DEPARTMENT OF COMPUTER SCIENCE
SERIES OF PUBLICATIONS C
REPORT C-2013-2

Rolling out trust management to cloud-based service ecosystems

Sini Ruohomaa, Lea Kutvonen

UNIVERSITY OF HELSINKI
FINLAND

Contact information

Postal address:
Department of Computer Science
P.O.Box 68 (Gustaf Hällströmin katu 2b)
FIN-00014 University of Helsinki
Finland

Email address: postmaster@cs.Helsinki.FI (Internet)

URL: http://www.cs.Helsinki.FI/

Telephone: +358 9 1911

Telefax: +358 9 191 51120

Computing Reviews (1998) Classification: C.2.4, H.5.3, K.4.4, K.6.5
Helsinki 2013
University of Helsinki, Department of Computer Science

DEPARTMENT OF COMPUTER SCIENCE
SERIES OF PUBLICATIONS C
REPORT C-2013-2

Rolling out trust management to cloud-based service ecosystems

Sini Ruohomaa, Lea Kutvonen

UNIVERSITY OF HELSINKI
FINLAND

Rolling out trust management to cloud-based service ecosystems

Sini Ruohomaa, Lea Kutvonen

Department of Computer Science
P.O. Box 26, FIN-00014 University of Helsinki, Finland
firstname.lastname@cs.helsinki.fi

Technical report, Series of Publications C, Report C-2013-2
Helsinki, Dec 2013, iii + 33 pages

Abstract

This report aims to determine how currently available technology satisfies the needs for trust management in cloud-based service ecosystems, and provides some recommendations for enterprises and infrastructure providers on how to proceed on the short to medium term.

Computing Reviews (1998) Categories and Subject Descriptors:
C.2.4 Computer-communication Networks: Distributed Systems
H.5.3 Information Interfaces and Presentation: Group and Organization Interfaces
K.6.5 Management of Computing and Information Systems: Security and Protection

General Terms:
Design, management, security, human factors

Additional Key Words and Phrases:
Inter-enterprise collaboration, trust management, cloud-based ecosystems

Contents

Chapter 1

Introduction

Inter-enterprise collaboration allows a set of independent service providers to focus on their key competences while providing a composed service to their end customers. Two major challenges in supporting the setup and management of such collaborations lies in ensuring service interoperability and finding trustworthy partners.

As cloud computing has taken off, the problem of finding trustworthy and interoperable partners has gained increasing visibility. State of the art solutions on the level of technical software construction concepts and low-level security mechanisms have trouble addressing business-level concerns, because the concepts and mechanisms are inherently technology-driven. The burden of the development and operational environment and ensuring interoperability are currently placed upon the concepts of SaaS (Software as a Service) and Paas (Platform as a Service) [5, 88].

The next necessary step forward from these pairwise connections between service consumer and provider (with SaaS) or service provider and platform provider (PaaS), is the adoption of multi-partner business-service collaborations supported by a "Collaboration as a Service" infrastructure (CaaS) [44, 47, 73]. The CaaS infrastructure enhances the SaaS and PaaS environments with service-oriented computing (SOC, extended SOA) and business process management (BPM) concepts and automated support for collaboration negotiation and administration. While these dimensions have not yet become fully acknowledged in the cloud context, the trend is also supported by work on reusable business process templates in the style of BPaaS (Business Process as a Service) [60], for example.

As an example, consider a travel agency connecting to a hotel, an airline and an online payment service, as depicted in Figure 1.1: the format of reserving a hotel and flights as well as handling the payment is reasonably fixed, but the actors implementing the process can change based on the customer's needs and the service providers' availability.

To address the interoperability issue, we have defined *cloud-based open service ecosystems* as environments for setting up, operating and managing service-oriented inter-enterprise collaborations [71, 46]. The collaborations are supported by infrastructure services such as service discovery, contract negotiation and monitoring, that are independent of technology platforms, and allows the collaborations to effortlessly cross the boundaries of multiple cloud-computing environments. Our focus is on public clouds from the point of view of management of the distributed processing, and fulfilling business-level governance needs.

The need for finding trustworthy partners is addressed through trust management systems, which collect, evaluate and aggregate information needed for decision-making. Trust decisions gauge an actor's willingness to collaborate with its potential collaboration partners, given an evaluation of the risks and incentives connected to the decision. Risk estimations are essentially a model of probable future behaviour, and one input to these estimates is past experience with the

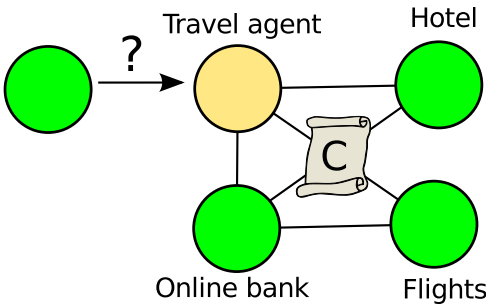

Figure 1.1: A collaboration pattern of four roles can be fulfilled by different service instances, and the actors decide whether they want to work together under the collaboration contract.

given target of trust, i.e. trustee. While rational actors can definitely change their behaviour at will, collecting and disseminating experience of actors' past actions serves two purposes:

First, the dissemination of experience and reaction to it reduces the payout of misbehaviour, as the potential targets receive warnings to steer clear. As a security measure, this provides a layer of defense, a stumbling block that makes attacks more expensive, rather than downright stopping misbehaviour. In an open environment, absolute defenses are not feasible.

Second, the experience serves as a collateral in unbalanced transactions: it creates peer pressure to respect contracts due to the threat of losing future business. Positive reputation becomes an investment the actor has in the ecosystem, and through this reduces their incentive to suddenly change behaviour and effectively exit the system. This underlying commitment required to build up a good reputation, more than the raw information of the past, provides the additional predictability that is valued with reputation systems.

This report analyzes how currently available technology satisfies the needs for trust management for cloud-based service ecosystems, and the gap between the state of the art and our vision of cloud-based open service ecosystems. Based on this differential, we provide some recommendations both from the enterprise and infrastructure provider point of views on how to proceed on the short to medium term.

Chapter 2 presents the key concepts, context and infrastructure of trust management. Chapter 3 describes the different trust-related ecosystem infrastructure services we have identified, and Chapter 4 analyzes the state of the art in standardization and adopted systems. Chapter 5 presents future directions and our recommendations. Chapter 6 concludes.

Chapter 2

Trust management in open service ecosystems

In this chapter, we first present the central concepts of trust management. We then discuss the trust needs created by outsourcing infrastructure, work, services and collaborations into cloud environments. In the third and final section, we take a look at the infrastructure context of where we see trust management services to fall into. This background will provide the basis for the next chapter, where we identify what kinds of trust management services would be beneficial.

2.1 Key concepts

The target of our research is **inter-enterprise collaborations** between business services that have a computational interface, and may connect to real-world resources. For example the aforementioned group of travel agency, hotel, airline and payments handler constitutes an inter-enterprise collaboration that may be in existence for the duration of a single transaction or continuously over months, with new transactions taking place regularly within it. The services make local trust decisions based on local policy when joining and whenever committing new resources into a collaboration.

We define **trust** as the willingness to enter into and continue in a given collaboration with a given partner, considering the risks and incentives involved. A **trust management system** collects, processes and applies the information trust decisions are based on. Although trusting beliefs and behaviour in a general sense are rather emotion-driven in human relationships, we focus on the calculative part of decision-making. Here 'trust' emphasizes an acceptance of a dependency that cannot be controlled, in contrast to 'security', which in the context of systems refers to the control mechanisms.

A key input to a trust management system is **reputation** information, which encodes how the potential partner services have behaved in collaborations in the past. It may also cover related supporting information such as certification of competence, compliance to standards or indicators on the service provider's financial stability, such as a credit rating. Reputation information is used when trying to estimate how the partner will behave in the future, to produce **risk** estimates. We use 'risk' to cover both positive and negative effects that are not controlled by the trustor; risk estimates assign probabilities to outcomes with different impacts with the goal of providing a measure to support decision-making. Risk is always compared to the risk tolerance of the particular action being decided on, and risk tolerance can also vary, based on the business importance of the particular action in a given context.

In addition to the predictive use of past experience, reputation also leverages social control: the danger of losing one's own reputation deters misbehaviour, as the cost of a reputation drop accumulates over time as lost business. This control effect, when connected to a competitive field of business, can have negative consequences through abuse [25, ch. 5]. It is therefore important that the punishment power of reputation is not freely abused. In order to limit the subjectivity of reputation reports, we propose binding them to explicit contracts [70] that define when a collaboration has been successful enough to warrant positive reports, for example.

A trust management system can automate routine trust decisions on whether to collaborate with a given partner, for example to approve the next transaction in the supply chain above as long as everything goes as expected. On the other hand, situations with insufficient information, particularly high risks and incentives or high uncertainty must be left for a human user to analyze and decide on. In these situations, the trust management system should provide information to support the human decision making, acting as an expert system [40] instead.

2.2 A new environment: Outsourcing in the cloud

For a small or medium-sized enterprise, operation in the cloud brings new threats as well as new opportunities.

On one hand, a startup case can be implemented at very low cost, as platform-level infrastructure can essentially be rented based on current need from PaaS cloud providers. This gives more freedom to business experimentation. Software as a Service (SaaS) provides a dissemination model where new services can be made available to clients overnight, accessed and operated online. This, in turn, shorten deployment cycles and makes it easier to react quickly to changes in the business situation. On top of this setting, service-oriented computing (SOC) can provide additional support for discovering new services, and establishing joint business processes for collaborations [59]; the Web Services standards stack is an example of how service-oriented computing is becoming realized in practice.

The new threats that we focus on can be roughly divided into two categories: loss of control and increased vulnerability from centralization.

Firstly, *loss of control* over low-level infrastructure, or external services that form a part of the overall business, appear in the dependency on the external cloud provider or partners within it. On one hand, the monetary cost of outsourcing critical data to a cloud storage is lower and backups may be handled better by the cloud provider than they would be within the organization. On the other hand, the importance of a small customer's data to the cloud provider is different as well: their business does not depend on it. When the cloud provider has an outage or network connections go down between the enterprise and provider, the enterprise operations grind to a halt. Further, if the change of providers has not been explicitly designed to be easy and low-cost, the enterprise client will have little chance but agree to any inevitable changes in pricing and terms during the operation of the service.

From the point of view of privacy and legislation, a startup storing sensitive customer information in the cloud may learn that an international cloud provider may routinely distribute information abroad that by law should stay within country borders. The cloud provider may itself collect information from its customer organizations; beyond direct use by the provider, the motivations to provide this information to third parties may range from industrial espionage, as has been reported for the ECHELON interception system [77] by the European Parliament back in 2001 already, to cooperating with law enforcement or other investigations, as for example the social networking site Facebook reportedly does [41, 24].

