

Welcome to Greenplum Database 4.0.4.0

Greenplum Database is a massively parallel processing (MPP) database server designed to support the next generation of data warehousing and large-scale analytics processing. It allows a cluster of servers to operate as a single database super computer — automatically partitioning data and parallelizing queries — to achieve performance tens or hundreds times faster than traditional databases. It supports SQL and MapReduce parallel processing and data volumes that range from hundreds of Gigabytes, to tens to hundreds of Terabytes, to multiple Petabytes.

Greenplum Database 4.0.4 is the fourth service pack release for version 4.0. This service pack release includes the contents of 4.0.1.x, 4.0.2.x and 4.0.3.x plus additional new issue resolutions. Please refer to the following sections for more information about this release:

- [Resolved Issues in Greenplum Database 4.0.4.x](#)
- [Known Issues in Greenplum Database 4.0.4.x](#)
- [Upgrading to Greenplum Database 4.0.4.0](#)
- [Greenplum Database Documentation](#)

For the contents of previous 4.0.x releases, please see the release notes section of each corresponding release:

- [Greenplum Database 4.0.3.x Release Notes](#)
- [Greenplum Database 4.0.2.x Release Notes](#)
- [Greenplum Database 4.0.1.x Release Notes](#)

Resolved Issues in Greenplum Database 4.0.4.x

This section lists the customer reported issues that are now resolved in Greenplum Database 4.0.4.x:

Table A.1 Resolved Issues in 4.0.4.x

Issue Number	Category	Description
11184, 11188	Standby Master	<p>“Master Mirroring Synchronization Lost” Error</p> <p>In prior 4.0.x releases, the standby master host could possibly get into a state where it was unable to synchronize or catch up to the transaction state of the primary master. When this occurred, the user would see a "Master mirroring synchronization lost" error message, requiring a resynchronization of the standby master using the <code>gpinitstandby</code> utility. This condition that caused this error has been resolved in this release.</p>
11352	Fault Detection and Recovery	<p>“PANIC”, “XX000”, “Append-Only Mirror Resync EOFs intent count...” Error</p> <p>In prior 4.0.x releases, if a system crash (such as a power failure) occurred while a primary and its mirror were in the process of resynchronizing append-only table data, it could cause both the primary and its mirror segment to become unavailable and unrecoverable through segment recovery utility (<code>gprecoverseg</code>). A symptom of this issue was Greenplum Database server and utility log messages such as:</p> <pre>PANIC", "XX000", "Append-Only Mirror Resync EOFs intent count... FATAL: DTM initialization: failure during startup/recovery... ... stderr:"failed: server did not respond with enough data"...</pre> <p>The condition that caused these errors has been resolved in this release.</p>
11354	Management Utilities	<p>Timeout Option Added to <code>gpstop</code></p> <p>When stopping Greenplum Database using <code>gpstop</code>, a segment may not complete shutdown in the allowed timeout of 60 seconds, especially if there are large transactions to roll back. To avoid shutdown errors, a new <code>-t</code> option was added to <code>gpstop</code> to allow the user to increase the allowed timeout. For example:</p> <pre>gpstop -t <timeout_in_seconds></pre>
11393	Data Loading	<p>COPY Errors when Loading Indexed Append-Only (AO) Tables</p> <p>In prior 4.0.x releases, using concurrent <code>COPY</code> commands to load data into append-only tables with indexes could fail with errors such as:</p> <pre>Unexpected internal error: Master process x received signal SIGSEGV... ...Error in copy_in_error_callback...</pre> <p>These errors could also cause primary segments to fail, requiring a segment recovery operation (<code>gprecoverseg</code>). This issue is now resolved in this release.</p>
11595	Fault Detection and Recovery	<p>Failed Segments Not Marked as Down by Fault Detection Process</p> <p>In prior 4.0.x releases, when a segment host went down due to file system problems (such as a corrupted/missing mount point or disk failure), the fault detection process of Greenplum Database did not mark all segments on that host down as expected. When a subsequent failure on another segment host occurred, Greenplum Database could potentially attempt to fail over to a mirror segment that was no longer available, leaving the system in an unresponsive state. This issue is now resolved in this release.</p>

Known Issues in Greenplum Database 4.0.4.x

This section lists the known issues in Greenplum Database 4.0.4.x. A work-around is provided where applicable.

Table 2 Known Issues in 4.0.4.x

Issue	Description
1589	<p>PostgreSQL Usage Statistics Views and Functions do not Work as Expected</p> <p>PostgreSQL has a number of views (<code>pg_stat_*</code>, <code>pg_statio_*</code>) for showing usage statistics. All of these views only report on the usage of the master (system catalogs), not the usage of user data on the segments. Many of the PostgreSQL statistics functions have the same problem. For example, <code>pg_stat_get_tuples_inserted()</code> shows only those inserts into the master (usually 0), not the number inserted into the table in question.</p>
3125, 3213	<p>Transaction Within a Function Not Recognized as a Sub-Transaction</p> <p>When a function containing multiple transaction blocks is run and an error occurs in one transaction block, the entire function exits with the errors:</p> <pre>ERROR: The distributed transaction 'Prepare' broadcast failed to one or more segments ERROR: current transaction is aborted, commands ignored until end of transaction block</pre>
5517	<p>Deadlock Detected when Concurrently Altering Tables with Bitmap Indexes</p> <p>Concurrent <code>ALTER TABLE</code> commands on tables that have bitmap indexes may encounter a deadlock and be unable to proceed. This issue occurs with concurrent <code>ALTER TABLE</code> commands that set distribution policies on tables that have bitmap indexes.</p> <p>Work-around: if you need to run concurrent <code>ALTER TABLE</code> commands on tables with bitmap indexes (as when running <code>gpexpand</code> with <code>-n</code> for multiple parallel processes), first drop the bitmap indexes and reinstate them when the concurrent <code>ALTER TABLE</code> commands are completed.</p>
5647	<p>Upgrade Utility Requires Standard Prefixes for Segment Data Directories</p> <p>The <code>gpmigrator</code> utility requires a single, standard segment prefix to perform a successful upgrade. System configurations that have been modified with multiple or omitted segment prefix values in segment data directories are not recognized as valid by <code>gpmigrator</code>, and cannot be upgraded using this utility.</p> <p>The naming convention for data directories in a Greenplum Database system is <code>SEG_PREFIXnumber</code> where number starts with 0 for segment instances (the master is always -1). So for example, if you choose the prefix <code>gp</code>, your master instance data directory would be named <code>gp-1</code>, and the segment instances would be named <code>gp0</code>, <code>gp1</code>, <code>gp2</code>, and so on.</p>
6279	<p>Append-Only Table Error After ALTER TABLE...DROP COLUMN</p> <p>After altering an append-only table to drop a column, users encounter the following error when trying to access the table:</p> <pre>ERROR: cache lookup failed for type 0</pre> <p>This issue is under investigation.</p> <p>Work-around: If you need to drop a column from an append-only table, do not use <code>ALTER TABLE</code>. Instead recreate the append-only table with the desired columns using <code>CREATE TABLE ...AS SELECT</code>.</p>
8445	<p>CREATE DATABASE ERROR: "template1" is being accessed by other users</p> <p>When the system has segments in <code>resynchronizing</code> mode, administrators will not be able to create a new database. Running a <code>CREATE DATABASE</code> command when segments are in the process of resynchronizing will result in the following error:</p> <pre>ERROR: source database "template1" is being accessed by other users...</pre> <p>Work-around: Use <code>gpstate -e</code> to check the status. When all segments are <code>synchronized</code>, you may retry the <code>CREATE DATABASE</code> command.</p>

