

RESTful Architecture

TABLE OF CONTENTS

Preface	1
Introduction	1
Definition of RESTful architecture	1
Benefits of using RESTful architecture	1
Characteristics of RESTful architecture	1
HTTP Verbs	2
HTTP verbs	2
Examples of each HTTP verb in action	3
Resources	5
Definition of a resource in RESTful architecture	5
How to identify resources	5
Request and Response Headers	7
Common headers in RESTful architecture	7
Status Codes	7
How to interpret status codes in responses	8
Common Status Codes	8
Authentication and Authorization	9
Explanation of authentication and authorization in RESTful architecture	9
Types of authentication and authorization (e.g. Basic, OAuth2)	9
How to implement authentication and authorization in a RESTful API	10
Hypermedia	11
Definition of hypermedia in RESTful architecture	11
How hypermedia links can be used to represent relationships between resources	11
Examples of hypermedia in action	12
Popular Frameworks	12
Express.js	12
Spring Boot	12
Django REST Framework	12
Ruby on Rails	12
Flask	12
ASP.NET Core	12

PREFACE

The web has become an integral part of our daily lives, and with it, web applications have also become increasingly popular. However, as the number of web applications has grown, so has the complexity of managing them. To tackle this complexity, a new architectural style was developed called RESTful architecture.

This cheatsheet is designed to serve as a quick reference guide for developers who are looking to build RESTful web applications. It provides a concise summary of the key concepts, principles, and best practices that are used in RESTful architecture. The cheatsheet covers topics such as resource identification, representation, and manipulation, HTTP methods, status codes, caching, security, and more.

INTRODUCTION

DEFINITION OF RESTFUL ARCHITECTURE

RESTful architecture is an architectural style for building web services. It is based on the principles of Representational State Transfer (REST), which define a set of constraints that web services should follow in order to be scalable, maintainable, and interoperable.

The main idea behind RESTful architecture is to treat web services as a set of resources that can be accessed using standard HTTP methods (GET, POST, PUT, DELETE, etc.). Each resource is identified by a unique URL (Uniform Resource Locator), and its state can be transferred between client and server in a standardized format, such as JSON or XML.

RESTful architecture is characterized by a number of constraints, including client-server architecture, statelessness, cacheability, layered system, uniform interface, and code on demand. These constraints make it easier to build web services that can be consumed by different clients, such as web browsers, mobile devices, and other applications, and that can scale to handle large amounts of traffic.

BENEFITS OF USING RESTFUL ARCHITECTURE

- **Scalability:** RESTful architecture makes it easier to build web services that can handle large amounts of traffic and scale as the demand for the service grows. This is because RESTful services are stateless and cacheable, which reduces the load on the server and makes it easier to distribute the workload across multiple servers.
- **Interoperability:** RESTful services can be consumed by a wide range of clients, including web browsers, mobile devices, and other applications, making them highly interoperable. This is because RESTful services use standard HTTP methods and formats, such as JSON or XML, which are widely supported by different platforms and programming languages.
- **Simplicity:** RESTful architecture is based on a simple set of principles and constraints, which makes it easier to design and develop web services that are easy to understand and maintain. This simplicity also makes it easier to debug and troubleshoot issues that may arise in the service.
- **Flexibility:** RESTful services can be designed to be flexible and adaptable to changing requirements, as they are not tied to any specific technology or platform. This makes it easier to evolve the service over time as new requirements emerge or as the needs of the users change.
- **Security:** RESTful services can be designed to be secure by implementing authentication and authorization mechanisms, such as OAuth2 or JSON Web Tokens. This makes it easier to protect sensitive data and ensure that only authorized users can access the service.

CHARACTERISTICS OF RESTFUL ARCHITECTURE

- **Client-server architecture:** RESTful architecture is based on a client-server model, where the client sends requests to the server and the server responds with the requested data. This separation of concerns allows for more flexible and scalable architectures.
- **Statelessness:** RESTful services are stateless,

meaning that each request contains all the information necessary for the server to process the request. This reduces the load on the server and makes it easier to distribute the workload across multiple servers.

