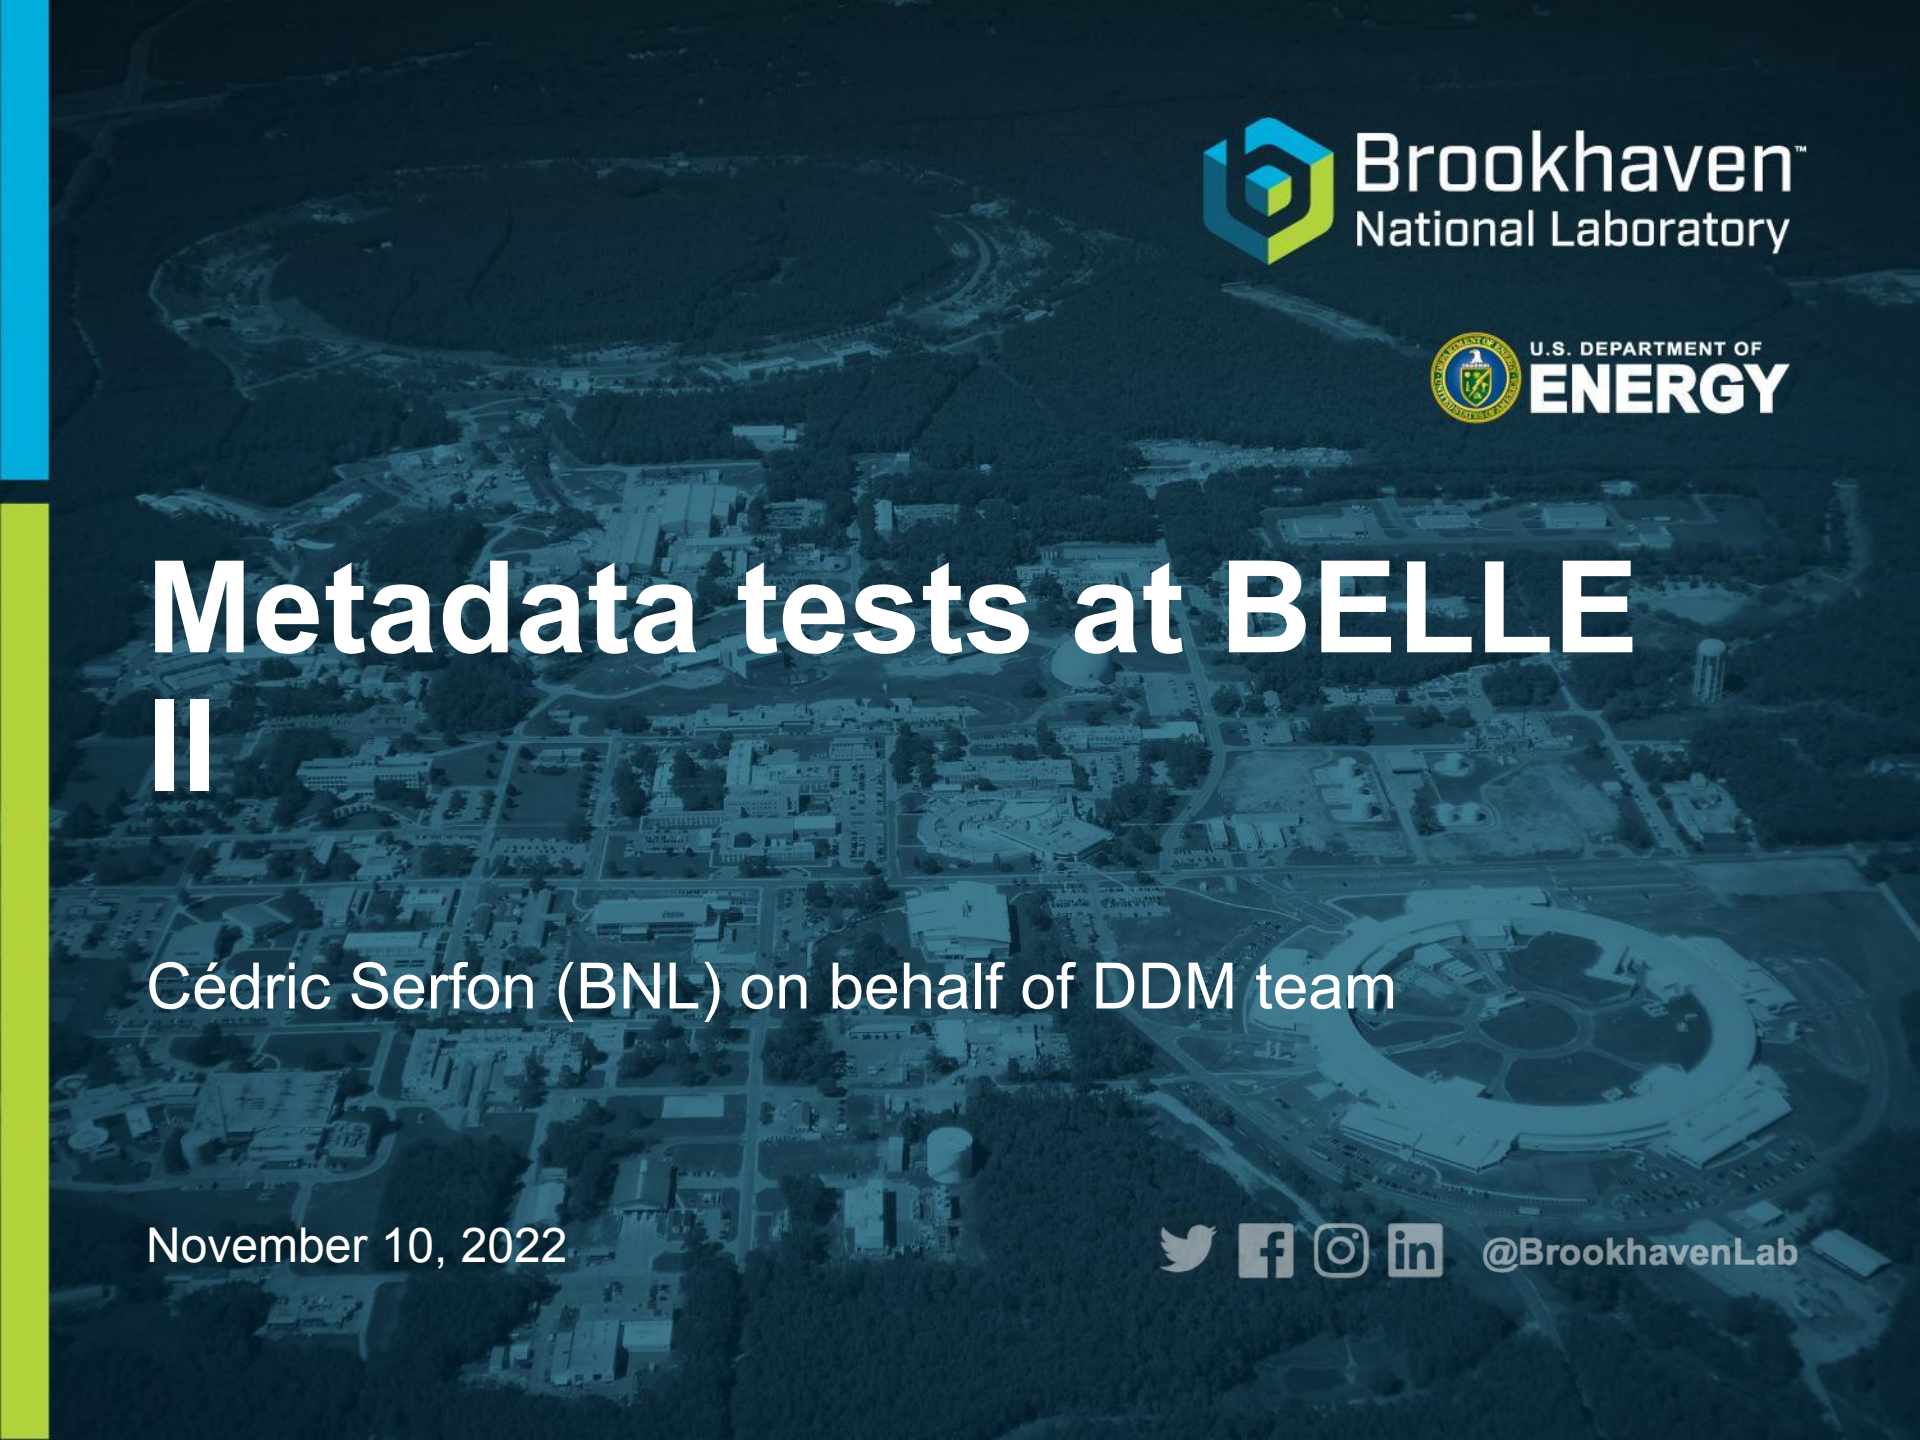


# Metadata tests at BELLE II

Cédric Serfon (BNL) on behalf of DDM team

November 10, 2022

# Introduction

- Belle II already uses column metadata to fill accounting Dashboard

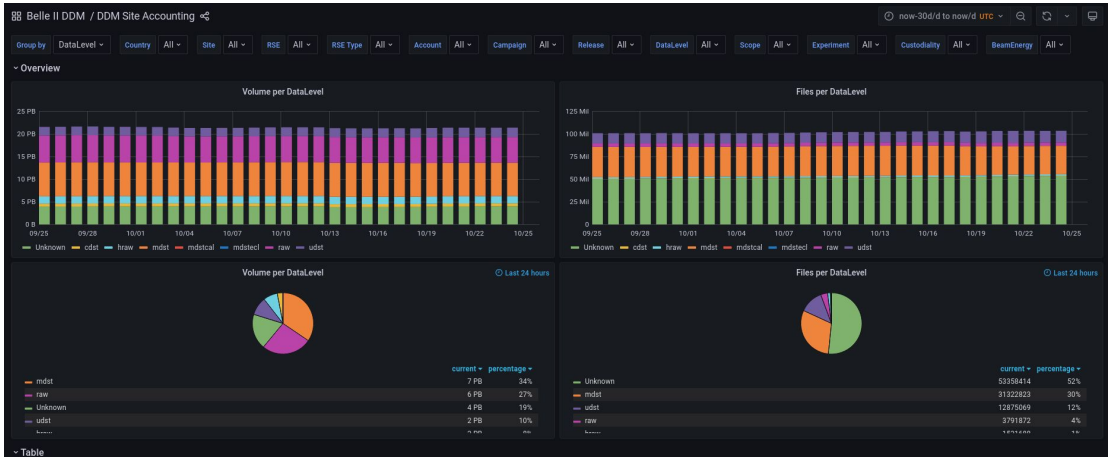


- Now interested in testing Rucio generic metadata
- This talk will details some of the developments and test performed during the last year

# Hierarchical namespace

- Belle II uses a hierarchical namespace
- There are different metadata at each level of the hierarchy

