BUILT FOR THE SPEED OF BUSINESS

Technical Overview of GPDB

Updated Q3, 2013 Pivotal Product Marketing

Pivotal

Pivotal Data Fabric Portfolio

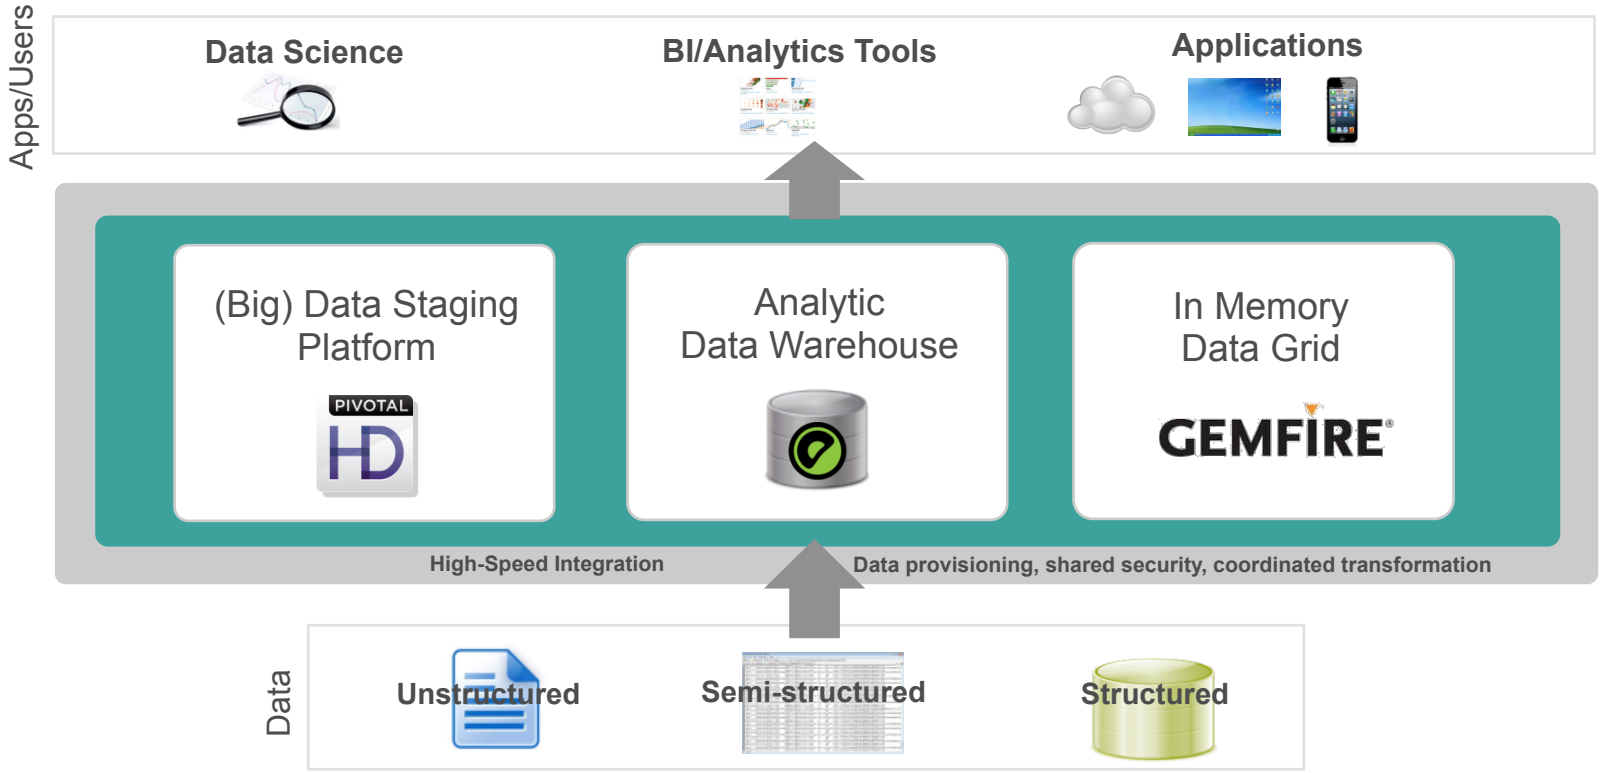

Pivotal Greenplum Database

The Analytic Data Warehouse

Pivotal Greenplum Database Overview

Pivotal

The Pivotal Greenplum Database is…

A Highly-Scalable, Shared-Nothing Database

- Leading MPP architecture, including a patented next-generation optimizer
- Optimized architecture and features for loading and queries
- Start small, scale as needed
- Polymorphic storage, compression, partitioning

A Platform for Advanced Analytics on Any (and All) Data

- Rich ecosystem (SAS, R, Chorus Studio, BI & ETL tools)
- In-DB Analytics (MADlib, Custom, languages: R, Java, Python, PERL, C, C++)
- High degree of SQL completeness so analysts can use a language they know
- Domain: Geospatial, Text processing (GPText)

An Enterprise Ready Platform Capable of Flexing With Your Needs

- Available as needed – either as an appliance or software
- Tightly integrated with other Pivotal products (Pivotal HD and Gemfire)
- Secures data in-place, in flight, and with authentication to suit
- Capable of managing a variety of mixed workloads with Pivotal VRP

Pivotal

The Pivotal Greenplum Database Overview

- A highly scalable shared-nothing database
- A platform for advanced analytics on any (and all) data
- An enterprise ready platform capable of flexing with your needs

Pivotal

MPP 101: Performance Through Parallelism

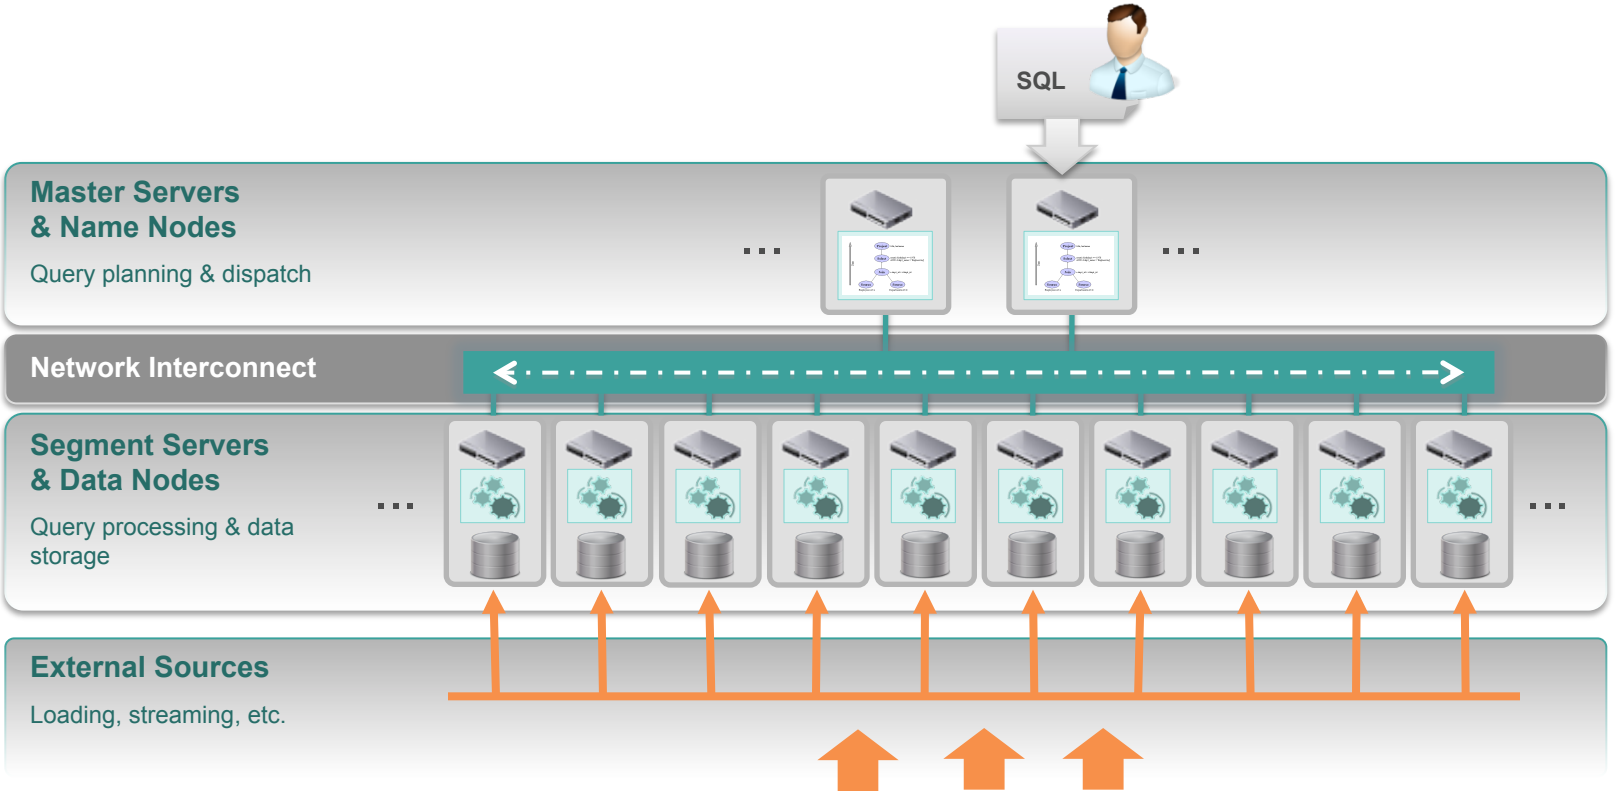

Pivotal

MPP 102: True High Speed Loading

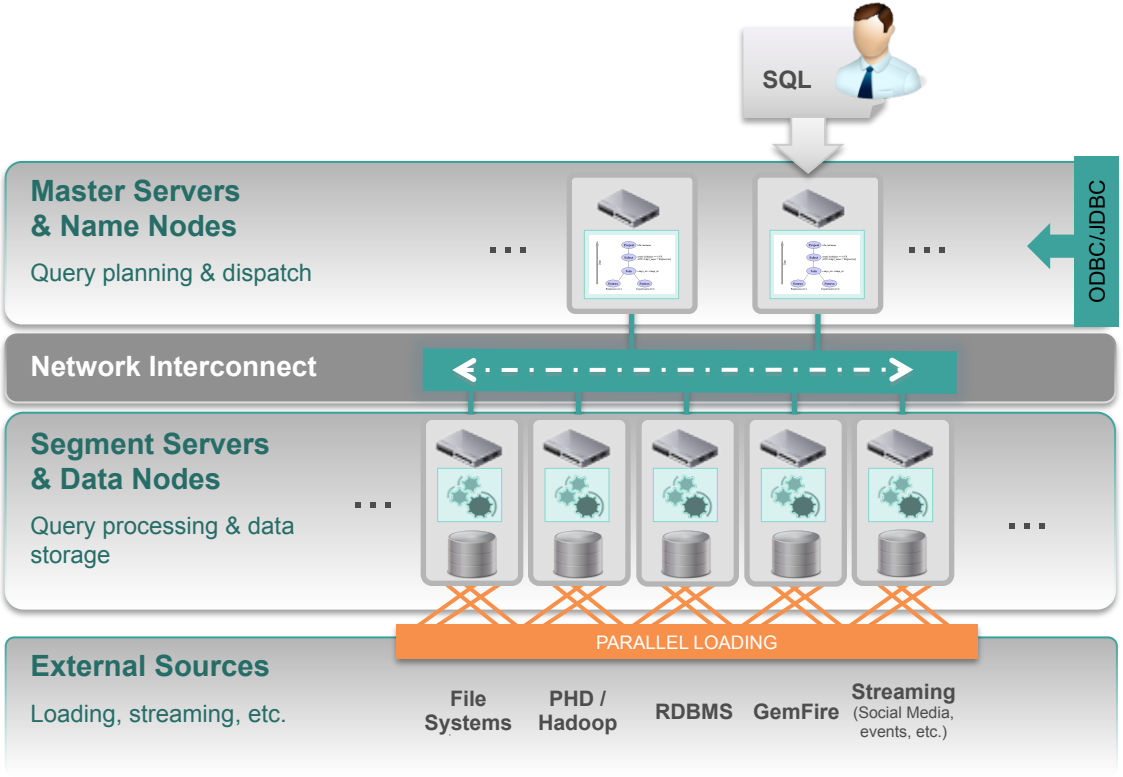

- Parallelizes Everything
  - All nodes can process loading requests
  - No subsequent “Data Reorganization” steps.
  - Scales at over 10+TB/hr. per rack.
  - Only constrained by the speed of the source
- Automates Parallelism
  - GPLoad utility automatically parallelizes file-based loading
  - Integrated with ETL products to parallelize ETL-based loading with minimal added effort

Pivotal

MPP 201: Start Small and Scale as Needed

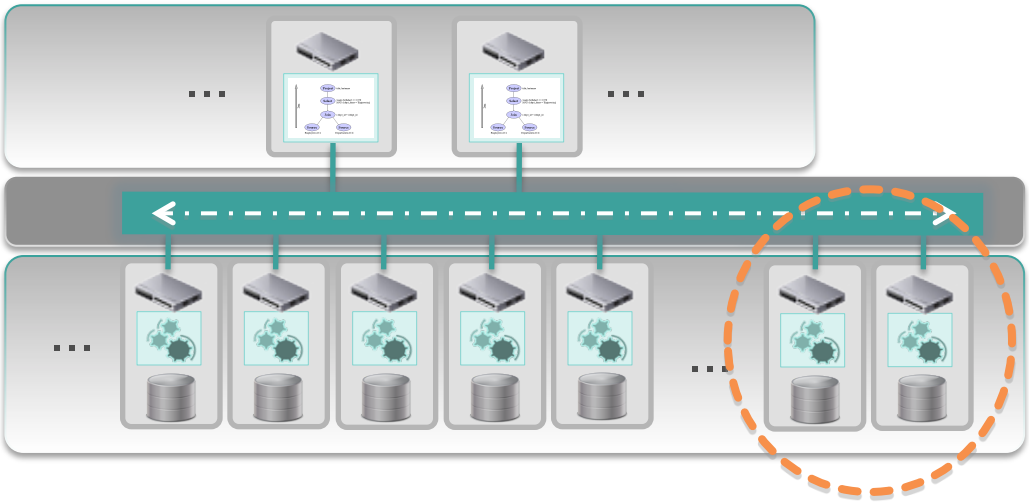

New Segment Servers

Query planning & dispatch

- Advantages:
  - Scale In-Place
  - No Forklifting
  - Immediately Usable
  - Simple Process

Pivotal

Advanced MPP: Polymorphic Storage™

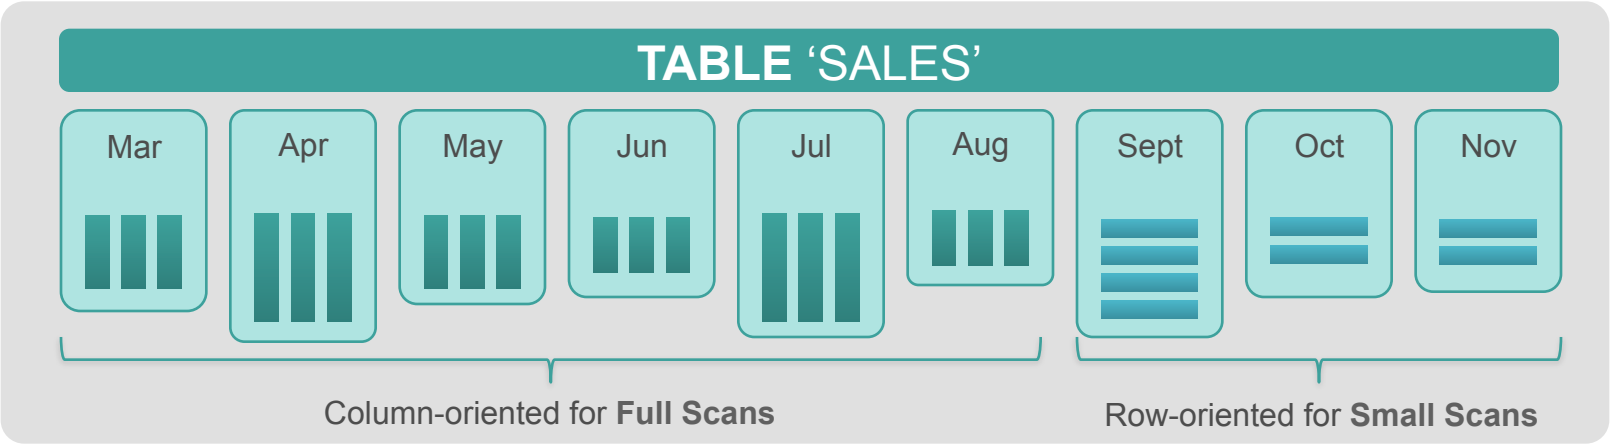

- Columnar storage is well suited to scanning a large percentage of the data
- Row storage excels at small lookups
- Most systems need to do both
- Row and column orientation can be mixed within a table or database

