

Changelog

Contents

pgmonitor	2
pgmonitor is your all-in-one tool to easily create an environment to visualize the health and performance of your PostgreSQL cluster.	2
Contents	2
Purpose	3
Supported Platforms	3
Operating Systems	3
PostgreSQL	3
Installation	4
1. Prometheus	4
2. exporter	4
3. Grafana	4
Usage	4
Advanced Usage	4
Roadmap	4
Version History	4
Sponsors	4
Legal Notices	4
Setting up exporters	4
Installation	4
Service Setup (RHEL/CENTOS 7)	5
Database Setup	5
postgresql.conf	5
Monitoring Queries File	5
GRANTS	5
Bloat setup	5
Startup services (RHEL/CENTOS 7)	6
Running multiple postgres exporters (RHEL7)	6
Note for packaging (RHEL/CENTOS 7)	6
Setup (RHEL/CENTOS 6)	6
Running multiple postgres exporters (RHEL6)	7

Setting up Prometheus	7
Installation	7
Setup (RHEL/CENTOS 7)	7
Start services (RHEL/CENTOS 7)	7
Note for packaging (RHEL/CENTOS 7)	8
Setup (RHEL/CENTOS 6)	8

Grafana	8
API Import	8
2.3	8
2.2	9
2.1	9
2.0	10
1.7	11
1.6	11
1.5	11
1.4	11
1.3	11
1.2	12
1.1	12
1.0	12

pgmonitor

pgmonitor is your all-in-one tool to easily create an environment to visualize the health and performance of your PostgreSQL cluster.

pgmonitor combines a suite of tools to facilitate the collection and visualization of important metrics that you need be aware of in your PostgreSQL database and your host environment, including:

- Connection counts: how busy is your system being accessed and if connections are hanging
- Database size: how much disk your cluster is using
- Replication lag: know if your replicas are falling behind in loading data from your primary
- Transaction wraparound: don't let your PostgreSQL database stop working
- Bloat: how much extra space are your tables and indexes using
- System metrics: CPU, Memory, I/O, uptime

pgmonitor is also highly configurable, and advanced users can design their own metrics, visualizations, and add in other features such as alerting.

Running pgmonitor will give you confidence in understanding how well your PostgreSQL cluster is performing, and will provide you the information to make calculated adjustments to your environment.

Contents

- Purpose
- Supported Platforms
- Operating Systems
- PostgreSQL
- Installation



Figure 1: pgmonitor

- [Usage](#)
- [Advanced Usage](#)
- [Roadmap](#)
- [Version History](#)
- [Sponsors](#)
- [Legal Notices](#)

Purpose

pgmonitor is an open-source monitoring solution for PostgreSQL and the systems that it runs on. pgmonitor came from the need to provide a way to easily create a visual environment to monitor all the metrics a database administrator needs to proactively ensure the health of the system.

pgmonitor combines multiple open-source software packages and necessary configuration to create a robust PostgreSQL monitoring environment. These include:

- [Prometheus](#) - an open-source metrics collector that is highly customizable.
- [Grafana](#) - an open-source data visualizer that allows you to generate many different kinds of charts and graphs.
- [Crunchy PostgreSQL Server Exporter](#) - an open-source data export to Prometheus that supports collecting metrics from any PostgreSQL server version 9.1 and above.

Supported Platforms

Operating Systems

- CentOS 6 or greater
- RHEL 6 or greater

PostgreSQL

- 9.4

- 9.5
- 9.6
- 10

Installation

Installation instructions for each package are provided in that packages subfolder. Each step in the installation process is listed here, with a link to additional to further installation instructions for each package.

1. [Prometheus](#)
2. [exporter](#)
3. [Grafana](#)

Usage

Advanced Usage

Roadmap

- Additional monitoring metrics out-of-the-box
- Improved visualizations
- Project build testing

Version History

For the [full history](#) of pgmonitor, please see the [CHANGELOG](#).

Sponsors



Figure 2: Crunchy Data

[Crunchy Data](#) is pleased to sponsors pgmonitor and many other [open-source projects](#) to help promote support the PostgreSQL community and software ecosystem.

