BRIDGE

Newsletter - February 2009

BUILDING RADIO FREQUENCY
IDENTIFICATION SOLUTIONS
FOR THE GLOBAL ENVIRONMENT

Welcome to the BRIDGE Project Newsletter !

This newsletter is published every two months to keep you updated on the happenings within the BRIDGE project. Each edition contains topical information arising from the various Work Packages within BRIDGE as well as other BRIDGE related information.

For this edition, we will report on the following topics:

- Serial-Level Supply Chain control - Progress on the track and trace analytics framework
- Security work package - Securing Collaborative Supply Chain Networks
- Calendar of events and latest news

Any feedback or questions contact emilie.danel@gs1.org

BRIDGE

SIXTH FRAMEWORK PROGRAMME

Information Society Technologies

The main aim of this work package is to gather information from the EPC network, analyse the information and utilise it to enhance existing supply chain business processes and/or decision-making processes.

Recently, the BRIDGE work package 3 has completed its second deliverable, consisting of a Track and Trace Analytics Framework implemented as a working software prototype. This has been the result of over 12 months of work in a collaborative effort involving researchers at BT, SAP and the University of Cambridge.

The Track & Trace Analytics Framework is intended to greatly ease the development of a new generation of business applications based upon the analysis of EPC-based tracking data. The framework is especially useful for applications requiring the analysis of event sequences (trace data).

To a considerable degree, the work develops the concept first successfully prototyped by ETH in their 'Supply Chain Visualizer' (part of WP4.6) more than a year ago. The key difference is that the Track & Trace Analytics Framework is a much more modular, detailed, flexible and extensible system. It has been developed with the specific intention of being highly re-usable, and open for development by many parties.

The key modules in the framework are shown below in Figure 1 and described in the following pages.