Table 2 Known Issues in 4.0.4.x

Issue	Description
9968	<p>Slow gp_dump / gpccrondump Performance</p> <p>Prior to executing a dump operation, the dump utilities need to look up metadata information in the system catalogs. On databases with thousands of tables and millions of column attributes, this portion of the dump operation can take a long time.</p> <p>Work-around: Run <code>VACUUM ANALYZE</code> on the system catalog tables before running a dump operation. Prior to running a dump operation, run the following commands for the role that is executing the dump utilities. For example:</p> <pre>ALTER ROLE gpadmin SET enable_nestloop = on; ALTER ROLE gpadmin SET random_page_cost = 10;</pre> <p>After the dump is complete, return to the default settings:</p> <pre>ALTER ROLE gpadmin RESET enable_nestloop; ALTER ROLE gpadmin RESET random_page_cost;</pre>
10028	<p>After Running gprecoverseg, Segments May Not Be Running in Their Preferred Role</p> <p>In 4.0.4, when a primary segment goes down, the mirror is activated and becomes the primary segment. After running <code>gprecoverseg</code>, the currently active segment remains the primary and the failed segment is then brought up as the mirror. The segment instances are not returned to the preferred role that they were given at system initialization time. This can leave the system in a potentially unbalanced state, as some segment hosts may still have more active segments that is optimal for top system performance.</p> <p>After recovering a failed segment, the roles may be reversed (the original mirror is now the primary and vice-versa). The <code>gpstate -e</code> command will show you if you have segments in this condition.</p> <p>Work-around: Restart Greenplum Database (<code>gpstop -r</code>). A restart will bring up all segments in their preferred role. The Greenplum Database system will be unavailable for a brief period while it restarts.</p>
10278	<p>pg_relation_size and pg_database_size do not Account for Custom Filespaces</p> <p>The <code>pg_relation_size</code> and <code>pg_database_size</code> functions are used to calculate the size of a relation or database in a distributed Greenplum Database system. These functions do not currently account for relations stored in any tablespaces besides the default <code>pg_system</code> tablespace.</p>
11667	<p>gpupgrademirror Errors when Upgrading Greenplum Database from 3.3.x to 4.0.x</p> <p>When running the Greenplum Database upgrade utility (<code>gpmigrator</code>), the <code>gpupgrademirror</code> sub-utility uses the system catalog tables to determine the database files it needs to copy from each primary segment to its respective mirror segment. In rare cases, index file records in the catalog may not match the actual index files found on the segments. When this type of mismatch occurs and an expected index file does not exist on a segment, users will see an error such as:</p> <p>Exception: There was a problem with one of the <code>gpupgrademirror</code> sub processes.</p> <p>Work-around: If you receive errors during the mirror upgrade process, Contact EMC Technical Support and provide them with the <code>/export/home/gpadmin/gpAdminLogs/gpupgrademirror.py_date.log</code> log file.</p>
11668	<p>Append Only Table Performance Issues on Linux Platforms</p> <p>XFS file system settings in the Red Hat Enterprise Linux 5.5 kernel can cause the fragmentation of the data files stored on disk for append-only (AO) tables. This fragmentation degrades query performance for queries that require AO table scan operations.</p> <p>Work-around: Update the Red Hat Enterprise Linux 5.5 kernel to version 2.6.18-194.26.1.el5 on all Greenplum hosts, and use the following recommended mount options for XFS:</p> <pre>rw,noatime,inode64,allocsize=16m</pre>

Table 2 Known Issues in 4.0.4.x

Issue	Description
11696	<p>gprecoverseg: Multi-Segment Recovery Process Unresponsive</p> <p>When the segment recovery utility (<code>gprecoverseg</code>) is run, it brings the failed segment instances back online one-by-one instead of in parallel. In cases where there are multiple segment primary/mirror pairs to recover (such as in a multi-host failure), the ordering of the sequential segment recovery can prevent a primary and its corresponding mirror from establishing a connection to each other in the allowed timeout. When this occurs, the <code>gprecoverseg</code> utility becomes unresponsive and all segments are not recovered as expected.</p> <p>Work-around: Run <code>gprecoverseg</code> again or increase the Greenplum configuration parameter <code>gp_segment_connect_timeout</code> to a value greater than the default (3 minutes).</p>
11707	<p>Cannot Start Greenplum Database when a NIC is Down</p> <p>The Greenplum Database startup utility, <code>gpstart</code>, fails to start Greenplum Database if a segment host has a failed network interface (NIC). When this occurs, the <code>gpstart</code> utility becomes unresponsive.</p> <p>Work-around: Repair the failed NIC and run <code>gpstart</code> again.</p>

Upgrading to Greenplum Database 4.0.4.0

For detailed upgrade procedures, see the following sections:

- [Upgrading From 4.0.x.x to 4.0.4.0](#)
- [Upgrading From 3.3.x.x to 4.0.4.0](#)
- [Upgrading From 3.2.x.x to 4.0.4.0](#)
- [Upgrading from Releases Prior to 3.2.x](#)

Upgrading From 4.0.x.x to 4.0.4.0

These instructions are for users currently running 4.0.x.x. An upgrade from 4.0.x.x to 4.0.4.0 involves removing the standby master from your Greenplum system (if configured), 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. If you have a standby master configured, remove it from your Greenplum configuration:

```
$ gpinitstandby -r
```

3. Perform a smart shutdown of your current Greenplum Database 4.0.0.x system (there can be no active connections to the database):

```
$ gpstop
```

4. Run the installer for 4.0.4.0 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-4.0.4.0`

5. 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-4.0.0.8/greenplum_path.sh
```

to:

```
source /usr/local/greenplum-db-4.0.4.0/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-4.0.4.0
/usr/local/greenplum-db
```

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

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

7. Log in as root and install the 4.0.4.0 binaries on all of the segment hosts. Tar the Greenplum installation directory, use `gpscp` to copy it to the segment hosts, and use `gpssh` to untar it at the segment hosts. Make sure that the `gpadmin` user owns the new installation directory. Also update any symbolic links to point to the new version. For example:

```
$ su -
# gtar -cvf gp.tar /usr/local/greenplum-db-4.0.4.0
# gpscp -f seghostname_file gp.tar =:/
# gpssh -f seghostname_file
=> gtar --directory / -xvf /gp.tar
=> chown -R gpadmin /usr/local/greenplum-db-4.0.4.0
=> rm /gp.tar
=> rm /usr/local/greenplum-db
=> ln -s /usr/local/greenplum-db-4.0.4.0
    /usr/local/greenplum-db
```

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

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

9. If your original system had a standby master, add the standby master back into your Greenplum configuration:

```
$ gpinitstandby -s standby_hostname
```

Upgrading From 3.3.x.x to 4.0.4.0

You must be on release 3.3.x.x in order to use the 4.0 upgrade utility. The Greenplum upgrade utility (`gpmigrator`) makes the required changes to the system catalogs without requiring users to dump and restore their databases.

Greenplum strongly recommends that you perform a backup of your databases before running the upgrade utility. If you find issues when testing your upgraded system, you can restore this backup.

- [Upgrade Checklist](#)
- [Upgrade Procedure](#)
- [Troubleshooting a Failed Upgrade](#)

Upgrade Checklist

This checklist provides a quick overview of all the steps required for an upgrade from 3.3.x.x to 4.0.4.0. Detailed upgrade instructions are also provided in the [Upgrade Procedure](#) section.

Pre-Upgrade Preparation (on your current system)

* 3.3.x.x system is up and available

- Log in to your master host as the `gpadmin` user (your Greenplum superuser).

<input type="checkbox"/>	Install the Greenplum Database 4.0.4.0 binaries on all Greenplum hosts.
<input type="checkbox"/>	Copy any custom modules from your current installation to your 4.0.4 installation on all Greenplum hosts. For example, shared library files for user-defined functions in <code>\$GPHOME/lib</code> or PostgreSQL add-on modules (such as <code>plr.so</code> or <code>pgcrypto.so</code>) in <code>\$GPHOME/lib/postgresql</code> .
<input type="checkbox"/>	Drop the <code>gp_jetpack</code> schema from all databases (if it exists). Also note that <code>gp_</code> is now a reserved prefix for schema names in Greenplum Database 4.0 (as is <code>pg_</code>). If you have any user-created schemas named with a <code>gp_</code> prefix, Greenplum recommends that you rename them prior to upgrading.
<input type="checkbox"/>	(Optional) Run <code>VACUUM</code> on all databases, and remove old server log files from <code>pg_log</code> in your master and segment data directories. This is not required, but will reduce the size of Greenplum Database files to be backed up and migrated.
<input type="checkbox"/>	Check for and recover any failed segments (<code>gpstate</code> , <code>gprecoverseg</code>).
<input type="checkbox"/>	(Optional) Run the <code>gpcheckcat</code> utility to validate your system catalogs. If any errors are reported, contact Greenplum Customer Support.
<input type="checkbox"/>	Run the <code>gpssh-exkeys</code> utility to exchange keys between all configured hostnames in your Greenplum Database system.
<input type="checkbox"/>	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.
Upgrade Execution	
<i>* The system will be locked down to all user activity during the upgrade process</i>	
<input type="checkbox"/>	Backup your current databases (<code>gpcrondump</code> or ZFS snapshots). If you find issues when testing your upgraded system, you can restore using this backup.
<input type="checkbox"/>	Remove the standby master from your system configuration (<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 your Greenplum Database 4.0.4 installation.
<input type="checkbox"/>	Run <code>gpmigrator</code>
<input type="checkbox"/>	After the upgrade process finishes successfully, your 4.0.4 system will be up and running.
Post-Upgrade (on your 4.0.4 system)	
<i>* The 4.0.4 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 by installing the version 4.0.4 Load Tools package.
<input type="checkbox"/>	Run the <code>gpbitmapreindex</code> utility to rebuild your bitmap indexes. The on-disk format of bitmap indexes has changed in 4.0.4.
<input type="checkbox"/>	Install the new 4.0.4 Performance Monitor Web Application and update your environment to point to the 4.0.4 Performance Monitor binaries.
<input type="checkbox"/>	Inform all database users of the completed upgrade. Tell users to update their environment to source the Greenplum Database 4.0.4 installation (if necessary).

Upgrade Procedure

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



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 3.3.x system)

Perform these steps on your current 3.3.x Greenplum Database system. This procedure is performed from your Greenplum master host and should be executed by the Greenplum superuser (`gpadmin`).

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


