Pivotal

Graph Analytics on Massively Parallel Processing Databases

Frank McQuillan

Feb 2017

MADlib

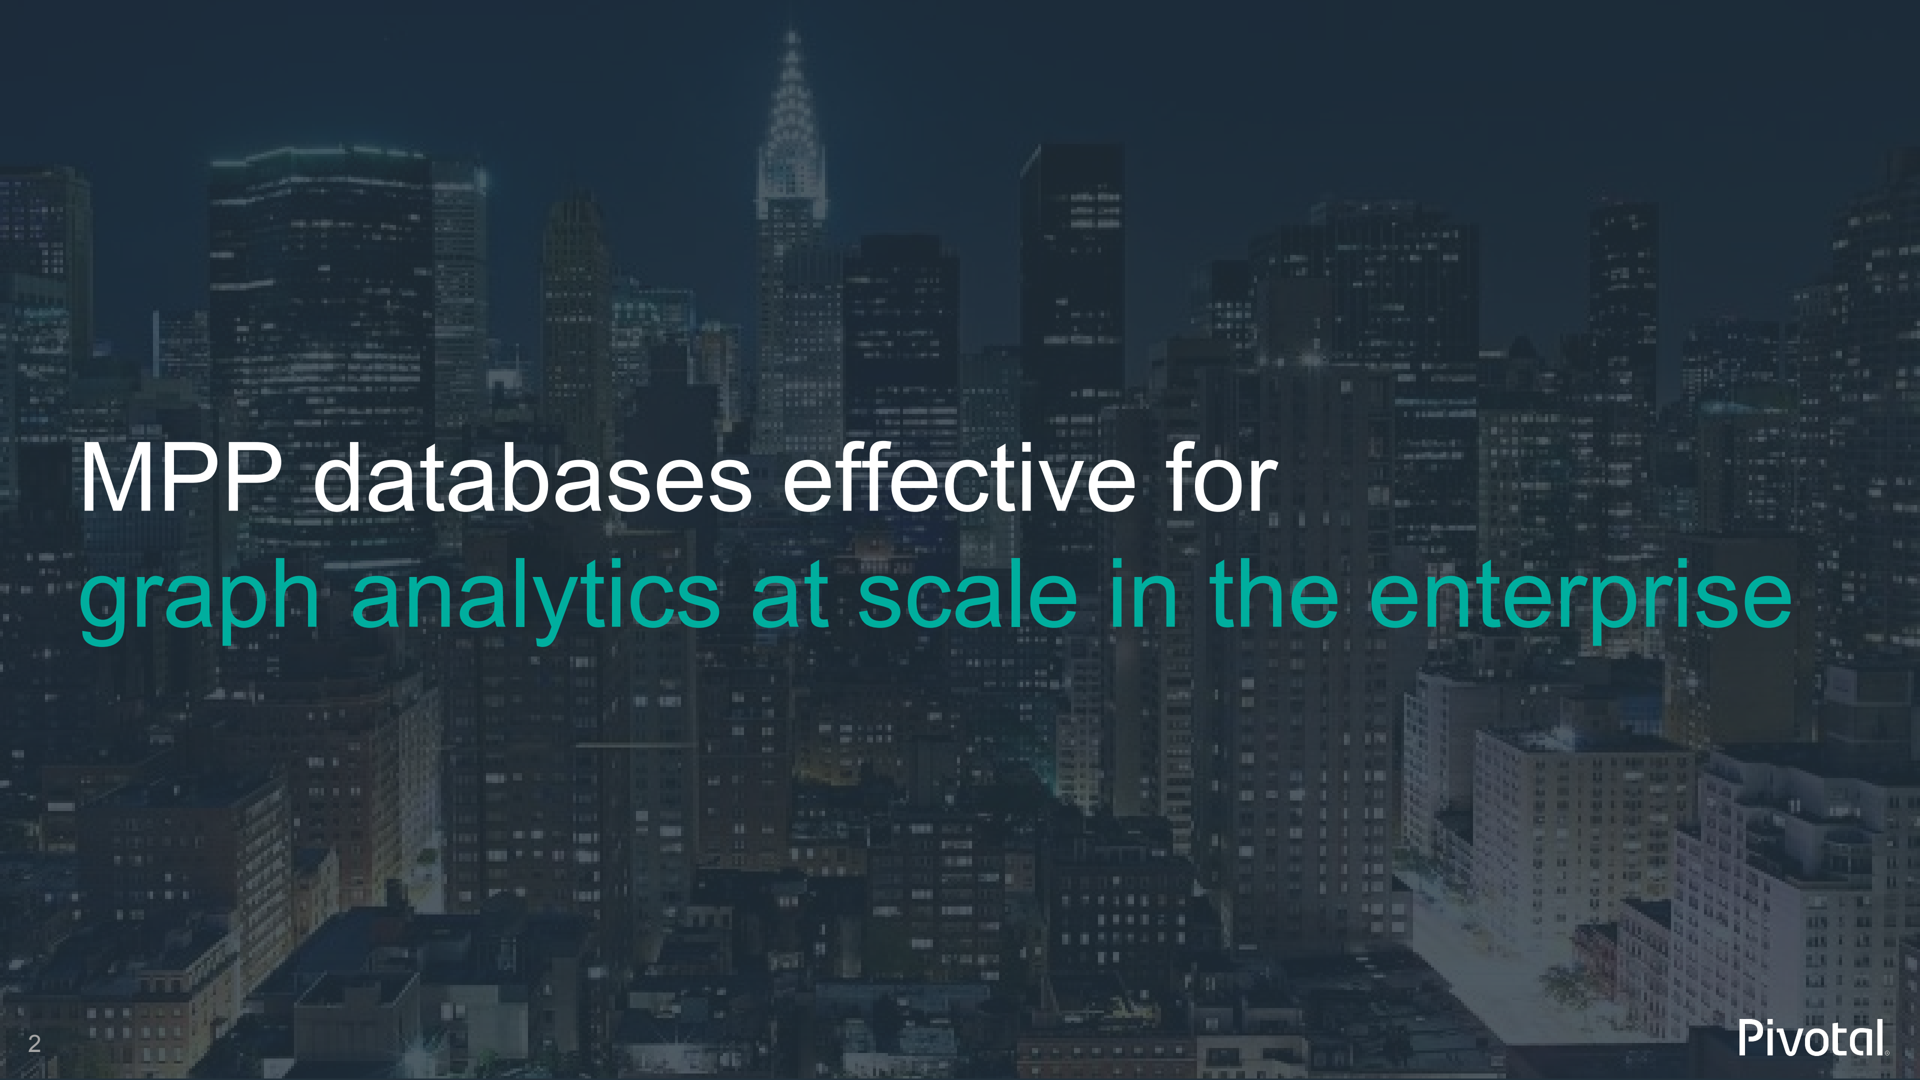

MPP databases effective for graph analytics at scale in the enterprise

Database Engine Popularity

315 systems in ranking, January 2017

Rank Jan 2017	Rank Dec 2016	Rank Jan 2016	DBMS	Database Model	Score Jan 2017	Score Dec 2016	Score Jan 2016
1.	1.	1.	Oracle	Relational DBMS	1416.72	+12.32	-79.36
2.	2.	2.	MySQL	Relational DBMS	1366.29	-8.12	+67.03
3.	3.	3.	Microsoft SQL Server	Relational DBMS	1220.95	-5.70	+76.89
4.	↑ 5.	4.	MongoDB	Document store	331.90	+3.22	+25.88
5.	↓ 4.	5.	PostgreSQL	Relational DBMS	330.37	+0.35	+47.97
6.	6.	6.	DB2	Relational DBMS	182.49	-1.85	-13.88
7.	7.	↑ 8.	Cassandra	Wide column store	136.44	+2.16	+5.49
8.	8.	↓ 7.	Microsoft Access	Relational DBMS	127.45	+2.75	-6.59
9.	9.	↑ 10.	Redis	Key-value store	118.70	-1.20	+17.54
10.	10.	↓ 9.	SQLite	Relational DBMS	112.38	+1.54	+8.64
21.	21.	21.	Neo4j	Graph DBMS	36.26	-0.56	+3.26
47.	47.	↓ 44.	Titan	Graph DBMS	5.50	+0.04	-0.15
112.	112.	↑ 118.	Giraph	Graph DBMS	1.04	+0.00	+0.11
180.	↑ 192.	↑ 235.	GraphDB	Multi-model	0.31	+0.09	+0.21

