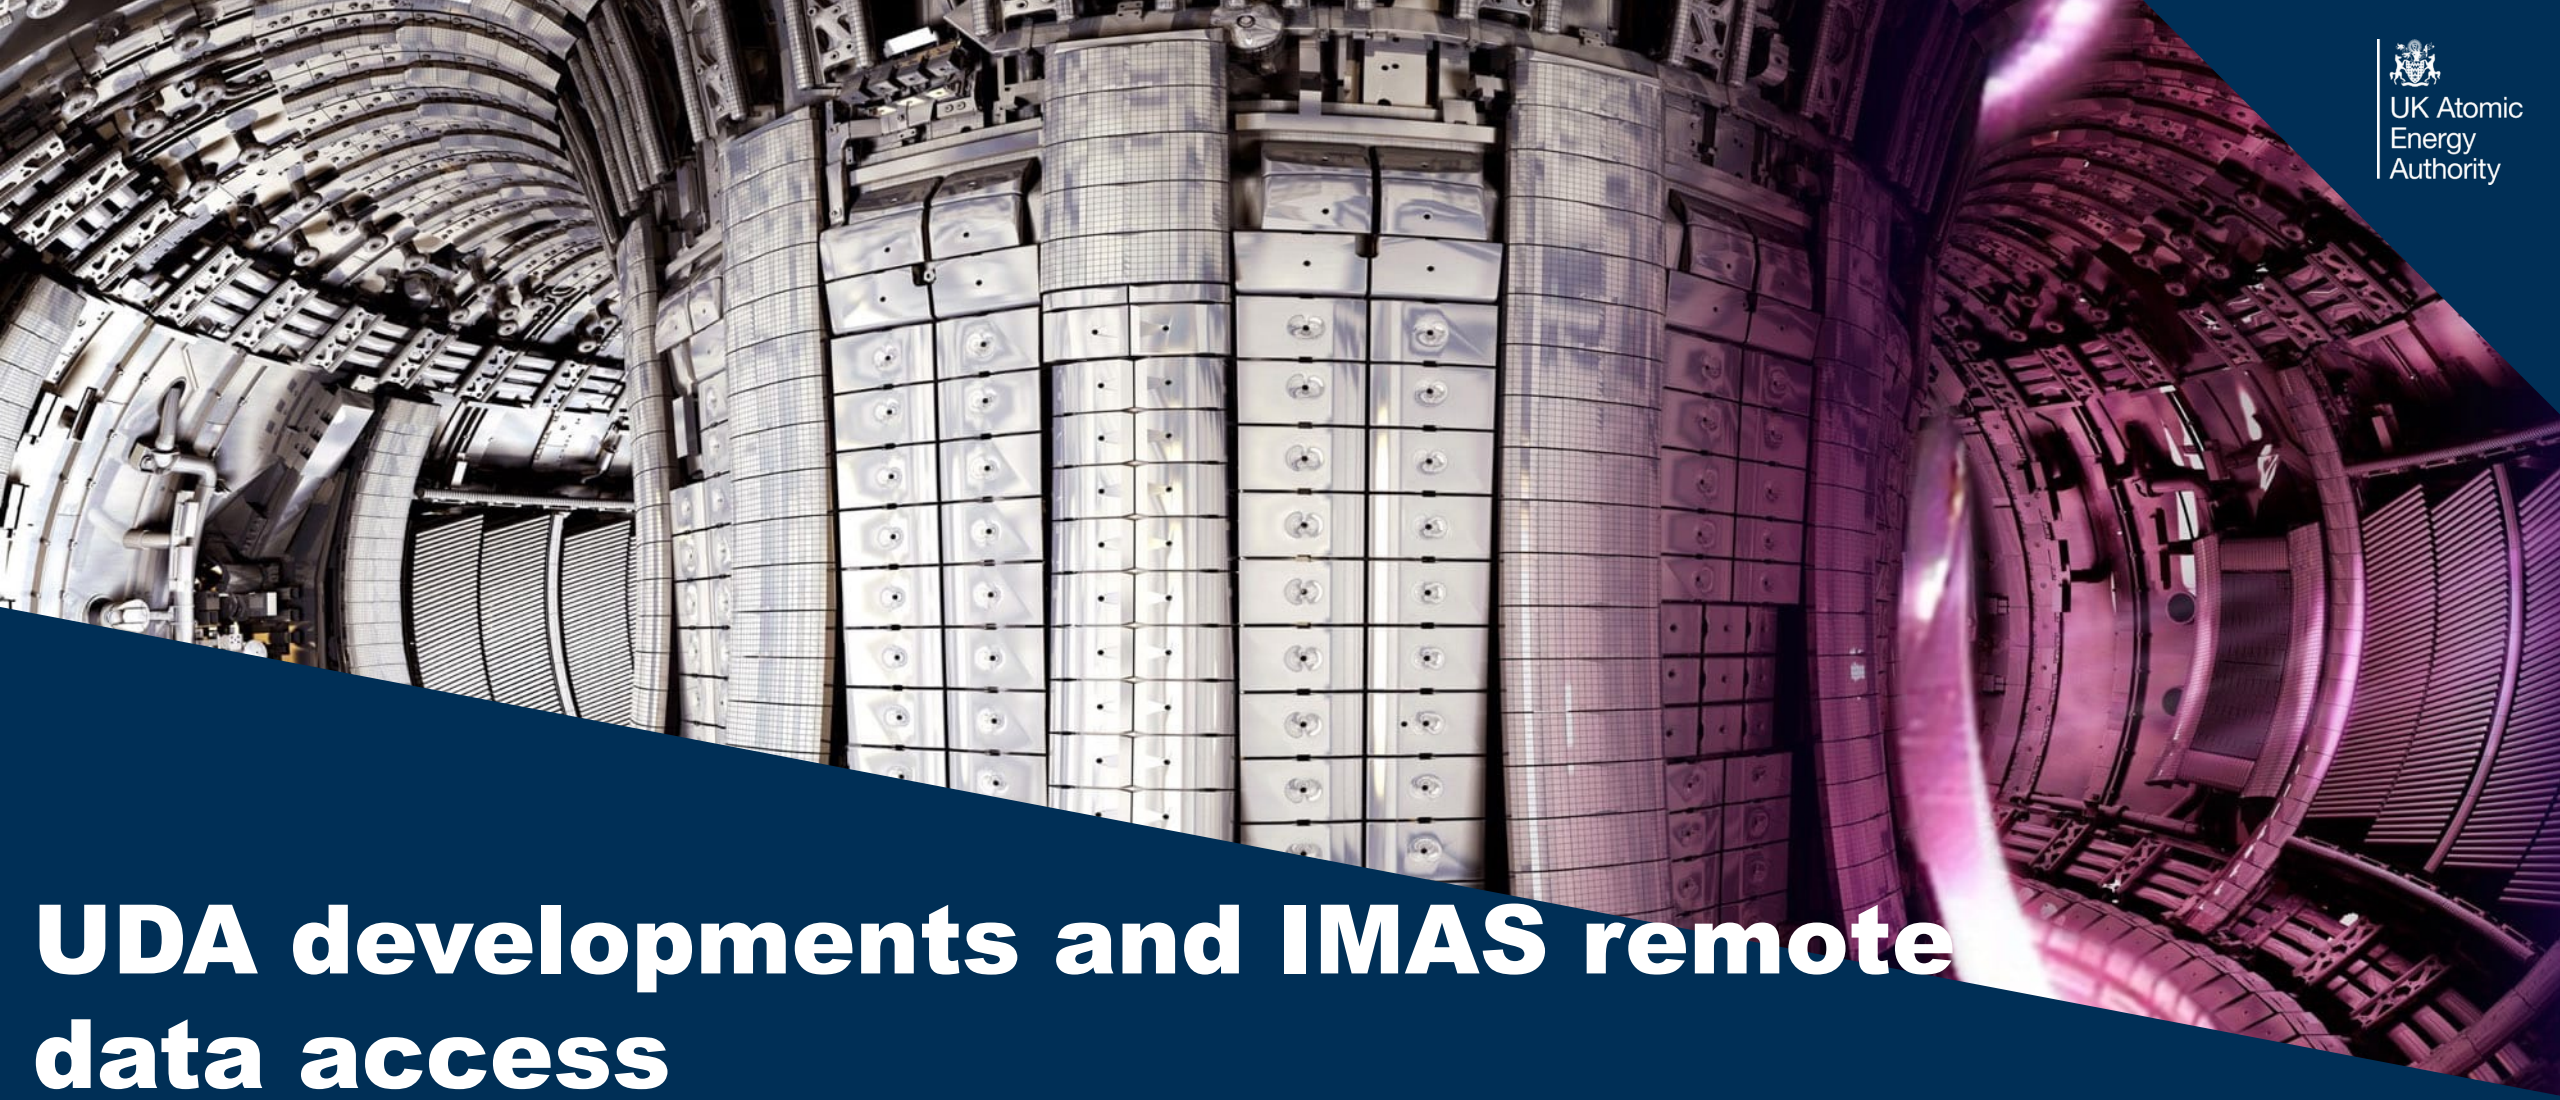

UK Atomic Energy Authority

UDA developments and IMAS remote data access

Jonathan Hollocombe (UKAEA)

DMP implementation working meeting 2023/05/12

iter

china eu india japan korea russia usa

Overview

UK Atomic Energy Authority

UDA developments and IMAS remote data access work done as part of ITER simulation catalogue development. Project finished March 2023 but support continuing under ITER Implementing Agreement with UKAEA.

- UDA developments
- UDA backend for AL5
- ITER server status
- Future developments

UDA developments

UK Atomic Energy Authority

- Open sourcing of UDA
- Batch request functionality
- SSL authentication
- Cap’n Proto serialization
- Other developments

Open sourcing of UDA

UK Atomic Energy Authority

- UDA is now available at https://github.com/ukaea/uda
- Open sourced under Apache-2.0 licence
- CI moved to GitHub actions
- Documentation being moved to https://ukaea.github.io/UDA (work in progress)

