Percona XtraDB Cluster-5.7 + ProxySQL

For your high availability and clustering needs

Ramesh Sivaraman

Krunal Bauskar

Agenda

- What is Good HA eco-system ?
- Understanding PXC-5.7
- Understanding ProxySQL
- PXC + ProxySQL = Complete HA solution
- Monitoring using PMM
- Q&A

HA eco-system

- **Characteristics of good HA eco-systems**
  - No single point failure
  - Scalable
  - Auto-failover
  - Auto/Seamless joining of new members
  - System represent unified view of data
  - Secure connection among members
  - Transparent handling of communication failure
  - Ease of doing maintenance w/o affecting active workload
  - Effectively control workload execution

Workload suitable for HA

- **Non-Suitable**
  - Things that can’t be rolled back.
  - Locks
  - Local Operation
  - Mix of transaction and non-transactional workload.
  - Order preserving replication. (ROW vs STMT)
  - Entity that can’t be uniquely represented across the cluster.
- **Suitable**
  - Non-overlapping workload
  - Batched DDL and DML
  - Replay-safe transactions
  - Tables with primary-keys

Understanding PXC-5.7

- **PXC-5.7**
  - Multi-master Solution
  - Virtually synchronous
  - Auto node provisioning (SST)
  - Fully Scalable
  - Ability to view updated data (avoid stable view)
  - Seamless Incremental Transfer
  - Completely secure: in-transit + at-rest
  - Operates across WAN
  - Compatible with MySQL Master-Slave
  - Integrated well with leading products: PMM, ProxySQL, Xtrabackup

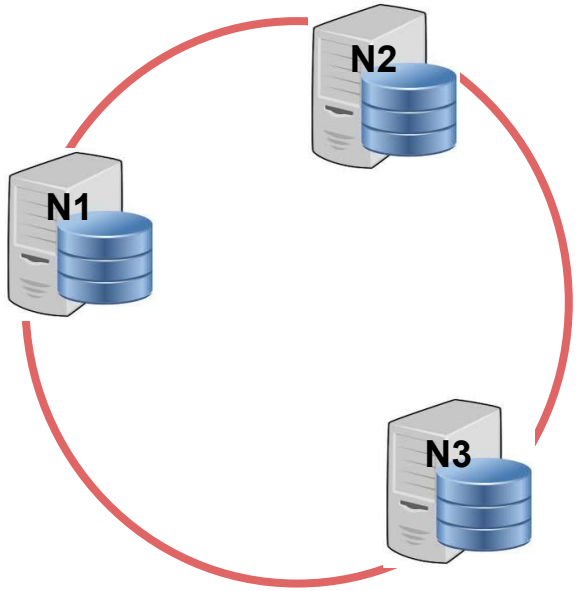

Understanding PXC-5.7

- **Does PXC-5.7 provides all the needed characteristics for good HA ?**
  - **pxc_strict_mode:** blocks non-safe workload
  - **sst/ist:** transparent node joining
  - **quorum:** protection against Split-Brain, Single Node.
  - **unified view:** wsrep_sync_wait
  - **security:** especially with geo-distributed nodes

Understanding PXC-5.7

- **Does PXC-5.7 provides all the needed characteristics for good HA ?**
  - **pxc_strict_mode:** blocks non-safe workload
  - **sst/ist:** transparent node joining
  - **quorum:** protection against Split-Brain, Single Node.
  - **unified view:** wsrep_sync_wait
  - **security:** especially with geo-distributed nodes

