

VOQL WG: Maria Nieto-Santisteban and Yuji Shirasaki
Notes, Session 2

Brian Thomas: Achieving a Unified Query Layer in the VO

- Selections against databases/SkyNodes require human interaction to understand contents; must expose local schema
- Solution: Object-Relational Mapping (ORM) layered above queries. Unified query is mapped onto local queries; query results are integrated into unified results.
- Implementation: XQuery on XML documents, with unifying data model in XML.
- Provides concurrent access to XML DB, SQL DB, and SkyNode
- Has tested this approach against Postgresql DB

Pedro Osuna: This corresponds to highest level VOQL.

BT: Chose standard, XQuery.

Masatoshi Ohishi: VOQL WG approach was always three-level. 1) Use SQL-based QL. 2) Develop more sophisticated, XML-based QLs.

BT: Yes, this is understood as the goal. Trying to build on VOQL.

Gerard Lemson: Russia VO has similar ideas, working with UML. But there is not really a standard for Object-Relational Mapping.

BT: Generally agree. General solution requires linking ORM with ontology. VOQL has to work with DM group. Semantics (UCD) group needs to define the objects. DM group's work on catalog/source DM is very important.

Pedro Osuna: Status of Catalog Data Model Group

- Brian's work fits well in the context of the catalog DM development.
- There is only one data model; serialization can be done in various ways.
- VOQL people encouraged to provide inputs.

Maria Nieto-Santisteban: VOQL group is responsible for QL and SkyNodes.

PO: Overlaps between DM, VOQL, and DAL are understood.

BT: Working closely with Catalog DM effort.

Andre Schaaf: Vizier Access to Catalogs

- ~5000 catalogs available through Vizier, by cgi, vizquery, and SOAP interfaces.
- Vizier is not a relational DB. Data are stored in binary, distributed on a cluster.
- ADQL access requires ASU extensions or other dedicated access methods.
- Discussion topics: column names/UCDs. Units. Built-in functions.

Francois Ochsenbein: Need to define how to express units. What is the syntax?

MN-S: Units are defined already in ADQL.

FO: Yes, it is there as option, but syntax is not fully specified.

MN-S: OK, need to introduce syntax for units into the ADQL specification.

- XMatch topics: having catalogs distributed on cluster enables quick cross-match on local tables

- VOTech project includes implementing Vizier catalogs as SkyNodes with ADQL access.

Bob Hanisch: Is the Vizier collection a single SkyNode, or is each catalog at SkyNode?

Sebastien Derriere: Problem for Registry – each access method leads to huge repetition. Suggests changing registry structure so that generic service can be registered and then access methods are listed.

BH: But is it one SkyNode or many?

SD: Should be one.

BT: Should be one.

GL: But this cannot always work.

BT: XML description gives hierarchy.

MN-S: But this is a database.

BT: Yes, exactly.

MN-S: SkyNodes need to move toward higher-level aggregation.

SD: Use reference to IVO identifier for type of service interface. This way each data collection can be registered once.

MN-S: This comes back to a question for the Registry WG.

BH: Yes, this is a Registry question, but also a question for the data provider. Most frequently used catalogs should perhaps be published as separate SkyNodes.

FO: Yes, we want to take a pragmatic approach.

Maria Nieto-Santisteban: SkyNode Specification

- Current draft still needs a lot of work.
- Basic SkyNodes must comply with the core ADQL.
- Full SkyNodes must comply with the core ADQL plus the cross-match function.
- Need a way to find out what ADQL extensions are supported by a SkyNode: use ADQLExtension() to give supported extensions.
- Added method Relation() to describe relationships between tables and columns in a database. Need to know the foreign keys.
- Added method Column() to indicate if a column is indexed.
- Need to decouple cascade and parallel in the execution plan
- XMatch algorithm specification will be removed from SkyNode specification into a separate document.
- Need to add a mechanism to transmit messages in addition to data (e.g., to indicate why no cross-match counterpart was found; no coverage, no object found)
- Update WSDL.

Inaki Ortiz: Should it be required to expose entire DB schema through SkyNode?

MN-S: SkyNode method allows query to determine tables and columns. But some SkyNode publishers may want to expose only some tables/columns through SkyNode without having to create specific views or tables in the DB.

Metadata in SkyNodes

- Column metadata needs to be accessible.

BT: Metadata should be in a wrapper. Astronomers should not be writing ADQL.
MN-S: Yes, but we need to see metadata for a given table behind a SkyNode.
BT and MN-S disagree about where metadata should be stored, though agree that it needs to be available.

Schedule

- Finish WDs by Nov 15. Follow-on intervals for RFCs, etc., lead to ~1 March 2006 target for SkyNode upgrade.

VOQL needs diversity of user interfaces. OpenSkyQuery is the only portal.