

Spring LDAP

Reference Documentation

Version 1.2-dev

Sometime 2007

Copyright © 2005-2006 Mattias Arthursson, Ulrik Sandberg

Copies of this document may be made for your own use and for distribution to others, provided that you do not charge any fee for such copies and further provided that each copy contains this Copyright Notice, whether distributed in print or electronically.

Table of Contents

Preface	
1. Introduction	
1.1. Overview	1
1.2. Packaging overview	2
1.3. Package structure	3
1.3.1. org.springframework.transaction.compensating	3
1.3.2. org.springframework.ldap	3
1.3.3. org.springframework.ldap.core	4
1.3.4. org.springframework.ldap.core.support	4
1.3.5. org.springframework.ldap.support	4
1.3.6. org.springframework.ldap.authentication	4
1.3.7. org.springframework.ldap.control	4
1.3.8. org.springframework.ldap.filter	4
1.3.9. org.springframework.ldap.transaction	4
1.3.10. org.springframework.ldap.transaction.support	5
1.4. Support	5
2. Basic Operations	
2.1. Search and Lookup Using AttributesMapper	6
2.2. Building Dynamic Filters	7
2.3. Building Dynamic Distinguished Names	8
2.4. Binding and Unbinding	9
2.4.1. Binding Data	9
2.4.2. Unbinding Data	9
2.5. Modifying	9
2.5.1. Modifying using rebind	10
2.5.2. Modifying using modifyAttributes	10
2.6. Sample applications	10
3. DirObjectFactory and DirContextAdapter	
3.1. Introduction	11
3.2. Search and Lookup Using ContextMapper	11
3.3. Binding and Modifying Using ContextMapper	11
3.3.1. Binding	11
3.3.2. Modifying	12
3.4. A Complete PersonDao Class	13
4. Transaction Support	
4.1. Introduction	15
4.2. Configuration	15
4.3. JDBC Transaction Integration	16
4.4. LDAP Compensating Transactions Explained	16
4.4.1. Renaming Strategies	17
5. Adding Missing Overloaded API Methods	
5.1. Implementing Custom Search Methods	19
5.2. Implementing Other Custom Context Methods	20
6. Processing the DirContext	
6.1. Custom DirContext Pre/Postprocessing	21
6.2. Implementing a Request Control DirContextProcessor	21
6.3. Paged Search Results	22
7. Configuration	

7.1. ContextSource Configuration	24
7.1.1. LDAP Server URLs	24
7.1.2. Authentication	24
7.1.3. Pooling	25
7.1.4. Advanced ContextSource Configuration	25
7.2. LdapTemplate Configuration	26
7.2.1. Ignoring PartialResultExceptions	26

Preface

The Java Naming and Directory Interface (JNDI) is for LDAP programming what Java Database Connectivity (JDBC) is for SQL programming. There are several similarities between JDBC and JNDI/LDAP (Java LDAP). Despite being two completely different APIs with different pros and cons, they share a number of less flattering characteristics:

- They require extensive plumbing code, even to perform the simplest of tasks.
- All resources need to be correctly closed, no matter what happens.
- Exception handling is difficult.

The above points often lead to massive code duplication in common usages of the APIs. As we all know, code duplication is one of the worst code smells. All in all, it boils down to this: JDBC and LDAP programming in Java are both incredibly dull and repetitive.

Spring JDBC, a part of the Spring framework, provides excellent utilities for simplifying SQL programming. We need a similar framework for Java LDAP programming.

Chapter 1. Introduction

1.1. Overview

Spring-LDAP (<http://www.springframework.org/ldap>) is a library for simpler LDAP programming in Java, built on the same principles as the [JdbcTemplate](#) in Spring JDBC. It completely eliminates the need to worry about creating and closing `LdapContext` and looping through `NamingEnumeration`. It also provides a more comprehensive unchecked Exception hierarchy, built on Spring's `DataAccessException`. As a bonus, it also contains classes for dynamically building LDAP filters and DN's (Distinguished Names).

Consider, for example, a method that should search some storage for all persons and return their names in a list. Using JDBC, we would create a *connection* and execute a *query* using a *statement*. We would then loop over the *result set* and retrieve the *column* we want, adding it to a list. In contrast, using Java LDAP, we would create a *context* and perform a *search* using a *search filter*. We would then loop over the resulting *naming enumeration* and retrieve the *attribute* we want, adding it to a list.

The traditional way of implementing this person name search method in Java LDAP looks like this, where the code marked as bold actually performs tasks related to the business purpose of the method:

```
package com.example.dao;

public class TraditionalPersonDaoImpl implements PersonDao {
    public List getAllPersonNames() {
        Hashtable env = new Hashtable();
        env.put(Context.INITIAL_CONTEXT_FACTORY, "com.sun.jndi.ldap.LdapCtxFactory");
        env.put(Context.PROVIDER_URL, "ldap://localhost:389/dc=example,dc=com");

        DirContext ctx;
        try {
            ctx = new InitialDirContext(env);
        } catch (NamingException e) {
            throw new RuntimeException(e);
        }

        LinkedList list = new LinkedList();
        NamingEnumeration results = null;
        try {
            SearchControls controls = new SearchControls();
            controls.setSearchScope(SearchControls.SUBTREE_SCOPE);
            results = ctx.search("", "(objectclass=person)", controls);

            while (results.hasMore()) {
                SearchResult searchResult = (SearchResult) results.next();
                Attributes attributes = searchResult.getAttributes();
                Attribute attr = attributes.get("cn");
                String cn = (String) attr.get();
                list.add(cn);
            }
        } catch (NameNotFoundException e) {
            // The base context was not found.
            // Just clean up and exit.
        } catch (NamingException e) {
            throw new RuntimeException(e);
        } finally {
            if (results != null) {
                try {
                    results.close();
                } catch (Exception e) {
                    // Never mind this.
                }
            }
            if (ctx != null) {
                try {
                    ctx.close();
                } catch (Exception e) {
                    // Never mind this.
                }
            }
        }
    }
}
```

```

    }
    }
    return list;
}

```

By using the Spring LDAP classes `AttributesMapper` and `LdapTemplate`, we get the exact same functionality with the following code:

```

package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;

    public void setLdapTemplate(LdapTemplate ldapTemplate) {
        this.ldapTemplate = ldapTemplate;
    }

    public List getAllPersonNames() {
        return ldapTemplate.search(
            "", "(objectclass=person)",
            new AttributesMapper() {
                public Object mapFromAttributes(Attributes attrs)
                    throws NamingException {
                    return attrs.get("cn").get();
                }
            });
    }
}

```

The amount of boiler-plate code is significantly less than in the traditional example. The `LdapTemplate` version of the search method performs the search, maps the attributes to a string using the given `AttributesMapper`, collects the strings in an internal list, and finally returns the list.