Second, *strong centralization* creates a new kind of vulnerability caused by monoculture: like in biology, differences between individual organisms provide improved group protection against external threats such as infections. While it should be noted that a small enterprise is likely to have less resources to spend on ensuring its systems are secure in the first place, heterogeneity does provide a layer of additional protection by increasing the cost of adjustment for the attacker, and through that reducing its payoffs. Other factors being equal, a large provider is a more attractive target for attacks as well due to the potential for larger rewards from getting through its defenses[1]. The content distribution service Akamai estimates on its website that it delivers between 15-30% of the volume of traffic on the Internet, at above 8 Tb/s [81], for example. Its real-time web monitor registers on the order of 200 attempted attacks every 24 hours against Akamai servers specifically [82].

The larger service providers may fail internally as well, both accidentally and deliberately. When the major web host and domain registrar GoDaddy went down for most of a day in September 2012, all its Domain Name System (DNS) services were down for hours, which led to its customer websites not being reachable during the 6-hour outage. The company cites the reason to have been internal rather than an attack, as was originally speculated in the media [30]. As an extreme example of deliberate denial of service, the news leak site Wikileaks was placed under global financial blockade by Mastercard, Visa, PayPal and other major finance companies, who refuse to handle any payments made to Wikileaks. The show of power from the dominant international online payment handlers was not entirely unprecedented in the field of banking [55], and predictably had devastating effects on the donation-based funding of the target[2].

The possibility of interaction between the customers of a cloud infrastructure provider are essentially one step forward from the previous threats of operating within the open Internet: in the cloud, independent actors operate within the same hardware, which has not been a typical scenario for the average enterprise. For example, through identifying a particular target virtual machine's likely location, researchers have been able to launch side channel attacks that leak information in third-party computational clouds such as Amazon's Elastic Compute Cloud (EC2) [68].

In summary, outsourcing brings dependencies that must be considered in the total cost estimations. The trend towards hosting services at large cloud providers seems to move the Internet towards a set of near-monocultures. The systemic shift in the power balance towards a low number of dominant service providers may lead to unexpected risks. From the point of view of a single service provider not in disagreement or competition with its cloud service provider, the main risk seems to be outside attacks, reliability and other users in the cloud.

2.3 Trust management within an ecosystem infrastructure

This section presents trust management as a part of a service-oriented architecture (SOA) [59], which is the paradigm adopted by the Pilarocs inter-enterprise collaboration infrastructure developed by Kutvonen et al. [47], as well as CrossWork [51], ECOLEAD [64] and Web Services[3]. This provides background on where the trust-related services fall into in a larger picture. For further discussion on the overall infrastructure we refer the reader to earlier work [69].

Figure 2.1 provides an overview of the connection between the lifecycle of a single collaboration in an ecosystem and shared information repositories that support multiple collaborations

[1] Symantec's yearly report in 2012 stated that in 2011, roughly half the attacks targeted smaller organizations, possibly as stepping stones against larger targets in the same partner ecosystem [86].

[2] A summary of the press coverage on the financial blockade can be found e.g. at `http://en.wikipedia.org/wiki/WikiLeaks` (Accessed 9. Nov 2012).

[3] For a comparison of the different approaches, we refer to earlier work [44].

within a domain. Individual collaborations depend on information repositories that are a part of the infrastructure to support them. The software development lifecycles connect to the collaborations through providing the software and model artefacts that are used and reused in inter-enterprise collaborations. The contracts, reputation and trust management support are Pilarcos additions to basic service-oriented architecture.

Reusable *collaboration templates* abstract away the individual services' attributes into common patterns, such as in our travel agency example. Best practices in the business domain can be encoded into such reusable models to speed up setting up new collaborations.

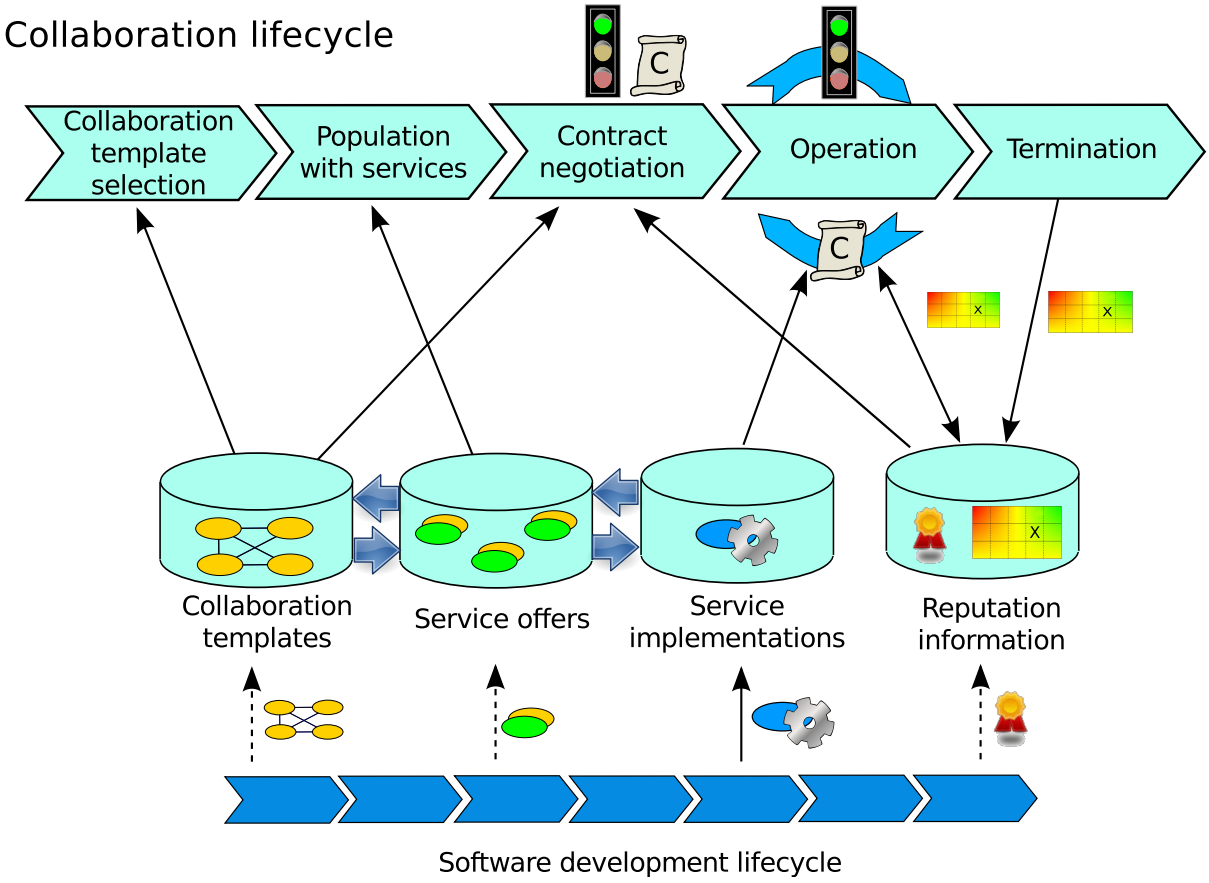

Figure 2.1: Collaboration lifecycles, information repositories and the connection to software development lifecycles in a service ecosystem.

Service offers are set up to fulfil standard roles in the templates. A service offer can indicate that Cloud Hotels, Incorporated provides an interface following the Example Hotel Reservation Standard version 3.0 for booking hotels, and other service-specific details about terms of service and pricing, for example. We expect that the interfaces follow de facto standards through service types [47] used in collaboration templates.

Service implementations, in turn, are fitted to these standards and are able to exhibit varying kinds of behaviour based on configuration. While the repository of implementations is a single item in the drawing, different service providers will naturally have their own implementations which may not be public or shared between organizations; only their interfaces are made public through service offers. Open source tools and Software-as-a-Service solutions may be available in

repository form as well, however.

Reputation information in the form of experience reports and certification is recorded into reputation systems, which are here abstracted into a single repository. In practice, different kinds of service providers are likely to have specialized certification services[4].

At the top of the picture, the collaboration lifecycle follows a high-level pattern where a collaboration initiator selects a template for the collaboration, also known as a business network model in Pilarocs [47].

It then matches available service offers to the roles in the template to populate it: for example our travel agency may have a customer in need of transport and lodgings in a small city in Western Germany, while the agency has existing pre-bookings and other contracts only for Northern Germany. It must therefore try out a service provider it has little or no personal experience on from before to fulfil the customer's request.

The initiator then contacts the prospective partners to negotiate the details of the collaboration contract. The collaboration template may name a set of standard terms that are widely in use in the business; this helps provide structural assurances for the actors that despite the unknown "faces", the general setting is business as usual. The details to negotiate, then, are pricing and other variable terms, such as whether the hotel provider has free rooms for the desired dates.

Each partner will make a trust decision on whether they are willing to take the chance of committing into the collaboration. As the actors do not all know each other beforehand, they take advantage of reputation information both in the form of certification and experiences of others whom they trust as information sources.

The risks taken are not symmetrical: Flight services may have terms of flight reservations that push the responsibility to the end customer, so they do not depend equally much on the travel agency. A large online bank, on the other hand, benefits from its size: problems with one travel payment do not rock its budget, while for a small travel agency the outcome of serving a single customer well may have notably higher reputational and financial impact. The unknown quality of the hotel and its service provide an information asymmetry for the agency: it would like to ensure not dealing the customer a bad hotel, as some of the customer's negative experience becomes associated with it as well. Meanwhile, the hotel may be considering the value of winning a bid with this particular agency in order to gain a new partner, and weighing it against the number of rooms it has free. It is important, therefore, that each actor makes their own decisions rather than a centralized service concluding what kind of reputation is "good enough" for doing business in a general sense.

While the collaboration is in operation, it is governed by the contract on one hand, and decisions made based on local policy in each service provider's end on the other hand. New commitments, such as new customers joining the initial order, lead to new trust decisions as they change the risk estimate. In the general case, this repetition of decisions provides a good basis for automation, as long as everything goes as planned. The collaboration may be in operation very briefly, such as for the length of a single travel or just a query for availability, or go on for months or more, in which case the costs of setting up the collaboration multiple times are saved in exchange for having to fine-tune the decision-making during operational time some more to handle changing contexts.

[4] In Finland, associations of construction companies have joined forces to set up a centralized actor, Rakentamisen Laatu (`http://www.rala.fi/`) to collect reputation and certification information on e.g. the financial status, references and experiences on service providers in its database. This aims to improve consumer trust in the entire business through increased transparency, and help clean up fraudulent actors from competitions for contract calls that focus very strongly on pricing. Some similarities can be drawn to software business, particularly from the asymmetry of information between consumer and provider.

In the termination stage, responsibilities for later claims must be divided, the appropriate handling of any shared assets, such as customer data, assured, and final experience reports can be provided. For long-term collaborations, intermediate experience reports can be given for example per booked trip instead of waiting until the last trip has been completed.

In the bottom of Figure 2.1 we contrast this process with software development lifecycles. Nonsurprisingly, service implementations are the most notable artifact that is provided into the ecosystem by this process. When a service provider designs an implementation, it may have specific collaboration types in mind where it wishes to hook it into. If the collaboration pattern is new, new versions of templates may be produced as a part of the software development, even. When the requirements for the new software are set, they reflect the service offers the provider plans to make, and deploying the service may involve publishing one or more offers. Finally, if the software development lifecycle involves compliance testing or other formal verification processes, it may also produce some kinds of certification-type reputation information as a side product. In the general sense, however, we do not tie certification, template generation and service offers specifically into software development, as they may emerge independently as well.

