



ICSE'09

3 RESTful Service Design

Cesare Pautasso
Faculty of Informatics
University of Lugano, Switzerland

c.pautasso@ieee.org
<http://www.pautasso.info>

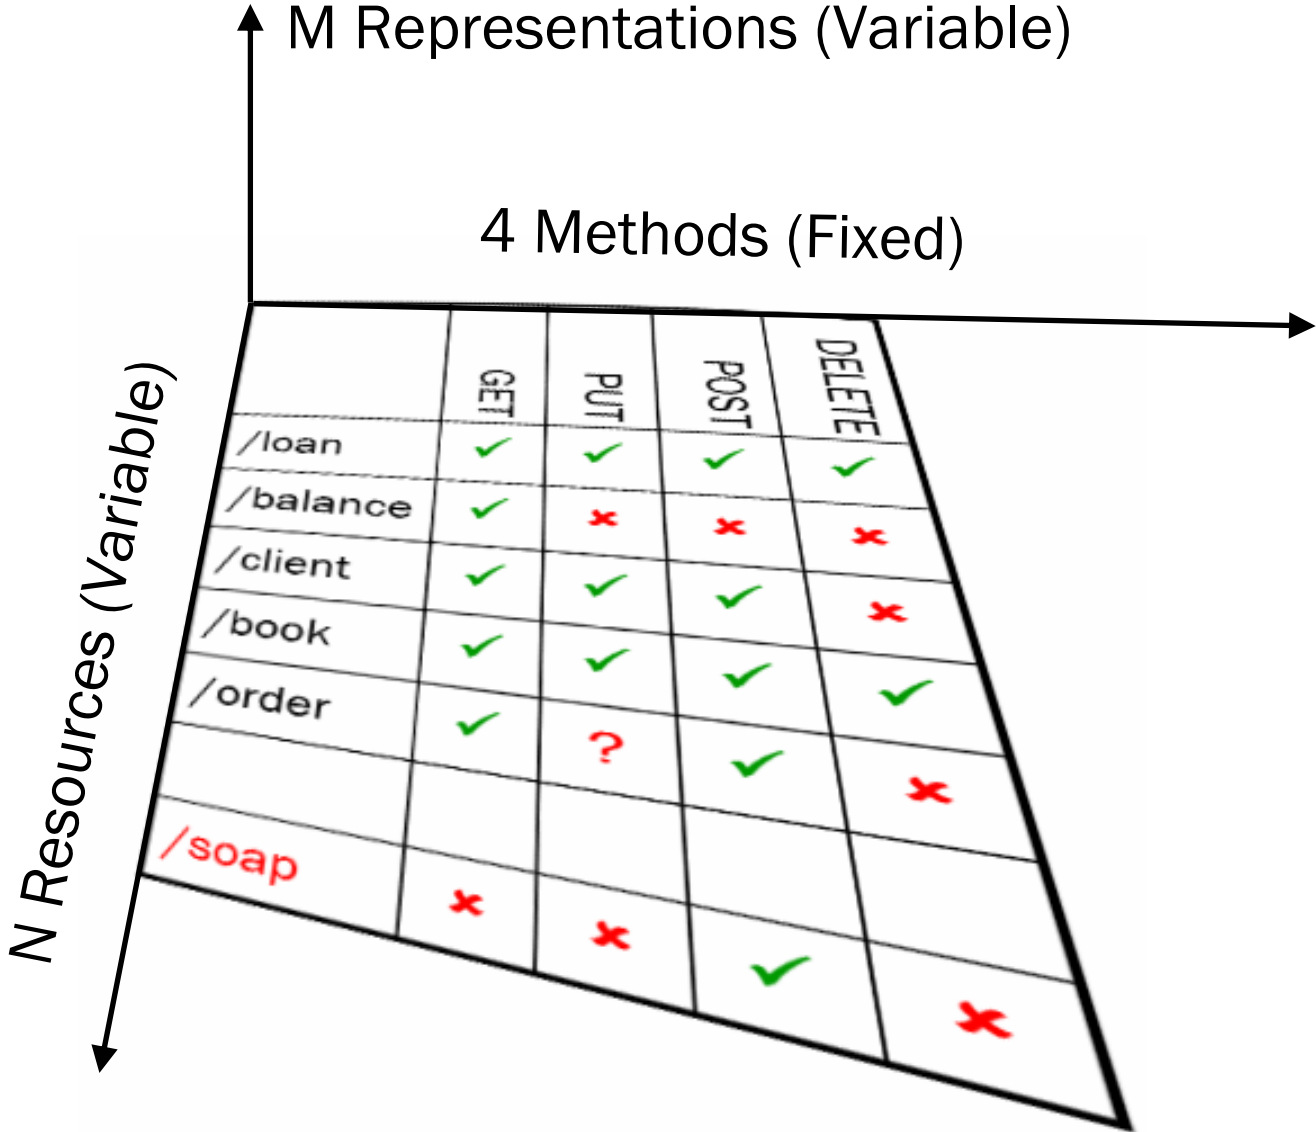
1. Resource Identification
2. Uniform Interface
GET, PUT, DELETE, POST
(HEAD, OPTIONS...)
3. Self-Describing Messages
4. Hypermedia Driving Application State
5. Stateless Interactions

