



EDB?

Simple, Reliable & Affordable

Roberto Giordano – Capital Markets Database Services Manager

Disclaimer

This content represents my own personal thoughts on these topics and are not meant to represent the views of London Stock Exchange.

The information contained here is not meant to be thought as a replacement of the official technical documentation provided by the vendor or any other professional service advice you may get.





**Our word is our bond.
True in 1801.
Just as true today.**

www.lseg.com

LSEG in one slide

London Stock Exchange Group (LSEG) is a global financial markets infrastructure business.

In Capital Markets, the Group operates a broad range of international equity, ETF, bond and derivatives markets.

In Information Services, through FTSE Russell, the Group is a global leader in financial indexing, data and analytics.

Post trade and risk management services are a significant part of the Group's business operations.

LSEG operates an open access model, offering choice and partnership to customers across all of its businesses. Headquartered in the United Kingdom, with significant operations in North America, Italy, France and Sri Lanka, the Group employs approximately 4,500 people.

Capital Markets Database Services

- Mission: Support Database Services of LSEG's Capital Markets business, no matter location or technology
- This means that we support more than one database technology for services based in different countries that belong to different legal entities within LSEG Capital Markets
- The Team implement a Follow The Sun model leveraging resources based on several locations in Asia, Europe and America.





Simple



Simple to Provision

- We provision new environments adapting Postgres roles and playbooks publicly available in Ansible Galaxy
- We install the binaries, create a new instance and its own standby, install and configure BART & PEM agents, etc all via Ansible

```
- name: Include EDB specific variables
  include_vars: "{{ inventory_dir }}/custom/edb_vars.yml"
  when: pg_type == "edb"
- import_tasks: install.yml
- import_tasks: create_directories.yml
- import_tasks: configure_environment.yml
- name: Create master cluster
  block:
    - import_tasks: initialize_master.yml
    - import_tasks: create_archivecommand.yml
    - import_tasks: create_autoconf.yml
    - import_tasks: create_pg_hba.yml
  when: pg_cluster_type == "master"
```



Simple to Maintain

- We developed patching automation who leverage Satellite and yum to deploy updated version of the binaries
- The automation takes care also of the database role (primary or standby) and put relevant resources under monitoring blackout to avoid fake alerts to be spammed

```
yum check-update $FILTER
```

```
yum -y update $FILTER
```

```
systemctl stop <service_name>
```

```
systemctl start <service_name>
```




Simple to Backup

- We use BART (Backup And Recovery Tool) as part of our backup policy.
- BART allows us to completely manage backup lifecycle, deletion of the old backups included
- BART is currently not integrated with our Media Manager, so the copy of the backup from disk to tape is managed outside BART (space to improve here)

```
bart -c $PGBACKUP/bart.cfg BACKUP -s $BART_SERVER -z
```

```
bart -c $PGBACKUP/bart.cfg VERIFY-CHKSUM -s $BART_SERVER
```

```
bart -c $PGBACKUP/bart.cfg SHOW-BACKUPS -s $BART_SERVER
```

```
bart -c $PGBACKUP/bart.cfg MANAGE -s $BART_SERVER
```

```
bart -c $PGBACKUP/bart.cfg MANAGE -d -s $BART_SERVER
```

Simple to Monitor

- We monitor our EDB targets using PEM (Postgres Enterprise Manager).
- PEM agents are deployed and configured as part of the provisioning process
- PEM is able to forward SNMP traps to our Central Monitoring System, some configuration is needed at that level to get meaningful alerts
- Monitoring and alerting for streaming replication for master – standby high availability architectures
- Easy to use blackout, retention and threshold management
- Secured agent connections and transaction pooling using pgbouncer module


Simple to Monitor - Dashboard

Home Global Overview ▾

Object Type System  Status N/A  Generated On 11/06/2020, 11:40:32  No. of alerts 10 (Acknowledged : 1)

▾ Enterprise Dashboard


Status




Simple to Monitor - Workload

[Dashboard](#)
[Properties](#)
[SQL](#)
[Statistics](#)
[Dependencies](#)
[Dependents](#)
[Monitoring](#)
[Global Overview](#)
[Manage Alerts](#)
[Manage Probes](#)


Server sessions




Transactions per second




Tuples in



Tuples out



Block I/O





Reliable

Database replica

- As part of our standard setup, every database instance has a standby database, this is also part of the automatic provisioning process
- If there're no specific requirements the Streaming Replication is configured using the default ASYNC mode
- In case of Automatic Failover needed and/or RPO=0 requirement the replica is configured as SYNC

```
$ cat $PGDATA/recovery.conf"
standby_mode = 'on'
primary_conninfo = 'host=999.999.999.999 port=9999 user=YYY password=ZZZ'
restore_command = 'cp /my_wal/my_pgwal_arch/%f %p'
trigger_file = '/tmp/trigger_dr_invocation'
recovery_target_timeline = 'latest'
```