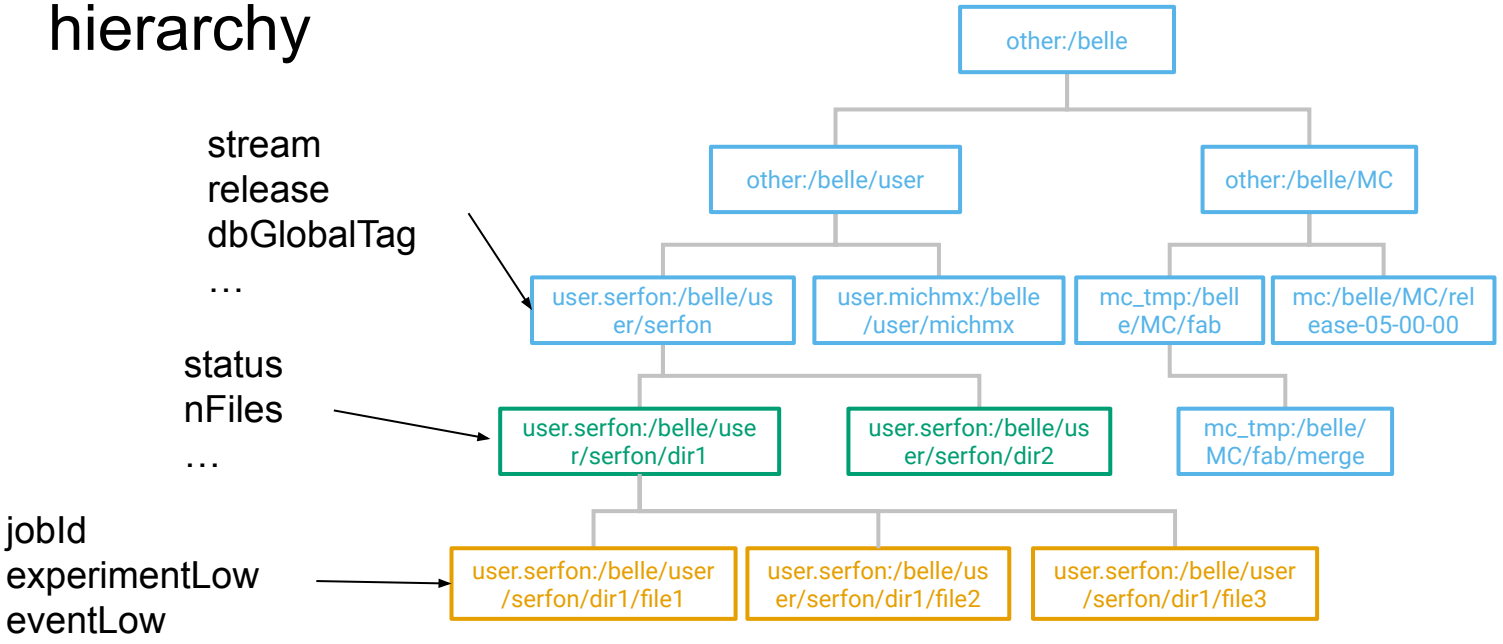


File → File

Directory containing files → Dataset

Directory containing directories → Container

# New features in Rucio 1

- In Dirac File Catalog (DFC), subdirs inherit the metadata of their parents but not Rucio File Catalog (RFC)
- New inherit option added to get_metadata_bulk :
  - When set to True, aggregate the metadata from the DIDs and all their parents/grandparents…

```
>>> from rucio.client import Client
>>> client = Client()
>>> for meta in client.get_metadata_bulk([{'scope': 'mc', 'name': '/belle/MC/release-05-02-00/DB00001330/MC14ri_d/prod00021640/s00/e1003/4S/r00000/charged/mdst
/sub00/mdst_000283_prod00021640_task10020000284.root'}]): print(meta)
...
{'scope': 'mc', 'name': '/belle/MC/release-05-02-00/DB00001330/MC14ri_d/prod00021640/s00/e1003/4S/r00000/charged/mdst/sub00/mdst_000283_prod00021640_task100200
00284.root', 'account': 'diracprd', 'did_type': 'FILE', 'is_open': None, 'monotonic': False, 'hidden': False, 'obsolete': False, 'complete': None, 'is_new': Fa
lse, 'availability': 'AVAILABLE', 'suppressed': False, 'bytes': 1919280007, 'length': None, 'md5': None, 'adler32': 'eaf6fbaf', 'expired_at': None, 'purge_repl
icas': True, 'deleted_at': None, 'events': 189000, 'guid': '608808b756f7563ff96d9f0dd2c0a4ef', 'project': '4S', 'datatype': 'mdst', 'run_number': 1003, 'stream
_name': None, 'prod_step': None, 'version': 'release-05-02-00', 'campaign': 'MC14ri_d', 'task_id': None, 'panda_id': None, 'lumiblocknr': None, 'provenance': N
one, 'phys_group': None, 'transient': False, 'accessed_at': None, 'closed_at': None, 'eol_at': None, 'is_archive': None, 'constituent': None, 'access_cnt': Non