- Design Methodology
 - Is URI Design part of REST?
- Simple Doodle Service Example
- REST Architectural Elements
- Design Tips
 - Understanding GET vs. POST vs. PUT
 - Multiple Representations
 - Content-Type Negotiation
 - Exception Handling
 - Idempotent vs. Unsafe
 - Dealing with Concurrency
 - Stateful or Stateless?
- Some REST AntiPatterns

Design Methodology

1. Identify resources to be exposed as services (e.g., yearly risk report, book catalog, purchase order, open bugs, polls and votes)
2. Model relationships (e.g., containment, reference, state transitions) between resources with hyperlinks that can be followed to get more details (or perform state transitions)
3. Define “nice” URIs to address the resources
4. Understand what it means to do a GET, POST, PUT, DELETE for each resource (and whether it is allowed or not)
5. Design and document resource representations
6. Implement and deploy on Web server
7. Test with a Web browser

	GET	PUT	POST	DELETE
/loan	✓	✓	✓	✓
/balance	✓	✗	✗	✗
/client	✓	✓	✓	✗
/book	✓	✓	✓	✓
/order	✓	?	✓	✗
/soap	✗	✗	✓	✗



- Internet Standard for resource naming and identification (originally from 1994, revised until 2005)
- Examples:

`http://tools.ietf.org/html/rfc3986`

URI Scheme Authority Path

`https://www.google.ch/search?q=rest&start=10#1`

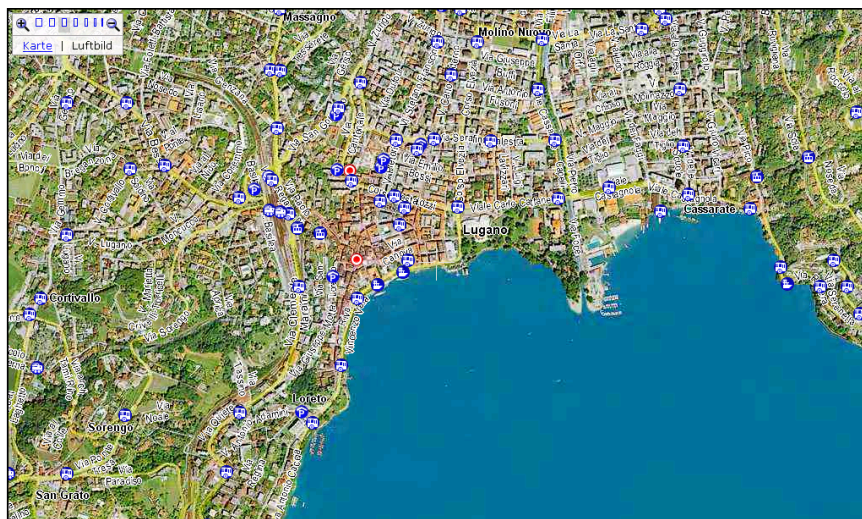
Query Fragment

- REST does not advocate the use of “nice” URIs
- In most HTTP stacks URIs cannot have arbitrary length (4Kb)

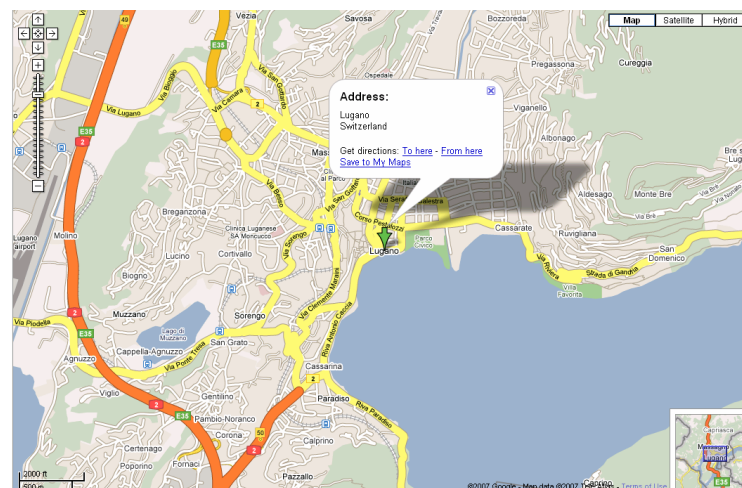
What is a “nice” URI?

A RESTful service is much more than just a set of nice URIs

<http://map.search.ch/lugano>



<http://maps.google.com/lugano>



<http://maps.google.com/maps?f=q&hl=en&q=lugano,+switzerland&layer=&ie=UTF8&z=12&om=1&iwloc=addr>

URI Design Guidelines

- Prefer Nouns to Verbs
- Keep your URIs short
- If possible follow a “positional” parameter-passing scheme for algorithmic resource query strings (instead of the `key=value&p=v` encoding)
- Some use URI postfixes to specify the content type
- Do not change URIs
- Use redirection if you really need to change them

GET `/book?isbn=24&action=delete`

DELETE `/book/24`


- **Note:** REST URIs are opaque identifiers that are meant to be discovered by following hyperlinks and *not constructed by the client*

- *This may break the abstraction*

- **Warning:** URI Templates introduce coupling between client and server

Simple Doodle API Example Design

1. Resources:
polls and votes
2. Containment Relationship:




	GET	PUT	POST	DELETE
/poll	✓	✗	✓	✗
/poll/{id}	✓	✓	✗	✓
/poll/{id}/vote	✓	✗	✓	✗
/poll/{id}/vote/{id}	✓	✓	✗	?

