pudl Documentation

Release 0.1.3

zulily, llc

December 14, 2015

Contents

pudl version: 0.1.3

`pudl` is a python package that wraps python-ldap and provides a somewhat-oo interface to Active Directory user, group and computer objects, retreived via LDAP with TLS. While not necessarily a replacement for existing client libraries and the ldapsearch binary, the api and bundled cli are perhaps simpler to work with than alternatives, for many common queries.

Contents:

CHAPTER 1

Getting Started

1.1 Installation

1.1.1 Prerequisites

To get up and running, the following must be installed:

- python 2.7.x
- python-dev
- libsasl2-dev
- libldap2-dev
- libyaml-dev

1.1.2 pip

From the top-level directory of the cloned repository:

```
pip install .
```

Note: This is typically performed with an active python virtual environment.

And for *optional* document generation with sphinx, install the following python packages as well:

```
pip install mock
pip install sphinx
pip install pygments
pip install sphinx_rtd_theme
pip install sphinx-argparse
```

1.2 TLS

pudl only communicates with Active Directory over TLS and by default, the remote server must meet strict criteria such as the commonname matching the hostname in the LDAP URL, and the remote server must present a certificate that is signed by a trusted authority.

Adding a trusted CA to a Linux system varies by distribution, but is rather simple with ubuntu:

1. Copy the CA certificate (.CER format) to /usr/share/ca-certificates/extra/
2. chown the CA certificate to root.root and chmod it to 444
3. Run sudo dpkg-reconfigure ca-certificates

If the only requirement is an encrypted channel and not verifying the identity of the remote server, the ADQuery constructor takes a tls_no_verify parameter, and the pudl CLI has a –tls-no-verify argument. Use these options with caution.

1.3 Basic Usage

Pull a user object (bhodges) and a few attributes to print out:

```
#! /usr/bin/env python

import sys
from pudl import *
from pudl.ad_query import ADQuery
from pudl.ad_user import ADUser

BASE_DN = 'OU=Departments,DC=example,DC=com'
LDAP_USER = 'jdupont'
PASSWORD = 'my_secret'

def main():
    """
    """
    adq = ADQuery(user=LDAP_USER, password=PASSWORD)
    adu = ADUser(adq)
    users = adu.users(base_dn=BASE_DN, attributes=['samaccountname', 'cn',
                      'title',], samaccountnames=['bhodges',])

    for user in users:
        print '{0}: {1}, {2}'.format(user.samaccountname, user.cn, user.title)

if __name__ == '__main__':
    sys.exit(main())
```

For additional usage examples, reviewing the pudl *cli* source is recommended

1.4 Why is this package named pudl?

While pronounced like “puddle”, the name is loosely related to ‘pudl’.replace(‘u’, ‘a’)[::-1].

CHAPTER 2

API Reference

2.1 ADQuery

class pudl.ad_query.**ADQuery**(*user*, *password*, *ldap_url='ldap://ldap:389'*, *tls_no_verify=False*, *page_size=300*)

Bases: object

Query Active directory with python-ldap. May be used directly, but is most commonly used indirectly via ADObject-based classes. All connections require TLS.

search(*base_dn*, *search_filter*, *attributes=()*)

Perform an AD search

Parameters

- **base_dn** (*str*) – The base DN to search within
- **search_filter** (*str*) – The search filter to apply, such as: *objectClass=person*
- **attributes** (*list*) – Object attributes to populate, defaults to all

2.2 ADUser

class pudl.ad_user.**ADUser**(*adq*)

Bases: pudl.ad_object.ADObject

A class to represent AD user objects. Includes a number of helper methods, particularly object-factory related.

ADUser objects have minimal depth, with attributes set to strings or lists. Available attributes are dependent on the results returned by the LDAP query.

group_samaccountnames(*base_dn*)

For the current ADUser instance, determine which groups the user is a member of and convert the group DistinguishedNames to sAMAccountNames. The resulting list of groups may not be complete if explicit_membership_only was set to True when the object factory method (user() or users()) was called.

Parameters **base_dn** (*str*) – The base DN to search within

Returns A list of groups (sAMAccountNames) for which the current ADUser instance is a member, sAMAccountNames

Return type list

is_member(*group_distinguishedname*)
For the current ADUser instance, determine if the user is a member of a specific group (the group DN is used). The result may not be accurate if explicit_membership_only was set to True when the object factory method (user() or users()) was called.