http://db-engines.com/en/ranking

Pivotal

Graph Engine Trends

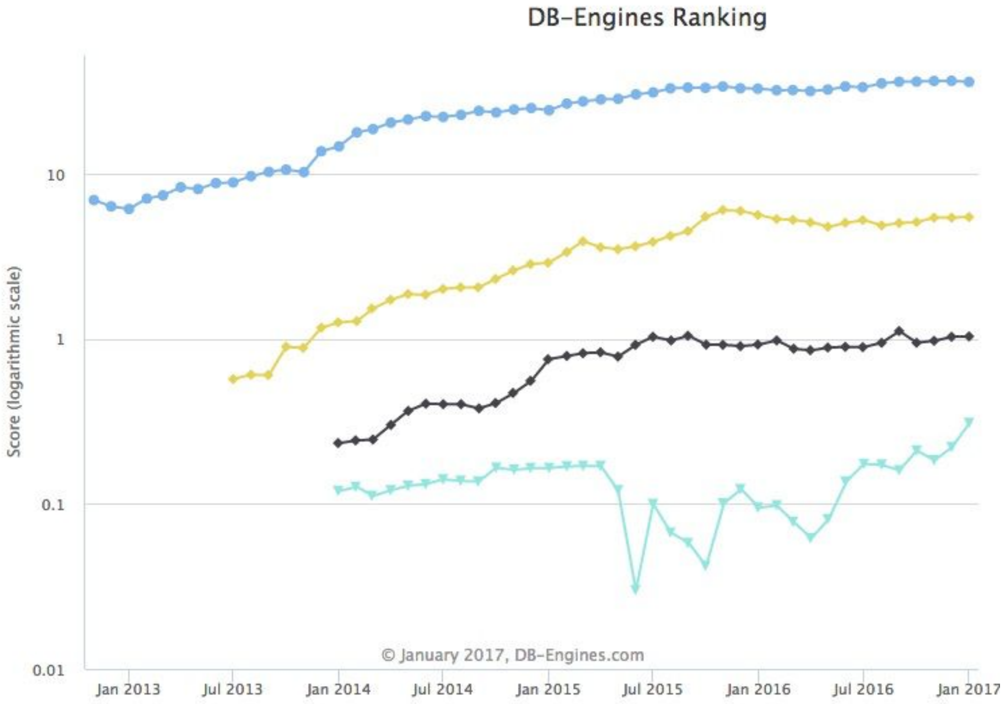

21. Neo4j

47. Titan

112. Giraph

180. GraphDB

Introduction to Graphs

- Graphs can be small...

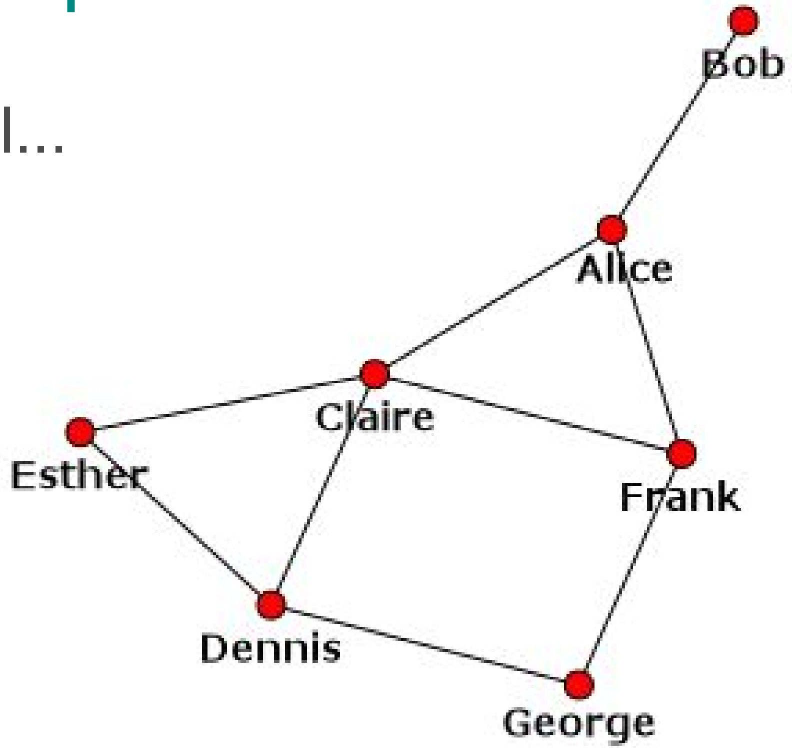

Introduction to Graphs

- ...but many real world graphs are very large

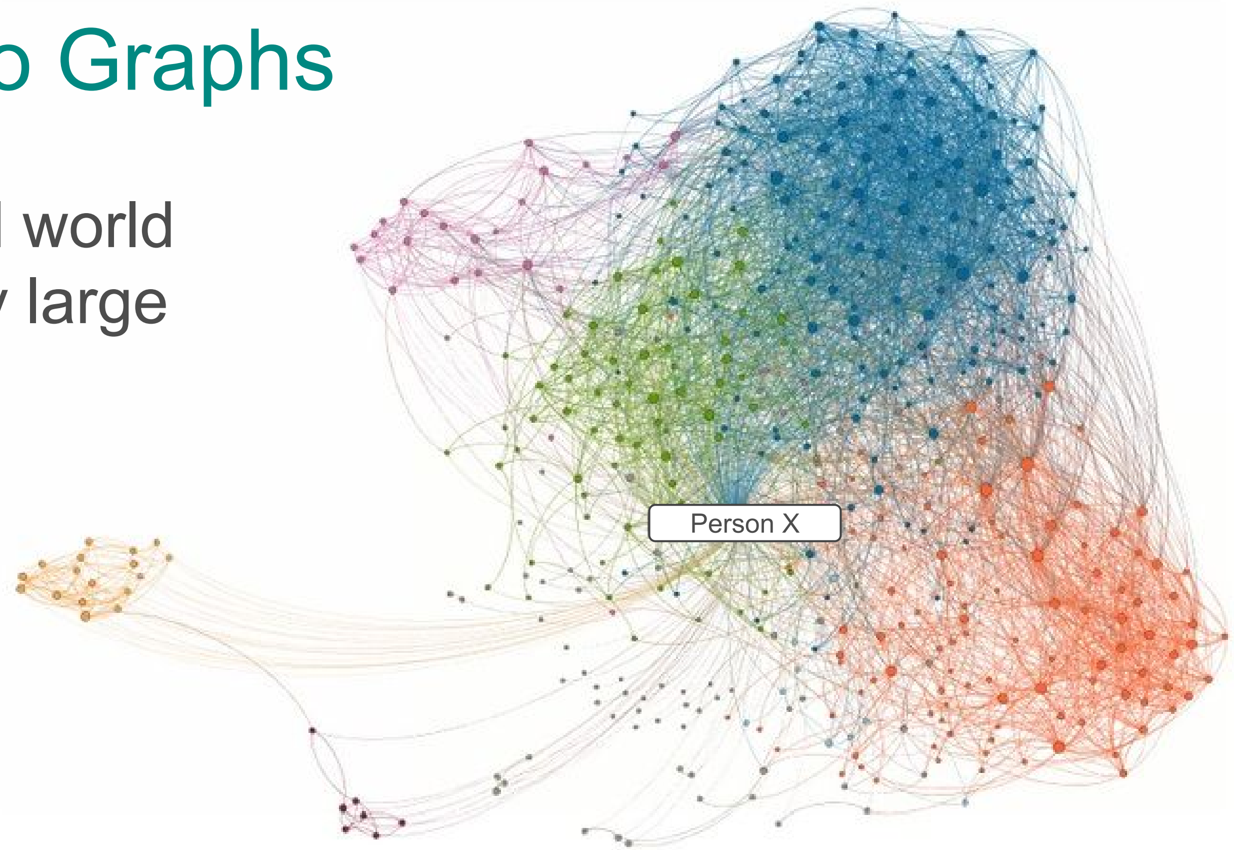

Sample LinkedIn social graph

Why Graph Analytics on MPP Databases?

- MPP is built for very large data sets
- Many enterprise use cases combine graph analytics with other techniques
- SQL
  - Most common workload in the enterprise
  - Widely used by analysts and data scientists
  - Ecosystem of business intelligence applications

Why Graph Analytics on MPP Databases?

- Data locality
  - Cost of replicating, moving and transforming data to an external system can be high
- Policy
  - Cost, deployment, oversight, support issues adding a new execution engine
  - Convince the CIO to use a specialized system in production

But...

Can graph analytic processing be
efficiently performed
on relational data in an MPP database?

