



Gateway Guide

/ ForgeRock Identity Gateway 5

Latest update: 5.0.0

Paul Bryan
Mark Craig
Jamie Nelson
Guillaume Sauthier
Joanne Henry

ForgeRock AS
201 Mission St., Suite 2900
San Francisco, CA 94105, USA
+1 415-599-1100 (US)
www.forgerock.com

Copyright © 2011-2017 ForgeRock AS.

Abstract

Instructions for installing and configuring the ForgeRock® Identity Gateway.



This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License.

To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/3.0/> or send a letter to Creative Commons, 444 Castro Street, Suite 900, Mountain View, California, 94041, USA.

ForgeRock® and ForgeRock Identity Platform™ are trademarks of ForgeRock Inc. or its subsidiaries in the U.S. and in other countries. Trademarks are the property of their respective owners.

UNLESS OTHERWISE MUTUALLY AGREED BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

DejaVu Fonts

Bitstream Vera Fonts Copyright

Copyright (c) 2003 by Bitstream, Inc. All Rights Reserved. Bitstream Vera is a trademark of Bitstream, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distribute the Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or characters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Bitstream" or the word "Vera".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been modified and is distributed under the "Bitstream Vera" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL BITSTREAM OR THE GNOME FOUNDATION BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the names of Gnome, the Gnome Foundation, and Bitstream Inc., shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Font Software without prior written authorization from the Gnome Foundation or Bitstream Inc., respectively. For further information, contact: fonts at gnome dot org.

Arev Fonts Copyright

Copyright (c) 2006 by Tavmjong Bah. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distribute the modifications to the Bitstream Vera Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or characters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Tavmjong Bah" or the word "Arev".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been modified and is distributed under the "Tavmjong Bah Arev" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL TAVMJONG BAH BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the name of Tavmjong Bah shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Font Software without prior written authorization from Tavmjong Bah. For further information, contact: tavmjong @ free . fr.

FontAwesome Copyright

Copyright (c) 2017 by Dave Gandy, <http://fontawesome.io>.

This Font Software is licensed under the SIL Open Font License, Version 1.1. This license is available with a FAQ at: <http://scripts.sil.org/OFL>.

Table of Contents

Preface	vii
1. About This Guide	vii
2. Formatting Conventions	viii
3. Accessing Documentation Online	viii
4. Using the ForgeRock.org Site	ix
5. Getting Support and Contacting ForgeRock	ix
1. About the ForgeRock Identity Gateway	1
1.1. About OpenIG	1
1.2. OpenIG Object Model	2
1.3. Configuration Directories and Files	3
1.4. Routing and Routes	3
1.5. Handlers, Filters, and Chains	4
1.6. Configuration Objects	8
1.7. Decorators	9
1.8. Configuration Parameters Declared as Property Variables	9
1.9. Using Comments in OpenIG Configuration Files	9
1.10. Understanding OpenIG APIs With API Descriptors	10
2. Getting Started	14
2.1. Before You Begin	14
2.2. Install OpenIG	14
2.3. Install the Sample Application	15
2.4. Configure OpenIG	16
2.5. Configure the Network	18
2.6. Try the Installation	19
3. Installation in Detail	22
3.1. Configuring Deployment Containers	22
3.2. Preparing the Network	28
3.3. Installing OpenIG	29
3.4. Changing the Default Location of the Configuration Folders	30
3.5. Preparing For Load Balancing and Failover	30
3.6. Configuring OpenIG For HTTPS (Client-Side)	32
3.7. Setting Up Keys For JWT Encryption	34
3.8. Making the Configuration Immutable	35
4. Getting Login Credentials From Data Sources	38
4.1. Before You Start	38
4.2. Log in With Credentials From a File	38
4.3. Log in With Credentials From a Database	41
5. Getting Login Credentials From Access Management	46
5.1. Detailed Flow	46
5.2. Setup Summary	48
5.3. Preparing the Tutorial	48
5.4. Setting Up OpenAM for Password Replay	50
5.5. Installing the OpenAM Policy Agent	51
5.6. Setting Up OpenIG for Password Replay	52

5.7. Testing the Setup	54
6. Enforcing Policy Decisions and Supporting Session Upgrade	55
6.1. About OpenIG As a PEP With OpenAM As PDP	55
6.2. Enforcing Policy Decisions From OpenAM	56
6.3. Upgrading a Session	60
7. Acting As a SAML 2.0 Service Provider	64
7.1. About SAML 2.0 SSO and Federation	64
7.2. Installation Overview	65
7.3. Preparing the Network	66
7.4. Configuring OpenAM As an IDP	67
7.5. Configuring OpenIG As an SP	68
7.6. Testing the Configuration	72
7.7. Example Federation Configuration Files	74
8. Acting As an OAuth 2.0 Resource Server	80
8.1. About OpenIG As an OAuth 2.0 Resource Server	80
8.2. Preparing the Tutorial	81
8.3. Setting Up OpenAM As an Authorization Server	82
8.4. Configuring OpenIG As a Resource Server	83
8.5. Testing the Configuration	85
9. Acting As an OAuth 2.0 Client or OpenID Connect Relying Party	88
9.1. About OpenIG As an OAuth 2.0 Client	88
9.2. About OpenIG As an OpenID Connect 1.0 Relying Party	88
9.3. Installation Overview	89
9.4. Setting Up OpenAM for OpenID Connect	90
9.5. Setting Up OpenIG As a Relying Party	91
9.6. Testing the Configuration	94
9.7. Adapting the Configuration to Authenticate Automatically to the Sample Application	95
9.8. Using OpenID Connect Discovery and Dynamic Client Registration	95
10. Transforming OpenID Connect ID Tokens Into SAML Assertions	101
10.1. About Token Transformation	101
10.2. Installation Overview	102
10.3. Setting Up OpenAM for Token Transformation	103
10.4. Setting Up OpenIG Routes for Token Transformation	106
11. Supporting UMA Resource Servers	113
11.1. About OpenIG in the UMA Resource Server Role	113
11.2. Preparing the Tutorial	117
11.3. Setting Up OpenAM As an Authorization Server	118
11.4. Setting Up OpenIG As a UMA Resource Server	121
11.5. Test the Configuration	125
11.6. Editing the Example to Match Custom Settings	126
12. Configuring Routers and Routes	127
12.1. Configuring Routers	127
12.2. Configuring Routes	128
12.3. Creating and Editing Routes Through Common REST	130
12.4. Creating Routes Through OpenIG Studio	132
12.5. Preventing the Reload of Routes	133

12.6. Accessing Reserved Routes	133
13. Configuration Templates	134
13.1. Proxy and Capture	134
13.2. Simple Login Form	135
13.3. Login Form With Cookie From Login Page	136
13.4. Login Form With Password Replay and Cookie Filters	137
13.5. Login Which Requires a Hidden Value From the Login Page	139
13.6. HTTP and HTTPS Application	141
13.7. OpenAM Integration With Headers	142
13.8. Microsoft Online Outlook Web Access	143
14. Extending the ForgeRock Identity Gateway	146
14.1. About Scripting	146
14.2. Scripting Dispatch	147
14.3. Scripting HTTP Basic Authentication	149
14.4. Scripting LDAP Authentication	151
14.5. Scripting SQL Queries	154
14.6. Developing Custom Extensions	157
15. Auditing and Monitoring	163
15.1. Monitoring Routes	163
15.2. Recording Audit Event Messages	166
16. Throttling the Rate of Requests to Protected Applications	173
16.1. Configuring a Simple Throttling Filter	173
16.2. Configuring a Mapped Throttling Filter	176
16.3. Configuring a Scriptable Throttling Filter	178
16.4. Dynamic Throttling Rate	181
17. Logging Events	183
17.1. Default Logging Behavior	183
17.2. Reference Logback Configuration	183
17.3. Changing the Logging Behavior	184
18. Troubleshooting	187
18.1. Troubleshooting the UMA Example	187
18.2. Can't Deploy Routes in OpenIG Studio	188
18.3. Object not found in heap	188
18.4. Extra or missing character / invalid JSON	188
18.5. The values in the flat file are incorrect	188
18.6. Problem accessing URL	189
18.7. StaticResponseHandler results in a blank page	189
18.8. OpenIG is not logging users in	189
18.9. Read timed out error when sending a request	189
18.10. OpenIG does not use new route configuration	190
18.11. Make OpenIG skip a route	190
A. SAML 2.0 and Multiple Applications	192
A.1. Installation Overview	192
A.2. Preparing the Network	193
A.3. Configuring the Circle of Trust	193
A.4. Configuring the Service Provider for Application One	193
A.5. Configuring the Service Provider for Application Two	199

A.6. Importing Service Provider Configurations Into OpenAM	199
A.7. Preparing OpenIG Configurations	199
A.8. Test the Configuration	203

Preface

ForgeRock Identity Platform™ is the only offering for access management, identity management, user-managed access, directory services, and an identity gateway, designed and built as a single, unified platform.

The platform includes the following components that extend what is available in open source projects to provide fully featured, enterprise-ready software:

- ForgeRock Access Management (AM)
- ForgeRock Identity Management (IDM)
- ForgeRock Directory Services (DS)
- ForgeRock Identity Gateway (IG)
- ForgeRock Identity Message Broker (IMB)

1. About This Guide

ForgeRock Identity Gateway integrates web applications, APIs, and microservices with the ForgeRock Identity Platform, without modifying the application or the container where they run. Based on reverse proxy architecture, it enforces security and access control in conjunction with the Access Management modules.

This guide is for access management designers and administrators who develop, build, deploy, and maintain ForgeRock Identity Gateway for their organizations. It helps you to get started quickly, and learn more as you progress through the guide.

This guide assumes basic familiarity with the following topics:

- Hypertext Transfer Protocol (HTTP), including how clients and servers exchange messages, and the role that a reverse proxy (gateway) plays
- JavaScript Object Notation (JSON), which is the format for ForgeRock Identity Gateway configuration files
- Managing services on operating systems and application servers
- Configuring network connections on operating systems
- Managing Public Key Infrastructure (PKI) used to establish HTTPS connections

- Access management for web applications

Depending on the features you use, you should also have basic familiarity with the following topics:

- Lightweight Directory Access Protocol (LDAP) if you use ForgeRock Identity Gateway with LDAP directory services
- Structured Query Language (SQL) if you use ForgeRock Identity Gateway with relational databases
- Configuring OpenAM if you use password capture and replay, or if you plan to follow the OAuth 2.0 or SAML 2.0 tutorials
- The Groovy programming language if you plan to extend ForgeRock Identity Gateway with scripts
- The Java programming language if you plan to extend ForgeRock Identity Gateway with plugins, and Apache Maven for building plugins

2. Formatting Conventions

Most examples in the documentation are created in GNU/Linux or Mac OS X operating environments. If distinctions are necessary between operating environments, examples are labeled with the operating environment name in parentheses. To avoid repetition file system directory names are often given only in UNIX format as in `/path/to/server`, even if the text applies to `C:\path\to\server` as well.

Absolute path names usually begin with the placeholder `/path/to/`. This path might translate to `/opt/`, `C:\Program Files\`, or somewhere else on your system.

Command-line, terminal sessions are formatted as follows:

```
$ echo $JAVA_HOME
/path/to/jdk
```

Command output is sometimes formatted for narrower, more readable output even though formatting parameters are not shown in the command.

Program listings are formatted as follows:

```
class Test {
    public static void main(String [] args) {
        System.out.println("This is a program listing.");
    }
}
```

3. Accessing Documentation Online

ForgeRock publishes comprehensive documentation online:

- The ForgeRock Knowledge Base offers a large and increasing number of up-to-date, practical articles that help you deploy and manage ForgeRock software.

While many articles are visible to community members, ForgeRock customers have access to much more, including advanced information for customers using ForgeRock software in a mission-critical capacity.

- ForgeRock product documentation, such as this document, aims to be technically accurate and complete with respect to the software documented. It is visible to everyone and covers all product features and examples of how to use them.

4. Using the ForgeRock.org Site

The [ForgeRock.org](https://www.forgerock.org) site has links to source code for ForgeRock open source software, as well as links to the ForgeRock forums and technical blogs.

If you are a *ForgeRock customer*, raise a support ticket instead of using the forums. ForgeRock support professionals will get in touch to help you.

5. Getting Support and Contacting ForgeRock

ForgeRock provides support services, professional services, training through ForgeRock University, and partner services to assist you in setting up and maintaining your deployments. For a general overview of these services, see <https://www.forgerock.com>.

ForgeRock has staff members around the globe who support our international customers and partners. For details, visit <https://www.forgerock.com>, or send an email to ForgeRock at info@forgerock.com.

Chapter 1

About the ForgeRock Identity Gateway

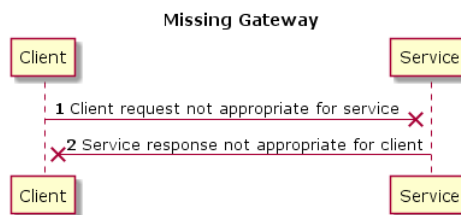
This chapter sets out the essentials of using OpenIG, including:

- What problems OpenIG solves and where it fits in your deployment
- How OpenIG acts on HTTP requests and responses
- How the configuration files for OpenIG are organized
- The roles played by routes, filters, handlers, and chains, which are the building blocks of an OpenIG configuration

1.1. About OpenIG

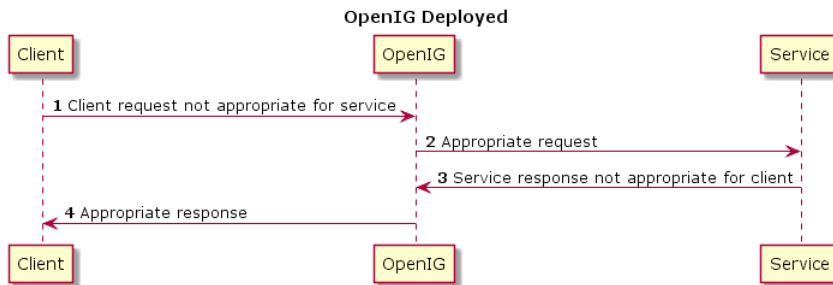
Most organizations have valuable existing services that are not easily integrated into newer architectures. These existing services cannot often be changed. Many client applications cannot communicate as they lack a gateway to bridge the gap. Figure 1.1, "Missing Gateway" illustrates one example of a missing gateway.

Figure 1.1. Missing Gateway



OpenIG works as an HTTP gateway, based on reverse proxy architecture. OpenIG is deployed on a network so it can intercept both client requests and server responses. Figure 1.2, "OpenIG Deployed" illustrates a OpenIG deployment.

Figure 1.2. OpenIG Deployed



Clients interact with protected servers through OpenIG. OpenIG can be configured to add new capabilities to existing services without affecting current clients or servers.

The list that follows features you can add to your solution by using OpenIG:

- Access management integration
- Application and API security
- Credential replay
- OAuth 2.0 support
- OpenID Connect 1.0 support
- Network traffic control
- Proxy with request and response capture
- Request and response rewriting
- SAML 2.0 federation support
- Single sign-on (SSO)

OpenIG supports these capabilities as out of the box configuration options. Once you understand the essential concepts covered in this chapter, try the additional instructions in this guide to use OpenIG to add other features.

1.2. OpenIG Object Model

OpenIG handles HTTP requests and responses in user-defined chains, making it possible to manage and to monitor processing at any point in a chain. The OpenIG object model provides both access to

the requests and responses that pass through each chain, and also context information associated with each request.

Contexts provide information about the client making the request, the session, the authentication or authorization identity of the principal, and any other state information associated with the request. Contexts provide a means to access state information throughout the duration of the HTTP session between the client and protected application, including when this involves interaction with additional services.

1.3. Configuration Directories and Files

By default, OpenIG configuration files are located under `$HOME/.openig` on Linux, macOS, and UNIX systems, and `%appdata%\OpenIG` on Windows systems. For information about how to change the default locations, see Section 3.4, "Changing the Default Location of the Configuration Folders".

OpenIG uses the following configuration directories:

- `$HOME/.openig/config`, `%appdata%\OpenIG\config`

OpenIG administration and gateway configuration files. For information, see `AdminHttpApplication(5)` in the *Configuration Reference* and `GatewayHttpApplication(5)` in the *Configuration Reference*.

- `$HOME/.openig/config/routes`, `%appdata%\OpenIG\config\routes`

OpenIG route configuration files. For more information see Chapter 12, "Configuring Routers and Routes".

- `$HOME/.openig/SAML`, `%appdata%\OpenIG\SAML`

OpenIG SAML 2.0 configuration files. For more information see Chapter 7, "Acting As a SAML 2.0 Service Provider".

- `$HOME/.openig/scripts/groovy`, `%appdata%\OpenIG\scripts\groovy`

OpenIG script files, for Groovy scripted filters and handlers. For more information see Chapter 14, "Extending the ForgeRock Identity Gateway".

- `$HOME/.openig/tmp`, `%appdata%\OpenIG\tmp`

OpenIG temporary files. This location can be used for temporary storage.

1.4. Routing and Routes

Routers are handlers that perform the following tasks:

- Define the routes directory and loads routes into the OpenIG configuration.

- Depending on the scanning interval, periodically scan the routes directory and updates the OpenIG configuration when routes are added, removed, or changed.
- Route requests to the first route in the OpenIG configuration whose condition is satisfied.

Routes are configuration files that you add to OpenIG to manage requests. They are flat files in JSON format. You can add routes in the following ways:

- Manually into the filesystem. Many of the examples in this guide, use this method.
- Through Common REST commands. For information, see Section 12.3, "Creating and Editing Routes Through Common REST".
- For some OpenIG features, through OpenIG Studio. For information, see Section 12.4, "Creating Routes Through OpenIG Studio".

Every route must call a handler to process the request and produce a response to a request.

When a route has a condition, it can handle only requests that meet the condition. When a route has no condition, it can handle any request.

Routes inherit settings from their parent configurations. This means that you can configure global objects in the `config.json` heap, for example, and then reference the objects by name in any other OpenIG configuration.

For examples of route configurations see Chapter 12, "*Configuring Routers and Routes*". For information about the parameters for routers and routes, see `Router(5)` in the *Configuration Reference* and `Route(5)` in the *Configuration Reference*.

1.5. Handlers, Filters, and Chains

Handlers and filters are chained together to modify a request, the response, or the context:

- *Handler*: Either delegates to another handler, or produces a response.

One way to produce a response is to send a request to and receive a response from an external service. In this case, OpenIG acts as a client of the service, often on behalf of the client whose request initiated the request.

Another way to produce a response is to build a response either statically or based on something in the context. In this case, OpenIG plays the role of server, generating a response to return to the client.

For more information, see `Handlers` in the *Configuration Reference*.

- *Filter*: Either transforms data in the request, response, or context, or performs an action when the request or response passes through the filter.

A filter can leave the request, response, and contexts unchanged. For example, it can log the context as it passes through the filter. Alternatively, it can change request or response. For example, it can generate a static request to replace the client request, add a header to the request, or remove a header from a response.

For more information, see [Filters](#) in the *Configuration Reference*.

- **Chain**: A type of handler that dispatches processing to an ordered list of filters, and then to the handler.

A **Chain** can be placed anywhere in a configuration that a handler can be placed. Filters process the incoming request, pass it on to the next filter, and then to the handler. After the handler produces a response, the filters process the outgoing response as it makes its way to the client. Note that the same filter can process both the incoming request and the outgoing response but most filters do one or the other.

For more information, see [Chain\(5\)](#) in the *Configuration Reference*.

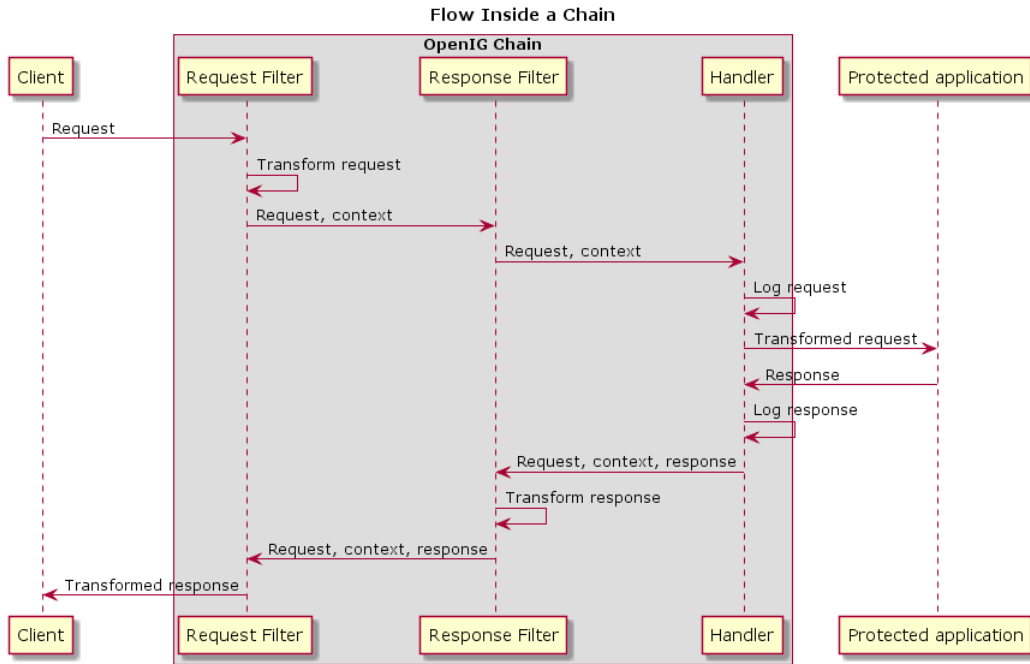
- **Chain of Filters**: A type of filter that dispatches processing to an ordered list of filters without then dispatching the request to a handler. Use this filter to assemble a list of filters into a single filter that you can then use in different places in the configuration.

A **ChainOfFilters** can be placed anywhere in a configuration that a filter can be placed.

For more information, see [ChainOfFilters\(5\)](#) in the *Configuration Reference*.

Figure 1.3, "Flow Inside a Chain" shows the flow inside a **Chain**, where a request filter transforms the request, a handler sends the request to a protected application, and then a response filter transforms the response. Notice how the flow traverses the filters in reverse order when the response comes back from the handler.

Figure 1.3. Flow Inside a Chain



The route configuration in Example 1.1, "Chain to a Protected Application" demonstrates the flow through a chain to a protected application. With OpenIG and the sample application set up as described in Chapter 2, "Getting Started", you can access this route on <http://openig.example.com:8080>.

Example 1.1. Chain to a Protected Application

```

{
  "handler": {
    "type": "Chain",
    "comment": "Base configuration defines the capture decorator",
    "config": {
      "filters": [
        {
          "type": "HeaderFilter",
          "comment": "Add a header to all requests",
          "config": {
            "messageType": "REQUEST",
            "add": {
              "MyHeaderFilter_request": [
                "Added by HeaderFilter to request"
              ]
            }
          }
        }
      ]
    }
  }
}

```

```
    ],
    "handler": {
      "type": "ClientHandler",
      "comment": "Log the request, pass it to the protected application,
                 and then log the response",
      "capture": "all",
      "baseURI": "http://app.example.com:8081"
    }
  }
}
```

The chain receives the request and context and processes it as follows:

- The first `HeaderFilter` adds a header to the incoming request.
- The second `HeaderFilter` is configured to manage responses, not requests, so it simply passes the request and context to the handler.
- The `ClientHandler` captures (logs) the request.
- The `ClientHandler` passes the transformed request to the protected application.
- The protected application passes a response to the `ClientHandler`.
- The `ClientHandler` captures (logs) the response.
- The second `HeaderFilter` adds a header added to the response.
- The first `HeaderFilter` is configured to manage requests, not responses, so it simply passes the response back to OpenIG.

Example 1.1, "Chain to a Protected Application" explained how a chain processes a request and its context. Example 1.2, "Requests and Responses in a Chain" illustrates the HTTP requests and responses captured as they flow through the chain.

Example 1.2. Requests and Responses in a Chain

```
### Original request from user-agent
GET http://openig.example.com:8080/ HTTP/1.1
Accept: */*
Host: openig.example.com:8080

### Add a header to the request (inside OpenIG) and direct it to the protected application
GET http://app.example.com:8081/ HTTP/1.1
Accept: */*
Host: openig.example.com:8080
MyHeaderFilter_request: Added by HeaderFilter to request

### Return the response to the user-agent
HTTP/1.1 200 OK
Content-Length: 1809
Content-Type: text/html; charset=ISO-8859-1

### Add a header to the response (inside OpenIG)
HTTP/1.1 200 OK
Content-Length: 1809
MyHeaderFilter_response: Added by HeaderFilter to response
```

1.6. Configuration Objects

Configuration objects have the following parts:

- *Name*: a unique string in the list of objects. When you declare inline objects, the name is not required.
- *Type*: the type name of the configuration object. OpenIG defines many object types for different purposes.
- *Config*: additional configuration settings. The content of the configuration object depends on its type.

If all of the configuration settings for the type are optional, the config field is also optional. The following configurations signify that the object uses default settings:

- Omitting the config field
- Setting the config field to an empty object, `"config": {}`
- Setting `"config": null`

Filters, handlers, and other objects whose configuration settings are defined by strings, integers, or booleans, can alternatively be defined by expressions that match the expected type.

1.7. Decorators

Decorators are additional heap objects to extend what another object can do. For example, a *CaptureDecorator* extends the capability of filters and handlers to log requests and responses. A *TimerDecorator* logs processing times. Decorate configuration objects with decorator names as field names.

OpenIG defines the following decorators: `audit`, `baseURI`, `capture`, and `timer`. You can use these decorators without configuring them explicitly.

You can log requests, responses, and processing times by adding decorations as shown in the following example:

```
{
  "handler": {
    "type": "Router",
    "capture": [ "request", "response" ],
    "timer": true
  }
}
```

For more information, see *Decorators* in the *Configuration Reference*.

1.8. Configuration Parameters Declared as Property Variables

Configuration parameters, such as host names, port numbers, and directories, can be declared as property variables in the OpenIG configuration or in an external JSON file. The variables can then be used in expressions in routes and in `config.json` to set the value of configuration parameters.

Properties can be inherited across the router, so a property defined in `config.json` can be used in any of the routes in the configuration.

Storing the configuration centrally and using variables for parameters that can be different for each installation makes it easier to deploy OpenIG in different environments without changing a single line in your route configuration.

For more information, see *Properties(5)* in the *Configuration Reference*.

1.9. Using Comments in OpenIG Configuration Files

The JSON format does not specify a notation for comments. If OpenIG does not recognize a JSON field name, it ignores the field. As a result, it is possible to use comments in configuration files.

Use the following conventions when commenting to ensure your configuration files are easier to read:

- Use `comment` fields to add text comments. Figure 1.4, "Using a Comment Field" illustrates a *CaptureDecorator* configuration that includes a text comment.

Figure 1.4. Using a Comment Field

```
{
  "name": "capture",
  "type": "CaptureDecorator",
  "comment": "Write request and response information to the logs",
  "config": {
    "captureEntity": true
  }
}
```

- Use an underscore (`_`) to comment a field temporarily. Figure 1.5, "Using an Underscore" illustrates a `CaptureDecorator` that has `"captureEntity": true` commented out. As a result, it uses the default setting (`"captureEntity": false`).

Figure 1.5. Using an Underscore

```
{
  "name": "capture",
  "type": "CaptureDecorator",
  "config": {
    "_captureEntity": true
  }
}
```

1.10. Understanding OpenIG APIs With API Descriptors

Common REST endpoints in OpenIG serve API descriptors at runtime. When you retrieve an API descriptor for an endpoint, a JSON that describes the API for that endpoint is returned.

To help you discover and understand APIs, you can use the API descriptor with a tool such as Swagger UI to generate a web page that helps you to view and test the different endpoints.

When you start OpenIG, or add or edit routes, registered endpoint locations for the routes hosted by the main router are written in `$HOME/.openig/logs/route-system.log`. Endpoint locations for subroutes are written to other log files. To retrieve the API descriptor for a specific endpoint, append one of the following query string parameters to the endpoint:

- `_api`, to represent the API accessible over HTTP. This OpenAPI descriptor can be used with endpoints that are complete or partial URLs.

The returned JSON respects the OpenAPI specification and can be consumed by Swagger tools, such as Swagger UI.

- `_crestapi`, to provide a compact representation that is independent of the transport protocol. This CREST API descriptor cannot be used with partial URLs.

The returned JSON respects a ForgeRock proprietary specification dedicated to describe CREST endpoints.

For more information about CREST API descriptors, see Section 5.5, "Common REST API Documentation" in the *Configuration Reference*.

Example 1.3. Retrieving API Descriptors for a Router

With OpenIG running as described in Chapter 2, "Getting Started", run the following query to generate a JSON that describes the router operations supported by the endpoint:

```
http://openig.example.com:8080/openig/api/system/objects/_router/routes?_api
{
  "swagger": "2.0",
  "info": {
    "version": "OpenIG version",
    "title": "OpenIG"
  },
  "host": "0:0:0:0:0:0:1",
  "basePath": "/openig/api/system/objects/_router/routes",
  "tags": [{
    "name": "Routes Endpoint"
  }],
  . . .
}
```

Alternatively, generate a CREST API descriptor by using the `?_crestapi` query string.

Example 1.4. Retrieving API Descriptors for the UMA Service

With the UMA tutorial running as described in Chapter 11, "Supporting UMA Resource Servers", run the following query to generate a JSON that describes the UMA share API:

```
http://openig.example.com:8080/openig/api/system/objects/_router/routes/00-uma/objects/umaservice/share?_api
{
  "swagger": "2.0",
  "info": {
    "version": "OpenIG version",
    "title": "OpenIG"
  },
  "host": "0:0:0:0:0:0:1",
  "basePath": "/openig/api/system/objects/_router/routes/00-uma/objects/umaservice/share",
  "tags": [{
    "name": "Manage UMA Share objects"
  }],
  . . .
}
```

Alternatively, generate a CREST API descriptor by using the `?_crestapi` query string.

Example 1.5. Retrieving API Descriptors for the Monitoring Endpoint of a Route

With monitoring enabled for a route, run a query to generate a JSON that exposes monitoring information for the route. For example:

```
http://openig.example.com:8080/openig/api/system/objects/_router/routes/00-uma/monitoring?_api
{
  "swagger": "2.0",
  "info": {
    "version": "OpenIG version",
    "title": "OpenIG"
  },
  "host": "openig.example.com:8080",
  "basePath": "/openig/api/system/objects/_router/routes/00-uma/monitoring",
  "tags": [{
    "name": "Monitoring endpoint"
  }],
  . . .
}
```

Alternatively, generate a CREST API descriptor by using the `?_crestapi` query string.

For information about how to set up monitoring for a route, see Procedure 15.1, "To Monitor a Route".

Example 1.6. Retrieving API Descriptors for the Main Router

Run a query to generate a JSON that describes the API for the main router and its subsequent endpoints. For example:

```
http://openig.example.com:8080/openig/api/system/objects/_router?_api
{
  "swagger": "2.0",
  "info": {
    "version": "OpenIG version",
    "title": "OpenIG"
  },
  "host": "openig.example.com:8080",
  "basePath": "/openig/api/system/objects/_router",
  "tags": [{
    "name": "Monitoring endpoint"
  }, {
    "name": "Manage UMA Share objects"
  }, {
    "name": "Routes Endpoint"
  }],
  . . .
}
```

Because the above URL is a partial URL, you cannot use the `?_crestapi` query string to generate a CREST API descriptor.

Example 1.7. Retrieving API Descriptors for an OpenIG Instance

Run a query to generate a JSON that describes the APIs provided by the OpenIG instance that is responding to a request. For example:

```
http://openig.example.com:8080/openig/api?_api
{
  "swagger": "2.0",
  "info": {
    "version": "OpenIG version",
    "title": "OpenIG"
  },
  "host": "openig.example.com:8080",
  "basePath": "/openig/api",
  "tags": [{
    "name": "Internal Storage for UI Models"
  }, {
    "name": "Monitoring endpoint"
  }, {
    "name": "Manage UMA Share objects"
  }, {
    "name": "Routes Endpoint"
  }, {
    "name": "Server Info"
  }],
  . . .
}
```

If routes are added after the request is performed, they are not included in the returned JSON.

Because the above URL is a partial URL, you cannot use the `?_crestapi` query string to generate a CREST API descriptor.

Chapter 2

Getting Started

In this chapter, you will learn to:

- Quickly set up OpenIG on Jetty
- Configure OpenIG to protect a sample application
- Prepare OpenIG so that you can follow all subsequent tutorials in the documentation

This chapter allows you to see how OpenIG works, and provides hands-on experience with a few key features. For more general installation and configuration instructions, start with Chapter 3, "*Installation in Detail*".

2.1. Before You Begin

Make sure you have a supported Java Development Kit installed. For details, see Section 2.2, "JDK Version" in the *Release Notes*.

2.2. Install OpenIG

You install OpenIG in the root context of a web application container. In this chapter, you use Jetty server as the web application container.

To perform initial installation, follow these steps:

1. Download and unzip a supported version of Jetty server.

Supported versions are listed in Section 2.3, "Web Application Containers" in the *Release Notes*.

2. Download **IG-5.0.0.war** from the ForgeRock BackStage download site .
3. Copy (or move) and rename the .war file as follows:

```
$ cp IG-5.0.0.war /path/to/jetty/webapps/root.war
```

Jetty automatically deploys OpenIG in the root context on startup.

4. Start Jetty in the background:

```
$ /path/to/jetty/bin/jetty.sh start
```

Or start Jetty in the foreground:

```
$ cd /path/to/jetty/  
$ java -jar start.jar
```

5. Make sure that you can see the OpenIG welcome page at <http://localhost:8080>.

When you start OpenIG without a configuration, requests to OpenIG default to a welcome page with a link to the documentation.

6. Make sure that you can see OpenIG Studio at <http://localhost:8080/openig/studio>.

For more information about OpenIG Studio, see Section 12.4, "Creating Routes Through OpenIG Studio".

7. Display the product version and build information at <http://localhost:8080/openig/api/info>.

8. Stop Jetty in the background:

```
$ /path/to/jetty/bin/jetty.sh stop
```

Or stop Jetty in the foreground by entering Ctrl+C in the terminal where Jetty is running.

2.3. Install the Sample Application

ForgeRock provides a mockup web application for testing OpenIG configurations. The sample application is used in many of the tutorials in this guide. This section takes you through the steps to download and run the sample application.

1. Download the sample application, [IG-doc-samples-5.0.0.jar](#), from the [ForgeRock BackStage](#) download site .
2. Run the sample application:


```
$ java -jar IG-doc-samples-5.0.0.jar
Preparing to listen for HTTP on port 8081.
Preparing to listen for HTTPS on port 8444.
The server will use a self-signed certificate not known to browsers.
When using HTTPS with curl for example, try --insecure.
Using OpenAM URL: http://openam.example.com:8088/openam/oauth2.
Starting server...
Sep 09, 2016 9:52:56 AM org.glassfish.grizzly.http.server.NetworkListener start
INFO: Started listener bound to [0.0.0.0:8444]
Sep 09, 2016 9:52:56 AM org.glassfish.grizzly.http.server.NetworkListener start
INFO: Started listener bound to [0.0.0.0:8081]
Sep 09, 2016 9:52:56 AM org.glassfish.grizzly.http.server.HttpServer start
INFO: [HttpServer] Started.
Press Ctrl+C to stop the server.
```

By default, this server listens for HTTP on port 8081, and for HTTPS on port 8444. If one or both of those ports are not free, specify other ports:

```
$ java -jar IG-doc-samples-5.0.0.jar 8888 8889
Preparing to listen for HTTP on port 8888.
Preparing to listen for HTTPS on port 8889.
The server will use a self-signed certificate not known to browsers.
When using HTTPS with curl for example, try --insecure.
Using OpenAM URL: http://openam.example.com:8088/openam/oauth2.
Starting server...
Sep 09, 2016 9:55:57 AM org.glassfish.grizzly.http.server.NetworkListener start
INFO: Started listener bound to [0.0.0.0:8889]
Sep 09, 2016 9:55:57 AM org.glassfish.grizzly.http.server.NetworkListener start
INFO: Started listener bound to [0.0.0.0:8888]
Sep 09, 2016 9:55:57 AM org.glassfish.grizzly.http.server.HttpServer start
INFO: [HttpServer] Started.
Press Ctrl+C to stop the server.
```

If you change the port numbers when starting the server, also account for the differences when using the examples.

3. Browse to <http://localhost:8081/home> to access the home page of the sample application.
Some information about the browser request is displayed.
4. Browse to <http://localhost:8081/login> to access the login page of the sample application, and then log in with username `demo` and password `changeit`.

The username, `demo`, and some information about the browser request is displayed.

2.4. Configure OpenIG

Now that you have installed both OpenIG and a sample application to protect, it is time to configure OpenIG.

Follow these steps to configure OpenIG to proxy traffic to the sample application:

1. Prepare the OpenIG configuration.

Add the following file as `$HOME/.openig/config/config.json`. By default, OpenIG looks for `config.json` in the `$HOME/.openig/config` directory:

```
{
  "handler": {
    "type": "Router",
    "name": "_router",
    "baseURI": "http://app.example.com:8081",
    "capture": "all"
  },
  "heap": [
    {
      "name": "JwtSession",
      "type": "JwtSession"
    },
    {
      "name": "capture",
      "type": "CaptureDecorator",
      "config": {
        "captureEntity": true,
        "_captureContext": true
      }
    }
  ]
}
```

```
$ mkdir -p $HOME/.openig/config
$ vi $HOME/.openig/config/config.json
```

On Windows, the configuration files belong in `%appdata%\OpenIG\config`. To locate the `%appdata%` folder for your version of Windows, open Windows Explorer, type `%appdata%` as the file path, and press Enter. You must create the `%appdata%\OpenIG\config` folder, and then copy the configuration files.

Important

If you plan to create routes through OpenIG Studio, make sure that this `config.json` contains a main router named `_router`. For information about OpenIG Studio, see Section 12.4, "Creating Routes Through OpenIG Studio".

If you adapt this base configuration for production use, make sure to adjust the log level, and to deactivate the `CaptureDecorator` that generates several log message lines for each request and response. Also consider editing the router based on recommendations described in Section 12.5, "Preventing the Reload of Routes".

2. Add the following default route configuration file as `$HOME/.openig/config/routes/zz-default.json`. By default, the router defined in the `config.json` looks for routes in the `$HOME/.openig/config/routes` directory:

```
{
```

```
"handler": "ClientHandler"  
}
```

```
$ mkdir $HOME/.openig/config/routes  
$ vi $HOME/.openig/config/routes/zz-default.json
```

On Windows, the file name should be `%appdata%\OpenIG\config\routes\zz-default.json`.

3. Start Jetty in the background:

```
$ /path/to/jetty/bin/jetty.sh start
```

Or start Jetty in the foreground:

```
$ cd /path/to/jetty/  
$ java -jar start.jar
```

2.5. Configure the Network

Because OpenIG uses reverse proxy architecture, you must configure the network so that traffic from the browser to the protected application goes through OpenIG.

If you followed the installation steps, you are running both OpenIG and the sample application on the same host as your browser (probably your laptop or desktop). Keep in mind that network configuration is an important deployment step. To encourage you to keep this in mind, the sample configuration for this chapter expects the sample application to be running on `app.example.com`, rather than `localhost`.

The quickest way to configure the network locally is to add an entry to your `/etc/hosts` file on UNIX systems or `%SystemRoot%\system32\drivers\etc\hosts` on Windows. See the Wikipedia entry, *Hosts (file)*, for more information on host files. If you are indeed running all servers in this chapter on the same host, add the following entry to the hosts file:

```
127.0.0.1    openig.example.com app.example.com
```

If you are running the browser and OpenIG on separate hosts, add the IP address of the host running OpenIG to the hosts file on the system running the browser, where the host name matches that of protected application. For example, if OpenIG is running on a host with IP address 192.168.0.15:

```
192.168.0.15    openig.example.com app.example.com
```

If OpenIG is on a different host from the protected application, also make sure that the host name of the protected application resolves correctly for requests from OpenIG to the application.

Restart Jetty to take the configuration changes into account.

Tip

Some browsers cache IP address resolutions, even after clearing all browsing data. Restart the browser after changing the IP addresses of named hosts.

The simplest way to make sure you have configured your DNS or host settings properly for remote systems is to stop OpenIG and then to make sure you cannot reach the target application with the host name and port number of OpenIG. If you can still reach it, double check your host settings.

Also make sure name resolution is configured to check host files before DNS. This configuration can be found in `/etc/nsswitch.conf` for most UNIX systems. Make sure `files` is listed before `dns`.

2.6. Try the Installation

`http://openig.example.com:8080/home` should take you to the home page of the sample application.

What just happened?

When your browser goes to `http://openig.example.com:8080/`, it is actually connecting to OpenIG deployed in Jetty. OpenIG proxies all traffic it receives to the protected application at `http://app.example.com:8081/`, and returns responses from the application to your browser. It does this based on the configuration that you set up.

Consider the base configuration file first, `config.json`. This file specifies a main router named `_router`. OpenIG calls this handler when it receives an incoming request.

The `baseURI` decoration in turn changes the request URI to point the request to the sample application to protect. The router captures the request on the way in, and captures the response on the way out.