Parameters **group_distinguishedname** (*str*) – The group DistinguishedName

Returns A boolean indicating whether or not the user is a member of the group

Return type bool

samaccountname(*base_dn*, *distinguished_name*)
Retrieve the sAMAccountName for a specific DistinguishedName

Parameters

- **base_dn** (*str*) – The base DN to search within
- **distinguished_name** (*list*) – The base DN to search within
- **attributes** (*list*) – Object attributes to populate, defaults to all

Returns A populated ADUser object

Return type *ADUser*

samaccountnames(*base_dn*, *distinguished_names*)
Retrieve the sAMAccountNames for the specified DNs

Parameters

- **base_dn** (*str*) – The base DN to search within
- **distinguished_name** (*list*) – A list of distinguished names for which to retrieve sAMAccountNames

Returns Key/value pairs mapping DistinguishedName to sAMAccountName

Return type dict

to_dict()
Prepare a minimal dictionary with keys mapping to attributes for the current instance.

user(*base_dn*, *samaccountname*, *attributes=()*, *explicit_membership_only=False*)
Produces a single, populated ADUser object through the object factory. Does not populate attributes for the caller instance.

Parameters

- **base_dn** (*str*) – The base DN to search within
- **samaccountname** (*str*) – The user's sAMAccountName
- **attributes** (*list*) – Object attributes to populate, defaults to all
- **explicit_membership_only** (*bool*) – If set True, memberof will only list groups for which the user is a directly referenced member

Returns A populated ADUser object

Return type *ADUser*

users(*base_dn*, *samaccountnames=()*, *attributes=()*, *explicit_membership_only=False*)
Gathers a list of ADUser objects

Parameters

- **base_dn** (*str*) – The base DN to search within

- **attributes** (*list*) – Object attributes to populate, defaults to all
- **samaccountnames** (*list*) – A list of usernames for which objects will be created, defaults to all users if unspecified
- **explicit_membership_only** (*bool*) – If set True, memberof will only list groups for which users are directly referenced members

Returns A list of populated ADUser objects

Return type list

2.3 ADGroup

class pudl.ad_group.**ADGroup**(*adq*)
Bases: pudl.ad_object.ADObject

A class to represent AD group objects. Includes a number of helper methods, particularly object-factory related.

ADGroup objects have minimal depth, with attributes set to strings or lists. Available attributes are dependent on the results returned by the LDAP query.

In its current implementation, the memberOf attribute is not expanded. The member attribute is however flattened out.

group(*base_dn*, *samaccountname*, *attributes=()*, *explicit_membership_only=False*)
Produces a single, populated ADGroup object through the object factory. Does not populate attributes for the caller instance.

sAMAccountName may not be present in group objects in modern AD schemas. Searching by common name and object class (group) may be an alternative approach if required in the future.

Parameters

- **base_dn** (*str*) – The base DN to search within
- **samaccountname** (*str*) – The group's sAMAccountName
- **attributes** (*list*) – Object attributes to populate, defaults to all

Returns A populated ADGroup object

Return type *ADGroup*

groups(*base_dn*, *samaccountnames=()*, *attributes=()*, *explicit_membership_only=False*)
Gathers a list of ADGroup objects

sAMAccountName may not be present in group objects in modern AD schemas. Searching by common name and object class (group) may be an alternative approach if required in the future.

Parameters

- **base_dn** (*str*) – The base DN to search within
- **samaccountnames** (*list*) – A list of group names for which objects will be created, defaults to all groups if unspecified
- **attributes** (*list*) – Object attributes to populate, defaults to all

Returns A list of populated ADGroup objects

Return type list

samaccountname (*base_dn*, *distinguished_name*)
Retrieve the sAMAccountName for a specific DistinguishedName

Parameters

- **base_dn** (*str*) – The base DN to search within
- **distinguished_name** (*list*) – The base DN to search within
- **attributes** (*list*) – Object attributes to populate, defaults to all

Returns A populated ADUser object

Return type *ADUser*

samaccountnames (*base_dn*, *distinguished_names*)
Retrieve the sAMAccountNames for the specified DNs

Parameters

- **base_dn** (*str*) – The base DN to search within
- **distinguished_name** (*list*) – A list of distinguished names for which to retrieve sAMAccountNames

Returns Key/value pairs mapping DistinguishedName to sAMAccountName

Return type dict

to_dict ()
Prepare a minimal dictionary with keys mapping to attributes for the current instance.