- Both types can be dramatically more efficient with compression
- Compression is definable column by column:
  - Blockwise: Gzip1-9 & QuickLZ
  - Streamwise: Run Length Encoding (RLE) (levels 1-4)
- Flexible indexing, partitioning enable more granular control and enable true ILM

Pivotal

The Pivotal Greenplum Database at a Glance

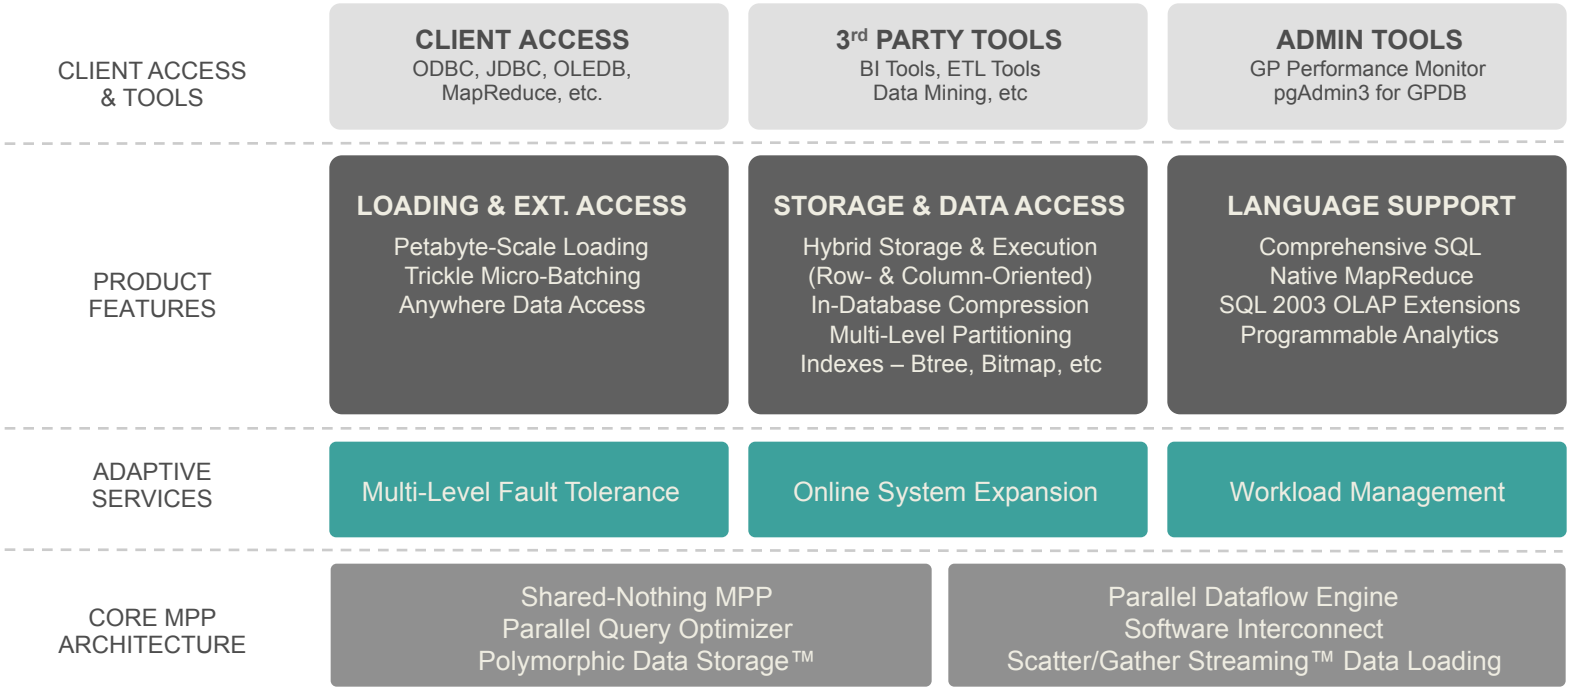

Pivotal

The Pivotal Greenplum Database Overview

- A highly scalable shared-nothing database
- A platform for advanced analytics on any (and all) data
- An enterprise ready platform capable of flexing with your needs

Pivotal

Analytical Architecture Overview

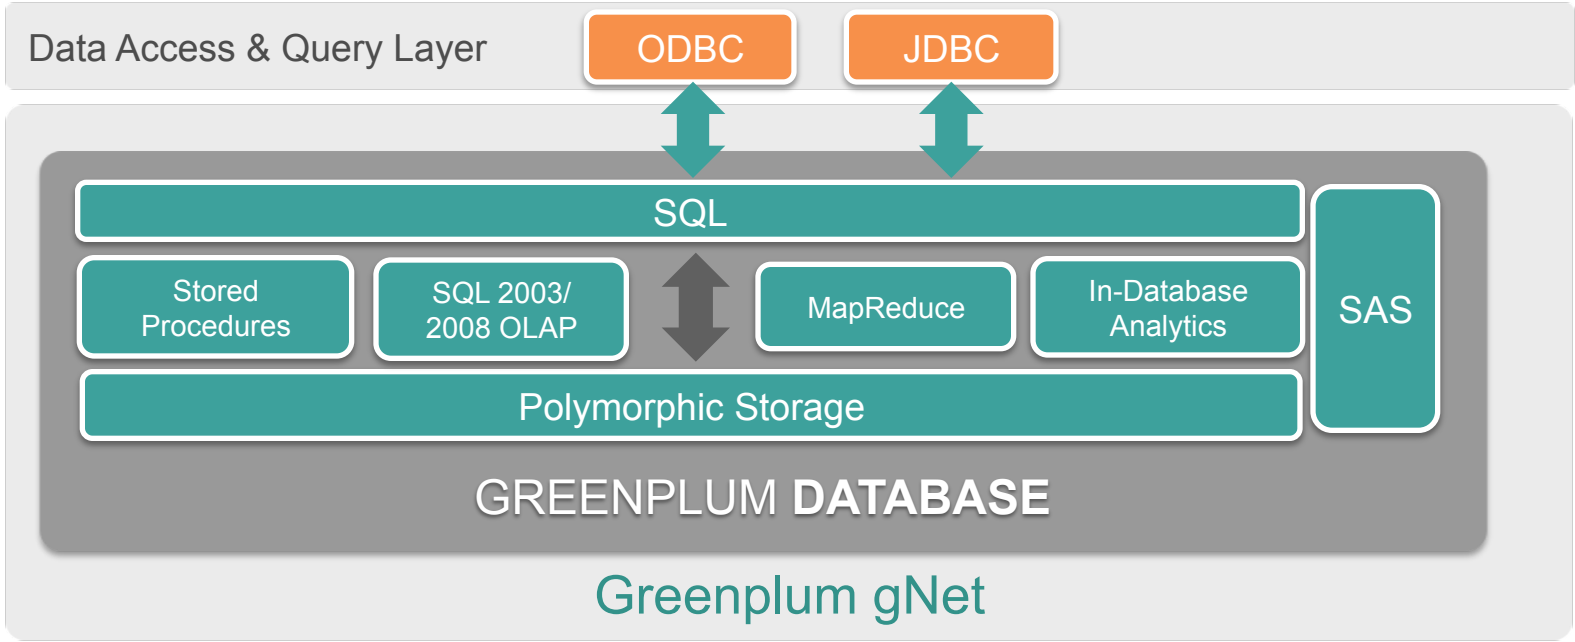

In-Database Analytics: Detail

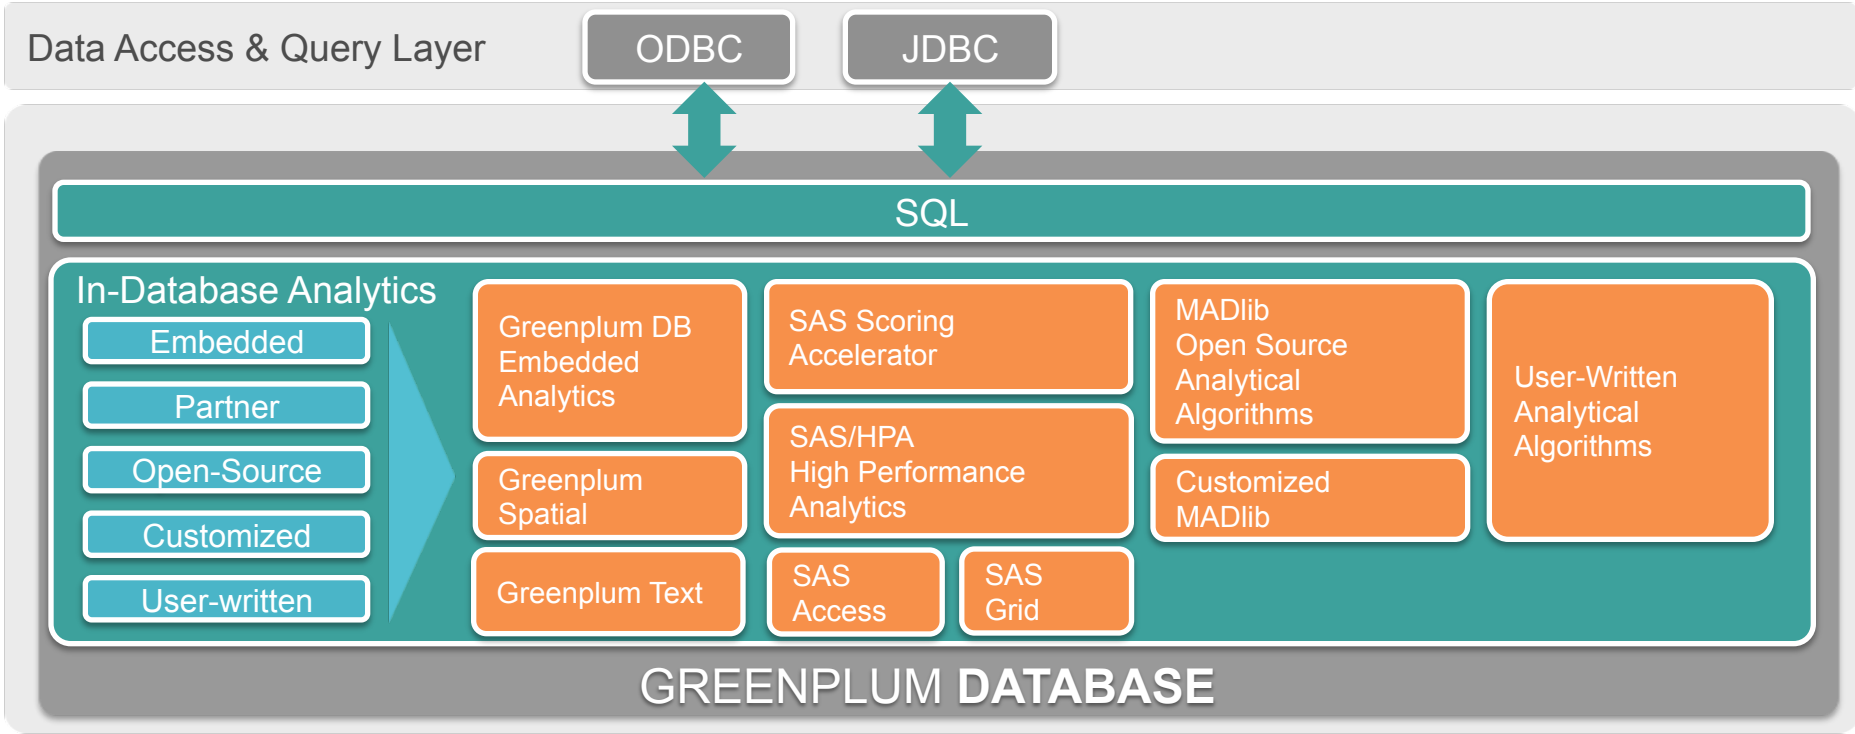

Chorus Analytics Studio

- Create, store, and share visual analytic workflows
- Build analytic flows for Greenplum, HAWQ, and Hadoop
- Powered by Alpine and MADlib
  - 75+ drag-and-drop operators for the entire analytics process
  - MADlib algorithms in-database

