

# **Keycloak**

# **Reference Guide**

**SSO for Web Apps and REST Services**

**1.1.0.Beta1**

---

---

---

Preface .....	vii
<b>1. License</b> .....	1
<b>2. Overview</b> .....	3
2.1. Key Concepts in Keycloak .....	4
2.2. How Does Security Work in Keycloak? .....	4
2.2.1. Permission Scopes .....	5
<b>3. Installation and Configuration of Keycloak Server</b> .....	7
3.1. Appliance Install .....	7
3.2. WAR Distribution Installation .....	7
3.3. Configuring the Server .....	8
3.3.1. Relational Database Configuration .....	9
3.3.2. MongoDB based model .....	11
3.3.3. AS7 specifics .....	13
3.3.4. SSL/HTTPS Requirement/Modes .....	14
3.3.5. SSL/HTTPS Setup .....	14
<b>4. Running Keycloak Server on OpenShift</b> .....	21
4.1. Create Keycloak instance with the web tool .....	21
4.2. Create Keycloak instance with the command-line tool .....	21
4.3. Next steps .....	22
<b>5. Master Admin Access Control</b> .....	23
5.1. Global Roles .....	23
5.2. Realm Specific Roles .....	23
<b>6. Per Realm Admin Access Control</b> .....	25
6.1. Realm Roles .....	25
<b>7. Adapters</b> .....	27
7.1. General Adapter Config .....	27
7.2. JBoss/Wildfly Adapter .....	30
7.2.1. Adapter Installation .....	30
7.2.2. Required Per WAR Configuration .....	33
7.2.3. Securing WARs via Keycloak Subsystem .....	34
7.3. Tomcat 7 Adapter .....	36
7.3.1. Adapter Installation .....	36
7.3.2. Required Per WAR Configuration .....	36
7.4. Pure Client Javascript Adapter .....	38
7.4.1. Session status iframe .....	40
7.4.2. JavaScript Adapter reference .....	41
7.5. Installed Applications .....	44
7.5.1. http://localhost .....	44
7.5.2. urn:ietf:wg:oauth:2.0:oob .....	44
7.6. Logout .....	44
7.7. Multi Tenancy .....	44
<b>8. Social</b> .....	47
8.1. Social Login Config .....	47
8.1.1. Enable social login .....	47

8.1.2. Social-only login .....	47
8.1.3. Social Callback URL .....	47
8.2. Facebook .....	47
8.3. GitHub .....	48
8.4. Google .....	48
8.5. Twitter .....	49
8.6. Social Provider SPI .....	49
<b>9. Themes .....</b>	<b>51</b>
9.1. Theme types .....	51
9.2. Configure theme .....	51
9.3. Default themes .....	51
9.4. Creating a theme .....	52
9.4.1. Stylesheets .....	52
9.4.2. Scripts .....	53
9.4.3. Images .....	53
9.4.4. Messages .....	53
9.4.5. Modifying HTML .....	54
9.5. SPIs .....	54
9.5.1. Theme SPI .....	54
9.5.2. Account SPI .....	54
9.5.3. Login SPI .....	54
<b>10. Email .....</b>	<b>57</b>
10.1. Email Server Config .....	57
10.1.1. Enable SSL or TLS .....	57
10.1.2. Authentication .....	58
<b>11. Application and Client Access Types .....</b>	<b>59</b>
<b>12. Roles .....</b>	<b>61</b>
12.1. Composite Roles .....	61
<b>13. Direct Access Grants .....</b>	<b>63</b>
<b>14. CORS .....</b>	<b>67</b>
<b>15. Cookie settings, Session Timeouts, and Token Lifespans .....</b>	<b>69</b>
15.1. Remember Me .....	69
15.2. Session Timeouts .....	69
15.3. Token Timeouts .....	69
<b>16. Admin REST API .....</b>	<b>71</b>
<b>17. Events .....</b>	<b>73</b>
17.1. Event types .....	73
17.2. Event Listener .....	73
17.3. Event Store .....	74
17.4. Configure Events Settings for Realm .....	74
<b>18. User Federation SPI and LDAP/AD Integration .....</b>	<b>77</b>
18.1. LDAP and Active Directory Plugin .....	77
18.1.1. Edit Mode .....	77
18.1.2. Other config options .....	78

---

18.2. Sync of LDAP users to Keycloak .....	78
18.3. Writing your own User Federation Provider .....	79
<b>19. Export and Import .....</b>	<b>81</b>
<b>20. Server Cache .....</b>	<b>85</b>
20.1. Disabling Caches .....	85
20.2. Clear Caches .....	86
20.3. Cache Config .....	86
<b>21. SAML SSO .....</b>	<b>87</b>
21.1. SAML Entity Descriptor .....	88
<b>22. Security Vulnerabilities .....</b>	<b>89</b>
22.1. SSL/HTTPS Requirement .....	89
22.2. CSRF Attacks .....	89
22.3. Clickjacking .....	90
22.4. Compromised Access Codes .....	90
22.5. Compromised access and refresh tokens .....	90
22.6. Open redirectors .....	90
22.7. Password guess: brute force attacks .....	91
22.8. Password database compromised .....	91
22.9. SQL Injection attacks .....	91
22.10. Limiting Scope .....	91
<b>23. Clustering .....</b>	<b>93</b>
23.1. Configure a shared database .....	93
23.2. Configure Infinispan .....	93
23.3. Enable realm and user cache invalidation .....	94
23.4. Enable distributed user sessions .....	94
23.5. Start in HA mode .....	95
23.6. Enabling cluster security .....	95
<b>24. Application Clustering .....</b>	<b>97</b>
24.1. Stateless token store .....	97
24.2. Relative URI optimization .....	98
24.3. Admin URL configuration .....	98
24.4. Registration of application nodes to Keycloak .....	99
24.5. Refresh token in each request .....	100
<b>25. Migration from older versions .....</b>	<b>101</b>
25.1. Migrate database .....	101
25.2. Migrate keycloak-server.json .....	101
25.3. Migrate providers .....	102
25.4. Migrate themes .....	102
25.5. Migrate application .....	102
25.6. Version specific migration .....	102
25.6.1. Migrating from 1.0.x.Final to 1.1.Beta1 .....	102
25.6.2. Migrating from 1.0 RC-1 to RC-2 .....	102
25.6.3. Migrating from 1.0 Beta 4 to RC-1 .....	102
25.6.4. Migrating from 1.0 Beta 1 to Beta 4 .....	103

25.6.5. Migrating from 1.0 Alpha 4 to Beta 1 .....	103
25.6.6. Migrating from 1.0 Alpha 2 to Alpha 3 .....	103
25.6.7. Migrating from 1.0 Alpha 1 to Alpha 2 .....	104

---

## Preface

In some of the example listings, what is meant to be displayed on one line does not fit inside the available page width. These lines have been broken up. A `\` at the end of a line means that a break has been introduced to fit in the page, with the following lines indented. So:

```
Let's pretend to have an extremely \  
long line that \  
does not fit  
This one is short
```

Is really:

```
Let's pretend to have an extremely long line that does not fit  
This one is short
```

---

# Chapter 1. License

Keycloak codebase is distributed under the ASL 2.0 license. It does not distribute any thirdparty libraries that are GPL. It does ship thirdparty libraries licensed under Apache ASL 2.0 and LGPL.



## Chapter 2. Overview

Keycloak is an SSO solution for web apps, mobile and RESTful web services. It is an authentication server where users can centrally login, logout, register, and manage their user accounts. The Keycloak admin UI can manage roles and role mappings for any application secured by Keycloak. The Keycloak Server can also be used to perform social logins via the user's favorite social media site i.e. Google, Facebook, Twitter etc.

Features:

- SSO and Single Log Out for browser applications
- Social Login. Enable Google, GitHub, Facebook, Twitter social login with no code required.
- LDAP and Active Directory support.
- Optional User Registration
- Password and TOTP support (via Google Authenticator). Client cert auth coming soon.
- Forgot password support. User can have an email sent to them
- Reset password/totp. Admin can force a password reset, or set up a temporary password.
- Not-before revocation policies per realm, application, or user.
- User session management. Admin can view user sessions and what applications/clients have an access token. Sessions can be invalidated per realm or per user.
- Pluggable theme and style support for user facing screens. Login, grant pages, account mgmt, and admin console all can be styled, branded, and tailored to your application and organizational needs.
- OAuth Bearer token auth for REST Services
- Integrated Browser App to REST Service token propagation
- OAuth Bearer token auth for REST Services
- OAuth 2.0 Grant requests
- OpenID Connect Support.
- SAML Support.
- CORS Support
- CORS Web Origin management and validation
- Completely centrally managed user and role mapping metadata. Minimal configuration at the application side

- Admin Console for managing users, roles, role mappings, applications, user sessions, allowed CORS web origins, and OAuth clients.
- Account Management console that allows users to manage their own account, view their open sessions, reset passwords, etc.
- Deployable as a WAR, appliance, or on Openshift. Completely clusterable.
- Multitenancy support. You can host and manage multiple realms for multiple organizations.
- Supports JBoss AS7, EAP 6.x, Wildfly, Tomcat 7 and Pure JavaScript applications. Plans to support Node.js, RAILS, GRAILS, and other non-Java deployments

### 2.1. Key Concepts in Keycloak

The core concept in Keycloak is a *Realm*. A realm secures and manages security metadata for a set of users, applications, and registered oAuth clients. Users can be created within a specific realm within the Administration console. Roles (permission types) can be defined at the realm level and you can also set up user role mappings to assign these permissions to specific users.

An *application* is a service that is secured by a realm. When a user browses an application's web site, the application can redirect the user agent to the Keycloak Server and request a login. Once a user is logged in, they can visit any other application managed by the realm and not have to re-enter credentials. This also hold true for logging out. Roles can also be defined at the application level and assigned to specific users. Depending on the application type, you may also be able to view and manage user sessions from the administration console.

An *oAuth client* is similar to an application in that it can request something like a login when a user visits the site of the oAuth client. The difference is that oAuth clients are not immediately granted all permissions of the user. In addition to requesting the login credentials of the user, the Keycloak Server will also display a grant page asking the user if it is ok to grant allowed permissions to the oAuth client.

### 2.2. How Does Security Work in Keycloak?

Keycloak uses *access tokens* to secure web invocations. Access tokens contains security metadata specifying the identity of the user as well as the role mappings for that user. The format of these tokens is a Keycloak extension to the *JSON Web Token* [<http://tools.ietf.org/html/draft-ietf-oauth-json-web-token-14>] specification. Each realm has a private and public key pair which it uses to digitally sign the access token using the *JSON Web Signature* [<http://tools.ietf.org/html/draft-ietf-jose-json-web-signature-19>] specification. Applications can verify the integrity of the digitally signed access token using the public key of the realm. The protocols used to obtain this token is defined by the *OAuth 2.0* [<http://tools.ietf.org/html/rfc6749>] specification.

The interesting thing about using these *smart* access tokens is that applications themselves are completely stateless as far as security metadata goes. All the information they need about the user is contained in the token and there's no need for them to store any security metadata locally other than the public key of the realm.

Signed access tokens can also be propagated by REST client requests within an `Authorization` header. This is great for distributed integration as applications can request a login from a client to obtain an access token, then invoke any aggregated REST invocations to other services using that access token. So, you have a distributed security model that is centrally managed, yet does not require a Keycloak Server hit per request, only for the initial login.

### 2.2.1. Permission Scopes

Each application and oauth client are configured with a set of permission scopes. These are a set of roles that an application or oauth client is allowed to ask permission for. Access tokens are always granted at the request of a specific application or oauth client. This also holds true for SSO. As you visit different sites, the application will redirect back to the Keycloak Server via the OAuth 2.0 protocol to obtain an access token specific to that application. The role mappings contained within the token are the intersection between the set of user role mappings and the permission scope of the application/oauth client. So, access tokens are tailor made for each application/oauth client and contain only the information required for by them.



# Chapter 3. Installation and Configuration of Keycloak Server

The Keycloak Server has two downloadable distributions.

- keycloak-appliance-dist-all-1.1.0.Beta1.zip
- keycloak-war-dist-all-1.1.0.Beta1.zip

## 3.1. Appliance Install

The `keycloak-appliance-dist-all-1.1.0.Beta1.zip` is quite large, but contains a complete server (backed by Wildfly) that runs out of the box. The only thing you'll have to enable and configure is SSL. Unzipping it, the directory layout looks something like this:

```
keycloak-appliance-dist-all-1.1.0.Beta1/  
  keycloak/  
    bin/  
      standalone.sh  
      standalone.bat  
      standalone/deployments/  
        auth-server.war/  
      standalone/configuration/  
        keycloak-server.json  
        themes/  
  adapters/  
    keycloak-as7-adapter-dist-1.1.0.Beta1.zip  
    keycloak-eap6-adapter-dist-1.1.0.Beta1.zip  
    keycloak-wildfly-adapter-dist-1.1.0.Beta1.zip  
  examples/  
  docs/
```

The `standalone.sh` or `standalone.bat` script is used to start the server. After executing that, log into the admin console at <http://localhost:8080/auth/admin/index.html> [<http://localhost:8080/auth/admin/index.html>]. Username: *admin* Password: *admin*. Keycloak will then prompt you to enter in a new password.

## 3.2. WAR Distribution Installation

The `keycloak-war-dist-all-1.1.0.Beta1.zip` contains just the bits you need to install keycloak on your favorite web container. We currently only support installing it on top of an existing Wildfly 8, JBoss EAP 6.x or JBoss AS 7.1.1 distribution. We may in the future provide directions

on how to install it on another web container like Tomcat or Jetty. If anybody in the community is interested in pulling this together, please contact us. Its mostly Maven pom work.

The directory structure of this distro looks like this:

```
keycloak-war-dist-all-1.1.0.Beta1/
  deployments/
    auth-server.war/
    keycloak-ds.xml
  configuration/
    keycloak-server.json
    themes/
  adapters/
    keycloak-as7-adapter-dist-1.1.0.Beta1.zip
    keycloak-eap6-adapter-dist-1.1.0.Beta1.zip
    keycloak-wildfly-adapter-dist-1.1.0.Beta1.zip
  examples/
  docs/
```

After unzipping this file, copy everything in `deployments` directory into the `standalone/deployments` of your JBoss or Wildfly distro. Also, copy everything in `configuration` directory into the `standalone/configuration` directory.

```
$ cd keycloak-war-dist-all-1.1.0.Beta1
$ cp -r deployments $JBOSS_HOME/standalone
$ cp -r configuration $JBOSS_HOME/standalone
```

After these steps you should also [install the client adapter](#) as this may contain modules the server needs (like Bouncycastle). You will also need to install the adapter to run the examples on the same server.

After booting up the JBoss or Wildfly distro, you can then make sure it is installed properly by logging into the admin console at <http://localhost:8080/auth/admin/index.html> [http://localhost:8080/auth/admin/index.html]. Username: `admin` Password: `admin`. Keycloak will then prompt you to enter in a new password.

If you have Keycloak on JBoss AS 7.1.1 [these steps](#).

### 3.3. Configuring the Server

Although the Keycloak Server is designed to run out of the box, there's some things you'll need to configure before you go into production. Specifically:

- Configuring Keycloak to use a production database.

- Setting up SSL/HTTPS
- Enforcing HTTPS connections

### 3.3.1. Relational Database Configuration

By default, Keycloak uses a relational database to store Keycloak data. This datasource is the `standalone/deployments/keycloak-ds.xml` file of your Keycloak Server installation if you used [Section 3.2, “WAR Distribution Installation”](#) or in `standalone/configuration/standalone.xml` if you used [Section 3.1, “Appliance Install”](#). File `keycloak-ds.xml` is used in WAR distribution, so that you have datasource available out of the box and you don't need to edit `standalone.xml` file. However a good thing is to always delete the file `keycloak-ds.xml` and move its configuration text into the centrally managed `standalone.xml` file. This will allow you to manage the database connection pool from the Wildfly/JBoss administration console. Here's what `standalone/configuration/standalone.xml` should look like after you've done this:

```
<subsystem xmlns="urn:jboss:domain:datasources:2.0">
  <datasources>
    <datasource jndi-name="java:jboss/datasources/ExampleDS"
                pool-name="ExampleDS" enabled="true" use-java-context="true">
      <connection-url>jdbc:h2:mem:test;DB_CLOSE_DELAY=-1;DB_CLOSE_ON_EXIT=FALSE</connection-url>
      <driver>h2</driver>
      <security>
        <user-name>sa</user-name>
        <password>sa</password>
      </security>
    </datasource>
    <datasource jndi-name="java:jboss/datasources/KeycloakDS"
                pool-name="KeycloakDS" enabled="true" use-java-context="true">
      <connection-url>jdbc:h2:${jboss.server.data.dir}/
keycloak;AUTO_SERVER=TRUE</connection-url>
      <driver>h2</driver>
      <security>
        <user-name>sa</user-name>
        <password>sa</password>
      </security>
    </datasource>
  </datasources>
  <drivers>
    <driver name="h2" module="com.h2database.h2">
      <xa-datasource-class>org.h2.jdbcx.JdbcDataSource</xa-datasource-class>
    </driver>
  </drivers>
</subsystem>
```

Besides moving the database config into the central `standalone.xml` configuration file you might want to use a better relational database for Keycloak like PostgreSQL or MySQL. You might also want to tweak the configuration settings of the datasource. Please see the [Wildfly](https://docs.jboss.org/author/display/WFLY8/DataSource+configuration) [https://docs.jboss.org/author/display/WFLY8/DataSource+configuration], [JBoss AS7](https://docs.jboss.org/author/display/AS71/DataSource+configuration) [https://docs.jboss.org/author/display/AS71/DataSource+configuration], or [JBoss EAP 6.x](https://docs.jboss.org/author/display/AS71/DataSource+configuration) [https://docs.jboss.org/author/display/AS71/DataSource+configuration] documentation on how to do this.