2.4 ADComputer

class pudl.ad_computer.**ADComputer** (*adq*)
Bases: pudl.ad_object.ADObject

A class to represent AD computer objects. Includes a number of helper methods, particularly object-factory related.

ADComputer objects have minimal depth, with attributes set to strings or lists. Available attributes are dependent on the results returned by the LDAP query.

computer (*base_dn*, *samaccountname*, *attributes=()*)
Produces a single, populated ADComputer object through the object factory. Does not populate attributes for the caller instance.

Parameters

- **base_dn** (*str*) – The base DN to search within
- **samaccountname** (*str*) – The computer's sAMAccountName
- **attributes** (*list*) – Object attributes to populate, defaults to all

Returns A populated ADComputer object

Return type *ADComputer*

computers (*base_dn*, *samaccountnames=()*, *attributes=()*)
Gathers a list of ADComputer objects

Parameters

- **base_dn** (*str*) – The base DN to search within

- **samaccountnames** (*list*) – A list of computer names for which objects will be created, defaults to all computers if unspecified
- **attributes** (*list*) – Object attributes to populate, defaults to all

Returns A list of populated ADComputer objects

Return type list

samaccountname(*base_dn*, *distinguished_name*)
Retrieve the sAMAccountName for a specific DistinguishedName

Parameters

- **base_dn** (*str*) – The base DN to search within
- **distinguished_name** (*list*) – The base DN to search within
- **attributes** (*list*) – Object attributes to populate, defaults to all

Returns A populated ADUser object

Return type *ADUser*

samaccountnames(*base_dn*, *distinguished_names*)
Retrieve the sAMAccountNames for the specified DNs

Parameters

- **base_dn** (*str*) – The base DN to search within
- **distinguished_name** (*list*) – A list of distinguished names for which to retrieve sAMAccountNames

Returns Key/value pairs mapping DistinguishedName to sAMAccountName

Return type dict

to_dict()
Prepare a minimal dictionary with keys mapping to attributes for the current instance.

2.5 Helper Functions

helper - a module containing a collection useful object manipulations

pudl.helper.**object_filter**(*objects*, *grep*)
Filter out any objects that do not have attributes with values matching *all* regular expressions present in grep (AND, essentially)

Parameters

- **ADObject** (*objects*) – A list of ADObjects
- **list** (*grep*) – A list of regular expressions that must match for filtering

Returns A list of filtered ADObjects

Return type list

pudl.helper.**serialize**(*ad_objects*, *output_format='json'*, *indent=2*, *attributes_only=False*)
Serialize the object to the specified format

Parameters

- **list** (*ad_objects*) – A list of ADObjects to serialize

- **str** (*output_format*) – The output format, json or yaml. Defaults to json
- **int** (*indent*) – The number of spaces to indent, defaults to 2
- **only** (*attributes*) – Only serialize the attributes found in the first record of the list of ADObjects

Returns A serialized, formatted representation of the list of ADObjects

Return type str

CHAPTER 3

pudl

3.1 Overview

pudl is a bundled command-line interface which wraps much of the functionality present in the pudl package modules. Not only does it demonstrate usage of the functionality present in pudl package modules, it also perhaps serves as a reasonable alternative to ldapsearch for the most common types of queries. With its simplified interface (and contrary to ldapsearch), there is no need to create custom ldap filters. Additionally, pudl has the added benefits of regular expression object filtering and object serialization, in json or yaml format. Note that all values returned are strings.

3.2 Environment Variables

To keep the pudl syntax as simple and minimal as possible, setting a few environment variables and adding them to an init file such as ~/.bashrc is advised:

- **PUDL_BASE_DN** - This is an important one to set, such as 'OU=Departments,DC=example,DC=com'.
- **PUDL_DOMAIN** - Also a key setting, the AD domain is prepended to the user name for authentication.
- **PUDL_PAGE_SIZE** - Adjusting the page size may result in faster queries, defaults to 300 results per page.
- **PUDL_TLS_NO_VERIFY** - Provides an encrypted communication channel with TLS, but does not verify the server's identity. Use with caution.