YES IT DOES

PXC in different Setups/Topologies

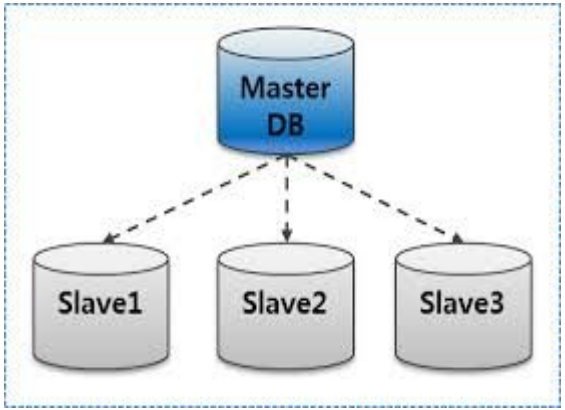

Traditional master slave topology

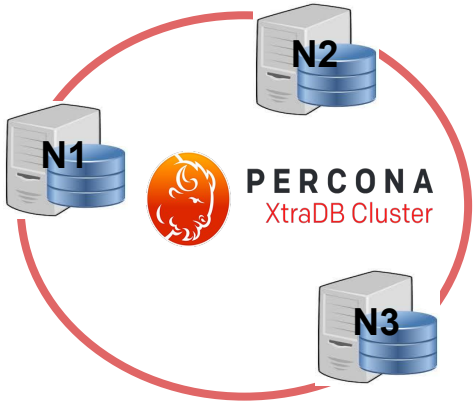

multi-master topology

PXC in different Setups/Topologies

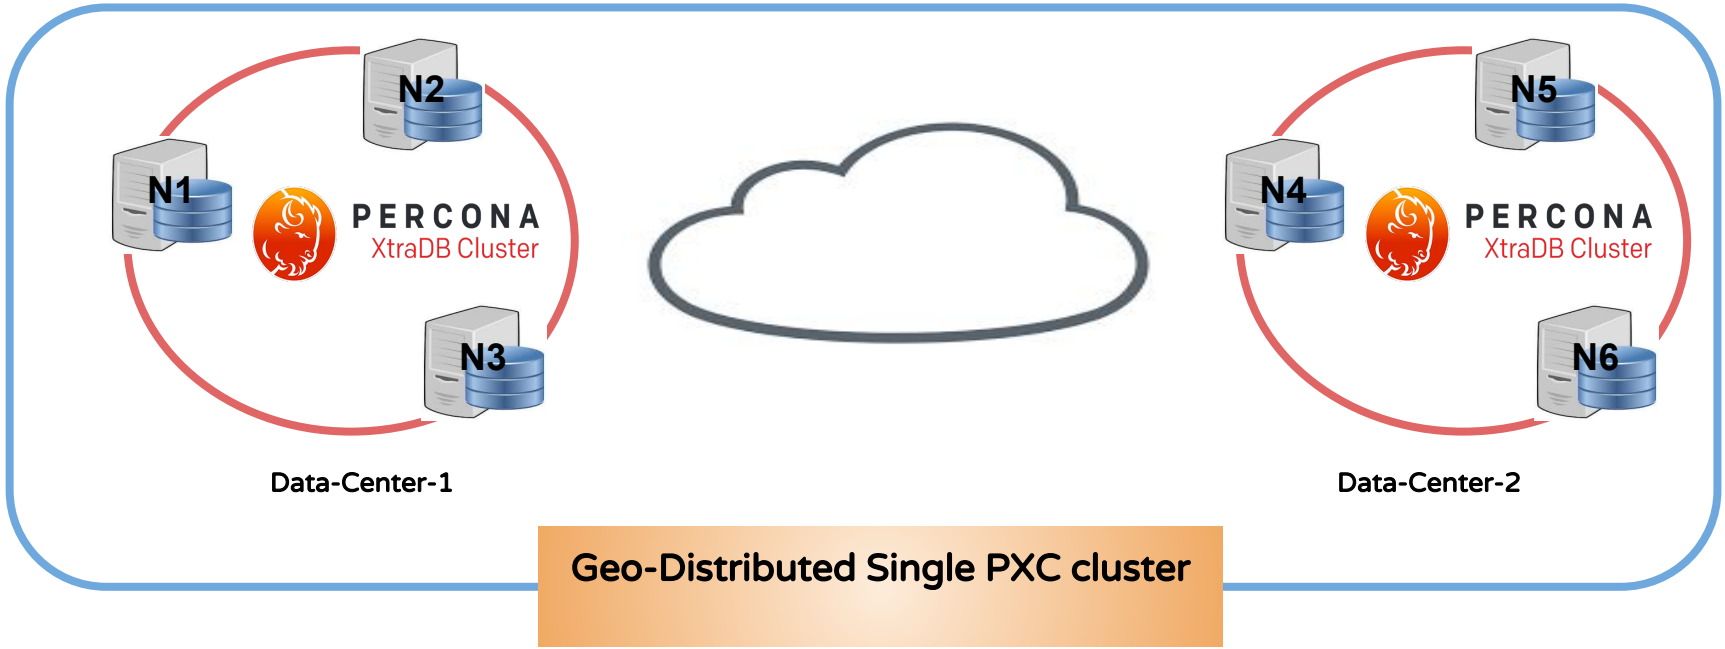

PXC in different Setups/Topologies

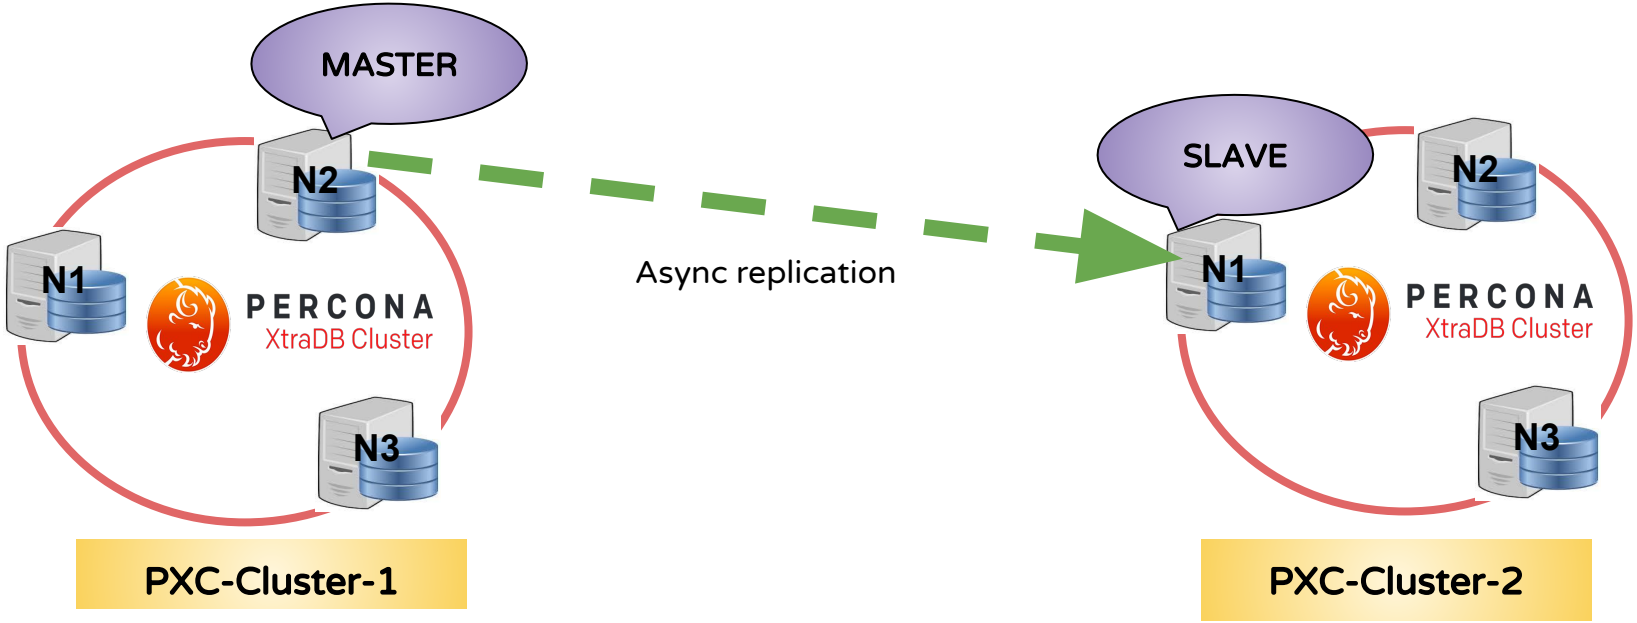

Cluster-Cluster replication

PXC in different Setups/Topologies

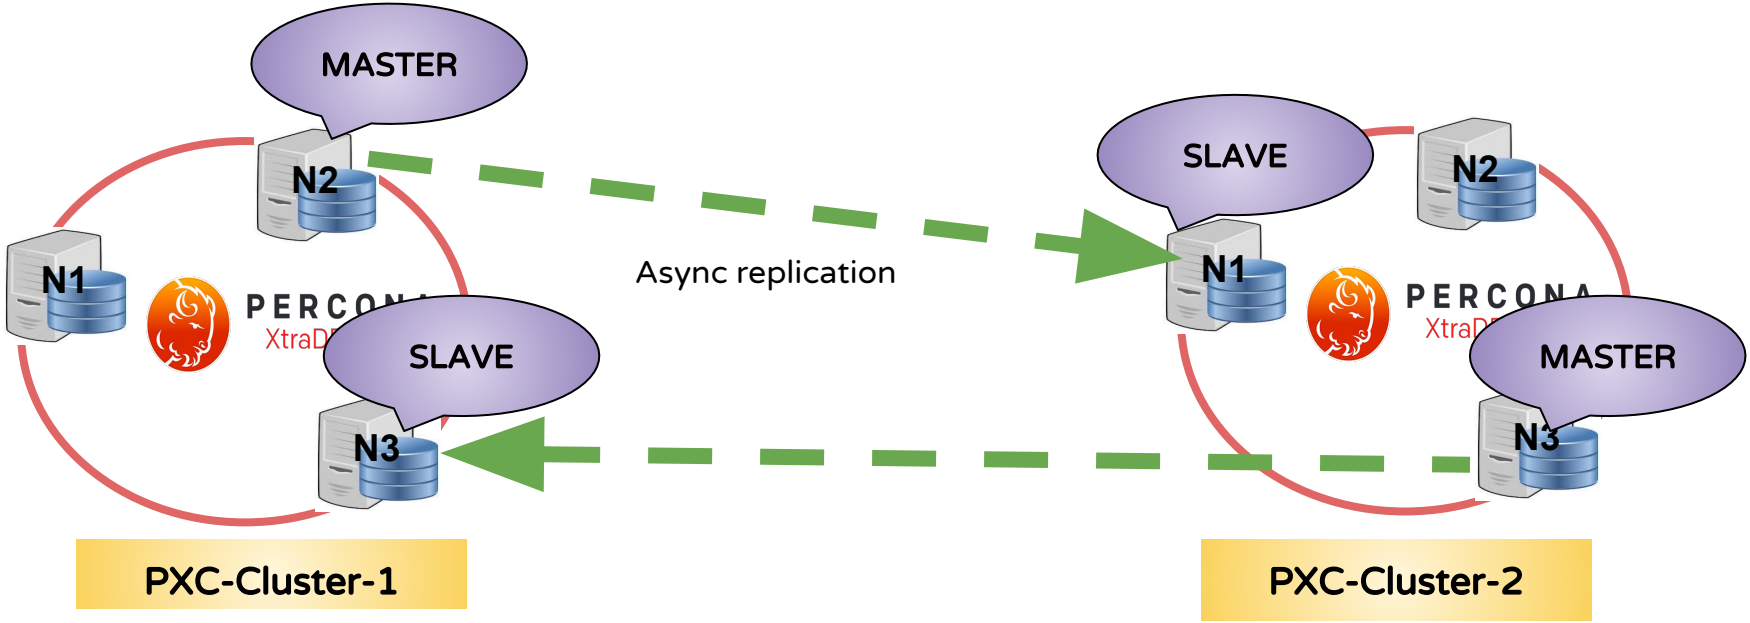

Cluster-Cluster replication

PXC in different Setups/Topologies

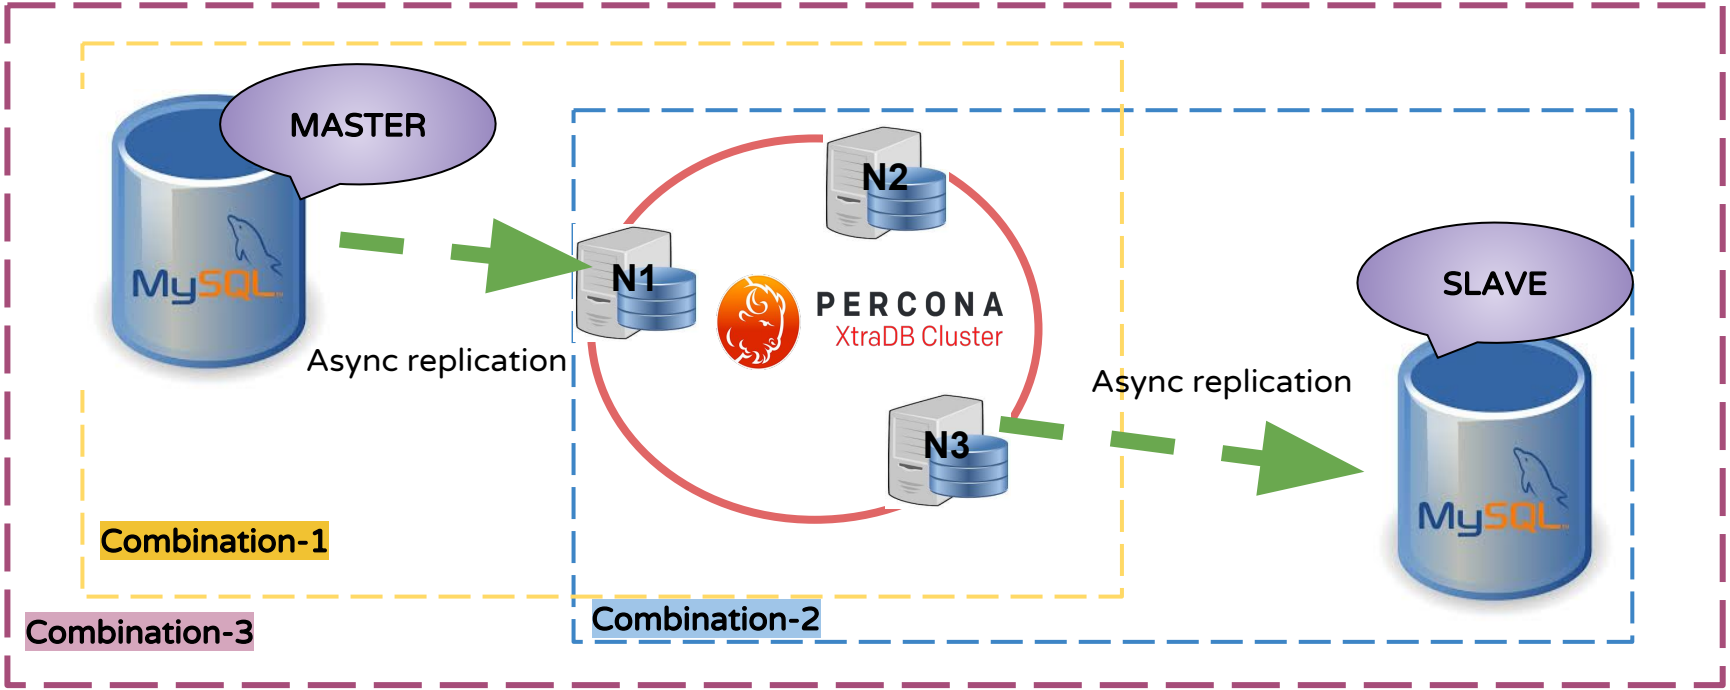

PXC in different Setups/Topologies

- PXC supports and works in different setup/topology environment.