Keycloak also runs on a Hibernate/JPA backend which is configured in the `standalone/configuration/keycloak-server.json`. By default the setting is like this:

```
"connectionsJpa": {
  "default": {
    "dataSource": "java:jboss/datasources/KeycloakDS",
    "databaseSchema": "update"
  }
},
```

Possible configuration options are:

### dataSource

JNDI name of the dataSource

### jta

boolean property to specify if datasource is JTA capable

### driverDialect

Value of Hibernate dialect. In most cases you don't need to specify this property as dialect will be autodetected by Hibernate.

### databaseSchema

Value of database schema (Hibernate property "hibernate.hbm2ddl.auto" ).

### showSql

Specify whether Hibernate should show all SQL commands in the console (false by default)

### formatSql

Specify whether Hibernate should format SQL commands (true by default)

### unitName

Allow you to specify name of persistence unit if you want to provide your own persistence.xml file for JPA configuration. If this option is used, then all other configuration options are ignored as you are expected to configure all JPA/DB properties in your own persistence.xml file. Hence you can remove properties "dataSource" and "databaseSchema" in this case.

For more info about Hibernate properties, see [Hibernate and JPA documentation](http://hibernate.org/orm/documentation/) [http://hibernate.org/orm/documentation/].

### 3.3.1.1. Tested databases

Here is list of RDBMS databases and corresponding JDBC drivers, which were tested with Keycloak. Note that Hibernate dialect is usually set automatically according to your database, but in some cases, you must manually set the proper dialect, as the default dialect may not work correctly. You can setup dialect by adding property `driverDialect` to the `keycloak-server.json` into `connectionsJpa` section (see above).

**Table 3.1. Tested databases**

Database	JDBC driver	Hibernate Dialect
H2 1.3.161	H2 1.3.161	auto
MySQL 5.5	MySQL Connector/J 5.1.25	auto
PostgreSQL 9.2	JDBC4 Postgresql Driver, Version 9.3-1100	auto
Oracle 11g R1	Oracle JDBC Driver v11.1.0.7	auto
Microsoft SQL Server 2012	Microsoft SQL Server JDBC Driver 4.0.2206.100	org.hibernate.dialect.SQLServer2008Dialect
Sybase ASE 15.7	JDBC(TM)/7.07 ESD #5 (Build 26792)/P/EBF20686	auto

### 3.3.2. MongoDB based model

Keycloak provides [MongoDB](http://www.mongodb.com) [http://www.mongodb.com] based model implementation, which means that your identity data will be saved in MongoDB instead of traditional RDBMS. To configure Keycloak to use Mongo open `standalone/configuration/keycloak-server.json` in your favourite editor, then change:

```
"eventsStore": {
  "provider": "jpa",
  "jpa": {
    "exclude-events": [ "REFRESH_TOKEN" ]
  }
},

"realm": {
  "provider": "jpa"
},

"user": {
  "provider": "${keycloak.user.provider:jpa}"
},
```

to:

```
"eventsStore": {
  "provider": "mongo",
  "mongo": {
    "exclude-events": [ "REFRESH_TOKEN" ]
  }
},

"realm": {
  "provider": "mongo"
},

"user": {
  "provider": "mongo"
},
```

And at the end of the file add the snippet like this where you can configure details about your Mongo database:

```
"connectionsMongo": {
  "default": {
    "host": "127.0.0.1",
    "port": "27017",
    "db": "keycloak",
    "connectionsPerHost": 100
  }
}
```

All configuration options are optional. Default values for host and port are localhost and 27017. Default name of database is `keycloak`. You can also specify properties `user` and `password` if you want authenticate against your MongoDB. If user and password are not specified, Keycloak will connect unauthenticated to your MongoDB.

Finally there is set of optional configuration options, which can be used to specify connection-pooling capabilities of Mongo client. Supported int options are: `connectionsPerHost`, `threadsAllowedToBlockForConnectionMultiplier`, `maxWaitTime`, `connectTimeout` `socketTimeout`. Supported boolean options are: `socketKeepAlive`, `autoConnectRetry`. Supported long option is `maxAutoConnectRetryTime`. See [Mongo documentation](http://api.mongodb.org/java/2.11.4/com/mongodb/MongoClientOptions.html) [http://api.mongodb.org/java/2.11.4/com/mongodb/MongoClientOptions.html] for details about those options and their default values.

### 3.3.3. AS7 specifics

If you have Keycloak on JBoss AS 7.1.1 you need to do few other things in order to have Keycloak up and running. This includes:

1. Remove or comment the element from file `KEYCLOAK_HOME/standalone/deployments/auth-server.war/WEB-INF/jboss-deployment-structure.xml`.

```
<exclude-subsystems>
  <subsystem name="webservices"/>
</exclude-subsystems>
```

2. Then either:

- Disable webservices extension and subsystem by remove (or comment) both in `KEYCLOAK_HOME/standalone/configuration/standalone.xml`.
- If you want to use webservices subsystem, the alternative is to remove 2 JAR files `KEYCLOAK_HOME/standalone/deployments/auth-server.war/WEB-INF/lib/keycloak-saml-protocol-1.1.0.Beta1.jar` and `KEYCLOAK_HOME/standalone/deployments/auth-server.war/WEB-INF/lib/picketlink-federation-2.7.0.CR1.jar` but note that [SAML](#) won't work.

#### 3.3.3.1. AS7/EAP6.x Logging

Accessing the admin console will get these annoying log messages:

```
WARN [org.jboss.resteasy.core.ResourceLocator] (http-/127.0.0.1:8080-3)
    Field providers of subresource xxx will not be injected
according to spec
```

These can be ignored by editing `standalone.xml` of your jboss installation:

```
<logger category="org.jboss.resteasy.core.ResourceLocator">
  <level name="ERROR"/>
</logger>
```

### 3.3.4. SSL/HTTPS Requirement/Modes



#### Warning

Keycloak is not set up by default to handle SSL/HTTPS in either the war distribution or appliance. It is highly recommended that you either enable SSL on the Keycloak server itself or on a reverse proxy in front of the Keycloak server.

Keycloak can run out of the box without SSL so long as you stick to private IP addresses like localhost, 127.0.0.1, 10.0.x.x, 192.168.x.x, and 172..16.x.x. If you try to access Keycloak from a non-IP address you will get an error.

Keycloak has 3 SSL/HTTPS modes which you can set up in the admin console under the Settings->Login page and the `Require SSL` select box. Each adapter config should mirror this server-side setting. See adapter config section for more details.

external

Keycloak can run out of the box without SSL so long as you stick to private IP addresses like localhost, 127.0.0.1, 10.0.x.x, 192.168.x.x, and 172..16.x.x. If you try to access Keycloak from a non-IP address you will get an error.

none

Keycloak does not require SSL.

all

Keycloak requires SSL for all IP addresses.

### 3.3.5. SSL/HTTPS Setup

First enable SSL on Keycloak or on a reverse proxy in front of Keycloak. Then configure the Keycloak Server to enforce HTTPS connections.

#### 3.3.5.1. Enable SSL on Keycloak

The following things need to be done

- Generate a self signed or third-party signed certificate and import it into a Java keystore using `keytool`.
- Enable JBoss or Wildfly to use this certificate and turn on SSL/HTTPS.

##### 3.3.5.1.1. Creating the Certificate and Java Keystore

In order to allow HTTPS connections, you need to obtain a self signed or third-party signed certificate and import it into a Java keystore before you can enable HTTPS in the web container you are deploying the Keycloak Server to.

### 3.3.5.1.1.1. Self Signed Certificate

In development, you will probably not have a third party signed certificate available to test a Keycloak deployment so you'll need to generate a self-signed one. Generating one is very easy to do with the `keytool` utility that comes with the Java JDK.

```
$ keytool -genkey -alias localhost -keyalg RSA -keystore keycloak.jks -
validity 10950
Enter keystore password: secret
Re-enter new password: secret
What is your first and last name?
[Unknown]: localhost
What is the name of your organizational unit?
[Unknown]: Keycloak
What is the name of your organization?
[Unknown]: Red Hat
What is the name of your City or Locality?
[Unknown]: Westford
What is the name of your State or Province?
[Unknown]: MA
What is the two-letter country code for this unit?
[Unknown]: US
Is CN=localhost, OU=Keycloak, O=Test, L=Westford, ST=MA, C=US correct?
[no]: yes
```

You should answer the `What is your first and last name?` question with the DNS name of the machine you're installing the server on. For testing purposes, `localhost` should be used. After executing this command, the `keycloak.jks` file will be generated in the same directory as you executed the `keytool` command in.

If you want a third-party signed certificate, but don't have one, you can obtain one for free at [cacert.org](http://cacert.org) [<http://cacert.org>]. You'll have to do a little set up first before doing this though.

The first thing to do is generate a Certificate Request:

```
$ keytool -certreq -alias yourdomain -keystore keycloak.jks > keycloak.careq
```

Where `yourdomain` is a DNS name for which this certificate is generated for. Keytool generates the request:

```
-----BEGIN NEW CERTIFICATE REQUEST-----
MIIC2jCCAcICAQAwZTElMAkGA1UEBhMCVVMxCzAJBgNVBAGTAK1BMREwDwYDVQQHEwhXZXN0Zm9y
ZDEQMA4GA1UEChMHUmVkiEhhdDEQMA4GA1UECmMHUmVkiEhhdDESMBAGA1UEAxMJbG9jYXVzbn0
MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAr7kck2TaavLEOGbcpi9c0rncY4HhdzmY
Ax2nZfq1eZEaIPqI5aTxwQZzzLdk9qbeAd8Ji79HzSqnRDxNYaZu7mAYhFKHgixsole3o5Yfzbw1
29Rvy+eUve+WZxv5oo9wolVVpdSINIMEL2LaFhtX/c1dqiYVpfnvFshZQaIg2nL8juzZcBjj4as
H98gIS7khql/dkZKsw9NLvyxgJvp7PaXurX29fnf3ihG+oFrL22oFyV54BWWXCKU/GPn61EGZGw
Ft2qSIGLdctpMD1aJR2bcnlhEjZKDksjQZoQ5YMXaAGkcYkG6QkgrocDE2YXDbi7GIdf9MegVJ35
2DQMpwIDAQABoDAwLgYJKoZIhvcNAQkOMSEwHzAdBgNVHQ4EFgQUQwLZJBA+fjiDdiVzaO9vrE/i
n2swDQYJKoZIhvcNAQELBQADggEBAC5FRvMkhal3q86tHPBYWBUttmcSjs4qUm6V6f63frhveWHf
PzRrI1xH272XUIeBk0gtzWo0nNZnf0mMctUBbHhhDcG82xolikfqibZiJoQZCiGiedVjHJFTniDQ
9bMDUOXEMQ7gHZg5q6mJfNG9MbMpQaUVEEFvFGEQQxbiFK7hRWU8S23/d80e8nExgQxdJWJ6vd0X
MzzFK6j4Dj55bJVuM7GFmfdNC52pNOD5vYe47Aqh8oaJHX9XTycVtPXl45rrWAH33ftbrS8SrZ2S
vqIFQeuLL3BaHwp13t7j2lMwCk1p801aAxEASib/fAwRRHpLHBXRcq6uALUOZl4Alt8=
-----END NEW CERTIFICATE REQUEST-----
```

Send this ca request to your CA. The CA will issue you a signed certificate and send it to you. Before you import your new cert, you must obtain and import the root certificate of the CA. You can download the cert from CA (ie.: root.crt) and import as follows:

```
$ keytool -import -keystore keycloak.jks -file root.crt -alias root
```

Last step is import your new CA generated certificate to your keystore:

```
$ keytool -import -alias yourdomain -keystore keycloak.jks -file your-
certificate.cer
```

### 3.3.5.1.2. Installing the keystore to WildFly

Now that you have a Java keystore with the appropriate certificates, you need to configure your Wildfly installation to use it. First step is to move the keystore file to a directory you can reference in configuration. I like to put it in `standalone/configuration`. Then you need to edit `standalone/configuration/standalone.xml` to enable SSL/HTTPS.

To the `security-realms` element add:

```
<security-realm name="UndertowRealm">
  <server-identities>
    <ssl>
```

```

        <keystore path="keycloak.jks" relative-to="jboss.server.config.dir"
keystore-password="secret" />
    </ssl>
</server-identities>
</security-realm>

```

Find the element `<server name="default-server">` (it's a child element of `<subsystem xmlns="urn:jboss:domain:undertow:1.0">`) and add:

```

<https-listener      name="https"      socket-binding="https"      security-
realm="UndertowRealm" />

```

Check the [Wildfly Undertow](https://docs.jboss.org/author/display/WFLY8/Undertow+(web)+subsystem+configuration) [https://docs.jboss.org/author/display/WFLY8/Undertow+(web)+subsystem+configuration] documentation for more information on fine tuning the socket connections.

### 3.3.5.1.3. Installing the keystore to JBoss EAP6/AS7

Now that you have a Java keystore with the appropriate certificates, you need to configure your JBoss EAP6/AS7 installation to use it. First step is to move the keystore file to a directory you can reference in configuration. I like to put it in `standalone/configuration`. Then you need to edit `standalone/configuration/standalone.xml` to enable SSL/HTTPS.

```

<subsystem xmlns="urn:jboss:domain:web:1.1" default-virtual-server="default-
host" native="false">
  <connector name="http" protocol="HTTP/1.1" scheme="http" socket-binding="http"
  redirect-port="443" />
    <connector name="https" scheme="https" protocol="HTTP/1.1" socket-
binding="https"
      enable-lookups="false" secure="true">
      <ssl name="localhost-ssl" password="secret" protocol="TLSv1"
      key-alias="localhost" certificate-key-file="{jboss.server.config.dir}/
keycloak.jks" />
    </connector>
    ...
</subsystem>

```

Check the [JBoss](https://docs.jboss.org/author/display/AS71/SSL+setup+guide) [https://docs.jboss.org/author/display/AS71/SSL+setup+guide] documentation for more information on fine tuning the socket connections.

### 3.3.5.2. Enable SSL on a Reverse Proxy

Follow the documentation for your web server to enable SSL and configure reverse proxy for Keycloak. It is important that you make sure the web server sets the `X-Forwarded-For` and `X-Forwarded-Proto` headers on the requests made to Keycloak. Next you need to enable `proxy-address-forwarding` on the Keycloak http connector. Assuming that your reverse proxy doesn't use port 8443 for SSL you also need to configure what port http traffic is redirected to.

#### 3.3.5.2.1. WildFly

Open `standalone/configuration/standalone.xml` in your favorite editor.

First add `proxy-address-forwarding` and `redirect-socket` to the `http-listener` element:

```
<subsystem xmlns="urn:jboss:domain:undertow:1.1">
  ...
  <http-listener name="default" socket-binding="http"
    proxy-address-forwarding="true" redirect-socket="proxy-https" />
  ...
</subsystem>
```

Then add a new `socket-binding` element to the `socket-binding-group` element:

```
<socket-binding-group name="standard-sockets" default-interface="public"
  port-offset="{jboss.socket.binding.port-offset:0}">
  ...
  <socket-binding name="proxy-https" port="443" />
  ...
</socket-binding-group>
```

Check the [WildFly](https://docs.jboss.org/author/display/WFLY8/Undertow+(web)+subsystem+configuration) [https://docs.jboss.org/author/display/WFLY8/Undertow+(web)+subsystem+configuration] documentation for more information.