Note that the `PersonDaoImpl` code simply assumes that it has an `LdapTemplate` instance, rather than looking one up somewhere. It provides a set method for this purpose. There is nothing Spring-specific about this "Inversion of Control". Anyone that can create an instance of `PersonDaoImpl` can also set the `LdapTemplate` on it. However, Spring provides a very flexible and easy way of [achieving this](#). The Spring container can be told to wire up an instance of `LdapTemplate` with its required dependencies and inject it into the `PersonDao` instance. This wiring can be defined in various ways, but the most common is through XML:

```

<beans>
    <bean id="contextSource" class="org.springframework.ldap.core.support.LdapContextSource">
        <property name="url" value="ldap://localhost:389" />
        <property name="base" value="dc=example,dc=com" />
        <property name="userDn" value="cn=Manager" />
        <property name="password" value="secret" />
    </bean>

    <bean id="ldapTemplate" class="org.springframework.ldap.core.LdapTemplate">
        <constructor-arg ref="contextSource" />
    </bean>

    <bean id="personDao" class="com.example.dao.PersonDaoImpl">
        <property name="ldapTemplate" ref="ldapTemplate" />
    </bean>
</beans>

```

1.2. Packaging overview

At a minimum, to use Spring LDAP you need:

- `spring-ldap` (the Spring LDAP library)
- `spring-core` (miscellaneous utility classes used internally by the framework)

- spring-beans (contains interfaces and classes for manipulating Java beans)
- spring-dao (basic transaction support)
- commons-logging (a simple logging facade, used internally)
- commons-lang (misc utilities, used internally)
- commons-collections (tools for working with collections, used internally)
- ldapbp (Sun LDAP Booster Pack, for required LDAP controls)

If your application is wired up using the Spring ApplicationContext, you also need:

- spring-context (adds the ability for application objects to obtain resources using a consistent API)

If you use the AcegiAuthenticationSource, you also need:

- acegi-security

1.3. Package structure

This section provides an overview of the logical package structure of the Spring LDAP codebase. The dependencies for each package are clearly noted. A package dependency noted as *optional* means that the dependency is needed to compile the package but is optionally needed at runtime (depending on your use of the package). For example, use of Spring LDAP together with Acegi Security entails use of the `org.acegisecurity` package.




Figure 1.1. Spring LDAP package structure

1.3.1. org.springframework.transaction.compensating

The `transaction.compensating` package contains the generic compensating transaction support. This is not LDAP-specific or JNDI-specific in any way.

- Dependencies: spring-dao, commons-logging

1.3.2. org.springframework.ldap

The *ldap* package contains the exceptions of the library. These exceptions form an unchecked hierarchy that mirrors the `NamingException` hierarchy.

- Dependencies: `spring-core`

1.3.3. `org.springframework.ldap.core`

The *ldap.core* package contains the central abstractions of the library. These abstractions include `AuthenticationSource`, `ContextSource`, `DirContextProcessor`, and `NameClassPairCallbackHandler`. This package also contains the central class `LdapTemplate`, plus various mappers and executors.

- Dependencies: `ldap`, `ldap.support`, `spring-beans`, `commons-collections`, `commons-lang`, `commons-logging`

1.3.4. `org.springframework.ldap.core.support`

The *ldap.core.support* package contains supporting implementations of some of the core interfaces.

- Dependencies: `ldap.core`, `ldap.support`, `spring-core`, `spring-beans`, `commons-lang`, `commons-logging`

1.3.5. `org.springframework.ldap.support`

The *ldap.support* package contains supporting utilities, like the exception translation mechanism.

- Dependencies: `ldap`, `spring-core`, `commons-logging`

1.3.6. `org.springframework.ldap.authentication`

The *ldap.authentication* package contains an implementation of the `AuthenticationSource` interface that can be used with [Acegi Security](#), as well as related helper classes.

- Dependencies: `ldap.core`, `acegi-security` (optional), `spring-beans`, `commons-lang`, `commons-logging`

1.3.7. `org.springframework.ldap.control`

The *ldap.control* package contains an abstract implementation of the `DirContextProcessor` interface that can be used as a basis for processing `RequestControls` and `ResponseControls`. There is also a concrete implementation that handles paged search results. The [LDAP Booster Pack](#) is used to get support for controls.

- Dependencies: `ldap`, `ldap.core`, `LDAP booster pack`, `spring-core`, `commons-lang`, `commons-logging`

1.3.8. `org.springframework.ldap.filter`

The *ldap.filter* package contains the `Filter` abstraction and several implementations of it.

- Dependencies: `ldap.core`, `commons-lang`

1.3.9. `org.springframework.ldap.transaction`

The *ldap.transaction* package contains the core LDAP-specific implementation of compensating transactions.

- Dependencies: `ldap`, `ldap.core`, `transaction.compensating`, `spring-core`, `spring-dao`, `commons-lang`,

commons-logging

1.3.10. org.springframework.ldap.transaction.support

The `ldap.transaction.support` package contains supporting transaction classes, like `ContextSourceTransactionManager`.

- Dependencies: `ldap.core`, `ldap.transaction`, `transaction.compensating`, `spring-dao`, `commons-logging`

For the exact list of jar dependencies, see the Spring LDAP [Ivy](#) dependency manager descriptor located within the Spring LDAP distribution at `spring-ldap/ivy.xml`

1.4. Support

Spring LDAP 1.2 is supported on Spring 1.2.8 or later, including 2.0.x.

The community support forum is located at <http://forum.springframework.org>, and the project web page is <http://www.springframework.org/ldap>.

Chapter 2. Basic Operations

2.1. Search and Lookup Using AttributesMapper

In this example we will use an `AttributesMapper` to easily build a List of all common names of all person objects.