The router routes processing to separate route configurations.

For now the only route available is the the default route you added, `zz-default.json`. The default route calls a `ClientHandler` with the default configuration. This `ClientHandler` simply proxies the request to and the response from the sample application to protect without changing either the request or the response. Therefore, the browser request is sent unchanged to the sample application and the response from the sample application is returned unchanged to your browser.

Now change the OpenIG configuration to log you in automatically with hard-coded credentials:

1. Add a route to automatically log you in as username `demo`, password `changeit`.

Add the following route configuration file as `$HOME/.openig/config/routes/01-static.json`:

```
{
  "handler": {
    "type": "Chain",
```

```
"config": {
  "filters": [
    {
      "type": "StaticRequestFilter",
      "config": {
        "method": "POST",
        "uri": "http://app.example.com:8081/login",
        "form": {
          "username": [
            "demo"
          ],
          "password": [
            "changeit"
          ]
        }
      }
    }
  ],
  "handler": "ClientHandler"
},
"condition": "${matches(request.uri.path, '^/static')}}"
}
```

On Windows, the file name should be `%appdata%\OpenIG\config\routes\01-static.json`.

2. Access the new route, `http://openig.example.com:8080/static`.

This time, OpenIG logs you in automatically.

Also view the information logged about requests and responses, which shows up in the Jetty log.

What's happening behind the scenes?

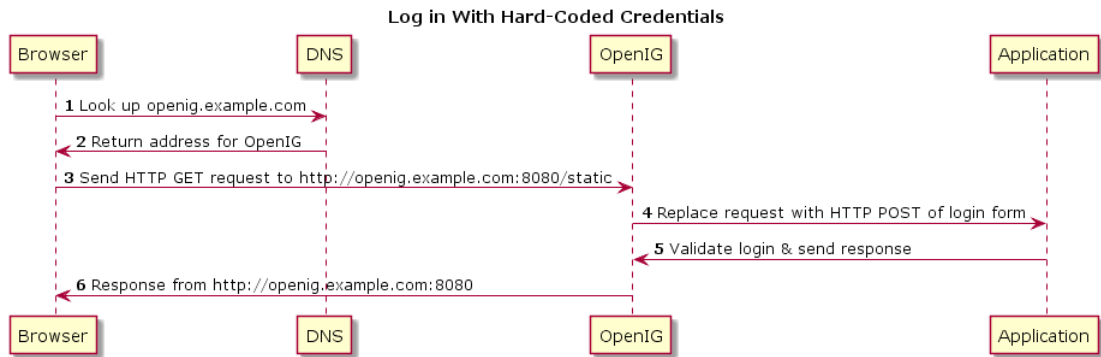
With the original configuration, OpenIG does not change requests or responses, but only proxies requests and responses, and captures request and response information.

After you change the configuration, OpenIG continues to capture request and response data. When your request does not go to the default route, but instead goes to `/static`, then the condition on the new route you added matches the request. OpenIG therefore uses the new route you added.

Using the route configuration in `01-static.json`, OpenIG replaces your browser's original HTTP GET request with an HTTP POST login request containing credentials to authenticate. As a result, instead of the login page with a login form, OpenIG logs you in directly, and the application responds with the page you see after logging in. OpenIG then returns this response to your browser.

Figure 2.1, "Log in With Hard-Coded Credentials" shows the steps.

Figure 2.1. Log in With Hard-Coded Credentials



1. The browser host makes a DNS request for the IP address of the HTTP server host, `app.example.com`.
2. DNS responds with the address for OpenIG.
3. Browser sends a request to the HTTP server.
4. OpenIG replaces the request with an HTTP POST request, including the login form with hard-coded credentials.
5. HTTP server validates the credentials, and responds with the profile page.
6. OpenIG passes the response back to the browser.

Chapter 3

Installation in Detail

This chapter describes how to do the following tasks:

- Prepare a deployment container for use with OpenIG (Section 3.1, "Configuring Deployment Containers").
- Prepare the network so that traffic passes through OpenIG (Section 3.2, "Preparing the Network").
- Download, deploy, and configure OpenIG (Section 3.3, "Installing OpenIG").
- Change the locations of the configuration files (Section 3.4, "Changing the Default Location of the Configuration Folders").
- Prepare for load balancing with OpenIG (Section 3.5, "Preparing For Load Balancing and Failover").
- Secure connections to and from OpenIG (Section 3.6, "Configuring OpenIG For HTTPS (Client-Side)").
- Use OpenIG JSON Web Token (JWT) Session cookies across multiple servers (Section 3.7, "Setting Up Keys For JWT Encryption").
- Prevent further updates to the configuration (Section 3.8, "Making the Configuration Immutable").

Before you begin to install or configure OpenIG, make sure that you are using a supported container and version of Java. For information about requirements for running OpenIG, see Chapter 2, "*Before You Install*" in the *Release Notes*.

3.1. Configuring Deployment Containers

This section provides installation and configuration tips that you need to run OpenIG in supported containers.

For the full list of supported containers see Section 2.3, "Web Application Containers" in the *Release Notes*.

For further information on advanced configuration for a particular container, see the container documentation.

3.1.1. About Securing Connections

OpenIG is often deployed to replay credentials or other security information. In a real world deployment, that information must be communicated over a secure connection using HTTPS, meaning in effect HTTP over encrypted Transport Layer Security (TLS). Never send real credentials, bearer tokens, or other security information unprotected over HTTP.

When OpenIG is acting as a server, the web application container where OpenIG runs is responsible for setting up TLS connections with client applications that connect to OpenIG. For details, see Section 3.1.3.2, "Configuring Jetty For HTTPS (Server-Side)" or Section 3.1.2.2, "Configuring Tomcat For HTTPS (Server-Side)".

When OpenIG is acting as a client, the `ClientHandler` configuration governs TLS connections from OpenIG to other servers. For details, see Section 3.6, "Configuring OpenIG For HTTPS (Client-Side)" and `ClientHandler(5)` in the *Configuration Reference*.

TLS depends on the use of digital certificates (public keys). In typical use of TLS, the client authenticates the server by its X.509 digital certificate as the first step to establishing communication. Once trust is established, then the client and server can set up a symmetric key to encrypt communications.

In order for the client to trust the server certificate, the client needs first to trust the certificate of the party who signed the server's certificate. This means that either the client has a trusted copy of the signer's certificate, or the client has a trusted copy of the certificate of the party who signed the signer's certificate.

Certificate Authorities (CAs) are trusted signers with well-known certificates. Browsers generally ship with many well-known CA certificates. Java distributions also ship with many well-known CA certificates. Getting a certificate signed by a well-known CA is often expensive.

It is also possible for you to self-sign certificates. The trade-off is that although there is no monetary expense, the certificate is not trusted by any clients until they have a copy. Whereas it is often enough to install a certificate signed by a well-known CA in the server keystore as the basis of trust for HTTPS connections, self-signed certificates must also be installed in all clients.

Like self-signed certificates, the signing certificates of less well-known CAs are also unlikely to be found in the default truststore. You might therefore need to install those signing certificates on the client side as well.

This guide describes how to install self-signed certificates, which are certainly fine for trying out the software and okay for deployments where you manage all clients that access OpenIG. If you need a well-known CA-signed certificate instead, see the documentation for your container for details on requesting a CA signature and installing the CA-signed certificate.

Once certificates are properly installed to allow client-server trust, also consider the cipher suites configured for use. The cipher suite used determines the security settings for the communication. Initial TLS negotiations bring the client and server to agreement on which cipher suite to use. Basically the client and server share their preferred cipher suites to compare and to choose. If you therefore have a preference concerning the cipher suites to use, you must set up your container to

use only your preferred cipher suites. Otherwise the container is likely to inherit the list of cipher suites from the underlying Java environment.

The Java Secure Socket Extension (JSSE), part of the Java environment, provides security services that OpenIG uses to secure connections. You can set security and system properties to configure the JSSE. For a list of properties you can use to customize the JSSE in Oracle Java, see the *Customization* section of the *JSSE Reference Guide*.

3.1.2. Configuring Apache Tomcat For OpenIG

This section describes essential Tomcat configuration that you need in order to run OpenIG.

Download and install a supported version of Tomcat from <http://tomcat.apache.org/>.

Important

If you use startup scripts to bootstrap the OpenIG web container, the scripts can start the container process with a different user. To prevent errors, make sure that the location of the OpenIG configuration is correct. Alternatively, adapt the startup scripts to specify the `OPENIG_BASE` env variable or `openig.base` system properties, taking care to set file permissions correctly.

If you start and stop the OpenIG web container yourself, the default location of the OpenIG configuration files is correct. By default, OpenIG configuration files are located under `$HOME/.openig` on Linux, Mac, and UNIX systems, and under `%appdata%\OpenIG` on Windows.

Configure Tomcat to use the same protocol as the application you are protecting with OpenIG. If the protected application is on a remote system, configure Tomcat to use the same port as well. If your application listens on both an HTTP and an HTTPS port, then you must configure Tomcat to do so, too.

To configure Tomcat to use an HTTP port other than 8080, modify the defaults in `/path/to/tomcat/conf/server.xml`. Search for the default value of 8080 and replace it with the new port number.

3.1.2.1. Configuring Tomcat Cookie Domains

To use OpenIG for multiple protected applications running on different hosts, set a cookie domain in Tomcat or `JwtSession`.

To set a cookie domain for an HTTP session (the default if you're not using a JWT session), add a context element to `/path/to/conf/Catalina/server/root.xml`, as in the following example, and then restart Tomcat to read the configuration changes:

```
<Context sessionCookieDomain=".example.com" />
```

To set a cookie domain for a JWT session, set the `cookieDomain` parameter in `JwtSession`. For information, see `JwtSession(5)` in the *Configuration Reference*. When the domain is set, a JWT cookie can be accessed from different hosts in that domain. When the domain is not set, the JWT cookie can be accessed only from the host where the cookie was created.

3.1.2.2. Configuring Tomcat For HTTPS (Server-Side)

To get Tomcat up quickly on an SSL port, add an entry similar to the following in `/path/to/tomcat/conf/server.xml`:

```
<Connector
  port="8443"
  protocol="HTTP/1.1"
  SSLEnabled="true"
  maxThreads="150"
  scheme="https"
  secure="true"
  address="127.0.0.1"
  clientAuth="false"
  sslProtocol="TLS"
  keystoreFile="/path/to/tomcat/conf/keystore"
  keystorePass="password"
/>
```

Also create a keystore holding a self-signed certificate:

```
$ keytool \  
-genkey \  
-alias tomcat \  
-keyalg RSA \  
-keystore /path/to/tomcat/conf/keystore \  
-storepass password \  
-keypass password \  
-dname "CN=openig.example.com,O=Example Corp,C=FR"
```

Notice the keystore file location and the keystore password both match the configuration. By default, Tomcat looks for a certificate with alias `tomcat`.

Restart Tomcat to read the configuration changes.

Browsers generally do not trust self-signed certificates. To work with a certificate signed instead by a trusted CA, see the Tomcat documentation on configuring HTTPS.

3.1.2.3. Configuring Tomcat to Access MySQL Over JNDI

If OpenIG accesses an SQL database, then you must configure Tomcat to access the database using Java Naming and Directory Interface (JNDI). To do so, you must add the driver `.jar` for the database, set up a JNDI data source, and set up a reference to that data source.

The following steps are for MySQL Connector/J:

1. Download the MySQL JDBC Driver Connector/J from <http://dev.mysql.com/downloads/connector/j>.
2. Copy the driver `.jar` to `/path/to/tomcat/lib/` so that it is on Tomcat's class path.
3. Add a JNDI data source for your MySQL server and database in `/path/to/tomcat/conf/context.xml`:

```
<Resource
  name="jdbc/forgerock"
  auth="Container"
  type="javax.sql.DataSource"
  maxActive="100"
  maxIdle="30"
  maxWait="10000"
  username="mysqladmin"
  password="password"
  driverClassName="com.mysql.jdbc.Driver"
  url="jdbc:mysql://localhost:3306/databasename"
/>
```

4. Add a resource reference to the data source in `/path/to/tomcat/conf/web.xml`:

```
<resource-ref>
  <description>MySQL Connection</description>
  <res-ref-name>jdbc/forgerock</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

5. Restart Tomcat to read the configuration changes.

3.1.3. Configuring Jetty For OpenIG

This section describes essential Jetty configuration that you need in order to run OpenIG.

Download and install a supported version of Jetty from <https://www.eclipse.org/jetty/download.html>.

Configure Jetty to use the same protocol as the application you are protecting with OpenIG. If the protected application is on a remote system, configure Jetty to use the same port as well. If your application listens on both an HTTP and an HTTPS port, then you must configure Jetty to do so as well.

To configure Jetty to use an HTTP port other than 8080, modify the defaults in `/path/to/jetty/etc/jetty.xml`. Search for the default value of 8080 and replace it with the new port number.

3.1.3.1. Configuring Jetty Cookie Domains

If you use OpenIG for more than a single protected application and the protected applications are on different hosts, then you must configure Jetty to set domain cookies. To do this, add a session domain handler element that specifies the domain to `/path/to/jetty/etc/webdefault.xml`, as in the following example:

```
<context-param>
  <param-name>org.eclipse.jetty.servlet.SessionDomain</param-name>
  <param-value>.example.com</param-value>
```

```
</context-param>
```

Restart Jetty to read the configuration changes.

3.1.3.2. Configuring Jetty For HTTPS (Server-Side)

To get Jetty up quickly on an SSL port, follow the steps in this section.

These steps involve replacing the built-in keystore with your own:

1. If you have not done so already, remove the built-in keystore:

```
$ rm /path/to/jetty/etc/keystore
```

2. Generate a new key pair with self-signed certificate in the keystore:

```
$ keytool \  
-genkey \  
-alias jetty \  
-keyalg RSA \  
-keystore /path/to/jetty/etc/keystore \  
-storepass password \  
-keypass password \  
-dname "CN=openig.example.com,O=Example Corp,C=FR"
```

3. Find the obfuscated form of the password:

```
$ java \  
-cp /path/to/jetty/lib/jetty-util-*.jar \  
org.eclipse.jetty.util.security.Password \  
password  
password  
0BF:1v2jluum1xtv1zej1zer1xtnl1uvk1v1v  
MD5:5f4dcc3b5aa765d61d8327deb882cf99
```

4. Edit the SSL Context Factory entry in the Jetty configuration file, `/path/to/jetty/etc/jetty-ssl.xml`:

```
<New id="sslContextFactory" class="org.eclipse.jetty.http.ssl.SslContextFactory">  
  <Set name="KeyStore"><Property name="jetty.home" default="." />/etc/keystore</Set>  
  <Set name="KeyStorePassword">0BF:1v2jluum1xtv1zej1zer1xtnl1uvk1v1v</Set>  
  <Set name="KeyManagerPassword">0BF:1v2jluum1xtv1zej1zer1xtnl1uvk1v1v</Set>  
  <Set name="TrustStore"><Property name="jetty.home" default="." />/etc/keystore</Set>  
  <Set name="TrustStorePassword">0BF:1v2jluum1xtv1zej1zer1xtnl1uvk1v1v</Set>  
</New>
```

5. Uncomment the line specifying that configuration file in `/path/to/jetty/start.ini`:

```
etc/jetty-ssl.xml
```

6. Restart Jetty.

7. Browse <https://openig.example.com:8443>.

You should see a warning in the browser that the (self-signed) certificate is not recognized.

3.1.3.3. Configuring Jetty to Access MySQL Over JNDI

If OpenIG accesses an SQL database, then you must configure Jetty to access the database over JNDI. To do so, you must add the driver .jar for the database, set up a JNDI data source, and set up a reference to that data source.

The following steps are for MySQL Connector/J:

1. Download the MySQL JDBC Driver Connector/J from <http://dev.mysql.com/downloads/connector/j>.
2. Copy the driver .jar to `/path/to/jetty/lib/jndi/` so that it is on Jetty's class path.
3. Add a JNDI data source for your MySQL server and database in `/path/to/jetty/etc/jetty.xml`:

```
<New id="jdbc/forgerock" class="org.eclipse.jetty.plus.jndi.Resource">
  <Arg></Arg>
  <Arg>jdbc/forgerock</Arg>
  <Arg>
    <New class="com.mysql.jdbc.jdbc2.optional.MysqlConnectionPoolDataSource">
      <Set name="Url">jdbc:mysql://localhost:3306/databasename</Set>
      <Set name="User">mysqladmin</Set>
      <Set name="Password">password</Set>
    </New>
  </Arg>
</New>
```

4. Add a resource reference to the data source in `/path/to/jetty/etc/webdefault.xml`:

```
<resource-ref>
  <description>MySQL Connection</description>
  <res-ref-name>jdbc/forgerock</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

5. Restart Jetty to read the configuration changes.

3.2. Preparing the Network

Because OpenIG uses reverse proxy architecture, you must configure the network so that that traffic from the browser to the protected application goes through OpenIG.

Modify DNS or host file settings so that the host name of the protected application resolves to the IP address of OpenIG on the system where the browser runs.

Restart the browser after making this change.

3.3. Installing OpenIG

Follow these steps to install OpenIG:

1. Download `IG-5.0.0.war` from the ForgeRock BackStage download site .
2. Copy (or move) and rename the `.war` file *to the root context* of the web application container:
 - For Jetty, name the `.war` file `root.war`. Jetty automatically deploys OpenIG in the root context on startup.
 - For Tomcat, name the `.war` file `ROOT.war`.

Important

If you use startup scripts to bootstrap the OpenIG web container, the scripts can start the container process with a different user. To prevent errors, make sure that the location of the OpenIG configuration is correct. Alternatively, adapt the startup scripts to specify the `OPENIG_BASE` env variable or `openig.base` system properties, taking care to set file permissions correctly.

If you start and stop the OpenIG web container yourself, the default location of the OpenIG configuration files is correct. By default, OpenIG configuration files are located under `$HOME/.openig` on Linux, Mac, and UNIX systems, and under `%appdata%\OpenIG` on Windows.

3. Prepare your OpenIG configuration files.

For information about configuration files, see Section 1.3, "Configuration Directories and Files". For information about how to change the default location of the configuration files, see Section 3.4, "Changing the Default Location of the Configuration Folders".

If you have not yet prepared configuration files, then start with the configuration described in Section 2.4, "Configure OpenIG".

Copy the template to `$HOME/.openig/config/config.json`. Replace the baseURI of the Dispatcher with that of the protected application.

On Windows, copy the template to `%appdata%\OpenIG\config\config.json`. To locate the `%appdata%` folder for your version of Windows, open Windows Explorer, type `%appdata%` as the file path, and press Enter. You must create the `%appdata%\OpenIG\config` folder, and then add the configuration file.

4. Start the web container where OpenIG is deployed.
5. Browse to the protected application.

OpenIG should now proxy all traffic to the application.

6. Make sure the browser is going through OpenIG.

Verify this in one of the following ways:

- Follow these steps:
 1. Stop the OpenIG web container.
 2. Verify that you cannot browse to the protected application.
 3. Start the OpenIG web container.
 4. Verify that you can now browse to the protected application again.
- Check the logs to see that traffic is going through OpenIG.

3.4. Changing the Default Location of the Configuration Folders

Change `$HOME/.openig` (or `%appdata%\OpenIG`) from the default location in the following ways:

- Set the `OPENIG_BASE` environment variable to the full path to the base location for OpenIG files:

```
# On Linux, macOS, and UNIX using Bash
$ export OPENIG_BASE=/path/to/openig

# On Windows
C:>set OPENIG_BASE=c:\path\to\openig
```

- Set the `openig.base` Java system property to the full path to the base location for OpenIG files when starting the web application container where OpenIG runs, as in the following example that starts Jetty server in the foreground:

```
$ java -Dopenig.base=/path/to/openig -jar start.jar
```

3.5. Preparing For Load Balancing and Failover

For a high scale or highly available deployment, you can prepare a pool of OpenIG servers with nearly identical configurations, and then load balance requests across the pool, routing around any servers that become unavailable. Load balancing allows the service to handle more load.

Before you spread requests across multiple servers, however, you must determine what to do with state information that OpenIG saves in the context, or retrieves locally from the OpenIG server system. If information is retrieved locally, then consider setting up failover. If one server becomes unavailable, another server in the pool can take its place. The benefit of failover is that a server failure can be invisible to client applications.

OpenIG can save state information in several ways:

- Handlers including a `SamlFederationHandler` or a custom `ScriptableHandler` can store information in the context. Most handlers depend on information in the context, some of which is first stored by OpenIG.
- Some filters, such as `AssignmentFilters`, `HeaderFilters`, `OAuth2ClientFilters`, `OAuth2ResourceServerFilters`, `ScriptableFilters`, `SqlAttributesFilters`, and `StaticRequestFilters`, can store information in the context. Most filters depend on information in the request, response, or context, some of which is first stored by OpenIG.

OpenIG can also retrieve information locally in several ways:

- Some filters and handlers, such as `FileAttributesFilters`, `ScriptableFilters`, `ScriptableHandlers`, and `SqlAttributesFilters`, can depend on local system files or container configuration.

By default, the context data resides in memory in the container where OpenIG runs. This includes the default session implementation, which is backed by the `HttpSession` that the container handles. You can opt to store session data on the user-agent instead, however. For details and to consider whether your data fits, see `JwtSession(5)` in the *Configuration Reference*.

When you use the `JwtSession` implementation with a `cookieDomain` set, be sure to share the encryption keys and the `sharedSecret` across all OpenIG configurations so that any server can read or update JWT cookies from any other server in the same `cookieDomain`.

If your data does not fit in an HTTP cookie, for example, because when encrypted it is larger than 4 KB, consider storing a reference in the cookie, and then retrieve the data by using another filter. OpenIG logs warning messages if the `JwtSession` cookie is too large. Using a reference can also work when a server becomes unavailable, and the load balancer must fail requests over to another server in the pool.

If some data attached to a context must be stored on the server side, then you have additional configuration steps to perform for session stickiness and for session replication. Session stickiness means that the load balancer sends all requests from the same client session to the same server. Session stickiness helps to ensure that a client request goes to the server holding the original session data. Session replication involves writing session data either to other servers or to a data store, so that if one server goes down, other servers can read the session data and continue processing. Session replication helps when one server fails, allowing another server to take its place without having to start the session over again. If you set up session stickiness but not session replication, when a server crashes, the client session information for that server is lost, and the client must start again with a new session.

How you configure session stickiness and session replication depends on your load balancer and on your container.

Tomcat can help with session stickiness, and a Tomcat cluster can handle session replication:

- If you choose to use the Tomcat connector (`mod_jk`) on your web server to perform load balancing, then see the *LoadBalancer HowTo* for details.

In the [HowTo](#), you configure the `jvmRoute` attribute in the Tomcat server configuration, `/path/to/tomcat/conf/server.xml`, to identify the server. The connector can use this identifier to achieve session stickiness.

- A Tomcat cluster configuration can handle session replication. When setting up a cluster configuration, the `ClusterManager` defines the session replication implementation.

Jetty has provisions for session stickiness, and also for session replication through clustering:

- Jetty's persistent session mechanism appends a node ID to the session ID in the same way Tomcat appends the `jvmRoute` value to the session cookie. This can be useful for session stickiness if your load balancer examines the session ID.
- [Session Clustering with a Database](#) describes how to configure Jetty to persist sessions over JDBC, allowing session replication.

Unless it is set up to be highly available, the database can be a single point of failure in this case.

- [Session Clustering with MongoDB](#) describes how to configure Jetty to persist sessions in MongoDB, allowing session replication.

The Jetty documentation recommends this implementation when session data is seldom written but often read.

3.6. Configuring OpenIG For HTTPS (Client-Side)

For OpenIG to connect to a server securely over HTTPS, OpenIG must be able to trust the server. The default settings rely on the Java environment truststore to trust server certificates. The Java environment default truststore includes public key signing certificates from many well-known Certificate Authorities (CAs). If all servers present certificates signed by these CAs, then you have nothing to configure.

If, however, the server certificates are self-signed or signed by a CA whose certificate is not trusted out of the box, then you can configure a `KeyStore` and a `TrustManager`, and optionally, a `KeyManager` to reference when configuring an `ClientHandler` to enable OpenIG to trust servers when acting as a client.

For details, see:

- `ClientHandler(5)` in the *Configuration Reference*
- `KeyManager(5)` in the *Configuration Reference*
- `KeyStore(5)` in the *Configuration Reference*
- `TrustManager(5)` in the *Configuration Reference*

The KeyStore holds the servers' certificates or the CA's signing certificate. The TrustManager allows OpenIG to handle the certificates in the KeyStore when deciding whether to trust a server certificate. The optional KeyManager allows OpenIG to present its certificate from the keystore when the server must authenticate OpenIG as client. The ClientHandler references whatever TrustManager and KeyManager you configure.

You can configure each of these either globally, for the OpenIG server, or locally, for a particular ClientHandler configuration.

The Java KeyStore holds the peer servers' public key certificates (and optionally, the OpenIG certificate and private key). For example, suppose you have a certificate file, `ca.crt`, that holds the trusted signer's certificate of the CA who signed the server certificates of the servers in your deployment. In that case, you could import the certificate into a Java Keystore file, `/path/to/keystore.jks`:

```
$ keytool \  
-import \  
-trustcacerts \  
-keystore /path/to/keystore \  
-file ca.crt \  
-alias ca-cert \  
-storepass changeit
```

You could then configure the following KeyStore for OpenIG that holds the trusted certificate. Notice that the url field takes an expression that evaluates to a URL, starting with a scheme such as `file:///`:

```
{  
  "name": "MyKeyStore",  
  "type": "KeyStore",  
  "config": {  
    "url": "file:///path/to/keystore",  
    "password": "changeit"  
  }  
}
```

The TrustManager handles the certificates in the KeyStore when deciding whether to trust the server certificate. The TrustManager references your KeyStore:

```
{  
  "name": "MyTrustManager",  
  "type": "TrustManager",  
  "config": {  
    "keystore": "MyKeyStore"  
  }  
}
```

The `ClientHandler` configuration has the following security settings:

"trustManager"

This references the `TrustManager`.

Recall that you must configure this when your server certificates are not trusted out of the box.

"hostnameVerifier"

This defines how the `ClientHandler` verifies host names in server certificates.

By default, host name verification is turned off.

"keyManager"

This references the optional `KeyManager`.

Configure this if servers request that OpenIG present its certificate as part of mutual authentication.

In that case, generate a key pair for OpenIG, and have the certificate signed by a well-known CA. For instructions, see the documentation for the Java `keytool` command. You can use a different keystore for the `KeyManager` than you use for the `TrustManager`.

The following `ClientHandler` configuration references `MyTrustManager` and sets strict host name verification:

```
{
  "name": "ClientHandler",
  "type": "ClientHandler",
  "config": {
    "hostnameVerifier": "STRICT",
    "trustManager": "MyTrustManager"
  }
}
```

3.7. Setting Up Keys For JWT Encryption

You can use a JSON Web Token (JWT) session, `JwtSession`, to configure OpenIG as described in `JwtSession(5)` in the *Configuration Reference*. A `JwtSession` stores session information in JWT cookies on the user-agent, rather than storing the information in the container where OpenIG runs.

In order to encrypt the JWTs, OpenIG needs cryptographic keys. OpenIG can generate its own key pair in memory, but that key pair disappears on restart and cannot be shared across OpenIG servers. Alternatively, OpenIG can use keys from a keystore.

The following procedure prepares a keystore for JWT encryption in a deployment with one OpenIG instance. For a deployment with more than one OpenIG instance, also configure a `sharedSecret` as described in `JwtSession(5)` in the *Configuration Reference*.

Procedure 3.1. To Prepare a Keystore for JWT Encryption

1. Generate the key pair in a new keystore file by using the Java `keytool` command.

The following command generates a Java Keystore format file, `/path/to/keystore.jks`, holding a key pair with alias `jwe-key`. Notice that both the keystore and the private key have the same password:

```
$ keytool \  
-genkey \  
-alias jwe-key \  
-keyalg rsa \  
-keystore /path/to/keystore.jks \  
-storepass changeit \  
-keypass changeit \  
-dname "CN=openig.example.com,O=Example Corp"
```

2. Add a KeyStore to your configuration that references the keystore file:

```
{  
  "name": "MyKeyStore",  
  "type": "KeyStore",  
  "config": {  
    "url": "file:///path/to/keystore.jks",  
    "password": "changeit"  
  }  
}
```

For details, see `KeyStore(5)` in the *Configuration Reference*.

3. Add a JwtSession to your configuration that references your KeyStore:

```
{  
  "name": "MyJwtSession",  
  "type": "JwtSession",  
  "config": {  
    "keystore": "MyKeyStore",  
    "alias": "jwe-key",  
    "password": "changeit",  
    "cookieName": "OpenIG",  
    "cookieDomain": ".example.com"  
  }  
}
```

For a deployment with more than one OpenIG instance, also configure a `sharedSecret` as described in `JwtSession(5)` in the *Configuration Reference*.

4. Specify your JwtSession object in the top-level configuration, or in the route configuration:

```
"session": "MyJwtSession"
```

3.8. Making the Configuration Immutable

OpenIG operates in development mode (mutable mode) and production mode (immutable mode):

Development mode

Use development mode to evaluate or demo OpenIG, or to develop configurations on a single instance. This mode is not suitable for production.

In development mode, by default all endpoints are open and accessible. You can create, edit, and deploy routes through OpenIG Studio, and manage routes through Common REST, without authentication or authorization.

To protect specific endpoints in development mode, configure an `ApiProtectionFilter` in `admin.json` and add it to the OpenIG configuration.

Production mode

After you have developed your configuration, switch to production mode to test the configuration, to run the software in pre-production or production, or to run multiple instances of the software with the same configuration.

In production mode, the `/routes` endpoint is not exposed or accessible. OpenIG Studio is effectively disabled, and you cannot manage, list, or even read routes through Common REST.

By default, other endpoints, such as `/monitoring`, `/share`, and `api/info` are exposed to the loopback address only. To change the default protection for specific endpoints, configure an `ApiProtectionFilter` in `admin.json` and add it to the OpenIG configuration.

Procedure 3.2. To Switch From Development Mode to Production Mode

By default, OpenIG operates in development mode. After creating your configuration in development mode, switch to production mode to prevent any unwanted changes:

1. Add the following file to the OpenIG configuration as `$HOME/.openig/config/admin.json`.

On Windows, add the file to `%appdata%\OpenIG\config`:

```
{
  "mode": "PRODUCTION"
}
```

The file changes the operating mode from the default value of `DEVELOPMENT`, used when you don't provide an `admin.json` file, to `PRODUCTION`.

For more information, see `AdminHttpApplication(5)` in the *Configuration Reference*.

2. Prevent routes from being reloaded after startup:
 - To prevent all routes in the configuration from being reloaded, edit the main router in `config.json`.
 - To prevent individual routes from being reloaded, edit the routers in those routes.

```
{
  "type": "Router",
  "config": {
    "scanInterval": "disabled"
  }
}
```

For more information, see Router(5) in the *Configuration Reference*.

3. Restart OpenIG.

When OpenIG starts up, the routes endpoint is not available and is not displayed in the log. When you try to access OpenIG Studio on <http://openig.example.com:8080/openig/studio>, an HTTP 404 is returned.

Chapter 4

Getting Login Credentials From Data Sources

In Chapter 2, "*Getting Started*" you learned how to configure OpenIG to proxy traffic and capture request and response data. You also learned how to configure OpenIG to use a static request to log in with hard-coded credentials. In this chapter, you will learn to:

- Configure OpenIG to look up credentials in a file
- Configure OpenIG to look up credentials in a relational database

4.1. Before You Start

Before you start this tutorial, prepare OpenIG and the sample application as shown in Chapter 2, "*Getting Started*".

OpenIG should be running in Jetty, configured to access the sample application as described in that chapter.

4.2. Log in With Credentials From a File

This sample shows you how to configure OpenIG to get credentials from a file.

The sample uses a comma-separated value file, `userfile`:

```
username,password,fullname,email
george,costanza,George Costanza,george@example.com
kramer,newman,Kramer,kramer@example.com
bjensen,hifalutin,Babs Jensen,bjensen@example.com
demo,changeit,Demo User,demo@example.com
kvaughan,bribery,Kirsten Vaughan,kvaughan@example.com
scarter,sprain,Sam Carter,scarter@example.com
```

OpenIG looks up the user credentials based on the user's email address. OpenIG uses a `FileAttributesFilter` to look up the credentials.

Follow these steps to set up log in with credentials from a file:

1. Add the user file on your system:

```

$ vi /tmp/userfile
$ cat /tmp/userfile
username,password,fullname,email
george,costanza,George Costanza,george@example.com
kramer,newman,Kramer,kramer@example.com
bjensen,hifalutin,Babs Jensen,bjensen@example.com
demo,changeit,Demo User,demo@example.com
kvaughan,bribery,Kirsten Vaughan,kvaughan@example.com
scarter,sprain,Sam Carter,scarter@example.com
    
```

On Windows systems, use an appropriate path such as `C:\Temp\userfile`.

2. Add a new route to the OpenIG configuration to obtain the credentials from the file.

To add the route, add the following route configuration file as `$HOME/.openig/config/routes/02-file.json`:

```

{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "PasswordReplayFilter",
          "config": {
            "loginPage": "${true}",
            "credentials": {
              "type": "FileAttributesFilter",
              "config": {
                "file": "/tmp/userfile",
                "key": "email",
                "value": "george@example.com",
                "target": "${attributes.credentials}"
              }
            }
          },
          "request": {
            "method": "POST",
            "uri": "http://app.example.com:8081/login",
            "form": {
              "username": [
                "${attributes.credentials.username}"
              ],
              "password": [
                "${attributes.credentials.password}"
              ]
            }
          }
        }
      ]
    }
  },
  "handler": "ClientHandler"
}
    
```



```
"condition": "${matches(request.uri.path, '^/file')}"
}
```

On Windows, the file name should be %appdata%\OpenIG\config\routes\02-file.json.

Notice the following features of the new route:

- The `FileAttributesFilter` specifies the file to access, the key and value to look up to retrieve the user's record, and where to store the results in the request context attributes map.
- The `PasswordReplayFilter` creates a request by retrieving the username and password from the attributes map and replacing your browser's original HTTP GET request with an HTTP POST login request that contains the credentials to authenticate.
- The route matches requests to `/file`.

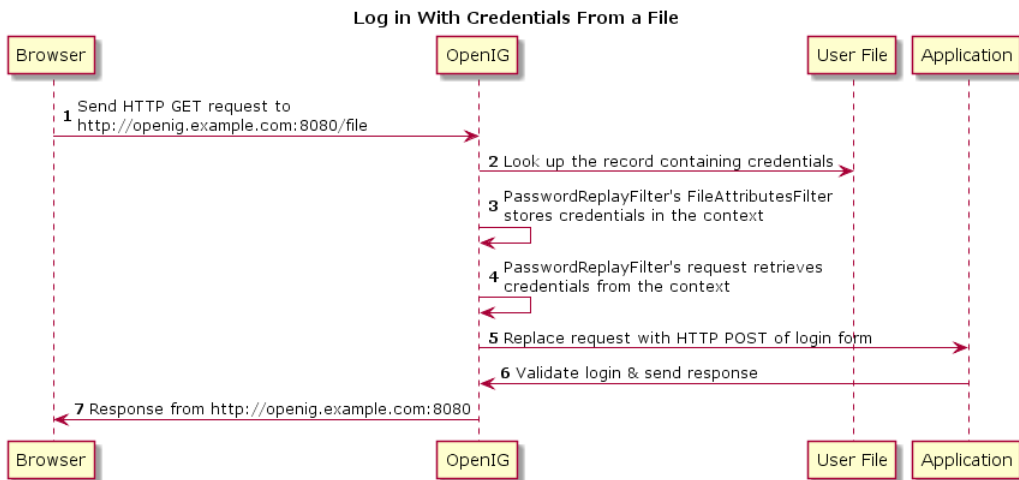
3. On Windows systems, edit the path name to the user file.

Now browse to `http://openig.example.com:8080/file`.

If everything is configured correctly, OpenIG logs you in as George.

What's happening behind the scenes?

Figure 4.1. Log in With Credentials From a File



OpenIG intercepts your browser's HTTP GET request. The request matches the new route configuration. The `PasswordReplayFilter`'s `FileAttributesFilter` looks up credentials in a file, and stores the credentials it finds in the request context attributes map. The `PasswordReplayFilter`'s request pulls the credentials out of the attributes map, builds the login form, and performs the HTTP POST request

to the HTTP server. The HTTP server validates the credentials, and responds with a profile page. OpenIG then passes the response from the HTTP server to your browser.

4.3. Log in With Credentials From a Database

This sample shows you how to configure OpenIG to get credentials from H2. This sample was developed with Jetty and with H2 1.4.178.

Although this sample uses H2, OpenIG also works with other database software. OpenIG relies on the application server where it runs to connect to the database. Configuring OpenIG to retrieve data from a database is therefore a question of configuring the application server to connect to the database, and configuring OpenIG to choose the appropriate data source, and to send the appropriate SQL request to the database. As a result, the OpenIG configuration depends more on the data structure than on any particular database drivers or connection configuration.

Procedure 4.1. Preparing the Database

Follow these steps to prepare the database:

1. On the system where OpenIG runs, download and unpack H2 database.
2. Start H2:

```
$ sh /path/to/h2/bin/h2.sh
```

H2 starts, listening on port 8082, and opens a browser console page.

3. In the browser console page, select Generic H2 (Server) under Saved Settings. This sets the Driver Class, `org.h2.Driver`, the JDBC URL, `jdbc:h2:tcp://localhost/~/test`, the User Name, `sa`.

In the Password field, type `password`.

Then click Connect to access the console.

4. Run a statement to create a users table based on the user file from Section 4.2, "Log in With Credentials From a File".

If you have not created the user file on your system, put the following content in `/tmp/userfile`:

```
username,password,fullname,email  
george,costanza,George Costanza,george@example.com  
kramer,newman,Kramer,kramer@example.com  
bjensen,hifalutin,Babs Jensen,bjensen@example.com  
demo,changeit,Demo User,demo@example.com  
kvaughan,bribery,Kirsten Vaughan,kvaughan@example.com  
scarter,sprain,Sam Carter,scarter@example.com
```

Then create the users table through the H2 console:

```
DROP TABLE IF EXISTS USERS;
CREATE TABLE USERS AS SELECT * FROM CSVREAD('/tmp/userfile');
```

On success, the table should contain the same users as the file. You can check this by running `SELECT * FROM users;` in the H2 console.

Procedure 4.2. Preparing Jetty's Connection to the Database

Follow these steps to enable Jetty to connect to the database:

1. Configure Jetty for JNDI.

For the version of Jetty used in this sample, stop Jetty and add the following lines to `/path/to/jetty/start.ini`:

```
# =====
# Enable JNDI
# -----
OPTIONS=jndi

# =====
# Enable additional webapp environment configurators
# -----
OPTIONS=plus
etc/jetty-plus.xml
```

For more information, see the Jetty documentation on *Configuring JNDI*.

2. Copy the H2 library to the classpath for Jetty:

```
$ cp /path/to/h2/bin/h2-*.jar /path/to/jetty/lib/ext/
```

3. Define a JNDI resource for H2 in `/path/to/jetty/etc/jetty.xml`:

```
<New id="jdbc/forgerock" class="org.eclipse.jetty.plus.jndi.Resource">
  <Arg></Arg>
  <Arg>jdbc/forgerock</Arg>
  <Arg>
    <New class="org.h2.jdbcx.JdbcDataSource">
      <Set name="Url">jdbc:h2:tcp://localhost/~/test</Set>
      <Set name="User">sa</Set>
      <Set name="Password">password</Set>
    </New>
  </Arg>
</New>
```

4. Add a resource reference to the data source in `/path/to/jetty/etc/webdefault.xml`:

```
<resource-ref>
  <res-ref-name>jdbc/forgerock</res-ref-name>
```

```

<res-type>javax.sql.DataSource</res-type>
<res-auth>Container</res-auth>
</resource-ref>
    
```

5. Restart Jetty to take the configuration changes into account.

Procedure 4.3. Preparing the OpenIG Configuration

Add a new route to the OpenIG configuration to look up credentials in the database:

1. To add the route, add the following route configuration file as `$HOME/.openig/config/routes/03-sql.json`:

```

{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "PasswordReplayFilter",
          "config": {
            "loginPage": "${true}",
            "credentials": {
              "type": "SqlAttributesFilter",
              "config": {
                "dataSource": "java:comp/env/jdbc/forgerock",
                "preparedStatement":
                  "SELECT username, password FROM users WHERE email = ?;",
                "parameters": [
                  "george@example.com"
                ],
                "target": "${attributes.sql}"
              }
            }
          },
          "request": {
            "method": "POST",
            "uri": "http://app.example.com:8081/login",
            "form": {
              "username": [
                "${attributes.sql.USERNAME}"
              ],
              "password": [
                "${attributes.sql.PASSWORD}"
              ]
            }
          }
        }
      ],
      "handler": "ClientHandler"
    }
  },
  "condition": "${matches(request.uri.path, '^/sql')}"
}
    
```

On Windows, the file name should be `%appdata%\OpenIG\config\routes\03-sql.json`.

