

DMP Implementation Status of IPP

2024-12-18

Testing the effect of “batch_size”: JET

- `./uda.py -u 'imas://uda.jetdata.eu:56565/uda?mapping=JET&verbose=0&shot=99399&batch_size=20' -c summary`
- `batch_size=1` Protocol 11 Error (Server Block #2)
- `batch_size=2` get = 58.567
- `batch_size=5` get = 29.951
- `batch_size=10` get = 19.043
- `batch_size=20` get = 15.444
- `batch_size=50` get = 12.017
- `batch_size=100` get = 10.326 ← **Best**
- `batch_size=200` get = 11.837
- `batch_size=500` get = 11.831
- `batch_size=1000` get = 11.190
- `batch_size=2000` get = 10.853
- `batch_size=5000` get = 11.280
- `batch_size=10000` get = 11.024
- `batch_size=20000` get = 10.710

Testing the effect of “batch_size”: AUG

- `./uda.py -u 'imas://uda.ipp.mpg.de:56565/uda?path=/root/public/imasdb/aug-dmp/3/41570/0&backend=hdf5&batch_size=100' -c summary`

- (default) get = 11.896 get = 10.118
- `batch_size=1` get = 165.911 get = 142.639
- `batch_size=10` get = 15.788 get = 22.492
- `batch_size=50` get = 5.932 get = 6.810
- `batch_size=100` get = 4.225 get = 3.476 ← **Best**
- `batch_size=200` get = 5.688 get = 5.222
- `batch_size=500` get = 3.859 get = 4.927
- `batch_size=1000` get = 4.398 get = 4.216
- `batch_size=10000` get = 6.069 get = 6.534

Suggest changing the default `batch_size` to 100

Testing the effect of “batch_size”: TCV

- `./uda.py -u 'imas://spcimasdata.epfl.ch:443/uda?path=/data/imas/public/imasdb/tcv/3/61010/1/&backend=mdsplus' -c summary`

● (default)	get = 4.272	get = 4.112	
● batch_size=1	get = 60.900	get = 64.296	
● batch_size=10	get = 7.285	get = 7.164	
● batch_size=50	get = 2.336	get = 2.337	
● batch_size=100	get = 1.852	get = 1.956	← Best
● batch_size=200	get = 2.101	get = 2.076	
● batch_size=500	get = 1.923	get = 1.937	
● batch_size=1000	get = 1.980	get = 2.014	

Suggest changing the default batch_size to 100

Testing the effect of “batch_size”: AUG

- `time idscp -s`
`'imas://uda.ipp.mpg.de:56565/uda?path=/root/public/imasdb/aug/3/32408/1&backend=hdf5' -d`
`"imas:hdf5?path=/tmp/aug_3_32408_1&backend=hdf5" -a`
- (default) 5.190u 4.161s 6:27.20 5.526u 5.204s 7:22.75
- `batch_size=100` 5.100u 4.897s 6:42.37 4.804u 4.550s 6:28.75
- `batch_size=200` 5.207u 4.598s 6:10.05 5.318u 5.547s 7:29.70
- `batch_size=500` 5.405u 4.602s 5:42.02 5.550u 4.525s 5:23.70
- `batch_size=1000` 5.087u 4.886s 7:22.39 5.270u 5.990s 8:48.58

Suggest changing the default `batch_size` to 100, possibly 500