Batch request functionality

UK Atomic Energy Authority

- Each UDA get(signal, source) request requires a round trip to the server, returning a DataBlock containing the result of the request
- For thousands of request (as generated by the IMAS UDA backend) this adds significant latency overhead
- Add a new get_batch(signals, sources)
- Allows for multiple requests to be sent at once, returning a list of DataBlocks, each one corresponding to a processed request

SSL authentication

UK Atomic Energy Authority

- To allow for remote access to ITER data an authenticated connection to UDA is required
- Functionality existed in UDA to use SSL certificates but wasn't fully implemented
- UDA can now be compiled with –DSSLAUTHENTICATION=ON CMake flag to enable SSL authentication
- Server is configured with SSL certificate, CA certificate and CRL
- Client certificates available for ITER from https://pkiuda.iter.org

Cap’n Proto serialisation

UK Atomic Energy Authority

- A new serialisation option was added to UDA to return hierarchical data structures
- Cap’n Proto (https://capnproto.org/) was chosen as an open source and performant option
- Additional option for plugins to use, does not change existing serialisation technology (XDR)
- New serialisation and deserialisation routines provided in the server which are used in the new IMAS plugin

Other developments

UK Atomic Energy Authority

- A UDA CLI has been created to simplify testing server status and printing result of requests to the terminal
- Continuing refactoring and rewriting efforts to improve performance and stability of UDA codebase

UDA backend for AL5

- URI for remote data
- New IMAS plugin
- Access to mapping plugins
- Local access of remote URIs

URI for remote data

UK Atomic Energy Authority

```
imas://<server>:<port>/uda?path=<path>&backend=<backend>
```

For example:

```
imas://uda.iter.org/uda?
path=/work/imas/shared/imasdb/ITER_MD/3/1180/17&backend=hdf5
```