3. URIs embed IDs of “child” instance resources
4. POST on the container is used to create child resources
5. PUT/DELETE for updating and removing child resources

Simple Doodle API Example

1. Creating a poll
(transfer the state of a new poll on the Doodle service)

/poll
 /poll/090331x
 /poll/090331x/vote



POST /poll
 <options>A, B, C</options>

201 Created
 Location: /poll/090331x

GET /poll/090331x

200 OK
 <options>A, B, C</options>
 <votes href="/vote"/>


2. Reading a poll
(transfer the state of the poll from the Doodle service)



Simple Doodle API Example

- Participating in a poll by creating a new vote sub-resource

/poll
/poll/090331x
/poll/090331x/vote
/poll/090331x/vote/1



POST /poll/090331x/vote

<name>C. Pautasso</name>

<choice>B</choice>

201 Created

Location:

/poll/090331x/vote/1

GET /poll/090331x

200 OK

<options>A,B,C</options>

<votes><vote id="1">

<name>C. Pautasso</name>

<choice>B</choice>


</vote></votes>

Simple Doodle API Example

- Existing votes can be updated (access control headers not shown)

```

/pol1
/pol1/090331x
/pol1/090331x/vote
/pol1/090331x/vote/1
  
```



```
PUT /pol1/090331x/vote/1
```

```
<name>C. Pautasso</name>
<choice>C</choice>
```

200 OK

```
GET /pol1/090331x
```

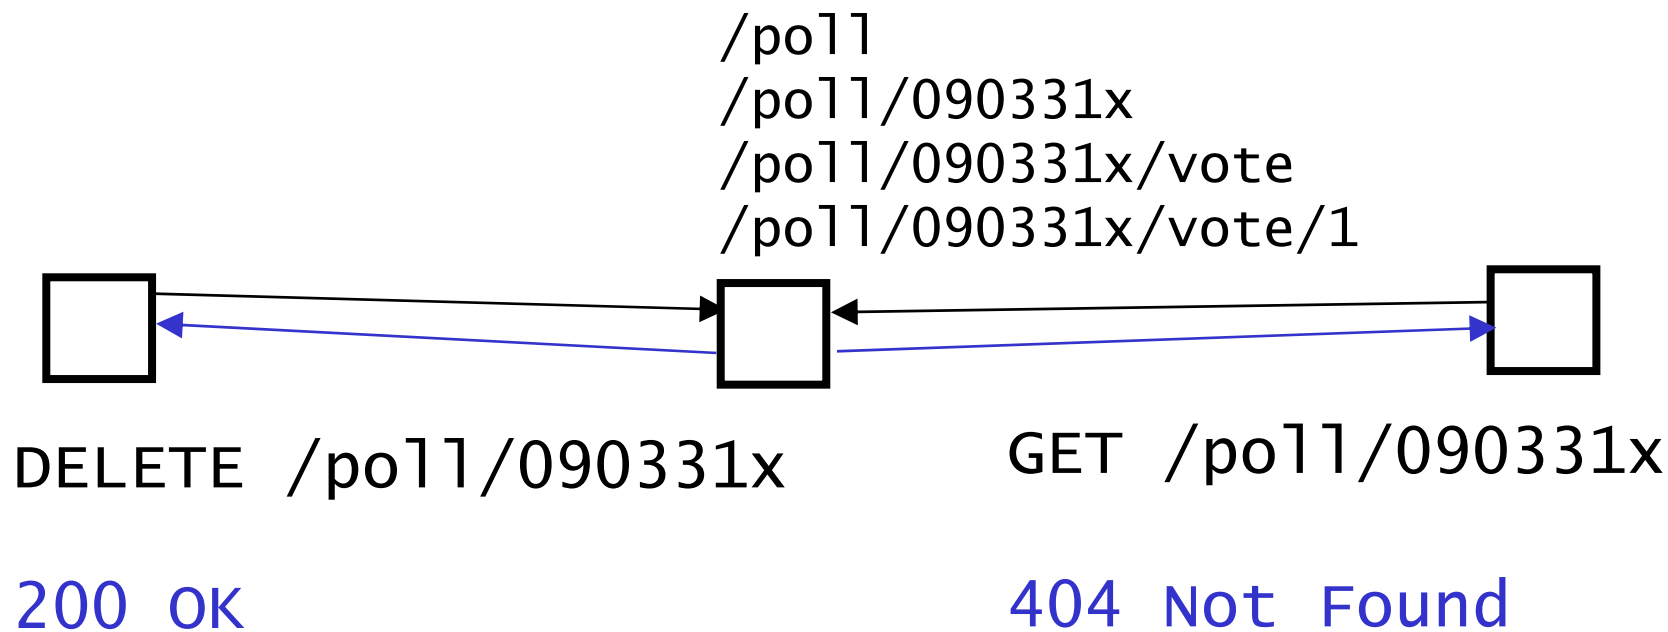
200 OK

```
<options>A,B,C</options>
<votes><vote id="/1">
<name>C. Pautasso</name>
<choice>C</choice>
</vote></votes>
```