- PXC is flexible and you can easily mix-match PXC cluster with most of the your existing setup.

Understanding ProxySQL

- ProxySQL is a high performance SQL proxy.
- ProxySQL runs as a daemon watched by a monitoring process.
- The monitoring process monitors the daemon and restarts it whenever it crashes (in order to minimize downtime).
- The daemon accepts incoming traffic from MySQL clients and forwards it to backend MySQL servers.

Understanding ProxySQL

- ProxySQL is designed to run on continuous basis without need of to restart.
- In order to allow this, most of configurations can be done at runtime.
- Configuration interface looks alike SQL queries making it simple to grasp for existing DB-user.
- Runtime parameters, server grouping, and traffic-related settings can all be changed at runtime.

Understanding ProxySQL

Configuration :

All configuration parameter changes are made in memory.

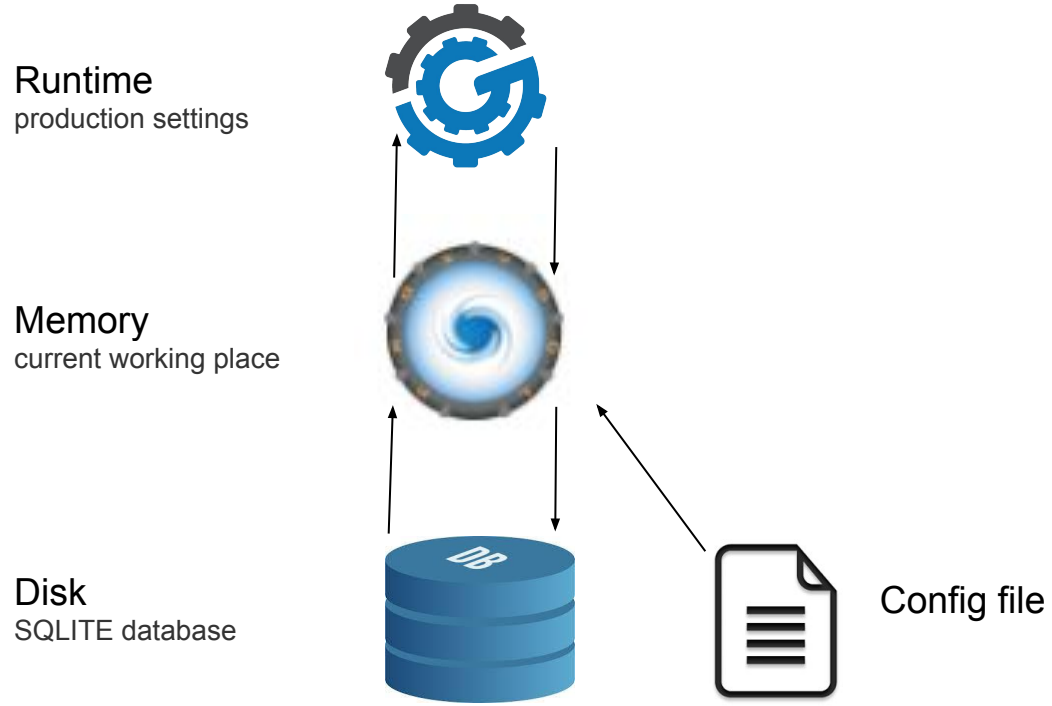

Understanding ProxySQL

Important features:

- Query caching
- Query Routing
- Supports failover
- Cross platform
- Advanced configuration with almost 0 downtime
- Advanced topology support
- Firewall

Understanding ProxySQL

Query caching:

Using advanced ProxySQL rule engine query results can be cached for a configurable timespan, in native MySQL packets format.

