

# DMP Implementation Status of IPP

2024-06-26

# Since last meeting ...

- Repeated my talk from the previous meeting at the ITER Data Mapping monthly meeting
  - Provided them with a copy of `uda.py` (via the repository)
- Also gave ITER access to the comparison tools I had mentioned at the 2nd meeting  
(<https://confluence.iter.org/pages/viewpage.action?pageId=613144222>) (via the repository <https://gitlab.eufus.psnc.pl/ets/ids-compare>)
- Updated the list of AUG DMP shots on the Wiki (a missed “/0” in all of the entries) at <https://wiki.euro-fusion.org/images/3/3e/Uris.txt>
- Further development of `uda.py`

# Discovered `imas.DBEntry()` cannot communicate problems

- Raised the issue on slack

- Maarten Sebregts 10:40 AM

- Hi David, this has been addressed with <https://jira.iter.org/browse/IMAS-4835> and will be part of the upcoming AL-Core 5.2.4 release.

- Engineered my way around this problem

- `subprocess.run` with a timeout set and output captured

DBEntry	get	close	Timeout	URI
1.51	7.672	0.038		<code>imas://uda.ipp.mpg.de:56565/uda?path=/root/public/imasdb/test/3/1000/1000&amp;backend=hdf5</code>
1.529	6.386	0.032		<code>imas://uda.ipp.mpg.de:56565/uda?path=/root/public/imasdb/aug-dmp/3/41570/0&amp;backend=hdf5</code>
1.584	7.144	0.035		<code>imas://uda.ipp.mpg.de:56565/uda?path=/root/public/imasdb/aug/3/17151/1&amp;backend=hdf5</code>
1.507	7.124	0.035		<code>imas://uda.ipp.mpg.de:56565/uda?path=/root/public/imasdb/aug/3/17151/3&amp;backend=hdf5</code>
1.343	8.608	0.041		<code>imas://data.mastu.ukaea.uk:56560/uda?mapping=MAST&amp;path=/&amp;verbose=0&amp;shot=30417&amp;batch_size=20</code>
			10.042	<code>imas://spcimasdata.epfl.ch:443/uda?path=/data/imas/public/imasdb/tcv/3/61010/1/&amp;backend=mdsplus</code>
			10.04	<code>imas://spcimasdata.epfl.ch:443/uda?path=/data/imas/public/imasdb/tcv/3/68340/2/&amp;backend=mdsplus</code>
2.776	6.842	0.035		<code>imas://irfm-uda.partenaires.cea.fr:8083/uda?path=/Imas_public/imas_simulation/public/imasdb/west_simu_preparation/3/57410/0&amp;backend=hdf5</code>

# SSLs for ITER now working

- “it was indeed a problem of outdated CRL on our server, it was updated and now access from outside shall work again (for 1 month before it expires again... so we have a few weeks to automate the process)”
  
- `setenv UDA_CLIENT_SSL_KEY ~/.uda/private.key`
- `setenv UDA_CLIENT_CA_SSL_CERT ~/.uda/ca-server-certificate.pem`
- `setenv UDA_CLIENT_SSL_CERT ~/.uda/certificate.pem`
- `setenv UDA_CLIENT_SSL_AUTHENTICATE 1`
- `./uda.py -u`  
`'imas://uda.iter.org/uda?path=/work/imas/shared/TEST/simulations/test/ed5a7c6a8e0211ee8f020894ef442712/&backend=mdsplus'`
- Accessing data from  
`imas://uda.iter.org/uda?path=/work/imas/shared/TEST/simulations/test/ed5a7c6a8e0211ee8f020894ef442712/&backend=mdsplus`
- Available IDses are: ['core\_profiles', 'core\_sources', 'core\_transport', 'dataset\_description', 'equilibrium', 'summary', 'transport\_solver\_numerics']
  
- Timing information
  - `DBentry = 4.561`
  - `get = 15.502`
  - `close = 0.035`

# Some timing information: data at ITER

- In all cases the command executed is
  - `./uda.py -u 'imas://uda.iter.org:56565/uda?path=/work/imas/shared/imasdb/ITER/3/134174/117&backend=mdsplus' -c summary`
- At ITER
  - With bypass enabled (`$IMAS_LOCAL_HOSTS == uda.iter.org`)
    - DBEntry = 1.012      DBEntry = 0.967      DBEntry = 0.972
    - get = 0.488      get = 0.125      get = 0.135
    - close = 0.002      close = 0.001      close = 0.002
  - With bypass disabled (`$IMAS_LOCAL_HOSTS == ""`)
    - DBEntry = 1.421      DBEntry = 1.325      DBEntry = 1.370
    - get = 94.286      get = 83.453      get = 80.465
    - close = 0.005      close = 0.004      close = 0.003
- At Gateway
  - DBEntry = 1.345      DBEntry = 1.384      DBEntry = 1.271
  - get = 96.223      get = 95.684      get = 100.322
  - close = 0.038      close = 0.036      close = 0.037