#### 3.3.5.2.2. AS7/EAP

Open `standalone/configuration/standalone.xml` in your favorite editor.

You need to add `redirect-port` to `http` connector element and add the `RemoteIpValve` valve:

```
<subsystem xmlns="urn:jboss:domain:web:1.5"
  default-virtual-server="default-host" native="false">
  <connector name="http" protocol="HTTP/1.1" scheme="http"
    socket-binding="http"
```

```
    redirect-port="443"/>
<virtual-server name="default-host" enable-welcome-root="true">
  <alias name="localhost"/>
  <alias name="example.com"/>
</virtual-server>
<valve name="remoteipvalve" module="org.jboss.as.web"
  class-name="org.apache.catalina.valves.RemoteIpValve">
  <param param-name="protocolHeader" param-value="x-forwarded-proto"/>
</valve>
</subsystem>
```



# Chapter 4. Running Keycloak Server on OpenShift

Keycloak provides a OpenShift cartridge to make it easy to get it running on OpenShift. If you don't already have an account or don't know how to create applications go to <https://www.openshift.com/> first. You can create the Keycloak instance either with the web tool or the command line tool, both approaches are described below.



## Warning

It's important that immediately after creating a Keycloak instance you open the Administration Console and login to reset the password. If this is not done anyone can easily gain admin rights to your Keycloak instance.

## 4.1. Create Keycloak instance with the web tool

Open <https://openshift.redhat.com/app/console/applications> and click on Add Application. Scroll down to the bottom of the page to find the Code Anything section. Insert `http://cartreflect-claytondev.rhcloud.com/github/keycloak/openshift-keycloak-cartridge` into the URL to a cartridge definition field and click on Next. Fill in the following form and click on Create Application.

Click on Continue to the application overview page. Under the list of applications you should find your Keycloak instance and the status should be Started. Click on it to open the Keycloak servers homepage.

## 4.2. Create Keycloak instance with the command-line tool

Run the following command from a terminal:

```
rhc app create <APPLICATION NAME> http://cartreflect-claytondev.rhcloud.com/github/keycloak/openshift-keycloak-cartridge
```

Replace <APPLICATION NAME> with the name you want (for example keycloak).

Once the instance is created the rhc tool outputs details about it. Open the returned URL in a browser to open the Keycloak servers homepage.

### 4.3. Next steps

The Keycloak servers homepage shows the Keycloak logo and `Welcome to Keycloak`. There is also a link to the `Administration Console`. Open that and log in using username `admin` and password `admin`. On the first login you are required to change the password.



#### Tip

On OpenShift Keycloak has been configured to only accept requests over `https`. If you try to use `http` you will be redirected to `https`.

# Chapter 5. Master Admin Access Control

You can create and manage multiple realms by logging into the `master` Keycloak admin console at `{keycloak-root}/admin/index.html`

Users in the Keycloak `master` realm can be granted permission to manage zero or more realms that are deployed on the Keycloak server. When a realm is created, Keycloak automatically creates various roles that grant fine-grain permissions to access that new realm. Access to The Admin Console and REST endpoints can be controlled by mapping these roles to users in the `master` realm. It's possible to create multiple super users as well as users that have only access to certain operations in specific realms.

## 5.1. Global Roles

There are two realm roles in the `master` realm. These are:

- `admin` - This is the super-user role and grants permissions to all operations on all realms
- `create-realm` - This grants the user permission to create new realms. A user that creates a realm is granted all permissions to the newly created realm.

To add these roles to a user select the `master` realm, then click on `Users`. Find the user you want to grant permissions to, open the user and click on `Role Mappings`. Under `Realm Roles` assign any of the above roles to the user by selecting it and clicking on the right-arrow.

## 5.2. Realm Specific Roles

Each realm in Keycloak is represented by an application in the `master` realm. The name of the application is `<realm name>-realm`. This allows assigning access to users for individual realms. The roles available are:

- `view-realm` - View the realm configuration
- `view-users` - View users (including details for specific user) in the realm
- `view-applications` - View applications in the realm
- `view-clients` - View clients in the realm
- `view-events` - View events in the realm
- `manage-realm` - Modify the realm configuration (and delete the realm)
- `manage-users` - Create, modify and delete users in the realm

- `manage-applications` - Create, modify and delete applications in the realm
- `manage-clients` - Create, modify and delete clients in the realm
- `manage-events` - Enable/disable events, clear logged events and manage event listeners

Manage roles includes permissions to view (for example a user with `manage-realm` role can also view the realm configuration).

To add these roles to a user select the `master` realm, then click on `Users`. Find the user you want to grant permissions to, open the user and click on `Role Mappings`. Under `Application Roles` select the application that represents the realm you're adding permissions to (`<realm name>-realm`), then assign any of the above roles to the user by selecting it and clicking on the right-arrow.

# Chapter 6. Per Realm Admin Access Control

Administering your realm through the `master` realm as discussed in [Chapter 5, Master Admin Access Control](#) may not always be ideal or feasible. For example, maybe you have more than one admin application that manages various admin aspects of your organization and you want to unify all these different "admin consoles" under one realm so you can do SSO between them. Keycloak allows you to grant realm admin privileges to users within that realm. These realm admins can participate in SSO for that realm and visit a keycloak admin console instance that is dedicated solely for that realm by going to the url: `/{keycloak-root}/admin/{realm}/console`

## 6.1. Realm Roles

Each realm has a built-in application called `realm-management`. This application defines roles that define permissions that can be granted to manage the realm.

- `realm-admin` - This is a composite role that grants all admin privileges for managing security for that realm.

These are more fine-grain roles you can assign to the user.

- `view-realm` - View the realm configuration
- `view-users` - View users (including details for specific user) in the realm
- `view-applications` - View applications in the realm
- `view-clients` - View clients in the realm
- `view-events` - View events in the realm
- `manage-realm` - Modify the realm configuration (and delete the realm)
- `manage-users` - Create, modify and delete users in the realm
- `manage-applications` - Create, modify and delete applications in the realm
- `manage-clients` - Create, modify and delete clients in the realm
- `manage-events` - Enable/disable events, clear logged events and manage event listeners

Manage roles includes permissions to view (for example a user with `manage-realm` role can also view the realm configuration).

To add these roles to a user select the realm you want. Then click on `Users`. Find the user you want to grant permissions to, open the user and click on `Role Mappings`. Under `Application`

Roles select `realm-management`, then assign any of the above roles to the user by selecting it and clicking on the right-arrow.

# Chapter 7. Adapters

Keycloak can secure a wide variety of application types. This section defines which application types are supported and how to configure and install them so that you can use Keycloak to secure your applications.

## 7.1. General Adapter Config

Each adapter supported by Keycloak can be configured by a simple JSON text file. This is what one might look like:

```
{
  "realm" : "demo",
  "resource" : "customer-portal",
  "realm-public-key" : "MIGfMA0GCSqGSIB3D...31LwIDAQAB",
  "auth-server-url" : "https://localhost:8443/auth",
  "ssl-required" : "external",
  "use-resource-role-mappings" : false,
  "enable-cors" : true,
  "cors-max-age" : 1000,
  "cors-allowed-methods" : [ "POST", "PUT", "DELETE", "GET" ],
  "bearer-only" : false,
  "expose-token" : true,
  "credentials" : {
    "secret" : "234234-234234-234234"
  },

  "connection-pool-size" : 20,
  "disable-trust-manager": false,
  "allow-any-hostname" : false,
  "truststore" : "path/to/truststore.jks",
  "truststore-password" : "geheim",
  "client-keystore" : "path/to/client-keystore.jks",
  "client-keystore-password" : "geheim",
  "client-key-password" : "geheim"
}
```

Some of these configuration switches may be adapter specific and some are common across all adapters. For Java adapters you can use `${...}` enclosure as System property replacement. For example `${jboss.server.config.dir}`. Also, you can obtain a template for this config file from the admin console. Go to the realm and application you want a template for. Go to the `Installation` tab and this will provide you with a template that includes the public key of the realm.

## Chapter 7. Adapters

---

Here is a description of each item:

### realm

Name of the realm representing the users of your distributed applications and services. This is *REQUIRED*.

### resource

Username of the application. Each application has a username that is used when the application connects with the Keycloak server to turn an access code into an access token (part of the OAuth 2.0 protocol). This is *REQUIRED*.

### realm-public-key

PEM format of public key. You can obtain this from the administration console. This is *REQUIRED*.

### auth-server-url

The base URL of the Keycloak Server. All other Keycloak pages and REST services are derived from this. It is usually of the form `https://host:port/auth`. This is *REQUIRED*.

### ssl-required

Ensures that all communication to and from the Keycloak server from the adapter is over HTTPS. This is *OPTIONAL*. The default value is *external* meaning that HTTPS is required by default for external requests. Valid values are 'all', 'external' and 'none'.

### use-resource-role-mappings

If set to true, the adapter will look inside the token for application level role mappings for the user. If false, it will look at the realm level for user role mappings. This is *OPTIONAL*. The default value is *false*.

### enable-cors

This enables CORS support. It will handle CORS preflight requests. It will also look into the access token to determine valid origins. This is *OPTIONAL*. The default value is *false*.

### cors-max-age

If CORS is enabled, this sets the value of the `Access-Control-Max-Age` header. This is *OPTIONAL*. If not set, this header is not returned in CORS responses.

### cors-allowed-methods

If CORS is enabled, this sets the value of the `Access-Control-Allow-Methods` header. This should be a JSON list of strings. This is *OPTIONAL*. If not set, this header is not returned in CORS responses.

### bearer-only

This tells the adapter to only do bearer token authentication. That is, it will not do OAuth 2.0 redirects, but only accept bearer tokens through the `Authorization` header. This is *OPTIONAL*. The default value is *false*.

**expose-token**

If `true`, an authenticated browser client (via a Javascript HTTP invocation) can obtain the signed access token via the URL `root/k_query_bearer_token`. This is *OPTIONAL*. The default value is *false*.

**credentials**

Specify the credentials of the application. This is an object notation where the key is the credential type and the value if the value of the credential type. Currently only `password` is supported. This is *REQUIRED*.

**connection-pool-size**

Adapters will make separate HTTP invocations to the Keycloak Server to turn an access code into an access token. This config option defines how many connections to the Keycloak Server should be pooled. This is *OPTIONAL*. The default value is `20`.

**disable-trust-manager**

If the Keycloak Server requires HTTPS and this config option is set to `true` you do not have to specify a truststore. While convenient, this setting is not recommended as you will not be verifying the host name of the Keycloak Server. This is *OPTIONAL*. The default value is *false*.

**allow-any-hostname**

If the Keycloak Server requires HTTPS and this config option is set to `true` the Keycloak Server's certificate is validated via the truststore, but host name validation is not done. This is not a recommended. This setting may be useful in test environments This is *OPTIONAL*. The default value is *false*.

**truststore**

This setting is for Java adapters. This is the file path to a Java keystore file. Used for outgoing HTTPS communications to the Keycloak server. Client making HTTPS requests need a way to verify the host of the server they are talking to. This is what the truststore does. The keystore contains one or more trusted host certificates or certificate authorities. You can create this truststore by extracting the public certificate of the Keycloak server's SSL keystore. This is *OPTIONAL* if `ssl-required` is `none` or `disable-trust-manager` is `true`. The default value is *false*.

**truststore-password**

Password for the truststore keystore. This is *REQUIRED* if `truststore` is set.

**client-keystore**

*Not supported yet, but we will support in future versions.* This setting is for Java adapters. This is the file path to a Java keystore file. This keystore contains client certificate for two-way SSL when the adapter makes HTTPS requests to the Keycloak server. This is *OPTIONAL*.

**client-keystore-password**

*Not supported yet, but we will support in future versions.* Password for the client keystore. This is *REQUIRED* if `client-keystore` is set.

client-key-password

*Not supported yet, but we will support in future versions.* Password for the client's key. This is **REQUIRED** if `client-keystore` is set.

auth-server-url-for-backend-requests

Alternative location of `auth-server-url` used just for backend requests. It must be absolute URI. Useful especially in cluster (see [Relative URI Optimization](#)) or if you would like to use `https` for browser requests but stick with `http` for backend requests etc.

always-refresh-token

If *true*, Keycloak will refresh token in every request. More info in [Refresh token in each request](#)

register-node-at-startup

If *true*, then adapter will send registration request to Keycloak. It's *false* by default as useful just in cluster (See [Registration of application nodes to Keycloak](#))

register-node-period

Period for re-registration adapter to Keycloak. Useful in cluster. See [Registration of application nodes to Keycloak](#) for details.

token-store

Possible values are *session* and *cookie*. Default is *session*, which means that adapter stores account info in HTTP Session. Alternative *cookie* means storage of info in cookie. See [Stateless token store](#) for details.

principal-attribute

OpenID Connection ID Token attribute to populate the `UserPrincipal` name with. If token attribute is null, defaults to `sub`. Possible values are `sub`, `preferred_username`, `email`, `name`, `nickname`, `given_name`, `family_name`.

## 7.2. JBoss/Wildfly Adapter

To be able to secure WAR apps deployed on JBoss AS 7.1.1, JBoss EAP 6.x, or Wildfly, you must install and configure the Keycloak Subsystem. You then have two options to secure your WARs. You can provide a keycloak config file in your WAR and change the `auth-method` to `KEYCLOAK` within `web.xml`. Alternatively, you don't have to crack open your WARs at all and can apply Keycloak via the Keycloak Subsystem configuration in `standalone.xml`. Both methods are described in this section.

### 7.2.1. Adapter Installation

There is a adapter zip file for AS7, EAP, and Wildfly in the `adapters/` directory in the Keycloak distribution.

Install on Wildfly:

```
$ cd $WILDFLY_HOME
$ unzip keycloak-wildfly-adapter-dist.zip
```

#### Install on JBoss EAP 6.x:

```
$ cd $JBOSS_HOME
$ unzip keycloak-eap6-adapter-dist.zip
```

#### Install on JBoss AS 7.1.1:

```
$ cd $JBOSS_HOME
$ unzip keycloak-as7-adapter-dist.zip
```

This zip file creates new JBoss Modules specific to the Wildfly Keycloak Adapter within your Wildfly distro.

After adding the Keycloak modules, you must then enable the Keycloak Subsystem within your app server's server configuration: `domain.xml` or `standalone.xml`.

For Wildfly:

```
<server xmlns="urn:jboss:domain:1.4">

  <extensions>
    <extension module="org.keycloak.keycloak-wildfly-subsystem"/>
    ...
  </extensions>

  <profile>
    <subsystem xmlns="urn:jboss:domain:keycloak:1.0"/>
    ...
  </profile>
```

For JBoss AS 7.1.1 and EAP 6.x:

```
<server xmlns="urn:jboss:domain:1.4">
```

```
<extensions>
  <extension module="org.keycloak.keycloak-as7-subsystem" />
  ...
</extensions>

<profile>
  <subsystem xmlns="urn:jboss:domain:keycloak:1.0" />
  ...
</profile>
```

Finally, for both AS7, EAP 6.x, and Wildfly installations you must specify a shared keycloak security domain. This security domain should be used with EJBs and other components when you need the security context created in the secured web tier to be propagated to the EJBs (other EE component) you are invoking. Otherwise this configuration is optional.

```
<server xmlns="urn:jboss:domain:1.4">
  <subsystem xmlns="urn:jboss:domain:security:1.2">
    <security-domains>
      ...
      <security-domain name="keycloak">
        <authentication>
          <login-module code="org.keycloak.adapters.jboss.KeycloakLoginModule"
            flag="required" />
        </authentication>
      </security-domain>
    </security-domains>
```

For example, if you have a JAX-RS service that is an EJB within your WEB-INF/classes directory, you'll want to annotate it with the `@SecurityDomain` annotation as follows:

```
import org.jboss.ejb3.annotation.SecurityDomain;
import org.jboss.resteasy.annotations.cache.NoCache;

import javax.annotation.security.RolesAllowed;
import javax.ejb.EJB;
import javax.ejb.Stateless;
import javax.ws.rs.GET;
import javax.ws.rs.Path;
import javax.ws.rs.Produces;
import java.util.ArrayList;
import java.util.List;
```

```

@Path("customers")
@Stateless
@SecurityDomain("keycloak")
public class CustomerService {

    @EJB
    CustomerDB db;

    @GET
    @Produces("application/json")
    @NoCache
    @RolesAllowed("db_user")
    public List<String> getCustomers() {
        return db.getCustomers();
    }
}

```

We hope to improve our integration in the future so that you don't have to specify the `@SecurityDomain` annotation when you want to propagate a keycloak security context to the EJB tier.

## 7.2.2. Required Per WAR Configuration

This section describes how to secure a WAR directly by adding config and editing files within your WAR package.