```
$ su - gpadmin
```
2. Drop the `gp_jetpack` schema from all databases (if it exists). This Greenplum administrative schema has been renamed to `gp_toolkit` in 4.0. The `gp_toolkit` schema will be installed in all databases by default after upgrading to 4.0.4. For example, to drop `gp_jetpack`:


```
psql template1 -c 'DROP SCHEMA IF EXISTS gp_jetpack
CASCADE;'
```

```
psql mydatabase -c 'DROP SCHEMA IF EXISTS gp_jetpack
CASCADE;'
```
3. (optional) Vacuum all databases prior to upgrade. For example:


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


```
$ gpssh -f seg_host_file -e 'rm /gpdata/*/gp*/pg_log/*.csv'
```
5. Install the Greenplum Database 4.0.4.0 binaries on all Greenplum hosts. See the *Greenplum Database 4.0 Installation Guide* for detailed instructions.

6. If your current installation is utilizing any custom modules or PostgreSQL add-on modules, make sure the required library files are copied into the correct directory of your 4.0.4.0 installation on all Greenplum hosts. Shared library files for user-defined functions go in `$GPHOME/lib` and PostgreSQL add-on modules (such as `plr.so` or `pgcrypto.so`) go in `$GPHOME/lib/postgresql`. For example:

```
$ gpscp -f all_hosts_file \
  /usr/local/greenplum-db-3.3.6.0/lib/postgresql/plr.so \
  =:/usr/local/greenplum-db-4.0.4.0/lib/postgresql/plr.so
```

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

```
$ gpstate
```

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

```
$ gprecoverseg
```

9. Run the `gpssh-exkeys` utility referencing a host file (`all_hosts_file`) that has the host name for each master, standby master, and segment host. Include all configured interface names for each host. For example:

```
$ gpssh-exkeys -f all_hosts_file
```

10. (optional) Run the `gpcheckcat` utility to validate your system catalogs. For example:

```
$ $GPHOME/bin/lib/gpcheckcat
```

If any errors are reported, contact Greenplum Customer Support.

Upgrade Execution

During the migration process, all client connections to the master will be locked out. The migration utility locks out all client access to the master, however it does not block direct utility mode access to the segments. To ensure a safe upgrade, make sure users are aware that any connections to a Greenplum segment are not safe during the upgrade time frame.

11. (optional but strongly recommended) Back up all databases in your Greenplum Database system. If you find issues when testing your upgraded system, you can restore this backup.

You should also make a backup of your database roles and server configuration files. For example, using the 3.3.x version of `gpcrondump` (make sure you move the backup files out of your segment data directories):

```
$ /usr/local/greenplum-db-3.3.7.2/bin/gpcrondump -x database
-c -g -G -p -u /backup_target_directory
```

(Solaris Users) By default, the upgrade utility will pause on Solaris systems to allow for a ZFS snapshot to be taken. Make sure you take a snapshot of all segment file systems and the master file system. All Greenplum data directories must be on ZFS file systems. See the *Sun ZFS Administration Guide* for more information: <http://opensolaris.org/os/community/zfs/docs/zfsadmin.pdf>. If you already have a recent set of snapshots and do not want `gpmigrator` to pause mid-upgrade, run `gpmigrator` with the `--nosnapshot` option.

- 12.** If your system has a standby master host configured, remove the standby master from your system configuration. For example:

```
$ gpinitstandby -r
```

- 13.** Perform a clean shutdown of your current Greenplum system. For example:

```
$ gpstop -f
```

- 14.** Update your environment so that it is sourcing your new 4.0.4.0 installation.

- a.** For example, update the `greenplum-db` symbolic link on the master and standby master to point to the new 4.0.4.0 installation directory. For example (as root):

```
# rm -rf /usr/local/greenplum-db
# ln -s /usr/local/greenplum-db-4.0.4.0
  /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-4.0.4.0
  /usr/local/greenplum-db
=> chown -R gpadmin /usr/local/greenplum-db
=> exit
```

- c.** If your user profile file (such as `.bashrc`) does not use the `greenplum-db` symbolic link, you will need to update your profile file to source the new 4.0.4.0 installation. For example, update the following line in your profile file:

```
source /usr/local/greenplum-db-4.0.4.0/greenplum_path.sh
```

And then source the updated file. For example:

```
source ~/.bashrc
```

- 15.** As `gpadmin`, run the 4.0.4.0 version of `gpmigrator` specifying your old and new `GPHOME` locations. You will also need to specify a port range for the new 4.0 *replication* processes (these must use different ports from your *database listener* processes), the mode by which to create the new 4.0 mirrors, and the path to your Greenplum Performance Monitor installation (if your system has the monitor program installed and enabled). For example:

```
$ su - gpadmin
```

```
$ which gpmigrator
/usr/local/greenplum-db-4.0.4.0/bin/gpmigrator
```

```
$ gpmigrator --replication_port_base 53000 \
  --mirror_replication_port_base 63000 \
  --mirror_mode redundant \
  --gpperfmonhome /usr/local/greenplum-perfmon-2.1.0.0 \
  /usr/local/greenplum-db-3.3.7.2 \
```

```
/usr/local/greenplum-db-4.0.4.0
```

- 16.** (*Solaris only*) By default, the upgrade utility will pause on Solaris systems to allow for a ZFS snapshot to be taken. To resume the upgrade after the snapshots have been taken:

```
$ gp migrator --resume
   /usr/local/greenplum-db-3.3.7.2 \
   /usr/local/greenplum-db-4.0.4.0
```

- 17.** After the `gp migrator` utility has completed successfully, your Greenplum Database 4.0 system will be up and running. The `gp migrator` utility will take a while to complete (the time is proportional to the amount of data and the number of segments in your system).

Post-Upgrade (on your 4.0.4 system)

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

```
$ gp initstandby -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 Greenplum Database 4.0.4 Load Tools package. Packages are available on <http://gpn.greenplum.com>.

- 20.** The on-disk format of bitmap indexes has changed in 4.0.4. If your system has bitmap indexes, they must be reindexed (or dropped and recreated) after upgrading. Greenplum provides the `gpbitmapreindex` utility to facilitate the upgrade of any bitmap indexes in your system.

- 21.** If you are using the Greenplum Performance Monitor, install the new 4.0.4 Performance Monitor Web Application and update your environment to point to the new 4.0.4 Performance Monitor binaries (source the `gpperfmon_path.sh` file from your new 4.0.4 installation). Greenplum Performance Monitor 4.0.4 Web Application Packages are available on <http://gpn.greenplum.com>. Note that the Performance Monitor data collection agents are integrated with the Greenplum Database server in 4.0.4, and are started automatically by the `postgres` server process. You no longer need a separate CRON job to monitor and restart the `gpmmon` data collection process.

- 22.** Note that resource queues are enabled by default in 4.0.4. If you have any roles that are not assigned to a resource queue, they will be assigned to the `pg_default` resource queue after upgrade. For more information on configuring resource queues for workload management, see *Chapter 8* of the *Greenplum Database 4.0.4 Administrator Guide*.

- 23.** Check your `postgresql.conf` server configuration files for any deprecated parameters you have set, and remove those settings. See “[Deprecated Parameters](#)” on page 24. Greenplum Database 4.0.4 will issue warning messages on startup if it finds any deprecated parameters set in your `postgresql.conf` files.

Troubleshooting a Failed Upgrade

If you experience issues during the migration process, contact Greenplum customer support at 1-866-410-6060 or open a support incident. Authorized Customer Administrators can log a support incident on the [support portal](#). If you are a Greenplum Administrator at your company, and do not have access, please contact entitlement@greenplum.com.

Be prepared to provide the following information:

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

Upgrading From 3.2.x.x to 4.0.4.0

Users on a release prior to 3.3.x.x cannot upgrade directly to 4.0.4. You must first upgrade from your current release to 3.3 (follow the upgrade instructions in the Greenplum Database 3.3 release notes available on <http://gpn.greenplum.com>). After you have upgraded to 3.3, follow the upgrade instructions for [Upgrading From 3.3.x.x to 4.0.4.0](#).

Upgrading from Releases Prior to 3.2.x

If you are running a Greenplum Database release prior to 3.2.x.x, contact Greenplum customer support at 1-866-410-6060 or open a support incident. Authorized Customer Administrators can log a support incident on the [support portal](#). If you are a Greenplum Administrator at your company, and do not have access, please contact entitlement@greenplum.com.