Legal Notices

Copyright © 2018 Crunchy Data Solutions, Inc.

CRUNCHY DATA SOLUTIONS, INC. PROVIDES THIS GUIDE “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Crunchy, Crunchy Data Solutions, Inc. and the Crunchy Hippo Logo are trademarks of Crunchy Data Solutions, Inc.

Setting up exporters

Installation

- Install latest `node_exporter` package from Crunchy Repository
- Install latest `postgres_exporter` package from Crunchy Repository
- Install latest `pgmonitor-pg###-extras` package for your major version of PostgreSQL
- Install latest `crunchy-pg_bloat-check` package if you need to monitor database bloat

Service Setup (RHEL/CENTOS 7)

- If necessary, modify `/etc/systemd/system/node_exporter.service.d/crunchy-node-exporter-service-el7.conf`. See notes in file for more details.
- If necessary, modify `/etc/sysconfig/node_exporter`. See notes in file for more details.
- If necessary, modify `/etc/sysconfig/postgres_exporter`. See notes in file for more details.
- Modify `/etc/postgres_exporter/###/crontab###.txt` to run relevant scripts and schedule the bloat check for off-peak hours. Add crontab entries manually to `ccp_monitoring` user (or user relevant for your environment).

Database Setup

postgresql.conf

Install contrib modules to provide additional monitoring capabilities. This requires a restart of the database if you would like these contrib modules installed.

```
shared_preload_libraries = 'pg_stat_statements,auto_explain'
```

`pg_stat_statements` requires running the following statement in the database(s) to be monitored

```
psql -d postgres -c "CREATE EXTENSION pg_stat_statements"
```

Monitoring Queries File

Install functions to all databases you will be monitoring in the cluster (if you don't have `pg_stat_statements` installed, you can ignore the error given). The queries common to all postgres versions are contained in `queries_common.yml`. Major version specific queries are contained in a relevantly named file. Queries for more specialized monitoring are contained in additional files. `postgres_exporter` only takes a single query file as an argument for custom queries, so cat together the queries necessary into a single file.

For example, to use just the common queries for PostgreSQL 9.6 do the following. Note the location of the final queries file is based on the major version installed. The exporter service will look in the relevant version folder in the `ccp_monitoring` directory:

```
cd /etc/postgres_exporter/96
cat queries_common.yml queries_per_db.yml queries_pg92-96.yml > queries.yml
psql -f /etc/postgres_exporter/96/functions_pg92-96.sql
```

As another example, to include queries for PostgreSQL 10 as well as `pg_stat_statements` and `bloat` do the following:

```
cd /etc/postgres_exporter/10
cat queries_common.yml queries_per_db.yml queries_pg10.yml queries_pg_stat_statements.yml
  queries_bloat.yml > queries.yml
psql -f /etc/postgres_exporter/10/functions_pg10.sql
```

For replica servers, the setup is the same except that the `functions_pg###.sql` file does not need to be run since writes cannot be done there and it was already run on the master.

GRANTS

The `ccp_monitoring` role (created by running the “`functions_pg###.sql`” file above) must be allowed to connect to all databases in the cluster. To do this, run the following command to generate the necessary GRANT statements:

```
SELECT 'GRANT CONNECT ON DATABASE "' || datname || "' TO ccp_monitoring;' FROM pg_database WHERE
  datallowconn = true;
```

This should generate one or more statements similar to the following:

```
GRANT CONNECT ON DATABASE "postgres" TO ccp_monitoring;
```

Bloat setup

Run script on the specific database(s) you will be monitoring bloat for in the cluster. See special note in `crontab.txt` concerning a superuser requirement for using this script

```
psql -d postgres -c "CREATE EXTENSION pgstattuple;"
/usr/bin/pg_bloat_check.py -c "host=localhost dbname=postgres user=postgres" --create_stats_table
psql -d postgres -c "GRANT SELECT ON bloat_indexes, bloat_stats, bloat_tables TO ccp_monitoring;"
```

Startup services (RHEL/CENTOS 7)

```
sudo systemctl enable node_exporter
sudo systemctl start node_exporter
sudo systemctl status node_exporter
```

