

Metadata Query Protocol

This page documents the consensus of a recent [thread of discussion](#) on the mailing list.

- The location syntax of per-entity metadata SHALL conform to the [Metadata Query Protocol](#) (md-query).
- An md-query server that supports arbitrary "ad hoc" queries is thought to be "a disaster waiting to happen" due to the unbounded resource requirements of such a server.
 - In what follows, a "metadata query" is **not** to be confused with an "arbitrary ad hoc query for metadata." The latter is **not** a goal advocated by this subgroup.
- An md-query server can avoid arbitrary "ad hoc" queries and still deliver significant functionality.
 - In particular, per-entity metadata can be pre-computed and stored in the file system. A simple script on the server maps the `entityID` in the query to a metadata file in the file system.
- It is believed that most (if not all) aggregates of interest can be known in advance and therefore be pre-computed. This precludes having to support arbitrary "ad hoc" queries.
- Let users use entity attributes (or some other mechanism) to pre-register their metadata queries (so that the target aggregate can be pre-computed).
- The location syntax of user-specified aggregates is an (interesting) open question. md-query allows metadata to be requested using arbitrary identifiers, but we may want to avoid potential clashes with entity identifiers.
- An online signing key (secured in a trusted HSM) is a separate issue.
 - From a security PoV, this approach has at least two advantages: 1) the validity window on metadata can be tightened, and 2) the human factor can be eliminated.