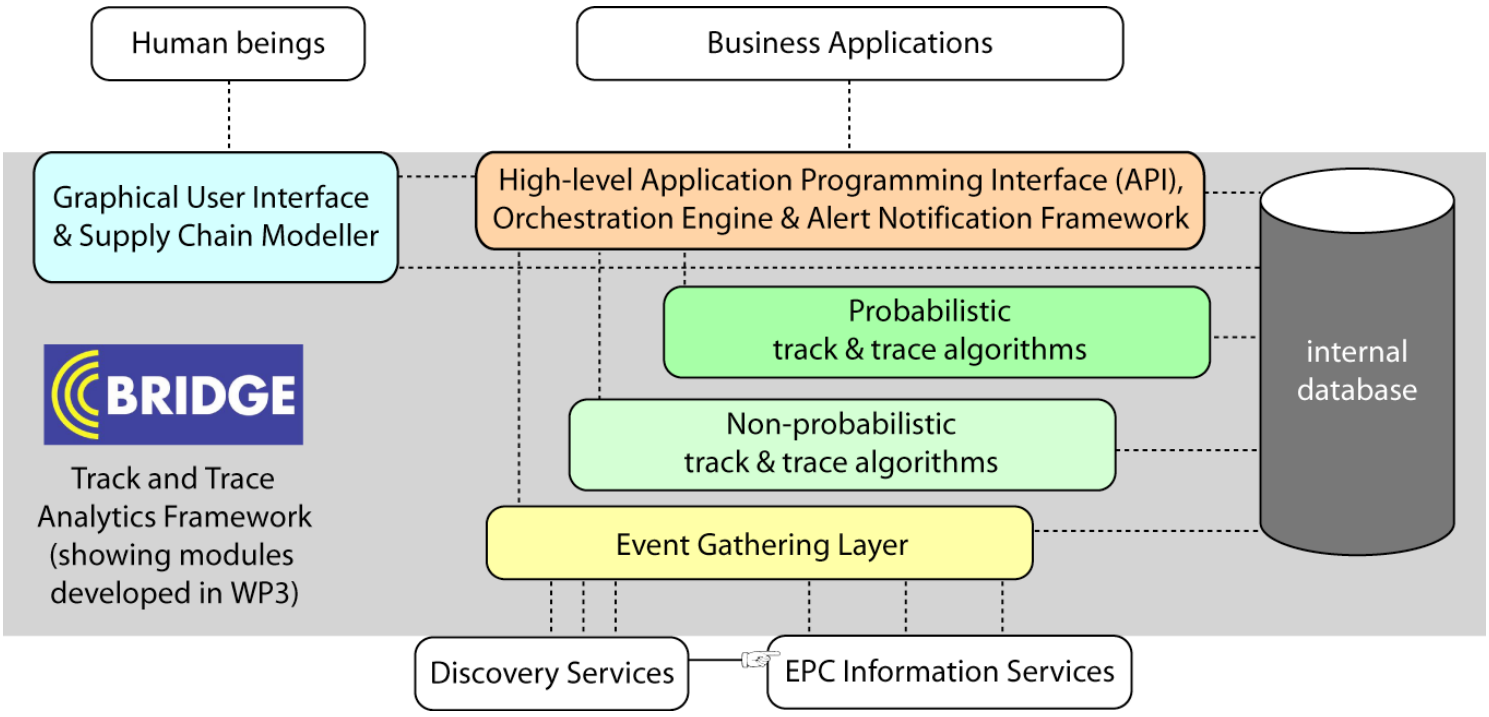

Figure 1: Conceptual diagram of Track & Trace Analytics Framework

Event Gathering Layer - convenient gathering of EPCIS events

The purpose of this is to collect together all the relevant event data for a particular analysis or business application requirement from multiple distributed databases that each provide an EPCIS query interface. For example, if a business application wished to determine the precise route a particular item had travelled through the supply chain, it would need to collect together not only all the read events associated with the item itself, but also those events associated with any containers (cases, pallets, or vehicles) in which it was placed. The Event Gathering Layer includes mechanisms for automatically following changes of aggregation, as well as inference techniques for reasoning about missing aggregation events. It also includes optimisations to avoid the need for unnecessary repeat requests to Discovery Services and EPCIS instances for event information that it has previously gathered. The development of the Event Gathering Layer was led by BT and researchers at both BT and SAP have developed implementations, specializing in particular aspects of the design.

Supply chain modeller - including automatic generation of the models

The supply chain model allows a particular supply chain route to be represented in detail, both from a hierarchical perspective, from the identity of a company down to the granularity of an individual shelf or dock door, as well as a supply-chain perspective, following the changes of location during the flow of the objects. It is extremely general and can immediately be applied to almost any conceivable real-world supply chain.