- **Cacheability:** RESTful services can be designed to be cacheable, which allows clients to reuse the response from a previous request. This can significantly reduce the number of requests made to the server, improving performance and scalability.
- **Layered system:** RESTful architecture can be designed as a layered system, where each layer performs a specific function. This allows for more flexibility and scalability, as each layer can be designed and scaled independently of the other layers.
- **Uniform interface:** RESTful architecture is based on a uniform interface, where resources are identified by a unique URL and are accessed using standard HTTP methods (GET, POST, PUT, DELETE, etc.). This simplifies the design of the service and makes it easier to consume by different clients.
- **Code on demand:** RESTful services can be designed to allow for code on demand, where the server can provide executable code, such as JavaScript, that can be executed by the client. This can provide additional functionality and flexibility to the service, but may also introduce security risks.

HTTP VERBS

HTTP VERBS

GET

GET is used to retrieve a resource from the server. It is a safe and idempotent operation, meaning that multiple requests for the same resource will return the same result and will not modify the state of the resource on the server. Use cases for GET include:

- Retrieving a web page from a server
- Fetching a list of blog posts from a server
- Downloading an image or file from a server

POST

POST is used to create a new resource on the server. It is not idempotent, meaning that multiple requests to create the same resource will result in multiple resources being created. POST requests may also modify the state of the resource on the server. Use cases for POST include:

- Creating a new user account on a server
- Uploading a new image or file to a server
- Submitting a form to a server to create a new resource

PUT

PUT is used to update an existing resource on the server. It is idempotent, meaning that multiple requests to update the same resource will result in the same state of the resource on the server. PUT requests may also create a new resource if the resource being updated does not exist. Use cases for PUT include:

- Updating the contents of a blog post on a server
- Modifying the details of a user account on a server
- Replacing a file or image on a server with a new version

DELETE

DELETE is used to delete a resource from the server. It is idempotent, meaning that multiple requests to delete the same resource will result in the same state of the resource on the server. Use cases for DELETE include:

- Removing a blog post from a server
- Deleting a user account from a server
- Removing a file or image from a server

HEAD

HEAD is used to retrieve the headers for a resource without retrieving the body of the resource. This can be useful for checking if a resource exists or for retrieving metadata about the resource. Use cases for HEAD include:

- Checking if a web page exists on a server

- Verifying if an image or file is available on a server
- Retrieving metadata about a resource without downloading the entire resource

OPTIONS

OPTIONS is used to retrieve information about the communication options available for a resource. This can include information about the supported HTTP methods, content types, and authentication options. Use cases for OPTIONS include:

- Retrieving information about the available API endpoints on a server
- Determining the supported content types for a resource on a server
- Checking the available authentication options for a resource on a server

PATCH

PATCH is used to make a partial update to an existing resource on the server. It allows for updates to be made to specific fields or properties of a resource without requiring the entire resource to be sent to the server. Use cases for PATCH include:

- Updating the title or description of a blog post on a server
- Changing the password or email address associated with a user account on a server
- Making a partial update to a file or image on a server

CONNECT

CONNECT is used to establish a network connection to a server over HTTPS. This is typically used in a proxy scenario, where a client wants to connect to a secure server through an intermediary server.

TRACE

TRACE is used to retrieve a diagnostic trace of the request-response cycle. This is mainly used for debugging purposes and is not commonly used in production.

PROPFIND

PROPFIND is used to retrieve the properties of a resource on a server, such as its metadata or attributes. This is mainly used in WebDAV (Web Distributed Authoring and Versioning) for managing files on a remote server.

PROPPATCH

PROPPATCH is used to modify the properties of a resource on a server. This is also mainly used in WebDAV for managing files on a remote server.

MKCOL

MKCOL is used to create a new collection (a group of related resources) on a server. This is also mainly used in WebDAV for managing files on a remote server.

COPY

COPY is used to create a copy of a resource on a server. This is also mainly used in WebDAV for managing files on a remote server.

MOVE

MOVE is used to move a resource from one location to another on a server. This is also mainly used in WebDAV for managing files on a remote server.

LOCK

LOCK is used to lock a resource on a server, preventing other clients from modifying the resource while it is locked. This is also mainly used in WebDAV for managing files on a remote server.

