux Enterprise Server 64-bit 11 SP2.
- Greenplum Database 4.3.1 supports DDBoost SDK 2.6.2.0.

See “[Supported Platforms](#)” for information about platform and SDK support.

Deprecated Features

Pivotal plans to deprecate the following item:

- In Greenplum Database 4.3.1, Pivotal is deprecating support for the TCP and UDP interconnect types or inter-process communication between Greenplum Database segments. The UDPIFC interconnect type will be the only supported interconnect type in a future release of Greenplum Database. The UDPIFC interconnect type provides better performance and stability.

The interconnect is the networking layer of Greenplum Database. The interconnect refers to the inter-process communication between segments and the network infrastructure on which this communication relies. The Greenplum Database interconnect type is controlled by the Greenplum Database server configuration parameter `gp_interconnect_type`.

For information about the Greenplum Database interconnect, see the *Greenplum Database Administrator Guide*. For information about the server configuration parameter `gp_interconnect_type`, see the *Greenplum Database Reference Guide*.

Starting with Greenplum Database 4.3.0.0, Solaris is no longer a supported operating system.

Please send any questions or comments about the deprecated items to gpdb@gopivotal.com.

Downloading Greenplum Database

The location for downloading Greenplum Database software and documentation has changed.

- Greenplum Database 4.3.x software is available from [Pivotal Network](#).
- Current release Greenplum Database documentation is available from the [Pivotal Documentation](#) site.

Previous release versions of Greenplum Database documentation, as well as other Greenplum Database documents, are available from [Support Zone](#)

Supported Platforms

Greenplum Database 4.3.1 runs on the following platforms:

- Red Hat Enterprise Linux 64-bit 5.5, 5.6, 5.7, 6.1, 6.2, 6.4, and 6.5
- SuSE Linux Enterprise Server 64-bit 10 SP4, 11 SP1, 11 SP2
- Solaris x86 64-bit v10 U7, U8, U9, U10
- Oracle Unbreakable Linux 64-bit 5.5
- CentOS 64-bit 5.5, 5.6, 5.7, 6.1, and 6.2

Note: Starting with Greenplum Database 4.3.0.0, Solaris is no longer a supported operating system. Please send any questions or comments about the changes to supported platforms to gpdb@pivotal.io.

Greenplum Database 4.3.1 supports Data Domain Boost on Red Hat Enterprise Linux. This table lists the versions of Data Domain Boost SDK and DDOS supported by Greenplum Database 4.3.x.

Table 1 Data Domain Boost Compatibility

Greenplum Database	Data Domain Boost	DDOS
4.3.1.0	2.6.2.0	5.2, 5.3, and 5.4
4.3.0.0	2.4.2.2	5.0.1.0, 5.1, and 5.2

Greenplum Database support on DCA:

- Greenplum Database 4.3.x, all versions, is supported on DCA V2, and requires DCA software version 2.1.0.0 or greater due to known DCA software issues in older DCA software versions.
- Greenplum Database 4.3.x, all versions, is supported on DCA V1, and requires DCA software version 1.2.2.2 or greater due to known DCA software issues in older DCA software versions.

Resolved Issues in Greenplum Database 4.3.1

The table below lists issues that are now resolved in Greenplum Database 4.3.1.

For issues resolved in prior releases, refer to the corresponding release notes available from EMC [Support Zone](#).

Table 2 Resolved Issues in 4.3.1

Issue Number	Category	Resolved in	Description
23757	Security	4.3.1	Greenplum Database software has been updated to use OpenSSL 0.9.8za in response to the OpenSSL Security Advisory [05 Jun 2014]. For information about the advisory, see http://www.openssl.org/news/secadv_20140605.txt .
22301	Replication: Master Mirroring	4.3.1	DCA customers who wished to use Greenplum Database 4.3 could not use the utility <code>dca_setup</code> . This issue has been resolved in Greenplum Database 4.3.1.
22281	Backup and Restore	4.3.1	For partitioned append-optimized tables, a partition was backed up even though it was not modified.
21591	Management Scripts Suite	4.3.1	The Greenplum Database utilities <code>gpstart</code> and <code>gprecoverseg</code> hung when checking the process ID in the <code>postmaster.pid</code> file and the ID matched a non-postgres running process.
23421	Locking, Signals, Processes	4.3.1	In some cases, concurrent CREATE TABLE and DROP TABLE operations caused Greenplum Database to hang due to incorrect lock handling.
13825	Functions and Languages, Transaction Management	4.3.1	In PL/PGSQL functions, exception blocks were not handled properly. Depending on where the exception is encountered during function execution, the improper block handling resulted in either the creation of catalog inconsistency between master and segment, or Greenplum Database issuing the following message: The distributed transaction 'Prepare' broadcast failed to one or more segments.
22655	Locking, Signals, Processes	4.3.1	Greenplum Database hung due to incorrect lock handling that caused a race condition. The lock handling issue was caused by a compiler optimization.
20924	Dispatch	4.3.1	For some queries that contained a window function and that executed on both the master and segments, the query would hang when executed from an ODBC/JDBC client.
21899	Backup and Restore	4.3.1	When performing an incremental backup, the <code>gpcrondump</code> utility backed up temporary tables that existed during the time of the backup. This caused a failure when performing a restore with the <code>gpdrestore</code> utility that used the incremental backup.
22293	Backup and Restore	4.3.1	Greenplum Database supports Data Domain DDOS 5.4. See Supported Platforms for information about supported versions of Data Domain Boost.
22442	Loaders: gpfdist	4.3.1	The Greenplum Database Load Tools for Windows installation did not include the <code>gssapi</code> and <code>auth</code> libraries. This issue has been resolved.
19476	Client Access Methods and Tools	4.3.1	Running multiple <code>gpload</code> sessions simultaneously that loaded data into the same table resulted in inconsistent data in the table. See the <code>gpload</code> information in Product Enhancements .

