FORGEROCK®

Configuration Reference

/ ForgeRock Identity Gateway 6.5

Latest update: 6.5.4

Paul Bryan
Mark Craig
Jamie Nelson
Joanne Henry

Copyright © 2011-2021 ForgeRock AS.

Abstract

Reference documentation for ForgeRock® Identity Gateway.

Table of Contents

Preface

ForgeRock Identity Platform™ serves as the basis for our simple and comprehensive Identity and Access Management solution. We help our customers deepen their relationships with their customers, and improve the productivity and connectivity of their employees and partners. For more information about ForgeRock and about the platform, see https://www.forgerock.com.

1. About This Guide

This guide describes in detail the configuration options for IG. It is for IG designers, developers, and administrators.

For API specifications, see the appropriate Javadoc.

2. Reserved Routes

By default, IG reserves all paths starting with `/openig` for administrative use, and only local client applications can access resources exposed under `/openig`.

To change the base for administrative routes, edit `admin.json`. For more information, see AdminHttpApplication(5).

3. Reserved Field Names

IG reserves all configuration field names that contain only alphanumeric characters.

If you must define your own field names, for example, in custom decorators, use names with dots, `.`, or dashes, `-`. Examples include `my-decorator` and `com.example.myDecorator`.

4. Field Value Conventions

IG configuration uses JSON notation.

This reference uses the following terms when referring to values of configuration object fields:

array

> JSON array.

boolean

Either `true` or `false`.

configuration token

Configuration tokens introduce variables into the server configuration. They can take values from Java system properties, environment variables, JSON and Java properties files held in specified directories, and from properties configured in routes.

For more information, see JSON Evaluation(5).

duration

A duration is a lapse of time expressed in English, such as `23 hours 59 minutes and 59 seconds`. Durations are not case sensitive, and negative durations are not supported. The following units can be used in durations:

- `indefinite`, `infinity`, `undefined`, `unlimited`: unlimited duration

- `zero`, `disabled`: zero-length duration

- `days`, `day`, `d`: days

- `hours`, `hour`, `h`: hours

- `minutes`, `minute`, `min`, `m`: minutes

- `seconds`, `second`, `sec`, `s`: seconds

- `milliseconds`, `millisecond`, `millisec`, `millis`, `milli`, `ms`: milliseconds

- `microseconds`, `microsecond`, `microsec`, `micros`, `micro`, `us`, `µs`: microseconds

- `nanoseconds`, `nanosecond`, `nanosec`, `nanos`, `nano`, `ns`: nanoseconds

expression

See Expressions(5).

configuration expression

Expression evaluated at configuration time, when routes are loaded.

Configuration expressions can refer to the system heap properties, the built-in functions listed in Functions(5), the `${env['variable']}`, and `${system['property']}`. Because configuration expressions are evaluated before any requests are made, they cannot refer to the runtime properties, `request`, `response`, or `context`. For more information, see Expressions(5).

runtime expression

Expression evaluated at runtime, for each request and response.

Runtime expressions can refer to the same information as configuration expressions, plus the following objects:

- `attributes`: org.forgerock.services.context.AttributesContext Map<String, Object>, obtained from `AttributesContext.getAttributes()`. For information, see AttributesContext(5).

- `context`: org.forgerock.services.context.Context object.

- `contexts`: map<string, context> object. For information, see Contexts(5).

- `request`: org.forgerock.http.protocol.Request object. For information, see Request(5).

- `response`: org.forgerock.http.protocol.Response object, available only when the expression is intended to be evaluated on the response flow. For information, see Response(5).

- `session`: org.forgerock.http.session.Session object, available only when the expression is intended to be evaluated for both request and response flow. For information, see SessionContext(5).

instant

An instantaneous point on the timeline, as a Java type. For more information, see Class Instant.

lvalue-expression

Expression yielding an object whose value is to be set.

number

JSON number.

object

JSON object where the content depends on the object's type.

pattern

A regular expression according to the rules for the Java Pattern class.

pattern-template

Template for referencing capturing groups in a pattern by using `$n`, where *n* is the index number of the capturing group starting from zero.

For example, if the pattern is `"\w+\s*=\s*(\w)+"`, the pattern-template is `"$1"`, and the text to match is `"key = value"`, the pattern-template yields `"value"`.

reference

References an object in the following ways:

- An inline configuration object, where the name is optional.

- A configuration expression that is a string or contains variable elements that evaluate to a string, where the string is the name of an object declared in the heap.

  For example, the following `temporaryStorage` object takes the value of the system property `openig.storage.ref`, which must a be string equivalent to the name of an object defined in the heap:

  ```
  {
    "temporaryStorage": "${system['openig.storage.ref']}"
  }
  ```

secret-id

String that references a secret managed the ForgeRock Commons Secrets Service, as described in Secrets.

The secret ID must conform to the following regex pattern: `Pattern.compile("[a-zA-Z0-9]+(\\.[a-zA-Z0-9]+)*");`

string

JSON string.

5. About ForgeRock Common REST

ForgeRock® Common REST is a common REST API framework. It works across the ForgeRock platform to provide common ways to access web resources and collections of resources. Adapt the examples in this section to your resources and deployment.

> **Note**
>
> This section describes the full Common REST framework. Some platform component products do not implement all Common REST behaviors exactly as described in this section. For details, refer to the product-specific examples and reference information in other sections of this documentation set.

5.1. Common REST Resources

Servers generally return JSON-format resources, though resource formats can depend on the implementation.

Resources in collections can be found by their unique identifiers (IDs). IDs are exposed in the resource URIs. For example, if a server has a user collection under `/users`, then you can access a user at `/users/user-id`. The ID is also the value of the `_id` field of the resource.

Resources are versioned using revision numbers. A revision is specified in the resource's `_rev` field. Revisions make it possible to figure out whether to apply changes without resource locking and without distributed transactions.

5.2. Common REST Verbs

The Common REST APIs use the following verbs, sometimes referred to collectively as CRUDPAQ. For details and HTTP-based examples of each, follow the links to the sections for each verb.

Create

Add a new resource.

This verb maps to HTTP PUT or HTTP POST.

For details, see "Create".

Read

Retrieve a single resource.

This verb maps to HTTP GET.

For details, see "Read".

Update

Replace an existing resource.

This verb maps to HTTP PUT.

For details, see "Update".

Delete

Remove an existing resource.

This verb maps to HTTP DELETE.

For details, see "Delete".

Patch

Modify part of an existing resource.

This verb maps to HTTP PATCH.

For details, see "Patch".

Action

Perform a predefined action.

This verb maps to HTTP POST.

For details, see "Action".

Query

Search a collection of resources.

This verb maps to HTTP GET.

For details, see "Query".

5.3. Common REST Parameters

Common REST reserved query string parameter names start with an underscore, `_`.

Reserved query string parameters include, but are not limited to, the following names:

```
_action
_api
_crestapi
_fields
_mimeType
_pageSize
_pagedResultsCookie
_pagedResultsOffset
_prettyPrint
_queryExpression
_queryFilter
_queryId
_sortKeys
_totalPagedResultsPolicy
```

> **Note**
>
> Some parameter values are not safe for URLs, so URL-encode parameter values as necessary.

Continue reading for details about how to use each parameter.

5.4. Common REST Extension Points

The *action* verb is the main vehicle for extensions. For example, to create a new user with HTTP POST rather than HTTP PUT, you might use `/users?_action=create`. A server can define additional actions. For example, `/tasks/1?_action=cancel`.

A server can define *stored queries* to call by ID. For example, `/groups?_queryId=hasDeletedMembers`. Stored queries can call for additional parameters. The parameters are also passed in the query string. Which parameters are valid depends on the stored query.

5.5. Common REST API Documentation

Common REST APIs often depend at least in part on runtime configuration. Many Common REST endpoints therefore serve *API descriptors* at runtime. An API descriptor documents the actual API as it is configured.

Use the following query string parameters to retrieve API descriptors:

_api

Serves an API descriptor that complies with the OpenAPI specification.

This API descriptor represents the API accessible over HTTP. It is suitable for use with popular tools such as Swagger UI.

_crestapi

Serves a native Common REST API descriptor.

This API descriptor provides a compact representation that is not dependent on the transport protocol. It requires a client that understands Common REST, as it omits many Common REST defaults.

> **Note**
>
> Consider limiting access to API descriptors in production environments in order to avoid unnecessary traffic.
>
> To provide documentation in production environments, see "To Publish OpenAPI Documentation" instead.

To Publish OpenAPI Documentation

In production systems, developers expect stable, well-documented APIs. Rather than retrieving API descriptors at runtime through Common REST, prepare final versions, and publish them alongside the software in production.

Use the OpenAPI-compliant descriptors to provide API reference documentation for your developers as described in the following steps:

1. Configure the software to produce production-ready APIs.

    In other words, the software should be configured as in production so that the APIs are identical to what developers see in production.

2. Retrieve the OpenAPI-compliant descriptor.

    The following command saves the descriptor to a file, `myapi.json`:

    ```
    $ curl -o myapi.json endpoint?_api
    ```

3. (Optional) If necessary, edit the descriptor.

    For example, you might want to add security definitions to describe how the API is protected.

    If you make any changes, then also consider using a source control system to manage your versions of the API descriptor.

4. Publish the descriptor using a tool such as Swagger UI.

    You can customize Swagger UI for your organization as described in the documentation for the tool.

5.6. Create

There are two ways to create a resource, either with an HTTP POST or with an HTTP PUT.

To create a resource using POST, perform an HTTP POST with the query string parameter `_action=create` and the JSON resource as a payload. Accept a JSON response. The server creates the identifier if not specified:

```
POST /users?_action=create HTTP/1.1
Host: example.com
Accept: application/json
Content-Length: ...
Content-Type: application/json
{ JSON resource }
```

To create a resource using PUT, perform an HTTP PUT including the case-sensitive identifier for the resource in the URL path, and the JSON resource as a payload. Use the `If-None-Match: *` header. Accept a JSON response:

```
PUT /users/some-id HTTP/1.1
Host: example.com
Accept: application/json
Content-Length: ...
Content-Type: application/json
If-None-Match: *
{ JSON resource }
```

The `_id` and content of the resource depend on the server implementation. The server is not required to use the `_id` that the client provides. The server response to the create request indicates the resource location as the value of the `Location` header.

If you include the `If-None-Match` header, its value must be `*`. In this case, the request creates the object if it does not exist, and fails if the object does exist. If you include the `If-None-Match` header with any value other than `*`, the server returns an HTTP 400 Bad Request error. For example, creating an object with `If-None-Match: revision` returns a bad request error. If you do not include `If-None-Match: *`, the request creates the object if it does not exist, and *updates* the object if it does exist.

Parameters

You can use the following parameters:

_prettyPrint=true

> Format the body of the response.

_fields=*field*[,*field*...]

> Return only the specified fields in the body of the response.

> The `field` values are JSON pointers. For example if the resource is `{"parent":{"child":"value"}}`, `parent/child` refers to the `"child":"value"`.

> If the `field` is left blank, the server returns all default values.

5.7. Read

To retrieve a single resource, perform an HTTP GET on the resource by its case-sensitive identifier (`_id`) and accept a JSON response:

```
GET /users/some-id HTTP/1.1
Host: example.com
Accept: application/json
```

Parameters

You can use the following parameters:

_prettyPrint=true

> Format the body of the response.

_fields=*field*[,*field*...]

> Return only the specified fields in the body of the response.

> The `field` values are JSON pointers. For example if the resource is `{"parent":{"child":"value"}}`, `parent/child` refers to the `"child":"value"`.

> If the `field` is left blank, the server returns all default values.

_mimeType=*mime-type*

> Some resources have fields whose values are multi-media resources such as a profile photo for example.

By specifying both a single *field* and also the *mime-type* for the response content, you can read a single field value that is a multi-media resource.

In this case, the content type of the field value returned matches the *mime-type* that you specify, and the body of the response is the multi-media resource.

The `Accept` header is not used in this case. For example, `Accept: image/png` does not work. Use the `_mimeType` query string parameter instead.

5.8. Update

To update a resource, perform an HTTP PUT including the case-sensitive identifier (`_id`) as the final element of the path to the resource, and the JSON resource as the payload. Use the `If-Match: _rev` header to check that you are actually updating the version you modified. Use `If-Match: *` if the version does not matter. Accept a JSON response:

```
PUT /users/some-id HTTP/1.1
Host: example.com
Accept: application/json
Content-Length: ...
Content-Type: application/json
If-Match: _rev
{ JSON resource }
```

When updating a resource, include all the attributes to be retained. Omitting an attribute in the resource amounts to deleting the attribute unless it is not under the control of your application. Attributes not under the control of your application include private and read-only attributes. In addition, virtual attributes and relationship references might not be under the control of your application.

Parameters

You can use the following parameters:

_prettyPrint=true

Format the body of the response.

_fields=*field*[,*field*...]

Return only the specified fields in the body of the response.

The `field` values are JSON pointers. For example if the resource is `{"parent":{"child":"value"}}`, `parent/child` refers to the `"child":"value"`.

If the `field` is left blank, the server returns all default values.

5.9. Delete

To delete a single resource, perform an HTTP DELETE by its case-sensitive identifier (`_id`) and accept a JSON response:

```
DELETE /users/some-id HTTP/1.1
Host: example.com
Accept: application/json
```

Parameters

You can use the following parameters:

`_prettyPrint=true`

Format the body of the response.

`_fields=field[,field...]`

Return only the specified fields in the body of the response.

The `field` values are JSON pointers. For example if the resource is `{"parent":{"child":"value"}}`, `parent/child` refers to the `"child":"value"`.

If the `field` is left blank, the server returns all default values.

5.10. Patch

To patch a resource, send an HTTP PATCH request with the following parameters:

- `operation`

- `field`

- `value`

- `from` (optional with copy and move operations)

You can include these parameters in the payload for a PATCH request, or in a JSON PATCH file. If successful, you'll see a JSON response similar to:

```
PATCH /users/some-id HTTP/1.1
Host: example.com
Accept: application/json
Content-Length: ...
Content-Type: application/json
If-Match: _rev
{ JSON array of patch operations }
```

PATCH operations apply to three types of targets:

- **single-valued**, such as an object, string, boolean, or number.

- **list semantics array**, where the elements are ordered, and duplicates are allowed.

- **set semantics array**, where the elements are not ordered, and duplicates are not allowed.

ForgeRock PATCH supports several different `operations`. The following sections show each of these operations, along with options for the `field` and `value`:

5.10.1. Patch Operation: Add

The `add` operation ensures that the target field contains the value provided, creating parent fields as necessary.

If the target field is single-valued, then the value you include in the PATCH replaces the value of the target. Examples of a single-valued field include: object, string, boolean, or number.

An `add` operation has different results on two standard types of arrays:

- **List semantic arrays**: you can run any of these `add` operations on that type of array:

  - If you `add` an array of values, the PATCH operation appends it to the existing list of values.

  - If you `add` a single value, specify an ordinal element in the target array, or use the `{-}` special index to add that value to the end of the list.

- **Set semantic arrays**: The list of values included in a patch are merged with the existing set of values. Any duplicates within the array are removed.

As an example, start with the following list semantic array resource:

```
{
    "fruits" : [ "orange", "apple" ]
}
```

The following add operation includes the pineapple to the end of the list of fruits, as indicated by the `-` at the end of the `fruits` array.

```
{
    "operation" : "add",
    "field" : "/fruits/-",
    "value" : "pineapple"
}
```

The following is the resulting resource:

```
{
    "fruits" : [ "orange", "apple", "pineapple" ]
}
```

5.10.2. Patch Operation: Copy

The copy operation takes one or more existing values from the source field. It then adds those same values on the target field. Once the values are known, it is equivalent to performing an `add` operation on the target field.

The following `copy` operation takes the value from a field named `mail`, and then runs a `replace` operation on the target field, `another_mail`.

```
[
  {
    "operation":"copy",
    "from":"mail",
    "field":"another_mail"
  }
]
```

If the source field value and the target field value are configured as arrays, the result depends on whether the array has list semantics or set semantics, as described in "Patch Operation: Add".

5.10.3. Patch Operation: Increment

The `increment` operation changes the value or values of the target field by the amount you specify. The value that you include must be one number, and may be positive or negative. The value of the target field must accept numbers. The following `increment` operation adds `1000` to the target value of `/user/payment`.

```
[
  {
    "operation" : "increment",
    "field" : "/user/payment",
    "value" : "1000"
  }
]
```

Since the `value` of the `increment` is a single number, arrays do not apply.

5.10.4. Patch Operation: Move

The move operation removes existing values on the source field. It then adds those same values on the target field. It is equivalent to performing a `remove` operation on the source, followed by an `add` operation with the same values, on the target.

The following `move` operation is equivalent to a `remove` operation on the source field, `surname`, followed by a `replace` operation on the target field value, `lastName`. If the target field does not exist, it is created.

```
[
  {
    "operation":"move",
    "from":"surname",
    "field":"lastName"
  }
]
```

To apply a `move` operation on an array, you need a compatible single-value, list semantic array, or set semantic array on both the source and the target. For details, see the criteria described in "Patch Operation: Add".

5.10.5. Patch Operation: Remove

The `remove` operation ensures that the target field no longer contains the value provided. If the remove operation does not include a value, the operation removes the field. The following `remove` deletes the value of the `phoneNumber`, along with the field.

```
[
  {
    "operation" : "remove",
    "field" : "phoneNumber"
  }
]
```

If the object has more than one `phoneNumber`, those values are stored as an array.

A `remove` operation has different results on two standard types of arrays:

- **List semantic arrays**: A `remove` operation deletes the specified element in the array. For example, the following operation removes the first phone number, based on its array index (zero-based):

  ```
  [
    {
      "operation" : "remove",
      "field" : "/phoneNumber/0"
    }
  ]
  ```

- **Set semantic arrays**: The list of values included in a patch are removed from the existing array.

5.10.6. Patch Operation: Replace

The `replace` operation removes any existing value(s) of the targeted field, and replaces them with the provided value(s). It is essentially equivalent to a `remove` followed by a `add` operation. If the arrays are used, the criteria is based on "Patch Operation: Add". However, indexed updates are not allowed, even when the target is an array.

The following `replace` operation removes the existing `telephoneNumber` value for the user, and then adds the new value of `+1 408 555 9999`.

```
[
  {
    "operation" : "replace",
    "field" : "/telephoneNumber",
    "value" : "+1 408 555 9999"
  }
]
```

A PATCH replace operation on a list semantic array works in the same fashion as a PATCH remove operation. The following example demonstrates how the effect of both operations. Start with the following resource:

```
{
    "fruits" : [ "apple", "orange", "kiwi", "lime" ],
}
```

Apply the following operations on that resource:

```
[
  {
    "operation" : "remove",
    "field" : "/fruits/0",
    "value" : ""
  },
  {
    "operation" : "replace",
    "field" : "/fruits/1",
    "value" : "pineapple"
  }
]
```

The PATCH operations are applied sequentially. The `remove` operation removes the first member of that resource, based on its array index, (`fruits/0`), with the following result:

```
[
  {
    "fruits" : [ "orange", "kiwi", "lime" ],
  }
]
```

The second PATCH operation, a `replace`, is applied on the second member (`fruits/1`) of the intermediate resource, with the following result:

```
[
  {
    "fruits" : [ "orange", "pineapple", "lime" ],
  }
]
```

5.10.7. Patch Operation: Transform

The `transform` operation changes the value of a field based on a script or some other data transformation command. The following `transform` operation takes the value from the field named `/objects`, and applies the `something.js` script as shown:

```
[
  {
    "operation" : "transform",
    "field" : "/objects",
    "value" : {
      "script" : {
        "type" : "text/javascript",
        "file" : "something.js"
      }
    }
  }
]
```

5.10.8. Patch Operation Limitations

Some HTTP client libraries do not support the HTTP PATCH operation. Make sure that the library you use supports HTTP PATCH before using this REST operation.

For example, the Java Development Kit HTTP client does not support PATCH as a valid HTTP method. Instead, the method `HttpURLConnection.setRequestMethod("PATCH")` throws `ProtocolException`.

Parameters

You can use the following parameters. Other parameters might depend on the specific action implementation:

`_prettyPrint=true`

Format the body of the response.

`_fields=field[,field...]`

Return only the specified fields in the body of the response.

The `field` values are JSON pointers. For example if the resource is `{"parent":{"child":"value"}}`, `parent/child` refers to the `"child":"value"`.

If the `field` is left blank, the server returns all default values.

5.11. Action

Actions are a means of extending Common REST APIs and are defined by the resource provider, so the actions you can use depend on the implementation.

The standard action indicated by `_action=create` is described in "Create".

Parameters

You can use the following parameters. Other parameters might depend on the specific action implementation:

_prettyPrint=true

> Format the body of the response.

_fields=*field*[,*field*...]

> Return only the specified fields in the body of the response.
>
> The `field` values are JSON pointers. For example if the resource is `{"parent":{"child":"value"}}`, `parent/child` refers to the `"child":"value"`.
>
> If the `field` is left blank, the server returns all default values.

5.12. Query

To query a resource collection (or resource container if you prefer to think of it that way), perform an HTTP GET and accept a JSON response, including at least a `_queryExpression`, `_queryFilter`, or `_queryId` parameter. These parameters cannot be used together:

```
GET /users?_queryFilter=true HTTP/1.1
Host: example.com
Accept: application/json
```

The server returns the result as a JSON object including a "results" array and other fields related to the query string parameters that you specify.

Parameters

You can use the following parameters:

_queryFilter=*filter-expression*

> Query filters request that the server return entries that match the filter expression. You must URL-escape the filter expression.
>
> The string representation is summarized as follows. Continue reading for additional explanation:

```
Expr          = OrExpr
OrExpr        = AndExpr ( 'or' AndExpr ) *
AndExpr       = NotExpr ( 'and' NotExpr ) *
NotExpr       = '!' PrimaryExpr | PrimaryExpr
PrimaryExpr   = '(' Expr ')' | ComparisonExpr | PresenceExpr | LiteralExpr
ComparisonExpr = Pointer OpName JsonValue
PresenceExpr  = Pointer 'pr'
LiteralExpr   = 'true' | 'false'
Pointer       = JSON pointer
OpName        = 'eq' |   # equal to
                'co' |   # contains
                'sw' |   # starts with
                'lt' |   # less than
                'le' |   # less than or equal to
                'gt' |   # greater than
                'ge' |   # greater than or equal to
                STRING  # extended operator
JsonValue     = NUMBER | BOOLEAN | '"' UTF8STRING '"'
STRING        = ASCII string not containing white-space
UTF8STRING    = UTF-8 string possibly containing white-space
```

JsonValue components of filter expressions follow RFC 7159: *The JavaScript Object Notation (JSON) Data Interchange Format*. In particular, as described in section 7 of the RFC, the escape character in strings is the backslash character. For example, to match the identifier `test\`, use `_id eq 'test\\'`. In the JSON resource, the `\` is escaped the same way: `"_id":"test\\"`.

When using a query filter in a URL, be aware that the filter expression is part of a query string parameter. A query string parameter must be URL encoded as described in RFC 3986: *Uniform Resource Identifier (URI): Generic Syntax* For example, white space, double quotes (`"`), parentheses, and exclamation characters need URL encoding in HTTP query strings. The following rules apply to URL query components:

```
query       = *( pchar / "/" / "?" )
pchar       = unreserved / pct-encoded / sub-delims / ":" / "@"
unreserved  = ALPHA / DIGIT / "-" / "." / "_" / "~"
pct-encoded = "%" HEXDIG HEXDIG
sub-delims  = "!" / "$" / "&" / "'" / "(" / ")"
              / "*" / "+" / "," / ";" / "="
```

ALPHA, DIGIT, and HEXDIG are core rules of RFC 5234: *Augmented BNF for Syntax Specifications*:

```
ALPHA       =  %x41-5A / %x61-7A   ; A-Z / a-z
DIGIT       =  %x30-39             ; 0-9
HEXDIG      =  DIGIT / "A" / "B" / "C" / "D" / "E" / "F"
```

As a result, a backslash escape character in a *JsonValue* component is percent-encoded in the URL query string parameter as `%5C`. To encode the query filter expression `_id eq 'test\\'`, use `_id +eq+'test%5C%5C'`, for example.

A simple filter expression can represent a comparison, presence, or a literal value.

For comparison expressions use *json-pointer comparator json-value*, where the *comparator* is one of the following:

`eq` (equals)
`co` (contains)
`sw` (starts with)
`lt` (less than)
`le` (less than or equal to)
`gt` (greater than)
`ge` (greater than or equal to)

For presence, use *json-pointer pr* to match resources where:

- The JSON pointer is present.

- The value it points to is not `null`.

Literal values include true (match anything) and false (match nothing).

Complex expressions employ `and`, `or`, and `!` (not), with parentheses, `(expression)`, to group expressions.

_queryId=*identifier*

Specify a query by its identifier.

Specific queries can take their own query string parameter arguments, which depend on the implementation.

_pagedResultsCookie=*string*

The string is an opaque cookie used by the server to keep track of the position in the search results. The server returns the cookie in the JSON response as the value of `pagedResultsCookie`.

In the request `_pageSize` must also be set and non-zero. You receive the cookie value from the provider on the first request, and then supply the cookie value in subsequent requests until the server returns a `null` cookie, meaning that the final page of results has been returned.

The `_pagedResultsCookie` parameter is supported when used with the `_queryFilter` parameter. The `_pagedResultsCookie` parameter is not guaranteed to work when used with the `_queryExpression` and `_queryId` parameters.

The `_pagedResultsCookie` and `_pagedResultsOffset` parameters are mutually exclusive, and not to be used together.

_pagedResultsOffset=*integer*

When `_pageSize` is non-zero, use this as an index in the result set indicating the first page to return.

The `_pagedResultsCookie` and `_pagedResultsOffset` parameters are mutually exclusive, and not to be used together.

_pageSize=*integer*

Return query results in pages of this size. After the initial request, use `_pagedResultsCookie` or `_pageResultsOffset` to page through the results.

_totalPagedResultsPolicy=*string*

When a `_pageSize` is specified, and non-zero, the server calculates the "totalPagedResults", in accordance with the `totalPagedResultsPolicy`, and provides the value as part of the response. The "totalPagedResults" is either an estimate of the total number of paged results (`_totalPagedResultsPolicy=ESTIMATE`), or the exact total result count (`_totalPagedResultsPolicy=EXACT`). If no count policy is specified in the query, or if `_totalPagedResultsPolicy=NONE`, result counting is disabled, and the server returns value of -1 for "totalPagedResults".

_sortKeys=[+-]*field*[,[+-]*field*...]

Sort the resources returned based on the specified field(s), either in `+` (ascending, default) order, or in `-` (descending) order.

Because ascending order is the default, including the `+` character in the query is unnecessary. If you do include the `+`, it must be URL-encoded as `%2B`, for example:

```
http://localhost:8080/api/users?_prettyPrint=true&_queryFilter=true&_sortKeys=%2Bname/givenName
```

The `_sortKeys` parameter is not supported for predefined queries (`_queryId`).

_prettyPrint=true

Format the body of the response.

_fields=*field*[,*field*...]

Return only the specified fields in each element of the "results" array in the response.

The `field` values are JSON pointers. For example if the resource is `{"parent":{"child":"value"}}`, `parent/child` refers to the `"child":"value"`.

If the `field` is left blank, the server returns all default values.

5.13. HTTP Status Codes

When working with a Common REST API over HTTP, client applications should expect at least the following HTTP status codes. Not all servers necessarily return all status codes identified here:

200 OK

The request was successful and a resource returned, depending on the request.

201 Created

The request succeeded and the resource was created.

204 No Content

The action request succeeded, and there was no content to return.

304 Not Modified

The read request included an `If-None-Match` header, and the value of the header matched the revision value of the resource.

400 Bad Request

The request was malformed.

401 Unauthorized

The request requires user authentication.

403 Forbidden

Access was forbidden during an operation on a resource.

404 Not Found

The specified resource could not be found, perhaps because it does not exist.

405 Method Not Allowed

The HTTP method is not allowed for the requested resource.

406 Not Acceptable

The request contains parameters that are not acceptable, such as a resource or protocol version that is not available.

409 Conflict

The request would have resulted in a conflict with the current state of the resource.

410 Gone

The requested resource is no longer available, and will not become available again. This can happen when resources expire for example.

412 Precondition Failed

The resource's current version does not match the version provided.

415 Unsupported Media Type

The request is in a format not supported by the requested resource for the requested method.

428 Precondition Required

The resource requires a version, but no version was supplied in the request.

500 Internal Server Error

The server encountered an unexpected condition that prevented it from fulfilling the request.

501 Not Implemented

The resource does not support the functionality required to fulfill the request.

503 Service Unavailable

The requested resource was temporarily unavailable. The service may have been disabled, for example.

6. 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.");
    }
}
```

7. 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.

8. Using the ForgeRock.org Site

The 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.

9. 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 on ForgeRock's support offering, including support plans and service level agreements (SLAs), visit https://www.forgerock.com/support.

Required Configuration

The configuration of IG is split into the following parts:

• AdminHttpApplication: the entry point for administrative requests

• GatewayHttpApplication: the entry point for gateway requests

Table of Contents

Name

AdminHttpApplication — initialize IG configuration

Description

The AdminHttpApplication serves requests on the administrative route, such as the creation of routes and the collection of monitoring information. The administrative route and its subroutes are reserved for administration endpoints.

The configuration is loaded from a JSON-encoded configuration file, expected by default at `$HOME/.openig/config/admin.json`. If you don't provide an `admin.json`, IG provides a file with the following defaults:

- The base of the administrative route is `/openig`. The product version and build information for a running instance of IG are at the `/openig/api/info` endpoint.

- IG runs in production mode.

The application creates the following objects by default:

- CaptureDecorator, to capture requests and response messages. The default CaptureDecorator is named capture. For details, see CaptureDecorator(5).

- SecretsService, to manage a store of secrets from files, system properties, and environment variables, by using ForgeRock Commons Secrets Service. The default SecretsService is a SystemAndEnvSecretStore with the default configuration. For information, see Secrets.

- TimerDecorator, to record time spent within filters and handlers. The default TimerDecorator is named timer. For details, see TimerDecorator(5).

- TemporaryStorage, to cache content during processing. To change the default values, add a TemporaryStorage object named `TemporaryStorage` and use non-default values.

The application declares default configurations in the heap for the following objects:

- ClientHandler, to communicate with with third-party services. For information, see ClientHandler(5).

- ReverseProxyHandler, to communicate with protected applications. For information, see ReverseProxyHandler(5).

- ForgeRockClientHandler, to send ForgeRock Common Audit transaction IDs when communicating with protected applications. The default object wraps a `ClientHandler`.

- ScheduledThreadPoolExecutor, to specify the number of threads in the pool.

- TransactionIdOutboundFilter, to insert the ID of a transaction in the header of the request.

Usage

```
{
    "heap": [ configuration object, ... ],
    "mode": enumeration,
    "prefix" : configuration expression<string>,
    "properties": JSON object,
    "secrets": configuration object,
    "temporaryStorage": TemporaryStorage reference,
    "preserveOriginalQueryString": configuration expression<boolean>
}
```

Properties

`mode`: *operating mode, optional*

Set the IG mode to `development` or `production`. The value is not case-sensitive.

- **Development mode (mutable mode)**

  Use development mode to evaluate or demo IG, 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 IG 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 IG configuration.

- **Production mode (immutable 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. Studio is effectively disabled, and you cannot manage, list, or even read routes through Common REST.

  By default, other endpoints, such as `/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 IG configuration.

If `mode` is not set, the IG mode can be set by the configuration token `ig.run.mode`. For more information, see Configuration Tokens(5).

Default: `production`

`"heap"`: *array of configuration objects, optional*

The heap object configuration, described in Heap Objects(5).

You can omit an empty array.

"prefix": *configuration expression<string>, optional*

The base of the route for administration requests. This route and its subroutes are reserved for administration endpoints.

Default: `openig`

"properties": *JSON object, optional*

Configuration parameters declared as property variables for use in the configuration. See also Properties(5).

Default: none

"secrets": *configuration object, optional*

An object that configures an inline array of one or more secret stores, as defined in Secrets.

```
{
  "secrets": {
    "stores": [ inline secret store declaration, ... ]
  }
}
```

"temporaryStorage": *TemporaryStorage reference, optional*

Cache content during processing based on this TemporaryStorage configuration.

Define the TemporaryStorage in one of the following ways:

- An inline TemporaryStorage configuration object.

- The name of a TemporaryStorage object defined in the heap.

- A configuration expression that evaluates to the name of a TemporaryStorage object defined in the heap.

Default: use the heap object named TemporaryStorage.

See also reference and TemporaryStorage(5).

"preserveOriginalQueryString": *configuration expression<boolean>, optional*

Process query strings in URLs, by applying or not applying a decode/encode process to the whole query string.

The following characters are disallowed in query string URL components: `"`, `{`, `}`, `<`, `>`, (space), and `|`. For more information about which query strings characters require encoding, see Uniform Resource Identifier (URI): Generic Syntax.

- `true`: Preserve query strings as they are presented.

  Select this option if the query string must not change during processing, for example, in signature verification.

  If a query string contains a disallowed character, the request produces a `400 Bad Request`.

- `false`: Tolerate disallowed characters in query string URL components by applying a decode/encode process to the whole query string.

  Select this option when a user agent or client produces query searches with disallowed characters. IG transparently encodes the disallowed characters before forwarding requests to the protected application.

  Characters in query strings are transformed as follows:

  - Allowed characters are not changed.

    For example, `sep=a` is not changed.

  - Percent-encoded values are re-encoded when the decoded value is an allowed character.

    For example, `sep=%27` is changed to `sep='`, because `'` is an allowed character.

  - Percent-encoded values are not changed when the decoded value is a disallowed character.

    For example, `sep=%22` is not changed, because `"` is a disallowed character.

  - Disallowed characters are encoded.

    For example, `sep=",` is changed to `sep=%22`, because `"` is a disallowed character.

  Default: `false`

Example

The following example shows an `admin.json` file configured to override the default `ApiProtectionFilter` that protects the reserved administrative route. This example is used in "To Set Up IG As an UMA Resource Server" in the *Gateway Guide*.

```
{
  "heap": [
    {
      "name": "ClientHandler",
      "type": "ClientHandler"
    },
    {
      "name": "ApiProtectionFilter",
      "type": "ScriptableFilter",
      "config": {
        "type": "application/x-groovy",
```

```
            "file": "CorsFilter.groovy"
        }
    }
  ],
  "prefix": "openig"
}
```

More Information

org.forgerock.openig.http.AdminHttpApplication

Name

GatewayHttpApplication — configure IG

Description

The GatewayHttpApplication is the entry point for all incoming gateway requests. It is responsible for initializing a heap of objects, described in Heap Objects(5), and providing the main Handler that receives all the incoming requests. The configuration is loaded from a JSON-encoded configuration file, expected by default at `$HOME/.openig/config/config.json`.

If you provide a `config.json`, the IG configuration is loaded from that file. If there is no file, the default configuration is loaded. For the default configuration, and the example `config.json` used in many of the examples in the documentation, see the Examples section of this page.

The 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.

> **Important**
>
> 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`.

The application creates the following objects by default:

- A BaseUriDecorator that you can use to override the scheme, host, and port of the existing request URI. The default BaseUriDecorator is named baseURI. For details, see BaseUriDecorator(5).

- A CaptureDecorator that you can use to capture requests and response messages. The default CaptureDecorator is named capture. For details, see CaptureDecorator(5).

- SecretsService, to manage a store of secrets from files, system properties, and environment variables, by using ForgeRock Commons Secrets Service. The default SecretsService is a SystemAndEnvSecretStore with the default configuration. For information, see Secrets.

- A TimerDecorator that you can use to record time spent within Filters and Handlers. The default TimerDecorator is named timer. For details, see TimerDecorator(5).

The application declares default configurations in the heap for the following objects:

- ClientHandler, to communicate with with third-party services. For information, see ClientHandler(5).

- ReverseProxyHandler, to communicate with protected applications. For information, see ReverseProxyHandler(5).

- ForgeRockClientHandler, to send ForgeRock Common Audit transaction IDs when communicating with protected applications. The default object wraps a `ClientHandler`.

- ScheduledThreadPoolExecutor, to specify the number of threads in the pool.

- TransactionIdOutboundFilter, to insert the ID of a transaction in the header of the request.

The application looks for an object named Session in the heap:

- If it finds an object named Session, it uses it as the default session producer.

  For example, to store session information in an HTTP cookie on the user-agent, you can define a JwtSession named Session in `config.json`. Stored session information must fit the constraints for storage in a JWT and in a cookie, as described in JwtSession(5).

- If it doesn't find an object named Session, the session is based on the Servlet HttpSession that is handled by the container where IG runs.

  Session information is stored in an IG cookie called by default `IG_SESSIONID`.

Usage

```
{
  "handler": Handler reference or inline Handler declaration,
  "heap": [ configuration object, ... ],
  "properties": JSON object,
  "secrets": configuration object,
  "temporaryStorage": TemporaryStorage reference
}
```

Properties

`"handler"`: *Handler reference, required*

Dispatch all requests to this handler.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

See also Handlers.

`"heap"`: *array of configuration objects, optional*

The heap object configuration, described in Heap Objects(5).

You can omit an empty array. If you only have one object in the heap, you can inline it as the handler value.

"properties": *JSON object, optional*

Configuration parameters declared as property variables for use in the configuration. See also Properties(5).

Default: none

"secrets": *configuration object, optional*

An object that configures an inline array of one or more secret stores, as defined in Secrets.

```
{
  "secrets": {
    "stores": [ inline secret store declaration, ... ]
  }
}
```

"temporaryStorage": *TemporaryStorage reference, optional*

Cache content during processing based on this TemporaryStorage configuration.

Provide either the name of a TemporaryStorage object defined in the heap, or an inline TemporaryStorage configuration object.

Default: use the heap object named TemporaryStorage. Otherwise use an internally-created TemporaryStorage object that is named TemporaryStorage and that uses default settings for a TemporaryStorage object.

See also TemporaryStorage(5).

Example Configuration Files

Default Configuration

When your configuration does not include a `config.json` file, the following configuration is provided by default.

```
{
    "heap": [
        {
            "name": "_router",
            "type": "Router",
            "config": {
                "defaultHandler": {
                    "type": "DispatchHandler",
                    "config": {
                        "bindings": [
                            {
                                "condition": "${request.method == 'GET' and request.uri.path == '/'}",
```

```
                              "handler": {
                                  "type": "WelcomeHandler"
                              }
                          },
                          {
                              "condition": "${request.uri.path == '/'}",
                              "handler": {
                                  "type": "StaticResponseHandler",
                                  "config": {
                                      "status": 405,
                                      "reason": "Method Not Allowed"
                                  }
                              }
                          },
                          {
                              "handler": {
                                  "type": "StaticResponseHandler",
                                  "config": {
                                      "status": 404,
                                      "reason": "Not Found"
                                  }
                              }
                          }
                      ]
                  }
              }
          }
      }
  ],
  "handler": "_router"
}
```

Notice the following features of the default configuration:

- The handler contains a main router named `_router`. When IG receives an incoming request, `_router` routes the request to the first route in the configuration whose condition is satisfied.

- If the request doesn't satisfy the condition of any route, it is routed to the defaultHandler. If the request is to access the IG welcome page, IG dispatches the request. Otherwise, IG returns an HTTP status 404 (Resource not found), because the requested resource does not exist.

Example config.json Used In the Doc

The following example of `config.json` is used in many of the examples in the documentation:

```
{
  "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
      }
    }
  ]
}
```

Notice the following features of the file:

- The handler contains a main router named `_router`. When IG receives an incoming request, `_router` routes the request to the first route in the configuration whose condition is satisfied.

- The `baseURI` changes the request URI to point the request to the sample application.

- The `capture` captures the body of the HTTP request and response.

- The JwtSession object in the heap can be used in routes to store the session information as JSON Web Tokens (JWT) in a cookie. For more information, see JwtSession(5).

More Information

org.forgerock.openig.http.GatewayHttpApplication

Name

Heap Objects — configure and initialize objects, with dependency injection

Description

A *heaplet* creates and initializes an object that is stored in a heap. A heaplet can retrieve objects it depends on from the heap.

A *heap* is a collection of associated objects created and initialized by heaplet objects. All configurable objects in IG are heap objects.

The heap configuration is included as an object in the configuration of the GatewayHttpApplication and AdminHttpApplication. For information, see AdminHttpApplication(5) and GatewayHttpApplication(5).

Usage

```
[
    {
        "name": string,
        "type": string,
        "config": {
            object-specific configuration
        }
    },
    ...
]
```

Properties

"name": *string, required except for inline objects*

The unique name to give the heap object in the heap. This name is used to resolve the heap object, for example, when another heap object names a heap object dependency.

"type": *string, required*

The class name of the object to be created. To determine the type name, see the object's documentation in this reference.

"config": *object, required*

The configuration that is specific to the heap object being created.

If all the fields are optional and the configuration uses only default settings, you can omit the config field instead of including an empty config object as the field value.

Provided Objects

IG creates the following objects when a filter with the name of the object is declared in `admin.json`:

"ApiProtectionFilter"

> The default filter used to protect administrative APIs on reserved routes. Reserved routes are described in "Reserved Routes".
>
> To override this filter, declare a different filter with the same name in `admin.json`. For more information, see AdminHttpApplication(5).
>
> Default: a filter that allows access only from the loopback address.

"MetricsProtectionFilter"

> A filter used to protect the monitoring endpoints.
>
> By default, the Prometheus Scrape Endpoint and Common REST Monitoring Endpoint are open and accessible. No special credentials or privileges are required to access the monitoring endpoints.
>
> To protect these monitoring endpoints, add an `admin.json` file to your configuration, with a filter declared in the heap and named `MetricsProtectionFilter`.
>
> For an example of how to manage access, see "Protecting the Monitoring Endpoints" in the *Gateway Guide*. For more information about `admin.json`, see AdminHttpApplication(5).

"StudioProtectionFilter"

> A filter to protect the Studio endpoint when IG is running in development mode.
>
> When IG is running in development mode, by default the Studio endpoint is open and accessible. Use this filter to restrict access to Studio in development mode.
>
> For an example configuration that requires users to authenticate with AM before they can access Studio, see "Restricting Access to Studio in Development Mode" in the *Gateway Guide*.
>
> For information about accessing Studio, see "*Configuring Routes With Studio*" in the *Getting Started Guide*.
>
> For information about `mode`, see AdminHttpApplication(5).

"TemporaryStorage"

> The default object to use for managing temporary buffers.
>
> Default: a TemporaryStorage object named "TemporaryStorage" with the default configuration is added to the top-level heap.

Routes can use this object without explicitly defining it. To override this object, create a TemporaryStorage heap object with the same name.

See also TemporaryStorage(5).

More Information

org.forgerock.openig.heap.Heap

Name

Configuration Settings — configure objects

Description

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.

Expressions can retrieve the values for configuration settings from system properties or environment variables. When IG starts up or when a route is reloaded, the expressions are evaluated. If you change the value of a system property or environment variable and then restart IG or reload the route, the configuration settings are updated with the new values.

If a configuration setting is required and the expression returns `null`, an error occurs when IG starts up or when the route is reloaded. If the configuration setting is optional, there is no error.

In the following example, `"timer"` is defined by an expression that recovers the environment variable `"ENABLE_TIMER"` and transforms it into a boolean. Similarly, `"numberOfRequests"` is defined by an expression that recovers the system property `"requestsPerSecond"` and transforms it into an integer:

```
{
  "name": "throttle-simple-expressions1",
  "timer" : "${bool(env['ENABLE_TIMER'])}",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/throttle-simple-expressions1')}",
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ThrottlingFilter",
          "name": "ThrottlingFilter-1",
          "config": {
            "requestGroupingPolicy": "",
            "rate": {
              "numberOfRequests": "${integer(system['requestsPerSecond'])}",
              "duration": "10 s"
            }
          }
        }
      ],
      "handler": "ReverseProxyHandler"
    }
  }
}
```

If `"requestsPerSecond"=6` and `"ENABLE_TIMER"=true`, after the expressions are evaluated IG views the example route as follows:

```
{
  "name": "throttle-simple-expressions2",
  "timer" : true,
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/throttle-simple-expressions2')}",
```

```
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ThrottlingFilter",
          "name": "ThrottlingFilter-1",
          "config": {
            "requestGroupingPolicy": "",
            "rate": {
              "numberOfRequests": 6,
              "duration": "10 s"
            }
          }
        }
      ],
      "handler": "ReverseProxyHandler"
    }
  }
}
```

For information about expressions, see Expressions(5).

Handlers

Handler objects process an HTTP request by producing an associated response.

Table of Contents

FORGEROCK

Name

Chain — dispatch the request to ordered list of filters and finally to a handler

Description

Dispatches a request to an ordered list of filters, and then finally to a handler.

Unlike `ChainOfFilters`, `Chain` finishes by dispatching the request to a handler. For more information, see ChainOfFilters(5).

Usage

```
{
    "name": string,
    "type": "Chain",
    "config": {
        "filters": [ Filter reference, ... ],
        "handler": Handler reference
    }
}
```

Properties

`"filters"`: *array of filter references, required*

An array of names of filter objects defined in the heap, and inline filter configuration objects.

The chain dispatches the request to these filters in the order they appear in the array.

See also Filters.

`"handler"`: *Handler reference, required*

Either the name of a handler object defined in the heap, or an inline handler configuration object.

The chain dispatches to this handler after the request has traversed all of the specified filters.

See also Handlers.

Example

```
{
    "name": "LoginChain",
    "type": "Chain",
    "config": {
        "filters": [ "LoginFilter" ],
        "handler": "ReverseProxyHandler"
    }
}
```

More Information

org.forgerock.openig.filter.ChainHandlerHeaplet

Name

ClientHandler — relay requests to third-party services

Description

Creates a response to a request by forwarding the request to a third-party service accessible through HTTP, and reconstructing the response from the received bytes.

When IG relays the request to the third-party service, IG is acting as a client of the service. IG is *client-side*.

A third-party service is one that IG calls for data, such as an HTTP API or AM, or one to which IG submits data, such as Splunk from an audit event handler.

If IG fails to connect to the third-party service, the ClientHandler propagates the error along the chain.

Use ClientHandler to submit requests to third-party services such as AM or HTTP APIs. Do not use it to proxy requests to a protected application. To proxy requests to a protected application, use a ReverseProxyHandler(5) instead.

Usage

```
{
    "name": string,
    "type": "ClientHandler",
    "config": {
        "connections": number,
        "disableReuseConnection": configuration expression<boolean>,
        "stateTrackingEnabled": configuration expression<boolean>,
        "hostnameVerifier": configuration expression<enumeration>,
        "soTimeout": duration string,
        "connectionTimeout": duration string,
        "connectionTimeToLive": duration string,
        "numberOfWorkers": number,
        "proxy": Server reference,
        "systemProxy": boolean,
        "temporaryStorage": string,
        "tls": TlsOptions reference,
        "asyncBehavior": enumeration,
        "retries": object,
        "websocket": object
    }
}
```

Properties

"connections": *number, optional*

> The maximum number of connections in the HTTP client connection pool.

Default: 64

"connectionTimeout": *duration string, optional*

Amount of time to wait to establish a connection, expressed as a duration

Default: 10 seconds

For information about supported formats for `duration`, see duration.

"connectionTimeToLive": *duration string, optional*

Amount of time before a reusable pooled connection expires.

Set this property to expire reusable pooled connections after a fixed duration. For example, to prevent the reuse of connections set this property in routes for applications where the IP address (baseURI) is not stable or can change.

Default: Unlimited

For information about supported formats for `duration`, see duration.

"disableReuseConnection": *configuration expression<boolean>, optional*

Whether to disable connection reuse.

Default: `false`

"stateTrackingEnabled": *configuration expression<boolean>, optional*

By default, the Apache HTTP Client does not allow connection reuse when a client certificate is used for authentication. However, because the client certificate is defined at the client level, it is acceptable for requests to the same target to share a client certificate.

Use in combination with `disableReuseConnection`:

disableReuseConnection	stateTrackingEnabled	Description
`false` (default)	`true` (default)	Do not allow connection reuse when a client certificate is used for authentication.
`false` (default)	`false`	Allow connection reuse when a client certificate is used for authentication.
`true`	`true` (default) or `false`	Do not allow connection reuse.

Default: `true`

"hostnameVerifier": *configuration expression<enumeration>, optional*

Way to handle hostname verification for outgoing SSL connections. Use one of the following values:

- `ALLOW_ALL`: Allow a certificate issued by a trusted CA for any hostname or domain to be accepted for a connection to any domain.

> **Caution**
>
> This setting allows a certificate issued for one company to be accepted as a valid certificate for another company.
>
> To prevent the compromise of TLS connections, use this setting in development mode only. In production, use `STRICT`.

- `STRICT`: Match the hostname either as the value of the the first CN, or any of the subject-alt names.

  A wildcard can occur in the CN, and in any of the subject-alt names. Wildcards match one domain level, so `*.example.com` matches `www.example.com` but not `some.host.example.com`.

Default: `STRICT`

`"numberOfWorkers"`: *number, optional*

The number of worker threads dedicated to processing outgoing requests.

Increasing the value of this attribute can be useful in deployments where a high number of simultaneous connections remain open, waiting for protected applications to respond.

Default: One thread per CPU available to the JVM.

`"proxy"`: *Server reference, optional*

A proxy server to which requests can be submitted. Use this property to relay requests to other parts of the network, for example, to relay requests from an internal network to the internet.

If both `proxy` and `systemProxy` are defined, `proxy` takes precedence.

`uri`: *configuration expression<uri string>, required*

URI of a server to use as a proxy for outgoing requests.

The result of the expression must be a string that represents a valid URI, but is not a real `java.net.URI` object.

`username`: *string, required if the proxy requires authentication*

Username to access the proxy server.

`"passwordSecretId"`: *configuration expression<string>, required if the proxy requires authentication*

The reference to the password required to access the proxy server.

The secret ID must conform to the regex pattern: `[a-zA-Z0-9]+(.[a-zA-Z0-9]+)*`. The route's secret service attempts to resolve the secret ID.

In the following example, the ClientHandler passes outgoing requests to the proxy server, which requires authentication:

```
"handler": {
  "type": "ClientHandler",
  "config": {
     "proxy" : {
        "uri": "http://proxy.example.com:3128",
        "username": "proxyuser",
        "passwordSecretId": "myproxy.secret.id"
     }
  }
}
```

`"soTimeout"`: *duration string, optional*

Socket timeout, after which stalled connections are destroyed, expressed as a duration

> **Tip**
>
> If `SocketTimeoutException` errors occur in the logs when you try to download large files, consider increasing `soTimeout`.

Default: 10 seconds

For information about supported formats for `duration`, see duration.

`"systemProxy"`: *boolean, optional*

Submit outgoing requests to a system-defined proxy, set by the following system properties or their HTTPS equivalents:

- `http.proxyHost`, the host name of the proxy server.

- `http.proxyPort`, the port number of the proxy server. The default is `80`.

- `http.nonProxyHosts`, a list of hosts that should be reached directly, bypassing the proxy.

This property can't be used with a proxy that requires a username and password. Use the property `proxy` instead.

If both `proxy` and `systemProxy` are defined, `proxy` takes precedence.

For more information, see *Java Networking and Proxies*.

Default: False.

`"temporaryStorage"`: *string, optional*

Specifies the heap object to use for temporary buffer storage.

Default: The temporary storage object named `TemporaryStorage`, declared in the top-level heap.

`tls`: *TlsOptions reference, optional*

Configure options for connections to TLS-protected endpoints, based on a TlsOptions(5) configuration. Define a TlsOptions object inline or in the heap.

Default: Connections to TLS-protected endpoints are not configured.

`"asyncBehavior"`: *enumeration, optional*

Specifies how the HTTP client behaves for asynchronous responses. Set the value to `streaming` or `non_streaming` (not case-sensitive):

* `streaming`: Responses are processed as soon as all headers are received. The entity content is downloaded in a background thread.

  Streaming mode reduces latency and is mandatory for Server-Sent Events (SSE) and the support of very large files (bigger than 2 GB). If thread starvation occurs, consider increasing `numberOfWorkers`, the number of worker threads dedicated to processing outgoing requests.

* `non_streaming`: Responses are processed when the entity content is entirely available.

  Non-streaming mode does not support SSE or very large files. However, it has higher latency, and does not cause thread starvation.

Default: `non_streaming`

`"retries"`: *object, optional*

Enable and configure retry for requests.

When a runtime error occurs while executing the request to the remote server, IG schedules a new execution of the request after the specified delay, until the allowed number of retries is reached or the execution succeeds.

`"enabled"`: *boolean, optional*

Enable retries.

Default: `true`

`"executor"`: *ScheduledExecutorService reference, optional*

The ScheduledExecutorService to use for scheduling delayed execution of the request.

Default: `ScheduledExecutorService`.

See also ScheduledExecutorService(5).

`"count"`: *number, optional*

The maximum number of retries to perform.

After this threshold is passed and if the request is still not successful, then the ClientHandler will propagate the failure.

Default: `5` retries.

`"delay"`: *duration, optional*

The delay to wait before retrying the request.

After a failure to send the request, if the number of retries is below the threshold, a new attempt is scheduled with the executor service after this delay.

Default: `10 seconds`.

For information about supported formats for `duration`, see duration.

The following example configures the handler to retry the request only once, after a 1-minute delay:

```
{
  "retries": {
    "count": 1,
    "delay": "1 minute"
  }
}
```

The following example configures the handler to retry the request at most 20 times, every second:

```
{
  "retries": {
    "count": 20,
    "delay": "1 second"
  }
}
```

The following example configures the handler to retry the request 5 times, every 10 seconds (default values), with a dedicated executor:

```
{
  "retries": {
    "executor": {
      "type": "ScheduledExecutorService",
      "config": {
        "corePoolSize": 20
      }
    }
  }
}
```

`"websocket"`: *object, optional*

Enable upgrade from HTTP or HTTPS protocol to WebSocket protocol.

For more information and an example of proxying WebSocket traffic, see "*Proxying WebSocket Traffic*" in the *Gateway Guide*

> **Important**
>
> When IG is running in the Jetty application container, it cannot proxy WebSocket traffic.

