
Pivotal®
Transforming How The World Builds Software

Pivotal Greenplum®

Pivotal Greenplum Roadmap

Ivan Novick

Legal Disclaimer

This presentation contains statements relating to Pivotal's expectations, projections, beliefs and prospects which are "forward-looking statements" about Pivotal's future which by their nature are uncertain. Such forward-looking statements are not guarantees of future performance, and you are cautioned not to place undue reliance on these forward-looking statements. Actual results could differ materially from those projected in the forward-looking statements as a result of many factors, including but not limited to: (i) adverse changes in general economic or market conditions; (ii) delays or reductions in information technology spending; (iii) risks associated with managing the growth of Pivotal's business, including operating costs; (iv) changes to Pivotal's software business model; (v) competitive factors, including pricing pressures and new product introductions; (vi) Pivotal's customers' ability to transition to new products and computing strategies such as cloud computing, the uncertainty of customer acceptance of emerging technologies, and rapid technological and market changes; (vii) Pivotal's ability to protect its proprietary technology; (viii) Pivotal's ability to attract and retain highly qualified employees; (ix) Pivotal's ability to execute on its plans and strategy; and (x) risks related to data and information security vulnerabilities. All information set forth in this presentation is current as of the date of this presentation. These forward-looking statements are based on current expectations and are subject to uncertainties and changes in condition, significance, value and effect as well as other risks disclosed previously and from time to time in documents filed by Dell Technologies Inc., the parent company of Pivotal, with the U.S. Securities and Exchange Commission. Dell and Pivotal assume no obligation to, and do not currently intend to, update any such forward-looking statements after the date of this presentation. The following is intended to outline the general direction of Pivotal's offerings. It is intended for information purposes only and may not be incorporated into any contract. Any information regarding pre-release of Pivotal offerings, future updates or other planned modifications is subject to ongoing evaluation by Pivotal and is subject to change. This information is provided without warranty or any kind, express or implied, and is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions regarding Pivotal's offerings. These purchasing decisions should only be based on features currently available. The development, release, and timing of any features or functionality described for Pivotal's offerings in this presentation remain at the sole discretion of Pivotal. Pivotal has no obligation to update forward-looking information in this presentation.

Pivotal Engineering Practices

concourse

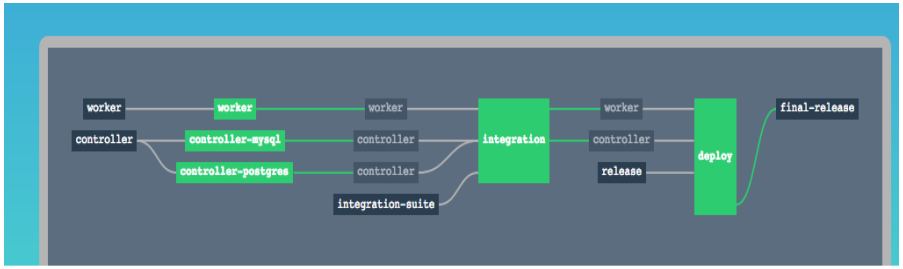

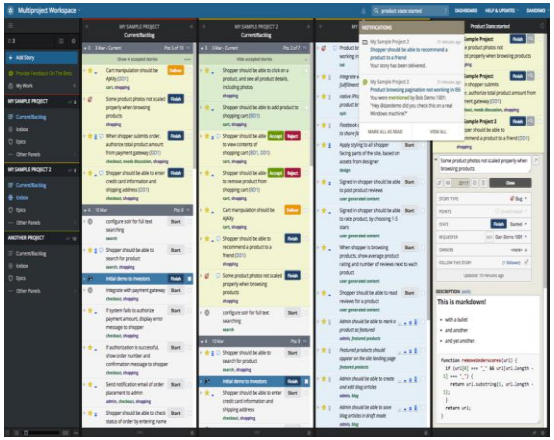

- Development teams in multiple geos
- Same methodology as all Pivotal teams
- Monthly Minor Greenplum Releases
- Annual Major Greenplum Releases

Pivotal

Pivotal Greenplum 5

TARGET: (Alpha: April, Beta: May, GA: June)