Greenplum Database Documentation

The following Greenplum Database documentation is available in the `$GPHOME/docs` directory of your Greenplum installation, or you can go to <http://gpn.greenplum.com> to download the latest documentation:

GPInstallGuide.pdf - *Greenplum Database Installation Guide*

GPAdminGuide.pdf - *Greenplum Database Administrator Guide*

GPPerfmonAdminGuide - *Greenplum Performance Monitor Administrator Guide*

Greenplum Database 4.0.3.x Release Notes

This section lists the customer reported issues that were resolved in Greenplum Database 4.0.3.x:

Table A.1 Resolved Issues in 4.0.3.x

Issue Number	Category	Description
11406	Append-Only Tables	<p>Append-Only Tables: Concurrent Transaction Conflicts</p> <p>In prior 4.0.x.x releases, multiple concurrent transactions that wrote to the same append-only (AO) table could potentially conflict with each other, especially on AO tables with indexes. When this occurred, the conflicting transaction would fail with an error such as:</p> <pre>ERROR "...tuple concurrently updated..."</pre> <p>This issue has been resolved in this release.</p>
11399	Fault Detection	<p>Intermittent Segment Failover when Loading Append-Only Tables</p> <p>In prior 4.0.x.x releases, loading an append-only table could intermittently trigger a segment failover. This would only happen in cases when the segment being updated was not running in its preferred role (for example, when the mirror segment was currently acting as the primary). When this type of segment failure occurred, administrators would see an error such as the following in the segment log file:</p> <pre>WARNING", "58P02", "mirror failure, could not link to 'pg_xlog...': File exists, failover requested..."</pre> <p>This issue has been resolved in this release.</p>
11378	Data Loading	<p>Greenplum Loaders Package Missing Files on AIX</p> <p>The 4.0.1 release of the Greenplum Loaders client package was missing certain library files that prevented gpload from running. This issue has been resolved in this release.</p>
10834	Management Utilities	<p>Shutdown in Progress Error After a gpstop</p> <p>After stopping Greenplum Database using <code>gpstop</code>, you will not be able to restart Greenplum Database until shutdown completes on all segments. In some cases, if there are large transactions to roll back, a segment may not complete shutdown in the allowed timeout of 60 seconds. When this occurs, <code>gpstop</code> will now print an informational message if all segments do not complete shutdown in the allowed timeout. To increase the allowed <code>gpstop</code> timeout for future shutdown operations, use <code>gpstop -t <timeout_in_seconds></code>.</p>
11008	Data Loading	<p>Backslashes in gpload Configuration File Cause Error on Windows</p> <p>In prior 4.0.x releases, backslash characters in the <code>gpload</code> configuration file (such as a backslash character in a Windows file path) would cause the following warning message, even though the use of the backslash character was legal in this context:</p> <pre>HINT: Use the escape string syntax for backslashes, e.g., E'\' WARNING: nonstandard use of \ in a string literal</pre> <p>This issue has been resolved in this release.</p>
11181	Standby Master	<p>Filespace/Tablespace DDL Commands not Synchronized to Standby Master</p> <p>In prior 4.0.x releases, <code>CREATE</code> and <code>DROP FILESPACE/TABLESPACE</code> commands did not always synchronize the change to the standby master. This issue has been resolved in this release.</p>

Table A.1 Resolved Issues in 4.0.3.x

Issue Number	Category	Description
11198	PL/Java	<p>PL/Java Resources Cannot Be Retrieved from the CLASSPATH</p> <p>In prior 4.0.x releases, resources (other than class files) could not be retrieved from the <code>CLASSPATH</code>. For example:</p> <pre>Example.class.getResource("somefile")</pre> <p>Would always return <code>null</code> for non-class files. This issue has been resolved in this release.</p>
11216	System Expansion	<p>Incorrect gpexpand Warning: The current system appears to be non-standard...</p> <p>In prior 4.0.x releases, using the <code>gpexpand</code> utility in interactive mode would always give the following warning, even when the system was using a standard configuration:</p> <pre>The current system appears to be non-standard...</pre> <p>This issue has been resolved in this release.</p>
11257	Table Partitioning	<p>Splitting a Default Partition Does Not Preserve Table Distribution Policy</p> <p>In prior 4.0.x releases, <code>ALTER TABLE...SPLIT DEFAULT PARTITION</code> did not propagate the table distribution key to the new partition that was created. This issue has been resolved in this release.</p>
11371	gp_toolkit	<p>gp_toolkit Log Views: ERROR: data line too long</p> <p>In prior 4.0.x releases, using the <code>gp_toolkit.gp_log_*</code> family of views would sometimes give the following error if the Greenplum Database log files contained large log messages:</p> <pre>ERROR: data line too long...</pre> <p>To reduce the likelihood of this error, the parameter <code>gp_max_csv_line_length</code> has an increased default of 1MB (1048576 bytes) and an increased allowed maximum of 4MB.</p>
11389	Management Utilities	<p>Add --config Option to gpcheck</p> <p>The <code>gpcheck</code> utility determines the platform on which you are running Greenplum Database and validates various platform-specific configuration settings. The checks are specified in the default file <code>\$GPHOME/etc/gpcheck.cnf</code>. The <code>gpcheck</code> utility now accepts a <code>--config</code> option so you can override the default file with your own config file.</p>

Greenplum Database 4.0.2.x Release Notes

This section lists the customer reported issues that were resolved in Greenplum Database 4.0.2. Greenplum Database 4.0.2 was a feature and service pack release intended for use on the EMC Greenplum Data Computing Appliance (DCA):

Table A.2 Resolved Issues in 4.0.2.x

Issue Number	Category	Description
10343	Performance Monitor	Performance Monitor Log File Rotation In prior releases of Greenplum Performance Monitor, log files were truncated when the configured size limit was reached. Logging has been changed so that log files are now rolled over. See the <i>Greenplum Performance Monitor Administrator Guide</i> for more information on log rotation and maintenance.
11079	Performance Monitor	Performance Monitor Console Online Help Prior releases of Greenplum Performance Monitor did not have online help available from within the Performance Monitor Console web application. Online help is now provided.
11103	DCA	EMC Greenplum DCA - ConnectEMC Integration For Greenplum Database instances running on the EMC Greenplum Data Computing Appliance (DCA), administrators now have the option of configuring ConnectEMC phone-home support. With ConnectEMC enabled, certain system events (such as a segment host failures) will be captured and automatically sent to EMC Customer Support.
10893	DCA	EMC Greenplum DCA - Performance Monitor Health Monitoring Tab For Greenplum Database instances running on the EMC Greenplum Data Computing Appliance (DCA), the Performance Monitor Console now has an additional Health Monitoring tab. This new tab shows the status of the various hardware components of the DCA.
11232	Management Utilities	New gpcheck Utility to Verify System Hardware and OS Settings This release introduces a new <code>gpcheck</code> utility that performs some pre-installation checks of the hardware and operating systems of hosts that comprise a Greenplum Database system. This utility replaces the <code>gpcheckos</code> utility.
10814	Management Utilities	Reload all postgresql.conf Files with gpstop -u In prior releases, a <code>gpstop -u</code> command would reload the <code>postgresql.conf</code> file of the master instance only. It now reloads all segment <code>postgresql.conf</code> files as well.

Greenplum Database 4.0.1.x Release Notes

Greenplum Database 4.0 is a major release which introduces a number of significant new features, performance and stability enhancements, and enhancements to the product architecture. Version 4.0.1 is the first general availability (GA) release of Greenplum Database 4.0, and contains all issue resolutions reported in the 4.0.0.0-4.0.0.8 controlled releases. Please refer to the following sections for more information about this release:

- [New Features in Greenplum Database 4.0](#)

- [Changed Features in Greenplum Database 4.0](#)
- [Resolved Issues in Greenplum Database 4.0.1.x](#)
- [Known Issues in Greenplum Database 4.0.4.x](#)
- [Upgrading to Greenplum Database 4.0.4.0](#)
- [Greenplum Database Documentation](#)
- [Copyright Information](#)

New Features in Greenplum Database 4.0

Greenplum Database 4.0 offers the following new features:

- [Enhanced Workload Management with Dynamic Query Prioritization](#)
- [Self Healing Fault Tolerance Model with Differential Online Recovery](#)
- [Direct Dispatch Performance Optimization of Single Row Operation](#)
- [MPP Tablespace Support for Non-Uniform and SSD Segment Storage](#)
- [B-Tree and Bitmap Indexes on Column-Oriented and Append-Only Tables](#)
- [Health Monitoring Infrastructure with Email and SNMP Alerting](#)
- [Writable External Tables for Parallel Data Output](#)
- [Object-level 'Metadata Management' Tracking and Querying](#)
- [Enhanced Global Statistics Collection](#)
- [MapReduce Support for C Language Functions](#)
- [Support for User Defined Functions in PL/Java](#)

Enhanced Workload Management with Dynamic Query Prioritization

Prior releases of Greenplum Database have included a range of workload management capabilities to allow database administrators (DBAs) to manage the resources allocated to query workloads. The primary mechanism has been role-based resource queues, which provide configurable query admission limits. By using resource queues to set limits on incoming queries, DBAs can control the number and complexity of active queries on the system at any given time, thereby protecting the system from over allocation of resources. Prior to 4.0, DBAs had to explicitly enable resource queues. In 4.0, resource queues are now always enabled.



