

Using PostgreSQL in Production Licensing and Professional Support

Koichi Suzuki

Nov. 13, 2020





Agenda page

- New formation
- Open source products and business use
- Professional service example
 - Support
 - Consultation
 - DBA/Monitoring





First of all ...

2ndQuadrant has been acquired by EnterpriseDB

- <https://www.enterprisedb.com/blog/how-edb-became-leader-postgres-market>
- Became the largest PostgreSQL Tech and Business company
- Will provide even more powerful and high quality product and support





About myself

- NTT Lab, NTT DATA, NTT DATA Intellilink
 - Unix L10N
 - Oracle
 - UniSQL
 - PostgreSQL (NTT OSS Center)
 - Postgres-XC, Postgres-XL
- Joined 2ndQuadrant
- Now EDB member





2ndQuadrant

Now a part of EnterpriseDB

- **Founded by an Enterprise Architect / Postgres Developer 15 years ago**
- **Contributed features to Postgres to make it Enterprise Ready**
 - Backup and Recovery
 - Point in Time Recovery
 - Streaming Replication
 - Logical Replication
 - Stored Procedures with Transactions
 - Performance improvements to partitioning
- **Funded by support of Postgres Servers**
- **Wrote the code, wrote the books**





Using Open Source Products in the business

It's free to use but ...

- **Issues while using in production**
 - Need specialist help
 - Too expensive to hire and too painful to foster
 - Obtain from communities
 - Many good people to provide help in best-effort basis
 - No SLA
 - Cannot expose sensitive information to communities
 - Not suitable for too individual issues
 - Database design
 - Information security policies



Top provider of Postgres Support

2nd Quadrant portion

- 100+ staff in over 20 countries worldwide
- Over 75% of revenue from PostgreSQL support
- Remaining Revenue comes from services
- Over 240 customers, including many fortune 100 companies
 - Major Credit Card Companies
 - Health Care Providers
 - Pharmaceuticals
 - Technology manufactures
- Over \$300 Billion per year of Financial Transactions flowing through 2ndQuadrant Supported Services

We will provide even better professional support through merging into EnterpriseDB



Support Plans

| | Diamond | Platinum | Gold |
|----------------------|------------|------------|----------------|
| Postgres BDR Support | yes | no | no |
| Response SLA | 15 minutes | 15 minutes | 1 hour |
| Workaround SLA | 4 hours | 4 hours | Ongoing effort |
| Bug Fix SLA | 24 hours | 24 hours | Ongoing effort |

- No other vendor offers guaranteed resolution times
- Staffed by Postgres Developers, Authors and DBAs
- Company is built around providing best possible support



Support Portal

The screenshot displays the EDB Support Portal interface. On the left is a dark sidebar with navigation links: Home, Knowledge base, Videos, Create support issue, Support, Company info, Products, Data collector, and Global DBaaS (GDS). The main content area is divided into several sections:

- User Profile:** A welcome message for Tom John Kincaid (ACME), last login on SSO 3 minutes ago. It lists service and newsletter subscriptions with icons for various products like PostgreSQL, pglogical, and BDR.
- Support Summary:** Two cards showing '1 Active issues' and '26 Issues in the last 6 months', each with a 'More info' link.
- Technical alerts:** A list of alerts, including one from Feb. 28, 2019, regarding PostgreSQL 9.6.11 or 9.5.15 upgrades, and another from Aug. 22, 2018, about binary compatibility for BDR and pglogical.
- Videos:** A featured video titled 'BDR 3.5 Release Overview' from 7 Mar 2019.
- Latest knowledge base docs:** A list of recent documents such as 'glibc locale data changes' and 'PITR with Barman'.
- Latest software releases:** A list of recent releases including 'BDR 3.6' and 'pglogical 3.6'.



2nd Quadrant Support

- L1 Engineers Follow the Sun – 4 shifts / day
- Ever growing knowledge base
- Always perform work in timezone
- Major Incident Manager for Sev. 1 calls
- Dozens of Sev. 1 Data Loss Bugs fixed
- Everybody does support



Some of our contributions to Postgres

- Online backup
- Point in Time Recovery
- Streaming Replication
- Logical Replication
- Many enhancements to Vacuum and MVCC
- Transaction controlled Stored Procedures
- Incremental Sort
- Many other features



Supported Software List

- 7 Years of support for all major Postgres releases
- Support for important infrastructure projects including ...

| Project | Description |
|-----------|---|
| repmgr | Managing and creating replicas and providing automatic failover |
| barman | Backup and Recover (including 0 RPO) |
| pglogical | Advanced Logical Replication between Postgres |
| pgaudit | Audit logging for Compliance |
| pgbouncer | Lightweight connection pooler |
| OmniDB | Multi-database management tool |



Training

- Delivered on-line with Advanced Video techniques
- Lab Machines fully provided
- Advanced Video and teaching techniques