```
{
  "websocket": {
    "enabled": boolean,
    "connectionTimeout": duration string,
    "soTimeout": duration string,
    "numberOfSelectors": number,
    "tls": TlsOptions reference
  }
}
```

For more information, see The WebSocket Protocol, RFC6455.

`"enabled"`: *boolean, optional*

Enable upgrade from HTTP protocol and HTTPS protocol to WebSocket protocol.

Default: `false`

`"connectionTimeout"`: *duration string, optional*

The maximum time allowed to establish a WebSocket connection.

Default: The value of handler's main `connectionTimeout`.

For information about supported formats for `duration`, see duration.

`"soTimeout"`: *duration string, optional*

The time after which stalled connections are destroyed.

> **Tip**
>
> If there can be long delays between messages, consider increasing this value. Alternatively, keep the connection active by using WebSocket ping messages in your application.

Default: The value of handler's main `soTimeout`.

For information about supported formats for `duration`, see duration.

`"numberOfSelectors"`: *number, optional*

The maximum number of worker threads.

In deployments with a high number of simultaneous connections, consider increasing the value of this property.

Default: `2`

FORGEROCK

`tls:` *TlsOptions reference, optional*

Configure options for connections to TLS-protected endpoints, based on a TlsOptions(5) configuration. Define a TlsOptions object inline or in the heap.

Default: Use TlsOptions defined for the handler

Example

The following object configures a `ClientHandler` named `Client`:

```
{
  "name": "Client",
  "type": "ClientHandler",
  "config": {
    "hostnameVerifier": "STRICT",
    "tls": {
      "type": "TlsOptions",
      "config": {
        "sslContextAlgorithm": "TLSv1.2",
        "keyManager": {
          "type": "KeyManager",
          "config": {
            "keystore": {
              "type": "KeyStore",
              "config": {
                "url": "file://${env['HOME']}/keystore.jks",
                "passwordSecretId": "keymanager.keystore.secret.id"
              }
            },
            "passwordSecretId": "keymanager.secret.id"
          }
        },
        "trustManager": {
          "type": "TrustManager",
          "config": {
            "keystore": {
              "type": "KeyStore",
              "config": {
                "url": "file://${env['HOME']}/truststore.jks",
                "passwordSecretId": "trustmanager.keystore.secret.id"
              }
            }
          }
        }
      }
    }
  }
}
```

More Information

org.forgerock.openig.handler.ClientHandler

Name

DesKeyGenHandler — generate a DES key

Description

Generates a DES key for use with AM as described in "To Create a DES Shared Key In IG" in the *Gateway Guide*.

Usage

```
{
    "name": string,
    "type": "DesKeyGenHandler"
}
```

More Information

org.forgerock.openig.handler.DesKeyGenHandler

Name

DispatchHandler — dispatch to a handler if a condition is met

Description

When a request is handled, the first condition in the list of conditions is evaluated. If the condition expression yields `true`, the request is dispatched to the associated handler with no further processing. Otherwise, the next condition in the list is evaluated.

Usage

```
{
    "name": string,
    "type": "DispatchHandler",
    "config": {
        "bindings": [
            {
                "condition": runtime expression<boolean>,
                "handler": Handler reference,
                "baseURI": runtime expression<uri string>,
            }, ...
        ]
    }
}
```

Properties

`"bindings"`: *array of objects, required*

A list of bindings of conditions and associated handlers to dispatch to.

`"condition"`: *runtime expression<boolean>, optional*

If the expression evaluates `true`, the request is dispatched to the associated handler. If no condition is specified, the request is dispatched to the associated handler unconditionally.

Default: No condition is specified.

`"handler"`: *Handler reference, required*

Dispatch to this handler if the associated condition yields `true`.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

See also Handlers.

"baseURI": *runtime expression<uri string>, optional*

A base URI that overrides the existing request URI. Only scheme, host, and port are used in the supplied URI.

The result of the expression must be a string that represents a valid URI, but is not a real `java.net.URI` object. For example, it would be incorrect to use `${request.uri}`, which is not a String but a MutableUri.

In the following example, the binding condition looks up the hostname of the request. If it finds a match, the value is used for the `baseURI`. Otherwise, the default value is used:

```
{
  "properties": {
    "uris": {
      "app1.example.com": {
        "baseURI": "http://backend1:8080/"
      },
      "app2.example.com": {
        "baseURI": "http://backend2:8080/"
      },
      "default": {
        "baseURI": "http://backend3:8080/"
      }
    }
  },
  "handler": {
    "type": "DispatchHandler",
    "config": {
      "bindings": [
        {
          "condition": "${not empty uris[contexts.router.originalUri.host]}",
          "baseURI": "${uris[contexts.router.originalUri.host].baseURI}",
          "handler": "ReverseProxyHandler"
        },
        {
          "baseURI": "${uris['default'].baseURI}",
          "handler": "ReverseProxyHandler"
        }
      ]
    }
  }
}
```

Default: No change to the base URI

Example

The following sample is from a SAML 2.0 federation configuration:

- If the incoming URI matches `/saml`, then IG dispatches to a SamlFederationHandler without further processing.

- If the previous condition is not met, and the user name is not set in the session context, then IG dispatches the request to a SPInitiatedSSORedirectHandler.

  In this case, the user has not authenticated with the SAML 2.0 Identity Provider, so the SPInitiatedSSORedirectHandler initiates SAML 2.0 SSO from the Service Provider, which is IG.

- If neither of the previous conditions are met, the request goes through a LoginChain handler.

```
{
    "name": "DispatchHandler",
    "type": "DispatchHandler",
    "config": {
        "bindings": [
            {
                "condition": "${matches(request.uri.path, '^/saml')}",
                "handler": "SamlFederationHandler"
            },
            {
                "condition": "${empty session.username}",
                "handler": "SPInitiatedSSORedirectHandler",
                "baseURI": "http://www.example.com:8081"
            },
            {

                "handler": "LoginChain",
                "baseURI": "http://www.example.com:8081"
            }
        ]
    }
}
```

More Information

org.forgerock.openig.handler.DispatchHandler

Expressions(5)

Name
ReverseProxyHandler — relay requests to proxied applications, return 502 on failure

Description

Proxy requests to protected applications.

When IG relays the request to the protected application, IG is acting as a client of the application. IG is *client-side*.

If the request is to upload or download a large file to or from the protected application, consider setting `asyncBehavior` to `streaming`, and increasing the value of `soTimeout`.

If IG fails to connect to the protected application, the ReverseProxyHandler does not propagate the error along the chain. Instead, it changes the runtime exception into a `502 Bad Gateway` response.

Use ReverseProxyHandler in a route to proxy requests to a protected application. To submit requests to third-party services, such as AM or HTTP APIs, use a ClientHandler(5) instead.

Usage

```
{
    "name": string,
    "type": "ReverseProxyHandler",
    "config": {
        "connections": number,
        "disableReuseConnection": configuration expression<boolean>,
        "stateTrackingEnabled": configuration expression<boolean>,
        "hostnameVerifier": configuration expression<enumeration>,
        "soTimeout": duration string,
        "connectionTimeout": duration string,
        "connectionTimeToLive": duration string,
        "numberOfWorkers": number,
        "proxy": Server reference,
        "systemProxy": boolean,
        "temporaryStorage": string,
        "tls": TlsOptions reference,
        "asyncBehavior": enumeration,
        "retries": object,
        "websocket": object
    }
}
```

Properties

`"connections"`: *number, optional*

The maximum number of connections in the HTTP client connection pool.

Default: 64

"connectionTimeout": *duration string, optional*

Amount of time to wait to establish a connection, expressed as a duration

Default: 10 seconds

For information about supported formats for `duration`, see duration.

"connectionTimeToLive": *duration string, optional*

Amount of time before a reusable pooled connection expires.

Set this property to expire reusable pooled connections after a fixed duration. For example, to prevent the reuse of connections set this property in routes for applications where the IP address (baseURI) is not stable or can change.

Default: Unlimited

For information about supported formats for `duration`, see duration.

"disableReuseConnection": *configuration expression<boolean>, optional*

Whether to disable connection reuse.

Default: `false`

"stateTrackingEnabled": *configuration expression<boolean>, optional*

By default, the Apache HTTP Client does not allow connection reuse when a client certificate is used for authentication. However, because the client certificate is defined at the client level, it is acceptable for requests to the same target to share a client certificate.

Use in combination with `disableReuseConnection`:

disableReuseConnection	stateTrackingEnabled	Description
`false` (default)	`true` (default)	Do not allow connection reuse when a client certificate is used for authentication.
`false` (default)	`false`	Allow connection reuse when a client certificate is used for authentication.
`true`	`true` (default) or `false`	Do not allow connection reuse.

Default: `true`

"hostnameVerifier": *configuration expression<enumeration>, optional*

Way to handle hostname verification for outgoing SSL connections. Use one of the following values:

- `ALLOW_ALL`: Allow a certificate issued by a trusted CA for any hostname or domain to be accepted for a connection to any domain.

  > **Caution**
  >
  > This setting allows a certificate issued for one company to be accepted as a valid certificate for another company.
  >
  > To prevent the compromise of TLS connections, use this setting in development mode only. In production, use `STRICT`.

- `STRICT`: Match the hostname either as the value of the the first CN, or any of the subject-alt names.

  A wildcard can occur in the CN, and in any of the subject-alt names. Wildcards match one domain level, so `*.example.com` matches `www.example.com` but not `some.host.example.com`.

Default: `STRICT`

`"numberOfWorkers"`: *number, optional*

The number of worker threads dedicated to processing outgoing requests.

Increasing the value of this attribute can be useful in deployments where a high number of simultaneous connections remain open, waiting for protected applications to respond.

Default: One thread per CPU available to the JVM.

`"proxy"`: *Server reference, optional*

A proxy server to which requests can be submitted. Use this property to relay requests to other parts of the network, for example, to submit requests from an internal network to the internet.

If both `proxy` and `systemProxy` are defined, `proxy` takes precedence.

`url`: *configuration expression<uri string>, required*

URI of a server to use as a proxy for outgoing requests.

The result of the expression must be a string that represents a valid URI, but is not a real `java.net.URI` object.

`username`: *string, required if the proxy requires authentication*

Username to access the proxy server.

`password`: *string, required if the proxy requires authentication*

Password to access the proxy server.

In the following example, the ReverseProxyHandler passes outgoing requests to the proxy server, which requires authentication:

```
"handler": {
  "type": "ReverseProxyHandler",
  "config": {
      "proxy" : {
          "uri": "http://proxy.example.com:3128",
          "username": "proxy-user",
          "password": "proxy-password"
      }
  }
}
```

`"soTimeout"`: *duration string, optional*

Socket timeout, after which stalled connections are destroyed, expressed as a duration

Default: 10 seconds

For information about supported formats for `duration`, see duration.

`"systemProxy"`: *boolean, optional*

Submit outgoing requests to a system-defined proxy, set by the following system properties or their HTTPS equivalents:

- `http.proxyHost`, the host name of the proxy server.

- `http.proxyPort`, the port number of the proxy server. The default is `80`.

- `http.nonProxyHosts`, a list of hosts that should be reached directly, bypassing the proxy.

This property can't be used with a proxy that requires a username and password. Use the property `proxy` instead.

If both `proxy` and `systemProxy` are defined, `proxy` takes precedence.

For more information, see *Java Networking and Proxies*.

Default: False.

`"temporaryStorage"`: *string, optional*

Specifies the heap object to use for temporary buffer storage.

Default: The temporary storage object named `TemporaryStorage`, declared in the top-level heap.

`tls`: *TlsOptions reference, optional*

Configure options for connections to TLS-protected endpoints, based on a TlsOptions(5) configuration. Define a TlsOptions object inline or in the heap.

Default: Connections to TLS-protected endpoints are not configured.

`"asyncBehavior"`: *enumeration, optional*

Specifies how the HTTP client behaves for asynchronous responses. Set the value to `streaming` or `non_streaming` (not case-sensitive):

- `streaming`: Responses are processed as soon as all headers are received. The entity content is downloaded in a background thread.

  Streaming mode reduces latency and is mandatory for Server-Sent Events (SSE) and the support of very large files (bigger than 2 GB). If thread starvation occurs, consider increasing `numberOfWorkers`, the number of worker threads dedicated to processing outgoing requests.

- `non_streaming`: Responses are processed when the entity content is entirely available.

  Non-streaming mode does not support SSE or very large files. However, it has higher latency, and does not cause thread starvation.

> **Tip**
>
> If timeout errors occur in the logs, consider setting `soTimeout` to limit the timeout, and setting `asyncBehavior` to `non_streaming`.

Default: `non_streaming`

`"retries"`: *object, optional*

Enable and configure retry for requests.

When a runtime error occurs while executing the request to the remote server, IG schedules a new execution of the request after the specified delay, until the allowed number of retries is reached or the execution succeeds.

`"enabled"`: *boolean, optional*

Enable retries.

Default: `true`

`"executor"`: *ScheduledExecutorService reference, optional*

The ScheduledExecutorService to use for scheduling delayed execution of the request.

Default: `ScheduledExecutorService`.

See also ScheduledExecutorService(5).

`"count"`: *number, optional*

The maximum number of retries to perform.

After this threshold is passed and if the request is still not successful, then the ReverseProxyHandler will return a `502 Bad Gateway` response.

Default: `5` retries.

`"delay"`: *duration, optional*

The delay to wait before retrying the request.

After a failure to send the request, if the number of retries is below the threshold, a new attempt is scheduled with the executor service after this delay.

Default: `10 seconds`.

For information about supported formats for `duration`, see duration.

The following example configures the handler to retry the request only once, after a 1-minute delay:

```
{
  "retries": {
    "count": 1,
    "delay": "1 minute"
  }
}
```

The following example configures the handler to retry the request at most 20 times, every second:

```
{
  "retries": {
    "count": 20,
    "delay": "1 second"
  }
}
```

The following example configures the handler to retry the request 5 times, every 10 seconds (default values), with a dedicated executor:

```
{
  "retries": {
    "executor": {
      "type": "ScheduledExecutorService",
      "config": {
        "corePoolSize": 20
      }
    }
  }
}
```

`"websocket"`: *object, optional*

Enable upgrade from HTTP or HTTPS protocol to WebSocket protocol.

For more information and an example of proxying WebSocket traffic, see "*Proxying WebSocket Traffic*" in the *Gateway Guide*

> **Important**
>
> When IG is running in the Jetty application container, it cannot proxy WebSocket traffic.

```
{
  "websocket": {
    "enabled": boolean,
    "connectionTimeout": duration string,
    "soTimeout": duration string,
    "numberOfSelectors": number,
    "tls": TlsOptions reference
  }
}
```

For more information, see The WebSocket Protocol, RFC6455.

`"enabled"`: *boolean, optional*

Enable upgrade from HTTP protocol and HTTPS protocol to WebSocket protocol.

Default: `false`

`"connectionTimeout"`: *duration string, optional*

The maximum time allowed to establish a WebSocket connection.

Default: The value of handler's main `connectionTimeout`.

For information about supported formats for `duration`, see duration.

`"soTimeout"`: *duration string, optional*

The time after which stalled connections are destroyed.

> **Tip**
>
> If there can be long delays between messages, consider increasing this value. Alternatively, keep the connection active by using WebSocket ping messages in your application.

Default: The value of handler's main `soTimeout`.

For information about supported formats for `duration`, see duration.

`"numberOfSelectors"`: *number, optional*

The maximum number of worker threads.

In deployments with a high number of simultaneous connections, consider increasing the value of this property.

Default: `2`

tls: *TlsOptions reference, optional*

Configure options for connections to TLS-protected endpoints, based on a TlsOptions(5) configuration. Define a TlsOptions object inline or in the heap.

Default: Use TlsOptions defined for the handler

More Information

org.forgerock.openig.handler.ReverseProxyHandler

Name

Route — JSON configuration files to handle requests and their context

Description

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

- Manually into the filesystem.

- Through Common REST commands. For information, see " Creating and Editing Routes Through Common REST " in the *Gateway Guide*.

- Through Studio. For information, see "*Configuring Routes With Studio*" in the *Getting Started Guide*.

Every route must call a handler to process requests and produce responses to requests.

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 configuration. 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 IG configuration.

For examples of route configurations see "*Configuring Routers and Routes*" in the *Gateway Guide*.

Usage

```
{
    "handler": Handler reference or inline Handler declaration,
    "heap": [ configuration object, ... ],
    "condition": runtime expression<boolean>,
    "name": string,
    "secrets": configuration object,
    "session": Session reference
}
```

Properties

`"handler"`: *Handler reference, required*

For this route, dispatch the request to this handler.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

See also Handlers.

`"heap"`: *array of configuration objects, optional*

Heap object configuration for objects local to this route.

Objects referenced but not defined here are inherited from the parent.

You can omit an empty array. If you only have one object in the heap, you can inline it as the handler value.

See also Heap Objects(5).

`"condition"`: *runtime expression<boolean>, optional*

If the expression evaluates to `true`, the request is dispatched to the route. If no condition is specified, the route accepts any request.

An external request can never match a condition that uses the reserved administrative route. Therefore, routes that use these conditions are effectively ignored. For example, if `/openig` is the administrative route, a route with the following condition would effectively be ignored: `${matches(request.uri.path, '^/openig/my/path')}`.

Default: No condition is specified.

For some example route conditions, see "Setting Route Conditions" in the *Gateway Guide*.

See also "Reserved Routes" and Expressions(5).

`"name"`: *string, optional*

Name for the route.

The Router uses the `name` property to order the routes in the configuration. If the route does not have a `name` property, the Router uses the route ID.

The route ID is managed as follows:

- When you add a route manually to the routes folder, the route ID is the value of the `_id` field. If there is no `_id` field, the route ID is the filename of the added route.

- When you add a route through the Common REST endpoint, the route ID is the value of the mandatory `_id` field.

- When you add a route through Studio, you can edit the default route ID.

> **Caution**
>
> The filename of a route cannot be `default.json`, and the route's `name` property and route ID cannot be `default`.

Default: route ID

"secrets": *configuration object, optional*

An object that configures an inline array of one or more secret stores, as defined in Secrets.

```
{
  "secrets": {
    "stores": [ inline secret store declaration, ... ]
  }
}
```

"session": *Session reference, optional*

Session storage implementation used by this route, such as a JwtSession as described in JwtSession(5).

Provide either the name of a session storage object defined in the heap, or an inline session storage configuration object.

Default: do not change the session storage implementation for session.

Metrics

This section describes the metrics that Route publishes to the Prometheus Scrape Endpoint and the Common REST Monitoring Endpoint.

Route Metrics at the Prometheus Scrape Endpoint

Route metrics at the Prometheus Scrape Endpoint have the following labels:

- name: Route name, for example, My Route.

  If the router was declared with a default handler, then its metrics are published through the route named default.

- route: Route identifier, for example, my-route.

- router: Fully qualified name of the router, for example, gateway.main-router.

The following table summarizes the recorded metrics:

Name	Type [a]	Description
ig_route_request_active	Gauge	Number of requests being processed.
ig_route_request_total	Counter	Number of requests processed by the router or route since it was deployed.
ig_route_response_error_total	Counter	Number of responses that threw an exception.
ig_route_response_null_total	Counter	Number of responses that were not handled by IG.

Name	Type [a]	Description
`ig_route_response_status_total`	Counter	Number of responses by HTTP status code family. The `family` label depends on the HTTP status code: • Informational (1xx) • Successful (2xx) • Redirection (3xx) • Client_error (4xx) • Server_error (5xx) • Unknown (status code >= 600)
`ig_route_response_time`	Summary	A summary of response time observations.

[a]As described in Monitoring Types(5)

For more information about the the Prometheus Scrape Endpoint, see "Prometheus Scrape Endpoint".

Route Metrics at the Common REST Monitoring Endpoint

Route metrics at the Common REST Monitoring Endpoint are published with an `_id` in the following pattern:

• *`heap.router-name.route.route-name.metric`*

The following table summarizes the recorded metrics:

Name	Type [a]	Description
`request`	Counter	Number of requests processed by the router or route since it was deployed.
`request.active`	Gauge	Number of requests being processed by the router or route at this moment.
`response.error`	Counter	Number of responses that threw an exception.
`response.null`	Counter	Number of responses that were not handled by IG.
`response.status.client_error`	Counter	Number of responses with an HTTP status code `400-499`, indicating client error.
`response.status.informational`	Counter	Number of responses with an HTTP status code `100-199`, indicating that they are provisional responses.
`response.status.redirection`	Counter	Number of responses with an HTTP status code `300-399`, indicating a redirect.

Name	Type [a]	Description
`response.status.server_error`	Counter	Number of responses with an HTTP status code 500-599, indicating server error.
`response.status.successful`	Counter	Number of responses with an HTTP status code 200-299, indicating success.
`response.status.unknown`	Counter	Number of responses with an HTTP status code 600-699, indicating that a request failed completely and was not executed.
`response.time`	Timer	Time-series summary statistics.

[a]As described in Monitoring Types(5)

For more information about the the Common REST Monitoring Endpoint, see "Common REST Monitoring Endpoint".

FORGEROCK

Name

Router — Route request processing to separate routes

Description

A router is a handler that performs the following tasks:

- Defines the routes directory and loads routes into the configuration.

- Depending on the scanning interval, periodically scans the routes directory and updates the IG configuration when routes are added, removed, or changed. The router updates the IG configuration without needing to restart IG or access the route.

- Manages an internal list of routes, where routes are ordered lexicographically by route name. If a route is not named, then the route ID is used instead. For more information, see Route(5).

- Routes requests to the first route in the internal list of routes, whose condition is satisfied.

  Because the list of routes is ordered lexicographically by route name, name your routes with this in mind:

  - If a request satisfies the condition of more than one route, it is routed to the first route in the list whose condition is met.

  - Even if the request matches a later route in the list, it might never reach that route.

  If a request does not satisfy the condition of any route, it is routed to the default handler if one is configured.

The router does not have to know about specific routes in advance - you can configure the router first and then add routes while IG is running.

> **Important**
>
> 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`.

Usage

```
{
    "name": "Router",
    "type": "Router",
    "config": {
        "defaultHandler": Handler reference,
        "directory": expression,
        "scanInterval": duration string or integer
    }
}
```

An alternative value for type is RouterHandler.

Properties

`"defaultHandler"`: *Handler reference, optional*

Handler to use when a request does not satisfy the condition of any route.

Provide either the name of a handler object defined in the heap, or an inline handler configuration object.

Default: If no default route is set either here or in the route configurations, IG aborts the request with an internal error.

See also Handlers.

`"directory"`: *expression, optional*

Base directory from which to load configuration files for routes.

Default: default base directory for route configuration files. For details, see "Installing IG" in the *Gateway Guide*.

The following example changes the file system location to look for routes:

```
{
    "type": "Router",
    "config": {
        "directory": "/path/to/safe/routes",
    }
}
```

> **Important**
>
> If you define a new router in the default base directory, then you must set the directory property to a different directory from the default base directory in order to avoid a circular reference to the new router.

See also Expressions(5).

`"scanInterval"`: *duration string or integer, optional*

Time interval at which IG scans the specified directory for changes to routes. When a route is added, removed, or changed, the router updates the IG configuration without needing to restart IG or access the route.

When an integer is used for the `scanInterval`, the time unit is seconds.

To load routes at startup only, and prevent changes to the configuration if the routes are changed, set the scan interval to `disabled`.

Default: 10 seconds

For information about supported formats for `duration`, see duration.

Metrics

This section describes the metrics that Router publishes to the Prometheus Scrape Endpoint and the Common REST Monitoring Endpoint.

Router Metrics at the Prometheus Scrape Endpoint

Router metrics at the Prometheus Scrape Endpoint have the following labels:

- `fully_qualified_name`: Fully qualified name of the router, for example, `gateway.main-router`.

- `heap`: Name of the heap in which this router is declared, for example, `gateway`.

- `name`: Simple name declared in router configuration, for example, `main-router`.

The following table summarizes the recorded metrics:

Name	Type [a]	Description
`ig_router_deployed_routes`	Gauge	Number of routes deployed in the configuration.

[a]As described in Monitoring Types(5)

For more information about the the Prometheus Scrape Endpoint, see "Prometheus Scrape Endpoint".

Router Metrics at the Common REST Monitoring Endpoint

Router metrics at the Common REST Monitoring Endpoint are JSON objects, with the following form:

- `[heap name].[router name].deployed-routes`

The following table summarizes the recorded metrics:

Name	Type [a]	Description
`deployed-routes`	Gauge	Number of routes deployed in the configuration.

[a]As described in Monitoring Types(5)

For more information about the the Common REST Monitoring Endpoint, see "Common REST Monitoring Endpoint".

More Information

org.forgerock.openig.handler.router.RouterHandler

Name
SamlFederationHandler — play the role of SAML 2.0 Service Provider

Description

A handler to play the role of SAML 2.0 Service Provider (SP).

> **Note**
>
> This handler does not support filtering. Specifically, do not use this as the handler for a Chain, which can include filters.
>
> More generally, do not use this handler when its use depends on something in the response. The response can be handled independently of IG, and can be `null` when control returns to IG. For example, do not use this handler in a `SequenceHandler` where the `postcondition` depends on the response.

Usage

```
{
    "name": string,
    "type": "SamlFederationHandler",
    "config": {
        "assertionMapping": object,
        "redirectURI": string,
        "assertionConsumerEndpoint": string,
        "authnContext": string,
        "authnContextDelimiter": string,
        "logoutURI": string,
        "sessionIndexMapping": string,
        "singleLogoutEndpoint": string,
        "singleLogoutEndpointSoap": string,
        "SPinitiatedSLOEndpoint": string,
        "SPinitiatedSSOEndpoint": string,
        "subjectMapping": string
    }
}
```

Properties

`"assertionMapping"`: *object, required*

The assertionMapping defines how to transform attributes from the incoming assertion to attribute value pairs in IG.

Each entry in the `assertionMapping` object has the form *localName: incomingName*, where:

- *localName* is the name of the attribute set in the session

- *incomingName* is the name of the attribute set in the incoming assertion

The following example is an assertionMapping object:

```
{
    "username": "mail",
    "password": "mailPassword"
}
```

If the incoming assertion contains the statement:

```
mail = george@example.com
```

```
mailPassword = C0stanza
```

Then the following values are set in the session:

```
username = george@example.com
```

```
password = C0stanza
```

For this to work, edit the <Attribute name="attributeMap"> element in the SP extended metadata file, `$HOME/.openig/SAML/sp-extended.xml`, so that it matches the assertion mapping configured in the SAML 2.0 Identity Provider (IDP) metadata.

Because the dot character (`.`) serves as a query separator in expressions, do not use dot characters in the *localName*.

To prevent different handlers from overwriting each others' data, use unique *localName* settings when protecting multiple service providers.

`"redirectURI"`: *string, required*

Set this to the page that the filter used to HTTP POST a login form recognizes as the login page for the protected application.

This is how IG and the Federation component work together to provide SSO. When IG detects the login page of the protected application, it redirects to the Federation component. Once the Federation handler validates the SAML exchanges with the IDP, and sets the required session attributes, it redirects back to the login page of the protected application. This allows the filter used to HTTP POST a login form to finish the job by creating a login form to post to the application based on the credentials retrieved from the session attributes.

`"assertionConsumerEndpoint"`: *string, optional*

Default: `fedletapplication` (same as the Fedlet)

If you modify this attribute you must change the metadata to match.

`"authnContext"`: *string, optional*

Name of the session field to hold the value of the authentication context. Because the dot character (`.`) serves as a query separator in expressions, do not use dot characters in the field name.

Use this setting when protecting multiple service providers, as the different configurations must not map their data into the same fields of `session`. Otherwise different handlers can overwrite each others' data.

As an example, if you set `"authnContext": "myAuthnContext"`, then IG sets `session.myAuthnContext` to the authentication context specified in the assertion. When the authentication context is password over protected transport, then this results in the session containing `"myAuthnContext":` `"urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport"`.

Default: map to `session.authnContext`

"authnContextDelimiter": *string, optional*

The authentication context delimiter used when there are multiple authentication contexts in the assertion.

Default: `|`

"logoutURI": *string, optional*

Set this to the URI to visit after the user is logged out of the protected application.

You only need to set this if the application uses the single logout feature of the Identity Provider.

"sessionIndexMapping": *string, optional*

Name of the session field to hold the value of the session index. Because the dot character (`.`) serves as a query separator in expressions, do not use dot characters in the field name.

Use this setting when protecting multiple service providers, as the different configurations must not map their data into the same fields of `session`. Otherwise different handlers can overwrite each others' data.

As an example, if you set `"sessionIndexMapping": "mySessionIndex"`, then IG sets `session.mySessionIndex` to the session index specified in the assertion. This results in the session containing something like `"mySessionIndex": "s24ccbbffe2bfd761c32d42e1b7a9f60ea618f9801"`.

Default: map to `session.sessionIndex`

"singleLogoutEndpoint": *string, optional*

Default: `fedletSLORedirect` (same as the Fedlet)

If you modify this attribute you must change the metadata to match.

"singleLogoutEndpointSoap": *string, optional*

Default: `fedletSloSoap` (same as the Fedlet)

If you modify this attribute you must change the metadata to match.

"SPinitiatedSLOEndpoint": *string, optional*

Default: `SPInitiatedSLO`

If you modify this attribute you must change the metadata to match.

"SPinitiatedSSOEndpoint": *string, optional*

Default: `SPInitiatedSSO`

If you modify this attribute you must change the metadata to match.

"subjectMapping": *string, optional*

Name of the session field to hold the value of the subject name. Because the dot character (`.`) serves as a query separator in expressions, do not use dot characters in the field name.

Use this setting when protecting multiple service providers, as the different configurations must not map their data into the same fields of `session`. Otherwise different handlers can overwrite each others' data.

As an example, if you set `"subjectMapping": "mySubjectName"`, then IG sets `session.mySubjectName` to the subject name specified in the assertion. If the subject name is an opaque identifier, then this results in the session containing something like `"mySubjectName": "vtOk+APj1s9Rr4yCka6V9pGUuzuL"`.

Default: map to `session.subjectName`

Example

The following sample configuration is corresponds to a scenario where IG receives a SAML 2.0 assertion from the IDP, and then logs the user in to the protected application using the username and password from the assertion:

```
{
    "name": "SamlFederationHandler",
    "type": "SamlFederationHandler",
    "config": {
        "assertionMapping": {
            "username": "mail",
            "password": "mailPassword"
        },
        "redirectURI": "/login",
        "logoutURI": "/logout"
    }
}
```

More Information

org.forgerock.openig.handler.saml.SamlFederationHandler

Name

ScriptableHandler — handle a request by using a Groovy script

Description

Handles a request by using a script.

Scripts must return either a Promise<Response, NeverThrowsException> or a Response.

This section describes the usage of ScriptableHandler. For information about script properties, available global objects, and automatically imported classes, see Scripts(5).

Usage

```
{
    "name": string,
    "type": "ScriptableHandler",
    "config": {
        "type": string,
        "file": expression,                    // Use either "file"
        "source": string or array of strings   // or "source", but not both
        "args": object,
        "clientHandler": Handler reference
    }
}
```

Properties

For information about properties for ScriptableHandler, see Scripts(5).

More Information

org.forgerock.openig.handler.ScriptableHandler

Name

SequenceHandler — process request through sequence of handlers

Description

Processes a request through a sequence of handlers. This allows multi-request processing such as retrieving a form, extracting form content (for example, nonce) and submitting in a subsequent request. Each `handler` in the `bindings` is dispatched to in order; the binding `postcondition` determines if the sequence should continue.

Usage

```
{
    "name": string,
    "type": "SequenceHandler",
    "config": {
        "bindings": [
            {
                "handler": Handler reference,
                "postcondition": runtime expression<boolean>
            }
        ]
    }
}
```

Properties

`"bindings"`: *array of objects, required*

A list of bindings of handler and postcondition to determine that sequence continues.

`"handler"`: *Handler reference, required*

Dispatch to this handler.

Either the name of the handler heap object to dispatch to, or an inline Handler configuration object.

See also Handlers.

`"postcondition"`: *runtime expression<boolean>, optional*

If the expression evaluates to `true`, the sequence continues. If a condition is not specified, the sequence continues unconditionally.

Default: No condition is specified.

See also Expressions(5).

More Information

org.forgerock.openig.handler.SequenceHandler

Name

StaticResponseHandler — create static response to a request

Description

Creates a static response to a request.

Usage

```
{
    "name": string,
    "type": "StaticResponseHandler",
    "config": {
        "status": number,
        "reason": string,
        "version": string,
        "headers": {
            name: [ runtime expression<string>, ... ], ...
        },
        "entity": runtime expression<string>
    }
}
```

Properties

"status": *number, required*

The response status code (for example, 200).

"reason": *string, optional*

The response status reason (for example, `"OK"`).

"version": *string, optional*

Protocol version. Default: `"HTTP/1.1"`.

"headers": *array of objects, optional*

The `name` specifies the header name. Its value is an array of runtime expressions to evaluate as header values.

"entity": *runtime expression<string>, optional*

The message entity to include in the response.

If present, it must conform to the `Content-Type` header and set the content length header automatically.

> **Important**
>
> Attackers during reconnaissance can use response messages to identify information about a deployment. For security, limit the amount of information in messages, and avoid using words that help identify IG.

Default: No message entity.

Example

```
{
    "name": "ErrorHandler",
    "type":"StaticResponseHandler",
    "config": {
        "status": 500,
        "reason": "Error",
        "entity": "<html>
                    <h2>Epic #FAIL</h2>
                    </html>",
        "headers": {
          "content-type": [
            "text/html"
          ]
        }
    }
}
```

More Information

org.forgerock.openig.handler.StaticResponseHandler

Filters

Filter objects intercept requests and responses during processing.

Table of Contents

Name

AssignmentFilter — assign values to a target if a condition is met

Description

Verifies that a specified condition is met. If the condition is met or if no condition is specified, the value is assigned to the target. Values can be assigned before the request is handled and after the response is handled.

Usage

```
{
    "name": string,
    "type": "AssignmentFilter",
    "config": {
        "onRequest": [
            {
                "condition": runtime expression<boolean>,
                "target": lvalue-expression,
                "value": runtime expression
            }, ...
        ],
        "onResponse": [
            {
                "condition": runtime expression<boolean>,
                "target": lvalue-expression,
                "value": runtime expression
            }, ...
        ]
    }
}
```

Properties

"onRequest": *array of objects, optional*

Defines a list of assignment bindings to evaluate before the request is handled.

"onResponse": *array of objects, optional*

Defines a list of assignment bindings to evaluate after the response is handled.

"condition": *runtime expression<boolean>, optional*

If the expression evaluates `true`, the value is assigned to the target. If no condition is specified, the value is assigned to the target unconditionally.

Default: No condition is specified.

See also Expressions(5).

FORGEROCK

`"target"`: *lvalue-expression, required*

Expression that yields the target object whose value is to be set.

See also Expressions(5).

`"value"`: *runtime expression, optional*

The value to be set in the target. The value can be a string, information from the context, or even a whole map of information.

See also Expressions(5).

Example

This is an example of how to capture credentials and store them in the IG session during a login request. Notice that the credentials are captured on the request but are not marked as valid until the response returns a positive 302. The credentials could then be used to log a user in to a different application:

```
{
    "name": "PortalLoginCaptureFilter",
    "type": "AssignmentFilter",
    "config": {
        "onRequest": [
            {
                "target": "${session.authUsername}",
                "value": "${request.form['username'][0]}",
            },
            {
                "target": "${session.authPassword}",
                "value": "${request.form['password'][0]}",
            },
            {
                "comment": "Authentication has not yet been confirmed.",
                "target": "${session.authConfirmed}",
                "value": "${false}",
            }
        ],
        "onResponse": [
            {
                "condition": "${response.status.code == 302}",
                "target": "${session.authConfirmed}",
                "value": "${true}",
            }
        ]
    }
}
```

More Information

org.forgerock.openig.filter.AssignmentFilter

Name

CapturedUserPasswordFilter — retrieve an AM password, decrypt it, and expose it in the context

Description

This filter makes an AM password available to IG in the following steps:

- Checks for the presence of the SessionInfoContext context, at `${contexts.amSession}`.

  - If the context is not present, or if `sunIdentityUserPassword` is `null`, the CapturedUserPasswordFilter collects session info and properties from AM.

    When `sunIdentityUserPassword` is included in the AM session properties whitelist, the password is included in the collected information.

  - If the context is present and `sunIdentityUserPassword` is not `null`, the CapturedUserPasswordFilter uses that value for the password.

- The CapturedUserPasswordFilter decrypts the password and stores it in the CapturedUserPasswordContext, at `${contexts.capturedPassword}`.

Supported with AM 5 and later versions, and with AM 6 and later versions when the `AES` keyType is used to decrypt the password.

Usage

```
{
  "name": string,
  "type": "CapturedUserPasswordFilter",
  "config": {
    "amService": AmService reference,
    "keySecretId": configuration expression<secret-id>,
    "keyType": configuration expression<string>,
    "ssoToken": runtime expression<string>
  }
}
```

Properties

`"amService"`: *AmService reference, required*

The AmService heap object to use for the following properties:

- `agent`, the credentials of an AM Java agent that acts on behalf of IG to authenticate with AM, get the user profile, and communicate WebSocket notifications from AM to IG.

- `url`, the URL of an AM service to use for session token validation and authentication.

- `amHandler`, the handler to use when communicating with AM to validate the token in the incoming request.

- `realm`: Realm of the AM Java agent that acts on behalf of IG.

- `version`: The version of the AM server.

This filter is compatible with AM version 5.5 or higher. If `version` is not set, the default version is AM 5 and an error is thrown.

See also, AmService(5).

"keySecretId": *configuration expression<secret-id>, required*

The secret ID for the key required decrypt the AM password.

For information about supported formats for `secret-id`, see secret-id.

"keyType": *configuration expression<enumeration>, required*

Algorithm to decrypt the AM password. Use one of the following values:

- `DES` for DES/ECB/NoPadding

- `AES` AES for JWT-based AES_128_CBC_HMAC_SHA_256 encryption, available from AM 6.

  For more information, see AES_128_CBC_HMAC_SHA_256 in the IETF *JSON Web Algorithms*.

Default: `DES`

"ssoToken": *runtime expression<string>, required*

Location of the AM SSO token.

Default: `${request.cookies['AmService-ssoTokenHeader'][0].value}`, where `AmService-ssoTokenHeader` is the name of the header or cookie where the AmService expects to find SSO tokens.

Examples

The following example route is used to get login credentials from AM in "*Getting Login Credentials From AM*" in the *Gateway Guide*.

```
{
  "name": "04-replay",
  "condition": "${matches(request.uri.path, '^/replay')}",
  "heap": [
    {
      "name": "AmService-1",
      "type": "AmService",
      "config": {
        "agent" : {
          "username" : "ig_agent",
          "passwordSecretId" : "agent.secret.id"
        },
        "url": "http://openam.example.com:8088/openam/",
        "version": "6.5"
```

```
        }
      },
      {
        "name": "CapturedUserPasswordFilter",
        "type": "CapturedUserPasswordFilter",
        "config": {
          "ssoToken": "${contexts.ssoToken.value}",
          "keySecretId": "DESKEY",
          "amService": "AmService-1"
        }
      }
    ],
    "handler": {
      "type": "Chain",
      "config": {
        "filters": [
          {
            "type": "SingleSignOnFilter",
            "config": {
              "amService": "AmService-1"
            }
          },
          {
            "type": "PasswordReplayFilter",
            "config": {
              "loginPage": "${true}",
              "credentials": "CapturedUserPasswordFilter",
              "request": {
                "method": "POST",
                "uri": "http://app.example.com:8081/login",
                "form": {
                  "username": [
                    "${contexts.ssoToken.info.uid}"
                  ],
                  "password": [
                    "${contexts.capturedPassword.value}"
                  ]
                }
              }
            }
          }
        ],
        "handler": "ReverseProxyHandler"
      }
    }
  }
}
```

More Information

org.forgerock.openig.openam.CapturedUserPasswordFilter

org.forgerock.openig.openam.CapturedUserPasswordContext

CapturedUserPasswordContext(5)

SessionInfoFilter(5)

Name

ConditionalFilter — dispatch a request to a delegate filter if a condition is met

Description

Verifies that a specified condition is met. If the condition is met, the request is dispatched to a delegate Filter. Otherwise, the delegate Filter is skipped.

Use `ConditionalFilter` to easily use or skip a Filter depending on whether a condition is met. To easily use or skip a set of Filters, use a `ChainOfFilters` as the delegate Filter and define a set of Filters. For information, see ChainOfFilters(5).

Usage

```
{
    "type": "ConditionalFilter",
    "config": {
        "condition": runtime expression<boolean>,
        "delegate" : filter reference
    }
}
```

Properties

`"condition"`: *runtime expression<boolean>, required*

If the expression evaluates to `true`, the request is dispatched to the delegate Filter. Otherwise the delegate Filter is skipped.

See also Expressions.

`"delegate"`: *filter reference, required*

Filter to treat the request if the condition expression evaluates as `true`.

Provide an inline Filter configuration object, or the name of a Filter object defined in the heap.

See also Filters.

Example

The following example tests whether a request finishes with `.js` or `.jpg`:

```
{
  "type": "Chain",
  "config": {
    "filters": [{
      "type": "ConditionalFilter",
      "config": {
        "condition": "${not matches ((request.uri.path, '.js$') or (request.uri.path, '.jpg$'))}",
        "delegate": "mySingleSignOnFilter"
      }
    }],
    "handler": "ReverseProxyHandler"
  }
}
```

If the request is to access a .js file or .jpg file, it skips the delegate SingleSignOnFilter filter declared in the heap, and passes straight to the ReverseProxyHandler.

If the request is to access another type of resource, it must pass through the delegate SingleSignOnFilter for authentication with AM before it can pass to the ReverseProxyHandler.

More Information

org.forgerock.http.filter.ConditionalFilter

Name

ConditionEnforcementFilter — continue the chain of execution if a condition is met

Description

Verifies that a specified condition is met. If the condition is met, the request continues to be executed. Otherwise, the request is referred to a failure handler, or IG returns 403 Forbidden and the request is stopped.

Usage

```
{
    "type": "ConditionEnforcementFilter",
    "config": {
        "condition": runtime expression<boolean>,
        "failureHandler": handler reference
    }
}
```

Properties

`"condition"`: *runtime expression<boolean>, required*

> If the expression evaluates to `true`, the request continues to be executed.
>
> See also Expressions.

`"failureHandler"`: *handler reference, optional*

> Handler to treat the request if the condition expression evaluates as `false`.
>
> Provide an inline handler configuration object, or the name of a handler object that is defined in the heap.
>
> See also Handlers.
>
> Default: HTTP 403 Forbidden, the request stops being executed.

Example

The following example tests whether a request contains a session username. If it does, the request continues to be executed. Otherwise, the request is dispatched to the `ConditionFailedHandler` failure handler.

```
{
    "name": "UsernameEnforcementFilter",
    "type": "ConditionEnforcementFilter",
    "config": {
        "condition": "${not empty (session.username)}",
        "failureHandler": "ConditionFailedHandler"
    }
}
```

More Information

org.forgerock.openig.filter.ConditionEnforcementFilter

FORGEROCK

Name

ChainOfFilters — dispatch the request to ordered list of filters

Description

Dispatches a request to an ordered list of filters. 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 used as a filter inside a chain.

Unlike `Chain`, `ChainOfFilters` does not finish by dispatching the request to a handler. For more information, see Chain(5).

Usage

```
{
    "name": string,
    "type": "ChainOfFilters",
    "config": {
        "filters": [ Filter reference, ... ]
    }
}
```

Properties

`"filters"`: *array of filter references, required*

An array of names of filter objects defined in the heap, and inline filter configuration objects.

The chain dispatches the request to these filters in the order they appear in the array.

See also Filters.

Example

```
{
    "name": "MyChainOfFilters",
    "type": "ChainOfFilters",
    "config": {
        "filters": [ "Filter1", "Filter2" ]
    }
}
```

More Information

org.forgerock.openig.filter.ChainFilterHeaplet

Name

CookieFilter — manage, suppress, relay cookies

Description

Manages, suppresses, and relays cookies as follows:

- **Manage**, to store cookies from the protected application in the IG session, and include them in later requests.

  For requests with a `Cookie` header, managed cookies are removed so that protected applications don't see them.

  For responses with a `Set-Cookie` header, managed cookies are removed and then added in a `Cookie` header to the next request that goes through that filter.

  Manage is the default action, and a common choice to manage cookies originating from the protected application.

- **Suppress**, to remove cookies from the request and response. Use this option to hide domain cookies, such as `iPlanetDirectoryPro`, that are used by IG but are not usually used by protected applications.

- **Relay**, to transmit cookies freely from the user agent to the remote server, and vice versa.

If a cookie does not appear in one of the three action parameters, then the default action is performed, controlled by setting the `defaultAction` parameter. If unspecified, the default action is to manage all cookies. In the event a cookie appears in more than one configuration parameter, then it will be selected in the order of precedence: managed, suppressed, relayed.

Usage

```
{
    "name": string,
    "type": "CookieFilter",
    "config": {
        "managed": [ string, ... ],
        "suppressed": [ string, ... ],
        "relayed": [ string, ... ],
        "defaultAction": string
    }
}
```

Properties

`"managed"`: *array of strings, optional*

A list of the names of cookies to be managed.

FORGEROCK

"suppressed": *array of strings, optional*

A list of the names of cookies to be suppressed.

"relayed": *array of strings, optional*

A list of the names of cookies to be relayed.

"defaultAction": *enumeration, optional*

Action to perform for cookies that do not match an action set. Set to `"MANAGE"`, `"RELAY"`, or `"SUPPRESS"`. Default: `"MANAGE"`.

More Information

org.forgerock.openig.filter.CookieFilter

Name

CrossDomainSingleSignOnFilter — allow SSO with an AM instance deployed on a different cookie domain

Description

When IG and AM are running in the same domain, the SingleSignOnFilter can be used for SSO. When IG and AM are running in different domains, AM cookies are not visible to IG because of the same-origin policy. The CrossDomainSingleSignOnFilter provides a mechanism to push tokens issued by AM to IG running in a different domain.

When this filter processes a request, it injects the CDSSO token, the session user ID, and the full claims set into the CdSsoContext(5). Should an error occur during authentication then the error details are captured in a CdSsoFailureContext(5).

For an example of how to configure CDSSO in AM and IG, and information about the flow of data between AM, IG, and a protected application, see "About CDSSO Using the CrossDomainSingleSignOnFilter" in the *Gateway Guide*.

Supported with AM 5.5 and later versions.

WebSocket Notifications for Sessions

When WebSocket notifications are set up for sessions, IG receives a notification from AM when a user logs out of AM, or when the AM session is modified, closed, or times out. IG then evicts entries that are related to the event from the `sessionCache`.

For information about setting up WebSocket notifications, using them to clear the session cache, and including them in the server logs, see "WebSocket Notification Service".

Usage

```
{
  "name": string,
  "type": "CrossDomainSingleSignOnFilter",
    "config": {
      "amService": AmService reference,
      "authCookie": object,
      "redirectEndpoint": runtime expression<uri>,
      "failureHandler": Handler reference
  }
}
```

Properties

`"amService"`: *AmService reference, required*

> The AmService heap object to use for the following properties:

- `agent`, the credentials of an AM Java agent that acts on behalf of IG to authenticate with AM, get the user profile, and communicate WebSocket notifications from AM to IG.

- `url`, the URL of an AM service to use for session token validation and authentication.

- `ssoTokenHeader`, the name of the cookie that contains the session token created by AM.

- `amHandler`, the handler to use when communicating with AM to validate the token in the incoming request.

- `realm`: Realm of the AM Java agent that acts on behalf of IG.

- `version`: The version of the AM server.

See also, AmService(5).

`"authCookie"`: *object, optional*

The configuration of the cookie used to store the authentication.

```
{
  "authCookie": {
    "name": configuration expression<string>,
    "domain": : configuration expression<string>,
    "httpOnly": configuration expression<boolean>,
    "path": : configuration expression<string>,
    "sameSite": configuration expression<enumeration>,
    "secure": configuration expression<boolean>
  }
}
```

`"name"`: *configuration expression<string>, optional*

Name of the cookie containing the authentication token from AM.

Default: `ig-token-cookie`

`"domain"`: *configuration expression<string>, optional*

Domain to which the cookie applies.

Default: Domain of the original request.

`"httpOnly"`: *configuration expression<boolean>, optional*

Flag to mitigate the risk of client-side scripts accessing protected cookies.

Default: `True`

`"path"`: *configuration expression<string>, optional*

Path protected by this authentication.

Default: `/`

"sameSite": *configuration expression<enumeration>, optional*

Manage the circumstances in which a cookie is sent to the server. Use one of the following options to reduce the risk of cross-site request forgery (CSRF) attacks:

- `STRICT`: Send the cookie only if the request was initiated from the cookie domain.

- `LAX`: Send the cookie only with GET requests in a first-party context, where the URL in the address bar matches the cookie domain.

The value is not case-sensitive.

Default: Null; send the cookie whenever a request is made to the cookie domain.

"secure": *configuration expression<boolean>, optional*

Flag to limit the scope of the cookie to secure channels.

Default: `False`

"redirectEndpoint": *runtime expression<uri>, required*

The URI to the endpoint for CDSSO.

This URI must also be configured in AM exactly as it is configured here. For more information about configuring policy agents in AM, see Implementing Cross-Domain Single Sign-On in the Access Management *Authentication and Single Sign-On Guide.*.

To make sure that the redirect is routed back to this filter, where the authentication can be validated, include the endpoint in the route condition in one of the following ways:

- As a sub-path of the condition path.

  For example, use the following route condition with the following endpoint:

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

  ```
  "redirectEndpoint": "/home/cdsso/callback"
  ```

- To match the route condition.

  For example, use the following route condition with the following endpoint:

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

  ```
  "redirectEndpoint": "/home/cdsso"
  ```

  > **Note**
  >
  > With this route condition, all POST requests on the condition path are treated as AM CDSSO callbacks. Any POST requests that aren't the result of an AM CDSSO callback will fail.

- As a specific path that is not related to the condition path.

To make sure that the redirect is routed back to this filter, include the redirectEndpoint as a path in the filter condition.

For example, use the following route condition with the following endpoint:

```
"condition": "${matches(request.uri.path, '^/home/cdsso/redirect') || matches(request.uri.path, '^/
ig/cdssoRedirectUri')}"
```

```
"redirectEndpoint": "/ig/cdssoRedirectUri"
```

`"failureHandler"`: *Handler reference, optional*

Failure handler to be invoked should an error occur during authentication.

Should an error occur during authentication, a CdSsoFailureContext(5) is populated with details of the error and any associated `Throwable`. This is available to the failure handler so that it can respond appropriately.

Be aware that the failure handler does not itself play a role in user authentication. It is only invoked if there is a problem that prevents user authentication from taking place.

A number of circumstances may cause the failure handler to be invoked, including:

- The redirect endpoint is invalid.

- The redirect endpoint is invoked without a valid CDSSO token.

- The redirect endpoint is invoked inappropriately.

- An error was reported by AM during authentication.

Should no failure handler be configured then the default failure handler is used.

See alsoHandlers.

Default: HTTP 200 Ok. The response entity contains details of the error.

Example

In the following example from "*Single Sign-On and Cross-Domain Single Sign-On*" in the *Gateway Guide*, IG uses authentication from AM on a different domain to process a request:

```
{
  "name": "cdsso",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/cdsso')}",
  "heap": [
    {
      "name": "AmService-1",
      "type": "AmService",
```

```
      "config": {
        "url": "http://openam.example.com:8088/openam",
        "realm": "/",
        "ssoTokenHeader": "iPlanetDirectoryPro",
        "version": "6.5",
        "agent": {
          "username": "ig_agent_cdsso",
          "password": "password"
        },
        "sessionCache": {
          "enabled": false
        }
      }
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "name": "CrossDomainSingleSignOnFilter-1",
          "type": "CrossDomainSingleSignOnFilter",
          "config": {
            "redirectEndpoint": "/home/cdsso/redirect",
            "authCookie": {
              "path": "/home",
              "name": "ig-token-cookie"
            },
            "amService": "AmService-1"
          }
        }
      ],
      "handler": "ReverseProxyHandler"
    }
  }
}
```

More Information

org.forgerock.openig.openam.SingleSignOnFilter

CdSsoContext(5)

CdSsoFailureContext(5)

SsoTokenContext(5)

Name
CryptoHeaderFilter — encrypt, decrypt headers

Description

Encrypts or decrypts headers in a request or response.

Usage

```
{
    "name": string,
    "type": "CryptoHeaderFilter",
    "config": {
        "messageType": configuration expression<enumeration>,
        "operation": configuration expression<enumeration>,
        "key": configuration expression<string>,
        "algorithm": configuration expression<string>,
        "charset": configuration expression<enumeration>,
        "keyType": configuration expression<enumeration>,
        "headers": [ configuration expression<string>, ... ]
    }
}
```

Properties

"messageType": *configuration expression<enumeration>, required*

Indicates the type of message whose headers to encrypt or decrypt.

Must be one of: "REQUEST", "RESPONSE".

"operation": *configuration expression<enumeration>, required*

Indicates whether to encrypt or decrypt.

Must be one of: "ENCRYPT", "DECRYPT".

"key": *configuration expression<string>, required*

Base64 encoded key value.

See also Expressions(5).

"algorithm": *configuration expression<string>, optional*

The algorithm name, mode, and padding used for encryption and decryption.

CryptoHeaderFilter does not support EC-based encryption. Use other cipher algorithm values given in Java Security Standard Algorithm Names.

For an example of how to generate a DES key, see "To Create a DES Shared Key In IG" in the *Gateway Guide*.

Default: `AES/ECB/PKCS5Padding`

"charset": *configuration expression<enumeration>, optional*

The name of the charset used to encrypt or decrypt values, as described in Class Charset.

Default: `UTF-8`

"keyType": *configuration expression<enumeration>, optional*

Algorithm name for the secret key.

Default: `AES`

"headers": *array of configuration expression<string>, optional*

The names of header fields to encrypt or decrypt.

Default: Do not encrypt or decrypt any headers

Example

```
{
    "name": "DecryptReplayPasswordFilter",
    "type": "CryptoHeaderFilter",
    "config": {
        "messageType": "REQUEST",
        "operation": "DECRYPT",
        "algorithm": "DES/ECB/NoPadding",
        "keyType": "DES",
        "key": "oqdP3DJdE1Q=",
        "headers": [
            "replaypassword"
        ]
    }
}
```

More Information

org.forgerock.openig.filter.CryptoHeaderFilter

Name

CsrfFilter — Prevents CSRF attacks

Description

Prevent Cross Site Request Forgery (CSRF) attacks when using cookie-based authentication, as follows:

- When a session is created or updated for a client, generate a CSRF token as a hash of the session cookie.

- Send the token in a response header to the client, and require the client to provide that header in subsequent requests.

- In subsequent requests, compare the provided token to the generated token.

- If the token is not provided or can't be validated, reject the request and return a valid CSRF token transparently in the response header.

  Rogue websites that attempt CSRF attacks operate in a different website domain to the targeted website. Because of same-origin policy, rogue websites can't access a response from the targeted website, and cannot, therefore, access the CSRF token.

Usage

```
{
  "name": string,
  "type": "CsrfFilter",
  "config": {
    "cookieName": configuration expression<string>,
    "headerName": configuration expression<string>,
    "excludeSafeMethods": configuration expression<boolean>,
    "failureHandler": handler reference
  }
}
```

Properties

"cookieName": *configuration expression<string>, required*

The name of the cookie used to store the session ID.

"headerName": *configuration expression<string>, optional*

The name of the header that carries the CSRF token. The same header is used to create and verify the token.

Default: `X-CSRF-Token`

"excludeSafeMethods": *configuration expression<boolean>, optional*

Whether to exclude GET, HEAD, and OPTION methods from CSRF testing. In most cases, these methods are assumed as safe from CSRF.

Default: `true`

"failureHandler": *handler reference, optional*

Handler to treat the request if the CSRF the token is not provided or can't be validated. Provide an inline handler declaration, or the name of a handler object defined in the heap.

Although IG returns the CSRF token transparently in the response header, this handler cannot access the CSRF token.

Default: Handler that generates `HTTP 403 Forbidden`.

Example

For an example of how to harden protection against CSRF attacks, see "Restricting Access to Studio" in the *Getting Started Guide*.