Important: Resource queues are required for all roles (users) in Greenplum Database 4.0.4. Any role not explicitly assigned to a resource queue will be assigned to the default resource queue, `pg_default`.

In addition to resource queues, Greenplum Database 4.0 adds a dynamic query prioritization infrastructure. Each query in the system has a priority value, which determines the relative share of system resources provided to it. The priority of a query is initially determined by the priority set on the resource queue through which it

enters. However, administrators also have the ability to adjust priority at runtime. This feature allows DBAs to control processing resources and ensure that important workloads can run with minimal interference from lower priority jobs.

Self Healing Fault Tolerance Model with Differential Online Recovery

In Greenplum Database 4.0, data redundancy (mirroring) is now performed using *physical block replication*. The primary and mirror segments are kept in sync at the physical disk block level, and changes to the primary are automatically applied to the mirror in a transactionally consistent manner. This new mirroring architecture offers a number of improvements over prior releases:

- **Automatic Failure Detection and Failover.** Should a segment server become unavailable, the system will automatically detect the failure and promote the necessary mirror segments to maintain full read/write operation. There is no longer a need to specify a fault action mode (*read-only* or *continue*).
- **Fast Differential Recovery.** Greenplum Database 4.0 keeps track of the changes that are made while a segment is down. When a failed segment becomes available again, only the modified disk blocks (as opposed to the entire contents) are copied over from the mirror. This ensures the fastest possible recovery time.
- **No Downtime for Segment Recovery.** Segment recovery takes place in the background while the system is fully online. The database is fully available and can support read/write operations while recovery is in progress.
- **Improved Write Performance for AO Tables.** Write transactions for compressed append-only tables are only processed once at the primary segments, and segment mirroring ensures that all modified disk blocks are synchronized to the mirrors.

In prior releases, Greenplum Database used *logical database replication* to maintain a mirror copy of a segment instance. This meant that a statement issued to Greenplum Database, such as an `INSERT`, was run on a primary segment first and then again on its corresponding mirror segment. While this was an effective technique for data redundancy, the new physical block replication infrastructure has a number of functional and performance advantages. This new infrastructure will also be the basis for future Greenplum Database high-availability and replication features.

Direct Dispatch Performance Optimization of Single Row Operation

Greenplum Database 4.0 introduces a performance enhancement to the query planning and dispatch process for small queries that only access data on a single segment (for example, a single-row `INSERT`, `UPDATE`, `DELETE` or `SELECT` statement). In queries such as these, the query plan is not dispatched to all segments, but is targeted to the segment that contains the affected row(s). This *direct dispatch* approach for this type of query dramatically reduces the response time and resource utilization of small queries.

MPP Tablespace Support for Non-Uniform and SSD Segment Storage

Greenplum Database 4.0 introduces support for tablespaces. Tablespaces allow database administrators to have multiple file systems per machine and decide how to best use their physical storage to store database objects. Tablespaces are useful for a

number of reasons, such as allowing different storage types for frequently versus infrequently used database objects, or controlling storage capacity and I/O performance on certain database objects. For example, highly utilized tables can be placed on file systems that use high performance solid-state drives (SSD), while the remaining tables utilize standard hard drives. This is an advanced feature for Greenplum system administrators who need greater control and flexibility over their database storage.

B-Tree and Bitmap Indexes on Column-Oriented and Append-Only Tables

In Greenplum Database 4.0, support for non-unique indexes has been added for append-only storage tables, including tables using compression and/or column-oriented storage. Indexes can greatly improve performance on compressed append-only tables for queries that return a targeted set of rows, as the optimizer now has the option to use an index access method rather than a full table scan when appropriate. For compressed data, an index access method means only the necessary rows are uncompressed.

Health Monitoring Infrastructure with Email and SNMP Alerting

Greenplum Database can now be configured to send email notifications to a system administrator whenever certain events occur, such as fatal server errors, segment failures, or system restarts.

Greenplum Database 4.0 also introduces support for SNMP. The Greenplum SNMP agent, `gpsnmpd`, can be configured to run on your Greenplum master host. This agent supports the standard relational database application management information base (`RDBMS-MIB.txt`) and can be polled by a network monitoring program, such as HP OpenView or Nagios. Greenplum Database can also be configured to send an SNMP notification to your network monitoring program when certain alert events occur (such as a segment failure). Greenplum Database supplies a custom management information base (`GPDB-MIB.txt`) to enable SNMP notifications for certain Greenplum Database events.

Writable External Tables for Parallel Data Output

Greenplum Database 4.0 now supports writable external tables, allowing users to perform high-speed parallel data output from a Greenplum Database instance to a file system, and ETL server, or other applications or databases. Writable external tables can be used in conjunction with Greenplum MapReduce to output job results to any external target. Writable external tables utilize the same *Scatter-Gather Streaming* infrastructure that is used when loading data.

Object-level 'Metadata Management' Tracking and Querying

Greenplum Database 4.0 now tracks metadata management information in its system catalogs about the objects stored in a database, such as tables, views, indexes and so on, as well as global objects such as roles and tablespaces. This allows administrators to examine information about an object, such as when it was created or what was the last operation performed. The system views `pg_stat_operations` and `pg_stat_partition_operations` can be used to look up actions performed on an object, such as a table. For example, you can use these views to see when a table was last vacuumed and analyzed.

Enhanced Global Statistics Collection

The `ANALYZE` command in Greenplum Database 4.0 now collects global database statistics from all active segments in the system, thereby providing the most accurate (and consistent) statistics for query planning and optimization. Previous versions of Greenplum Database would select a single segment to use as the basis for statistical data analysis, which did not always represent the true statistical variance of the data. Existing customers will be able to see the benefit of improved statistics collection the first time they run `ANALYZE` on a table after upgrading to 4.0. No other additional configuration is necessary.

MapReduce Support for C Language Functions

Greenplum MapReduce allows programmers who are familiar with the MapReduce programming paradigm to write map and reduce functions and submit them to the Greenplum Database parallel data flow engine for processing. Prior releases of Greenplum MapReduce provided language support for Perl and Python. In 4.0, developers can also use C functions. Both user-defined C functions and built-in database functions are supported.

Support for User Defined Functions in PL/Java

PL/Java is a PostgreSQL language extension that allows users to write custom database functions in the Java programming language. Implementation details can be found in `$GPHOME/share/postgresql/pljava/PLJAVA_README` of your 4.0 installation. Note that PL/Java support for Greenplum MapReduce is not included in this release.

Changed Features in Greenplum Database 4.0

This section describes the new and changed SQL commands, utilities and server configuration parameters for 4.0. Refer to the *Greenplum Database Administrator Guide* for detailed descriptions of items listed in this section.

- [SQL Commands](#)
- [Client Utilities](#)
- [Management Utilities](#)
- [Server Configuration Parameters](#)
- [System Catalogs](#)

SQL Commands

- [Changed SQL Commands](#)
- [New SQL Commands](#)

Changed SQL Commands

Table 2 Changed SQL Commands in 4.0

SQL Command	Description of Change
ALTER RESOURCE QUEUE	New syntax for setting resource queue attributes: WITH (<i>queue_attribute=value</i>) Deprecated (but still accepted) syntax: ACTIVE THRESHOLD, COST THRESHOLD OVERCOMMIT NOOVERCOMMIT, IGNORE THRESHOLD
ALTER ROLE	New role attributes to grant permission to create an external table: CREATEEXTERNAL NOCREATEEXTERNAL
ALTER TABLE	SET WITH (OIDS=TRUE) has been disabled for partitioned tables and append-only column-oriented tables.
COPY	New syntax to declare the newline format used in a load file: NEWLINE '[LF CR CRLF]'
CREATE DATABASE	SQL_ASCII disabled as a database ENCODING.
CREATE EXTERNAL TABLE	New syntax for WRITABLE external tables.
CREATE RESOURCE QUEUE	New resource queue attribute to set priority. New syntax for setting resource queue attributes: WITH (<i>queue_attribute=value</i>) Deprecated (but still accepted) syntax: ACTIVE THRESHOLD, COST THRESHOLD OVERCOMMIT NOOVERCOMMIT, IGNORE THRESHOLD The CREATE RESOURCE QUEUE command cannot be run within a transaction in 4.0.
CREATE ROLE	New role attributes to grant permission to create an external table: CREATEEXTERNAL NOCREATEEXTERNAL
CREATE TABLE	WITH OIDS (OIDS=TRUE) clause has been disabled for partitioned tables and append-only column-oriented tables.
CREATE INDEX	Added support for non-unique indexes on append-only tables. Dropped support for Hash indexes. Dropped support for GIN indexes. Dropped support for CONCURRENTLY keyword (building indexes concurrently).