Table 2 Resolved Issues in 4.3.1

Issue Number	Category	Resolved in	Description
22863	DDL and Utility Statements	4.3.1	When > (greater than) was used in the CREATE OPERATOR CLASS command as an operator name, this error was returned. operator > is not a valid ordering operator when using operator classes
22219	Query Planner	4.3.1	In certain queries that contain the median function and a GROUP BY clause, the query planner produced an incorrect plan in which some necessary columns were not projected in the operator nodes. This caused an error when trying to look up the missing columns.
22084	OS Abstraction	4.3.1	Improved handling of situations where Greenplum Database encounters segment violation errors.
17995	DDL and Utility Statements	4.3.1	In some cases, the functions pg_cancel_backend() and pg_terminate_backend() did not terminate sessions.
17773	DDL and Utility Statements	4.3.1	Greenplum Database did not properly check privileges during certain RESET ALL operations.
17481	Catalog and Metadata, DDL and Utility Statements	4.3.1	Queries on the system view <i>pg_partitions</i> could fail to return when DDL statements on partitioned tables were running concurrently.
15834	Loaders: Copy/External Tabs	4.3.1	A COPY command cancel request (Ctrl+c) followed by another COPY command and a cancel request caused the Greenplum Database session to hang. When cancel request was attempted again, a SIGSEGV error occurred.
14367	DDL and Utility Statements	4.3.1	ALTER TABLE ADD COLUMN with default NULL was not supported for append-optimized tables. This syntax is now supported.
21522	Backup and Restore	4.3	The Greenplum Database utility pg_dump printed information-level messages (messages with the label [INFO]) to stderr that were not printed in previous releases. These messages were printed even when pg_dump completes without errors.

Known Issues in Greenplum Database 4.3.1

This section lists the known issues in Greenplum Database 4.3.1. A workaround is provided where applicable.

For known issues discovered in previous releases, including patch releases to Greenplum Database 4.2.x, 4.1 or 4.0.x, see the corresponding release notes, available from EMC [Support Zone](#):

Table 3 All Known Issues in 4.3.1

Issue	Category	Description
19660	Authentication	An issue in Greenplum Database prevents LDAPS (LDAP over SSL) from functioning on the standard secure port 636.
23824	Authentication	In some cases, LDAP client utility tools cannot be used after running the source command <pre>source \$GPHOME/greenplum_path.sh</pre> because the LDAP libraries included with Greenplum Database are not compatible with the LDAP client utility tools that are installed with operating system. Workaround: The LDAP tools can be used without running the source command in the environment.
22328	Management Scripts	The process of updating a Greenplum Database package includes removing all previous versions of the system objects related to the package. For example, previous versions of shared libraries are removed. After the update process, a database function will fail when it is called if the function references a package file that has been removed.
23227	Client Access Methods and Tools	When using Kerberos and the GSSAPI authentication method, the Greenplum Database role property Valid Until is ignored. This property is used to control access to a Greenplum database.
23568	Backup and Restore	When backing up a Greenplum database with the Greenplum Database gpcrondump utility and specifying an NFS directory with the -u option, the gpcrondump utility creates an empty db_dumps directory in the master and segment data directories.
23637	Backup and Restore	When restoring a Greenplum database with the Greenplum Database gpcrondump utility, the utility performs an ANALYZE operation on the entire database. Workaround: When restoring Greenplum database with the Greenplum Database gpcrondump utility, specify the --noanalyze option, and then run the ANALYZE command on the tables that require updated statistics.
23485	Transaction Management	When a single session to Greenplum Database runs transactions, temporary files were not removed after the transaction completed. If a the session ran a large number of transactions, the temporary files required a large amount of disk space.
23417	Transaction Management	Some SQL queries against an append-optimized table that has compression enabled and that contains a column with an unknown data type cause a Greenplum Database SIGSEGV error.
22205	Replication: Segment Mirroring	In some cases, running the Greenplum Database command <code>gprecoverseg -r</code> to rebalance segment instances fails and causes database catalog corruption.

Table 3 All Known Issues in 4.3.1

Issue	Category	Description
23525	Query Planner	Some SQL queries that contain sub-selects fail with this error. ERROR: Failed to locate datatype for paramid 0
22792	Build	Greenplum Database is not certified on Red Hat Enterprise Linux 5.10.
22215	Build	Greenplum Database is not certified with these connectivity drivers: <ul style="list-style-type: none"> • Data Direct v 7.022; PowerExchange for Greenplum 9.5.1 • 32-bit Microstrategy ODBC for Greenplum Wire Protocol 6.10.01.80 • Open source ODBC 9.01.0100 and JDBC 9.1.902 Type 4 • SAS/ACCESS 9.3 driver provided with SAS software2
23366	Resource Management	In Greenplum Database 4.2.7.0 and later, the priority of some running queries, cannot be dynamically adjusted with the <code>gp_adjust_priority()</code> function. The attempt to execute this request might silently fail. The return value of the <code>gp_adjust_priority()</code> call indicates success or failure. If 1 is returned, the request was not successfully executed. If a number greater than 1 is returned, the request was successful. If the request fails, the priority of all running queries are unchanged, they remain as they were before the <code>gp_adjust_priority()</code> call.
23492	Backup and Restore,	A backup from a Greenplum Database 4.3.x system that is created with a Greenplum Database back up utility, for example <code>gpcrondump</code>, cannot be restored to a Greenplum Database 4.2.x system with the <code>psql</code> utility or the corresponding restore utility, for example <code>gpdbrestore</code>.
23521	Client Access Methods and Tools	Hadoop YARN based on Hadoop 2.2 or later does not work with Greenplum Database. Workaround: For Hadoop distributions based on Hadoop 2.2 or later that are supported by Greenplum Database, the classpath environment variable and other directory paths defined in <code>\$GPHOME/lib/hadoop/hadoop_env.sh</code> must be to be modified so that the paths point to the appropriate JAR files.
21917	Replication: Segment Mirroring	In some rare cases after the Greenplum Database utility <code>gprecoverseg</code> was run, some append-optimized tables and a persistent table were detected having less data on a mirror segment corresponding to a primary segment.
20453	Query Planner	For SQL queries of either of the following forms: <code>SELECT columns FROM table WHERE table.column NOT IN subquery;</code> <code>SELECT columns FROM table WHERE table.column = ALL subquery;</code> tuples that satisfy both of the following conditions are not included in the result set: <ul style="list-style-type: none"> • <code>table.column</code> is NULL. • <code>subquery</code> returns the empty result.
21724	Query Planner	Greenplum Database executes an SQL query in two stages if a scalar subquery is involved. The output of the first stage plan is fed into the second stage plan as a external parameter. If the first stage plan generates zero tuples and directly contributes to the output of the second stage plan, incorrect results might be returned.
21838	Backup and Restore	When restoring sets of tables with the Greenplum Database utility <code>gpdbrestore</code>, the table schemas must be defined in the database. If a table's schema is not defined in the database, the table is not restored. When performing a full restore, the database schemas are created when the tables are restored. Workaround: Before restoring a set of tables, create the schemas for the tables in the database.