The first thing you must do is create a `keycloak.json` adapter config file within the `WEB-INF` directory of your WAR. The format of this config file is describe in the [general adapter configuration](#) section.

Next you must set the `auth-method` to `KEYCLOAK` in `web.xml`. You also have to use standard servlet security to specify role-base constraints on your URLs. Here's an example pulled from one of the examples that comes distributed with Keycloak.

```

<web-app xmlns="http://java.sun.com/xml/ns/javaee"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://java.sun.com/xml/ns/javaee http://java.sun.com/
xml/ns/javaee/web-app_3_0.xsd"
    version="3.0">

    <module-name>customer-portal</module-name>

    <security-constraint>

```

```
<web-resource-collection>
  <web-resource-name>Admins</web-resource-name>
  <url-pattern>/admin/*</url-pattern>
</web-resource-collection>
<auth-constraint>
  <role-name>admin</role-name>
</auth-constraint>
</security-constraint>
<security-constraint>
  <web-resource-collection>
    <web-resource-name>Customers</web-resource-name>
    <url-pattern>/customers/*</url-pattern>
  </web-resource-collection>
  <auth-constraint>
    <role-name>user</role-name>
  </auth-constraint>
</security-constraint>

<security-constraint>
  <web-resource-collection>
    <url-pattern>/*</url-pattern>
  </web-resource-collection>
  <user-data-constraint>
    <transport-guarantee>CONFIDENTIAL</transport-guarantee>
  </user-data-constraint>
</security-constraint>

<login-config>
  <auth-method>KEYCLOAK</auth-method>
  <realm-name>this is ignored currently/realm-name>
</login-config>

<security-role>
  <role-name>admin</role-name>
</security-role>
<security-role>
  <role-name>user</role-name>
</security-role>
</web-app>
```

### 7.2.3. Securing WARs via Keycloak Subsystem

You do not have to crack open a WAR to secure it with Keycloak. Alternatively, you can externally secure it via the Keycloak Subsystem. While you don't have to specify KEYCLOAK as an `auth-method`, you still have to define the `security-constraints` in `web.xml`. You do not, however, have to create a `WEB-INF/keycloak.json` file. This metadata is instead defined within XML in your server's `domain.xml` or `standalone.xml` subsystem configuration section.

```

<server xmlns="urn:jboss:domain:1.4">

  <profile>
    <subsystem xmlns="urn:jboss:domain:keycloak:1.0">
      <secure-deployment name="WAR MODULE NAME.war">
        <realm>demo</realm>
        <realm-public-key>MIGfMA0GCSqGSIB3DQEBAQUAA</realm-public-key>
        <auth-server-url>http://localhost:8081/auth</auth-server-url>
        <ssl-required>external</ssl-required>
        <resource>customer-portal</resource>
        <credential name="secret">password</credential>
      </secure-deployment>
    </subsystem>
  </profile>

```

The `security-deployment` name attribute identifies the WAR you want to secure. Its value is the `module-name` defined in `web.xml` with `.war` appended. The rest of the configuration corresponds pretty much one to one with the `keycloak.json` configuration options defined in [general adapter configuration](#). The exception is the `credential` element.

To make it easier for you, you can go to the Keycloak Administration Console and go to the Application/Installation tab of the application this WAR is aligned with. It provides an example XML file you can cut and paste.

There is an additional convenience format for this XML if you have multiple WARs you are deployment that are secured by the same domain. This format allows you to define common configuration items in one place under the `realm` element.

```

<subsystem xmlns="urn:jboss:domain:keycloak:1.0">
  <realm name="demo">
    <realm-public-key>MIGfMA0GCSqGSIB3DQEBA</realm-public-key>
    <auth-server-url>http://localhost:8080/auth</auth-server-url>
    <ssl-required>external</ssl-required>
  </realm>
  <secure-deployment name="customer-portal.war">
    <realm>demo</realm>
    <resource>customer-portal</resource>
    <credential name="secret">password</credential>
  </secure-deployment>
  <secure-deployment name="product-portal.war">
    <realm>demo</realm>
    <resource>product-portal</resource>
    <credential name="secret">password</credential>
  </secure-deployment>

```

```
</secure-deployment>
<secure-deployment name="database.war">
  <realm>demo</realm>
  <resource>database-service</resource>
  <bearer-only>true</bearer-only>
</secure-deployment>
</subsystem>
```

### 7.3. Tomcat 7 Adapter

To be able to secure WAR apps deployed on Tomcat 7 you must install the Keycloak Tomcat 7 adapter into your Tomcat installation. You then have to provide some extra configuration in each WAR you deploy to Tomcat. Let's go over these steps.

#### 7.3.1. Adapter Installation

There is a adapter zip file for Tomcat 7 in the `adapters/` directory in the Keycloak appliance or war distribution. You must unzip this file into Tomcat's `lib/` directory. Including adapter's jars within your `WEB-INF/lib` directory will not work! The Keycloak adapter is implemented as a Valve and valve code must reside in Tomcat's main `lib/` directory.

```
$ cd $TOMCAT_HOME/lib
$ unzip keycloak-tomcat7-adapter-dist.zip
```

#### 7.3.2. Required Per WAR Configuration

This section describes how to secure a WAR directly by adding config and editing files within your WAR package.

The first thing you must do is create a `META-INF/context.xml` file in your WAR package. This is a Tomcat specific config file and you must define a Keycloak specific Valve.

```
<Context path="/your-context-path">
  <Valve className="org.keycloak.adapters.tomcat7.KeycloakAuthenticatorValve"/>
</Context>
```

Next you must create a `keycloak.json` adapter config file within the `WEB-INF` directory of your WAR. The format of this config file is describe in the [general adapter configuration](#) section.

Finally you must specify both a `login-config` and use standard servlet security to specify role-base constraints on your URLs. Here's an example:

```
<web-app xmlns="http://java.sun.com/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/javaee http://java.sun.com/
xml/ns/javaee/web-app_3_0.xsd"
  version="3.0">

  <module-name>customer-portal</module-name>

  <security-constraint>
    <web-resource-collection>
      <web-resource-name>Customers</web-resource-name>
      <url-pattern>/*</url-pattern>
    </web-resource-collection>
    <auth-constraint>
      <role-name>user</role-name>
    </auth-constraint>
  </security-constraint>

  <security-constraint>
    <web-resource-collection>
      <url-pattern>/*</url-pattern>
    </web-resource-collection>
    <user-data-constraint>
      <transport-guarantee>CONFIDENTIAL</transport-guarantee>
    </user-data-constraint>
  </security-constraint>

  <login-config>
    <auth-method>BASIC</auth-method>
    <realm-name>this is ignored currently/realm-name>
  </login-config>

  <security-role>
    <role-name>admin</role-name>
  </security-role>
  <security-role>
    <role-name>user</role-name>
  </security-role>
</web-app>
```

### 7.4. Pure Client Javascript Adapter

The Keycloak Server comes with a Javascript library you can use to secure pure HTML/Javascript applications. It works in the same way as other application adapters except that your browser is driving the OAuth redirect protocol rather than the server.

The disadvantage of using this approach is that you end up having a non-confidential, public client. This can be mitigated by registering valid redirect URLs. You are still vulnerable if somebody hijacks the IP/DNS name of your pure HTML/Javascript application though.

To use this adapter, you must first configure an application (or client) through the `Keycloak Admin Console`. You should select `public` for the `Client Type` field. As public clients can't be verified with a client secret you are required to configure one or more valid redirect uris as well. Once you've configured the application click on the `Installation` tab and download the `keycloak.json` file. This file should be hosted in your web-server at the same root as your HTML pages. Alternatively you can either specify the URL for this file, or manually configure the adapter.

Next you have to initialize the adapter in your application. An example on how to do this is shown below.

```
<head>
  <script src="http://<keycloak server>/auth/js/keycloak.js"></script>
  <script>
    var keycloak = Keycloak();
    keycloak.init().success(function(authenticated) {
      alert(authenticated ? 'authenticated' : 'not authenticated');
    }).error(function() {
      alert('failed to initialize');
    });
  </script>
</head>
```

To specify the location of the `keycloak.json` file:

```
var keycloak = Keycloak('http://localhost:8080/myapp/keycloak.json');
```

Or finally to manually configure the adapter:

```
var keycloak = Keycloak({
  url: 'http://keycloak-server/auth',
  realm: 'myrealm',
  clientId: 'myapp'
});
```

You can also pass `login-required` or `check-sso` to the `init` function. Login required will redirect to the login form on the server, while `check-sso` will redirect to the auth server to check if the user is already logged in to the realm. For example:

```
keycloak.init({ onLoad: 'login-required' })
```

After you login, your application will be able to make REST calls using bearer token authentication. Here's an example pulled from the `customer-portal-js` example that comes with the distribution.

```
<script>
  var loadData = function () {
    document.getElementById('username').innerText = keycloak.username;

    var url = 'http://localhost:8080/database/customers';

    var req = new XMLHttpRequest();
    req.open('GET', url, true);
    req.setRequestHeader('Accept', 'application/json');
    req.setRequestHeader('Authorization', 'Bearer ' + keycloak.token);

    req.onreadystatechange = function () {
      if (req.readyState == 4) {
        if (req.status == 200) {
          var users = JSON.parse(req.responseText);
          var html = '';
          for (var i = 0; i < users.length; i++) {
            html += '<p>' + users[i] + '</p>';
          }
          document.getElementById('customers').innerHTML = html;
          console.log('finished loading data');
        }
      }
    }

    req.send();
  };
</script>
```

```
var loadFailure = function () {
    document.getElementById('customers').innerHTML = '<b>Failed to load
data. Check console log</b>';
};

var reloadData = function () {
    keycloak.updateToken().success(loadData).error(loadFailure);
}
</script>

<button onclick="loadData()">Submit</button>
```

The `loadData()` method builds an HTTP request setting the `Authorization` header to a bearer token. The `keycloak.token` points to the access token the browser obtained when it logged you in. The `loadFailure()` method is invoked on a failure. The `reloadData()` function calls `keycloak.onValidAccessToken()` passing in the `loadData()` and `loadFailure()` callbacks. The `keycloak.onValidAccessToken()` method checks to see if the access token hasn't expired. If it hasn't, and your oauth login returned a refresh token, this method will refresh the access token. Finally, if successful, it will invoke the success callback, which in this case is the `loadData()` method.

To refresh the token if it's expired call the `updateToken` method. This method returns a promise object which can be used to invoke a function on success or failure. This method can be used to wrap functions that should only be called with a valid token. For example the following method will refresh the token if it expires within 30 seconds, and then invoke the specified function. If the token is valid for more than 30 seconds it will just call the specified function.

```
keycloak.updateToken(30).success(function() {
    // send request with valid token
}).error(function() {
    alert('failed to refresh token');
});
```

### 7.4.1. Session status iframe

By default the JavaScript adapter creates a non-visible iframe that is used to detect if a single-sign out has occurred. This does not require any network traffic, instead the status is retrieved from a special status cookie. This feature can be disabled by setting `checkLoginIframe: false` in the options passed to the `init` method.

## 7.4.2. JavaScript Adapter reference

### 7.4.2.1. Constructor

```
new Keycloak();
new Keycloak('http://localhost/keycloak.json');
new Keycloak({ url: 'http://localhost/auth', realm: 'myrealm', clientId:
  'myApp' });
```

### 7.4.2.2. Properties

- `authenticated` - true if the user is authenticated
- `token` - the base64 encoded token that can be sent in the `Authorization` header in requests to services
- `tokenParsed` - the parsed token
- `subject` - the user id
- `idToken` - the id token if claims is enabled for the application, null otherwise
- `realmAccess` - the realm roles associated with the token
- `resourceAccess` - the resource roles associated with the token
- `refreshToken` - the base64 encoded token that can be used to retrieve a new token
- `refreshTokenParsed` - the parsed refresh token

### 7.4.2.3. Methods

#### `init(options)`

Called to initialize the adapter.

Options is an Object, where:

- `onLoad` - specifies an action to do on load, can be either 'login-required' or 'check-sso'
- `token` - set an initial value for the token
- `refreshToken` - set an initial value for the refresh token
- `checkLoginIframe` - set to enable/disable monitoring login state (default is true)

- `checkLoginInterval` - set the interval to check login state (default is 5 seconds)

Returns promise to set functions to be invoked on success or error.

### **login(options)**

Redirects to login form on (options is an optional object with `redirectUri` and/or `prompt` fields)

Options is an Object, where:

- `redirectUri` - specifies the uri to redirect to after login
- `prompt` - can be set to 'none' to check if the user is logged in already (if not logged in a login form is not displayed)
- `loginHint` - used to pre-fill the username/email field on the login form

### **createLoginUrl(options)**

Returns the url to login form on (options is an optional object with `redirectUri` and/or `prompt` fields)

Options is an Object, where:

- `redirectUri` - specifies the uri to redirect to after login
- `prompt` - can be set to 'none' to check if the user is logged in already (if not logged in a login form is not displayed)

### **logout(options)**

Redirects to logout

Options is an Object, where:

- `redirectUri` - specifies the uri to redirect to after logout

### **createLogoutUrl(options)**

Returns logout out

Options is an Object, where:

- `redirectUri` - specifies the uri to redirect to after logout

### **accountManagement()**

Redirects to account management

### **createAccountUrl()**

Returns the url to account management

### **hasRealmRole(role)**

Returns true if the token has the given realm role

### **hasResourceRole(role, resource)**

Returns true if the token has the given role for the resource (resource is optional, if not specified clientId is used)

### **loadUserProfile()**

Loads the users profile

Returns promise to set functions to be invoked on success or error.

### **isTokenExpired(minValidity)**

Returns true if the token has less than minValidity seconds left before it expires (minValidity is optional, if not specified 0 is used)

### **updateToken(minValidity)**

If the token expires within minValidity seconds (minValidity is optional, if not specified 0 is used) the token is refreshed. If the session status iframe is enabled, the session status is also checked.

Returns promise to set functions that can be invoked if the token is still valid, or if the token is no longer valid. For example:

```
keycloak.updateToken(5).success(function(refreshed) {
  if (refreshed) {
    alert('token was successfully refreshed');
  } else {
    alert('token is still valid');
  }
}).error(function() {
  alert('failed to refresh the token, or the session has expired');
});
```

## **7.4.2.4. Callback Events**

The adapter supports setting callback listeners for certain events. For example:

```
keycloak.onAuthSuccess = function() { alert('authenticated'); }
```

- `onReady(authenticated)` - called when the adapter is initialized
- `onAuthSuccess` - called when a user is successfully authenticated
- `onAuthError` - called if there was an error during authentication
- `onAuthRefreshSuccess` - called when the token is refreshed
- `onAuthRefreshError` - called if there was an error while trying to refresh the token
- `onAuthLogout` - called if the user is logged out (will only be called if the session status iframe is enabled, or in Cordova mode)

## 7.5. Installed Applications

Keycloak provides two special redirect uris for installed applications.

### 7.5.1. `http://localhost`

This returns the code to a web server on the client as a query parameter. Any port number is allowed. This makes it possible to start a web server for the installed application on any free port number without requiring changes in the `Admin Console`.

### 7.5.2. `urn:ietf:wg:oauth:2.0:oob`

If its not possible to start a web server in the client (or a browser is not available) it is possible to use the special `urn:ietf:wg:oauth:2.0:oob` redirect uri. When this redirect uri is used Keycloak displays a page with the code in the title and in a box on the page. The application can either detect that the browser title has changed, or the user can copy/paste the code manually to the application. With this redirect uri it is also possible for a user to use a different device to obtain a code to paste back to the application.

## 7.6. Logout

There are multiple ways you can logout from a web application. For Java EE servlet containers, you can call `HttpServletRequest.logout()`. For any other browser application, you can point the browser at the url `http://auth-server/auth/realms/{realm-name}/tokens/logout?redirect_uri=encodedRedirectUri`. This will log you out if you have an SSO session with your browser.

## 7.7. Multi Tenancy

Multi Tenancy, in our context, means that one single target application (WAR) can be secured by a single (or clustered) Keycloak server, authenticating its users against different realms. In practice,

this means that one application needs to use different `keycloak.json` files. For this case, there are two possible solutions:

- The same WAR file deployed under two different names, each with its own Keycloak configuration (probably via the Keycloak Subsystem). This scenario is suitable when the number of realms is known in advance or when there's a dynamic provision of application instances. One example would be a service provider that dynamically creates servers/deployments for their clients, like a PaaS.
- A WAR file deployed once (possibly in a cluster), that decides which realm to authenticate against based on the request parameters. This scenario is suitable when there are an undefined number of realms. One example would be a SaaS provider that have only one deployment (perhaps in a cluster) serving several companies, differentiating between clients based on the hostname (`client1.acme.com`, `client2.acme.com`) or path (`/app/client1/`, `/app/client2/`) or even via a special HTTP Header.