New SQL Commands

Table 3 New SQL Commands in 4.0

SQL Command	Description of Change
DROP FILESPACE	New command for tablespace support.
CREATE FILESPACE	New command for tablespace support.
ALTER FILESPACE	New command for tablespace support.
CREATE TABLESPACE	Now supported in Greenplum Database. New syntax for declaring storage locations: FILESPACE <i>filespace_name</i>
ALTER EXTERNAL TABLE	New command for altering an external table definition (partial support).

Client Utilities

Greenplum Database has merged in client side changes from PostgreSQL 8.3, which has changed the following client utilities:

Table 4 PostgreSQL Client Utility Changes in 4.0

Utility Name	Description
clusterdb	NEW OPTIONS: <code>-v --verbose</code> , <code>-w --no-password</code>
createdb	NEW OPTIONS: <code>-w --no-password</code> DEPRECATED OPTIONS: <code>-q --quiet</code>
createlang	NEW OPTIONS: <code>-w --no-password</code>
createuser	NEW OPTIONS: <code>-w --no-password</code> DEPRECATED OPTIONS: <code>-q --quiet</code>
dropdb	NEW OPTIONS: <code>-w --no-password</code> DEPRECATED OPTIONS: <code>-q --quiet</code>
droplang	NEW OPTIONS: <code>-w --no-password</code>
dropuser	NEW OPTIONS: <code>-w --no-password</code> DEPRECATED OPTIONS: <code>-q --quiet</code>
pg_dumpall	NEW OPTIONS: <code>-f --filespaces</code> , <code>-r --resource-queues</code>
reindexdb	NEW OPTIONS: <code>-w --no-password</code>
vacuumdb	NEW OPTIONS: <code>-w --no-password</code> , <code>-F --freeze</code>

Management Utilities

- [Deprecated Utilities](#)
- [New Utilities](#)
- [Changed Utilities](#)

Deprecated Utilities

Table 5 Deprecated Management Utilities in 4.0

Utility Name	Description
gpchecknet	The <code>gpchecknet</code> utility has been removed. The same functionality is provided in the <code>gpcheckperf</code> utility.
gpcheckos	The <code>gpcheckos</code> utility has been deprecated and replaced by the <code>gpcheck</code> utility.
gpdemo	The <code>gpdemo</code> program has been deprecated (located in <code>\$GPHOME/demo/gpdemo.tar.gz</code> in prior releases). This program initialized a single-node Greenplum system, which is still possible using the sample single-node configuration file is supplied in: <code>\$GPHOME/docs/cli_help/gp_init_singlenode_example</code> .
gpbuildsystem	Recommended process to rebuild a Greenplum system is to reinitialize the array using <code>gpinitssystem</code> and then restore your data from backup files.

Table 5 Deprecated Management Utilities in 4.0

Utility Name	Description
gp_sizecalc	The <code>gp_sizecalc</code> utility was used in prior releases to check the size of a database, schema, or table. Similar functionality is now supplied in the <code>gp_toolkit</code> administrative schema.
gp_skew	The <code>gp_skew</code> utility was used in prior releases to check the data distribution of a table. Similar functionality is now supplied in the <code>gp_toolkit</code> administrative schema.

New Utilities

Table 6 New Management Utilities in 4.0

Utility Name	Description
gp_filespace	New utility for tablespace support.
gp_check	New utility for verifying system and OS settings.
gp_config	New utility for editing <code>postgresql.conf</code> files.
gp_perfmon_install	New utility to install the <code>gp_perfmon</code> database and optionally enable the Performance Monitor data collection agents.
gp_snmpd	New utility for SNMP support.

Changed Utilities

The following utilities have deprecated, new, or changed options:

Table 7 Changed Management Utilities in 4.0

Utility Name	Description of Change
gp_dump	CHANGED OPTIONS: <code>--gp-s=dbid</code> (no more <code>p</code> (rimary), <code>m</code> (irror), <code>i</code> (ndividual) choices)
gp_addmirrors	DEPRECATED OPTIONS: <code>-D</code> (debug) NEW OPTION: <code>-v</code> (verbose) CHANGED OPTIONS: <code>-p port_offset</code> (calculation of mirror ports has changed), <code>-i mirror_config</code> (configuration file format has changed)
gp_checkperf	NEW OPTIONS: <code>-r M</code> (full matrix network test), <code>--duration</code> (duration for network test), <code>--netperf</code> (use <code>netperf</code> binary for network test instead of Greenplum network test) CHANGED OPTIONS: <code>-r n N</code> (The network tests now use a Greenplum-provided network benchmark test. <code>netperf</code> is no longer distributed with Greenplum Database).
gp_crondump	DEPRECATED OPTIONS: <code>-m</code> (dump mirrors only), <code>-p</code> (dump primaries only)
gp_detective	NEW OPTIONS: <code>--start_date</code> , <code>--end_date</code> , <code>--diagnostics</code> , <code>--logs</code> , <code>--cores</code> , <code>--pg_dumpall</code> , <code>--pg_dump_options</code> , <code>--tempdir</code> , <code>--connect</code>
gp_expand	CHANGED OPTIONS: <code>-i expansion_config</code> (configuration file format has changed)
gp_initstandby	DEPRECATED OPTIONS: <code>-i</code> (do not start standby master synchronization process)

Table 7 Changed Management Utilities in 4.0

Utility Name	Description of Change
gpinitssystem	NEW OPTIONS: <code>--locale</code> , <code>--lc-collate</code> , <code>--lc-ctype</code> , <code>--lc-numeric</code> , <code>--lc-monetary</code> , <code>--lc-time</code> , <code>--lc-messages</code> DEPRECATED OPTION: <code>-r</code> (serial mode), <code>-i</code> (do not start standby master synchronization process)
gpload	NEW OPTIONS: <code>--gpfdist-timeout</code>
gpmapproduce	The YAML control file format for defining MapReduce jobs has been modified to allow support for built-in database functions and user-defined C functions.
gprecoverseg	DEPRECATED OPTIONS: <code>-S seg_dbid</code> , <code>-z seg_data_dir:seg_hostname</code> , <code>-f</code> (force restart), <code>-D</code> (debug) NEW OPTION: <code>-v</code> (verbose) CHANGED OPTION: <code>-F</code> (full recovery)
gpstart	DEPRECATED OPTION: <code>--recover</code> NEW OPTION: <code>-t timeout_in_seconds</code>
gpstate	DEPRECATED OPTION: <code>-t</code> (show default utility settings) NEW OPTION: <code>-e</code> (show segments with error conditions) CHANGED OPTION: <code>-v</code> (was 'version'; now 'verbose')
gpstop	DEPRECATED OPTION: <code>--recover</code> NEW OPTION: <code>-t timeout_in_seconds</code> CHANGED OPTIONS: <code>-u</code> (upload configuration files - now reloads <code>postgresql.conf</code> and <code>pg_hba.conf</code> files of master and all segments), <code>-r</code> (a restart now returns all segments to their preferred role)

Server Configuration Parameters

- [Deprecated Parameters](#)
- [New Parameters](#)

Deprecated Parameters

Table 8 Deprecated Server Configuration Parameters in 4.0

Parameter Names	Description
<code>gp_external_grant_privileges</code>	Marked for future deprecation. The <code>CREATE ROLE</code> and <code>ALTER ROLE</code> commands now have the ability to grant the permission to create an external table to an individual role. If this parameter is set to <code>on</code> , then pre-4.0 behavior will remain (any role can create an external table regardless of the individual permissions granted to the role).
<code>gp_fault_action</code>	Removed. With the new mirroring architecture, the system is always in <code>continue</code> (read/write) mode.
<code>gp_hashagg_compress_spill_files</code>	Replaced by <code>gp_workfile_compress_algorithm</code>
<code>resource_scheduler</code>	Removed. Resource queues are now always on.
<code>stats_block_level</code> <code>stats_row_level</code>	Combined and renamed to <code>track_counts</code> .
<code>stats_command_string</code>	Renamed to <code>track_activities</code> .

Table 8 Deprecated Server Configuration Parameters in 4.0

Parameter Names	Description
stats_reset_server_on_start	Removed. Similar functionality provided by the <code>pg_stat_reset()</code> function.
stats_start_collector	Removed. Statistics collection is always started in 4.0.

