

Unix Cluster Plugin

The Unix Cluster plugin manages Unix Accounts for CO People.

⚠ This plugin is considered Experimental.

- [Installation](#)
- [Configuration](#)
 - [Setting a Default Group For Each Unix Account](#)
 - [Understanding Identifier Management](#)
 - [Mapping CO Groups to Unix Cluster Groups](#)
- [Account Status](#)
- [Supported Provisioners](#)
- [User Self Service](#)

Installation

1. This is a non-core plugin, see [Installing and Enabling Registry Plugins](#) for more information.

Configuration

The Unix Cluster Plugin can be instantiated multiple times, with each instantiation representing a single (logical) Unix Cluster. A CO Person may then have one or more Unix Cluster Accounts within each Unix Cluster.

Unix Accounts may be created completely manually, in which case no particular rules are applied, or they may be created automatically, in accordance with the Unix Cluster configuration:

- **Sync Mode:** Whether or not updates to the CO Person record propagate to the Unix Cluster Account record after initial creation. See *Understanding Identifier Management* below for more details.
- **Username Type:** When a Unix Account is created, the username will be obtained from the CO Person's Identifier of the specified type. This identifier might have been automatically generated via [Identifier Assignment](#). It is recommended, but not required, to use the *Network* Identifier type.
- **UID Type:** When a Unix Account is created, the numeric UID will be obtained from the CO Person's Identifier of the specified type. The identifier must be a positive integer. This identifier might have been automatically generated via [Identifier Assignment](#), in which case a format of (#) with a *Minimum Value* of 1000 (or as otherwise appropriate for the local deployment) is recommended. It is recommended, but not required, to use the *UI* Identifier type.
- **Default Login Shell:** Unix Accounts will by default have the specified shell. This is also the default value used during manual account creation.
- **Home Directory Prefix:** This string will be suffixed with a slash (/) and the username (as specified above) to construct the Unix Account's home directory. For example, the prefix `/home` would generate a home directory like `/home/jsmith`. Note the Unix Account Plugin does not create home directories, it is expected that either a Provisioner Plugin or automatic home directory creation will be used.
- **Home Directory Subdivisions:** Home directories can be sorted into subtrees based on the username, up to the specified number of subdivisions. For example, at level 2 the home directory would become `/home/j/s/jsmith`.
- **Group Name Type:** When a CO Group is mapped to a Unix Cluster Group (see below), the CO Group Identifier of this type will be used as the group's Unix name. The identifier should consist of characters and length suitable for the target servers.
- **GID Type:** When a CO Group is mapped to a Unix Cluster Group (see below), the CO Group Identifier of this type will be used as the group's numeric ID. The identifier must be a positive integer, and a minimum value of 1000 is recommended.
- **Default Group:** Unix Accounts will by default be placed into the Unix Cluster Group (see below) associated with this Unix Cluster Group.
 - ⓘ In order to be available for use as a Default Group, a CO Group must first be mapped to a Unix Cluster Group, as described below.
 - A Default Group can also be created for each Unix Account, see below for more information.

Note the GECOS field is currently populated using the Primary Name of the CO Person. The generation of this field is subject to change in a future release.

The following situations will cause the generation of the Unix Account to fail:

- If a Username is not found in the CO Person record in accordance with the configuration
- If the Username is already associated with another Unix Account within this Unix Cluster
- If a UID is not found in accordance with the configuration
- If the UID is already associated with another Unix Account within this Unix Cluster
- If a Default Group cannot be assigned, because
 - The CO Person is not a member of the Default CO Group (if set)
 - A CO Person-specific CO Group could not be created

Note that automatically generating a Unix Account can only create one unique Unix Account for a CO Person based on the configuration. Any additional Unix Accounts should be created manually, or via custom scripting (using a plugin or the REST API).

Setting a Default Group For Each Unix Account

By setting *Default Group* to *Create a New Group For Each Account*, the following behavior will take place:

- The Plugin will look for a CO Group with an identifier of *Group Name Type* with a value equal to the *Username Type* and, if found, assign that as the default group. For example, for the user `jsmith` the Plugin will attempt to find a group with *Group Name Type* identifier of `jsmith`. Note this is an Identifier attached to the CO Group, and not the name of the CO Group.

