JTB phone call support (12 May 2003)

GLUE Domains and Network Service Schema implementation

Augusto Ciuffoletti and Paul Mealor

next: Reminder

index

Reminder

During a GLUE conference call in March we concluded that the GLUE **domains architecture** and the GLUE **schema** for network characteristics observations were mature, and the merge of the two was appropriate.

In order to perform this step, we designed an interface between the two parts: Paul Mealor implemented the API needed to access the R-GMA archive using the addressing implemented by the GlueDomains database, for which I implemented the API and the database structure.

The target language was Perl, using its object oriented features to obtain a clean program structure. GlueDomains is based on MySQL, as well as R-GMA.

next: Work done

up: JTB phone call support (12 May 2003)

index

Work done

The APIs have been designed to give a minimal yet useful toolset: one query to obtain the last observation of a certain characteristic, and one query to store a new observation of a certain characteristic.

These functions alone allow the experimentation of two scenarios:

- an application reads a set of observations of relevant Network Services, and decides which is the more appropriate for a given task;
- a monitoring tool obtains an observation of a Network Service and publishes it in the database.

next: Basic features

previous: Reminder

up: JTB phone call support (12 May 2003)

index

Basic features

The two databases are independent, and their interface is limited: just the data structure that contains the keys is in common between the two.

This improves the modularity of the design. As a consequence, for instance:

- the part that describes the database structure is independent from observation schemas (which are subject to refinements);
- the design of the two parts can be optimized for the specific operations (replicas, RO vs. RW, distribution);
- the technologies used for the two parts are independent: one can switch one from SQL to WML or LDAP without touching the other.

We envision two kinds of scenarios, and have designed the Perl APIs to enable a simplified access in Perl.

next: Scenario 1

previous: Work done

up: JTB phone call support (12 May 2003)

index

Scenario 1

A **Grid Application** wants to obtain network monitoring observations between a pair of Grid Services (Computing or Storage).

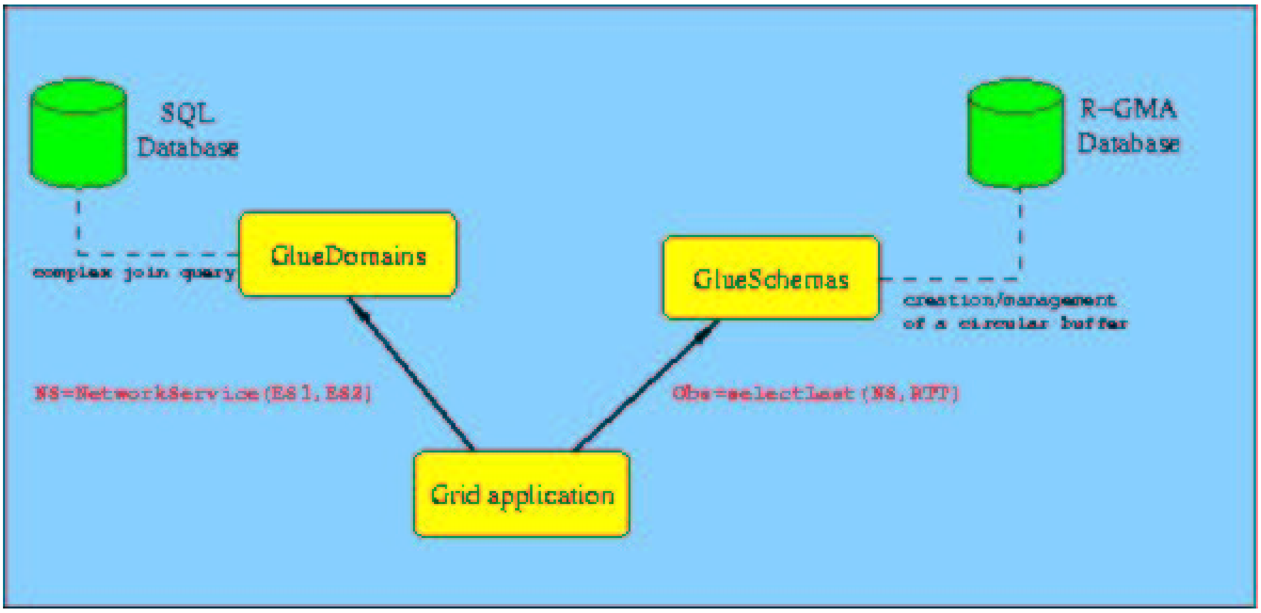

next: Scenario 2

previous: Basic features

up: JTB phone call support (12 May 2003)

index

Scenario 2

A Monitoring Tool wants to record a network monitoring observation.