Table 3 All Known Issues in 4.3.1

Issue	Category	Description
21129	DDL and Utility Statements	SSL is only supported on the master host. It is not supported on segment hosts.
20822	Backup and Restore	Special characters such as !, \$, #, and @ cannot be used in the password for the Data Domain Boost user when specifying the Data Domain Boost credentials with the <code>gpcrondump</code> options <code>--ddboost-host</code> and <code>--ddboost-user</code>.
18247	DDL and Utility Statements	TRUNCATE command does not remove rows from a sub-table of a partitioned table. If you specify a sub-table of a partitioned table with the TRUNCATE command, the command does not remove rows from the sub-table and its child tables. Workaround: Use the ALTER TABLE command with the TRUNCATE PARTITION clause to remove rows from the sub-table and its child tables.
19788	Replication: Resync, Transaction Management	In some rare circumstances, performing a full recovery with <code>gprecoverseg</code> fails due to inconsistent LSN. Workaround: Stop and restart the database. Then perform a full recovery with <code>gprecoverseg</code> .
19705	Loaders: gpload	gpload fails on Windows XP with Python 2.6. Workaround: Install Python 2.5 on the system where gpload is installed.
19493 19464 19426	Backup and Restore	The <code>gpcrondump</code> and <code>gpdbrestore</code> utilities do not handle errors returned by DD Boost or Data Domain correctly. These are two examples: <ul style="list-style-type: none"> • If invalid Data Domain credentials are specified when setting the Data Domain Boost credentials with the <code>gpcrondump</code> utility, the error message does not indicate that invalid credentials were specified. • Restoring a Greenplum database from a Data Domain system with <code>gpdbrestore</code> and the <code>--ddboost</code> option indicates success even though segment failures occurred during the restore. Workaround: The errors are logged in the master and segment server backup or restore status and report files. Scan the status and report files to check for error messages.
19278	Backup and Restore	When performing a selective restore of a partitioned table from a full backup with <code>gpdbrestore</code>, the data from leaf partitions are not restored. Workaround: When doing a selective restore from a full backup, specify the individual leaf partitions of the partitioned tables that are being restored. Alternatively, perform a full backup, not a selective backup.
15692 17192	Backup and Restore	Greenplum Database's implementation of RSA lock box for Data Domain Boost changes backup and restore requirements for customers running SUSE. The current implementation of the RSA lock box for Data Domain Boost login credential encryption only supports customers running on Red Hat Enterprise Linux. Workaround: If you run Greenplum Database on SUSE, use NFS as your backup solution. See the <i>Greenplum Database Administrator Guide</i> for information on setting up a NFS backup.
18850	Backup and Restore	Data Domain Boost credentials cannot be set up in some environments due to the absence of certain libraries (for example, <code>libstdc++</code>) expected to reside on the platform. Workaround: Install the missing libraries manually on the system.

Table 3 All Known Issues in 4.3.1

Issue	Category	Description
18851	Backup and Restore	When performing a data-only restore of a particular table, it is possible to introduce data into Greenplum Database that contradicts the distribution policy of that table. In such cases, subsequent queries may return unexpected and incorrect results. To avoid this scenario, we suggest you carefully consider the table schema when performing a restore.
18774	Loaders	External web tables that use IPv6 addresses must include a port number.
18713	Catalog and Metadata	Drop language plpgsql cascade results in a loss of <code>gp_toolkit</code> functionality. Workaround: Reinstall <code>gp_toolkit</code> .
18710	Management Scripts Suite	Greenplum Management utilities cannot parse IPv6 IP addresses. Workaround: Always specify IPv6 hostnames rather than IP addresses
18703	Loaders	The <code>bytenum</code> field (byte offset in the load file where the error occurred) in the error log when using <code>gpfdist</code> with data in text format errors is not populated, making it difficult to find the location of an error in the source file.
12468	Management Scripts Suite	<code>gpexpand --rollback</code> fails if an error occurs during expansion such that it leaves the database down <code>gpstart</code> also fails as it detects that expansion is in progress and suggests to run <code>gpexpand --rollback</code> which will not work because the database is down. Workaround: Run <code>gpstart -m</code> to start the master and then run <code>rollback</code> ,
18785	Loaders	Running <code>gpload</code> with the <code>--ssl</code> option and the relative path of the source file results in an error that states the source file is missing. Workaround: Provide the full path in the yml file or add the loaded data file to the certificate folder.
18414	Loaders	Unable to define external tables with fixed width format and empty line delimiter when file size is larger than <code>gpfdist</code> chunk (by default, 32K).
14640	Backup and Restore	<code>gpdbrstore</code> outputting incorrect non-zero error message. When performing single table restore, <code>gpdbrstore</code> gives warning messages about non-zero tables however prints out zero rows.
17285	Backup and Restore	NFS backup with <code>gpcrondump -c</code> can fail. In circumstances where you haven't backed up to a local disk before, backups to NFS using <code>gpcrondump</code> with the <code>-c</code> option can fail. On fresh systems where a backup has not been previously invoked there are no dump files to cleanup and the <code>-c</code> flag will have no effect. Workaround: Do not run <code>gpcrondump</code> with the <code>-c</code> option the first time a backup is invoked from a system.
17837	Upgrade/ Downgrade	Major version upgrades internally depend on the <code>gp_toolkit</code> system schema. The alteration or absence of this schema may cause upgrades to error out during preliminary checks. Workaround: To enable the upgrade process to proceed, you need to reinstall the <code>gp_toolkit</code> schema in all affected databases by applying the SQL file found here: <code>\$GPHOME/share/postgresql/gp_toolkit.sql</code> .