Many parts of a collaboration lifecycle are similar enough independently of the specific type of collaboration, and can be supported by infrastructure services [46]. For example a service for finding matching interoperable service offers can simplify the population stage so that every collaboration initiator does not need to figure out a way to navigate the offer repository — or relevant repositories, in case there are multiple domains involved. Technically solid collaboration templates and the involved processes can be designed by domain experts, for example to represent the interests of larger coalitions of service providers. Similarly, contract templates can be designed by external specialists to minimize the need for case-by-case legal analysis particularly for small service providers. We even envision that the technical support for trust decisions can be outsourced, although the policy-setting power should always remain at the enterprise who is choosing its partners with the decisions.

In the next chapter, we will look into what infrastructure services are related to trust management specifically.

Chapter 3

Infrastructure services for trust management

In this chapter, we divide trust management into a set of infrastructure services that operate in the open service ecosystem. The services can be divided into four categories: identity management, a reputation service for experience sharing, a notary service for addressing subjectivity of experience reports, certification services, and finally, a trust decision service or trust brokering using the aforementioned set as its input. The dependencies between these different services are shown in Figure 3.1.

To contrast these abstract services to reality, in the next chapter, we provide an overview of existing standards and commercial services. While the currently available services do not cover the services described in this section specifically, they share some characteristics that can be learned from. Currently the dominant type of ecosystem service is recommender services based on aggregating customer feedback, i.e. end-user reviews of businesses, or customer-to-customer electronic commerce (eCommerce). This can be seen as an early precursor to the reputation service indicated here.

3.1 Identity management for organizations and services

In order to set up legally binding contracts in an electronic environment, a form of identity management is required that maps an organization as a legal entity into the digital credentials of an online business service.

This form of infrastructure service is a minimum requirement for signing contracts between the services, and as such is not a special requirement for establishing reputation-based trust management in the ecosystem. It has strong implications on the operation of the reputation system, however, and therefore warrants discussion here.

We expect that acquiring a new organizational identity is similar in cost to setting up a new legal entity, i.e. registering an organization. The providers of this level of electronic identities can be expected to be connected to the public body that controls the creation of legal entities in the first place.

On the other hand, we also maintain that reputation must be collected per business service, as the granularity of service provider reputation is too coarse. A large company can operate hundreds of services, and it may have strong expertise in some of them and perform poorly with others.

In addition, some services offered in the ecosystem may be composite services themselves, provided by multiple organizations and combined under a single interface. The reputation of such

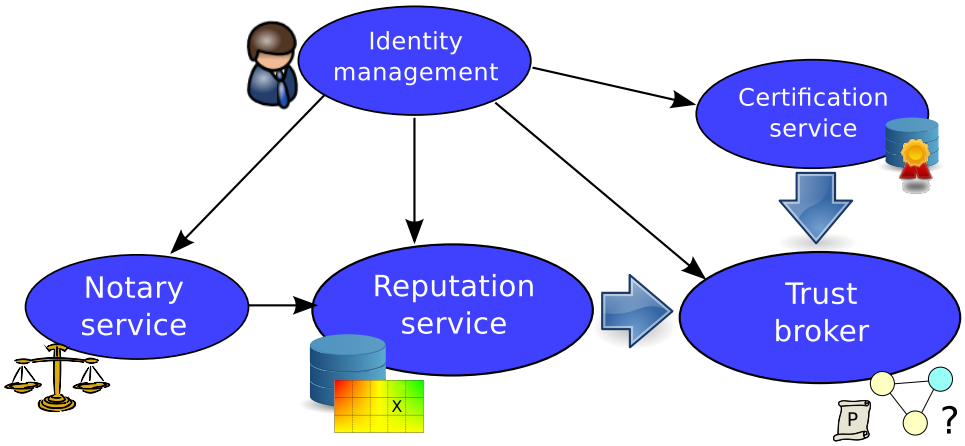

Figure 3.1: The interdependencies of different trust management infrastructure services. Small arrows indicate functional support, while large arrows indicate the service provides input for another.

composite services cannot be directly calculated from the reputation of the participating service providers, as one of them may cover the core service and others simply have supporting role. If a large credit card company covers the payment traffic of an online shop, we cannot say much about the reliability of the shop based on the reputation of the credit card company as such.

Recording experiences by service means that new identities in a reputation system are straightforward to set up. However, connecting the service identities to their providers makes the situation slightly different from ad hoc networks or even electronic commerce between private people operating under pseudonyms: the services are still connected to their provider organization, and therefore an organization establishing a thousand drone services to overwhelm a reputation system in a Sybil attack [18] should be detected and treated differently from the reports of a thousand services that are all from different organizations.

The need for identity management is based on mapping electronic entities to real-world legal entities, which is also at the core of why this service would have users.

The service provider must be trusted, so the task would fall naturally to large ecosystem infrastructure organizations such as operators. This trust can be problematic. As an example of a related but different service, Verisign and other businesses providing public key certificates currently verify by their signature that a given company operates under a given domain (DNS) name when providing the certificate. Their trusted root certificates are then installed to web browsers in an attempt to simplify the process of public key exchange between the user and the website they are visiting. This model is criticized for providing a false sense of security for multiple reasons [39]: 1) a certificate authority can accidentally fail to verify that the person applying for a certificate actually represents the organization, which breaks the chain of binding a DNS name to an organizational name, 2) a certificate authority can deliberately cooperate with an attacker or have their private signing keys compromised, which breaks the authentication procedure itself, and most notably, 3) it can easily mislead a user into thinking that they are visiting the *correct site* rather than that their connection is encrypted. This correctness is not actually verified in the process at all. In reality, advanced "typosquatting" with similar DNS names chosen for phishing sites (for example a Hawaiian bank has a DNS name `www.hawaiifcu.com`, while a phishing site registered a DNS name `www.hawaiiusafcuhb.com`) can still get their certificates signed by Verisign, because the certificate authorities do not claim responsibility for what the names resemble, just what their exact form is. For an overview of the problem, we refer to an article by

Jøsang, which provides examples of all three categories [39].

3.2 Reputation as a service

A reputation service produces reputation information that can be used as input for the service providers' local risk analysis to complement their first-hand experiences.

The motivation for a reputation service is twofold: On one hand, estimating the risks of a collaboration benefits from information on past behaviour. On the other hand, the existence of visible reputation acts as a form of social control: the act of measuring affects actors' behaviour, particularly when it is used as a basis of decisions.

The value proposition of a reputation service is that background checking new partners is laborous and few good information sources are available for the task (see Chapter 4). The high uncertainty stifles urges to experiment on new partners, and as we discussed earlier, the cloud otherwise provides some increased support for positively opportunistic trying new things out. In addition, we estimate that providing structured reputation information that it can be used for automating trust decisions would be worth a subscription payment. A reputation service can provide the experience information raw, or provide added-value services based on providing analyses of the results.

A reputation system can also a part of the services of an ecosystem provider, in which case it brings value to the ecosystem itself. For example online auction websites provide reputation information as a part of their business, where the income comes from the completed sales instead.

While the reputation system sits at the very core of reputation-based trust management, it is not unproblematic to establish. We would recommend that reputation systems be operated by "neutral", democratically run bodies such as networks of service providers in a given domain, that a common contract be set up on submission of a certain baseline of reputation information to make it essentially compulsory rather than a volunteer activity, and that all collaboration contracts map different kinds of outcomes explicitly to reputation reports in order to set up a basis for punishing false reports.

The following discussion provides rationales for these recommendations. We discuss

- the potential effects of introducing a new explicit measurement to the ecosystem,
- the importance of the neutrality of the reputation service,
- providing a motivation for participating in submitting experience information, and
- the need for controlling subjectivity in reputation reports.

Measurements are always an abstract representation of what goes on in reality: a symbolic system that tries to describe an analog world. Research on a related field of making systems auditable [62, 63] indicates that actors can react to outside measurement that affects them in different ways: In one extreme, the measurement may begin to direct behaviour very strongly. For example in science, an increase in publishing just for the sake of publishing is a side effect of measuring a scientist's work based on number of publications generated from it. It affects a scientist's reputation, in other words. Representing reputation as a counter for number of collaborations that have gone well, for example, will motivate service providers to chop all their joint activities into minimal experience units to maximize their number, and to enter into trivial collaborations.

In the other extreme, the measurement can become abstracted far enough to separate it from reality. For example a product's carbon neutrality, in the sense of the production tying down

as much carbon as producing it releases, has experienced a first step of abstraction in that the product owner may just give money to an organization that promises to affect carbon emissions elsewhere. If the activities of this organization, in turn, are not tracked in any way that relates to actual changes in carbon emissions, carbon neutrality becomes a term that means "have paid money for the privilege of carbon emissions", and taken to an extreme this can turn into receipt sale circles that have nothing to do with carbon any more. In the context of reputation, two or more organizations can agree on collaborations that require minimal effort from either of them, and simply mass-produce positive experiences for each other.

In both cases, reputation systems have failed to provide additional benefit to the ecosystem. The reaction should therefore be somewhere in between: reputation information should provide a way for promoting the trustworthiness of an organization.

The neutrality of the reputation service is important because the producer and handler of reputation information uses power in an ecosystem, somewhat similarly to how the judicial system and police are granted power over other actors in order to maintain social order. In this case, it essentially advices collaborators on who to do business with, somewhat similarly to credit rating companies.

This means that the reputation service provider must be trusted, and it should have an interest to remain fair. The conflicts of interest between the provider and its users can be subtle: it has been argued, for example, that reputation and trust systems set up with the interest of making the ecosystem more attractive for transascting in do not necessarily aim primarily at the informed decision-making of their users, but simply encourage more transacting also when it may not be the most sensible course of action [19]. We discuss the observed problems of reputation extortion and conflicts of interest further in the next chapter.

Providing a motivation for participating in the reputation system is necessary, because the service is dependent on collaborators submitting information to it. Feeding in information to a third-party service takes effort with no direct returns, while taking advantage of everyone else's efforts is the main reward. This setting can easily lead to a tragedy of the commons, where the service is desired by all actors but no one is sufficiently motivated to provide the support for it because it would benefit a freerider majority more than itself.

We find that payment in return for submitting information is not likely to work out by itself, because of the same misdirection of effort as we argued before: submitting as many reports as possible (or reports that simply repeat the majority view in order to appear "correct", for that matter) should not be the primary motivation.

In addition to cost in terms of effort, the reputation system itself is a privacy tradeoff for service providers. If all experiences are public or even submitted to a third party service, they provide a basis for traffic analysis: As an example, it may not be necessary to hear the exact content of the phone calls to figure out something is going on if corporate leaders in Microsoft and Nokia suddenly start to call each other much more often than before. Revealing long-term key partners may be undesirable for an organization as well: as a naive popular example, busy parents may not want to advertise their excellent babysitter too much to not have to compete with everyone else for the service.