To most easily allow the possibility of multiple postgres exporters and avoid maintaining many similar service files, a systemd template service file is used. The name of the sysconfig EnvironmentFile to be used by the service is passed as the value after the “@” and before “.service” in the service name. The default exporter’s EnvironmentFile is named “postgres_exporter”.

```
sudo systemctl enable crunchy_postgres_exporter@postgres_exporter.service
sudo systemctl start crunchy_postgres_exporter@postgres_exporter
sudo systemctl status crunchy_postgres_exporter@postgres_exporter
```

Running multiple postgres exporters (RHEL7)

Certain metrics are not cluster-wide, so in that case multiple exporters must be run to collect all relevant metrics. The queries_per_db.yml file contains these queries and the secondary exporter(s) can use this file to collect those metrics and avoid duplicating cluster-wide metrics. Note that some other metrics are per database as well (bloat). You can then define multiple targets for that job in Prometheus so that all the metrics are collected together. Note that the “functions_*.sql” file does not need to be run on these additional databases.

```
cd /etc/postgres_exporter/96
cat queries_per_db.yml queries_bloat.yml > queries_mydb.yml
```

You’ll need to create a new sysconfig environment file for the second exporter service. You can just copy the existing ones and modify the relevant lines, mainly being the port, database name, and query file

```
cp /etc/sysconfig/postgres_exporter /etc/sysconfig/postgres_exporter_mydb

OPT="--web.listen-address=0.0.0.0:9188
    --extend.query-path=/etc/postgres_exporter/96/queries_mydb.yml"
DATA_SOURCE_NAME="postgresql://ccp_monitoring@localhost:5432/mydb?sslmode=disable"
```

Since a systemd template is used for the postgres_exporter services, all you need to do is pass the sysconfig file name as part of the new service name.

```
sudo systemctl enable crunchy_postgres_exporter@postgres_exporter_mydb.service
sudo systemctl start crunchy_postgres_exporter@postgres_exporter_mydb
sudo systemctl status crunchy_postgres_exporter@postgres_exporter_mydb
```

Lastly, update the Prometheus auto.d target file to include the new exporter in the same one you already had running for this system

Note for packaging (RHEL/CENTOS 7)

The service override file(s) must be placed in the relevant drop-in folder to override the default service files.

```
/etc/systemd/system/node_exporter.service.d/*.conf
```

After a daemon-reload, systemd should automatically find these files and the crunchy services should work as intended.

Setup (RHEL/CENTOS 6)

The node_exporter and postgres_exporter services on RHEL6 require the “daemonize” package that is part of the EPEL repository. This can be turned on by running:

```
sudo yum install epel-release
```

All setup for the exporters is the same on RHEL6 as it was for 7 with the exception of the base service files. Whereas RHEL7 uses systemd, RHEL6 uses init.d. The RHEL6 packages will create the base service files for you

```
/etc/init.d/crunchy-node-exporter
/etc/init.d/crunchy-postgres-exporter
```

Note that these service files are managed by the package and any changes you make to them could be overwritten by future updates. If you need to customize the service files for RHEL6, it’s recommended making a copy and editing/using those.

The same /etc/sysconfig files that are used in RHEL7 above are also used in RHEL6, so follow guidance above concerning them and the notes that are contained in the files themselves.

Once the files are in place, set the service to start on boot, then manually start it

```
sudo chkconfig crunchy-node-exporter on
sudo service crunchy-node-exporter start
sudo service crunchy-node-exporter status

sudo chkconfig crunchy-postgres-exporter on
sudo service crunchy-postgres-exporter start
sudo service crunchy-postgres-exporter status
```

Running multiple postgres exporters (RHEL6)

If you need to run multiple `postgres_Exporter` services, follow the same instructions as RHEL7 for making a new `queries_XX.yml` file to only gather database specific metrics. Then follow the steps below:

```
- Make a copy of the /etc/sysconfig file with a new name
- Update --web.listen-address in the new sysconfig file to use a new port number
- Update --extend.query-path in the new sysconfig file to point to the new query file generated
- Update the DATA_SOURCE_NAME in the new sysconfig file to point to the name of the database to be monitored
- Make a copy of the /etc/init.d/crunchy-postgres-exporter with a new name
- Update the SYSCONFIG variable in the new init.d file to match the new sysconfig file
- Update the Prometheus auto.d target file to include the new exporter in the same one you already had running for this system
```

Remaining steps to initialize service at boot and start it up should be the same as above for the default service.

