

Background: The Grid Arms Race

Until recently, Grid systems were developed primarily to run over dedicated, high-speed, private networks between supercomputer centres. As a result, they did not need to deal with firewall perimeter defenses that police connections over public network.

Grids use high-speed protocols (that may look like DoS attacks), and externally generated notifications, both of which originally used non-standard protocols that are normally blocked by firewalls.

To overcome these barriers, Grid developers now use firewall-friendly protocols like HTTP(S). However, this only bypasses the network defences, removing control from the network administrator who is responsible for security. When these "tunnels" start to be targeted by malicious users, network administrators will close the loopholes and Grid systems and applications will stop working.

Our goal is to overcome this "arms race" between Grid developers, network administrators and malicious crackers.

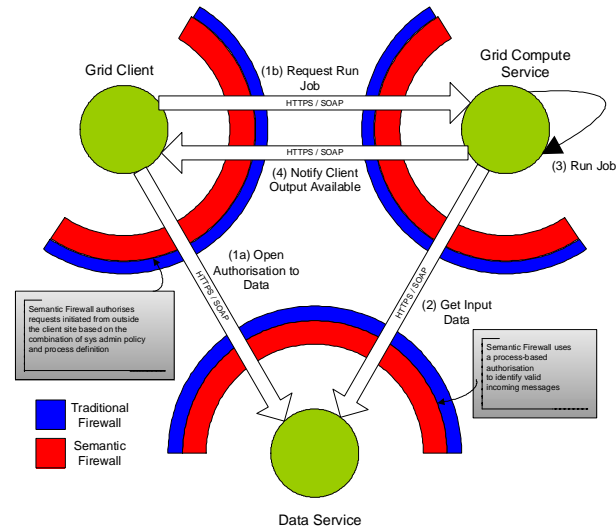
Project Goals:

We are trying to devise a new approach to network perimeter security for Grid systems and applications. Our approach is to implement a Semantic Firewall with the following properties:

- Access to services within the trusted domain is regulated according to dynamic control mechanisms.
- These control mechanisms can adapt to the needs of legitimate applications.
- The device is managed by a network administrator, who can keep control of the overall access control policy for the domain.
- The device uses semantic reasoning to identify and flag (or possibly resolve) policy conflicts, and to determine specific access policies that should apply in a given case.

The project aims to investigate semantic representations needed to express application needs and domain policies, to develop a proof-of-concept version of the Semantic Firewall device, and evaluate it against a real-world Grid application scenario.

Motivating Scenario



In the first 6 months of work, we have identified a useful motivating scenario for the project involving three Grid participants: a Grid client user, a data service and a computational service.

In this scenario:

- The client requests to use the compute service to process data from the data service, and asks to be notified when this is complete.
- The compute service requests the data from the data service.
- Optionally, the data service may notify the client before giving access to the data (not shown).
- The compute service receives the data and carries out the requested computation.
- The compute service notifies the client using a call-back to its domain.

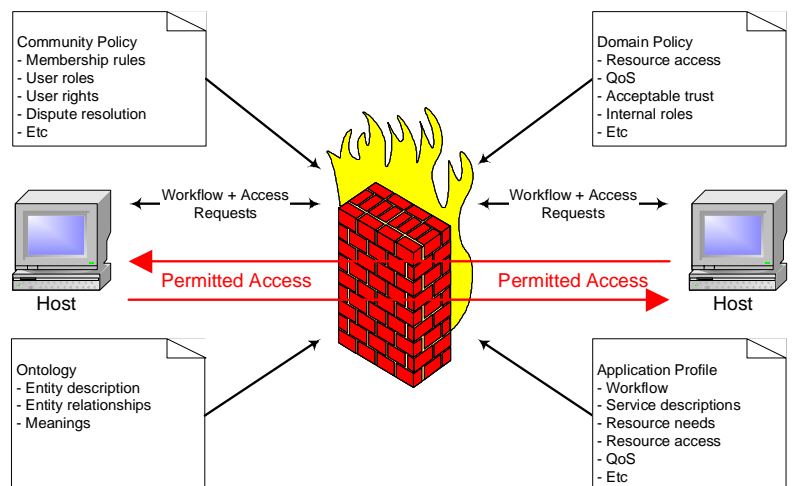
In this scenario we have a good range of "hard" problems: multiple interacting sites operating as a virtual organisation to meet a need, delegation of rights by the client to the compute service, and at least one call-back notification event initiated by an external actor.