Yes!

- Graph analytic processing on Greenplum database using Apache MADlib can solve for a wide range of real world use cases

Apache MADlib (incubating)

Scalable, In-Database Machine Learning

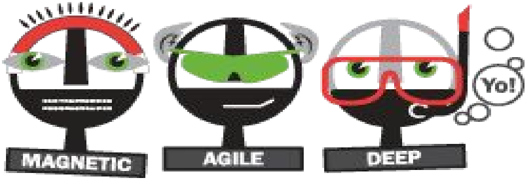

Apache MADlib (incubating): Big Data Machine Learning in SQL for Data Scientists

Open source, commercially friendly Apache license

Supports PostgreSQL, Greenplum Database™, and Apache HAWQ (incubating)

Powerful analytics for big data

- Open source https://github.com/apache/incubator-madlib
- Downloads and docs http://madlib.incubator.apache.org/
- Wiki https://cwiki.apache.org/confluence/display/MADLIB/

History

MADlib project was initiated in 2011 by EMC/Greenplum architects and Joe Hellerstein from Univ. of California, Berkeley.

UrbanDictionary.com:
mad *(adj.): an adjective used to enhance a noun.*

1- dude, you got skills.

*2- dude, you got **mad** skills.*

MADlib Functions

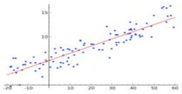

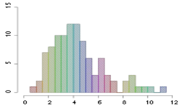

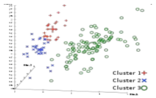

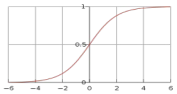

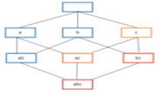

Generalized Linear Models
- Linear Regression
- Logistic Regression
- Multinomial Logistic Regression
- Ordinal Regression
- Cox Proportional Hazards Regression
- Elastic Net Regularization
- Robust Variance (Huber-White), Clustered Variance, Marginal Effects

Matrix Factorization
- Singular Value Decomposition (SVD)
- Low Rank

Linear Systems
- Sparse and Dense Solvers
- Linear Algebra

Graph
- Single Source Shortest Path

New in v1.10, more to come

Other Machine Learning Algorithms
- Principal Component Analysis (PCA)
- Association Rules (Apriori)
- Topic Modeling (Parallel LDA)
- Decision Trees
- Random Forest
- Conditional Random Field (CRF)
- Clustering (K-means)
- Cross Validation
- Naïve Bayes
- Support Vector Machines (SVM)
- Prediction Metrics
- K-Nearest Neighbors

Time Series
- ARIMA

Path Functions
- Operations on Pattern Matches

Descriptive Statistics

Sketch-Based Estimators
- CountMin (Cormode-Muth.)
- FM (Flajolet-Martin)
- MFV (Most Frequent Values)

Correlation and Covariance

Summary

Inferential Statistics

Hypothesis Tests

Utility Modules

Array and Matrix Operations

Sparse Vectors

Random Sampling

Probability Functions

Data Preparation

PMML Export

Conjugate Gradient

Stemming

Sessionization

Pivot

Example Usage

Train a model

```
SELECT madlib.linregr_train('houses',                   --- Input table
                            'houses_out',               --- Output table
                            'price',                    --- Variable to predict
                            'ARRAY[1, tax, bath, size]', --- Features in data
                            'bedroom'                   --- Group data to create
                                                        ---    multiple models
                            )
```

Predict for new data

```
SELECT houses.*,
    madlib.linregr_predict(ARRAY[1, tax, bath, size],   --- Use same features
                           model.coef)as predict
FROM houses_test, houses_out as model;                  --- Combine test data
                                                        ---    and model table
```