```
{
  "name": CsrfFilter-1,
  "type": "CsrfFilter",
  "config": {
    "cookieName": "openig-jwt-session",
    "headerName": "X-CSRF-Token",
    "excludeSafeMethods": true
  }
}
```

More Information

org.forgerock.openig.filter.CsrfFilterHeaplet

Name

EntityExtractFilter — extract pattern from message entity

Description

Extracts regular expression patterns from a message entity. The extraction results are stored in a "target" object. For a given matched pattern, as described in Patterns(5), the value stored in the object is either the result of applying its associated pattern template (if specified) or the match result itself otherwise.

Usage

```
{
    "name": string,
    "type": "EntityExtractFilter",
    "config": {
        "messageType": string,
        "charset": string,
        "target": lvalue-expression,
        "bindings": [
            {
                "key": string,
                "pattern": pattern,
                "template": pattern-template
            }, ...
        ]
    }
}
```

Properties

`"messageType"`: *string, required*

The message type to extract patterns from.

Must be one of: REQUEST, RESPONSE.

`"charset"`: *string, optional*

Overrides the character set encoding specified in message.

Default: the message encoding is used.

`"target"`: *lvalue-expression, required*

Expression that yields the target object that contains the extraction results.

The bindings determine what type of object is stored in the target location.

The object stored in the target location is a Map<String, String>. You can then access its content with `${target.key}` or `${target['key']}`.

See also Expressions(5).

`"key"`: *string, required*

Name of element in target object to contain an extraction result.

`"pattern"`: *pattern, required*

The regular expression pattern to find in the entity.

See also Patterns(5).

`"template"`: *pattern-template, optional*

The template to apply to the pattern and store in the named target element.

Default: store the match result itself.

See also Patterns(5).

Examples

Extracts a nonce from the response, which is typically a login page, and sets its value in the attributes context to be used by the downstream filter posting the login form. The nonce value would be accessed using the following expression: `${attributes.extract.wpLoginToken}`.

The pattern finds all matches in the HTTP body of the form `wpLogintoken value="abc"`. Setting the template to `$1` assigns the value `abc` to `attributes.extract.wpLoginToken`:

```
{
    "name": "WikiNoncePageExtract",
    "type": "EntityExtractFilter",
    "config": {
        "messageType": "response",
        "target": "${attributes.extract}",
        "bindings": [
            {
                "key": "wpLoginToken",
                "pattern": "wpLoginToken\"\\s.*value=\"(.*)\"",
                "template": "$1"
            }
        ]
    }
}
```

The following example reads the response looking for the AM login page. When found, it sets `isLoginPage = true` to be used in a SwitchFilter to post the login credentials:

```
{
    "name": "FindLoginPage",
    "type": "EntityExtractFilter",
    "config": {
        "messageType": "response",
        "target": "${attributes.extract}",
        "bindings": [
            {
                "key": "isLoginPage",
                "pattern": "OpenAM\s\(Login\)",
                "template": "true"
            }
        ]
    }
}
```

More Information

org.forgerock.openig.filter.EntityExtractFilter

Name

FileAttributesFilter — retrieve record from a file

Description

Retrieves and exposes a record from a delimiter-separated file. Lookup of the record is performed using a specified `key`, whose `value` is derived from an expression. The resulting record is exposed in an object whose location is specified by the `target` expression. If a matching record cannot be found, then the resulting object is empty.

The retrieval of the record is performed lazily; it does not occur until the first attempt to access a value in the `target`. This defers the overhead of file operations and text processing until a value is first required. This also means that the value expression is not evaluated until the object is first accessed.

Usage

```
{
    "name": string,
    "type": "FileAttributesFilter",
    "config": {
        "file": expression,
        "charset": string,
        "separator": string,
        "header": boolean,
        "fields": [ string, ... ],
        "target": lvalue-expression,
        "key": string,
        "value": runtime expression<string>
    }
}
```

For an example see "Log in With Credentials From a File" in the *Gateway Guide*.

Properties

`"file"`: *expression, required*

The file containing the record to be read.

See also Expressions(5).

`"charset"`: *string, optional*

The character set in which the file is encoded.

Default: `"UTF-8"`.

`"separator"`: *separator identifier string, optional*

The separator character, which is one of the following:

COLON

> Unix-style colon-separated values, with backslash as the escape character.

COMMA

> Comma-separated values, with support for quoted literal strings.

TAB

> Tab separated values, with support for quoted literal strings.

Default: `COMMA`

`"header"`: *boolean, optional*

The setting to treat or not treat the first row of the file as a header row.

When the first row of the file is treated as a header row, the data in that row is disregarded and cannot be returned by a lookup operation.

Default: `true`.

`"fields"`: *array of strings, optional*

A list of keys in the order they appear in a record.

If `fields` is not set, the keys are assigned automatically by the column numbers of the file.

`"target"`: *lvalue-expression, required*

Expression that yields the target object to contain the record.

The target object is a `Map<String, String>`, where the fields are the keys. For example, if the target is `${attributes.file}` and the record has a `username` field and a `password` field mentioned in the fields list, Then you can access the user name as `${attributes.file.username}` and the password as `${attributes.file.password}`.

See also Expressions(5).

`"key"`: *string, required*

The key used for the lookup operation.

`"value"`: *runtime expression<string>, required*

The value to be looked-up in the file.

See also Expressions(5).

More Information

org.forgerock.openig.filter.FileAttributesFilter

Name

HeaderFilter — remove and add headers

Description

Removes headers from and adds headers to request and response messages. Headers are added to any existing headers in the message. To replace a header, remove the header and then add it again.

Usage

```
{
    "name": string,
    "type": "HeaderFilter",
    "config": {
        "messageType": enumeration,
        "remove": [ string, ... ],
        "add": {
            name: [ runtime expression<string>, ... ], ...
        }
    }
}
```

Properties

`"messageType"`: *enumeration, required*

Indicates the type of message to filter headers for. Must be one of: `"REQUEST"`, `"RESPONSE"`.

`"remove"`: *array of strings, optional*

The names of header fields to remove from the message.

`"add"`: *object, optional*

Header fields to add to the message. The header name is specified by `name`. The header values are specified by an array of runtime expressions that evaluate to strings.

Examples

Replace the host header on the incoming request with the value `myhost.com`:

```
{
    "name": "ReplaceHostFilter",
    "type": "HeaderFilter",
    "config": {
        "messageType": "REQUEST",
        "remove": [ "host" ],
        "add": {
            "host": [ "myhost.com" ]
        }
    }
}
```

Add a Set-Cookie header in the response:

```
{
    "name": "SetCookieFilter",
    "type": "HeaderFilter",
    "config": {
        "messageType": "RESPONSE",
        "add": {
            "Set-Cookie": [ "mysession=12345" ]
        }
    }
}
```

Add headers `custom1` and `custom2` to the request:

```
{
    "name": "SetCustomHeaders",
    "type": "HeaderFilter",
    "config": {
        "messageType": "REQUEST",
        "add": {
            "custom1": [ "12345", "6789" ],
            "custom2": [ "abcd" ]
        }
    }
}
```

Add the value of session's policy enforcement token to the `pef_sso_token` header in the response:

```
{
    "type": "HeaderFilter",
    "config": {
        "messageType": "RESPONSE",
        "add": {
            "pef_sso_token": ["${session.pef_token}"]
        }
    }
}
```

More Information

org.forgerock.openig.filter.HeaderFilter

Name

HttpBasicAuthFilter — perform HTTP Basic authentication

Description

Performs authentication through the HTTP Basic authentication scheme. For more information, see RFC 2617.

If challenged for authentication via a `401 Unauthorized` status code by the server, this filter retries the request with credentials attached. After an HTTP authentication challenge is issued from the remote server, all subsequent requests to that remote server that pass through the filter include the user credentials.

If authentication fails (including the case where no credentials are yielded from expressions), then processing is diverted to the specified authentication failure handler.

Usage

```
{
    "name": string,
    "type": "HttpBasicAuthFilter",
    "config": {
        "username": runtime expression<string>,
        "password": runtime expression<string>,
        "failureHandler": Handler reference,
        "cacheHeader": boolean
    }
}
```

Properties

`"username"`: *runtime expression<string>, required*

    The username to supply during authentication.

    See also Expressions(5).

`"password"`: *runtime expression<string>, required*

    The password to supply during authentication.

    See also Expressions(5).

`"failureHandler"`: *Handler reference, required*

    Dispatch to this Handler if authentication fails.

    Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

See also Handlers.

"cacheHeader": *boolean, optional*

Whether or not to cache credentials in the session after the first successful authentication, and then replay those credentials for subsequent authentications in the same session.

With `"cacheHeader": false`, the filter generates the header for each request. This is useful, for example, when users change their passwords during a browser session.

Default: `true`

Example

```
{
    "name": "TomcatAuthenticator",
    "type": "HttpBasicAuthFilter",
    "config": {
        "username": "tomcat",
        "password": "tomcat",
        "failureHandler": "TomcatAuthFailureHandler",
        "cacheHeader": false
    }
}
```

More Information

org.forgerock.openig.filter.HttpBasicAuthFilter

Name

JwtBuilderFilter — pass identity or other runtime information downstream

Description

This filter collects data at runtime, packs it in a JSON Web Token (JWT), and places the resulting JWT into the JwtBuilderContext(5). Use this filter with a HeaderFilter for a flexible way to pass identity or other runtime information to the protected application.

Configure JwtBuilderFilter to create an unsigned JWT, a signed JWT, or a signed then encrypted JWT:

- Sign the JWT so that an application can validate the authenticity of the claims/data. The JWT can be signed with a shared secret or private key, and verified with a shared secret or corresponding public key.

- Encrypt the JWT to reduce the risk of a data breach.

To help with development, the sample app includes a `/jwt` endpoint to display the JWT, verify its signature, and decrypt the JWT.

Usage

```
{
  "name": string,
  "type": "JwtBuilderFilter",
  "config": {
    "template": map or runtime expression<map>,
    "signature" : object
  }
}
```

Properties

`"template"`: *map or runtime expression<map>, required*

A map of information taken from the request or associated contexts in IG.

If this property is a map, the structure must have the format `Map<String, Object>`. The JSON object structure is evaluated as an expression, as follows:

- Where the object is a string: `"e": "${expression}"`

- Where the object is a map that contains a value to evaluate as an expression: `{ "e": { "n": "${nested.expression}" } }`

If this property is an expression, its evaluation must give an object of type `Map<String, Object>`.

See also Expressions(5).

"signature": *object, optional*

A JWT signature to allow the authenticity of the claims/data to be validated. A signed JWT can be encrypted.

```
{
  "signature": {
    "secretId":  configuration expression<secret-id>,
    "algorithm": configuration expression<string>,
    "encryption": object
  }
}
```

"secretId": *configuration expression<secret-id>, required if `signature` is used*

The secret ID of the key used to sign the JWT.

For information about supported formats for `secret-id`, see secret-id.

"algorithm": *expression<string>, optional*

The algorithm with which to sign the JWT.

The following algorithms are supported but not necessarily tested in IG:

- Algorithms described in *Cryptographic Algorithms for Digital Signatures and MACs*.

  For RSASSA-PSS, you must install Bouncy Castle. For information, see *The Legion of the Bouncy Castle*.

- From IG 6.1, `Ed25519` described in *CFRG Elliptic Curve Diffie-Hellman (ECDH) and Signatures*.

Default: RS256

"encryption": *object, optional*

Encrypt the JWT.

```
{
  "encryption": {
    "secretId": configuration expression<secret-id>,
    "algorithm": configuration expression<string>,
    "method": configuration expression<enumeration>
  }
}
```

"secretId": *configuration expression<secret-id>, optional*

The secret ID of the key used to encrypt the JWT. The value is mapped to key `aliases` in KeyStoreSecretStore(5).

For information about supported formats for `secret-id`, see secret-id.

`"algorithm"`: *expression<string>, required*

The algorithm used to encrypt the JWT.

For information about available algorithms, see "*alg*" *(Algorithm) Header Parameter Values for JWE*.

`"method"`: *configuration expression<enumeration>, required*

The method used to encrypt the JWT.

For information about available methods, see "*enc*" *(Encryption Algorithm) Header Parameter Values for JWE*.

Example

Packing Data Into JWT

In this example, the JwtBuilderFilter takes the username and email from the UserProfileContext, and stores them in an unsigned JWT.

Before you start:

- Prepare IG and the sample app as described in "*First Steps*" in the *Getting Started Guide*

- Install and configure AM on http://openam.example.com:8088/openam, using the default configuration.

1. In AM, add a Java agent as described in "To Set Up a Java Agent in AM" in the *Gateway Guide*.

2. Set an environment variable for the Java agent password, and then restart IG:

   ```
   $ export AGENT_SECRET_ID='cGFzc3dvcmQ='
   ```

   The password is retrieved by the default SystemAndEnvSecretStore, and must be base64-encoded.

3. Add the following route to the configuration as `$HOME/.openig/config/routes/jwt-builder.json` (on Windows, `%appdata%\OpenIG\config\routes\jwt-builder.json`):

   ```json
   {
     "name": "jwt",
     "heap": [{
       "name": "AmService-1",
       "type": "AmService",
       "config": {
         "agent" : {
           "username" : "ig_agent",
           "passwordSecretId": "agent.secret.id"
         },
         "url": "http://openam.example.com:8088/openam",
         "version": "6.5"
       }
   ```

```
    }],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [{
        "name": "SingleSignOnFilter",
        "type": "SingleSignOnFilter",
        "config": {
          "amService": "AmService-1"
        }
      }, {
        "name": "UserProfileFilter",
        "type": "UserProfileFilter",
        "config": {
          "username": "${contexts.ssoToken.info.uid}",
          "userProfileService": {
            "type": "UserProfileService",
            "config": {
              "amService": "AmService-1"
            }
          }
        }
      }, {
        "name": "JwtBuilderFilter-1",
        "type": "JwtBuilderFilter",
        "config": {
          "template": {
            "name": "${contexts.userProfile.commonName}",
            "email": "${contexts.userProfile.rawInfo.mail[0]}"
          }
        }
      }, {
        "name": "HeaderFilter-1",
        "type": "HeaderFilter",
        "config": {
          "messageType": "REQUEST",
          "add": {
            "x-openig-user": ["${contexts.jwtBuilder.value}"]
          }
        }
      }],
      "handler": "ReverseProxyHandler"
    }
  },
  "condition": "${matches(request.uri.path, '/jwt$')}",
  "baseURI": "http://app.example.com:8081"
}
```

Notice the following features of the route:

- The route matches requests to `/jwt`.

- An AmService object is available in the heap. It uses the default secrets service configuration to resolve the `passwordSecretId` field.

- If the request does not have a valid iPlanetDirectoryPro cookie, the SingleSignOnFilter redirects the request to authenticate with the AM server declared in the heap.

If the request already has a valid iPlanetDirectoryPro cookie, or after the user authenticates with AM, the SingleSignOnFilter passes the request to the next filter, and stores the cookie value in an SsoTokenContext.

- The UserProfileFilter retrieves the AM user profile data from the SsoTokenContext, and stores it in the UserProfileContext.

- The JwtBuilderFilter takes the username and email from the UserProfileContext, and stores them in a JWT in the JwtBuilderContext.

- The HeaderFilter retrieves the JWT from the JwtBuilderContext, and adds it to the header field `x-openig-user` in the request.

- The ClientHandler passes the request to the sample app, which displays the JWT.

4. Browse to http://openig.example.com:8080/jwt.

5. Log in to AM as user `demo`, password `Ch4ng31t`, or as another user.

   The sample app displays the JWT along with its header and payload, and labels it as unsigned.

Packing Data Into a Signed JWT

This example builds on the previous example to sign the JWT with a key set in the KeyStore. Run the previous example before you try this.

1. Generate a PKCS12 KeyStore that contains an RSA key:

```
$ keytool \
  -genkeypair \
  -keyalg RSA \
  -keysize 1024 \
  -alias signature-key \
  -keystore /path/to/secrets/jwtbuilderkeystore.pkcs12 \
  -storepass password \
  -storetype pkcs12 \
  -dname "CN=openig.example.com,O=Example Corp,C=FR"
```

2. Set an environment variable for the KeyStore `storepass`:

```
$ export KEYSTORE_SECRET_ID='cGFzc3dvcmQ='
```

   The default SystemAndEnvSecretStore is used to manage the secret. Because the default store is configured to read base64 values, the value of `keystore.secret.id` must be base64-encoded.

3. Add the following route to the configuration as `$HOME/.openig/config/routes/jwt-builder-signature.json` (on Windows, `%appdata%\OpenIG\config\routes\jwt-builder-signature.json`):

```
{
  "name": "jwt-signature",
  "heap": [{
    "name": "AmService-1",
```

```
      "type": "AmService",
      "config": {
        "agent" : {
          "username" : "ig_agent",
          "passwordSecretId": "agent.secret.id"
        },
        "url": "http://openam.example.com:8088/openam",
        "version": "6.5"
      }
  }],
  "secrets": {
    "stores": [
      {
        "name":"myKeyStoreSecretStore",
        "type": "KeyStoreSecretStore",
        "config": {
          "file": "/path/to/secrets/jwtbuilderkeystore.pkcs12",
          "storeType": "PKCS12",
          "storePassword": "keystore.secret.id",
          "keyEntryPassword": "keystore.secret.id",
          "mappings": [{
            "secretId": "id.key.for.signing.jwt",
            "aliases": [ "signature-key" ]
          }]
        }
      }
    ]
  },
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [{
        "name": "SingleSignOnFilter-1",
        "type": "SingleSignOnFilter",
        "config": {
          "amService": "AmService-1"
        }
      }, {
        "name": "UserProfileFilter-1",
        "type": "UserProfileFilter",
        "config": {
          "username": "${contexts.ssoToken.info.uid}",
          "userProfileService": {
            "type": "UserProfileService",
            "config": {
              "amService": "AmService-1"
            }
          }
        }
      }, {
        "name": "JwtBuilderFilter-1",
        "type": "JwtBuilderFilter",
        "config": {
          "template": {
            "name": "${contexts.userProfile.commonName}",
            "email": "${contexts.userProfile.rawInfo.mail[0]}"
          },
          "signature": {
            "secretId": "id.key.for.signing.jwt"
```

```
        }
      }
    }, {
      "name": "HeaderFilter-1",
      "type": "HeaderFilter",
      "config": {
        "messageType": "REQUEST",
        "add": {
          "x-openig-user": ["${contexts.jwtBuilder.value}"]
        }
      }
    }],
    "handler": "ReverseProxyHandler"
  }
},
"condition": "${matches(request.uri.path, '/jwt-signature$')}",
"baseURI": "http://app.example.com:8081"
}
```

Notice the following features of the route compared to the previous route:

- The route matches requests to `/jwt-signature`.

- The `secrets` object configures a KeyStoreSecretStore to manage secrets for cryptographic keys and certificates. The KeyStoreSecretStore includes a secret ID mapping to the alias of the signing key.

- The JwtBuilderFilter adds a `signature` property to refer to the mapping for the signature key in KeyStoreSecretStore.

4. Browse to http://openig.example.com:8080/jwt-signature.

5. Log in to AM as user `demo`, password `Ch4ng31t`, or as another user.

   The sample app displays the signed JWT along with its header and payload.

6. Enter information about the KeyStore, and select Verify to verify the signature.

Packing Data Into a Signed and Encrypted JWT

This example builds on the previous example to encrypt the signed JWT, using a new key in the `pkcs12` KeyStore. Run the previous example before you try this.

1. Generate another key in the KeyStore:

```
$ keytool \
  -genkeypair \
  -keyalg RSA \
  -keysize 1024 \
  -alias encryption-key \
  -keystore /path/to/secrets/jwtbuilderkeystore.pkcs12 \
  -storepass password \
  -storetype pkcs12 \
  -dname "CN=openig.example.com,O=Example Corp,C=FR"
```

2. Add the following route to the configuration as `$HOME/.openig/config/routes/jwt-builder-encrypt.json` (on Windows, `%appdata%\OpenIG\config\routes\jwt-builder-encrypt.json`):

```
{
  "name": "jwt-encryption",
  "heap": [{
    "name": "AmService-1",
    "type": "AmService",
    "config": {
      "agent" : {
        "username" : "ig_agent",
        "passwordSecretId": "agent.secret.id"
      },
      "url": "http://openam.example.com:8088/openam",
      "version": "6.5"
    }
  }],
  "secrets": {
    "stores": [
      {
        "name":"myKeyStoreSecretStore",
        "type": "KeyStoreSecretStore",
        "config": {
          "file": "/path/to/secrets/jwtbuilderkeystore.pkcs12",
          "storeType": "PKCS12",
          "storePassword": "keystore.secret.id",
          "keyEntryPassword": "keystore.secret.id",
          "mappings": [
            {
              "secretId": "id.key.for.signing.jwt",
              "aliases": [ "signature-key" ]
            },
            {
              "secretId": "id.key.for.encrypting.jwt",
              "aliases": [ "encryption-key" ]
            }
          ]
        }
      }
    ]
  },
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [{
        "name": "SingleSignOnFilter-1",
        "type": "SingleSignOnFilter",
        "config": {
          "amService": "AmService-1"
        }
      }, {
        "name": "UserProfileFilter-1",
        "type": "UserProfileFilter",
        "config": {
          "username": "${contexts.ssoToken.info.uid}",
          "userProfileService": {
            "type": "UserProfileService",
            "config": {
              "amService": "AmService-1"
```

```
            }
          }
        }
      }, {
        "name": "JwtBuilderFilter-1",
        "type": "JwtBuilderFilter",
        "config": {
          "template": {
            "name": "${contexts.userProfile.commonName}",
            "email": "${contexts.userProfile.rawInfo.mail[0]}"
          },
          "signature": {
            "secretId": "id.key.for.signing.jwt",
            "encryption": {
              "secretId": "id.key.for.encrypting.jwt",
              "algorithm": "RSA-OAEP-256",
              "method": "A128CBC-HS256"
            }
          }
        }
      }, {
        "name": "HeaderFilter-1",
        "type": "HeaderFilter",
        "config": {
          "messageType": "REQUEST",
          "add": {
            "x-openig-user": ["${contexts.jwtBuilder.value}"]
          }
        }
      }],
      "handler": "ReverseProxyHandler"
    }
  },
  "condition": "${matches(request.uri.path, '/jwt-encryption$')}",
  "baseURI": "http://app.example.com:8081"
}
```

Notice the following features of the route compared to the previous route:

- The route matches requests to `/jwt-encryption`.

- The KeyStoreSecretStore includes a secret ID mapping to the alias of the encryption key.

- The JwtBuilderFilter adds a `signature`/`encryption` property to refer to the mapping for the encryption key in KeyStoreSecretStore.

3. Browse to http://openig.example.com:8080/jwt-encryption.

4. Log in to AM as user `demo`, password `Ch4ng31t`, or as another user.

   The sample app displays the signed JWT along with its header and encryption info. Because the JWT is encrypted, the sample app doesn't display the payload.

More Information

org.forgerock.openig.filter.JwtBuilderFilter

org.forgerock.openig.filter.JwtBuilderContext

Name

LocationHeaderFilter — rewrites Location headers

Description

For a response that generates a redirect to the proxied application, this filter rewrites the Location header on the response to redirect the user to IG.

Usage

```
{
    "name": string,
    "type": "LocationHeaderFilter",
    "config": {
        "baseURI": runtime expression<uri string>
    }
}
```

An alternative value for type is RedirectFilter.

Properties

`"baseURI"`: *runtime expression<uri string>, optional*

The base URI of the IG instance. This is used to rewrite the Location header on the response.

The result of the expression must be a string that represents a valid URI, but is not a real `java.net.URI` object. For example, it would be incorrect to use `${request.uri}`, which is not a String but a MutableUri.

Default: Redirect to the original URI specified in the request.

See also Expressions(5).

Example

In the following example, IG listens on `https://openig.example.com:443` and the application it protects listens on `http://app.example.com:8081`. The filter rewrites redirects that would normally take the user to locations under `http://app.example.com:8081` to go instead to locations under `https://openig.example.com:443`.

```
{
    "name": "LocationRewriter",
    "type": "LocationHeaderFilter",
    "config": {
        "baseURI": "https://openig.example.com:443/"
     }
}
```

More Information

org.forgerock.openig.filter.LocationHeaderFilter

Name

OAuth2ClientFilter — authenticate an end user with OAuth 2.0 delegated authorization

Description

An OAuth2ClientFilter is a filter that authenticates an end user using OAuth 2.0 delegated authorization. The filter can act as an OpenID Connect relying party as well as an OAuth 2.0 client.

The client filter does not include information about identity providers, or information about static registration with identity providers. For information about an identity provider, see Issuer(5). For information about registration with an identity provider, see ClientRegistration(5).

In the case where all users share the same identity provider, you can configure the filter as a client of a single provider by referencing a single client registration name for the filter. You can also configure the filter to work with multiple providers, taking the user to a login handler page—often full of provider logos, and known as a *Nascar page*. The name comes from Nascar race cars, some of which are covered with sponsors' logos—to choose a provider.

What an OAuth2ClientFilter does depends on the incoming request URI. In the following list *clientEndpoint* represents the value of the clientEndpoint in the filter configuration:

clientEndpoint*/login/?discovery=*user-input*&goto=*url

Using the *user-input* value, discover and register dynamically with the end user's OpenID Provider or with the client registration endpoint as described in RFC 7591.

Upon successful registration, redirect the end user to the provider for authentication and authorization consent before redirecting the user-agent back to the callback client endpoint.

The goto URL must use the same scheme, host, and port as the original URI, or be a relative URI (just the path). Otherwise, the request fails with an error.

Supported with OpenAM 13 and later versions, and AM 5 and later versions.

clientEndpoint*/login?registration=*clientId*&issuer=*issuerName*&goto=*url

Redirect the end user for authorization with the specified *registration*, defined by the ClientRegistration's `clientId` and `issuerName`. For information, see ClientRegistration(5).

The provider corresponding to the registration then authenticates the end user and obtains authorization consent before redirecting the user-agent back to the callback client endpoint.

Ultimately if the entire process is successful, the filter saves the authorization state in the context and redirects the user-agent to the specified URL.

The goto URL must use the same scheme, host, and port as the original URI, or be a relative URI (just the path). Otherwise, the request fails with an error.

clientEndpoint/**logout?goto=***url*

> Remove the authorization state for the end user and redirect to the specified URL.
>
> The goto URL must use the same scheme, host, and port as the original URI, or be a relative URI (just the path). Otherwise, the request fails with an error.
>
> If no goto URL is specified in the request, the `defaultLogoutGoto` is used.

clientEndpoint/**callback**

> Handle the callback from the OAuth 2.0 authorization server that occurs as part of the authorization process.
>
> If the callback is handled successfully, the filter saves the authorization state in the context at the specified target location and redirects to the URL during login.

Other request URIs

> Restore authorization state in the specified target location and call the next filter or handler in the chain.

Usage

```
{
  "name": string,
  "type": "OAuth2ClientFilter",
  "config": {
    "clientEndpoint": runtime expression<uri sting>,
    "failureHandler": Handler reference,
    "discoveryHandler": Handler reference,
    "loginHandler": Handler reference,
    "registrations": [ ClientRegistration reference(s) ],
    "metadata": dynamic registration client metadata object,
    "cacheExpiration": duration string,
    "executor": executor,
    "target": expression,
    "defaultLoginGoto": runtime expression<uri string>,
    "defaultLogoutGoto": runtime expression<uri string>,
    "requireHttps": boolean,
    "requireLogin": boolean,
    "issuerRepository": Issuer repository reference
  }
}
```

Properties

"clientEndpoint": *runtime expression<uri string>, required*

> The base URI for the filter. For example, if you set `"clientEndpoint": "/openid"`, the service URIs for this filter on your IG server are `/openid/login`, `/openid/logout`, and `/openid/callback`. These endpoints are implicitly reserved. Attempts to access them directly can cause undefined errors.

The result of the expression must be a string that represents a valid URI, but is not a real `java.net` `.URI` object. For example, it would be incorrect to use `${request.uri}`, which is not a String but a MutableUri.

See also Expressions(5).

`"failureHandler"`: *handler reference, required*

Provide an inline handler configuration object, or the name of a handler object that is defined in the heap.

When the OAuth 2.0 Resource Server denies access to a resource, the failure handler can be invoked only if the error response contains a WWW-Authenticate header (meaning that there was a problem with the OAuth 2.0 exchange). All other responses are forwarded to the user-agent without invoking the failure handler.

If the value of the WWW-Authenticate header is `invalid_token`, the OAuth2ClientFilter tries to refresh the access token:

- If the token is refreshed, the OAuth2ClientFilter tries again to access the protected resource.

- If the token is not refreshed, or if the second attempt to access the protected resource fails, the OAuth2ClientFilter invokes the failure handler.

When the failure handler is invoked, the target in the context can be populated with information such as the exception, client registration, and error. The failure object in the target is a simple map, similar to the following example:

```
{
    "client_registration": "ClientRegistration name string",
    "error": {
        "realm": "optional string",
        "scope": [ "optional scope string (required by the client)", ... ],
        "error": "optional string",
        "error_description": "optional string",
        "error_uri": "optional string"
    },
    "access_token": "string",
    "id_token": "string",
    "token_type": "Bearer",
    "expires_in": "number",
    "scope": [ "optional scope string", ... ],
    "client_endpoint": "URL string",
    "exception": exception
}
```

In the failure object, the following fields are not always present. Their presence depends on when the failure occurs:

- "access_token"

- "id_token"

- "token_type"

- "expires_in"

- "scope"

- "client_endpoint"

See also Handlers.

"discoveryHandler": *Handler reference, optional*

Invoke this HTTP client handler to communicate with the OpenID Provider for OpenID Connect Discovery.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

Usually set this to the name of a ClientHandler configured in the heap, or a chain that ends in a ClientHandler.

Default: IG uses the default ClientHandler.

See also Handlers, ClientHandler(5).

"loginHandler": *Handler reference, required if there are zero or multiple client registrations, optional if there is one client registration*

Use this Handler when the user must choose an identity provider. When `registrations` contains only one client registration, this Handler is optional but is displayed if specified.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

The route in "Example Configuration For Multiple Identity Providers" includes a login handler that allows the user to choose from two client registrations, `openam` and `google`. For an example of a login handler where no client registrations are defined, see "Preparing IG for Discovery and Dynamic Registration" in the *Gateway Guide*.

See also Handlers.

"registrations": *Array of ClientRegistration references or inline ClientRegistration declarations, optional*

List of client registrations that authenticate IG to the identity providers. The list must contain all client registrations that are to be used by the client filter.

The value represents a static client registration with an identity provider as described in ClientRegistration(5).

"metadata": *client metadata object, required for dynamic client registration and ignored otherwise*

This object holds client metadata as described in *OpenID Connect Dynamic Client Registration 1.0*, and optionally a list of scopes. See that document for additional details and a full list of fields.

This object can also hold client metadata as described in RFC 7591, *OAuth 2.0 Dynamic Client Registration Protocol*. See that RFC for additional details.

The following partial list of metadata fields is not exhaustive, but includes metadata that is useful with AM as OpenID Provider:

"redirect_uris": *array of URI strings, required*

The array of redirection URIs to use when dynamically registering this client.

"client_name": *string, optional*

Name of the client to present to the end user.

"scope": *space separated string, optional*

Space separated string of scopes to request of the OpenID Provider, for example:

```
"scope": "openid profile"
```

This property is available for dynamic client registration with AM 5.5 and later versions, or with identity providers that support RFC 7591, *OAuth 2.0 Dynamic Client Registration Protocol*

"scopes": *array of strings, optional*

Array of scope strings to request of the OpenID Provider, for example:

```
"scopes": [
  "openid",
  "profile",
  "email"
]
```

This property is available for dynamic client registration with AM 5.5 and earlier versions only.

"cacheExpiration": *duration string, optional*

Duration for which to cache user-info resources.

IG lazily fetches user info from the OpenID provider. In other words, IG only fetches the information when a downstream Filter or Handler uses the user info. Caching allows IG to avoid repeated calls to OpenID providers when reusing the information over a short period.

Default: 10 minutes

Set this to disabled or zero to disable caching. When caching is disabled, user info is still lazily fetched.

For information about supported formats for `duration`, see duration.

"executor": *executor, optional*

An executor service to schedule the execution of tasks, such as the eviction of entries in the OpenID Connect user information cache.

Default: `ScheduledExecutorService`

See also ScheduledExecutorService(5).

"target": *expression, optional*

An expression that yields the target object. Downstream filters and handlers can use data in the target to enrich the existing request or create a new request.

- If the authorization process completes successfully, the OAuth2ClientFilter injects the authorization state data into the target.

  In the following example, a downstream StaticRequestFilter retrieves the username and password from the target to log the user in to the sample application.

  ```
  {
    "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}"
        ]
      }
    }
  }
  ```

  For information about setting up this example, see "Authenticating Automatically to the Sample Application" in the *Gateway Guide*.

- If the failure handler is invoked, the target can be populated with information such as the exception, client registration, and error, as described in "failureHandler" in this reference page.

Default: `${attributes.openid}`

See also Expressions(5).

"defaultLoginGoto": *runtime expression<uri string>, optional*

After successful authentication and authorization, if the user accesses the *clientEndpoint*/login endpoint without providing a landing page URL in the `goto` parameter, the request is redirected to this URI.

The goto URL must use the same scheme, host, and port as the original URI, or be a relative URI (just the path). Otherwise, the request fails with an error.

The result of the expression must be a string that represents a valid URI, but is not a real `java.net.URI` object. For example, it would be incorrect to use `${request.uri}`, which is not a String but a MutableUri.

Default: return an empty page.

See also Expressions(5).

`"defaultLogoutGoto"`: *runtime expression<uri string>, optional*

If the user accesses the `clientEndpoint/logout` endpoint without providing a goto URL, the request is redirected to this URI.

The goto URL must use the same scheme, host, and port as the original URI, or be a relative URI (just the path). Otherwise, the request fails with an error.

The result of the expression must be a string that represents a valid URI, but is not a real `java.net.URI` object. For example, it would be incorrect to use `${request.uri}`, which is not a String but a MutableUri.

Default: return an empty page.

See also Expressions(5).

`"requireHttps"`: *boolean, optional*

Whether to require that original target URI for the request (`originalUri` in UriRouterContext(5)) uses the HTTPS scheme.

Default: true.

`"requireLogin"`: *boolean, optional*

Whether to require authentication for all incoming requests.

Default: true.

`"issuerRepository"`: *Issuer repository reference, optional*

A repository of OAuth 2.0 issuers, built from discovered issuers and the IG configuration.

Provide the name of an IssuerRepository object defined in the heap.

Default: Look up an issuer repository named `IssuerRepository` in the heap. If none is explicitly defined, then a default one named `IssuerRepository` is created in the current route.

See also IssuerRepository(5).

Example Configuration For Multiple Identity Providers

The following example builds on the AM and IG configuration in "Setting Up IG As a Relying Party" in the *Gateway Guide*.

The client registration for the AM provider is now declared in the heap. A second client registration defines Google as an alternative identity provider. The NascarPage helps the user to choose an identity provider.

When a request is made to `/home/id_token/login`, the OAuth2ClientFilter directs the user to the NascarPage to choose an identity provider, and then handles negotiation for authorization with the chosen provider.

If the authorization process completes successfully, the filter injects the authorization state data into `attributes.openid`. The ReverseProxyHandler then passes the request to the home page of the sample application.

If a request is made to `/home/id_token/logout` without specifying a goto URL, the request is redirected to www.forgerock.com by the `defaultLogoutGoto` property.

```
{
  "heap": [
    {
      "name": "openam",
      "type": "ClientRegistration",
      "config": {
        "clientId": "oidc_client",
        "clientSecret": "password",
        "issuer": {
          "name": "Issuer",
          "type": "Issuer",
          "config": {
            "wellKnownEndpoint": "http://openam.example.com:8088/openam/oauth2/.well-known/openid-
configuration"
          }
        },
        "scopes": [
          "openid",
          "profile",
          "email"
        ],
        "tokenEndpointAuthMethod": "client_secret_basic"
      }
    },
    {
      "name": "google",
      "type": "ClientRegistration",
      "config": {
        "clientId": "myClientId",
        "clientSecret": "myClientSecret",
        "issuer": {
          "name": "accounts.google.com",
          "type": "Issuer",
          "config": {
            "wellKnownEndpoint": "https://accounts.google.com/.well-known/openid-configuration"
```

```
          }
        },
        "scopes": [
          "openid",
          "profile"
        ]
      }
    }
  },
  {
    "name": "NascarPage",
    "type": "StaticResponseHandler",
    "config": {
      "status": 200,
      "entity": "<html><body><p><a href='/home/id_token/login?
registration=oidc_client&issuer=Issuer&goto=${urlEncodeQueryParameterNameOrValue('http://
openig.example.com:8080/home/id_token')}'>AM Login</a></p><p><a href='/home/id_token/login?
registration=googleClientId&issuer=accounts.google.com&goto=${urlEncodeQueryParameterNameOrValue('http://
openig.example.com:8080/home/id_token')}'>Google Login</a></p></body></html>"
    }
  }
],
"name": "07-openid-nascar",
"baseURI": "http://app.example.com:8081",
"condition": "${matches(request.uri.path, '^/home/id_token')}",
"handler": {
  "type": "Chain",
  "config": {
    "filters": [
      {
        "type": "OAuth2ClientFilter",
        "config": {
          "target" : "${attributes.openid}",
          "clientEndpoint": "/home/id_token",
          "loginHandler" : "NascarPage",
          "registrations": [ "openam", "google" ],
          "defaultLogoutGoto" : "http://www.forgerock.com",
          "requireHttps": false,
          "requireLogin": true,
          "failureHandler": {
            "type": "StaticResponseHandler",
            "config": {
              "comment": "Trivial failure handler for debugging only",
              "status": 500,
              "reason": "Error",
              "entity": "${attributes.openid}"
            }
          }
        }
      }
    ],
    "handler": "ReverseProxyHandler"
  }
}
}
```

Use the following hints to set up Google credentials for this example:

1. Go to https://console.cloud.google.com/apis/credentials.

2. Create credentials for an OAuth 2.0 client ID with the following options:

- Application type: `Web application`

- Authorized redirect URI: `http://openig.example.com:8080/home/id_token/callback`

3. In the IG route, replace `myClientId` and `myClientSecret` with the returned credentials.

For information about configuring providers, see Issuer(5) and ClientRegistration(5).

Notice that this configuration is for development and testing purposes only, and is not secure ("requireHttps": false). Make sure you do require HTTPS in production environments.

More Information

org.forgerock.openig.filter.oauth2.client.OAuth2ClientFilter

Issuer(5), ClientRegistration(5)

The OAuth 2.0 Authorization Framework

OAuth 2.0 Bearer Token Usage

OpenID Connect site, in particular the list of standard OpenID Connect 1.0 scope values

Name

OAuth2ResourceServerFilter — validate a request containing an OAuth 2.0 access_token

Description

This filter validates a request containing an OAuth 2.0 access_token. The filter expects an OAuth 2.0 token from the HTTP Authorization header of the request, such as the following example header, where the OAuth 2.0 access_token is `1fc0e143-f248-4e50-9c13-1d710360cec9`:

```
Authorization: Bearer 1fc0e143-f248-4e50-9c13-1d710360cec9
```

The filter extracts the access_token, validates it with the configured access token resolver, and stores the access_token in the `oauth2` context. For information, see OAuth2Context(5).

The following table summarizes the errors that can occur when IG tries to validate the access_token:

Errors During access_token Validation

Error	Response from the filter to the user-agent
Combination of the filter configuration and access_token result in an invalid request to the authorization server.	`HTTP 400 Bad Request`
No access_token in the request header	`HTTP 401 Unauthorized WWW-Authenticate: Bearer realm="IG"`
Access-token not valid, for example, because it has expired	`HTTP 401 Unauthorized`
Scopes for the access_token don't match the specified required scopes	`HTTP 403 Forbidden`

Usage

```
{
  "name": string,
  "type": "OAuth2ResourceServerFilter",
  "config": {
    "accessTokenResolver": AccessTokenResolver reference,
    "cache": object,
    "executor": executor,
    "requireHttps": boolean,
    "realm": string,
    "scopes": [ runtime expression<string>, ... ] or object
  }
}
```

An alternative value for type is OAuth2RSFilter.

Properties

"accessTokenResolver": *AccessTokenResolver reference, required*

Resolves an access_token against an authorization server. Configure one of the following access_token resolvers:

- TokenIntrospectionAccessTokenResolver(5)

- StatelessAccessTokenResolver(5)

- OpenAmAccessTokenResolver(5)

- ConfirmationKeyVerifierAccessTokenResolver(5)

- ScriptableAccessTokenResolver(5)

"cache": *object, optional*

Configuration of caching for OAuth 2.0 access_tokens. By default, access_tokens are not cached.

When an access_token is cached, IG can reuse the token information without repeatedly asking the authorization server to verify the access_token. When caching is disabled, IG must ask the authorization server to verify the access_token for each request.

The following example caches access_tokens for these times:

- One hour, if the access_token doesn't provide a valid expiry value.

- The duration specified by the token expiry value, when the token expiry value is shorter than one day.

- One day, when the token expiry value is longer than one day.

```
"cache": {
  "enabled": true,
  "defaultTimeout": "1 hour",
  "maxTimeout": "1 day"
}
```

enabled: *boolean, optional*

Enable or disable caching.

Default: `false`

defaultTimeout: *configuration expression<duration string>, optional*

The duration for which to cache an OAuth 2.0 access_token if it doesn't provide a valid expiry value.

If an access_token provides an expiry value that falls *before* the current time plus the `maxTimeout`, IG uses the token expiry value.

Default: `1 minute`

`maxTimeout:` *configuration expression<duration string>, optional*

The maximum duration for which to cache OAuth 2.0 access_tokens.

If an access_token provides an expiry value that falls *after* the current time plus the `maxTimeout`, IG uses the `maxTimeout`.

The duration cannot be `zero` or `unlimited`.

For information about supported formats for `duration`, see duration.

`"executor":` *executor, optional*

An executor service to schedule the execution of tasks, such as the eviction of entries in the access_token cache.

Default: `ScheduledExecutorService`

See also ScheduledExecutorService(5).

`"requireHttps":` *boolean, optional*

Whether to require that original target URI for the request (`originalUri` in UriRouterContext(5)) uses the HTTPS scheme.

Default: true

`"realm":` *string, optional*

HTTP authentication realm to include in the WWW-Authenticate response header field when returning an HTTP 401 Unauthorized status to a user-agent that need to authenticate.

Default: IG

`"scopes":` *array of runtime expression<string> or ScriptableResourceAccess object, required*

A list of one of more scopes required in an OAuth 2.0 access_token to access a protected resource.

array of runtime expression<string>, required if ScriptableResourceAccess is not used

A string, array of strings, runtime expression<string>, or array of runtime expression<string> to represent one or more scopes; mutually exclusive with `ScriptableResourceAccess`.

ScriptableResourceAccess, required if "array of runtime expression<string>" is not used

A script that produces a list of one or more required scopes; mutually exclusive with "array of runtime expression<string>".

The script evaluates each request dynamically and returns the scopes that request needs to access the protected resource. The script must return a Promise<Set, ResponseException> or a Set<String>.

For information about the properties of ScriptableResourceAccess, see Scripts(5). For an example of how to use this property, see the examples section on this page.

```
{
  "name": ScriptableResourceAccess-1,
  "type": "ScriptableResourceAccess",
  "config": {
    "type": string,
    "file": expression,                  // Use either "file"
    "source": string or array of strings,   // or "source", but not both.
    "args": object,
    "clientHandler": Handler reference
  }
}
```

Examples

Configuring an OAuth 2.0 Resource Filter

The following example validates an access_token against the AM token info endpoint. The filter expects to retrieve the scopes `mail` and `employeenumber` in the token, and returns an HTTP 403 Forbidden status if the scopes are not present.

```
{
  "name": "ProtectedResourceFilter",
  "type": "OAuth2ResourceServerFilter",
  "config": {
    "scopes": [
      "mail",
      "employeenumber"
    ],
    "requireHttps": false,
    "accessTokenResolver": {
      "type": "OpenAmAccessTokenResolver",
      "config": {
        "amService": "MyAmService"
      }
    }
  }
}
```

For information about how to set up examples using OAuth2ResourceServerFilter, see "*Acting as an OAuth 2.0 Resource Server*" in the *Gateway Guide*.

Defining Required Scopes by Using a Script

The following example uses a ScriptableResourceAccess object to define the scopes required in a request:

- If the request path is `/rs-tokeninfo`, only the scopes `mail` is required.

- If the request path is `/rs-tokeninfo/employee`, the scopes `mail` and `employeenumber` are required.

1. Set up the example in "Validating Access-Tokens Through the Token Info Endpoint" in the *Gateway Guide*, until the step for Scope configuration.

   Instead of configuring scopes Statically, select Dynamically, and enter the following script to evaluate requests dynamically and return required scopes:

   ```
   // Minimal set of required scopes
   def scopes = [ 'mail' ] as Set
   if (request.uri.path =~ /employee$/) {
     // Require another scope to access this resource
     scopes += 'employeenumber'
   }
   return scopes
   ```

   For notes about entering scripts in Studio, see "Scripting in Studio" in the *Gateway Guide*.

2. Finish the setup as in "Validating Access-Tokens Through the Token Info Endpoint" in the *Gateway Guide*. The following route should be displayed:

   ```
   {
     "name": "rs-tokeninfo",
     "baseURI": "http://app.example.com:8081",
     "condition": "${matches(request.uri.path, '^/rs-tokeninfo')}",
     "heap": [
       {
         "name": "AmService-1",
         "type": "AmService",
         "config": {
           "url": "http://openam.example.com:8088/openam",
           "realm": "/",
           "ssoTokenHeader": "iPlanetDirectoryPro",
           "version": "6.5",
           "agent": {
             "username": "ig_agent",
             "password": "password"
           },
           "sessionCache": {
             "enabled": false
           }
         }
       }
     ],
     "handler": {
       "type": "Chain",
       "config": {
         "filters": [
           {
             "name": "OAuth2ResourceServerFilter-1",
             "type": "OAuth2ResourceServerFilter",
             "config": {
               "scopes": {
                 "name": "myscript",
                 "type": "ScriptableResourceAccess",
                 "config": {
                   "type": "application/x-groovy",
   ```

```
                "source": [
                  "// Minimal set of required scopes",
                  "def scopes = [ 'mail' ] as Set",
                  "if (request.uri.path =~ /employee$/) {",
                  "  // Require another scope to access this resource",
                  "  scopes += 'employeenumber'",
                  "}",
                  "return scopes"
                ]
              }
            },
            "requireHttps": false,
            "realm": "OpenIG",
            "accessTokenResolver": {
              "name": "token-resolver-1",
              "type": "OpenAmAccessTokenResolver",
              "config": {
                "amService": "AmService-1"
              }
            }
          }
        }
      }
    ],
    "handler": {
      "type": "StaticResponseHandler",
      "config": {
        "status": 200,
        "entity": "<html><body><h2>Decoded access_token: ${contexts.oauth2.accessToken.info}</h2></
body></html>"
      }
    }
  }
 }
}
```

3. In a terminal window, use a **curl** command similar to the following to retrieve an access_token:

```
$ mytoken=$(curl -s
 \
--user "client-application:password"
 \
--data "grant_type=password&username=george&password=C0stanza&scope=mail" \
http://openam.example.com:8088/openam/oauth2/access_token | jq -r ".access_token")
```

4. Validate the access_token for the `/rs-tokeninfo` path and confirm that the content of the access_token is returned:

```
$ curl -v http://openig.example.com:8080/rs-tokeninfo --header "Authorization: Bearer ${mytoken}"

{
  access_token = XdoJ6O3Bo - PIrTzQTg5bXYLzOtE,
    mail = george@example.com,
    grant_type = password,
    auth_level = 0,
    scope = [mail],
    realm = /,
  token_type=Bearer,
    expires_in=3598,
    client_id=client-application
}
```

5. Validate the access_token for the `/rs-tokeninfo/employee` path, and confirm the request returns an HTTP 403 because the `employee` scope is not present:

```
$ curl -v http://openig.example.com:8080/rs-tokeninfo/employee --header "Authorization: Bearer
 ${mytoken}"
```

6. Repeat the previous steps, but this time retrieve an access_token with both the `mail` and `employeenumber` scopes, and confirm that the content of the access_token is returned for the `/rs-tokeninfo/employee` path.