2. Notice the following features of the new route:

- The `SqlAttributesFilter` specifies the data source to access, a prepared statement to look up the user's record, a parameter to pass into the statement, and where to store the search results in the request context attributes map.
- The `PasswordReplayFilter`'s request retrieves the username and password from the attributes map and replaces your browser's original HTTP GET request with an HTTP POST login request that contains the credentials to authenticate.

Notice that the request is for `username`, `password`, and that H2 returns the fields as `USERNAME` and `PASSWORD`. The configuration reflects this difference.

- The route matches requests to `/sql`.

Procedure 4.4. To Try Logging in With Credentials From a Database

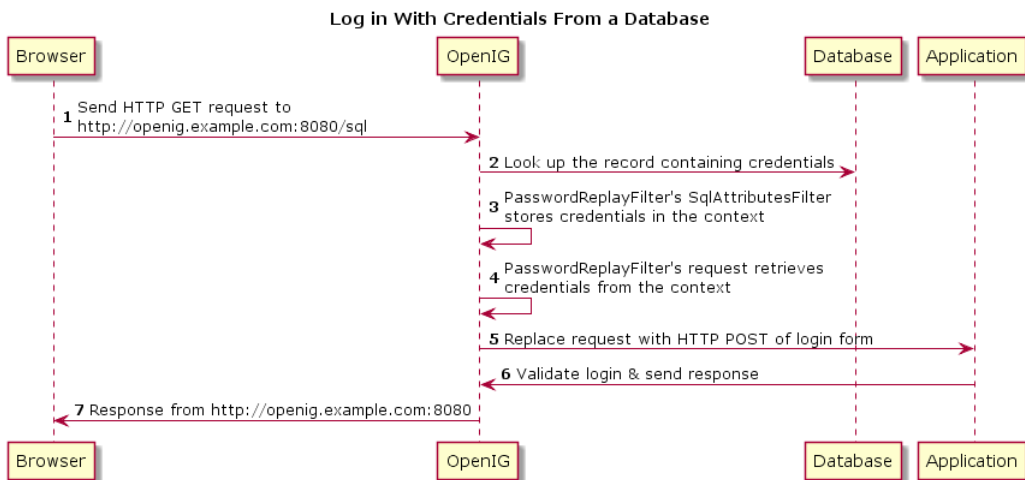
With H2, Jetty, and OpenIG correctly configured, you can try it out:

- Access the new route, `http://openig.example.com:8080/sql`.

OpenIG logs you in automatically as George.

What's happening behind the scenes?

Figure 4.2. Log in With Credentials From a Database



OpenIG intercepts your browser's HTTP GET request. The request matches the new route configuration. The `PasswordReplayFilter`'s `SqlAttributesFilter` looks up credentials in H2, and stores the credentials it finds in the request context attributes map. The `PasswordReplayFilter`'s request pulls the credentials out of the attributes map, builds the login form, and performs the HTTP POST request to the HTTP server. The HTTP server validates the credentials, and responds with a profile page. OpenIG then passes the response from the HTTP server to your browser.

Chapter 5

Getting Login Credentials From Access Management

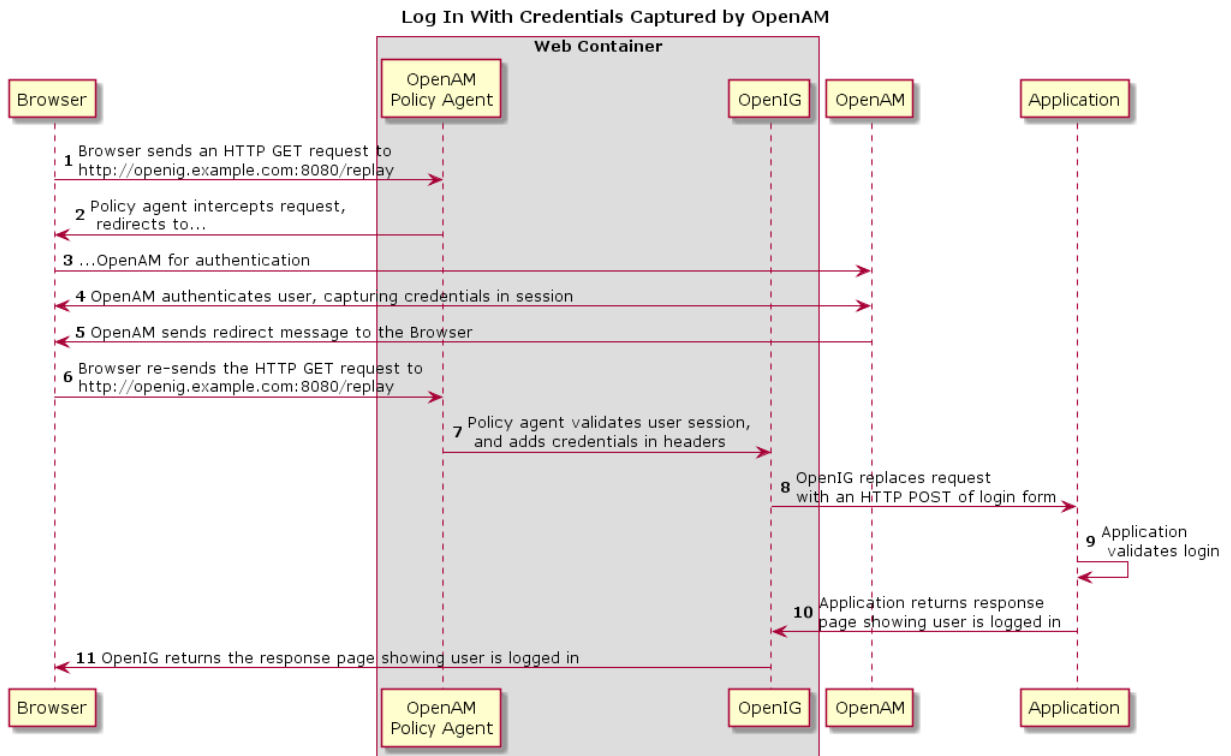
OpenIG helps integrate applications with OpenAM's password capture and replay feature. This feature of OpenAM is typically used to integrate with Microsoft Outlook Web Access (OWA) or SharePoint by capturing the password during OpenAM authentication, encrypting it, and adding to the session, which is later decrypted and used for Basic Authentication to OWA or SharePoint. In this chapter, you will learn:

- How OpenAM password capture and replay works
- To configure OpenIG to obtain credentials from OpenAM authentication
- To use the credentials to log the user in to a protected application

5.1. Detailed Flow

The following figure illustrates the flow of requests for a user who is not yet logged into OpenAM accessing a protected application. After successful authentication, the user is logged into the application with the username and password from the OpenAM login session.

Figure 5.1. Log in With Credentials Captured by OpenAM



1. The user sends a browser request to access a protected application.
2. The OpenAM policy agent protecting OpenIG intercepts the request.
3. The policy agent redirects the browser to OpenAM.
4. OpenAM authenticates the user, capturing the login credentials, storing the password in encrypted form in the user's session.
5. After authentication, OpenAM redirects the browser...
6. ...back to the protected application.
7. The OpenAM policy agent protecting OpenIG intercepts the request, validates the user session with OpenAM (not shown here), adds the username and encrypted password to headers in the request, and passes the request to OpenIG.

8. OpenIG retrieves the credentials from the headers, and uses the username and decrypted password to replace the request with an HTTP POST of the login form.
9. The application validates the login credentials.
10. The application sends a response back to OpenIG.
11. OpenIG passes the response from the application back to the user's browser.

5.2. Setup Summary

Table 5.1. Tasks for Configuring Policy Enforcement

Task	See Section(s)
Create a file containing the password that is to be captured by the policy agent and shared with OpenIG.	Procedure 5.1, "To Create a Password File"
Create a DES shared key to decrypt the password shared by OpenAM and OpenIG.	Procedure 5.2, "To Create a DES Shared Key In OpenIG"
Set up a sample user in OpenAM.	Procedure 5.3, "To Set Up a Sample User in OpenAM"
Create a profile for a J2EE policy agent in OpenAM.	Procedure 5.4, "To Create a Policy Agent Profile in OpenAM"
Configure password capture and add the DES key to the OpenAM configuration.	Procedure 5.5, "To Configure Password Capture in OpenAM"
Install the OpenAM policy agent alongside OpenIG.	Procedure 5.6, "To Install the Policy Agent"
Create a route in OpenIG to handle the requests.	Procedure 5.7, "To Configure OpenIG for Password Replay"
Try the setup.	Procedure 5.8, "To Test the Setup"

5.3. Preparing the Tutorial

Before you start:

- Prepare OpenIG and the sample application as described in Chapter 2, "Getting Started".
- Install and configure OpenAM on <http://openam.example.com:8088/openam>, using the default configuration. If you use a different configuration, substitute in the tutorial accordingly.

Procedure 5.1. To Create a Password File

Create a text file containing the password specified in Procedure 5.4, "To Create a Policy Agent Profile in OpenAM".

1. Create the password file:

```
$ echo password > /tmp/pwd.txt
```

On Windows:

```
C:\> echo password > pwd.txt
```

2. Protect the password file:

```
$ chmod 400 /tmp/pwd.txt
```

In Windows Explorer, right-click the created password file, select **Read-Only**, and then select **OK**.

Procedure 5.2. To Create a DES Shared Key In OpenIG

In this procedure you generate a DES shared key in OpenIG that you then use later in the OpenAM and OpenIG configuration.

When you configure password capture and replay, an OpenAM policy agent shares captured passwords with OpenIG. Before communicating passwords to OpenIG, however, OpenAM encrypts them with a DES shared key. OpenIG uses the DES shared key to decrypt the shared passwords.

The shared key is sensitive information. If it is possible for others to inspect the response, make sure you use HTTPS to protect the communication.

For more information, see `DesKeyGenHandler(5)` in the *Configuration Reference*.

1. Add the following route to the OpenIG configuration as `$HOME/.openig/config/routes/04-keygen.json`.

On Windows, add the route as `%appdata%\OpenIG\config\routes\04-keygen.json`:

```
{
  "handler": {
    "type": "DesKeyGenHandler"
  },
  "condition": "${matches(request.uri.path, '^/keygen')
    and (matches(contexts.client.remoteAddress, ':1')
    or matches(contexts.client.remoteAddress, '127.0.0.1'))}"
}
```

2. Call the route to generate a key:

```
$ curl http://localhost:8080/keygen
{"key": "1A+BCdEfGhI="}
```

In this example, the key is `1A+BCdEfGhI=`.

3. Make a note of the key for later.

5.4. Setting Up OpenAM for Password Replay

Procedure 5.3. To Set Up a Sample User in OpenAM

If you haven't set up the user George Costanza in a previous tutorial, do so now.

1. In the OpenAM console, select the top level realm and browse to Subjects.
2. Click New and create a user with the following values:
 - ID: `george`
 - Last Name: `costanza`
 - Full Name: `george costanza`
 - Password: `costanza`

Procedure 5.4. To Create a Policy Agent Profile in OpenAM

For more information about configuring policy agents in OpenAM, see [Configuring Java EE Policy Agents](#) in the *OpenAM Java EE Policy Agent User's Guide*.

1. In the top-level realm, select Applications > Agents > J2EE.
2. Add an agent with the following values:
 - Name: `JavaEE`
 - Password: `password`
 - Server URL: `http://openam.example.com:8088/openam`
 - Agent URL: `http://openig.example.com:8080/agentapp`
3. Edit the agent profile to add these settings, making sure to save your changes in each tab:
 - On the Global settings tab, select General, and change the Agent Filter Mode from `ALL` to `SSO_ONLY`.
 - On the Application tab, select Session Attributes Processing, and change the following values:
 - Session Attribute Fetch Mode: Change from `NONE` to `HTTP_HEADER`.
 - Session Attribute Mapping: Add `[UserToken]=username` and `[sunIdentityUserPassword]=password`.

Note

If you plan to use ForgeRock Common REST to create or edit routes, on the Application tab, under Not Enforced URI Processing, add `/openig/api/*`.

This step adds the URL pattern for Common REST requests to the list of not-enforced URIs. The policy agent bypasses requests that match this URL pattern, granting them access without requiring authentication.

For information about managing routes through Common REST, see Section 12.3, "Creating and Editing Routes Through Common REST".

Procedure 5.5. To Configure Password Capture in OpenAM

For more information about configuring policy agents in OpenAM, see *Configuring Java EE Policy Agents* in the *OpenAM Java EE Policy Agent User's Guide*.

1. Update the Authentication Post Processing Classes for password replay.
 - a. In the top level realm, select Authentication > Settings > Post Authentication Processing.
 - b. Add the following class to Authentication Post Processing Classes, and then save the changes:
`com.sun.identity.authentication.spi.ReplayPasswd`.
2. Add the key you created in Procedure 5.2, "To Create a DES Shared Key In OpenIG" to the OpenAM configuration:
 - a. In the OpenAM console, select DEPLOYMENT > Servers, and then select the OpenAM server name.

In some earlier releases, select Configuration > Servers and Sites.
 - b. Select Advanced, and add the property `com.sun.am.replaypasswd.key` with the value of the DES shared key.
3. Restart OpenAM.

5.5. Installing the OpenAM Policy Agent

Procedure 5.6. To Install the Policy Agent

For more information about how to install the OpenAM policy agent, see *To Install the Policy Agent into Jetty* in the *OpenAM Java EE Policy Agent User's Guide*.

1. Make sure that OpenIG is stopped and OpenAM is running.
2. Download and then unzip the OpenAM Java EE policy agent in the directory alongside OpenIG. For example, in the `/path/to` directory.

A `/path/to/j2ee_agents` directory is created.

3. Install the agent with a command similar to this:

```
$ /path/to/j2ee_agents/jetty_v7_agent/bin/agentadmin --install --acceptLicense
```

You are prompted to enter information about the installation.

4. Answer the prompts with the following hints:

- Jetty Server Config Directory: `/path/to/jetty/etc`
- Jetty installation directory: `/path/to/jetty`
- OpenAM server URL: `http://openam.example.com:8088/openam`
- Agent URL: `http://openig.example.com:8080/agentapp`
- Agent Profile Name: `JavaEE`
- Agent Profile Password file name : `/tmp/pwd.txt`

5. Add the following filter configuration to `/path/to/jetty/etc/webdefault.xml`:

```
<filter>
  <filter-name>Agent</filter-name>
  <display-name>Agent</display-name>
  <description>OpenAM Policy Agent Filter</description>
  <filter-class>com.sun.identity.agents.filter.AmAgentFilter</filter-class>
</filter>

<filter-mapping>
  <filter-name>Agent</filter-name>
  <url-pattern>/replay</url-pattern>
  <dispatcher>REQUEST</dispatcher>
  <dispatcher>INCLUDE</dispatcher>
  <dispatcher>FORWARD</dispatcher>
  <dispatcher>ERROR</dispatcher>
</filter-mapping>
```

6. To test the setup of the policy agent, start OpenIG, and browse to `http://openig.example.com:8080/replay`.

You should be redirected to OpenAM for authentication, but do not log in yet.

5.6. Setting Up OpenIG for Password Replay

Procedure 5.7. To Configure OpenIG for Password Replay

1. Add the following route to OpenIG configuration as `$HOME/.openig/config/routes/04-replay.json`:

On Windows, add the route as `%appdata%\OpenIG\config\routes\04-replay.json`.

```

{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "PasswordReplayFilter",
          "config": {
            "loginPage": "${true}",
            "headerDecryption": {
              "algorithm": "DES/ECB/NoPadding",
              "key": "DESKEY",
              "keyType": "DES",
              "charSet": "utf-8",
              "headers": [
                "password"
              ]
            },
            "request": {
              "method": "POST",
              "uri": "http://app.example.com:8081/login",
              "form": {
                "username": [
                  "${request.headers['username']}[0]"
                ],
                "password": [
                  "${request.headers['password']}[0]"
                ]
              }
            }
          }
        },
        {
          "type": "HeaderFilter",
          "config": {
            "messageType": "REQUEST",
            "remove": [
              "password",
              "username"
            ]
          }
        }
      ],
      "handler": "ClientHandler"
    }
  },
  "condition": "${matches(request.uri.path, '^/replay')}"
}

```

2. Edit the route to change **DESKEY** to the actual key value that you generated in Procedure 5.2, "To Create a DES Shared Key In OpenIG".

Notice the following features of the new route:

- The route matches requests to `/replay`.

- The `PasswordReplayFilter` uses the `headerDecryption` information to decrypt the password that OpenAM captured and encrypted, and that the OpenAM policy agent included in the headers for the request.

The `headerDecryption` object uses the DES shared key value that you generated.

- The `PasswordReplayFilter` retrieves the username and password from the context and replaces your browser's original HTTP GET request with an HTTP POST login request that contains the credentials to authenticate.
- The `HeaderFilter` removes the username and password headers before continuing to process the request.

5.7. Testing the Setup

Procedure 5.8. To Test the Setup

1. Log out of OpenAM if you are logged in.
2. Access the route on `http://openig.example.com:8080/replay`.

You should be redirected to the OpenAM login page.

3. Log in with username `george`, password `costanza`.

You should be redirected to the sample application.

Chapter 6

Enforcing Policy Decisions and Supporting Session Upgrade

This chapter describes how to set up OpenIG as a policy enforcement point, with OpenAM as a policy decision point. It provides an example of how to enforce a policy decision from OpenAM, and an example of upgrading a session to a higher authentication level.

For more information about authentication and session upgrade, see OpenAM's *Authentication and Single Sign-On Guide*.

6.1. About OpenIG As a PEP With OpenAM As PDP

The following terms are used in access management:

- *Policy Decision Point (PDP)*: Entity that evaluates access rights and then issues authorization decisions.
- *Policy Enforcement Point (PEP)*: Entity that intercepts a request for a resource and then enforces policy decisions from a PDP.

OpenIG as a PEP intercepts requests for a resource, and provides information about the request to OpenAM as a PDP.

OpenAM evaluates requests based on their context and the configured policies. OpenAM then returns decisions that indicate what actions are allowed or denied, as well as any advices, subject attributes, or static attributes for the specified resources.

After a policy decision, OpenIG continues to process requests as follows:

- If the request is allowed, processing continues.
- If the request is denied with advice, OpenIG checks whether it can respond to the advice. If OpenIG can respond to the advice, it processes the advice.
- If the request is denied without advice, or if OpenIG cannot respond to the advice, OpenIG forwards the request to a `failureHandler` declared in the `PolicyEnforcementFilter`. If there is no `failureHandler`, OpenIG returns a 403 Forbidden.
- If an error occurs during the process, OpenIG returns 500 Internal Server Error.

Attributes and advices returned by the policy decision are stored in the `attributes` and `advices` properties of the `#{contexts.policyDecision}` context. They are available to the `PolicyEnforcementFilter`'s downstream filters and handlers.

In OpenAM, administrators can maintain centralized, fine-grained, declarative policies to manage who can access what resources, and under what conditions. Policies can be managed separately by OpenAM realm and by OpenAM application.

OpenAM provides a REST API for authorized users to request policy decisions. OpenIG provides a `PolicyEnforcementFilter` that uses the REST API. For information, see `PolicyEnforcementFilter(5)` in the *Configuration Reference*.

6.2. Enforcing Policy Decisions From OpenAM

This section gives an example of how to set up OpenIG as a PEP, requesting policy decisions from OpenAM as a PDP.

Table 6.1. Tasks for Configuring Policy Enforcement

Task	See Section(s)
Set up a policy in OpenAM to allow authenticated users to access the sample application.	Procedure 6.1, "To Create a Policy in OpenAM"
Create credentials and privileges in OpenAM for a policy administrator. When OpenIG requests policy decisions, it must use these credentials.	Procedure 6.2, "To Create a Policy Administrator in OpenAM"
Configure OpenIG as a PEP in OpenIG Studio.	Procedure 6.3, "To Set Up OpenIG as a PEP"
Test the setup.	Procedure 6.4, "To Test the Setup"

6.2.1. Setting Up OpenAM As a PDP

This section describes how to create a policy in OpenAM and configure a user who can request policy decisions.

Before you start, install and configure OpenAM on <http://openam.example.com:8088/openam>, with the default configuration. If you use a different configuration, substitute in the tutorial accordingly.

Procedure 6.1. To Create a Policy in OpenAM

1. Log in to the OpenAM console as administrator.
2. In the top-level realm, select Authorization > Policy Sets.
3. Add a new policy set with the following values, and then select Create:

- Id: `PEP policy set`
 - Resource Types: `URL`
4. In the new policy set, add a policy with the following values, and then select Create:
 - Name: `OpenIG Policy`
 - Resource Type: `URL`
 - Resource pattern: `*://*:*/*`
 - Resource value: `http://app.example.com:8081/home/pep`

This policy protects the home page of the sample application.
 5. Select the Actions tab, add an action to allow HTTP `GET`, and then save your changes.
 6. Select the Subjects tab, remove any default subject conditions, add a subject condition for all `Authenticated Users`, and then save your changes.

Procedure 6.2. To Create a Policy Administrator in OpenAM

1. In the top-level realm, create a subject with ID `policyAdmin` and password `password`.
2. Create a `policyAdmins` group and add the user you created.
3. In the privileges configuration for the `policyAdmins` group, add the privilege `REST calls for policy evaluation`.

This privilege allows a subject in that group to request policy decisions.

4. Make sure all the changes are saved.

6.2.2. Setting Up OpenIG as a PEP

This section describes how to set up OpenIG to configure policy enforcement, where the user agent is redirected to OpenAM for authentication.

To configure OpenIG without using OpenIG Studio, add the route in Example 6.1, "Route for OpenIG As a PEP" to the OpenIG configuration as `$HOME/.openig/config/routes/04-pep.json`. On Windows, add the route as `%appdata%\OpenIG\config\routes\04-pep.json`.

Procedure 6.3. To Set Up OpenIG as a PEP

Before you start, prepare OpenIG and the sample application as described in Chapter 2, "Getting Started". Make sure that your `config.json` has a main router named `_router`.

1. Access OpenIG Studio on `http://openig.example.com:8080/openig/studio`, and select Protect an Application.
2. In the Create a route window, select Advanced Options and enter the following information, and then select Create route:
 - Base URI: `http://app.example.com:8081`
 - Condition: Path: `/home/pep`
 - Name: `04-pep`
3. Select Authentication, and then enable authentication.
4. Select Single Sign-On, enter the following information, and then select Save:
 - OpenAM URL: `http://openam.example.com:8088/openam`

Leave the other fields with their default values.

5. Select Authorization, and then enable authorization.
6. Select the following options to reflect the configuration in Section 6.2.1, "Setting Up OpenAM As a PDP", and then save the settings:
 - OpenAM configuration:
 - OpenAM URL: `http://openam.example.com:8088/openam`
 - Policy administrator ID: `policyAdmin`
 - Policy administrator password: `password`
 - OpenAM policy endpoint:
 - Policy set: `PEP policy set`
 - OpenAM SSO token ID: `${contexts.ssoToken.value}`

Leave the other fields with their default values, or empty if they are empty.

7. On the top-right of the screen, select `# > Display`, and review the route. The route in Example 6.1, "Route for OpenIG As a PEP" should be displayed.

Notice the following features of the new route:

- The route matches requests to `/home/pep`.
- The first filter in the chain is the `SingleSignOnFilter`. For information, see `SingleSignOnFilter(5)` in the *Configuration Reference*.

If the request does not have a valid `iPlanetDirectoryPro` cookie, the `SingleSignOnFilter` redirects the request to OpenAM for authentication.

If the request already has a valid `iPlanetDirectoryPro` cookie, or after authenticating with OpenAM to get a valid `iPlanetDirectoryPro` cookie, the `SingleSignOnFilter` passes the request to the next filter. The `SingleSignOnFilter` stores the cookie value in an `SsoTokenContext`.

- The next filter in the chain is the `PolicyEnforcementFilter`. For information, see `PolicyEnforcementFilter(5)` in the *Configuration Reference*.

The `PolicyEnforcementFilter` retrieves the token value from the `SsoTokenContext` to identify the subject making the request, and then calls the OpenAM policy endpoint for a policy decision.

8. Select Deploy to push the route to the OpenIG configuration.

You can check the `$HOME/.openig/config/routes` folder to see that the route is there.

Example 6.1. Route for OpenIG As a PEP

```
{
  "name": "04-pep",
  "baseURI": "http://app.example.com:8081",
  "monitor": false,
  "condition": "${matches(request.uri.path, '^/home/pep')}",
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "SingleSignOnFilter",
          "name": "SingleSignOn",
          "config": {
            "openamUrl": "http://openam.example.com:8088/openam",
            "realm": "/",
            "cookieName": "iPlanetDirectoryPro"
          }
        },
        {
          "type": "PolicyEnforcementFilter",
          "name": "PEPFilter",
          "config": {
            "openamUrl": "http://openam.example.com:8088/openam",
            "pepUsername": "policyAdmin",
            "pepPassword": "password",
            "pepRealm": "/",
            "realm": "/",
            "application": "PEP policy set",
            "ssoTokenSubject": "${contexts.ssoToken.value}"
          }
        }
      ]
    }
  },
  "handler": "ClientHandler"
}
```

```
}  
}
```

6.2.3. Testing the Setup

Procedure 6.4. To Test the Setup

1. Log out of OpenAM.
2. Browse to <http://openig.example.com:8080/home/pep>.

Because you have not previously authenticated to OpenAM, the request does not contain a cookie with an SSO token. The SingleSignOnFilter redirects you to OpenAM for authentication.

3. Log in to OpenAM as user `demo`, password `changeit`.

When you have authenticated, OpenAM redirects you back to the request URL, and OpenIG requests a policy decision using the iPlanetDirectoryPro cookie value.

OpenAM returns a policy decision that grants access to the sample application.

6.3. Upgrading a Session

This section builds on the example in Section 6.2, "Enforcing Policy Decisions From OpenAM". The authentication level required to access the sample application is increased to level 1. After authenticating with OpenAM, the user must upgrade the session by entering a verification code before they can access the sample application.

Table 6.2. Tasks for Configuring Session Upgrade

Task	See Section(s)
Set up OpenIG as a policy enforcement point and OpenAM as a policy decision point, as described in the previous section.	Section 6.2, "Enforcing Policy Decisions From OpenAM"
Add an environment condition to the OpenIG policy to require authentication level 1 to access the sample application.	Procedure 6.5, "To Add an Environment Condition to the OpenAM Policy"
Create a scripted authentication module and scripts in OpenAM to deliver a token with authentication level 1.	Procedure 6.6, "To Set Up an Authentication Module"
Try the setup.	Procedure 6.7, "To Test the Setup"

6.3.1. Setting Up OpenAM for Session Upgrade

This section provides an example of session upgrade. After authenticating with OpenAM, the user must also provide a verification code to access the sample application.

To run this example, configure OpenIG and OpenAM as described in Section 6.2, "Enforcing Policy Decisions From OpenAM". No additional configuration of OpenIG is required; the following procedures build on the OpenAM configuration.

Procedure 6.5. To Add an Environment Condition to the OpenAM Policy

This section changes the policy you created in Procedure 6.1, "To Create a Policy in OpenAM".

1. Log in to the OpenAM console as administrator.
2. In the top-level realm, select Authorization > Policy Sets > PEP policy set.
3. In the OpenIG policy, select Environments, and add an environment condition to require authentication level greater than or equal to 1.

Procedure 6.6. To Set Up an Authentication Module

This section sets up an OpenAM authentication module and scripts to upgrade a session's authentication level to 1.

1. In the top level realm, select Scripts > Scripted Module - Client Side, and replace the default script with the following script:

```
spinner.hideSpinner();
autoSubmitDelay = 60000;
$(document).ready(function() {
  fs = $(document.forms[0]).find("fieldset");
  strUI = '<div class="form-group"> \
    <label class="sr-only separator" for="answer"> \
      Verification Code</label><input onchange="s=$(\'#clientScriptOutputData\')[0]; \
      if (!s.value) s.value=\'{}\'; d=JSON.parse(s.value); d[\'answer\']=value; \
      s.value=JSON.stringify(d);" id="answer" class="form-control input-lg" type="text" \
      placeholder="Enter your verification code" value="" name="answer"></input></div>';
  $(fs).prepend(strUI);
});
```

Leave all other values as default.

This client side script adds a field to the OpenAM form, in which the user is required to enter a verification code. The script formats the entered code as a JSON object, as required by the server side script.

2. Select Scripts > Scripted Module - Server Side, and replace the default script with the following script:

```
username = 'demo'
logger.error('username: ' + username)

// Test whether the user 'demo' enters the correct validation code
data = JSON.parse(clientScriptOutputData);
answer = data.answer;

if (answer !== '123456') {
  logger.error('Authentication Failed !!!')
  authState = FAILED;
} else {
  logger.error('Authenticated !!!')
  authState = SUCCESS;
}
```

Leave all other values as default.

This server side script tests that the user `demo` has entered `123456` as the verification code.

3. In the top level realm, select Authentication > Modules, and add a module with the following characteristics:
 - Name: `VerificationCodeLevel1`
 - Type: `Scripted Module`
4. In the authentication module, enable the option for client-side script, and select the following options:
 - Client-side Script: `Scripted Module - Client Side`
 - Server-side Script: `Scripted Module - Server Side`
 - Authentication Level: `1`

6.3.2. Testing the Setup

Procedure 6.7. To Test the Setup

1. Log out of OpenAM.
2. Browse to `http://openig.example.com:8080/home/pep`.

If you have not previously authenticated to OpenAM, the `SingleSignOnFilter` redirects the request to OpenAM for authentication.

3. Log in to OpenAM as user `demo`, password `changeit`.

OpenAM creates a session with an authentication level 0, and OpenIG requests a policy decision. The updated policy requires authentication level 1, which is higher than the session's current authentication level.

OpenAM steps up the authentication level by using the authentication module and scripts you just created.

4. In the session upgrade window, enter the verification code **123456**.

OpenAM returns a policy decision that grants access to the sample application. You can now log in to the sample application.

Chapter 7

Acting As a SAML 2.0 Service Provider

The federation component of OpenIG is a standards-based authentication service used by OpenIG to validate users and log them in to applications that OpenIG protects. The federation component implements SAML 2.0. In this chapter, you will learn:

- How OpenIG works as a SAML 2.0 service provider.
- How to configure OpenIG as a service provider for a single application.

Appendix A, "*SAML 2.0 and Multiple Applications*" describes how to set up OpenIG as a SAML 2.0 service provider for two applications, using OpenAM as the identity provider.

7.1. About SAML 2.0 SSO and Federation

Federation allows organizations to share identities and services without giving away their identity information or the services they provide. Federation depends on standards to orchestrate interaction and exchange information between providers.

SAML 2.0 is a standard that describes the messages that providers exchange and the way that they exchanged them. SAML 2.0 enables web single sign-on (SSO), for example, where the service managing the user's identity does not belong to the same organization and does not use the same software as the service that the user wants to access.

The following terms are used for federation and SAML:

- *Identity Provider* (IDP): The service that manages the user identity.
- *Service Provider* (SP): The service that a user wants to access.
- *Circle of trust* (CoT): An identity provider and service provider that participate in the federation.

When an identity provider and a service provider participate in federation, they agree on what security information to exchange, and mutually configure access to each others' services. The metadata that identity providers and service providers share is in an XML format defined by the SAML 2.0 standard.

7.1.1. Steps in SSO

SSO can be initiated by the SP (*SP initiated SSO*) or by the IDP (*IDP initiated SSO*).

Before SSO can be initiated by an IDP, the IDP must be configured with links that refer to the SP, and the user must be authenticated to the IDP. Instead of accessing `app.example.com` directly on the SP, the user accesses a link on the IDP that refers to the remote SP. The IDP provides SAML assertions for the user to the SP.

When SSO is initiated by the SP, the user attempts to access `app.example.com` directly on the SP. Because the user's federated identity is managed by the IDP, the SP sends a SAML authentication request to the IDP. After authenticating the user, the IDP provides SAML assertions for the user to the SP.

In both cases, the SAML assertion sent from the IDP to the SP attests which user is authenticated, when the authentication succeeded, how long the assertion is valid, and so on. The assertions can optionally contain additional and configurable attribute values, such as user meta-information or anything else provided by the IDP.

The SP uses the SAML assertions from the IDP to make authorization decisions, for example, to let an authenticated user complete a purchase that gets charged to the user's account at the identity provider.

SAML assertions can optionally be signed and encrypted.

7.1.2. OpenIG As a SAML 2.0 SP

OpenIG acts as a SAML 2.0 SP for SSO, providing users with an interface to applications that don't support SAML 2.0.

When SSO is initiated by the IDP, the IDP sends an unsolicited authentication statement to OpenIG. When SSO is initiated by OpenIG, OpenIG calls the federation component to initiate SSO with the IDP. In both cases, the job of the federation component is to authenticate the user and to pass the required attributes to OpenIG so that OpenIG can log the user into protected applications.

7.2. Installation Overview

This tutorial assumes that you are familiar with SAML 2.0 federation and with the components involved, including OpenAM. For information about OpenAM, read the documentation for your version of OpenAM.

This tutorial does not address PKI configuration for validation and encryption, although OpenIG is capable of handling both, just as any OpenAM Fedlet can handle both.

This tutorial takes you through the steps to set up OpenAM as an IDP and OpenIG as an SP to protect an application. To set up multiple SPs, read this chapter, work through the samples, and then consider the explanation in [Appendix A, "SAML 2.0 and Multiple Applications"](#).

Table 7.1. Tasks for Configuring SAML 2.0 SSO and Federation

Task	See Section(s)
Prepare the network.	Section 7.3, "Preparing the Network"
Configure OpenAM As an IDP.	Section 7.4.1, " Setting Up OpenAM for This Tutorial " Section 7.4.2, " Setting Up a Hosted Identity Provider " Section 7.4.3, " Setting up a Fedlet "
Configuring OpenIG as a SP.	Section 7.5.1, " Retrieve the Fedlet Configuration Files " Section 7.5.2, " Adding a Route for Credential Injection " Section 7.5.3, " Adding a Route for SAML Federation "

Table 7.2. Fedlet Configuration Files

File	Description
<code>fedlet.cot</code>	Circle of trust for OpenIG and the identity provider.
<code>idp.xml</code>	Standard metadata (usually generated by the IDP).
<code>idp-extended.xml</code>	Metadata extensions (usually generated by the IDP).
<code>sp.xml</code>	Standard metadata for the OpenIG SP (usually generated by the IDP).
<code>sp-extended.xml</code>	Metadata extensions for the OpenIG SP (usually generated by the IDP).

For examples of the federation configuration files, see Section 7.7, " Example Federation Configuration Files ". You can copy and edit these files to create new configurations.

7.3. Preparing the Network

Configure the network so that browser traffic to the application hosts is proxied through OpenIG. The example in this chapter uses the host name `sp.example.com`.

Add the following address to your hosts file: `sp.example.com`.

```
127.0.0.1    localhost openam.example.com openig.example.com app.example.com sp.example.com
```

7.4. Configuring OpenAM As an IDP

7.4.1. Setting Up OpenAM for This Tutorial

Procedure 7.1. To Set Up OpenAM for This Tutorial

1. Install and configure OpenAM on <http://openam.example.com:8088/openam>, with the default configuration. If you use a different configuration, make sure you substitute in the tutorial accordingly.
2. In the top level realm, browse to Subject and create a user with following credentials:
 - ID (Username): `george`
 - Last Name: `costanza`
 - Full Name: `george costanza`
 - Password: `costanza`
3. Edit the user to add the following information:
 - Email Address: `george`
 - Employee Number: `costanza`

For simplicity, this tutorial uses `mail` to hold the username, and `employeenumber` to hold the password. Both attributes are in the standard user profile with the default OpenAM configuration, and neither is needed for anything else in this tutorial. In a real deployment, you would use other attributes to represent real user profiles.

4. Test that you can log in to OpenAM with this username and password.

7.4.2. Setting Up a Hosted Identity Provider

Procedure 7.2. To Set Up a Hosted Identity Provider

1. Select the top level realm and browse to Configure SAMLv2 Provider > Create Hosted Identity Provider.
A configuration page for the IDP is displayed.
2. In metadata > Name, change <http://openam.example.com:8088/openam> to `openam`.
This makes it easier to refer to OpenAM as the IDP later.
3. In metadata > Signing Key, select `test`.

4. In Circle of Trust, select an existing circle of trust (CoT) or select Add and enter the name of a new CoT. In this example, the CoT is called `Circle of Trust`.
5. In Attribute Mapping, map the `mail` attribute to `mail`, and map the `employeenumber` attribute to `employeenumber`.

The SAML 2.0 attribute mapping indicates that OpenIG (the SP) wants OpenAM (the IDP) to get the value of these attributes from the user profile and send them to OpenIG. OpenIG can use the attribute values to log the user in to the application it protects.

6. Select Configure.

A confirmation page is displayed. You can start to create a Fedlet from this page or go back to the top level realm, as described in the following procedure.

7.4.3. Setting up a Fedlet

A Fedlet is an example web application that acts as a lightweight SAML v2.0 SP. When you create a Fedlet, the federation configuration files are created in a directory similar to this: `$HOME/openam/myfedlets/openig-fedlet/Fedlet.zip`.

Procedure 7.3. To Set Up a Fedlet

1. In the top level realm, browse to Create Fedlet.
2. In Name, enter a name for the Fedlet. In this tutorial, the Fedlet is named `sp`.
3. In Destination URL, enter the following URL for the SP: `http://sp.example.com:8080/saml`.
4. In Attribute Mapping, map the `mail` attribute to `mail`, and map the `employeenumber` attribute to `employeenumber`.
5. Select Create.

After successfully creating the Fedlet, OpenAM displays the location of the configuration files. Depending on your version of OpenAM, the configuration files are in a `war` directory or `.zip` file.

The `.zip` file is named something like `$HOME/openam/myfedlets/sp/Fedlet.zip` on the system where OpenAM runs.

7.5. Configuring OpenIG As an SP

Before you start, prepare OpenIG and the sample application as shown in Chapter 2, "Getting Started". Getting the basic setup to work before you configure federation makes it simpler to troubleshoot if anything goes wrong.

To test your setup, access the login page of the sample application through OpenIG at <http://openig.example.com:8080/login>. Log in as username `george`, password `costanza`. You should see a page showing the username and some information about the request.

7.5.1. Retrieve the Fedlet Configuration Files

Procedure 7.4. To Retrieve the Fedlet Configuration Files

1. Unpack the configuration files from the Fedlet you created in Section 7.4.3, "Setting up a Fedlet". For example, unpack the .zip file as follows:

```
$ cd $HOME/openam/myfedlets/sp
$ unzip Fedlet.zip
$ mkdir $HOME/.openig/SAML
$ cp conf/* $HOME/.openig/SAML
$ ls -l $HOME/.openig/SAML
```

```
FederationConfig.properties
fedlet.cot
idp-extended.xml
idp.xml
sp-extended.xml
sp.xml
```

2. If the following fedlet adapter header is defined in `sp-extended.xml`, comment it out to prevent issues with timeout:

```
<!--
  <Attribute name="fedLetAdapter">
    <Value>com.sun.identity.saml2.plugins.DefaultFedletAdapter</Value>
  </Attribute>
-->
```

3. Restart OpenIG.

7.5.2. Adding a Route for Credential Injection

Create the configuration file `$HOME/.openig/config/routes/05-saml.json`.

On Windows, the file name should be `%appdata%\OpenIG\config\routes\05-saml.json`.

```
{
  "handler": {
    "type": "SamlFederationHandler",
    "config": {
      "assertionMapping": {
        "username": "mail",
        "password": "employeenumber"
      }
    }
  },
```

```

    "subjectMapping": "sp-subject-name",
    "redirectURI": "/federate"
  },
  "condition": "${matches(request.uri.path, '^/saml')}",
  "session": "JwtSession"
}

```

The route injects credentials into the context, based on attribute values from the SAML assertion returned on successful authentication. Note the following features of the route:

- The route matches requests to `/saml`.
- The `SamLFederationHandler` extracts the mail and employee number from the SAML assertion and maps them to the session fields `session.username` and `session.password`.

The handler stores the subject name as a string in the session field `session.sp-subject-name`, which is named by the `subjectMapping` property. By default, the subject name is stored in the session field `session.subjectName`.

The handler then redirects the request to the `/federate` route.

- The route uses the `JwtSession` implementation, meaning it stores encrypted session information in a browser cookie. The name is a reference to the `JwtSession` object defined in `config.json`. For details, see `JwtSession(5)` in the *Configuration Reference*.

7.5.3. Adding a Route for SAML Federation

Create the configuration file `$HOME/.openig/config/routes/05-federate.json`.