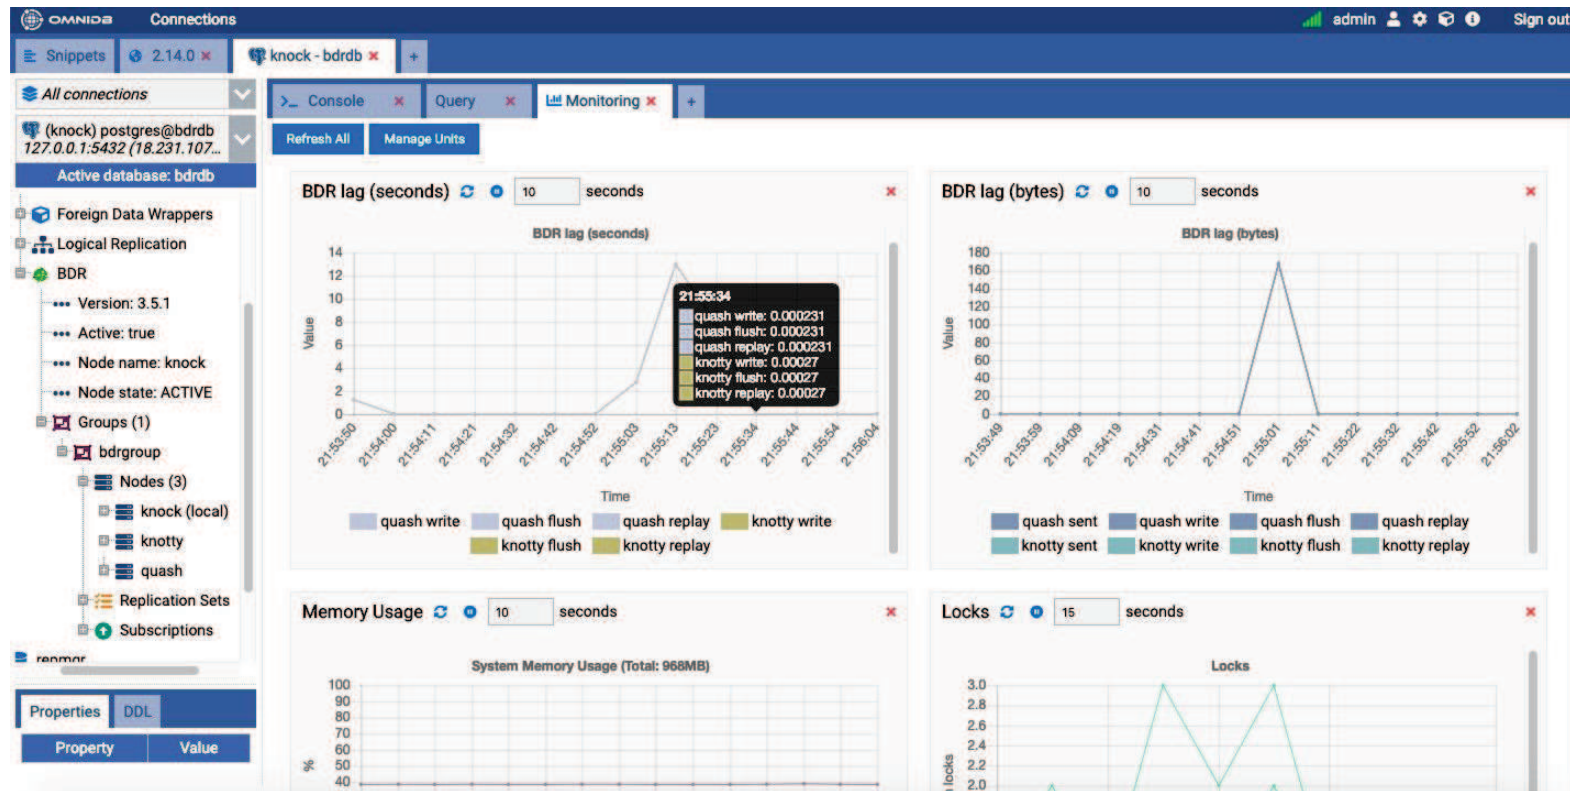
2nd Quadrant Training

| Title | Modules | Days | Description |
|---------------------------|---------|------|---|
| Postgres Development | 22 | 3 | For developers new to Postgres |
| Advanced Development | 11 | 3 | In depth development topics and architecture |
| Administration | 10 | 2 | Basics of PG Administration |
| DR / HA | 17 | 3 | HA Disaster Recovery Administration |
| Immersion | 27 | 5 | Zero to hero (customized for team) |
| Linux for PostgreSQL DBAs | 8 | 1 | Gain Linux skills to Administer Postgres |
| BDR | 2 | 1 | Designing and deploying applications with BDR |
| Oracle Migration | 5 | 1 | How to migrate from Oracle to Postgres |