Privacy concerns may not become an overwhelming issue for all types of collaborations. Advertising reference customers is likely to be less problematic than other dependencies, for example. For fields where privacy is a major issue, the problem could be alleviated by for example only reporting collaborations that caused negative experiences in detail, or hiding away some partner information through a pseudonym system, for example.

Towards the main purpose of providing a motivation for participation, we find that demanding explicit contracts for participation in order to use the system are necessary. The motivation for

not breaking a contract builds then on contract law, and the required compensation and other consequences for breaches. On an abstract level, service providers essentially sign up for an improved ecosystem with social control, and commit to doing their part in supporting this. As a result, reputation information will not be entirely public either — unless it is limited in visibility to those who enter into the reputation network contract, the motivation for signing up falls down to altruism and general benevolence again.

Explicit agreements are also required between the collaborators to control the subjectivity of reputation information. In a contract-based reputation system, service providers can agree in their collaboration contract that the output of providing the service as agreed is a specific kind of experience report.

This contract should include an agreement on which reputation system or systems the report should be submitted to. It serves the target's interests to have positive reports disseminated widely, and the source's interests to have the impact of negative reports as high as possible to ensure the experience provides a motivation to respect the contract.

We see two major approaches in research for producing reputation information: experienced based on subjective opinion, and experience based on documentable contractual agreements. Subjectivity is defined as depending on the subject, in this case, the information source, while information denoted as objective should be independent of the observer [78].

The common approach, inherited from reputation systems between private people and customer feedback to companies, involves voluntary reports based on subjective opinion (see e.g. Farmer and Glass [25]). In the originating systems, which are further discussed in Section 4.2, there is minimal infrastructure available and the agreements between the actors are more implicit than explicit. As a result, the truthfulness of a report is difficult or impossible to argue about.

Subjective opinion cannot be proven to be false. Systems based on subjective statements are therefore also particularly susceptible to manipulation: misbehaviour as a recommender in providing false reports is hard to punish when there is no agreed-upon way to measure the accuracy of a report.

Contract-based approaches are less subjective, and we find them therefore better suited for inter-enterprise collaboration. Besides documenting what the service provision agreement between the two actors is, the literal contract agreement should document what kind of experience report different outcomes should provide and what kind of evidence should be presented to support it [70]. This kind of experience is from the point of view of subjective satisfaction more rigid in the sense that it cannot capture unforeseen service failures, such as unexpected delays when bounds have not been explicitly given. On the other hand, contract-based semantics give the experience a meaning that can be communicated to other actors: it fulfils a basic sense of objectivity.

Due to the possibility of interpreting even explicit contract text in slightly different ways, we can never guarantee absolute objectivity of an experience report. Despite this, we can strive towards metrics that the trustor and trustee agree on beforehand. This should improve the predictability of the content of the experience report, and make it possible to argue whether a report is supported by a given audit trail of evidence or not.

3.3 Notarization as a service

A notary service is a supporting service; essentially it is a security mechanism needed for producing objective reputation reports when the interests of the collaborators are opposite. For example when a service delivery fails, it is in the interest of the provider to avoid all repercussions, while the service user would rather demand compensation.

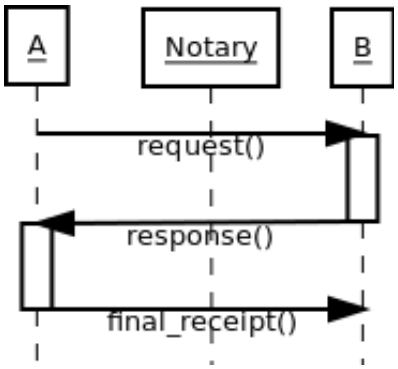

Figure 3.2: A basic message exchange with a final receipt. The notary can act as a middle-man to verify that the protocol was followed.

One of the repercussions that objective reputation systems bring is negative experience reports, and notary services aim to enable that.

Objective experience reports require evidence to support them, and an audit trail of exchanged receipts is the basic building block of such evidence. A basic request-response protocol with a final receipt is presented in Figure 3.2.

This protocol is by nature asymmetric: no matter how many additional layers we add on it, the sender of the final receipt that matters can always deny it received the previous message and refuse to send the receipt. As a result, the service provider B who has provided a service in the response() step will be without a receipt of fulfilling its part of the contract, and A can as a result submit a negative experience report that B cannot prove wrong.

The receipts must also be cryptographically signed, otherwise A can claim that B simply wrote the receipt itself. This together with timestamps that identify the specific transaction can be used to make the receipts nonrepudiable, i.e. make it technically credible that the receipt was in fact issued by the party that signed it, and concerns the transaction it refers to. The topic is discussed at more length in a separate publication [70].

In essence, to get around this asymmetry, we propose to pass nonrepudiable receipts on optimistic fair exchange protocols [7]. This requires a notary to be available to observe the transaction as needed - in the optimistic model, the notary is only called in to observe the *replayed* transaction and to provide evidence that one participant failed to complete it.

The notary needs to be paid for by a subscription model; most notably, it cannot be funded through monetary punishments for observed misbehaviour. Micali and Li et al. explain how the money flows affect the notary's motivation to remain honest [52].

We expect that paying a small fee for the threat of involving the notary if a transaction protocol is not followed is worth it — it essentially eliminates the motivation of misbehaving on the protocol level, and therefore needs to only be involved when the actors are either malfunctioning or considering the misbehaviour worth the reputation loss, in which case they do not really try to gain from violating the protocol itself.

The notary service would be a natural additional service provided at a low cost by the ecosystem provider or e.g. an operator, as it improves the viability of the ecosystem itself and makes it more attractive, but requires very little effort to maintain. In Pilarcos, protocols like the receipt exchange are defined by electronic collaboration contracts, so all the notary needs to do is to see if the message flow is correct and to send out a signed message saying how far it got in case one of the participants fails to complete it.

As a main downside, notarization has limited power if misbehaviour and compensation pro-

cesses in the contract are not possible to make visible on the protocol level. For example not making a physical delivery is invisible to a network-protocol-level notary. The goods can be given to a trusted logistics service to deliver, in which case the receipts from the logistics service can be sent over the protocol. Even then, delivering low-quality goods is invisible to the notary, as long as they are delivered on time. The recipient can make an official statement that the goods were not acceptable, but the notary does not really have the means to verify this. Eventually, extreme misbehaviour has to still be solved through external means, such as a complaint process with the help of a trusted third party, or eventually a lawsuit.

It should also be kept in mind that when the protocol has been followed, the experience must be "positive" even if the recipient was not directly happy with the outcome. This is because the objective reputation system only tracks contract fulfilment or contract breaches. Locally, it can of course make a note to not ever use this service again, and companies can still apply out-of-band means of communicating "I have had problems with this service, I do not recommend it".

3.4 Domain-specific certification as a service

Positive experience reports can make it easier for businesses who have been in the domain for a while to attract new partners. Newcomers will not benefit from a reputation system in this way, however: they have no reports of their past behaviour before they are able to attract the first few partners.

They may take two approaches to enter the market: compete aggressively with other service attributes, such as very low prices, or ask for supporting recommendations from third parties. Providing these recommendations to unknown services is a threat to the credibility of the recommender, but valuable enough as a service that this kind of certification already exists in multiple forms: credit rating services, which look at the financial records of a business and assign recommendations for lending money, compliance certification, which are based on audits[1] or other types of fulfilling criteria, such as participating in specific training.

These certification stamps often have value either based on associations to known names, such as when a large company certifies an expert on knowing their product, or large enough standards that the requirements are known by the trustor evaluating the certified trustee. Both standards-based auditers and enterprises providing their own certificates charge for their effort of analyzing the target system appropriately in order to not risk their credibility in the market, and on the other hand actors entering the market pay for the certification as an investment in future business opportunities.

The service itself is not new, the only new addition that trust management in the cloud may bring is automating the processing of these kinds of certificates once they have been issued. For example certified compliance to a given standard can be considered equivalent to a number of positive experiences in reputation analysis due to the third-party evaluation performed, or simply reduce the risks because certain security mechanisms are verified to be in place.

From a risk management perspective, certification does not imply that the certified service or its provider remain trustworthy, but it does increase the price of entry to the service market — and with more investment made, the more the service provider has to lose by gaining negative reputation and not being chosen into collaborations any more.

The downside of certification is that if a market depends too much on a specific provider for deciding who can enter the market, the trusted third party can use this power to selectively keep newcomers out of the market or charge exorbitant prices for it.

[1] Ries et al. mention SAS 70 II, FISMA and ISO 27001 as examples of auditing in the cloud [33].

3.5 Trust decisions as a service

A trust broker service extends the reputation service by outsourcing also the policy implementation of trust decisions. For small enterprises in particular, it may not be attractive to invest in trust decision infrastructure but they are still interested in collaborating with previously unknown actors and have a need for expert recommendations on whom to work with.

The "brokering", as we see it, comes from supporting the decision-making systems of the involved actors to manage trust between the collaboration participants. Beyond providing a reputation report, the trust broker estimates the risk to the trustor and matches this with its idea of the trustor's risk tolerance.

For this service, the trustor service expresses a trust broker its own trust decision policy and expects it to be upheld. The policy can be provided by the trust broker itself, e.g. by providing a few options to choose from. From the point of view of a small company, this may not be that different from buying virus and spam filtering from their operator.

Like with the other services, the main value proposition is based on automating something that otherwise takes unbearably much effort. In other words, it allows collaborations to be formed that are not otherwise practical, but too risky to leave the background investigation out altogether. The funding could be handled directly through a subscription payment model or some form of payment per consultation.

The trust broker service relies on some form of reputation information as input, which can either be collected by the service itself or subscribed to through an external service.

Our ongoing research gauges the feasibility of trust management services from the point of view of business [72]. For this, we have identified two service concepts for trust management, taking different approaches to reach different target users:

- A low-cost service aimed for small actors. This approach is implemented as web-based Software-as-a-Service (SaaS) solution whose configurability is limited.

- The second approach is a personalized service for larger business, which stand to benefit enough from trust management solution to invest more money and effort into configuring it to their needs specifically. It is implemented as a solution deployed according to the customer's specific needs, involving consultation support and a software solution that can be configured (and continuously reconfigured) to follow the client's own trust decision preferences.

In the next chapter, we contrast the different service types described in this chapter to what kinds of services, standards and systems already exist.

Chapter 4

Standards and services

In this chapter, we look at the current state of standardization and available commercial services in the field. The latter are divided into first-generation reputation measures in business-to-consumer settings, recommender systems for product and service reviews, online image management services and a set of examples that we find to be closest to the services we have described.

4.1 Standardization

In this section, we discuss what the current state of the art in standardization in the field has to offer. WS-Policy provides a high-level grammar for expressing capabilities of services. WS-Trust is a slightly misleadingly named standard for handling security tokens used in SOAP messages. Finally, ORMS aims to standardize the data format for sharing accumulated reputation scores between systems.

