

Setting up a PGPool II Cluster

Thursday, 18 March 2021 10:05 (25 minutes)

Databases have to fulfil a variety of requirements in an operational system. They should be highly-available, redundant, suffer minimal downtime during maintenance/upgrade works and be easily recoverable in case of critical system failure.

All of these requirements can be realized with a PGPool II cluster that uses PostgreSQL as backends. The high-availability of the backends are provided by the PGPool II frontend while the redundancy is provided by the PostgreSQL backends. Backups of the backends can be managed via pgBackRest. To further reduce downtimes after critical system failures all nodes are puppetized allowing for easy re-installations if necessary.

This talk will introduce the setup and commissioning phase of the described components and pitfalls identified during our tests.

Speaker release

Yes

Desired slot length

Primary author: HUBNER, Michael (University of Bonn (DE))

Co-authors: WIENEMANN, Peter (University of Bonn (DE)); FREYERMUTH, Oliver (University of Bonn (DE))

Presenter: HUBNER, Michael (University of Bonn (DE))

Session Classification: Basic IT Services

Track Classification: Basic IT Services