Postgres-BDR Plugin for OmniDB

Visual Administration

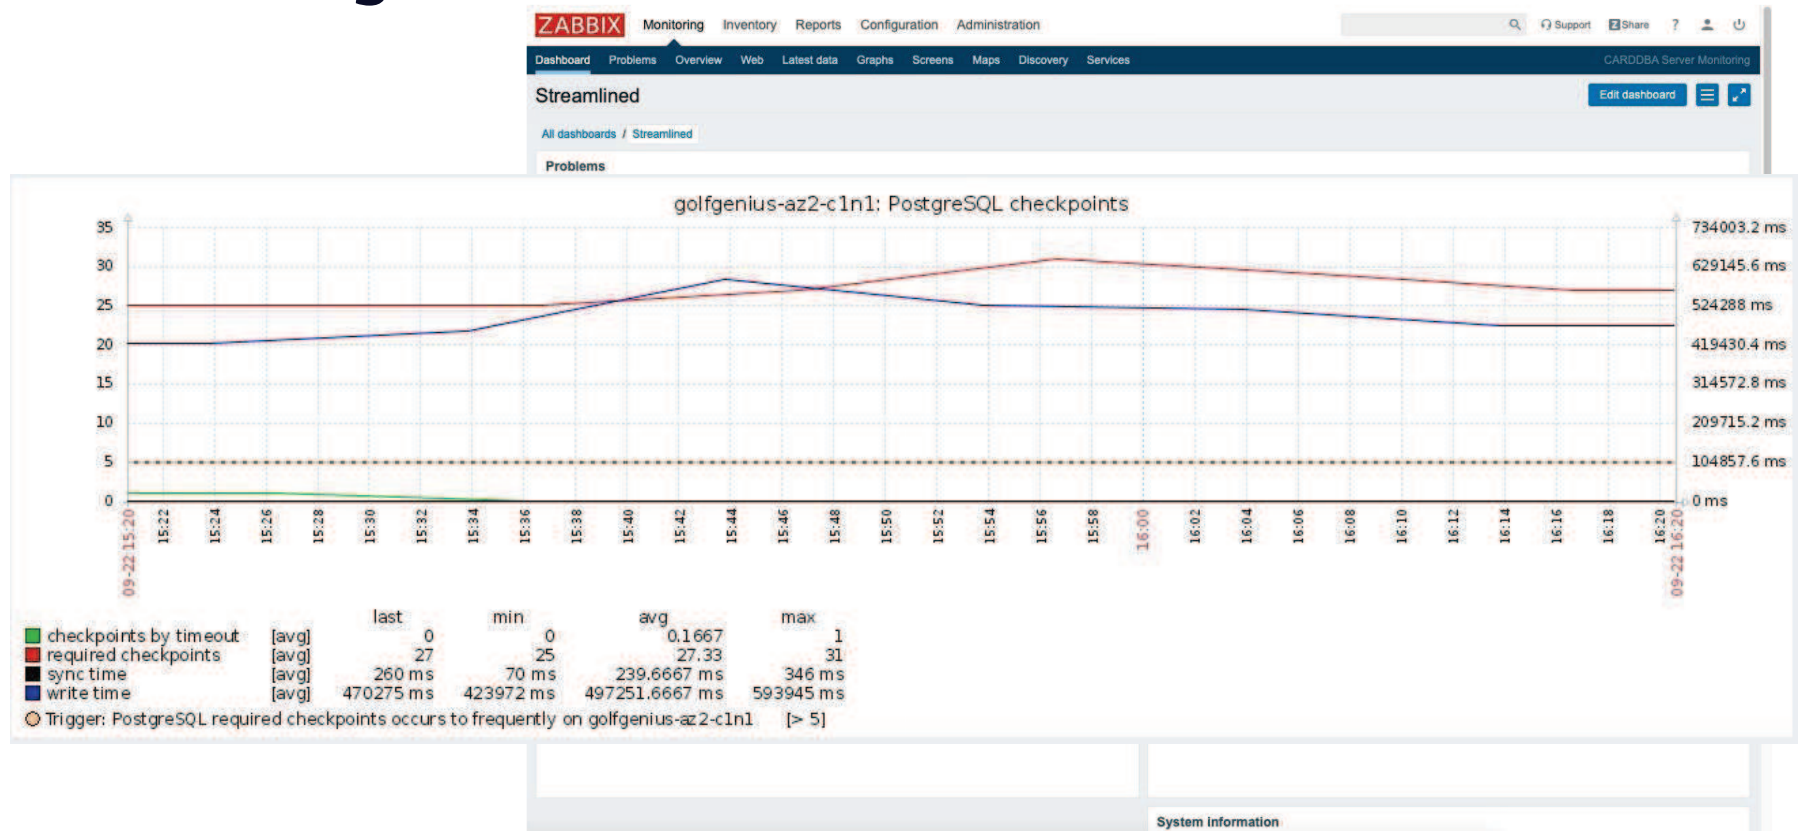




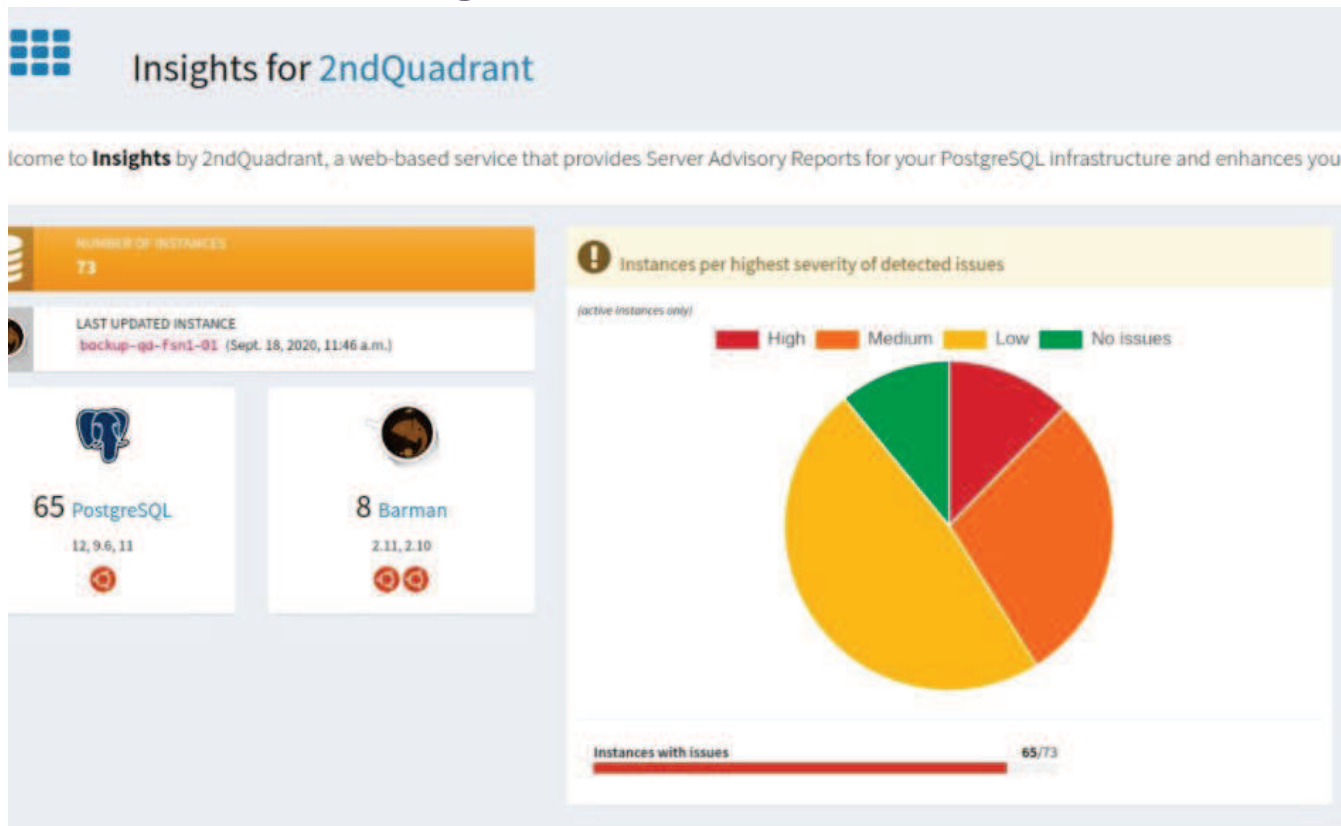
Monitoring

- 2ndQuadrant can setup and train on Zabbix Monitoring
 - A completely open source solution
 - Support available from Zabbix
- 2ndQuadrant has extensive experience around monitoring production servers.
 - Will be announcing a new monitoring solution for our customers shortly

EDB[™] Monitoring Zabbix



EDB[™] 2ndQuadrant Insights



EDB[™] 2ndQuadrant Insights

The screenshot displays the 2ndQuadrant Insights web application. The top navigation bar includes the logo, '2ndQuadrant Portal', a search bar, and a 'Clock-in' button. A dark sidebar on the left contains navigation options like Home, Knowledge base, Videos, Support, Company info, Products, Data Collector, Insights, Sentinel, and PGM Service. The main content area features a header 'Insights for 2ndQuadrant' and a breadcrumb trail 'All instances > Tests list'. The primary section is titled 'Test: Resource stealing processes'. It contains a warning message: 'For the best and most consistent performance it is recommended that the system where a PostgreSQL instance is installed is entirely dedicated to Postgres. Please make sure that processes such as web servers, application servers, window managers, or even other database management systems are not running on the same system. This is a test that was partially carried out by 2ndQuadrant's Insights and that requires human input and/or judgement in order declare the outcome of the test.' Below this is a green button that says 'I want to acknowledge this issue permanently.' To the right, a 'PostgreSQL Server:' box identifies the instance as 'codescanner-qa-fsri-01.2ndquadrant.it (standby) at 2ndQuadrant'. A 'Summary of issues' section follows, listing five error messages. The first four are identical: 'Resource stealing process detected: /usr/sbin/apache2 -k start'. The fifth error message is a detailed log snippet: 'Resource stealing process detected: /srv/coverity/cov-platform/jre/bin/java -Djava.awt.headless=true -Dces.home=/srv/coverity/cov-platform -Ddir.log=/var/lib/coverity/cov-platform/logs -Dcatalina.home=/srv/coverity/cov-platform/server/home -Dcatalina.base=/srv/coverity/cov-platform/server/base -Djava.io.tmpdir=/srv/coverity/cov-platform/server/base/temp -Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager -Djava.util.logging.config.file=/srv/coverity/cov-platform/server/base/conf/logging.properties --add-opens=java.base/java.lang=ALL-UNNAMED --add-opens=java.base/java.io=ALL-UNNAMED --add-opens=java.rmi/sun.rmi.transport=ALL-UNNAMED -XX:HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/var/lib/coverity/cov-platform/logs -Xlog:gc*:info:file=/var/lib/coverity/cov-platform/logs/gc_%t_%p.log*:utctime,uptime,tags,level,filecount=50,filesize=40M -Djava.security.properties=/srv/coverity/cov-platform/extras/jdk/java.security.additional-XX:+UseG1GC -Xms512m -Xmx5g -server -cp /srv/coverity/cov-platform/server/home/bin/bootstrap.jar:/srv/coverity/cov-platform/server/base/bin/tomcat-juli.jar -Djava.library.path=/srv/coverity/cov-platform/server/base/webapps/ROOT/WEB-INF/lib/org.apache.catalina.startup.Bootstrap start'.

EDB™ 2ndQuadrant Insights

◀ All instances

Test: Unused indexes

An index can be either **mandatory** or **optional**.

Mandatory indexes are created automatically by PostgreSQL as part of **UNIQUE** or **PRIMARY KEY** constraints, or with a **CREATE UNIQUE INDEX** statement; the user cannot remove them without dropping the corresponding constraint.

Optional indexes are not associated to any constraint; they are created by the user with an explicit **CREATE INDEX** command, and they only affect performance.

By creating an index, the user adds a (usually small) overhead on some or all the write operations on the corresponding table, plus some other secondary costs affecting other actions. On the other hand, the existence of an index might make certain queries faster.

In some cases an index improves query speed by several orders of magnitude; in other cases none of the main queries are able to benefit from that index.

Mandatory indexes cannot be dropped. However, you can (and should) drop an optional index in case its maintenance cost is not justified by its performance improvements.

The `idx_scan` column of the `pg_stat_user_indexes` catalog table records the number of times an index was used by a query; it can be used to detect indexes that do not provide any performance improvements. Those indexes can be removed, with two caveats:

1. The query planner decides whether to use an index based on current information such as statistics and settings, meaning that an optional index might stop being useful at any time.
2. Physical standby instances are binary copies of their upstream instance; they have exactly the same indexes that exist there. Because of this, sometimes users create an index on the primary node because it is useful for some queries launched on one of the standbys.

This is a test that was carried out automatically by 2ndQuadrant's Insights and does not require any human intervention or supervision.

I want to acknowledge this issue permanently.