EFM – EDB Failover Manager

- If Automatic Failover and/or RPO=0 are required, we rely on EFM and SYNC Streaming Replica to achieve that.
- We moved from OS Clustering to EFM for simplicity (KISS)
- EFM can provide both Intra and Inter site High Availability capability, best practices requires one witness per site

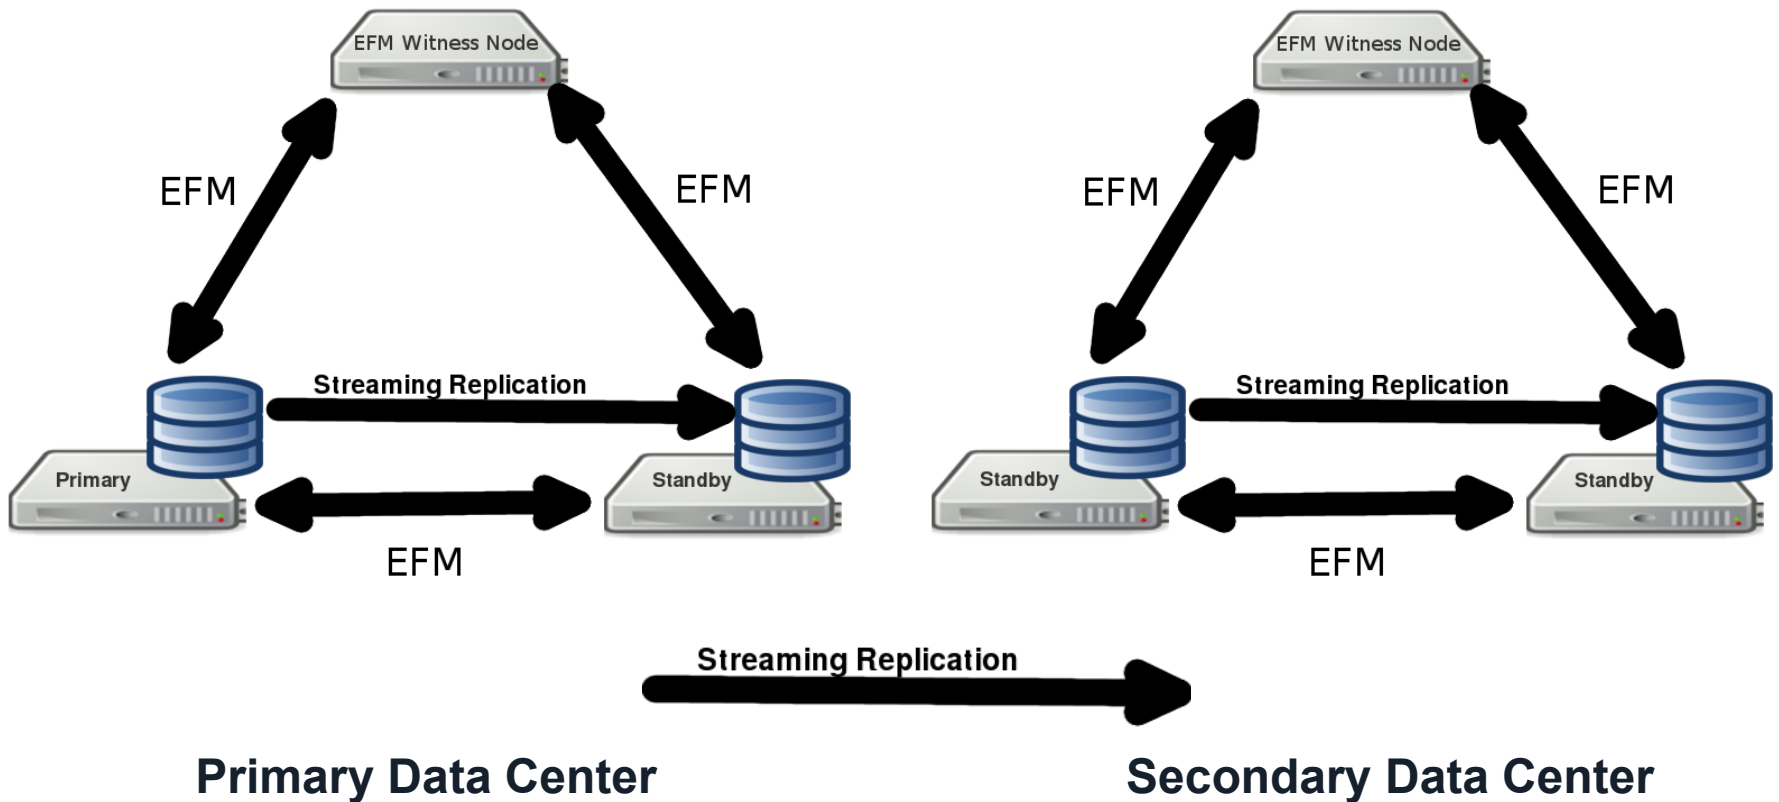
```
$ /usr/edb/efm-3.9/bin/efm cluster-status efm
```

```
Cluster Status: efm
```

```
Agent Type Address Agent DB VIP
```

```
-----  
Master 192.168.56.112 UP UP  
Standby 192.168.56.113 UP UP  
Witness 192.168.56.114 UP N/A
```