Table 3 All Known Issues in 4.3.1

Issue	Category	Description
17513	Management Scripts Suite	Running more than one <code>gpfilespace</code> command concurrently with itself to move either temporary files (<code>--movetempfilespace</code>) or transaction files (<code>--movetransfilespace</code>) to a new filespace can in some circumstances cause OID inconsistencies. Workaround: Do not run more than one <code>gpfilespace</code> command concurrently with itself. If an OID inconsistency is introduced <code>gpfilespace --movetempfilespace</code> or <code>gpfilespace --movetransfilespace</code> can be used to revert to the default filespace.
17780	DDL/DML: Partitioning	ALTER TABLE ADD PARTITION inheritance issue When performing an <code>ALTER TABLE ADD PARTITION</code> operation, the resulting parts may not correctly inherit the storage properties of the parent table in cases such as adding a default partition or more complex subpartitioning. This issue can be avoided by explicitly dictating the storage properties during the <code>ADD PARTITION</code> invocation. For leaf partitions that are already afflicted, the issue can be rectified through use of <code>EXCHANGE PARTITION</code> .
17795	Management Scripts Suite	Under some circumstances, <code>gppkg</code> on SUSE is unable to correctly interpret error messages returned by <code>rpm</code> . On SUSE, <code>gppkg</code> is unable to operate correctly under circumstances that require a non-trivial interpretation of underlying <code>rpm</code> commands. This includes scenarios that result from overlapping packages, partial installs, and partial uninstalls.
17604	Security	A Red Hat Enterprise Linux (RHEL) 6.x security configuration file limits the number of processes that can run on <code>gpadm</code> . RHEL 6.x contains a security file (<code>/etc/security/limits.d/90-nproc.conf</code>) that limits available processes running on <code>gpadm</code> to 1064. Workaround: Remove this file or increase the processes to 131072.
17415	Installer	When you run <code>gppkg -q -info<some gppkg></code> , the system shows the GPDBversion as main build dev.
17334	Management Scripts Suite	You may see warning messages that interfere with the operation of management scripts when logging in. Greenplum recommends that you edit the <code>/etc/motd</code> file and add the warning message to it. This will send the messages to be redirected to <code>stdout</code> and not <code>stderr</code> . You must encode these warning messages in UTF-8 format.
17221	Resource Management	Resource queue deadlocks may be encountered if a cursor is associated with a query invoking a function within another function.
17113	Management Scripts Suite	Filespaces are inconsistent when the Greenplum database is down. Filespaces become inconsistent in case of a network failure. Greenplum recommends that processes such as moving a filespace be done in an environment with an uninterrupted power supply.
17189	Loaders: <code>gpfdist</code>	<code>gpfdist</code> shows the error "Address already in use" after successfully binding to socket IPv6. Greenplum supports IPv4 and IPv6. However, <code>gpfdist</code> fails to bind to socket IPv4, and shows the message "Address already in use", but binds successfully to socket IPv6.

Table 3 All Known Issues in 4.3.1

Issue	Category	Description
16519	Backup and Restore	<p>Limited data restore functionality and/or restore performance issues can occur when restoring tables from a full database backup where the default backup directory was not used.</p> <p>In order to restore from backup files not located in the default directory you can use the <code>-R</code> to point to another host and directory. This is not possible however, if you want to point to a different directory on the same host (NFS for example).</p> <p>Workaround: Define a symbolic link from the default dump directory to the directory used for backup, as shown in the following example:</p> <ol style="list-style-type: none"> 1. Perform a full Database Backup to a specific NFS directory: <pre>\$ gpccrondump -x <db_name> -z -u /backup/DCA-93 -a</pre> 2. Create a file listing the segment servers: <pre>\$ vi /home/gpadmin/segments sdw1 sdw2 sdw3 ...</pre> 3. Remove the relevant date folder from <code>db_dumps</code> directories on the master and segments: <pre>\$ rm -r /data/master/gpseg-1/db_dumps/20120619 \$ gpssh -f segments 'rm -r /data1/primary/gpseg*/db_dumps/20120619' \$ gpssh -f segments 'rm -r /data2/primary/gpseg*/db_dumps/20120619'</pre> 4. Create a symbolic link between the master and segment directories and the directory to which you backed up in step 1. Only the master and <code>sdw1</code> was shown here, write a script for the remaining segments: <pre>\$ ln -s /backup/DCA-93/db_dumps/20120619 /data/master/gpseg-1/db_dumps/20120619 \$ gpssh -h sdw1 'ln -s /backup/DCA-93/db_dumps/20120619 /data1/primary/gpseg0/db_dumps/20120619' \$ gpssh -h sdw1 'ln -s /backup/DCA-93/db_dumps/20120619 /data1/primary/gpseg1/db_dumps/20120619' \$ gpssh -h sdw1 'ln -s /backup/DCA-93/db_dumps/20120619 /data1/primary/gpseg2/db_dumps/20120619' \$ gpssh -h sdw1 'ln -s /backup/DCA-93/db_dumps/20120619 /data2/primary/gpseg3/db_dumps/20120619' \$ gpssh -h sdw1 'ln -s /backup/DCA-93/db_dumps/20120619 /data2/primary/gpseg4/db_dumps/20120619' \$ gpssh -h sdw1 'ln -s /backup/DCA-93/db_dumps/20120619 /data2/primary/gpseg5/db_dumps/20120619'</pre> 5. Restore from backup files: <pre>\$ gpdbrestore -t 20120619061835 -T <schema.table> -a</pre> 6. Remove the symbolic links <pre>\$ rrm -r /data/master/gpseg-1/db_dumps/20120619 \$ gpssh -f segments 'rm -r /data1/primary/gpseg*/db_dumps/20120619' \$ gpssh -f segments 'rm -r /data2/primary/gpseg*/db_dumps/20120619'</pre>

Table 3 All Known Issues in 4.3.1

Issue	Category	Description
16064	Backup and Restore	Restoring a compressed dump with the <code>--ddbboost</code> option displays incorrect dump parameter information. When using <code>gpdbrestore --ddbboost</code> to restore a compressed dump, the restore parameters incorrectly show "Restore compressed dump = Off". This error occurs even if <code>gpdbrestore</code> passes the <code>--gp-c</code> option to use <code>gunzip</code> for in-line de-compression.
15899	Backup and Restore	When running <code>gpdbrestore</code> with the <code>list (-L)</code> option, external tables do not appear; this has no functional impact on the restore job.

Upgrading to Greenplum Database 4.3.1

The upgrade path supported for this release is Greenplum Database 4.2.x.x to Greenplum Database 4.3.1. The minimum recommended upgrade path for this release is from Greenplum Database version 4.2.x.x. If you have an earlier major version of the database, you must first upgrade to version 4.2.x.x.

For detailed upgrade procedures and information, see the following sections:

- [Upgrading from 4.3 to 4.3.1](#)
- [Upgrading from 4.2.x.x to 4.3.1](#)
- [For Users Running Greenplum Database 4.1.x.x](#)
- [For Users Running Greenplum Database 4.0.x.x](#)
- [For Users Running Greenplum Database 3.3.x.x](#)
- [Troubleshooting a Failed Upgrade](#)

If you are utilizing Data Domain Boost, you have to re-enter your DD Boost credentials after upgrading from Greenplum Database 4.2.x.x to 4.3 as follows:

```
gpcrondump --ddboost-host ddboost_hostname --ddboost-user
ddboost_user
```

Note that if you do not reenter your login credentials after an upgrade, your backup will never start because the Greenplum Database cannot connect to the Data Domain system. You will receive an error advising you to check your login credentials.