4.1.1 WS-Policy: A grammar for policy expression

WS-policy provides a flexible and extensible grammar for expressing capabilities, requirements and general characteristics of web services [10]. These properties can be expressed as policies.

An example policy that can be expressed with WS-Security could be of the form "service must support one of the following: X and Y", where X and Y could be for example RSA or Triple-DES for encryption purposes, or different version numbers for a transaction protocol. It does not specify how to make X and Y understood between the different services, or provide more complex attribute-based expression, for example a specification for how large the keys must at minimum be for the RSA or Triple-DES algorithms. WS-Policy refers to capabilities by name, and therefore referring to specific encryption algorithms with different key length requirements would all require their own unique policy names. The semantics of the capabilities themselves are outside the scope of the proposal.

The suitability of WS-Policy for expressing policies for reputation-based trust management seems to be limited. It represents technical-level choices between technical capabilities such as specific algorithms and protocols. It does not support more pragmatic policies or the priorization between contract terms and local behaviour rules, for example.

4.1.2 WS-Trust: Certification-based access control

WS-Trust is an OASIS standard that provides extensions to WS-Security. Its scope is the issuing, renewal, requesting and validation of security tokens as defined in WS-Security [8], as well as

mechanisms for key exchange. WS-Security, in turn, deals with message signing, encryption and attaching security tokens to messages to provide evidence of the sender's identity [8].

While its name may suggest otherwise, WS-Trust has very little to do with reputation-based trust management.

In mid-1990s, Blaze, Feigenbaum et al. coined the term "trust management" to refer to an emerging certification-based approach to access control [13, 12]. The idea was to make a separation from existing approaches to the same access control task, but the term has later been found a misnomer that adds unnecessary confusion to the field [31].

The certification-based approach known as trust management has its roots in flexible access policies proposed by Blaze, Feigenbaum and others. The three most well-known systems in the field are PolicyMaker [13], KeyNote [12] and REFEREE [16]. In the Web Services context, the approach is also adopted by the Web Services Trust Language WS-Trust [57].

The evaluation of a certification-based policy is based on the exchange of cryptographically signed certificates or security tokens, which can be used to express authorization to access a specific service directly, or more abstract notions such as identity, group membership or delegation of authority. These kinds of general-purpose certificates can be used as a basis for access policies such as "user can present sufficient proof that it is a business service of Partner A, or of a subcontractor whom A has authorized to operate in its stead for this purpose". This way, it is not necessary for the service provider to produce new authorization certificates every time Partner A decides to authorize a different subcontractor, nor does the service provider even need to know who exactly is subcontracting for A at the moment.

4.1.3 ORMS: Open Reputation Management Systems

In 2008, OASIS launched an Open Reputation Management Systems (ORMS) technical committee aiming to develop common data formats for representing reputation data and standard definitions of reputation scores [56].

The scope of the work included gathering requirements through use cases, developing a framework for reputation data gathering, and performing security risk analysis and establishing profiles for best practices [48]. The framework included common data models, XML Schemas for representing reputation management system data and reputation values [74], and standard ways of exchanging reputation claims between systems. The definition of algorithms for computing the scores was out of the scope of the committee.

The 2009 Working Draft [48] of the framework has a strong focus on numerical reputation data, i.e. composed reputation scores. Attributes represented in the reputation data format include "score", "mean", "standard deviation", "sample size" and identifiers of the subject and evaluator (reputor).

The standard presents a syntactic frame for exchanging reputation data, while the semantical interpretation of the scores, i.e. what a reputation score of "4.5" actually means, remains outside the scope of the work and will have to be specified by each reputation system independently. Rating scores by themselves are not sufficient to convey the meaning behind them [3].

For addressing the common understanding of reputation, ontologies for reputation information has been analyzed by e.g. AlNemr et al. [2, 1]. Translations between different reputation score scales of two reputation systems and compensating for different bias within them has been studied by Gal-Oz et al.q [28].

4.2 Commercial services

This section provides an overview of the commercial services currently available that are aimed towards fulfilling similar needs as those addressed by this report. The broad categorization is chosen due to a lack of offerings for reputation services that would be directly usable for inter-enterprise collaboration.

4.2.1 First-generation reputation measures

In the brick-and-mortar business environment, available reputation measures have generally been connected with return of investment aspects. The reputation of a company in terms of a target for borrowing money or buying stock may have been estimated through their share price at the stock market and various ranking lists, such as the Fortune 500 [26], which ranks corporations by their gross revenue in the United States.

The reputation of a company in terms of investing to it and its good behaviour as a business partner are separate matters. At best, if this kind of information can be reliably tied to the service provider, it can be treated as positive evidence that the service provider company exists on the market and is less likely to be a fraudster out to make quick money.

In general, evidence that an organization providing a service is a serious actor and has something to lose may assure a trustor that the service provider will not suddenly desert the ecosystem altogether. On the other hand, unless the ecosystem forms a notable part of the revenue of the organization, its continued existence and continued presence in a given ecosystem are not strongly connected. For example large technology corporations are well-known for discontinuing experimental projects that do not fulfill expectations[1]. In order for reputation to have a social pressure effect, the participants of the ecosystem must anticipate continued interactions with other participants. This creates a "shadow of the future" [9], which is essential for the effectiveness of reputation-based sanctioning [66].

In terms of debt handling, another financial reputation service is provided by credit rating companies, which use a measure based solely on local judgement and experience to translate public and private information into a credit rating. The scales in use are extensions of a discrete scale from D to A, D being a company in default and A doing well; for example the "AAA" grade of Standard & Poor's [79] is better than "A", while "A+" is slightly better than "A" but below "AA-". Again, while this measure targets a different category of behaviour, it can be used as background information: a company at risk of default can be considered a more risky partner particularly in a venture where the partner buys services on credit.

In the credit rating business, the trustee generally pays the evaluator to rate its debt issues, which may cause a conflict of interest: the rating may be artificially boosted in order to please the paying customer. The reputation of the credit rating agency becomes particularly relevant in such a model; for example Hill analyzes the need for regulation of the credit rating business from a legal point of view [35]. In some cases an agency may provide an unsolicited rating as well; in the case of a negative rating, the agency may face legal retaliation [38].

As a third category, one approach to handling disputes in service provision between consumers and business has been adopted by accreditor and consumer protection listing organizations such as the Better Business Bureau [11] in North America, the National Consumer Agency of Ireland [54],

[1] The topic is attractive to report on. For reference, a Google search for "google pulls plug" provides over 60 000 results, one high-profile discontinuation being the Wave service [27]. The equivalent search on Yahoo! provides 7500 results. Microsoft gives almost 70 000 results and Apple 45 000, but a smaller portion of the news is about online services.

and TrustMark [84].

The Better Business Bureau provides consumer dispute resolution services and accredits businesses. On one hand, it forwards complaints to organizations, and as a form of reputation information on how the organizations follow their contractual duties, it shows the complaints and their resolution status for the past 36 months in response to queries about the organization.

The National Consumer Agency of Ireland is a public body that provides consumer protection lists. In other words, misbehaving service providers are blacklisted by the agency. This service provides only negative reputation information, which has been processed by the agency into essentially expressed trust decisions: do not trust these organizations. The basis of the blacklist campaigning has been in naming retailers who are caught breaking consumer law. The agency has also intervened in high-profile consumer disputes.

TrustMark is a government-licensed not-for-profit certification service operating in the United Kingdom. It specializes on the field of home repair, maintenance and improvement. Within this field, TrustMark provides accredition to service providers; its self-expressed goal is to protect consumers from rogue traders. The consumer law enforcement body, Office of Fair Trading [58] is a separate actor which does not actively provide positive or negative reputation information to outside actors.

Legislation-based blacklisting services can provide a strong deterrent against violating laws, but the issuers are by necessity local actors. As such organizations essentially exercise policing rights, they receive their authorization from the local government. Accreditors can in theory operate both under government support and independently; for example large corporations can provide accreditation to partners in customizing their (software) products or providing technical support for them, similarly to brand-based car maintenance services.

The Better Business Bureau disseminates the least preprocessed reputation information in the form of dispute information, although it also provides summaries such as an organization having been a "BBB accredited business" since a given time; failure to handle disputes according to the accreditor's rules leads to the loss of accreditation. The dispute information database, while not based on explicit contracts of any kind, bears the closest semblance to our proposal for reputation systems based on objective reputation information [70].

4.2.2 Recommender systems

A number of recommender systems have arisen from consumer ratings on goods and businesses, such as the Epinions product review site [23], the Tripadvisor review site for hotels, flights, vacation rentals and restaurants [83], the eBay reputation system for consumer-to-consumer and other small sellers [20], and the Amazon store and recommendation system focusing on books and other media for sale on the site [4].

Epinions is a third-party review system of products; recommendations are made based on numerical ratings in different categories, such as ease of use and durability for household appliances. It has a feature for users to indicate which reviewers they trust, i.e. whose reviews they consistently find to be valuable. The trust information is visible to other users as well, which has attracted research on the influence of trust networks of Epinions users on the system [50, 49]. Summaries of the trust expressed by other users, specialization areas and number of reviews written by the author of a review are shown next to each detailed review, which aims to help users estimate the credibility of the review.

Tripadvisor is a worldwide third-party review site on hotels and other services for travellers, which reached a total of 75 million recorded reviews in 2012 [83]. Like with Epinions, reviews consist of numeric ratings in multiple categories, such as food, service, value and atmosphere for

hotels. The detailed rating pages make it possible to see beyond rating averages to how strongly different reviews disagree with each other and to read the textual comments. On the other hand, the rankings of hotels and other services are quite prevalent on the site, and they are based on calculated averages which are the same for all users. These scores make attractive targets for manipulation [36].

The eBay online marketplace originated as an online timed auction system, but has since expanded to more traditional shopping as well. It provides a seller reputation system to protect the interests of buyers who are required to first send in the money before receiving the bought goods. While there are many competitors for the site that hold large market shares locally, such as Taobao [80] in China and Huuto.net [37] in Finland, worldwide eBay is among the largest online retailers. Its size has attracted elaborate attacks, such as setting up a fake car auction with a high-reputation vendor to attract a cash buyer in to be mugged [34]. In this case, the reputation system, which can be manipulated, creates a false sense of security.

Issues with eBay have also made the value of reputation apparent: feedback extortion [21] has emerged as a genuine problem. Problems with e.g. punishing negative seller feedback with negative buyer feedback, and threatening negative feedback if additional services were not provided lead to system changes to reduce reciprocation misuses in 2008 [76], which in turn created an impetus for frustrated sellers to launch a competing reputation system for eBay users [22]. As we have argued before, subjective reputation reports are problematic as a power balance mechanism in online commerce, because they become an attractive target for fraud themselves [70].

The Amazon recommendation service rates products that are for sale at the Amazon online store; the business model of the service is to increase sales. In addition to editorial reviews, customers can provide reviews of the products, with a single five-star numerical rating used to calculate an average. As with Tripadvisor, he detailed review summary provides a way to see how much reviewers disagreed, and some minor reputation elements of review writers are provided, such as "Top 500 reviewer" based on other users' votes that the provided review was helpful. Again, the large market share of the site [75, 14] attracts manipulation, such as buying reviews [61] and the beneficiaries such as authors and publishers writing positive reviews for their products [36].