Setting up Prometheus

Installation

- Install latest Prometheus package from Crunchy Repository
- Install latest Alertmanager package from Crunchy Repository
- Install latest `pgmonitor-prometheus-extras` package
- Install latest `pgmonitor-alertmanager-extras` package

Setup (RHEL/CENTOS 7)

- If necessary, modify `/etc/systemd/system/prometheus.service.d/crunchy-prometheus-service-el7.conf`. See notes in example file for more details.
- If necessary, modify `/etc/systemd/system/alertmanager.service.d/crunchy-alertmanager-service-el7.conf`. See notes in example file for more details.
- If necessary, modify `/etc/sysconfig/prometheus` to set prometheus startup properties. See notes within the file itself for recommendations.
- If necessary, modify `/etc/sysconfig/alertmanager` to set alertmanager startup properties. See notes within the file itself for recommendations.
- Modify `/etc/prometheus/crunchy-prometheus.yml` to set scrape interval if different from default. Activate alert rules and alertmanager by uncommenting lines when set as needed. Default service expects config file to be named `crunchy-prometheus.yml`.
- Modify `/etc/prometheus/crunchy-alertmanager.yml` and setup alert target (smtp, sms, etc), receiver and route information. Default service expects config file to be named `crunchy-alertmanager.yml`
- Modify `/etc/prometheus/crunchy-alert-rules.yml` and update rules as needed. Default prometheus config expects file to be named `crunchy-alert-rules.yml`.
- Modify `/etc/prometheus/auto.d/*.yml.example` file(s) to point to exporter services to auto-discover. Copy example file to create as many additional targets as needed. Remove `.example` suffix when configuration is final and Prometheus will auto-discover.

Start services (RHEL/CENTOS 7)

```
sudo systemctl enable prometheus
sudo systemctl start prometheus
sudo systemctl status prometheus

sudo systemctl enable alertmanager
```

```
sudo systemctl start alertmanager
sudo systemctl status alertmanager
```

Note for packaging (RHEL/CENTOS 7)

The service override files must be placed in the relevant drop-in folder to override the default service files.

```
/etc/systemd/system/prometheus.service.d/crunchy-prometheus-service.conf
/etc/systemd/system/alertmanager.service.d/crunchy-alertmanager-service.conf
```

After a daemon-reload, systemd should automatically find these files and the crunchy services should work as intended.

Setup (RHEL/CENTOS 6)

TODO

Grafana

The Grafana RPM Package can be downloaded and installed from <https://grafana.com/grafana/download>.

The steps to access the customized dashboards are as follows:

- Connect to Grafana via <https://>ip-address<:3000>
- Login as admin/admin
- Change admin password
- Add Prometheus datasource
- Import all 5 dashboards
- PostgreSQL.json
- PostgreSQLDetails.json
- BloatDetails.json
- CRUD_Details.json
- TableSize_Details.json

API Import

It is possible to import these graphs through the “import” HTTP API using the following curl command to add some required wrapper information to each json blob.

```
curl --user username:password -S "http://localhost:3000/api/dashboards/import" -X POST -H
'Content-Type: application/json;charset=UTF-8' --data-binary "{ \"dashboard\" : $(cat
PostgreSQL.json) , \"overwrite\":true, \"inputs\":[ { \"name\": \"DS_PROMETHEUS\",
\"type\": \"datasource\", \"pluginId\": \"prometheus\", \"value\": \"PROMETHEUS\" } ] }"
```

You will likely have to edit the following parts of the above command:

- Username
- Password
- URL (and port if different from default)
- The file to import goes in the `$cat()` parentheses. In the above example this is “PostgreSQL.json”.
- In the above example, “value: PROMETHEUS” is your grafana datasource name for this dashboard. Replace “PROMETHEUS” with the proper datasource name.