Upgrading from 4.3 to 4.3.1

An upgrade from 4.3 to 4.3.1 involves stopping Greenplum Database, updating the Greenplum Database software binaries, and restarting Greenplum Database.

1. Log in to your Greenplum Database master host as the Greenplum administrative user:

```
$ su - gadmin
```

2. Perform a smart shutdown of your current Greenplum Database 4.3 system (there can be no active connections to the database):

```
$ gpstop
```

3. Run the installer for 4.3.1 on the Greenplum Database master host. When prompted, choose an installation location in the same base directory as your current installation. For example:

```
/usr/local/greenplum-db-4.3.1.0
```

4. Edit the environment of the Greenplum Database superuser (gadmin) and make sure you are sourcing the `greenplum_path.sh` file for the new installation. For example change the following line in `.bashrc` or your chosen profile file:

```
source /usr/local/greenplum-db-4.3.0.0/greenplum_path.sh
to:
```

```
source /usr/local/greenplum-db-4.3.1.0/greenplum_path.sh
```

Or if you are sourcing a symbolic link (`/usr/local/greenplum-db`) in your profile files, update the link to point to the newly installed version. For example:

```
$ rm /usr/local/greenplum-db
$ ln -s /usr/local/greenplum-db-4.3.1.0
  /usr/local/greenplum-db
```

5. Source the environment file you just edited. For example:

```
$ source ~/.bashrc
```

6. Run the `gpsegininstall` utility to install the 4.3.1 binaries on all the segment hosts specified in the `hostfile`. For example:

```
$ gpsegininstall -f hostfile
```

7. After all segment hosts have been upgraded, you can log in as the `gpadmin` user and restart your Greenplum Database system:

```
$ su - gpadmin
$ gpstart
```

8. If you are utilizing Data Domain Boost, you have to re-enter your DD Boost credentials after upgrading from Greenplum Database 4.3 to 4.3.1 as follows:

```
gpcrondump --ddboost-host ddboost_hostname --ddboost-user
ddboost_user
```

Note that if you do not reenter your login credentials after an upgrade, your backup will never start because the Greenplum Database cannot connect to the Data Domain system. You will receive an error advising you to check your login credentials.

Upgrading from 4.2.x.x to 4.3.1

This section describes how you can upgrade from Greenplum Database 4.2.x.x or later to Greenplum Database 4.3.1. For users running versions prior to 4.2.x.x of Greenplum Database, see the following:

- [For Users Running Greenplum Database 4.1.x.x](#)
- [For Users Running Greenplum Database 4.0.x.x](#)
- [For Users Running Greenplum Database 3.3.x.x](#)

Planning Your Upgrade

Before you begin your upgrade, make sure the master and all segments (data directories and filesystem) have at least 2GB of free space.

Prior to upgrading your database, Pivotal recommends that you run a pre-upgrade check to verify your database is healthy.

You can perform a pre-upgrade check by executing the `gpmigrator` (`_mirror`) utility with the `--check-only` option.

For example:

```
source $new_gphome/greenplum_path.sh;
gpmigrator_mirror --check-only $old_gphome $new_gphome
```


Note: Performing a pre-upgrade check of your database with the `gpmigrator` (`_mirror`) utility should be done during a database maintenance period. When the utility checks the database catalog, users cannot access the database.

Migrating a Greenplum Database That Contains AO Tables

The migration process updates AO tables that are in a Greenplum Database to UAO tables. For a database that contains a large number of AO tables, the conversion to UAO tables might take a considerable amount of time.

Upgrade Procedure

This section divides the upgrade into the following phases: pre-upgrade preparation, software installation, upgrade execution, and post-upgrade tasks.

We have also provided you with an [Upgrade Checklist](#) that summarizes this procedure.

IMPORTANT: Carefully evaluate each section and perform all required and conditional steps. Failing to perform any of these steps can result in an aborted upgrade, placing your system in an unusable or even unrecoverable state.

Pre-Upgrade Preparation (on your 4.2.x system)

Perform these steps on your current 4.2.x Greenplum Database system. This procedure is performed from your Greenplum master host and should be executed by the Greenplum superuser (`gadmin`).

1. Log in to the Greenplum Database master as the `gadmin` user:

```
$ su - gadmin
```
2. (optional) Vacuum all databases prior to upgrade. For example:

```
$ vacuumdb database_name
```
3. (optional) Clean out old server log files from your master and segment data directories. For example, to remove log files from 2011 from your segment hosts:

```
$ gpssh -f seg_host_file -e 'rm /gpdata/*/gp*/pg_log/gpdb-2011-*.csv'
```

Note: Running Vacuum and cleaning out old logs files is not required, but it will reduce the size of Greenplum Database files to be backed up and migrated.

4. Run `gpstate` to check for failed segments.

```
$ gpstate
```
5. If you have failed segments, you must recover them using `gprecoverseg` before you can upgrade.

```
$ gprecoverseg
```

Note: It might be necessary to restart the database if the preferred role does not match the current role; for example, if a primary segment is acting as a mirror segment or a mirror segment is acting as a primary segment.
6. Copy or preserve any additional folders or files (such as backup folders) that you have added in the Greenplum data directories or `$GPHOME` directory. Only files or folders strictly related to Greenplum Database operations are preserved by the migration utility.

Install the Greenplum Database 4.3 Software Binaries

1. Download or copy the installer file to the Greenplum Database master host.
2. Unzip the installer file. For example:

```
# unzip greenplum-db-4.3.1.0-PLATFORM.zip
```
3. Launch the installer using bash. For example:

```
# /bin/bash greenplum-db-4.3.1.0-PLATFORM.bin
```
4. The installer will prompt you to accept the Greenplum Database license agreement. Type `yes` to accept the license agreement.
5. The installer will prompt you to provide an installation path. Press `ENTER` to accept the default install path (for example: `/usr/local/greenplum-db-4.3.1.0`), or enter an absolute path to an install location. You must have write permissions to the location you specify.
6. The installer installs the Greenplum Database software and creates a `greenplum-db` symbolic link one directory level above your version-specific Greenplum installation directory. The symbolic link is used to facilitate patch maintenance and upgrades between versions. The installed location is referred to as `$GPHOME`.
7. Source the path file from your new 4.3.1 installation. For example:

```
$ source /usr/local/greenplum-db-4.3.1.0/greenplum_path.sh
```
8. Run the `gpsegininstall` utility to install the 4.3.1 binaries on all the segment hosts specified in the `hostfile`. For example:

```
$ gpsegininstall -f hostfile
```

Upgrade Execution

During upgrade, all client connections to the master will be locked out. Inform all database users of the upgrade and lockout time frame. From this point onward, users should not be allowed on the system until the upgrade is complete.