New Parameters

Table 9 New Server Configuration Parameters in 4.0

Parameter Names	Description
gp_contentid	New read-only parameter that shows the local segment content id.
gp_dbid	New read-only parameter that shows the local segment id.
gp_email_smtp_server gp_email_smtp_userid gp_email_smtp_password gp_email_from gp_email_to	New parameters to configure email alerts.
gp_enable_gpperfmon gpperfmon_port	Greenplum Performance Monitor parameters are now integrated into the default server <code>postgresql.conf</code> files.
gp_max_databases gp_max_filespaces gp_max_tablespace	New parameters to control the number of databases, tablespaces, and tablespaces allowed in a Greenplum Database system.
gp_num_contents_in_cluster	New read-only parameter that shows the number of primary segments in the Greenplum system.
gp_resqueue_priority gp_resqueue_priority_cpucore_per_segment gp_resqueue_priority_sweeper_interval	New parameters to enable and configure query prioritization.
gp_snmp_monitor_address gp_snmp_community gp_snmp_use_inform_or_trap	New parameters to enable SNMP notifications.
gp_vmem_idle_resource_timeout	New parameter to manage system resource utilization for idle database sessions.
gp_workfile_compress_algorithm	Replaces the <code>gp_hashagg_compress_spill_files</code> parameter.
max_work_mem	New parameter to control memory usage at the query processing level.
pljava_classpath pljava_statement_cache_size pljava_release_lingering_savepoints pljava_vmoptions	New parameters for PL/Java implementation.
ssl_ciphers	New parameter to specify the list of SSL ciphers that are allowed to be used on secure connections.
track_activities	Replaces <code>stats_command_string</code> parameter.
track_counts	Replaces <code>stats_block_level</code> and <code>stats_row_level</code> parameters.

Parameters with Changed Defaults

Table 10 Parameters with Changed Default Values in 4.0

Parameter	Old Value	New Value
join_collapse_limit	8	16
from_collapse_limit	8	16
gp_max_csv_line_length	65536	1048576

System Catalogs

Deprecated System Catalogs

Table A.1 Deprecated System Catalogs in 4.0

Catalog Name	Description
gp_configuration	Replaced by <i>gp_segment_configuration</i>
gp_jetpack (schema)	Replaced by the <i>gp_toolkit</i> schema, which is now installed in all databases by default.

New System Catalogs

Table 2 New System Catalogs in 4.0

Catalog Name	Description
gp_fastsequence gp_global_sequence gp_persistent_database_node gp_persistent_filespace_node gp_persistent_relation_node gp_persistent_tablespace_node gp_relation_node	New catalogs to support the new physical file block replication mirroring architecture.
gp_segment_configuration	New Greenplum system configuration table. Replaces <i>gp_configuration</i> .
pg_appendonly_alter_column	New catalog to support adding columns to append-only tables.
pg_filespace pg_filespace_entry	New catalogs for tablespace support.
pg_resourcetype pg_resqueuecapability	New catalogs to support query prioritization and future resource queue enhancements.
pg_stat_last_operation pg_stat_last_shoperation	New catalogs to support meta-data tracking.

New System Views

Table 3 New System Views in 4.0

View Names	Description
pg_stat_operations pg_stat_partition_operations	New views for 'last operation' tracking.
pg_resqueue_attributes	New view for examining the attributes of a resource queue.

New System Schemas

Table 4 New System Schemas in 4.0

Schema Name	Description
gp_toolkit	Replaces <code>gp_jetpack</code> (the optional 3.3.x Greenplum Administrative Schema). <code>gp_toolkit</code> is automatically installed in all databases in 4.0. Also <code>gp_</code> is now a reserved prefix for schema and tablespace names.

Built-in Database Functions

Table 5 New Built-in Functions in 4.0

Function	Description
<code>gp_adjust_priority(session_id, statement_count, priority)</code>	New function to adjust the priority of a running query. This function replaces some customer-specific (and unsupported) utilities that are no longer available in Greenplum Database.
<code>gp_eelog('message_text', boolean)</code>	New function that can be used to trigger an email or SNMP alert in Greenplum Database.

Resolved Issues in Greenplum Database 4.0.1.x

This section lists the 3.3.x customer reported issues that were resolved in Greenplum Database 4.0.1.x. For customers who participated in the 4.0 controlled release beta program, release 4.0.1 contains all of the controlled release fixes reported in 4.0.0.0-4.0.0.8, plus the additional fixes described in this section:

Table A.1 Resolved Issues in 4.0.1.x

Issue Number	Category	Description
10343	Performance Monitor	Performance Monitor Log File Rotation In prior releases of Greenplum Performance Monitor, log files were truncated when the configured size limit was reached. Logging has been changed so that log files are now rolled over. See the <i>Greenplum Performance Monitor Administrator Guide</i> for more information on log rotation and maintenance.
11079	Performance Monitor	Performance Monitor Console Online Help Prior releases of Greenplum Performance Monitor did not have online help available from within the Performance Monitor Console web application. Online help is now provided.
11103	DCA	EMC Greenplum DCA - ConnectEMC Integration For Greenplum Database instances running on the EMC Greenplum Data Computing Appliance (DCA), administrators now have the option of configuring ConnectEMC phone-home support. With ConnectEMC enabled, certain system events (such as a segment host failures) will be captured and automatically sent to EMC Customer Support.
10893	DCA	EMC Greenplum DCA - Performance Monitor Health Monitoring Tab For Greenplum Database instances running on the EMC Greenplum Data Computing Appliance (DCA), the Performance Monitor Console now has an additional Health Monitoring tab. This new tab shows the status of the various hardware components of the DCA.
11232	Management Utilities	New gpcheck Utility to Verify System Hardware and OS Settings This release introduces a new <code>gpcheck</code> utility that performs some pre-installation checks of the hardware and operating systems of hosts that comprise a Greenplum Database system. This utility replaces the <code>gpcheckos</code> utility.
10814	Management Utilities	Reload all postgresql.conf Files with gpstop -u In prior releases, a <code>gpstop -u</code> command would reload the <code>postgresql.conf</code> file of the master instance only. It now reloads all segment <code>postgresql.conf</code> files as well.
2976	Functions (Built-In)	Cannot Insert pg_relation_size Results into a Table The use of the <code>pg_relation_size</code> function has been disabled for <code>CREATE TABLE AS SELECT</code> and <code>INSERT INTO SELECT</code> queries. For example: <pre>CREATE TABLE mytbl_size AS SELECT pg_relation_size(mytable) FROM pg_tables WHERE tablename = 'mytable';</pre> In prior releases, the <code>pg_relation_size</code> function would not return results when used in these types of queries.
4835	Functions (User-Defined)	Use of SECURITY DEFINER in User-Defined Functions In prior releases, Greenplum Database did not properly handle execution of user-defined functions created with <code>SECURITY DEFINER</code> declared. <code>SECURITY DEFINER</code> specifies that the function is to be executed with the privileges of the user that created it, including superuser privileges. This issue has been resolved in this release.
4974	Performance Monitor	Performance Monitor - ERROR: missing data for column "t0_val" In prior releases of Greenplum Performance Monitor, the data collection agents would fail with the following error when loading certain query history data into the <code>iterators_history</code> table of the <code>gpperfmon</code> database: <pre>ERROR: missing data for column "t0_val"</pre> This issue has been resolved in this release.

Table A.1 Resolved Issues in 4.0.1.x

Issue Number	Category	Description
6379	Table Partitioning	Added Partitions Do Not Inherit Primary Keys and Indexes of Parent Table If a user creates a new partitioned table with a primary key or an index, the child table partitions will inherit the primary keys and indexes of the parent table. However, in prior releases, adding a new partition using <code>ALTER TABLE...ADD PARTITION</code> did not inherit primary keys or indexes as expected. This issue has been resolved in this release.
6400	Client Authentication	Format of <code>pg_hba.conf</code> File Changed for LDAP and PAM Authentication This release of Greenplum Database uses the PostgreSQL 8.4 client authentication libraries (not the 8.2 version used in Greenplum releases prior to 3.3.x). This means that the format of the <code>pg_hba.conf</code> client authentication file requires different syntax when declaring LDAP or PAM authentication methods. If the old format is used, Greenplum Database will not start. See the <i>Greenplum Database Administrator Guide</i> for more information on the <code>pg_hba.conf</code> file format.
6567	Data Loading	Unencrypted Password Authentication Support Added to <code>gpload</code> In prior releases, <code>gpload</code> did not support unencrypted password authentication. If the user running <code>gpload</code> had password authentication configured in <code>pg_hba.conf</code> configuration file of Greenplum Database, their load operations would fail. Password authentication for <code>gpload</code> is now supported in this release.
6940	Memory Control	Improved Query-Level Memory Control One challenge Greenplum system administrators face is preventing user workloads from causing out-of-memory errors on the segment hosts, while still allowing users to tune memory usage for a particular query. To address this challenge, the server configuration parameter <code>max_work_mem</code> has been added to control memory usage at the query processing level. As a complimentary parameter to <code>work_mem</code> , it sets the maximum amount of memory that can be used by a single segment worker process during query processing.
6957	Data Loading	Function Expression Support Added to <code>gpload</code> In prior releases, <code>gpload</code> did not support the use of function expressions in the <code>gpload</code> control file (for example, using a function to encrypt a column in the target table). This is now supported in this release.
6958	Concurrency	Resource Utilization of Idle Database Sessions In prior releases, idle sessions would maintain a hold on their allocated resources (such as shared memory). While this is often desirable behavior, it can result in a lower number of concurrent sessions allowed to the system. In 4.0.x, if a database session is idle for longer than 180 seconds, the session will release its resources but remain connected. This idle timeout can be configured using the server configuration parameter: <code>gp_vmem_idle_resource_timeout</code> .
8340	Standby Master	Changes Made by <code>gpactivestandby</code> Not Transactional In prior releases, if the <code>gpactivestandby</code> utility exited or was cancelled mid-operation, it did not properly roll back any changes it made to the Greenplum Database catalogs. When this occurred, any subsequent <code>gpactivestandby</code> or <code>gpstop/start</code> operations would fail when trying to connect to the current master instance. This issue has been resolved in this release.