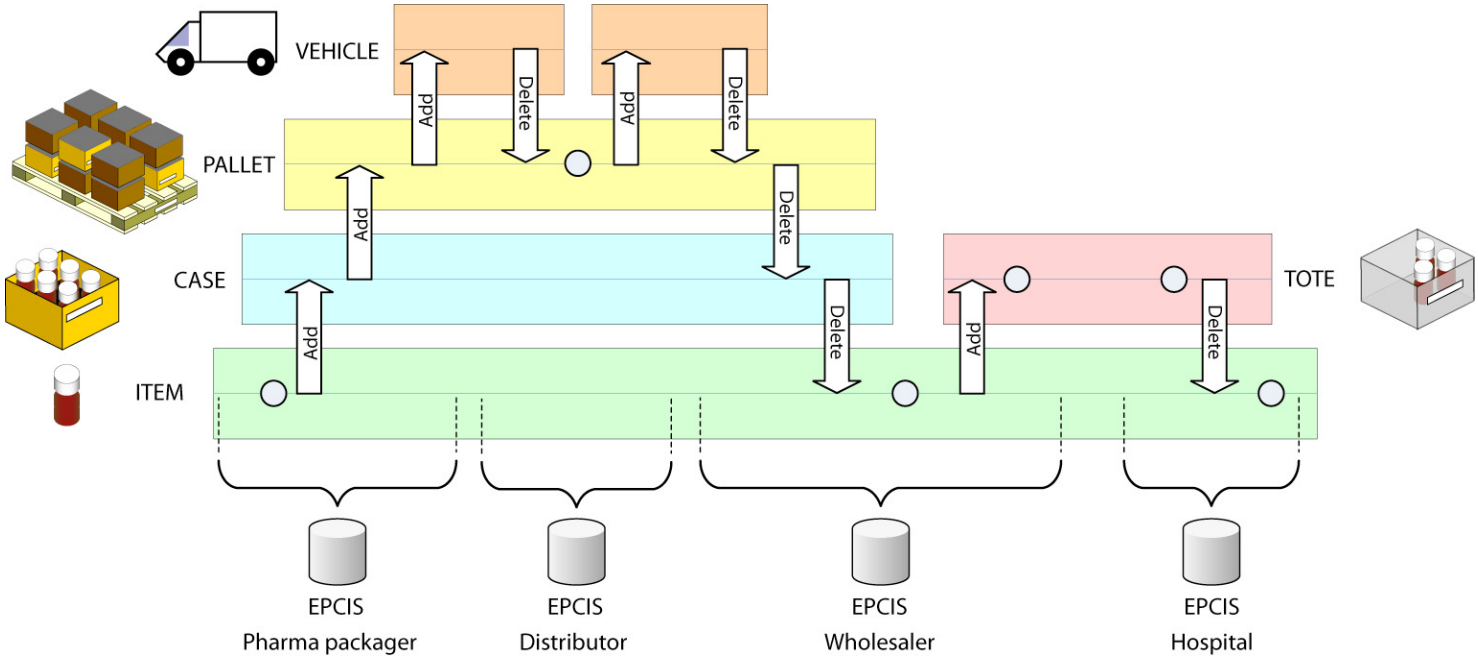

Figure 2: The Event Gathering Layer can be used to automatically follow changes of aggregation, such as these changes observed in the WP6 pharmaceutical traceability pilot

A very powerful feature of this component is the flexible manner in which the model can be formed, or populated. If the supply chain routes are already well known by an organisation, then the model can be manually specified and configured via a graphical user interface. Analysis algorithms can then make comparisons (and detect deviations) between the actual routes travelled by the goods (as represented by the measured event sequences) and the expected route travelled (as represented by the supply chain model itself).

Alternatively, if the supply chain routes are known to be complex, or not well understood, then the supply chain modeller can automatically deduce a supply chain from the measured event sequences themselves. We would expect this to be a very useful feature, and to allow business analysts a far more accurate and deeper insight into their supply chains than has ever been possible before.

The Supply Chain Modeller and its graphical interface were developed by researchers at SAP.