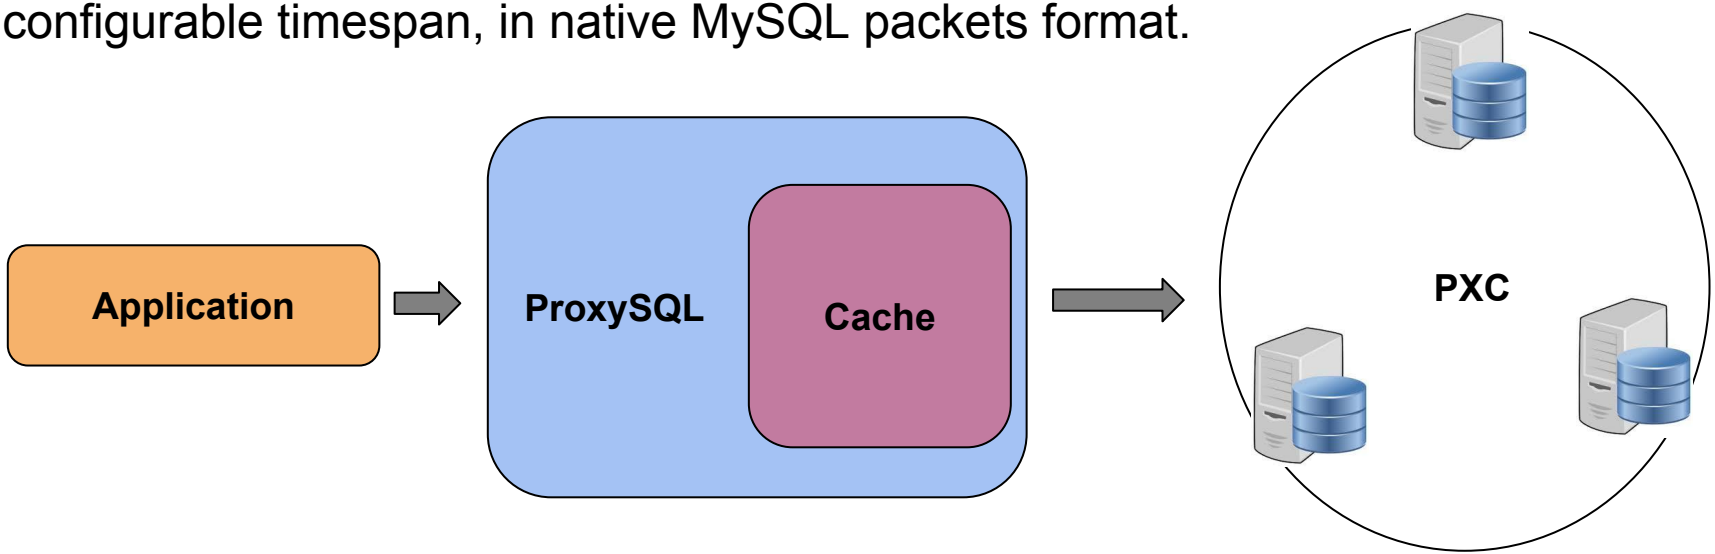

Understanding ProxySQL

Query Routing:

With the help of ProxySQL's hostgroup concept. it is easy to route queries transparently towards the destination cluster that can execute them most efficiently.

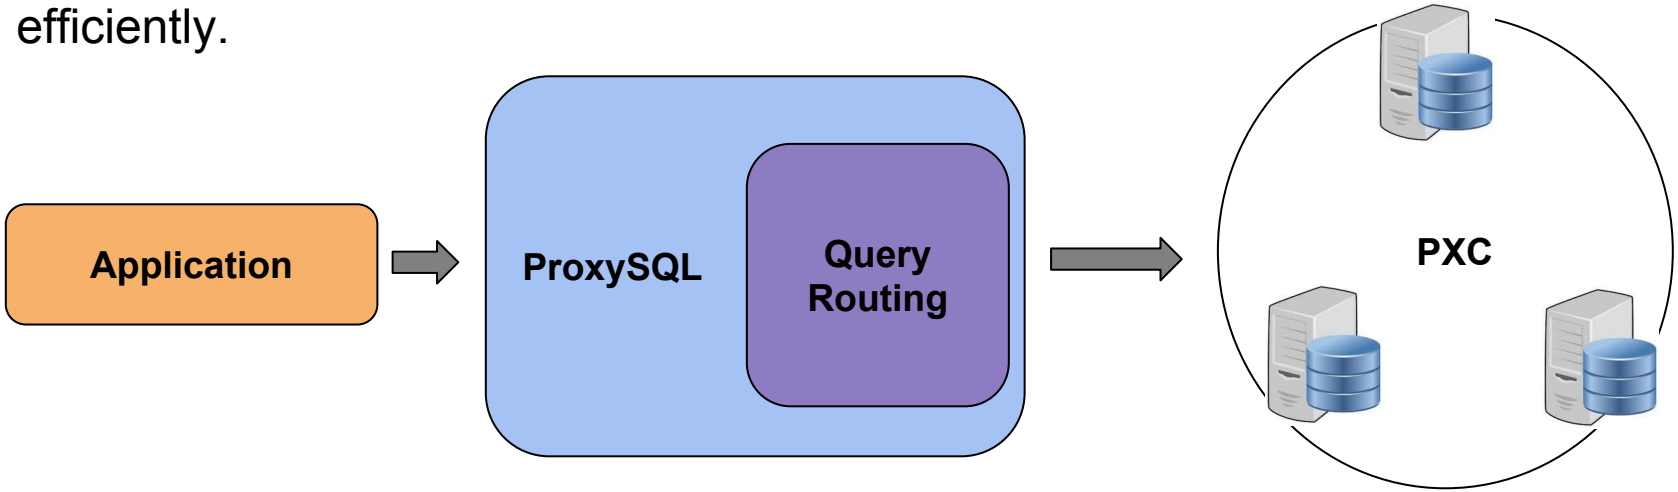

Understanding ProxySQL

Supports failover:

While ProxySQL does not offer support failover as a feature, it collaborates smoothly with the existing tools that enable it. It monitors the health of the backends it communicates with and is able to temporarily shun them based on configurable error rates.