3.3 Example Usage

3.3.1 Pull Specific AD User Objects

Pulls two full user objects. Outputs json by default, with yaml being another supported output format.

```
$ pudl user bhodges jdupont
```

3.3.2 Pull a Paired-Down User Object

Pull a single user object with just three specific attributes, output results as yaml.

```
$ pudl user -a samaccountname -a title -a memberof --output-format=yaml bhodges
```

3.3.3 AD Object grep!

Pull all user objects with just two specific attributes, but filter down to only those that that match a regular expression. Note that matching is case-insensitive.

```
$ pudl user -a samaccountname -a title --grep="[iI]nfrastruct.re"
```

3.3.4 Retrieve AD Group Objects

Pull all attributes for three groups. Note that while member attribute items are fully expanded by default, the memberOf attribute is not currently flattened.

```
$ pudl group HR Finance Technology
```

3.3.5 List AD Object Attributes

Return a list of all attribute names for the first returned object in the results set

```
$ pudl user --attributes-only bhodges
```

3.4 Command-line Reference

A script for interacting with Active Directory, leveraging python-ldap

```
usage: pudl [-h] [-V] {user,group,computer} ...
```

Options:

-V, --version Print the version number and exit

Sub-commands:

user Pull user objects from AD

```
usage: pudl user [-h] [--user USER] [--password PASSWORD] [--host HOST]
                 [--port PORT] [--page-size PAGE_SIZE] [--base-dn BASE_DN]
                 [--attribute ATTRIBUTE] [--grep GREP] [--attributes-only]
                 [--output-format {json,yaml}] [--verbose] [--debug]
                 [--tls-no-verify] [--explicit-membership-only]
                 [samaccountnames [samaccountnames ...]]
```

Positional arguments:

samaccountnames sAMAccountNames for any user objects that are to be looked up. If unspecified, returns all users under the base DN provided

Options:

--user=EXAMPLE\docs, -u=EXAMPLE\docs The ldap user (bind dn) to connect as. The full DN will work, or often, just the CN may be sufficient, such as "John Smith", or more commonly, specify the domain and sAMAccountName. Defaults to EXAMPLE\username. The domain portion may be overridden with PUDL_DOMAIN

--password, -p The connecting user's password

--host=ldap, -H=ldap The AD/LDAP host, defaults to ldap

--port=389, -P=389 The ldap port, defaults to 389. 389 is is the standard port

--page-size=300, -s=300 The ldap results are paged, specify the number of results per page, defaults to 300. May be overridden with PUDL_PAGE_SIZE

--base-dn=OU=Departments,DC=example,DC=com, -b=OU=Departments,DC=example,DC=com The Base DN to use, defaults to OU=Departments,DC=example,DC=com. May be overridden with PUDL_BASE_DN

--attribute, -a Attributes to include in results objects. Note that any nested objects return all attributes. Maybe be used multiple times, and if not specified, all attributes are included in top-level objects

--grep, -g Filter results to only those matching the specified regular expression (compares against all attributes). May be used multiple times

--attributes-only=False, -A=False Only display a list of attributes that are present for the object type returned by the LDAP query

--output-format=json, -f=json Output format, defaults to json.

Possible choices: json, yaml

--verbose=False, -v=False Turn on verbose output

--debug=False, -d=False Print out debugging information, very chatty

--tls-no-verify=False, -V=False Don't verify the authenticity of the server's certificate, defaults to False and may be overridden with PUDL_TLS_NO_VERIFY

--explicit-membership-only=False, -e=False Only show membership for users that is explicit, not taking into account group nesting. Defaults to False

group Pull group objects from AD

```
usage: pudl group [-h] [--user USER] [--password PASSWORD] [--host HOST]
                  [--port PORT] [--page-size PAGE_SIZE] [--base-dn BASE_DN]
                  [--attribute ATTRIBUTE] [--grep GREP] [--attributes-only]
                  [--output-format {json,yaml}] [--verbose] [--debug]
                  [--tls-no-verify] [--explicit-membership-only]
                  [samaccountnames [samaccountnames ...]]
```