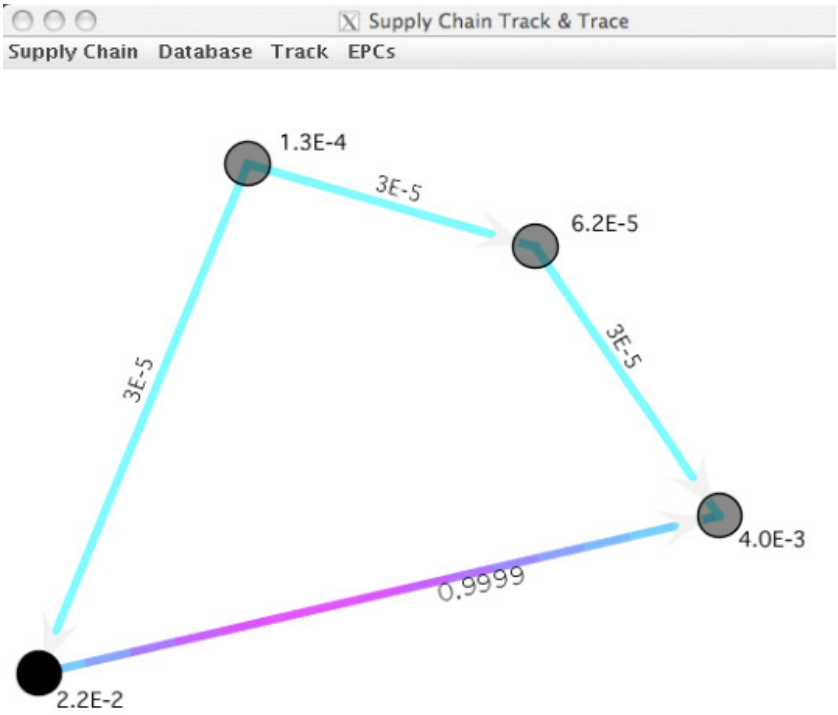

Figure 3
The supply chain modeller allows organisations and locations to be modelled as a network of nodes. Here we see results of the probabilistic tracking algorithm superimposed on the model. The numbers next to the nodes and arcs correspond to the probability that the object is currently on each node or arc. The colour gradient indicates the current probability distribution along the most probable arc.

Analysis Algorithms - converting event data into information

Having used the Event Gathering Layer to collect all the necessary event sequences, various general purpose analytical algorithms can then be used to analyse the data. This is another, or possibly *the* key element of re-use of the Framework. Rather than expecting every single distinct business application to write its own analysis algorithms, we recognised that many different business applications could be built from just a relatively small number of individual algorithms. The analysis algorithms are primarily of two sorts: non-probabilistic, and probabilistic algorithms.

Non-probabilistic algorithms - analysis of event data

The non-probabilistic algorithms deal with information concretely and analytically derived from the raw data event sequence. To date, we have implemented non-probabilistic track and trace query methods that provide convenient wrapper methods for retrieving results from the eventlog database that was populated by the Event Gathering Layer.
However, further non-probabilistic algorithms could be developed, including algorithms that automatically determine how many containers (and levels of container) an item has been placed in throughout its journey, or how many (of a particular item) are typically transported on a pallet - or which objects were stored or transported near a particular object of interest. Non-probabilistic algorithms also include those that are used to check the integrity of certain event sequences (e.g. to ensure that a particular processing sequence has been followed, or that all the items aggregated into a container are correctly reported to have been disaggregated from it). Such algorithms are being developed in WP4 (Security) - task 4.6 (Data integrity) - and could be incorporated within the Track & Trace Analytics Framework.

Probabilistic algorithms - learning and predicting the flow patterns

The second form of algorithms is probabilistic in nature and uses analysis of historical flow patterns in order to predict the most probable current and future locations of objects of a particular type. They can be used to answer questions like:

- what time *do we expect* a particular batch of goods to arrive at location X?
- what is the *probability* that my supplies will reach me on time?
- which is the *most likely* route a particular item has travelled?

These algorithms tend to be more sophisticated in nature and are described in our first deliverable D3.1, with some further refinements described in D3.2.2.
These techniques should demonstrate to business users the significant business benefits of adopting serialised identification combined with automatic identification technologies such as RFID.

The probabilistic and non-probabilistic track and trace algorithms were developed by researchers at the University of Cambridge. It is anticipated that the Track & Trace Analytics Framework could also accommodate the data integrity algorithms developed in task 4.6 and anti-counterfeit algorithms developed in WP5.

Graphical Visualisation Tools

Visualisation tools allow the supply chain model and its features to be easily represented to human analysts. To some degree, this component really represents an application-level component, rather than the genuine middleware elements described previously. However, a general-purpose visualisation tool is of such utility that it makes sense to provide such a tool within the Analytics Framework itself. It is again a highly re-usable component that can aid the development of many distinct business applications.

The researchers at SAP developed a visualisation tool that is able to represent the supply chain either in a 'hierarchical view' that emphasises the internal, hierarchical, tree-like structures of the individual companies involved in a particular supply chain or in a 'supply chain view' that emphasises the routes and the movement of products between locations – and can optionally omit the details of processes within individual companies.
Researchers at the University of Cambridge have developed a further extension, to enable the results of the probabilistic track and trace algorithms to be superimposed on the visualisation tool.