```
$ mytoken=$(curl -s
 \
--user "client-application:password"
 \
--data "grant_type=password&username=george&password=C0stanza&scope=mail%20employeenumber" \
http://openam.example.com:8088/openam/oauth2/access_token | jq -r ".access_token")
```

More information

org.forgerock.openig.filter.oauth2.OAuth2ResourceServerFilterHeaplet

org.forgerock.http.oauth2.OAuth2Context

org.forgerock.http.oauth2.AccessTokenInfo

OAuth2Context(5)

ConfirmationKeyVerifierAccessTokenResolver(5)

OpenAmAccessTokenResolver(5)

TokenIntrospectionAccessTokenResolver(5)

StatelessAccessTokenResolver(5)

ScriptableAccessTokenResolver(5)

The OAuth 2.0 Authorization Framework

OAuth 2.0 Bearer Token Usage

Name

PasswordReplayFilter — replay credentials with a single filter

Description

Replays credentials in a single composite filter for the following cases:

• When the request is for a login page

• When the response contains a login page

When the response contains a login page, a PasswordReplayFilter can extract values from the response entity and reuse the values when replaying credentials.

A PasswordReplayFilter does not retry failed authentication attempts.

Usage

```
{
    "name": string,
    "type": "PasswordReplayFilter",
    "config": {
        "request": request configuration object,
        "loginPage": runtime expression<boolean>,
        "loginPageContentMarker": pattern,
        "credentials": Filter reference,
        "headerDecryption": crypto configuration object,
        "loginPageExtractions": [ extract configuration object, ... ]
    }
}
```

Properties

`"request"`: *request configuration object, required*

The request that replays the credentials.

The request configuration object has the following fields:

`"method"`: *string, required*

The HTTP method to be performed on the resource such as `GET` or `POST`.

`"uri"`: *string, required*

The fully qualified URI of the resource to access such as `http://www.example.com/login`.

`"entity"`: *expression, optional*

The entity body to include in the request.

This setting is mutually exclusive with the `form` setting when the `method` is set to `POST`.

See also Expressions(5).

"form": *object, optional*

A form to include in the request.

The `param` specifies the form parameter name. Its value is an array of expressions to evaluate as form field values.

This setting is mutually exclusive with the `entity` setting when the `method` is set to `POST`.

"headers": *object, optional*

Header fields to set in the request.

The `name` specifies the header name. Its value is an array of expressions to evaluate as header values.

"version": *string, optional*

The HTTP protocol version.

Default: `"HTTP/1.1"`.

The implementation uses a StaticRequestFilter. The fields are the same as those described in StaticRequestFilter(5).

"loginPage": *runtime expression<boolean>, required unless loginPageContentMarker is defined*

An expression that is true when a login page is requested, false otherwise.

For example, the following expression specifies that an HTTP GET to the path `/login` is a request for a login page:

```
${matches(request.uri.path, '/login') and (request.method == 'GET')}
```

IG evaluates the expression only for the request, not for the response.

See also Expressions(5).

"loginPageContentMarker": *pattern, required unless loginPage is defined*

A pattern that matches when a response entity is that of a login page.

For an example route that uses this property, see "Login Form With Password Replay and Cookie Filters" in the *Gateway Guide*.

See also Patterns(5).

"credentials": *Filter reference, optional*

Filter that injects credentials, making them available for replay. Consider using a `FileAttributesFilter` or a `SqlAttributesFilter`.

When this is not specified, credentials must be made available to the request by other means.

See also Filters.

"headerDecryption": *crypto configuration object, optional*

Object to decrypt request headers that contain credentials to replay.

The crypto configuration object has the following fields:

"key": *expression, required*

Base64 encoded key value.

See also Expressions(5).

"algorithm": *string, optional*

Algorithm used for decryption.

Use the same algorithm that is used to send the encrypted credentials.

Default: `AES/ECB/PKCS5Padding`

"keyType": *string, optional*

Algorithm name for the secret key.

Default: `AES`

"headers": *array of strings, optional*

The names of header fields to decrypt.

Default: Do not decrypt any headers.

"loginPageExtractions": *extract configuration array, optional*

Object to extract values from the login page entity.

For an example route that uses this property, see "Login Which Requires a Hidden Value From the Login Page" in the *Gateway Guide*.

The extract configuration array is a series of configuration objects. To extract multiple values, use multiple extract configuration objects. Each object has the following fields:

"name": *string, required*

Name of the field where the extracted value is put.

The names are mapped into `attributes.extracted`.

For example, if the name is `nonce`, the value can be obtained with the expression `${attributes.extracted.nonce}`.

The name `isLoginPage` is reserved to hold a boolean that indicates whether the response entity is a login page.

"pattern": *pattern, required*

The regular expression pattern to find in the entity.

The pattern must contain one capturing group. (If it contains more than one, only the value matching the first group is placed into `attributes.extracted`.)

For example, suppose the login page entity contains a nonce required to authenticate, and the nonce in the page looks like `nonce='n-0S6_WzA2Mj'`. To extract `n-0S6_WzA2Mj`, set `"pattern": " nonce='(.*)'"`.

See also Patterns(5).

Examples

The following example route authenticates requests using static credentials whenever the request is for `/login`. This PasswordReplayFilter example does not include any mechanism for remembering when authentication has already been successful. It simply replays the authentication every time that the request is for `/login`:

```
{
    "handler": {
        "type": "Chain",
        "config": {
            "filters": [
                {
                    "type": "PasswordReplayFilter",
                    "config": {
                        "loginPage": "${request.uri.path == '/login'}",
                        "request": {
                            "method": "POST",
                            "uri": "https://www.example.com:8444/login",
                            "form": {
                                "username": [
                                    "MY_USERNAME"
                                ],
                                "password": [
                                    "MY_PASSWORD"
                                ]
                            }
                        }
                    }
                }
            ],
            "handler": "ReverseProxyHandler"
        }
    }
}
```

For additional examples, see "*Configuration Templates*" in the *Gateway Guide*, and the Javadoc for the PasswordReplayFilter class.

More Information

org.forgerock.openig.filter.PasswordReplayFilterHeaplet

Name

PolicyEnforcementFilter — enforce policy decisions from AM

Description

This filter requests policy decisions from AM, which allows or denies the request based on the request context, the request URI, and the AM policies.

Supported with AM 5 and later versions.

Processing Requests After a Policy Decision

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

- If the request is allowed, processing continues.

- If the request is denied with advices, IG checks whether it can respond to the advices. If IG can respond, it sends a redirect and information about how to meet the conditions in the advices.

  By default, the request is redirected to AM. If the SingleSignOnFilter property `loginEndpoint` is configured, the request is redirected to that endpoint.

- If the request is denied without advice, or if IG cannot respond to the advice, IG forwards the request to a `failureHandler` declared in the `PolicyEnforcementFilter`. If there is no `failureHandler`, IG returns a 403 Forbidden.

- If an error occurs during the process, IG returns 500 Internal Server Error.

Using Advices From Policy Decisions

Attributes and advices returned by the policy decision are stored in the `policyDecision` context. For information, see PolicyDecisionContext(5).

IG responds to the following advice types from AM:

- `AuthLevel`: The minimum authentication level at which a user-agent must authenticate to access a resource.

- `AuthenticateToService`: The name of an authorization chain or service to which a user-agent must authenticate to access a resource.

- `AuthenticateToRealm`: The name of a realm to which a user-agent must authenticate to access a resource.

- `AuthScheme`: The name of an authentication module to which a user-agent must authenticate to access a resource, the policy set name, and the authentication timeout.

- `Transaction`: The additional actions that a user-agent must perform before having a one-time access to the protected resource. Supported with AM 5.5 and later versions.

Notes on Configuring Policies in AM

In the AM policy, remember to configure the `Resources` parameter with the URI of the protected application.

The request URI from IG must match the `Resources` parameter defined in the AM policy. If the URI of the incoming request is changed before it enters the policy filter (for example, by rebasing or scripting), remember to change the `Resources` parameter in AM policy accordingly.

WebSocket Notifications for Policy Changes

When WebSocket notifications are set up for changes to policies, IG receives a notification from AM when a policy decision is created, deleted, or updated.

For information about setting up WebSocket notifications, using them to clear the policy cache, and including them in the server logs, see "WebSocket Notification Service".

Usage

```
{
  "name": string,
  "type": "PolicyEnforcementFilter",
  "config": {
    "amService": AmService reference,
    "pepRealm": string,
    "ssoTokenSubject": runtime expression<string>,
    "jwtSubject": runtime expression<string>,
    "claimsSubject": map or runtime expression<map>,
    "cache": object,
    "application": string,
    "environment": map or runtime expression<map>,
    "failureHandler": handler reference,
    "resourceUriProvider": ResourceUriProvider reference
  }
}
```

Properties

"amService": *AmService reference, required*

The AmService object to use for the following properties:

- `agent`, the credentials of an AM Java agent that acts on behalf of IG to authenticate with AM, get the user profile, and communicate WebSocket notifications from AM to IG.

- `realm`: Realm of the AM Java agent that acts on behalf of IG.

- `url`, the URL of an AM service to use for session token validation and authentication.

- `ssoTokenHeader`, the name of the HTTP header that provides the session token created by AM

- `amHandler`, the handler to use when requesting policy decisions from AM.

- `version`: The version of the AM server.

See also, AmService(5).

"pepRealm": *string, optional*

The AM realm where the policy set is located.

Default: The realm declared for `amService`.

"ssoTokenSubject": *runtime expression<string>, required if neither of the following properties are present:* **"jwtSubject", "claimsSubject"**

The AM SSO or CDSSO token ID string for the subject making the request to the protected resource.

`ssoTokenSubject` can take the following values:

- `${contexts.ssoToken.value}`, when the SingleSignOnFilter is used for authentication

- `${contexts.ssoToken.value}` or `${contexts.cdsso.token}`, when the CrossDomainSingleSignOnFilter is used for authentication

When no SSO is involved (API protection), then `ssoTokenSubject` usually points to a header value (`${request.headers.iPlanetDirectoryPro[0]}`).

See also Expressions(5).

"jwtSubject": *runtime expression<string>, required if neither of the following properties are present:* **"ssoTokenSubject", "claimsSubject"**

The JWT string for the subject making the request to the protected resource.

To use the raw id_token (base64, not decoded) returned by the OpenID Connect Provider during authentication, place an `OAuth2ClientFilter` filter before the PEP filter, and then use `${attributes.openid.id_token}` as the expression value.

See also OAuth2ClientFilter(5) and Expressions(5).

"claimsSubject": *map or runtime expression<map>, required if neither of the following properties are present:* **"jwtSubject", "ssoTokenSubject"**

A representation of JWT claims for the subject.

The claim `"sub"` must be specified. Other claims are optional.

When this property is an expression, its evaluation must give an object of type `Map<String, Object>`.

```
"claimsSubject": "${attributes.openid.id_token_claims}"
```

When this property is a map, the structure must have the format `Map<String, Object>`. The value is evaluated as an expression.

```
"claimsSubject": {
  "sub": "${attributes.subject_identifier}",
  "iss": "openam.example.com"
}
```

For an example of using `claimsSubject` as a map, see "Example Policy Enforcement Using claimsSubject" on this reference page.

`"application"`: *string, optional*

The AM policy set to use when requesting policy decisions.

Default: `Default Policy Set`

`cache`: *object, optional*

Enable and configure caching of policy decisions from AM, based on *Caffeine*. For more information, see the GitHub entry, *Caffeine*.

```
{
  "cache": {
    "enabled": configuration expression<boolean>,
    "defaultTimeout": configuration expression<duration>,
    "executor": executor service reference,
    "maximumSize": configuration expression<number>,
    "maximumTimeToCache": configuration expression<duration>,
    "onNotificationDisconnection": configuration expression<enumeration>,
  }
}
```

Default: Policy decisions are not cached.

> **Note**
>
> Policy decisions that contain advices are never cached.

The following code example caches AM policy decisions without advices for these times:

- One hour, when the policy decision doesn't provide a `ttl` value.

- The duration specified by the `ttl`, when `ttl` is shorter than one day.

- One day, when `ttl` is longer than one day.

```
"cache": {
  "enabled": true,
  "defaultTimeout": "1 hour",
  "maximumTimeToCache": "1 day"
}
```

For information about supported formats for `duration`, see duration.

enabled: *configuration expression<boolean>, optional*

Enable or disable caching of policy decisions.

When a policy decision is cached, IG can reuse the policy decision without repeatedly asking AM for a new policy decision. When caching is disabled, IG must ask AM to make a decision for each request.

Default: `false`

defaultTimeout: *configuration expression<duration>, optional*

The default duration for which to cache AM policy decisions.

If an AM policy decision provides a valid `ttl` value to specify the time until which the policy decision remains valid, IG uses that value or the `maxTimeout`.

Default: `1 minute`

For information about supported formats for `duration`, see duration.

executor: *executor service reference, optional*

An executor service to schedule the execution of tasks, such as the eviction of entries in the cache.

Default: `ForkJoinPool.commonPool()`

"maximumSize": *configuration expression<number>, optional*

The maximum number of entries the cache can contain.

Default: Unlimited/unbound.

maximumTimeToCache: *configuration expression<duration>, optional*

The maximum duration for which to cache AM policy decisions.

If the `ttl` value provided by the AM policy decision is after the current time plus the `maximumTimeToCache`, IG uses the `maximumTimeToCache`.

The duration cannot be `zero` or `unlimited`.

For information about supported formats for `duration`, see duration.

FORGEROCK

onNotificationDisconnection: *configuration expression<enumeration>, optional*

The strategy to manage the cache when the WebSocket notification service is disconnected, and IG receives no notifications for AM events, such as policy changes. If the cache is not cleared it can become outdated, and IG can enforce outdated policy decisions.

Cached entries that expire naturally while the notification service is disconnected are removed from the cache.

Use one of the following values:

- NEVER_CLEAR

  - When the notification service is disconnected:

    - Continue to use the existing cache.

    - Deny access for requests that are not cached, but do not update the cache with these requests.

  - When the notification service is reconnected:

    - Continue to use the existing cache.

    - Query AM for incoming requests that are not found in the cache, and update the cache with these requests.

- CLEAR_ON_DISCONNECT

  - When the notification service is disconnected:

    - Clear the cache.

    - Deny access to all requests, but do not update the cache with these requests.

  - When the notification service is reconnected:

    - Query AM for all requests that are not found in the cache. (Because the cache was cleared, the cache is empty after reconnection.)

    - Update the cache with these requests.

- CLEAR_ON_RECONNECT

  - When the notification service is disconnected:

    - Continue to use the existing cache.

    - Deny access for requests that are not cached, but do not update the cache with these requests.

- When the notification service is reconnected:

  - Query AM for all requests that are not found in the cache. (Because the cache was cleared, the cache is empty after reconnection.)

  - Update the cache with these requests.

  Default: `CLEAR_ON_DISCONNECT`

`"environment"`: *map or runtime expression<map>, optional*

A list of strings to forward to AM with the request for a policy decision, that represent the environment (or context) of the request. Forward any HTTP header or any value that the AM policy definition can use.

AM can use the environment conditions to set the circumstances under which a policy applies. For example, environment conditions can specify that the policy applies only during working hours or only when accessing from a specific IP address.

If this property is a map, the structure must have the format `Map<String, List<String>>`. In the following example, the `PolicyEnforcementFilter` forwards standard and non-standard request headers, an ID token, and the IP address of the subject making the request:

```
"environment": {
  "H-Via": "${request.headers['Via']}",
  "H-X-Forwarded-For": "${request.headers['X-Forwarded-For']}",
  "H-myHeader": "${request.headers['myHeader']}",
  "id_token": [
    "${attributes.openid.id_token}"
  ],
    "IP": [
      "${contexts.client.remoteAddress}"
  ]
}
```

If this property is an expression, its evaluation must give an object of type `Map<String, List<String>>`.

```
"environment": "${attributes.my_environment}"
```

`"failureHandler"`: *handler reference, optional*

Handler to treat requests when they are denied by the policy decision.

In the following example, the `failureHandler` is a chain with a scriptable filter. If there are some advices with the policy decision, the script recovers the advices for processing. Otherwise, it passes the request to the `StaticResponseHandler` to display a message.

```
"failureHandler":{
  "type": "Chain",
  "config": {
    "filters": [
      {
        "type": "ScriptableFilter",
        "config": {
          "type": "application/x-groovy",
          "source": [
            "if (contexts.policyDecision.advices['MyCustomAdvice'] != null) {",
            "  return handleCustomAdvice(context, request)",
            "} else {",
            "  return next.handle(context, request)",
            "}"
          ]
        }
      }
    ],
    "handler" : {
      "type": "StaticResponseHandler",
      "config": {
        "status": 403,
        "entity": "Restricted area. You do not have sufficient privileges."
      }
    }
  }
}
```

Provide an inline handler configuration object, or the name of a handler object that is defined in the heap.

See also Handlers.

Default: HTTP 403 Forbidden, the request stops being executed.

"resourceUriProvider": *ResourceUriProvider reference, optional*

Return the resource URL to include in policy decision requests to AM. Configure one of the following ResourceUriProviders inline or in the heap:

- `RequestResourceUriProvider`

- `ScriptableResourceUriProvider`

Default: `RequestResourceUriProvider`

RequestResourceUriProvider

Return the resource URL to include in policy decision requests to AM by using the original URI of the request or the `baseURI` of the route.

```
"resourceUriProvider": {
  "type": "RequestResourceUriProvider",
  "config": {
    "useOriginalUri": runtime expression<boolean>,
    "includeQueryParams": runtime expression<boolean>
  }
}
```

useOriginalUri: *runtime expression<boolean>, optional*

In policy decision requests to AM, use the original URI of the request as the resource URL.

- `true`: Use the original URI of the request as the resource URL when requesting policy decisions from AM.

- `false`: Use the `baseURI` of the route as the resource URL when requesting policy decisions from AM.

Default: `false`

includeQueryParams: *runtime expression<boolean>, optional*

Include query parameters in the resource URL when requesting policy decisions from AM:

- `true`: Include query parameters in the resource URL. For example, use the following URL with a query parameter: `http://openig.example.com:8080/login?demo=capture`.

- `false`: Exclude query parameters from the resource URL. For example, exclude the query parameter from the previous example: `http://openig.example.com:8080/login`.

  For AM policies that specify resource URLs without query parameters, use this option to reduce the amount of cached information. For information about resource URLs with query parameters, see AM's Specifying Resource Patterns with Wildcards.

Default: `true`

ScriptableResourceUriProvider

Use a script to return the resource URL to include in policy decision requests to AM. The result of the script must be a string that represents the resource to be used in the policy decision request sent to AM.

```
"resourceUriProvider": {
  "type": "ScriptableResourceUriProvider",
  "config": {
    "type": string,
    "file": expression,                    // Use either "file"
    "source": string or array of strings,  // or "source", but not both.
    "args": object,
    "clientHandler": Handler reference
  }
}
```

For information about these properties, see Scripts(5).

The following example script replaces existing query parameters with a single parameter:

```
"resourceUriProvider": {
  "type": "ScriptableResourceUriProvider",
  "config": {
    "type": "application/x-groovy",
    "source": [
      "request.uri.setQuery('fromIG=true')",
      "return request.uri.toASCIIString()"
    ]
  }
}
```

Example Policy Enforcement Using claimsSubject

This route is a variant of the route in "Enforcing AM Policy Decisions In the Same Domain" in the *Gateway Guide*. It enforces a policy decision from AM, using `claimsSubject` instead of `ssoTokenSubject` to identify the subject.

To use this route:

1. Set up AM as described in "To Set Up AM As a PDP" in the *Gateway Guide*.

2. In AM, select the policy you created in the previous step, and add a new resource:

    *   Resource Type: `URL`

    *   Resource pattern: `*://*:*/*`

    *   Resource value: `http://app.example.com:8081/home/pep-claims`

3. In the same policy, add the following subject condition:

    *   `Any of`

    *   Type: `OpenID Connect/JwtClaim`

    *   claimName: `iss`

    *   claimValue: `openam.example.com`

4. Set an environment variable for the Java agent password, and then restart IG:

    ```
    $ export AGENT_SECRET_ID='cGFzc3dvcmQ='
    ```

    The password is retrieved by the default SystemAndEnvSecretStore, and must be base64-encoded.

5. Add the following route to the IG configuration as `$HOME/.openig/config/routes/04-pep-claims.json`. On Windows, add the route as `%appdata%\OpenIG\config\routes\04-pep-claims.json`:

```
{
  "name": "pep-claims",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/pep-claims')}",
  "heap": [
    {
      "name": "AmService-1",
      "type": "AmService",
      "config": {
        "url": "http://openam.example.com:8088/openam",
        "realm": "/",
        "ssoTokenHeader": "iPlanetDirectoryPro",
        "agent": {
          "username": "ig_agent",
          "passwordSecretId" : "agent.secret.id"
        },
        "version": "6.5"
      }
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "name": "SingleSignOnFilter-1",
          "type": "SingleSignOnFilter",
          "config": {
            "amService": "AmService-1"
          }
        },
        {
          "name": "PolicyEnforcementFilter-1",
          "type": "PolicyEnforcementFilter",
          "config": {
            "pepRealm": "/",
            "application": "PEP-SSO" ,
            "claimsSubject": {
              "sub": "${contexts.ssoToken.info.uid}",
              "iss": "openam.example.com"
            },
            "amService": "AmService-1"
          }
        }
      ],
      "handler": "ReverseProxyHandler"
    }
  }
}
```

6. Browse to http://openig.example.com:8080/home/pep-claims.

7. Log in to AM as user `demo`, password `Ch4ng31t`.

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

More Information

org.forgerock.openig.openam.PolicyEnforcementFilter

org.forgerock.openig.openam.PolicyDecisionContext

PolicyDecisionContext(5)

AM Authorization Guide

Name

ScriptableFilter — process requests and responses by using a script

Description

Processes requests and responses by using a Groovy script.

When a ScriptableFilter processes a request, it can execute `return next.handle(context, request)` to call the next filter or handler in the current chain and return the value from the call. Actions on the response must be performed in the Promise's callback methods.

Scripts must return a Promise<Response, NeverThrowsException> or a Response.

This section describes the usage of ScriptableFilter, and refers to the following sections of the documentation:

- For information about script properties, available global objects, and automatically imported classes, see Scripts(5).

- For information about creating scriptable objects in Studio, see "Scripting in Studio" in the *Gateway Guide* and " Configuring a Scriptable Throttling Filter " in the *Gateway Guide*.

Usage

```
{
    "name": string,
    "type": "ScriptableFilter",
    "config": {
        "type": string,
        "file": expression,                    // Use either "file"
        "source": string or array of strings,  // or "source", but not both.
        "args": object,
        "clientHandler": Handler reference
    }
}
```

Properties

For information about properties for ScriptableFilter, see Scripts(5).

Examples

See the following examples of scriptable filters:

- For an example scriptable filter that recovers policy advices from AM, see the `failureHandler` property of PolicyEnforcementFilter(5).

- For an example scriptable filter that calls a groovy file to include headers for cross-origin requests in UMA, see "To Set Up IG As an UMA Resource Server" in the *Gateway Guide*.

More Information

Scripts(5)

org.forgerock.openig.filter.ScriptableFilter

Name

SessionInfoFilter — Capture Information About the AM Session and Store It In the Context

Description

This filter calls the AM endpoint for session information, and makes the data available as a new context to downstream IG filters and handlers. For information, see SessionInfoContext(5). Session properties that are whitelisted in AM are available.

Supported with AM 5 and later versions.

For more information about AM sessions, see Using Sessions in the AM *Authentication and Single Sign-On Guide*.

WebSocket Notifications for Sessions

When WebSocket notifications are set up for sessions, IG receives a notification from AM when a user logs out of AM, or when the AM session is modified, closed, or times out. IG then evicts entries that are related to the event from the `sessionCache`.

For information about setting up WebSocket notifications, using them to clear the session cache, and including them in the server logs, see "WebSocket Notification Service".

Usage

```
{
  "name": string,
  "type": "SessionInfoFilter",
  "config": {
    "amService": AmService reference,
    "ssoToken": runtime expression<string>
  }
}
```

Properties

`"amService"`: *AmService reference, required*

The AmService heap object to use for the following properties:

- `agent`, the credentials of an AM Java agent that acts on behalf of IG to authenticate with AM, get the user profile, and communicate WebSocket notifications from AM to IG.

- `url`, the URL of an AM service to use for session token validation and authentication.

- `amHandler`, the handler to use when communicating with AM to validate the token in the incoming request.

- `realm`: Realm of the AM Java agent that acts on behalf of IG.

- `version`: The version of the AM server.

This filter is compatible with AM version 5 or higher.

See also, AmService(5).

"ssoToken": *runtime expression<string>, optional*

Location of the AM SSO token. For example, with `${request.headers['mySsoToken'].values[0]}` the SSO token is the first value of the mySsoToken header in the request.

Default: `${request.cookies['AmService-ssoTokenHeader'][0].value}`, where `AmService-ssoTokenHeader` is the name of the header or cookie where the AmService expects to find SSO tokens.

Examples

In the following example, the SingleSignOnFilter requires authentication with AM before passing the request along the chain. The SessionInfoFilter collects session info from AM and stores it in the `${contexts.amSession}` context. Then the HeaderFilter adds headers containing some of the session info to the request. The session info is then displayed on the home page of the sample application:

```
{
  "name" : "session-info",
  "baseURI" : "http://app.example.com:8081",
  "condition" : "${matches(request.uri.path, '^/home/session-info')}",
  "heap": [
    {
      "name": "AmService-1",
      "type": "AmService",
      "config": {
        "agent" : {
          "username" : "ig_agent",
          "password" : "password"
        },
        "url": "http://openam.example.com:8088/openam",
        "version": "6.5"
      }
    }
  ],
  "handler" : {
    "type" : "Chain",
    "config" : {
      "filters" : [
        {
          "name": "SingleSignOnFilter",
          "type": "SingleSignOnFilter",
          "config": {
            "amService": "AmService-1"
          }
        },
        {
          "name": "SessionInfoFilter",
```

```
        "type": "SessionInfoFilter",
        "config": {
          "amService": "AmService-1"
        }
      },
      {
        "name": "HeaderFilter",
        "type": "HeaderFilter",
        "config": {
          "messageType": "REQUEST",
          "add": {
            "X-IG-SessionInfo": [ "username: ${contexts.amSession.username}, realm:
 ${contexts.amSession.realm}, properties: ${contexts.amSession.properties}" ]
          }
        }
      }
    ],
    "handler": "ReverseProxyHandler"
    }
  }
}
```

To try this example:

1. In AM, select Applications > Agents > Java, and add a Java agent with the following values:

    *   Name: `ig_agent`

    *   Password: `password`

    *   Server URL: `http://openam.example.com:8088/openam`

    *   Agent URL: `http://openig.example.com:8080/agentapp`

2. With AM, IG, and the sample application running, access the route on http://openig.example.com:8080/home/session-info, and then log in to AM as user `demo`, password `Ch4ng31t`.

    The home page of the sample application is displayed.

3. Note that the header `x-ig-sessioninfo` and its values are displayed:

    ```
    x-ig-sessioninfo: username: demo, realm: /, properties: {}
    ```

To capture additional session properties from AM:

1. On AM, add the property `user-status` to the Session Property Whitelist Service. For information, see Session Property Whitelist Service .

2. On IG, access the route and note that `user-status` is now displayed:

    ```
    x-ig-sessioninfo: username: demo, realm: /, properties: {user-status=}
    ```

To add a value to a property and capture it:

1. Post the value `gold` to `user-status`:

```
$ curl --request POST --header "iPlanetDirectoryPro: <token>"
 \
--header 'Accept-API-Version: resource=2.0'
 \
--header "Content-Type: application/json" \
'http://openam.example.com:8088/openam/json/sessions/?_action=updateSessionProperties'
 \
--data '{"user-status":"gold"}'
```

   Where <token> is the value of the iPlanetDirectoryPro cookie displayed on the home page of the sample application.

2. On IG, access the route again and note that the value of `user-status` is displayed:

```
x-ig-sessioninfo: username: demo, realm: /, properties: {user-status=gold}
```

More Information

org.forgerock.openig.openam.SessionInfoFilter

org.forgerock.openig.openam.SessionInfoContext

SessionInfoContext(5)

AM Authorization Guide

Name
SingleSignOnFilter — require AM authentication before processing a request

Description

This filter tests for the presence and validity of an SSO token in the cookie header of a request:

- If an SSO token is present, the filter calls AM to validate the SSO token. If the SSO token is valid, the request continues along the chain. The token value and additional information are stored in the `ssoToken` context. For information, see SsoTokenContext(5).

- If the SSO token is not present, or is empty or invalid, IG checks the `goto` query parameter for the presence of the `_ig` marker parameter:

  - If the marker parameter is present (`_ig=true`), IG fails the request because the cookie domain is incorrectly configured.

  - If the marker parameter is not present (`_ig=false`), IG redirects the user-agent for authentication to the AM login page or another provided login page.

Supported with AM 5 and later versions.

For more information about SSO, see "About SSO Using the SingleSignOnFilter" in the *Gateway Guide*.

> **Tip**
>
> To prevent issues with performance when accessing large resources, such as .jpg and .js files, consider using the SingleSignOnFilter with the following options:
>
> - The `sessionCache`, so that IG can reuse session token information without repeatedly asking AM to verify the session token.
>
> - A ConditionalFilter, so that requests to access large resources skip the SingleSignOnFilter. For an example configuration, see the example in ConditionalFilter(5).

> **Note**
>
> Each time the SingleSignOnFilter validates an SSO token on AM, the idle timeout for the AM stateful session is reset. As a result, validating a session token triggers a write operation to the AM Core Token Service token store.

For information about token validation in AM, see Using Sessions in the AM *Authentication and Single Sign-On Guide*.

WebSocket Notifications for Sessions

When WebSocket notifications are set up for sessions, IG receives a notification from AM when a user logs out of AM, or when the AM session is modified, closed, or times out. IG then evicts entries that are related to the event from the `sessionCache`.

For information about setting up WebSocket notifications, using them to clear the session cache, and including them in the server logs, see "WebSocket Notification Service".

Usage

```
{
  "name": string,
  "type": "SingleSignOnFilter",
    "config": {
      "amService": AmService reference,
      "realm": string,
      "defaultLogoutLandingPage": expression<uri>,
      "loginEndpoint": expression<uri>,
      "logoutExpression": runtime expression<boolean>
  }
}
```

Properties

`"amService"`: *AmService reference, required*

An AmService object to use for the following properties:

- `agent`, the credentials of an AM Java agent that acts on behalf of IG to authenticate with AM, get the user profile, and communicate WebSocket notifications from AM to IG.

- `realm`: Realm of the AM Java agent that acts on behalf of IG.

- `url`, the URL of an AM service to use for session token validation and authentication when `loginEndpoint` is not specified.

- `ssoTokenHeader`, the name of the cookie that contains the session token created by AM.

- `amHandler`, the handler to use when communicating with AM to validate the token in the incoming request.

- `sessionCache`, the configuration of a cache for session information from AM.

- `version`: The version of the AM server.

See also, AmService(5).

"realm": *string, optional*

The AM realm where the user is authenticated.

Default: The realm declared for `amService`.

"defaultLogoutLandingPage": *expression<uri>, optional*

The URI to which a request is redirected after the user logs out of AM.

This parameter is effective only when `logoutEndpoint` is specified.

Default: None. Processing continues.

"loginEndpoint": *URI expression, optional*

The URL of a service instance for the following tasks:

- Managing authentication and the location to which the request is redirected after authentication.

- Processing policy advices after an AM policy decision denies a request with supported advices. The PolicyEnforcementFilter redirects the request to this URL, with information about how to meet the conditions in the advices.

  For examples of different advice types, and the conditions that cause AM to return advices, see Policy Decision Advice in the AM *Authorization Guide*. For information about supported advice types in IG, see "Using Advices From Policy Decisions".

Default: The value of `url` in `amService`

Authentication can be performed in the following ways:

- Directly through AM, with optional authentication parameters in the query string, such as `service`, `module`, and `realm`. For a list of authentication parameters that you can include in the query string, see Authentication Parameters in the AM *Authentication and Single Sign-On Guide*.

  The value must include a redirect with a `goto` parameter.

  The following example uses AM as the authentication service, and includes the `service` authentication parameter:

```
"loginEndpoint" : "https://openam.example.com/openam?
service=TwoFactor&goto=${urlEncodeQueryParameterNameOrValue(contexts.router.originalUri)}"
```

- Through the URL of another application, with optional authentication parameters in the query string, such as `service`, `module`, and `realm`. The application must create a session with an AM instance to set an SSO token and return the request to the redirect location.

The value can optionally include a redirect with a `goto` parameter or different parameter name.

The following example uses an authentication service that is not AM, and includes a redirect parameter:

```
"loginEndpoint" : "https://authservice.example.com/auth?
redirect=${urlEncodeQueryParameterNameOrValue(contexts.router.originalUri)}"
```

When using this option, review the cookie domains to make sure that cookies set by the authentication server are properly conveyed to the IG instance.

`"logoutExpression"` *runtime expression<boolean>, optional*

An expression to define a condition for logout, based on the request. If the expression evaluates to `true`, the AM authentication token for the end-user is revoked.

If a `defaultLogoutLandingPage` is specified, the request is redirected to that page. Otherwise, the request continues to be processed.

The following example expressions can be used to trigger revocation of the end-user token:

- The request URI contains `/logout`:
  ```
  ${matches(request.uri, '/logout')}
  ```

- The request path ends in `/logout`:
  ```
  ${matches(request.uri.path, '^/logout')}
  ```

- The request query includes the `logOff=true` query parameter:
  ```
  ${matches(request.uri.query, 'logOff=true')}
  ```

Default: Logout is not managed by this filter.

More Information

org.forgerock.openig.openam.SingleSignOnFilter

org.forgerock.openig.openam.SsoTokenContext

SsoTokenContext(5)

Name

SqlAttributesFilter — execute SQL query

Description

Executes a SQL query through a prepared statement and exposes its first result. Parameters in the prepared statement are derived from expressions. The query result is exposed in an object whose location is specified by the `target` expression. If the query yields no result, then the resulting object is empty.

The execution of the query is performed lazily; it does not occur until the first attempt to access a value in the target. This defers the overhead of connection pool, network and database query processing until a value is first required. This also means that the parameters expressions is not evaluated until the object is first accessed.

Usage

```
{
    "name": string,
    "type": "SqlAttributesFilter",
    "config": {
        "dataSource": string,
        "preparedStatement": string,
        "parameters": [ runtime expression<string>, ... ],
        "target": lvalue-expression
    }
}
```

Properties

`"dataSource"`: *string, required*

The JNDI name of the factory for connections to the physical data source.

`"preparedStatement"`: *string, required*

The parameterized SQL query to execute, with `?` parameter placeholders.

`"parameters"`: *array of runtime expressions<string>, optional*

The parameters to evaluate and include in the execution of the prepared statement.

See also Expressions(5).

`"target"`: *lvalue-expression, required*

Expression that yields the target object that will contain the query results.

See also Expressions(5).

Example

Using the user's session ID from a cookie, query the database to find the user logged in and set the profile attributes in the attributes context:

```
{
        "name": "SqlAttributesFilter",
        "type": "SqlAttributesFilter",
        "config": {
                "target": "${attributes.sql}",
                "dataSource": "java:comp/env/jdbc/mysql",
                "preparedStatement": "SELECT f.value AS 'first', l.value AS
                  'last', u.mail AS 'email', GROUP_CONCAT(CAST(r.rid AS CHAR)) AS
                  'roles'
                  FROM sessions s
                  INNER JOIN users u
                  ON ( u.uid = s.uid AND u.status = 1 )
                  LEFT OUTER JOIN profile_values f
                  ON ( f.uid = u.uid AND f.fid = 1 )
                  LEFT OUTER JOIN profile_values l
                  ON ( l.uid = u.uid AND l.fid = 2 )
                  LEFT OUTER JOIN users_roles r
                  ON ( r.uid = u.uid )
                  WHERE (s.sid = ? AND s.uid <> 0) GROUP BY s.sid;",
                "parameters": [ "${request.cookies
                  [keyMatch(request.cookies,'JSESSION1234')]
                  [0].value}" ]
        }
}
```

Lines are folded for readability in this example. In your JSON, keep the values for `"preparedStatement"` and `"parameters"` on one line.

More Information

org.forgerock.openig.filter.SqlAttributesFilter

Name

StaticRequestFilter — create new request

Description

Creates a new request, replacing any existing request. The request can include an entity specified in the `entity` parameter. Alternatively, the request can include a form, specified in the `form` parameter, which is included in an entity encoded in `application/x-www-form-urlencoded` format if request method is `POST`, or otherwise as (additional) query parameters in the URI. The `form` and `entity` parameters cannot be used together when the `method` is set to `POST`.

Usage

```
{
    "name": string,
    "type": "StaticRequestFilter",
    "config": {
        "method": string,
        "uri": runtime expression<string>,
        "version": string,
        "headers": {
            configuration expression<string>: [ runtime expression<string>, ... ], ...
        },
        "form": {
            configuration expression<string>: [ runtime expression<string>, ... ], ...
        },
        "entity": runtime expression<string>
    }
}
```

Properties

`"method"`: *string, required*

The HTTP method to be performed on the resource (for example, `"GET"`).

`"uri"`: *runtime expression<string>, required*

The fully-qualified URI of the resource to access (for example, `"http://www.example.com/resource.txt"`).

The result of the expression must be a string that represents a valid URI, but is not a real `java.net.URI` object. For example, it would be incorrect to use `${request.uri}`, which is not a String but a MutableUri.

`"version"`: *string, optional*

Protocol version. Default: `"HTTP/1.1"`.

`"headers"`: *object, optional*

Header fields to set in the request, with the format `name: [ value, ... ]`.

The `name` field of `headers` specifies the header name. It can be defined by a configuration expression or string. If the configuration expressions for multiple `name` resolve to the same final string, multiple values are associated with the `name`.

The `value` field of `headers` is an array of runtime expressions to evaluate as header values.

In the following example, the name of the header is the value of the configuration time, system variable defined in `cookieHeaderName`. The value of the header is the runtime value stored in `contexts.ssoToken.value`:

```
"headers": {
  "${application['header1Name']}": [
    "${application['header1Value'}"
  ]
```

`"form"`: *object, optional*

A form to include in the request, with the format `param: [ value, ... ]`.

The `param` field of `form` specifies the name of the form parameter. It can be defined by a configuration expression or string. If the configuration expressions for multiple `param` resolve to the same final string, multiple values are associated with the `param`.

The `value` field of `form` is an array of runtime expressions to evaluate as form field values.

When the `method` is set to `POST`, this setting is mutually exclusive with the `entity` setting.

In the following example, the names of the field parameters and the values are hardcoded in the form:

```
"form": {
  "username": [
    "demo"
  ],
  "password": [
    "Ch4ng31t"
  ]
```

In the following example, the names of the field parameters are hardcoded. The values take the first value of `username` and `password` provided in the session:

```
"form": {
  "username": [
    "${session.username[0]}"
  ],
  "password": [
    "${session.password[0]}"
  ]
```

In the following example, the name of the first field param take the value of the expression `${application['formName']}` when it is evaluated at startup. The values take the first value of `username` and `password` provided in the session:

```
"form": {
  "${application['formName']}": [
    "${session.username[0]}"
  ],
  "${application['formPassword']}": [
    "${session.password[0]}"
  ]
```

`"entity"`: *runtime expression<string>, optional*

The entity body to include in the request.

This setting is mutually exclusive with the `form` setting when the `method` is set to `POST`.

See also Expressions(5).

Example

In the following example, IG replaces the browser's original HTTP GET request with an HTTP POST login request containing credentials to authenticate to the sample application. For information about how to set up and test this example, see "Trying IG With a Simple Configuration" in the *Getting Started Guide*.

```
{
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "StaticRequestFilter",
          "config": {
            "method": "POST",
            "uri": "http://app.example.com:8081/login",
            "form": {
              "username": [
                "demo"
              ],
              "password": [
                "Ch4ng31t"
              ]
            }
          }
        }
      ],
      "handler": "ReverseProxyHandler"
    }
  },
  "condition": "${matches(request.uri.path, '^/static')}"
}
```

More Information

org.forgerock.openig.filter.StaticRequestFilter

Name
SwitchFilter — divert requests to another handler if a condition is met

Description

Verifies that a specified condition is met. If the condition is met or no condition is specified, the request is diverted to the associated handler, with no further processing by the switch filter.

Usage

```
{
    "name": string,
    "type": "SwitchFilter",
    "config": {
        "onRequest": [
            {
                "condition": runtime expression<boolean>,
                "handler": Handler reference,
            },
            ...
        ],
        "onResponse": [
            {
                "condition": runtime expression<boolean>,
                "handler": Handler reference,
            },
            ...
        ]
    }
}
```

Properties

`"onRequest"`: *array of objects, optional*

Conditions to test (and handler to dispatch to, if `true`) before the request is handled.

`"onResponse"`: *array of objects, optional*

Conditions to test (and handler to dispatch to, if `true`) after the response is handled.

`"condition"`: *runtime expression<boolean>, optional*

If the expression evaluates to `true`, the request is dispatched to the handler. If no condition is specified, the request is dispatched to the handler unconditionally.

Default: No condition is specified.

See also Expressions(5).

`"handler"`: *Handler reference, required*

Dispatch to this handler if the condition yields `true`.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

See also Handlers.

Example

This example intercepts the response if it is equal to 200 and executes the LoginRequestHandler. This filter might be used in a login flow where the request for the login page must go through to the target, but the response should be intercepted in order to send the login form to the application. This is typical for scenarios where there is a hidden value or cookie returned in the login page, which must be sent in the login form:

```
{
    "name": "SwitchFilter",
    "type": "SwitchFilter",
    "config": {
        "onResponse": [
            {
                "condition": "${response.status.code == 200}",
                "handler": "LoginRequestHandler"
            }
        ]
    }
}
```

More Information

org.forgerock.openig.filter.SwitchFilter

Name

TokenTransformationFilter — transform a token issued by AM to another type

Description

This filter transforms a token issued by AM to another token type.

The TokenTransformationFilter makes the result of the token transformation available to downstream handlers in the `sts` context. For information, see StsContext(5).

The current implementation uses REST Security Token Service (STS) APIs to transform an OpenID Connect ID Token (`id_token`) into a SAML 2.0 assertion. The subject confirmation method is Bearer, as described in *Profiles for the OASIS Security Assertion Markup Language (SAML) V2.0*.

The TokenTransformationFilter makes the result of the token transformation available to downstream handlers in the `issuedToken` property of the `${contexts.sts}` context.

The TokenTransformationFilter configuration references a REST STS instance that must be set up in AM before the TokenTransformationFilter can be used. The REST STS instance exposes a preconfigured transformation under a specific REST endpoint. For information about setting up a REST STS instance, see the AM documentation.

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

Supported with OpenAM 13.5, and AM 5 and later versions.

Usage

```
{
    "name": "string",
    "type": "TokenTransformationFilter",
    "config": {
        "amService": AmService reference,
        "idToken": runtime expression<string>,
        "instance": "expression"
    }
}
```

Properties

`"amService"`: *AmService reference, required*

The AmService heap object to use for the following properties:

- `agent`, the credentials of an AM Java agent to authenticate IG as an AM REST STS client, and to communicate WebSocket notifications from AM to IG. This credentials are evaluated when the route is initialized

- `url`, the URL of an AM service to use for session token validation and authentication. Authentication and REST STS requests are made to this service.

- `realm`, the AM realm containing the following information:

  - The AM application that can make the REST STS request and whose credentials are the username and password.

  - The STS instance described by the instance field.

- `ssoTokenHeader`, the name of the HTTP header that provides the SSO token for the REST STS client subject.

- `amHandler`, the handler to use for authentication and STS requests to AM.

See also, AmService(5).

`"idToken"`: *runtime expression<string>, required*

The value of the OpenID Connect ID token. The expected value is a string that is the JWT encoded `id_token`.

See also Expressions(5).

`"instance"`: *expression, required*

An expression evaluating to the name of the REST STS instance.

This expression is evaluated when the route is initialized, so the expression cannot refer to `request` or `contexts`.

See also Expressions(5).

Example

The following example shows a configuration for a TokenTransformationFilter:

```
{
  "type": "TokenTransformationFilter",
  "config": {
    "amService": "MyAmService",
    "idToken": "${attributes.openid.id_token}",
    "instance": "openig"
  }
}
```

For an example of how to set up and test the TokenTransformationFilter, see "*Transforming OpenID Connect ID Tokens Into SAML Assertions*" in the *Gateway Guide*.

More Information

org.forgerock.openig.openam.TokenTransformationFilter

org.forgerock.openig.openam.StsContext

StsContext(5)

Name

UmaFilter — protect access as an UMA resource server

Description

This filter acts as a policy enforcement point, protecting access as a User-Managed Access (UMA) resource server. Specifically, this filter ensures that a request for protected resources includes a valid requesting party token with appropriate scopes before allowing the response to flow back to the requesting party.

UMA 2.0 is supported with AM 5.5 and later versions. UMA 1.0 is supported with AM 5.1 and later versions.

Usage

```
{
    "type": "UmaFilter",
    "config": {
        "protectionApiHandler": Handler reference,
        "umaService": UmaService reference,
        "realm": string
    }
}
```

Properties

"protectionApiHandler": *Handler reference, required*

The handler to use when interacting with the UMA authorization server for token introspection and permission requests, such as a ClientHandler capable of making an HTTPS connection to the server.

For details, see Handlers.

"umaService": *UmaService reference, required*

The UmaService to use when protecting resources.

For details, see UmaService(5).

"realm": *string, optional*

The UMA realm set in the response to a request for a protected resource that does not include a requesting party token enabling access to the resource.

Default: uma

More Information

User-Managed Access (UMA) Profile of OAuth 2.0

org.forgerock.openig.uma.UmaResourceServerFilter

Name
UserProfileFilter — retrieve user profile attributes and store them in the context

Description

This filter queries AM to retrieve the profile attributes of an user identified by their `username`.

Only profile attributes that are enabled in AM can be returned by the query. The `roles` field is not returned.

The data is made available to downstream IG filters and handlers, in the context UserProfileContext(5).

Supported with AM 5 and later versions.

Usage

```
{
  "name": string,
  "type": "UserProfileFilter",
  "config": {
    "username": configuration expression<string>,
    "userProfileService": UserProfileService reference
  }
}
```

Properties

`"username"`: *configuration expression<string>, required*

The username of an AM subject. This filter retrieves profile attributes for the subject.

`"userProfileService"`: *UserProfileService reference, required*

The service to retrieve profile attributes from AM, for the subject identified by `username`.

```
"userProfileService": {
  "type": "UserProfileService",
  "config": {
    "amService": AmService reference,
    "cache": object,
    "profileAttributes": array of runtime expression<string>,
    "realm": configuration expression<string>
  }
}
```

`"amService"`: *AmService reference, required*

The AmService heap object to use for the following properties:

- `agent`, the credentials of an AM Java agent that acts on behalf of IG to authenticate with AM, get the user profile, and communicate WebSocket notifications from AM to IG.

- `url`: URL of the AM server where the user is authenticated.

- `amHandler`: Handler to use when communicating with AM to fetch the requested user's profile.

- `realm`: Realm of the AM Java agent that acts on behalf of IG.

- `version`: The version of the AM server.

`cache`: *object, optional*

Caching of AM user profiles, based on *Caffeine*. For more information, see the GitHub entry, *Caffeine*.

When caching is enabled, IG can reuse cached profile attributes without repeatedly querying AM. When caching is disabled, IG must query AM for each request, to retrieve the required user profile attributes.

Default: No cache.

`enabled`: *boolean, optional*

Enable or disable caching of user profiles. When `false`, the cache is disabled but the cache configuration is maintained.

Default: `true` when `cache` is configured

`executor`: *executor, optional*

An executor service to schedule the execution of tasks, such as the eviction of entries in the cache.

Default: `ForkJoinPool.commonPool()`

`"maximumSize"`: *configuration expression<number>, optional*

The maximum number of entries the cache can contain.

Default: Unlimited/unbound

`maximumTimeToCache`: *duration, required*

The maximum duration for which to cache user profiles.

The duration cannot be `zero`.

For information about supported formats for `duration`, see duration.

profileAttributes: *array of runtime expression<string>, optional*

List of one or more fields to return and store in UserProfileContext.

Field names are defined by the underlying repository in AM. When AM is installed with the default configuration, the repository is ForgeRock Directory Server.

The following convenience accessors are provided for commonly used fields:

- `cn`: Retrieved through `${contexts.userProfile.commonName}`

- `dn`: Retrieved through `${contexts.userProfile.distinguishedName}`

- `realm`: Retrieved through `${contexts.userProfile.realm}`

- `username`: Retrieved through `${contexts.userProfile.username}`

All other available fields can be retrieved through `${contexts.userProfile.rawInfo}` and `${contexts.userProfile.asJsonValue()}`.

When `profileAttributes` is configured, the specified fields **and** the following fields are returned: `username`, `_id`, and `_rev`.

Default: All available fields are returned.

"realm": *configuration expression<string>, optional*

The AM realm where the subject is authenticated.

Default: The realm declared for `amService`.

Examples

In the following examples, the UserProfileFilter retrieves user profile attributes and stores them in the context.

The `userProfile` property of AmService is configured to retrieve `employeeNumber` and `mail`. When the property is not configured, all available attributes in `rawInfo` or `asJsonValue()` are displayed.

Retrieving Profile Attributes for a User Authenticated With an SSO Token

In this example the user is authenticated with AM through the SingleSignOnFilter, which stores the SSO token and its validation information in the `SsoTokenContext`. The UserProfileFilter retrieves the user's mail and employee number, as well as the `username`, `_id`, and `_rev`, from that context.

1. In AM, add the following items, and then log out of AM:

   • A subject, as described in "To Set Up a Sample User In AM" in the *Gateway Guide*

   • A Java agent, as described in "To Set Up a Java Agent in AM" in the *Gateway Guide*

2. In IG, add the following route to the configuration as `$HOME/.openig/config/routes/user-profile-sso.json` (on Windows, `%appdata%\OpenIG\config\routes\user-profile-sso.json`):

```
{
  "name" : "user-profile-sso",
  "condition" : "${matches(request.uri.path, '^/user-profile-sso')}",
  "heap": [
    {
      "name": "AmService-1",
      "type": "AmService",
      "config": {
        "url": "http://openam.example.com:8088/openam",
        "realm": "/",
        "ssoTokenHeader": "iPlanetDirectoryPro",
        "version": "6.5",
        "agent": {
          "username": "ig_agent",
          "password": "password"
        },
        "amHandler": "ForgeRockClientHandler"
      }
    }
  ],
  "handler" : {
    "type" : "Chain",
    "config" : {
      "filters" : [
        {
          "name": "SingleSignOnFilter",
          "type": "SingleSignOnFilter",
          "config": {
            "amService": "AmService-1"
          }
        },
        {
          "name": "UserProfileFilter-1",
          "type": "UserProfileFilter",
          "config": {
            "username": "${contexts.ssoToken.info.uid}",
            "userProfileService": {
              "type": "UserProfileService",
              "config": {
                "amService": "AmService-1",
                "profileAttributes": [ "employeeNumber", "mail" ]
              }
            }
          }
        }
      ],
      "handler": {
        "type": "StaticResponseHandler",
        "config": {
```

```
         "status": 200,
         "headers": {
           "Content-Type": ["text/html"]
         },
         "entity": "<html><body>username: ${contexts.userProfile.username}<br><br>rawInfo: <pre>
${contexts.userProfile.rawInfo}</pre></body></html>"
       }
     }
   }
 }
}
```

3. Browse to http://openig.example.com:8080/user-profile-sso.

4. Log in to AM as user `george`, password `C0stanza`.

   The UserProfileFilter retrieves George's profile data and stores it in the UserProfileContext. The StaticResponseHandler displays the username and the profile data that is available in `rawInfo`.

```
username: george

rawInfo:
{_id=george, _rev=273001616, employeeNumber=[123], mail=[george@example.com], username=george}
```

Retrieving a Username From sessionInfo Context

This example shows how the UserProfileFilter retrieves AM profile information for the user identified by the SessionInfoContext, at `${contexts.amSession.username}`.

Consider using the UserProfileFilter with a SessionInfoFilter to protect APIs. The SessionInfoFilter validates an SSO token without redirecting the request to an authentication page.

1. In AM, add the following items, and then log out of AM:

   - A subject, as described in "To Set Up a Sample User In AM" in the *Gateway Guide*

   - A Java agent, as described in "To Set Up a Java Agent in AM" in the *Gateway Guide*

2. Add the following route to the configuration as `$HOME/.openig/config/routes/user-profile-ses-info` (on Windows, `%appdata%\OpenIG\config\routes\user-profile-ses-info.json`):

```
{
  "name" : "user-profile-ses-info",
  "condition" : "${matches(request.uri.path, '^/user-profile-ses-info')}",
  "heap": [
    {
      "name": "AmService-1",
      "type": "AmService",
      "config": {
        "url": "http://openam.example.com:8088/openam",
        "realm": "/",
        "ssoTokenHeader": "iPlanetDirectoryPro",
        "version": "6.5",
```

```
          "agent": {
            "username": "ig_agent",
            "password": "password"
          },
          "amHandler": "ForgeRockClientHandler"
        }
      }
    ],
    "handler" : {
      "type" : "Chain",
      "capture": "all",
      "config" : {
        "filters" : [
          {
            "name": "SessionInfoFilter-1",
            "type": "SessionInfoFilter",
            "config": {
              "amService": "AmService-1"
            }
          },
          {
            "name": "UserProfileFilter-1",
            "type": "UserProfileFilter",
            "config": {
              "username": "${contexts.amSession.username}",
              "userProfileService": {
                "type": "UserProfileService",
                "config": {
                  "amService": "AmService-1",
                  "profileAttributes": [ "employeeNumber", "mail" ]
                }
              }
            }
          }
        ],
        "handler": {
          "type": "StaticResponseHandler",
          "config": {
            "status": 200,
            "headers": {
              "Content-Type": [ "application/json" ]
            },
            "entity": "{ \"username\": \"${contexts.userProfile.username}\", \"user_profile\":
${contexts.userProfile.asJsonValue()} }"
          }
        }
      }
    }
  }
}
```

3. In a terminal window, retrieve an SSO token:

```
$ curl --request POST
 \
--url http://openam.example.com:8088/openam/json/realms/root/authenticate
 \
--header 'accept-api-version: resource=2.0'
 \
--header 'content-type: application/json'
 \
--header 'x-openam-username: george'
 \
--header 'x-openam-password: C0stanza'
 \
--data '{}'

{"tokenId":"AQIC5wM2LY . . . Dg5AAJTMQAA*","successUrl":"/openam/console"}
```

4. Access the route, providing the token ID retrieved in the previous step:

```
$ curl --cookie 'iPlanetDirectoryPro=<id_token>' http://openig.example.com:8080/user-profile-ses-info
 | jq .

{
  "username": "george",
  "user_profile": {
    "_id": "george",
    "_rev": "273001616",
    "employeeNumber": [
      "123"
    ],
    "mail": [
      "george@example.com"
    ],
    "username": "george"
  }
}
```

The UserProfileFilter retrieves George's profile data and stores it in the UserProfileContext. The StaticResponseHandler displays the username and the profile data that is available in `asJsonValue()`.

Retrieving a Username From the OAuth2Context

In this example the OAuth2ResourceServerFilter validates a request containing an OAuth 2.0 access token, using the introspection endpoint, and injects the token into the OAuth2Context context. The UserProfileFilter retrieves AM profile information for the user identified by this context.

1. In AM, add the following items, and then log out of AM:

   • A subject, as described in "To Set Up a Sample User In AM" in the *Gateway Guide*

   • A Java agent, as described in "To Set Up a Java Agent in AM" in the *Gateway Guide*

   • An OAuth 2.0 Authorization Server and profiles, as described in "*Acting as an OAuth 2.0 Resource Server*" in the *Gateway Guide*.

2. Add the following route to the configuration as `$HOME/.openig/config/routes/user-profile-oauth.json` (on Windows, `%appdata%\OpenIG\config\routes\user-profile-oauth.json`):

```json
{
  "name": "user-profile-oauth",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/user-profile-oauth')}",
  "heap": [
    {
      "name": "AmService-1",
      "type": "AmService",
      "config": {
        "url": "http://openam.example.com:8088/openam",
        "realm": "/",
        "ssoTokenHeader": "iPlanetDirectoryPro",
        "version": "6.5",
        "agent": {
          "username": "ig_agent",
          "password": "password"
        },
        "amHandler": "ForgeRockClientHandler"
      }
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "name": "OAuth2ResourceServerFilter-1",
          "type": "OAuth2ResourceServerFilter",
          "config": {
            "scopes": [
              "mail",
              "employeenumber"
            ],
            "requireHttps": false,
            "realm": "OpenIG",
            "accessTokenResolver": {
              "name": "token-resolver-1",
              "type": "TokenIntrospectionAccessTokenResolver",
              "config": {
                "endpoint": "http://openam.example.com:8088/openam/oauth2/introspect",
                "providerHandler": {
                  "type": "Chain",
                  "config": {
                    "filters": [
                      {
                        "type": "HeaderFilter",
                        "config": {
                          "messageType": "request",
                          "add": {
                            "Authorization": [
                              "Basic ${encodeBase64('resource-server:password')}"
                            ]
                          }
                        }
                      }
                    ],
```

```
                        "handler": "ForgeRockClientHandler"
                      }
                    }
                  }
                }
              }
            },
            {
              "name": "UserProfileFilter-1",
              "type": "UserProfileFilter",
              "config": {
                "username": "${contexts.oauth2.accessToken.info.sub}",
                "userProfileService": {
                  "type": "UserProfileService",
                  "config": {
                    "amService": "AmService-1",
                    "profileAttributes": [ "employeeNumber", "mail" ]
                  }
                }
              }
            }
          ],
          "handler": {
            "type": "StaticResponseHandler",
            "config": {
              "status": 200,
              "headers": {
                "Content-Type": [ "application/json" ]
              },
              "entity": "{ \"username\": \"${contexts.userProfile.username}\", \"user_profile\":
 ${contexts.userProfile.asJsonValue()} }"
            }
          }
        }
      }
    }
  }
}
```

3. In a terminal window, retrieve an access token:

```
$ curl \
--user "client-application:password"
 \
--data "grant_type=password&username=george&password=C0stanza&scope=mail%20employeenumber" \
http://openam.example.com:8088/openam/oauth2/access_token

{
  "access_token":"<access_token>",
  "scope":"employeenumber mail",
  "token_type":"Bearer",
  "expires_in":3599
}
```

4. Access the route, providing the access token retrieved in the previous step:

```
$ curl --header 'Authorization: Bearer <access_token>' http://openig.example.com:8080/user-profile-
oauth | jq .