Example 2.1. AttributesMapper that returns a single attribute

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;

    public void setLdapTemplate(LdapTemplate ldapTemplate) {
        this.ldapTemplate = ldapTemplate;
    }

    public List getAllPersonNames() {
        return ldapTemplate.search(
            "", "(objectclass=person)",
            new AttributesMapper() {
                public Object mapFromAttributes(Attributes attrs)
                    throws NamingException {
                    return attrs.get("cn").get();
                }
            });
    }
}
```

The inline implementation of `AttributesMapper` just gets the desired attribute value from the `Attributes` and returns it. Internally, `LdapTemplate` iterates over all entries found, calling the given `AttributesMapper` for each entry, and collects the results in a list. The list is then returned by the search method.

Note that the `AttributesMapper` implementation could easily be modified to return a full `Person` object:

Example 2.2. AttributesMapper that returns a Person object

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;
    ...
    private class PersonAttributesMapper implements AttributesMapper {
        public Object mapFromAttributes(Attributes attrs) throws NamingException {
            Person person = new Person();
            person.setFullName((String)attrs.get("cn").get());
            person.setLastName((String)attrs.get("sn").get());
            person.setDescription((String)attrs.get("description").get());
            return person;
        }
    }

    public List getAllPersons() {
        return ldapTemplate.search("", "(objectclass=person)", new PersonAttributesMapper());
    }
}
```

If you have the distinguished name (dn) that identifies an entry, you can retrieve the entry directly, without

searching for it. This is called a *lookup* in Java LDAP. The following example shows how a lookup results in a Person object:

Example 2.3. A lookup resulting in a Person object

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;
    ...
    public Person findPerson(String dn) {
        return (Person) ldapTemplate.lookup(dn, new PersonAttributesMapper());
    }
}
```

This will look up the specified dn and pass the found attributes to the supplied AttributesMapper, in this case resulting in a Person object.

2.2. Building Dynamic Filters

We can build dynamic filters to use in searches, using the classes from the `org.springframework.ldap.filter` package. Let's say that we want the following filter: `(&(objectclass=person)(sn=?))`, where we want the `?` to be replaced with the value of the parameter `lastName`. This is how we do it using the filter support classes:

Example 2.4. Building a search filter dynamically

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;
    ...
    public List getPersonNamesByLastName(String lastName) {
        AndFilter filter = new AndFilter();
        filter.and(new EqualsFilter("objectclass", "person"));
        filter.and(new EqualsFilter("sn", lastName));
        return ldapTemplate.search(
            "", filter.encode(),
            new AttributesMapper() {
                public Object mapFromAttributes(Attributes attrs)
                    throws NamingException {
                    return attrs.get("cn").get();
                }
            });
    }
}
```

To perform a wildcard search, it's possible to use the `WhitespaceWildcardsFilter`:

Example 2.5. Building a wildcard search filter

```
AndFilter filter = new AndFilter();
filter.and(new EqualsFilter("objectclass", "person"));
filter.and(new WhitespaceWildcardsFilter("cn", cn));
```



Note

In addition to simplifying building of complex search filters, the `Filter` classes also provide proper escaping of any unsafe characters. This prevents "ldap injection", where a user might use such characters to inject unwanted operations into your LDAP operations.

2.3. Building Dynamic Distinguished Names

The standard [Name](#) interface represents a generic name, which is basically an ordered sequence of components. The `Name` interface also provides operations on that sequence; e.g., `add` or `remove`. `LdapTemplate` provides an implementation of the `Name` interface: `DistinguishedName`. Using this class will greatly simplify building distinguished names, especially considering the sometimes complex rules regarding escapings and encodings. As with the `Filter` classes this helps preventing potentially malicious data being injected into your LDAP operations.

The following example illustrates how `DistinguishedName` can be used to dynamically construct a distinguished name:

Example 2.6. Building a distinguished name dynamically

```
package com.example.dao;

import org.springframework.ldap.core.support.DistinguishedName;
import javax.naming.Name;

public class PersonDaoImpl implements PersonDao {
    public static final String BASE_DN = "dc=example,dc=com";
    ...
    protected Name buildDn(Person p) {
        DistinguishedName dn = new DistinguishedName(BASE_DN);
        dn.add("c", p.getCountry());
        dn.add("ou", p.getCompany());
        dn.add("cn", p.getFullname());
        return dn;
    }
}
```

Assuming that a `Person` has the following attributes:

country	Sweden
company	Some Company
fullname	Some Person

The code above would then result in the following distinguished name:

```
cn=Some Person, ou=Some Company, c=Sweden, dc=example, dc=com
```

In Java 5, there is an implementation of the `Name` interface: [LdapName](#). If you are in the Java 5 world, you might as well use `LdapName`. However, you may still use `DistinguishedName` if you so wish.

2.4. Binding and Unbinding

2.4.1. Binding Data

Inserting data in Java LDAP is called binding. In order to do that, a distinguished name that uniquely identifies the new entry is required. The following example shows how data is bound using `LdapTemplate`:

Example 2.7. Binding data using Attributes

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;
    ...
    public void create(Person p) {
        Name dn = buildDn(p);
        ldapTemplate.bind(dn, null, buildAttributes(p));
    }

    private Attributes buildAttributes(Person p) {
        Attributes attrs = new BasicAttributes();
        BasicAttribute ocattr = new BasicAttribute("objectclass");
        ocattr.add("top");
        ocattr.add("person");
        attrs.put(ocattr);
        attrs.put("cn", "Some Person");
        attrs.put("sn", "Person");
        return attrs;
    }
}
```

The Attributes building is--while dull and verbose--sufficient for many purposes. It is, however, possible to simplify the binding operation further, which will be described in Chapter 3, *DirObjectFactory* and *DirContextAdapter*.

2.4.2. Unbinding Data

Removing data in Java LDAP is called unbinding. A distinguished name (dn) is required to identify the entry, just as in the binding operation. The following example shows how data is unbound using `LdapTemplate`:

Example 2.8. Unbinding data

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;
    ...
    public void delete(Person p) {
        Name dn = buildDn(p);
        ldapTemplate.unbind(dn);
    }
}
```

2.5. Modifying

In Java LDAP, data can be modified in two ways: either using *rebind* or *modifyAttributes*.

2.5.1. Modifying using *rebind*

A rebind is a very crude way to modify data. It's basically an unbind followed by a bind. It looks like this:

Example 2.9. Modifying using *rebind*

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;
    ...
    public void update(Person p) {
        Name dn = buildDn(p);
        ldapTemplate.rebind(dn, null, buildAttributes(p));
    }
}
```

2.5.2. Modifying using *modifyAttributes*

If only the modified attributes should be replaced, there is a method called *modifyAttributes* that takes an array of modifications:

Example 2.10. Modifying using *modifyAttributes*

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;
    ...
    public void updateDescription(Person p) {
        Name dn = buildDn(p);
        Attribute attr = new BasicAttribute("description", p.getDescription())
        ModificationItem item = new ModificationItem(DirContext.REPLACE_ATTRIBUTE, attr);
        ldapTemplate.modifyAttributes(dn, new ModificationItem[] {item});
    }
}
```

Building *Attributes* and *ModificationItem* arrays is a lot of work, but as you will see in Chapter 3, *DirObjectFactory* and *DirContextAdapter*, the update operations can be simplified.