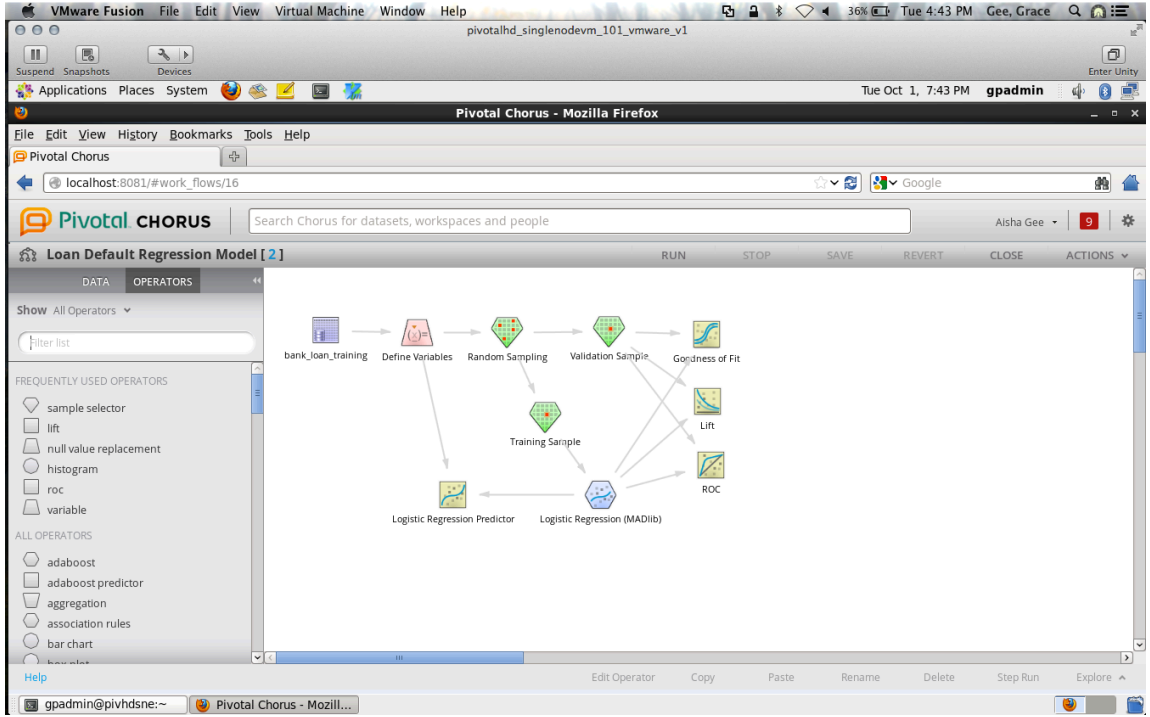

Data & Analytics Technology Ecosystem

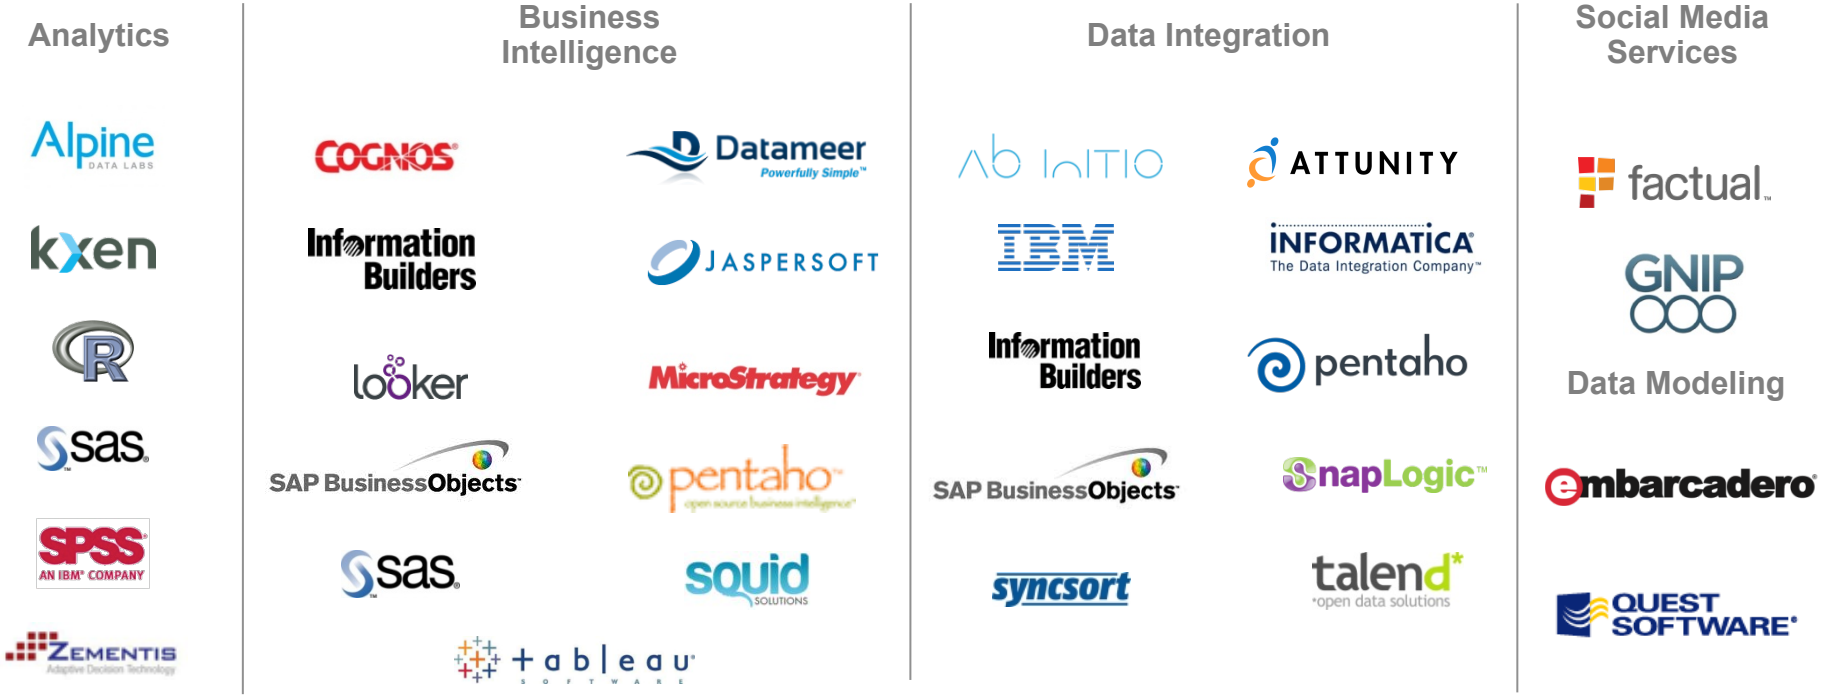

Pivotal

Solving Real World Analytics Problems at Scale

GPDB's performance and scale, combined with built-in analytics and data science expertise, has solved very tough business problems:

Predict Security Prices | RNA Sequencing | Predict Customer Engagement | Network Intrusion Detection

Predict Traffic Patterns | Credit Stress Testing | Predict Churn (1600x Faster)

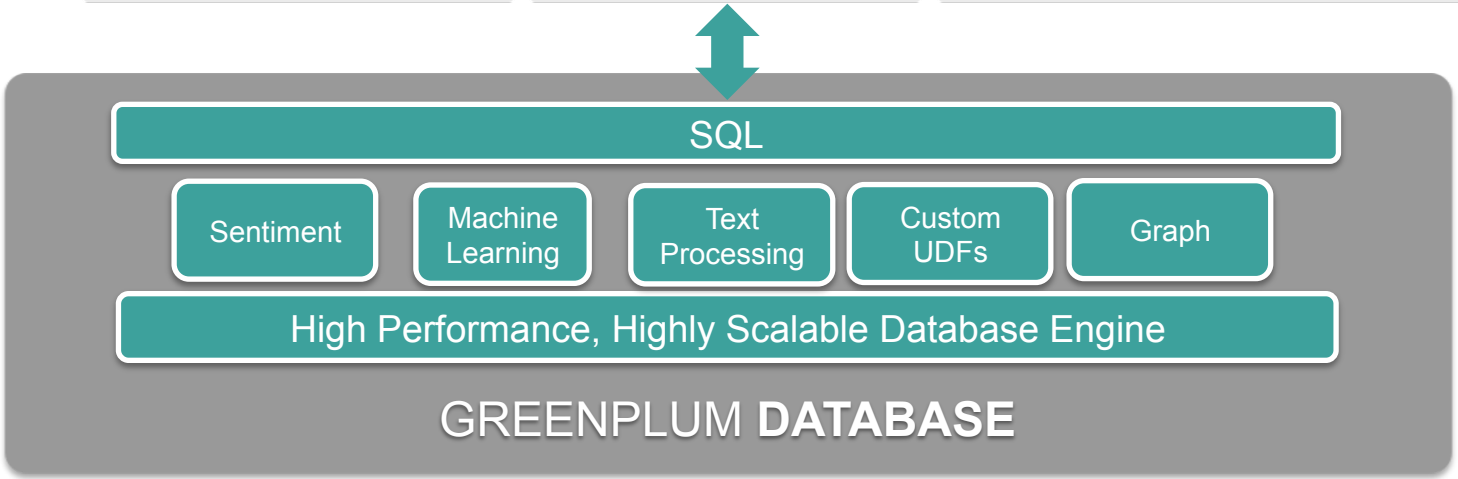

Pivotal

The Pivotal Greenplum Database Overview

- A highly scalable shared-nothing database
- A platform for advanced analytics on any (and all) data
- An enterprise ready platform capable of flexing with your needs

Pivotal

Deployment Choice & Flexibility

The Data Computing Appliance (DCA)

- Modular Flexibility
- Database, Hadoop and ETL Modules
- Future Partner-Specific Modules
- Common Admin and Network Mgmt.
- Incremental Scalability
- Rapid Deployment

PIVOTAL

DATABASE

PIVOTAL HD

Software Only

- Deploy on your x86 hardware
- Certified Configurations
- Perpetual or Subscription License
- Community Editions

Pivotal

The Pivotal Data Computing Appliance: Modular Options

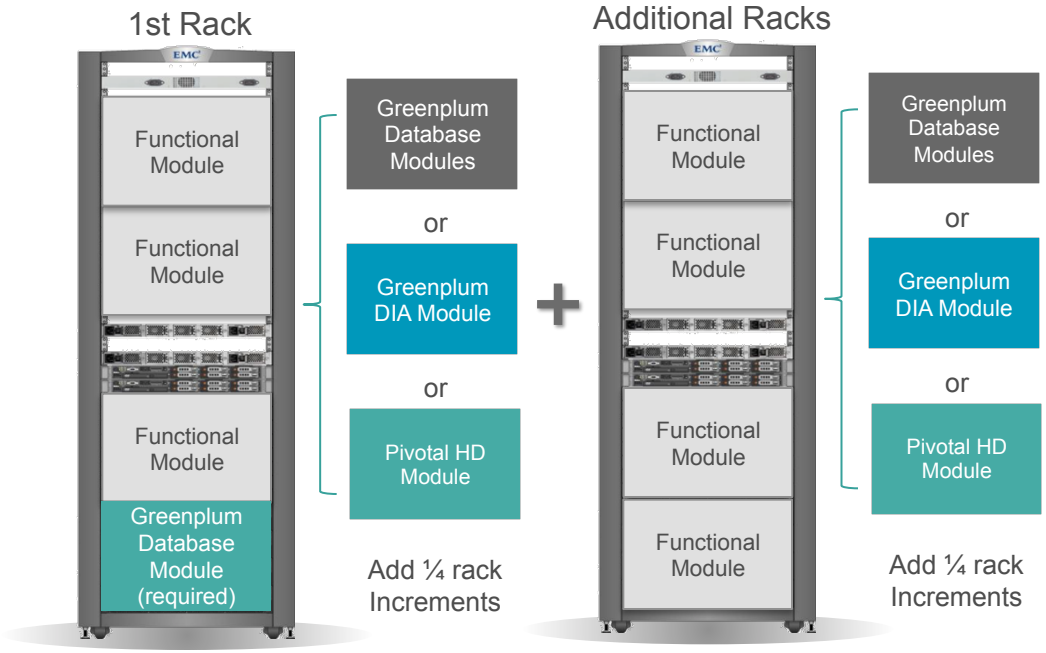

Modules:

- Pivotal Greenplum Database
- Pivotal HD
- Pivotal Data Integration Accelerator
- Multiple module versions available tailored to workload requirements
- ¼ Rack Minimum
- Incremental Scale

Pivotal

High Performance Integration with Hadoop

Parallel Query Access

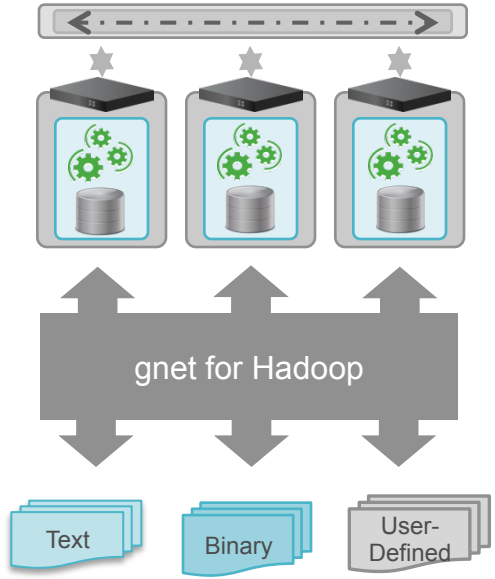

- Connect any data set in Hadoop to GPDB SQL Engine
- Process Hadoop data in place
- Parallelize movement of data to/from Hadoop thanks to GPDB market leading data sharing performance
- Supported formats:
  - Text (compressed and uncompressed)
  - binary
  - proprietary/user-defined
- Support for Pivotal HD, MapR, Hortonworks, Cloudera

Pivotal

Comprehensive High Availability

- Master and Segment Mirroring with block level replication
  - Low resource consumption
  - Differential resynch capable for fast recovery
  - Minimize interdependencies!
- Segment servers support multiple database instances
  - Primary instances that actively process queries
  - Standby mirror instances

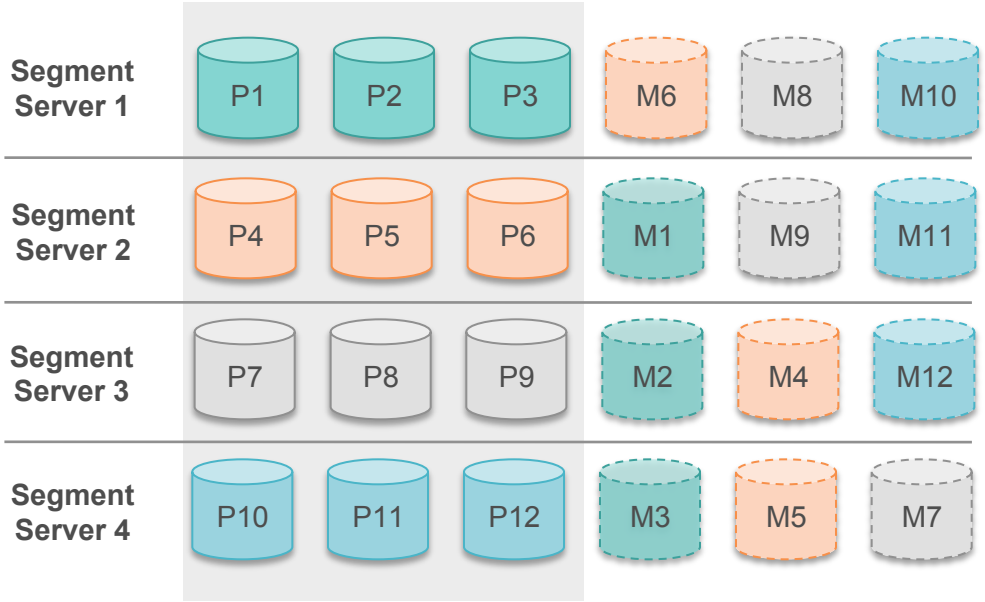

Set of Active Segment Instances

Pivotal

Comprehensive Backup/Restore

- Full and Incremental backup support with in-database tools
- Incremental backup
  - Only changed partitions are pulled for the backup
  - Restore to any point-in-time through support of “synthetic restores”
  - Synthetic restores automatically assemble the right backup based on the point-in-time specified: manual backup specification is not required
- Deep support for Data Domain
  - WAN replication of backup sets to remote DR sites
  - Granular delta-only backup support

Pivotal

Comprehensive Management Tools

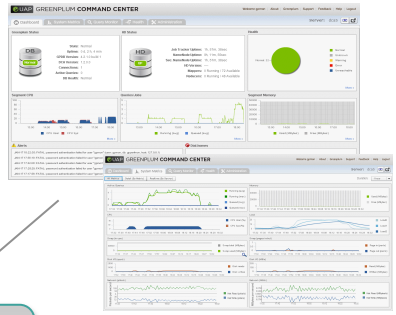

Single console for software plus DCA

Supports easy deployment of database extensions

Manage the Database with Command Center

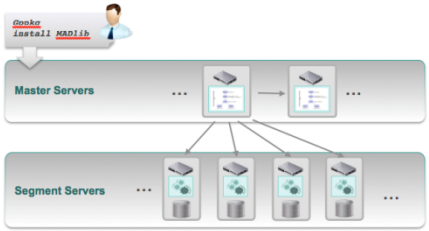

Manage Workloads with PivotalVRP

Manage Analytics & Extensions with gpPkg

Best in class workload monitoring, management and performance tuning

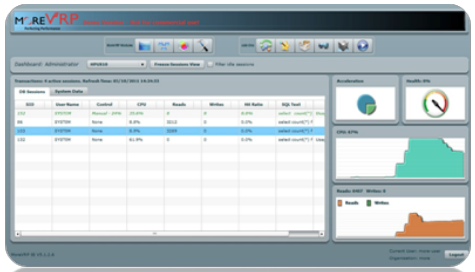

Pivotal

Demo

Data Loading Methods

Pivotal

Data Load Options

- SQL INSERT
  - Standard Row by row insert – slowest method
    - INSERT into tableX VALUES ('John', 'Doe', 'Manager')
  - All data is passed through Master server
- PostgreSQL Copy command
  - Inserts data from a file or stdin (another query) – faster than SQL INSERT
    - COPY tableX FROM {file | STDIN}
  - All data is passed through Master server
- Parallel loading with gpfdist/gpload
  - Segment servers connect directly to external files served via gpfdist
  - Load bypasses Master server
  - Segment servers load in parallel
  - External tables point to the streamed files
    - CREATE EXTERNAL TABLE ext_table LOCATION (gpfdist://dir/*)
    - CREATE TABLE tableY AS SELECT * FROM ext_table
  - Integrated with Informatica PowerExchange and Pentaho

