

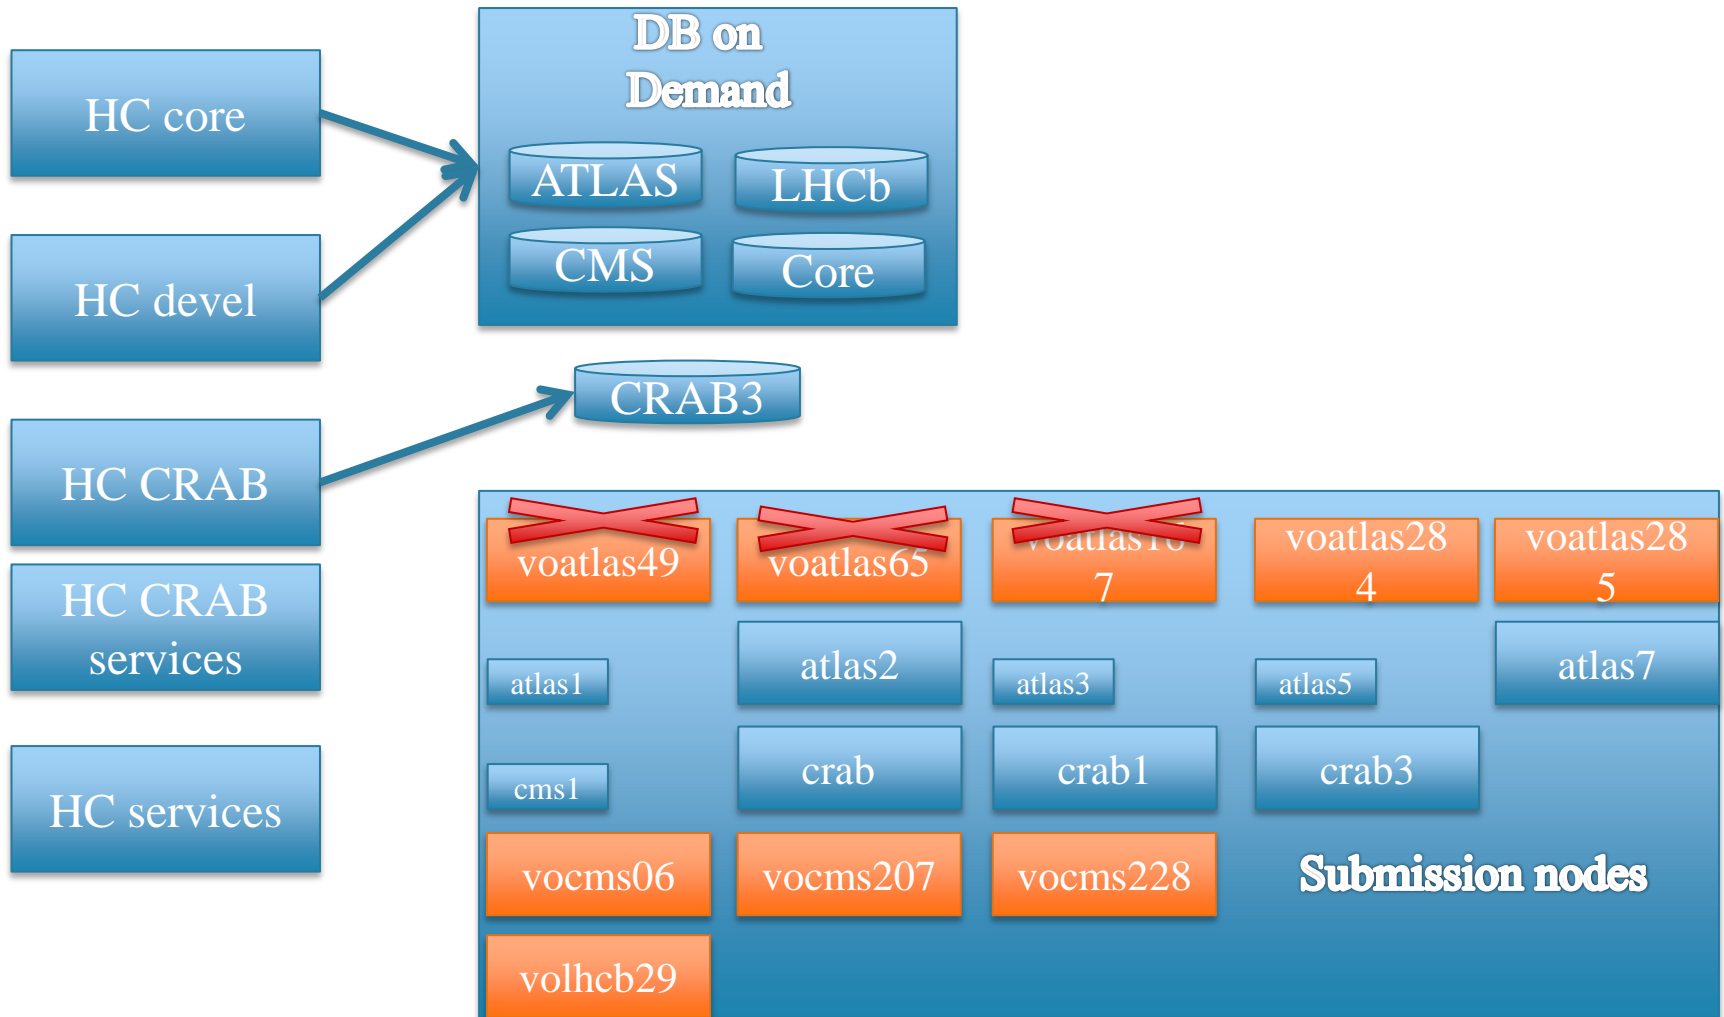


Hammercloud deployment

Pablo Saiz
CERN



Current deployment



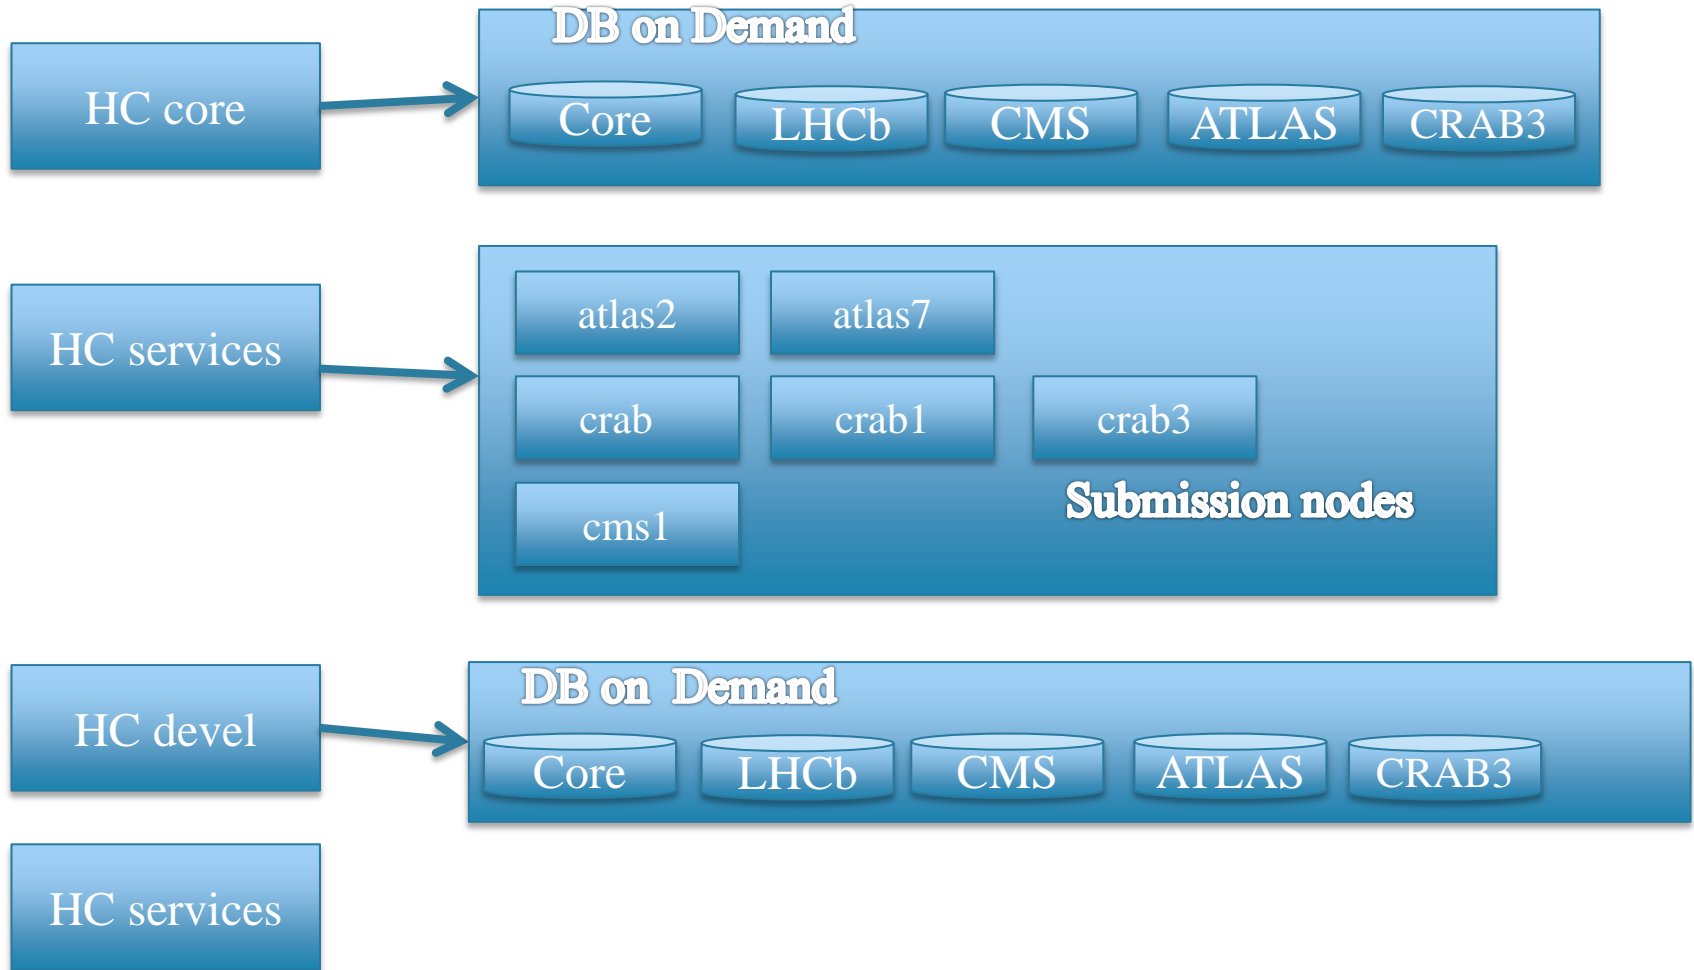
Overview of current setup

- Clear distinction of roles
- Multiple nodes for submission
 - With different sizes
 - Some of them in quattor
- Development machine for web UI
 - Using same DB
- Using openstack project
 - With multiple images: SLC5, SLC6, HC v4, atlas-snap
- Software installed on core machine
 - Using git checkout
 - And some manual modifications!
 - Nfs-mounted on submission nodes
 - Very fast deployment cycle