PostgreSQL Server:

ss0-db-ga-fsn1-02.2ndquadrant.it (standby) at 2ndQuadrant

Useful links

[Examining Index Usage in the official documentation](#)

Summary of issues

 Notice

There are 64 unused indexes in this instance. The listed indexes are not unique keys and have not been used in an index scan according to the collected statistics. Therefore, they can most likely be dropped, unless they are used in a



Remote DBA monitoring by Sentinel

The dashboard displays the following information:

- Host checks:** 195 Successful Heartbeat received, 1 Unknown No heartbeat (High).
- Unit checks:** 6735 Successful, 68 Unknown Failed / No data. Breakdown: Critical (0), High (3), Medium (69), Low (69). Last state: Critical (1), High (1), Other (47).
- Checks Summary:**
 - 16 Autovacuum Freeze, 8 Autovacuum Workers Usage, 4 Bloated Indexes
 - 8 Bloated Tables, 1 Custom, 14 Disk Usage, 2 Long Autovacuums
 - 1 Long Idle Transactions, 4 Long Queries, 3 Long Transactions, 20 Autovacuum Freeze
 - 1 Autovacuum Workers Usage, 2 Bloated Indexes, 4 Bloated Tables, 1 Long Autovacuums
 - 1 Long Queries, 2 Long Transactions, 38 pg_stat_statements Usage
- Notifications:** A list of alerts with details such as severity, source, and status.



Remote DBA monitoring -- All servers per customer

| Host | Status | Time Elapsed | Actions | Max Severity | Hostname | IPs | Role | Services | Notes |
|-----------------------|--------|----------------------------|---------|--------------|----------------------------|-----------------------------|---------|--------------------|---|
| 2ndQ Website | ok | 05/10 22:22 14d 5hr 3m | | ok | keener.mkt.2ndquadrant.com | 10.33.62.125 | | | |
| 2ndQ-Blog | ok | 05/10 22:22 14d 5hr 3m | | ok | quirk | 10.33.62.123 | | | |
| 2ndQ-QA Website | ok | 05/10 22:22 14d 5hr 3m | | ok | knife.mkt.2ndquadrant.com | 10.33.62.124 | | | |
| 2q PgConf | ok | 05/10 22:22 14d 5hr 3m | | ok | 2qpgconf.com | 10.0.0.8 | | | |
| Postgres London | ok | 05/10 22:22 14d 5hr 3m | | ok | postgreslondon.org | 10.33.62.117 | | | |
| bdr-trial-mgmt-backup | ok | 05/10 22:22 14d 5hr 3m | | ok | bdr-trial-mgmt-backup | 10.20.1.10 | | barman | |
| bdr-trial-mgmt-db-1 | ok | 05/10 22:22 14d 5hr 3m | | ok | bdr-trial-mgmt-db1 | 10.20.1.11 | Primary | postgresql | |
| bdr-trial-mgmt-db-2 | ok | 05/10 22:22 14d 5hr 3m | | ok | bdr-trial-mgmt-db2 | 10.20.1.12 | Standby | postgresql | |
| sentinel01 | ok | 12/08 22:48 68d 4hr 37m | | medium | sentinel01 | 34.239.49.100, 10.39.0.14 | | postgresql | |
| sentinel02 | ok | 12/08 21:43 68d 5hr 42m | | medium | sentinel02 | 3.89.199.7, 10.39.0.33 | | postgresql | |
| sentinel03 | ok | 22/07 20:47 89d 6hr 38m | | low | sentinel03 | 34.227.96.196, 10.39.0.170 | | postgresql, barman | |
| sentinel04 | ok | 20/08 19:20 60d 8hr 5m | | medium | sentinel04 | 35.176.130.231, 10.39.0.115 | | | |
| sentinel05 | ok | 20/08 20:37 60d 6hr 48m | | medium | sentinel05 | 3.9.16.208, 10.39.0.62 | | postgresql | |
| sentinel06 | ok | 25/08 10:33 55d 16hr 52m | | ok | sentinel06 | 35.178.44.71, 10.39.0.197 | | postgresql, barman | |
| staging | ok | 05/10 22:22 14d 5hr 3m | | ok | ip-172-26-6-2 | 172.26.6.2 | Primary | postgresql | this is a test system , used to test new features |



Remote DBA monitoring -- All servers per customer

| Host | Status | Time Elapsed | Actions | Max Severity | Hostname | IPs | Role | Services | Notes |
|-----------------------|--------|----------------------------|---------|--------------|----------------------------|-----------------------------|---------|--------------------|---|
| 2ndQ Website | ok | 05/10 22:22 14d 5hr 3m | | ok | keener.mkt.2ndquadrant.com | 10.33.62.125 | | | |
| 2ndQ-Blog | ok | 05/10 22:22 14d 5hr 3m | | ok | quirk | 10.33.62.123 | | | |
| 2ndQ-QA Website | ok | 05/10 22:22 14d 5hr 3m | | ok | knife.mkt.2ndquadrant.com | 10.33.62.124 | | | |
| 2q PgConf | ok | 05/10 22:22 14d 5hr 3m | | ok | 2qpgconf.com | 10.0.0.8 | | | |
| Postgres London | ok | 05/10 22:22 14d 5hr 3m | | ok | postgreslondon.org | 10.33.62.117 | | | |
| bdr-trial-mgmt-backup | ok | 05/10 22:22 14d 5hr 3m | | ok | bdr-trial-mgmt-backup | 10.20.1.10 | | barman | |
| bdr-trial-mgmt-db-1 | ok | 05/10 22:22 14d 5hr 3m | | ok | bdr-trial-mgmt-db1 | 10.20.1.11 | Primary | postgresql | |
| bdr-trial-mgmt-db-2 | ok | 05/10 22:22 14d 5hr 3m | | ok | bdr-trial-mgmt-db2 | 10.20.1.12 | Standby | postgresql | |
| sentinel01 | medium | 12/08 22:48 68d 4hr 37m | | medium | sentinel01 | 34.239.49.100, 10.39.0.14 | | postgresql | |
| sentinel02 | medium | 12/08 21:43 68d 5hr 42m | | medium | sentinel02 | 3.89.199.7, 10.39.0.33 | | postgresql | |
| sentinel03 | low | 22/07 20:47 89d 6hr 38m | | low | sentinel03 | 34.227.96.196, 10.39.0.170 | | postgresql, barman | |
| sentinel04 | medium | 20/08 19:20 60d 8hr 5m | | medium | sentinel04 | 35.176.130.231, 10.39.0.115 | | | |
| sentinel05 | medium | 20/08 20:37 60d 6hr 48m | | medium | sentinel05 | 3.9.16.208, 10.39.0.62 | | postgresql | |
| sentinel06 | ok | 25/08 10:33 55d 16hr 52m | | ok | sentinel06 | 35.178.44.71, 10.39.0.197 | | postgresql, barman | |
| staging | ok | 05/10 22:22 14d 5hr 3m | | ok | ip-172-26-6-2 | 172.26.6.2 | Primary | postgresql | this is a test system , used to test new features |

