

CES Release Notes

- 1.0.0

1.0.0

HA Cluster

The master/slave architecture of the cluster allows it to scale horizontally by introducing more workers and also ensures availability. We call the master 'Coordinator' and the slave nodes 'Workers'.

Coordinator

- Load balances /extract requests
- Load balances dictionary updates
- Manages update feeds
- Manages worker dictionary queries

Worker

- Keeps track of its dictionaries fingerprint
- Creates a pool of pipelines
- Serves extraction requests
- Ability to update it's Gazetteer and Metadata PR SPARQL queries
- Ability to reload the whole pipeline dictionary

Dynamic updates

- Multiple retries on failed update, full dictionary reload is triggered after a limit is hit (configurable)
- EUF plug-in keeps track of updated entities and serves changelists through special SPARQL queries
- Updates are kept track of via GraphDB's fingerprints