WHITE PAPER

Pivotal Greenplum 5: The Next-Generation Data Platform

Open Source, Multi-Cloud, and Built for Advanced Analytics

By Keaton Adams, Dan Baskette, and Cesar Rojas

Pivotal

Table of Contents

About this White Paper

Pivotal recently announced the availability of Pivotal Greenplum 5, the world's first open source, multi-cloud data platform built for advanced analytics. This white paper focuses on the core features of Greenplum 5 and the ecosystem that has developed around it over the years.

Summary

Pivotal Greenplum is infrastructure agnostic, which means it is a fully portable analytic database software solution that can be deployed in a multi-cloud environment (public and private clouds), as well as a variety of on-premises configurations. At the heart of the massively parallel processing (MPP) SQL design is a next-generation query optimizer, known as GPORCA. Developed specifically to address advanced analytics in a multi-structured data environment, GPORCA is capable of handling a wide range of complex queries with concurrent mixed workloads, providing massively improved query performance over more traditional RDBMS query optimizers normally used in legacy MPP databases.

Pivotal Greenplum 5: Next-Generation Data Platform

As a major new release, Pivotal Greenplum 5 brings with it a number of product enhancements and additional features that most customers will find highly useful as they manage data and apply data science, analytics, reporting, and data insight methods to the information stored in the database. Greenplum is architected as a solution to manage very complex queries and to provide powerful analytics enhancements to ANSI-compliant SQL. By automatically partitioning data and running parallel queries, it allows a cluster of servers to operate as a single data supercomputer, performing tens or hundreds of times faster than either a traditional database or alternative platforms. It supports ANSI SQL with a number of analytic extensions and has several built-in languages and additional features provided through packaged extensions. Greenplum has the capabilities to manage data volumes ranging from hundreds of gigabytes to petabytes of data.

