

XSPEED - High Performance XML Processing

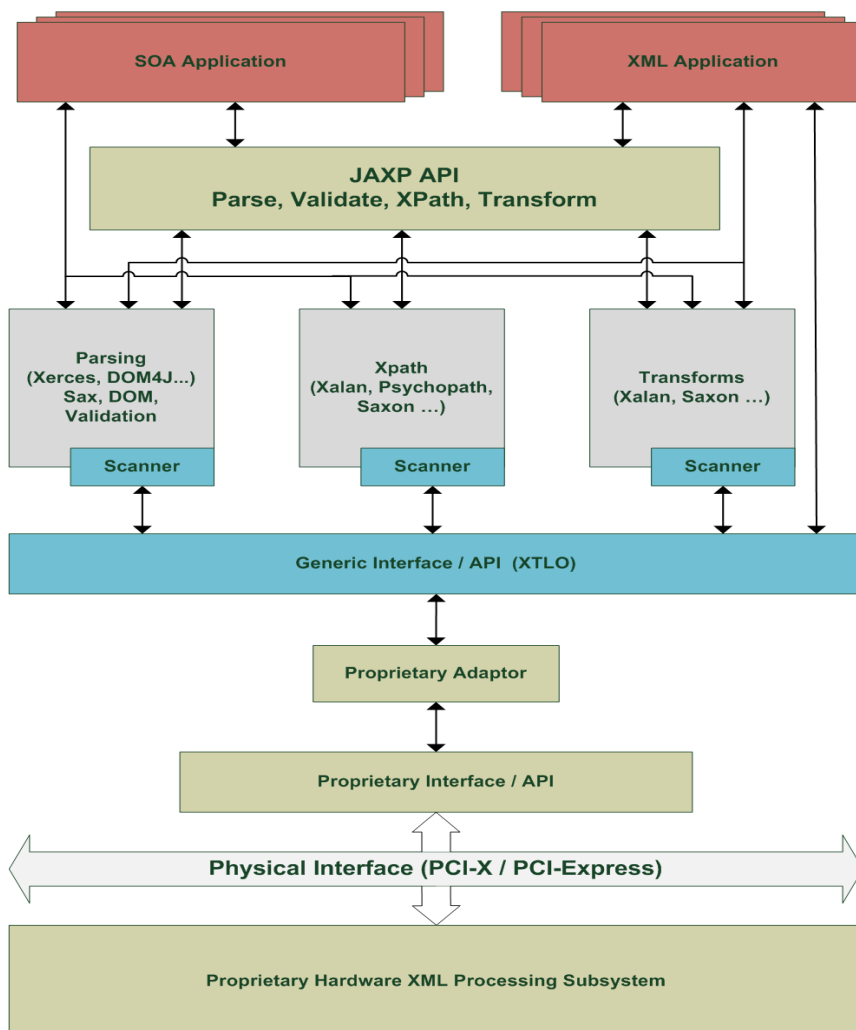
XSPEED aims to contribute to the adoption of hardware acceleration for XML processing in Service Oriented Computing (SOC), by defining, implementing and evaluating a generic interface to which hardware vendors can readily adapt their products, and that will allow application designers to take advantage of XML processing hardware while maintaining vendor independence.

Main focus

As the data format underlying Service Oriented Architecture (SOA), XML brings simplicity, extensibility, interoperability and

openness. However XML and its associated security protocols carry large computing overheads that can lead to unacceptable reductions in message throughput and increases in latency as SOA installations scale.

Hardware acceleration can offer a solution, but the absence of a standard interface, through which applications use XML acceleration, gives rise to three issues that block its more widespread deployment. Firstly, in a typical corporate SOA, only those systems subject to unsustainable xml traffic need acceleration. Secondly, software needs to be modified to accom-



XSPEED

Project ID: CP5-057

Start Date: August 2009

Closure date: September 2011

Partners:

Cybersoft, Turkey

Dublin City University, Ireland

MKK, Turkey

Co-ordinator:

Seamus Farrell

Dublin City University

E-mail: sfarrell@computing.dcu.ie

Project Website

www.celtic-initiative.org/projects/XSPEED

modate a proprietary interface, and thirdly, modifying software in this manner generates a single-supplier dependency.