- Open Source Based Database
- PG 8.2 → PG 8.3
- PostgreSQL based Analyze (faster)
- Asynchronous Dispatcher
- GPORCA is Default Optimizer
- JSON Type & Functions
- Improved XML Type/Functions
- UUID Type
- Anonymous Code Blocks (Do statement)
- UDF default and Variadic parameters
- DBLink
- Raster PostGIS
- Python 2.7.12
- Heap Only Tuples (improves catalog maintenance)
- Lazy XID (less frequent xid wrap around)
- Gem Connector (Additional Modes)
- Resource Groups (CPU Targets) (post GA)
- GP Spark Driver (post GA)
- gpload multi-byte delimiters (post GA)
- PXF For Hadoop (post GA)

Pivotal Greenplum 6+

Current Active Projects

- Accelerate to Annual Major Release Cadence
- Improved Major Upgrade Framework
- One or more PostgreSQL upgrades
- Segment Write Ahead Log Replication
- Full Text Search
- GIN Indices (text, json, xml)
- Column Level Permissions
- Recursive CTEs

Longer Term Priorities

Blocked on dependency projects

- Cluster to Cluster Replication
- Point in Time Recovery (based on WAL replication)
- Disk Space Quotas
- Replicated Table Types
- JSONB Data Type
- Foreign Data Wrappers
- Virtual Segments for External Tables

Pivotal Greenplum Backup Restore

- Rewrite gpcrondump for improved user satisfaction
  - Released in 5.x series
  - Restore to different topologies
  - Reduced Lock Contention on pg_class
  - Progress Reporting & Improved Error Reporting
  - Pluggable Backends (Data Domain, NetBackup, S3, Other…)

- Continuous Archiving & Point in Time Recovery
  - Dependent on WAL Replication project (GP 6+)
  - Less intrusive backup system
  - Full cluster restore
  - Restore to point in time

Greenplum Command Center Roadmap

DBA Management Console

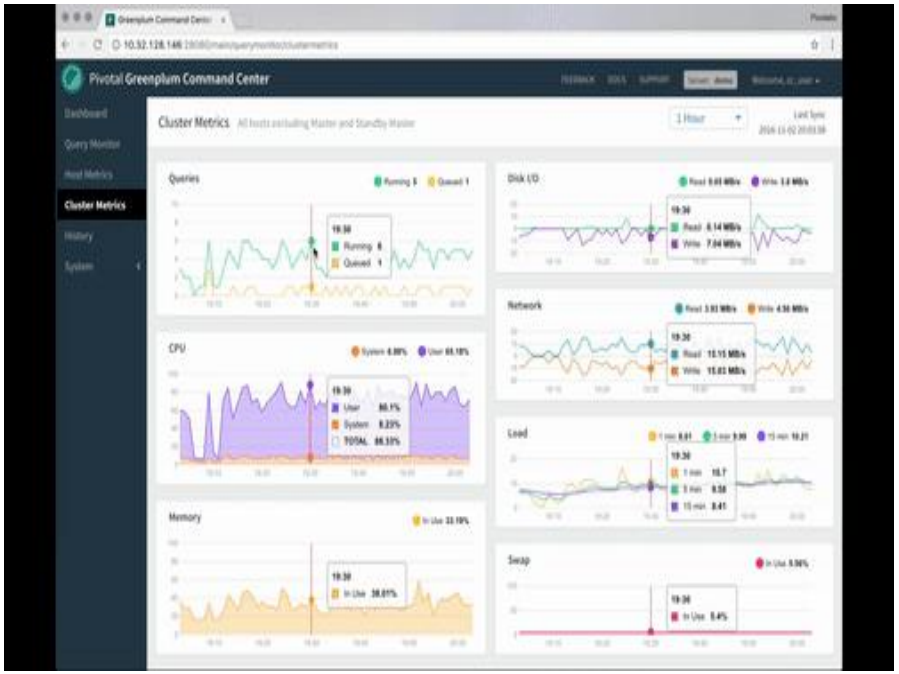

- WLM Visual Rule Screen
- WLM Rule Creation
- WLM Rule Editing
- WLM Events and Records View
- Improved Installation Experience
- Kerberos Single Sign On