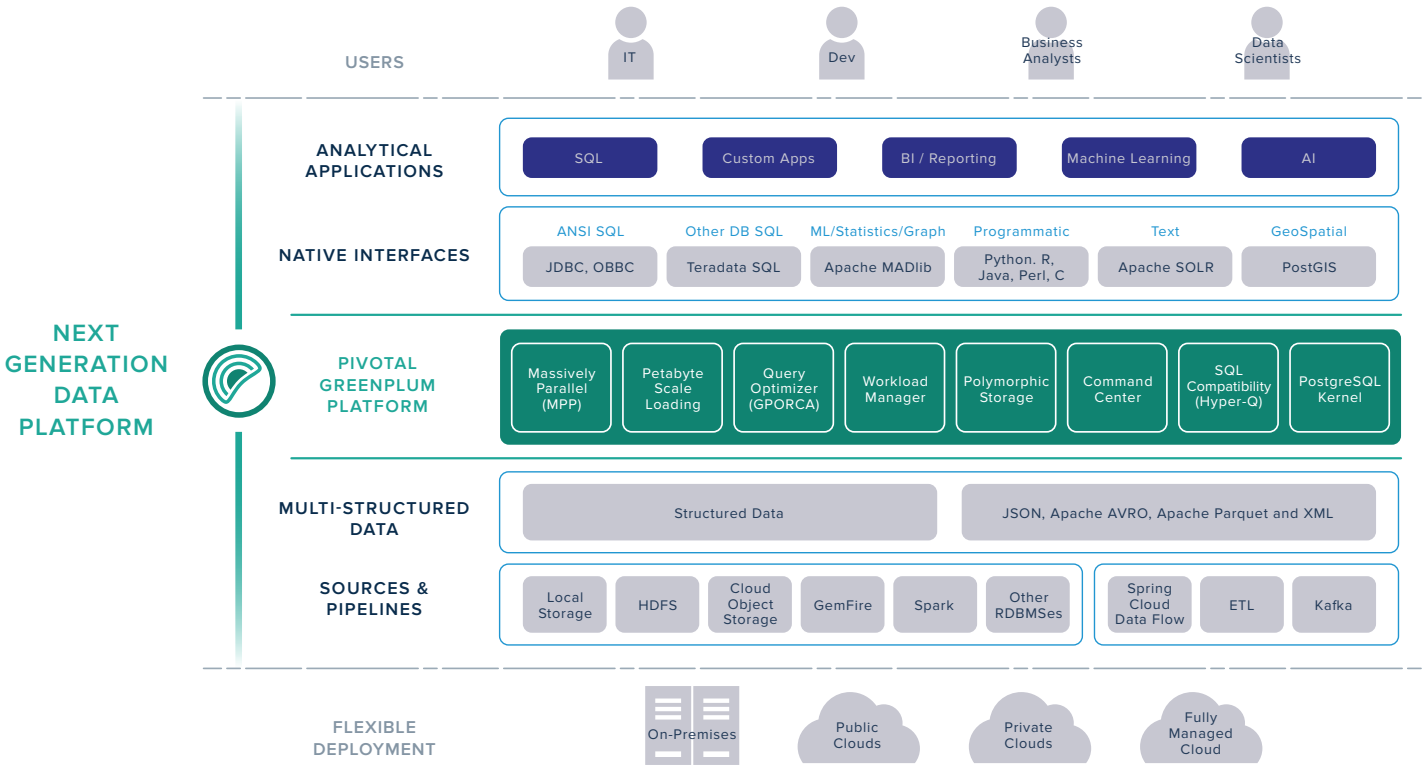

Image 1: Pivotal Greenplum 5: Next-Generation Data Platform.

Open Source Innovation as an Alternative to Proprietary Analytical Environments

To support continued growth of Greenplum, Pivotal made the decision in 2015 to open source its product as the Greenplum Database. The most positive outcome has been a rapid surge in a community that contributes to the development of the core components, and which has a vested interest in its long-term success in the marketplace. It has also led to a rapid and reliable rate of innovation, with an updated release of Greenplum every month since the beginning of 2017. What comes with the new Greenplum 5 release is a fully featured, dynamic, and innovative analytical data platform, with a healthy and vibrant product roadmap that is sure to please customers, both in the near term and for many years to come.

An Infrastructure-Agnostic, Multi-Cloud Data Platform

Pivotal Greenplum 5 is the first release whose core is 100% based on the open source Greenplum Database project and its related source code. This means that, for the first time, both the version that can be downloaded and compiled from the greenplum.org website and the version that is packaged and distributed through the Pivotal Network will have the same core (with a few minor exceptions). The effort is a culmination of two years of work to integrate with PostgreSQL 8.3.23, in order to expand and embrace the ecosystem and community around Greenplum. The major refactoring of the code base to better align with the PostgreSQL community provides a path to more easily incorporate additional PostgreSQL features from their latest releases (PostgreSQL 9.X and 10 in the future).

As a software-only data platform, providing high performance without the dependence on dedicated hardware acceleration, Greenplum runs anywhere a customer needs it to. This “infrastructure-agnostic” approach enables the same type of analytical database to be deployed on premises, or in a multi-cloud environment (private or public cloud).

Most of the advantages of an infrastructure-agnostic approach are equally powerful in the commercial Pivotal Greenplum or the open source Greenplum Database. When deploying analytics on Greenplum Database, users have some extra benefits:

- **Greenplum Database** eliminates platform/vendor lock-in. Services and support for Greenplum can be obtained from different vendors.
- **Greenplum Database** is developed using a community/customer-focused development model. There are more open and accessible methods of influencing the overall product direction, which in turn accelerates product innovation.

Customers have the ability to deploy Pivotal Greenplum on an initial set of servers in the cluster, with the ability to expand the number of servers in the configuration as data storage and user demand requirements increase, without the need to unload and reload the data. This flexibility increases several fold as more customers are moving their production datasets to the public cloud. Pivotal Greenplum is now available on Amazon Web Services, Microsoft Azure, and Google Cloud Platform in both a bring-your-own-license (BYOL) and a per-hour, on-demand configuration. Simply select the number of data nodes to provision in the cluster (up to a predefined maximum size), choose a series of deployment options, and the cluster will be up and running for customer use in no time.

