

Accounting sensors and client important features

These are functionalities that the current DGAS client implements and are considered important by the DGAS community. These functionalities/use cases are therefore perceived to be important requirements for the Emi accounting client.

1. The Accounting client Transport Protocol must be defined as a project standard. It must be based on a commonly agreed messaging infrastructure. It must be possible for the accounting server to consume the record messages leveraging existing and commonly available APIs.
2. The accounting client must be able to send the Usage Records to multiple end points and Usage Records can be routed to different endpoints on the basis of their contents.
 - a. *Example: Usage Records of Local jobs submitted directly to the queue manager get archived in a Record Repository local to the site.*
 - b. *Example: Usage Record of WLCG VO Jobs get archived in Record Repository A while other VO records in record repository B.*
 - c. *Example: Usage Records get archived on both the Central Repository and in a Local or National one.*
 - d. *Example: A copy of the usage records produced on a given CE is digested by a monitoring system (E.G. HLRMon to generate graphical usage reports)*
3. Local usage must be accounted and differentiated from the grid accounting one.
4. Sensors must be capable to optionally map local user or group usage to a given VO. Mapping must be configurable by the system manager.
 - a. *Example: the site manager who gives local access to Users A AND B Belonging to a known VO maps their local uid, or gid, to that VO so that the information is available in the accounting repository.*
5. It must be possible for the system manager to trigger the production and republishing of Usage Records starting from a given date.
6. The accounting client must be capable to transport Usage Records generated by a third party sensor.

Usage Record

This is an analysis of the current APEL record format and a set of requirements on which fields in the EMI-CAR record must be filled in order for DGAS services to properly work.

APEL SSM Usage Record mappings

These fields are already present in the APEL SSM record. These should however be mapped in the corresponding EMI-CAR XML fields, with proper format conversion when needed.

1. Site, maps to <Site type="gocdb"> (EMI-CAR Par. 2.4.18)
2. SubmitHost, maps to <SubmitHost> (EMI-CAR Par. 2.4.16)
3. LocalJobId, maps to <LocalJobId> (EMI-CAR Par. 2.4.4.2)
4. LocalUserId, maps to <LocalUserId> (EMI-CAR Par. 2.4.5.1)

5. GlobalUserName, maps to <GlobalUserName> (EMI-CAR Par. 2.4.5.2)
6. FQAN, maps to <GroupAttribute type="FQAN"> (EMI-CAR Par. 2.4.5.5)
7. WallDuration, maps to <WallDuration> but requires format change to ISO 8601:2004 (EMI-CAR Par. 2.4.10)
8. CpuDuration, maps to <CpuDuration usageType="all"> but requires format change to ISO 8601:2004 (EMI-CAR Par. 2.4.11 and 2.4.11.1)
9. Processors, maps to <Processors metric="max"> (EMI-CAR Par. 2.5.4)
10. NodeCount, maps to <NodeCount> (EMI-CAR Par. 2.5.3)
11. StartTime, maps to <StartTime> but requires format change to ISO 8601:2004 (EMI-CAR Par. 2.4.12)
12. EndTime, maps to <EndTime> but requires format change to ISO 8601:2004 (EMI-CAR Par. 2.4.12)
13. MemoryReal, maps to <Memory type="Physical" metric="max" storageUnit="KB"> (EMI-CAR Par. 2.5.1 and 2.5.1.1, 2.5.1.2, 2.5.1.3)
14. MemoryVirtual, maps to <Memory type="Shared" metric="max" storageUnit="KB"> (EMI-CAR Par. 2.5.1 and 2.5.1.1, 2.5.1.2, 2.5.1.3)
15. ServiceLeveltype (HEPSPEC | Si2k), maps to <ServiceLevel type="Si2k | HEPSPEC"> (EMI-CAR Par. 2.5.6 and 2.5.6.1)
16. ServiceLevel, maps to <ServiceLevel type="Si2k | HEPSPEC"> (EMI-CAR Par. 2.5.6 and 2.5.6.1)

APEL Usage Record missing elements

Not all the fields in EMI-CAR are mandatory, but in order not to lose functionalities DGAS services needs that the following optional fields in EMI-CAR be published:

1. recordId (EMI-CAR 1.0 Par. 2.4.3.1),
2. createTime (EMI-CAR 1.0 Par. 2.4.3.2),
3. Status (EMI-CAR 1.0 Par. 2.4.8),
4. Infrastructure (EMI-CAR 1.0 Par. 2.4.19),
5. Host - WorkerNodeProcessor (Multiple instances of EMI-CAR 1.0 Par. 2.4.15),
6. LocalGroup (EMI-CAR 1.0 Par. 2.4.5.3),
7. Group (VO) (EMI-CAR 1.0 Par. 2.4.5.4),
8. Queue (EMI-CAR 1.0 Par. 2.4.17);

Then there are other fields in the record, which are not fundamental but should really be taken into consideration:

1. TimeInstant - job created (EMI-CAR 1.0 Par. 2.5.5),
2. TimeInstant - job queued (EMI-CAR 1.0 Par. 2.5.5),
3. TimeInstant - job eligible to run (EMI-CAR 1.0 Par. 2.5.5),
4. ServiceLevel sf2k (EMI-CAR 1.0 Par. 2.5.6),
5. JobIdentity GlobalJobId (EMI-CAR 1.0 Par. 2.4.4.1)
6. CE X509 DN (can be mapped to <Site type="X509"> EMI-CAR 1.0 Par. 2.6.3)