EXAMPLES OF EACH HTTP VERB IN ACTION

GET

- **Example:** Retrieving a list of blog posts from a server
- **Request:** `GET /api/posts`
- **Response:** Returns a list of blog posts in the response body

POST

- **Example:** Creating a new user account on a server
- **Request:** `POST /api/users`
- **Body:** { "name": "John Smith", "email": "john@example.com", "password": "password123" }
- **Response:** Returns the newly created user account in the response body

PUT

- **Example:** Updating the details of a user account on a server
- **Request:** `PUT /api/users/1`
- **Body:** { "name": "Jane Doe", "email": "jane@example.com" }
- **Response:** Returns the updated user account in the response body

DELETE

- **Example:** Removing a blog post from a server
- **Request:** `DELETE /api/posts/1`
- **Response:** Returns a success message in the response body

HEAD

- **Example:** Checking if a web page exists on a server
- **Request:** `HEAD /example-page.html`
- **Response:** Returns the headers for the web page in the response, but without the body content

OPTIONS

- **Example:** Retrieving information about the available API endpoints on a server
- **Request:** `OPTIONS /api`
- **Response:** Returns a list of available HTTP methods and other options for the specified resource

PATCH

- **Example:** Updating the title of a blog post on a server
- **Request:** `PATCH /api/posts/1`
- **Body:** { "title": "New Title" }
- **Response:** Returns the updated blog post in the response body

CONNECT

- **Example:** Establishing a secure network connection to a server over HTTPS via a proxy server
- **Request:** `CONNECT secure.example.com:443`
- **Response:** Establishes a secure connection with the target server via the proxy server

TRACE

- **Example:** Retrieving a diagnostic trace of the request-response cycle for a specific resource
- **Request:** `TRACE /api/posts/1`
- **Response:** Returns a trace of the request-response cycle for the specified resource

PROPFIND

- **Example:** Retrieving the properties of a file on a remote server
- **Request:** `PROPFIND /webdav/files/document.txt`
- **Response:** Returns a list of properties for the specified file on the remote server

PROPPATCH

- **Example:** Updating the properties of a file on a remote server
- **Request:** `PROPPATCH /webdav/files/document.txt`
- **Body:** { "Author": "John Doe", "LastModified": "2022-03-08T12:00:00Z" }
- **Response:** Returns the updated properties for the specified file on the remote server

MKCOL

- **Example:** Creating a new collection on a remote server

- **Request:** `MKCOL /webdav/collection`
- **Response:** Creates a new collection on the remote server and returns a success message in the response body

COPY

- **Example:** Creating a copy of a file on a remote server
- **Request:** `COPY /webdav/files/document.txt`
`HTTP/1.1 Destination: webdav/files/document_copy.txt`
- **Response:** Creates a copy of the specified file on the remote server and returns a success message in the response body

MOVE

- **Example:** Moving a file from one location to another on a remote server
- **Request:** `MOVE /webdav/files/document.txt`
`HTTP/1.1 Destination: /webdav/new`

LOCK

- **Example:** Locking a resource on a remote server to prevent concurrent access
- **Request:** `LOCK /webdav/files/document.txt`
`HTTP/1.1 If: <etag-value>`
- **Response:** Locks the specified resource on the remote server and returns a success message in the response body

RESOURCES

DEFINITION OF A RESOURCE IN RESTFUL ARCHITECTURE

In RESTful architecture, a resource is a piece of information or data that can be accessed and manipulated through a unique identifier, called a Uniform Resource Identifier (URI). Resources can be anything that can be named, such as a document, an image, a video, a user account, or any other type of data.

Each resource is uniquely identified by a URI, which provides a way to locate and manipulate the resource. A URI is composed of a scheme (e.g. "http"), a host (e.g. "example.com"), a path (e.g.

"/users/123"), and optional query parameters (e.g. "?limit=10").

Resources are the fundamental building blocks of a RESTful API, and the interactions with them are defined by the HTTP verbs, such as GET, POST, PUT, and DELETE. By using a uniform and standardized interface, RESTful APIs can be easily integrated with a wide range of clients and servers, making it a popular choice for building scalable and interoperable web services.

HOW TO IDENTIFY RESOURCES

Identifying resources in a RESTful API involves defining the different types of data that can be represented by the API, and assigning each type of data a unique URI. Here are some general guidelines for identifying resources in a RESTful architecture:

Identify the types of data

Start by identifying the different types of data that your API will represent. This could include users, products, orders, invoices, or any other type of data that your application needs to manage.