On Windows, the file name should be `%appdata%\OpenIG\config\routes\05-federate.json`.

```

{
  "handler": {
    "type": "DispatchHandler",
    "config": {
      "bindings": [
        {
          "condition": "${empty session.username}",
          "handler": {
            "type": "StaticResponseHandler",
            "config": {
              "status": 302,
              "reason": "Found",
              "headers": {
                "Location": [
                  "http://sp.example.com:8080/saml/SPInitiatedSSO"
                ]
              }
            }
          }
        ]
      }
    }
  }
}

```

```

    }
  },
  {
    "handler": {
      "type": "Chain",
      "config": {
        "filters": [
          {
            "type": "StaticRequestFilter",
            "config": {
              "method": "POST",
              "uri": "http://app.example.com:8081/login",
              "form": {
                "username": [
                  "${session.username[0]}"
                ],
                "password": [
                  "${session.password[0]}"
                ]
              }
            }
          }
        ],
        "handler": "ClientHandler"
      }
    }
  ]
},
"condition": "${matches(request.uri.path, '^/federate')}",
"session": "JwtSession"
}

```

Notice the following features of the route:

- The route matches requests to `/federate`. This is the route you use to test the configuration.
- If the username has not been populated in the context, the user has not yet authenticated with the IDP. In this case,
 - The `DispatchHandler` dispatches requests to the `StaticResponseHandler`.
 - The `StaticResponseHandler` redirects to the SP-initiated SSO end point to initiate SAML 2.0 web browser SSO.
 - After authentication is successful, the `SamlFederationHandler` injects credentials into the session.

If the credentials have been inserted into the context, or after successful authentication in the previous step, the `DispatchHandler` dispatches requests to the `Chain` to log the user in to the protected application.

- The `StaticRequestFilter` retrieves the first value for the username and password attributes of the SAML assertion.

Note

The attributes of a SAML assertion can contain one or more values, which are made available as a list of strings. Even if an attribute contains a single value, it is made available as a list of strings.

The `StaticRequestFilter` then replaces the browser's original HTTP GET request with an HTTP POST login request that contains the credentials to authenticate.

- The route uses the `JwtSession` implementation, meaning it stores encrypted session information in a browser cookie. The name is a reference to the `JwtSession` object defined in `config.json`. For details, see `JwtSession(5)` in the *Configuration Reference*.

Tip

If more dynamic control is needed for the URL where the user agent is redirected, then use the `RelayState` query string parameter in the URL of the redirect `Location` header. The `RelayState` query string parameter specifies where to redirect the user when the SAML 2.0 web browser SSO process is complete. The `RelayState` overrides the `redirectURI` set in the `SamlFederationHandler`. The `RelayState` value must be URL-encoded. When using an expression, use the `urlEncode()` function to encode the value. For example: `#{urlEncodeQueryParameterNameOrValue(contexts.router.originalUri)}`. In the following example, the user is finally redirected to the original URI from the request:

```
"headers": {
  "Location": [
    "http://openig.example.com:8080/saml/SPInitiatedSSO?RelayState=
    #{urlEncodeQueryParameterNameOrValue(contexts.router.originalUri)}"
  ]
}
```

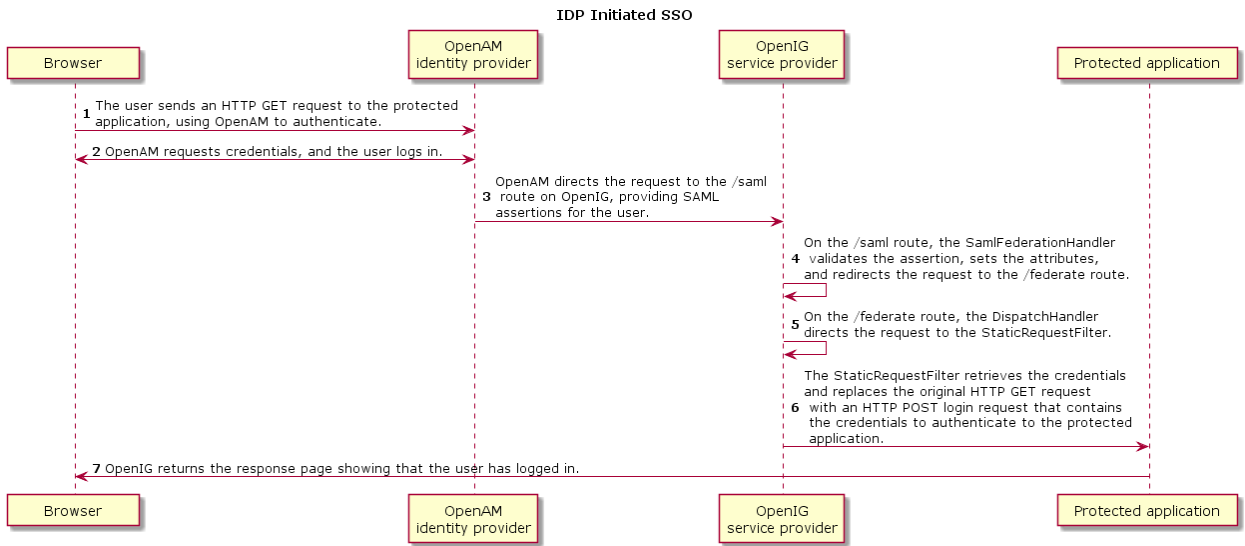
7.6. Testing the Configuration

7.6.1. Testing IDP-initiated SSO

- Log out of the OpenAM console and select this link for IDP-initiated SSO. The OpenAM login page is displayed.
- Log in to OpenAM with username `george` and password `costanza`. OpenIG returns the response page showing that the user has logged in.

The following sequence diagram shows what just happened.

Figure 7.1. IDP-Initiated SSO

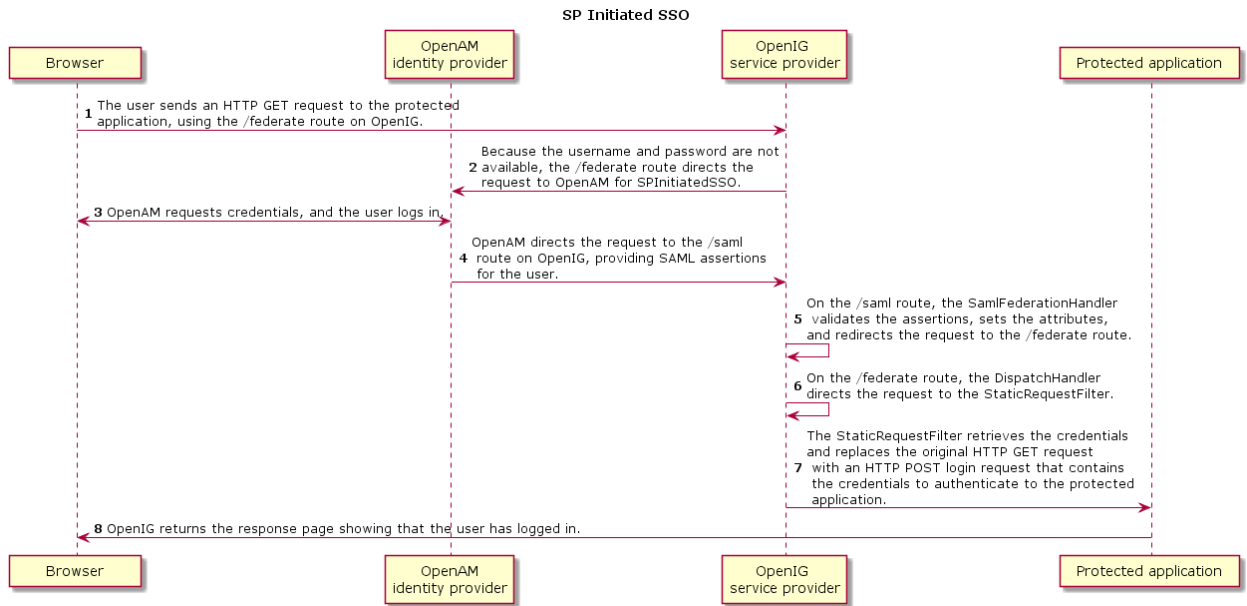


7.6.2. Testing SP-initiated SSO

- Log out of the OpenAM console, and browse to the URL for the route at `http://openig.example.com:8080/federate`. The OpenAM login page is displayed.
- Log in to OpenAM with the username `george` and password `costanza`. OpenIG returns the response page showing that the user has logged in.

The following sequence diagram shows what just happened.

Figure 7.2. SP-Initiated SSO



7.7. Example Federation Configuration Files

7.7.1. Circle of Trust

The following example of `$HOME/.openig/SAML/fedLet.cot` defines a CoT between OpenAM as the IDP and an OpenIG SP:

```

cot-name=Circle of Trust
sun-fm-cot-status=Active
sun-fm-trusted-providers=openam,sp
sun-fm-saml2-readerservice-url=
sun-fm-saml2-writerservice-url=

```

7.7.2. SAML Configuration File

The following example of `$HOME/.openig/SAML/sp.xml` defines a SAML configuration file for an OpenIG service provider, `sp`:

```

<EntityDescriptor

```

```

entityID="sp"
xmlns="urn:oasis:names:tc:SAML:2.0:metadata">
<SPSSODescriptor
  AuthnRequestsSigned="false"
  WantAssertionsSigned="false"
  protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
  <SingleLogoutService
    Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect"
    Location="http://sp.example.com:8080/saml/fedletSloRedirect"
    ResponseLocation="http://sp.example.com:8080/saml/fedletSloRedirect"/>
  <SingleLogoutService
    Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
    Location="http://sp.example.com:8080/saml/fedletSloPOST"
    ResponseLocation="http://sp.example.com:8080/saml/fedletSloPOST"/>
  <SingleLogoutService
    Binding="urn:oasis:names:tc:SAML:2.0:bindings:SOAP"
    Location="http://sp.example.com:8080/saml/fedletSloSoap"/>
  <NameIDFormat>urn:oasis:names:tc:SAML:2.0:nameid-format:transient</NameIDFormat>
  <AssertionConsumerService
    isDefault="true"
    index="0"
    Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
    Location="http://sp.example.com:8080/saml/fedletapplication"/>
  <AssertionConsumerService
    index="1"
    Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifact"
    Location="http://sp.example.com:8080/saml/fedletapplication"/>
</SPSSODescriptor>
<RoleDescriptor
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:query="urn:oasis:names:tc:SAML:metadata:ext:query"
  xsi:type="query:AttributeQueryDescriptorType"
  protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
</RoleDescriptor>
<XACMLAuthzDecisionQueryDescriptor
  WantAssertionsSigned="false"
  protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
</XACMLAuthzDecisionQueryDescriptor>
</EntityDescriptor>
    
```

7.7.3. Extended Configuration File

The following example of `$HOME/.openig/SAML/sp-extended.xml` defines a SAML configuration file for an OpenIG service provider, `sp`:

```

<EntityConfig xmlns="urn:sun:fm:SAML:2.0:entityconfig"
  xmlns:fm="urn:sun:fm:SAML:2.0:entityconfig"
  hosted="1"
  entityID="sp">

  <SPSSOConfig metaAlias="/sp">
    <Attribute name="description">
      <Value></Value>
    </Attribute>
    <Attribute name="signingCertAlias">
      <Value></Value>
    </Attribute>
  </SPSSOConfig>
</EntityConfig>
    
```

```

</Attribute>
<Attribute name="encryptionCertAlias">
  <Value></Value>
</Attribute>
<Attribute name="basicAuthOn">
  <Value>>false</Value>
</Attribute>
<Attribute name="basicAuthUser">
  <Value></Value>
</Attribute>
<Attribute name="basicAuthPassword">
  <Value></Value>
</Attribute>
<Attribute name="autofedEnabled">
  <Value>>false</Value>
</Attribute>
<Attribute name="autofedAttribute">
  <Value></Value>
</Attribute>
<Attribute name="transientUser">
  <Value>anonymous</Value>
</Attribute>
<Attribute name="spAdapter">
  <Value></Value>
</Attribute>
<Attribute name="spAdapterEnv">
  <Value></Value>
</Attribute>
<!--
<Attribute name="fedletAdapter">
  <Value>com.sun.identity.saml2.plugins.DefaultFedletAdapter</Value>
</Attribute>
-->
<Attribute name="fedletAdapterEnv">
  <Value></Value>
</Attribute>
<Attribute name="spAccountMapper">
  <Value>com.sun.identity.saml2.plugins.DefaultLibrarySPAccountMapper</Value>
</Attribute>
<Attribute name="useNameIDAsSPUserID">
  <Value>>false</Value>
</Attribute>
<Attribute name="spAttributeMapper">
  <Value>com.sun.identity.saml2.plugins.DefaultSPAttributeMapper</Value>
</Attribute>
<Attribute name="spAuthncontextMapper">
  <Value>com.sun.identity.saml2.plugins.DefaultSPAuthnContextMapper</Value>
</Attribute>
<Attribute name="spAuthncontextClassrefMapping">
  <Value>urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport|0|default</Value>
</Attribute>
<Attribute name="spAuthncontextComparisonType">
  <Value>exact</Value>
</Attribute>
<Attribute name="attributeMap">
  <Value>employeenumber=employeenumber</Value>
  <Value>mail=mail</Value>
</Attribute>
<Attribute name="saml2AuthModuleName">

```

```

        <Value></Value>
    </Attribute>
    <Attribute name="localAuthURL">
        <Value></Value>
    </Attribute>
    <Attribute name="intermediateUrl">
        <Value></Value>
    </Attribute>
    <Attribute name="defaultRelayState">
        <Value></Value>
    </Attribute>
    <Attribute name="appLogoutUrl">
        <Value>http://sp1.example.com:8080/saml/logout</Value>
    </Attribute>
    <Attribute name="assertionTimeSkew">
        <Value>300</Value>
    </Attribute>
    <Attribute name="wantAttributeEncrypted">
        <Value></Value>
    </Attribute>
    <Attribute name="wantAssertionEncrypted">
        <Value></Value>
    </Attribute>
    <Attribute name="wantNameIDEncrypted">
        <Value></Value>
    </Attribute>
    <Attribute name="wantPOSTResponseSigned">
        <Value></Value>
    </Attribute>
    <Attribute name="wantArtifactResponseSigned">
        <Value></Value>
    </Attribute>
    <Attribute name="wantLogoutRequestSigned">
        <Value></Value>
    </Attribute>
    <Attribute name="wantLogoutResponseSigned">
        <Value></Value>
    </Attribute>
    <Attribute name="wantMNIRequestSigned">
        <Value></Value>
    </Attribute>
    <Attribute name="wantMNIResponseSigned">
        <Value></Value>
    </Attribute>
    <Attribute name="responseArtifactMessageEncoding">
        <Value>URI</Value>
    </Attribute>
    <Attribute name="cotlist">
        <Value>Circle of Trust</Value></Attribute>
    <Attribute name="saeAppSecretList">
    </Attribute>
    <Attribute name="saeSPUrl">
        <Value></Value>
    </Attribute>
    <Attribute name="saeSPLogoutUrl">
    </Attribute>
    <Attribute name="ECPRequestIDPLISTFinderImpl">
        <Value>com.sun.identity.saml2.plugins.ECPIDPFinder</Value>
    </Attribute>

```

```

<Attribute name="ECPRequestIDPList">
  <Value></Value>
</Attribute>
<Attribute name="ECPRequestIDPListGetComplete">
  <Value></Value>
</Attribute>
<Attribute name="enableIDPPProxy">
  <Value>>false</Value>
</Attribute>
<Attribute name="idpProxyList">
  <Value></Value>
</Attribute>
<Attribute name="idpProxyCount">
  <Value>0</Value>
</Attribute>
<Attribute name="useIntroductionForIDPPProxy">
  <Value>>false</Value>
</Attribute>
<Attribute name="spSessionSyncEnabled">
  <Value>>false</Value>
</Attribute>
<Attribute name="relayStateUrlList">
  </Attribute>
</SPSSOConfig>
<AttributeQueryConfig metaAlias="/attrQuery">
  <Attribute name="signingCertAlias">
    <Value></Value>
  </Attribute>
  <Attribute name="encryptionCertAlias">
    <Value></Value>
  </Attribute>
  <Attribute name="wantNameIDEncrypted">
    <Value></Value>
  </Attribute>
  <Attribute name="cotlist">
    <Value>Circle of Trust</Value>
  </Attribute>
</AttributeQueryConfig>
<XACMLAuthzDecisionQueryConfig metaAlias="/pep">
  <Attribute name="signingCertAlias">
    <Value></Value>
  </Attribute>
  <Attribute name="encryptionCertAlias">
    <Value></Value>
  </Attribute>
  <Attribute name="basicAuthOn">
    <Value>>false</Value>
  </Attribute>
  <Attribute name="basicAuthUser">
    <Value></Value>
  </Attribute>
  <Attribute name="basicAuthPassword">
    <Value></Value>
  </Attribute>
  <Attribute name="wantXACMLAuthzDecisionResponseSigned">
    <Value>>false</Value>
  </Attribute>
  <Attribute name="wantAssertionEncrypted">
    <Value>>false</Value>
  </Attribute>

```

```
</Attribute>  
<Attribute name="cotlist">  
  <Value>Circle of Trust</Value>  
</Attribute>  
</XACMLAuthzDecisionQueryConfig>  
</EntityConfig>
```


Chapter 8

Acting As an OAuth 2.0 Resource Server

OpenIG helps integrate applications into OAuth 2.0 deployments. In this chapter, you will learn to use OpenIG as an OAuth 2.0 Resource Server.

This chapter explains how OpenIG acts as an OAuth 2.0 Resource Server, and follows with a tutorial that shows you how to use OpenIG as a resource server.

8.1. About OpenIG As an OAuth 2.0 Resource Server

The OAuth 2.0 Authorization Framework describes a way of allowing a third-party application to access a user's resources without having the user's credentials. When resources are protected with OAuth 2.0, users can use their credentials with an OAuth 2.0-compliant identity provider, such as OpenAM, Facebook, Google and others to access the resources, rather than setting up an account with yet another third-party application.

In OAuth 2.0, there are four entities involved:

- *Resource owner*: The user who owns protected resources on a resource server.

For example, a resource owner has photos stored in a web service.

- *Resource server*: Provides the user's protected resources to authorized client applications.

In OAuth 2.0, an authorization server grants the client application authorization based on the resource owner's consent.

For example, a web service holds user's photos.

- *Client*: The application that needs access to the protected resources.

For example, a photo printing service needs access to the user's photos.

- *Authorization server*: The service responsible for authenticating resource owners and obtaining their consent to allow client applications to access their resources.

For example, OpenAM can act as the OAuth 2.0 authorization server to authenticate resource owners and obtain their consent. Other services can play this role as well. Google and Facebook, for example, provide OAuth 2.0 authorization services.

In OAuth 2.0, there are different grant mechanisms whereby the client can obtain authorization. One grant mechanism involves the client redirecting the resource owner's browser to the authorization

server to complete authentication and authorization. You might have experienced this grant mechanism yourself when logging in with your identity provider account to access a web service, rather than creating a new account directly with the web service. Whatever the grant mechanism, the client's aim is to get an OAuth 2.0 *access token* from the authorization server.

Access tokens are the credentials used to access protected resources. An access token is a string given by the authorization server that represents the authorization to access protected resources. An access token, like cash, is a bearer token. This means that anyone who has the access token can use it to get the resources. Access tokens therefore must be protected, so requests involving them must go over HTTPS. The advantage of access tokens over passwords or other credentials is that access tokens can be granted and revoked without exposing the user's credentials.

When the client requests access to protected resources, it supplies the access token to the resource server housing the resources. The resource server must then validate the access token. If the access token is found to be valid, then the resource server can let the client have access to the resources.

When OpenIG acts as an OAuth 2.0 resource server, its role is to validate access tokens. How an access token is validated is technically not covered in the specifications for OAuth 2.0. Typically the resource server validates an access token by submitting the token to a token information endpoint. The token information endpoint typically returns the time until the token expires, the OAuth 2.0 *scopes* associated with the token, and potentially other information. In OAuth 2.0, the token scopes are strings that can identify the scope of access authorized to the client, but can also be used for other purposes. For example, OpenAM maps them to user profile attribute values by default, and also allows custom scope handling plugins.

In the tutorial that follows, you configure OpenIG as a resource server, and use OpenAM as the OAuth 2.0 authorization server.

8.2. Preparing the Tutorial

Chapter 2, "*Getting Started*" describes how to configure OpenIG to proxy traffic and capture request and response data. You also learned how to configure OpenIG to use a static request to log in with hard-coded credentials.

You will learn how OpenIG can act as an OAuth 2.0 resource server, validating OAuth 2.0 access tokens and including token info in the context.

This tutorial relies on OpenAM as an OAuth 2.0 authorization server. As an OAuth 2.0 client of OpenAM, you get an access token. You then submit the access token to OpenIG, and OpenIG acts as the resource server.

Before you start this tutorial, prepare OpenIG and the sample application as described in Chapter 2, "*Getting Started*".

OpenIG should be running in Jetty, configured to access the sample application as described in that chapter.

Edit `config.json` to make sure that the `CaptureDecorator` also captures the context. After you make the changes, the object declaration appears as follows:

```
{
  "name": "capture",
  "type": "CaptureDecorator",
  "config": {
    "captureEntity": true,
    "captureContext": true
  }
}
```

Restart Jetty for the changes to take effect. This allows you to view the token information that OpenAM returns.

8.3. Setting Up OpenAM As an Authorization Server

For more complete information about configuring the OAuth 2.0 authorization service in OpenAM, see [Configuring the OAuth 2.0 Authorization Service](#) in the *OpenAM OAuth 2.0 Guide*.

Before you start, install and configure OpenAM on <http://openam.example.com:8088/openam>, with the default configuration. If you use a different configuration, substitute in the tutorial accordingly.

Procedure 8.1. To Set Up OpenAM As an Authorization Server

1. Log in to the OpenAM console as administrator.
2. Configure an OAuth 2.0 Authorization Server:
 - a. In the top level realm, select `Configure OAuth Provider > Configure OAuth 2.0`.
 - b. Accept all of the default values and select `Create`.
3. Configure an OAuth 2.0 Client profile to request OAuth 2.0 access tokens on behalf of the client:
 - a. In the top level realm, select `Applications > OAuth 2.0`.
 - b. In the Agent table, create an OAuth 2.0 agent with the following values:
 - Name: `OpenIG`
 - Password `password`

The name is the OAuth 2.0 `client_id`, and the password is the `client_secret`.

4. Edit the `OpenIG` client profile to add `mail` and `employeenumber` scopes to the Scope(s) list, and then save your work.

Here, you overload these profile settings to pass credentials to OpenIG. This tutorial uses `mail` and `employeenumber` for the sake of simplicity. Both of those attributes are part of a user's profile out of the box with the default OpenAM configuration. Neither of the attributes are needed for anything else in this tutorial.

So, this tutorial uses `mail` to hold the username, and `employeenumber` to hold the password. In a real deployment, you would no doubt use other attributes that depend on how the real user profiles are configured.

5. Create a user whose additional credentials you set in the Email Address and Employee Number fields if you have not already done so for another tutorial:
 1. In the top level realm, select Subjects and create a new subject.
 2. Set the ID to `george`, the password to `costanza`, and fill the other required fields as you like before clicking OK.
 3. Click the user name to edit the profile again, setting Email Address to `george` and Employee Number to `costanza` before clicking Save.
 4. When finished, log out of OpenAM console.

8.4. Configuring OpenIG As a Resource Server

To configure OpenIG as an OAuth 2.0 resource server, you use an `OAuth2ResourceServerFilter` as described in `OAuth2ResourceServerFilter(5)` in the *Configuration Reference*.

The filter expects an OAuth 2.0 access token in an incoming `Authorization` request header, such as the following:

```
Authorization: Bearer 7af41ddd-47a4-40dc-b530-a9aa9f7ceda9
```

The filter then uses the access token to validate the token and to retrieve token information from the authorization server.

On successful validation, the filter creates a new context for the authorization server response, at `contexts.oauth2`.

The context is named `oauth2` and can be reached at `contexts.oauth2` or `contexts['oauth2']`.

The context contains data such as the access token, which can be reached at `contexts.oauth2.accessToken` or `contexts['oauth2'].accessToken`.

Filters and handlers placed after the `OAuth2ResourceServerFilter` in the chain, can access the token info through the context.

If no access token is present in the request, or token validation does not complete successfully, the filter returns an HTTP error status to the user-agent, and OpenIG does not continue processing the request. This is done as specified in the RFC, OAuth 2.0 Bearer Token Usage.

To configure OpenIG as an OAuth 2.0 resource server, add a new route to the OpenIG configuration by including the following route configuration file as `$HOME/.openig/config/routes/06-rs.json`:

```
{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "OAuth2ResourceServerFilter",
          "config": {
            "providerHandler": "ClientHandler",
            "scopes": [
              "mail",
              "employeenumber"
            ],
            "tokenInfoEndpoint": "http://openam.example.com:8088/openam/oauth2/tokeninfo",
            "requireHttps": false
          },
          "capture": "filtered_request",
          "timer": true
        },
        {
          "type": "AssignmentFilter",
          "config": {
            "onRequest": [
              {
                "target": "${session.username}",
                "value": "${contexts.oauth2.accessToken.info.mail}"
              },
              {
                "target": "${session.password}",
                "value": "${contexts.oauth2.accessToken.info.employeenumber}"
              }
            ]
          },
          "timer": true
        },
        {
          "type": "StaticRequestFilter",
          "config": {
            "method": "POST",
            "uri": "http://app.example.com:8081/login",
            "form": {
              "username": [
                "${session.username}"
              ],
              "password": [
                "${session.password}"
              ]
            }
          }
        }
      ]
    }
  },
  "timer": true
}
```

```
        "timer": true
      }
    ],
    "handler": "ClientHandler"
  }
},
"condition": "${matches(request.uri.path, '^/rs')}",
"timer": true
}
```

On Windows, the file name should be `%appdata%\OpenIG\config\routes\06-rs.json`.

Notice the following features of the new route:

- The `OAuth2ResourceServerFilter` includes a client handler to perform the following tasks:
 - Send access token validation requests.
 - Provide the list of scopes that the filter expects to find in access tokens.
 - Provide the OpenAM token info endpoint used to validate access tokens.
 - Set `"requireHttps": false` to allow testing without having to set up keys and certificates. (In production environments, do use HTTPS to protect access tokens.)

After successfully using the token info endpoint to validate an access token, the `OAuth2ResourceServerFilter` creates a new context for the authorization server response, at `${contexts.oauth2.accessToken}`. The context contains the access token and the other info returned by the token info endpoint.

- The `AssignmentFilter` accesses the token info through the context, and injects the credentials from the user profile in OpenAM into `session`.
- The `StaticRequestFilter` retrieves the username and password from `session`, and replaces the original HTTP GET request with an HTTP POST login request that contains the credentials to authenticate.
- The route matches requests to `/rs`.

8.5. Testing the Configuration

To try your configuration you get an access token from OpenAM and use it to access OpenIG, which uses the OAuth 2.0 resource owner password credentials authorization grant.

Procedure 8.2. To Test the Configuration

1. In a terminal window, use a `curl` command similar to the following to retrieve the access token:

```
$ curl \
--user "OpenIG:password" \
--data "grant_type=password&username=george&password=costanza&scope=mail%20employeeenumber" \
http://openam.example.com:8088/openam/oauth2/access_token

{
  "access_token":"aba19a55-468d-45e2-b1c4-decc7202faea",
  "scope":"employeeenumber mail",
  "token_type":"Bearer",
  "expires_in":3599
}
```

2. In the following command, replace `<access_token>` with the access token returned by the previous step, and then run the command:

```
$ curl \
--header "Authorization: Bearer <access_token>" \
http://openig.example.com:8080/rs

...
<h1>USER INFORMATION</h1>
...
  <th>Username</th>
  <td>george</td>
</tr>
...
<h1>REQUEST INFORMATION</h1>
...
  <th>Method</th>
  <td>POST</td>
</tr>
...
  <th>URI</th>
  <td>/login</td>
</tr>
...
  <th>Headers</th>
  <td><b>content-length</b>: 33<br><b>content-type</b>: ...
```

What is happening behind the scenes?

After OpenIG gets the `curl` request, the resource server filter validates the access token with OpenAM, and creates a new context for the authorization server response, at `${contexts.oauth2.accessToken}`. If the access token had been missing or invalid, then the resource server filter would have returned an error status to the user-agent, and the OAuth 2.0 client would then have had to deal with the error.

OpenIG captures the token information into the log, and the `AssignmentFilter` injects the credentials into the session context.

Finally, the `StaticRequestFilter` uses the credentials to log the user in to the sample application, which responds with the user information page.

Chapter 9

Acting As an OAuth 2.0 Client or OpenID Connect Relying Party

OpenIG helps integrate applications into OAuth 2.0 and OpenID Connect deployments. In this chapter, you will:

- Configure OpenIG as an OpenID Connect 1.0 relying party
- Configure OpenIG to use OpenID Connect discovery and dynamic client registration

9.1. About OpenIG As an OAuth 2.0 Client

As described in [Chapter 8, "Acting As an OAuth 2.0 Resource Server"](#), an OAuth 2.0 client is a third-party application that needs access to a user's protected resources.

OpenIG can act as an OAuth 2.0 client when you configure an `OAuth2ClientFilter` as described in `OAuth2ClientFilter(5)` in the *Configuration Reference*. The `OAuth2ClientFilter` handles the process of allowing the user to select a provider, and redirecting the user through the authentication and authorization steps of an OAuth 2.0 authorization code grant. The code grant results in the authorization server returning an access token to the filter.

When an authorization grant succeeds, the `OAuth2ClientFilter` injects the access token data into a configurable target in the context so that subsequent filters and handlers have access to the access token. Subsequent requests can use the access token without reauthentication. If an authorization grant fails, the `failureHandler` is invoked.

If the protected application is an OAuth 2.0 resource server, then OpenIG can send the access token with the resource request.

9.2. About OpenIG As an OpenID Connect 1.0 Relying Party

The specifications available through the [OpenID Connect site](#) describe the OpenID Connect 1.0 authentication layer built on OAuth 2.0.

OpenID Connect 1.0 is a specific implementation of OAuth 2.0, where the identity provider holds the protected resource that the third-party application aims to access.

OpenID Connect 1.0 has the following key entities:

- *End user*: An OAuth 2.0 resource owner whose user information the application needs to access.

The end user wants to use an application through an existing identity provider account without signing up and creating credentials for another web service.

- *Relying Party (RP)*: An OAuth 2.0 client that needs access to the end user's protected user information.

For example, an online mail application needs to know which end user is accessing the application in order to present the correct inbox.

As another example, an online shopping site needs to know which end user is accessing the site in order to present the right offerings, account, and shopping cart.

- *OpenID Provider (OP)*: An OAuth 2.0 authorization server and also resource server that holds the user information and grants access.

The OP has the end user consent to provide the RP with access to some of its user information. Because OpenID Connect 1.0 defines unique identification for an account (subject identifier + issuer identifier), the RP can use this as a key to its own user profile.

For the online mail application, this key could be used to access the mailboxes and related account information. For the online shopping site, this key could be used to access the offerings, account, shopping cart and others. The key makes it possible to serve users as if they had local accounts.

- *UserInfo*: The protected resource that the third-party application aims to access. The information about the authenticated end-user is expressed in a standard format. The user-info endpoint is hosted on the authorization server and is protected with OAuth 2.0.

When OpenIG acts as an OpenID Connect relying party, its ultimate role is to retrieve user information from the OpenID provider, and then to inject that information into the context for use by subsequent filters and handlers.

9.3. Installation Overview

This tutorial takes you through the steps to set up OpenAM as an OpenID Connect provider and OpenIG as an OpenID Connect relying party. For an example configuration with two identity providers, OpenAM and Google, and a login handler, see "Example Configuration For Multiple Identity Providers" in the *Configuration Reference*.

Table 9.1. Tasks for Configuring OpenID Connect

Task	See Section(s)
Set up OpenAM as an OpenID Connect provider.	Section 9.4, "Setting Up OpenAM for OpenID Connect"
Set up OpenIG as a relying party for browser requests to the home page of the sample application.	Section 9.5, "Setting Up OpenIG As a Relying Party"

Task	See Section(s)
Test the configuration.	Section 9.6, "Testing the Configuration"

9.4. Setting Up OpenAM for OpenID Connect

In this part, you create a sample user and set up OpenAM as an OpenID Connect provider.

Before you start, install and configure OpenAM on <http://openam.example.com:8088/openam>, with the default configuration. If you use a different configuration, substitute in the tutorial accordingly.

Procedure 9.1. To Set Up a Sample User

If you haven't set up the user George Costanza in a previous tutorial, do so now.

1. In the OpenAM console, select the top level realm and browse to Subjects.
2. Click New and create a user with the following values:
 - ID: `george`
 - Last Name: `costanza`
 - Full Name: `george costanza`
 - Password: `costanza`
3. In the User window, select the new user and set Email Address: `george@example.com`.

Procedure 9.2. To Set Up OpenAM as an OpenID Connect Provider

1. Configure OpenAM as an OAuth 2.0 Authorization Server and OpenID Connect Provider.
 - a. In the OpenAM console, select the top-level realm and browse to Configure OAuth Provider > Configure OpenID Connect.
 - b. Accept the default values and click Create.
2. Register an OpenID Connect relying party. This step enables OpenIG to communicate as an OAuth 2.0 relying party with OpenAM.
 - a. In the top level realm, select Applications > OAuth 2.0.
 - b. In the Agent table, create an OAuth 2.0 agent with the following values:
 - Name: `oidc_client`
 - Password: `password`

You use these values for `clientId` and `clientSecret` in Section 10.3.2, "Creating a Bearer Module".

- c. Edit the `oidc_client` profile to add the following values:
 - Redirection URIs: `http://openig.example.com:8080/home/id_token/callback`
 - Scope(s): `openid`, `profile`, and `email`.
 - ID Token Signing Algorithm: `HS256`, `HS384`, or `HS512`.

Because the algorithm must be HMAC, the value of `RS256` is not okay.

- d. Click Save and log out of OpenAM.

9.5. Setting Up OpenIG As a Relying Party

This section describes how to use OpenIG Studio to configure OpenIG as a relying party for browser requests to the home page of the sample application. The example refers to the provider configuration in Section 9.4, "Setting Up OpenAM for OpenID Connect".

To configure OpenIG without using OpenIG Studio, add the route in Example 9.1, "Route for OpenIG As a Relying Party" to the OpenIG configuration as `$HOME/.openig/config/routes/07-openid.json`. On Windows, add the route as `%appdata%\OpenIG\config\routes\07-openid.json`.

Procedure 9.3. To Set Up OpenID Connect in OpenIG Studio

Before you start, prepare OpenIG and the sample application as described in Chapter 2, "Getting Started". Make sure that your `config.json` has a main router named `_router`.

1. Access OpenIG Studio on `http://openig.example.com:8080/openig/studio`, and select Protect an Application.
2. In the Create a route window, select Advanced Options and enter the following information:
 - Base URI: `http://app.example.com:8081`
 - Condition: Path: `/home/id_token`
 - Name: `07-openid`
3. Select Create route.
4. Select Authentication, and then enable authentication.
5. Select OpenID Connect, and then select the following options to reflect the configuration in Procedure 9.2, "To Set Up OpenAM as an OpenID Connect Provider":

- Client Filter:
 - Client Endpoint: `/home/id_token`
 - Require HTTPS: Deselect this option
 - Client Registration:
 - Client ID: `oidc_client`
 - Client secret: `password`
 - Scopes: `openid, profile, email`
 - Basic authentication: Select this option
 - Issuer:
 - Well-known Endpoint: `http://openam.example.com:8088/openam/oauth2/.well-known/openid-configuration`
6. Select Save.

A pencil icon in the OpenID Connect panel indicates that OpenID Connect is configured. Select the panel again to edit your settings.

7. On the top-right of the screen, select # > Display, and review the route. The route in Example 9.1, "Route for OpenIG As a Relying Party" should be displayed.
8. Select Deploy to push the route to the OpenIG configuration.

You can check the `$HOME/.openig/config/routes` folder to see that the route is there.

Example 9.1. Route for OpenIG As a Relying Party

```
{
  "name": "07-openid",
  "baseURI": "http://app.example.com:8081",
  "monitor": false,
  "condition": "${matches(request.uri.path, '^/home/id_token')}",
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "OAuth2ClientFilter",
          "name": "OAuth2Client",
          "config": {
            "clientEndpoint": "/home/id_token",
            "failureHandler": {
```


- Requests to `/home/id_token/callback` are expected as redirects from the OAuth 2.0 Authorization Server (OpenID Connect provider). This is why the redirect URI in the client profile in OpenAM is set to `http://openig.example.com:8080/home/id_token/callback`.
- Requests to `/home/id_token/logout` remove the authorization state for the end user, and redirect to the specified URL if a `goto` parameter is provided.

These endpoints are implicitly reserved. Attempts to access them directly can cause undefined errors.

- For convenience in this test, `"requireHttps"` is false. In production environments, set it to true. So that you see the delegated authorization process when you make a request, `"requireLogin"` has the default value `true`.
- The target for storing authorization state information is `${attributes.openid}`. This is where subsequent filters and handlers can find access tokens and user information.

9.6. Testing the Configuration

Procedure 9.4. To Test the Configuration

1. With OpenIG running, access `http://openig.example.com:8080/home/id_token`.

The OpenAM login screen is displayed.

2. Log in to OpenAM with the username `george` and password `costanza`.

An OpenID Connect request to access private information is displayed.

3. Select Allow.

The home page of the sample application is displayed.

What's happening behind the scenes?

- When OpenIG gets the browser request, the `OAuth2ClientFilter` redirects you to authenticate with OpenAM and consent to authorize access to user information. After you authorize access, OpenAM returns an access token to the filter.
- The `OAuth2ClientFilter` uses the access token to get the user information, and then injects the authorization state information into `attributes.openid`.
- The `ClientHandler` redirects the request to the home page of the sample application.

9.7. Adapting the Configuration to Authenticate Automatically to the Sample Application

To authenticate automatically to the sample application, edit `$HOME/.openig/config/routes/07-openid.json` manually as follows:

- In OpenAM and OpenIG, change the endpoints from `/home/openid` to `/openid`. This endpoint directs requests to the login page of the sample application instead of the home page.
- Add a `StaticRequestFilter` like the following to the end of the chain in `07-openid.json`:

```
{
  "type": "StaticRequestFilter",
  "config": {
    "method": "POST",
    "uri": "http://app.example.com:8081/login",
    "form": {
      "username": [
        "${attributes.openid.user_info.sub}"
      ],
      "password": [
        "${attributes.openid.user_info.family_name}"
      ]
    }
  }
}
```

The `StaticRequestFilter` retrieves the username and password from the context, and replaces the original HTTP GET request with an HTTP POST login request containing credentials.

9.8. Using OpenID Connect Discovery and Dynamic Client Registration

OpenID Connect defines mechanisms for discovering and dynamically registering with an identity provider that is not known in advance. These mechanisms are specified in [OpenID Connect Discovery](#) and [OpenID Connect Dynamic Client Registration](#). OpenIG supports discovery and dynamic registration. In this section you will learn how to configure OpenIG to try these features with OpenAM.

Although this tutorial focuses on OpenID Connect dynamic registration, OpenIG also supports dynamic registration as described in RFC 7591, *OAuth 2.0 Dynamic Client Registration Protocol*.

9.8.1. Preparing to Try Discovery and Dynamic Client Registration

This short tutorial builds on the previous tutorial in this chapter. If you have not already done so, start by performing the steps described in Section 9.3, "Installation Overview". This tutorial requires

a recent sample application, as the newer versions include a small WebFinger service that is used here.

When ready, complete preparations for OpenID Connect discovery and dynamic client registration:

- Procedure 9.5, "Preparing OpenAM for OpenID Connect Dynamic Registration"
- Procedure 9.6, "Preparing OpenIG for Discovery and Dynamic Registration"

Procedure 9.5. Preparing OpenAM for OpenID Connect Dynamic Registration

By default, OpenAM does not allow dynamic registration without an access token.

After carrying out the steps described in Section 9.4, "Setting Up OpenAM for OpenID Connect", also perform these steps:

1. Log in to OpenAM console as administrator.
2. In the top-level realm, browse to the Services configuration and display the OAuth2 Provider configuration.
3. Select Advanced OpenID Connect, and enable Allow Open Dynamic Client Registration.
4. Save your work, and log out of OpenAM console.