9. Source the path file from your old 4.2.x.x installation. For example:

```
$ source /usr/local/greenplum-db-4.2.6.3/greenplum_path.sh
```
10. (*optional but strongly recommended*) Back up all databases in your Greenplum Database system using `gpcrondump` (or `zfs` snapshots on Solaris systems). See the *Greenplum Database Administrator Guide* for more information on how to do backups using `gpcrondump`. Make sure to secure your backup files in a location outside of your Greenplum data directories.
11. If your system has a standby master host configured, remove the standby master from your system configuration. For example:

```
$ gpinitstandby -r
```
12. Perform a clean shutdown of your current Greenplum Database 4.2.x.x system. For example:

```
$ gpstop
```

- 13.** Source the path file from your new 4.3.1.0 installation. For example:
- ```
$ source /usr/home/greenplum-db-4.3.1.0/greenplum_path.sh
```
- 14.** Update the Greenplum Database environment so it is referencing your new 4.3.1 installation.

- a.** For example, update the `greenplum-db` symbolic link on the master and standby master to point to the new 4.3.1 installation directory. For example (as root):

```
rm -rf /usr/local/greenplum-db
ln -s /usr/local/greenplum-db-4.3.1.0
 /usr/local/greenplum-db
chown -R gpadmin /usr/local/greenplum-db
```

- b.** Using `gpssh`, also update the `greenplum-db` symbolic link on all of your segment hosts. For example (as root):

```
gpssh -f segment_hosts_file
=> rm -rf /usr/local/greenplum-db
=> ln -s /usr/local/greenplum-db-4.3.1.0
 /usr/local/greenplum-db
=> chown -R gpadmin /usr/local/greenplum-db
=> exit
```

- 15.** (*optional but recommended*) Prior to running the migration, perform a pre-upgrade check to verify that your database is healthy by executing the 4.3.1 version of the `gpmigrator` utility with the `--check-only` option. For example:

```
gpmigrator_mirror --check-only
 /usr/local/greenplum-db-4.2.6.3
 /usr/local/greenplum-db-4.3.1
```

- 16.** As `gpadmin`, run the 4.3.1 version of the migration utility specifying your old and new `GPHOME` locations. If your system has mirrors, use `gpmigrator_mirror`. If your system does not have mirrors, use `gpmigrator`. For example on a system with mirrors:

```
$ su - gpadmin
$ gpmigrator_mirror /usr/local/greenplum-db-4.2.6.3
 /usr/local/greenplum-db-4.3.1.0
```

Note: If the migration does not complete successfully, contact Customer Support (see [“Troubleshooting a Failed Upgrade”](#) on page 23).

- 17.** The migration can take a while to complete. After the migration utility has completed successfully, the Greenplum Database 4.3.1 system will be running and accepting connections.

**Note:** After the migration utility has completed, the resynchronization of the mirror segments with the primary segments continues. Even though the system is running, the mirrors are not active until the resynchronization is complete.

#### Post-Upgrade (on your 4.3.1 system)

- 18.** If your system had a standby master host configured, reinitialize your standby master using `gpinitstandby`:

```
$ gpinitstandby -s standby_hostname
```

- 19.** If your system uses external tables with `gpfdist`, stop all `gpfdist` processes on your ETL servers and reinstall `gpfdist` using the compatible Greenplum Database 4.3.1 Load Tools package. Application Packages are available at [Pivotal Network](#).
- 20.** Rebuild any custom modules against your 4.3.1 installation (for example, any shared library files for user-defined functions in `$GPHOME/lib`).
- 21.** Use the Greenplum Database `gppkg` utility to install Greenplum Database extensions. If you were previously using any Greenplum Database extensions such as `pgcrypto`, `PL/R`, `PL/Java`, `PL/Perl`, and `PostGIS`, download the corresponding packages from [Pivotal Network](#), and install using this utility. See the *Greenplum Database Utility Guide 4.3* for usage details.
- 22.** If you want to utilize the Greenplum Command Center management tool, install the latest Command Center Console and update your environment variable to point to the latest Command Center binaries (source the `gpperfmon_path.sh` file from your new installation).  
  
Note that the Greenplum Command Center management tool replaces Greenplum Performance Monitor.  
  
Command Center Console packages are available from [Pivotal Network](#).
- 23.** Inform all database users of the completed upgrade. Tell users to update their environment to source the Greenplum Database 4.3.1 installation (if necessary).

## Upgrade Checklist

This checklist provides a quick overview of all the steps required for an upgrade from 4.2.x.x to 4.3.1. Detailed upgrade instructions are provided in the [Upgrade Procedure](#) section.

| <b>Pre-Upgrade Preparation (on your current system)</b>                                 |                                                                                                                       |
|-----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------|
| <i>* 4.2.x.x system is up and available</i>                                             |                                                                                                                       |
| <input type="checkbox"/>                                                                | Log in to your master host as the <code>gpadmin</code> user (your Greenplum superuser).                               |
| <input type="checkbox"/>                                                                | (Optional) Run <code>VACUUM</code> on all databases,                                                                  |
| <input type="checkbox"/>                                                                | (Optional) Remove old server log files from <code>pg_log</code> in your master and segment data directories.          |
| <input type="checkbox"/>                                                                | Check for and recover any failed segments ( <code>gpstate</code> , <code>gprecoverseg</code> ).                       |
| <input type="checkbox"/>                                                                | Copy or preserve any additional folders or files (such as backup folders).                                            |
| <input type="checkbox"/>                                                                | Install the Greenplum Database 4.3 binaries on all Greenplum hosts.                                                   |
| <input type="checkbox"/>                                                                | Inform all database users of the upgrade and lockout time frame.                                                      |
| <b>Upgrade Execution</b>                                                                |                                                                                                                       |
| <i>* The system will be locked down to all user activity during the upgrade process</i> |                                                                                                                       |
| <input type="checkbox"/>                                                                | Backup your current databases.                                                                                        |
| <input type="checkbox"/>                                                                | Remove the standby master ( <code>gpinitstandby -r</code> ).                                                          |
| <input type="checkbox"/>                                                                | Do a clean shutdown of your current system ( <code>gpstop</code> ).                                                   |
| <input type="checkbox"/>                                                                | Update your environment to source the new Greenplum Database 4.3.1 installation.                                      |
| <input type="checkbox"/>                                                                | Run the upgrade utility ( <code>gpmigrator_mirror</code> if you have mirrors, <code>gpmigrator</code> if you do not). |
| <input type="checkbox"/>                                                                | After the upgrade process finishes successfully, your 4.3.1 system will be up and running.                            |
| <b>Post-Upgrade (on your 4.3 system)</b>                                                |                                                                                                                       |
| <i>* The 4.2.x.x system is up</i>                                                       |                                                                                                                       |
| <input type="checkbox"/>                                                                | Reinitialize your standby master host ( <code>gpinitstandby</code> ).                                                 |