e, 'created_at': datetime.datetime(2021, 10, 22, 20, 58, 16), 'updated_at': datetime.datetime(2022, 9, 29, 19, 31, 13)}
>>> for meta in client.get_metadata_bulk([{'scope': 'mc', 'name': '/belle/MC/release-05-02-00/DB00001330/MC14ri_d/prod00021640/s00/e1003/4S/r00000/charged/mdst
/sub00/mdst_000283_prod00021640_task10020000284.root'}], inherit=True): print(meta)
...
{'scope': 'mc', 'name': '/belle/MC/release-05-02-00/DB00001330/MC14ri_d/prod00021640/s00/e1003/4S/r00000/charged/mdst/sub00/mdst_000283_prod00021640_task100200
00284.root', 'date': datetime.datetime(2021, 10, 22, 20, 56, 54), 'site': 'LCG.KEK2.jp', 'size': 1919280007, 'jobId': 210927648, 'runLow': 0, 'status': 'good',
 'nEvents': 189000, 'runHigh': 0, 'checksum': 'eaf6fbaf', 'eventLow': 4419002, 'eventHigh': 53829001, 'parentGuids': '', 'checksumType': 'Adler32', 'experiment
Low': 1003, 'experimentHigh': 1003, 'mc': 'MC14ri_d', 'owner': 'g:belle_mcprod', 'logLfn': None, 'nFiles': 286, 'stream': 0, 'dataset': '/belle/MC/release-05-0
2-00/DB00001330/MC14ri_d/prod00021640/s00/e1003/4S/r00000/charged/mdst', 'release': 'release-05-02-00', 'dataType': 'mc', 'dataLevel': 'mdst', 'beamEnergy': '4
S', 'lastUpdate': datetime.datetime(2021, 10, 24, 0, 41, 19), 'sourceCode': None, 'dbGlobalTag': 'DB00001330', 'description': 'MC14ri_d (Jira: BIIDP-4802) - Ea
rly phase 3 Y(4S) charged BBbar (BGx1), 100/fb equivalent', 'mcEventType': 'charged', 'creationDate': datetime.datetime(2021, 10, 19, 12, 31, 30), 'productionI
d': 21640, 'steeringFile': 'MC/MC14/MC14ri_d/release-05-02-00/DB00001330/4S/charged/charged_eph3.py', 'skimDecayMode': None, 'parentDatasets': None, 'generalSk
imName': None, 'transformationId': 919020, 'sourceCodeRevision': None, 'steeringFileRevision': None, 'int_luminosity': 1000.0}
```