Simple Doodle API Example

- Polls can be deleted once a decision has been made



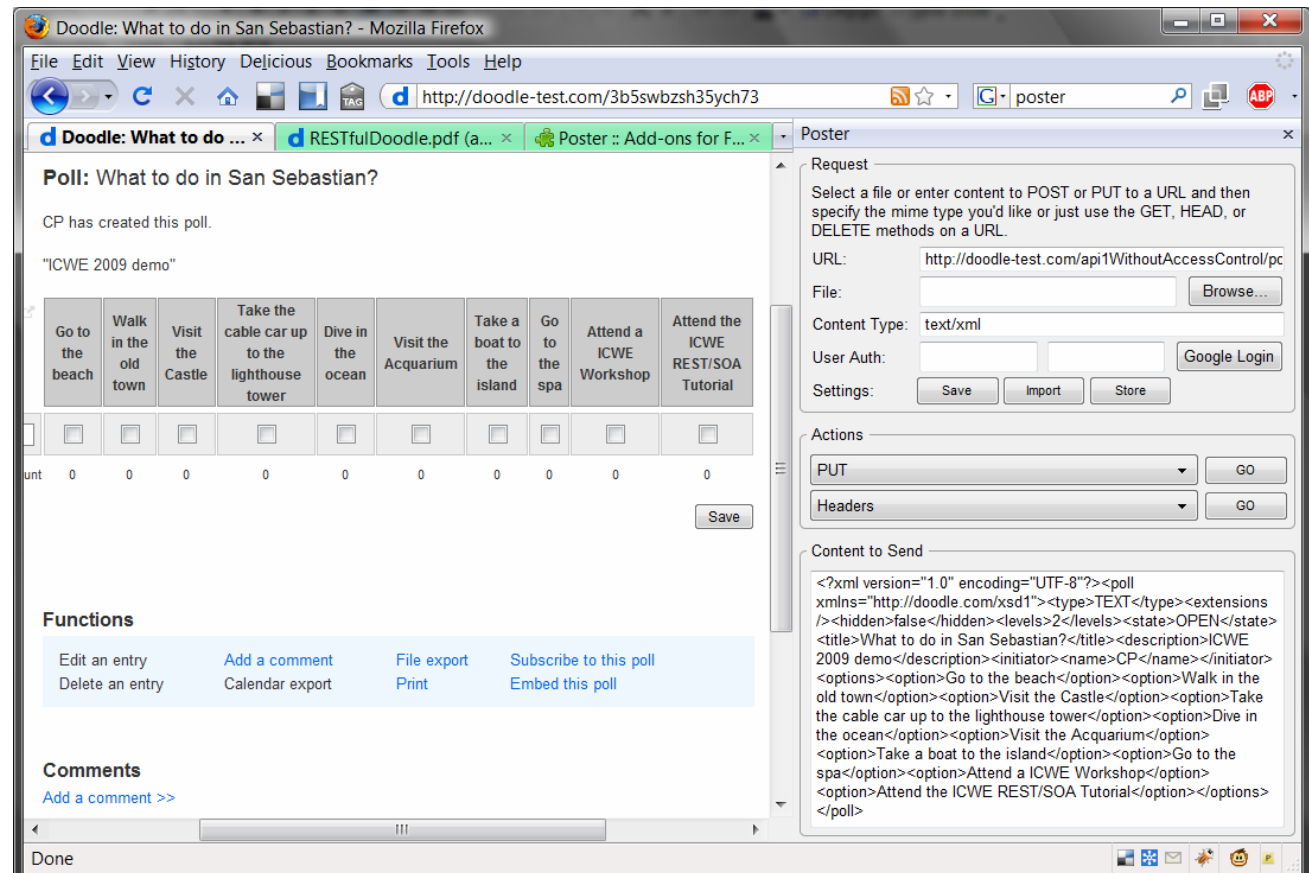
Real Doodle Demo

- Info on the real Doodle API:

<http://doodle.com/xsd1/RESTfulDoodle.pdf>

- Lightweight demo with Poster Firefox Extension:

<http://addons.mozilla.org/en-US/firefox/addon/2691>



The screenshot shows a Mozilla Firefox browser window with the following elements:

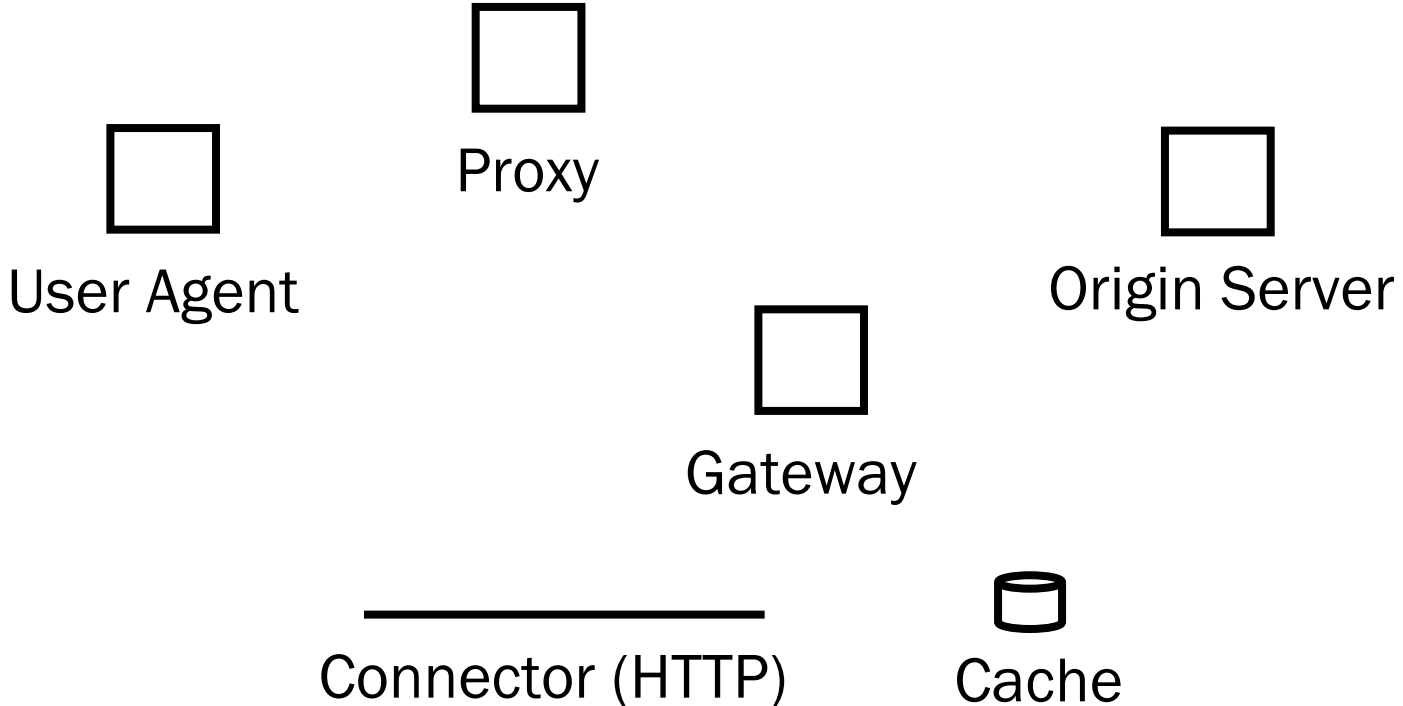
- Browser Tab:** "Doodle: What to do in San Sebastian? - Mozilla Firefox"
- Address Bar:** "http://doodle-test.com/3b5swbzh35ych73"
- Extension:** "Poster" is installed and active.
- Main Content:**
 - Poll:** "What to do in San Sebastian?"
 - Text:** "CP has created this poll." and "ICWE 2009 demo"
 - Options Table:**

Go to the beach	Walk in the old town	Visit the Castle	Take the cable car up to the lighthouse tower	Dive in the ocean	Visit the Acquarium	Take a boat to the island	Go to the spa	Attend a ICWE Workshop	Attend the ICWE REST/SOA Tutorial
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
0	0	0	0	0	0	0	0	0	0
 - Buttons:** "Save" (bottom right of table), "Add a comment >>" (bottom left)
 - Functions:** Edit an entry, Delete an entry, Add a comment, Calendar export, File export, Print, Subscribe to this poll, Embed this poll
 - Comments:** "Add a comment >>"
- Poster Extension Panel:**
 - Request:** Select a file or enter content to POST or PUT to a URL...
 - URL:** http://doodle-test.com/api1WithoutAccessControl/pc
 - File:** [Browse...]
 - Content Type:** text/xml
 - User Auth:** [Google Login]
 - Settings:** [Save] [Import] [Store]
 - Actions:** [PUT] [GO]
 - Headers:** [GO]
 - Content to Send:**


```
<?xml version="1.0" encoding="UTF-8"?><poll xmlns="http://doodle.com/xsd1"><type>TEXT</type><extensions /><hidden>false</hidden><levels>2</levels><state>OPEN</state><title>What to do in San Sebastian?</title><description>ICWE 2009 demo</description><initiator><name>CP</name></initiator><options><option>Go to the beach</option><option>Walk in the old town</option><option>Visit the Castle</option><option>Take the cable car up to the lighthouse tower</option><option>Dive in the ocean</option><option>Visit the Acquarium</option><option>Take a boat to the island</option><option>Go to the spa</option><option>Attend a ICWE Workshop</option><option>Attend the ICWE REST/SOA Tutorial</option></options></poll>
```

REST Architectural Elements


Client/Server Layered Stateless Communication Cache



Basic Setup




Adding Caching




Proxy or Gateway?

Intermediaries forward (and may translate) requests and responses

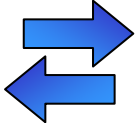
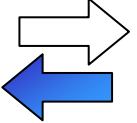
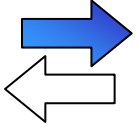
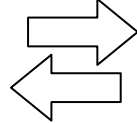


A proxy is chosen by the Client (for caching, or access control)



The use of a gateway (or reverse proxy) is imposed by the server

Uniform Interface Principle


CRUD	REST	
CREATE	POST 	Create a sub resource
READ	GET 	Retrieve the <i>current state</i> of the resource
UPDATE	PUT 	Initialize or update the state of a resource at the given URI
DELETE	DELETE 	Clear a resource, after the URI is no longer valid

POST vs. GET


- GET is a **read-only** operation. It can be repeated without affecting the state of the resource (idempotent) and can be cached.