Image 2: Pivotal Greenplum 5: Infrastructure-Agnostic Software Architecture.

Integrated Analytics: New and Improved Analytical Interfaces

It has always been possible to do advanced analytics in Pivotal Greenplum, providing a way to push application logic down to where the data resides, executing an analytical function, or building a data model in massively parallel fashion. Greenplum 5 now supports the most comprehensive, advanced analytical packages and extensions for data discovery and data science work.

Greenplum 5 provides easy-to-use installers for the most popular add-on libraries for Python and R.

- The Python data science libraries and package supported in Greenplum 5 are: Tensorflow, NumPy, SciPy, scikit-learn, Pandas, NLTK, Pattern-en, Statsmodels, gensim, pyldavis, lifelines, spaCy, XGBoost, BeautifulSoup, lxml, Keras, and PyMC3 (Tensorflow and Keras require RHEL 7 as host OS).
- The R data science libraries and package supported in Greenplum 5 are: BH, DBI, MASS, MCMCpack, Matrix, R2jags, R6, RColorBrewer, ROCR, RPostgreSQL, Rcpp, RcppEigen, RobustRankAggreg, SparseM, abind, adabag, arm, assertthat, bitops, caTools, car, caret, coda, colorspace, compHclust, curl, data.table, dichromat, digest, dplyr, e1071, flashClust, forecast, foreign, gdata, ggplot2, glmnet, gplots, gtable, gtools, hms, hybridHclust, igraph, labeling, lattice, lazyeval, lme4, lmtest, magrittr, minqa, MTS, munsell, neuralnet, nloptr, nnet, pbkrtest, plyr, quantreg, randomForest, readr, reshape2, rjags, rpart, sandwich, scales, stringi, stringr, survival, tibble, tseries, and zoo.

Greenplum 5 also supports the latest Apache MADlib release, which provides machine learning and graph analytics in SQL, as well as GPText, which is an in-database implementation of Apache Solr for indexing and searching, including custom tokenizers for international text and social media text, and a Universal Query Processor that accepts queries with mixed syntax from supported Solr query processors. The PostGIS package is a spatial database extension for PostgreSQL that allows Geographic Information Systems (GIS) objects to be stored in the database. The Pivotal Greenplum PostGIS extension includes support for GiST-based R-Tree spatial indexes and functions for analysis and processing of GIS objects.

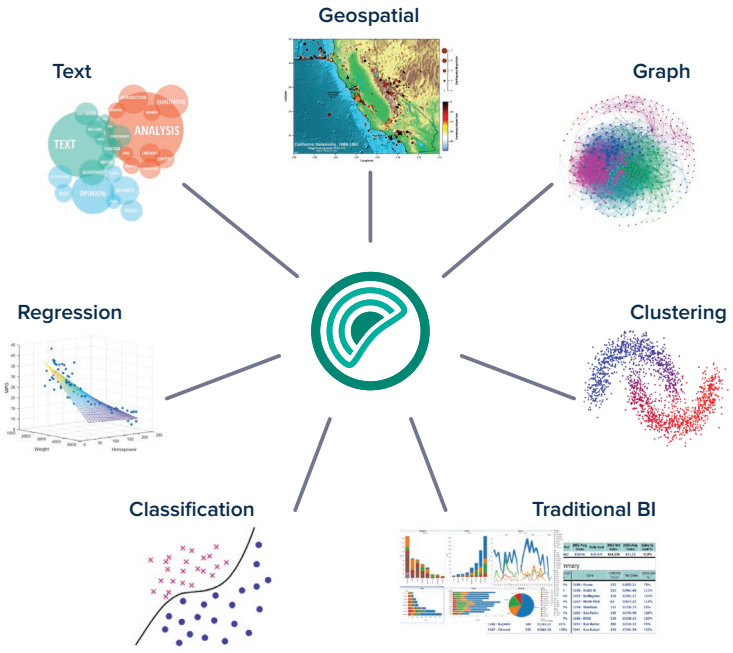

Image 3: Pivotal Greenplum 5: Integrated Analytics.

