



Release Notes - TigerGraph Server 3.2

Release Date: September 21, 2021

TigerGraph

Release Notes for Previous Versions

- [Release Notes - TigerGraph 3.1](#)
- [Release Notes - TigerGraph 3.0](#)
- [Release Notes - TigerGraph 2.6](#)
- [Release Notes - TigerGraph 2.5](#)
- [Release Notes - TigerGraph 2.4](#)
- [Release Notes - TigerGraph 2.3](#)
- [Release Notes - TigerGraph 2.2](#)
- For v2.1 and older, contact TigerGraph Support

For the running log of bug fixes, see [Change Log](#).

Compatibility with TigerGraph 3.1

The following changes were made to the built-in roles in TigerGraph's Role-based Access Control

- The built-in role `queryreader` can no longer run queries that include updates to the database.
 - To emulate the old `queryreader` role, create a role with all `queryreader` privileges, and also grant the `WRITE_DATA` privilege to the new role.
- The built-in role `admin` can no longer create users
 - To emulate the old `admin` role, create a global role with all `admin` privileges, and also grant the `WRITE_USER` privilege to the new role.
- To learn more about role management and the privileges of built-in roles, see:
 - [Role Management](#)
 - [List of Privileges](#)

Compatibility with TigerGraph 2

Major revisions (e.g., from TigerGraph 2 to TigerGraph 3) are the opportunity to deliver significant improvements. While we make every effort to maintain backward compatibility, in selected cases APIs have changed or deprecated features have been dropped, in order to advance the overall product.

Data migration: A tool is available to migrate the data in TigerGraph 2.6 to TigerGraph 3.0. Please contact TigerGraph Support for assistance.

Query and API compatibility:

- Some gadmin syntax has changed. Notably, `gadmin set config` is now `gadmin config set`. Please see [Managing with gadmin](#).
- Some features which were previously deprecated have been dropped. Please see [V3.0 Removal of Previously Deprecated Features](#) for a detailed list.

New Features

User-defined Roles

Privileges are introduced as the atomic unit for managing database user access. Database administrators can now define their own roles with customizable collections of privileges.

For details on this feature, see [User Access Management](#).

Cross-region replication

TigerGraph 3.2 provides complete native support for all data and metadata cross-region replication including automated schema changes, user and query management.

For details on this feature, see [Disaster Recovery](#).

Containerization of services

TigerGraph 3.2 allows users to deploy TigerGraph single servers and clusters on Kubernetes.

Running applications in containers on Kubernetes provides rapid spin-up and repeatability across environments.

For details on this feature, see [Kubernetes](#).

Cluster resizing

TigerGraph 3.2 provides built-in cluster management features that allow users to expand, shrink, and repartition their TigerGraph clusters.

For details on this feature, see [Cluster resizing](#).

Workload Management

TigerGraph 3.2 provides the `GSQL-REPLICA` and `GSQL-THREAD-LIMIT` header to specify the replica for a query to run on and a thread limit that a query is allowed to use.

For details on this feature, see [Run a query](#).

File output policy

Starting in TigerGraph 3.2, GSQL has a [file output policy](#) that contains a whitelist and a blacklist. GSQL queries will only write to the whitelist and are forbidden from writing to the blacklist.

Log viewing with Elasticsearch, Kibana, and Filebeat

Starting with TigerGraph 3.2, TigerGraph provides a gadmin command that can generate Filebeat configuration files for a TigerGraph cluster. Read our [step-by-step guide](#) to set up Elasticsearch, Kibana, and Filebeat to view TigerGraph logs.

General Improvements

Expanded library of built-in GSQL functions

TigerGraph 3.2 added over 30 built-in functions to the GSQL query language.

Relaxed variable declaration rules

Starting in TigerGraph 3.2, base type variables and accumulators can be declared anywhere in a query and are block-scoped.

Parameter passing with JSON

Starting in TigerGraph 3.2, users can pass in parameters to a GSQL query with a JSON object.

Subquery return anonymous tuples

Starting in TigerGraph 3.2, subqueries in GSQL can return an anonymous tuple or a collection of anonymous tuples.

Query installation improvements

TigerGraph 3.2 implemented the following improvements for query installation:

- Updating a subquery will no longer require reinstalling all dependent queries
- Schema change will no longer trigger reinstalling all queries of the graph
- Concurrent query installation between graphs is now supported
- When installing queries on a cluster, TigerGraph will now utilize the computing power of multiple nodes to compile the queries, greatly improving installation performance

WCAG-compliant accessibility features

The user interfaces of GraphStudio and Admin Portal - TigerGraph's GUI are improved to meet WCAG accessibility criteria. More users across a wider range of physical abilities will now be able to work effectively with GraphStudio and the visual Admin Portal.

Edit mode for graph exploration in GraphStudio

Starting with TigerGraph 3.2, users need to enter Edit Mode in the Graph Exploration Panel in order to write to the graph.

Enhanced cluster service status for AdminPortal

Cluster service status is moved from the footer of the Admin Portal page to the bottom of the navigation menu.

Improved GSQL auto-complete and syntax highlighting in GraphStudio

We made improvements to GraphStudio's auto-complete and syntax highlighting features so users have a better experience writing and editing queries in GraphStudio.

Substring search within graph exploration result in GraphStudio

Starting with TigerGraph 3.2, users can search substring to locate the vertices in addition to using exact match.

Function overloading

Starting in TigerGraph 3.2, function overloading is now available. Query UDFs with the same name but different signatures can be defined in the UDF library.

Known Issues

Cluster expansion

If you have ever deleted a vertex or edge attribute before conducting an expansion, cluster expansion may lead to issues reading attributes of that vertex or edge type.

For details, see [Cluster Expansion](#).

Engine

- Concurrent read and write operations on the same vertex and its connected edges during data loading could cause Graph Processing Engine (GPE) dysfunction in rare cases.

Applications

- **GraphStudio**
 - [Known Issues for GraphStudio](#)
 - The No-Code Data Migration feature is in Alpha release. Your feedback would be appreciated.
 - The No-Code Visual Query Builder is in Beta release. Your feedback would be appreciated.
- **AdminPortal**
 - [Known Issues for AdminPortal](#)

GSQL

- **Multiple (Conjunctive) Path Patterns:**
 - There are no known functional problems, but the performance has not been optimized. Your feedback would be appreciated.
- **DML type check error in V2 Syntax:**
 - GSQL will report a wrong type check error for Query block with multiple POST-ACCUM clauses and Delete/Update attribute operation.
- **Turn on GSQL HA manually when upgrading from 3.0.x**
 - Users who are upgrading from 3.0.X need to manually start GSQL HA service. Please reach out to support for help with the process documented in: <https://tigergraph.freshdesk.com/a/solutions/articles/5000865072>
- **Stale data visible after Deletes using index**
 - Queries that use secondary index may still see the vertices being deleted until after the snapshots are fully rebuilt.

Resources

Support

Community

Social

Linkedin

Facebook

Twitter

Legal

Privacy Policy

Terms of Use

Sitemap

Developer Site

Test Drive