Pivotal

SQL Insert Method

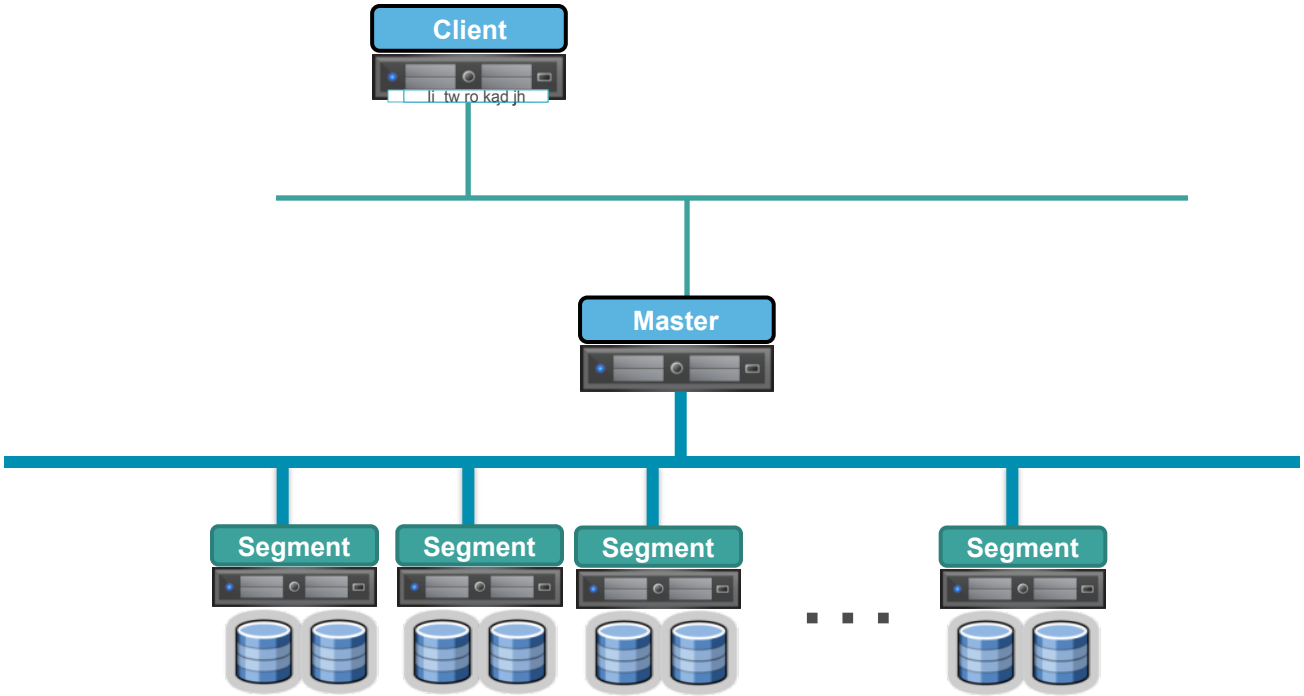

PostgreSQL Copy Method

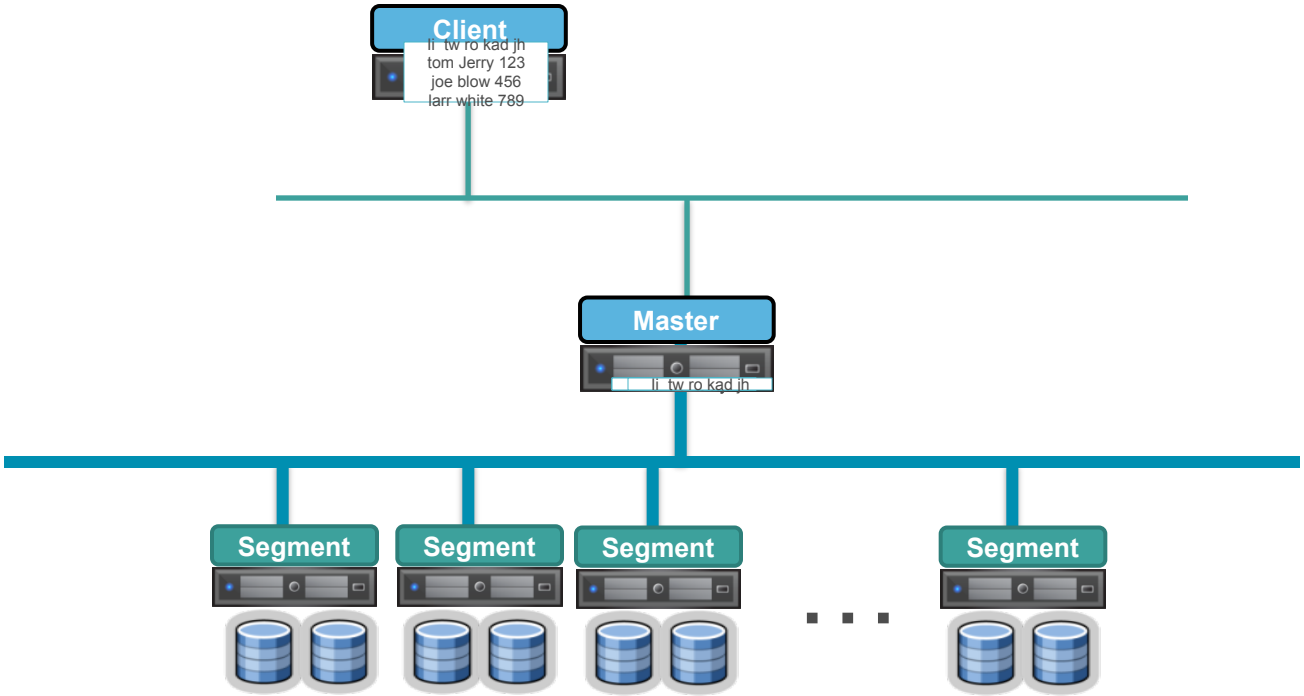

Pivotal

Parallel Load with gpfdist

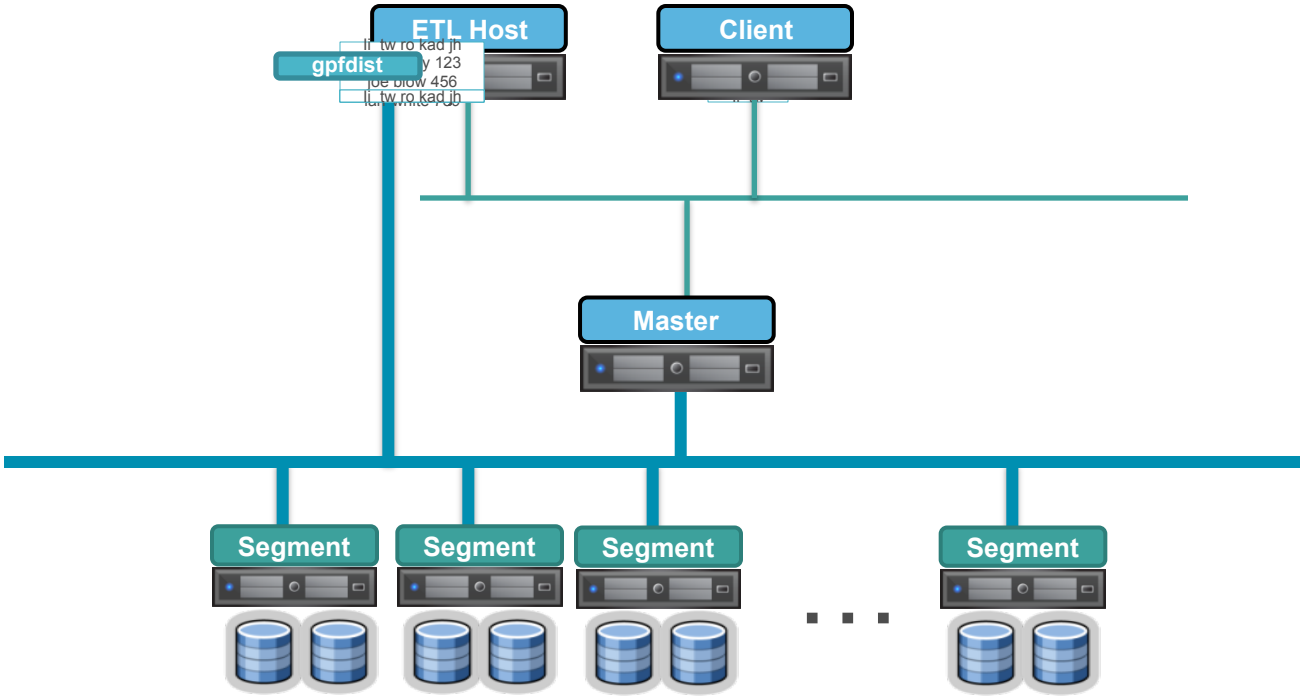

Parallel Load with gpfdist

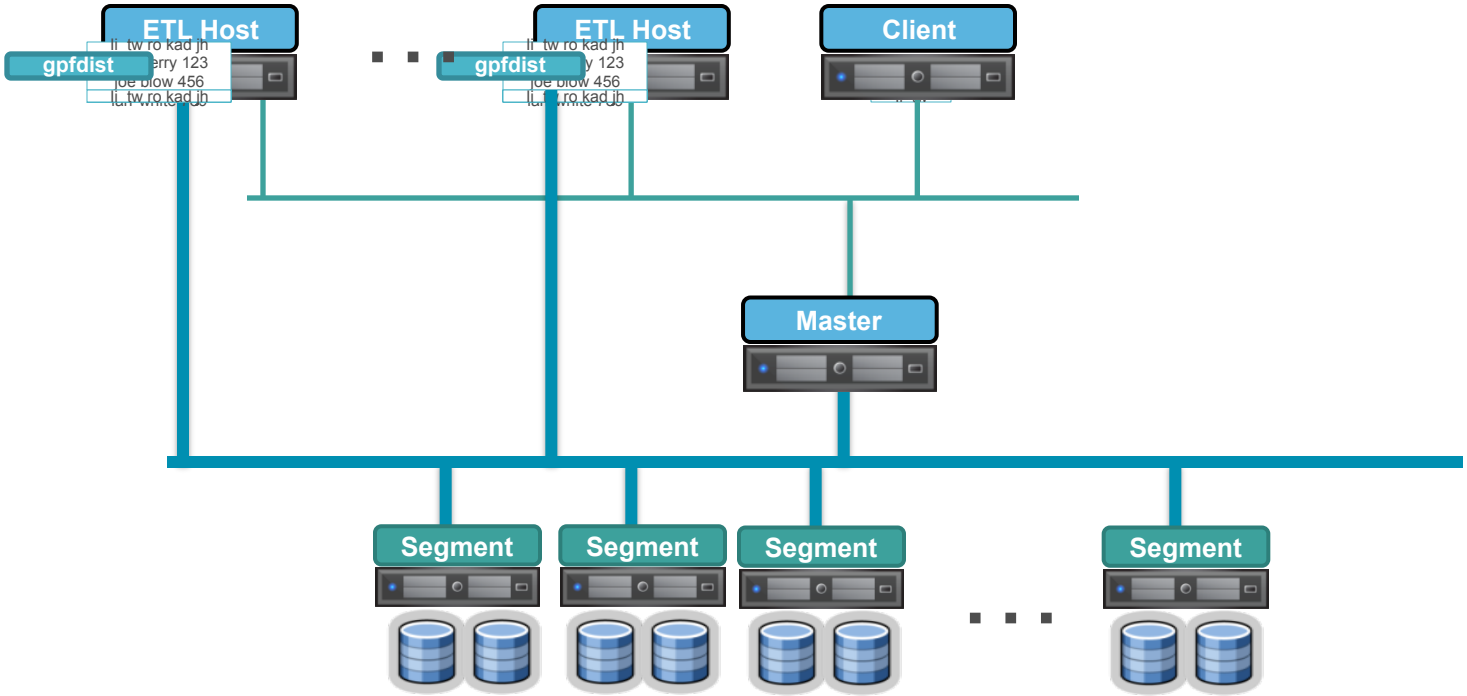

Parallel Load with Informatica, Pentaho, etc.

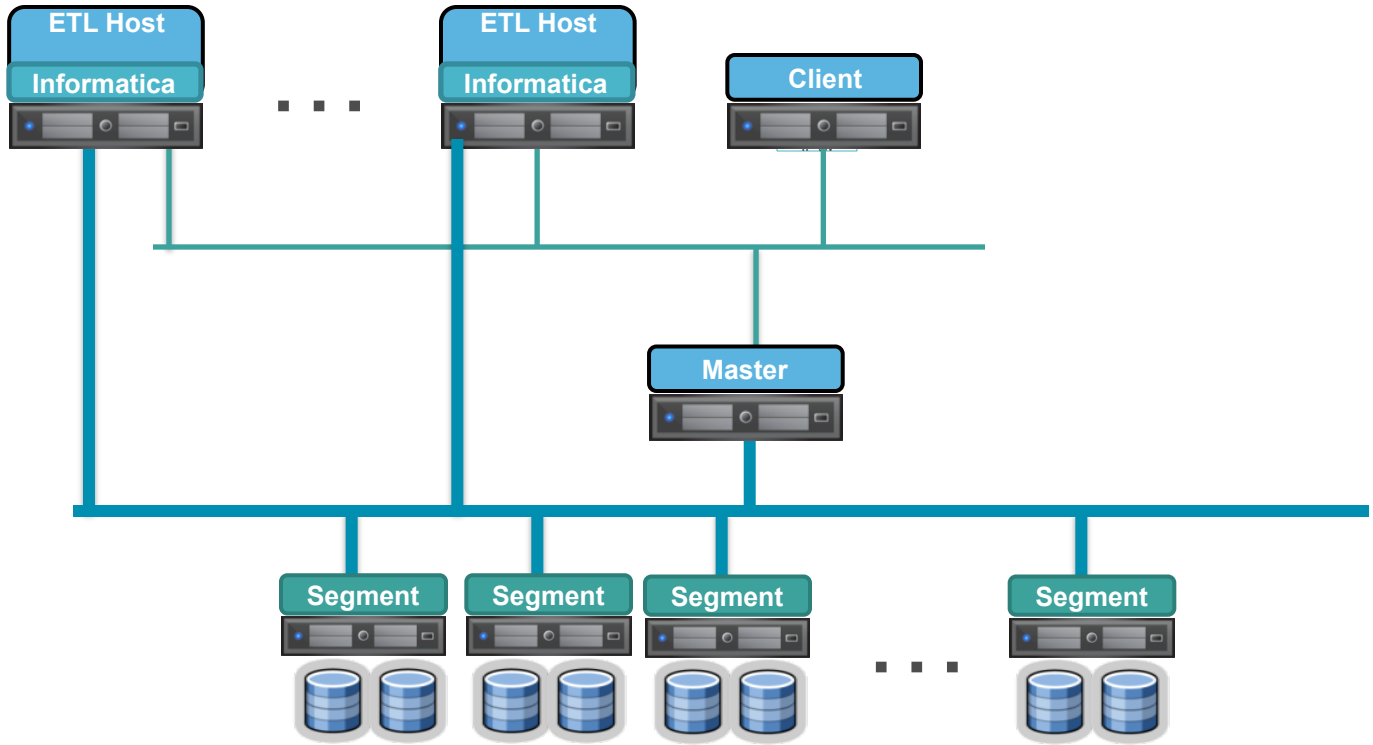

Pivotal

Data Distribution and Partitioning

Pivotal

Table and Data Distribution

- Technique for balancing workload across all nodes
  - All database tables are distributed across all Segments
  - Distribution Policy determines how the data in the tables is distributed
- Hash Distribution (Default)
  - One or more columns used as distribution key
    - Hashed into “Buckets” for each Segment
  - Unique keys assure even distribution
  - When key is unspecified - uses Primary key or first column
- Random Distribution
  - Round-Robin distribution of rows to Segments
  - Evenly distributes data but may be less efficient at query time