GPORCA: Pivotal Query Optimizer Updates

Until recently, Greenplum used what is referred to as the legacy query optimizer (LQO). This is a derivative of the original PostgreSQL planner that was adapted to the Greenplum code base. The PostgreSQL planner was originally built for single-node PostgreSQL optimized for OLTP queries, as opposed to the long-running queries seen in analytic data platforms. Although features like join ordering were carefully thought out, the architecture and design choices make maintenance and adding new features increasingly difficult.[1]

At the end of 2010, Greenplum began an internal effort to produce a modern query optimizer, which made its first appearance in Greenplum version 4.3.5. as GPORCA. What makes GPORCA particularly useful is its ability to generate more efficient code for some of the complex situations that commonly arise in analytic data warehouses. Previously, the legacy query optimizer was set as the default, but as of Greenplum 5, GPORCA is the default query optimizer. It is capable of handling a wide range of complex queries with concurrent mixed workloads with improved query performance.[2]

It enables large teams to work in parallel on multiple analytics use cases with advanced analytics and diverse workloads, delivering high analytical query performance on large data volumes. The power of GPORCA is in its ability to compute a large number of possible query plans for a submitted SQL statement in parallel fashion. To generate the fastest plan, GPORCA considers thousands of alternative query execution plans and makes a cost-based decision. It also improves optimization time by eliminating unnecessary profiling steps. Compared to the legacy query optimizer, GPORCA considers many more plan alternatives, which allows it to optimize a wider range of queries.[3]

Modern data analytics and business intelligence (BI) often produce SQL with correlated subqueries, where the inner subquery requires knowledge of the outer query. GPORCA generates a de-correlated plan, which is done only once. The intermediate results are then joined with the master table to generate a resultset that meets the user's criteria. With these and other optimizations, SQL optimized by GPORCA can achieve increases in speed of a factor of 10 or more. There are other queries—albeit a small number—for which GPORCA has not yet produced an improvement in performance. As more capabilities are added to GPORCA over time, it will be the rare case for which the LQO provides better performance.[4]

Greenplum Resource Groups and Workload Manager

Managing concurrency and user resource distribution is a key feature of Greenplum. Workload Manager has been enhanced in this release, and a new way of managing database queries, called Resource Groups, provides more control by a database administrator over user activity, especially in the areas of CPU and memory management. Once resource groups are defined, it places everything under management, including super users. Any statements executed by a superuser are routed to a default resource group, which can be tuned with additional or fewer system resources, based on demand. If there is no open slot in the queue for a query to occupy, or if there is not enough available memory to run the query, it will automatically be queued until resources become available for successful execution.

1. Marshall Presser, Data Warehousing with Greenplum: Open Source Massively Parallel Data Analytics (Sebastopol, CA: O'Reilly Media, 2017), 71-76.
2. Ibid., 71-76.
3. Ibid., 71-76.
4. Ibid., 71-76.

The features of Workload Manager are enhanced in Greenplum 5, allowing more customizable rule creation and better ways to monitor query activity and the resources they consume. It monitors and detects memory, CPU, and disk I/O skew while a query is actively running. Workload Manager then logs when a query exceeds the defined threshold for a resource, and if necessary as defined in the rule, may terminate a query that exceeds the resource limit definition. Rules can be written to be time-of-day and day-of-week sensitive to allow varying types of workloads (e.g., ETL vs. report processing) to run uninterrupted.

