

ESAC TAP Updates

IVOA Northern Spring 2024 Interoperability Meeting

Jose Osinde (Starion for ESA)

C. Rios (Starion for ESA), M. Hendar (Starion for ESA), J. Ballester (Starion for ESA)

R. Parejo (Starion for ESA), R. Bhatawdekar (ESA)

SCO-08: Archives Software Development

ESAC

Camino Bajo del Castillo s/n, Urb. Villafranca Del Castillo

28692 Villanueva de la Cañada (Madrid) Spain

Release 9.0.0: Library stabilization.

Making the library more stable by fixing bugs, optimizing performance and addressing standard issues:

- ADQL as an independent module of the library
- Adopt latest ADQL 2.1 and DALI 1.1 recommendations
- Adding IT coverage for most of the TAP main functionality

Release 9.1.0 till 9.9.0 (1/3)

- TAP/DataLink authentication mechanism now aligns with IVOA recommendations:
 - IVOA Single-Sign-On Profile: Authentication Mechanisms
Version 2.0
IVOA Recommendation 2017-05-24
 - GWS WG IVOA Interop 27 April 2022
 - <https://www.ivoa.net/documents/SSO/20170524/index.html>
 - <https://github.com/ivoa-std/SSO/>

Release 9.1.0 till 9.9.0 (2/3)

- New commands for updating tables already listed in a TAP service
- Improved support to the Observation Data Model Core (ObsCore) standard
 - Automatically check datatypes, ucd, utypes, xtype on all the table columns
- Adds recommended ADQL 2.1 IVO healpix-related functions
- Improved support to Greenplum databases

Release 9.1.0 till 9.9.0 (3/3)

- More robust session handling behaviour
- Support of POST request methods in the TAP Data Distribution
 - Remove current limit for the arguments size
- Direct Download Strategy
 - First implemented by Euclid team and now integrated in the common library
 - Avoid the generation of big bundles in the server
 - Save server resources (cpu/disk space/memory)
 - User can decide what to download and when

Medium-term planning:

- Migrate to a Spring framework and annotation-based implementation
 - Required step before adopting OpenAPI
- TAP stateless version

Long-term planning:

- Improved support to different databases
- Support of new HEALPix Multi-Order Coverage (MOC) functions and types
- Authentication & Authorisation

Migrate to a Spring framework and annotation-based implementation (1/2):

- **Reduced Boilerplate Code:** Annotation-based configuration in Spring reduces the need for verbose XML configuration files, resulting in cleaner and more concise code. This makes development faster and easier to maintain.
- **Increased Productivity:** Annotations provide a more declarative and intuitive way to configure components, reducing the time required for setup and configuration. Developers can focus more on business logic rather than infrastructure concerns.
- **Simplified Dependency Injection (DI):** Spring's dependency injection mechanism, facilitated by annotations like `@Autowired`, simplifies the management of object dependencies. This allows for loose coupling between components, making the codebase more modular and easier to test.

Migrate to a Spring framework and annotation-based implementation (2/2)

- **Integrated Testing Support:** Spring provides robust support for unit and integration testing. Annotations like `@RunWith` and `@SpringBootTest` facilitate the creation of test cases, while features like Dependency Injection make it easier to mock dependencies for isolated testing.
- **Built-in Security Features:** Spring Security, integrated with Spring MVC, provides comprehensive security features such as authentication, authorization, and session management. Annotations like `@Secured`, `@PreAuthorize`, and `@PostAuthorize` make it easy to secure web endpoints and control access to resources.
- **Support for Web MVC:** Spring MVC offers a robust model-view-controller (MVC) framework for building web applications. Annotations like `@Controller`, `@RequestMapping`, and `@ResponseBody` simplify the development of RESTful APIs and web services.

TAP “stateless” version (1/2)

The advantages of a stateless server that we want for our TAP are:

- **Scalability:** Stateless servers are inherently easier to scale horizontally because they don't store any session state locally. This means you can add more instances of the server to handle increased load without worrying about synchronizing session state between instances.
- **Fault Tolerance:** Since there's no session state stored on the server, if a server instance fails, another instance can seamlessly take over without affecting the user experience. This improves fault tolerance and resilience in the system.

TAP “stateless” version (2/2)

The existing TAP service does not operate as a stateless service. Certain information must be kept and synchronized across all instances:

- Unique identifiers: Over all the different nodes of the cluster
- Job Metadata: State information and job results need to be consistently shared among all instances.
- User Quota: A stateless service requires shared access to user quota information from different nodes in the system to enforce defined restrictions regardless the instance being used
- Session handling, shared filesystem, review database write access (transactions), UWS events, ...

A related but different issue is to connect more than one TAP service to the same database to avoid data duplication and allow to share job information between TAPs

Managing different databases

- Reasoning: Is not always possible to provide a common solution for all our archives
- Current databases supported:
 - PostgreSQL (9-14): A powerful, open-source object-relational database
 - Greenplum: A massively parallel processing (MPP) database management system designed for analytics and data warehousing based on open-source PostgreSQL
- Take into account the extensions: q3c, PgSphere, ...
- Requirements:
 - Transparent for the developer using our library
 - Optimized code for all the configurations
 - Integration tests covering all different configurations

TAP Authentication & Authorisation

- Define public, private and shared tables.
- Logic implemented in the server -> makes it a bit more complex
 - tap_schema.tables does not provide the full picture: It shows the public tables but not the ones also accessible to the logged user
- Implement row level authorisation in a table
 - On-the-fly update of the resulting SQL sentence
- User groups/roles, one user included in one or more groups
 - Independent from database user and privileges

Thank you for your attention