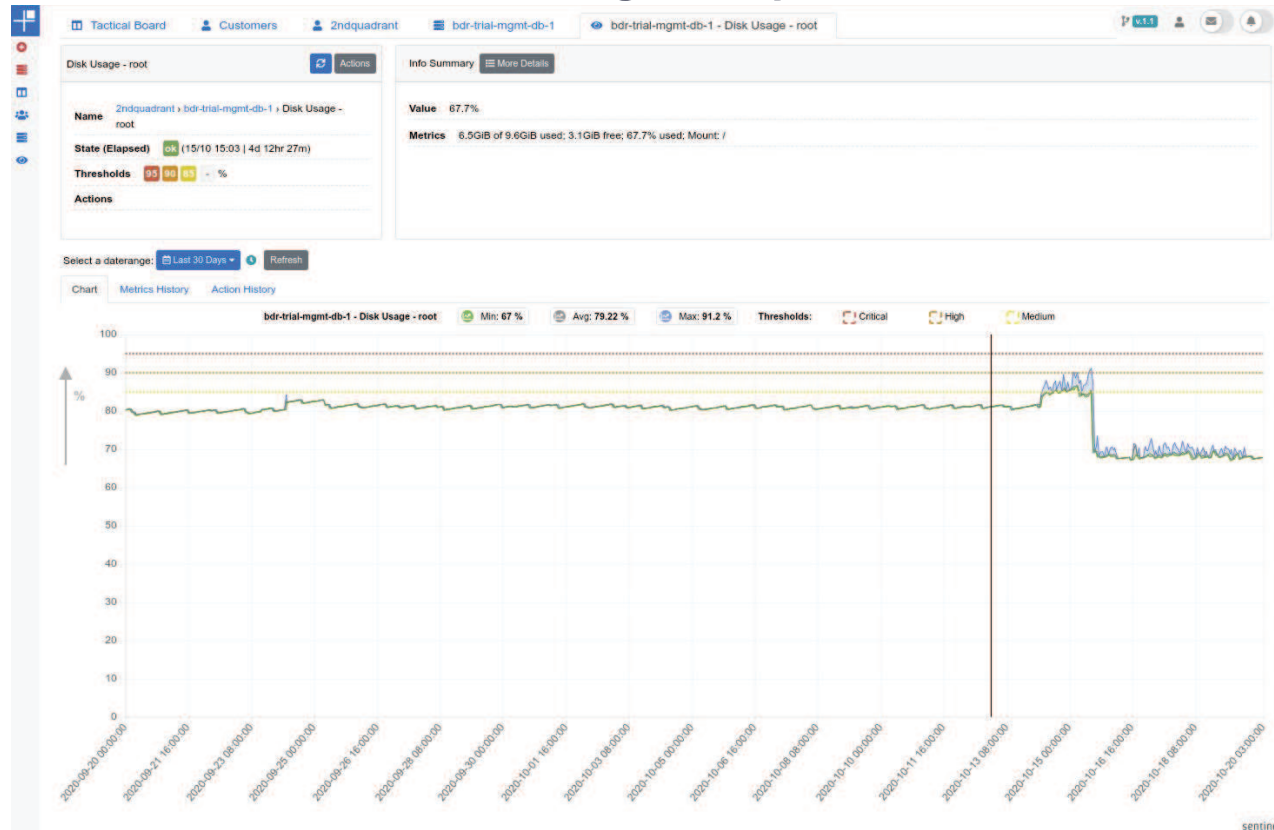


Remote DBA monitoring -- monitoring each server

| Unit Check | State | Severity | Time Elapsed | Actions | Thresholds (C / H / M / L) | Value | Metrics | No |
|----------------------------------|---------|----------|-----------------------------|---------|-------------------------------|-------------------------|--|----|
| Autovacuum Freeze | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / 90 / 80 % | 0.0% | 0 tables above 80%; Highest age at 0.0%; autovacuum_freeze_max_age: 200000000 | |
| Autovacuum Workers Usage | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / 100 / 70 % | 0.0% | 0 of 3 workers busy; 0.0% | |
| Backend Usage | success | ok | 05/10 22:22 14d 5hr 6m | | 95 / 90 / 85 / - % | 3.6% | 9 of 250 backends used; 3.6% | |
| Bloated Indexes | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / 1000 / 500 x | 0x | 0 indexes above 500x; Highest at 0x / 0.0B; Total estimated wasted space: 0.0B | |
| Bloated Tables | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / 300 / 150 x | 0x | 0 tables above 150x; Highest at 0x / 0.0B; Total estimated wasted space: 0.0B | |
| Blocked Locks | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / - / 500 | 0 | 0 blocked lock(s) | |
| Disk Usage - root | success | ok | 15/10 15:03 4d 12hr 25m | | 95 / 90 / 85 / - % | 67.7% | 6.5GiB of 9.6GiB used; 3.1GiB free; 67.7% used; Mount: / | |
| Exclusive Locks | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / - / - | 1 | 1 exclusive lock(s) | |
| Host up | success | ok | 05/10 22:22 14d 5hr 6m | | 1 / - / - / - | Up | bdr-trial-mgmt-db-1 is Up | |
| Inactive Slots | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / - / - | | No inactive slots | |
| Long Autovacuum | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / 43200 / 21600 seconds | 0 seconds | 0 autovacuum running for more than 21600 seconds; Longest: 0 seconds | |