Integration and Testing

The Framework components have so far been integrated into two slightly different configurations, at Cambridge and SAP. Both of these form successful working prototypes of the overall Track & Trace Analytics Framework.
Each has been configured to test and to integrate with somewhat different systems:
The focus at Cambridge has been on the testing of the 'probabilistic algorithms' with data acquired from both the WP6 real-world pharmaceutical pilot, and from their own RFID-enabled automated manufacturing testbed laboratory.
The focus at SAP has been on integration with their implementation of a Discovery Service in WP2, with the testing of Integrity-checking algorithms, and possibly with their anti-counterfeiting work in WP5.
Researchers at Cambridge, SAP and BT have all been involved in analysis of data received from the WP6 pharmaceutical traceability pilot. This is a particularly interesting dataset, because of the large number of different organisations, the existence of multiple routes that converge, as well as several changes of aggregation during the distribution process. From analysis, a number of interesting features have been detected - and WP3 intends to demonstrate the Track & Trace Analytics Framework at the next BRIDGE review meeting (1-2 April 2009), using the WP6 data to illustrate the purpose of each component, as well as highlighting interesting features in the WP6 data.

Further development

It is anticipated that task 3.5 will result in a further enhancement of the Track & Trace Analytics Framework. Detailed architecture designs are being developed for an Alert & Notification Component, which will allow business users to be notified about problems occurring within their supply chains. Several distribution channels will be supported, including e-mail and SMS text messages, as well as live updating of visualisation tools and logging to files. Additionally, a number of user-configurable convenience methods are being developed, which allow business users to indicate the criteria for such alerts and the thresholds and parameters of interest. For example, it is expected that a method will be developed for monitoring delays of objects within the supply chain. This can also form the basis of further methods, such as detection of shrinkage. Task 3.5 is being led by SAP, with contributions from researchers at BT and the University of Cambridge.

For more information on this topic, contact:
Mark Harrison
(Auto-ID Labs Cambridge, WP3 leader)
Tel: +44 (0)1223 338178
E-mail: mark.harrison@cantab.net

The objective of this work package is to secure RFID data services by increasing the control over data sharing. RFID services developed in harmony with security features will enable effective and safe deployment of new applications within enterprises and across global supply chains.

WP4 is a technical WP on RFID security. This means balancing the needs of applications for visibility of RFID and related data, against requirements for the confidentiality of information. Since critical business decisions are made as a result of RFID data, the integrity of the data is also of utmost importance.
Recently, WP4 has completed the design of a security framework for the Discovery Service developed in BRIDGE WP2 and for other RFID networked services such as the EPCIS. This has been the result of over 2 years of collaboration among teams in SAP, AT4 wireless, ETH Zurich and BT. The work has developed the technology required to introduce fine-grained access control within the Discovery Service, looked at realistic policies and management scenarios, and implemented the concepts in a Discovery Service working prototype.

Need for Collaborative Information Sharing

Organising collaborative supply chains is becoming increasingly complex. Every problem, from delays and shortage of stock to production outages, affecting one business function in one company can affect the whole value chain. Traditional supply chain analysis has focused on the single operations of companies. RFID has the opportunity to change supply chain models that have been untouched for years where supply chain networks run as closed networks of trusted peers.

End-to-end visibility over open supply chains has the opportunity to provide a range of benefits that cannot be realised by a single partner. Such benefits may include:

- *More information, choice and safety for the consumer*
- *Greater trust in retail & manufacturer brands*
- *Reduction of danger/costs of illicit goods (e.g. pedigree, anti-counterfeiting)*
- *Condition monitoring (e.g. cold chain, blood, vaccines)*
- *End-to-end supply chain efficiencies*
- *Correct apportionment of RTI costs*
- *Early detection of irregularities & exception handling*
- *Regulatory & practice compliance*

Our research in the BRIDGE Security Work Package has studied mechanisms to discover and extract serialised supply chain data in order to share that data securely across organisations. The Discovery Service developed in collaboration with BRIDGE WP-2 is a mechanism to find serial-level data from previously unknown parties and to optimize the communications load on the information services, the network and the requesting client applications.
Sharing RFID supply chain data will not succeed unless we can guarantee that we can maintain the value of the critical information that we are sharing. In technology terms this means that we want the owner of the information to have security mechanisms to control who has access and to control the integrity of data.

