Trust Domains

Cloud & Security Lab

Information control within the cloud

Trust Domains is a three year, Technology Strategy Board and EPSRC funded collaborative project involving HP and Perpetuity along with Oxford University, Birmingham University and Aberdeen University.

As the internet has grown to be an increasingly important business tool, companies are looking to use online services to support their collaborative working and information sharing. Collaboration may be internal to a company, with business partners or even part of a dynamic coalition formed to solve a particular problem. The need for online collaboration is growing with globalisation, the adoption of cloud computing, and partners becoming more geographically spread and virtual.

The problem with collaborative working is that we need to share information and hence trust others with our information, persuade others that we are trustworthy stewards of their information and have ways to trust information coming from the collaboration. This suggests that we need ways of deploying systems where information can be shared in a trustable manner and where we can ensure the integrity and provenance of information as well as ensuring it is kept confidential yet available. We describe such a solution as a trust domain.

There are many different technologies, security processes and policies that could be used to create trust domains providing some level of control and assurance. We believe that recent research into trusted infrastructure combining virtualization and trusted computing offers the best choice. Virtualisation is used to create distributed compartments controlled by policies within the infrastructure. These compartments can stretch between machines running services within the data centre to the desktops, laptops and smartphones using information. Trusted computing provides a hardware root of trust for attesting components and ultimately assurance that policies are being maintained. We expect trusted infrastructure to emerge over the next few years enabling new architectures to become possible.

Collaboration can involve many different tasks and with different organisations with different

Example Trust Domain Scenario

An **emergency response** requires multiple parties to share infrastructure and information. When these groups are from different countries or are regulated to keep certain information as "need to know", how do they decide what trust policies should be associated with their assets and how can they quickly stand up services and establish operations?

levels of trust. Information sharing solutions for governments and enterprises are likely to consist of a number of composed trust domains. When solution architects come to deploy such systems they need to address questions such as how do we best configure our systems to allow access to information for those who need it but not to those who do not. How do we achieve the best tradeoffs between risk, productivity and the cost of the solution?

Within the trust domains project we are pulling on and expanding research in trusted infrastructure, security analytics and trust economics with the aim of joining up the security management life-cycle between risk, policy setting, operations and assurance. In doing so we aim to provide tools and technologies to aid information sharing and understand the associated risks.

Within the project we are working on a number of areas:

Empirical studies

We are working to gain a better understand companies' attitudes to trust and information sharing. This will influence the direction of the other work packages as well as providing sample scenarios.

Modelling

Here we are looking at building on techniques developed within the TSB funded trust economics and cloud stewardship economics projects to produce methodologies to model and understand information flow and the risks associated with different trust domain policies and implementations.

As we look to modelling and reasoning about different implementations we need ways of describing the technologies and their management interfaces. As such we are looking to develop a taxonomy to describe trust domains and the underlying technical mechanisms. As we describe the different mechanisms we are looking to methods to understand and verify the properties of components.

Trusted Infrastructure

Work on combining trusted computing and virtualisation to create trusted infrastructure is still in its infancy especially from a practical perspective. Within the project we are looking at what components are necessary to build trust domains, which exist and which we need to build. In doing so we are also analysing the overall stack looking for security issues and ways to improve it.

Example Scenario

An important part of this project is to show how a variety of different research areas can lead to joined up solutions for information sharing. Hence as the project proceeds we will look at various scenarios.

Contact

For further information about HP Labs please visit: www.hpl.hp.com

For further information about the Trust Domains project please contact: adrian.baldwin@hp.com