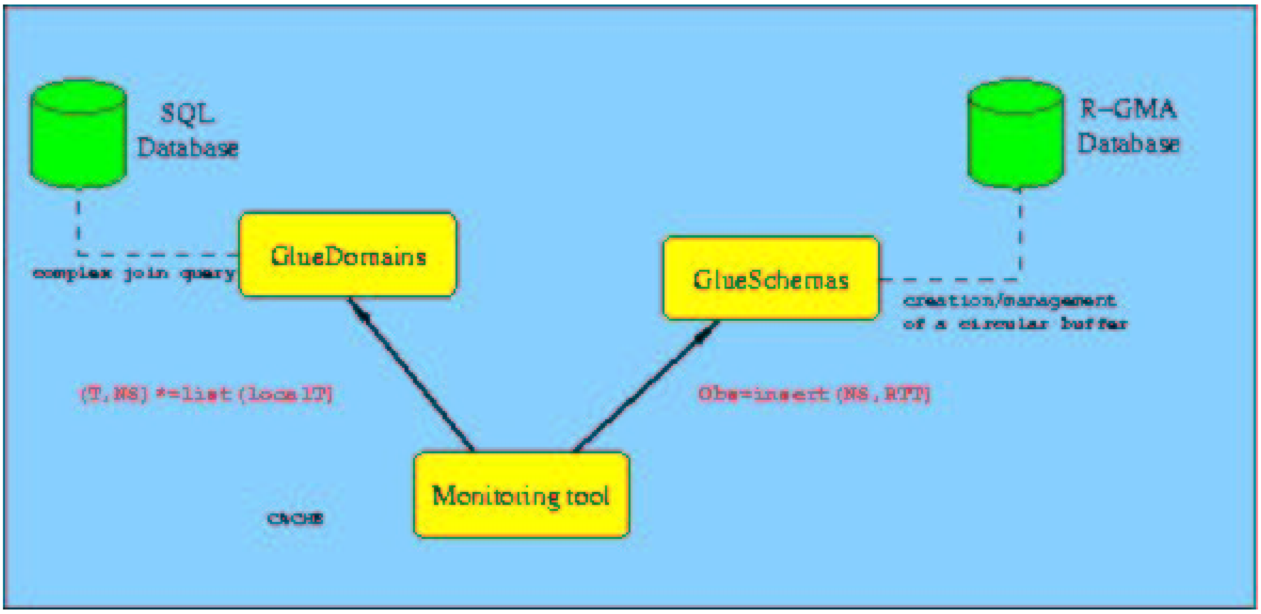

next: GlueDomains Architecture

previous: Scenario 1

up: JTB phone call support (12 May 2003)

index

GlueDomains Architecture

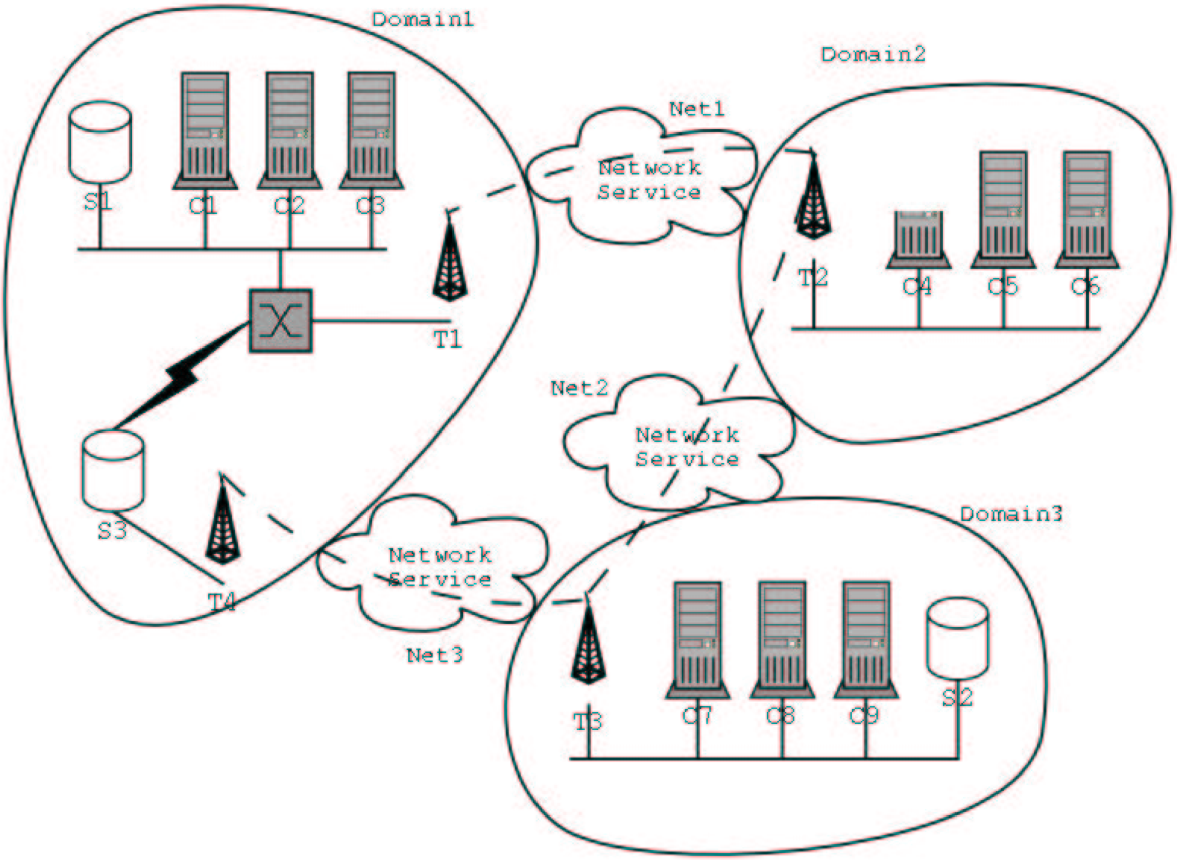

Implementation and concept paper at GlueDomains Home Page

next: Status

previous: Scenario 2

up: JTB phone call support (12 May 2003)

index

Status

The two APIs have been implemented separately (Paul Mealor implemented the part that I call GlueSchemas, I implemented GlueDomains).

We are setting up a testbed using R-GMA as a support for GlueSchemas, and a replicated SQL database for GlueDomains.

next: Index of the slides

previous: GlueDomains Architecture

up: JTB phone call support (12 May 2003)

index