{
  "username": "george",
  "user_profile": {
    "_id": "george",
    "_rev": "-576067013",
    "employeeNumber": [
      "123"
    ],
    "mail": [
      "george@example.com"
    ],
    "username": "george"
  }
}
```

The UserProfileFilter retrieves George's profile data and stores it in the UserProfileContext. The StaticResponseHandler displays the username and the profile data that is available in `asJsonValue()`.

More Information

org.forgerock.openig.openam.UserProfileFilter

org.forgerock.openig.tools.userprofile.UserProfileService

org.forgerock.openig.openam.UserProfileContext

UserProfileContext(5)

AM Authorization Guide

Decorators

Decorators are additional heap objects that let you extend what another object can do. For example, a *CaptureDecorator* extends the capability of filters and handlers to log requests and responses.

1. Predefined Decorators

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

2. Guidelines for Naming Decorators

To prevent unwanted behavior, avoid naming decorators as follows:

- For heap objects, avoid decorators named `config`, `name`, and `type`.

- For routes, avoid decorators named `auditService`, `baseURI`, `condition`, `globalDecorators`, `heap`, `handler`, `name`, and `session`.

- In `config.json`, avoid decorators named `temporaryStorage`.

- In general, avoid decorators named `comment` or `comments`.

Use names that do not clash with field names for the decorated objects. The best way to avoid a clash is to avoid IG reserved field names, which include all purely alphanumeric field names. Instead use dots in your decorator names, such as `my.decorator`.

3. Local, Route, and Global Decorators

IG applies decorations in this order:

1. Local decoration - decorators declared on an object

2. globalDecoration - decorators declared on a parent route are applied before decorators declared in the current route.

3. Route decoration: decorators declared on a route handler

Decorations are inherited as follows:

- Local decorations that are part of an object's declaration are inherited wherever the object is used.

• The globalDecorations on a route are inherited on child routes.

To prevent loops, decorators themselves cannot be decorated. Instead, decorators apply only to specific types of objects such as filters and handlers.

Decorations can apply more than once. For example, if you set a decoration both on a route and also on an object defined within the route, then IG can apply the decoration twice. The following route results in the request being captured twice:

```
{
    "handler": {
        "type": "ReverseProxyHandler",
        "capture": "request"
    },
    "capture": "all"
}
```

4. Deterministic Order of Decorators

When an object has multiple decorations, the decorations are applied in the order they appear in the JSON.

In the following route, the handler is decorated with a `baseURI` first, and a `capture` second:

```
{
  "name": "myroute",
  "baseURI": "http://app.example.com:8081",
  "capture": "all",
  "handler": {
    "type": "StaticResponseHandler",
    "config": {
      "status": 200,
      "reason": "OK",
      "entity": "Hello, world from myroute!"
    }
  },
  "condition": "${matches(request.uri.path, '^/myroute1')}"
}
```

The decoration can be represented as `capture[ baseUri[ handler ] ]`. When a request is processed, it is captured, and then rebased, and then processed by the handler: The log for this route shows that the capture occurs before the rebase:

```
2018-09-10T13:23:18,990Z | INFO  | http-nio-8080-exec-1 | o.f.o.d.c.C.c.top-level-handler | @myroute |

--- (request) id:f792d2ad-4409-4907-bc46-28e1c3c19ac3-7 --->

GET http://openig.example.com:8080/myroute HTTP/1.1
...
```

Conversely, in the following route, the handler is decorated with a `capture` first, and a `baseURI` second:

```
{
  "name": "myroute",
  "capture": "all",
  "baseURI": "http://app.example.com:8081",
  "handler": {
    "type": "StaticResponseHandler",
    "config": {
      "status": 200,
      "reason": "OK",
      "entity": "Hello, world from myroute1!"
    }
  },
  "condition": "${matches(request.uri.path, '^/myroute')}"
}
```

The decoration can be represented as `baseUri[ capture[ handler ] ]`. When a request is processed, it is rebased, and then captured, and then processed by the handler. The log for this route shows that the rebase occurs before the capture:

```
2018-09-10T13:07:07,524Z | INFO  | http-nio-8080-exec-1 | o.f.o.d.c.C.c.top-level-handler | @myroute |

--- (request) id:3c26ab12-3cc0-403e-bec6-43bf5621f657-7 --->

GET http://app.example.com:8081/myroute HTTP/1.1
...
```

5. Decorating Individual Objects

To decorate individual objects, add the decorator's name value as a top-level field of the object, next to `type` and `config`.

6. Decorating the Handler for a Route

To decorate the handler for a route, add the decorator as a top-level field of the route. The following route includes a timer decoration on the handler. This configuration decorates the ReverseProxyHandler only for the current route. It does not decorate other uses of ReverseProxyHandler in other routes:

```
{
    "handler": "ReverseProxyHandler",
    "timer": true
}
```

7. Decorating All Objects in a Route's Heap

To decorate all applicable objects defined in a route's heap, configure globalDecorators as a top-level field of the Route. The globalDecorators field takes a map of the decorations to apply.

In the following example, the route has audit and capture decorations that apply to the Chain, HeaderFilter, and StaticResponseHandler. The decorations apply to all objects in this route's heap:

```
{
    "globalDecorators": {
        "timer": true,
        "capture": "all"
    },
    "handler": {
        "type": "Chain",
        "config": {
            "filters": [
                {
                    "type": "HeaderFilter",
                    "config": {
                        "messageType": "RESPONSE",
                        "add": [
                            {
                                "X-Powered-By": [
                                    "IG"
                                ]
                            }
                        ]
                    }
                }
            ],
            "handler": {
                "type": "StaticResponseHandler",
                "config": {
                    "status": 200,
                    "entity": "Hello World"
                }
            }
        }
    },
    "condition": "${matches(request.uri.path, '^/static')}"
}
```

8. Decorating Individual Uses of a Named Filter or Handler

When a named filter or handler is configured in `config.json` or in the heap, it can be used many times in the configuration. To decorate each use of the filter or handler individually, use a delegate filter or handler, and decorate the delegate. In this way you can decorate the filter or handler differently each time you use it in the configuration.

```
{
    "filter or handler": {
        "type": "Delegate",
        "config": {
            "delegate": [object reference]
        },
        [decorator reference, ...]
    }
}
```

A delegate can be used to decorate a filter, a handler, or any other object type.

In the following example, an AmService heap object configures an `amHandler` to delegate tasks to `ForgeRockClientHandler`, and capture all requests and responses passing through the handler.

```
{
  "name": "AmService-1",
  "type": "AmService",
  "config": {
    "agent" : {
      "username" : "ig_agent",
      "password" : "password"
    },
    "amHandler": {
      "type": "Delegate",
      "config": {
        "delegate": "ForgeRockClientHandler"
      },
      "capture": "all"
    },
    "url": "http://openam.example.com:8088/openam",
    "version": "6.0"
  }
}
```

For information about which filters can use the AmService heap object, see AmService(5).

You can use the same `ForgeRockClientHandler` in another part of the configuration, in a different route for example, without adding a capture decorator. Requests and responses that pass through that use of the handler are not captured.

Table of Contents

Name

BaseUriDecorator — override scheme, host, and port of request URI

Description

Overrides the scheme, host, and port of the existing request URI, rebasing the URI and so making requests relative to a new base URI. Rebasing changes only the scheme, host, and port of the request URI. Rebasing does not affect the path, query string, or fragment.

Decorator Usage

```
{
    "name": string,
    "type": "BaseUriDecorator"
}
```

A BaseUriDecorator does not have configurable properties.

IG creates a default BaseUriDecorator named baseURI at startup time in the top-level heap, so you can use baseURI as the decorator name without adding the decorator declaration explicitly.

Decorated Object Usage

```
{
    "name": string,
    "type": string,
    "config": object,
    decorator name: string
}
```

`"name"`: **string, required except for inline objects**

The unique name of the object, just like an object that is not decorated

`"type"`: **string, required**

The class name of the decorated object, which must be either a Filter or a Handler.

See also Filters and Handlers.

`"config"`: **object, required unless empty**

The configuration of the object, just like an object that is not decorated

decorator name: **string, required**

A string representing the scheme, host, and port of the new base URI. The port is optional when using the defaults (80 for HTTP, 443 for HTTPS).

IG ignores this setting if the value is not a string.

Examples

Add a custom decorator to the heap named myBaseUri:

```
{
    "name": "myBaseUri",
    "type": "BaseUriDecorator"
}
```

Set a Router's base URI to `https://www.example.com:8443`:

```
{
    "name": "Router",
    "type": "Router",
    "myBaseUri": "https://www.example.com:8443/"
}
```

More Information

org.forgerock.openig.decoration.baseuri.BaseUriDecorator

Name

CaptureDecorator — capture request and response messages

Description

Captures request and response messages in log files.

> **Important**
>
> During debugging, consider using a CaptureDecorator to capture the entity and context of requests and responses. However, increased logging consumes resources, such as disk space, and can cause performance issues. In production, reduce logging by disabling the CaptureDecorator properties `captureEntity` and `captureContext`, or setting `maxEntityLength`.

For information about how to set up capture in Studio, see "Capturing Log Messages for Routes" in the *Gateway Guide*.

Decorator Usage

```
{
    "name": string,
    "type": "CaptureDecorator",
    "config": {
        "captureEntity": boolean,
        "captureContext": boolean,
        "maxEntityLength": number
    }
}
```

Captured information is written to SLF4J logs, and named in this format:

```
org.forgerock.openig.decoration.capture.CaptureDecorator.<decoratorName>.<decoratedObjectName>
```

If the decorated object is not named, the object path is used in the log.

The decorator configuration has these properties:

`"captureEntity"`: *boolean, optional*

Whether the message entity should be captured. The message entity is the body of the HTTP message, which can be a JSON document, XML, HTML, image, or other information.

The filter omits binary entities, instead writing a `[binary entity]` marker to the file.

Default: false

`"captureContext"`: *boolean, optional*

Whether the context should be captured as JSON. The context chain is used when processing the request inside IG in the filters and handlers.

Default: false

`"maxEntityLength"`: *number, optional*

The maximum number of bytes that can be captured for an entity. This property is used when `captureEntity` is `true`.

If the captured entity is bigger than `maxEntityLength`, everything up to `maxEntityLength` is captured, and an `[entity truncated]` message is written in the log.

Set `maxEntityLength` to be big enough to allow capture of normal entities, but small enough to prevent excessive memory use or `OutOfMemoryError` errors. Setting `maxEntityLength` to 2 GB or more causes an exception at startup.

Default: 524 288 bytes (512 KB)

Decorated Object Usage

```
{
    "name": string,
    "type": string,
    "config": object,
    decorator name: capture point(s)
}
```

`"name"`: string, required except for inline objects

The unique name of the object, just like an object that is not decorated

`"type"`: string, required

The class name of the decorated object, which must be either a Filter or a Handler.

See also Filters and Handlers.

`"config"`: object, required unless empty

The configuration of the object, just like an object that is not decorated

***decorator name*: capture point(s), optional**

The *decorator name* must match the name of the CaptureDecorator. For example, if the CaptureDecorator has `"name": "capture"`, then *decorator name* is capture.

The capture point(s) are either a single string, or an array of strings. The strings are documented here in lowercase, but are not case-sensitive:

`"all"`

Capture at all available capture points.

`"none"`

Disable capture.

If `none` is configured with other capture points, `none` takes precedence.

"request"

Capture the request as it enters the Filter or Handler.

"filtered_request"

Capture the request as it leaves the Filter.

Only applies to Filters.

"response"

Capture the response as it enters the Filter or leaves the Handler.

"filtered_response"

Capture the response as it leaves the Filter.

Only applies to Filters.

Examples

Decorator configured to log the entity:

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

Decorator configured not to log the entity:

```
{
    "name": "capture",
    "type": "CaptureDecorator"
}
```

Decorator configured to log the context in JSON format, excluding the request and the response:

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

To capture requests and responses with the entity before sending the request and before returning the response, do so as in the following example:

```
{
    "heap": [
        {
            "name": "capture",
            "type": "CaptureDecorator",
            "config": {
                "captureEntity": true
            }
        },
        {
            "name": "ReverseProxyHandler",
            "type": "ReverseProxyHandler",
            "capture": [
                "request",
                "response"
            ]
        }
    ],
    "handler": "ReverseProxyHandler"
}
```

To capture all transformed requests and responses as they leave filters, decorate the Route as in the following example. This Route uses the default CaptureDecorator:

```
{
    "handler": {
        "type": "Chain",
        "config": {
            "filters": [
                {
                    "type": "HeaderFilter",
                    "config": {
                        "messageType": "REQUEST",
                        "add": {
                            "X-RequestHeader": [
                                "Capture at filtered_request point",
                                "And at filtered_response point"
                            ]
                        }
                    }
                },
                {
                    "type": "HeaderFilter",
                    "config": {
                        "messageType": "RESPONSE",
                        "add": {
                            "X-ResponseHeader": [
                                "Capture at filtered_response point"
                            ]
                        }
                    }
                }
            ],
            "handler": {
                "type": "StaticResponseHandler",
                "config": {
                    "status": 200,
                    "reason": "OK",
```

```
                        "entity": "<html><body><p>Hello, World!</p></body></html>"
                    }
                }
            }
        },
        "capture": [
            "filtered_request",
            "filtered_response"
        ]
}
```

To capture the context as JSON, excluding the request and response, before sending the request and before returning the response, do so as in the following example:

```
{
    "heap": [
        {
            "name": "capture",
            "type": "CaptureDecorator",
            "config": {
                "captureContext": true
            }
        },
        {
            "name": "ReverseProxyHandler",
            "type": "ReverseProxyHandler",
            "capture": [
                "request",
                "response"
            ]
        }
    ],
    "handler": "ReverseProxyHandler"
}
```

More Information

org.forgerock.openig.decoration.capture.CaptureDecorator

Name

TimerDecorator — record times to process filters and handlers

Description

Records time in milliseconds to process filters and handlers.

Decorator Usage

```
{
    "name": string,
    "type": "TimerDecorator",
    "config": {
        "timeUnit": string
    }
}
```

IG configures a default TimerDecorator named `timer`. You can use `timer` as the decorator name without explicitly declaring a decorator named timer.

"timeUnit": *duration string, optional*

Unit of time used in the decorator output. The unit of time can be any unit allowed in the `duration` field.

Default: `ms`

For information about supported formats for `duration`, see duration.

Decorated Object Usage

```
{
    "name": string,
    "type": string,
    "config": object,
    decorator name: boolean
}
```

"name": string, required except for inline objects

The unique name of the object, just like an object that is not decorated.

"type": string, required

The class name of the decorated object, which must be either a filter or a handler.

See also Filters and Handlers.

"config": object, required unless empty

The configuration of the object, just like an object that is not decorated.

`decorator name`: boolean, required

> IG looks for the presence of the *decorator name* field for the TimerDecorator.
>
> To activate the timer, set the value of the *decorator name* field to `true`.
>
> To deactivate the TimerDecorator temporarily, set the value to `false`.

Timer Metrics At the Prometheus Scrape Endpoint

This section describes the timer metrics recorded at the Prometheus Scrape Endpoint. For more information about metrics, see Monitoring Endpoints(5).

When IG is set up as described in "*First Steps*" in the *Getting Started Guide*, the endpoint is http://openig.example.com:8080/openig/metrics/prometheus.

Each timer metric is labelled with the following fully qualified names:

- `decorated_object`

- `heap`

- `name` (decorator name)

- `route`

- `router`

Timer Metrics At the Prometheus Scrape Endpoint

Name	Type [a]	Description
`ig_timerdecorator_handler_elapsed_seconds`	Summary	Time to process the request and response in the decorated handler.
`ig_timerdecorator_filter_elapsed_seconds`	Summary	Time to process the request and response in the decorated filter **and** its downstream filters and handler.
`ig_timerdecorator_filter_internal_seconds`	Summary	Time to process the request and response in the decorated filter.
`ig_timerdecorator_filter_downstream_seconds`	Summary	Time to process the request and response in filters and handlers that are downstream of the decorated filter.

[a]As described in Monitoring Types(5)

Timer Metrics At the Common REST Monitoring Endpoint

This section describes the metrics recorded at the the ForgeRock Common REST Monitoring Endpoint. For more information about metrics, see Monitoring Endpoints(5).

When IG is set up as described in "*First Steps*" in the *Getting Started Guide*, the endpoint is http://openig.example.com:8080/openig/metrics/api?_queryFilter=true.

Metrics are published with an `_id` in the following pattern:

- *heap.router-name.route-name.decorator-name.object*

Timer Metrics At the Common REST Monitoring Endpoint

Name	Type [a]	Description
elapsed	Timer	Time to process the request and response in the decorated handler, or in the decorated filter **and** its downstream filters and handler.
internal	Timer	Time to process the request and response in the decorated filter.
downstream	Timer	Time to process the request and response in filters and handlers that are downstream of the decorated filter.

[a]As described in Monitoring Types(5)

Timer Metrics in SLF4J Logs

SLF4J logs are named in this format:

```
<className>.<decoratorName>.<decoratedObjectName>
```

If the decorated object is not named, the object path is used in the log.

When a route's top-level handler is decorated, the timer decorator records the elapsed time for operations traversing the whole route:

```
2018-09-04T12:16:08,994Z | INFO  | I/O dispatcher 17 | o.f.o.d.t.T.t.top-level-handler | @myroute |
  Elapsed time: 13 ms
```

When an individual handler in the route is decorated, the timer decorator records the elapsed time for operations traversing the handler:

```
2018-09-04T12:44:02,161Z | INFO  | http-nio-8080-exec-8 | o.f.o.d.t.T.t.StaticResponseHandler-1 | @myroute
  | Elapsed time: 1 ms
```

More Information

org.forgerock.openig.decoration.timer.TimerDecorator

Audit Framework

IG uses the ForgeRock common audit framework to log system boundary events using an implementation that is common across the ForgeRock platform.

Audit logs use timestamps in UTC format (for example, 2018-07-18T08:48:00.160Z), a unified standard that is not affected by time changes for daylight savings. The timestamps format is not configurable.

Table of Contents

Name
AuditService — configure common audit service for a route

Description

Configure the audit service for a route. The audit service uses the ForgeRock common audit event framework.

The route is decorated with an `auditService` field whose value references the configuration, either inline or from the heap.

By default, the following fields are included in an audit message:

- `eventName`
- `timestamp`
- `transactionId`
- `server`
  - `ip`
  - `port`
- `client`
  - `ip`
  - `port`
- `http`
  - `request`
    - `secure`
    - `method`
    - `path`
    - `queryParameters`
    - `headers`
    - `cookies`
  - `response`
    - `status`

- statusCode

- elapsedTime

- elapsedTimeUnits

- _id

> **Important**
>
> In your AuditService configuration, consider using the following filter policy to exclude sensitive data from log files:

```
{
  "name": "MyAuditService",
  "type": "AuditService",
  "config": {
    "config": {
      "filterPolicies": {
        "field": {
          "excludeIf": [
            "/access/http/request/cookies/iPlanetDirectoryPro",
            "/access/http/request/headers/iPlanetDirectoryPro",
            "/access/http/request/headers/AMAuthCookie",
            "/access/http/request/headers/authorization",
            "/access/http/request/headers/proxy-authorization",
            "/access/http/request/headers/X-OpenAM-Password",
            "/access/http/request/headers/X-OpenIDM-Password",
            "/access/http/request/queryParameters/access_token",
            "/access/http/request/queryParameters/id_token_hint",
            "/access/http/request/queryParameters/IDToken1",
            "/access/http/request/queryParameters/Login.Token1",
            "/access/http/request/queryParameters/redirect_uri",
            "/access/http/request/queryParameters/requester",
            "/access/http/request/queryParameters/sessionUpgradeSSOTokenId",
            "/access/http/request/queryParameters/tokenId",
            "/access/http/response/headers/Set-Cookie"
          ]
        }
      }
    },
    "event-handlers": MyEventHandler
  }
}
```

If you use a different name for `iPlanetDirectoryPro` or `AMAuthCookie`, edit the list to include the correct name.

Usage

```
{
  "name": string,
  "type": "AuditService",
  "config": {
    "config": object,
    "event-handlers": array
  }
}
```

Properties

"config": *object, required*

Configures the audit service itself, rather than event handlers. If the configuration uses only default settings, you can omit the field instead of including an empty object as the field value.

The configuration object has the following fields:

"handlerForQueries": *string, optional*

The name of the event handler to use when querying audit event messages over REST.

"availableAuditEventHandlers": *array of strings, optional*

A list of fully qualified event handler class names for event handlers available to the audit service.

"caseInsensitiveFields": *array of strings, optional*

A list of audit event fields to be considered as case-insensitive for filtering. The fields are referenced using JSON pointer syntax. The list can be `null` or empty.

Default: `/access/http/request/headers` and `/access/http/response/headers` fields are considered case-insensitive for filtering. All other fields are considered case-sensitive.

"filterPolicies": *object, optional*

Specifies fields to exclude from audit event messages.

"field": *object, optional*

Audit event fields use JSON pointer notation, and are taken from the JSON schema for the audit event content. This property specifies which fields to exclude:

Default: Include all fields

"excludeIf": *array of strings, optional*

A list of audit event fields to exclude from log files. Use this property to exclude passwords, authentication cookies or other sensitive information from the logs.

The following example excludes audit event messages about all cookies in the request:

```
"filterPolicies": {
  "field": {
    "excludeIf": [ "/access/http/request/cookies" ]
  }
}
```

The following example excludes audit event messages about a cookie named `session-cookie`, but leaves messages for other cookies:

```
"filterPolicies": {
  "field": {
    "excludeIf": [ "/access/http/request/cookies/session-cookie" ]
  }
}
```

"event-handlers": *array of configuration objects, required*

This array of audit event handler configuration objects defines the event handlers that deal with audit events.

Each event handler configuration depends on type of the event handler.

IG supports the following audit event handlers:

- CsvAuditEventHandler(5)

- ElasticsearchAuditEventHandler(5)

- JdbcAuditEventHandler(5)

- JmsAuditEventHandler(5)

- JsonAuditEventHandler(5)

- JsonStdoutAuditEventHandler(5)

- SplunkAuditEvenHandler(5)

- SyslogAuditEventHandler(5)

Example

The following example configures an audit service to log access event messages in a comma-separated variable file, named `/path/to/audit/logs/access.csv`:

```
{
  "name": "AuditService",
  "type": "AuditService",
  "config": {
    "config": {},
    "event-handlers": [
      {
        "class": "org.forgerock.audit.handlers.csv.CsvAuditEventHandler",
        "config": {
          "name": "csv",
          "logDirectory": "/path/to/audit/logs",
          "topics": [
            "access"
          ]
        }
      }
    ]
  }
}
```

The following example route uses the audit service:

```
{
  "handler": "ReverseProxyHandler",
  "auditService": "AuditService"
}
```

More Information

org.forgerock.audit.AuditService

Name

CsvAuditEventHandler — log audit events to CSV format files

Description

An audit event handler that responds to events by logging messages to files in comma-separated variable (CSV) format.

The configuration is declared in an audit service configuration. For details, see AuditService(5).

> **Important**
>
> The CSV handler does not sanitize messages when writing to CSV log files.
>
> Do not open CSV logs in spreadsheets and other applications that treat data as code.

Usage

```
{
  "class": "org.forgerock.audit.handlers.csv.CsvAuditEventHandler",
  "config": {
    "name": string,
    "logDirectory": string,
    "topics": array,
    "enabled": boolean,
    "formatting": {
      "quoteChar": single-character string,
      "delimiterChar": single-character string,
      "endOfLineSymbols": string
    },
    "buffering": {
      "enabled": boolean,
      "autoFlush": boolean
    },
    "security": {
      "enabled": boolean,
      "filename": string,
      "password": string,
      "signatureInterval": duration
    },
    "fileRetention": {
      "maxDiskSpaceToUse": number,
      "maxNumberOfHistoryFiles": number,
      "minFreeSpaceRequired": number,
      "rotationRetentionCheckInterval": duration
    },
    "fileRotation": {
      "rotationEnabled": boolean,
      "maxFileSize": number,
      "rotationFilePrefix": string,
      "rotationFileSuffix": string,
      "rotationInterval": duration,
      "rotationTimes": array
    }
```

```
    }
}
```

The values in this configuration object can use expressions as long as they resolve to the correct types for each field. For details about expressions, see Expressions(5).

Configuration

The `"config"` object has the following properties:

`"name"`: *string, required*

The name of the event handler.

`"logDirectory"`: *string, required*

The file system directory where log files are written.

`"topics"`: *array of strings, required*

The topics that this event handler intercepts.

IG handles access events that occur at the system boundary, such as arrival of the initial request and departure of the final response.

Set this to `"topics": [ "access" ]`.

`"enabled"`: *boolean, optional*

Whether this event handler is active.

Default: true

`"formatting"`: *object, optional*

Formatting settings for CSV log files.

The formatting object has the following fields:

`"quoteChar"`: *single-character string, optional*

The character used to quote CSV entries.

Default: `"`

`"delimiterChar"`: *single-character string, optional*

The character used to delimit CSV entries.

Default: `,`

`"endOfLineSymbols"`: *string, optional*

The character or characters that separate a line.

Default: system-dependent line separator defined for the JVM

`"buffering"`: *object, optional*

Buffering settings for writing CSV log files. The default is for messages to be written to the log file for each event.

The buffering object has the following fields:

`"enabled"`: *boolean, optional*

Whether log buffering is enabled.

Default: false

`"autoFlush"`: *boolean, optional*

Whether events are automatically flushed after being written.

Default: true

`"security"`: *object, optional*

Security settings for CSV log files. These settings govern tamper-evident logging, whereby messages are signed. By default tamper-evident logging is not enabled.

The security object has the following fields:

`"enabled"`: *boolean, optional*

Whether tamper-evident logging is enabled.

Default: false

Tamper-evident logging depends on a specially prepared keystore. For details, see "Preparing a Keystore for Tamper-Evident Logs".

`"filename"`: *string, required*

File system path to the keystore containing the private key for tamper-evident logging.

The keystore must be a keystore of type `JCEKS`. For details, see "Preparing a Keystore for Tamper-Evident Logs".

`"password"`: *string, required*

The password for the keystore for tamper-evident logging.

This password is used for the keystore and for private keys. For details, see "Preparing a Keystore for Tamper-Evident Logs".

`"signatureInterval"`: *duration, required*

The time interval after which to insert a signature in the CSV file. This duration must not be zero, and must not be unlimited.

For information about supported formats for `duration`, see duration.

"fileRetention": *object, optional*

File retention settings for CSV log files.

The file retention object has the following fields:

"maxDiskSpaceToUse": *number, optional*

The maximum disk space in bytes the audit logs can occupy. A setting of 0 or less indicates that the policy is disabled.

Default: 0

"maxNumberOfHistoryFiles": *number, optional*

The maximum number of historical log files to retain. A setting of -1 disables pruning of old history files.

Default: 0

"minFreeSpaceRequired": *number, optional*

The minimum free space in bytes that the system must contain for logs to be written. A setting of 0 or less indicates that the policy is disabled.

Default: 0

"fileRotation": *object, optional*

File rotation settings for log files.

The file rotation object has the following fields:

"rotationEnabled": *configuration expression<boolean>, optional*

Whether file rotation is enabled for log files.

Default: false.

"maxFileSize": *configuration expression<number>, optional*

The maximum file size of an audit log file in bytes. A setting of 0 or less indicates that the policy is disabled.

Default: 0.

"rotationFilePrefix": *configuration expression<string>, optional*

The prefix to add to a log file on rotation.

This has an effect when time-based file rotation is enabled.

"rotationFileSuffix": *configuration expression<string>, optional*

The suffix to add to a log file on rotation, possibly expressed in SimpleDateFormat.

This has an effect when time-based file rotation is enabled.

Default: `-yyyy.MM.dd-HH.mm.ss`, where *yyyy* characters are replaced with the year, *MM* characters are replaced with the month, *dd* characters are replaced with the day, *HH* characters are replaced with the hour (00-23), *mm* characters are replaced with the minute (00-60), and *ss* characters are replaced with the second (00-60).

"rotationInterval": *duration, optional*

The time interval after which to rotate log files. This duration must not be zero.

This has the effect of enabling time-based file rotation.

For information about supported formats for `duration`, see duration.

"rotationTimes": *array of durations, optional*

The durations, counting from midnight, after which to rotate files.

The following example schedules rotation six and twelve hours after midnight:

```
"rotationTimes": [ "6 hours", "12 hours" ]
```

This has the effect of enabling time-based file rotation.

For information about supported formats for `duration`, see duration.

"rotationRetentionCheckInterval": *duration, optional*

The time interval after which to check file rotation and retention policies for updates.

Default: 5 seconds

For information about supported formats for `duration`, see duration.

Preparing a Keystore for Tamper-Evident Logs

Tamper-evident logging depends on a public key/private key pair and on a secret key that are stored together in a JCEKS keystore. Follow these steps to prepare the keystore:

1. Generate a key pair in the keystore.

    The CSV event handler expects a JCEKS-type keystore with a key alias of `Signature` for the signing key, where the key is generated with the `RSA` key algorithm and the `SHA256withRSA` signature algorithm:

```
$ keytool \
 -genkeypair \
 -keyalg RSA \
 -sigalg SHA256withRSA \
 -alias "signature" \
 -dname "CN=openig.example.com,O=Example Corp,C=FR" \
 -keystore /path/to/audit-keystore \
 -storetype JCEKS \
 -storepass password \
 -keypass password
```

> **Note**
>
> Because KeyStore converts all characters in its key aliases to lower case, use only lowercase in alias definitions of a KeyStore.

2. Generate a secret key in the keystore.

   The CSV event handler expects a JCEKS-type keystore with a key alias of `Password` for the symmetric key, where the key is generated with the `HmacSHA256` key algorithm and 256-bit key size:

```
$ keytool \
 -genseckey \
 -keyalg HmacSHA256 \
 -keysize 256 \
 -alias "password" \
 -keystore /path/to/audit-keystore \
 -storetype JCEKS \
 -storepass password \
 -keypass password
```

3. Verify the content of the keystore:

```
$ keytool \
 -list \
 -keystore /path/to/audit-keystore \
 -storetype JCEKS \
 -storepass password

Keystore type: JCEKS
Keystore provider: SunJCE

Your keystore contains 2 entries

signature, Nov 27, 2015, PrivateKeyEntry,
Certificate fingerprint (SHA1): 4D:CF:CC:29:...:8B:6E:68:D1
password, Nov 27, 2015, SecretKeyEntry,
```

Example

For instructions on recording audit events in a CSV file, see "To Record Audit Events in a CSV File" in the *Gateway Guide*.

The following example configures a CSV audit event handler to write a log file, `/path/to/audit/logs/access.csv`, that is signed every 10 seconds to make it tamper-evident:

```
{
  "name": "csv",
  "topics": [
    "access"
  ],
  "logDirectory": "/path/to/audit/logs/",
  "security": {
    "enabled": "true",
    "filename": "/path/to/audit-keystore",
    "password": "password",
    "signatureInterval": "10 seconds"
  }
}
```

More Information

org.forgerock.audit.handlers.csv.CsvAuditEventHandler

Name

ElasticsearchAuditEventHandler — log audit events in the Elasticsearch search and analytics engine

Description

An audit event handler that responds to events by logging messages in the Elasticsearch search and analytics engine.

The configuration is declared in an audit service configuration. For information, see AuditService(5).

For Elasticsearch downloads and installation instructions, see the Elasticsearch *Getting Started* document.

A special client handler called `ElasticsearchClientHandler` can be defined to send audit events to Elasticsearch. You can use this client handler to capture the exchange between the audit service and Elasticsearch, or to wrap the search with a filter, for example, the `OAuth2ClientFilter`.

To define an `ElasticsearchClientHandler`, create the following object in the heap for the Elasticsearch audit event handler

```
{
  "name": "ElasticsearchClientHandler",
  "type": "ClientHandler",
  "config": {},
}
```

Usage

```
{
  "class": "org.forgerock.audit.handlers.elasticsearch.ElasticsearchAuditEventHandler",
  "config": {
    "name": configuration expression<string>,
    "topics" : [ configuration expression<string>, ... ],
    "connection" : {
      "host" : configuration expression<string>,
      "port" : configuration expression<number>,
      "useSSL" : configuration expression<boolean>,
      "username" : configuration expression<string>,
      "password" : configuration expression<string>
    },
    "indexMapping" : {
      "indexName" : configuration expression<string>
    },
    "buffering" : {
      "enabled" : configuration expression<boolean>,
      "writeInterval" : configuration expression<duration>,
      "maxSize" : configuration expression<number>,
      "maxBatchedEvents" : configuration expression<number>
    }
  }
}
```

The values in this configuration object can use expressions if they resolve to the correct types for each field. For information about expressions, see Expressions(5).

Properties

The `"config"` object has the following properties:

`"name"`: *configuration expression<string>, required*

The name of the event handler.

`"topics"`: *array of configuration expressions<string>, required*

The topics that this event handler intercepts.

IG handles access events that occur at the system boundary, such as arrival of the initial request and departure of the final response.

Set this to `"topics": [ "access" ]`.

`"connection"`: *object, optional*

Connection settings for sending messages to Elasticsearch. If this object is not configured, it takes default values for its fields. This object has the following fields:

`"host"`: *configuration expression<string>, optional*

Hostname or IP address of Elasticsearch.

Default: `localhost`

`"port"`: *configuration expression<number>, optional*

The port used by Elasticsearch. The value must be between 0 and 65535.

Default: `9200`

`"useSSL"`: *configuration expression<boolean>, optional*

Setting to use or not use SSL/TLS to connect to Elasticsearch.

Default: `false`

`"username"`: *configuration expression<string>, optional*

Username when basic authentication is enabled through Elasticsearch Shield.

`"password"`: *configuration expression<string>, optional*

Password when basic authentication is enabled through Elasticsearch Shield.

"indexMapping": *object, optional*

Defines how an audit event and its fields are stored and indexed.

"indexName": *configuration expression<string>, optional*

The index name. Set this parameter if the default name `audit` conflicts with an existing Elasticsearch index.

Default: `audit`.

"buffering": *object, optional*

Settings for buffering events and batch writes.

"enabled": *configuration expression<boolean>, optional*

Setting to use or not use log buffering.

Default: false.

"writeInterval": *configuration expression<duration>*

The interval at which to send buffered event messages to Elasticsearch. If buffering is enabled, this interval must be greater than 0.

Default: 1 second

For information about supported formats for `duration`, see duration.

"maxBatchedEvents": *configuration expression<number>, optional*

The maximum number of event messages in a batch write to Elasticsearch for each `writeInterval`.

Default: 500

"maxSize": *configuration expression<number>, optional*

The maximum number of event messages in the queue of buffered event messages.

Default: 10000

Example

In the following example, an Elasticsearch audit event handler logs audit events for access. For an example of setting up and testing this configuration, see "To Record Audit Events in Elasticsearch" in the *Gateway Guide*.

```
{
```

```
  "name": "30-elasticsearch",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/audit')}",
  "auditService": {
    "name": "AuditService",
    "type": "AuditService",
    "config": {
      "event-handlers": [
        {
          "class": "org.forgerock.audit.handlers.elasticsearch.ElasticsearchAuditEventHandler",
          "config": {
            "name": "elasticsearch",
            "indexMapping": {
              "indexName": "audit"
            },
            "connection": {
              "host": "localhost",
              "port": 9200,
              "useSSL": false
            },
            "topics": [
              "access"
            ]
          }
        }
      ]
    }
  },
  "handler": "ReverseProxyHandler"
}
```

More Information

org.forgerock.audit.handlers.elasticsearch.ElasticsearchAuditEventHandler

Name

JdbcAuditEventHandler — log audit events to relational database

Description

An audit event handler that responds to events by logging messages to an appropriately configured relational database table.

The configuration is declared in an audit service configuration. For details, see AuditService(5).

Usage

```
{
    "class": "org.forgerock.audit.handlers.jdbc.JdbcAuditEventHandler",
    "config": {
        "name": string,
        "topics": array,
        "databaseType": string,
        "enabled": boolean,
        "buffering": {
            "enabled": boolean,
            "writeInterval": duration,
            "autoFlush": boolean,
            "maxBatchedEvents": number,
            "maxSize": number,
            "writerThreads": number
        },
        "connectionPool": {
            "driverClassName": string,
            "dataSourceClassName": string,
            "jdbcUrl": string,
            "username": string,
            "password": string,
            "autoCommit": boolean,
            "connectionTimeout": number,
            "idleTimeout": number,
            "maxLifetime": number,
            "minIdle": number,
            "maxPoolSize": number,
            "poolName": string
        },
        "tableMappings": [
            {
                "event": string,
                "table": string,
                "fieldToColumn": {
                    "event-field": "database-column"
                }
            }
        ]
    }
}
```

The values in this configuration object can use expressions as long as they resolve to the correct types for each field. For details about expressions, see Expressions(5).

Configuration

The `"config"` object has the following properties:

`"name"`: *string, required*

> The name of the event handler.

`"topics"`: *array of strings, required*

> The topics that this event handler intercepts.
>
> IG handles access events that occur at the system boundary, such as arrival of the initial request and departure of the final response.
>
> Set this to `"topics": [ "access" ]`.

`"databaseType"`: *string, required*

> The database type name.
>
> Built-in support is provided for `oracle`, `mysql`, and `h2`. Unrecognized database types rely on a GenericDatabaseStatementProvider.

`"enabled"`: *boolean, optional*

> Whether this event handler is active.
>
> Default: true.

`"buffering"`: *object, optional*

> Buffering settings for sending messages to the database. The default is for messages to be written to the log file for each event.
>
> The buffering object has the following fields:
>
> > **`"enabled"`:** *boolean, optional*
> >
> > > Whether log buffering is enabled.
> > >
> > > Default: false.
> >
> > **`"writeInterval"`:** *duration, required*
> >
> > > The interval at which to send buffered event messages to the database.

This interval must be greater than 0 if buffering is enabled.

For information about supported formats for `duration`, see duration.

"autoFlush": *boolean, optional*

Whether the events are automatically flushed after being written.

Default: true.

"maxBatchedEvents": *number, optional*

The maximum number of event messages batched into a PreparedStatement.

Default: 100.

"maxSize": *number, optional*

The maximum size of the queue of buffered event messages.

Default: 5000.

"writerThreads": *number, optional*

The number of threads to write buffered event messages to the database.

Default: 1.

"connectionPool": *object, required*

Connection pool settings for sending messages to the database.

The connection pool object has the following fields:

"driverClassName": *string, optional*

The class name of the driver to use for the JDBC connection. For example, with MySQL Connector/J, the class name is `com.mysql.jdbc.Driver`.

"dataSourceClassName": *string, optional*

The class name of the data source for the database.

"jdbcUrl": *string, required*

The JDBC URL to connect to the database.

"username": *string, required*

The username identifier for the database user with access to write the messages.

`"password":` *number, optional*

The password for the database user with access to write the messages.

`"autoCommit":` *boolean, optional*

Whether to commit transactions automatically when writing messages.

Default: true.

`"connectionTimeout":` *number, optional*

The number of milliseconds to wait for a connection from the pool before timing out.

Default: 30000.

`"idleTimeout":` *number, optional*

The number of milliseconds to allow a database connection to remain idle before timing out.

Default: 600000.

`"maxLifetime":` *number, optional*

The number of milliseconds to allow a database connection to remain in the pool.

Default: 1800000.

`"minIdle":` *number, optional*

The minimum number of idle connections in the pool.

Default: 10.

`"maxPoolSize":` *number, optional*

The maximum number of connections in the pool.

Default: 10.

`"poolName":` *string, optional*

The name of the connection pool.

`"tableMappings":` *array of objects, required*

Table mappings for directing event content to database table columns.

A table mappings object has the following fields:

`"event":` *string, required*

The audit event that the table mapping is for.

Set this to `access`.

`"table"`: *string, required*

The name of the database table that corresponds to the mapping.

`"fieldToColumn"`: *object, required*

This object maps the names of audit event fields to database columns, where the keys and values are both strings.

Audit event fields use JSON pointer notation, and are taken from the JSON schema for the audit event content.

Example

The following example configures a JDBC audit event handler using a local MySQL database, writing to a table named `auditaccess`:

```
{
    "class": "org.forgerock.audit.handlers.jdbc.JdbcAuditEventHandler",
    "config": {
        "databaseType": "mysql",
        "name": "jdbc",
        "topics": [
            "access"
        ],
        "connectionPool": {
            "jdbcUrl": "jdbc:mysql://localhost:3306/audit?allowMultiQueries=true&characterEncoding=utf8",
            "username": "audit",
            "password": "audit"
        },
        "tableMappings": [
            {
                "event": "access",
                "table": "auditaccess",
                "fieldToColumn": {
                    "_id": "id",
                    "timestamp": "timestamp_",
                    "eventName": "eventname",
                    "transactionId": "transactionid",
                    "userId": "userid",
                    "trackingIds": "trackingids",
                    "server/ip": "server_ip",
                    "server/port": "server_port",
                    "client/host": "client_host",
                    "client/ip": "client_ip",
                    "client/port": "client_port",
                    "request/protocol": "request_protocol",
                    "request/operation": "request_operation",
                    "request/detail": "request_detail",
                    "http/request/secure": "http_request_secure",
                    "http/request/method": "http_request_method",
                    "http/request/path": "http_request_path",
```

```
                "http/request/queryParameters": "http_request_queryparameters",
                "http/request/headers": "http_request_headers",
                "http/request/cookies": "http_request_cookies",
                "http/response/headers": "http_response_headers",
                "response/status": "response_status",
                "response/statusCode": "response_statuscode",
                "response/elapsedTime": "response_elapsedtime",
                "response/elapsedTimeUnits": "response_elapsedtimeunits"
            }
        }
    ]
  }
}
```

Examples including statements to create tables are provided in the JDBC handler library, `forgerock-audit-handler-jdbc-version.jar`, that is built into the IG .war file. Unpack the library, then find the examples under the `db/` folder.

More Information

org.forgerock.audit.handlers.jdbc.JdbcAuditEventHandler

Name

JmsAuditEventHandler — record messages between a JMS message broker and clients

Description

The Java Message Service (JMS) is a Java API for sending asynchronous messages between clients. It wraps audit events in JMS messages and publishes them in a JMS broker, which then delivers the messages to the appropriate destination.

The JMS API architecture includes a *JMS provider* and *JMS clients*, and supports the *publish/subscribe* messaging pattern. For more information, see *Basic JMS API Concepts*.

The JMS audit event handler does not support queries. To support queries, also enable a second handler that supports queries.

The ForgeRock JMS audit event handler supports JMS communication, based on the following components:

- JMS message broker, to provide clients with connectivity, message storage, and message delivery functionality.

- JMS messages, which follow the format in "Example".

- Destinations, maintained by a message broker. A destination can be a JMS topic, using publish/subscribe to take the ForgeRock JSON for an audit event, wrap it into a JMS TextMessage, and send it to the broker.

- JMS clients, to produce and/or receive JMS messages.

Depending on the configuration, some or all of these components are included in JMS audit log messages.

> **Important**
>
> The example in this section is based on *Apache ActiveMQ*, but you can choose a different JMS message broker.
>
> Make sure that the `.jar` files required by the JMS message broker are available in the IG web container.

Declare the configuration in an audit service, as described in AuditService(5). To define an `JmsAuditEventHandler`, create the following object in the heap:

```
{
  "name": string,
  "type": "AuditService",
  "config": {},
}
```

Usage

```
{
  "type": "AuditService",
  "config": {
    "config": {},
    "event-handlers": [
    {
      "class": "org.forgerock.audit.handlers.jms.JmsAuditEventHandler",
      "config": {
        "name": string,
        "topics": array of strings,
        "deliveryMode": string,
        "sessionMode": string,
        "jndi": {
          "contextProperties": map,
          "topicName": string,
          "connectionFactoryName": string
        }
      }
    }]
  }
}
```

The values in this configuration object can use configuration expressions, as described in "Configuration and Runtime Expressions".

Configuration

For a full list of properties in the `"config"` object, see Configuration Properties for the JMS Audit Event Handler in the *IDM Integrator's Guide*.

`"name"`: *string, required*

The name of the event handler.

`"topics"`: *array of strings, required*

The topics that this event handler intercepts.

IG handles access events that occur at the system boundary, such as arrival of the initial request and departure of the final response.

Set this to `"topics": [ "access" ]`.

`"deliveryMode"`: *string, required*

Delivery mode for messages from a JMS provider. Set to `PERSISTENT` or `NON_PERSISTENT`.

`"sessionMode"`: *string, required*

Acknowledgement mode in sessions without transactions. Set to `AUTO`, `CLIENT`, or `DUPS_OK`.

"contextProperties": *map, optional*

Settings with which to populate the initial context.

The following properties are required when ActiveMQ is used as the message broker:

- `java.naming.factory.initial`

  For example, `"org.apache.activemq.jndi.ActiveMQInitialContextFactory"`.

  To substitute a different JNDI message broker, change the JNDI context properties.

- `java.naming.provider.url`

  For example, `"tcp://127.0.0.1:61616"`.

  To configure the message broker on a remote system, substitute the associated IP address.

  To set up SSL, set up keystores and truststores, and change the value of the `java.naming.provider.url` to:

  ```
  ssl://127.0.0.1:61617?daemon=true&socket.enabledCipherSuites=SSL_RSA_WITH_RC4_128_SHA
  ,SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
  ```

- `topic.audit`

  For example, `"audit"`.

To use the JMS resources provided by your application server, leave this field empty. The values for `topicName` and `connectionFactoryName` are then JNDI names that depend on the configuration of your application server.

"topicName": *string, required*

JNDI lookup name for the JMS topic.

For ActiveMQ, this property must be consistent with the value of `topic.audit` in `contextProperties`.

"connectionFactoryName": *string, required*

JNDI lookup name for the JMS connection factory.

Example

In the following example, a JMS audit event handler delivers audit events in batches. The handler is configured to use the ActiveMQ JNDI message broker, on port 61616. For an example of setting up and testing this configuration, see "To Record Audit Events With a JMS Audit Event Handler" in the *Gateway Guide*.

```
{
  "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"
    }
  },
  "condition" : "${request.uri.path == '/activemq_event_handler'}"
}
```

The following code shows an example message produced by the JMS audit event handler.

```
{
  "auditTopic": "access",
  "event": {
    "eventName": "OPENIG-HTTP-ACCESS",
    "timestamp": "2016-11-28T14:39:30.004Z",
    "transactionId": "882918f9-f7c3-47ee-9f87-5e3cfcfb98be-37",
    "server": {
      "ip": "0:0:0:0:0:0:0:1",
      "port": 8080
    },
    "client": {
      "ip": "0:0:0:0:0:0:0:1",
      "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"
  }
}
```

More Information

org.forgerock.audit.handlers.jms.JmsAuditEventHandler

Name

JsonAuditEventHandler — log events as JSON objects to a set of JSON files

Description

The JSON audit event handler logs events as JSON objects to a set of JSON files. This is the preferred file-based audit event handler.

Declare the configuration in an audit service, as described in AuditService(5). To define an `JsonAuditEventHandler`, create the following object in the heap:

```
{
  "name": runtime expression<string>,
  "type": "AuditService",
  "config": {},
}
```

Usage

```
{
  "type": "AuditService",
  "config": {
    "config": {},
    "event-handlers": [
    {
      "class": "org.forgerock.audit.handlers.json.JsonAuditEventHandler",
      "config": {
        "name": runtime expression<string>,
        "topics": array of runtime expressions<string>,
        "logDirectory" : expression,
        "elasticsearchCompatible" : runtime expression<boolean>,
        "fileRotation": {
            "rotationEnabled": runtime expression<boolean>,
            "maxFileSize": runtime expression<number>,
            "rotationFilePrefix": runtime expression<string>,
            "rotationFileSuffix": runtime expression<string>,
            "rotationInterval": runtime expression<duration>,
            "rotationTimes": array of runtime expressions<duration>
        },
        "fileRetention": {
            "maxNumberOfHistoryFiles": runtime expression<number>,
            "maxDiskSpaceToUse": runtime expression<number>,
            "minFreeSpaceRequired": runtime expression<number>,
            "rotationRetentionCheckInterval": runtime expression<duration>
        },
        "buffering" : {
          "writeInterval" : runtime expression<duration>,
          "maxSize" : runtime expression<number>
        }
      }
    }]
  }
}
```

Configuration

`"name"`: *runtime expression<string>, required*

The name of the event handler.

`"topics"`: *array of runtime expressions<string>, required*

The topics that this event handler intercepts.

IG handles access events that occur at the system boundary, such as arrival of the initial request and departure of the final response.

Set this to `"topics": [ "access" ]`.

`"logDirectory"`: *runtime expression<string>, required*

The file system directory where log files are written.

`elasticsearchCompatible`: *boolean true or false, optional*

Set to `true` to enable compatibility with ElasticSearch JSON format. For more information, see the ElasticSearch documentation.

Default: `false`

`"fileRotation"`: *object, optional*

File rotation settings for log files.

The file rotation object has the following fields:

`"rotationEnabled"`: *configuration expression<boolean>, optional*

Whether file rotation is enabled for log files.

Default: false.

`"maxFileSize"`: *configuration expression<number>, optional*

The maximum file size of an audit log file in bytes. A setting of 0 or less indicates that the policy is disabled.

Default: 0.

`"rotationFilePrefix"`: *configuration expression<string>, optional*

The prefix to add to a log file on rotation.

This has an effect when time-based file rotation is enabled.

"rotationFileSuffix": *configuration expression<string>, optional*

The suffix to add to a log file on rotation, possibly expressed in SimpleDateFormat.

This has an effect when time-based file rotation is enabled.

Default: `-yyyy.MM.dd-HH.mm.ss`, where *yyyy* characters are replaced with the year, *MM* characters are replaced with the month, *dd* characters are replaced with the day, *HH* characters are replaced with the hour (00-23), *mm* characters are replaced with the minute (00-60), and *ss* characters are replaced with the second (00-60).

"rotationInterval": *duration, optional*

The time interval after which to rotate log files. This duration must not be zero.

This has the effect of enabling time-based file rotation.

For information about supported formats for `duration`, see duration.

"rotationTimes": *array of durations, optional*

The durations, counting from midnight, after which to rotate files.

The following example schedules rotation six and twelve hours after midnight:

```
"rotationTimes": [ "6 hours", "12 hours" ]
```

This has the effect of enabling time-based file rotation.

For information about supported formats for `duration`, see duration.

"fileRetention": *object, optional*

File retention settings for log files.

The file retention object has the following fields:

"maxNumberOfHistoryFiles": *configuration expression<number>, optional*

The maximum number of historical audit files that can be stored. If the number exceeds this maximum, older files are deleted. A value of `-1` disables purging of old log files.

Default: 0.

"maxDiskSpaceToUse": *runtime expression<number>, optional*

The maximum disk space in bytes that can be used for audit files. If the audit files use more than this space, older files are deleted. A negative or zero value indicates that this policy is disabled, and historical audit files can use unlimited disk space.

`"minFreeSpaceRequired"`: *runtime expression<string>, optional*

> The minimum free disk space in bytes required on the system that houses the audit files. If the free space drops below this minimum, older files are deleted. A negative or zero value indicates that this policy is disabled, and no minimum space requirements apply.

`"rotationRetentionCheckInterval"`: *runtime expression<string>, optional*

Interval at which to periodically check file rotation and retention policies. The interval must be a duration, for example, 5 seconds, 5 minutes, or 5 hours.

`"buffering"`: *object, optional*

Settings for buffering events and batch writes.

> `"writeInterval"`: *duration*
>
> > The interval at which to send buffered event messages. If buffering is enabled, this interval must be greater than 0.
> >
> > Default: 1 second
> >
> > For information about supported formats for `duration`, see duration.
>
> `"maxSize"`: *number, optional*
>
> > The maximum number of event messages in the queue of buffered event messages.
> >
> > Default: 10000

Example

In the following example, a JSON audit event handler logs audit events for access. For an example of setting up and testing this configuration, see "Recording Audit Events in JSON" in the *Gateway Guide*.