|                          |                                                                                                                                  |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| <input type="checkbox"/> | Upgrade <code>gpfdist</code> on all of your ETL hosts.                                                                           |
| <input type="checkbox"/> | Rebuild any custom modules against your 4.3.1 installation.                                                                      |
| <input type="checkbox"/> | Download and install any Greenplum Database extensions.                                                                          |
| <input type="checkbox"/> | (Optional) Install the latest Command Center Console and update your environment to point to the latest Command Center binaries. |
| <input type="checkbox"/> | Inform all database users of the completed upgrade.                                                                              |

### For Users Running Greenplum Database 4.1.x.x

Users on a release prior to 4.1.x.x cannot upgrade directly to 4.3.1.

1. Upgrade from your current release to 4.2.x.x (follow the upgrade instructions in the latest Greenplum Database 4.2.x.x release notes available at [Pivotal Documentation](#)).
2. Follow the upgrade instructions in these release notes for [Upgrading from 4.2.x.x to 4.3.1](#).

### For Users Running Greenplum Database 4.0.x.x

Users on a release prior to 4.1.x.x cannot upgrade directly to 4.3.1.

1. Upgrade from your current release to 4.1.x.x (follow the upgrade instructions in the latest Greenplum Database 4.1.x.x release notes available on [Support Zone](#)).
2. Upgrade from the current release to 4.2.x.x (follow the upgrade instructions in the latest Greenplum Database 4.2.x.x release notes available at [Pivotal Documentation](#)).
3. Follow the upgrade instructions in these release notes for [Upgrading from 4.2.x.x to 4.3.1](#).

### For Users Running Greenplum Database 3.3.x.x

Users on a release prior to 4.0.x.x cannot upgrade directly to 4.3.1.

1. Upgrade from your current release to the latest 4.0.x.x release (follow the upgrade instructions in the latest Greenplum Database 4.0.x.x release notes available on [Support Zone](#)).
2. Upgrade the 4.0.x.x release to the latest 4.1.x.x release (follow the upgrade instructions in the latest Greenplum Database 4.1.x.x release notes available on [Support Zone](#)).
3. Upgrade from the 4.1.1 release to the latest 4.2.x.x release (follow the upgrade instructions in the latest Greenplum Database 4.2.x.x release notes available at [Pivotal Documentation](#)).

4. Follow the upgrade instructions in these release notes for [Upgrading from 4.2.x.x to 4.3.1](#).

### Troubleshooting a Failed Upgrade

If you experience issues during the migration process and have active entitlements for Greenplum Database that were purchased through Pivotal, contact Pivotal Support. Information for contacting Pivotal Support is at <https://support.gopivotal.com>.

#### Be prepared to provide the following information:

- A completed [Upgrade Procedure](#).
- Log output from `gpmigrator` and `gpcheckcat` (located in `~/gpAdminLogs`)

## Greenplum Database Tools Compatibility

### Client Tools

Greenplum releases a number of client tool packages on various platforms that can be used to connect to Greenplum Database and the Greenplum Command Center management tool. The following table describes the compatibility of these packages with this Greenplum Database release.

Tool packages are available from [Pivotal Network](#).

**Table 4** Greenplum Database Tools Compatibility

| Client Package           | Description of Contents                                                                           | Client Version | Server Versions |
|--------------------------|---------------------------------------------------------------------------------------------------|----------------|-----------------|
| Greenplum Clients        | Greenplum Database Command-Line Interface (psql)<br>Greenplum MapReduce (gmapreduce) <sup>1</sup> | 4.3            | 4.3             |
| Greenplum Connectivity   | Standard PostgreSQL Database Drivers (ODBC, JDBC)<br>PostgreSQL Client C API (libpq)              | 4.3            | 4.3             |
| Greenplum Loaders        | Greenplum Database Parallel Data Loading Tools (gpfdist, gpload)                                  | 4.3            | 4.3             |
| Greenplum Command Center | Greenplum Database management tool.                                                               | 1.2.0.1        | 4.3             |

1. gmapreduce is not available on Windows.

The Greenplum Database Client Tools, Load Tools, and Connectivity Tools are supported on the following platforms:

- AIX 5.3L (32-bit)
- AIX 5.3L and AIX 6.1 (64-bit)
- Apple OSX on Intel processors (32-bit)
- HP-UX 11i v3 (B.11.31) Intel Itanium (Client and Load Tools only)
- Red Hat Enterprise Linux i386 (RHEL 5)
- Red Hat Enterprise Linux x86\_64 (RHEL 4)
- Red Hat Enterprise Linux x86\_64 (RHEL 5 and RHEL 6)
- SUSE Linux Enterprise Server x86\_64 (SLES 10 and SLES 11)
- Solaris 10 SPARC32
- Solaris 10 SPARC64
- Solaris 10 i386
- Solaris 10 x86\_64
- Solaris 9 SPARC32
- Windows 7 (32-bit and 64-bit)
- Windows Server 2003 R2 (32-bit and 64-bit)
- Windows Server 2008 R2 (64-bit)



- Windows XP (32-bit and 64-bit)

### Greenplum GPText

GPText enables processing mass quantities of raw text data (such as social media feeds or e-mail databases) into mission-critical information that guides business and project decisions. GPText joins the Greenplum Database massively parallel-processing database server with Apache Solr enterprise search.

GPText requires Greenplum Database. See the GPText release notes for the required version of Greenplum Database.

---

## Greenplum Database Extensions Compatibility

Greenplum Database delivers an agile, extensible platform for in-database analytics, leveraging the system's massively parallel architecture. Greenplum Database enables turn-key in-database analytics with Greenplum extensions.

You can download Greenplum extensions packages from [Pivotal Network](#) and install them using the Greenplum Packager Manager (`gppkg`). See the *Greenplum Database Utility Guide* for details.

Note that Greenplum Package Manager installation files for extension packages may release outside of standard Database release cycles. Therefore, for the latest install and configuration information regarding any supported database package/extension, go to the [Support](#) site and download [Primus Article 288189](#) from our knowledge base (Requires a valid login to the EMC Support site).

The following table provides information about the compatibility of the Greenplum Database Extensions and their components with this Greenplum Database release.

