2ndQuadrant®
PostgreSQL

Pglogical and Postgres-BDR

14th Feb 2019

Nikhil Sontakke
PostgreSQL Consultant, 2ndQuadrant

pglogical

PostgreSQL to PostgreSQL replication

- Logical replication (one direction only)
- De-facto tool for
  - Data Movement
  - Migration
  - Low Downtime Upgrade
- Pglogical 2.2 available as binary and also now available on Amazon RDS by popular demand

Pglogical

Releases and Plans

- Pglogical 2.2 available now
  - Supports PostgreSQL 11 and earlier
  - Will continue to be supported for 3 years
  - Much community interest and patches - thanks!
- Pglogical 3
  - Future architecture for PostgreSQL 12+
  - Many new features

Pglogical 3

- New architecture for performance
  - Split Receiver/Apply processes, so now 2 procs
  - Similar to WALReceiver/Startup process in physical streaming replication
- Pluggable Architecture
  - Allows heterogenous Source(s)/Target(s)
  - Kafka, RabbitMQ targets available
  - Bulk-mode for high volume data loads
- Parallel initial copy on roadmap for pglogical 3

Postgres-BDR3

Future of multi-master database for PostgreSQL

- **BDR is Bi-Directional Replication (BDR)**
  2-way/Multi-way logical replication

- **Postgres-BDR3** runs on **PgLogical3**
- Coordinated code bases, operating as a full stack
- Streamlines feature development and support
- New feature releases quarterly

- BDR and pglogical 3.5.1 now available

Postgres-BDR

Trusted multi-master database for PostgreSQL

- Postgres-BDR 3
  - **Future architecture**

- Postgres-BDR 2
  - Upgrade path to BDR3 due 3Q2019
  - De-supported by mid 2020
- Postgres-BDR 1
  - Upgrade path to BDR3 due 1Q2019
  - De-supported by end 2019

Postgres-BDR3

Future of multi-master database for PostgreSQL

- Postgres-BDR Standard Edition
  - Runs as an *Extension* on Community PostgreSQL 10+
- Postgres-BDR Enterprise Edition
  - Runs as an Extension, using ***minimal*** additional features in 2ndQuadrant PostgreSQL 11+

Using Postgres-BDR

Transparent multi-master database for PostgreSQL

- Connect to any node
- Read and Write SQL like standard PostgreSQL
- Create/Manage tables like standard PostgreSQL

- Options to control which tables are replicated
- Some minor restrictions on DDL
- Design considerations for distributed database

Writing to Postgres-BDR

Distributed database options

- **Post**-Commit Synchronization
  Resolve issues *after* COMMIT
  - Conflict-Free Custom Datatypes (CRDTs)
  - Row-level Conflict Handling by default
  - Column-level Conflict Handling option
  - Logging and resolution of issues
  - Conflict Triggers

Writing to Postgres-BDR

Distributed database options

- **Pre**-Commit Synchronization
  Eager Replication *avoids* conflicts
  - All Nodes

- Avoids issues at COMMIT
  - Additional latency not desirable in many cases
  - Some transaction aborts in conflict cases

BDR Cluster Management

Enhancements for Distributed PostgreSQL

- Node Roles
  - Send/Receive Node (default)
  - Receive-Only Node
- Rep Sets (node subsets)
- *Rolling Cross-Version Upgrades*
  - One-by-one upgrades with zero downtime
  - Minimal APIs designed to allow upgrades
  - Across BDR releases **and** Postgres major releases
  - Fully tested to ensure compatibility

BDR Security

Cloud and High Security use cases

- BDR changes are applied as the table-owning user
- BDR functions do not need to be executed as superuser, except for installation
- Default roles allow principle of least privilege
- Secure access to BDR catalog tables for non-admin users

BDR Robustness and Supportability

When it's critical, you can count on us!