```
{
  "name": "30-json",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/audit')}",
  "auditService": {
    "name": "AuditService",
    "type": "AuditService",
    "config": {
      "event-handlers": [
        {
          "class": "org.forgerock.audit.handlers.json.JsonAuditEventHandler",
          "config": {
            "name": "json",
            "logDirectory": "/tmp/logs",
            "elasticsearchCompatible": false,
```

```
            "topics": [
              "access"
            ],
            "fileRetention": {
              "rotationRetentionCheckInterval": "1 minute"
            },
            "buffering": {
              "maxSize": 100000,
              "writeInterval": "100 ms"
            }
          }
        }
      ]
    }
  },
  "handler": "ReverseProxyHandler"
}
```

More Information

org.forgerock.audit.handlers.json.JsonAuditEventHandler

Name

JsonStdoutAuditEventHandler — log events to JSON standard output

Description

Logs events to JSON standard output (stdout).

Declare the configuration in an audit service, as described in AuditService(5). To define a
`JsonStdoutAuditEventHandler`, create the following object in the heap:

```
{
  "name": runtime expression<string>,
  "type": "AuditService",
  "config": {},
}
```

Usage

```
{
  "type": "AuditService",
  "config": {
    "config": {},
    "event-handlers": [
    {
      "class": "org.forgerock.audit.handlers.json.stdout.JsonStdoutAuditEventHandler",
      "config": {
        "name": runtime expression<string>,
        "topics": array of runtime expressions<string>,
        "elasticsearchCompatible" : runtime expression<boolean>,
      }
    }
  }
}
}
```

Configuration

"name": *runtime expression<string>, required*

> The name of the event handler.

"topics": *array of runtime expressions<string>, required*

> The topics that this event handler intercepts.
>
> IG handles access events that occur at the system boundary, such as arrival of the initial request
> and departure of the final response.
>
> Set this to `"topics": [ "access" ]`.

`elasticsearchCompatible`: ***boolean true or false, optional***

Set to `true` to enable compatibility with ElasticSearch JSON format. For more information, see the ElasticSearch documentation.

Default: `false`

Example

In the following example, a JsonStdoutAuditEventHandler logs audit events. For an example of setting up and testing this configuration, see "Recording Audit Events to Standard Output" in the *Gateway Guide*.

```
{
  "name": "30-jsonstdout",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/jsonstdout-audit')}",
  "auditService": {
    "name": "AuditService-1",
    "type": "AuditService",
    "config": {
      "event-handlers": [
        {
          "class": "org.forgerock.audit.handlers.json.stdout.JsonStdoutAuditEventHandler",
          "config": {
            "name": "jsonstdout",
            "elasticsearchCompatible": false,
            "topics": ["access"]
          }
        }
      ]
    }
  },
  "handler": "ReverseProxyHandler"
}
```

More Information

org.forgerock.audit.handlers.json.stdout.JsonStdoutAuditEventHandler

Name
SyslogAuditEventHandler — log audit events to the system log

Description

An audit event handler that responds to events by logging messages to the UNIX system log as governed by RFC 5424, *The Syslog Protocol*.

The configuration is declared in an audit service configuration. For details, see AuditService(5).

Usage

```
{
    "class": "org.forgerock.audit.handlers.syslog.SyslogAuditEventHandler",
    "config": {
        "name": string,
        "topics": array,
        "protocol": string,
        "host": string,
        "port": number,
        "connectTimeout": number,
        "facility": "string",
        "buffering": {
            "enabled": boolean,
            "maxSize": number
        },
        "severityFieldMappings": [
            {
                "topic": string,
                "field": string,
                "valueMappings": {
                    "field-value": "syslog-severity"
                }
            }
        ]
    }
}
```

The values in this configuration object can use expressions as long as they resolve to the correct types for each field. For details about expressions, see Expressions(5).

Configuration

The `"config"` object has the following properties:

`"name"`: *string, required*

The name of the event handler.

`"topics"`: *array of strings, required*

The topics that this event handler intercepts.

IG handles access events that occur at the system boundary, such as arrival of the initial request and departure of the final response.

Set this to `"topics": [ "access" ]`.

`"protocol"`: *string, required*

The transport protocol used to send event messages to the Syslog daemon.

Set this to `TCP` for Transmission Control Protocol, or to `UDP` for User Datagram Protocol.

`"host"`: *string, required*

The hostname of the Syslog daemon to which to send event messages. The hostname must resolve to an IP address.

`"port"`: *number, required*

The port of the Syslog daemon to which to send event messages.

The value must be between 0 and 65535.

`"connectTimeout"`: *number, required when using TCP*

The number of milliseconds to wait for a connection before timing out.

`"facility"`: *string, required*

The Syslog facility to use for event messages.

Set this to one of the following values:

`kern`

Kernel messages

`user`

User-level messages

`mail`

Mail system

`daemon`

System daemons

auth

Security/authorization messages

syslog

Messages generated internally by `syslogd`

lpr

Line printer subsystem

news

Network news subsystem

uucp

UUCP subsystem

cron

Clock daemon

authpriv

Security/authorization messages

ftp

FTP daemon

ntp

NTP subsystem

logaudit

Log audit

logalert

Log alert

clockd

Clock daemon

local0

Local use 0

local1

> Local use 1

local2

> Local use 2

local3

> Local use 3

local4

> Local use 4

local5

> Local use 5

local6

> Local use 6

local7

> Local use 7

"buffering": *object, optional*

Buffering settings for writing to the system log facility. The default is for messages to be written to the log for each event.

The buffering object has the following fields:

"enabled": *boolean, optional*

> Whether log buffering is enabled.

> Default: false.

"maxSize": *number, optional*

> The maximum number of buffered event messages.

> Default: 5000.

"severityFieldMappings": *object, optional*

Severity field mappings set the correspondence between audit event fields and Syslog severity values.

The severity field mappings object has the following fields:

`"topic"`: *string, required*

The audit event topic to which the mapping applies.

Set this to `access`.

`"field"`: *string, required*

The audit event field to which the mapping applies.

Audit event fields use JSON pointer notation, and are taken from the JSON schema for the audit event content.

`"valueMappings"`: *object, required*

The map of audit event values to Syslog severities, where both the keys and the values are strings.

Syslog severities are one of the following values:

`emergency`

System is unusable.

`alert`

Action must be taken immediately.

`critical`

Critical conditions.

`error`

Error conditions.

`warning`

Warning conditions.

`notice`

Normal but significant condition.

`informational`

Informational messages.

`debug`

Debug-level messages.

Example

The following example configures a Syslog audit event handler that writes to the system log daemon on `syslogd.example.com`, port `6514` over TCP with a timeout of 30 seconds. The facility is the first one for local use, and response status is mapped to Syslog informational messages:

```
{
    "class": "org.forgerock.audit.handlers.syslog.SyslogAuditEventHandler",
    "config": {
        "protocol": "TCP",
        "host": "https://syslogd.example.com",
        "port": 6514,
        "connectTimeout": 30000,
        "facility": "local0",
        "severityFieldMappings": [
            {
                "topic": "access",
                "field": "response/status",
                "valueMappings": {
                    "FAILED": "INFORMATIONAL",
                    "SUCCESSFUL": "INFORMATIONAL"
                }
            }
        ]
    }
}
```

More Information

org.forgerock.audit.handlers.syslog.SyslogAuditEventHandler

Name

SplunkAuditEvenHandler — log events to a Splunk system

Description

The Splunk audit event handler logs IG events to a Splunk system. For an example of setting up and testing Splunk, see "Recording Audit Events in Splunk" in the *Gateway Guide*.

Declare the configuration in an audit service, as described in AuditService(5). To define an `SplunkAuditEvenHandler`, create the following object in the heap:

```
{
  "name": string,
  "type": "AuditService",
  "config": {},
}
```

Usage

```
{
  "type": "AuditService",
  "config": {
    "config": {},
    "event-handlers": [
    {
      "class": "org.forgerock.audit.handlers.splunk.SplunkAuditEventHandler",
      "config": {
        "name" : string,
        "topics" : array of strings,
        "enabled" : boolean
        "connection" : {
            "useSSL" : boolean,
            "host" : string,
            "port" : number
        },
        "buffering" : {
            "maxSize" : number,
            "writeInterval" : duration,
            "maxBatchedEvents" : number
        },
        "authzToken" : string,
    }
}
```

The values in this configuration object can use configuration expressions, as described in "Configuration and Runtime Expressions".

Configuration

"name": *string, required*

The name of the event handler.

"topics": *array of strings, required*

The topics that this event handler intercepts.

IG handles access events that occur at the system boundary, such as arrival of the initial request and departure of the final response.

Set this to `"topics": [ "access" ]`.

"enabled": *boolean expression, required*

Specifies whether this audit event handler is enabled.

"connection": *object, optional*

Connection settings for sending messages to the Splunk system. If this object is not configured, it takes default values for its fields. This object has the following fields:

"useSSL": *boolean, optional*

Specifies whether IG should connect to the audit event handler instance over SSL.

Default: `false`

"host": *string, optional*

Hostname or IP address of the Splunk system.

Default: `localhost`

"port": *number, optional*

The dedicated Splunk port for HTTP input.

Before you install Splunk, make sure that this port is free. Otherwise, change the port number in Splunk and in the IG routes that use Splunk.

Default: `8088`

"buffering": *object, optional*

Settings for buffering events and batch writes. If this object is not configured, it takes default values for its fields. This object has the following fields:

"maxSize": *number, optional*

The maximum number of event messages in the queue of buffered event messages.

Default: 10000

`"maxBatchedEvents"`: *number, optional*

The maximum number of event messages in a batch write to this event handler for each `writeInterval`.

Default: 500

`"writeInterval"`: *duration*

The delay after which the writer thread is scheduled to run after encountering an empty event buffer.

Default: 100 ms (units of 'ms' or 's' are recommended)

For information about supported formats for `duration`, see duration.

`"authzToken"`: *string, required*

The authorization token associated with the configured HTTP event collector.

Example

In the following example, IG events are logged to a Splunk system.

```
{
  "name": "30-splunk",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/audit')}",
  "auditService": {
    "name": "AuditService",
    "type": "AuditService",
    "config": {
      "event-handlers": [
        {
          "class": "org.forgerock.audit.handlers.splunk.SplunkAuditEventHandler",
          "config": {
            "name": "splunk",
            "enabled": true,
            "authzToken": "<splunk-authorization-token>",
            "connection": {
              "host": "localhost",
              "port": 8088,
              "useSSL": false
            },
            "topics": [
              "access"
            ],
            "buffering": {
              "maxSize": 10000,
              "maxBatchedEvents": 500,
              "writeInterval": "100 ms"
```

```
                }
              }
            }
          ]
        }
    },
    "handler": "ReverseProxyHandler"
}
```

For an example of setting up and testing this configuration, see "Recording Audit Events in Splunk" in the *Gateway Guide*.

More Information

org.forgerock.audit.handlers.splunk.SplunkAuditEventHandler

Monitoring

This section describes monitoring metrics and endpoints exposed by IG. For examples of how to set up monitoring, see "*Monitoring*" in the *Gateway Guide*.

Table of Contents

Name

Monitoring Data Types — data types used in monitoring

This section describes the data types used in monitoring:

Counter

Cumulative metric for a numerical value that only increases.

Gauge

Metric for a numerical value that can increase or decrease.

Summary

Metric that samples observations, providing a count of observations, sum total of observed amounts, average rate of events, and moving average rates across a sliding time window.

The Prometheus view does not provide time-based statistics, as rates can be calculated from the time-series data. Instead, the Prometheus view includes summary metrics whose names have the following suffixes or labels:

- `_count`: number of events recorded

- `_total`: sum of the amounts of events recorded

- `{quantile="0.5"}`: 50% at or below this value

- `{quantile="0.75"}`: 75% at or below this value

- `{quantile="0.95"}`: 95% at or below this value

- `{quantile="0.98"}`: 98% at or below this value

- `{quantile="0.99"}`: 99% at or below this value

- `{quantile="0.999"}`: 99.9% at or below this value

Timer

Metric combining time-series summary statistics.

Common REST views show summaries as JSON objects. JSON summaries have the following fields:

```
{
 "count": number,             // events recorded for this metric
 "max": number,              // maximum duration recorded
 "mean": number,             // total/count, or 0 if count is 0
 "min": number,              // minimum duration recorded for this metric
 "mean_rate": number,        // average rate
 "p50": number,              // 50% at or below this value
 "p75": number,              // 75% at or below this value
 "p95": number,              // 95% at or below this value
 "p98": number,              // 98% at or below this value
 "p99": number,              // 99% at or below this value
 "p999": number,             // 99.9% at or below this value
 "stddev": number,           // standard deviation of recorded durations
 "m15_rate": number,         // fifteen-minute average rate
 "m5_rate": number,          // five-minute average rate
 "m1_rate": number,          // one-minute average rate
 "duration_units": string,   // time unit used in durations
 "rate_units": string,       // event count unit and time unit used in rate
 "total": number             // sum of the durations of events recorded
}
```

Name

Monitoring Endpoints — endpoints to expose metrics

IG collects metrics on requests and responses, and exposes them over the following HTTP endpoints:

• Prometheus Scrape Endpoint

• Common REST Monitoring Endpoint

By default, no special credentials or privileges are required to access monitoring endpoints. For information about how to restrict access, see "Protecting the Monitoring Endpoints" in the *Gateway Guide*.

Metrics are available for each router, subrouter, and route in the configuration. When a TimerDecorator is configured, timer metrics are also available.

Prometheus Scrape Endpoint

All ForgeRock products automatically expose a monitoring endpoint where Prometheus can scrape metrics, in a standard Prometheus format. For information about configuring Prometheus to scrape metrics, see the Prometheus web site.

When IG is set up as described in "*First Steps*" in the *Getting Started Guide*, the endpoint is http://openig.example.com:8080/openig/metrics/prometheus.

For information about available metrics, see:

• "Route Metrics at the Prometheus Scrape Endpoint"

• "Router Metrics at the Prometheus Scrape Endpoint"

• TimerDecorator(5)

For an example that queries the Prometheus Scrape Endpoint, see "To Monitor the Prometheus Scrape Endpoint" in the *Gateway Guide*.

Common REST Monitoring Endpoint

All ForgeRock products expose a monitoring endpoint where metrics are exposed as a JSON format monitoring resource.

When IG is set up as described in "*First Steps*" in the *Getting Started Guide*, the endpoint is http://openig.example.com:8080/openig/metrics/api?_queryFilter=true.

For information about available metrics, see:

• "Route Metrics at the Common REST Monitoring Endpoint"

- "Router Metrics at the Common REST Monitoring Endpoint"

- TimerDecorator(5)

To access metrics for certain routes or handlers, or for any other subset of metrics, run queries on the endpoint. For some examples, see "To Monitor the Common REST Monitoring Endpoint" in the *Gateway Guide*.

Throttling Filters and Policies

To protect applications from being overused by clients, use a throttling filter to limit how many requests clients can make in a defined time.

Table of Contents

Name

ThrottlingFilter — limit the rate of requests

Description

Limits the rate that requests pass through a filter. The maximum number of requests that a client is allowed to make in a defined time is called the *throttling rate*.

When the throttling rate is reached, IG issues an HTTP status code 429 `Too Many Requests` and a `Retry-After` header, whose value is rounded up to 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
```

Usage

```
{
    "name": string,
    "type": "ThrottlingFilter",
    "config": {
        "requestGroupingPolicy": runtime expression<string>,
        "throttlingRatePolicy": reference or inline declaration,   //Use either "throttlingRatePolicy"
        "rate": {                                                   //or "rate", but not both.
            "numberOfRequests": integer,
            "duration": duration string
        },
        "cleaningInterval": duration string,
        "executor": executor
    }
}
```

Properties

"requestGroupingPolicy": *runtime expression<string>, optional*

An expression to identify the partition to use for the request. In many cases the partition identifies an individual client that sends requests, but it can also identify a group that sends requests. The expression can evaluate to the client IP address or user ID, or an OpenID Connect subject/issuer.

Default: Empty string. The value for this expression must not be null.

See also Expressions(5).

"throttlingRatePolicy": *reference or inline declaration, required if* **"rate"** *is not used*

A reference to or inline declaration of a policy to apply for throttling rate. The following policies can be used:

- MappedThrottlingPolicy(5)

- ScriptableThrottlingPolicy(5)

- DefaultRateThrottlingPolicy(5)

This value for this parameter must not be null.

"rate": *rate object, required if* **"throttlingRatePolicy"** *is not used*

The throttling rate to apply to requests. The rate is calculated as the number of requests divided by the duration:

"numberOfRequests": *integer, required*

The number of requests allowed through the filter in the time specified by `"duration"`.

"duration": *duration string, required*

A time interval during which the number of requests passing through the filter is counted.

For information about supported formats for `duration`, see duration.

"cleaningInterval": *duration, optional*

The time to wait before cleaning outdated partitions. The value must be more than zero but not more than one day.

"executor": *executor, optional*

An executor service to schedule the execution of tasks, such as the clean up of partitions that are no longer used.

Default: `ScheduledExecutorService`

See also ScheduledExecutorService(5).

Examples

The following links provide examples of how the throttling policies are implemented:

- "Example of a Mapped Throttling Policy"

- "Example of a Scriptable Throttling Policy"

The following route defines a throttling rate of 6 requests/10 seconds to requests. For information about how to set up and test this example, see " Configuring a Simple Throttling Filter " in the *Gateway Guide*.

```
{
  "name": "00-throttle-simple",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/throttle-simple')}",
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "ThrottlingFilter",
          "name": "ThrottlingFilter-1",
          "config": {
            "requestGroupingPolicy": "",
            "rate": {
              "numberOfRequests": 6,
              "duration": "10 s"
            }
          }
        }
      ],
      "handler": "ReverseProxyHandler"
    }
  }
}
```

More Information

org.forgerock.openig.filter.throttling.ThrottlingFilterHeaplet

Name

MappedThrottlingPolicy — map throttling rates to groups of requests

Description

Maps different throttling rates to different groups of requests, according to the evaluation of `throttlingRateMapper`.

Usage

```
{
    "type": "ThrottlingFilter",
    "config": {
        "requestGroupingPolicy": expression,
        "throttlingRatePolicy": {
            "type": "MappedThrottlingPolicy",
            "config": {
                "throttlingRateMapper": runtime expression<string>,
                "throttlingRatesMapping": {
                    "mapping1": {
                        "numberOfRequests": integer,
                        "duration": duration string
                    },
                    "mapping2": {
                        "numberOfRequests": integer,
                        "duration": duration string
                    }
                },
                "defaultRate": {
                    "numberOfRequests": integer,
                    "duration": duration string
                }
            }
        }
    }
}
```

Properties

`"throttlingRateMapper"`: *runtime expression<string>, required*

An expression to categorize requests for mapping to a throttling rate in the `throttlingRatesMapping`.

If this parameter is null or does not match any specified mappings, the default throttling rate is applied.

`"throttlingRatesMapping"`: *object, required*

A map of throttling rate by request group. Requests are categorized into groups by the evaluation of the expression `"throttlingRateMapper"`.

"*mapping1*" and "*mapping2*": *string, required*

The evaluation of the expression `"throttlingRateMapper"`.

The number of mappings is not limited to two.

"numberOfRequests": *integer, required*

The number of requests allowed through the filter in the time specified by `"duration"`.

"duration": *duration string, required*

A time interval during which the number of requests passing through the filter is counted.

For information about supported formats for `duration`, see duration.

"defaultRate": *object, required*

The default throttling rate to apply if the evaluation of the expression `"throttlingRateMapper"` is null or is not mapped to a throttling rate.

"numberOfRequests": *integer, required*

The number of requests allowed through the filter in the time specified by `"duration"`.

"duration": *duration string, required*

A time interval during which the number of requests passing through the filter is counted.

For information about supported formats for `duration`, see duration.

Example of a Mapped Throttling Policy

In the following example, requests from users with different statuses are mapped to different throttling rates. For information about how to set up and test this example, see "Configuring a Mapped Throttling Filter" in the *Gateway Guide*.

```
{
  "name": "00-throttle-mapped",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/throttle-mapped')}",
  "heap": [
    {
      "name": "AmService-1",
      "type": "AmService",
      "config": {
        "url": "http://openam.example.com:8088/openam",
        "realm": "/",
        "ssoTokenHeader": "iPlanetDirectoryPro",
        "version": "6.5",
        "agent": {
          "username": "ig_agent",
          "password": "password"
        },
```

```
          "sessionCache": {
            "enabled": false
          }
        }
      }
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "name": "OAuth2ResourceServerFilter-1",
          "type": "OAuth2ResourceServerFilter",
          "config": {
            "scopes": [
              "mail",
              "employeenumber"
            ],
            "requireHttps": false,
            "realm": "OpenIG",
            "accessTokenResolver": {
              "name": "token-resolver-1",
              "type": "OpenAmAccessTokenResolver",
              "config": {
                "amService": "AmService-1"
              }
            }
          }
        },
        {
          "name": "ThrottlingFilter-1",
          "type": "ThrottlingFilter",
          "config": {
            "requestGroupingPolicy": "${contexts.oauth2.accessToken.info.mail}",
            "throttlingRatePolicy": {
              "name": "MappedPolicy",
              "type": "MappedThrottlingPolicy",
              "config": {
                "throttlingRateMapper": "${contexts.oauth2.accessToken.info.status}",
                "throttlingRatesMapping": {
                  "gold": {
                    "numberOfRequests": 6,
                    "duration": "10 s"
                  },
                  "silver": {
                    "numberOfRequests": 3,
                    "duration": "10 s"
                  },
                  "bronze": {
                    "numberOfRequests": 1,
                    "duration": "10 s"
                  }
                },
                "defaultRate": {
                  "numberOfRequests": 1,
                  "duration": "10 s"
                }
              }
            }
          }
```

```
            }
          }
        ],
        "handler": "ReverseProxyHandler"
      }
    }
  }
}
```

More Information

org.forgerock.openig.filter.throttling.MappedThrottlingPolicyHeaplet

Name

ScriptableThrottlingPolicy — map throttling rates by using a Groovy script

Description

Uses a script to look up the throttling rates to apply to groups of requests.

The script can store the mapping for the throttling rate in memory, and can use a more complex mapping mechanism than that used in the `MappedThrottlingPolicy`. For example, the script can map the throttling rate for a range of IP addresses. The script can also query an LDAP directory, an external database, or read the mapping from a file.

Scripts must return a Promise<ThrottlingRate, Exception> or a ThrottlingRate.

This section describes the usage of ScriptableThrottlingPolicy, and refers to the following sections of the documentation:

- For information about script properties, available global objects, and automatically imported classes, see Scripts(5).

- For an example of how to create a ScriptableThrottlingPolicy in Studio, see " Configuring a Scriptable Throttling Filter " in the *Gateway Guide*.

Usage

```
{
    "type": "ThrottlingFilter",
    "config": {
        "requestGroupingPolicy": expression,
        "throttlingRatePolicy": {
            "name": string,
            "type": "ScriptableThrottlingPolicy",
            "config": {
                "type": string,
                "file": string,                      // Use either "file"
                "source": string or array of strings,  // or "source", but not both
                "args": object
            }
        }
    }
}
```

Properties

For information about properties for ScriptableThrottlingPolicy, see Scripts(5).

Example of a Scriptable Throttling Policy

In the following example, the `DefaultRateThrottlingPolicy` delegates the management of throttling to the scriptable throttling policy. For information about how to set up and test this example, see " Configuring a Scriptable Throttling Filter " in the *Gateway Guide*.

```
{
  "name": "00-throttle-scriptable",
  "baseURI": "http://app.example.com:8081",
  "condition": "${matches(request.uri.path, '^/home/throttle-scriptable')}",
  "heap": [
    {
      "name": "AmService-1",
      "type": "AmService",
      "config": {
        "url": "http://openam.example.com:8088/openam",
        "realm": "/",
        "ssoTokenHeader": "iPlanetDirectoryPro",
        "version": "6.5",
        "agent": {
          "username": "ig_agent",
          "password": "password"
        },
        "sessionCache": {
          "enabled": false
        }
      }
    }
  ],
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "name": "OAuth2ResourceServerFilter-1",
          "type": "OAuth2ResourceServerFilter",
          "config": {
            "scopes": [
              "mail",
              "employeenumber"
            ],
            "requireHttps": false,
            "realm": "OpenIG",
            "accessTokenResolver": {
              "name": "token-resolver-1",
              "type": "OpenAmAccessTokenResolver",
              "config": {
                "amService": "AmService-1"
              }
            }
          }
        },
        {
          "name": "ThrottlingFilter-1",
          "type": "ThrottlingFilter",
          "config": {
            "requestGroupingPolicy": "${contexts.oauth2.accessToken.info.mail}",
```

```
            "throttlingRatePolicy": {
              "type": "DefaultRateThrottlingPolicy",
              "config": {
                "delegateThrottlingRatePolicy": {
                  "name": "ScriptedPolicy",
                  "type": "ScriptableThrottlingPolicy",
                  "config": {
                    "type": "application/x-groovy",
                    "source": [
                      "if (contexts.oauth2.accessToken.info.status == status) {",
                      "   return new ThrottlingRate(rate, duration)",
                      "} else {",
                      "   return null",
                      "}"
                    ],
                    "args": {
                      "status": "gold",
                      "rate": 6,
                      "duration": "10 seconds"
                    }
                  }
                },
                "defaultRate": {
                  "numberOfRequests": 1,
                  "duration": "10 s"
                }
              }
            }
          }
        }
      }
    ],
    "handler": "ReverseProxyHandler"
  }
 }
}
```

More Information

org.forgerock.openig.filter.throttling.ScriptableThrottlingPolicy.Heaplet

Name

DefaultRateThrottlingPolicy — default policy for throttling rate

Description

Provides a default throttling rate if the delegating throttling policy returns `null`.

Usage

```
{
    "type": "ThrottlingFilter",
    "config": {
        "requestGroupingPolicy": expression,
        "throttlingRatePolicy": {
            "type": "DefaultRateThrottlingPolicy",
            "config": {
                "delegateThrottlingRatePolicy" : reference or inline declaration,
                "defaultRate": {
                    "numberOfRequests": integer,
                    "duration": duration string
                }
            }
        }
    }
}
```

Properties

`"delegateThrottlingRatePolicy"`: *reference, required*

The policy to which the default policy delegates the throttling rate. The `DefaultRateThrottlingPolicy` delegates management of throttling to the policy specified by `delegateThrottlingRatePolicy`.

If `delegateThrottlingRatePolicy` returns `null`, the `defaultRate` is used.

For information about policies to use, see MappedThrottlingPolicy(5) and ScriptableThrottlingPolicy(5).

`"defaultRate"`: *object, required*

The default throttling rate to apply if the delegating policy returns `null`.

`"numberOfRequests"`: *integer, required*

The number of requests allowed through the filter in the time specified by `"duration"`.

`"duration"`: *duration string, required*

A time interval during which the number of requests passing through the filter is counted.

For information about supported formats for `duration`, see duration.

Example

For an example of how this policy is used, see "Example of a Scriptable Throttling Policy" .

More Information

org.forgerock.openig.filter.throttling.DefaultRateThrottlingPolicyHeaplet

Miscellaneous Heap Objects

Table of Contents

Name

AmService — hold configuration information about an instance of AM

Description

An AmService holds information about the configuration of an instance of AM. It is available to IG filters that communicate with that instance.

Usage

```
{
  "name": string,
  "type": "AmService",
  "config": {
    "agent": object,
    "notifications": object,
    "realm": configuration expression<string>,
    "amHandler": handler reference,
    "sessionCache": object,
    "ssoTokenHeader": configuration expression<string>,
    "url": configuration expression<uri string>,
    "version": configuration expression<string>
  }
}
```

WebSocket Notification Service

AmService configures the WebSocket notification service, where IG receives WebSocket notifications from AM for the following events:

- AM session tokens are revoked- IG receives a notification when a user logs out of AM, or when the AM session is modified, closed, or times out.

- AM policies are changed- IG receives a notification when AM creates, deletes, or changes a policy decision.

By default, WebSocket notifications are enabled. To disable them, set `enabled` to `false` in `notifications`.

If the WebSocket connection is lost for a time, notifications that occur during that time are lost. IG waits for the delay specified by `reconnectDelay` in `notifications`, and then attempts to re-establish the connection. If subsequent attempts fail, IG waits and tries again an unlimited number of times. Depending on the `onNotificationDisconnection` configured in AmService and the PolicyEnforcementFilter, IG can clear the session cache and policy cache.

IG can use WebSocket notifications to evict entries from the session cache and policy cache. For an example of setting up cache eviction, see "Using WebSocket Notifications to Evict the Session Cache" in the *Gateway Guide* and "Using WebSocket Notifications to Evict the Policy Cache" in the *Gateway Guide*.

To include WebSocket notifications in your server log files, add the following line to your `logback.xml` as described in " Reference Logback Configuration " in the *Gateway Guide*:

```
<logger name="org.forgerock.openig.tools.notifications" level="TRACE"/>
```

Notifications like this are included in the log files:

```
2018-... | TRACE | Grizzly(1) | o.f.o.t.n.w.WebSocketClientEndpoint | @system | Received a message:
  { "topic": "/agent/policy",...
```

Properties

`"agent"`: *object, required*

The credentials of an AM Java agent that acts on behalf of IG to authenticate with AM, get the user profile, and communicate WebSocket notifications from AM to IG.

`"username"`: *configuration expression<string>, required*

Java agent name.

`"passwordSecretId"`: *configuration expression<string>, required*

The reference to the Java agent password.

The secret ID must conform to the regex pattern: `[a-zA-Z0-9]+(.[a-zA-Z0-9]+)*`. The route's secret service attempts to resolve the secret ID.

`"realm"`: *configuration expression<string>, optional*

The realm of the AM Java agent that acts on behalf of IG. This is the default realm for SingleSignOnFilter, PolicyEnforcementFilter, and UserProfileFilter when they do not specify a realm.

Default: `/` (top level realm).

`"amHandler"`: *handler reference, optional*

The handler to use for communicating with AM. In production, use a ClientHandler that is capable of making an HTTPS connection to AM.

> **Tip**
>
> To facilitate auditing, configure this handler with a `ForgeRockClientHandler`, which sends a ForgeRock Common Audit transaction ID when it communicates with protected applications.
>
> Alternatively, configure this handler as a chain containing a `TransactionIdOutboundFilter`, as in the following configuration:

```
"amHandler" : {
  "type" : "Chain",
  "config" : {
    "handler": "MySecureClientHandler",
    "filters": [ "TransactionIdOutboundFilter" ]
  }
}
```

Default: `ForgeRockClientHandler`

See also Handlers, ClientHandler(5).

`"notifications"`: *object, optional*

Configure WebSocket notification service. For information about WebSocket notifications, see "WebSocket Notification Service".

```
{
  "notifications": {
    "enabled": boolean,
    "reconnectDelay": object,
    "tls": TlsOptions reference
  }
}
```

`enabled`: *boolean, optional*

Enable or disable WebSocket notifications. Set to `false` to disable WebSocket notifications.

Default: `true`

`reconnectDelay`: *duration, optional*

The time between attempts to re-establish a lost WebSocket connection.

When a WebSocket connection is lost, IG waits for this delay and then attempts to re-establish the connection. If subsequent attempts fail, IG waits and tries again an unlimited number of times.

Default: `5 seconds`

For information about supported formats for `duration`, see duration.

`tls`: *TlsOptions reference, optional*

Configure options for WebSocket connections to TLS-protected endpoints, based on a TlsOptions(5) configuration. Define a TlsOptions object inline or in the heap.

Default: Connections for WebSocket connections to TLS-protected endpoints not configured.

"url": *configuration expression<uri string>, required*

The URI of the AM server instance, for example: `https://openam.example.com/openam`.

"sessionCache": *optional*

Enable and configure caching of session information from AM, based on *Caffeine*. For more information, see the GitHub entry, *Caffeine*.

When session information is cached, IG can reuse session token information without repeatedly asking AM to verify the session token. When caching is disabled, IG must ask AM to verify the session token for each request.

If WebSocket notifications are enabled, the cache evicts entries based on session notifications from AM.

Supported with AM 5.5 when the user manually whitelists the `AMCtxId` session property, and with AM 6 and later versions (where the `AMCtxId` session property is whitelisted by default).

In AM, if the realm includes a customized session property whitelist, include `AMCtxId` in the list of properties. The customized session property whitelist overrides the global session property whitelist.

```
{
  "sessionCache": {
    "enabled": configuration expression<boolean>,
    "executor": executor service reference,
    "maximumSize": configuration expression<number>,
    "maximumTimeToCache": configuration expression<duration>,
    "onNotificationDisconnection": configuration expression<enumeration>
  }
}
```

enabled: *configuration expression<boolean>, optional*

Enable caching.

Default: `false`

executor: *executor service reference, optional*

An executor service to schedule the execution of tasks, such as the eviction of entries in the cache.

Default: `ForkJoinPool.commonPool()`

"maximumSize": *configuration expression<number>, optional*

The maximum number of entries the cache can contain.

Default: Unlimited/unbound.

`maximumTimeToCache`: *configuration expression<duration string>, optional*

The maximum duration for which to cache session info.

If `maximumTimeToCache` is longer than `maxSessionExpirationTime`, `maxSessionExpirationTime` is used.

Default: `maxSessionExpirationTime`, from SessionInfoContext.

For information about supported formats for `duration`, see duration.

`onNotificationDisconnection`: *configuration expression<enumeration>, optional*

The strategy to manage the cache when the WebSocket notification service is disconnected, and IG receives no notifications for AM events, such as session revocation. If the cache is not cleared it can become outdated, and IG can allow requests on revoked sessions.

Cached entries that expire naturally while the notification service is disconnected are removed from the cache.

Use one of the following values:

- `NEVER_CLEAR`

  - When the notification service is disconnected:

    - Continue to use the existing cache.

    - Deny access for requests that are not cached, but do not update the cache with these requests.

  - When the notification service is reconnected:

    - Continue to use the existing cache.

    - Query AM for incoming requests that are not found in the cache, and update the cache with these requests.

- `CLEAR_ON_DISCONNECT`

  - When the notification service is disconnected:

    - Clear the cache.

    - Deny access to all requests, but do not update the cache with these requests.

  - When the notification service is reconnected:

    - Query AM for all requests that are not found in the cache. (Because the cache was cleared, the cache is empty after reconnection.)

    - Update the cache with these requests.

- CLEAR_ON_RECONNECT

  - When the notification service is disconnected:

    - Continue to use the existing cache.

    - Deny access for requests that are not cached, but do not update the cache with these requests.

  - When the notification service is reconnected:

    - Query AM for all requests that are not found in the cache. (Because the cache was cleared, the cache is empty after reconnection.)

    - Update the cache with these requests.

  Default: CLEAR_ON_DISCONNECT

"ssoTokenHeader": *configuration expression<string>, optional*

The name of the header or cookie where this AM server expects to find SSO tokens.

Default: iPlanetDirectoryPro

"version": *configuration expression<string>, optional*

The version number of the AM server. Specify 5 or higher.

Default: AM 5.0.

Example

For examples where AmService is used, see the example routes at the end of OAuth2ResourceServerFilter(5), PolicyEnforcementFilter(5), SingleSignOnFilter(5), TokenTransformationFilter(5), and UserProfileFilter(5).

More Information

org.forgerock.openig.tools.am.AmService

Name

ClientRegistration — Hold OAuth 2.0 client registration information

Description

A ClientRegistration holds information about registration with an OAuth 2.0 authorization server or OpenID Provider.

The configuration includes the client credentials that are used to authenticate to the identity provider. The client credentials can be included directly in the configuration, or retrieved in some other way using an expression, described in Expressions(5).

Usage

```
{
  "name": string,
  "type": "ClientRegistration",
  "config": {
    "clientId": expression,
    "clientSecretId": configuration expression<string>,
    "issuer": Issuer reference,
    "registrationHandler": Handler reference,
    "scopes": [ expression, ...],
    "tokenEndpointAuthMethod": enumeration,
    "tokenEndpointAuthSigningAlg": string,
    "keystore": reference,
    "privateKeyJwtAlias": string,
    "privateKeyJwtPassword": string,
    "claims": map or runtime expression<map>,
    "jwtExpirationTimeout": duration
  }
}
```

Properties

The client registration configuration object properties are as follows:

"name": *string, required*

A name for the client registration.

"clientId": *expression, required*

The `client_id` obtained when registering with the authorization server.

See also Expressions(5).

"clientSecretId": *configuration expression<string>, required if `tokenEndpointAuthMethod` is* `client_secret_basic` **or** `client_secret_post`

The reference to the client secret required to authenticate the client to the authorization server.

The secret ID must conform to the regex pattern: `[a-zA-Z0-9]+(.[a-zA-Z0-9]+)*`. The route's secret service attempts to resolve the secret ID.

"issuer": *Issuer reference, required*

The provider configuration to use for this client registration.

Provide either the name of a Issuer object defined in the heap, or an inline Issuer configuration object.

See also Issuer(5).

"registrationHandler": *Handler reference, optional*

Invoke this HTTP client handler to communicate with the authorization server.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

Usually set this to the name of a ClientHandler configured in the heap, or a chain that ends in a ClientHandler.

Default: IG uses the default ClientHandler.

See also Handlers, ClientHandler(5).

"scopes": *array of expressions, optional*

OAuth 2.0 scopes to use with this client registration.

See also Expressions(5).

"tokenEndpointAuthMethod": *enumeration, optional*

The authentication method with which a client authenticates to the authorization server or OpenID provider at the token endpoint. The following client authentication methods are allowed:

- `client_secret_basic`: Clients that have received a `client_secret` value from the authorization server authenticate with the authorization server by using the HTTP Basic authentication scheme, as in the following example:

```
POST /oauth2/token HTTP/1.1
Host: as.example.com
Authorization: Basic ....
Content-Type: application/x-www-form-urlencoded

grant_type=authorization_code&
  code=...
```

- `client_secret_post`: Clients that have received a `client_secret` value from the authorization server authenticate with the authorization server by including the client credentials in the request body, as in the following example:

```
POST /oauth2/token HTTP/1.1
Host: as.example.com
Content-Type: application/x-www-form-urlencoded

grant_type=authorization_code&
  client_id=...&
  client_secret=...&
  code=...
```

- `private_key_jwt`: Clients send a signed JSON Web Token (JWT) to the authorization server. IG builds and signs the JWT, and prepares the request as in the following example:

```
POST /token HTTP/1.1
Host: as.example.com
Content-Type: application/x-www-form-urlencoded

grant_type=authorization_code&
  code=...&
  client_id=<clientregistration_id>&
  client_assertion_type=urn%3Aietf%3Aparams%3Aoauth%3Aclient-assertion-type%3Ajwt-bearer&
  client_assertion=PHNhbWxwOl ... ZT
```

Some providers accept more than one authentication method. Where a provider strictly enforces how the client must authenticate, align the configuration with the provider.

If the configured method is not supported by the provider, then according to RFC 6749 *The OAuth 2.0 Authorization Framework*, section 5.2 the provider sends an HTTP 400 Bad Request response with an `invalid_client` error message as in the following example response:

```
HTTP/1.1 400 Bad Request
Content-Type: application/json;charset=UTF-8
Cache-Control: no-store
Pragma: no-cache

{
  "error":"invalid_client"
}
```

If the configured method is invalid, an `IllegalArgumentException` is thrown.

Default: `client_secret_basic`

For information about client authentication methods, see OpenID Client Authentication.

"tokenEndpointAuthSigningAlg": *string, optional*

The JSON Web Algorithm (JWA) used to sign the JWT that is used to authenticate the client at the token endpoint. The property is used when `private_key_jwt` is used for authentication.

Use one of the following algorithms:

- `RS256`: RSA using SHA-256

- `ES256`: ECDSA with SHA-256 and NIST standard P-256 elliptic curve

- `ES384`: ECDSA with SHA-384 and NIST standard P-384 elliptic curve

- `ES512`: ECDSA with SHA-512 and NIST standard P-521 elliptic curve

Default: `RS256`

"keystore": *reference, required if `private_key_jwt` is used*

The Java KeyStore containing the private key that is used to sign the JWT.

Provide the name of a KeyStore object defined in the heap, or an inline KeyStore configuration object.

For more information, see KeyStore(5).

"privateKeyJwtAlias": *string, required if `private_key_jwt` is used*

Name of the private key contained in the KeyStore.

"privateKeyJwtPassword": *string, required if `private_key_jwt` is used*

Password to access the private key contained in the KeyStore.

"claims": *map or runtime expression<map>, optional*

When `private_key_jwt` is used for authentication, this property specifies the claims used in the authentication. If this property is a map, the structure must have the format `Map<String, Object>`.

The JWT can contain the following claim value and other optional claims, where claims that are not understood are ignored:

"aud": *string or array of strings, optional*

The URI of the authorization server that is the intended audience of the token.

Default: URL of the authorization server token endpoint

In the following example, the claims include the value `aud`, which is the URI of the authorization server that is the audience of the token:

```
"claims": {
  "aud": "https://myapp.authentication.example.com"
}
```

If this property is an expression, its evaluation must yield an object of type `Map<String, Object>`. In the following example, `overrideAudience` is declared in the properties and then included in an expression in the claims declaration:

```
{
  "properties": {
    "overrideAudience": {
      "aud": "https://myapp.authentication.example.com"
    }
  }
}
```

```
"claims": "${overrideAudience}"
```

"jwtExpirationTimeout": *duration, optional*

When `private_key_jwt` is used for authentication, this property specifies the duration for which the JWT is valid.

Default: 1 minute

For information about supported formats for `duration`, see duration.

Example

The following example shows a client registration for AM. In this example client credentials are replaced with `**********`. In the actual configuration either include the credentials and protect the configuration file or obtain the credentials from the environment in a safe way:

```
{
    "name": "registration",
    "type": "ClientRegistration",
    "config": {
        "clientId": "**********",
        "clientSecretId": "**********",
        "issuer": {
            "type": "Issuer",
            "config": {
                "wellKnownEndpoint": "http://openam.example.com:8088/openam/oauth2/.well-known/openid-configuration"
            }
        },
        "scopes": [
          "openid",
          "profile",
          "email"
        ]
    }
}
```

More Information

org.forgerock.openig.filter.oauth2.client.ClientRegistration

Issuer(5), OAuth2ClientFilter(5)

The OAuth 2.0 Authorization Framework

OAuth 2.0 Bearer Token Usage

OpenID Connect

Name

JwtSession — store sessions in encrypted JWT cookies

Description

A JwtSession object holds settings for storing session information in JSON Web Tokens (JWT) that are encrypted and signed, and then placed in a cookie.

A JWT session is built or rebuilt when there is a write operation in the session (for example, when a user successfully logs in with SAML, assertions are stored in the session). The JWT contains a JSON representation of the session information. The supporting cookie can be configured as a session cookie or persistent cookie. By default it is a session cookie.

The JWT session information contains the session attributes (serialized as JSON), and a marker that IG uses to control the session timeout. When IG receives a JWT cookie, it checks that the current date and time are before the session timeout.

When using this storage implementation, you must use data types for session information that can be mapped to JavaScript Object Notation (JSON). JSON allows strings, numbers, `true`, `false`, `null`, as well as arrays and JSON objects composed of the same primitives. Java and Groovy types that can be mapped include Java primitive types and `null`, String and CharSequence objects, as well as List and Map objects.

Browser cookie storage capacity is limited to 4 KB. Because encryption adds overhead, take care to limit the size of any JSON that you store. Rather than storing larger data in the session information, consider storing a reference instead.

When a request enters a route that uses a new session type, the scope of the session information becomes limited to the route. IG builds a new session object and does not propagate any existing session information to the new object. `session` references the new session object. When the response then exits the route, the session object is closed, and serialized to a JWT cookie in this case, and `session` references the previous session object. Session information set inside the route is no longer available.

An HTTP client that performs multiple requests in a session that modify the content of its session can encounter inconsistencies in the session information. This is because IG does not share JwtSessions across threads. Instead, each thread has its own JwtSession objects that it modifies as necessary, writing its own session to the JWT cookie regardless of what other threads do.

> **Important**
>
> The security of your system depends on how well your keys and keystores are secured.

Usage

```
{
  "name": string,
  "type": "JwtSession",
  "config": {
    "encryptionSecretId": configuration expression<secret-id>,
    "cookie": object,
    "sessionTimeout": duration,
    "persistentCookie": boolean,
    "signatureSecretId": configuration expression<secret-id>,
  }
}
```

An alternative value for type is JwtSessionFactory.

Properties

"encryptionSecretId": *configuration expression<secret-id>, optional*

The secret ID of the encryption key used to encrypt the JWT.

For information about supported formats for `secret-id`, see secret-id.

"cookie" *object, optional*

The configuration of the cookie used to store the encrypted JWT.

Default: The cookie is treated as a host-based cookie.

```
{
  "cookie": {
    "name": configuration expression<string>,
    "domain": : configuration expression<string>,
    "httpOnly": configuration expression<boolean>,
    "path": : configuration expression<string>,
    "sameSite": configuration expression<enumeration>,
    "secure": configuration expression<boolean>
  }
}
```

name *configuration expression<string>, optional*

Name of the JWT cookie stored on the user-agent.

Default: `openig-jwt-session`

`domain` *configuration expression<string>, optional*

Name of the domain from which the JWT cookie can be accessed.

When the domain is specified, a JWT cookie can be accessed from different hosts in that domain.

Default: The domain is not specified. The JWT cookie can be accessed only from the host where the cookie was created.

`"httpOnly":` *configuration expression<boolean>, optional*

Flag to mitigate the risk of client-side scripts accessing protected cookies.

Default: `true`

`"path":` *configuration expression<string>, optional*

Path protected by this session.

Default: `/`

`"sameSite":` *configuration expression<enumeration>, optional*

Manage the circumstances in which a cookie is sent to the server. Use one of the following options to reduce the risk of cross-site request forgery (CSRF) attacks:

- `STRICT`: Send the cookie only if the request was initiated from the cookie domain.

- `LAX`: Send the cookie only with GET requests in a first-party context, where the URL in the address bar matches the cookie domain.

The value is not case-sensitive.

Default: Null; send the cookie whenever a request is made to the cookie domain.

`"secure":` *configuration expression<boolean>, optional*

Flag to limit the scope of the cookie to secure channels.

Default: `false`

`"sessionTimeout"` *duration, optional*

The duration for which a JWT session is valid. If the supporting cookie is persistent, this property also defines the expiry of the cookie.

The value must be above zero. The maximum value is 3650 days (approximately 10 years). If you set a longer duration, IG truncates the duration to 3650 days.

Default: 30 minutes

For information about supported formats for `duration`, see duration.

"persistentCookie" *boolean, optional*

Whether or not the supporting cookie is persistent:

- `true`: the supporting cookie is a persistent cookie. Persistent cookies are re-emitted by the user-agent until their expiration date or until they are deleted.

- `false`: the supporting cookie is a session cookie. IG does not specify an expiry date for session cookies. The user-agent is responsible for deleting them when it considers that the session is finished (for example, when the browser is closed).

Default: `false`

"signatureSecretId" *configuration expression<secret-id>, optional*

The secret ID of the JWT signature required to sign and verify the JWTs. HMAC-SHA-256 is used to sign JWTs.

The value of this attribute must be:

- Base64-encoded.

- At least 32 bytes (32 characters)/256 bits long after base64 decoding. If the provided key is too short, an error message is created.

For information about supported formats for `secret-id`, see secret-id.

Default: Random data is generated as the key, and the IG instance can verify only the sessions it has created.

Example

For pointers on how to configure a JwtSession for encryption, and use `encryptionSecretId` for the encryption key, see "To Set Up Keys for JWT Encryption" in the *Gateway Guide*.

For pointers on how to add `signatureSecretId` to manage multiple instances of IG in the same deployment, see "To Set Up Shared Secrets for Multiple Instances of IG" in the *Gateway Guide*.

More Information

org.forgerock.openig.jwt.JwtSessionManager

Name

KeyManager — configure a Java Secure Socket Extension KeyManager

Description

This represents the configuration for a Java Secure Socket Extension KeyManager, which manages the keys used to authenticate an SSLSocket to a peer. The configuration references the keystore that actually holds the keys.

Usage

```
{
    "name": string,
    "type": "KeyManager",
    "config": {
        "keystore": KeyStore reference,
        "passwordSecretId": configuration expression<string>
        "alg": string
    }
}
```

Properties

"keystore": *KeyStore reference, optional*

The keystore that references the store for the actual keys.

Provide either the name of the KeyStore object defined in the heap, or the inline KeyStore configuration object inline.

See also KeyStore(5).

"passwordSecretId": *configuration expression<string>, required*

The reference to the password required to read private keys from the KeyStore.

The secret ID must conform to the regex pattern: `[a-zA-Z0-9]+(.[a-zA-Z0-9]+)*`. The route's secret service attempts to resolve the secret ID.

"alg" *string, optional*

The certificate algorithm to use.

Default: the default for the platform, such as `SunX509`.

See also Expressions(5).

Example

The following example configures a key manager that depends on a KeyStore configuration. The keystore takes a password supplied as a Java system property when starting the container where IG runs, as in `-Dkeypass=password`. This configuration uses the default certificate algorithm:

```
{
  "name": "MyKeyManager",
  "type": "KeyManager",
  "config": {
    "keystore": {
      "type": "KeyStore",
      "config": {
        "url": "file://${env['HOME']}/keystore.jks",
        "passwordSecretId": "keymanager.keystore.secret.id"
      }
    },
    "passwordSecretId": "keymanager.secret.id"
  }
}
```

More Information

org.forgerock.openig.security.KeyManagerHeaplet

JSSE Reference Guide, KeyStore(5), TrustManager(5)

Name

KeyStore — configure a Java KeyStore

Description

This represents the configuration for a Java KeyStore, which stores cryptographic private keys and public key certificates.

Usage

```
{
    "name": name,
    "type": "KeyStore",
    "config": {
        "url": configuration expression<uri string>,
        "passwordSecretId": configuration expression<string>,
        "type": configuration expression<string>
    }
}
```

Properties

`"url"`: *configuration expression<uri string>, required*

URL to the keystore file.

See also Expressions(5).

`"passwordSecretId"`: *configuration expression<string>, optional*

The reference to the password required to read private keys from the KeyStore.

The secret ID must conform to the regex pattern: `[a-zA-Z0-9]+(.[a-zA-Z0-9]+)*`. The route's secret service attempts to resolve the secret ID.

If the KeyStore is used as a truststore to store only public key certificates of peers and no password is required to do so, then you do not have to specify this field.

Default: No password is set.

See also Expressions(5).

`"type"`: *configuration expression<string>, optional*

The KeyStore format.

Default: the default for the platform, such as `JKS`.

Example

The following example configures a KeyStore that references a Java KeyStore file, `$HOME/keystore.jks`. The KeyStore takes a password supplied as a Java system property when starting the container where IG runs, as in `-Dkeypass=password`. As the KeyStore file uses the default format, no type is specified:

```
{
    "name": "MyKeyStore",
    "type": "KeyStore",
    "config": {
        "url": "file://${env['HOME']}/keystore.jks",
        "passwordSecretId": "${system['keypass']}"
    }
}
```

More Information

org.forgerock.openig.security.KeyStoreHeaplet

JSSE Reference Guide, KeyManager(5), TrustManager(5)

Name
Issuer — Describe an Authorization Server or OpenID Provider

Description

An Issuer describes an OAuth 2.0 Authorization Server or an OpenID Provider that IG can use as a OAuth 2.0 client or OpenID Connect relying party.

An Issuer is generally referenced from a ClientRegistration, described in ClientRegistration(5).

Usage

```
{
  "name": string,
  "type": "Issuer",
  "config": {
    "wellKnownEndpoint": URL string,
    "authorizeEndpoint": URI expression,
    "registrationEndpoint": URI expression,
    "tokenEndpoint": URI expression,
    "userInfoEndpoint": URI expression,
    "issuerHandler": Handler reference,
    "issuerRepository": Issuer repository reference,
    "supportedDomains": [ domain pattern, ... ]
  }
}
```

Properties

If the provider has a well-known configuration URL as defined for OpenID Connect 1.0 Discovery that returns JSON with at least authorization and token endpoint URLs, then you can specify that URL in the provider configuration. Otherwise, you must specify at least the provider authorization and token endpoint URLs, and optionally the registration endpoint and user info endpoint URLs.

The provider configuration object properties are as follows:

`"name"`: *string, required*

   A name for the provider configuration.

`"wellKnownEndpoint"`: *URL string, required unless authorizeEndpoint and tokenEndpoint are specified*

   The URL to the well-known configuration resource as described in OpenID Connect 1.0 Discovery.

`"authorizeEndpoint"`: *expression, required unless obtained through wellKnownEndpoint*

   The URL to the provider's OAuth 2.0 authorization endpoint.

See also Expressions(5).

"registrationEndpoint": *expression, optional*

The URL to the provider's OpenID Connect dynamic registration endpoint.

See also Expressions(5).

"tokenEndpoint": *expression, required unless obtained through wellKnownEndpoint*

The URL to the provider's OAuth 2.0 token endpoint.

See also Expressions(5).

"userInfoEndpoint": *expression, optional*

The URL to the provider's OpenID Connect UserInfo endpoint.

Default: no UserInfo is obtained from the provider.

See also Expressions(5).

"issuerHandler": *Handler reference, optional*

Invoke this HTTP client handler to communicate with the authorization server.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

Usually set this to the name of a ClientHandler configured in the heap, or a chain that ends in a ClientHandler.

Default: IG uses the default ClientHandler.

See also Handlers, ClientHandler(5).

"issuerRepository": *Issuer repository reference, optional*

A repository of OAuth 2.0 issuers, built from discovered issuers and the IG configuration.

Provide the name of an IssuerRepository object defined in the heap.

Default: Look up an issuer repository named `IssuerRepository` in the heap. If none is explicitly defined, then a default one named `IssuerRepository` is created in the current route.

See also IssuerRepository(5).

"supportedDomains": *array of patterns, optional*

List of patterns matching domain names handled by this issuer, used as a shortcut for OpenID Connect discovery before performing OpenID Connect dynamic registration.

In summary when the OpenID Provider is not known in advance, it might be possible to discover the OpenID Provider Issuer based on information provided by the user, such as an email address. The OpenID Connect discovery specification explains how to use WebFinger to discover the issuer. IG can discover the issuer in this way. As a shortcut IG can also use supported domains lists to find issuers already described in the IG configuration.