Pivotal

Distribution

```
Create Table tableX …
Distributed by
CustomerID;
```

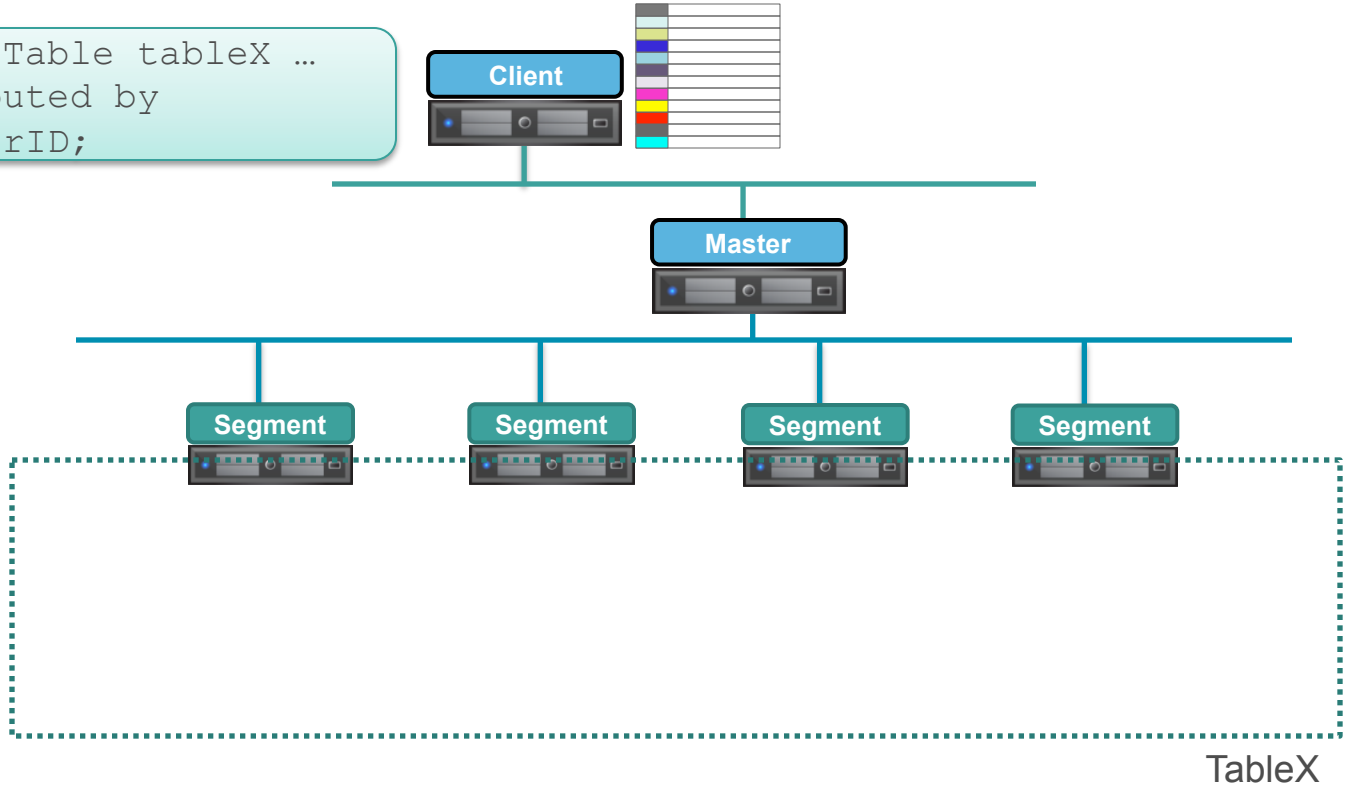

Distribution

```
Create Table tableX …
Distributed by
CustomerID;
```

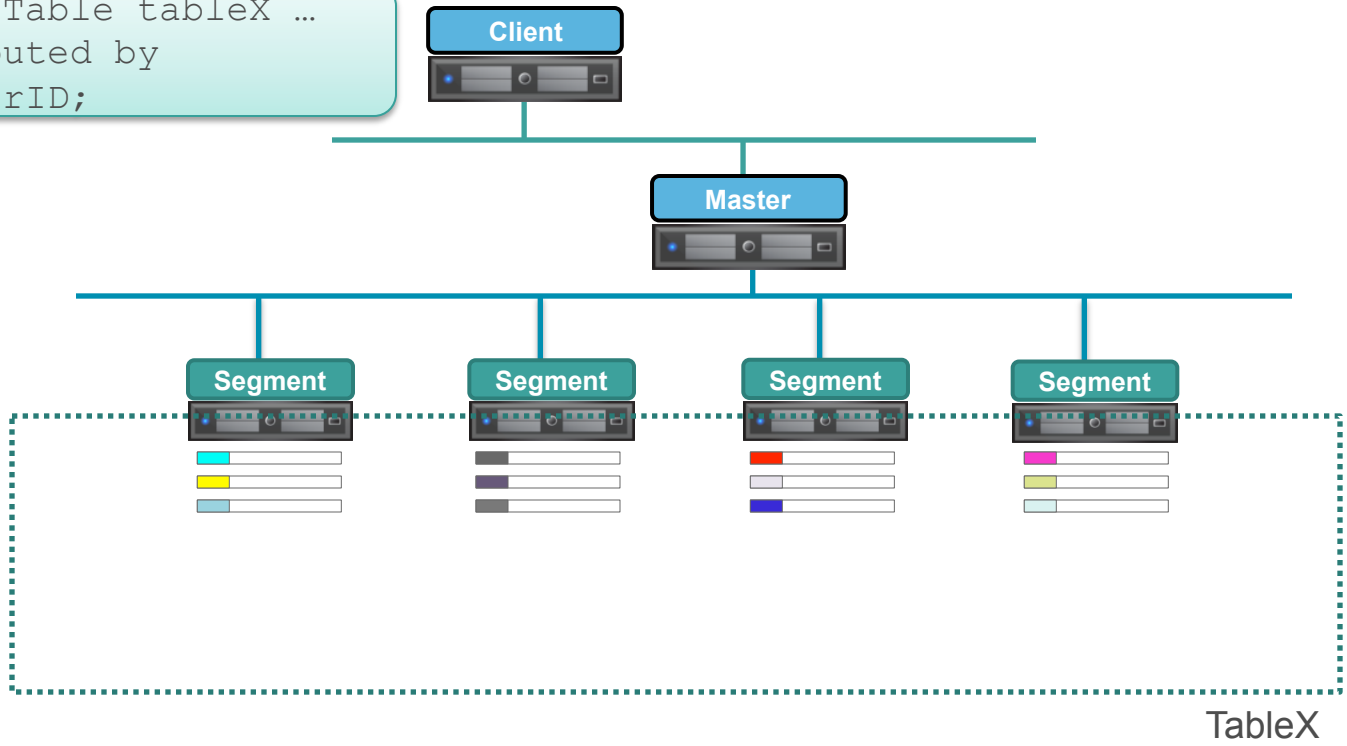

Distribution

```
Create Table tableX …
Distributed by Gender;
```

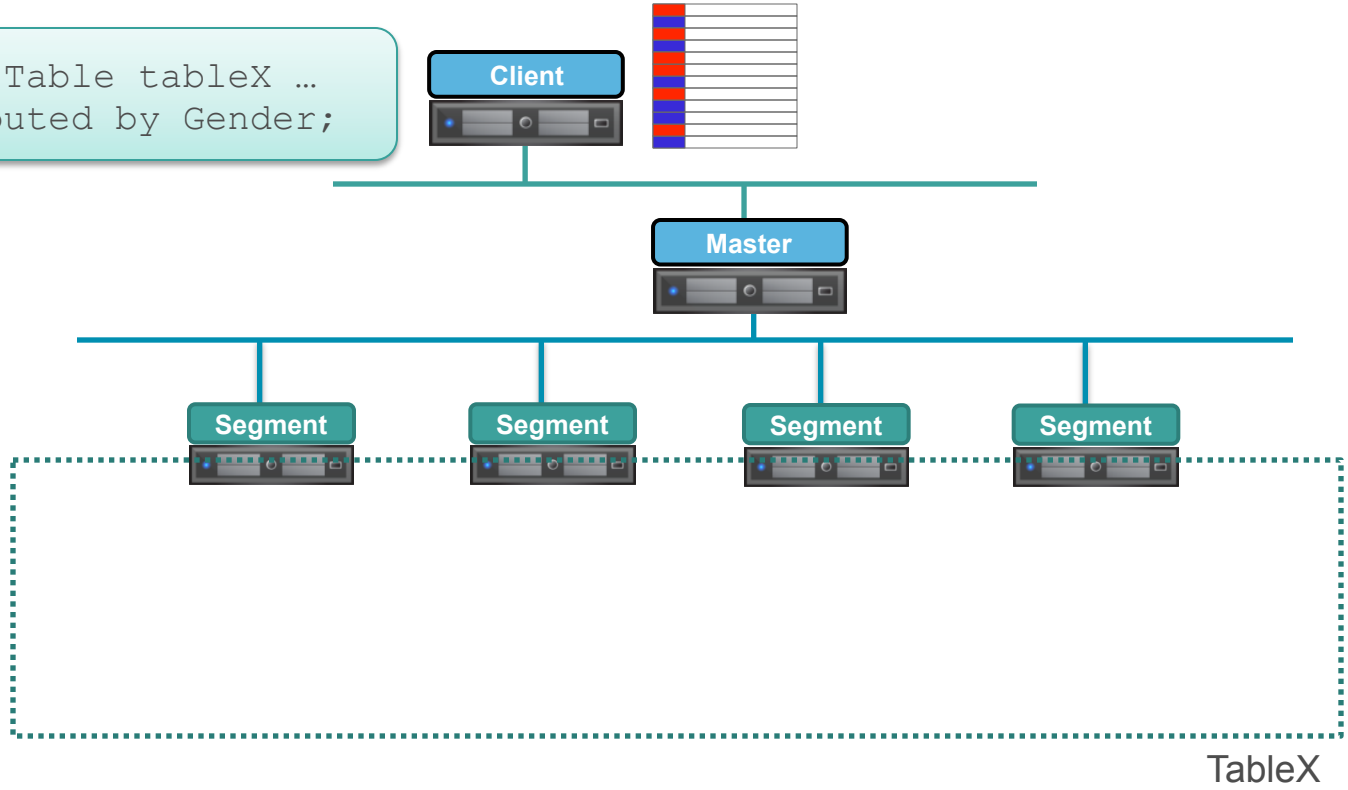

Distribution

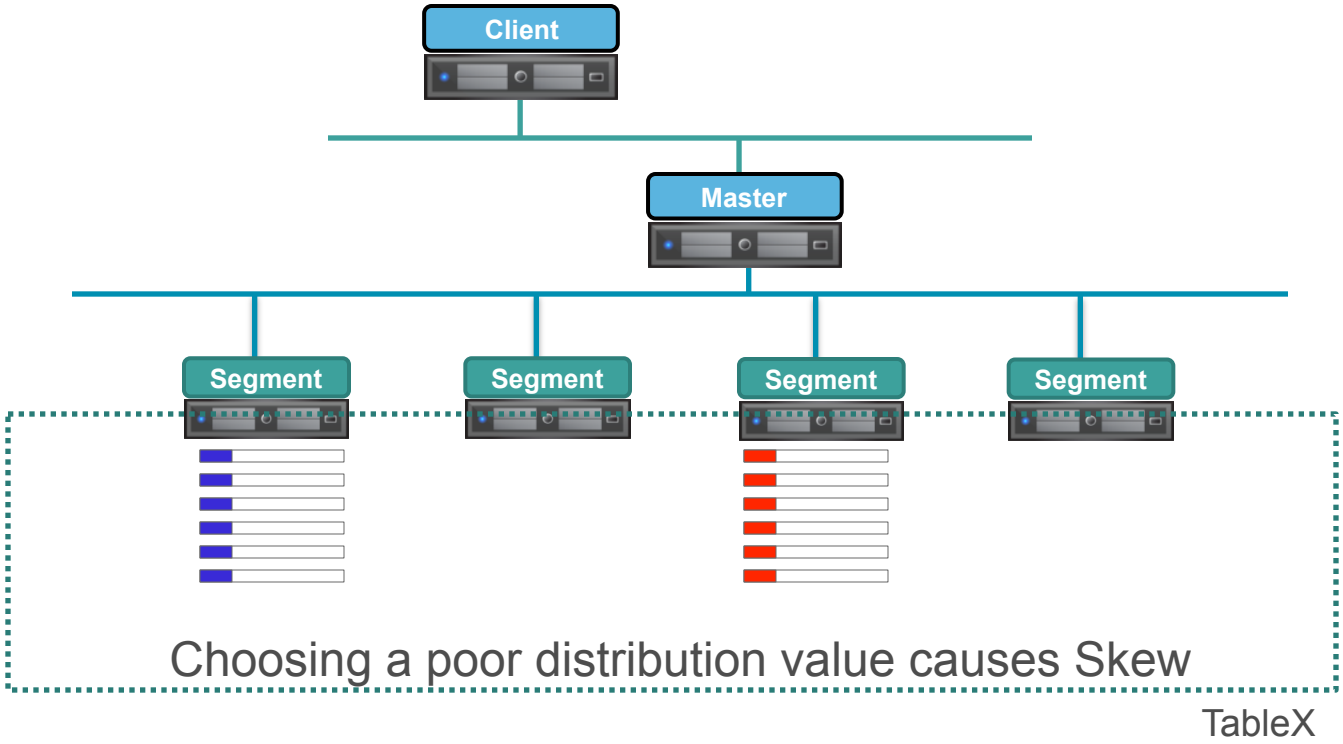

Table Partitioning

- Technique for eliminating rows at query time
- Logically divides a large table into smaller parts
- Significantly improves query performance
- Facilitates database maintenance
- Two types
  - Range partitioning
  - List partitioning
- Partitioning in Greenplum works using table inheritance and CHECK table constraints
- Partitioned tables also have distribution keys