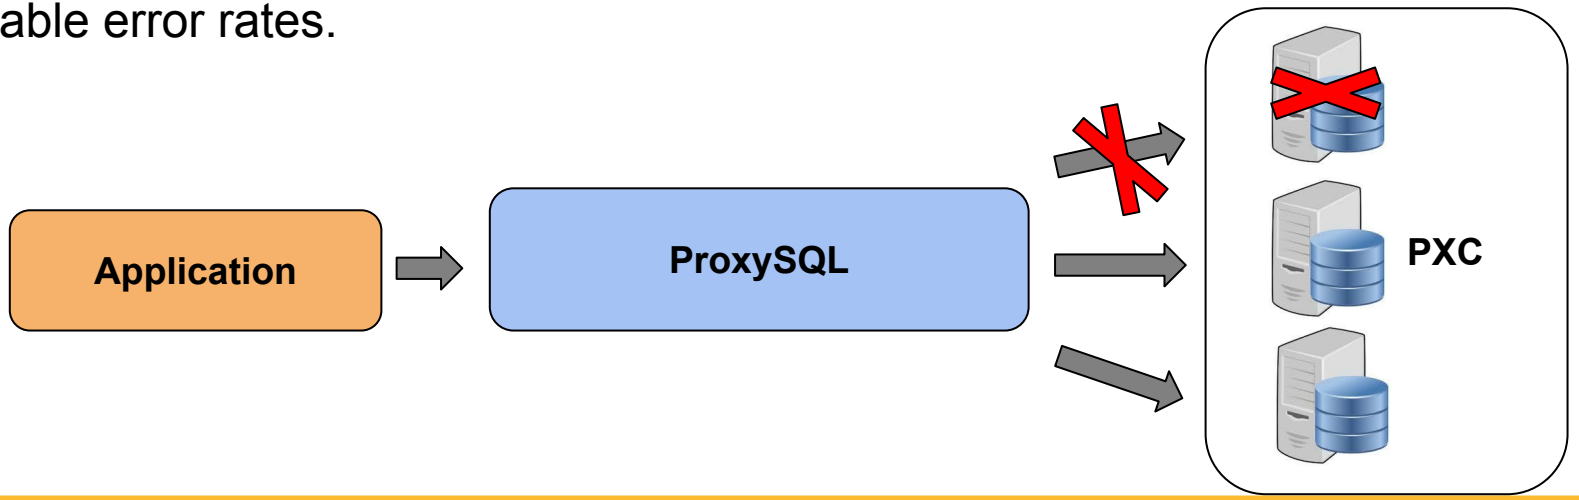

Understanding ProxySQL

Cross platform:

ProxySQL works with multiple flavors of Linux and FreeBSD. Moreover, it should be easily portable to other platforms where a gcc-compatible toolchain is available.

Understanding ProxySQL

Advanced configuration with almost 0 downtime:

ProxySQL's configuration system is inspired from routers. You can dynamically configure everything, persist the configuration and modify it. All with 0 downtime.

Advanced topology support:

Cascading proxies for extra availability and flexibility. Complex MySQL topologies support, involving replication and failover or query mirroring. All of it performed effortlessly by ProxySQL.

Understanding ProxySQL

Advanced topology support:

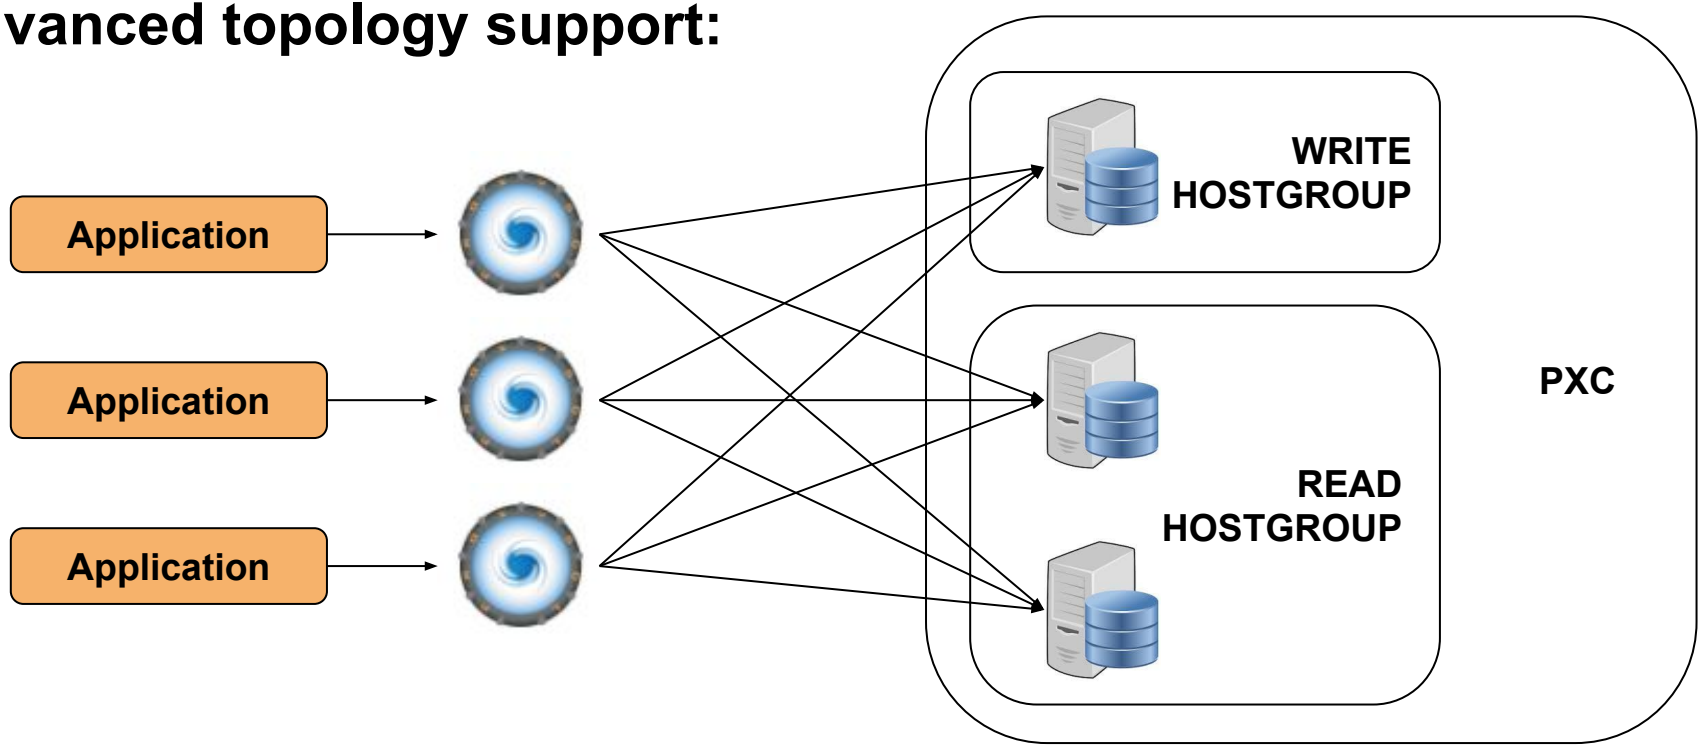

Understanding ProxySQL

Firewall:

In case of offending queries that cause problems to the DB (SQL injection or inefficient retrieval of information via SELECT * without WHERE, for example), ProxySQL acts as a gatekeeper between the application and the DB, allowing DBAs to react quickly.