- Using cvmfs for PANDA dependencies

Possible improvements

- Decoupling test/production
 - Like done for CRAB3
 - (Decoupling of VO?)
- All submission nodes on AI
 - Use puppet, and default images
 - Extra configuration in manifests
 - Hostgroups for core/services/submission node
- Install software on submission nodes
 - Avoid nfs mounted area
- Split Hammercloud core and VO specific repo
- Automate release procedure
 - Introduce Koji as build system
 - Use standard ai repo (**Latency**: 15m ai-testing, 4h ai-stable)
 - Use automatic versioning number

Suggested deployment



Git repo

- Current structure:
 - Single repo: hammercloud (1GB)
- Suggested:
 - Hammercloud-core (6.5 MB)
 - Plus MANIFEST.in, setup.cfg, setup.py, post-install
 - Hammercloud-atlas (1GB)
 - Apps/atlas + web/[templates/src/media]/atlas
 - (apps/atlas/inputfiles/templates/HWWNtupleCode-00-02-07/test/HWWlvivCode/Run/__panda_rootCoreWorkDir 170 MB!)
 - Hammercloud-cms (1 MB)
 - Apps/cms + web/[templates/src/media]/cms
 - Each repo also contains:
 - Files for building rpm with python distutils: MANIFEST.in, setup.py, setup.cfg, post_install