A standard generic interface can address these problems. The presence of an accelerator can be transparent to application level software, allowing the acceleration to be applied only in those installations where it is needed. Such a standard should encourage and eventually oblige hardware vendors to conform, thus removing vendor dependencies.

Approach

The project begins with research into the current state of the art with respect to XML acceleration subsystems. The intention is to explore the common characteristics of such hardware subsystems with a view to designing a suitable generic interface to which any hardware vendor may readily adapt a proprietary solution. This study will focus on plug-in hardware rather than dedicated network appliance offload solutions which are outside the scope of the project. The most popular XML processing libraries, specifically those for which source code is available, will be studied with a view to a design for an efficient interface to underlying acceleration. Selected SOA and other XML intensive application software will be analysed to establish which kind of xml processing represents the greatest threat to effective performance, and this will contribute to the requirements of the proposed interface. The actual XML libraries that

are to be part of the implementation will be agreed.

Based on this study, a Generic XML processing Interface will be designed. Once the definition has been agreed, development work will proceed to an implementation. Dajeil Limited¹ has provided access to its XML acceleration technology, so the first implementation will be an adaptor interposed between Dajeil's proprietary interface and the Generic Interface. At the same time, a scanner providing the link from the Generic Interface to the selected parsing library will be developed. Depending on the outcome of the earlier analysis, scanners / adaptors may also be developed for XPath and XSLT libraries.

Once initial prototype adaptors are available, work may be done to adapt the selected XML applications to use the modified libraries, or to directly use the Generic Interface as appropriate.

The effectiveness of the Generic Interface will be measured through performance tests conducted in scenarios with and without acceleration. Further tests will establish the efficiency with which a vendor specific proprietary interface may be adapted to the Generic Interface.

Main results

- ◆ Investigation into and implementation of an efficient mechanism to deliver the performance uplift of hardware acceleration at a level of abstraction useful to

SOA and other XML intensive applications.

- ◆ Definition of a Generic XML Hardware Subsystem Interface
- ◆ An application of XML hardware acceleration to one or more real-world XML intensive software, connecting the software to the hardware through the generic interface.

Impact

Perhaps the biggest factor limiting the more widespread adoption of XML acceleration has been the absence of a standard interface between application software and the low-level proprietary interfaces presented by hardware acceleration products. A standard interface supported by hardware vendors would give application designers freedom to employ acceleration without being locked in to a particular hardware vendor. Furthermore, for a given application, not all installations will encounter sufficient XML traffic to require the use of hardware acceleration.

A suitably designed Standard Generic Interface would allow hardware acceleration to be used only where it is necessary, and, with the exception of its effect on performance, the presence or not of the acceleration would be transparent to the application. There would no longer be a dependence on a specific hardware vendor. All in all, the presence of a standard would reduce the reluctance of software vendors to use acceleration and so benefit hardware vendors.

As an additional benefit, effective and appropriate use of hardware acceleration should reduce the number of general purpose server systems needed to satisfy a given level of XML traffic. Reducing the number of servers also reduces space and power consumption, both environmentally worthwhile savings.

About Celtic

Celtic is a European research and development programme, designed to strengthen Europe's competitiveness in telecommunications through short and medium term collaborative R&D projects. Celtic is currently the only European R&D programme fully dedicated to end-to-end telecommunication solutions.

Timeframe: 8 years, from 2004 to 2011

Clusterbudget: in the range of 1 billion euro, shared between governments and private participants

Participants: small, medium and large companies from telecommunications industry, universities, research institutes, and local authorities from all 35 Eureka countries.

Celtic Office

c/o Eurescom, Wieblingen Weg 19/4,

69123 Heidelberg, Germany

Phone: +49 6221 989 405, e-mail: office@celtic-initiative.org

www.celtic-initiative.org



¹Dajeil Limited (<http://www.dajeil.com>) is an Irish company with a number of products that address XML optimisation. It is providing acceleration hardware and expertise to the project consortium.