Operations IG Discussion Topics

Mark Taylor (University of Bristol)

Operations IG
IVOA Interop
Bologna

10 May 2023

$Id: ops-topics.tex,v 1.10 2023/05/10 08:01:16 mbt Exp $

Overview

- HTTP $\rightarrow$ HTTPS transition
- Validation of authenticated services
- Operational requirements for authentication

... Discussion ...

HTTP → HTTPS Transition (1)

Many services are migrating from HTTP to HTTPS equivalents

- Often driven by institutional policy

Not always as simple as it seems...

- Certificate management required
  - ▷ acquisition
  - ▷ hostname aliases
  - ▷ expiry (including link rot)
- Legacy clients looking for HTTP variant may stop working
  - ▷ May be a short-term issue, but still sometimes problematic
  - ▷ Reduce impact: ensure URLs under service control point to the right place
    - Registry entries
    - VOSI `capabilities`
    - Links on web pages
    - Redirection `Location` headers in HTTPS responses (201, 3xx)
  - ▷ Some clients may still have hard-coded/outdated URLs
    - How best to mitigate this?

HTTP → HTTPS Transition (2)

Options for HTTP→HTTPS transition, implications for legacy clients using HTTP URL

- Shut down HTTP service
  - ▷ Hard-coded references to HTTP service will fail
- Parallel HTTPS service running alongside HTTP service
  - ▷ Should be OK
- HTTP 3xx redirect to HTTPS
  - ▷ Wholesale 301/302/307/308 from `http://example.com/path` → `https://example.com/path`
  - ▷ Which 3xx code to use?
  - ▷ Works OK for GET and HEAD (any 3xx)
  - ▷ But problematic for POST (e.g. TAP)
    - ... especially if you're POSTing a 100Mb VOTable
    - Recommended behaviour has changed over the years
    - RFC 2616 Sec 10.3 (1999) says MUST NOT just re-POST to redirected location on 301/302/307
    - RFC 9110 Sec 15.4 (2022, obsoletes RFC 2616 via RFC 7231) suggests it's OK for 307/308
    - 308 is probably most correct — but may not be recognised by older clients
  - ▷ Java library code does not follow protocol-changing (e.g. `http`→`https`) redirects
    - (`java.net.URL.openStream()` issue, even with `setFollowRedirects` — see Java bug #4620571)
    - Application code can work round it (if it knows it has to), though not always straightforward
    - Other languages OK?

Validation of Authenticated Services

Authentication can get in the way of validation

- Central validators run a test query on an authenticated service
- Query fails because validator does not authenticate (401/403)
- Validator may log a failed query

What to do?

- Services only advertise accessible data
  - ▷ e.g. VOSI/TAP_SCHEMA table metadata queries only lists tables which currently authenticated user is authorized to see
  - ▷ Is this a better user experience anyway?
  - ▷ ... but won't work for registry
  - ▷ ... or fine-grained restrictions (e.g. auth constraints on certain rows)
- Validators log authentication failures differently than others
  - ▷ Probably validators should do this
  - ▷ I will consider `taplint` updates, but need some (partially) auth-protected services to play with — volunteers?
  - ▷ ... but it may mask genuine problems
- Provide validators (self-run and/or central) with authentication tokens

Other options?

Operational Requirements for Authentication

- Increasing need to track service usage by user
- Driven by resource management rather than data rights
  - *"an increasing frequency of 'denial of service' storms of heavy activity from enthusiastic users doing more with our API's"*
    — Steve Groom, IRSA
  - Concerns about science platform resource usage, e.g. CPU cycles
- Authentication is a GWS WG concern
  - But Ops IG may have input into requirements

Opinions? Experiences? Suggestions?