Greenplum Database

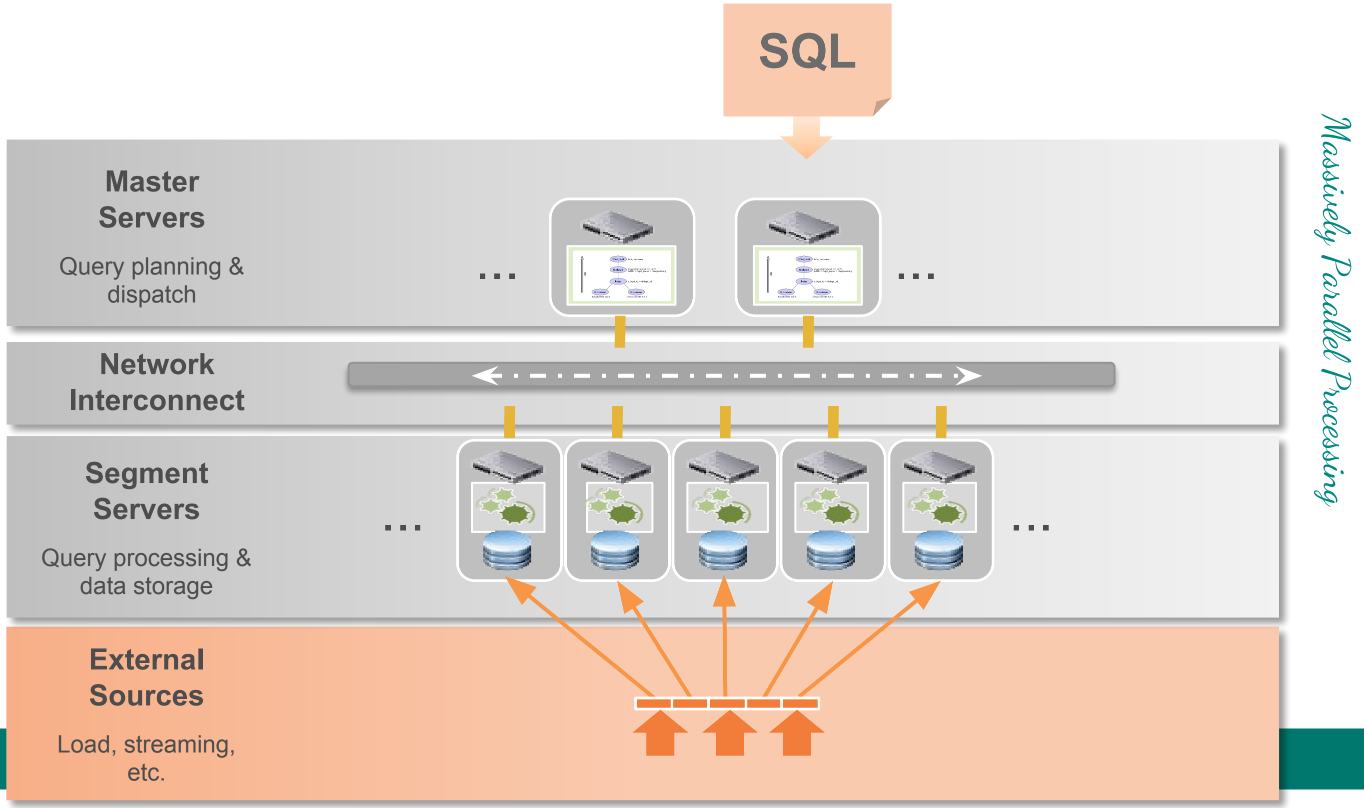

Pivotal™

MADlib on Greenplum

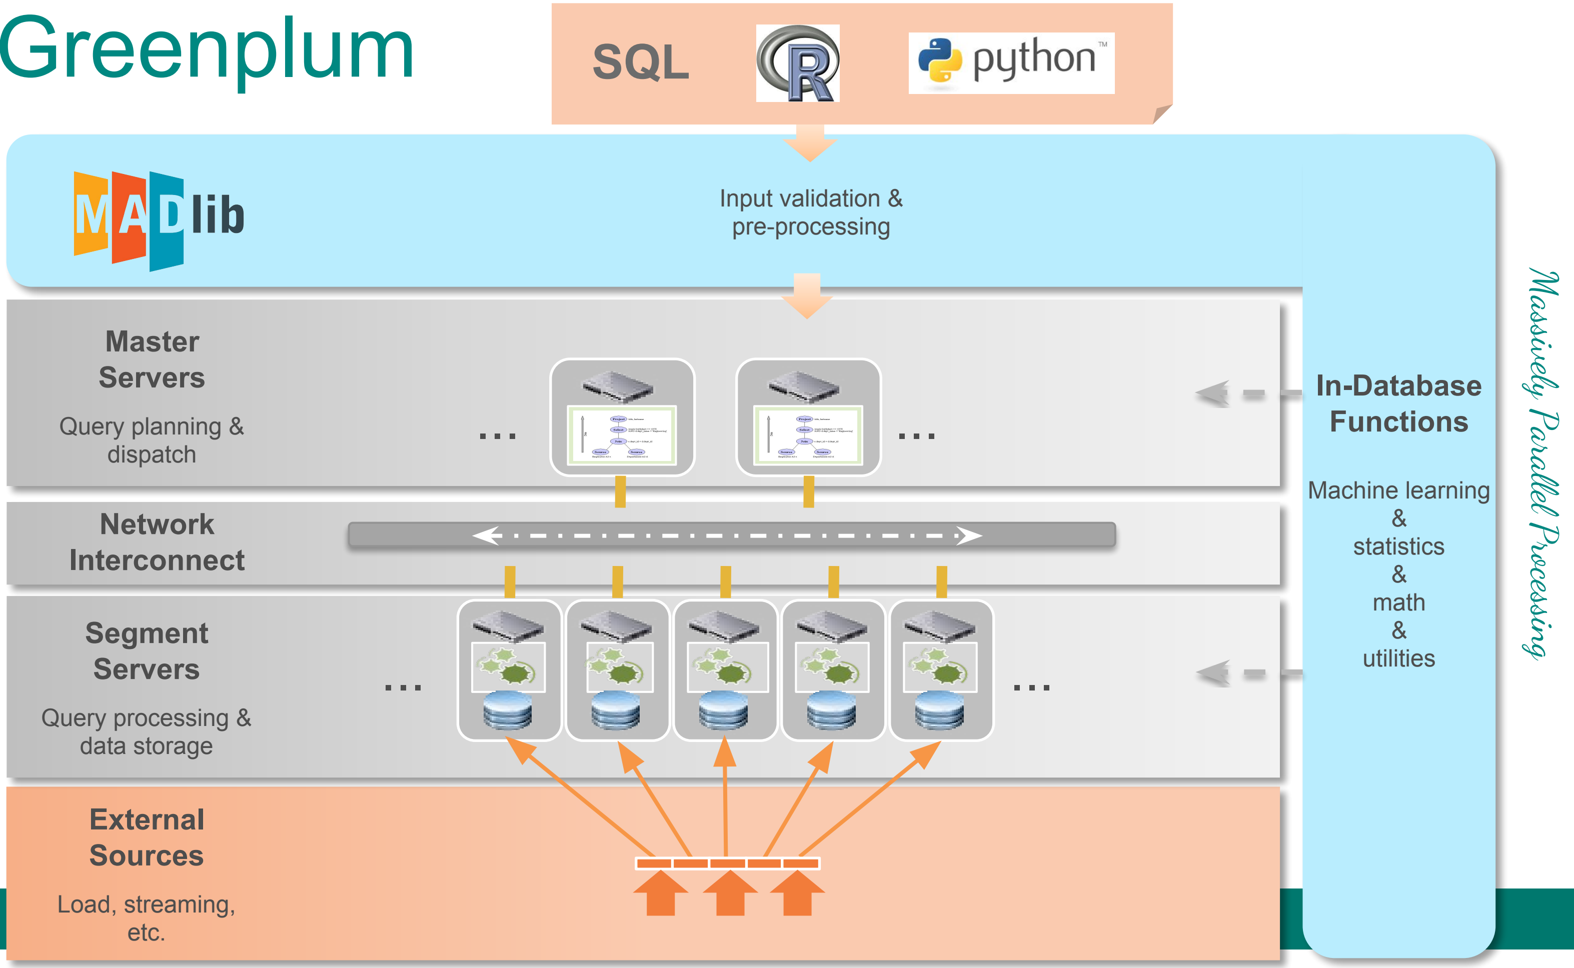

Graph Representation in MADlib

Directed graph
(digraph)

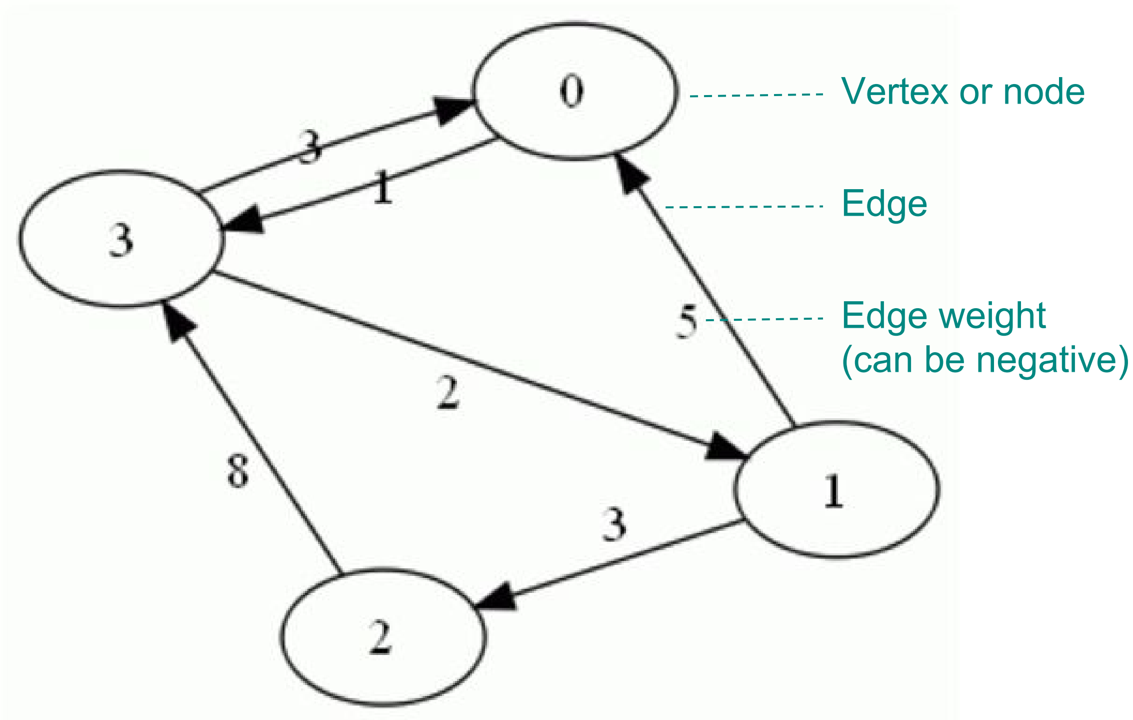

Graph Representation in MADlib

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	...	
. . .				

Single Source Shortest Path

- Given a graph and a source vertex, find a path to every vertex such that the sum of the weights of its constituent edges is minimized

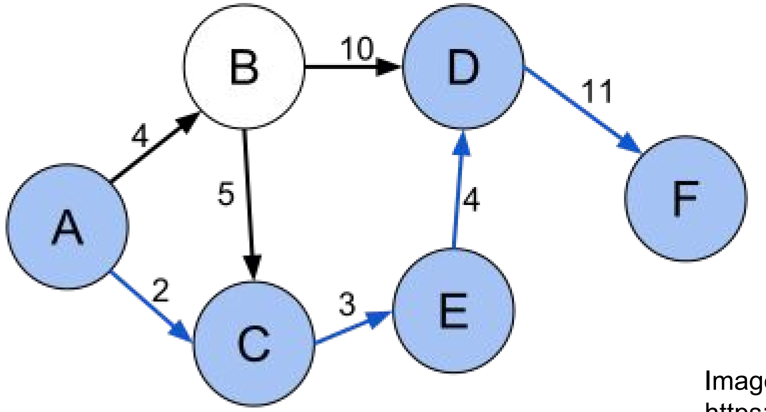

Shortest path (A, C, E, D, F) between vertices A and F in the weighted directed graph

Image from
https://en.wikipedia.org/wiki/Shortest_path_problem

Pivotal

Single Source Shortest Path

- Use cases
  - Vehicle routing/navigation
  - Degrees of separation in a social network
  - Min-delay path in a telecommunications network
  - Plant and facility layout
  - VLSI design