Positional arguments:

samaccountnames sAMAccountNames for any group objects that are to be looked up. If unspecified, returns all groups under the base DN provided. sAMAccountName may not be present in group objects in modern AD schemas

Options:

--user=EXAMPLE\docs, -u=EXAMPLE\docs The ldap user (bind dn) to connect as. The full DN will work, or often, just the CN may be sufficient, such as "John Smith", or more commonly, specify the domain and sAMAccountName. Defaults to EXAM-

PLE\username. The domain portion may be overridden with PUDL_DOMAIN

--password, -p The connecting user's password

--host=ldap, -H=ldap The AD/LDAP host, defaults to ldap

--port=389, -P=389 The ldap port, defaults to 389. 389 is is the standard port

--page-size=300, -s=300 The ldap results are paged, specify the number of results per page, defaults to 300. May be overridden with PUDL_PAGE_SIZE

--base-dn=OU=Departments,DC=example,DC=com, -b=OU=Departments,DC=example,DC=com The Base DN to use, defaults to OU=Departments,DC=example,DC=com. May be overridden with PUDL_BASE_DN

--attribute, -a Attributes to include in results objects. Note that any nested objects return all attributes. Maybe be used multiple times, and if not specified, all attributes are included in top-level objects

--grep, -g Filter results to only those matching the specified regular expression (compares against all attributes). May be used multiple times

--attributes-only=False, -A=False Only display a list of attributes that are present for the object type returned by the LDAP query

--output-format=json, -f=json Output format, defaults to json.

Possible choices: json, yaml

--verbose=False, -v=False Turn on verbose output

--debug=False, -d=False Print out debugging information, very chatty

--tls-no-verify=False, -V=False Don't verify the authenticity of the server's certificate, defaults to False and may be overridden with PUDL_TLS_NO_VERIFY

--explicit-membership-only=False, -e=False Only show membership for users that is explicit, not taking into account group nesting. Defaults to False

computer Pull computer objects from AD

```
usage: pudl computer [-h] [--user USER] [--password PASSWORD] [--host HOST]
                     [--port PORT] [--page-size PAGE_SIZE] [--base-dn BASE_DN]
                     [--attribute ATTRIBUTE] [--grep GREP] [--attributes-only]
                     [--output-format {json,yaml}] [--verbose] [--debug]
                     [--tls-no-verify]
                     [samaccountnames [samaccountnames ...]]
```

Positional arguments:

samaccountnames sAMAccountNames for any computer objects that are to be looked up. If unspecified, returns all computers under the base DN provided.

Options:

--user=EXAMPLE\docs, -u=EXAMPLE\docs The ldap user (bind dn) to connect as. The full DN will work, or often, just the CN may be

sufficient, such as "John Smith", or more commonly, specify the domain and sAMAccountName. Defaults to EXAMPLE\username. The domain portion may be overridden with PUDL_DOMAIN

--password, -p The connecting user's password

--host=ldap, -H=ldap The AD/LDAP host, defaults to ldap

--port=389, -P=389 The ldap port, defaults to 389. 389 is is the standard port

--page-size=300, -s=300 The ldap results are paged, specify the number of results per page, defaults to 300. May be overridden with PUDL_PAGE_SIZE

--base-dn=OU=Departments,DC=example,DC=com, -b=OU=Departments,DC=example,DC=com The Base DN to use, defaults to OU=Departments,DC=example,DC=com. May be overridden with PUDL_BASE_DN

--attribute, -a Attributes to include in results objects. Note that any nested objects return all attributes. Maybe be used multiple times, and if not specified, all attributes are included in top-level objects

--grep, -g Filter results to only those matching the specified regular expression (compares against all attributes). May be used multiple times

--attributes-only=False, -A=False Only display a list of attributes that are present for the object type returned by the LDAP query

--output-format=json, -f=json Output format, defaults to json.

Possible choices: json, yaml

--verbose=False, -v=False Turn on verbose output

--debug=False, -d=False Print out debugging information, very chatty

--tls-no-verify=False, -V=False Don't verify the authenticity of the server's certificate, defaults to False and may be overridden with PUDL_TLS_NO_VERIFY

CHAPTER 4

License

Apache License, version 2.0. Please see LICENSE

CHAPTER 5

Indices and tables

- genindex

p

A

C

G

I

O

P

S

T

U