Pivotal

Partitions

Single Partition

```
Create Table tableX …
Distributed by
CustomerID;
```

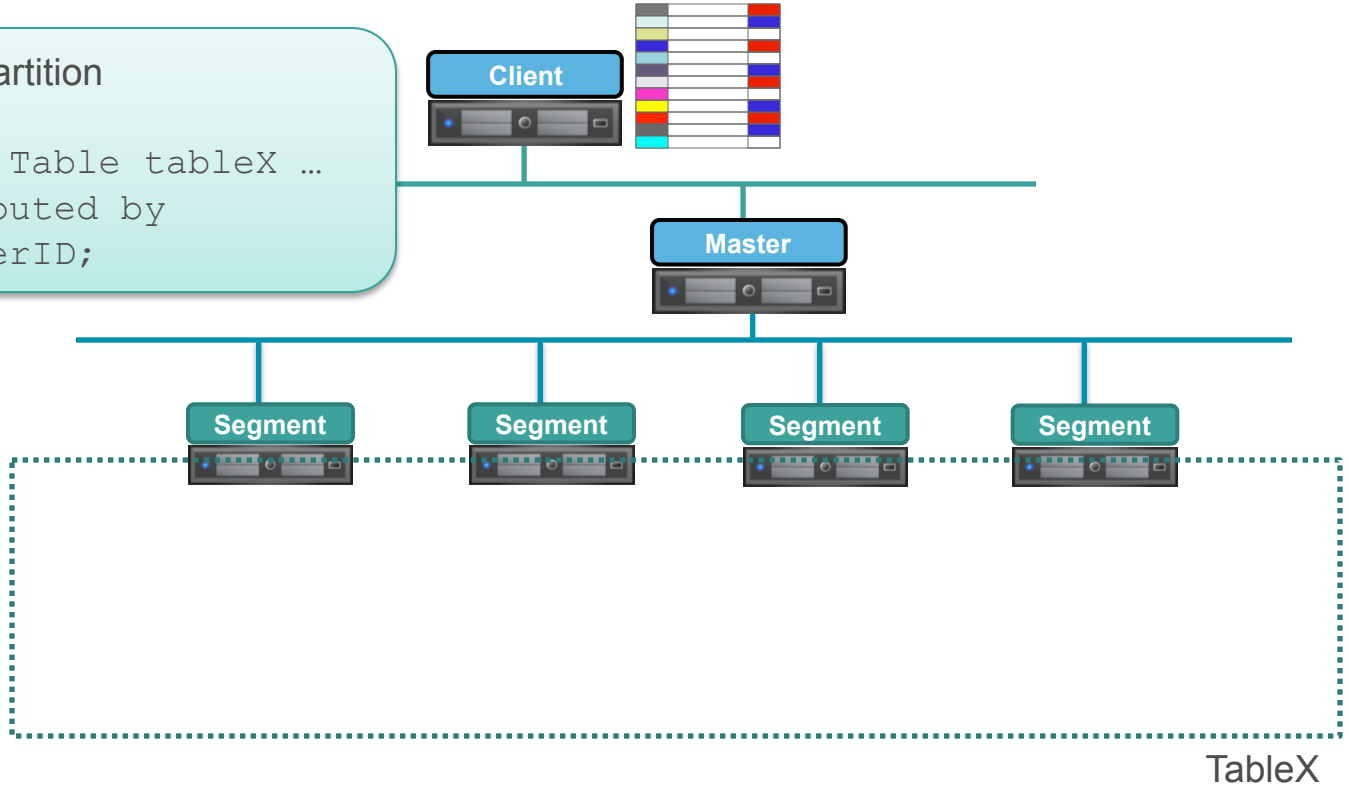

Partitions

Single Partition

```
Create Table tableX …
Distributed by
CustomerID;
```

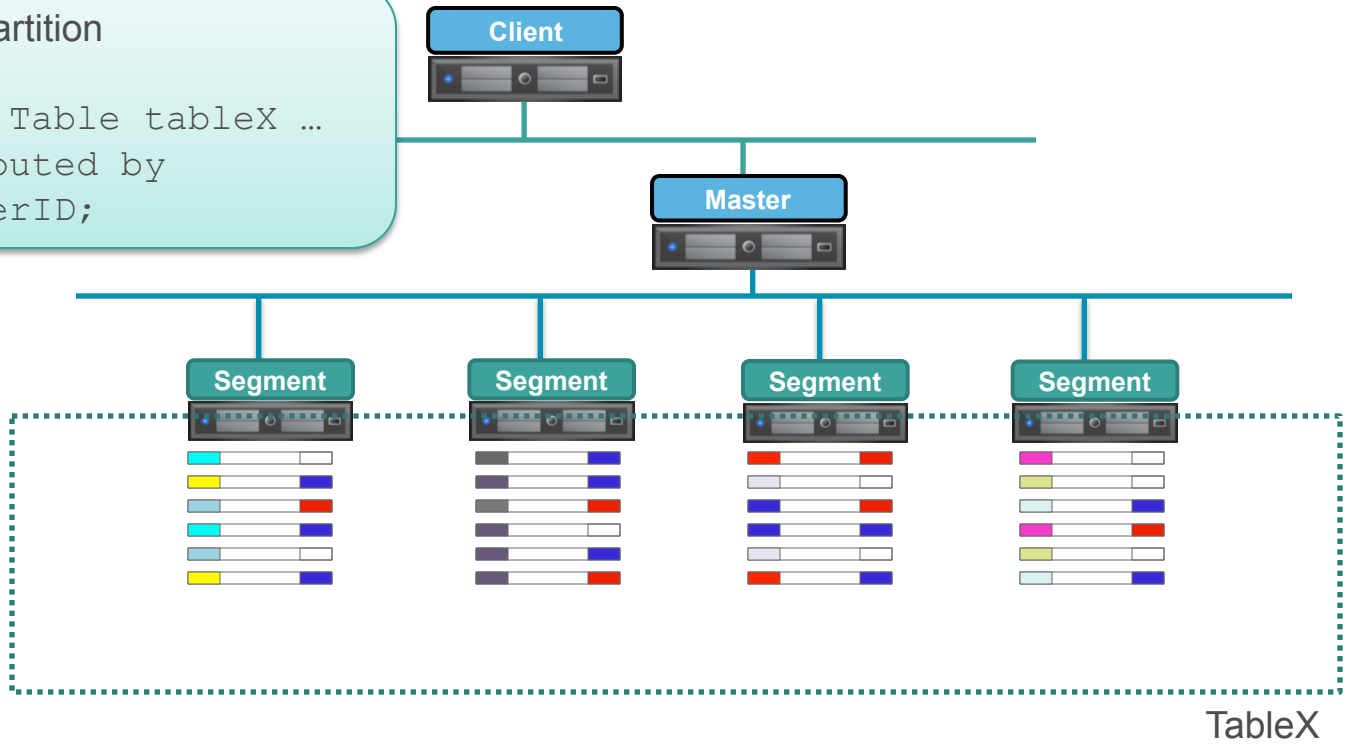

Partitions

Query

```
Select * from tableX …
Where Color = 'Blue';
```

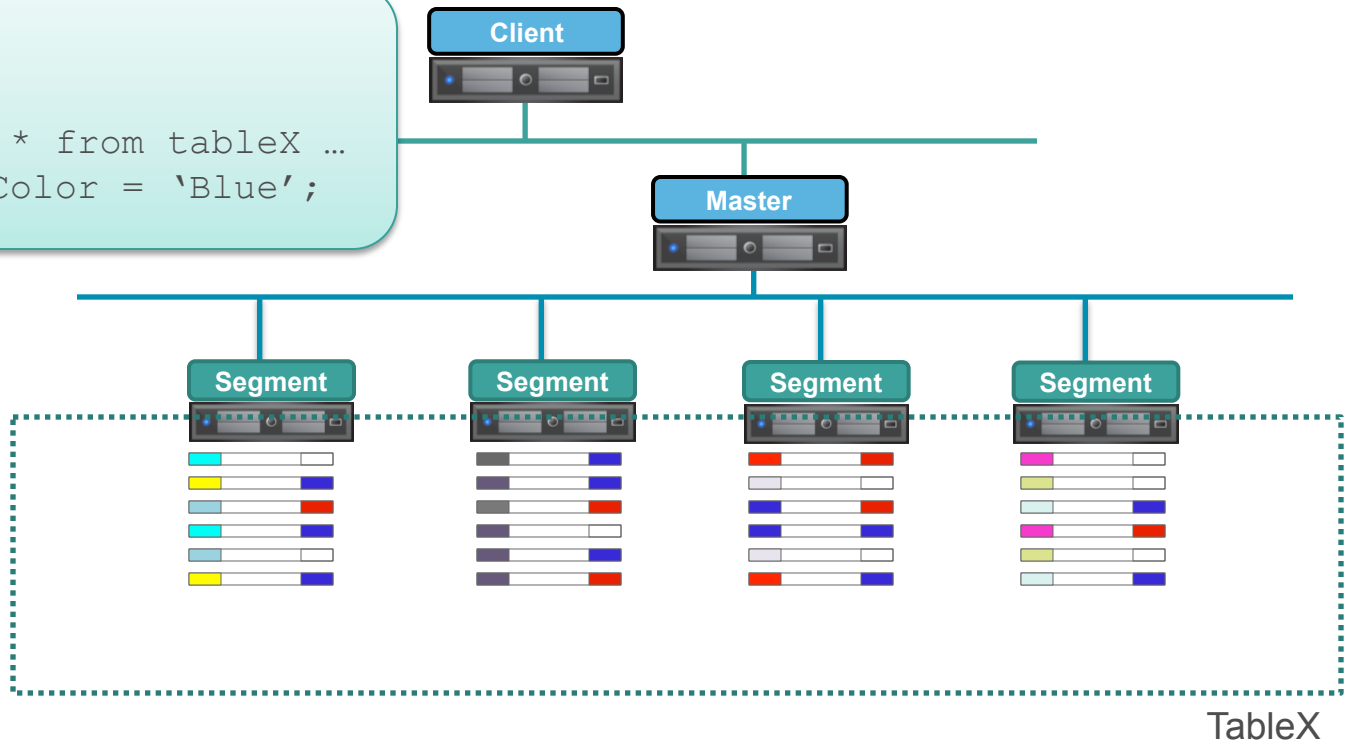

Partitions

Multiple Partitions

```
Create Table tableX …
Distributed by
CustomerID Partition
by Color;
```

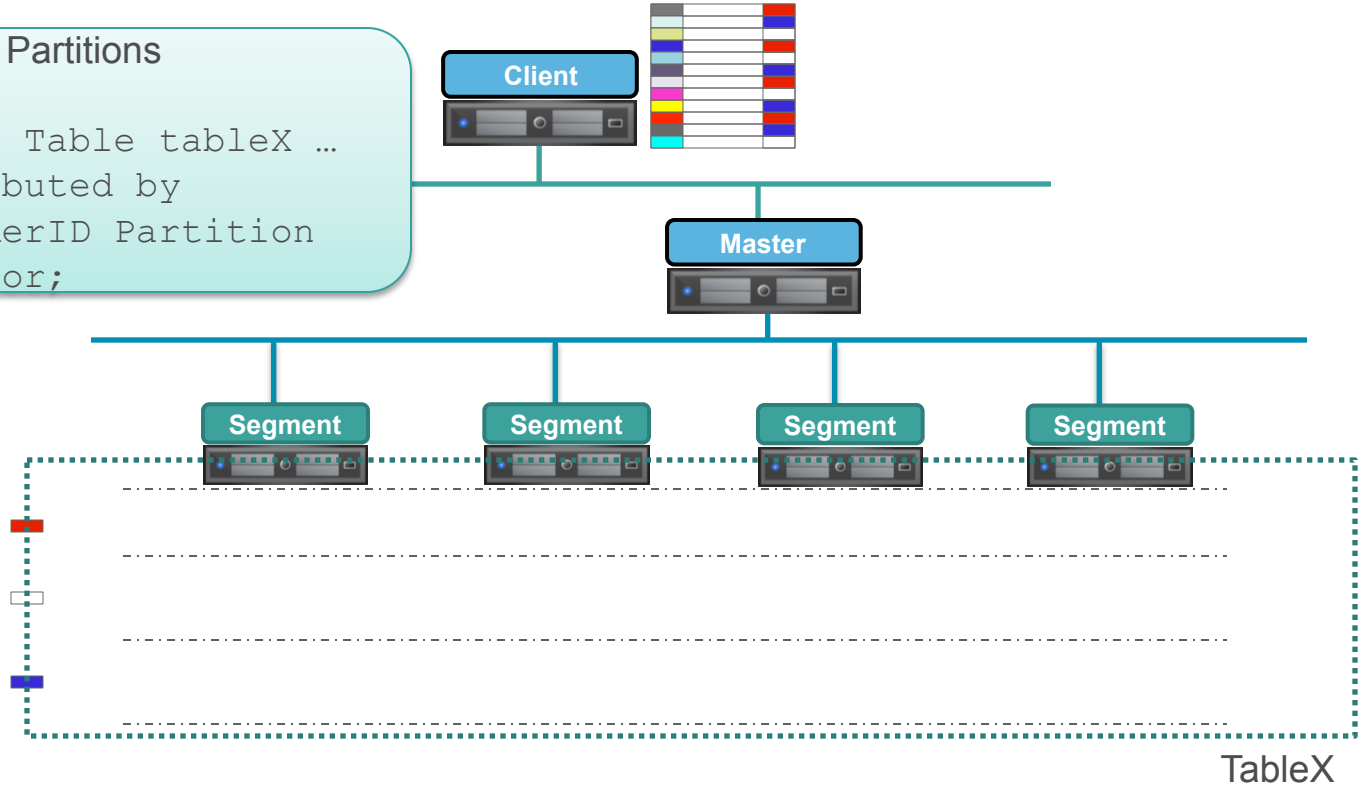

Partitions

Multiple Partitions

```
Create Table tableX …
Distributed by
CustomerID Partition
by Color;
```

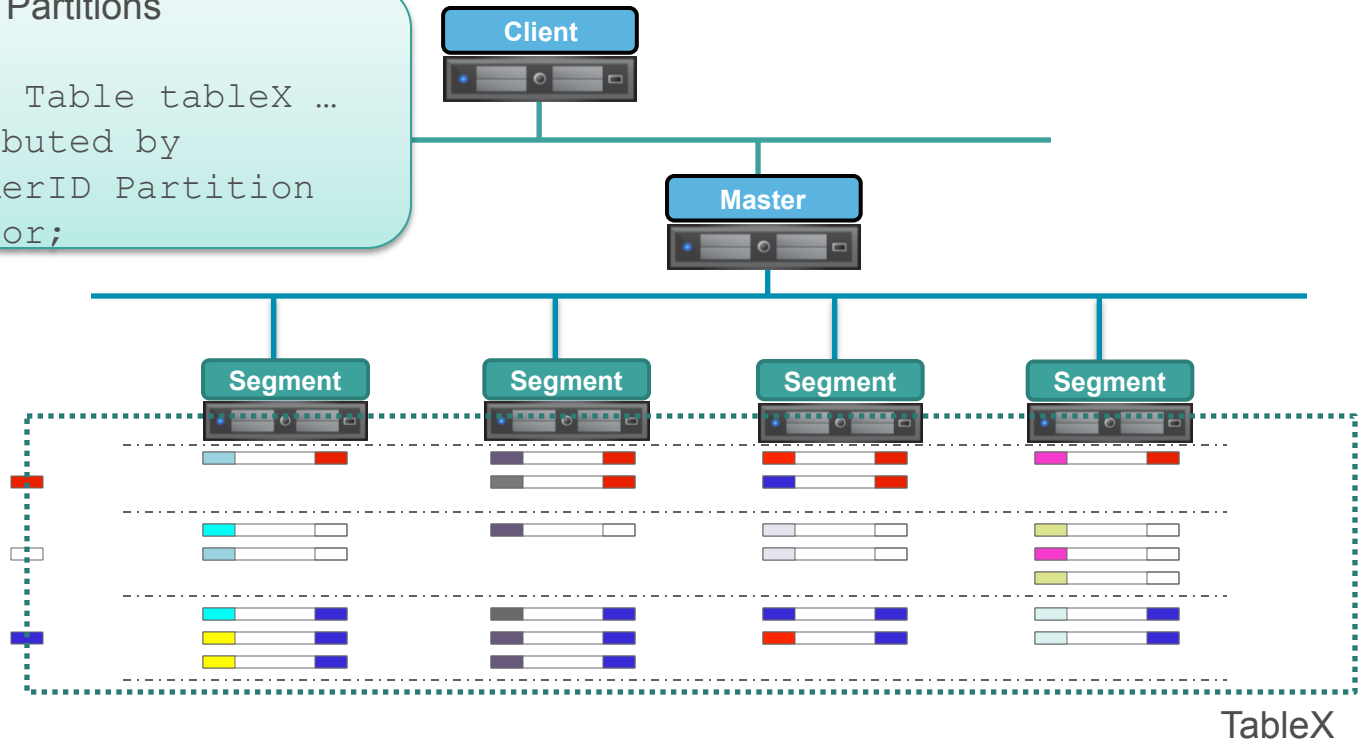

Pivotal

Partitions

Query

```
Select * from tableX …
Where Color = 'Blue';
```

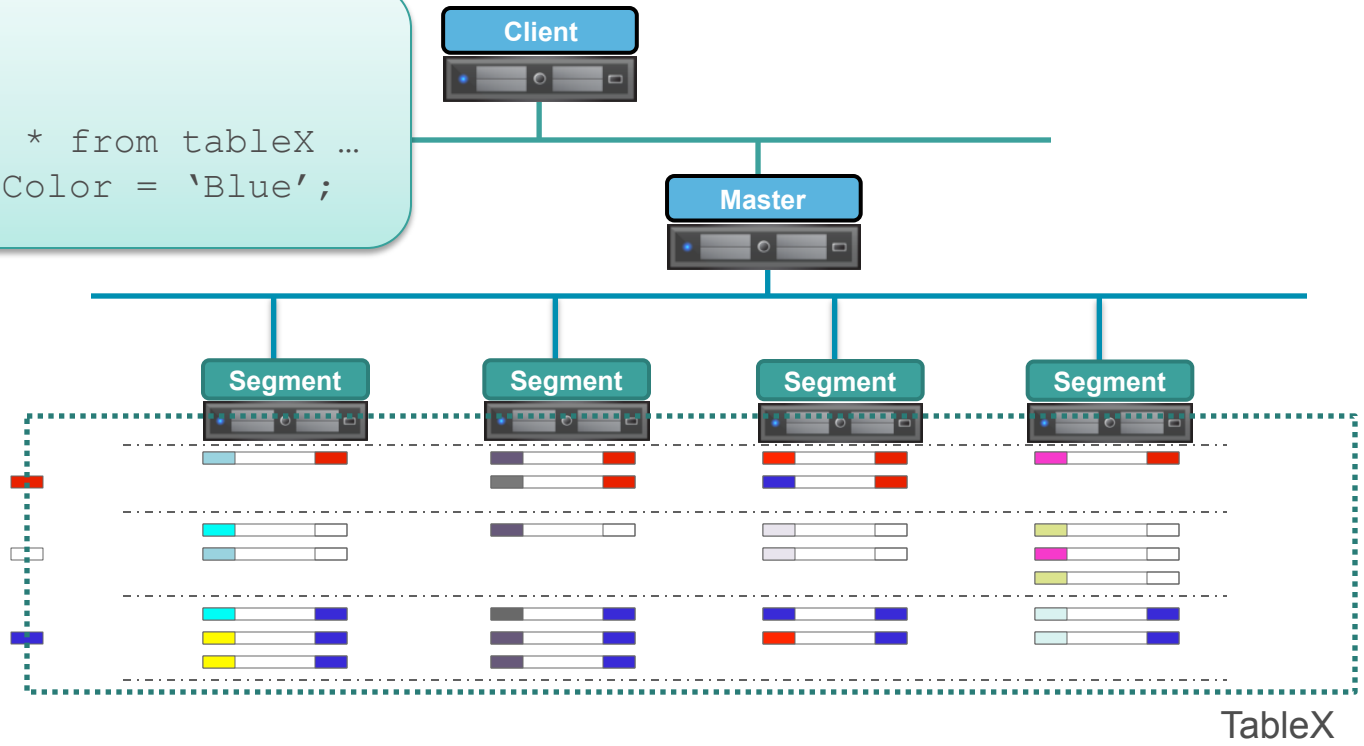

When To Partition?