Check out 3.1.0 with pg_hba management view

PGAdmin 4

SQL Developer Console

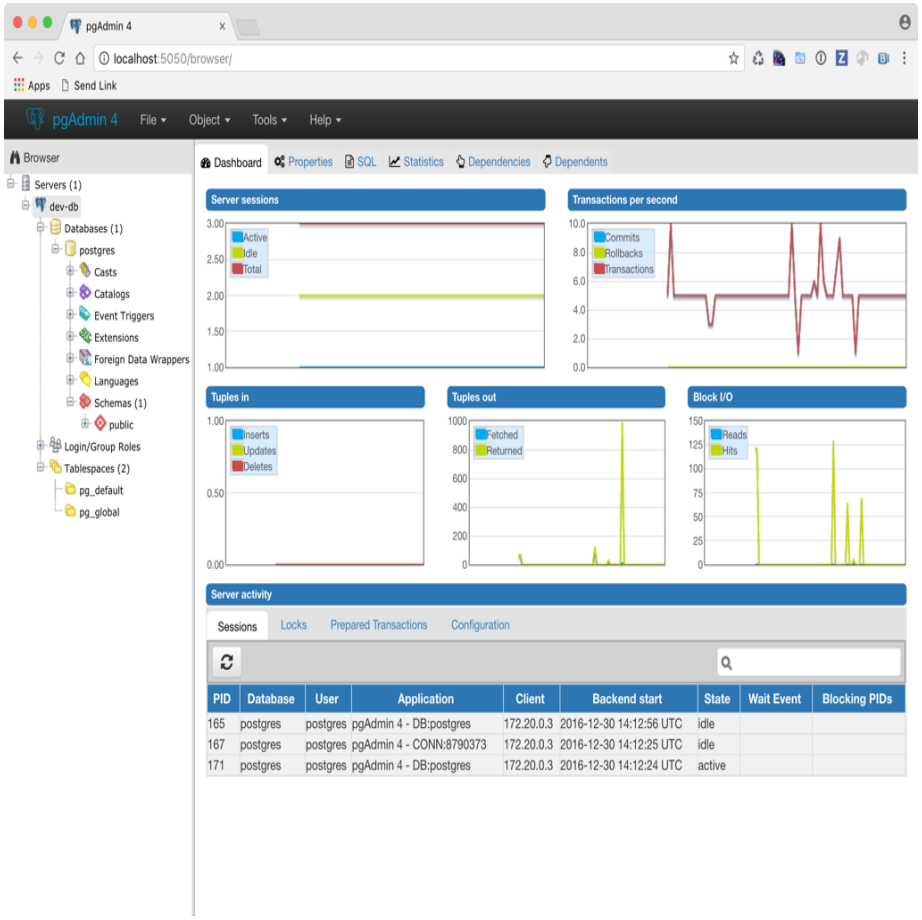

- Pivotal Collaborating with PostgreSQL community
- Download from PivNet
- Greenplum 5.x Support
- Developer UI
- Schema Browser
- Interactive Query Window

Graphs Analytical Processing in Pivotal Greenplum

Yes!

- Graph analytic processing on Pivotal Greenplum's MPP architecture can solve for a wide range of real world use cases
- Reduced operational challenges of dedicated graph database engine
- The case against specialized graph analytics engines
  http://cidrdb.org/cidr2015/Papers/CIDR15_Paper20.pdf: University of Wisconsin
- This software is being developed today in Apache MADlib (incubating)