To use this shortcut, IG extracts the domain from the user input, and looks for an issuer whose supported domains list contains a match.

Supported domains patterns match host names with optional port numbers. Do not specify a URI scheme such as HTTP. IG adds the scheme. For instance, `*.example.com` matches any host in the `example.com` domain. You can specify the port number as well as in `host.example.com:8443`. Patterns must be valid regular expression patterns according to the rules for the Java Pattern class.

Examples

The following example shows an AM issuer configuration for AM. AM exposes a well-known endpoint for the provider configuration, but this example demonstrates use of the other fields:

```
{
    "name": "openam",
    "type": "Issuer",
    "config": {
        "authorizeEndpoint":
          "https://openam.example.com:8443/openam/oauth2/authorize",
        "registration_endpoint":
          "https://openam.example.com:8443/openam/oauth2/connect/register",
        "tokenEndpoint":
          "https://openam.example.com:8443/openam/oauth2/access_token",
        "userInfoEndpoint":
          "https://openam.example.com:8443/openam/oauth2/userinfo",
        "supportedDomains": [ "mail.example.*", "docs.example.com:8443" ]
    }
}
```

The following example shows an issuer configuration for Google:

```
{
    "name": "google",
    "type": "Issuer",
    "config": {
        "wellKnownEndpoint":
          "https://accounts.google.com/.well-known/openid-configuration",
        "supportedDomains": [ "gmail.*", "googlemail.com:8052" ]
    }
}
```

More Information

org.forgerock.openig.filter.oauth2.client.Issuer

Name

IssuerRepository — Store discovered and built OAuth2 issuers in a repository

Description

A repository to store OAuth2 issuers that are discovered or built from the configuration.

It is not normally necessary to change this object. Change it only for the following tasks:

- To isolate different repositories in the same route.

- To view the interactions of the well-known endpoint, for example, if the `issuerHandler` is delegating to another handler.

Usage

```
{
    "name": string,
    "type": "IssuerRepository",
    "config": {
      "issuerHandler": Handler reference
    }
}
```

Properties

The object properties are as follows:

`"issuerHandler"`: *handler reference, optional*

The default handler to fetch OAuth2 issuer configurations from the well-known endpoint.

Provide the name of a Handler object defined in the heap, or an inline Handler configuration object.

Default: ForgeRockClientHandler

More Information

org.forgerock.openig.filter.oauth2.client.IssuerRepository

Name

ScheduledExecutorService — schedule the execution of tasks

Description

An executor service to schedule tasks for execution after a delay or for repeated execution with a fixed interval of time in between each execution. You can configure the number of threads in the executor service and how the executor service is stopped.

The `ScheduledExecutorService` is shared by all downstream components that use an executor service.

Usage

```
{
    "name": string,
    "type": "ScheduledExecutorService",
    "config": {
        "corePoolSize":  integer or expression<integer>,
        "gracefulStop":  boolean or expression<boolean>,
        "gracePeriod" :  duration string or expression<duration string>
    }
}
```

Properties

`"corePoolSize"`: *integer or expression<integer>, optional*

The minimum number of threads to keep in the pool. If this property is an expression, the expression is evaluated as soon as the configuration is read.

The value must be an integer greater than zero.

Default: 1

`"gracefulStop"`: *boolean or expression<boolean> , optional*

Defines how the executor service stops. If this property is an expression, the expression is evaluated as soon as the configuration is read.

If true, the executor service does the following:

• Blocks the submission of new jobs.

• Allows running jobs to continue.

• If a grace period is defined, waits for up to that maximum time for running jobs to finish before it stops.

If false, the executor service does the following:

- Blocks the submission of new jobs.

- Removes submitted jobs without running them.

- Attempts to end running jobs.

- If a grace period is defined, ignores it.

Default: true

`"gracePeriod"`: *duration string or expression<duration string>, optional*

The maximum time that the executor service waits for running jobs to finish before it stops. If this property is an expression, the expression is evaluated as soon as the configuration is read.

If all jobs finish before the grace period, the executor service stops without waiting any longer. If jobs are still running after the grace period, the executor service stops anyway and prints a message.

When `gracefulStop` is `false`, the grace period is ignored.

Default: 10 seconds

For information about supported formats for `duration`, see duration.

Example

The following example creates a thread pool to execute tasks. When the executor service is instructed to stop, it blocks the submission of new jobs, and waits for up to 10 seconds for submitted and running jobs to complete before it stops. If any jobs are still submitted or running after 10 seconds, the executor service stops anyway and prints a message.

```
{
    "name": "ExecutorService",
    "comment": "Default service for executing tasks in the background.",
    "type": "ScheduledExecutorService",
    "config": {
        "corePoolSize": 5,
        "gracefulStop": true,
        "gracePeriod": "10 seconds"
    }
}
```

More Information

org.forgerock.openig.thread.ScheduledExecutorServiceHeaplet

Name

SecretsProvider — provide a secrets service for StatelessAccessTokenResolver

Description

Provide a secrets service for the StatelessAccessTokenResolver(5), that uses specified secret stores to resolve access_tokens.

If SecretsProvider is not configured, the StatelessAccessTokenResolver uses the global secrets service defined for the route, which looks for keys across the whole configuration.

Usage

```
{
  "name": string,
  "type": "SecretsProvider",
  "config": {
    "stores": [ secret store declaration, ... ]
  }
}
```

Properties

`"stores"`: *array of secret store declarations, optional*

One or more secret stores to provide a secrets service for the StatelessAccessTokenResolver. Configure secret stores described in Secrets.

Default: No stores. When resolving a token, use the global secrets service to resolve access_tokens.

Example

For an example that uses SecretsProvider, see "Validating Stateless Access_Tokens With the StatelessAccessTokenResolver" in the *Gateway Guide*.

More Information

StatelessAccessTokenResolver(5)

Secrets

org.forgerock.secrets.SecretsProvider

Name

TemporaryStorage — cache streamed content

Description

Allocates temporary buffers for caching streamed content during request processing. Initially uses memory; when the memory limit is exceeded, switches to a temporary file.

Usage

```
{
    "name": string,
    "type": "TemporaryStorage",
    "config": {
        "initialLength": number,
        "memoryLimit": number,
        "fileLimit": number,
        "directory": string
    }
}
```

Properties

`"initialLength"`: *number, optional*

Initial size of the memory buffer.

Default: 8 192 bytes (8 KB). Maximum: The value of `"memoryLimit"`.

`"memoryLimit"`: *number, optional*

Maximum size of the memory buffer. When the memory buffer is full, the content is transferred to a temporary file.

Default: 65 536 bytes (64 KB). Maximum: 2 147 483 647 bytes (2 GB).

`"fileLimit"`: *number, optional*

Maximum size of the temporary file. If the downloaded file is bigger than this value, an OverflowException is thrown.

Default: 1 073 741 824 bytes (1 GB). Maximum: 2 147 483 647 bytes (2 GB).

`"directory"`: *string, optional*

The directory where temporary files are created.

Default: The value of the system property `java.io.tmpdir`, typically `/tmp` on Unix systems, or `/var/tmp` on Linux.

More Information

org.forgerock.openig.io.TemporaryStorageHeaplet

Name

TlsOptions — Configure connections to TLS-protected endpoints

Description

Configure connections to TLS-protected endpoints. Use in ClientHandler(5), ReverseProxyHandler(5), and AmService(5).

Usage

```
{
  "type": TlsOptions,
  "config": {
    "keyManager": KeyManager reference or [ KeyManager reference, ...],
    "sslCipherSuites": [configuration expression<string>, ...],
    "sslContextAlgorithm": configuration expression<string>,
    "sslEnabledProtocols": [configuration expression<string>, ...],
    "trustManager": TrustManager reference or [ TrustManager reference, ...]
  }
}
```

Properties

`"keyManager"`: *KeyManager reference or array of KeyManager references, optional*

One or more key managers to handle this client's keys and certificates.

Provide either the name(s) of KeyManager object(s) defined in the heap, or specify the configuration object(s) inline.

Specify a single KeyManager or an array of KeyManagers:

- `"keyManager": "MyKeyManager"`

- `"keyManager": [ "FirstKeyManager", "SecondKeyManager" ]`

If keyManager is not configured, the client can't authenticate itself in a mutual authentication scenario.

See also KeyManager(5).

`"sslCipherSuites"`: *array of configuration expression<string>, optional*

Array of cipher suite names, used to restrict the cipher suites allowed when negotiating transport layer security for an HTTPS connection.

For information about the available cipher suite names, see the documentation for the Java virtual machine (JVM) used by the container where you run IG. For Oracle Java, see the list of *JSSE Cipher Suite Names*.

Default: Allow any cipher suite supported by the JVM.

`"sslContextAlgorithm"`: *configuration expression<string>, optional*

The `SSLContext` algorithm name, as listed in the table of *`SSLContext` Algorithms* for the Java Virtual Machine used by the container where IG runs.

Default: `TLS`

`"sslEnabledProtocols"`: *array of configuration expression<string>, optional*

Array of protocol names, used to restrict the protocols allowed when negotiating transport layer security for an HTTPS connection.

For information about the available protocol names, see the documentation for the Java virtual machine (JVM) used by the container where you run IG. For Oracle Java, see the list of *Additional JSSE Standard Names*.

Default: Allow any protocol supported by the JVM.

`"trustManager"`: *TrustManager reference or array of TrustManager references, optional*

One or more trust managers that handle peers' public key certificates.

Provide either the name(s) of TrustManager object(s) defined in the heap, or specify the configuration object(s) inline.

Specify a single TrustManager or an array of TrustManager:

- `"trustManager": "TrustManager"`

- `"trustManager": [ "FirstTrustManager", "SecondTrustManager" ]`

If you do not configure a trust manager, then the client uses only the default Java truststore. The default Java truststore depends on the Java environment. For example, `$JAVA_HOME/lib/security/cacerts`.

See also TrustManager(5).

Example

```
{
  "tls": {
    "type": "TlsOptions",
    "config": {
      "sslContextAlgorithm": "TLSv1.2",
      "keyManager": {
        "type": "KeyManager",
        "config": {
          "keystore": {
            "type": "KeyStore",
            "config": {
```

```
            "url": "file://${env['HOME']}/keystore.jks",
            "passwordSecretId": "keymanager.keystore.secret.id"
          }
        },
        "passwordSecretId": "keymanager.secret.id"
      }
    },
    "trustManager": {
      "type": "TrustManager",
      "config": {
        "keystore": {
          "type": "KeyStore",
          "config": {
            "url": "file://${env['HOME']}/truststore.jks",
            "passwordSecretId": "trustmanager.keystore.secret.id"
          }
        }
      }
    }
  }
}
```

Name

TrustManager — configure a Java Secure Socket Extension TrustManager

Description

This represents the configuration for a Java Secure Socket Extension TrustManager, which manages the trust material (typically X.509 public key certificates) used to decide whether to accept the credentials presented by a peer. The configuration references the keystore that actually holds the trust material.

Usage

```
{
    "name": string,
    "type": "TrustManager",
    "config": {
        "keystore": KeyStore reference,
        "alg": string
    }
}
```

Properties

`"keystore"`: *KeyStore reference, optional*

The KeyStore that references the store for public key certificates.

Provide either the name of the KeyStore object defined in the heap, or the inline KeyStore configuration object inline.

See also KeyStore(5).

`"alg"` *string, optional*

The certificate algorithm to use.

Default: the default for the platform, such as `SunX509`.

Example

The following example configures a trust manager that depends on a KeyStore configuration. This configuration uses the default certificate algorithm:

```
{
    "name": "MyTrustManager",
    "type": "TrustManager",
    "config": {
        "keystore": {
            "type": "KeyStore",
            "config": {
                "url": "file://${env['HOME']}/keystore.jks",
                "passwordSecretId": "${system['keypass']}"
            }
        }
    }
}
```

More Information

org.forgerock.openig.security.TrustManagerHeaplet

JSSE Reference Guide, KeyManager(5), KeyStore(5)

Name

TrustAllManager — a TrustManager that blindly trusts all servers

Description

The TrustAllManager blindly trusts all server certificates presented the servers for protected applications. It can be used instead of a TrustManager(5) in test environments to trust server certificates that were not signed by a well-known CA, such as self-signed certificates.

The TrustAllManager is not safe for production use. Use a properly configured TrustManager(5) instead.

Usage

```
{
    "name": string,
    "type": "TrustAllManager"
}
```

Example

The following example configures a client handler that blindly trusts server certificates when IG connects to servers over HTTPS:

```
{
    "name": "BlindTrustClientHandler",
    "type": "ReverseProxyHandler",
    "config": {
        "trustManager": {
            "type": "TrustAllManager"
        }
    }
}
```

More Information

org.forgerock.openig.security.TrustAllManager

Name
UmaService — represent an UMA resource server configuration

Description

The UmaService includes a list of resource patterns and associated actions that define the scopes for permissions to matching resources. When creating a share using the REST API described below, you specify a path matching a pattern in a resource of the UmaService.

UMA 2.0 is supported with AM 5.5 and later versions. UMA 1.0 is supported with AM 5.1 and later versions.

Usage

```
{
    "type": "UmaService",
    "config": {
        "protectionApiHandler": Handler reference,
        "amService": AmService reference,    // Use either "amService"
        "wellKnownEndpoint": URI string,     // or "wellKnownEndpoint", but not both.
        "resources": [ resource, ... ]
    }
}
```

Properties

`"protectionApiHandler"`: *Handler reference, required*

The handler to use when interacting with the UMA authorization server to manage resource sets, such as a ClientHandler capable of making an HTTPS connection to the server.

For more information, see Handlers.

`"amService"`: *AmService reference, required if `"wellKnownEndpoint"` is not configured*

The AmService heap object to use for the URI to the well-known endpoint for this UMA authorization server. The endpoint is extrapolated from the `url` property of the AmService, and takes the realm into account.

If the UMA authorization server is AM, use this property to define the endpoint.

If `amService` is configured, it takes precedence over `wellKnownEndpoint`.

For more information about the endpoint for UMA configuration, see Discovering UMA Configuration in the AM *User-Managed Access (UMA) 2.0 Guide*.

See also, AmService(5).

`"wellKnownEndpoint"`: *URI string, required if `"amService"` is not configured*

The URI to the well-known endpoint for this UMA authorization server.

If the UMA authorization server is not AM, use this property to define the endpoint.

If `amService` is configured, it takes precedence over `wellKnownEndpoint`.

In this example, the UMA configuration is in the default realm of AM:

```
https://openam.example.com:8088/openam/uma/.well-known/uma2-configuration
```

In this example, the UMA configuration is in a European customer realm:

```
https://openam.example.com:8088/openam/uma/realms/root/realms/customer/realms/europe/.well-known/uma2-
configuration
```

For more information about the endpoint for UMA configuration, see Discovering UMA
Configuration in the AM *User-Managed Access (UMA) 2.0 Guide*.

`"resources"`: *array of resources, required*

Resource objects matching the resources the resource owner wants to share.

Each resource object has the following form:

```
{
    "pattern": resource pattern,
    "actions": [
        {
            "scopes": [ scope string, ... ],
            "condition": runtime expression<boolean>
        },
        {
            ...
        }
    ]
}
```

Each resource pattern can be seen to represent an application, or a consistent set of endpoints
that share scope definitions. The actions map each request to the associated scopes. This
configuration serves to set the list of scopes in the following ways:

1. When registering a resource set, IG uses the list of actions to provide the aggregated,
    exhaustive list of all scopes that can be used.

2. When responding to an initial request for a resource, IG derives the scopes for the ticket
    based on the scopes that apply according to the request.

3. When verifying the RPT, IG checks that all required scopes are encoded in the RPT.

A description of each field follows:

`"pattern"`: *resource pattern, required*

> A pattern matching resources to be shared by the resource owner, such as `.*` to match any resource path, and `/photos/.*` to match paths starting with `/photos/`.
>
> See also Patterns(5).

`"actions"`: *array of action objects, optional*

> A set of actions on matching resources that the resource owner can authorize.
>
> When granting permission, the resource owner specifies the action scope. Conditions specify what the scopes mean in concrete terms. A given scope matches a requesting party operation when the corresponding condition evaluates to `true`.

`"scopes"`: *array of scope strings, optional*

> Scope strings to identify permissions.
>
> For example, `#read` (read access on a resource).

`"condition"`: *runtime expression<boolean>, required*

> A boolean expression representing the meaning of a scope.
>
> For example, `${request.method == 'GET'}` (true when reading a resource).
>
> See also Expressions(5).

The REST API for Shares

The REST API for UMA shares is exposed at a registered endpoint. IG logs the paths to registered endpoints when the log level is `INFO` or finer. Look for messages such as the following in the log:

```
UMA Share endpoint available at
  '/openig/api/system/objects/_router/routes/00-uma/objects/umaservice/share'
```

To access the endpoint over HTTP or HTTPS, prefix the path with the IG scheme, host, and port to obtain a full URL, such as `http://localhost:8080/openig/api/system/objects/_router/routes/00-uma/objects/umaservice/share`.

The UMA REST API supports create (POST only), read, delete, and query (`_queryFilter=true` only). For an introduction to common REST APIs, see "About ForgeRock Common REST".

In the present implementation, IG does not have a mechanism for persisting shares. When the IG container stops, the shares are discarded.

For information about API descriptors for the UMA share endpoint, see "Understanding IG APIs With API Descriptors" in the *Gateway Guide*. For information about Common REST, see "About ForgeRock Common REST".

A share object has the following form:

```
{
    "path": pattern,
    "pat": UMA protection API token (PAT) string,
    "id": unique identifier string,
    "resource_id": unique identifier string,
    "user_access_policy_uri": URI string
}
```

The fields are as follows:

`"path"`: *pattern, required*

A pattern matching the path to protected resources, such as `/photos/.*`.

This pattern must match a pattern defined in the UmaService for this API.

See also Patterns(5).

`"pat"`: *PAT string, required*

A PAT granted by the UMA authorization server given consent by the resource owner.

In the present implementation, IG has access only to the PAT, not to any refresh tokens.

`"id"`: *unique identifier string, read-only*

This uniquely identifies the share. This value is set by the service when the share is created, and can be used when reading or deleting a share.

`"resource_id"`: *unique identifier string, read-only*

This uniquely identifies the UMA resource set registered with the authorization server. This value is obtained by the service when the resource set is registered, and can be used when setting access policy permissions.

`"user_access_policy_uri"`: *URI string, read-only*

This URI indicates the location on the UMA authorization server where the resource owner can set or modify access policies. This value is obtained by the service when the resource set is registered.

More Information

User-Managed Access (UMA) 2.0 Grant for OAuth 2.0 Authorization

org.forgerock.openig.uma.UmaSharingService

Property Value Substitution

In an environment with multiple IG instances, you can require similar but not identical configurations across the different instances.

Property value substitution enables you to do the following:

- Define a configuration that is specific to a single instance, for example, setting the location of the keystore on a particular host.

- Define a configuration whose parameters vary between different environments, for example, the URLs and passwords for test, development, and production environments.

- Disable certain capabilities on specific nodes.

Property value substitution uses *configuration tokens* to introduce variables into the server configuration. For information, see Configuration Tokens(5).

The substitution follows a process of token resolution, JSON evaluation, and data transformation, as described in the following sections:

- JSON Evaluation(5)

- Token Resolution(5)

- Transformations(5)

Table of Contents

Name

Configuration Tokens — introduce variables into the server configuration

Description

A configuration token is a simple reference to a value. When configuration tokens are resolved, the result is always a string. Transformation described in Transformations(5) can be used to coerce the output type.

Configuration Tokens for File System

IG provides `ig.instance.dir` and `ig.instance.url` to define the file system directory and URL for configuration files.

Their values are computed at startup, and evaluate to a directory such as `$HOME/.openig` (`%appdata% \OpenIG`). You can use these tokens in your configuration without explicitly setting their values.

For information about how to change the default values, see "Changing the Default Location of the Configuration Folders" in the *Gateway Guide*.

Configuration Token for Run Mode

If `mode` is not set in `admin.json`, the provided configuration token `ig.run.mode` can be resolved at startup to define the run mode.

Set the value to `development` or `production`, the value is not case-sensitive:

- **Development mode (mutable mode)**

  Use development mode to evaluate or demo IG, 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 IG 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 IG configuration.

- **Production mode (immutable 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. Studio is effectively disabled, and you cannot manage, list, or even read routes through Common REST.

By default, other endpoints, such as `/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 IG configuration.

For more information about `admin.json`, see AdminHttpApplication(5).

Syntax

Configuration tokens follow the syntax `&{token[|default]}`, as follows:

- Are preceded by an ampersand, `&`

- Are enclosed in braces, `{}`

- Define default values with a vertical bar (`|`) after the configuration token

- Are in lowercase

- Use the period as a separator, `.`

When a configuration token is supplied in a configuration parameter, it is always inside a string enclosed in quotation marks, as shown in the following example:

```
"&{listen.port|8080}"
```

To escape a string with the syntax of a configuration token, use a backslash (`\`). The following string is treated as normal text:

```
"\&{listen.port|8080}"
```

A configuration property can include a mix of static values and expressions, as shown in the following example:

```
"&{hostname}.example.com"
```

Configuration tokens can be nested inside other configuration tokens as shown in the following example:

```
"&{&{protocol.scheme}.port}"
```

Default values or values in the property resolver chain can be nested, as shown in the following example:

```
"&{&{protocol.scheme|http}.port|8080}"
```

Name

JSON Evaluation — evaluation JSON files to substitute tokens and transform data types

Description

JSON evaluation is the process of substituting configuration tokens and transforming JSON nodes for an entire JSON configuration. After JSON evaluation, all configuration tokens and transformations in the configuration are replaced by values.

At startup, IG evaluates the configuration tokens in `config.json` and `admin.json`. When routes are deployed, IG evaluates the configuration tokens in the route.

Configuration tokens are matched with tokens available in the chain of resolvers, and the configuration token is substituted with the value available in the resolver. For information about each of the resolvers mentioned in the following section, see Token Resolution(5).

IG searches for matching tokens in the chain of resolvers, using the following order of precedence:

1. Local resolver:

   • The route resolver for the route being deployed

2. Intermediate resolver:

   • All intermediate route resolvers (for example, for parent routes to the route being deployed) up to the bootstrap resolver

3. Bootstrap resolver:

   a. Environment variables resolver

   b. System properties resolver

   c. Token source file resolvers

   d. Hardcoded default values

The first resolver that matches the token returns the value of the token.

If the token can't be resolved, IG uses the default value defined with the configuration token. If there is no default value, the token can't be resolved and an error occurs:

• If the configuration token is in `config.json` or `admin.json`, IG fails to start up.

• If the configuration token is in a route, the route fails to load.

When configuration tokens are nested inside other configuration tokens, the tokens are evaluated bottom-up, or leaf-first. For example, if the following configuration token takes only the default values, it is resolved as follows:

1. `"&{&{protocol.scheme|http}.port|8080}"`

2. `"&{http.port|8080}"`

   When `&{protocol.scheme|http}` takes the default value `http`.

3. `"8080"`

   When `&{http.port|8080}` takes the default value `8080`.

If the configuration includes a transformation, IG applies the transformation after the token is substituted. When transformations are nested inside other transformations, the transformations are applied bottom-up, or leaf-first. For more information, see Transformations(5).

FORGEROCK

Name

Token Resolution — build a chain of token resolvers

Description

At startup, the bootstrap resolver builds a chain of resolvers to resolve configuration tokens included in `config.json` and `admin.json`. When a route is deployed, *route resolvers* build on the chain to add resolvers for the route.

Route Token Resolvers

When a route is deployed in IG a route resolver is created to resolve the configuration tokens for the route. The resolvers uses token values defined in the `properties` section of the route.

If the token can't be resolved locally, the route resolver accesses token values recursively in a parent route.

For more information, about route properties, see Properties(5).

Environment Variables Resolver

When the bootstrap resolver resolves a configuration token to an environment variable, it replaces the lowercase and periods (`.`) in the token to match the convention for environment variables.

Environment variable keys are transformed as follows:

• Periods (.) are converted to underscores

• All characters are transformed to uppercase

The following example sets the value of an environment variable for the port number:

```
$ export LISTEN_PORT=8080
```

In the following IG configuration, the value of `port` is `8080`:

```
{
  "port": "&{listen.port}"
}
```

System Properties Resolver

The system property name must match a configuration token exactly. The following example sets a system property for a port number:

```
$ java -Dlisten.port=8080 -jar start.jar
```

In the following IG configuration, the value of `port` is `8080`:

```
{
  "port": "&{listen.port}"
}
```

Token Source File Resolvers

Token source files have the `.json` or `.properties` extension. The bootstrap resolver uses the files to add file resolvers to the chain of resolvers:

- **JSON file resolvers**

  Token source files with the `.json` extension take a JSON format. The token name is mapped either to the JSON attribute name or to the JSON path.

  Each of the following `.json` files set the value for the configuration token `product.listen.port` :

  ```
  {
    "product.listen.port": 8080
  }
  ```

  ```
  {
    "product.listen {
      "port": 8080
    }
  }
  ```

  ```
  {
    "product" {
      "listen" {
        "port": 8080
      }
    }
  }
  ```

- **Properties file resolvers**

  Token source files with the `.properties` extension are Java properties files. They contain a flat list of key/value pairs, and keys must match tokens exactly.

  The following `.properties` file also sets the value for the tokens `listen.port` and `listen.address`:

  ```
  listen.port: 8080
  listen.address: 192.168.0.10
  ```

Token source files are stored in one or more directories defined by the environment variable `IG_ENVCONFIG_DIRS` or the system property `ig.envconfig.dirs`.

If token source files are in multiple directories, each directory must be specified in a comma-separated list. IG doesn't scan subdirectories. The following example sets an environment variable to define two directories that hold token source files:

```
$ export IG_ENVCONFIG_DIRS="/myconfig/directory1, /myconfig/directory2"
```

At startup, the bootstrap resolver scans the directories in the specified order, and adds a resolver to the chain of resolvers for each token source file in the directories.

Although the bootstrap resolver scans the directories in the specified order, within a directory it scans the files in a nondeterministic order.

Note the following constraints for using the same configuration token more than once:

- Do not define the same configuration token more than once in a single file. There is no error, but you won't know which token is used.

- Do not define the same configuration token in more than one file in a single directory. An error occurs.

> **Important**
>
> This constraint implies that you can't have backup `.properties` and `.json` files in a single directory if they define the same tokens.

- You can define the same configuration token once in several files that are located in different directories, but the first value that IG reads during JSON evaluation is used.

> **Note**
>
> When logging is enabled at the DEBUG level for token resolvers, the origin of the token value is logged.
>
> If you are using the logback implementation in " Reference Logback Configuration " in the *Gateway Guide*, add the following line to your `logback.xml` to enable logging:
>
> ```
> <logger name="org.forgerock.config.resolvers" level="DEBUG"/>
> ```

Name

Transformations — built-in transformations to coerce strings to other data types

Description

A set of built-in transformations are available to coerce strings to other data types. The transformations can be applied to any string, including strings resulting from the resolution of configurations tokens.

After transformation, the JSON node representing the transformation is replaced by the result value.

Usage

```
{
    "$transformation": string or transformation
}
```

A transformation is a JSON object with a required main attribute, starting with a `$`. The following example transforms a string to an integer:

```
{"$int": string}
```

.

The value of a transformation value can be a JSON string or another transformation that results in a string. The following example shows a nested transformation:

```
{
  "$array": {
    "$base64:decode": string
  }
}
```

The input string must match the format expected by the transformation. In the previous example, because the final transformation is to an array, the input string must be a string that represents an array, such as `"[ \"one\", \"two\" ]"`.

In the first transformation, the encoded string is transformed to a base64-decoded string. In the second, the string is transformed into a JSON array, for example, `[ "one", "two" ]`.

array

```
{"$array": string}
```

Returns a JSON array of the argument.

Argument

string

String representing a JSON array.

Returns

array

> JSON array.

Example

The transformation `{"$array": "[ \"one\", \"two\" ]"}` results in the JSON array `[ "one", "two" ]`.

bool

```
{"$bool": string}
```

Returns a Boolean with a value represented by the argument.

Returns `true` if the input value equals `"true"` (ignoring case). Otherwise, returns `false`.

Argument

string

> String containing the boolean representation.

Returns

boolean

> Boolean value represented by the argument.

Example

If the configuration token `&{capture.entity}"` resolves to `"true"`, the transformation `{"$bool": "&{capture.entity}"}` results in the value `true`.

decodeBase64

```
{
  "$base64:decode": string,
  "$charset": "charset"
}
```

Transforms a base64-encoded string into a decoded string.

If `$charset` is specified, the decoded value is interpreted with the character set.

Argument

string

> Base64-encoded string.

Parameters

$charset

> The name of a Java character set, as described in Class Charset.

Returns

string

> Base64-decoded string in the given character set.

Example

The transformation

```
{
 "$base64:decode": "SGVsbG8=",
 "$charset": "UTF-8"
}
```

returns the `Hello` string.

encodeBase64

```
{
  "$base64:encode": string,
  "$charset": "charset"
}
```

Transforms a string into a base64-encoded string. Transforms to `null` if the string is `null`.

If `$charset` is specified, the string is encoded with the character set.

Argument

string

> String to encode with the given character set.

Parameters

$charset

The name of a Java character set, as described in Class Charset.

Returns

string

Base64-encoded string.

int

```
{"$int": string}
```

Transforms a string into an integer.

If the parameter is not a valid number in radix 10, returns `null`.

Argument

string

String containing the integer representation.

Returns

int

Integer value represented by the argument.

Example

The transformation `{"$int": "1234"}`, results in the integer `1234`.

list

```
{$list: string}
```

Transforms a comma-separated list of strings into a JSON array of strings

Argument

string

A string representing a comma-separated list of strings.

Returns

array

> The JSON array of the provided argument. Values are not trimmed of leading spaces.

Example

The transformation `{"$list": "Apple,Banana,Orange,Strawberry"}`, results in the array of strings `["Apple"` `,"Banana","Orange","Strawberry"]`.

The transformation `{"$list": "Apple, Banana, Orange, Strawberry"}` results in the array of strings `["Apple"` `," Banana"," Orange"," Strawberry"]`, including the untrimmed spaces.

The transformation `{"$list": "1,2,3,4"}` results in the array of strings `["1","2","3","4"]`, and not an array of JSON numbers `[1,2,3,4]`.

number

```
{$number: string}
```

Transform a string into a Java number, as defined in Class Number.

Argument

strings

> A string containing the number representation.

Returns

number

> The number value represented by the argument.

Example

The transformation `{"$number": ".999"}`, results in the number `0.999`.

object

```
{"$object": string}
```

Transforms a string representation of a JSON object into a JSON object.

FORGEROCK

Argument

string

String representation of a JSON object.

Returns

object

JSON object of the argument.

Example

The transformation `{"$object": "{\"ParamOne\":{\"InnerParamOne\":\"InnerParamOneValue\",\"InnerParamTwo\": false}}"}`, results in the following JSON object:

```
{
  "ParamOne": {
    "InnerParamOne": "myValue",
    "InnerParamTwo": false
  }
}
```

Expressions

Many configuration parameters support dynamic expressions.

Table of Contents

Name

Expressions — expression configuration parameter values

Description

Expressions are specified as configuration parameter values for a number of built-in objects. Such expressions conform to the Universal Expression Language as specified in JSR-245.

General Syntax

All expressions follow standard Universal Expression Language syntax: `${expression}`. The expression can be a simple reference to a value, a function call, or arbitrarily complex arithmetic, logical, relational and conditional operations. When supplied within a configuration parameter, an expression is always a string enclosed in quotation marks, for example: `"${request.method}"`.

Configuration and Runtime Expressions

Expressions are evaluated at configuration time (when routes are loaded), or at runtime (when IG is running).

When expressions are evaluated, they access the current environment through the implicit object `openig`. The object has the following properties:

- `baseDirectory`, the path to the base location for IG files. The default location is `$HOME/.openig` (`%appdata%\OpenIG`).

- `configDirectory`, the path to the IG configuration files. The default location is `$HOME/.openig/config` (`%appdata%\OpenIG\config`).

- `tempDirectory`, the path to the IG temporary files. The default location is `$HOME/.openig/tmp` (`%appdata%\OpenIG\tmp`).

For information about how to change the default values, see "Changing the Default Location of the Configuration Folders" in the *Gateway Guide*.

configuration expression

> Expression evaluated at configuration time, when routes are loaded.

> Configuration expressions can refer to the system heap properties, the built-in functions listed in Functions(5), the `${env['`*`variable`*`']}`, and `${system['`*`property`*`']}`. Because configuration expressions are evaluated before any requests are made, they cannot refer to the runtime properties, `request`, `response`, or `context`. For more information, see Expressions(5).

runtime expression

> Expression evaluated at runtime, for each request and response.

Runtime expressions can refer to the same information as configuration expressions, plus the following objects:

- `attributes`: org.forgerock.services.context.AttributesContext Map<String, Object>, obtained from `AttributesContext.getAttributes()`. For information, see AttributesContext(5).

- `context`: org.forgerock.services.context.Context object.

- `contexts`: map<string, context> object. For information, see Contexts(5).

- `request`: org.forgerock.http.protocol.Request object. For information, see Request(5).

- `response`: org.forgerock.http.protocol.Response object, available only when the expression is intended to be evaluated on the response flow. For information, see Response(5).

- `session`: org.forgerock.http.session.Session object, available only when the expression is intended to be evaluated for both request and response flow. For information, see SessionContext(5).

Value Expressions

A value expression references a value relative to the scope supplied to the expression. For example, `"${request.method}"` references the method of an incoming HTTP request.

An *lvalue-expression* is a specific type of value expression that references a value to be written. For example, `"${session.gotoURL}"` specifies a session attribute named `gotoURL` to write a value to. Attempts to write values to read-only values are ignored.

Indexed Properties

Properties of values are accessed using the `.` and `[]` operators, and can be nested arbitrarily.

The value expressions `"${request.method}"` and `"${request['method']}"` are equivalent.

To prevent errors, in property names containing characters that are also expression operators, use the `[]` operator instead of the `.` operator. For example, use `contexts.amSession.properties['a-b-c']` instead of `contexts.amSession.properties.a-b-c`.

In the case of arrays, the index of an element in the array is expressed as a number in brackets. For example, `"${request.headers['Content-Type'][0]}"` references the first `Content-Type` header value in a request. If a property does not exist, then the index reference yields a `null` (empty) value.

Operators

Universal Expression Language provides the following operators:

- Index property value: `[]`, `.`

- Change precedence: `()`

- Arithmetic: `+` (binary), `-` (binary), `*`, `/`, `div`, `%`, `mod`, `-` (unary)

- Logical: `and`, `&&`, `or`, `||`, `not`, `!`

- Relational: `==`, `eq`, `!=`, `ne`, `<`, `lt`, `>`, `gt`, `<=`, `le`, `>=`, `ge`

- Empty: `empty`

  Prefix operation that can be used to determine whether a value is null or empty.

- Conditional: `?`, `:`

Operators have the following precedence, from highest to lowest, and from left to right:

- `[] .`

- `()`

- `-` (unary) `not ! empty`

- `* / div % mod`

- `+` (binary) `-` (binary)

- `< > <= >= lt gt le ge`

- `== != eq ne`

- `&& and`

- `|| or`

- `? :`

System Properties and Environment Variables

You can use expressions to retrieve Java system properties, and to retrieve environment variables.

For system properties, `${system['`*`property`*`']}` yields the value of *property*, or `null` if there is no value for *property*. For example, `${system['user.home']}` yields the home directory of the user running the application server for IG.

For environment variables, `${env['`*`variable`*`']}` yields the value of *variable*, or `null` if there is no value for *variable*. For example, `${env['HOME']}` yields the home directory of the user running the application server for IG.

Token Resolution

Runtime expressions have access to evaluated configuration tokens described in JSON Evaluation(5). For example, the following boolean expression returns `true` if the configuration token `my.status.code` resolves to `200`:

```
${integer(_token.resolve('my.status.code', '404')) == 200}
```

Functions

A number of built-in functions described in Functions(5) can be called within an expression.

The syntax is `${function(parameter, ...)}`, where zero or more parameters are supplied to the function. For example:

- `"${bool(env['ENABLE_TIMER'])}"` recovers the environment variable `"ENABLE_TIMER"` and transforms it into a boolean

- `"${toLowerCase(request.method)}"` yields the method of the request, converted to lower case.

Functions can be operands for operations, and can yield parameters for other function calls.

Escaping Literal Expressions

The character `\` is treated as an escape character when it is followed by `${` or `#{`. For example, the expression `${true}` normally evaluates to `true`. To include the string `${true}` in an expression, write `${true}`

When the character `\` is followed by any other character sequence, it is not treated as an escape character. For example, `${` evaluates to `${`, but `\$` evaluates to `\$`.

Embedding Expressions

Although an expression cannot be embedded as `${expression}` inside another expression, embedding system property, environment variable, and function expressions within each other is fine. Do not enclose the embedded elements in `${}`.

The following single line example embeds an `env` environment variable expression and the Java `String.concat()` method in the argument to a `read()` function:

```
"entity" : "${read('&{ig.instance.dir}/html/defaultResponse.html')}"
```

In the example the entity property value is set to the contents of the file `$HOME/.openig/html/defaultResponse.html`.

Extensions

IG offers a plugin interface for extending expressions. See "Key Extension Points" in the *Gateway Guide*.

If your deployment uses expression plugins, read the plugin documentation about the additional expressions you can use.

Examples

```
"${request.uri.path == '/wordpress/wp-login.php'
   and request.form['action'][0] != 'logout'}"

"${request.uri.host == 'wiki.example.com'}"

"${request.cookies[keyMatch(request.cookies,'^SESS.*')][0].value}"

"${toString(request.uri)}"

"${request.method == 'POST' and request.uri.path == '/wordpress/wp-login.php'}"

"${request.method != 'GET'}"

"${request.headers['cookie'][0]}"

"${request.uri.scheme == 'http'}"

"${not (response.status.code == 302 and not empty session.gotoURL)}"

"${response.headers['Set-Cookie'][0]}"

"${request.headers['host'][0]}"

"${not empty system['my-variable'] ? system['my-variable'] : '/path/to'}/logs/gateway.log"
```

More Information

Contexts(5), Functions(5), Request(5), Response(5)

Name

Functions — built-in functions to call within expressions

Description

A set of built-in functions that can be called from within expressions, which are described in Expressions(5).

array

```
array(strings...)
```

Returns an array of the strings given as argument.

Parameters

strings

the strings to put in the array.

Returns

array

the resulting array of containing the given strings.

boolean

```
bool(string)
```

Returns a Boolean with a value represented by the specified string.

The returned Boolean represents a true value if the string argument is not `null` and is equal to the string `"true"`, ignoring case.

Parameters

string

the string containing the boolean representation.

Returns

Boolean

the Boolean value represented by the string.

contains

```
contains(object, value)
```

Returns `true` if the object contains the specified value. If the object is a string, a substring is searched for the value. If the object is a collection or array, its elements are searched for the value.

Parameters

object

> the object to be searched for the presence of.

value

> the value to be searched for.

Returns

true

> if the object contains the specified value.

decodeBase64

```
decodeBase64(string)
```

Returns the base64-decoded string, or `null` if the string is not valid Base64.

Parameters

string

> The base64-encoded string to decode.

Returns

string

> The base64-decoded string.

encodeBase64

```
encodeBase64(string)
```

Returns the base64-encoded string, or `null` if the string is `null`.

Parameters

string

> The string to encode into Base64.

Returns

string

> The base64-encoded string.

fileToUrl

```
fileToUrl(file)
```

Converts a `java.io.File` into a string representation for the URL of the file or directory.

Parameters

file

> The file or directory for which to build the URL.
>
> For example, `${fileToUrl(openig.configDirectory)}/myProperties.json`.

Returns

file

> The string representation for the URL of the file or directory, or `null` if the file or directory is `null`.
>
> For example, `file:///home/gcostanza/.openig/config/myProperties.json`.

formDecodeParameterNameOrValue

```
formDecodeParameterNameOrValue(string)
```

Returns the string that results from decoding the provided form encoded parameter name or value as per `application/x-www-form-urlencoded`, which can be `null` if the input is `null`.

Parameters

string

> the parameter name or value

FORGEROCK

Returns

string

The string resulting from decoding the provided form encoded parameter name or value as per `application/x-www-form-urlencoded`.

formEncodeParameterNameOrValue

```
formEncodeParameterNameOrValue(string)
```

Returns the string that results from form encoding the provided parameter name or value as per `application/x-www-form-urlencoded`, which can be `null` if the input is `null`.

Parameters

string

the parameter name or value

Returns

string

The string resulting from form encoding the provided parameter name or value as per `application/x-www-form-urlencoded`.

indexOf

```
indexOf(string, substring)
```

Returns the index within a string of the first occurrence of a specified substring.

Parameters

string

the string in which to search for the specified substring.

substring

the value to search for within the string.

Returns

number

the index of the first instance of substring, or -1 if not found.

The index count starts from 1, not 0.

integer

```
integer(string)
```

Transforms the string parameter into an integer. If the parameter is not a valid number in radix 10, returns null.

Parameters

string

the string containing the integer representation.

Returns

integer

the integer value represented by the string.

integerWithRadix

```
integer(string, radix)
```

Uses the radix as the base for the string, and transforms the string into a base-10 integer. For example:

- `("20", 8)`: transforms `20` in base `8`, and returns `16`

- `("11", 16)` transforms `11` in base `16`, and returns `17`

If either parameter is not a valid number, returns null.

Parameters

string

a string containing the integer representation, and an integer containing the radix representation.

Returns

integer

the integer value in base-10

join

```
join(strings, separator)
```

Joins an array of strings into a single string value, with a specified separator.

Parameters

separator

the separator to place between joined elements.

strings

the array of strings to be joined.

Returns

string

the string containing the joined strings.

keyMatch

```
keyMatch(map, pattern)
```

Returns the first key found in a map that matches the specified regular expression pattern, or `null` if no such match is found.

Parameters

map

the map whose keys are to be searched.

pattern

a string containing the regular expression pattern to match.

Returns

string

the first matching key, or `null` if no match found.

length

```
length(object)
```

Returns the number of items in a collection, or the number of characters in a string.

Parameters

object

the object whose length is to be determined.

Returns

number

the length of the object, or 0 if length could not be determined.

matchingGroups

```
matchingGroups(string, pattern)
```

Returns an array of matching groups for the specified regular expression pattern applied to the specified string, or `null` if no such match is found. The first element of the array is the entire match, and each subsequent element correlates to any capture group specified within the regular expression.

Parameters

string

the string to be searched.

pattern

a string containing the regular expression pattern to match.

Returns

array

an array of matching groups, or `null` if no such match is found.

matches

```
matches(string, pattern)
```

Returns `true` if the string contains a match for the specified regular expression pattern.

Parameters

string

the string to be searched.

pattern

a string containing the regular expression pattern to find.

Returns

true

if the string contains the specified regular expression pattern.

pathToUrl

```
pathToUrl(path)
```

Converts the given path into the string representation of its URL.

Parameters

path

The path of a file or directory as a string.

For example, `${pathToUrl(system['java.io.tmpdir'])}`.

Returns

string

The string representation for the URL of the path, or `null` if the path is `null`.

For example, `file:///var/tmp`.

read

```
read(string)
```

Takes a file name as a string, interprets the content of the file with the UTF-8 character set, and returns the content of the file as a plain string.

Provides the absolute path to the file, or a path relative to the location of the Java system property `user.dir`.

Parameters

string

> The name of the file to read.

Returns

string

> The content of the file, or `null` on error.

readProperties

```
readProperties(string)
```

Takes a Java Properties file name as a `string`, and returns the content of the file as a key/value map of properties, or `null` on error (due to the file not being found, for example).

Either provide the absolute path to the file, or a path relative to the location of the Java system property `user.dir`.

For example, to get the value of the `key` property in the properties file `/path/to/my.properties`, use `${readProperties('/path/to/my.properties')['key']}`.

Parameters

string

> The name of the Java Properties file to read.

Returns

object

> The key/value map of properties or `null` on error.

readWithCharset

```
readWithCharset(string, charset)
```

Takes a file name as a string, interprets the content of the file with the specified Java character set, and returns the content of the file as a plain string.

Parameters

string

The name of the file to read.

Provides the absolute path to the file, or a path relative to the location of the Java system property `user.dir`.

charset

The name of a Java character set with which to interpret the file, as described in Class Charset.

Returns

string

The content of the file, or `null` on error.

split

```
split(string, pattern)
```

Splits the specified string into an array of substrings around matches for the specified regular expression pattern.

Parameters

string

the string to be split.

pattern

the regular expression to split substrings around.

Returns

array

the resulting array of split substrings.

toJson

```
toJson(JSON string)
```

Converts a string to a JSON structure.

Parameters

JSON string

> JSON string representing a JavaScript object.
>
> For example, the string `foo` contains the JSON object `{"email":"test@example.com"}`.

Returns

JSON structure

> JSON structure, or null if there is a parsing error.
>
> In the expression `"${toJson(foo).email}"`, the string value of `foo` is treated as JSON. The expression evaluates to `test@example.com`.

toLowerCase

```
toLowerCase(string)
```

Converts all of the characters in a string to lower case.

Parameters

string

> the string whose characters are to be converted.

Returns

string

> the string with characters converted to lower case.

toString

```
toString(object)
```

Returns the string value of an arbitrary object.

Parameters

object

the object whose string value is to be returned.

Returns

string

the string value of the object.

toUpperCase

```
toUpperCase(string)
```

Converts all of the characters in a string to upper case.

Parameters

string

the string whose characters are to be converted.

Returns

string

the string with characters converted to upper case.

trim

```
trim(string)
```

Returns a copy of a string with leading and trailing whitespace omitted.

Parameters

string

the string whose white space is to be omitted.

Returns

string

the string with leading and trailing white space omitted.

urlDecode

```
urlDecode(string)
```

Returns the URL decoding of the provided string.

This is equivalent to "formDecodeParameterNameOrValue".

Parameters

string

> The string to be URL decoded, which may be null.

Returns

string

> The URL decoding of the provided string, or null if string was null.

urlEncode

```
urlEncode(string)
```

Returns the URL encoding of the provided string.

This is equivalent to "formEncodeParameterNameOrValue".

Parameters

string

> The string to be URL encoded, which may be null.

Returns

string

> The URL encoding of the provided string, or null if string was null.

urlDecodeFragment

```
urlDecodeFragment(string)
```

Returns the string that results from decoding the provided URL encoded fragment as per RFC 3986, which can be null if the input is null.

Parameters

string

the fragment

Returns

string

The string resulting from decoding the provided URL encoded fragment as per RFC 3986.

urlDecodePathElement

```
urlDecodePathElement(string)
```

Returns the string that results from decoding the provided URL encoded path element as per RFC 3986, which can be `null` if the input is `null`.

Parameters

string

the path element

Returns

string

The string resulting from decoding the provided URL encoded path element as per RFC 3986.

urlDecodeQueryParameterNameOrValue

```
urlDecodeQueryParameterNameOrValue(string)
```

Returns the string that results from decoding the provided URL encoded query parameter name or value as per RFC 3986, which can be `null` if the input is `null`.

Parameters

string

the parameter name or value

Returns

string

> The string resulting from decoding the provided URL encoded query parameter name or value as per RFC 3986.

urlDecodeUserInfo

```
urlDecodeUserInfo(string)
```

Returns the string that results from decoding the provided URL encoded userInfo as per RFC 3986, which can be null if the input is null.

Parameters

string

> the userInfo

Returns

string

> The string resulting from decoding the provided URL encoded userInfo as per RFC 3986.

urlEncodeFragment

```
urlEncodeFragment(string)
```

Returns the string that results from URL encoding the provided fragment as per RFC 3986, which can be null if the input is null.

Parameters

string

> the fragment

Returns

string

> The string resulting from URL encoding the provided fragment as per RFC 3986.

urlEncodePathElement

```
urlEncodePathElement(string)
```

Returns the string that results from URL encoding the provided path element as per RFC 3986, which can be `null` if the input is `null`.

Parameters

string

the path element

Returns

string

The string resulting from URL encoding the provided path element as per RFC 3986.

urlEncodeQueryParameterNameOrValue

```
urlEncodeQueryParameterNameOrValue(string)
```

Returns the string that results from URL encoding the provided query parameter name or value as per RFC 3986, which can be `null` if the input is `null`.

Parameters

string

the parameter name or value

Returns

string

The string resulting from URL encoding the provided query parameter name or value as per RFC 3986.

urlEncodeUserInfo

```
urlEncodeUserInfo(string)
```

Returns the string that results from URL encoding the provided userInfo as per RFC 3986, which can be `null` if the input is `null`.

Parameters

string

> the userInfo

Returns

string

> The string resulting from URL encoding the provided userInfo as per RFC 3986.

More Information

Some functions are provided by org.forgerock.openig.el.Functions.

Other functions are provided by org.forgerock.http.util.Uris.

Name

Patterns — regular expression patterns

Description

Patterns in configuration parameters and expressions use the standard Java regular expression Pattern class. For more information on regular expressions, see Oracle's tutorial on Regular Expressions.

Pattern Templates

A regular expression pattern template expresses a transformation to be applied for a matching regular expression pattern. It may contain references to capturing groups within the match result. Each occurrence of `$g` (where *g* is an integer value) is substituted by the indexed capturing group in a match result. Capturing group zero `"$0"` denotes the entire pattern match. A dollar sign or numeral literal immediately following a capture group reference can be included as a literal in the template by preceding it with a backslash (`\`). Backslash itself must be also escaped in this manner.

More Information

Java Pattern class

Regular Expressions tutorial

Scripts

Use scripts with the following scriptable object types:

- ScriptableFilter(5), to customize flow of requests and responses

- ScriptableHandler(5), to customize creation of responses

- ScriptableThrottlingPolicy(5), to customize throttling rates

- ScriptableAccessTokenResolver(5) to customize resolution and validation of OAuth 2.0 access tokens

- `ScriptableResourceAccess` in OAuth2ResourceServerFilter(5), to customize the list of OAuth 2.0 scopes required in an OAuth 2.0 access_token

After updating a script that is used in a route, leave at least one second before processing a request. The Groovy interpreter needs time to detect and take the update into account.

> **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.

Table of Contents

Name

Scripts — customize IG behaviour with scripts

Usage

```
  "name": string,
  "type": scriptable object type,
  "config": {
    "type": string,
    "file": expression,                // Use either "file"
    "source": string or array of strings,   // or "source", but not both.
    "args": object,
    "clientHandler": Handler reference
  }
}
```

Properties

"type": *string, required*

> The Internet media type (formerly MIME type) of the script, `"application/x-groovy"` for Groovy

"file": *expression*

> Path to the file containing the script; mutually exclusive with `"source"`.
>
> Relative paths are with respect to to the base location for scripts. The base location depends on the configuration. For details, see "Installing IG" in the *Gateway Guide*.
>
> The base location for Groovy scripts is on the classpath when the scripts are executed. If 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 *openig-base*/scripts/groovy/com/example/groovy/.

"source": *string or array of strings, required if* **"file"** *is not used*

> The script as a string or array of strings; mutually exclusive with `"file"`.
>
> The following example shows the source of a script as an array of strings:

```
"source" : [
  "Response response = new Response(Status.OK)",
  "response.entity = 'foo'",
  "return response"
]
```

"args": *map, optional*

> Parameters passed from the configuration to the script.
>
> • The following example configures arguments as a map whose values can are scalars, arrays, and objects:

```
{
  "args": {
    "title": "Coffee time",
    "status": 418,
    "reason": [
      "Not Acceptable",
      "I'm a teapot",
      "Acceptable"
    ],
    "names": {
      "1": "koffie",
      "2": "kafe",
      "3": "cafe",
      "4": "kafo"
    }
  }
}
```

A script can access the args parameters in the same way as other global objects. The following example sets the response status to `I'm a teapot`:

```
response.status = Status.valueOf(418, reason[1])
```

For information about the 418 status code see RFC 7168, Section 2.3.3 *418 I'm a Teapot*.

- The following example configures arguments as strings and numbers for a ScriptableThrottlingPolicy:

```
"args" : {
  "status" : "gold",
  "rate" : 6,
  "duration": "10 seconds"
}
```

The following lines set the throttling rate to 6 requests each 10 seconds when the response status is `gold`:

```
if (attributes.rate.status == status) {
  return new ThrottlingRate(rate, duration)
}
```

- The following example configures arguments that reference a SampleFilter defined in the heap:

```
{
  "heap": [
    {
      "name": "SampleFilter",
      "type": "SampleFilter",
      "config": {
        "name": "X-Greeting",
        "value": "Hello world"
      }
    }
  ]
}
```

The following line uses an expression in the args parameter to pass SampleFilter to the script:

```
{
  "args": {
    "filter": "${heap['SampleFilter']}"
  }
}
```

The script can then reference SampleFilter as `filter`.

`"clientHandler"`, *ClientHandler reference, optional*

A Handler for making outbound HTTP requests to third-party services. In a script, `clientHandler` is wrapped within the global object `http`.

Default: The default ClientHandler.

For details, see Handlers.

For more information, see org.forgerock.http.Client.

Available Objects

The following global objects are available to scripts:

Any parameters passed as args

You can use the configuration to pass parameters to the script by specifying an args object.

The args object is a map whose values can be scalars, arrays, and objects. The args object can reference objects defined in the heap by using expressions, for example, `"${heap['ObjectName']}"`.

The values for script arguments can be defined as configuration expressions, and evaluated at configuration time.

Script arguments cannot refer to `context` and `request`, but `context` and `request` variables can be accessed directly within scripts.

Take care when naming keys in the args object. If you reuse the name of another global object, cause the script to fail and IG to return a response with HTTP status code 500 Internal Server Error.

All heap objects

The heap object configuration, described in Heap Objects(5).

openig

An implicit object that provides access to the environment when expressions are evaluated.

FORGEROCK

attributes

The attributes object provides access to a context map of arbitrary attributes, which is a mechanism for transferring transient state between components when processing a single request.