2.3

- Fixed bug in Prometheus alerts that was causing some of them to be stuck in PENDING mode indefinitely and never firing. This unfortunately removes the current alert value from the Grafana Prometheus Alerts dashboard.
- If you can't simply overwrite your current alerts configuration file with the one provided, remove the following option from every alert: `alert_value: '{{ $value }}'`
- Added feature to monitor pgbackrest backups (<https://pgbackrest.org>)

- Separate metrics exist to monitor for the latest full, incremental and/or differential backups. Note that a full will always count as both an incremental and diff and a diff will always count as an incremental.
- Another metric can monitor the runtime of the latest backup of each type per stanza.
- Run the `setup_pg###.sql` file again in the database that your exporter(s) connect to to install the new, required function: `monitor.pgbackrest_info()`. It has security definer so execution privileges can be granted as needed, but it must be owned by a superuser.
- New metrics are located in the `exporter/postgres/queries_backrest.yml` file. Add the one(s) you want to the main queries file being used by your currently running exporter(s) and restart.
- Example alert rules for different backup scenarios have been added to the `prometheus/crunchy-alert-rules.yml` file. They are commented out to avoid false alarms until valid backup settings for your environment are in place.
- Added new feature to monitor for failing `archive_command` calls.
 - New metric `ccp_archive_command_status` is located in `exporter/postgres/queries_common.yml`. Add this to the main queries file being used by your currently running exporter(s) and restart.
 - A new alert rule `PGArchiveCommandStatus` has been added to the `prometheus/crunchy-alert-rules.yml` file.
- Added new feature to monitor for sequence exhaustion
 - Requires installation of a new function located in the `setup_pg###.yml` file for your relevant major version of PostgreSQL. Must be installed by a superuser.
 - New metric `ccp_sequence_exhaustion` located in `exporter/postgres/queries_common.yml`. Add this to the main queries file being used by your currently running exporter(s) and restart. A new alert rule `PGSequenceExhaustion` has been added to the `prometheus/crunchy-alert-rules.yml` file.
- The `setup_pg###.sql` file now has logic to avoid throwing errors when the `ccp_monitoring` role already exists. Also always attempts to drop the functions it manages first to account for when the function signature changes in ways that `OR REPLACE` doesn't handle. All this allows easier re-running of the script when new features are added or used in automation systems. Thanks to Jason O'Donnell for role logic.

2.2

- Fixed broken `ccp_wal_activity` check for PostgreSQL 9.4 & 9.5. Updated check is located in the relevant `exporter/postgres/queries_pg###.yml` file
- Fixed broken service files for `postgres_exporter` on RHEL6 systems.
- Removed explicit `public` schema in `ccp_bloat_check` query so that it will properly use the `search_path` in case bloat tables were installed in another schema
- Removed query files for PostgreSQL versions no longer supported by `pgmonitor` (9.2 & 9.3)

2.1

- **IMPORTANT UPGRADE NOTE FOR CRUNCHY PACKAGE USERS:** In version 2.0, the Crunchy provided extras for `node_exporter` were split out from the `pgmonitor-pg###-extras` package. A dependency was kept between these packages to make upgrading easier. For 2.1, the dependency between these packages has been removed. When upgrading from 1.7 or earlier, if you have `node_exporter` and `postgres_exporter` running on the same systems, ensure that you install the separate `pgmonitor-node_exporters_extras` package after the update. See the README for the full package name(s).
- Minimum required versions of software used in `pgmonitor` have been updated to:
 - Prometheus 2.5.0
 - Prometheus Alertmanager 0.15.3
 - `postgres_exporter` 0.4.7 (enables full PG11 support)
 - Grafana 5.3.4.
- Fixed Grafana data source to use the `proxy` mode instead of `direct` with default install. Should fix connection issues encountered during default setup between Grafana & Prometheus.
- Renamed `functions_pg###.sql` file to `setup_pg###.sql` to better clarify what it's for (and because it's not just functions).
- Added `ccp_wal_activity` metric to help monitor WAL generation rate.
- For all PG versions, provides total current size of WAL directory. For PG10+, it also provides the size of WAL generated in the last 5 minutes