Introduction to Graphs

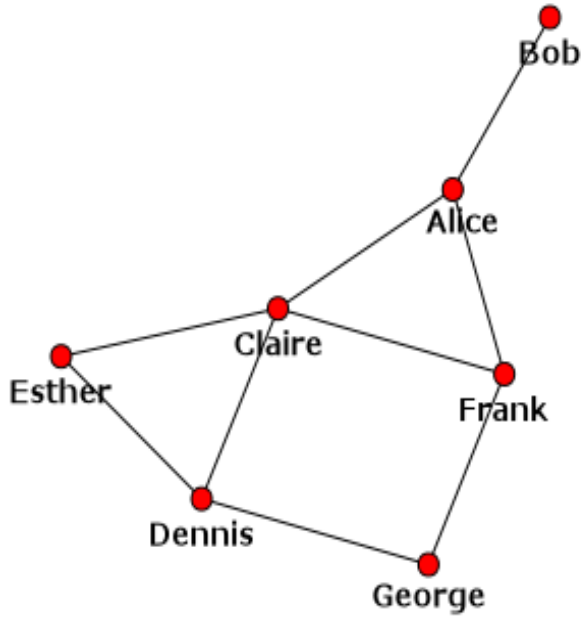

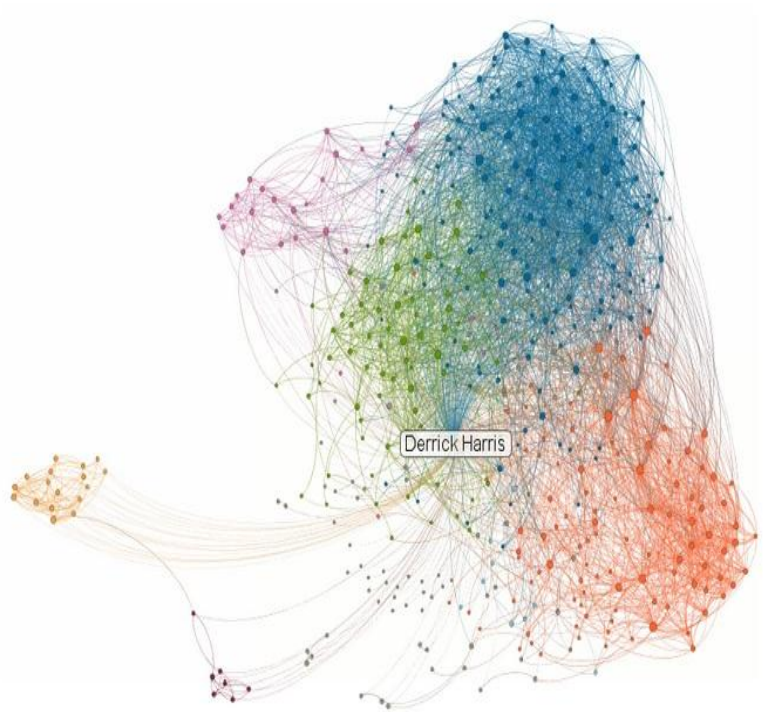

Use cases: Social Network, Computer Networks, Security, etc….

Graphs Representation in Pivotal Greenplum

Vertex Table

Vertex	Vertex Params	...
0	...	
1	...	
2	...	
3	...	

Edge Table

Source Vertex	Dest Vertex	Edge Weight	Edge Params	...
0	3	1.0	...	
1	0	5.0	...	
1	2	3.0	...	
2	3	8.0	...	
3	0	3.0	...	
3	1	2.0		

Calling a graph function in Pivotal Greenplum

Apache Madlib Example Function

Single Source Shortest Path

```
graph_sssp( vertex_table,     -- vertex table
            vertex_id,                -- col in vertex table containing vertex IDs
            edge_table,               -- edge table
            edge_args,        -- source, dest and edge weights col in the edge table
            source_vertex,    -- source vertex for the algorithm to start
            sssp_table                -- output table of SSSP for all dest vertices
                    );
```

Path retrieval

```
graph_sssp_get_path( sssp_table,                -- sssp table
                     dest_vertex                -- dest of the path of interest
                   );
```

Madlib Graph Roadmap

- PageRank - measuring importance of vertices
- Graph cut - partition a graph into two disjoint subsets
- Connected components - resiliency measure
- Betweenness - influencer nodes and edges
- Graph search - traversal algorithms
- Export in form for viz - e.g. https://gephi.org/

We are open to discuss algorithms and priority
with any Pivotal customer interested in graph analytics.

GPText Roadmap

90% of unstructured data is text! Integrate Text into a Data Warehouse

- Integrated Partitioning with Greenplum
- Raw Document Formats: PDF, Word, etc
- Need Feedback: Machine Log Analytics?