Procedure 9.6. Preparing OpenIG for Discovery and Dynamic Registration

Follow these steps to add a route demonstrating OpenID Connect discovery and dynamic client registration:

1. Add a new route to the OpenIG configuration, by including the following route configuration file as `$HOME/.openig/config/routes/07-discovery.json`:

```
{
  "heap": [
    {
      "name": "DiscoveryPage",
      "type": "StaticResponseHandler",
      "config": {
        "status": 200,
        "reason": "OK",
        "entity":
          "<!doctype html>
          <html>
          <head>
            <title>OpenID Connect Discovery</title>
            <meta charset='UTF-8'>
          </head>
          <body>
            <form id='form' action='/discovery/login?'>
              Enter your user ID or email address:
              <input type='text' id='discovery' name='discovery'
                placeholder='george or george@example.com' />
            </form>
          </body>
          </html>
          </doctype html>
```

```

<input type='hidden' name='goto'
      value='${contexts.router.originalUri}' />
</form>
<script>
// The sample application handles the WebFinger request,
// so make sure the request is sent to the sample app.
window.onload = function() {
  document.getElementById('form').onsubmit = function() {
    // Fix the URL if not using the default settings.
    var sampleAppUrl = 'http://app.example.com:8081/';
    var discovery = document.getElementById('discovery');
    discovery.value = sampleAppUrl + discovery.value.split('@', 1)[0];
  };
};
</script>
</body>
</html>"
}
},
"handler": {
  "type": "Chain",
  "config": {
    "filters": [
      {
        "name": "DynamicallyRegisteredClient",
        "type": "OAuth2ClientFilter",
        "config": {
          "clientEndpoint": "/discovery",
          "requireHttps": false,
          "requireLogin": true,
          "target": "${attributes.openid}",
          "failureHandler": {
            "type": "StaticResponseHandler",
            "config": {
              "comment": "Trivial failure handler for debugging only",
              "status": 500,
              "reason": "Error",
              "entity": "${attributes.openid}"
            }
          }
        }
      },
      {
        "name": "DiscoveryPage",
        "type": "Page",
        "config": {
          "metadata": {
            "client_name": "My Dynamically Registered Client",
            "redirect_uris": [
              "http://openid.example.com:8080/discovery/callback"
            ],
            "scopes": [
              "openid",
              "profile"
            ]
          }
        }
      }
    ]
  }
}
},
"handler": {
  "type": "Chain",
  "config": {
    "filters": [

```

```

{
  "type": "StaticRequestFilter",
  "config": {
    "method": "POST",
    "uri": "http://app.example.com:8081/login",
    "form": {
      "username": [
        "${attributes.openid.user_info.sub}"
      ],
      "password": [
        "${attributes.openid.user_info.family_name}"
      ]
    }
  }
},
"handler": "ClientHandler"
}
},
"condition": "${matches(request.uri.path, '^/discovery')}",
"baseURI": "http://openid.example.com:8080"
}

```

On Windows, the file name should be `%appdata%\OpenIG\config\routes\07-discovery.json`.

2. Consider the differences with `07-openid.json`:

- For discovery and dynamic client registration, no issuer or client registration is defined. Instead a `StaticResponseHandler` is used as a login handler for the client filter.

The static response handler serves an HTML page that provides important pieces of information to OpenIG:

- The value of a `discovery` parameter.

OpenIG uses the value to perform OpenID Connect discovery. Examples from the specification include `acct:joe@example.com`, `https://example.com:8080/`, and `https://example.com/joe`. First, OpenIG extracts the domain host and port from the value, and attempts to find a match in the `supportedDomains` lists for any issuers that are already configured for the route. If it finds a match, then it can potentially use the issuer's registration end point and avoid an additional request to look up the user's issuer using the `WebFinger` protocol. If there is no match in the supported domains lists, OpenIG uses the `discovery` value as the `resource` for a `WebFinger` request according to the OpenID Connect Discovery protocol.

On success, OpenIG has either found an appropriate issuer in the configuration, or found the issuer using the `WebFinger` protocol. OpenIG can thus proceed to dynamic client registration.

The small JavaScript function in the HTML page transforms user input into a useful `discovery` value for OpenIG. This is not a requirement for deployment, only a convenience for the purposes of this example. Alternatives are described in the discovery protocol specification.

- The value of a `goto` parameter.

The `goto` parameter takes a URI that tells OpenIG where to redirect the end user's browser once the process is complete and OpenIG has injected the OpenID Connect user information into the context. In this case, the user is redirected back to this route so that the innermost chain of the configuration can log the user in to the protected application.

- The OAuth 2.0 client filter specifies a login handler, and dynamic client registration metadata, including a client name, redirection URIs, and scopes.

The login handler points to the login page described above.

OpenIG uses the metadata to prepare the dynamic registration request.

As set out in OAuth2 and OpenID RFCs, the redirection URIs are mandatory for dynamic client registration, to represent an array of redirection URIs used by the client. One of the registered redirection URI values **must** exactly match the `clientEndpoint/callback` URI.

OpenIG also needs the scopes that are required for your application.

- `07-discovery.json` uses the path `/discovery`, whereas `07-openid.json` uses `/home/id_token`.

This distinction makes it easy to keep traffic separate on the two routes with a simple condition as in the following:

```
"condition": "${matches(request.uri.path, '^/discovery')}"
```

9.8.2. Trying OpenID Connect Discovery and Dynamic Client Registration

After following the steps described in Section 9.8.1, "Preparing to Try Discovery and Dynamic Client Registration", test your configuration by browsing to OpenIG at `http://openig.example.com:8080/discovery`.

Provide the following email address: `george@example.com`.

When redirected to the OpenAM login page, log in as user `george`, password `costanza`, and then allow the application access to user information.

If successful, OpenIG logs you in to the sample application as George Costanza, and the sample application returns George's page.

What is happening behind the scenes?

After OpenIG gets the browser request, it returns the example page for discovery. You provide a user ID or email address, and the page transforms that into a `discovery` value. The value is tailored to let OpenIG use the sample application as a WebFinger server. (In the real world the WebFinger server is more likely a service on the issuer's domain, not part of the protected application. For the purposes of

this tutorial the WebFinger service has been embedded in the sample application to avoid leaving you with another server to manage during the tutorial.)

OpenIG learns from the WebFinger service that OpenAM is the issuer for the user. OpenIG retrieves the OpenID Provider configuration from OpenAM, and registers itself dynamically with OpenAM, using the redirection URIs and scopes specified in the OAuth 2.0 client filter configuration.

Once the issuer and client registration are properly configured, the OAuth 2.0 client filter redirects the browser to OpenAM for authentication and authorization to access to the user information. The rest is the same as the previous tutorial in this chapter. For details, see Section 9.6, "Testing the Configuration".

OpenIG reuses issuer and client registration configurations that it builds after discovery and dynamic registration. These dynamically generated configuration objects are held in memory, and do not persist when OpenIG is restarted.

Chapter 10

Transforming OpenID Connect ID Tokens Into SAML Assertions


OpenIG provides a token transformation filter to transform OpenID Connect ID tokens (`id_tokens`) issued by OpenAM into SAML 2.0 assertions. The implementation uses the OpenAM REST Security Token Service (STS) APIs, where the subject confirmation method is Bearer.

Many enterprises use existing or legacy, SAML 2.0-based SSO, but many mobile and social applications are managed by OAuth 2.0/OpenID. Use the OpenIG token transformation filter to bridge OAuth 2.0 and SAML 2.0 frameworks.

10.1. About Token Transformation

The following figure illustrates the basic flow of information between a request, OpenIG, the OpenAM OAuth 2.0 and STS modules, and an application. For a more detailed view of the flow, see Figure 10.2, "Flow of Events".

Figure 10.1. Token Transformation



The basic process is as follows:

1. A user tries to access to a protected resource.
2. If the user is not authenticated, the OAuth2ClientFilter redirects the request to OpenAM. After authentication, OpenAM asks for the user's consent to give OpenIG access to private information.
3. If the user consents, OpenAM returns an id_token to the OAuth2ClientFilter. The filter opens the id_token JWT and makes it available in `attributes.openid.id_token` and `attributes.openid.id_token_claims` for downstream filters.
4. The TokenTransformationFilter calls the OpenAM STS to transform the id_token into a SAML 2.0 assertion.
5. The STS validates the signature, decodes the payload, and verifies that the user issued the transaction. The STS then issues a SAML assertion to OpenIG on behalf of the user.
6. The TokenTransformationFilter makes the result of the token transformation available to downstream handlers in the `issuedToken` property of the `${contexts.sts}` context.

10.2. Installation Overview

This tutorial takes you through the steps to set up OpenAM as an OAuth 2.0 provider and OpenIG as an OpenID Connect relying party.

When a user makes a request, the `OAuth2ClientFilter` returns an `id_token`. The `TokenTransformationFilter` references a REST STS instance that you set up in OpenAM to transform the `id_token` into a SAML 2.0 assertion.

You need to be logged in to OpenAM as administrator to set up the configuration for token transformation, but you do not need special privileges to request a token transformation. For more information, see `TokenTransformationFilter(5)` in the *Configuration Reference*.

Before you start this tutorial:

- Prepare OpenIG as described in Chapter 2, "Getting Started".
- Install and configure OpenAM on `http://openam.example.com:8088/openam`, with the default configuration. If you use a different configuration, substitute in the tutorial accordingly.

Table 10.1. Tasks for Configuring Token Transformation

Task	See Section(s)
Set up OpenAM as an OpenID Connect provider and OpenIG as an an OpenID Connect relying party.	Section 9.4, " Setting Up OpenAM for OpenID Connect "
Create a Bearer authentication module to validate the id_token and OpenAM user.	Section 10.3.2, " Creating a Bearer Module "

Task	See Section(s)
Create an STS REST instance to transform the <code>id_token</code> into a SAML assertion.	Section 10.3.3, "Creating an Instance of STS REST"
Create the routes in OpenIG to send the requests and test the setup.	Section 10.4, " Setting Up OpenIG Routes for Token Transformation "

10.3. Setting Up OpenAM for Token Transformation

10.3.1. Setting Up OpenAM as an OpenID Connect Provider

Set up OpenAM as described in Section 9.4, " Setting Up OpenAM for OpenID Connect ". OpenIG authenticates with OpenAM and retrieves an OpenID Connect ID token (`id_token`) to be transformed by STS.

10.3.2. Creating a Bearer Module

The OpenID Connect ID token bearer module expects an `id_token` in an HTTP request header. It validates the `id_token`, and, if successful, looks up the OpenAM user profile corresponding to the end user for whom the `id_token` was issued. Assuming the `id_token` is valid and the profile is found, the module authenticates the OpenAM user.

You configure the token bearer module to specify how OpenAM gets the information to validate the `id_token`, which request header contains the `id_token`, the identifier for the provider who issued the `id_token`, and how to map the `id_token` claims to an OpenAM user profile.

Important

For OpenAM13.0.0, create the bearer with a **curl** command as described in Procedure 10.1, "To Create a Bearer Module for the `id_token` by Using curl For OpenAM13.0.0". For newer versions of OpenAM, follow the instructions in Procedure 10.2, " To Create a Bearer Module for the `id_token` by Using the OpenAM Console ".

Procedure 10.1. To Create a Bearer Module for the `id_token` by Using curl For OpenAM13.0.0

1. In a terminal window, use a **curl** command similar to the following to retrieve the SSO token for your OpenAM installation.


```

$ curl \
--request POST \
--header "Content-Type: application/json" \
--header "X-OpenAM-Username: amadmin" \
--header "X-OpenAM-Password: password" \
--data "{}" \
http://openam.example.com:8088/openam/json/authenticate

"tokenId": "AQIC5w...NTcy*", "successUrl": "/openam/console" . . .
    
```

For more information about using **curl** for OpenAM authentication, see the OpenAM Developer's Guide.

2. Replace `<tokenId>` in the following command with the `tokenId` returned by the previous step, and then run the command:

```

$ curl -X POST -H "Content-Type: application/json" \
-H "iplanetDirectoryPro: <tokenId>" \
-d \
'{
"cryptoContextValue": "password",
"jwtToLdapAttributeMappings": ["sub=uid", "email=mail"],
"principalMapperClass": "org.forgerock.openam.authentication.modules.oidc.JwtAttributeMapper",
"acceptedAuthorizedParties": ["oidc_client"],
"idTokenHeaderName": "oidc_id_token",
"accountProviderClass": "org.forgerock.openam.authentication.modules.common.mapping.DefaultAccountProvider",
"idTokenIssuer": "http://openam.example.com:8088/openam/oauth2",
"cryptoContextType": "client_secret",
"audienceName": "oidc_client",
" id": "oidc"
}' \
http://openam.example.com:8088/openam/json/realm-config/authentication/modules/openidconnect?_action=create

http://openam.example.com:8088/openam/json/realm-config/authentication/modules/openidconnect?_action=create
{"principalMapperClass": "org.forgerock.openam.authentication.modules.oidc.JwtAttributeMapper", . . .
    
```

The Bearer module is created in OpenAM.

Procedure 10.2. To Create a Bearer Module for the `id_token` by Using the OpenAM Console

1. Log in to the OpenAM console as administrator.

2. In the top level realm, select Authentication > Modules.
3. Select Add Module and create a new bearer module with the following characteristics:
 - Module name: `oidc`
 - Type: `OpenID Connect id_token bearer`
4. In the configuration page, enter the following values and leave the other fields with the default values:
 - Audience name: `oidc_client`, the name OAuth 2.0/OpenID Connect client.
 - List of accepted authorized parties: `oidc_client`.
 - OpenID Connect validation configuration type: `client_secret`
 - OpenID Connect validation configuration value: `password`.
This is the password of the OAuth 2.0/OpenID Connect client.
 - Name of OpenID Connect ID Token Issuer: `http://openam.example.com:8088/openam/oauth2`
 - Client Secret (available from AM 5.0): `password`
5. Select Save Changes.

10.3.3. Creating an Instance of STS REST

The REST STS instance exposes a preconfigured transformation under a specific REST endpoint.

For more information about setting up a REST STS instance, see the *OpenAM Security Token Service Guide* .

For information about configuring keystores, see *Configuring Keystores* in the *OpenAM Setup and Maintenance Guide*.

1. In the top level realm, select STS.
2. Create a Rest STS instance with the following characteristics:
 - Deployment Configuration
 - Deployment Url Element: `openig`
This value identifies the STS instance and is used by the `instance` parameter in the `TokenTransformationFilter`.
 - Issued SAML2 Token Configuration

- SAML2 issuer Id: `OpenAM`
- Service Provider Entity Id: `openig_sp`
- NameIdFormat: Select `nameid:format:transient`
- Attribute Mappings: Add `password=mail` and `userName=uid`.
- Keystore Password: Retrieve the value from your OpenAM installation, at `/path/to/openam/openam/.storepass`

This property specifies the password used to decrypt the keystore.

- Signature Key Alias: By default, the value is `test`

This property specifies the private key alias in the keystore used to sign the assertion. It is configured in the Certificate Alias field of the OpenAM Security Key Store Tab.

- Signature Key Password: `changeit`

This property specifies the password of the private key used to sign the assertion.

Note

For STS, it isn't necessary to create a SAML SP configuration in OpenAM.

- OpenIdConnect Token Configuration
 - The id of the OpenIdConnect Token Provider: `oidc`
 - Token signature algorithm: The value must be consistent with the one you selected in Procedure 9.2, " To Set Up OpenAM as an OpenID Connect Provider ", `HMAC SHA 256`
 - Client secret (for HMAC-signed-tokens): `password`
 - The audience for issued tokens: `oidc_client`.

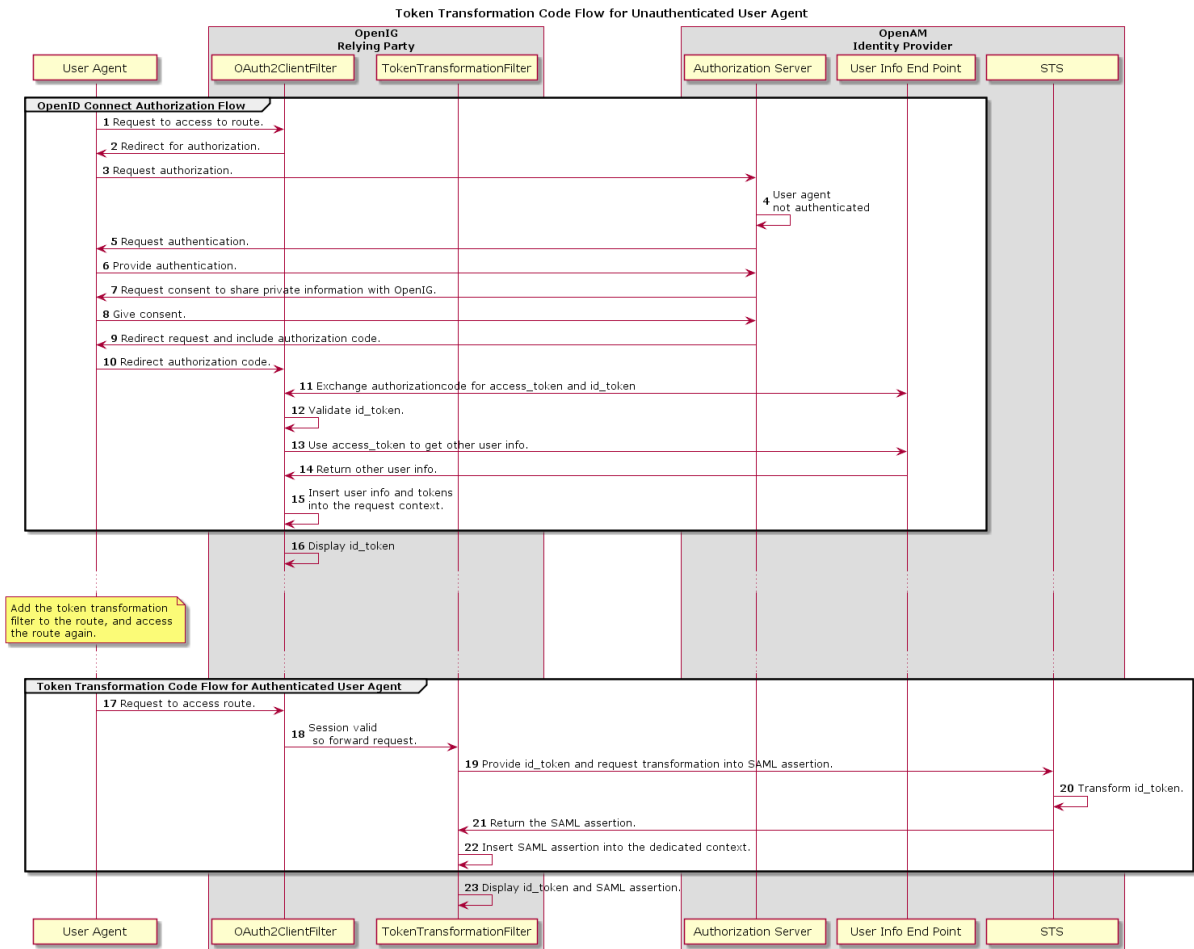
3. Select Create.

4. Log out of OpenAM.

10.4. Setting Up OpenIG Routes for Token Transformation

The following sequence diagram shows what happens when you set up and access routes for token transformation.

Figure 10.2. Flow of Events



Procedure 10.3. To Set Up Routes to Create an id_token

Errors that occur during the token transformation cause an error response to be returned to the client and an error message to be logged for the OpenIG administrator.

1. Edit `config.json` to comment the baseURI in the top-level handler. The handler declaration appears as follows:

```

{
  "handler": {
    "type": "Router",
    "name": "_router",
    "_baseURI": "http://app.example.com:8081",
    "capture": "all"
  }
}

```

Restart OpenIG for the changes to take effect.

2. Add the following route to the OpenIG configuration as `$HOME/.openig/config/routes/50-idthoken.json`

On Windows, add the route as `%appdata%\OpenIG\config\routes\50-idthoken.json`.

```

{
  "heap": [
    {
      "name": "openam",
      "type": "ClientRegistration",
      "config": {
        "clientId": "oidc_client",
        "clientSecret": "password",
        "issuer": {
          "type": "Issuer",
          "config": {
            "wellKnownEndpoint": "http://openam.example.com:8088/openam/oauth2/.well-known/openid-configuration"
          }
        },
      },
      "scopes": [
        "openid",
        "profile",
        "email"
      ]
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "OAuth2ClientFilter",
          "config": {
            "clientEndpoint": "/home/id_token",
            "requireHttps": false,
            "requireLogin": true,
            "registrations": "openam",
            "target": "${attributes.openid}",
            "failureHandler": {
              "type": "StaticResponseHandler",
              "config": {
                "entity": "OAuth2ClientFilter failed...",
                "reason": "NotFound",

```

```

        "status": 500
      }
    }
  },
  ],
  "handler": {
    "type": "StaticResponseHandler",
    "config": {
      "entity": "{\|\"id_token\|":\|n\|\"${attributes.openid.id_token}\|"} \|n\|n\|n{\|\"saml_assertions\|":\|n\|
\|\"${contexts.sts.issuedToken}\|\"}",
      "reason": "Found",
      "status": 200
    }
  }
},
},
"condition": "${matches(request.uri.path, '^/home/id_token')}"
}

```

Notice the following features of the route:

- The route matches requests to `/home/id_token`.
 - The client registration is defined in the heap. The registration holds configuration parameters provided during Procedure 9.2, "To Set Up OpenAM as an OpenID Connect Provider", and OpenIG uses these parameters to connect with OpenAM.
 - The `OAuth2ClientFilter` enables OpenIG to act as an OpenID Connect relying party.
 - The client endpoint is set to `/home/id_token`, so the service URIs for this filter on the OpenIG server are `/home/id_token/login`, `/home/id_token/logout`, and `/home/id_token/callback`.
 - For convenience in this test, `requireHttps` is false. In production environments, set it to true. So that you see the delegated authorization process when you make a request, `requireLogin` is true.
 - The target for storing authorization state information is `${attributes.openid}`. This is where subsequent filters and handlers can find access tokens and user information.
 - When the request succeeds, a `StaticResponseHandler` retrieves and displays the `id_token` from the target `{attributes.openid.id_token}`.
3. To prevent conflicts with the route you set up in Chapter 9, "*Acting As an OAuth 2.0 Client or OpenID Connect Relying Party*", rename the route `07-openid.json` to `07-openid.json.old`
 4. With OpenIG running, access `http://openig.example.com:8080/home/id_token`.
The OpenAM login screen is displayed.
 5. Log in to OpenAM with the username `george` and password `costanza`.
An OpenID Connect request to access private information is displayed.

6. Select Allow.

The `id_token` is displayed above an empty placeholder for the SAML assertion.

```
{
  "id_token":
  "eyJhdHlwIjogIkpXVCIsICJhbGciOiAiA1SFMjNTYiIH0.eyJYXRfaGFzaCI6ICJ . . ."
}

{"saml_assertions":
""}
```

Procedure 10.4. To Edit the Route to Transform the `id_token` Into a SAML Assertion

1. Add the following filter in the chain of `50-idtoken.json`, after the `OAuth2ClientFilter`. An example of the edited route is at the end of this procedure.

```
{
  "type": "TokenTransformationFilter",
  "config": {
    "openamUri": "http://openam.example.com:8088/openam",
    "username": "oidc_client",
    "password": "password",
    "idToken": "${attributes.openid.id_token}",
    "instance": "openig",
    "ssoTokenHeader": "iPlanetDirectoryPro"
  }
}
```

Notice the following features of the new filter:

- Requests from this filter are made to `http://openam.example.com:8088/openam`.
- The username and password are for the OpenAM subject set up in Procedure 9.2, "To Set Up OpenAM as an OpenID Connect Provider".
- The `id_token` parameter defines where this filter gets the `id_token` created by the `OAuth2ClientFilter`.

The `TokenTransformationFilter` makes the result of the token transformation available to downstream handlers in the `issuedToken` property of the `${contexts.sts}` context.

- The `instance` parameter must match the `Deployment URL Element` for the REST STS instance.

2. With OpenIG running, access `http://openig.example.com:8080/home/id_token`.

The SAML assertions are displayed under the `id_token`.

```
{
  "id_token":
  "eyJhdHlwIjogIkpXVCIsICJhbGciOiAiA1SFMjNTYiIH0.eyJYXRfaGFzaCI6ICJ . . ."
}

{"saml_assertions":
<"saml:Assertion xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion" Version= . . ."}
}
```

Example of the final `50_idtoken.json`:

```
{
  "heap": [
    {
      "name": "openam",
      "type": "ClientRegistration",
      "config": {
        "clientId": "oidc_client",
        "clientSecret": "password",
        "issuer": {
          "type": "Issuer",
          "config": {
            "wellKnownEndpoint": "http://openam.example.com:8088/openam/oauth2/.well-known/openid-configuration"
          }
        },
        "scopes": [
          "openid",
          "profile",
          "email"
        ]
      }
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "OAuth2ClientFilter",
          "config": {
            "clientEndpoint": "/home/id_token",
            "requireHttps": false,
            "requireLogin": true,
            "registrations": "openam",
            "target": "${attributes.openid}",
            "failureHandler": {
              "type": "StaticResponseHandler",
              "config": {
                "entity": "OAuth2ClientFilter failed...",
                "reason": "NotFound",
                "status": 500
              }
            }
          }
        },
        {
          "type": "TokenTransformationFilter",
          "config": {
            "openamUri": "http://openam.example.com:8088/openam",
            "username": "oidc_client",
            "password": "password",
            "idToken": "${attributes.openid.id_token}",
            "instance": "openig",
            "ssoTokenHeader": "iPlanetDirectoryPro"
          }
        }
      ]
    }
  }
}
```


Chapter 11

Supporting UMA Resource Servers

OpenIG provides experimental support for building a User-Managed Access (UMA) resource server. In this chapter, you will learn:

- Where OpenIG fits in the UMA picture
- How to configure OpenIG to allow a resource owner to register UMA resource sets
- How to configure OpenIG to protect access to resources using UMA

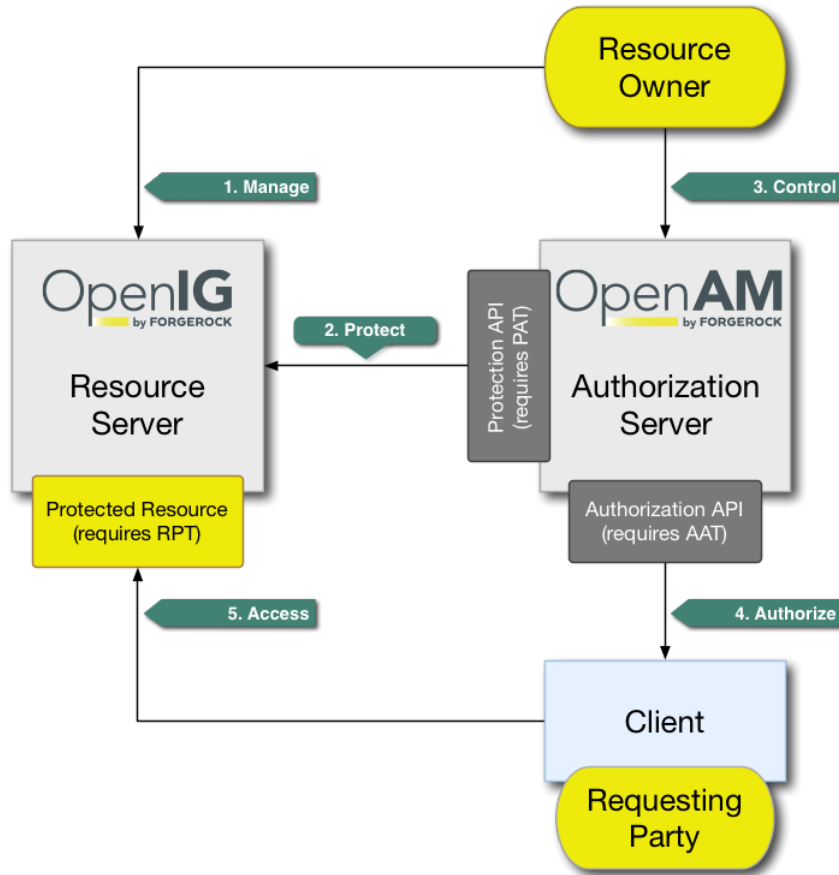
11.1. About OpenIG in the UMA Resource Server Role

This section covers the role OpenIG plays as UMA resource server.

11.1.1. About UMA

User-Managed Access (UMA) Profile of OAuth 2.0 defines a workflow that allows resource owners to share their protected resources with requesting parties. Figure 11.1, "UMA Workflow" illustrates the relationships where OpenIG protects the resource server.

Figure 11.1. UMA Workflow



The actions that form the UMA workflow are as follows:

1. Manage

The resource owner manages their resources on the resource server.

When using OpenIG to protect the resources, OpenIG creates the *resource sets* that describe what the resource owner shares. Resource set registration is covered in *OAuth 2.0 Resource Set Registration*.

2. Protect

The resource owner links their resource server and chosen authorization server, such as OpenAM.

The authorization server provides a protection API so that the resource server can register sets of resources. Use of the protection API requires a *protection API token* (PAT), an OAuth 2.0 token with scope `uma_protection`.

3. Control

The resource owner controls who has access to their registered resources by creating policies on the authorization server.

Only a resource owner can create policies for their registered resources.

4. Authorize

The client, acting on behalf of the requesting party, uses the authorization server's authorization API to acquire a *requesting party token* (RPT). The requesting party or client may need further interaction with the authorization server at this point, for example, to supply identity claims. Use of the authorization API requires an *authorization API token* (AAT), an OAuth 2.0 token with scope `uma_authorization`.

5. Access

The client presents the RPT to the resource server, which verifies its validity with the authorization server and, if both valid and containing sufficient permissions, returns the protected resource to the requesting party.

11.1.2. Sharing Protected Resources

When acting as a UMA resource server, OpenIG helps the resource owner register resource sets with the authorization server. The resource owner then interacts with the authorization server to authorize access to registered resources.

This process of sharing protected resources includes the following steps:

1. The OpenIG administrator configures a route with the following:

- An `UmaService` that describes OpenIG registration as an OAuth 2.0 client of the authorization server and the resource sets to share, including resource path patterns and scopes.

The `UmaService` exposes a REST API to use when managing resource sets.

For details, see `UmaService(5)` in the *Configuration Reference*.

- An `UmaFilter` that acts as a policy enforcement point, protecting access to resources on the route.

For details, see `UmaFilter(5)` in the *Configuration Reference*.

2. The resource owner obtains a PAT from the authorization server.
3. The resource owner provides the PAT and a resource path to OpenIG, which registers a corresponding resource set with the authorization server.

OpenIG responds with the resource set identifier and a link where the resource owner can set up access permissions.

4. The resource owner creates policies on the authorization server to authorize requesting parties to access protected resources.

11.1.3. Accessing Protected Resources

When acting as an UMA resource server, OpenIG interacts with the UMA client and the authorization server. OpenIG challenges the UMA client to gain authorization with the authorization server, and enforces policy for protected resources according to policy decisions by the authorization server.

The process of accessing protected resources can start after the process of sharing resources is successfully completed. The process of accessing a protected resource includes the following steps:

1. The requesting party attempts to access the resource without an RPT.

OpenIG responds with an UMA `WWW-Authenticate` header, and a ticket that the requesting party can use to get an RPT.

2. The requesting party gets an AAT from the authorization server.

This step lets the authorization server authenticate the requesting party.

3. The requesting party uses AAT from the authorization server, and the ticket from OpenIG to obtain an RPT from the authorization server.

4. The requesting party uses the RPT to access the resource as originally intended.

11.1.4. Understanding the UMA API With an API Descriptor

The UMA share endpoint serves API descriptors at runtime. When you retrieve an API descriptor for the endpoint, a JSON that describes the API for the endpoint is returned.

You can use the API descriptor with a tool such as Swagger UI to generate a web page that helps you to view and test the endpoint. For information, see Section 1.10, "Understanding OpenIG APIs With API Descriptors".

11.1.5. Limitations of This Implementation

Keep the following points in mind when using OpenIG as an UMA resource server:

- OpenIG depends on the resource owner for the PAT.

When a PAT expires, no refresh token is available to OpenIG. The resource owner must perform the entire share process again with a new PAT in order to authorize access to protected resources. The resource owner should delete the old resource and create a new one.

- Data about PATs and shared resources is held in memory.

OpenIG has no mechanism for persisting the data across restarts. When OpenIG stops and starts again, the resource owner must perform the entire share process again.

- UMA client applications for sharing and accessing protected resources must deal with UMA error conditions and OpenIG error conditions.
- OpenIG exposes a REST API to manage share objects that is not protected by default.
- When matching protected resource paths with share patterns, OpenIG takes the longest match.

For example, if resource owner Alice shares `/photos/.*` with Bob, and `/photos/vacation.png` with Charlie, and then Bob attempts to access `/photos/vacation.png`, OpenIG applies the sharing permissions for Charlie, not Bob. As a result, Bob can be denied access.

11.2. Preparing the Tutorial

This tutorial guides you through one way to set up OpenIG as an UMA resource server. It uses OpenAM as an authorization server for OAuth 2.0 and for UMA, and uses the sample application as a resource to protect and for files that serve as a basic UMA client.

Table 11.1. Tasks for Configuring Policy Enforcement

Task	See Section(s)
Modify the OpenAM configuration to allow cross-site access.	Procedure 11.1, "To Enable CORS Support for OpenAM"
Configure OpenAM as an authorization server.	Procedure 11.2, "To Configure OpenAM As an OAuth 2.0 Authorization Server and UMA Authorization Server"
Register client profiles in OpenAM for OAuth 2.0 and UMA.	Procedure 11.3, "To Register Client Profiles in OpenAM"
Create a subject to act as a resource owner and a subject to act as a requesting party.	Procedure 11.4, "To Create a Resource Owner and Requesting Party"
Set up the OpenIG configuration for an UMA resource server	Procedure 11.5, "To Set Up OpenIG As a UMA Resource Server"
If you use a configuration that is different from that described in this chapter, adjust the sample to your configuration.	Section 11.6, "Editing the Example to Match Custom Settings"

11.3. Setting Up OpenAM As an Authorization Server

This section describes the following tasks to set up OpenAM as an authorization server:

- Enabling cross-origin resource sharing (CORS) support in OpenAM
- Configuring OpenAM as an authorization server
- Registering UMA client profiles with OpenAM
- Setting up a resource owner (Alice) and requesting party (Bob)

Before you start, install and configure OpenAM on <http://openam.example.com:8088/openam>, with the default configuration. If you use a different configuration, substitute in the tutorial accordingly.

Procedure 11.1. To Enable CORS Support for OpenAM

For information about CORS support, see the OpenAM product documentation. This procedure describes how to modify the OpenAM configuration to allow cross-site access.

Caution

The settings in this section are suggestions for this tutorial, and are not intended as documentation for setting up OpenAM CORS support on a server in production.

In production, do not add `Content-Type` to the list of allowed headers in the CORS filter. Doing so can expose OpenAM to Cross-Site Request Forgery (CSRF) attacks.

1. In the `WEB-INF/web.xml` file of OpenAM, edit the URL pattern for `CORSFilter` to match all endpoints:

```
<filter-mapping>
  <filter-name>CORSFilter</filter-name>
  <url-pattern>/*</url-pattern>
</filter-mapping>
```

2. In the same file, edit the `CORSFilter` configuration to authorize cross-site access for origins, hosts, and headers, as shown in the following excerpt:

```
<filter>
  <filter-name>CORSFilter</filter-name>
  <filter-class>org.forgerock.openam.cors.CORSFilter</filter-class>
  <init-param>
    <description>
      Accepted Methods (Required):
      A comma separated list of HTTP methods for which to accept CORS requests.
    </description>
    <param-name>methods</param-name>
    <param-value>POST,GET,PUT,DELETE,PATCH,OPTIONS</param-value>
  </init-param>
  <init-param>
    <description>
```

```

        Accepted Origins (Required):
        A comma separated list of origins from which to accept CORS requests.
    </description>
    <param-name>origins</param-name>
    <param-value>http://app.example.com:8081,http://openig.example.com:8080,http://
openam.example.com:8088</param-value>
</init-param>
<init-param>
    <description>
        Allow Credentials (Optional):
        Whether to include the Vary (Origin)
        and Access-Control-Allow-Credentials headers in the response.
        Default: false
    </description>
    <param-name>allowCredentials</param-name>
    <param-value>true</param-value>
</init-param>
<init-param>
    <description>
        Allowed Headers (Optional):
        A comma separated list of HTTP headers
        which can be included in the requests.
    </description>
    <param-name>headers</param-name>
    <param-value>
        Authorization,Content-Type,iPlanetDirectoryPro,X-OpenAM-Username,X-OpenAM-
        Password,Accept,Accept-Encoding,Connection,Content-Length,Host,Origin,User-Agent,Accept-
        Language,Referer,Dnt,Accept-Api-Version,If-None-Match
    </param-value>
</init-param>
<init-param>
    <description>
        Expected Hostname (Optional):
        The name of the host expected in the request Host header.
    </description>
    <param-name>expectedHostname</param-name>
    <param-value>openam.example.com:8088</param-value>
</init-param>
<init-param>
    <description>
        Exposed Headers (Optional):
        The comma separated list of headers
        which the user-agent can expose to its CORS client.
    </description>
    <param-name>exposeHeaders</param-name>
    <param-value>Access-Control-Allow-Origin,Access-Control-Allow-Credentials,Set-Cookie,WWW-
Authenticate</param-value>
</init-param>
<init-param>
    <description>
        Maximum Cache Age (Optional):
        The maximum time that the CORS client can cache
        the pre-flight response, in seconds.
        Default: 600
    </description>
    <param-name>maxAge</param-name>
    <param-value>600</param-value>
</init-param>

```



```
</filter>
```

Procedure 11.2. To Configure OpenAM As an OAuth 2.0 Authorization Server and UMA Authorization Server

1. Log in to the OpenAM console as administrator.
2. In the top-level realm, select Configure OAuth Provider > Configure OAuth 2.0, accept the default values, and select Create.

An OpenAM policy named `OAuth2ProviderPolicy` is created for the authorization end point. The PAT and AAT are obtained through the OAuth 2.0 access token endpoint. The RPT is obtained through the UMA endpoint.

If you plan to build your own examples or modify the sample clients, consider extending the default token lifetimes to 3600 seconds.

3. Select Configure OAuth Provider > Configure User Managed Access, accept the default values, and select Create.
4. Select Services > UMA Provider, deselect Require Trust Elevation, and save the changes.

Procedure 11.3. To Register Client Profiles in OpenAM

Follow these steps to register client profiles for OAuth 2.0 and UMA:

1. In the top level realm, select Applications > OAuth 2.0.
2. In the Agent table, create an agent to use when sharing resources:
 - Name (client_id): `OpenIG`
 - Password (client_secret): `password`
 - Scope: `uma_protection`
3. Create an agent to use when accessing resources, with the following values:
 - Name (client_id): `UmaClient`
 - Password (client_secret): `password`
 - Scope: `uma_authorization`

Procedure 11.4. To Create a Resource Owner and Requesting Party

Follow these steps to create subjects in the top-level realm:

1. In the top-level realm, select Subjects.

2. Add a subject to act as a resource owner, with the following values:
 - ID: `alice`
 - First Name: `Alice`
 - Last Name: `User`
 - Full Name: `Alice User`
 - Password: `password`
 - User Status: Active
3. Add a subject to act as a requesting party, with the following values:
 - ID: `bob`
 - First Name: `Bob`
 - Last Name: `User`
 - Full Name: `Bob User`
 - Password: `password`
 - User Status: Active

11.4. Setting Up OpenIG As a UMA Resource Server

Procedure 11.5. To Set Up OpenIG As a UMA Resource Server

Before you start, prepare OpenIG and the sample application as described in Chapter 2, "Getting Started".

1. Add the following file as `$HOME/.openig/config/admin.json`.

On Windows, add the route as `%appdata%\OpenIG\config\admin.json`:

```
{
  "heap": [
    {
      "name": "ClientHandler",
      "type": "ClientHandler"
    },
    {
      "name": "ApiProtectionFilter",
      "type": "ScriptableFilter",
      "config": {
        "type": "application/x-groovy",

```

```

        "file": "CorsFilter.groovy"
    }
}
],
"prefix": "openig"
}

```

To allow CORS support for the UMA share API, this route overrides the default `ApiProtectionFilter` that protects the reserved administrative route. By default, the administrative route is under `/openig`. For information, see `AdminHttpApplication(5)` in the *Configuration Reference*.

2. Add the following route to the OpenIG configuration as `$HOME/.openig/config/routes/00-uma.json`.

On Windows, add the route as `%appdata%\OpenIG\config\routes\00-uma.json`.

```

{
  "heap": [
    {
      "name": "UmaService",
      "type": "UmaService",
      "config": {
        "protectionApiHandler": "ClientHandler",
        "authorizationServerUri": "http://openam.example.com:8088/openam/",
        "clientId": "OpenIG",
        "clientSecret": "password",
        "resources": [
          {
            "comment": "Protects all resources matching the following pattern.",
            "pattern": ".*",
            "actions": [
              {
                "scopes": [
                  "#read"
                ],
                "condition": "${request.method == 'GET'}"
              },
              {
                "scopes": [
                  "#create"
                ],
                "condition": "${request.method == 'POST'}"
              }
            ]
          }
        ]
      }
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ScriptableFilter",

```

```

        "config": {
            "type": "application/x-groovy",
            "file": "CorsFilter.groovy"
        }
    },
    {
        "type": "UmaFilter",
        "config": {
            "protectionApiHandler": "ClientHandler",
            "umaService": "UmaService"
        }
    }
],
"handler": "ClientHandler"
}
},
"condition": "${request.uri.host == 'app.example.com'}"
}

```

Notice the following features of the route:

- The `UmaService` describes the resources that a resource owner can share. It uses OpenAM as the authorization server, relying on one of the client profiles you created (`client_id`: OpenIG).

The `UmaService` provides a REST API to manage sharing of resource sets.

- The handler for the route chains together the CORS filter, the `UmaFilter`, and the default handler.

The `UmaFilter` manages requesting party access to protected resources, using the `UmaService`. Protected resources are on the sample application, which responds to requests on port 8081.

- The route matches requests to `app.example.com`.