2.6. Sample applications

It is recommended that you review the Spring LDAP sample applications included in the release distribution for best-practice illustrations of the features of this library. A description of each sample is provided below:

1. spring-ldap-person - the sample demonstrating most features.
2. spring-ldap-article - the sample application that was written to accompany a [java.net article](#) about Spring LDAP.

Chapter 3. DirObjectFactory and DirContextAdapter

3.1. Introduction

A little-known--and probably underestimated--feature of the Java LDAP API is the ability to register a `DirObjectFactory` to automatically create objects from found contexts. One of the reasons why it is seldom used is that you will need an implementation of `DirObjectFactory` that creates instances of a meaningful implementation of `DirContext`. The Spring LDAP library provides the missing pieces: a default implementation of `DirContext` called `DirContextAdapter`, and a corresponding implementation of `DirObjectFactory` called `DefaultDirObjectFactory`. Used together with `DefaultDirObjectFactory`, the `DirContextAdapter` can be a very powerful tool.

3.2. Search and Lookup Using ContextMapper

The `DefaultDirObjectFactory` is registered with the `ContextSource` by default, which means that whenever a context is found in the LDAP tree, its `Attributes` and `Distinguished Name (DN)` will be used to construct a `DirContextAdapter`. This enables us to use a `ContextMapper` instead of an `AttributesMapper` to transform found values:

Example 3.1. Searching using a ContextMapper

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    ...
    private static class PersonContextMapper implements ContextMapper {
        public Object mapFromContext(Object ctx) {
            DirContextAdapter context = (DirContextAdapter)ctx;
            Person p = new Person();
            p.setFullName(context.getStringAttribute("cn"));
            p.setLastName(context.getStringAttribute("sn"));
            p.setDescription(context.getStringAttribute("description"));
            return p;
        }
    }

    public Person findByPrimaryKey(
        String name, String company, String country) {
        Name dn = buildDn(name, company, country);
        return ldapTemplate.lookup(dn, new PersonContextMapper());
    }
}
```

The above code shows that it is possible to retrieve the attributes directly by name, without having to go through the `Attributes` and `BasicAttribute` classes.

3.3. Binding and Modifying Using ContextMapper

The `DirContextAdapter` can also be used to hide the `Attributes` when binding and modifying data.

3.3.1. Binding

This is an example of an improved implementation of the create DAO method. Compare it with the previous implementation in Section 2.4.1, “Binding Data”.

Example 3.2. Binding using DirContextAdapter

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    ...
    public void create(Person p) {
        Name dn = buildDn(p);
        DirContextAdapter context = new DirContextAdapter(dn);

        context.setAttributeValues("objectclass", new String[] {"top", "person"});
        context.setAttributeValue("cn", p.getFullname());
        context.setAttributeValue("sn", p.getLastname());
        context.setAttributeValue("description", p.getDescription());

        ldapTemplate.bind(dn, context, null);
    }
}
```

Note that we use the retrieved `DirContextAdapter` as the second parameter to `bind`, which should be a `Context`. The third parameter is `null`, since we're not using any `Attributes`.

3.3.2. Modifying

The code for a `rebind` would be pretty much identical to Example 3.2, except that the method called would be `rebind`. However, let's say that you don't want to remove and re-create the entry, but instead update only the attributes that have changed. The `DirContextAdapter` has the ability to keep track of its modified attributes. The following example takes advantage of this feature:

Example 3.3. Modifying using DirContextAdapter

```
package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    ...
    public void update(Person p) {
        Name dn = buildDn(p);
        DirContextAdapter context = (DirContextAdapter)ldapTemplate.lookup(dn);

        context.setAttributeValues("objectclass", new String[] {"top", "person"});
        context.setAttributeValue("cn", p.getFullname());
        context.setAttributeValue("sn", p.getLastname());
        context.setAttributeValue("description", p.getDescription());

        ldapTemplate.modifyAttributes(dn, context.getModificationItems());
    }
}
```

The observant reader will see that we have duplicated code in the `create` and `update` methods. This code maps from a domain object to a context. It can be extracted to a separate method:

Example 3.4. Binding and modifying using DirContextAdapter

```
package com.example.dao;
```



```

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;

    ...

    public void create(Person p) {
        Name dn = buildDn(p);
        DirContextAdapter context = new DirContextAdapter(dn);
        mapToContext(p, context);
        ldapTemplate.bind(dn, context, null);
    }

    public void update(Person p) {
        Name dn = buildDn(p);
        DirContextAdapter context = (DirContextAdapter)ldapTemplate.lookup(dn);
        mapToContext(person, context);
        ldapTemplate.modifyAttributes(dn, context.getModificationItems());
    }

    protected void mapToContext (Person p, DirContextAdapter context) {
        context.setAttributeValues("objectclass", new String[] {"top", "person"});
        context.setAttributeValue("cn", p.getFullName());
        context.setAttributeValue("sn", p.getLastName());
        context.setAttributeValue("description", p.getDescription());
    }
}

```

3.4. A Complete PersonDao Class

To illustrate the power of Spring LDAP, here is a complete Person DAO implementation for LDAP in just 68 lines:

Example 3.5. A complete PersonDao class

```

package com.example.dao;

import java.util.List;

import javax.naming.Name;
import javax.naming.NamingException;
import javax.naming.directory.Attributes;

import org.springframework.ldap.core.AttributesMapper;
import org.springframework.ldap.core.ContextMapper;
import org.springframework.ldap.core.LdapTemplate;
import org.springframework.ldap.core.DirContextAdapter;
import org.springframework.ldap.core.support.DistinguishedName;
import org.springframework.ldap.filter.EqualsFilter;

public class PersonDaoImpl implements PersonDao {
    private LdapTemplate ldapTemplate;

    public void setLdapTemplate(LdapTemplate ldapTemplate) {
        this.ldapTemplate = ldapTemplate;
    }

    public void create(Person person) {
        DirContextAdapter context = new DirContextAdapter();
        mapToContext(person, context);
        ldapTemplate.bind(buildDn(person), context, null);
    }

    public void update(Person person) {
        Name dn = buildDn(person);
        DirContextAdapter context = (DirContextAdapter)ldapTemplate.lookup(dn);
        mapToContext(person, context);
        ldapTemplate.modifyAttributes(dn, context.getModificationItems());
    }

    public void delete(Person person) {