Note: this does not mean that the same representation will be returned every time.

- POST is a **read-write** operation and may change the state of the resource and provoke side effects on the server.



Web browsers warn you when refreshing a page generated with POST



POST vs. PUT

What is the right way of creating resources (initialize their state)?

→ **PUT /resource/{id}**

← **201 Created**

Problem: How to ensure resource {id} is unique?
(Resources can be created by multiple clients concurrently)

Solution 1: let the client choose a unique id (e.g., GUID)

→ **POST /resource**

← **301 Moved Permanently**

Location: /resource/{id}

Solution 2: let the server compute the unique id

Problem: Duplicate instances may be created if requests are repeated due to unreliable communication

Negotiating the message format does not require to send more messages (the added flexibility comes for free)

⇒ **GET /resource**

**Accept: text/html, application/xml,
application/json**

1. The client lists the set of understood formats (MIME types)

← **200 OK**

Content-Type: application/json

2. The server chooses the most appropriate one for the reply (status 406 if none can be found)

Quality factors allow the client to indicate the relative degree of preference for each representation (or media-range).

Media/Type; q=X

If a media type has a quality value $q=0$, then content with this parameter is not acceptable for the client.

Accept: text/html, text/*; q=0.1

The client prefers to receive HTML (but any other text format will do with lower priority)

**Accept: application/xhtml+xml; q=0.9,
text/html; q=0.5, text/plain; q=0.1**

The client prefers to receive XHTML, or HTML if this is not available and will use Plain Text as a fall back

Forced Content Negotiation

The generic URI supports content negotiation

GET /resource

**Accept: text/html, application/xml,
application/json**

The specific URI points to a specific representation format using the postfix (extension)

GET /resource.html

GET /resource.xml

GET /resource.json

Warning: This is a conventional practice, not a standard.

What happens if the resource cannot be represented in the requested format?

Content Negotiation is very flexible and can be performed based on different dimensions (each with a specific pair of HTTP headers).

Request Header	Example Values	Response Header
Accept:	application/xml, application/json	Content-Type:
Accept-Language:	en, fr, de, es	Content-Language:
Accept-Charset:	iso-8859-5, unicode-1-1	Charset parameter fo the Content-Type header
Accept-Encoding:	compress, gzip	Content-Encoding:

Exception Handling

Learn to use HTTP Standard Status Codes

100 Continue
200 OK
201 Created
202 Accepted
203 Non-Authoritative
204 No Content
205 Reset Content
206 Partial Content
300 Multiple Choices
301 Moved Permanently
302 Found
303 See Other
304 Not Modified
305 Use Proxy
307 Temporary Redirect

4xx Client's fault



400 Bad Request
401 Unauthorized
402 Payment Required
403 Forbidden
404 Not Found
405 Method Not Allowed
406 Not Acceptable
407 Proxy Authentication Required
408 Request Timeout
409 Conflict
410 Gone
411 Length Required
412 Precondition Failed
413 Request Entity Too Large
414 Request-URI Too Long
415 Unsupported Media Type
416 Requested Range Not Satisfiable
417 Expectation Failed

500 Internal Server Error
501 Not Implemented
502 Bad Gateway
503 Service Unavailable
504 Gateway Timeout
505 HTTP Version Not Supported

5xx Server's fault



Idempotent vs. Unsafe

- Idempotent requests can be processed multiple times without side-effects

GET /book

PUT /order/x

DELETE /order/y

- If something goes wrong (server down, server internal error), the request can be simply replayed until the server is back up again
- Safe requests are idempotent requests which do not modify the state of the server (can be cached)

GET /book

- Unsafe requests modify the state of the server and cannot be repeated without additional (unwanted) effects:

Withdraw(200\$) //unsafe

Deposit(200\$) //unsafe

- Unsafe requests require special handling in case of exceptional situations (e.g., state reconciliation)