Note that the PL/Python database extension is already included with the standard Greenplum Database distribution.

**Table 5** Greenplum Database Extensions Compatibility

| Greenplum Database Extension               | Extension Components |                    |
|--------------------------------------------|----------------------|--------------------|
|                                            | Name                 | Version            |
| PostGIS 2.0 for Greenplum Database 4.3.x.x | PostGIS              | 2.0.3              |
|                                            | Proj                 | 4.8.0              |
|                                            | Geos                 | 3.3.8              |
| PostGIS 1.0 for Greenplum Database         | PostGIS              | 1.4.2              |
|                                            | Proj                 | 4.7.0              |
|                                            | Geos                 | 3.2.2              |
| PL/Java 1.1 for Greenplum Database 4.3.x.x | PL/Java              | Based on 1.4.0     |
|                                            | Java JDK             | 1.6.0_26 Update 31 |

**Table 5** Greenplum Database Extensions Compatibility

| Greenplum Database Extension                | Extension Components |                                                  |
|---------------------------------------------|----------------------|--------------------------------------------------|
|                                             | Name                 | Version                                          |
| PL/R 1.0 for Greenplum Database 4.3.x.x     | PL/R                 | 8.3.0.12                                         |
|                                             | R                    | 2.13.0                                           |
| PL/Perl 1.2 for Greenplum Database 4.3.x.x  | PL/Perl              | Based on PostgreSQL 9.1                          |
|                                             | Perl                 | 5.12.4 on RHEL 6.x<br>5.5.8 on RHEL 5.x, SUSE 10 |
| PL/Perl 1.1 for Greenplum Database          | PL/Perl              | Based on PostgreSQL 9.1                          |
|                                             | Perl                 | 5.12.4 on RHEL 5.x, SUSE 10                      |
| PL/Perl 1.0 for Greenplum Database          | PL/Perl              | Based on PostgreSQL 9.1                          |
|                                             | Perl                 | 5.12.4 on RHEL 5.x, SUSE 10                      |
| Pgcrypto 1.1 for Greenplum Database 4.3.x.x | Pgcrypto             | Based on PostgreSQL 8.3                          |
| MADlib 1.5 for Greenplum Database 4.3.x.x   | MADlib               | Based on MADlib version 1.8                      |

Greenplum Database 4.3 supports these minimum Greenplum Database extensions package versions.

**Table 6** Greenplum Database 4.3 Package Version

| Greenplum Database Extension | Minimum Package Version |
|------------------------------|-------------------------|
| PostGIS                      | 2.0.3                   |
| PL/Java                      | 1.1                     |
| PL/Perl                      | 1.2                     |
| PL/R                         | 1.0                     |
| Pgcrypto                     | 1.1                     |
| MADlib                       | 1.5                     |

### Package File Naming Convention

For Greenplum Database 4.3, this is the package file naming format.

*pkgname-ver\_pvpkg-version\_gpdbrel-os-version-arch.gppkg*

This example is the package name for a postGIS package.

*postgis-ossv2.0.3\_pv2.0\_gpdb4.3-rhel5-x86\_64.gppkg*

*pkgname-ver* - The package name and optional version of the software that was used to create the package extension. If the package is based on open source software, the version has format *ossvversion*. The *version* is the version of the open source software that the package is based on. For the postGIS package, *ossv2.0.3* specifies that the package is based on postGIS version 2.0.3.

*pvpkg-version* - The package version. The version of the Greenplum Database package. For the postGIS package, *pv2.0* specifies that the Greenplum Database package version is 2.0.

*gpdbr-el-OS-version-arch* - The compatible Greenplum Database release. For the postGIS package, *gpdbr4.3-rhel15-x86\_64* specifies that package is compatible with Greenplum Database 4.3 on Red Hat Enterprise Linux version 5.x, x86 64-bit architecture.

## Hadoop Distribution Compatibility

This table lists the Hadoop extensions compatibility matrix:

**Table 7** Hadoop Extensions Compatibility

| Hadoop Distribution | Version                     |
|---------------------|-----------------------------|
| Pivotal HD          | Pivotal HD 1.0 <sup>1</sup> |
| Greenplum HD        | Greenplum HD 1.1            |
|                     | Greenplum HD 1.2            |
| Cloudera            | cdh3u2                      |
|                     | cdh3u4                      |
|                     | CDH4.1 with MRv1            |
| Greenplum MR        | Greenplum MR 1.0            |
|                     | Greenplum MR 1.2            |

1. A distribution of Hadoop 2.0

**Note:** With Greenplum Database 4.3.1 and later, installing the gNet package is not required to use the gphdfs protocol that accesses data from Hadoop file systems. The jar file for the extensions, the libraries, and the documentation for the gphdfs extensions required to connect to Hadoop file systems is bundled with Greenplum Database. The files are installed in `$GPHOME/lib/hadoop`.

## Greenplum Database 4.3.1 Documentation

For the latest Greenplum Database documentation go to [Pivotal Documentation](#). Greenplum documentation is provided in PDF format.

**Table 8** Greenplum Database Documentation

| Title                                                   | Revision |
|---------------------------------------------------------|----------|
| Greenplum Database 4.3.1 Release Notes                  | A04      |
| Greenplum Database 4.3 Installation Guide               | A02      |
| Greenplum Database 4.3 Administrator Guide <sup>1</sup> | A01      |
| Greenplum Database 4.3 Reference Guide                  | A02      |
| Greenplum Database 4.3 Utility Guide                    | A02      |
| Greenplum Database 4.3 Client Tools for UNIX            | A02      |
| Greenplum Database 4.3 Client Tools for Windows         | A02      |
| Greenplum Database 4.3 Connectivity Tools for UNIX      | A02      |
| Greenplum Database 4.3 Connectivity Tools for Windows   | A02      |
| Greenplum Database 4.3 Load Tools for UNIX              | A02      |
| Greenplum Database 4.3 Load Tools for Windows           | A02      |
| Greenplum Command Center 1.2.2 Administrator Guide      | A01      |

1. For Greenplum Database release 4.3.1, the *Greenplum Database Administrator Guide* combines the information from the *Greenplum Database Database Administrator Guide* and the *Greenplum Database System Administrator Guide* that are available for previous Greenplum Database releases.

**Copyright © 2015 Pivotal Software, Inc. All rights reserved.**

Pivotal Software, Inc. believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." PIVOTAL SOFTWARE, INC. ("Pivotal") MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any Pivotal software described in this publication requires an applicable software license.

All trademarks used herein are the property of Pivotal or their respective owners.