- Have a large (fact) table
- Query predicates have identifiable access patterns
  - e.g., WHERE trans_date >= ’07/01/2012’) AND trans_date < ’08/01/2012’
- Need to maintain a “rolling window” of data
- Data can be divided into somewhat equal parts based on some defining criteria
- Experiencing unsatisfactory performance

Pivotal

Backup and Recovery

Pivotal

Backup in a nutshell

- Option 1: custom external tables
  - Good control over which tables/data to backup
  - Enables incremental backup
  - Doesn't include other objects, such as roles, resource queues, etc.
- Option 2: pgdump
  - Free utility
  - Not parallelized
  - Creates one dump file on the master
- Option 3: gpcrondump
  - Free utility
  - Parallelized backup
  - Creates SQL files on the master and segment hosts
  - Must restore to same number of hosts/segments
  - Incremental backup not supported
- Option 4: Data Domain
  - See next slide

Pivotal

Efficient Backup/Restore with EMC Data Domain

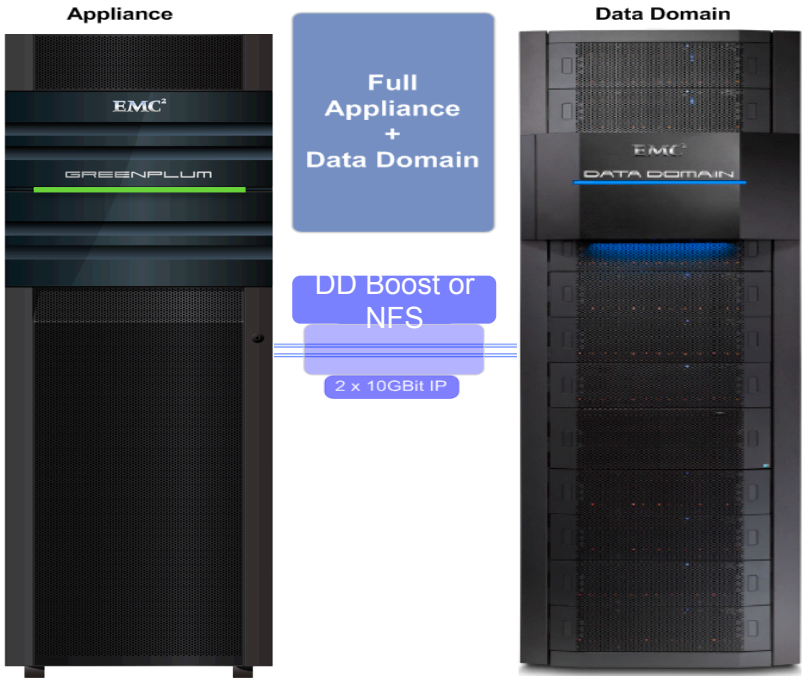

- Back up protocol options
  - NFS: Data Domain device mounted as NFS storage. Direct IO option available in 4.1.2.3 or 4.2.2 or higher
  - DD Boost: Native, client-side deduplication, dedicated communication. Requires 4.2.1 or higher
- Data Domain deduplication is an ideal fit for Greenplum
  - Integrates seamlessly into standard Greenplum full backup data export and data restore procedures
- Drastic reduction in backup storage requirement
- Backup all segment servers in parallel directly to Data Domain
- Backup/Restore at Full or Table level

Pivotal

High Availability

Pivotal

Greenplum Segments

- A database instance
- Works in parallel with other Segments (instances) to process SQL (loads and queries)

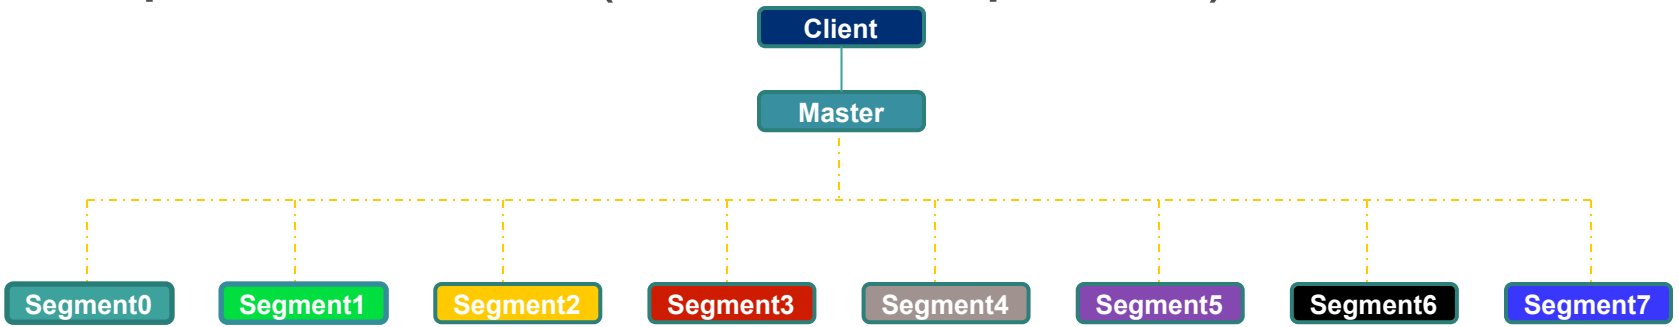

Pivotal

Segment Hosts

- Physical servers that holds some number of Segments (instances)
- Dedicated CPU and storage that is not shared with other Hosts
- High speed interconnect between Segment Hosts

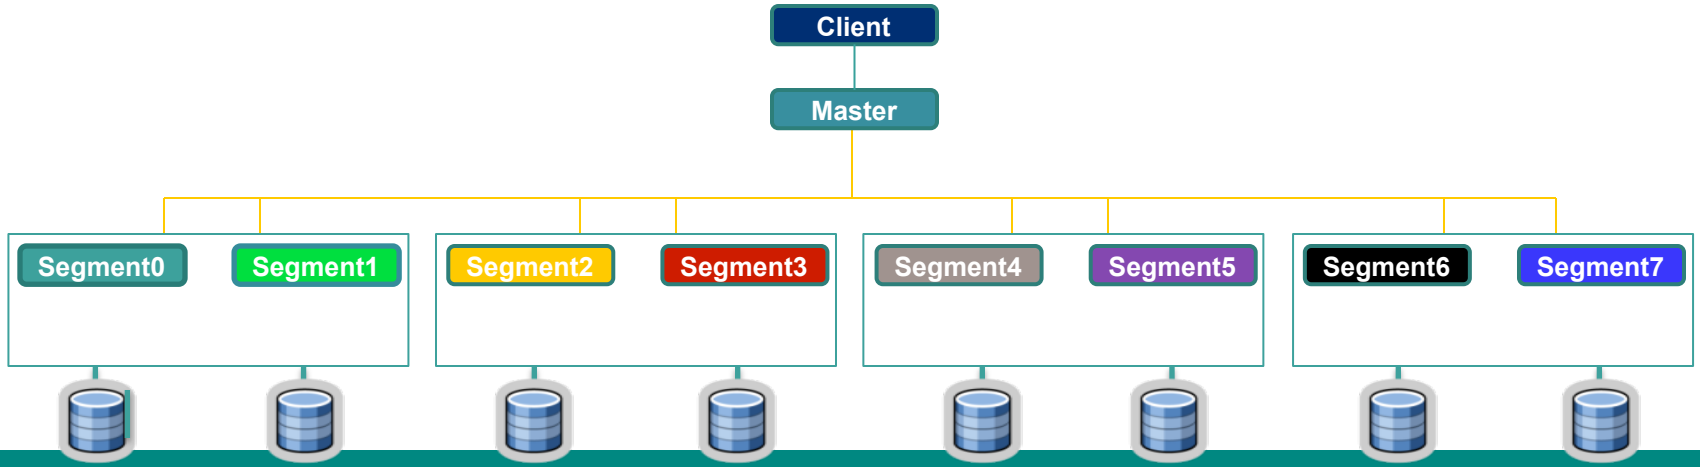

Pivotal

Mirroring

- Protects Segment Instances from host failures
- Mirror Segments
  - Warm standby for Primary Segment
  - Continuously updated with data from Primary
  - Spread across Segment hosts

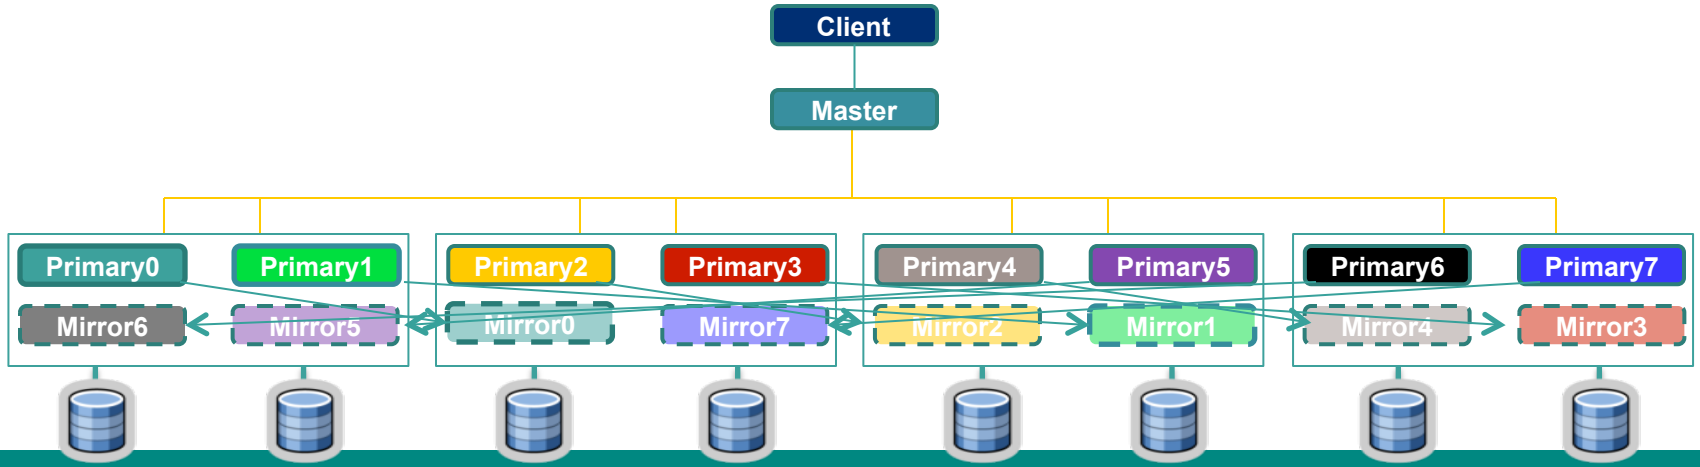

Pivotal

Failover

- Master continuously monitors segments
  - If no response, initiates failover to Mirror Segments

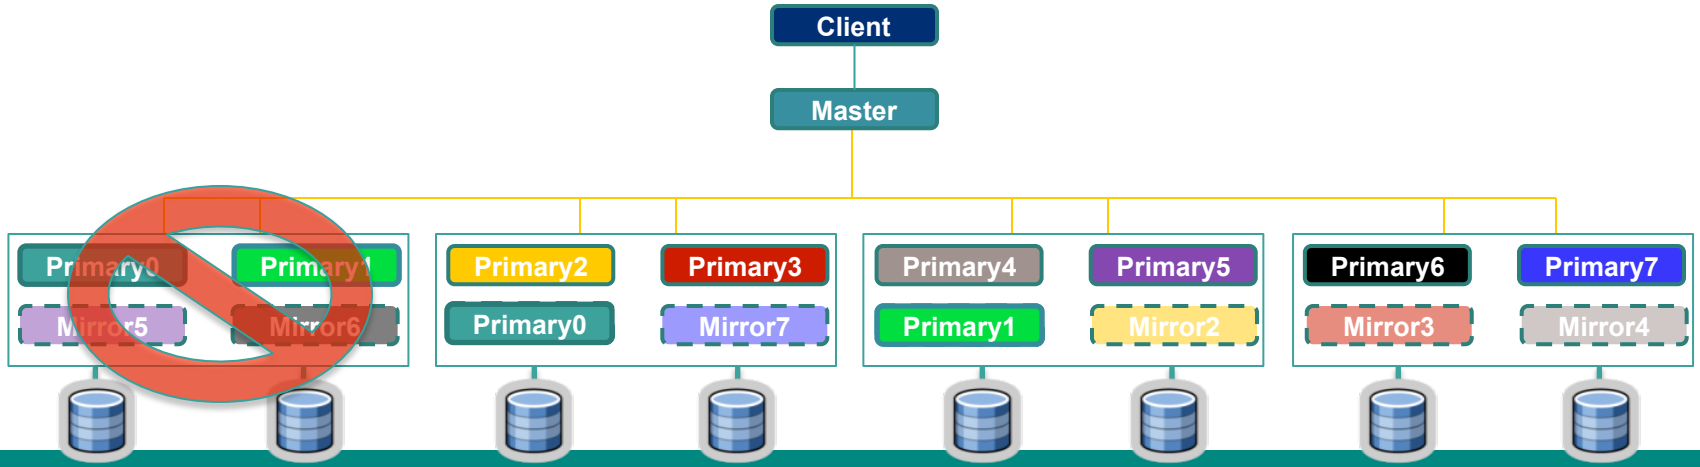

Workload Management

Pivotal

Workload Management

- Database Roles
  - Users and Groups
    - Controls permissions and access
- Resource Queues
  - Controls number of concurrent queries
  - Controls allocation of processing power to queries

Pivotal

Database Roles

- Manages access to objects and operations
  - `CREATE ROLE jsmith WITH LOGIN;`
  - `GRANT INSERT ON mytable TO jsmith;`
- Roles can also be used as Groups
  - `CREATE ROLE admin CREATEROLE CREATEDB;`
  - `GRANT admin TO jsmith, sally;`

Pivotal

Roles

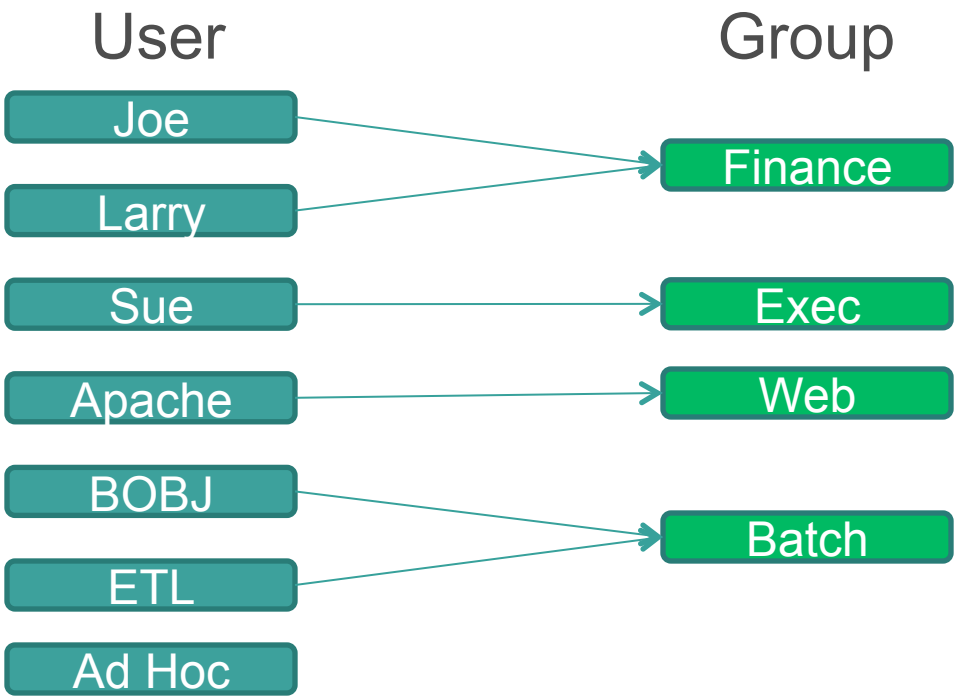

Pivotal

Resource Queues

- Controls query concurrency and priority
- Roles are assigned to Resource Queues
  - Queues can have multiple roles assigned
  - A Role can only be assigned to one Resource Queue
- Queue Limit options
  - Count Limit – Max number of active queries in the queue
  - Cost Limit – Max Query Planner cost of active queries in the queue
  - Minimum Cost Limit – Queries below this cost are not queued but run immediately
    - Can be specified with either Max Count limit or Max Cost limits
- Queue Priority
  - Specifies Min, Low, Medium, High, Max Resource utilization