```

```

        ldapTemplate.unbind(buildDn(person));
    }

    public Person findByPrimaryKey(String name, String company, String country) {
        Name dn = buildDn(name, company, country);
        return (Person) ldapTemplate.lookup(dn, getContextMapper());
    }

    public List findAll() {
        EqualsFilter filter = new EqualsFilter("objectclass", "person");
        return ldapTemplate.search(DistinguishedName.EMPTY_PATH, filter.encode(), getContextMapper());
    }

    protected ContextMapper getContextMapper() {
        return new PersonContextMapper();
    }

    protected Name buildDn(Person person) {
        return buildDn(person.getFullname(), person.getCompany(), person.getCountry());
    }

    protected Name buildDn(String fullname, String company, String country) {
        DistinguishedName dn = new DistinguishedName();
        dn.add("c", country);
        dn.add("ou", company);
        dn.add("cn", fullname);
        return dn;
    }

    protected void mapToContext(Person person, DirContextAdapter context) {
        context.setAttributeValues("objectclass", new String[] {"top", "person"});
        context.setAttributeValue("cn", person.getFullName());
        context.setAttributeValue("sn", person.getLastName());
        context.setAttributeValue("description", person.getDescription());
    }

    private static class PersonContextMapper implements ContextMapper {
        public Object mapFromContext(Object ctx) {
            DirContextAdapter context = (DirContextAdapter)ctx;
            Person person = new Person();
            person.setFullName(context.getStringAttribute("cn"));
            person.setLastName(context.getStringAttribute("sn"));
            person.setDescription(context.getStringAttribute("description"));
            return person;
        }
    }
}

```



Note

In several cases the Distinguished Name (DN) of an object is constructed using properties of the object. E.g. in the above example, the country, company and full name of the Person are used in the DN, which means that updating any of these properties will actually require moving the entry in the LDAP tree using the `rename()` operation in addition to updating the Attribute values. Since this is highly implementation specific this is something you'll need to keep track of yourself - either by disallowing the user to change these properties or performing the `rename()` operation in your `update()` method if needed.

Chapter 4. Transaction Support

4.1. Introduction

Programmers used to working with relational databases coming to the LDAP world often express surprise to the fact that there is no notion of transactions. It is not specified in the protocol, and thus no servers support it. Recognizing that this may be a major problem, Spring LDAP provides support for client-side, compensating transactions on LDAP resources.

LDAP transaction support is provided by `ContextSourceTransactionManager`, a `PlatformTransactionManager` implementation that manages Spring transaction support for LDAP operations. Along with its collaborators it keeps track of the LDAP operations performed in a transaction, making record of the state before each operation and taking steps to restore the initial state should the transaction need to be rolled back.

In addition to the actual transaction management, Spring LDAP transaction support also makes sure that the same `DirContext` instance will be used throughout the same transaction, i.e. the `DirContext` will not actually be closed until the transaction is finished, allowing for more efficient resources usage.



Note

It is important to note that while the approach used by Spring LDAP to provide transaction support is sufficient for many cases it is by no means "real" transactions in the traditional sense. The server is completely unaware of the transactions, so e.g. if the connection is broken there will be no hope to rollback the transaction. While this should be carefully considered it should also be noted that the alternative will be to operate without any transaction support whatsoever; this is pretty much as good as it gets.



Note

The client side transaction support will add some overhead in addition to the work required by the original operations. While this overhead should not be something to worry about in most cases, if your application will not perform several LDAP operations within the same transaction (e.g. a `modifyAttributes` followed by a `rebind`), or if transaction synchronization with a JDBC data source is not required (see below) there will be nothing to gain by using the LDAP transaction support.

4.2. Configuration

Configuring Spring LDAP transactions should look very familiar if you're used to configuring Spring transactions. You will create a `TransactionManager` instance and wrap your target object using a `TransactionProxyFactoryBean`. In addition to this, you will also need to wrap your `ContextSource` in a `TransactionAwareContextSourceProxy`.

```
<beans>
...
<bean id="contextSourceTarget" class="org.springframework.ldap.core.support.LdapContextSource">
  <property name="url" value="ldap://localhost:389" />
  <property name="base" value="dc=example,dc=com" />
  <property name="userDn" value="cn=Manager" />
  <property name="password" value="secret" />
</bean>

<bean id="contextSource"
```

```

        class="org.springframework.ldap.transaction.compensating.manager.TransactionAwareContextSourceProxy"
        <constructor-arg ref="contextSourceTarget" />
    </bean>

    <bean id="ldapTemplate" class="org.springframework.ldap.core.LdapTemplate">
        <constructor-arg ref="contextSource" />
    </bean>

    <bean id="transactionManager"
        class="org.springframework.ldap.transaction.compensating.manager.ContextSourceTransactionManager">
        <constructor-arg ref="contextSource" />
    </bean>

    <bean id="myDataAccessObjectTarget" class="com.example.MyDataAccessObject">
        <property name="ldapTemplate" ref="ldapTemplate" />
    </bean>

    <bean id="myDataAccessObject"
        class="org.springframework.transaction.interceptor.TransactionProxyFactoryBean">
        <property name="transactionManager" ref="transactionManager" />
        <property name="target" ref="myDataAccessObjectTarget" />
        <property name="transactionAttributes">
            <props>
                <prop key="*">PROPAGATION_REQUIRES_NEW</prop>
            </props>
        </property>
    </bean>
    ...

```

In a real world example you would probably apply the transactions on the service object level rather than the DAO level; the above serves as an example to demonstrate the general idea.



Note

You'll notice that the actual ContextSource and DAO instances get ids with a "Target" suffix. The beans you will actually refer to are the Proxies that are created around the targets; contextSource and myDataAccessObject

4.3. JDBC Transaction Integration

A common use case when working against LDAP is that some of the data is stored in the LDAP tree, but other data is stored in a relational database. In this case, transaction support becomes even more important, since the update of the different resources should be synchronized.

While actual XA transactions is not supported, support is provided to conceptually wrap JDBC and LDAP access within the same transaction using the ContextSourceAndDataSourceTransactionManager. A DataSource and a ContextSource is supplied to the ContextSourceAndDataSourceTransactionManager, which will then manage the two transactions, virtually as if they were one. When performing a commit, the LDAP part of the operation will always be performed first, allowing both transactions to be rolled back should the LDAP commit fail. The JDBC part of the transaction is managed exactly as in DataSourceTransactionManager, except that nested transactions is not supported.