This chapter of the reference guide focus on this second scenario.

Keycloak provides an extension point for applications that need to evaluate the realm on a request basis. During the authentication and authorization phase of the incoming request, Keycloak queries the application via this extension point and expects the application to return a complete representation of the realm. With this, Keycloak then proceeds the authentication and authorization process, accepting or refusing the request based on the incoming credentials and on the returned realm. For this scenario, an application needs to:

- Add a context parameter to the `web.xml`, named `keycloak.config.resolver`. The value of this property should be the fully qualified name of the a class extending `org.keycloak.adapters.KeycloakConfigResolver`.
- A concrete implementation of `org.keycloak.adapters.KeycloakConfigResolver`. Keycloak will call the `resolve(org.keycloak.adapters.HttpFacade.Request)` method and expects a complete `org.keycloak.adapters.KeycloakDeployment` in response. Note that Keycloak will call this for every request, so, take the usual performance precautions.

An implementation of this feature can be found on the examples.



## Chapter 8. Social

Keycloak makes it easy to let users log in to your application using an existing account with a social network. Currently Facebook, Google and Twitter is supported with more planned for the future. There's also a Social Provider SPI that makes it relatively simple to add additional social networks.

### 8.1. Social Login Config

To enable log in with a social network you need to enable social login for your realm and configure one or more social providers.

#### 8.1.1. Enable social login

To configure social login, open the `Keycloak Admin Console`, select your realm from the drop-down box in the top left corner. In the `Login Options` section click on `Social login` to set it to `ON`. Click `save settings`, then click on `Social` in the menu at the top.

To enable a social provider select the provider you want from the drop-down and click on `Add Provider`. Then continue to the section below that provides specific instructions for the provider you are adding.

#### 8.1.2. Social-only login

It's possible to configure a realm to only allow social login. To do this open the `Keycloak Admin Console`, select your realm from the drop-down box in the top left corner. Click the `Credentials` tab, and click on the `x` next to `password` in the `Required User Credentials`. This will disable login with username and password.

#### 8.1.3. Social Callback URL

There is a single callback url used by all realms and social providers. This makes it possible to share the configuration for a social network between multiple realms. An example callback url is `http://localhost:8080/auth/rest/social/callback`. To get the callback url for your server replace `http://localhost:8080` with the base address of your server. You can also find the callback url in the `Keycloak Admin Console` under social settings.

### 8.2. Facebook

To enable login with Facebook you first have to create an app in the [Facebook Developer Console](https://developers.facebook.com/) [https://developers.facebook.com/]. Then you need to copy the client id and secret into the `Keycloak Admin Console`.

1. Log in to the [Facebook Developer Console](https://developers.facebook.com/) [https://developers.facebook.com/]. Click `Apps` in the menu and select `Create a New App`. Use any value for `Display Name` and `Category` you want, then click the `Create App` button. Wait for the project to be created (this may take

a while). If after creating the app you are not redirected to the app settings, click on `Apps` in the menu and select the app you created.

2. Once the app has been created click on `Settings` in sidebar on the left. You must specify a contact email. Save your changes. Then click on `Advanced`. Under `Security` make sure `Client OAuth Login` is enabled. In `Valid OAuth redirect URIs` insert the [social callback url](#). Scroll down and click on the `Save Changes` button.
3. Click `Status & Review` and select `YES` for `Do you want to make this app and all its live features available to the general public?`. You will not be able to set this until you have provided a contact email in the general settings of this application.
4. Click `Basic`. Copy `App ID` and `App Secret` (click `show`) from the [Facebook Developer Console](#) [<https://developers.facebook.com/>] into the settings page in the Keycloak Admin Console as the `Key` and `Secret`. Then click `Save` in the Keycloak Admin Console to enable login with Facebook.

### 8.3. GitHub

To enable login with GitHub you first have to create an application in [GitHub Settings](#) [<https://github.com/settings/applications>]. Then you need to copy the client id and secret into the Keycloak Admin Console.

1. Log in to [GitHub Settings](#) [<https://github.com/settings/applications>]. Click the `Register new application` button. Use any value for `Application name`, `Homepage URL` and `Application Description` you want. In `Authorization callback URL` enter the [social callback url](#) for your realm. Click the `Register application` button.
2. Copy `Client ID` and `Client secret` from the [GitHub Settings](#) [<https://github.com/settings/applications>] into the settings page in the Keycloak Admin Console as the `Key` and `Secret`. Then click `Save` in the Keycloak Admin Console to enable login with Google.

### 8.4. Google

To enable login with Google you first have to create a project and a client in the [Google Developer Console](#) [<https://cloud.google.com/console/project>]. Then you need to copy the client id and secret into the Keycloak Admin Console.

1. Log in to the [Google Developer Console](#) [<https://cloud.google.com/console/project>]. Click the `Create Project` button. Use any value for `Project name` and `Project ID` you want, then click the `Create` button. Wait for the project to be created (this may take a while).
2. Once the project has been created click on `APIs & auth` in sidebar on the left. To retrieve user profiles the `Google+ API` has to be enabled. Scroll down to find it in the list. If its status is `OFF`, click on `OFF` to enable it (it should move to the top of the list and the status should be `ON`).
3. Now click on the `Consent screen` link on the sidebar menu on the left. You must specify a project name and choose an email for the consent screen. Otherwise users will get a login

error. There's other things you can configure here like what the consent screen looks like. Feel free to play around with this.

4. Now click `Credentials` in the sidebar on the left. Then click `Create New Client ID`. Select `Web application` as `Application type`. Empty the `Authorized Javascript origins` textarea. In `Authorized redirect URI` enter the [social callback url](#) for your realm. Click the `Create Client ID` button.
5. Copy `Client ID` and `Client secret` from the [Google Developer Console](#) [https://cloud.google.com/console/project] into the settings page in the Keycloak Admin Console as the `Key` and `Secret`. Then click `Save` in the Keycloak Admin Console to enable login with Google.

## 8.5. Twitter

To enable login with Twitter you first have to create an application in the [Twitter Developer Console](#) [https://dev.twitter.com/apps]. Then you need to copy the consumer key and secret into the Keycloak Admin Console.

1. Log in to the [Twitter Developer Console](#) [https://dev.twitter.com/apps]. Click the `Create a new application` button. Use any value for `Name`, `Description` and `Website` you want. Insert the `social callback url` in `Callback URL`. Then click `Create your Twitter application`.
2. Now click on `Settings` and tick the box `Allow this application to be used to Sign in with Twitter`, then click on `Update this Twitter application's settings`.
3. Now click `API Keys` tab. Copy `API key` and `API secret` from the [Twitter Developer Console](#) [https://dev.twitter.com/apps] into the settings page in the Keycloak Admin Console as the `Key` and `Secret`. Then click `Save` in the Keycloak Admin Console to enable login with Twitter.



### Tip

Twitter doesn't allow `localhost` in the redirect URI. To test on a local server replace `localhost` with `127.0.0.1`.

## 8.6. Social Provider SPI

Keycloak provides an SPI to make it easy to add additional social providers. This is done by implementing `org.keycloak.social.SocialProvider` in `social/core` and adding a provider configuration file (`META-INF/services/org.keycloak.social.SocialProvider`).

A good reference for implementing a Social Provider is the Google provider which you can find in `social/google` on GitHub or in the source download.



# Chapter 9. Themes

Keycloak provides theme support for login forms and account management. This allows customizing the look and feel of end-user facing pages so they can be integrated with your brand and applications.

## 9.1. Theme types

There are several types of themes in Keycloak:

- Account - Account management
- Admin - Admin console
- Common - Shared resources for themes
- Email - Emails
- Login - Login forms
- Welcome - Welcome pages

## 9.2. Configure theme

All theme types, except welcome, is configured through `Keycloak Admin Console`. To change the theme used for a realm open the `Keycloak Admin Console`, select your realm from the drop-down box in the top left corner. Under `Settings` click on `Theme`.

To change the welcome theme you need to edit `standalone/configuration/keycloak-server.json` and add `welcomeTheme` to the theme element, for example:

```
"theme": {  
  ...  
  "welcomeTheme": "custom-theme"  
}
```

## 9.3. Default themes

Keycloak comes bundled with default themes in `standalone/configuration/themes`. It is not recommended to edit these themes directly. Instead you should create a new theme to extend a default theme. A good reference is to copy the keycloak themes as these extend the base theme to add styling.

### 9.4. Creating a theme

A theme consists of:

- *FreeMarker* [<http://freemarker.org>] templates
- Stylesheets
- Scripts
- Images
- Message bundles
- Theme properties

A theme can extend another theme. When extending a theme you can override individual files (templates, stylesheets, etc.). The recommended way to create a theme is to extend the base theme. The base theme provides templates and a default message bundle. It should be possible to achieve the customization required by styling these templates.

To create a new theme, create a folder in `.../standalone/configuration/themes/<theme type>`. The name of the folder is the name of the theme. Then create a file `theme.properties` inside the theme folder. The contents of the file should be:

```
parent=base
```

You have now created your theme. Check that it works by configuring it for a realm. It should look the same as the base theme as you've not added anything to it yet. The next sections will describe how to modify the theme.

#### 9.4.1. Stylesheets

A theme can have one or more stylesheets, to add a stylesheet create a file inside `resources/css` (for example `resources/css/styles.css`) inside your theme folder. Then registering it in `theme.properties` by adding:

```
styles=css/styles.css
```

The `styles` property supports a space separated list so you can add as many as you want. For example:

```
styles=css/styles.css css/more-styles.css
```

## 9.4.2. Scripts

A theme can have one or more scripts, to add a script create a file inside `resources/js` (for example `resources/js/script.js`) inside your theme folder. Then registering it in `theme.properties` by adding:

```
scripts=js/script.js
```

The `scripts` property supports a space separated list so you can add as many as you want. For example:

```
scripts=js/script.js js/more-script.js
```

## 9.4.3. Images

To make images available to the theme add them to `resources/img`. They can then be used through stylesheets. For example:

```
body {  
    background-image: url('../img/image.jpg');  
}
```

Or in templates, for example:

```

```

## 9.4.4. Messages

Text in the templates are loaded from message bundles. Currently internationalization isn't supported, but that will be added in a later release. A theme that extends another theme will inherit all messages from the parents message bundle, but can override individual messages. For example to replace `Username` on the login form with `Your Username` create the file `messages/messages.properties` inside your theme folder and add the following content:

```
username=Your Username
```

### 9.4.5. Modifying HTML

Keycloak uses *Freemarker Templates* [<http://freemarker.org>] in order to generate HTML. These templates are defined in `.ftl` files and can be overridden from the base theme. Check out the Freemarker website on how to form a template file.

## 9.5. SPIs

For full control of login forms and account management Keycloak provides a number of SPIs.

### 9.5.1. Theme SPI

The Theme SPI allows creating different mechanisms to load themes for the default FreeMarker based implementations of login forms and account management. To create a theme provider you will need to implement `org.keycloak.freemarker.ThemeProviderFactory` and `org.keycloak.freemarker.ThemeProvider`.

Keycloak comes with two theme providers, one that loads themes from the classpath (used by default themes) and another that loads themes from a folder (used by custom themes). Looking at these would be a good place to start to create your own theme provider. You can find them inside `forms/common-themes` on GitHub or the source download.

### 9.5.2. Account SPI

The Account SPI allows implementing the account management pages using whatever web framework or templating engine you want. To create an Account provider implement `org.keycloak.account.AccountProviderFactory` and `org.keycloak.account.AccountProvider`.

Once you have deployed your account provider to Keycloak you need to configure `keycloak-server.json` to specify which provider should be used:

```
"account": {
  "provider": "custom-provider"
}
```

### 9.5.3. Login SPI

The Login SPI allows implementing the login forms using whatever web framework or templating engine you want. To create a Login forms provider implement `org.keycloak.login.LoginFormsProviderFactory` and `org.keycloak.login.LoginFormsProvider` in `forms/login-api`.

Once you have deployed your account provider to Keycloak you need to configure `keycloak-server.json` to specify which provider should be used:

```
"login": {  
  "provider": "custom-provider"  
}
```



# Chapter 10. Email

Keycloak sends emails to users to verify their email address. Emails are also used to allow users to safely restore their username and passwords.

## 10.1. Email Server Config

To enable Keycloak to send emails you need to provide Keycloak with your SMTP server settings. If you don't have a SMTP server you can use one of many hosted solutions (such as Sendgrid or smtp2go).

To configure your SMTP server, open the `Keycloak Admin Console`, select your realm from the drop-down box in the top left corner. Then click on `Email` in the menu at the top.

You are required to fill in the `Host` and `Port` for your SMTP server (the default port for SMTP is 25). You also have to specify the sender email address (`From`). The other options are optional.

The screenshot below shows a simple example where the SMTP server doesn't use SSL or TLS and doesn't require authentication.

### acme-inc Email Server Settings

---

#### Required Settings

Host \*

Port \*

From \*

Enable SSL

 OFF

Enable StartTLS

 OFF

#### 10.1.1. Enable SSL or TLS

As emails are used for recovering usernames and passwords it's recommended to use SSL or TLS, especially if the SMTP server is on an external network. To enable SSL click on `Enable SSL`

or to enable TLS click on `Enable TLS`. You will most likely also need to change the `Port` (the default port for SSL/TLS is 465).

### 10.1.2. Authentication

If your SMTP server requires authentication click on `Enable Authentication` and insert the `Username` and `Password`.

# Chapter 11. Application and Client Access Types

When you create an Application or OAuth Client you may be wondering what the "Access Types" are.

## confidential

Confidential access type is for clients that need to perform a browser login and that you want to require a client secret when they turn an access code into an access token, (see [Access Token Request](http://tools.ietf.org/html/rfc6749#section-4.1.3) [http://tools.ietf.org/html/rfc6749#section-4.1.3] in the OAuth 2.0 spec for more details). The advantages of this is that it is a little extra security. Since Keycloak requires you to register valid redirect-uris, I'm not exactly sure what this little extra security is though. :) The disadvantages of this access type is that confidential access type is pointless for pure Javascript clients as anybody could easily figure out your client's secret!

## public

Public access type is for clients that need to perform a browser login and that you feel that the added extra security of confidential access type is not needed. FYI, Pure javascript clients are by nature public.

## bearer-only

Bearer-only access type means that the application only allows bearer token requests. If this is turned on, this application cannot participate in browser logins.

## direct access only

For OAuth clients, you would also see a "Direct Access Only" switch when creating the OAuth Client. This switch is for oauth clients that only use the [Direct Access Grant](#) protocol to obtain access tokens.



## Chapter 12. Roles

In Keycloak, roles (or permissions) can be defined globally at the realm level, or individually per application. Each role has a name which must be unique at the level it is defined in, i.e. you can have only one "admin" role at the realm level. You may have that a role named "admin" within an Application too, but "admin" must be unique for that application.

The description of a role is displayed in the OAuth Grant page when Keycloak is processing a browser OAuth Grant request. Look for more features being added here in the future like internationalization and other fine grain options.

### 12.1. Composite Roles

Any realm or application level role can be turned into a Composite Role. A Composite Role is a role that has one or more additional roles associated with it. I guess another term for it could be Role Group. When a composite role is mapped to the user, the user gains the permission of that role, plus any other role the composite is associated with. This association is dynamic. So, if you add or remove an associated role from the composite, then all users that are mapped to the composite role will automatically have those permissions added or removed. Composites can also be used to define Application or OAuth Client scopes.

Composite roles can be associated with any type of role Realm or Application. In the admin console simply flip the composite switch in the Role detail, and you will get a screen that will allow you to associate roles with the composite.



## Chapter 13. Direct Access Grants

Keycloak allows you to make direct REST invocations to obtain an access token. (See [Resource Owner Password Credentials Grant](http://tools.ietf.org/html/rfc6749#section-4.3) [http://tools.ietf.org/html/rfc6749#section-4.3] from OAuth 2.0 spec). To use it, Direct Access Grants must be allowed by your realm. This is a configuration switch in the admin console under Settings->General, specifically the "Direct Grant API" switch. You must also have registered a valid OAuth Client or Application to use as the "client\_id" for this grant request.



### Warning

It is highly recommended that you do not use Direct Access Grants to write your own login pages for your application. You will lose a lot of features that Keycloak has if you do this. Specifically all the account management, remember me, lost password, account reset features of Keycloak. Instead, if you want to tailor the look and feel of Keycloak login pages, you should create your own *theme*.

