Technical guide

Full LDAP synchronisation for ENTs

Script architecture and operation

All the command scripts are located in the component's /cli directory. These scripts are designed to optimise mass synchronisations on a large number of Moodle instances in parallel (VMoodle), notably in academic installations. They are designed to optimise synchronisation processing time by spreading it across several server processors. These scripts are of three kinds:

- Synchronisation task scripts
- “Workers”
- “Mass launchers” (bulks)

The synchronisation tasks are all compatible with VMoodle virtualisation.

The synchronisation scripts

They run a simple synchronisation on one of the supported business entities on a given Moodle host (by default, the main installation).

The workers

They chain a series of synchronisation tasks over a specified series of hosts.

The mass launchers (bulks)

They distribute the hosts to synchronise over workers. When the distributed processing option is used, the workers are launched in parallel in order to benefit from a multi-processor architecture.

The mass launcher `bulk_sync_hosts.php` is a super launcher that processes all the active business entities one after the other (users, role assignments, cohorts and course groups)

An example of a processing chain

```
bulk_sync_cohorts.php ==> sync_cohorts_worker.php ==> sync_cohorts.php
```

Configuration

The ENT Installer configuration is an essential aspect of using this component. Indeed the main use of this component is the regular automated synchronisation of Moodle's main business entities based on these configurations.

Data source

The data is extracted from an LDAP directory as defined by the LDAP authentication plugin, **even if it is not used or enabled** for user authentication. If another authentication method with LDAP settings is used, the same parameters must be copied into the auth/ldap plugin configuration for the ENT_Installer's use.

General principle of the configurations

An LDAP directory is a data tree made of data records classified in a tree. Each node of the tree can have its own attributes at most. Moreover an attribute can contain a simple or composite piece of data that itself has a certain construction format.

The data definitions and schemas are generally designed by the IT teams, with standards and general concerns of IT urbanisation. They are not, in general, dedicated to any particular application, but are rather organised and formatted to serve several heterogeneous user applications.

There is therefore a priori a chance that the directory data does not directly match the formats Moodle needs to be fed.

the ENT installer works on a general processing basis that can be split into 5 concepts:

- Find the data (the attribute) in the LDAP
- Find the significant portion of value in the attribute content
- Optionally establish a correspondence between these extracted significant values and values expected by Moodle
- Find which part of the data (translated or not) is needed by Moodle
- Find where the data must be stored in Moodle

General aspects

- General parameters and behaviour options
- General precautions on the LDAP directory
- Additional LDAP reading API

Synchronised entities

The main Moodle business entities are synchronised by a task associated with a command-line CLI script.

- User account synchronisation
- Course and course category synchronisation
- Cohort synchronisation
- Course group synchronisation
- Role assignment synchronisation:todo:

Related documentation

- Injection of the initial categories

Capabilities

- Capabilities guide

Back to the component index - Back to plugins index

From:
https://docsen.activeprolearn.com/ - **Documentation Moodle ActiveProLearn**

Permanent link:
https://docsen.activeprolearn.com/doku.php?id=local:entinstaller:technique

Last update: **2026/11/13 20:37**