Resource Type	Example
Users	/users/123
Products	/products/456
Orders	/orders/789
Invoices	/invoices/1011
Blog posts	/posts/1213
Photos	/photos/1415

Assign a unique URI to each resource

Assigning a unique URI to each resource is a key aspect of RESTful architecture. The URI should be composed of a scheme, host, and path that uniquely identifies the resource. Here's an example of how to assign a URI to a user resource:

- Choose a scheme: The scheme specifies the protocol used to access the resource. For example, "http" is a commonly used scheme for web resources. Let's use "https" as our scheme for this example.
- Choose a host: The host specifies the domain

name or IP address of the server where the resource is hosted. For example, if the resource is hosted on a server at example.com, the host would be "example.com". Let's use "api.example.com" as our host for this example.

- **Choose a path:** The path specifies the unique identifier for the resource. For example, to identify a user resource with an ID of 123, we might use the path "/users/123". Here's the full URI:

```
https://api.example.com/users/123
```

This URI uniquely identifies the user resource with an ID of 123 within the API. By following this approach, you can assign a unique URI to each resource in your API, making it easy to locate and manipulate the resources using HTTP verbs such as GET, POST, PUT, and DELETE.

Use nouns instead of verbs

- **Verb-based URI:** Let's say we have a resource that represents a list of users. A verb-based URI might look something like this:

```
GET /get_users
```

This URI uses the verb **GET** to describe the action being performed, and "users" to identify the type of resource being retrieved.

- **Noun-based URI:** Now let's rephrase the same URI using a noun-based approach:

```
GET /users
```

This URI uses the noun "users" to identify the type of resource being retrieved. By using a noun-based approach, the URI is more expressive and easier to understand.

Using nouns instead of verbs is a common convention in RESTful architecture, and it helps to make the URIs more consistent and intuitive. It also makes it easier for clients to work with the API, since they can infer the meaning of the resource from the URI.

Use plural nouns for collections

Using plural nouns for collections is a widely accepted convention in RESTful architecture. This helps to differentiate collections from individual resources and makes the URIs more consistent and intuitive. Here's an example to illustrate this convention:

- **Singular noun for individual resource:** Let's say we have a resource that represents a user. A URI for an individual user might look something like this:

```
GET /user/123
```

This URI uses the singular noun "user" to identify the individual resource with ID 123.

- **Plural noun for collection of resources:** Now let's say we have a collection of users. A URI for the collection of users would use a plural noun, like this:

```
GET /users
```

This URI uses the plural noun "users" to identify the collection of resources. It indicates that the API should return a list of all users, rather than a single user.

Using plural nouns for collections helps to make the URIs more consistent and intuitive. It also makes it easier for clients to work with the API, since they can infer the meaning of the resource from the URI.

Use query parameters for filtering and sorting

Using query parameters for filtering and sorting is a common convention in RESTful architecture. Query parameters allow clients to customize the data returned by the API without the need for separate URIs for each possible combination of filtering and sorting options. Here's an example to illustrate this convention:

- **Filtering with query parameters:** Let's say we have a collection of users, and we want to filter the collection to only return users with a particular status. We can use a query parameter to achieve this:

```
GET /users?status=active
```

This URI uses the `status` query parameter to filter the collection of users to only include those with an `active` status.

- **Sorting with query parameters:** Now let's say we want to sort the collection of users by their last name. We can use a query parameter to achieve this:

```
GET /users?sort=last_name
```

This URI uses the `sort` query parameter to sort the collection of users by their last name.

Using query parameters for filtering and sorting allows clients to customize the data returned by the API without the need for separate URIs for each possible combination of filtering and sorting options. It also makes the URIs more consistent and easier to understand, since the same URI can be used to retrieve different subsets of the resource by simply changing the query parameters.

REQUEST AND RESPONSE HEADERS

COMMON HEADERS IN RESTFUL ARCHITECTURE

HTTP headers are used to provide additional information about a request or response in RESTful architecture. Here are some common headers used in RESTful APIs and their explanations:

Authorization

The `Authorization` header can be used to provide authentication information, such as an access token, to the server. This allows the server to verify that the client is authorized to access the requested resource.

Content-Type

The `Content-Type` header can be used to specify the media type of the data being sent in the request or response. This allows the client and server to agree on the format of the data being exchanged.

Accept

This header is used to specify the media types that the client can accept in the response. If the server cannot provide a response in one of the specified media types, it will return an error.