Note

Once again it should be noted that the provided support is all client side. The wrapped transaction is not an XA transaction. No two-phase as such commit is performed, as the LDAP server will be unable to vote on its outcome. Once again, however, for the majority of cases the supplied support will be sufficient.

4.4. LDAP Compensating Transactions Explained

Spring LDAP manages compensating transactions by making record of the state in the LDAP tree before each modifying operation (bind, unbind, rebind, modifyAttributes, and rename).

This enables the system to perform compensating operations should the transaction need to be rolled back. In many cases the compensating operation is pretty straightforward. E.g. the compensating rollback operation for a bind operation will quite obviously be to unbind the entry. Other operations however require a different, more complicated approach because of some particular characteristics of LDAP databases. Specifically, it is not always possible to get the values of all Attributes of an entry, making the above strategy insufficient for e.g. an unbind operation.

This is why each modifying operation performed within a Spring LDAP managed transaction is internally split up in four distinct operations - a recording operation, a preparation operation, a commit operation, and a rollback operation. The specifics for each LDAP operation is described in the table below:

Table 4.1.

LDAP Operation	Recording	Preparation	Commit	Rollback
bind	Make record of the DN of the entry to bind.	Bind the entry.	No operation.	Unbind the entry using the recorded DN.
rename	Make record of the original and target DN.	Rename the entry.	No operation.	Rename the entry back to its original DN.
unbind	Make record of the original DN and calculate a temporary DN.	Rename the entry to the temporary location.	Unbind the temporary entry.	Rename the entry from the temporary location back to its original DN.
rebind	Make record of the original DN and the new Attributes, and calculate a temporary DN.	Rename the entry to a temporary location.	Bind the new Attributes at the original DN, and unbind the original entry from its temporary location.	Rename the entry from the temporary location back to its original DN.
modifyAttributes	Make record of the DN of the entry to modify and calculate compensating ModificationItems for the modifications to be done.	Perform the modifyAttributes operation.	No operation.	Perform a modifyAttributes operation using the calculated compensating ModificationItems.

A more detailed description of the internal workings of the Spring LDAP transaction support is available in the javadocs.

4.4.1. Renaming Strategies

As described in the table above, the transaction management of some operations require the original entry affected by the operation to be temporarily renamed before the actual modification can be made in the commit. The manner in which the temporary DN of the entry is calculated is managed by a `TempEntryRenamingStrategy` supplied to the `ContextSourceTransactionManager`. Two implementations are supplied with Spring LDAP, but if specific behaviour is required a custom implementation can easily be implemented by the user. The provided `TempEntryRenamingStrategy` implementations are:

- `DefaultTempEntryRenamingStrategy` (the default). Adds a suffix to the least significant part of the entry DN. E.g. for the DN `cn=john doe, ou=users`, this strategy would return the temporary DN `cn=john doe_temp, ou=users`. The suffix is configurable using the `tempSuffix` property
- `DifferentSubtreeTempEntryRenamingStrategy`. Takes the least significant part of the DN and appends a subtree DN to this. This makes all temporary entries be placed at a specific location in the LDAP tree. The temporary subtree DN is configured using the `subtreeNode` property. E.g., if `subtreeNode` is `ou=tempEntries` and the original DN of the entry is `cn=john doe, ou=users`, the temporary DN will be `cn=john doe, ou=tempEntries`. Note that the configured subtree node needs to be present in the LDAP tree.

Chapter 5. Adding Missing Overloaded API Methods

5.1. Implementing Custom Search Methods

While `LdapTemplate` contains several overloaded versions of the most common operations in `DirContext`, we have not provided an alternative for each and every method signature, mostly because there are so many of them. We have, however, provided a means to call whichever `DirContext` method you want and still get the benefits that `LdapTemplate` provides.

Let's say that you want to call the following `DirContext` method:

```
NamingEnumeration search(Name name, String filterExpr, Object[] filterArgs, SearchControls ctls)
```

There is no corresponding overloaded method in `LdapTemplate`. The way to solve this is to use a custom `SearchExecutor` implementation:

```
public interface SearchExecutor {  
    public NamingEnumeration executeSearch(DirContext ctx) throws NamingException;  
}
```

In your custom executor, you have access to a `DirContext` object, which you use to call the method you want. You then provide a handler that is responsible for mapping attributes and collecting the results. You can for example use one of the available implementations of `CollectingNameClassPairCallbackHandler`, which will collect the mapped results in an internal list. In order to actually execute the search, you call the search method in `LdapTemplate` that takes an executor and a handler as arguments. Finally, you return whatever your handler has collected.

Example 5.1. A custom search method using `SearchExecutor` and `AttributesMapper`

```
package com.example.dao;  
  
public class PersonDaoImpl implements PersonDao {  
    ...  
    public List search(final Name base, final String filter, final String[] params,  
        final SearchControls ctls) {  
        SearchExecutor executor = new SearchExecutor() {  
            public NamingEnumeration executeSearch(DirContext ctx) {  
                return ctx.search(base, filter, params, ctls);  
            }  
        };  
  
        CollectingNameClassPairCallbackHandler handler =  
            new AttributesMapperCallbackHandler(new PersonAttributesMapper());  
  
        ldapTemplate.search(executor, handler);  
        return handler.getList();  
    }  
}
```

If you prefer the `ContextMapper` to the `AttributesMapper`, this is what it would look like:

Example 5.2. A custom search method using `SearchExecutor` and `ContextMapper`

```
package com.example.dao;
```

```

public class PersonDaoImpl implements PersonDao {
    ...
    public List search(final Name base, final String filter, final String[] params,
        final SearchControls ctls) {
        SearchExecutor executor = new SearchExecutor() {
            public NamingEnumeration executeSearch(DirContext ctx) {
                return ctx.search(base, filter, params, ctls);
            }
        };

        CollectingNameClassPairCallbackHandler handler =
            new ContextMapperCallbackHandler(new PersonContextMapper());

        ldapTemplate.search(executor, handler);
        return handler.getList();
    }
}

```



Note

When using the ContextMapperCallbackHandler you must make sure that you have called setReturningObjFlag(true) on your SearchControls instance.