- If no such CO Group is found, a new CO Group will be created as follows:
 - The CO Group name will be constructed using the language key `pl.unixcluster.fd.co_group_id.new.name`
 - An Identifier of *Group Name Type* will be created, with the same value as the CO Person's *Username Type*
 - An Identifier of *GID Type* will be created, with the same value as the CO Person's *UID Type*
 - ⚠ In order to ensure GIDs can be successfully assigned to matching UIDs, GIDs for other groups should be assigned in a different numeric space. As an example, if a CO Person has a UID of 1001, and a GID for the "Users" CO Group was also assigned 1001, then automatic GID assignment for the CO Person will fail, since the GID 1001 is not available.
 - The CO Person will be added as both Member and Owner of the new CO Group
 - The CO Group will be added as a Unix Cluster Group (see below)

Understanding Identifier Management

Username and UIDs must be unique within a Unix Cluster.

Although not required, it is recommended that Usernames and UIDs be consistent across Unix Clusters to simplify management. (An exception might be for managing legacy clusters.) Keep in mind that Identifiers must be unique within a CO (within all CO People or CO Groups as appropriate). In order for (eg) two Unix Cluster Groups in different Unix Clusters to use the same GID Number (with different names), each Unix Cluster must have a different GID Type.

When a Unix Cluster Account is created automatically, subsequent management of these identifiers depends on the Unix Cluster's *Sync Mode*. More specifically, when an Account is automatically created it is given the Sync Mode setting of the Cluster, and so the individual Account's Sync Mode setting governs subsequent behavior. Changing the Cluster's setting will not affect the settings of existing Accounts.

If the Sync Mode is set to *Manual*, then there is no subsequent automatic linkage between Identifiers attached to CO Person records and identifiers stored within a Unix Cluster Account record. However, if the Sync Mode is set to *Track CO Person Attributes* then changes made to the CO Person record will also update the associated Unix Cluster Account record as follows:

- An update to the CO Person's Primary Name will cause an update to the Unix Cluster Account's GECOS field.
- An update to a CO Person Identifier of type *UID Type* will cause an update to the Unix Cluster Account's UID field.
- An update to a CO Person Identifier of type *Username Type* will cause an update to the Unix Cluster Account's Username field, **and** cause a recalculation of the Account's Home Directory.
- If the Default Group setting is *Create a New Group For Each Account*, then
 - An update to a CO Person Identifier of type *UID Type* will cause the CO Group Identifier of type *GID Type* to also be updated, but only if it has the same value as the original CO Person Identifier value.
 - An update to a CO Person Identifier of type *Username Type* will cause an update to the CO Group Name, **and** cause the CO Group Identifier of type *Group Name Type* to also be updated, but each update only happens if the original has the same value as the original CO Person Identifier value.

These automatic changes are made even if the Unix Cluster Account is currently suspended. If an update results in a conflict (eg: the new Identifier is already in use by another Unix Cluster Account) the automatic update will fail.

Additionally, manually assigning an identifier (username or UID) on a Unix Account record will not prevent the same Identifier from being created on a CO Person record. For example, if the UID 1500 is manually assigned to an account, and that would be the next sequential identifier assigned for a new CO Person record, if that new CO Person record is subsequently made eligible for a Unix Account, the Unix Account auto-generation will fail as described above.

Mapping CO Groups to Unix Cluster Groups

While Unix Accounts are attached to (and managed via) CO Person records, Unix Cluster Groups are managed directly via the Unix Cluster configuration, via *Manage Unix Cluster Groups*. A CO Group becomes available to a Unix Cluster when

1. It is attached, via *Manage Unix Cluster Groups*
2. The CO Group has identifiers of the types (*Group Name Type* and *GID Type*) required by the configuration

A Unix Account will be a member of a given Unix Cluster Group when the CO Person associated with the Unix Account is a CO Group Member of the CO Group that, in turn, is associated with the Unix Cluster Group. Note that this means all Unix Accounts associated with the same CO Person will be members of the same Unix Cluster Groups on a given Unix Cluster.

Account Status

Unix Cluster Accounts currently support a limited set of statuses. Accounts in Active or GracePeriod statuses will be made available to provisioners, but otherwise status is ignored.

Supported Provisioners

The [LDAP Provisioning Plugin](#) supports writing Unix Cluster Account information to various schema.

User Self Service

User Self Service (eg: to change a shell or GECOS) is not currently supported.