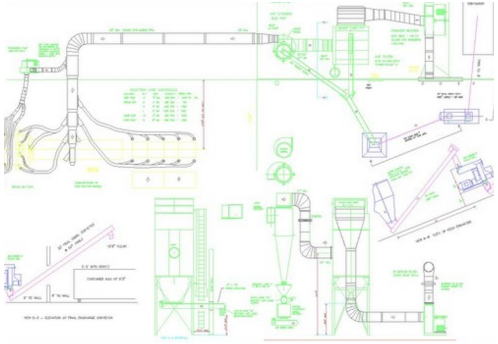

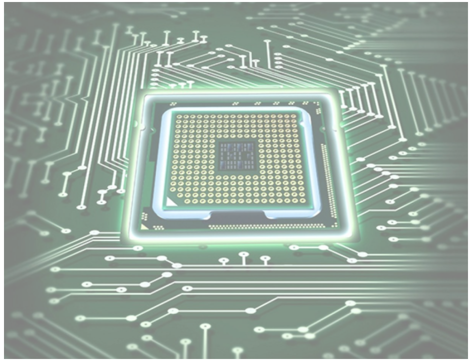

SSSP Performance on Greenplum Database

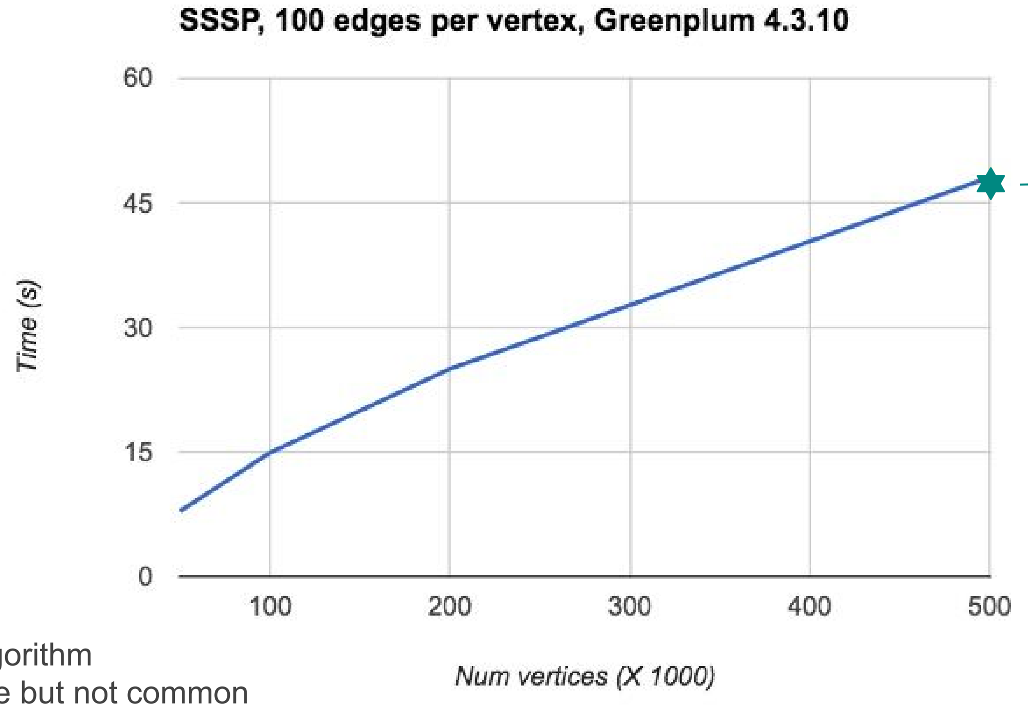

50M edges

Bellman-Ford algorithm
O(VE) worst case but not common

Greenplum cluster:

- 1 master
- 4 segment hosts with 6 segments per host

Single Source Shortest Path in MADlib

SSSP

```
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
                   );
```

Pivotal™

Implementation Considerations

- Relationships
  - Not a 1st class citizen in relational databases (unlike certain graph databases)
  - `JOIN` operations are compute and memory intensive so want to minimize
- Table scans
  - Depth first search involves more table scans (expensive) than breadth first search
  - Greedy algorithms that do not take advantage of query optimizer will be slower

Implementation Considerations

- Database limits
  - PostgreSQL limits maximum field size to 1GB

MADlib Graph Roadmap (Near Term)*

Algorithm	Uses
All pairs shortest path (APSP)	● *$O(V^3)$* Floyd-Warshall ● Betweenness and closeness centrality measures to identify influencers ● Graph diameter
Page rank	● Identify importance of vertices
Connected components	● Clustering common components ● Measure of resilience in network flow problems
Graph cut	● Partition a graph into two disjoint subsets

*Subject to community interest and contribution, and subject to change at any time without notice.

Cybersecurity Example

Lateral Movement Detection

Perimeter Defense Inadequate

- Defending the perimeter **no longer enough**
- No **100%, fool-proof way to keep bad actors out**
- Some **threats come from within**
- The idea of a **perimeter becoming obsolete** with mobile, cloud, IoT
- Need better methods for **threat detection inside the network**

APT Kill Chain

Advanced Persistent Threat (APT)

1. **Phishing and Zero Day Attack**
   A handful of users are targeted by two phishing attacks: one user opens Zero day payload (CVE-02011-0609)

2. **Back Door**
   The user machine is accessed remotely by Poison Ivy tool

3. **Lateral Movement**
   Attacker elevates access to important user, service and admin accounts, and specific systems

4. **Data Gathering**
   Data is acquired from target servers and staged for exfiltration

5. **Ex-filtrate**
   Data is exfiltrated via encrypted files over ftp to external, compromised machine at a hosting provider

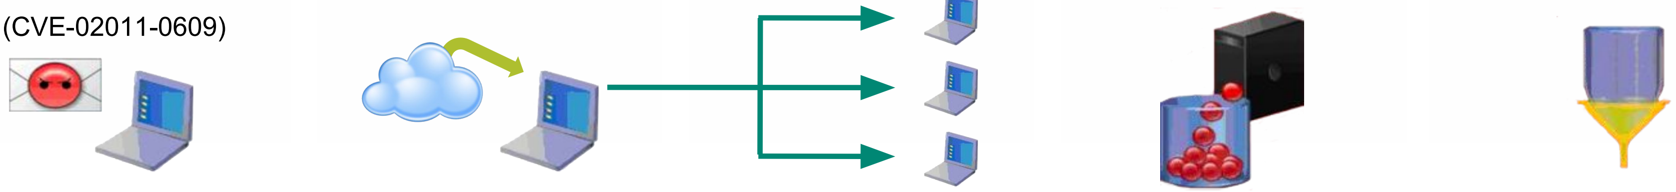

Pivotal

Lateral Movement Detection

What: Identify anomalous user-level access to hosts

How: Look at people & machines

- **Users (user behavior models)**
- **Network, servers (user peer models)**

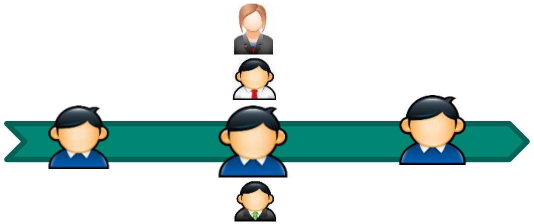

Scenarios:

Network reconnaissance from remote adversary on hijacked device

Ill-intentioned activities by legitimate employee

Access policy abuse

Business values:

Immediate security alert generation

Enhanced SIEM alert queue prioritization

Focused monitoring

Future integration with other analytic models for 360° attack view

Lateral Movement Detection (LMD) – Flow Diagram

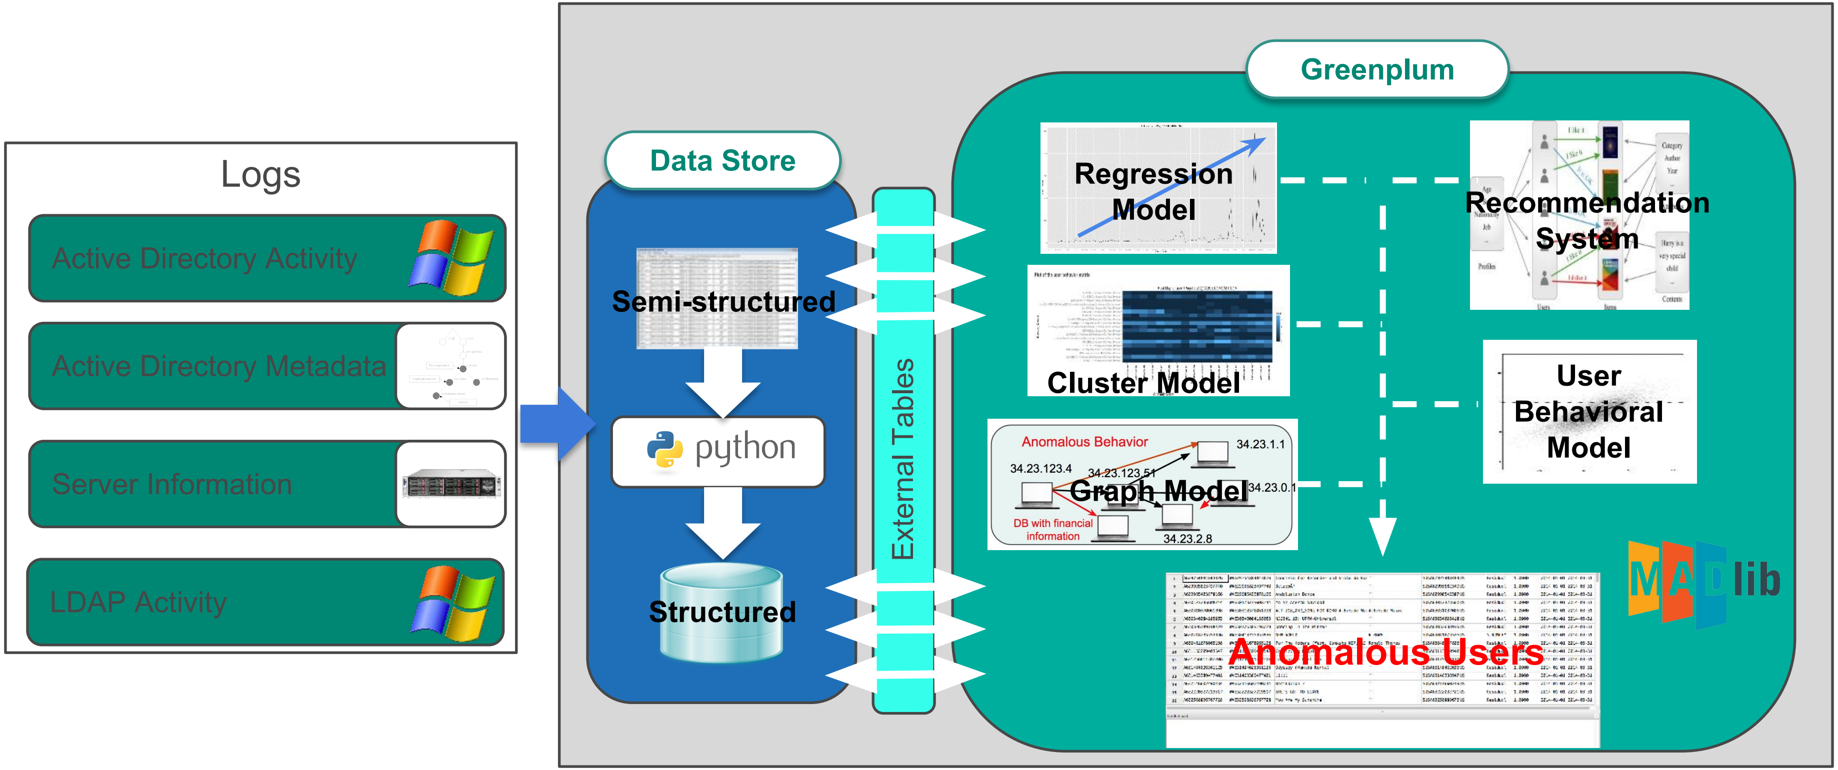

Pivotal

Regression-Based Model

- Model to identify users with unusual variation in the number of servers accessed over time
- Build a regression model for each user *(Y = aX + b)*
- No. of servers accessed each week *(Y)* ~ Week Index *(X)*
- Find the slope of the regression line for each user *(a)*
- Identify users who have a high positive or negative slope to find users with unusual activity

Regression plot of number of servers for a user

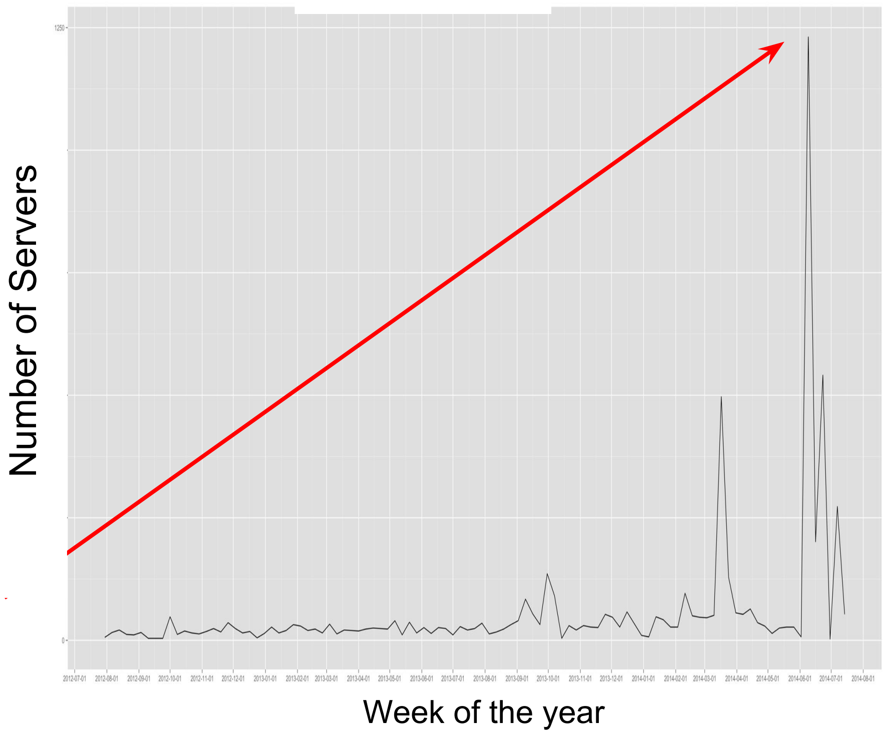

User Behavior Models (UBM)

Build historical behavioral profile for each user based on following features:

- Servers accessed
- IP addresses logged in from
- Geographical information of login

Models stress individual user/job log-in frequency

Multiple **Feature Generations** reduce false alarms:

- Aggregate servers to respective server group
- Incorporate server criticality
- Assign more weight to less popular servers and IP addresses
- E.g. print servers are low-weighted
- Use recommendation engine to suggest servers to users based on job roles and peers

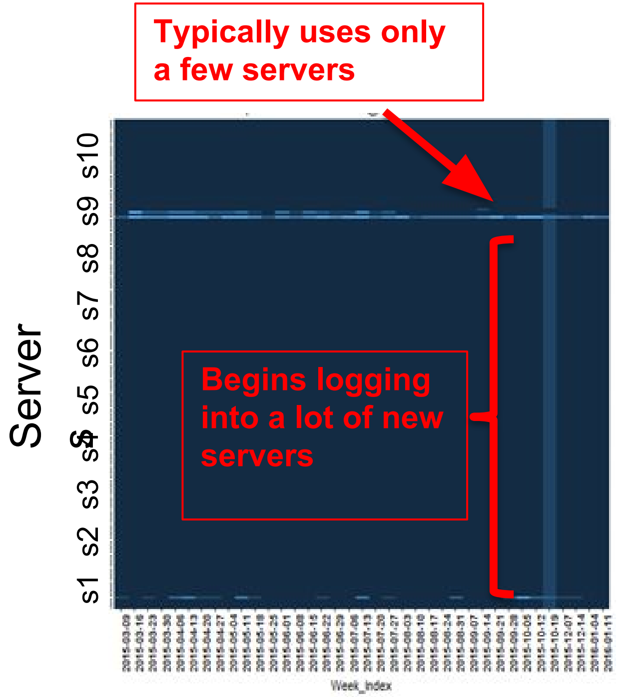

Pivotal

Graph Model

Using historical windows events data to build graphs* of typical user behavior

- Which machines does the user log into?
- Which machines does the user log in from?
- How often?
- In which order?

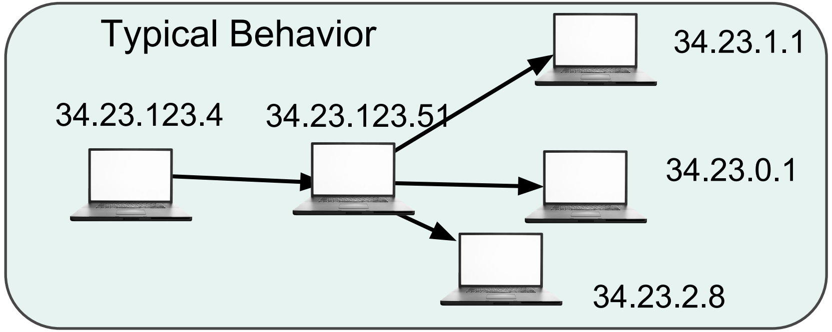

Ask if this behavior is typical

- Is it typical for this user?
- Is it typical for someone in a particular department?
- Is this typical for someone in the user's job role?