5.2. Implementing Other Custom Context Methods

In the same manner as for custom search methods, you can actually execute any method in DirContext by using a ContextExecutor.

```

public interface ContextExecutor {
    public Object executeWithContext(DirContext ctx) throws NamingException;
}

```

When implementing a custom ContextExecutor, you can choose between using the executeReadOnly() or the executeReadWrite() method. Let's say that we want to call this method:

```
Object lookupLink(Name name)
```

It's available in DirContext, but there is no matching method in LdapTemplate. It's a lookup method, so it should be read-only. We can implement it like this:

Example 5.3. A custom DirContext method using ContextExecutor

```

package com.example.dao;

public class PersonDaoImpl implements PersonDao {
    ...
    public Object lookupLink(final Name name) {
        ContextExecutor executor = new ContextExecutor() {
            public Object executeWithContext(DirContext ctx) {
                return ctx.lookupLink(name);
            }
        };

        return ldapTemplate.executeReadOnly(executor);
    }
}

```

In the same manner you can execute a read-write operation using the executeReadWrite() method.

Chapter 6. Processing the DirContext

6.1. Custom DirContext Pre/Postprocessing

In some situations, one would like to perform operations on the `DirContext` before and after the search operation. The interface that is used for this is called `DirContextProcessor`:

```
public interface DirContextProcessor {  
    public void preProcess(DirContext ctx) throws NamingException;  
    public void postProcess(DirContext ctx) throws NamingException;  
}
```

The `LdapTemplate` class has a search method that takes a `DirContextProcessor`:

```
public void search(SearchExecutor se, NameClassPairCallbackHandler handler,  
    DirContextProcessor processor) throws DataAccessException;
```

Before the search operation, the `preProcess` method is called on the given `DirContextProcessor` instance. After the search has been executed and the resulting `NamingEnumeration` has been processed, the `postProcess` method is called. This enables a user to perform operations on the `DirContext` to be used in the search, and to check the `DirContext` when the search has been performed. This can be very useful for example when handling request and response controls.

There are also a few convenience methods for those that don't need a custom `SearchExecutor`:

```
public void search(Name base, String filter,  
    SearchControls controls, NameClassPairCallbackHandler handler, DirContextProcessor processor)  
  
public void search(String base, String filter,  
    SearchControls controls, NameClassPairCallbackHandler handler, DirContextProcessor processor)  
  
public void search(Name base, String filter,  
    SearchControls controls, AttributesMapper mapper, DirContextProcessor processor)  
  
public void search(String base, String filter,  
    SearchControls controls, AttributesMapper mapper, DirContextProcessor processor)  
  
public void search(Name base, String filter,  
    SearchControls controls, ContextMapper mapper, DirContextProcessor processor)  
  
public void search(String base, String filter,  
    SearchControls controls, ContextMapper mapper, DirContextProcessor processor)
```

6.2. Implementing a Request Control DirContextProcessor

The LDAPv3 protocol uses Controls to send and receive additional data to affect the behavior of predefined operations. In order to simplify the implementation of a request control `DirContextProcessor`, Spring LDAP provides the base class `AbstractRequestControlDirContextProcessor`. This class handles the retrieval of the current request controls from the `LdapContext`, calls a template method for creating a request control, and adds it to the `LdapContext`. All you have to do in the subclass is to implement the template method `createRequestControl`, and of course the `postProcess` method for performing whatever you need to do after the search.

```
public abstract class AbstractRequestControlDirContextProcessor implements  
    DirContextProcessor {  
  
    public void preProcess(DirContext ctx) throws NamingException {
```

```

    }
    ...
    public abstract Control createRequestControl();
}

```

This is what it can look like when you implement your own request control `DirContextProcessor`:

Example 6.1. A request control `DirContextProcessor` implementation

```

package com.example.control;

public class MyCoolRequestControl extends AbstractRequestControlDirContextProcessor {
    private static final boolean CRITICAL_CONTROL = true;
    private MyCoolCookie cookie;
    ...
    public MyCoolCookie getCookie() {
        return cookie;
    }

    public Control createRequestControl() {
        return new SomeCoolControl(cookie.getCookie(), CRITICAL_CONTROL);
    }

    public void postProcess(DirContext ctx) throws NamingException {
        LdapContext ldapContext = (LdapContext) ctx;
        Control[] responseControls = ldapContext.getResponseControls();

        for (int i = 0; i < responseControls.length; i++) {
            if (responseControls[i] instanceof SomeCoolResponseControl) {
                SomeCoolResponseControl control = (SomeCoolResponseControl) responseControls[i];
                this.cookie = new MyCoolCookie(control.getCookie());
            }
        }
    }
}

```



Note

Make sure you use `LdapContextSource` when you use Controls. The [Control](#) interface is specific for LDAPv3 and requires that `LdapContext` is used instead of `DirContext`. If an `AbstractRequestControlDirContextProcessor` subclass is called with an argument that is not an `LdapContext`, it will throw an `IllegalArgumentException`.

6.3. Paged Search Results

Some searches may return large numbers of results. When there is no easy way to filter out a smaller amount, it would be convenient to have the server return only a certain number of results each time it is called. This is known as *paged search results*. Each "page" of the result could then be displayed at the time, with links to the next and previous page. Without this functionality, the client must either manually limit the search result into pages, or retrieve the whole result and then chop it into pages of suitable size. The former would be rather complicated, and the latter would be consuming unnecessary amounts of memory.

Some LDAP servers have support for the `PagedResultsControl`, which requests that the results of a search operation are returned by the LDAP server in pages of a specified size. The user controls the rate at which the pages are returned, simply by the rate at which the searches are called. However, the user must keep track of a *cookie* between the calls. The server uses this cookie to keep track of where it left off the previous time it was called with a paged results request.

Spring LDAP provides support for paged results by leveraging the concept for pre- and postprocessing of an `LdapContext` that was discussed in the previous sections. It does so by providing two classes: `PagedResultsRequestControl` and `PagedResultsCookie`. The `PagedResultsRequestControl` class creates a `PagedResultsControl` with the requested page size and adds it to the `LdapContext`. After the search, it gets the `PagedResultsResponseControl` and retrieves two pieces of information from it: the estimated total result size and a cookie. This cookie is a byte array containing information that the server needs the next time it is called with a `PagedResultsControl`. In order to make it easy to store this cookie between searches, Spring LDAP provides the wrapper class `PagedResultsCookie`.