Need to Define Standards to Enable Global Visibility

WP4 works closely with IETF ESDS and EPCglobal Data Discovery JRG working groups to ensure that any efforts to develop Discovery Service standards fully consider the implications of security during the requirements and early design phases.
In an EPC global network, EPCIS allows the pair -wise exchange of information only when both parties know each other's address, and that relevant information exists within such repositories. Therefore, a mechanism is needed to allow each organization to register that it holds some information about an individual object, so that another organization may discover who holds information about it. This is the role of Discovery Services.
A Discovery Service acts as a directory or network of individual serial-level information resources. As such it will only have value if organisations can be persuaded to participate. This value is determined by the business efficiencies and new processes and applications that may be enabled from external data, and offset by the risk of trusting external data sources for often critical supply chain processes. Any organisation participating in a DS (so that others can find their supply chain or other serial-level data) must be able to protect their commercial interests. Similarly, those looking for data must be able to protect their own confidentiality and retrieve (and ultimately use) data with confidence.

Preserve both the integrity and confidentiality of Supply Chain Information

A DS can be implemented over different communications models - for example query propagation or directory approaches. While these different approaches have some security implications, what is critical is the ability to perform authentication and access control to preserve both the integrity and confidentiality of any data held within, or passed through, the DS. The control of information is where our work has focused.
A DS must protect data (about information resources) from multiple parties. Each party must be able to specify their own access controls. The security permissions for data related to one specific EPC number will be different for the various players. For example a retailer wants only to share inventory information with its own suppliers, and only for the supplier of an individual product. A DS could also specify emergency policies that only apply in exceptional circumstances (e.g. food contamination, product recall, legal investigation, government emergencies).

Designing, Integration and Testing

A working prototype is currently being developed between BT and AT4 wireless. The access control framework is controlled by distributed policies. For example, the Discovery Service will release records only when allowed to do so by both the Discovery Service policy and policies delegated to the Discovery Service by the publishers (owners) of the information records. It is expected that we will use, wherever possible, standard security mechanisms already developed for general Internet services, such as those under OASIS.
Access control policies can be implemented by using a flexible access control policy language such as XACML. With our research, the EPC global network can now be protected by a security framework that enables companies to specify automated policies to protect their supply chain.

Our security framework is based on Web Service Security open standards providing components such as token issuing and validation points along with policy enforcement and decision points.

The key modules of the framework are shown in Figure 1:

1) *Policy Enforcement Point (PEP) is the logical entity within our Security Framework that enforces policy for the admission control in response to a request from a user wanting to access a record on the Discovery Service.*
2) *Policy Decision Point (PDP) has the job to decide whether or not to authorize the access based on the user's attributes and the action they are attempting to perform of the Discovery Service.. The decision is then returned to the PEP.*
3) *Policy Administration Point (PAP) provides administration, management and monitoring of policies. Where policies are written by individual resource owners they will operate their own PAPs alongside that of the Discovery Service operator. Both access control and delegated policy rights will be used to allow legitimate publishers to control access to their own records (and only their own records).*