- Active-Active architecture offers Very High Availability
- Each master has multiple protection options
  - Logical Standby(s)
  - Physical Standby(s) (or both at once)
- Synchronous and asynchronous Replication options
- **Commit At Most Once option**
- Trace facility for investigation of complex bugs in distributed system

BDR Application Maintenance

DevOps Ease of Use

- Transparent handling of DDL for partitioned tables
- Optional modes for very detailed specification of node targets, allowing DDL even with down nodes
- Allows differing indexes on different nodes, allowing nodes for Business Intelligence and New Application-Version Testing
- Stream Triggers for filtering and transformation

BDR Operations

Practical experience

- Clearly identify node state, even when disconnected
- Monitor current lag time and lag bytes, even for down nodes
- Options for controlling down-node effects
- Error handling and controls

- Postgres-BDR plugin for OmniDB also available

- Postgres Cloud Manager with cluster visualization

Postgres-BDR Plugin for OmniDB

Visual Administration

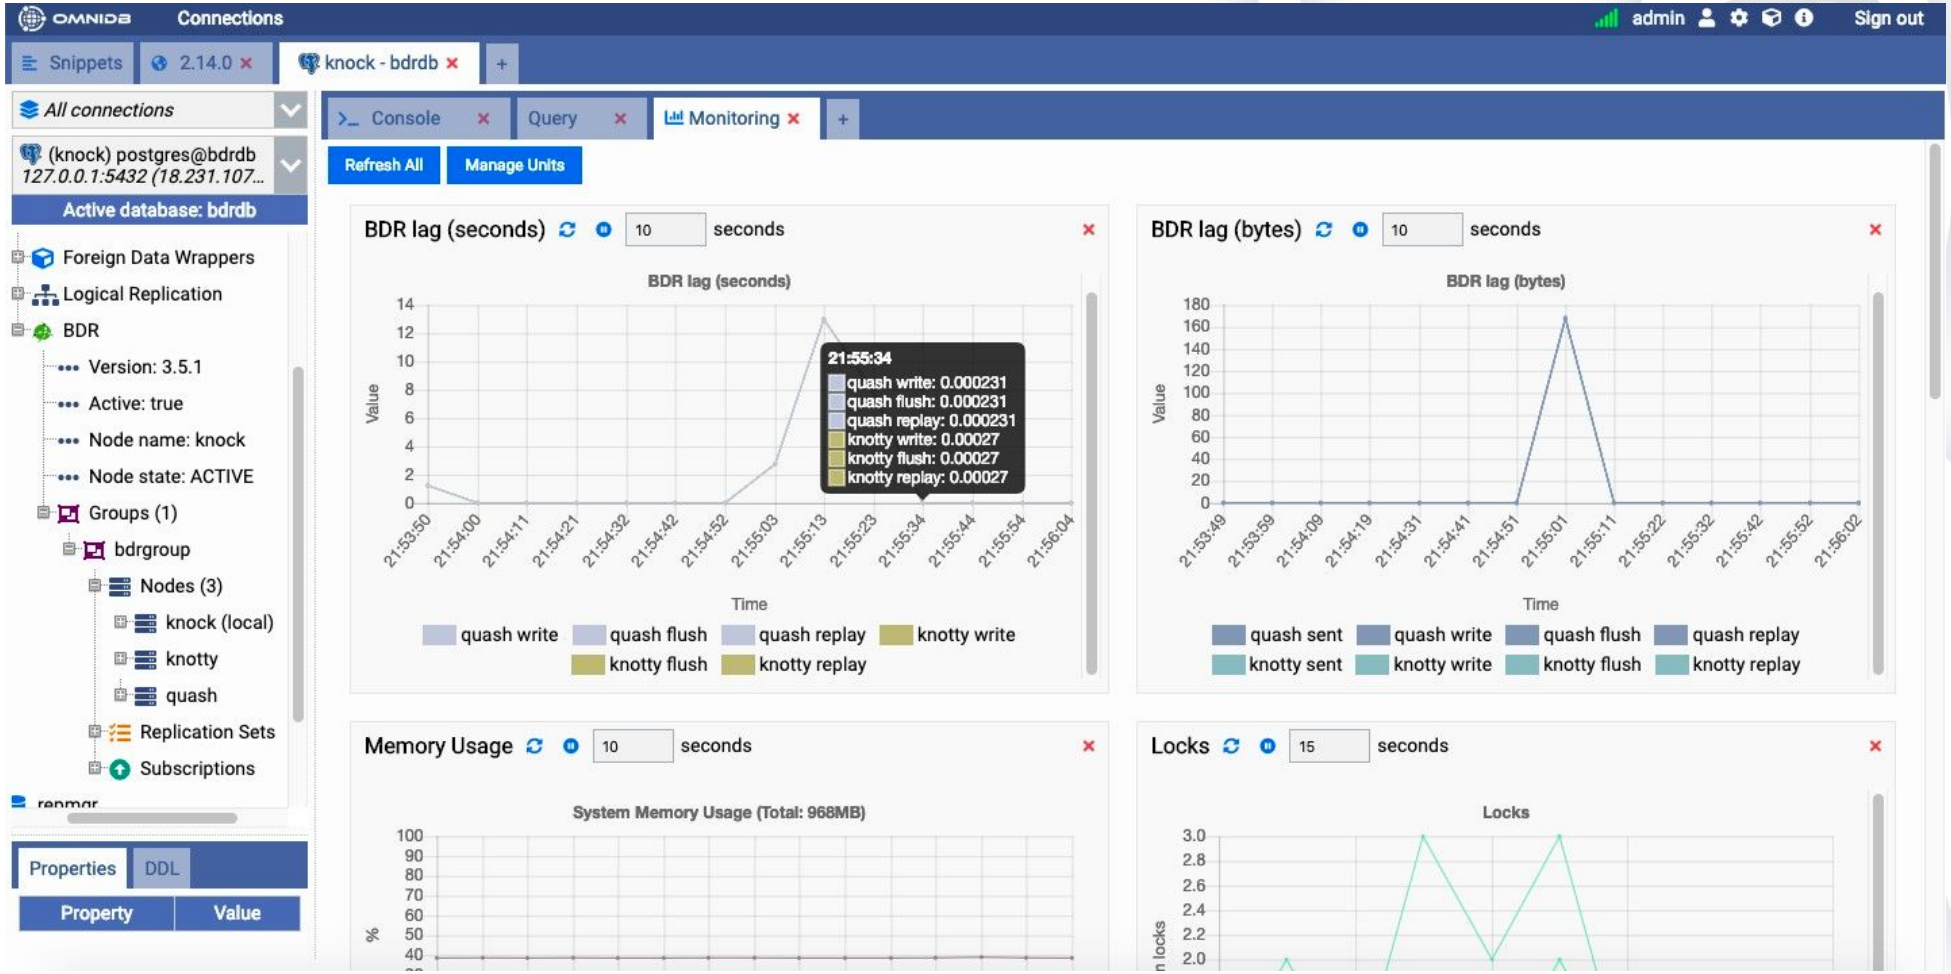

BDR Performance

Real-World Production Performance

- Efficient logical replication
- Streaming of large transactions
- Efficient distributed sequences
- Fast node addition
- Choice of options for selecting appropriate robustness and performance trade-offs

- Put data where its needed,
  keep it fresh easily and automatically

BDR Multi-node Query

Consistency and Performance

- Timestamp-based snapshots
- Allow consistent queries across nodes even with real-time replication of data
- Data verification between nodes

- Multi-node parallel query (MPP) across
  - Local clusters with remote DR nodes
  - Geo-distributed clusters

Full Stack Integration

Working together as one

- Fully integrated stack
  - Full automatic testing of operations and upgradability
  - 258 pages of detailed documentation
- Postgres-BDR, pglogical, 2QPG
- Repmgr, Barman
- PCISecurityPack and other tools
- OmniDB

- TPAexec Cloud/On-Premise Orchestration

Postgres-BDR3 Summary

Major new functionality, born out of production experience

- Cluster Management
- Security
- Robustness and Supportability
- Application Maintenance
- Operations
- Performance

- Rapid, agile development from many user requests

2ndQuadrant PostgreSQL Solutions

Website https://www.2ndquadrant.com/
Blog https://blog.2ndquadrant.com/
Email info@2ndQuadrant.com