Use `session` for maintaining state between successive requests from the same logical client.

context

The processing context.

This context is the leaf of a chain of contexts. It provides access to other Context types, such as SessionContext, AttributesContext, and ClientContext, through the `context.asContext(`*`ContextClass`*`.class)` method.

contexts

a map<string, context> object. For information, see Contexts(5).

request

The HTTP request.

globals

This object is a Map that holds variables that persist across successive invocations.

http

An embedded client for making outbound HTTP requests, which is an org.forgerock.http.Client.

If a `"clientHandler"` is set in the configuration, then that Handler is used. Otherwise, the default ClientHandler configuration is used.

For details, see Handlers.

ldap

The ldap object provides an embedded LDAP client.

Use this client to perform outbound LDAP requests, such as LDAP authentication.

logger

The logger object provides access to a unique SLF4J logger instance for scripts, where the logger instance is named with the script name.

For information about logging for scripts, see "Logging for Scripts" in the *Gateway Guide*.

next

The next object refers to the next handler in the filter chain.

session

> The session object provides access to the session context, which is a mechanism for maintaining state when processing a successive requests from the same logical client or end-user.
>
> Use `attributes` for transferring transient state between components when processing a single request.

Imported Classes

The following classes are imported automatically for Groovy scripts:

- org.forgerock.http.Client

- org.forgerock.http.Filter

- org.forgerock.http.Handler

- org.forgerock.http.filter.throttling.ThrottlingRate

- org.forgerock.http.util.Uris

- org.forgerock.util.AsyncFunction

- org.forgerock.util.Function

- org.forgerock.util.promise.NeverThrowsException

- org.forgerock.util.promise.Promise

- org.forgerock.services.context.Context

- org.forgerock.http.protocol.*

- org.forgerock.http.oauth2.AccessTokenInfo

- org.forgerock.json.JsonValue, and all its static methods, including `json(Object)`, `array(Object...)`, `object(fields...)`, and `field(String, Object)`

More Information

- ScriptableFilter(5), org.forgerock.openig.filter.ScriptableFilter, and org.forgerock.http.Filter

- ScriptableHandler(5), org.forgerock.openig.handler.ScriptableHandler, and org.forgerock.http.Handler

- ScriptableThrottlingPolicy(5), org.forgerock.openig.filter.throttling.ScriptableThrottlingPolicy.Heaplet, and org.forgerock.http.filter.throttling.ThrottlingPolicy

FORGEROCK

- `ScriptableResourceAccess` in OAuth2ResourceServerFilter(5), org.forgerock.openig.filter.oauth2.ScriptableResourceAccess, and org.forgerock.http.oauth2.ResourceAccess

- `ScriptableAccessTokenResolver` in OAuth2ResourceServerFilter(5), org.forgerock.openig.filter.oauth2.ScriptableAccessTokenResolver, and org.forgerock.http.oauth2.AccessTokenResolver

Properties

Configuration parameters can be declared as properties in the IG configuration or in an external JSON file.

Table of Contents

Name

Properties — declare configuration parameters as property variables in a route

Description

Configuration parameters, such as host names, port numbers, and directories, can be declared as property variables in the IG 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 IG in different environments without changing a single line in your route configuration.

Usage

Simple Property Configured Inline

```
{
    "properties": {
        "<variable name>": "valid JSON value"
    }
}
```

Group Property Configured Inline

```
{
    "properties": {
        "<group name>": {
            ["<variable name>": "valid JSON value", ... ]
        }
    }
}
```

Properties Configured in One or More External Files

```
{
    "properties": {
        "$location": expression
    }
}
```

In this example, `description1` and `description2` prefix the variable names contained in the external file.

```
{
    "properties": {
        "description1" {
            "$location": expression
        {
        "description2" {
            "$location": expression
        {
    }
}
```

Properties

"<variable name>": *string*

The name of a variable to use in the IG configuration. The variable can be used in expressions in routes or in `config.json` to assign the value of a configuration parameter.

The value assigned to the variable can be any valid JSON value: string, number, boolean, array, object, or null.

- In the following example from `config.json`, the URL of an application is declared as a property variable named `appLocation`. The variable is then used by the `baseURI` parameter of the handler, and can be used again in other routes in the configuration.

```
{
    "properties": {
        "appLocation": "http://app.example.com:8081"
    },
    "handler": {
        "type": "Router",
        "baseURI": "${appLocation}",
        "capture": "all"
    }
}
```

- The following example adds the property variable `ports` to define an array of port numbers used by the configuration. The `ports` variable is referenced in the `appLocation` variable, and is resolved at runtime with the value in the ports array:

```
{
    "properties" : {
        "ports": [8080, 8081, 8088],
        "appLocation": "http://app.example.com:${ports[1]}"
    },
    "handler": {
        "type": "Router",
        "baseURI": "${appLocation}",
        "capture": "all"
    }
}
```

"\<group name>": *string, required*

The name of a group of variables to use in the IG configuration. The group name and variable name are combined using dot notation in an expression.

In the following example from `config.json`, the property group `directories` contains two variables that define the location of files:

```
{
    "properties" : {
        "directories": {
            "config": "${openig.configDirectory.path}",
            "auditlog": "/tmp/logs"
        }
    }
}
```

The group name and variable name are combined using dot notation in the following example to define the directory where the audit log is stored:

```
{
    "type": "AuditService",
    "config": {
        "event-handlers": [
            {
            "class": "org.forgerock.audit.handlers.csv.CsvAuditEventHandler",
            "config": {
                "name": "csv",
                "logDirectory": "${directories.auditlog}",
                . . .
```

"$location": *expression, required*

The location and name of one or more JSON files where property variables are configured.

Files must be `.json` files, and contain property variables with a key/value format, where the key cannot contain the period (`.`) separator.

For example, this file is correct:

```
{
  "openamLocation": "http://openam.example.com:8088/openam/",
  "portNumber": 8081
}
```

This file would cause an error:

```
{
  "openam.location": "http://openam.example.com:8088/openam/",
  "port.number": 8081
}
```

Property Variables Configured In One File

In the following example, the location of the file that contains the property variables is defined as an expression:

```
{
    "properties" : {
        "$location": "${fileToUrl(openig.configDirectory)}/myProperties.json"
    }
}
```

In the following example, the location of the file that contains the property variables is defined as a string:

```
{
    "properties" : {
        "$location": "file:///Users/user-id/.openig/config/myProperties.json"
    }
}
```

The file location can be defined as any real URL.

The file `myProperties.json` contains the base URL of an AM service and the port number of an application.

```
{
    "openamLocation": "http://openam.example.com:8088/openam/",
    "appPortNumber": 8081
}
```

Properties Variables Configured In Multiple Files

In the following example, the property variables are contained in two files, defined as a set of strings:

```
{
    "properties": {
        "urls": {
            "$location": "file://path-to-file/myUrlProperties.json"
        },
        "ports": {
            "$location": "file://path-to-file/myPortProperties.json"
        }
    }
}
```

The file `myUrlProperties.json` contains the base URL of the sample application:

```
{
    "appUrl": "http://app.example.com"
}
```

The file `myPortProperties.json` contains the port number of an application:

```
{
    "appPort": 8081
}
```

The base config file, `config.json`, can use the properties as follows:

```
{
  "properties": {
    "urls": {
        "$location": "file:///Users/user-id/.openig/config/myUrlProperties.json"
    },
    "ports": {
        "$location": "file:///Users/user-id/.openig/config/myPortProperties.json"
    }
  },
  "handler": {
    "type": "Router",
    "name": "_router",
    "baseURI": "${urls.appUrl}:${ports.appPort}",
    . . .
```

Requests, Responses, and Contexts

This part of the reference describes the IG object model. The top-level objects are request, response, and contexts.

Table of Contents

Name

AttributesContext — context for request attributes

Description

Provides a map for request attributes. When IG processes a single request, it injects transient state information about the request into this context. Attributes stored when processing one request are not accessible when processing a subsequent request.

IG automatically provides access to the `attributes` field through the `attributes` bindings in expressions. For example, to access a username with an expression, use `${attributes.credentials.username}` instead of `${contexts.attributes.attributes.credentials.username}`

Use SessionContext(5) to maintain state between successive requests from the same logical client.

Properties

`"attributes"`: **map**

Map of information conveyed between filters and handlers. Cannot be null.

More Information

org.forgerock.services.context.AttributesContext

Name

CapturedUserPasswordContext — context for the decrypted AM password of the current user

Description

Provides the decrypted AM password of the current user. When the CapturedUserPasswordFilter(5) processes a request, it injects the decrypted password from AM into this context.

Properties

The context name is `capturedPassword`. The context has the following properties:

`"raw"`: **byte[]**

The decrypted password as bytes.

`"value"`: **string**

The decrypted password as a UTF-8 string.

More Information

org.forgerock.openig.openam.CapturedUserPasswordContext

Name

ClientContext — context for the HTTP client information

Description

Information about the client sending a request. When IG receives a request, it injects information about the client sending the request into this context.

Properties

"certificates": array

List of X.509 certificates presented by the client

If the client does not present any certificates, IG returns an empty list.

Never `null`.

"isExternal": boolean

True if the client connection is external.

"isSecure": boolean

True if the client connection is secure.

"localAddress": string

The IP address of the interface that received the request

"localPort": number

The port of the interface that received the request

"remoteAddress": string

The IP address of the client (or the last proxy) that sent the request

"remotePort": number

The source port of the client (or the last proxy) that sent the request

"remoteUser": string

The login of the user making the request, or `null` if unknown

This is likely to be `null` unless you have deployed IG with a non-default deployment descriptor that secures the IG web application.

`"userAgent"`: string

> The value of the User-Agent HTTP header in the request if any, otherwise `null`

More Information

org.forgerock.services.context.ClientContext

Name
Contexts — HTTP request contexts

Description

The root object for request context information.

Contexts is a map of available contexts, which implement the Context interface. The contexts map's keys are strings and the values are context objects. A context holds type-safe information useful for processing requests and responses. The `contexts` map is populated dynamically when creating bindings for evaluation of expressions and scripts.

All context objects use their version of the following properties:

"*context-Name*": string

    Name of the context.

"*context-ID*": string

    Read-only string uniquely identifying the context object.

"*context-rootContext*": boolean

    True if the context object is a RootContext (has no parent).

"*context-Parent*": Context object

    Parent of this context object.

Properties

The contexts object can provide access to the following contexts for each request:

- AttributesContext(5)
- ClientContext(5)
- SessionContext(5)
- UriRouterContext(5)
- TransactionIdContext(5)

The contexts object can provide access to the following contexts when related filters are used:

- CapturedUserPasswordContext(5)
- CdSsoContext(5)

- CdSsoFailureContext(5)

- OAuth2Context(5)

- PolicyDecisionContext(5)

- SessionInfoContext(5)

- SsoTokenContext(5)

- StsContext(5)

More Information

org.forgerock.services.context.Context

FORGEROCK

Name

CdSsoContext — context for cross-domain SSO properties

Description

Provides the cross-domain SSO properties for the CDSSO token, the user ID of the session, and the full claims set. When the CrossDomainSingleSignOnFilter(5) processes a request, it injects the information in this context.

Properties

The context is named `cdsso`, and is accessible at `${contexts.cdsso}`. The context has the following properties:

`"claimsSet"`: *JwtClaimsSet object, required*

Full claims set for the identity of the authenticated user. Cannot be null.

`"cookieInfo"`: *object*

Configuration data for the CDSSO authentication cookie, with the following attributes:

- `name`: Cookie name (string)
- `domain`: Cookie domain (optional string)
- `path`: Cookie path (string)

None of the attributes can be null.

`redirectEndpoint"`: *string*

Redirect endpoint URI configured for communication with AM. Cannot be null.

`"sessionUid"`: *string*

Universal session ID. Cannot be null.

`"token"`: *string*

Value of the CDSSO token. Cannot be null.

More Information

org.forgerock.openig.openam.CdSsoContext

Name

CdSsoFailureContext — context for cross-domain SSO authentication error details

Description

Contains the error details for any error that occured during cross-domain SSO authentication. When the CrossDomainSingleSignOnFilter(5) processes a request, should an error occur that prevents authentication, the error details are captured in this context.

Properties

The context is named `cdssoFailure`, and is accessible at `${contexts.cdssoFailure}`. The context has the following properties:

`"error"`: *The error (string)*

> The error that occurred during authentication. Cannot be null.

`"description"`: *Error description (string)*

> A description of the error that occurred during authentication. Cannot be null.

`"throwable"`: *Throwable*

> Any `Throwable` associated with the error that occured during authentication. Can be null.

More Information

org.forgerock.openig.openam.CdSsoFailureContext

Name

JwtBuilderContext — context for data from JwtBuilderFilter

Description

When the JwtBuilderFilter(5) processes a request, it stores provided data in this context. This context returns the JWT as string for downstream use.

Properties

The context is named `jwtBuilder`, and is accessible at `${contexts.jwtBuilder}`. The context has the following properties:

`"value"`: *string*

The base64url encoded UTF-8 parts of the JWT, containing name-value pairs of data. Cannot be null.

More Information

org.forgerock.openig.filter.JwtBuilderFilter

org.forgerock.openig.filter.JwtBuilderContext

Name

OAuth2Context — context for OAuth 2.0 access tokens

Description

Provides OAuth 2.0 access tokens. When the OAuth2ResourceServerFilter(5) processes a request, it injects the access token into this context.

Properties

The context name is `oauth2`, and the context has the following property:

"accessToken": AccessTokenInfo

>   The AccessTokenInfo is built from the following properties:
>
>   **"info": (map<string>, object)**
>
>   >   Raw JSON as a map.
>
>   **"token": string**
>
>   >   Access token identifier issued from the authorization server.
>
>   **"scopes": space separated string**
>
>   >   Scopes associated to this token.
>
>   **"expiresAt": long**
>
>   >   Timestamp (in milliseconds since epoch) when the token expires.

More Information

org.forgerock.http.oauth2.OAuth2Context

org.forgerock.http.oauth2.AccessTokenInfo

Name

PolicyDecisionContext — context for attributes and advices returned by AM policy decisions

Description

Provides attributes and advices returned by AM policy decisions. When the PolicyEnforcementFilter(5) processes a request, it injects the attributes and advices into this context.

Properties

The context name is `policyDecision`. The context has the following properties:

`"attributes"`: unmodifiable map

The map of attributes provided in the policy decision. Can be empty, but not null.

`"jsonAttributes"`: JsonValue

The map of attributes provided in the policy decision. Cannot be null.

`"advices"`: map

The map of advices provided in the policy decision. Can be empty, but not null.

`"jsonAdvices"`: JsonValue

The map of advices provided in the policy decision. Cannot be null.

More Information

org.forgerock.openig.openam.PolicyDecisionContext.html

Name

Request — HTTP request

Description

An HTTP request message.

Properties

"method": *string*

> The method to be performed on the resource. Example: `"GET"`.

"uri": *object*

> The fully-qualified URI of the resource being accessed. Example: `"http://www.example.com/resource .txt"`.
>
> See also URI(5).

"version": *string*

> Protocol version. Example: `"HTTP/1.1"`.

"headers": *object*

> Exposes message header fields as name-value pairs, where name is header name and value is an array of header values.

"cookies": *object*

> Exposes incoming request cookies as name-value pairs, where name is cookie name and value is an array of string cookie values.

"form": *object*

> Exposes query parameters and/or `application/x-www-form-urlencoded` entity as name-value pairs, where name is the field name and value is an array of string values.

"entity": *object*

> The message entity body.
>
> Methods are provided for accessing the entity as byte, string, or JSON content. For information, see Entity.

More Information

org.forgerock.http.protocol.Request

org.forgerock.http.protocol.Entity

Name

Response — HTTP response

Description

An HTTP response message.

Properties

"cause": *Exception object*

> The cause of an error if the status code is in the range 4xx-5xx. Possibly null.

"status": *Status object*

> The response status.
>
> For details, see Status(5).

"version": *string*

> Protocol version. Example: `"HTTP/1.1"`.

"headers": *object*

> Exposes message header fields as name-value pairs, where name is header name and value is an array of header values.

"entity": *object*

> The message entity body.
>
> Methods are provided for accessing the entity as byte, string, or JSON content. For information, see Entity.

More Information

org.forgerock.http.protocol.Response

org.forgerock.http.protocol.Entity

Name

SessionContext — context for the HTTP session

Description

Provides access to the HTTP session. When IG processes a request, an HTTP session is automatically created. Subsequent requests that resend the same session cookie share the session content.

When processing a single request, use AttributesContext(5) to transfer transient state between components and prevent IG from creating additional sessions.

IG automatically provides access to the `session` field through the `session` bindings in expressions. For example, to access a username with an expression, use `${session.username}` instead of `${contexts.session.session.username}`

Properties

`"session"`: **map**

Provides access to the HTTP session. The session is a map with attributes that are name-value pairs of the format `Map<String, Object>`.

For web container sessions, any object can be stored in the session. For JwtSession, only JSON compatible types (such as primitive JSON structures, lists, arrays, and maps) can be stored in the session.

By default, sessions are web container sessions.

More Information

org.forgerock.http.session.SessionContext

Name

SessionInfoContext — context for AM session information and properties

Description

Provides AM session information and properties. When the SessionInfoFilter(5) processes a request, it injects info and properties from the AM session into this context.

Properties

The context name is `amSession`. The context has the following properties:

`"asJsonValue()"`: JsonValue

    Raw JSON.

`"latestAccessTime"`: instant

    The timestamp of when the session was last used. Can be null if the DN is not resident on the SSO token, or if the time cannot be obtained from the session.

`"maxIdleExpirationTime"`: instant

    The timestamp of when the session would time out for inactivity. Can be null if the DN is not resident on the SSO token, or if the time cannot be obtained from the session.

`"maxSessionExpirationTime"`: instant

    The timestamp of when the session would time out regardless of activity. Can be null if the DN is not resident on the SSO token, or if the time cannot be obtained from the session.

`"properties"`: map

    The read-only map of properties bound to the session. Can be empty, but not null.

    Only session properties that are whitelisted in AM are available. For information, see Session Property Whitelist Service .

`"realm"`: string

    The realm as specified by AM, in a user-friendly slash (/) separated format. Can be null if the DN is not resident on the SSO token.

`"sessionHandle"`: string

    The handle to use for logging out of the session. Can be null if the handle is not available for the session.

`"universalId"`: string

The DN that AM uses to uniquely identify the user. Can be null if it cannot be obtained from the SSO token.

`"username"`: string

A user-friendly version of the username. Can be null if the DN is not resident on the SSO token, or empty if it cannot be obtained from the DN.

More Information

org.forgerock.openig.openam.SessionInfoContext

Name

SsoTokenContext — context for SSO tokens and their validation information

Description

Provides SSO tokens and their validation information. When the SingleSignOnFilter(5) processes a request, it injects the value of the SSO token and additional information in this context.

Properties

The context name is `ssoToken`. The context has the following properties:

`"info"`: map

Information associated with the SSO token, such as `realm` or `uid`. Cannot be null.

`"loginEndpoint"`: url

URL for the login endpoint, evaluated from the configuration of SingleSignOnFilter.

`"value"`: string

The value of the SSO token. Cannot be null.

More Information

org.forgerock.openig.openam.SsoTokenContext

Name

Status — HTTP response status

Description

Represents an HTTP response status. For details, see *RFC 7231: HTTP/1.1 Semantics and Content*, Section 6.1. Overview of Status Codes.

Properties

"code": *integer*

Three-digit integer reflecting the HTTP status code.

"family": *enum*

Family Enum value representing the class of response that corresponds to the code:

Family.INFORMATIONAL

Status code reflects a provisional, informational response: 1xx.

Family.SUCCESSFUL

The server received, understood, accepted and processed the request successfully. Status code: 2xx.

Family.REDIRECTION

Status code indicates that the client must take additional action to complete the request: 3xx.

Family.CLIENT_ERROR

Status code reflects a client error: 4xx.

Family.SERVER_ERROR

Status code indicates a server-side error: 5xx.

Family.UNKNOWN

Status code does not belong to one of the known families: 600+.

"reasonPhrase": *string*

The human-readable reason-phrase corresponding to the status code.

For details, see *RFC 7231: HTTP/1.1 Semantics and Content*, Section 6.1. Overview of Status Codes.

"isClientError": *boolean*

True if Family.CLIENT_ERROR.

"isInformational": *boolean*

True if Family.INFORMATIONAL.

"isRedirection": *boolean*

True if Family.REDIRECTION.

"isServerError": *boolean*

True if Family.SERVER_ERROR.

"isSuccessful": *boolean*

True if Family.SUCCESSFUL.

More Information

org.forgerock.http.protocol.Status

Name

StsContext — context for the result of a token transformation

Description

Provides the result of a token transformation. When the TokenTransformationFilter(5) processes a request, it injects the result into this context.

Properties

The context name is `sts`. The context has the following properties:

"issuedToken": string

> The result of the token transformation.

More Information

org.forgerock.openig.openam.StsContext

Name

TransactionIdContext — context for the transaction ID of a request

Description

The transaction ID of a request. When IG receives a request, it injects the transaction ID into this context.

Properties

The context name is `transactionId`. The context has the following properties:

`"transactionId"`: TransactionId

> The ID of the transaction.

More Information

org.forgerock.services.TransactionIdContext

org.forgerock.services.context.TransactionIdContext

Name

URI — Uniform Resource Identifier

Description

Represents a Uniform Resource Identifier (URI) reference.

Properties

"scheme": *string*

> The scheme component of the URI, or `null` if the scheme is undefined.

"authority": *string*

> The decoded authority component of the URI, or `null` if the authority is undefined.
>
> Use "rawAuthority" to access the raw (encoded) component.

"userInfo": *string*

> The decoded user-information component of the URI, or `null` if the user information is undefined.
>
> Use "rawUserInfo" to access the raw (encoded) component.

"host": *string*

> The host component of the URI, or `null` if the host is undefined.

"port": *number*

> The port component of the URI, or `null` if the port is undefined.

"path": *string*

> The decoded path component of the URI, or `null` if the path is undefined.
>
> Use "rawPath" to access the raw (encoded) component.

"query": *string*

> The decoded query component of the URI, or `null` if the query is undefined.

> **Note**
>
> The query key and value is decoded. However, because a query value can be encoded more than once in a redirect chain, even though it is decoded it can contain unsafe ASCII characters.

> Use "rawQuery" to access the raw (encoded) component.

`"fragment"`: *string*

> The decoded fragment component of the URI, or `null` if the fragment is undefined.
>
> Use "rawFragment" to access the raw (encoded) component.

More Information

org.forgerock.http.MutableUri

Name

UriRouterContext — context for HTTP request routing

Description

Provides routing information associated with a request. When IG routes a request, it injects information about the routing into this context.

Properties

The context name is `router`. The context has the following properties:

"baseUri": string

> The portion of the request URI which has been routed so far.

"matchedUri": string

> The portion of the request URI that matched the URI template.

"originalUri": URI

> The original target URI for the request, as received by the web container.

> The value of this field is read-only.

"remainingUri": string

> The portion of the request URI that is remaining to be matched.

"uriTemplateVariables": map

> An unmodifiable map, where the keys and values are strings. The map contains the parsed URI template variables keyed on the URI template variable name.

More Information

org.forgerock.http.routing.UriRouterContext

Name

UserProfileContext — context for user profile information from AM

Description

When the UserProfileFilter(5) processes a request, it injects the user profile information into this context. This context provides raw JSON representation, and convenience accessors that map commonly used LDAP field names to a context names.

Properties

The context name is `userProfile`, available at `${contexts.userProfile}` with the following properties:

`"username"`: string

User-friendly version of the username. This field is always fetched. If the underlying data store doesn't include `username`, this field is null.

Example of use: `${contexts.userProfile.username}`

`"realm"`: string

Realm as specified by AM, in a user-friendly slash (/) separated format. Can be null.

Example of use: `${contexts.userProfile.realm}`

`"distinguishedName"`: string

Distinguished name of the user. Can be null.

Example of use: `${contexts.userProfile.distinguishedName}`

`"commonName"`: string

Common name of the user. Can be null.

Example of use: `${contexts.userProfile.commonName}`

`"rawInfo"`: (map<string>, object)

Unmodifiable map of the user profile information.

This context contains the object structure of the AM user profile. Any individual field can be retrieved from the map. Depending on the requested fields, the context can be empty or values can be null.

Examples of use: `${contexts.userProfile.rawInfo}`, `${contexts.userProfile.rawInfo.username}`, `${contexts.userProfile.rawInfo.employeeNumber[0]}`.

`"asJsonValue()"`: JsonValue

Raw JSON of the user profile information.

Example of use: `${contexts.userProfile.asJsonValue()}`

More Information

org.forgerock.openig.openam.UserProfileContext

UserProfileFilter(5)

Access Token Resolvers

This section describes the objects available to resolve access token against an authorization server. For information about where these objects are used, see OAuth2ResourceServerFilter(5).

Table of Contents

Name

TokenIntrospectionAccessTokenResolver — resolve access tokens at introspection endpoint

Description

In OAuth2ResourceServerFilter, use the token introspection endpoint, `/oauth2/introspect`, to resolve access tokens and retrieve metadata about the token. The endpoint typically returns the time until the token expires, the OAuth 2.0 *scopes* associated with the token, and potentially other information.

The introspection endpoint is defined as a standard method for resolving access tokens, in RFC-7662, *OAuth 2.0 Token Introspection*.

Usage

Use this resolver with the `accessTokenResolver` property of OAuth2ResourceServerFilter.

```
"accessTokenResolver": {
  "type": "TokenIntrospectionAccessTokenResolver",
  "config": {
    "amService": AmService reference,   // Use either "amService"
    "endpoint": URI string,             // or "endpoint", but not both.
    "providerHandler": Handler reference
  }
}
```

Properties

`"amService"`: *AmService reference, required if `"endpoint"` is not configured*

The AmService heap object to use for the token introspection endpoint. The endpoint is extrapolated from the `url` property of the AmService.

When the authorization server is AM, use this property to define the token introspection endpoint.

If `amService` is configured, it takes precedence over `endpoint`.

See also, AmService(5).

`"endpoint"`: *URI string, required if `"amService"` is not configured*

The URI for the token introspection endpoint. Use `/oauth2/introspect`.

When the authorization server is not AM, use this property to define the token introspection endpoint.

If `amService` is configured, it takes precedence over `endpoint`.

`"providerHandler"`: *Handler reference, optional*

Invoke this HTTP client handler to send token info requests.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

Default: ForgeRockClientHandler

If you use the AM token introspection endpoint, this handler can be a `Chain` containing a `HeaderFilter` to add the authorization to the request header, as in the following example:

```
"providerHandler": {
  "type": "Chain",
  "config": {
    "filters": [
      {
        "type": "HeaderFilter",
        "config": {
          "messageType": "request",
          "add": {
            "Authorization": [ "Basic ${encodeBase64('<client_id>:<client_secret>')}" ]
          }
        }
      }
    ],
    "handler": "ForgeRockClientHandler"
  }
}
```

Example

For an example route that uses the token introspection endpoint, see "Validating Access-Tokens Through the Introspection Endpoint" in the *Gateway Guide*.

More Information

org.forgerock.http.oauth2.resolver.TokenIntrospectionAccessTokenResolver

OAuth2ResourceServerFilter(5)

Name

StatelessAccessTokenResolver — resolve stateless access_tokens without referring to AM

Description

Locally resolve and validate stateless access_tokens issued by AM, without referring to AM.

AM can be configured to secure access_tokens by signing or encrypting. The StatelessAccessTokenResolver must be configured for signature or encryption according to the AM configuration.

Supported with OpenAM 13.5, and AM 5 and later versions.

Usage

Use this resolver with the `accessTokenResolver` property of OAuth2ResourceServerFilter.

```
"accessTokenResolver": {
  "type": "StatelessAccessTokenResolver",
  "config": {
    "issuer": URI string,
    "secretsProvider": SecretsProvider reference,
    "verificationSecretId": configuration expression<secret-id>, // Use "verificationSecretId" or
    "decryptionSecretId": configuration expression<secret-id>    // "decryptionSecretId", but not both
  }
}
```

Properties

`"issuer"`: *Issuer reference, required*

> URI of the AM instance responsible for issuing access_tokens.

`"secretsProvider"`: *SecretsProvider reference, optional*

> The SecretsProvider(5) to query when resolving an access_token. SecretsProvider resolves access_tokens with the specified secret stores only. If SecretsProvider is not configured, the global secrets service is used to resolve access_tokens.
>
> Provide either the name of a SecretsProvider object defined in the heap, or specify a SecretsProvider object inline.
>
> Default: No SecretsProvider.

`"verificationSecretId"`: *configuration expression<secret-id>, required if AM secures access_tokens with a signature*

> The secret ID for the secret used to verify the signature, to validate the message authenticity.

Use one of the following values:

- `null`: No signature verification is required.

- A `kid` as a string: Signature verification is required with the provided `kid`. The StatelessAccessTokenResolver searches for the matching `kid` in the SecretsProvider or global secrets service.

Use either `verificationSecretId` or `decryptionSecretId`, according to the configuration of the token provider in AM. If AM is configured to sign **and** encrypt tokens, encryption takes precedence over signing.

For information about supported formats for `secret-id`, see secret-id.

"decryptionSecretId": *configuration expression<secret-id>, required if AM secures access_tokens with encryption*

The secret ID for the secret used to decrypt the JWT, for confidentiality.

Use either `verificationSecretId` or `decryptionSecretId`, according to the configuration of the token provider in AM. If AM is configured to sign **and** encrypt the token, encryption takes precedence over signing.

For information about supported formats for `secret-id`, see secret-id.

Examples

For examples of how to set up and use StatelessAccessTokenResolver to resolve signed and encrypted access_tokens, see "Validating Signed Access_Tokens With the StatelessAccessTokenResolver" in the *Gateway Guide* and "Validating Encrypted Access_Tokens With the StatelessAccessTokenResolver" in the *Gateway Guide*.

More Information

org.forgerock.openig.filter.oauth2.StatelessAccessTokenResolver

OAuth2ResourceServerFilter(5)

Name
OpenAmAccessTokenResolver — resolve access tokens at token info endpoint

Description

In OAuth2ResourceServerFilter, use the AM token info endpoint, `/oauth2/tokeninfo`, to resolve access tokens and retrieve information. The endpoint typically returns the time until the token expires, the OAuth 2.0 *scopes* associated with the token, and potentially other information.

Usage

Use this resolver with the `accessTokenResolver` property of OAuth2ResourceServerFilter.

```
"accessTokenResolver": {
  "type": "OpenAmAccessTokenResolver",
  "config": {
    "amService": AmService reference,
    "providerHandler": Handler reference
  }
}
```

Properties

`"amService"`: *AmService reference, required*

The AmService heap object to use for the token info endpoint. The endpoint is extrapolated from the `url` property of the AmService.

See also, AmService(5).

`"providerHandler"`: *Handler reference, optional*

Invoke this HTTP client handler to send token info requests.

Provide either the name of a Handler object defined in the heap, or an inline Handler configuration object.

> **Tip**
>
> To facilitate auditing, configure this handler with a `ForgeRockClientHandler`, which sends a ForgeRock Common Audit transaction ID when it communicates with protected applications.
>
> Alternatively, configure this handler as a chain containing a `TransactionIdOutboundFilter`, as in the following configuration:

```
providerHandler : {
  "type" : "Chain",
  "config" : {
    "handler": "MySecureClientHandler",
    "filters": [ "TransactionIdOutboundFilter" ]
  }
}
```

Default: ForgeRockClientHandler

Example

For an example route that uses this access token resolver, see "Validating Access-Tokens Through the Token Info Endpoint" in the *Gateway Guide*.

More Information

org.forgerock.http.oauth2.resolver.OpenAmAccessTokenResolver

OAuth2ResourceServerFilter(5)

Name

ConfirmationKeyVerifierAccessTokenResolver — verify certificate-bound OAuth 2.0 bearer tokens presented by clients

Description

Supported with AM 6.5.1 and later versions.

In OAuth2ResourceServerFilter, use the ConfirmationKeyVerifierAccessTokenResolver to verify that certificate-bound OAuth 2.0 bearer tokens presented by clients use the same mTLS-authenticated HTTP connection.

When a client obtains an access_token from AM by using mTLS, AM can optionally use a confirmation key to bind the access_token to a certificate. When the client connects to IG using that certificate, the ConfirmationKeyVerifierAccessTokenResolver verifies that the confirmation key corresponds to the certificate.

This proof-of-possession interaction ensures that only the client in possession of the key corresponding to the certificate can use the access_token to access protected resources.

To use the ConfirmationKeyVerifierAccessTokenResolver, the following configuration is required in AM:

- OAuth 2.0 clients must be registered using an X.509 certificate, that is self-signed or signed in public key infrastructure (PKI)

- The AM client authentication method must be `self_signed_client_auth` or `tls_client_auth`.

- AM must be configured to bind a confirmation key to each client certificate.

The ConfirmationKeyVerifierAccessTokenResolver delegates the token resolution to a specified AccessTokenResolver, which retrieves the token information. The ConfirmationKeyVerifierAccessTokenResolver verifies the confirmation keys bound to the access_token, and then acts as follows:

- If there is no confirmation key, pass the request down the chain.

- If the confirmation key matches the client certificate, pass the request down the chain.

- If the confirmation key doesn't match the client certificate, throw an error.

- If the confirmation key method is not supported by IG, throw an error.

For an example that uses the ConfirmationKeyVerifierAccessTokenResolver, see "Validating Access_Tokens Obtained Through mTLS" in the *Gateway Guide*.

For information about issuing certificate-bound OAuth 2.0 access_tokens, see AM's Certificate-Bound Proof-of-Possession. For information about authenticating an OAuth 2.0 client using mTLS certificates, see AM's Authenticating Clients Using Mutual TLS.

Usage

Use this resolver with the `accessTokenResolver` property of OAuth2ResourceServerFilter(5).

```
"accessTokenResolver": {
  "type": "ConfirmationKeyVerifierAccessTokenResolver",
  "config": {
    "delegate": accessTokenResolver reference
  }
}
```

Properties

"delegate": *accessTokenResolver reference, required*

The access token resolver to use for resolving access_tokens. Use any access token resolver described in Access Token Resolvers.

Examples

For an example that uses the ConfirmationKeyVerifierAccessTokenResolver with the following route, see "Validating Access_Tokens Obtained Through mTLS" in the *Gateway Guide*.

More Information

org.forgerock.openig.filter.oauth2.ConfirmationKeyVerifierAccessTokenResolver

OAuth 2.0 Mutual TLS Client Authentication and Certificate Bound Access Tokens

OAuth2ResourceServerFilter(5)

Name

ScriptableAccessTokenResolver — resolve access tokens by using a Groovy script

Description

In OAuth2ResourceServerFilter, use a Groovy script to resolve access tokens against an authorization server.

Receive a string representing an access token and use a Groovy script to create an instance or promise of `org.forgerock.http.oauth2.AccessTokenInfo`.

Usage

Use this resolver with the `accessTokenResolver` property of OAuth2ResourceServerFilter.

```
"accessTokenResolver": {
  "type": "ScriptableAccessTokenResolver",
  "config": {
    "type": string,
    "file": expression,                 // Use either "file"
    "source": string or array of strings,   // or "source", but not both.
    "args": object,
    "clientHandler": Handler reference
  }
}
```

Properties

For information about properties for ScriptableAccessTokenResolver, see Scripts(5).

More Information

org.forgerock.openig.filter.oauth2.ScriptableAccessTokenResolver

OAuth2ResourceServerFilter(5)

Secrets

IG uses the ForgeRock Commons Secrets Service to manage secrets, such as passwords and cryptographic keys.

A secrets service instance is automatically created in each route, in `config.json`, and in `admin.json`. Each secrets service instance can be configured independently.

When a secrets store is not declared in a route or its parent routes, in `config.json`, or in `admin.json`, a default SystemAndEnvSecretStore is available in the local secrets service. When a secrets stores is declared, the default is not installed in the local secrets service.

Each secrets service object configures an inline array of one or more secret stores:

```
{
  "secrets": {
    "stores": [ inline secret store declaration, ... ]
  }
}
```

The following example configures two stores:

```
{
  "secrets": {
    "stores": [
      {
        "type": "FileSystemSecretStore",
        "config": {
          "directory": "/path/to/secrets",
          "format": "BASE64"
        }
      },
      {
        "type": "SystemAndEnvSecretStore",
        "config": {
          "format": "PLAIN"
        }
      }
    ]
  }
}
```

When secrets in a route are queried, their values are searched in the following order until they are found:

- Locally in the route, starting with the first secret store in the list, up to the last.

- In ascending parent routes, starting with the first secret store in each list, up to the last.

- In `config.json`, starting with the first secret store in the list, up to the last.

- If a secrets service instance is not configured in `config.json` the secret is queried in a default SystemAndEnvSecretStore(5), and a base64-encoded value is expected.

• If a secret is not resolved, the route fails to load and an error is produced.

Secrets stores defined in `admin.json` can be accessed only by heap objects in `admin.json`.

Table of Contents

Name

FileSystemSecretStore — manage a store of secrets from a directory

Description

Manage a store of secrets held in files in a specified directory. Secrets are read lazily from the filesystem.

For each secret ID in the IG configuration, this object checks the specified directory for a file whose name matches the ID followed by the specified suffix. If there is a match, the content of the file is stored as the value of the secret, and then cached indefinitely.

This object checks the specified directory, but not its subdirectories. Each file must contain only one secret, in the `format` declared in the configuration.

Usage

```
{
  "name": string,
  "type": "FileSystemSecretStore",
  "config": {
    "directory": configuration expression<string>,
    "format": configuration expression<enumeration>,
    "suffix": configuration expression<string>
  }
}
```

Properties

`"directory"`: *configuration expression<string>, required*

> File path to a directory containing secret files.

`"format"`: *configuration expression<enumeration>, optional*

> Format in which the secret is stored. Use one of the following values:
>
> - `BASE64`: Base64-encoded
>
> - `PLAIN`: Plain text
>
> Default: `BASE64`

`"suffix"`: *configuration expression<string>, optional*

> File suffix.
>
> When set, the FileSystemSecretStore will append that suffix to the secret ID and try to find a file with the mapped name.
>
> Default: None

Example

In the following example, an AmService is configured to communicate WebSocket notifications from AM to IG. An AM Java agent is configured with a password, and IG uses the FileSystemSecretStore to retrieve the password from a file in the specified directory. When the route is deployed, a WebSocket notification is written to the IG log.

To Retrieve a Secret From a File

Before you start this tutorial:

- Prepare IG as described in "*First Steps*" in the *Getting Started Guide*.

- Install and configure AM on http://openam.example.com:8088/openam, with the default configuration. If you use a different configuration, substitute in the tutorial accordingly.

1. Set up AM:

    a. (For AM 6.5.3 and earlier versions) Select ▣ Identities > demo, and set the demo user password to `Ch4ng31t`.

    b. (For AM 6.5.3 and later versions) Select Services, and add a Validation Service with the following URLs for IG:

        - `http://openig.example.com:8080/*`

        - `http://openig.example.com:8080/*?*`

    c. (For AM 6.5.3 and later versions) Select Applications > Agents > Identity Gateway, and add an agent with the following values:

        - Agent ID: `ig_agent`

        - Password: `password`

        Leave all other values as default.

        (For AM 6.5.2 and earlier versions) Set up an agent as described in "To Set Up a Java Agent in AM" in the *Gateway Guide*.

2. Set up IG:

    a. Add a file called `agent.secret.id` to the secret store directory, and add the AM Java agent password to the file:

        ```
        $ echo cGFzc3dvcmQ= > /path/to/secrets/agent.secret.id
        ```

        The value is the base64-encoded value of `password`, which is the password of the AM Java agent.

FORGEROCK

FileSystemSecretStore

b. In IG, add the following route as `$HOME/.openig/config/routes/filesystemsecret.json` (on Windows, `%appdata%\OpenIG\config\routes\filesystemsecret.json`):

```json
{
  "heap":[
    {
      "name": "AmService-1",
      "type": "AmService",
      "config": {
        "url": "http://openam.example.com:8088/openam",
        "agent": {
          "username": "ig_agent",
          "passwordSecretId": "agent.secret.id"
        },
        "version": "6.5",
        "notifications": {
          "enabled": true
        }
      }
    }],
  "name": "filesystemsecret",
  "secrets": {
    "stores": [
      {
        "type": "FileSystemSecretStore",
        "config": {
          "format": "BASE64",
          "directory": "/path/to/secrets"
        }
      }
    ]
  },
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [
        {
          "type": "SingleSignOnFilter",
          "config": {
            "amService": "AmService-1"
          }
        }
      ],
      "handler": "ReverseProxyHandler"
    }
  },
  "condition": "${matches(request.uri.path, '/home/filesystemsecret')}"
}
```

Notice the following features of the route:

- The route matches requests to `/filesystemsecret`.

- The route uses a FileSystemSecretStore to retrieve the value of a secret from a file in `/path/to/secrets/`.

Configuration Reference ForgeRock Identity Gateway 6.5 (2022-03-07T09:02:59.882275)
Copyright © 2011-2021 ForgeRock AS. All rights reserved.
409

- The heap defines an AmService Java agent, that acts on behalf of IG to authenticate with AM. The agent password is defined by the `passwordSecretId`.

  The password is stored in the file `agent.secret.id`. Because the store is configured with base64 format, the password must be base64-encoded.

- The SingleSignOnFilter manages redirects to AM for authentication, using the AM Java agent in `AmService-1`.

3. Test the setup:

   a. If you are logged in to AM, log out.

   b. Go to http://openig.example.com:8080/home/filesystemsecret.

      The SingleSignOnFilter redirects the request to AM for authentication.

   c. Log in to AM as user `demo`, password `Ch4ng31t`.

      When you have authenticated, the SingleSignOnFilter passes the request to the sample app, which returns the profile page.

In the following example, the FileSystemSecretStore is configured to look for suffixed secrets in the `secrets/` directory.

```
{
  "directory": "&{ig.instance.dir}/secrets",
  "suffix": ".base64"
  "format": "BASE64"
}
```

Given that the directory contains both `keystore.pass` and `keystore.pass.base64`, when the `keystore.pass` secret is trying to be resolved, the secret stored in `keystore.pass.base64` will be returned.

More Information

Secrets

org.forgerock.openig.secrets.FileSystemSecretStoreHeaplet

Name

HsmSecretStore — manage a store of secrets with an HSM

Description

Manage a store of secrets with a hardware security module (HSM) device or a software emulation of an HSM device, such as SoftHSM.

Usage

```
{
  "name": string,
  "type": "HsmSecretStore",
  "config": {
    "providerName": configuration expression<string>,
    "storePassword": configuration expression<secret-id>,
    "mappings": [ configuration object, ... ]
  }
}
```

Properties

`"providerName"`: *configuration expression<enumeration>, required*

The name of the pre-installed Java Security Provider supporting an HSM. Use a physical HSM device, or a software emulation of an HSM device, such as SoftHSM.

For the SunPKCS11 provider, concatenate `"providerName"` with the prefix `SunPKCS11-`. For example, declare the following for the name `FooAccelerator`:

```
"providerName": "SunPKCS11-FooAccelerator"
```

`"storePassword"`: *configuration expression<secret-id>, required*

The secret ID of the password to access the HsmSecretStore.

IG searches for the value of this password until it finds it, first locally, then in parent routes, then in `config.json`.

For information about supported formats for `secret-id`, see secret-id.

`"mappings"`: *array of objects, required*

One or more mappings of one secret ID to one or more aliases. The secret store uses the mappings as follows:

- When the secret is used to create signatures or encrypt values, the secret store uses the *active secret*, the first alias in the list.

- When the secret is used to verify signatures or decrypt data, the secret store tries all of the mapped aliases in the list, starting with the first, and stopping when it finds a secret that can successfully verify signature or decrypt the data.

The following example maps a secret ID to two aliases:

```
"mappings": [
  {
    "secretId": "global.pcookie.crypt",
    "aliases": [ "rsapair72-1", "rsapair72-2" ]
  }
]
```

`secretId`: *configuration expression<secret-id>, required*

The ID of the secret used in your configuration.

For information about supported formats for `secret-id`, see secret-id.

`aliases`: *array of configuration expression<string>, required*

One or more aliases for the secret ID.

Example

The example in this section is based on the example in JwtBuilderFilter(5), replacing the KeyStoreSecretStore in that example with an HsmSecretStore.

To set up this example:

1. Create a provider config file, as specified in the PKCS#11 Reference Guide.

2. Depending on your version of Java, create a `java.security.ext` file for the IG instance, with the following content:

   ```
   security.provider.<number>=<provider-name> <path-to-provider-cfg-file>
   ```

   or

   ```
   security.provider.<number>=<class-name> <path-to-provider-cfg-file>
   ```

3. Start the IG JVM with the following system property that points to the provider config file:

   ```
   -Djava.security.properties=file://path-to-security-extension-file
   ```

The following example route is based on the examples in JwtBuilderFilter(5), replacing the KeyStoreSecretStore with an HsmSecretStore:

```
{
  "name": "hsm-jwt-signature",
  "heap": [{
    "name": "AmService-1",
    "type": "AmService",
    "config": {
      "agent" : {
        "username" : "ig_agent",
```

```
        "passwordSecretId": "agent.secret.id"
      },
      "url": "http://openam.example.com:8088/openam",
      "version": "6.5"
    }
  }],
  "secrets": {
    "stores": [
      {
        "type": "HsmSecretStore",
        "config": {
          "providerName": "SunPKCS11-SoftHSM",
          "storePassword": "hsm.pin",
          "mappings": [{
            "secretId": "id.key.for.signing.jwt",
            "aliases": [ "signature-key" ]
          }]
        }
      }
    ]
  },
  "handler": {
    "type": "Chain",
    "config": {
      "filters": [{
        "name": "SingleSignOnFilter-1",
        "type": "SingleSignOnFilter",
        "config": {
          "amService": "AmService-1"
        }
      }, {
        "name": "UserProfileFilter-1",
        "type": "UserProfileFilter",
        "config": {
          "username": "${contexts.ssoToken.info.uid}",
          "userProfileService": {
            "type": "UserProfileService",
            "config": {
              "amService": "AmService-1"
            }
          }
        }
      }, {
        "name": "JwtBuilderFilter-1",
        "type": "JwtBuilderFilter",
        "config": {
          "template": {
            "name": "${contexts.userProfile.commonName}",
            "email": "${contexts.userProfile.rawInfo.mail[0]}"
          },
          "signature": {
            "secretId": "id.key.for.signing.jwt"
          }
        }
      },
      {
        "name": "HeaderFilter-1",
        "type": "HeaderFilter",
        "config": {
```

```
            "messageType": "REQUEST",
            "add": {
              "x-openig-user": ["${contexts.jwtBuilder.value}"]
            }
          }
        }],
        "handler": "ReverseProxyHandler"
      }
    },
    "condition": "${matches(request.uri.path, '/hsm-jwt-signature$')}",
    "baseURI": "http://app.example.com:8081"
}
```

More Information

Secrets

org.forgerock.openig.secrets.HsmSecretStoreHeaplet

Name

KeyStoreSecretStore — manage a KeyStore for cryptographic keys and certificates

Description

Manages a secret store for cryptographic keys and certificates, based on a standard Java KeyStore.

The KeyStore is typically file-based PKCS12 KeyStore. Legacy proprietary formats such as JKS and JCEKS are supported, but implement weak encryption and integrity protection mechanisms. Consider not using them for new functionality.

Usage

```
{
  "name": string,
  "type": "KeyStoreSecretStore",
  "config": {
    "file": configuration expression<string>,
    "storeType": configuration expression<string>,
    "storePassword": configuration expression<secret-id>,
    "keyEntryPassword": configuration expression<secret-id>,
    "mappings": [ configuration object, ... ]
  }
}
```

Properties

`"file"`: *configuration expression<string>, required*

> The path to the KeyStore file.

`"storeType"`: *storeType reference, optional*

> The KeyStore type.
>
> For a list of KeyStore types, see KeyStore Types
>
> Default: The store type used by the platform.

`"storePassword"`: *configuration expression<secret-id>, required*

> The secret ID of the password to access the KeyStore.
>
> IG searches for the value of the password until it finds it, first locally, then in parent routes, then in `config.json`.
>
> For information about supported formats for `secret-id`, see secret-id.

`"keyEntryPassword"`: *configuration expression<secret-id>, optional*

> The secret ID of the password to access entries in the KeyStore.

When this property is used, the password must be the same for all entries in the KeyStore. If JKS uses different password for entries, `keyEntryPassword` doesn't work.

For information about supported formats for `secret-id`, see secret-id.

Default: The value of `storePassword`

`"mappings"`: *array of objects, required*

One or more mappings of one secret ID to one or more aliases. The secret store uses the mappings as follows:

- When the secret is used to create signatures or encrypt values, the secret store uses the *active secret*, the first alias in the list.

- When the secret is used to verify signatures or decrypt data, the secret store tries all of the mapped aliases in the list, starting with the first, and stopping when it finds a secret that can successfully verify signature or decrypt the data.

```
"mappings": [
  {
    "secretId": "id.key.for.signing.jwt",
    "aliases": [ "SigningKeyAlias", "AnotherSigningKeyAlias" ]
  },
  {
    "secretId": "id.key.for.encrypting.jwt",
    "aliases": ["EncryptionKeyAlias"]
  }
]
```

`secretId`: *configuration expression<secret-id>, required*

The ID of the secret used in your configuration.

For information about supported formats for `secret-id`, see secret-id.

`aliases`: *array of configuration expression<string>, required*

One or more aliases for the secret ID.

Example

For examples of routes that use KeyStoreSecretStore, see the examples in JwtBuilderFilter(5).

More Information

org.forgerock.secrets.keystore.KeyStoreSecretStore

org.forgerock.openig.secrets.KeyStoreSecretStoreHeaplet

Name

SystemAndEnvSecretStore — manage a store of secrets from system properties and environment variables

Description

Manage a store of secrets from system properties and environment variables.

For each secret ID in the IG configuration, this object checks for a system property or environment variable with the same name as the ID. If there is a match, the value of the system property or environment variable is stored as the value of the secret.

A secret ID must conform to the regex pattern: `Pattern.compile("[a-zA-Z0-9]+(\\.[a-zA-Z0-9]+)*");`. The reference is then transformed as follows, to match the environment variable name:

• Periods (.) are converted to underscores

• All characters are transformed to uppercase

For example, `my.secret.id` is transformed to `MY_SECRET_ID`.

Usage

```
{
  "name": string,
  "type": "SystemAndEnvSecretStore",
  "config": {
    "format": configuration expression<enumeration>
  }
}
```

Properties

`"format"`: *configuration expression<enumeration>, optional*

Format in which the secret is stored. Use one of the following values:

• `BASE64`: Base64-encoded

• `PLAIN`: Plain text

Default: `BASE64`

Example

In the following example an AmService is configured to authenticate with AM. An AM Java agent is configured with a password, and IG uses the SystemAndEnvSecretStore to retrieve the password from an environment variable.

To Retrieve a Secret From an Environment Variable

Before you start:

- Prepare IG and the sample app as described in "*First Steps*" in the *Getting Started Guide*

- Install and configure AM on http://openam.example.com:8088/openam, using the default configuration.

1. Set up AM:

   a. (For AM 6.5.3 and earlier versions) Select ▣ Identities > demo, and set the demo user password to `Ch4ng31t`.

   b. (For AM 6.5.3 and later versions) Select Services, and add a Validation Service with the following URLs for IG:

      - http://openig.example.com:8080/*

      - http://openig.example.com:8080/*?*

   c. (For AM 6.5.3 and later versions) Select Applications > Agents > Identity Gateway, and add an agent with the following values:

      - Agent ID: `ig_agent`

      - Password: `password`

      Leave all other values as default.

      (For AM 6.5.2 and earlier versions) Set up an agent as described in "To Set Up a Java Agent in AM" in the *Gateway Guide*.

2. Set up IG:

   a. Set an environment variable for the Java agent password:
      ```
      $ export AGENT_SECRET_ID='password'
      ```

   b. In IG, add the following route as `$HOME/.openig/config/routes/systemandenvsecret.json` (on Windows, `%appdata%\OpenIG\config\routes\systemandenvsecret.json`):
      ```
      {
        "heap": [
          {
            "name": "AmService-1",
            "type": "AmService",
            "config": {
              "url": "http://openam.example.com:8088/openam",
              "agent": {
                "username": "ig_agent",
                "passwordSecretId": "agent.secret.id"
      ```

```
      },
      "version": "6.5",
      "notifications": {
        "enabled": true
      }
    }
  }
 }
],
"name": "systemandenvsecret",
"secrets": {
  "stores": [
    {
      "type": "SystemAndEnvSecretStore",
      "config": {
        "format": "PLAIN"
      }
    }
  ]
},
"handler": {
  "type": "Chain",
  "config": {
    "filters": [
      {
        "type": "SingleSignOnFilter",
        "config": {
          "amService": "AmService-1"
        }
      }
    ],
    "handler": "ReverseProxyHandler"
  }
},
"condition": "${matches(request.uri.path, '/home/systemandenvsecret')}"
}
```

Notice the following features of the route:

- The route matches requests to `/home/systemandenvsecret`.

- The route uses a SystemAndEnvSecretStore to retrieve the value of a secret from an environment variable or system property.

- The heap defines an AmService Java agent, that acts on behalf of IG to authenticate with AM. The agent password is defined by the `passwordSecretId agent.secret.id`, that is mapped to the environment variable named `AGENT_SECRET_ID`.

- The SingleSignOnFilter manages redirects to AM for authentication, using the AM Java agent in `AmService-1`.

3. Test the setup:

   a. If you are logged in to AM, log out.

   b. Go to http://openig.example.com:8080/home/systemandenvsecret.

The SingleSignOnFilter redirects the request to AM for authentication.

c. Log in to AM as user `demo`, password `Ch4ng31t`.

When you have authenticated, the SingleSignOnFilter passes the request to sample app, which returns the profile page.

More Information

Secrets

org.forgerock.openig.secrets.SystemAndEnvSecretStoreHeaplet