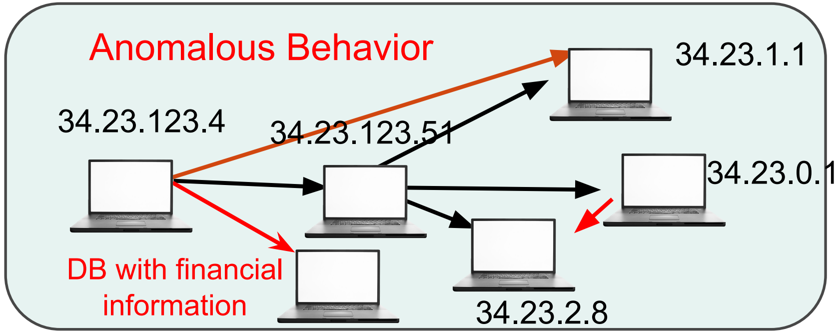

Graph models are sensitive to direction, order, and frequency

*Reference: Alexander D. Kenta, Lorie M. Liebrockb, Joshua C. Neila. Authentication graphs: Analyzing user behavior within an enterprise network.

- 4th Apache MADlib (incubating) release Feb 2017
- Project is moving toward top level status

You are welcome to join us!!!

MPP databases effective for graph analytics at scale in the enterprise

References

[1] The case against specialized graph analytics engines

http://cidrdb.org/cidr2015/Papers/CIDR15_Paper20.pdf

http://pages.cs.wisc.edu/~jignesh/publ/Grail-slides.pdf

[2] MADlib papers

http://db.cs.berkeley.edu/papers/vldb09-madskills.pdf

https://www2.eecs.berkeley.edu/Pubs/TechRpts/2012/EECS-2012-38.pdf

[3] Bellman-Ford algorithm
R. Bellman, “On a routing problem,” Quarterly of applied mathematics (1958), pp. 87–90.
L. R. Ford Jr, “Network flow theory,” Tech. rep. DTIC Document, 1956.

[4] Alexander D. Kenta, Lorie M. Liebrock, Joshua C. Neila, “Authentication graphs: Analyzing user behavior within an enterprise network”

Apache MADlib Resources

- Web site
  - http://madlib.incubator.apache.org/
- Wiki
  - http://incubator.apache.org/projects/madlib.html
- User docs
  - http://madlib.incubator.apache.org/docs/latest/index.html
- Technical docs
  - http://madlib.incubator.apache.org/design.pdf
- Pivotal commercial site
  - http://pivotal.io/madlib
- Mailing lists and JIRAs
  - https://mail-archives.apache.org/mod_mbox/incubator-madlib-dev/
  - http://mail-archives.apache.org/mod_mbox/incubator-madlib-user/
  - https://issues.apache.org/jira/browse/MADLIB
- PivotalR
  - https://cran.r-project.org/web/packages/PivotalR/index.html
- Github
  - https://github.com/apache/incubator-madlib
  - https://github.com/pivotalsoftware/PivotalR

Thank you!

Backup Slides

MADlib Execution Flow

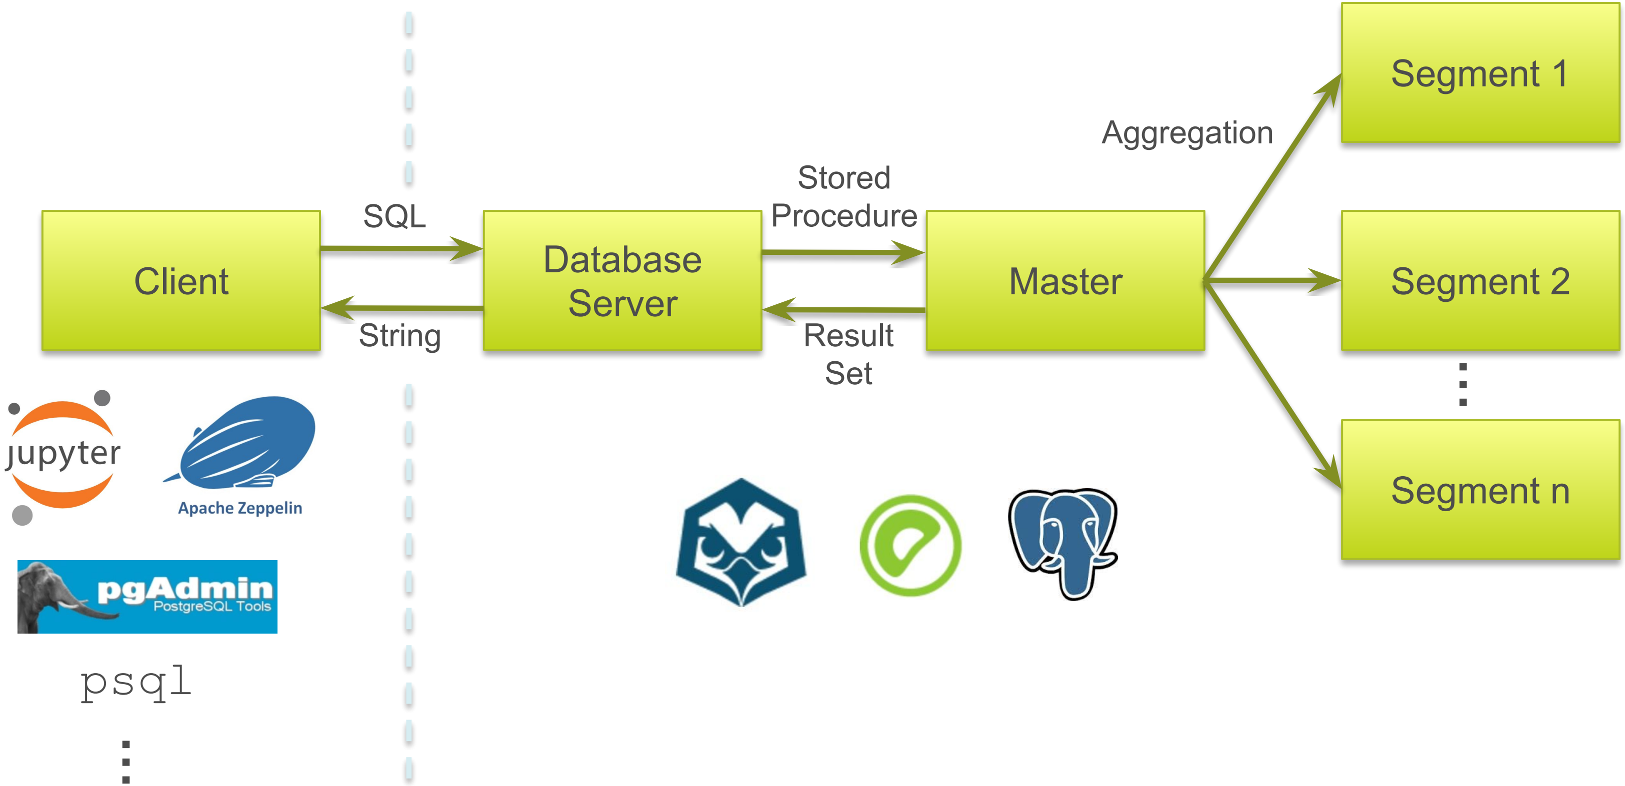

Iterative Model Execution

Stored Procedure for Model

```
model = init(…)
WHILE model not converged
  model =
    SELECT
      model.aggregation(…)
    FROM
      data table
ENDWHILE
```