Priority	Weight
Min	100
Low	200
Medium	500
High	1000
Max	1,000,000

Pivotal

Resource Queues

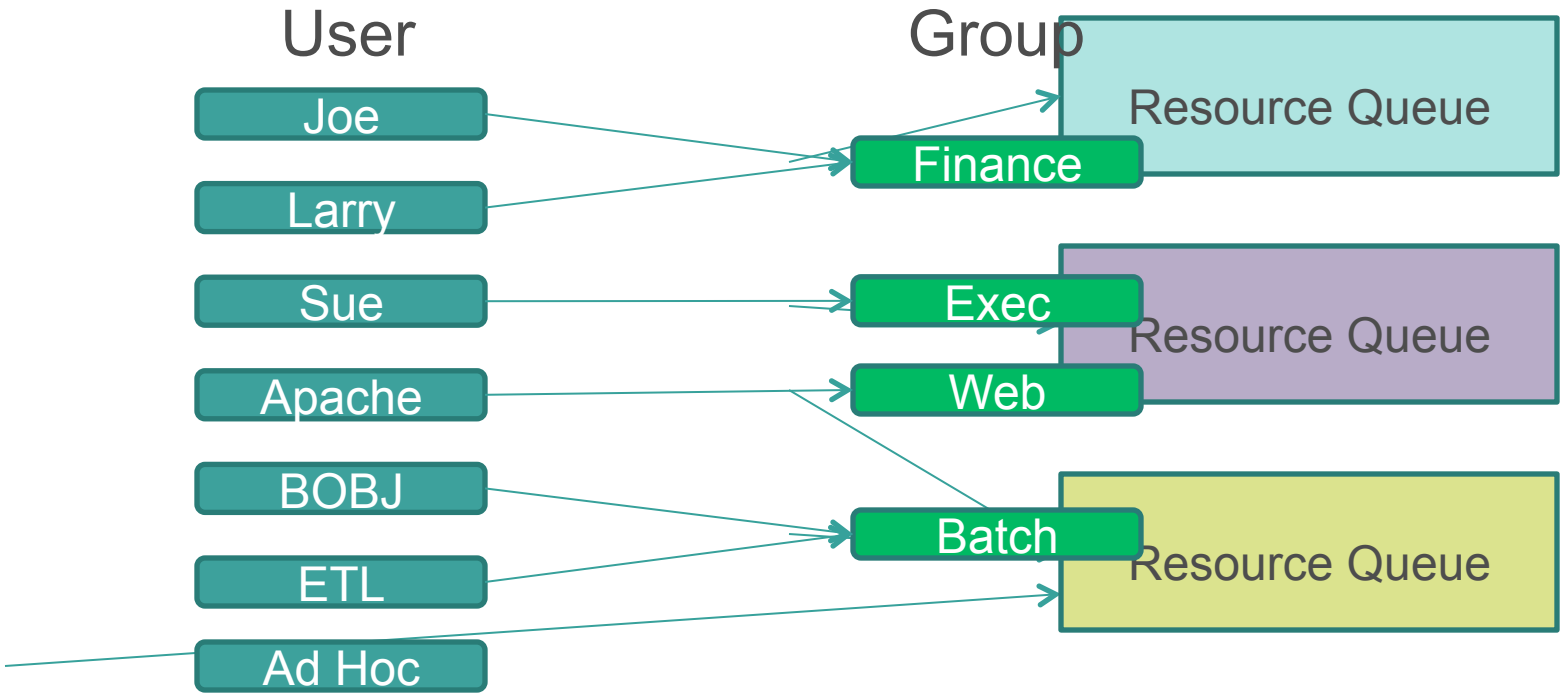

Queue Limit: Active Statement Count

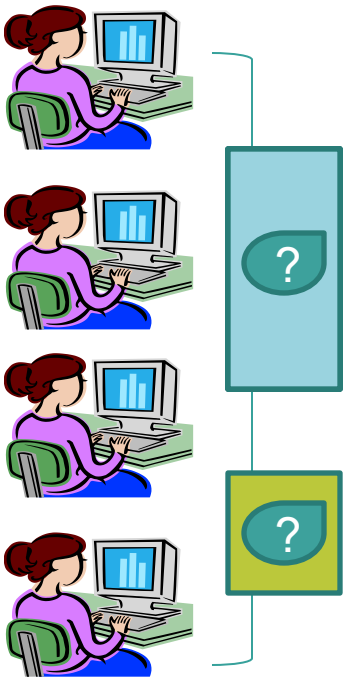

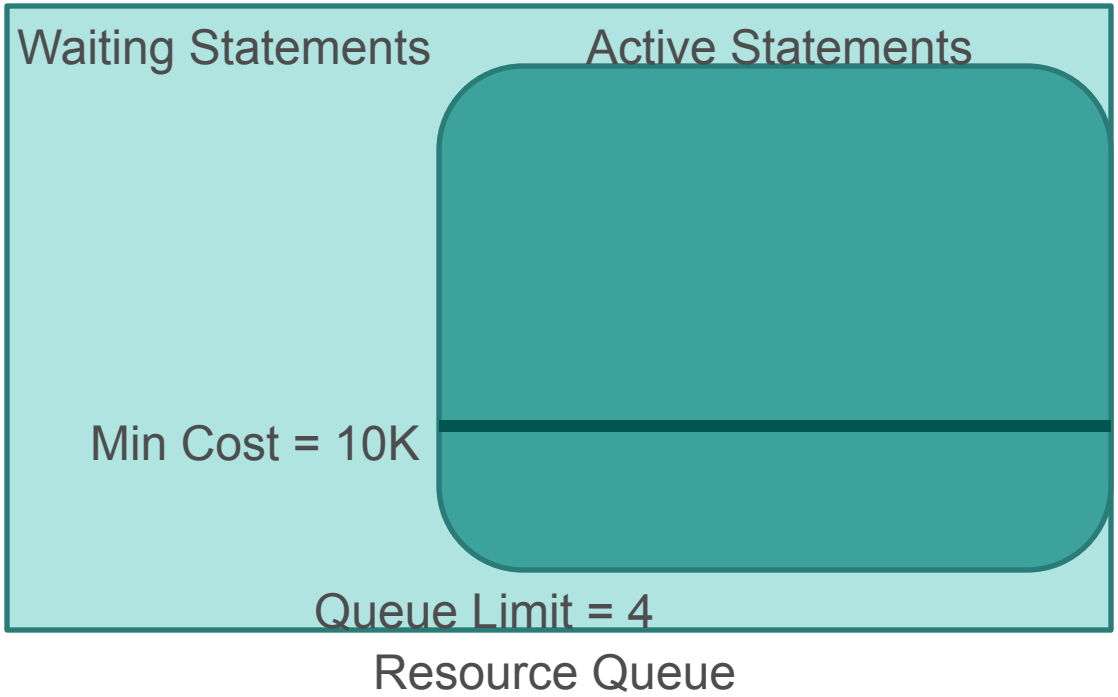

Queue Limit: Active Statement Cost

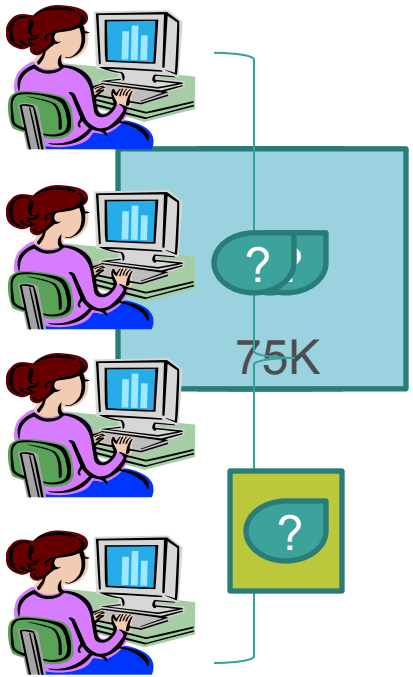

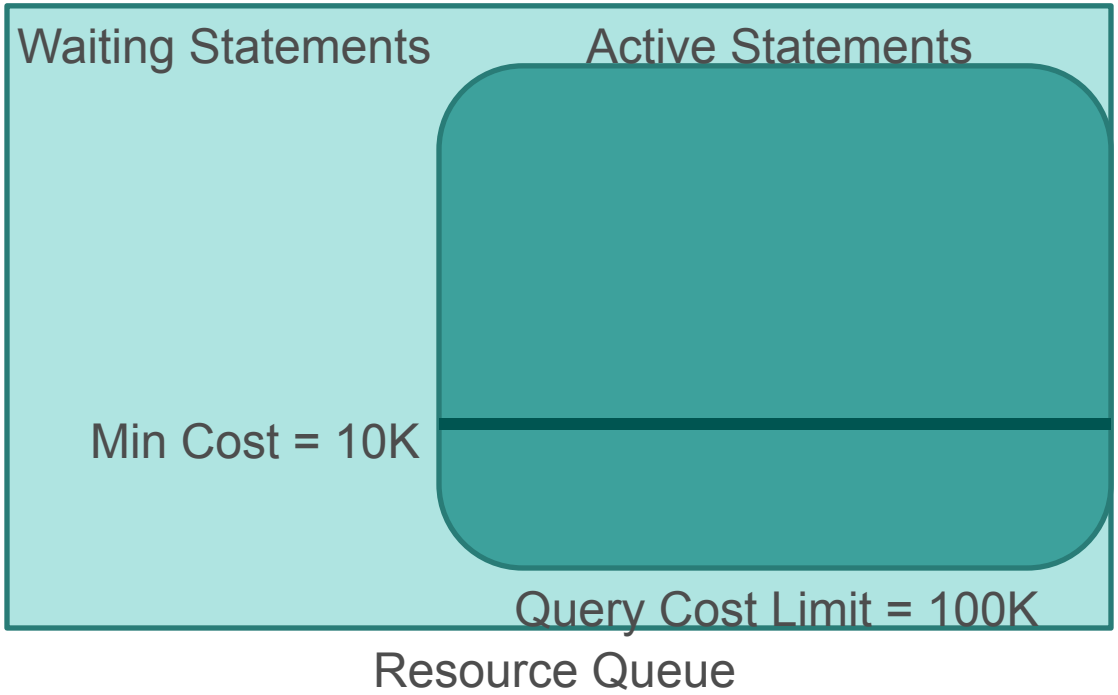

Queue Limit: Active Statement Cost

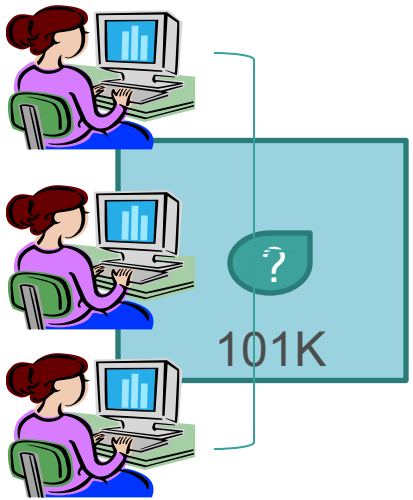

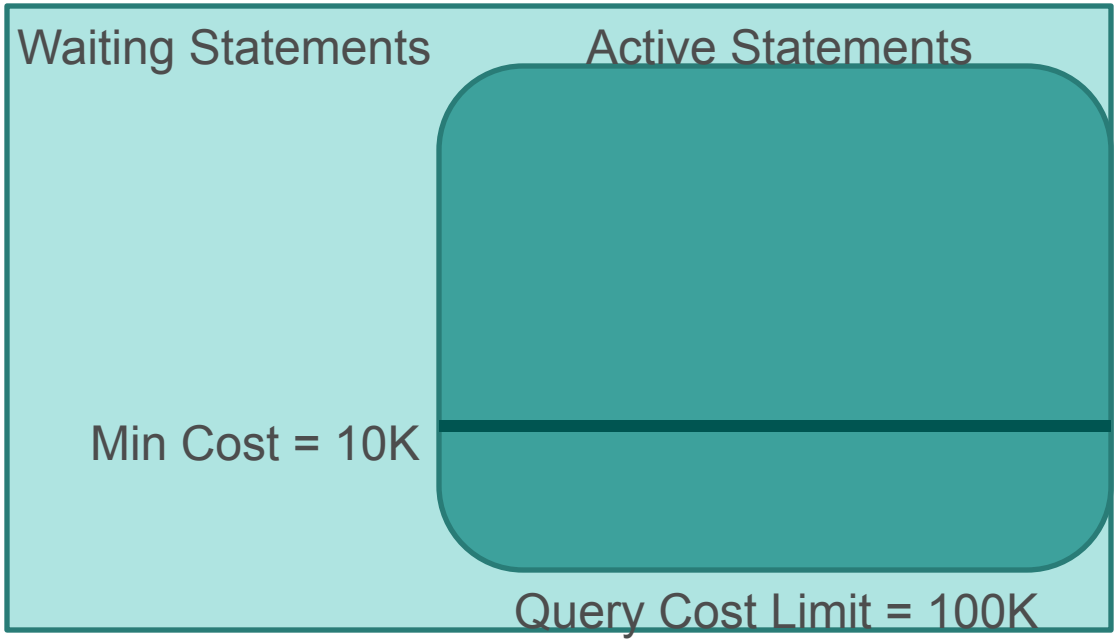

Resource Queue

Queue Priority

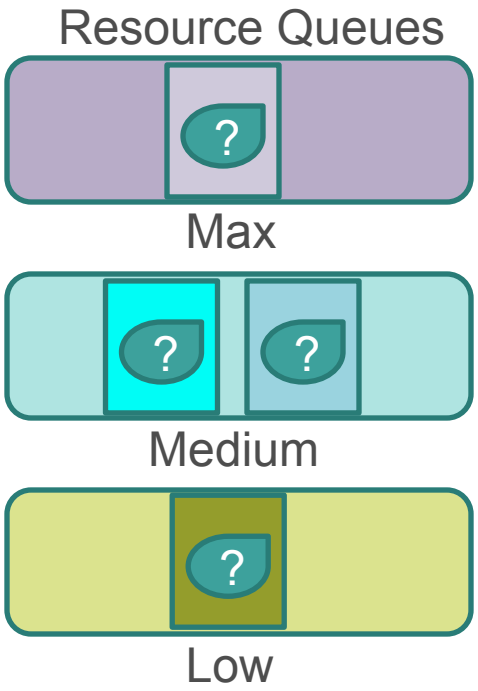

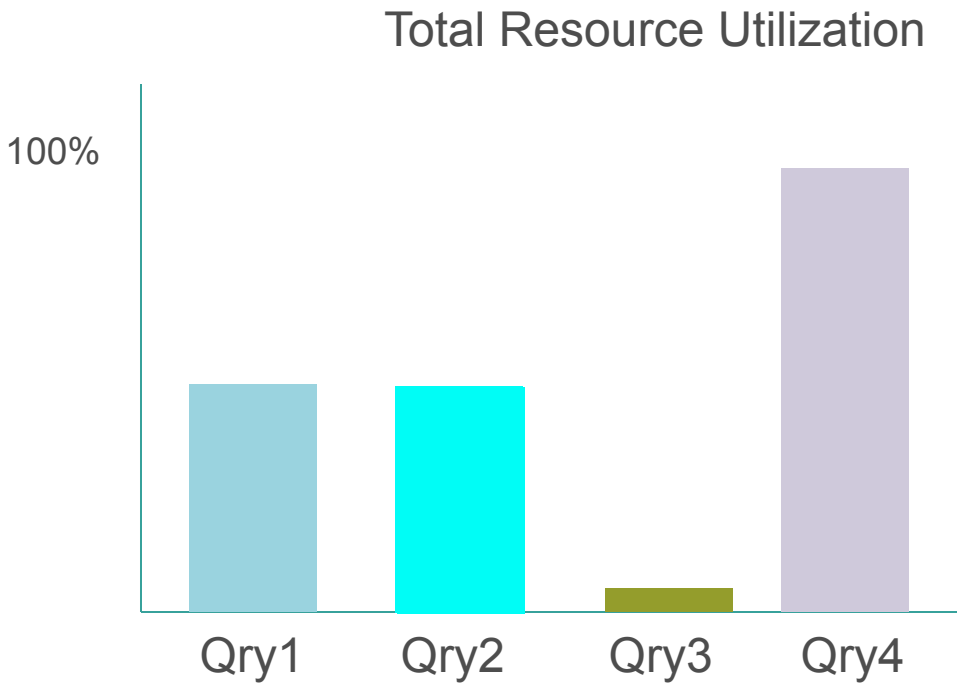

Pivotal

Align Queue Resource to Business Priorities

Queue Function	*Priority*	*Active*	*Min Cost*	*Max Cost*	*Over-commit*
Batch ETL / ETLT	Low	25			
Trickle Feed ETL / ETLT	Low	10	10k		
Web App / BI Parameterized Reports	High	50	5k		
Ad-Hoc, Power Users	Med	5	10k	150m	Yes
“VIP” Queue (with tripwire)	Max	3	10k	50m	Yes
Data Mining	Min			500m	Yes
“Cartesian Threat” or “Penalty Box”	Min			150m	No
Default Queue, if none assigned	Med	20			No

Pivotal

BUILT FOR THE SPEED OF BUSINESS