POST /order/x/payment

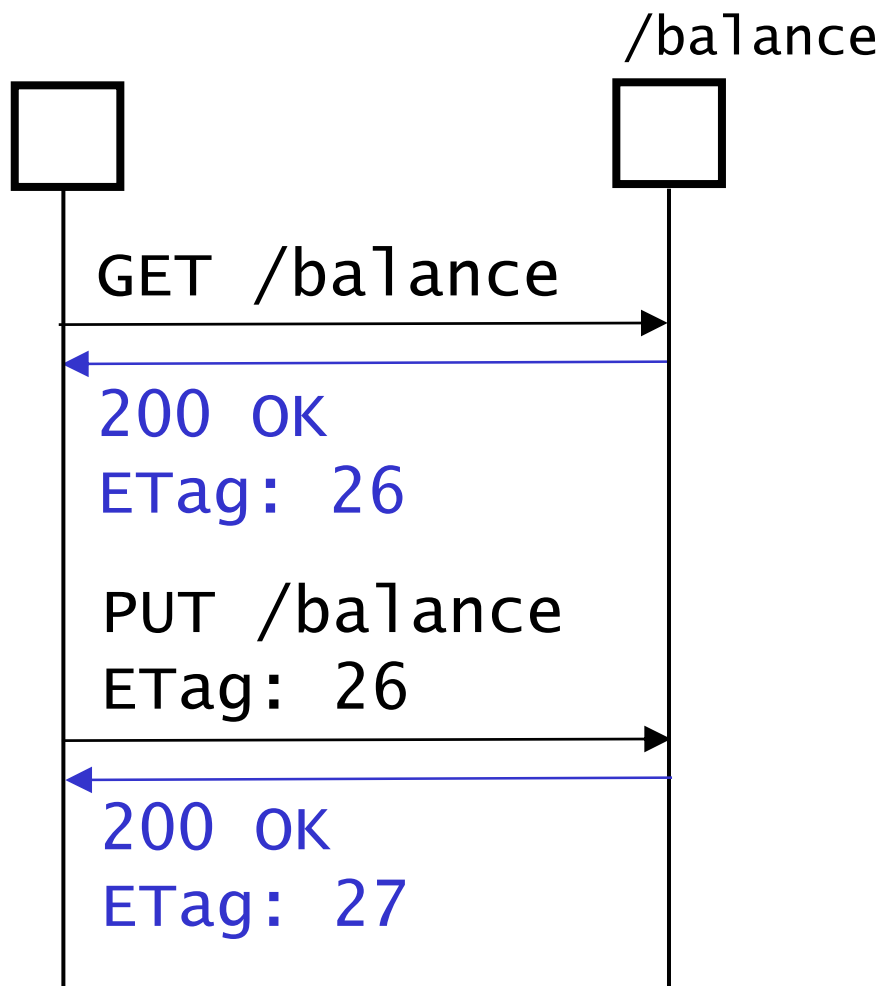
- In some cases the API can be redesigned to use idempotent operations:

B = GetBalance() //safe

B = B + 200\$ //local

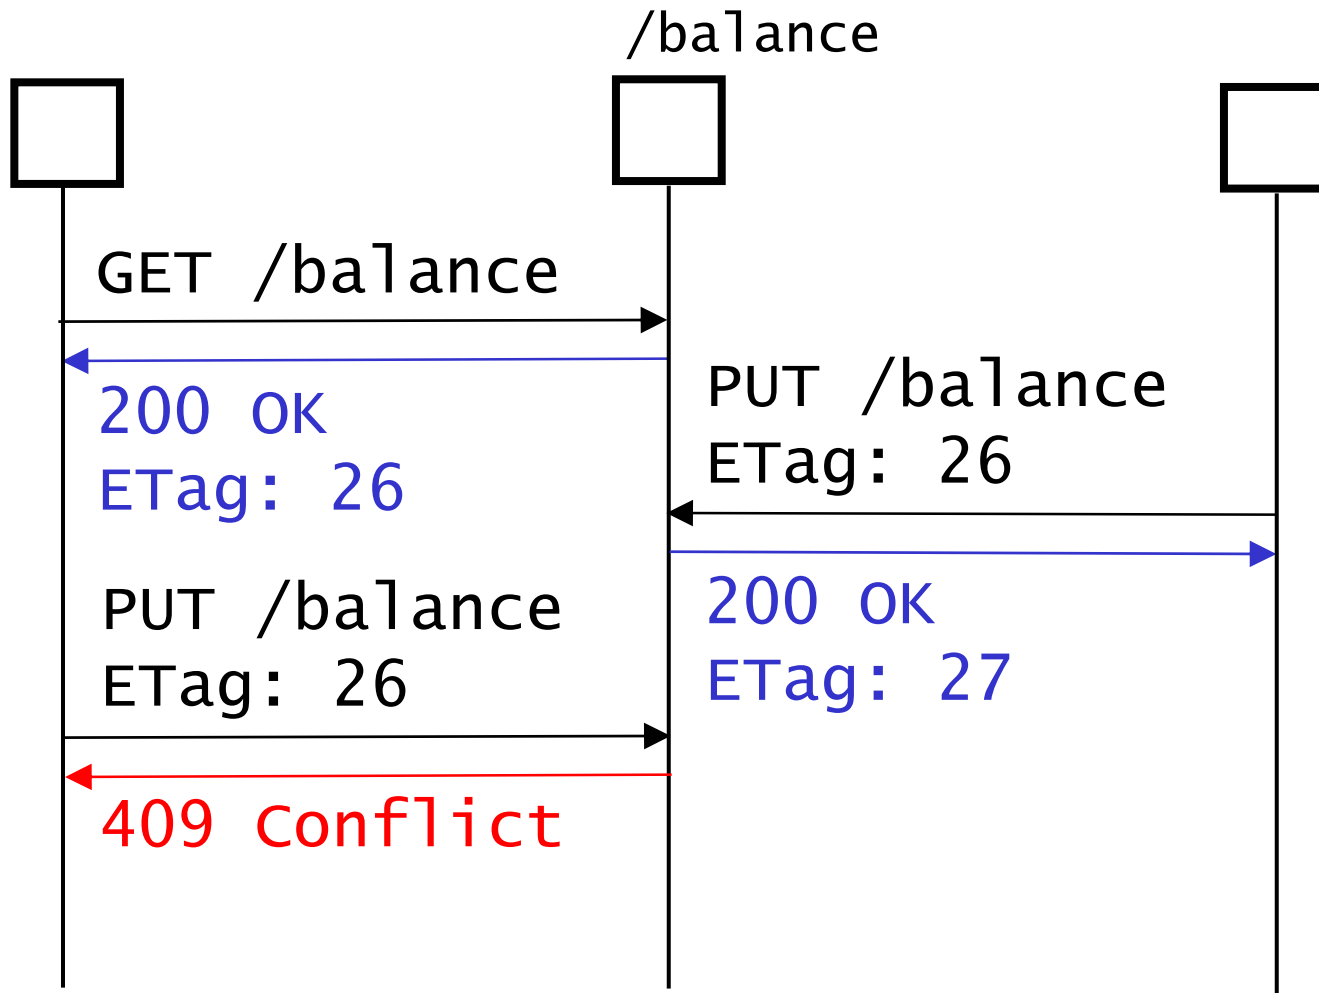
SetBalance(B) //idempotent

Dealing with Concurrency



- Breaking down the API into a set of idempotent requests helps to deal with temporary failures.
- But what about if another client concurrently modifies the state of the resource we are about to update?
- Do we need to create an explicit **/balance/lock** resource? (Pessimistic Locking)
- Or is there an optimistic solution?


Dealing with Concurrency




The 409 status code can be used to inform a client that his request would render the state of the resource inconsistent

Blocking or Non-Blocking?

- HTTP is a synchronous interaction protocol. However, it does not need to be blocking.



- A Long running request may time out.
- The server may answer it with 202 Accepted providing a URI from which the response can be retrieved later.
- Problem: how often should the client do the polling? `/slow/x` could include an estimate of the finishing time if not yet completed



- Tunnel through one HTTP Method

GET /api?method=addCustomer&name=wilde

GET /api?method=deleteCustomer&id=42

GET /api?method=getCustomerName&id=42


GET /api?method=findCustomers&name=wilde*

- Everything through GET

- Advantage: Easy to test from a Browser address bar (the “action” is represented in the resource URI)

- **Problem: GET should only be used for read-only (= idempotent and safe) requests.**

What happens if you bookmark one of those links?




- Limitation: Requests can only send up to approx. 4KB of data (414 Request-URI Too Long)

Antipatterns – HTTP as a tunnel

- Tunnel through one HTTP Method

- Everything through POST



- Advantage: Can upload/download an arbitrary amount of data (this is what SOAP or XML-RPC do)
 - Problem: POST is not idempotent and is unsafe (cannot cache and should only be used for “dangerous” requests)

POST /service/endpoint

```
<soap:Envelope>
  <soap:Body>
    <findCustomers>
      <name>wild* </name>
    </findCustomers>
  </soap:Body>
</soap:Envelope>
```



- Are Cookies RESTful or not?
 - It depends. REST is about stateless communication (without establishing any session between the client and the server)
- 1. Cookies can also be self-contained
 - carry all the information required to interpret them with every request/response
- 2. Cookies contain references to the application state (not maintained as a resource)
 - they only carry the so-called “session-key”
 - Advantage: less data to transfer
 - Disadvantage: the request messages are no longer self-contained as they refer to some context that the server needs to maintain. Also, some garbage collection mechanism for cleaning up inactive sessions is required. More expensive to scale-up the server.

Stateless or Stateful?

- RESTful Web services are not stateless. The very name of “Representational State Transfer” is centered around how to deal with state in a distributed system.

Client State

- The client interacts with resources by “navigating hyperlinks” and its state captures the current position in the hypertext.
- The server may influence the state transitions of the client by sending different representations (containing hyperlinks to be followed) in response to GET requests



Resource State

- The state of resources captures the persistent state of the service.
- This state can be accessed by clients under different representations
- The client manipulates the state of resources using the uniform interface CRUD-like semantics (PUT, DELETE, POST)

Stateless or Stateful?

- RESTful Web services are not stateless. The very name of “Representational State Transfer” is centered around how to deal with state in a distributed system.

Client State



GET /resource

</1>
</2>

GET /1

200 OK
<xml>




Resource State

/resource

</1>
</2>

/1
<xml>



Leonard Richardson,
Sam Ruby,
RESTful Web Services,
O'Reilly, May 2007



Raj Balasubramanians, Benjamin
Carlyle, Thomas Erl, Cesare Pautasso,
SOA with REST,
Prentice Hall, End of 2009



Mashups09 @ OOPSLA

<http://mashup.inf.unisi.ch/mashups09>

3rd International Workshop on
Web APIs and Services Mashups



ICWE'09

From SOA to REST

<http://dret.net/netdret/docs/soa-rest-icwe2009/>

Cesare Pautasso
Faculty of Informatics
University of Lugano, Switzerland

c.pautasso@ieee.org
<http://www.pautasso.info>

Erik Wilde
School of Information
UC Berkeley

dret@berkeley.edu
<http://dret.net/netdret/>