3. Add the following script to the OpenIG configuration as `$HOME/.openig/scripts/groovy/CorsFilter.groovy`.

On Windows, add the script as `%appdata%\OpenIG\scripts\groovy\CorsFilter.groovy`.

```

import org.forgerock.http.protocol.Response
import org.forgerock.http.protocol.Status

if (request.method == 'OPTIONS') {
    /**
     * Supplies a response to a CORS preflight request.
     *
     * Example response:
     *
     * HTTP/1.1 200 OK
     * Access-Control-Allow-Origin: http://app.example.com:8081
     * Access-Control-Allow-Methods: POST
     * Access-Control-Allow-Headers: Authorization
     * Access-Control-Allow-Credentials: true
     * Access-Control-Max-Age: 3600
    */
}

```

```
*/
def origin = request.headers['Origin']?.firstValue
def response = new Response(Status.OK)

// Browsers sending a cross-origin request from a file might have Origin: null.
response.headers.put("Access-Control-Allow-Origin", origin)
request.headers['Access-Control-Request-Method']?.values.each() {
    response.headers.add("Access-Control-Allow-Methods", it)
}
request.headers['Access-Control-Request-Headers']?.values.each() {
    response.headers.add("Access-Control-Allow-Headers", it)
}
response.headers.put("Access-Control-Allow-Credentials", "true")
response.headers.put("Access-Control-Max-Age", "3600")

return response
}

return next.handle(context, request)
/**
 * Adds headers to a CORS response.
 */
*/
    .thenOnResult({ response ->
if (response.status.isServerError()) {
    // Skip headers if the response is a server error.
} else {
    def headers = [
        "Access-Control-Allow-Origin": request.headers['Origin']?.firstValue,
        "Access-Control-Allow-Credentials": "true",
        "Access-Control-Expose-Headers": "WWW-Authenticate"
    ]
    response.headers.addAll(headers)
}
})
```

This script adds a CORS filter to include headers for cross-origin requests.

The tutorial involves JavaScript clients that are served by the sample application, and so not from the same origin as OpenAM or OpenIG. The route uses a CORS filter to include appropriate response headers for cross-origin requests.

The CORS filter handles pre-flight (HTTP OPTIONS) requests, and responses for all HTTP operations. The logic for the filter is provided through the script.

The filter adds the appropriate headers to CORS requests. Pre-flight requests are diverted to a dedicated handler, which returns the response directly to the user agent. For all other requests, the headers are added to the response.

For information about scripting filters and handlers, see Chapter 14, *"Extending the ForgeRock Identity Gateway"*.

4. Restart OpenIG to reload the configuration.

11.5. Test the Configuration

Follow these steps to test the configuration and demonstrate OpenIG acting as an UMA resource server:

1. Log out of OpenAM.
2. Browse to <http://app.example.com:8081/uma/>.

If you used the settings described in this chapter and Section 2.3, "Install the Sample Application", your configuration should match the displayed configuration. If you used other settings, you might need to edit the sample files to match your settings. For information, see Section 11.6, "Editing the Example to Match Custom Settings".

3. Select the link to demonstrate Alice sharing resources.
4. On Alice's page, select Share with Bob to simulate Alice sharing resources as described in Section 11.1.2, "Sharing Protected Resources".

The following items are displayed:

- The PAT that Alice receives
- The metadata for the resource set that Alice registers through OpenIG
- The result of Alice authenticating with OpenAM in order to create a policy
- The successful result `{}` when Alice creates the policy

If the step fails, get help in Section 18.1, "Troubleshooting the UMA Example".

5. Go back to the first page, and select the link to demonstrate Bob accessing resources.
6. On Bob's page, select Get Alice's resources to simulate Bob accessing one of Alice's resources as described in Section 11.1.3, "Accessing Protected Resources".

The following items are displayed:

- The ticket returned initially
- The AAT that Bob gets to obtain the RPT
- The RPT that Bob gets in order to request the resource again
- The final response containing the body of the resource

11.6. Editing the Example to Match Custom Settings

If you use a configuration that is different to that described in this chapter, consider the following tasks to adjust the sample to your configuration:

1. Unpack the UMA files from the sample application described in Section 2.3, "Install the Sample Application" to temporary folder:

```
$ mkdir /tmp/uma
$ cd /tmp/uma
$ jar -xvf /path/to/IG-doc-samples-5.0.0.jar uma
created: uma/
inflated: uma/alice.html
inflated: uma/bob.html
inflated: uma/common.js
inflated: uma/index.html
inflated: uma/style.css
```

2. Edit the configuration in `common.js`, `alice.html`, and `bob.html` to match your settings.
3. Repack the UMA sample client files and then restart the sample application:

```
$ jar -uvf /path/to/IG-doc-samples-5.0.0.jar uma
adding: uma/(in = 0) (out= 0)(stored 0%)
adding: uma/index.html(in = 1698) (out= 880)(deflated 48%)
adding: uma/common.js(in = 4265) (out= 1319)(deflated 69%)
adding: uma/bob.html(in = 5427) (out= 1811)(deflated 66%)
adding: uma/style.css(in = 1403) (out= 696)(deflated 50%)
adding: uma/alice.html(in = 5494) (out= 1762)(deflated 67%)
```

4. If necessary, adjust the CORS settings for OpenAM.

Chapter 12

Configuring Routers and Routes

OpenIG provides routers and routes to handle requests and their context. In this chapter, you will learn about:

- How routers and routes are configured
- How to read, create, edit, or delete routes through Common REST or OpenIG Studio
- How to prevent changes to routes when OpenIG is running

12.1. Configuring Routers

When you set up the first tutorial, you configured a `Router` in the top-level `config.json` file, which is shown here again in the following listing:

```
{
  "handler": {
    "type": "Router",
    "name": "_router",
    "baseURI": "http://app.example.com:8081",
    "capture": "all"
  },
  "heap": [
    {
      "name": "JwtSession",
      "type": "JwtSession"
    },
    {
      "name": "capture",
      "type": "CaptureDecorator",
      "config": {
        "captureEntity": true,
        "_captureContext": true
      }
    }
  ]
}
```

In this configuration, all requests are passed with the default settings to the router. The router scans `$HOME/.openig/config/routes` at startup, and rescans the directory every 10 seconds. If routes have been added, deleted or changed, the router applies the changes.

The main router and any subrouters are used to build the monitoring endpoints. For information about monitoring endpoints, see Chapter 15, "Auditing and Monitoring". For information about the parameters of a router, see Router(5) in the *Configuration Reference*.

12.2. Configuring Routes

Routes are JSON configuration files that handle requests and their context, and then hand off any request they accept to a handler. Another way to think of a route is like an independent dispatch handler, as described in Dispatcher(5) in the *Configuration Reference*.

Routes can have a base URI to change the scheme, host, and port of the request.

For information about the parameters of routes, see Route(5) in the *Configuration Reference*.

12.2.1. Configuring Objects Inline or In the Heap

If you use an object only once in the configuration, you can declare it inline in the route and do not need to name it. However, when you need use an object multiple times, declare it in the heap, and then reference it by name in the route.

The following route shows an inline declaration for a handler. The handler is a router to route requests to separate route configurations:

```
{
  "handler": {
    "type": "Router"
  }
}
```

The following example shows a named router in the heap, and a handler references the router by its name:

```
{
  "handler": "My Router",
  "heap": [
    {
      "name": "My Router",
      "type": "Router"
    }
  ]
}
```

Notice that the heap takes an array. Because the heap holds all configuration objects at the same level, you can impose any hierarchy or order when referencing objects. Note that when you declare all objects in the heap and reference them by name, neither hierarchy nor ordering are obvious from the structure of the configuration file alone.

12.2.2. Setting Route Conditions

When a route has a condition, it can handle only requests that meet the condition. When a route has no condition, it can handle any request.

A condition can be based on almost any characteristic of the request, context, or OpenIG runtime environment. Conditions are defined using OpenIG expressions, as described in [Expressions\(5\)](#) in the *Configuration Reference*.

The following example shows a route condition that is met when the request path is `/login`:

```
"condition": "${matches(request.uri.path, '^/login')}}"
```

The following example shows a route condition that is met only by requests with `api.example.com` as the host portion of the URI:

```
"condition": "${request.uri.host == 'api.example.com'}"
```

The following example shows a route with no condition. This route accepts any request:

```
{
  "name": "default",
  "handler": {
    "type": "ClientHandler"
  }
}
```

Because routes define the conditions on which they accept a request, the router does not have to know about specific routes in advance. In other words, you can configure the router first and then add routes while OpenIG is running.

12.2.3. Configuring Route Names, IDs, and Filenames

Route filenames have the extension `.json`, in lowercase. A router scans the routes folder for files with the `.json` extension.

When a route has a `name` configuration object, the name is used by the router to order the routes lexicographically in the configuration. If the route is not named, the route ID is used.

When you add a route manually to the routes folder or add it through Common REST, the route ID is the same as the filename of the route you add. When you add a route through OpenIG Studio, you can edit the default route ID.

12.3. Creating and Editing Routes Through Common REST

Note

When OpenIG is in production mode, you cannot manage, list, or even read routes through Common REST. For more information, see Section 3.8, "Making the Configuration Immutable".

Note

If an OpenAM policy agent is configured in the same container as OpenIG, by default the policy agent intercepts requests to manage routes. When you try to add a route through Common REST, the policy agent redirects the request to OpenAM and the route is not added.

To override this behavior, add the URL pattern `/openig/api/*` to the list of not-enforced URI in the policy agent profile. For information, see Procedure 5.4, "To Create a Policy Agent Profile in OpenAM".

Through Common REST, you can read, add, delete, and edit routes on OpenIG without manually accessing the file system. You can also list the routes in the order that they are loaded in the configuration, and set fields to filter the information about the routes.

The following examples show some ways to manage routes through Common REST. For more information, see Section 5, "About ForgeRock Common REST" in the *Configuration Reference*.

Procedure 12.1. To Manage Routes Through Common REST

Before you start, prepare OpenIG as described in Chapter 2, "Getting Started".

1. Add the following route to the OpenIG configuration as `$HOME/.openig/config/routes/00-crest.json`.

On Windows, add the route as `%appdata%\OpenIG\config\routes\00-crest.json`.

```
{
  "name": "crest",
  "handler": {
    "type": "StaticResponseHandler",
    "config": {
      "status": 200,
      "reason": "OK",
      "entity": "Hello, world!"
    }
  },
  "condition": "${matches(request.uri.path, '^/crest')}"
}
```

To check that the route is working, access the route on: `http://openig.example.com:8080/crest`.

2. To read a route through Common REST:
 - Enter the following command in a terminal window:

```
$ http GET http://openig.example.com:8080/openig/api/system/objects/_router/routes/00-crest
```

The route is displayed. Note that the route's `_id` is displayed in the JSON of the route.

3. To add a route through Common REST:

- Move `$HOME/.openig/config/routes/00-crest.json` to `/tmp/00-crest.json`.
- Check in `$HOME/.openig/logs/route-system.log` that the route has been removed from the configuration. To double check, access `http://openig.example.com:8080/crest`. You should get an HTTP 404.
- Enter the following command in a terminal window:

```
$ http PUT http://openig.example.com:8080/openig/api/system/objects/_router/routes/00-crest < /tmp/00-crest.json
```

This command posts the file in `/tmp/00-crest.json` to the `routes` directory.

- Check in `$HOME/.openig/logs/route-system.log` that the route has been added to configuration. To double-check, access `http://openig.example.com:8080/crest`. You should see the "Hello, world!" message.

4. To edit a route through CREST:

- Edit `/tmp/00-crest.json` to change the message displayed by the response handler in the route.
- Enter the following command in a terminal window:

```
$ http PUT http://openig.example.com:8080/openig/api/system/objects/_router/routes/00-crest If-Match:* < /tmp/00-crest.json
```

This command deploys the route with the new configuration. Because the changes are persisted into the configuration, the existing `$HOME/.openig/config/routes/00-crest.json` is replaced with the edited version in `/tmp/00-crest.json`.

- Check in `$HOME/.openig/logs/route-system.log`, that the route has been updated. To double-check, access `http://openig.example.com:8080/crest` to confirm that the displayed message has changed.

5. To delete a route through CREST:

- Enter the following command in a terminal window:

```
$ http DELETE http://openig.example.com:8080/openig/api/system/objects/_router/routes/00-crest
```

- Check in `$HOME/.openig/logs/route-system.log` that the route has been removed from the configuration. To double-check, access `http://openig.example.com:8080/crest`. You should get an HTTP 404.

6. To list the routes deployed on the router, in the order that they are tried by the router:

- Enter the following command in a terminal window:

```
$ http "http://openig.example.com:8080/openig/api/system/objects/_router/routes?_queryFilter=true"
```

The list of loaded routes is displayed.

12.4. Creating Routes Through OpenIG Studio

Note

When OpenIG is in production mode, OpenIG Studio is effectively disabled. For more information, see Section 3.8, "Making the Configuration Immutable".

OpenIG Studio is a user interface to configure and deploy routes in OpenIG. You can use OpenIG Studio to create routes for tasks such as authenticating users, and authorizing access to APIs, throttling the rate of requests to protected applications, capturing messages, and collecting statistics. New features will be added as OpenIG Studio evolves.

When OpenIG is installed and running as described in this guide, access OpenIG Studio on `http://openig.example.com:8080/openig/studio`.

For examples of how to use OpenIG Studio to configure routes, see the following sections of this guide:

- To configure OpenIG to enforce OpenAM policy decisions, see Procedure 6.3, "To Set Up OpenIG as a PEP".
- To configure OpenIG as a relying party for OpenID Connect 1.0, see Section 9.5, "Setting Up OpenIG As a Relying Party".
- To configure a simple throttling filter, see Procedure 16.1, "To Configure a Simple Throttling Filter".

Important

OpenIG Studio has the current limitations:

- OpenIG Studio deploys and undeploys routes through a main router named `_router`, which is the name of the main router in the default configuration. If you use a custom `config.json`, make sure that it contains a main router named `_router`.
- From OpenIG Studio, you can view and edit only routes that have been created and edited exclusively in OpenIG Studio. You cannot use OpenIG Studio to view or edit other routes in the `$HOME/.openig/config/routes` folder.
- After you create and deploy a route in OpenIG Studio, you can continue to edit it in OpenIG Studio and deploy the changed route. However, after you edit the route in the `$HOME/.openig/config/routes` folder, you can't import the changed configuration into OpenIG Studio. You must make any further edits in the `$HOME/.openig/config/routes` folder.

12.5. Preventing the Reload of Routes

To prevent routes from being reloaded after startup, stop OpenIG, edit the router `scanInterval`, and restart OpenIG. When the interval is set to `disabled`, routes are loaded only at startup:

```
{
  "name": "Router",
  "type": "Router",
  "config": {
    "scanInterval": "disabled"
  }
}
```

The following example changes the location where the router looks for the routes:

```
{
  "name": "Router",
  "type": "Router",
  "config": {
    "directory": "/path/to/safe/routes",
    "scanInterval": "disabled"
  }
}
```

12.6. Accessing Reserved Routes

OpenIG uses an `ApiProtectionFilter` to protect the reserved routes. By default, the filter allows access to reserved routes only from the loopback address. To override this behavior, declare a custom `ApiProtectionFilter` in the top-level heap. For an example, see the CORS filter described in Procedure 11.5, "To Set Up OpenIG As a UMA Resource Server".

Chapter 13

Configuration Templates

This chapter contains template routes for common configurations.

Before you use one of the templates here, install and configure OpenIG with a router and default route as described in Chapter 2, "Getting Started".

Next, take one of the templates and then modify it to suit your deployment. Read the summary of each template to find the right match for your application.

When you move to use OpenIG in production, be sure to turn off DEBUG level logging, and to deactivate `CaptureDecorator` use to avoid filling up disk space. Also consider locking down the `Router` configuration.

13.1. Proxy and Capture

If you installed and configured OpenIG with a router and default route as described in Chapter 2, "Getting Started", then you already proxy and capture both the application requests coming in and the server responses going out.

The route shown in Example 13.1, "Proxy and Capture" uses a `DispatchHandler` to change the scheme to HTTPS on login. To use this template change the baseURI settings to match those of the target application.

When connecting to the protected application over HTTPS, the `ClientHandler` must be configured to trust the application's public key server certificate. If the certificate was signed by a well-known Certificate Authority, then there should be no further configuration to do. Otherwise, use a `ClientHandler` that references a truststore holding the certificate.

Example 13.1. Proxy and Capture

```
{
  "handler": {
    "type": "DispatchHandler",
    "config": {
      "bindings": [
        {
          "condition": "${request.uri.path == '/login'}",
          "handler": "ClientHandler",
          "comment": "Must be able to trust the server cert for HTTPS",
        }
      ]
    }
  }
}
```

```

    "baseURI": "https://app.example.com:8444"
  },
  {
    "condition": "${request.uri.scheme == 'http'}",
    "handler": "ClientHandler",
    "baseURI": "http://app.example.com:8081"
  },
  {
    "handler": "ClientHandler",
    "baseURI": "https://app.example.com:8444"
  }
]
},
"capture": "all",
"condition": "${matches(request.uri.query, 'demo=capture')}"
}

```

To try this example with the sample application, save the file as `$HOME/.openig/config/routes/20-capture.json`, and browse to `http://openig.example.com:8080/login?demo=capture`.

To use this as a default route with a real application, remove the route-level condition on the handler that specifies a `demo` query string parameter.

13.2. Simple Login Form

The route in Example 13.2, "Simple Login Form" logs the user into the target application with hard-coded user name and password. The route intercepts the login page request and replaces it with the login form. Adapt the `uri`, `form`, and `baseURI` settings as necessary.

Example 13.2. Simple Login Form

```

{
  "heap": [
    {
      "name": "ClientHandler",
      "type": "ClientHandler",
      "comment": "Testing only: blindly trust the server cert for HTTPS.",
      "config": {
        "trustManager": {
          "type": "TrustAllManager"
        }
      }
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {

```



```

    "type": "PasswordReplayFilter",
    "config": {
      "loginPage": "${request.uri.path == '/login'}",
      "request": {
        "method": "POST",
        "uri": "https://app.example.com:8444/login",
        "form": {
          "username": [
            "MY_USERNAME"
          ],
          "password": [
            "MY_PASSWORD"
          ]
        }
      }
    },
    "handler": "ClientHandler"
  },
  "condition": "${matches(request.uri.query, 'demo=simple')}"
}

```

The parameters in the `PasswordReplayFilter` form, `MY_USERNAME` and `MY_PASSWORD`, can use strings or expressions. When connecting to the protected application over HTTPS, the `ClientHandler` must be configured to trust the application's public key server certificate.

To try this example with the sample application, save the file as `$HOME/.openig/config/routes/21-simple.json`, replace `MY_USERNAME` with `demo` and `MY_PASSWORD` with `changeit`, and browse to `http://openig.example.com:8080/login?demo=simple`.

To use this as a default route with a real application, use a `ClientHandler` that does not blindly trust the server certificate, and remove the route-level condition on the handler that specifies a `demo` query string parameter.

13.3. Login Form With Cookie From Login Page

Some applications expect a cookie from the login page to be sent in the login request form. OpenIG can manage the cookies. The route in Example 13.3, "Login Form With Cookie From Login Page" allows the login page request to go through to the target, and manages the cookies set in the response rather than passing the cookie through to the browser.

Example 13.3. Login Form With Cookie From Login Page

```

{
  "handler": {
    "type": "Chain",

```

```
"config": {
  "filters": [
    {
      "type": "PasswordReplayFilter",
      "config": {
        "loginPage": "${request.uri.path == '/login'}",
        "request": {
          "method": "POST",
          "uri": "https://app.example.com:8444/login",
          "form": {
            "username": [
              "MY_USERNAME"
            ],
            "password": [
              "MY_PASSWORD"
            ]
          }
        }
      }
    },
    {
      "type": "CookieFilter"
    }
  ],
  "handler": "ClientHandler"
},
"condition": "${matches(request.uri.query, 'demo=cookie')}}"
}
```

The parameters in the `PasswordReplayFilter` form, `MY_USERNAME` and `MY_PASSWORD`, can use strings or expressions. A `CookieFilter` with no specified configuration manages all cookies that are set by the protected application. When connecting to the protected application over HTTPS, the `ClientHandler` must be configured to trust the application's public key server certificate.

To try this example with the sample application, save the file as `$HOME/.openig/config/routes/22-cookie.json`, replace `MY_USERNAME` with `kramer` and `MY_PASSWORD` with `newman`, and browse to `http://openig.example.com:8080/login?demo=cookie`.

To use this as a default route with a real application, remove the route-level condition on the handler that specifies a `demo` query string parameter.

13.4. Login Form With Password Replay and Cookie Filters

The route in Example 13.4, "Login Form With Password Replay and Cookie Filters" works with an application that returns the login page when the user tries to access a page without a valid session. This route shows how to use a `PasswordReplayFilter` to find the login page with a pattern that matches a mock OpenAM Classic UI page.

Note

The route uses a `CookieFilter` to manage cookies, ensuring that cookies from the protected application are included with the appropriate requests. The side effect of OpenIG managing cookies is none of the cookies are sent to the browser, but are managed locally by OpenIG.

Example 13.4. Login Form With Password Replay and Cookie Filters

```
{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "PasswordReplayFilter",
          "config": {
            "loginPageContentMarker": "OpenAM\\|s\\|(Login\\|)",
            "request": {
              "comments": [
                "An example based on OpenAM classic UI: ",
                "uri is for the OpenAM login page; ",
                "IDToken1 is the username field; ",
                "IDToken2 is the password field; ",
                "host takes the OpenAM FQDN:port.",
                "The sample app simulates OpenAM."
              ],
              "method": "POST",
              "uri": "http://app.example.com:8081/openam/UI/Login",
              "form": {
                "IDToken0": [
                  ""
                ],
                "IDToken1": [
                  "demo"
                ],
                "IDToken2": [
                  "changeit"
                ],
                "IDButton": [
                  "Log+In"
                ],
                "encoded": [
                  "false"
                ]
              },
              "headers": {
                "host": [
                  "app.example.com:8081"
                ]
              }
            }
          }
        },
        {
          "type": "CookieFilter"
        }
      ]
    }
  }
}
```

```

    }
  ],
  "handler": "ClientHandler"
}
},
"condition": "${matches(request.uri.query, 'demo=classic')}"
}

```

The parameters in the `PasswordReplayFilter` form can use strings or expressions.

To try this example with the sample application, save the file as `$HOME/.openig/config/routes/23-classic.json`, and use the `curl` command to check that it works as in the following example, which shows that the `CookieFilter` has removed cookies from the response except for the session cookie added by the container:

```

$ curl -D- http://openig.example.com:8080/login?demo=classic
HTTP/1.1 200
OK
...
Set-Cookie: JSESSIONID=1gwp5h0ugkciv1g200c9hid4sp;Path=/
Content-Length: 15
Content-Type: text/plain;charset=ISO-8859
-1
...

Welcome, demo!

```

To use this as a default route with a real application, remove the route-level condition on the handler that specifies a `demo` query string parameter, and adjust the `PasswordReplayFilter` as necessary.

13.5. Login Which Requires a Hidden Value From the Login Page

Some applications call for extracting a hidden value from the login page and including the value in the login form POSTed to the target application. The route in Example 13.5, "Login Which Requires a Hidden Value From the Login Page" extracts a hidden value from the login page, and posts a static form including the hidden value.

Example 13.5. Login Which Requires a Hidden Value From the Login Page

```

{
  "heap": [
    {
      "name": "ClientHandler",
      "type": "ClientHandler",
      "comment": "Testing only: blindly trust the server cert for HTTPS.",
      "config": {

```

```

    "trustManager": {
      "type": "TrustAllManager"
    }
  },
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "PasswordReplayFilter",
          "config": {
            "loginPage": "${request.uri.path == '/login'}",
            "loginPageExtractions": [
              {
                "name": "hidden",
                "pattern": "loginToken\\s+value=\\\"(.*)\\\"\"
              }
            ],
            "request": {
              "method": "POST",
              "uri": "https://app.example.com:8444/login",
              "form": {
                "username": [
                  "MY_USERNAME"
                ],
                "password": [
                  "MY_PASSWORD"
                ],
                "hiddenValue": [
                  "${attributes.extracted.hidden}"
                ]
              }
            }
          }
        ]
      }
    },
    "handler": "ClientHandler"
  },
  "condition": "${matches(request.uri.query, 'demo=hidden')}"
}

```

The parameters in the `PasswordReplayFilter` form, `MY_USERNAME` and `MY_PASSWORD`, can have string values, and they can also use expressions. When connecting to the protected application over HTTPS, the `ClientHandler` must be configured to trust the application's public key server certificate.

To try this example with the sample application, save the file as `$HOME/.openig/config/routes/24-hidden.json`, replace `MY_USERNAME` with `scarter` and `MY_PASSWORD` with `sprain`, and browse to `http://openig.example.com:8080/login?demo=hidden`.

To use this as a default route with a real application, use a `ClientHandler` that does not blindly trust the server certificate, and remove the route-level condition on the handler that specifies a `demo` query string parameter.

13.6. HTTP and HTTPS Application

The route in Example 13.6, "HTTP and HTTPS Application" proxies traffic to an application with both HTTP and HTTPS ports. The application uses HTTPS for authentication and HTTP for the general application features. Assuming all login requests are made over HTTPS, you must add the login filters and handlers to the chain.

Example 13.6. HTTP and HTTPS Application

```
{
  "handler": {
    "type": "DispatchHandler",
    "config": {
      "bindings": [
        {
          "condition": "${request.uri.scheme == 'http'}",
          "handler": "ClientHandler",
          "baseURI": "http://app.example.com:8081"
        },
        {
          "condition": "${request.uri.path == '/login'}",
          "handler": {
            "type": "Chain",
            "config": {
              "comment": "Add one or more filters to handle login.",
              "filters": [],
              "handler": "ClientHandler"
            }
          },
          "baseURI": "https://app.example.com:8444"
        },
        {
          "handler": "ClientHandler",
          "baseURI": "https://app.example.com:8444"
        }
      ]
    },
    "condition": "${matches(request.uri.query, 'demo=https')}"
  }
}
```

When connecting to the protected application over HTTPS, the `ClientHandler` must be configured to trust the application's public key server certificate.

To try this example with the sample application, save the file as `$HOME/.openig/config/routes/25-https.json`, and browse to `http://openig.example.com:8080/login?demo=https`.

To use this as a default route with a real application, remove the route-level condition on the handler that specifies a `demo` query string parameter.

13.7. OpenAM Integration With Headers

The route in Example 13.7, "OpenAM Integration With Headers" logs the user into the target application using the headers such as those passed in from an OpenAM policy agent. If the header passed in contains only a user name or subject and requires a lookup to an external data source, you must add an attribute filter to the chain to retrieve the credentials.

Example 13.7. OpenAM Integration With Headers

```
{
  "heap": [
    {
      "name": "ClientHandler",
      "type": "ClientHandler",
      "comment": "Testing only: blindly trust the server cert for HTTPS.",
      "config": {
        "trustManager": {
          "type": "TrustAllManager"
        }
      }
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "PasswordReplayFilter",
          "config": {
            "loginPage": "${request.uri.path == '/login'}",
            "request": {
              "method": "POST",
              "uri": "https://app.example.com:8444/login",
              "form": {
                "username": [
                  "${request.headers['username']}[0]"
                ],
                "password": [
                  "${request.headers['password']}[0]"
                ]
              }
            }
          }
        }
      ]
    },
    "handler": "ClientHandler"
  }
},
"condition": "${matches(request.uri.query, 'demo=headers')}"
}
```

When connecting to the protected application over HTTPS, the `ClientHandler` must be configured to trust the application's public key server certificate.

To try this example with the sample application, save the file as `$HOME/.openig/config/routes/26-headers.json`, and use the `curl` command to simulate the headers being passed in from an OpenAM policy agent as in the following example:

```
$ curl \
  --header "username: kvaughan" \
  --header "password: bribery" \
  http://openig.example.com:8080/login?demo=headers
...
<title id="welcome">Howdy, kvaughan</
title>
...
```

To use this as a default route with a real application, use a `ClientHandler` that does not blindly trust the server certificate, and remove the route-level condition on the handler that specifies a `demo` query string parameter.

13.8. Microsoft Online Outlook Web Access

The route in Example 13.8, "Microsoft Online Outlook Web Access" logs the user into Microsoft Online Outlook Web Access (OWA). The example shows how you would use OpenIG and the OpenAM password capture feature to integrate with OWA. Follow the example in Chapter 5, "Getting Login Credentials From Access Management", and substitute this template as a replacement for the default route.

Example 13.8. Microsoft Online Outlook Web Access

```
{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "PasswordReplayFilter",
          "config": {
            "loginPage": "${request.uri.path == '/owa/auth/logon.aspx'}",
            "headerDecryption": {
              "algorithm": "DES/ECB/NoPadding",
              "key": "DESKEY",
              "keyType": "DES",
              "charSet": "utf-8",
              "headers": [
                "password"
              ]
            },
            "request": {
```

To try this example, save the file as `$HOME/.openig/config/routes/27-owa.json`. Change `DESKEY` to the actual key value that you generated when following the instructions in Procedure 5.5, "To Configure Password Capture in OpenAM".

To use this as a default route with a real application, remove the route-level condition on the handler that specifies a `demo` query string parameter.

Chapter 14

Extending the ForgeRock Identity Gateway

This chapter describes how to extend OpenIG, taking you through the steps to:

- Write scripts to create custom filters and handlers
- Plug additional Java libraries into OpenIG for further customization

To extend filter and handler functionality, OpenIG supports the Groovy dynamic scripting language through the use of `ScriptableFilter` and `ScriptableHandler` objects.

For when you can't achieve complex server interactions or intensive data transformations with scripts or existing handlers, filters, or expressions, OpenIG allows you to develop custom extensions in Java and provide them in additional libraries that you build into OpenIG. The libraries allow you to develop custom extensions to OpenIG.

Important

When you are writing scripts or Java extensions, never use a `Promise` blocking method, such as `get()`, `getOrThrow()`, or `getOrThrowUninterruptibly()`, to obtain the response.

A promise represents the result of an asynchronous operation. Therefore, using a blocking method to wait for the result can cause deadlocks and/or race issues.

14.1. About Scripting

Scriptable filters and handlers are added to the configuration in the same way as standard filters and handlers. Each takes as its configuration the script's Internet media type and either a source script included in the JSON configuration, or a file script that OpenIG reads from a file. The configuration can optionally supply arguments to the script.

The following example defines a `ScriptableFilter`, written in the Groovy language, and stored in a file named `$HOME/.openig/scripts/groovy/SimpleFormLogin.groovy` (`%appdata%\OpenIG\scripts\groovy\SimpleFormLogin.groovy` on Windows):

```
{
  "name": "SimpleFormLogin",
  "type": "ScriptableFilter",
  "config": {
    "type": "application/x-groovy",
    "file": "SimpleFormLogin.groovy"
  }
}
```

Relative paths in the file field depend on how OpenIG is installed. If OpenIG is installed in an application server, then paths for Groovy scripts are relative to `$HOME/.openig/scripts/groovy`.

This base location `$HOME/.openig/scripts/groovy` is on the classpath when the scripts are executed. If therefore some Groovy scripts are not in the default package, but instead have their own package names, they belong in the directory corresponding to their package name. For example, a script in package `com.example.groovy` belongs under `$HOME/.openig/scripts/groovy/com/example/groovy/`.

OpenIG provides several global variables to scripts at runtime. As well as having access to Groovy's built-in functionality, scripts can access the request and the context, store variables across executions, write messages to logs, make requests to a web service or to an LDAP directory service, and access responses returned in promise callback methods. For information about scripting in OpenIG, see `ScriptableFilter(5)` in the *Configuration Reference* and `ScriptableHandler(5)` in the *Configuration Reference*.

Before trying the scripts shown in this chapter, first install and configure OpenIG as described in Chapter 2, "Getting Started".

When developing and debugging your scripts, consider configuring a capture decorator to log requests, responses, and context data in JSON form. You can then turn off capturing when you move to production. For details, see `CaptureDecorator(5)` in the *Configuration Reference*.

14.2. Scripting Dispatch

In order to route requests, especially when the conditions are complicated, you can use a `ScriptableHandler` instead of a `DispatchHandler` as described in `DispatchHandler(5)` in the *Configuration Reference*.

The following script demonstrates a simple dispatch handler:

```
/*
 * Copyright 2014-2017 ForgeRock AS. All Rights Reserved
 *
 * Use of this code requires a commercial software license with ForgeRock AS.
 * or with one of its affiliates. All use shall be exclusively subject
 * to such license between the licensee and ForgeRock AS.
 */

/*
```

```

* This simplistic dispatcher matches the path part of the HTTP request.
* If the path is /mylogin, it checks Username and Password headers,
* accepting bjensen:hifalutin, and returning HTTP 403 Forbidden to others.
* Otherwise it returns HTTP 401 Unauthorized.
*/

// Rather than return a Promise of a response from an external source,
// this script returns the response itself.
response = new Response(Status.OK);

switch (request.uri.path) {
  case "/mylogin":
    if (request.headers.Username.values[0] == "bjensen" &&
        request.headers.Password.values[0] == "hifalutin") {
      response.status = Status.OK
      response.entity = "<html><p>Welcome back, Babs!</p></html>"
    } else {
      response.status = Status.FORBIDDEN
      response.entity = "<html><p>Authorization required</p></html>"
    }
    break
  default:
    response.status = Status.UNAUTHORIZED
    response.entity = "<html><p>Please <a href='./mylogin'>log in</a>.</p></html>"
    break
}

// Return the locally created response, no need to wrap it into a Promise
return response

```

To try this handler, save the script as `$HOME/.openig/scripts/groovy/DispatchHandler.groovy` (`%appdata%\OpenIG\scripts\groovy\DispatchHandler.groovy` on Windows).

Next, add the following route to your configuration as `$HOME/.openig/config/routes/98-dispatch.json` (`%appdata%\OpenIG\config\routes\98-dispatch.json` on Windows):

```

{
  "heap": [
    {
      "name": "DispatchHandler",
      "type": "DispatchHandler",
      "config": {
        "bindings": [{
          "condition": "${matches(request.uri.path, '/mylogin')}",
          "handler": {

```

```

    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "HeaderFilter",
          "config": {
            "messageType": "REQUEST",
            "add": {
              "Username": [
                "bjensen"
              ],
              "Password": [
                "hifalutin"
              ]
            }
          }
        },
        {
          "handler": "Dispatcher"
        }
      ],
      "handler": "Dispatcher"
    },
    {
      "handler": "Dispatcher",
      "condition": "${matches(request.uri.path, '/dispatch')}"
    }
  ],
  {
    "name": "Dispatcher",
    "type": "ScriptableHandler",
    "config": {
      "type": "application/x-groovy",
      "file": "DispatchHandler.groovy"
    }
  },
  {
    "handler": "DispatchHandler",
    "condition": "${matches(request.uri.path, '^/dispatch') or matches(request.uri.path, '^/mylogin')}"
  }
}

```

The route sets up the headers required by the script when the user logs in.

To try it out, browse to <http://openig.example.com:8080/dispatch>.

The response from the script says, "Please log in." When you click the log in link, the `HeaderFilter` sets `Username` and `Password` headers in the request, and passes the request to the script.

The script then responds, `Welcome back, Babs!`

14.3. Scripting HTTP Basic Authentication

HTTP Basic authentication calls for the user agent such as a browser to send a user name and password to the server in an `Authorization` header. HTTP Basic authentication relies on an encrypted

connection to protect the user name and password credentials, which are base64-encoded in the **Authorization** header, not encrypted.

The following script, for use in a **ScriptableFilter**, adds an **Authorization** header based on a username and password combination:

```

/*
 * Copyright 2014-2017 ForgeRock AS. All Rights Reserved
 *
 * Use of this code requires a commercial software license with ForgeRock AS.
 * or with one of its affiliates. All use shall be exclusively subject
 * to such license between the licensee and ForgeRock AS.
 */

/*
 * Perform basic authentication with the user name and password
 * that are supplied using a configuration like the following:
 *
 * {
 *   "name": "BasicAuth",
 *   "type": "ScriptableFilter",
 *   "config": {
 *     "type": "application/x-groovy",
 *     "file": "BasicAuthFilter.groovy",
 *     "args": {
 *       "username": "bjensen",
 *       "password": "hifalutin"
 *     }
 *   }
 * }
 */

def userPass = username + ":" + password
def base64UserPass = userPass.getBytes().encodeBase64()
request.headers.add("Authorization", "Basic ${base64UserPass}" as String)

// Credentials are only base64-encoded, not encrypted: Set scheme to HTTPS.

/*
 * When connecting over HTTPS, by default the client tries to trust the server.
 * If the server has no certificate
 * or has a self-signed certificate unknown to the client,
 * then the most likely result is an SSLPeerUnverifiedException.
 *
 * To avoid an SSLPeerUnverifiedException,
 * set up HTTPS correctly on the server.
 * Either use a server certificate signed by a well-known CA,
 * or set up the gateway to trust the server certificate.
 */
request.uri.scheme = "https"

// Calls the next Handler and returns a Promise of the Response.
// The Response can be handled with asynchronous Promise callbacks.
next.handle(context, request)

```

To try this filter, save the script as `$HOME/.openig/scripts/groovy/BasicAuthFilter.groovy` (%appdata%\openIG\scripts\groovy\BasicAuthFilter.groovy on Windows).

Next, add the following route to your configuration as `$HOME/.openig/config/routes/09-basic.json` (`%appdata%\OpenIG\config\routes\09-basic.json` on Windows):

```
{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ScriptableFilter",
          "config": {
            "type": "application/x-groovy",
            "file": "BasicAuthFilter.groovy",
            "args": {
              "username": "bjensen",
              "password": "hifalutin"
            }
          },
          "capture": "filtered_request"
        }
      ],
      "handler": {
        "type": "StaticResponseHandler",
        "config": {
          "status": 200,
          "reason": "OK",
          "entity": "Hello, Babs!"
        }
      }
    }
  },
  "condition": "${matches(request.uri.path, '^/basic')}"
}
```

When the request path matches `/basic` the route calls the `Chain`, which runs the `ScriptableFilter`. The capture setting captures the request as updated by the `ScriptableFilter`. Finally, OpenIG returns a static page.

To try it out, browse to `http://openig.example.com:8080/basic`.

The captured request in the console log shows that the scheme is now HTTPS, and that the `Authorization` header is set for HTTP Basic:

```
GET https://openig.example.com:8080/basic HTTP/1.1
Authorization: Basic YmplbnNlbjpoaWZhbHV0aW4=
```

14.4. Scripting LDAP Authentication

Many organizations use an LDAP directory service to store user profiles including authentication credentials. The LDAP directory service securely stores user passwords in a highly-available, central service capable of handling thousands of authentications per second.

The following script, for use in a [ScriptableFilter](#), performs simple authentication against an LDAP server based on request form fields `username` and `password`:

```

/*
 * Copyright 2014-2017 ForgeRock AS. All Rights Reserved
 *
 * Use of this code requires a commercial software license with ForgeRock AS.
 * or with one of its affiliates. All use shall be exclusively subject
 * to such license between the licensee and ForgeRock AS.
 */

import org.forgerock.ldap.*

/*
 * Perform LDAP authentication based on user credentials from a form.
 *
 * If LDAP authentication succeeds, then return a promise to handle the response.
 * If there is a failure, produce an error response and return it.
 */

username = request.form?.username[0]
password = request.form?.password[0]

// For testing purposes, the LDAP host and port are provided in the context's attributes.
// Edit as needed to match your directory service.
host = attributes.ldapHost ?: "localhost"
port = attributes.ldapPort ?: 1389

client = ldap.connect(host, port as Integer)
try {

    // Assume the username is an exact match of either
    // the user ID, the email address, or the user's full name.
    filter = "(|(uid=%s)(mail=%s)(cn=%s))"

    user = client.searchSingleEntry(
        "ou=people,dc=example,dc=com",
        ldap.scope.sub,
        ldap.filter(filter, username, username, username))

    client.bind(user.name as String, password?.toCharArray())

    // Authentication succeeded.

    // Set a header (or whatever else you want to do here).
    request.headers.add("Ldap-User-Dn", user.name.toString())

    // Most LDAP attributes are multi-valued.
    // When you read multi-valued attributes, use the parse() method,
    // with an AttributeParser method
    // that specifies the type of object to return.
    attributes.cn = user.cn?.parse().asSetOfString()

    // When you write attribute values, set them directly.
    user.description = "New description set by my script"

    // Here is how you might read a single value of a multi-valued attribute:

```

```

attributes.description = user.description?.parse().asString()

// Call the next handler. This returns when the request has been handled.
return next.handle(context, request)

} catch (AuthenticationException e) {

    // LDAP authentication failed, so fail the response with
    // HTTP status code 403 Forbidden.

    response = new Response(Status.FORBIDDEN)
    response.entity = "<html><p>Authentication failed: " + e.message + "</p></html>"

} catch (Exception e) {

    // Something other than authentication failed on the server side,
    // so fail the response with HTTP 500 Internal Server Error.

    response = new Response(Status.INTERNAL_SERVER_ERROR)
    response.entity = "<html><p>Server error: " + e.message + "</p></html>"

} finally {
    client.close()
}

// Return the locally created response, no need to wrap it into a Promise
return response

```

For the list of methods to specify which type of objects to return, see the OpenDJ LDAP SDK Javadoc for [AttributeParser](#).