It is even highly recommended that you use the browser to log in for native mobile applications! Android and iPhone applications allow you to redirect to and from the browser. You can use this to redirect the user from your native mobile app to the web browser to perform login, then the browser will redirect back to your native application.

The REST URL to invoke on is `{keycloak-root}/realms/{realm-name}/tokens/grants/access`. Invoking on this URL is a POST request and requires you to post the username and credentials of the user you want an access token for. You must also pass along the "client\_id" of the application or oauth client you are creating an access token for. This "client\_id" is the application or oauth client name (not it's id!). Depending on whether your application/oauth client is *"public"* or *"confidential"*, you may also have to pass along it's client secret as well.

For public applications or oauth client's, the POST invocation requires form parameters that contain the username, credentials, and client\_id of your application. For example:

```
POST /auth/realms/demo/tokens/grants/access
Content-Type: application/x-www-form-urlencoded

username=bburke&password=geheim&client_id=customer-portal
```

The response would be this [standard JSON document](http://tools.ietf.org/html/rfc6749#section-4.3.3) [http://tools.ietf.org/html/rfc6749#section-4.3.3] from the OAuth 2.0 specification.

```
HTTP/1.1 200 OK
Content-Type: application/json;charset=UTF-8
Cache-Control: no-store
Pragma: no-cache
```

```
{
  "access_token": "2YotnFZFEjrlzCsicMWpAA",
  "token_type": "bearer",
  "expires_in": 3600,
  "refresh_token": "tGzv3JOkF0XG5Qx2TlKWIA",
  "id_token": "tGzv3JOkF0XG5Qx2TlKWIA",
  "session-state": "234234-234234-234234"
}
```

For confidential applications or oauth client's, you must create a `Basic Auth Authorization` header that contains the `client_id` and `client secret`. And pass in the form parameters for username and for each user credential. For example:

```
POST /auth/realms/demo/tokens/grants/access
Authorization: Basic atasdf02312312023
Content-Type: application/x-www-form-urlencoded

username=bburke&password=geheim
```

Here's a Java example using Apache HTTP Client and some Keycloak utility classes.:

```
HttpClient client = new HttpClientBuilder()
    .disableTrustManager().build();

try {
    HttpPost post = new HttpPost(
        KeycloakUriBuilder.fromUri("http://localhost:8080/auth")
            .path(ServiceUrlConstants.TOKEN_SERVICE_DIRECT_GRANT_PATH).build("demo"));
    List <NameValuePair> formparams = new ArrayList <NameValuePair>();
    formparams.add(new BasicNameValuePair("username", "bburke"));
    formparams.add(new BasicNameValuePair("password", "password"));

    if (isPublic()) { // if client is public access type
        formparams.add(new BasicNameValuePair(OAuth2Constants.CLIENT_ID,
            "customer-portal"));
    }
}
```

---

```

    } else {
        String authorization = BasicAuthHelper.createHeader("customer-portal",
"secret-secret-secret");
        post.setHeader("Authorization", authorization);
    }
    UrlEncodedFormEntity form = new UrlEncodedFormEntity(formparams, "UTF-8");
    post.setEntity(form);

    HttpResponse response = client.execute(post);
    int status = response.getStatusLine().getStatusCode();
    HttpEntity entity = response.getEntity();
    if (status != 200) {
        throw new IOException("Bad status: " + status);
    }
    if (entity == null) {
        throw new IOException("No Entity");
    }
    InputStream is = entity.getContent();
    try {
        AccessTokenResponse tokenResponse = JsonSerialization.readValue(is,
AccessTokenResponse.class);
    } finally {
        try {
            is.close();
        } catch (IOException ignored) { }
    }
} finally {
    client.getConnectionManager().shutdown();
}

```

Once you have the access token string, you can use it in REST HTTP bearer token authorized requests, i.e

```

GET /my/rest/api
Authorization: Bearer 2YotnFZFEjrlzCsicMWpAA

```

To logout you must use the refresh token contained in the AccessTokenResponse object.

```

List<NameValuePair> formparams = new ArrayList<NameValuePair>();
if (isPublic()) { // if client is public access type

```

```
formparams.add(new BasicNameValuePair(OAuth2Constants.CLIENT_ID, "customer-portal"));
} else {
    String authorization = BasicAuthHelper.createHeader("customer-portal",
"secret-secret-secret");
    post.setHeader("Authorization", authorization);
}
formparams.add(new BasicNameValuePair(OAuth2Constants.REFRESH_TOKEN,
tokenResponse.getRefreshToken()));
HttpResponse response = null;
URI logoutUri = KeycloakUriBuilder.fromUri(getBaseUrl(request) + "/auth")
    .path(ServiceUrlConstants.TOKEN_SERVICE_LOGOUT_PATH)
    .build("demo");
HttpPost post = new HttpPost(logoutUri);
UrlEncodedFormEntity form = new UrlEncodedFormEntity(formparams, "UTF-8");
post.setEntity(form);
response = client.execute(post);
int status = response.getStatusLine().getStatusCode();
HttpEntity entity = response.getEntity();
if (status != 204) {
    error(status, entity);
}
if (entity == null) {
    return;
}
InputStream is = entity.getContent();
if (is != null) is.close();
```

## Chapter 14. CORS

CORS stands for Cross-Origin Resource Sharing. If executing browser Javascript tries to make an AJAX HTTP request to a server's whose domain is different than the one the Javascript code came from, then the request uses the [CORS protocol](http://www.w3.org/TR/cors/) [http://www.w3.org/TR/cors/]. The server must handle CORS requests in a special way, otherwise the browser will not display or allow the request to be processed. This protocol exists to protect against XSS and other Javascript-based attacks. Keycloak has support for validated CORS requests.

Keycloak's CORS support is configured per application and oauth client. You specify the allowed origins in the application's or oauth client's configuration page in the admin console. You can add as many you want. The value must be what the browser would send as a value in the `Origin` header. For example `http://example.com` is what you must specify to allow CORS requests from `example.com`. When an access token is created for the application or OAuth client, these allowed origins are embedded within the token. On authenticated CORS requests, your application's Keycloak adapter will handle the CORS protocol and validate the `Origin` header against the allowed origins embedded in the token. If there is no match, then the request is denied.

To enable CORS processing in your application's server, you must set the `enable-cors` setting to `true` in your [adapter's configuration file](#). When this setting is enabled, the Keycloak adapter will handle all CORS preflight requests. It will validate authenticated requests (protected resource requests), but will let unauthenticated requests (unprotected resource requests) pass through.



# Chapter 15. Cookie settings, Session Timeouts, and Token Lifespans

Keycloak has a bunch of fine-grain settings to manage browser cookies, user login sessions, and token lifespans. Sessions can be viewed and managed within the admin console for all users, and individually in the user's account management pages. This chapter goes over configuration options for cookies, sessions, and tokens.

## 15.1. Remember Me

If you go to the admin console page of Settings->General, you should see a `Remember Me` on/off switch. Your realm sets a SSO cookie so that you only have to enter in your login credentials once. This `Remember Me` admin config option, when turned on, will show a "Remember Me" checkbox on the user's login page. If the user clicks this, the realm's SSO cookie will be persistent. This means that if the user closes their browser they will still be logged in the next time they start up their browser.

## 15.2. Session Timeouts

If you go to the Sessions and Tokens->Timeout Settings page of the Keycloak administration console there is a bunch of fine tuning you can do as far as login session timeouts go.

The `SSO Session Idle Timeout` is the idle time of a user session. If there is no activity in the user's session for this amount of time, the user session will be destroyed, and the user will become logged out. The idle time is refreshed with every action against the keycloak server for that session, i.e.: a user login, SSO, a refresh token grant, etc.

The `SSO Session Max Lifespan` setting is the maximum time a user session is allowed to be alive. This max lifespan countdown starts from when the user first logs in and is never refreshed. This works great with `Remember Me` in that it allow you to force a relogin after a set timeframe.

## 15.3. Token Timeouts

The `Access Token Lifespan` is how long an access token is valid for. An access token contains everything an application needs to authorize a client. It contains roles allowed as well as other user information. When an access token expires, your application will attempt to refresh it using a refresh token that it obtained in the initial login. The value of this configuration option should be however long you feel comfortable with the application not knowing if the user's permissions have changed. This value is usually in minutes.

The `Client login timeout` is how long an access code is valid for. An access code is obtained on the 1st leg of the OAuth 2.0 redirection protocol. This should be a short time limit. Usually seconds.

The `Login user action lifespan` is how long a user is allowed to attempt a login. When a user tries to login, they may have to change their password, set up TOTP, or perform some other action before they are redirected back to your application as an authenticated user. This value is relatively short and is usually measured in minutes.

## Chapter 16. Admin REST API

The Keycloak Admin Console is implemented entirely with a fully functional REST admin API. You can invoke this REST API from your Java applications by obtaining an access token. You must have the appropriate permissions set up as describe in [Chapter 5, Master Admin Access Control](#) and [Chapter 6, Per Realm Admin Access Control](#)

The documentation for this REST API is auto-generated and is contained in the distribution of keycloak under the docs/rest-api/overview-index.html directory, or directly from the docs page at the keycloak website.

There are a number of examples that come with the keycloak distribution that show you how to invoke on this REST API. `examples/preconfigured-demo/admin-access-app` shows you how to access this api from java. `examples/cors/angular-product-app` shows you how to invoke on it from Javascript.



# Chapter 17. Events

Keycloak provides an Events SPI that makes it possible to register listeners for user related events, for example user logins. There are two interfaces that can be implemented, the first is a pure listener, the second is an events store which listens for events, but is also required to store events. An events store provides a way for the admin and account management consoles to view events.

## 17.1. Event types

Login events:

- Login - A user has logged in
- Register - A user has registered
- Logout - A user has logged out
- Code to Token - An application/client has exchanged a code for a token
- Refresh Token - An application/client has refreshed a token

Account events:

- Social Link - An account has been linked to a social provider
- Remove Social Link - A social provider has been removed from an account
- Update Email - The email address for an account has changed
- Update Profile - The profile for an account has changed
- Send Password Reset - A password reset email has been sent
- Update Password - The password for an account has changed
- Update TOTP - The TOTP settings for an account has changed
- Remove TOTP - TOTP has been removed from an account
- Send Verify Email - A email verification email has been sent
- Verify Email - The email address for an account has been verified

For all events there is a corresponding error event.

## 17.2. Event Listener

Keycloak comes with an Email Event Listener and a JBoss Logging Event Listener. The Email Event Listener sends an email to the users account when an event occurs. The JBoss Logging Event Listener writes to a log file when an event occurs.

The Email Event Listener only supports the following events at the moment:

- Login Error
- Update Password
- Update TOTP
- Remove TOTP

You can exclude one or more events by editing `standalone/configuration/keycloak-server.json` and adding for example:

```
"eventListener": {
  "email": {
    "exclude-events": [ "UPDATE_TOTP", "REMOVE_TOTP" ]
  }
}
```

### 17.3. Event Store

Event Store listen for events and is expected to persist the events to make it possible to query for them later. This is used by the admin console and account management to view events. Keycloak includes providers to persist events to JPA and Mongo.

You can specify events to include or exclude by editing `standalone/configuration/keycloak-server.json`, and adding for example:

```
"eventsStore": {
  "jpa": {
    "exclude-events": [ "LOGIN", "REFRESH_TOKEN", "CODE_TO_TOKEN" ]
  }
}
```

### 17.4. Configure Events Settings for Realm

To enable persisting of events for a realm you first need to make sure you have a event store provider registered for Keycloak. By default the JPA event store provider is registered. Once you've done that open the admin console, select the realm you're configuring, select `Events`. Then click on `Config`. You can enable storing events for your realm by toggling `Save Events` to `ON`. You can also set an expiration on events. This will periodically delete events from the database that are older than the specified time.

To configure listeners for a realm on the same page as above add one or more event listeners to the `Listeners` select box. This will allow you to enable any registered event listeners with the realm.



# Chapter 18. User Federation SPI and LDAP/AD Integration

Keycloak can federate external user databases. Out of the box we have support for LDAP and Active Directory. Before you dive into this, you should understand how Keycloak does federation.

Keycloak performs federation a bit differently than other products/projects. The vision of Keycloak is that it is an out of the box solution that should provide a core set of feature irregardless of the backend user storage you want to use. Because of this requirement/vision, Keycloak has a set data model that all of its services use. Most of the time when you want to federate an external user store, much of the metadata that would be needed to provide this complete feature set does not exist in that external store. For example your LDAP server may only provide password validation, but not support TOTP or user role mappings. The Keycloak User Federation SPI was written to support these completely variable configurations.

The way user federation works is that Keycloak will import your federated users on demand to its local storage. How much metadata that is imported depends on the underlying federation plugin and how that plugin is configured. Some federation plugins may only import the username into Keycloak storage, others might import everything from name, address, and phone number, to user role mappings. Some plugins might want to import credentials directly into Keycloak storage and let Keycloak handle credential validation. Others might want to handle credential validation themselves. The goal of the Federation SPI is to support all of these scenarios.

## 18.1. LDAP and Active Directory Plugin

Keycloak comes with a built-in LDAP/AD plugin. Currently it is set up only to import username, email, first and last name. It supports password validation via LDAP/AD protocols and different user metadata synchronization modes. To configure a federated LDAP store go to the admin console. Click on the `Users` menu option to get you to the user management page. Then click on the `Federation` submenu option. When you get to this page there is an "Add Provider" select box. You should see "ldap" within this list. Selecting "ldap" will bring you to the ldap configuration page.

### 18.1.1. Edit Mode

Edit mode defines various synchronization options with your LDAP store depending on what privileges you have.

#### READONLY

Username, email, first and last name will be unchangable. Keycloak will show an error anytime anybody tries to update these fields. Also, password updates will not be supported.

#### WRITABLE

Username, email, first and last name, and passwords can all be updated and will be synchronized automatically with your LDAP store.

### UNSYNCED

Any changes to username, email, first and last name, and passwords will be stored in Keycloak local storage. It is up to you to figure out how to synchronize back to LDAP.

### 18.1.2. Other config options

#### Display Name

Name used when this provider is referenced in the admin console

#### Priority

The priority of this provider when looking up users or for adding registrations.

#### Sync Registrations

If a new user is added through a registration page or admin console, should the user be eligible to be synchronized to this provider.

#### Other options

The rest of the configuration options should be self explanatory. You can use tooltips in admin console to see some more details about them.

## 18.2. Sync of LDAP users to Keycloak

LDAP Federation Provider will automatically take care of synchronization (import) of needed LDAP users into Keycloak database. For example once you first authenticate LDAP user `john` from Keycloak UI, LDAP Federation provider will first import this LDAP user into Keycloak database and then authenticate against LDAP password.

Thing is that Federation Provider import just requested users by default, so if you click to `View all users` in Keycloak admin console, you will see just those LDAP users, which were already authenticated/requested by Keycloak.

If you want to sync all LDAP users into Keycloak database, you may configure and enable Sync, which is in admin console on same page like the configuration of Federation provider itself. There are 2 types of sync:

#### Full sync

This will synchronize all LDAP users into Keycloak DB. Those LDAP users, which already exist in Keycloak and were changed in LDAP directly will be updated in Keycloak DB (For example if user `Mary Kelly` was changed in LDAP to `Mary Doe`).

#### Changed users sync

This will check LDAP and it will sync into Keycloak just those users, which were created or updated in LDAP from the time of last sync.

In usual cases you may want to trigger full sync at the beginning, so you will import all LDAP users to Keycloak just once. Then you may setup periodic sync of changed users, so Keycloak will

periodically ask LDAP server for newly created or updated users and backport them to Keycloak DB. Also you may want to trigger full sync again after some longer time or setup periodic full sync as well.

In admin console, you can trigger sync directly or you can enable periodic changed or full sync.

## 18.3. Writing your own User Federation Provider

The keycloak examples directory contains an example of a simple User Federation Provider backed by a simple properties file. See `examples/providers/federation-provider`. Most of how to create a federation provider is explain directly within the example code, but some information is here too.

Writing a User Federation Provider starts by implementing the `UserFederationProvider` and `UserFederationProviderFactory` interfaces. Please see the Javadoc and example for complete details on on how to do this. Some important methods of note: `getUserByUsername()` and `getUserByEmail()` require that you query your federated storage and if the user exists create and import the user into Keycloak storage. How much metadata you import is fully up to you. This import is done by invoking methods on the object returned `KeycloakSession.userStorage()` to add and import user information. The `proxy()` method will be called whenever Keycloak has found an imported `UserModel`. This allows the federation provider to proxy the `UserModel` which is useful if you want to support external storage updates on demand.