Cache-Control

The `Cache-Control` header can be used to specify caching behavior for the response. This allows the client and server to agree on how long the response should be cached and whether it can be cached at all

If-None-Match

The `If-None-Match` header can be used to perform conditional requests. This allows the client to include an entity tag in the request, and the server will only return a response if the entity tag has changed since the last request

User-Agent

This header is used to identify the client making the request, such as the browser or operating system being used.

X-Requested-With

This header is used to indicate that a request was made using Ajax. It is often used by server-side frameworks to differentiate between traditional HTTP requests and Ajax requests.

Custom headers

Custom headers can be used to provide additional information that is specific to the application or API being developed. For example, a custom `X-Forwarded-For` header can be used to indicate the IP address of the original client when requests are proxied through a load balancer.

Location

This header is used in a 201 (Created) response to indicate the URL of the newly created resource.

STATUS CODES

HOW TO INTERPRET STATUS CODES IN RESPONSES

HTTP status codes are an important part of RESTful APIs, as they provide information about the outcome of a request to a client. Here's how to interpret status codes in responses:

Pattern	Meaning
2xx codes	These indicate successful responses, with 200 indicating a successful response with content, 201 indicating a successful response with a newly created resource, and 204 indicating a successful response with no content
3xx codes	These indicate redirection, with 301 indicating a permanent redirection, 302 indicating a temporary redirection, and 304 indicating that the resource has not been modified and the cached version can be used
4xx codes	These indicate client errors, with 400 indicating a bad request due to invalid syntax or missing parameters, 401 indicating that authentication is required and has failed or has not been provided, 403 indicating that the client does not have permission to access the requested resource, and 404 indicating that the requested resource was not found

Pattern	Meaning
5xx codes	These indicate server errors, with 500 indicating that an unexpected error occurred on the server

It's important to note that not all status codes are returned for all requests. For example, a GET request that successfully returns a resource will typically return a 200 status code, while a DELETE request that successfully deletes a resource will typically return a 204 status code. Similarly, some APIs may return additional status codes for specific types of requests or errors.

Interpreting status codes is an important part of working with RESTful APIs, as it allows clients to respond appropriately to errors and other conditions.

COMMON STATUS CODES

HTTP status codes are used to indicate the outcome of an HTTP request in RESTful architecture. Here are some common status codes used in RESTful APIs and their explanations:

Status Code	Explanation
200 OK	This status code indicates that the request was successful and the server has returned the requested resource
201 Created	This status code indicates that the request was successful and a new resource has been created on the server. The URL of the newly created resource should be included in the <code>Location</code> header of the response

Status Code	Explanation
204 No Content	This status code indicates that the request was successful, but there is no content to return in the response. This is often used for PUT and DELETE requests
400 Bad Request	This status code indicates that the server was unable to understand the request due to invalid syntax or missing parameters
401 Unauthorized	This status code indicates that the request requires authentication and the client has not provided valid authentication credentials
403 Forbidden	This status code indicates that the client does not have permission to access the requested resource
404 Not Found	This status code indicates that the server was unable to locate the requested resource
405 Method Not Allowed	This status code indicates that the requested HTTP method is not supported for the requested resource
500 Internal Server Error	This status code indicates that an unexpected error occurred on the server

By using these and other HTTP status codes in responses, RESTful APIs can provide information about the outcome of requests to clients, enabling them to react appropriately to errors and other conditions.

AUTHENTICATION AND AUTHORIZATION

EXPLANATION OF AUTHENTICATION AND AUTHORIZATION IN RESTFUL ARCHITECTURE

Both authentication and authorization are critical for securing RESTful APIs and protecting sensitive data. By requiring authentication and properly authorizing requests, APIs can ensure that only authorized users and clients are able to access protected resources.

Authentication

Authentication is the process of verifying the identity of a user or client. In the context of RESTful APIs, this often involves sending credentials (such as a username and password) to the server to prove identity. The server will then validate the credentials and return an authentication token that can be used to access protected resources. Common authentication methods for RESTful APIs include Basic Auth, OAuth, and JWT (JSON Web Tokens).

Authorization

Authorization is the process of determining whether a user or client has permission to access a particular resource or perform a particular action. In the context of RESTful APIs, this often involves checking the permissions associated with the user or client's authentication token. If the user or client has sufficient permissions, the request will be allowed to proceed. If not, the request will be denied with an appropriate error code (such as 401 Unauthorized or 403 Forbidden).