A key problem with current service and product recommender systems based on end-user feedback and global scores is susceptibility to manipulation; further examples are provided for example by Hu et al. [36]. Larger recommender systems become attractive for more resourceful attackers until the information in the system becomes so untrustworthy that the usefulness of the system as a whole degrades. All open systems relying on equal-weight votes to produce global rating values that are the same for all users are vulnerable to Sybil attacks, where new user accounts are set up to stuff the ballots and gain more influence in the global scores [15].

Currently systems measuring customer satisfaction do not measure the credibility of the provided information, beyond trying to detect obvious spam postings. A handful of research proposals have been made to address this (e.g. [85]). The Epinions approach of indicating trusted reviewers and giving more weight to them to produce a personalized reviews is a step towards more selective use of recommenders, which is also more resistant to ballot stuffing and other recommender system spam; similar proposals have been made in research by e.g. Gal-Oz et al. [29, 32] based on user similarity and by Kinateder et al. [43, 42] based on an evaluation of the usefulness of a recommender's earlier recommendations.

While the quality of recommendation information remains a lucrative target for manipulation, this gives rise to a second layer of business: recommender systems manipulation services.

4.2.3 Online image management

For many businesses, the need for online image management is as at least as burning as traditional needs for visibility and marketing. The market is vast and the offerings, such as information sources, are scattered and even unreliable, so this field too attracts service providers to sell products whose implementation can fall into a gray area of manipulating popular information sources. This kind of market pressure, in turn, makes the information sources even less reliable.

As popular recommender systems have emerged, the traditional example of search engine optimization (SEO) has been expanded to related areas such as recommender system monitoring. Search engine optimization focuses on improving the visibility of a website or web page in search engines' search results, specifically the algorithmic ("organic") results as opposed to the paid advertiser locations. One approach to this is to organize the organization's website in a way that is easier for search engine crawlers to parse; this generally benefits all parties. Another, manipulative approach ("Black hat SEO") is to produce link spam that makes the website look more relevant than it is; this reduces the usability of the search engine, and can attract retaliation from the search engine service.

As an extension to search engine optimization, monitoring and improving results also in the aforementioned recommender systems has become a target for commercial service provision. The Reputation.com [65] online service falls within this category. Besides search engine optimization, it offer includes monitoring a number of sites collecting customer satisfaction information to collect an overview of reviews the company's products are getting, and tools to solicit reviews from customers.

In a somewhat more disturbing trend, recommendation services themselves may start to sell online image management services, which concentrates the power of storing actors' reputation and the profit-making of manipulating it into the same actor, while a separation of concerns would be important here to keep the incentives of the service provider and the market aligned. As an example, three class action lawsuits were filed against the local business recommendation site Yelp! [87] in close succession in 2010, on the basis that the site offered ways to hide negative customer reviews to those businesses that bought advertising from it, and allegedly also marketed this ability aggressively enough for the practice to be considered reputation extortion [17]. In the case of Yelp!, its automated review filter, which is supposed to identify fraudulent customer reviews and not allow them to skew the scores of a business, has added a layer of unpredictability and opacity to the reputation service, and due to heavy criticism the company eventually agreed to make also the filtered reviews accessible to users, among other things [53]. The service itself is still criticized for essentially interfering with business opportunities without consent and shared norms; this reflects how powerful actors simple online reputation systems with a critical mass of users may become in the market[2]

While it should be noted that not all service providers in these fields are participating in the vicious loop of information deterioration, the structure of the market itself remains problematic: centralized, easy-to-use solutions to online reputation are also the most prone to and attractive targets for manipulation. As reputation attracts business, it can be seen as a form of currency, yet unlike the currency market, computational reputation is still missing even localized regulation: anyone can set up a new "reputation bank", there are no standards for transfers between two banks, and false currency printing is commonplace.

These issues must be addressed before computational reputation can gain wide adoption for any serious application areas. It will also require legislative support as a final resort, as technology alone cannot solve this level of social issues.

[2] In 2009, the value of Yelp! was informally estimated to be around 500 million USD according to reports on Google and Yelp! negotiating on a buyout [6]; Yelp! reportedly rejected the offer in the end.

Chapter 5

Recommendations for the future

Our goal is to roll out reputation-based trust management into open service ecosystems, with a focus on the current trend of cloud-based ecosystems. The motivation for reputation-based trust management is twofold: it acts both as a behaviour-based security mechanism from the point of view of single enterprises, and a mechanism for social control form the point of view of the ecosystem.

The key requirements we have identified for making open service ecosystems scalable with respect to social control are i) persistent identities that trace back into valid legal actors such as individuals or registered organizations, ii) collaboration contracts that provide the vessel for binding and explicit mutual agreements missing from purely technology-driven solutions, and iii) contract-based objectivity in reputation, i.e. moving from the subjective recommendations visible in current-day recommender systems for consumers to reports on whether agreements were followed. For bringing the benefits to the individual enterprises, the additional requirement is iv) making explicit subjective trust decisions — instead of assuming implicit trust between everyone in the ecosystem. The fulfilment of these requirements makes it possible for service ecosystems to have a way to recover from fraud, contract violations and related misbehaviour.

In addition to the trust management addition, we find that the cloud infrastructure needs to adopt additional service-oriented architecture support for computing, communication services and collaboration network structure.

The aim of this chapter is to provide direction for developments in the near future in order to increase the maturity of cloud-based service ecosystems. The research vision for open service ecosystems with automated and semi-automated processes and easily set up collaborations is not right around the corner, but in the last years the evolutionary movement towards it has accelerated.

In order to bring real-world practices closer to the vision, we identify the following key directions for consideration: identity management, reputation sources, service orientation, standardization and legislation.

Identity management for services and contracting. In the Finnish environment, all commercial organizations have a centrally set identity (Y-tunnus) for taxation purposes, and private people have a uniquely identifying social security ID which is used to identify them across various different services[1]. Legislation is in place for accepting legally binding digital signatures through the VETUMA federation, for example using banks' login identification [67]. This means that the basic infrastructure is in place for online environments to support strong authentication of private people and service providers for the purposes of contracting.

[1] Finland operates with a single-identity model built around state authorities. For example in Great Britain, individuals have in practice multiple context-bound identities from different sources, so this approach is not universal.

An identity certification scheme for services should build on the identity of the service provider due to the provider being the actual legal entity behind a service. Any automated entry into new collaborations eventually falls back to the provider and a framework-establishing base contract it has manually signed.

The identities of single services have their own value, however, because reputation information should primarily be bound to individual services and only inferred between different services of a single provider. This is because different services operate under different policies and in practice can be established and maintained independently within an organization [69].

Service identification for the purposes of reputation may be connected to the same naming scheme that is necessary for unambiguously reaching them, although a resource-based naming scheme (such as DOIs) is likely to be preferrable to routing-based naming (such as URLs).

Reputation systems demand that identities should be long-lived in order to create an expectation of changes in reputation affecting future transactions, which is needed to support the social control effect [66]. We find that the organizational identity is sufficiently long-lived by nature, as mass-producing new organizations with their organizational taxation IDs is a nontrivial. Nothing stops an organization from releasing a massive number of instances of the same service, however. This problem we expect to alleviate by binding services to their providers, and using the reputation of the provider's other services as an inferred estimate for a new, unknown service's reputation. Similarly, the baseline reputation of a service composed from multiple individual services can be inferred from that of its members, but from that baseline on, it should also have its own reputation.

Establishing and identifying reputation sources. Trust management depends on its reputation information input. As long as reputation systems for inter-enterprise collaborations do not yet exist, we cannot expect to start rolling out automated trust management from scratch.

In other words, the first step is establishing sources of reputation information that can be used manually, to help reduce the time spent doing background checking on a service provider, for example by looking up what shows up in search engines.

We see that associations of service providers within a domain could be a potential source for collecting references and potentially partner feedback in one place. The main argument for this is maintaining joint interest in the quality of the reputation information: entirely independent and external review collections need a way to make money, and as we have seen in the previous chapter, connecting advertisement-based funding or "image consultation services" directly with the reputation system tend to cause problematic conflicts of interest.

On the other hand, a centralized association is only a real option when the domain is both open enough that expected membership in the association or its reputation network does not become an overwhelming barrier for entry to the market, and limited enough that all actors can be expected to consider joining it worth the costs. For example a global association of restaurant services is highly unlikely, which is part of the reason the review services for the domain will most likely remain independent companies. Certain parts of product delivery chains for providing supplies to the restaurants, on the other hand, already have a notably less broad and more localized customer base.

We are beginning to see this kind of direction becoming reality. In the Finnish construction business, Rakentamisen Laatu (RaLa, the name means literally Quality of Construction) was established to help deal with the asymmetry of information between the customer offering a construction work contract, and the service provider. The field has complex networks of subcontracting with small and medium-sized enterprises, reasonably well-structured agreements and processes such as standardized contract templates, high competition despite being reasonably location-dependent, and a widely identified need for internal control due to questionable practices in e.g. not paying the legally required worker social security fees in order to compete with pric-

ing. One service provided by RaLa supports the financial background checks of registered service providers[2], while another collects references of completed partnerships and records feedback on them. While references are public, the feedback system and reports of e.g. how the feedback relates to other providers in the same field is limited to members.

Service-orienting the cloud. The establishment of a basic cloud stack of low-level communication patterns is not yet sufficient by itself for supporting easy setup of collaborations. The cloud infrastructure will need to import principles from service oriented computing for handling communication, service discovery and collaborative business process establishment (see e.g. [59]).

All computational services, as they are depicted in this report, should form logical units that have a well-defined access interface and that can be loosely coupled with other services. Infrastructure services are needed for finding them, monitoring their behaviour and governing them through collaboration contracts and local policy.

In order to couple services together automatically, additional structure such as strong service typing becomes necessary. Service typing forms the basis for establishing reusable collaboration patterns as well. The basic idea is not too different from existing paradigms: in programming, different operations are available and make sense to strings (e.g. concatenation), integers (e.g. addition) and data structure objects, and for services, typing essentially specifies how to call the service, with what parameters, what kinds of things can be expected to come out of it, and what kind of configuration options the service has that should be considered in the coupling.

Standardization of shared information: reputation, contracts, collaboration patterns and service offers. Standards and de facto standards are central drivers for interoperability: in a simplified sense, they represent agreements between different active or dominant actors of a field on common practices. These agreements reduce the adjustment and integration effort of the average service provider and through that greatly improve the usability of the service ecosystem.

The immediate targets for it are information that is expected to be reused among different organizations, such as reputation reports, the format of electronic contracts, and the supporting artefacts we have identified, such as the the business network models representing collaboration patterns and joint business processes, and the service offers relating to them.