After your code is written you must package up all your classes within a JAR file. This jar file must contain a file called `org.keycloak.models.UserFederationProviderFactory` within the `META-INF/services` directory of the JAR. This file is a list of fully qualified classnames of all implementations of `UserFederationProviderFactory`. This is how Keycloak discovers which providers have been deployment. Place the JAR in the keycloak WAR deployment in the `WEB-INF/lib` directory.



## Chapter 19. Export and Import

Export/import is useful especially if you want to migrate your whole Keycloak database from one environment to another or migrate to different database (For example from MySQL to Oracle). You can trigger export/import at startup of Keycloak server and it's configurable with System properties right now. The fact it's done at server startup means that no-one can access Keycloak UI or REST endpoints and edit Keycloak database on the fly when export or import is in progress. Otherwise it could lead to inconsistent results.

You can export/import your database either to:

- Encrypted ZIP file on local filesystem
- Directory on local filesystem
- Single JSON file on your filesystem

Encrypted ZIP is recommended as export contains many sensitive informations like passwords of your users (even if they are hashed), but also their email addresses, and especially private keys of the realms. Directory and Single JSON file are useful especially for testing as data in the files are not protected. On the other hand, it's useful if you want to look at all your data in JSON files directly.

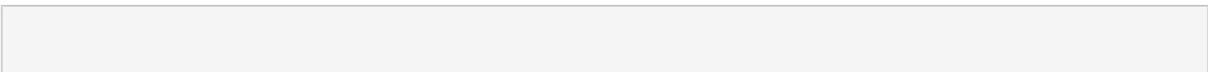
If you import to ZIP or Directory, you can specify also the number of users to be stored in each JSON file. So if you have very large amount of users in your database, you likely don't want to import them into single file as the file might be very big. Processing of each file is done in separate transaction as exporting/importing all users at once could also lead to memory issues.

So to export the content of your Keycloak database into encrypted ZIP, you can execute Keycloak server with the System properties like:

```
bin/standalone.sh -Dkeycloak.migration.action=export
-Dkeycloak.migration.provider=zip -Dkeycloak.migration.zipFile=<FILE TO EXPORT
TO>
-Dkeycloak.migration.zipPassword=<PASSWORD TO DECRYPT EXPORT>
```

Then you can move or copy the encrypted ZIP file into second environment and you can trigger import from it into Keycloak server with the same command but use `-Dkeycloak.migration.action=import` instead of `export`.

To export into unencrypted directory you can use:



```
bin/standalone.sh -Dkeycloak.migration.action=export
-Dkeycloak.migration.provider=dir -Dkeycloak.migration.dir=<DIR TO EXPORT TO>
```

And similarly for import just use `-Dkeycloak.migration.action=import` instead of `export` .

To export into single JSON file you can use:

```
bin/standalone.sh -Dkeycloak.migration.action=export
-Dkeycloak.migration.provider=singleFile -Dkeycloak.migration.file=<FILE TO
EXPORT TO>
```

Here's an example of importing:

```
bin/standalone.sh -Dkeycloak.migration.action=import
-Dkeycloak.migration.provider=singleFile -Dkeycloak.migration.file=<FILE TO
IMPORT>
-Dkeycloak.migration.strategy=OVERWRITE_EXISTING
```

Other available options are:

**-Dkeycloak.migration.realmName**

can be used if you want to export just one specified realm instead of all. If not specified, then all realms will be exported.

**-Dkeycloak.migration.usersExportStrategy**

can be used to specify for ZIP or Directory providers to specify where to import users. Possible values are:

- **DIFFERENT\_FILES** - Users will be exported into more different files according to maximum number of users per file. This is default value
- **SKIP** - exporting of users will be skipped completely
- **REALM\_FILE** - All users will be exported to same file with realm (So file like "foo-realm.json" with both realm data and users)
- **SAME\_FILE** - All users will be exported to same file but different than realm (So file like "foo-realm.json" with realm data and "foo-users.json" with users)

**-Dkeycloak.migration.usersPerFile**

can be used to specify number of users per file (and also per DB transaction). It's 5000 by default. It's used only if `usersExportStrategy` is `DIFFERENT_FILES`

---

-Dkeycloak.migration.strategy

is used during import. It can be used to specify how to proceed if realm with same name already exists in the database where you are going to import data. Possible values are:

- IGNORE\_EXISTING - Ignore importing if realm of this name already exists
- OVERWRITE\_EXISTING - Remove existing realm and import it again with new data from JSON file. If you want to fully migrate one environment to another and ensure that the new environment will contain same data like the old one, you can specify this.



## Chapter 20. Server Cache

By default, Keycloak caches realm metadata and users. There are two separate caches, one for realm metadata (realm, application, client, roles, etc...) and one for users. These caches greatly improves the performance of the server.

### 20.1. Disabling Caches

The realm and user caches can be disabled through configuration or through the management console. To manually disable the realm or user cache, you must edit the `keycloak-server.json` file in your distribution. Here's what the config looks like initially.

```
"realmCache": {
  "provider": "${keycloak.realm.cache.provider:mem}"
},

"userCache": {
  "provider": "${keycloak.user.cache.provider:mem}",
  "mem": {
    "maxSize": 20000
  }
},
```

You must then change it to:

```
"realmCache": {
  "provider": "${keycloak.realm.cache.provider:none}"
},

"userCache": {
  "provider": "${keycloak.user.cache.provider:none}"
},
```

You can also disable either of the caches at runtime through the Keycloak admin console Realm Settings page. This will not permanently disable the cache. If you reboot the server, the cache will be re-enabled unless you manually disable the cache in the `keycloak-server.json` file.

### 20.2. Clear Caches

To clear the realm or user cache, go to the Keycloak admin console Realm Settings->Cache Config page. Disable the cache you want. Save the settings. Then re-enable the cache. This will cause the cache to be cleared.

### 20.3. Cache Config

Cache configuration is done within `keycloak-server.json`. Changes to this file will not be seen by the server until you reboot. Currently you can only configure the max size of the user cache.

```
"userCache": {
  "provider": "${keycloak.user.cache.provider:mem}",
  "mem": {
    "maxSize": 20000
  }
},
```

## Chapter 21. SAML SSO

Keycloak supports SAML 2.0 for registered applications. Both POST and Redirect bindings are supported. You can choose to require client signature validation and can have the server sign and/or encrypt responses as well. We do not yet support logout via redirects. All logouts happen via a background POST binding request to the application that will be logged out. We do not support SAML 1.1 either. If you want support for either of those, please log a JIRA request and we'll schedule it.

When you create an application in the admin console, you can choose which protocol the application will log in with. In the application create screen, choose `saml` from the protocol list. After that there are a bunch of configuration options. Here is a description of each item:

### Include AuthnStatement

SAML login responses may specify the authentication method used (password, etc.) as well as a timestamp of the login. Setting this to on will include that statement in the response document.

### Multi-valued Roles

If this switch is off, any user role mappings will have a corresponding attribute created for it. If this switch is turned on, only one role attribute will be created, but it will have multiple values within it.

### Sign Documents

When turned on, Keycloak will sign the document using the realm's private key.

### Sign Assertions

With the `Sign Documents` switch signs the whole document. With this setting you just assign the assertions of the document.

### Signature Algorithm

Choose between a variety of algorithms for signing SAML documents.

### Encrypt Assertions

Encrypt assertions in SAML documents with the realm's private key. The AES algorithm is used with a key size of 128 bits.

### Client Signature Required

Expect that documents coming from a client are signed. Keycloak will validate this signature using the client keys set up in the `Application Keys` submenu item.

### Force POST Binding

By default, Keycloak will respond using the initial SAML binding of the original request. By turning on this switch, you will force Keycloak to always respond using the SAML POST Binding even if the original request was a the Redirect binding.

You have to specify an admin URL if you want logout to work. This should be a URL that will except single logout requests from the Keycloak server. You should also specify a default redirect url. Keycloak will redirect to this url after single logout is complete.

One thing to note is that roles are not treated as a hierarchy. So, any role mappings will just be added to the role attributes in the SAML document using their basic name. So, if you have multiple application roles you might have name collisions. You can use the Scope Mapping menu item to control which role mappings are set in the response.

### 21.1. SAML Entity Descriptor

If you go into the admin console in the application list menu page you will see an `Import` button. If you click on that you can import SAML Service Provider definitions using the [Entity Descriptor](http://docs.oasis-open.org/security/saml/v2.0/saml-metadata-2.0-os.pdf) [http://docs.oasis-open.org/security/saml/v2.0/saml-metadata-2.0-os.pdf] format described in SAML 2.0. You should review all the information there to make sure everything is set up correctly.

Each realm has a URL where you can view the XML entity descriptor for the IDP. `root/realms/{realm}/protocol/saml/descriptor`

## Chapter 22. Security Vulnerabilities

This chapter discusses possible security vulnerabilities Keycloak could have, how Keycloak mitigates those vulnerabilities, and what steps you need to do to configure Keycloak to mitigate some vulnerabilities. A good list of potential vulnerabilities and what security implementations should do to mitigate them can be found in the [OAuth 2.0 Thread Model](http://tools.ietf.org/html/rfc6819) [http://tools.ietf.org/html/rfc6819] document put out by the IETF. Many of those vulnerabilities are discussed here.

### 22.1. SSL/HTTPS Requirement

If you do not use SSL/HTTPS for all communication between the Keycloak auth server and the clients it secures you will be very vulnerable to man in the middle attacks. OAuth 2.0/OpenID Connect uses access tokens for security. Without SSL/HTTPS, attackers can sniff your network and obtain an access token. Once they have an access token they can do any operation that the token has been given permission for.

Keycloak has [three modes for SSL/HTTPS](#). SSL can be hard to set up, so out of the box, Keycloak allows non-HTTPS communication over private IP addresses like localhost, 192.168.x.x, and other private IP addresses. In production, you should make sure SSL is enabled and required across the board.

On the adapter/client side, Keycloak allows you to turn off the SSL trust manager. The trust manager ensures identity the client is talking to. It checks the DNS domain name against the server's certificate. In production you should make sure that each of your client adapters is configured to use a truststore. Otherwise you are vulnerable to DNS man in the middle attacks.

### 22.2. CSRF Attacks

Cross-site request forgery (CSRF) is a web-based attack whereby HTTP requests are transmitted from a user that the web site trusts or has authenticated (e.g., via HTTP redirects or HTML forms). Any site that uses cookie based authentication is vulnerable for these types of attacks. These attacks are mitigated by matching a state cookie against a posted form or query parameter.

OAuth 2.0 login specification requires that a state cookie be used and matched against a transmitted state parameter. Keycloak fully implements this part of the specification so all logins are protected.

The Keycloak administration console is a pure Javascript/HTML5 application that makes REST calls to the backend Keycloak admin API. These calls all require bearer token authentication and are made via Javascript Ajax calls. CSRF does not apply here. The admin REST API can also be configured to validate CORS origins as well.

The only part of Keycloak that really falls into CSRF is the user account management pages. To mitigate this Keycloak sets a state cookie and also embeds the value of this state cookie within hidden form fields or query parameters in action links. This query or form parameter is checked against the state cookie to verify that the call was made by the user.

## 22.3. Clickjacking

With clickjacking, a malicious site loads the target site in a transparent iFrame overlaid on top of a set of dummy buttons that are carefully constructed to be placed directly under important buttons on the target site. When a user clicks a visible button, they are actually clicking a button (such as an "Authorize" button) on the hidden page. An attacker can steal a user's authentication credentials and access their resources.

By default, every response by Keycloak sets some specific browser headers that can prevent this from happening specifically *X-FRAME\_OPTIONS* [<http://tools.ietf.org/html/rfc7034>] and *Content-Security-Policy* [<http://www.w3.org/TR/CSP/>]. You should take a look at both of these headers. In the admin console you can specify the values these headers will have. By default, Keycloak only sets up a same-origin policy for iframes.

## 22.4. Compromised Access Codes

It would be very hard for an attacker to compromise Keycloak access codes. Keycloak generates a cryptographically strong random value for its access codes so it would be very hard to guess an access token. An access code can only be turned into an access token once so it can't be replayed. In the admin console you can specify how long an access token is valid for. This value should be really short. Like a seconds. Just long enough for the client to make the request to turn the code into an token.

## 22.5. Compromised access and refresh tokens

There's a few things you can do to mitigate access tokens and refresh tokens from being stolen. Most importantly is to enforce SSL/HTTPS communication between Keycloak and its clients and applications. Short lifespans (minutes) for access tokens allows Keycloak to check the validity of a refresh token. Making sure refresh tokens always stay private to the client and are never transmitted ever is very important as well.

If an access token or refresh token is compromised, the first thing you should do is go to the admin console and push a not-before revocation policy to all applications. This will enforce that any tokens issued prior to that date are now invalid. You can also disable specific applications, clients, and users if you feel that any one of those entities is completely compromised.

## 22.6. Open redirectors

An attacker could use the end-user authorization endpoint and the redirect URI parameter to abuse the authorization server as an open redirector. An open redirector is an endpoint using a parameter to automatically redirect a user agent to the location specified by the parameter value without any validation. An attacker could utilize a user's trust in an authorization server to launch a phishing attack.

Keycloak requires that all registered applications and clients register at least one redirection uri pattern. Any time a client asks Keycloak to perform a redirect (on login or logout for example),

Keycloak will check the redirect uri vs. the list of valid registered uri patterns. It is important that clients and applications register as specific a URI pattern as possible to mitigate open redirector attacks.

## 22.7. Password guess: brute force attacks

A brute force attack happens when an attacker is trying to guess a user's password. Keycloak has some limited brute force detection capabilities. If turned on, a user account will be temporarily disabled if a threshold of login failures is reached. The downside of this is that this makes Keycloak vulnerable to denial of service attacks. Eventually we will expand this functionality to take client IP address into account when deciding whether to block a user.

Another thing you can do to prevent password guessing is to point a tool like [Fail2Ban](http://fail2ban.org) [http://fail2ban.org] to the Keycloak server's log file. Keycloak logs every login failure and client IP address that had the failure.

In the admin console, per realm, you can set up a password policy to enforce that users pick hard to guess passwords.

Finally, the best way to mitigate against brute force attacks is to require user to set up a one-time-password (OTP).

## 22.8. Password database compromised

Keycloak does not store passwords in raw text. It stores a hash of them. Because of performance reasons, Keycloak only hashes passwords once. While a human could probably never crack a hashed password, it is very possible that a computer could. The security community suggests around 20,000 (yes thousand!) hashing iterations to be done to each password. This number grows every year due to increasing computing power (It was 1000 12 years ago). The problem with this is that password hashing is a huge performance hit as each login would require the entered password to be hashed that many times and compared to the stored hash. So, its up to the admin to configure the password hash iterations. This can be done in the admin console password policy configuration. Again, the default value is 1 as we thought it might be more important for Keycloak to scale out of the box. There's a lot of other measures admins can do to protect their password databases.

## 22.9. SQL Injection attacks

At this point in time, there is no knowledge of any SQL injection vulnerabilities in Keycloak

## 22.10. Limiting Scope

Using the `scope` menu in the admin console for oauth clients or applications, you can control exactly which role mappings will be included within the token sent back to the client or application. This allows you to limit the scope of permissions given to the application or client which is great if the client isn't very trusted and is known to not being very careful with its tokens.



## Chapter 23. Clustering

To improve availability and scalability Keycloak can be deployed in a cluster.

It's fairly straightforward to configure a Keycloak cluster, the steps required are:

- Configure a shared database
- Configure Infinispan
- Enable realm and user cache invalidation
- Enable distributed user sessions
- Start in HA mode

### 23.1. Configure a shared database

Keycloak doesn't replicate realms and users, but instead relies on all nodes using the same database. This can be a relational database or Mongo. To make sure your database doesn't become a single point of failure you may also want to deploy your database to a cluster.

### 23.2. Configure Infinispan

Keycloak uses *Infinispan* [<http://www.infinispan.org/>] caches to share information between nodes.

For realm and users Keycloak uses a invalidation cache. An invalidation cache doesn't share any data, but simply removes stale data from remote caches. This reduces network traffic, as well as preventing sensitive data (such as realm keys and password hashes) from being sent between the nodes.

User sessions and login failures supports either distributed caches or fully replicated caches. We recommend using a distributed cache.

To configure the required Infinispan caches open `standalone/configuration/standalone-ha.xml` and add:

```
<subsystem xmlns="urn:jboss:domain:infinispan:2.0">
  <cache-container name="keycloak" jndi-name="infinispan/Keycloak"
    start="EAGER">
    <invalidation-cache name="realms" mode="SYNC" />
    <invalidation-cache name="users" mode="SYNC" />
    <distributed-cache name="sessions" mode="SYNC" owners="1" />
    <distributed-cache name="loginFailures" mode="SYNC" owners="1" />
  </cache-container>
```

```
...
</subsystem>
```

For more advanced options refer to the [Infinispan Subsystem](http://docs.jboss.org/author/display/WFLY8/Infinispan+Subsystem) [http://docs.jboss.org/author/display/WFLY8/Infinispan+Subsystem] and [Infinispan](http://www.infinispan.org/docs/6.0.x/user_guide/user_guide.html) [http://www.infinispan.org/docs/6.0.x/user\_guide/user\_guide.html] documentation.