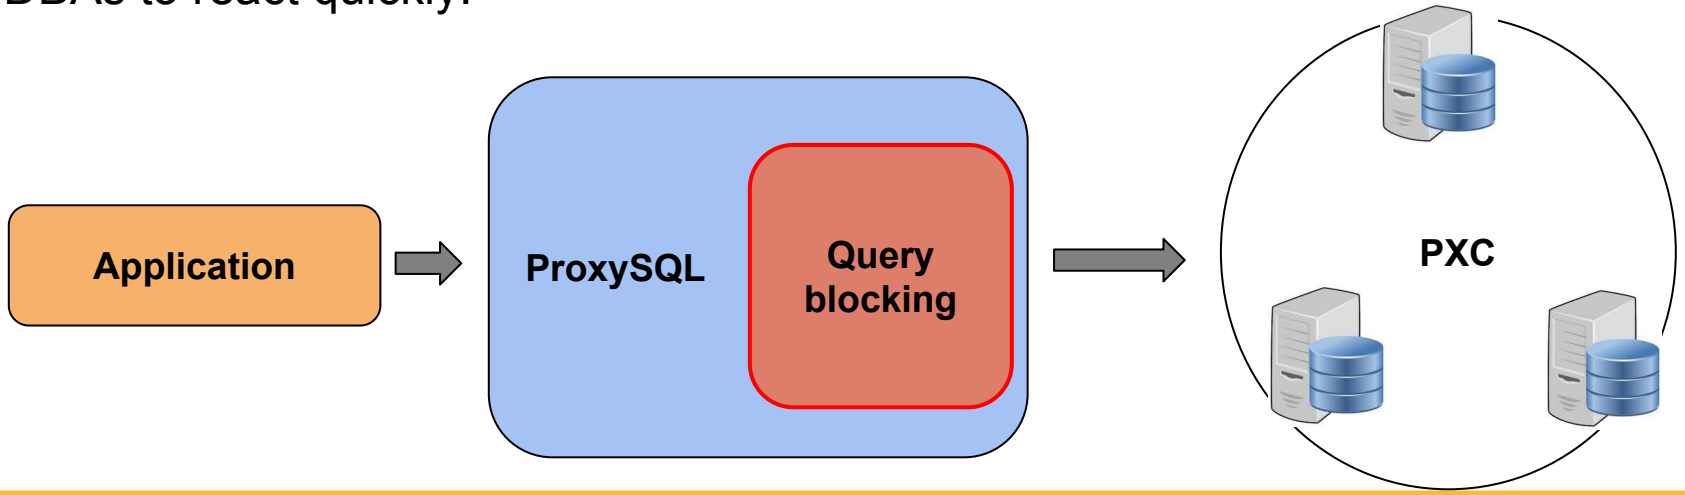

Understanding ProxySQL

- **Does ProxySQL provides all the needed characteristics for good HA ?**
  - Dynamically route traffic on demand.
  - Built-in monitoring.
  - Runtime reconfigurable.
  - Load Balancing.
  - Read/Write split and sharing.

Understanding ProxySQL

- **Does ProxySQL provides all the needed characteristics for good HA ?**
  - Dynamically route traffic on demand.
  - Built-in monitoring.
  - Runtime reconfigurable.
  - Load Balancing.
  - Read/Write split and sharing.

YES IT DOES

PXC + ProxySQL = HA

- PXC and ProxySQL are fully compatible.
- ProxySQL assisted PXC maintenance mode
- ProxySQL provides single-step setup process that can help auto-discover PXC nodes and add them to ProxySQL tracking table. (proxysql-admin tool)
- Customizable monitoring scripts

PXC + ProxySQL = HA

pxc_maint_mode

pxc_maint_transition_period

Transparent rerouting of workload without affecting active workload/application*

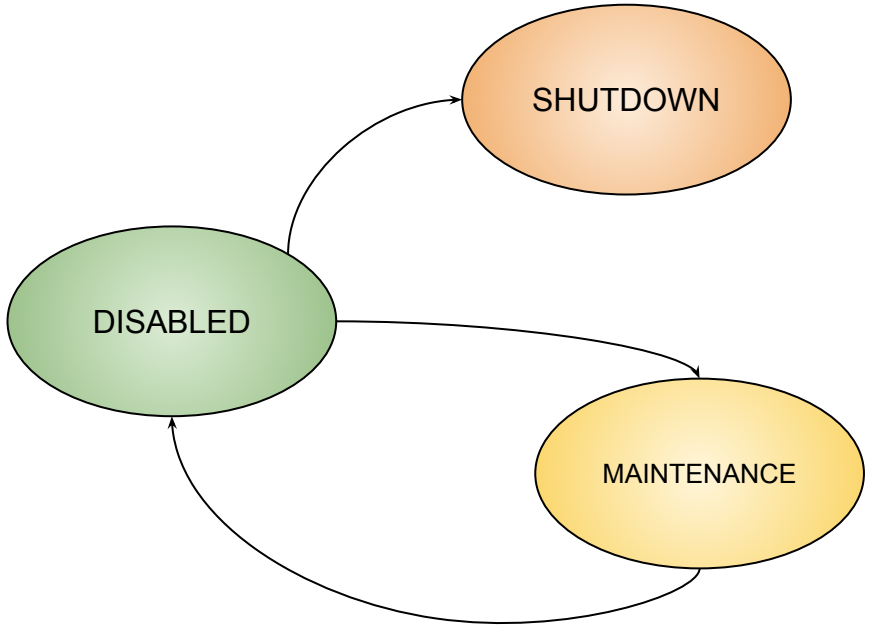

PXC + ProxySQL = HA

- **ProxySQL Admin Tool**

ProxySQL Admin (proxysql-admin) is a powerful tool for configuring Percona XtraDB Cluster nodes into ProxySQL. The proxysql-admin tool comes with the ProxySQL package from Percona apt/yum repositories.