To try the LDAP authentication script, follow these steps:

1. Install an LDAP directory server such as [ForgeRock Directory Services](#).

Either import some sample users who can authenticate over LDAP, or generate sample users at installation time.

2. Save the script as `$HOME/.openig/scripts/groovy/LdapAuthFilter.groovy` (`%appdata%\OpenIG\scripts\groovy\LdapAuthFilter.groovy` on Windows).

If the directory server installation does not match the assumptions made in the script, adjust the script to use the correct settings for your installation.

3. Add the following route to your configuration as `$HOME/.openig/config/routes/10-ldap.json` (`%appdata%\OpenIG\config\routes\10-ldap.json` on Windows):

```

{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ScriptableFilter",
          "config": {

```

```

        "type": "application/x-groovy",
        "file": "LdapAuthFilter.groovy"
    }
},
"handler": {
    "type": "ScriptableHandler",
    "config": {
        "type": "application/x-groovy",
        "source": [
            "dn = request.headers['Ldap-User-Dn'].values[0]",
            "entity = '<html><body><p>Ldap-User-Dn: ' + dn + '</p></body></html>'",
            "",
            "response = new Response(Status.OK)",
            "response.entity = entity",
            "return response"
        ]
    }
},
"condition": "${matches(request.uri.path, '^/ldap')}"
}

```

The route calls the `LdapAuthFilter.groovy` script to authenticate the user over LDAP. On successful authentication, it responds with the the bind DN.

To test the configuration, browse to a URL where query string parameters specify a valid username and password, such as `http://openig.example.com:8080/ldap?username=user.0&password=password`.

The response from the script shows the DN: `Ldap-User-Dn: uid=user.0,ou=People,dc=example,dc=com`.

14.5. Scripting SQL Queries

You can use a `ScriptableFilter` to look up information in a relational database and include the results in the request context.

The following filter looks up user credentials in a database given the user's email address, which is found in the form data of the request. The script then sets the credentials in headers, making sure the scheme is HTTPS to protect the request when it leaves OpenIG:

```

/*
 * Copyright 2014-2017 ForgeRock AS. All Rights Reserved
 *
 * Use of this code requires a commercial software license with ForgeRock AS.
 * or with one of its affiliates. All use shall be exclusively subject
 * to such license between the licensee and ForgeRock AS.
 */

/*
 * Look up user credentials in a relational database
 * based on the user's email address provided in the request form data,
 * and set the credentials in the request headers for the next handler.
 */

```

```

*/
def client = new SqlClient()
def credentials = client.getCredentials(request.form?.mail[0])
request.headers.add("Username", credentials.Username)
request.headers.add("Password", credentials.Password)

// The credentials are not protected in the headers, so use HTTPS.
request.uri.scheme = "https"

// Calls the next Handler and returns a Promise of the Response.
// The Response can be handled with asynchronous Promise callbacks.
next.handle(context, request)

```

The previous script demonstrates a `ScriptableFilter` that uses a `SqlClient` class defined in another script. The following code listing shows the `SqlClient` class:

```

/*
 * Copyright 2014-2017 ForgeRock AS. All Rights Reserved
 *
 * Use of this code requires a commercial software license with ForgeRock AS.
 * or with one of its affiliates. All use shall be exclusively subject
 * to such license between the licensee and ForgeRock AS.
 */

import groovy.sql.Sql

import javax.naming.InitialContext
import javax.sql.DataSource

/**
 * Access a database with a well-known structure,
 * in particular to get credentials given an email address.
 */
class SqlClient {

    // Get a DataSource from the container.
    InitialContext context = new InitialContext()
    DataSource dataSource = context.lookup("jdbc/forgerock") as DataSource
    def sql = new Sql(dataSource)

    // The expected table is laid out like the following.

    // Table USERS
    // -----
    // | USERNAME | PASSWORD | EMAIL | ... |
    // -----
    // | <username> | <passwd> | <mail@...> | ... |
    // -----

    String tableName = "USERS"
    String usernameColumn = "USERNAME"
    String passwordColumn = "PASSWORD"
    String mailColumn = "EMAIL"

    /**
     * Get the Username and Password given an email address.
     */
}

```

```

* @param mail Email address used to look up the credentials
* @return Username and Password from the database
*/
def getCredentials(mail) {
    def credentials = [:]
    def query = "SELECT " + usernameColumn + ", " + passwordColumn +
        " FROM " + tableName + " WHERE " + mailColumn + "='$mail';"

    sql.eachRow(query) {
        credentials.put("Username", it."$usernameColumn")
        credentials.put("Password", it."$passwordColumn")
    }
    return credentials
}
}

```

To try the script, follow these steps:

1. Follow the tutorial in Section 4.3, "Log in With Credentials From a Database".

When everything in that tutorial works, you know that OpenIG can connect to the database, look up users by email address, and successfully authenticate to the sample application.

2. Save the scripts as `$HOME/.openig/scripts/groovy/SqlAccessFilter.groovy` (`%appdata%\OpenIG\scripts\groovy\SqlAccessFilter.groovy` on Windows), and as `$HOME/.openig/scripts/groovy/SqlClient.groovy` (`%appdata%\OpenIG\scripts\groovy\SqlClient.groovy` on Windows).
3. Add the following route to your configuration as `$HOME/.openig/config/routes/l1-db.json` (`%appdata%\OpenIG\config\routes\l1-db.json` on Windows):

```

{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ScriptableFilter",
          "config": {
            "type": "application/x-groovy",
            "file": "SqlAccessFilter.groovy"
          }
        },
        {
          "type": "StaticRequestFilter",
          "config": {
            "method": "POST",
            "uri": "http://app.example.com:8081/login",
            "form": {
              "username": [
                "${request.headers['Username']}[0]"
              ],
              "password": [
                "${request.headers['Password']}[0]"
              ]
            }
          }
        }
      ]
    }
  }
}

```

```
    }  
    ],  
    "handler": "ClientHandler"  
  }  
},  
"condition": "${matches(request.uri.path, '^/db')}"  
}
```

The route calls the `ScriptableFilter` to look up credentials over SQL. It then uses calls a `StaticRequestFilter` to build a login request. Although the script sets the scheme to HTTPS, the `StaticRequestFilter` ignores that and resets the URI. This makes it easier to try the script without additional steps to set up HTTPS.

To try the configuration, browse to a URL where a query string parameter specifies a valid email address, such as `http://openig.example.com:8080/db?mail=george@example.com`.

If the lookup and authentication are successful, you see the profile page of the sample application.

14.6. Developing Custom Extensions

OpenIG includes a complete Java application programming interface to allow you to customize OpenIG to perform complex server interactions or intensive data transformations that you cannot achieve with scripts or the existing handlers, filters, and expressions described in *Expressions(5)* in the *Configuration Reference*.

14.6.1. Key Extension Points

Interface Stability: Evolving (For information, see Section A.2, "ForgeRock Product Interface Stability" in the *Configuration Reference*.)

The following interfaces are available:

Decorator

A `Decorator` adds new behavior to another object without changing the base type of the object.

When suggesting custom `Decorator` names, know that OpenIG reserves all field names that use only alphanumeric characters. To avoid clashes, use dots or dashes in your field names, such as `my-decorator`.

ExpressionPlugin

An `ExpressionPlugin` adds a node to the `Expression` context tree, alongside `env` (for environment variables), and `system` (for system properties). For example, the expression `${system['user.home']}` yields the home directory of the user running the application server for OpenIG.

In your `ExpressionPlugin`, the `getKey()` method returns the name of the node, and the `getObject()` method returns the unified expression language context object that contains the values needed to resolve the expression. The plugins for `env` and `system` return Map objects, for example.

When you add your own `ExpressionPlugin`, you must make it discoverable within your custom library. You do this by adding a services file named after the plugin interface, where the file contains the fully qualified class name of your plugin, under `META-INF/services/org.forgerock.openig.el.ExpressionPlugin` in the `.jar` file for your customizations. When you have more than one plugin, add one fully qualified class name per line. For details, see the reference documentation for the Java class `ServiceLoader`. If you build your project using Maven, then you can add this under the `src/main/resources` directory. As described in Section 14.6.4, "Embedding the Customization in OpenIG", you must add your custom libraries to the `WEB-INF/lib/` directory of the OpenIG `.war` file that you deploy.

Be sure to provide some documentation for OpenIG administrators on how your plugin extends expressions.

Filter

A `Filter` serves to process a request before handing it off to the next element in the chain, in a similar way to an interceptor programming model.

The `Filter` interface exposes a `filter()` method, which takes a `Context`, a `Request`, and the `Handler`, which is the next filter or handler to dispatch to. The `filter()` method returns a `Promise` that provides access to the `Response` with methods for dealing with both success and failure conditions.

A filter can elect not to pass the request to the next filter or handler, and instead handle the request itself. It can achieve this by merely avoiding a call to `next.handle(context, request)`, creating its own response object and returning that in the promise. The filter is also at liberty to replace a response with another of its own. A filter can exist in more than one chain, therefore should make no assumptions or correlations using the chain it is supplied. The only valid use of a chain by a filter is to call its `handle()` method to dispatch the request to the rest of the chain.

Handler

A `Handler` generates a response for a request.

The `Handler` interface exposes a `handle()` method, which takes a `Context`, and a `Request`. It processes the request and returns a `Promise` that provides access to the `Response` with methods for dealing with both success and failure conditions. A handler can elect to dispatch the request to another handler or chain.

14.6.2. Implementing a Customized Sample Filter

The `SampleFilter` class implements the `Filter` interface to set a header in the incoming request and in the outgoing response. The following sample filter adds an arbitrary header:

```

package org.forgerock.openig.doc;

import org.forgerock.http.Filter;
import org.forgerock.http.Handler;
import org.forgerock.http.protocol.Request;
import org.forgerock.http.protocol.Response;
import org.forgerock.openig.heap.GenericHeaplet;
import org.forgerock.openig.heap.HeapException;
import org.forgerock.services.context.Context;
import org.forgerock.util.promise.NeverThrowsException;
import org.forgerock.util.promise.Promise;
import org.forgerock.util.promise.ResultHandler;

/**
 * Filter to set a header in the incoming request and in the outgoing response.
 */
public class SampleFilter implements Filter {

    /** Header name. */
    String name;

    /** Header value. */
    String value;

    /**
     * Set a header in the incoming request and in the outgoing response.
     * A configuration example looks something like the following.
     *
     * <pre>
     * {
     *   "name": "SampleFilter",
     *   "type": "SampleFilter",
     *   "config": {
     *     "name": "X-Greeting",
     *     "value": "Hello world"
     *   }
     * }
     * </pre>
     *
     * @param context      Execution context.
     * @param request      HTTP Request.
     * @param next         Next filter or handler in the chain.
     * @return A {@code Promise} representing the response to be returned to the client.
     */
    @Override
    public Promise<Response, NeverThrowsException> filter(final Context context,
                                                         final Request request,
                                                         final Handler next) {

        // Set header in the request.
        request.getHeaders().put(name, value);

        // Pass to the next filter or handler in the chain.
        return next.handle(context, request)
            // When it has been successfully executed, execute the following callback
            .thenOnResult(new ResultHandler<Response>() {
                @Override
                public void handleResult(final Response response) {

```



```

        // Set header in the response.
        response.getHeaders().put(name, value);
    });
}

/**
 * Create and initialize the filter, based on the configuration.
 * The filter object is stored in the heap.
 */
public static class Heaplet extends GenericHeaplet {

    /**
     * Create the filter object in the heap,
     * setting the header name and value for the filter,
     * based on the configuration.
     *
     * @return The filter object.
     * @throws HeapException Failed to create the object.
     */
    @Override
    public Object create() throws HeapException {

        SampleFilter filter = new SampleFilter();
        filter.name = config.get("name").as(evaluatedWithHeapProperties()).required().asString();
        filter.value = config.get("value").as(evaluatedWithHeapProperties()).required().asString();

        return filter;
    }
}

```

When you set the sample filter type in the configuration, you need to provide the fully qualified class name, as in `"type": "org.forgerock.openig.doc.SampleFilter"`. You can however implement a class alias resolver to make it possible to use a short name instead, as in `"type": "SampleFilter"`:

```

package org.forgerock.openig.doc;

import org.forgerock.openig.alias.ClassAliasResolver;

import java.util.HashMap;
import java.util.Map;

/**
 * Allow use of short name aliases in configuration object types.
 *
 * This allows a configuration with {@code "type": "SampleFilter"}
 * instead of {@code "type": "org.forgerock.openig.doc.SampleFilter"}.
 */
public class SampleClassAliasResolver implements ClassAliasResolver {

    private static final Map<String, Class<?>> ALIASES =
        new HashMap<>();

    static {
        ALIASES.put("SampleFilter", SampleFilter.class);
    }
}

```

```
/**
 * Get the class for a short name alias.
 *
 * @param alias Short name alias.
 * @return      The class, or null if the alias is not defined.
 */
@Override
public Class<?> resolve(String alias) {
    return ALIASES.get(alias);
}
}
```

When you add your own resolver, you must make it discoverable within your custom library. You do this by adding a services file named after the class resolver interface, where the file contains the fully qualified class name of your resolver, under `META-INF/services/org.forgerock.openig.alias.ClassAliasResolver` in the `.jar` file for your customizations. When you have more than one resolver, add one fully qualified class name per line. If you build your project using Maven, then you can add this under the `src/main/resources` directory. The content of the file in this example is one line:

```
org.forgerock.openig.doc.SampleClassAliasResolver
```

The corresponding heap object configuration then looks as follows:

```
{
  "name": "SampleFilter",
  "type": "SampleFilter",
  "config": {
    "name": "X-Greeting",
    "value": "Hello world"
  }
}
```

14.6.3. Configuring the Heap Object for the Customization

Objects are added to the heap and supplied with configuration artifacts at initialization time. To be integrated with the configuration, a class must have an accompanying implementation of the `Heaplet` interface. The easiest and most common way of exposing the heaplet is to extend the `GenericHeaplet` class in a nested class of the class you want to create and initialize, overriding the heaplet's `create()` method.

Within the `create()` method, you can access the object's configuration through the `config` field.

14.6.4. Embedding the Customization in OpenIG

After building your customizations into a `.jar` file, you can include them in the OpenIG `.war` file for deployment. You do this by unpacking `IG-5.0.0.war`, including your `.jar` library in `WEB-INF/lib`, and then creating a new `.war` file.

For example, if your .jar file is in a project named `sample-filter`, and the development version is `1.0.0-SNAPSHOT`, you might include the file as in the following example:

```
$ mkdir root && cd root
$ jar -xf ~/Downloads/IG-5.0.0.war
$ cp ~/Documents/sample-filter/target/sample-filter-1.0.0-SNAPSHOT.jar WEB-INF/lib
$ jar -cf ../custom.war *
```

In this example, the resulting `custom.war` contains the custom sample filter. You can deploy the custom .war file as you would deploy `IG-5.0.0.war`.

Chapter 15

Auditing and Monitoring

For each route, OpenIG can collect statistics on the number of requests and responses passed through the route since startup, and on throughput time and response time. The information is exposed as JSON format monitoring resource over an HTTP endpoint.

You can add an audit service to a route, and the service can then publish messages to a consumer such as a CSV file, a relational database, or the Syslog facility. In this chapter, you will learn to:

- Enable monitoring for a route
- Read monitoring statistics for a route as a JSON format monitoring resource
- Add an audit service to a route to integrate with the ForgeRock common audit event framework, sometimes referred to as Common Audit

15.1. Monitoring Routes

To collect statistics for a route, set the route's `monitor` attribute to `true`. Alternatively, use an appropriate boolean expression to enable or disable monitoring with an environment variable or system property.

When monitoring is enabled, statistics for the route are exposed as a JSON format monitoring resource that you can access over a Common REST endpoint. For information about what information is monitored, see "The REST API for Monitoring" in the *Configuration Reference*.

Monitoring endpoints serve API descriptors at runtime. When you retrieve an API descriptor for a monitoring endpoint, a JSON that describes the API for the endpoint is returned. You can use the API descriptor with a tool such as [Swagger UI](#) to generate a web page that helps you to view and test the endpoint. For information, see Section 1.10, "Understanding OpenIG APIs With API Descriptors".

When you start OpenIG, or add or edit routes, the endpoints are written in `$HOME/.openig/logs/route-system.log`. The endpoints are constructed from the base URL of OpenIG, the object name of the main router configured in `config.json`, and the route ID. When a route contains a subrouter, its endpoint includes the main router and the subrouter.

The routes endpoint is defined by the presence and content of `config.json`, as follows:

- When `config.json` is not provided, the routes endpoint includes the name of the main router in the default configuration, `_router`.

- When `config.json` is provided with an unnamed main router, the routes endpoint includes the main router name `router-handler`.
- When `config.json` is provided with a named main router, the routes endpoint includes the provided name or the transformed, URL-friendly name.

Procedure 15.1. To Monitor a Route

Before you start, prepare OpenIG and the sample application as described in Chapter 2, "Getting Started".

1. Access OpenIG Studio on `http://openig.example.com:8080/openig/studio`, and select Protect an Application.
2. In the Create a route window, select Advanced Options and enter the following information, and then select Create route:
 - Base URI: `http://app.example.com:8081`
 - Condition: Path: `/home/monitor`
 - Name: `00-monitor`
3. Select Statistics, and then enable statistics.
4. In Percentiles, remove the default percentiles, add the percentiles `0.25`, `0.5`, and `0.75`, and then select Save.

These are example values. Alternatively, use the default values or other values.

5. On the top-right of the screen, select `# > Display`, and review the route. The following route should be displayed.

```
{
  "name": "00-monitor",
  "baseURI": "http://app.example.com:8081",
  "monitor": {
    "enabled": true,
    "percentiles": [
      0.25,
      0.5,
      0.75
    ]
  },
  "condition": "${matches(request.uri.path, '^/home/monitor')}",
  "handler": "ClientHandler"
}
```

Notice the following features of the new route:

- The route matches requests to `/home/monitor`.

- Each time the route is accessed, the ClientHandler passes the request to the sample application and OpenIG collects statistics.
6. Select Deploy to push the route to the OpenIG configuration.
You can check the `$HOME/.openig/config/routes` folder to see that the route is there.
 7. Access the route a few times at `http://openig.example.com:8080/home/monitor`.
 8. Go to the monitoring endpoint for the route, at `http://openig.example.com:8080/openig/api/system/objects/_router/routes/00-monitor/monitoring`.

The monitoring resource displays statistics for requests and responses, as in the following example:

```
{
  "requests": {
    "active": 0,
    "total": 3
  },
  "responseTime": {
    "mean": 0.108,
    "median": 0.109,
    "percentiles": {
      "0.25": 0.085,
      "0.5": 0.109,
      "0.75": 0.13
    },
    "standardDeviation": 0.018,
    "total": 0
  },
  "responses": {
    "clientError": 0,
    "errors": 0,
    "info": 0,
    "null": 0,
    "other": 0,
    "redirect": 0,
    "serverError": 0,
    "success": 3,
    "total": 3
  },
  "throughput": {
    "last15Minutes": 0.0,
    "last5Minutes": 0.0,
    "lastMinute": 0.0,
    "mean": 0.1
  }
}
```

15.2. Recording Audit Event Messages

The ForgeRock common audit framework is a platform-wide infrastructure to handle audit events by using common audit event handlers. The handlers record events by logging them into files, relational databases, or syslog.

OpenIG provides audit event handlers to write messages to the following formats:

- JSON audit event handler, to log audit topics to JSON files.

For information about the JSON audit event handler, see `JsonAuditEventHandler(5)` in the *Configuration Reference*.

- The Elasticsearch search and analytics engine.

For an example of how to record audit events in elasticsearch, see Procedure 15.3, "To Record Audit Events In Elasticsearch ". For information about how to download and install Elasticsearch, see the Elasticsearch *Getting Started* document.

- CSV files, with support for retention, rotation, and tamper-evident logs.

For an example of how to record audit events in a CSV file, see Procedure 15.2, "To Record Audit Events In a CSV File".

- Relational databases using JDBC.
- The UNIX system log (Syslog) facility.
- Java Message Service (JMS), to send asynchronous messages between clients.

For an example of how to record audit events with a JMS audit event handler, see Procedure 15.4, "To Record Audit Events With a JMS Audit Event Handler". For more information about JMS audit event handlers, see `JmsAuditEventHandler(5)` in the *Configuration Reference*.

Each audit event is identified by a unique transaction ID that can be communicated across products and recorded for each local event. By using the transaction ID, requests can be tracked as they traverse the platform, making it easier to monitor activity and to enrich reports.

Transaction IDs from other services in the ForgeRock platform are sent as `X-ForgeRock-TransactionId` header values.

By default, OpenIG does not trust transaction ID headers from client applications.

Note

If you trust transaction IDs sent by client applications, and want monitoring and reporting systems consuming the logs to allow correlation of requests as they traverse multiple servers, then set the boolean system property

`org.forgerock.http.TrustTransactionHeader` to `true` in the Java command to start the container where OpenIG runs.

To enable the audit framework for a route, you specify an audit service and configure an audit event handler. The following procedures describe how to record audit events in a CSV file and to the Elasticsearch search and analytics engine. For more information about recording audit events, see *Audit Framework* in the *Configuration Reference*.

Procedure 15.2. To Record Audit Events In a CSV File

Before you start, prepare OpenIG and the sample application as shown in Chapter 2, "Getting Started".

1. Add the following route to the OpenIG as `$HOME/.openig/config/routes/30-audit.json`.

On Windows, add the route as `%appdata%\OpenIG\config\routes\30-audit.json`.

```
{
  "handler": "ForgeRockClientHandler",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/audit')}",
  "auditService": {
    "type": "AuditService",
    "config": {
      "config": {},
      "event-handlers": [
        {
          "class": "org.forgerock.audit.handlers.csv.CsvAuditEventHandler",
          "config": {
            "name": "csv",
            "logDirectory": "/tmp/logs",
            "buffering": {
              "enabled": "true",
              "autoFlush": "true"
            },
            "topics": [
              "access"
            ]
          }
        }
      ]
    }
  }
}
```

The route calls an audit service configuration for publishing log messages to the CSV file, `/tmp/logs/access.csv`. When a request matches `audit`, audit events are logged to the CSV file.

The route uses the `ForgeRockClientHandler` as its handler, to send the `X-ForgeRock-TransactionId` header with its requests to external services.

2. Access the route on `http://openig.example.com:8080/home/audit`.

The home page of the sample application should be displayed and the file `/tmp/logs/access.csv` should be updated.

Procedure 15.3. To Record Audit Events In Elasticsearch

Before you start, make sure that Elasticsearch is installed and running. For Elasticsearch downloads and installation instructions, see the Elasticsearch *Getting Started* document. For information about configuring the Elasticsearch event handler, see `ElasticsearchAuditEventHandler(5)` in the *Configuration Reference*.

1. Add the following route to the OpenIG as `$HOME/.openig/config/routes/30-elasticsearch.json`.

On Windows, add the route as `%appdata%\OpenIG\config\routes\30-elasticsearch.json`.

```
{
  "MyCapture": "all",
  "auditService": {
    "name": "audit-service",
    "type": "AuditService",
    "config": {
      "config": {},
      "enabled": true,
      "event-handlers": [
        {
          "class": "org.forgerock.audit.handlers.elasticsearch.ElasticsearchAuditEventHandler",
          "config": {
            "name": "elasticsearch",
            "topics": [
              "access"
            ],
            "connection": {
              "useSSL": false,
              "host": "localhost",
              "port": 9200
            },
            "indexMapping": {
              "indexName": "audit"
            },
            "buffering": {
              "enabled": true,
              "maxSize": 10000,
              "writeInterval": "250 millis",
              "maxBatchedEvents": 500
            }
          }
        }
      ]
    },
    "condition": "${matches(request.uri.path, '^/elasticsearch')}",
    "handler": {
      "type": "StaticResponseHandler",
      "config": {
```

```

    "entity": "View audit events in Elasticsearch at\rhttp://localhost:9200/audit/access/_search?q='\"OPENIG-HTTP-ACCESS\"'",
    "reason": "found",
    "status": 200,
    "headers": {
      "content-type": [
        "text/plain"
      ]
    }
  }
}
}
}
}

```

The route calls an audit service configuration for publishing log messages in Elasticsearch. When a request matches the `/elasticsearch` route, audit events are logged to the `ElasticsearchAuditEventHandler`.

The URL where you can view the messages logged by Elasticsearch is displayed. The URL is constructed from the host, port, index name, and topics defined in the event handler.

2. Access the route on `http://openig.example.com:8080/elasticsearch`.

The audit events are logged in Elasticsearch and the URL where you can view the messages is displayed.

3. Access the URL `http://localhost:9200/audit/access/_search?q='\"OPENIG-HTTP-ACCESS\"'`.

The audit events logged in Elasticsearch are displayed.

4. Repeat the previous two steps again to access the OpenIG route and then the Elasticsearch URL.

Each time you access the OpenIG route, the audit events logged in Elasticsearch should be updated.

Procedure 15.4. To Record Audit Events With a JMS Audit Event Handler

Important

This procedure is an example of how to record audit events with a JMS audit event handler configured to use the ActiveMQ message broker. This example is not tested on all configurations, and can be more or less relevant to your configuration.

For information about configuring the JMS event handler, see `JmsAuditEventHandler(5)` in the *Configuration Reference*.

Before you start, prepare OpenIG as described in Chapter 2, "Getting Started".

1. Add ActiveMQ client dependencies to OpenIG:
 - a. Download the following `.jar` files from *Apache ActiveMQ* :

- `geronimo-j2ee-management_1.1_spec-1.0.1.jar`
- `hawtbuf-1.11.jar`
- `activemq-client-5.13.3.jar`

b. Add the `.jar` files to the OpenIG container, at `/path/to/jetty/webapps/ROOT/WEB-INF/lib`.

2. Download and install the ActiveMQ message broker from <http://activemq.apache.org/>. For help, see the the ActiveMQ documentation on the same site.
3. Create a consumer that subscribes to the `audit` topic.

From the ActiveMQ installation directory, run the following command:

```
$ ./bin/activemq consumer --destination topic://audit
```

4. Add the following route to the OpenIG as `$HOME/.openig/config/routes/30-jms.json`.

On Windows, add the route as `%appdata%\OpenIG\config\routes\30-jms.json`:

```
{
  "MyCapture" : "all",
  "auditService" : {
    "config" : {
      "event-handlers" : [
        {
          "class" : "org.forgerock.audit.handlers.jms.JmsAuditEventHandler",
          "config" : {
            "name" : "jms",
            "topics" : [ "access" ],
            "deliveryMode" : "NON_PERSISTENT",
            "sessionMode" : "AUTO",
            "jndi" : {
              "contextProperties" : {
                "java.naming.factory.initial"
: "org.apache.activemq.jndi.ActiveMQInitialContextFactory",
                "java.naming.provider.url" : "tcp://openam.example.com:61616",
                "topic.audit" : "audit"
              },
              "topicName" : "audit",
              "connectionFactoryName" : "ConnectionFactory"
            }
          }
        }
      ],
      "config" : { }
    },
    "type" : "AuditService"
  },
  "handler" : {
    "type" : "StaticResponseHandler",
    "config" : {
      "status" : 200,
      "headers" : {

```

```

    "content-type" : [ "text/plain" ]
  },
  "reason" : "found",
  "entity" : "Message from audited route"
}
},
"monitor" : true,
"condition" : "${request.uri.path == '/activemq_event_handler'}"
}

```

When a request matches the `/activemq_event_handler` route, this configuration publishes JMS messages containing audit event data to an ActiveMQ managed JMS topic, and the `StaticResponseHandler` displays a message.

5. Access the route on `http://openig.example.com:8080/activemq_event_handler`.

Depending on how ActiveMQ is configured, audit events are displayed on the ActiveMQ console or written to file. For example, the following log message can be written to a log file in the folder where you installed ActiveMQ:

```

{
  "auditTopic": "access",
  "event": {
    "eventName": "OPENIG-HTTP-ACCESS",
    "timestamp": "2016-11-28T14:39:30.004Z",
    "transactionId": "882918f9-f7c3-47ee-9f87-5e3cfcfb98be-37",
    "server": {
      "ip": "0:0:0:0:0:0:1",
      "port": 8080
    },
    "client": {
      "ip": "0:0:0:0:0:0:1",
      "port": 56095
    },
    "http": {
      "request": {
        "secure": false,
        "method": "GET",
        "path": "http://openig.example.com:8080/activemq_event_handler",
        "queryParameters": {},
        "headers": {
          "accept": ["*/*"],
          "accept-encoding": ["gzip, deflate"],
          "Connection": ["keep-alive"],
          "host": ["openig.example.com:8080"],
          "user-agent": ["python-requests/2.9.1"]
        },
        "cookies": {}
      },
      "response": {
        "headers": {
          "Content-Length": ["26"],
          "Content-Type": ["text/plain"]
        }
      }
    }
  }
},

```

```
"response": {
  "status": "SUCCESSFUL",
  "statusCode": "200",
  "elapsedTime": 73,
  "elapsedTimeUnits": "MILLISECONDS"
},
"_id": "882918f9-f7c3-47ee-9f87-5e3cfcfb98be-38"
}
```

Chapter 16

Throttling the Rate of Requests to Protected Applications

To protect applications from being overused by clients, use a throttling filter to limit how many requests can be made in a defined time. This section describes how to set up two throttling filters. For more configuration options, see `ThrottlingFilter(5)` in the *Configuration Reference*.

The throttling filter limits the rate that requests pass through a filter. The maximum number of requests that are allowed in a defined time is called the *throttling rate*.

By default, the throttling filter uses a strategy based on the token bucket algorithm, which allows some bursts.

When the throttling rate is reached, OpenIG issues an HTTP status code 429 `Too Many Requests` and a `Retry-After` header like the following, where the value is the number of seconds to wait before trying the request again:

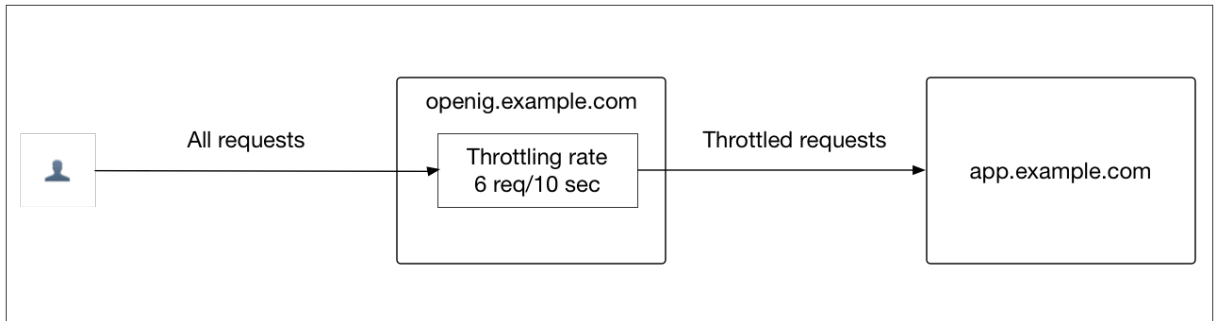
```
GET http://openig.example.com:8080/home/throttle-scriptable HTTP/1.1
. . .
HTTP/1.1 429 Too Many Requests
Retry-After: 10
```

16.1. Configuring a Simple Throttling Filter

This section describes how to use OpenIG Studio to configure a simple throttling filter that applies a throttling rate of 6 requests/10 seconds to requests. When an application is protected by this throttling filter, no more than 6 requests, irrespective of their origin, can access the sample application in a 10 second period.

To configure OpenIG without using OpenIG Studio, add the route in Example 16.1, "Route for Simple Throttling" to the OpenIG configuration as `$HOME/.openig/config/routes/simple-throttling.json`. On Windows, add the route as `%appdata%\OpenIG\config\routes\simple-throttling.json`.

Figure 16.1. Simple Throttling Filter



Example 16.1. Route for Simple Throttling

```

{
  "name": "00-throttle-simple",
  "baseURI": "http://app.example.com:8081",
  "monitor": false,
  "condition": "${matches(request.uri.path, '^/home/throttle-simple')}",
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ThrottlingFilter",
          "name": "Throttling",
          "config": {
            "rate": {
              "numberOfRequests": 6,
              "duration": "10 s"
            }
          }
        }
      ]
    }
  },
  "handler": "ClientHandler"
}

```

Procedure 16.1. To Configure a Simple Throttling Filter

This procedure configures a simple throttling filter in OpenIG Studio. The other procedures in this chapter create throttling filters manually, in the file system.

Before you start, prepare OpenIG and the sample application as described in Chapter 2, "Getting Started". Make sure that your `config.json` has a main router named `_router`.

1. Access OpenIG Studio on <http://openig.example.com:8080/openig/studio>, and select Protect an Application.
2. In the Create a route window, select Advanced Options and enter the following information, and then select Create route:
 - Base URI: <http://app.example.com:8081>
 - Condition: Path: </home/throttle-simple>
 - Name: [00-throttle-simple](#)
3. Select Throttling, and then enable throttling.
4. Select options to allow 6 requests each 10 seconds, and then save the settings.
5. On the top-right of the screen, select # > Display, and review the route. The route in Example 16.1, "Route for Simple Throttling" should be displayed.
6. Select Deploy to push the route to the OpenIG configuration.

You can check the [\\$HOME/.openig/config/routes](#) folder to see that the route is there.

Procedure 16.2. To Test the Setup

1. With OpenIG and the sample application running, use **curl**, a bash script, or another tool to access the following route in a loop: <http://openig.example.com:8080/home/simple-throttle>.

Accessing the route in a loop runs the request multiple times in quick succession, allowing you to test the throttling rate.

```
$ curl -v http://openig.example.com:8080/home/throttle-simple/[01-10] \  
> /tmp/simple-throttle.txt 2>&1
```

2. Search the output file to see the result:

```
$ grep "< HTTP/1.1" /tmp/simple-throttle.txt | sort | uniq -c  
6 < HTTP/1.1 200 OK  
4 < HTTP/1.1 429 Too Many Requests
```

Notice that the first six requests returned a success response, and the following four requests returned an HTTP 429 **Too Many Requests**. This result demonstrates that the throttling filter has allowed only six requests to access the application, and has blocked the other requests.

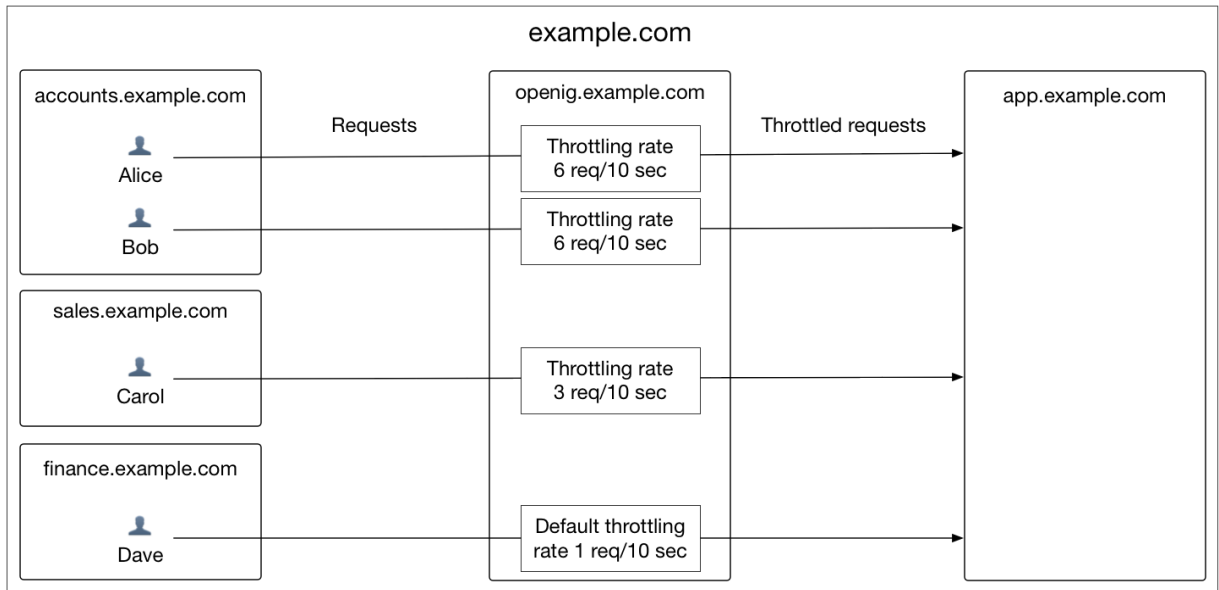
16.2. Configuring a Mapped Throttling Filter

The following example route uses a mapped throttling policy to map requests from users in the accounts and sales departments of `example.com` to different throttling rates. Requests from other departments use the default throttling rate.

The throttling rate is applied to groups according to the evaluation of `requestGroupingPolicy`. In the example, this parameter evaluates to the first `UserID` in the request header, representing each user.

The value of the throttling rate is assigned according to the evaluation of `throttlingRateMapper`. In the example, this parameter evaluates to the value of the request header `X-Forwarded-For`, representing the hostname of the department.

Figure 16.2. Mapped Throttling Filter



```

{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ThrottlingFilter",
          "config": {
            "requestGroupingPolicy": "${request.headers['UserId']}[0]",
          }
        }
      ]
    }
  }
}
  
```

```

    "throttlingRatePolicy": {
      "type": "MappedThrottlingPolicy",
      "config": {
        "throttlingRateMapper": "${request.headers['X-Forwarded-For']}[0]}",
        "throttlingRatesMapping": {
          "accounts.example.com": {
            "numberOfRequests": 6,
            "duration": "10 seconds"
          },
          "sales.example.com": {
            "numberOfRequests": 3,
            "duration": "10 seconds"
          }
        },
        "defaultRate": {
          "numberOfRequests": 1,
          "duration": "10 seconds"
        }
      }
    },
    "handler": "ClientHandler"
  },
  "condition": "${matches(request.uri.path, '^/home/throttle-mapped')}"
}

```

Procedure 16.3. To Configure a Mapped Throttling Filter

Before you start, prepare OpenIG and the sample application as shown in Chapter 2, "Getting Started".

1. Add the example route to the OpenIG configuration as `$HOME/.openig/config/routes/00-throttle-mapped.json`.

On Windows, add the route as `%appdata%\OpenIG\config\routes\00-throttle-mapped.json`.

2. With OpenIG and the sample application running, access the following route in a loop: `http://openig.example.com:8080/home/throttle-mapped`.

Accessing the route in a loop runs the request multiple times in quick succession, allowing you to test the throttling rate. You can use a command-line tool such as **curl** to run the following commands in quick succession or in a bash script:

```
$ curl -v http://openig.example.com:8080/home/throttle-mapped/[01-10] \
--header "X-Forwarded-For:accounts.example.com" \
--header "UserId:Alice" \
> /tmp/Alice.txt 2>&1

$ curl -v http://openig.example.com:8080/home/throttle-mapped/[01-10] \
--header "X-Forwarded-For:accounts.example.com" \
--header "UserId:Bob" \
> /tmp/Bob.txt 2>&1

$ curl -v http://openig.example.com:8080/home/throttle-mapped/[01-10] \
--header "X-Forwarded-For:sales.example.com" \
--header "UserId:Carol" \
> /tmp/Carol.txt 2>&1

$ curl -v http://openig.example.com:8080/home/throttle-mapped/[01-10] \
--header "X-Forwarded-For:finance.example.com" \
--header "UserId:Dave" \
> /tmp/Dave.txt 2>&1
```

3. Search the output files to see the result for each user and each organization:

```
$ grep "< HTTP/1.1" /tmp/Alice.txt | sort | uniq -c

6 < HTTP/1.1 200 OK
4 < HTTP/1.1 429 Too Many Requests

$ grep "< HTTP/1.1" /tmp/Bob.txt | sort | uniq -c

6 < HTTP/1.1 200 OK
4 < HTTP/1.1 429 Too Many Requests

$ grep "< HTTP/1.1" /tmp/Carol.txt | sort | uniq -c

3 < HTTP/1.1 200 OK
7 < HTTP/1.1 429 Too Many Requests

$ grep "< HTTP/1.1" /tmp/Dave.txt | sort | uniq -c

1 < HTTP/1.1 200 OK
9 < HTTP/1.1 429 Too Many Requests
```

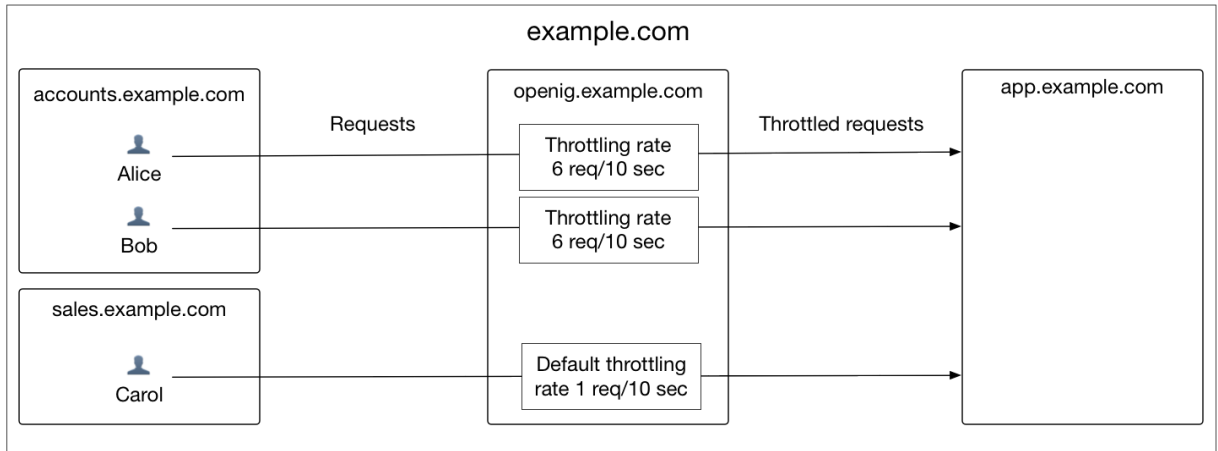
Notice that the first six requests from Alice and Bob in accounts are successful, and the first three requests from Carol in sales are successful, consistent with the mapping in `00-throttle-mapped.json`. Requests from finance are not mapped, and therefore receive the default rate.

