



Science and
Technology
Facilities Council

HEPscore in the accounting system: Development Status

Adrian Coveney

HEPscore Workshop, 20 September 2022



Science and
Technology
Facilities Council

Outline

1. Requirements
2. Existing workflow
3. External changes
4. Sample message extract
5. APEL client and server changes
6. Timeline

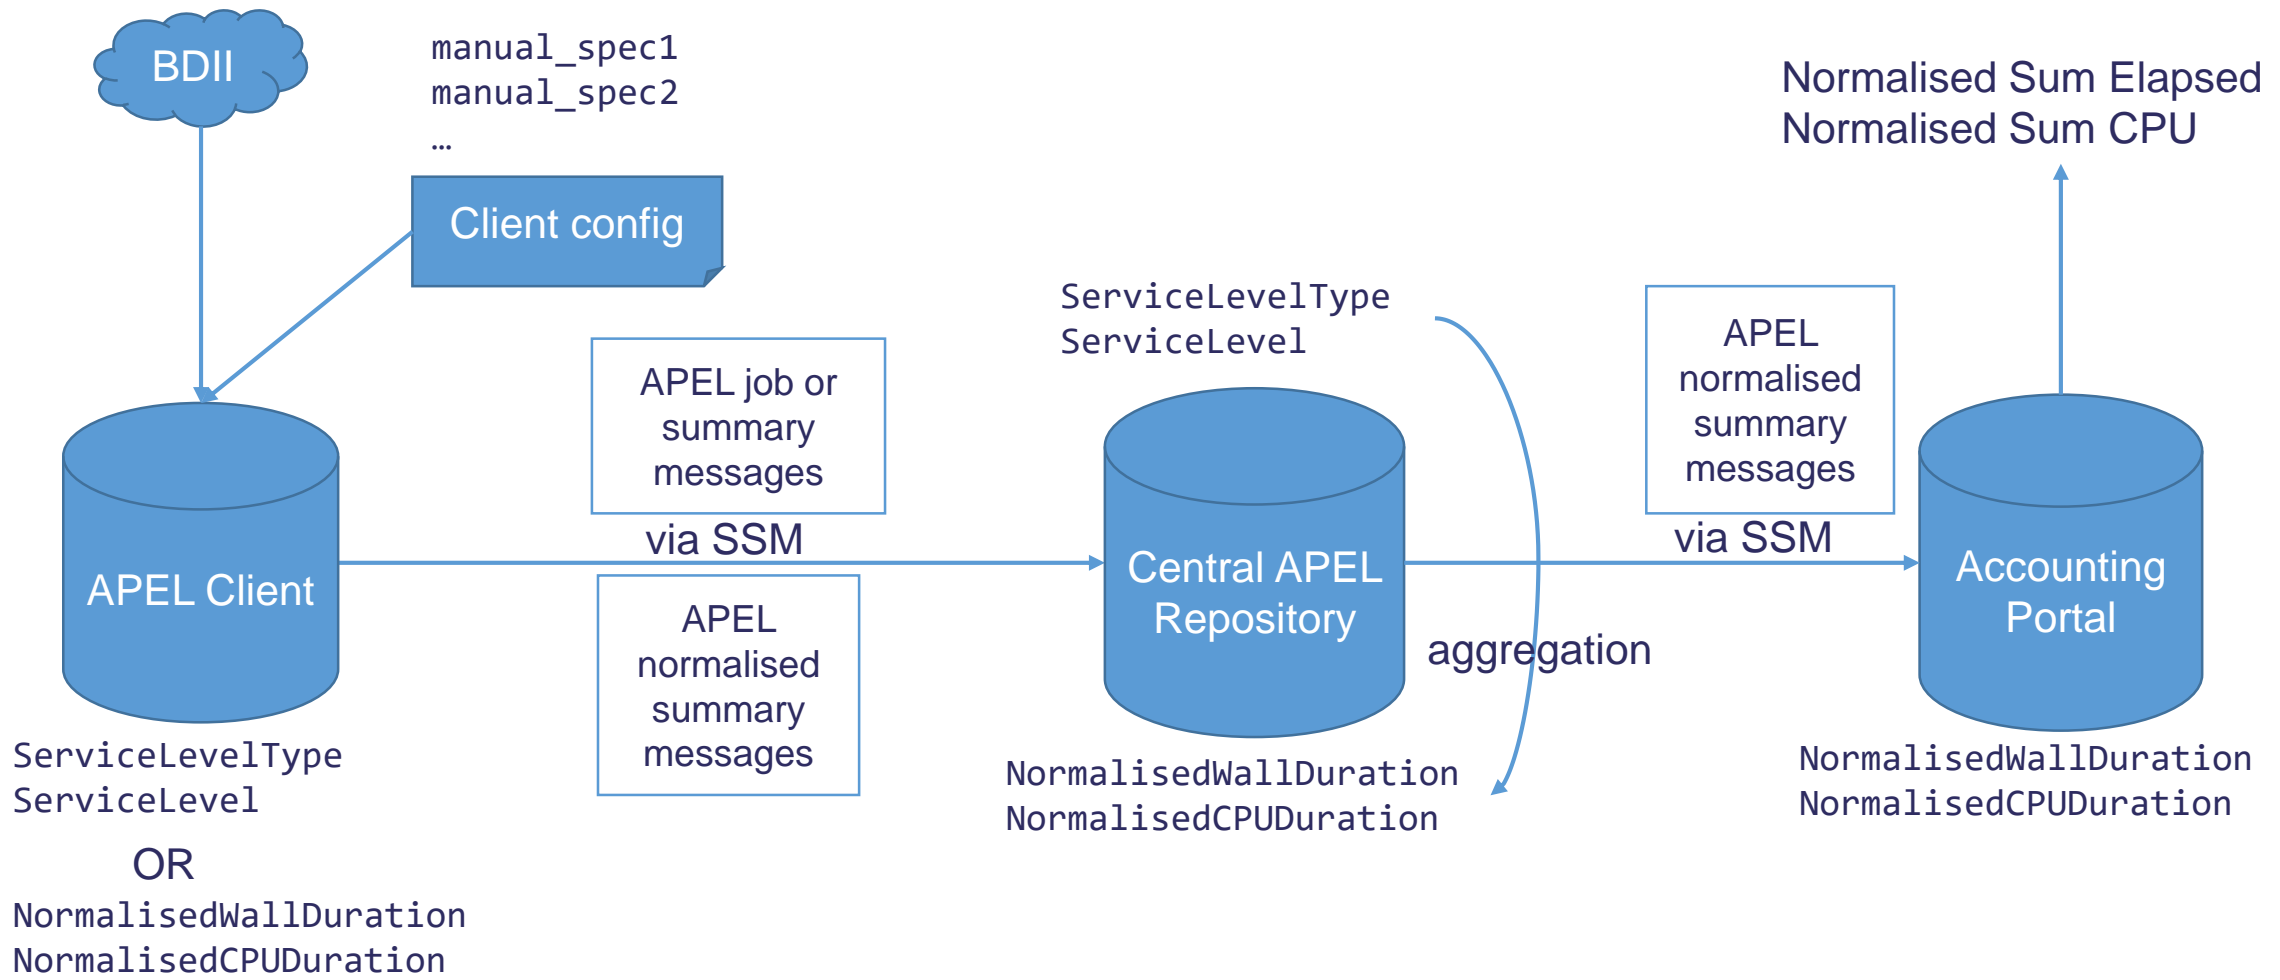
Requirements

- Support for new benchmark metrics in clients, message formats (job, summary, normalised summary), Repository and Portal - alongside the existing benchmarks
- Simplify any future changes in benchmark (including versioning)
- Method to query which sites are publishing the new benchmark (to support decision making on timing of switchover)
- Ability to display consumption depending on the benchmark selected in the Accounting Portal
- WLCG Accounting Utility should be modified to provide comparison of the consumption in both benchmarks

Existing workflow

- Most sites include benchmarks as part of APEL accounting records (job or summary)
- Many sites only provide final usage \times benchmark values in normalised summaries
- Benchmarks can be pulled in from BDIIIs or set manually in local APEL accounting client configuration
- Sites apply a scaling factor between execute nodes so that they can use a single benchmark per cluster

GLUE_CE_CAPABILITY
GLUE_HOST_BENCHMARK



External changes

- Update record messaging formats to allow multiple benchmarks or values normalised against multiple benchmarks
 - Only the one or two benchmarks that we're interested in for centralised accounting will be parsed (e.g. HEPSPEC06 and HEPscore). Any other values ignored
 - This gives the messaging format maximum flexibility but simplifies the internal changes
- Normalised accounting messages to include benchmark type that the records were normalised using