- Note that for PG96 and lower, a new security definer function must be added (can just run `setup_pg###.sql` again).
- New metric definition is located in the `queries_pg###.yml` file.
- No default rules have been added since this is very use-case dependent.
- Improved accuracy of “Idle In Transaction” monitoring times in PostgreSQL. Base the time measured on the state change of the session vs the total transaction runtime.
- Split `setup_pg92-96.sql` and `queries_pg92-96.sql` into individual files per major version.
- Added commented out example prometheus alert rule for checking if a postgres database has changed from replica to primary or vice versa. Must be set on a per system basis since you have to tell it if a system is supposed to be a primary or replica.
- Removed `pg_stat_statements` prometheus metric and security definer function from setup script. We highly recommend having `pg_stat_statements` installed on a database, and we still include its installation in the documentation, but we currently don’t have any useful metric recommendations from it to collect in prometheus.
- Added some default filters for the bloat check cronjob to avoid unnecessary waste in the prometheus storage of bloat metrics.
- Update documentation.

2.0

- Recommended version of Prometheus is now 2.3.2. Recommended version of Alertmanager is 0.15.1. Recommended version of `postgres_exporter` is 0.4.6.
- Upgrade required version of `node_exporter` to minimum of 0.16.0. Note that many of the metrics that are used in Grafana and Prometheus alerting have had their names changed.
- This version adds these new metrics into Grafana graphs without removing the old metric names on most, but not all, graphs. This allows trending history to be kept. Note that line colors will change in graphs and legend names will be duplicated until the old metric data is expired out.
- Prometheus alerts have been set to use the new metric names since the alerts are based only on recent values.
- **IMPORTANT:** A future `pgmonitor` update will remove these old metric names from Grafana graphs, so please ensure these changes are accounted for in your architecture.
- See full release notes for 0.16.0 - https://github.com/prometheus/node_exporter/releases/tag/v0.16.0
- The `postgres_exporter` service no longer uses a symlink in `/etc/sysconfig` to point to a default “`postgres_exporter`” file. This was causing issues with several upgrade scenarios. New installation instructions now have the service pointing directly to the relevant `sysconfig` file for the major PostgreSQL version.
- **IMPORTANT:** If you are using the default `postgres_exporter` service, you will need to update your service name so it uses the proper `sysconfig` file. See the README file for the new default service name in the “Enable Services” section and run the “enable” command found there. You should then also disable/remove the old service so it doesn’t try to start again in the future.
- The additional Crunchy provided configurations for `node_exporter` have been split out from the `pgmonitor-pg###-extras` package to the `pgmonitor-node_exporter-extras` package. This was done to allow multiple versions of the `pg###-extras` package to be installed with different major versions of Postgres. There is still currently a dependency that the node extras packages must be installed with the `pg###-extras` so that upgrading doesn’t break existing systems. This dependency will be revisited in the future.
- Removed the requirement for a shell script to monitor if the database is up and its status as either a primary or replica. Up status is now using the native “`pg_up`” metric from `postgres_exporter` and a new metric query was written for checking the recovery status of a system (`ccp_is_in_recovery`).
- The PostgreSQL.json overview dashboard that used this metric has been redesigned. Unfortunately it can no longer be colored RED for down systems, only go colorless and say “DOWN”. This is a known limitation of handling null metric values in Grafana and part of a larger fix coming in future versions - <https://github.com/grafana/grafana/issues/11418>
- Upgrade required version of Grafana to minimum of 5.2.1.
- All provided dashboards require this minimum version to work.
- If you notice that links between the dashboards are broken after the upgrade, clear your browser’s cache. The 301 redirects used between dashboards can get cached and they have changed in the new major version.
- See extensive release notes for major version changes in Grafana - <https://community.grafana.com/t/release-notes-v-5-1-x>
- Change Grafana datasource and dashboard installation to use provisioning vs manual setup via the web interface. Note this means that future updates to the provided datasources and dashboards must be done through config files as well. Or they can be saved as a new dashboard for more extensive customization.
- Change recommended configuration for Grafana to use PostgreSQL as database backend. Updated installation documentation.
- Added Prometheus Alerts Dashboard. Shows both active alerts and 1 week history in table format.
- Removed Gauges from PostgreSQLDetails Dashboard. “Current” value was not being shown properly and gauges were misleading in their values depending on the time range chosen. For a quick glance to see if there are any problems, be sure to set your alert thresholds properly and use the new Prometheus Alerts Dashboard.
- Added `max_query_time` metric to track long running queries in general. Also added an alert for that metric to crunchy prometheus alerts.