Next open `standalone/configuration/keycloak-server.json` and add:

```
"connectionsInfinispan": {
  "default" : {
    "cacheContainer" : "java:jboss/infinispan/Keycloak"
  }
}
```

### 23.3. Enable realm and user cache invalidation

To reduce number of requests to the database Keycloak caches realm and user data. In cluster mode Keycloak uses an Infinispan invalidation cache to make sure all nodes re-load data from the database when it is changed. Using an invalidation cache instead of a replicated cache reduces the network traffic generated by the cluster, but more importantly prevents sensitive data from being sent.

To enable realm and user cache invalidation open `keycloak-server.json` and change the `realmCache` and `userCache` providers to `infinispan`:

```
"realmCache": {
  "provider": "infinispan"
},

"userCache": {
  "provider": "infinispan"
}
```

### 23.4. Enable distributed user sessions

To help distribute the load of user sessions Keycloak uses an Infinispan distributed cache. A distributed cache splits user sessions into segments where each node holds one or more segment. It is possible to replicate each segment to multiple nodes, but this is not strictly necessary since the failure of a node will only result in users having to log in again. If you need to prevent node

failures from requiring users to log in again, set the `owners` attribute to 2 or more for the `sessions` cache (see [Configure Infinispan](#)).

To enable the Infinispan user sessions provider open `keycloak-server.json` and change the `userSessions` provider to `infinispan`:

```
"userSessions": {  
  "provider": "infinispan"  
}
```

## 23.5. Start in HA mode

To start the server in HA mode, start it with:

```
# bin/standalone --server-config=standalone-ha.xml
```

Alternatively you can copy `standalone/config/standalone-ha.xml` to `standalone/config/standalone.xml` to make it the default server config.

## 23.6. Enabling cluster security

By default there's nothing to prevent unauthorized nodes from joining the cluster and sending potentially malicious messages to the cluster. However, as there's no sensitive data sent there's not much that can be achieved. For realms and users all that can be done is to send invalidation messages to make nodes load data from the database more frequently. For user sessions it would be possible to modify existing user sessions, but creating new sessions would have no affect as they would not be linked to any access tokens. There's not to much that can be achieved by modifying user sessions. For example it would be possible to prevent sessions from expiring, by changing the creation time. However, it would for example have no effect adding additional permissions to the sessions as these are rechecked against the user and application when the token is created or refreshed.

In either case your cluster nodes should be in a private network, with a firewall protecting them from outside attacks. Ideally isolated from workstations and laptops. You can also enable encryption of cluster messages, this could for example be useful if you can't isolate cluster nodes from workstations and laptops on your private network. However, encryption will obviously come at a cost of reduced performance.

To enable encryption of cluster messages you first have to create a shared keystore (change the key and store passwords!):

```
# keytool -genseckey -alias keycloak -keypass <PASSWORD> -storepass <PASSWORD> \
-keyalg Blowfish -keysize 56 -keystore defaultStore.keystore -storetype JCEKS
```

Copy this keystore to all nodes (for example to standalone/configuration). Then configure JGroups to encrypt all messages by adding the `ENCRYPT` protocol to the JGroups sub-system (this should be added after the `pbcast.GMS` protocol):

```
<subsystem xmlns="urn:jboss:domain:jgroups:2.0" default-stack="udp">
  <stack name="udp">
    ...
    <protocol type="pbcast.GMS"/>
    <protocol type="ENCRYPT">
      <property name="key_store_name">
        ${jboss.server.config.dir}/defaultStore.keystore
      </property>
      <property name="key_password">PASSWORD</property>
      <property name="store_password">PASSWORD</property>
      <property name="alias">keycloak</property>
    </protocol>
    ...
  </stack>
  <stack name="tcp">
    ...
    <protocol type="pbcast.GMS"/>
    <protocol type="ENCRYPT">
      <property name="key_store_name">
        ${jboss.server.config.dir}/defaultStore.keystore
      </property>
      <property name="key_password">PASSWORD</property>
      <property name="store_password">PASSWORD</property>
      <property name="alias">keycloak</property>
    </protocol>
    ...
  </stack>
  ...
</subsystem>
```

See the *JGroups manual* [<http://www.jgroups.org/manual/index.html#ENCRYPT>] for more details.

# Chapter 24. Application Clustering

This chapter is focused on clustering support for your own AS7, EAP6 or Wildfly applications, which are secured by Keycloak. We support various deployment scenarios according if your application is:

- stateless or stateful
- distributable (replicated http session) or non-distributable and just relying on sticky sessions provided by loadbalancer
- deployed on same or different cluster hosts where keycloak servers are deployed

The situation is a bit tricky as application communicates with Keycloak directly within user's browser (for example redirecting to login screen), but there is also backend (out-of-bound) communication between keycloak and application, which is hidden from end-user and his browser and hence can't rely on sticky sessions.

## 24.1. Stateless token store

By default, the servlet web application secured by Keycloak uses HTTP session to store information about authenticated user account. This means that this info could be replicated across cluster and your application will safely survive failover of some cluster node.

However if you don't need or don't want to use HTTP Session, you may alternatively save all info about authenticated account into cookie. This is useful especially if your application is:

- stateless application without need of HTTP Session, but with requirement to be safe to failover of some cluster node
- stateful application, but you don't want sensitive token data to be saved in HTTP session
- stateless application relying on loadbalancer, which is not aware of sticky sessions (in this case cookie is your only way)

To configure this, you can add this line to configuration of your adapter in `WEB-INF/keycloak.json` of your application:

```
"token-store": "cookie"
```

Default value of `token-store` is `session`, hence saving data in HTTP session. One disadvantage of cookie store is, that whole info about account is passed in cookie

KEYCLOAK\_ADAPTER\_STATE in each HTTP request. Hence it's not the best for network performance.

## 24.2. Relative URI optimization

In many deployment scenarios will be Keycloak and secured applications deployed on same cluster hosts. For this case Keycloak already provides option to use relative URI as value of option *auth-server-url* in `WEB-INF/keycloak.json`. In this case, the URI of Keycloak server is resolved from the URI of current request.

For example if your loadbalancer is on `https://loadbalancer.com/myapp` and *auth-server-url* is `/auth`, then relative URI of Keycloak is resolved to be `https://loadbalancer.com/myapp`.

For cluster setup, it may be even better to use option *auth-server-url-for-backend-request*. This allows to configure that backend requests between Keycloak and your application will be sent directly to same cluster host without additional round-trip through loadbalancer. So for this, it's good to configure values in `WEB-INF/keycloak.json` like this:

```
"auth-server-url": "/auth",  
"auth-server-url-for-backend-requests": "http://${jboss.host.name}:8080/auth"
```

This would mean that browser requests (like redirecting to Keycloak login screen) will be still resolved relatively to current request URI like `https://loadbalancer.com/myapp`, but backend (out-of-bound) requests between keycloak and your app are send always to same cluster host with application.

Note that additionally to network optimization, you may not need "https" in this case as application and keycloak are communicating directly within same cluster host.

## 24.3. Admin URL configuration

Admin URL for particular application can be configured in Keycloak admin console. It's used by Keycloak server to send backend requests to application for various tasks, like logout users or push revocation policies.

For example logout of user from Keycloak works like this:

1. User sends logout request from one of applications where he is logged.
2. Then application will send logout request to Keycloak
3. Keycloak server logout user in itself, and then it re-sends logout request by backend channel to all applications where user is logged. Keycloak is using admin URL for this. So logout is propagated to all apps.

You may again use relative values for admin URL, but in cluster it may not be the best similarly like in [previous section](#) .

Some examples of possible values of admin URL are:

`http://${jboss.host.name}:8080/myapp`

This is best choice if "myapp" is deployed on same cluster hosts like Keycloak and is distributable. In this case Keycloak server sends logout request to itself, hence no communication with loadbalancer or other cluster nodes and no additional network traffic.

Note that since the application is distributable, the backend request sent by Keycloak could be served on any application cluster node as invalidation of HTTP Session on *node1* will propagate the invalidation to other cluster nodes due to replicated HTTP sessions.

`http://${application.session.host}:8080/myapp`

Keycloak will track hosts where is particular HTTP Session served and it will send session invalidation message to proper cluster node.

For example application is deployed on `http://node1:8080/myapp` and `http://node2:8080/myapp` . Now HTTP Session *session1* is sticky-session served on cluster node *node2* . When keycloak invalidates this session, it will send request directly to `http://node2:8080/myapp` .

This is ideal configuration for distributable applications deployed on different host than keycloak or for non-distributable applications deployed either on same or different nodes than keycloak. Good thing is that it doesn't send requests through load-balancer and hence helps to reduce network traffic.

## 24.4. Registration of application nodes to Keycloak

Previous section describes how can Keycloak send logout request to proper application node. However in some cases admin may want to propagate admin tasks to all registered cluster nodes, not just one of them. For example push new notBefore for realm or application, or logout all users from all applications on all cluster nodes.

In this case Keycloak should be aware of all application cluster nodes, so it could send event to all of them. To achieve this, we support auto-discovery mechanism:

1. Once new application node joins cluster, it sends registration request to Keycloak server
2. The request may be re-sent to Keycloak in configured periodic intervals
3. If Keycloak won't receive re-registration request within specified timeout (should be greater than period from point 2) then it automatically unregister particular node
4. Node is also unregistered in Keycloak when it sends unregistration request, which is usually during node shutdown or application undeployment. This may not work properly for forced shutdown when undeployment listeners are not invoked, so here you need to rely on automatic unregistration from point 3 .

Sending startup registrations and periodic re-registration is disabled by default, as its main usecase is just cluster deployment. In `WEB-INF/keycloak.json` of your application, you can specify:

```
"register-node-at-startup": true,  
"register-node-period": 600,
```

which means that registration is sent at startup (accurately when 1st request is served by the application node) and then it's resent each 10 minutes.

In Keycloak admin console you can specify the maximum node re-registration timeout (makes sense to have it bigger than *register-node-period* from adapter configuration for particular application). Also you can manually add and remove cluster nodes in admin console, which is useful if you don't want to rely on adapter's automatic registration or if you want to remove stale application nodes, which weren't unregistered (for example due to forced shutdown).

### 24.5. Refresh token in each request

By default, application adapter tries to refresh access token when it's expired (period can be specified as *Access Token Lifespan*). However if you don't want to rely on the fact, that Keycloak is able to successfully propagate admin events like logout to your application nodes, then you have possibility to configure adapter to refresh access token in each HTTP request.

In `WEB-INF/keycloak.json` you can configure:

```
"always-refresh-token": true
```

Note that this has big performance impact. It's useful just if performance is not priority, but security is critical and you can't rely on logout and push `notBefore` propagation from Keycloak to applications.

# Chapter 25. Migration from older versions

To upgrade to a new version of Keycloak first download and install the new version of Keycloak. You then have to migrate the database, keycloak-server.json, providers, themes and applications from the old version.

## 25.1. Migrate database

Keycloak provides automatic migration of the database. It's highly recommended that you backup your database prior to upgrading Keycloak.

To enable automatic upgrading of the database if you're using a relational database make sure `databaseSchema` is set to `update` for `connectionsJpa`:

```
"connectionsJpa": {
  "default": {
    ...
    "databaseSchema": "update"
  }
}
```

For MongoDB do the same, but for `connectionsMongo`:

```
"connectionsMongo": {
  "default": {
    ...
    "databaseSchema": "update"
  }
}
```

When you start the server with this setting your database will automatically be migrated if the database schema has changed in the new version.

## 25.2. Migrate keycloak-server.json

You should copy `standalone/configuration/keycloak-server.json` from the old version to make sure any configuration changes you've done are added to the new installation. The version specific section below will list any changes done to this file that you have to do when upgrading from one version to another.

## 25.3. Migrate providers

If you have implemented any SPI providers you need to copy them to the new server. The version specific section below will mention if any of the SPI's have changed. If they have you may have to update your code accordingly.

## 25.4. Migrate themes

If you have created a custom theme you need to copy them to the new server. The version specific section below will mention if changes have been made to themes. If there is you may have to update your themes accordingly.

## 25.5. Migrate application

If you deploy applications directly to the Keycloak server you should copy them to the new server. For any applications including those not deployed directly to the Keycloak server you should upgrade the adapter. The version specific section below will mention if any changes are required to applications.

## 25.6. Version specific migration

### 25.6.1. Migrating from 1.0.x.Final to 1.1.Beta1

- RealmModel JPA and Mongo storage schema has changed
- UserSessionModel JPA and Mongo storage schema has changed as these interfaces have been refactored
- Upgrade your adapters, old adapters are not compatible with Keycloak 1.1. We interpreted JSON Web Token and OIDC ID Token specification incorrectly. 'aud' claim must be the client id, we were storing the realm name in there and validating it.

### 25.6.2. Migrating from 1.0 RC-1 to RC-2

- A lot of info level logging has been changed to debug. Also, a realm no longer has the jboss-logging audit listener by default. If you want log output when users login, logout, change passwords, etc. enable the jboss-logging audit listener through the admin console.

### 25.6.3. Migrating from 1.0 Beta 4 to RC-1

- logout REST API has been refactored. The GET request on the logout URI does not take a session\_state parameter anymore. You must be logged in in order to log out the session. You can also POST to the logout REST URI. This action requires a valid refresh token to perform

the logout. The signature is the same as refresh token minus the grant type form parameter. See documentation for details.

#### 25.6.4. Migrating from 1.0 Beta 1 to Beta 4

- LDAP/AD configuration is changed. It is no longer under the "Settings" page. It is now under Users->Federation. Add Provider will show you an "ldap" option.
- Authentication SPI has been removed and rewritten. The new SPI is UserFederationProvider and is more flexible.
- `ssl-not-required` property in adapter config has been removed. Replaced with `ssl-required`, valid values are `all` (require SSL for all requests), `external` (require SSL only for external request) and `none` (SSL not required).
- DB Schema has changed again.
- Created applications now have a full scope by default. This means that you don't have to configure the scope of an application if you don't want to.
- Format of JSON file for importing realm data was changed. Now role mappings is available under the JSON record of particular user.

#### 25.6.5. Migrating from 1.0 Alpha 4 to Beta 1

- DB Schema has changed. We have added export of the database to Beta 1, but not the ability to import the database from older versions. This will be supported in future releases.
- For all clients except bearer-only applications, you must specify at least one redirect uri. Keycloak will not allow you to log in unless you have specified a valid redirect uri for that application.
- Resource Owner Password Credentials flow is now disabled by default. It can be enabled by setting the toggle for `Direct Grant API ON` under realm config in the admin console.
- Configuration is now done through `standalone/configuration/keycloak-server.json`. This should mainly affect those that use MongoDB.
- JavaScript adapter has been refactored. See the [JavaScript adapter](#) section for more details.
- The "Central Login Lifespan" setting no longer exists. Please see the [Session Timeout](#) section for me details.

#### 25.6.6. Migrating from 1.0 Alpha 2 to Alpha 3

- `SkeletonKeyToken`, `SkeletonKeyScope`, `SkeletonKeyPrincipal`, and `SkeletonKeySession` have been renamed to: `AccessToken`, `AccessScope`, `KeycloakPrincipal`, and `KeycloakAuthenticatedSession` respectively.

- `ServeOAuthClient.getBearerToken()` method signature has changed. It now returns an `AccessTokenResponse` so that you can obtain a refresh token too.
- Adapters now check the access token expiration with every request. If the token is expired, they will attempt to invoke a refresh on the auth server using a saved refresh token.
- Subject in `AccessToken` has been changed to the User ID.

### 25.6.7. Migrating from 1.0 Alpha 1 to Alpha 2

- DB Schema has changed. We don't have any data migration utilities yet as of Alpha 2.
- JBoss and Wildfly adapters are now installed via a JBoss/Wildfly subsystem. Please review the adapter installation documentation. Edits to `standalone.xml` are now required.
- There is a new credential type "secret". Unlike other credential types, it is stored in plain text in the database and can be viewed in the admin console.
- There is no longer required Application or OAuth Client credentials. These client types are now hard coded to use the "secret" credential type.
- Because of the "secret" credential change to Application and OAuth Client, you'll have to update your `keycloak.json` configuration files and regenerate a secret within the Application or OAuth Client credentials tab in the administration console.