Structured Query Language Performance Enhancements

Pivotal Greenplum 5 introduces a number of improvements to SQL query processing. A popular SQL structure—the correlated subquery—is a subquery (i.e., a query nested inside another query) that uses values from the outer query. This is probably the most significant enhancement in GPORCA, because of the heavy use of subqueries by the major BI/reporting tools in the industry. In large datasets where the subquery has to be evaluated once for each row processed by the outer query, execution may never end. GPORCA has been architected to handle these types of complex queries in a much more efficient way by removing unnecessary nesting, decorrelating the subquery by pulling up the subquery predicates, and by converting subqueries into much more efficient table joins.[5]

Another area of improvement is with Common Table Expressions (CTE), which are temporary tables that are used for just one query and typically heavily utilized in analytical workloads. For instance, in the TPC-DS, 46 of the 111 queries leverage CTE. GPORCA introduces a new producer-consumer model for the WITH clause. The model allows evaluating a complex expression once, and consuming its output by multiple operators. This gives Pivotal Greenplum the ability to work with much more complex CTE because it is not forced to fully expand them, but instead deals with them dynamically. On average, CTE plans generated by GPORCA are seven times faster over the traditional Planner optimizer.[6]

A common method of organizing data in a large table is the use of partitions. Greenplum 5, with GPORCA as the default query optimizer, improves dynamic partition elimination by identifying cost-effective ways to eliminate partitions that are not relevant to the result. This is achieved by the introduction of three new query operators that work together in a producer/consumer model to perform scans over partitioned tables: PartitionSelector, DynamicScan, and Sequence. The placement of these PartitionSelectors in the query plan allow GPORCA to support more complex patterns, such as partition selection based on equality and range predicates, as well as dynamic partition elimination.[7]

The ANALYZE command in Greenplum 5 uses the faster PostgreSQL implementation to gather table statistics, improving its performance for both heap and append-optimized tables. A sample of rows is collected in a single query, and a calculation of statistics for each column is performed in memory. Previously, separate queries for each column were issued. During the analyze operation, a table is no longer created to hold the sample. These improvements produce better overall statistics for queries that involve large tables, and greatly improve the efficiency for running ANALYZE on smaller tables in the database.

5. https://content.pivotal.io/blog/greenplum-database-adds-the-pivotal-query-optimizer
6. https://content.pivotal.io/blog/greenplum-database-adds-the-pivotal-query-optimizer
7. https://content.pivotal.io/blog/greenplum-database-adds-the-pivotal-query-optimizer

Additional Features and Functionality

Pivotal Greenplum 5 adds support for a number of new built-in datatypes, and improves upon existing data types. Added types include Universally Unique Identifiers (UUID) that conform to the RFC 4122 and ISO/IEC 9834-8:2005 standards. The JSON type allows for the storage and manipulation of variable, unlimited-length JSON data, including new built-in functions to query and parse JSON records. Similar to JSON, HSTORE allows for semistructured data to be stored and queried in the database. XML support is enhanced, as new PostgreSQL 9.1 functions and SQL commands have been added into Greenplum. These new data types and related functions help to provide an environment for storage and analysis of structured, semi-structured, and unstructured data all in one place, without silos. Other features include dblink, which is a new module for establishing connections to other databases, either on the same database host, or on a remote host. Pivotal Greenplum provides dblink support for database users to perform short ad hoc queries in other databases. Finally, Greenplum 5 upgrades the installed Python version to 2.7. PL/Python and core Python management utilities are now based on version 2.7.

Conclusion

Greenplum 5 is a next-generation data platform offered by Pivotal, which brings with it several new features and improvements to a product with over ten years of engineering investment, used by organizations to manage their enterprise data warehouse and advanced analytics requirements. This new version expands the customer's deployment options by certifying numerous cloud platform alternatives, in addition to the currently supported on-premises platforms. It helps customers eliminate analytical silos by providing a single environment for traditional and new analytical workloads. For data scientists and business analysts, it offers diverse analytics through an integrated environment that helps accelerate innovation. Pivotal Greenplum 5 brings with it the latest generation of data management and data analysis features by way of modern development practices, continuous delivery, and an open source approach to software distribution that will help to ensure its success for a long time to come.

PVWP-Greenplum-1f2 / CRJM0917