Examples of potential international standardization bodies can be divided roughly in two based on whether the standardization relates to service-oriented architectures or advancing cloud security mechanisms. In terms of service orientation, the global consortium OASIS drives adoption of e-business and web service standards, such as WS-Trust and WS-Policy described in the previous chapter. The Pilarcos model of the service ecosystem and the necessary infrastructure for it builds heavily on the reference model of open distributed processing (RM-ODP) [45], which is a joint effort between ISO, IEC and ITU-T.

The Cloud Security Alliance (CSA) formed in 2008 aims to promote use of best practices for providing security assurance within cloud computing. The US National Institute of Standards and Technology (NIST) and the European Network and Information Security Agency (ENISA) are active in topics related to cloud security as well.

Advancement of legislation. Automated contract negotiations, handling of breaches of reputation system contracts and so on will need improved backing in legislation. This kind of "societal standardization" follows technological advances and the lessons learned from their adoption, as the use patterns stabilize enough to be both possible and desirable targets for legal modelling.

[2]Tilaajavastuu, http://www.tilaajavastuu.fi/

Chapter 6

Conclusion

We have presented a vision of cloud-based open service ecosystem as environments for setting up, operating and managing service-oriented inter-enterprise collaborations in an automated, but human-controllable way.

These collaborations are supported by infrastructure services for discovering services, negotiating collaboration contracts based on templates that reflect best practices in the field, service behaviour monitoring, and reputation-based, subjective trust decisions that balance the expected benefits of the collaboration against its estimated risks in the business context.

Our review on the existing services and standards relevant for trust management identified a gap: trust management and reputation services robust enough for supporting inter-enterprise collaborations are essentially missing. The offered solutions are mainly targeted for opinion-based recommendations for eCommerce environments where the business models, clients and shopkeepers involved are far more limited.

Based on research domain solutions, we provide short and mid-term recommendations for further steps towards mature service ecosystems.

As a general trend, we expect that — like many other fields of technology — inter-enterprise collaborations in the cloud will follow three waves of development (cf. [45]): At first, the new approach or methodology is being explored, and all solutions for it are unique, manual and not very reusable: collaborations require some manual integration effort and instead of standards, smaller-scale agreements emerge. In the second wave, the experiences from the initial attempts can be gathered into best practices and models that allow the generation of interoperable solutions. Solutions of this phase rely on environments with implicit overall trust based on membership in the collaboration or ecosystem. This is what we have at present. In the third wave, configurable systems, standardized interfaces and technology convergence and allow the tools to be distilled into infrastructure services that can be reused between ecosystems, and their behaviour adjusted through reconfiguration rather than re-deployment. Solutions of this phase require explicit mechanisms to be included for automated management of trust-related knowledge and distributed decision making on trust.

The evolution towards mature service ecosystems requires the adoption of SOA/SOC and business process management (BPM) as an additional collaboration as a service (CaaS) layer over existing cloud computing solutions. This involves joint consortia work for creating business process vocabularies and contract template standards in a dynamic way (in repositories) to boost existing and new business domains, such as healthcare, education and production.

The development of any trust management system that can be used in the inter-enterprise collaborations context requires standards and technical solutions for identity management, in such a way that permanent, traceable identities of enterprises, groups, individuals and collaborations can

be used within and between ecosystems. Further, it requires appropriate standards and technical solutions for reputation management systems in such a way that experience information can be distributed freely within ecosystems and mediated by trusted repository holders.

Finally, as the best practices emerge from early adoption efforts, legislation enhancements should follow to better support software-based commitments and different aspects of experience information sharing.

Acknowledgements

This work has been performed at the Department of Computer Science of the University of Helsinki, in the Collaborative and Interoperable Computing Group (CINCO). It has been inspired by collaboration with the TEKES Cloud Software program of DIGILE (Finnish Strategic Centre for Science, Technology and Innovation in the field of ICT and digital business).

Bibliography

[1] ALNEMR, R., KOENIG, S., EYMANN, T., AND MEINEL, C. Enabling usage control through reputation objects: A discussion on e-commerce and the internet of services environments. *Journal for Theoretical and Applied Electronic Commerce Research 5*, 2 (Aug. 2010), 59–76.

[2] ALNEMR, R., AND MEINEL, C. From reputation models and systems to reputation ontologies. In *Trust Management V; 5th IFIP WG 11.11 International Conference, IFIPTM 2011; Proceedings* (Copenhagen, Denmark, June/July 2011), vol. 358 of *IFIP AICT*, Springer, pp. 98–116.

[3] ALNEMR, R., AND MEINEL, C. Why rating is not enough: A study on online reputation systems. In *7th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)* (Orlando, FL, USA, Oct. 2011), IEEE, pp. 415–421.

[4] The Amazon store and recommendation system website, 2007. `http://www.amazon.com/`.

[5] ARMBRUST, M., FOX, A., GRIFFITH, R., JOSEPH, A. D., KATZ, R., KONWINSKI, A., LEE, G., PATTERSON, D., RABKIN, A., STOICA, I., AND ZAHARIA, M. A view of cloud computing. *Communications of the ACM 53* (Apr. 2010), 50–58.

[6] ARRINGTON, M. Google in discussions to acquire Yelp for a half billion dollars or more. *TechCrunch* (Dec. 2009). `http://techcrunch.com/2009/12/17/google-acquire-buy-yelp/` (Visited 1.11.2012).

[7] ASOKAN, N., SHOUP, V., AND WAIDNER, M. Asynchronous protocols for optimistic fair exchange. In *Proceedings of the IEEE Symposium on Research in Security and Privacy* (May 1998), IEEE Computer Society, pp. 86–99.

[8] ATKINSON, B., DELLA-LIBERA, G., HADA, S., HONDO, M., HALLAM-BAKER, P., KALER, C., ET AL. *Web Services Security (WS-Security) Version 1.0*, Apr. 2002. `http://schemas.xmlsoap.org/specs/ws-security/ws-security.htm`.

[9] AXELROD, R., AND HAMILTON, W. D. The evolution of cooperation. *Science 211*, 4489 (1981), 1390–1396.

[10] BAJAJ, S., BOX, D., CHAPPELL, D., CURBERA, F., DANIELS, G., HALLAM-BAKER, P., HONDO, M., ET AL. *Web Services Policy Framework (WS-Policy) Version 1.2*, Mar. 2006. `http://specs.xmlsoap.org/ws/2004/09/policy/ws-policy.pdf`.

[11] Better Business Bureau, 2013. `http://www.bbb.org/`.

[12] BLAZE, M., FEIGENBAUM, J., AND KEROMYTIS, A. D. KeyNote: Trust management for public-key infrastructures (position paper). In *Proceedings of Security Protocols: 6th International Workshop, Cambridge, UK, April 1998* (Apr. 1998), Springer-Verlag, LNCS 1550/1998, pp. 59–63.

[13] BLAZE, M., FEIGENBAUM, J., AND LACY, J. Decentralized trust management. In *Proceedings of the IEEE Symposium on Security and Privacy* (Oakland, California, May 1996), IEEE, pp. 164–173.

[14] BUSINESS WIRE. Fifty percent of global online retail visits were to Amazon, eBay and Alibaba in June 2011: companiesandmarkets.com. `http://www.businesswire.com/news/home/20110823005719/en/Fifty-Percent-Global-Online-Retail-Visits-Amazon` (Visited 21.8.2012).

[15] CHENG, A., AND FRIEDMAN, E. Sybilproof reputation mechanisms. In *Proc. of the 2005 ACM SIGCOMM Workshop on Economics of Peer-to-Peer Systems* (2005), ACM, pp. 128–132.

[16] CHU, Y.-H., FEIGENBAUM, J., LAMACCHIA, B., RESNICK, P., AND STRAUSS, M. REFEREE: Trust management for Web applications. *Computer Networks and ISDN Systems 29*, 8–13 (Sept. 1997), 953–964.

[17] CITIZEN MEDIA LAW PROJECT. Cats and Dogs Animal Hospital, Inc. v. Yelp! Inc., 2011. `http://www.citmedialaw.org/threats/cats-and-dogs-animal-hospital-inc-v-yelp-inc`.

[18] DOUCEUR, J. R. The Sybil attack. In *Electronic Proceedings of the 1st International Workshop on Peer-to-Peer systems (IPTPS'02)* (Cambridge, MA, USA, Mar. 2002), p. 101.

[19] DWYER, N., BASU, A., AND MARSH, S. Reflections on measuring the trust empowerment qualities of digital designs. In *Trust Management VII* (June 2013), vol. 401 of *IFIP Advances in Information and Communication Technology*, Springer Berlin Heidelberg, pp. 127–135.

[20] The eBay online marketplace, 2013. `http://www.ebay.com/`.

[21] EBAY RULES AND POLICIES. Feedback extortion policy, 2012. `http://pages.ebay.com/help/policies/feedback-extortion.html` (Visited 21.8.2012.).

[22] eBuyer Feedback website, 2008. `http://www.ebuyer-feedback.com/` (Visited 21.8.2012).

[23] The Epinions recommendation system website, 2007. `http://www.epinions.com/`.

[24] FACEBOOK. Information for law enforcement authorities, 2012. `http://www.facebook.com/safety/groups/law/guidelines/` (Visited 9.11.2012).

[25] FARMER, F. R., AND GLASS, B. *Building Web Reputation Systems*. O'Reilly, 2010.

[26] FORTUNE. Fortune 500, a ranking of America's largest corporations. `http://www.fortune.com/500` (Visited 17.8.2012).

[27] FRIED, I. Google pulls plug on Google Wave. *CNET News* (Aug. 2010). `http://news.cnet.com/8301-13860_3-20012698-56.html` (Visited 17.8.2012).

[28] GAL-OZ, N., GRINSHPOUN, T., GUDES, E., AND FRIESE, I. TRIC: An infrastructure for trust and reputation across virtual communities. In *Proceedings of the Fifth International Conference on Internet and Web Applications and Services (ICIW 2010)* (Barcelona, Spain, May 2010), IEEE, pp. 43–50.

[29] GAL-OZ, N., GUDES, E., AND HENDLER, D. A robust and knot-aware trust-based reputation model. In *Trust Management II* (Pisa, Italy, May 2008), vol. 263 of *IFIP International Federation for Information Processing*, Springer, pp. 167–182.

[30] GODADDY. News releases: Go Daddy site outage investigation completed. `http://www.godaddy.com/newscenter/release-view.aspx?news_item_id=410` (Visited 9.11.2012).

[31] GOLLMANN, D. From access control to trust management, and back — a petition. In *Trust Management V; 5th IFIP WG 11.11 International Conference, IFIPTM 2011; Proceedings* (Copenhagen, Denmark, June/July 2011), vol. 358 of *IFIP AICT*, Springer, pp. 1–8.

[32] GUDES, E., GAL-OZ, N., AND GRUBSHTEIN, A. Methods for computing trust and reputation while preserving privacy. In *Data and Applications Security XXIII* (2009), vol. 5645 of *Springer LNCS*, pp. 291–298.

[33] HABIB, S., RIES, S., HAUKE, S., AND MUHLHAUSER, M. Fusion of opinions under uncertainty and conflict – application to trust assessment for cloud marketplaces. In *Trust, Security and Privacy in Computing and Communications (TrustCom), 2012 IEEE 11th International Conference on* (June 2012), pp. 109 –118.