| Long Idle Transactions | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / 86400 / 43000 seconds | 0 seconds | 0 transactions idle for more than 43000 seconds; Longest: 0 seconds | |
| Long Queries | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / 3600 / 1800 seconds | 0 seconds | 0 queries running for more than 1800 seconds; Longest: 0 seconds | |
| Long Transactions | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / 10800 / 7200 seconds | 0 seconds | 0 transactions kept open for more than 7200 seconds; Longest: 0 seconds | |
| Postgres Instance | success | ok | 05/10 22:22 14d 5hr 6m | | 1 / - / - / - | Up | Connected to PostgreSQL 11.7 (Ubuntu 11.7-2.pgdg18.04+1) | |
| Replay Paused | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / - / - | | Replay NOT paused | |
| Replication - bdr-trial-mgmt-db2 | success | ok | 04/05 14:50 168d 12hr 38m | | - / - / 5000 / 1000 MIB | Streaming / Lag=0.0 MIB | bdr-trial-mgmt-db2 in sync (lag=0.0 MIB); | |
| Replication Lag Standby | fail | unknown | 03/05 20:51 169d 6hr 37m | 🚫 | - / - / 900 / 300 seconds | Check failed | Unable to determine lag | |
| Role Change | success | ok | 05/10 22:22 14d 5hr 6m | | - / 1 / - / - | Same role | Same role: primary | |
| Sequences | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / 90 / 80 % | 0.0% | 0 sequences above 80%; Highest at 0.0% | |
| Share Locks | success | ok | 05/10 22:22 14d 5hr 6m | | - / - / - / - | 1 | 1 share lock(s) | |



Remote DBA monitoring -- specific measure



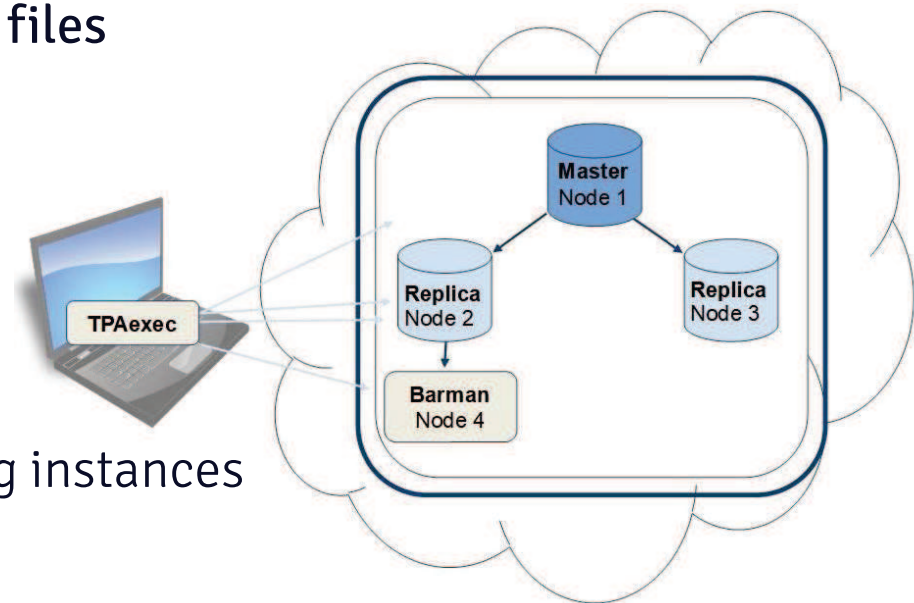


Trusted Postgres Architecture

- Ansible based technology for quickly deploying repeatable Postgres architectures
- 2ndQuadrant is a top contributor to Ansible outside Redhat
- Create clusters based on 2ndQuadrant best practices
- Entirely repeatable

TPAExec

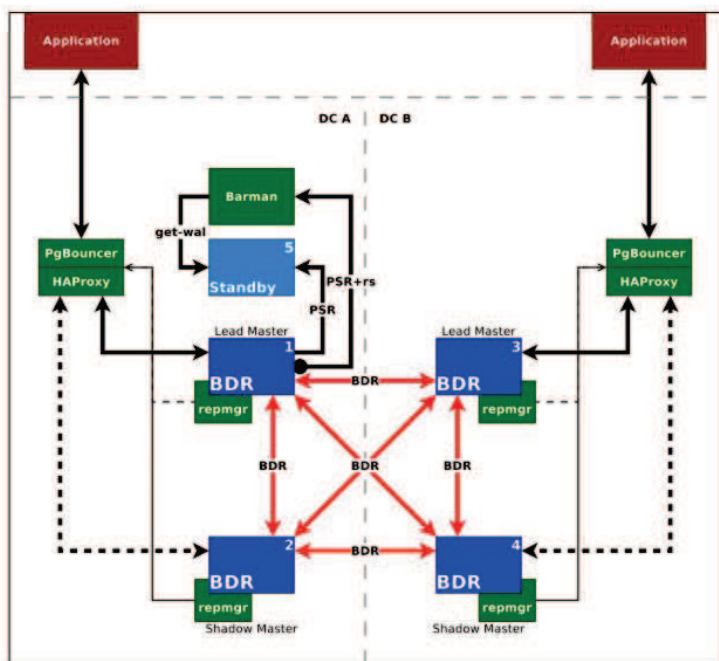
- Uses Ansible - config done via .yml files
- Platform independent - currently:
 - AWS
 - GCP
 - docker
 - vagrant on VirtualBox
 - bare metal - used for existing instances
- Platforms in the pipeline:
- Kubernetes/docker