Table A.1 Resolved Issues in 4.0.1.x

Issue Number	Category	Description
8422	Standby Master	Do Not Allow Activation of an Out-of-Date Standby Master In prior releases, the <code>gpactivatestandby</code> utility did not safe-guard against activating an out-of-date standby master. Activating a standby master that is not in sync with the transactional state of the segments may introduce catalog and data inconsistencies that can render your Greenplum Database instance unusable. In this release, <code>gpactivatestandby</code> will not promote the standby master to active until first checking its synchronization status. If problems are detected, the utility will give the user an error.
8458	Table Partitioning	Failed ALTER TABLE...SPLIT PARTITION on an Open-Ended Range Partition In 3.3.x releases, users could not split a range partition if the partition was created with an open-ended range (did not have both a <code>START</code> and an <code>END</code> criteria defined). Running an <code>ALTER TABLE...SPLIT PARTITION</code> command on such a partition would result in the following error: <code>ERROR: Unexpected internal error...</code> This issue has been resolved in this release.
8471	Data Loading	SQL_ASCII Disabled as a Database Encoding In prior releases, the <code>SQL_ASCII</code> character set was allowed as a server-side encoding. This can potentially cause problems if users load any non-ASCII data into the database, as this encoding does not validate non-ASCII characters, and may allow invalid encoding sequences to be loaded. To prevent problems when converting between server and client-side encodings, Greenplum has disabled the use of <code>SQL_ASCII</code> as a database encoding.
8559	Backup and Restore	gpcrondump Fails when Using Multiple Tables with the -T Option The <code>-T</code> option for <code>gpcrondump</code> allows you to specify one or more table names to exclude from the database dump. Using <code>-T</code> with a comma-separated list of table names caused <code>gpcrondump</code> to fail. In this release, you can specify the <code>-T</code> option multiple times (for example, <code>-T table1_name -T table2_name -T table3_name</code>).
8623	Data Loading	Add Role Permissions for CREATE EXTERNAL TABLE In 3.3.x releases, only superusers had the ability to create an external table by default. The <code>gp_external_grant_privileges</code> server configuration parameter was added to allow all other roles the ability to create external tables, however this ability could not be granted on a per-role basis. In 4.0, the <code>CREATE ROLE</code> and <code>ALTER ROLE</code> commands now include the <code>CREATEEXTABLE/NOCREATEEXTABLE</code> attributes which can be granted to individual roles. To utilize this new feature, the <code>gp_external_grant_privileges</code> parameter must be <code>off</code> in the master <code>postgresql.conf</code> file. If set to <code>on</code> , then pre-4.0 behavior will remain (any role can create an external table regardless of the individual permissions granted to the role).
8658	Data Loading	gpfdist Does Not Read Appended Data in gzip-Compressed Load Files In prior releases, <code>gpfdist</code> did not read data that was appended to an existing <code>gzip</code> compressed file. This would result in appended data not being loaded. This issue has been resolved in this release.
8854	DDL and Utility Statements	ALTER TABLE - ERROR: ZLIB Compress2 Failed (Detail: 'Insufficient Memory'...) In 3.3.x releases, adding or altering a column in a compressed append-only table would potentially fail due to insufficient memory. This issue has been resolved in this release.

Table A.1 Resolved Issues in 4.0.1.x

Issue Number	Category	Description
9429	System Initialization	<p>Improved Error Message for gpinitssystem</p> <p>In prior releases, <code>gpinitssystem</code> would fail with a confusing error message if it could not start a <code>postgres</code> database listener process. For example:</p> <p>The program "postgres" is needed by <code>initdb</code> but was not found in the same directory as "<code>\$GPHOME/bin/initdb</code>"</p> <p>The error message has been changed to provide more information to help users troubleshoot the problem.</p>
9823	Table Partitioning	<p>ALTER TABLE...ADD/DROP COLUMN on Partitioned Table not Propagated to Child Tables</p> <p>In 3.3.0 and later releases, when adding or dropping columns on a partitioned parent table, the columns did not fully propagate to the child table partitions. This could cause some column data to not appear in the child tables when data was loaded through the parent table. This issue has been resolved in this release.</p>
10361	Query Execution	<p>ERROR: unsupported call to mark position of Motion operator</p> <p>In 3.3.x releases, certain queries involving a merge-join operation in their query plan could fail to execute with the following error:</p> <p>ERROR: unsupported call to mark position of Motion operator...</p> <p>This issue has been resolved in this release.</p>
10532	User Defined Functions	<p>PL/Java Support for User Defined Functions</p> <p>The PostgreSQL PL/Java language extension is now supported in Greenplum Database, which allows users to write custom database functions in the Java programming language. Implementation details can be found in <code>\$GPHOME/share/postgresql/pljava/PLJAVA_README</code> of your 4.0 installation.</p>
10597	Data Loading	<p>New gpload Option to Set Allowed Timeout</p> <p>In prior releases, <code>gpload</code> operations could fail if they did not establish a connection to the <code>gpfdist</code> file distribution process before the timeout threshold. On systems with significant network traffic, the default timeout threshold of 10 seconds could be insufficient. <code>gpload</code> now has a <code>--gpfdist-timeout</code> option to allow users to increase the timeout threshold.</p>
10603	Query Execution	<p>Query Planner Performance on Queries Involving Partitioned Tables</p> <p>In prior releases, certain queries involving large partitioned tables could take longer than expected in the query planning phase, especially when most of the partitions did not contain any data. This issue has been resolved in this release.</p>
10753	Table Partitioning	<p>ALTER TABLE...SPLIT PARTITION Improvements</p> <p>In prior 3.3.x releases, the <code>ALTER TABLE...SPLIT PARTITION</code> command had a number of issues. First, users would encounter the following error when reusing an existing partition name in an <code>ALTER TABLE...SPLIT PARTITION</code> command:</p> <p>ERROR: invalid use of boundary specification for DEFAULT partition "<name>"</p> <p>This restriction has been removed in this release. The existing partition name can now be reused in a <code>SPLIT PARTITION</code> command.</p> <p>Second, a split on a partitioned table that had a primary key constraint would sometimes fail with the following error:</p> <p>ERROR: new partition overlaps existing partition...</p> <p>This issue has been resolved in this release.</p>

Table A.1 Resolved Issues in 4.0.1.x

Issue Number	Category	Description
10911	Data Loading	ERROR missing chunk number 0 for toast value... In prior 3.3.x releases, users sometimes encountered the following error when accessing certain large rows from data load error tables: <code>ERROR missing chunk number 0 for toast value...</code> This issue has been resolved in this release.
10927	Query Execution	Certain Date Values Not Recognized in WHERE Clause In prior 3.3.x releases, when a <code>WHERE</code> clause contained an array of date values, certain dates with years prior to 1910 and later than 2089 would not always be recognized by the query planner. This issue has been resolved in this release.
10419	Performance Monitor	Performance Monitor: 'ERROR: data line too long...' When the Greenplum Performance Monitor agents attempted to save query information to the Performance Monitor database, certain long queries would fail to load with the following error: <code>ERROR: data line too long. likely due to invalid csv data</code> This issue has been resolved in this release.
10848, 10853	Management Utilities	Management Utilities Should Place Mirror on Different Subnet than its Primary In prior releases, the Greenplum Database management utilities such as <code>gpinitssystem</code> , <code>gpaddmirrors</code> and <code>gpexpand</code> did not take into consideration the network address when deciding where to place a mirror segment. For maximum redundancy, a primary and its corresponding mirror should be configured to use different hosts as well as different subnets (interconnect networks). This issue has been resolved in this release.
8948	DDL and Utility Statements	reindexdb: 'ERROR: Invalid distributed snapshot received...' In prior 3.3.x releases, users occasionally encountered the following error when running the <code>reindexdb</code> client utility (rebuilding all indexes in a database) when the database contained TOAST tables (oversized attribute storage tables): <code>ERROR: Invalid distributed snapshot received...</code> This issue has been resolved in this release.
11074	Client Tools - gpmapproduce	gpmapproduce: C Function Fatal Error In prior 4.0.x controlled releases, running a Greenplum MapReduce program written in C could result in the following fatal error if the C function encountered compiler warnings: <code>Program terminated with signal 11, Segmentation fault.</code> This issue has been resolved in this release.

Copyright Information

Copyright © 2010 EMC Corporation. All rights reserved.

EMC 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." EMC CORPORATION 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 EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

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