Towards Dynamic Security Perimeters for Virtual Collaborative Networks

Ivan Djordjevic[*] and Theo Dimitrakos[**]

[*]Queen Mary University of London, Mile End Road, London E1 4NS, UK

[**]Central Laboratory of the Research Councils, Rutherford Appleton Laboratory, OX11 0QX, UK

Extended Abstract

The Internet provides a ubiquitous, standards-based substrate for global communications of all kinds. Rapid advances are now being made in agreeing protocols and machine-processible message/document formats that will soon enable open application-application communication and brings about the prospect of *ad hoc* integration of systems across organisational boundaries to support collaborations that may last for a single transaction or evolve dynamically over many years. Effectively, we will witness on-demand creation of *dynamically-evolving, scalable Virtual Organisations (VO)* spanning national and enterprise borders, where the participating entities pool resources, capabilities and information to achieve common objectives.

As a motivating example consider the following scenario (Figure 1). A team of engineers at a University is working on recommendations for a new aircraft wing. As a part of the work, researcher Alice needs to perform material analysis. This is conducted on line by using specialised services provided by different Application Service Providers (ASP1 and ASP2). Such services may include analysis tools (hosted at another institution SH1), pre-existing data sets (held by a remote data archive SH2), additional computational power outsourced to a supercomputing centre acting as ASP1. (For the purpose of this example we assume that the University has fixed contracts with ASP1 and ASP2, and the latter subcontracts with SH1 and SH2 which we do not examine further, and we treat the terms Application Service Provider and Service Host as simplifying abstractions.)

Alice belongs to team of researchers administered by the local administrator at the University. The main activities of the material analysis are executed by end-to-end services CSI1 provided by ASP1 and CS2 provided by ASP2. We assume that CSI1 is using subservices executed in house at ASP1 who is responsible for administering CSI1 and its subsesrvices. Whereas ASP2 is effectively outsourcing some of the subservices needed for executing CSI2 to different service hosts SH1 (analysis tools) and SH2 (data set). Each administrator wants to protect its local "private" resources from the general "public" which may include hostile agents. At the same time seamless interaction between Alice and the end-to-end services, as well as CSI2 and its outsourced subservices, is highly desirable in order to facilitate collaboration objectives, i.e., material analysis. The goal is, as the material analysis proceeds, to create overlaying security perimeters which are protecting the different virtual collaboration teams that may exist at a time (as a firewall would do in a fixed topology) while ensuring the security of each member as defined by the local manager administering it.

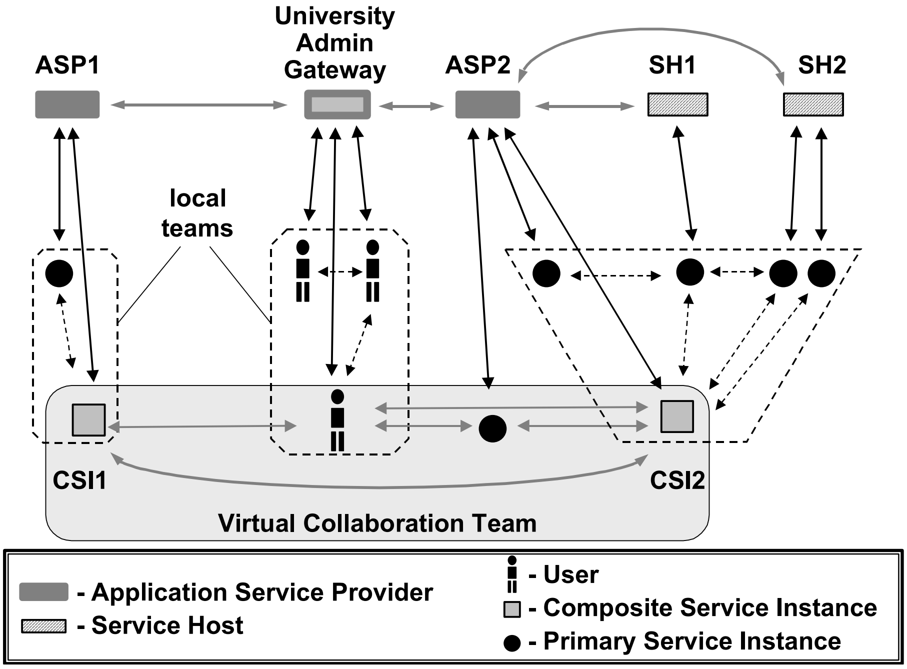
University Admin Gateway
ASP1
ASP2
SH1
SH2
local teams
CSI1
CSI2
Virtual Collaboration Team
- Application Service Provider
- Service Host
- User
- Composite Service Instance
- Primary Service Instance

Figure 1: A motivating scenario[1]

This scenario highlights several issues related to secure collaboration in dynamic virtual organisations.

- Collaboration of resources that are controlled by different institutions. Each institution will have their own policies on access control and conditions of use.
- Resources may be called upon to participate in the task without previous knowledge of the other participants. Trust between resources has to be established in real time on a P2P basis.
- Resources need to be protected from their collaborators and the whole collaboration team has to be protected from outsiders including other entities residing with the participating institutions.
- The same resource may interact in different collaborations. A separation between those interactions has to be achieved.
- Different security conditions may be applied for different parts of the resource, including restrictions on data.
- Collaborating resources may play different roles in their organisation and various collaborations, and different (potentially conflicting) security policies may apply.

[1] Vertical arrows denote association to a local manager. Dashed boxes indicate security perimenters of CCTs. Dashed arrows indicate P2P interactions within CCT. The shadowed box represents the CCT responsible for performing the material analysis.

- There is no centralised administrative point. Security has to be achieved via a devolved policy management scheme combined with distributed enforcement at a peer level.
- Complex trust relationships may hold between collaborating resources (users or services) and their managers: On the one hand, the trust of the other resources in a resource performing its role as expected within one collaboration may evolve as over time, depending on the direct observations of its collaborators, which are collected and correlated by their managers, and also on their reputation base on performance in related collaborations, as communicated by the corresponding managers administering those collaborations. On the other hand, changes of the trust level in a manager may reflect on the rust level in the resources it manages, and vice versa.

A suitable architecture needs to be able to provide a security and trust management infrastructure that meets these requirements. In this talk we will introduce the basic elements of such an architecture, which is currently under development, and gradually explain how it aims to address the above requirements, emphasising on security management and trust/distrust formation and propagation issues.

The architecture provides mechanisms where secure collaboration groups can be dynamically altered in terms of membership and policy constraints. The interaction model of the proposed architecture integrates a layered peer-to-peer model (between collaborating resources and between managers administering resources), with a centralised community management model (between members and their local managers) and a master/slave model (between security managers and enforcement agents). It supports on-demand creation and management of dynamic virtual collaborations in the form of secure groups of peers (users, services, resources, etc.) that cut across geographical and enterprise boundaries. The proposed architecture is being developed with two main goals in mind:

- Enabling communication within dynamically created collaboration groups, that is: secure, scalable, accountable, robust and independent of network topology.
- Enforcing security perimeters, which adapt to the highly dynamic evolution of a collaboration group (in terms of membership and security policy).

These goals are addressed through the following means:

- Certificates to manage CCT membership and privileges.
- Role based security policies describing permissions, prohibitions and obligations within CCTs, set by, and negotiated between, the community mangers.

- Integrated distributed firewall / intrusion prevention and detection mechanism to protect individual members within a collaboration group and the collaboration group as a whole.
- A framework for assessing and propagating trust within and across collaboration groups as a means of assessing confidence in a network entity on the basis of evidence about its observed behaviour and its reputation in different collaborations.