

Project Profile

Making the most of massive data flows

Real-time publish/subscribe middleware infrastructure

The PubSub4RT project is developing a middleware infrastructure for processing massive data flows making use of publish/subscribe technology. The objective is to enable real-time processing of large data streams over wide area networks (WAN). The infrastructure will process a wide variety of data sources such as company email, the output of large sensor networks, general network traffic in an organisation, credit-card payments, financial data and phone calls. Potential markets include banking, security, Internet service-provider (ISP) services and sensor networks.

Many applications now need to process massive data flows. Examples include:

- Anti-spam and anti-virus filters for company email systems;
- Handling the output of large sensor networks;
- Protecting an organisation's network traffic against distributed denial of service (DDoS) attacks;
- Anti-fraud services for credit-card payments;
- Financial services such as dealing with stock quotes;
- Real-time processing of call description records to fight fraud in mobile phone networks; and
- Real-time processing of business information when carrying out business intelligence.

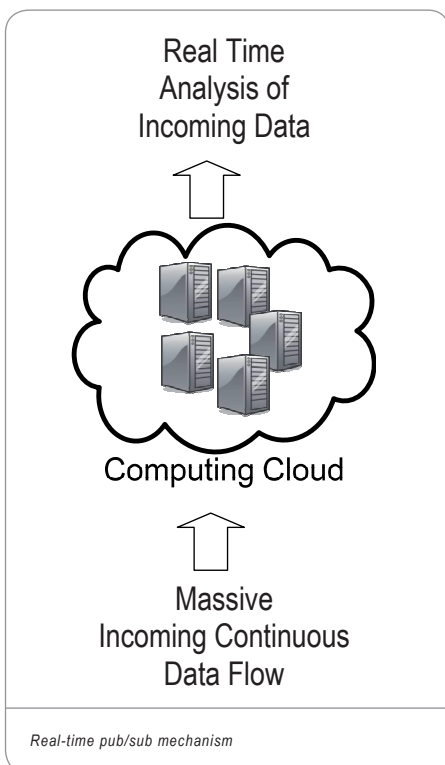
The goal of PubSub4RT is to develop a middleware infrastructure able to process these massive flows of data and to create a new breed of applications providing services based on processing of these flows.

TWO-TIER STRUCTURE

PubSub4RT proposes a two-tier middleware infrastructure.

A core tier will enable the real-time processing of the massive data flows, while an upper tier will make it possible to supply on-line publishing/subscription services to client applications – for example subscribe to events announcing a potentially fraudulent credit-card payment, a potential DDoS attack, a potential spam message or an interesting stock-market position.

The main novelty of this ITEA project is that the proposed platform aims at providing support for both services on a single domain and in multiple domains separated by a WAN – that is a federation of clusters.



PubSub4RT (ITEA 2 ~ 07004)

■ Partners

Atos Origin
Universidad Politécnica de Madrid
FORTH
Virtual Trip

■ Countries involved

Greece
Spain

■ Project start

March 2008

■ Project end

December 2010

■ Contact

Project Leader :
Pedro Soria-Rodriguez
Atos Origin

Email :
pedro.soria@atosorigin.com

Project Profile

TARGETING SEVERAL DOMAINS

Market relevance is particularly high, including:

Banking: a typical problem is fraud detection in credit-card payments, where real-time processing of payment authorisation is currently impossible. With the PubSub4RT platform it will be possible to characterise payment flow in real time and detect deviations.

ISP services: a number of enhanced services could be provided in the ISP sector – particularly improved spam filters and anti-virus software. For example, the proposed platform will be able to process all an ISP's email traffic, enabling correlation of email contents and improving the detection of spam.

e-service protection: safeguarding organisation that make business through e-services – such as flight ticket sellers, book sellers, auction sites and Yellow pages services – from DDoS attacks. Such attacks can disrupt a company system for several hours, severely damaging the company's image and creating huge economic losses.

Sensor network services: many applications process the output of large sensor networks – such as maritime traffic surveillance to protect maritime borders from illegal immigration, register the route taken by commercial traffic and monitor commercial fishing. The output of such networks is huge. The PubSub4RT platform would enable a much higher levels of border control, reduce maritime pollution and control of overfishing.

Economic and social intelligence: the PubSub4RT platform would make it much easier to handle vast amounts of data – enabling a company for example to react quickly to a crisis by real-time opinion tracking and society in general to predict open crises before they happen.

BUILDING ON EXISTING RESEARCH

PubSub4RT is building on existing research results in three areas to contribute to an area of significant commercial importance in emerging applications domains:

1. Parallelisation of individual data-

streaming operators across a cluster of sites, based on parallel and distributed query processing – here the key challenge is the real time nature of data streaming that is quite different from traditional relational databases;

2. Achieving high performance and scalability using the latest communications subsystems technologies based on low-latency, high-bandwidth interconnects. PubSub4RT will contribute application-specific extensions to the communications subsystem used on top of such interconnects and examine novel communications layer functionality specific to data streaming, including support for synchronisation, processing offloading and fault tolerance.
3. Building on existing research results in scalable middleware for correlating events, such as in database technology. PubSub4RT will use these results to build a middleware layer specific to data-stream processing. This middleware will deal with performance and scalability issues in the presence of data dependencies.

USABLE WITH ANY SYSTEM

PubSub4RT will take full benefit of experience in developing the Global Remote-Write Architecture (GRWA). This software architecture is composed of a set of independent modules which are all associated with an active element in the communications path. This approach will ensure that all developments performed during the ITEA project will be usable on any network system – and also any processor family and/or operating system – with very little further development.

Expected results include: development of an initial prototype of the PubSub4RT middleware including requirements, software and documentation; and development of several demonstrators as proof of concept of some application domains: requirement analysis, software, and documentation. PubSub4RT will prototype and evaluate the proposed solutions using real applications and practical real-life setups.

ITEA 2 Office

High Tech Campus 69 - 3

5656 AG Eindhoven

The Netherlands

Tel : +31 88 003 6136

Fax : +31 88 003 6130

Email : itea2@itea2.org

Web : www.itea2.org

■ ITEA 2 – Information Technology for European Advancement – is Europe's premier co-operative R&D programme driving pre-competitive research on embedded and distributed software-intensive systems and services. As a EUREKA strategic Cluster, we support co-ordinated national funding submissions and provide the link between those who provide finance, technology and software engineering. Our aim is to mobilise a total of 20,000 person-years over the full eight-year period of our programme from 2006 to 2013.

■ ITEA 2-labelled projects are industry-driven initiatives building vital middleware and preparing standards to lay the foundations for the next generation of products, systems, appliances and services. Our programme results in real product innovation that boosts European competitiveness in a wide range of industries. Specifically, we play a key role in crucial application domains where software dominates, such as aerospace, automotive, consumer electronics, healthcare/medical systems and telecommunications.

■ ITEA 2 projects involve complementary R&D from at least two companies in two countries. We issue annual Calls for Projects, evaluate projects and help bring research partners together. Our projects are open to partners from large industrial companies and small and medium-sized enterprises (SMEs) as well as public research institutes and universities.



Σ! 3674