We can configure multiple clusters in proxysql through proxysql-admin tool. ProxySQL 1.4.6 is now available from the Percona repositories.

PXC + ProxySQL = HA

- Proxysql-admin can handle read/write operations in two modes
  - **`singlewrite`** : This mode accepts writes only on one single node. All other remaining nodes will be read-only and accept only read statements.

    With **`--write-node`** option we can control a priority order of what host is most desired to be the writer at any given time. When used the feature will create the config file which is by default "/var/lib/proxysql/host_priority.conf", this is configurable in proxysql-admin.cnf. Servers can be specified comma delimited - 10.0.0.51:3306, 10.0.0.52:3306 - The 51 node will always be in the writer hostgroup if it is ONLINE, if it is OFFLINE the 52 node will go into the writer hostgroup, and if it goes down a node from the remaining nodes will be randomly chosen for the writer hostgroup.

PXC + ProxySQL = HA

- **`loadbal`** : With this mode, read/write transactions will split across all nodes.

You can download proxysql-admin tool from: https://github.com/percona/proxysql-admin-tool
OR ProxySQL package from percona repo.

PXC + ProxySQL = HA

- Customizable monitoring scripts:
    - **proxysql_node_monitor**: checks cluster node membership and re-configures ProxySQL if the cluster membership changes.
    - **proxysql_galera_checker**: checks desynced nodes, and temporarily deactivates them.

Note : These scripts will make sure to setup a good HA solution

Monitoring through PMM

- Important to monitor HA system using a single window monitoring system.
  - PMM helps provide this.
- Fully integrated with PXC and ProxySQL to monitor crucial status allowing single window tracking.
- Track items like
  - Is the node hitting Flow-control
  - What is average write-set packet size
  - Receive/Send queue
  - How long can node sustain w/o doing SST
  - Conflicts

ProxySQL PMM Demo :
https://pmmdemo.percona.com/graph/dashboard/db/proxysql-overview?orgId=1&from=now-1h&to=now&var-interval=$__auto_interval&var-proxysql=pxc57-1&var-hostgroup=11

Monitoring through PMM

PXC + ProxySQL + PMM = HA

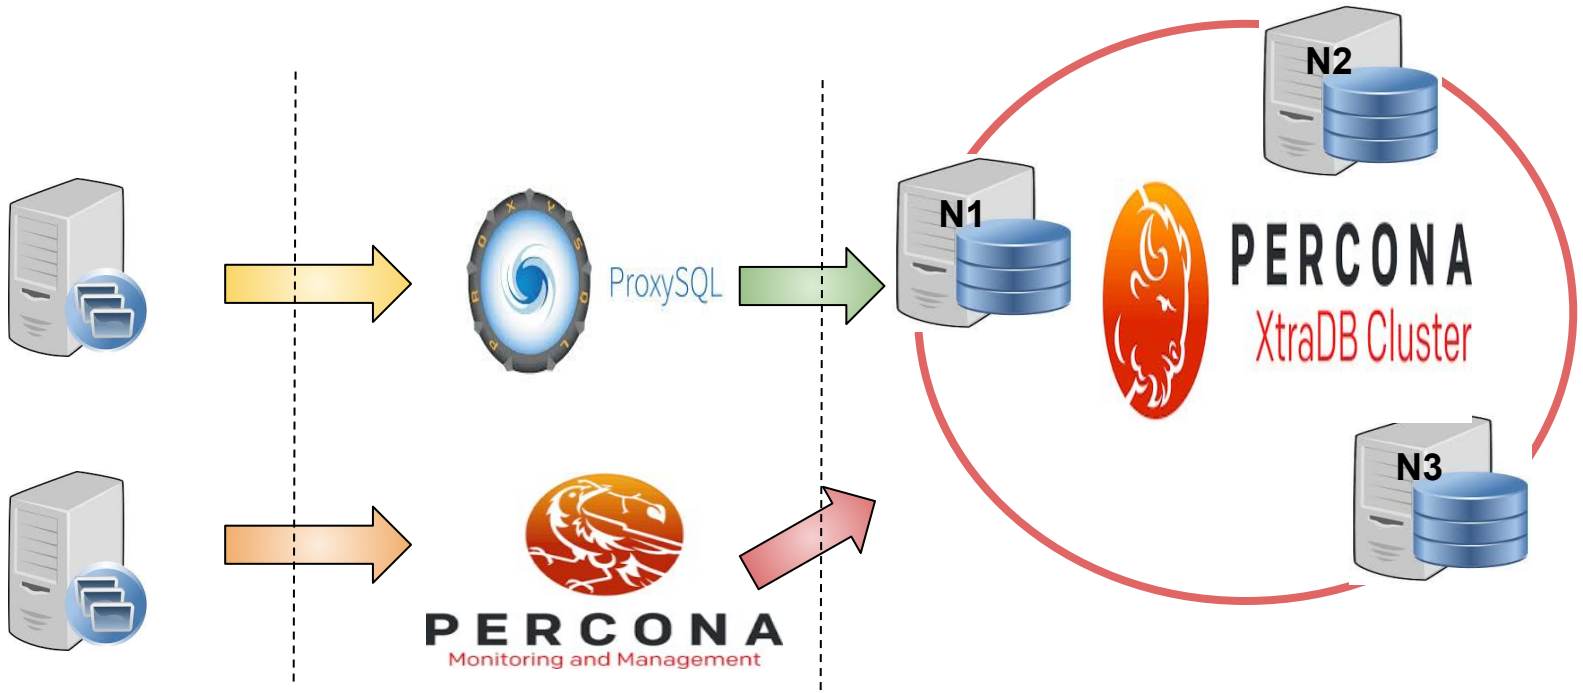

Questions/Comments/Queries

You can also write to us:
krunal.bauskar@percona.com
ramesh.sivaraman@percona.com
Percona XtraDB Cluster Forum(https://goo.gl/pby3EU**)**
ProxySQL-admin issue tracker(https://goo.gl/QDVLBN**)**