TPAExec - Use proven configurations

```
tpaexec configure ~/tpa/test_cluster \  
--architecture BDR-Always-ON .....
```



2ndQuadrant
Trusted PostgreSQL Architecture:
Postgres-BDR AlwaysOn

Version 4.3
13 March 2019
2ndQuadrant



Kubernetes

- 2ndQuadrant is a Silver Member of the CNCF
- 2Q specific Kubernetes operators for BDR and PostgreSQL
- Moving to fully in-house Kubernetes infra using above
- OpenTracing built into BDR3.7 for end-end observability
- Prometheus storage plugin for BDR AutoScale 3.7
- Fluentd integration via syslog input



kubernetes



Postgres-BDR



Postgres-BDR3

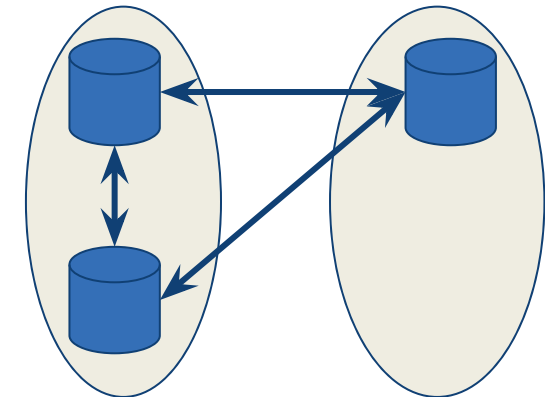

Future of multi-master database for PostgreSQL

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



BDR AlwaysOn

- One main node in heavy use
 - All nodes Active
- Fast switch to alternate nodes
- Run maintenance on “off” node
- 66 ms failover with no transaction loss





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

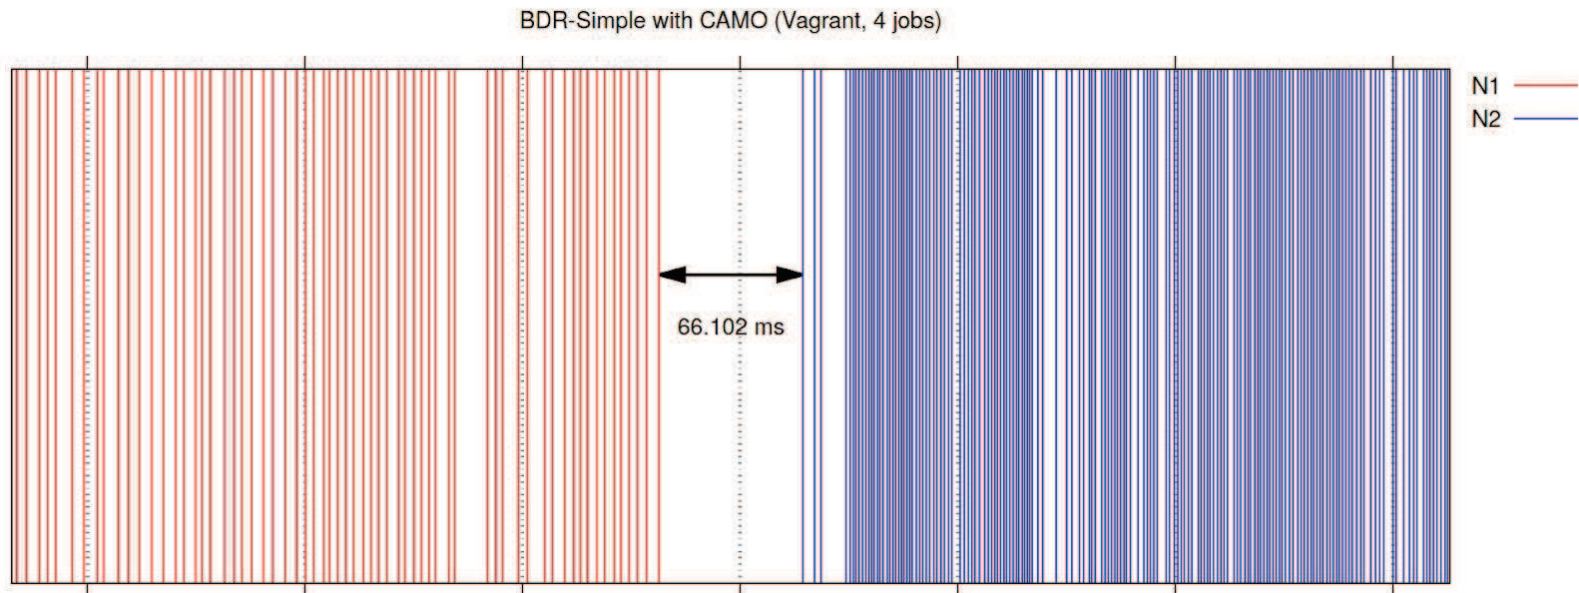


Writing to Postgres-BDR

- 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
- Low latency, suitable for widely distributed nodes



EDB™ BDR Fast Switchover



- Execute on **node1** until failure, fast failover **node2**
- Compare 30-90s for single master failover against <100ms for AlwaysOn failover



Rolling System Upgrades

- Rolling upgrades start with least used node and roll across all nodes slowly
- System Upgrades can upgrade BDR and/or main PostgreSQL releases
 - e.g. PG10 to PG11
 - E.g. BDR3.5.5 to BDR3.6.2
 - Nodes re-negotiate their protocols to ensure compatibility





World Wide Deployments

Conclusion

- Professional service boosts your database
 - Stability
 - Availability
 - Scalability
 - Productivity
 - Robustness
 - Problem solving
 - Day-by-day operation



Thank you very much

Koichi Suzuki

koichi.suzuki@enterprisedb.com