This is an example of how the paged search results functionality can be used:

Example 6.2. Example of an integration test for paged search results

```
public class LdapTemplatePagedSearchITest extends TestCase {
    private LdapTemplate tested;
    ...
    // LDAP contains 5 persons matching the filter. Page size is 3.
    // Expects two batches of 3 and 2 persons respectively.
    public void testPagedResult() {
        SearchControls searchControls = new SearchControls();
        searchControls.setSearchScope(SearchControls.SUBTREE_SCOPE);
        String base = "dc=example,dc=com";
        String filter = "(&(objectclass=person)(cn=Some Person*))";
        PersonAttributesMapper mapper = new PersonAttributesMapper();
        CollectingNameClassPairCallbackHandler handler =
            tested.new AttributesMapperCallbackHandler(mapper);

        PagedResultsRequestControl requestControl;
        requestControl = new PagedResultsRequestControl(3);
        tested.search(base, filter, searchControls, handler, requestControl);
        PagedResultsCookie cookie = requestControl.getCookie();
        assertNotNull("Cookie should not be null yet", cookie.getCookie());
        assertEquals(3, callbackHandler.getList().size());

        // Prepare for second and last search
        requestControl = new PagedResultsRequestControl(3, cookie);
        tested.search(base, filter, searchControls, handler, requestControl);
        cookie = requestControl.getCookie();
        assertNull("Cookie should be null now", cookie.getCookie());
        assertEquals(5, callbackHandler.getList().size());
    }
}
```



Note

Important to note here is that we use the same `CollectingNameClassPairCallbackHandler` for both searches. This means that the results are appended to the same list. The second batch of two are added to the first three, giving a total of five after the second search.



Note

When using the `PagedResultsRequestControl` it is imperative that you keep track of the cookie returned from an operation and supply the same instance to subsequent calls. This means that your Dao method will typically need to wrap the result list together with the cookie in the value returned to the higher tiers.

Chapter 7. Configuration

7.1. ContextSource Configuration

There are several properties in `AbstractContextSource` (superclass of `DirContextSource` and `LdapContextSource`) that can be used to modify its behaviour.

7.1.1. LDAP Server URLs

The URL of the LDAP server is specified using the `url` property. The URL should be in the format `ldap://myserver.example.com:389`. For SSL access, use the `ldaps` protocol and the appropriate port, e.g. `ldaps://myserver.example.com:636`

It is possible to configure multiple alternate LDAP servers using the `urls` property. In this case, supply all server urls in a `String` array to the `urls` property.

7.1.2. Authentication

Authenticated contexts are created for both read-only and read-write operations by default. You specify `userDn` and `password` of the LDAP user to be used for authentication on the `ContextSource`.



Note

The `userDn` needs to be the full Distinguished Name (DN) of the user.

Some LDAP server setups allow anonymous read-only access. If you want to use anonymous Contexts for read-only operations, set the `anonymousReadOnly` property to `true`.

7.1.2.1. Custom Authentication Using Acegi

While the user name (i.e. user DN) and password used for creating an authenticated Context are static by default - the ones set on the `ContextSource` on startup will be used throughout the lifetime of the `ContextSource` - there are however several cases in which this is not the desired behaviour. A common scenario is that the principal and credentials of the current user should be used when executing LDAP operations for that user. The default behaviour can be modified by supplying a custom `AuthenticationSource` implementation to the `ContextSource` on startup, instead of explicitly specifying the `userDn` and `password`. The `AuthenticationSource` will be queried by the `ContextSource` for principal and credentials each time an authenticated Context is to be created.

To use the authentication information of the currently logged in user using [Acegi Security](#), use the `AcegiAuthenticationSource`:

Example 7.1. The Spring bean definition for an `AcegiAuthenticationSource`

```
<beans>
...
<bean id="contextSource" class="org.springframework.ldap.core.support.LdapContextSource">
  <property name="url" value="ldap://localhost:389" />
  <property name="base" value="dc=example,dc=com" />
  <property name="acegiAuthenticationSource" ref="authenticationSource" />
</bean>
```

```
<bean id="acegiAuthenticationSource"
      class="org.springframework ldap.authentication.AcegiAuthenticationSource" />
...
</beans>
```



Note

We don't specify any userDn or password to our ContextSource when using an AuthenticationSource - these properties are needed only when the default behaviour is used.



Note

When using the AcegiAuthenticationSource you need to use Acegi's LdapAuthenticationProvider to authenticate the users against LDAP.

7.1.2.2. Default Authentication

When using AcegiAuthenticationSource, authenticated contexts will only be possible to create once the user is logged in using Acegi. To use default authentication information when no user is logged in, use the DefaultValuesAuthenticationSourceDecorator:

Example 7.2. Configuring a DefaultValuesAuthenticationSourceDecorator

```
<beans>
...
<bean id="contextSource" class="org.springframework ldap.core.support.LdapContextSource">
  <property name="url" value="ldap://localhost:389" />
  <property name="base" value="dc=example,dc=com" />
  <property name="authenticationSource" ref="authenticationSource" />
</bean>

<bean id="authenticationSource"
      class="org.springframework ldap.authentication.DefaultValuesAuthenticationSourceDecorator">
  <property name="target" ref="acegiAuthenticationSource" />
  <property name="defaultUser" value="cn=myDefaultUser" />
  <property name="defaultPassword" value="pass" />
</bean>

<bean id="acegiAuthenticationSource"
      class="org.springframework ldap.authentication.AcegiAuthenticationSource" />
...
</beans>
```

7.1.3. Pooling

LDAP connection pooling can be turned on/off using the pooled flag. Default is true. The configuration of LDAP connection pooling is managed using System properties, so this needs to be handled manually. Details of pooling configuration can be found [here](#).

7.1.4. Advanced ContextSource Configuration

7.1.4.1. Alternate ContextFactory

It is possible to configure the ContextFactory that the ContextSource is to use when creating Contexts using

the `contextFactory` property. The default value is `com.sun.jndi.ldap.LdapCtxFactory`.

7.1.4.2. Custom `DirObjectFactory`

As described in Chapter 3, *DirObjectFactory* and *DirContextAdapter*, a `DirObjectFactory` can be used to translate the `Attributes` of found `Contexts` to a more useful `DirContext` implementation. This can be configured using the `dirObjectFactory` property. You can use this property if you have your own, custom `DirObjectFactory` implementation.

The default value is `DefaultDirObjectFactory`.

7.1.4.3. Custom `DirContext` Environment Properties

In some cases the user might want to specify additional environment setup properties in addition to the ones directly configurable from `AbstractContextSource`. Such properties should be set in a `Map` and supplied to the `baseEnvironmentProperties` property.

7.2. LdapTemplate Configuration

7.2.1. Ignoring `PartialResultExceptions`

Some Active Directory (AD) servers are unable to automatically following referrals, which often leads to a `PartialResultException` being thrown in searches. You can specify that `PartialResultException` is to be ignored by setting the `ignorePartialResultException` property to `true`.



Note

This causes all referrals to be ignored, and no notice will be given that a `PartialResultException` has been encountered. There is currently no way of manually following referrals using `LdapTemplate`.