The Semantic Firewall

Our vision is a firewall device that:

- Has a set of domain policies governing access - e.g. forcing users to negotiate for resources before running computations.
- Has a set of workflows that an application must carry out - these could be service workflows or client-initiated workflows involving call-back notification events.
- Has access to a community policy (if required) defining the rules of engagement with other sites in the community.
- Has access to a global ontology through which the above policies and workflows can be defined.

The firewall can then process internal and external access requests, and regulate access according to the policies it has been given.



Project Phases

The Semantic Firewall project started on 01 Oct 2003 and will run for two years. The work is broken into four six-month phases:

Phase 1: Investigate requirements and semantic representations that could be used to express policies and application requirements.

Phase 2: Develop a simple prototype platform incorporating access control mechanisms, logging and other basic features.

Phase 3: Incorporate semantic reasoning to produce a proof-of-concept implementation of a Semantic Firewall.

Phase 4: Evaluate this device using a real-world Grid application.

We are now six months into the project, and are approaching the end of Phase 1. This poster gives a preliminary indication of our direction and initial findings.

Progress: Semantics and Ontology

We have almost concluded our investigation of requirements and semantic representations for the project. We have investigated several process-oriented representations for the project:

- DAML-S: good on semantic reasoning, but doesn't handle workflows involving multiple services and clients
- BPEL4WS: over-complex for our purposes
- eBXML: good business oriented process descriptions, but no practical open-source software available.
- WS-Choreography: possible candidate, still investigating open source implementations.

At this stage, we believe WS-Choreography is closest to our needs, but the representation of relationships and constraints is much more explicit than we would like. However, all these languages are still in development and we are feeding input to the DAML-S developments in particular. Our initial choice is therefore provisional at this stage.

Progress: SFW Collaboration

The project started out as a collaboration between IT Innovation and IAM, with support from the EC GEMSS project regarding application examples (from U.Sheffield and NEC).

The network of organisations involved with the Semantic Firewall is now much larger, and includes strong collaborations in the US with Dr Grit Denker at SRI International and Jeff Bradshaw at IHMC, and through them to the wider DARPA community.

We are also now receiving additional industrial support from a major manufacturer!

Progress: SFW Design

We have identified several use cases:

- defining or amending site policies
- deploying services
- registering application workflows (involving clients as well as services)
- processing events and incoming messages

A system design with four main subsystems is proposed:

SFW Knowledge Base

Site policies
Service descriptions
Application workflows

SFW Reasoner

Policy conflict detection
Workflow analysis
Message interpretation

SFW Administration

Static configuration
Logging system
Administration GUI

SFW Communications

Event handling
Message enforcement

The communication subsystem will be based on a Process-Based Access Control (PBAC) model. This was developed in the UK e-Science Pilot Comb-e-Chem, and applied to industrial Grid-based virtual organisations in the EC GRIA project.

The main problem with our existing PBAC system is that it doesn't handle client-side issues like notification, and it is driven at a very low level by Grid services that have to know about their relationships: e.g. a compute service has to know that users must negotiate with a resource allocation service before running computations.

The Semantic Firewall will allow us to express such relationships via site policies, and drive the PBAC system at a much higher level.

Progress: Publications

The first publication from the project is *Towards a Semantic Web Security Infrastructure*, Ron Ashri, Terry Payne, Darren Marvin, Mike SurrIDGE and Steve Taylor, Procs AAAI Symposium, California, 22-24 Mar 2004.

Contacts and Credits

At IT Innovation: **Dr Mike SurrIDGE** (Principal Investigator) **Darren Marvin** and **Dr Steve Taylor** Contact ms@it-innovation.soton.ac.uk.
At IAM: **Dr Terry Payne** (Co-Investigator) and **Dr Ron Ashri** Contact trp@ecs.soton.ac.uk.

The project website is now at <http://www.semanticfirewall.org>.

The Semantic Firewall project receives and gratefully acknowledges funding from EPSRC under the Semantic Grid and Autonomic Computing Programme, and support in kind from U.Sheffield Hospital, NEC and a growing band of industrial and academic collaborators.