# Some timing information: data at ITER

	from ITER using bypass			from ITER using UDA			from gateway using UDA		
	DBentry	get	close	DBentry	get	close	DBentry	get	close
	1.012	0.488	0.002	1.421	94.286	0.005	1.345	96.223	0.038
	0.967	0.125	0.001	1.325	83.453	0.004	1.384	95.684	0.036
	0.972	0.135	0.002	1.37	80.465	0.003	1.271	100.322	0.037
mean	0.984	0.249	0.002	1.372	86.068	0.004	1.333	97.410	0.037
std dev	0.025	0.207	0.001	0.048	7.272	0.001	0.057	2.537	0.001
ITER UDA/Direct	1.39	<b>345.19</b>	2.40						
UDA GW/ITER	0.97	<b>1.13</b>	9.25						

- For the most time consuming part (at least when using UDA) of getting the data
  - **Direct access is more than 300 times faster**
  - **Remote access occurs a penalty of about 13% compared to local UDA access**

# Some timing information: data at ITER (larger case)

- `scripts/idssize -u 'imas:mdsplus?path=/work/imas/shared/imasdb/ITER/3/134174/117'`
- Reading 16.909 MB of data for `core_profiles/0` took 1.00 seconds
- Reading 30.368 MB of data for `core_sources/0` took 2.56 seconds
- Reading 24.921 MB of data for `core_transport/0` took 1.50 seconds
- Reading 337.822 MB of data for `edge_profiles/0` took 79.72 seconds
- Reading 153.546 MB of data for `edge_sources/0` took 76.50 seconds
- Reading 85.356 MB of data for `edge_transport/0` took 81.46 seconds
- Reading 48.138 MB of data for `equilibrium/0` took 2.01 seconds ←
- Reading 0.061 MB of data for `summary/0` took 0.14 seconds
- Processing... \_\_\_\_\_ 100% 0:04:40
- Total reading time = 244.90 s
- Total data size = 697.1 MB
- Fractions of the total size for `imas:mdsplus?path=/work/imas/shared/imasdb/ITER/3/134174/117`
- % bytes    IDS
- 2.43 %    `core_profiles/0`
- 4.36 %    `core_sources/0`
- 3.57 %    `core_transport/0`
- 48.46 %    `edge_profiles/0`
- 22.03 %    `edge_sources/0`
- 12.24 %    `edge_transport/0`
- 6.91 %    `equilibrium/0`
- 0.01 %    `summary/0`

# Some timing information: data at ITER (larger case)

- `./uda.py -u 'imas://uda.iter.org:56565/uda?path=/work/imas/shared/imasdb/ITER/3/134174/117&backend=mdsplus' -c equilibrium`
- From ITER
- with `$IMAS_LOCAL_HOSTS == uda.iter.org`
  - DBentry = 1.029
  - get = 0.794
  - close = 0.000
- with `$IMAS_LOCAL_HOSTS == "`
  - DBentry = 1.591
  - get = 16284.980
  - close = 0.003
- From Gateway
- Timing information
  - DBentry = 2.039
  - get = 24115.763
  - close = 0.036
  
- UDA (local) / local : **20510.05**
- UDA (Gateway) / UDA (local) : **1.481**



# Trying logrotate on the docker UDA VMs

- On the VMs
  - `dnf install logrotate`
  - `cat <<EOF > /etc/logrotate.d/uda`
  - `/opt/uda/etc/xinetd.logfile`
  - `/opt/uda/etc/Access.log`
  - `/opt/uda/etc/Error.log`
  - `/opt/uda/etc/Debug.log {`
  - `missingok`
  - `notifempty`
  - `rotate 53`
  - `compress`
  - `delaycompress`
  - `weekly`
  - `create`
  - `dateext`
  - `}`
  - `EOF`
- On the docker host, add crontab entries
  - `0 1 * * * docker exec -it uda-server-with-ssl logrotate /etc/logrotate.d/uda`
  - `0 1 * * * docker exec -it uda-server-without-ssl logrotate /etc/logrotate.d/uda`