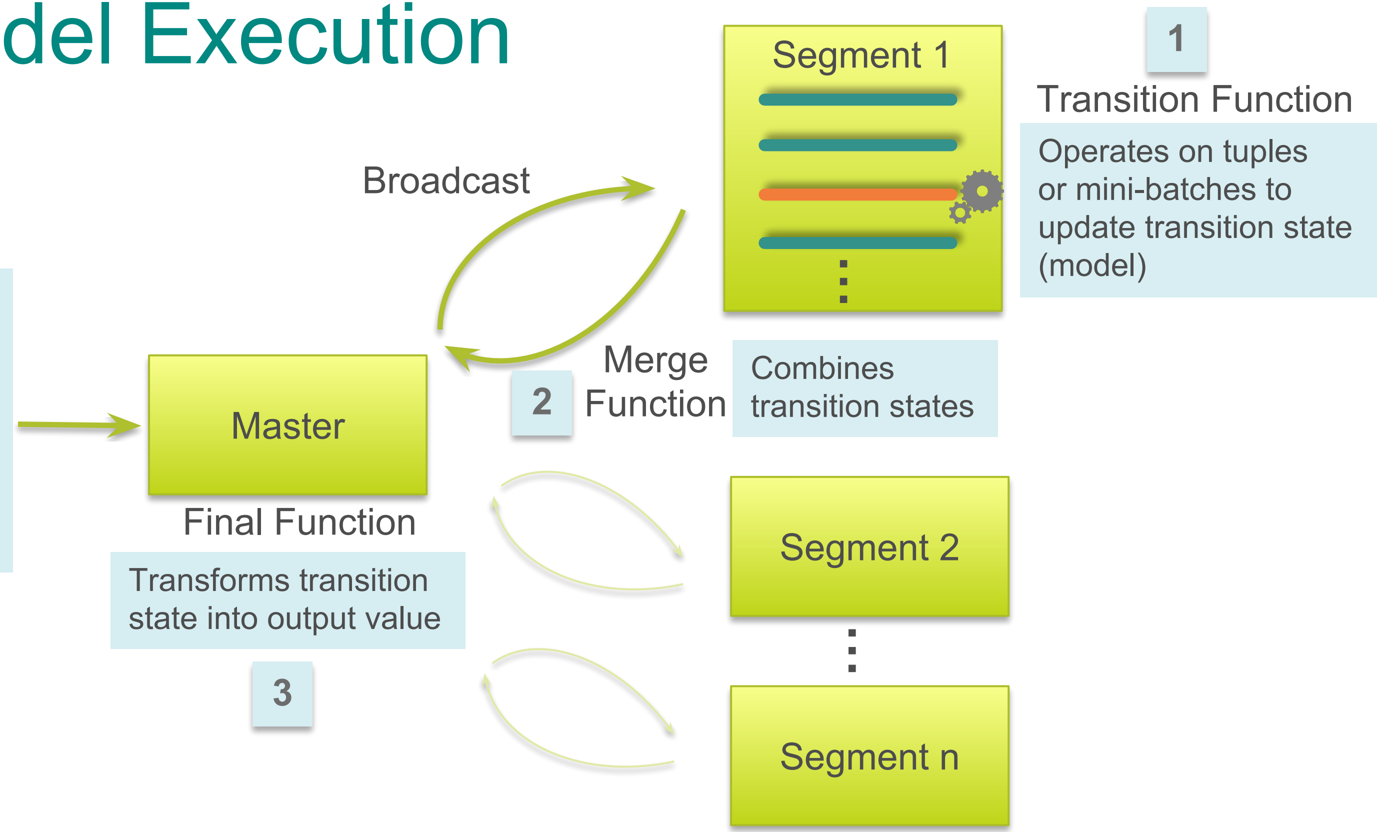

MADlib Architecture

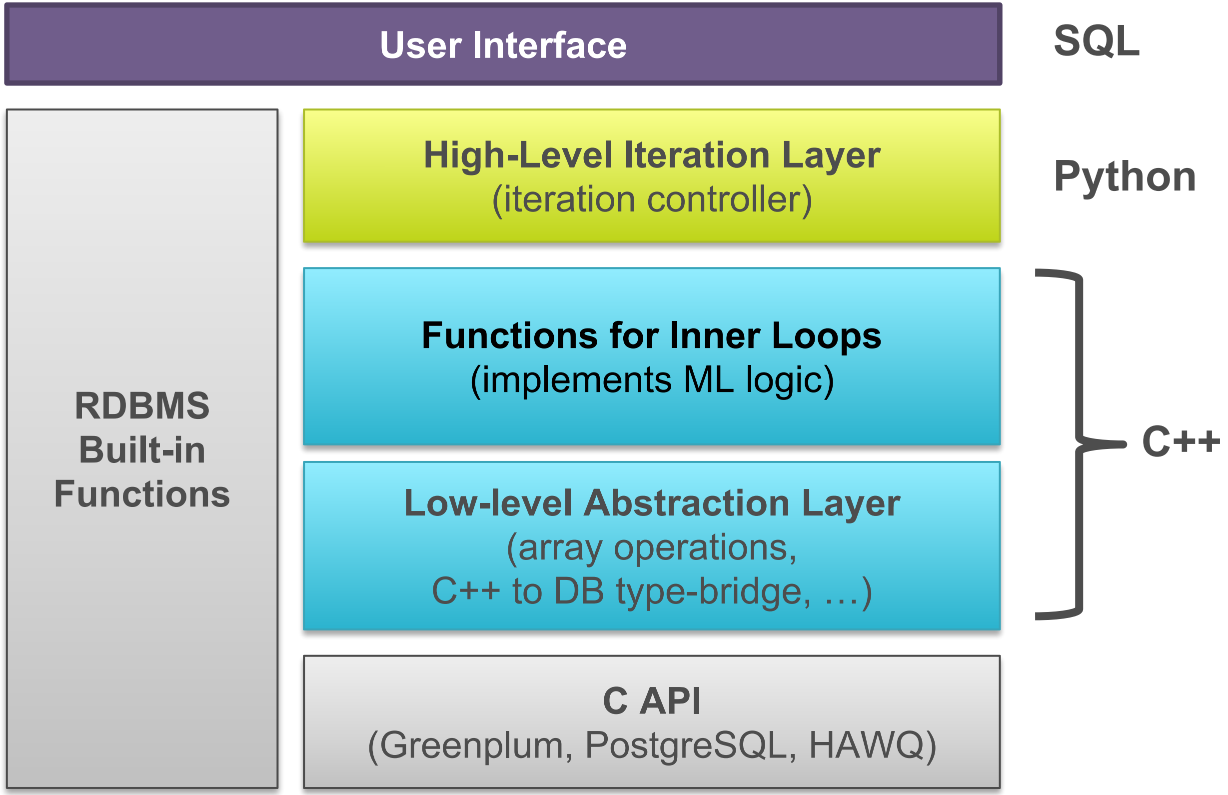

Greenplum Database

SYSTEM ACCESS

CLIENT ACCESS
PSQL, ODBC, JDBC

BULK LOAD/UNLOAD
GPLoad, GPFdist,
External Tables, GPHDFS

ADMIN TOOLS
GP Perfmon, GP Support

3rd PARTY TOOLS
Compatible with Industry
Standard BI & ETL Tools

DATA PROCESSING

SQL STANDARD COMPLIANCE

Workload Management
Resource Queues
GP Workload Manger

Big Data Query Processing
GPORCA Optimizer
MPP Query Execution

IN-DATABASE PROGRAMMING LANGUAGES
PL/pgSQL, PL/Python,
PL/R, PL/Perl, PL/Java,
PL/C

IN-DATABASE ANALYTICS & EXTENSIONS
MADlib, PostGIS,
PGCrypto

DATA STORAGE

FULLY ACID COMPLIANT TRANSACTIONAL DATABASE

POLYMORPHIC STORAGE
HEAP, Append Only,
Columnar, External,
Compression

MULTI-VERSION CONCURRENCY CONTROL (MVCC)

INDEXES
B-Tree, Bitmap,
GiST

Pivotal Query Optimizer

- Applies broad set of optimization strategies at once
  - Considers many more plan alternatives
  - Optimizes a wider range of queries
  - Optimizes memory usage
- Significant improvements for demanding queries

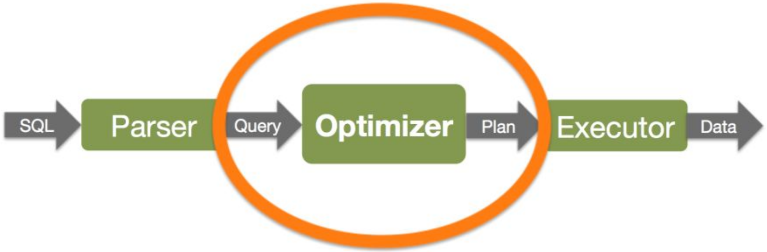

Turns a SQL query into an execution plan

Cost of Cybercrime on the Rise

- Cybercrime costs average US enterprise **$17m per year***
- Cost **grew at 15% CAGR** over last three years
- Any given cybercrime can cost significantly more
- Target’s 2014 hack cost company approximately **$162m**
- Costs not just financial, also reputational

*Source: 2016 Cost of Cyber Crime Study & the Risk of Business Innovation, Ponemon Institute