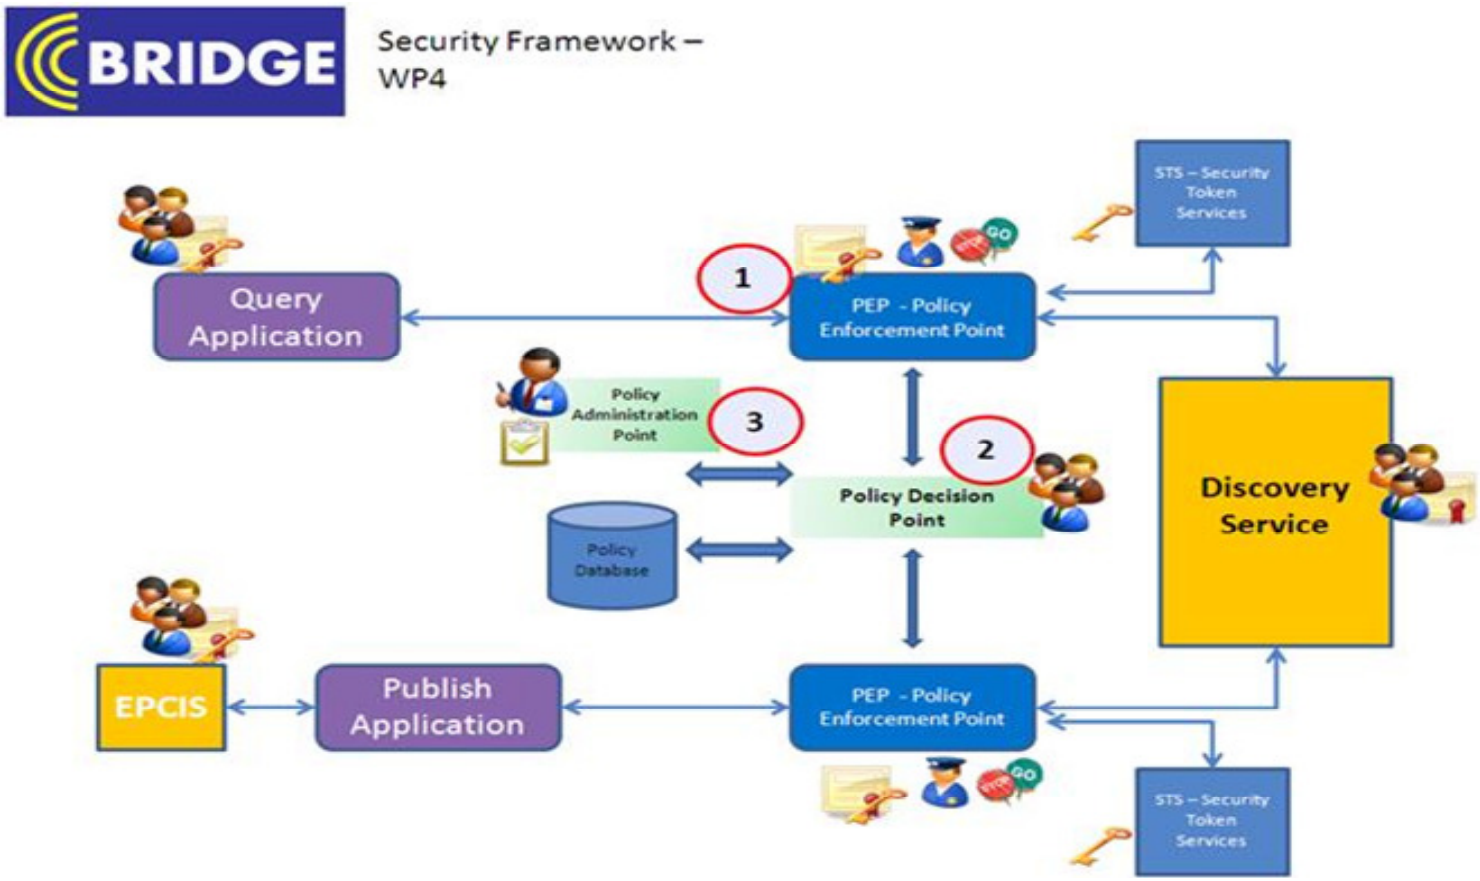

By integrating the Discovery Service into a network access control framework, the security framework provides a high level of security. We enable detection of any unauthorised query but we are also able to prevent any unauthorised disclosure of confidential information. WP4 has published an architecture document and further information on Discovery Service security is available on the BRIDGE EU project website.

Further Work

While most of the technical work in terms of design and architecture of the Security Framework has been done within WP4 more work is still required to evaluate the level of trust and security required in real supply chain environment.