- Added “IO Time Per Device in Seconds” graph to Filesystems dashboard.
- Fixed Memory and Swap Graphs on PostgreSQLDetails dashboard to more accurately show used resources. History for these graphs before this upgrade is not being shown since it is no longer graphing the same data.
- Crontabs are no longer PostgreSQL major version dependent at this time. Consolidated down to a single crontab file for all versions.
- Removed unnecessary functions from functions_pg10.sql. All queries in queries_pg10.yml currently only require the pg_monitor system role to be granted and have been updated with this assumption.
- Changed default cron runtime of pg_bloat_check to once a week on early morning weekend.
- Change PostgreSQL overview dashboard to use background colors instead of gauges for better visibility.
- Fixed permission issues with /etc/postgres_exporter folder to allow ccp_monitoring system user better control.

1.7

- Fixed duplicate and incorrect replication byte lag queries. The one contained in queries_common.yml should not have been there. It should be in queries_pg92-96.yml, but there was also one already there. However, the one already in pg92-96 was incorrect since prior to PG10, it requires superuser/security definer to fully access replication statistics. Corrected the version specific file to have the correct query. Made the query in the pg10 file consistent. Ensure you update your generated queries.yml file with the new queries.
- Fixed the PostgreSQLDetails.json dashboard to use the correct replication byte lag metric (referencing above fix). The easiest way to fix this is to delete this dashboard and re-import it. Otherwise, if you’ve made customizations you don’t want to lose, you can grab the correct metric query from the updated dashboard gauge and edit your existing dashboard to use it.
- The combination of the above two fixes corrects the pgmonitor setup being able to properly handle there being multiple replicas from a single primary. Previously this would cause postgres_exporter to throw duplicate metric errors.
- Fixed the query in queries_bloat.yml to be able to properly handle if there was a bloat amount larger than max int4 bytes. Ensure you update your generated queries.yml file with the new query.

1.6

- Fixed formatting bug in crunchy-prometheus.yml. Thanks to Doug Hunley for reporting the issue.

1.5

- Add support for disabling built in queries in postgres_exporter 0.4.5. Also explicitly ignore these metrics via a prometheus filter so they’re not ingested even if new option isn’t used. This means that v1.5 of pgmonitor now requires 0.4.5 of postgres_exporter by default.
- Improved exporter down alert to avoid unnecessary alerts for brief outages that resolve themselves quickly.
- Added new FilesystemDetails dashboard for grafana that is linked to from the Filesystem graph on PostgreSQLDetails.
- Top level PostgreSQL grafana dashboard now identifies whether a system is read/write or readonly to better distinguish primary/replica systems.
- Added instructions for non-packaged installation using pgmonitor configuration files.
- Revised and better formatted README documentation

1.4

- Fixed filesystem graphs in PostgreSQLDetails dashboard
- Cosmetic changes to PostgreSQLDetails dashboard
- Added instructions for importing dashboards via Grafana API

1.3

- Fixed error in PG10 queries file.
- Fixed disk usage alert for prometheus to work better when there are many jobs with similar mountpoints. Also fixed syntax error in warning alert.
- Moved connection stats query from common to version specific queries due to PG10 differences. Clarified naming of files for which versions they work for.
- Added dropdown for the Job to the lower level drill down dashboards in Grafana. Allows selecting of a specific system from the dashboard itself without having to click through on a higher level.
- Removed pg_stat_statements graph from PostgreSQLDetails dashboard. Needs refinement to make it more useful.

1.2

- Change service and sysconfig files to use single OPT environment variable instead of one variable per cmd option
- Fix error in PG10 monitoring functions file
- Initial version of Prometheus 2.0 job deletion script. Requires API call not available yet in 2.0.0 for full functionality

1.1

- Implement rpmnew/rpmsave feature instead of using .example files to prevent package overwriting user changes to configs

1.0

- Initial stable release