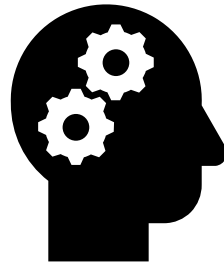
EFM – EDB Failover Manager





Support Service

- Easy to engage
- Very responsive
- Your gate to the Dev Team (less responsive than Support)
- Always ready to help, eventually 'overstepping'

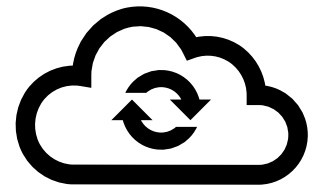
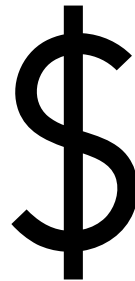
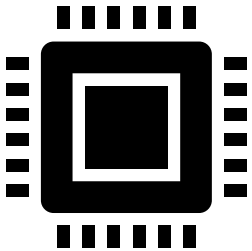




Affordable

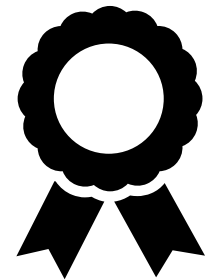
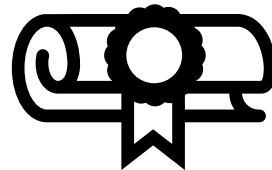
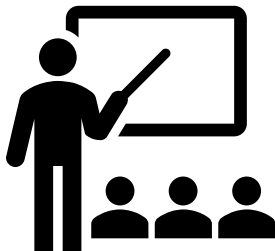
Clear & Flexible License Model

- Licensing model is just simple and clear
- You pay for the number of UniCores you use, no matter how you deploy them
- No big entrance fee, just a perpetual costs
- **Almost** a Cloud like model



Ad-Hoc Training

- EDB is capable on setting up ad-hoc training remotely or on live classes
- This can really boost the knowledge within your team with an quite affordable cost
- Certification sessions can be included

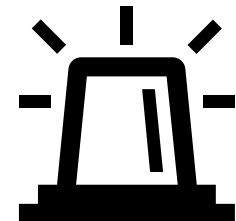
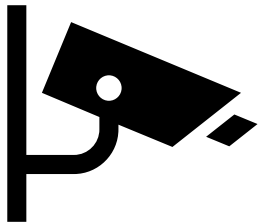




... but ...

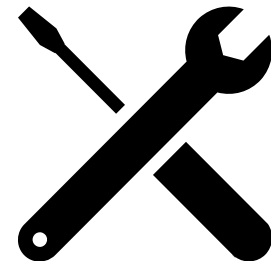
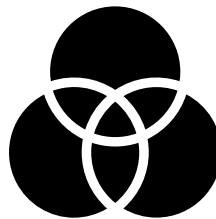
PEM has no external connector

- Unfortunately we recently ran into a big limitation of PEM compared to its competitors
- EDB doesn't provide any external connector to any other Central Monitoring Tools, just standard SNMP traps are available
- We believe this is big limitation as it means that a considerable amount of work is needed to configure SNMP traps to be meaningful



.. and no Active-Active yet

- We believe this is still the main gap between EDB and market leaders
- EDB has been working on an Active-Active solution (or Multi Master replica) for a while but, as far as we know, the results are not quite comparable
- BTW we're trying to understand this better ... so stay tuned!





Q&A

Roberto Giordano – roberto.giordano@lseg.com
<https://www.linkedin.com/in/roberto-giordano-a200a24/>