


Project Results

Highly scalable platform boosts service dependability

Secure real-time processing of massive data flows opens up novel financial and communications applications



Many applications in communications, banking, security, Internet services and sensor networks involve treating massive amounts of data over wide geographic areas. The ITEA 2 PubSub4RT project has developed an infrastructure to handle such flows in real time using publish/subscribe technology. This opens up a range of novel services, particularly in the financial sector, business intelligence and telecommunications. The platform is highly scalable and has been successfully demonstrated in wide-area networks for real-time credit-card fraud detection.

Until now, data-streaming processing infrastructures have had only modest scalability, limiting their application. Such infrastructures are characterised by having to process the whole information flow in each node, which become individual bottlenecks. PubSub4RT developed an architecture that enables information flows to be processed in parallel between a subset of nodes, avoiding concentrating the information flow in a single node.

The parallel data-streaming system enables aggregation of hundreds of nodes to process a single query and handle a million events a second. Such distributed processing results in a highly scalable infrastructure able to tackle a new class of applications. PubSub4RT middleware can also help solve many of the security problems facing financial entities and Internet service providers – such as denial-of-service attacks, spamming and phishing.

TWO-TIER MIDDLEWARE

A two-tier middleware infrastructure involves a central core with a data-streaming engine that can handle massive data flows in real time, while an additional upper layer provides publish/subscribe on-line services to the client applications.

A major novelty is the ability to support services in both single and multiple domains separated by a wide area network. Thus the PubSub4RT approach is particular relevant to telecommunications, banking and financial services, network security, Internet service provider (ISP) services and sensor networks such as those used to monitor road traffic, the quality of running water in large cities or environmental conditions.

The platform includes two self-management features: self-optimisation – dynamic load balancing makes it possible to maximise the throughput for a given configuration; and self-provisioning – new nodes are provisioned and decommissioned to

PubSub4RT (ITEA 2 ~ 07004)

Partners

Atos Origin
ICS-FORTH
Universidad Politecnica de Madrid

Countries involved

Greece
Spain

Project start

June 2008

Project end

December 2010

Contact

Project leader :

Pedro Soria-Rodriguez, Atos Origin

Email :

pedro.soria@atosorigin.com

Project website :

<http://pubsub4rt.atosorigin.es/en>

Project Results

minimise the resources required to process the incoming load.

Data-streaming middleware provides a SQL-like language to handle queries. This offers 'stateless' operators, which provide output based on each individual input string, and 'stateful' operators, with output based on a sliding time window. The network overhead is reduced by cutting the number of copies made in the network path, introducing flow control and batching data. The result is a highly scalable data-streaming platform which enables events to be processed at very high rates.

DETECTING CREDIT-CARD FRAUD

The PubSub4RT approach was demonstrated successfully in a credit-card fraud detection application. Current systems are only able to detect fraud after the occurrence – and this can be hours or even days after such a fraud has been committed with corresponding financial loss. PubSub4RT enables real-time detection of fraud before payment authorisation, avoiding losses.

The resulting platform opens up a wide range of potential applications:

- Banking: anti-fraud, anti-money laundering and market surveillance;
- Financial services: automating and optimising buying/selling share options;
- Telecommunications: anti-fraud operations in mobile phone networks as

PubSub4RT can process all calls in real time to detect fraud patterns;

- ISP services: particularly anti-spam and anti-virus filters;
- Network security: detection and mitigation of distributed denial-of-service attacks as well as general network monitoring; and
- Business intelligence: offering real-time approaches that permit fast detection of market trends.

Other possible applications include: air, road and sea traffic control; border surveillance; environmental monitoring; e-health systems; and cloud computing.

FAST EXPLOITATION OF RESULTS

Partner Atos Origin is already deploying PubSub4RT in its business risk and information security activities. A particular application is to provide security in the information and communication technology services management of future Olympic Games as part of its event-management activity.

Universidad Politecnica de Madrid has filed a US patent on the parallel data streaming engine and is exploring telecommunications applications. It already has a contract with Ericsson to explore the data-streaming technology for deep packet inspection and is providing training in the data-streaming area with a course at Ericsson.

Major project outcomes

DISSEMINATION

- 4 presentations of PubSub4RT at industry and research fairs
- 1 course on data streaming technologies
- 6 paper submissions

EXPLOITATION

- Atos Origin is testing PubSub4RT technologies within a project to be tested during the Olympic Games (Atos Origin is Worldwide IT Partner for the Olympic Games and Top sponsor)

PATENTS

- 1 patent application filed in US by Universidad Politécnica de Madrid for the PubSub4RT parallel data streaming engine

ITEA 2 Office

High Tech Campus 69 - 3
5656 AG Eindhoven
The Netherlands

Tel : +31 88 003 6136
Fax : +31 88 003 6130
Email : info@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.