- 3rd-party clients (e.g. ARC, HTCCondor) will need to make changes to support the new messaging formats, though the old format will continue to be accepted

Sample message extracts

- Job records (similar in Summaries)
 - Accept an associative array of benchmarks
 - Only read benchmarks we're interested in for centralised accounting

```
APEL-individual-job-message: v0.4
...[etc.]
WallDuration: 47248
CpuDuration: 46871
ServiceLevel: {hepspec:11.4,
                HEPscore2X: 15.3}
%%
...[other records]
```

- Normalised summaries
 - Renamed header
 - Accept an associative array of normalised values

```
APEL-normalised-summary-message:
v0.4
...[etc.]
WallDuration: 47248
CpuDuration: 46871
NormalisedWallDuration: {hepspec:
                          519728, hepscore2X: 708720}
NormalisedCpuDuration: {hepspec:
                         515581, hepscore2X: 703065}
%%
...[other records]
```

APEL client and server changes

- Add fields to store an additional benchmark value and benchmark type in APEL client and server schemas for job, summary and normalised summary records
 - Server schema used by both Repository and Portal
 - Not as flexible as messaging record change but simpler to implement based on existing structure
- Update summarising process
- Allows direct comparisons to be made between old and new benchmarks if both supplied

Timeline

- Message format changes documented and being reviewed
- September/October Repository development to support change in record format
- Prototype Repository ready for testing in October – check message flows, record loading and aggregation with test site(s)
- November/December make changes as necessary following testing and start on Repository-WAU interface if that is required



Science and
Technology
Facilities Council

Thank you