This causes the following entry to be opened on the uda.iter.org server:

```
imas:hdf5?path=/work/imas/shared/imasdb/ITER_MD/3/1180/17
```

New IMAS plugin

UK Atomic Energy Authority

- IMAS remote data plugin has been rewritten for AL5
- Replaces old plugins IMAS_REMOTE, IMAS_MAPPING, etc.
- Start of request (i.e. get or getSlice) causes a pre-fetch of data
- Makes use of UDA get_batch functionality to reduce round trips to server
- Uses new Cap'n Proto serialisation to reduce serialisation overhead
- When use partial_get only the required data is pre-fetched

New IMAS plugin (cont.)

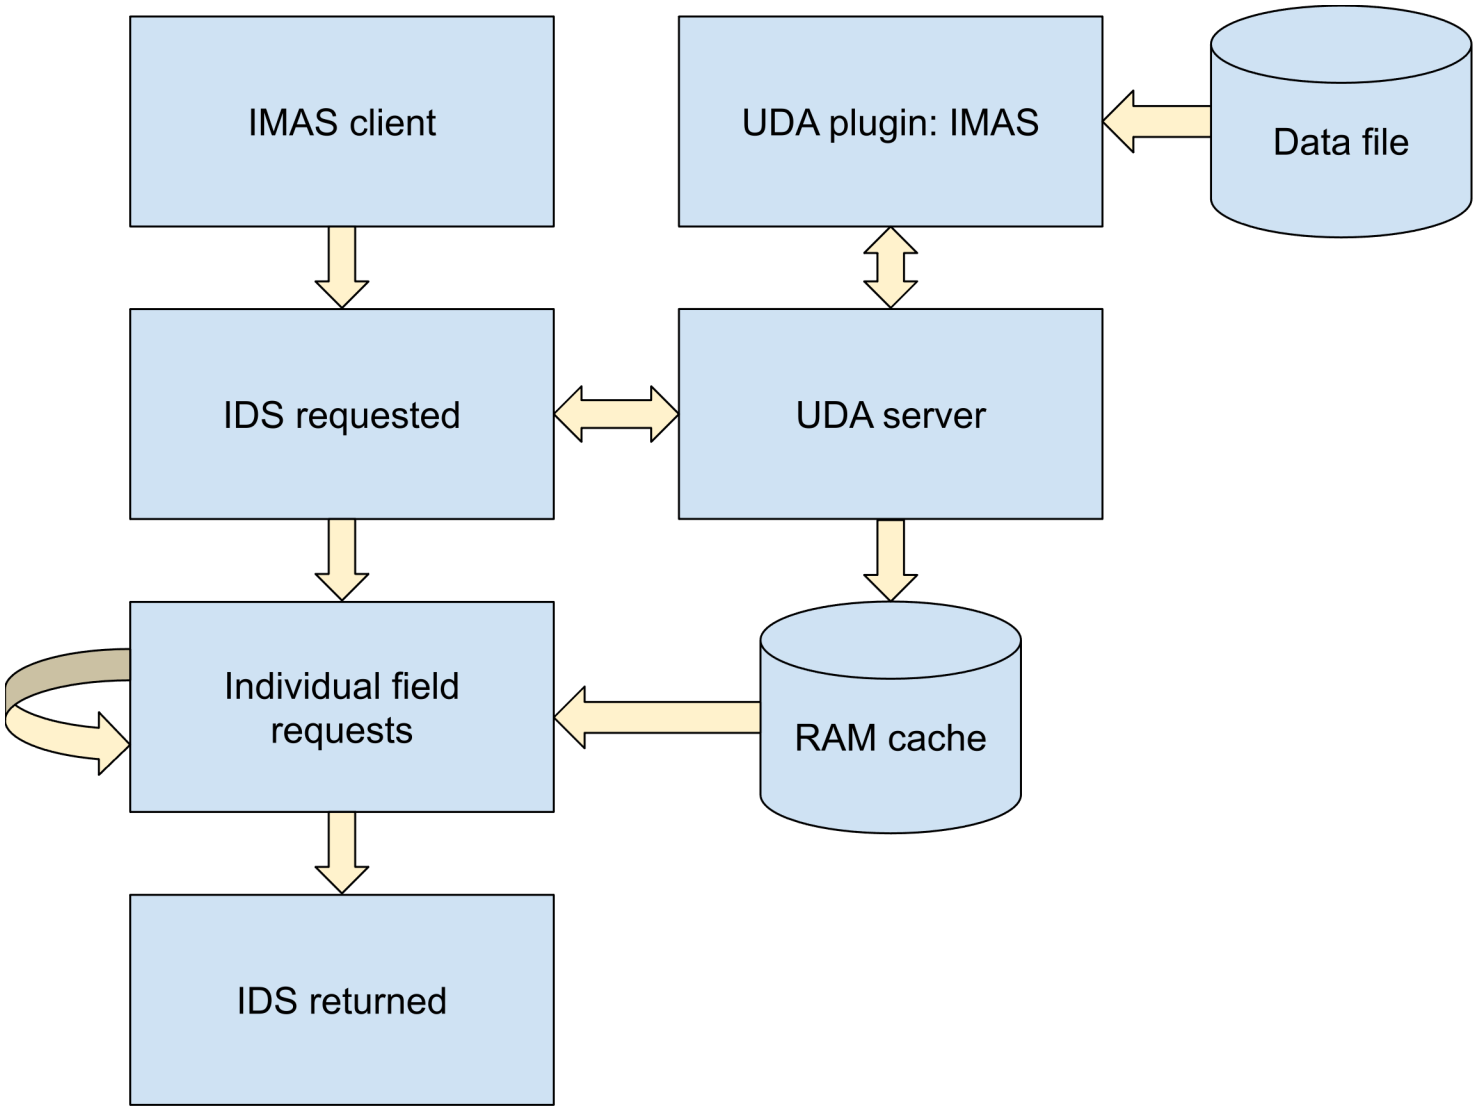

Local access of remote URIs

UK Atomic Energy Authority

- When accessing an IDS using a URI, the data may be available locally so would be inefficient to go via the UDA server
- When the ”server” element of the URI is found in the IMAS_LOCAL_HOST environmental variable then the local backend will be used instead of the UDA backend
- For example, if io-ls-uda01.iter.org is found in $IMAS_LOCAL_HOSTS:

```
imas://io-ls-uda01.iter.org/uda?
path=/work/imas/shared/imasdb/ITER_MD/3/1180/17&backend=hdf5
```

Will use the hdf5 backend and access the data directly, as if using URI

```
imas:hdf5?path=/work/imas/shared/imasdb/ITER_MD/3/1180/17
```

Access to mapping plugins

UK Atomic Energy Authority

Access to mapped experimental data is available using the new IMAS plugin by using mapped URI format:

```
imas://<server>:<port>/uda?machine=<machine>&<machine-specific-args>
```

For example:

```
imas://data.mastu.ukaea.uk/uda?machine=MASTU&shot=45125&pass=0
```

This will pass the request onto the new JSON plugin on the UDA server specified.

ITER server status

UK Atomic Energy Authority

- UDA server installed and running on io-ls-uda01.iter.org (external alias uda.iter.org)
- Internal server on port 56565 – this does not use SSL authentication but only accessible from SDCC
- External server on port 56563 but mapped to external port 56565 – this uses SSL authentication and requires client certificate from pkiuda.iter.org
- Access to uda.iter.org is being opened to all IPs – currently only accessible from whitelisted IPs
- Access has been validated on all IDSs currently available to AL5

Future developments

UK Atomic Energy Authority

- All SimDB URIs will be updated to use AL5 URIs
- Authenticated access to IMAS data will be tested and demonstrated at various sites
- More technical details will be disseminated at IAEA workshop in Ghent and in technical report
- Integration with new mapping plugin to provide mapped data for AL5
- UDA server and plugin training workshop to arrange