[34] HAINES, L. Gang robs eBay car buyer at gunpoint. *The Register* (2006). `http://www.theregister.co.uk/2006/08/29/ebay_scam_robbery/` (Visited 21.8.2012).

[35] HILL, C. Regulating the rating agencies. *American Law & Economics Association Annual Meetings* (2004), 42–95.

[36] HU, N., BOSE, I., KOH, N. S., AND LIU, L. Manipulation of online reviews: An analysis of ratings, readability, and sentiments. *Decision Support Systems 52* (2012), 674–684.

[37] The Huuto.net online marketplace site (in Finnish), 2012. `http://www.huuto.net/`.

[38] Jefferson County School District No. R-1 v. Moody's Investor's Services, Inc., May 1999. 175 F.3d 848, `http://bulk.resource.org/courts.gov/c/F3/175/175.F3d.848.97-1157.html`.

[39] JØSANG, A. Trust extortion on the internet. In *Proceedings of the 7th international conference on Security and Trust Management (STM'11)* (2012), vol. 7170 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 6–21.

[40] KAUR, P., RUOHOMAA, S., AND KUTVONEN, L. Enabling user involvement in trust decision making for inter-enterprise collaborations. *International Journal On Advances In Intelligent Systems 5*, 3&4 (Dec. 2012), 533–552.

[41] KELLY, H. Police embrace social media as crime-fighting tool. *CNN.com* (Aug. 2012). `http://www.cnn.co.uk/2012/08/30/tech/social-media/fighting-crime-social-media/index.html` (Visited 9.11.2012).

[42] KINATEDER, M., BASCHNY, E., AND ROTHERMEL, K. Towards a generic trust model - comparison of various trust update algorithms. In *Proceedings of Trust Management: Third International Conference, iTrust 2005, Paris, France, May 23–26, 2005* (Apr. 2005), P. Herrmann, V. Issarny, and S. Shiu, Eds., vol. 3477 of *LNCS*, Springer-Verlag, pp. 177–192.

[43] KINATEDER, M., AND ROTHERMEL, K. Architecture and algorithms for a distributed reputation system. In *Proceedings of Trust Management: First International Conference, iTrust 2003, Heraklion, Crete, Greece, May 28-30, 2003* (May 2003), P. Nixon and S. Terzis, Eds., vol. 2692 of *LNCS*, Springer-Verlag, pp. 1–16.

[44] KUTVONEN, L. Multi-tier agent architecture for open service ecosystems. In *Proceedings of First International Conference on Agreement Technologies* (Dubrovnik, Croatia, Oct. 2012).

[45] KUTVONEN, L. ODP RM reflections on open service ecosystems. *Computer Standards & Interfaces 35* (Mar. 2013), 294–312.

[46] KUTVONEN, L., AND RUOHOMAA, S. Behavioural evaluation of reputation-based trust systems — from strategic networks to mature ecosystems, 2013. Submitted manuscript.

[47] KUTVONEN, L., RUOKOLAINEN, T., RUOHOMAA, S., AND METSO, J. Service-oriented middleware for managing inter-enterprise collaborations. In *Global Implications of Modern Enterprise Information Systems: Technologies and Applications* (Dec. 2008), Advances in Enterprise Information Systems (AEIS), IGI Global, pp. 209–241.

[48] MAHALINGAM MANI (ED.). ORMS: Use-cases version 0.24b, working draft 21. Tech. rep., OASIS, Jan. 2009. `https://www.oasis-open.org/committees/download.php/30812/orms-uc-v0_24Bar.doc`.

[49] MASSA, P., AND AVESANI, P. Trust metrics in recommender systems. In *Computing with Social Trust* (2009), pp. 259–285.

[50] MASSA, P., AND BHATTACHARJEE, B. Using trust in recommender systems: An experimental analysis. In *Proceedings of Trust Management: Second International Conference, iTrust 2004, Oxford, UK, March 29–April 1, 2004* (Mar. 2004), Springer-Verlag, LNCS 2995/2004, pp. 221–235.

[51] MEHANDIEV, N., AND GREFEN, P., Eds. *Dynamic Business Process Formation for Instant Virtual Enterprises*. Advanced Information and Knowledge Processing. Springer, June 2010.

[52] MICALI, S. Simple and fast optimistic protocols for fair electronic exchange. In *Proceedings of the twenty-second annual symposium on Principles of distributed computing (PODC '03)* (2003), ACM, pp. 12–19.

[53] MILLER, C. C. Yelp makes changes in response to small-business owners. *The New York Times: Bits Blog* (Apr. 2010). `http://bits.blogs.nytimes.com/2010/04/06/yelp-makes-changes-to-appease-small-business-owners/` (Visited 31.10.2012).

[54] National consumer agency of Ireland, 2010. `http://www.consumerconnect.ie/`.

[55] NEW YORK TIMES EDITORIAL. Banks and WikiLeaks. *The New York Times* (Dec. 2010). `http://www.nytimes.com/2010/12/26/opinion/26sun3.html?_r=0` (Visited 9.11.2012).

[56] OASIS ORMS TC. *OASIS Open Reputation Management Systems (ORMS) Technical Committee charter*, 2008. https://www.oasis-open.org/committees/orms/charter.php (Visited 20.8.2012).

[57] OASIS WEB SERVICE SECURE EXCHANGE TC. *WS-Trust 1.3 OASIS Standard*. OASIS, Mar. 2007. http://docs.oasis-open.org/ws-sx/ws-trust/v1.3/ws-trust.pdf.

[58] United Kingdom Office of Fair Trading, 2010. http://www.oft.gov.uk/.

[59] PAPAZOGLOU, M. P. *Web Services & SOA Principles and Technology. Second Edition*. Pearson Education, Jan. 2012.

[60] PAPAZOGLOU, M. P., AND VAN DEN HEUVEL, W.-J. Blueprinting the cloud. *IEEE Internet Computing 15*, 6 (2011), 74–79.

[61] PARSA, A. Exclusive: Belkin's development rep is hiring people to write fake positive Amazon reviews. *The Daily Background* (Jan. 2009). http://www.thedailybackground.com/2009/01/16/exclusive-belkins-development-rep-is-hiring-people-to-write-fake-positive-amazon-reviews/ (Visited 16.8.2012).

[62] PENTLAND, B. T. Will auditors take over the world? program, technique and the verification of everything. *Accounting, Organizations and Society 25* (2000), 307–312.

[63] POWER, M. *The Audit Society: Rituals of Verification*. Oxford university press, 1999.

[64] RABELO, R. J., GUSMEROLI, S., ARANA, C., AND NAGELLEN, T. The ECOLEAD ICT infrastructure for collaborative networked organizations. In *Network-Centric Collaboration and Supporting Frameworks* (2006), vol. 224, Springer, pp. 451–460.

[65] Reputation.com online image control website, 2012. http://www.reputation.com/.

[66] RESNICK, P., ZECKHAUSER, R., FRIEDMAN, E., AND KUWABARA, K. Reputation systems. *Communications of the ACM 43*, 12 (Dec. 2000), 45–48.

[67] RISSANEN, T. Electronic identity in Finland: ID cards vs. bank IDs. *Identity in the Information Society 3*, 1 (July 2010), 175–194.

[68] RISTENPART, T., TROMER, E., SCHACHAM, H., AND SAVAGE, S. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *Proceedings of the 16th ACM conference on Computer and communications security (CSS'09)* (New York, NY, USA, 2009), ACM, pp. 199–212.

[69] RUOHOMAA, S. *The effect of reputation on trust decisions in inter-enterprise collaborations*. PhD thesis, University of Helsinki, Department of Computer Science, May 2012.

[70] RUOHOMAA, S., KAUR, P., AND KUTVONEN, L. From subjective reputation to verifiable experiences - augmenting peer-control mechanisms for open service ecosystems. In *Trust Management VI* (Surat, India, May 2012), vol. 374 of *IFIP Advances in Information and Communication Technology*, pp. 142–157.

[71] RUOHOMAA, S., AND KUTVONEN, L. Towards trust management for cloud-based ecosystems. *Discussion Paper, Communications of Cloud Software* (2013).

[72] RUOHOMAA, S., LUOMA, E., AND KUTVONEN, L. Trust broker service, 2013. Submitted manuscript.

[73] RUOKOLAINEN, T. *A Model-Driven Approach to Service Ecosystem Engineering*. PhD thesis, University of Helsinki, Department of Computer Science, Feb. 2013.

[74] SAKIMURA, N., AND OASIS ORMS TC. Open Reputation Data (ORD) draft version 0.1. Tech. rep., OASIS, Sept. 2010. https://www.oasis-open.org/committees/download.php/39509/Open_Reputation_Management_System_OR_.doc.

[75] SAVITZ, E. Amazon: Now one-third of all U.S. e-commerce. *Barrons Tech Trader Daily* (Apr. 2009). http://blogs.barrons.com/techtraderdaily/2009/04/14/amazon-now-one-third-of-all-us-e-commerce/ (Visited 21.8.2012).

[76] SCHIFFMAN, B. EBay pulls feedback option for sellers. *Wired* (May 2008). http://www.wired.com/business/2008/05/ebay-feedback/ (Visited 21.8.2012).

[77] SCHMID, G. Report on the existence of a global system for the interception of private and commercial communications (ECHELON interception system), part 1. Tech. rep., European Parliament: Temporary Committee on the ECHELON Interception System, 2001.

[78] SOLHAUG, B., AND STØLEN, K. Uncertainty, subjectivity, trust and risk: How it all fits together. In *STM 2011* (2012), vol. 7170 of *Lecture Notes in Computer Science*, pp. 1–5.

[79] Standard & Poor's website, 2010. http://www.standardandpoors.com.

[80] The Taobao Marketplace site (in Chinese), 2012. http://www.taobao.com/.

[81] TECHNOLOGIES", A. Akamai.com: Facts & figures. http://www.akamai.com/html/about/facts_figures.html (Visited 9.11.2012).

[82] TECHNOLOGIES", A. Akamai.com: Methodology & data collection. http://www.akamai.com/html/technology/realtime_web_methodology.html (Visited 9.11.2012).

[83] Tripadvisor: hotel, flight and vacation rental website, 2012. http://www.tripadvisor.com/.

[84] Trustmark tradesman certification, 2010. http://www.trustmark.org.uk/.

[85] VU, L.-H., PAPAIOANNOU, T. G., AND ABERER, K. Impact of trust management and information sharing to adversarial cost in ranking systems. In *IFIPTM 2010* (2010), vol. 321 of *IFIP AICT*, Springer, pp. 108–124.

[86] WOOD, P., ET AL. Internet security threat report. Tech. rep., Symantec, Apr. 2012. http://www.symantec.com/threatreport/ (Visited 9.11.2012).

[87] Yelp: a local business consumer review website, 2013. http://www.yelp.com/.

[88] ZHANG, Q., CHENG, L., AND BOUTABA, R. Cloud computing: state-of-the-art and research challenges. *Journal of Internet services and applications 1*, 1 (Apr. 2010), 7–18.