Workshop on Distributed Systems and Mobile Code

André van der Hoek, Dennis Heimbigner, and Alexander L. Wolf

Department of Computer Science, CB 430
University of Colorado
Boulder, Colorado 80309 USA
{andre,dennis,alw}@cs.colorado.edu

In a situation where a number of geographically dispersed development teams are working on some set of inter-dependent software tools, software release management plays a crucial role. On one hand, one would like to provide the development teams with a single, distributed, and structured repository where they can store all software releases, their descriptive information, and—most importantly—their inter-dependencies. On the other hand, one would like to have as large an audience as possible, and publicize the released software as complete packages through a ubiquitous medium such as the internet.

This scenario was the driving force for the design and implementation of SRM [2], a Software Release Management tool. SRM uses NUCM [3, 1], a CORBA based distributed repository for Configuration Management, to implement its database. In addition to a Tcl interface through which software can be released/retrieved from the database, SRM has a WWW interface through which released software, including its dependencies, can be retrieved. Thus, the architecture of SRM has two layers of distribution: the CORBA based distributed database, and the WWW for retrieval of information from the database. A collection of CGI scripts provides a WWW browser access to the repository. Of particular interest is the CGI script that retrieves a software release and its dependents. It uses CORBA to contact the NUCM database, retrieves the release and its dependents from the various geographically distributed sites, collects all software in a single TAR file, and ships the TAR file back to the user of the WWW interface of SRM. As such, the CGI script bridges the gap between the HTTP protocol used by the WWW, and the CORBA protocol used by SRM.

In building SRM, we have learned a number of valuable lessons about WWW, CORBA, and their interaction:

1. *Static web pages:* it can be advantageous to maintain a set of up to date WWW pages parallel to the structured data in the repository, instead of generating a WWW page from the data in the repository for each incoming request. However, such a solution is not always feasible.
2. *Viewing structured information:* as soon as a structured view of the information in the database contains multiple parts (e.g., a dependency graph in a WWW page with descriptive information), CGI scripts become significantly more complex to implement due to the need for a caching mechanism.
3. *Protected CGI scripts:* due to the fact that CGI scripts are executed in a limited and protected environment at our site (and expectedly so at many other sites), work-arounds in the CGI scripts and adjustments in the repository implementation are a necessity.
4. *Firewalls:* existing implementations of the CORBA standard are vulnerable to firewalls, and therefore the metaphor of (almost) transparent distribution in the implementation of a wide area structured repository does not hold up very well.

It is clear that a combination of a structured CORBA repository and a WWW based access mechanism is very appealing, but, as the above points illustrate, difficult issues remain that need to be resolved before wide-spread deployment of such an architecture can take place.

This material is based upon work sponsored by the Air Force Material Command, Rome Laboratory, and the Advanced Research Projects Agency under Contract Number F30602-94-C-0253. The content of the information does not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred.

References

1. Software Engineering Research Laboratory University of Colorado. Distributed Configuration Management. Available on the world wide web at *http://www.cs.colorado.edu/home/serl/projects/NUCM.html*.
2. Software Engineering Research Laboratory University of Colorado. Software Release Management. Available on the world wide web at *http://www.cs.colorado.edu/home/serl/projects/SRM.html*.
3. A. van der Hoek, D. Heimbigner, and A.L. Wolf. A Generic, Peer-to-Peer Repository for Distributed Configuration Management. In *Proceedings of the 18th International Conference on Software Engineering*. Association for Computer Machinery, March 1996. To appear.