We have worked with the partners in WP9 and WP6 to analyze the business case for RTI and pharmaceutical scenarios. The research team in SAP has examined what critical information is, exploring how composite information released from multiple accesses to various systems can be used to attack business confidentiality. For the future, we welcome the opportunity to work with other organizations outside the BRIDGE project to test and evaluate the potential mismatch between the criteria used to make security decisions and the desired behaviour of the system. We will be interested to test our security framework in the case of security policies for exceptional events - for example a product recall situation where a large number of supply chain partners need to be notified in a very short period of time but where access control still needs to be managed.

For more information on this topic, contact:
Andrea Soppera
(BT, WP4 leader)
E-mail: andrea.2.soppera@bt.com

Calendar of events

CeBIT 2009

3-8 March 2009 - Hannover, Germany
http://www.cebit.de/

RFID Journal LIVE! 2009

27-29 April 2009, Orlando, FL
http://www.rfidjournalevents.com/live/

Second Transatlantic Symposium on the Societal Benefits of RFID

6 May 2009 Brussels, Belgium
http://ec.europa.eu/information_society/policy/rfid/index_en.htm

The Internet of Things Europe 2009 - Emerging Technologies for the Future

7-8 May 2009 - Central Brussels
www.internetofthings2009.eu

Future of Internet Conference

11-13 May 2009, Prague, Czech Republic
http://www.future-internet.eu

IEEE Wireless VITAE 2009

17-20 May 2009, Aalborg, Denmark
www.wirelessvitae2009.org

RFIDSec09 - Workshop on RFID Security 2009

July 1-3, 2009, Leuven
http://www.cosic.esat.kuleuven.be/rfidsec09/index.html

BRIDGE is a European Union funded 3-year Integrated Project addressing ways to resolve the barriers to the implementation of RFID and EPCglobal technologies in Europe. Seven Business work packages have been set up to identify the opportunities, establish the business cases and perform trials and implementations in various sectors including anti-counterfeiting, pharmaceuticals, textile, manufacturing, re-usable assets, products in service and retail non-food items. The project includes an important research and development program in various aspects of RFID hardware, software, network and security. A series of horizontal activities is providing training and dissemination services, enabling the adoption of the technology on a large scale in Europe for the sectors addressed by BRIDGE and beyond.

*** Latest news ***

Erratum - BRIDGE Newsletter October 2008 - « BRIDGE presents an innovative sensor tag »
On page 2 of the newsletter, The CAEN Gen2 RFID Temperature Logger semi-passive tag was described as being compatible with ISO/IEC 18000-6c. For clarity, it is worth mentioning that the standard we are referring to is : ISO/IEC 18000-6C:2004 / Amd.1: 2006. The sensor and battery protocols are proprietary, pending the availability of a standard addressing these functionalities.

World's first drug traceability pilot across global supply chain - BRIDGE pilot demonstrates technology is available today to combat counterfeit drugs and improve patient safety in healthcare
As part of the BRIDGE (Building Radio frequency IDentification solutions for the Global Environment) project, the healthcare sector has demonstrated a complete track and trace system for pharmaceutical products in a live operating, international supply chain environment through its successfully completed Pharma Traceability Pilot.

"The success of the pilot demonstrates that the technology required to implement a full international supply chain traceability system is available today. There is no doubt that traceability systems such as that demonstrated by this pilot will in the future have a significant positive impact on the security and safety of the pharmaceutical supply chain. The experience gained by this project will be invaluable when helping our clients to exploit these opportunities" commented John Jenkins, Managing Director, JJ Associates who managed the project.

More information about BRIDGE and the Pharma Traceability Pilot can be found at **www.bridge-project.eu** and at the pilot's information dissemination website, **www.bridgewp6.eu**

http://www.bridge-project.eu
If you have questions regarding BRIDGE
contact: info@bridge-project.eu

BRIDGE

SIXTH FRAMEWORK PROGRAMME

Information Society Technologies