# New features in Rucio 2

- Introduction of set_dids_metadata_bulk → Done
- Possibility to archive metadata for deleted DIDs → In Review
- Possible new development : Introduce a new method in rucio.core.dirac to create new file + namespace hierarchy + metadata in an atomic way → To be discussed

# Implementation of metadata methods in BelleDirac

- Rucio API is never used directly by DIRAC. If we want to get/set/update/delete metadata these methods need to be implemented into RFC in (Belle)DIRAC
- New methods were introduced in BelleDIRAC :
  - getFileMetadata, getDirectoryMetadata, getFileUserMetadata, getFileUserMetadataBulk
  - setMetadataBulk, setMetadata
  - removeMetadata
- Most of these methods are similar to the ones implemented in DFC. Some extra ones (e.g. setMetadataBulk) were added

# Read/write tests

- Read/write tests were performed to check all the new methods and if our current infrastructure scales
- Snapshot of Belle II production instance was done and deployed on a test instance. About ~100M files imported
- Test consisted in submitting multiple jobs on BNL cluster that read and register random generic data at high frequency

# Read/write tests

2M / 3 hours ~ **185** metadata rows /s
or
Since the test writes 7 metadata per LFN.
it is 1.3KHz metadata /s


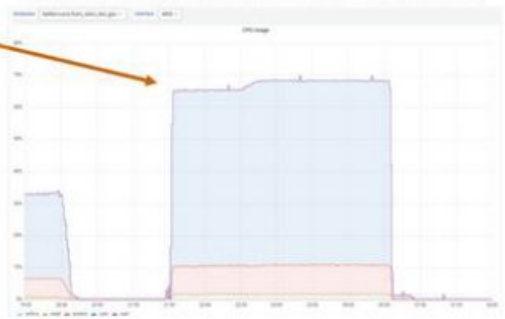




Only Apache CPU shows high load.
Easily remedied by more front end server.

**Hiro**

# Import tests

- Dump of all data stored in our Metadata service was done a few weeks ago
- Metadata then imported to Rucio :
  - All accounting metadata (5 metadata) imported as DID column metadata
  - Import of generic (i.e. json metadata) for ~1M files performed as a test. Size ~1 kB / file

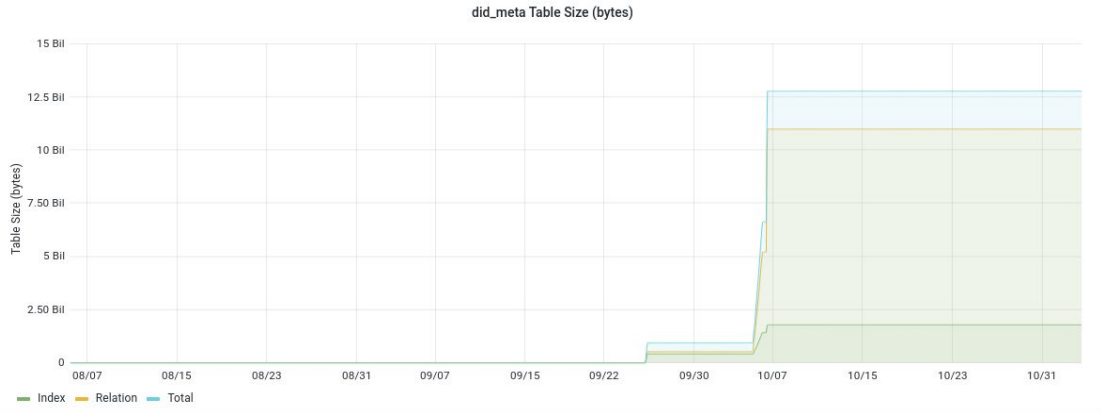


Brookhaven National Laboratory

# Conclusion

- A few new developments were done to serve Belle II metadata workflows
- The tests performed show that Rucio is able to deliver the performance needed for Belle II
- Next steps are being discussed internally