16.3. Configuring a Scriptable Throttling Filter

In this example, the `DefaultRateThrottlingPolicy` delegates the management of throttling to the scriptable throttling policy.

The script applies a throttling rate of 6 requests/10 seconds to requests from the accounts department of `example.com`. For all other requests, the script returns `null`. When the script returns `null`, the default rate of 1 request/10 seconds is applied.

Figure 16.3. Scriptable Throttling Policy



```

{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ThrottlingFilter",
          "config": {
            "requestGroupingPolicy": "${request.headers['UserId']}[0]",
            "throttlingRatePolicy": {
              "type": "DefaultRateThrottlingPolicy",
              "config": {
                "delegateThrottlingRatePolicy": {
                  "type": "ScriptableThrottlingPolicy",
                  "config": {
                    "type": "application/x-groovy",
                    "source": [
                      "if (request.headers['X-Forwarded-For'].values[0] == 'accounts.example.com') {",
                      "  return new ThrottlingRate(6, '10 seconds')",
                      "} else {",
                      "  return null",
                      "}"
                    ]
                  }
                }
              }
            }
          }
        },
        {
          "type": "DefaultRateThrottlingPolicy",
          "config": {
            "defaultRate": {
              "numberOfRequests": 1,
            }
          }
        }
      ]
    }
  }
}
  
```

```

        "duration": "10 seconds"
      }
    }
  }
},
"handler": "ClientHandler"
}
},
"condition": "${matches(request.uri.path, '^/home/throttle-scriptable')}"
}

```

Procedure 16.4. To Configure a Scriptable Throttling Filter

Before you start, prepare OpenIG and the sample application as shown in Chapter 2, "Getting Started".

1. Add the example route to the OpenIG configuration as `$HOME/.openig/config/routes/00-throttle-scriptable.json`.

On Windows, add the route as `%appdata%\OpenIG\config\routes\00-throttle-scriptable.json`.

2. With OpenIG and the sample application running, access the following route in a loop: `http://openig.example.com:8080/home/throttle-scriptable`.

Accessing the route in a loop runs the request multiple times in quick succession, allowing you to test the throttling rate. You can use a command-line tool such as **curl** to run the following commands in quick succession or in a bash script:

```

$ curl -v http://openig.example.com:8080/home/throttle-scriptable/[01-10] \
  --header "X-Forwarded-For:accounts.example.com" \
  --header "UserId:Alice" \
  > /tmp/Alice.txt 2>&1

$ curl -v http://openig.example.com:8080/home/throttle-scriptable/[01-10] \
  --header "X-Forwarded-For:accounts.example.com" \
  --header "UserId:Bob" \
  > /tmp/Bob.txt 2>&1

$ curl -v http://openig.example.com:8080/home/throttle-scriptable/[01-10] \
  --header "X-Forwarded-For:sales.example.com" \
  --header "UserId:Carol" \
  > /tmp/Carol.txt 2>&1

```

3. Search the output files to see the result for each user and each organization:

```
$ grep "< HTTP/1.1" /tmp/Alice.txt | sort | uniq -c
6 < HTTP/1.1 200 OK
4 < HTTP/1.1 429 Too Many Requests

$ grep "< HTTP/1.1" /tmp/Bob.txt | sort | uniq -c
6 < HTTP/1.1 200 OK
4 < HTTP/1.1 429 Too Many Requests

$ grep "< HTTP/1.1" /tmp/Carol.txt | sort | uniq -c
1 < HTTP/1.1 200 OK
9 < HTTP/1.1 429 Too Many Requests
```

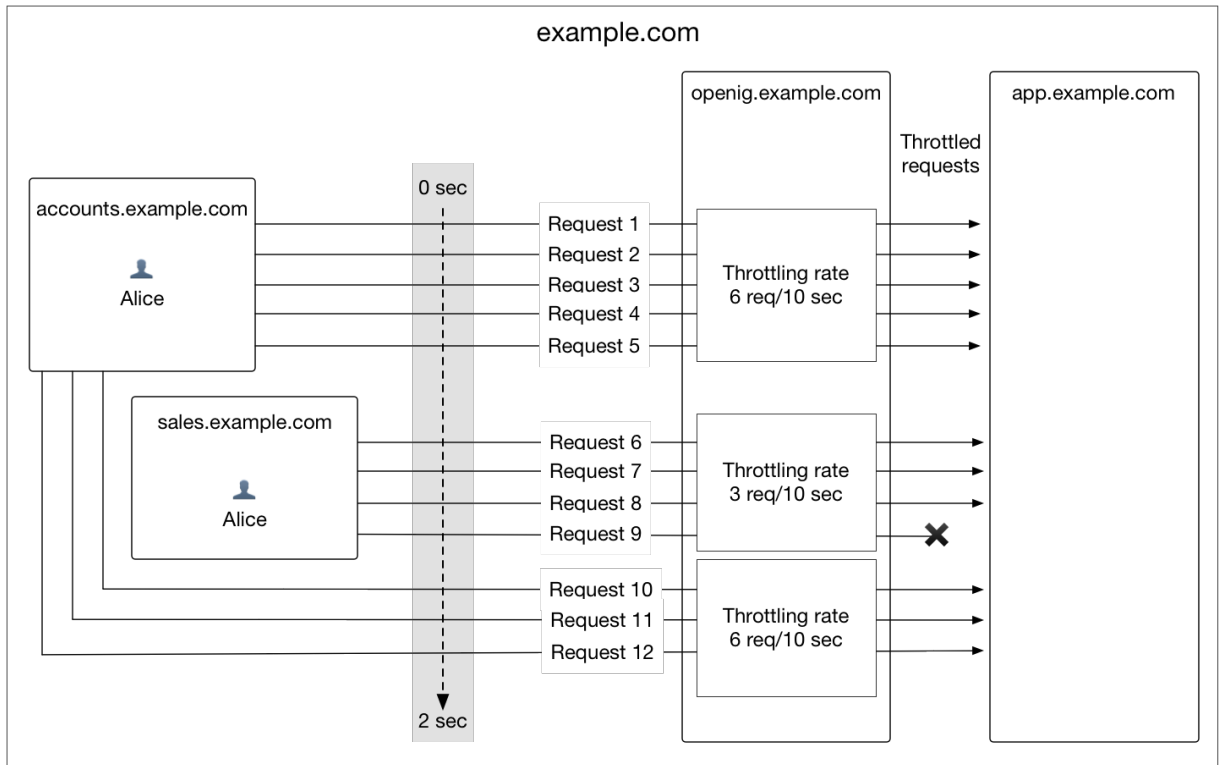
Notice that the first six requests from Alice and Bob in accounts are successful, consistent with the value in `ThrottlingScript.groovy`. The script returns `null` for requests from Carol in sales, so those requests receive the default throttling rate.

16.4. Dynamic Throttling Rate

In Section 16.2, "Configuring a Mapped Throttling Filter", requests from the same user were always sent from the same department in `example.com`. This example shows what happens to the throttling rate when a user sends requests from more than one department.

The throttling rate is applied to users according to the evaluation of `requestGroupingPolicy`, and different throttling rates are mapped to different departments of `example.com` according to the evaluation of `throttlingRateMapper`.

Figure 16.4. Dynamic Throttling Rate



In the example, Alice sends five requests from the accounts department, quickly followed by four requests from sales, and then three more requests from accounts.

After making five requests from accounts, Alice has almost reached the throttling rate. When she switches to sales, the number of requests she has already made is disregarded and the full throttling rate for sales is applied. Alice can now make three more requests even though she had nearly reached her throttling rate for accounts.

After making three requests from sales, Alice has reached her throttling rate. When she makes a fourth request from sales, the request is refused. Alice switches back to accounts and can now make six more requests even though she had reached her throttling rate for sales.

When you configure `requestGroupingPolicy` and `throttlingRateMapper`, bear in mind what happens when requests from the same `requestGroupingPolicy` can be mapped to different throttling rates by the `throttlingRateMapper`.

Chapter 17

Logging Events

Log messages in OpenIG and third-party dependencies are recorded using the Logback implementation of the Simple Logging Facade for Java (SLF4J) API.

17.1. Default Logging Behavior

Log messages are recorded with the following default configuration:

- When OpenIG starts, log messages for OpenIG and third-party dependencies, such as the ForgeRock Common Audit framework, are displayed on the console and written to `$HOME/.openig/logs/route-system.log`.
- When a route is accessed, log messages for requests passing through the route are written to a separate log file. The file is in `$HOME/.openig/logs` and is named by the route's name or filename. A separate log file is created for each route that is accessed.
- By default, log messages with the level `INFO` or higher are recorded, with the titles and the top line of the stack trace. The messages are highlighted with a color related to their logging level.

17.2. Reference Logback Configuration

The content and format of logs is defined by the reference `logback.xml` delivered with OpenIG. This file defines the following configuration items for logs:

- A root logger to set the overall level to `INFO`, and to write all log messages to the `SIFT` and `STDOUT` appenders.
- A `STDOUT` appender to define the format of log messages on the console.
- A `SIFT` appender to separate log messages according to the key `routeId`, to define when log files are rolled, and to define the format of log messages in the file.
- An exception logger, called `LogAttachedExceptionFilter`, to write log messages for exceptions attached to responses.

```
<?xml version="1.0" encoding="UTF-8"?><!--  
  Copyright 2016-2017 ForgeRock AS. All Rights Reserved  
  
  Use of this code requires a commercial software license with ForgeRock AS.
```



```

or with one of its affiliates. All use shall be exclusively subject
to such license between the licensee and ForgeRock AS.
--><configuration>

  <appender name="STDOUT" class="ch.qos.logback.core.ConsoleAppender">
    <encoder>
      <pattern>%nopex[%thread] %highlight(%-5level) %boldWhite(%logger{35}) - %message%n
%highlight(%rootException{short})</pattern>
    </encoder>
  </appender>

  <appender name="SIFT" class="ch.qos.logback.classic.sift.SiftingAppender">
    <discriminator>
      <key>routeId</key>
      <defaultValue>system</defaultValue>
    </discriminator>
    <sift>
      <!-- Create a separate log file for each <key> -->
      <appender name="FILE-${routeId}" class="ch.qos.logback.core.rolling.RollingFileAppender">
        <file>${openig.base}/logs/route-${routeId}.log</file>

        <rollingPolicy class="ch.qos.logback.core.rolling.SizeAndTimeBasedRollingPolicy">
          <!-- Rotate files daily -->
          <fileNamePattern>${openig.base}/logs/route-${routeId}-%d{yyyy-MM-dd}.%i.log</fileNamePattern>

          <!-- each file should be at most 100MB, keep 30 days worth of history, but at most 3GB -->
          <maxFileSize>100MB</maxFileSize>
          <maxHistory>30</maxHistory>
          <totalSizeCap>3GB</totalSizeCap>
        </rollingPolicy>

        <encoder>
          <pattern>%d{HH:mm:ss:SSS} | %-5level | %thread | %logger{20} | %message%n%n%xException</pattern>
        </encoder>
      </appender>
    </sift>
  </appender>

  <!-- Disable logs of exceptions attached to responses by defining 'level' to OFF -->
  <logger name="org.forgerock.openig.filter.LogAttachedExceptionFilter" level="INHERITED"/>

  <root level="INFO">
    <appender-ref ref="SIFT"/>
    <appender-ref ref="STDOUT"/>
  </root>
</configuration>

```

17.3. Changing the Logging Behavior

The Logback configuration is very flexible, providing a wide range of options for logging. For a full description of its parameters, see [the Logback website](#). The following examples show some simple changes that you can make.

To change the logging behavior, create a new Logback file at `$HOME/.openig/config/logback.xml`. This Logback file overrides the default configuration.

To take into account edits to `logback.xml`, stop and restart OpenIG or edit the `configuration` parameter to add a scan and interval:

```
<configuration scan="true" scanPeriod="5 seconds">
```

The configuration `scan="true"` requires `logback.xml` to be scanned for changes. The file is scanned after both of the following criteria are met:

- The specified number of logging operations have occurred, where the default is 16.
- The scan period has elapsed, where the example specifies 5 seconds.

Note

If your custom `logback.xml` contains errors, messages like these are displayed on the console but log messages are not recorded:

```
14:38:59,667 |-ERROR in ch.qos.logback.core.joran.spi.Interpreter@20:72 ...  
14:38:59,690 |-ERROR in ch.qos.logback.core.joran.action.AppenderRefAction ...
```

17.3.1. Formatting Log Messages

You can format log messages in many ways. For example, to add the date to the log message, edit `logback.xml` to change the pattern of the log messages in the `encoder` part of the SIFT appender:

```
%d{yyyyMMdd-HH:mm:ss} | %-5level | %thread | %logger{20} | %message%n%xException
```

17.3.2. Logging for Different Object Types

You can change the log messages for a single object type without changing them for the rest of the configuration.

For example, to record log messages with the level `ERROR` or higher for the `ClientHandler`, edit `logback.xml` to add a logger defined by the fully qualified class name of the `ClientHandler`, and to set its logging level to `ERROR`:

```
<logger name="org.forgerock.openig.handler.ClientHandler" level="ERROR"/>
```

Log messages with a level lower than `ERROR` are no longer recorded for the `ClientHandler` but continue to be recorded for the rest of the configuration.

17.3.3. Logging for the BaseUriDecorator

During setup and configuration, it can be helpful to display log messages from the `BaseUriDecorator`.

For example, to record a log message each time a request URI is rebased, edit `logback.xml` to add a logger defined by the fully qualified class name of the `BaseUriDecorator` appended by the name of the `baseURI` decorator:

```
<logger name="org.forgerock.openig.decoration.baseuri.BaseUriDecorator.baseURI" level="TRACE"/>
```

Each time a request URI is rebased, a log message similar to this is created:

```
12:27:40| TRACE | http-nio-8080-exec-3 | o.f.o.d.b.B.b.{Router}/handler  
| Rebasing request to http://app.example.com:8081
```

17.3.4. Switching Off Exception Logging

To stop recording log messages for exceptions, edit `logback.xml` to set the level to `OFF`:

```
<logger name="org.forgerock.openig.filter.LogAttachedExceptionFilter" level="OFF"/>
```

Chapter 18

Troubleshooting

This chapter covers common problems and their solutions.

To help with troubleshooting, get the product version and build information for the running instance of OpenIG from the `/api/info` endpoint. If OpenIG is set up as described in Chapter 2, "Getting Started", access the information at `http://openig.example.com:8080/openig/api/info`.

18.1. Troubleshooting the UMA Example

You have set up and are testing the example in Section 11.1, "About OpenIG in the UMA Resource Server Role".

If you have problems creating shares for Alice, perform the following steps to see if you can get a PAT from OpenAM:

1. With OpenAM running, run the following command to get a tokenID:

```
$ curl --request POST \
  --header "X-OpenAM-Username: alice" \
  --header "X-OpenAM-Password: password" \
  --header "Content-Type: application/json" \
  --data "{}" \
  http://openam.example.com:8088/openam/json/authenticate

{"tokenId":"AQIC5wM2LY . . . Dg5AAJTMQAA*","successUrl":"/openam/console"}
```

2. In the following command, replace `tokenID` with the value you got in the previous step:

```
$ curl -X POST -H "Cache-Control: no-cache" -H "Cookie: iPlanetDirectoryPro=tokenID" \
  -H "Content-Type: application/x-www-form-urlencoded" \
  -d \
  'grant_type=password&scope=uma_protection&username=alice&password=password&client_id=OpenIG&client_secret=pass' \
  http://openam.example.com:8088/openam/oauth2/access_token

{"access_token":"AQIC5wM2LY . . . Dg5AAJTMQAA*","scope":"uma_protection","token_type":"Bearer"
,"expires_in":3599}
```

An access token should be displayed. If you fail to get an access token, check that OpenAM is configured as described in Section 11.3, "Setting Up OpenAM As an Authorization Server".

If you continue to have problems, make sure that your OpenIG configuration matches that shown when you are running the test on <http://app.example.com:8081/uma/>.

18.2. Can't Deploy Routes in OpenIG Studio

OpenIG Studio deploys and undeploys routes through a main router named `_router`, which is the name of the main router in the default configuration. If you use a custom `config.json`, make sure that it contains a main router named `_router`.

For information about OpenIG Studio, see Section 12.4, "Creating Routes Through OpenIG Studio".

18.3. Object not found in heap

```
org.forgerock.json.fluent.JsonValueException: /handler:  
object Router2 not found in heap  
at org.forgerock.openig.heap.HeapImpl.resolve(HeapImpl.java:351)  
at org.forgerock.openig.heap.HeapImpl.resolve(HeapImpl.java:334)  
at org.forgerock.openig.heap.HeapImpl.getHandler(HeapImpl.java:538)
```

You have specified `"handler": "Router2"` in `config.json`, but no handler configuration object named Router2 exists. Make sure you have added an entry for the handler and that you have correctly spelled its name.

18.4. Extra or missing character / invalid JSON

When the JSON for a route is not valid, OpenIG does not load the route. Instead, a description of the error appears in the log:

```
16:09:50 | ERROR | openig.example.com-startStop-1 | o.f.o.h.r.RouterHandler |  
The file '/Users/me/.openig/config/routes/zz-default.json' is not a valid route configuration.
```

Use a JSON editor or JSON validation tool such as JSONLint to make sure that your JSON is valid.

18.5. The values in the flat file are incorrect

Ensure the flat file is readable by the user running the container for OpenIG. Values are all characters including space and tabs between the separator, so make sure the values are not padded with spaces.

18.6. Problem accessing URL

```
HTTP ERROR 500
```

```
Problem accessing /myURL . Reason:
```

```
java.lang.String cannot be cast to java.util.List
Caused by:
java.lang.ClassCastException: java.lang.String cannot be cast to java.util.List
```

This error is typically encountered when using an `AssignmentFilter` as described in `AssignmentFilter(5)` in the *Configuration Reference* and setting a string value for one of the headers. All headers are stored in lists so the header must be addressed with a subscript.

For example, rather than trying to set `request.headers['Location']` for a redirect in the response object, you should instead set `request.headers['Location'][0]`. A header without a subscript leads to the error above.

18.7. StaticResponseHandler results in a blank page

Define an entity for the response as in the following example:

```
{
  "name": "AccessDeniedHandler",
  "type": "StaticResponseHandler",
  "config": {
    "status": 403,
    "reason": "Forbidden",
    "entity": "<html><body><p>User does not have permission</p></body></html>"
  }
}
```

18.8. OpenIG is not logging users in

If you are proxying to more than one application in multiple DNS domains, you must make sure your container is enabled for domain cookies. For details on your specific container, see Section 3.1, "Configuring Deployment Containers".

18.9. Read timed out error when sending a request

If a `baseURI` configuration setting causes a request to come back to OpenIG, OpenIG never produces a response to the request. You then observe the following behavior.

You send a request and OpenIG seems to hang. Then you see a failure message, `HTTP Status 500 - Read timed out`, accompanied by OpenIG throwing an exception, `java.net.SocketTimeoutException: Read timed out`.

To fix this issue, make sure that `baseURI` configuration settings use a different host and port than the host and port for OpenIG.

18.10. OpenIG does not use new route configuration

OpenIG loads all configuration at startup. By default, it then periodically reloads changed route configurations.

If you make changes to a route that result in an invalid configuration, OpenIG logs errors, but it keeps the previous, correct configuration, and continues to use the old route.

OpenIG only uses the new configuration after you save a valid version or when you restart OpenIG.

Of course, if you restart OpenIG with an invalid route configuration, then OpenIG tries to load the invalid route at startup and logs an error. In that case, if there is no default handler to accept any incoming request for the invalid route, then you see an error, `No handler to dispatch to`.

18.11. Make OpenIG skip a route

If you have copied routes from another OpenIG server, those routes might depend on environment or container configuration that you have not yet configured locally.

You can work around this problem by changing the route file extension. A router ignores route files that do not have the `.json` extension.

For example, suppose you copy route all sample route configurations from the documentation, and then start OpenIG without first configuring your container. This can result in an error such as the following:

```

/handler/config/filters/0/config/dataSource: javax.naming.NameNotFoundException;
  remaining name 'jdbc/forgerock'
[   JsonValueException] > /handler/config/filters/0/config/dataSource:
  javax.naming.NameNotFoundException; remaining name 'jdbc/forgerock'
[   NameNotFoundException] > null

org.forgerock.json.fluent.JsonValueException:
/handler/config/filters/0/config/dataSource:
  javax.naming.NameNotFoundException; remaining name 'jdbc/forgerock'
at org.forgerock.openig.filter.SqlAttributesFilter$Heaplet.create(
  SqlAttributesFilter.java:211)
at org.forgerock.openig.heap.GenericHeaplet.create(GenericHeaplet.java:81)
at org.forgerock.openig.heap.HeapImpl.extract(HeapImpl.java:316)
at org.forgerock.openig.heap.HeapImpl.get(HeapImpl.java:281)
...

```

This arises from the route in `03-sql.json`, which defines an `SqlAttributesFilter` that depends on a JNDI data source configured in the container:

```

{
  "type": "SqlAttributesFilter",
  "config": {
    "dataSource": "java:comp/env/jdbc/forgerock",
    "preparedStatement":
      "SELECT username, password FROM users WHERE email = ?;",
    "parameters": [
      "george@example.com"
    ],
    "target": "${attributes.sql}"
  }
}

```

To prevent OpenIG from loading the route configuration until you have had time to configure the container, change the file extension to render the route inactive:

```
$ mv ~/.openig/config/routes/03-sql.json ~/.openig/config/routes/03-sql.inactive
```

If necessary, restart the container to force OpenIG to reload the configuration.

When you have configured the data source in the container, change the file extension back to `.json` to render the route active again:

```
$ mv ~/.openig/config/routes/03-sql.inactive ~/.openig/config/routes/03-sql.json
```


Appendix A. SAML 2.0 and Multiple Applications

Chapter 7, "*Acting As a SAML 2.0 Service Provider*" describes how to set up OpenIG as a SAML 2.0 service provider for a single application, using OpenAM as the identity provider. This chapter describes how to set up OpenIG as a SAML 2.0 service provider for two applications, still using OpenAM as the identity provider.

Before you try the samples described here, familiarize yourself with OpenIG SAML 2.0 support by reading and working through the examples in Chapter 7, "*Acting As a SAML 2.0 Service Provider*". Before you start, you should have OpenIG protecting the sample application as a SAML 2.0 service provider, with OpenAM working as identity provider configured as described in that tutorial.

A.1. Installation Overview

In this chapter you use the Fedlet configuration from Chapter 7, "*Acting As a SAML 2.0 Service Provider*" to create a configuration for each new protected application. You then import the new configurations as SAML 2.0 entities in OpenAM. If you subsequently edit a configuration, import it again.

In the following examples, the first application has entity ID `sp1` and runs on the host `sp1.example.com`, the second application has entity ID `sp2` and runs on the host `sp2.example.com`. To prevent unwanted behavior, the applications must have different values.

Table A.1. Tasks for Configuring SAML 2.0 SSO and Federation

Task	See Section(s)
Prepare the network.	Section A.2, "Preparing the Network"
Prepare the configuration for two OpenIG service providers.	Section A.3, "Configuring the Circle of Trust" Section A.4, "Configuring the Service Provider for Application One" Section A.5, "Configuring the Service Provider for Application Two"
Import the service provider configurations into OpenAM.	Section A.6, "Importing Service Provider Configurations Into OpenAM"
Add OpenIG routes.	Section A.7.1, "Preparing the Base Configuration File" Section A.7.2, "Preparing Routes for Application One" Section A.7.3, "Preparing Routes for Application Two"

A.2. Preparing the Network

Configure the network so that browser traffic to the application hosts is proxied through OpenIG.

Add the following addresses to your hosts file: sp1.example.com and sp2.example.com.

```
127.0.0.1 localhost openam.example.com openig.example.com app.example.com sp1.example.com
sp2.example.com
```

A.3. Configuring the Circle of Trust

Edit the `$HOME/.openig/SAML/fedlet.cot` file you created in Chapter 7, "Acting As a SAML 2.0 Service Provider" to include the entity IDs `sp1` and `sp2`, as in the following example:

```
cot-name=Circle of Trust
sun-fm-cot-status=Active
sun-fm-trusted-providers=openam, sp1, sp2
sun-fm-saml2-readerservice-url=
sun-fm-saml2-writerservice-url=
```

A.4. Configuring the Service Provider for Application One

To configure the service provider for application one, you can use the example files Example A.1, "Configuration File for Application One" and Example A.2, "Extended Configuration File for

Application One", saving them as `sp1.xml` and `sp1-extended.xml`. Alternatively, follow the steps below to use the files you created in Chapter 7, "Acting As a SAML 2.0 Service Provider".

Procedure A.1. To Configure the Service Provider for Application One By Using Files Created In Chapter 7, "Acting As a SAML 2.0 Service Provider"

1. Copy the SAML configuration files `sp.xml` and `sp-extended.xml` you created in Chapter 7, "Acting As a SAML 2.0 Service Provider", and save them as `$HOME/.openig/SAML/sp1.xml` and `$HOME/.openig/SAML/sp1-extended.xml`.
2. Make the following changes in `sp1.xml`:
 - For `entityID`, change `sp` to `sp1`. The `entityID` must match the application.
 - On each line that starts with `Location` or `ResponseLocation`, change `sp.example.com` to `sp1.example.com`, and add `/metaAlias/sp1` at the end of the line.

For an example of how this file should be, see Example A.1, "Configuration File for Application One".

3. Make the following changes in `sp1-extended.xml`:
 - For `entityID`, change `sp` to `sp1`.
 - For `SPSSOConfig metaAlias`, change `sp` to `sp1`.
 - For `appLogoutUrl`, change `sp` to `sp1`.
 - For `hosted=`, make sure that the value is `1`.

For an example of how this file should be, see Example A.2, "Extended Configuration File for Application One".

Example A.1. Configuration File for Application One

```
<!--
- sp1.xml.txt
- Set the entityID
- Set metaAlias/<sp-name> at the end of each of the following lines:
  - Location
  - ResponseLocation
- Note that AssertionConsumerService Location attributes include the metaAlias.
-->
<EntityDescriptor
  entityID="sp1"
  xmlns="urn:oasis:names:tc:SAML:2.0:metadata">
  <SPSSODescriptor
    AuthnRequestsSigned="false"
    WantAssertionsSigned="false"
    protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
    <SingleLogoutService
```

```

Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect"
Location="http://sp1.example.com:8080/saml/fedletSloRedirect/metaAlias/sp1"
ResponseLocation="http://sp1.example.com:8080/saml/fedletSloRedirect/metaAlias/sp1"/>
<SingleLogoutService
  Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
  Location="http://sp1.example.com:8080/saml/fedletSloPOST/metaAlias/sp1"
  ResponseLocation="http://sp1.example.com:8080/saml/fedletSloPOST/metaAlias/sp1"/>
<SingleLogoutService
  Binding="urn:oasis:names:tc:SAML:2.0:bindings:SOAP"
  Location="http://sp1.example.com:8080/saml/fedletSloSoap/metaAlias/sp1"/>
<NameIDFormat>urn:oasis:names:tc:SAML:2.0:nameid-format:transient</NameIDFormat>
<AssertionConsumerService
  isDefault="true"
  index="0"
  Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST"
  Location="http://sp1.example.com:8080/saml/fedletapplication/metaAlias/sp1"/>
<AssertionConsumerService
  index="1"
  Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifact"
  Location="http://sp1.example.com:8080/saml/fedletapplication/metaAlias/sp1"/>
</SPSSODescriptor>
<RoleDescriptor
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:query="urn:oasis:names:tc:SAML:metadata:ext:query"
  xsi:type="query:AttributeQueryDescriptorType"
  protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
</RoleDescriptor>
<XACMLAuthzDecisionQueryDescriptor
  WantAssertionsSigned="false"
  protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
</XACMLAuthzDecisionQueryDescriptor>
</EntityDescriptor>

```

Example A.2. Extended Configuration File for Application One

```

<!--
- sp1-extended.xml
- Set the entityID.
- Set the SPSSOConfig metaAlias attribute.
- Set the value of appLogoutUrl.
- Set the value of hosted to 1.
- Comment out the attribute "com.sun.identity.saml2.plugins.DefaultFedletAdapter".
-->
<EntityConfig xmlns="urn:sun:fm:SAML:2.0:entityconfig"
  xmlns:fm="urn:sun:fm:SAML:2.0:entityconfig"
  hosted="1"
  entityID="sp1">

  <SPSSOConfig metaAlias="/sp1">
    <Attribute name="description">
      <Value></Value>
    </Attribute>
    <Attribute name="signingCertAlias">
      <Value></Value>
    </Attribute>
    <Attribute name="encryptionCertAlias">

```

```

    <Value></Value>
  </Attribute>
  <Attribute name="basicAuthOn">
    <Value>>false</Value>
  </Attribute>
  <Attribute name="basicAuthUser">
    <Value></Value>
  </Attribute>
  <Attribute name="basicAuthPassword">
    <Value></Value>
  </Attribute>
  <Attribute name="autofedEnabled">
    <Value>>false</Value>
  </Attribute>
  <Attribute name="autofedAttribute">
    <Value></Value>
  </Attribute>
  <Attribute name="transientUser">
    <Value>anonymous</Value>
  </Attribute>
  <Attribute name="spAdapter">
    <Value></Value>
  </Attribute>
  <Attribute name="spAdapterEnv">
    <Value></Value>
  </Attribute>
  <!--
  <Attribute name="fedletAdapter">
    <Value>com.sun.identity.saml2.plugins.DefaultFedletAdapter</Value>
  </Attribute>
  -->
  <Attribute name="fedletAdapterEnv">
    <Value></Value>
  </Attribute>
  <Attribute name="spAccountMapper">
    <Value>com.sun.identity.saml2.plugins.DefaultLibrarySPAccountMapper</Value>
  </Attribute>
  <Attribute name="useNameIDAsSPUserID">
    <Value>>false</Value>
  </Attribute>
  <Attribute name="spAttributeMapper">
    <Value>com.sun.identity.saml2.plugins.DefaultSPAttributeMapper</Value>
  </Attribute>
  <Attribute name="spAuthncontextMapper">
    <Value>com.sun.identity.saml2.plugins.DefaultSPAuthnContextMapper</Value>
  </Attribute>
  <Attribute name="spAuthncontextClassrefMapping">
    <Value>
      urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport|0|default
    </Value>
  </Attribute>
  <Attribute name="spAuthncontextComparisonType">
    <Value>exact</Value>
  </Attribute>
  <Attribute name="attributeMap">
    <Value>employeenumber=employeenumber</Value>
    <Value>mail=mail</Value>
  </Attribute>
  <Attribute name="saml2AuthModuleName">

```

```
<Value></Value>
</Attribute>
<Attribute name="localAuthURL">
  <Value></Value>
</Attribute>
<Attribute name="intermediateUrl">
  <Value></Value>
</Attribute>
<Attribute name="defaultRelayState">
  <Value></Value>
</Attribute>
<Attribute name="appLogoutUrl">
  <Value>http://sp1.example.com:8080/saml/logout</Value>
</Attribute>
<Attribute name="assertionTimeSkew">
  <Value>300</Value>
</Attribute>
<Attribute name="wantAttributeEncrypted">
  <Value></Value>
</Attribute>
<Attribute name="wantAssertionEncrypted">
  <Value></Value>
</Attribute>
<Attribute name="wantNameIDEncrypted">
  <Value></Value>
</Attribute>
<Attribute name="wantPOSTResponseSigned">
  <Value></Value>
</Attribute>
<Attribute name="wantArtifactResponseSigned">
  <Value></Value>
</Attribute>
<Attribute name="wantLogoutRequestSigned">
  <Value></Value>
</Attribute>
<Attribute name="wantLogoutResponseSigned">
  <Value></Value>
</Attribute>
<Attribute name="wantMNIRequestSigned">
  <Value></Value>
</Attribute>
<Attribute name="wantMNIResponseSigned">
  <Value></Value>
</Attribute>
<Attribute name="responseArtifactMessageEncoding">
  <Value>URI</Value>
</Attribute>
<Attribute name="cotlist">
  <Value>Circle of Trust</Value></Attribute>
<Attribute name="saeAppSecretList">
  <Value></Value>
</Attribute>
<Attribute name="saeSPUrl">
  <Value></Value>
</Attribute>
<Attribute name="saeSPLogoutUrl">
  <Value></Value>
</Attribute>
<Attribute name="ECPRequestIDPLISTFinderImpl">
  <Value>com.sun.identity.saml2.plugins.ECPIDPFinder</Value>
</Attribute>
```

```

<Attribute name="ECPRequestIDPList">
  <Value></Value>
</Attribute>
<Attribute name="ECPRequestIDPListGetComplete">
  <Value></Value>
</Attribute>
<Attribute name="enableIDPPProxy">
  <Value>>false</Value>
</Attribute>
<Attribute name="idpProxyList">
  <Value></Value>
</Attribute>
<Attribute name="idpProxyCount">
  <Value>0</Value>
</Attribute>
<Attribute name="useIntroductionForIDPPProxy">
  <Value>>false</Value>
</Attribute>
<Attribute name="spSessionSyncEnabled">
  <Value>>false</Value>
</Attribute>
<Attribute name="relayStateUrlList">
  </Attribute>
</SPSSOConfig>
<AttributeQueryConfig metaAlias="/attrQuery">
  <Attribute name="signingCertAlias">
    <Value></Value>
  </Attribute>
  <Attribute name="encryptionCertAlias">
    <Value></Value>
  </Attribute>
  <Attribute name="wantNameIDEncrypted">
    <Value></Value>
  </Attribute>
  <Attribute name="cotlist">
    <Value>Circle of Trust</Value>
  </Attribute>
</AttributeQueryConfig>
<XACMLAuthzDecisionQueryConfig metaAlias="/pep">
  <Attribute name="signingCertAlias">
    <Value></Value>
  </Attribute>
  <Attribute name="encryptionCertAlias">
    <Value></Value>
  </Attribute>
  <Attribute name="basicAuthOn">
    <Value>>false</Value>
  </Attribute>
  <Attribute name="basicAuthUser">
    <Value></Value>
  </Attribute>
  <Attribute name="basicAuthPassword">
    <Value></Value>
  </Attribute>
  <Attribute name="wantXACMLAuthzDecisionResponseSigned">
    <Value>>false</Value>
  </Attribute>
  <Attribute name="wantAssertionEncrypted">
    <Value>>false</Value>
  </Attribute>

```

```
</Attribute>
<Attribute name="cotlist">
  <Value>Circle of Trust</Value>
</Attribute>
</XACMLAuthzDecisionQueryConfig>
</EntityConfig>
```

A.5. Configuring the Service Provider for Application Two

Procedure A.2. To Configure the Service Provider for Application Two

1. Copy the SAML configuration files `sp1.xml` and `sp1-extended.xml` you created in Section A.4, "Configuring the Service Provider for Application One", and save them as `$HOME/.openig/SAML/sp2.xml` and `$HOME/.openig/SAML/sp2-extended.xml`.
2. In both files, replace all incidences of `sp1` with `sp2`. To prevent unwanted behavior, application two must have different values to application one.

A.6. Importing Service Provider Configurations Into OpenAM

For each new protected application, import a SAML 2.0 entity into OpenAM. If you subsequently edit a service provider configuration, import it again.

Procedure A.3. To Import the Service Provider Configurations Into OpenAM

1. Log in to OpenAM console as administrator.
2. On the Federation tab, select the Entity Providers table and click Import Entity.

The Import Entity Provider page is displayed.

3. For the metadata file, select File and upload `sp1.xml`. For the extended data file, select File and upload `sp1-extended.xml`.
4. Repeat the previous step to upload `sp2.xml` and `sp2-extended.xml` for `sp2`.
5. Log out of the OpenAM console.

A.7. Preparing OpenIG Configurations

For each new protected application, prepare an OpenIG configuration. The configurations in this section follow the example in Chapter 7, "Acting As a SAML 2.0 Service Provider".

Tip

To prevent unspecified behavior, use different keys for session-stored values in the routes for each application. For example, use different keys for `session.sp1Username` and `session.sp2Username`.

To prevent configurations from overwriting each others' data, use the `subjectMapping` property of `SamlFederationHandler` to define a different session field for the subject name of each application. The two applications must not map data into the same session field.

A.7.1. Preparing the Base Configuration File

Edit the base configuration file, `$HOME/.openig/config/routes/config.json`, so that it does not rebase incoming URLs. The following example file differs from that used in earlier tutorials:

```
{
  "handler": {
    "type": "Router"
  },
  "heap": [
    {
      "name": "capture",
      "type": "CaptureDecorator",
      "config": {
        "captureEntity": true,
        "captureContext": true
      }
    }
  ]
}
```

Restart OpenIG to put the configuration changes into effect.

A.7.2. Preparing Routes for Application One

Set up the following routes for application one:

- `$HOME/.openig/config/routes/05-federate-sp1.json`, to redirect the request for SAML authentication. After authentication, this route logs the user in to the application.
- `$HOME/.openig/config/routes/05-saml-sp1.json`, to map attributes from the SAML assertion into the context, and then redirect the request back to the first route.

Example A.3. `05-federate-sp1.json`

```
{
  "handler": {
    "type": "DispatchHandler",
    "config": {
      "bindings": [
        {

```

```

        "condition": "${empty session.sp1Username}",
        "handler": {
            "type": "StaticResponseHandler",
            "config": {
                "status": 302,
                "reason": "Found",
                "headers": {
                    "Location": [
                        "http://sp1.example.com:8080/saml/SPInitiatedSSO?metaAlias=/sp1"
                    ]
                }
            }
        },
        {
            "handler": {
                "type": "Chain",
                "config": {
                    "filters": [
                        {
                            "type": "StaticRequestFilter",
                            "config": {
                                "method": "POST",
                                "uri": "http://app.example.com:8081/login",
                                "form": {
                                    "username": [
                                        "${session.sp1Username}"
                                    ],
                                    "password": [
                                        "${session.sp1Password}"
                                    ]
                                }
                            }
                        }
                    ],
                    "handler": "ClientHandler"
                }
            }
        }
    ],
    "condition": "${matches(request.uri.host, 'sp1.example.com') and not matches(request.uri.path, '^/saml')}"
}

```

Example A.4. `05-saml-sp1.json`

```

{
    "handler": {
        "type": "SamlFederationHandler",
        "config": {
            "comment": "Use unique session properties for this SP.",
            "assertionMapping": {
                "sp1Username": "mail",
                "sp1Password": "employeenumber"
            }
        }
    }
}

```

```

    },
    "authnContext": "sp1AuthnContext",
    "sessionIndexMapping": "sp1SessionIndex",
    "subjectMapping": "sp1SubjectName",
    "redirectURI": "/sp1"
  }
},
"condition": "${matches(request.uri.host, 'sp1.example.com') and matches(request.uri.path, '^/saml')}"
}

```

A.7.3. Preparing Routes for Application Two

Set up the following routes for application two:

- `$HOME/.openig/config/routes/05-federate-sp2.json`, to redirect the request for SAML authentication. After authentication, this route logs the user in to the application.
- `$HOME/.openig/config/routes/05-saml-sp2.json`, to map attributes from the SAML assertion into the context, and then redirect the request back to the first route.

Example A.5. `05-federate-sp2.json`

```

{
  "handler": {
    "type": "DispatchHandler",
    "config": {
      "bindings": [
        {
          "condition": "${empty session.sp2Username}",
          "handler": {
            "type": "StaticResponseHandler",
            "config": {
              "status": 302,
              "reason": "Found",
              "headers": {
                "Location": [
                  "http://sp2.example.com:8080/saml/SPInitiatedSSO?metaAlias=/sp2"
                ]
              }
            }
          }
        ]
      }
    },
    {
      "handler": {
        "type": "Chain",
        "config": {
          "filters": [
            {
              "type": "StaticRequestFilter",
              "config": {
                "method": "POST",
                "uri": "http://app.example.com:8081/login",
                "form": {
                  "username": [

```


- The link for SP-initiated SSO.
- The link for IDP-initiated SSO.

If you have not configured the examples exactly as shown in this guide, then adapt the SSO links accordingly.