TYPES OF AUTHENTICATION AND AUTHORIZATION (E.G. BASIC, OAUTH2)

There are various types of authentication and authorization mechanisms that can be used in RESTful architecture, depending on the specific security requirements of an API. Here are some common types:

Authentication

Basic Authentication

In this method, the user’s credentials (username and password) are sent in the header of the HTTP request in plain text, usually in the form of a Base64-encoded string. This method is not very secure and should only be used with HTTPS.

Token-Based Authentication

This method involves the server generating a token (such as a JWT) after the user has authenticated. The token is then sent to the client, who includes it in the header of subsequent requests. The server then validates the token to ensure that the user is authorized to access the requested resources.

OAuth

This is a widely used authentication standard that allows users to grant third-party applications access to their resources on a server without giving them their login credentials. It works by the user authorizing the third-party application to access their resources using an access token that is issued by the server.

Authorization

Role-Based Access Control (RBAC)

In this method, permissions are assigned to roles, and users are assigned to roles based on their job responsibilities or other criteria. Users can then only access resources that are authorized for their assigned roles.

Attribute-Based Access Control (ABAC)

In this method, permissions are based on attributes (such as user attributes or resource attributes) and access is granted or denied based on the attributes associated with the user and the resource.

Rule-Based Access Control (RBAC)

In this method, access is granted or denied based on a set of rules that are defined by an administrator. The rules can be based on a wide range of criteria, such as user roles, time of day, location, or any other criteria that the administrator chooses.

HOW TO IMPLEMENT AUTHENTICATION AND AUTHORIZATION IN A RESTFUL API

Choose an appropriate authentication and authorization mechanism

There are many authentication and authorization mechanisms available, as we discussed in the previous answer. Choose the mechanism that best meets your security requirements and the needs of your users.

Implement the authentication mechanism on the server

This involves adding code to the server that can verify user credentials and generate authentication tokens. This code should be designed to securely handle user credentials and prevent unauthorized access.

Implement the authorization mechanism on the server

This involves adding code to the server that can check the permissions associated with the user’s authentication token and determine whether the requested resource can be accessed by the user.

Add authentication and authorization headers to API requests

Once the authentication and authorization mechanisms are in place on the server, the client can add authentication and authorization headers to their API requests. These headers should include the appropriate authentication token and any other relevant information needed to authorize the request.

To add authentication and authorization headers to API requests, you typically need to follow these steps:

- Choose an authentication mechanism that is appropriate for your use case, such as Basic Authentication, Token-Based Authentication, OAuth2, or OpenID Connect.
- Implement the authentication mechanism on the server-side. This typically involves setting up user accounts, generating tokens or credentials, and providing an API for clients to authenticate themselves.

- Generate the appropriate headers on the client-side based on the authentication mechanism you are using. For example, if you are using Basic Authentication, you need to base64 encode the username and password and set the "Authorization" header to the resulting value. If you are using Token-Based Authentication, you need to generate a token and set the "Authorization" header to the token.
- Include any additional headers required by your authentication mechanism. For example, OAuth2 requires a "Bearer" token prefix in the "Authorization" header.

Here's an example of adding a Basic Authentication header to an API request using Python and the Requests library:

```
import requests
from requests.auth import
HTTPBasicAuth

url =
'https://api.example.com/resource'
username = 'myusername'
password = 'mypassword'

response = requests.get(url,
auth=HTTPBasicAuth(username,
password))
```

In this example, the `auth` parameter is used to specify the authentication mechanism and the `HTTPBasicAuth` class is used to encode the username and password into the appropriate format.

Handle authentication and authorization errors

When authentication or authorization fails, the server should return an appropriate error code (such as 401 Unauthorized or 403 Forbidden) to the client. The client should then handle these errors appropriately (for example, by prompting the user to enter valid credentials or displaying an error message).

HYPERMEDIA

DEFINITION OF HYPERMEDIA IN RESTFUL ARCHITECTURE

In RESTful architecture, hypermedia refers to the practice of including links to related resources within API responses. Hypermedia allows clients to discover and navigate related resources without relying on hardcoded URLs, which can be brittle and difficult to maintain.

Hypermedia can be implemented in a variety of ways, but the most common approach is to use a format like HTML or JSON-LD to represent API responses, which include links to related resources. The links typically include a URI for the related resource, as well as metadata describing the relationship between the resources.

