# hswaw - Bugless #14

## k0: figure out a better postgres story for high-traffic OLTP uses

02/08/2021 07:53 PM - q3k

| | | |
|---|---|---|
| **Status:** | New | |
| **Priority:** | Normal | |
| **Assignee:** | | |
| **Category:** | hscloud | |

**Description**

We currently setup our postgres instances via kube/postgres.libsonnet, which places them on a single-instance deployment backed in Ceph.

This is fine for simple software, but obviously suboptimal for high traffic usecases:

- ceph eats IOPS for breakfast, so the effective IOPS available to postgres are tiny, thereby limiting our ability to do sustained writes
- recovery from a failed node takes O(minutes) until Kube decides that the node is lost
- the backup story isn't great, as we do ext4 dumps via benji, and these generally are dirty

Some better strategy is needed, either using one of the Well Known Postgres Operatoros, or NIHing our own. We don't even need sharding or autoplacement, just some ability to quickly and reliably fail over from a leader that ended up in a dead/unreachable node.

### History

**#1 - 02/08/2021 07:53 PM - q3k**

*- Description updated*

**#2 - 07/04/2022 01:08 PM - q3k**

*- Category set to hscloud*