The use of hypermedia can help make APIs more flexible and adaptable to changing requirements, as clients can follow links to related resources as needed, rather than relying on hardcoded URLs. Hypermedia can also help improve the discoverability of APIs, as clients can use links to explore the API and discover new resources that may be useful to them.

HOW HYPERMEDIA LINKS CAN BE USED TO REPRESENT RELATIONSHIPS BETWEEN RESOURCES

Representing one-to-many relationships

In a one-to-many relationship, one resource is associated with many related resources. For example, a blog post may be associated with many comments. In this case, hypermedia links can be used to represent the relationship between the blog post and its comments. The blog post resource might include a "comments" link that points to a collection of comments associated with that post.

Representing many-to-many relationships

In a many-to-many relationship, multiple resources are associated with multiple related resources. For example, a user may be associated with many groups, and a group may have many users. In this case, hypermedia links can be used to represent the relationship between users and groups. The user resource might include a "groups" link that points to a collection of groups associated with that user, and the group resource might include a "users" link

that points to a collection of users associated with that group.

Representing hierarchical relationships

In a hierarchical relationship, resources are organized in a tree-like structure, with parent resources containing child resources. For example, a file system might be organized into directories, with each directory containing files and subdirectories. In this case, hypermedia links can be used to represent the hierarchical relationship between resources. The parent directory resource might include links to its child files and subdirectories.

EXAMPLES OF HYPERMEDIA IN ACTION

HTML web pages

HTML is a common format for representing hypermedia links. Web pages are a good example of how hypermedia links can be used to represent relationships between resources. For example, a blog post might include links to related posts, comments, and author information.

HAL (Hypertext Application Language)

HAL is a format for representing hypermedia links in JSON. In HAL, each resource includes links to related resources, along with metadata describing the relationship between the resources. For example, a blog post resource might include links to related comments, along with metadata describing the relationship between the post and its comments.

Siren

Siren is another format for representing hypermedia links in JSON. Siren uses a more expressive format than HAL, which allows for more complex relationships between resources. For example, a user resource might include links to related groups, along with metadata describing the relationship between the user and the groups.

HATEOAS (Hypermedia as the Engine of Application State)

HATEOAS is a principle of RESTful architecture that emphasizes the use of hypermedia links to drive the

application state. In HATEOAS, API clients follow hypermedia links to navigate between resources and perform actions, rather than relying on hardcoded URLs.

POPULAR FRAMEWORKS

EXPRESS.JS

[Express.js](#) is a popular framework for building Node.js applications, including RESTful APIs. It provides a simple and flexible API for handling HTTP requests and responses and can be used with a variety of middleware and plugins.

SPRING BOOT

[Spring Boot](#) is a framework for building Java applications, including RESTful APIs. It provides a powerful and flexible API for handling HTTP requests and responses, along with a variety of plugins and libraries for handling authentication, security, and other common tasks.

DJANGO REST FRAMEWORK

[Django REST Framework](#) is a framework for building RESTful APIs using the Python Django web framework. It provides a powerful and flexible API for handling HTTP requests and responses, along with a variety of plugins and libraries for handling authentication, security, and other common tasks.

RUBY ON RAILS

[Ruby on Rails](#) is a popular web framework for building web applications, including RESTful APIs. It provides a simple and flexible API for handling HTTP requests and responses and can be used with a variety of plugins and libraries.

FLASK

[Flask](#) is a lightweight framework for building Python web applications, including RESTful APIs. It provides a simple and flexible API for handling HTTP requests and responses and can be used with a variety of plugins and libraries.

ASP.NET CORE

[ASP.NET Core](#): ASP.NET Core is a cross-platform

framework for building web applications and APIs using the .NET platform. It provides a powerful and flexible API for handling HTTP requests and responses, along with a variety of plugins and libraries for handling authentication, security, and other common tasks. It can be used with a variety of programming languages, including C#, F#, and Visual Basic.



JCG delivers over 1 million pages each month to more than 700K software developers, architects and decision makers. JCG offers something for everyone, including news, tutorials, cheat sheets, research guides, feature articles, source code and more.

Copyright © 2014 Exelisis Media P.C. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by means electronic, mechanical, photocopying, or otherwise, without prior written permission of the publisher.

CHEATSHEET FEEDBACK
WELCOME
support@javacodegeeks.com

SPONSORSHIP
OPPORTUNITIES
sales@javacodegeeks.com